CCIE Service Provider Lab Workbook v4.0 (http://labs.ine.com/workbook/toc/service-provider-v4) » CCIE SP v4 Advanced Technology Labs - Services

LISP - Basic overlay IPv4 to IPv4

« Inter-AS MPLS TE (/workbook/view/service-provider-v4/task/inter-as-mpls-te-Mjg4OQ%3D%3D) | LISP - IPv6 over IPv4 core (/workbook/view/service-provider-v4/task/lisp-ipv6-over-ipv4-core-Mjk1Mg%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

- Configure a LISP overlay to provide IPv4 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 IPv4 EID registrations from the LISP sites as follows:
  - LISP Site A:
    - EID 10.1.0.0/16 - allow more specifics
    - EID 1.1.1.1/32
    - Use authentication key "SITEA_xTR-R2"
  - LISP Site B:
    - EID 10.19.0.0/16 - allow more specifics
    - EID 20.20.20.20/32
    - 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:
    - 10.1.2.0/24
    - 1.1.1.1/32
  - 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

CONTENTS

  - 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:
    - 10.19.20.0/24
    - 20.20.20.20/32
  - 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 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 1.1.1.1/32 locator-set SITEA_RLOC_SET
 database-mapping 10.1.2.0/24 locator-set SITEA_RLOC_SET
 ipv4 itr map-resolver 9.9.9.9
 ipv4 itr
 ipv4 etr map-server 9.9.9.9 key SITEA_xTR-R2
 ipv4 etr
 exit



!
! Map Server config
!

R9:
router lisp
 site SITE_A
  authentication-key SITEA_xTR-R2
  eid-prefix 1.1.1.1/32
  eid-prefix 10.1.0.0/16 accept-more-specifics
  exit
 !
 site SITE_B
  authentication-key SITEB_xTR-XR1
  eid-prefix 10.19.0.0/16 accept-more-specifics
  eid-prefix 20.20.20.20/32
  exit
 !
 ipv4 map-server
 ipv4 map-resolver
 exit



!
! LISP SiteB xTR config
!

XR1:
router lisp
 address-family ipv4 unicast
 !
 locator-set SITEB_RLOC_SET
  20.5.19.19 priority 1 weight 50
```

CONTENTS ›

https://t.me/learningnets

```
  20.6.19.19 priority 1 weight 50
 !
 eid-table default instance-id 0
  address-family ipv4 unicast
   etr map-server 9.9.9.9 key encrypted 122A2C233729331C1E1969100164
   etr
   itr map-resolver 9.9.9.9
   itr
   database-mapping 10.19.20.0/24 locator-set SITEB_RLOC_SET
   database-mapping 20.20.20.20/32 locator-set SITEB_RLOC_SET
  !
 !
 locator-table default
!
```

Verification

SP Core Network

Before we begin diving into LISP and verifying this task, lets take a moment to understand how the current network is setup. It will be important to understand how the underlay routing is configured in this network in order to fully appreciate how the LISP overlay is providing connectivity.

The SP Core consists of R3-6, R9, and R10. All of these devices are dual stacked, and are running OSPFv2 for IPv4 and OSPFv3 for IPv6 connectivity.

CONTENTS ›

```
R3#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
9.9.9.9           1   FULL/DR         00:00:37    20.3.9.9        GigabitEthernet1.39
6.6.6.6           1   FULL/DR         00:00:31    20.3.6.6        GigabitEthernet1.36
4.4.4.4           1   FULL/DR         00:00:34    20.3.4.4        GigabitEthernet1.34

R3#show ipv6 ospf neighbor

            OSPFv3 Router with ID (3.3.3.3) (Process ID 100)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
9.9.9.9           1   FULL/DR         00:00:34    12              GigabitEthernet1.39
6.6.6.6           1   FULL/DR         00:00:30    12              GigabitEthernet1.36
4.4.4.4           1   FULL/DR         00:00:37    13              GigabitEthernet1.34


R3#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/2] via 20.3.4.4, 3d00h, GigabitEthernet1.34
      5.0.0.0/32 is subnetted, 1 subnets
O        5.5.5.5 [110/3] via 20.3.6.6, 3d00h, GigabitEthernet1.36
                 [110/3] via 20.3.4.4, 3d00h, GigabitEthernet1.34
      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/2] via 20.3.6.6, 3d00h, GigabitEthernet1.36
      9.0.0.0/32 is subnetted, 1 subnets
O        9.9.9.9 [110/2] via 20.3.9.9, 3d00h, GigabitEthernet1.39
      10.0.0.0/32 is subnetted, 1 subnets
O        10.10.10.10 [110/3] via 20.3.6.6, 3d00h, GigabitEthernet1.36
      20.0.0.0/8 is variably subnetted, 17 subnets, 2 masks
O        20.4.5.0/24 [110/2] via 20.3.4.4, 3d00h, GigabitEthernet1.34
O        20.4.6.0/24 [110/2] via 20.3.6.6, 3d00h, GigabitEthernet1.36
                     [110/2] via 20.3.4.4, 3d00h, GigabitEthernet1.34
O        20.5.6.0/24 [110/2] via 20.3.6.6, 3d00h, GigabitEthernet1.36
O        20.6.10.0/24 [110/2] via 20.3.6.6, 3d00h, GigabitEthernet1.36


R3#show ipv6 route ospf
IPv6 Routing Table - default - 24 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
```

https://t.me/learningnets

```
       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
O   2002::4:4:4:4/128 [110/1]
     via FE80::250:56FF:FE9E:1302, GigabitEthernet1.34
O   2002::5:5:5:5/128 [110/2]
     via FE80::250:56FF:FE9E:1302, GigabitEthernet1.34
     via FE80::250:56FF:FE9E:5CEC, GigabitEthernet1.36
O   2002::6:6:6:6/128 [110/1]
     via FE80::250:56FF:FE9E:5CEC, GigabitEthernet1.36
O   2002::9:9:9:9/128 [110/1]
     via FE80::250:56FF:FE9E:7492, GigabitEthernet1.39
O   2002::10:10:10:10/128 [110/2]
     via FE80::250:56FF:FE9E:5CEC, GigabitEthernet1.36
O   2002:20:4:5::/64 [110/2]
     via FE80::250:56FF:FE9E:1302, GigabitEthernet1.34
O   2002:20:4:6::/64 [110/2]
     via FE80::250:56FF:FE9E:5CEC, GigabitEthernet1.36
     via FE80::250:56FF:FE9E:1302, GigabitEthernet1.34
O   2002:20:5:6::/64 [110/2]
     via FE80::250:56FF:FE9E:5CEC, GigabitEthernet1.36
O   2002:20:6:10::/64 [110/2]
     via FE80::250:56FF:FE9E:5CEC, GigabitEthernet1.36


R6#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.10.10.10       1   FULL/DR         00:00:35    20.6.10.10      GigabitEthernet1.610
5.5.5.5           1   FULL/BDR        00:00:32    20.5.6.5        GigabitEthernet1.56
4.4.4.4           1   FULL/BDR        00:00:30    20.4.6.4        GigabitEthernet1.46
3.3.3.3           1   FULL/BDR        00:00:36    20.3.6.3        GigabitEthernet1.36

R6#show ipv6 ospf neighbor

            OSPFv3 Router with ID (6.6.6.6) (Process ID 100)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
10.10.10.10       1   FULL/DR         00:00:30    12              GigabitEthernet1.610
5.5.5.5           1   FULL/BDR        00:00:33    13              GigabitEthernet1.56
4.4.4.4           1   FULL/BDR        00:00:35    15              GigabitEthernet1.46
3.3.3.3           1   FULL/BDR        00:00:36    14              GigabitEthernet1.36



R6#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

CONTENTS ›

https://t.me/learningnets

```
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set


      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/2] via 20.3.6.3, 3d00h, GigabitEthernet1.36
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/2] via 20.4.6.4, 3d00h, GigabitEthernet1.46
      5.0.0.0/32 is subnetted, 1 subnets
O        5.5.5.5 [110/2] via 20.5.6.5, 3d00h, GigabitEthernet1.56
      9.0.0.0/32 is subnetted, 1 subnets
O        9.9.9.9 [110/3] via 20.3.6.3, 3d00h, GigabitEthernet1.36
      10.0.0.0/32 is subnetted, 1 subnets
O        10.10.10.10 [110/2] via 20.6.10.10, 3d00h, GigabitEthernet1.610
      20.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
O        20.3.4.0/24 [110/2] via 20.4.6.4, 3d00h, GigabitEthernet1.46
                     [110/2] via 20.3.6.3, 3d00h, GigabitEthernet1.36
O        20.3.9.0/24 [110/2] via 20.3.6.3, 3d00h, GigabitEthernet1.36
O        20.4.5.0/24 [110/2] via 20.5.6.5, 3d00h, GigabitEthernet1.56
                     [110/2] via 20.4.6.4, 3d00h, GigabitEthernet1.46



R6#show ipv6 route ospf
IPv6 Routing Table - default - 23 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
O   2002::3:3:3:3/128 [110/1]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.36
O   2002::4:4:4:4/128 [110/1]
     via FE80::250:56FF:FE9E:1302, GigabitEthernet1.46
O   2002::5:5:5:5/128 [110/1]
     via FE80::250:56FF:FE9E:962, GigabitEthernet1.56
O   2002::9:9:9:9/128 [110/2]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.36
O   2002::10:10:10:10/128 [110/1]
     via FE80::250:56FF:FE9E:5F10, GigabitEthernet1.610
O   2002:20:3:4::/64 [110/2]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.36
     via FE80::250:56FF:FE9E:1302, GigabitEthernet1.46
O   2002:20:3:9::/64 [110/2]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.36
O   2002:20:4:5::/64 [110/2]
     via FE80::250:56FF:FE9E:1302, GigabitEthernet1.46
     via FE80::250:56FF:FE9E:962, GigabitEthernet1.56
```

CONTENTS ›

```
R9#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           1   FULL/BDR        00:00:34    20.3.9.3        GigabitEthernet1.39
R9#show ipv6 ospf neighbor

            OSPFv3 Router with ID (9.9.9.9) (Process ID 100)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
3.3.3.3           1   FULL/BDR        00:00:35    15              GigabitEthernet1.39



R9#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/2] via 20.3.9.3, 3d00h, GigabitEthernet1.39
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/3] via 20.3.9.3, 3d00h, GigabitEthernet1.39
      5.0.0.0/32 is subnetted, 1 subnets
O        5.5.5.5 [110/4] via 20.3.9.3, 3d00h, GigabitEthernet1.39
      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/3] via 20.3.9.3, 3d00h, GigabitEthernet1.39
      10.0.0.0/32 is subnetted, 1 subnets
O        10.10.10.10 [110/4] via 20.3.9.3, 3d00h, GigabitEthernet1.39
      20.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
O        20.3.4.0/24 [110/2] via 20.3.9.3, 3d00h, GigabitEthernet1.39
O        20.3.6.0/24 [110/2] via 20.3.9.3, 3d00h, GigabitEthernet1.39
O        20.4.5.0/24 [110/3] via 20.3.9.3, 3d00h, GigabitEthernet1.39
O        20.4.6.0/24 [110/3] via 20.3.9.3, 3d00h, GigabitEthernet1.39
O        20.5.6.0/24 [110/3] via 20.3.9.3, 3d00h, GigabitEthernet1.39
O        20.6.10.0/24 [110/3] via 20.3.9.3, 3d00h, GigabitEthernet1.39



R9#show ipv6 route ospf
IPv6 Routing Table - default - 20 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
```

https://t.me/learningnets

```
       a - Application
O   2002::3:3:3:3/128 [110/1]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002::4:4:4:4/128 [110/2]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002::5:5:5:5/128 [110/3]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002::6:6:6:6/128 [110/2]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002::10:10:10:10/128 [110/3]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002:20:3:4::/64 [110/2]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002:20:3:6::/64 [110/2]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002:20:4:5::/64 [110/3]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002:20:4:6::/64 [110/3]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002:20:5:6::/64 [110/3]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
O   2002:20:6:10::/64 [110/3]
     via FE80::250:56FF:FE9E:6E6A, GigabitEthernet1.39
```

BGP AS 100 is running between all devices in the SP Core - note that there is no MPLS configured. R3 is configured as the Route Reflector for for IPv4 and IPv6 address-families. The SP is originating all of the Customer facing links into BGP via connected redistribution. Additionally, R9 and R10 have network statements for their Loopback0 prefix.

CONTENTS ›

https://t.me/learningnets

```
R3#show bgp ipv4 unicast summary

BGP router identifier 3.3.3.3, local AS number 100

BGP table version is 106, main routing table version 106

11 network entries using 2728 bytes of memory

11 path entries using 1320 bytes of memory

6/5 BGP path/bestpath attribute entries using 1488 bytes of memory

2 BGP AS-PATH entries using 48 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 5584 total bytes of memory

BGP activity 16/0 prefixes, 16/0 paths, scan interval 60 secs


Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd

 4.4.4.4         4          100    4845    4862      106    0    0 3d00h           3

 5.5.5.5         4          100    4851    4860      106    0    0 3d00h           3

 6.6.6.6         4          100    4855    4862      106    0    0 3d00h           1

 9.9.9.9         4          100    4815    4858      106    0    0 3d00h           1

 10.10.10.10     4          100    4819    4874      106    0    0 3d00h           1

20.2.3.2        4          200    1420    1429      106    0    0 21:17:19        0



R3#show bgp ipv6 unicast summary

BGP router identifier 3.3.3.3, local AS number 100

BGP table version is 8, main routing table version 8

5 network entries using 1360 bytes of memory

5 path entries using 720 bytes of memory

3/3 BGP path/bestpath attribute entries using 744 bytes of memory

2 BGP AS-PATH entries using 48 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 2872 total bytes of memory

BGP activity 16/0 prefixes, 16/0 paths, scan interval 60 secs


Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd

 2002::4:4:4:4   4          100    4932    4812        8    0    0 3d00h           3

 2002::5:5:5:5   4          100    4948    4809        8    0    0 3d00h           2

 2002::6:6:6:6   4          100    4807    4814        8    0    0 3d00h           0

 2002::9:9:9:9   4          100    4808    4811        8    0    0 3d00h           0

 2002::10:10:10:10

                4          100    4808    4807        8    0    0 3d00h           0
```

The SP has links to multiple customers. LISP Site A and LISP Site B are dual homed into the SP network, and run BGP AS 200 and 1900 respectively. R2 and R3 are the PEs for LISP Site A, and R5 and R6 are the PEs for LISP Site B. The provider is peering with both of these sites, and is advertising the PE/CE links into BGP on their behalf. Notice that neither site is advertising routes into the SP. Also notice that only IPv4 is enabled to both of these sites.

› CONTENTS

```
R3#show bgp ipv4 unicast summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 106, main routing table version 106
11 network entries using 2728 bytes of memory
11 path entries using 1320 bytes of memory
6/5 BGP path/bestpath attribute entries using 1488 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5584 total bytes of memory
BGP activity 16/0 prefixes, 16/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4          100    4851    4868      106    0    0 3d00h           3
5.5.5.5         4          100    4857    4866      106    0    0 3d00h           3
6.6.6.6         4          100    4862    4868      106    0    0 3d00h           1
9.9.9.9         4          100    4821    4865      106    0    0 3d00h           1
10.10.10.10     4          100    4826    4881      106    0    0 3d00h           1
20.2.3.2        4          200    1426    1436      106    0    0 21:23:11        0



R4#show bgp ipv4 unicast summary
BGP router identifier 4.4.4.4, local AS number 100
BGP table version is 200, main routing table version 200
11 network entries using 2728 bytes of memory
11 path entries using 1320 bytes of memory
6/5 BGP path/bestpath attribute entries using 1488 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5744 total bytes of memory
BGP activity 16/0 prefixes, 16/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4          100    4868    4851      200    0    0 3d00h           7
20.2.4.2        4          200    1410    1437      200    0    0 21:23:02        0
20.4.7.7        4          700    4863    4863      200    0    0 3d00h           1



R5#show bgp ipv4 unicast summary
BGP router identifier 5.5.5.5, local AS number 100
BGP table version is 200, main routing table version 200
11 network entries using 2728 bytes of memory
11 path entries using 1320 bytes of memory
6/5 BGP path/bestpath attribute entries using 1488 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5744 total bytes of memory
BGP activity 16/0 prefixes, 16/0 paths, scan interval 60 secs
```

CONTENTS ›

https://t.me/learningnets

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4          100    4867    4859      200    0    0 3d00h           7
20.5.8.8        4          800    4857    4854      200    0    0 3d00h           1
 20.5.19.19      4         1900    1330    1477      200    0    0 22:00:48        0



R6#show bgp ipv4 unicast summary
BGP router identifier 6.6.6.6, local AS number 100
BGP table version is 109, main routing table version 109
11 network entries using 2728 bytes of memory
11 path entries using 1320 bytes of memory
6/5 BGP path/bestpath attribute entries using 1488 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5744 total bytes of memory
BGP activity 16/0 prefixes, 16/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4          100    4870    4863      109    0    0 3d00h           9
 20.6.19.19      4         1900    1330    1478      109    0    0 22:00:40        0
```

The SP Core also peers with two other Non LISP sites: BGP AS 700 and 800. Both of these sites are dual stacked, and the SP provides IPv4 and IPv6 transport for these sites.

CONTENTS ›

```
R7#show bgp ipv4 unicast summary

BGP router identifier 7.7.7.7, local AS number 700

BGP table version is 99, main routing table version 99

10 network entries using 2480 bytes of memory

10 path entries using 1200 bytes of memory

5/5 BGP path/bestpath attribute entries using 1240 bytes of memory

2 BGP AS-PATH entries using 64 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 4984 total bytes of memory

BGP activity 15/0 prefixes, 15/0 paths, scan interval 60 secs


Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd

 20.4.7.4        4          100    4869    4869       99    0    0 3d00h           9


R7#show bgp ipv6 unicast summary

BGP router identifier 7.7.7.7, local AS number 700

BGP table version is 444, main routing table version 444

5 network entries using 1360 bytes of memory

5 path entries using 720 bytes of memory

4/4 BGP path/bestpath attribute entries using 992 bytes of memory

2 BGP AS-PATH entries using 64 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 3136 total bytes of memory

BGP activity 15/0 prefixes, 15/0 paths, scan interval 60 secs


Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd

 2002:20:4:7::4  4          100    4867    4860      444    0    0 3d00h           4
```

```
R8#show bgp ipv4 unicast summary

BGP router identifier 8.8.8.8, local AS number 800

BGP table version is 99, main routing table version 99

10 network entries using 2480 bytes of memory

10 path entries using 1200 bytes of memory

5/5 BGP path/bestpath attribute entries using 1240 bytes of memory

2 BGP AS-PATH entries using 64 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 4984 total bytes of memory

BGP activity 15/0 prefixes, 15/0 paths, scan interval 60 secs


Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd

 20.5.8.5        4          100    4859    4863       99    0    0 3d00h           9


R8#show bgp ipv6 unicast summary

BGP router identifier 8.8.8.8, local AS number 800

BGP table version is 443, main routing table version 443

5 network entries using 1360 bytes of memory

5 path entries using 720 bytes of memory

4/4 BGP path/bestpath attribute entries using 992 bytes of memory

2 BGP AS-PATH entries using 64 bytes of memory
```

CONTENTS ›

https://t.me/learningnets

```
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3136 total bytes of memory
BGP activity 15/0 prefixes, 15/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
 2002:20:5:8::5  4          100    4863    4857      443    0    0 3d00h           4
```

With this underlying infrastructure, each site's edge router is provided with connectivity into the provider network, from now on referred to as the "underlay".

```
R2#show ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

      7.0.0.0/32 is subnetted, 1 subnets
B        7.7.7.7 [20/0] via 20.2.4.4, 21:37:53
      8.0.0.0/32 is subnetted, 1 subnets
B        8.8.8.8 [20/0] via 20.2.4.4, 21:37:53
      9.0.0.0/32 is subnetted, 1 subnets
B        9.9.9.9 [20/0] via 20.2.4.4, 21:37:53
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
B        10.10.10.10/32 [20/0] via 20.2.4.4, 21:37:53
      20.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
B        20.4.7.0/24 [20/0] via 20.2.3.3, 21:37:45
B        20.5.8.0/24 [20/0] via 20.2.4.4, 21:37:53
B        20.5.19.0/24 [20/0] via 20.2.4.4, 21:37:53
B        20.6.19.0/24 [20/0] via 20.2.4.4, 21:37:53


RP/0/0/CPU0:XR1#show route ipv4 bgp
Tue Jul 21 23:00:57.356 UTC

B    7.7.7.7/32 [20/0] via 20.5.19.5, 22:14:18
B    8.8.8.8/32 [20/0] via 20.5.19.5, 22:14:18
B    9.9.9.9/32 [20/0] via 20.5.19.5, 21:51:36
B    10.10.10.10/32 [20/0] via 20.5.19.5, 21:51:06
B    20.2.3.0/24 [20/0] via 20.5.19.5, 22:14:18
B    20.2.4.0/24 [20/0] via 20.5.19.5, 22:14:18
B    20.4.7.0/24 [20/0] via 20.5.19.5, 22:14:18
B    20.5.8.0/24 [20/0] via 20.6.19.6, 22:14:18
```

R2 and XR1 have connectivity between each other's uplinks into the SP, as well as to routes advertised by AS 700 and 800, and the Loopbacks of R9 and R10. All sites connecting into the SP have connectivity all all routes into the BGP table.

CONTENTS ›

https://t.me/learningnets

```
R2#ping 20.5.19.19
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.5.19.19, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/18 ms
R2#ping 20.6.19.19
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.6.19.19, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/19 ms
R2#ping 9.9.9.9
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 9.9.9.9, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/6/19 ms
R2#ping 7.7.7.7
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/18 ms


RP/0/0/CPU0:XR1#ping 20.2.4.2
Tue Jul 21 23:06:17.424 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.2.4.2, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/6/29 ms
RP/0/0/CPU0:XR1#ping 10.10.10.10
Tue Jul 21 23:06:22.543 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.10, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/19 ms
RP/0/0/CPU0:XR1#ping 8.8.8.8
Tue Jul 21 23:06:31.563 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/19 ms
```

CONTENTS

Site Routing

LISP Site A and B run EIGRP within the site. R2 and XR1 originate a default route into EIGRP so that R1 and XR2 have a way to route out of the site. R2 and XR1 have a static default route to null0 which is redistributed into EIGRP.

https://t.me/learningnets

```
R2#show ip eigrp neighbors
EIGRP-IPv4 VR(LISP_SITE_A) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   10.1.2.1                Gi1.12                   13 23:28:20    1   100  0  7

R2#show ip route eigrp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

      1.0.0.0/32 is subnetted, 1 subnets
D        1.1.1.1 [90/10880] via 10.1.2.1, 23:28:30, GigabitEthernet1.12


R2#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
   Known via "static", distance 1, metric 0 (connected), candidate default path
  Redistributing via eigrp 100
   Advertised by eigrp 100 route-map DEFAULT_EIGRP
  Routing Descriptor Blocks:
  * directly connected, via Null0
      Route metric is 0, traffic share count is 1
```

```
R1#show ip route eigrp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is 10.1.2.2 to network 0.0.0.0

D*EX  0.0.0.0/0 [170/568320] via 10.1.2.2, 00:12:04, GigabitEthernet1.12
      2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/10880] via 10.1.2.2, 23:28:46, GigabitEthernet1.12


RP/0/0/CPU0:XR1#show eigrp ipv4 neighbors
Tue Jul 21 23:10:06.998 UTC
```

CONTENTS ›

https://t.me/learningnets

```
IPv4-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   10.19.20.20             Gi0/0/0/0.1920    14 22:51:08   27   200  0  12


RP/0/0/CPU0:XR1#show route ipv4 eigrp

Tue Jul 21 23:10:18.877 UTC


D    20.20.20.20/32 [90/10752] via 10.19.20.20, 22:51:08, GigabitEthernet0/0/0/0.1920


RP/0/0/CPU0:XR1#show route ipv4 0.0.0.0

Tue Jul 21 23:11:27.103 UTC


Routing entry for 0.0.0.0/0

   Known via "static", distance 1, metric 0 (connected), candidate default path

  Installed Jul 21 22:58:41.595 for 00:12:45

  Routing Descriptor Blocks

    directly connected, via Null0

      Route metric is 0

  No advertising protos.


RP/0/0/CPU0:XR2#show route ipv4 eigrp

Tue Jul 21 23:09:29.855 UTC


 D*EX 0.0.0.0/0 [170/568320] via 10.19.20.19, 00:11:18, GigabitEthernet0/0/0/0.1920

D    19.19.19.19/32 [90/10752] via 10.19.20.19, 22:51:30, GigabitEthernet0/0/0/0.1920
```

Neither R2 or XR1 are advertising any of the site's routes (the EIGRP routes) into BGP. LISP will be used to provide reachability between these prefixes using an overlay.

LISP

The main idea behind LISP, the Locator ID Separation Protocol, is to decouple EID (Endpoint ID) prefixes (the site owned prefixes, consisting of either Provider Independent IP space in the public internet, or RFC-1918 space in a private VPN) from the RLOCs (Resource Locators), consisting of the IP addresses used to reach the EIDs. For example, a company with Provider Independent IP space would be able to treat that block of addresses as their EIDs - these are the addresses that define the endpoints within their network. The IP addressing on the links used to connect to this company's internet/service provider, which are usually /30s assigned by the provider, would make up the company's RLOCs. These Resource Locators are what the rest of the internet uses to reach the company's EIDs. However, both the RLOC and EID prefixes are advertised into the same routing system - traditionally BGP. The company we are referring to in this example would normally advertise its "EID" prefixes into BGP, and the provider would also advertise the /30s (aggregated at some point) into BGP.

CONTENTS ›

Our network has the following EID and RLOC breakdown:

- LISP Site A:
  - EIDs:
    - 10.1.0.0/16
    - 1.1.1.1/32
  - RLOCs:
    - 20.2.3.2
    - 20.2.4.2
- LISP Site B:

https://t.me/learningnets

  - EIDs:
    - 10.19.0.0/16
    - 20.20.20.20/32
  - RLOCs:
    - 20.5.19.19
    - 20.6.19.19

Note that the RLOC prefixes consist of the IP addresses used by R2 and XR1 to connect into the underlay, and EIDs consist of the site owned prefixes.

LISP is a tunneling/encapsulation approach that provides a level of indirection between the EID and RLOC space by using the following model:

- Remove the EID space from the global routing system (BGP), and advertise them into a a DNS-like database, called the LISP Mapping System, that can scale and uses a "pull" based model vs a "push" model offered by BGP.
  - Leave only the RLOC prefixes in the traditional routing system.
  - Doing this cleans up the global BGP table from all the pollution caused by recent de-aggregation of provider independent advertisements.
  - Yields several other benefits and use cases covered in sections to follow.
- The LISP Mapping System database that stores EIDs has one mayor function:
  - Map EID prefixes to the corresponding RLOCs for that site.
  - Hosts needing to reach an EID prefix can query the Mapping System for the site's RLOCs.

Traditional routing protocols build a topology from the information gathered from the routing protocol, compute and install best paths, and advertise this information to all neighbors. The information is being "pushed" towards all neighbors, without each neighbor actually requesting the routing information. LISP uses a pull model - a router needing to make a forwarding decision on a packet will "ask" the mapping database for the RLOCs associated with the destination EID prefix. The database responds with an answer and the router caches it. This is the "pull" model often described by LISP - where it resembles the architecture employed by DNS. Once the information is cached and programmed into the forwarding table, the router will perform normal packet forwarding for subsequent packets being forwarded to that same destination. In order words, the router requests the forwarding information it needs, instead of being given all of the routing information (BGP table) implicitly.

Since LISP is an encapsulation/tunneling technology used to build overlays, it can use any underlying transport, as long as there is IPv4/IPv6 reachability in the underlay. This is accomplished in our network as the SP core is running BGP with each site, and is redistributing the connected interfaces to each site into BGP. Note that this "underlay" is what is referred to as the RLOC space - the global BGP table. Like GRE or other tunneling techniques, we can tunnel one address-family over another. LISP uses a simple header, consisting of the outer IPv4/IPv6 header (RLOC source/destination), and a UDP + LISP header. For user data-traffic being encapsulated, LISP uses a destination UDP port of 4341. For the control-plane messaging between LISP nodes, UDP port 4342 is used. From an MTU perspective, LISP adds 36 Bytes of overhead for IPv4 (20 for IPv4 header, 8 for UDP, 8 for LISP), and 56 Bytes for IPv6.

CONTENTS ›

LISP hashes the payload of the inner user packet into the source UDP port of the LISP header, which accomplishes what is sometimes referred to as "payload entropy". This allows the core underlay network to perform efficient load sharing if there are port-channels/LAGs or ECMP throughout the core using standard 5-tuple load sharing (src & dst ip, protocol, src & dst port). If both source/dest LISP header UDP ports were fixed, the underlay network's load sharing hashing algorithm would not perform as well, as all traffic between two RLOCs would look as if it were a single UDP flow in the core.

Instance ID - portion of the LISP header that can be used to add a layer of virtualization to the LISP dataplane. This is similar to what is accomplished in a L3VPN network with MPLS and VRFs, however the customer, instead of the service provider, is in charge of creating the virtual networks. The instance-id has multiple use cases - a simple example is emmulating a "vlan tag".

Role Definitions

https://t.me/learningnets

There are a few roles that are defined by LISP:

- **iTR** - Ingress tunnel router. This device is in charge of taking LISP packet from the site and LISP encapsulating them towards their destination. An iTR is also in charge of querying the mapping database when trying to make a forwarding decision. The query is sent to the Map Resolver using a Map-Request message. Just like in DNS, queries are sent to DNS resolvers.
- **eTR** - Egress tunnel router - This device is in charge of receiving LISP encapsulated packets from the underlay network, decapsulating them, and forwarding them towards the final destination inside the site. A second forwarding decision is made by this router after the decapsulation. An eTR is in charge of registering the site's EID prefixes with the Map-Server. This is done by building a local database on the router, commonly populated by the use of `database-mapping` commands. For example, R2 has `database-mapping` entries for 10.1.2.0/24, which is the EID prefix belonging to the site. Note that this prefix does not need to be directly connected to the eTR - R2 could have learned this via the site's IGP. Note that R2 is able to register R1's Loopback0 into the mapping system. A newer method exists of populating the eTR's local database, which consists of redistributing the routes installed via the site's IGP in the RIB into LISP. On IOS this is done using the `ipv4 route-import database` command. This allows the population of the database to be more dynamic, instead of relying on static entries. For example, R2 could `import` the EIGRP learned routes into LISP, and thus cause them to get registered with the mapping-system without manually entering them in.

As soon as the eTR's local database is populated, it is then registered with the Map Server using a LISP Map-Register packet. The Map-Register message maps a site's EIDs to RLOCs, with the corresponding RLOC weights and priorities. Just like in DNS, the site making the registration is authoritative for the EIDs its registering.

- **MS** - Map Server - This device holds all of the EID to RLOC mappings for all of the sites it is responsible for. The Map Server takes registrations from eTR routers via Map-Register messages.
- **MR** - Map Resolver - This device is queried by the iTRs using Map-Request messages. The queries consist of iTRs asking for the mapping between EID and RLOC.

There are a two other roles that will be discussed in later sections, namely PiTR and PeTR - Proxy Ingress Tunnel Router, and Proxy Egress Tunnel Router.

Note:

A Note iTR and eTR: Usually iTR and eTR functions are collocated on the same router, collectively called xTR. This is the case for R2 and XR1. **A Note MS and MR**: Usually MS and MR functions are also collocation on the same device (router/server). R9 in our network is acting as a MS and MR.

xTR Configuration Breakdown

Enable the LISP process, the same way any other routing protocol would be enabled. More than one LISP process can be enabled, specified by the process ID.

```
router lisp
```

Create a `locator-set` specifying the RLOCs used to reach the EIDs at the LISP site. Each locator is configured with its corresponding `weight` and `priority` value. Note that the interface name could be used instead of the IP address. This is useful for sites using dynamic IP addressing for RLOCs (DLS, Cable, etc)

CONTENTS ›

https://t.me/learningnets

```
locator-set SITEA_RLOC_SET
 20.2.3.2 priority 1 weight 50
 20.2.4.2 priority 1 weight 50
 exit
!
```

Create a local database mapping of the site owned EIDs. Associate an RLOC, or a `locator-set` with each database entry. These entries will be advertised to the Map Server.

```
database-mapping 1.1.1.1/32 locator-set SITEA_RLOC_SET
database-mapping 10.1.2.0/24 locator-set SITEA_RLOC_SET
```

Configure the xTR roles: iTR and eTR. Point the iTR towards the map-resolver that will be queried for remote EIDs, and pint to eTR towards the map-server that will be used to register the site's EID/RLOC mappings.

```
ipv4 itr map-resolver 9.9.9.9
ipv4 itr
ipv4 etr map-server 9.9.9.9 key SITEA_xTR-R2
ipv4 etr
exit
```

LISP Verification

Lets begin the verification process by testing reachability between R1 and XR2 over LISP. At this point, both devices should have connectivity between their connected interfaces and their loopback0.

R2 and XR1 start out with an empty map-cache. That is, they have no information about remote EIDs. Both devices have a 'default' entry, with an action stating `send map-request` . This will cause R2 and XR1 to query the mapping system for unknown destinations.

```
R2#show ip lisp map-cache
LISP IPv4 Mapping Cache for EID-table default (IID 0), 1 entries

 0.0.0.0/0 , uptime: 00:00:29, expires: never, via static  send map-request
   Negative cache entry, action: send-map-request



RP/0/0/CPU0:XR1#show lisp ipv4 map-cache
Sat Jul 25 15:13:36.341 UTC


LISP IPv4 Mapping Cache for EID-table default (IID 0), 1 entries

 0.0.0.0/0 , uptime: 00:00:03, expires: never, via static  send map-request
   Negative cache entry, action: send-map-request
```

Pings between the devices are successful. Notice that the first two packets are dropped when we first try to use LISP for forwarding towards a particular destination.

CONTENTS ›

https://t.me/learningnets

```
R1#ping 20.20.20.20 source loopback 0

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 20.20.20.20, timeout is 2 seconds:

Packet sent with a source address of 1.1.1.1

 ..!!!

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


R1#ping 10.19.20.20

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.19.20.20, timeout is 2 seconds:

 ..!!!

Success rate is 60 percent (3/5), round-trip min/avg/max = 3/3/4 ms


R1#traceroute 20.20.20.20 source loopback 0

Type escape sequence to abort.

Tracing the route to 20.20.20.20

VRF info: (vrf in name/id, vrf out name/id)

  1 10.1.2.2 4 msec 1 msec 1 msec

  2 20.2.4.4 2 msec 2 msec 2 msec

  3 20.4.5.5 1 msec 7 msec 12 msec

  4 20.5.19.19 13 msec 15 msec 16 msec

  5 10.19.20.20 16 msec *  4 msec
```

The map-caches on R2 and XR1 have been populated with the remote EID entries. This is what allowed the forwarding to occur.

CONTENTS ›

```
R2#show ip lisp map-cache
LISP IPv4 Mapping Cache for EID-table default (IID 0), 3 entries


0.0.0.0/0, uptime: 00:04:32, expires: never, via static send map-request
  Negative cache entry, action: send-map-request
 10.19.20.0/24 , 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
 20.20.20.20/32 , uptime: 00:00:13, expires: 23:59:46, via map-reply, complete
  Locator     Uptime    State      Pri/Wgt
  20.5.19.19  00:00:13  up           1/50
  20.6.19.19  00:00:13  up           1/50


RP/0/0/CPU0:XR1#show lisp ipv4 map-cache
Sat Jul 25 15:16:56.378 UTC




LISP IPv4 Mapping Cache for EID-table default (IID 0), 3 entries


0.0.0.0/0, uptime: 00:03:23, expires: never, via static send map-request
  Negative cache entry, action: send-map-request
 1.1.1.1/32 , uptime: 00:00:24, expires: 23:59:37, via map-reply, complete
  Locator   Uptime    State      Pri/Wgt
  20.2.3.2  00:00:24  up           1/50
  20.2.4.2  00:00:24  up           1/100
 10.1.2.0/24 , uptime: 00:00:18, expires: 23:59:43, via map-reply, complete
  Locator   Uptime    State      Pri/Wgt
  20.2.3.2  00:00:18  up           1/100
  20.2.4.2  00:00:18  up           1/100
```

Note however that the SP network (RLOC space) does not have forwarding information for these routes.

```
R4#show ip route 20.20.20.20
 % Subnet not in table


R4#show ip cef 20.20.20.20
0.0.0.0/0
   no route
```

To break down the underling control-plane messaging that enabled forwarding between LISP sites to take place, we can begin by validating the registrations from each LISP site to the map server (R9).

›
CONTENTS

https://t.me/learningnets

```
R9#show lisp site

LISP Site Registration Information


Site Name      Last      Up   Who Last            Inst     EID Prefix

               Register       Registered          ID

SITE_A         00:00:08   yes  20.2.4.2                     1.1.1.1/32

               never     no   --                           10.1.0.0/16

               00:00:08   yes  20.2.4.2                     10.1.2.0/24

SITE_B         never     no   --                           10.19.0.0/16

               00:00:55   yes  20.5.19.19                   10.19.20.0/24

               00:00:55   yes  20.5.19.19                   20.20.20.20/32
```

R9 sees the registrations from each site properly. Note that there are entries for 10.1.0.0/16 and 10.19.0.0/16, the aggregate EID blocks of each site. The Mapping System was configured to allow any more specific entries to be registered from each site's /16, as per the task requirements. This helps keeps the Mapping Server configuration small.

A more verbose output shows us both RLOCs associated with a particular EID, and their corresponding weights and priorities (covered later in this section).

CONTENTS ›

https://t.me/learningnets

```
R9#show lisp site name SITE_A 10.1.2.0/24
LISP Site Registration Information

Site name: SITE_A
Allowed configured locators: any
Requested EID-prefix:
  EID-prefix: 10.1.2.0/24
    First registered:     2d23h
    Routing table tag:    0
    Origin:               Dynamic, more specific of 10.1.0.0/16
    Merge active:         No
    Proxy reply:          No
     TTL:                 1d00h
     State:               complete
    Registration errors:
      Authentication failures:   0
      Allowed locators mismatch: 0
     ETR 20.2.4.2, last registered 00:00:17 , no proxy-reply, map-notify
                  TTL 1d00h, no merge, hash-function sha1, nonce 0x0408D118-0x5409EE53
                  state complete, no security-capability
                  xTR-ID 0xFB2B2554-0x18B820E2-0xCE764CBD-0x36F371B7
                  site-ID unspecified
      Locator   Local  State      Pri/Wgt  Scope
      20.2.3.2  yes    up           1/50   IPv4 none
      20.2.4.2  yes    up           1/50   IPv4 none


R9#show lisp site name SITE_B 10.19.20.0/24
LISP Site Registration Information

Site name: SITE_B
Allowed configured locators: any
Requested EID-prefix:
  EID-prefix: 10.19.20.0/24
    First registered:     2d22h
    Routing table tag:    0
    Origin:               Dynamic, more specific of 10.19.0.0/16
    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 00:01:03 , no proxy-reply, map-notify
                    TTL 1d00h, no merge, hash-function sha1, nonce 0xA95FDAB4-0x2F801B19
                    state complete, no security-capability
                    xTR-ID 0x0404430A-0x83067718-0x0D293195-0xAE564F7A
                    site-ID unspecified
      Locator     Local  State      Pri/Wgt  Scope
      20.5.19.19  yes    up           1/50   IPv4 none
      20.6.19.19  yes    up           1/50   IPv4 none
```

CONTENTS ›

https://t.me/learningnets

The xTR routers create local databases for the EID prefixes being registered. These local database mappings match what the Map Server is seeing.

```
R2#show ip lisp database

LISP ETR IPv4 Mapping Database for EID-table default (IID 0), LSBs: 0x3, 2 entries


 1.1.1.1/32 , 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

 10.1.2.0/24 , 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 ipv4 database

Thu Jul 23 23:48:02.694 UTC




LISP ETR IPv4 Mapping Database for EID-table default (IID 0), LSBs: 0x3, 2 entries


 10.19.20.0/24 , 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

 20.20.20.20/32 , 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
```

To summarize, each xTR advertises to the mapping system which EIDs it owns and what RLOCs each EID is reachable via.

Packet Walk - LISP to LISP

A packet flow between Site-A and Site-B will be explained below. For this exercise, R1 is going to send a packet to XR2 we saw in the previous output. The source will be 10.1.2.1 towards 10.19.20.20.

Note:

Assumption: There is reachability in the core network - R2 can reach XR1, and both R1 and XR1 can reach the mapping system through the core.

CONTENTS ›

1. R1 sends a ping to 10.19.20.19, which is forwarded towards R2 following the default route. R2 does not have a route for this destination in the RIB, and it does not have a LISP map-cache entry for it either, so its iTR LISP role sends a Map-Request towards the Map-Resolver - 4.4.4.4 (R4).
    - While the request is made to the Mapping System, packets are dropped as R2 does not have enough information to forward the packet. This is what caused the first two packets to be dropped in our earlier output.
2. The Map Resolver receives this Map-Request and passes it to the Map-Server portion of the LISP code for processing.

3. The Map Server performs a longest match lookup in its database against 10.19.20.20 in order to determine the site owning the EID, and finds that XR1's site registered EID 10.19.20.0/24. Since XR1 is authoritative for this EID, the Map Server forwards the Map-Request to XR1 so that it can give an authoritative answer.
4. XR1 receives the Map-Request, and its eTR role directly replies back to iTR router R2 using a Map-Reply message, with the mapping between 10.19.20.0/24 and XR1's RLOCs - 20.5.19.19 and 20.6.19.19. Both of them having a priority of 1 and a weight of 50 - signaling a desired 1:1 ratio. The Map-Reply also has a TTL field, denoting the amount of time that R2 should cache the entry for (similar to a DNS TTL).
5. iTR R2 receives the Map-Reply from the eTR XR1 and caches an entry, as seen by `show ip lisp map-cache` . This entry gets programmed in the hardware forwarding tables [FIB].
6. During all this control-plane signaling taking place, R1's pings towards 10.19.20.20 have still been going on. The first 2 pings timed out, but the last 3 were successful. After R2 caches the entry, it is able to encapsulate the packets towards XR1's RLOCs, and normal packet forwarding takes place.

This same list of steps also took place in the reverse direction. XR2 send the echo reply to 10.1.2.1, but XR1 had no routing information or map-cache entry for it, so it engages the mapping system and 'pulls' down the mapping for EID prefix 10.1.2.0/24.

Taking a look at the cached entry, we can see that R2 learned about 10.19.20.0/24 and 20.20.20.20/32 after it requested this information from the mapping system.

```
R2#show ip lisp map-cache
LISP IPv4 Mapping Cache for EID-table default (IID 0), 3 entries

0.0.0.0/0, uptime: 3d00h, expires: never, via static send map-request
  Negative cache entry, action: send-map-request
 10.19.20.0/24 , uptime: 00:23:33, expires: 23:36:26,  via map-reply, complete
  Locator     Uptime    State      Pri/Wgt
   20.5.19.19  00:23:33  up           1/50
   20.6.19.19  00:23:33  up           1/50
 20.20.20.20/32 , uptime: 00:23:54, expires: 23:36:06,  via map-reply, complete
  Locator     Uptime    State      Pri/Wgt
   20.5.19.19  00:23:54  up           1/50
   20.6.19.19  00:23:54  up           1/50
```

Similarly XR1 queried the mapping system for 10.1.2.0/24 and 1.1.1.1/32 as XR2 replied to R1's pings.

CONTENTS ›

https://t.me/learningnets

```
RP/0/0/CPU0:XR1#show lisp ipv4 map-cache
Fri Jul 24 00:06:38.048 UTC



LISP IPv4 Mapping Cache for EID-table default (IID 0), 3 entries


0.0.0.0/0, uptime: 2d23h, expires: never, via static send map-request
  Negative cache entry, action: send-map-request
 1.1.1.1/32 , uptime: 00:26:30, expires: 23:35:46,  via map-reply, complete
  Locator   Uptime    State      Pri/Wgt
   20.2.3.2  00:26:30  up           1/50
   20.2.4.2  00:26:30  up           1/50
 10.1.2.0/24 , uptime: 00:26:09, expires: 23:36:04,  via map-reply, complete
  Locator   Uptime    State      Pri/Wgt
   20.2.3.2  00:26:09  up           1/50
   20.2.4.2  00:26:09  up           1/50
```

This information is then programmed in CEF, allowing the router to perform normal hardware forwarding.

```
R2#show ip cef 10.19.20.0/24 detail
10.19.20.0/24, epoch 2, flags [attached, subtree context, check lisp eligibility, default route]
  SC owned,sourced: LISP remote EID - locator status bits 0x00000003
   LISP remote EID : 4 packets 472 bytes  fwd action encap
  LISP source path list
     nexthop 20.5.19.19 LISP0
     nexthop 20.6.19.19 LISP0
  2 IPL sources [active source]
    Dependent covered prefix type inherit, cover 0.0.0.0/0
  recursive via 0.0.0.0/0
    attached to Null0
```

XR1 shows more information about what is programmed in hardware.

CONTENTS >

https://t.me/learningnets

```
RP/0/0/CPU0:XR1#show cef 10.1.2.0/24 detail
Fri Jul 24 00:09:02.968 UTC
10.1.2.0/24, version 0, internal 0x4000001 0x0 (ptr 0xa0edc474) [1], 0x0 (0x0), 0x0 (0x0)
 Updated Jul 23 23:40:28.375
 Prefix Len 24, traffic index 0, precedence n/a, priority 0
  gateway array (0xa0d30388) reference count 2, flags 0x4010, source lisp lo (7), 0 backups
                [1 type 3 flags 0x10048289 (0xa0df1938) ext 0x0 (0x0)]
  LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
   via 20.2.3.2, 4 dependencies, recursive [flags 0x1000]
    path-idx 0 NHID 0x0 [0xa162f050 0x0]
     next hop 20.2.3.2 via 20.2.3.2
     IP ENCAP (3 headers), locks: 4
      Transport: ipv4 (vrf 'default': tbl ID 0xe0000000)
      Parent: DATA_TYPE_LEAF, 0xa0edc274
       Payload AF/MTU: ipv4/1456
      HDR: IP4
        src: 20.5.19.19
        dst: 20.2.3.2
        prot: 17, ttl: dyn, 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 20.2.4.2, 4 dependencies, recursive [flags 0x1000]
    path-idx 1 NHID 0x0 [0xa162f0e4 0x0]
     next hop 20.2.4.2 via 20.2.4.2
     IP ENCAP (3 headers), locks: 4
      Transport: ipv4 (vrf 'default': tbl ID 0xe0000000)
      Parent: DATA_TYPE_LEAF, 0xa0edc2f4
      Payload AF/MTU: ipv4/1456
      HDR: IP4
        src: 20.5.19.19
        dst: 20.2.4.2
        prot: 17, ttl: dyn, 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                 20.2.3.2
    1     Y   recursive                 20.2.4.2
```

When R2 forwards traffic towards 10.19.20.0/24, it uses the RLOCs as the destination of the outer header. To the core network this just looks like a UDP packet destined to one of XR1's uplinks.

https://t.me/learningnets

Lets perform a packet capture on R3 to see the details. 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.

```
R2#clear ip lisp map-cache


RP/0/0/CPU0:XR1#clear lis ipv4 map-cache
Sat Jul 25 15:36:33.437 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 20.20.20.20 source loopback 0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.20.20.20, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 2/3/4 ms
R1#


R3#monitor capture LISP_CAP export ftp://cisco:cisco@169.254.254.1/lisp.v4.to.v4.001.pcap
Writing lisp.v4.to.v4.001.pcap
Exported Successfully
Notice the sequence of events detailed in the packet capture.
```

CONTENTS ›

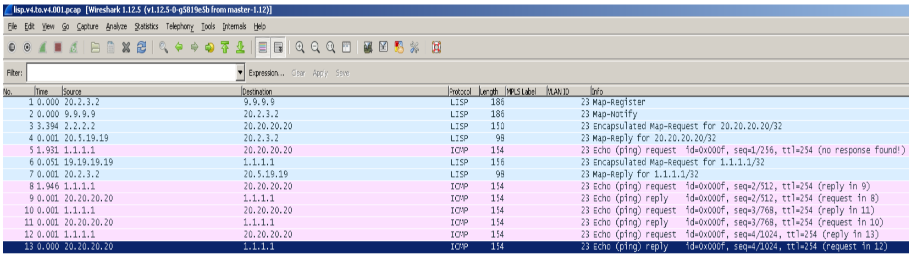

```
No.  Time  Source       Destination  Protocol Length MPLS Label VLAN ID Info
  1 0.000 20.2.3.2      9.9.9.9       LISP   186   23 Map-Register
  2 0.000 9.9.9.9       20.2.3.2      LISP   186   23 Map-Notify
  3 3.394 2.2.2.2       20.20.20.20   LISP   150   23 Encapsulated Map-Request for 20.20.20.20/32
  4 0.001 20.5.19.19    20.2.3.2      LISP    98   23 Map-Reply for 20.20.20.20/32
  5 1.931 1.1.1.1       20.20.20.20   ICMP   154   23 Echo (ping) request  id=0x000f, seq=1/256, ttl=254 (no response found!)
  6 0.051 19.19.19.19   1.1.1.1       LISP   156   23 Encapsulated Map-Request for 1.1.1.1/32
  7 0.001 20.2.3.2      20.5.19.19    LISP    98   23 Map-Reply for 1.1.1.1/32
  8 1.946 1.1.1.1       20.20.20.20   ICMP   154   23 Echo (ping) request  id=0x000f, seq=2/512, ttl=254 (reply in 9)
  9 0.001 20.20.20.20   1.1.1.1       ICMP   154   23 Echo (ping) reply    id=0x000f, seq=2/512, ttl=254 (request in 8)
 10 0.001 1.1.1.1       20.20.20.20   ICMP   154   23 Echo (ping) request  id=0x000f, seq=3/768, ttl=254 (reply in 11)
 11 0.001 20.20.20.20   1.1.1.1       ICMP   154   23 Echo (ping) reply    id=0x000f, seq=3/768, ttl=254 (request in 10)
 12 0.001 1.1.1.1       20.20.20.20   ICMP   154   23 Echo (ping) request  id=0x000f, seq=4/1024, ttl=254 (reply in 13)
 13 0.000 20.20.20.20   1.1.1.1       ICMP   154   23 Echo (ping) reply    id=0x000f, seq=4/1024, ttl=254 (request in 12)
```

Packet #	**Description**
1	R2 Registers its EIDs with the mapping-system
3	R2 sends Map-Request to the Map-Server for 20.20.20.20
4	XR1, the authoritative owner of 20.20.20.20, replies directly to R2 with a Map-Reply
5	ICMP Echo is able to be forwarded from R1 to XR2, after R2 has the RLOCS for 20.20.20.20
6	Map-Request for 1.1.1.1 is sent to R2
7	R2 replies back directly to XR1, since R2 is the authoritative owner of 1.1.1.1
8-13	ICMP echos/replies can now be forwarded

Zooming in on packet #13, we can see the encapsulated ICMP/IP headers from the user payload. Just like a L3VPN or any other VPN type solution, the payload is sent in clear text unless otherwise encrypted.

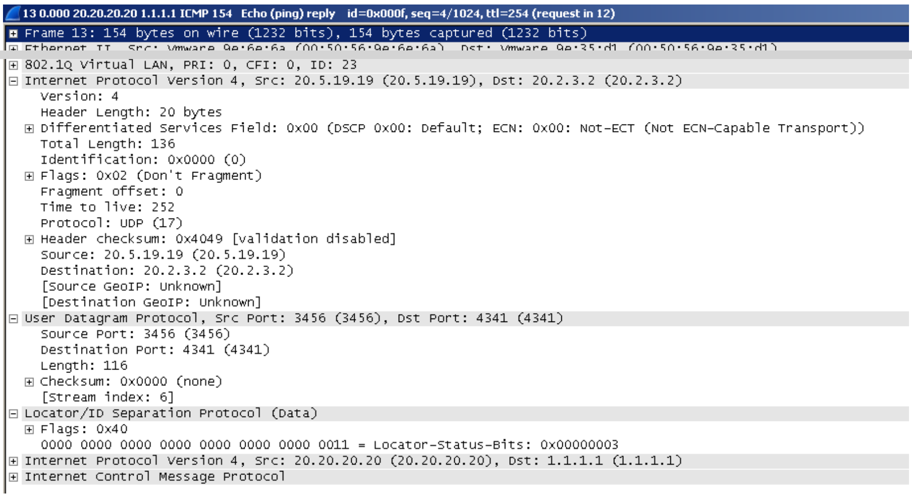

```
13 0.000 20.20.20.20 1.1.1.1 ICMP 154  Echo (ping) reply   id=0x000f, seq=4/1024, ttl=254 (request in 12)
⊞ Frame 13: 154 bytes on wire (1232 bits), 154 bytes captured (1232 bits)
⊞ Ethernet II, Src: Vmware_9e:6e:6a (00:50:56:9e:6e:6a), Dst: Vmware_9e:35:d1 (00:50:56:9e:35:d1)
⊞ 802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 23
⊟ Internet Protocol Version 4, Src: 20.5.19.19 (20.5.19.19), Dst: 20.2.3.2 (20.2.3.2)
    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: 0x0000 (0)
  ⊞ Flags: 0x02 (Don't Fragment)
    Fragment offset: 0
    Time to live: 252
    Protocol: UDP (17)
  ⊞ Header checksum: 0x4049 [validation disabled]
    Source: 20.5.19.19 (20.5.19.19)
    Destination: 20.2.3.2 (20.2.3.2)
    [Source GeoIP: Unknown]
    [Destination GeoIP: Unknown]
⊟ User Datagram Protocol, Src Port: 3456 (3456), Dst Port: 4341 (4341)
    Source Port: 3456 (3456)
    Destination Port: 4341 (4341)
    Length: 116
  ⊞ Checksum: 0x0000 (none)
    [Stream index: 6]
⊟ Locator/ID Separation Protocol (Data)
  ⊞ Flags: 0x40
    0000 0000 0000 0000 0000 0000 0000 0011 = Locator-Status-Bits: 0x00000003
⊞ Internet Protocol Version 4, Src: 20.20.20.20 (20.20.20.20), Dst: 1.1.1.1 (1.1.1.1)
⊞ Internet Control Message Protocol
```

It is important to note that LISP forwarding can coexist with traditional forwarding as well. If an xTR receives a packet from its local site destined to a prefix known via a non default route (via IGP/BGP, etc), then it will simply forward the packet using the traditional longest match based forwarding paradigm. In order for the route look-up to be picked up by the LISP sub-system, an xTR must either:

A) Have NO routing information to the destination - which is true in our case. OR: B) Have a default route as a 'catch-all' for that destination. Instead of running IGP/BGP with the core underlay, an xTR simply use default static routes pointing towards their upstream next-hop, and this will be enough to trigger the LISP code to pick up the look-up and engage the LISP system for forwarding.

CONTENTS ›

Note:

Note: The last option of only having a default route works on IOS, IOS-XE, and NX-OS. IOS-XR does not seem to have the same forwarding behavior in the code versions used during this lab (5.2.0).

This greatly simplifies simple deployments consisting of a dual homed sites to multiple providers. Instead of running BGP with each provider and dealing with the complexities of traffic engineering and other best practices, the site could simply use static default routes out towards the provider and use LISP for forwarding. There are other LISP components discussed in later sections which allow a LISP site to do this, even if it requires connectivity to both LISP and non-LISP sites (the internet for example). This is accomplished by using a PxTR (Proxy Ingress/Egress Tunnel Router), either hosted by a provider and offered as a service, or hosted by the LISP site with the gear in another facility, such as colo site.

LISP Ingress Traffic Engineering

Since the LISP forwarding model is based on a pull rather than push model, and this information comes from a common central mapping database, performing traffic engineering becomes much simpler. Take the following example:

XR1 connects into the network via its Gig0/0/0/0.619 and Gig0/0/0/0.519. For the sake of the argument, lets say both of these links are full Gig links. Since both links are equal, we would want traffic leaving the site via XR1, and traffic ingressing the site via XR1 to be load shared 50/50 (or 1:1) between the two links. R2 connects into the underlay network using Gig1.12 and Gig1.23. Lets pretend Gig1.24 is a 100 Mbps link, and Gig1.23 is a 50 Mbps link. Since these links are not equal, traditionally a site would either do an active/standby type of design for outbound/inbound flows, where only the 100Mbps is used unless there is a failure, or would try to do unequal cost load sharing using a 1:2 ratio - for every 2 flows sent out of the 100Mbps link, send 1 out of the 50 Mbps link. Although accomplishing this 1:2 split is technically possible, it is overly complex to accomplish using traditional tools. If using BGP, one way of accomplishing this would consist of receiving a default via one link, and a few specific routes on the other link. There are other methods (PfR for example), but its very difficult to arrive close to the desired 1:2. In the inbound direction, one could advertise a few /24s out of one link, and more /24s out of the other link, while advertising the aggregate of all the subnets out of both links. The problem again is that arriving at a desired 1:2 ratio is very difficult and complex. The site would need to own multiple PI subnets to begin with.

LISP aims at solving the ingress load sharing problem. When a site registers its EID to RLOC mappings to the mapping database, it can associate different priorities/weights (i,e - metrics) with each RLOC to signal to other sites the desired ingress traffic ratio. For example, when XR1's site wants to reach 10.1.2.0/24 EID prefix located behind R2's RLOCs, it queries the mapping database for the locators that XR1 needs to encapsulate towards. The mapping system responds with 20.2.3.2 and 20.2.4.2, the two RLOCs that R2 registered for the EID. If R2 registered these RLOCS with metrics resembling the desired ingress load sharing state, such as Gig1.12 with a weight of 100 and Gig1.23 with a weight of 50 (1:2 ratio), XR1 will forward traffic towards the EID following this ratio. Out of 150 flows destined to 10.1.2.0/24, XR1 will encapsulate 100 of them towards 20.2.4.2, and 50 towards 20.2.3.2.

Using this model, it becomes very simple for any site to signal how it wants to receive packets to the rest of the sites in the LISP network, and not have to worry about the complexities of traditional ingress traffic engineering policies.

Lets make this change on R2 and inspect the forwarding tables on XR1.

CONTENTS ›

```
R2(config)#router lisp
R2(config-router-lisp)# locator-set SITEA_RLOC_SET
R2(config-router-lisp-locator-set)#20.2.4.2 priority 1 weight 100
R2(config-router-lisp-locator-set)#end



RP/0/0/CPU0:XR1#clear lisp ipv4 map-cache
Fri Jul 24 01:01:26.953 UTC


RP/0/0/CPU0:XR2#ping 10.1.2.1
Fri Jul 24 01:01:30.557 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.2.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/5/19 ms
```

Notice that the change is reflected when XR1 refreshes its cache.

```
RP/0/0/CPU0:XR1#show lisp ipv4 map-cache
Fri Jul 24 01:02:44.567 UTC


LISP IPv4 Mapping Cache for EID-table default (IID 0), 2 entries

0.0.0.0/0, uptime: 00:00:22, expires: never, via static send map-request
  Negative cache entry, action: send-map-request
 10.1.2.0/24 , uptime: 00:00:14, expires: 23:59:47, via map-reply, complete
  Locator   Uptime    State      Pri/Wgt
   20.2.3.2  00:00:14  up           1/50
   20.2.4.2  00:00:14  up           1/100
```

R2's ingress load balancing intent is obeyed by XR1, and it distributes the flows in a perfect 1:2 ratio. R2 is able to influence how traffic is received at LISP Site A by simply setting the priorities on its RLOCs.

CONTENTS ›

https://t.me/learningnets

```
RP/0/0/CPU0:XR1#show cef 10.1.2.0/24 detail
Fri Jul 24 01:04:21.591 UTC
10.1.2.0/24, version 0, internal 0x4000001 0x0 (ptr 0xa0edc574) [1], 0x0 (0x0), 0x0 (0x0)
 Updated Jul 24 01:02:30.729
 Prefix Len 24, traffic index 0, precedence n/a, priority 0
  gateway array (0xa0d30388) reference count 1, flags 0x4010, source lisp lo (7), 0 backups
                [1 type 3 flags 0x10048089 (0xa0df1938) ext 0x0 (0x0)]
  LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
   via 20.2.3.2, 3 dependencies, recursive [flags 0x1000]
    path-idx 0 NHID 0x0 [0xa162f050 0x0]
    next hop 20.2.3.2 via 20.2.3.2
    IP ENCAP (3 headers), locks: 3
      Transport: ipv4 (vrf 'default': tbl ID 0xe0000000)
      Parent: DATA_TYPE_LEAF, 0xa0edc274
      Payload AF/MTU: ipv4/1456
      HDR: IP4
        src: 20.5.19.19
        dst: 20.2.3.2
        prot: 17, ttl: dyn, 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 20.2.4.2, 3 dependencies, recursive [flags 0x1000]
    path-idx 1 NHID 0x0 [0xa162f0e4 0x0]
    next hop 20.2.4.2 via 20.2.4.2
    IP ENCAP (3 headers), locks: 3
      Transport: ipv4 (vrf 'default': tbl ID 0xe0000000)
      Parent: DATA_TYPE_LEAF, 0xa0edc2f4
      Payload AF/MTU: ipv4/1456
      HDR: IP4
        src: 20.5.19.19
        dst: 20.2.4.2
        prot: 17, ttl: dyn, 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 100, normalized_weight 2, class 0

    Load distribution: 0 1 1 (refcount 1)

    Hash  OK  Interface                 Address
     0     Y   recursive                 20.2.3.2
     1     Y   recursive                 20.2.4.2
     2     Y   recursive                 20.2.4.2
```

CONTENTS ›

https://t.me/learningnets

Priorities can be configured between 0-255, with a lower value being more preferred. If both priorities are equal, then both RLOCs are used. If the priorities are not equal, then only the RLOC with the lowest priority is used (Active/Standby). When the priorities are equal, weights can be configured to signal the desired ingress load sharing. Weights can be configured between 0-100.

« Inter-AS MPLS TE (/workbook/view/service-provider-v4/task/inter-as-mpls-te-Mjg4OQ%3D%3D) | LISP - IPv6 over IPv4 core (/workbook/view/service-provider-v4/task/lisp-ipv6-over-ipv4-core-Mjk1Mg%3D%3D) »

CONTENTS ›