Breaking Down the Task:

- For each host, ping its respective gateway for Edge node host discovery. In a real network the gratuitous ARP would be done by the host when it boots, so this is only required in a lab setting.
- Verify each Edge node has a LISP database entry for the attached host (omitted as this was shown earlier)
- On the Border-CP routers, run the command *show lisp site* and ensure that each host at its site is registered and has an RLOC entry (omitted as this was shown earlier)

This task is just ensuring the attached hosts are live, active, and discovered by the Edge nodes in preparation for sending traffic between them. If the Edge node ARP entry for the host times out or the host otherwise becomes unreachable, the Edge node will un-register the EID.

- Ping between the hosts at Site 1, and the hosts at Site 2. Each host should be able to reach the other host at its own site at this time using LISP as the control plane and VXLAN data plane encapsulation. This should be successful.

Here it is, the moment of truth. Will LISP with VXLAN work to send traffic between hosts attached to different Edge nodes? Remember that the Edge nodes are not aware of how to reach the hosts directly and despite sharing an Anycast gateway are not connected to both hosts by way of a switch or intervening Layer 2 extension.

How will this process work?

To investigate further, we should look at the LISP process at work and capture some packets.

Here is the capture from SITE1-EDGE 1 interface facing the client when the client sends an ARP request for the other host in the same subnet.

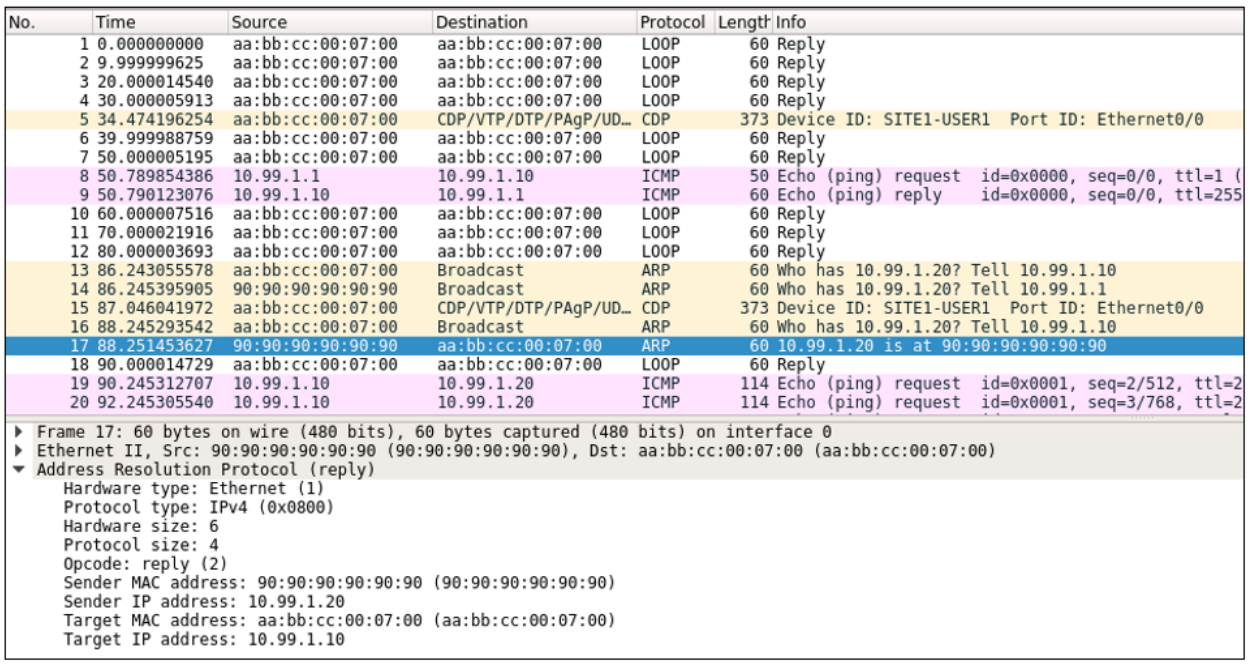

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
2	9.999999625	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
3	20.000014540	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
4	30.000005913	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
5	34.474196254	aa:bb:cc:00:07:00	CDP/VTP/DTP/PAgP/UD…	CDP	373	Device ID: SITE1-USER1 Port ID: Ethernet0/0
6	39.999988759	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
7	50.000005195	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
8	50.789854386	10.99.1.1	10.99.1.10	ICMP	50	Echo (ping) request id=0x0000, seq=0/0, ttl=1 (
9	50.790123076	10.99.1.10	10.99.1.1	ICMP	60	Echo (ping) reply id=0x0000, seq=0/0, ttl=255
10	60.000007516	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
11	70.000021916	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
12	80.000003693	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
13	86.243055578	aa:bb:cc:00:07:00	Broadcast	ARP	60	Who has 10.99.1.20? Tell 10.99.1.10
14	86.245395905	90:90:90:90:90:90	Broadcast	ARP	60	Who has 10.99.1.20? Tell 10.99.1.1
15	87.046041972	aa:bb:cc:00:07:00	CDP/VTP/DTP/PAgP/UD…	CDP	373	Device ID: SITE1-USER1 Port ID: Ethernet0/0
16	88.245293542	aa:bb:cc:00:07:00	Broadcast	ARP	60	Who has 10.99.1.20? Tell 10.99.1.10
17	88.251453627	90:90:90:90:90:90	aa:bb:cc:00:07:00	ARP	60	10.99.1.20 is at 90:90:90:90:90:90
18	90.000014729	aa:bb:cc:00:07:00	aa:bb:cc:00:07:00	LOOP	60	Reply
19	90.245312707	10.99.1.10	10.99.1.20	ICMP	114	Echo (ping) request id=0x0001, seq=2/512, ttl=2
20	92.245305540	10.99.1.10	10.99.1.20	ICMP	114	Echo (ping) request id=0x0001, seq=3/768, ttl=2

```
▸ Frame 17: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
▸ Ethernet II, Src: 90:90:90:90:90:90 (90:90:90:90:90:90), Dst: aa:bb:cc:00:07:00 (aa:bb:cc:00:07:00)
▾ Address Resolution Protocol (reply)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: reply (2)
    Sender MAC address: 90:90:90:90:90:90 (90:90:90:90:90:90)
    Sender IP address: 10.99.1.20
    Target MAC address: aa:bb:cc:00:07:00 (aa:bb:cc:00:07:00)
    Target IP address: 10.99.1.10
```

The host ARPs directly for the other host IP instead of just sending the packet to the default gateway because that is the default behavior of two hosts in the same network subnet. We know that the ARP will never reach that host because they are not attached to the same Layer 2 segment.

This is why the Edge nodes perform proxy ARP and inform the host that the MAC address for the other host is its own gateway MAC, and this is what makes Anycast gateway (and LISP mobility) work in the campus fabric setup.

The Edge node will consult its local LISP map-cache and LISP database, and, finding no entry for the other host in either, will send a Map Request to the configured LISP MR, the Border-CP node.

Now let us look at the capture from the Border-CP node where the LISP Map Request would be sent, and where the LISP Map Reply will be generated.

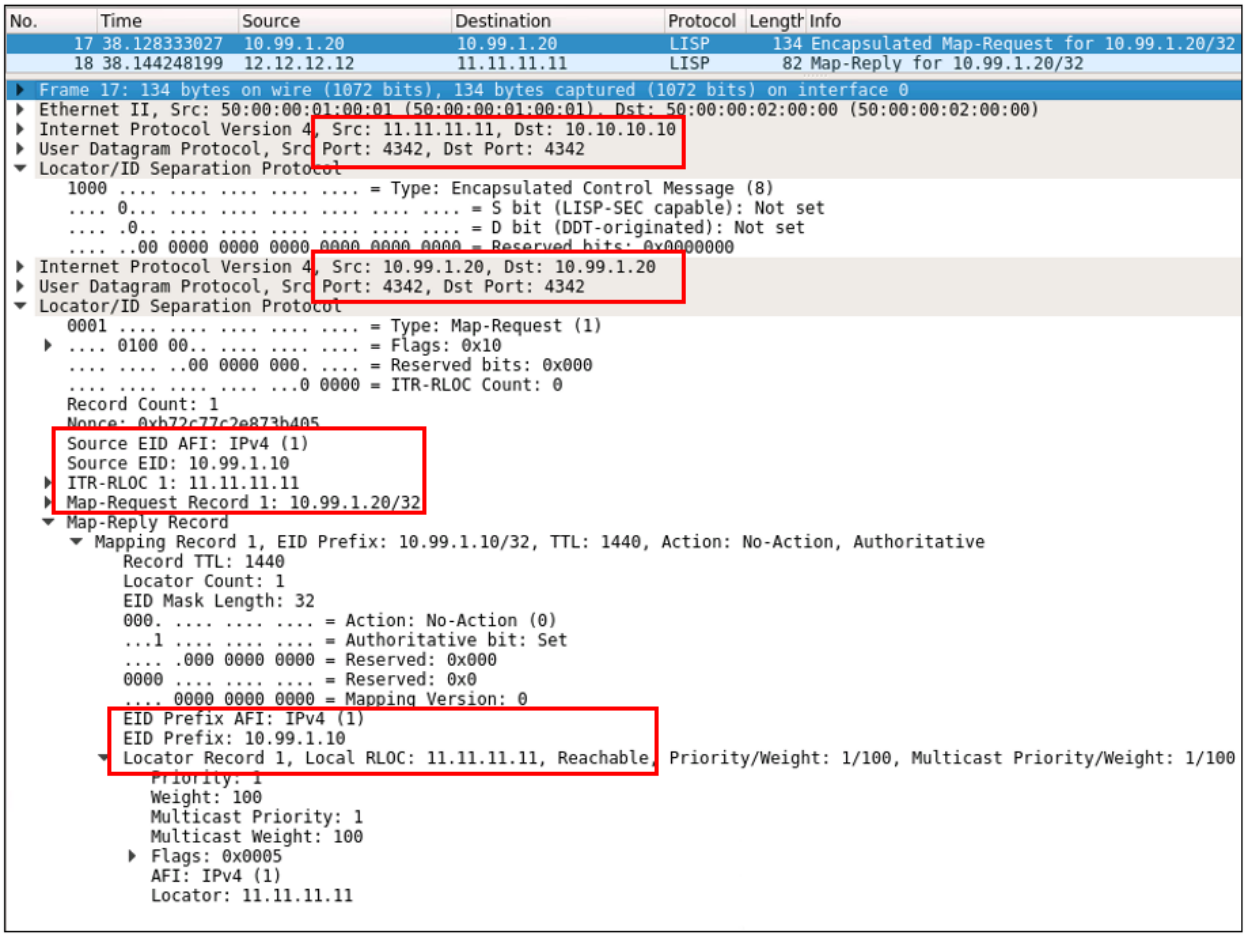

There are several things to take note of in this Map Request packet from SITE1-EDGE1.

1. The outside IP header is RLOC to RLOC between SITE-1EDGE1 and SITE1-BORDER-CP.
2. The inside IP header is the requested EID, listed as both source and destination IP, and because it is within a LISP header marked as an encapsulated control plane packet, the LISP process understands what to do with this information.
3. The Map Request packet includes what EID record is being requested, the source EID sending traffic to the destination EID referenced in the Map Request as well as the RLOC of that EID.
4. The source EID prefix information is included with this message as a Map Reply record to facilitate return traffic from the remote Edge node.

Further reading on the specifics of how Map Request message formats must be used is available within the LISP RFC: https://tools.ietf.org/html/rfc6830#section-6.1.3