Risoluzione dei problemi ARP in SD-Access Fabric

Introduzione

In questo documento viene descritto come risolvere i problemi relativi al protocollo ARP (Address Resolution Protocol) nell'infrastruttura SD-Access.

Topologia

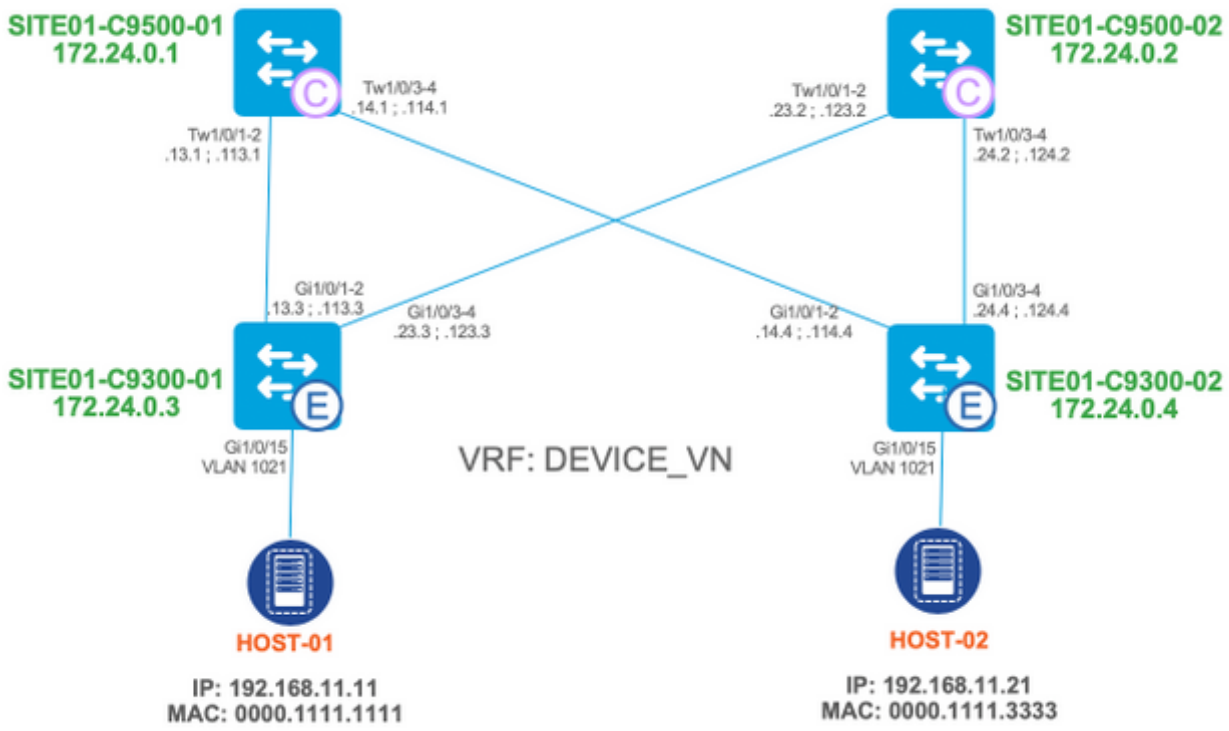

dove:

- SITE01-C9500-01 e SITE01-C9500-02 sono nodi di controllo.
- SITE01-C9300-01 e SITE01-C9300-02 sono nodi perimetrali.

L'attenzione è posta sulla comunicazione host-host (principalmente ARP-Request), anche se le stesse tecniche di risoluzione dei problemi possono essere usate per risolvere la risoluzione MAC del gateway predefinito e per ARP-Reply.

In questo documento vengono presentati due casi di utilizzo principali:

1. Risolvere l'indirizzo MAC del gateway predefinito in base agli endpoint (HOST-01 e HOST-02).
2. Risolvere l'indirizzo MAC dell'host remoto connesso alla stessa VLAN dell'host di origine: Richiesta ARP inviata da HOST-01 a HOST-02.Risposta ARP inviata da HOST-02 a HOST-01.

Stato iniziale

Si presume che:

- Entrambi gli host (SITE01-HOST-01 e SITE02-HOST-02) sono connessi alla VLAN1021 e possono raggiungere il gateway predefinito anycast locale corrispondente (192.168.11.254).

```
site01-host-01#ping 192.168.11.254
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.11.254, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/118/204 ms

site01-host-02#ping 192.168.11.254
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.11.254, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

> **Suggerimento**: Si consiglia di eseguire il ping tra il gateway predefinito e l'host finale verso lo switch sul lato, non il contrario, per evitare che conclusioni false relative ai pacchetti ICMP vengano scartati da un firewall sul dispositivo host finale.

- Edge Node ha aggiunto entrambi gli host ai database locali di rilevamento dispositivi e LISP:

HOST-01:

```
site01-c9300-01#show device-tracking database interface gi1/0/15
portDB has 1 entries for interface Gi1/0/15, 1 dynamic
<SNIP>
    Network Layer Address              Link Layer Address Interface        vlan prlvl  age
state     Time left
DH4 192.168.11.11                          0000.1111.1111  Gi1/0/15       1021  0025   34s
REACHABLE  210 s try 0(42651 s)


site01-c9300-01#show lisp eid-table vlan 1021 ethernet database 0000.1111.1111
LISP ETR MAC Mapping Database for EID-table Vlan 1021 (IID 8189), LSBs: 0x1
Entries total 2, no-route 0, inactive 0

0000.1111.1111/48, dynamic-eid Auto-L2-group-8189, inherited from default locator-set
rloc_aba7a76a-fadd-4f6e-a44e-ef4258a1c129
  Locator     Pri/Wgt  Source     State
  172.24.0.3   10/10   cfg-intf   site-self, reachable


site01-c9300-01#show lisp eid-table vlan 1021 ethernet database address-resolution
LISP ETR Address Resolution for EID-table Vlan 1021 (IID 8189)
(*) -> entry being deleted

Hardware Address       Host Address         L3 InstID
0000.1111.1111         192.168.11.11/32          4100


site01-c9300-01#show lisp eid-table vrf DEVICE_VN ipv4 database 192.168.11.11/32
LISP ETR IPv4 Mapping Database for EID-table vrf DEVICE_VN (IID 4100), LSBs: 0x1
Entries total 2, no-route 0, inactive 0
```

```
192.168.11.11/32, dynamic-eid 192_168_11_0-DEVICE_VN-IPV4, inherited from default locator-set
rloc_aba7a76a-fadd-4f6e-a44e-ef4258a1c129
  Locator     Pri/Wgt  Source     State
  172.24.0.3   10/10   cfg-intf   site-self, reachable
```

HOST-02:

```
site01-c9300-02#show device-tracking database interface gi1/0/15
<SNIP>
    Network Layer Address              Link Layer Address Interface        vlan prlvl  age
state     Time left
DH4 192.168.11.21                          0000.1111.3333  Gi1/0/15       1021  0025  109s
REACHABLE  142 s try 0(22885 s)


site01-c9300-02#show lisp eid-table vlan 1021 ethernet database 0000.1111.3333
LISP ETR MAC Mapping Database for EID-table Vlan 1021 (IID 8189), LSBs: 0x1
Entries total 2, no-route 0, inactive 0

0000.1111.3333/48, dynamic-eid Auto-L2-group-8189, inherited from default locator-set
rloc_1ee7629b-33d1-425f-82f6-60104ffbd8da
  Locator     Pri/Wgt  Source     State
  172.24.0.4   10/10   cfg-intf   site-self, reachable

site01-c9300-02#show lisp eid-table vlan 1021 ethernet database address-resolution
LISP ETR Address Resolution for EID-table Vlan 1021 (IID 8189)
(*) -> entry being deleted

Hardware Address       Host Address         L3 InstID
0000.1111.3333         192.168.11.21/32          4100


site01-c9300-02#show lisp eid-table vrf DEVICE_VN ipv4 database 192.168.11.21/32
LISP ETR IPv4 Mapping Database for EID-table vrf DEVICE_VN (IID 4100), LSBs: 0x1
Entries total 2, no-route 0, inactive 0

192.168.11.21/32, dynamic-eid 192_168_11_0-DEVICE_VN-IPV4, inherited from default locator-set
rloc_1ee7629b-33d1-425f-82f6-60104ffbd8da
  Locator     Pri/Wgt  Source     State
  172.24.0.4   10/10   cfg-intf   site-self, reachable
```

- Entrambi gli host sono stati registrati correttamente sul control-plane fabric (nodi di controllo - SITE01-C9500-01 e SITE01-C9500-02):

```
site01-c9500-01#show lisp instance-id 8189 ethernet server
LISP Site Registration Information
* = Some locators are down or unreachable
# = Some registrations are sourced by reliable transport

Site Name      Last      Up     Who Last             Inst     EID Prefix
               Register         Registered           ID
site_uci       never     no     --                   8189     any-mac
               00:28:04  yes#   172.24.0.3:16882     8189     0000.1111.1111/48
               3d23h     yes#   172.24.0.4:19075     8189     0000.1111.3333/48


site01-c9500-01#show lisp instance-id 8189 ethernet server address-resolution

Address-resolution data for router lisp 0 instance-id 8189

L3 InstID    Host Address                              Hardware Address
     4100    192.168.11.11/32                          0000.1111.1111
```

```
     4100    192.168.11.21/32                          0000.1111.3333
```

```
site01-c9500-01#show lisp instance-id 4100 ipv4 server
LISP Site Registration Information
* = Some locators are down or unreachable
# = Some registrations are sourced by reliable transport

Site Name      Last      Up     Who Last             Inst     EID Prefix
               Register         Registered           ID
site_uci       never     no     --                   4100     192.168.11.0/24
               00:27:47  yes#   172.24.0.3:16882     4100     192.168.11.11/32
               3d23h     yes#   172.24.0.4:19075     4100     192.168.11.21/32
               never     no     --                   4100     192.168.12.0/24
               never     no     --                   4100     192.168.16.0/24
```

Richiesta ARP - Risoluzione dei problemi relativi al flusso

Ingress Edge Node (SITE01-C9300-01)

In primo luogo, è sempre utile verificare che il pacchetto ARP sia effettivamente ricevuto dal nodo Ingress Edge:

```
site01-c9300-01#monitor capture cap interface gi1/0/15 in match any buffer size 1
site01-c9300-01#monitor capture cap start
Started capture point : cap
!
! trigger ping / communication between both end-points
!
site01-c9300-01#monitor capture cap stop
Capture statistics collected at software:
Capture duration - 26 seconds
Packets received - 5
Packets dropped - 0
Packets oversized - 0

Bytes dropped in asic - 0

Capture buffer will exists till exported or cleared

Stopped capture point : cap
site01-c9300-01#show monitor capture cap buffer display-filter arp
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

  3 10.098559 00:00:11:11:11:11 -> ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.11.21? Tell
192.168.11.11

site01-c9300-01#show monitor capture cap buffer display-filter arp dump
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

0000  ff ff ff ff ff ff 00 00 11 11 11 11 08 06 00 01   ................
0010  08 00 06 04 00 01 00 00 11 11 11 11 c0 a8 0b 0b   ................
0020  00 00 00 00 00 00 c0 a8 0b 15 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00 00 00 00 00               ............
```

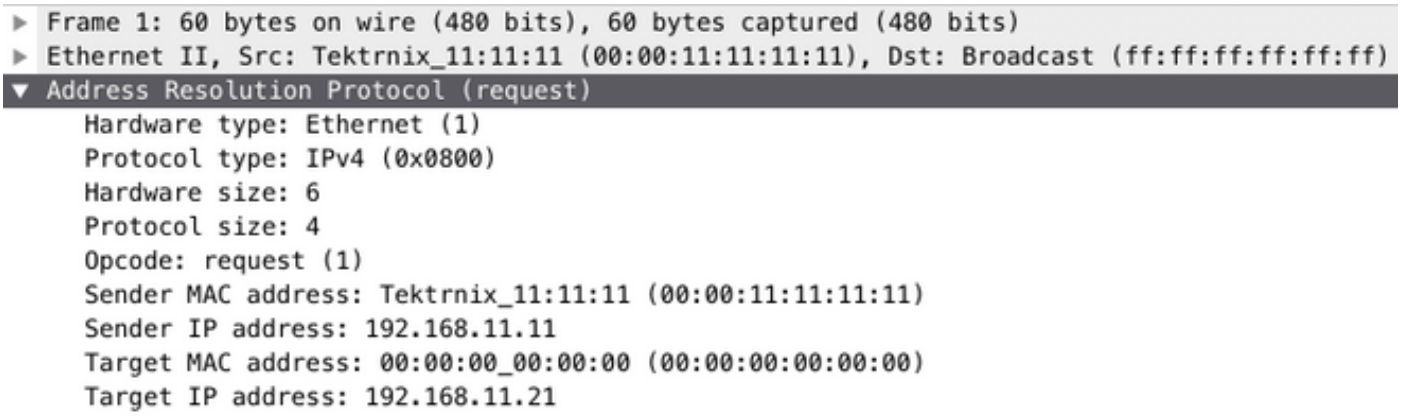

immagine 1: ricevuta richiesta ARP (nodo perimetrale in entrata)

Il pacchetto ARP-Request ricevuto viene inviato alla CPU in modo che il protocollo LISP possa essere attivato per identificare la posizione dell'indirizzo MAC di destinazione in modo che il pacchetto ARP possa essere inviato come unicast (nell'intestazione VXLAN) verso l'host remoto.

Per verificare che la richiesta ARP sia stata inoltrata correttamente alla CPU, l'acquisizione del pacchetto può essere effettuata sul control plane (il pacchetto acquisito ha esattamente lo stesso aspetto di un'acquisizione sull'interfaccia in entrata):

```
site01-c9300-01#monitor capture cpu control-plane in match any buffer size 1
site01-c9300-01#monitor capture cpu start
!
! trigger ping / communication between end-points
!
site01-c9300-01#monitor capture cpu stop
site01-c9300-01#show monitor capture cpu buffer display-filter arp
Starting the packet display ........ Press Ctrl + Shift + 6 to exit
521  43.928372 00:00:11:11:11:11 -> ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.11.21? Tell
192.168.11.11
```

Il pacchetto ARP viene indirizzato alla CPU e viene ulteriormente elaborato da diversi processi interni, tra cui:

- DAI (Dynamic ARP Inspection).
- Framework Device-Tracking / SISF (Switched Integrated Security Features).
- LISP

```
site01-c9300-01#debug platform dai all
site01-c9300-01#debug device-tracking parser
site01-c9300-01#debug lisp control-plane all
```

Il pacchetto ARP viene quindi elaborato nel software (spiegazione dettagliata fornita in linea):

```
!
! 1. ARP packet is received by control-plane (DAI process) and is passed to SISF framework:
! 000276: Mar 26 09:44:05.046: Enqueued packet in dai software queue 000277: Mar 26
09:44:05.046:  DAI processing: SMAC = 0000.1111.1111{mac} and SRC_ADDR = 192.168.11.11{ipv4}
DMAC = ffff.ffff.ffff{mac} and DST_ADDR = 192.168.11.21{ipv4}vlan: 1021, if_input: Gi1/0/15
000278: Mar 26 09:44:05.046: SISF[PRS]: ARP-REQUEST target set to 192.168.11.21
000279: Mar 26 09:44:05.046: SISF[PRS]: Gi1/0/15 vlan 1021 Arp sender LLA set to 0000.1111.1111
000280: Mar 26 09:44:05.046: SISF[PRS]: ARP sender L3 address set to 192.168.11.11
000281: Mar 26 09:44:05.047: SISF[PRS]: Gi1/0/15 vlan 1021 Advertise from access: default action
is update entry
```

000282: Mar 26 09:44:05.047: SISF[PRS]: Gi1/0/15 vlan 1021 Arp sender LLA set to 0000.1111.1111
000283: Mar 26 09:44:05.047: SISF[PRS]: Gi1/0/15 vlan 1021 Source and LLA match
000284: Mar 26 09:44:05.047: SISF[PRS]: Gi1/0/15 vlan 1021 preference level set 5
!
! 2a. LISP process (Ethernet instance: 8189) is invoked to send LISP MAP-REQUEST message to find a MAC address that corresponds with target IP address: 192.168.11.21/32
!
000285: Mar 26 09:44:05.047: [XTR] LISP-0: Remote EID IID 8189 prefix 192.168.11.21/32, Change state to incomplete (sources: <watch>, state: unknown, rlocs: 0). 000286: Mar 26 09:44:05.047: [XTR] LISP-0: Remote EID IID 8189 prefix 192.168.11.21/32, [incomplete] Scheduling map requests delay 00:00:00 min_elapsed 00:00:01 (sources: <watch>, state: incomplete, rlocs: 0). 000287: Mar 26 09:44:05.047: [XTR] LISP-0: Remote EID IID 8189 prefix 192.168.11.21/32, Starting idle timer (delay 00:02:30) (sources: <watch>, state: incomplete, rlocs: 0). 000288: Mar 26 09:44:05.176: LISP-0: IID 8189 Request processing of remote EID prefix map requests to IPv4. 000289: Mar 26 09:44:05.176: [XTR] LISP: Send map request type AR **000290: Mar 26 09:44:05.176: [XTR] LISP: Send map request for EID prefix IID 8189 192.168.11.21/32**
000291: Mar 26 09:44:05.176: [XTR] LISP-0: Remote EID IID 8189 prefix 192.168.11.21/32, Send map request (1) (sources:

000292: Mar 26 09:44:05.176: LISP-0: EID-AF IPv4, Sending map-request from 192.168.11.21 to 192.168.11.21 for EID 192.168.11.21/32, ITR-RLOCs 1, **nonce 0xDD902BBE-0x062F365F** (encap src 172.24.0.3, dst 172.24.0.2).!
!
! 2b. LISP process (Ethernet instance: 8189) receives LISP MAP-REPLY with the information about target MAC address: 0000.1111.3333.
!
000293: Mar 26 09:44:05.177: [MS] LISP: Processing received Map-Reply(2) message on GigabitEthernet1/0/4 from 172.24.0.2:4342 to 172.24.0.3:4342
000294: Mar 26 09:44:05.177: [MS] LISP: Received map reply nonce **0xDD902BBE-0x062F365F**, records 1
000295: Mar 26 09:44:05.177: [MS] LISP: Parsing mapping record for EID prefix IID 8189 192.168.11.21/32
000296: Mar 26 09:44:05.177: [MS] LISP-0: This is a Address Resolution message.
000297: Mar 26 09:44:05.177: [XTR] LISP: Processing Map-Reply mapping record for IID 8189 SVC_VLAN_IAF_ARP 192.168.11.21/32 LCAF 53, ttl 1440, action none, not authoritative, 1 locator 0000.1111.3333 pri/wei=1/100 lpr
000298: Mar 26 09:44:05.177: [XTR] LISP-0: Map Request IID 8189 prefix 192.168.11.21/32 AR[LL], Received reply with rtt 1ms.
000299: Mar 26 09:44:05.177: [XTR] LISP: Processing mapping information for EID prefix IID 8189 192.168.11.21/32
000300: Mar 26 09:44:05.177: [XTR] LISP-0: Remote EID IID 8189 prefix 192.168.11.21/32, Change state to reused (sources: <map-rep>, state: incomplete, rlocs: 0).
000301: Mar 26 09:44:05.177: [XTR] LISP-0: IAF IID 8189 SVC_VLAN_IAF_ARP, Persistent db: ignore writing request, ITR disabled.
000302: Mar 26 09:44:05.177: [XTR] LISP-0: Remote EID IID 8189 prefix 192.168.11.21/32, Change state to complete (sources:

000304: Mar 26 09:44:05.177: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, [incomplete] Scheduling map requests delay 00:00:00 min_elapsed 00:00:01 (sources: <watch>, state: incomplete, rlocs: 0).
000305: Mar 26 09:44:05.177: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, Starting idle timer (delay 00:02:30) (sources: <watch>, state: incomplete, rlocs: 0).
!
! 3a. LISP process (Ethernet instance: 8189) is invoked again to send LISP MAP-REQUEST to determine RLOC for discovered MAC: 0000.1111.3333
!
000306: Mar 26 09:44:05.305: LISP-0: IID 8189 Request processing of remote EID prefix map requests to IPv4.

000307: Mar 26 09:44:05.306: [XTR] LISP: Send map request type remote EID prefix **000308: Mar 26 09:44:05.306: [XTR] LISP: Send map request for EID prefix IID 8189 0000.1111.3333/48**
000309: Mar 26 09:44:05.306: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, Send map request (1) (sources:

000310: Mar 26 09:44:05.306: LISP-0: EID-AF MAC, Sending map-request from 0.0.0.0 to 0.0.0.0 for EID 0000.1111.3333/48, ITR-RLOCs 1, nonce **0xB518EE02-0x9E2DF283** (encap src 172.24.0.3, dst 172.24.0.2).
000311: Mar 26 09:44:05.307: [XTR] LISP: Processing received Map-Reply(2) message on GigabitEthernet1/0/4 from 172.24.0.2:4342 to 172.24.0.3:4342
!
! 3b. LISP process (Ethernet instance: 8189) receives LISP MAP-REPLY with the information about RLOC 172.24.0.4 for target MAC: 0000.1111.3333.
!
000312: Mar 26 09:44:05.307: [XTR] LISP: Received map reply nonce **0xB518EE02-0x9E2DF283**, records 1
000313: Mar 26 09:44:05.307: [XTR] LISP: Processing Map-Reply mapping record for IID 8189 SVC_VLAN_IAF_MAC 0000.1111.3333/48 LCAF 2, ttl 1440, action none, not authoritative, 1 locator 172.24.0.4 pri/wei=10/10 lpR
000314: Mar 26 09:44:05.307: [XTR] LISP-0: Map Request IID 8189 prefix 0000.1111.3333/48 remote EID prefix[LL], Received reply with rtt 1ms.
000315: Mar 26 09:44:05.307: [XTR] LISP: Processing mapping information for EID prefix IID 8189 0000.1111.3333/48
000316: Mar 26 09:44:05.307: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, Change state to reused (sources: <map-rep>, state: incomplete, rlocs: 0).
000317: Mar 26 09:44:05.307: [XTR] LISP-0: IAF IID 8189 SVC_VLAN_IAF_MAC, Persistent db: ignore writing request, disabled.
000318: Mar 26 09:44:05.307: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, Change state to complete (sources: <map-rep>, state: reused, rlocs: 0).
000319: Mar 26 09:44:05.307: [XTR] LISP: RIB Watch Group default 172.24.0.4/32 , created.
000320: Mar 26 09:44:05.307: [XTR] LISP: RIB Watch Group default 172.24.0.4/32 , scheduling RIB update.
000321: Mar 26 09:44:05.308: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, RLOCs pending rwatch update, defer fwd update (sources: <map-rep>, state: complete, rlocs: 0).
000322: Mar 26 09:44:05.308: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, 1 RLOCs pending rwatch update, defer fwd update (sources: <map-rep>, state: complete, rlocs: 0).
000323: Mar 26 09:44:05.308: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, Recalculated RLOC status bits from 0x0 to 0x1 (sources: <map-rep>, state: complete, rlocs: 1).
000324: Mar 26 09:44:05.308: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, 1 RLOCs pending rwatch update, defer fwd update (sources: <map-rep>, state: complete, rlocs: 1).
000325: Mar 26 09:44:05.308: [XTR] LISP: RIB Watch Group default 172.24.0.4/32 , installing in RIB.
000326: Mar 26 09:44:05.308: [XTR] LISP-0: Remote shrRLOC 172.24.0.4, Reachability notification, up* allow* remote.
000327: Mar 26 09:44:05.308: [XTR] LISP-0: Remote EID IID 8189 prefix 0000.1111.3333/48, No more RLOCs pending rwatch update, schedule deferred fwd update (sources: <map-rep>, state: complete, rlocs: 1).
000328: Mar 26 09:44:05.308: [XTR] LISP: MAC, SISF L2 event: ignoring event CREATED for remote host.
000329: Mar 26 09:44:05.309: [XTR] LISP: IPv4, SISF L3 event: ignoring event UPDATED for remote host.
000330: Mar 26 09:44:05.309: [XTR] LISP: IPv4, SISF L3 event: ignoring event STATE_CHANGE for remote host.
000331: Mar 26 09:44:05.309: [XTR] LISP: MAC, SISF L2 event: ignoring event VERIFIED for remote host.
000332: Mar 26 09:44:05.309: [XTR] LISP: MAC, SISF L2 event: ignoring event ACTIVE for remote host.
000333: Mar 26 09:44:05.309: [XTR] LISP: IPv4, SISF L3 event: ignoring event CREATED for remote host.

Dopo la convergenza del control-plane, le tabelle MAC / LISP devono contenere informazioni sulla posizione dell'HOST remoto.

```
site01-c9300-01#show lisp instance-id 8189 ethernet map-cache 0000.1111.3333
LISP MAC Mapping Cache for EID-table Vlan 1021 (IID 8189), 1 entries

0000.1111.3333/48, uptime: 00:31:06, expires: 23:28:53, via map-reply, complete
  Sources: map-reply
  State: complete, last modified: 00:31:06, map-source: 172.24.0.4
  Idle, Packets out: 0(0 bytes)
  Encapsulating dynamic-EID traffic
  Locator     Uptime    State      Pri/Wgt     Encap-IID
  172.24.0.4  00:31:06  up          10/10        -
    Last up-down state change:         00:31:06, state change count: 1
    Last route reachability change:    00:31:06, state change count: 1
    Last priority / weight change:     never/never
    RLOC-probing loc-status algorithm:
      Last RLOC-probe sent:            00:31:06 (rtt 1ms)

site01-c9300-01#show mac address-table dynamic | in 0000.1111.3333
1021    0000.1111.3333    CP_LEARN    Tu0
```

In questa fase, il piano dati inoltra il pacchetto ARP alla destinazione finale (notare che il primo pacchetto ARP originale non viene scartato, ma memorizzato nel buffer per il momento in cui il tempo di controllo converge, per evitare un potenziale impatto negativo su alcuni endpoint, come i telefoni IP):

```
site01-c9300-01#monitor capture cpu control-plane in match any buffer size 1
site01-c9300-01#monitor capture cpu start
!
! trigger ping / communication between end-points
!
site01-c9300-01#monitor capture cpu stop
site01-c9300-01#show monitor capture cpu buffer display-filter arp
Starting the packet display ........ Press Ctrl + Shift + 6 to exit
793  53.106637 00:00:11:11:11:11 -> 00:00:11:11:33:33 ARP 42 Who has 192.168.11.21? Tell
192.168.11.11

site01-c9300-01#show monitor capture cpu buffer display-filter frame.number==793 dump
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

0000  00 00 11 11 33 33 00 00 11 11 11 11 08 06 00 01   ....33..........
0010  08 00 06 04 00 01 00 00 11 11 11 11 c0 a8 0b 0b   ................
0020  00 00 00 00 00 00 c0 a8 0b 15                     ..........
```

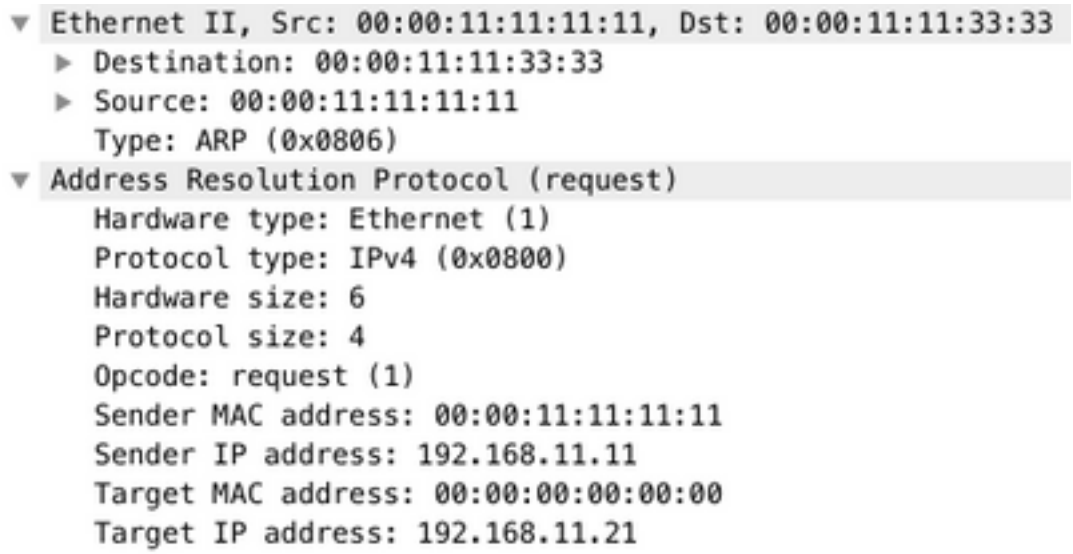

Come mostrato nell'immagine, l'indirizzo MAC di destinazione Ethernet viene modificato da broadcast a unicast con l'indirizzo MAC di destinazione.

Allo stesso tempo, l'indirizzo MAC di destinazione nell'intestazione ARP stessa non viene modificato.

Il pacchetto ARP viene inviato come unicast all'hardware dove viene incapsulato nell'intestazione VXLAN unicast basata sulla ricerca L2 verso la RLOC 172.24.0.4 tramite l'adiacenza correlata:

```
site01-c9300-01#show platform software fed switch 1 matm macTable vlan 1021
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
riHandle           diHandle             *a_time  *e_time  ports
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------
1021   0000.0c9f.f45c      0x8002      0  98140     64  0x7fab2cc44f78      0x7fab2cc43c48
0x0                0x0                          0        0  Vlan1021
1021   7c21.0d1c.f8f5      0x8002      0  98140     64  0x7fab2cc46518      0x7fab2cc43c48
0x0                0x0                          0        0  Vlan1021
1021   0000.1111.1111         0x1     26      0      0  0x7fab2cf1cb88      0x7fab2cbaec48
0x0                0x7fab2ca137b8            300       15  GigabitEthernet1/0/15
1021   0000.1111.3333   0x1000001      0      0     64  0x7fab2cd7d0f8      0x7fab2cd66908
0x7fab2cb76b68      0x0                          0       15  RLOC 172.24.0.4 adj_id 116

site01-c9300-01#show platform software fed switch 1 matm adjacencies
VLAN   ADJ_ID      ADJ_KEY          Encap  Link   siHandle         riHandle         diHandle
Obj_type
-------------------------------------------------------------------------------------------
--------------
1021   116         0x100000074      VXLAN  V4     0x7fab2cd66908   0x7fab2cb76b68   0x0
CP

site01-c9300-01#show platform software fed switch active matm adjacencies adjkey 0x100000074
ADJ_ID      IF_NUMBER  VNI        Len    Vlan   Encap  Link   Source IP         Dest IP
siHandle             riHandle             diHandle             Obj_type
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------
116         0x420011   8189       40     1021   VXLAN  V4     172.24.0.3        172.24.0.4
0x7fab2cd66908       0x7fab2cb76b68       0x0                  CP
```

Stato interfaccia:

```
site01-c9300-01#show platform software fed switch 1 ifm if-id 0x420011
Interface IF_ID         : 0x0000000000420011
Interface Name          : L2LISP0.8189
Interface Block Pointer : 0x7fab2cbdfa28
Interface Block State   : READY
Interface State         : Enabled
Interface Status        : ADD
Interface Ref-Cnt       : 2
Interface Type          : L2_LISP
Is top interface    : FALSE
Asic_num            : 0
Switch_num          : 0
AAL port Handle     : ae000063
Parent interface id : 43
Multicast Tunnel IP : 0.0.0.0
Mcast Tunnel Handle : NULL
Vlan Id             : 1021
Instance Id         : 8189
Dest Port           : 4789
SGT                 : Enable
Underlay VRF (V4)   : 0
Underlay VRF (V6)   : 0
Flood Access-tunnel : Disable
Flood unknown ucast : Disable
Broadcast           : Enable
Multicast Flood     : Disable

Port Information
Handle ............ [0xae000063]
        Type .............. [L2-LISP-sub]
Identifier ........ [0x420011]
Unit .............. [4325393]
 L2 LISP Sub-interface Subblock
Parent iif id     : 0x43
Switch Num        : 1
Asic Num          : 0
Encap PORT LE handle    : 0x7fab2ca9d1c8
Decap PORT LE handle    : 0x7fab2ca9d018
L3IF LE handle    : 0x7fab2ca9d698
SI handle decap   : 0x7fab2cc00948
DI handle         : 0x7fab2c311688
RI handle         : 0x7fab2cbfdff8
RCP Service ID    : 0x0
TRANS HTM handle     : 0x7fab2cbfe5a8
TRANS CATCH ALL handle : 0x7fab2c3125a8
Port L2 Subblock
Enabled ............. [No]
Allow dot1q ......... [No]
Allow native ........ [No]
Default VLAN ........ [0]
Allow priority tag ... [No]
Allow unknown unicast  [No]
Allow unknown multicast[No]
Allow unknown broadcast[No]
Allow unknown multicast[Enabled]
Allow unknown unicast  [Enabled]
Protected ............ [No]
                IPv4 ARP snoop ....... [No]
IPv6 ARP snoop ....... [No]
Jumbo MTU ............ [0]
Learning Mode ........ [0]
Vepa ................. [Disabled]
Port QoS Subblock
Trust Type ................... [0x7]
```

```
Default Value ................. [0]
Ingress Table Map ............. [0x0]
Egress Table Map .............. [0x0]
Queue Map ..................... [0x0]
Port Netflow Subblock
Port Policy Subblock
List of Ingress Policies attached to an interface
List of Egress Policies attached to an interface
Port CTS Subblock
Disable SGACL ................. [0x0]
Trust ......................... [0x0]
Propagate ..................... [0x0]
Port SGT ...................... [0xffff]
Ref Count : 2 (feature Ref Counts + 1)
IFM Feature Ref Counts
FID : 96 (AAL_FEATURE_L2_MULTICAST_IGMP), Ref Count : 1
No Sub Blocks Present
```

Tutti i gestori interni possono essere controllati (dopo aver abilitato il **servizio interno** in modalità di configurazione) nei dettagli:

Indice stazione:

```
site01-c9300-01#show platform hardware fed switch 1 fwd-asic abstraction print-resource-handle
0x7fab2cd66908 1
Handle:0x7fab2cd66908 Res-Type:ASIC_RSC_SI Res-Switch-Num:255 Asic-Num:255 Feature-
ID:AL_FID_L2_WIRELESS Lkp-ftr-id:LKP_FEAT_INVALID ref_count:1
priv_ri/priv_si Handle: 0x7fab2cb76b68Hardware Indices/Handles: index0:0xd4
mtu_index/l3u_ri_index0:0x0  index1:0xd4  mtu_index/l3u_ri_index1:0x0
Features sharing this resource:58 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 fd 03 00 00 00 00 00 00 00 00 00 00 07 00 74 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Detailed Resource Information (ASIC# 0)
---------------------------------------


Station Index (SI) [0xd4]
RI = 0x3c
DI = 0x5012
stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: LD

Detailed Resource Information (ASIC# 1)
---------------------------------------


Station Index (SI) [0xd4]
RI = 0x3c
DI = 0x5013
stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: LD
```

==========

Riscrivi:

```
site01-c9300-01#show platform hardware fed switch 1 fwd-asic abstraction print-resource-handle
0x7fab2cb76b68 1
Handle:0x7fab2cb76b68 Res-Type:ASIC_RSC_RI Res-Switch-Num:255 Asic-Num:255 Feature-
ID:AL_FID_L2_WIRELESS Lkp-ftr-id:LKP_FEAT_INVALID ref_count:1
priv_ri/priv_si Handle: 0x7fab2cd66b38Hardware Indices/Handles: index0:0x3c
mtu_index/l3u_ri_index0:0x0  index1:0x3c  mtu_index/l3u_ri_index1:0x0
Features sharing this resource:58 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 fd 03 00 00 00 00 00 00 00 00 00 00 07 00 74 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Detailed Resource Information (ASIC# 0)
---------------------------------------

Detailed Resource Information (ASIC# 1)
---------------------------------------
```

==========

Indice destinazione:

```
site01-c9300-01#show platform hardware fed switch 1 fwd-asic resource asic all destination-index
range 0x5012 0x5013
ASIC#0:

index = 0x5012
pmap = 0x00000000 0x00000000
cmi = 0x0
rcp_pmap = 0x1
al_rsc_cmi
CPU Map Index (CMI) [0]
ctiLo0 = 0
ctiLo1 = 0
ctiLo2 = 0
cpuQNum0 = 0
cpuQNum1 = 0
cpuQNum2 = 0
npuIndex = 0
stripSeg = 0
copySeg = 0

index = 0x5013
pmap = 0x00000000 0x00000000
cmi = 0x0
rcp_pmap = 0x0
al_rsc_cmi
CPU Map Index (CMI) [0]
ctiLo0 = 0
ctiLo1 = 0
ctiLo2 = 0
cpuQNum0 = 0
cpuQNum1 = 0
cpuQNum2 = 0
npuIndex = 0
stripSeg = 0
copySeg = 0

ASIC#1:
```

```
index = 0x5012
pmap = 0x00000000 0x00000000
cmi = 0x0
rcp_pmap = 0x0
al_rsc_cmi
CPU Map Index (CMI) [0]
ctiLo0 = 0
ctiLo1 = 0
ctiLo2 = 0
cpuQNum0 = 0
cpuQNum1 = 0
cpuQNum2 = 0
npuIndex = 0
stripSeg = 0
copySeg = 0

index = 0x5013
pmap = 0x00000000 0x00000000
cmi = 0x0
rcp_pmap = 0x1
al_rsc_cmi
CPU Map Index (CMI) [0]
ctiLo0 = 0
ctiLo1 = 0
ctiLo2 = 0
cpuQNum0 = 0
cpuQNum1 = 0
cpuQNum2 = 0
npuIndex = 0
stripSeg = 0
copySeg = 0
```

Nodo perimetrale in uscita (SITE01-C9300-02)

Quando tutte le operazioni su Ingress Edge Node hanno esito positivo, il pacchetto ARP viene ricevuto come unicast (e VXLAN incapsulata) sul corretto Edge Node in uscita.

Èpossibile acquisire il pacchetto ARP sul nodo del perimetro di uscita (dove gli indirizzi IP utilizzati per l'acquisizione sono Loopback0 di entrambi i nodi del perimetro), come mostrato:

```
site01-c9300-02#monitor capture uplink interface range gi1/0/1 - 4 in match ipv4 host 172.24.0.3
host 172.24.0.4 buffer size 1
site01-c9300-02#monitor capture uplink start
!
! trigger ping / communication between both end-points
!
site01-c9300-02#monitor capture uplink stop
site01-c9300-02#show monitor capture uplink buffer display-filter arp Starting the packet
display ........ Press Ctrl + Shift + 6 to exit

  1   0.000000 00:00:11:11:11:11 -> 00:00:11:11:33:33 ARP 110 Who has 192.168.11.21? Tell
192.168.11.11

site01-c9300-02#show monitor capture uplink buffer display-filter arp dump
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

0000  7c 21 0d 1d 6e f6 4c e1 75 17 6d 9f 08 00 45 00   |!..n.L.u.m...E.
0010  00 60 00 0b 40 00 fd 11 25 4a ac 18 00 03 ac 18   .`..@...%J......
0020  00 04 ff 49 12 b5 00 4c 00 00 88 00 00 0d 00 1f   ...I...L........
0030  fd 00 00 00 11 11 33 33 00 00 11 11 11 11 08 06   ......33........
0040  00 01 08 00 06 04 00 01 00 00 11 11 11 11 c0 a8   ................
```

```
0050  0b 0b 00 00 00 00 00 00 c0 a8 0b 15 00 00 00 00   ................
0060  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
▶ Frame 1: 110 bytes on wire (880 bits), 110 bytes captured (880 bits)
▶ Ethernet II, Src: 4c:e1:75:17:6d:9f, Dst: 7c:21:0d:1d:6e:f6
▶ Internet Protocol Version 4, Src: 172.24.0.3, Dst: 172.24.0.4
▶ User Datagram Protocol, Src Port: 65353, Dst Port: 4789
▼ Virtual eXtensible Local Area Network
  ▶ Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    Group Policy ID: 13
    VXLAN Network Identifier (VNI): 8189
    Reserved: 0
▶ Ethernet II, Src: 00:00:11:11:11:11, Dst: 00:00:11:11:33:33
▼ Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: 00:00:11:11:11:11
    Sender IP address: 192.168.11.11
    Target MAC address: 00:00:00:00:00:00
    Target IP address: 192.168.11.21
```

Nell'immagine viene mostrata l'acquisizione di pacchetti decodificati, in cui la richiesta ARP unicast è stata ricevuta come frame VXLAN incapsulato con VNI: 8189 (che corrisponde a LISP Ethernet Instance-id) e ID Criteri di gruppo: 13 (che è il valore SGT associato all'origine HOST-01).

Successivamente, il pacchetto viene ricircolato nell'hardware (per rimuovere l'intestazione VXLAN) e puntato alla CPU - gli output sottostanti possono essere generati da SPF (Show Platform Forward sotto la supervisione di Cisco TAC):

```
Input Packet Details:
###[ Ethernet ]###
  dst       = 7c:21:0d:1d:6e:f6
  src=4c:e1:75:17:6d:9f
  type      = 0x800
###[ IP ]###
     version   = 4L
     ihl       = 5L
     tos       = 0x0
     len       = 96
     id        = 28
     flags     = DF
     frag      = 0L
     ttl       = 253
     proto     = udp
     chksum    = 0x2539
     src=172.24.0.3
     dst       = 172.24.0.4
     options   = ''
###[ UDP ]###
```

```
        sport     = 65353
        dport     = 4789
        len       = 76
        chksum    = 0x0
###[ VXLAN ]###
           flags     = I+N
           vxlanSGT  = 0xdL
           vxlanNetworkIdentifier= 0x1ffdL
           reserved  = 0x0L
###[ Ethernet ]###
              dst       = 00:00:11:11:33:33
              src=00:00:11:11:11:11
              type      = 0x806
###[ ARP ]###
                 hwtype    = 0x1
                 ptype     = 0x800
                 hwlen     = 6
                 plen      = 4
                 op        = who-has
                 hwsrc=00:00:11:11:11:11
                 psrc=192.168.11.11
                 hwdst     = 00:00:00:00:00:00
                 pdst      = 192.168.11.21
###[ Padding ]###
                    load      = '00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00'
Ingress:
   Port                      : GigabitEthernet1/0/4
   Global Port Number        : 4
   Local Port Number         : 4
   Asic Port Number          : 3
   Asic Instance             : 1
   Vlan                      : 4095
   Mapped Vlan ID            : 1
   STP Instance              : 1
   BlockForward              : 0
   BlockLearn                : 0
   L3 Interface              : 40
       IPv4 Routing          : enabled
       IPv6 Routing          : enabled
       Vrf Id                : 0
   Adjacency:
       Station Index         : 203
       Destination Index     : 20498
       Rewrite Index         : 59399
       Replication Bit Map   : 0x11   ['remoteData', 'coreData']
   Decision:
       Destination Index     : 20498  [DI_RCP_PORT1]
       Rewrite Index         : 59399
       Dest Mod Index        : 0      [IGR_FIXED_DMI_NULL_VALUE]
       CPU Map Index         : 0      [CMI_NULL]
       Forwarding Mode       : 3      [Other or Tunnel]
       Replication Bit Map   :        ['remoteData', 'coreData']
       Winner                :        LISPVXLANINSTTRANSLATION LOOKUP
       Qos Label             : 1
       SGT                   : 0
       DGTID                 : 0
EXCEPTION:
   Datapath: Expected packet not replicated on interface
******************************************************************************
```

Per verificare che venga ricevuto/rimandato indietro decapsulato dal control-plane, è possibile effettuare un'ulteriore acquisizione del pacchetto sul control-plane in direzione di entrata:

```
site01-c9300-02#monitor capture cpu control-plane both match any buffer size 1
site01-c9300-02#monitor capture cpu start
!
! trigger ping / communication between both end-points
!
site01-c9300-02#monitor capture cpu stop
site01-c9300-02#show monitor cap cpu buffer display-filter arp
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

523  22.146501 00:00:11:11:11:11 -> 00:00:11:11:33:33 ARP 60 Who has 192.168.11.21? Tell
192.168.11.11   <-- punt from Hardware to CPU
524  22.146667 00:00:11:11:11:11 -> 00:00:11:11:33:33 ARP 60 Who has 192.168.11.21? Tell
192.168.11.11   <-- inject from CPU to Hardware

site01-c9300-02#show monitor capture cpu buffer display-filter frame.number==523 dump
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

0000  00 00 11 11 33 33 00 00 11 11 11 11 08 06 00 01   ....33..........
0010  08 00 06 04 00 01 00 00 11 11 11 11 c0 a8 0b 0b   ................
0020  00 00 00 00 00 00 c0 a8 0b 15 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00 00 00 00 00               ............
```

```
▶ Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
▶ Ethernet II, Src: 00:00:11:11:11:11, Dst: 00:00:11:11:33:33
▼ Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: 00:00:11:11:11:11
    Sender IP address: 192.168.11.11
    Target MAC address: 00:00:00:00:00:00
    Target IP address: 192.168.11.21
```

Il pacchetto ARP viene gestito dal processo software e inviato all'hardware. Questo comportamento può essere confermato tramite i debug:

```
site01-c9300-02#debug platform dai all
site01-c9300-02#show logging
584813848: Mar 26 13:29:08.721:  Enqueued packet in dai software fromCAPWAP or Access Tunnel or
LISP interface queue
584813849: Mar 26 13:29:08.721:  DAI processing: SMAC = 0000.1111.1111{mac} and SRC_ADDR =
192.168.11.11{ipv4} DMAC = 0000.1111.3333{mac} and DST_ADDR = 192.168.11.21{ipv4}vlan: 1021,
if_input: Tu0
584813850: Mar 26 13:29:08.721:  Hijacking ARP packet on LISP interface vlan: 1021, if_input:
Tu0, link_type: 1
584813851: Mar 26 13:29:08.721:  Packet consumed
584813852: Mar 26 13:29:08.722:  Enqueued packet in dai software queue
```

Il pacchetto viene inviato all'hardware nella VLAN1021 per la ricerca finale:

```
site01-c9300-02#monitor capture vlan interface vlan 1021 in match any buffer size 1
site01-c9300-02#monitor capture vlan start
!
! trigger ping / communication between both end-points
!
site01-c9300-02#monitor capture vlan stop
```

```
site01-c9300-02#show monitor capture vlan buffer display-filter arp
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

  5  21.471664 00:00:11:11:11:11 -> 00:00:11:11:33:33 ARP 60 Who has 192.168.11.21? Tell
192.168.11.11

site01-c9300-02#show monitor capture vlan buffer display-filter arp dump
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

0000  00 00 11 11 33 33 00 00 11 11 11 11 08 06 00 01   ....33..........
0010  08 00 06 04 00 01 00 00 11 11 11 11 c0 a8 0b 0b   ................
0020  00 00 00 00 00 00 c0 a8 0b 15 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00 00 00 00 00               ............
```

```
▶ Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
▶ Ethernet II, Src: 00:00:11:11:11:11, Dst: 00:00:11:11:33:33
▼ Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: 00:00:11:11:11:11
    Sender IP address: 192.168.11.11
    Target MAC address: 00:00:00:00:00:00
    Target IP address: 192.168.11.21
```

La ricerca finale si basa sull'indirizzo MAC di destinazione:

```
site01-c9300-02#show mac address-table dynamic address 0000.1111.3333
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
1021    0000.1111.3333    DYNAMIC     Gi1/0/15
Total Mac Addresses for this criterion: 1

site01-c9300-02#show platform software fed switch 1 matm macTable vlan 1021
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
riHandle            diHandle              *a_time  *e_time  ports
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------
1021   0000.1111.3333         0x1   1554      0      0  0x7fe044d9ece8      0x7fe044a34518
0x0                 0x7fe044948588            300       11  GigabitEthernet1/0/15
```

Il pacchetto ARP viene inoltrato verso la porta di uscita sul nodo del perimetro di uscita.

Risposta ARP - Risoluzione dei problemi relativi al flusso

La risposta ARP può essere utilizzata per la risoluzione dei problemi esattamente come la richiesta ARP presentata in questo documento.