

# What You Make Possible



# Advanced OTV – Configure, Verify and Troubleshoot OTV in Your Network

BRKDCT - 3103

# Build The Bridge



# The OTV Bridge



# Agenda

- **OTV Introduction**
- **Configuration**
  - Multicast Transport
  - Unicast-only Transport
- **Verification**
  - Adjacency
  - Unicast Forwarding
  - Multicast Forwarding
  - ARP
- **Troubleshooting**

# Introduction


## Overlay Transport Virtualisation (OTV) in a Nutshell

- OTV is a MAC-in-IP method that extends Layer 2 connectivity across a transport network infrastructure
- OTV supports both multicast and unicast-only transport networks
- OTV uses ISIS as the control protocol
- OTV on Nexus7000 does not encrypt encapsulated payload

# Introduction

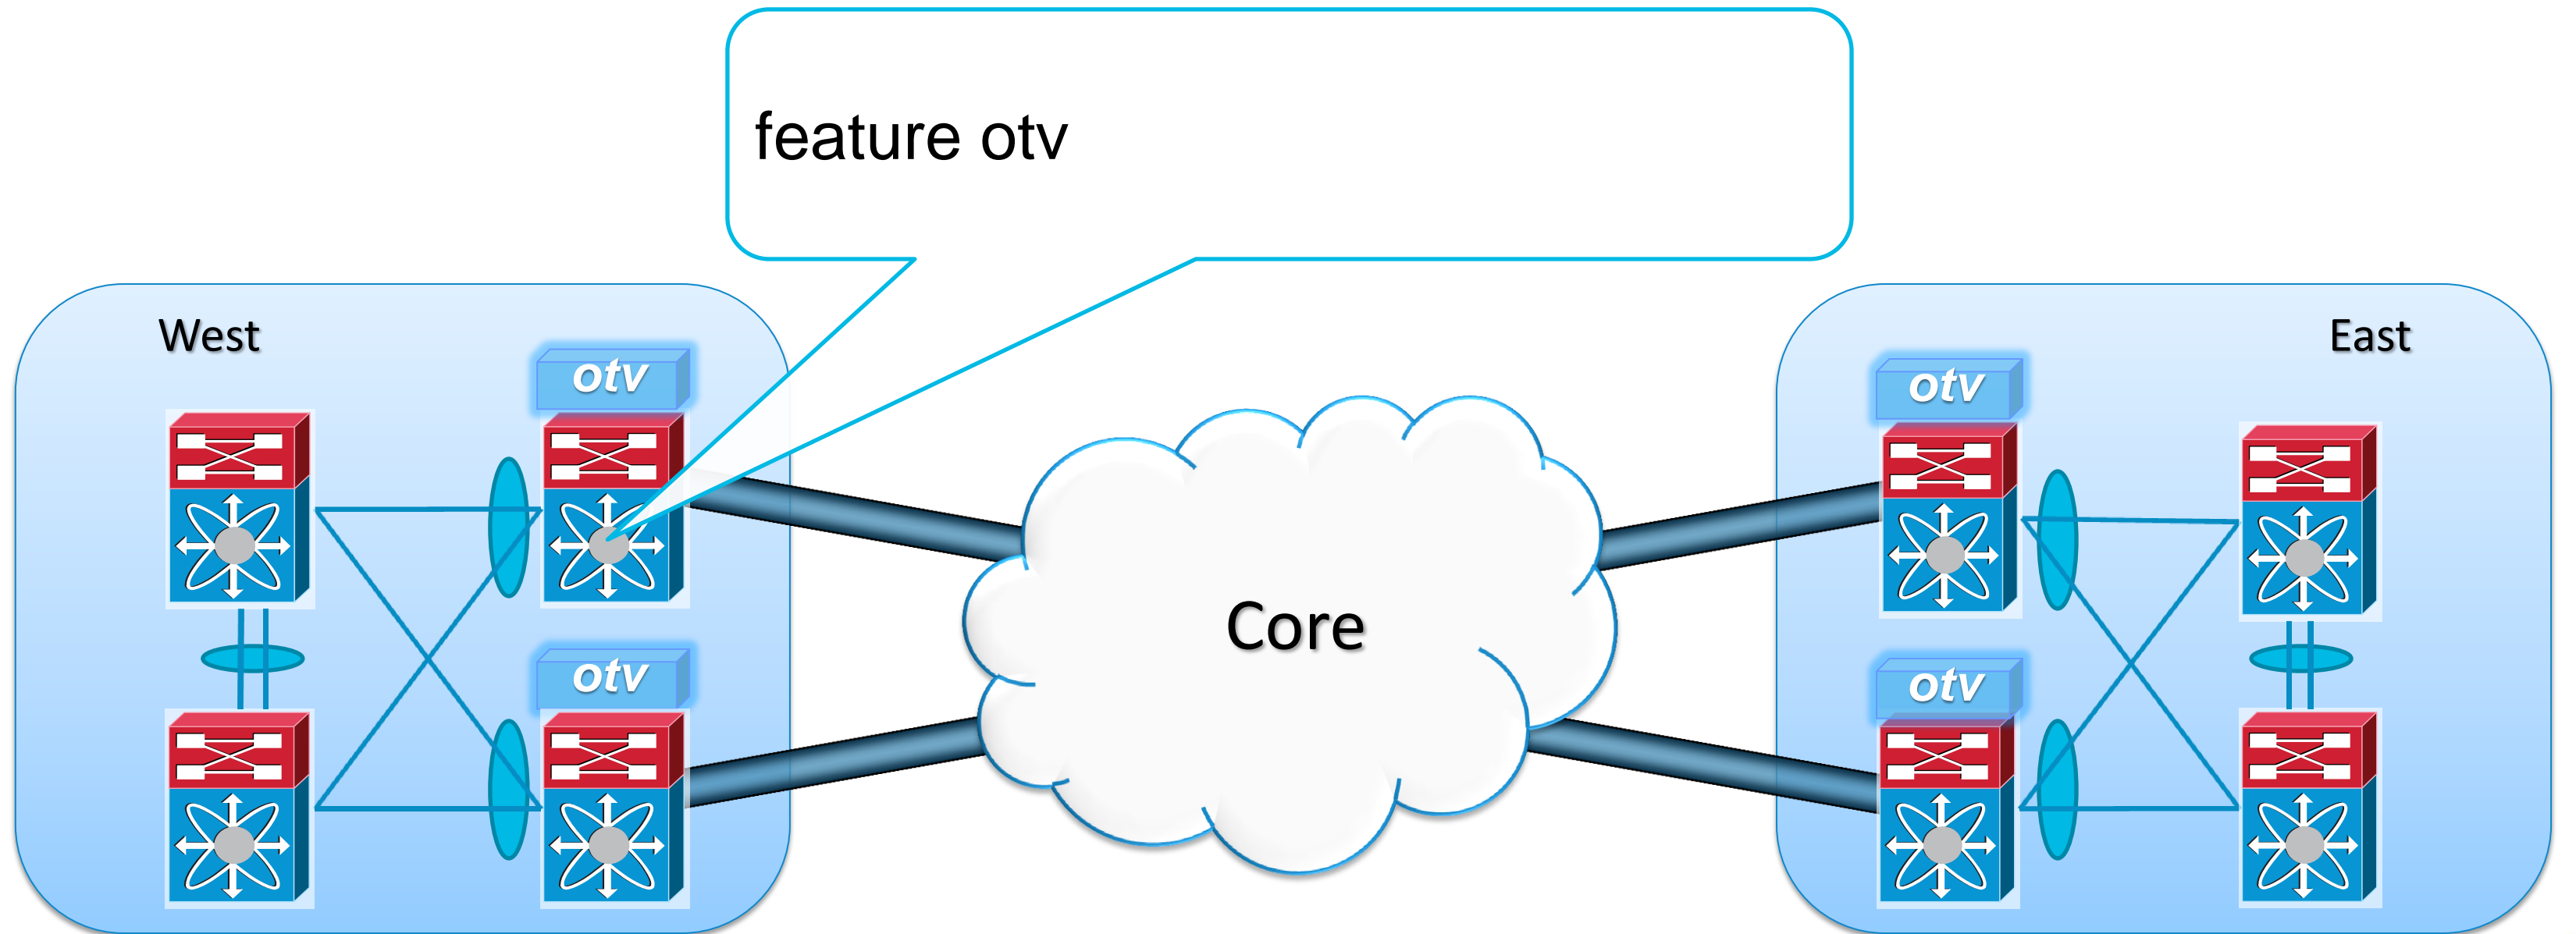
## Terminology: Edge Device

- Performs OTV functions
- Support multiple OTV devices per site
- OTV requires the Transport Services (TRS) license
- Creating non default VDC's requires Advanced Services license



# Configuration

## Enable OTV Feature





# Configuration

## Site VLAN and Site Identifier

- **Site VLAN** needs to be configured and active even if you do not have multiple OTV devices in the same site
  - site VLAN should not be extended across overlay
- **Site Identifier** can be any number between 0000.0000.0001 and ffff.ffff.ffff. Value will always be displayed in MAC format
- **Site Identifier must be unique for each site**
- **Site Identifier** is required in 5.2(1) and above for overlay to come up. This must be kept in mind when performing an ISSU from a pre-5.2(1)


Service Impacting

# Configuration

## Site VLAN and Site Identifier

otv site-vlan 210  
otv site-identifier 0001.0001.0001


otv site-vlan 210  
otv site-identifier 0002.0002.0002



# Configuration

## Join Interface

```
interface port-channel 100  
  mtu 9216  
  ip address 172.16.1.34/30  
  ip igmp version 3
```



# Agenda

- **OTV Introduction**

- **Configuration**

  - Multicast Transport

  - Unicast-only Transport

- **Verification**

  - Adjacency

  - Unicast Forwarding

  - Multicast Forwarding


  - ARP

- **Troubleshooting**

# Configuration

## Multicast Transport: Overlay

```
interface Overlay1
  otv join-interface port-channel100
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/24
  otv extend-vlan 200-209
```



# Configuration

## Multicast Transport: Full Picture

### WEST\_OTVA

```
feature otv
otv site-vlan 210
otv site-identifier 0001.0001.0001
```

```
interface Overlay1
  otv join-interface port-channel100
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/24
  otv extend-vlan 200-209
  no shutdown
```

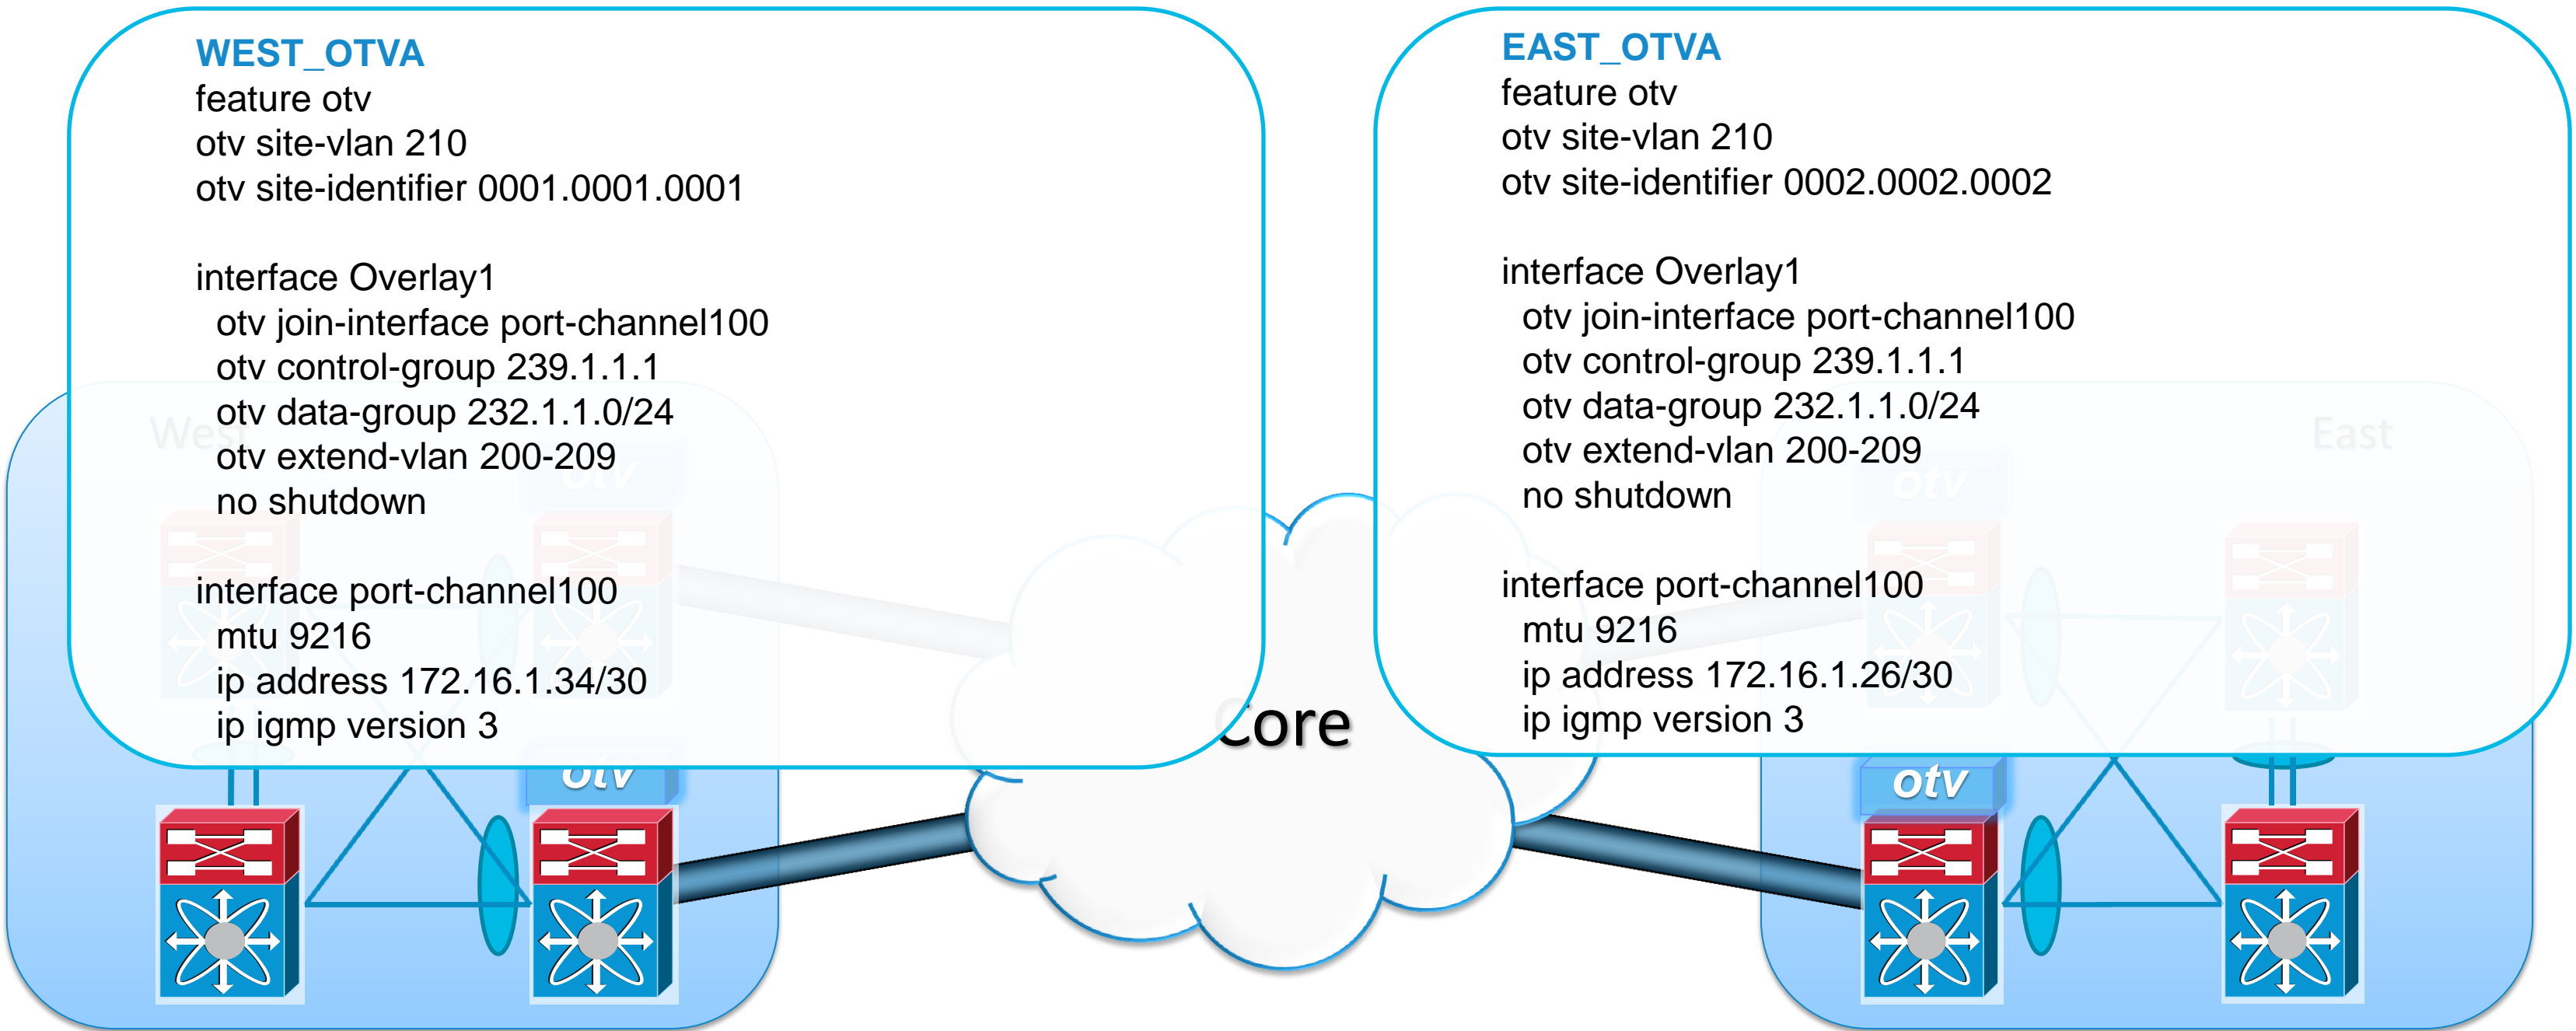
```
interface port-channel100
  mtu 9216
  ip address 172.16.1.34/30
  ip igmp version 3
```

### EAST\_OTVA

```
feature otv
otv site-vlan 210
otv site-identifier 0002.0002.0002
```

```
interface Overlay1
  otv join-interface port-channel100
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/24
  otv extend-vlan 200-209
  no shutdown
```

```
interface port-channel100
  mtu 9216
  ip address 172.16.1.26/30
  ip igmp version 3
```



# Agenda

- **OTV Introduction**

## Configuration

- Multicast Transport

- Unicast-only Transport

- **Verification**

- Adjacency

- Unicast Forwarding

- Multicast Forwarding

- ARP

- **Troubleshooting**

# Configuration

## Unicast Transport: Overlay


- OTV can run across a unicast only transport
- Unicast Transport requires the configuration of one or more adjacency servers. OTV devices register with the adjacency server which in turn provides each with an OTV Neighbour List (oNL).
- Think of the adjacency server as a special process running on a generic OTV edge device
- A primary and secondary adjacency server can be configured for redundancy



# Configuration

## Unicast Transport: Primary Adjacency Server Overlay

```
interface Overlay1
  otv join-interface port-channel100
  otv extend-vlan 200-209
  otv adjacency-server unicast-only
```

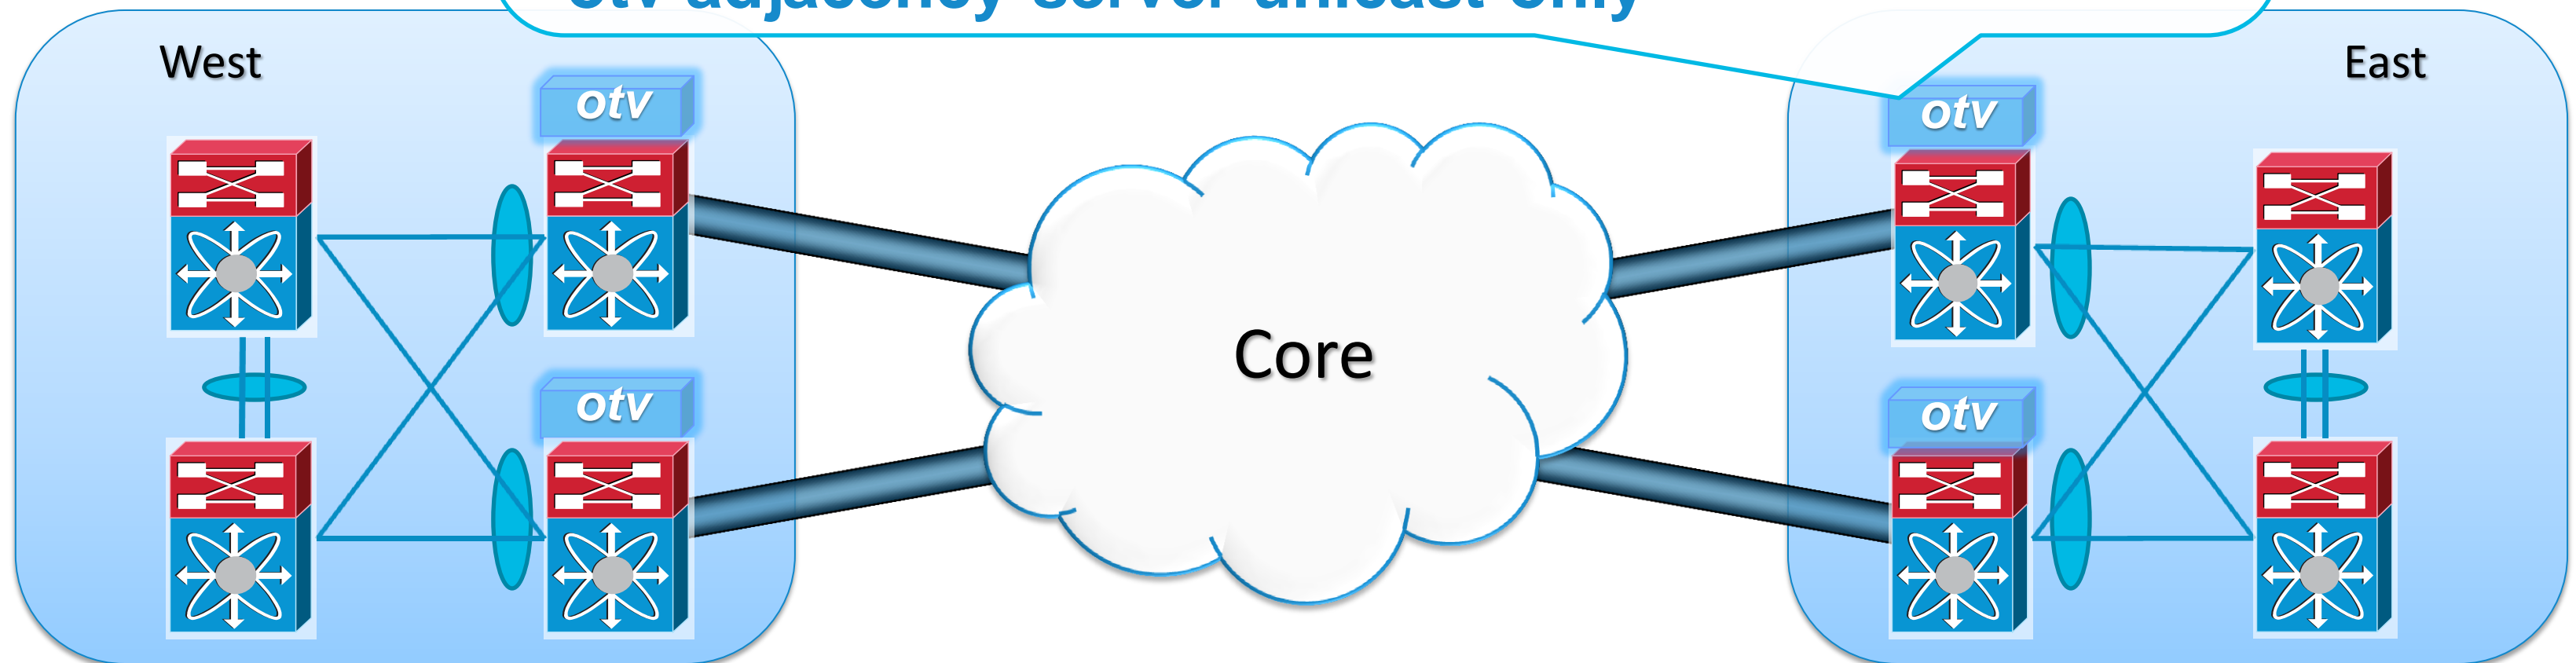


# Configuration

## Unicast Transport: Secondary Adjacency Server Overlay

```
interface Overlay1
  otv join-interface port-channel100
  otv extend-vlan 200-209
  otv use-adjacency-server 172.16.1.34 unicast-only
  otv adjacency-server unicast-only
```

Primary Server




# Configuration

## Unicast Transport: Client Overlay

```
interface Overlay1
  otv join-interface port-channel100
  otv extend-vlan 200-209
  otv use-adjacency-server 172.16.1.34 172.16.1.26 unicast-only
```

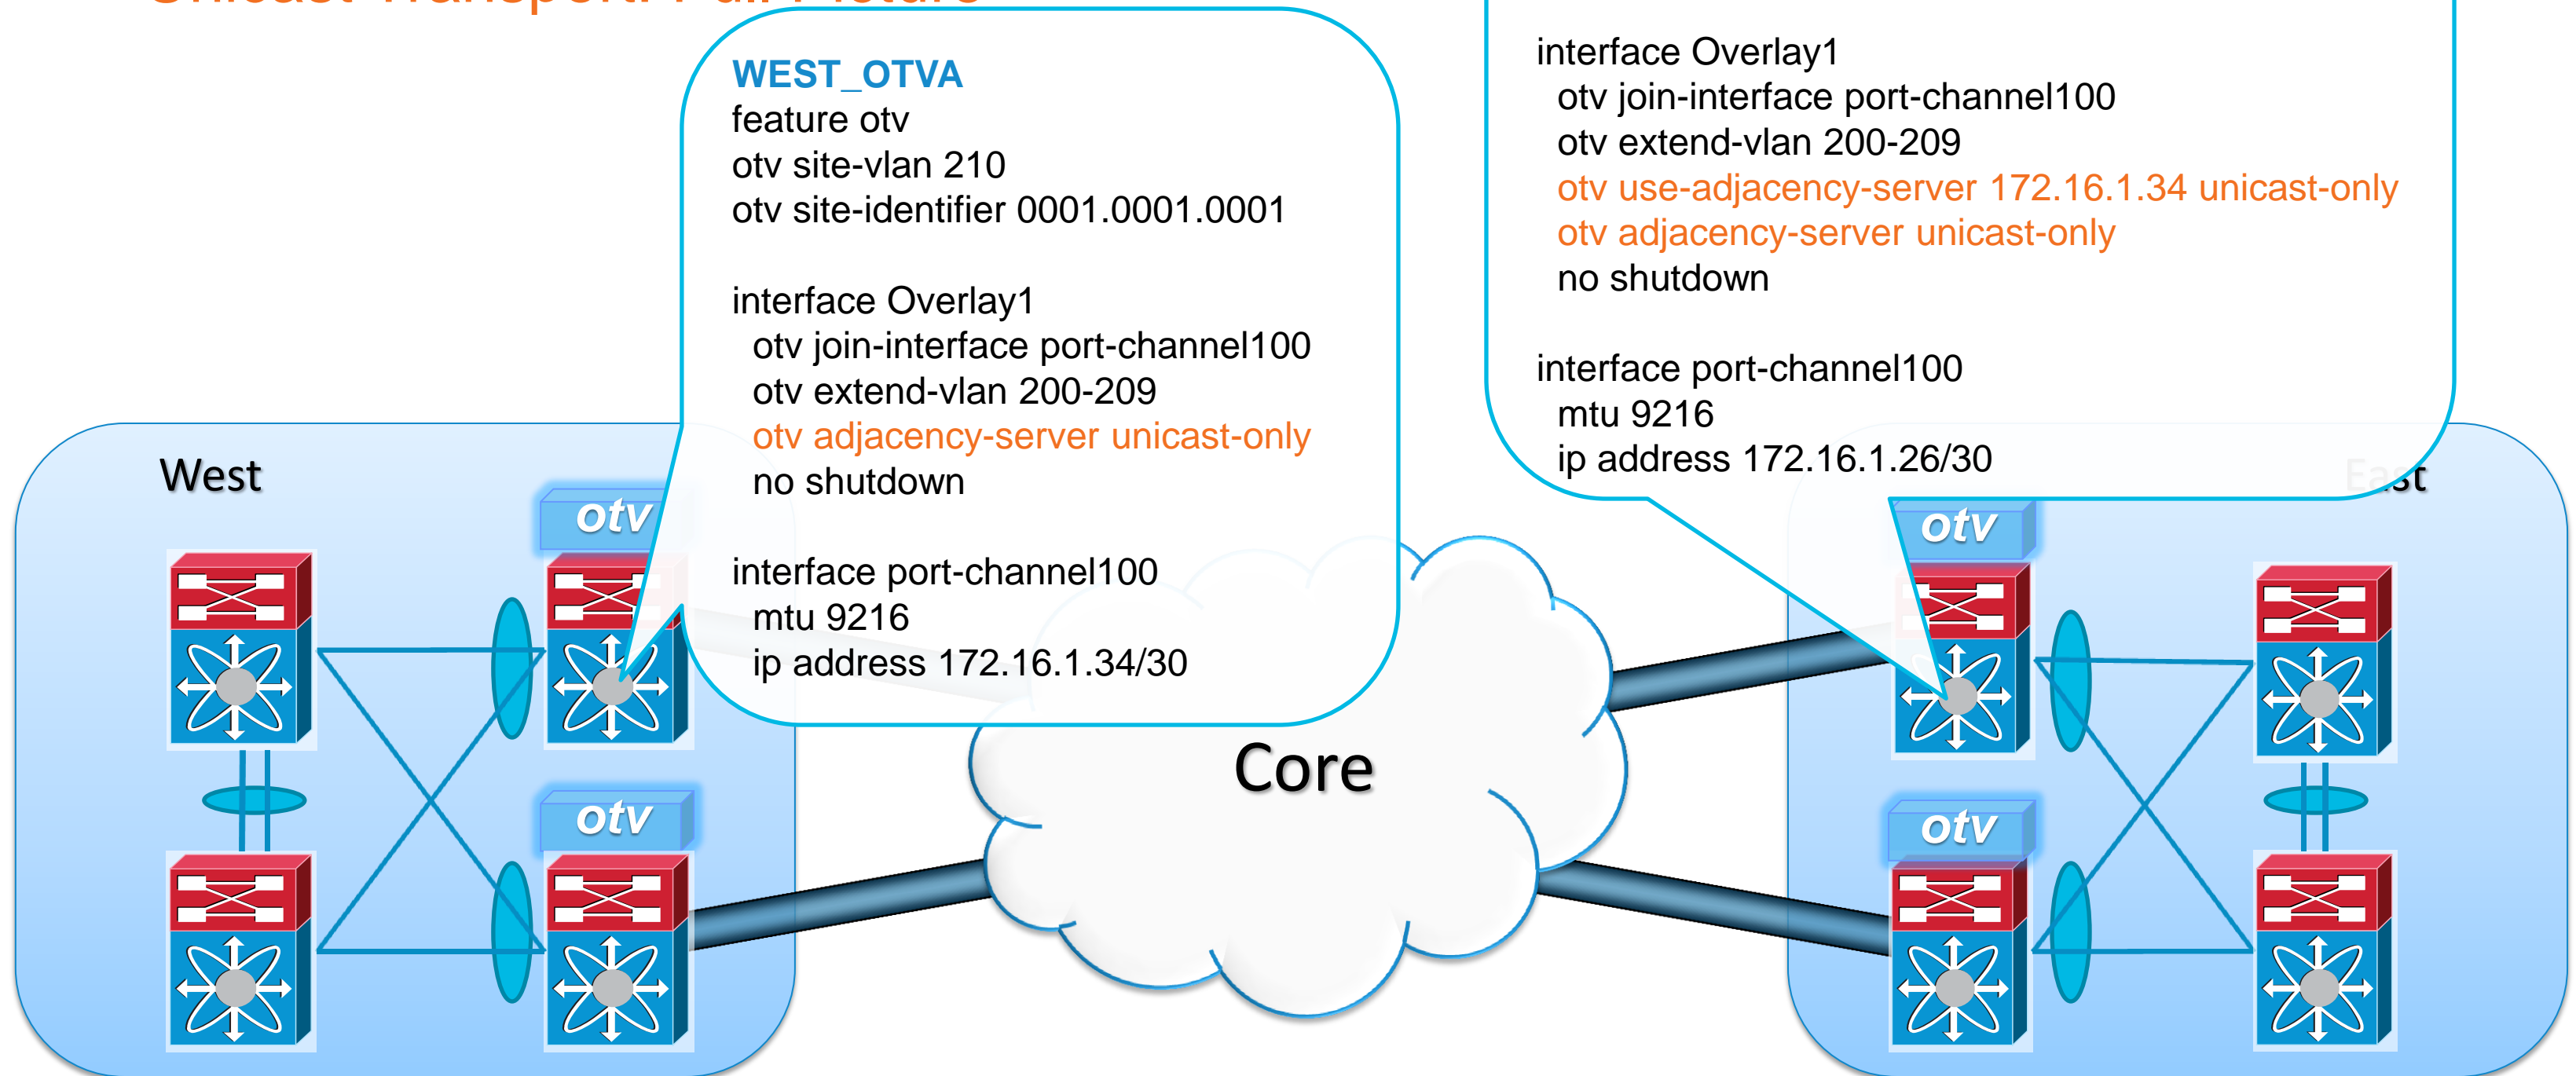
Primary Server

Secondary Server



# Configuration

## Unicast Transport: Full Picture



# Configuration

## Unicast Transport: Full Picture

### WEST\_OTVB

```
feature otv
otv site-vlan 210
otv site-identifier 0001.0001.0001
```

### interface Overlay1

```
otv join-interface port-channel100
otv extend-vlan 200-209
```

```
otv use-adjacency-server 172.16.1.34 172.16.1.26 unicast-only
no shutdown
```

### interface port-channel100

```
mtu 9216
ip address 172.16.1.38/30
```

### EAST\_OTVB

```
feature otv
otv site-vlan 210
otv site-identifier 0002.0002.0002
```

### interface Overlay1

```
otv join-interface port-channel100
otv extend-vlan 200-209
```

```
otv use-adjacency-server 172.16.1.34 172.16.1.26 unicast-only
no shutdown
```

### interface port-channel100

```
mtu 9216
ip address 172.16.1.30/30
```

Core

# Configuration

## Authentication

- OTV supports authentication of Hello messages along with authentication of PDU's

### *! Configure OTV key chain*

```
key chain OTVKeys  
key 1  
key-string 0 cisco
```

### *! Apply md5 authentication to OTV Hellos*

```
interface Overlay1  
otv isis authentication-type md5  
otv isis authentication key-chain OTVKeys
```

### *! Apply md5 authentication to OTV PDUs*

```
otv-isis default  
vpn Overlay1  
authentication-check  
authentication-type md5  
authentication key-chain OTVKeys
```

# Agenda

- **OTV Introduction**
- **Configuration**
  - Multicast Transport
  - Unicast-only Transport
- **Verification**
  - Adjacency
  - Unicast Forwarding
  - Multicast Forwarding
  - ARP
- **Troubleshooting**

# Verification

## Adjacency: IP Connectivity

- For both multicast and unicast transports, adjacencies cannot be formed without IP connectivity between the join interfaces of each OTV edge device

```
WEST_OTVA# ! Ping EAST_OTVA join interface
WEST_OTVA# ping 172.16.1.26 count 1
PING 172.16.1.26 (172.16.1.26): 56 data bytes
64 bytes from 172.16.1.26: icmp_seq=0 ttl=251 time=1.287 ms
--- 172.16.1.26 ping statistics ---
1 packets transmitted, 1 packets received, 0.00% packet loss

WEST_OTVA# ! Ping EAST_OTVB join interface
WEST_OTVA# ! Ping WEST_OTVB join interface
```



# Verification

## Adjacency: Overlay

- Verify overlay and site-vlan are up

WEST\_OTVA# `show otv` **Multicast Transport**

OTV Overlay Information  
Site Identifier 0001.0001.0001

Overlay interface Overlay1

```
VPN name           : Overlay1
VPN state         : UP
Extended vlans     : 200-209 (Total:10)
Control group      : 239.1.1.1
Data group range(s) : 232.1.1.0/24
Join interface(s)  : Po100 (172.16.1.34)
Site vlan          : 210 (up)
AED-Capable       : Yes
Capability         : Multicast-Reachable
```

WEST\_OTVA# `show otv` **Unicast Transport**

OTV Overlay Information  
Site Identifier 0001.0001.0001

Overlay interface Overlay1

```
VPN name           : Overlay1
VPN state         : UP
Extended vlans     : 200-209 (Total:10)
Join interface(s)  : Po100 (172.16.1.34)
Site vlan          : 210 (up)
AED-Capable       : Yes
Capability         : Unicast-Only
Is Adjacency Server : Yes
Adjacency Server(s) : [None] / [None]
```

# Verification

## Adjacency: ISIS Hello (IIH) statistics

```
WEST_OTVA# show otv isis site statistics | begin PDU
OTV-IS-IS PDU statistics for site-vlan:
```

| PDU            | Received | Sent  | RcvAuthErr | OtherRcvErr | ReTransmit |
|----------------|----------|-------|------------|-------------|------------|
| <b>LAN-IIH</b> | 91697    | 91700 | 0          | 0           | n/a        |
| CSNP           | 0        | 8013  | 0          | 0           | n/a        |
| PSNP           | 1        | 0     | 0          | 0           | n/a        |
| LSP            | 155      | 312   | 0          | 0           | 0          |

```
WEST_OTVA# show otv isis traffic
```

```
OTV-IS-IS process: default
```

```
VPN: Overlay1
```

```
OTV-IS-IS Traffic:
```

| PDU            | Received | Sent  | RcvAuthErr | OtherRcvErr | ReTransmit |
|----------------|----------|-------|------------|-------------|------------|
| <b>LAN-IIH</b> | 85530    | 23298 | 0          | 0           | n/a        |
| CSNP           | 3        | 8015  | 0          | 0           | n/a        |
| PSNP           | 17       | 2     | 0          | 0           | n/a        |
| LSP            | 896      | 393   | 0          | 0           | 0          |

# Verification

## Adjacency: ISIS Hello over Multicast Transport

The image shows a Wireshark packet capture of an ISIS Hello packet. The packet list pane shows three packets, with packet 32 selected. The packet details pane shows the following structure:

- Frame 32 (1456 bytes on wire, 1456 bytes captured)
- Ethernet II, Src: 6c:9c:ed:40:17:44 (6c:9c:ed:40:17:44), Dst: IPv4mcast\_01:01:01:00:00:00
- Internet Protocol Version 4, Src: 172.16.1.34 (172.16.1.34), Dst: 239.1.1.1 (239.1.1.1)
- Differentiated Services Field: 0x50 (Class Selector 6; ECN: 0x00)
- Total Length: 1442
- Identification: 0xd965 (55653)
- Flags: 0x02 (Don't Fragment)
- Fragment offset: 0
- Time to live: 254
- Protocol: GRE (0x2f)
- Header checksum: 0xffd1 [correct]
- Source: 172.16.1.34 (172.16.1.34)
- Destination: 239.1.1.1 (239.1.1.1)
- Generic Routing Encapsulation (0x8848 - unknown)
- MultiProtocol Label Switching Header, Label: 1 (Router Alert), Exp: 0, S: 1, TTL: 255
- PW Ethernet Control word
- IEEE 802.3 Ethernet
- Logical-Link Control
- Data (895 bytes)

The packet bytes pane shows the following data:

```
0000 01 00 5e 01 01 01 6c 9c ed 40 17 44 08 00 45 c0 ..^...l. .@.D..E.
0010 05 a2 d9 65 40 00 fe 2f ff d1 ac 10 01 22 ef 01 ...e@../ .....".
0020 01 01 00 00 88 48 00 00 11 ff 01 00 0c df df df .....H. ....
0030 6c 9c ed 40 17 42 05 78 fe fe 03 83 1b 01 00 0f l..@.B.x .....
0040 01 00 00 01 6c 9c ed 40 17 42 00 0c 05 75 40 6c ....l..@ .B...u@l
0050 9c ed 40 17 42 01 01 02 01 00 81 02 cc 8e 84 04 ..@.B...
0060 ac 10 01 22 06 12 64 a0 e7 41 c8 41 64 a0 e7 41 ...u..
0070 c8 42 6c 9c ed 40 17 41 8f 0d 00 00 fa 09 00 01 .Bl..@.A .....
0080 00 01 00 01 00 00 02 d3 01 00 08 ff 00 00 00 00 .....

```

ISIS Hellos have packet size of 1442 Bytes through multicast transport  
*1400B ISIS + 42B OTV header*

OTV ISIS sent on multicast control-group, sourced from join interface

System ID and Site Identifier included in Hello



# Verification

## Adjacency: ISIS Hello over Unicast Transport

| No.  | size | Time            | Source      | Destination | dscp | Protocol | Info                                |
|------|------|-----------------|-------------|-------------|------|----------|-------------------------------------|
| 645  | 134  | 10:23:17.968202 | 172.16.1.34 | 172.16.1.30 | 48   | UDP      | Source port: 58518 Destination port |
| 646  | 134  | 10:23:17.968290 | 172.16.1.34 | 172.16.1.38 |      |          |                                     |
| 1249 | 1464 | 10:23:19.017068 | 172.16.1.26 | 172.16.1.34 |      |          |                                     |
| 1250 | 1464 | 10:23:19.017196 | 172.16.1.26 | 172.16.1.30 |      |          |                                     |
| 1252 | 1464 | 10:23:19.017352 | 172.16.1.26 | 172.16.1.38 |      |          |                                     |

Frame 1249 (1464 bytes on wire, 1464 bytes captured)

Ethernet II, Src: 6c:9c:ed:40:17:43 (6c:9c:ed:40:17:43), Dst: Cisco\_be:52:c1 (00:13:00:00:00:00)

Internet Protocol, Src: 172.16.1.26 (172.16.1.26), Dst: 172.16.1.34 (172.16.1.34)

- Version: 4
- Header length: 20 bytes
- Differentiated Services Field: 0xc0000000: Class Selector 6; ECN: 0x00
- Total Length: 1450
- Identification: 0x1b98 (7064)
- Flags: 0x02 (Don't Fragment)
- Fragment offset: 0
- Time to live: 254
- Protocol: UDP (0x11)
- Header checksum: 0x008e [correct]
- Source: 172.16.1.26 (172.16.1.26)
- Destination: 172.16.1.34 (172.16.1.34)

User Datagram Protocol, Src Port: 59706 (59706), Dst Port: otv (8472)

Data (1422 bytes)

Data: 000000010000000001000CDFDFDF6C9CED4017410004FEFE...

[Length: 1422]

```
0000 00 1a e2 be 52 c1 6c 9c ed 40 17 43 08 00 45 c0 .....R.l. .@.C..E.
0010 05 aa 1b 98 40 00 fe 11 00 8e ac 10 01 1a ac 10 .....@... ..
0020 01 22 e9 3a 21 18 05 96 00 00 00 00 00 01 00 00 ..":!... ..
0030 00 00 01 00 0c df df df 6c 9c ed 40 17 41 00 04 .....l..@.A..
0040 fe fe 03 83 1b 01 00 0f 01 00 00 01 6c 9c ed 40 .....l @
0050 17 41 00 0c 05 75 40 6c 9c ed 40 17 42 01 01 02 .....
0060 01 00 81 02 cc 8e 84 04 ac 10 01 1a 06 12 64 a0 .....d.
0070 e7 41 c8 41 64 a0 e7 41 c8 42 6c 9c ed 40 17 42 .A.Ad..A .Bl..@.B
0080 8f 1e 00 00 fa 09 00 02 00 02 00 02 00 00 03 fd .....
0090 0f ac 10 01 1e 01 ac 10 01 26 01 ac 10 01 22 01 .....&....".
00a0 d3 01 00 08 ff 00 00 00 00 00 00 00 00 00 00 00 .....

```

ISIS Hellos have packet size of 1450 Bytes through unicast transport  
*1400B ISIS + 42B OTV header + 8B UDP*

OTV ISIS via unicast, sourced and destined between join interfaces

OTV ISIS Hello sent encapsulated in UDP unicast on port 8472

System ID and Site Identifier included in Hello



# Verification


## Adjacency: ISIS Overlay Adjacencies

```
WEST_OTVA# show otv isis adjacency
```

```
OTV-IS-IS process: default VPN: Overlay1
```

```
OTV-IS-IS adjacency database:
```

| System ID        | SNPA           | Level | State | Hold Time | Interface | Site-ID        |
|------------------|----------------|-------|-------|-----------|-----------|----------------|
| <b>EAST_OTVB</b> | 64a0.e741.c841 | 1     | UP    | 00:00:11  | Overlay1  | 0002.0002.0002 |
| <b>WEST_OTVB</b> | 64a0.e741.c842 | 1     | UP    | 00:00:09  | Overlay1  | 0001.0001.0001 |
| <b>EAST_OTVA</b> | 6c9c.ed40.1741 | 1     | UP    | 00:00:13  | Overlay1  | 0002.0002.0002 |



# Verification

## Adjacency: OTV Overlay Adjacencies

```
WEST_OTVA# show otv adjacency
```

```
Overlay Adjacency database
```

```
Overlay-Interface Overlay1 :
```

```
Hostname
```

```
System-ID
```

```
Dest Addr
```

```
Up Time
```

```
State
```

```
EAST_OTVA
```

```
6c9c.ed40.1741
```

```
172.16.1.26
```

```
19:34:34
```

```
UP
```

```
EAST_OTVB
```

```
64a0.e741.c841
```

```
172.16.1.30
```

```
19:34:30
```

```
UP
```


```
WEST_OTVB
```

```
64a0.e741.c842
```

```
172.16.1.38
```

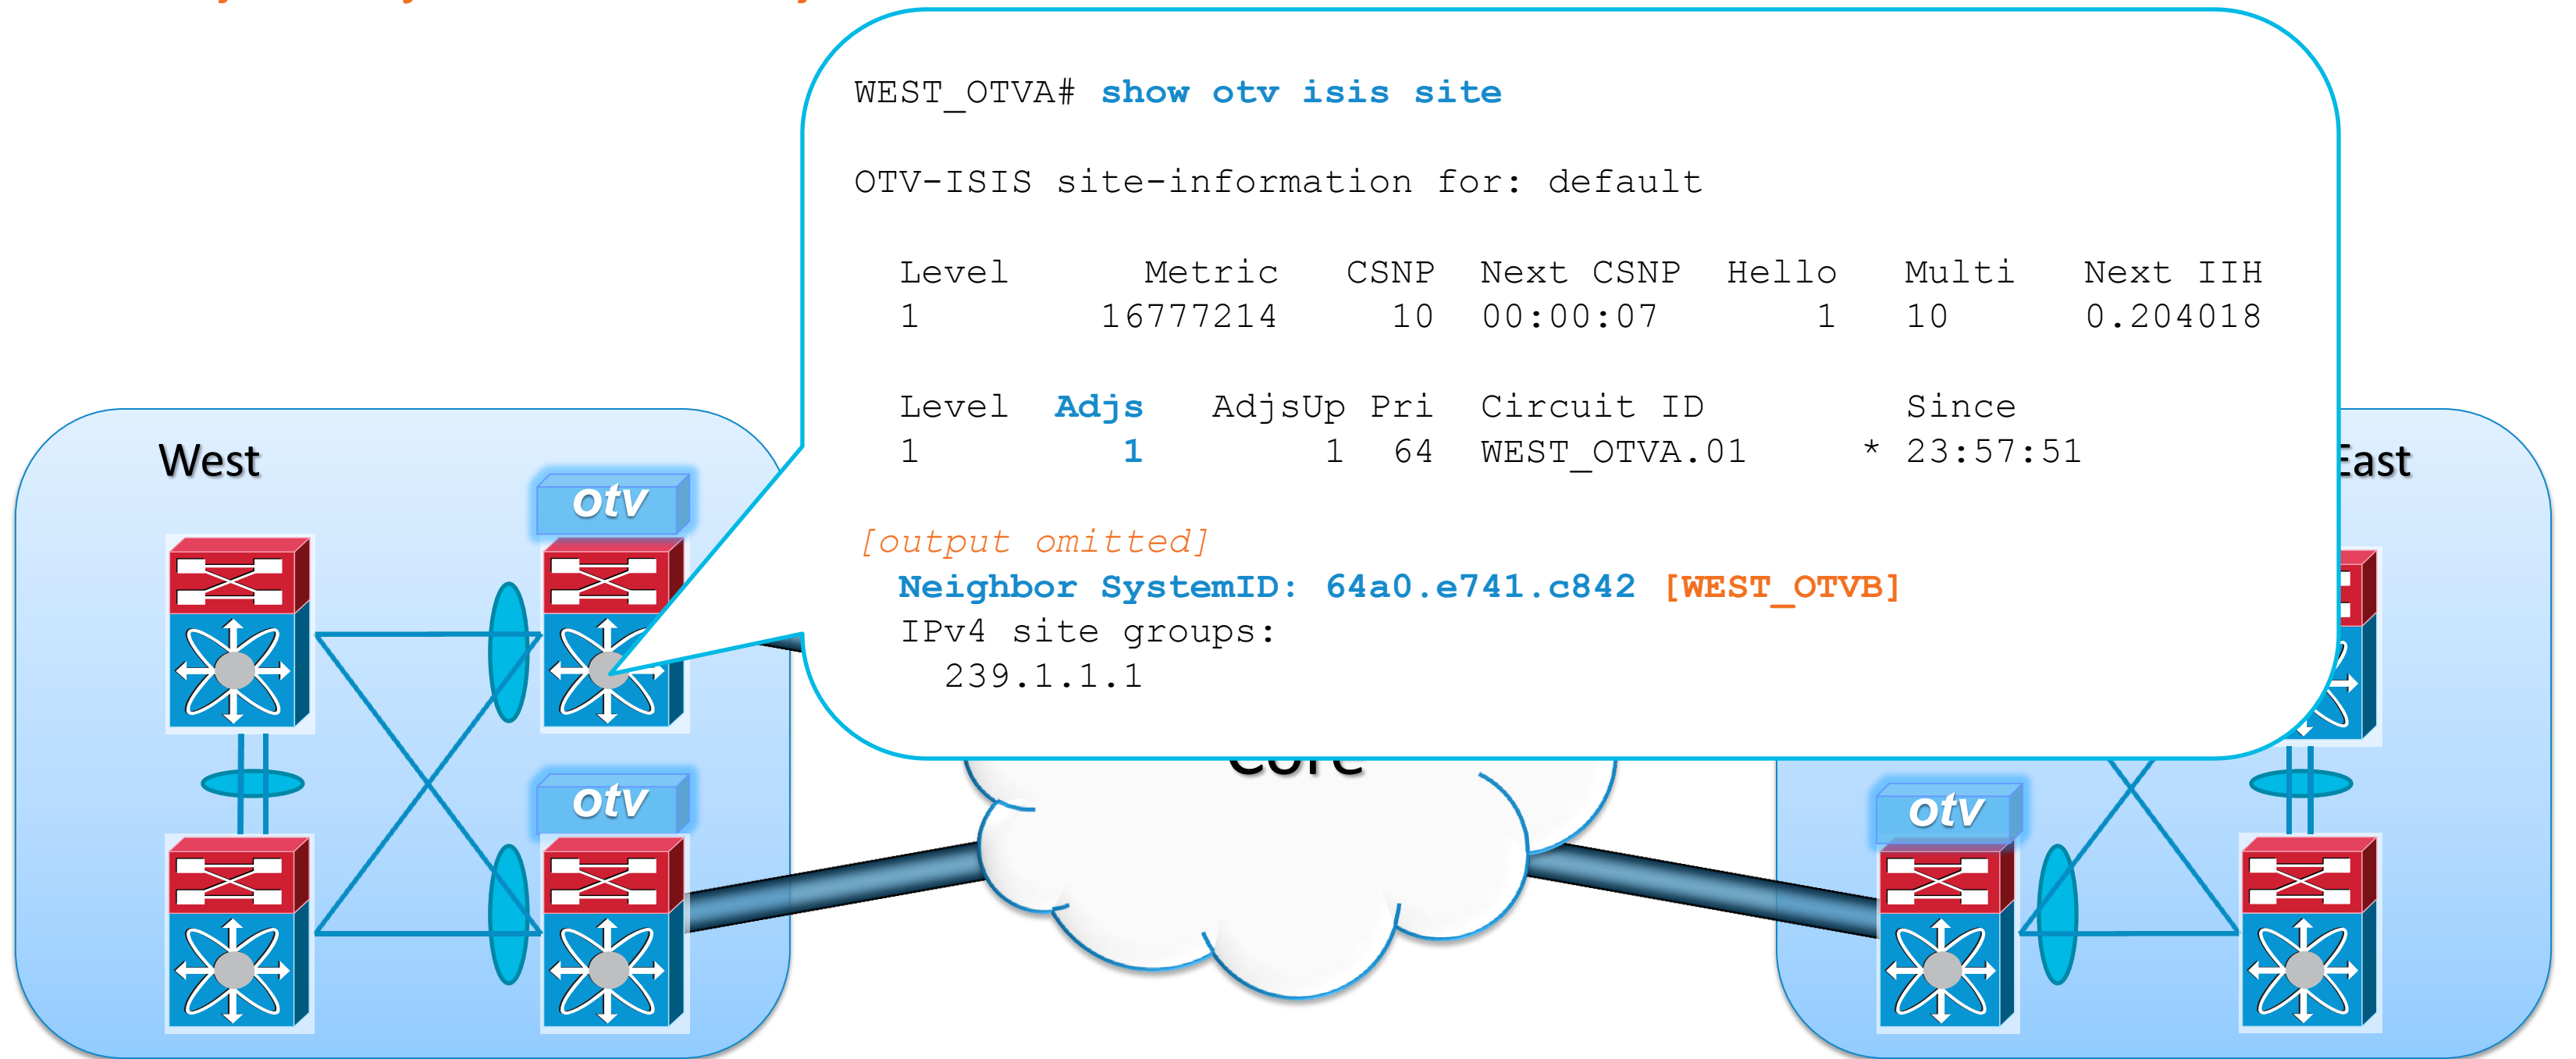
```
19:34:30
```

```
UP
```



# Verification

## Adjacency: ISIS Site Adjacencies



```
WEST_OTVA# show otv isis site
```

```
OTV-ISIS site-information for: default
```

| Level | Metric   | CSNP | Next CSNP | Hello | Multi | Next IIH |
|-------|----------|------|-----------|-------|-------|----------|
| 1     | 16777214 | 10   | 00:00:07  | 1     | 10    | 0.204018 |

| Level | Adjs | AdjsUp | Pri | Circuit ID   | Since      |
|-------|------|--------|-----|--------------|------------|
| 1     | 1    | 1      | 64  | WEST_OTVA.01 | * 23:57:51 |

*[output omitted]*

```
Neighbor SystemID: 64a0.e741.c842 [WEST_OTVB]
```

```
IPv4 site groups:  
239.1.1.1
```

# Verification

## Adjacency: OTV Site Adjacencies

```
WEST_OTVA# show otv site
```

```
Dual Adjacency State Description
```

- Full - Both site and overlay adjacency up
- Partial - Either site/overlay adjacency down
- Down - Both adjacencies are down (Neighbor is down/unreachable)
- (!) - Site-ID mismatch detected

```
Local Edge Device Information:
```

```
Hostname WEST_OTVA  
System-ID 6c9c.ed40.1742  
Site-Identifier 0001.0001.0001  
Site-VLAN 210 State is Up
```

```
Site Information for Overlay1:
```

```
Local device is AED-Capable
```

```
Neighbor Edge Devices in Site: 1
```

| Hostname         | System-ID             | Adjacency-State | Adjacency-Uptime | AED-Capable |
|------------------|-----------------------|-----------------|------------------|-------------|
| -----            |                       |                 |                  |             |
| <b>WEST_OTVB</b> | <b>64a0.e741.c842</b> | <b>Full</b>     | 23:57:51         | <b>Yes</b>  |




# Verification

## Authoritative Edge Device (AED)

```
WEST_OTVA# show otv vlan 200-201
```

```
OTV Extended VLANs and Edge Device State Information (* - AED)
```

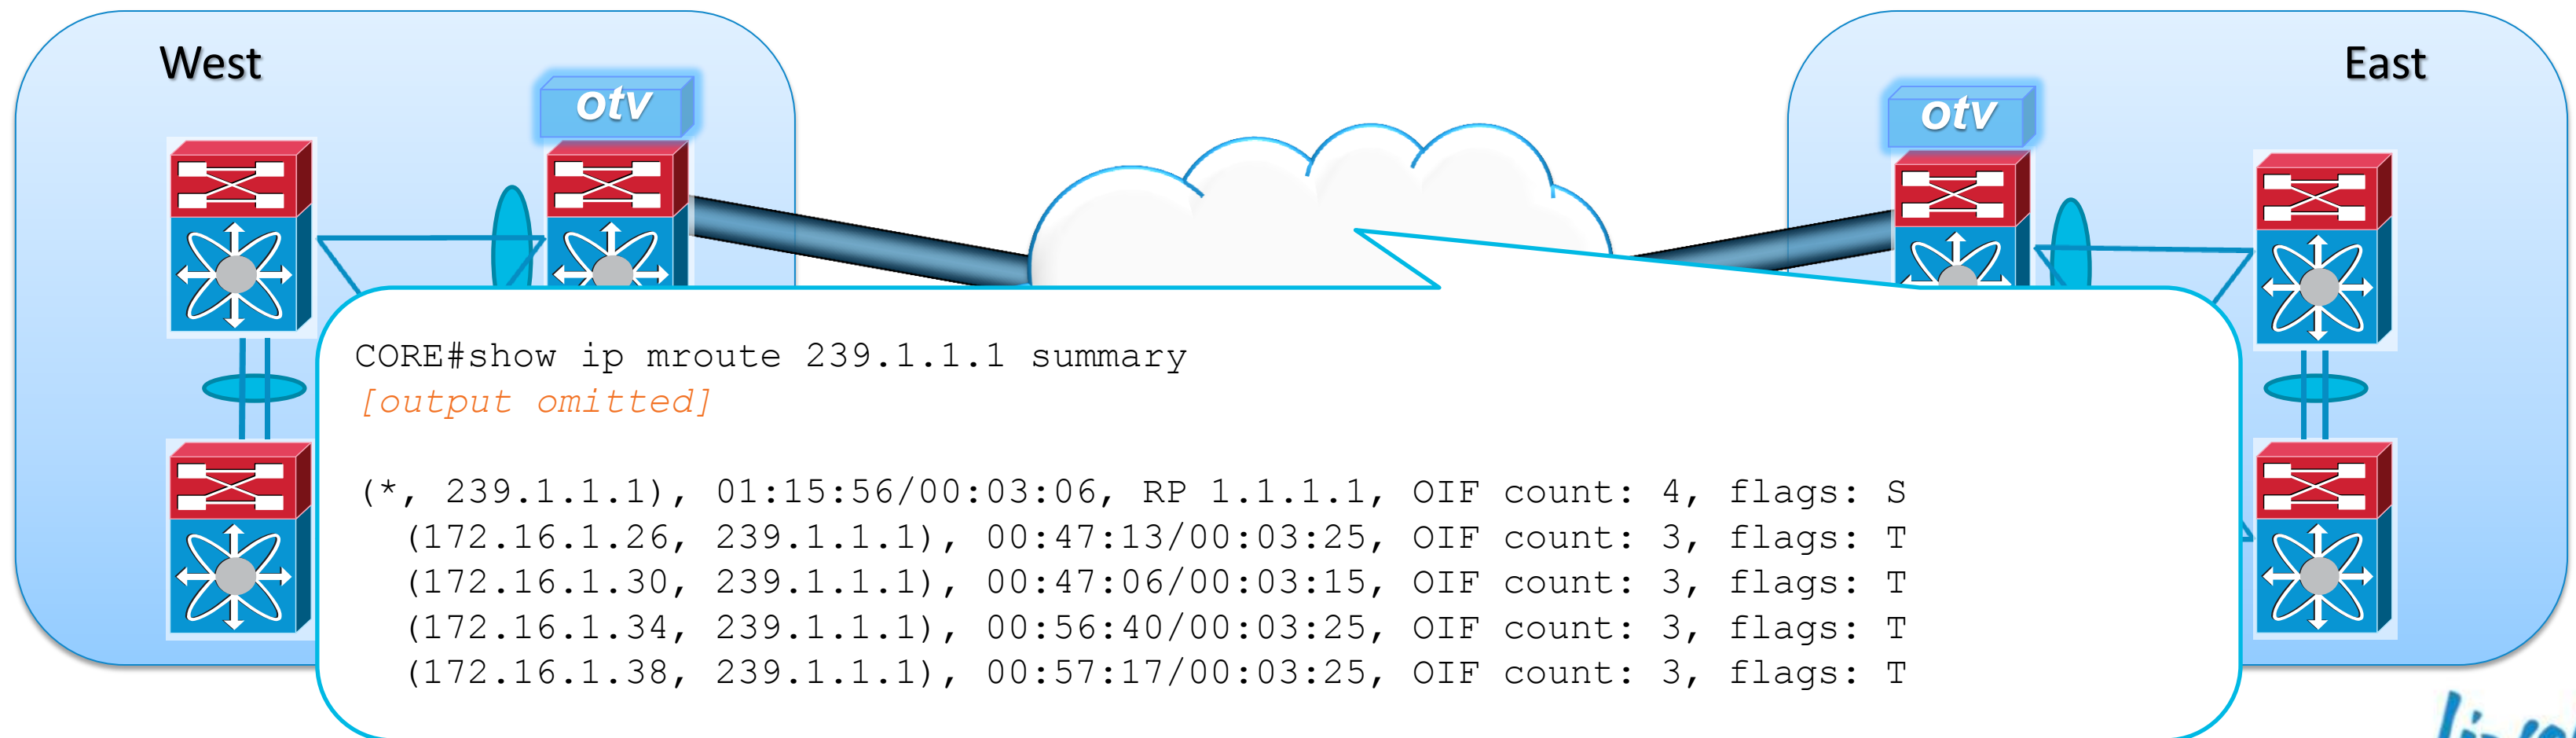
| VLAN | Auth. Edge Device | Vlan State                | Overlay  |
|------|-------------------|---------------------------|----------|
| 200  | WEST_OTVB         | <b>inactive (Non AED)</b> | Overlay1 |
| 201* | WEST_OTVA         | <b>active</b>             | Overlay1 |



# Verification

## Adjacency: OTV Overlay Adjacencies for Multicast Transport

- For multicast transport, OTV join interfaces are configured with IGMPv3. Therefore, from the transport's perspective, the OTV edge devices appear as host sending and requesting traffic from the control-group

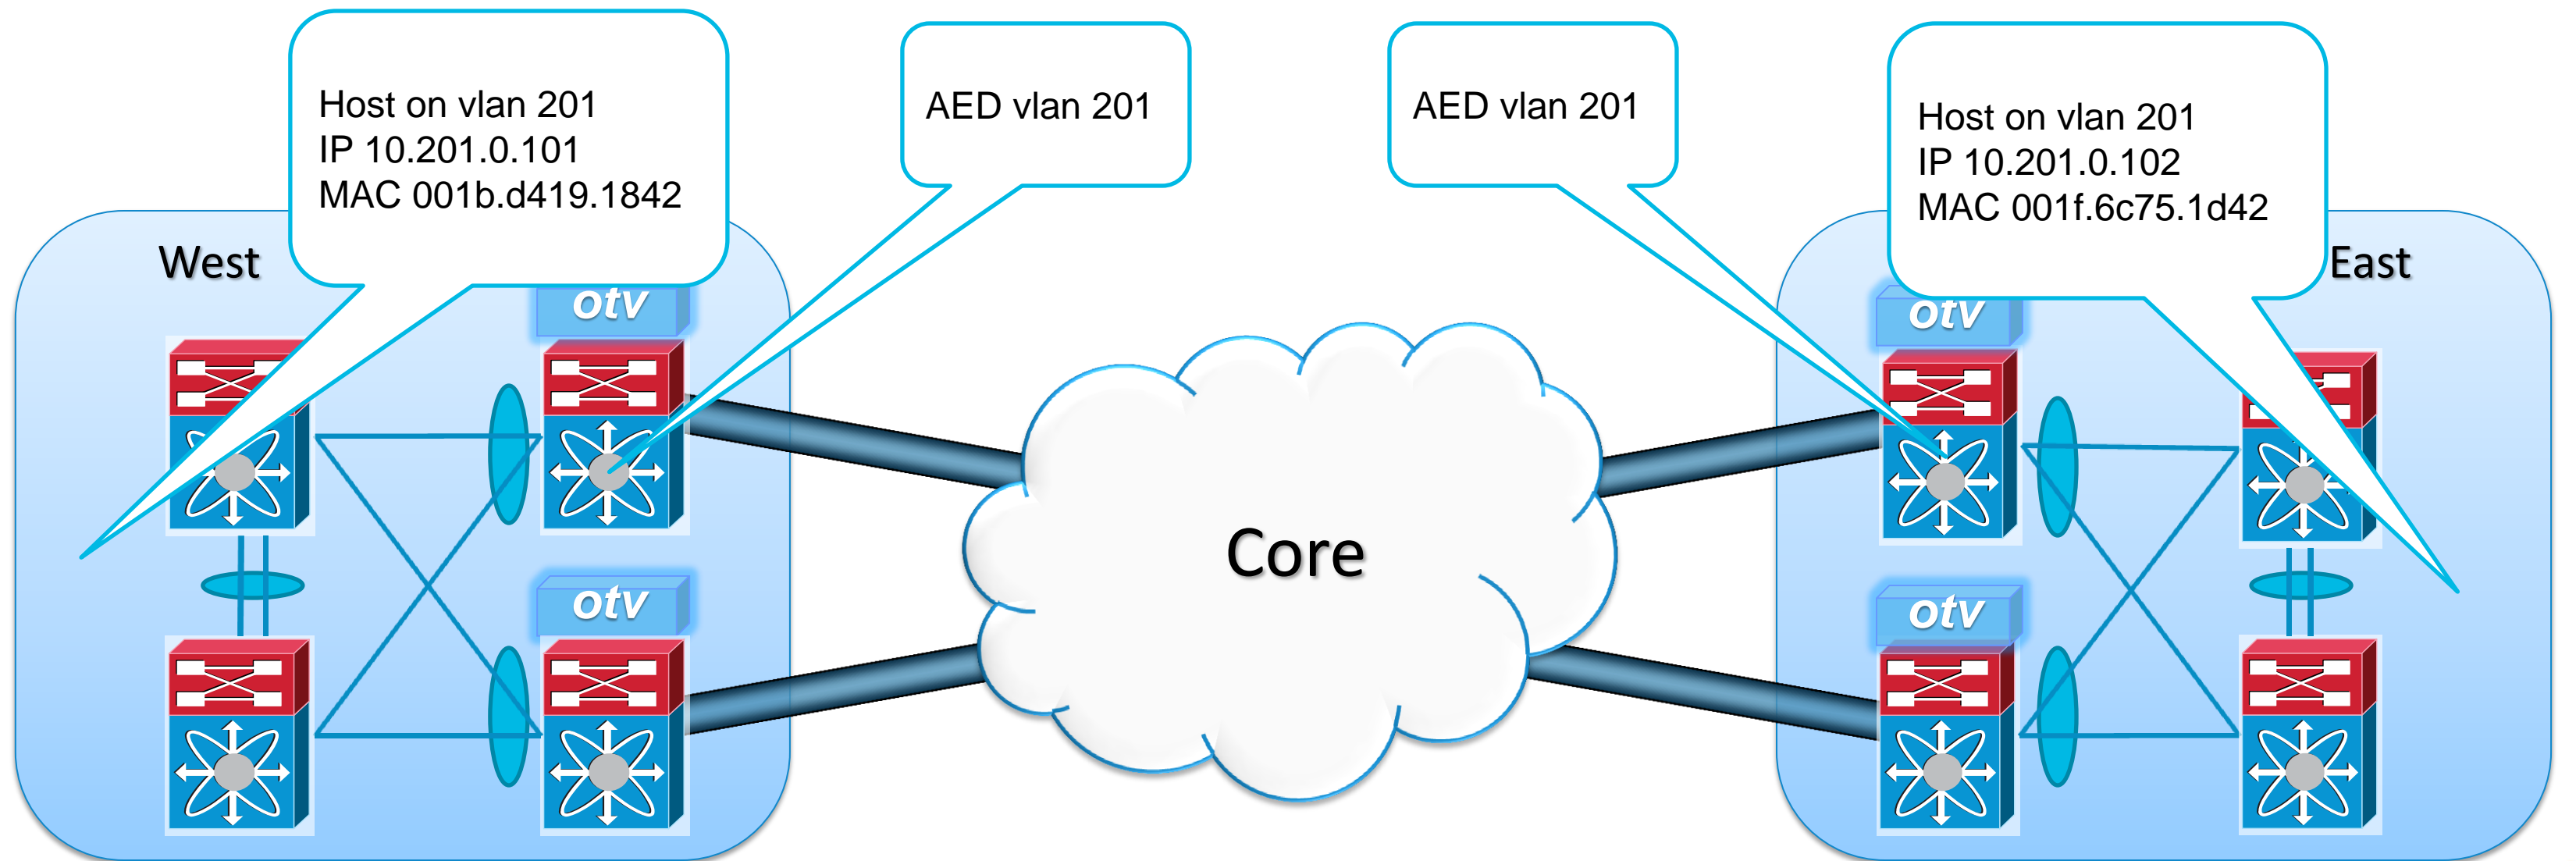


# Agenda

- **OTV Introduction**
- **Configuration**
  - Multicast Transport
  - Unicast-only Transport
- **Verification**
  - Adjacency
  - Unicast Forwarding
  - Multicast Forwarding
  - ARP
- **Troubleshooting**

# Verification


## Unicast Forwarding



# Verification

## Unicast Forwarding

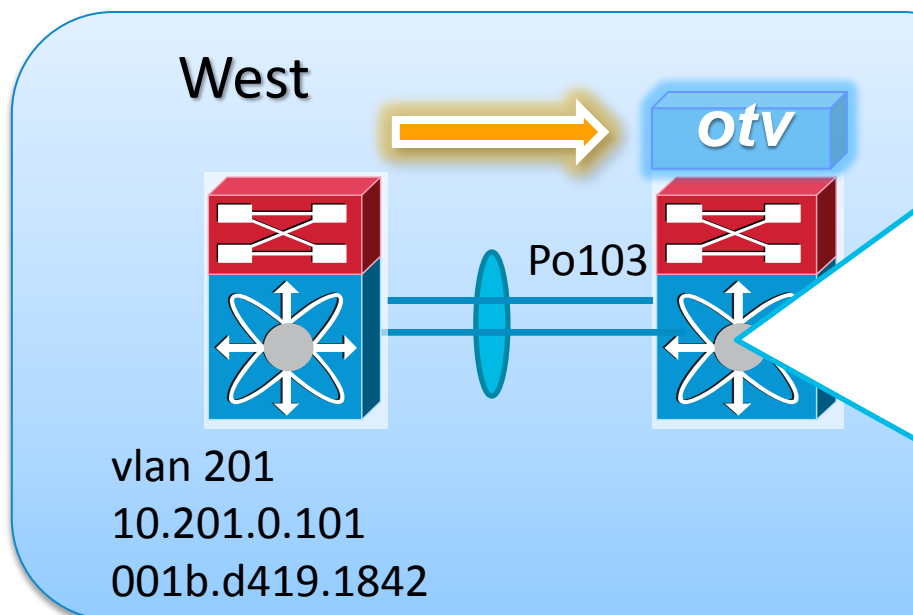
- Simplified topology based on AED for vlan 201
- Let's assume that ARP has already resolved between hosts



# Verification

## Unicast Forwarding: User on Site West Sends Unicast Packet to Site East

- Packet is received on internal interface on OTV AED
- Verify CAM entry and OTV route



```
WEST_OTVA# show mac address-table address 001f.6c75.1d42 vlan 201
```

Legend:

\* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC  
age - seconds since last seen,+ - primary entry using vPC Peer-Link

| VLAN  | MAC Address    | Type    | age | Secure | NTFY | Ports/SWID.SSID.LID |
|-------|----------------|---------|-----|--------|------|---------------------|
| O 201 | 001f.6c75.1d42 | dynamic | 0   | F      | F    | Overlay1            |

```
WEST_OTVA# show otv route 001f.6c75.1d42
```

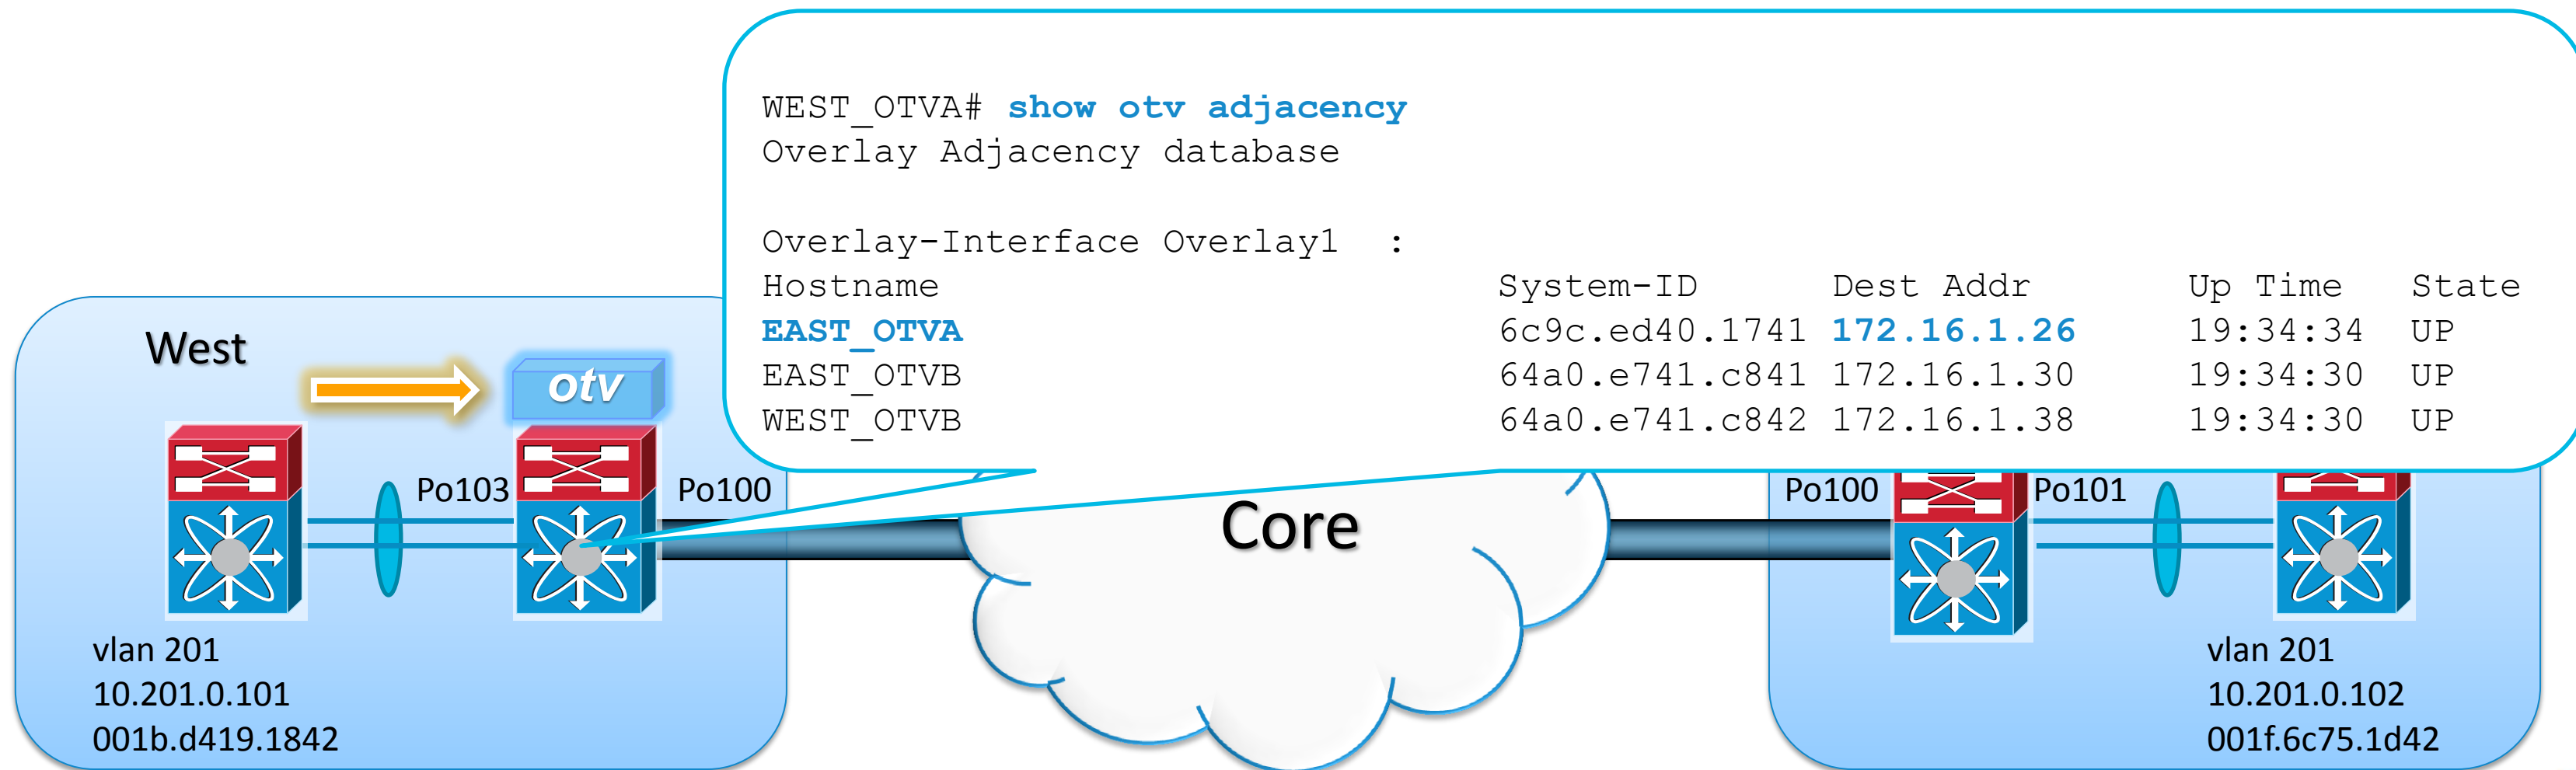
OTV Unicast MAC Routing Table For Overlay1

| VLAN | MAC-Address    | Metric | Uptime   | Owner   | Next-hop(s) |
|------|----------------|--------|----------|---------|-------------|
| 201  | 001f.6c75.1d42 | 42     | 00:02:25 | overlay | EAST_OTVA   |

# Verification

## Unicast Forwarding: User on Site West Sends Unicast Packet to Site East

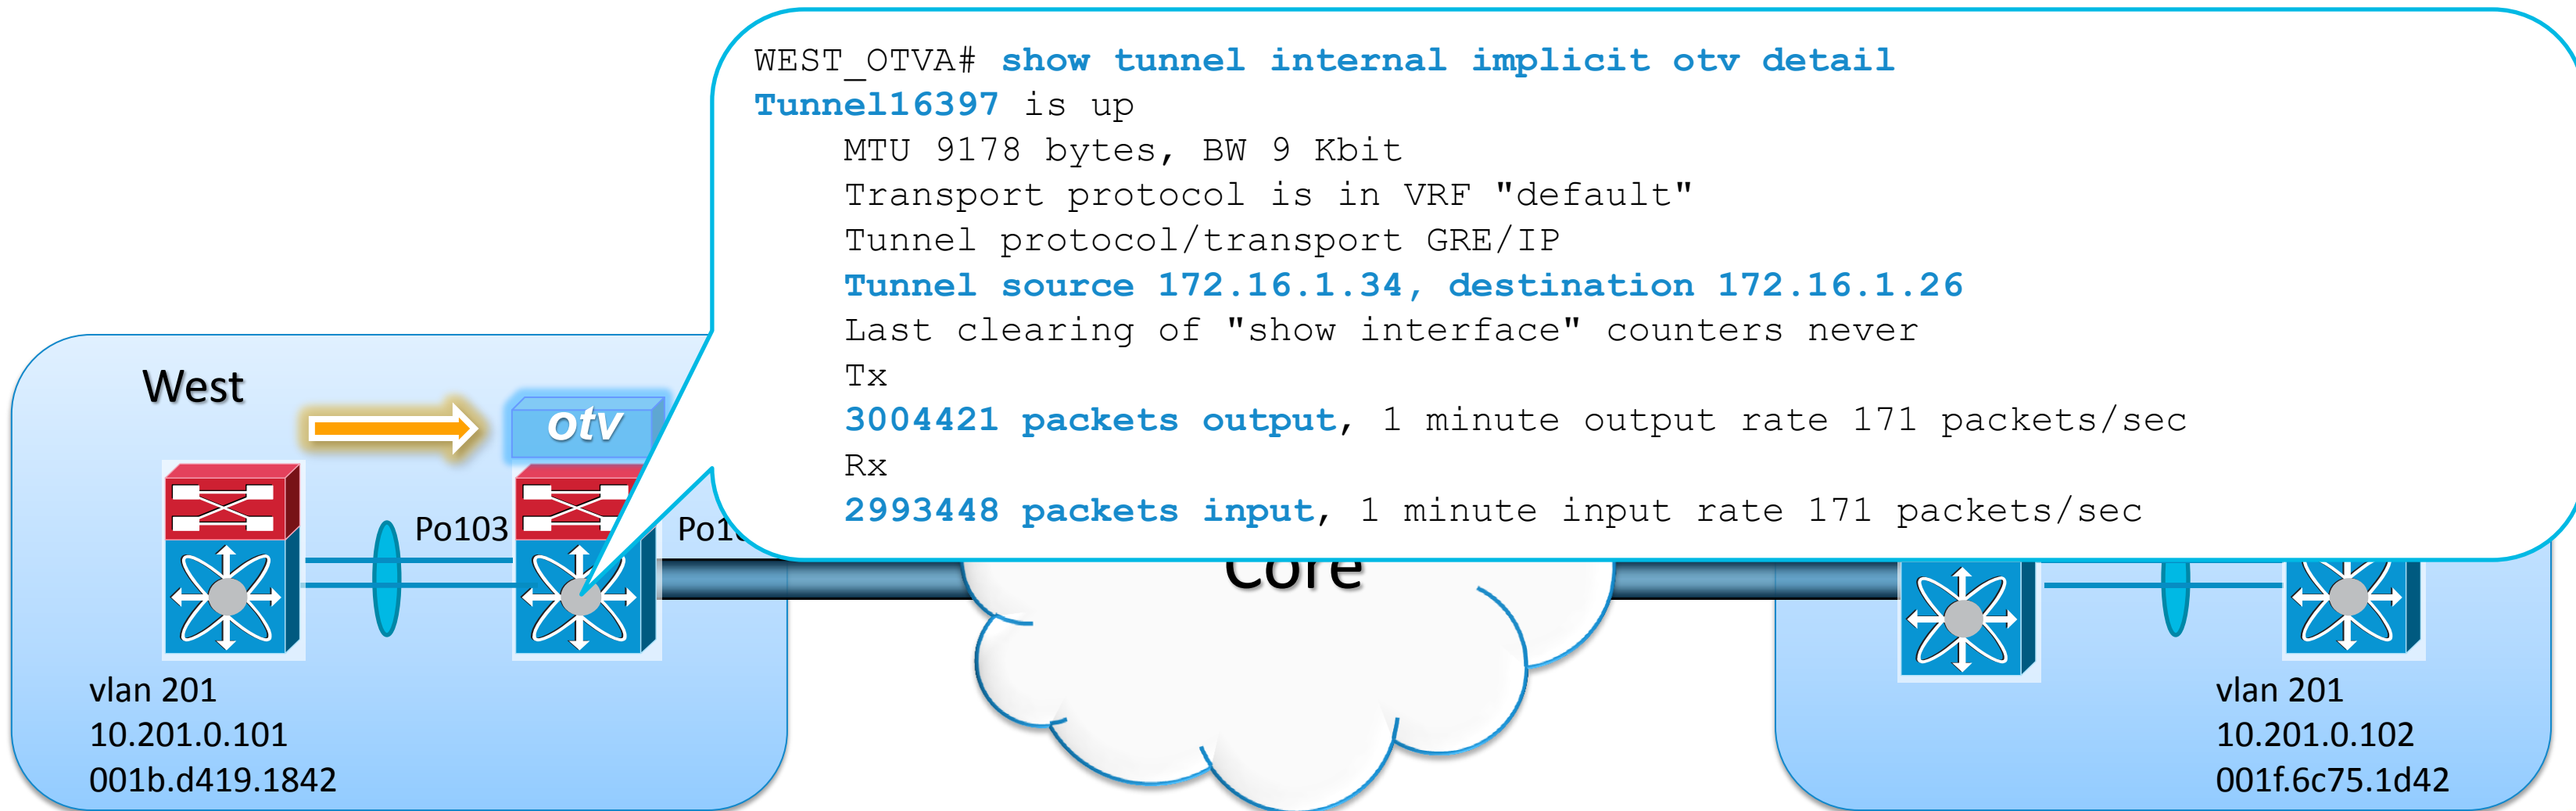
- Verify next hop IP address for MAC



# Verification

## Unicast Forwarding: User on Site West Sends Unicast Packet to Site East

- Verify incrementing counters on tunnel interface





# Verification


## Unicast Forwarding: User on Site West Sends Unicast Packet to Site East

- Verify hardware adjacency and label

OTV shim includes an MPLS label to identify the encapsulated VLAN.  
MPLS Label = 32 + vlan

$$32 + 201 = 233 = 0xe9$$

West



vlan 201  
10.201.0.101  
001b.d419.1842

```
WEST_OTVA# show system internal forwarding otv overlay 1 vlan 201 detail
Dev No: 1  MPLS Label: 0xe9  VPN id: 10
          TCAM Idx: 0x1ffa9  Adj Idx: 0x4302a  Type: EoMPLS
          rc: 1  ccc: 6  pv: 0  l2_fwd: 1
          Packets: 0  Bytes: 0

(BD: 21, Vlan Id: 201, Peer Id: 1)
DevNo: 1  TCAM Idx: 0x1ff03
Adj Idx: 0x4303d  Egress Lif: 0x8004  RDT: 1  Egress Lif Base: 5
ccc: 0  pv: 1  l2_fwd: 1  rc: 1
label hi: 0x0  label low: 0xe90  Packets: 0  Bytes: 0
Tunnel adj index: 0x43037  GRE: YES
DIP: 172.16.1.26  SIP: 172.16.1.34
LIF: 0x4074  DI: 0x421  ccc: 6  L2_FWD: NO  RDT: YES
Packets: 0  Bytes: 0  zone enforce: 0
```

10.201.0.102  
001f.6c75.1d42

# Verification

## Unicast Forwarding: Encapsulated Packet

| No. | size | Time            | Source       | Destination  | dscp | Protocol | Info                |
|-----|------|-----------------|--------------|--------------|------|----------|---------------------|
| 8   | 138  | 06:21:02.090802 | 10.201.0.102 | 10.201.0.101 | 0    | ICMP     | Echo (ping) request |
| 9   | 156  | 06:21:02.093258 | 10.201.0.101 | 10.201.0.102 | 0    | ICMP     | Echo (ping) reply   |
| 10  | 156  | 06:21:02.095519 | 10.201.0.102 | 10.201.0.101 | 0    | ICMP     | Echo (ping) request |
| 11  | 156  | 06:21:02.097489 | 10.201.0.101 | 10.201.0.102 | 0    | ICMP     | Echo (ping) reply   |
| 12  | 156  | 06:21:02.099450 | 10.201.0.102 | 10.201.0.101 | 0    | ICMP     | Echo (ping) request |

Frame 9 (156 bytes on wire, 156 bytes captured)

- Ethernet II, Src: 6c:9c:ed:40:17:44 (6c:9c:ed:40:17:44), Dst: Cisco\_be:52:c3 (00:1a:e2:b3:52:c3)
- Internet Protocol, Src: 172.16.1.34 (172.16.1.34), Dst: 172.16.1.26 (172.16.1.26)
  - Version: 4
  - Header length: 20 bytes
  - Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
  - Total Length: 142
  - Identification: 0x3bad (15277)
  - Flags: 0x02 (Don't Fragment)
  - Fragment offset: 0
  - Time to live: 253
  - Protocol: GRE (0x2f)
  - Header checksum: 0xe736 [correct]
  - Source: 172.16.1.34 (172.16.1.34)
  - Destination: 172.16.1.26 (172.16.1.26)
- Generic Routing Encapsulation (MPLS label switched packet)
- MultiProtocol Label Switching Header, Label: 233, Exp: 0, S: 1, TTL: 254
- Ethernet II, Src: Cisco\_19:18:42 (00:1b:d4:19:18:42), Dst: Cisco\_75:1d:42 (00:1f:6c:75:1d:42)
- Internet Protocol, Src: 10.201.0.101 (10.201.0.101), Dst: 10.201.0.102 (10.201.0.102)
  - Version: 4
  - Header length: 20 bytes
  - Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
  - Total Length: 100
  - Identification: 0x099c (2460)
  - Flags: 0x00
  - Fragment offset: 0
  - Time to live: 255
  - Protocol: ICMP (0x01)
  - Header checksum: 0x9ba0 [correct]
  - Source: 10.201.0.101 (10.201.0.101)
  - Destination: 10.201.0.102 (10.201.0.102)
- Internet Control Message Protocol

42 Byte OTV header overhead  
Increases packet size to 142 Bytes

Packet sent unicast through  
transport with endpoint IP's of  
AED Join interface

Label to identify vlan  
 $32+201 = 233$

Original IP header maintained




# Verification

Unicast Forwarding: User on site West sends unicast packet to site East

- Encapsulated packet is received on Join interface of East AED
- For Decapsulation, verify hits against internal ACL

```
EAST_OTVA# show system internal access-list output statistics | begin Tcam
Tcam 0 resource usage:
-----
Label_a = 0x802
Bank 0
-----
IPv4 Class
Policies: Tunnel Decap(Tunnel Decap on VRF default)
Entries:
[Index] Entry [Stats]
-----
[0006] redirect(0x4307d) 47 172.16.1.30/32 172.16.1.26/32 [0]
[0007] redirect(0x4307b) 47 172.16.1.38/32 172.16.1.26/32 [0]
[0009] redirect(0x43079) 47 172.16.1.34/32 172.16.1.26/32 [18505]
[0010] permit ip 0.0.0.0/0 0.0.0.0/0 [1444]
```



# Verification

Unicast Forwarding: User on site West sends unicast packet to site East

- Verify OTV route for entry points to local site
- Verify CAM entry for destination MAC points out internal interface

```
EAST_OTVA# show otv route 001f.6c75.1d42 vlan 201
```

```
OTV Unicast MAC Routing Table For Overlay1
```

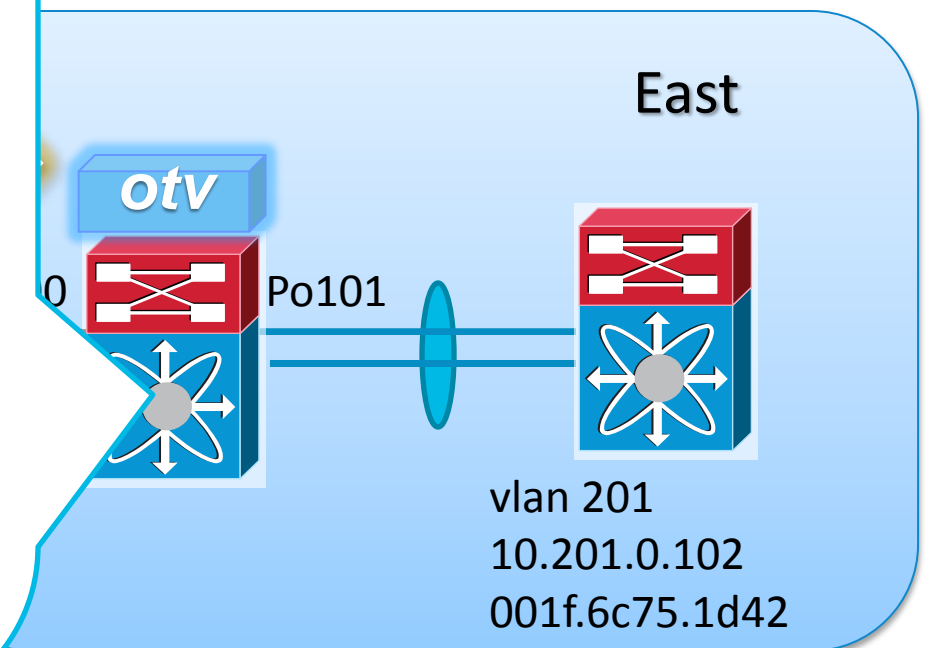
| VLAN | MAC-Address    | Metric | Uptime   | Owner | Next-hop(s)     |
|------|----------------|--------|----------|-------|-----------------|
| 201  | 001f.6c75.1d42 | 1      | 00:01:50 | site  | port-channel101 |

```
EAST_OTVA# show mac address-table address 001f.6c75.1d42 vlan 201
```

Legend:

\* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC  
age - seconds since last seen, + - primary entry using vPC Peer-Link


| VLAN  | MAC Address    | Type    | age | Secure | NTFY | Ports/SWID.SSID.LID |
|-------|----------------|---------|-----|--------|------|---------------------|
| * 201 | 001f.6c75.1d42 | dynamic | 0   | F      | F    | Po101               |



# Verification

Unicast Forwarding: User on site West sends unicast packet to site East

- Packet will be sent out internal interface at site East and L2 switched to the host
- Return path from East to West will be the same




# Agenda

- **OTV Introduction**
- **Configuration**
  - Multicast Transport
  - Unicast-only Transport
- **Verification**
  - Adjacency
  - Unicast Forwarding
  - Multicast Forwarding
  - ARP
- **Troubleshooting**

# Verification

## Multicast Forwarding: IGMP Join from Client

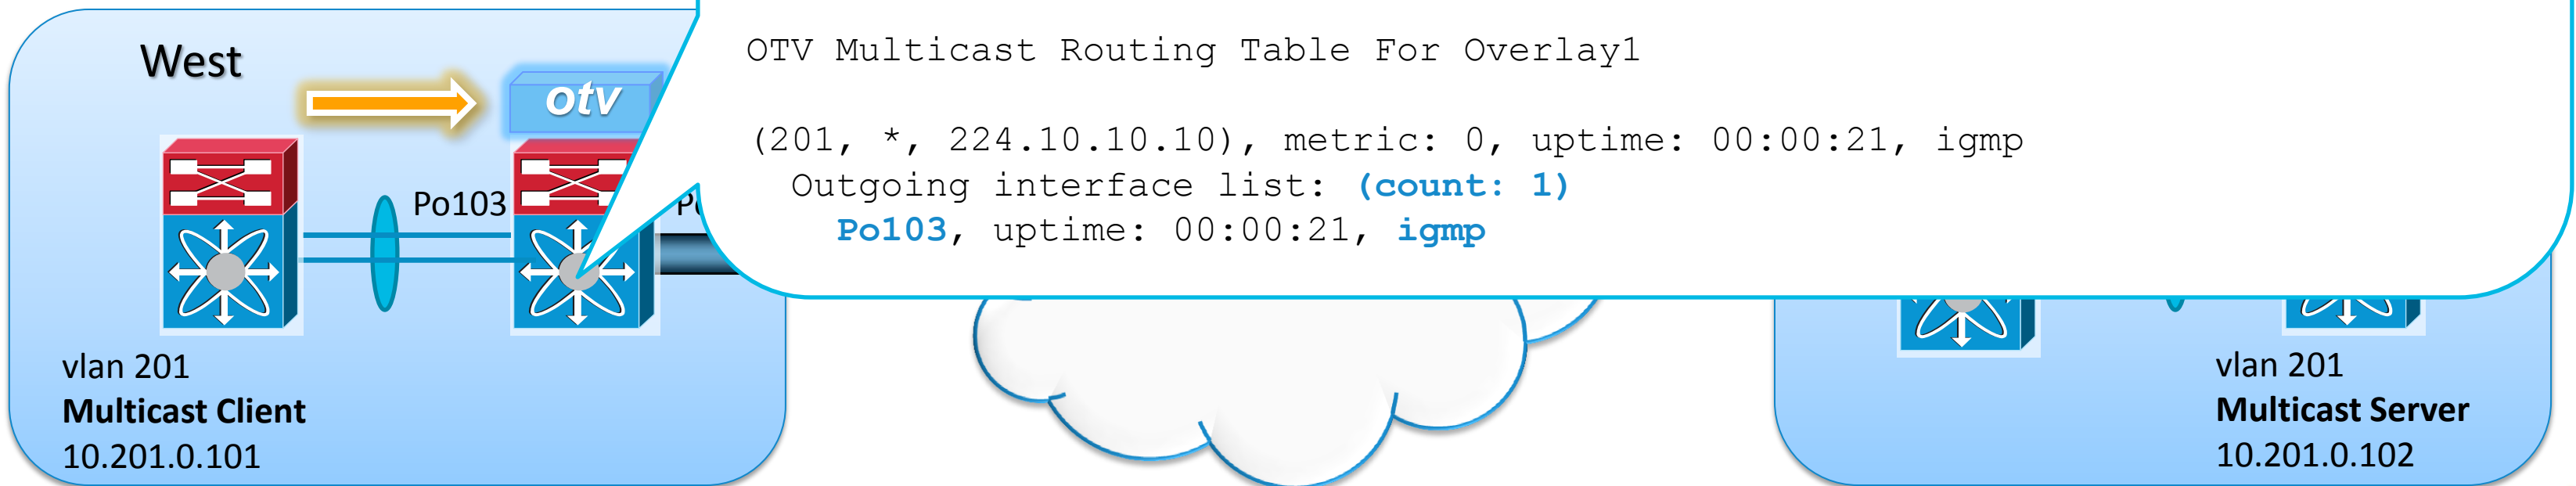
- Client at West site requests traffic for multicast group 224.10.10.10
- There is no server sending traffic for this group



# Verification

## Multicast Forwarding: IGMP Join from Client

- Multicast Transport



```
WEST_OTVA# show ip igmp snooping groups vlan 201
```

```
Type: S - Static, D - Dynamic, R - Router port, F - Fabricpath core port
```

| Vlan | Group Address | Ver | Type | Port list |
|------|---------------|-----|------|-----------|
| 201  | 224.10.10.10  | v2  | D    | Po103     |

```
WEST_OTVA# show otv mroute vlan 201 detail
```

```
OTV Multicast Routing Table For Overlay1
```

```
(201, *, 224.10.10.10), metric: 0, uptime: 00:00:21, igmp
```

```
Outgoing interface list: (count: 1)
```

```
Po103, uptime: 00:00:21, igmp
```



# Verification

## Multicast Forwarding: IGMP Join from Client

- Multicast Transport

(r) means there is a receiver that exists across the overlay

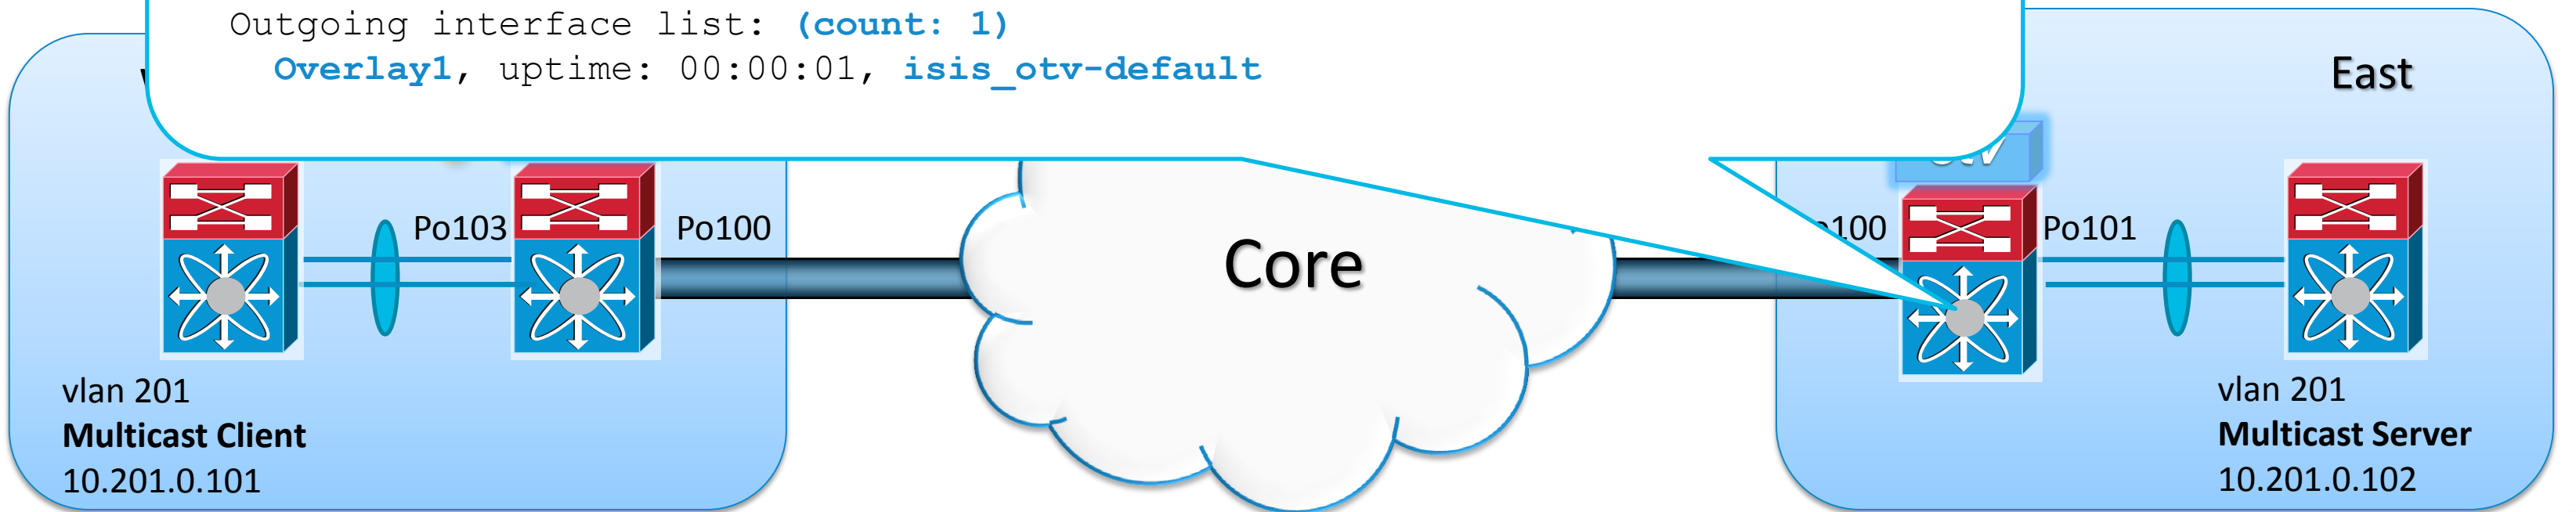
```
EAST_OTVA# show otv mroute vlan 201
```

```
OTV Multicast Routing Table For Overlay1
```

```
(201, *, 224.10.10.10), metric: 0, uptime: 00:00:01, overlay (r)
```

```
Outgoing interface list: (count: 1)
```

```
Overlay1, uptime: 00:00:01, isis_otv-default
```



# Verification

## Multicast Forwarding: IGMP Join from Client

- Unicast Transport

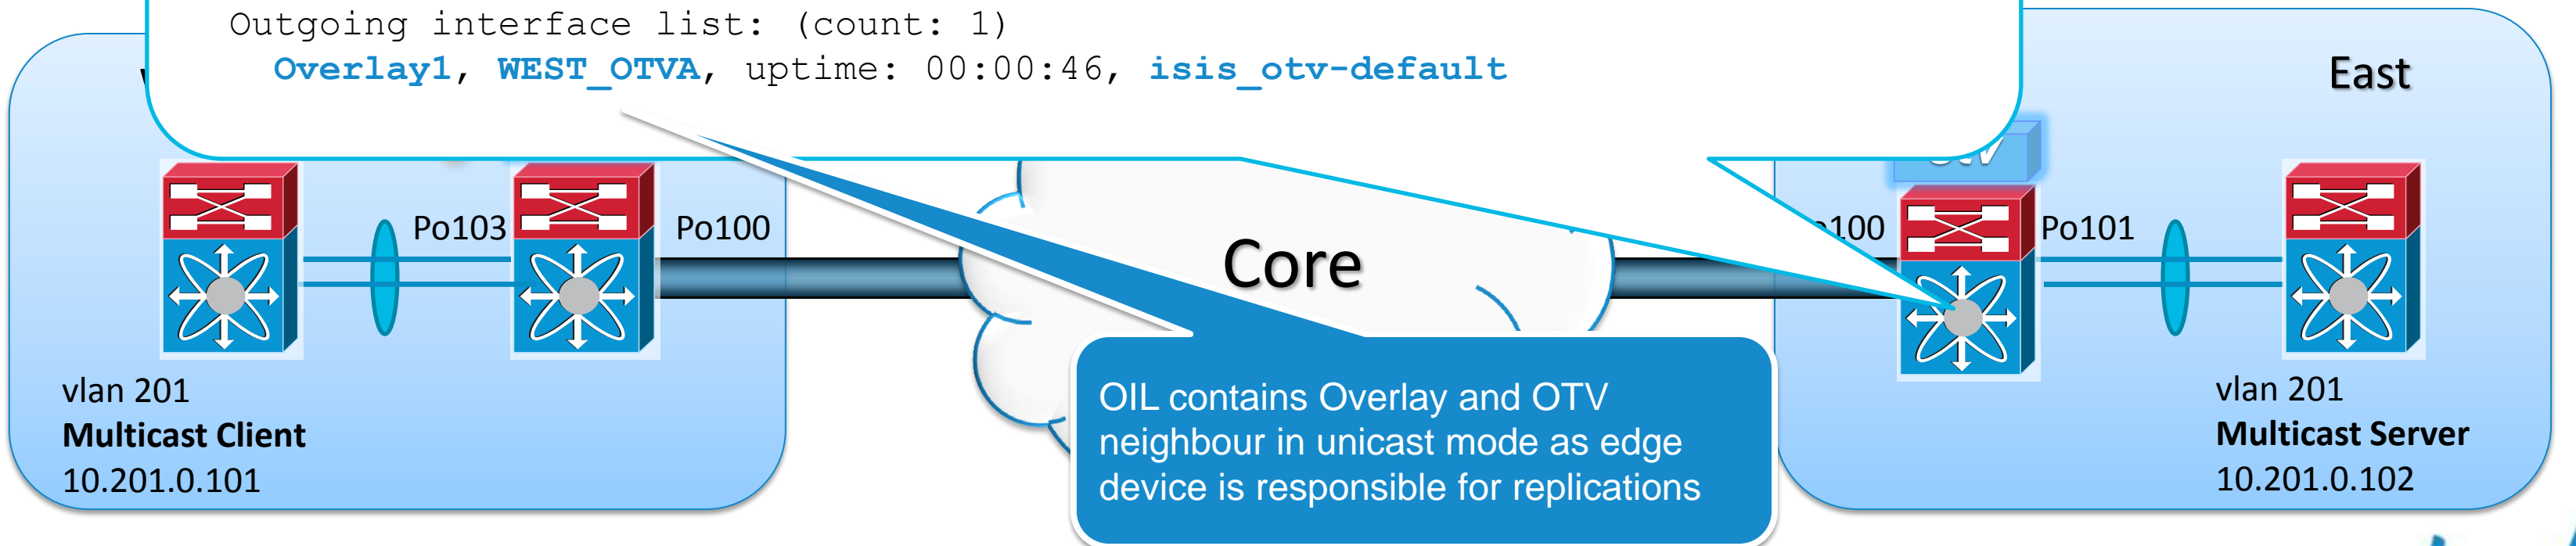
```
EAST_OTVA# show otv mroute vlan 201 detail
```

```
OTV Multicast Routing Table For Overlay1
```

```
(201, *, 224.10.10.10), metric: 0, uptime: 00:00:46, overlay(r)
```

```
Outgoing interface list: (count: 1)
```


```
Overlay1, WEST_OTVA, uptime: 00:00:46, isis_otv-default
```



# Verification

## Multicast Forwarding: Multicast Server Discovery

- Server at East site starts sending traffic to group 224.10.10.10
- There are no clients requesting the stream



# Verification

## Multicast Forwarding: Multicast Server Discovery

### ■ Multicast Transport

```
EAST_OTVA# show otv mroute vlan 201 detail
```

```
OTV Multicast Routing Table For Overlay1
```


```
(201, 10.201.0.102, 224.10.10.10), metric: 0, uptime: 00:00:30, site  
  Outgoing interface list: (count: 0)  
    Local Delivery: s = 172.16.1.26, g = 232.1.1.0
```

```
EAST_OTVA# show otv data-group vlan 201
```

```
Local Active Sources for Overlay1
```

| VLAN | Active-Source | Active-Group | Delivery-Source | Delivery-Group | Join-IF | State |
|------|---------------|--------------|-----------------|----------------|---------|-------|
| 201  | 10.201.0.102  | 224.10.10.10 | 172.16.1.26     | 232.1.1.0      | Po100   | Local |

```
vlan 201  
Multicast Client  
10.201.0.101
```



# Verification

## Multicast Forwarding: Multicast Server Discovery

### ■ Multicast Transport

(s) means there is a source that exists across the overlay

```
WEST_OTVA# show otv mroute vlan 201 detail
```

```
OTV Multicast Routing Table For Overlay1
```


```
(201, 10.201.0.102, 224.10.10.10), metric: 0, uptime: 00:00:53, overlay(s)  
Outgoing interface list: (count: 0)  
Remote Delivery: s = 172.16.1.26, g = 232.1.1.0
```

```
WEST_OTVA# show otv data-group vlan 201
```

```
Remote Active Sources for Overlay1
```

| VLAN | Active-Source | Active-Group | Delivery-Source | Delivery-Group | Joined-I/F |
|------|---------------|--------------|-----------------|----------------|------------|
| 201  | 10.201.0.102  | 224.10.10.10 | 172.16.1.26     | 232.1.1.0      | -          |

West



vlan 201  
Multicast Client  
10.201.0.101

otv

Po103



Multicast Server  
10.201.0.102

# Verification

## Multicast Forwarding: Multicast Server Discovery

### ■ Unicast Transport

```
EAST_OTVA# show otv mroute vlan 201 detail
```

```
OTV Multicast Routing Table For Overlay1
```

```
(201, 10.201.0.102, 224.10.10.10), metric: 0, uptime: 00:00:11, site  
  Outgoing interface list: (count: 0)  
  Local Delivery: s = 0.0.0.0, g = 0.0.0.0
```


```
EAST_OTVA# show otv data-group vlan 201
```

```
Local Active Sources for Overlay1
```

| VLAN | Active-Source | Active-Group | Delivery-Source | Delivery-Group | Join-IF | State |
|------|---------------|--------------|-----------------|----------------|---------|-------|
| 201  | 10.201.0.102  | 224.10.10.10 | 0.0.0.0         | 0.0.0.0        |         | Local |

```
Multicast Client  
10.201.0.101
```

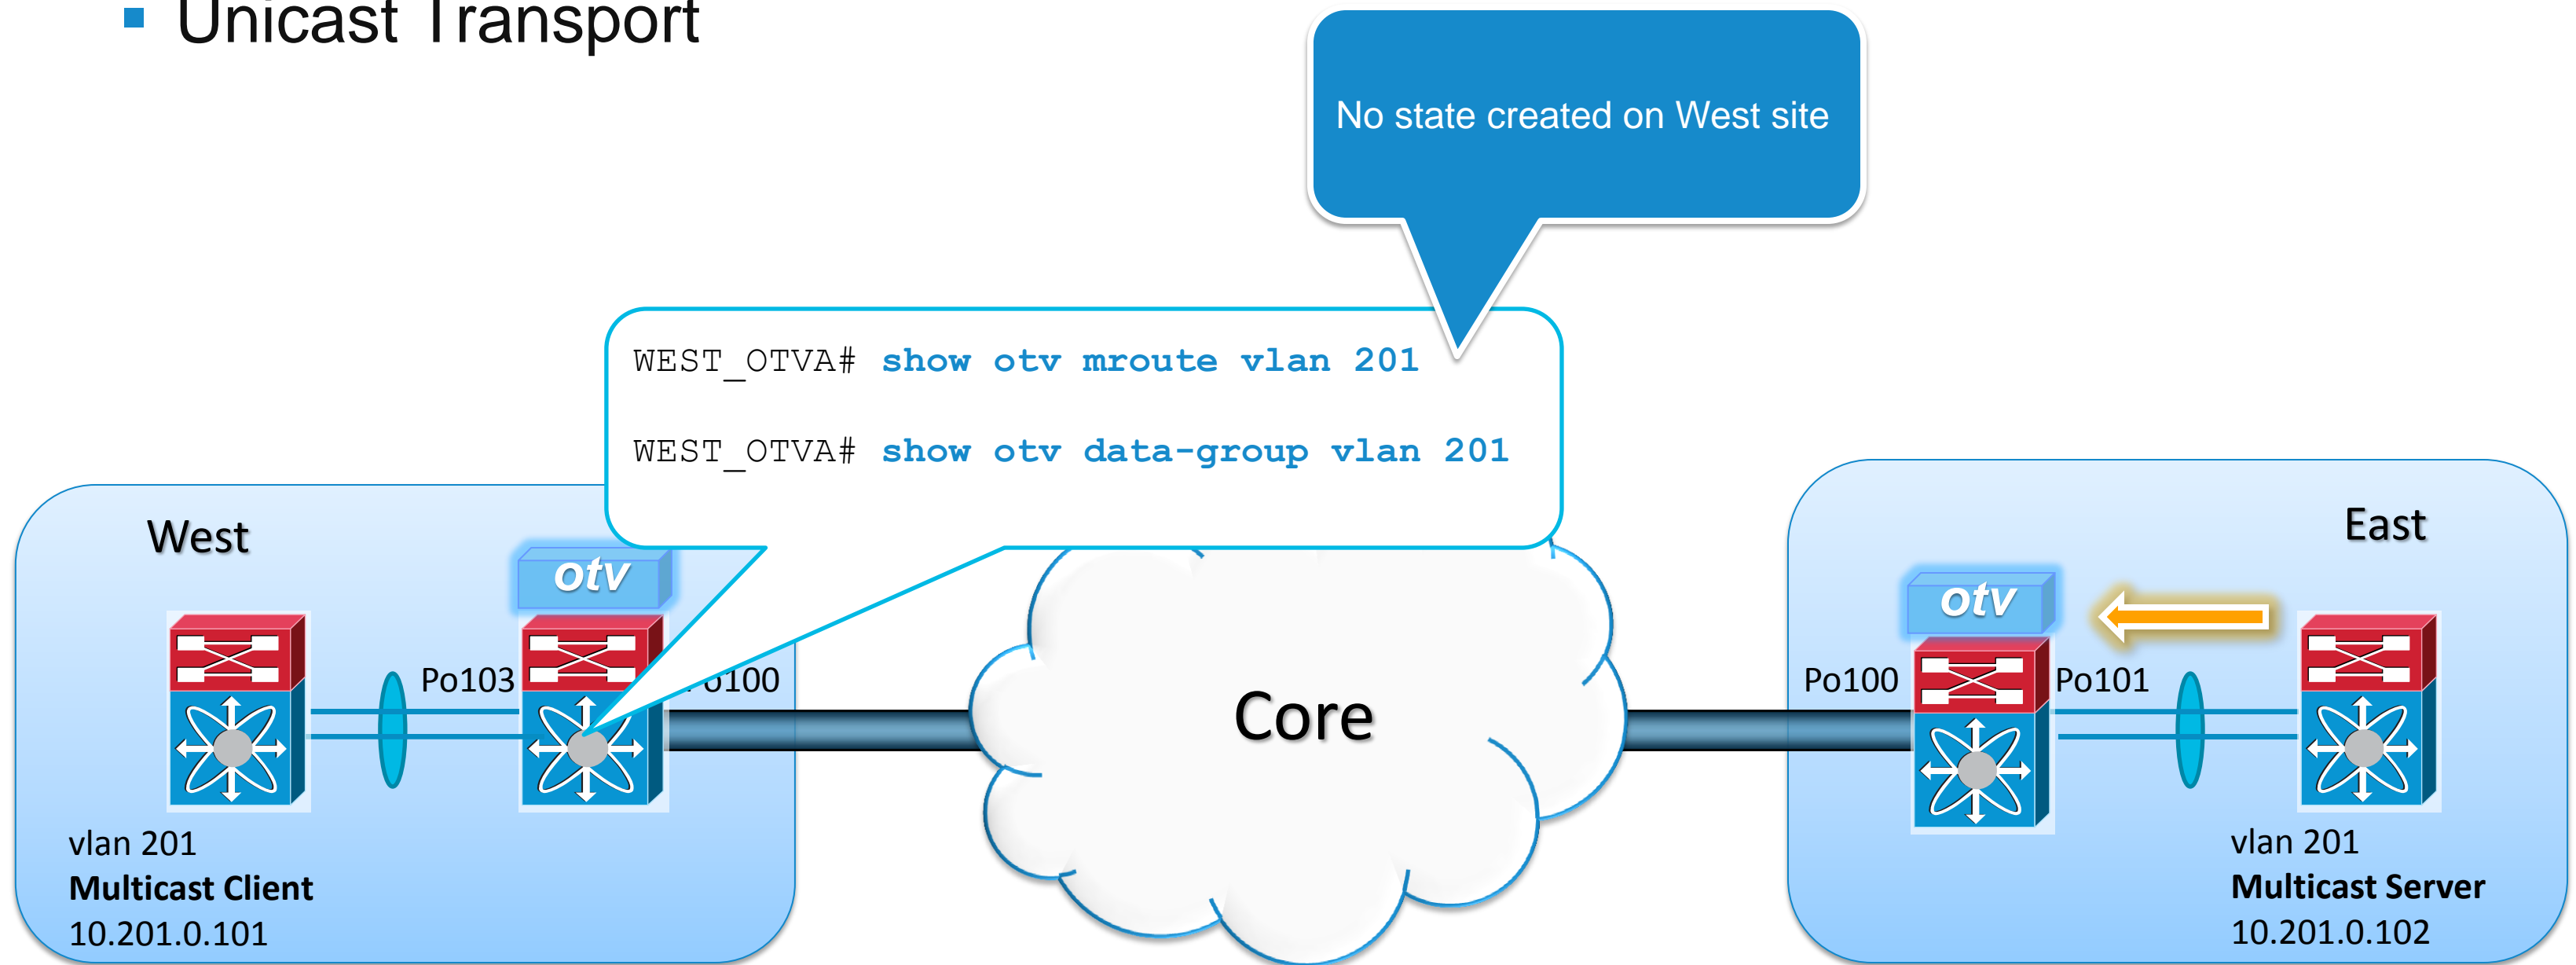
No delivery group in unicast mode



# Verification

## Multicast Forwarding: Multicast Server Discovery

- Unicast Transport



# Verification

## Multicast Forwarding: Source and Client Present

### ■ Multicast Transport

West Edge Device sends an IGMPv3 SSM join for the Delivery Source and the Delivery Group on its Join interface


```
WEST_OTVA# show otv mroute vlan 201 detail
OTV Multicast Routing Table For Overlay1
```

```
(201, *, 224.10.10.10), metric: 0, uptime: 00:01:50, igmp
  Outgoing interface list: (count: 1)
    Po103, uptime: 00:01:50, igmp
```

```
(201, 10.201.0.102, 224.10.10.10), metric: 0, uptime: 00:04:47, overlay(s
  Outgoing interface list: (count: 0)
  Remote Delivery: s = 172.16.1.26, g = 232.1.1.0
```

```
WEST_OTVA# show otv data-group
Remote Active Sources for Overlay1
```

| VLAN | Active-Source | Active-Group | Delivery-Source | Delivery-Group | Joined-I/F |
|------|---------------|--------------|-----------------|----------------|------------|
| 201  | 10.201.0.102  | 224.10.10.10 | 172.16.1.26     | 232.1.1.0      | Po100      |





# Verification

## Multicast Forwarding: Source and Client Present

### ■ Multicast Transport

```
EAST_OTVA# show otv mroute vlan 201 detail
```

```
OTV Multicast Routing Table For Overlay1
```

```
(201, *, 224.10.10.10), metric: 0, uptime: 00:01:42, overlay(r)
```

```
  Outgoing interface list: (count: 1)
```

```
    Overlay1, uptime: 00:01:42, isis_otv-default
```

```
(201, 10.201.0.102, 224.10.10.10), metric: 0, uptime: 00:04:42, site
```

```
  Outgoing interface list: (count: 1)
```


```
    Overlay1, uptime: 00:01:42, otv
```

```
      Local Delivery: s = 172.16.1.26, g = 232.1.1.0
```

```
EAST_OTVA# show otv data-group
```

```
Local Active Sources for Overlay1
```

| VLAN | Active-Source | Active-Group | Delivery-Source | Delivery-Group | Join-IF | State |
|------|---------------|--------------|-----------------|----------------|---------|-------|
| 201  | 10.201.0.102  | 224.10.10.10 | 172.16.1.26     | 232.1.1.0      | Po100   | Local |



# Verification

## Multicast Forwarding: Encapsulated Packet, Multicast Transport

| No. . | size | Time            | Source       | Destination  | dscp | Protocol | Info                               |
|-------|------|-----------------|--------------|--------------|------|----------|------------------------------------|
| 10    | 102  | 08:56:27.309463 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 11    | 102  | 08:56:27.329463 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 12    | 102  | 08:56:27.349463 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 13    | 102  | 08:56:27.369463 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 14    | 102  | 08:56:27.389463 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 15    | 102  | 08:56:27.409463 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |

Frame 10 (102 bytes on wire, 102 bytes captured)

- Ethernet II, Src: 6c:9c:ed:40:17:43 (6c:9c:ed:40:17:43), Dst: IPv4mcast\_01:01:00 (01:00:5e:00:01:00)
- Internet Protocol, Src: 172.16.1.26 (172.16.1.26), Dst: 232.1.1.0 (232.1.1.0)
  - Version: 4
  - Header length: 20 bytes
  - Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
  - Total Length: 88
  - Identification: 0x3bde (15326)
  - Flags: 0x02 (Don't Fragment)
  - Fragment offset: 0
  - Time to live: 254
  - Protocol: GRE (0x2f)
  - Header checksum: 0xaa6c [correct]
  - Source: 172.16.1.26 (172.16.1.26)
  - Destination: 232.1.1.0 (232.1.1.0)
- Generic Routing Encapsulation (0x8848 - unknown)
- MultiProtocol Label Switching Header, Label: 233, Exp: 0, S: 1, TTL: 255
- Ethernet II, Src: Cisco\_75:1d:42 (00:11:6c:75:1d:42), Dst: IPv4mcast\_0a:0a:0a (01:00:5e:0a:0a:0a)
- Internet Protocol, Src: 10.201.0.102 (10.201.0.102), Dst: 224.10.10.10 (224.10.10.10)
  - Version: 4
  - Header length: 20 bytes
  - Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
  - Total Length: 46
  - Identification: 0x0000 (0)
  - Flags: 0x00
  - Fragment offset: 0
  - Time to live: 64
  - Protocol: UDP (0x11)
  - Header checksum: 0x857c [correct]
  - Source: 10.201.0.102 (10.201.0.102)
  - Destination: 224.10.10.10 (224.10.10.10)
- User Datagram Protocol, Src Port: complex-main (5000), Dst Port: complex-link (5001)
- Data (18 bytes)

Sourced from Join interface, and destined to first address in data-group

Label to identify vlan  $32+201 = 233$

Original IP header maintained



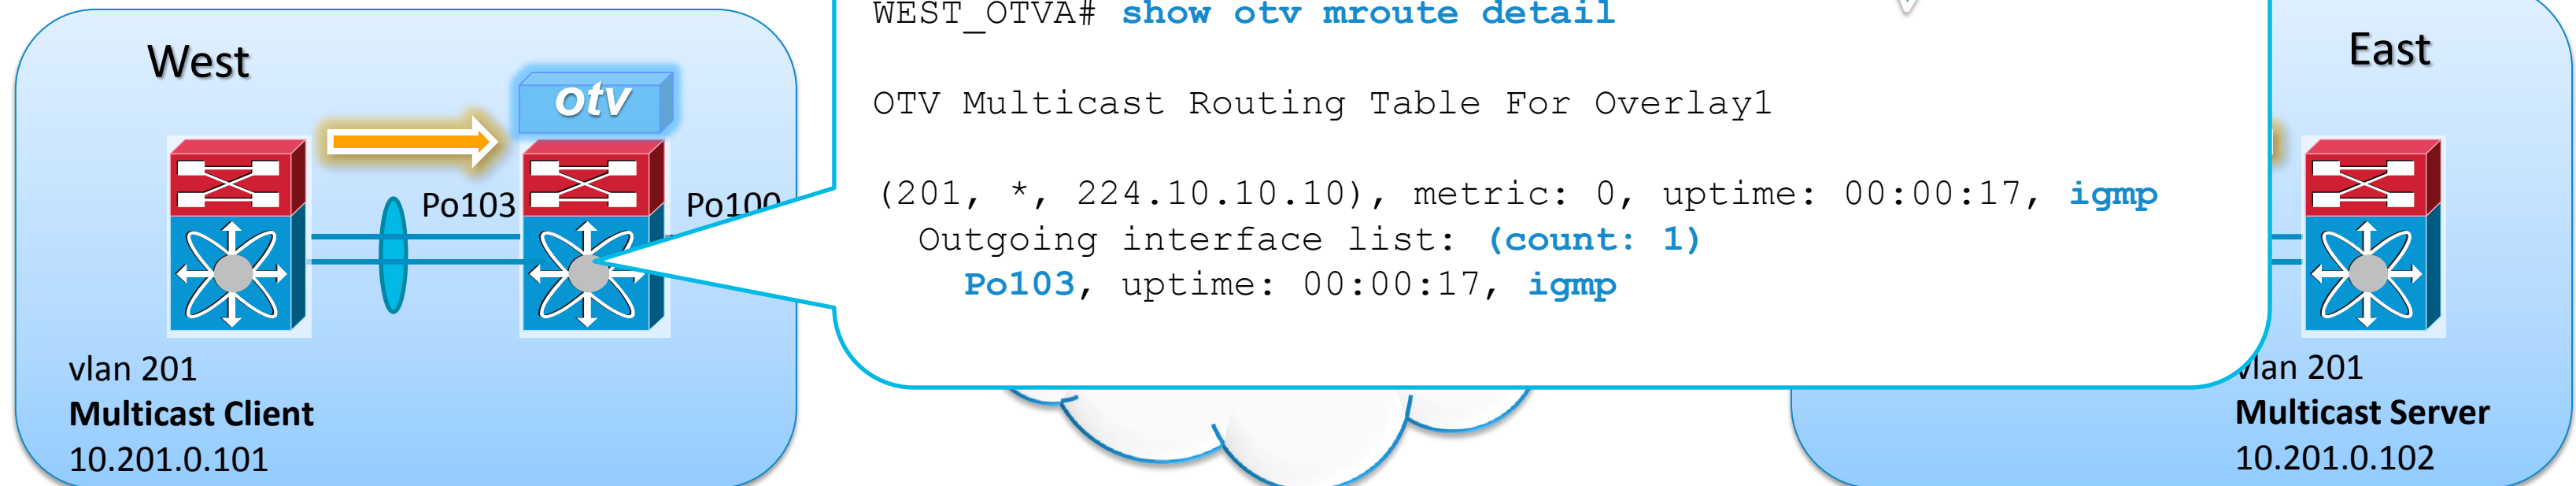
# Verification

## Multicast Forwarding: Source and Client Present

- Unicast Transport

Since core does not support multicast, West site cannot send SSM join for group. Instead, West needs only to communicate to East that it has a receiver and it will receive the group via unicast.

Only the \*,G is created in unicast mode on client site.



# Verification

## Multicast Forwarding: Source and Client Present

### ■ Unicast Transport


```
EAST_OTVA# show otv mroute vlan 201
OTV Multicast Routing Table For Overlay1

(201, *, 224.10.10.10), metric: 0, uptime: 00:00:46, overlay(r)
  Outgoing interface list: (count: 1)
    Overlay1, WEST_OTVA, uptime: 00:00:46, isis_otv-default

(201, 10.201.0.102, 224.10.10.10), metric: 0, uptime: 00:01:54, site
  Outgoing interface list: (count: 1)
    Overlay1, WEST_OTVA, uptime: 00:00:46, otv

EAST_OTVA# show forwarding otv multicast route vlan 201
! Some output omitted
(10.201.0.102/32, 224.10.10.10/32), RPF Interface: NULL, flags:
  Received Packets: 59188 Bytes: 3788032
  Number of Outgoing Interfaces: 1
  Outgoing Interface List Index: 25
  Tunnel16407 Outgoing Packets:41096 Bytes:3369872
  OTV unicast tunnel end-points: (172.16.1.26, 172.16.1.34)
  vlan: 201
```

Each multicast group is sent via unicast to each site with receivers present



# Verification

## Multicast Forwarding: Encapsulated Packet, Unicast Transport

| No. . | size | Time            | Source       | Destination  | dscp | Protocol | Info                               |
|-------|------|-----------------|--------------|--------------|------|----------|------------------------------------|
| 1     | 102  | 10:32:25.345410 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 2     | 102  | 10:32:25.365410 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 3     | 102  | 10:32:25.385410 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 4     | 102  | 10:32:25.405410 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 5     | 102  | 10:32:25.425410 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |
| 6     | 102  | 10:32:25.445410 | 10.201.0.102 | 224.10.10.10 | 0    | UDP      | Source port: complex-main Destinat |

Frame 1 (102 bytes on wire, 102 bytes captured)

- Ethernet II, Src: 6c:9c:ed:40:17:43 (6c:9c:ed:40:17:43), Dst: Cisco\_be:52:c1 (00:1a:e2:be:52:c1)
- Internet Protocol, Src: 172.16.1.26 (172.16.1.26), Dst: 172.16.1.34 (172.16.1.34)
  - Version: 4
  - Header length: 20 bytes
  - Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
  - Total Length: 88
  - Identification: 0x38a3 (14499)
  - Flags: 0x02 (Don't Fragment)
  - Fragment offset: 0
  - Time to live: 253
  - Protocol: GRE (0x2f)
  - Header checksum: 0xea76 [correct]
    - Source: 172.16.1.26 (172.16.1.26)
    - Destination: 172.16.1.34 (172.16.1.34)
- Generic Routing Encapsulation (0x8848 - unknown)
- MultiProtocol Label Switching Header, Label: 233, Exp: 0, S: 1, TTL: 254
- Ethernet II, Src: Cisco\_75:1d:42 (00:11:6c:75:1d:42), Dst: IPv4mcast\_0a:0a:0a (01:00:5e:0a:0a:0a)
- Internet Protocol, Src: 10.201.0.102 (10.201.0.102), Dst: 224.10.10.10 (224.10.10.10)
  - Version: 4
  - Header length: 20 bytes
  - Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
  - Total Length: 46
  - Identification: 0x0000 (0)
  - Flags: 0x00
  - Fragment offset: 0
  - Time to live: 64
  - Protocol: UDP (0x11)
  - Header checksum: 0x857c [correct]
    - Source: 10.201.0.102 (10.201.0.102)
    - Destination: 224.10.10.10 (224.10.10.10)
- User Datagram Protocol, Src Port: complex-main (5000), Dst Port: complex-link (5001)
- Data (18 bytes)

Source and Destination IP between Join Interfaces

Label to identify vlan  
 $32+201 = 233$

Original IP header maintained


# Agenda

- **OTV Introduction**
- **Configuration**
  - Multicast Transport
  - Unicast-only Transport
- **Verification**
  - Adjacency
  - Unicast Forwarding
  - Multicast Forwarding
  - ARP
- **Troubleshooting**

# Verification

## Address Resolution Protocol (ARP)

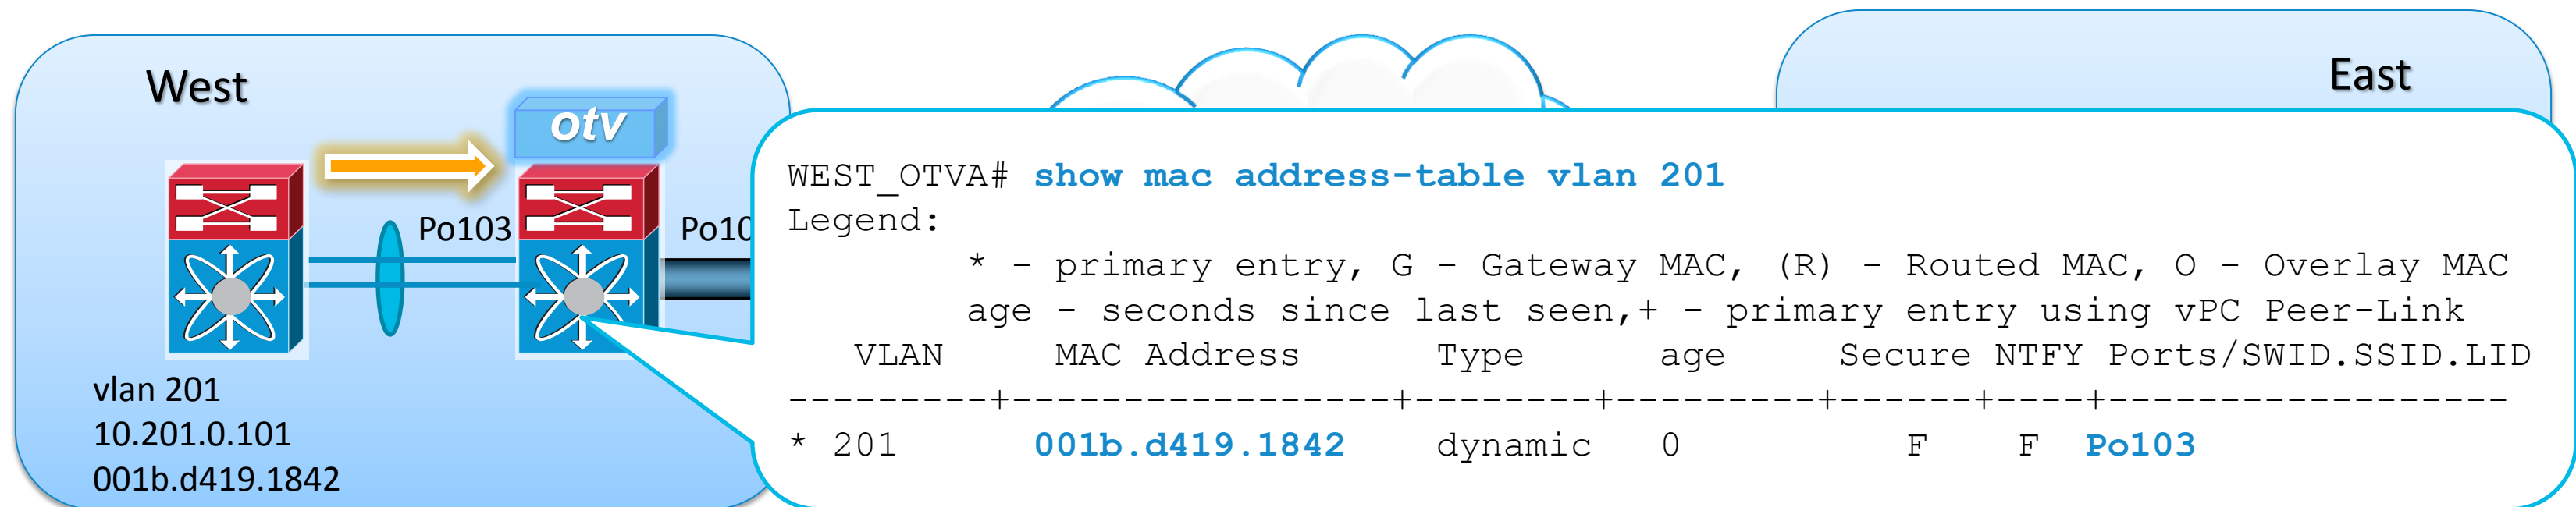
- We will assume that none of the devices in the topology have ARP or CAM entries for the hosts in vlan 201



# Verification

## ARP: Host at West Site Sends ARP Request for Host at East

1. Since it's a broadcast packet, it is forwarded to both the OTV devices at West site
2. Non AED at West site drops the broadcast packet (loop prevention)
3. AED learns the MAC address on its internal interface

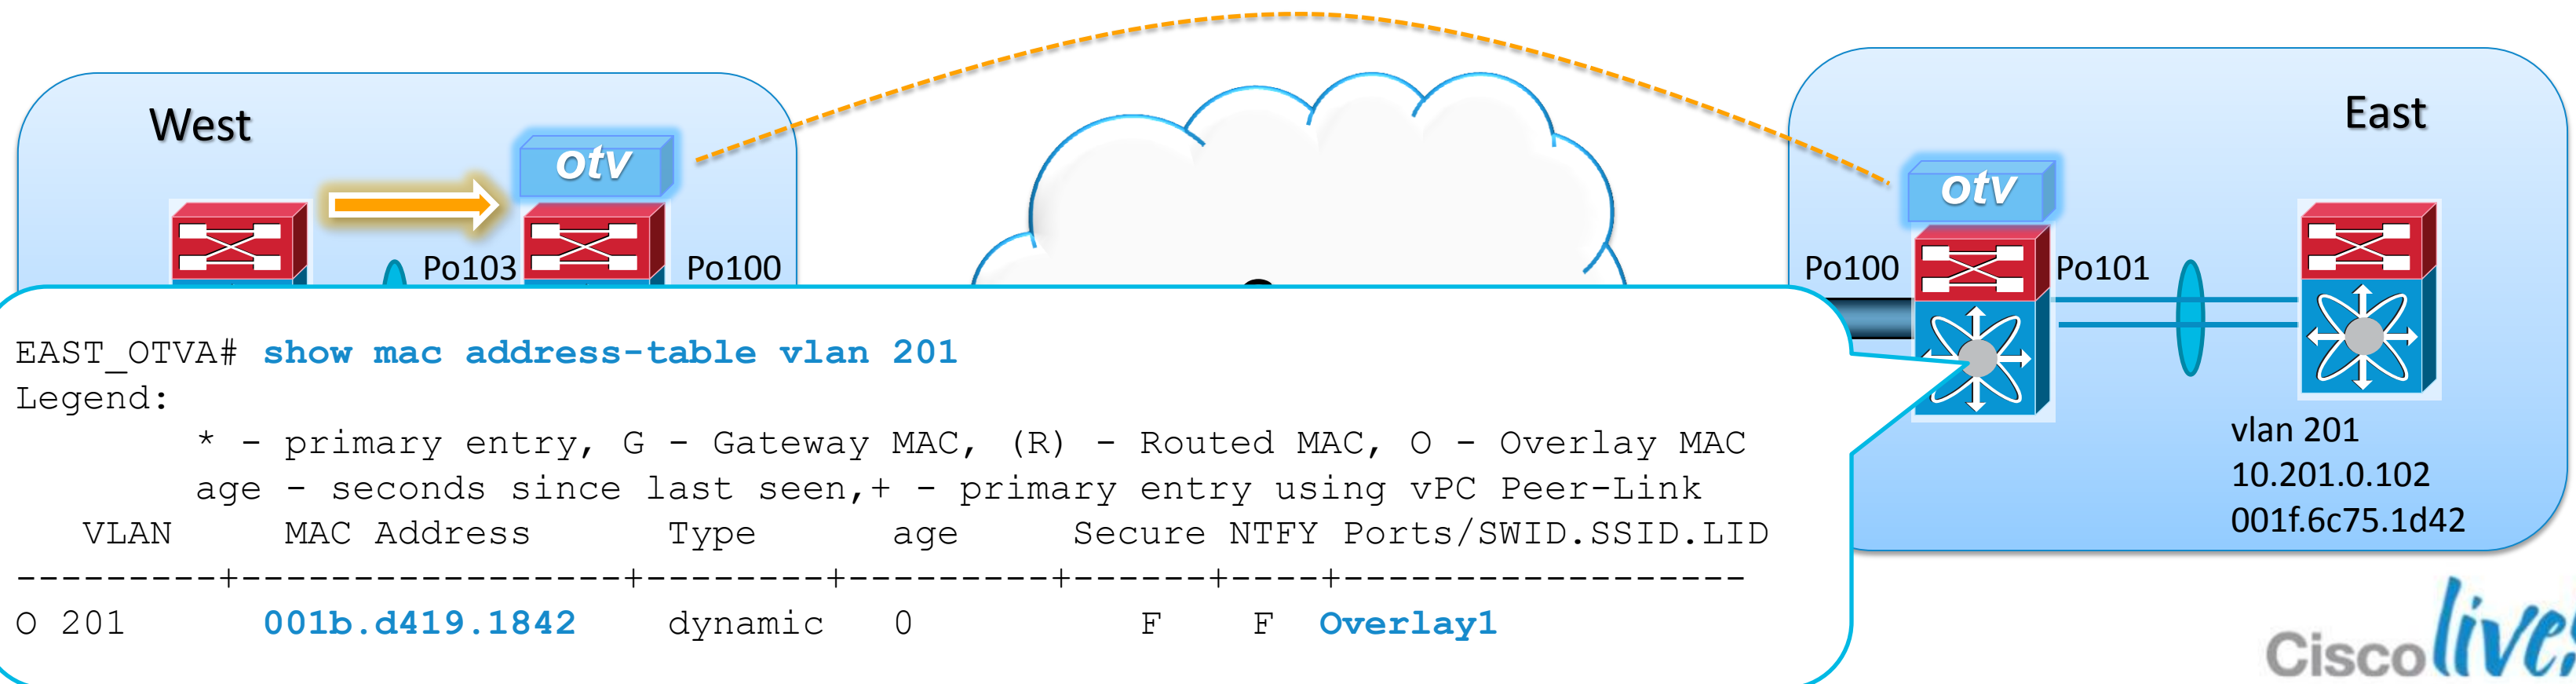




# Verification

## ARP: Host at West Site Sends ARP Request for Host at East


4. On learning new MAC, West AED sends ISIS update to all OTV devices
  - Single packet on multicast control group (Multicast Transport)
  - Or, unicast to each adjacency (Unicast Transport)
5. Only AED at remote sites program new MAC into OTV route and CAM tables



# Verification

## ARP: Host at West Site Sends ARP Request for Host at East


6. West AED performs lookup in ARP-ND (ARP- IPv6 Neighbour Discover) cache for East Host IP
  - If an entry were present, West could send ARP reply (proxy) to local host without forwarding packet across overlay



# Verification

## ARP: Host at West Site Sends ARP Request for Host at East


7. Since there is no entry present in cache, West encapsulates ARP broadcast and sends to all OTV devices
- Single packet on multicast control group (Multicast Transport)
  - Or, unicast to each adjacency (Unicast Transport)



# Verification

## ARP: Host at West Site Sends ARP Request for Host at East

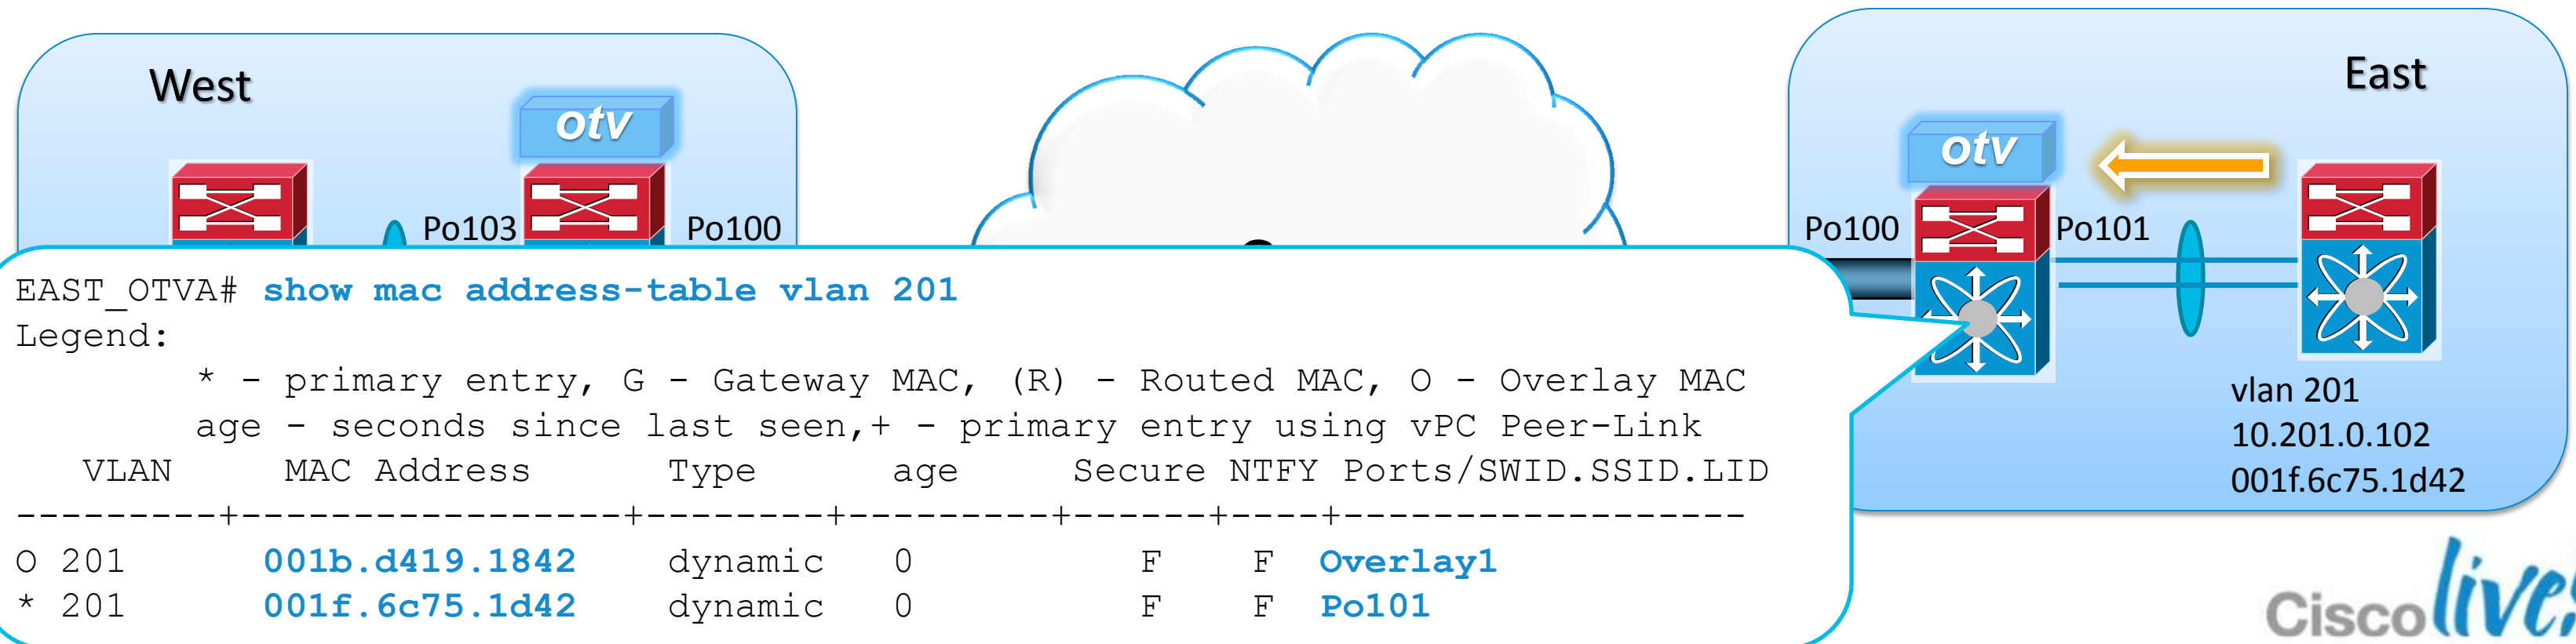
8. AED at East site receives packet on Join interface, decapsulates and sends it on internal interface toward host
  - Non AED at East will also receive packet but will not forward



# Verification

## ARP: Host at East Site Sends ARP Reply for Host at West

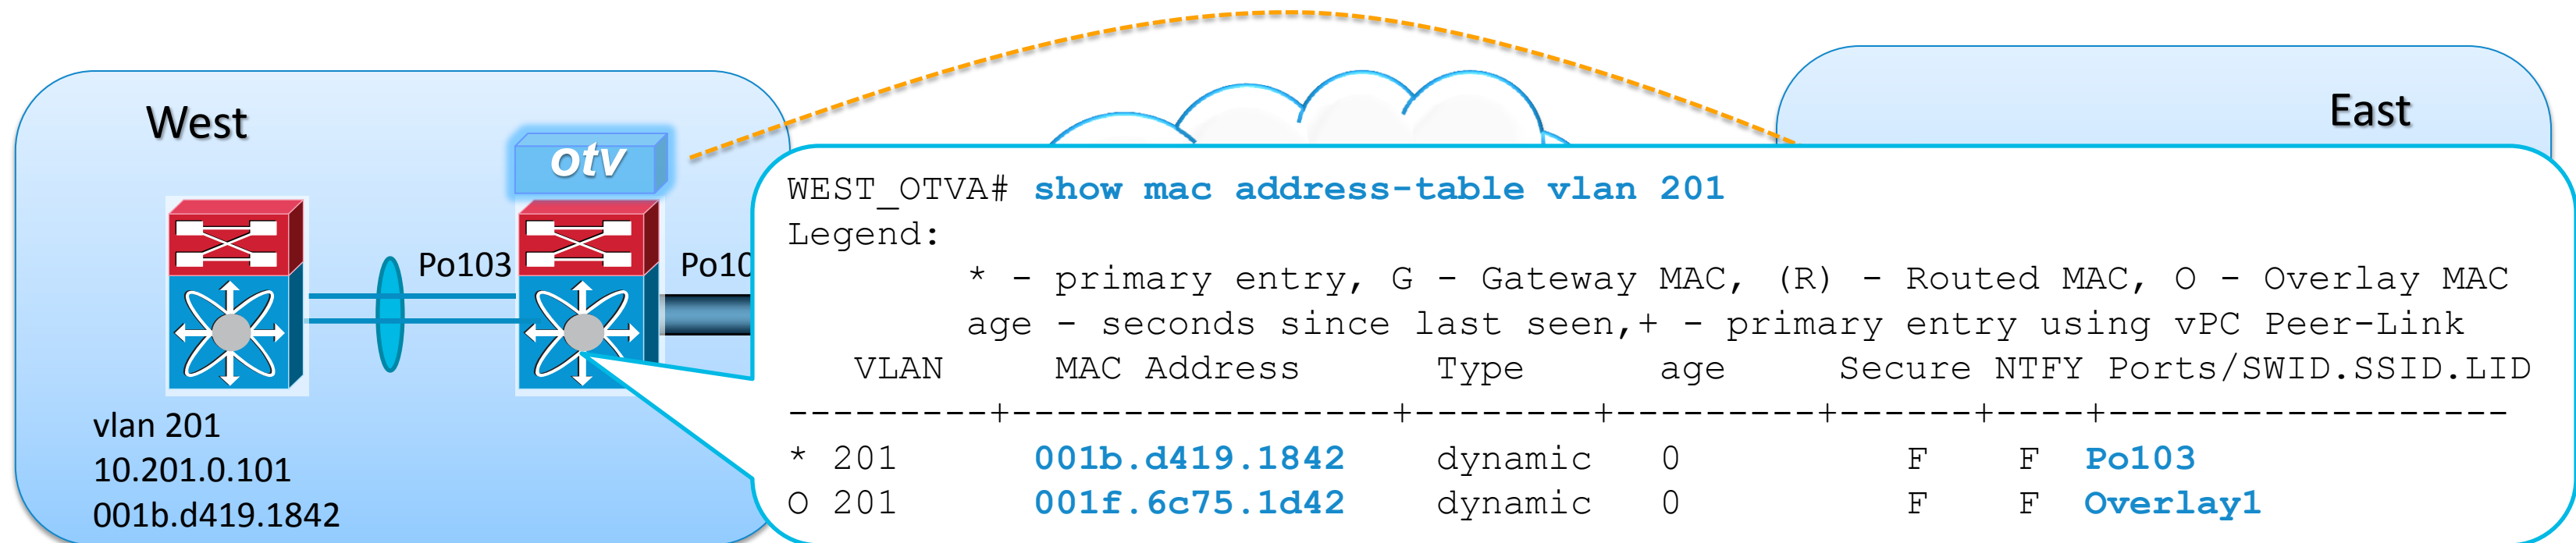
1. AED at East receives unicast Reply on its internal interface
2. East updates its CAM table with the MAC address pointing out its internal interface



# Verification

## ARP: Host at East Site Sends ARP Reply for Host at West


3. On learning new MAC, East sends ISIS update to all OTV devices
  - Single packet on multicast control group (Multicast Transport)
  - Or, unicast to each adjacency (Unicast Transport)
4. Only AED at remote sites program new MAC into OTV route and CAM tables



# Verification

## ARP: Host at East Site Sends ARP Reply for Host at West

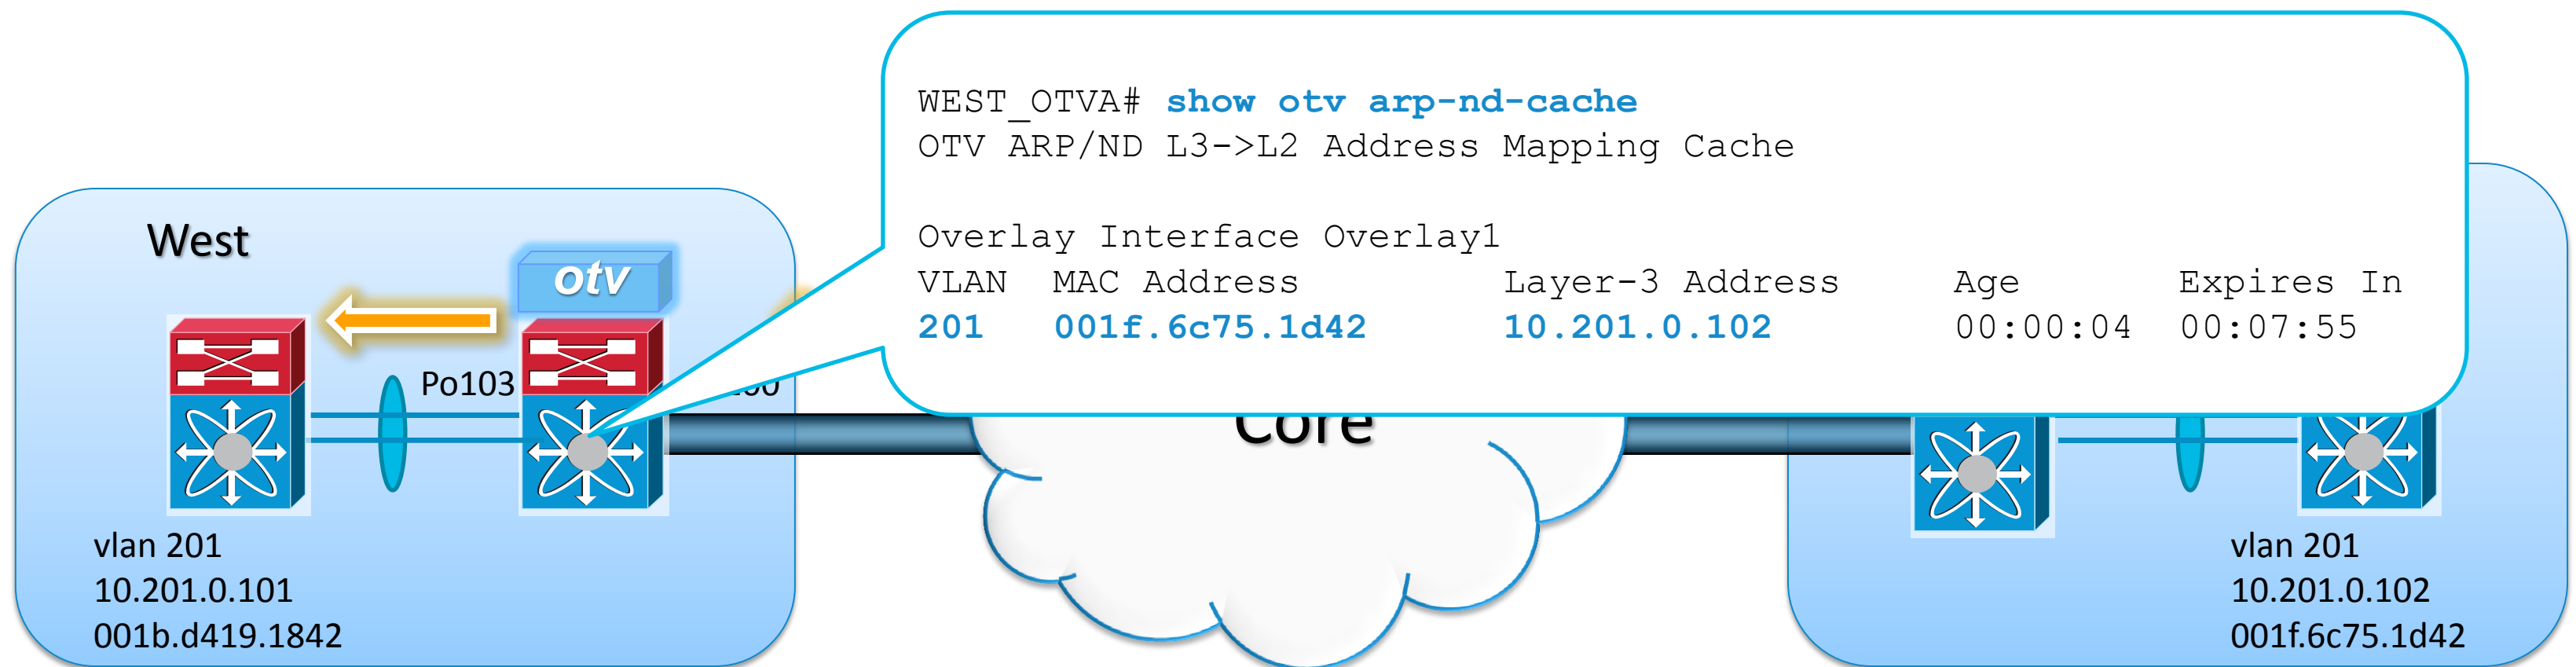
5. East performs lookup in its CAM for unicast destination. Because of previous ISIS update, East finds an entry pointing out overlay toward West
6. East encapsulates ARP reply and sends via unicast to West



# Verification

## ARP: Host at East Site Sends ARP Reply for Host at West

7. West receives packet on Join Interface, decapsulates packet and sends out internal interface toward host
8. West updates ARP-ND cache for East Host from ARP reply





# Agenda

- **OTV Introduction**
- **Configuration**
  - Multicast Transport
  - Unicast-only Transport
- **Verification**
  - Adjacency
  - Unicast Forwarding
  - Multicast Forwarding
  - ARP
- **Troubleshooting**

# Troubleshooting

## MTU

- Verify appropriate MTU via ping between OTV join interfaces
  - ‘packet-size’ in NxOS represents size of data in ICMP packet
  - To test MTU, must account for 8 Byte ICMP header, 20 Byte IP header
- Example:
  1.  $1442 = 1414 + 20 + 8$ , use packet-size of 1414
  2.  $1450 = 1422 + 20 + 8$ , use packet-size of 1422
  3.  $1542 = 1514 + 20 + 8$ , use packet-size of 1514

```
WEST_OTVA# ! Verify transport supports MTU of 1542
WEST_OTVA# ping 172.16.1.26 packet-size 1514 df-bit
PING 172.16.1.26 (172.16.1.26): 1514 data bytes
1522 bytes from 172.16.1.26: icmp_seq=0 ttl=251 time=2.333 ms
```



# Troubleshooting

## Partial Adjacency

- (!) Flag implies a mismatch site-id due
  1. Receiving same site-id across overlay without site adjacency
  2. Receiving different site-id across site adjacency

```
WEST_OTVA# show otv site
```

```
Dual Adjacency State Description
```

```
Full      - Both site and overlay adjacency up
```

```
Partial   - Either site/overlay adjacency down
```

```
Down      - Both adjacencies are down (Neighbor is down/unreachable)
```

```
(!)       - Site-ID mismatch detected
```

*! Some output omitted*

| Hostname  | System-ID      | Adjacency-<br>State | Adjacency-<br>Uptime | AED-<br>Capable |
|-----------|----------------|---------------------|----------------------|-----------------|
| -----     | -----          | -----               | -----                | -----           |
| WEST_OTVB | 64a0.e741.c842 | <b>Partial (!)</b>  | 00:15:16             | No              |

# Troubleshooting


## Partial Adjacency

- Partial Adjacency implies either the site or overlay adjacency is down.
- Both are required to maintain a full adjacency
- **Most common reason for 'down' overlay adjacency is PIM misconfiguration in transport (multicast transport) or insufficient MTU**
- Ensure
  1. Matching site-id configuration at each site
  2. Sufficient MTU through transport
  3. IP connectivity between each edge device
  4. Correct PIM configuration through core for multicast transport

# Troubleshooting

## ARP and CAM timer issue

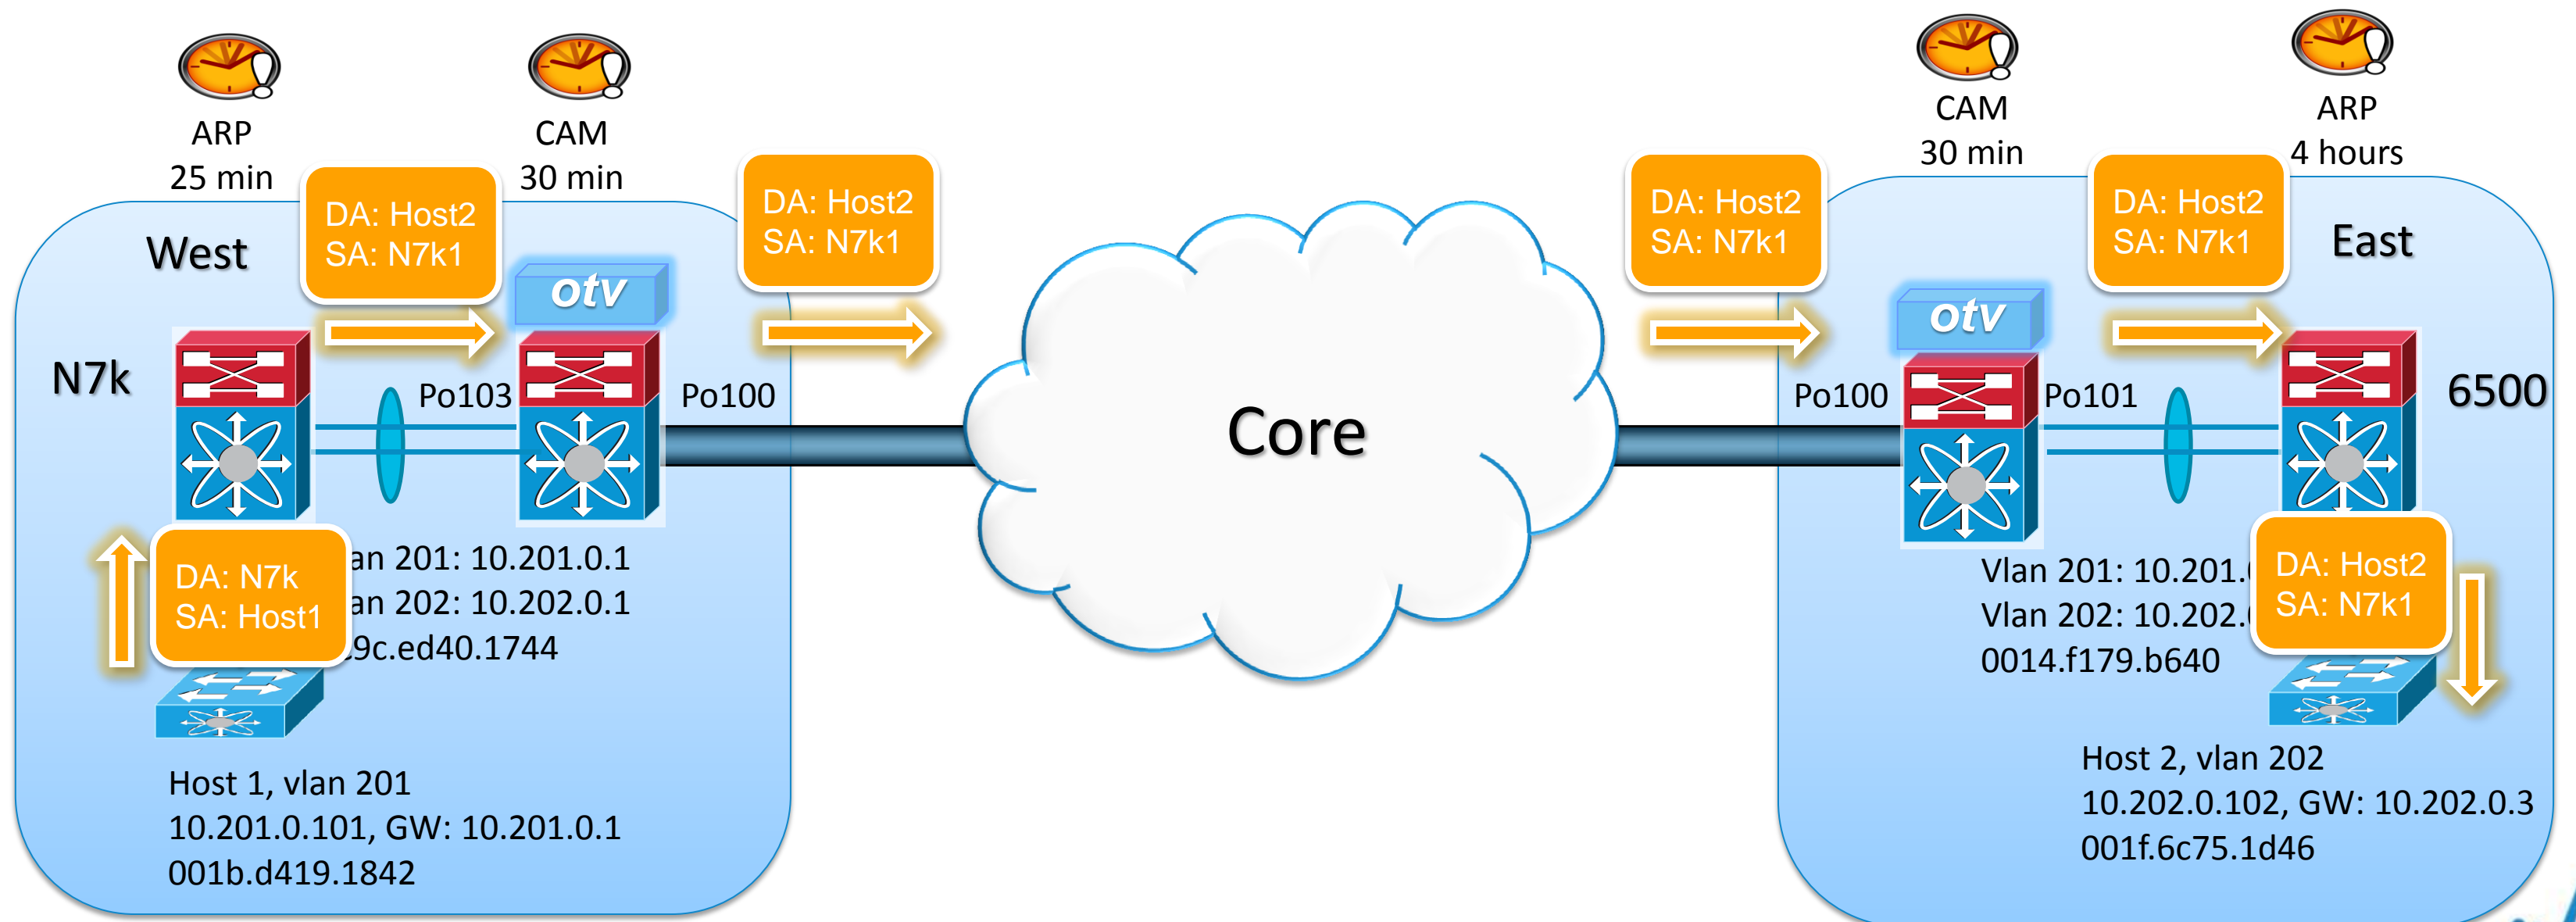
- Asymmetrical routing with mis-match ARP timers can cause traffic to black-hole across OTV



# Troubleshooting

## ARP and CAM timer issue

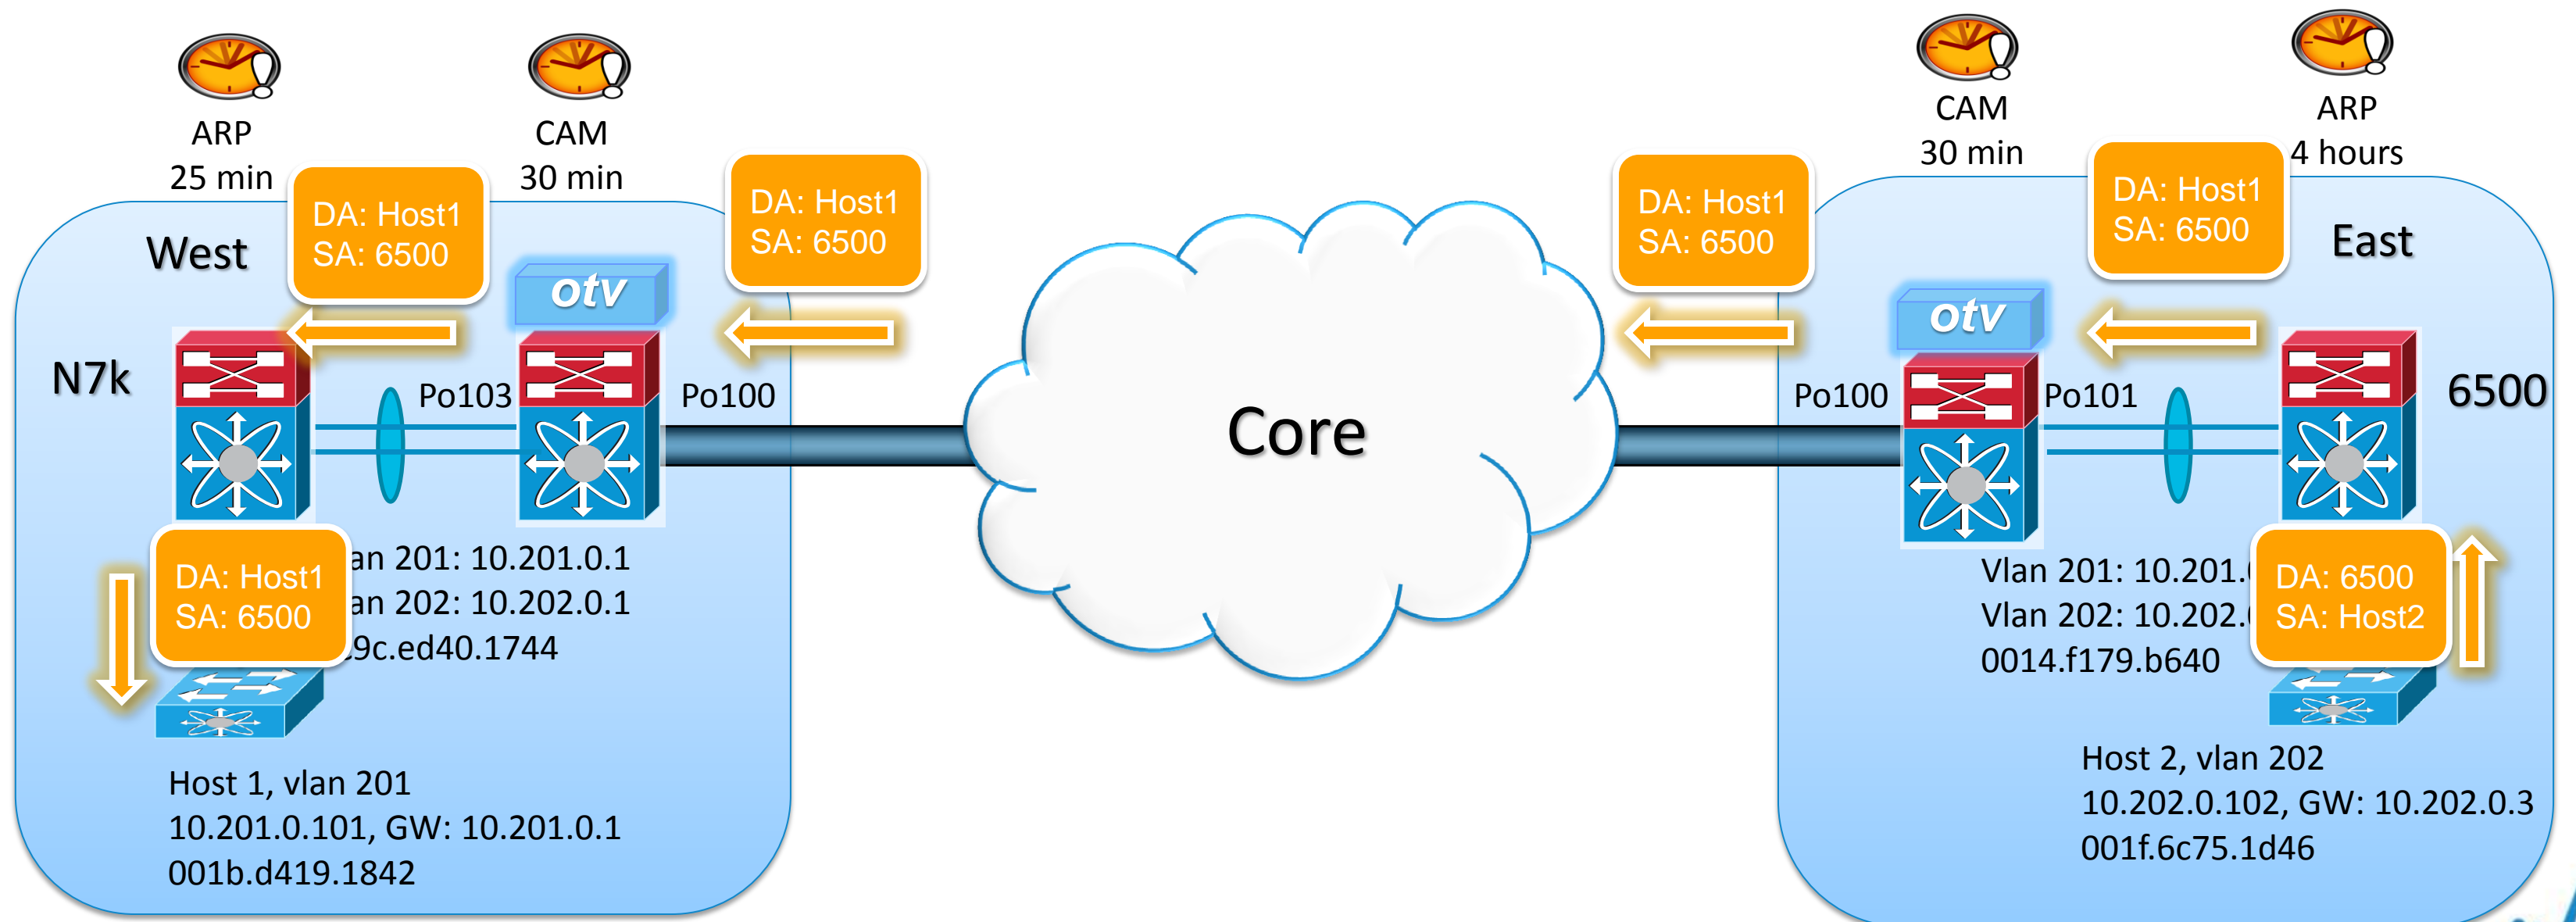
- Since the traffic flow between Host1 and Host2 is routed traffic, OTV will only see source MAC of the gateways and destination of the Hosts



# Troubleshooting

## ARP and CAM timer issue

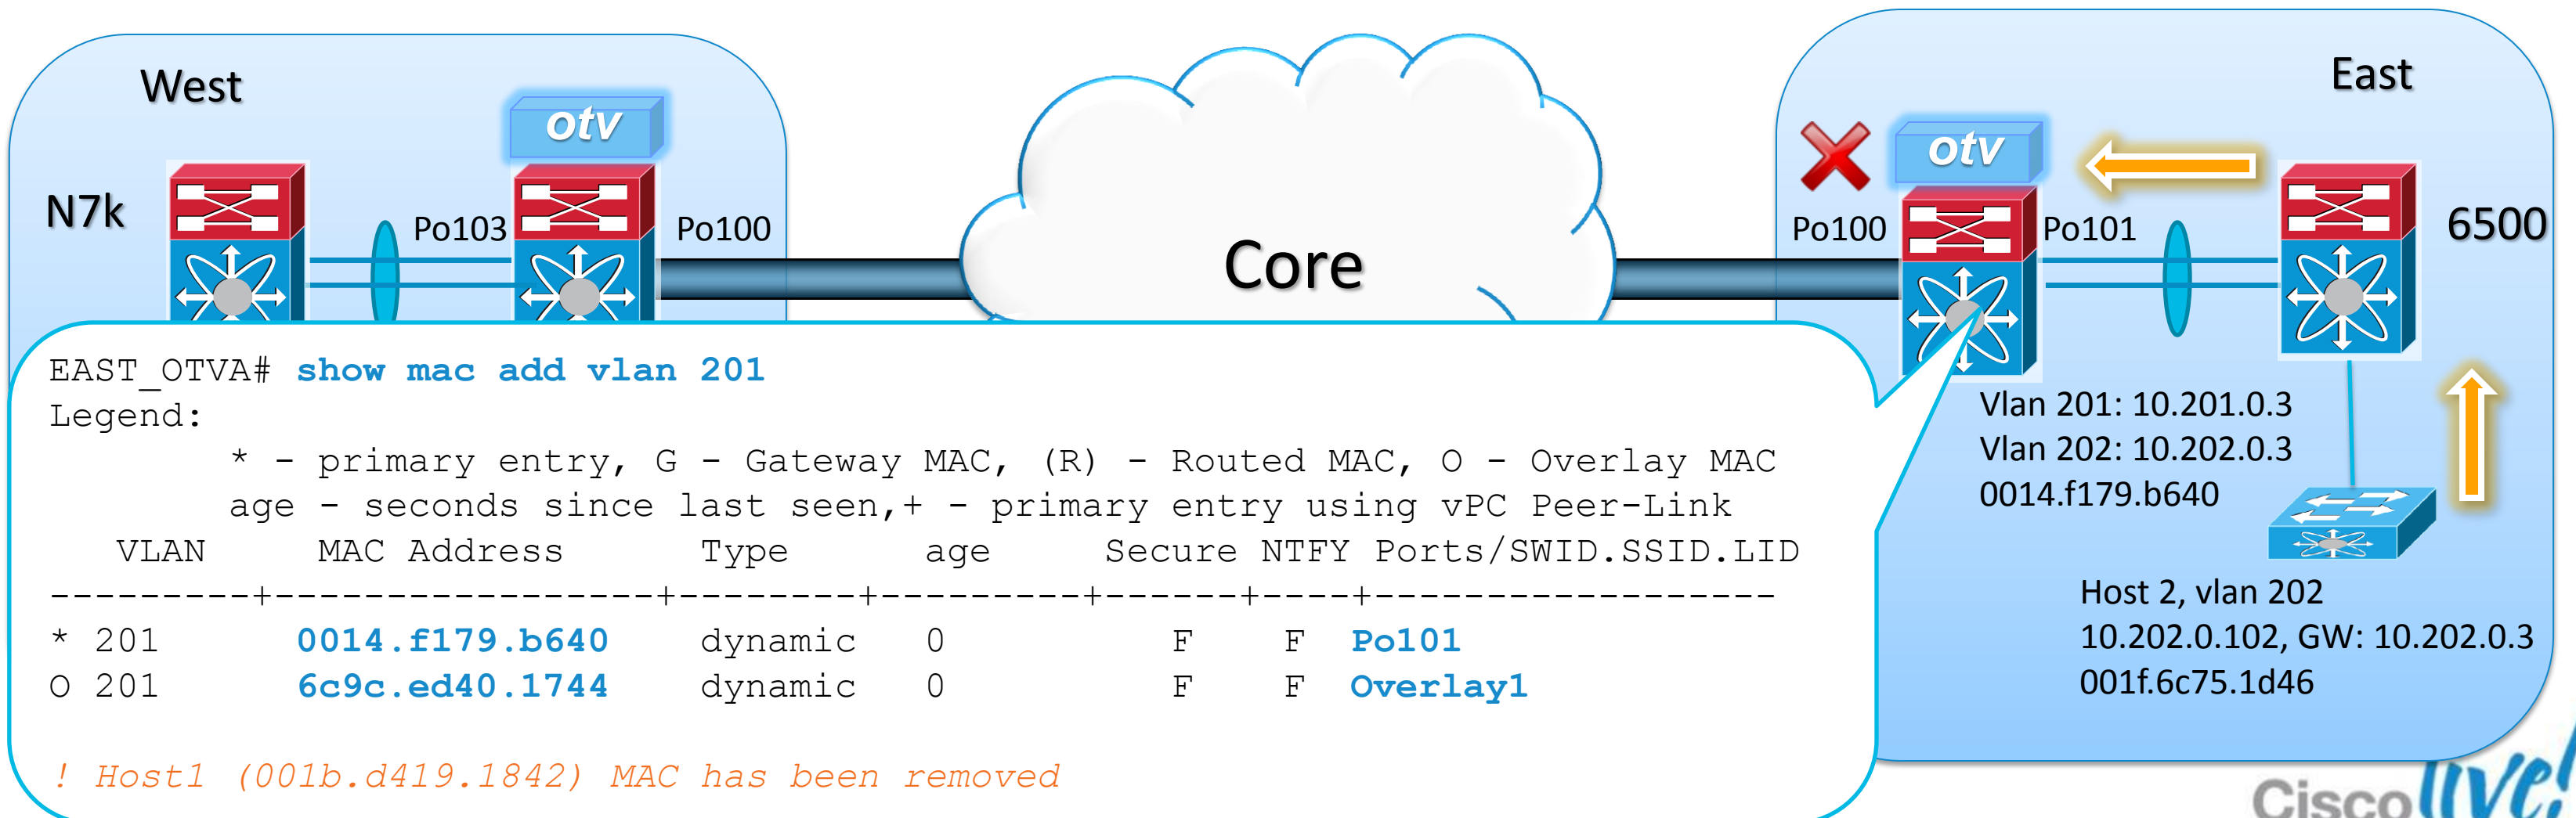
- Since the traffic flow between Host1 and Host2 is routed traffic, OTV will only see source MAC of the gateways and destination of the Hosts



# Troubleshooting

## ARP and CAM timer issue

- OTV does not send unknown unicast traffic across Overlay
- Subsequent packets from East toward Host1 will be dropped until Host1 MAC is relearned on West






# Troubleshooting

## ARP and CAM timer issue - Solution

- **Solution:** Change 6500 ARP timer to be less than OTV CAM timer



# Troubleshooting

## Network Load Balancer Services

- Some network load balancer services (NLBS) rely on flooding to reach all devices in the cluster
- Clusters that rely on a unicast IP to multicast MAC will be forwarded across overlay in same fashion as a broadcast packet without any additional configurations
  - Encapsulated within the control group (multicast transport)
  - Unicast to each OTV neighbour (unicast transport)
- Clusters that rely on a unicast IP to unicast MAC will be dropped

# Troubleshooting

## Network Load Balancer Services - Solution

- A static unicast MAC entry can be configured at a **single** site

Future support  
planned for selective  
unicast flooding

```
EAST_OTVA# show run | i static  
mac address-table static 0200.0ac9.00a2 vlan 201 interface port-channel101
```

```
EAST_OTVA# show otv route vlan 201
```

OTV Unicast MAC Routing Table For Overlay1

| VLAN | MAC-Address    | Metric | Uptime   | Owner | Next-hop(s)     |
|------|----------------|--------|----------|-------|-----------------|
| 201  | 0200.0ac9.00a2 | 1      | 00:00:11 | site  | port-channel101 |

West

otv

Po103

Po100

Po100

Po101

```
WEST_OTVA# show otv route vlan 201
```

OTV Unicast MAC Routing Table For Overlay1

| VLAN | MAC-Address    | Metric | Uptime   | Owner   | Next-hop(s) |
|------|----------------|--------|----------|---------|-------------|
| 201  | 0200.0ac9.00a2 | 42     | 00:04:34 | overlay | EAST_OTVA   |

# Summary

- **OTV Introduction**
- **Configuration**
  - Multicast Transport
  - Unicast-only Transport
- **Verification**
  - Adjacency
  - Unicast Forwarding
  - Multicast Forwarding
  - ARP
- **Troubleshooting**

# Q & A



# Complete Your Online Session Evaluation

## Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site [www.ciscoliveaustralia.com/mobile](http://www.ciscoliveaustralia.com/mobile)
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm



Cisco *live!* 365

Don't forget to activate your Cisco Live 365 account for access to all session material,

communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.

[www.ciscoliveaustralia.com/portal/login.ww](http://www.ciscoliveaustralia.com/portal/login.ww)


Cisco *live!*



# Appendix

## ASR 1000

- Support beginning in 3.5S
- Advance Enterprise Image or Advance IP Service (AES or AIS) to have the cli enabled
- Extended and site VLANs configured via EFP's and bridge-domains
- Multi-homing ASR and N7k OTV at **same site is not supported** (must be located at different sites)
- Support for multicast transport only (unicast transport planned in future release)
- [http://www.cisco.com/en/US/docs/ios-xml/ios/wan\\_otv/configuration/xe-3s/wan-otv-config.html](http://www.cisco.com/en/US/docs/ios-xml/ios/wan_otv/configuration/xe-3s/wan-otv-config.html)





# Appendix

## ASR 1000 – Configuration Internal Interface


- Site-ID and Site Bridge-Domain required
- Bridge-Domain must be forwarding on internal interface before adjacencies will be built

```
otv site bridge-domain 210
otv site-identifier 0003.0003.0003

interface GigabitEthernet1/0/2
no ip address
cdp enable
service instance 201 ethernet
encapsulation dot1q 201
bridge-domain 201
service instance 210 ethernet
encapsulation dot1q 210
bridge-domain 210
```

Bridge-domain for an extended VLAN

Site Bridge-domain must be active on internal interface



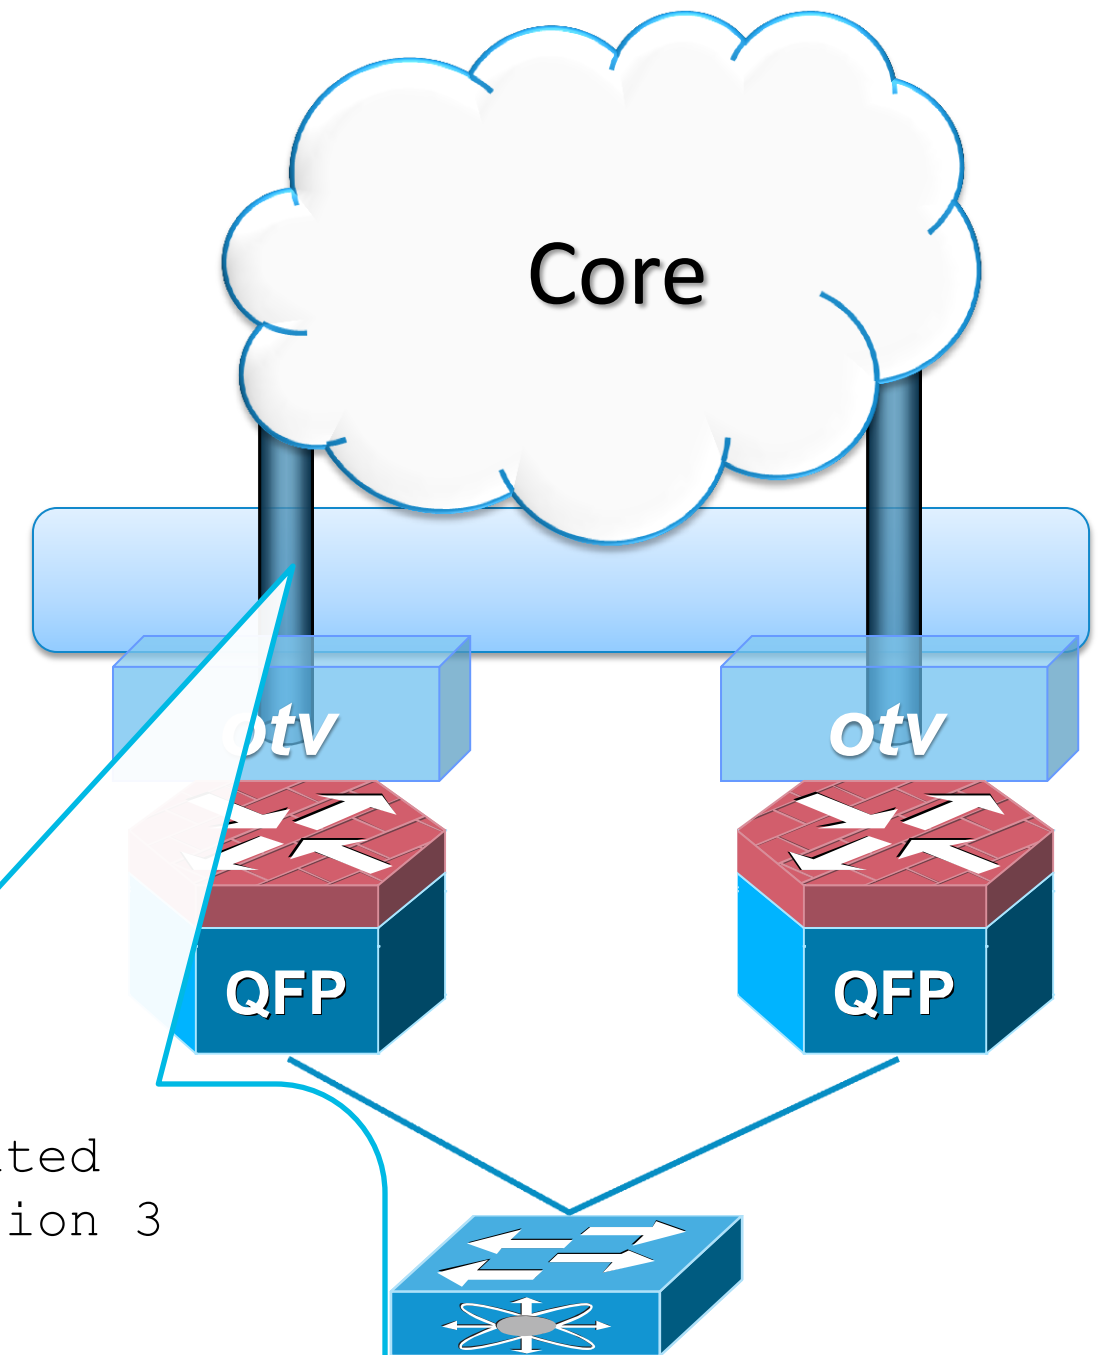
# Appendix

## ASR 1000 – Configuration Join Interface

- Join Interface must be configured with IGMPv3 for multicast transport.
- Multicast routing must be enabled
- Enable IGMP snooping querier
- Configure PIM Passive mode on Join Interface

```
ip multicast-routing distributed
ip igmp snooping querier version 3
ip igmp snooping querier

interface GigabitEthernet1/0/1
mtu 9000
ip address 172.16.1.18 255.255.255.252
ip pim passive
ip igmp version 3
```




# Appendix

## ASR 1000 – Configuration Overlay

- Configure control and data-groups
- Specify join-interface
- Create service instance for each bridge-domain that should be extended across overlay
- Do not extend site bridge-domain

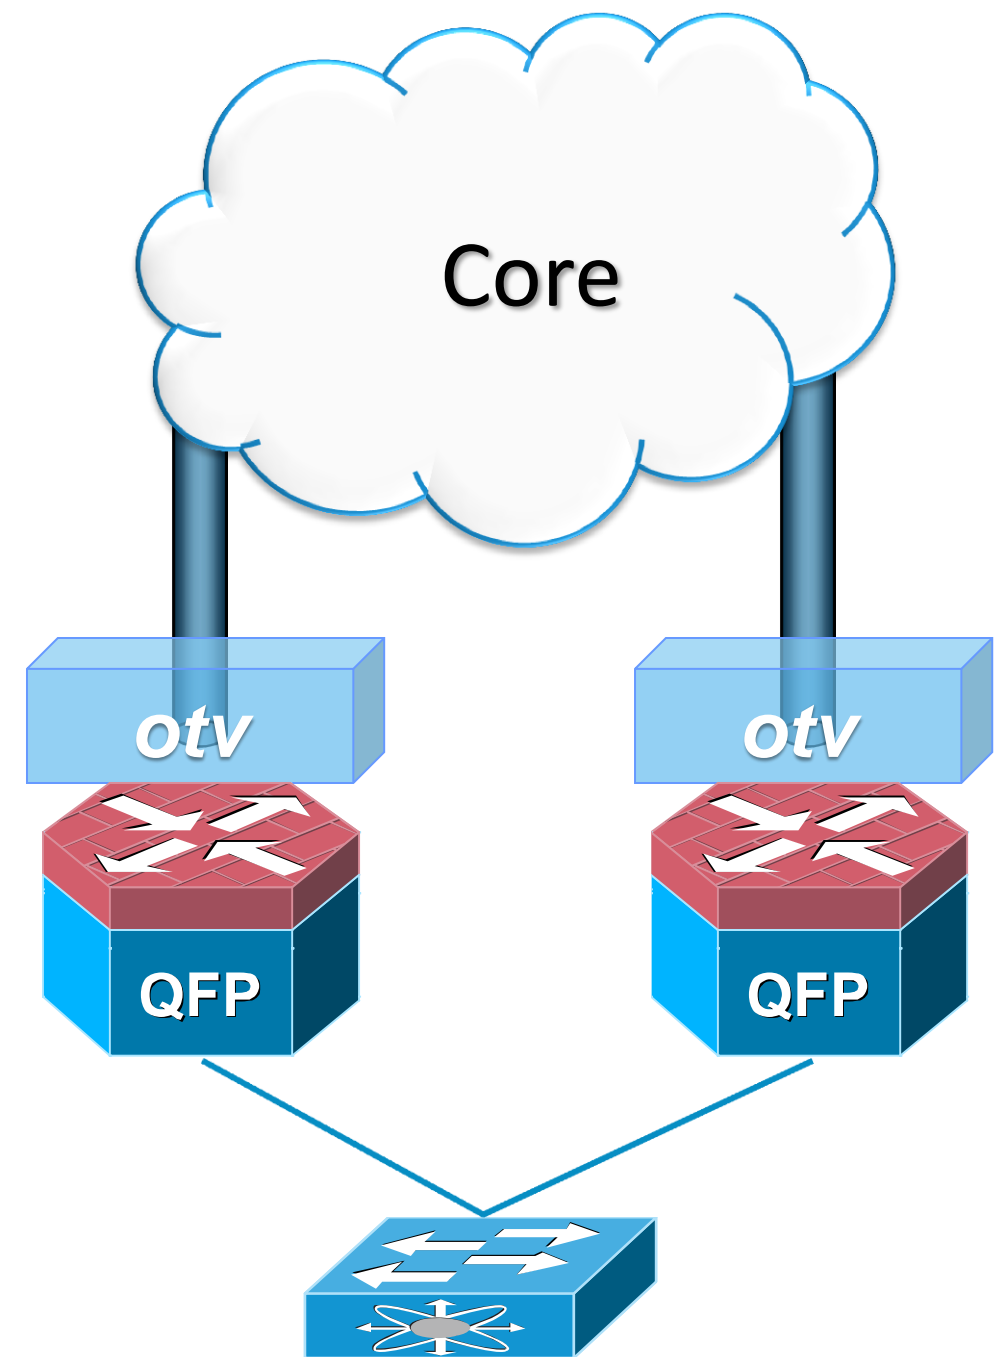
```
interface Overlay1
no ip address
otv control-group 239.1.1.1
otv data-group 232.1.1.0/24
otv join-interface GigabitEthernet1/0/1
service instance 201 ethernet
encapsulation dot1q 201
bridge-domain 201
```



# Appendix

## ASR 1000 – Verify Overlay is UP

```
SOUTH_OTVA#show otv
Overlay Interface Overlay1
VPN name                : None
VPN ID                  : 1
State                   : UP
AED Capable             : Yes
IPv4 control group      : 239.1.1.1
Mcast data group range(s) : 232.1.1.0/24
Join interface(s)       : GigabitEthernet1/0/1
Join IPv4 address       : 172.16.1.18
Tunnel interface(s)     : Tunnel0
Encapsulation format    : GRE/IPv4
Site Bridge-Domain      : 210
Capability               : Multicast-reachable
Is Adjacency Server     : No
Adj Server Configured   : No
Prim/Sec Adj Svr(s)     : None
```



# Appendix

## ASR 1000 – Verify Site Adjacency and AED

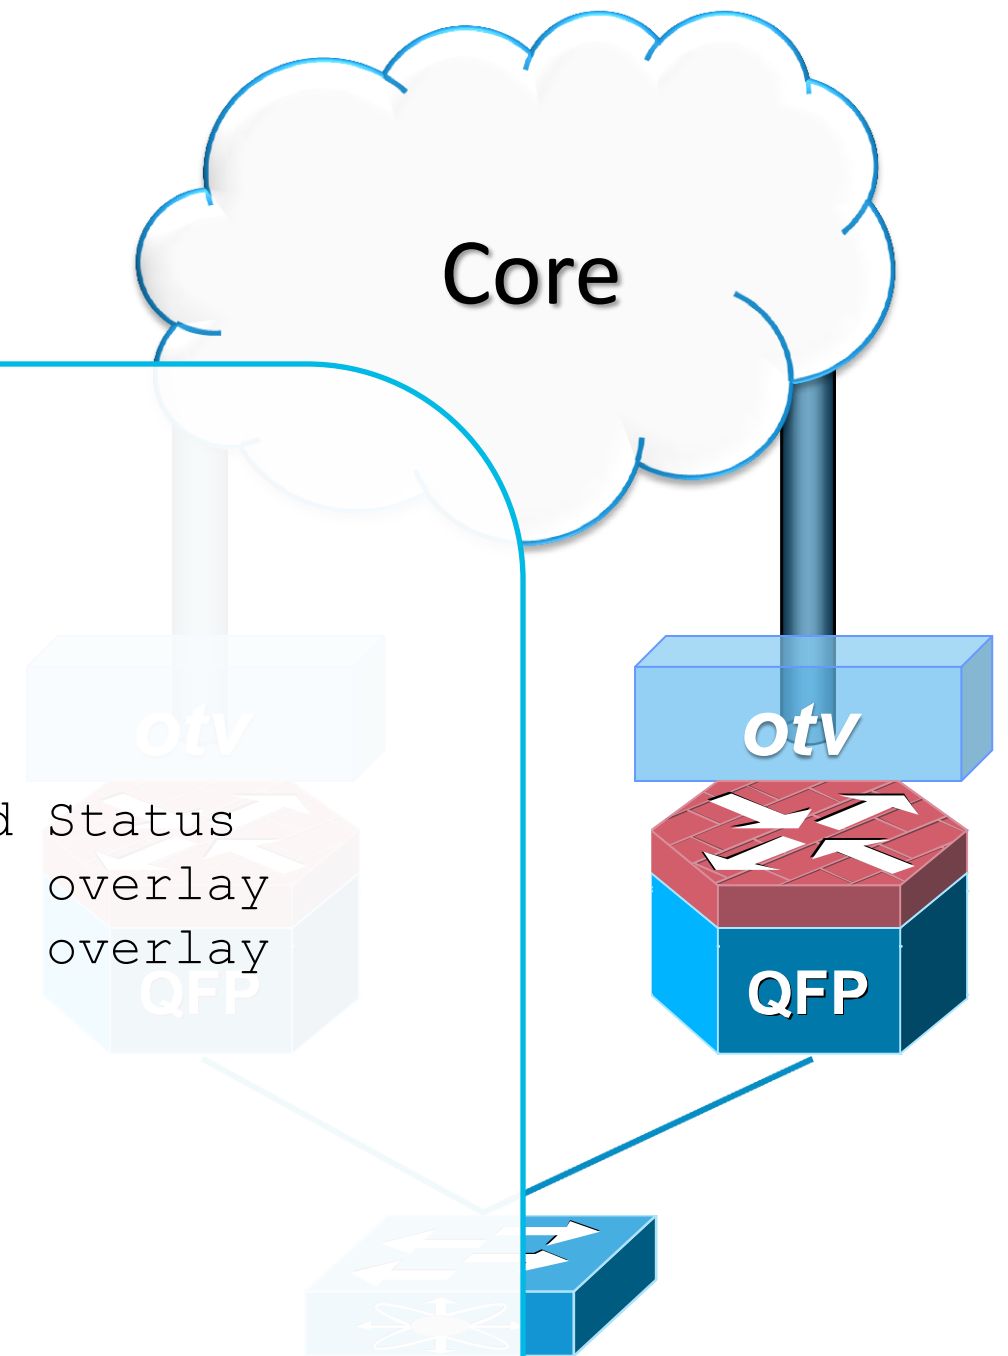
```
SOUTH_OTVA#show otv site
Site Adjacency Information (Site Bridge-Domain: 210)
```

```
Overlay1 Site-Local Adjacencies (Count: 2)
```

| Hostname     | System ID      | Last Change | Ordinal | AED Enabled | Status  |
|--------------|----------------|-------------|---------|-------------|---------|
| SOUTH_OTVB   | 001D.707E.1B00 | 01:02:23    | 0       | site        | overlay |
| * SOUTH_OTVA | 001D.707E.3A00 | 00:42:08    | 1       | site        | overlay |

```
SOUTH_OTVA#show otv vlan
Key:  SI - Service Instance
```

```
Overlay 1 VLAN Configuration Information
Inst  VLAN  Bridge-Domain  Auth  Site Interface(s)
0     201    201             yes   Gi1/0/2:SI201
Total VLAN(s): 1
Total Authoritative VLAN(s): 1
```



# Appendix


## ASR 1000 – Verify Overlay Adjacencies

```
SOUTH_OTVA#show otv adjacency
```

```
Overlay 1 Adjacency Database
```

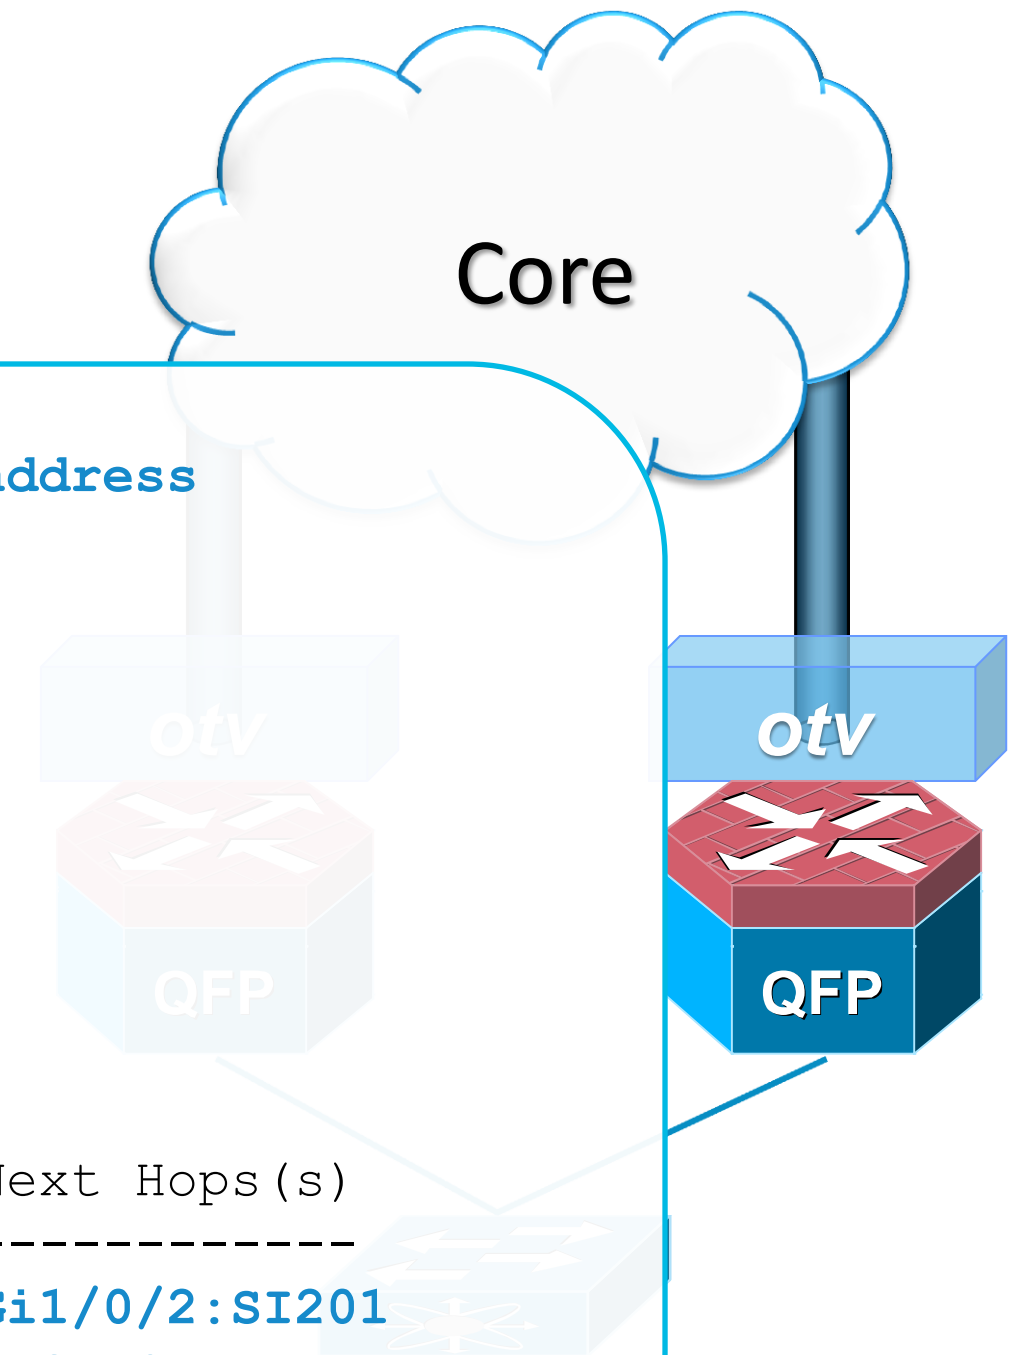
| Hostname   | System-ID      | Dest Addr   | Up Time  | State |
|------------|----------------|-------------|----------|-------|
| EAST_OTVB  | 64a0.e741.c841 | 172.16.1.30 | 01:06:21 | UP    |
| WEST_OTVB  | 64a0.e741.c842 | 172.16.1.38 | 01:06:21 | UP    |
| WEST_OTVA  | 6c9c.ed40.1742 | 172.16.1.34 | 01:06:21 | UP    |
| EAST_OTVA  | 6c9c.ed40.1741 | 172.16.1.26 | 01:06:21 | UP    |
| SOUTH_OTVB | 001d.707e.1b00 | 172.16.1.22 | 00:29:48 | UP    |

Peering between ASR and N7k between sites is supported.



# Appendix

## ASR 1000 – Verify Locally and Remotely Learned Routes



```
SOUTH_OTVA#show bridge-domain 201 mac dynamic address
```

```
Port                MAC Address
Gi1/0/2 ServInst 201 001a.e2be.52cd
```

```
SOUTH_OTVA#show otv route vlan 201
```

```
Codes: BD - Bridge-Domain, AD - Admin-Distance,
       SI - Service Instance, * - Backup Route
```

```
OTV Unicast MAC Routing Table for Overlay1
```

| Inst | VLAN | BD  | MAC Address    | AD | Owner  | Next Hops (s) |
|------|------|-----|----------------|----|--------|---------------|
| 0    | 201  | 201 | 001a.e2be.52cd | 40 | BD Eng | Gi1/0/2:SI201 |
| 0    | 201  | 201 | 001b.d419.1842 | 50 | ISIS   | WEST_OTVA     |
| 0    | 201  | 201 | 001f.6c75.1d42 | 50 | ISIS   | EAST_OTVA     |

Locally Learned MAC  
from Bridge-domain 201

MAC learned via ISIS  
from OTV peer

# Appendix

## ASR 1000 – Multicast Local Receiver

```
SOUTH_OTVA#show otv mroute
```

```
OTV Multicast Routing Table for Overlay1
```

```
Bridge-Domain = 201, s = *, g = *
```

```
Outgoing interface list:
```

```
Default, NoRedist
```

```
Incoming interface count = 0, Outgoing interface count = 1
```

```
Bridge-Domain = 201, s = *, g = 224.10.10.10
```

```
Outgoing interface list:
```

```
Service Instance 201, GigabitEthernet1/0/2
```

```
Incoming interface count = 0, Outgoing interface count = 1
```

```
Bridge-Domain = 201, s = 10.201.0.102, g = 224.10.10.10
```


```
Incoming interface list:
```

```
Service Instance 201, Overlay1, 001f.6c75.1d42
```

```
Incoming interface count = 1, Outgoing interface count = 0
```

(Bridge-domain, \*,G)  
programmed based on  
IGMP join from client

(Bridge-domain, S,G)  
created to deliver to local  
receiver once received  
from overlay





# Appendix

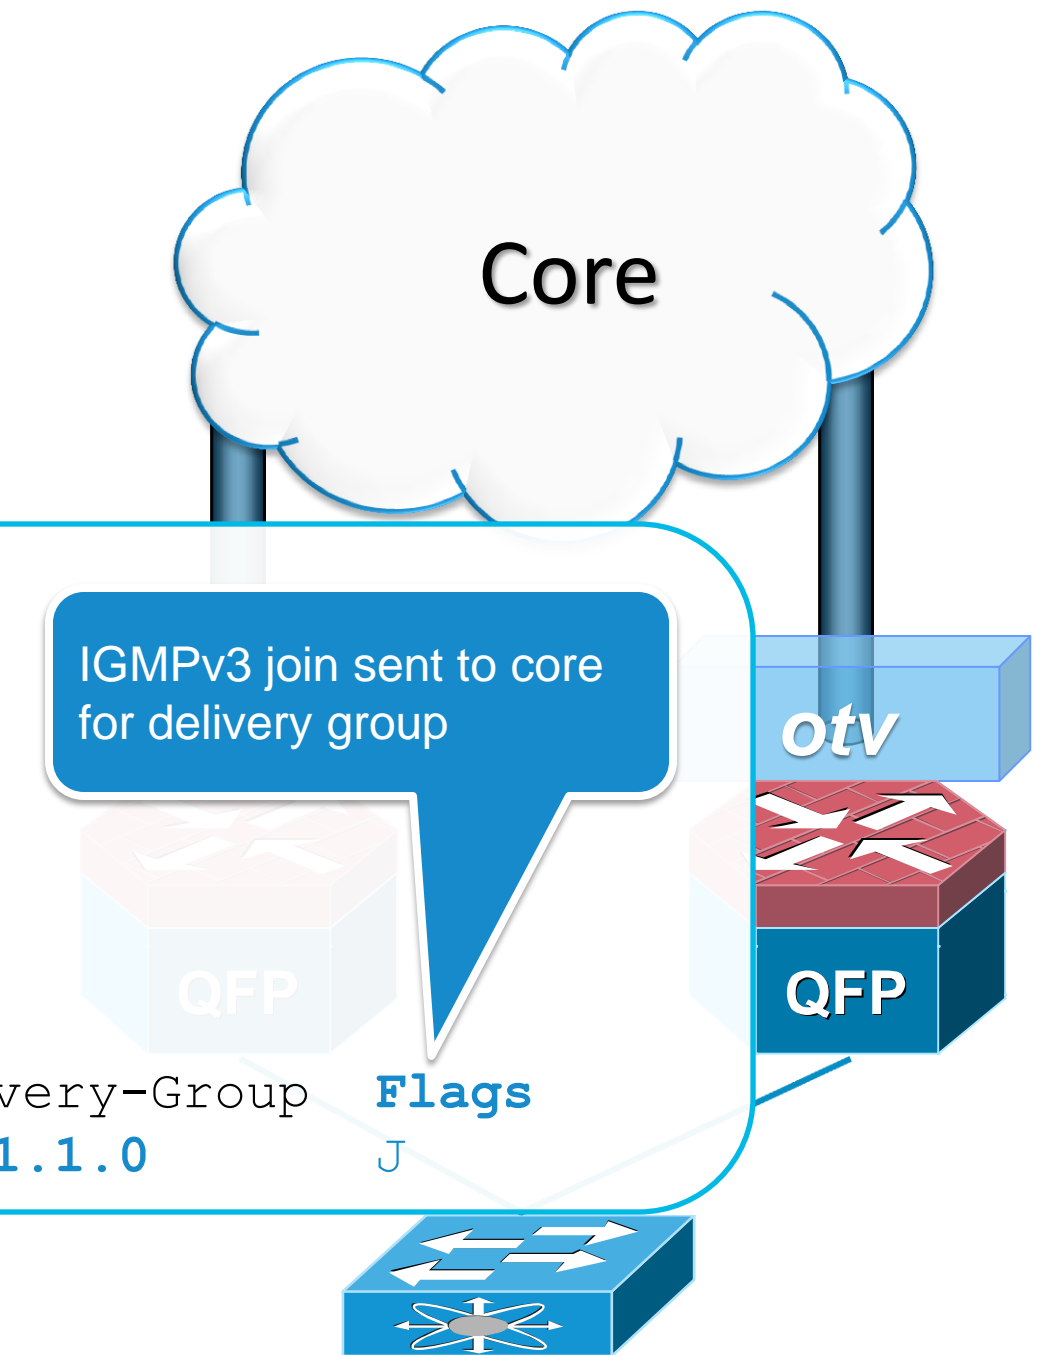
## ASR 1000 – Multicast Local Receiver

```
SOUTH_OTVA#show otv data-group
```

```
Flags:  D - Local active source dynamically detected  
        S - Local active source statically configured  
        J - Data group has been joined in the core  
        U - Data group has not been joined in the core
```

```
Remote Active Sources for Overlay1
```

| BD  | Active-Source | Active-Group | Delivery-Source | Delivery-Group | Flags |
|-----|---------------|--------------|-----------------|----------------|-------|
| 201 | 10.201.0.102  | 224.10.10.10 | 172.16.1.26     | 232.1.1.0      | J     |



# Appendix

## ASR 1000 – Multicast Local Source

```
SOUTH_OTVA#show otv data-group
```

```
No remote data-group mappings
```

```
Flags:  D - Local active source dynamically detected  
        S - Local active source statically configured  
        J - Data group has been joined in the core  
        U - Data group has not been joined in the core
```

```
Local Active Sources for Overlay1
```

| BD  | Active-Source | Active-Group | Delivery-Source | Delivery-Group | Flags |
|-----|---------------|--------------|-----------------|----------------|-------|
| 201 | 10.201.0.105  | 224.50.50.50 | 172.16.1.18     | 232.1.1.0      | D     |

```
SOUTH_OTVA#show otv mroute
```

```
! Some output omitted
```

```
Bridge-Domain = 201, s = *, g = 224.50.50.50
```

```
Outgoing interface list:
```

```
Overlay1, EAST_OTVA
```

```
Incoming interface count = 0, Outgoing interface count = 1
```

```
Bridge-Domain = 201, s = 10.201.0.105, g = 224.50.50.50
```

```
Incoming interface list:
```

```
Service Instance 201, GigabitEthernet1/0/2, 001a.e2be.52cd
```

```
Incoming interface count = 1, Outgoing interface count = 0
```

Local Source Flag

Join advertised from remote site across overlay

Local source on internal interface

