# SD-WAN OMP Unlocked: From Session Bring-Up to Advanced Routing, Multicast, and L2VPN

Waqas Daar
Customer Success Specialist
SDWAN/Thousand Eyes

cisco Live !

# Webex App

**Questions?**

Use Webex App to chat with the speaker after the session

**How**

1. Find this session in the Cisco Events App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

**Webex spaces will be moderated by the speaker until February 27, 2026.**

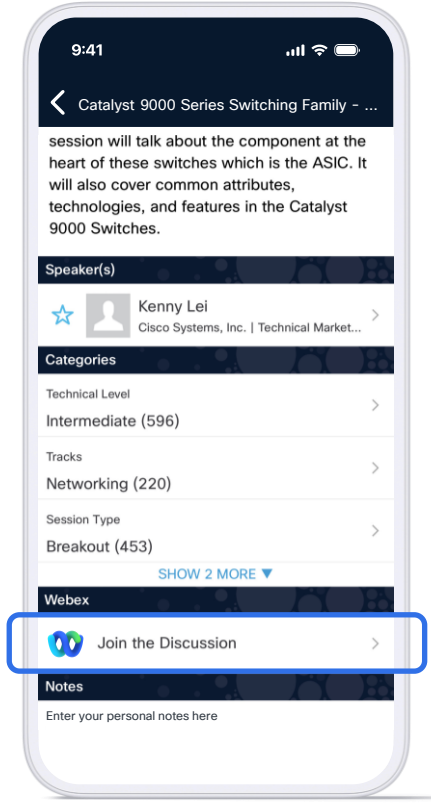

# “If you cannot explain it simply, you don’t understand it well enough.”

# Who am I?



# Quick SD-WAN Recap

# Cisco SD-WAN Fabric overview

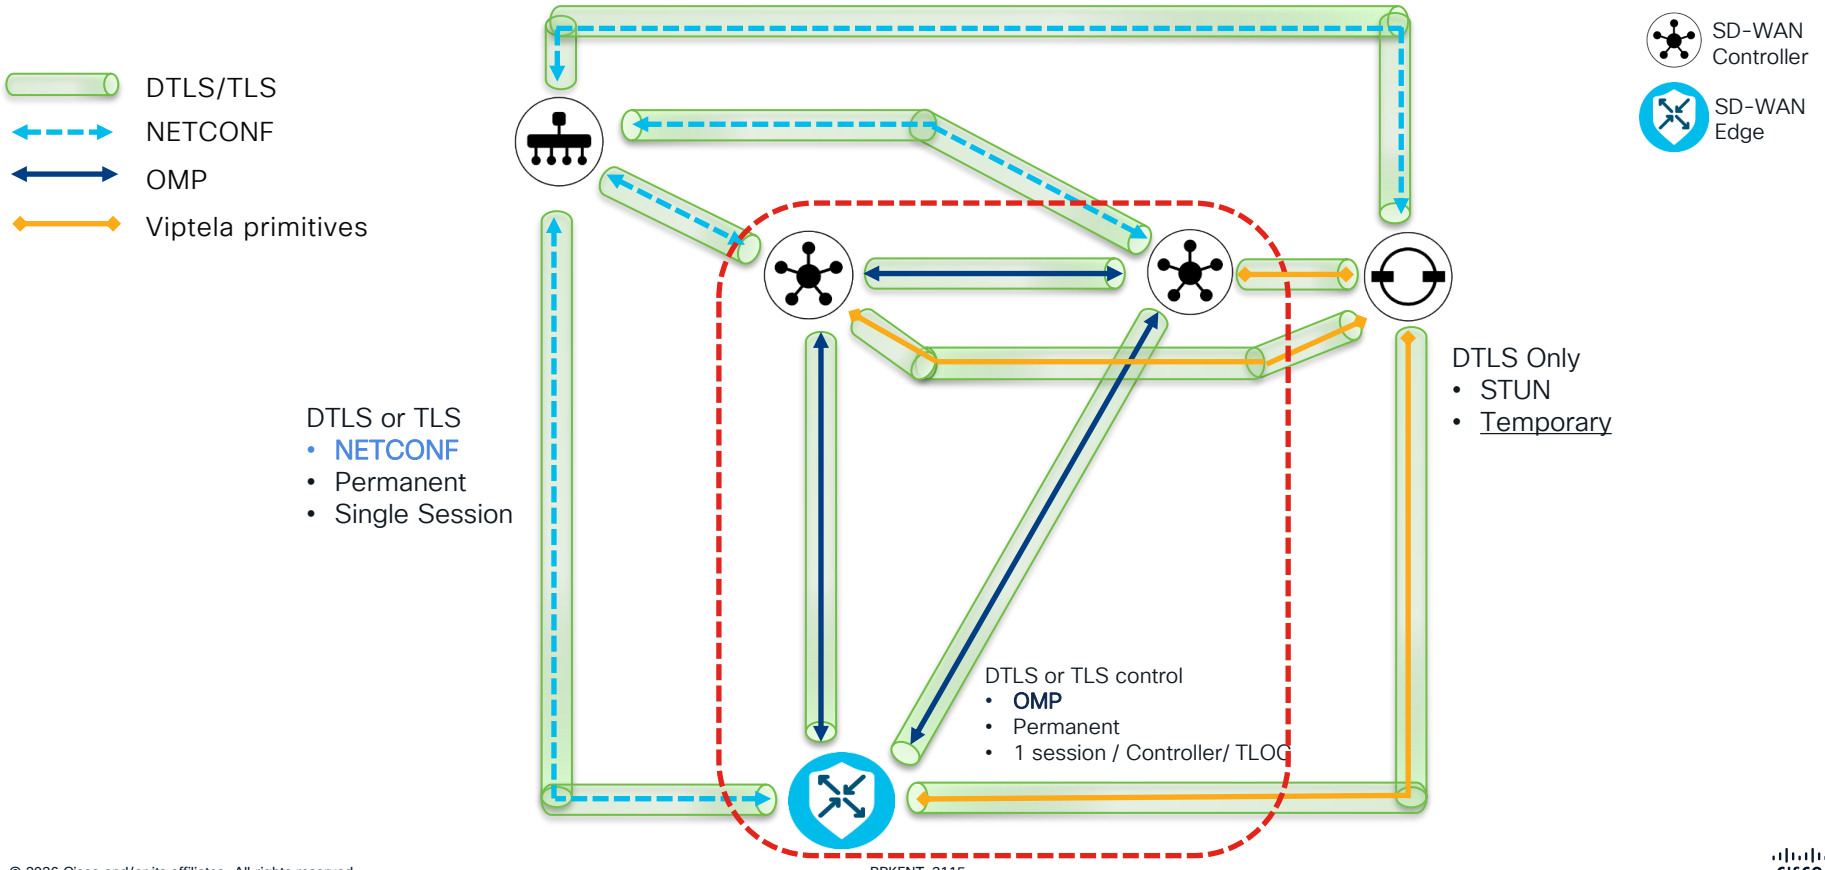

# How SD-WAN Control Connections (CC) are established? 1/3

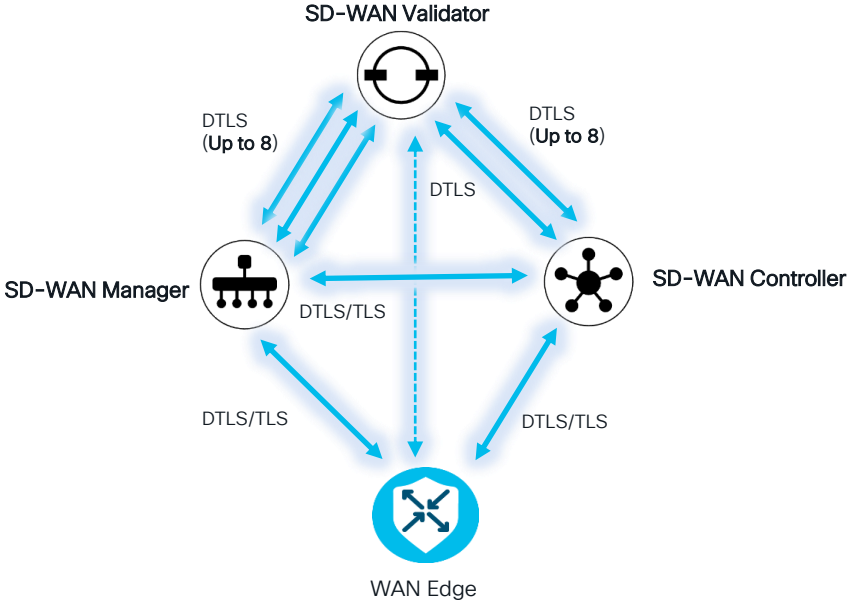


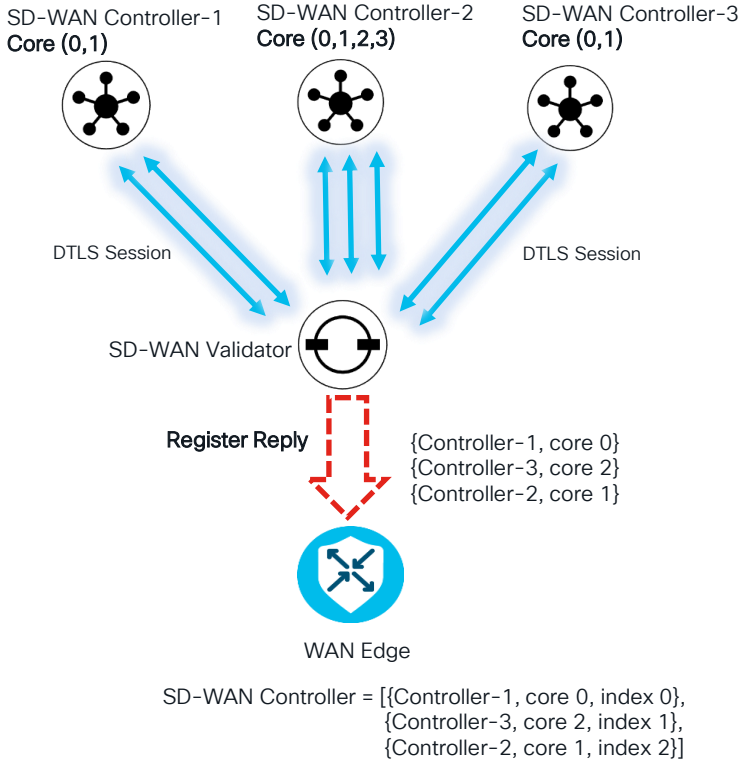

# How SD-WAN Control Connections (CC) are established? 2/3

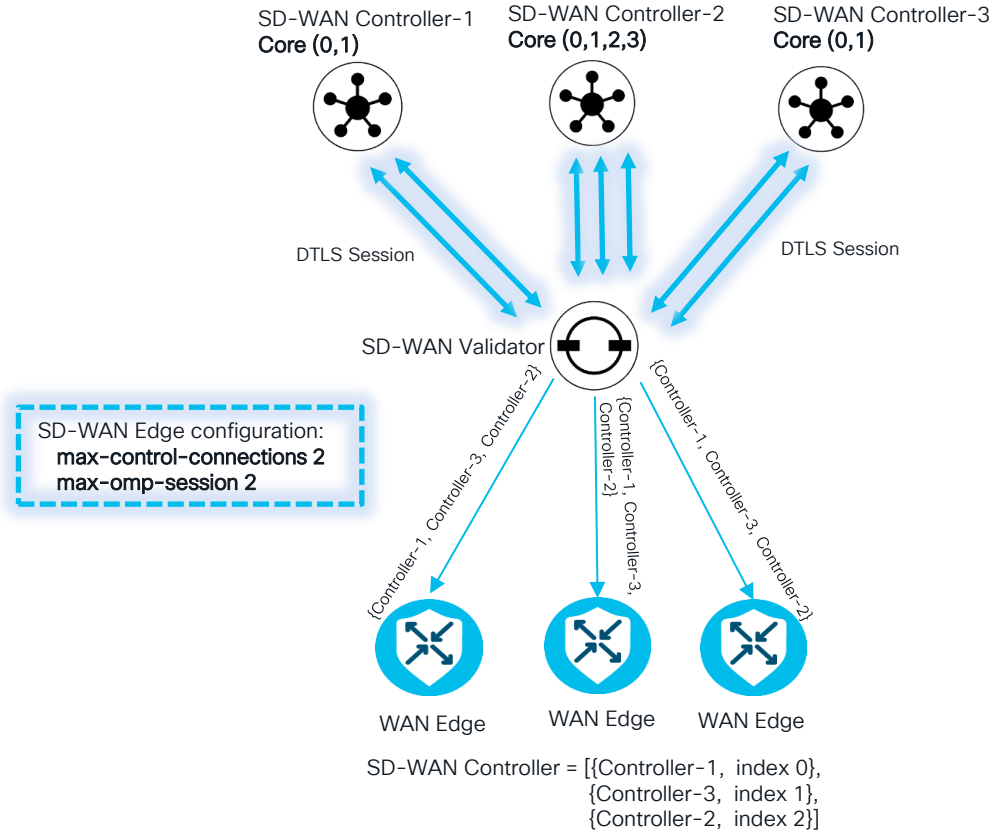

# How SD-WAN Control Connections (CC) are established? 3/3

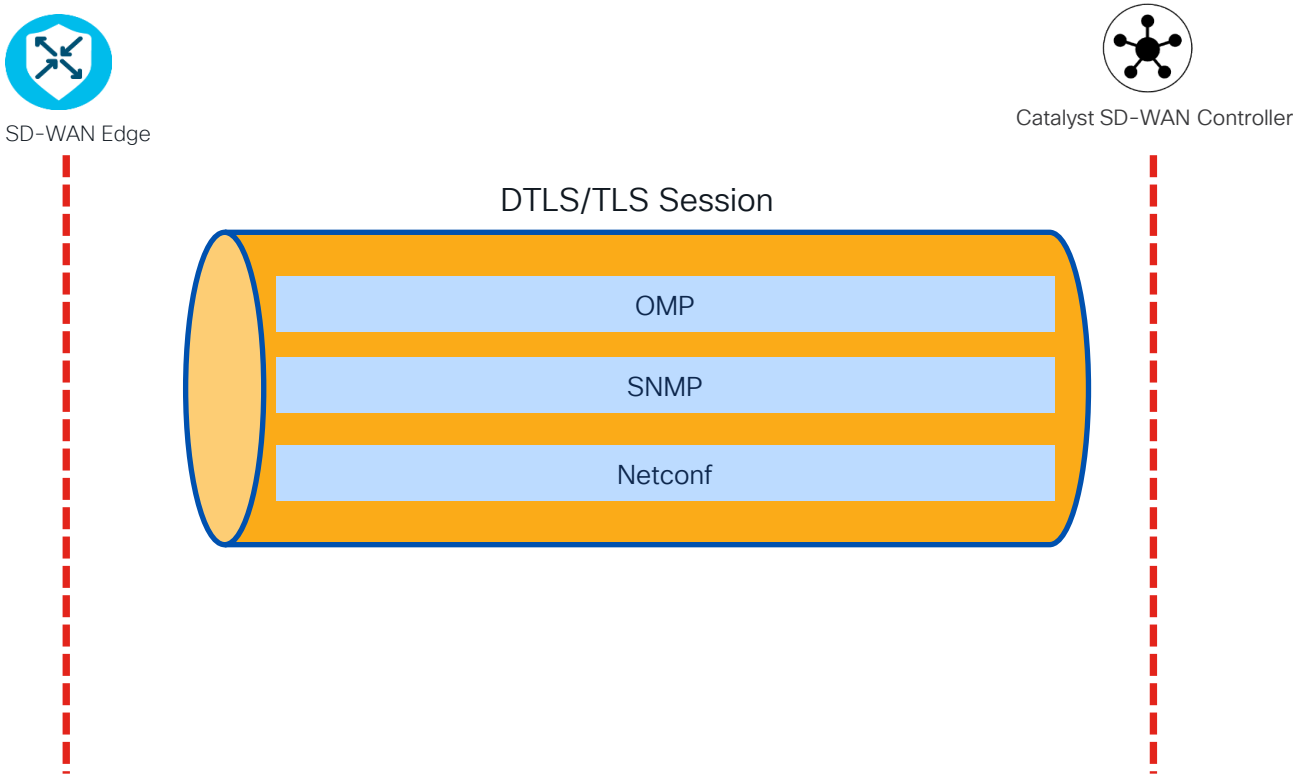

# Overlay Management Protocol (OMP)

# Overlay Management Protocol (OMP)

OMP is a path vector routing protocol derived from BGP.

01.

OMP peering operates over the control connection (TLS/DTLS) via TCP port 17900.

02.

Leverages address families for reachability advertisement.

03.

OMP provides native loop-avoidance capabilities.

04.

# What information is propagated by OMP ?

| vRoutes | TLOCs | Policies | Services |
| --- | --- | --- | --- |
| Service-side Routing info redistributed to Overlay + Attributes | WAN Attachments: system-ip link color encapsulation + Attributes | Data Policy App-Route Policy VPN Membership cFlowd Template | Services: Type of Service, Location (TLOC), Forwarding Information |

# Demystifying OMP Peering

# OMP Peering Overview

- OMP peering uses System-IP addresses inside the control connection (CC).
- Only one OMP peering session per WAN Edge ↔ SD-WAN Controller pair (regardless of multiple DTLS/TLS control connections).
- OMP peering requires a TCP connection to be established first between peers.
- TCP Roles (Active vs. Passive):
  - WAN Edge to Controller: WAN Edge is always Active; Controller is Passive.
  - Controller to Controller**:** The device with the lower System-IP is Active.

# OMP Packet details

| | |
|---|---|
| HANDSHAKE | Once the TCP session is established, the first packet which is exchange between peers is HANDSHAKE and it contains information like SITE-ID, HOLD-TIME and capability information like what it suports Multi-Protocol (MP) IPv4, MP IPv6, Multi-cast, services etc.. just like BGP OPEN message. |
| HELLO | OMP peers periodicly exchange HELLO packets to indicate each peer is alive and reachable. |
| UPDATE | This message is used to transfer routing information between peers. OMP update message is used to advertise feasible routes that share common path attributes to a peer or to withdraw multiple unfeasible routes. |
| QUERY | This message is used to send a request for a specific route for which an aggregate or else specific route exists. Query message is ONLY send by the WAN Edge router once it finds out that group of prefixes received is equipped with a query attribute. |

# OMP Packet details

| | |
|---|---|
| ALERT | Once the error condition has been deducted then peer used ALERT message to notify the opposite peer. The format and intention of the message is like BGP NOTIFICATION message. |
| INFORM | This OMP message is associated with the OMP GRACEFUL RESTART feature. For example, once OMP session went down and came back up then peer send End-of-RIB EOR marker, that is an OMP INFORM packet. |
| POLICY | This OMP message entails all the policies including control policies, centralized data policies, Application Aware Routing (AAR), cFlow template etc… |

# Finite State Machine for OMP Peer Establishment

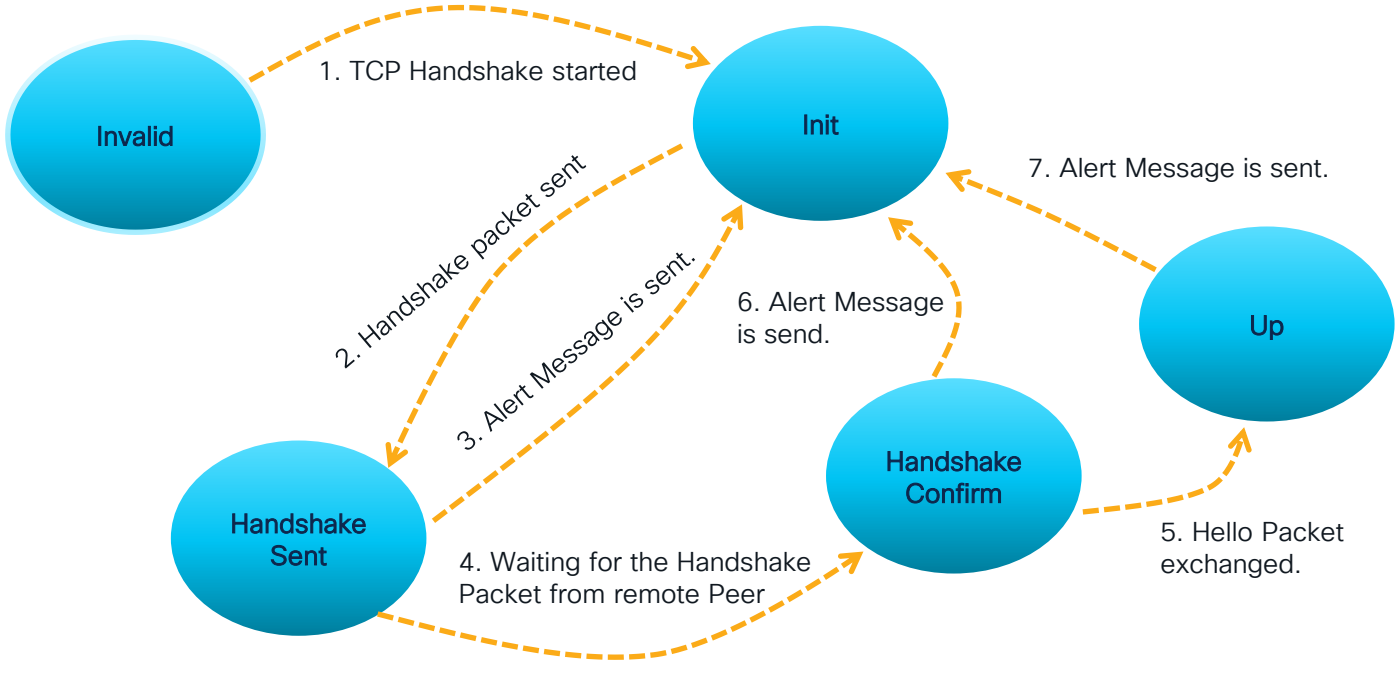

# Conditions for Establishing OMP Peer Relationship

- When Handshake packet is received from the peer, following conditions met:
  - WAN Edge and Catalyst controllers belong to same domain (**domain-id**).
  - Atleast ONE address family matches.
  - Peer capability need to matches.
- If any of the conditions are not met, an ALERT packet is sent with the appropriate error code, and the OMP peering session is terminated.
- Upon receiving a Hello packet from a peer, the peer validates the correctness of the peer address. If there is a discrepancy, an ALERT packet is sent with an error code, leading to the termination of the OMP peering session.

# OMP Packet Insights and commonalities with BGP

# Exploring Key Similarities between BGP and OMP

**BGP** employs TCP as its transport protocol (**port 179**) to ensure all the transport reliability is taken care of by TCP.

**BGP** speakers exchange **KEEPALIVE** packets to maintain the connection keep alive.

**BGP** provides a mechanism to gracefully close a connection with the peer by sending a **NOTIFICATION error** message.

When **BGP** session has been established all candidate BGP routes are exchanged, then after that **only incremental updates** are sent.

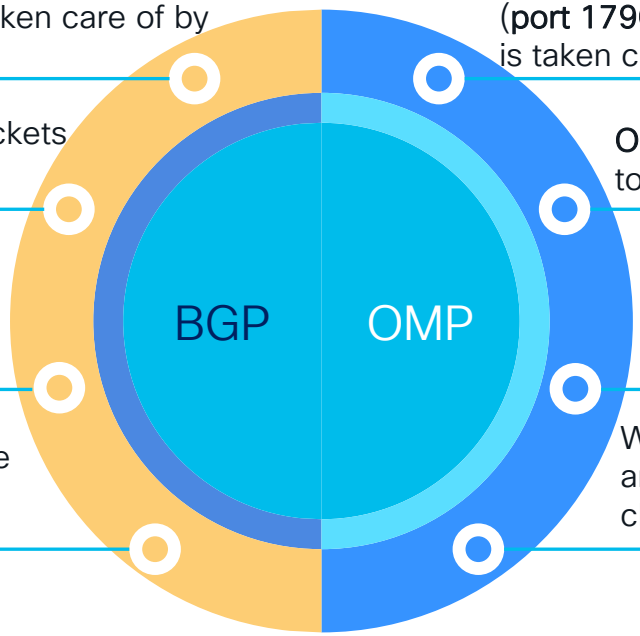


**OMP** also employs TCP as its transport protocol (**port 17900**) to ensure all the transport reliability is taken care by TCP.

**OMP** peers also exchange **HELLO packets** to maintain the peering alive.

**OMP** also provides a mechanism to gracefully close a connection with the peer by sending a **ALERT error** message.

When **OMP peering** has been established and routes are exchanged then if any change **only incremental updates** are sent.

# Let’s recap classic BGP Routing Process

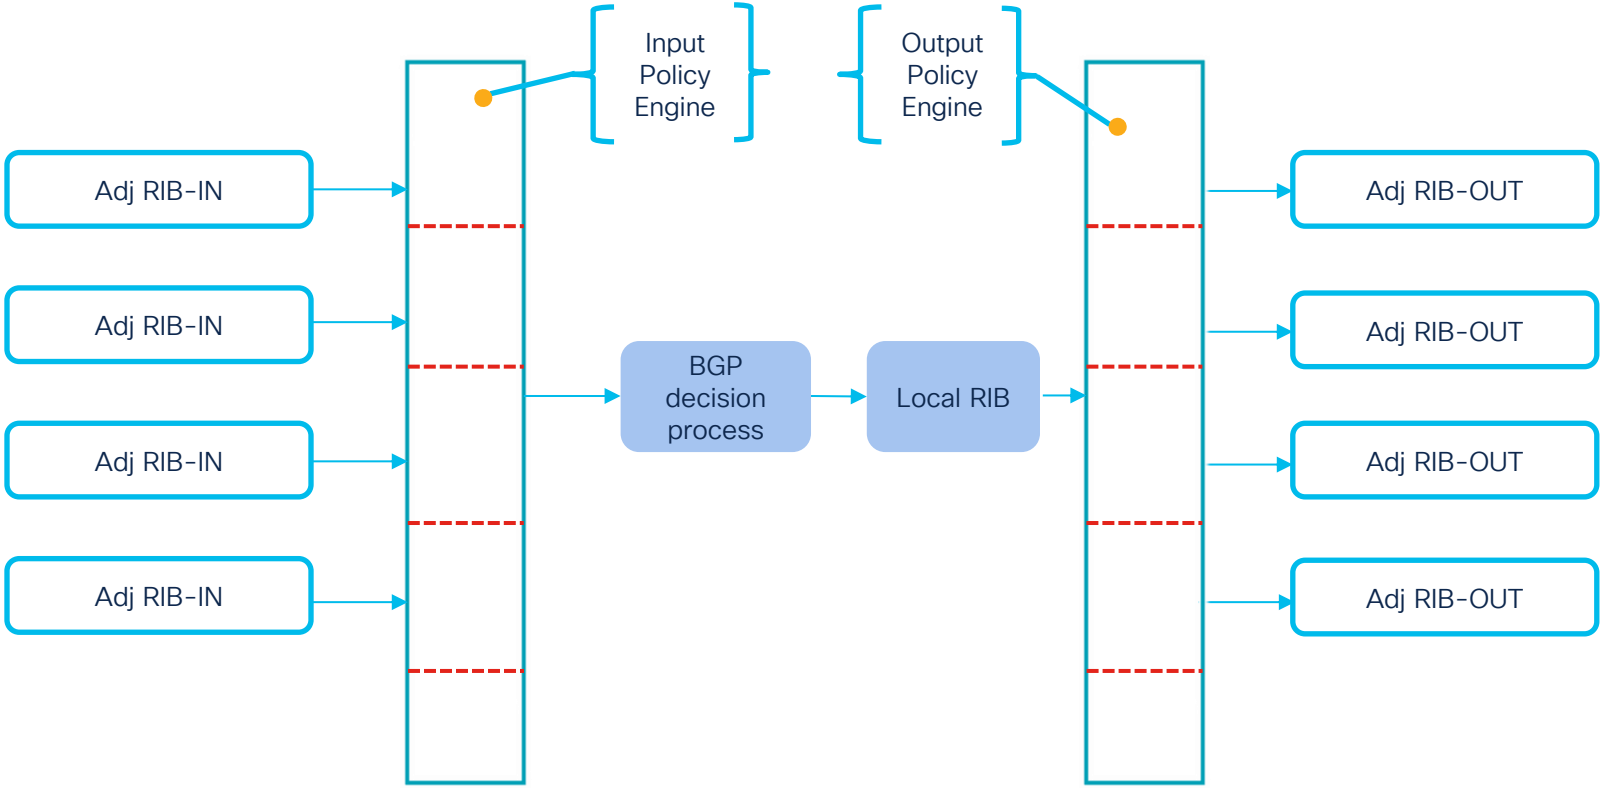


**Source**: Internet Routing Architectures by Sam Halabi

# SD-WAN Controller OMP Routing Process

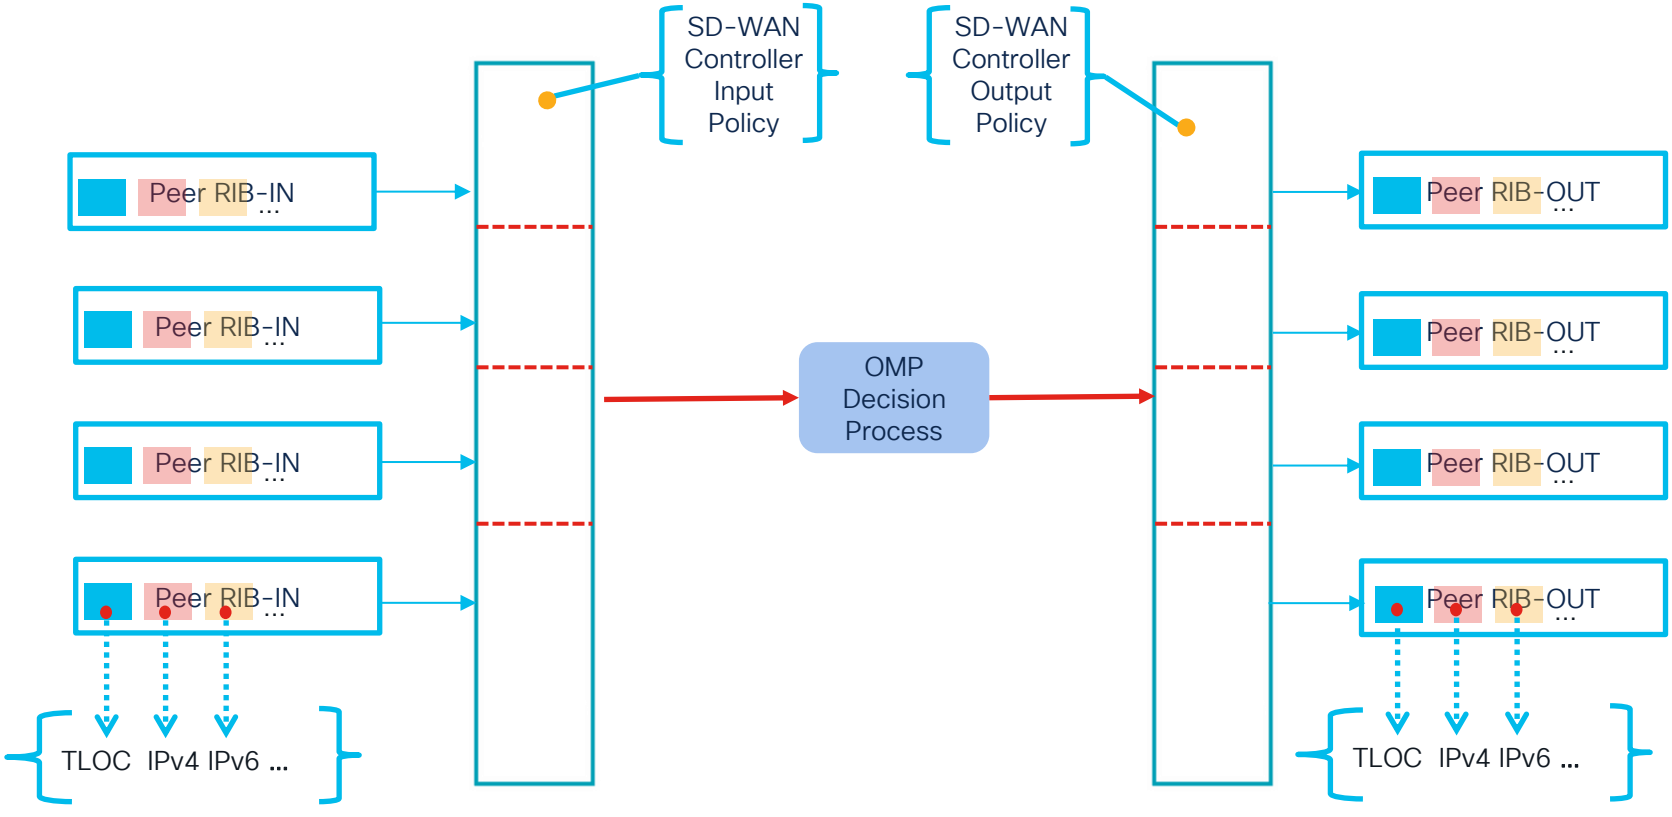

# Key points to remember for SD-WAN Controller OMP Routing Process

SD-WAN controller have separate RIB-IN and RIB-OUT for each OMP peer.

SD-WAN controller have further separate RIB-IN and RIB-OUT for each address family for example TLOC/prefix/service etc.. for each OMP peer.

From SD-WAN control policy perspective, when SD-WAN control policy is applied in IN bound direction then what ever action we perform either on TLOC/prefix/service it will be stored and shown in RIB-IN for that address-family for that peer and FLAG will be set accordingly.

If we apply SD-WAN policy in OUT bound direction, then if action is REJECT either for TLOC/prefix/service etc.., then RIB-OUT will NOT be generated for that address family, so it will NOT be advertised to other peers.

If we apply SD-WAN control policy in OUT bound direction and change/modified either TLOC/prefix/service etc.. the attributes, then RIB-OUT will be generated as per our control policy/data policy and advertised to other peers.

# Optimizing OMP vRoute Advertisements with System-IP

# Inside the OMP RIB-Out Architecture: Current Implementation

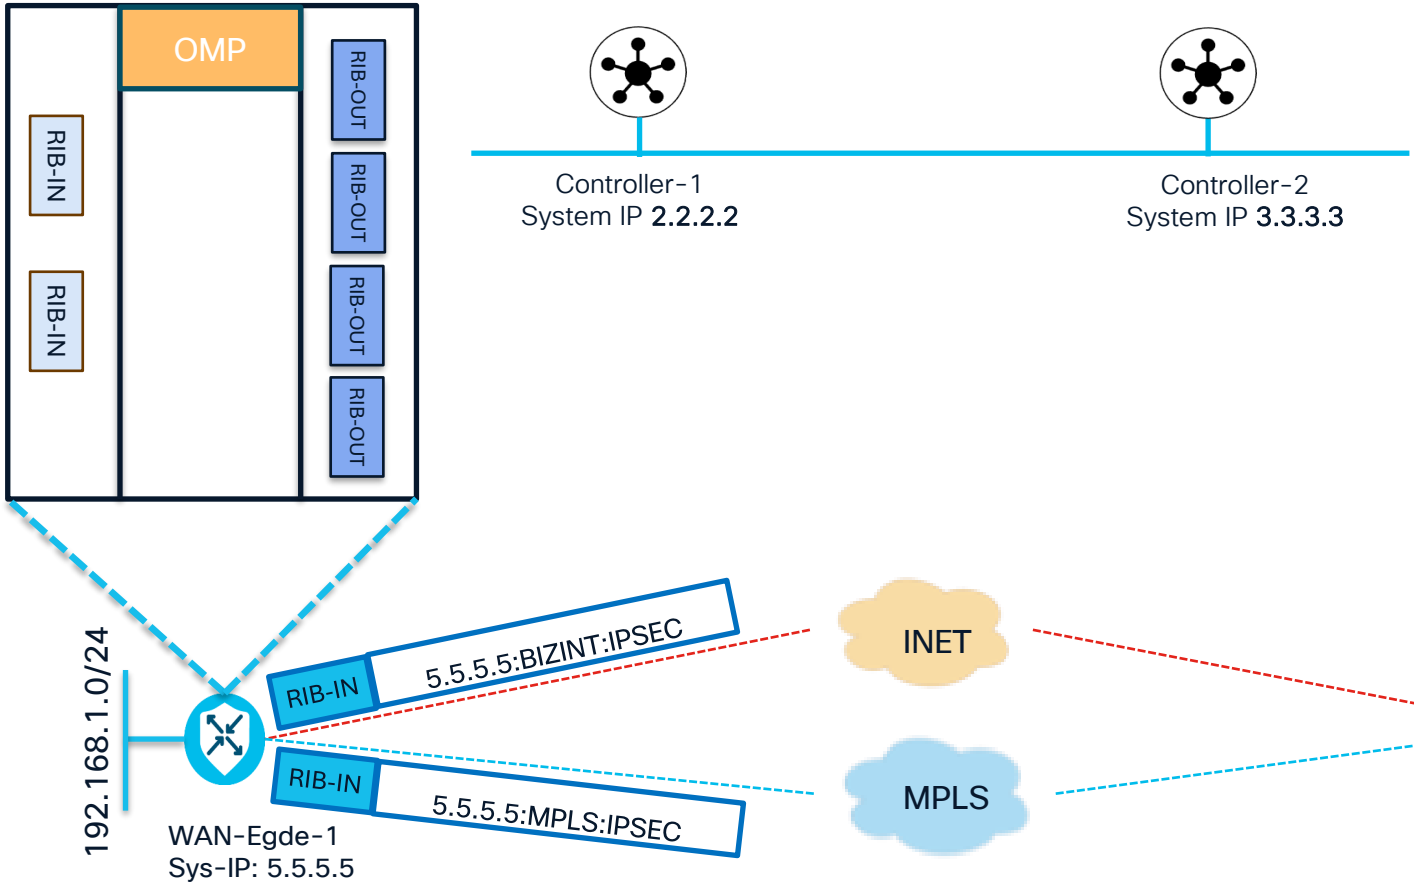

# Inside the OMP RIB-Out Architecture: Current Implementation

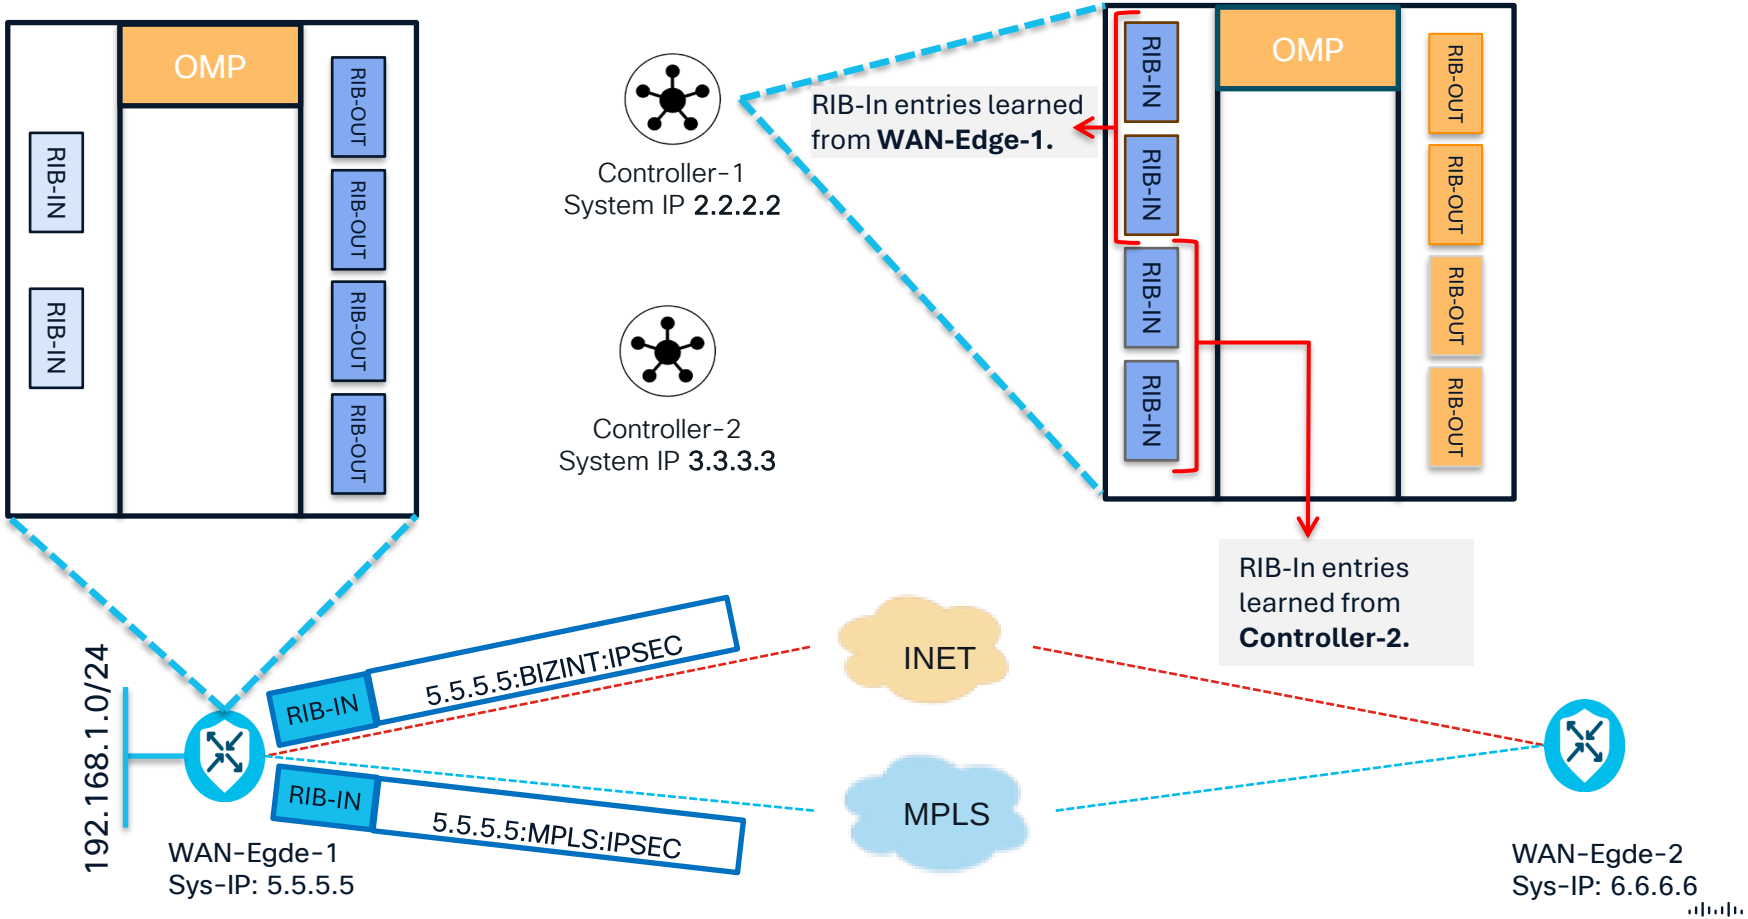

# Inside the OMP RIB-Out Architecture: Current Implementation

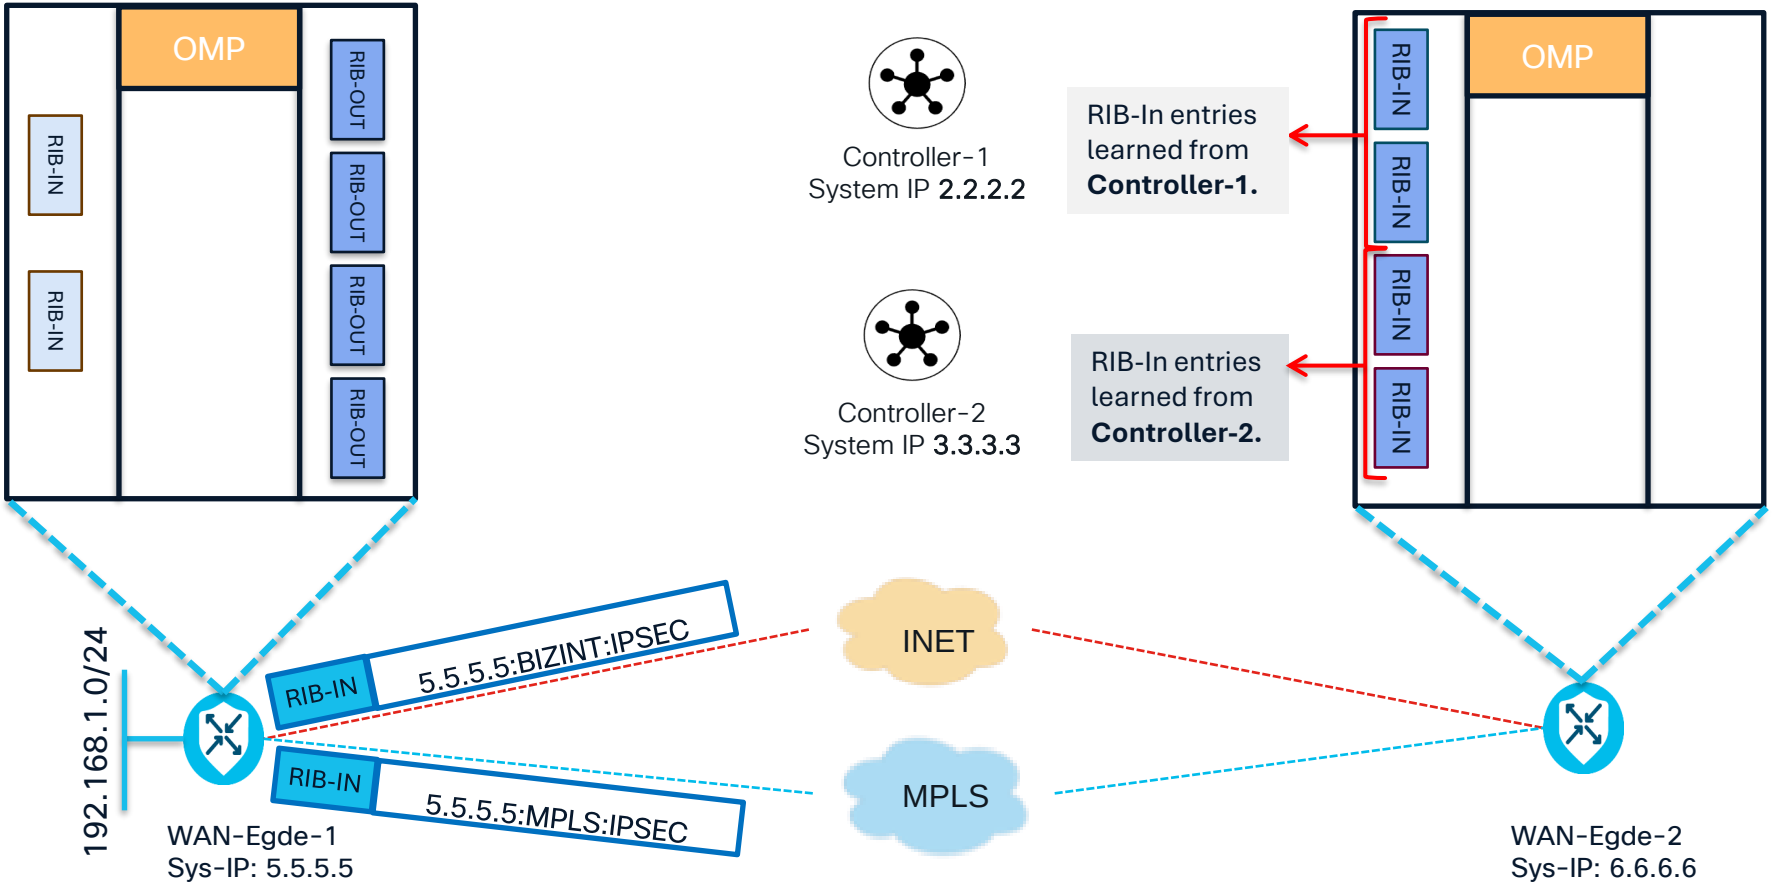






cisco

# OMP vRoute(v4/v6) Advertisement Optimization Using System IP (20.18.2 / 17.18.2 onward)

# Next-Gen OMP RIB-Out Optimization: Enhancement Overview

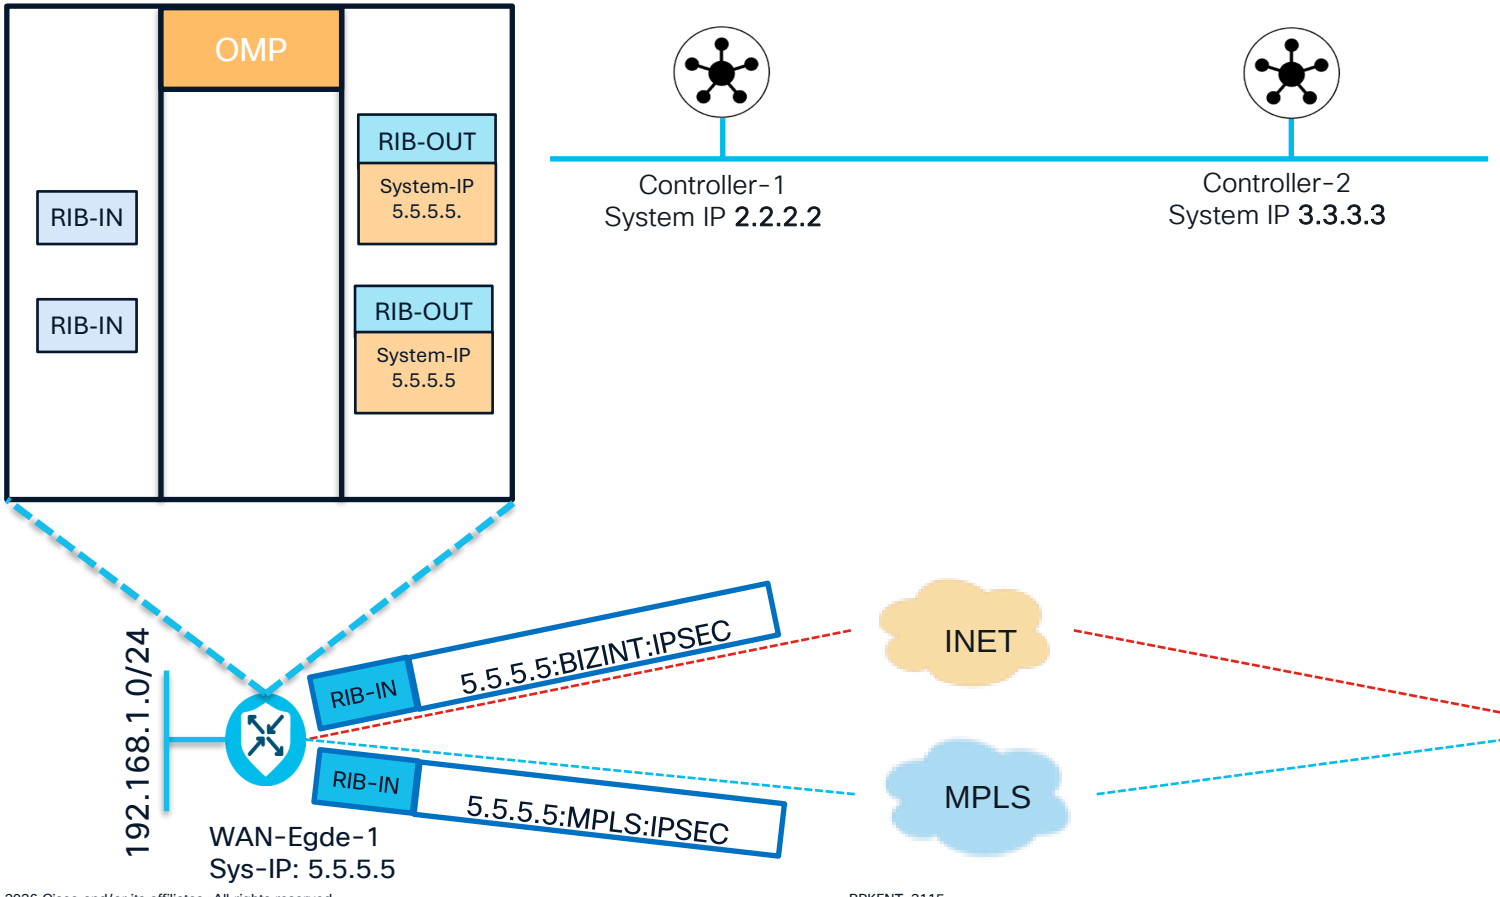

# Next-Gen OMP RIB-Out Optimization: Enhancement Overview

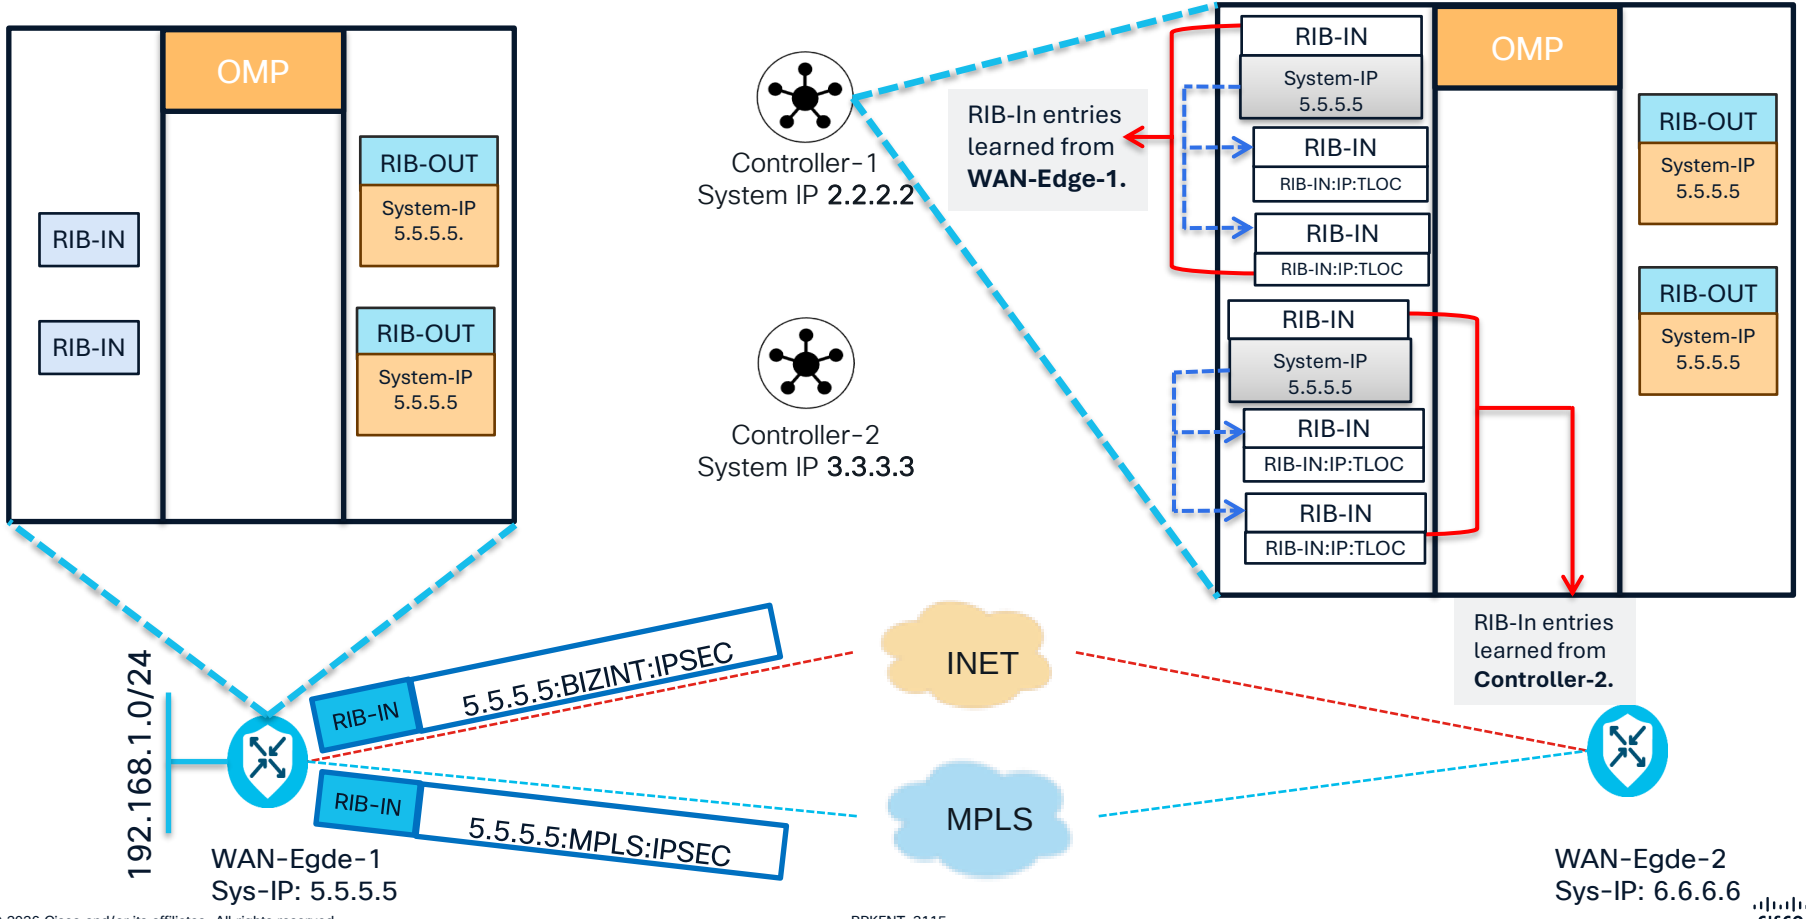

# SDWAN Controller RIB-IN/RIB-OUT Demo

# SD-WAN Demo Topology

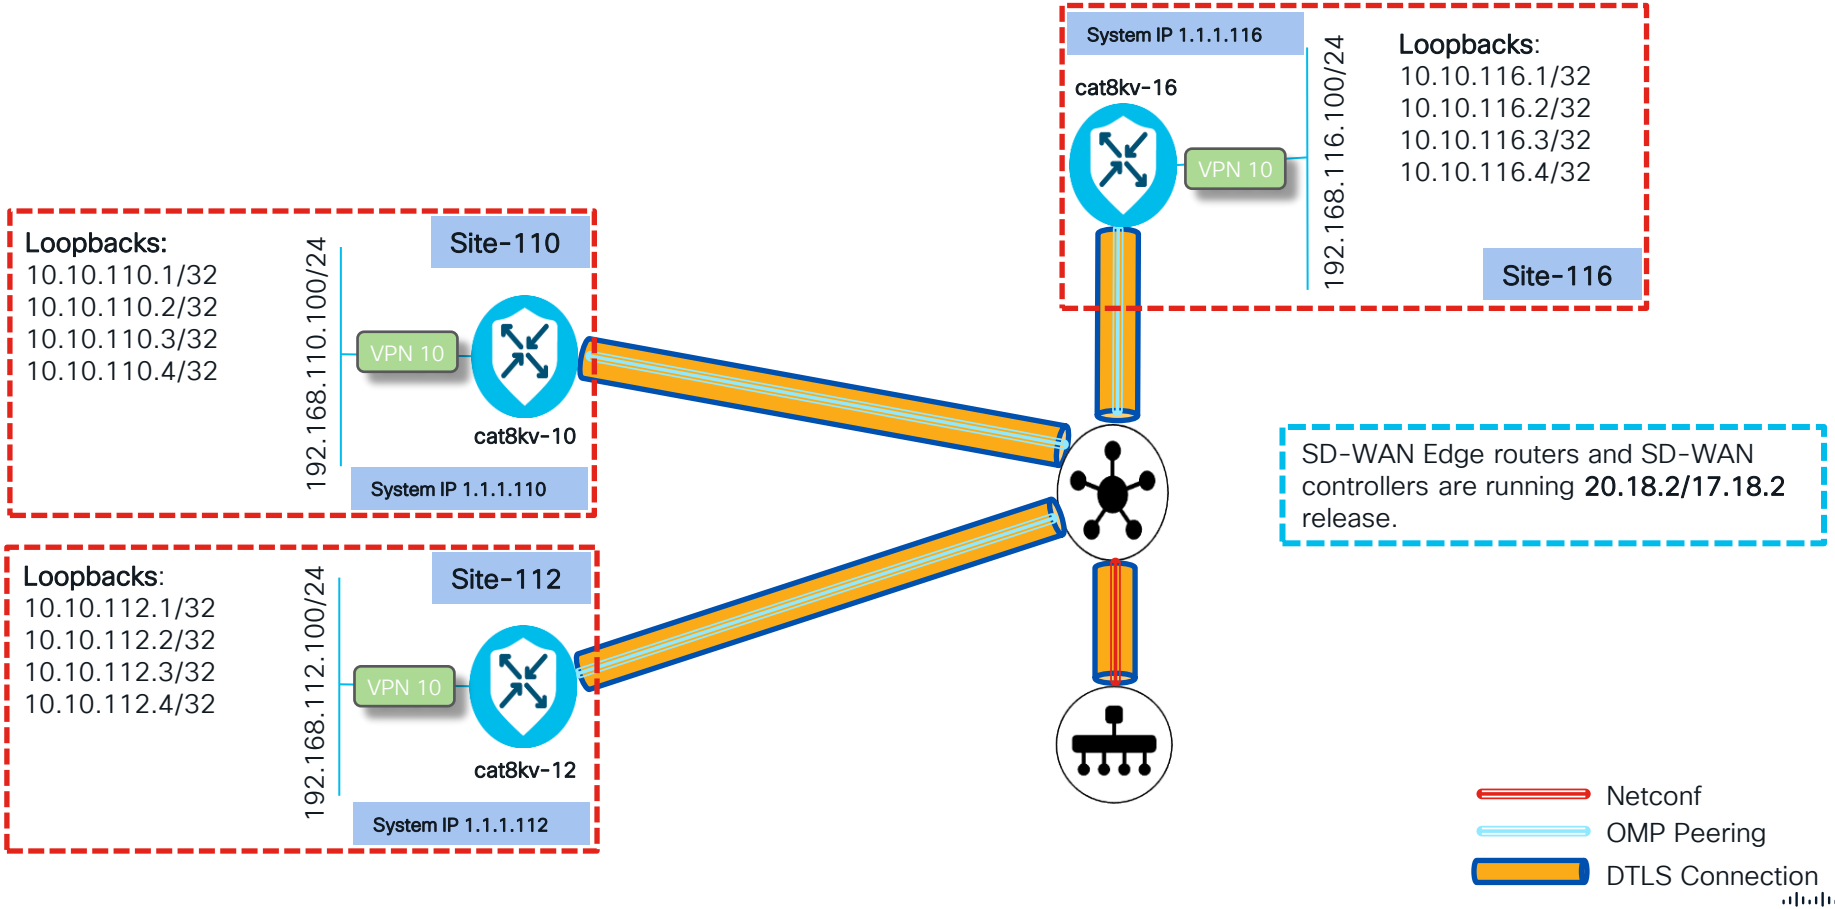

# SD-WAN Demo Topology with Control Policy

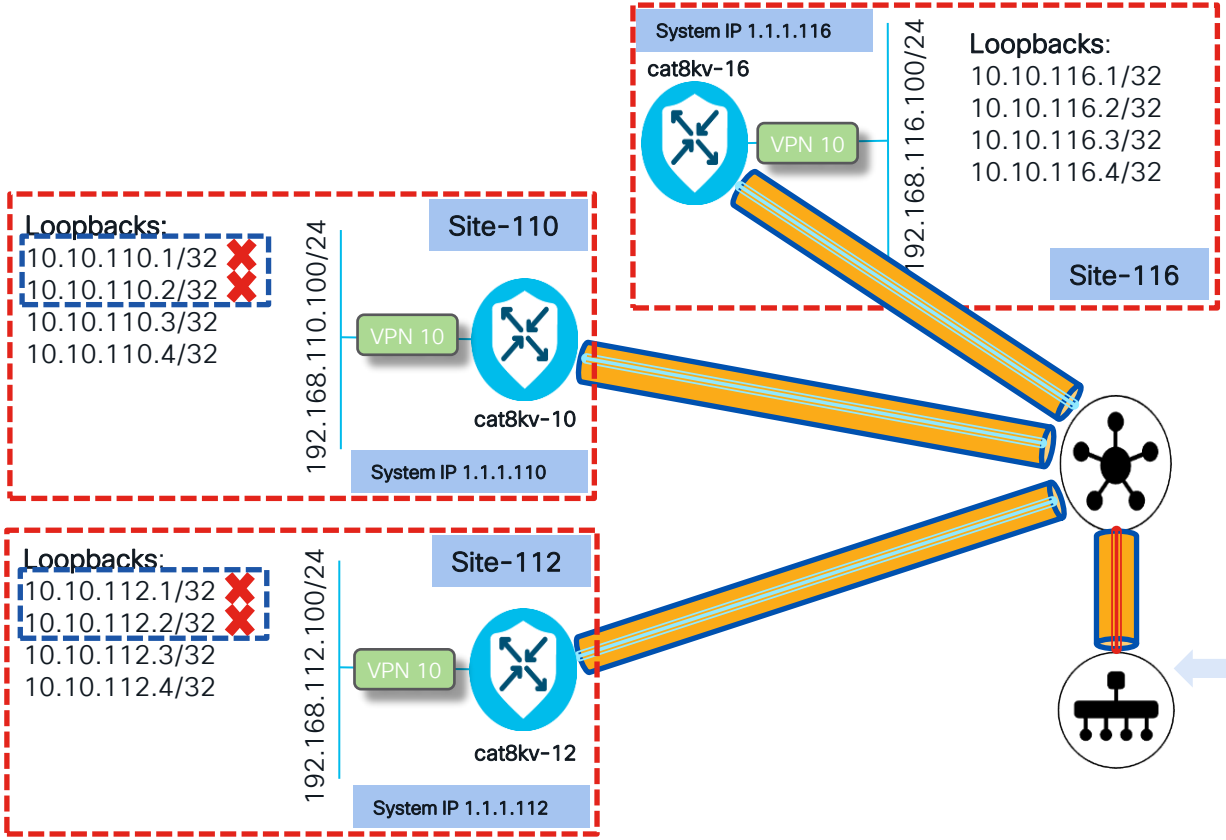


```
control-policy Control-Policy-In
  sequence 10
   match route
    prefix-list Site-110-Prefix-List
    vpn-list    vpn_WAN-Edge-Service-VPN-10
   !
   action reject
   !
  !
  default-action accept
!
control-policy Control-Policy-Out
 sequence 10
  match route
   prefix-list Site-112-Prefix-List
   vpn-list    vpn_WAN-Edge-Service-VPN-10
  !
  action reject
  !
 !
 default-action accept
!
```

Netconf
OMP Peering
DTLS Connection

# Centralized Control Policy

```
!
control-policy Control-Policy-In
  sequence 10
   match route
    prefix-list Site-110-Prefix-List
    vpn-list    vpn_WAN-Edge-Service-VPN-10
   !
   action reject
   !
  !
  default-action accept
 !
 control-policy Control-Policy-Out
  sequence 10
   match route
    prefix-list Site-112-Prefix-List
    vpn-list    vpn_WAN-Edge-Service-VPN-10
   !
   action reject
   !
  !
  default-action accept
 !
```

```
!
lists
  vpn-list vpn_WAN-Edge-Service-VPN-10
   vpn 10
   !
   site-list site-2047270480
    site-id 110
   !
   site-list site1335570788
    site-id 116
   !
   prefix-list Site-110-Prefix-List
    ip-prefix 10.10.110.1/32
    ip-prefix 10.10.110.2/32
   !
   prefix-list Site-112-Prefix-List
    ip-prefix 10.10.112.1/32
    ip-prefix 10.10.112.2/32
   !
apply-policy
 site-list site-2047270480
  control-policy Control-Policy-In in
 !
 site-list site1335570788
  control-policy Control-Policy-Out out
 !
```

# OMP Routes state before applying control policy

```
Controller1# show omp routes vpn 10 received | tab | begin SYS
```

| VPN | PREFIX | FROM PEER | SYS PATH ID | PATH ID | LABEL | STATUS | ATTRIBUTE TYPE | TLOC IP | COLOR | ENCAP | PREFERENCE |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 10 | 10.10.110.1/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |
| 10 | 10.10.110.2/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |
| 10 | 10.10.110.3/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |
| 10 | 10.10.110.4/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |
| 10 | 10.10.112.1/32 | 1.1.1.112 | 0 | 0 | 1003 | R | installed | 1.1.1.112 | - | - | - |
| | | 1.1.1.112 | 0 | 66 | 1003 | C,R | installed | 1.1.1.112 | mpls | ipsec | - |
| | | 1.1.1.112 | 0 | 81 | 1003 | C,R | installed | 1.1.1.112 | private1 | ipsec | - |
| 10 | 10.10.112.2/32 | 1.1.1.112 | 0 | 0 | 1003 | R | installed | 1.1.1.112 | - | - | - |
| | | 1.1.1.112 | 0 | 66 | 1003 | C,R | installed | 1.1.1.112 | mpls | ipsec | - |
| | | 1.1.1.112 | 0 | 81 | 1003 | C,R | installed | 1.1.1.112 | private1 | ipsec | - |
| 10 | 10.10.112.3/32 | 1.1.1.112 | 0 | 0 | 1003 | R | installed | 1.1.1.112 | - | - | - |
| | | 1.1.1.112 | 0 | 66 | 1003 | C,R | installed | 1.1.1.112 | mpls | ipsec | - |
| | | 1.1.1.112 | 0 | 81 | 1003 | C,R | installed | 1.1.1.112 | private1 | ipsec | - |
| 10 | 10.10.112.4/32 | 1.1.1.112 | 0 | 0 | 1003 | R | installed | 1.1.1.112 | - | - | - |
| | | 1.1.1.112 | 0 | 66 | 1003 | C,R | installed | 1.1.1.112 | mpls | ipsec | - |
| | | 1.1.1.112 | 0 | 81 | 1003 | C,R | installed | 1.1.1.112 | private1 | ipsec | - |
| 10 | 10.10.116.1/32 | 1.1.1.116 | 0 | 0 | 1003 | R | installed | 1.1.1.116 | - | - | - |
| | | 1.1.1.116 | 0 | 66 | 1003 | C,R | installed | 1.1.1.116 | mpls | ipsec | - |
| | | 1.1.1.116 | 0 | 81 | 1003 | C,R | installed | 1.1.1.116 | private1 | ipsec | - |
| 10 | 10.10.116.2/32 | 1.1.1.116 | 0 | 0 | 1003 | R | installed | 1.1.1.116 | - | - | - |
| | | 1.1.1.116 | 0 | 66 | 1003 | C,R | installed | 1.1.1.116 | mpls | ipsec | - |
| | | 1.1.1.116 | 0 | 81 | 1003 | C,R | installed | 1.1.1.116 | private1 | ipsec | - |
| 10 | 10.10.116.3/32 | 1.1.1.116 | 0 | 0 | 1003 | R | installed | 1.1.1.116 | - | - | - |
| | | 1.1.1.116 | 0 | 66 | 1003 | C,R | installed | 1.1.1.116 | mpls | ipsec | - |
| | | 1.1.1.116 | 0 | 81 | 1003 | C,R | installed | 1.1.1.116 | private1 | ipsec | - |
| 10 | 10.10.116.4/32 | 1.1.1.116 | 0 | 0 | 1003 | R | installed | 1.1.1.116 | - | - | - |
| | | 1.1.1.116 | 0 | 66 | 1003 | C,R | installed | 1.1.1.116 | mpls | ipsec | - |
| | | 1.1.1.116 | 0 | 81 | 1003 | C,R | installed | 1.1.1.116 | private1 | ipsec | - |

# OMP Routes state before applying control policy

- **SYS-PATH-ID** for a parent System-IP **RIB-IN** is **always** = 0.
- **SYS-PATH-ID** for a child TLOC **RIB-IN = PATH-ID** for the parent System-IP **RIB-IN**.
- **PATH-ID** for a child TLOC **RIB-IN** is locally generated based on an internal hash of **<TLOC-Color, TLOC-Encap>**.
- **SYS-PATH-ID** & **PATH-ID** are used to associate such children **TLOC RIB-INs** with the parent System-IP **RIB-IN**.

```
Controller1# show omp routes vpn 10 received | tab | begin SYS
```

| VPN | PREFIX | FROM PEER | SYS PATH ID | PATH ID | LABEL | STATUS | ATTRIBUTE TYPE | TLOC IP | COLOR | ENCAP | PREFERENCE |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 10 | 10.10.110.1/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |
| 10 | 10.10.110.2/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |
| 10 | 10.10.110.3/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |
| 10 | 10.10.110.4/32 | 1.1.1.110 | 0 | 0 | 1003 | R | installed | 1.1.1.110 | - | - | - |
| | | 1.1.1.110 | 0 | 66 | 1003 | C,R | installed | 1.1.1.110 | mpls | ipsec | - |
| | | 1.1.1.110 | 0 | 81 | 1003 | C,R | installed | 1.1.1.110 | private1 | ipsec | - |

# OMP Routes state after applying Inbound control policy

```
Controller1# show omp routes vpn 10 received | tab | begin SYS
                                        SYS
                                        PATH    PATH                                ATTRIBUTE
VPN     PREFIX              FROM PEER   ID      ID      LABEL   STATUS      TYPE    TLOC IP         COLOR           ENCAP   PREFERENCE
-------------------------------------------------------------------------------------------------------------------------------------
10      10.10.110.1/32      1.1.1.110   0       0       1003    Rej,R,Inv   installed 1.1.1.110     -               -       -
                            1.1.1.110   0       66      1003    Rej,R,Inv   installed 1.1.1.110     mpls            ipsec   -
                            1.1.1.110   0       81      1003    Rej,R,Inv   installed 1.1.1.110     private1        ipsec   -
10      10.10.110.2/32      1.1.1.110   0       0       1003    Rej,R,Inv   installed 1.1.1.110     -               -       -
                            1.1.1.110   0       66      1003    Rej,R,Inv   installed 1.1.1.110     mpls            ipsec   -
                            1.1.1.110   0       81      1003    Rej,R,Inv   installed 1.1.1.110     private1        ipsec   -
10      10.10.110.3/32      1.1.1.110   0       0       1003    R           installed 1.1.1.110     -               -       -
                            1.1.1.110   0       66      1003    C,R         installed 1.1.1.110     mpls            ipsec   -
                            1.1.1.110   0       81      1003    C,R         installed 1.1.1.110     private1        ipsec   -
10      10.10.110.4/32      1.1.1.110   0       0       1003    R           installed 1.1.1.110     -               -       -
                            1.1.1.110   0       66      1003    C,R         installed 1.1.1.110     mpls            ipsec   -
                            1.1.1.110   0       81      1003    C,R         installed 1.1.1.110     private1        ipsec   -
10      10.10.112.1/32      1.1.1.112   0       0       1003    R           installed 1.1.1.112     -               -       -
                            1.1.1.112   0       66      1003    C,R         installed 1.1.1.112     mpls            ipsec   -
                            1.1.1.112   0       81      1003    C,R         installed 1.1.1.112     private1        ipsec   -
10      10.10.112.2/32      1.1.1.112   0       0       1003    R           installed 1.1.1.112     -               -       -
                            1.1.1.112   0       66      1003    C,R         installed 1.1.1.112     mpls            ipsec   -
                            1.1.1.112   0       81      1003    C,R         installed 1.1.1.112     private1        ipsec   -
10      10.10.112.3/32      1.1.1.112   0       0       1003    R           installed 1.1.1.112     -               -       -
                            1.1.1.112   0       66      1003    C,R         installed 1.1.1.112     mpls            ipsec   -
                            1.1.1.112   0       81      1003    C,R         installed 1.1.1.112     private1        ipsec   -
10      10.10.112.4/32      1.1.1.112   0       0       1003    R           installed 1.1.1.112     -               -       -
                            1.1.1.112   0       66      1003    C,R         installed 1.1.1.112     mpls            ipsec   -
                            1.1.1.112   0       81      1003    C,R         installed 1.1.1.112     private1        ipsec   -
10      10.10.116.1/32      1.1.1.116   0       0       1003    R           installed 1.1.1.116     -               -       -
                            1.1.1.116   0       66      1003    C,R         installed 1.1.1.116     mpls            ipsec   -
                            1.1.1.116   0       81      1003    C,R         installed 1.1.1.116     private1        ipsec   -
10      10.10.116.2/32      1.1.1.116   0       0       1003    R           installed 1.1.1.116     -               -       -
                            1.1.1.116   0       66      1003    C,R         installed 1.1.1.116     mpls            ipsec   -
                            1.1.1.116   0       81      1003    C,R         installed 1.1.1.116     private1        ipsec   -
10      10.10.116.3/32      1.1.1.116   0       0       1003    R           installed 1.1.1.116     -               -       -
                            1.1.1.116   0       66      1003    C,R         installed 1.1.1.116     mpls            ipsec   -
                            1.1.1.116   0       81      1003    C,R         installed 1.1.1.116     private1        ipsec   -
10      10.10.116.4/32      1.1.1.116   0       0       1003    R           installed 1.1.1.116     -               -       -
                            1.1.1.116   0       66      1003    C,R         installed 1.1.1.116     mpls            ipsec   -
                            1.1.1.116   0       81      1003    C,R         installed 1.1.1.116     private1        ipsec   -
```

# OMP Routes state after applying Inbound control policy

```
Controller1# show omp routes vpn 10 advertised | tab | begin VPN
VPN    PREFIX                  TO PEER
-----------------------------------------------
10     10.10.110.3/32          1.1.1.112
                               1.1.1.114
                               1.1.1.116
10     10.10.110.4/32          1.1.1.112
                               1.1.1.114
                               1.1.1.116
```

# OMP Routes state after applying Inbound control policy

```
Controller1# show support omp rib vroute 10:10.10.110.1/32
Looking up vroute 10.10.110.1/32 in 10
RIB-Entry: ROUTE-IPV4 Flags: (0x0) , recv-attr-count 3, adv-attr-count 0
VPN-ID: 10, Prefix: 10.10.110.1/32
Region-Info Region-ID: 65534
RIB-IN: Peer: 1.1.1.110, ID: 0x3, updated: Sun Jan 18 11:50:03 2026
Path-id: 66, Sys-Path-id: 0, Label: 1003 BUM-Label: 0 Affinity Number: 0 TLOC-pref: 0 TLOC-stale: 0 version: 1
Lost-to-peer: ::, Lost-to-path-id: 0, Lost-to-sys-path-id: 0, Loss-Reason: None(0)
Flags: (0x8028) REJECT RESOLVED EXPANDED
Attribute: ROUTE-IPV4, Length: 1208, Ref: 8
Distance: 0, Site-ID: 110, Carrier: 0, Query: 0, Gen-ID: 0x0, Border: 0 Overlay: 1 Site-Type: 0 0 0 0
Originator: 1.1.1.110
RIB-IN: Peer: 1.1.1.110, ID: 0x2, updated: Sun Jan 18 11:50:03 2026
Path-id: 81, Sys-Path-id: 0, Label: 1003 BUM-Label: 0 Affinity Number: 0 TLOC-pref: 0 TLOC-stale: 0 version: 1
Lost-to-peer: 1.1.1.110, Lost-to-path-id: 66, Lost-to-sys-path-id: 0, Loss-Reason: TLOC ID(10)
Flags: (0x8028) REJECT RESOLVED EXPANDED
Distance: 0, Site-ID: 110, Carrier: 0, Query: 0, Gen-ID: 0x0, Border: 0 Overlay: 1 Site-Type: 0 0 0 0
Originator: 1.1.1.110
TLOC: 1.1.1.110 : private1 : ipsec
SysIP RIB-IN-LIST:
RIB-IN: Peer: 1.1.1.110, ID: 0x1, updated: Sun Jan 18 11:50:03 2026
Path-id: 0, Sys-Path-id: 0, Label: 1003 BUM-Label: 0 Affinity Number: 0 TLOC-pref: 0 TLOC-stale: 0 version: 0
Lost-to-peer: ::, Lost-to-path-id: 0, Lost-to-sys-path-id: 0, Loss-Reason: None(0)
Flags: (0x28) REJECT RESOLVED
Attribute: ROUTE-IPV4, Length: 1208, Ref: 8
Distance: 0, Site-ID: 110, Carrier: 0, Query: 0, Gen-ID: 0x0, Border: 0 Overlay: 1 Site-Type: 0 0 0 0
Originator: 1.1.1.110
Origin: Protocol: ospf[4], Sub-Type: intra-area[1], Metric: 11
System-IP: ((nil)) 1.1.1.110
```

# OMP Routes state after applying Inbound control policy

```
Controller1# show support omp rib vroute 10:10.10.110.3/32 | include RIB-OUT\|RIB-IN
  RIB-IN:  Peer: 1.1.1.110, ID: 0x6, updated: Sun Jan 18 11:50:03 2026
  RIB-IN:  Peer: 1.1.1.110, ID: 0x5, updated: Sun Jan 18 11:50:03 2026
SysIP RIB-IN-LIST:
  RIB-IN:  Peer: 1.1.1.110, ID: 0x4, updated: Sun Jan 18 11:50:03 2026
  RIB-OUT: RI-ID: 0x101016e0101016e, Peer: 1.1.1.112 Path-id: 3, Label: 1003, Flags: (0x1) ADV  Common Ref: 3
  RIB-OUT: RI-ID: 0x101016e0101016e, Peer: 1.1.1.114 Path-id: 3, Label: 1003, Flags: (0x1) ADV  Common Ref: 3
  RIB-OUT: RI-ID: 0x101016e0101016e, Peer: 1.1.1.116 Path-id: 3, Label: 1003, Flags: (0x1) ADV  Common Ref: 3
```

# OMP Routes state after applying Outbound control policy

```
Controller1# show omp routes vpn 10 advertised | begin VPN
VPN    PREFIX                TO PEER
-----------------------------------------
10     10.10.110.3/32        1.1.1.112
                             1.1.1.114
                             1.1.1.116
10     10.10.110.4/32        1.1.1.112
                             1.1.1.114
                             1.1.1.116
10     10.10.112.1/32        1.1.1.110
                             1.1.1.114
10     10.10.112.2/32        1.1.1.110
                             1.1.1.114
10     10.10.112.3/32        1.1.1.110
                             1.1.1.114
                             1.1.1.116
10     10.10.112.4/32        1.1.1.110
                             1.1.1.114
                             1.1.1.116
```

# OMP Routes state after applying Outbound control policy

```
Controller1# show support omp rib vroute 10:10.10.112.1/32 rib-out-peer-ip 1.1.1.112 | include RIB
RIB-Entry: ROUTE-IPV4 Flags: (0x0) , recv-attr-count 3, adv-attr-count 2
  RIB-IN:  Peer: 1.1.1.112, ID: 0x3, updated: Tue Jan 13 22:14:10 2026
  RIB-IN:  Peer: 1.1.1.112, ID: 0x2, updated: Tue Jan 13 22:14:10 2026
SysIP RIB-IN-LIST:
  RIB-IN:  Peer: 1.1.1.112, ID: 0x1, updated: Tue Jan 13 22:14:10 2026
  RIB-OUT: RI-ID: 0x101017001010170, Peer: 1.1.1.110 Path-id: 3, Label: 1003, Flags: (0x1) ADV  Common Ref: 2
  RIB-OUT: RI-ID: 0x101017001010170, Peer: 1.1.1.114 Path-id: 3, Label: 1003, Flags: (0x1) ADV  Common Ref: 2
```

# OMP and BGP Packet Insights

# Commonalities between OMP and BGP Packets

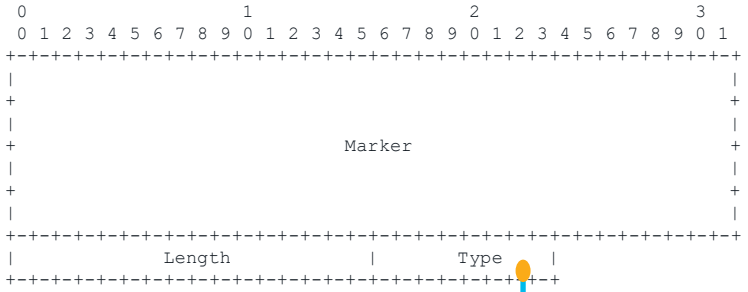


BGP Message Header Format

- OPEN
- UPDATE
- NOTIFICATION
- KEEPALIVE

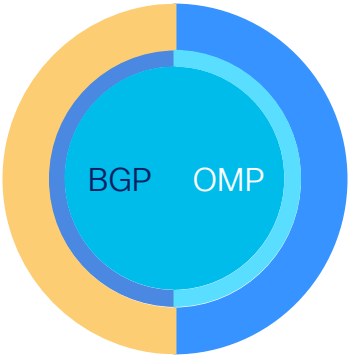


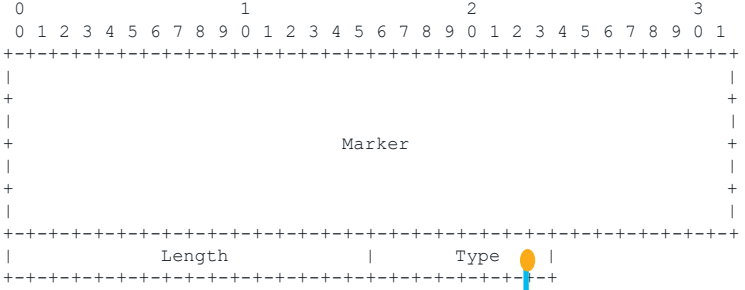


OMP Message Header Format

- HANDSHAKE
- UPDATE
- ALERT
- HELLO
- QUERY
- POLICY
- INFORM

# Commonalities between OMP and BGP Packets

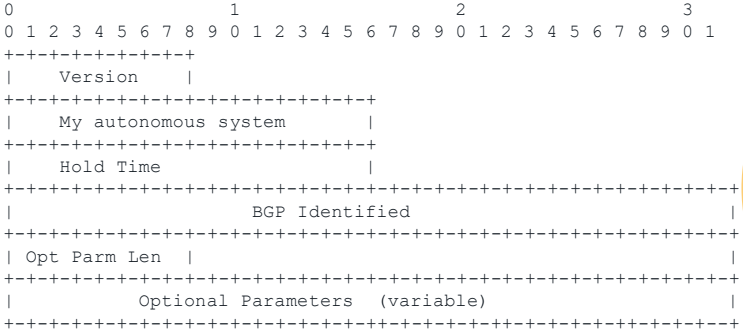


BGP Open Packet Format

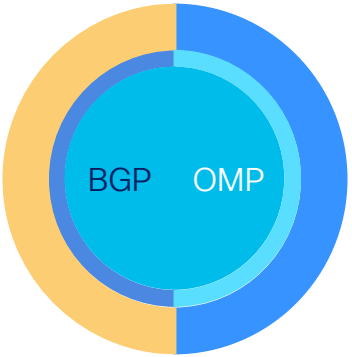


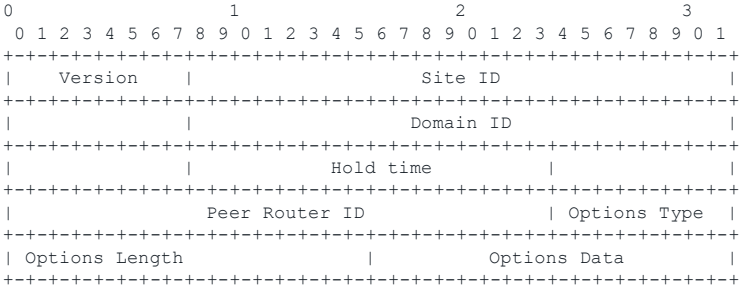


OMP Handshake Packet Format

# Commonalities between OMP and BGP Packets

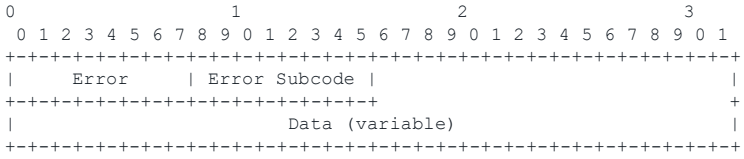


BGP Notification Packet Format

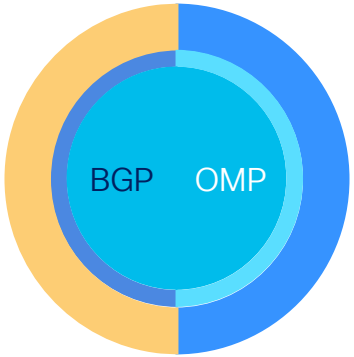


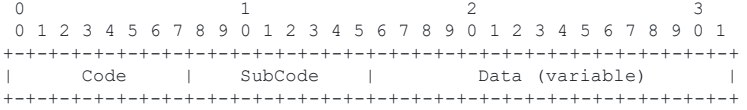


OMP Alert Packet Format

# OMP Policy Packet Format

- SD-WAN policies are send in OMP Policy Packet.
- Maximum Policy packet size is 4000 bytes and if policy is huge then policy will be send in fragments.
- If policy data is less than 3980 (4000-20 TCP header) bytes, then Type field is set with Policy Complete.
- If policy data is more than 3980 bytes then it will be send in fragments with Type field is set to Policy Start, when first packet is send and with last fragment Type field is set to Policy End.

# Walkthrough of OMP Session establishment at the WAN Edge Process level

# SD-WAN OMP Discovery

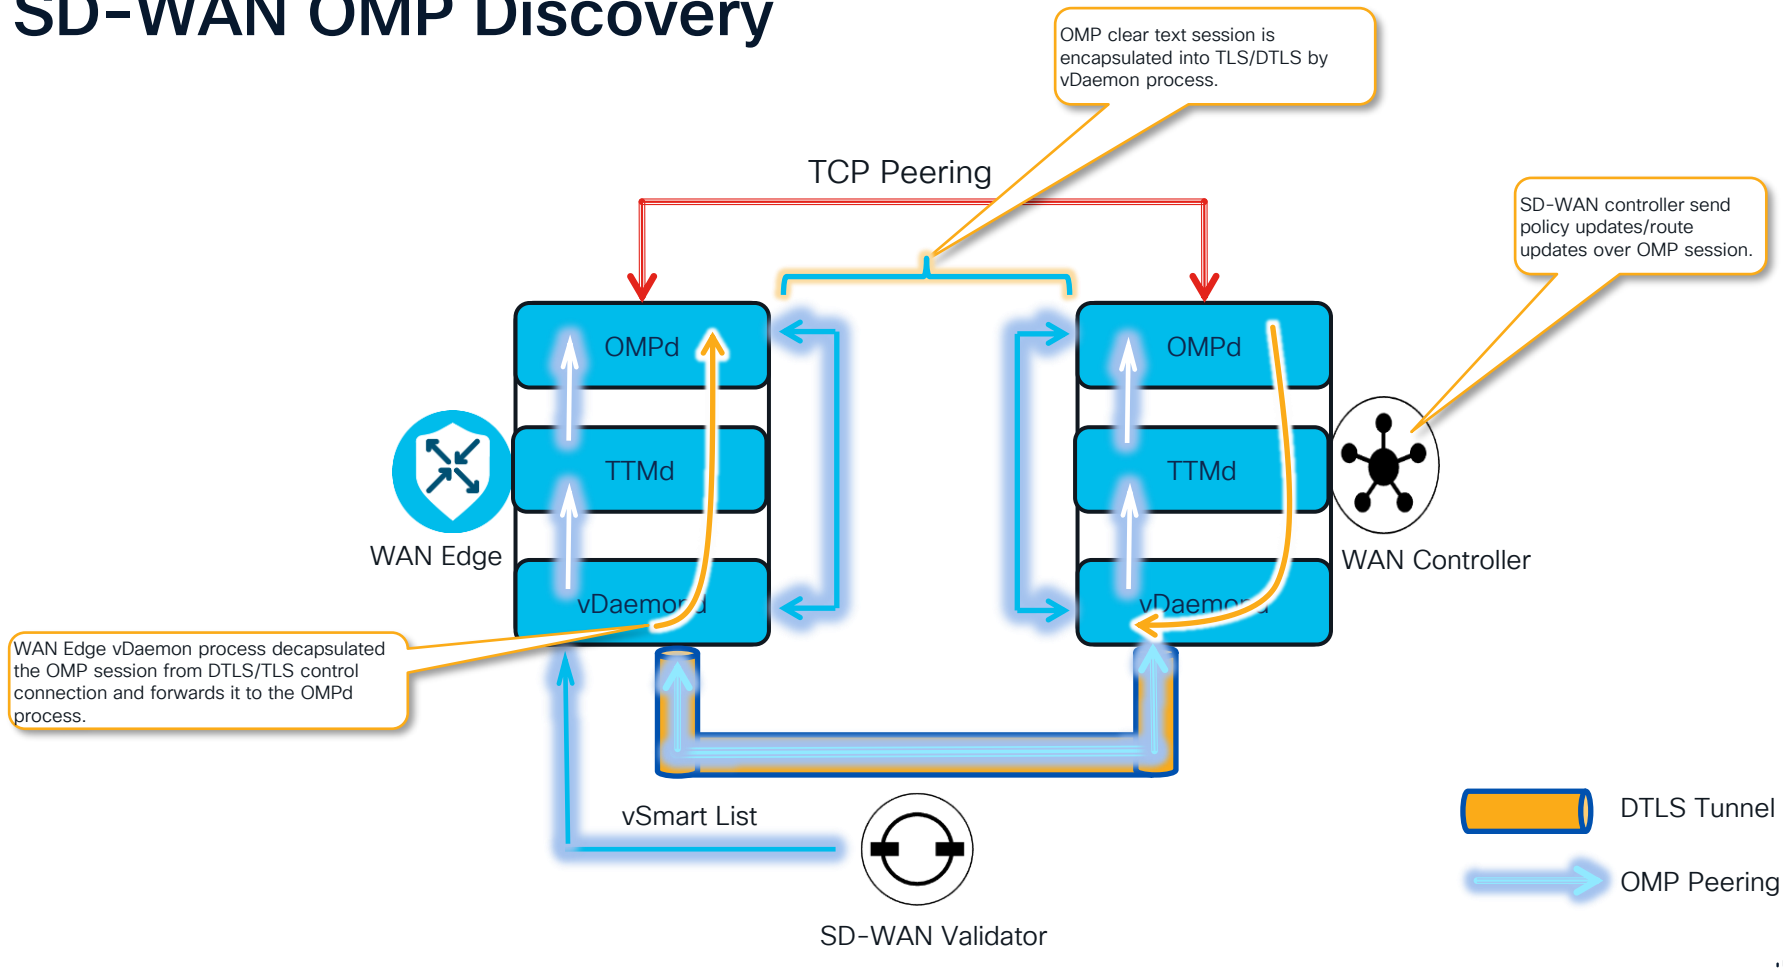

# Gaining Visibility into OMP: Debugs on Controllers and WAN Edges

# How can I begin navigating through the OMP logs? 1/3

- From SD-WAN Edge perspective:
  - We can enable OMP debugging by using flags to filter peer OMP negotiation or OMP packets.
    - `[no]debug platform software sdwan omp packets "direction both peer-address <PEER-IP-ADDRESS> packet-type all"`

SD-WAN debug information does not appear in the output of the '**show debugging**' command.

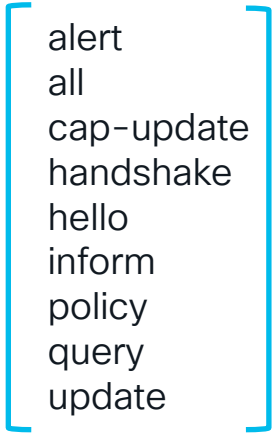


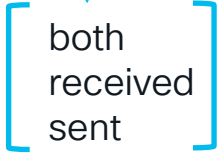

# How can I begin navigating through the OMP logs? 2/3

- From SD-WAN Edge perspective:
  - Starting from the 17.12 release, we need to enable traces for OMP process along with the module name and then we can view the messages exchanged between SD-WAN controller and SD-WAN Edge devices.
    - `set platform software trace ompd RP active ompd-pkt verbose`
    - `set platform software trace ompd RP active ompd-oper verbose`
    - `set platform software trace ompd RP active ompd-event verbose`
  - By default, all modules for OMP processes are set to the "Notice" level.
  - We can verify with the following command if OMP traces are enabled or not.
    - `show platform software trace level ompd rp active`
  - The following command will exhibit what flags are enabled for OMP debugging.
    - `show platform software sdwan omp debug`
  - To view the OMP debug messages:
    - `show logging process ompd internal`

# How can I begin navigating through the OMP logs? 3/3

- From SD-WAN controller perspective:
  - The following debug can be enabled to view the OMP packet exchange between SD-WAN controllers and WAN-Edges.

```
Controller-1# debug omp
Possible completions:
  best-path          Debug OMP best-path calculation
  cxp                Debug OMP cloudexpress
  events             Debug OMP events
  graceful-restart   Debug OMP Graceful Restart
  identity           Debug OMP Identity
  ipcs               Debug OMP IPCs
  packets            Debug OMP packets
  policy             Debug OMP policy
```

  - All the debug logs are written by the SD-WAN controller in the following files.

```
Controller-1# vshell
Controller-1:~$ tail -f /var/log/vdebug
```

```
Controller-1# vshell
Controller-1:~$ tail -f /var/log/tmplog/vdebug
```

# My OMP Peering is down, where to start?

# What steps should be taken if OMP peering fails to establish?

- Check if Control Connection (**CC**) is up between SD-WAN Edge and SD-WAN controller?
- If NOT, then we need to troubleshoot first, why CC is not coming up?

- If CC is UP but still OMP peering is not coming UP, then verify that **system-ip** is **NOT** overlapping in SD-WAN overlay.
- If OMP peering remains in the **Init** state, verify that the **tun-tap** interface route has been installed by the **vDaemon** process by running "*route -n*" from the Controller's shell.

```
Controller01# vshell
Controller01:~$ route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
4.4.4.4         127.0.1.254     255.255.255.255 UGH   0      0        0 tun_0_0
172.16.0.0      0.0.0.0         255.255.255.0   U     0      0        0 eth1
172.16.0.0      172.16.0.254    255.255.0.0     UG    0      0        0 eth1
```

# What steps should be taken if OMP peering fails to establish?

- Check if Control Connection (**CC**) is up between SD-WAN Edge and SD-WAN controller?
- If NOT, then we need to troubleshoot first, why CC is not coming up?

- If CC is UP but still OMP peering is not coming UP, then verify that **system-ip** is **NOT** overlapping in SD-WAN overlay.
- If OMP peering remains in the **Init** state, verify that the **tun-tap** interface route has been installed by the **vDaemon** process by running "*route -n*" from the Controller's shell.

- OMPd process crashed that cause OMP peering failed to establish.
- We can check the status if OMPd process crash/rebooted with the following commands
  - On SD-WAN Controller:
    - `show system status`
    - `show support omp peer peer-ip` **`<PEER-IP>`**
  - On SD-WAN Router:
    - `show platform software sdwan omp peer` **`<PEER-IP>`**
    - `show sdwan crash`
- Memory consumption/leakage can also cause OMP peering issue and can be seen with the following commands:
    - `[`**`SD-WAN Edge`**`] show platform software process memory r0 all sorted | include ompd`
    - `[`**`SD-WAN Controller`**`] show support omp memory-statistics`

# Interesting SHOW commands for OMP debugging !!!


- SD-WAN Controller
  - `show omp peers`
  - `show omp peers <PEER-IP> detail`
  - `show omp summary`
  - `show support omp context`
  - `show support omp daemon`
  - `show support policy route-policy`
  - `show support omp peer peer-ip <PEER-IP>`
  - `show support omp rib vroute <VPN-ID:Prefix/Length>`
  - `show support omp sysip-tlocs`
  - `show support omp memory-statistics`
  - `show support omp label-db`
- SD-WAN Edge
  - `show sdwan omp summary`
  - `show sdwan omp peers`
  - `show platform software sdwan omp omp daemon`
  - `show platform software sdwan omp peer`
  - `show platform software sdwan omp peer-ip <PEER-IP>`
  - `show platform software sdwan omp rib vroute vpn <VPN> <Prefix/Length>`
  - `show platform software sdwan omp rib tloc <TLOC-IP> <COLOR> <Encap>`
  - `show platform software sdwan omp sysip-tlocs`
  - `show platform software sdwan omp memory`

# What is OMP Graceful Restart?

Once control connections (CC) are **down**, and **OMP Hold timer expires**, OMP peering moves to Graceful Restart(GR) state.

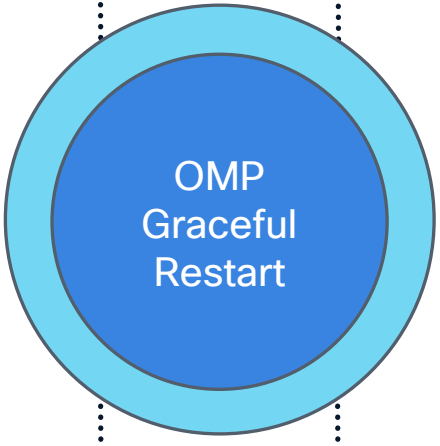


Graceful Restart(GR) timers are configured **independently** per OMP peering endpoint.

When OMP peering is Up, routes are refreshed and marked **non-stale** while the GR timer runs until **EOR** is received.

EOR is sent per AF after route refresh; stale routes are removed ONLY after **EOR** is received **for all AFs**.

# OMP Graceful Restart - Off

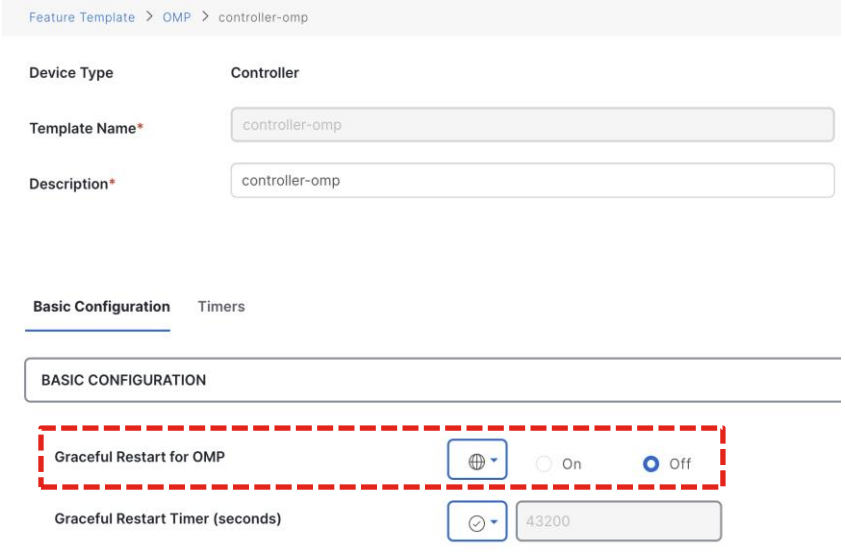


- Graceful Restart for OMP = Off at all SD-WAN Controllers means Graceful Restart is disabled on all SD-WAN Edges.
- OMP Routing Table and BFD Session are lost immediately when SD-WAN Edges loses all OMP peers.

```
Controller1# show running-config omp
omp
 no shutdown
 filter-route
  no outbound affinity-group-preference
  no outbound tloc-color
 exit
 no graceful-restart
 outbound-policy-caching
!
```

```
Controller2# show running-config omp
omp
 no shutdown
 filter-route
  no outbound affinity-group-preference
  no outbound tloc-color
 exit
 no graceful-restart
 outbound-policy-caching
!
```

```
cat8kv-10#show sdwan omp peers detail | include graceful|peer|state|type
peer                    2.2.2.2
type                    vsmart
state                   up
graceful-restart        not-supported
graceful-restart-interval
peer                    3.3.3.3
type                    vsmart
state                   up
graceful-restart        not-supported
graceful-restart-interval
peer                    5.5.5.5
type                    vsmart
state                   up
graceful-restart        not-supported
graceful-restart-interval
```

```
cat8kv-10# !! -- #### shutdown Controller1, Controller2 and Controller3 ####
-- !!
Site400-cE1#
2026/01/26 08:45:20.988025744 {iosrp_R0-0}{255}: [iosrp] [17261]: (info):
*Jan 26 08:45:20.988: %Cisco-SDWAN-cat8kv-10-OMPD-6-INFO-1400002:
Notification: 2026/01/26 08:45:20 omp-number-of-vsmarts-change severity-
level:major host-name:"cat8kv-10" system-ip:1.1.1.110 tenant-
name:"[Default]" tenant-global-id:0 number-of-vsmarts:0
cat8kv-10#show sdwan bfd sessions
cat8kv-10#
```

# OMP Graceful Restart Timer

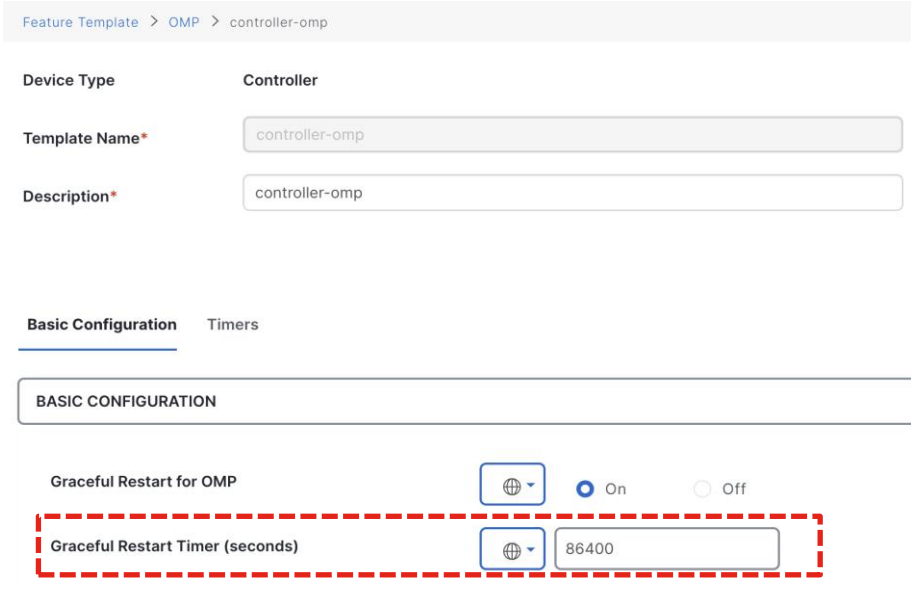


- Graceful Restart timer configured on SD-WAN controllers is applied to SD-WAN Edge, and conversely.
- If any change to an OMP graceful restart configuration is made, the OMP session between the Cisco SD-WAN controllers and the device is flapped.

```
Controller1# show running-config omp
omp
 no shutdown
 filter-route
  no outbound affinity-group-preference
  no outbound tloc-color
 exit
 graceful-restart
 outbound-policy-caching
 timers
  graceful-restart-timer 86400
 exit
!
```

```
Controller1#show omp peers 1.1.1.110 detail | include
peer\|state\|graceful
peer                          1.1.1.110
state                         up
graceful-restart              supported
graceful-restart-interval 43200
```

```
cat8kv-10#show sdwan omp peers detail | include
peer|state|graceful
peer                          2.2.2.2
state                         up
graceful-restart              supported
graceful-restart-interval 86400
peer                          3.3.3.3
state                         up
graceful-restart              supported
graceful-restart-interval 86400
```

# Why OMP Routes Don’t Loop?

# How does OMP ensure loop avoidance?

- OMP loop avoidance is based on originator System-IP.
- Native built-in loop prevention mechanisims when OMP interacts with OSPF, EIGRP, RIP and BGP.

**OSPF**

- OSPF uses "Down Bit" (RFC 4577). When redistributed form OMP into OSPF on WAN Edge, it is set.
- When LSA distributed through sevice side network gets to the other WAN Edge, as DN bit is set so route is not installed into RIB on WAN Edge as SDWAN-Dnbit flag is set.

**EIGRP**

- EIGRP uses "**External Protocol**" ID field. It is set to a value of "**OMP-Agent**".
- When the other WAN Edge on the same site receives such update, it installs the route into the EIGRP topology table, sets "**SDWAN-Down**" flag and then install into the RIB with **Administrative Distance (AD)** to **252**. This, in turn, makes OMP the preferred route because it has an **AD** of **251**.

**RIP**

- RIP uses "**OMP-ROUTE-TAG**" with value **44270,** which is **not configurable**.
- When the other WAN Edge on the same site receives RIP update, it is not get installed because of "**OMP-ROUTE-TAG**" tag. This route will get installed with **AD** of **252**, ONLY when OMP route gets withdrawn, thus avoiding routing loop.

# How does OMP ensure loop avoidance?

- OMP loop avoidance is based on originator System-IP.
- Native built-in loop prevention mechanisims when OMP interacts with OSPF, EIGRP, RIP and BGP.

BGP

- BGP uses **SoO**, extended community which value is set to the OMP **site ID**.
- When the other WAN Edge receives the BGP update from the service-side network and there **SoO** community matches its **own site ID**, then route will **not be installed into RIB**.
- BGP peers at site must send **BGP extended communities** and have the **same site ID**.

# How Down-bit (DN) works? 1/2

```
route-map OSPF-Route-Policy permit 1
 set metric 100
 set metric-type type-1
 match ip address prefix-list OSPF-Prefix
!
route-map OSPF-Route-Policy permit 65535
!
ip prefix-list OSPF-Prefix seq 5 permit 192.168.100.100/32
!
router ospf 10 vrf 10
 auto-cost reference-bandwidth 100
 compatible rfc1583
 distance ospf intra-area 110 inter-area 110 external 110
 no local-rib-criteria
 max-lsa 50000
 redistribute maximum-prefix 10240
 redistribute omp
 redistribute omp route-map OSPF-Route-Policy
 table-map OSPF-Route-Policy filter
 timers throttle spf 200 1000 10000
!
```

```
cEdge-3#show ip ospf rib 192.168.100.100 255.255.255.255

            OSPF Router with ID (10.120.10.1) (Process ID 10)


                                Base Topology (MTID 0)

OSPF local RIB
Codes: * - Best, > - Installed in global RIB
LSA: type/LSID/originator

*   192.168.100.100/32, Ext1, cost 103, fwd cost 3, tag 0
     SPF Instance 46, age 00:08:19
      contributing LSA: 5/192.168.100.100/10.110.10.1
     Flags: HiPrio, SDWAN-DNbit
      via 10.120.10.2, GigabitEthernet3, label 1048578, strict label
1048578
       Flags: none
       LSA: 5/192.168.100.100/10.110.10.1
       Source: 10.110.10.1 (area 0)
```

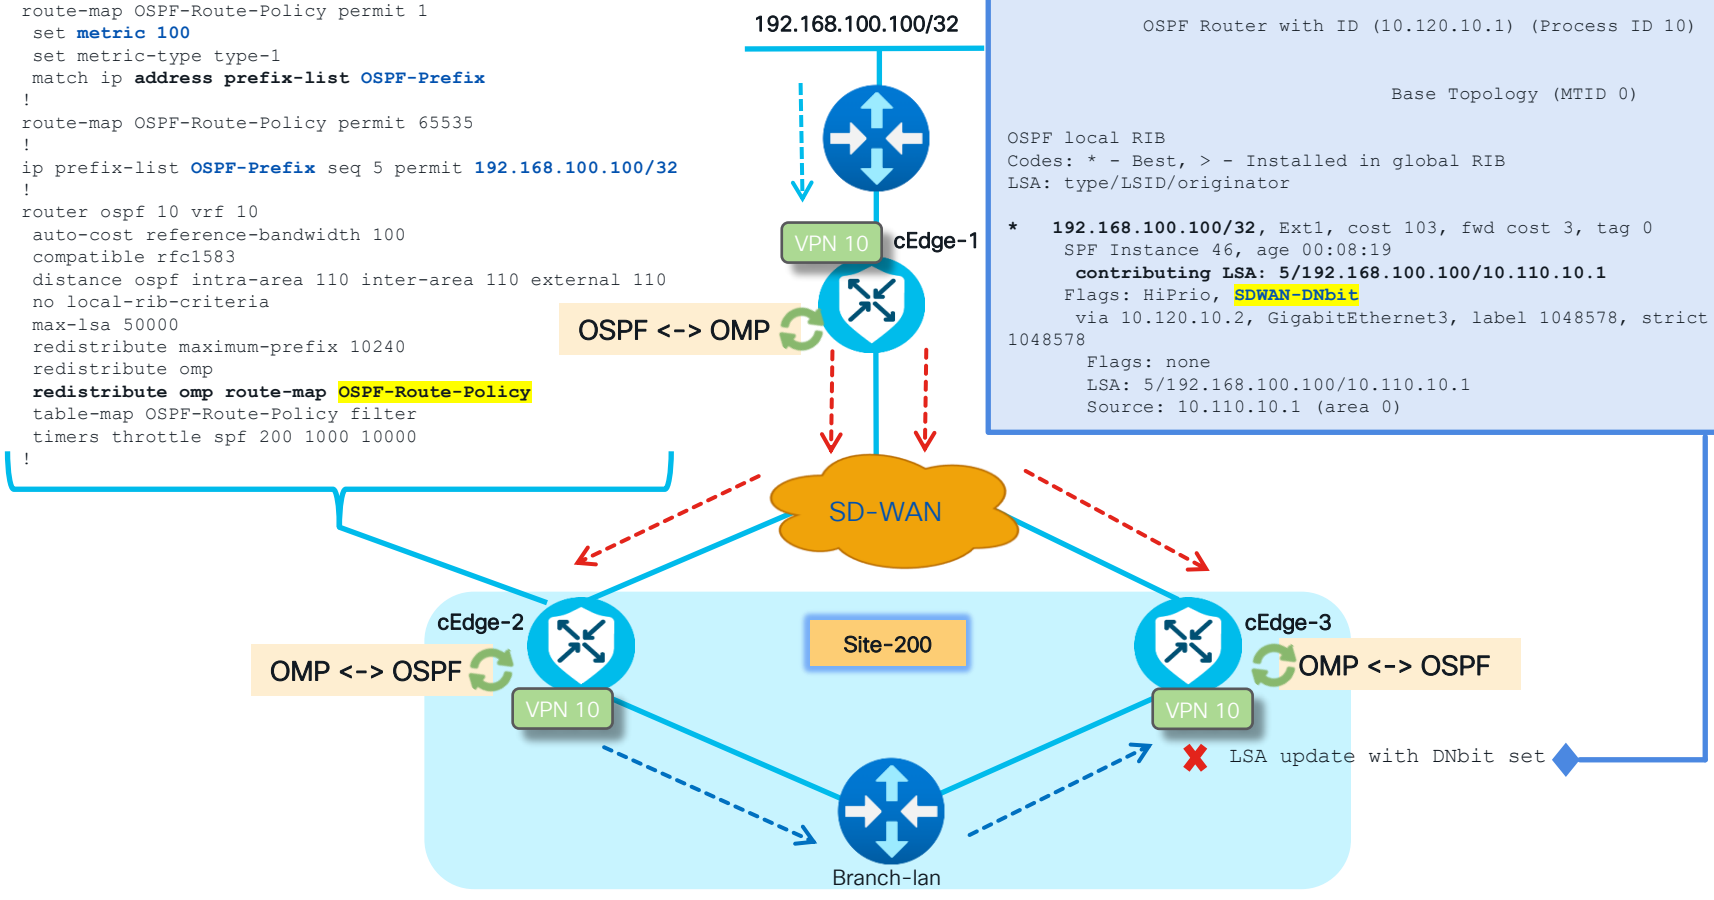

# How Down-bit (DN) works? 2/2

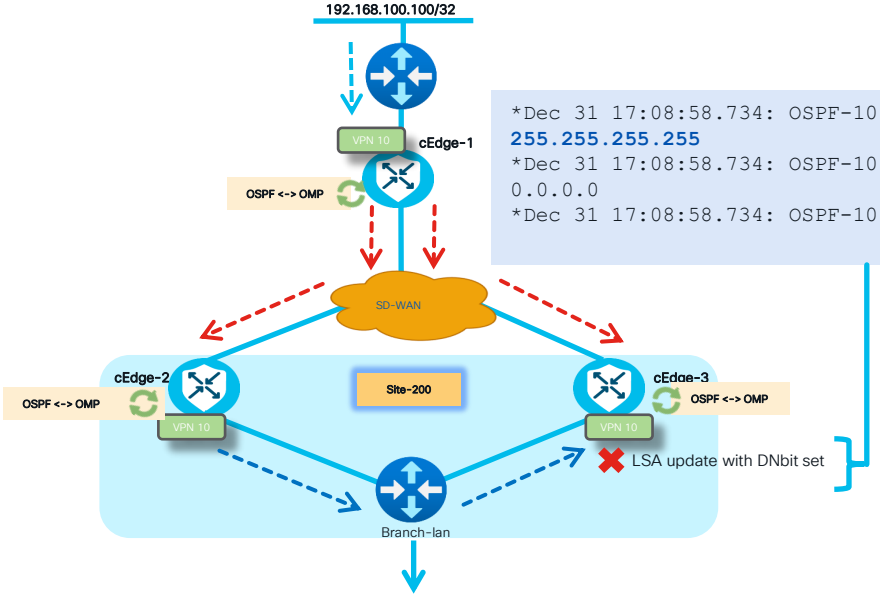


```
*Dec 31 17:08:58.734: OSPF-10 EXTER: Start processing AS External LSA 5/192.168.100.100/10.110.10.1, mask
255.255.255.255
*Dec 31 17:08:58.734: OSPF-10 EXTER:  age 885, seq 0x80000028, lsa_metric 100, metric-type 1, fw-addr
0.0.0.0
*Dec 31 17:08:58.734: OSPF-10 EXTER: Downward bit with SDWAN, ignoring the LSA
```

```
Branch-Lan#show ip route 192.168.100.100 255.255.255.255
Routing entry for 192.168.100.100/32
  Known via "ospf 10", distance 110, metric 101, type extern 1
  Last update from 10.110.10.1 on GigabitEthernet0/3, 03:35:49 ago
  Routing Descriptor Blocks:
  * 10.110.10.1, from 10.110.10.1, 03:35:49 ago, via GigabitEthernet0/3
      Route metric is 101, traffic share count is 1
```

```
LSA-type 5 (AS-External-LSA (ASBR)), len 36
    .000 0000 1110 0101 = LS Age (seconds): 229
    0... .... .... .... = Do Not Age Flag: 0
  Options: 0xa0, DN, (DC) Demand Circuits
    1... .... = DN: Set
    .0.. .... = (O) Opaque: Not set
    ..1. .... = (DC) Demand Circuits: Supported
    ...0 .... = (L) LLS Data block: Not Present
    .... 0... = (N) NSSA: Not supported
    .... .0.. = (MC) Multicast: Not capable
    .... ..0. = (E) External Routing: Not capable
    .... ...0 = (MT) Multi-Topology Routing: No
  LS Type: AS-External-LSA (ASBR) (5)
  Link State ID: 192.168.100.100
  Advertising Router: 10.110.10.1
  Sequence Number: 0x80000028
  Checksum: 0xa9c9
  Length: 36
  Netmask: 255.255.255.255
  0... .... = External Type: Type 1 (metric is specified in the same units as interface cost)
  .000 0000 = TOS: 0
  Metric: 100
  Forwarding Address: 0.0.0.0
  External Route Tag: 0
```

# How loop avoidance works in EIGRP 1/2?

```
External Route  =   192.168.200.200/32
  Type: External Route (0x0603)
  Length: 65
  Topology: 0
  AFI: IPv4 (1)
  RouterID: 10.110.10.1
> Wide Metric
  NextHop: 0.0.0.0
v External Data
    Originating RouterID: 10.110.10.1
    Originating A.S.: 0
    Administrative Tag: 0
    External Metric: 4294967294
    Reserved: 0
    External Protocol ID: Unknown (17)
  > External Flags
  Prefix Length: 32
  Destination: 192.168.200.200
```

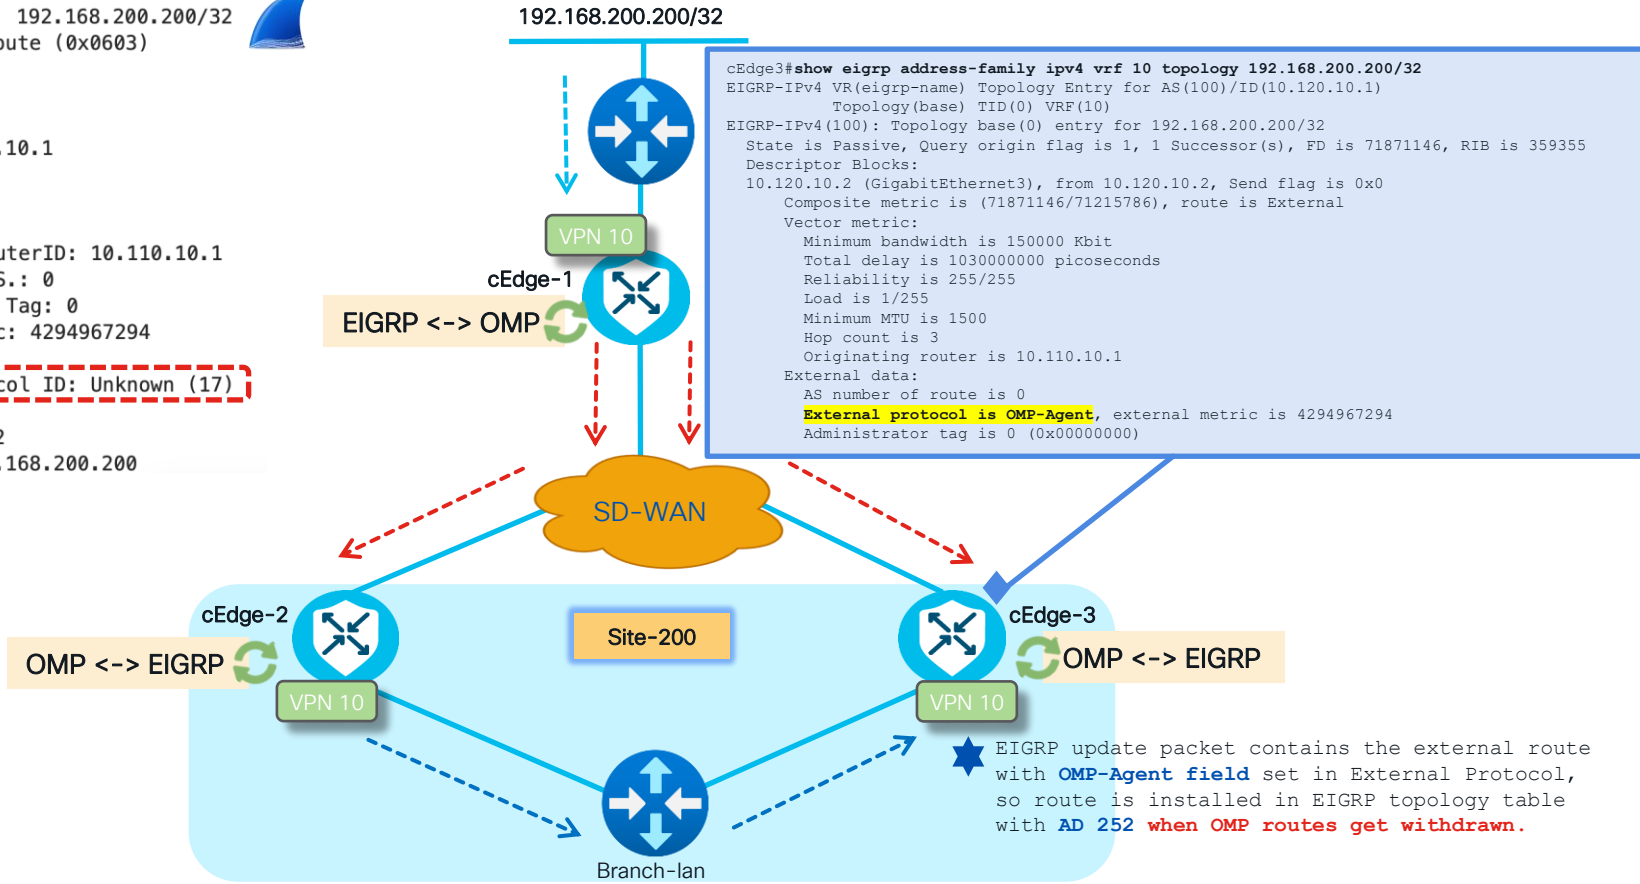


```
cEdge3#show eigrp address-family ipv4 vrf 10 topology 192.168.200.200/32
EIGRP-IPv4 VR(eigrp-name) Topology Entry for AS(100)/ID(10.120.10.1)
          Topology(base) TID(0) VRF(10)
EIGRP-IPv4(100): Topology base(0) entry for 192.168.200.200/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 71871146, RIB is 359355
  Descriptor Blocks:
  10.120.10.2 (GigabitEthernet3), from 10.120.10.2, Send flag is 0x0
      Composite metric is (71871146/71215786), route is External
      Vector metric:
        Minimum bandwidth is 150000 Kbit
        Total delay is 1030000000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 3
        Originating router is 10.110.10.1
      External data:
        AS number of route is 0
        External protocol is OMP-Agent, external metric is 4294967294
        Administrator tag is 0 (0x00000000)
```

EIGRP update packet contains the external route with **OMP-Agent field** set in External Protocol, so route is installed in EIGRP topology table with **AD 252 when OMP routes get withdrawn.**

# How loop avoidance works in EIGRP 2/2?

```
External Route  =   192.168.200.200/32
  Type: External Route (0x0603)
  Length: 65
  Topology: 0
  AFI: IPv4 (1)
  RouterID: 10.110.10.1
> Wide Metric
  NextHop: 0.0.0.0
v External Data
    Originating RouterID: 10.110.10.1
    Originating A.S.: 0
    Administrative Tag: 0
    External Metric: 4294967294
    Reserved: 0
    External Protocol ID: Unknown (17)
  > External Flags
  Prefix Length: 32
  Destination: 192.168.200.200
```

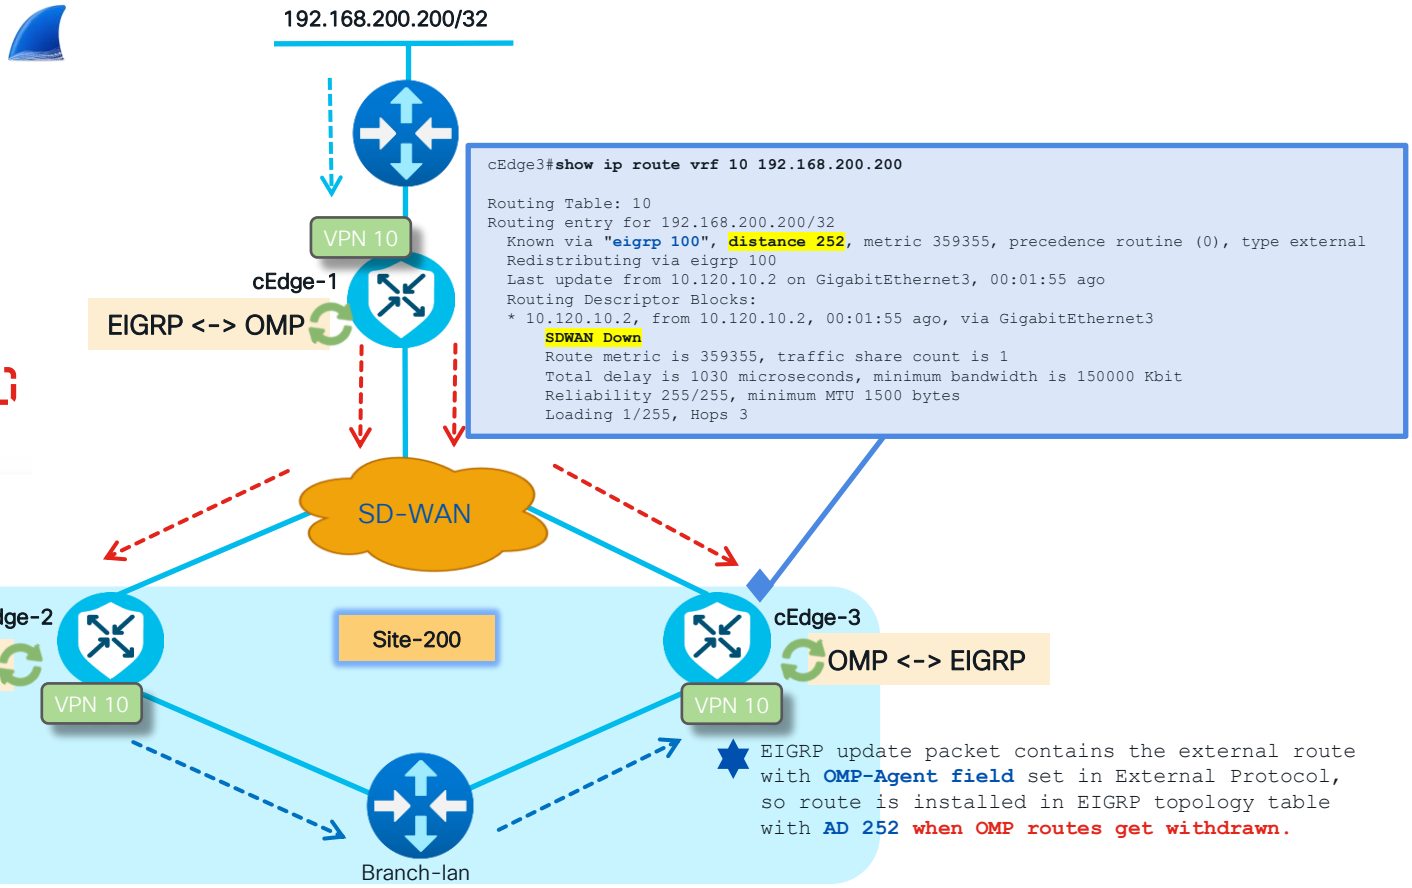

# How loop avoidance works in RIP?

```
IP Address: 192.168.200.200, Metric: 2
  Address Family: IP (2)
  Route Tag: 44270
  IP Address: 192.168.200.200
  Netmask: 255.255.255.255
  Next Hop: 10.120.10.1
  Metric: 2
```

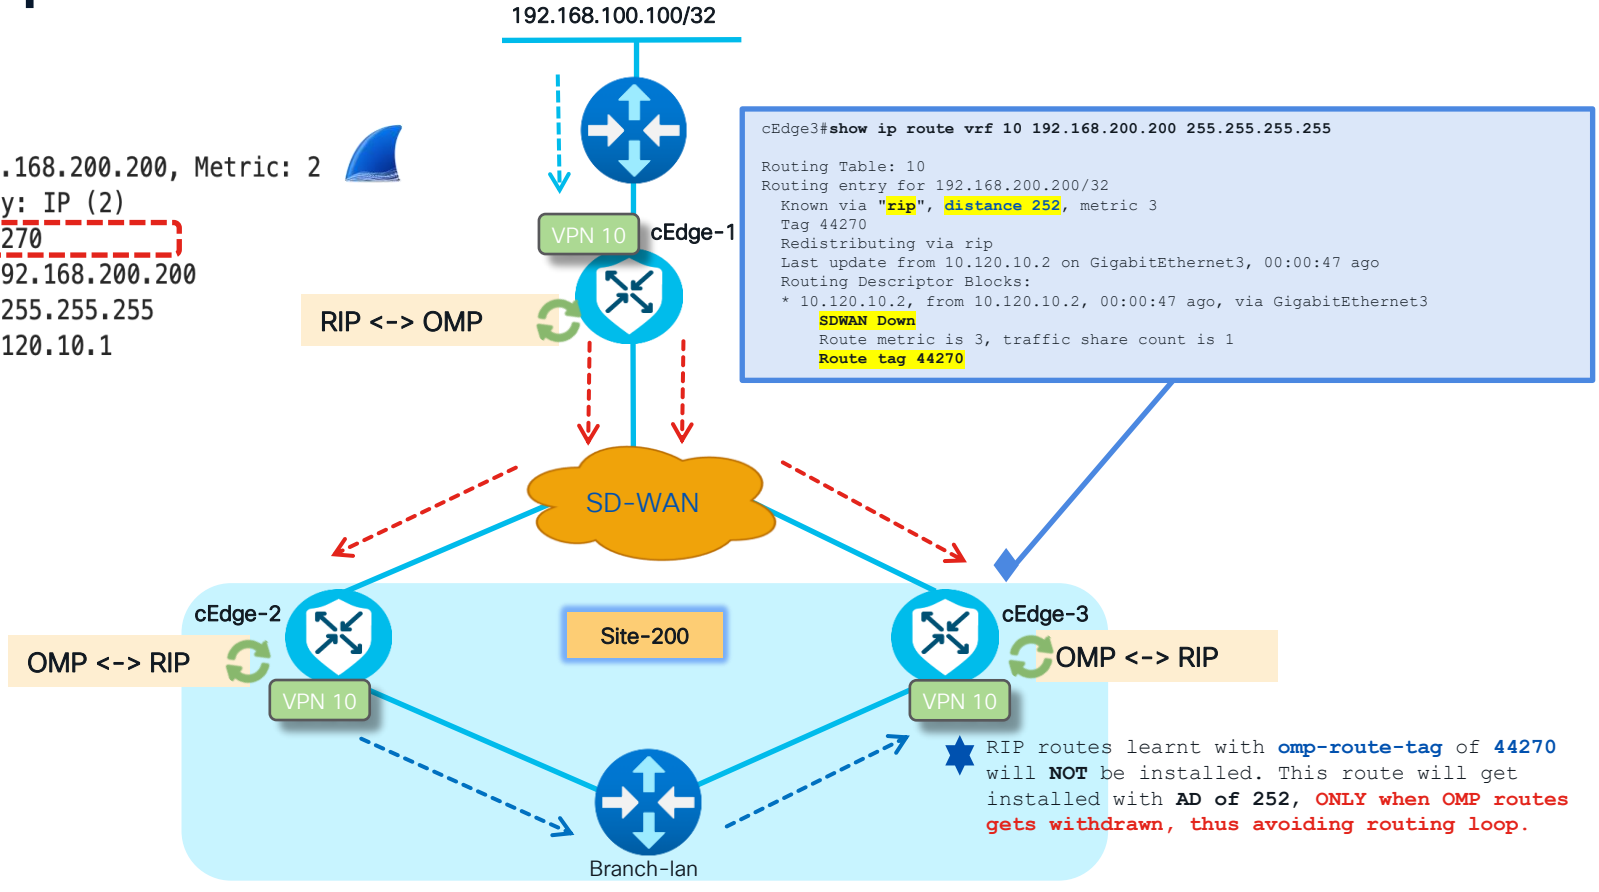


```
cEdge3#show ip route vrf 10 192.168.200.200 255.255.255.255

Routing Table: 10
Routing entry for 192.168.200.200/32
  Known via "rip", distance 252, metric 3
  Tag 44270
  Redistributing via rip
  Last update from 10.120.10.2 on GigabitEthernet3, 00:00:47 ago
  Routing Descriptor Blocks:
  * 10.120.10.2, from 10.120.10.2, 00:00:47 ago, via GigabitEthernet3
      SDWAN Down
      Route metric is 3, traffic share count is 1
      Route tag 44270
```

RIP routes learnt with **omp-route-tag** of **44270** will **NOT** be installed. This route will get installed with **AD of 252**, **ONLY when OMP routes gets withdrawn, thus avoiding routing loop.**

# How loop avoidance works in BGP 1/2?

```
Border Gateway Protocol - UPDATE Message
  Marker: ffffffffffffffffffffffffffffffff
  Length: 92
  Type: UPDATE Message (2)
  Withdrawn Routes Length: 0
  Total Path Attribute Length: 50
  Path attributes
    Path Attribute - ORIGIN: INCOMPLETE
    Path Attribute - AS_PATH: 65500 64500
    Path Attribute - NEXT_HOP: 10.101.10.1
    Path Attribute - MULTI_EXIT_DISC: 1000
    Path Attribute - EXTENDED_COMMUNITIES
      Flags: 0xc0, Optional, Transitive, Complete
      Type Code: EXTENDED_COMMUNITIES (16)
      Length: 16
      Carried extended communities: (2 communities)
        Route Origin: 0:100 [Transitive 2-Octet AS-Specific]
          Type: Transitive 2-Octet AS-Specific (0x00)
          Subtype (AS2): Route Origin (0x03)
          2-Octet AS: 0
          4-Octet AN: 100
        Route Target: 65500:10 [Transitive 2-Octet AS-Specific]
  Network Layer Reachability Information (NLRI)
    10.0.100.0/24
    10.110.10.101/32
    10.110.10.102/32
    192.168.200.200/32
      NLRI prefix length: 32
      NLRI prefix: 192.168.200.200
```

```
cEdge3#show bgp vpnv4 unicast vrf 10 192.168.200.200 255.255.255.255
BGP routing table entry for 1:10:192.168.200.200/32, version 265
Paths: (1 available, best #1, table 10)
  Advertised to update-groups:
     3
  Refresh Epoch 1
  64500
    10.110.110.110 (via default) from 0.0.0.0 (10.101.10.1)
      Origin incomplete, metric 1000, localpref 50, valid, sourced, best
      Extended Community: SoO:0:100 RT:65500:10
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 3 2024 10:08:13 UTC


cEdge3#show ip route vrf 10 192.168.200.200 255.255.255.255

Routing Table: 10
Routing entry for 192.168.200.200/32
  Known via "omp", distance 251, metric 0, type omp
  Redistributing via bgp 65500
  Advertised by bgp 65500
  Last update from 10.110.110.110 on Sdwan-system-intf, 00:54:37 ago
  Routing Descriptor Blocks:
  * 10.110.110.110 (default), from 10.110.110.110, 00:54:37 ago, via Sdwan-system-intf
      opaque_ptr 0x71E1A66566E0
      Route metric is 0, traffic share count is 1
```

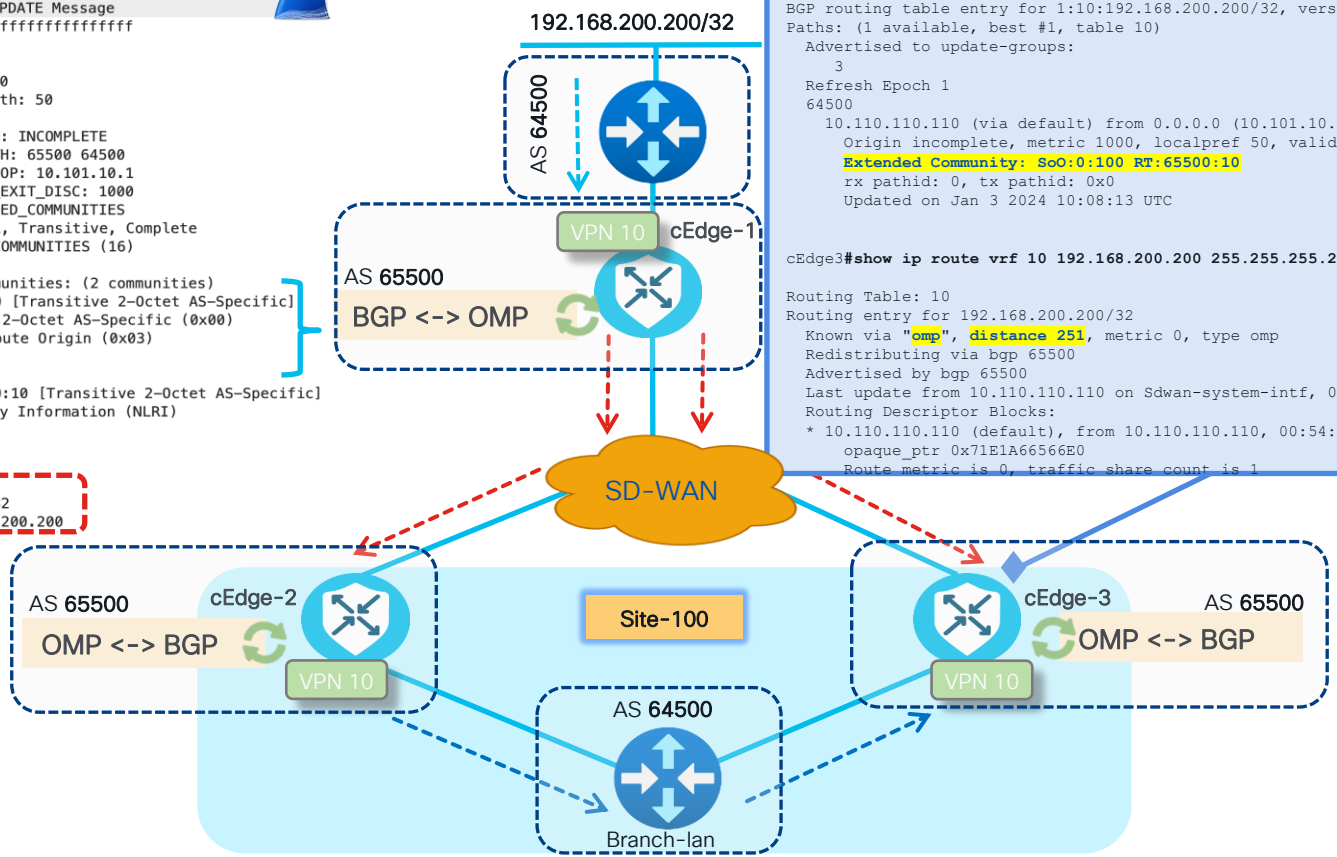

# How loop avoidance works in BGP 2/2?

```
Border Gateway Protocol - UPDATE Message
  Marker: ffffffffffffffffffffffffffffffff
  Length: 92
  Type: UPDATE Message (2)
  Withdrawn Routes Length: 0
  Total Path Attribute Length: 50
  Path attributes
    Path Attribute - ORIGIN: INCOMPLETE
    Path Attribute - AS_PATH: 65500 64500
    Path Attribute - NEXT_HOP: 10.101.10.1
    Path Attribute - MULTI_EXIT_DISC: 1000
    Path Attribute - EXTENDED_COMMUNITIES
      Flags: 0xc0, Optional, Transitive, Complete
      Type Code: EXTENDED_COMMUNITIES (16)
      Length: 16
      Carried extended communities: (2 communities)
        Route Origin: 0:100 [Transitive 2-Octet AS-Specific]
          Type: Transitive 2-Octet AS-Specific (0x00)
          Subtype (AS2): Route Origin (0x03)
          2-Octet AS: 0
          4-Octet AN: 100
        Route Target: 65500:10 [Transitive 2-Octet AS-Specific]
  Network Layer Reachability Information (NLRI)
    10.0.100.0/24
    10.110.10.101/32
    10.110.10.102/32
    192.168.200.200/32
      NLRI prefix length: 32
      NLRI prefix: 192.168.200.200
```

```
cEdge3#show bgp vpnv4 unicast vrf 10 192.168.200.200
BGP routing table entry for 1:10:192.168.200.200/32, version 379
Paths: (1 available, best #1, table 10)
  Not advertised to any peer
  Refresh Epoch 1
  64500 65500 64500
    10.101.10.2 (via vrf 10) from 10.101.10.2 (10.101.10.2)
      Origin incomplete, localpref 100, valid, external, best
      Extended Community: SoO:0:100 RT:65500:10
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 3 2024 14:02:13 UTC

cEdge3#show ip route vrf 10 192.168.200.200

Routing Table: 10
Routing entry for 192.168.200.200/32
  Known via "bgp 65500", distance 252, metric 0
  Tag 64500, type external
  Last update from 10.101.10.2 00:00:37 ago
  Routing Descriptor Blocks:
  * 10.101.10.2, from 10.101.10.2, 00:00:37 ago
      opaque_ptr 0x71E1AA80CEB0
      SDWAN Down
      Route metric is 0, traffic share count is 1
      AS Hops 3
      Route tag 64500
      MPLS label: none
```

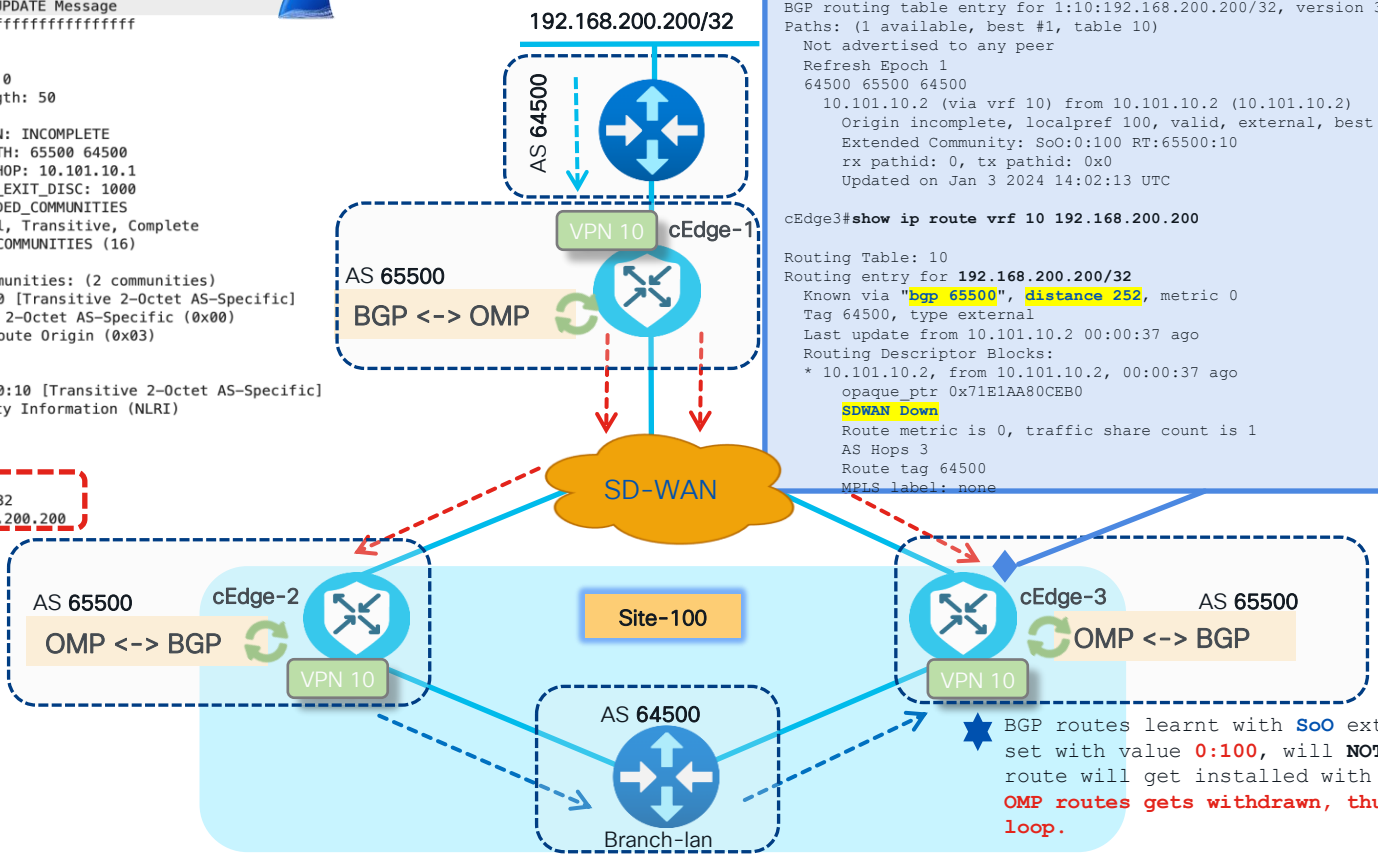


BGP routes learnt with **SoO** ext-community attribute set with value **0:100**, will **NOT** be installed. This route will get installed with **AD of 252**, **ONLY when OMP routes gets withdrawn, thus avoiding routing loop.**

# Prefer Direct e/iBGP over OMP-Redistributed Routes

```
cEdge-2#show bgp vpnv4 unicast all

     Network          Next Hop            Metric LocPrf Weight Path

Route Distinguisher: 1:10 (default for vrf 10)

*>   192.168.200.200/32

                      10.0.0.1                 1000      50      0 64500 ?
```

- OMP routes redistributed into BGP are assigned **LOCAL_PREF 50** and **MED 1000** by default, ensuring they are less preferred than locally learned routes via **iBGP (higher local preference)** and **eBGP (lower MED)**.

| Path | Local Pref | MED | Result |
|---|---|---|---|
| OMP → BGP | 50 | 1000 | ❌ Loses |
| Direct eBGP | 100 | 0 | ✅ Wins |

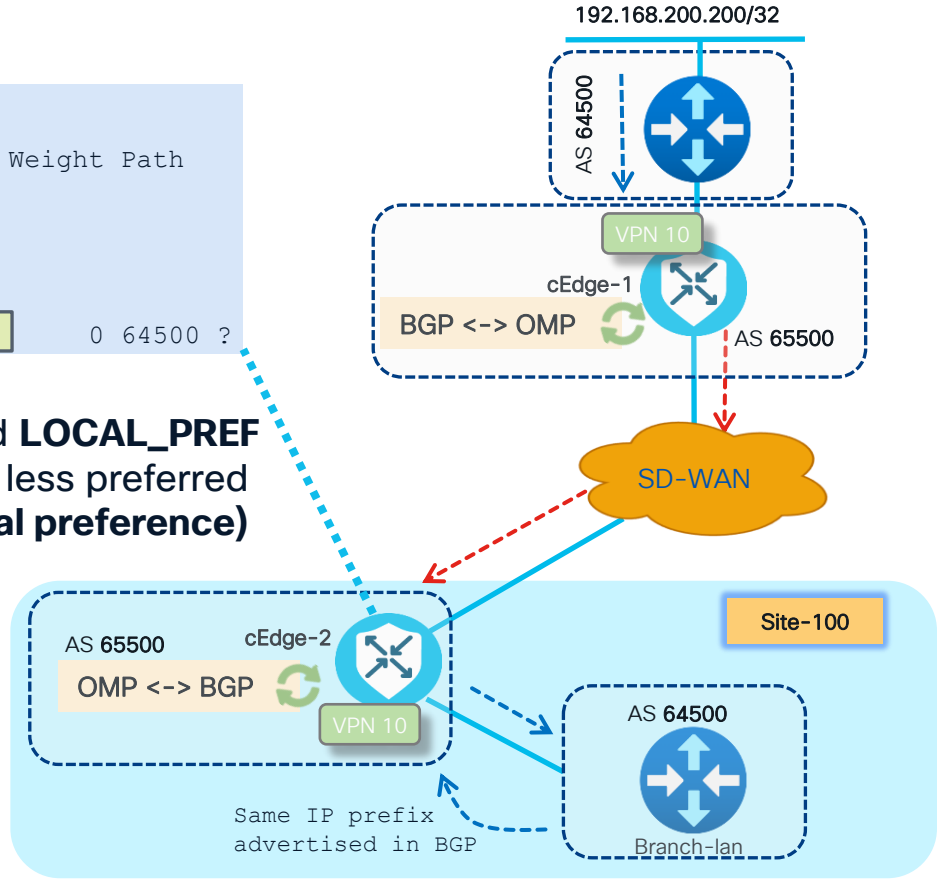

# SD-WAN Multicast

# OMP Multicast Auto-Discovery

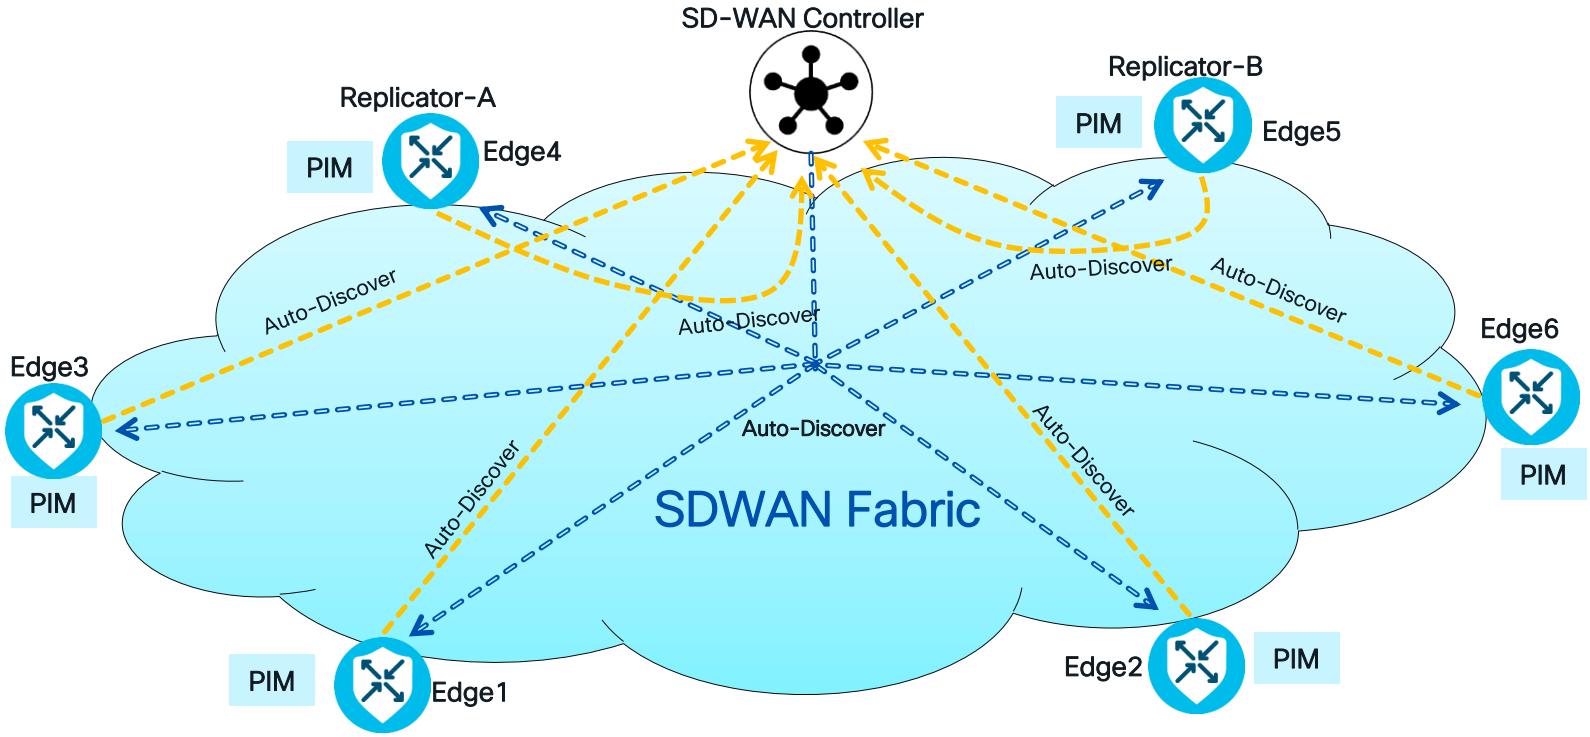

# OMP Multicast Auto-Discovery

- OMP Multicast Auto-discovery packet contains:
  - List of WAN Edge devices are configured as multicast capable
  - List of devices are configured as replicator
  - Replicator capacity, GPS location etc.
- The multicast Autodiscover routes indicate whether the router has PIM enabled and whether it is a replicator.

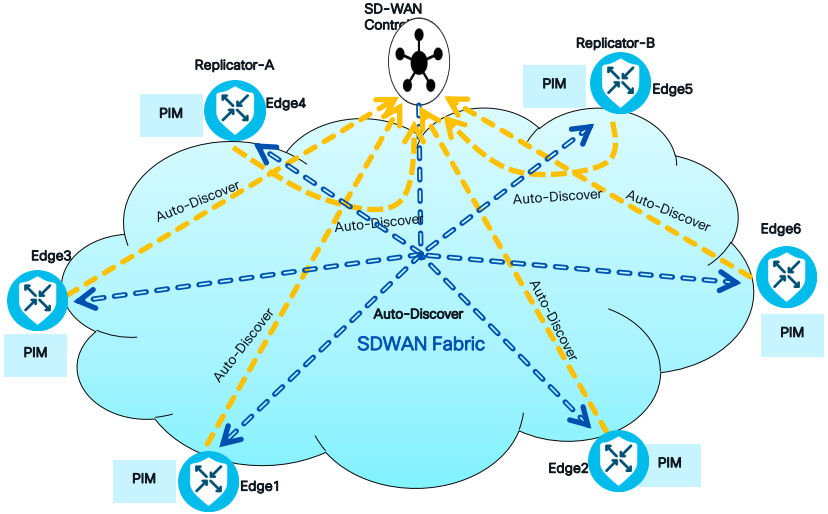

# SD-WAN Multicast control packet flow - ASM

- Source register itself to an RP
- Receiver sends the (*,G) join
- First Join gets forwarded to the SD-WAN controller as an OMP packet and then forwarded to the replicator
- Replicator forwards (*,G) to the RP
- RP forwards it to the source
- Stream is forwarded to the receiver through the replicator. Stream NEVER goes to SD-WAN controller.
- Once receiver has the source information, it will then join using (S,G)
- First (S,G) join gets forwarded as an OMP control packet to the SD-WAN controller and then to replicator
- Replicator then forward the (S,G) to the source.

# OMP Multicast Message

| Route Type | OMP Overlay Multicast Route Type |
|---|---|
| 5 | Source Active |
| 6 | Shared Tree Join |
| 7 | Source Tree Join |

# SD-WAN Multicast control packet flow – ASM 1/3

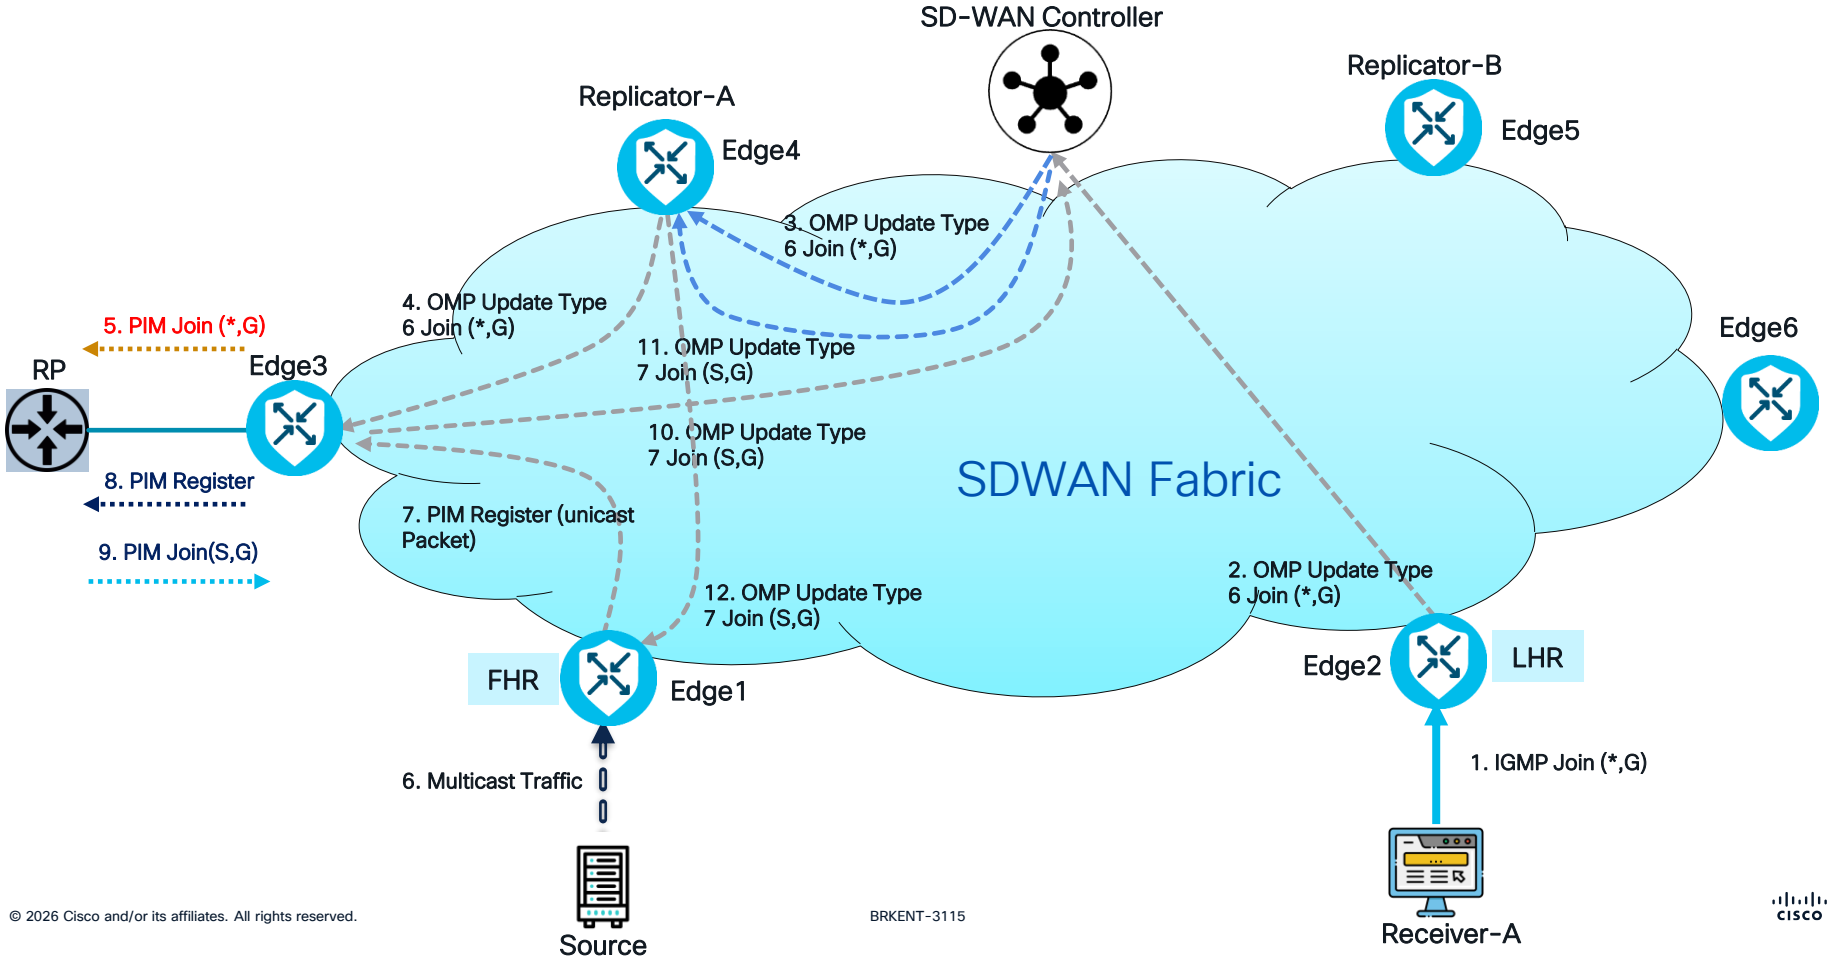

# SD-WAN Multicast control packet flow – ASM 2/3

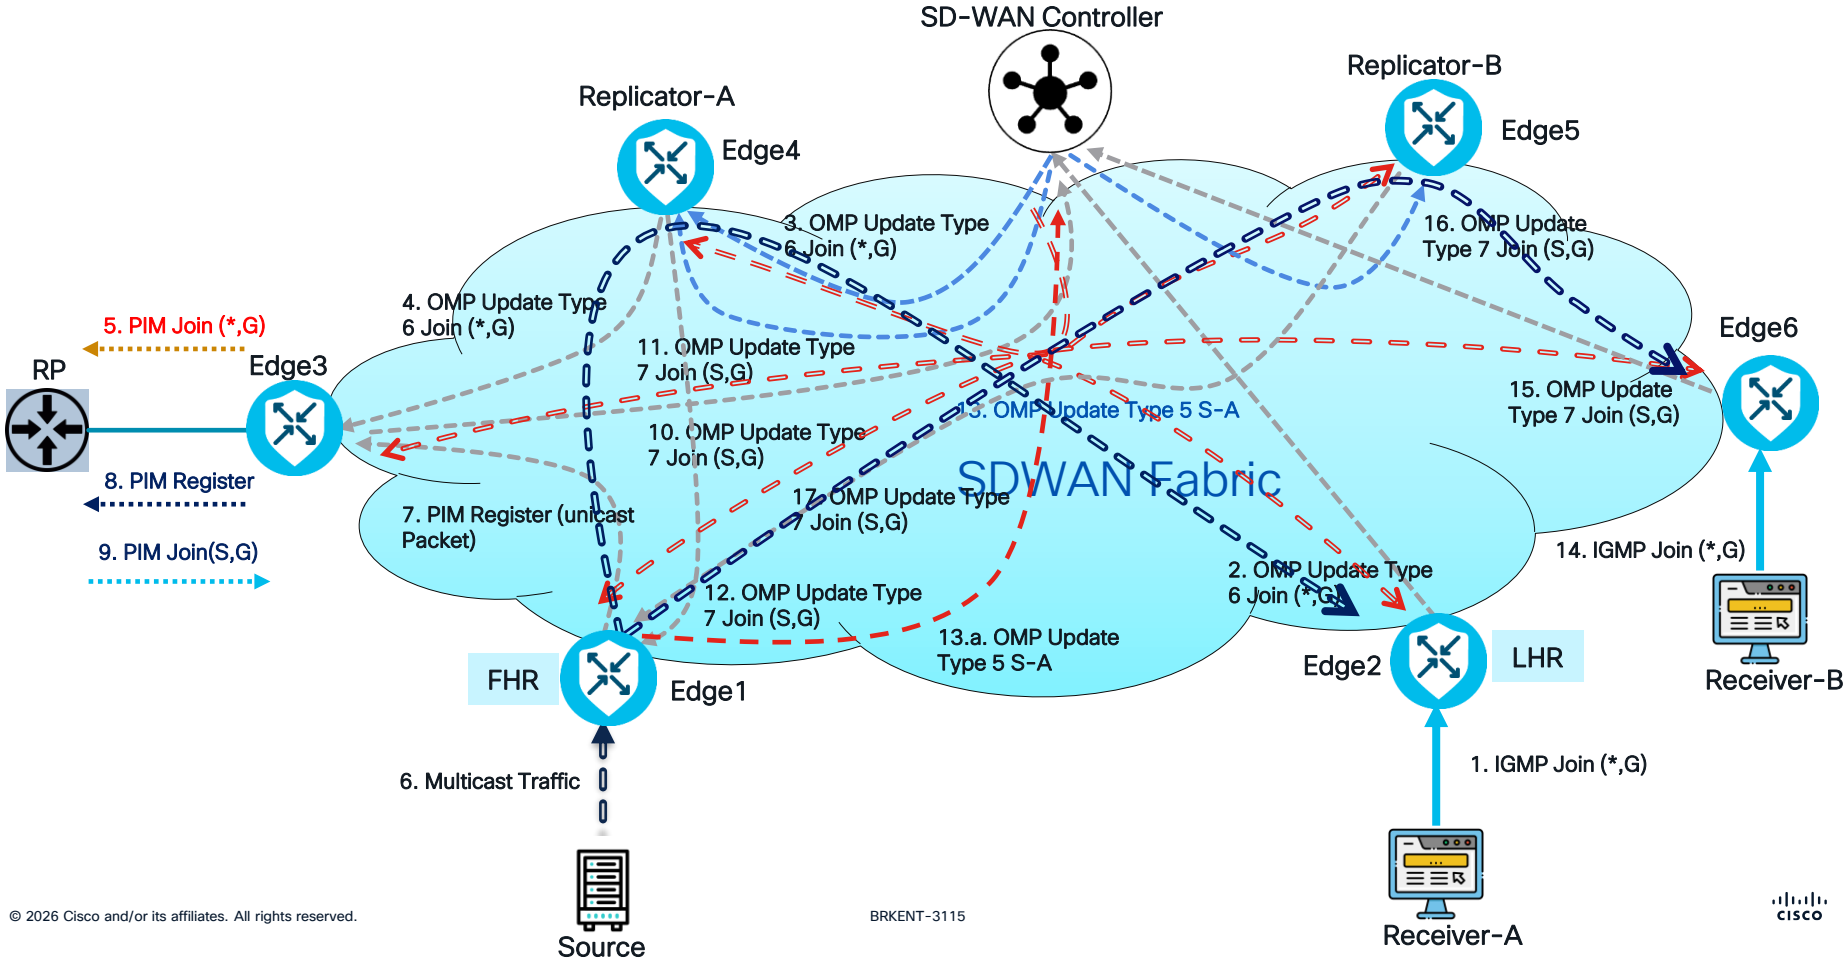

# SD-WAN Multicast control packet flow – ASM 3/3

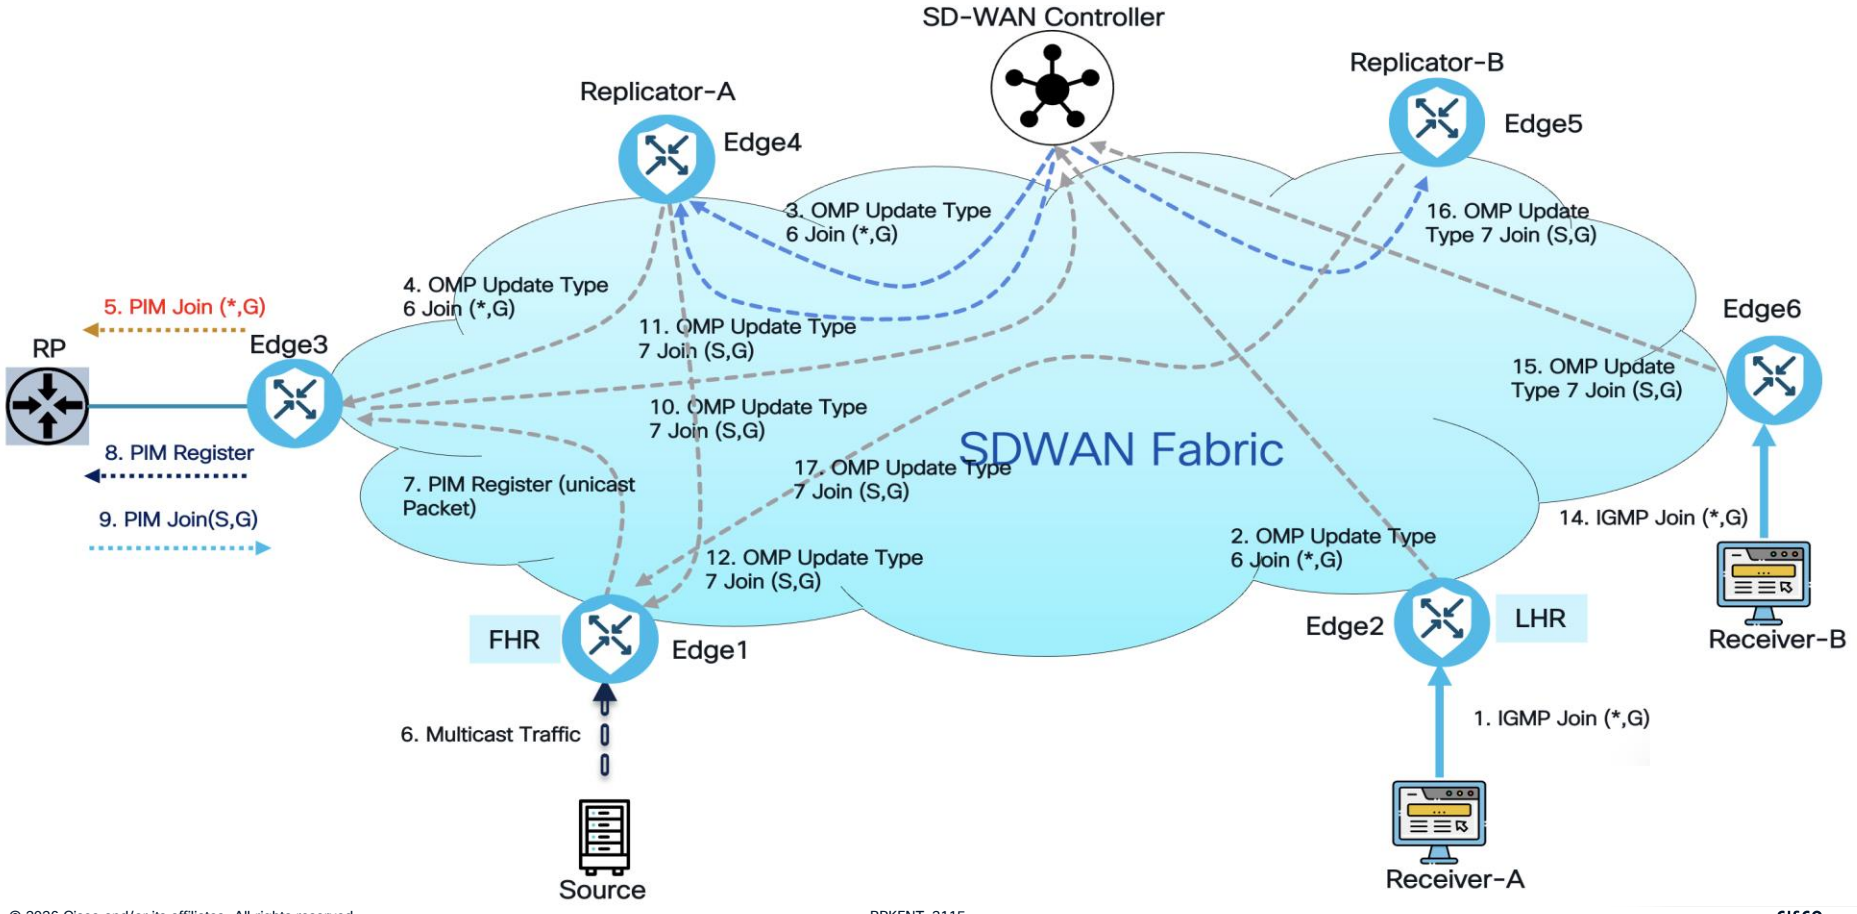

# SD-WAN Multicast control packet flow – ASM 3/3

For Reference

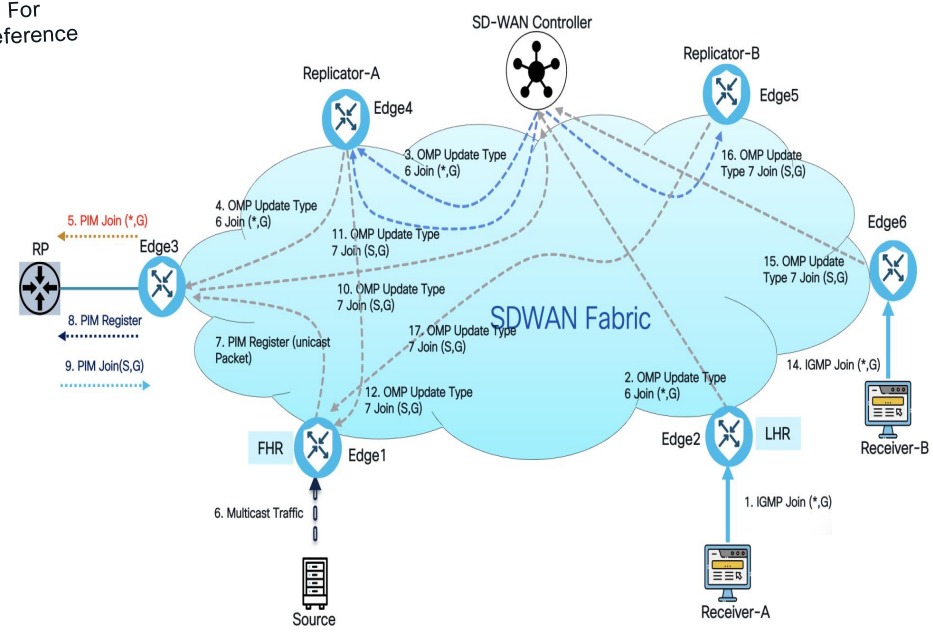


- IGMP join from Receiver is sent to LHR
- Type 6 (*,G) OMP join from LHR is sent to SD-WAN controller and SD-WAN controller forwards it to Replicator.
- Type 6 (*,G) OMP join from Replicator is sent to RP.
- Source streams the traffic to FHR.
- FHR sends PIM Registration message to RP.
- Type 7 (S,G) join from RP is sent to SD-WAN controller and then SD-WAN controller send it to Replicator.
- Type 7 (S,G) join from Replicator to is sent to FHR.
- Type 5 (S,A) message is sent from FHR to SD-WAN controller and SD-WAN controller sends to all the WAN Edges including LHR, RP and Replicator.
- Type 7 (S,G) join from LHR is sent to Replicator
- FHR transmits the traffic to the Replicator, which subsequently forwards it to LHR, and ultimately, the receiver obtains the Multicast stream.





cisco

# SD-WAN Multicast Key points 1/2

- The significance of replicator selection lies solely from the perspective of the LHR (receivers' point of view).
- When multiple replicators are involved, there is no synchronous mechanism in place. In the event that a replicator becomes inactive for a specific stream, the receiver will be required to re-join the tree using an alternative replicator.
- In scenarios where there are no receivers for the multicast traffic, a PIM Register unicast message will be sent to RP. However, since there are no receivers at that point, the RP will respond by transmitting a PIM Register STOP message back to the source.
- Enabling the "spt-mode" helps minimize the exchange of control plane packets for SDWAN multicast, while still allowing multicast data traffic to be replicated through the chosen replicator.
- It is crucial to ensure that "spt-mode" is enabled on every WAN edge router within the SDWAN fabric.

# SD-WAN Multicast Key points 2/2

- Data policy interworking with SDWAN overlay multicast is NOT officially tested before and not supported.

```
data-prefix-list Unicast_IPv4
   ip-prefix 0.0.0.0/1
   ip-prefix 128.0.0.0/2
   ip-prefix 192.0.0.0/3
  !
data-policy _VPN_1_AS-Set-local-tloc-policy_v2
  vpn-list VPN_1
    sequence 1
```

## Restrictions for Multicast Overlay Routing

Multicast overlay routing does not support the following features:

- MSDP/Anycast-RP on Cisco Catalyst SD-WAN routers
- IPv6 overlay and IPv6 underlay
- Dynamic BFD tunnel for multicast
- Multicast with asymmetric unicast routing
- Multicast overlay working does not support Data Policy. In case data policy is configured, then only required traffic is matched and not multicast traffic.

```
     set
      local-tloc-list
       color mpls
       encap ipsec
      !
     !
    !
  default-action accept
```

# SD-WAN Multicast Key points 2/2

- Data policy interworking with SDWAN overlay multicast is NOT officially tested before and not supported.
- Even localized QoS policy, it might count on data policy to set different forwarding class.
- While configuring data policy or localized QoS policy, we need to make sure ONLY unicast traffic is matched.
- Support for **Hub and Spoke** topology begins with the 17.15.1/20.15.1 software release.
  - Following configuration must be performed on spoke sites using **CLI-Add-on**.

```
!
multicast
 address-family ipv4 vrf <vrf-id>
  spoke
 !
```

```
data-prefix-list Unicast_IPv4
   ip-prefix 0.0.0.0/1
   ip-prefix 128.0.0.0/2
   ip-prefix 192.0.0.0/3
  !
data-policy _VPN_1_AS-Set-local-tloc-policy_v2
  vpn-list VPN_1
    sequence 1
     match
      source-ip 0.0.0.0/0
      dscp 10
      destination-data-prefix-list Unicast_IPv4 ← MATCH ONLY Unicast Traffic
     !
     action accept
      count dscp10_counter_1526348700
      set
       local-tloc-list
        color public-internet
        encap ipsec
      !
     !
    !
    sequence 11
     match
      source-ip 0.0.0.0/0
      destination-data-prefix-list Unicast_IPv4 ← MATCH ONLY Unicast Traffic
     !
     action accept
      count all_traffic_counter_1526348700
      set
       local-tloc-list
        color mpls
        encap ipsec
      !
     !
    !
  default-action accept
```

# SD-WAN Multicast Key points 2/2

- Data policy interworking with SDWAN overlay multicast is NOT officially tested before and not supported.
- Even localized QoS policy, it might count on data policy to set different forwarding class.
- While configuring data policy or localized QoS policy, we need to make sure ONLY unicast traffic is matched.
- Support for **Hub and Spoke** topology begins with the 17.15.1/20.15.1 software release.
  - Configure the routing metric (Administrative Distance) to direct all multicast traffic through **SD-WAN paths** on all hubs, rather than non-SD-WAN paths.

```
router bgp <asn>
address-family ipv4 unicast vrf <vrfid>
 distance <252 or higher> <bgp-neighbor-ip> <mask>
     <access-list-number>
 exit-address-family
!
```

```
data-prefix-list Unicast_IPv4
   ip-prefix 0.0.0.0/1
   ip-prefix 128.0.0.0/2
   ip-prefix 192.0.0.0/3
  !
data-policy _VPN_1_AS-Set-local-tloc-policy_v2
  vpn-list VPN_1
    sequence 1
     match
      source-ip 0.0.0.0/0
      dscp 10
      destination-data-prefix-list Unicast_IPv4 ← MATCH ONLY Unicast Traffic
     !
     action accept
      count dscp10_counter_1526348700
      set
       local-tloc-list
        color public-internet
        encap ipsec
      !
     !
    !
    sequence 11
     match
      source-ip 0.0.0.0/0
      destination-data-prefix-list Unicast_IPv4 ← MATCH ONLY Unicast Traffic
     !
     action accept
      count all_traffic_counter_1526348700
      set
       local-tloc-list
        color mpls
        encap ipsec
      !
     !
    !
default-action accept
```

# Layer 2 VPN (L2VPN)

# What Cisco Catalyst SD-WAN Layer 2 VPN (L2VPN) offers?

- Architecture Context
  - **Extends** Cisco Catalyst SD-WAN overlay to support **Layer 2 services**
  - Designed for applications requiring **L2 adjacency across WAN**
  - Coexists with native Layer-3 SD-WAN services, segmentation, and security
- Supported L2VPN Services
  - **Point-to-Point** (P2P) and **Point-to-Multipoin**t (P2MP) Ethernet VPN
  - **Full-mesh** and **hub-and-spoke** L2VPN topologies
  - **Single-homing** and **multi-homing** for resiliency and load sharing
- Forwarding & Control Plane Behaviour
  - Ingress replication for **Broadcast**, **Unknown-Unicast**, and **Multicast** (**BUM**).
  - MAC learning via **OMP** (**control plane**) instead of data-plane flooding.

# Demo L2VPN Topology

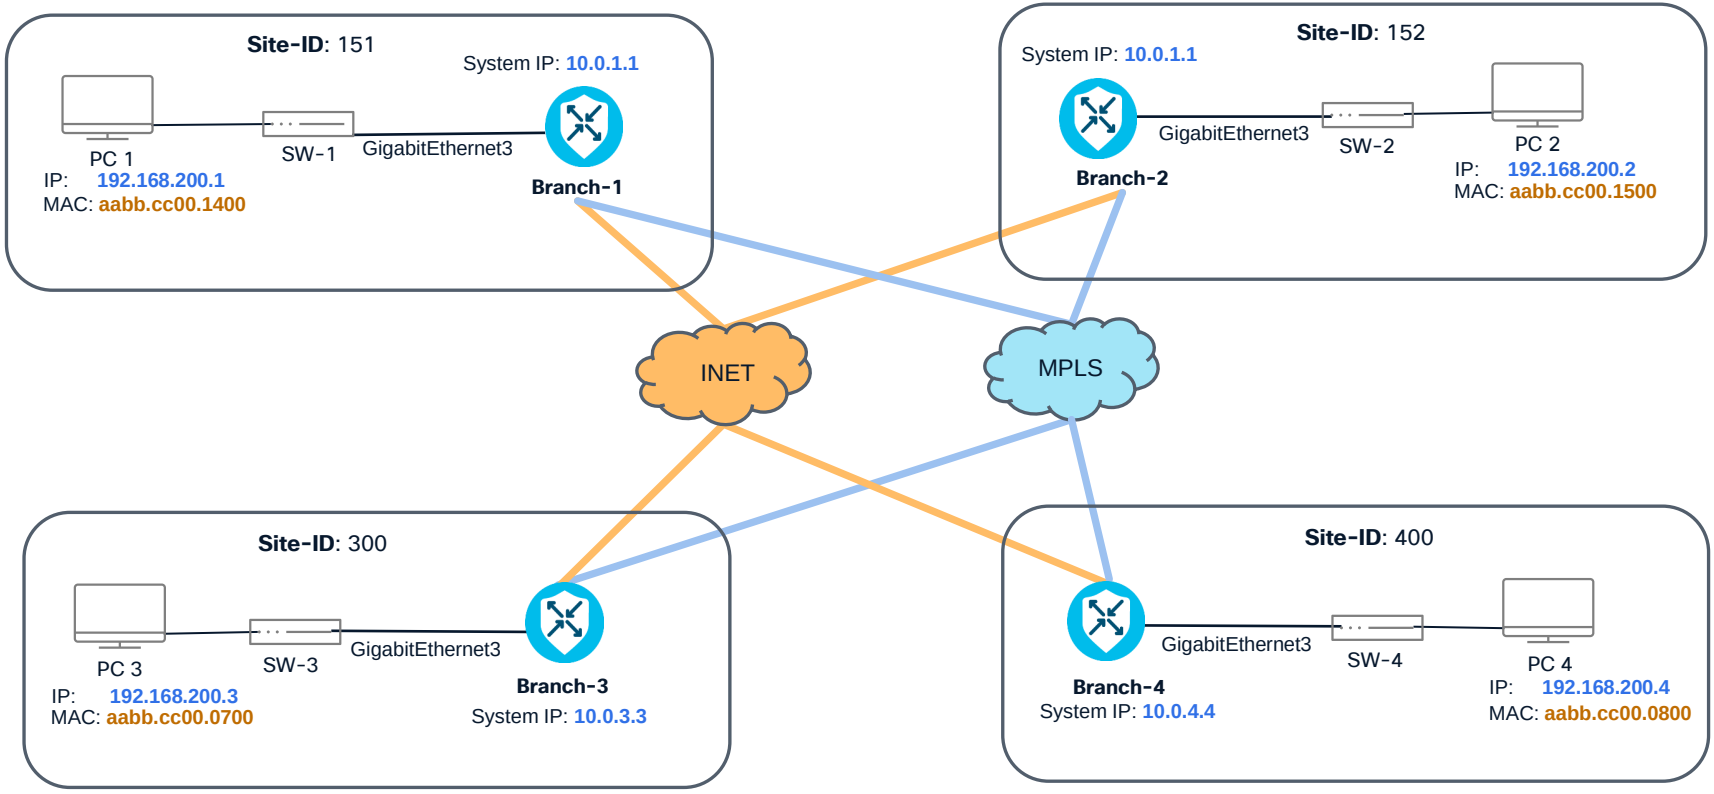

# L2VPN Point-to-Point (P2P)

# How is Point-to-Point (P2P) L2VPN configured?

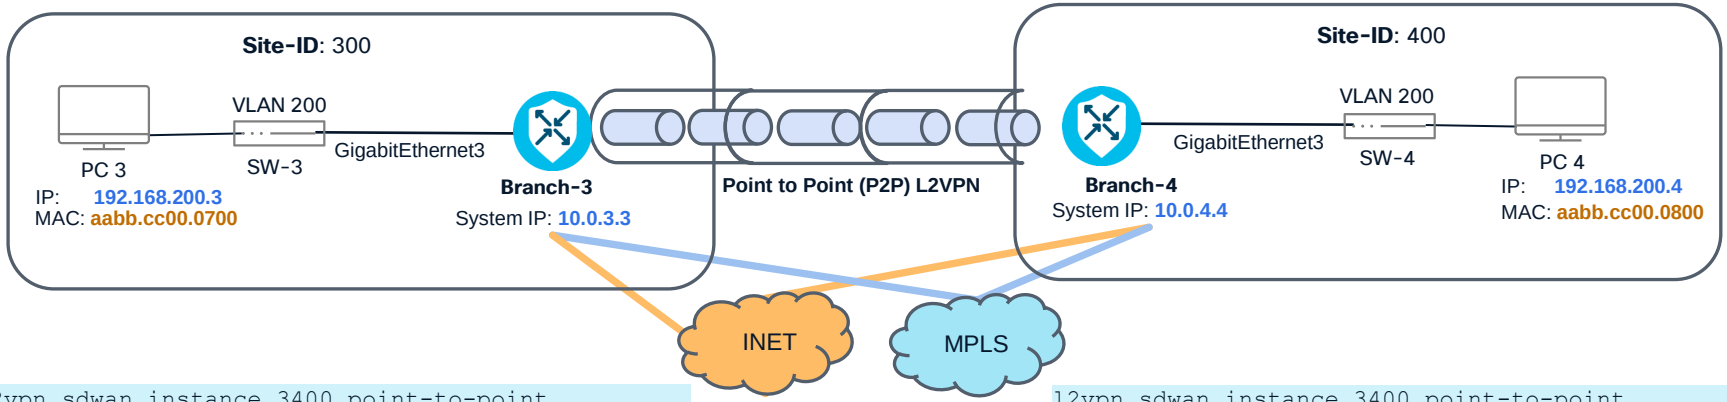


```
l2vpn sdwan instance 3400 point-to-point
!
bridge-domain 900
 member GigabitEthernet3 service-instance 300
 !
 member sdwan-instance 3400 remote-site-id 400
 vc-id 1000 single-homing
!
interface GigabitEthernet3
 no ip address
  service instance 300 ethernet
  encapsulation dot1q 200
 !
```

Currently, configuration is available exclusively via the **CLI Add-On**.

```
l2vpn sdwan instance 3400 point-to-point
!
bridge-domain 900
 member GigabitEthernet3 service-instance 400
 !
 member sdwan-instance 3400 remote-site-id 300
 vc-id 1000 single-homing
!
interface GigabitEthernet3
 no ip address
  service instance 400 ethernet
  encapsulation dot1q 200
 !
```

# L2VPN Point-to-Point(P2P) configuration breakdown

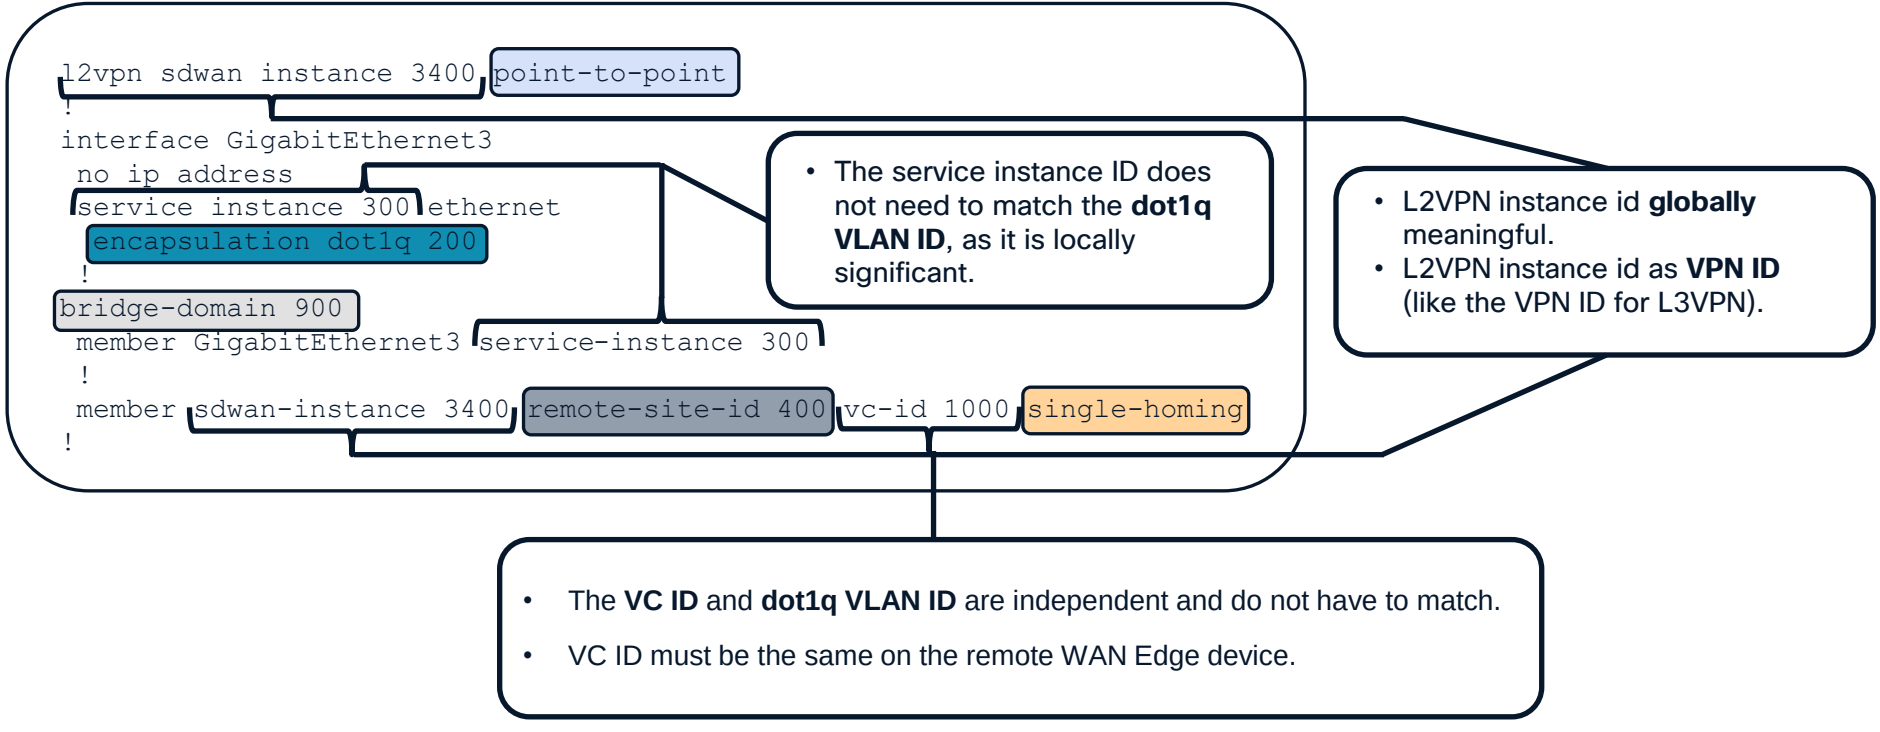

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 1/8

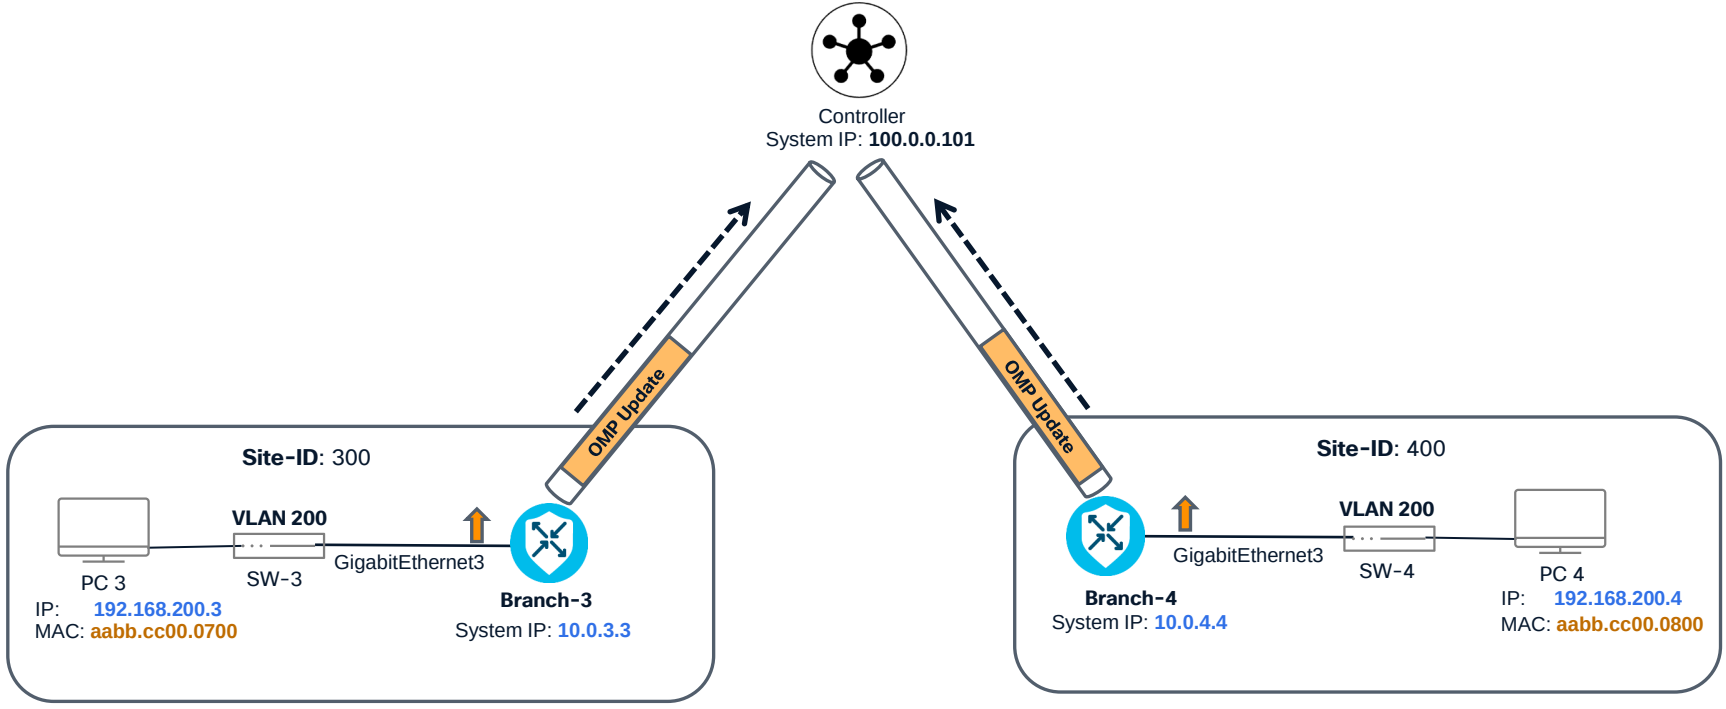

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 2/8

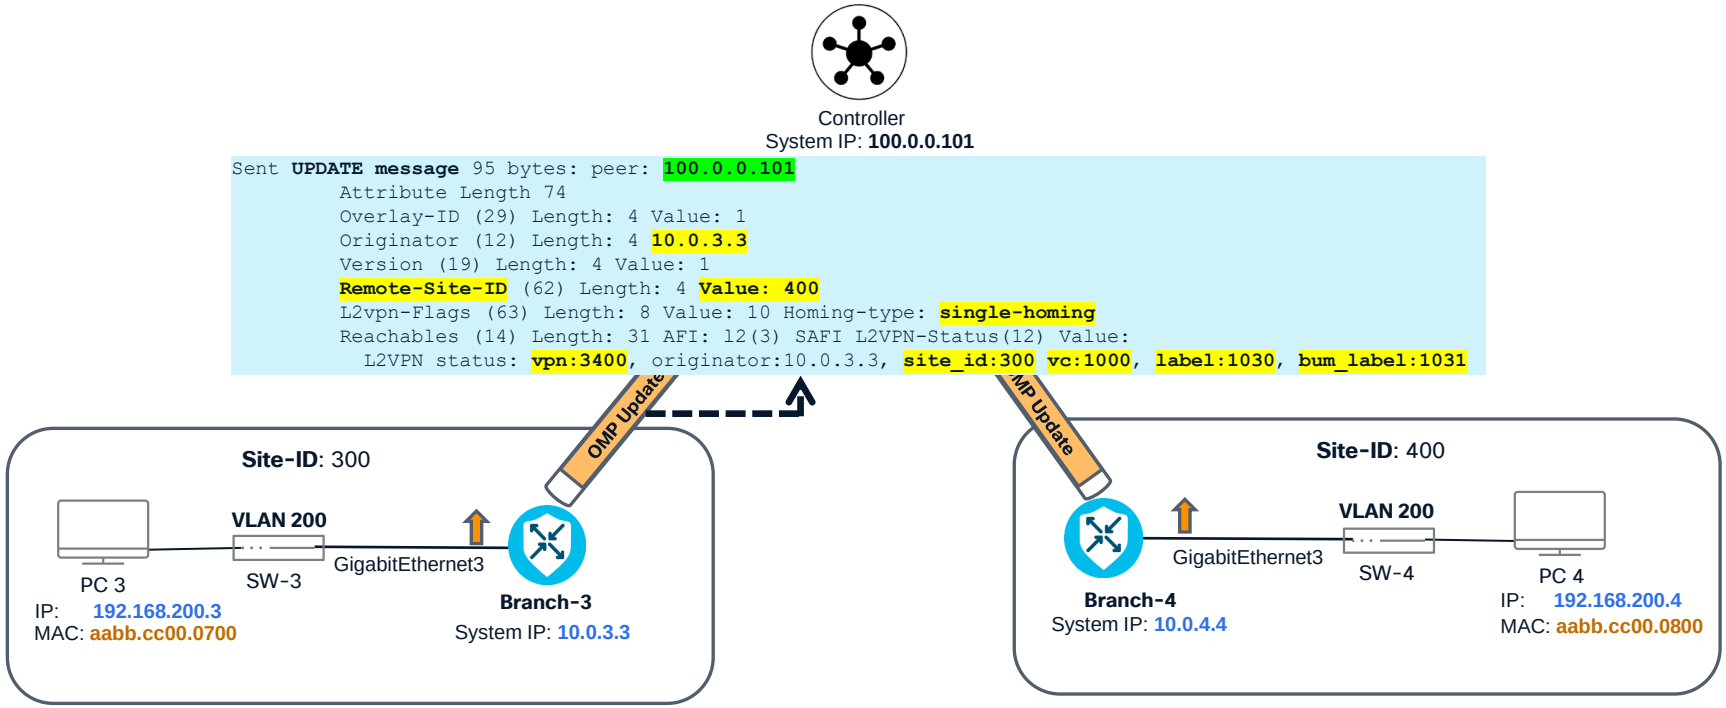

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 3/8

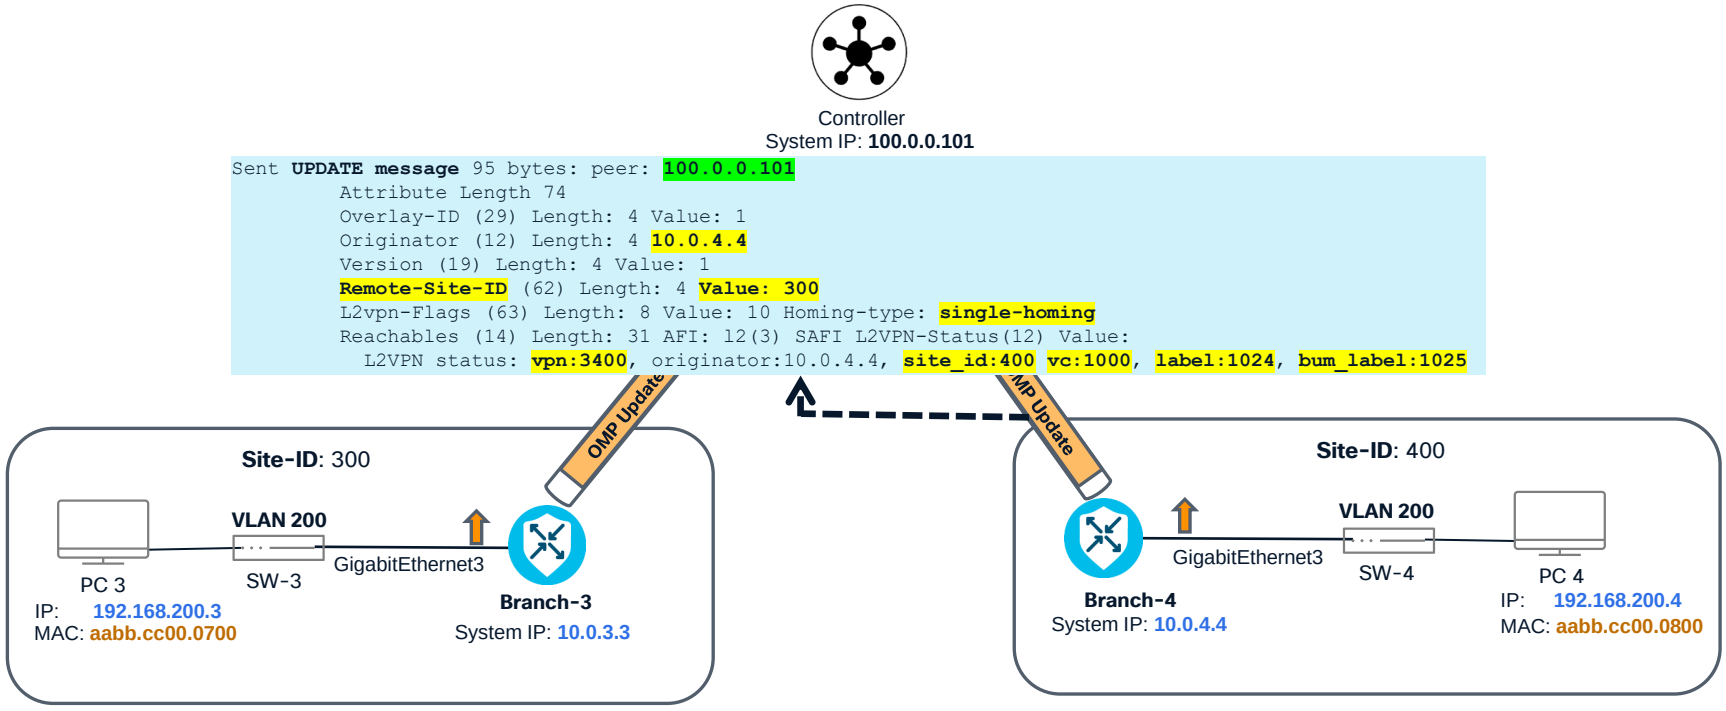

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 4/8

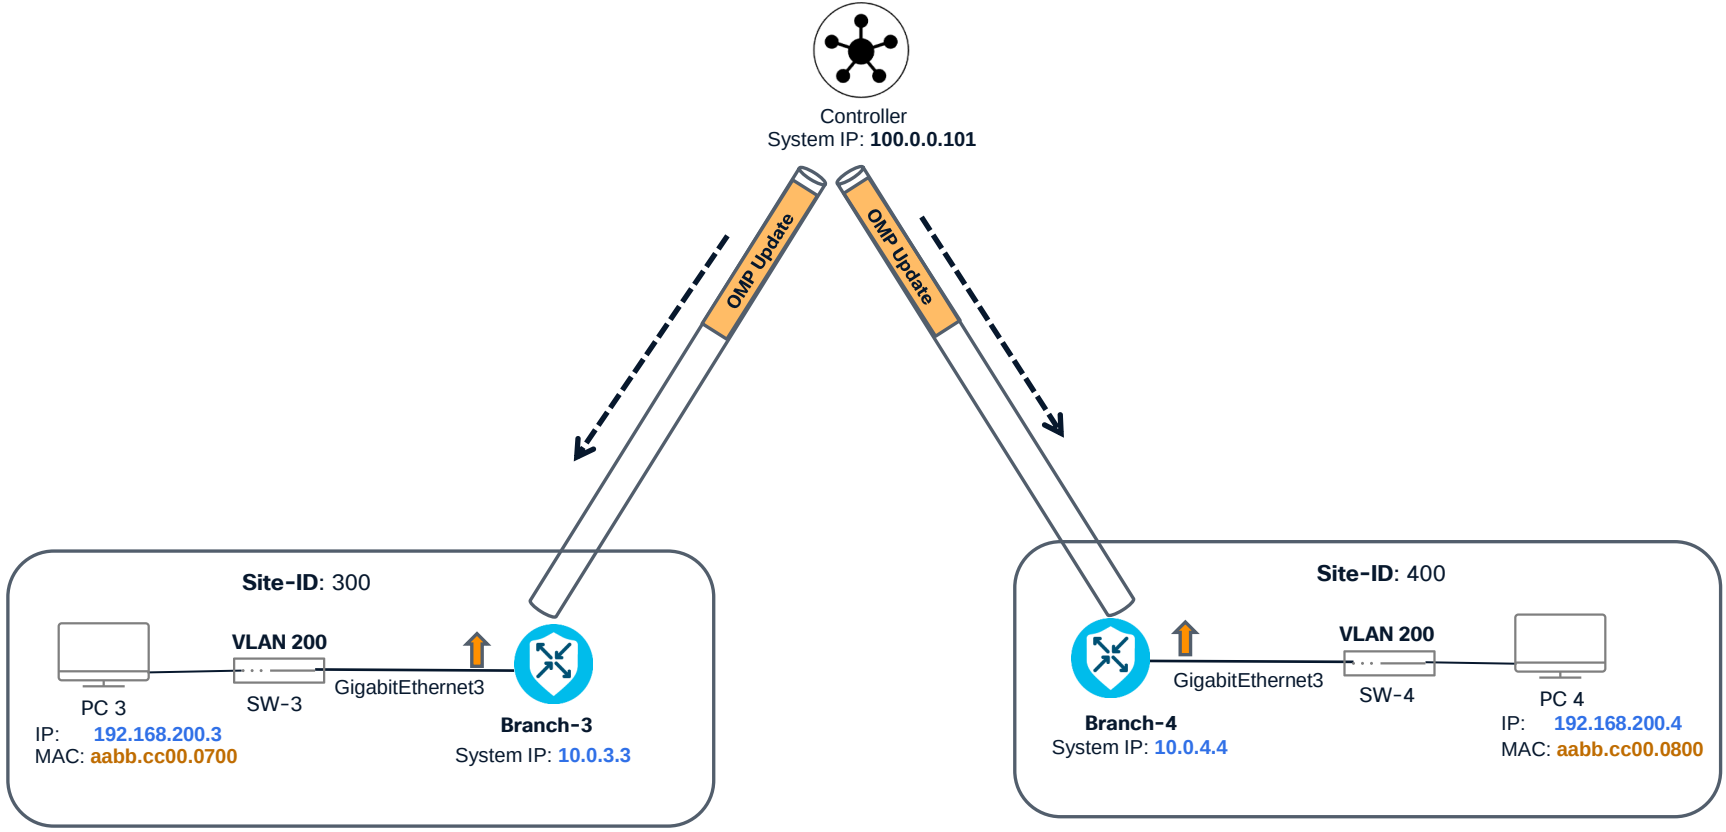

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 5/8

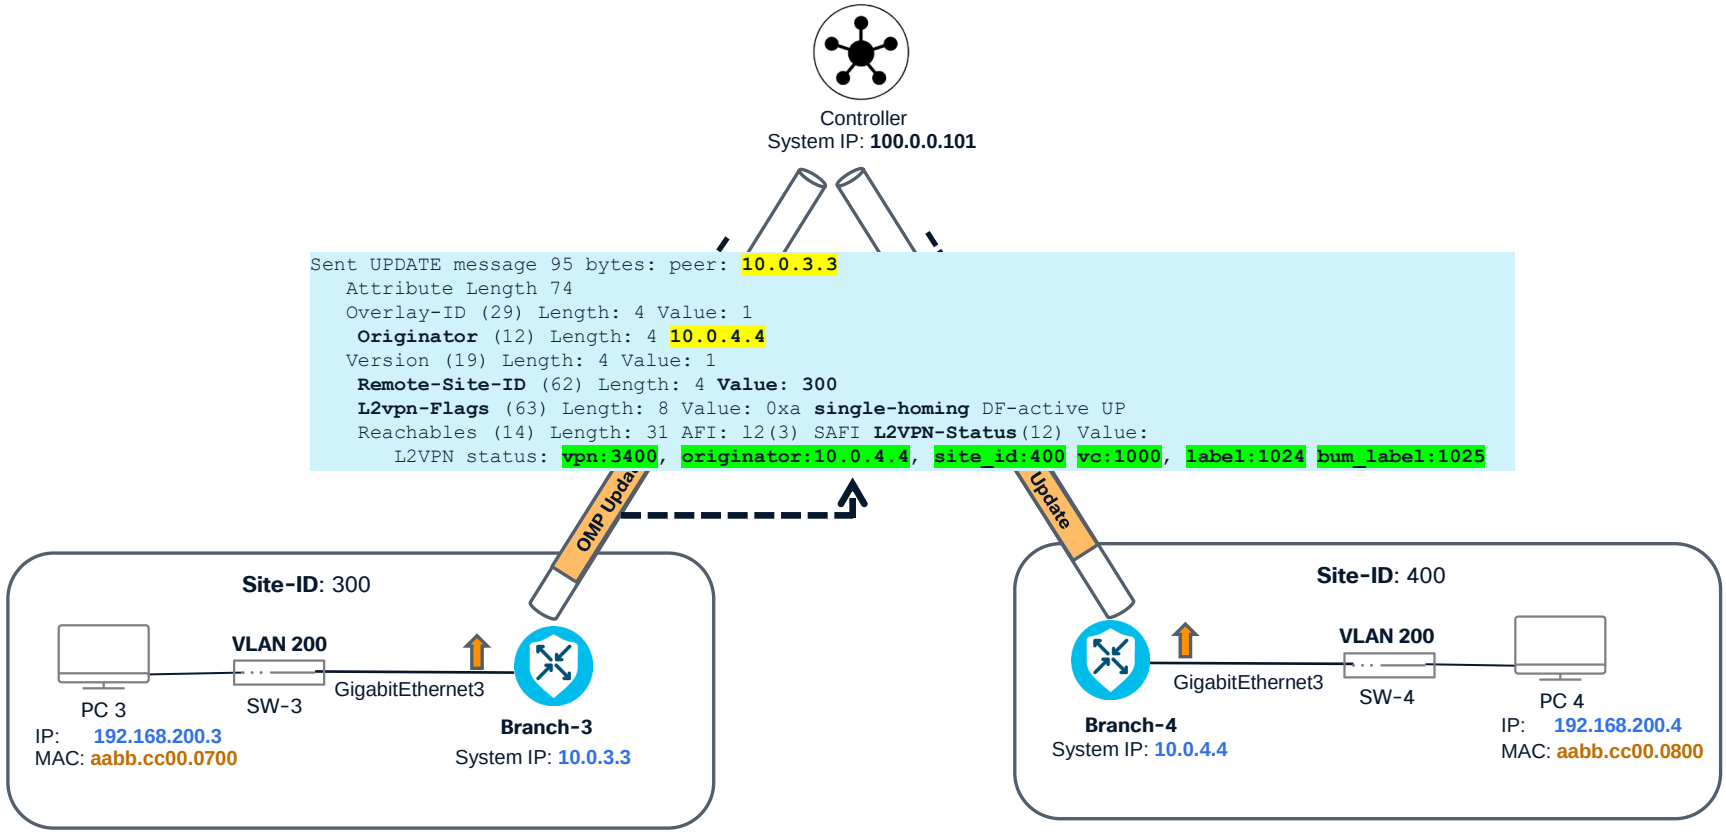

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 6/8

```
Sent UPDATE message 96 bytes: peer: 10.0.3.3
   Attribute Length 75
   Site-ID (9) Length: 4 Value: 400
   Overlay-ID (29) Length: 4 Value: 1
    TLOC (1) Length: 6 10.0.4.4 : mpls : ipsec
    Originator (12) Length: 4 10.0.4.4
    Remote-Site-ID (62) Length: 4 Value: 300
    Reachables (14) Length: 34 AFI: l2(3) SAFI vroute(1) Value:
      originator: 10.0.4.4 site-id: 400 vpn: 3400 vc: 1000 path-id: 2 label: 1024 route-type: vpn

 Sent UPDATE message 96 bytes: peer: 10.0.3.3
    Attribute Length 75
    Site-ID (9) Length: 4 Value: 400
    Overlay-ID (29) Length: 4 Value: 1
     TLOC (1) Length: 6 10.0.4.4 : biz-internet : ipsec
     Originator (12) Length: 4 10.0.4.4
     Remote-Site-ID (62) Length: 4 Value: 300
     Reachables (14) Length: 34 AFI: l2(3) SAFI vroute(1) Value:
       originator: 10.0.4.4 site-id: 400 vpn: 3400 vc: 1000 path-id: 1 label: 1024 route-type: vpn
```

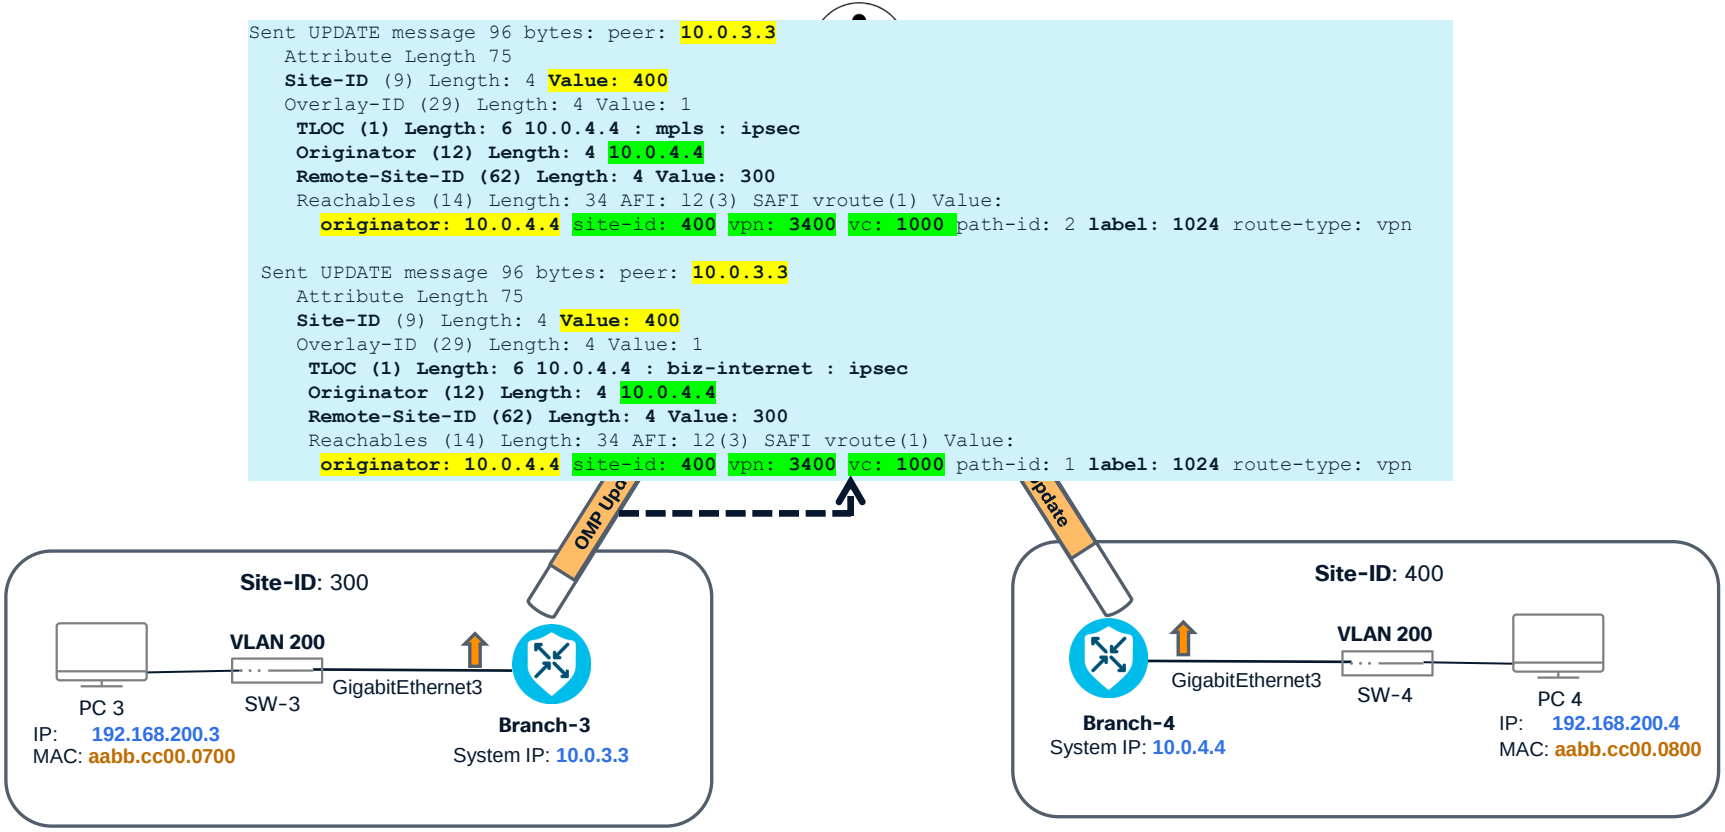

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 7/8

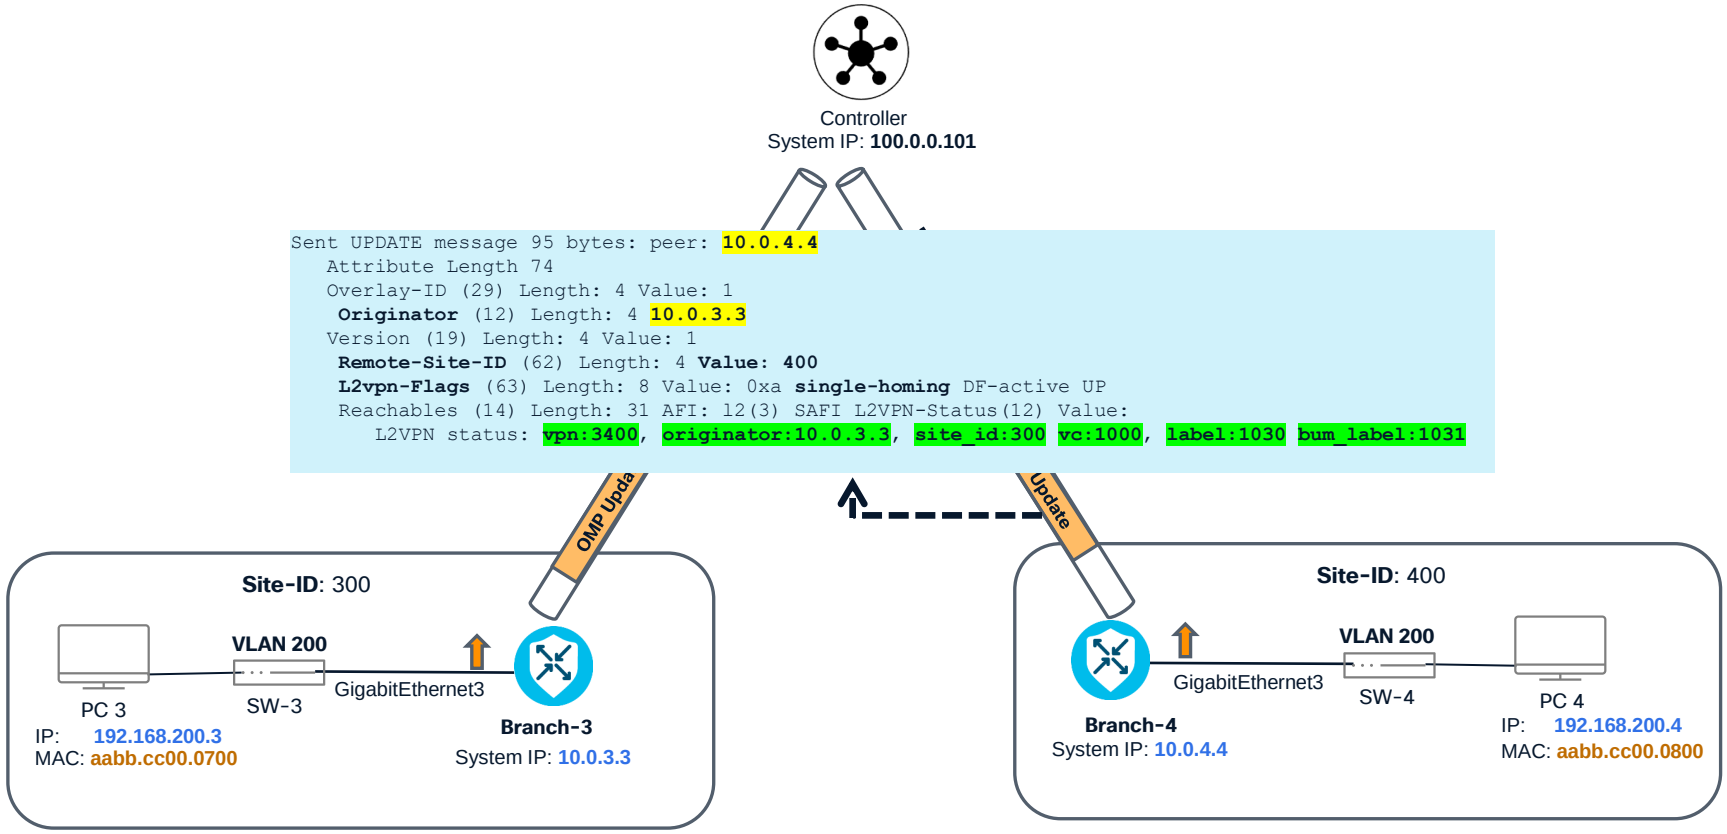

# How OMP Builds Point-to-Point (P2P) L2VPN Services? 8/8

```
Sent UPDATE message 96 bytes: peer: 10.0.4.4
   Attribute Length 75
   Site-ID (9) Length: 4 Value: 300
   Overlay-ID (29) Length: 4 Value: 1
    TLOC (1) Length: 6 10.0.3.3 : mpls : ipsec
    Originator (12) Length: 4 10.0.3.3
    Remote-Site-ID (62) Length: 4 Value: 400
    Reachables (14) Length: 34 AFI: l2(3) SAFI vroute(1) Value:
      originator: 10.0.3.3 site-id: 300 vpn: 3400 vc: 1000 path-id: 2 label: 1030 route-type: vpn

 Sent UPDATE message 96 bytes: peer: 10.0.4.4
    Attribute Length 75
    Site-ID (9) Length: 4 Value: 300
    Overlay-ID (29) Length: 4 Value: 1
     TLOC (1) Length: 6 10.0.3.3 : biz-internet : ipsec
     Originator (12) Length: 4 10.0.3.3
     Remote-Site-ID (62) Length: 4 Value: 400
     Reachables (14) Length: 34 AFI: l2(3) SAFI vroute(1) Value:
       originator: 10.0.3.3 site-id: 300 vpn: 3400 vc: 1000 path-id: 1 label: 1030 route-type: vpn
```

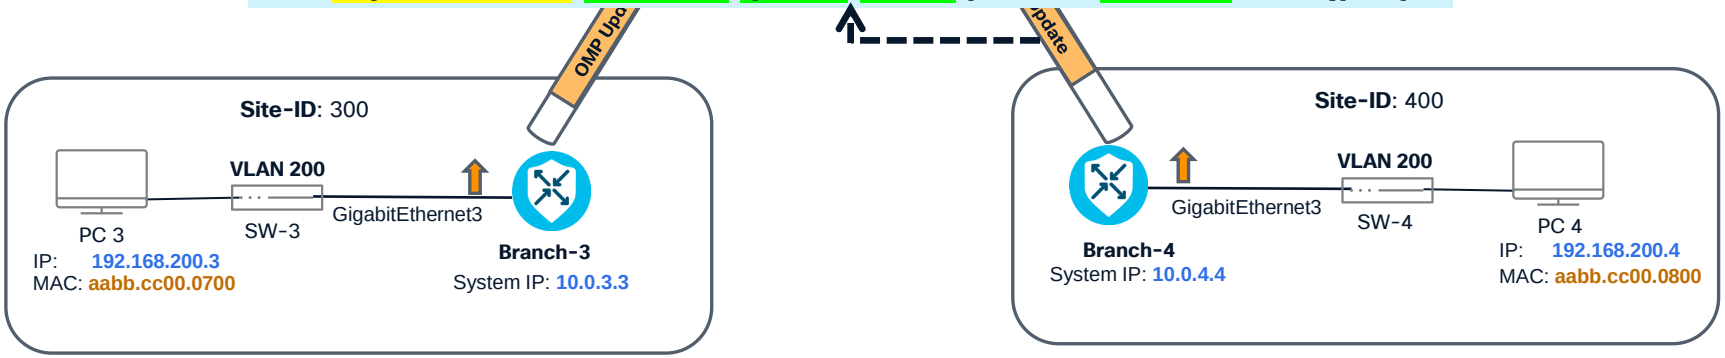

# L2VPN Status and OMP Control-Plane Validation

```
Controller01# show omp l2-statuses | begin VPN | tab

                                                                                L2VPN    HOMING   DF       BUM
VPN    VC ID       ORIGINATOR      SITE ID     FROM PEER       PATH ID          STATUS   TYPE     STATE    LABEL    LABEL    STATUS
-------------------------------------------------------------------------------------------------------------------------------
3400   1000        10.0.3.3        300         10.0.3.3        0                up       single   active   1031     1030     C,R
3400   1000        10.0.4.4        400         10.0.4.4        0                up       single   active   1025     1024     C,R
```

# How traffic flows between P2P L2VPN? 1/2

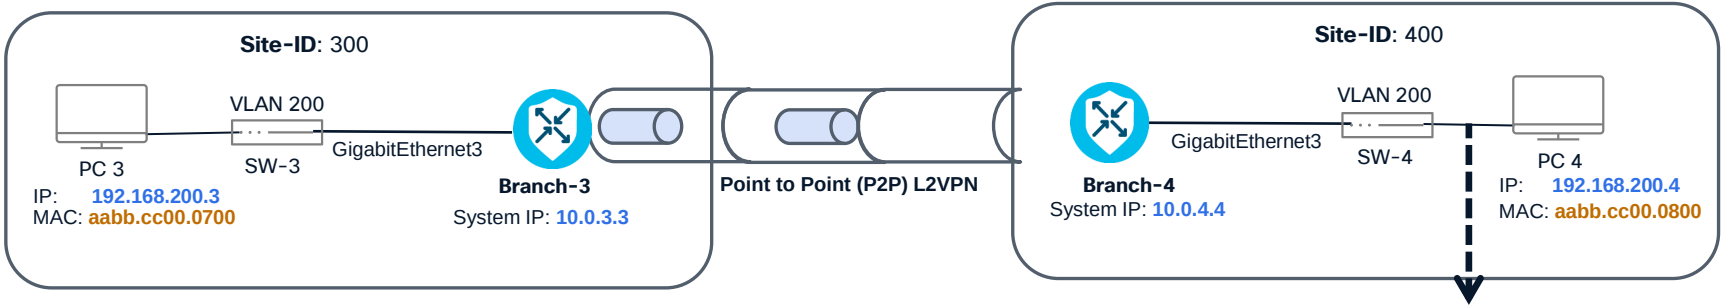


```
> Frame 17: Packet, 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
> Ethernet II, Src: aa:bb:cc:00:07:00 (aa:bb:cc:00:07:00), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
v Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: aa:bb:cc:00:07:00 (aa:bb:cc:00:07:00)
    Sender IP address: 192.168.200.3
    Target MAC address: 00:00:00_00:00:00 (00:00:00:00:00:00)
    Target IP address: 192.168.200.4
```

# How traffic flows between P2P L2VPN? 2/2

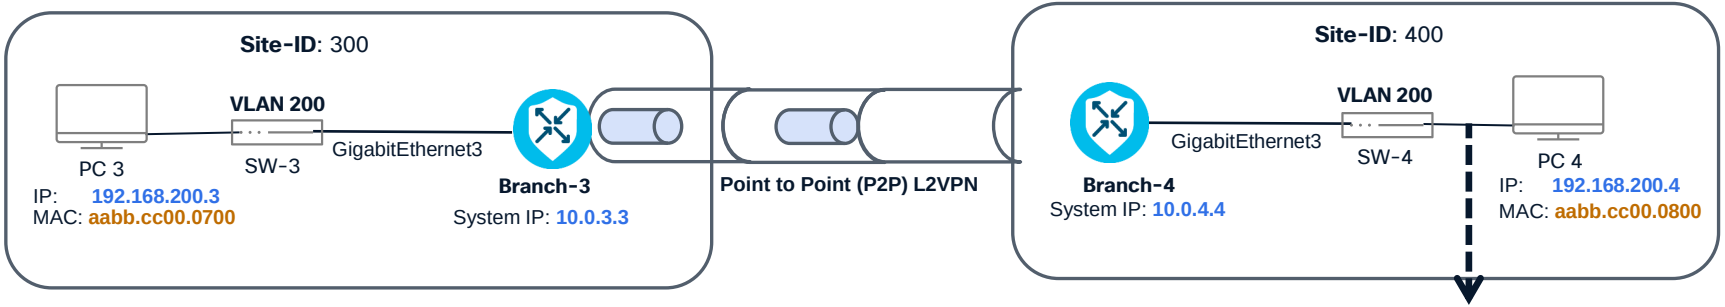


```
PC-3#ping 192.168.200.4

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.200.4, timeout is
2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max =
6/7/9 ms
```

```
> Frame 18: Packet, 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
> Ethernet II, Src: aa:bb:cc:00:08:00 (aa:bb:cc:00:08:00), Dst: aa:bb:cc:00:07:00 (aa:bb:cc:00:07:00)
v Address Resolution Protocol (reply)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: reply (2)
    Sender MAC address: aa:bb:cc:00:08:00 (aa:bb:cc:00:08:00)
    Sender IP address: 192.168.200.4
    Target MAC address: aa:bb:cc:00:07:00 (aa:bb:cc:00:07:00)
    Target IP address: 192.168.200.3
```

# L2VPN Point-to-Multipoint (P2MP)

# Point-to-Multipoint(P2MP) L2VPN Demo Topology

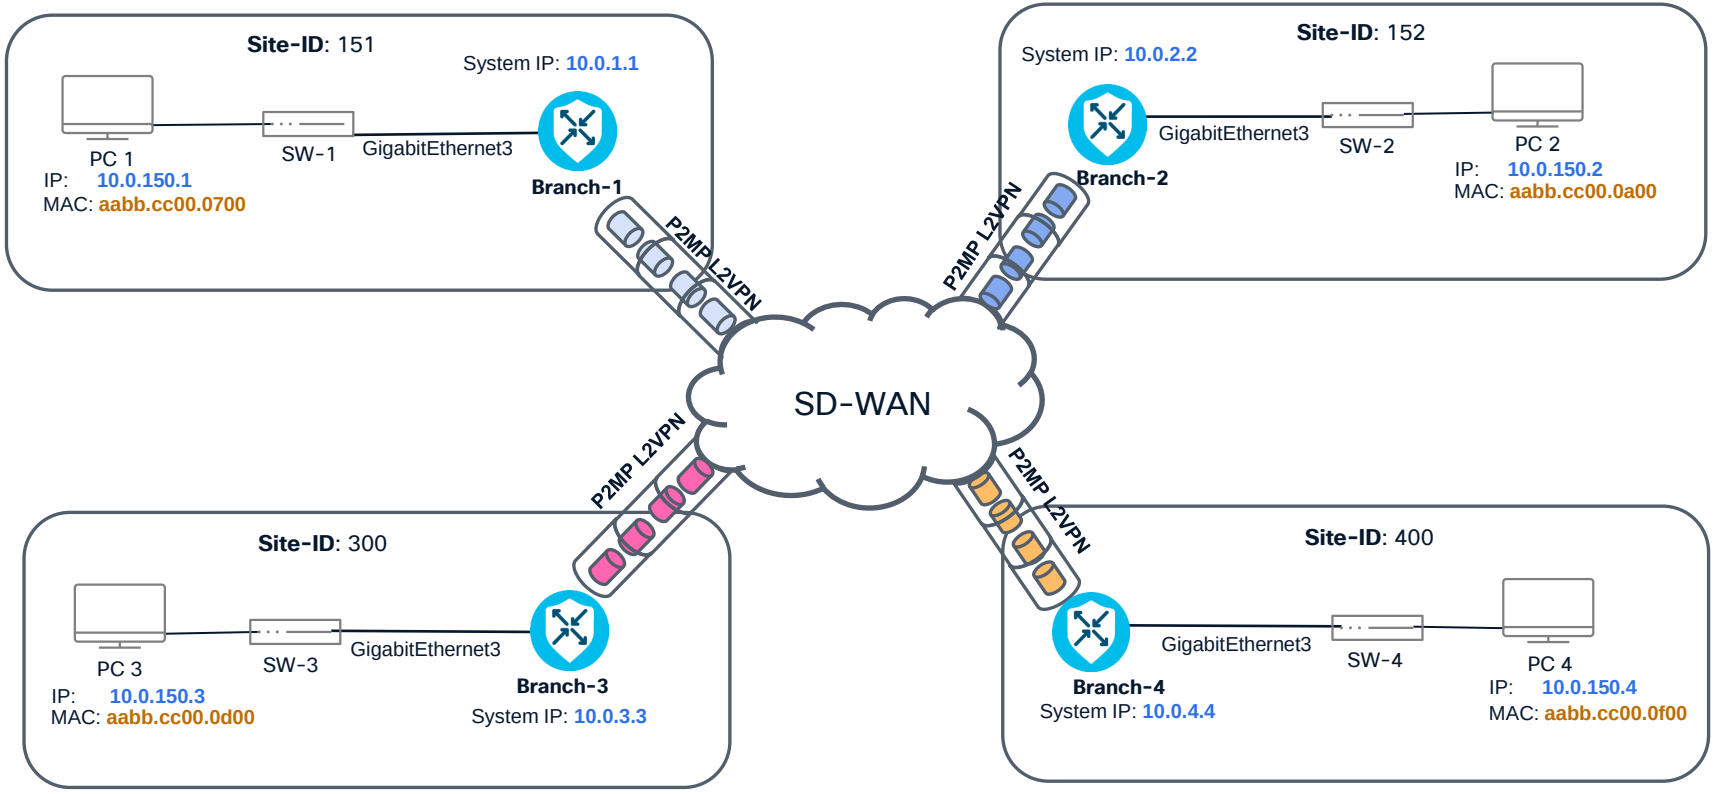

# How is Point-to-Multipoint (P2MP) L2VPN configured?

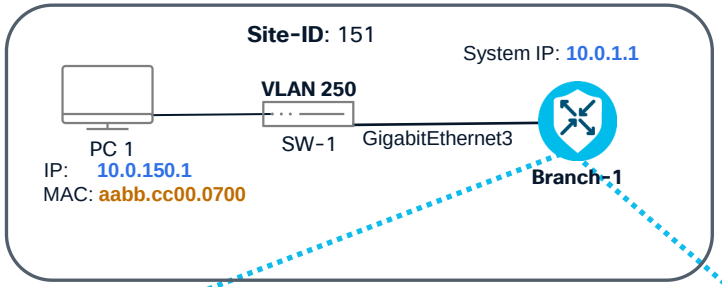


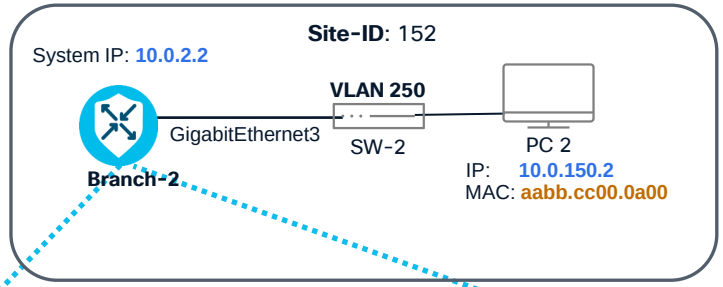


```
l2vpn sdwan instance 12340 multipoint
!
bridge-domain 950
 member GigabitEthernet3 service-instance 1510
 member sdwan-instance 12340 vc-id 2000 single-homing
!
interface GigabitEthernet3
 service instance 1510 ethernet
  encapsulation dot1q 250
```

```
l2vpn sdwan instance 12340 multipoint
!
bridge-domain 950
 member GigabitEthernet3 service-instance 1510
 member sdwan-instance 12340 vc-id 2000 single-homing
!
interface GigabitEthernet3
 service instance 1510 ethernet
  encapsulation dot1q 250
```

# How OMP Builds Point-to-MultiPoint (P2MP) L2VPN Services? 1/6

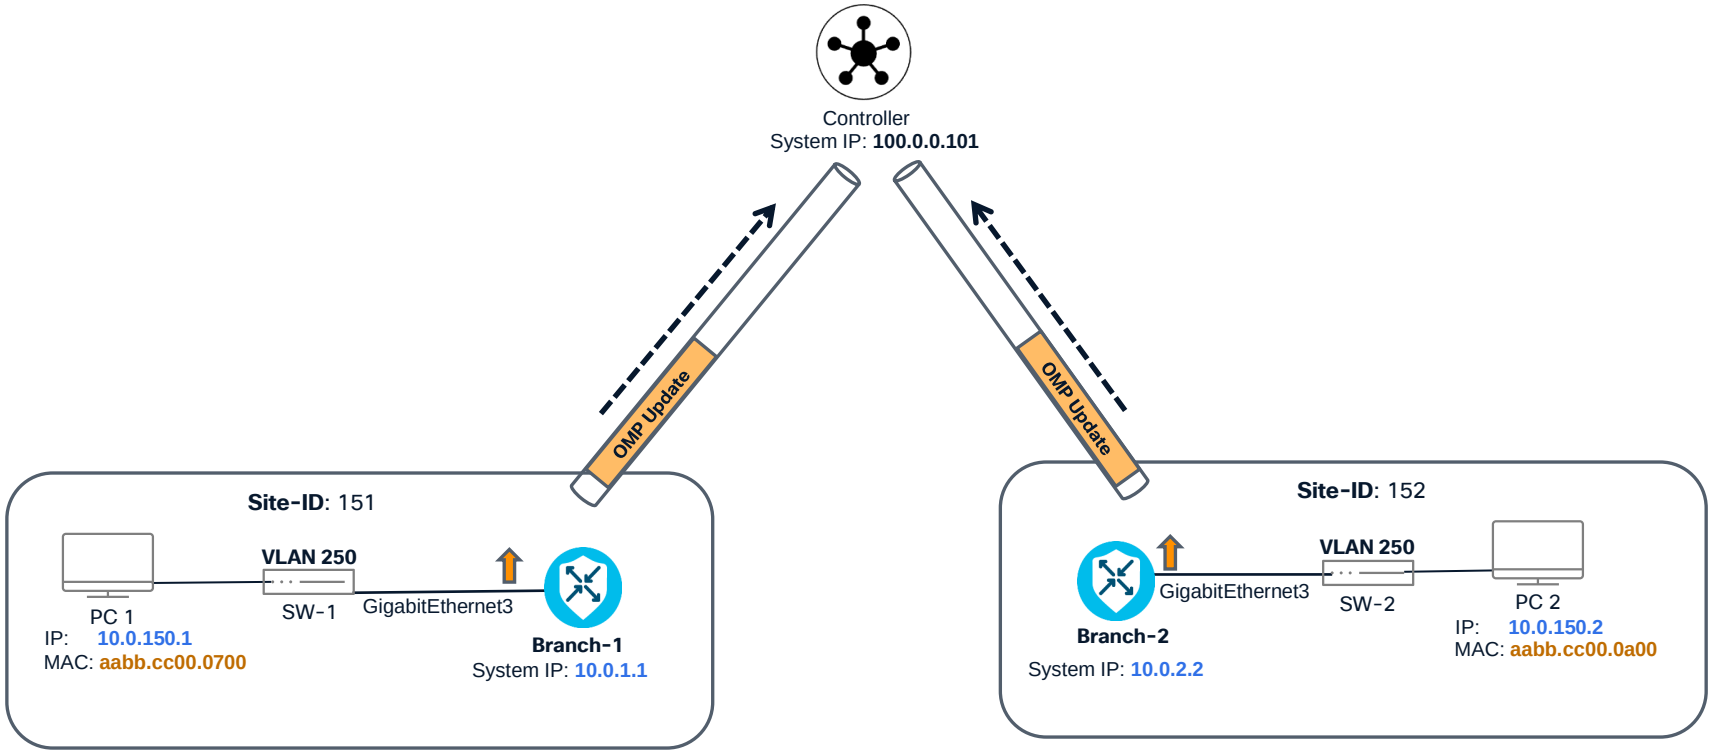

# How OMP Builds Point-to-MultiPoint (P2MP) L2VPN Services? 2/6

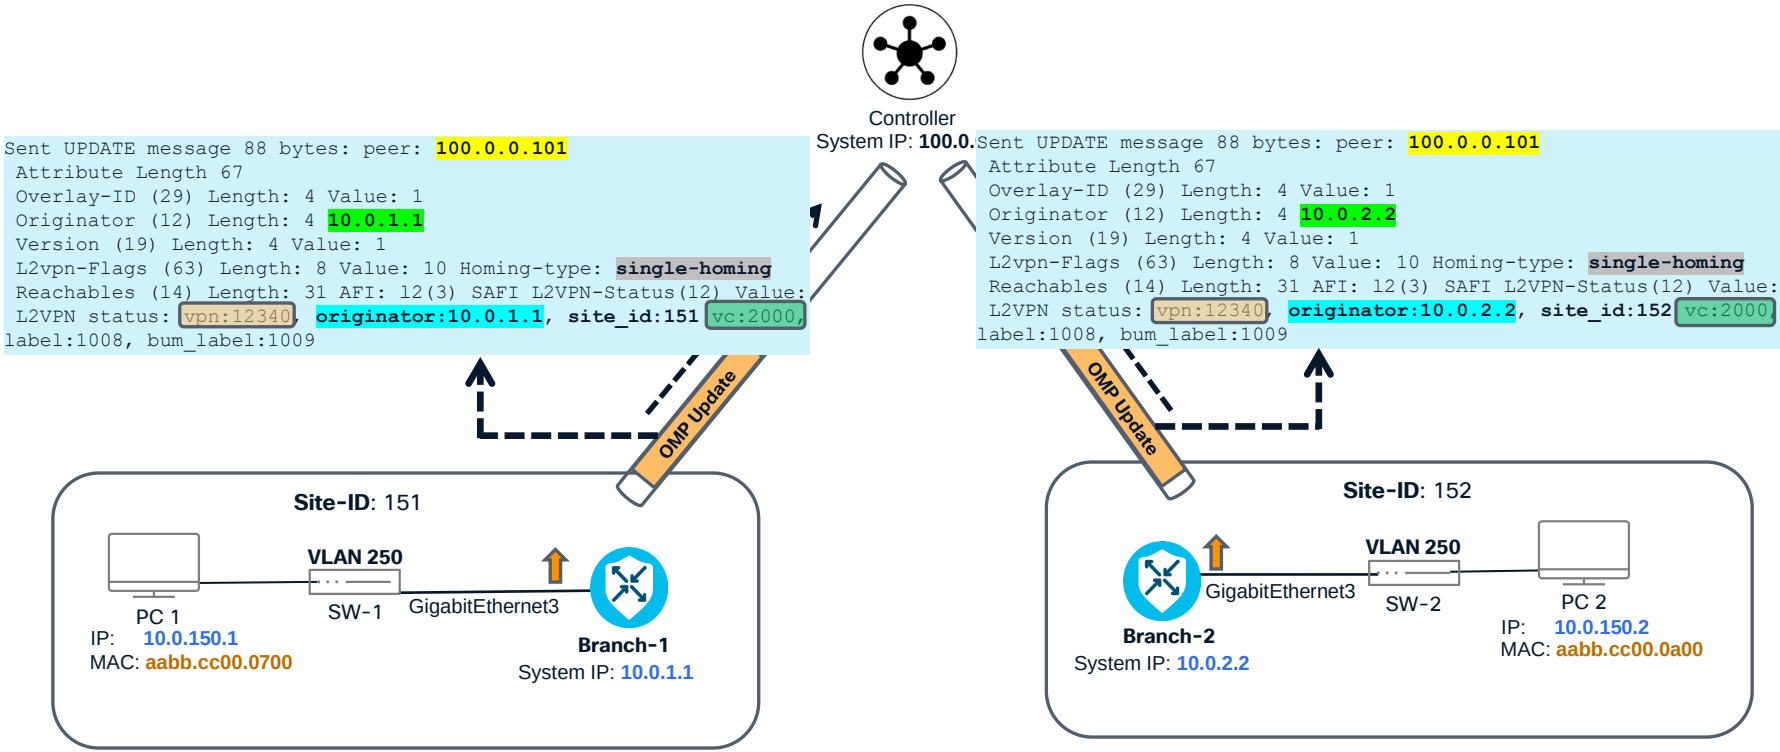

# How OMP Builds Point-to-MultiPoint (P2MP) L2VPN Services? 3/6

```
Sent UPDATE message 79 bytes: peer: 100.0.0.101
 Attribute Length 58
 TLOC (1) Length: 6 10.0.1.1 : mpls : ipsec
 Overlay-ID (29) Length: 4 Value: 1
 Originator (12) Length: 4 10.0.1.1
 Reachables (14) Length: 31 AFI: 12(3) SAFI service(3) Value:
 site-id: 151 vpn: 12340 vc: 2000 service: vpn(0) vpn-type 2
path-id: 66 label: 1008

Sent UPDATE message 79 bytes: peer: 100.0.0.101
 Attribute Length 58
 TLOC (1) Length: 6 10.0.1.1 : biz-internet : ipsec
 Overlay-ID (29) Length: 4 Value: 1
 Originator (12) Length: 4 10.0.1.1
 Reachables (14) Length: 31 AFI: 12(3) SAFI service(3) Value:
 site-id: 151 vpn: 12340 vc: 2000 service: vpn(0) vpn-type 2
path-id: 68 label: 1008
```

```
Sent UPDATE message 79 bytes: peer: 100.0.0.101
 Attribute Length 58
 TLOC (1) Length: 6 10.0.2.2 : mpls : ipsec
 Overlay-ID (29) Length: 4 Value: 1
 Originator (12) Length: 4 10.0.2.2
 Reachables (14) Length: 31 AFI: 12(3) SAFI service(3) Value:
 site-id: 152 vpn: 12340 vc: 2000 service: vpn(0) vpn-type 2
path-id: 66 label: 1008

Sent UPDATE message 79 bytes: peer: 100.0.0.101
 Attribute Length 58
 TLOC (1) Length: 6 10.0.2.2 : biz-internet : ipsec
 Overlay-ID (29) Length: 4 Value: 1
 Originator (12) Length: 4 10.0.2.2
 Reachables (14) Length: 31 AFI: 12(3) SAFI service(3) Value:
 site-id: 152 vpn: 12340 vc: 2000 service: vpn(0) vpn-type 2
path-id: 68 label: 1008
```

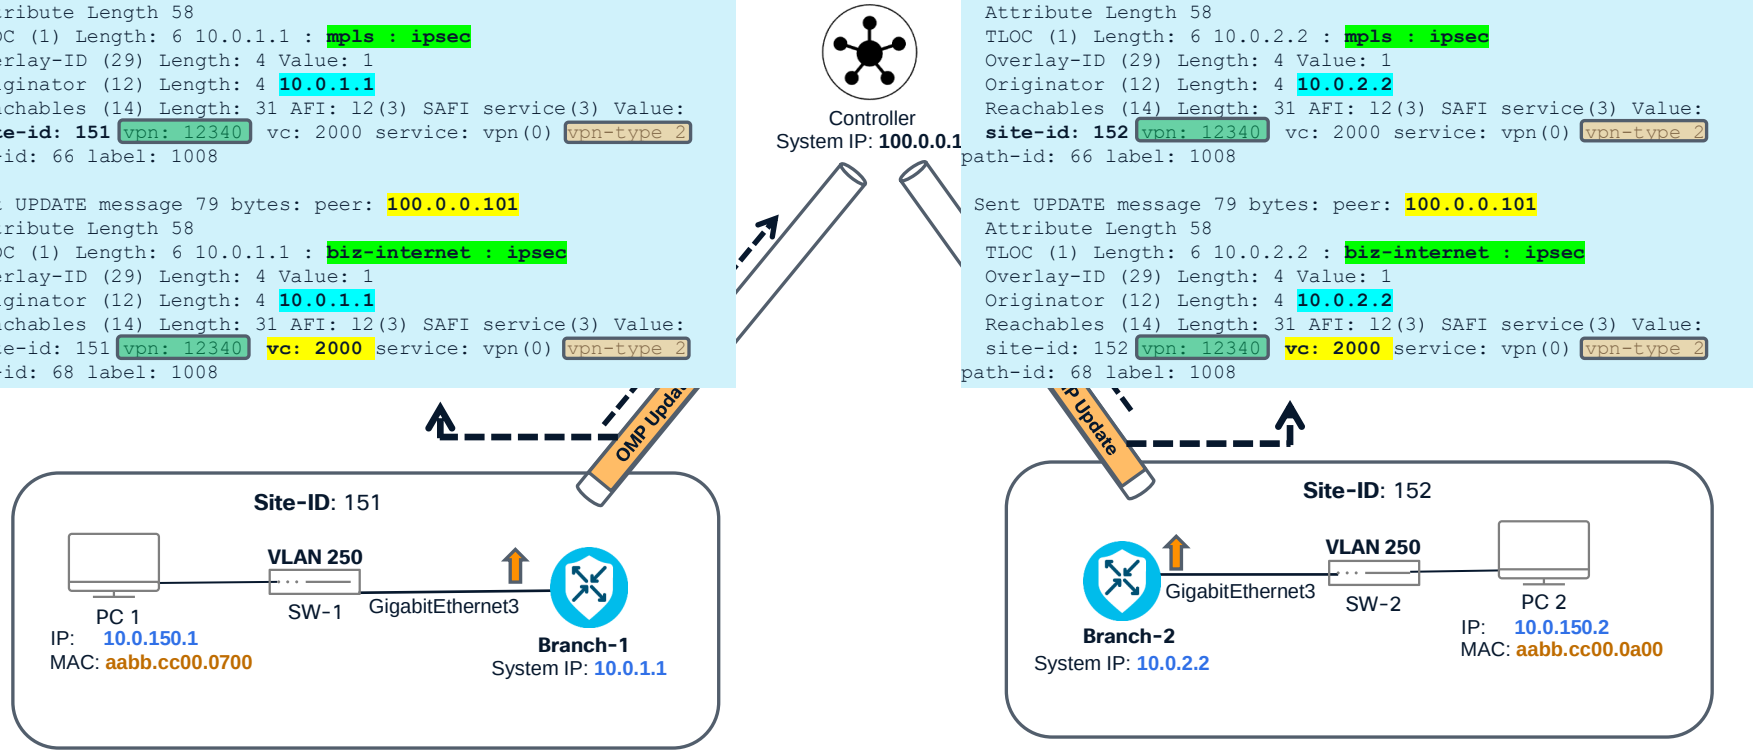

# How OMP Builds Point-to-MultiPoint (P2MP) L2VPN Services? 4/6

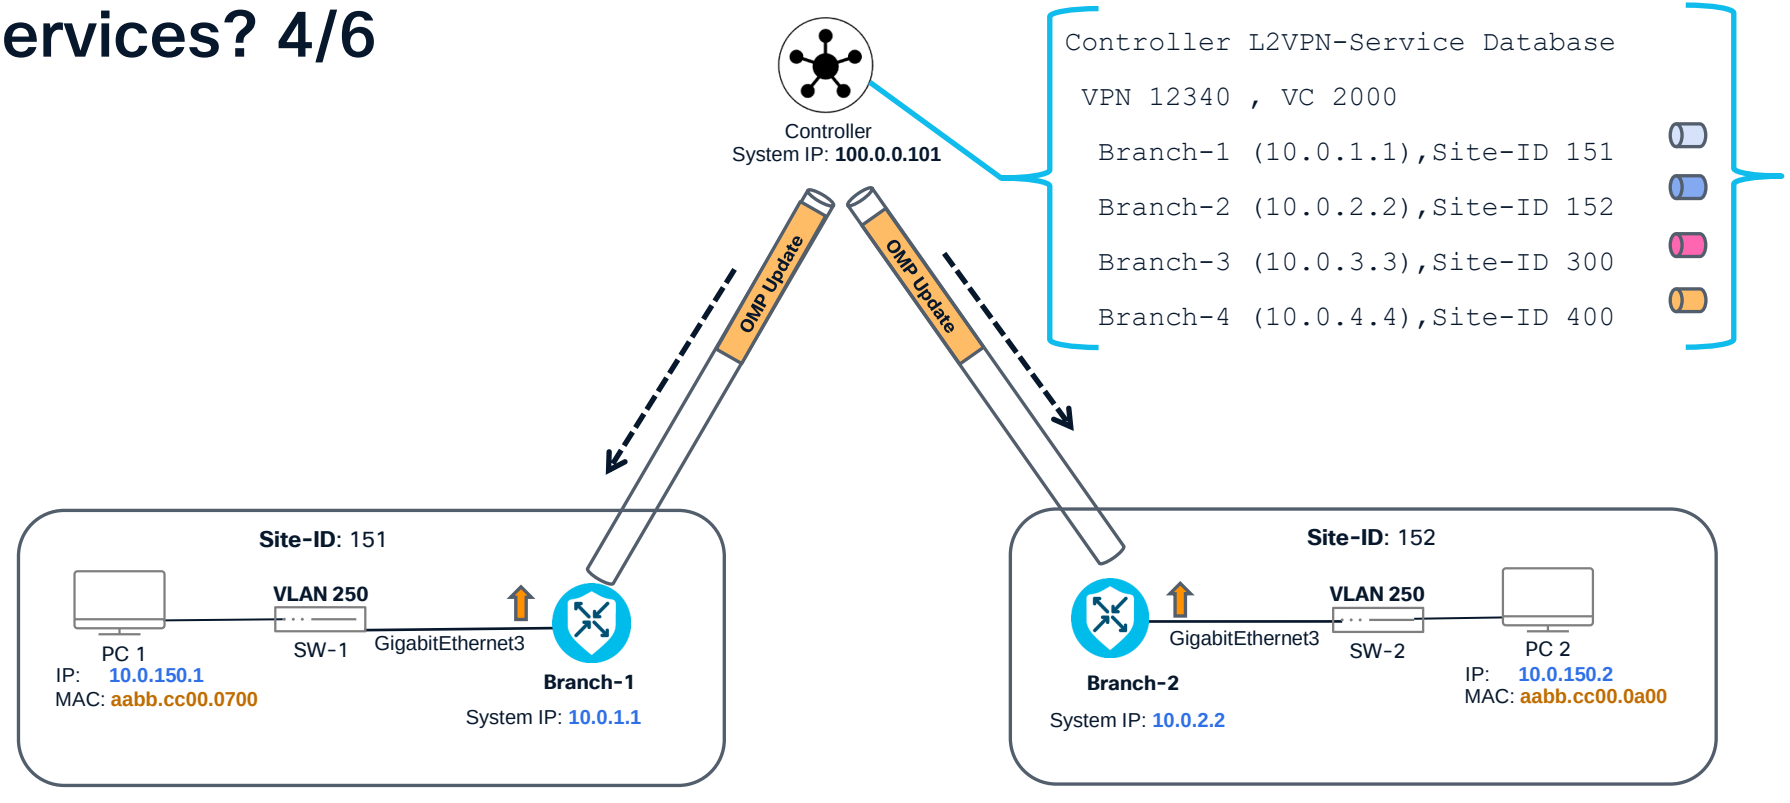

# How OMP Builds Point-to-MultiPoint (P2MP) L2VPN Services? 5/6

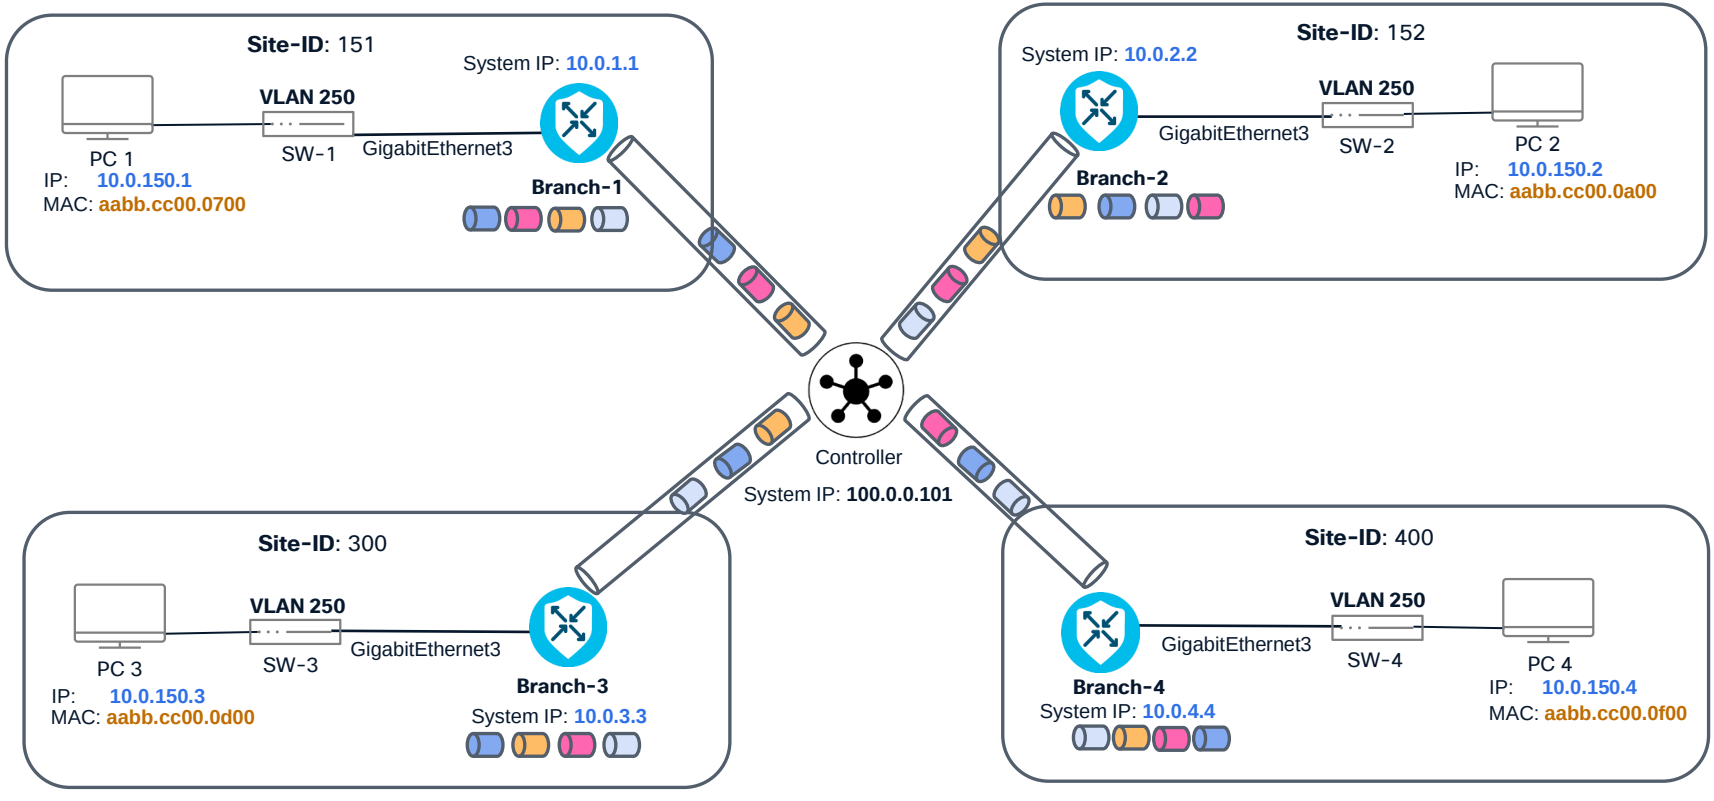

# How OMP Builds Point-to-MultiPoint (P2MP) L2VPN Services? 6/6

```
Branch-1# show l2vpn sdwan instance 12340 vc-id 2000 all

Bridge-Domain: 950

Local state: up

Homing type: single-homing

Homing role: active

Dual-Homing peer: no peer

Service EFP: Gi3 ServInst 1510

Remote Site: 152

  Originator       State          Up/Down    Homing-Role  Hub/Spoke  Color           Encap  Installed
  10.0.2.2         up             1d11h      active       no         mpls            ipsec  yes(1008)
                                                                     biz-internet    ipsec  yes(1008)

Remote Site: 300

  Originator       State          Up/Down    Homing-Role  Hub/Spoke  Color           Encap  Installed
  10.0.3.3         up             1d11h      active       no         mpls            ipsec  yes(1008)
                                                                     biz-internet    ipsec  yes(1008)

Remote Site: 400

  Originator       State          Up/Down    Homing-Role  Hub/Spoke  Color           Encap  Installed
  10.0.4.4         up             1d11h      active       no         mpls            ipsec  yes(1008)
                                                                     biz-internet    ipsec  yes(1008)
```

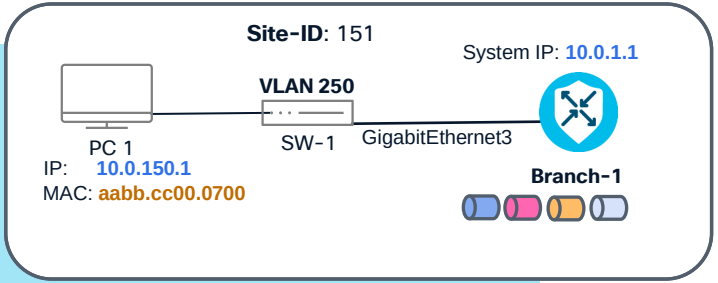

# How traffic flows between P2MP L2VPN? 1/9

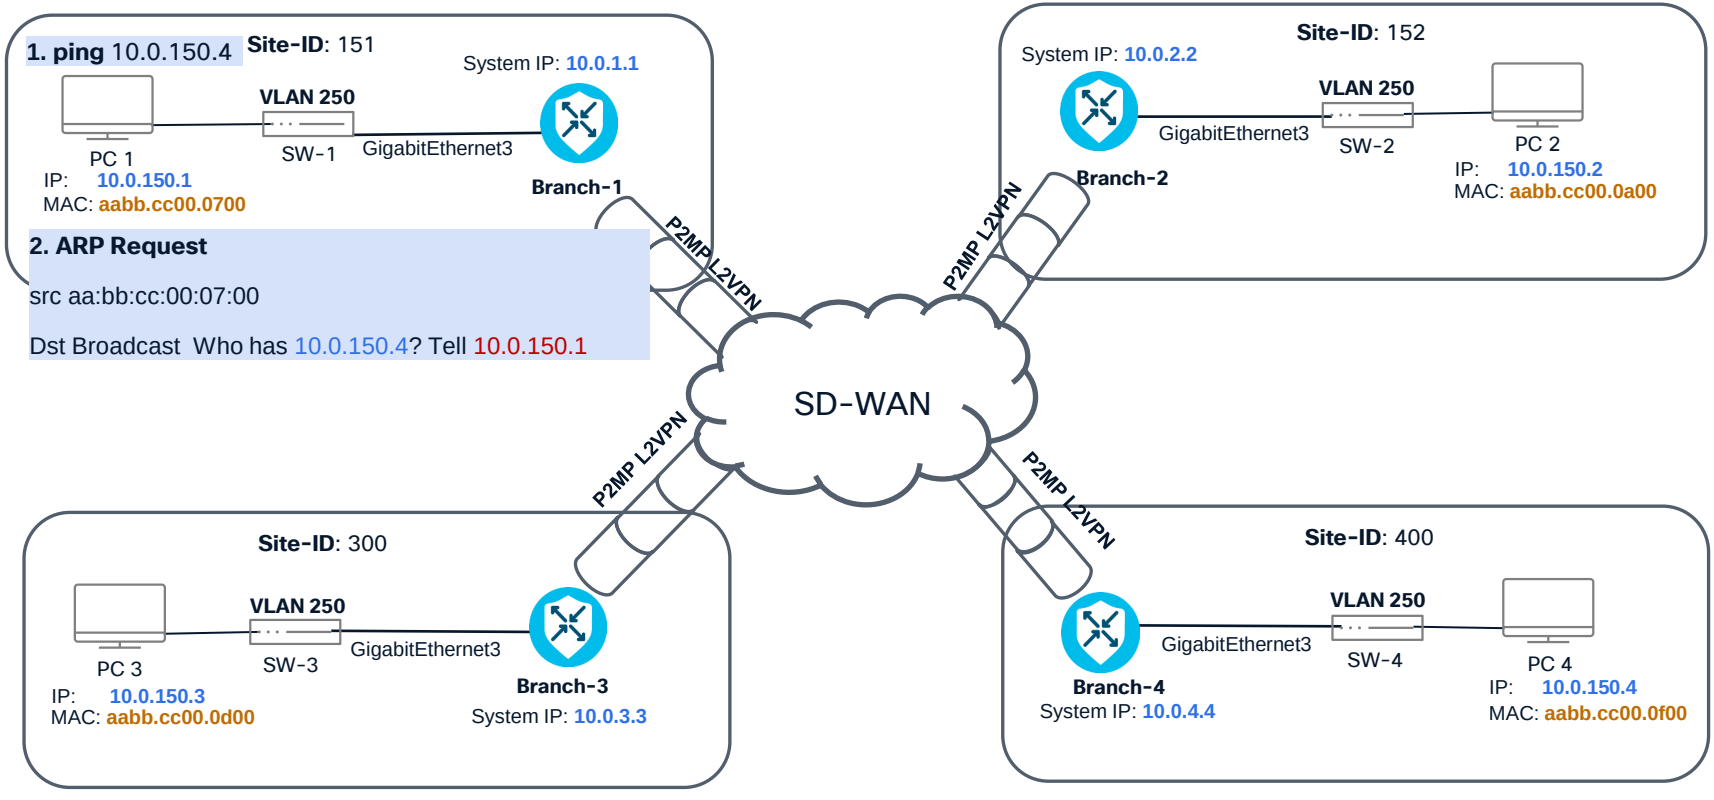

# How traffic flows between P2MP L2VPN? 2/9

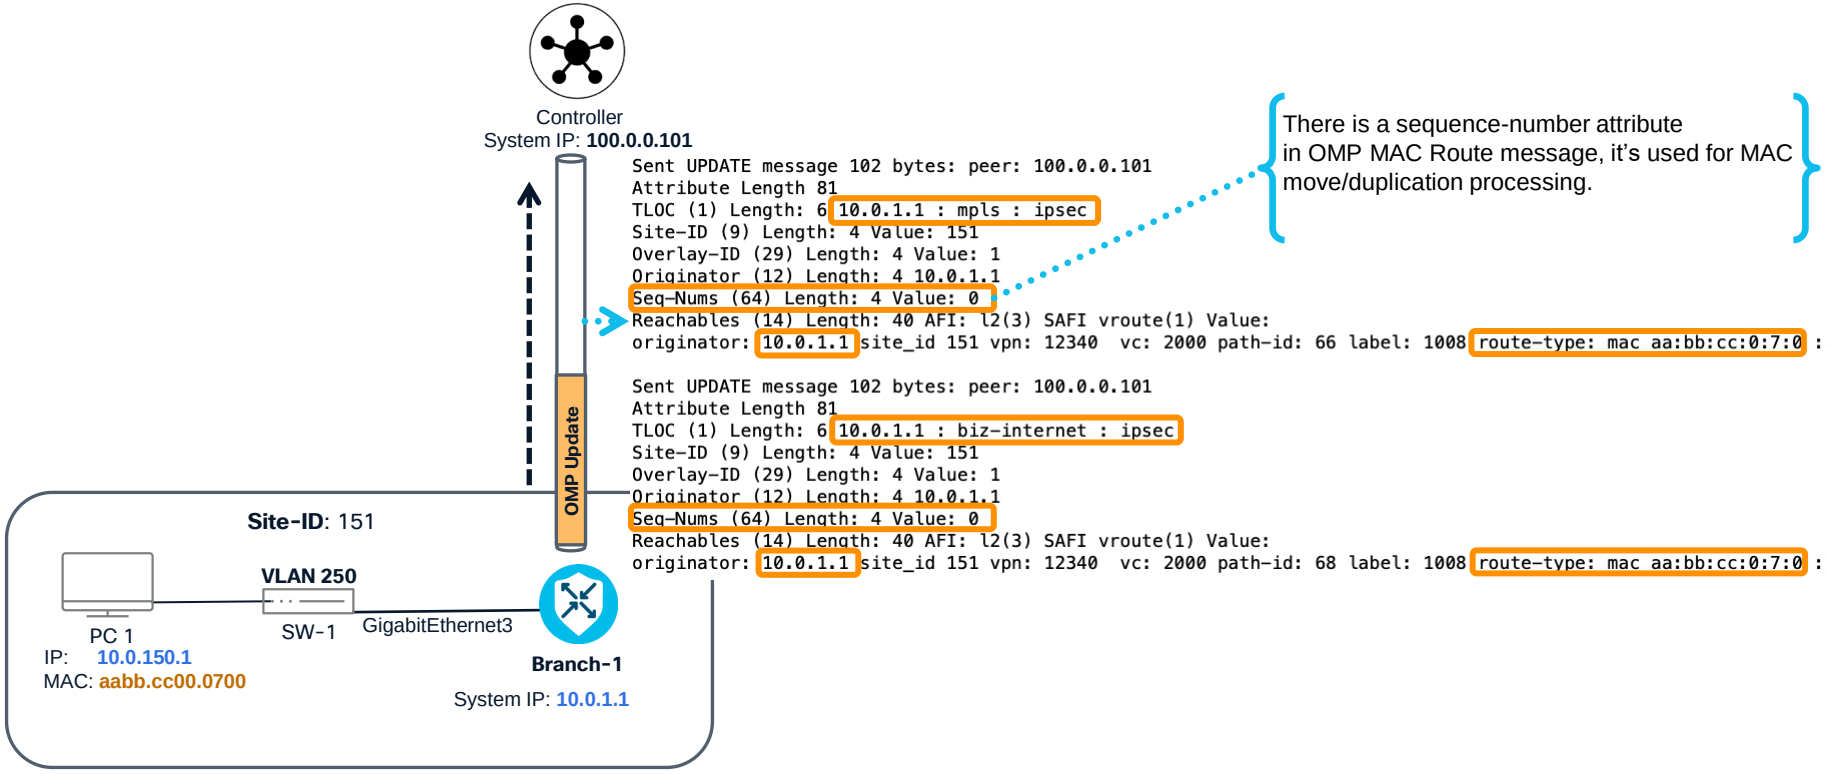

# How traffic flows between P2MP L2VPN? 3/9

Controller
System IP: **100.0.0.101**

```
Sent UPDATE message 102 bytes: peer: 10.0.2.2
   Attribute Length 81
   Site-ID (9) Length: 4 Value: 151
   Overlay-ID (29) Length: 4 Value: 1
    TLOC (1) Length: 6 10.0.1.1 : mls : ipsec
    Originator (12) Length: 4 10.0.1.1
    Seq-Nums (64) Length: 4 Value: 0
    Reachables (14) Length: 40 AFI: l2(3) SAFI vroute(1) Value:
        originator: 10.0.1.1 site-id: 151 vpn: 12340 vc: 2000 path-id: 1 label: 1008 route-type: mac aa:bb:cc:0:7:0 :
Sent UPDATE message 102 bytes: peer: 10.0.2.2
   Attribute Length 81
   Site-ID (9) Length: 4 Value: 151
   Overlay-ID (29) Length: 4 Value: 1
    TLOC (1) Length: 6 10.0.1.1 : biz-internet : ipsec
    Originator (12) Length: 4 10.0.1.1
    Seq-Nums (64) Length: 4 Value: 0
    Reachables (14) Length: 40 AFI: l2(3) SAFI vroute(1) Value:
]:      originator: 10.0.1.1 site-id: 151 vpn: 12340 vc: 2000 path-id: 2 label: 1008 route-type: mac aa:bb:cc:0:7:0 :
```

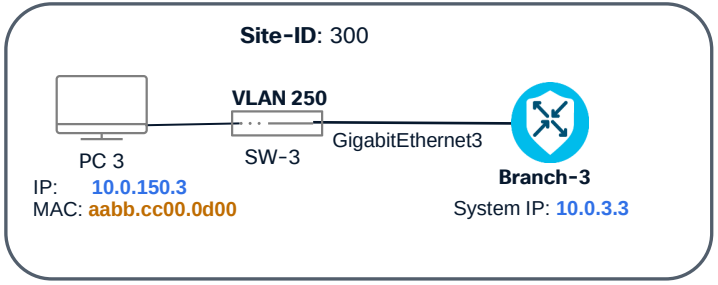


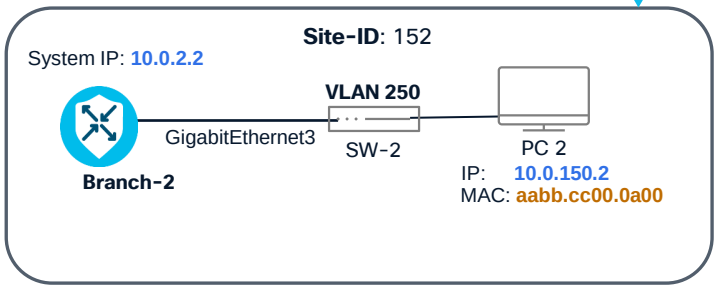

# How traffic flows between P2MP L2VPN? 4/9

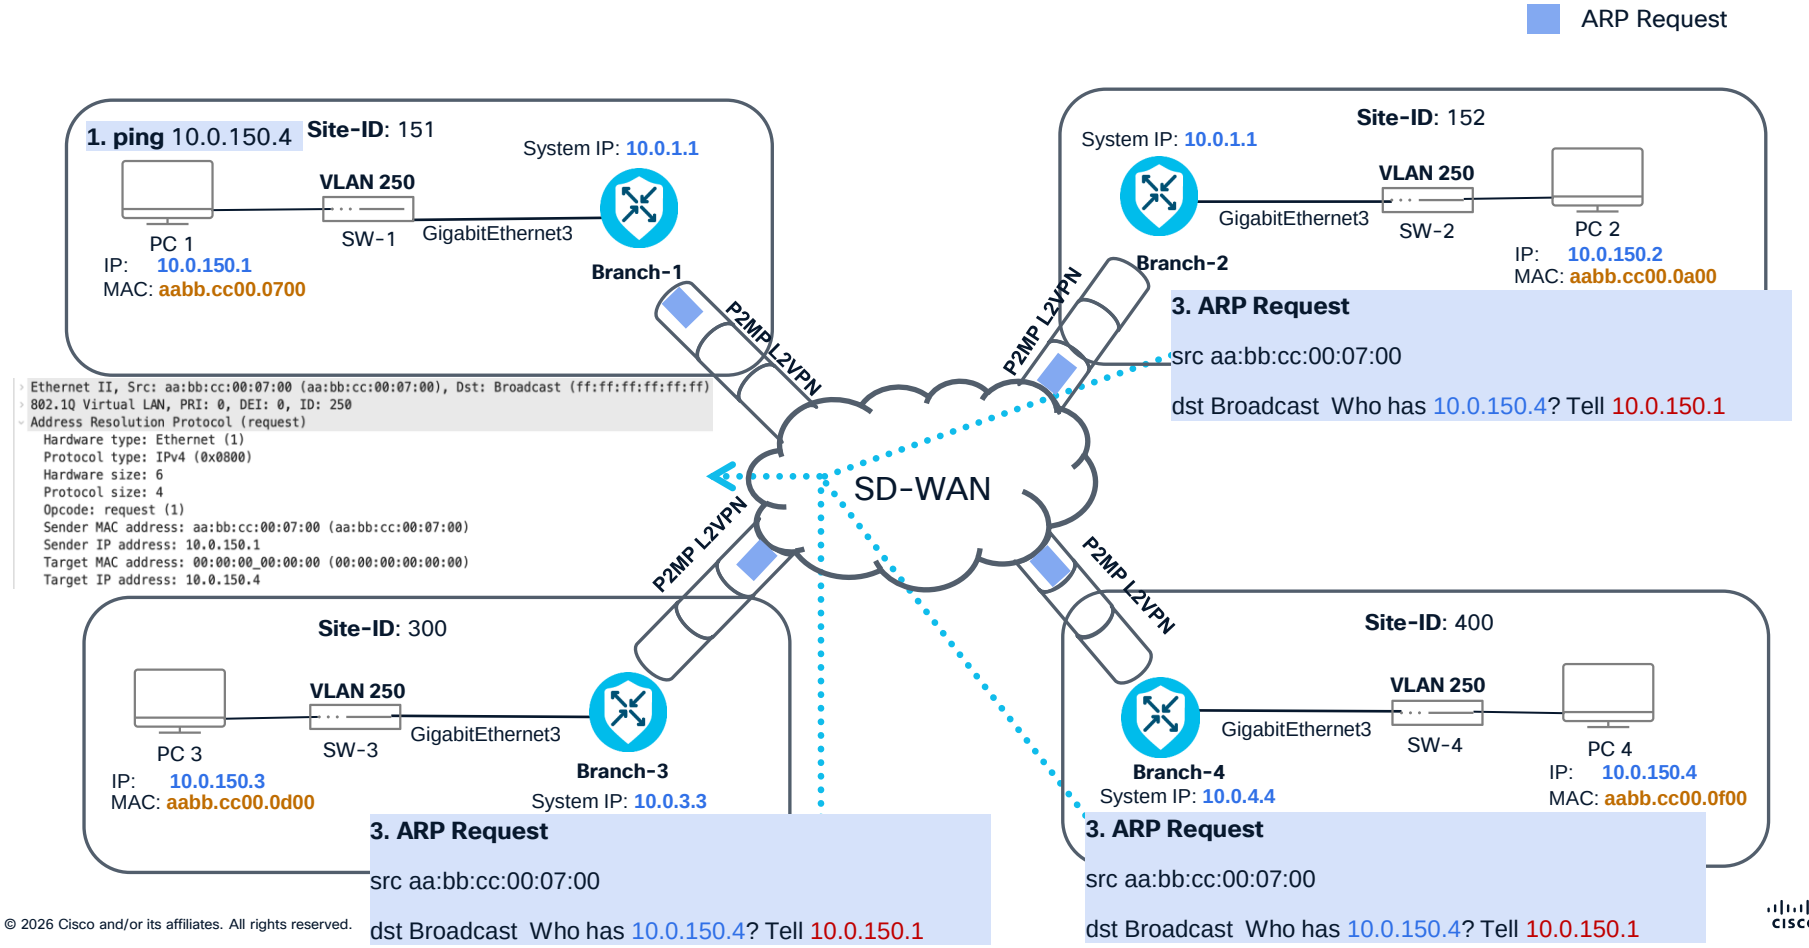

# How traffic flows between P2MP L2VPN? 5/9

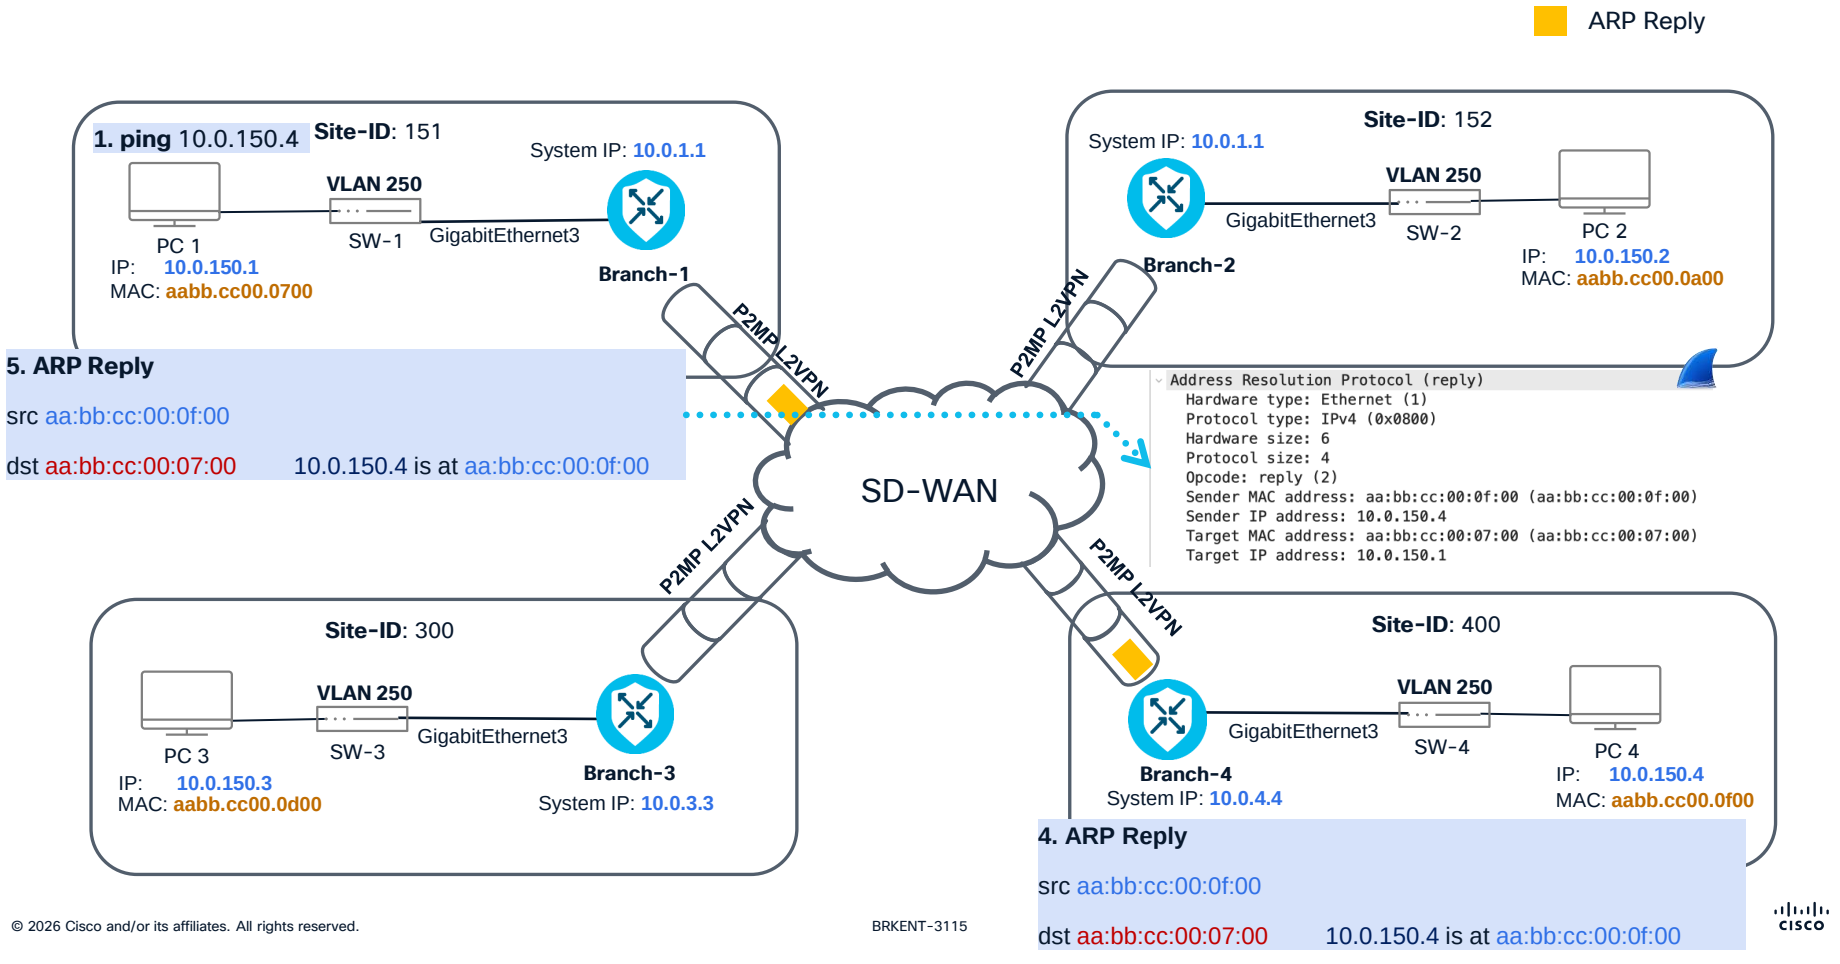

# How traffic flows between P2MP L2VPN? 6/9

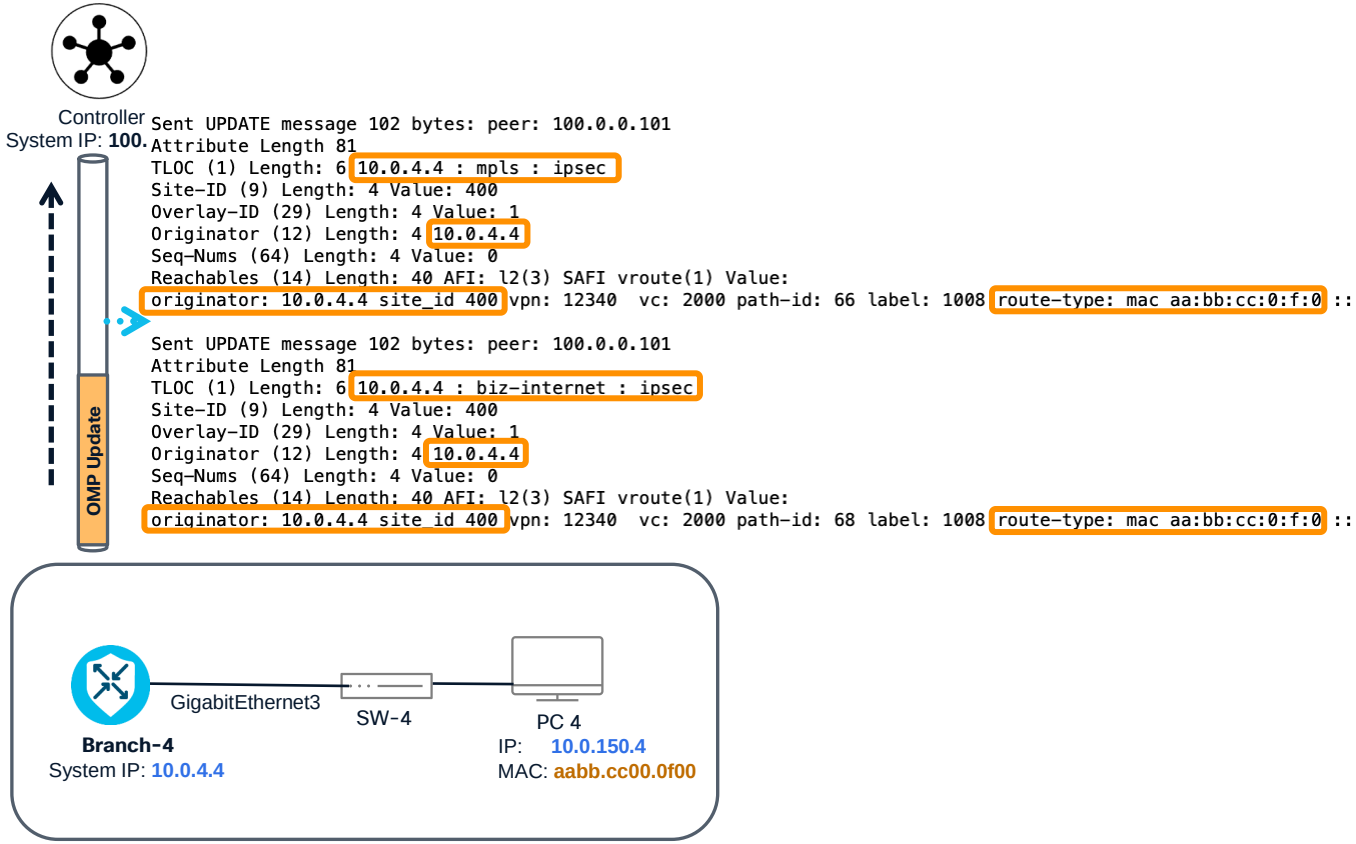

# How traffic flows between P2MP L2VPN? 7/9

Controller
System IP: **100.0.0.101**

```
Sent UPDATE message 102 bytes: peer: 10.0.2.2
   Attribute Length 81
   Site-ID (9) Length: 4 Value: 400
   Overlay-ID (29) Length: 4 Value: 1
    TLOC (1) Length: 6 10.0.4.4 : biz-internet : ipsec
    Originator (12) Length: 4 10.0.4.4
    Seq-Nums (64) Length: 4 Value: 0
    Reachables (14) Length: 40 AFI: l2(3) SAFI vroute(1) Value:
         originator: 10.0.4.4 site-id: 400 vpn: 12340 vc: 2000 path-id: 2 label: 1008 route-type: mac aa:bb:cc:0:f:0 ::
Sent UPDATE message 102 bytes: peer: 10.0.2.2
   Attribute Length 81
   Site-ID (9) Length: 4 Value: 400
   Overlay-ID (29) Length: 4 Value: 1
    TLOC (1) Length: 6 10.0.4.4 : mpls : ipsec
    Originator (12) Length: 4 10.0.4.4
    Seq-Nums (64) Length: 4 Value: 0
    Reachables (14) Length: 40 AFI: l2(3) SAFI vroute(1) Value:
         originator: 10.0.4.4 site-id: 400 vpn: 12340 vc: 2000 path-id: 1 label: 1008 route-type: mac aa:bb:cc:0:f:0 ::
```

**Site-ID**: 300

PC 3
IP: 10.0.150.3
MAC: aabb.cc00.0d00

VLAN 250
SW-3
GigabitEthernet3

**Branch-3**
System IP: 10.0.3.3

**Site-ID**: 152

**Branch-2**
System IP: 10.0.2.2

GigabitEthernet3
VLAN 250
SW-2

PC 2
IP: 10.0.150.2
MAC: aabb.cc00.0a00

# How traffic flows between P2MP L2VPN? 8/9

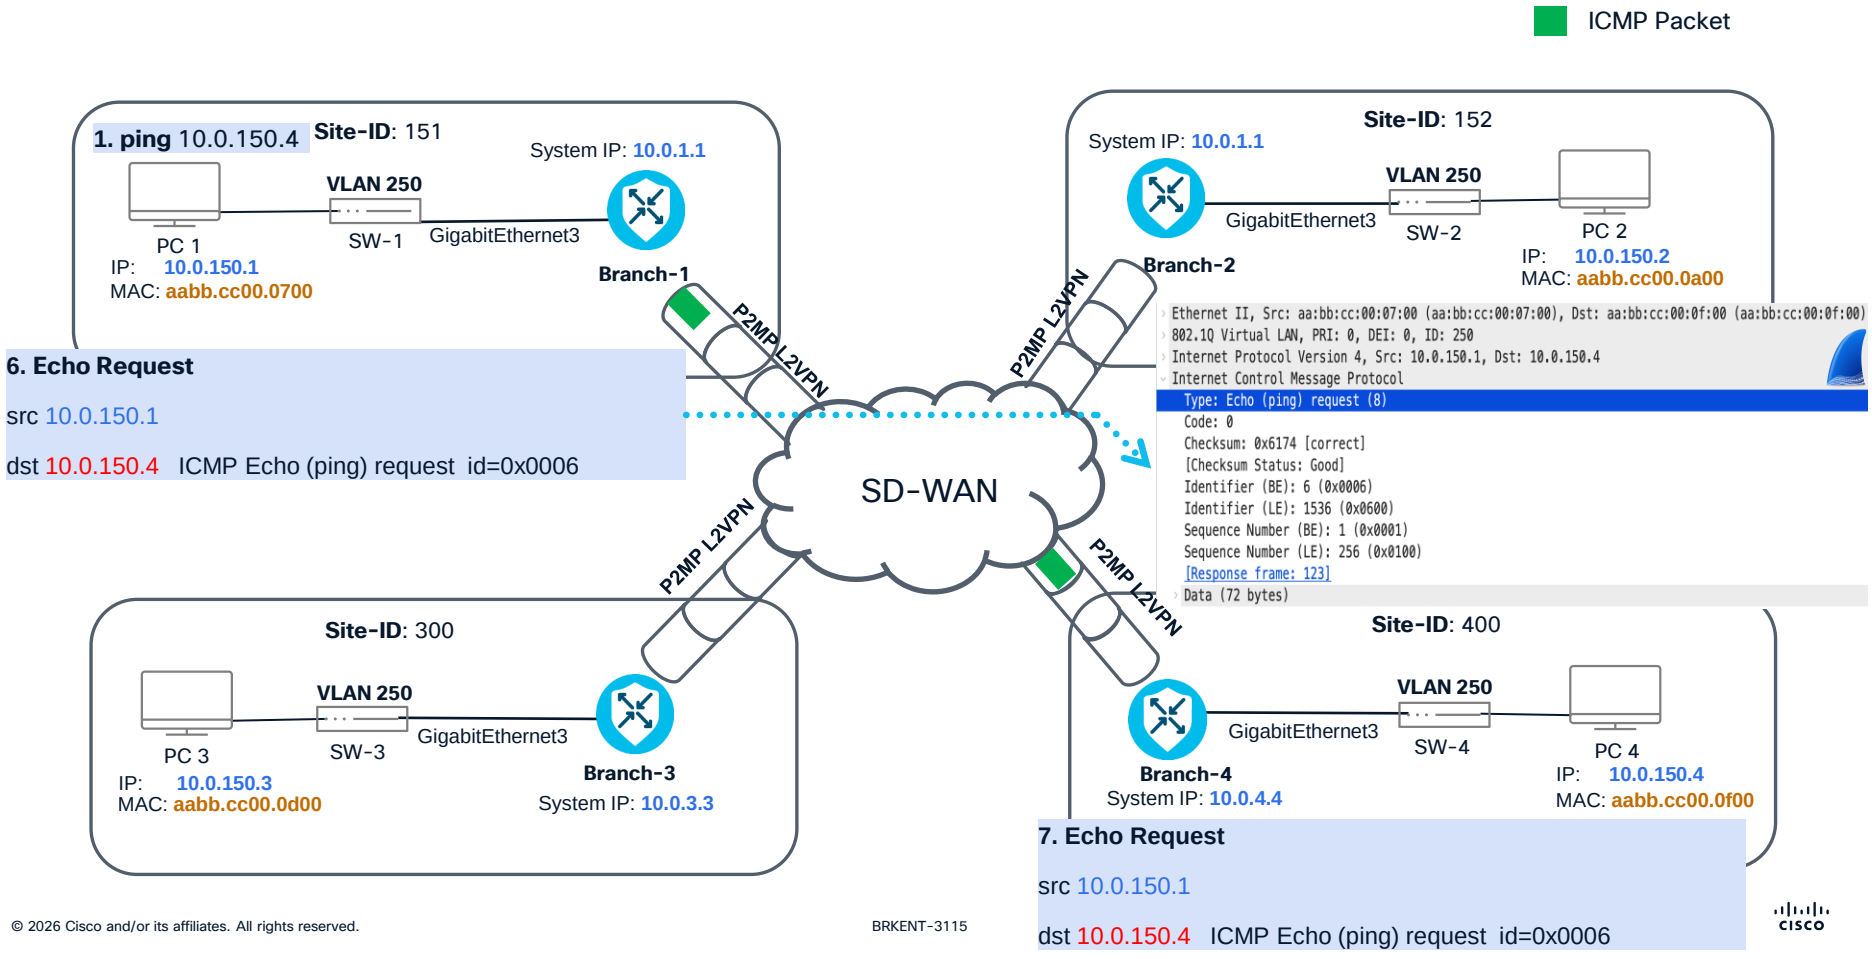

# How traffic flows between P2MP L2VPN? 9/9

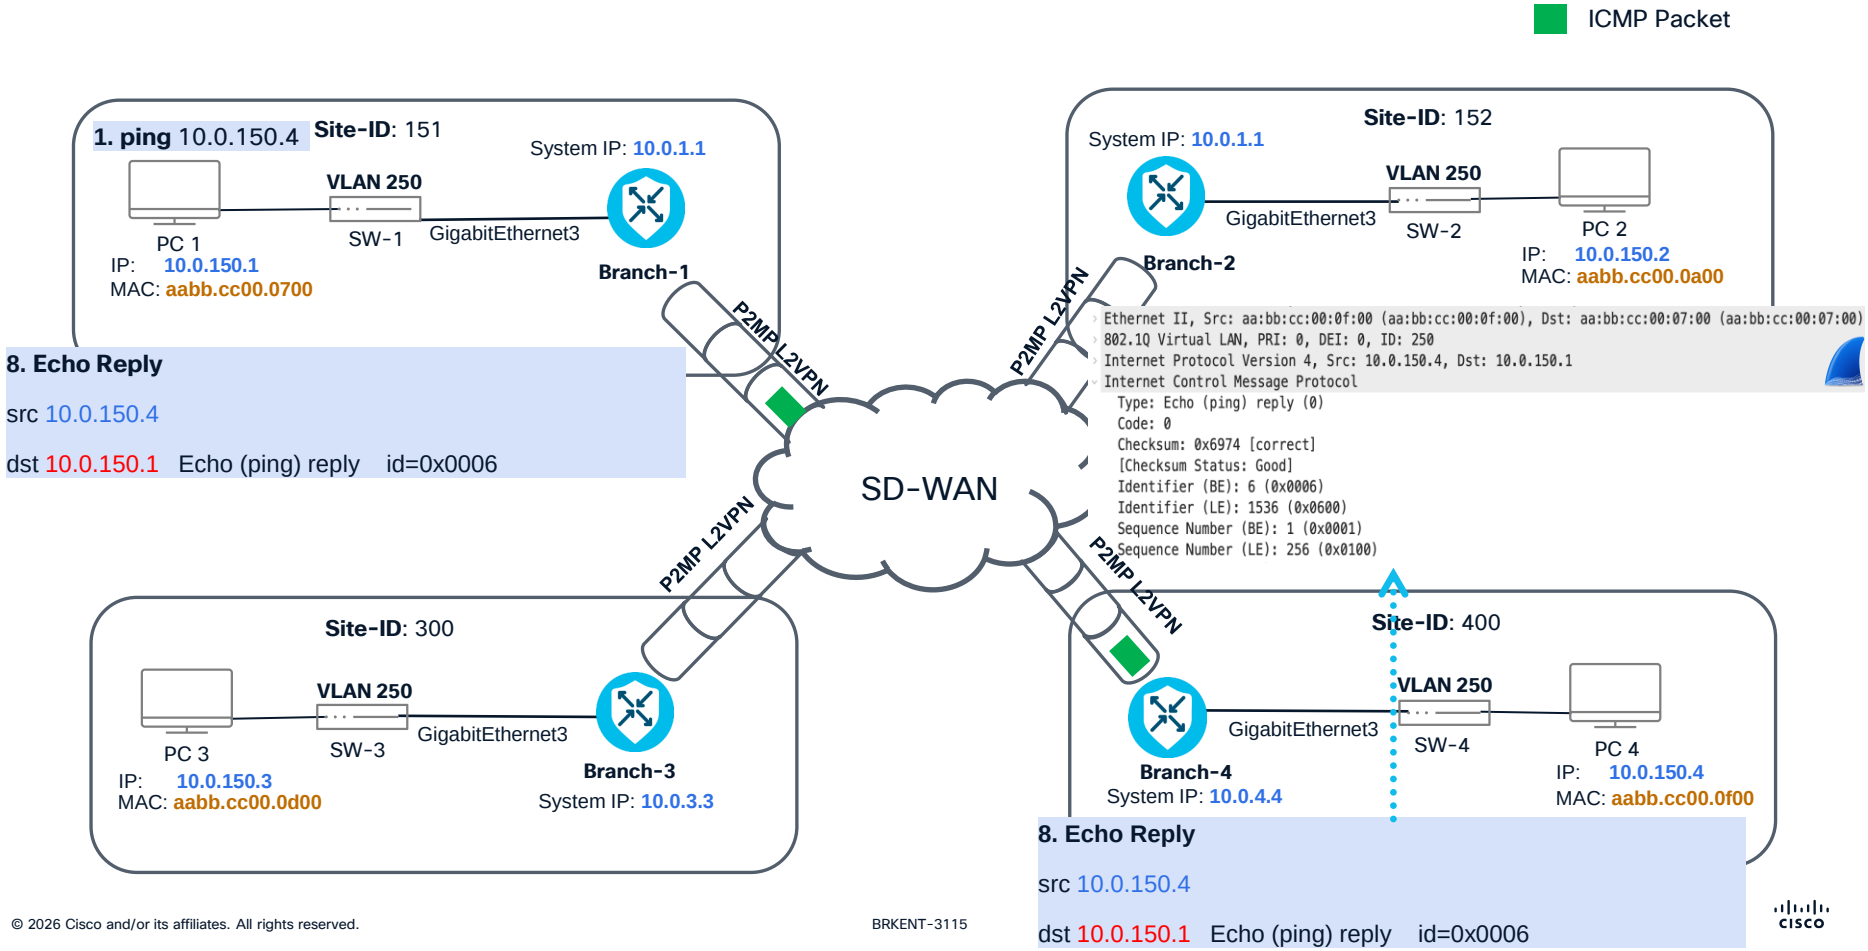

# P2MP L2VPN status and MAC-Learning on WAN-Edge

```
Branch-1# show l2vpn sdwan instance 12340 vc-id 2000 all
Bridge-Domain: 950
Local state: up
Homing type: single-homing
Homing role: active
Dual-Homing peer: no peer
Service EFP: Gi3 ServInst 1510
Remote Site: 152
  Originator       State          Up/Down    Homing-Role  Hub/Spoke  Color            Encap  Installed
  10.0.2.2         up             00:02:39   active       no         mpls             ipsec  yes(1008)
                                                                     biz-internet     ipsec  yes(1008)
Remote Site: 300
  Originator       State          Up/Down    Homing-Role  Hub/Spoke  Color            Encap  Installed
  10.0.3.3         up             00:02:38   active       no         mpls             ipsec  yes(1008)
                                                                     biz-internet     ipsec  yes(1008)
Remote Site: 400
  Originator       State          Up/Down    Homing-Role  Hub/Spoke  Color            Encap  Installed
  10.0.4.4         up             00:02:39   active       no         mpls             ipsec  yes(1008)
                                                                     biz-internet     ipsec  yes(1008)

Local MAC Info(1):
  MAC-Address     RecvTime
  aabb.cc00.0700  00:01:37
Remote MAC Info(3):
  MAC-Address     Originator       Site        RecvTime   Tloc-NH          Color            Encap  Label  Installed
  aabb.cc00.0a00  10.0.2.2         152         00:01:29   10.0.2.2         mpls             ipsec  1008   yes
                                                          10.0.2.2         biz-internet     ipsec  1008   yes

  aabb.cc00.0d00  10.0.3.3         300         00:01:26   10.0.3.3         mpls             ipsec  1008   yes
                                                          10.0.3.3         biz-internet     ipsec  1008   yes

  aabb.cc00.0f00  10.0.4.4         400         00:01:24   10.0.4.4         mpls             ipsec  1008   yes
                                                          10.0.4.4         biz-internet     ipsec  1008   yes
```

# Data Policy support for L2VPN

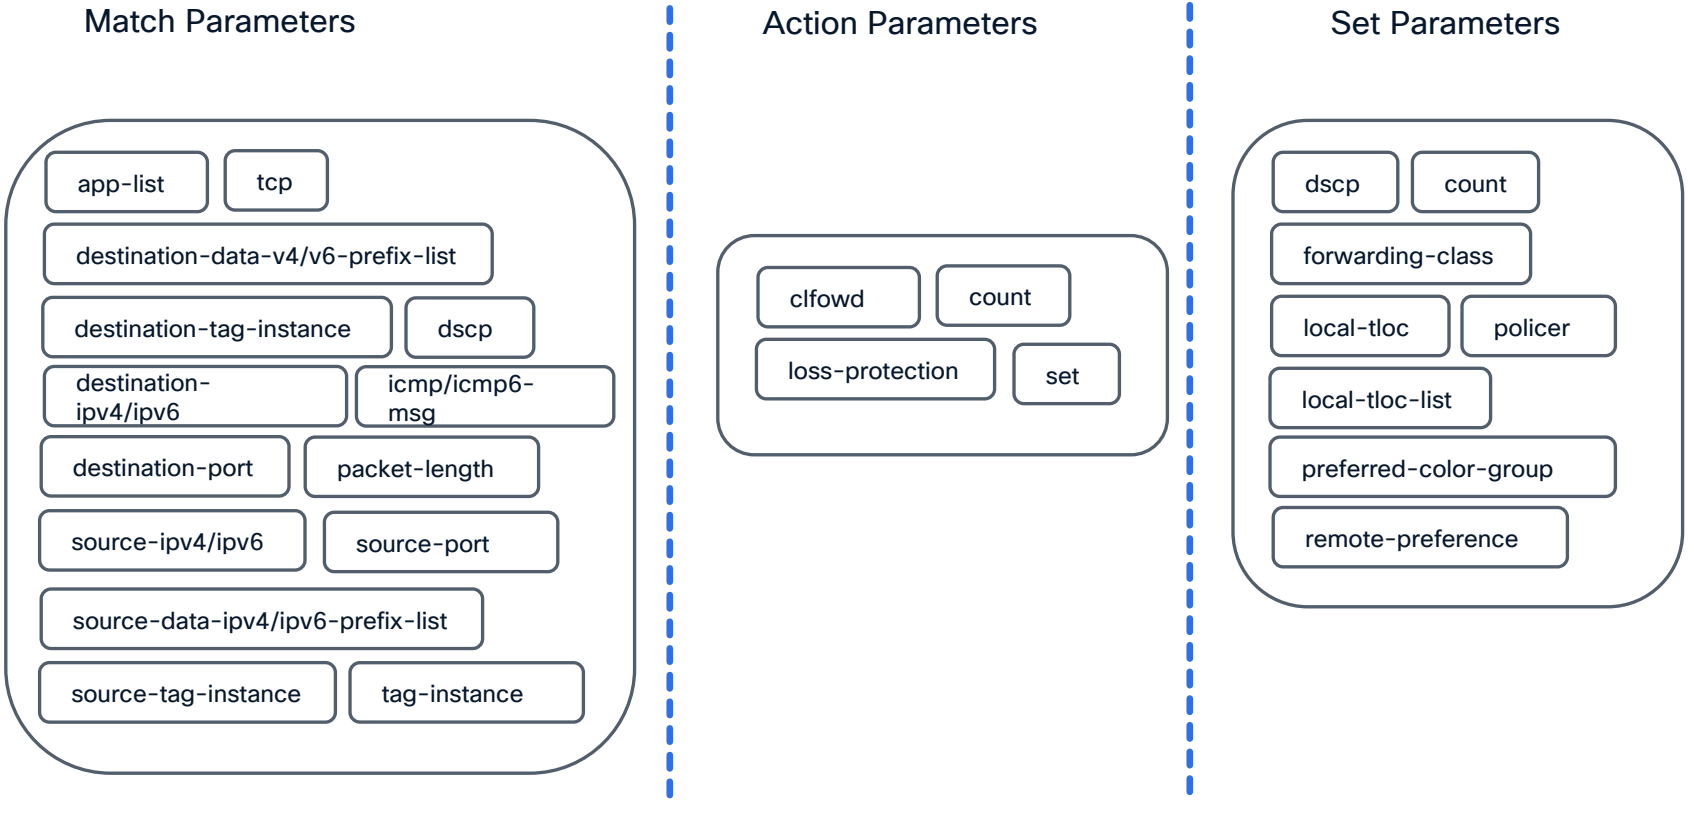

# Cisco Catalyst SD-WAN L2VPN Limitations

- **Hub-and-spoke** topology is supported in SD-WAN L2VPN deployments, but only the **intent-based hub-and-spoke** topology is supported.
- **Point-to-Point (P2P) L2VPN** service is **not** supported **between spoke sites**.
- For **BUM traffic** (Broadcast, Unknown Unicast, and Multicast) **originating** from a **spoke**, the **hub replicates** and forwards it to the other spokes.
- SD-WAN L2VPN supports **dual-homing** with per-VLAN single-active multi-homing, where for each VLAN one WAN Edge operates as **active** while the other remains in **standby**. On Standby, bi-direction traffic is blocked for that VLAN.

# What L2VPN Troubleshooting commands available?

For Reference

- Following commands can be employed on SD-WAN controller.
  - `show omp l2-routes | begin VPN | tab`
  - `show omp l2-services | begin VPN | tab`
  - `show support omp rib l2-routes vpn-id <VPN-ID> vc-id <VC-ID> originator <ORIGINATOR-System-IP> route-type vpn rib-out-peer-ip <PEER-IP-System-IP>`
- Following commands can be employed on SD-WAN WAN-Edge.
  - `show sdwan omp l2-routes`
  - `show sdwan omp l2-statuses`
  - `show l2vpn sdwan instance <L2VPN-Instance> vc-id <VC-ID> all`
  - `show platform software sdwan ftmd bridge-domain <Bridge-Domain-ID>`
  - `show sdwan ftm next-hop binos-nh-id <binosId-ID-Hex>`
  - `show platform hardware qfp active feature bridge-domain datapath <Bridge-Domain-ID>`

# How OMP Enables Seamless Service Chaining?

# What is Service Chaining?

In Cisco SD-WAN, a **service chain** is used to define the sequence of network services (like firewalls, load balancers, etc.) that traffic must pass through before reaching its destination.

- The short hairpin serves as the base model.
- Long & Exit are cases of short hairpin.
- Service Chain can be in any device in any topology: **full mesh**, **hub-spoke**, **MRF**.
- Max **16 Service Chain** types.
- Max **4 services** in a **Service Chain**(SC).

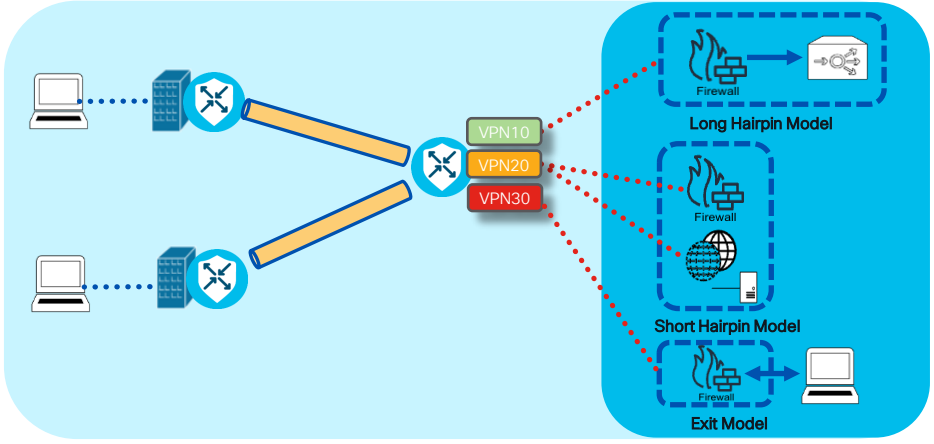


**UX1.0**

Traffic Steering to a SC "**set service-chain <>**" action in:

- Centralized **Control Policy**
- Centralized **Data Policy**
  - Co-located
  - Remote
- **Interface ACL**
  - Always **co-located.**

**UX2.0**

Traffic Steering to a SC "**set service-chain <>**" action in:

- Policy Group (**Traffic Policy**)
- Topology Group

# How Does OMP Relay Service Chaining Information?

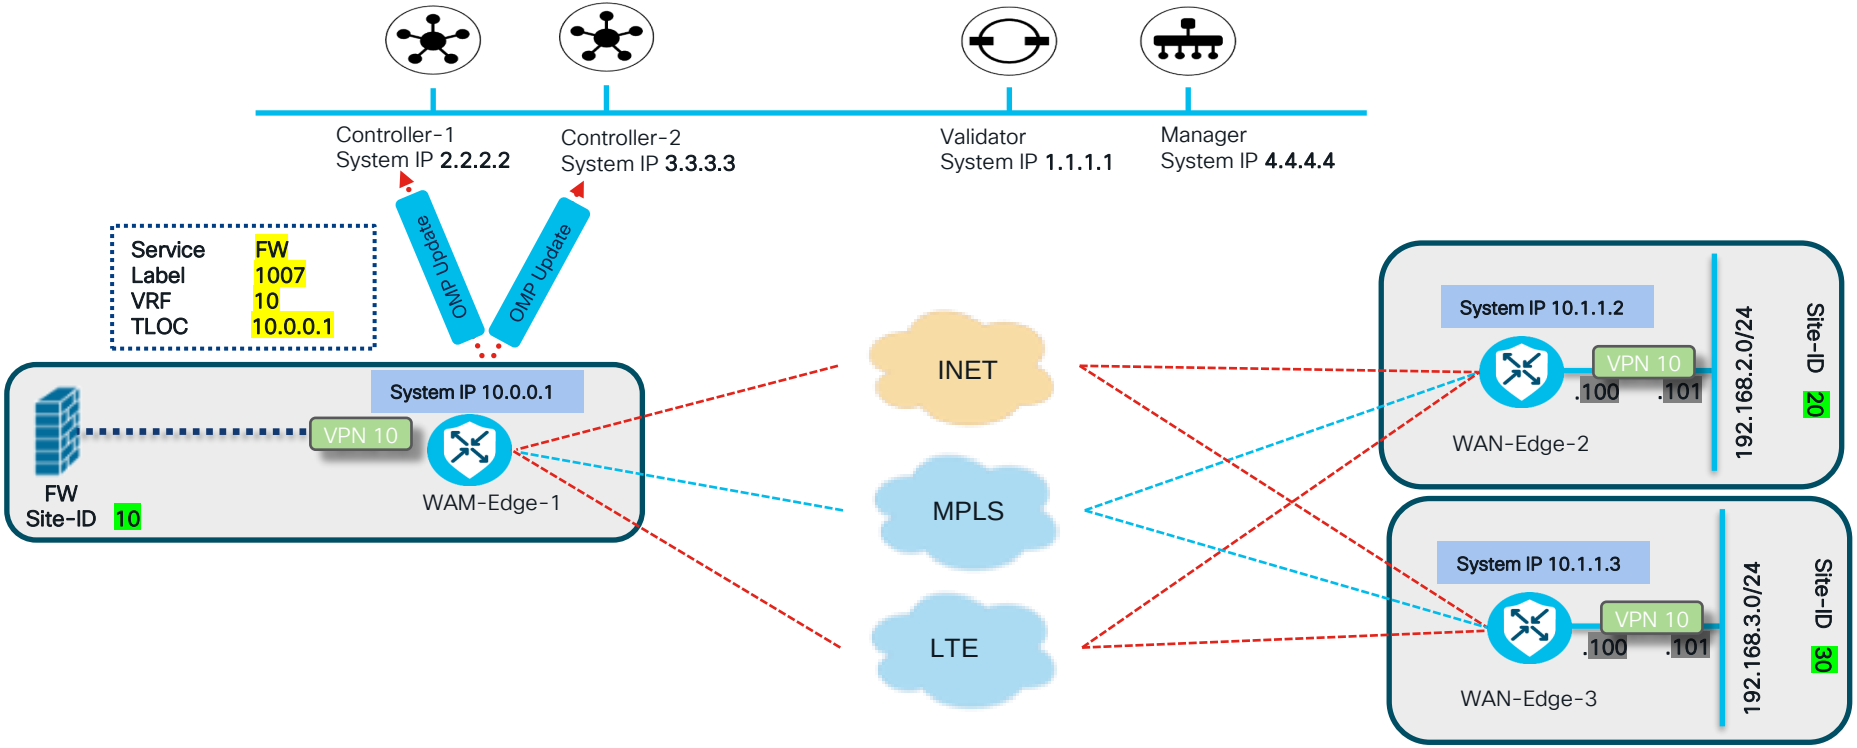

# How Does OMP Relay Service Chaining Information?

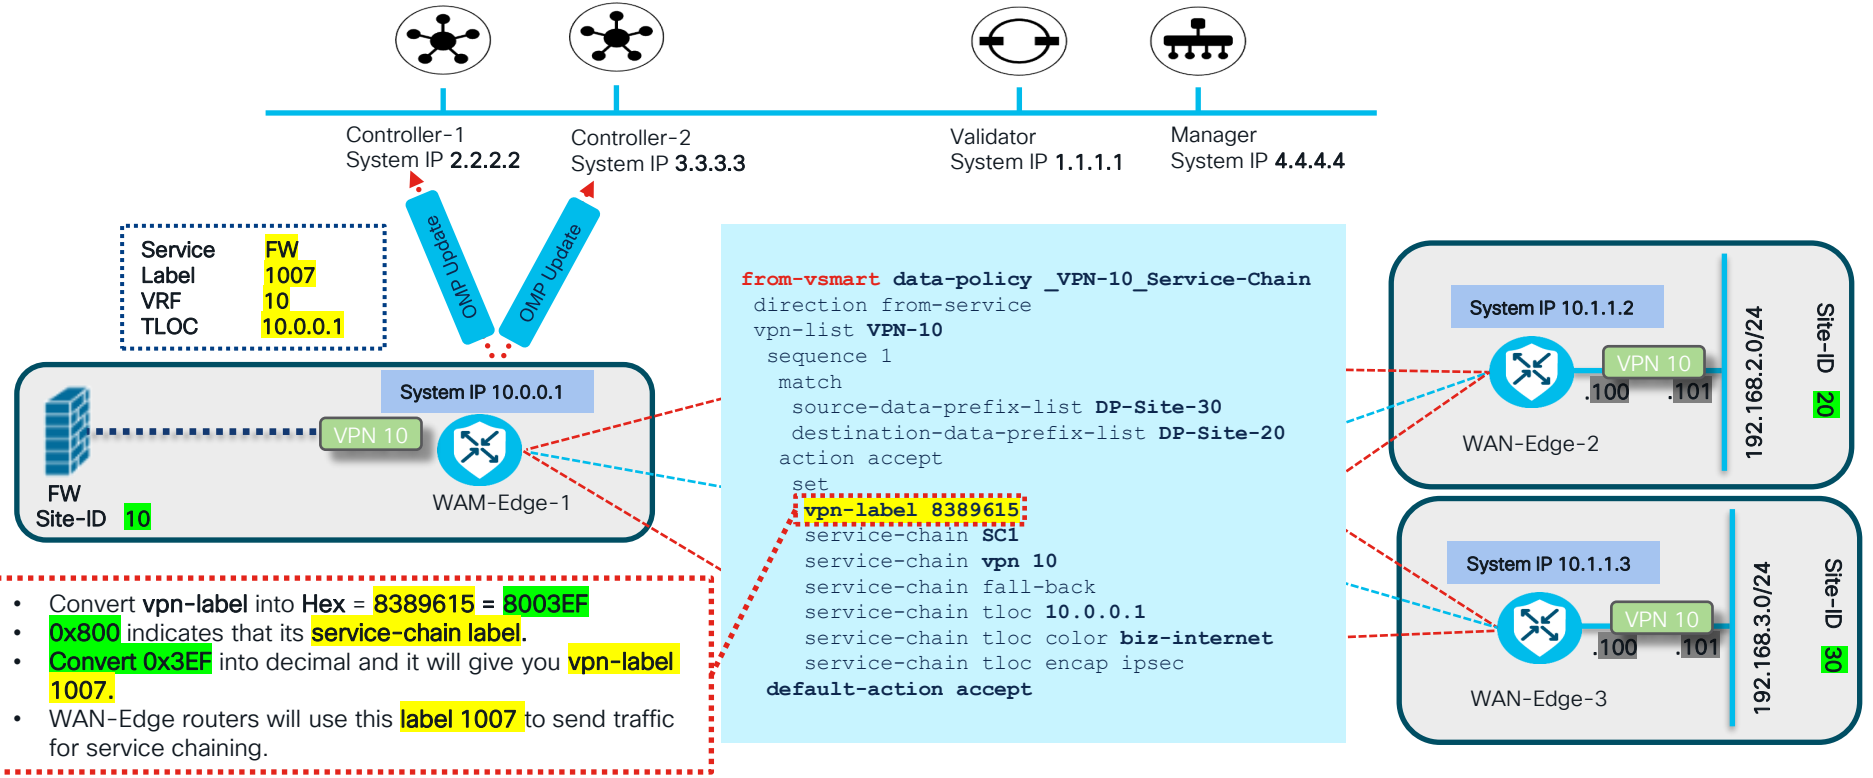


- Convert **vpn-label** into **Hex** = 8389615 = 8003EF
- 0x800 indicates that its **service-chain label**.
- Convert 0x3EF into decimal and it will give you **vpn-label** 1007.
- WAN-Edge routers will use this **label 1007** to send traffic for service chaining.

# Demo Topology

SD-WAN Edge routers and SD-WAN controllers are running **20.18.2/17.18.2** release.

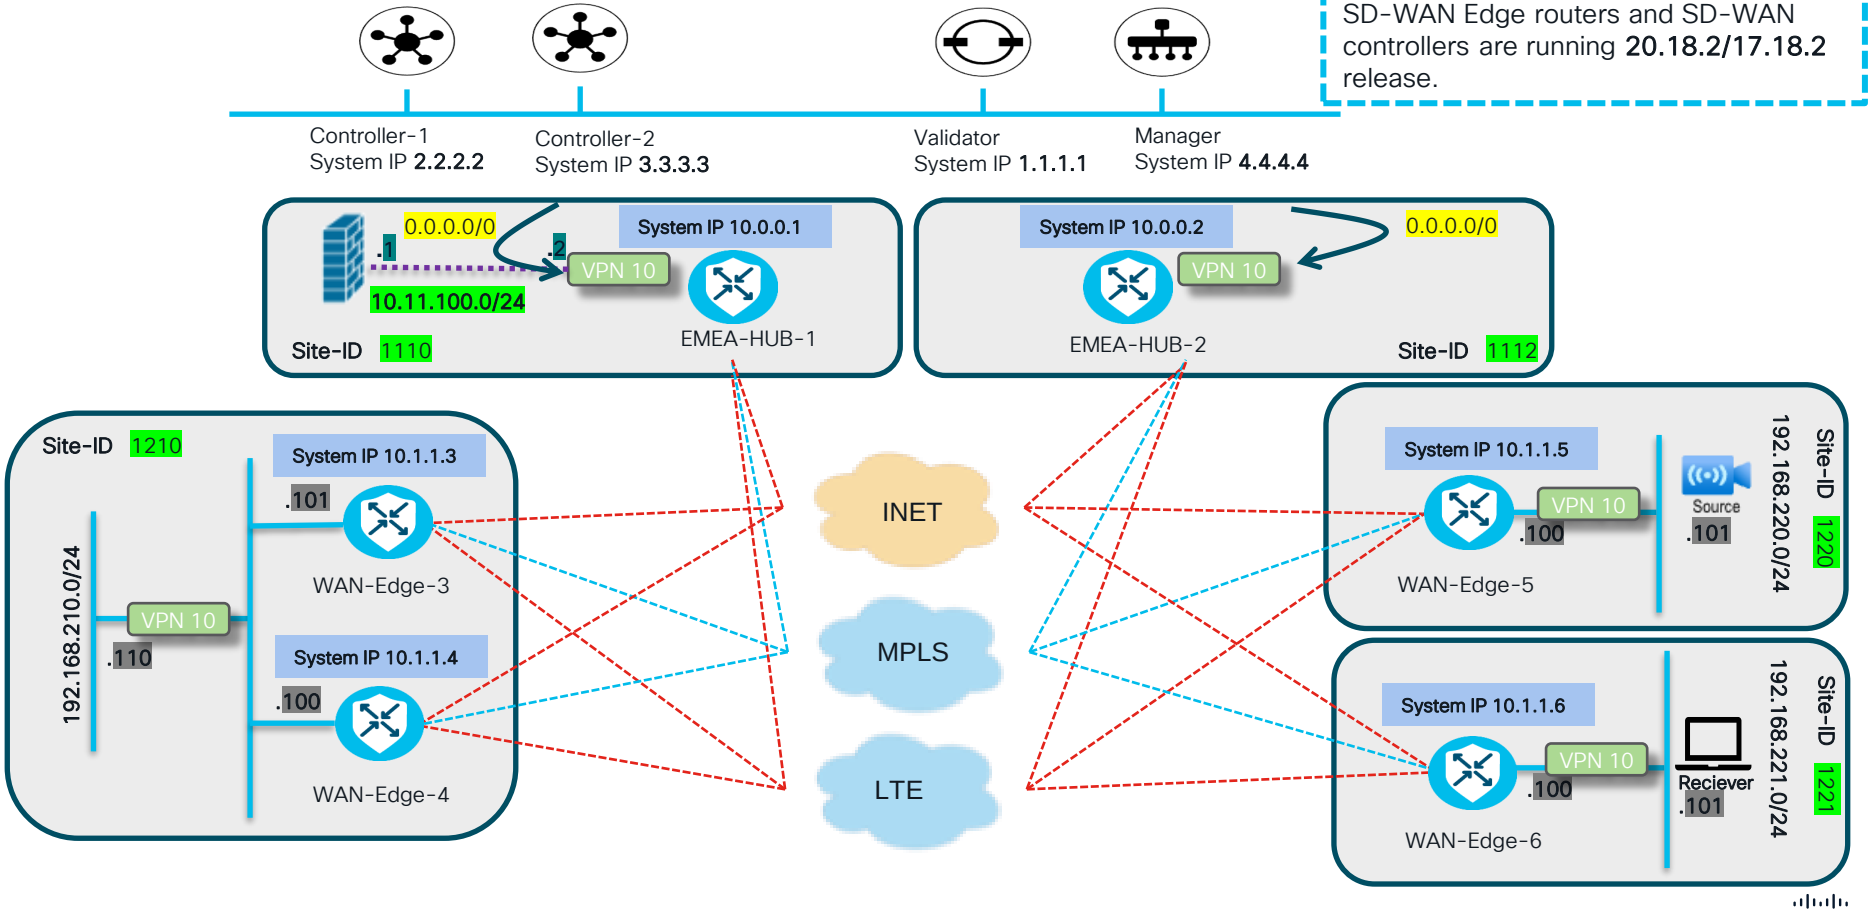

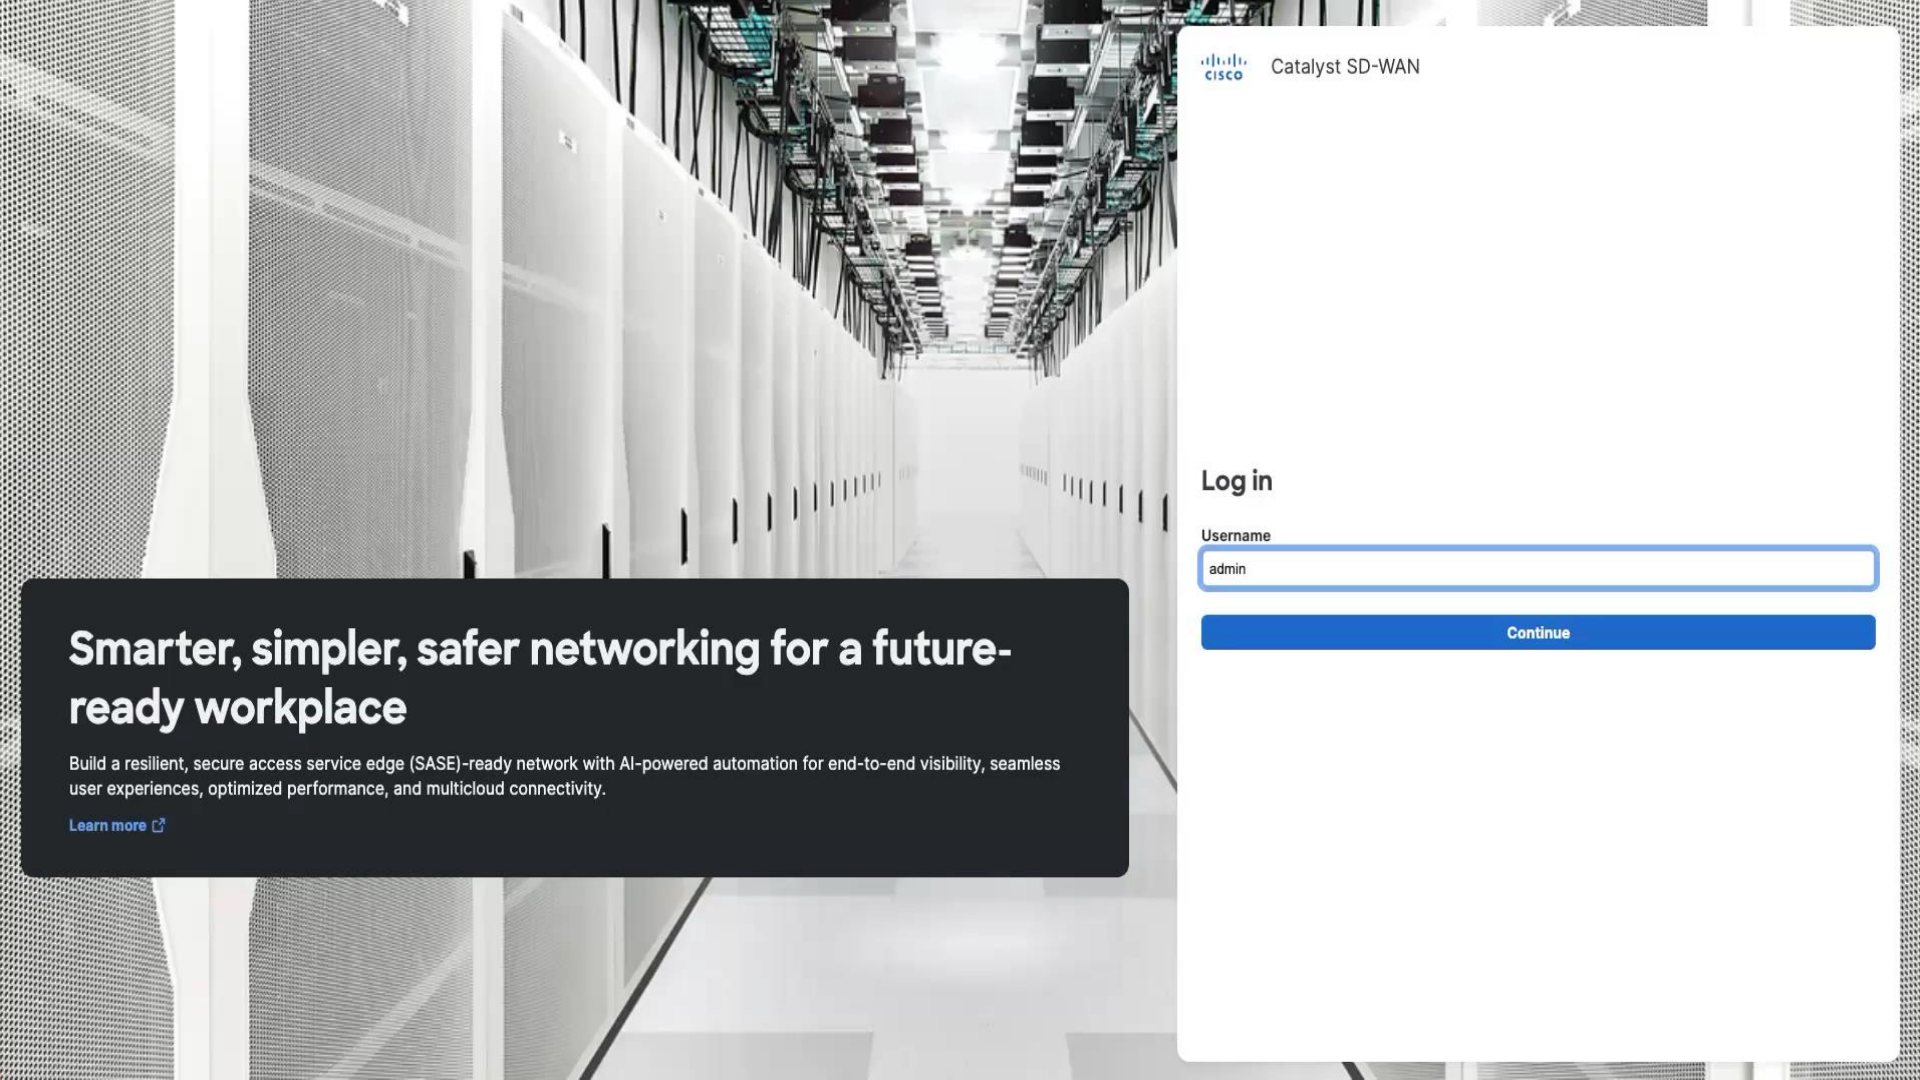
Catalyst SD-WAN
Log in
Username
admin
Continue
Smarter, simpler, safer networking for a future-ready workplace
Build a resilient, secure access service edge (SASE)-ready network with AI-powered automation for end-to-end visibility, seamless user experiences, optimized performance, and multicloud connectivity.
Learn more

# Quiz Time

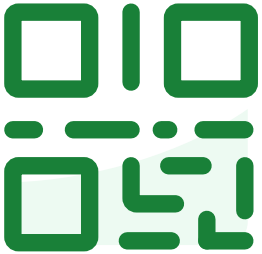

# Join at slido.com
# #3115

# Which statement is True about OMP peering?

# In SD-WAN Multicast, which statement is True, when there is an active source but no receivers?

# Which statement is TRUE about RIB-IN/RIB-OUT on SD-WAN controller?

# What is the primary objective of the OMP vRoute RIB-Out optimization introduced in Cisco Catalyst SD-WAN Release 17.18.2 / 20.18.2?

# Which two statements are TRUE about service-chaining?

# SD-WAN

Learn how to confidently deploy and operate Cisco's SD-WAN solution in a new or existing network. These sessions provide a journey from the foundation to latest Cisco SD-WAN innovations focusing on design, innovations, and integrations with Cloud, SASE, and Assurance/Analytics.

**START**

Tuesday, February 10 | 11:45 a.m.
**BRKENT-1313**
Making Catalyst SD-WAN Easy: Operational Simplification and User Experience

Tuesday, February 10 | 3:45 p.m.
**BRKENT-2126**
Three Steps to Gain Actionable Visibility in the Cisco Catalyst SD-WAN using ThousandEyes

Wednesday, February 11 | 5:00 p.m.
**BRKTRS-3475**
Automation and In-Depth Troubleshooting of Cisco Catalyst 8000, ASR 1000, ISR, and SD-WAN Edge

Wednesday, February 11 | 11:30 a.m.
**BRKTRS-3050**
Cisco SD-WAN, Hidden Complexity Revealed: How Cisco TAC Addresses Really Tricky Problems

Wednesday, February 11 | 3:15 p.m.
**BRKENT-3115**
SD-WAN OMP Unlocked: From Session Bring-Up to Advanced Routing, Multicast, and L2VPN

Thursday, February 12 | 5:00 p.m.
**BRKENT-3797**
Advanced Catalyst SD-WAN Policies Troubleshooting

Friday, February 13 | 9:00 a.m.
**BRKENT-2609**
Solving Global WAN Challenges with Multi-Region Fabric

Friday, February 13 | 11:00 a.m.
**FINISH**
**BRKTRS-2595**
Expedite your Troubleshooting with SD-WAN Manager Tools

If you are unable to attend a live session, you can watch it in the On-demand library.

# Become a power user

**Cisco Live Amsterdam Exclusive:**
Save 25% on CML-Personal and CML-Personal Plus.

Visit the Learning & Certification booth to learn more.

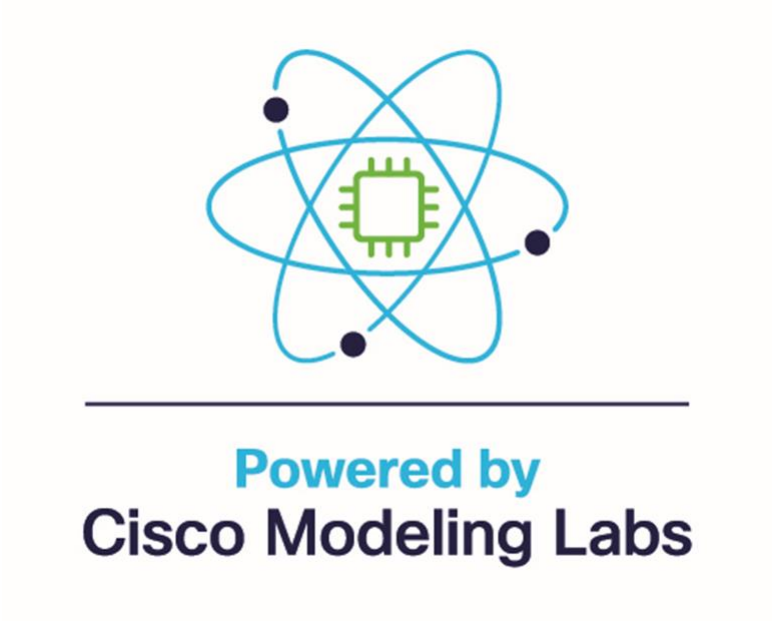

# Complete your session surveys

**Complete your surveys** in the Cisco Events App.

**Complete** a minimum of 4 session surveys and the overall event survey to receive a unique Cisco Live t-shirt.
(from 11:30 on Thursday, while supplies last)

# Continue your education

**Visit** the Cisco Showcase for related demos

**Book** your one-on-one Meet the Engineer meeting

**Visit** the Technical Solutions Clinics to discuss your technical questions

**Attend** the interactive education with DevNet, Capture the Flag, and Walk-in Labs

**Visit** the On-Demand Library for more sessions at CiscoLive.com/On-Demand

**Contact me at**: LinkedIn (https://www.linkedin.com/in/waqasdaar/)

# Thank you

cisco Live !


CISCO