
cisco Live!
ALL IN

The bridge to possible

Advanced SD-WAN routing troubleshooting

Lessons learned from the field

Eugene Khabarov, BU Escalation Engineer, CCIE #51348
@ekhabaro

BRKENT-3793

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile 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

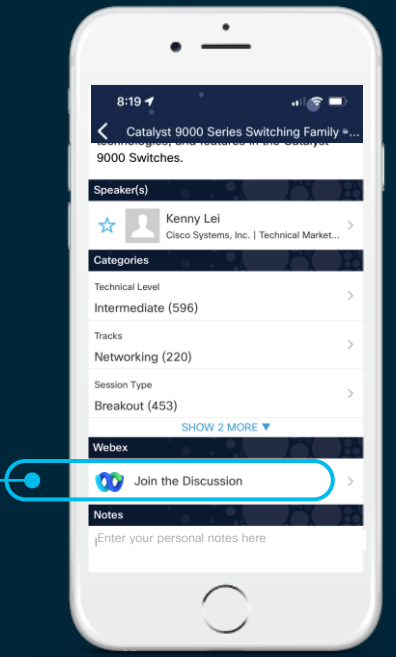

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKENT-3793

Webex spaces will be moderated by the speaker until June 17, 2022.

About me

and why I'm the right person to present about OMP troubleshooting

Eugene Khabarov

- Engineer's Degree in Computer Systems Networking and Telecommunications (VorSTU, Russia)
- 15+ years in ICT as support engineer, network engineer, consulting engineer, architect
- CCIE #51348 since 2015
- Cisco TAC engineer since 2017 at Cisco Systems Belgium
- EMEA SD-WAN TAC Team Lead since 2019
- Enterprise BU (SD-WAN) escalation engineer since 2021
- LinkedIn: https://www.linkedin.com/in/enk/
- GitHub: https://github.com/enk37/

Baseline and Objectives

- Cisco SD-WAN operational level knowledge required
- This is advanced level session, so basics are not covered
- The session main objective is to share experience about some typical failures seen in the field to help you avoid them in your network
- There are always more ☺
- Consider this session as a ”cookbook” for SD-WAN routing failures, but not a “Tour de Force”
- The session is OMP protocol oriented mainly, no service-side routing, AAR, multicast routing or control policies issues
- Topics touched:
    - Implicit ACL and underlay routing
    - OMP tuning
    - OMP path selection issues

“It’s good to learn from your mistakes. It’s better to learn from other people’s mistakes.”
– Warren Buffett

Agenda

- Part 1: SD-WAN Routing Troubleshooting basics
- Part 2: Issues seen in the field
  - VPN 0 (GRT) Routing Troubleshooting Cases
    - Case 1. Can not establish BGP peering with my ISP in the underlay
    - Case 2. BGP session established while it should not
  - OMP Troubleshooting Cases
    - Case 3. vSmart does not advertise any OMP routes
    - Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs
    - Case 5. OMP Path selection and global scalability
    - Case 6. Disjoined underlay active-standby redundancy fault scenario (origin metric)
    - Case 7. OMP double failure scenario

Part 1: SD-WAN Routing Troubleshooting Basics

What should we know to troubleshoot various issues with routing and forwarding

Cisco SD-WAN Routing

SD-WAN Fabric Components Quick Recap

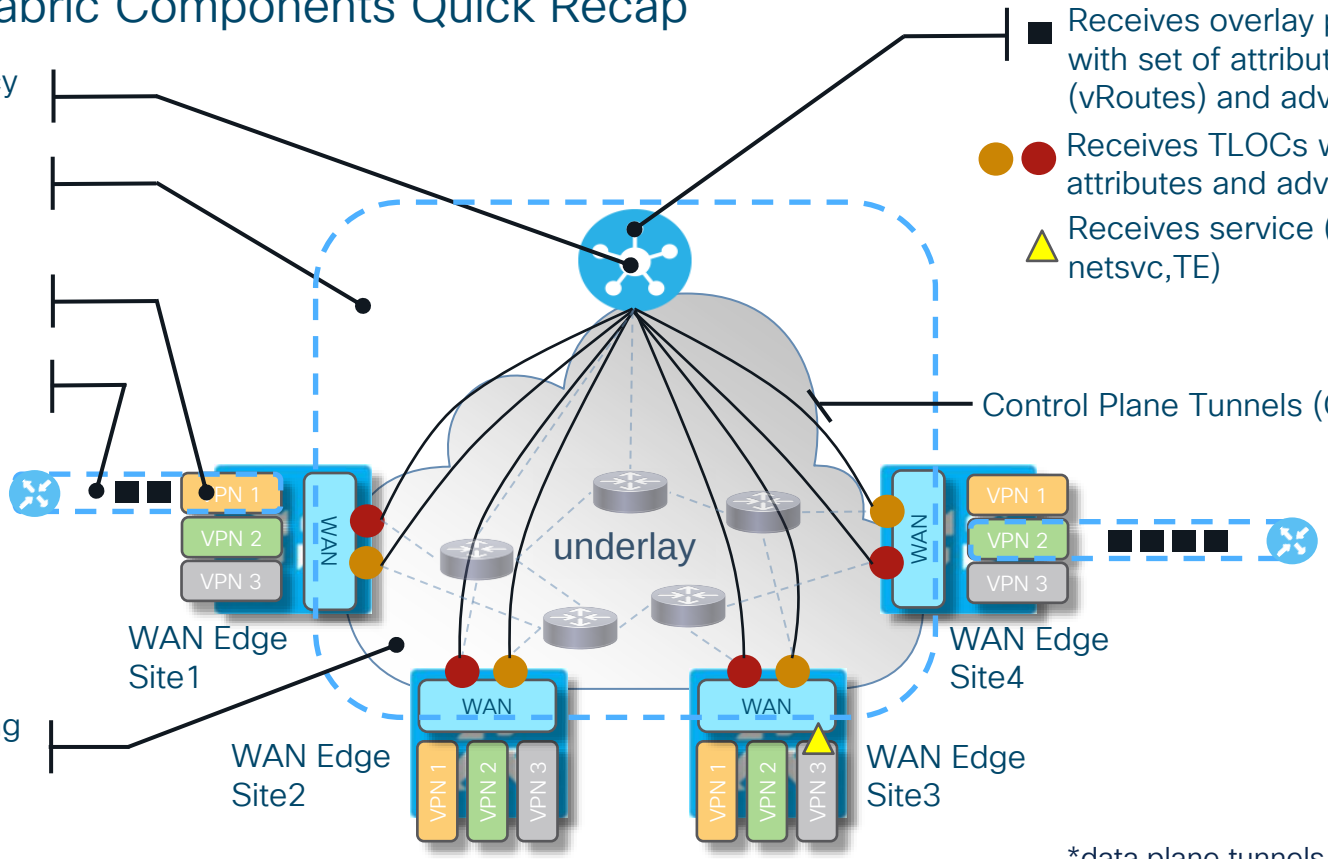

*data plane tunnels are not shown

First things first: OMP Peering

To exchange routing information, first of all peering with and between vSmart controllers must be established

```
vEdge1# show omp peers
                        DOMAIN   OVERLAY   SITE
PEER            TYPE    ID       ID        ID     STATE  UPTIME       R/I/S
-------------------------------------------------------------------------
10.255.255.78   vsmart  1         1         2     up     9:08:32:22   74/74/2
```

OMP Peering

If there is no OMP peering between WAN Edge and vSmart

- Check control connections and transport (VPN 0 underlay routing)
- Check for possible duplicate System ID
- Check OMP protocol status (`show [sdwan] omp summary`)

```
cEdge1#show sdwan omp peer
                        DOMAIN     OVERLAY    SITE
PEER             TYPE   ID         ID         ID         STATE    UPTIME       R/I/S
-------------------------------------------------------------------------------------
10.255.255.78    vsmart 1          1          2          init-in-gr            80/0/0
```

Graceful Restart Timer

If all vSmart controller connections are lost, the WAN Edge router continues to operate with the latest control plane information for the duration of configured OMP `graceful-restart` timer.

```
vsmart# show omp peer
                        DOMAIN   OVERLAY   SITE
PEER              TYPE  ID       ID        ID        STATE      UPTIME     R/I/S
-------------------------------------------------------------------------------
10.255.241.121    vedge 1        1         2111002   down-in-gr           4/0/0
```

```
vEdge1# show omp peer
                        DOMAIN   OVERLAY   SITE
PEER              TYPE  ID       ID        ID        STATE      UPTIME     R/I/S
-------------------------------------------------------------------------------
10.255.255.78     vsmart 1       1         2         init-in-gr           82/82/0
```

OMP Routing Basics – WAN Edge
From service-side and out to vSmart

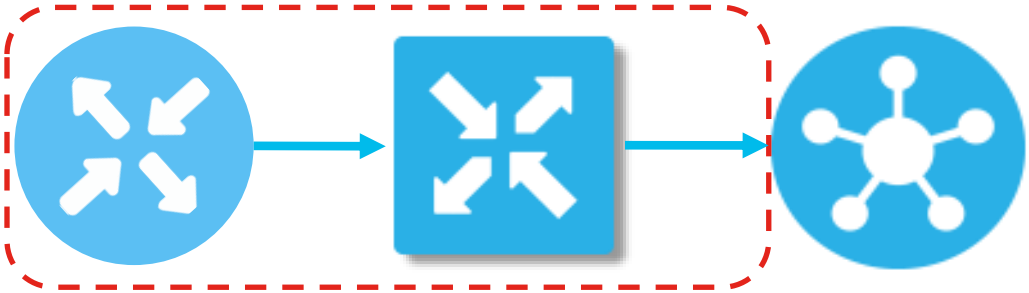

- By default, static, connected, and OSPF internal routes are automatically redistributed into OMP.

  All other (BGP, EIGRP, OSPF external) redistribution must be configured (into OMP and from OMP into another routing protocol)
- 4 best equal-cost paths are advertised to vSmart by default (configurable via `send-path-limit [1-16]`)
- Only the best paths are advertised to vSmart (routes installed into FIB) (non-best paths still can be sent if `send-backup-paths` configured)

OMP Routing Basics – vSmart

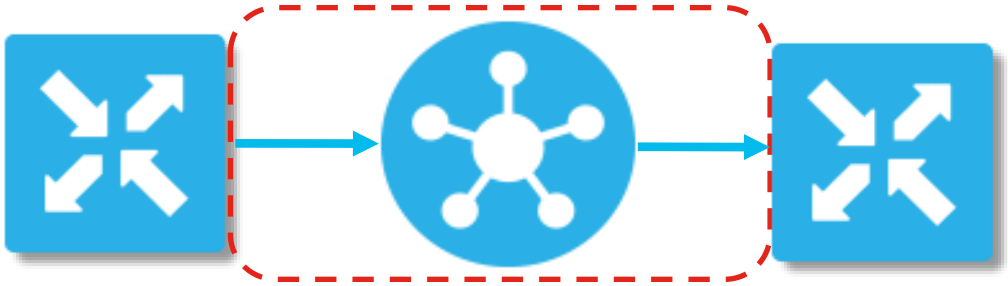

1. Applies incoming control policy on routes
2. Performs best paths selection (TLOC preference and TLOC reachability is not considered on vSmart)
3. Stores best paths in an ordered (sorted) list
4. Applies outgoing control policy and advertises best paths to WAN Edges
5. By default, only 4 best* equal-cost paths are advertised (configurable via `send-path-limit [1-16]` and `controller-send-path-limit [4-128]`)

*non-best paths can be sent also if `send-backup-paths` configured; 16 ECMP paths sent by default between controller starting from 20.5 (behavior change)

OMP Routing Basics – WAN Edge
From vSmart

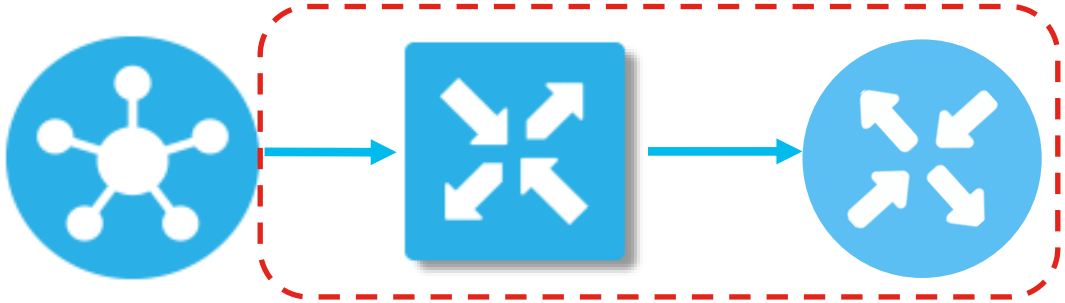

To decide which OMP routes are installed into the routing (RIB) and forwarding (FIB) tables:

- OMP performs loop avoidance and best path selection
- Implements localized policy

Up to 4 ECMP paths can be installed by default (configurable via `ecmp-limit [1-16]`)

- Installs route into RIB/FIB table if TLOC is active and there is a BFD session associated with it in the “up” state (route resolved).

Default OMP Administrative Distance is 250/251 for vEdge/cEdge and cannot be changed

Routing/Forwarding Information Base (RIB/FIB)

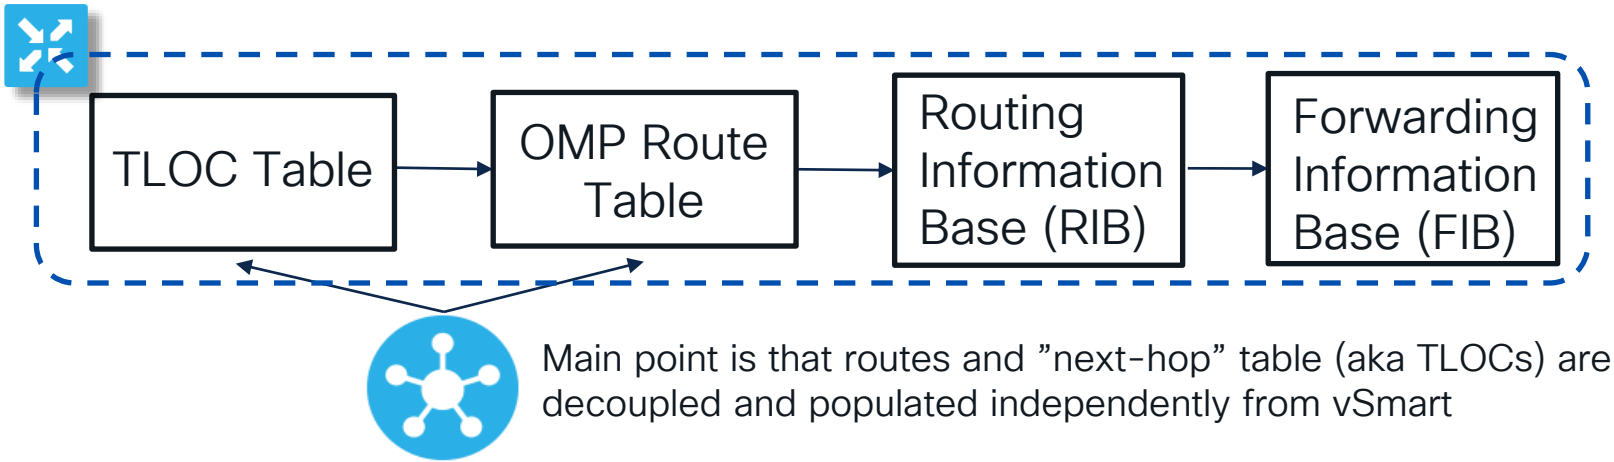

- RIB: Control plane view on what WAN Edge router decided to install from all routing protocol tables like OMP and use for unicast forwarding
- FIB: Forwarding plane version of the RIB that includes full “next-hop” information

OMP Routing Basics – Best Path Selection

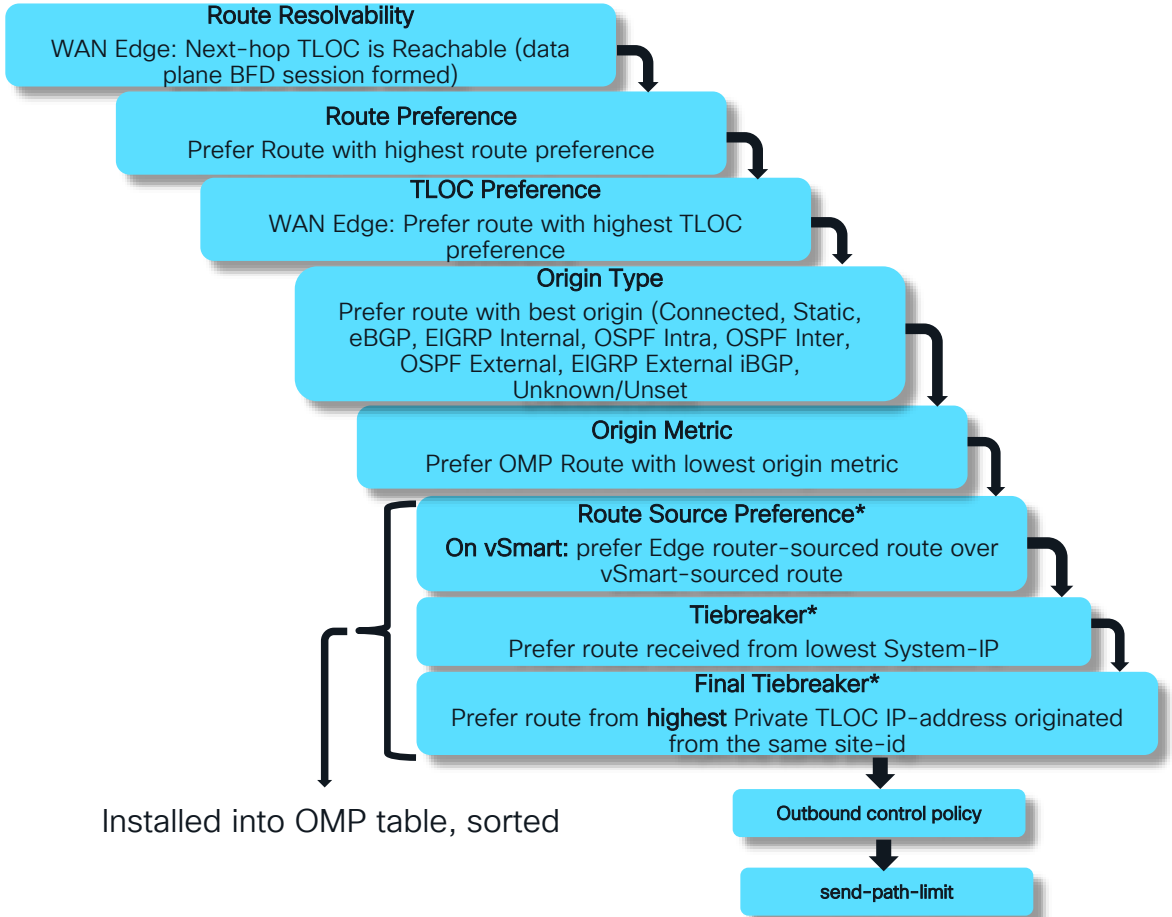

Missing Route(s) troubleshooting algorithm

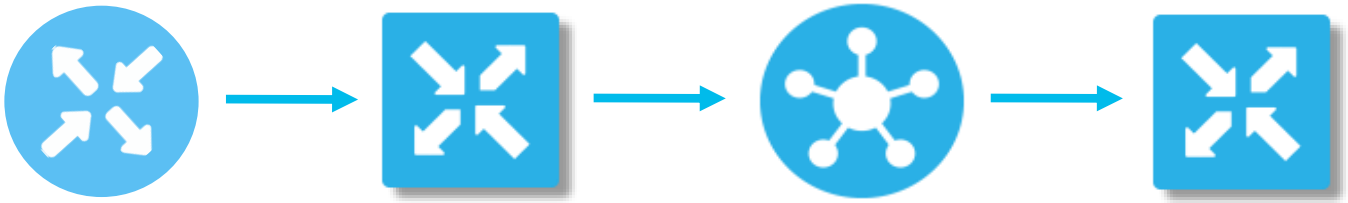

Check on WAN Edge:

1. RIB/FIB (`show ip route/show [sdwan] ip fib`)
2. OMP table if route is not in RIB (`show [sdwan] omp route`)
3. TLOC information presented (`show [sdwan] omp tloc`)
4. BFD session with remote TLOC (`show [sdwan] bfd sessions`) -> troubleshoot data plane
5. Local policy filtering on redistribution to/from OMP table (`show run policy/show run route-map`)

Check on vSmart:

1. OMP route and TLOC tables on vSmart (`show omp route, show omp tloc`)
2. Centralized control policies filters (`show run policy, show run apply-policy, test policy`)

RIB and FIB Tables (vEdge)

```
vEdge1# show ip route vpn 1 10.4.3.0/24

                                PROTOCOL  NEXTHOP  NEXTHOP NEXTHOP
VPN   PREFIX       PROTOCOL     SUB TYPE  IF NAME  ADDR    VPN  TLOC IP        COLOR        ENCAP  STATUS
-------------------------------------------------------------------------------------------------------
1     10.4.3.0/24  omp            -         -        -       -   10.255.241.101 mpls         ipsec  F,S
1     10.4.3.0/24  omp            -         -        -       -   10.255.241.101 biz-internet ipsec  F,S
```

```
vEdge1# show ip fib vpn 1 10.4.3.0/24

                           NEXTHOP     NEXTHOP          NEXTHOP  SA
VPN  PREFIX                IF NAME     ADDR             LABEL    INDEX    TLOC IP          COLOR
-----------------------------------------------------------------------------------------------
1    10.4.3.0/24           ipsec       10.4.1.2         1003     531      10.255.241.101   mpls
1    10.4.3.0/24           ipsec       64.100.1.23      1003     532      10.255.241.101   biz-internet
```

RIB and FIB Tables (IOS XE SD-WAN)

```
cEdge1#show ip route vrf 1 10.4.3.0 255.255.255.0
Routing Table: 1
Routing entry for 10.4.3.0/24
  Known via "omp", distance 251, metric 0, type omp
  Last update from 10.255.241.102 05:42:07 ago
  Routing Descriptor Blocks:
   10.255.241.102 (default), from 10.255.241.102, 05:42:07 ago
      Route metric is 0, traffic share count is 1
 * 10.255.241.101 (default), from 10.255.241.101, 05:42:07 ago
      Route metric is 0, traffic share count is 1
```

```
cEdge1#show sdwan ip fib vpn 1 | include 10.4.3.0
1    10.4.3.0/24         ipsec       10.4.1.2        1003     1385     10.255.241.101   mpls
1    10.4.3.0/24         ipsec       10.4.2.2        1006     1390     10.255.241.102   mpls
1    10.4.3.0/24         ipsec       64.100.1.23     1003     1404     10.255.241.101   biz-internet
1    10.4.3.0/24         ipsec       64.100.1.24     1006     1405     10.255.241.102   biz-internet
```

`show sdwan ip fib` – shows only SDWAN part of FIB on IOS-XE, non-SDWAN routes are in CEF table (`show ip cef`)

OMP Routes and Associated TLOCs

- TLOCs are uniquely identified with 3 parameters: System ID, color, encapsulation

```
vEdge1# show omp route
                                        PATH                ATTRIBUTE
VPN   PREFIX        FROM PEER           ID      LABEL       STATUS    TYPE        TLOC IP          COLOR          ENCAP
-------------------------------------------------------------------------------------------------------------------------
1     0.0.0.0/0     10.255.255.78       181     1003        C,I,R     installed   10.255.241.101   mpls           ipsec  -
                    10.255.255.78       182     1003        C,I,R     installed   10.255.241.101   biz-internet   ipsec  -
```

```
vEdge1# show omp tloc ip 10.255.241.101 color mpls encap ipsec
                                                                  PUBLIC              PRIVATE
                          PSEUDO                  PUBLIC                    PRIVATE   PUBLIC   IPV6     PRIVATE   IPV6     BFD
FROM PEER       STATUS    KEY    PUBLIC IP        PORT      PRIVATE IP      PORT      IPV6     PORT     IPV6      PORT     STATUS
-------------------------------------------------------------------------------------------------------------------------------
10.255.255.78   C,I,R     1      64.100.1.34      12369     10.4.1.2        12369     ::       0        ::        0        up
```

```
vEdge1# show omp tloc ip 10.255.241.101 color biz-internet encap ipsec
                                                                  PUBLIC              PRIVATE
                          PSEUDO                  PUBLIC                    PRIVATE   PUBLIC   IPV6     PRIVATE   IPV6     BFD
FROM PEER       STATUS    KEY    PUBLIC IP        PORT      PRIVATE IP      PORT      IPV6     PORT     IPV6      PORT     STATUS
-------------------------------------------------------------------------------------------------------------------------------
-10.255.255.78  C,I,R     1      64.100.1.23      12389     10.4.1.6        12389     ::       0        ::        0        up
```

IOS-XE SD-WAN: `show sdwan omp route`
`show sdwan omp tlocs`

OMP Routing Typical Issue – TLOC resolvability

Missing Routes from DC in WAN Edge RIB. Expected 4 default routes, but only 2 installed into RIB.

```
vEdge1# show ip route vpn 1
                                                PROTOCOL  NEXTHOP   NEXTHOP        NEXTHOP
VPN     PREFIX          PROTOCOL                SUB TYPE  IF NAME   ADDR           VPN      TLOC IP          COLOR      ENCAP  STATUS
-------------------------------------------------------------------------------------------------------------------------------------
1       0.0.0.0/0         omp                     -         -         -              -        10.255.241.101   mpls       ipsec  F,S
1       0.0.0.0/0         omp                     -         -         -              -        10.255.241.102   mpls       ipsec  F,S
```

```
vEdge1# show omp route vpn 1  | begin PATH
                                        PATH                          ATTRIBUTE
VPN     PREFIX      FROM PEER           ID      LABEL     STATUS      TYPE        TLOC IP          COLOR             ENCAP
-----------------------------------------------------------------------------------------------------------------------------

1       0.0.0.0/0   10.255.255.78       63      1006      C,I,R       installed   10.255.241.102   mpls              ipsec  -
                    10.255.255.78       64      1006      Inv,U       installed   10.255.241.102   biz-internet      ipsec  -
                    10.255.255.78       95      1003      C,I,R       installed   10.255.241.101   mpls              ipsec  -
                    10.255.255.78       96      1003      Inv,U       installed   10.255.241.101   biz-interne       ipsec  -
```

```
vEdge1# show omp tlocs ip 10.255.241.101 color biz-internet

                          PSEUDO                PUBLIC                PRIVATE  PUBLIC  IPV6    PRIVATE  IPV6    BFD
FROM PEER        STATUS   KEY     PUBLIC IP     PORT    PRIVATE IP    PORT     IPV6    PORT    IPV6     PORT    STATUS
-----------------------------------------------------------------------------------------------------------------------
10.255.255.78    C,I,R    1       64.100.1.23   12389   10.4.1.6      12389    ::      0       ::       0       down
```

OMP Routing Typical Issue – TLOC resolvability

There is new and better command available to do the same in one go starting from 20.8.1 - `show omp verify-routes`

```
vedge1# show omp verify-routes vpn 1 0.0.0.0/8 | b PATH
                  PATH                        ATTRIBUTE                                                    TLOC                  BFD     RIB
FROM PEER         ID      LABEL     STATUS    TYPE       TLOC IP          COLOR            ENCAP  STATUS     PREFERENCE  STATUS  STATUS
-------------------------------------------------------------------------------------------------------------------------------------------
169.254.206.4     70      1003      C,I,R     installed  169.254.206.12   mpls             ipsec  C,I,R      -           up      F,S
169.254.206.4     71      1003      Inv,U     installed  169.254.206.12   biz-internet     ipsec  C,I,R      -           down    -
169.254.206.4     72      1003      C,I,R     installed  169.254.206.11   mpls             ipsec  C,I,R      -           up      F,S
169.254.206.4     73      1003      Inv,U     installed  169.254.206.11   biz-internet     ipsec  C,I,R      -           down    -
```

Useful Commands and Debugs

IOS XE SD-WAN:

```
show ip route vrf <>

show sdwan ip fib vpn <>
show sdwan omp tlocs
show sdwan omp routes
show sdwan omp peers
show sdwan omp summary

debug platform software sdwan omp best-path
debug platform software sdwan omp packets
debug platform software sdwan omp events

show logging process ompd internal <>
```

Packet-tracer:

```
debug platform packet-trace
debug platform condition
```

Embedded Packet Capture: `monitor capture`

vSmart:

```
show omp tlocs <> [detail]
show omp routes [detail]
show omp peers <> [detail]
test policy match control-policy
```

vEdge:

```
show ip route vpn <>
show ip fib vpn <>
show omp tlocs
show omp routes
show omp peers
show omp summary
show omp verify-routes

debug omp best-path
debug omp packets
debug omp events
debug omp policy

show log /var/log/tmplog/vdebug tail -f
```

Packet-tracer (20.4): `debug packet-trace condition`

Packet Capture: `tcpdump <>`

vManage – Monitor OMP routing

Monitor>Network>*[Device]*>Real Time

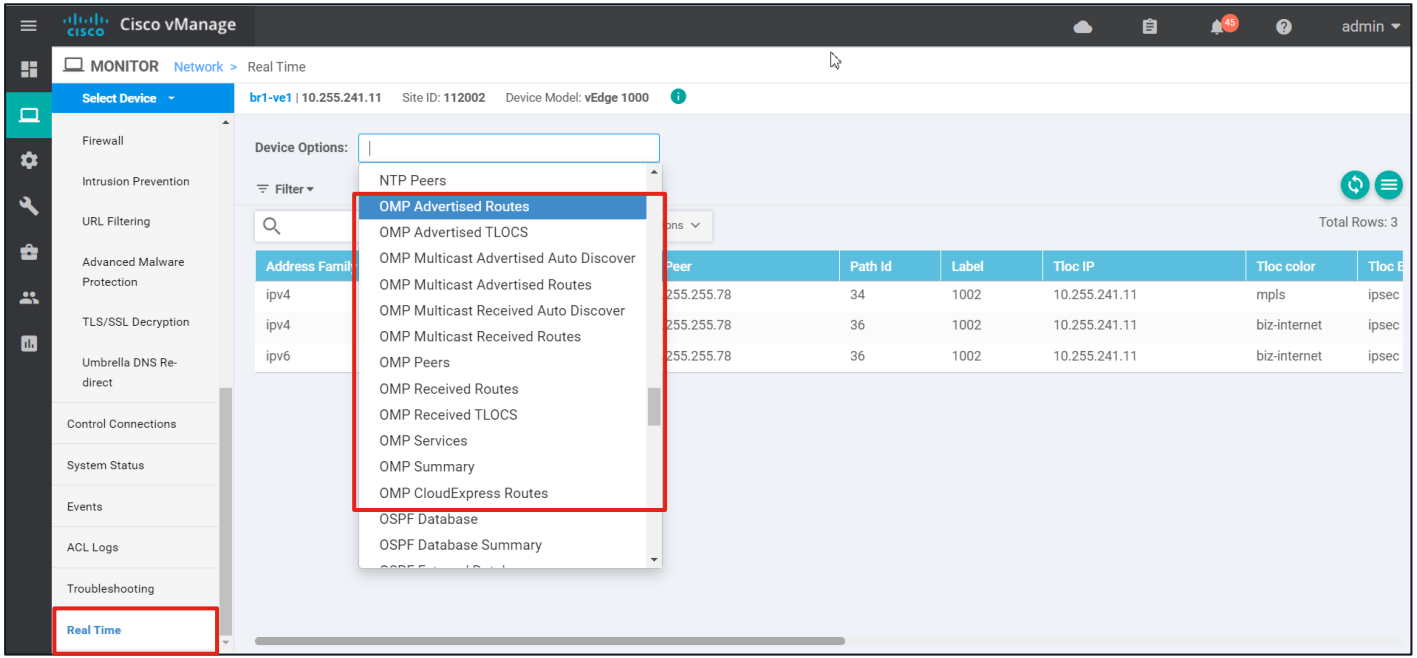

Part 2: Issues Seen in the Field

VPN 0 (GRT) Routing Cases underlay routing issues

Case 1. Can not establish BGP peering with my ISP in the underlay

Case 1. Can not establish BGP peering with my ISP in the underlay

Symptoms. BGP peer is stuck in “connect” state

```
vEdge1# show bgp neighbor

                         MSG   MSG   OUT                    LAST    AFI
VPN  PEER ADDR     AS    RCVD  SENT  Q    UPTIME  STATE     UPTIME  ID   AFI
-------------------------------------------------------------------------------
0    192.168.9.1   65000 0     0     0    -       connect   -       0    ipv4-unicast
```

BGP peer is reachable via icmp

```
vEdge1# ping 192.168.9.1 rapid count 5
Ping in VPN 0
!!!!!
--- 192.168.9.1 statistics ---
5 packets transmitted, 5 received, 0% packet loss
```

But can not connect to tcp/179:

```
vEdge1:~$ telnet 192.168.9.1 179

telnet: can't connect to remote host (192.168.9.1): Connection timed out
```

Your ISP denies problem on their side...

Case 1. Can not establish BGP peering with my ISP in the underlay

`debug bgp events vpn 0` and `debug bgp packets vpn 0` - not very helpful:

```
vEdge1# vshell
vEdge1:~$ tail /var/log/tmplog/vdebug
local7.debug: Sep 30 19:55:45 vEdge1 BGP[1386]: Import timer expired.
local7.debug: Sep 30 19:55:46 vEdge1 BGP[1386]: 192.168.9.1 [Event] Connect start to 192.168.9.1 fd 15
local7.debug: Sep 30 19:55:47 vEdge1 ZEBRA[1410]: RTM_DELROUTE ipv6 unicast proto
local7.debug: Sep 30 19:55:49 vEdge1 BGP[1386]: Performing BGP general scanning
local7.debug: Sep 30 19:55:49 vEdge1 BGP[1386]: scanning IPv4 Unicast routing tables
local7.debug: Sep 30 19:55:49 vEdge1 BGP[1386]: scanning IPv6 Unicast routing tables
local7.info: Sep 30 19:55:53 vEdge1 CFGMGR[1387]: %Viptela-vEdge1-cfgmgr-6-INFO-1400002: Notification: 9/30/2020 17:55:53
bgp-peer-state-change severity-level:major host-name:"vEdge1" system-ip:10.10.10.232 vpn-id:0 peer:192.168.9.1 bgp-new-
state:idle local-address:0.0.0.0 local-routerid:10.10.10.232 peer-routerid:0.0.0.0
local7.debug: Sep 30 19:55:54 vEdge1 BGP[1386]: Performing BGP general scanning
local7.debug: Sep 30 19:55:54 vEdge1 BGP[1386]: scanning IPv4 Unicast routing tables
local7.debug: Sep 30 19:55:54 vEdge1 BGP[1386]: scanning IPv6 Unicast routing tables
```

Case 1. Can not establish BGP peering with my ISP in the underlay

Form a theory. What would we check further?

Facts:

- It’s not transport issue, peer is reachable via icmp
- It’s not BGP misconfiguration, we are trying to connect to the right peer
- It’s not ISP issue, we trust in our ISP

Case 1. Can not establish BGP peering with my ISP in the underlay

`packet-trace` tool is our universal troubleshooting helper:

```
vEdge1# debug packet-trace condition source-ip 192.168.9.1 bidirectional vpn 0
vEdge1# debug packet-trace condition start

vEdge1# show packet-trace statistics

                                SOURCE                          DESTINATION SOURCE       DESTINATION
TRACE ID    SOURCE IP           PORT        DESTINATION IP      PORT        INTERFACE    INTERFACE        DECISION   DURATION
-------------------------------------------------------------------------------------------------------------------------------
0           192.168.9.1         179         192.168.9.10        56414       ge0_0        loop0.0          DROP       42
1           192.168.9.10        56414       192.168.9.1         179         loop0.0      ge0_0            FORWARD    80
2           192.168.9.1         179         192.168.9.10        56414       ge0_0        loop0.0          DROP       26
```

Case 1. Can not establish BGP peering with my ISP in the underlay

Packets details won't give us more info on vEdge because packet processed in Fast Path:

```
vEdge1# show packet-trace details 0
============================================================================================
Pkt-id  src_ip(ingress_if)          dest_ip(egress_if)             Duration    Decision    Protocol
============================================================================================
0       192.168.9.1:179 (ge0_0)     192.168.9.10:56414 (loop0.0)  42 us       DROP         6
INGRESS_PKT:
00 0c 29 ba 66 26 c4 b2 39 de cc 04 08 00 45 c0 00 2c e2 ad 40 00 01 06 03 03 c0 a8 09 01 c0
a8 09 0a 00 b3 dc 5e f6 c9 61 db 35 56 d5 c4 60 12 40 00 84 14 00 00 02 04 05 88 00 00 00
00 00 00 00 00 8c fe fd 3c 5f 73 67 c8 42 20 1e ac 7a f7 74 aa 26 7b b6 18 cc 0b c1 a2 34
f3 73 bb 7d 00
EGRESS_PKT:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
Feature Data
-------------------------------------
TOUCH : fp_proc_packet
core_id: 1
DSCP: 48
-------------------------------------
TOUCH : fp_hw_x86_pkt_free
core_id: 1
DSCP: 48
```

Case 1. Can not establish BGP peering with my ISP in the underlay

But system-wide drop counters must help to narrow-down to common suspect:

```
vEdge1# show system statistics diff | i drop
                  rx_implicit_acl_drops  :     4
                        rx_non_ip_drops  :     3
                   rx_arp_non_local_drops  :     2
```

Let's enable `implicit-acl-logging` to confirm the root cause:

```
vEdge1(config)# policy
vEdge1(config-policy)# implicit-acl-logging
vEdge1(config-policy)# commit
Commit complete.
vEdge1(config-policy)# end
vEdge1# vshell
vEdge1:~$ grep 179.*Implicit-ACL /var/log/messages
local7.notice: Sep 30 19:51:48 vEdge1 FTMD[1394]: %Viptela-vEdge1-FTMD-5-NTCE-1000026: FLOW LOG vpn-0 192.168.9.1/179
192.168.9.10/34813 tcp: tos: 192  inbound-acl, Implicit-ACL, Result: denyPkt count 1: Byte count 60 Ingress-Intf ge0/2
Egress-intf cpu
```

Case 1. Can not establish BGP peering with my ISP in the underlay

Solution 1: allow BGP in implicit ACL:

```
vEdge1(config)# vpn 0
vEdge1(config-vpn-0)# interface ge0/2
vEdge1(config-interface-ge0/2)# tunnel-interface
vEdge1(config-tunnel-interface)# allow-service bgp
vEdge1(config-tunnel-interface)# commit
Commit complete.
vEdge1(config-tunnel-interface)# end
vEdge1# show bgp summary
vpn                     0
bgp-router-identifier   10.10.10.232
local-as                65001
rib-entries             1
rib-memory              112
total-peers             1
peer-memory             4816
Local-soo               SoO:0:232
ignore-soo
                        MSG       MSG       OUT                    PREFIX  PREFIX  PREFIX
NEIGHBOR          AS    RCVD      SENT      Q      UPTIME          RCVD    VALID   INSTALLED  STATE
-------------------------------------------------------------------------------------------------
192.168.9.1       65000 6         11        0      0:00:00:04      1       1       1          established
```

Case 1. Can not establish BGP peering with my ISP in the underlay

Solution 2: configure explicit ACL as in the example:

Implicit ACL	SD-WAN Explicit ACL	SD-WAN Explicit ACL (default)	Result
✓	X		X
✓		X	✓
X	✓		✓
X		✓	X

```
policy
access-list ALLOW_BGP_PEER
  sequence 10
   match
    source-ip         192.168.9.1/32
    destination-port 179
   !
   action accept
   !
  !
  sequence 20
   match
    source-ip   192.168.9.1/32
    source-port 179
   !
   action accept
   !
  !
  default-action drop
 !
vpn 0
 interface ge0/2
  access-list ALLOW_BGP_PEER in
 !
!
```

- Control plane tunnels are not affected
- Data plane tunnels (“BFD”) are not affected
- Overlay traffic is not affected

Case 2. Why BGP established even if “no allow-service bgp” configured?

Case 2. Why BGP established even if “no allow-service bgp” configured?

Symptoms. BGP peer has established session:

```
cE1#show ip bgp summary | b Neighbor
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.9.1     4        65000      12      11        2    0    0 00:04:21        1
```

But BGP is not allowed under “tunnel-interface” section

```
cE1#show sdwan running-config "sdwan interface GigabitEthernet0/0/0.9 tunnel-interface allow-service"
sdwan
 interface GigabitEthernet0/0/0.9
  tunnel-interface
   no allow-service all
   no allow-service bgp
   no allow-service dhcp
   allow-service dns
   allow-service icmp
   no allow-service sshd
   no allow-service netconf
   no allow-service ntp
   no allow-service ospf
   no allow-service stun
   no allow-service https
   no allow-service snmp
  exit
 exit
!
```

Case 2. Why BGP established even if “no allow-service bgp” configured?

If session is cleared, it got established again:

```
cE1#clear ip bgp *

cE1#show ip bgp summary | b Neighbor
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.9.1     4        65000       0       0        1    0    0 00:00:13 Idle

cE1#show ip bgp summary | b Neighbor
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.9.1     4        65000       6       2        1    0    0 00:00:02        1
```

With packet-trace you can even see few “SdwanImplicitAclDrop” like in Case 1

```
cE1#debug platform condition ipv4 192.168.9.1/32 both
cE1#debug platform packet-trace packet 1024
Please remember to turn on 'debug platform condition start' for packet-trace to work
cE1#show platform packet-trace summary | b SdwanImplicitAclDrop
1     Gi0/0/0.9                Gi0/0/0.9                DROP   479 (SdwanImplicitAclDrop)
2     INJ.2                    Gi0/0/0.9                FWD
3     Gi0/0/0.9                Gi0/0/0.9                DROP   479 (SdwanImplicitAclDrop)
4     Gi0/0/0.9                internal0/0/rp:0         PUNT   11  (For-us data)
5     INJ.2                    Gi0/0/0.9                FWD
6     Gi0/0/0.9                internal0/0/rp:0         PUNT   11  (For-us data)
…
```

Case 2. Why BGP established even if “no allow-service bgp” configured?

And this is indeed BGP packet from the neighbor dropped due to Implicit ACL not allowing BGP service:

```
cE1#show platform packet-trace packet 3 | begin Summary
Summary
  Input     : GigabitEthernet0/0/0.9
  Output    : GigabitEthernet0/0/0.9
  State     : DROP 479 (SdwanImplicitAclDrop)
State     : DROP 479 (SdwanImplicitAclDrop)
Timestamp
    Start   : 965430525669342 ns (05/23/2022 16:08:35.119054 UTC)
    Stop    : 965430525683208 ns (05/23/2022 16:08:35.119068 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : GigabitEthernet0/0/0.9
    Output      : <unknown>
    Source      : 192.168.9.1
    Destination : 192.168.9.23
    Protocol    : 6 (TCP)
      SrcPort   : 179
      DstPort   : 53399
  Feature: SDWAN Implicit ACL
    Action : DISALLOW
    Reason : SDWAN_SERV_BGP
```

Case 2. Why BGP established even if “no allow-service bgp” configured?

Form a theory. What would we check further?

Facts:

- It’s not BGP misconfiguration
- There is no explicit ACL configured to override implicit ACL
- Issue is reproducible every time and consistent across the versions and platfroms (unlikely a bug)
- Only some packets are dropped, not all of them

Case 2. Why BGP established even if “no allow-service bgp” configured?

Let’s check next BGP packet (nr.5) that was allowed.

Compared to previous packets, this one is self-generated by cE1 (injected)

Destination is BGP peer

```
show platform packet-trace packet 5
Packet: 5           CBUG ID: 13887
Summary
  Input     : INJ.2
  Output    : GigabitEthernet0/0/0.9
  State     : FWD
Timestamp
    Start   : 965431957078314 ns (05/23/2022 16:08:36.550463 UTC)
    Stop    : 965431957550849 ns (05/23/2022 16:08:36.550936 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : internal0/0/rp:0
    Output      : <unknown>
    Source      : 192.168.9.23
    Destination : 192.168.9.1
    Protocol    : 6 (TCP)
      SrcPort   : 11600
      DstPort   : 179
  Feature: SDWAN Internal Intf
    VRF ID      : 0 (Global VPN 0)
    Encap Type  : unknown
    IP DSCP     : 48
    IP Version  : 4
    IP Protocol : 6
    Dst Port    : 179
    Is Marked High Priority         : NO
    Is SDWAN Control Tunnel Traffic : NO
    Set HIGH_QUEUE                  : NO (NOT marked high priority, NOT SDWAN
control tunnel traffic)
    Skip SDWAN Policy               : TRUE
```

Case 2. Why BGP established even if “no allow-service bgp” configured?

Let’s check one more packet (nr.6) that was allowed:

```
cE1#show platform packet-trace packet 6
Packet: 6           CBUG ID: 14142
Summary
  Input     : GigabitEthernet0/0/0.9
  Output    : internal0/0/rp:0
  State     : PUNT 11  (For-us data)
  Timestamp
    Start   : 965431958123058 ns (05/23/2022 16:08:36.551508 UTC)
    Stop    : 965431958153208 ns (05/23/2022 16:08:36.551538 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : GigabitEthernet0/0/0.9
    Output      : <unknown>
    Source      : 192.168.9.1
    Destination : 192.168.9.23
    Protocol    : 6 (TCP)
      SrcPort   : 179
      DstPort   : 5062
  Feature: SDWAN Implicit ACL
    Action : ALLOW
    Reason : SDWAN_NAT_DIA
```

This packet is from the BGP peer and intended for local router (punted to CPU) and the reason to allow the packet is clearly shown (NAT DIA)

Case 2. Why BGP established even if “no allow-service bgp” configured?

In this case router with ”implicit ACL” acts like a stateful firewall. And it works that way only if NAT enabled (e.g. for DIA):

```
ip nat route vrf 1 0.0.0.0 0.0.0.0 global
ip nat inside source list nat-dia-vpn-hop-access-list interface GigabitEthernet0/0/1 overload
!
interface GigabitEthernet0/0/1
 ip nat outside
```

Simply said, NAT entry has a precedence over implicit ACL. Explicit SD-WAN ACL can still override this.

Case 2. Why BGP established even if “no allow-service bgp” configured?

Egress self-generated packet (nr.5) creates NAT table entry for BGP session and return traffic is allowed by this entry:

```
cE1#show ip nat translations
Pro  Inside global         Inside local          Outside local         Outside global
tcp  192.168.9.23:5063     192.168.9.23:60456    192.168.9.1:179       192.168.9.1:179
Total number of translations: 1

cE1#show sdwan nat-fwd ip-nat-translation-verbose | b 179
nat-fwd ip-nat-translation-verbose 192.168.9.23 192.168.9.1 60456 179 0 6
 inside-global-addr  192.168.9.23
 outside-global-addr 192.168.9.1
 inside-global-port  5063
 outside-global-port 179
 flags               2113536
 application-type    0
 entry-id            0x3279dbd0
 in_mapping_id       1
 out_mapping_id      0
 create_time         "Mon May 23 16:08:49 2022"
 last_used_time      "Mon May 23 16:22:34 2022"
 pkts_in             40
 pkts_out            37
 timeout             "86381 seconds"
 usecount            1
 input-idb           internal0/0/rp:0
 output-idb          GigabitEthernet0/0/0.9
 bytes_in            1211
 bytes_out           1304
```

Preceding packets are dropped legitimately when remote peer tries to init a session

OMP Troubleshooting Cases

Overlay routing issues

Case 3. vSmart does not advertising any OMP routes

Case 3. Why my vSmart does not advertise any routes?

Symtoms. In the OMP table we can see **only** locally-originated routes (0.0.0.0):

```
ce1#show sdwan omp routes | count C,I,R
Number of lines which match regexp = 0
ce1#show sdwan omp routes | b PATH
                                                PATH                        ATTRIBUTE
VPN     PREFIX                  FROM PEER       ID      LABEL     STATUS    TYPE        TLOC IP     COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------------------------------
1       192.168.40.0/24         0.0.0.0         66      1002      C,Red,R   installed   10.0.0.1    mpls             ipsec  -
                                0.0.0.0         68      1002      C,Red,R   installed   10.0.0.1    biz-internet     ipsec  -
```

Control Connections (CC) and Peering with vSmart is up, but nothing received:

```
ce1#show sdwan control connections | i vsmart|ID|PEER|---
                                                    PEER                    PEER
CONTROLLER
PEER    PEER PEER         SITE   DOMAIN PEER            PRIV  PEER            PUB
TYPE    PROT SYSTEM IP    ID     ID     PRIVATE IP      PORT  PUBLIC IP       PORT  LOCAL COLOR     PROXY STATE UPTIME      GROUP ID
-----------------------------------------------------------------------------------------------------------------------------------
vsmart  dtls 10.0.0.100   1      1      192.168.20.213  12346 192.168.20.213  12346 biz-internet    No    up    0:00:11:49  10
vsmart  dtls 10.0.0.100   1      1      192.168.20.213  12346 192.168.20.213  12346 mpls            No    up    0:00:11:49  10

ce1#show sdwan omp peers
R -> routes received
I -> routes installed
S -> routes sent

                        DOMAIN    OVERLAY   SITE
PEER         TYPE    ID         ID        ID        STATE    UPTIME          R/I/S
-----------------------------------------------------------------------------------
10.0.0.100   vsmart  1          1         1         up       0:00:02:35      0/0/2
```

Case 3. Why my vSmart does not advertise any routes?

Form a theory. What would we check further?

Facts:

- It's not transport or control connections issue because control connections are up and stable
- OMP has zero-line config requirements for basic operations and OMP peering is up
- There is no centralized control policy applied on vSmart in this case (no filtering applied)

Case 3. Why my vSmart does not advertise any routes?

Next let's check on vSmart, peer is "up", and we got 2 routes:

```
vsmart1# show omp peers 10.0.0.1
R -> routes received
I -> routes installed
S -> routes sent

                        DOMAIN    OVERLAY   SITE
PEER       TYPE    ID     ID        ID       STATE    UPTIME           R/I/S
-------------------------------------------------------------------------------
10.0.0.1   vedge   1      1         101      up       0:00:25:16       2/0/0
```

But they are rejected:

```
vsmart1# show omp routes received | include 10.0.0.1
2      192.168.70.0/24     10.0.0.1   66     1003     Rej,Inv,U,Stg installed  10.0.0.1   mpls            ipsec  -
                           10.0.0.1   68     1003     Rej,Inv,U,Stg installed  10.0.0.1   biz-internet    ipsec  -
```

And that router is in "staging" mode:

```
vsmart1# show control valid-vedges | i CSR-5B62E0D4-B8F5-4939-B77C-A9B822019921
CSR-5B62E0D4-B8F5-4939-B77C-A9B822019921 B6159EFC                          staging   CALO - 100589  N/A
```

Case 3. Why my vSmart does not advertise any routes?

Let's check on vManage, certificate is Valid, but...

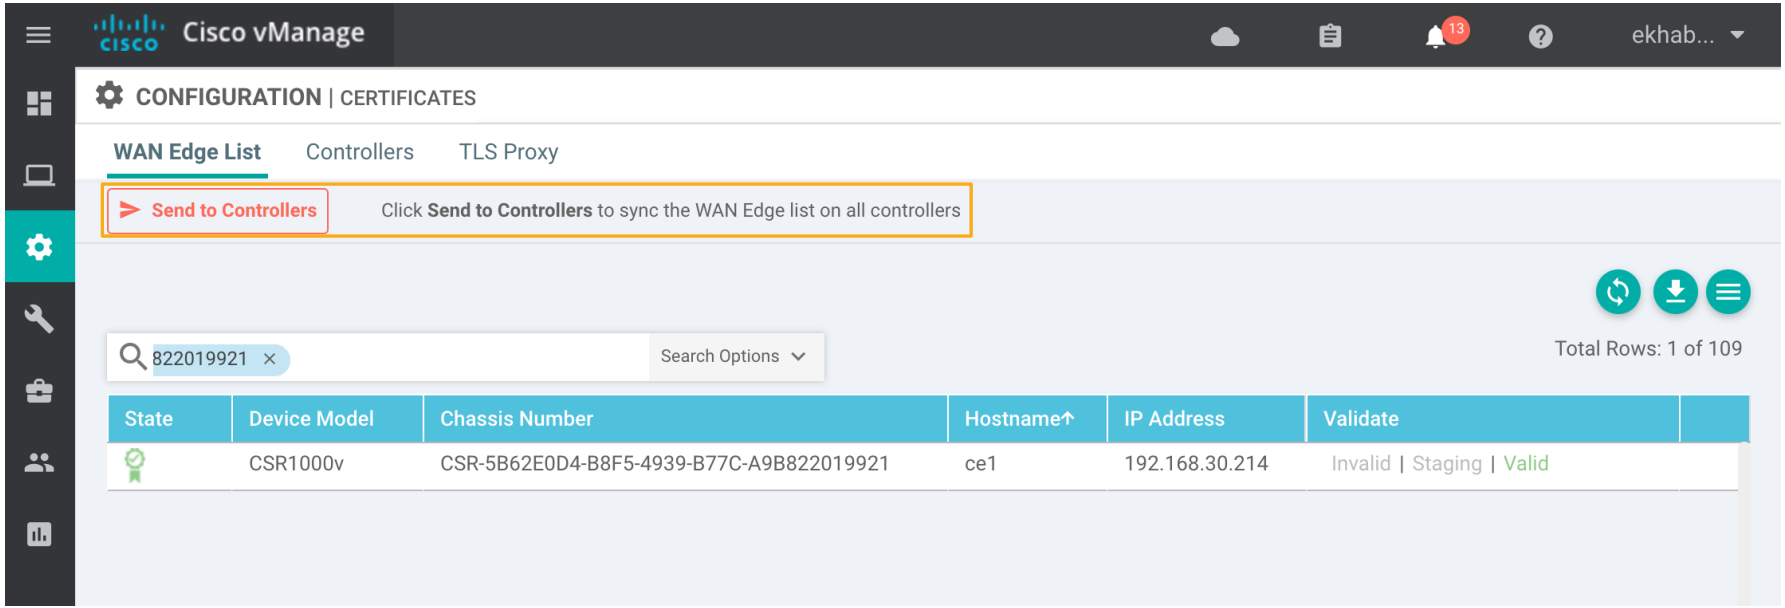

Operator forgot to click "Send to Controllers" after onboarding

Case 4. Traffic is not load-balanced over ECMP paths sent by 2 active-active hubs

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs

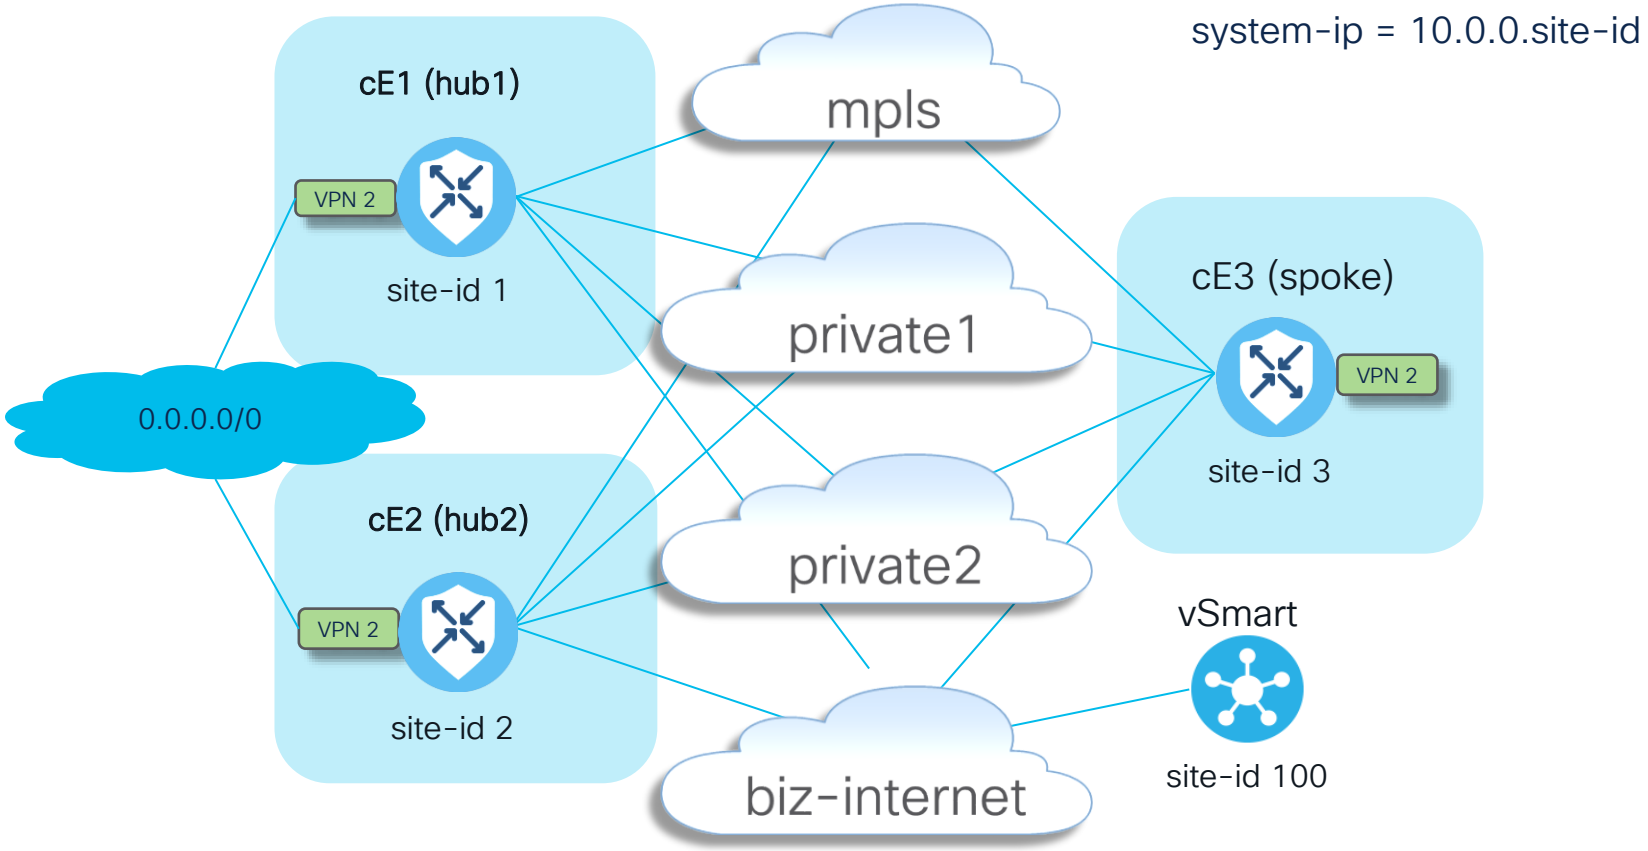

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs

Symproms. Branch router cE3 installs default route only via cE1 (hub1):

```
ce3#show ip route vrf 2 | begin Gateway
Gateway of last resort is 10.0.0.1 to network 0.0.0.0

m*    0.0.0.0/0 [251/0] via 10.0.0.1, 00:08:30, sdwan_system_ip
      192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.3.0/24 is directly connected, Loopback2
L        192.168.3.103/32 is directly connected, Loopback2
```

This is because cE3 receives only 4 paths 0.0.0.0/0 and all resolvable via the same TLOC IP 10.0.0.1

```
ce3#show sdwan omp routes vpn 2 | begin PATH | exclude C,Red,R
                                    PATH                        ATTRIBUTE
VPN     PREFIX              FROM PEER     ID      LABEL    STATUS    TYPE       TLOC IP          COLOR            ENCAP  PREFERENCE
-------------------------------------------------------------------------------------------------------------------------------------
2       0.0.0.0/0           10.0.0.100    61614   1003     C,I,R     installed  10.0.0.1   mpls             ipsec  -
                            10.0.0.100    61615   1003     C,I,R     installed  10.0.0.1   biz-internet     ipsec  -
                            10.0.0.100    61616   1003     C,I,R     installed  10.0.0.1   private1         ipsec  -
                            10.0.0.100    61617   1003     C,I,R     installed  10.0.0.1   private2         ipsec  -
```

4 paths

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs

At the same time vSmart has all 8 routes (4 routes for each TLOC color from each hub):

```
vsmart1# show omp routes vpn 2 | b PATH
                                        PATH                                ATTRIBUTE
VPN     PREFIX              FROM PEER   ID      LABEL     STATUS    TYPE        TLOC IP     COLOR            ENCAP  PREFERENCE
-------------------------------------------------------------------------------------------------------------------------------
2       0.0.0.0/0           10.0.0.1    66      1003      C,R       installed   10.0.0.1    mpls             ipsec  -
                            10.0.0.1    68      1003      C,R       installed   10.0.0.1    biz-internet     ipsec  -
                            10.0.0.1    81      1003      C,R       installed   10.0.0.1    private1         ipsec  -
                            10.0.0.1    82      1003      C,R       installed   10.0.0.1    private2         ipsec  -
                            10.0.0.2    66      1003      C,R       installed   10.0.0.2    mpls             ipsec  -
                            10.0.0.2    68      1003      C,R       installed   10.0.0.2    biz-internet     ipsec  -
                            10.0.0.2    81      1003      C,R       installed   10.0.0.2    private1         ipsec  -
                            10.0.0.2    82      1003      C,R       installed   10.0.0.2    private2         ipsec  -
```

8 paths

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs

- Result: If default route from cE1 (hub1) is lost, spoke router installs route from cE2 (hub2), hence there is no active-active redundancy and rather active-standby with cE1 acting as a primary router.
- We can also check which egress path is taken for specific traffic flow:

```
ce3#show sdwan policy service-path vpn 2  interface Loopback2 source-ip 192.168.3.3 dest-ip 192.168.12.1 protocol 6 source-port 53453
dest-port 22 dscp 48 app ssh
Next Hop: IPsec
  Source: 192.168.9.3 12347 Destination: 192.168.10.1 12427 Local Color: biz-internet Remote Color: mpls Remote System IP: 10.0.0.1
```

- In order to see all available paths for specific traffic type, use **all** keyword:

```
ce3#show sdwan policy service-path vpn 2  interface Loopback2 source-ip 192.168.103.103 dest-ip 192.168.2.1 protocol 6 source-port 53453
dest-port 22 dscp 48 app ssh all
Number of possible next hops: 4
Next Hop: IPsec
  Source: 192.168.9.3 12347 Destination: 192.168.10.1 12427 Local Color: biz-internet Remote Color: mpls Remote System IP: 10.0.0.1
Next Hop: IPsec
  Source: 192.168.8.3 12367 Destination: 192.168.8.1 12407 Local Color: private2 Remote Color: private2 Remote System IP: 10.0.0.1
Next Hop: IPsec
  Source: 192.168.7.3 12367 Destination: 192.168.7..1 12407 Local Color: private1 Remote Color: private1 Remote System IP: 10.0.0.1
Next Hop: IPsec
  Source: 192.168.9.3 12347 Destination: 192.168.9.1 12387 Local Color: biz-internet Remote Color: biz-internet Remote System IP:
10.0.0.1
```

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs

Form a theory. What would we check further?

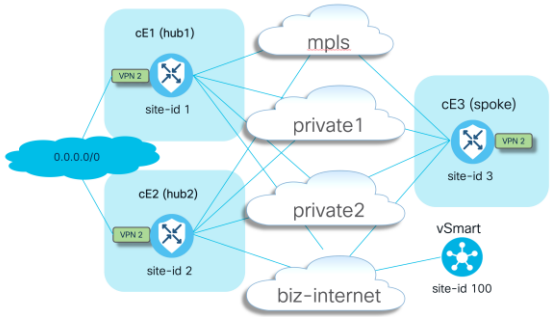

Facts:

- vSmart is a control plane “brain” and route distribution control/filtering/enforcement point
- There is no centralized control policy applied on vSmart to filter something in our case
- OMP has zero-line config requirements for basic operations (and hence has some default settings)

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs

If you check what exactly vSmart advertises, you can find only 4 routes advertised toward cE3:

```
vsmart1# show omp routes vpn 2 0.0.0.0/0 detail | nomore | exclude not\ set | b ADVERTISED\ TO: | b "peer     10.0.0.3" | exclude
label|path-id|overlay|origin
peer     10.0.0.3
    Attributes:
     originator       10.0.0.1
     tloc             10.0.0.1, private2, ipsec
     site-id          1
    Attributes:
     originator       10.0.0.1
     tloc             10.0.0.1, mpls, ipsec
     site-id          1
    Attributes:
     originator       10.0.0.1
     tloc             10.0.0.1, private1, ipsec
     site-id          1
    Attributes:
     originator       10.0.0.1
     tloc             10.0.0.1, biz-internet, ipsec
     site-id          1
```

4 total

And all of them from cE1 (hub1) only

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs. Solution

As you may know or remember already, by default vSmart sends only 4 routes. Let's fix this:

```
vsmart1# conf t
Entering configuration mode terminal
vsmart1(config)# omp
vsmart1(config-omp)# send-path-limit 8
vsmart1(config-omp)# commit
Commit complete.
vsmart1(config-omp)# end
vsmart1# show run omp
omp
 no shutdown
 send-path-limit  8
 graceful-restart
!
vsmart1#
```

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs. Solution

Then, all 8 routes will be advertised by vSmart and received on the branch router:

```
ce3#show sdwan omp routes vpn 2 | begin PATH
                                        PATH                        ATTRIBUTE
VPN     PREFIX              FROM PEER   ID     LABEL   STATUS       TYPE        TLOC IP    COLOR            ENCAP  PREFERENCE
-------------------------------------------------------------------------------------------------------------------------
-
2       0.0.0.0/0           10.0.0.100  61626  1003    C,I,R        installed   10.0.0.1   mpls             ipsec  -
                            10.0.0.100  61627  1003    C,I,R        installed   10.0.0.1   biz-internet     ipsec  -
                            10.0.0.100  61628  1003    C,I,R        installed   10.0.0.1   private1         ipsec  -
                            10.0.0.100  61629  1003    C,I,R        installed   10.0.0.1   private2         ipsec  -
                            10.0.0.100  61637  1003    C,R          installed   10.0.0.2   mpls             ipsec  -
                            10.0.0.100  61638  1003    C,R          installed   10.0.0.2   biz-internet     ipsec  -
                            10.0.0.100  61639  1003    C,R          installed   10.0.0.2   private1         ipsec  -
                            10.0.0.100  61640  1003    C,R          installed   10.0.0.2   private2         ipsec  -
```

8 total

But still branch routers install routes only via cE1 (hub1):

```
ce3#sh ip route vrf 2 0.0.0.0

Routing Table: 2
Routing entry for 0.0.0.0/0, supernet
  Known via "omp", distance 251, metric 0, candidate default path, type omp
  Last update from 10.0.0.1 on sdwan_system_ip, 01:11:26 ago
  Routing Descriptor Blocks:
  * 10.0.0.1 (default), from 10.0.0.1, 01:11:26 ago, via sdwan_system_ip
      Route metric is 0, traffic share count is 1
```

`show sdwan policy service-path` will confirm the same and hence the output is not provided for brevity

Case 4. Traffic is not load-balanced over ECMP paths advertised by 2 active-active hubs. Final Solution

Reason - there are also default settings.
WAN Edge router installs only first 4 equal paths into the routing table. Let's change this:

```
ce3#config-t

admin connected from 127.0.0.1 using console on ce3
ce3(config)# sdwan
ce3(config-sdwan)# omp
ce3(config-omp)# ecmp-limit 8
ce3(config-omp)# commit
Commit complete.
```

`show ip route` confirms both routes via both hubs are installed now:

```
ce3#sh ip ro vrf 2 | b Gateway
Gateway of last resort is 10.0.0.2 to network 0.0.0.0

m*    0.0.0.0/0 [251/0] via 10.0.0.2, 00:00:37, sdwan_system_ip
                [251/0] via 10.0.0.1, 00:00:37, sdwan_system_ip
      192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.3.0/24 is directly connected, Loopback2
L        192.168.3.3/32 is directly connected, Loopback2
```

Case 5. OMP Path selection and global scalability

Case 5. OMP Path selection and global scalability

Case 5. OMP Path selection and global scalability. Objective.

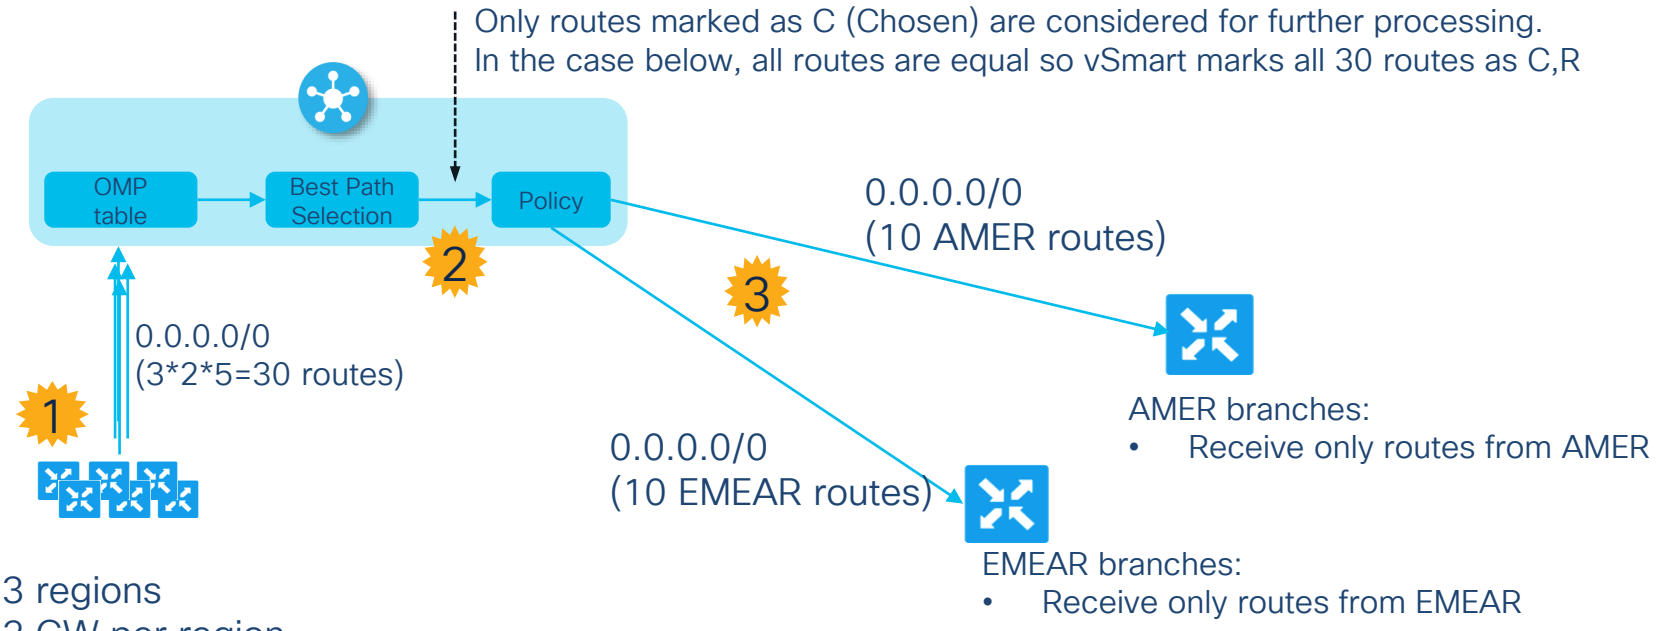

- 3 regions
- 2 GW per region
- each GW with 5 TLOCs
- each advertising 0.0.0.0/0

Case 5. OMP Path selection and global scalability. Output from vSmart.

VPN	PREFIX	FROM PEER	ID	LABEL	STATUS	TYPE	TLOC IP	COLOR	ENCAP	PREFERENCE
1	0.0.0.0/0	10.0.0.10	81	1003	C,R	installed	10.0.0.10	private1	ipsec	-
		10.0.0.10	82	1003	C,R	installed	10.0.0.10	private2	ipsec	-
		10.0.0.10	83	1003	C,R	installed	10.0.0.10	private3	ipsec	-
		10.0.0.10	84	1003	C,R	installed	10.0.0.10	private4	ipsec	-
		10.0.0.10	85	1003	C,R	installed	10.0.0.10	private5	ipsec	-
		10.0.0.11	81	1003	C,R	installed	10.0.0.11	private1	ipsec	-
		10.0.0.11	82	1003	C,R	installed	10.0.0.11	private2	ipsec	-
		10.0.0.11	83	1003	C,R	installed	10.0.0.11	private3	ipsec	-
		10.0.0.11	84	1003	C,R	installed	10.0.0.11	private4	ipsec	-
		10.0.0.11	85	1003	C,R	installed	10.0.0.11	private5	ipsec	-
		10.0.0.12	81	1003	C,R	installed	10.0.0.12	private1	ipsec	-
		10.0.0.12	82	1003	C,R	installed	10.0.0.12	private2	ipsec	-
		10.0.0.12	83	1003	C,R	installed	10.0.0.12	private3	ipsec	-
		10.0.0.12	84	1003	C,R	installed	10.0.0.12	private4	ipsec	-
		10.0.0.12	85	1003	C,R	installed	10.0.0.12	private5	ipsec	-
		10.0.0.13	81	1003	C,R	installed	10.0.0.13	private1	ipsec	-
		10.0.0.13	82	1003	C,R	installed	10.0.0.13	private2	ipsec	-
		10.0.0.13	83	1003	C,R	installed	10.0.0.13	private3	ipsec	-
		10.0.0.13	84	1003	C,R	installed	10.0.0.13	private4	ipsec	-
		10.0.0.13	85	1003	C,R	installed	10.0.0.13	private5	ipsec	-
		10.0.0.14	81	1003	C,R	installed	10.0.0.14	private1	ipsec	-
		10.0.0.14	82	1003	C,R	installed	10.0.0.14	private2	ipsec	-
		10.0.0.14	83	1003	C,R	installed	10.0.0.14	private3	ipsec	-
		10.0.0.14	84	1003	C,R	installed	10.0.0.14	private4	ipsec	-
		10.0.0.14	85	1003	C,R	installed	10.0.0.14	private5	ipsec	-
		10.0.0.15	81	1003	C,R	installed	10.0.0.15	private1	ipsec	-
		10.0.0.15	82	1003	C,R	installed	10.0.0.15	private2	ipsec	-
		10.0.0.15	83	1003	C,R	installed	10.0.0.15	private3	ipsec	-
		10.0.0.15	84	1003	C,R	installed	10.0.0.15	private4	ipsec	-
		10.0.0.15	85	1003	C,R	installed	10.0.0.15	private5	ipsec	-

EMEAR

APAC

AMER

Case 5. OMP Path selection and global scalability. Outputs from Edge router.

`omp ecmp-limit 16` configured to install all 16 equal paths into RIB.

On EMEAR branch router:

VPN	PREFIX	FROM PEER	ID	LABEL	STATUS	TYPE	TLOC IP	COLOR	ENCAP	PREFERENCE
1	0.0.0.0/0	10.0.0.100	1	1003	C,I,R	installed	10.0.0.10	private1	ipsec	-
		10.0.0.100	2	1003	C,I,R	installed	10.0.0.10	private2	ipsec	-
		10.0.0.100	3	1003	C,I,R	installed	10.0.0.10	private3	ipsec	-
		10.0.0.100	4	1003	C,I,R	installed	10.0.0.10	private4	ipsec	-
		10.0.0.100	5	1003	C,R	installed	10.0.0.10	private5	ipsec	-
		10.0.0.100	6	1003	C,R	installed	10.0.0.11	private1	ipsec	-
		10.0.0.100	7	1003	C,R	installed	10.0.0.11	private2	ipsec	-
		10.0.0.100	8	1003	C,R	installed	10.0.0.11	private3	ipsec	-
		10.0.0.100	9	1003	C,R	installed	10.0.0.11	private4	ipsec	-
		10.0.0.100	10	1003	C,R	installed	10.0.0.11	private5	ipsec	-

From EMEAR GW

On AMER branch router:

VPN	PREFIX	FROM PEER	ID	LABEL	STATUS	TYPE	TLOC IP	COLOR	ENCAP	PREFERENCE
1	0.0.0.0/0	10.0.0.100	1	1003	C,I,R	installed	10.0.0.14	private1	ipsec	-
		10.0.0.100	2	1003	C,I,R	installed	10.0.0.14	private2	ipsec	-
		10.0.0.100	3	1003	C,I,R	installed	10.0.0.14	private3	ipsec	-
		10.0.0.100	4	1003	C,I,R	installed	10.0.0.14	private4	ipsec	-
		10.0.0.100	5	1003	C,R	installed	10.0.0.14	private5	ipsec	-
		10.0.0.100	6	1003	C,R	installed	10.0.0.15	private1	ipsec	-
		10.0.0.100	7	1003	C,R	installed	10.0.0.15	private2	ipsec	-
		10.0.0.100	8	1003	C,R	installed	10.0.0.15	private3	ipsec	-
		10.0.0.100	9	1003	C,R	installed	10.0.0.15	private4	ipsec	-
		10.0.0.100	10	1003	C,R	installed	10.0.0.15	private5	ipsec	-

From AMER GW

Case 5. OMP Path selection and scalability. Multiple vSmarts – Failure scenario.

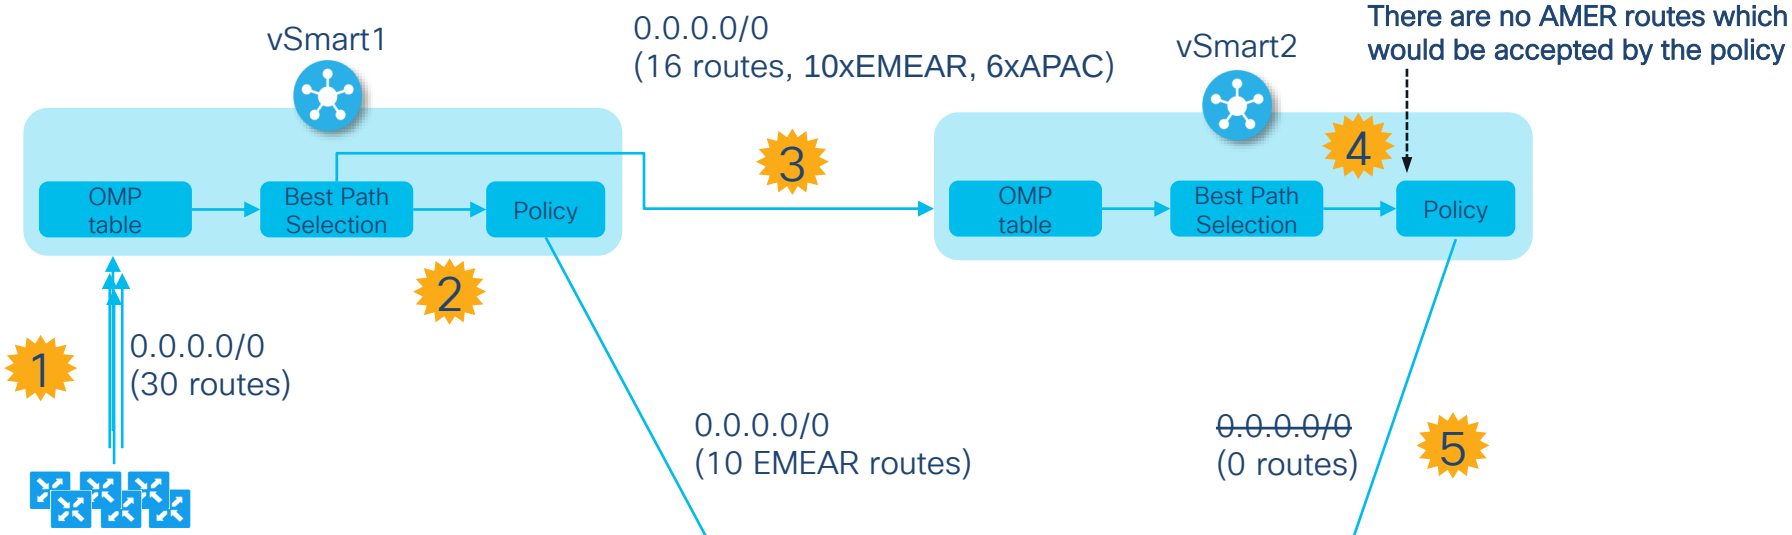

Problem:

- All gateways (GW) are suddenly connected to vSmart1 **only**. This may happen due to communication failure or affinity misconfiguration
- Some branches connected to vSmart2 only

EMEAR branch:

- Connected only to vSmart1

AMER branch:

- Connected only to vSmart2

Case 5. OMP Path selection and scalability. Solution.

- Properly plan affinity groups and redundancy.
- Increase `max-control-connections` (default is 2) to peer with all vSmarts. 🤔 Questionable
- Increase `controller-send-path-limit` (available starting from 20.5). Best option.

```
omp
  controller-send-path-limit [4-128]
```

Case 6: Disjoined underlay active-standby redundancy fault scenario (origin metric)

Case 6. Disjoint Underlay and Origin Metric.

Main objective is to ensure R1 is a preferred path to DC

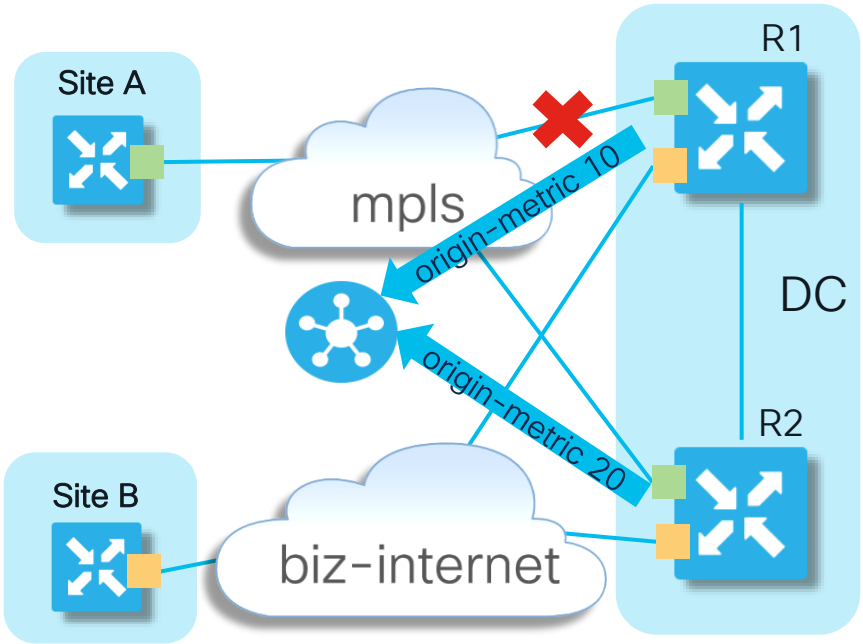

Case 6. Disjoint Underlay and Origin Metric.

Failure scenario. R1 lost MPLS link connection, hence only one route remains and can't be resolved via biz-internet TLOC (no BFD session between SiteA and R1 in DC):

```
SiteA# show omp routes vpn 1 0.0.0.0/0 | nomore | exclude not\ set

-------------------------------------------------
omp route entries for vpn 1 route 0.0.0.0/0
-------------------------------------------------
            RECEIVED FROM:
peer            10.10.10.10
path-id         40
label           1002
status          Inv,U
loss-reason     invalid
lost-to-peer    10.10.10.10
lost-to-path-id 51
    Attributes:
     originator       10.0.0.1
     type             installed
     tloc             10.0.0.1, biz-internet, ipsec
     overlay-id        1
     site-id          1
     preference       0
     origin-proto     static
     origin-metric    10
```

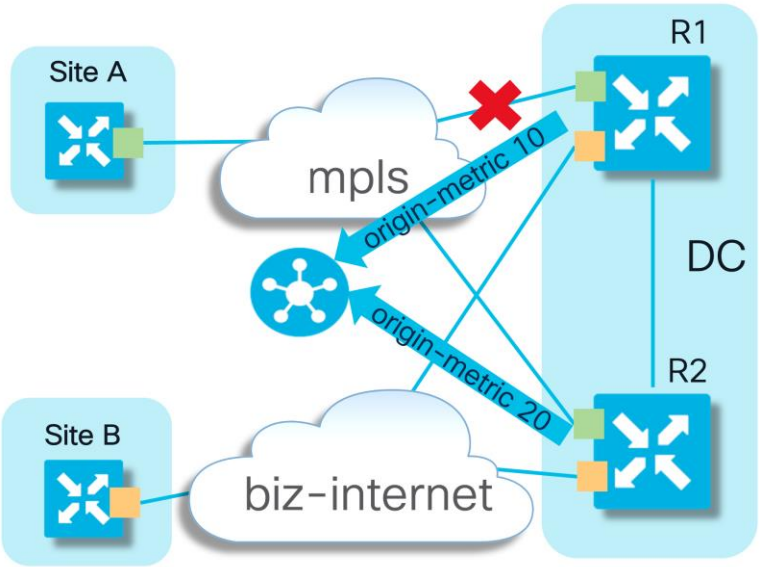

Case 6. Disjoint Underlay and Origin Metric.

Why problem happened?

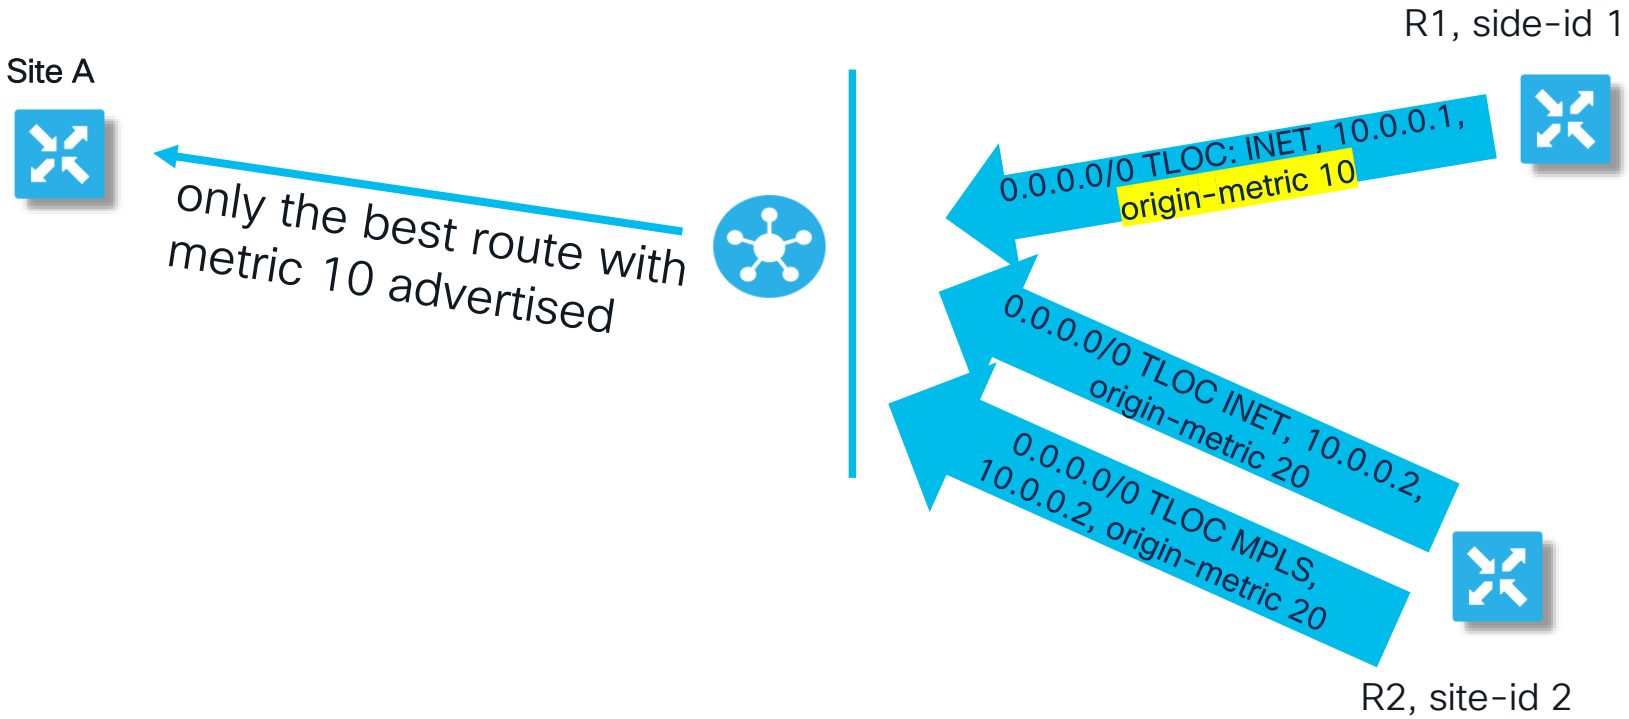

Case 6. Solution 1.

Solution is to remove metric and influence preference only after best path selection (e.g. outbound control policy to set preference instead of origin metric)

```
policy
 lists
  site-list ALL_BRANCHES
   site-id 1
   site-id 2
  !
 !
!
control-policy PREFER_R1
  sequence 10
   match route
    site-id 1
   !
   action accept
    set
     preference 200
    !
   !
  !
  default-action accept
 !
!
apply-policy
 site-list ALL_BRANCHES
  control-policy PREFER_R1 out
 !
```

Case 6. Solution 1 (cont.)

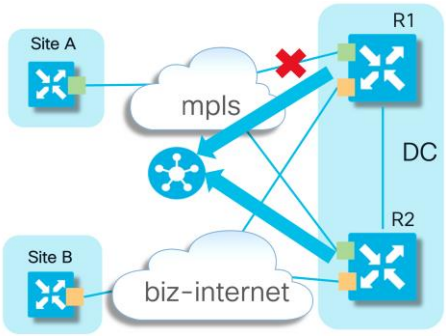

Outbound to branches control-policy configured to prefer R1 and origin-metric removed.

Normal conditions:

```
SiteA# show omp routes vpn 1 0.0.0.0/0 | tab | b PATH
                  PATH                        ATTRIBUTE
FROM PEER         ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.10.10.10       40     1002     Inv,U       installed   10.0.0.1         biz-internet     ipsec  200
10.10.10.10       64     1004     R           installed   10.0.0.2         mpls             ipsec  -
10.10.10.10       65     1004     Inv,U       installed   10.0.0.2         biz-internet     ipsec  -
10.10.10.10       72     1002     C,I,R       installed   10.0.0.1         mpls             ipsec  200
```

Link failure scenario. R1 lost MPLS connection, all fine, we switched to R2!

```
SiteA# show omp routes vpn 1 0.0.0.0/0 | tab | b PATH
                  PATH                        ATTRIBUTE
FROM PEER         ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.10.10.10       40     1002     Inv,U       installed   10.0.0.1         biz-internet     ipsec  200
10.10.10.10       64     1004     C,I,R       installed   10.0.0.2         mpls             ipsec  -
10.10.10.10       65     1004     Inv,U       installed   10.0.0.2         biz-internet     ipsec  -
```

Case 6. Solution 2.

TLOC preference configured on R1 instead of control policy (i.e. directly on interface) and no origin-metric

```
R1(config)# show configuration
vpn 0
 interface ge0/0
  tunnel-interface
   encapsulation ipsec preference 200
   color mpls
  !
 !
 interface ge0/1
  tunnel-interface
   encapsulation ipsec preference 200
   color biz-internet
  !
 !
!
```

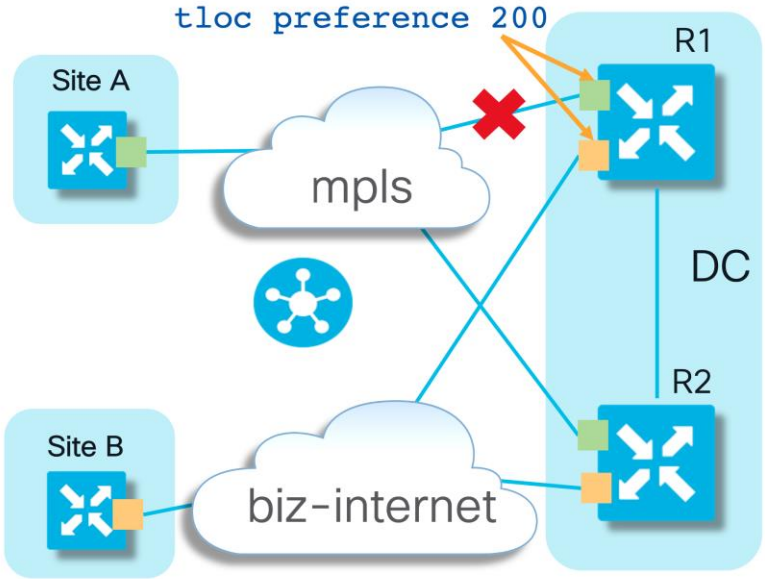

Case 6. Solution 2. (cont.)

Normal conditions:

```
SiteA# show omp route vpn 1 0.0.0.0/0 | nomore | exc "(not\ set|metric|overlay|label|path-id)"
```

```
-----------------------------------------------------
omp route entries for vpn 1 route 0.0.0.0/0
-----------------------------------------------------
           RECEIVED FROM:
peer            10.10.10.10
status          Inv,U
loss-reason     tloc-id
lost-to-peer    10.10.10.10
    Attributes:
     originator       10.0.0.1
     type             installed
     tloc             10.0.0.1, biz-internet, ipsec
     site-id          10
     origin-proto     static
           RECEIVED FROM:
peer            10.10.10.10
status          R
loss-reason     tloc-preference
lost-to-peer    10.10.10.10
    Attributes:
     originator       10.0.0.2
     type             installed
     tloc             10.0.0.2, mpls, ipsec
     site-id          20
     origin-proto     static
```

```
           RECEIVED FROM:
peer            10.10.10.10
status          Inv,U
loss-reason     tloc-id
lost-to-peer    10.10.10.10
    Attributes:
     originator       10.0.0.2
     type             installed
     tloc             10.0.0.2, biz-internet, ipsec
     site-id          20
     origin-proto     static
           RECEIVED FROM:
peer            10.10.10.10
status          C,I,R
    Attributes:
     originator       10.0.0.1
     type             installed
     tloc             10.0.0.1, mpls, ipsec
     site-id          10
     origin-proto     static
```

Preference for routes is **not** set. Path selection enforced with TLOC preference.

Case 6. Solution 2 (cont.)

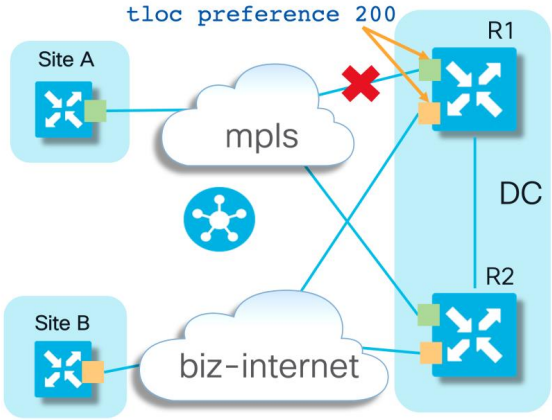

Under normal conditions, Site A receives all TLOCs and preference set to 200 for TLOCs from R1 having System-IP 10.0.0.1

```
SiteA# show omp tlocs received detail | i installed

PSEUDO  ATTRIBUTE  BFD                                                                                                UNKNOWN
KEY     TYPE       STATUS  DOMAIN ID  SITE ID  PREFERENCE  TAG  STALE  CARRIER  GROUPS  BORDER  ATTRIBUTE  LEN  WEIGHT  GEN ID      VERSION  ORIGINATOR  RESTRICT
OVERLAY ID
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------
---------
1       installed  up      -          10       200         -    -      default  [ 0 ]   -       -          1    0x80000012  2        10.0.0.1    1       -
1       installed  down    -          10       200         -    -      default  [ 0 ]   -       -          1    0x80000012  2        10.0.0.1    0       -
1       installed  up      -          20       0           -    -      default  [ 0 ]   -       -          1    0x80000033  2        10.0.0.2    1       -
1       installed  down    -          20       0           -    -      default  [ 0 ]   -       -          1    0x80000033  2        10.0.0.2    0       -
1       installed  up      -          1        0           -    -      default  [ 0 ]   -       -          1    0x8000002f  2        10.0.0.10   1       -
1       installed  down    -          2        0           -    -      default  [ 0 ]   -       -          1    0x8000001c  2        10.0.0.20   0       -
```

Case 6. Solution 2 (cont.)

R1 MPLS link failure scenario, R2 route selected because it's the only resolved, all OK:

```
SiteA# show omp routes vpn 1 0.0.0.0/0 | nomore | exclude "(not\ set|metric|overlay|label|path-id)"
```

```
-----------------------------------------------
omp route entries for vpn 1 route 0.0.0.0/0
-----------------------------------------------
            RECEIVED FROM:                                      RECEIVED FROM:
peer            10.10.10.10                         peer            10.10.10.10
status          Inv,U                               status          Inv,U
loss-reason     tloc-id                             loss-reason     tloc-id
lost-to-peer    10.10.10.10                         lost-to-peer    10.10.10.10
    Attributes:                                         Attributes:
     originator       10.0.0.1                          originator       10.0.0.2
     type             installed                         type             installed
     tloc             10.0.0.1, biz-internet, ipsec     tloc             10.0.0.2, biz-internet, ipsec
     site-id          10                                site-id          20
     origin-proto     static                            origin-proto     static
            RECEIVED FROM:
peer            10.10.10.10
status          C,I,R
    Attributes:
     originator       10.0.0.2
     type             installed
     tloc             10.0.0.2, mpls, ipsec
     site-id          20
     origin-proto     static
```

Case 6. Solution 3.

Send non-best paths as well (and keep origin-metric):

```
vsmart1# show running-config omp
omp
 no shutdown
 send-backup-paths
```

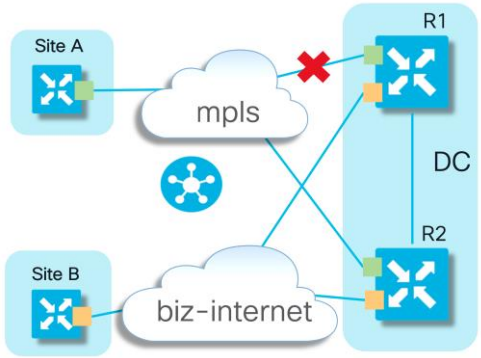

Under normal conditions:

```
SiteA# show omp routes vpn 1 0.0.0.0/0 | tab | b PATH
                PATH                        ATTRIBUTE
FROM PEER       ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-------------------------------------------------------------------------------------------------------------
10.10.10.10     40     1002     Inv,U       installed   10.0.0.1         biz-internet     ipsec  -
10.10.10.10     41     1004     R           installed   10.0.0.2         mpls             ipsec  -
10.10.10.10     42     1004     Inv,U       installed   10.0.0.2         biz-internet     ipsec  -
10.10.10.10     49     1002     C,I,R       installed   10.0.0.1         mpls             ipsec  -
```

4 paths

R1 MPLS link failure scenario, R2 selected, all OK:

```
SiteA# show omp routes vpn 1 0.0.0.0/0 | tab | b PATH
                PATH                        ATTRIBUTE
FROM PEER       ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-------------------------------------------------------------------------------------------------------------
10.10.10.10     40     1002     Inv,U       installed   10.0.0.1         biz-internet     ipsec  -
10.10.10.10     41     1004     C,I,R       installed   10.0.0.2         mpls             ipsec  -
10.10.10.10     42     1004     Inv,U       installed   10.0.0.2         biz-internet     ipsec  -
```

Case 7 vSmart double failure scenario

or why you may not want to enable graceful-restart

Case 7 vSmart double failure scenario

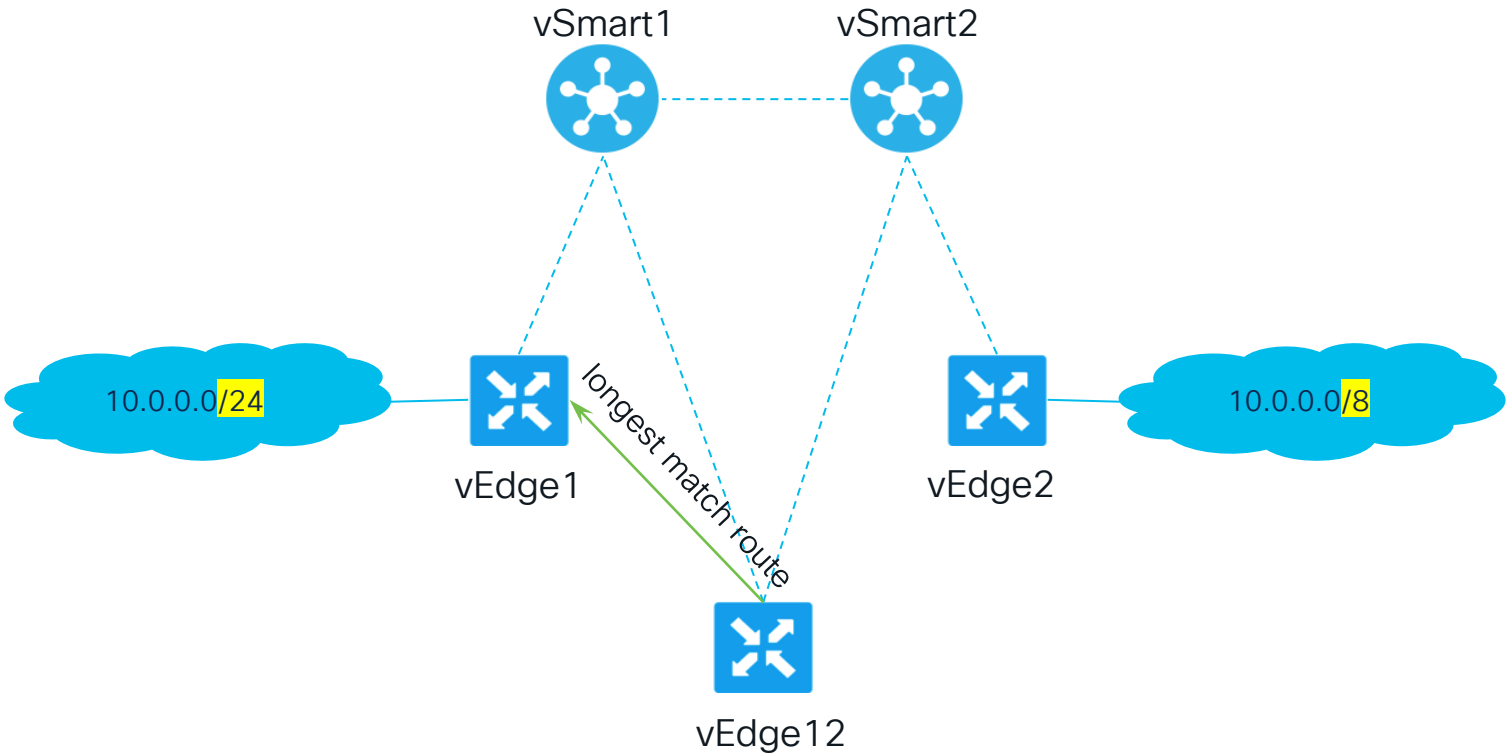

Case 7 vSmart double failure scenario

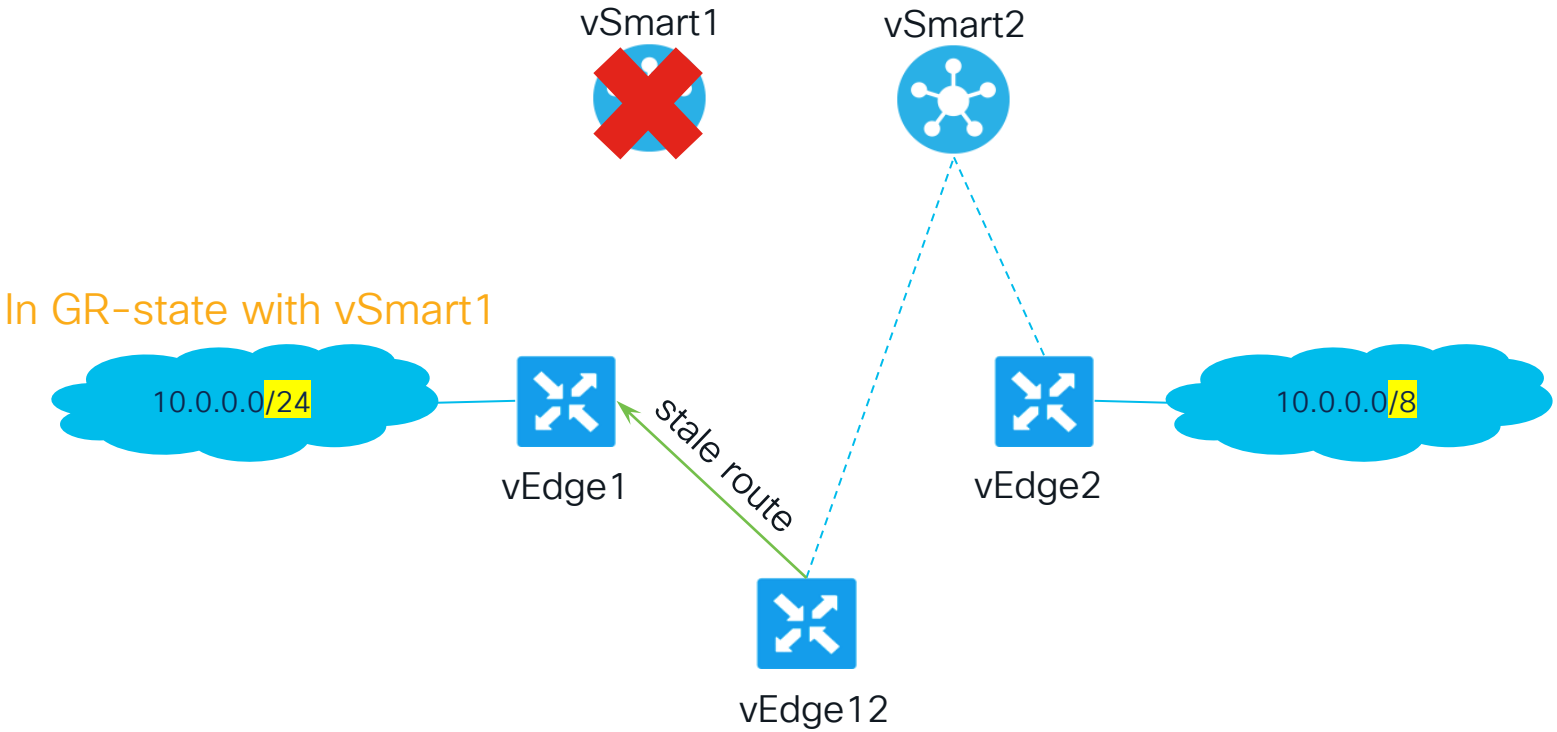

Case 7 vSmart double failure scenario

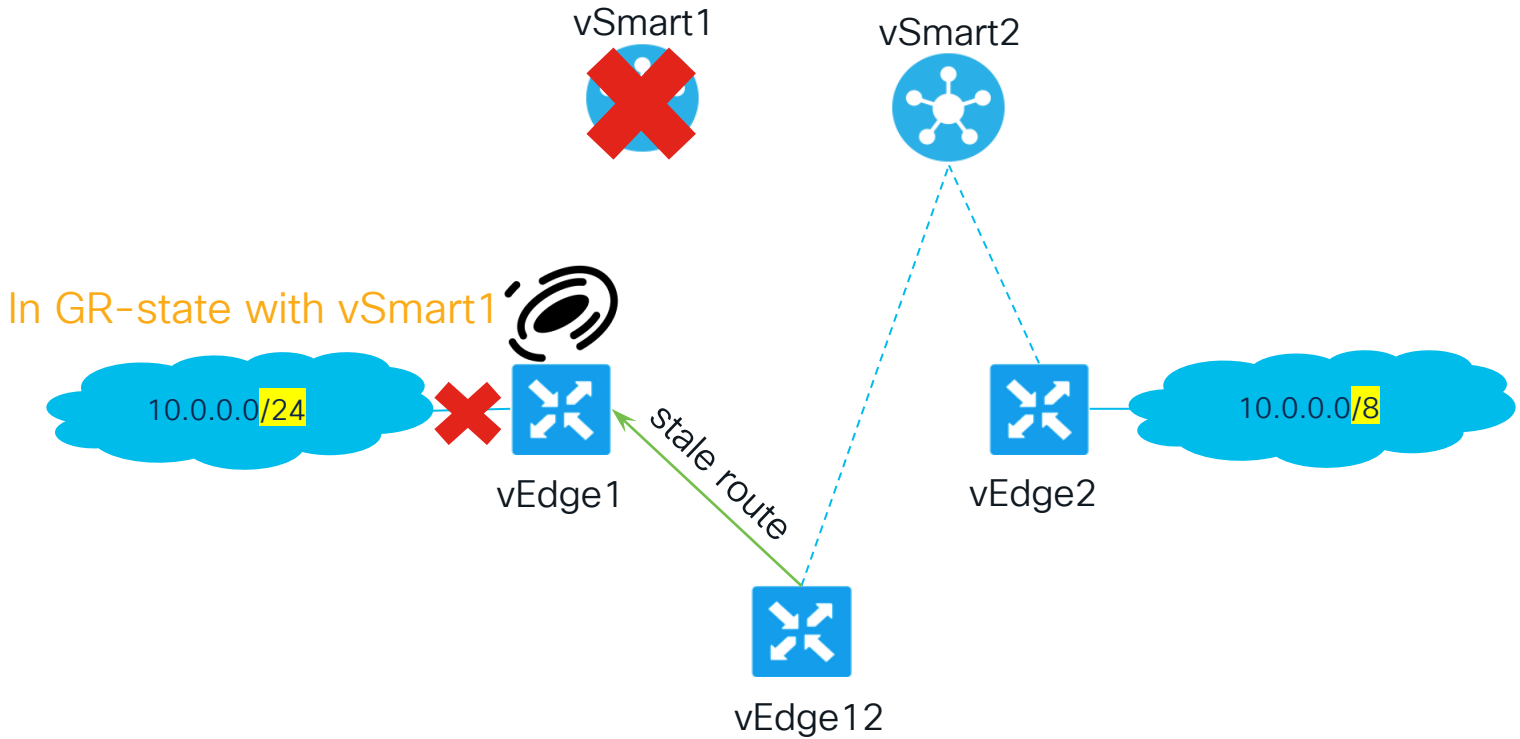

SD-WAN can't handle such scenarios with GR configured, this is expected

Case 7 vSmart double failure scenario. Solutions.

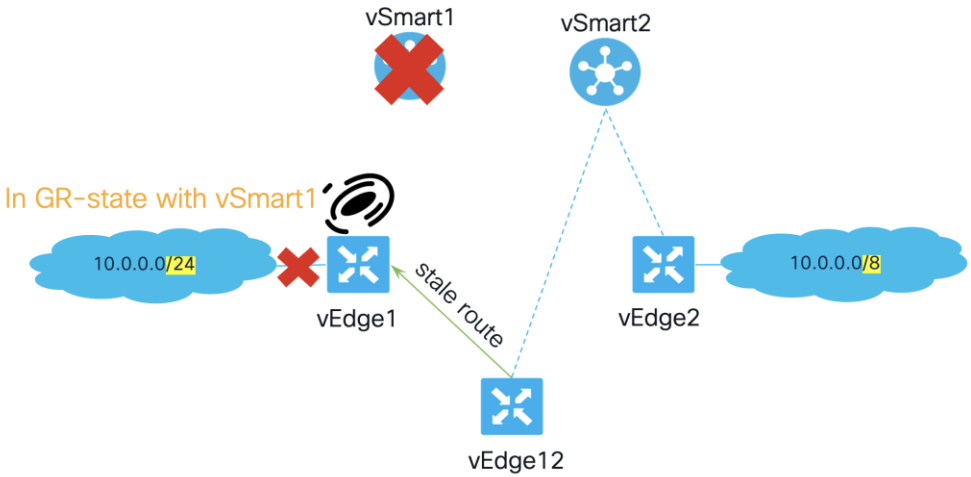

1. Be careful with summarization.
2. Ensure vSmart redundancy, geo-reservation and underlay paths diversity.
3. Properly plan contollers affinity if you use it, mind (2) also.
4. Don’t use GR 🤔 Questionable.

Q&A

References
and recommended resources

- Cisco Troubleshooting Tech Notes: https://www.cisco.com/c/en/us/support/routers/sd-wan/products-tech-notes-list.html

Cisco Networking Bot

https://cnb.cisco.com/

Empowering Users by Digitizing Cisco Product & Adoption Experiences

EoS/EoL

Product Migration/Adoption

Config / Tshoot Guides

Security Advisories

Release Recommendations

HW-SW Compatibility

Transforming Customer Experience with Cisco AI Chatbots

Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.

Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

Learn

Cisco U.
IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning
Subscription-based product, technology, and certification training

Cisco Modeling Labs
Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network
Resource community portal for certifications and learning

Train

Cisco Training Bootcamps
Intensive team & individual automation and technology training programs

Cisco Learning Partner Program
Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training
Accelerated curriculum of product, technology, and certification courses

Certify

Cisco Certifications and Specialist Certifications
Award-winning certification program empowers students and IT Professionals to advance their technical careers

Cisco Guided Study Groups
180-day certification prep program with learning and support

Cisco Continuing Education Program
Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

cisco Live!

Thank you


cisco Live!
ALL IN