Resolve BGP-OMP Routing Loop when BGP Routing and SoO Are Used

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Background Information

Configuration

Verify

Troubleshoot

- Solution 1
  - Overlay-AS Case 1
  - Overlay-AS Case 2
- Solution 2

SoO Loop Prevention Explanation

Related Information

Introduction

This document describes how to avoid a routing loop in SD-WAN fabric when Border Gateway Protocol (BGP) routing and Site of Origin (SoO) is used.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Basic understanding of Overlay Management Protocol (OMP)
- Basic understanding of BGP
- SD-WAN components and interaction between them

Components Used

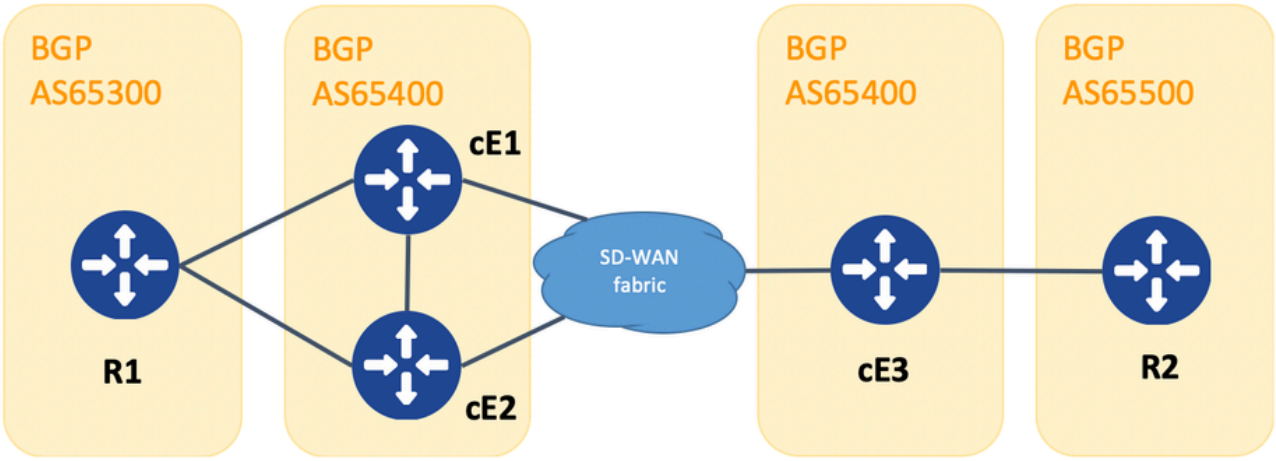
The information in this document is based on these software and hardware versions:

- 3 Cisco IOS® XE CSR1000v routers with Software Release 17.2.1v that run in controller mode (SD-WAN)
- 2 Cisco IOS XE CSR1000v routers with Software Release 16.7.3

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

For the purpose of this document, this topology is used:



Topology

R1 and R2 are generic Cisco IOS XE routers (or any other router capable of running BGPv4). cE1, cE2, and cE3 run Cisco IOS XE in controller (SD-WAN) mode. Here you can find a summary of assigned site-id and system-ip parameters to each SD-WAN router:

SD-WAN router	site-id	system-ip
cE1	214	192.168.30.214
cE2	215	192.168.30.215
cE3	216	192.168.30.216

Here is a set of events that took place initially:

1. R1 and R2 establish eBGP peering correspondingly with cE1, cE2, and cE3. cE1 and cE2 establish iBGP peering.
2. R2 originates BGP route 10.1.1.0/24 and advertises it via eBGP to cE3.
3. cE3 receives this BGP route on the service side in the VRF 1 address family and then redistributes this route into OMP.
4. cE3 advertises the 10.1.1.0/24 OMP route to the SD-WAN overlay (vSmart controllers are responsible for routing information dissemination via the OMP protocol to all other Edge routers joined to the SD-WAN overlay).
5. cE1 and cE2 receive the OMP route and redistribute it back via eBGP in VRF 1 to R1.

Configuration

Here is the relevant configuration of cE1. Note that **send-comminity** is not configured for neighbor 192.168.160.215:

```
router bgp 65401
```

```
 bgp log-neighbor-changes
 distance bgp 20 200 20
 !
 address-family ipv4 vrf 1
  redistribute omp
  propagate-aspath
  neighbor 192.168.140.10 remote-as 65300
  neighbor 192.168.140.10 activate
  neighbor 192.168.140.10 send-community both
  neighbor 192.168.160.215 remote-as 65400
  neighbor 192.168.160.215 activate
 exit-address-family
 !
sdwan
 omp
  no shutdown
  send-path-limit  4
  ecmp-limit       4
  graceful-restart
  no as-dot-notation
  timers
   holdtime               60
   advertisement-interval 1
   graceful-restart-timer 43200
   eor-timer              300
  exit
  address-family ipv4 vrf 1
   advertise bgp
  !
  address-family ipv4
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise connected
   advertise static
```

cE2:

```
router bgp 65401
 bgp log-neighbor-changes
 distance bgp 20 200 20
 !
 address-family ipv4 vrf 1
  redistribute omp
  propagate-aspath
  neighbor 192.168.150.10 remote-as 65300
  neighbor 192.168.150.10 activate
  neighbor 192.168.150.10 send-community both
  neighbor 192.168.160.214 remote-as 65401
  neighbor 192.168.160.214 activate
  neighbor 192.168.160.214 send-community both
 exit-address-family
 !
sdwan
 omp
  no shutdown
  send-path-limit  4
```

```
  ecmp-limit       4
  graceful-restart
  no as-dot-notation
  timers
   holdtime               60
   advertisement-interval 1
   graceful-restart-timer 43200
   eor-timer              300
  exit
  address-family ipv4 vrf 1
   advertise bgp
  !
  address-family ipv4
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise connected
   advertise static
```

cE3:

```
router bgp 65401
 bgp log-neighbor-changes
 timers bgp 5 15
 !
 address-family ipv4 vrf 1
  redistribute omp
  propagate-aspath
  neighbor 192.168.60.11 remote-as 65500
  neighbor 192.168.60.11 activate
 exit-address-family
!
sdwan
 omp
  no shutdown
  send-path-limit  4
  ecmp-limit       4
  graceful-restart
  no as-dot-notation
  timers
   holdtime               60
   advertisement-interval 1
   graceful-restart-timer 43200
   eor-timer              300
  exit
  address-family ipv4 vrf 1
   advertise bgp
  !
  address-family ipv4
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise connected
   advertise static
  !
```

Verify

1. In the initial state, the route is advertised from cE3 and learned by cE1 and cE2 via OMP. Both redistribute the route to BGP and announce to each other and to R1:

cE1#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 342041
Paths: (2 available, best #2, table 1)
  Advertised to update-groups:
     4          5
  Refresh Epoch 1
  65500
    192.168.160.215 (via vrf 1) from 192.168.160.215 (192.168.109.215)
      Origin incomplete, metric 1000, localpref 50, valid, internal
      Extended Community: SoO:0:215 RT:1:1
      rx pathid: 0, tx pathid: 0
      Updated on Aug 21 2020 11:23:32 GMT
  Refresh Epoch 1
  65500
    192.168.30.216 (via default) from 0.0.0.0 (192.168.109.214)
      Origin incomplete, metric 1000, localpref 50, valid, sourced, best
      Extended Community: SoO:0:214 RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Aug 21 2020 11:23:32 GMT
```

cE2#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 327810
Paths: (2 available, best #2, table 1)
  Advertised to update-groups:
     5          6
  Refresh Epoch 1
  65500
    192.168.160.214 (via vrf 1) from 192.168.160.214 (192.168.109.214)
      Origin incomplete, metric 1000, localpref 50, valid, internal
      Extended Community: RT:1:1
      rx pathid: 0, tx pathid: 0
      Updated on Aug 21 2020 11:23:32 GMT
  Refresh Epoch 1
  65500
    192.168.30.216 (via default) from 0.0.0.0 (192.168.109.215)
      Origin incomplete, metric 1000, localpref 50, valid, sourced, best
      Extended Community: SoO:0:215 RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Aug 21 2020 11:23:32 GMT
```

2. The WAN interface is disconnected or connectivity to SD-WAN fabric is lost on cE2, hence OMP peers (vSmart connections) go down. Only one route remains learned from iBGP:

```
ce2(config)#
interface GigabitEthernet 2
ce2(config-if)#
shutdown
ce2(config-if)#
end
Uncommitted changes found, commit them? [yes/no/CANCEL] yes
Commit complete.
```

ce2#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 345276
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     6
  Refresh Epoch 1
  65500
    192.168.160.214 (via vrf 1) from 192.168.160.214 (192.168.109.214)
      Origin incomplete, metric 1000, localpref 50, valid, internal, best
      Extended Community: RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Aug 21 2020 11:23:32 GMT

cE1 still prefers the route via OMP (this is the only route that remains) originated by cE3:
```

ce1#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 342041
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     4          5
```

```
  Refresh Epoch 1
  65500
    192.168.30.216 (via default) from 0.0.0.0 (192.168.109.214)
      Origin incomplete, metric 1000, localpref 50, valid, sourced, best
      Extended Community: SoO:0:214 RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Aug 21 2020 11:23:32 GMT
```

3. Connectivity on the WAN interface of cE2 is established again. The route is still preferred from cE1 via iBGP because of better Administrative Distance (AD).

```
ce2(config)#
interface GigabitEthernet 2
ce2(config-if)#
no shutdown
ce2(config-if)#
end
Uncommitted changes found, commit them? [yes/no/CANCEL] yes
Commit complete.
```

ce2#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 345276
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     6
  Refresh Epoch 1
  65500
    192.168.160.214 (via vrf 1) from 192.168.160.214 (192.168.109.214)
      Origin incomplete, metric 1000, localpref 50, valid, internal, best
      Extended Community: RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Aug 21 2020 11:23:32 GMT
```

cE1 still prefers the route via OMP originated by cE3. Keep in mind that cE1 redistributes OMP into BGP:

ce1#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 569358
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     4          5
  Refresh Epoch 1
  65500
    192.168.30.216 (via default) from 0.0.0.0 (192.168.109.214)
      Origin incomplete, metric 1000, localpref 50, valid, sourced, best
      Extended Community: SoO:0:214 RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Aug 21 2020 15:13:09 GMT
```

4. Something happens with cE3 connectivity to R2. To test, the interface is shut down, and the R2 BGP peer is lost:

```
ce3(config)#
interface GigabitEthernet 6
ce3(config-if)#
shutdown
ce3(config-if)#
commit
```

5. As a result, the routing loop is formed between cE1 and cE2 (cE2 redistributes the route from OMP and advertises to cE1 via BGP, cE1 redistributes BGP to OMP and advertises to cE2):

ce1#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 732548
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     5
  Refresh Epoch 1
  65500
    192.168.160.215 (via vrf 1) from 192.168.160.215 (192.168.109.215)
      Origin incomplete, metric 1000, localpref 50, valid, internal, best
      Extended Community: SoO:0:215 RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Aug 21 2020 15:38:47 GMT
```

ce2#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1/24

BGP routing table entry for 1:1:10.1.1.1/24, version 639650
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     5          6
  Refresh Epoch 1
  65500
    192.168.30.214 (via default) from 0.0.0.0 (192.168.109.215)
      Origin incomplete, metric 1000, localpref 50, valid, sourced, best
      Extended Community: SoO:0:215 RT:1:1
      rx pathid: 1, tx pathid: 0x0
      Updated on Aug 21 2020 15:38:47 GMT
```

Troubleshoot

There are two possible solutions.

Solution 1

Configure **overlay-as** for OMP. Then some Autonomous System (AS) number is assigned to the OMP overlay itself. For example:

```
config-transaction
 sdwan
  omp
   overlay-as 64512
   exit
```

By default, OMP is transparent to BGP even if **propagate-aspath** is configured. **overlay-as** is a feature that prepends AS specified as a parameter of this command to the BGP AS_PATH attribute of routes exported from OMP to BGP. If you configure the same overlay AS number on multiple devices in the overlay network, all these devices are considered to be part of the same AS. As a result, they do not forward any routes that contain the overlay AS number, hence the routing loop is prevented.

Keep in mind that **overlay-as** and **propagate-aspath** are dependent on each other. This feature is discussed in detail.

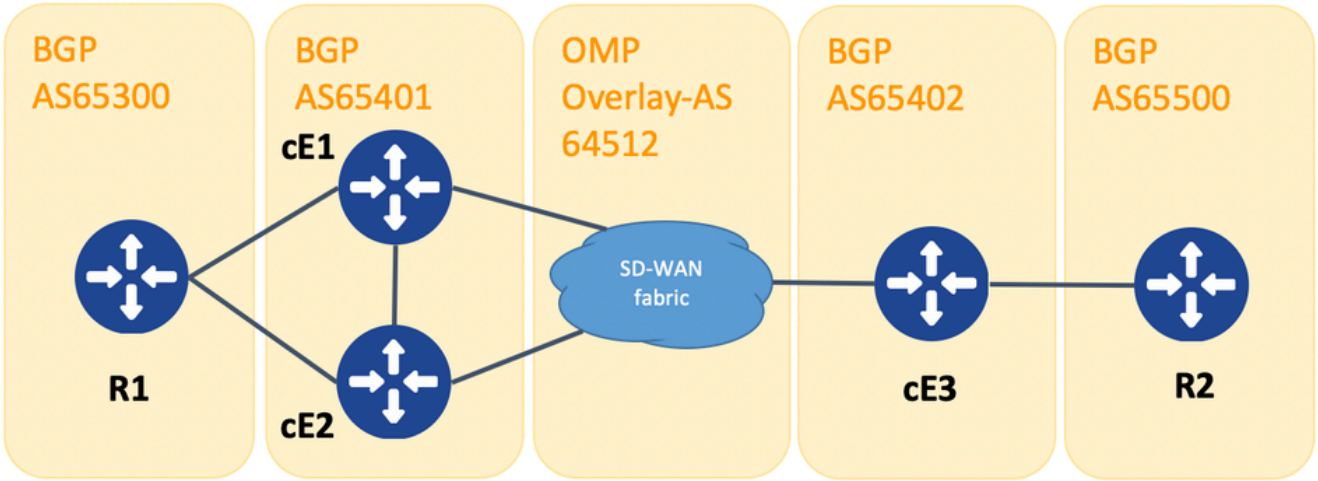
There are two cases that exist.

Overlay-AS Case 1

overlay-as configured on the global level under **sdwan omp** section and **propagate-aspath** is not configured (rest configuration is the same as described initially: **advertise bgp** is enabled under **omp address-family ipv4 vrf 1** section,

redistribute omp configured under **router bgp <AS> address-family ipv4 vrf 1** section).

overlay-as 64512, configured on cE1/cE2 and cE3.



Topology for overlay-as demonstration

For the purpose of demonstration, BGP AS on cE1, cE2, and cE3 changed.

R1 - cE1/cE2 still peer via eBGP, AS **65300** and **65401** are used respectively.

cE3 - R2 still peer via eBGP, AS **65402** and **65500** are used respectively.

R1 sends route (for example, 192.168.41.11/32) to cE1/cE2. cE1/cE2 redistributes this route into OMP, without any AS_PATH attribute.

cE3 receives it and advertises it into BGP towards R2, only with its own AS (normal eBGP behavior).

Route route1 on R2 has AS_PATH: **65402**.

R2#

```
sh ip bgp | i 192.168.41.11/32

 *>   192.168.41.11/32 192.168.60.216        1000             0 65402 ?
```

Overlay-AS Case 2

propagate-aspath configured under **router bgp** section for the particular service side VPN (**address-family ipv4 vrf 1**). Here are sub-cases as well.

Case 2.1. With **overlay-as** enabled on cE3, **propagate-aspath** is also enabled under **router bgp 65401 address-family ipv4 vrf 1** on cE1/cE2.

R1 sends route route1 to cE1/cE2. cE1/cE2 redistributes this route into OMP with an as-path that comes from the R1 site.

OMP route on vSmart has AS-Path: **65300**.

vsmart1#

```
show omp routes vpn 1 192.168.41.11/32 | nomore | exclude not\ set

-------------------------------------------------
omp route entries for vpn 1 route 192.168.41.11/32
-------------------------------------------------
            RECEIVED FROM:
peer            192.168.30.214
path-id         81
label           1001
status          C,R
    Attributes:
     originator       192.168.30.214
     type             installed
     tloc             192.168.30.214, biz-internet, ipsec
     overlay-id        1
     site-id          25
     origin-proto     eBGP
     origin-metric    0
     as-path          "65300"
            RECEIVED FROM:
peer            192.168.30.215
path-id         68
label           1002
status          C,R
    Attributes:
     originator       192.168.30.215
     type             installed
     tloc             192.168.30.215, biz-internet, ipsec
     overlay-id        1
     site-id          25
     origin-proto     eBGP
     origin-metric    0
     as-path          "65300"
```

Case 2.1.a. With **propagate-aspath** disabled on cE3, cE3 receives it as an OMP route and advertises it into BGP, ignores any as-path attribute, overlays as, towards R2, and adds only its own BGP AS (normal eBGP behavior).

Route route1 on R2 AS-path: **65402**.

R2#

```
sh ip bgp | i 192.168.41.11/32

 *>   192.168.41.11/32 192.168.60.216        1000             0 65402 ?
```

Case 2.1.b. With **propagate-aspath** enabled on cE3, cE3 receives it as an OMP route and advertises it into BGP,

prepends the received as-path attribute, towards R2 then adds the Overlay-AS followed by its own BGP AS.

Route route1 on R2 AS-path: **65402 64512 65300**.

R2#

```
sh ip bgp | i 192.168.41.11/32

 *>   192.168.41.11/32 192.168.60.216        1000             0 65402 64512 65300 ?
```

Case 2.1.c. With **propagate-aspath** disabled on cE1/cE2, cE3 receives it as an OMP route without any as-path attribute and advertises it into BGP, towards R2, prepends the Overlay-AS and adds only its own BGP AS.

Route route1 on R2 AS-path: **6540264512**.

R2#

```
sh ip bgp | i 192.168.41.11/32

 *>   192.168.41.11/32 192.168.60.216        1000             0 65402 64512 ?
```

Case 2.2. Without **overlay-as** configured on cE3, **propagate-aspath** is enabled under **router bgp 65401 address-family ipv4 vrf 1** on cE1/cE2.

Case 2.2.a. With **propagate-aspath** disabled on cE3 only, cE3 receives it as an OMP route and advertises it into BGP, ignoring any AS_PATH attribute, towards R2, adds its own BGP AS (normal eBGP behavior).

Route route1 on R2 AS-path: **65402**.

R2#

```
sh ip bgp | i 192.168.41.11/32

 *>   192.168.41.11/32 192.168.60.216        1000             0 65402 ?
```

Case 2.2.b. When **propagate-aspath** is enabled on cE3, cE3 receives it as an OMP route and advertises it into BGP, prepends the received AS_PATH attribute, towards R2 then adds its own AS.

Route route1 on R2 AS-path: **6540265300**.

R2#

```
sh ip bgp | i 192.168.41.11/32
```

```
 *>   192.168.41.11/32 192.168.60.216        1000             0 65402 65300 ?
```



Note: When you send the AS-Path attribute into OMP, the Edge router does not add its own AS (as demonstrated in the article vEdge Does Not Advertise Its Own AS When BGP Routes Are Advertised Into OMP). If the remote Edge router receives an OMP route with its own AS in the AS_PATH attribute, it does not perform loop detection and it sends the route with the received AS path to the router on the service side.

Solution 2

Configure the same **site-id** on routers cE1 and cE2. Although the vSmart advertises routes back to the site with the same site-id as in the route itself, since the originator attribute of the route is different, loop prevention is not triggered, but the control plane routing loop does not form because the OMP route is not installed into the RIB. This is because the OMP route stays in the Inv,U (Invalid,Unresolved) state. By default, data plane tunnels cannot be established between sites that have the same site-id unless **allow-same-site-tunnels** is configured. If the data plane tunnel BFD session is in the down state, TLOC remains unresolved. In the example here, **site-id 214215** was configured on both routers ce1 and ce2. Route 10.0.0.2/32 advertised by cE2 and cE1 does not install it into the routing table because no data plane sessions exist between cE1 and cE2:

ce1#

```
show sdwan omp route 10.0.0.2/32 det | exc not set

-------------------------------------------------
omp route entries for vpn 3 route 10.0.0.2/32
-------------------------------------------------
            RECEIVED FROM:
peer            192.168.30.113
path-id         3
label           1004
status          Inv,U
    Attributes:
     originator       192.168.30.215
     type             installed
     tloc             192.168.30.215, mpls, ipsec
     overlay-id        1
     site-id          214215
     origin-proto     connected
     origin-metric    0
            RECEIVED FROM:
peer            192.168.30.113
path-id         4
label           1004
status          Inv,U
loss-reason     tloc-id
lost-to-peer    192.168.30.113
lost-to-path-id 3
    Attributes:
     originator       192.168.30.215
```

```
     type             installed
     tloc             192.168.30.215, biz-internet, ipsec
     overlay-id        1
     site-id          214215
     origin-proto     connected
     origin-metric    0
```

ce1#

```
show sdwan omp tlocs "ip 192.168.30.215" | exclude not set

----------------------------------------------------
tloc entries for 192.168.30.215
                 mpls
                 ipsec
----------------------------------------------------
            RECEIVED FROM:
peer            192.168.30.113
status          C,I,R
    Attributes:
     attribute-type    installed
     encap-proto       0
     encap-spi         256
     encap-auth        sha1-hmac,ah-sha1-hmac
     encap-encrypt     aes256
     public-ip         192.168.110.215
     public-port       12347
     private-ip        192.168.110.215
     private-port      12347
     public-ip         ::
     public-port       0
     private-ip        ::
     private-port      0
     bfd-status        down
     site-id           214215
     preference        0
     weight            1
     version           3
    gen-id             0x80000026
     carrier           default
     restrict          0
     groups            [ 0 ]
     bandwidth         0
     qos-group         default-group

----------------------------------------------------
tloc entries for 192.168.30.215
                 biz-internet
                 ipsec
----------------------------------------------------
            RECEIVED FROM:
peer            192.168.30.113
status          C,I,R
    Attributes:
     attribute-type    installed
     encap-proto       0
```

```
     encap-spi          256
     encap-auth         sha1-hmac,ah-sha1-hmac
     encap-encrypt      aes256
     public-ip          192.168.109.215
     public-port        12347
     private-ip         192.168.109.215
     private-port       12347
     public-ip          ::
     public-port        0
     private-ip         ::
     private-port       0
     bfd-status         down
     site-id            214215
     preference         0
     weight             1
     version            3
    gen-id              0x80000026
     carrier            default
     restrict           0
     groups             [ 0 ]
     bandwidth          0
     qos-group          default-group
```

ce1#

You can check this command on the vSmart controller in order to understand which routes receive the particular prefix (see "ADVERTISED TO" section):

vsmart1#

```
show omp routes 10.1.1.0/24 detail | nomore | exclude not\ set

-------------------------------------------------
omp route entries for vpn 1 route 10.1.1.0/24
-------------------------------------------------
            RECEIVED FROM:
peer            192.168.30.216
path-id         68
label           1002
status          C,R
    Attributes:
     originator       192.168.30.216
     type             installed
     tloc             192.168.30.216, biz-internet, ipsec
     overlay-id        1
     site-id          216
     origin-proto     eBGP
     origin-metric    0
     as-path          65500
            ADVERTISED TO:
peer    192.168.30.214
    Attributes:
     originator       192.168.30.216
     label            1002
     path-id          5525
     tloc             192.168.30.216, biz-internet, ipsec
```

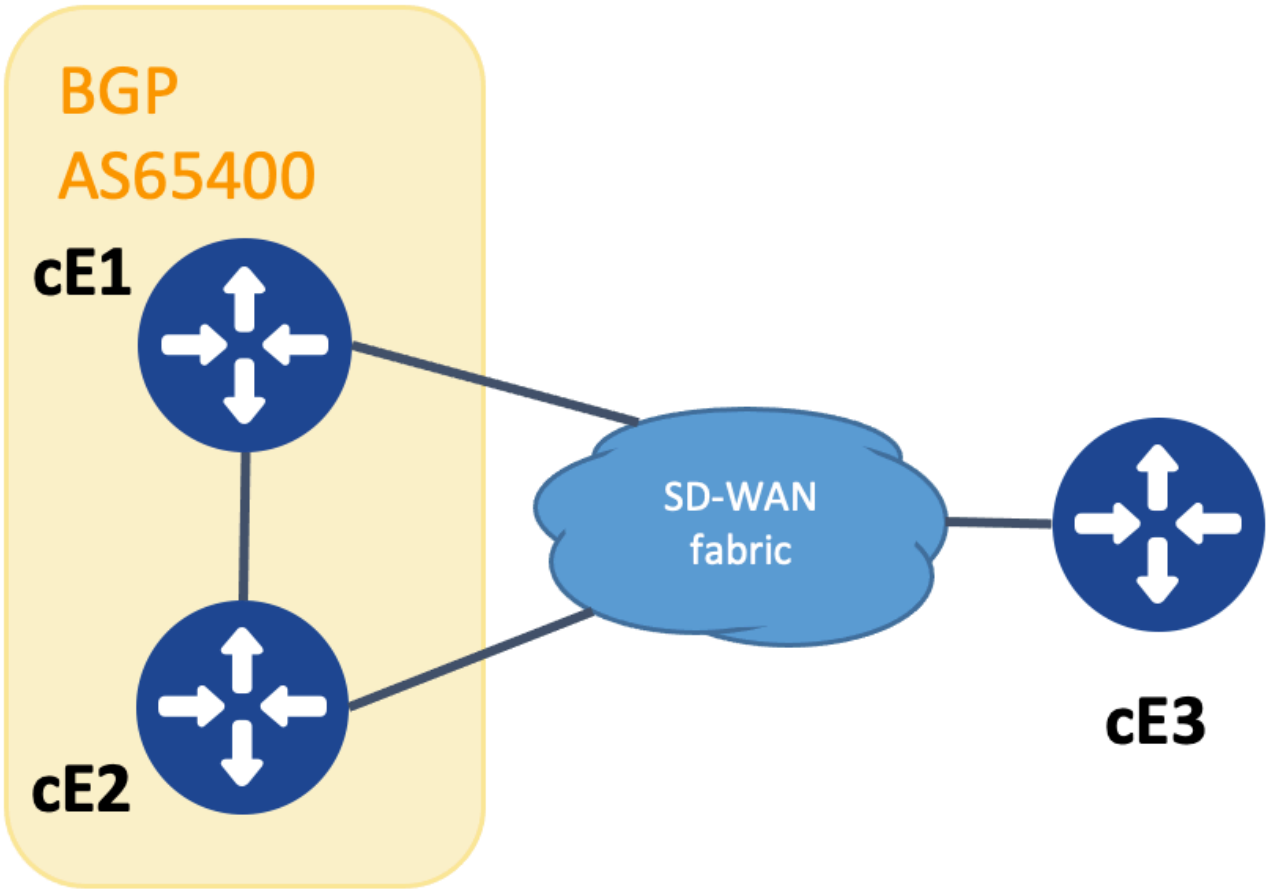
```
     site-id          216
     overlay-id        1
     origin-proto     eBGP
     origin-metric    0
     as-path          65500
            ADVERTISED TO:
peer    192.168.30.215
    Attributes:
     originator       192.168.30.216
     label            1002
     path-id          5287
     tloc             192.168.30.216, biz-internet, ipsec
     site-id          216
     overlay-id        1
     origin-proto     eBGP
     origin-metric    0
     as-path          65500
```

The**site-id** is also preserved as BGP site-of-origin (SoO) extended community attribute (you can notice **SoO:0:<site-id>** in the previous outputs). That is used to identify routes that have originated from a site so that the re-advertisement of that prefix back can be prevented. For this to function correctly, routers must send extended communities. Configure cE1 to send extended communities to router cE2:

```
router bgp 65401
 address-family ipv4 vrf 1
 neighbor 192.168.160.215 send-community both
```

SoO Loop Prevention Explanation

For the case where two routers at the same site are iBGP neighbors, SD-WAN has a built-in loop-prevention mechanism in order to prevent a routing loop from OMP to BGP and back from BGP to OMP. In order to demonstrate this, the topology was slightly updated and the same site-id 214215 was configured on both routers that run BGP AS65400 (cE1/cE2). In this example, a 10.1.1.0/24 prefix is advertised into OMP from the remote site (cE3) and learned in OMP at Site 214215 (cE1-cE2).



Topology for SoO demontration

In order to accomplish loop-prevention, the BGP extended community SoO is used to show which site originated the prefix. This community is added to a prefix when it is redistributed from OMP into BGP.

The **send-community <both|extended>** command must be configured on the neighbor statement in both devices as shown, in order for this functionality to work correctly.

cEdge1#

```
show run | sec router bgp

router bgp 65400
 bgp log-neighbor-changes
 !
 address-family ipv4 vrf 1
  redistribute omp
  neighbor 192.168.160.215 remote-as 65400
  neighbor 192.168.160.215 activate
  neighbor 192.168.160.215 send-community both
 exit-address-family
```

cEdge2#

```
show run | sec router bgp

router bgp 65400
 bgp log-neighbor-changes
 !
 address-family ipv4 vrf 1
  neighbor 192.168.160.214 remote-as 65400
  neighbor 192.168.160.214 activate
  neighbor 192.168.160.214 send-community both
 exit-address-family
```

The extended community can be seen with the output of **show bgp vpnv4 unicast vrf 1 <prefix>** from either the advertising or receiving site.

Example

cEdge1#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1

BGP routing table entry for 1:10:10.1.1.1/24, version 4
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    192.168.30.215 (via default) from 0.0.0.0 (192.168.109.215)
      Origin incomplete, metric 1000, localpref 50, valid, sourced, best
      Extended Community: SoO:0:214215 RT:1:1
      rx pathid: 0, tx pathid: 0x0
      Updated on Jul 5 2152 23:30:55 UTC
```

On the router which advertises the prefix from OMP into BGP (cEdge1 in this example), only the OMP route must be present in the RIB.

Example

cEdge1#

```
show ip route vrf 1 10.1.1.1

Routing Table: 1
Routing entry for 10.1.1.1/32
  Known via "omp", distance 251, metric 0, type omp
  Redistributing via bgp 65400
  Advertised by bgp 65400
  Last update from 192.168.30.215 on Sdwan-system-intf, 15:59:54 ago
  Routing Descriptor Blocks:
  * 192.168.30.215 (default), from 192.168.30.215, 15:59:54 ago, via Sdwan-system-intf
      Route metric is 0, traffic share count is 1
```

However, it can happen that a race condition occur on the second router which receives the advertised prefix and causes the BGP route to get installed in the RIB prior to the OMP route getting learned.

On cEdge2, the output of **sh bpg vpnv4 unicast vrf 1 <prefix>** shows these:

1. Not advertised to any peer.
2. Extended Community includes the site-id 214215 which is the same site this router is at.

Example

cEdge2#

```
show bgp vpnv4 unicast vrf 1 10.1.1.1

BGP routing table entry for 1:1:10.1.1.1/24, version 32
Paths: (1 available, best #1, table 1)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    192.168.160.214 (via vrf 1) from 192.168.160.214 (192.168.54.11)
      Origin incomplete, metric 1000, localpref 50, valid, internal, best
      Extended Community: SoO:0:214215 RT:65512:10
      rx pathid: 0, tx pathid: 0x0
      Updated on Jul 6 2152 17:26:19 UTC
```

On cEdge2, the output of **sh ip route vrf <vrf_number> <prefix>** shows these:

1. The flag “SDWAN Down” is seen to show that this has been detected to have originated from the same site.
2. The administrative distance of the route is 252 (higher than OMP and different than the expected iBGP AD 200).

Example

cEdge2#

```
show ip route vrf 1 10.1.1.1
Routing Table: 1
Routing entry for 10.1.1.0/24
Known via "bgp 65400",

distance 252
, metric 1000, type internal
Redistributing via omp
Last update from 192.168.160.214 00:15:13 ago
Routing Descriptor Blocks:
* 192.168.160.214, from 192.168.160.214, 00:15:13 ago
opaque_ptr 0x7F9DD0B86818

SDWAN Down
Route metric is 1000, traffic share count is 1
AS Hops 0
MPLS label: none
```

When a site router detects that a BGP learned route originates from the same site-id, the route is not advertised back into OMP.

Related Information

- **vEdge Does Not Advertise Its Own AS When BGP Routes Are Advertised Into OMP**
- **IP Routing: BGP Configuration Guide**
- **Routing Configuration Guide, Cisco Catalyst SD-WAN Releases 17.x**
- **Cisco SD-WAN Command Reference - overlay-as**
- **Cisco Technical Support & Downloads**