

vEdge Does Not Advertise Its Own AS When BGP Routes Are Advertised Into OMP

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Configure](#)

[Network Diagram](#)

[Verify](#)

[Troubleshoot](#)

[Conclusion](#)

[Related Information](#)

Introduction

This document describes why vEdge router does not advertise its own Autonomous System (AS) number when you advertise Border Gateway Protocol (BGP) routes into Overlay Management Protocol (OMP).

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Basic understanding of Viptela Software Defined Wide Area network (SDWAN) solution
- OMP

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created with help of 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.

Configure

Network Diagram

Here is a simple topology diagram that describes the setup:

192.168.41.0/24---|R1---vedge1---Overlay---vedge2---R2|---192.168.51.0/24



	R1	vedge1		vedge2	R2
A	Local	65156	AS-PATH	65156	65002 64500 65156 ?
B	65001 64500 65157 ?	65157		65157	Local

And the problem can be depicted here:

For prefix A on R2 - you expect to see AS-PATH: 65002, 64500, 65001, 65156 but only saw 65002 64500 65156

For prefix B on R1 - you expect to see AS-PATH: 65001, 64500, 65002, 65157 but only saw 65001 64500 65157

Verify

Use this section in order to confirm that your configuration works properly.

1. Take a look at the prefix B (192.168.51.0/24) on vedge2.

```
vedge2# show bgp routes 192.168.51.0/24 detail
bgp routes-table vpn 40 192.168.51.0/24
best-path 1
info 0
nexthop      192.168.50.115
metric       0
weight       0
origin       igp
as-path      65157
ri-peer      192.168.50.115
ri-routerid  2.2.2.1
path-status  valid,best,external
tag          0
```

As you notice, the AS path here consists of a single AS of the peer router R2 and this is expected behavior. You see this prefix on vedge2 the way it was advertised.

You can observe similar behavior as ordinary Cisco Internetworking Operating Systems (IOS®) routers as well.

2. Advertising BGP routes via OMP.

```
omp
 overlay-as      64500
 advertise bgp
!
```

Advertise bgp configuration is just like a well-known Cisco IOS® **redistribute** command. OMP is similar to BGP, but in reality, it is another protocol developed specifically for the SDWAN solution. So, you preserve BGP prefix attributes that you had in BGP table (refer to step 1).

3. Let's take a look at prefix B (192.168.51.0/24) on vedge1 after you have passed through the overlay.

```
vedge1# show omp routes 192.168.51.0/24 detail
```

```
-----
omp route entries for vpn 40 route 192.168.51.0/24
-----
```

```
                RECEIVED FROM:
peer            192.168.30.103
path-id        12
label          1003
status         C,I,R
loss-reason    not set
lost-to-peer   not set
lost-to-path-id not set
  Attributes:
    originator  192.168.30.105
    type        installed
    tloc        192.168.30.105, mpls, ipsec
    ultimate-tloc not set
    domain-id   not set
    overlay-id  1
    site-id     50
    preference  not set
    tag         not set
    origin-proto eBGP
    origin-metric 0
    as-path     "65157"
    unknown-attr-len not set
```

As you notice, attributes of this prefix were preserved similar to how it would be done for any other dynamic routing protocol. AS-path of BGP is just another one attribute of foreign routing protocol that has no meaning for OMP. Only one AS is seen in the path. It is expected behavior because vedge2 didn't do any advertisement to another BGP neighbor. It did redistribution to OMP from BGP, so it must not prepend its own AS to this prefix.

4. Perform redistribution of OMP into BGP AS 65001, but now on vedge1.

```
omp
no shutdown
overlay-as 64500 advertise bgp ! ... ! vpn 40 router bgp 65001 propagate-aspath address-family
ipv4-unicast redistribute omp ! neighbor 192.168.40.114 no shutdown remote-as 65156 ! ! !
```

Here is the place where overlay-AS actually steps in and plays it's role, but unlike on ordinary Cisco IOS® you won't see changes to prefix as it is advertised to R1, because, on vEdge there is no command similar to Cisco IOS® **show ipv4 unicast bgp neighbors 192.168.40.114 advertised-routes** that can help see newly constructed AS-path.

5. So, you are only left with the possibility to check the AS-path on the receiving router (in this case R1).

```
R1#show bgp vpnv4 unicast vrf vEdge1_18.3 neighbors 192.168.40.104 routes
BGP table version is 11, local router ID is 192.168.41.10
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 1:183 (default for vrf vEdge1_18.3)
*>  192.168.51.0       192.168.40.104      1000              0 65001 64500 65157 ?
```

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Conclusion

OMP shares some similarities with BGP, but this fact should not make any confusion when two protocols interact with each other. AS-path is one of the topics that commonly causes such misunderstanding.

Related Information

- [Configure OMP](#)
- [Technical Support & Documentation - Cisco Systems](#)