

AOS-CX 10.6 Update
November, 2020

aruba
a Hewlett Packard
Enterprise company

BGP redistribute local loopback

Aruba Switching TME



Redistribute local loopback

Agenda



- 1 Overview
- 2 Use Cases
- 3 Details / Caveats
- 4 Configuration
- 5 Best Practices
- 6 Troubleshooting
- 7 Demo

Overview

aruba

a Hewlett Packard
Enterprise company

Overview

Redistribution of loopback IP address

- Before 10.6, redistribution of loopback /32 IP address or /128 IPv6 address was not well supported.

A workaround was required:

- Workaround#1: consists in sourcing the loopback /32 IP address in BGP with a network statement command:
- Workaround#2: consists in redistributing a static route with destination route being the IP address of the loopback pointing to the loopback interface itself, and then redistributing the static route into BGP.
- Workaround#3: consists in using a non /32 IP address (like /31) instead of a /32 traditional IP address for loopback so as to create a non /32 subnet entry that would be redistributed through connected redistribution into BGP.

```
8325-1(config-bgp-vrf-ipv4-uc)# network  
A.B.C.D/M Configure the IP network to import into BGP
```

- **Why redistribute connected could not redistribute loopback ?**

Reason: any /32 IP address is seen as a “**local**” route hosted by the switch itself and not as a “**connected**” route. Consequently, **redistribute connected** will not redistribute any loopback /32 IP address.

- **Since 10.6**, on all CX platforms, a new BGP redistribution command is introduced to fix this gap:

```
router bgp 65003  
  address-family ipv4 unicast  
    redistribute connected  
    redistribute local loopback
```

```
8325-1# show ip route  
A.B.C.D Display longest prefix match  
A.B.C.D/M Display exact route match  
all-vrfs all vrf information  
bgp Show bgp routes only  
connected Show connected routes only  
local Show local routes only  
ospf Show ospf routes only  
rip Show rip routes only  
static Show static routes only  
summary Display the aggregate count of routes per routing protocol  
vrf Specify the VRF name  
vsx-peer Displays VSX peer switch information  
<cr>  
8325-1# show ip route local  
  
Displaying ipv4 routes selected for forwarding  
'[x/y]' denotes [distance/metric]  
  
192.168.1.3/32, vrf default  
  via loopback0, [0/0], local  
192.168.3.1/32, vrf default  
  via 1/1/23, [0/0], local  
192.168.3.9/32, vrf default  
  via 1/1/24, [0/0], local  
192.168.3.200/32, vrf default  
  via vlan2, [0/0], local  
192.168.11.3/32, vrf default  
  via loopback1, [0/0], local
```



Use Cases

aruba

a Hewlett Packard
Enterprise company

Use cases

- This new redistribution command is restricted to BGP.
- No OSPF support for this command yet (for future release).
Currently, injecting Loopback IP address in OSPF is addressed with:

```
interface loopback 0
ip address 192.168.1.3/32
ip ospf 1 area 0.0.0.0
```
- Redistribution of Loopback IP address is important for:
 - In-band management for SSH/HTTPS/SNMP access to Loopback IP address
 - Traffic initiated from the switch (like NTP, syslog, RADIUS/TACACS...) that is predefined in Firewall rules with source IP address range being the allocated subnet for Loopback (usual best practice).
 - BGP EVPN fabric: loopback anycast IP address must be in BGP routing table for underlay VTEP reachability.

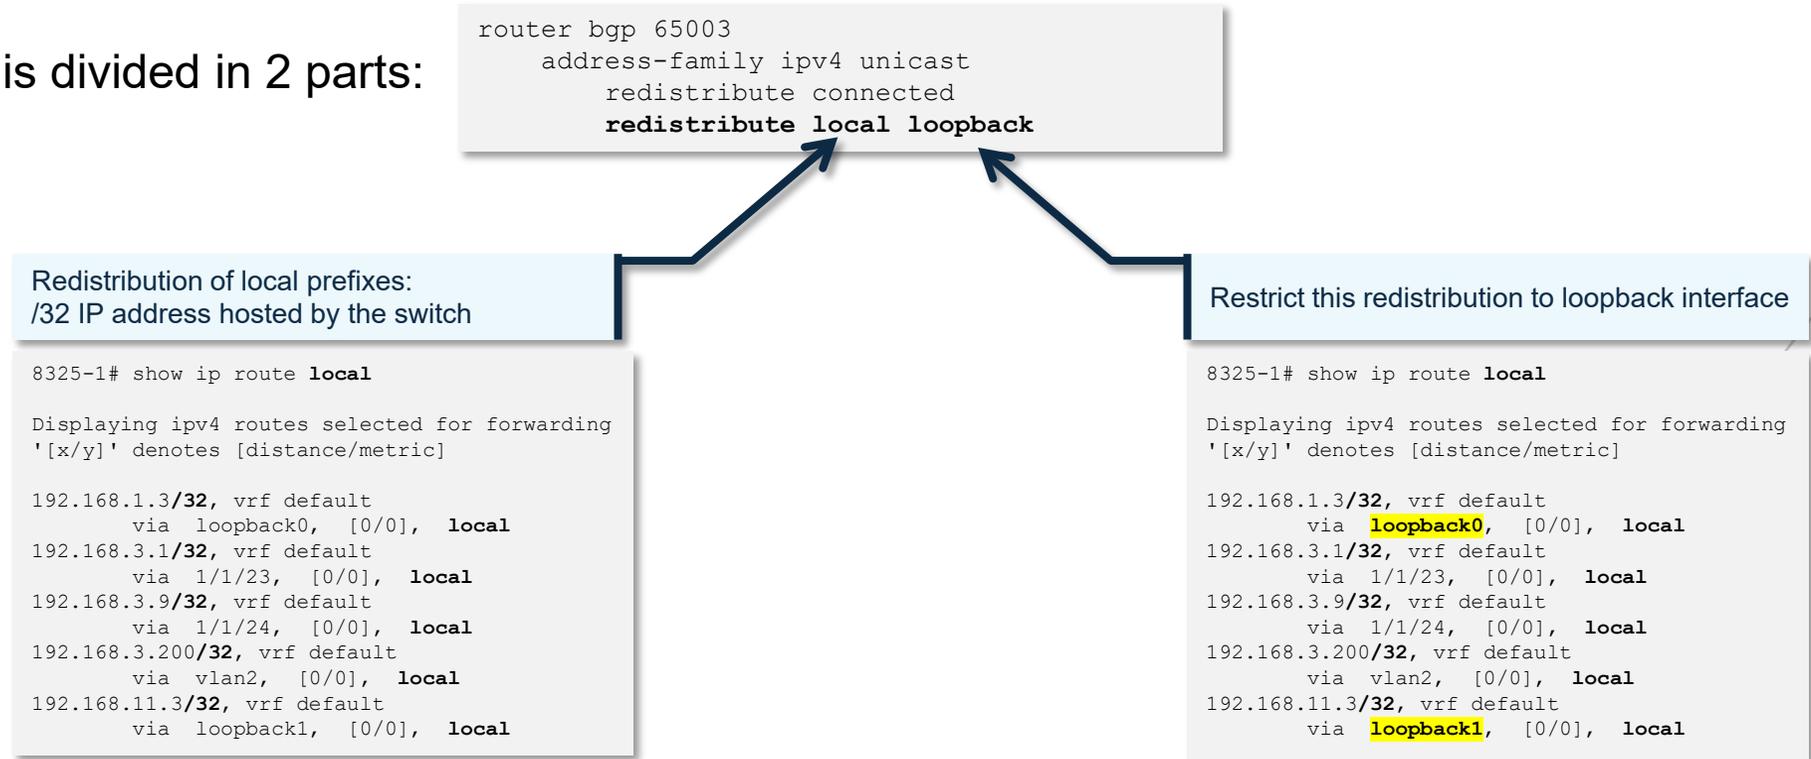
Details

aruba

a Hewlett Packard
Enterprise company

Redistribution Details

- The redistribution is divided in 2 parts:



- No other “**local**” prefixes besides those corresponding to loopback interfaces can be redistributed.
- The “**redistribute connected**” and “**redistribute local loopback**” commands are independent. “**Redistribute connected**” is still required to redistribute connected subnet, and is not a prerequisite for “**redistribute local loopback**” command.

Configuration

aruba

a Hewlett Packard
Enterprise company

Redistribution local configuration

- A specific BGP redistribution command is available since 10.6:

```
router bgp 65003
  address-family ipv4 unicast
  redistribute connected
  redistribute local loopback
```

- Only local /32 IPv4 prefixes or /128 IPv6 prefixes related to loopback interfaces can be redistributed:

```
switch(config-bgp-vrf-ipv6-uc)# redistribute local
  loopback Redistribute Local routes on loopback interfaces
```

- The redistribution can be controlled with an optional route-map:

```
switch(config-bgp-vrf-ipv6-uc)# redistribute local loopback
  route-map Apply route-map policy for redistribution
  <cr>
```

- The redistribution is configured:
 - per Address Family: IPv4 and/or IPv6
 - per VRF:

```
router bgp 65003
  address-family ipv4 unicast
    redistribute connected
    redistribute local loopback
  exit-address-family
  address-family ipv6 unicast
    redistribute connected
    redistribute local loopback
  exit-address-family
  vrf VRF1
    address-family ipv4 unicast
      redistribute connected
      redistribute local loopback
    exit-address-family
  address-family ipv6 unicast
    redistribute connected
    redistribute local loopback
  exit-address-family
```

Default VRF

VRF1 VRF



aruba

a Hewlett Packard
Enterprise company

Best Practices

Redistribution Best Practices

- Best practices are:
 1. redistribute both connected and loopback.
 2. Use route-map with prefix-list to strictly control the redistribution (to avoid human errors).
 3. Same route-map can be used for connected and local loopback to simplify administration.

Best Practice Example

```
interface loopback 0
  ip address 192.168.1.3/32
  ipv6 address fd00:192:168:1::3/128
  ip ospf 1 area 0.0.0.0
  ipv6 ospfv3 1 area 0.0.0.0
!
router ospf 1
  router-id 192.168.1.3
  passive-interface default
  area 0.0.0.0
!
router ospfv3 1
  router-id 192.168.1.3
  passive-interface default
  area 0.0.0.0
!
ip prefix-list endpoint seq 5 permit 192.168.1.0/24 ge 32
ip prefix-list endpoint seq 10 permit 10.1.0.0/16 le 32
ipv6 prefix-list v6-endpoint seq 5 permit fd00:192:168:1::/64 ge 128
ipv6 prefix-list v6-endpoint seq 10 permit fd00:10:1::/64 le 128
!
route-map connected-bgp permit seq 10
  match ip address prefix-list endpoint
route-map connected-bgp permit seq 20
  match ipv6 address prefix-list v6-endpoint
!
router bgp 65001
  bgp router-id 192.168.1.3
  <.. omitted for focus on redistribution topic ..>
  address-family ipv4 unicast
    redistribute connected route-map connected-bgp
    redistribute local loopback route-map connected-bgp
  exit-address-family
  address-family ipv6 unicast
    redistribute connected route-map connected-bgp
    redistribute local loopback route-map connected-bgp
  exit-address-family
```

Loopback IP address is injected in OSPF through OSPF command in interface context

Loopback in OSPFv2

Loopback in OSPFv3

Prefix-list and route-map for redistribution control

redistribution



aruba

a Hewlett Packard
Enterprise company

Troubleshooting

Redistribution Troubleshooting

IPv4 - On originating router: show bgp

show ip route

```
8325-1# show ip route vrf VRF1
Displaying ipv4 routes selected for forwarding
'[x/y]' denotes [distance/metric]
0.0.0.0/0, vrf VRF1
  via 192.168.4.0[vrf SERVICES], [200/0], bgp
10.1.10.0/24, vrf VRF1
  via vlan10, [0/0], connected
10.1.10.2/32, vrf VRF1
  via vlan10, [0/0], local
<..omitted ..>
192.168.101.1/32, vrf VRF1
  via loopback10, [0/0], local
```

```
8325-3# show ip route 192.168.101.1/32 vrf VRF1
Displaying ipv4 routes selected for forwarding
'[x/y]' denotes [distance/metric]
192.168.101.1/32, vrf VRF1, tag 0
  via 192.168.11.3[vrf default], [200/0], bgp, encap vxlan, l3vni 100001
```

show bgp for IPv4 AF

```
8325-1# show bgp vrf VRF1 ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : VRF1
Local Router-ID 192.168.1.3

   Network          Nexthop          Metric      LocPrf      Weight Path
   Route Distinguisher: 192.168.1.3:1
*> 0.0.0.0/0         192.168.4.0      0           100         0         65100 i
*> 10.1.10.0/24     0.0.0.0          0           100         0         ?
<.. Omitted ..>
*> 192.168.101.1/32 0.0.0.0          0           100         0         ?
Total number of entries 8
```

/32 loopback IP address, locally originated

Redistribution Troubleshooting

IPv6 - On originating router: show bgp

show ipv6 route

```
8360# show ipv6 route vrf VRF1
Displaying ipv6 routes selected for forwarding
'[x/y]' denotes [distance/metric]
fd00:10:1:10::/64, vrf VRF1
  via vlan10, [0/0], connected
fd00:10:1:10::2/128, vrf VRF1
  via vlan10, [0/0], local
<..omitted ..>
fd00:192:168:101::1/128, vrf VRF1
  via loopback10, [0/0], local
```

show bgp for IPv6 AF

```
8325-1# show bgp vrf VRF1 ipv6 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : VRF1
Local Router-ID 192.168.1.3

   Network          Nexthop          Metric    LocPrf    Weight Path
   Route Distinguisher: 192.168.1.3:1
*> fd00:10:1:10::/64  ::                0         100      0        ?
<..omitted ..>
*> fd00:192:168:101::1/128
   ::                0         100      0        ?
Total number of entries 3
```

/128 loopback IP address, locally originated

Demo



The Aruba logo consists of the word "aruba" in a lowercase, bold, orange sans-serif font. The letters are closely spaced, and the 'a' and 'u' have a distinctive shape with a small gap at the top.

a Hewlett Packard
Enterprise company

Thank you

vincent.giles@hpe.com