

AOS-CX 10.6 Update
November, 2020



Outbound Route Filtering

Aruba Switching TME



Outbound Route Filtering (ORF)

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

Definitions

Acronyms

- MP-BGP **M**ulti-**P**rotocol **B**order **G**ateway **P**rotocol
- AF **A**ddress **F**amily (Ex: IPv4, IPv6 or EVPN address families used in MP-BGP)
- EVPN **E**thernet **V**irtual **P**rivate **N**etwork
- L2VPN **L**ayer2 **V**irtual **P**rivate **N**etwork
- MB-BGP EVPN Refers to the EVPN address family in MP-BGP
- ORF **O**utbound **R**oute **F**iltering
- CE **C**ustomer **E**dge router
- PE **P**rovider **E**dge router

Overview

Route Filtering

- **Route filtering** is the mechanism that removes specific routes from being added in the local routing database, or not advertised to neighbors.
- Reasons for route filtering:
 - **Economic**: large BGP tables consumes memory and require CPU cycles to compute
 - **Routing engineering**: some paths due to available bandwidth might not be used
 - **Security**: control exit point; anti-spoofing.
- Filter direction: in or/and out

```
switch(config-bgp-ipv4-uc)# neighbor 192.168.25.1 route-map test
in    Inbound direction of the route-map
out   Outbound direction of the route-map
```

- **Problem statement:**
The PE is advertising 1000 routes to the CE and the CE requires only 10-20 of them.
The CE has to go through all the prefixes and deny them individually which increases the cpu/mem load on the CE.
- In order to reduce the computation, it is better for the **PE to send only the necessary prefixes** to the CE.

Overview

Route Filtering Computation



Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	0	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.2.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.3.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.4.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.5.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.6.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.7.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.8.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.9.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.11.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.12.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.13.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.14.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.15.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.16.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.17.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.18.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.19.0.0/16	192.168.3.0	0	0	0	65001 ?
Total number of entries 20					

```
CE# show bgp ipv4 unicast neighbors 192.168.3.0 routes
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 : default
Local Router-ID 192.168.1.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	100	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	100	0	65001 ?
Total number of entries 3					

```
ip prefix-list prefix-listA seq 10 permit 0.0.0.0/0
ip prefix-list prefix-listA seq 20 permit 10.1.0.0/16
ip prefix-list prefix-listA seq 30 permit 10.10.0.0/16
!
route-map filtering1 permit seq 10
    match ip address prefix-list prefix-listA
!
router bgp 65002
    neighbor 192.168.3.0 remote-as 65001
    address-family ipv4 unicast
        neighbor 192.168.3.0 activate
        neighbor 192.168.3.0 route-map filtering1 in
    exit-address-family
```

```
CE# show bgp ipv4 unicast neighbors 192.168.3.0 received-routes
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 : default
Local Router-ID 192.168.1.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	0	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.2.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.3.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.4.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.5.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.6.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.7.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.8.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.9.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.11.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.12.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.13.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.14.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.15.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.16.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.17.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.18.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.19.0.0/16	192.168.3.0	0	0	0	65001 ?
Total number of entries 20					

All prefixes are received and processed !



Only 3 routes are selected in BGP RIB

Reminder

How to see all the received-routes ?



Reminder

Soft-reconfiguration in



```
router bgp 65002
  bgp router-id 192.168.1.2
  neighbor 192.168.3.0 remote-as 65001
  address-family ipv4 unicast
    neighbor 192.168.3.0 activate
    neighbor 192.168.3.0 route-map filtering1 in
    neighbor 192.168.3.0 soft-reconfiguration inbound
```

```
CE# show bgp ipv4 unicast neighbors 192.168.3.0 received-routes
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 : default
Local Router-ID 192.168.1.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	0	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.2.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.3.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.4.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.5.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.6.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.7.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.8.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.9.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.11.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.12.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.13.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.14.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.15.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.16.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.17.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.18.0.0/16	192.168.3.0	0	0	0	65001 ?
e 10.19.0.0/16	192.168.3.0	0	0	0	65001 ?

Total number of entries 20

```
CE# show bgp ipv4 unicast neighbors 192.168.3.0
Codes: ^ inherited from peer-group
```

```
VRF : default
```

```
BGP Neighbor 192.168.3.0 (External)
```

Description	Peer-group	Remote Router Id	Remote AS	Remote Port	State	Conn. Established	Passive	Cfg. Hold Time	Neg. Hold Time	Up/Down Time	Local-AS Prepend	BFD	Password	Last Err Sent	Last SubErr Sent	Last Err Rcvd	Last SubErr Rcvd	Graceful-Restart	Gr. Stalepath Time	TTL	Weight	Confederation-Peers	Local Router Id	Local AS	Local Port	Admin Status	Conn. Dropped	Update-Source	Cfg. Keep Alive	Neg. Keep Alive	Alt. Local-AS
		192.168.1.1	65001	179	Established	1	No	180	180	00h:01m:59s	No	Disabled		No Error	No Error	No Error	No Error	Enabled	300	1	0	No	192.168.1.2	65002	55498	Up	0		60	60	0

Message statistics	Sent	Rcvd
Open	1	1
Notification	0	0
Updates	2	2
Keepalives	3	3
Route Refresh	0	0
Total	6	6

Capability	Advertised	Received
Route Refresh	Yes	Yes
Graceful Restart	Yes	Yes
Add-Path	No	No
Four Octet ASN	Yes	Yes
Address family IPv4 Unicast	Yes	Yes
Address family IPv6 Unicast	No	No
Address family L2VPN EVPN	No	No

```
Address Family : IPv4 Unicast
```

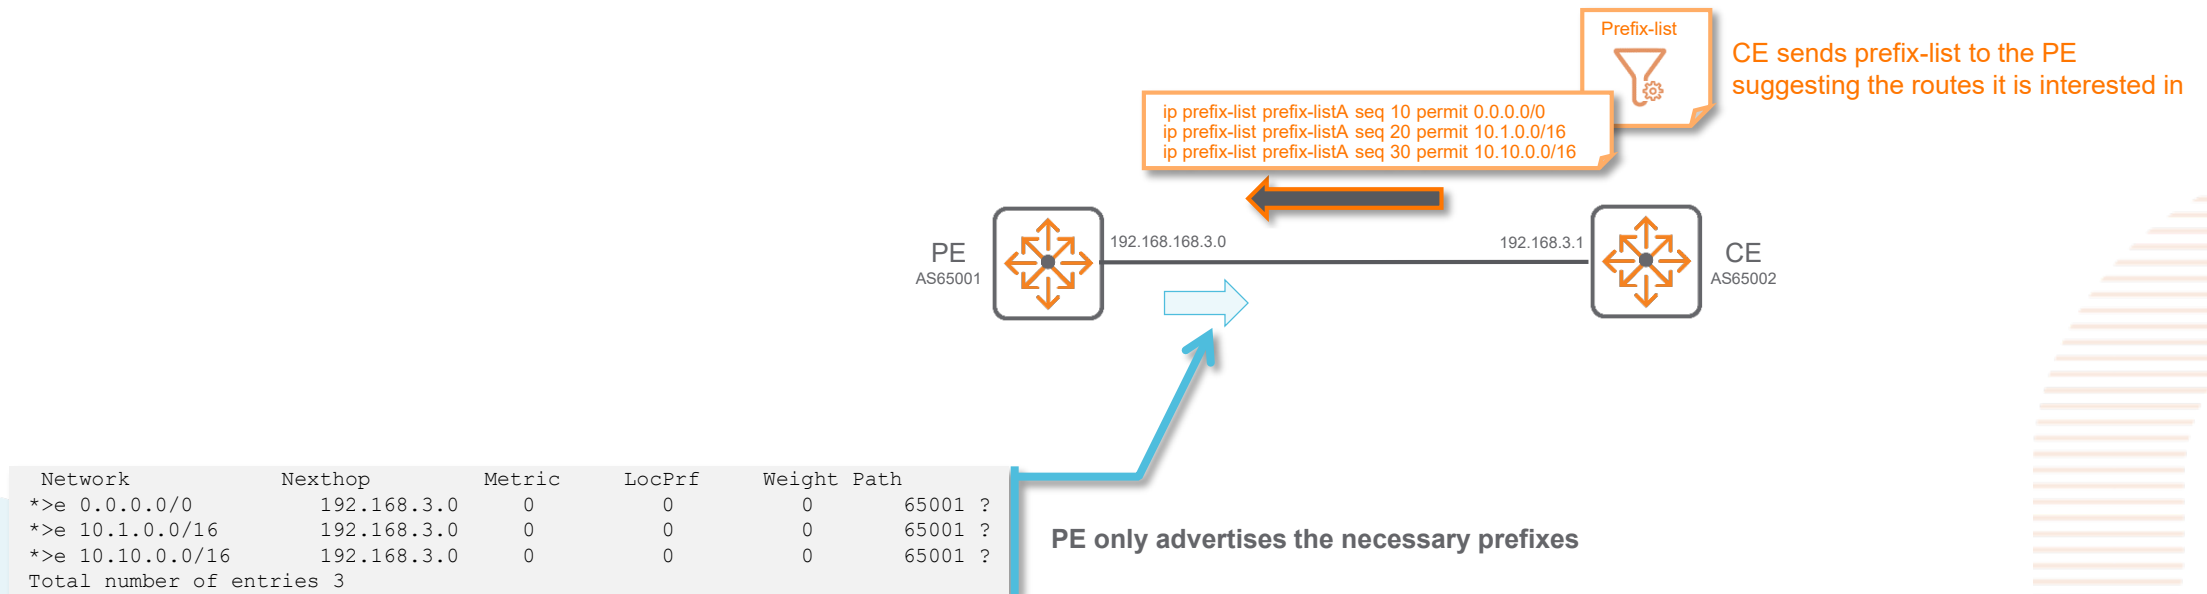
Rt. Reflect. Client	Allow-AS in	Max. Prefix	Nexthop-Self	Cfg. Add-Path	Neg. Add-Path	Send Community	Advt. Interval	Soft Reconfig In	Default-Originate
No	0	64000			Disable		30	Yes	

```
Routemap In : filtering1
Routemap Out :
ORF type :
ORF capability :
```


Overview

Outbound Route Filtering (ORF)

- The **Outbound Route Filtering** capability provides a mechanism for a BGP speaker to **send to its BGP peer** a set of Outbound Route Filters (ORFs) that can be **used by its peer to filter its outbound routing updates** to the requesting neighbor.
- This is a filtering method used to **reduce the computation** on the router receiving the route updates.
- Inbound route-map is technically no longer required to deny these prefixes based on prefix-list filter.




Use Cases

aruba

a Hewlett Packard
Enterprise company

Use cases



- The ORF capability is a “Cooperative” route filtering capability and must be supported on both peers:
 - Local BGP router must support ORF capability (send)
 - BGP neighbor must support ORF capability (receive)
- Very useful to control routes sent from a router belonging to an **other administrative domain**. All the filtering customization is not done on the PE, but on the CE, bringing lot of flexibility to control routing.
- Main use-cases:
 - for **large routing table**, to minimize massive in-bound route filtering.
 - for **operational enforcement**: received routes are filtered at two places: in-bound filter (traditional) + out-bound from the sender. This provides extra routing control in cases of human errors (mistake should happen on both enforcement points).
- Usually, the prefix-list defined for in-bound route-map includes the entries of the prefix-list used for ORF. Content of both prefix-lists could be the same. (see Best Practices).
- Although not restricted to eBGP, there is little interest to use ORF in the same Autonomous System (same administrative domain).

Details / Caveats

Platform Support

10.6 – ORF support

ORF support	6200	6300	6400	8360	8320	8325	8400
MP-BGP IPv4 AF	N/A	Yes	Yes	Yes	Yes	Yes	Yes
MP-BGP IPv6 AF		Yes	Yes	Yes	Yes	Yes	Yes
MP-BGP EVPN AF		No	No	No	No	No	No

ORF Details / Caveats

- Prefix-list size:

- The maximum number of prefix-list entries in the ORF prefix-list being sent is 128.
- The maximum number of prefix-list entries in the ORF prefix-list being received is 128.

WARNING: if a **third party ORF sender** requests to **AOS-CX ORF receiver** to filter more than 128 entries in the ORF prefix-list, it might work but it is not tested and consequently not supported.

- Applies to IPv4 and IPv6 address families. **It does not apply to the EVPN AF.**

In AOS-CX, route-map mechanism is used internally to process ORF filter on the sender and on the receiver. As route-map is not yet supported for EVPN AF, ORF is not yet supported in EVPN AF.

- **Warning: inbound soft reconfiguration** prevents ORF from the PE. *Verification pending.*

ORF filtering still happens but inbound on the CE (just like regular inbound route-map).

With inbound soft reconfiguration, the entire neighbor RIB is stored in the BGP neighbor RIB cache, regardless of the current inbound policy.



Important Caveat

How to refresh routes after ORF filter modification

- An ORF prefix-list change will **NOT** trigger a **new ROUTE REFRESH message**.
- BGP session must be “soft reset”.
- BGP supports two methods of clearing a BGP session:
 - A hard-reset, which tears down the BGP session, removes BGP routes from the peer, and is disruptive.
 - A soft-reset, which invalidates the BGP cache and requests a full advertisement from its BGP peer, and is not disruptive
- Route-refresh capability is **always enabled in AOS-CX** (as most of the industry vendors).
- If route-refresh capability is supported on the **ORF receiver**, a soft-reset of the BGP session on the **ORF sender** will trigger a **new ROUTE REFRESH message** and the **ORF receiver** will send route updates according to the new received ORF prefix-list.
- If route-refresh capability is not supported on the **ORF receiver**, the BGP session must be “hard reset”. This will be disruptive.

ORF Caveats

- Same ORF prefix-list can be used for multiple neighbors.
- A unique ORF prefix-list can be used per neighbor.
- If route-map inbound is also applied on multiple neighbors along with ORF, then the route-map name must be different for each neighbor.
- If route-map inbound is also applied on a BGP neighbor on an IPv6 AF along with ORF, then the route-map sequence number 1 can not be used (configuration might take it without error).
- ORF interoperability with Cisco is under testing.

ORF Details

BGP OPEN message from PE

- PE should be capable of “**receiving**” the route-filter.
- If ORF is enabled on both IPv4 and IPv6 AFs, then capability in BGP OPEN message is **not decoded properly** by **Wireshark**. Decoding must be done manually based on RFC specification.



ORF Details

BGP OPEN message from CE

- CE should be capable of **“sending”** the route-filter.
- If ORF is enabled on both IPv4 and IPv6 AFs, then capability in BGP OPEN message is **not decoded properly** by **Wireshark**. Decoding must be done manually based on RFC specification.




ORF Details

BGP ROUTE-REFRESH message

- CE sends the route-filter.

0000	08 00 09 3b 20 61 08 00	09 8b 0c b3 08 00 45 c0	...; a... ..E..
0010	00 b1 aa 29 40 00 01 06	47 0c c0 a8 03 01 c0 a8	...)@... G... ..
0020	03 00 00 b3 c0 be 90 da	fb 7b a1 79 20 83 80 18{.y... ..
0030	00 e3 bb b5 00 00 01 01	08 0a f5 af a8 4f f8 6a0.j... ..
0040	69 c8 ff ff ff ff ff ff	ff ff ff ff ff ff ff ff	i... ..
0050	ff ff 00 37 05 00 01 00	01 01 40 00 1c 00 00 00	...7.... .@....
0060	00 1e 00 00 10 0a 0a 00	00 00 00 14 00 00 10 0a
0070	01 00 00 00 00 0a 00 00	00 ff ff ff ff ff ff ff
0080	ff ff ff ff ff ff ff ff	ff 00 2f 02 00 00 00 14/.
0090	40 01 01 02 40 02 06 02	01 00 00 fd ea 40 03 04	@...@... ..@..
00a0	c0 a8 03 01 18 ac 10 01	ff ff ff ff ff ff ff ff
00b0	ff ff ff ff ff ff ff ff	00 17 02 00 00 00 00 00

- The BGP REFRESH message is **not decoded properly** by **Wireshark**. Decoding must be done manually based on RFC specification.



No.	Time	Source	Destination	Protocol	Length	Info
16	8.907730	192.168.3.0	192.168.3.1	TCP	66	49342 → 179 [ACK] Seq=86 Ack=86 Win=29200 Len=0 TSval=
17	8.909702	192.168.3.1	192.168.3.0	BGP	191	ROUTE-REFRESH Message, UPDATE Message, UPDATE Message
18	8.912454	192.168.3.0	192.168.3.1	TCP	66	49342 → 179 [ACK] Seq=86 Ack=211 Win=29200 Len=0 TSval=

> Frame 17: 191 bytes on wire (1528 bits), 191 bytes captured (1528 bits) on interface -, id 0
> Ethernet II, Src: HewlettP_8b:0c:b3 (08:00:09:8b:0c:b3), Dst: HewlettP_3b:20:61 (08:00:09:3b:20:61)
> Internet Protocol Version 4, Src: 192.168.3.1, Dst: 192.168.3.0
> Transmission Control Protocol, Src Port: 179, Dst Port: 49342, Seq: 86, Ack: 86, Len: 125
> Border Gateway Protocol - ROUTE-REFRESH Message
Marker: ffffffffffffffffffffffffffffffffff
Length: 55
Type: ROUTE-REFRESH Message (5)
Address family identifier (AFI): IPv4 (1)
Subtype: Normal route refresh request [RFC2918] with/without ORF [RFC5291] (0)
Subsequent address family identifier (SAFI): Unicast (1)
ORF information
ORF flag: Immediate (1)
ORF type: Unknown (64)
[Expert Info (Error/Unknown (2097152)): ORFEntry-Unknown (type 64)]
[ORFEntry-Unknown (type 64)]
[Severity level: Error]
[Group: Unknown]
ORF length: 28
> Border Gateway Protocol - UPDATE Message
Marker: ffffffffffffffffffffffffffffffffff
Length: 47
Type: UPDATE Message (2)
Withdrawn Routes Length: 0
Total Path Attribute Length: 20
Path attributes
> Path Attribute - ORIGIN: INCOMPLETE
> Path Attribute - AS_PATH: 65002
> Path Attribute - NEXT_HOP: 192.168.3.1
> Network Layer Reachability Information (NLRI)
> 172.16.1.0/24
> Border Gateway Protocol - UPDATE Message
Marker: ffffffffffffffffffffffffffffffffff
Length: 23
Type: UPDATE Message (2)
Withdrawn Routes Length: 0
Total Path Attribute Length: 0

ORF Details

Tcpdump on PE

PE
AS65001



192.168.168.3.0

192.168.3.1



CE
AS65002

capability

update

```
tcpdump: listening on any, link-type LINUX_SLL (Linux cooked), capture size 262144 bytes
1 18:55:26.778005 IP (tos 0xc0, ttl 1, id 4673, offset 0, flags [DF], proto TCP (6), length 52)
gateway.44978 > CE.bgp: Flags [F.], cksum 0xfeff (correct), seq 274038905, ack 1866370842, win 30016, options [nop,nop,TS val 4154842770 ecr 4109004921], length 0
2 18:55:26.807707 IP (tos 0xc0, ttl 1, id 11089, offset 0, flags [DF], proto TCP (6), length 60)
gateway.49124 > CE.bgp: Flags [S], cksum 0x51ce (correct), seq 4115361970, win 29200, options [mss 1460,sackOK,TS val 4154842778 ecr 0,nop,wscale 0], length 0
3 18:55:26.807925 IP (tos 0xc0, ttl 1, id 4674, offset 0, flags [DF], proto TCP (6), length 52)
gateway.44978 > CE.bgp: Flags [F.], cksum 0x523d (correct), seq 1, ack 2, win 30016, options [nop,nop,TS val 4154842793 ecr 4109049123], length 0
4 18:55:26.824179 IP (tos 0xc0, ttl 1, id 11090, offset 0, flags [DF], proto TCP (6), length 52)
gateway.49124 > CE.bgp: Flags [F.], cksum 0xd31f (correct), seq 4115361971, ack 2354712259, win 29200, options [nop,nop,TS val 4154842805 ecr 4109049151], length 0
5 18:55:26.829442 IP (tos 0xc0, ttl 1, id 11091, offset 0, flags [DF], proto TCP (6), length 118)
gateway.49124 > CE.bgp: Flags [P.], cksum 0xbb3f (correct), seq 0:66, ack 1, win 29200, options [nop,nop,TS val 4154842813 ecr 4109049151], length 66: BGP
Open Message (1), length: 66
  Version 4, my AS 65001, Holdtime 180s, ID 192.168.1.1
  Optional parameters, length: 37
  Option Capabilities Advertisement (2), length: 35
    Multiprotocol Extensions (1), length: 4
      AFI IPv4 (1), SAFI Unicast (1)
      0x0000: 0001 0001
    Route Refresh (2), length: 0
    Route Refresh (Cisco) (128), length: 0
    Enhanced Route Refresh (70), length: 0
      no decoder for Capability 70
    Graceful Restart (64), length: 6
      Restart Flags: [none], Restart Time 120s
      AFI IPv4 (1), SAFI Unicast (1), Forwarding state preserved: no
      0x0000: 0078 0001 0100
    Cooperative Route Filtering (3), length: 7
      no decoder for Capability 3
      0x0000: 0001 0001 0140 01
      4 Byte AS 65001
      0x0000: 0000 fde9
6 18:55:26.836603 IP (tos 0xc0, ttl 1, id 11092, offset 0, flags [DF], proto TCP (6), length 52)
gateway.49124 > CE.bgp: Flags [F.], cksum 0xd268 (correct), seq 66, ack 67, win 29200, options [nop,nop,TS val 4154842835 ecr 4109049172], length 0
7 18:55:26.843624 IP (tos 0xc0, ttl 1, id 11093, offset 0, flags [DF], proto TCP (6), length 71)
gateway.49124 > CE.bgp: Flags [P.], cksum 0xce37 (correct), seq 66:85, ack 67, win 29200, options [nop,nop,TS val 4154842838 ecr 4109049172], length 19: BGP
  Keepalive Message (4), length: 19
8 18:55:26.886356 IP (tos 0xc0, ttl 1, id 11094, offset 0, flags [DF], proto TCP (6), length 52)
gateway.49124 > CE.bgp: Flags [F.], cksum 0xd209 (correct), seq 85, ack 86, win 29200, options [nop,nop,TS val 4154842880 ecr 4109049184], length 0
9 18:55:26.889724 IP (tos 0xc0, ttl 1, id 11095, offset 0, flags [DF], proto TCP (6), length 52)
gateway.49124 > CE.bgp: Flags [F.], cksum 0xce21 (correct), seq 85, ack 221, win 30016, options [nop,nop,TS val 4154842884 ecr 4109049229], length 0
10 18:55:26.896003 IP (tos 0xc0, ttl 1, id 11096, offset 0, flags [DF], proto TCP (6), length 103)
gateway.49124 > CE.bgp: Flags [P.], cksum 0x8f3f (correct), seq 85:138, ack 221, win 30016, options [nop,nop,TS val 4154842893 ecr 4109049229], length 53: BGP
  Update Message (2), length: 53
    Origin (1), length: 1, Flags [T]: Incomplete
    0x0000: 02
    AS Path (2), length: 6, Flags [T]: 65001
    0x0000: 0201 0000 fde9
    Next Hop (3), length: 4, Flags [T]: gateway
    0x0000: c0a8 0300
    Updated routes:
      0.0.0.0/0
      10.1.0.0/16
      10.10.0.0/16
      10.11.0.0/16
11 18:55:26.907352 IP (tos 0xc0, ttl 1, id 11097, offset 0, flags [DF], proto TCP (6), length 75)
gateway.49124 > CE.bgp: Flags [P.], cksum 0xcb9a (correct), seq 138:161, ack 221, win 30016, options [nop,nop,TS val 4154842900 ecr 4109049241], length 23: BGP
  Update Message (2), length: 23
  End-of-Rib Marker (empty NLRI)
```



Tcpdump on CE

```
CE# diag utilities tcpdump verbosity level4 destination-ip 192.168.3.0
tcpdump: listening on any, link-type LINUX_SLL (Linux cooked), capture size 262144 bytes
1 22:02:38.251713 IP (tos 0xc0, ttl 64, id 0, offset 0, flags [DF], proto TCP (6), length 60)
CE.bgp > 192.168.3.0.49310: Flags [S.], cksum 0x4cf5 (correct), seq 3013020605, ack 4220967622, win 28960, options [mss 1460,sackOK,TS val 4120280594 ecr 4166074249,nop,wscale 7], length 0
2 22:02:38.263343 IP (tos 0xc0, ttl 1, id 11448, offset 0, flags [DF], proto TCP (6), length 118)
CE.bgp > 192.168.3.0.49310: Flags [P.], cksum 0xd40d (correct), seq 1:67, ack 1, win 227, options [nop,nop,TS val 4120280606 ecr 4166074258], length 66: BGP
Open Message (1), length: 66
Version 4, my AS 65002, Holdtime 180s, ID 192.168.1.2
Optional parameters, length: 37
Option Capabilities Advertisement (2), length: 35
Multiprotocol Extensions (1), length: 4
AFI IPv4 (1), SAFI Unicast (1)
0x0000: 0001 0001
Route Refresh (2), length: 0
Route Refresh (Cisco) (128), length: 0
Enhanced Route Refresh (70), length: 0
no decoder for Capability 70
Graceful Restart (64), length: 6
Restart Flags: [none], Restart Time 120s
AFI IPv4 (1), SAFI Unicast (1), Forwarding state preserved: no
0x0000: 0078 0001 0100
Cooperative Route Filtering (3), length: 7
no decoder for Capability 3
0x0000: 0001 0001 0140 02
32-Bit AS Number (65), length: 4
4 Byte AS 65002
0x0000: 0000 fdea
3 22:02:38.272900 IP (tos 0xc0, ttl 1, id 11449, offset 0, flags [DF], proto TCP (6), length 52)
CE.bgp > 192.168.3.0.49310: Flags [.], cksum 0xeb4f (correct), seq 67, ack 67, win 227, options [nop,nop,TS val 4120280616 ecr 4166074270], length 0
4 22:02:38.278378 IP (tos 0xc0, ttl 1, id 11450, offset 0, flags [DF], proto TCP (6), length 71)
CE.bgp > 192.168.3.0.49310: Flags [P.], cksum 0xe71c (correct), seq 67:86, ack 67, win 227, options [nop,nop,TS val 4120280621 ecr 4166074270], length 19: BGP
Keepalive Message (4), length: 19
5 22:02:38.321861 IP (tos 0xc0, ttl 1, id 11451, offset 0, flags [DF], proto TCP (6), length 52)
CE.bgp > 192.168.3.0.49310: Flags [.], cksum 0xeafl (correct), seq 86, ack 86, win 227, options [nop,nop,TS val 4120280664 ecr 4166074278], length 0
22:02:38.321902 IP (tos 0xc0, ttl 1, id 11452, offset 0, flags [DF], proto TCP (6), length 187)
CE.bgp > 192.168.3.0.49310: Flags [P.], cksum 0xa2a4 (correct), seq 86:221, ack 86, win 227, options [nop,nop,TS val 4120280665 ecr 4166074321], length 135: BGP
Route Refresh Message (5), length: 65
AFI IPv4 (1), SAFI Unicast (1)
0x0000: ffff ffff ffff ffff ffff ffff ffff ffff
0x0010: 0041 0500 0100 0101 4000 2600 0000 0028
0x0020: 0000 100a 0b00 0000 001e 0000 100a 0a00
0x0030: 0000 0014 0000 100a 0100 0000 000a 0000
0x0040: 00
Update Message (2), length: 47
Origin (1), length: 1, Flags [T]: Incomplete
0x0000: 02
AS Path (2), length: 6, Flags [T]: 65002
0x0000: 0201 0000 fdea
Next Hop (3), length: 4, Flags [T]: CE
0x0000: c0a8 0301
Updated routes:
172.16.1.0/24
Update Message (2), length: 23
End-of-Rib Marker (empty NLRI)
7 22:02:38.335178 IP (tos 0xc0, ttl 1, id 11453, offset 0, flags [DF], proto TCP (6), length 52)
CE.bgp > 192.168.3.0.49310: Flags [.], cksum 0xe9ef (correct), seq 221, ack 139, win 227, options [nop,nop,TS val 4120280678 ecr 4166074334], length 0
8 22:02:38.342190 IP (tos 0xc0, ttl 1, id 11454, offset 0, flags [DF], proto TCP (6), length 52)
CE.bgp > 192.168.3.0.49310: Flags [.], cksum 0xe9ca (correct), seq 221, ack 162, win 227, options [nop,nop,TS val 4120280685 ecr 4166074341], length 0
9 22:02:42.788615 ARP, Ethernet (len 6), IPv4 (len 4), Request who-has 192.168.3.0 tell 192.168.3.0, length 28
10 22:03:29.774528 IP (tos 0xc0, ttl 1, id 11455, offset 0, flags [DF], proto TCP (6), length 71)
CE.bgp > 192.168.3.0.49310: Flags [P.], cksum 0x1cb4 (correct), seq 221:240, ack 162, win 227, options [nop,nop,TS val 4120332117 ecr 4166074341], length 19: BGP
Keepalive Message (4), length: 19

^C
10 packets captured
10 packets received by filter
0 packets dropped by kernel
```

ORF prefix-list



ORF prefix-list exchange

With soft-reconfiguration in

```
6 10:41:28.106833 IP (tos 0xc0, ttl 1, id 60855, offset 0, flags [DF],
proto TCP (6), length 145)
  CE.bgp > 192.168.3.0.48460: Flags [P.], cksum 0xfc4b (correct), seq
86:179, ack 86, win 227, options [nop,nop,TS val 3551286768 ecr
3854295444], length 93: BGP
```

Route Refresh Message (5), length: 23

AFI IPv4 (1), SAFI Unicast (1)

0x0000: ffff ffff ffff ffff ffff ffff ffff ffff

0x0010: 0017 0500 0100 01

Update Message (2), length: 47

Origin (1), length: 1, Flags [T]: Incomplete

0x0000: 02

AS Path (2), length: 6, Flags [T]: 65002

0x0000: 0201 0000 fdea

Next Hop (3), length: 4, Flags [T]: CE

0x0000: c0a8 0301

Updated routes:

172.16.1.0/24

Update Message (2), length: 23

End-of-Rib Marker (empty NLRI)

Without soft-reconfiguration in

```
5 10:49:06.718840 IP (tos 0xc0, ttl 1, id 47103, offset 0, flags [DF], proto TCP (6),
length 107)
  CE.bgp > 192.168.3.0.48470: Flags [P.], cksum 0xfcb4 (correct), seq 86:141, ack 86,
win 227, options [nop,nop,TS val 3551745380 ecr 3854754046], length 55: BGP
```

Route Refresh Message (5), length: 55

AFI IPv4 (1), SAFI Unicast (1)

0x0000: ffff ffff ffff ffff ffff ffff ffff ffff

0x0010: 0037 0500 0100 0101 4000 1c00 0000 001e

0x0020: 0000 100a 0a00 0000 0014 0000 100a 0100

0x0030: 0000 000a 0000 00

```
6 10:49:06.730110 IP (tos 0xc0, ttl 1, id 47104, offset 0, flags [DF], proto TCP
(6), length 122)
  CE.bgp > 192.168.3.0.48470: Flags [P.], cksum 0xbd15 (correct), seq 141:211, ack 136,
win 227, options [nop,nop,TS val 3551745391 ecr 3854754063], length 70: BGP
```

Update Message (2), length: 47

Origin (1), length: 1, Flags [T]: Incomplete

0x0000: 02

AS Path (2), length: 6, Flags [T]: 65002

0x0000: 0201 0000 fdea

Next Hop (3), length: 4, Flags [T]: CE

0x0000: c0a8 0301

Updated routes:

172.16.1.0/24

Update Message (2), length: 23

End-of-Rib Marker (empty NLRI)

Configuration

aruba

a Hewlett Packard
Enterprise company

ORF configuration

2 steps per IPv4 or IPv6 AF

- This feature consists in two commands:

1. **Enable the ORF capability:** the router notifies its neighbor about its ORF capability:

- Typically a PE is configured to receive ORF requests
- Typically a CE is configured to send ORF requests
- Currently the filter supported is of type prefix-list only (RFC specifies other filters as well; ex: AS-path filter..).

2. **Configure the prefix-list filter:**

- Configure the prefix-list on the CE for “in-bound” direction, to be sent to the PE.
- The CE sends this prefix-list to PE in a **ROUTE REFRESH message**.
- The prefix-list received at the PE is used for outbound filtering towards the CE.

CE

```
CE(config-bgp-ipv4-uc)# neighbor 192.168.3.0 capability
  orf Outbound Route Filtering
CE(config-bgp-ipv4-uc)# neighbor 192.168.3.0 capability orf
  prefix-list Filter based on the prefix-list
CE(config-bgp-ipv4-uc)# neighbor 192.168.3.0 capability orf prefix-list
  both Configure the capability to send and receive the ORF for the
  neighbor
  receive Configure the capability to receive the ORF from the neighbor
  send Configure the capability to send the ORF to the neighbor
CE(config-bgp-ipv4-uc)# neighbor 192.168.3.0 capability orf prefix-list send
```

capability

PE

```
PE(config-bgp-ipv4-uc)# neighbor 192.168.3.1 capability
  orf Outbound Route Filtering
PE(config-bgp-ipv4-uc)# neighbor 192.168.3.1 capability orf
  prefix-list Filter based on the prefix-list
PE(config-bgp-ipv4-uc)# neighbor 192.168.3.1 capability orf prefix-list
  both Configure the capability to send and receive the ORF for the
  neighbor
  receive Configure the capability to receive the ORF from the neighbor
  send Configure the capability to send the ORF to the neighbor
PE(config-bgp-ipv4-uc)# neighbor 192.168.3.1 capability orf prefix-list receive
```

prefix-list

```
CE(config-bgp-ipv4-uc)# neighbor 192.168.3.0 orf-prefix-list test1
  in Apply the ORF prefix-filter on received NLRI's from the neighbor
CE(config-bgp-ipv4-uc)# neighbor 192.168.3.0 orf-prefix-list test1 in
```



Baseline

Without ORF



Baseline

Without ORF



PE
AS65001



192.168.168.3.0

192.168.3.1



CE
AS65002

```
PE# sh bgp ipv4 unicast neighbors 192.168.3.1
Codes: ^ Inherited from peer-group
```

VRF : default

BGP Neighbor 192.168.3.1 (External)

Description	:		
Peer-group	:		
Remote Router Id	:	192.168.1.2	Local Router Id : 192.168.1.1
Remote AS	:	65002	Local AS : 65001
Remote Port	:	179	Local Port : 44448
State	:	Established	Admin Status : Up
Conn. Established	:	5	Conn. Dropped : 4
Passive	:	No	Update-Source :
Cfg. Hold Time	:	180	Cfg. Keep Alive : 60
Neg. Hold Time	:	180	Neg. Keep Alive : 60
Up/Down Time	:	00h:00m:15s	Alt. Local-AS : 0
Local-AS Prepend	:	No	
BFD	:	Disabled	
Password	:		
Last Err Sent	:	No Error	
Last SubErr Sent	:	No Error	
Last Err Rcvd	:	Cease	
Last SubErr Rcvd	:	Administrative Shutdown	
Graceful-Restart	:	Enabled	Gr. Restart Time : 120
Gr. Stalepath Time	:	300	Remove Private-AS : No
TTL	:	1	Local Cluster-ID :
Weight	:	0	Fall-over : No
Confederation-Peers	:	No	

Message statistics	Sent	Rcvd
Open	5	5
Notification	0	2
Updates	10	9
Keepalives	18	19
Route Refresh	0	0
Total	33	35

Capability	Advertised	Received
Route Refresh	Yes	Yes
Graceful Restart	Yes	Yes
Add-Path	No	No
Four Octet ASN	Yes	Yes
Address family IPv4 Unicast	Yes	Yes
Address family IPv6 Unicast	No	No
Address family L2VPN EVPN	No	No

Address Family : IPv4 Unicast

Rt. Reflect. Client	:	No	Send Community	:	
Allow-AS in	:	0	Advt. Interval	:	30
Max. Prefix	:	64000	Soft Reconfig In	:	
NextHop-Self	:		Default-Originate	:	
Cfg. Add-Path	:				
Neg. Add-Path	:	Disable			

Routemap In	:	
Routemap Out	:	
ORF type	:	Prefix-list
ORF capability	:	

```
CE# sh bgp ipv4 unicast neighbors 192.168.3.0
Codes: ^ Inherited from peer-group
```

VRF : default

BGP Neighbor 192.168.3.0 (External)

Description	:		
Peer-group	:		
Remote Router Id	:	192.168.1.1	Local Router Id : 192.168.1.2
Remote AS	:	65001	Local AS : 65002
Remote Port	:	179	Local Port : 55498
State	:	Established	Admin Status : Up
Conn. Established	:	1	Conn. Dropped : 0
Passive	:	No	Update-Source :
Cfg. Hold Time	:	180	Cfg. Keep Alive : 60
Neg. Hold Time	:	180	Neg. Keep Alive : 60
Up/Down Time	:	00h:01m:59s	Alt. Local-AS : 0
Local-AS Prepend	:	No	
BFD	:	Disabled	
Password	:		
Last Err Sent	:	No Error	
Last SubErr Sent	:	No Error	
Last Err Rcvd	:	No Error	
Last SubErr Rcvd	:	No Error	
Graceful-Restart	:	Enabled	Gr. Restart Time : 120
Gr. Stalepath Time	:	300	Remove Private-AS : No
TTL	:	1	Local Cluster-ID :
Weight	:	0	Fall-over : No
Confederation-Peers	:	No	

Message statistics	Sent	Rcvd
Open	1	1
Notification	0	0
Updates	2	2
Keepalives	3	3
Route Refresh	0	0
Total	6	6

Capability	Advertised	Received
Route Refresh	Yes	Yes
Graceful Restart	Yes	Yes
Add-Path	No	No
Four Octet ASN	Yes	Yes
Address family IPv4 Unicast	Yes	Yes
Address family IPv6 Unicast	No	No
Address family L2VPN EVPN	No	No

Address Family : IPv4 Unicast

Rt. Reflect. Client	:	No	Send Community	:	
Allow-AS in	:	0	Advt. Interval	:	30
Max. Prefix	:	64000	Soft Reconfig In	:	Yes
NextHop-Self	:		Default-Originate	:	
Cfg. Add-Path	:				
Neg. Add-Path	:	Disable			

Routemap In	:	filtering1
Routemap Out	:	
ORF type	:	Prefix-list
ORF capability	:	

Baseline

Without ORF



```
PE# show bgp ipv4 unicast neighbors 192.168.3.1 advertised-routes
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 : default
Local Router-ID 192.168.1.1
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	0	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.2.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.3.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.4.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.5.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.6.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.7.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.8.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.9.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.11.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.12.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.13.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.14.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.15.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.16.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.17.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.18.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.19.0.0/16	192.168.3.0	0	0	0	65001 ?

Total number of entries 20

```
CE# sh bgp ipv4 unicast neighbors 192.168.3.0 routes
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 : default
Local Router-ID 192.168.1.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	100	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.2.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.3.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.4.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.5.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.6.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.7.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.8.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.9.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.11.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.12.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.13.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.14.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.15.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.16.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.17.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.18.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.19.0.0/16	192.168.3.0	0	100	0	65001 ?

Total number of entries 20

ORF Configuration

ORF capability + prefix-list



```
router bgp 65001
  bgp router-id 192.168.1.1
  neighbor 192.168.3.1 remote-as 65002
  address-family ipv4 unicast
    neighbor 192.168.3.1 activate
    neighbor 192.168.3.1 capability orf prefix-list receive

  redistribute static
  exit-address-family
```

```
ip prefix-list prefix-listA seq 10 permit 0.0.0.0/0
ip prefix-list prefix-listA seq 20 permit 10.1.0.0/16
ip prefix-list prefix-listA seq 30 permit 10.10.0.0/16
!
!
router bgp 65002
  bgp router-id 192.168.1.2
  neighbor 192.168.3.0 remote-as 65001
  address-family ipv4 unicast
    neighbor 192.168.3.0 activate
    neighbor 192.168.3.0 capability orf prefix-list send
    neighbor 192.168.3.0 orf-prefix-list prefix-listA in
  redistribute static
  exit-address-family
```

ORF Configuration

ORF capability

PE
AS65001



192.168.168.3.0

192.168.3.1



CE
AS65002

```
PE# show bgp ipv4 unicast neighbors 192.168.3.1
Codes: ^ Inherited from peer-group
```

VRF : default

BGP Neighbor 192.168.3.1 (External)

Description	:		
Peer-group	:		
Remote Router Id	:	192.168.1.2	Local Router Id : 192.168.1.1
Remote AS	:	65002	Local AS : 65001
Remote Port	:	179	Local Port : 44742
State	:	Established	Admin Status : Up
Conn. Established	:	1	Conn. Dropped : 0
Passive	:	No	Update-Source :
Cfg. Hold Time	:	180	Cfg. Keep Alive : 60
Neg. Hold Time	:	180	Neg. Keep Alive : 60
Up/Down Time	:	00h:10m:45s	Alt. Local-AS : 0
Local-AS Prepend	:	No	
BFD	:	Disabled	
Password	:		
Last Err Sent	:	No Error	
Last SubErr Sent	:	No Error	
Last Err Rcvd	:	No Error	
Last SubErr Rcvd	:	No Error	
Graceful-Restart	:	Enabled	Gr. Restart Time : 120
Gr. Stalepath Time	:	300	Remove Private-AS : No
TTL	:	1	Local Cluster-ID :
Weight	:	0	Fall-over : No
Confederation-Peers	:	No	

Message statistics	Sent	Rcvd
Open	1	1
Notification	0	0
Updates	2	2
Keepalives	13	13
Route Refresh	0	1
Total	16	17

Capability	Advertised	Received
Route Refresh	Yes	Yes
Graceful Restart	Yes	Yes
Add-Path	No	No
Four Octet ASN	Yes	Yes
Address family IPv4 Unicast	Yes	Yes
Address family IPv6 Unicast	No	No
Address family L2VPN EVPN	No	No

Address Family : IPv4 Unicast

Rt. Reflect. Client	:	No	Send Community	:	
Allow-AS in	:	0	Advt. Interval	:	30
Max. Prefix	:	64000	Soft Reconfig In	:	
NextHop-Self	:		Default-Originate	:	
Cfg. Add-Path	:				
Neg. Add-Path	:	Disable			

Routemap In	:	
Routemap Out	:	
ORF type	:	Prefix-list
ORF capability	:	receive

```
CE# show bgp ipv4 un nei 192.168.3.0
Codes: ^ Inherited from peer-group
```

VRF : default

BGP Neighbor 192.168.3.0 (External)

Description	:		
Peer-group	:		
Remote Router Id	:	192.168.1.1	Local Router Id : 192.168.1.2
Remote AS	:	65001	Local AS : 65002
Remote Port	:	44742	Local Port : 179
State	:	Established	Admin Status : Up
Conn. Established	:	2	Conn. Dropped : 1
Passive	:	No	Update-Source :
Cfg. Hold Time	:	180	Cfg. Keep Alive : 60
Neg. Hold Time	:	180	Neg. Keep Alive : 60
Up/Down Time	:	00h:12m:21s	Alt. Local-AS : 0
Local-AS Prepend	:	No	
BFD	:	Disabled	
Password	:		
Last Err Sent	:	No Error	
Last SubErr Sent	:	No Error	
Last Err Rcvd	:	No Error	
Last SubErr Rcvd	:	No Error	
Graceful-Restart	:	Enabled	Gr. Restart Time : 120
Gr. Stalepath Time	:	300	Remove Private-AS : No
TTL	:	1	Local Cluster-ID :
Weight	:	0	Fall-over : No
Confederation-Peers	:	No	

Message statistics	Sent	Rcvd
Open	2	2
Notification	0	0
Updates	4	4
Keepalives	16	16
Route Refresh	2	0
Total	24	22

Capability	Advertised	Received
Route Refresh	Yes	Yes
Graceful Restart	Yes	Yes
Add-Path	No	No
Four Octet ASN	Yes	Yes
Address family IPv4 Unicast	Yes	Yes
Address family IPv6 Unicast	No	No
Address family L2VPN EVPN	No	No

Address Family : IPv4 Unicast

Rt. Reflect. Client	:	No	Send Community	:	
Allow-AS in	:	0	Advt. Interval	:	30
Max. Prefix	:	64000	Soft Reconfig In	:	
NextHop-Self	:		Default-Originate	:	
Cfg. Add-Path	:				
Neg. Add-Path	:	Disable			

Routemap In	:	
Routemap Out	:	
ORF type	:	Prefix-list
ORF capability	:	send

ORF Configuration

Outcome



```
PE# sh bgp ipv4 unicast neighbors 192.168.3.1 advertised-routes
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 : default
Local Router-ID 192.168.1.1
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	0	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.2.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.3.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.4.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.5.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.6.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.7.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.8.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.9.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.11.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.12.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.13.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.14.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.15.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.16.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.17.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.18.0.0/16	192.168.3.0	0	0	0	65001 ?
*>e 10.19.0.0/16	192.168.3.0	0	0	0	65001 ?
Total number of entries 20					

```
CE# show bgp ipv4 unicast neighbors 192.168.3.0 routes
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 : default
Local Router-ID 192.168.1.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 0.0.0.0/0	192.168.3.0	0	100	0	65001 ?
*>e 10.1.0.0/16	192.168.3.0	0	100	0	65001 ?
*>e 10.10.0.0/16	192.168.3.0	0	100	0	65001 ?
Total number of entries 3					

Best Practices

aruba

a Hewlett Packard
Enterprise company

ORF Best Practices

- **For operational enforcement:**
Use different prefix-list names for ORF and for route-map in-bound filtering.
In case of a human error occurs on one of the prefix-list, the other one should still enforce the routing control appropriately.
- **Add anti-spoofing rules in ORF filter:**
The CE should not received prefixes advertisement for routes that are supposed to be behind the administrative domain of the CE.

aruba

a Hewlett Packard
Enterprise company

Troubleshooting

ORF Troubleshooting

On the PE: show bgp received orf-prefix-filter



```
show bgp [{ipv4 unicast | ipv6 unicast| ipv4 unicast {vrf <vrf-name>} | all unicast} neighbors] <ip4/6 address | all> received orf-prefix-filter
```

```
PE# show bgp ipv4 unicast neighbors 192.168.3.1 received orf-prefix-list
ip prefix-list 192.168.3.1: 3 entries
  seq 10 permit 0.0.0.0/0
  seq 20 permit 10.1.0.0/16
  seq 30 permit 10.10.0.0/16
```

ORF Troubleshooting

show bgp neighbors

```
switch# show bgp ipv4 unicast neighbors 10.1.1.2
Codes: ^ Inherited from peer-group
```

VRF : default

BGP Neighbor 10.1.1.2 (Internal)

Description	:		
Peer-group	:		
Remote Router Id	:	10.1.1.2	Local Router Id : 1.0.0.1
Remote AS	:	1	Local AS : 1
Remote Port	:	0	Local Port : 0
State	:	Idle	Admin Status : Up
Conn. Established	:	0	Conn. Dropped : 0
Passive	:	No	Update-Source :
Cfg. Hold Time	:	180	Cfg. Keep Alive : 60
Neg. Hold Time	:	0	Neg. Keep Alive : 0
Up/Down Time	:	00h:00m:00s	Connect-Retry Time : 120
Local-AS Prepend	:	No	Alt. Local-AS : 0
BFD	:	Disabled	
Password	:		
Last Err Sent	:	No Error	
Last SubErr Sent	:	No Error	
Last Err Rcvd	:	No Error	
Last SubErr Rcvd	:	No Error	
Graceful-Restart	:	Enabled	Gr. Restart Time : 120
Gr. Stalepath Time	:	150	Remove Private-AS : No
TTL	:	255	Local Cluster-ID :
Weight	:	0	Fall-over : No

Message statistics	Sent	Rcvd
Open	0	0
Notification	0	0
Updates	0	0
Keepalives	0	0
Route Refresh	0	0
Total	0	0

Capability	Advertised	Received
Route Refresh	Yes	No
Graceful Restart	Yes	No
Four Octet ASN	Yes	No
Add-Path	Yes	Yes
Address family IPv4 Unicast	Yes	No
Address family IPv6 Unicast	No	No
Address family L2VPN EVPN	No	No

Address Family : IPv4 Unicast

Rt. Reflect. Client	: No	Send Community	:
Allow-AS in	: 0	Advt. Interval	: 30
Max. Prefix	: 64000	Soft Reconfig In	:
Nexthop-Self	:	Default-Originate	:

Routemap In	:
Routemap Out	:

ORF Type : Prefix-list
ORF Capability : Receive

Cfg. Add-Path	:
Neg. Add-Path	: Disable



ORF Advanced Troubleshooting

Diag dump + MIB dump and Fastlogs

- Below MIB tables should be verified for the configured **ORF capability** and **prefix-list** values.
These table are captured as part of “diag bgp dump mib” command.
 - python /etc/mib.py get localhost bgpPeerOrfCapabilityTable
 - python /etc/mib.py get localhost bgpRouteMapTable
 - python /etc/mib.py get localhost bgpPeerAfiSafiTable
- **Fastlogs:** ovs-appctl -t hpe-routing fastlog show bgp_dump
 - BGP ORF Capability[1] set for peer[10.10.10.2] afi[1]
 - Route Map delete ORF_RMAP_IN:_10.10.10.2
 - rmap_in = ORF_RMAP_IN:_10.10.10.2 : rmap_out = :

ORF Troubleshooting

Event logs

- There is no event logs for ORF as this is not required for this feature.

Demo



Demo

CE

```
hostname CE
router ospf 1
  area 0.0.0.0
interface 1/1/1
  no shutdown
  ip address 10.10.10.1/24
      ip ospf 1 area 0.0.0.0
!
ip prefix-list pl seq 10 permit 10.10.10.0/24
ip prefix-list pl seq 20 permit 20.20.20.0/24
!
!
!
router bgp 40
  neighbor 10.10.10.2 remote-as 100
  address-family ipv4 unicast
    neighbor 10.10.10.2 activate
    neighbor 10.10.10.2 capability orf prefix-list send
    neighbor 10.10.10.2 orf-prefix-list pl in
    redistribute local loopback
  exit-address-family
!
```

PE

```
hostname PE
interface 1/1/1
  no shutdown
  ip address 10.10.10.2/24
!
interface 1/1/2
  no shutdown
  ip address 20.20.20.2/24
!
interface 1/1/3
  no shutdown
  ip address 30.30.30.2/24
!
router bgp 100
  neighbor 10.10.10.1 remote-as 40
  address-family ipv4 unicast
    neighbor 10.10.10.1 activate
    neighbor 10.10.10.1 capability orf prefix-list receive
    redistribute connected
  exit-address-family
!
```

Thank you

vincent.giles@hpe.com