Updates and Discussion on RD-ORF Solutions

draft-wang-idr-rd-orf-06
draft-wang-idr-vpn-routes-control-analysis-03

Aijun Wang (China Telecom)
Wei Wang (China Telecom)
Gyan Mishra (Verizon)
Haibo Wang (Huawei)
Shunwan Zhuang (Huawei)
Jie Dong (Huawei)

IETF Interim, August 2021

Motivation of This Presentation

- ✓ Describes the scenarios and solutions to control excessive VPN routes within one AS or across different ASes
- ✓ Reaches consensus on the proposed solutions
- ✓ Forwards the updated solution draft(if necessary, also the scenario draft)

Scenario-1 and Solution
(Intra-AS, Different RD, one RT)

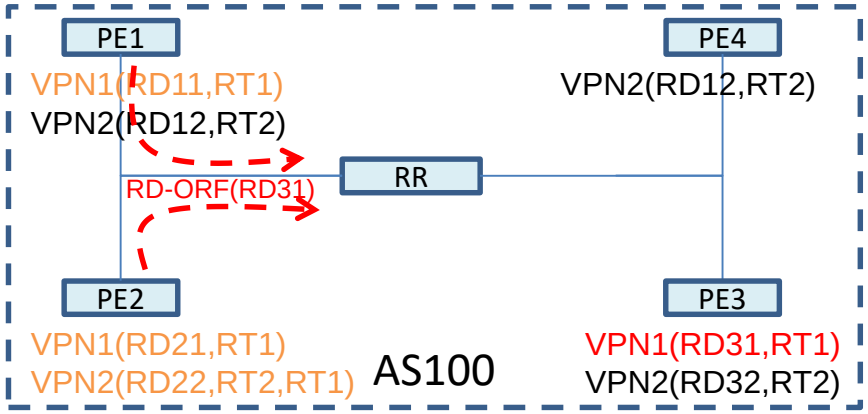

1. Shared BGP session between RR and PE for VRFs
   - ① RD is allocated per VPN/per PE
   - ② PE3 send excessive VPN routes with RT1
   - ③ PE1、PE2 will be influenced with the excessive VPN routes
2. PE/RR should have some mechanisms to identify and control the advertisement of specified excessive VPN routes.

Proposed Solution

① Once PE1 detects the VPN1 VRF is overflowed , and:
- ✓ The excessive VPN routes has RD31, associated RT is RT1
- ✓ No other VRFs on it to import the VPN routes with RT1

PE1 triggers the RD-ORF message to RR(RD field is set to RD31)

② Once PE2 detects the VPN1 VRF is overflowed, and:
- ✓ The excessive VPN routes has RD31, associated RT is RT1
- ✓ There is other VRF on it to import the VPN routes with RT1

PE2 triggers the RD-ORF message to RR(RD field is set to RD31) only when all the VRFs that import RT1 are overflowed. Else, it discards the overflowed VPN routes locally.

Scenario-2 and Solution
(Intra-AS, Different RD, Multiple RTs)

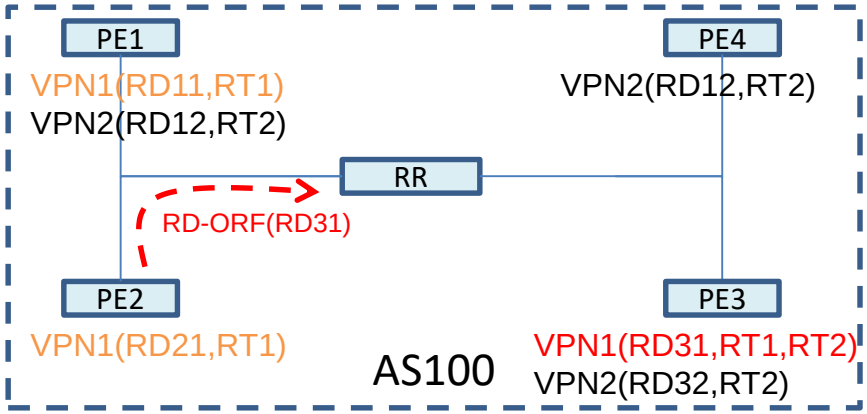

① RD is allocated per VPN/per PE
② Multiple RTs are associated with such VPN routes, and be imported into different VRFs in other devices(PE1)
③ PE3 send excessive VPN routes with RT1, RT2.

Proposed Solution

① Once PE1 detects the VPN1 VRF is overflowed , and:
- ✓ The excessive VPN routes has RD31, associated with RT1, RT2
- ✓ There are different VRFs on it import the VPN routes respectively with RT1, RT2

② PE1 triggers the RD-ORF message to RR(RD field is set to RD31) only when all these VRFs are overflowed; else, it discards the overflowed VPN routes locally.

③ In this example, PE1 will not trigger RD-ORF, only PE2 will trigger RD-ORF(RD31).

Scenario-3 and Solution (Intra-AS, Unique RD)

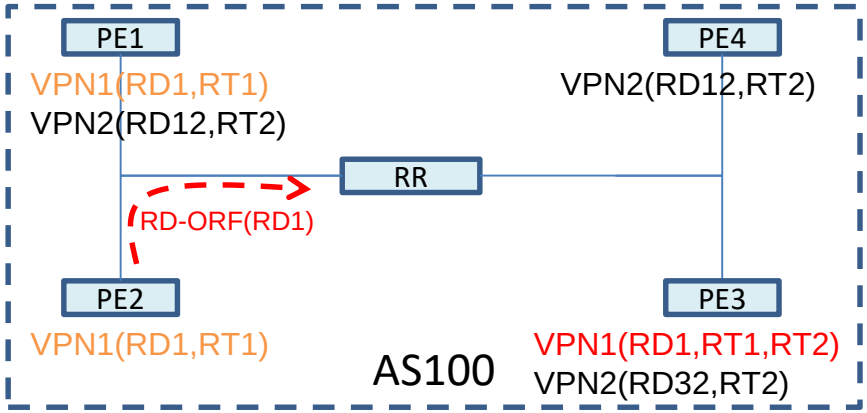

① RD is allocated per VPN
② One/Multiple RTs are associated with such VPN routes, and be imported into different VRFs in other devices(PE1)
③ PE3 send excessive VPN routes with RD1 and attached RT1, RT2.

Proposed Solution

① Once PE2 detects the VPN1 VRF is overflowed , and:
  - ✓ The excessive VPN routes has RD1
  - ✓ There is no other VRF on it import the VPN routes with RT1,RT2
② PE2 triggers the RD-ORF message to RR(RD field is set to RD1)
③ RR withdraws and stops to advertise such excessive VPN routes to PE2
④ The communication among PE2 with other PEs(PE1, PE3) for VPN1 will be influenced. It is acceptable.

Scenario-4 and Solution (Inter-AS)

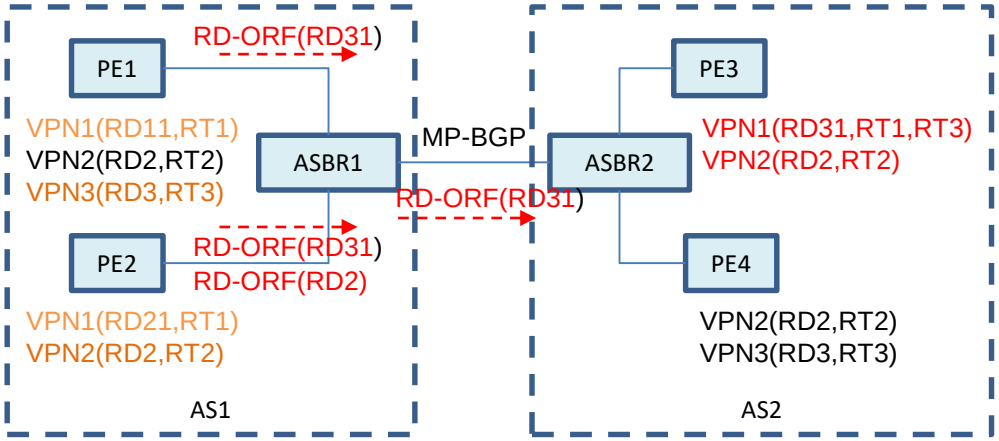

1. BGP session is shared between PE/ASBR and ASBR1/ASBR2
2. Overwhelming VPN routes in one VRF can certainly influence the control/forward behavior of the PE for other VRFs.
   a) RD may be allocated per VRF per PE, as that in intra-AS(VPN1)
   b) RD may also be allocated per VRF only(VPN2/VPN3)
3. The PE/ASBR should have some mechanism to control the advertisement of excessive VPN routes.

Proposed Solution

1. If the excessive VPN routes are from VPN1 or VPN2 on PE3, once the PE1 detects the overflow of VPNs on it:
   - ✓ PE1 will trigger the RD-ORF(RD field is set to RD31 or RD2) message to ASBR1
2. When ASBR1 receives such RD-ORF message, it checks:
   - ✓ If all its downstream peers sent the same message, or the process of excessive VPN routes have exceed its capabilities, it will send such message to upstream peer(ASBR2)
   - ✓ Or else, it will filter the excessive VPN routes on its side, on behalf of the trigger device(PE2)

3. For example, in above figure:
   - ✓ If PE1 and PE2 all sent the RD-ORF(with RD field set to RD31) message, the ASBR1 can send out the RD-ORF(with RD field set to RD31) message then.
   - ✓ If only PE2 sent the RD-ORF(with RD field set to RD2) message, ASBR1 will filter the excessive VPN routes to PE2. PE1 can still receive such routes.

Solution Summary

1. Control message for the specified VPN routes can be triggered automatically upon the excessive VPN routes
2. Control message should be sent only out the device:
   - ① For PE: when all the VRFs on it don't want to process it
   - ② For RR: when all its BGP clients don't want to process it
   - ③ For ASBR: when all its BGP peers within one AS don’t want to process it
   - ④ Or for all of them: the process of such excessive routes has exceed its own capability.

3. The removal of RD-ORF message should be manual to avoid the possible flapping of excessive VPN routes advertisement.
4. RD information is enough, no need to add RT.
   - ① The same RT may be imported by several VRFs.
   - ② Within one PE device, RT can’t uniquely identify one VPN. RD can accomplish this.

Further Action

- Comments?
- Is this clear to describe the problem and solution?
- If so, forward the draft(adopt directly or second WG adoption call?)

wangaj3@chinatelecom.cn
wangw36@chinatelecom.cn
gyan.s.mishra@verizon.com
rainsword.wang@huawei.com
zhuangshunwan@huawei.com
jie.dong@huawei.com
IETF Interim(111&112)