

Inter AS option D
(draft-mapathak-interas-option-d-00)

Manu Pathak

Keyur Patel

Arjun Sreekantiah

November 2012

Problem Statement

- Address scalability and security concerns of option A and option B in a new VPN InterAS option (option D)
- option A
 - provides for per VRF isolation, security and IP QoS
 - scalability concerns - requires a bgp session per VRF does not scale well as the number of VRFs increase.
- option B
 - single bgp session for VPN prefixes- scalable for a large number of VRFs
 - lacks per VRF isolation, security and IP QoS

InterAS option D

- New InterAS Option D
 - Hybrid model combining benefits of option A and B
 - Still retains VRF configuration from option A but does not require a BGP session per VRF. Single MP-BGP session is used to signal VPN prefixes, similar to option B. Thus, **there is no need to run BGP sessions on every sub-interface.**
 - Forwarding connections are still maintained on a per-VRF basis, similar to option A, so **separation of VPN traffic is possible.**
 - Provides for per VRF isolation, security and IP QoS while being scalable as the number of VRFs increase.

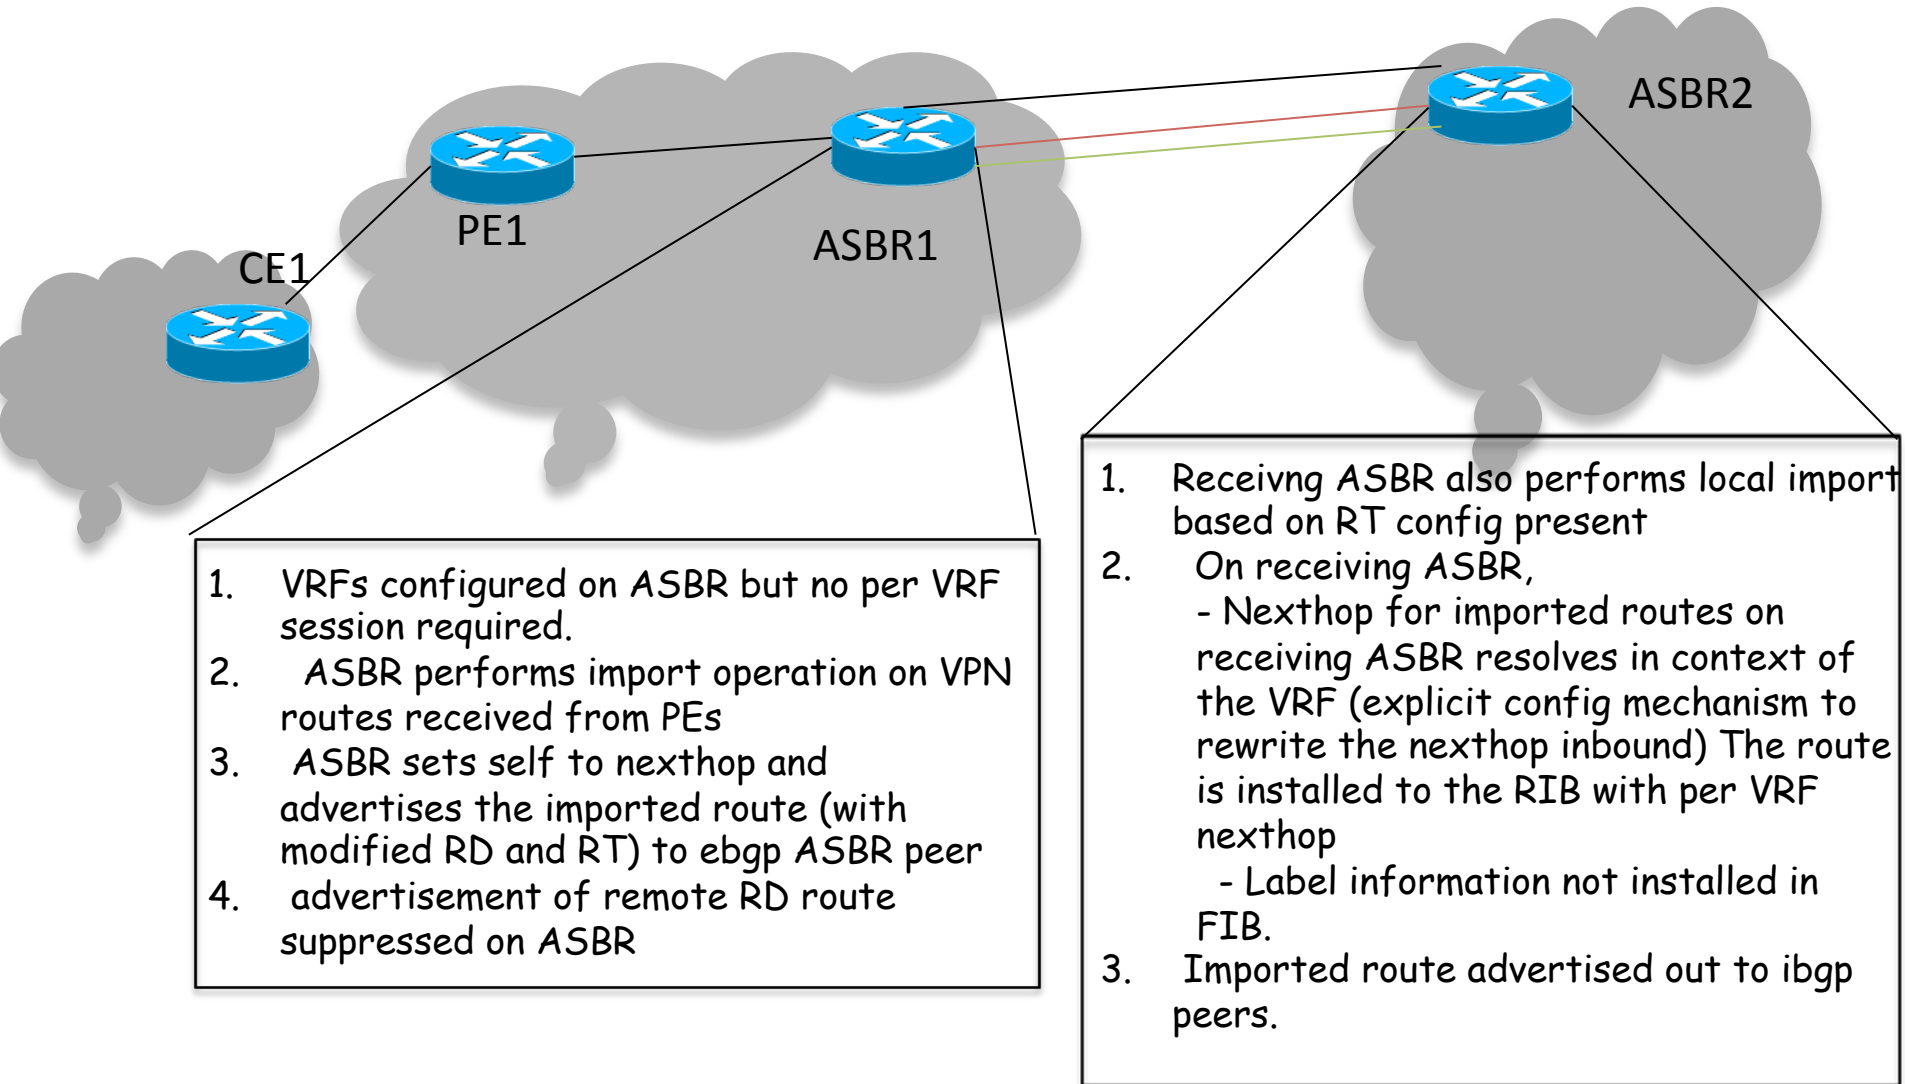
Option D operation

- VRFs configured on ASBR but no per VRF session required.
- ASBR performs import operation on VPN routes received from PEs
- ASBR sets self to nexthop and advertises the imported route (with modified RD and RT) with an allocated local label to ebgp ASBR peer
- Advertisement of received route (with unmodified RD) suppressed on ASBR

Option D operation

- Receiving ASBR also performs local import based on RT config present
- On receiving ASBR while processing the VPN route,
 - Nexthop for imported routes on receiving ASBR resolves in context of the VRF interface (either the VRF interface IP address is configured to be the same as that of the ASBR local address on both ASBRs or explicit config mechanism to rewrite the nexthop inbound). The route is installed in the VRF RIB with a per VRF nexthop.
 - Label information is not installed in FIB.
- Imported route advertised out to IBGP peers.

InterAS option D operation



InterAS option D forwarding

- Forwarding proceeds similar to option A
- CE sends a packet destined for N to the PE.
- PE encapsulates the packet with the VPN label allocated by ASBR and the IGP label (if any) needed to tunnel the packet to ASBR.
- The packet arrives on ASBR with the VPN Label, ASBR pops the VPN Label and sends the packet as IP to remote ASBR on the VRF interface.
- The IP packet arrives at remote ASBR on the VRF interface. Remote ASBR then encapsulates the packet with the VPN Label allocated by remote PE and the IGP label needed to tunnel the packet to remote PE.
- The packet arrives on remote PE with the VPN label; remote PE disposes the VPN label and forwards the IP packet to remote CE.

Option D: Other Considerations

- The draft also provides for a shared interface forwarding mode in which forwarding is similar to option B, but with local import for RD and RT rewrite. More details in the draft.
- Deployment of Option D is simplified if VRF next-hop is same as global next-hop, but this is not required.
- Route-targets need to be properly configured on PEs and ASBRs, and local import operation in option D can provide for application of per VRF policies at ASBRs.