

# Aggregation PE

draft-kulmala-l3vpn-aggregation-pe-00.txt

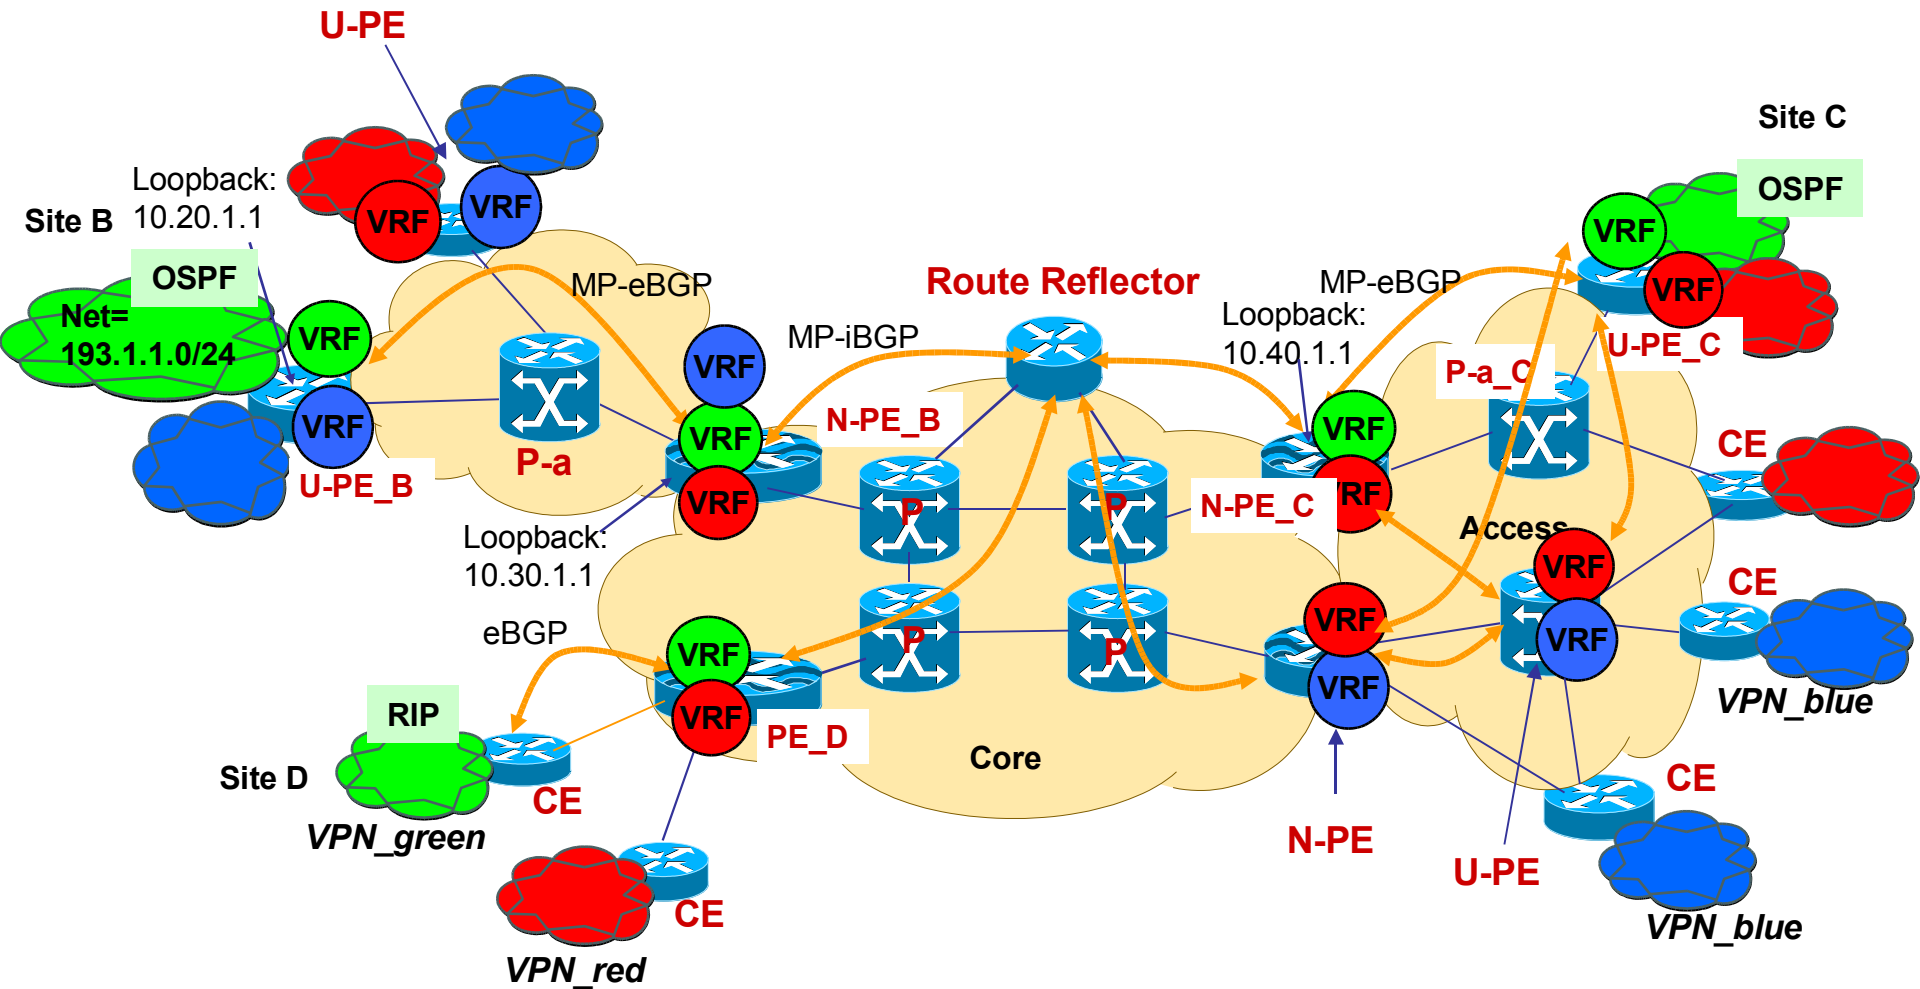
# Overview

- Use multiple sub-AS to separate access networks from core
  - Reduce number of LSPs
  - Allow TE with multiple areas
- Aggregation PE uses mechanism that is a hybrid of RFC2547bis ch 10 options a and b
  - Hides core network structure
  - Allows load balancing
  - N-PE has access to IP packets (->aggregation, filtering, QoS)
  - Provides implicit filtering of advertised routes

# Operation at aggregation PE

- Aggregation PE router will install VPN-IPv4 routes into VRFs
  - Aggregation PE will re-advertise the VPN-IPv4 routes that are imported into VRFs by re-exporting them from VRF RIBs
    - BGP next hop is set to itself
    - RD is replaced with the RD of the VRF the route is currently associated in aggregation PE
    - RTs are replaced with export route-targets of the VRF the route is currently associated in aggregation PE
- > Instead of “set of VRFs” (option b), complete VPNs are advertised

# Network structure



Connection below IP/MPLS

# Benefits of aggregation PE compared to 2547bis 10b

- Routes are associated to VPNs and IP packets are accessible
  - Per VPN QoS handling is possible
  - Per VPN aggregation is possible
  - Per VPN filtering is possible
  - Per VPN PSN tunnel selection is possible (e.g. separate some VPNs to separate RSVP tunnels in access, and use LDP in core)
    - End-to-end MPLS TE is possible
    - Can use TE selectively in some parts of network (e.g. LDP core, RSVP-TE access)
- Hides core network structure (RTs and RDs)
  - No unnecessary copies of same route with different RTs
  - If each N-PE assigns unique RD, load balancing between networks is possible
- Easy to use – implicit filtering is provided by configured VRFs

# Benefits of access networks in separate AS

- No full mesh LSPs between PEs --> RSVP TE is reasonable (-->stricter QoS guarantees for VPNs)
- LSPs limited to area -->inter-area TE can be accomplished without complex inter-area TE extension and without losing possibility of end-to-end TE control
  - As VRF information is available at N-PE, tunnels can be re-selected per VRF+prefix+qos basis on network boundaries
  - Exit point from access network can be easily selected using normal BGP methods
- Traffic within access network does not go via N-PE
- Divides routing process load (e.g. OSPF) between multiple U-PEs.

# Scalability

- Removes full mesh LSP requirement between PEs in different access networks
  - Very granular traffic engineering is possible in access (where it is most needed).
- N-PE needs to have VPN routes of the connected U-PEs
  - Having routes in VRFs does not consume more resources than having them in RIBs --> same BGP scalability properties with 2547bis option B
  - Removing multiple copies and aggregation makes this slightly more scalable than option B
- N-PE does not need to run customer IGP instances
- Unlike L2 access, this method can route traffic within access networks (traffic in local VPNs might never visit core).

# Pros and cons

- No LSPs needed between packet's ingress and egress PE routers that may be operated by different SPs: Security is enhanced.
- Allows using MPLS between ASBRs which eliminates multiple logical interfaces and multiple eBGP sessions between ASBRs.
- Allows MPLS network between ASBRs.
- Backward compatible – N-PE:s can be added to existing core (and connected to existing VPNs)
- Needs to keep VPN routes in ASBR routers, which may become scalability issue if lot of VRFs are needed.

# Next steps

- Comments ? Questions ?
- Discussion on the list.

THANK YOU !