

# VS (Virtual Subnet)

## draft-xu-virtual-subnet-03

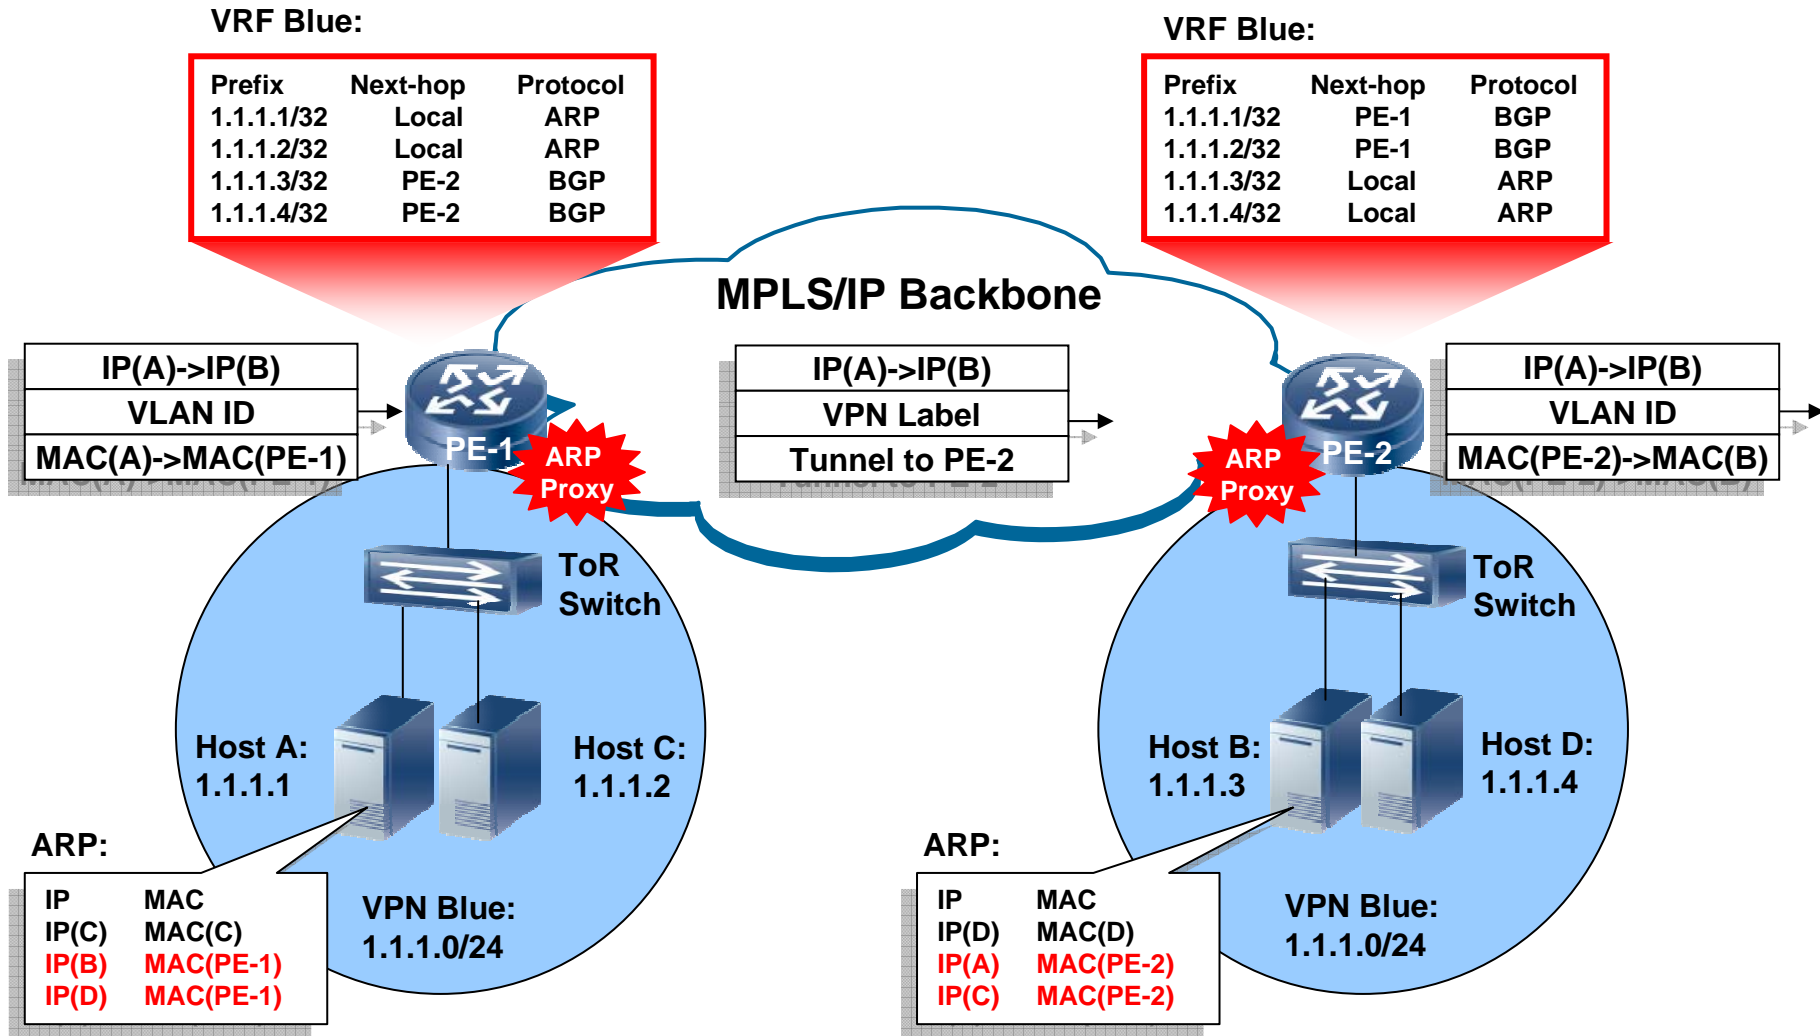
Xiaohu Xu

[<xuxh@huawei.com>](mailto:xuxh@huawei.com)

IETF 79, Beijing

# Virtual Subnet

one subnet spanning across multiple locations



# Local CE Host Discovery

- Local CE hosts are discovered through ARP learning.
  - PE sends unicast ARP requests to those learnt local CE hosts periodically to keep their corresponding ARP entries from expiring.
- To ensure the PE has learnt all local CE hosts, especially in the event of rebooting, ARP scan should be performed at least once after rebooting:
  - Option 1 (available today):
    - PE sends to its local site an ARP request for each IP address within the configured IP subnet in turn.
  - Option 2 (extensions to existing ARP needed):
    - PE sends to its local site an ARP request for a directed broadcast address (i.e., 255.255.255.255) or an ALL-Systems multicast group address (i.e., 224.0.0.1).
    - Any CE host receiving such ARP request should respond with an ARP reply containing its IP and MAC addresses.

# ARP Reduction

- Besides ARP learning, PE should perform the ARP proxy [RFC 925] function:
  - For an ARP request for a local CE host, discards it.
  - For an ARP request for a remote CE host, return its own MAC as an ARP reply.
  - For an ARP request for an unknown CE host (i.e., no matching VRF entry found), discards it.
- ARP broadcast traffic from CE hosts is limited to local VPN sites
  - ARP broadcast traffic would not be flooded across PEs.
  - ARP update for a CE host (e.g., triggered by VM mobility) would not trigger any BGP update as long as that CE host is still attached to its original PE and VRF instance (e.g., VM mobility within the VPN site).

# CE Migration (e.g., VM Mobility)

- CE migration within a VPN site.
  - PE just needs to update the corresponding ARP entry.
  - No BGP update is triggered.
- CE migration across VPN sites.
  - Upon learning a host route for a given local CE host via BGP, PE should immediately send an ARP request to that host to determine whether that host is still connected to it.
    - If not, PE should delete the corresponding ARP entry and host route for that CE host, and withdraw the corresponding BGP route advertised before.
    - Otherwise, it is judged as CE multi-homing.

# Comparison

	IPLS	VS
CE reachability Information Distribution	MAC reachability advertisement via LDP	IP reachability advertisement via BGP
ARP reduction mechanism	ARP cache/snooping (return a real MAC of the requested CE).	ARP proxy (return the MAC of the ARP proxy)
Eliminating ARP/unknown unicast flooding across PEs	No	Yes
CE multi-homing	Not support	Support natively
MAC table capacity pressure on Intermediary bridges	Need to learn MACs of both local and remote CEs. Not aging out learned MAC entries worsen such pressure.	Only need to learn local CE hosts' MAC addresses.

Comments?