BGP L3VPN Virtual PE Framework

draft-fang-l3vpn-virtual-pe-framework-01

L. Fang, D. Ward, R. Fernando, D. Rao (Cisco)

M. Napierala (AT&T)

N. Bitar (Verizon)

11/09/2012

IETF 85, Atlanta

Motivation

- Service Providers and large enterprises want to build virtual Private Clouds (vPCs) by leveraging existing BGP IP VPNs [RFC 4364], and extending it deep into the Service networks.
- Why IP VPN?
  - Is most deployed, extend to service network to complete vPC
  - Supports routing isolation and traffic separation at the edge
  - Is mature technology with 14 yrs wide deployment and large customer bases
  - Is proven to Scale (e.g. 7+ Million vpn routes in some network)
  - Support end-to-end vPC, especially applications are mostly IP
  - Rich policies for value added service creation beyond simple connectivity

Virtual PE Definition

- Virtual PE (vPE)

  A PE software instance which can reside in any network or compute devices. A common place for vPE can be a service network end device, e.g., a server which supports multiple client/application Virtual Machines (VMs), or a Top-of-Rack switch (ToR) in the Data Center. Another example can be a service node in a 3GPP network.

  The control and forwarding components of the vPE are decoupled, they may reside in the same physical device or in different physical devices.

vPE Architecture
Example of IP VPN vPE at the End Device

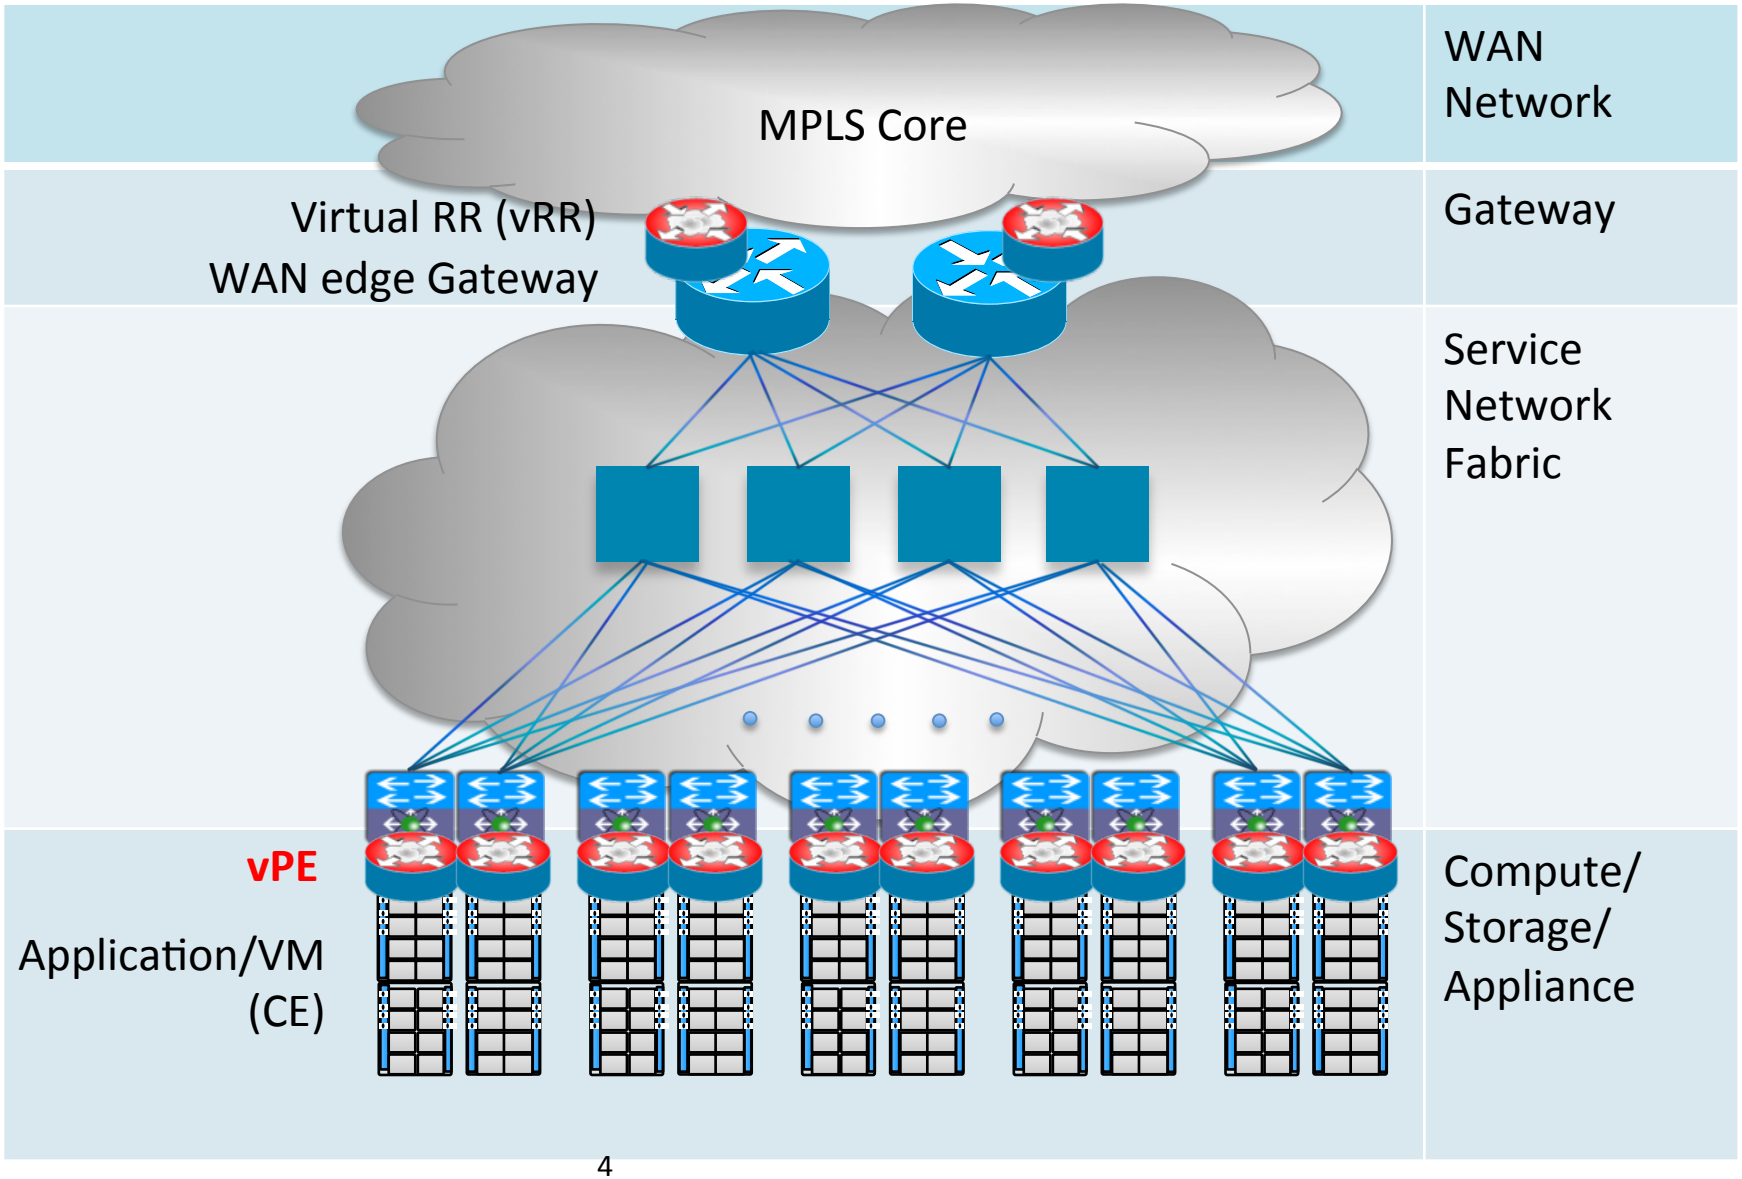

vPE on a Server Example

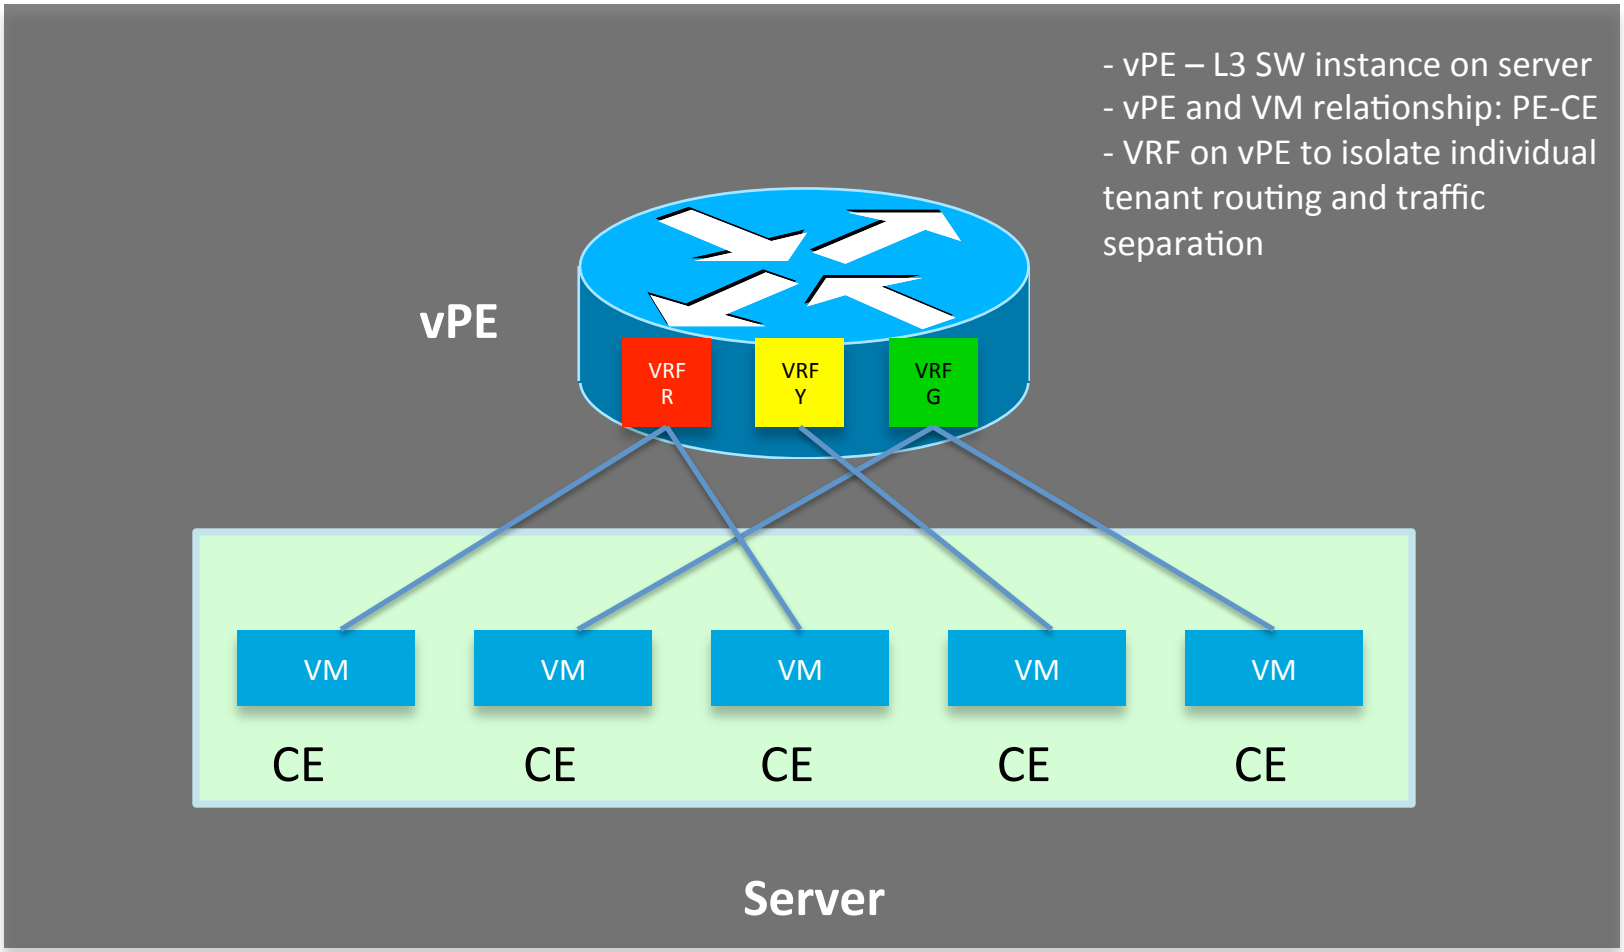

End-to-end L3VPN Overlay from Enterprise to SP DC with vPE in End Device

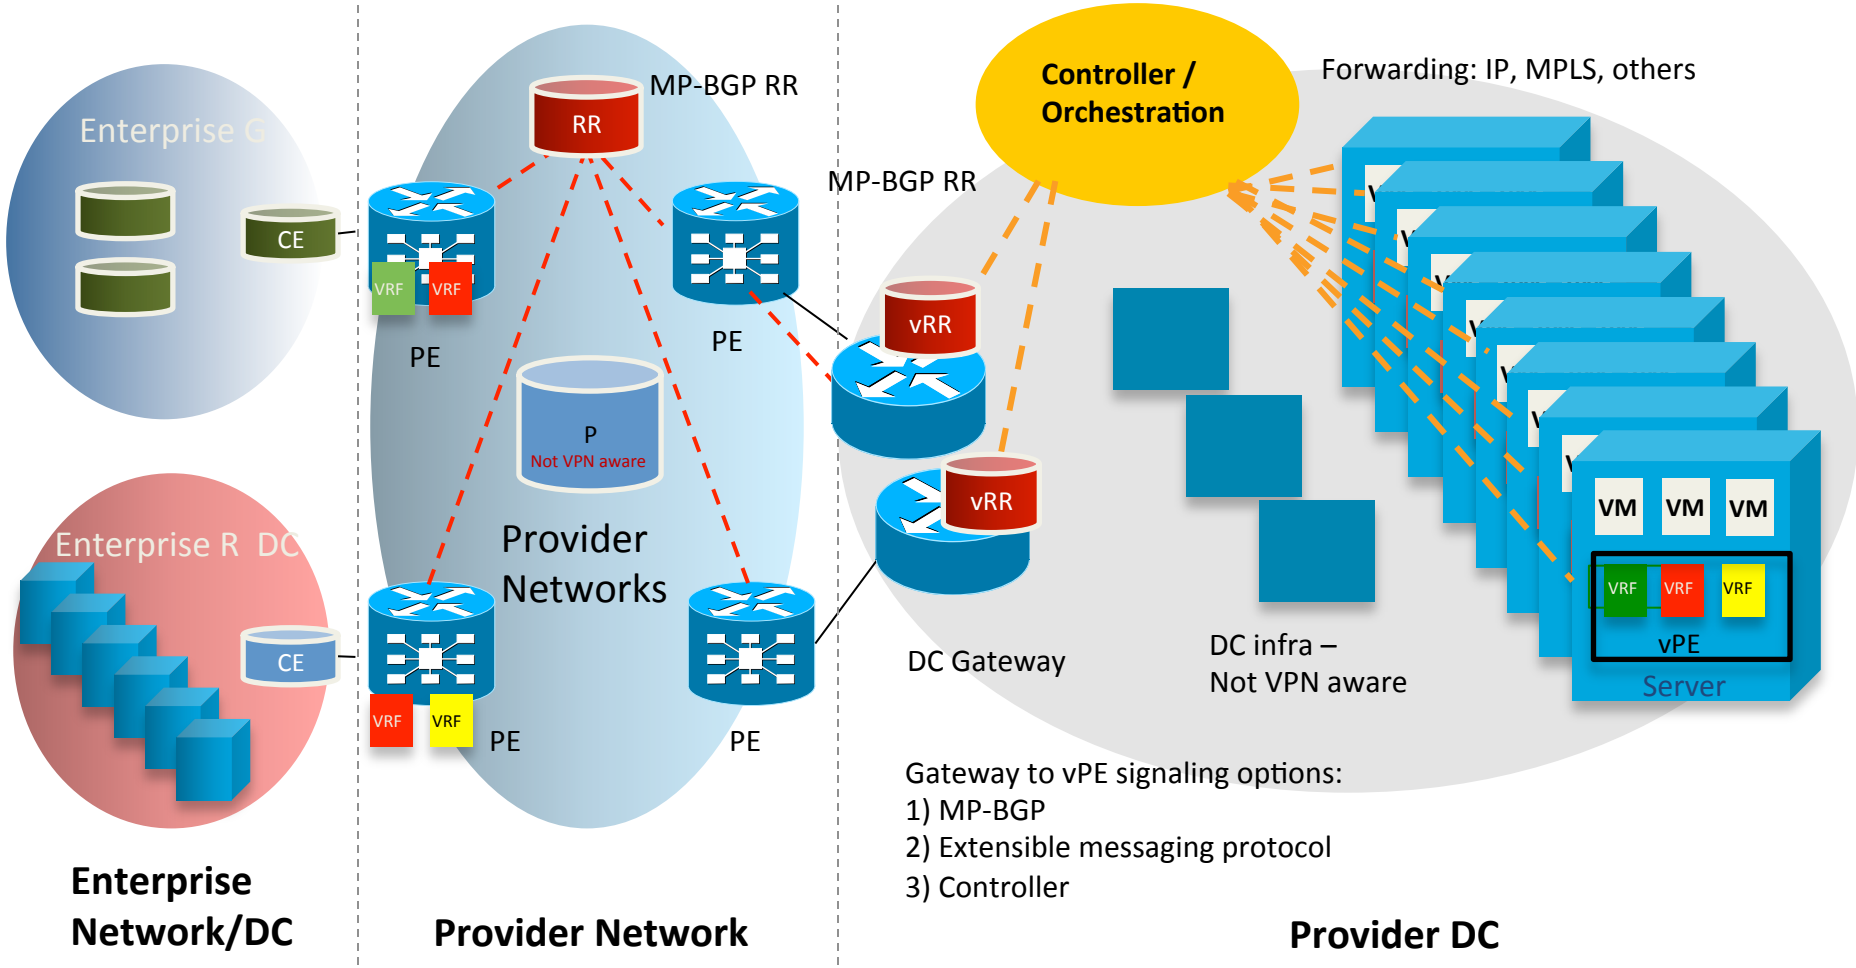

Control Plane – Route Server

1. Distributed IP VPN control plane
   - vPE participates in overlay BGP IP VPN control protocol: MP-BGP [RFC4364].
   - using extensible signaling messaging protocols can be alternative, such technologies, e.g. XMPP, [I-D.ietf-l3vpn-end-system].
2. Centralized routing controller
   - vPE control plane and data plane are physically decoupled. The control plane directing the data flow may reside elsewhere, such a centralized controller.
   - The controller can be used for routing information distribution, or directly insert the entries into FIB (details will be in next revision).
3. Route server can be anywhere
4. Use RR and RT Constrain [RFC4684] to scale.

Data Plane – VPN Forwarder

1. The VPN forwarder location options:

   1) within the end device where the CE (e.g., application/VMs) are.

   2) in an external device which the end device connect to, for example, a Top of the Rack (ToR) in a data center.

2. Considerations in design:
   - Device capability
   - Overall solution economics
   - QoS/firewall/NAT placement
   - Optimal forwarding
   - Latency and performance
   - Operational impact

3. Encapsulation
   1) MPLS
   2) IP / GRE tunnel [RFC4023]
   3) Other IP network overlay encapsulations eg. VXLAN, NVGRE.

Inter-connection considerations

- Must support connections with non-IP VPN solutions.
- A clear demarcation should be preserved at the inter-connecting points. Problems encountered in one domain should not impact other domains.
- Consider this scenario: Legacy layer 2 connectivity must be supported in certain sites/CEs/VMs (which is a small percentage of total sites), and the rest sites/CEs/VMs need only L3 connectivity.
  - One may consider to use combined vPE and CE/virtual CE solution to solved the problem. Use L3VPN vPE for all sites with IP connectivity, and use a physical or virtual CE (vCE, may reside on the end device) to aggregate the L2 sites which, for example, may be a single container in a data center. The CE/vCE can be considered as inter-connecting point, where the L2 network are terminated and the corresponding routes for connectivity of the L2 network are inserted into L3VPN VRF. The L2 aspect is transparent to the L3VPN in this case.

Next Steps

- Address all comments on the list and in the meeting
- Issuing next revision soon after the IETF 85
- Ask the WG if folks are interested to have this work as WG item