BGP L3VPN Virtual CE

draft-fang-l3vpn-virtual-ce-02

Luyuan Fang
John Evans
David Ward
Rex Fernando
John Mullooly
Ning So
Nabil Bitar
Maria Napierala

IETF 88 Vancouver, Nov. 2013

Update

- More editing since last version
- Several SPs thought it is useful draft to them
- Need to hear more feedback and move forward
- Ask the WG to check interest for adopting this work as WG item
- The following is content overview

Motivation

- Architecture re-design for virtualized DC
  - Goal: simplicity, routing/forwarding optimization, and easier service chaining.
  - A virtualized container: It includes virtual CE, virtual appliances, application VMs, as co-residents on virtualized servers.
    - virtual CE can interconnect the virtual appliances (e.g., FW, LB, NAT), applications (e.g., Web, App., and DB) in a co-located fashion.
    - Virtualizing L3-L7 on a per-tenant basis provides simplicity for managing per tenant service orchestration, tenant container creation and moves, capacity planning across tenants and per-tenant policies.
- Leverage the SP strength in l3vpn in the WAN
  - Inter-connecting through l3vpn in the WAN
  - Cloud extension for managed l3vpn services

Virtual CE Definition

- Virtual CE (vCE): a software instance of IP VPN CE function which can reside in any network or compute devices.
  - For example, a vCE may reside in an end device, such as a server in a DC, where the application VMs reside.
  - The CE functionality and management models remain the same as defined in [RFC4364].

Characteristics of vCE

- Same as a physical CE, a virtual CE supports a single tenant.
- A single tenant can use multiple physical or virtual CEs.
- An end device, such as a server, can support one or more vCE(s).
- Virtual CE and virtual PE are complimentary approaches for extending IP VPN into tenant containers.

vCE Reference Model

vCE in the end device, e.g. a VM in a server

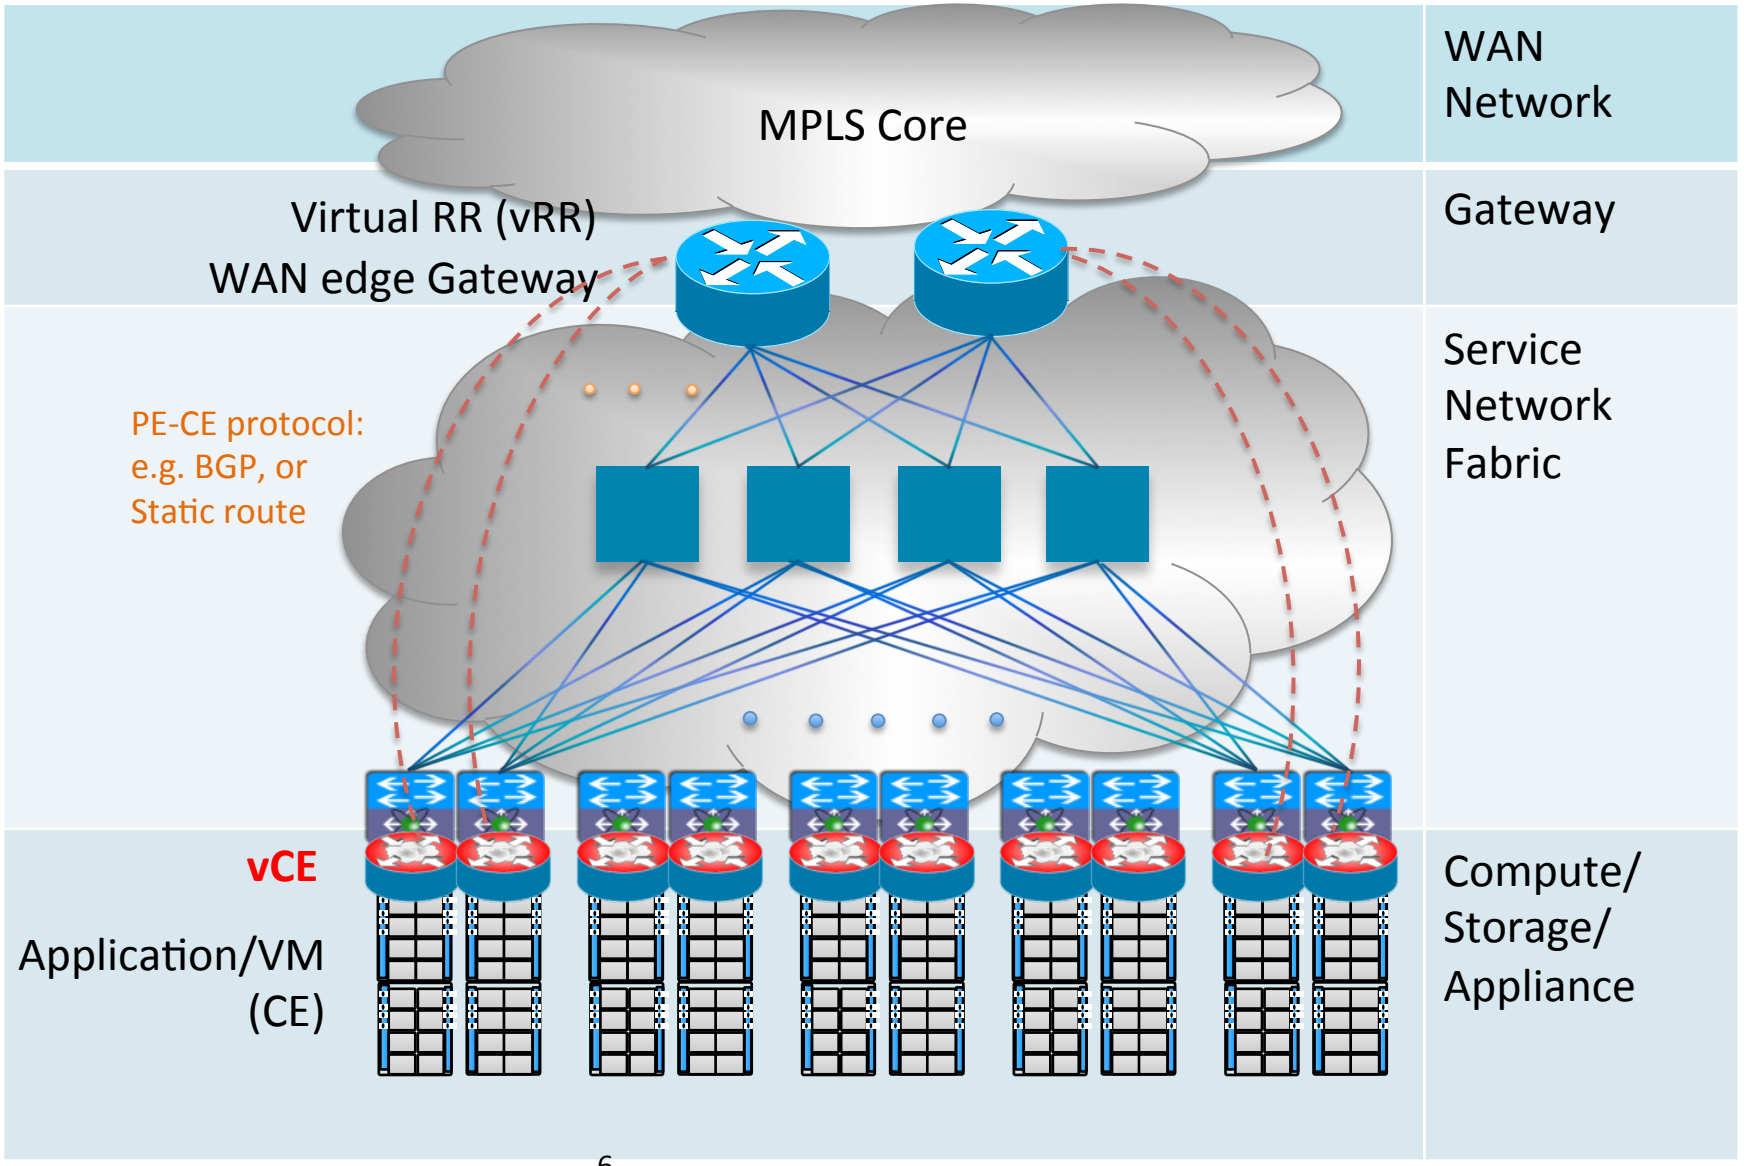

vCE Service Architecture

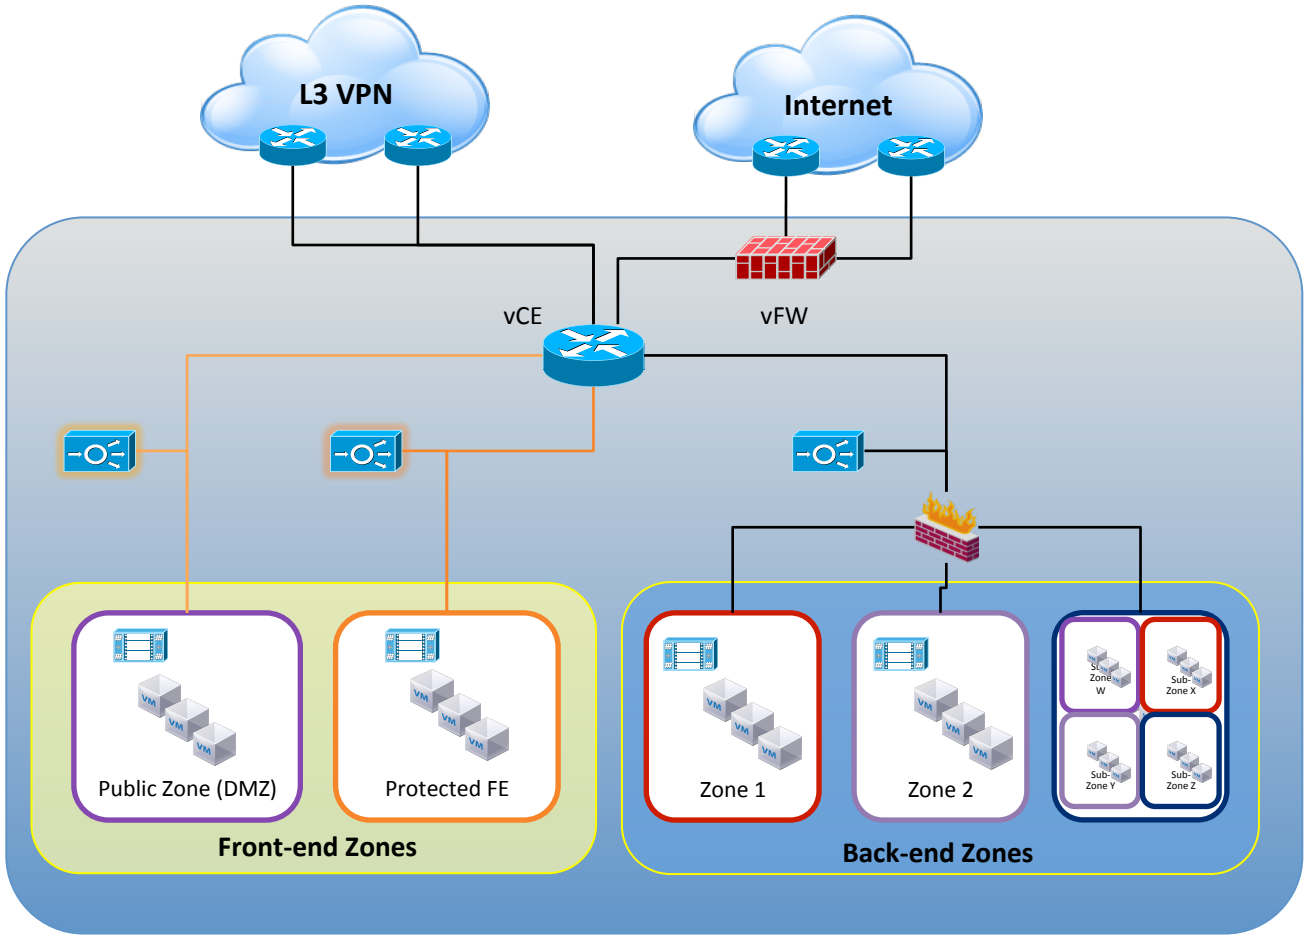

- A Virtualized Container with vCE in an End Device

Control Plane

1. Use distributed control protocol, e.g., BGP
   - BGP is policy rich, a helps to avoid single point of failure
   - But the vCE must support BGP
2. Use Static routing
   - Simple
   - But it does not provide rich policy and may have scaling issues.
3. Use Controller approach
   - MUST use standard interfaces

Data Plane

1. If the vCE and the application VM which the vCE is connecting are co-located in the same server, the connection is internal to the server, no external protocol involved.

2. If the vCE and the application VM which the vCE is connecting are located in different devices, standard external protocols are needed. The forwarding can be native or overlay techniques.

QoS

- Differentiated Services [RFC2475] Quality of Service (QoS) is standard functionality for physical CEs and MUST be supported on vCE.
- It is important to ensure seamless end-to-end SLA from IP VPN in the WAN into service network/Data center.

Management plane

- Network abstraction and management
  - vCE North bound interface SHOULD be standards based.
  - vCE element management MUST be supported, it can be in the similar fashion as for physical CE, without the hardware aspects.
- Service VM Management
  - Service VM Management SHOULD be hypervisor agnostic, e.g. On demand service VMs turning-up should be supported.
  - The management tool SHOULD be open standards.

Orchestration

DC Instance to WAN IP VPN instance "binding" Requirements

- MUST support service activation in the physical and virtual environment, assign VLAN to correct VRF.
- MUST support per VLAN Authentication, Authorization, and Accounting (AAA).
- MUST be able to apply other policies to VLAN. e.g. , per VLAN QOS, ACLs.
- MUST ensure that WAN IP VPN state and Data Center state are dynamically synchronized.
- Ensure that there is no possibility of customer being connected to the wrong VRF.
- MUST integrate with existing WAN IP VPN provisioning processes.
- MUST scale to at least 10,000 tenant service instances.
- MUST cope with rapid tenant mobility.
- MAY support Automated cross provisioning accounting correlation between WAN IP VPN and cloud/DC for the same tenant.
- MAY support Automated cross provisioning state correlation between WAN IP VPN and cloud/DC/extended Data Center for the same tenant.

vCE Push

- Process
  1. DC orchestration configures vCE
  2. Orchestration initiates WAN provisioning; passes VLAN / VXLAN + tenant context
  3. WAN provisioning system provisions PE VRF + other policies as per normal
- DC Orch or WAN provisioning needs to know the topology connecting the DC and WAN, i.e. which int on core switch connects to which int on DC PE
- Requires offline state correlation
- Requires offline accounting correlation
- Requires per SP integration

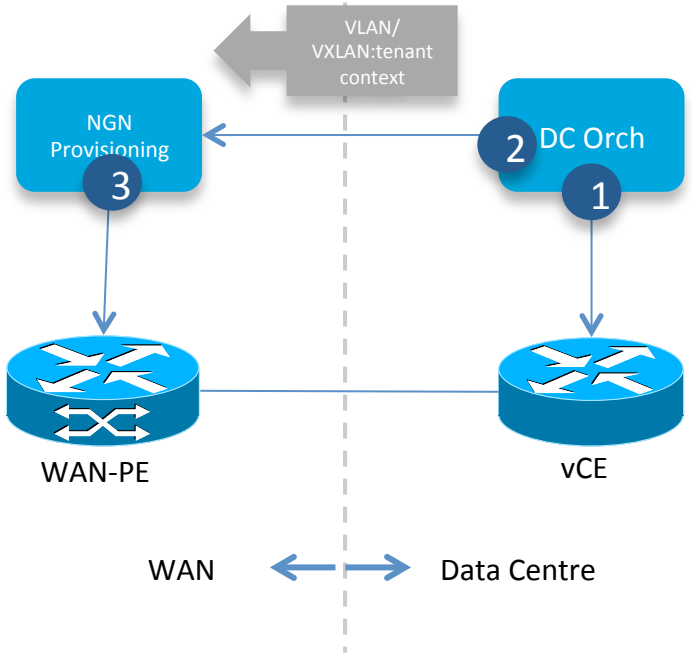

vCE Pull

- Process
  1. DC orchestration configures vCE
  2. Orchestration primes NGN provisioning/ AAA for new service, i.e. passes VLAN / VXLAN + tenant context
  3. DC PE detects new VLAN; Radius Access-Request
  4. Radius Access-Accept with VRF + other policies
- Requires VLAN/VLAN: Tenant context to passed on a per transaction basis
  - In practise may just be DC orch updating LDAP directory
- Auto state correlation
- Auto accounting correlation

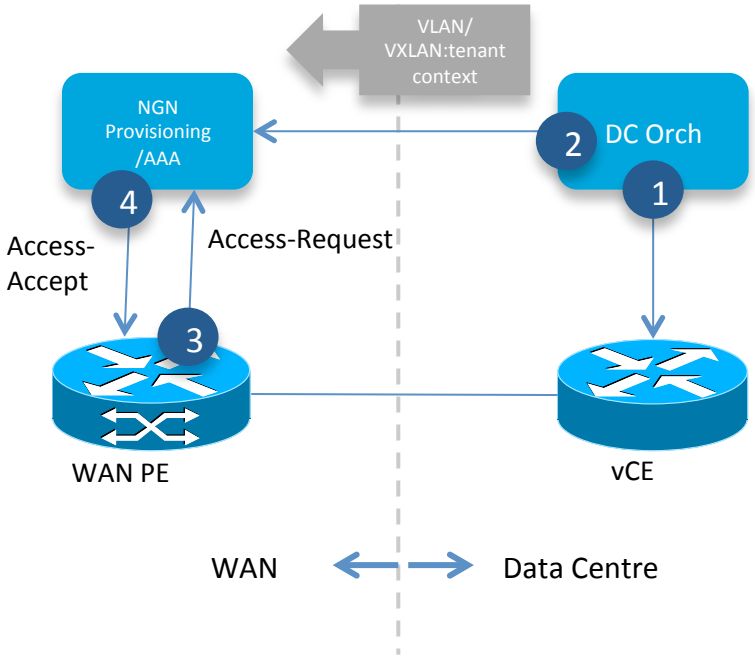

Next Steps

- Address all comments on the list, in the meeting, and off-line discussions.
- Submit a new version