

87th IETF @ Berlin

MPLS Big Label and BGP Signaling

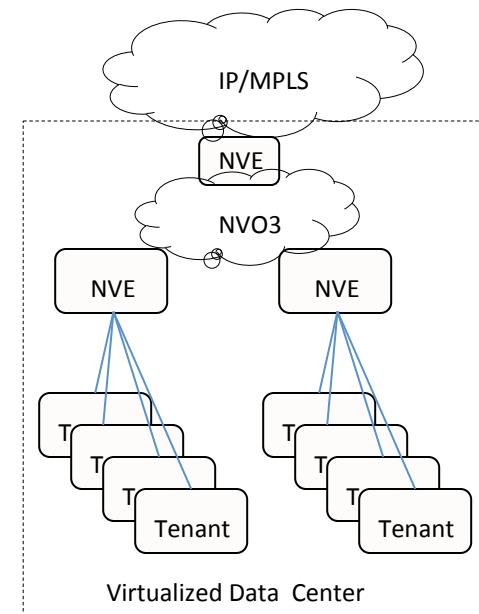

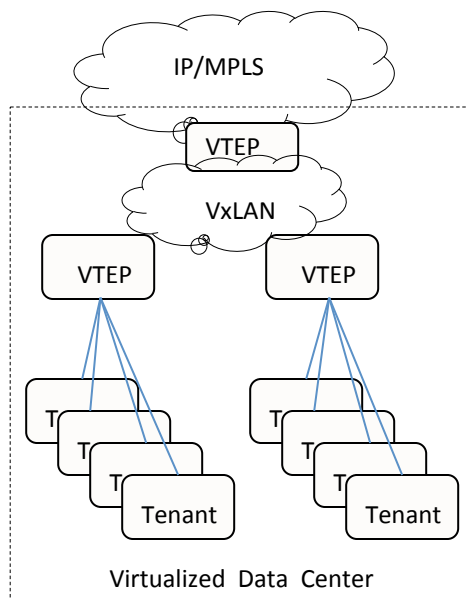
draft-renwei-mpls-big-label-00.txt

draft-renwei-mpls-bgp-big-label-00.txt

Richard Li (renwei.li@huawei.com)

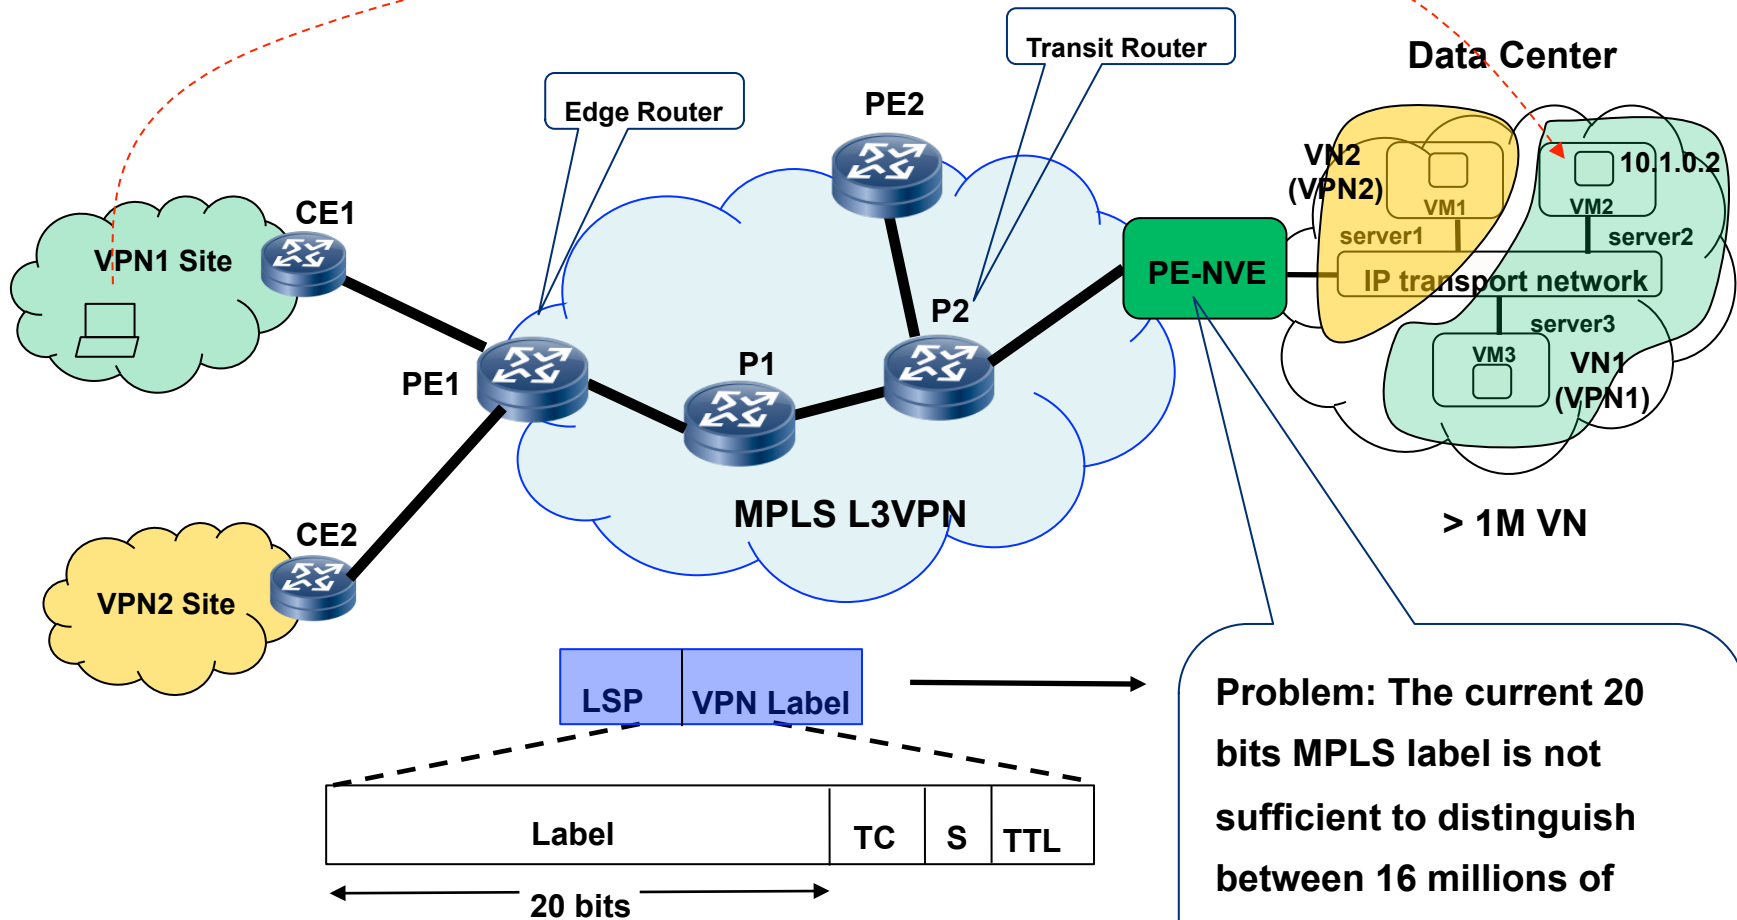
Background

- Supporting large scale network virtualization in data center has become a commonly agreed requirement
- VXLAN, NVGRE and NVO3 are currently being standardized, with all supporting up to 16 millions of virtual overlaid networks



Problem Statement

How to access VM2 in VN1 using L3VPN




NVE : Network Virtualization Edge

VN: Virtual Network


Problem: The current 20 bits MPLS label is not sufficient to distinguish between 16 millions of virtual networks hosted in a Data Center

Context-Based Solution



- Need a total of 16 contexts for 16M of VPNs
- The same VPN labels may appear in different contexts for different VPNs
- Two labels jointly serve to simulate one label
- It needs **two** labels, **two** NHLFEs, **two** table lookups
 - Implementation complexity in both hardware and software: a hierarchical or embedded data structure is in need.
 - Slow down the data plane because of two label lookups and the associated actions.
- In disagreement with old L3VPN (old L3VPN requires one VPN label, and at least three commercial implementations assumes so)
- Association between contexts and VPN labels are arbitrary: one can swap VPN labels and their associated NHLFEs between different contexts
- BGP needs to distribute both contexts and VPN labels, and thus adversely affect the re-convergence time.
- A 20 bits value is generally represented by a 32 bits integer, and thus 64 bits are needed in the VPN ILM for a single VPN (remember we are talking about 16M VPNs)


Proposed Solution: Big Labels



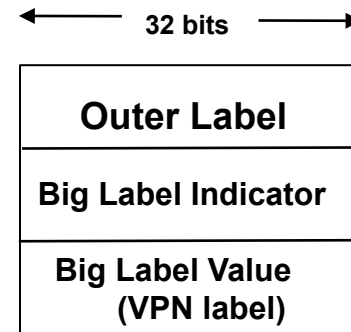
- They are viewed as one and only one label
- No need for two NHLFEs and two table lookups.
- Compliant with the BGP/MPLS L3VPN architecture: the BGP/MPLS L3VPN uses ONE label to represent the VPN/VRF but doesn't use TWO labels for that purpose.
- BGP just needs to distribute the Big Label Value instead of two labels needed in the context-based solution
- A 20 bits value is generally represented by a 32 bits integer. And thus the above format doesn't add complexity to the data structures in software (I am aware that three commercial operating systems represent the 20 bits label with 32 bits integer)

Big Label in MPLS IP VPN

- The outer label is unchanged
- The inner label (VPN label) can be either the common 20 bits label or the current proposed big label
- The Big Label Format is compatible with the MPLS Architecture and L3VPN model, but is slightly different from the current one with respect to the de-capsulation of the VPN label.



Two Entries with Common Labels



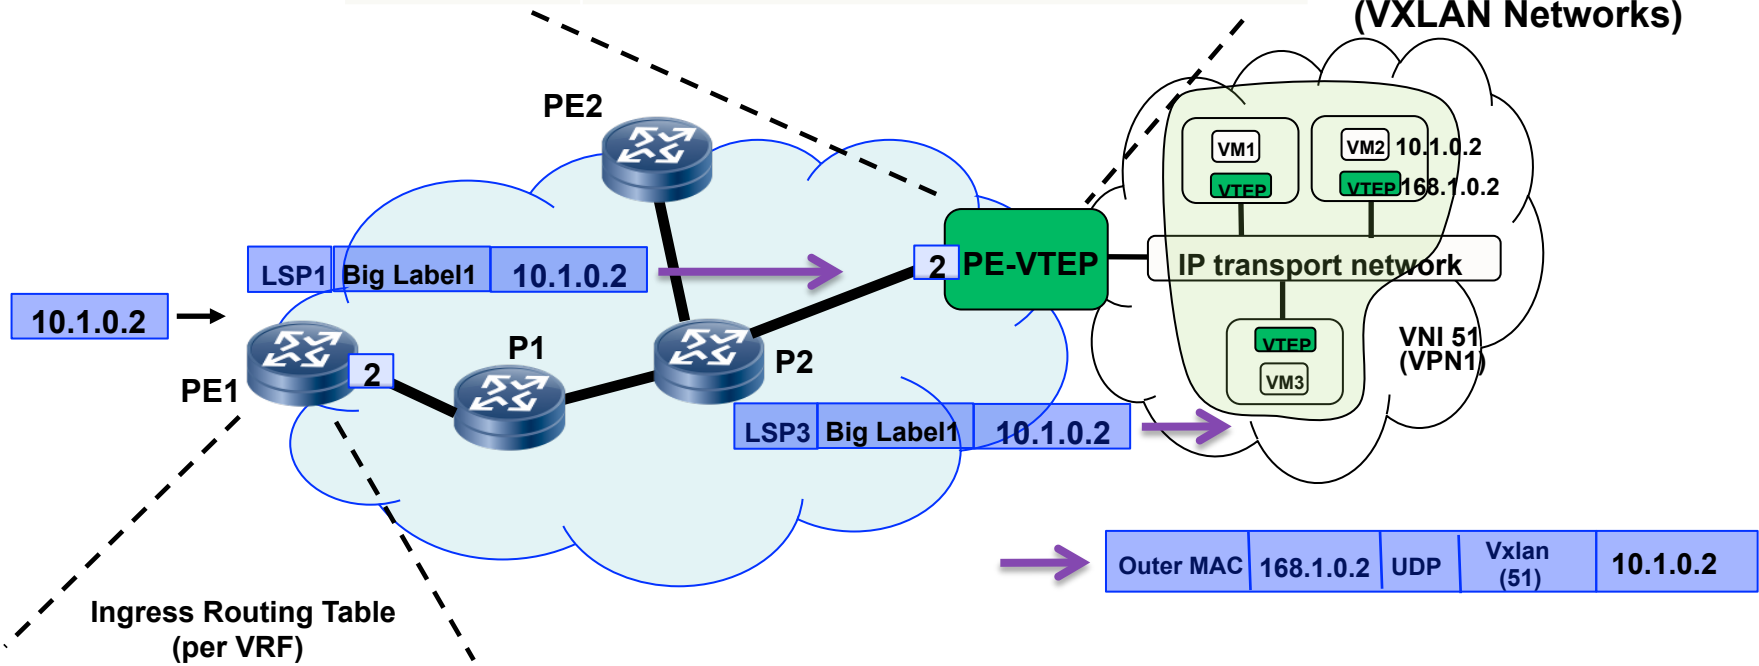
Two Entries with Big Label

Packet Life: From Customer to DC

Egress PE-VTEP FW Table

Destination	Out
10.1.0.0/24	(168.1.0.2, VTEP_MAC1, VNI1(51), VM_MAC1)

Data Center
(VXLAN Networks)




Destination	Next Hop
10.1.0.0/24	(2, LSP1, Big Label1)


Outer MAC	168.1.0.2	UDP	Vxlan (51)	10.1.0.2
-----------	-----------	-----	------------	----------

Extension of MP-BGP

- **New SAFI is introduced for big label signaling**
- **BGP capability advertisement for new SAFI if MPLS big label is used for NLRI**
- **New NLRI Format**



Current NLRI for Common Label



New NLRI for Big Label

Next Step

- **Seeking comments from you**
- **Experimenting and prototyping**
- **Exploring and comparing with other solutions**