

INTERNET-DRAFT
Intended Status: Standards track
Expires: January 4, 2015

Luyuan Fang
Microsoft
John Evans
David Ward
Rex Fernando
Cisco
Ning So
Vinci Systems
Nabil Bitar
Verizon
Maria Napierala
AT&T

July 4, 2014

**BGP IP MPLS VPN Virtual CE
draft-fang-l3vpn-virtual-ce-03**

Abstract

This document describes the architecture and solutions of using virtual Customer Edge (vCE) of BGP IP MPLS VPN. The solution is aimed at providing efficient service delivery capability through CE virtualization, and is especially beneficial in virtual Private Cloud (vPC) environments when extending IP MPLS VPN into tenant virtual Data Center containers. This document includes: BGP IP MPLS VPN virtual CE architecture; Control plane and forwarding options; Data Center orchestration processes; integration with existing WAN enterprise VPNs; management capability requirements; and security considerations. The solution is generally applicable to any BGP IP VPN deployment. The virtual CE solution is complementary to the virtual PE solutions.

Today's data center's require multi-tenancy and mechanisms to establish overlay network connectivity. This document describes one approach to enabling data center network connectivity.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#) [4](#)
- [1.1 Terminology](#) [4](#)
- [1.2 Problem statement](#) [6](#)
- [1.3 Scope of the document](#) [6](#)
- [2. Virtual CE Architecture and Reference Model](#) [7](#)
- [2.1 Virtual CE](#) [7](#)
- [2.2 Architecture](#) [8](#)
- [3. Control Plane](#) [10](#)
- [3.1 vCE Control Plane](#) [10](#)
- [4. Forwarding Plane](#) [11](#)
- [4.1 Forwarding between vCE and PE/vPE](#) [11](#)
- [4.2 Forwarding between vCE and VM](#) [11](#)
- [5. Addressing and QoS](#) [11](#)
- [5.1 Addressing](#) [11](#)
- [5.2 QoS](#) [12](#)
- [6. Management plane](#) [12](#)
- [6.1 Network abstraction and management](#) [12](#)
- [6.2 Service VM Management](#) [12](#)

Expires <January 4, 2015>

[Page 2]

- [7. Orchestration and IP VPN inter-provisioning](#) [12](#)
- 7.1 DC Instance to WAN IP VPN instance "binding" Requirements . 12
- [7.2. Provisioning/Orchestration](#) [13](#)
- [7.2.1 vCE Push model](#) [13](#)
- [7.2.1.1 Inter-domain provisioning vCE Push Model](#) [14](#)
- [7.2.1.2 Cross-domain provisioning vCE Push Model](#) [14](#)
- [7.1.1 vCE Pull model](#) [15](#)
- 8. vCE and vPE interaction [16](#)
- [8.1 Traditional vCE-PE connectivity](#) [16](#)
- [8.2 vCE-vPE connectivity](#) [16](#)
- [8.2.1 Co-located vCE-vPE connectivity with vPE Model 1](#) [17](#)
- [8.2.2 Co-located vCE-vPE connectivity with vPE Model 2](#) [18](#)
- 8. Security Considerations [18](#)
- 9. IANA Considerations [18](#)
- 10. References [18](#)
- [10.1 Normative References](#) [18](#)
- [10.2 Informative References](#) [19](#)
- 11. Acknowledgement [20](#)
- Authors' Addresses [20](#)

Expires <January 4, 2015>

[Page 3]

1. Introduction

In the typical enterprise BGP/MPLS IP VPN [RFC4364] deployment, the Provider Edge (PE) and Customer Edge (CE) are physical routers which support the PE and CE functions. With the recent development of cloud services, using virtual instances of PE or CE functions, which reside in a compute device such as a server, can be beneficial to emulate the same logical functions as the physical deployment model but now achieved via cloud based network virtualization principles.

This document describes IP VPN virtual CE (vCE) solutions, while Virtual PE (vPE) concept and implementation options are discussed in [I-D.fang-l3vpn-virtual-pe], [I-D.ietf-l3vpn-end-system]. vPE and vCE solutions provide two avenues to realize network virtualization.

1.1 Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table with 2 columns: Term and Definition. Terms include AAA, ACL, 3GPP, AS, ASBR, BFD, BGP, CE, DB, DMZ, ED, FE, FIB, Forwarder, FRR, FTP, GRE, HTTP, Hypervisor, I2RS, LDAP, and MP-BGP.

Expires <January 4, 2015>

[Page 4]

NVGRE	Network Virtualization using GRE
OSPF	Open Shortest Path First
PE	Provider Edge
QinQ	Provider Bridging, stacked VLANs
RR	Route Reflector
SDN	Software Defined Network
SLA	Service Level Agreement
SMTP	Simple Mail Transfer Protocol
ToR	Top of the Rack switch
VI	Virtual Interface
vCE	virtual Customer Edge Router
vLB	virtual Load Balancer
VM	Virtual Machine
VLAN	Virtual Local Area Network
vPC	virtual Private Cloud
vPE	virtual Provider Edge Router
VPN	Virtual Private Network
vRR	virtual Route Reflector
vSG	virtual Security Gateway
VXLAN	Virtual eXtensible Local Area Network
WAN	Wide Area Network

Definitions:

Virtual CE (vCE): A virtual instance of the Customer Edge (CE) routing function which resides in one or more network or compute devices. For example, the vCE data plane may reside in an end device, such as a server, and as co-resident with application Virtual Machines (VMs) on the server; the vCE control plane may reside in the same device or in a separate entity such as a controller.

Network Container/Tenant Container: An abstraction of a set of network and compute resources which can be physical and virtual, providing the cloud services for a tenant. One tenant can have more than one Tenant Containers.

Zone: A logical grouping of VMs and service assets within a tenant container. Different security policies may be applied within and between zones.

DMZ: Demilitarized zone, a.k.a. perimeter networking. It is often a machine or a small subnet that sits between a trusted internal network, such as a corporate private LAN, and an un-trusted external network, such as the public Internet. Typically, the DMZ contains devices accessible to Internet traffic, such as Web (HTTP) servers, FTP servers, SMTP (e-mail) servers and DNS servers.

Expires <January 4, 2015>

[Page 5]

1.2 Problem statement

With the growth of cloud services and the increase in the number of CE devices, routers/switches, and appliances, such as Firewalls (FWs) and Load Balancers (LBs), that need to be supported, there are benefits to virtualize the Data Center tenant container. The virtualized container can increase resource sharing, optimize routing and forwarding of inter-segment and inter-service traffic, and simplify design, provisioning, and management.

The following two aspects of the virtualized Data Center tenant container for the IP VPN CE solution are discussed in this document.

1. Architecture re-design for virtualized DC.

The optimal architecture of the virtualized container includes virtual CE, virtual appliances, application VMs. All these functions are co-residents on virtualized servers. In this arrangement, CEs and appliances can be created and removed easily on demand, and the virtual CE can interconnect the virtual appliances (e.g., FW, LB, NAT), applications (e.g., Web, App., and DB) in a co-located fashion for simplicity, routing/forwarding optimization, and easier service chaining. Virtualizing these functions on a per-tenant basis provides simplicity for the network operator in regards to managing per tenant service orchestration, tenant container moves, capacity planning across tenants and per-tenant policies.

2. Provisioning/orchestration. Two issues need to be addressed:

a) The provisioning/orchestration system of the virtualized data center need to support VM life cycle and VM migration.

b) The provisioning/orchestration systems of the DC and the WAN networks need to be coordinated to support end-to-end IP VPN from DC to DC or from DC to enterprise remote office in the same VPN. The DC and the WAN network are often operated by separate departments, even if they belong to the same provider. Today, the process of inter-connecting is slow and painful, and automation is highly desirable.

1.3 Scope of the document

It is assumed that the readers are familiar with BGP/MPLS IP VPN [[RFC4364](#)] terms and technologies, the base technology and its operation are not reviewed in details in this document.

As the majority (all in some networks) of applications are IP, this vCE solution is focusing on IP VPN solutions to cover the most common cases and keep matters as simple as possible.

Expires <January 4, 2015>

[Page 6]

2. Virtual CE Architecture and Reference Model

2.1 Virtual CE

As described in [[RFC4364](#)], IP uses a "peer model" - the customers' edge routers (CE routers) exchange routes with the Service Provider's edge routers (PE routers); the CEs do not peer with each other. MP-BGP [[RFC4271](#), [RFC4760](#)] is used between the PEs (often with RRs) which have a particular VPN attached to them to exchange the VPN routes. A CE sends IP packets to the PE; no VPN labels for packets forwarded between CE and PE.

A virtual CE (vCE) as defined in this document is 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 Data Center, where the application VMs reside. The CE functionality and management models remain the same as defined in [[RFC4364](#)] regardless of whether the CE is physical or virtual.

Using the virtual CE model, the CE functions CAN easily co-located with the VM/applications, e.g., in the same server. This allows tenant inter-segment and inter-service routing to be optimized. Likewise the vCE can be in a separate server (in the same DC rack or across racks) than the application VMs, in which case VMs would typically use standard L2 technologies to access the vCE via the DC network.

Similar to the virtual PE solution, the control and forwarding of a virtual CE can be on the same device, or decoupled and reside on different physical devices. The provisioning of a virtual CE, associated applications, and the tenant network container can be supported through DC orchestration systems.

Unlike a physical or virtual PE which can support multi-tenants, a physical or virtual CE supports a single tenant only. A single tenant CAN use multiple physical or virtual CEs. An end device, such as a server, CAN support one or more vCE(s). While the vCE is defined as a single tenant device, each tenant can have multiple logical departments which are under the tenants administrative control, requiring logical separation, this is the same model as today's physical CE deployments.

Virtual CE and virtual PE are complimentary approaches for extending IP VPN into tenant containers. In the vCE solution, there is no IP VPN within the data center or other type of service network, the vCE can connect to the PE which is a centralized IP VPN PE/Gateway/ASBR, or connect to distributed vPE on a server or on the Top of the Rack switch (ToR). Virtual CE can be used to extend the SP managed CE

Expires <January 4, 2015>

[Page 7]

solution to create new cloud enabled services and provide the same topological model and features that are consistent with the physical CE systems.

2.2 Architecture

Figure 1 illustrates the topology where vCE is resident in the servers where the applications are hosted.

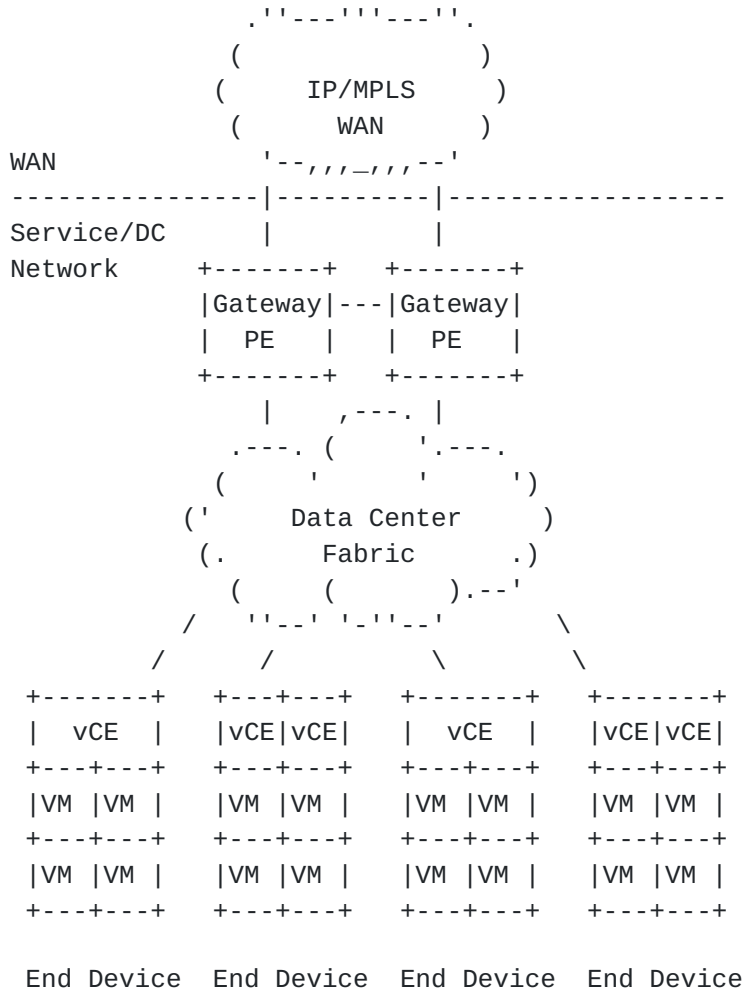


Figure 1. Virtualized Data Center with vCE

Figure 1 shows above vCE solution in a virtualized Data Center with application VMs on the servers. One or more vCEs MAY be used on each server.

The vCEs logically connect to the PEs/Gateway PEs to join the particular IP VPN which the tenant belongs to. Gateway PEs connect to the IP MPLS WAN network for inter-DC and DC to enterprise VPN sites

Expires <January 4, 2015>

[Page 8]

connection. The server physically connects to the DC Fabric for packet forwarding.

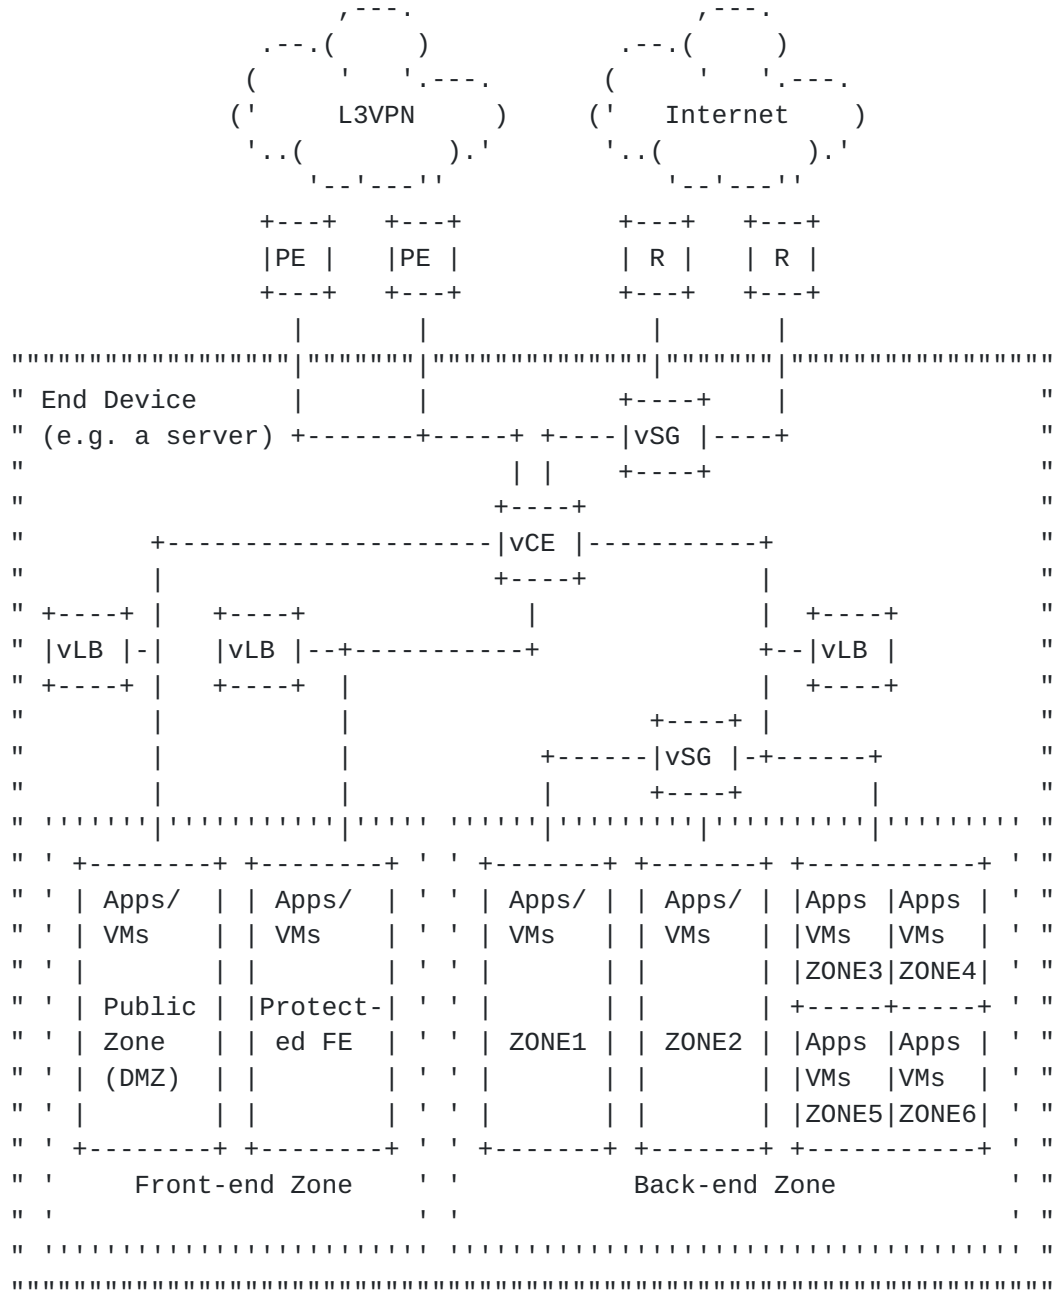


Figure 2. A Virtualized Container with vCE in an End Device

An end device shown in Figure 2 is a physical server supporting multiple virtualized appliances and application, and hosts multiple client VMs. An end device shown in Figure 2 is a physical server supporting multiple In the traditional deployment, the topology often involves multiple physical CEs, physical Security Gateways and Load

Expires <January 4, 2015>

[Page 9]

Balancers residing in the same Data Center.

The virtualized approach provides the benefit of reduced number of physical devices, simplified management, optimal routing due to the co-location of vCE, services, and client VMs.

While the above diagram represents a simplified view of all of the tenant service and application VMs residing in the same physical server, the above model can also be represented with the VMs spread across many physical servers and the DC network would provide the physical inter-connectivity while the vCE and the VMs connected to the vCE form the logical connections.

3. Control Plane

3.1 vCE Control Plane

The vCE control plane can be distributed or centralized.

1) Distributed control plane

vCE CAN exchange BGP routes with PE or vPE for the particular IP VPN as described in [[RFC4364](#)].

The vCE needs to support BGP if this approach is used.

The advantage of distributed protocols is to avoid single point of failure and bottleneck. Service chaining can be easily and efficiently supported in this approach.

BGP as PE-CE protocol is used in about 70% of cases in typical Enterprise IP VPN PE-CE connections. BGP supports rich policy compared to other alternatives.

2) Static routing. It is used in about 30% of cases in Enterprise IP VPN PE-CE connections. It MAY be used if the operator prefers.

2. Using controller approach

Using controller is the Software Defined Network (SDN) approach. A controller can be distributed or centralized. The central controller performs the control plane functions, and sends instructions to the vCE on the end devices to configure the data plane.

This requires standard interface to routing system (I2RS). The Interface to Routing System (I2RS) is work in progress in IETF [[I-D.ward-irs-framework](#)], [[I-D.rfernando-irs-framework-requirement](#)].

Expires <January 4, 2015>

[Page 10]

4. Forwarding Plane

4.1 Forwarding between vCE and PE/vPE

No MPLS forwarding is required between PE and CE in typical PE-CE connection scenarios, though MPLS label forwarding is required for implementing Carriers' Carrier (CSC) model.

IPv4 and IPv6 packet forwarding MUST be supported.

Native fabric CAN be used to support isolation between vCEs to PE connections.

Examples of native fabric include:

- VLANs [IEEE 802.1Q], Virtual Local Area Network- IEEE 802.1ad [IEEE 802.1ad]/QinQ, Provider Bridge

Or overlay segmentation with better scalability:

- VXLANs [I-D.mahalingam-dutt-dcops-vxlan], Virtual Extensible LAN- NVGRE [I-D.sridharan-virtualization-nvgre], Network Virtualization using GRE

Note the the above references for overlay network are currently work in progress in IETF.

4.2 Forwarding between vCE and VM

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

If the vCE and the VM the vCE is connecting are located in different devices, standard external protocols are needed. The forwarding can be native or overlay techniques as listed in the above sub-section.

5. Addressing and QoS

5.1 Addressing

IPv4 and IPv6 addressing MUST be supported.

IP address allocation for vCEs and applications/client:

- 1) IP address MAY be assigned by central management/provisioning with predetermined blocks through planning process.

Expires <January 4, 2015>

[Page 11]

2) IP address MAY be obtained through DHCP server.

Address space separation: The IP addresses used for clients in the IP VPNs in the Data Center SHOULD be in separate address blocks outside the blocks used for the underlay infrastructure of the Data Center. The purpose is to protect the Data Center infrastructure from being attacked if the attacker gain access of the tenant VPNs.

5.2 QoS

Differentiated Services [[RFC2475](#)] Quality of Service (QoS) is standard functionality for physical CEs and MUST be supported on vCE. This is important to ensure seamless end-to-end SLA from IP VPN in the WAN into service network/Data center. The use of MPLS Diffserv tunnel model Pipe Mode ([RFC3270](#)) with explicit null LSP must be supported.

6. Management plane

6.1 Network abstraction and management

The use of vCE with single tenant virtual service instances can simplify management requirements as there is no need to discover device capabilities, track tenant dependencies and manage service resources.

vCE North bound interface SHOULD be standards based.

The Interface to Routing System (I2RS) is work in progress in IETF [[I-D.ward-irs-framework](#)], [[I-D.rfernando-irs-framework-requirement](#)].

vCE element management MUST be supported, it can be in the similar fashion as for physical CE, without the hardware aspects.

6.2 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.

7. Orchestration and IP VPN inter-provisioning

7.1 DC Instance to WAN IP VPN instance "binding" Requirements

- MUST support service activation in the physical and virtual environment.

Expires <January 4, 2015>

[Page 12]

For example, assign VLAN to correct VRF.

- MUST support per VLAN Authentication, Authorization, and Accounting (AAA).

The PE function is an OA&M boundary.

- MUST be able to apply other policies to VLAN.

For example, per VLAN QOS, ACLs.

- MUST ensure that WAN IP VPN state and Data cCentre state are dynamically synchronized.

Ensure that there is no possibility of customer being connected to the wrong VRF. For example, remove all tenant state when service instance terminated.

- MUST integrate with existing WAN IP VPN provisioning processes.
- MUST scale to at least 10,000 tenant service instances.
- MUST cope with rapid (sub minute) 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.

7.2. Provisioning/Orchestration

There are two primary approaches for IP VPN provisioning - push and pull, both CAN be used for provisioning/orchestration.

7.2.1 vCE Push model

Push model: It is a top down approach - push IP VPN provisioning from network management system or other central control provisioning systems to the IP VPN network elements.

This approach supports service activation and it is commonly used in the existing IP VPN enterprise deployment. When existing the IP VPN solution into the cloud/data center or separate Data Center, it MUST support off-line accounting correlation between the WAN IP VPN and the cloud/DC IP VPN for the tenant, the systems SHOULD be able to bind interface accounting to particular tenant. It MAY requires

Expires <January 4, 2015>

[Page 13]

offline state correlation as well, for example, bind interface state to tenant.

7.2.1.1 Inter-domain provisioning vCE Push Model

Provisioning process:

- 1) Cloud/DC orchestration configures vCE.
- 2) Orchestration initiates WAN IP VPN provisioning; passes connection IDs (e.g., of VLAN/VXLAN) and tenant context to WAN IP VPN provisioning systems.
- 3) WAN IP VPN provisioning system provisions PE VRF and other policies per normal enterprise IP VPN provisioning processes.

This model requires the following:

- The DC Orchestration system or the WAN IP VPN provisioning system know the topology inter-connecting the DC and WAN VPN. For example, which interface on the WAN core device connects to which interface on the DC PE.
- Offline state correlation.
- Offline accounting correlation.
- Per SP integration.

Dynamic BGP session between PE/vPE and vCE MAY be used to automate the PE provisioning in the PE-vCE model, that will remove the needs for PE configuration. Other protocols can be used for this purpose as well, for example, use Enhanced Interior Gateway Routing Protocol (EIGRP) for dynamic neighbour relationship establishment.

The dynamic routing Prevents the need to configure the PEs in PE-vCE model.

Caution: This is only under the assumption that the DC provisioning system is trusted and could support dynamic establishment of PE-vCE BGP neighbor relationships, for example, the WAN network and the cloud/DC belongs to the same Service Provider.

7.2.1.2 Cross-domain provisioning vCE Push Model

Provisioning Process:

- 1) Cross-domain orchestration system initiates DC orch.

Expires <January 4, 2015>

[Page 14]

- 2) DC orchestration system configures vCE
- 3) DC orchestration system passes back VLAN/VXLAN and tenant context to Cross-domain orchestration system
- 4) Cross-domain orchestration system initiates WAN IP VPN provisioning
- 5) WAN IP VPN provisioning system provisions PE VRF and other policies as per normal enterprise IP VPN provisioning processes.

This model requires the following:

- Cross-domain orchestration system knows the topology connecting the DC and WAN IP VPN, for example, which interface on core device connects to which interface on DC PE.- Offline state correlation.
- Offline accounting correlation.
- Per SP integration.

7.1.1 vCE Pull model

Pull model: It is a bottom-up approach - pull from network elements to network management/AAA based upon data plane or control plane activity. It supports service activation, this approach is often used in broadband deployment. Dynamic accounting correlation and dynamic state correlation are supported. For example, session based accounting is implicitly includes tenant context state correlation, as well as session based state which implicitly includes tenant context.

Inter-domain Provisioning:

Process:

- 1) Cloud/DC orchestration system configures vCE
- 2) Cloud/DC Orchestration system primes WAN IP VPN provisioning/AAA for new service, passes connection IDs (e.g., VLAN/VXLAN) and tenant context WAN IP VPN provisioning systems.
- 3) Cloud/DC PE detects new VLAN, send Radius Access-Request.
- 4) Radius Access-Accept with VRF and other policies.

This model requires VLAN/VLAN information and tenant context to be passed on a per transaction basis. In practice, it may simplify to

Expires <January 4, 2015>

[Page 15]

use DC orchestration updating LDAP directory

Auto accounting correlation and auto state correlation is supported.

8. vCE and vPE interaction

A vPE ([[I-D.fang-13vpn-virtual-pe](#)] [[I-D.ietf-13vpn-end-system](#)]) is treating the VMs in the server as a virtual CE. In this section, the relationship between the vPE and such vCE is discussed. vPE can support one of the following two models:

Model 1: a limited control-plane functionality that advertises local VPN routes to a controller and receive VPN routes from the controller.

Model 2: a control plane component physically separated from the forwarding component that fully performs the control plane routing functionality and communicate FIB entries to the vPE forwarding entity implemented on servers.

A vCE provides subnet routing, firewalling or SLB services to host VMs. The underlying connectivity between the vCE and these VMs can be at layer 2 or layer 3. In addition, the vCE can be connected to other vCEs over Layer 2 or using an IP VPN infrastructure. In this section, the focus is on IP VPN connectivity and more importantly on the interaction between a vCE, a traditional PE (simply referred to as PE), and between a vCE and a vPE.

8.1 Traditional vCE-PE connectivity

This connectivity is described in BGP/MPLS IPVPN [[RFC4364](#)]. The only distinction being that the VE is a virtual CE. The vCE attaches to the layer 3 PE using a layer2 logical connection, e.g., Ethernet VLAN, or a tunnel (e.g., IP/GRE, VXLAN) that are presented as IP interfaces to a corresponding VRF at the PE. Routing between the vCE and PE can be static or based on a dynamic routing protocol (e.g., OSPF, BGP). A routing protocol, in addition to enabling the exchange of routing information between the PE and vCE, provides liveliness check between the vCE and the PE. In the absence of a dynamic routing protocol, the vCE must support a mechanism that provides for liveliness check, or an out-of-band mechanism must be implemented to monitor the liveliness of a vCE and a connected PE, and effect routing changes upon a failure. Options for in-band liveliness check include IP BFD [[RFC5880](#)], Ethernet Continuity Check (CC) [IEEE 802.1ag], and IP ping [[RFC4560](#)]. IP BFD must be supported while the other mechanisms are optional.

8.2 vCE-vPE connectivity

Expires <January 4, 2015>

[Page 16]

In this model, the vCE and vPE forwarding plane can be: (1) co-located on the same end device, e.g., a server, or (2) located on different servers. In addition, the control plane interaction differs between vPE model 1 and model 2.

8.2.1 Co-located vCE-vPE connectivity with vPE Model 1

In vPE Model 1, there is a control plane component of the vPE implemented on the end-server (e.g., [[I-D.ietf-l3vpn-end-system](#)], [[I-D.fang-l3vpn-virtual-pe](#)]). In addition, there is a control plane component implemented on a separate control plane entity (out-of-band) that enables the exchange of routing information among vPEs. In [[I-D.ietf-l3vpn-end-system](#)], the out-of-band control plane component is referred to as router server; in [[I-D.fang-l3vpn-virtual-pe](#)], it is referred as vPE-C. There are two cases that must be considered:

Case 1-A: vCE to vPE local route exchange on a server / vPE-C

Case 1-B: vCE to route server / vPE-C route exchange.

In these two cases, the vPE control plane or route server must send the CE a default route with next hop being the co-located vPE forwarding plane entity.

In case 1-A, the vCE must send local routes to the vPE control plane with itself being the next hop. The vPE control plane entity in turn updates the out-of-band control entity (e.g., route server) with routes reachable via the local CE, as VPN routes, with itself being the next hop for these routes. The vPE also receives from the route server VPN routes reachable via other vPEs [[end-system](#)]. It should be noted in this case, that the vCE must be able support one or more routing contexts, each with separate attachment circuit to the vPE. Each such routing context must be associated with a VPN and one or more VPNs must be supported.

In case 1-B, the vCE must have a control channel with a route server. There must be a control channel per vCE routing context or alternatively must allow the unambiguous multiplexing of routes that belong to different routing context on the same channel. The vCE sends routes reachable via the vCE to the route server with itself being the next hop. The route server must learn from the co-located vPE control plane component reachability to the local vCE IP address used as next hop. This IP address must be exchanged between the vCE and vPE in-band over a corresponding attachment circuit that identifies the routing context . Alternatively, the route server/vPE-C must be programmed with the association of the vCE control channel, a VPN and an end-device IP address. As a result, the route server/vPE-C must populate the vPE distributed control plane with the

Expires <January 4, 2015>

[Page 17]

corresponding routes as non-VPN routes and the vPE must respond with VPN routes that correspond to each of these routes. Alternatively, routes reachable via a vCE must be defined via in portal per routing context and therefore VPN, and then correlated upon instantiation of the vCE on an end-system with the end-system IP address and the appropriate VRF on that end-system. In addition, the vCE must be configured with default routes per routing context with the next hop being the vPE.

8.2.2 Co-located vCE-vPE connectivity with vPE Model 2

In this model, there is no control plane routing component implemented on the end-system. That, is the end-system does not generate VPN routes and only receives VPN FIB entries from the out-of-band control plane component for routes reachable locally and for remote routes. The vCE-control plane interaction is similar to that of the interaction in Model 1 case 1-B described in the previous section whereby route population is management-driven.

8. Security Considerations

vCE creation on server - is server owned by the the operator? is this managed CE model? how to authenticate?

vCE in DC connecting VPN in WAN IP - are the DC and WAN IP VPN belong to the same SP or different? How much info are permitted to pass through auto-provisioning? How to authenticate connections, especially in pull models?

How vCE protects itself from attach from client VMs?

Additional security procedures in all virtualized cloud/DC environment, FW placement. All virtualized appliances need to be protected against attack.

Three tier (Web, App, DB) interaction access control.

Details to be added.

9. IANA Considerations

None.

10. References

10.1 Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Expires <January 4, 2015>

[Page 18]

Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [RFC4560] Quittek, J., Ed., and K. White, Ed., "Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations", [RFC 4560](#), June 2006.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), June 2010.
- [I-D.ietf-l3vpn-end-system] Marques, P., Fang, L., Pan, P., Shukla, A., Napierala, M., "BGP-signaled end-system IP/VPNs", [draft-ietf-l3vpn-end-system](#), work in progress.
- [IEEE 802.1ad] IEEE, "Provider Bridges", 2005.
- [IEEE 802.1q] IEEE, "802.1Q - Virtual LANs", 2006.
- [IEEE 802.1ag] IEEE "802.1ag - Connectivity Fault Management", 2007.

[10.2](#) Informative References

- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Service", [RFC 2475](#), December 1998.
- [I-D.fang-l3vpn-virtual-pe] Fang, L., Ward, D., Fernando, R., Napierala, M., Bitar, N., Rao, D., Rijsman, B., So, N., "BGP IP VPN Virtual PE", [draft-fang-l3vpn-virtual-pe](#), work in progress.
- [I-D.ward-irs-framework] Atlas, A., Nadeau, T., Ward, D., "Interface to the Routing System Framework", [draft-ward-irs-framework](#), work in progress.

Expires <January 4, 2015>

[Page 19]

[I-D.rfernando-irs-framework-requirement] Fernando, R., Medved, J., Ward, D., Atlas, A., Rijsman, B., "IRS Framework Requirements", [draft-rfernando-irs-framework-requirement-00](#), work in progress.

[I-D.mahalingam-dutt-dcops-vxlan]: Mahalingam, M, Dutt, D., et al., "A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks" [draft-mahalingam-dutt-dcops-vxlan](#), work in progress.

[I-D.sridharan-virtualization-nvgre]: SridharanNetwork, M., et al., "Virtualization using Generic Routing Encapsulation", [draft-sridharan-virtualization-nvgre](#), work in progress.

11. Acknowledgement

The authors would like to thank Vaughn Suazo for his review and comments.

Authors' Addresses

Luyuan Fang
Microsoft
5600 148th Ave NE
Redmond, WA 98052
US
Email: lufang@microsoft.com

John Evans
Cisco
16-18 Finsbury Circus
London, EC2M 7EB
UK
Email: joevans@cisco.com

David Ward
Cisco
170 W Tasman Dr
San Jose, CA 95134
US
Email: wardd@cisco.com

Rex Fernando
Cisco
170 W Tasman Dr
San Jose, CA

Expires <January 4, 2015>

[Page 20]

L. Fang et al.

BGP IP VPN Virtual CE

<July 4, 2014>

US

Email: rex@cisco.com

Ning So

Vinci Systems

Email: ning.so@vinci-systems.com

Nabil Bitar

Verizon

40 Sylvan Road

Waltham, MA 02145

Email: nabil.bitar@verizon.com

Maria Napierala

AT&T

200 Laurel Avenue

Middletown, NJ 07748

Email: mnapierala@att.com

Expires <January 4, 2015>

[Page 21]