International Journal of Applied Information Systems (IJAIS) – ISSN : 2249-0868
Foundation of Computer Science FCS, New York, USA
Volume 5 – No. 8, June 2013 – www.ijais.org

Campus Network Virtualization using Multiprotocol Label Switching Virtual Private Networks (MPLS-VPNs)

Frank Ibikunle
Electrical and Information Engineering Dept,
Covenant University
Ota, Nigeria

Segun Oloja
Dept. of Maths/Stat & Computer,
Achiever's University, Owo

Esi Philip O.
Electrical and Information Engineering Dept.
Covenant University
Ota, Nigeria

ABSTRACT
Campus Networks are ever evolving. With each evolution comes a need for increased availability, scalability, flexibility and security of the network. This work presents how the aforementioned needs can be met via network virtualization. It also goes further to show how network virtualization can be achieved using Multiprotocol Label Switching-Virtual Private Networks (MPLS-VPNS). Finally, the paper proposed how MPLS-VPNs (a technology generally assumed to be limited to network service providers) can be implemented in campus networks. This proposed solution as evident from the outcome of this research leads to better network design and improved network efficiency in terms of bandwidth management and network delay.

Keywords
Distribution blocks, MPLS, VRF, MP-BGP, PE, CE, LER, LSR, LSP

1. INTRODUCTION
Emerging applications such as voice and video over IP and wireless networks are built upon the campus foundation. Much like the construction of a house, if the engineering work is poorly done at the foundation level, the house will crack and eventually fail. If the foundation services and reference design in an enterprise network are not rock-solid, applications that depend on the services offered by the network will eventually suffer performance and reliability challenges. To continue the analogy, if a reliable foundation is engineered and built, the house will stand for years, growing with the owner through alterations and expansions to provide safe and reliable service throughout its life cycle. The same is true for an enterprise campus network [1]. For the purposes of high availability and fast convergence, redundant (preferably equal cost) paths should exist between network segments via redundant devices. In reality, doing this physically may lead to the skyrocketing of operating and capital expenses which in turn may be detrimental to business targets. Network virtualization, allows virtual networks to be created on the existing physical network infrastructure and while maintaining the reliability and resiliency of a physical network..

2. LITERATURE REVIEW
While some rightly describe a network as an interconnection between computer nodes (or just nodes in some cases), a broader assessment of the concept behind this description is needed. In any given organization, there is a flow of information. It is this flow of information that keeps the said organization alive. A communication network is a set of nodes that are interconnected to permit the flow of information [2]. However, upon close observation, one would notice that this information flow is not random but actually deterministic in nature. The information flows between persons or automations known to have access to certain resources (for instance, skill or certain files) and persons in need of such resources.

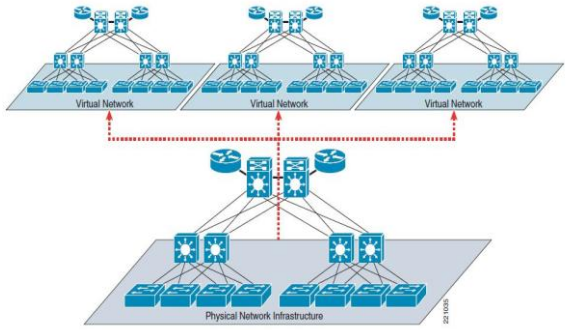

Fig.1. Network Virtualization [1]

Given this, a network is more rightly described as an interconnection between resources and resource seekers. In modern times, these resources have either been placed on computers (for non human resources) or can only be easily reached by computers (for human resources) hence the need for a computer network. Businesses, realizing the power of the computer network, have achieved improving levels of productivity and competitive advantages thus leading to an explosion in the demand on and for computer networks [3].

2.1 The Network
A computer network, often simply referred to as a network, is a collection of hardware components and computers interconnected by communication channels that allow sharing of resources and information [4]. Where at least one process in one device is able to send/receive data to/from at least one process residing in a remote device, then the two devices are said to be in a network. [5]. Computer networks (or networks as we shall henceforth refer to them) are of varying sizes and depend also on the size and need of an organization [6]. They can be generically classified as:

- Small Campus Network
- Medium Campus Network

- Enterprise Campus Network

The term "campus" is a building or group of buildings all connected into one enterprise network that consists of many local area networks (LAN). It is generally a portion of a company (or the whole company) constrained to a fixed geographic area. It can also be viewed as that portion of the computing infrastructure that provides access to network communication services and resources to end users spread over a single geographic location. It might span a single floor, building or even a large group of buildings spread over an extended geographic area [7].

Focusing on the enterprise network which shall henceforth be referred to as the "campus network" is as shown in Figure 2. It can be observed that the network has evolved to become the foundation of the business computing and communication infrastructure. Due to the ever increasing complexity of business and network requirements, a fixed model no longer describes the capabilities and services that make up the campus network today.

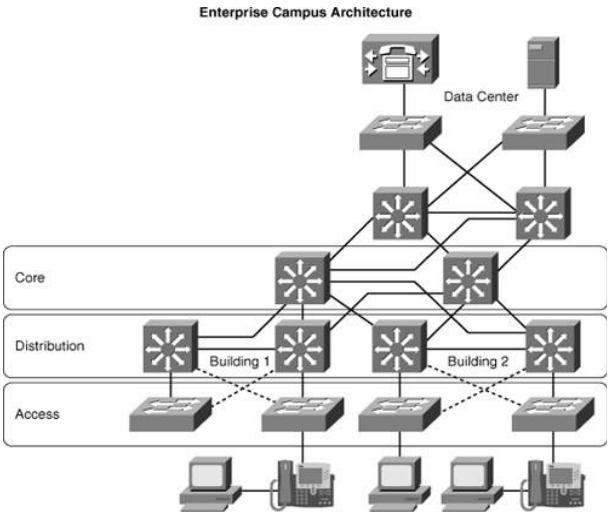

Fig.2. Enterprise Campus Architecture [3]

Instead, an ever evolving model consisting of ever evolving technologies and ever evolving devices is needed. Now, there is a dilemma which every network designer must face. The network has been widely accepted to be a critical tool for information flow, yet it is being pressured to cost less and still support the emerging applications, higher number of users and provide improved performance. So, network designer is faced with the problem of optimizing network availability at the lowest possible cost.

When designing networks, three objectives have to be considered. They are:

- Environmental givens: This consists of issues such as the location of hosts, servers, terminals and other nodes, the projected traffic for the environment and the projected cost for delivering different service levels.
- Performance constraints: This consists of issues such as network reliability, traffic throughput and host/client computer speeds.
- Internetworking variables: This consists of issues such as network topology, line capacities and packet flow assignments [7].

A healthy balance has to be maintained across these three objectives for a campus network design to be viable. However, there are scenarios in which sacrifices have to be made to meet up with business requirements. In such cases, proper opportunity costing should be carried out along with equipment hardware and software costs, performance tradeoff costs, installation costs, expansion costs, support costs, cost of downtime and sunken costs [8]. Examining more carefully the campus network as shown in Figure 3, it can be dissected into five following parts:

i) The Core part- is the backbone for campus connectivity. It serves as an aggregation point for and connects distribution blocks together while providing high speed switching (preferably layer 3 switching) between them. It is usually characterized by high redundancy, high speed links (10 GigE) and intelligent high level protocols. The kind of equipment typically found here are Layer 3 switches (or in some cases routers) like the Cisco 6500 Catalyst Switch.

ii) The Campus part- is that portion of the computing infrastructure that provides access to network communication services and resources to end users spread over a single geographic location. It might span a single floor, a building or even a large group of buildings spread over an extended geographic area.

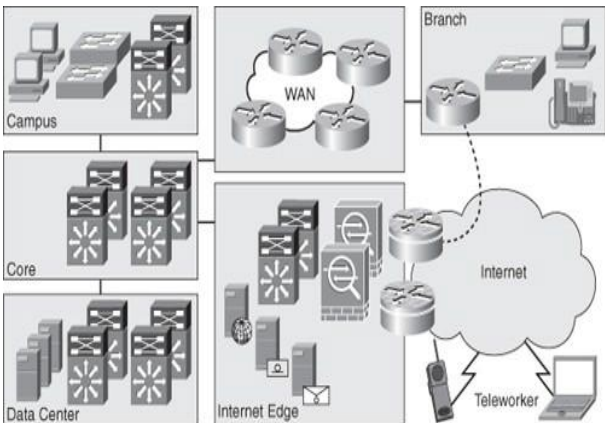

Fig.3. High-Level view of the enterprise network

This is where typical LAN technologies such as Ethernet, Fast Ethernet, Gigabit Ethernet and rarely 10GigE exist. Equipment found here include legacy hubs (these rarely exist these days) and layer 2 switches. Some designs involve having layer3 switches at the access layer also. In this segment of the network, different design considerations come into play. Network designers often find themselves in a quandary having to choose between legacy layer 2 designs and modern hierarchical and routed access designs.

iii) The data center part- is based on a layered approach to improve scalability, performance, flexibility, resiliency, and maintenance.

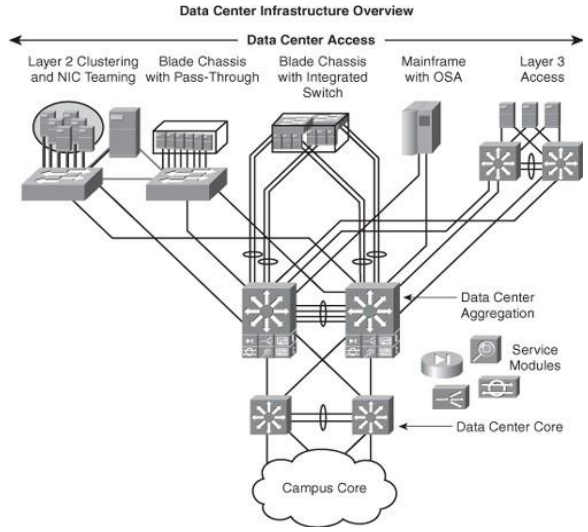

Fig.4. Data Center Infrastructure

There are three layers of the data center design as depicted in Figure 4. They are:

• **Core layer:** Provides a high-speed packet switching backplane for all flows going in and out of the data center.
• **Aggregation layer:** Provides important functions, such as service module integration, Layer 2 domain definitions, spanning tree processing, and default gateway redundancy.
• **Access layer:** Connects servers physically to the network. Multitier HTTP-based applications supporting web, application, and database tiers of servers dominate the Multi-tier data center model. The access layer network infrastructure can support both Layer 2 and Layer3 topologies, and Layer 2 adjacency requirements fulfilling the various server broadcast domain or administrative requirements. Layer 2 in the access layer is more prevalent in the data center because some applications support low-latency via Layer 2 domains. Most servers in the data center consist of single and dual attached one rack unit (RU) servers, blade servers with integrated switches, blade servers with pass-through cabling, clustered servers, and mainframes with a mix of oversubscription requirements.
iv)Wide area network (WAN) communication part occurs between geographically separated areas. WANs connect campuses together as shown in Figure 3.
v) The Internet edge part- is the network infrastructure that provides connectivity to the Internet and that acts as the gateway for the enterprise to the rest of the cyberspace. The Internet edge serves other building blocks that are present in a typical enterprise network. This modular building-block approach enables flexibility and customization in network design to meet the needs of customers and business models of differing sizes and requirements.

2.2 Multi Protocol Label Switching (MPLS)

Traditional IP packet forwarding analyzes the destination IP address contained in the network layer header of each packet as the packet travels from its source to its final destination. A router analyzes the destination IP address independently at each hop in the network. Dynamic routing protocols or static configuration builds the database needed to analyze the destination IP address (the routing table). The process of implementing traditional IP routing also is called hop-by-hop destination-based unicast routing [9, 11, 13]. Although successful, and obviously widely deployed, certain restrictions, which have been realized for some time, exist for this method of packet forwarding that diminishes its flexibility. New techniques are therefore required to address and expand the functionality of an IP-based network infrastructure [9].

MPLS defines protocols that create a different paradigm for how routers forward packets. Instead of forwarding packets based on the packets' destination IP address, MPLS defines how routers can forward packets based on an MPLS label. By disassociating the forwarding decision from the destination IP address, MPLS allows forwarding decisions based on other factors, such as traffic engineering, QoS requirements, and the privacy requirements for multiple customers connected to the same MPLS network, while still considering the traditional information learned using routing protocols as shown in Figure 5. The MPLS technology combines the richness of IP routing and the simplicity of hop-by-hop label switching of Frame Relay or ATM to provide the seamless integration of the connection-oriented forwarding with the IP world. Due to their dual nature (they operate on both the IP layer as well as the label-switching layer), the MPLS devices are called label switch routers (LSRs). All devices in an MPLS network run IP routing protocols on their control plane to build IP routing tables. In MPLS devices that support IP forwarding, the IP routing tables are used to build IP forwarding tables, also called forwarding information base (FIB) [10]. After the IP routing tables have been built, MPLS labels are assigned to individual entries in the IP routing table (individual IP prefixes) and propagated to adjacent MPLS devices through a Label Distribution Protocol (LDP). Each MPLS device uses its own local label space; globally unique labels or centralized label assignment is unnecessary, making MPLS extremely robust and scalable. Every label assigned by an MPLS device is entered as an input label in its label forwarding information base (LFIB), which is the forwarding table used for label switching [10].

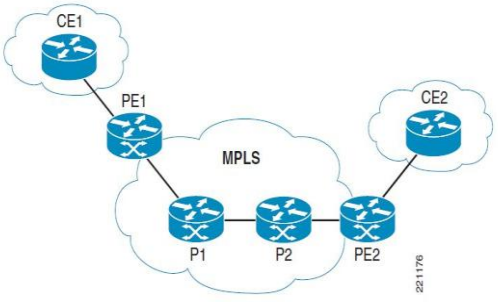

Fig.5. Basic MPLS Network Architecture [1]

Most label assignments, both local, as well as those made by adjacent devices, are entered into a table called the label information base (LIB). The label that the IP next-hop assigns for a particular IP prefix is entered as an output label in the local LFIB to enable pure label forwarding. In devices that support IP forwarding, such a label is also entered into the FIB to support IP-to-label forwarding. After the IP routing tables, IP forwarding tables, and label forwarding tables have been built, the MPLS devices can start to forward IP traffic. All MPLS devices must support label forwarding; whenever

they receive a labeled packet, they perform a label lookup in the LFIB, replace the input label with the output label, and forward the labeled packet to the next-hop LSR. Some MPLS devices (ingress LSRs) can receive IP datagrams, perform a lookup in the FIB, insert an MPLS label stack in front of the IP datagram based on information stored in the FIB, and forward the labeled packet to the next-hop LSR. The PE router within the MPLS VPN architecture is an example of such a device. Other MPLS devices (egress LSR) can receive labeled packets, perform an LFIB lookup, and (based on the absence of an output label in the LFIB) remove the label from the ingress labeled datagram and forward the IP datagram to the next-hop IP router.

2.2.1 MPLS Device Roles

Customer edge (CE) router: This is traditionally the network device at the customer location that interfaces with the service provider. In Figure 5, CE1 and CE2 represent the routers at the customer remote locations that need to be interconnected via the MPLS service provider network.
Provider edge (PE) router: This is the device at the edge of the service provider network that interfaces with the customer devices. The PE devices are often also called label switching routers edge (LSR-Edge), because they sit at the edge of the MPLS-enabled network.
Provider (P) router: These are the devices building the core of the MPLS-enabled network. Their main functionality is to label switch traffic based on the most external MPLS tag imposed to each packet and for this reason are often referred to as label switching routers (LSRs) [1].

2.2.2 The MPLS-VPN

Multiprotocol Label Switching has traditionally been viewed as a service provider (SP) routing technology: SPs have commonly used MPLS-VPN to create tunnels across their backbone networks for multiple customers. In that way, individual customer traffic is carried on a common service provider network infrastructure. Using the same principle, MPLS-VPN can be deployed inside the enterprise network to logically isolate traffic between users belonging to separate groups (as for example guest, contractors, and employees) and to provide a technical answer to business problems [1, 12]. MPLS-VPN facilitates full mesh of connectivity inside each provided segment (or logical partition) with the speed of provisioning and scalability found in no other protocol. In this way, MPLS-VPN allows the consolidation of separate logical partitions into a common network infrastructure.

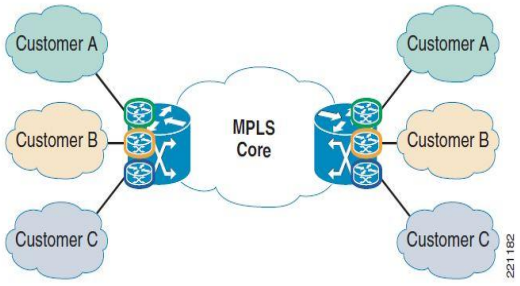

Fig.7. MPLS VPN

The key technology that simplifies the deployment of MPLS VPN is VRF (Virtual Routing and Forwarding). Defining distinct VRF instances on each PE device allows separating the traffic belonging to different customers, allowing for logical isolation and independent transport across the common MPLS core of the network. Notice that the VRF definition is required only on the PE devices, whereas the P routers in the core of the network have no knowledge of VRFs; they simply label-switch traffic based on the most external MPLS label. From a control plane perspective, an additional component now needs to be added to the IGP and LDP protocols previously discussed: Multi-Protocol BGP (MP-BGP), which is used as the mechanism to exchange VPN routes between PE devices [1, 14].

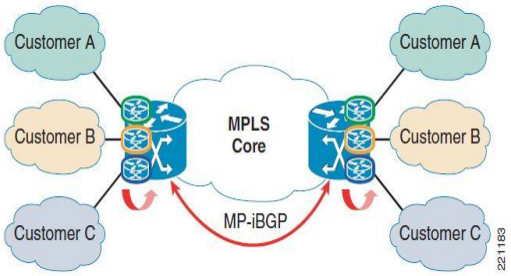

Fig.6. Control Plane For MPLS VPN

From a control plane perspective, the following two important elements need to be defined to perform the exchange of VPN routes through MP-BGP:

• Route distinguisher (RD)- Represents a 64-bit field (unique for each defined VRF) added to each 32-bit IPv4 address to come up with a unique 96-bit VPN IPv4 prefix. This ensures the uniqueness of address prefixes across different VPNs, allowing support for overlapping IPv4 addresses.

• Route target- Represents an extended attribute exchanged through MP-BGP and allows the PE devices to know which routes need to be inserted into which VRF. Every VPN route is tagged with one or more route targets when it is exported from a VRF (to be offered to other VRFs). It is also possible to associate a set of route targets with a VRF, so that all the routes tagged with at least one of those route targets are inserted into the VRF.

From a data plane perspective, the packets belonging to each VPN are labeled with two tags: the internal tag uniquely identifies the specific VPN the packets belong to, whereas the external tag is used to label-switch the traffic along the LSP connecting the ingress PE toward the egress PE [1].

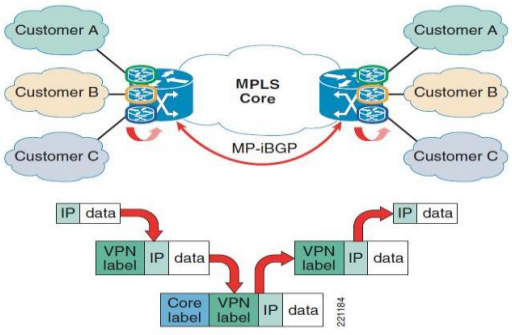

Fig.8. Data Plane for MPLS VPN

Finally, the last element that needs to be considered for an MPLS VPN deployment is the route reflector (RR) as shown in Figure 9. Each route reflector peers with every PE device (in a hub-and-spoke fashion), contributing to the overall stability of the design. Also, deploying route reflectors eases the addition of new sites, because only a new peering with the route reflector needs to be established without modifying the configuration of the remaining PE devices [1].

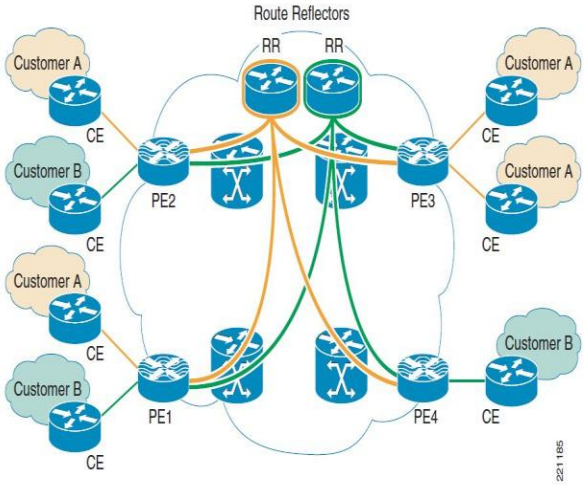

Fig.9. Deployment of Route Reflectors

In designing networks, the network designer must not lose sight of the budget restraints set by the customer and must at the same time create a design that meets the customer's needs. By tapping into the capabilities of an MPLS enabled network, a network designer can design cost effective business oriented "super" networks.

In conclusion, this section has shown that organizations that have the most resilient, highly available and high performance networks must have a network designs and implementations that are modular, a scalable and highly available networks, and a network that are relatively future proof (that is, they are built with future technological expansions in mind).

3. NETWORK DESIGN MODEL

"Best" network design practices recommend that a network be built in a hierarchical and modular way as shown in Figure 10. This is to ensure scalability, redundancy and limitation of the size of failure domains should they occur.

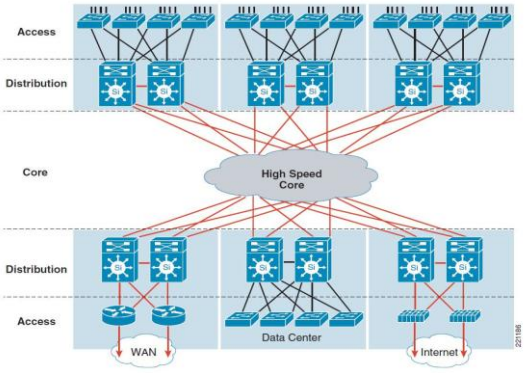

Fig.10. Hierarchical and Modular Network Model

However, the total cost of ownership of such a network and the incurred operating and capital expenditures makes it in many cases, unfeasible to build. To manage the aforementioned costs, it is possible to build just one "super" physical network and virtualizes as many networks as needed on the existing network infrastructure.

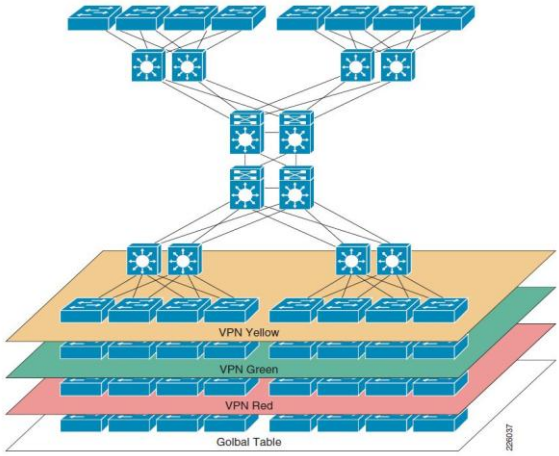

Fig.11. Virtualized Hierarchical and Modular Network Model

In Figure 11, the singular physical network infrastructure now plays host to four different networks. With network virtualization, it is important to note that all "rules" of networking still apply. Hence, redundancy, resiliency and high availability must be kept in the virtual networks. For the deployment of MPLS VPN in a campus network, a few modifications have to be made to the device roles. Most notable of these is the removal of the role of the CE. The CE is however not totally unwanted in the campus MPLS VPN architecture. The reason for its removal most times is because, the PE device is usually the first layer 3 hop from the edge of the network. This in turn, maps to the distribution layer switch with the core switches playing the P role. As such, the CE role in a normal distribution block (except in routed access topologies) is obsolete. However, in more complex distribution blocks such as the Internet Edge, the CE role is often found.

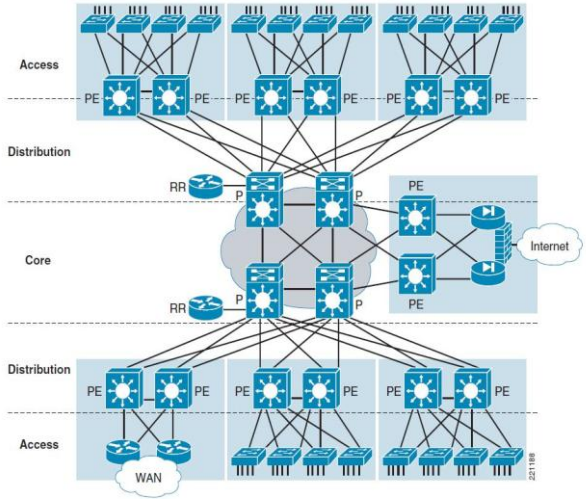

Fig.12. MPLS Device Roles in a Campus Network [1]

The test bed network utilizes a layer 3 routed core to which the other architectural building blocks are connected.

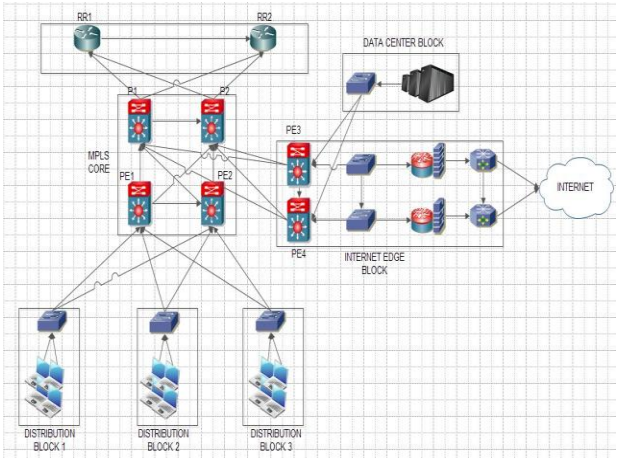

Fig.13. Physical Test Bed Network

Figure 13 shows the physical test bed network. It can be seen that distribution blocks 1 to 3 physically share two distribution switches which in turn act as PE devices in the MPLS core. Also, the data center and internet edge blocks share two distribution switches.

It can be seen in Figure 14 that logically, three pairs of distribution switches exist for distribution blocks 1 through to 3 and a pair of distribution switches now exist for the data center block. As noted earlier that, the cost of physically building the network depicted in Figure 14 can be outrageous hence the need to virtualizes the network. The test bed network in Figure 13 consists of:

- 6 Cisco 3725 Ether-Switch routers running c3725-advipservicesk9-mz.124-4.XC4 IOS acting as PE routers at the distribution layer.
- 6 Cisco 3725 Ether-Switch routers running c3725-advipservicesk9-mz.124-4.XC4 IOS acting as access layer switches. Cisco 2691 routers running c2691-adventerprisek9 _ivs-mz.124-15.T12 IOS acting as route reflectors

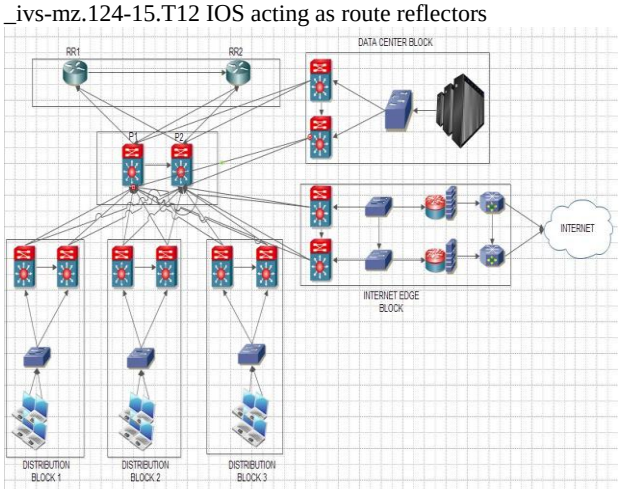

Fig.14. Logical Test Bed Network

- 4 Cisco 2691 routers running c2691-adventerprisek9_ivs-mz.124-15.T12 IOS acting as hosts and a server.
- 4 Cisco 7200VXR routers running c7200-adventerprisek9-mz.124-24.T acting as NAT and Edge routers.
- Cisco Enhanced Interior Gateway Routing Protocol (EIGRP), used between all devices in the layer 3 domain (between the distribution, core and route-reflector layers).

Before beginning virtualization, basic layer 3 connectivity must be established in the core and between distribution switches. Then, MPLS is enabled on all interfaces connected to P routers and on all interfaces interconnecting the distribution layer switches and globally in the router configuration.

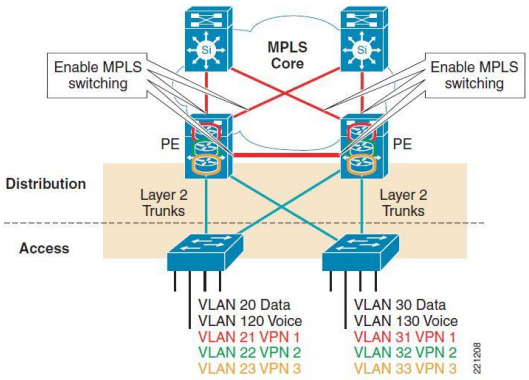

Fig.15. Enabling MPLS [1]

```
P1(config-if-range)#interface range f1/2 - 5
P1(config-if-range)#mpls ip
P1(config-if-range)#
```

Fig.16. Configuration for Enabling MPLS

Now, MPLS labels are being exchanged for all routes in the routers IPv4 routing table. This can be seen from the MPLS forwarding table of any PE router. From section 2.2.2, it can be seen that for MPLS-VPNs to work appropriately, the control plane and data plane have to be successfully built. BGP is used for the buildup of the control plane. The control plane holds all routes advertised in and into the routing domain. It is chosen because due to its possession of extended communities, larger-than-32bit routes can be sent over the network. In other words, it is capable of carrying overlapping IP addresses unlike other routing protocols.

```
PE1#sh mpls forwarding-table
Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
16     Pop tag     10.10.10.36/30   0          Fa1/0       10.10.10.33
17     Pop tag     10.10.10.52/30   0          Fa1/1       10.10.10.49
18     20          192.168.10.6/32  0          Fa1/1       10.10.10.49
       27          192.168.10.6/32  0          Fa1/0       10.10.10.33
19     Pop tag     10.10.10.24/30   0          Fa1/1       10.10.10.49
       Pop tag     10.10.10.24/30   0          Fa1/0       10.10.10.33
20     Pop tag     10.10.10.28/30   0          Fa1/1       10.10.10.49
       Pop tag     10.10.10.28/30   0          Fa1/0       10.10.10.33
21     Pop tag     10.10.10.8/30    0          Fa1/0       10.10.10.33
22     Pop tag     10.10.10.16/30   0          Fa1/0       10.10.10.33
23     Pop tag     10.10.10.40/30   0          Fa1/0       10.10.10.33
24     Pop tag     10.10.10.44/30   0          Fa1/0       10.10.10.33
25     Pop tag     192.168.10.3/32  20424      Fa1/0       10.10.10.33
26     27          10.10.10.72/30   0          Fa1/1       10.10.10.49
       16          10.10.10.72/30   0          Fa1/0       10.10.10.33
27     28          10.10.10.76/30   0          Fa1/1       10.10.10.49
       17          10.10.10.76/30   0          Fa1/0       10.10.10.33
28     Pop tag     10.10.10.60/30   0          Fa1/1       10.10.10.49
29     29          192.168.10.8/32  0          Fa1/1       10.10.10.49
```

Fig.17. Build-up of MPLS forwarding table

When VRFs are implemented, this characteristic becomes invaluable. MP-BGP peering is made between every PE router and the route-reflectors (which are redundant in our case). The core is kept BGP free. In fact, the purpose of the P routers is to label switch routes and provide a high speed connection between PE routers. Now that the control plane has been built, the data plane for each virtual network (VRF) needs to be created so that it can be populated. This is done by the creation of VRF instances on the PE/distribution switches and then the mapping of the respective VLANS in each distribution block to the appropriate VPN/VRF. All these processes are shown in Figures 18 and 19 respectively.

```
router bgp 65512
 no synchronization
 bgp cluster-id 3
 bgp log-neighbor-changes
 neighbor RRCLIENTS peer-group
 neighbor RRCLIENTS remote-as 65512
 neighbor RRCLIENTS update-source Loopback0
 neighbor RRCLIENTS route-reflector-client
 neighbor RRCLIENTS next-hop-self
 neighbor RRCLIENTS send-community both
 neighbor 192.168.10.2 peer-group RRCLIENTS
 neighbor 192.168.10.5 peer-group RRCLIENTS
 neighbor 192.168.10.6 peer-group RRCLIENTS
 neighbor 192.168.10.7 peer-group RRCLIENTS
 neighbor 192.168.10.8 peer-group RRCLIENTS
 no auto-summary
 !
 address-family vpnv4
  neighbor RRCLIENTS send-community both
  neighbor RRCLIENTS route-reflector-client
  neighbor RRCLIENTS next-hop-self
  neighbor 192.168.10.2 activate
  neighbor 192.168.10.5 activate
  neighbor 192.168.10.6 activate
  neighbor 192.168.10.7 activate
  neighbor 192.168.10.8 activate
 exit-address-family
```

Fig.18. Configuration of MP-BGP on Route Reflectors for Control Plane Buildup

When VRFs have been created, routes are pulled into them and taken out of them by the definition of route targets. Using an analogy, imagine the control plane built by MP-BGP as a marketplace. All routes in all VRFs are by default exported into the marketplace with a price tag of the export route target defined in the VRF. The routes are kept unique in the market place by the route distinguisher defined in the VRF. Now, for a VRF to "buy" or import routes it needs, it must have the required amount of cash which is the import route target defined in the VRF. If the VRF doesn't have an import route target figure that matches the target routes export route target figure, it cannot import the route.

```
PE1#sh run | sec ip vrf
ip vrf A
 rd 65512:4
 route-target export 122:1
 route-target import 122:1
 route-target import 122:6
 route-target import 122:7
ip vrf B
 rd 65512:5
 route-target export 122:2
 route-target import 122:6
 route-target import 122:7
 route-target import 122:2
ip vrf C
 rd 65512:6
 route-target export 122:3
 route-target import 122:3
 route-target import 122:6
 route-target import 122:7
 route-target import 122:8
```

Fig.19. Configuration of PE1 showing VRFs and assigned Route Targets

In the test bed network, 5 VRFs have been defined. Now, traffic must not flow between VRFs A, B and C, but can flow between each of them and VRF D and each of them and the default route only coming from the Internet VRF. To achieve this, the route targets of the excluded VRFs will not be imported. After, the VRFs have been created and respective import and export route targets have been defined, the access layer VLANs are then mapped to the required VRFs. In the test bed network, VRFA is mapped to VLAN 10, VRF B is mapped to VLAN 20, VRF C is mapped to VLAN 30, VRF D is mapped to VLAN 40 and VRF Internet is mapped to VLAN 7. Note that in the internet edge block, CE devices are used. This is because the internet edge devices have to dynamically learn about networks in the relevant distribution blocks. Figures 20 and 21 explain this.

```
NAT_FW1#sh ip route eigrp 123
     172.17.0.0/24 is subnetted, 1 subnets
D EX    172.17.10.0 [170/53760] via 172.20.10.2, 03:33:08, FastEthernet1/0
     172.16.0.0/24 is subnetted, 1 subnets
D EX    172.16.10.0 [170/53760] via 172.20.10.2, 03:33:08, FastEthernet1/0
     172.18.0.0/24 is subnetted, 1 subnets
D EX    172.18.10.0 [170/53760] via 172.20.10.2, 03:33:08, FastEthernet1/0
NAT_FW1#
```

Fig.20. VRF Routes Received By CE Router

Another requirement is that, although users are to access the internet, that's all they access in the internet edge. A similar requirement happens in the datacenter block, and to do this, there will be modification of the way export route targets are specified. An export map is used here in Figure 20.

```
ip vrf INTERNET
 rd 65512:3
 export map DEFAULT
 route-target import 122:1
 route-target import 122:2
 route-target import 122:3
```

Fig.21. Declaring an Export Map

```
route-map DEFAULT permit 10
 match ip address prefix-list DEFAULT
 set extcommunity rt  122:7
```

```
ip prefix-list DEFAULT seq 5 permit 0.0.0.0/0
```

Fig.22. Matching and Tagging the Default Route

The export map in turn references a route map that in turn references a prefix list that specifies the interested network which in this case is the default route 0.0.0.0/0. The route map sets the desired route target which can then be imported by interested VRFs as seen in Figure 22. All these processes lead to the successful deployment of an MPLS-VPN in a campus network.

4. RESULTS AND DISCUSSION

Having deployed the MPLS-VPN solution for the campus test bed network, the followings can be verified that: the successful building of VRF specific routes from the MP-BGP routes as depicted in Figure 23; Routes VRFs A and B are the only destinations that can be reached by hosts in this VRFs as shown in Figure 24; Minimal path jitter between distribution

```
PE2#sh ip bgp vpn all
BGP table version is 45, local router ID is 192.168.10.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 65512:1 (default for vrf A)
* i0.0.0.0          192.168.10.8             0    100      0 i
*>i                 192.168.10.7             0    100      0 i
* i11.11.11.11/32   192.168.10.7             0    100      0 ?
*>i                 192.168.10.7             0    100      0 ?
* i172.16.10.0/24   192.168.10.5             0    100      0 ?
* i                 192.168.10.5             0    100      0 ?
*>                  0.0.0.0                  0         32768 ?
* i172.17.10.0/24   192.168.10.5             0    100      0 ?
*>                  0.0.0.0                  0         32768 ?
* i172.18.10.0/24   192.168.10.5             0    100      0 ?
*>                  0.0.0.0                  0         32768 ?
* i172.19.10.0/24   192.168.10.7             0    100      0 ?
*>i                 192.168.10.7             0    100      0 ?
Route Distinguisher: 65512:2 (default for vrf B)
* i0.0.0.0          192.168.10.8             0    100      0 i
*>i                 192.168.10.7             0    100      0 i
   Network          Next Hop            Metric LocPrf Weight Path
* i11.11.11.11/32   192.168.10.7             0    100      0 ?
*>i                 192.168.10.7             0    100      0 ?
* i                 192.168.10.8             0    100      0 ?
* i                 192.168.10.8             0    100      0 ?
* i172.16.10.0/24   192.168.10.5             0    100      0 ?
*>                  0.0.0.0                  0         32768 ?
* i172.17.10.0/24   192.168.10.5             0    100      0 ?
* i                 192.168.10.5             0    100      0 ?
*>                  0.0.0.0                  0         32768 ?
Route Distinguisher: 65512:3 (default for vrf C)
* i0.0.0.0          192.168.10.8             0    100      0 i
* i                 192.168.10.8             0    100      0 i
* i                 192.168.10.7             0    100      0 i
*>i                 192.168.10.7             0    100      0 i
* i11.11.11.11/32   192.168.10.7             0    100      0 ?
*>i                 192.168.10.7             0    100      0 ?
* i172.16.10.0/24   192.168.10.5             0    100      0 ?
*>                  0.0.0.0                  0         32768 ?
* i172.18.10.0/24   192.168.10.5             0    100      0 ?
* i                 192.168.10.5             0    100      0 ?
*>                  0.0.0.0                  0         32768 ?
* i172.19.10.0/24   192.168.10.7             0    100      0 ?
   Network          Next Hop            Metric LocPrf Weight Path
*>i                 192.168.10.7             0    100      0 ?
Route Distinguisher: 65512:4
* i0.0.0.0          192.168.10.8             0    100      0 i
```

Fig.23. Routes in MP-BGP

blocks and a destination on the internet in the MPLS-VPN solution as compared to a normal campus network deployment.

```
PE2#sh ip route vrf A

Routing Table: A
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 192.168.10.8 to network 0.0.0.0

     172.17.0.0/24 is subnetted, 1 subnets
B       172.17.10.0 is directly connected, 00:39:25, Vlan20
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.10.0 is directly connected, Vlan10
     172.18.0.0/24 is subnetted, 1 subnets
B       172.18.10.0 is directly connected, 00:37:25, Vlan30
     11.0.0.0/32 is subnetted, 1 subnets
B       11.11.11.11 [200/0] via 192.168.10.7, 04:40:34
B*   0.0.0.0/0 [200/0] via 192.168.10.8, 04:40:34
               [200/0] via 192.168.10.7, 04:40:36
PE2#sh ip route vrf B

Routing Table: B
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 192.168.10.8 to network 0.0.0.0

     172.17.0.0/24 is subnetted, 1 subnets
C       172.17.10.0 is directly connected, Vlan20
     172.16.0.0/24 is subnetted, 1 subnets
B       172.16.10.0 is directly connected, 00:38:33, Vlan10
     11.0.0.0/32 is subnetted, 1 subnets
B       11.11.11.11 [200/0] via 192.168.10.8, 04:40:42
                    [200/0] via 192.168.10.7, 04:40:42
B*   0.0.0.0/0 [200/0] via 192.168.10.8, 04:40:42
               [200/0] via 192.168.10.7, 04:40:42
```

Fig.24. VRF specific routes

Successful ping to the 3.3.3.3 IP address on the Internet is as shown in Figure 25. To fully appreciate the ping and the power of MPLS and MPLS-VPNs, a traceroute is run from host A to the 3.3.3.3 IP address.

```
HOST_A(config)#do ping 3.3.3.3 re 10

Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 236/385/576 ms
```

Fig.25. Ping to the 3.3.3.3 address on the Internet

.

```
HOST_A(config)#do trace 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3

  1 172.16.10.1 104 msec 80 msec 116 msec
  2 10.10.10.53 [MPLS: Labels 29/41 Exp 0] 700 msec 264 msec 288 msec
  3 172.20.10.2 [MPLS: Label 41 Exp 0] 204 msec 184 msec 240 msec
  4 172.20.10.3 316 msec 244 msec 320 msec
  5 140.10.10.6 320 msec 616 msec 1284 msec
  6 150.10.10.1 536 msec 552 msec 744 msec
```

Fig.26. Traceroute to the 3.3.3.3 address on the Internet

Figure 26 shows the path that the packets take to get to the 3.3.3.3 IP address, and how the MPLS labels are swapped between the P and PE routers. The "29" label is the next-hop label while the "41" label is the VPN label. When the packet gets to 172.20.10.2, a penultimate-Hop-Pop (PHP) is done on the packet before it is forwarded on to its destination. The PHP simply refers to the removal of the MPLS label from the IP packet. This is usually done by the label edge router (LER). After the label is removed, the packet is then IP forwarded.
On a normal hierarchical campus network that has no virtualization whatsoever, the same tests are run. From Figure 31, we can see that there is a slight increase in the path jitter between Host A and 3.3.3.3 on the internet.

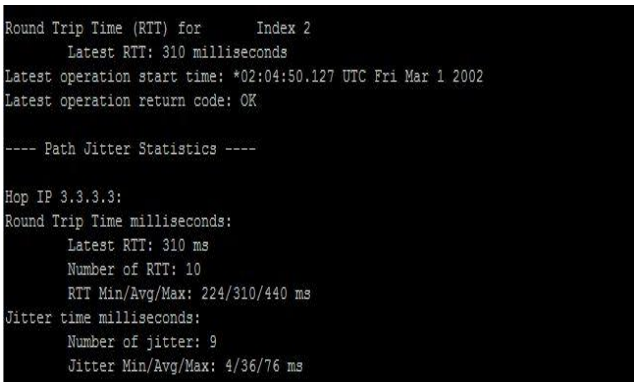

Fig.31. Path Jitter in a non-virtualized network (Between Host A and 3.3.3.3)

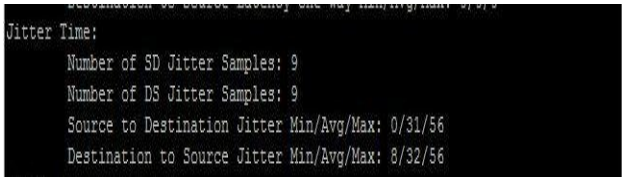

Fig.27. ICMP Jitter in a non-virtualized network (Between Two Distribution Blocks)

From the results shown and discussed, it can be seen from the various values obtained for each performance measures in Table 1 and the bar chart in Figure 28 that network virtualization via MPLS-VPNs provides the same benefits with multiple physical networks, especially in terms of network performance without incurring outrageous costs.

Table 1: Collated Network Metrics Results

Indices	Virtualized Network	Physical Network
Jitter	23ms	36ms
Delay	37ms	72ms
Round Trip Time (RTT)	400ms	310ms

Fig.28. Measured network metrics

5. CONCLUSION

The need for the deployment of business specific networks which are separate from the network used for everyday business is paramount. However, when the implied costs are considered, network designers often either boycott the business need altogether or provide very poor performing solutions. In this paper, we have shown that Campus Network virtualization using MPLS-VPNs allows for the creation and isolation of multiple networks over a shared physical infrastructure while improving bandwidth utilization and reducing network delay. We have also shown that the proposed solution allows for flexibility in the control of route propagation between all created networks and indeed the global network.

To further this research work, considerations can be given to the implication of extending the campus MPLS network to a branch over a service provider MPLS network. The researcher must also consider the benefits and implications of applying MPLS-Traffic Engineering in the network. MPLS-TE enables constraint based routing capabilities in the network which in turn helps increase the overall efficiency of the network. In conclusion, the potential of the MPLS technology is seriously untapped in some developing countries with respect to the services it can provide. These services include MPLS-VPNs and MPLS Traffic Engineering to mention a few. Enterprises and service providers alike can experience a boost in the rate of achievements of business targets by engaging MPLS in their networks.

6. REFERENCES

[1] Network Virtualization- Path Isolation Design Guide, Cisco Validated Design, 2009.

[2] Jean Warland, Communication Networks: A first Course, 2nd Edition, McGraw Hill, 1998.

[3] Froom, Frahim and Sivasubramanain, Implementing Cisco IP Switched Networks, CiscoPress, 2010.

[4] http://en.wikipedia.org/wiki/Computer_network

[5] http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap6.html

[6] Cisco Community College LAN Deployment Guide, Cisco Validated Design, 2009

[7] CCIE Fundamentals: Network Design

[8] Cisco Campus network for High Availability Design Guide, Cisco Validated Design, 2008.

[9] Luc De Ghein, MPLS Fundamentals, Cisco Press, 2007.

[10] Guichard and Pepelnjak, MPLS and VPN Architectures, Volume I, Cisco Press, 2001.

[11] Donohue, Healy and Odom, CCIE Routing and Switching Certification Guide, 4th Edition, Cisco Press, 2009.

[12] Guichard and Pepelnjak, MPLS and VPN Architectures, Volume II, Cisco Press, 2001.

[13] Callon, Rosen and Visawanathan, RFC 3031, Multiprotocol Label Switching Architecture, Jan 2001.

[14] Rekhter and Rosen, RFC 2547, BGP/MPLS IP Virtual Private Networks (VPNs), Febuary 2006