

Borderless Campus Network Virtualization Path Isolation Design Fundamentals

Last Updated: May 9, 2012



Building Architectures to Solve Business Problems



About the Authors



Mike Jessup

Mike Jessup, Solution Manager/Technical Marketing Engineer, Systems Development Unit (SDU), Cisco Systems

Mike Jessup joined Cisco in 1996 and is a Solution Manager/Technical Marketing Engineer within Cisco's Systems Development Unit. Mike's primary focus is identifying and qualifying Borderless Networks system architectures and working with Solutions Architects in validating and documenting these systems in Cisco Validated Designs. Prior to joining SDU in November of 2011, Mike worked in the US Sales Organization as an Enterprise System Engineer specializing in Routing and Switching, Data Center, Application Optimization, and Network Management. Prior to joining Cisco, Mike worked in positions as a Field Engineer and Systems Engineer in various Information Systems companies and Partner organizations since 1982.



Rahul Kachalia

Rahul Kachalia, Technical Marketing Engineer, Systems Development Unit (SDU), Cisco Systems

Rahul Kachalia is a technical marketing engineer in Cisco's Systems Development Unit, helping to create the design guidance that will help build the next-generation borderless network infrastructure. Kachalia has more than 15 years of broad internetworking experience, primarily in service provider core and edge focused products and technologies including broadband, MPLS, VPN, and managed services. He has led many key initiatives to develop solutions that can deliver network design guidance and accelerate infrastructure and advanced services deployments for Enterprise networks. In the System Architecture and Strategy group he has also worked on designing next-generation unified virtual campus networks to enable borderless services for large enterprise customers. Kachalia holds CCNP, CCNA, MCSE, MCP, CNE, and CCIE certifications (11740, Routing & Switching and Service Provider). He holds a bachelor's degree from Mumbai University, India.

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit <http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <http://www.cisco.com/go/trademarks>. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Borderless Campus Network Virtualization—Path Isolation Design Fundamentals

© 2012 Cisco Systems, Inc. All rights reserved.



Borderless Campus Network Virtualization— Path Isolation Design Fundamentals

Introduction

As an organization's network architects and engineers consider designing their company's next-generation campus networks, they must be able to address many critical, anticipated and unanticipated requirements. With the explosion of time sensitive, and potentially bandwidth intensive, voice and video used collaboratively in videoconferencing as well as internal communications and training and the ever-changing requirements for network access, the networks being built today have a very different set of requirements than those in the not so distant past. Four prime areas of concern are:

- **Scalability**—The network must be built taking a modular approach that allows easy expansion of the network as employees, applications, or services are added. It should also easily scale to accommodate growing bandwidth requirements through equipment upgrades rather than wholesale replacement.
- **Availability**—The network must be capable of providing non-stop operation with negligible data loss during unexpected outages. Availability is addressed not only through redundant hardware and network paths, but through the use of resilient protocols ensuring user access to the network regardless of how they connect.
- **Mobility**—Users and devices must be able to move and attach seamlessly throughout the network, regardless of the type of device or how they attach.
- **Security**—The network itself must be capable of supporting secure, authenticated access to users and devices attaching to it. Furthermore, the network must be capable of serving as a policy enforcement point either granting or restricting access to an organization's assets based on who or what is accessing the network as well as from where they are accessing it.

For design guidance around building scalable and highly available campus networks, an excellent reference is the Borderless Campus 1.0 Design Guide at:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cOverall_design.html.

Additional guidance for mobility and security can be found in the Unified Access Design Guide at:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_unified_access.html.



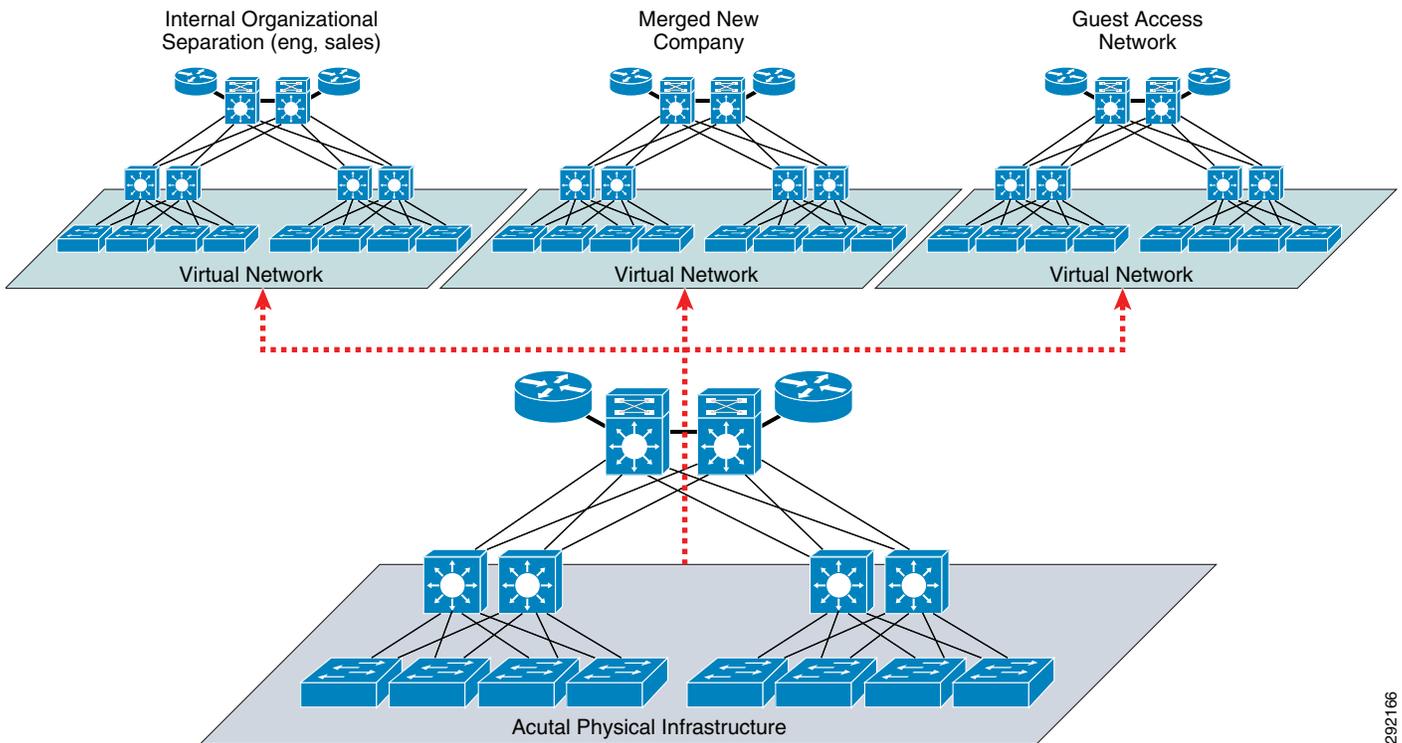
Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Copyright © 2011 Cisco Systems, Inc. All rights reserved.

Overview

The purpose of this document is to focus on one important aspect of providing a secure infrastructure through the use of network virtualization. Network virtualization may be described as the ability to logically define individual, virtual networks (VNs) whose data planes for forwarding traffic and control planes for processing unicast and multicast routing information are completely isolated from each other. Effectively, each virtual network can consist of its own set of devices, interfaces, and services, allowing complete isolation from other virtual networks or, as policy dictates, the ability to share common services between virtual networks. Figure 1 depicts three virtual networks defined on a single physical infrastructure.

Figure 1 The Concept of a Virtual Network



292166

Organizations may have business requirements that call for the separation of various classes of traffic or the isolation of specific groups of users from one another. Typically, all corporate data or voice traffic would flow across the infrastructure at large or, as we discuss later, the default or “global” network, while specific types of users, such as guests and contractors, or applications, such as IP video surveillance, would be assigned to a virtual network that is isolated from the rest of the network.

Some possible use cases where network virtualization may be useful include:

- Education—College campuses divided into administrative and student (dormitory) networks as well as guest access and potentially departmental isolation.
- Retail—Separate network required for all point-of-sale (POS) equipment for Payment Card Industry (PCI) compliance.
- Energy—Isolation of critical infrastructure communications (supervisory control and data acquisition [SCADA]), such as power generation and transmission data from administrative traffic for critical infrastructure protection (CIP) compliance.

- Manufacturing—Isolation of machine-to-machine (M2M) traffic from typical administrative traffic.
- Government—Isolation of various hosted government agencies from one another on a common infrastructure.
- Healthcare—Networks defined for imaging, patient guest access, robotics communications, and HIPAA compliance.
- Financial Services—Isolation of production and non-production networks, physical security, contractor access, Teller 21 applications, and regulatory compliance mandated by OCC, FDIC, SEC, and SOX.

In these organizations, the ability to create separate, logical networks and place specific classes of traffic within the confines of the virtual network construct can result in:

- Lower total cost of ownership (TCO) through elimination of duplicate hardware.
- Less administrative overhead through reduction of policy enforcement points implementing access control lists (ACLs) or firewalls.
- Ease of integration of two networks during merger and acquisition where overlapping address spaces may exist.

The fundamental building block of a virtual network is the Virtual Routing and Forwarding (VRF) instance. VRF is a technology that allows multiple instances of a routing table to exist on a single router with each instance isolated from the other. A VRF consists of its own IP routing and forwarding table as well as a set of physical interfaces assigned to the VRF that use the derived forwarding table. VRFs are discussed in more detail in [Virtual Routing and Forwarding](#).

Network virtualization is enabled by technologies such as VRF-Lite (also known as Multi-VRF) and Easy Virtual Network (EVN), either exclusively or in addition to MPLS. The use of one or all of these technologies allows organizations to implement a scalable virtualization strategy providing secure, logical isolation within a single physical network. VRF-Lite has been, and is today, commonly utilized in networks where only a handful of virtual networks are required. MPLS provides the additional scalability to define hundreds and even thousands of virtual networks within an organization.

In the past, VRF-Lite has been the most easily deployed technology when less than eight VNs are required. VRF-Lite relies on common routing protocols such as OSPF and EIGRP and a hop-by-hop configuration of the VNs by first defining them on the networking device, such as a switch or router, and then configuring the method by which the devices interconnect. This interconnection can be accomplished through the use of 802.1q Trunking or Generic Route Encapsulation (GRE)/Multipoint GRE (mGRE), and mapping each virtual network (VNET) to a VLAN over an Ethernet trunk or tunnel in the case of GRE/mGRE.

Recently, VRF-Lite has been enhanced through features in the new technology known as Easy Virtual Network or EVN. EVN makes configuration easier and improves scalability to 32 VNs per device. In addition to incorporating features found in Multi-VRF, EVN introduces new concepts such as VNET Trunks, VNET Tags, and a simplified approach to route replication between virtual networks.

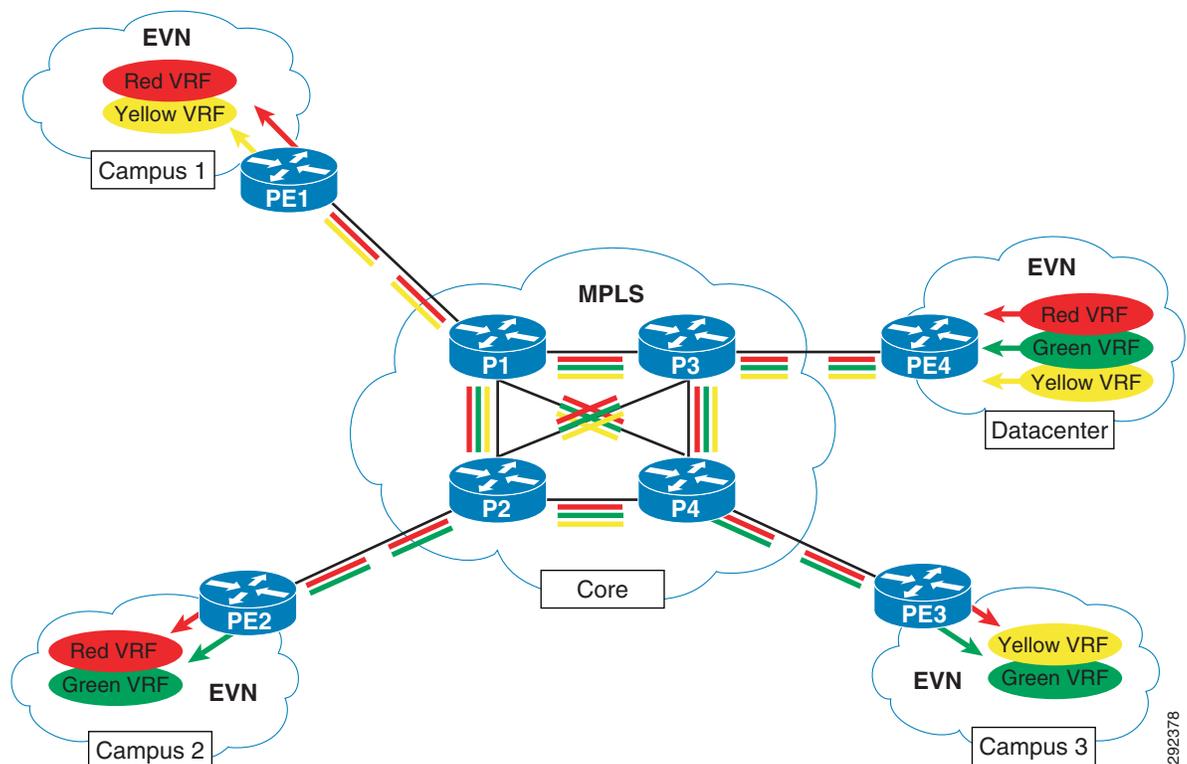
When more than 32 VNs are required, an organization can use MPLS to provide any-to-any virtualized connectivity. The primary advantage of MPLS is that it provides end-to-end connectivity, without the need to define connections on a hop-by-hop basis, through the use of label switching enabled by the Label Distribution Protocol and Multi-Protocol BGP (MP-BGP) for the exchange of VRF information at each Layer 3 device. Optionally, one might use MPLS within the core of the network and deploy EVN or Multi-VRF within campus distribution blocks or in remote sites over the WAN. For some organizations, MPLS may pose too many challenges and prove to be too complex as it requires skills that their staff may not have. This is where the extended capabilities of EVN can be a valuable alternative.

Pre-Deployment Considerations

Prior to planning and deploying a network virtualization strategy, it is imperative that you first review the existing network design. Issues with Layer 2 and, in particular, Layer 3 designs and any potential issue that may exist will be replicated and most likely compounded as the number of control and data planes increase with each virtual network. As a result, these issues will certainly complicate troubleshooting when problems occur.

The first consideration should include a review of the network architecture. Many organizations have a clearly defined model which includes an access layer for user and device connectivity, a distribution layer for access aggregation, a core for interconnecting distribution blocks, and edge services supporting Internet, third party, and WAN aggregation. Others may have simply deployed a collapsed distribution/core. The point is that regardless of the model followed, it is easier to develop and deploy a virtualization strategy if there are distinct boundaries between areas of the network such that the approach taken to virtualization can be tailored to suit the needs of each “block”. For example, as depicted in [Figure 2](#), one might look at a large campus with a number of buildings where a technology such as EVN may be implemented within that building or distribution block because only a couple of VNs may be required. Subsequently each block might then be interconnected to the core via MPLS and MP-BGP for increased scalability and communications between instances of a VRF present in multiple distribution blocks. This recommendation should not be ignored by smaller organizations where a collapsed core has been deployed with multiple wiring closets and a shared services segment or compute block attached. In this example, the organization may use EVN throughout and, by clearly delineating where access and compute or services meet, redistribution of routes between the VNs can occur there rather than being spread throughout the infrastructure.

Figure 2 Campus Virtualization Using EVN Interconnected with MPLS Core



292378

An analysis of the Layer 2 and Layer 3 protocols should be conducted as a precautionary measure to ensure that a sound foundation is present upon which network virtualization can be a simple overlay. From the Layer 2 perspective, consideration should be given to using either Rapid-PVST+ or 802.1s MST/802.1w as the Spanning Tree protocol of choice, static definition of STP domain relative to Root placement, and various Layer 2 extensions such as Uni-Directional Link Detection (UDLD), Root Guard, and BPDU Guard. From a Layer 3 perspective, one should review how the routing protocol has been configured to optimize the addressing strategy for summarization where possible and the protocol timers for fast and efficient convergence after a failure.

In addition to the considerations discussed here, refer to the Borderless Campus 1.0 Design Guide (http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/Borderless_Campus_Network_1.0/Borderless_Campus_1.0_Design_Guide.html) which provides design guidance for building a robust, scalable, and extremely resilient campus network.

Network Virtualization Technologies

Within the campus, network virtualization began with the concept of VLANs. VLANs provide the most basic means of isolating network traffic at Layer 2 in a bridged/broadcast domain and require a Layer 3 device to route between those domains. As we discuss network virtualization today, the VLAN is still considered to be the most basic element by which we group users and devices.

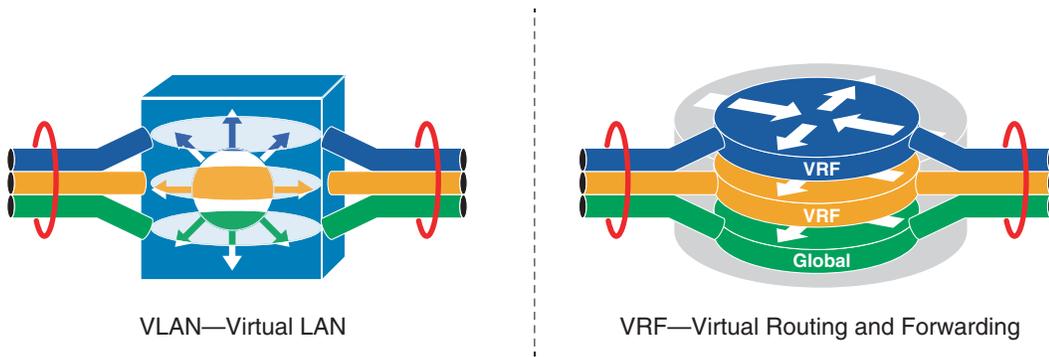
VLAN Assignment

In the past, the most common approach to VLAN assignment would be to manually configure a port to be a member of a specific VLAN and potentially define a voice VLAN for that port as well. Another method which is becoming much more common today is through the enhanced security capabilities of Flexible Authentication Sequencing using 802.1X, MAC Authentication and Bypass (MAB), or Webauth as alternate means to first authenticate a user against a Radius Server or a Policy Enforcement Server, such as the Cisco Identity Services Engine (ISE), for network access. Once authenticated, by using Radius attributes communicated between the Radius Server and access switch the switchport is dynamically changed to the appropriate VLAN and, optionally, an ACL can be pushed down to the switch enforcing specific access to the network. It is beyond the scope of this document to detail these technologies; refer to the TrustSec Phased Deployment Configuration Guide at: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/Phased_Deploy/Phased_Dep_Guide.html#wp392175.

Virtual Routing and Forwarding

We discussed how VLANs are the most basic path isolation technique for Layer 2. However as the goal of every solid network design is to minimize the extent of the broadcast domain and exposure to Spanning Tree loops, a method to translate the Layer 2 VLAN to a Layer 3 virtual network or VPN is required. This Layer 3 VN must be capable of supporting its own unique control plane complete with its own addressing structure and routing tables for data forwarding completely isolated from any other Layer 3 VPN on that device and in the network. The technology enabling this type of functionality is known as Virtual Routing and Forwarding (VRF) instance. [Figure 3](#) draws the comparison between Layer 2 VLANs and Layer 3 VRFs.

Figure 3 Comparison of Layer 2 VLAN and Layer 3 VRF

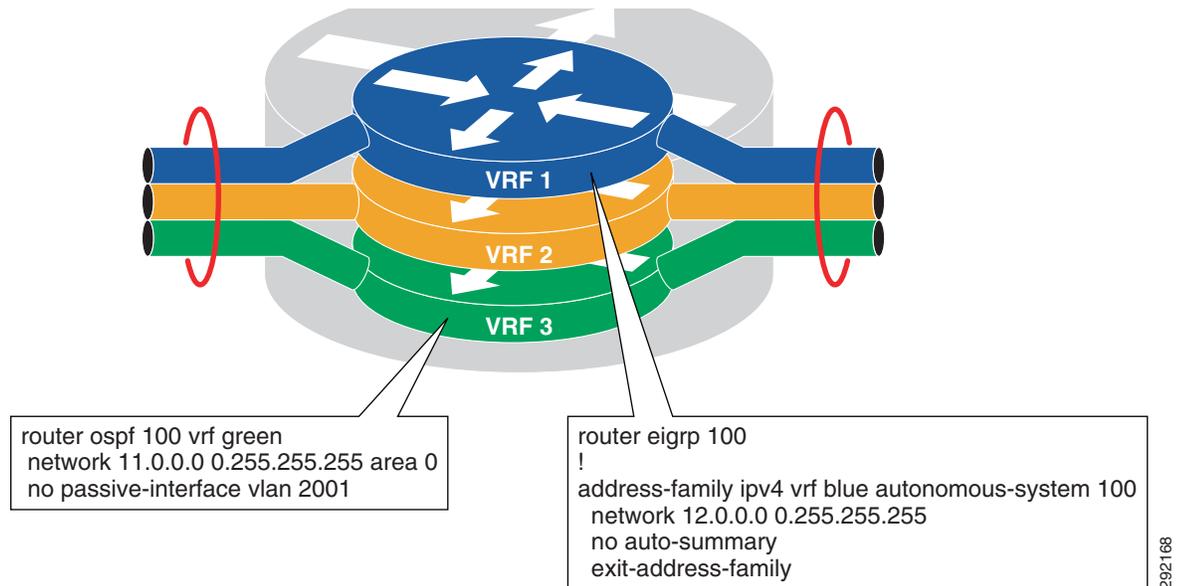


The VRF is defined on a networking device that serves as the boundary between the Layer 2, client-side VLANs, and the Layer 3 network. Each VRF instance consists of an IP routing table, a forwarding table, and interface(s) assigned to it. Common routing protocols such as OSPF, EIGRP, BGP, and RIPv2 can be used to advertise and learn routes to populate the routing table unique to each virtual network. This routing information is then used to populate the CEF table using those interfaces, either logical (SVI) or interfaces and sub-interfaces that have been specifically allocated to that VRF through device configuration. VRFs exist on top of a global routing table consisting of IPv4 prefixes and interfaces that have not been assigned to a VRF.

VRF instances can be compared to virtual routers co-resident on a single Layer 3 switch or router. However, the comparison stops inasmuch as the VRF does not carve out any dedicated compute or memory from the physical device. Figure 4 depicts three VRFs residing on a single physical device.

As we discuss characteristics of routing in a virtualized environment, it is important to understand that Cisco routers support a number of routing protocols and individual processes per router. Some routing protocols such as BGP only support a single instance of that process and hence the concept of routing contexts was developed. These contexts were designed to support isolated copies of the routing protocol running per VRF. An example of this concept is depicted in Figure 4.

Figure 4 Routing Contexts through Processes and Address Families



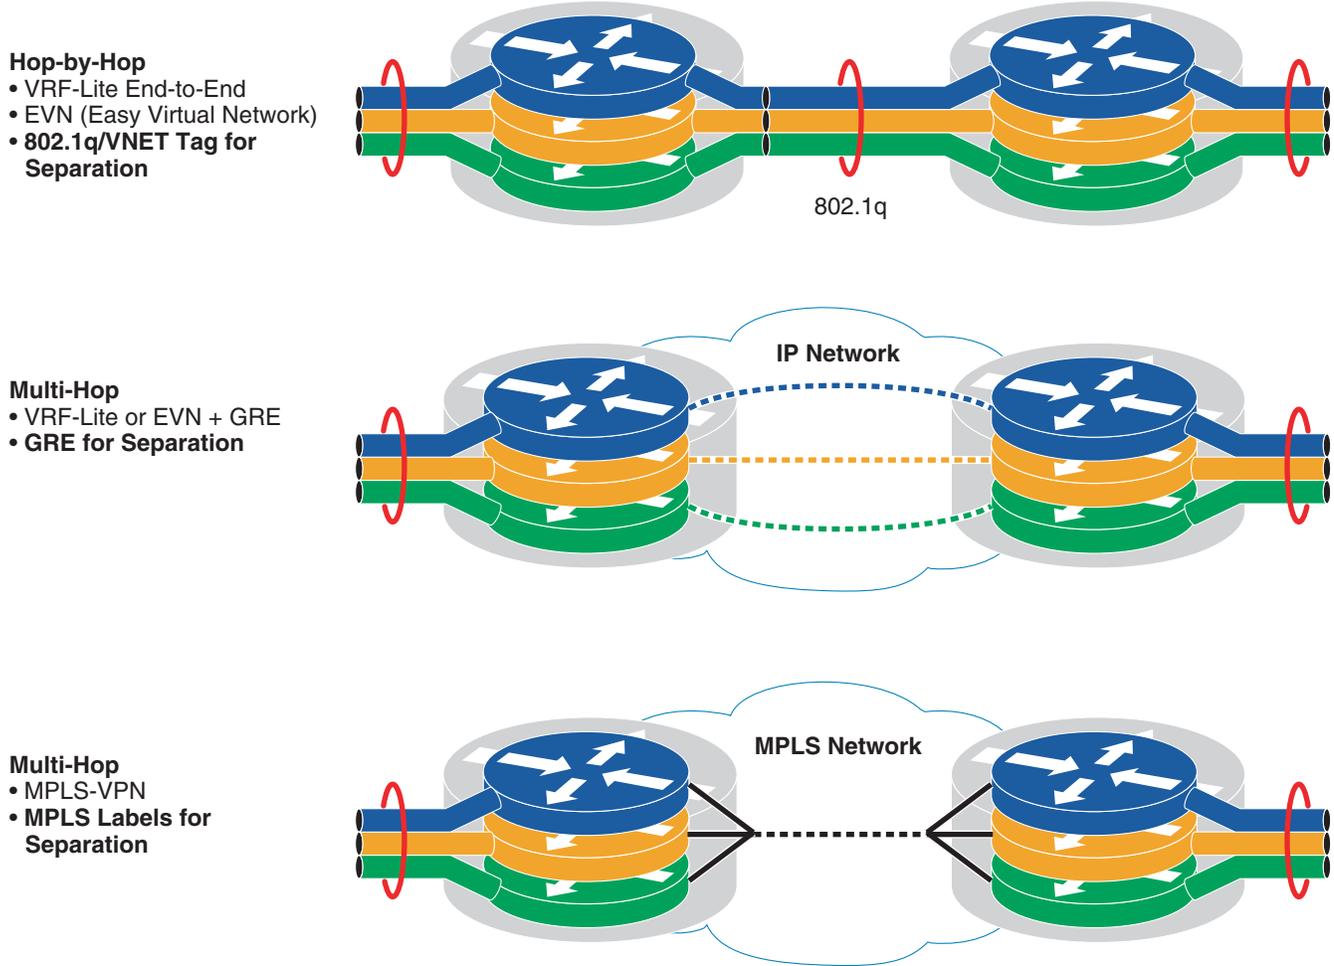
An important concept to understand is that VRFs and the virtual networks they represent, complete with their own routing tables and forwarding paths, are an overlay on top of a base or non-virtualized IP infrastructure. This non-virtualized infrastructure has its own “Global” routing table and data forwarding plane. Typically, as an organization approaches network virtualization, all user/device traffic flows across the non-virtualized infrastructure using the global routing table. In most cases, a virtual network is only defined to accommodate the special requirements of a specific type of traffic or policy surrounding a group of users such as “Guest” or “Contractor” and what those users can access. As such the vast majority of an organization’s user and device traffic may remain within the global table. Once a specific type of traffic has been migrated to a virtual network, it is possible to “leak” routes between the virtual networks or the global domain to grant specific access to resources within other VNs or the global table. Obviously, it is possible to remove all user/device traffic from the global domain and place it into a virtual network, but a great deal of consideration must be given to providing sufficient access between virtual networks so as to not unintentionally block access to other resources.

Path Isolation

Having discussed the method by which the control and data plane of the networking equipment are virtualized, we now cover path isolation and the means by which the routers and switches along with the VRFs configured on them are connected together .

The VRF instance on a networking device is an isolated object that must be linked to other instances of the same VRF on other devices throughout the network. As touched on briefly in [Overview](#), there are several means by which this is accomplished today. Should any-to-any connectivity between numerous sites be required, MPLS may provide the best alternative. If however a hop-by-hop, multi-hop, or hub-and-spoke approach is sufficient, a subset of the functionality found in MPLS and implemented in a technology known as Multi-VRF or VRF-Lite in conjunction with 802.1Q or GRE/mGRE may be used. Moving forward, a new technology that we discuss in this document known as EVN will likely be used as an alternative to VRF-Lite. A depiction of these technologies can be found in [Figure 5](#).

Figure 5 Path Isolation Techniques



Whereas VRF-Lite can still be used, it is anticipated that with the new enhanced capabilities found in EVN that it will become the preferred method for providing connectivity within organizations requiring 32 virtual networks or less. EVN introduces several new capabilities including VNET Tags, VNET Trunks, and route replication which minimizes the amount of work required to configure the virtual network and, more importantly, eliminates the need for MP-BGP, as is the case even with Multi-VRF, to leak routes between VRFs. It is worth noting that EVN is still backward compatible with VRF-Lite when utilizing 802.1Q trunks from non-EVN devices. A more detailed discussion follows in subsequent sections.

To determine which technology to use, a couple of factors should be considered. There are no set rules to determine the correct approach as a combination of methods might be suitable.

If an organization wants to deploy less than 32 virtual networks and does not have any specific requirements for MPLS, traffic engineering, and BGP, EVN may prove to be more than sufficient. Additionally this assumes that the virtual networks do not require a full mesh type of topology providing a substantial amount of any-to-any connectivity over a WAN. This is not to say that it is impossible to build a full-mesh topology, just that it may become impractical from an administrative point as one could easily construct multiple mGRE tunnels between sites.

On the other hand, should you require more than 32 virtual networks per device or require traffic engineering and optimal path selection with Fast Reroute capabilities along with the ability to provide extremely scalable full mesh connectivity, MPLS is the correct choice. One consideration that should not

be overlooked are the capabilities for Traffic Engineering and Fast ReRoute found within MPLS. Traffic Engineering provides the capability to efficiently utilize multiple paths through an MPLS network based on the type of traffic as well as the resource requirements for a given class of traffic while using RSVP to establish and maintain a TE Tunnel throughout the MPLS core. Additionally, on platforms supporting Fast Rerouting, Tunnel traffic can be easily switched to another tunnel in approximately 50 milliseconds. To support Traffic Engineering, the only other consideration is that the IGP in the MPLS core is either OSPF or IS-IS. However as pointed out earlier, MPLS and features such as Traffic Engineering require additional skills, including a fairly sophisticated understanding of Multi-Protocol BGP and Label Distribution Protocol. It is beyond the scope of this document to detail the functionality of traffic engineering, however you can get additional information at:

http://www.cisco.com/en/US/tech/tk436/tk428/tech_design_guides_list.html.

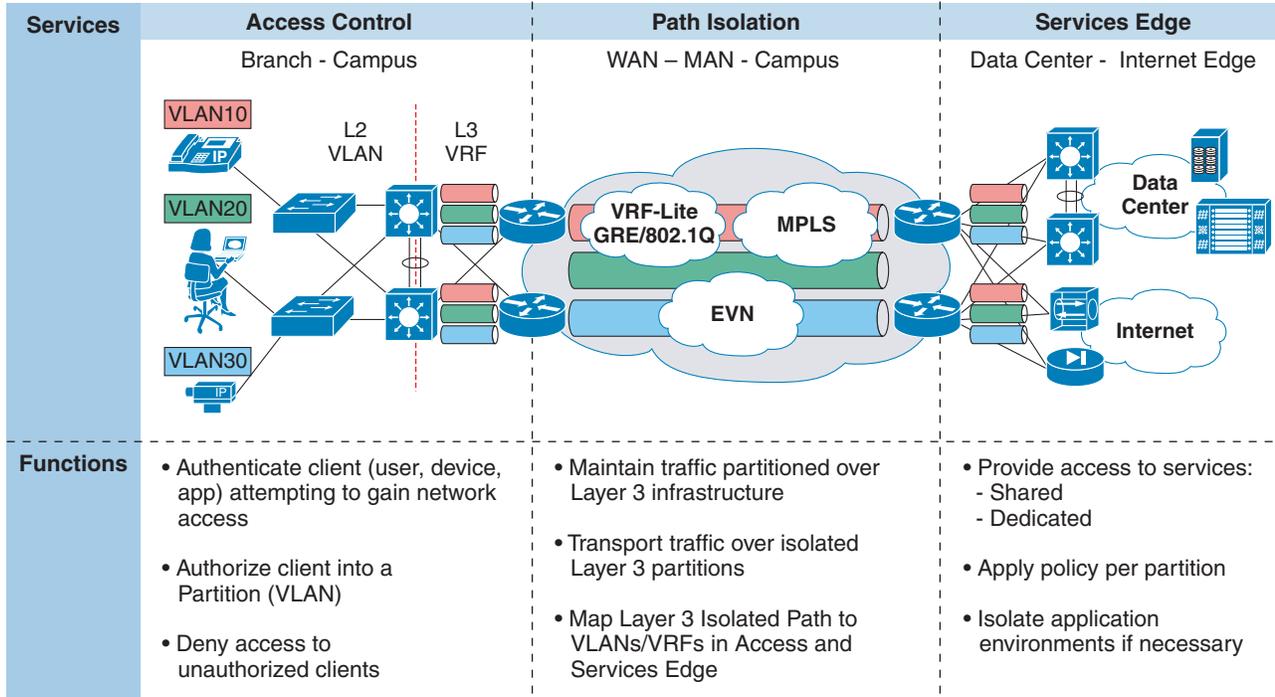
Table 1 Network Virtualization Technique Comparison Chart

	Multi-VRF (Hop-by-Hop VN)	EVN (Hop-by-Hop VN)	MPLS (Edge-to-Edge VN)
Network Design	Small to mid-size networks based on VN scalability requirements		Large size
VN Scale Limit	8 ¹	32	4000
Technology Integration	Non-disruptive—No design and infrastructure change required.	Disruptive—No design change but requires EVN-capable hardware and software.	Disruptive—Requires MPLS-capable core and edge systems.
Global Infrastructure	IP-based		IP/MPLS-based
Operational Complexity	Simple. Complex as network expands.	Simple	Complex. Minimized with simplified network design.
VN Manageability	Disruptive—Requires hop-by-hop VN provisioning. However EVN simplifies the manual sub-interface configuration procedure.		Non-disruptive—Required only at edge.
Shared Services	BGP-based inter-VRF route leaking.	EVN-based inter-VRF route replication.	BGP-based inter-VRF route leaking.
VN Expansion	Limited scale. Can evolve to MPLS.		High
Layer 2 VPN Extension	Not supported by VRF-Lite/EVN.		Supported
Virtualization over WAN	Heterogeneous with VRF enabled GRE, DMVPN.		MLS, MPLSoGRE
Supported Campus Platforms	Cisco Catalyst 3xxx Cisco Catalyst 4500-E Cisco Catalyst 4500-X Cisco Catalyst 6500-E Cisco Nexus 7000 Cisco ASR 1000	Cisco Catalyst 4500-E Cisco Catalyst 4500-X Cisco Catalyst 6500-E (Sup2T) Cisco ASR 1000	Cisco Catalyst 6500-E Cisco Nexus 7000 Cisco ASR 1000

1. VRF-lite scalability on each supported campus platform varies. However scaling beyond 8 VRF on a system may become complex from a management, operational, and troubleshooting perspective.

In the following sections, we look at each of the virtualization technologies in greater detail. The network virtualization concepts discussed in this section are summarized in Figure 6.

Figure 6 Network Virtualization Architecture



292170

VRF-Lite

Multi-VRF or VRF-Lite has been used for several years as an alternate method to MPLS for implementing path isolation in a network on a limited scale. Going forward as EVN support extends beyond the ASR100, Catalyst 6500, and Catalyst 4500, it will likely be adopted over VRF lite as the preferred method to deploy network virtualization due to the simplified configuration it introduces. In the event that EVN is not supported on a platform, VRF-Lite can still be used in conjunction with EVN as it is fully backward compatible.

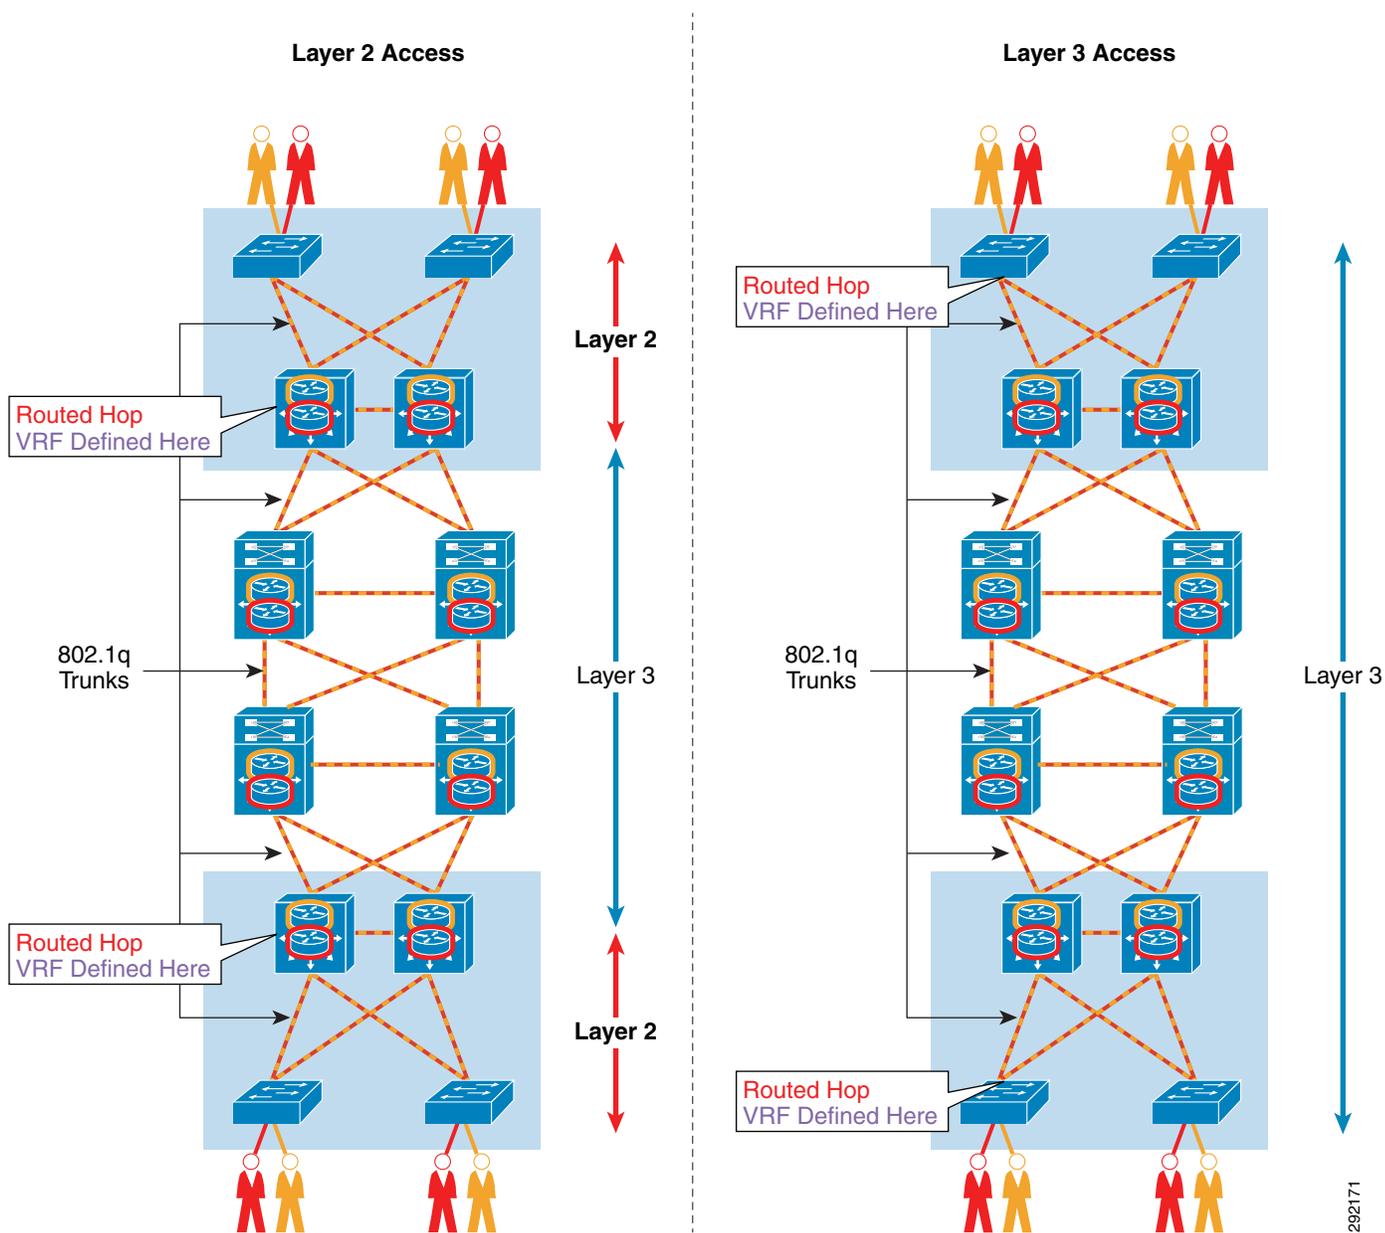
VRF-Lite combines the use of VRF Instances with either 802.1q trunking for hop-by-hop path isolation or Generic Route Encapsulation (GRE)/Multipoint GRE (mGRE) for multi-hop path isolation. VRF-Lite with 802.1Q trunking would be found in a campus network where the IP routing is completely under the control of the organization deploying it and would typically make use of OSPF or EIGRP as the IGP for routing. VRF-Lite with GRE/mGRE would typically be found when a VRF must be extended over an IP network that is not under the control of the organization or in instances where several devices with the same VRF(s) need to be interconnected over a non-virtualized Core network.

VRF-Lite End-to-End (Hop-by-Hop)

VRF-Lite deployed on a hop-by-hop basis in a campus makes use of 802.1Q trunks to interconnect the devices configured for VRFs. If a campus network were a completely flat Layer 2 network with a single spanning tree instance, a VLAN would obviously be used to provide path isolation. However, as it is

against all campus best practices to build a network in this fashion, Spanning Tree domains are restricted to the smallest diameter possible while using Layer 3 links/trunks to connect devices together. As defined in the Borderless Campus 1.0 Design Guide (http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cOverall_design.html), there are two models for connecting the access layer to a distribution or a collapsed distribution/core. The first model is where the access layer is connected to the distribution or collapsed core via Layer 2 links and the second is through the use of Layer 3 links. The VRFs are defined where the Layer 2 VLANs border the Layer 3 network. Therefore, if the access layer is connected to aggregation via Layer 2, the VRFs are defined on the distribution or collapsed core switches aggregating the access layer. If however the access layer is Layer 3 connected, the VRFs are defined on the access switch itself. Figure 7 depicts these relationships and where the VRFs are defined.

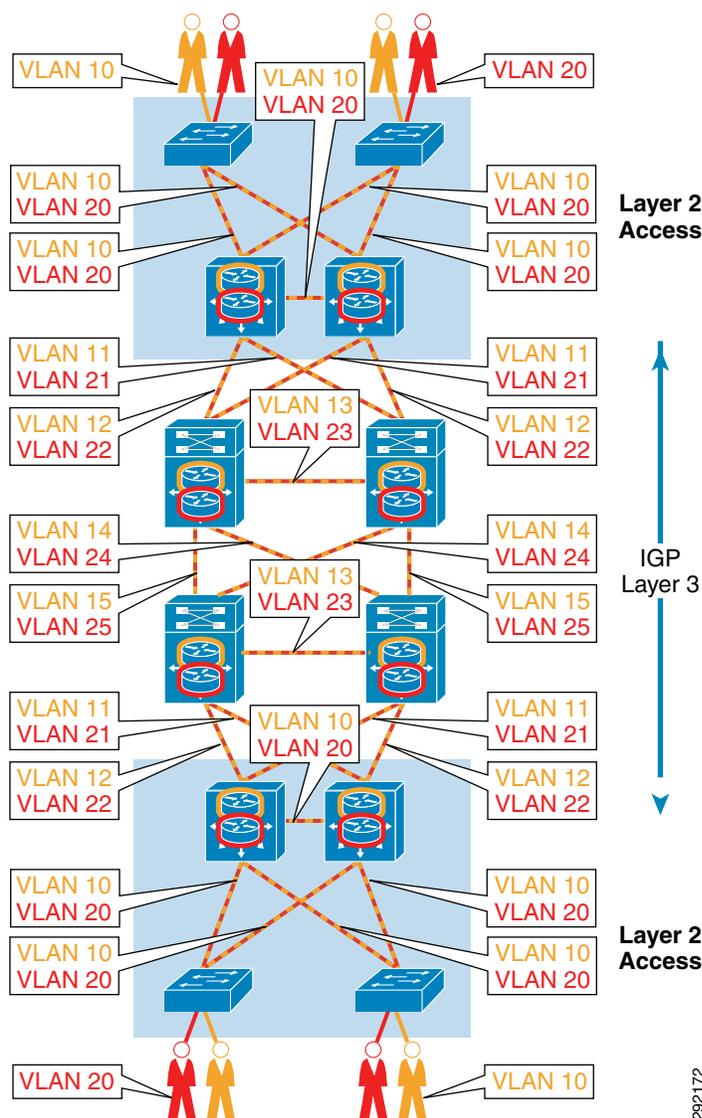
Figure 7 VRF Definition



If the access is Layer2, the link between the access switch might belong to a single VLAN or, more likely, if multiple VLANs exist, an 802.1q trunk is used. Regardless, the SVI used for gateway services and routing these VLANs is configured at the distribution switches. After having defined the required VRFs on these distribution switches, the SVI is then configured to be a member of the VRF. Likewise, if the access switch is configured for Layer 3 access to distribution or collapsed core, the SVI is defined on the access switch as well as the VRFs. The SVI used for gateway services and routing for the access VLANs (users and devices) is then assigned as a member of a VRF.

Once the VLAN to VRF mapping has been completed at either the access or distribution networking hardware, the core-facing interfaces must then be configured. These interfaces can potentially be configured in two different ways if a Catalyst 6500 is used as the access switch. The first approach would be to use a VLAN and its associated SVI, which would then be assigned to the appropriate VRF. The second would be to use sub-interfaces on the core-facing interface with each sub-interface assigned to the appropriate VRF. Although the use of sub-interfaces is preferred, it must be noted that the Catalyst 6500 is the only Catalyst switching platform supporting routed sub-interfaces and thus the Catalyst 4000 and Catalyst 3000 platforms require the use of SVIs. For each VRF, a unique VLAN ID must be used for every physical link interconnecting network devices as each is considered a routed hop. [Figure 8](#) depicts VLAN assignment in a Layer 2 access scenario.

Figure 8 VLAN Definition Example for VRF-Lite End-to-End



The example in [Figure 8](#) depicts two VRFs, Red and Orange. As depicted above, the Red traffic only traverses those VLANs that have been defined as members of the Red VRF and never "leaks" into the Orange VRF unless explicitly configured to do so. Note that VLAN IDs have been reused throughout the infrastructure, as in the case of the access VLANs as well as 802.1q trunks or sub-interfaces interconnecting Layer 3 devices. This is possible as these are routed and not bridged hops with SVIs or sub-interfaces defined for each VLAN.

VRF-Lite and GRE (Multi-Hop)

VRF-Lite can be used in conjunction with GRE or Multipoint GRE (mGRE) when it becomes necessary to extend a virtual network across a Layer 3 infrastructure or domain where virtualization is either not required or, as in the case of a service provider WAN, is beyond the control of the organization. When using GRE to interconnect VRFs, all configuration of the virtualized infrastructure remains the same as

previously discussed in the end-to-end model. The major difference however is that rather than using 802.1q for interconnecting devices, GRE is used. In addition to GRE/mGRE, MPLS can also be used to extend VRFs across an IP infrastructure as well. MPLS is discussed in a following section.

GRE, when used in conjunction with VRF-Lite, can be configured as either point-to-point or multipoint, as in the case of a WAN branch network where the headend or hub router is configured for mGRE and each branch has one or two GRE tunnels pointing to one or more hub routers. The use of VRF-Lite over GRE is typically, although not limited to, use at the WAN or Internet edge. In a campus scenario, VRF-Lite in conjunction with 802.1q can be used to virtualize the multi-tiered infrastructure up to the edge of the network. The edge router or Layer 3 switch then has a unique GRE or mGRE tunnel defined for each VRF requiring IP transit. Here on these devices, the VRF is defined on the access-facing SVI or sub-interface in the case of a router as well as the GRE or mGRE tunnel to be associated with that VRF.

The first consideration in deploying VRF-Lite over GRE is platform selection. As of this writing, only Cisco routers, Catalyst 6500, and Catalyst 4500 support GRE. In most scenarios this should not be an issue as one would anticipate that a router or switch such as a Catalyst 6500 would be found at the WAN edge.

The second consideration lies in the decision to use GRE or mGRE at the hub as, also of this writing, only Cisco routers and the Catalyst 6500 support Multipoint GRE as well as Next Hop Routing Protocol required for the hub devices to dynamically learn about the spoke devices and build a GRE tunnel to each of them. Configuration of the hub devices is greatly simplified due to this ability to dynamically learn the spokes as opposed to having to configure each one individually. The use of mGRE is thus the preferred choice if more than a few spoke devices exist or tunnels are required.

[Figure 9](#) depicts a campus network connecting over a non-virtualized core or SP cloud to an Internet edge DMZ. In [Figure 9](#) you see that each hub device has a unique GRE tunnel built to each spoke. In turn, each spoke has two GRE tunnels, one to each hub for high availability.

Figure 9 VRF-Lite Used in Conjunction with GRE at the Hub Routers

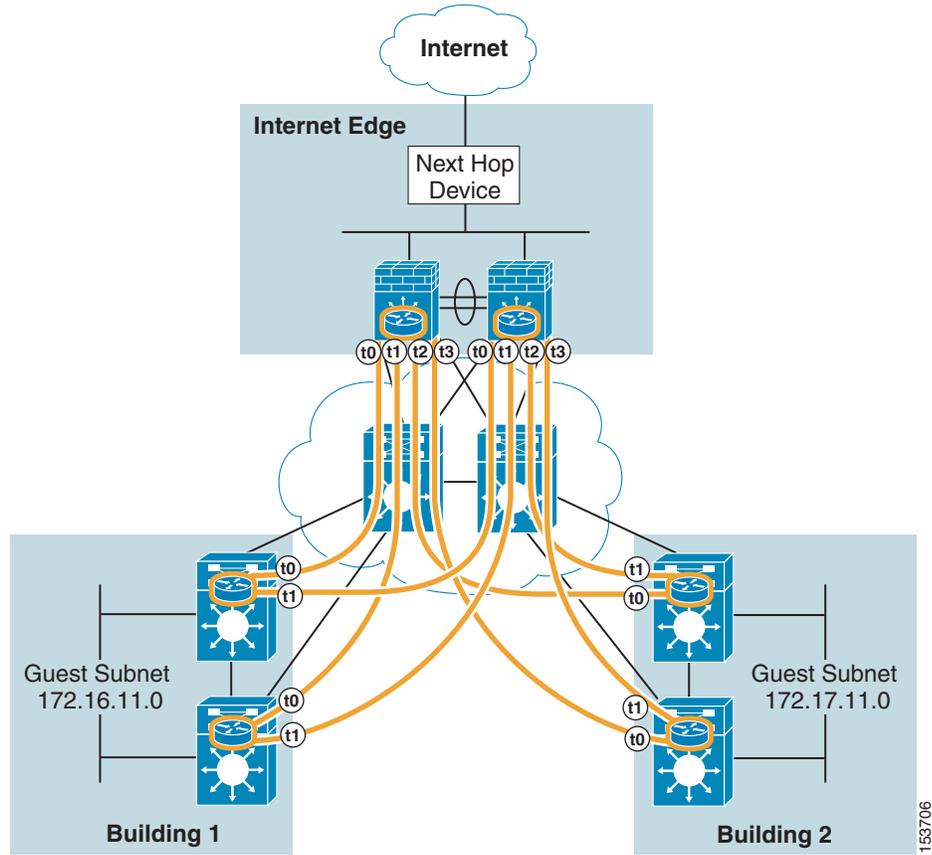
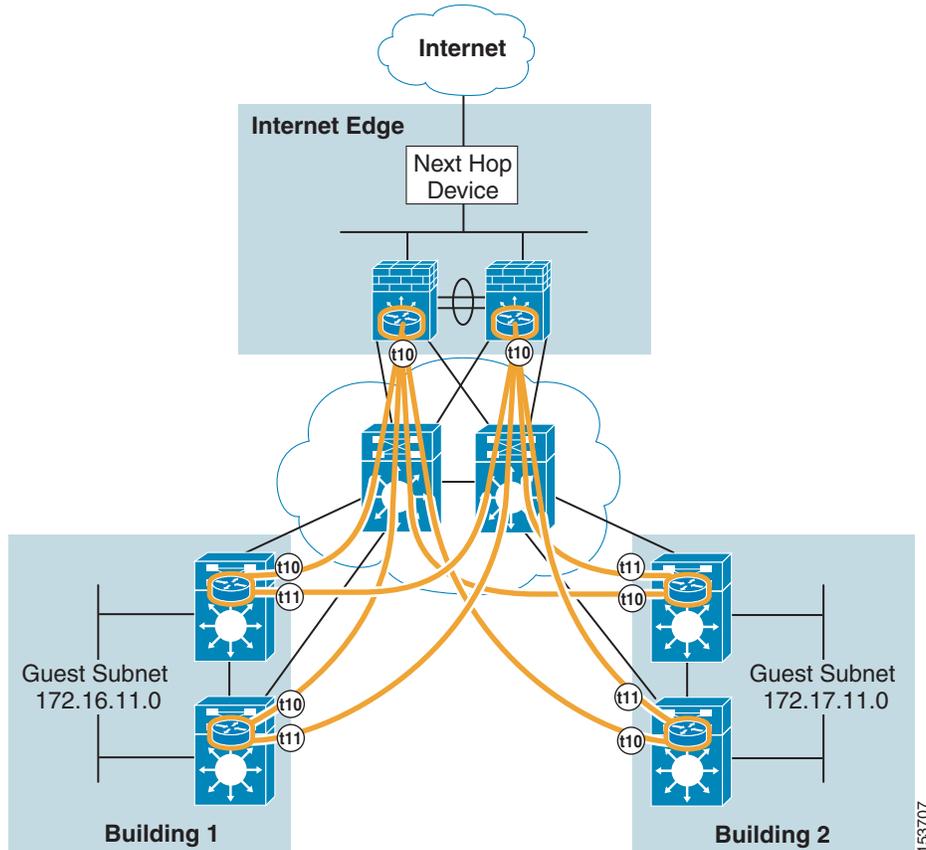


Figure 10 depicts the same campus network except a mGRE Tunnel is defined at hub device.

Figure 10 VRF-Lite Used in Conjunction with mGRE at the Hub Routers



In both [Figure 9](#) and [Figure 10](#), we see the GRE/mGRE tunnels that have been defined for forwarding traffic belonging to the Orange VRF. Should additional VRFs be required, it would be necessary to build another, complete set of interconnecting tunnels for that traffic.

VRF-Lite and Routing

After having defined the overlay network for path isolation through the use of 802.1q trunks and GRE/mGRE tunnels, the routing protocol(s) must be configured. As pointed out earlier, in addition to the global routing table which exists on the physical infrastructure, each VRF has its own control plane with its own routing table. The virtual networks that have been defined use the routing protocol that has been configured for the physical infrastructure (global table).

It is possible but impractical to run a separate routing protocol for each virtual network and typically a single protocol such as EIGRP or OSPF is most commonly be used. This is possible through the use of “Address Families” in EIGRP or unique processes in OSPF.

With EIGRP, a routing process is defined on each Layer 3 device. The process is first configured for the global table and then an address family is defined for each respective VRF. The following example depicts an EIGRP configuration.

```
router eigrp 100
network 10.0.0.0
  eigrp router-id 10.122.137.1
  no auto-summary
!
```

```

address-family ipv4 vrf Green
  network 11.0.0.0
  no auto-summary
  autonomous-system 100
  eigrp router-id 10.122.138.1
  exit-address-family
!
address-family ipv4 vrf Red
  network 12.0.0.0
  no auto-summary
  autonomous-system 100
  eigrp router-id 10.122.139.1
  exit-address-family

```

In the example for EIGRP above, make note of the `autonomous-system` command defined within each VRF as this is one of the primary differences between the global table and the address family. It is required to activate the address family and associates it with the autonomous system ID of 100 defined for EIGRP on the device. As with standard EIGRP routing, it is necessary for the autonomous system ID to match between neighboring devices to build an adjacency.

OSPF is similar to EIGRP, however instead of using address families, it uses separate OSPF processes as in the following example. Note that the global process is ID 100.

```

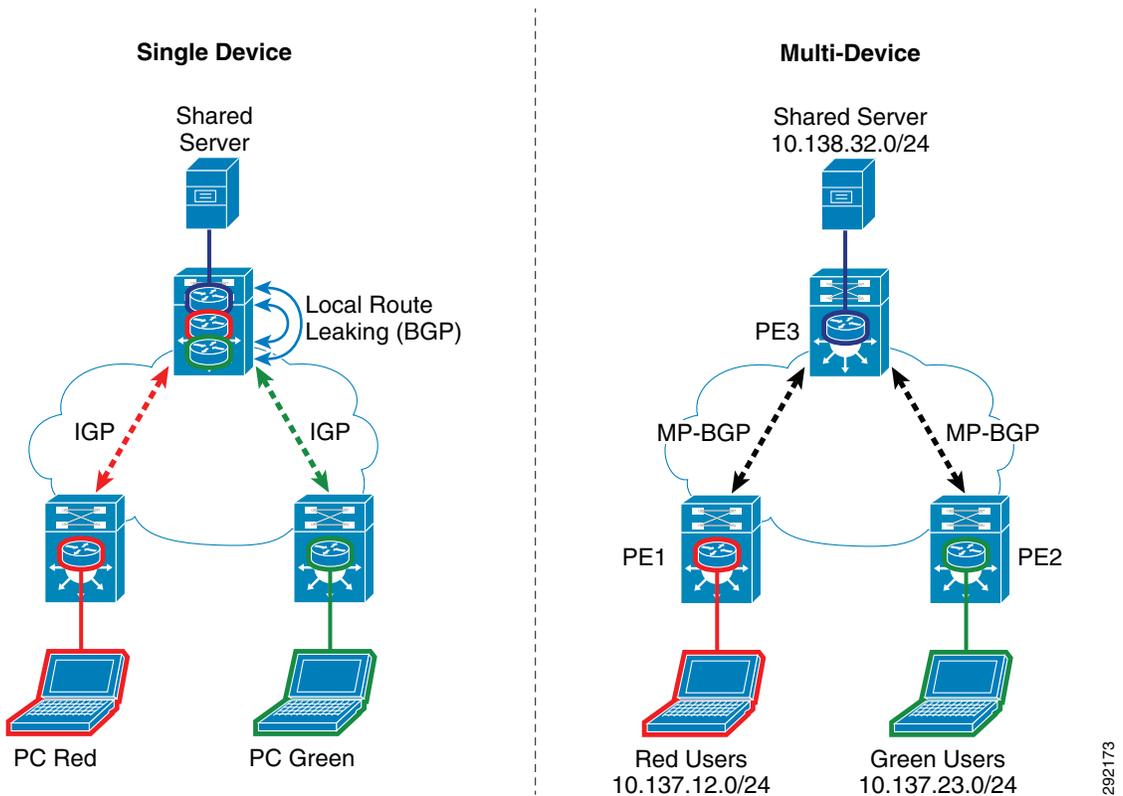
router ospf 1 vrf Green
  network 11.0.0 0.255.255.255 area 0
  router-id 10.122.138.1
!
router ospf 2 vrf Red
  network 12.0.0 0.255.255.255 area 0
  router-id 10.122.139.1
!
router ospf 100
  network 10.0.0 0.255.255.255 area 0
  router-id 10.122.137.1

```

The one drawback to Multi-VRF which has been simplified dramatically within EVN and is discussed in the next section, is the ability to “leak” routes between VRFs. As discussed earlier, traffic traversing a VRF cannot access any routes or devices within the global table or another VRF. To access these resources, the routes from the global or VRF to those resources must be advertised through a redistribution or “leakage” between the tables. For Multi-VRF, the means by which this is accomplished is through the use of the route target attribute found within a Multi-Protocol BGP Process (MP-BGP) and the ability to define a route target value that can be exported (advertised) from each VRF as well as a route target value from other VRFs that can be imported (learned) within a VRF. Route maps can be used to explicitly define which prefixes will be associated with a route target, thereby filtering what routes should be “leaked”. MP-BGP can be configured on either a single device where this redistribution is desired or on multiple devices upon which MP-BGP has been configured and neighbor relationships established for sharing BGP routing information. Once the MP-BGP process has been configured, it is then possible to redistribute the leaked routes into the IGP for that specific VRF or global table.

Typically, this route leaking process is configured at the Internet edge to leak the default route into the various VRFs or in the data center for advertisement of servers/services located there, which might include DNS, DHCP, Microsoft AD, and other application or Web servers. [Figure 11](#) depicts both the single and multi-device deployments. It is important to understand that although a route may have been leaked at one networking device, this routing information is not propagated to other devices without either redistributing that route into the IGP and advertising it or configuring a similar MP-BGP process on the other devices requiring the leaked route.

Figure 11 Route Leaking on a Single Device or Multiple Devices



It is beyond the scope of this document to further detail the configuration steps required for MP-BGP and route leaking, however, for complete information as to how to accomplish this, refer to the Network Virtualization—Services Edge Design Guide at: http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/ServEdge.html.

Summary

For additional information that provides much greater detail as well as best practices for configuring network virtualization and Multi-VRF, refer to the Network Virtualization—Path Isolation Design Guide at: http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/PathIsol.html#wp277683.

For additional information that goes into much greater detail about configuring MP-BGP and configuring shared services within a Multi-VRF environment, refer to the Network Virtualization—Services Edge Design Guide at: http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/ServEdge.html.

Easy Virtual Network

Multi-VRF was one of the key edge technologies in the overall service provider offered managed VPN network designs. Each VRF uniquely builds secure, trusted, and connection-oriented relationships with enterprise edge systems. This design helps build a scalable, connection-less solution for enterprises to

interconnect different remote sites. With its simplicity and security, Multi-VRF quickly became a proven and efficient overlay technique to solve segmentation challenges in enterprise campus and branch network designs. The initial VRFs deployments were simple, but as IT learns about virtualization benefits it becomes challenging to deploy and manage the expansion of additional VNs. The campus core requires a full-mesh network design to provide optimal data load balancing and redundancy. To integrate Multi-VRF in such a resilient campus design, every in-path system and network path requires redundant configurations and control plane.

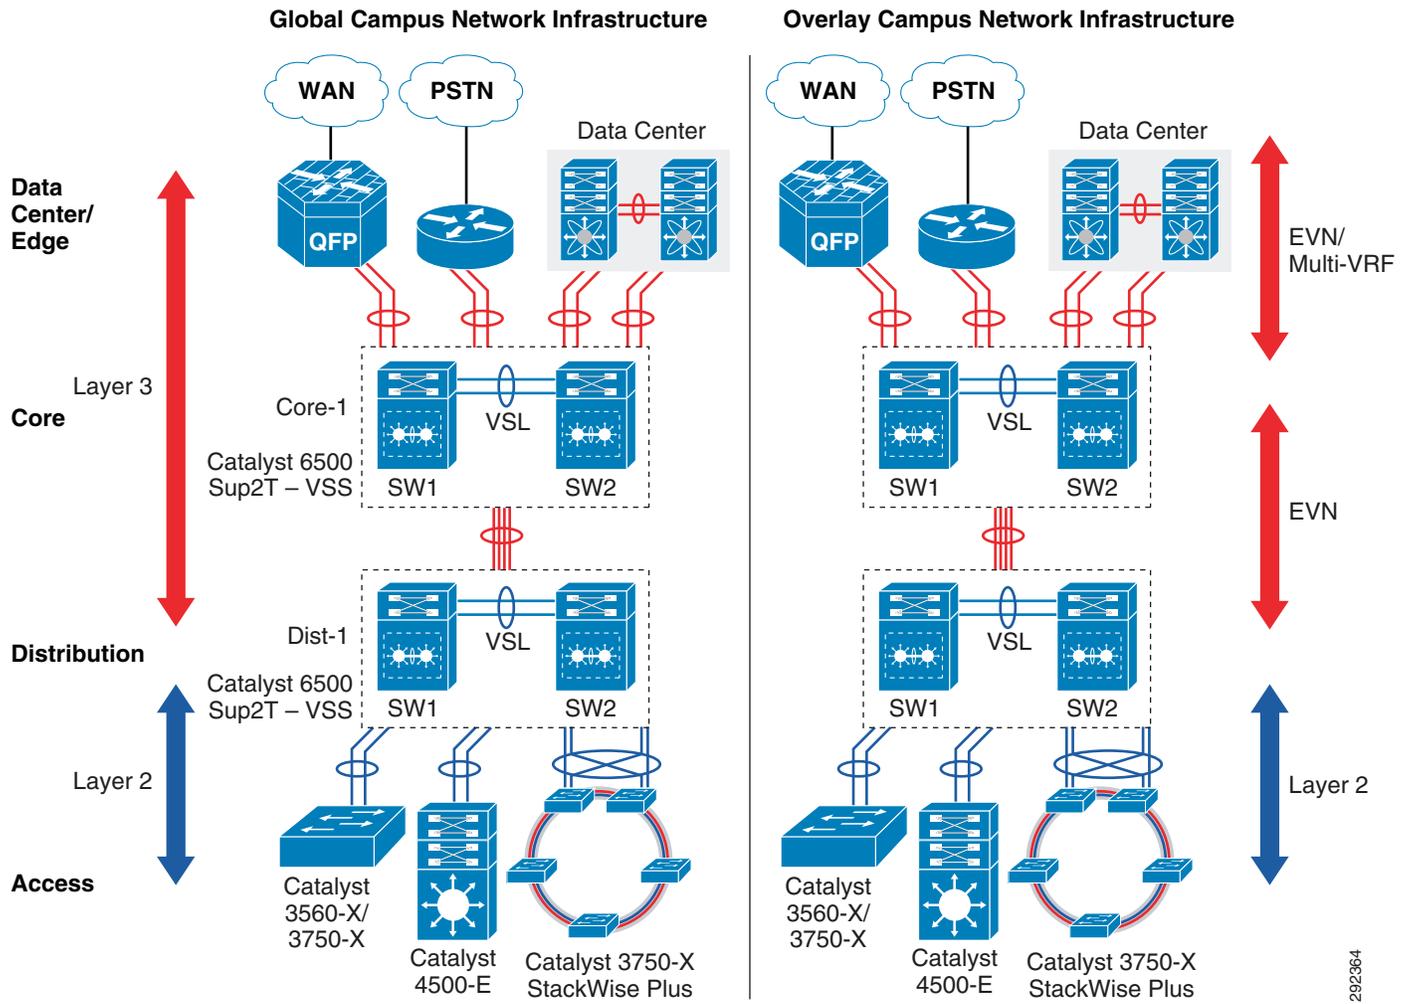
To mitigate such operational challenges, Cisco IOS software implemented Easy Virtual Network (EVN) to significantly simplify the operational tasks required to provision, unprovision, and manage VNs in campus networks. The Cisco EVN is a new virtualization technology that can be considered an evolution of Multi-VRF in that new constructs have been added, such as VNET Tags, VNET Trunks, and Route Replication, which greatly simplify the configuration of virtual networks. Technically the concept of VRFs has not changed whatsoever with EVN and in fact EVN is completely backward compatible with Multi-VRF. Whereas Multi-VRF can scale to at least eight VNs to efficiently operate the network, EVN eliminates operational complexity and provides additional scalability up to 32 VNs. In the campus switching portfolio, Cisco EVN technology is supported on the next-generation Cisco Catalyst 6500-E with Supervisor 2T (Sup2T) starting with 15.0(SY1) and the Cisco Catalyst 4500-E and Cisco Catalyst 4500-X starting with the 15.1(1)SG IOS release. To extend the EVN operational boundary up to the enterprise edge, the Cisco ASR 1000 series routers support EVN starting with the IOS XE 3.2S release.

Path Isolation with EVN in Campus

Cisco EVN technology does not change the fundamentals of the hop-by-hop virtualization technique. Each Layer 3 campus system involved to support VPN routing infrastructure requires provisioning VRF and per-VRF routing adjacency. The data plane segmentation with EVN remains consistent as Multi-VRF by uniquely imposing an IEEE 802.1Q header for each packet as they traverse across the Ethernet network. By retaining existing per-VN control plane signaling and 802.1Q-based data plane forwarding mechanism, it provides the flexibility to integrate EVN and pair with Multi-VRF systems or even interconnect to the MPLS network edge boundary.

Virtualization allows you to integrate several overlay networks in an existing foundational network. The network architect can design the virtual network environment by retaining the existing global network infrastructure hierarchy and the boundary between Layer 2 and Layer 3 domain. [Figure 12](#) illustrates a campus-wide end-to-end virtualized network environment with EVN capable system and interoperability with other Layer 2 and Layer 3 systems.

Figure 12 Hop-by-Hop Path Isolation with EVN in Campus



292364

EVN VRF Design

EVN does not change the fundamentals of enabling VRF in the system. In Cisco IOS software the network administrator can create VRF using legacy and new definition-based models. Functionally both methods are valid, however the new definition introduces a multi-protocol address family concept that simplifies using a single VRF to enable virtualization for IPv4 and IPv6 networks.

6500-E—Core and Distribution

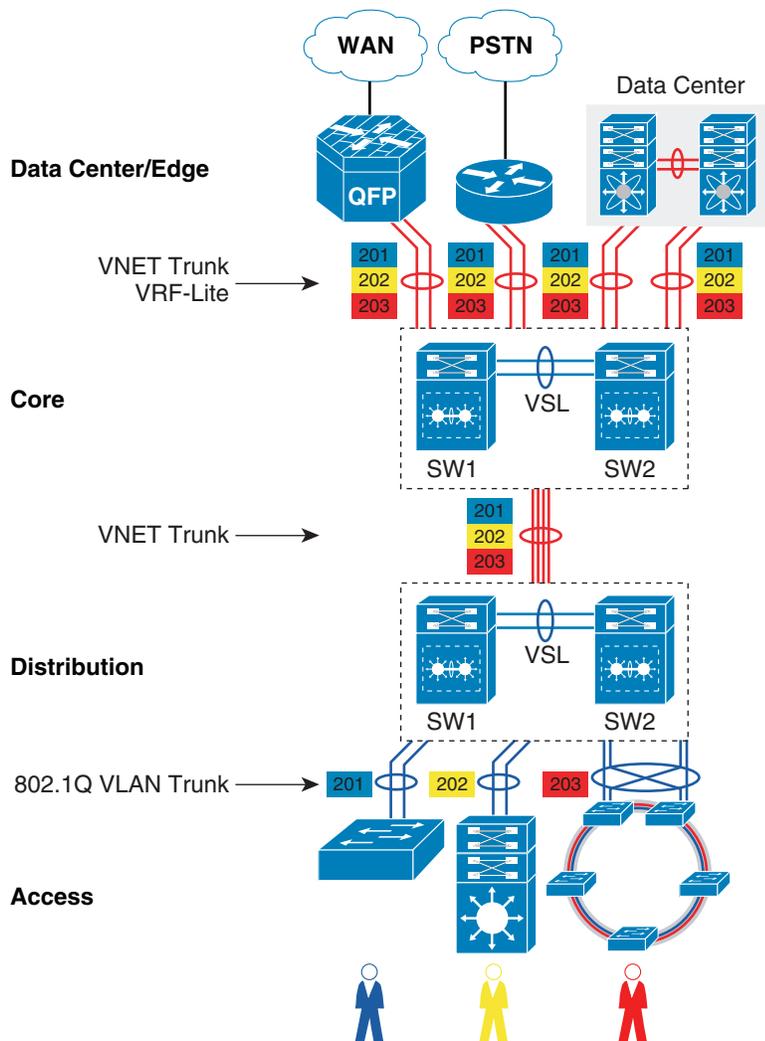
```
cr23-6500-1#config terminal
cr23-6500-1#vrf definition Blue-VRF
cr23-6500-1(config-vrf)# address-family ipv4
```

EVN VNET Tag

The concept of building isolated network paths by defining per-port unique sub-interfaces and assigning different 802.1Q VLAN is changing with Cisco EVN technology. As illustrated in Figure 8, the Multi-VRF 802.1Q implementation remained link-local between two Layer 3 systems within campus, so

it becomes a prerequisite for the network administrator to identify a unique 802.1Q tag that they want to assign for each VN. The VN provisioning and un-provisioning implementation method on a hop-by-hop and link-by-link basis is now simplified with VNET Tag. The VNET Tag can be identified as an IEEE 802.1Q ID that the administrator plans to implement for given VRF, however the provisioning mechanism is now system-wide with global significance. From the network design and operation perspective, this capability provides significant advantages over the traditional solution. Instead of identifying a unique 802.1Q for each VN on a per-hop basis, now with Cisco EVN technology the administrator can define a network-wide unique VNET Tag ID.

Figure 13 Campus VNET Tag Design



On an EVN-capable system the VNET trunk is supported on any physical Layer 3 Ethernet and Layer 3 Port-Channel interface that supports 802.1q encapsulation and it is these VNET trunks that are used to interconnect EVN devices. Rather than defining a new encapsulation header, the 802.1q VLAN ID field was reused to carry the VNET Tag over a VNET Trunk. The Cisco IOS software supports a VNET Tag range of 2-4094. EVN requires a single, network-wide VNET tag for each VRF and supports up to 32 VNETs/VRFs in the range 2-4094.

In such a network design, the VNET ID must be end-to-end available and consistent across all systems. During the EVN planning phase the network administrator must evaluate if an identified VNET Tag for each VRF is available. This best practice facilitates a seamless deployment with a peering switch if it cannot reuse the same VLAN across multiple interfaces. The Cisco Catalyst 6500-E internal hardware, software, and switching architecture broadly differs from other switching platforms. When deploying EVN on the Cisco Catalyst 6500-E with Sup2T, the network administrator needs to perform a two-step evaluation to identify an available VNET Tag:

-
- Step 1** Identify an available VLAN ID from the VLAN database. It is recommended to use unique VNET ID or VRF instead re-using same VLAN that may be use for other purpose. For example, verify VLAN database if VLAN 201 is available for VNET Tag 201 for Blue-VRF:

```
cr23-6500-1#show vlan id 201
VLAN id 201 not found in current VLAN database
```

- Step 2** Identify if the planned VNET Tag is not used by the switch itself or internal mechanisms. The 6500-E system architecture automatically uses several VLANs from the global VLAN database for internal communication. Assigning VNET tags in the range 2-1000 does not create any conflicts with internal VLANs. However if it is necessary to use VNET Tag in the range 1001-4094, then it is recommended to verify the internal VLAN usage table with the command **show internal vlan** and ensure that the intended VNET Tag ID does not conflict with an internal VLAN that is already in use.

```
cr23-6500-1#show vlan internal usage | include 201
!No Result yield confirms VLAN 201 is available to use VNET Tag
```

Once a network-wide available VNET ID is identified, then the network administrator can associate the same VNET Tag into the VRF definition on all EVN-capable systems. Each VN will have a unique VNET Tag, since the VNET Tag is used to segment the data plane between different VNs; Cisco IOS does not allow the reuse of the same VNET Tag for different VRFs. To interoperate with non-EVN-capable peering systems, the network administrator must assign the same VLAN ID as the VNET Tag on manually created sub-interfaces for each VRF. Defining a common VNET Tag and VLAN ID enables successful control and data plane communication over the same 802.1Q VLAN between two systems. [Table 2](#) illustrates enabling VNET Tag and interoperability with non-EVN capable systems:

Table 2 EVN-Capable and Non-EVN-Capable Systems

EVN-Capable Systems	Non-EVN-Capable Systems
<pre>vrf definition Blue-VRF vnet tag 201 ! vrf definition Yellow-VRF vnet tag 202 ! vrf definition Red-VRF vnet tag 203 ! interface Port-Channel 10 no switchport ip address 10.125.11.0 255.255.255.254 vnet trunk !</pre>	<pre>vrf definition Blue-VRF ! vrf definition Yellow-VRF ! vrf definition Red-VRF ! interface Port-Channel 10 no switchport ip address 10.125.11.1 255.255.255.254 ! interface Port-Channel 10.201 description Blue-VRF Sub-Interface ip address 10.125.11.1 255.255.255.254 encapsulation dot1q 201 ! interface Port-Channel 10.202 description Yellow-VRF Sub-Interface ip address 10.125.11.1 255.255.255.254 encapsulation dot1q 202 ! interface Port-Channel 10.203 description Red-VRF Sub-Interface ip address 10.125.11.1 255.255.255.254 encapsulation dot1q 203 !</pre>

The EVN VNET trunk simplicity is derived with new software intelligence in Cisco IOS software. Most of the value between two Layer 3 systems is link local, such as IP addressing, per-protocol stateful connections, security parameters such as authentication, etc. These are commonly deployed scenarios as it eases the VRF sub-interface configuration. In Multi-VRF design, the user must manually create the sub-interfaces, however on EVN-capable systems this complex task is eliminated. Once the interface is configured as VNET Trunk port, then with the EVN software enhancement the sub-interfaces are dynamically generated and automatically associated to VRFs configured in global configuration mode. Even the sub-interface configurations are automatically inherited from parent or physical interfaces.

Depending on network design and scale, the dynamically generated sub-interfaces could be large, which prevents them from being visible with normal operational commands, such as **show run**. The parent or primary interface configuration with VNET trunk can be viewed using normal **show** commands, however the derived configuration on each sub-interface can be verified using **show derived-config interface**, as shown below:

Parent/Primary Interface

```
cr23-6500-1#show running-config interface Port-Channel10
interface Port-channel10
description Connected to Core1
vnet trunk
ip address 10.125.12.0 255.255.255.254
ip pim sparse-mode
ip ospf message-digest-key 1 md5 7 13061E010803
ip ospf network point-to-point
!
```

EVN-generated sub-interface

```
cr23-6500-1#show derived-config interface Port-Channel10.201
```

```

interface Port-channel10.201
  description Subinterface for VNET Blue-VRF
  encapsulation dot1Q 201
  vrf forwarding Blue-VRF
  ip address 10.125.12.0 255.255.255.254
  ip pim sparse-mode
  ip ospf message-digest-key 1 md5 7 13061E010803
  ip ospf network point-to-point
!
cr23-6500-1#show derived-config interface Port-Channel10.202
interface Port-channel10.202
  description Subinterface for VNET Yellow-VRF
  encapsulation dot1Q 202
  vrf forwarding Yelllow-VRF
  ip address 10.125.12.0 255.255.255.254
  ip pim sparse-mode
  ip ospf message-digest-key 1 md5 7 13061E010803
  ip ospf network point-to-point
!
cr23-6500-1#show derived-config interface Port-Channel10.203
interface Port-channel10.203
  description Subinterface for VNET Red-VRF
  encapsulation dot1Q 203
  vrf forwarding Red-VRF
  ip address 10.125.12.0 255.255.255.254
  ip pim sparse-mode
  ip ospf message-digest-key 1 md5 7 13061E010803
  ip ospf network point-to-point
!

```

The VNET Tag ID, associations to an interface, and advanced information such as IPv4 or IPv6 address family information can be verified using the command **show vnet**. To maintain a consistent operational model, the **show vnet** command includes EVN specific information, however most of the other VRF attributes are illustrated in a traditional way. The following sample output shows the Blue-VRF is associated to different physical and logical interfaces, a user-defined unique VNET ID:

```

cr23-6500-1#show vnet detail Blue-VRF
VRF Blue-VRF (VRF Id = 21); default RD <not set>; default VPNID <not set>
  Interfaces:
    Po10.3133          Vl2101
  Lists:
    cr24-4500-2-VRF
  VNET:
    Tag 201
Address family ipv4 (Table ID = 21 (0x15)):
  No Export VPN route-target communities
  No Import VPN route-target communities
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-vrf (No Label)
Address family ipv6 not active

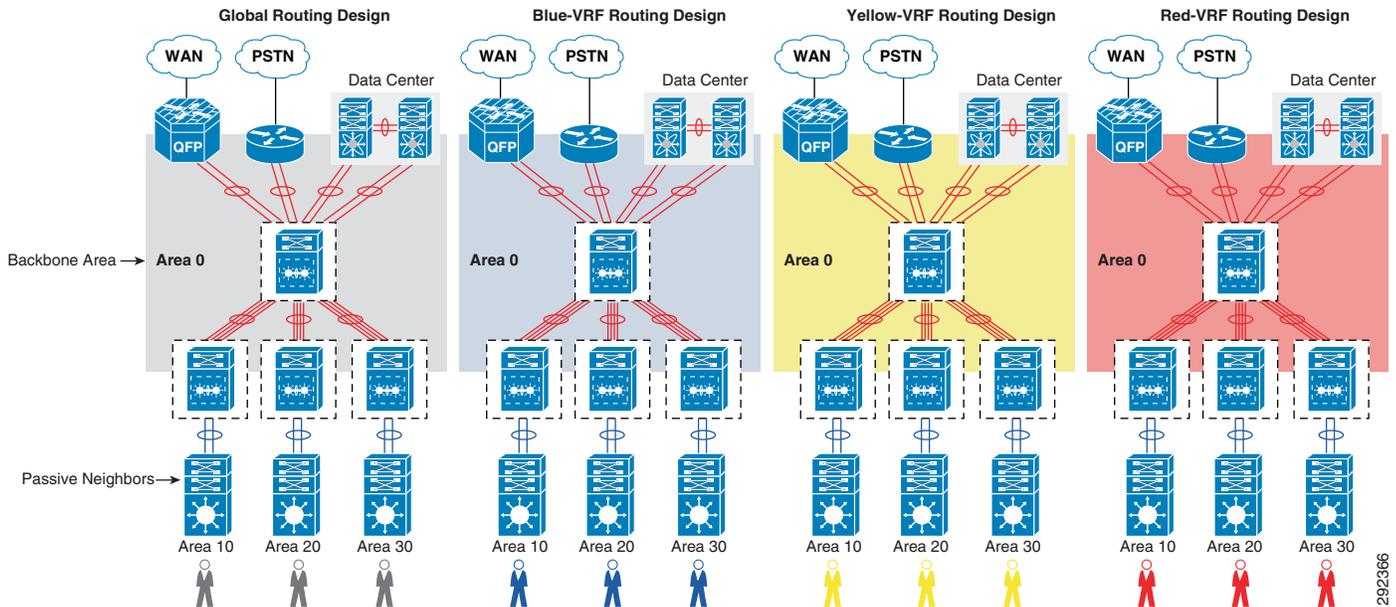
```

EVN Routing Design

The VPN routing infrastructure design in a virtualized network can use the same hierarchy as the global routing infrastructure. To support an isolated control plane, the routing protocol fundamentals and operation differ slightly from the global infrastructure. Each system supporting the end-to-end network virtualization builds a unified control plane, routing database, and tables on a per-VRF basis. This secured routing model assists in decoupling a virtual network into multiple partitions that can coexist but are logically separated to provide end-to-end data protection.

The network administrator can follow the global routing hierarchy and principles to design and build secure virtualized routing for each VRF in the campus core.

Figure 14 EVN OSPF Routing Design



The OSPF virtual routing implementation is the same, whether implementing EVN or not. For logical routing adjacencies and topology database separation, the network administrator must create unique OSPF routing instance on each VRF on each system. The process id uniqueness is required for each OSPF instance; however most of the parameters within the routing process can be implemented based on a common template. The following sample configuration shows the implementation of virtualized OSPF routing adjacencies and applies whether implementing EVN or not.

EVN and VRF-Lite Layer 3 Systems

```

cr23-6500-1(config)#router ospf 3001 vrf Blue-VRF
cr23-6500-1(config-router)# router-id 102.65.1.201
cr23-6500-1(config-router)# log-adjacency-changes detail
cr23-6500-1(config-router)# nsf
cr23-6500-1(config-router)# capability vrf-lite
cr23-6500-1(config-router)# area 10 stub no-summary
cr23-6500-1(config-router)# passive-interface default
cr23-6500-1(config-router)# no passive-interface Port-channel10.201
cr23-6500-1(config-router)# network 10.102.255.0 0.0.0.15 area 10
cr23-6500-1(config-router)# network 10.125.0.0 0.0.255.255 area 0
!
cr23-6500-1(config)#router ospf 3002 vrf Yellow-VRF
cr23-6500-1(config-router)# router-id 102.65.1.202
cr23-6500-1(config-router)# log-adjacency-changes detail
cr23-6500-1(config-router)# nsf
cr23-6500-1(config-router)# capability vrf-lite
cr23-6500-1(config-router)# area 10 stub no-summary
cr23-6500-1(config-router)# passive-interface default
cr23-6500-1(config-router)# no passive-interface Port-channel10.202
cr23-6500-1(config-router)# network 10.102.255.0 0.0.0.15 area 10
cr23-6500-1(config-router)# network 10.125.0.0 0.0.255.255 area 0
!
cr23-6500-1(config)#router ospf 3003 vrf Red-VRF
cr23-6500-1(config-router)# router-id 102.65.1.203

```

```

cr23-6500-1(config-router)# log-adjacency-changes detail
cr23-6500-1(config-router)# nsf
cr23-6500-1(config-router)# capability vrf-lite
cr23-6500-1(config-router)# area 10 stub no-summary
cr23-6500-1(config-router)# passive-interface default
cr23-6500-1(config-router)# no passive-interface Port-channel10.203
cr23-6500-1(config-router)# network 10.102.255.0 0.0.0.15 area 10
cr23-6500-1(config-router)# network 10.125.0.0 0.0.255.255 area 0
!
```

The OSPF adjacencies status for global and VRF routing instance can be verified using a single command in exec mode as shown below:

```
cr23-6500-1#show ip ospf neighbor
```

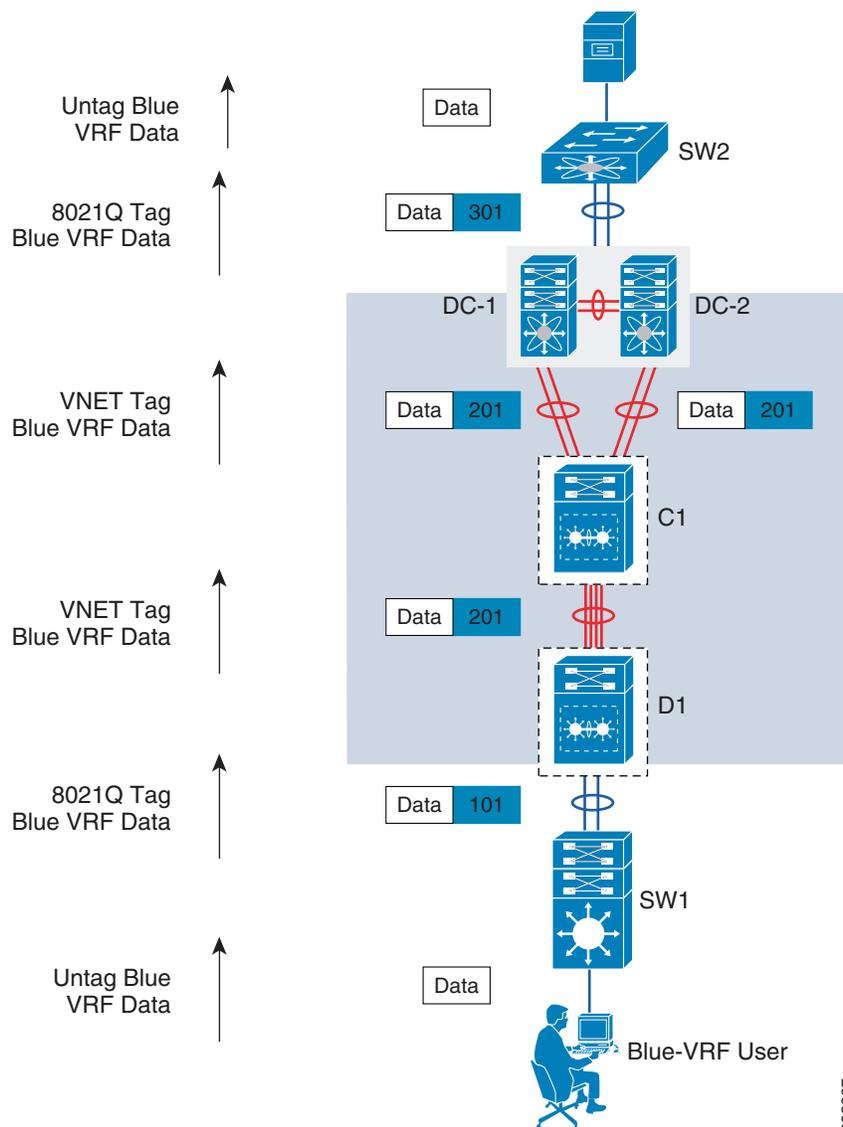
Neighbor ID	Pri	State	Dead Time	Address	Interface
100.65.0.1	0	FULL/ -	00:00:35	10.125.12.1	Port-channel10
100.65.0.201	0	FULL/ -	00:00:33	10.125.12.1	Port-channel10.201
100.65.0.202	0	FULL/ -	00:00:30	10.125.12.1	Port-channel10.202
100.65.0.203	0	FULL/ -	00:00:32	10.125.12.1	Port-channel10.203

EVN Forwarding Design

As in Multi-VRF, VLANs and their associated Layer 3 SVIs are mapped to a VRF so as users connect to a switchport that has had a VLAN statically defined or is dynamically assigned through the process of 802.1x authentication, they are mapped into the appropriate VRF. As user traffic enters an “Edge Interface”, the appropriate VNET Tag is applied to that traffic and forwarded out the VNET Trunk after the routing lookup has been performed within the relevant VRF's routing table. Traffic coming in from a VLAN or interface that is not associated with a VRF is left untagged is forwarded using the global routing table. As with Multi-VRF, each virtual network has its own routing table. EVN supports static, OSPFv2 and EIGRP for edge interfaces and OSPFv2 and EIGRP for VNET Trunks. Refer to [Figure 15](#) to understand how end-to-end virtualization and isolation is maintained based on a per-VRF forwarding information lookup and tag swapping design.

The following provides an example of traffic flow through the EVN topology depicted in [Figure 15](#).

Figure 15 Traffic Flow Through EVN Network Design



1. The untagged user packets arrive on the ingress edge interface of SW1, the access layer switch. The ingress edge interface is associated with a particular VLAN through static or dynamic configuration. Prior to sending data to upstream distribution, the access switch imposes an 802.1Q header with the same ID used on the ingress port.
2. The D1 distribution layer switch should have configured the VLAN and Layer 3 logical SVI interface mapped into appropriate VRF, i.e., Blue VRF. Note that the EVN VNET Tag is used to tag Layer 3 routed frames, so the VLAN ID on Layer 2 trunk interface can be unique.
 - a. Based on the destination IP address on the ingress packet, the D1 does route and egress path lookup within the associated virtual routing table.
 - b. D1 swaps 802.1Q VLAN 101 encapsulation and rewrites the 802.1Q VLAN 201 prior to sending it out to the C1 switch.
3. The C1 core layer switch receives the packet over the VNET Trunk and it uses the VNET Tag 201 to identify the VRF.

- a. Based on the destination IP address on the ingress packet, the C1 does route and egress path lookup within the associated virtual routing table.
 - b. C1 swaps the ingress 802.1Q header with a new egress header, even though the VNET ID is the same. This is part of hardware programming to update TTL information and update the new encapsulation with checksum.
4. Based on load-balancing results, either the DC1 or DC2 switch can receive traffic. In this design it can be between the Layer 2 and Layer 3 network boundary. It can receive the packet on the VNET Trunk or Layer 3 Ethernet interface, but the egress 802.1Q ID could be different.
 - a. DC1 or DC2 receives the routed packet and performs destination IP address and local egress path lookup. The egress VLAN could be different; the interface may be a Layer 3 SVI that is trunked to another Layer 2 switch.
 - b. The DC1 or DC2 switch swaps the VNET Tag from the routed frame with a Layer 2 301 VLAN header prior to sending out to SW2 access layer switch.
5. The SW2 access layer switch receives 802.1Q 301 tag traffic and identifies the destination mac-address and egress interface based on a per-VLAN MAC address table.
6. The server receives the untagged original IP packet.

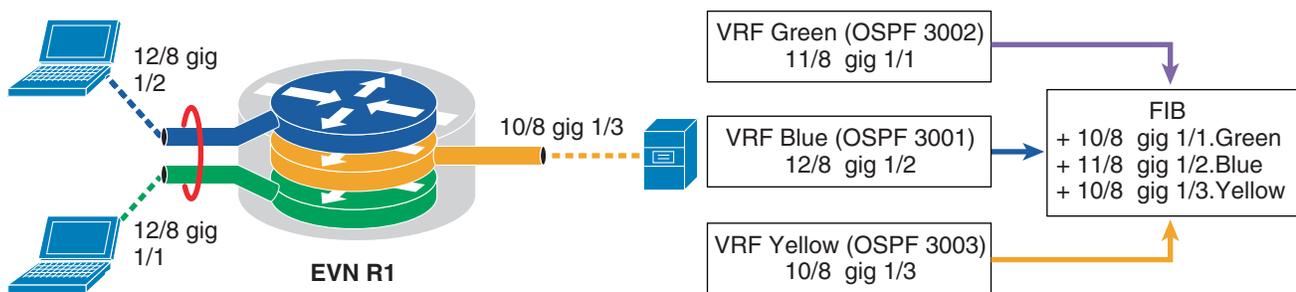
EVN VRF-List

Within EVN it is possible to restrict what traffic can traverse a VNET Trunk based on the VNET Tag. EVN introduces the concept of a “VRF List”. Essentially the VRF List is created and the various tags that are permitted are configured as a “member”. Once the complete list has been defined, the list is applied to the VNET Trunk interface. Traffic for those VNs that have not been configured as a member is not forwarded across that link.

The major difference between EVN and Multi-VRF routing lies in the method used to “leak” routes between virtual networks. With Multi-VRF, as described in [VRF-Lite](#), MP-BGP is required on minimally one or perhaps multiple devices through BGP peering to leak routes for access to shared services. With EVN however, this configuration has been greatly simplified by a process known as route replication.

Through route replication, EVN VNs can share VRF-specific routes with the Routing Information Base (RIB) of the EVN device. Only route prefixes and not next hop information is installed in the RIB. These routes have a “+” sign in front of them to denote a replicated and connected route within the same system. [Figure 17](#) provides a depiction of this process.

Figure 17 EVN Route Replication



The following example shows the configuration required on an EVN device to replicate shared services routes in the Yellow VRF with users in the Green and Blue VRFs. Notice that replication is a bidirectional process requiring replication of the shared services prefixes into the user VRFs and the user prefixes into the shared services VRF.

```
vrf definition Yellow-VRF
address-family ipv4
route-replicate from vrf Green all route-map greenuser-prefix-map
route-replicate from vrf Blue all route-map blueuser-prefix-map
!
vrf definition Green-VRF
address-family ipv4
route-replicate from vrf Yellow all route-map yellowserver-prefix-map
!
vrf definition Blue-VRF
address-family ipv4
route-replicate from vrf Yellow all route-map yellowserver-prefix-map
```

Routes that have been replicated between VNs are not redistributed into the VN’s IGP automatically and hence must be redistributed manually so that remote devices dynamically learn the new routes. The following sample configuration illustrates how to conditionally redistribute Vlan2 and Loopback interface to propagate the imported route to peer EVN or Multi-VRF neighbors:

```
6500E#show ip route vrf Green-VRF connected
<snip>
C + 110.98.101.0/28 is directly connected (Extranet-MVRF-1), Vlan2
C + 110.126.254.254/32 is directly connected (Extranet-MVRF-1), Loopback101

6500E(config)#router ospf 3001 vrf Blue-VRF
6500E(config-router)# redistribute vrf Extranet-MVRF-1 connected subnets route-map
Extranet-MVRF-1-RP
!
6500E(config)#router ospf 3002 vrf Green-VRF
6500E(config-router)# redistribute vrf Extranet-MVRF-1 connected subnets route-map
Extranet-MVRF-1-RP
!
```

```

6500E(config)#route-map Extranet-MVRF-1-RP permit 10
6500E(config-route-map)#match ip address <acl>
!
6500E(config)#access-list <range> permit <source-IP-network> <mask>
6500E(config)#access-list <range> permit <RP-loopback-network> <mask>
!

```

EVN High Availability

In enterprises, network availability and reliability is a top concern as users demand that networks be constantly functional. Any system or network anomalies have a direct impact on end user productivity, end-to-end application performance, and business revenue. As the network becomes more intelligent and secure, the enterprise IT may expand networking services in existing infrastructure to support additional network users and devices. In such network designs, the magnitude of network users and devices in the campus may increase several folds, so campus high availability must be strategically designed by following three key design principles, as illustrated in [Figure 18](#).

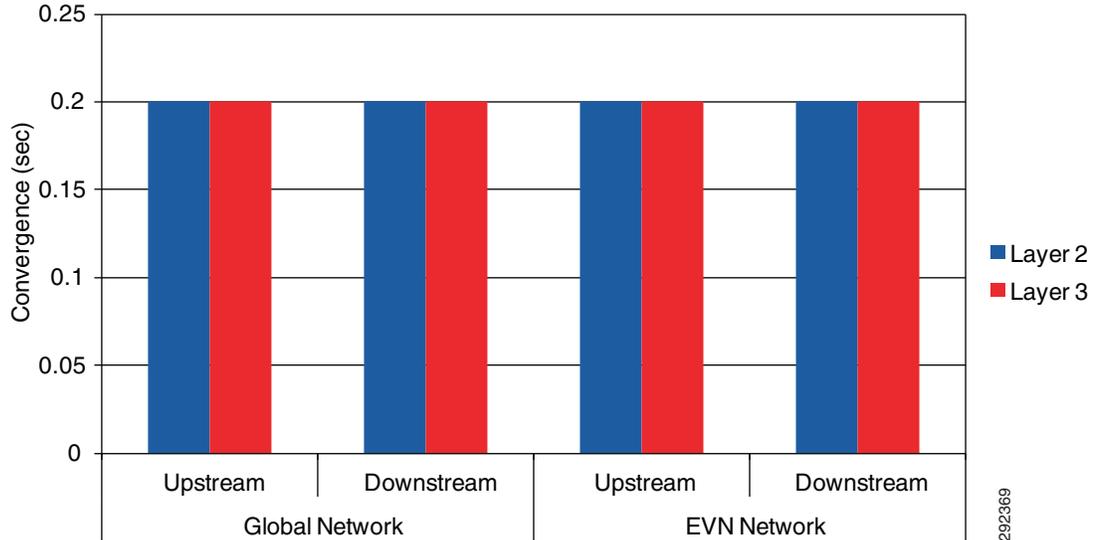
Figure 18 High-Availability Goals, Strategies, and Technologies

Resilient Goal	Network Service Availability		
Resilient Strategies	Network Resiliency	Device Resiliency	Operational Resiliency
Resilient Technologies	EtherChannel/MEC UDLD IP Event Dampening	NSF/SSO Stack Wise	ISSU eFSU

292368

The key high availability advantage of this design strategy is to keep available and operational the campus network infrastructure and implemented services to mitigate the impact to business. In mission critical networks performance is unaffected with this recommended design as it delivers deterministic network recovery during any type of planned or un-planned network outage. The deterministic network operation is achieved by simplifying the campus network design with system and path virtualization technique such as VSS and Multi-Chassis EtherChannel (MEC). The hardware and software design in Cisco Catalyst switches are enhanced to rapidly detect faults and initialize recovery to alternate paths without requiring decisions from the complex and slow upper layer routing infrastructure. As the campus systems and network become intelligent, network recovery time from fault detection to re-programming hardware with alternate paths remains consistently in the sub-seconds.

Implementing virtualization in the campus may increase the control and data plane factor, however the resiliency that Cisco Catalyst switching offers does not change its fundamental design and principles. Since such designs offer protocol-independent network recovery, the end-to-end application recovery remains consistently sub-second for data flowing through the global or virtual routing infrastructure. The analysis shown in [Figure 19](#) confirms that inducing a fault at the Layer 2 or Layer 3 domain does not introduce substantial difference as the recovery mechanism is hardware driven rather than protocol driven.

Figure 19 *Link Failure HA Analysis*

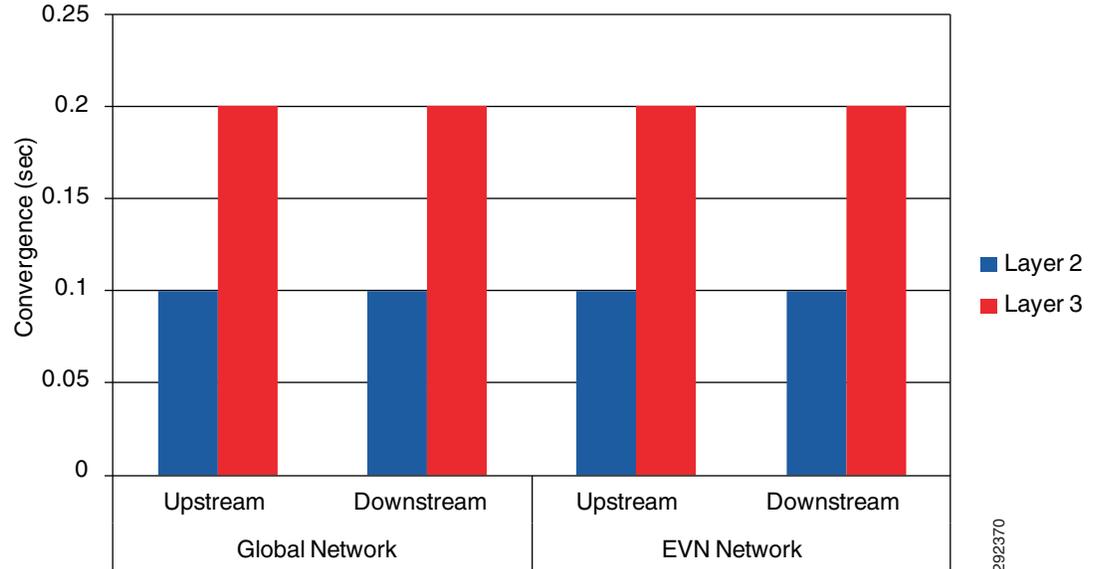
The Stateful Switch Over (SSO) and Non-Stop-Forwarding (NSF) operation in system is independent of services implemented on Cisco Catalyst switches. When dual supervisor modules are deployed in redundant configuration mode in Cisco Catalyst 6500-E and 4500-E systems, the control plane and forwarding information of each VN is fully synchronized to the standby supervisor module. The NSF helper function is by default enabled to support graceful recovery for a peering-capable switch, however the network administrator must explicitly implement NSF capability on each VRF selected routing protocol, as shown below:

```
6500E(config)#router ospf 3001 vrf Blue-VRF
6500E(config-router)# nsf
!
6500E(config)#router ospf 3002 vrf Green-VRF
6500E(config-router)# nsf
!
```

The NSF capability can be verified using the command **show ip ospf <pid>**, as shown below:

```
6500E(config)#show ip ospf 3001
<snip>
Non-Stop Forwarding enabled
IETF NSF helper support enabled
Cisco NSF helper support enabled
```

During supervisor switchover the new active supervisor initializes the graceful recovery process with global peers and each VN routing peer. The graceful recovery process exchanges control plane reset information and routing database information to maintain synchronization. The complete control plane recovery time may depend on the number of VNs implemented and the routing database size, however with NSF the forwarding plane remains completely intact. Network availability and recovery from catastrophic failure is un-compromised because of the innovative Cisco Catalyst switching architecture and the best practices used to design the campus infrastructure. During supervisor failure, network recovery for end-to-end applications running in the global and virtual network remains within a sub-second boundary. The internal fabric and switching architecture of the modular Cisco Catalyst 4500-E and 6500-E are different, hence the network recovery and convergence time may differ by a few milliseconds. [Figure 20](#) compares the supervisor switchover convergence times for global and virtual traffic and for Layer 2 domain in the access layer and Layer 3 routing domain in the aggregation or core layer.

Figure 20 Supervisor Failure HA Analysis

Summary

EVN is an evolutionary Cisco technology built upon Multi-VRF allowing complete backward compatibility while introducing new constructs to greatly simplify network virtualization within campus environments. With Multi-VRF the need to deploy MPLS and Label Distribution Protocol, along with an MP-BGP control plane in the infrastructure, was virtually eliminated with the exception that if, as in most cases, shared services were required, a pocket of MP-BGP was still required to leak routes between VNs. With EVN and route replication, this requirement has been removed to provide a solution entirely free of MPLS and MP-BGP. Due to the fact that EVN is not only backward compatible with Multi-VRF as well as MPLS, it is anticipated that EVN will quickly become one of the most popular strategies when network virtualization is required.

For further information about EVN, see: <http://www.cisco.com/go/evn>.

For configuration information, see:

<http://www.cisco.com/en/US/docs/ios-xml/ios/evn/configuration/xe-3s/evn-xe-3s-book.html>.

For the EVN Command Reference, see:

<http://www.cisco.com/en/US/docs/ios-xml/ios/evn/command/evn-cr-book.html>.

MPLS

Although first deployed as a service provider technology, MPLS has increasingly found its way into enterprise networks as an extremely scalable and mature technology for providing network virtualization. One of the primary benefits of deploying an MPLS infrastructure is the ability to provide dynamic any-to-any mesh connectivity through the use of Multi-Protocol BGP and the Label Distribution Protocol (LDP).

MPLS should not be viewed as a solution intended for WANs alone. Many enterprise organizations have successfully deployed MPLS throughout the distribution and core of campus networks as well as data center, WAN, and Internet edges. For example, many organizations make use of MPLS in the distribution

and core of their networks to provide either guest access or, for example, though peering with the WAN edge, traffic isolation for PCI/Point of Sale applications, Financial Teller 21 applications, and even SCADA control traffic found in the utilities.

Another application for the use of MPLS as previously discussed may be to provide transport over an IP infrastructure outside of the control of an organization, thus extending and interconnecting virtual networks built using EVN within various campuses or large branches. This is possible inasmuch as the VRF, as the fundamental building block of EVN, serves the same purpose within MPLS.

There are literally volumes of information dedicated to the subject and advanced capabilities of MPLS, so it is beyond the scope of this document to provide anything more than a high level overview of this technology. See [References](#) for links to information about MPLS.

MPLS Fundamentals

Multi-Protocol Label Switching, as its name implies, is built on the concept of switching packets or cells based on a label as opposed to using the IP/CEF forwarding table exclusively. That is not to say that the IP forwarding table is no longer needed, as nothing could be further than the truth, just that when an MPLS label switching device receives a frame with a label imposed, the forwarding path has already been determined through a process, to be discussed later, and is able to switch that frame out the appropriate interface based on the label. Refer to [Figure 21](#) which identifies the MPLS label format.

Figure 21 *MPLS Label Format*

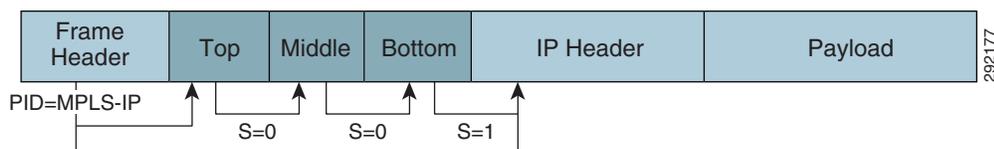


The structure is as follows:

- MPLS Label—20-bit field used for label switching the packet and is replaced at every hop in the MPLS network
- EXP—3-bit field that is used to indicate the class of service (CoS) of the MPLS packet (similar to the CoS field in Ethernet frames)
- S—Bit used to indicate the bottom of the stack when more than one MPLS label is imposed on the packet as seen subsequently in the case in the MPLS VPN scenario
- TTL—8-bit time-to-live value (having the same functions of loop detections as the homonymous IP field)

The MPLS label is placed after the Layer 2 headers for a packet. Notice that a packet can have multiple MPLS labels prepended to it; this is referred to as the label stack. Each MPLS label has a specific meaning for the node that pushed the label onto the packet and the node that pops that label from the stack. The Label Switch Routers (LSR) in the network forward packets based only on the outermost label. The lower labels are taken into account only when they become the outermost label after the previous outermost label has been popped. MPLS labels are pushed onto packets starting with the original frame and additional labels are added on top of the outermost label. MPLS labels are popped starting with the outermost label, the last one pushed onto the label stack. Refer to [Figure 22](#).

Figure 22 *MPLS Label Stacking*

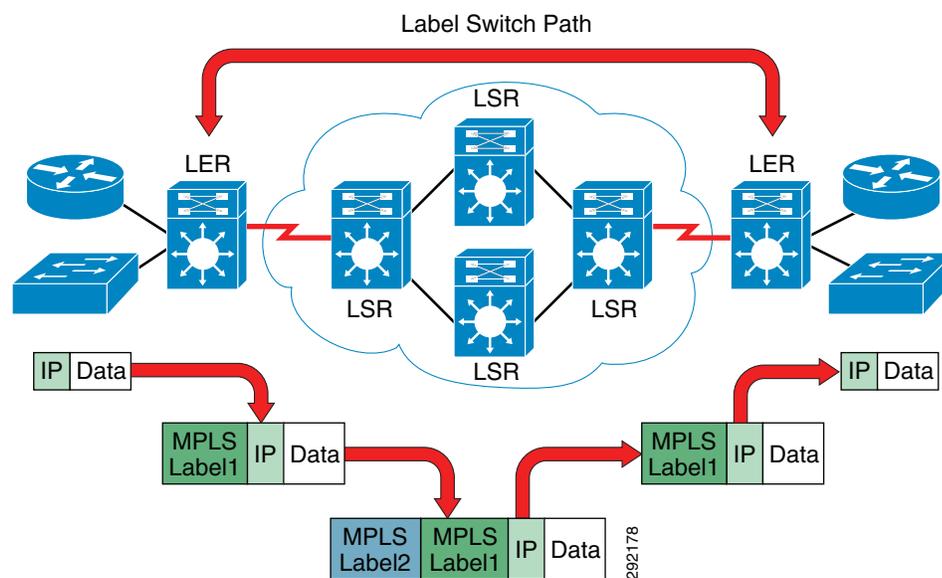


Note the following from [Figure 22](#). The Protocol Identifier (PID) in the Layer 2 header specifies that label(s) follow the frame header. Within each label, there is a Bottom of Stack Bit that indicates whether the next header is a label or an IP header. If the stack bit is 1, the next header is an IP header.

MPLS Device Roles

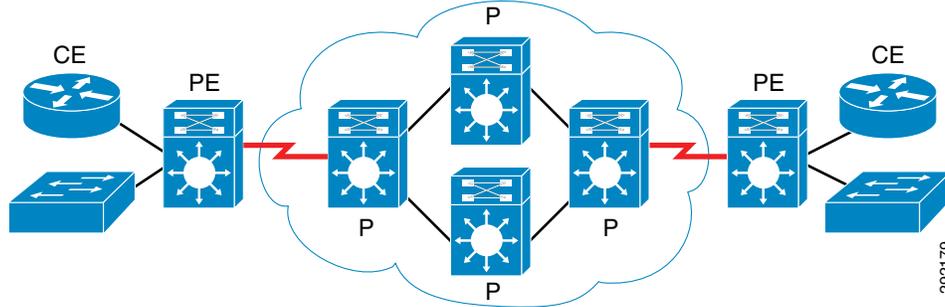
The Label Switch Router (LSR) is any router or switch that runs the Label Distribution Protocol (LDP) and is capable of forwarding packets based on an imposed label. Several different types of LSRs exist, but the one that is most important to this discussion is the “Edge-LSR”, more commonly known as the Label Edge Router (LER). The LER, as its name implies, is connected at the edge of the MPLS network and as such has non-MPLS neighbors attached on one side and pure LSRs attached within the MPLS network on the other. The LER is responsible for label “imposition” (“push” action), which is the act of prepending a label or stack of labels to an IP packet entering the MPLS network, as well as label “disposition” (“pop” action), which is the act of removing the last label at the egress of the MPLS network. Refer to [Figure 23](#).

Figure 23 *MPLS LSR and LER*



In addition to the terms LER and LSR, three other more commonly used terms are used to classify the different equipment roles in an MPLS network. These are CE, PE, and P routers. The MPLS legacy started with service providers and so in reviewing the naming convention of the various roles comprising an MPLS network, one can easily see that influence. Refer to [Figure 24](#).

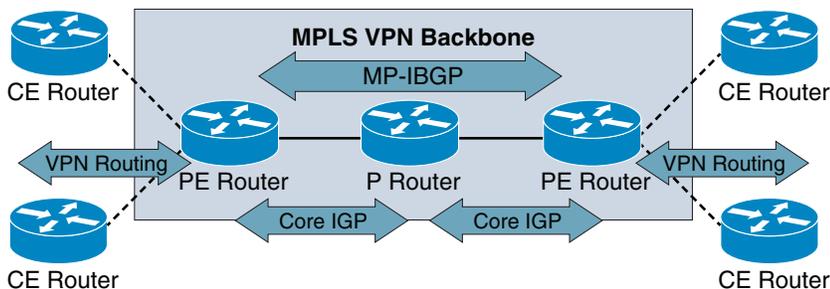
Figure 24 *MPLS CE, PE, and P roles*



The roles are as follows:

- **Provider (P) router**—These are the devices building the core of the MPLS-enabled network. Their main function is to label switch traffic based on the most external MPLS tag imposed to each packet and for this reason are also referred to as label switching routers (LSRs). P routers run a backbone IGP with the PE routers and exchange information about global subnets (core links and loopbacks).
- **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 edge label switching routers (Edge-LSR) or LERs because they sit at the edge of the MPLS-enabled network. Refer to [Figure 25](#). LERs:
 - Exchange VRF routes with CE routers via per-VRF routing protocols.
 - Exchange core routes with P routers and PE routers via core IGP.
 - Exchange VPNv4 routes with other PE routers via MP-iBGP sessions. The VPNv4 routes are representative of the VRFs for the virtual networks and are advertised within MP-BGP address families configured at the PE.
- **Customer edge (CE) router**—This is traditionally the network device at the customer location that interfaces with the service provider. In [Figure 25](#), the CE routers represent the customer remote locations that need to be interconnected via the MPLS service provider network. In an enterprise network, these could just as easily represent Layer 3, non-MPLS devices attached to the PE and forming an Layer 3 IGP adjacency with the PE.
- The switches underneath the CE routers are there to demonstrate that a Layer 2 switch could be easily connected to the PE router. This is the typical deployment model in many enterprises today.

Figure 25 *Routing Protocols in MPLS*

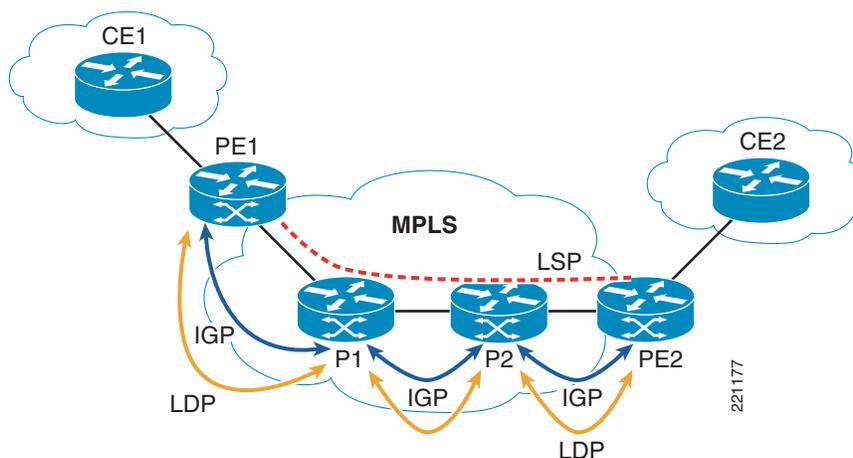


From a control plane point of view, an MPLS-enabled network uses two separate protocols: first, an IGP running in the core of the network and providing connectivity between the various network devices and second, a Label Distribution Protocol (LDP) providing a standard dynamic methodology for hop-by-hop label distribution in the MPLS network. LDP is enabled on interfaces and works by assigning labels to

routes that have been chosen by the underlying IGP routing protocol. The resulting labeled paths, shown in [Figure 26](#) and called label switched paths (LSPs), forward label traffic across an MPLS backbone to particular destinations.

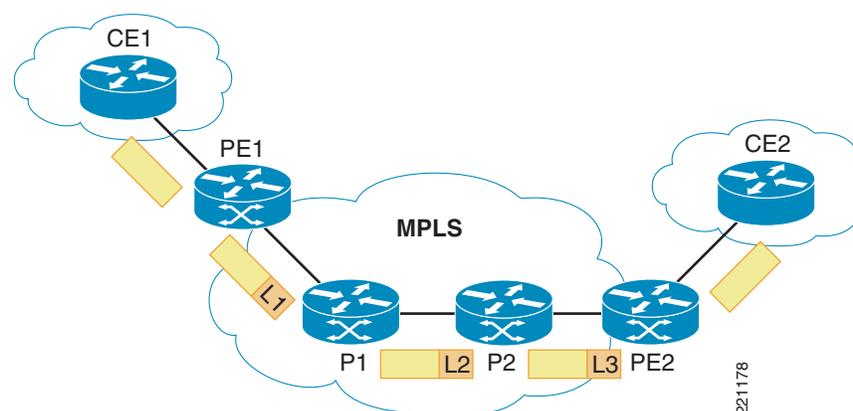
When an interface is enabled for label switching (as shown in the previous section), the LDP process starts and tries to discover other MPLS-enabled neighbors (either PE or P devices) by sending LDP hello packets. When a neighbor has been discovered, an LDP session is established with it by setting up a TCP session on the well-known port 646. As a consequence, IP connectivity is required between neighbors to be able to successfully establish the LDP session. After the LDP session has been established, keepalive messages are exchanged between the neighbor devices (by default every 60 seconds), as highlighted in [Figure 26](#).

Figure 26 *MPLS Control Plane*



From the point of view of data forwarding, traffic that needs to be sent between remote customer sites is label-switched along the LSP, as shown in [Figure 27](#).

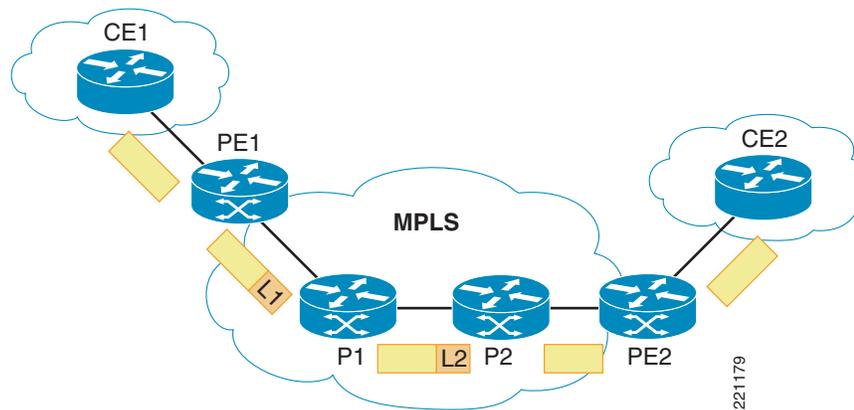
Figure 27 *Label Switch Path*



Each device along the LSP switches the traffic based on the incoming MPLS label; a new tag is imposed before the packet is sent to the next device. Notice that the behavior shown in [Figure 27](#) may be in reality slightly different because of functionality called Penultimate Hop Pop (PHP). By default, the egress PE device explicitly informs the neighbor P not to tag packets directed to it, so that the PE can switch the

packet based only on IP information without having to do a double lookup (the first one for the MPLS tag and the second one for the IP information). Figure 28 shows the same network above when using PHP.

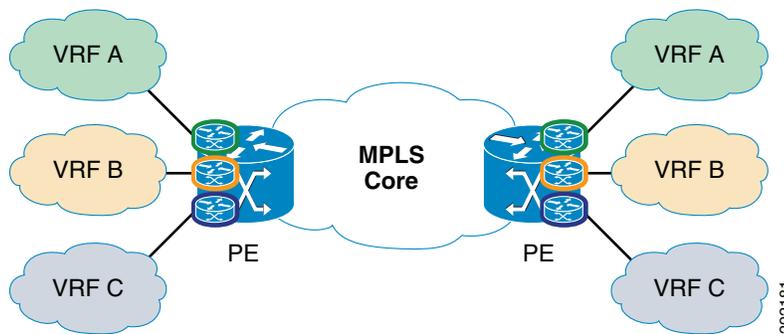
Figure 28 Penultimate Hop Popping (PHP)



MPLS VPN

The VRF, as one would expect, is at the very foundation of MPLS VPN and has been discussed earlier in [Network Virtualization Technologies](#). As shown in Figure 29, defining distinct VRF instances on each PE device allows for 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 outermost MPLS label.

Figure 29 MPLS VPN



The routing protocols (usually OSPF and EIGRP in a campus environment) running in the global table learn routes from the routing peers and install those routes into the routing database of the PE. After the routing database has been populated, the CEF process takes the information in the database and populates the forwarding table (FIB).

The IGP used in the global table has a double functionality. On one side, it allows the establishing of MP-iBGP sessions between the PE devices deployed at the edge of the MPLS domain and the exchange of MPLS labels through a specific LDP protocol. At the same time, it is also used to allow network connectivity to the entities that remain in the global table. As already mentioned, the current recommendation is to use virtual networks only for specific purposes. This means that most of the

internal enterprise traffic still remains switched in the global table. This represents the first differentiation from the SP-like MPLS VPN deployment, because in that case the global table is usually used to provide only PE-PE connectivity and does not extend to the edge of the network but only remains in the core.

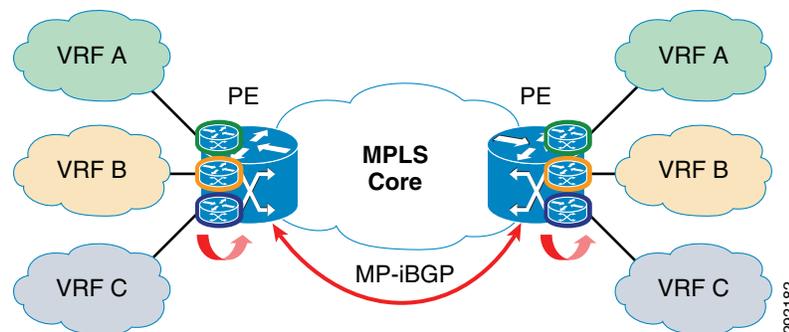
In addition to the IGP, after MPLS is enabled on the device, there is an additional control plane represented by the Label Distribution Protocol that can be thought of as a routing protocol for MPLS because it provides neighbor devices with information about MPLS labels. The label information received from the neighbors is loaded into the label database. Once again, the CEF process running on a device takes that information and builds a second label database known as the LFIB. Notice in the example below that this data structure contains v4 routes, v6 routes, and MPLS forwarding entries, and those MPLS forwarding entries basically form part of it.

```
cr20-6500-1#sh mpls forwarding-table
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag	Outgoing interface	Next Hop
16	Pop tag	192.168.100.19/32	0		Te1/1	10.122.5.30
17	Pop tag	10.122.5.10/31	0		Te1/2	10.122.5.26
	Pop tag	10.122.5.10/31	0		Te1/1	10.122.5.30
18	Pop tag	10.122.5.6/31	0		Te1/2	10.122.5.26

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. As shown in [Figure 30](#), for this to work an MP-iBGP session needs to be established between all the PE devices (in a fully-meshed fashion).

Figure 30 Control Plane for MPLS VPN



In an MPLS VPN design, the exchange of VPN routes is achieved by using an additional control plane element called Multi-Protocol BGP (MP-BGP), which is an extension of the existing BGP-4 protocol. As described here, MP-BGP is introduced only as an overlay protocol to provide the capabilities for exchanging VPN routes. Very large networks can be deployed as separate autonomous systems (AS) and in such scenarios the use of BGP may be required also to connect these separate AS and exchange global table routes. The recommended design discussed here is instead constituted by a single AS and an IGP deployed end-to-end, so that there is no requirement for BGP in global table.

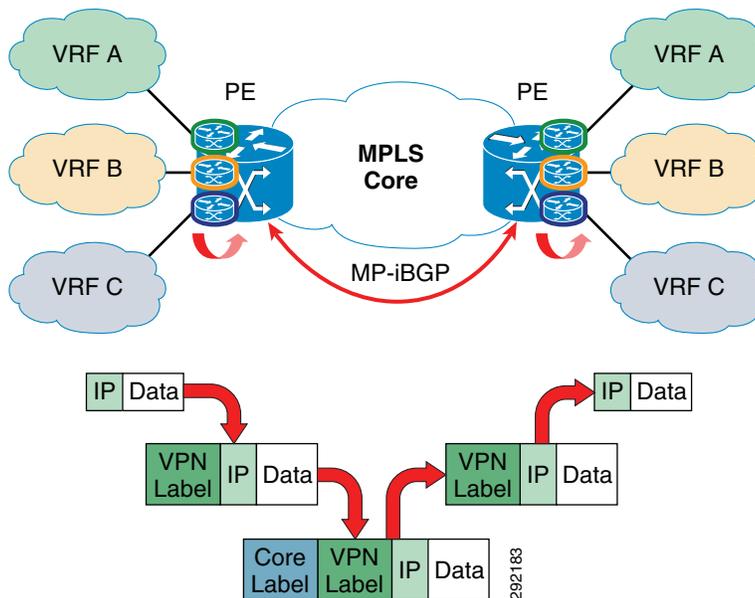
As a consequence, MP-BGP needs to be configured only between the PE devices, because they are the only ones containing VPN routes in the various VRF routing tables. A direct consequence of the fact that the main MPLS VPN strength is to provide any-to-any connectivity inside each defined VPN is the requirement for the PE devices to establish MP-iBGP connections between them in a fully-meshed fashion. By deploying route reflectors, it is possible to relax this requirement, thus improving the scalability of the overall solution as depicted in [Figure 32](#).

MP-iBGP is required within the MPLS VPN architecture because the BGP updates exchanged between PE devices need to carry more information than just an IPv4 address. At a high level, the following three pieces of information are critical to the MPLS VPN functionality and are exchanged through MP-iBGP:

- **VPNv4 addresses**—Unique 96 bit address prefixes defined in the context of each VPN that needs to be communicated between the various PE devices to provide connectivity inside each VPN. A VPNv4 address is achieved by concatenating together the IPv4 prefix and a 64-bit entity called a route distinguisher (RD). A unique, arbitrarily defined RD needs to be used for each VRF defined on the PE device. The RD uniqueness contributes to the uniqueness of each VPNv4 prefix, allowing the support of overlapping IPv4 prefixes between separate VPNs.
- **MPLS VPN label information**—Each PE allocates a specific MPLS label for each defined VPN prefix. This is the more internal label that is pushed in each MPLS packet before sending it to the MPLS core and is used by the receiving PE to determine in which VPN to route the packet.
- **Extended BGP communities**—The most important of these extended communities is called the route target and represents a 64-bit value that is attached to each BGP route. The value of the route target determines how the VPN routes are exported and imported into each VPN. Basically, every VPNv4 route received by a PE may have one or more route targets associated to it. Depending on the route targets configured locally on the receiving PE for each VRF, the route is either imported or ignored for that specific VRF. Likewise, within VRF definition, the export command identifies what routes are exported from a given VRF. Using route targets provides great flexibility to provision many different VPN topologies.

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. This concept is highlighted in [Figure 31](#).

Figure 31 Data Plane for MPLS VPN

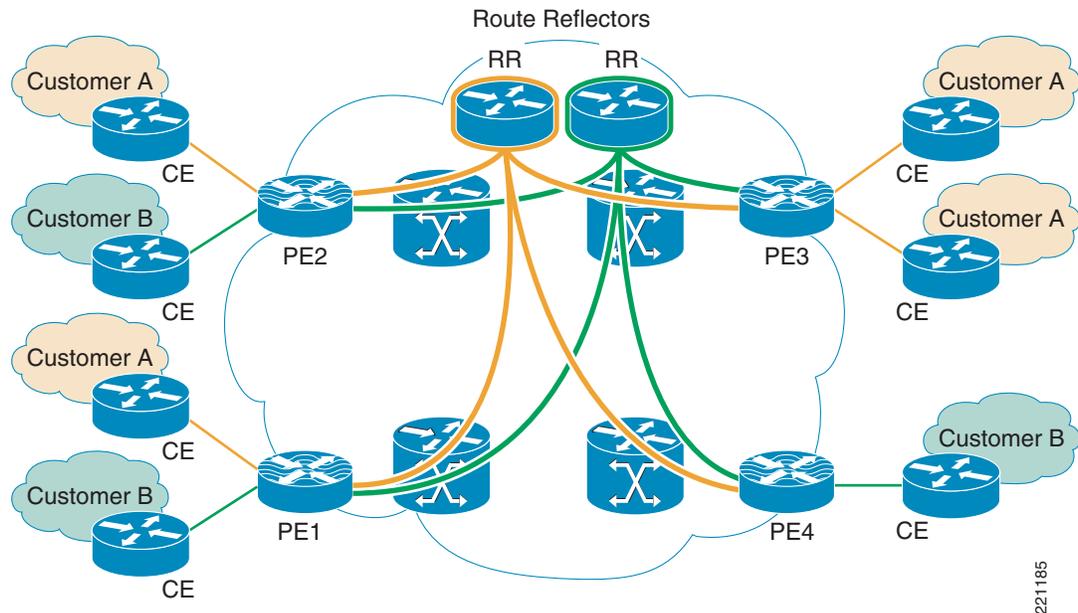


As shown in [Figure 31](#), when the IP packet is received at the ingress PE, the first VPN label is imposed on it. The information on what VPN label to apply has been received from the egress PE (destination) via MP-iBGP. Before sending the packet to the MPLS core, the ingress PE must also impose a second

tag (the most external one), which is used to label switch the packet along the LSP connecting the ingress PE to the egress PE. When the egress PE receives the packet, it is able to look at the VPN label and based on that specific label, send the traffic in the proper VPN.

As just mentioned, the last element that needs to be considered for an MPLS VPN deployment is the route reflector (RR). Because MP-iBGP sessions need to be established between the different PEs defined at the edge of the MPLS network, Cisco usually recommends not deploying a full mesh of iBGP connections but instead using several devices as route reflector routers as shown in Figure 32.

Figure 32 Deployment of Route Reflectors



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.

Summary

MPLS today still provides the most flexibility and scalability for the deployment of network virtualization. The main drawback has been the requirement for MP-BGP and LDP and the ensuing skills required to efficiently design, implement, and maintain the network. However, once the initial design of the network has been completed and implemented, it becomes a routine process to add to or change the operational characteristics of the network. From a compliance perspective where path isolation may be mandated, MPLS is a standardized and accepted approach to securely isolate sensitive traffic from insecure networks.

For additional information that provides much greater detail as well as best practices for configuring MPLS, refer to the Network Virtualization Path Isolation Cisco Validated Design (CVD): http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/PathIsol.html#wp277683.

Additionally, in the following section are URLs to several Cisco Press books that cover MPLS in great detail.

References

Borderless Campus Design and Unified Access

- For design guidance around building scalable and highly available Campus Networks, an excellent reference is the Borderless Campus 1.0 Design Guide (CVD)
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cOverall_design.html
- Additional guidance for mobility and security can be found in the Unified Access Design Guide (CVD)
http://www.cisco.com/en/US/partner/solutions/ns340/ns414/ns742/ns815/landing_unified_access.html

Cisco Validated Designs and Design Zone

http://www.cisco.com/en/US/netsol/ns742/networking_solutions_program_category_home.html

Authentication and Secure Access

- Cisco TrustSec Solutions Page on Cisco.com
<http://www.cisco.com/en/US/netsol/ns1051/index.html>
- TrustSec Phased Deployment Configuration Guide
http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/Phased_Deploy/Phased_Dep_Guide.html#wp392175

Network Virtualization Design Guidance

- Network Virtualization—Path Isolation CVD
http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/PathIsol.html#wp277683
- Network Virtualization—Services Edge CVD
http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/ServEdge.html

Easy Virtual Networking

- EVN Reference information on Cisco.com
<http://www.cisco.com/go/evn>
- EVN Configuration Guide
<http://www.cisco.com/en/US/docs/ios-xml/ios/evn/configuration/xe-3s/evn-xe-3s-book.html>
- EVN Command Reference
http://www.cisco.com/en/US/netsol/ns742/networking_solutions_program_category_home.html

Multi-Protocol Label Switching

- MPLS Reference materials on CCO
http://www.cisco.com/en/US/products/ps6557/products_ios_technology_home.html
- MPLS Design Guides on CCO
http://www.cisco.com/en/US/tech/tk436/tk428/tech_design_guides_list.html

MPLS Reference Books

- MPLS Fundamentals
<http://www.ciscopress.com/bookstore/product.asp?isbn=1587051974>
- MPLS and VPN Architectures
<http://www.ciscopress.com/bookstore/product.asp?isbn=1587050021>
- MPLS and VPN Architectures II
<http://www.ciscopress.com/bookstore/product.asp?isbn=1587051125>
- MPLS and Next-Generation Networks: Foundations for NGN and Enterprise Virtualization
<http://www.ciscopress.com/bookstore/product.asp?isbn=1587201208>