



Network Virtualization for the Campus

As the demands placed on campus networks have grown in complexity, so has the need for scalable solutions to separate groups of network users and resources into logical partitions. Virtualization of the network provides multiple solutions for centralizing services and security policies while preserving the high-availability, manageability, security, and scalability benefits of the existing campus design. To be effective, these solutions must address the three primary aspects of network virtualization: access control, path isolation, and services edge. Implementing these solutions enables network virtualization to coalesce with the Cisco Systems® Service-Oriented Network Architecture (SONA), creating a solid framework for enterprises to migrate to an Intelligent Information Network.

Utilizing Network Admission Control (NAC) and the IEEE 802.1x protocol, a SONA network delivers identity services, which provide optimal access control. After users gain access to the network, three solutions for path isolation—generic routing encapsulation (GRE) tunnels, Virtual Route Forwarding (VRF)-lite, and Multiprotocol Label Switching (MPLS) VPNs—preserve the benefits of today’s campus design while introducing the capability of separating the network into secure virtual networks by overlaying partitioning mechanisms onto the existing LAN. These solutions address the problems associated with deploying services and security policies in a distributed manner. Finally, the centralization of shared services and security policy enforcement dramatically reduces the capital and operational expenses of maintaining different groups’ security policies and services within a campus. This centralization also enables consistent policy enforcement throughout the campus.

Whatever their size or security needs, enterprises today can enjoy the benefits of a virtualized campus network with many closed user groups, all on a single physical network.

CHALLENGE

The ability to provide flexible connectivity options while keeping closed user groups (groups of users and resources isolated from non-group members) is vital in today’s campus networks. Network virtualization—giving multiple groups access to the same physical network while keeping them logically separate to a degree that they have no visibility into other groups—is a requirement that has challenged network managers for many years.

In the 1990s, campus LANs were characterized by broad deployments of Layer 2 switches. Campus LAN design has evolved significantly since the introduction of Ethernet switching. Many technological improvements and a wealth of expertise, combined with evolving customer requirements, have shaped today’s network design philosophy. With the growth of campus LANs has come the need to partition the network more effectively—by users, location, or function. All different groups need to retain their privacy while sharing the same physical network. For end users on the network, the experience should be that of using totally separate physical networks that can be securely interconnected.

Initially, virtual LANs (VLANs) and the Spanning Tree Protocol offered a mechanism to divide the LAN into virtual networks to address the need for separate workgroups within a common network. This solution was effective and secure, but it did not scale well, nor was it easy to manage as these campus LANs grew to encompass a large number of switches and global locations.

Compelling Factors

Many factors promote the requirement to create closed user groups. The following are just a few examples:

- **Varying levels of access privileges within an enterprise:** Almost every enterprise needs solutions for granting different levels of access to customers, vendors, and partners as well as employees on the campus LAN. A visiting vendor, for instance, needs only to connect to the Internet while on campus, and the organization needs to help ensure that the vendor cannot gain access to corporate resources.
- **Regulatory compliance:** Some businesses are required by laws or rules to separate segments of a larger organization. For example, in a financial company, banking needs to remain separate from trading. Office buildings commonly require the separation of different departments, such as human resources and customer service in a corporate setting, or a police department and a fire department in a municipal office building.
- **Network simplification for very large enterprises:** In the case of very large campus networks such as airports, hospitals, or universities, the need for security between different groups or departments has in the past required the building and management of separate physical networks, an undertaking that is costly and difficult to manage.
- **Network consolidation:** In mergers and acquisitions, there is often a need to integrate the acquired company's network expeditiously.
- **Outsourcing:** As outsourcing and offshoring proliferate, subcontractors must demonstrate absolute isolation of information between clients. This is especially critical when a contractor services competing companies.
- **Enterprises providing network services:** Often retail chains support kiosks for other companies or on-location Internet access for patrons; similarly, airports serving multiple airlines and retailers can provide isolated and common services with a single network.

The need to create closed user groups can be mapped to enterprise networks of various vertical markets. Some examples for different verticals are listed in Table 1.

Table 1. Application Examples of Network Virtualization in Individual Verticals

Vertical	Examples of Cases for Network Virtualization
Manufacturing	Production plants (robots, automation of production environment, and so on), administration, sales, video surveillance.
Finance	Trading floors, administration, mergers.
Government	Shared buildings and facilities supporting different departments. In some countries the law mandates separate networks between such departments.
Healthcare	General trend toward hotel service with medical treatment. Separation among medical staff, magnetic resonance imaging (MRI) and other technical equipment, Internet access for patients, media services such as radio and television for patients.
Commercial Real Estate: Multibusiness Campus	Some resources are shared among groups. Multiple companies on the same campus where different buildings belong to different groups, but all rely on the same core and Internet access. Building automation is administered by the owner and spans across all buildings.
Retail	Kiosks, public wireless LAN (PWLAN) in branches, RF identification, WLAN devices (for example, older WLAN barcode readers that do not support any WLAN security).
Education	Separation among students, professors, administrators, and external research groups. Alternatively, individual departments that spread across multiple buildings might require access to their respective server areas. Some resources (Internet, e-mail, and news, for example) might be shared or accessed through a services zone. Building automation, too, must be separated.

Campus VLANs

In years past, when Layer 2 switches defined the campus network environment, network virtualization was achieved with the use of VLANs. A separate VLAN could be configured for every closed user group on the network. These VLANs would then be configured to span the entire campus, an approach commonly referred to as *campuswide* VLANs. However, although this method was reasonably simple to implement, it introduced a number of complications, including:

- Limited network scalability resulting from the use of the Spanning Tree Protocol across a large number of switches
- Impaired client and network performance resulting from increases in broadcast and multicast traffic
- Complexity of troubleshooting
- Risk of problem propagation

Among these problems, the Spanning Tree Protocol represented the greatest limitation, from a Layer 2 topology point of view. The risk of a Layer 2 loop, or broadcast storm, grew in parallel with the number of bridges in a Layer 2 domain. Furthermore, the network diameter could become a gating factor in midsize to large topologies.

As the number of interfaces in a VLAN increased, the frequency of broadcasts also increased. These broadcasts resulted in higher CPU load for client and network devices, as well as reduced performance for applications. The Spanning Tree Protocol goal of providing a loop-free topology inherently prevented networks from having multiple active paths between any two local destinations and therefore limited the available network bandwidth. At the network edge, this characteristic did not present a major problem, but the core of the network could become a bottleneck.

Troubleshooting large Layer 2 topologies required a significant amount of technical expertise and could be time-intensive. When a Layer 2 loop occurred, it not only resulted in the loss of client connectivity, but also could affect remote network administration.

Layer 3 Campus

In recent years, with the advent of multilayer switches capable of performing routing tasks, scalability and availability gained importance as design criteria. Today's campus LAN combines the services of Layer 2 switched and routed platforms to achieve the best-possible modularity, scalability, and high availability.

The use of Layer 3 switching in the core and distribution layers can eliminate the scalability, performance, and troubleshooting drawbacks associated with the VLAN-based approach to network virtualization. Layer 3-based campus networks built over the past several years have proven to be scalable, robust, and high performing.

However, a characteristic behavior of a Layer 3 switch—to route between all networks in the forwarding table—runs counter to the need for partitioning and closed user groups. Although access control lists (ACLs) and policy-based routing are possible approaches to traffic partitioning, the anticipated number of closed user groups and distribution zones is an important factor to consider in network planning. With increasing numbers of closed user groups come increases in administrative and operational tasks and the associated operational costs. One error in an ACL configuration for a single location could result in a dangerous security compromise, exposing the entire network to any number of threats.

Network addressing structures should be carefully considered when using ACLs or policy-based routing. Optimized group IP address ranges simplify the configuration process; however, this can present a drawback because end-system addressing often needs to be changed. Making this change affects not only the network group within an organization, but also the client/server administrators for individual closed user groups. One point to note is that with VLAN-based separation, address space cannot be reused.

Campus design recommendations have lacked an elegant way of isolating traffic in the network to provide a secure and independent environment for closed user groups within the campus.

THE SOLUTION: NETWORK VIRTUALIZATION

A scalable solution is needed for keeping groups of users totally separate and centralizing services and security policies while preserving the high-availability, security, and scalability benefits of the campus design. To address this solution, the network design needs to effectively solve the following challenges:

- **Access control:** Help ensure that legitimate users and devices are recognized, classified, and authorized entry to their assigned portions of the network.
- **Path isolation:** Help ensure that the substantiated user or device is mapped to the correct secure set of available resources—effectively, the right VPN.
- **Services edge:** Help ensure that the right services are accessible to the legitimate set or sets of users and devices, with centralized policy enforcement.

The answer is network virtualization, which can be achieved in several ways. Virtualization technologies enable a single physical device or resource to act like it is multiple physical versions of itself and be shared across the network. Network virtualization is a crucial element of the Cisco® SONA framework. Cisco SONA uses virtualization technologies to increase use of networked assets such as servers and storage-area networks (SANs). For example, one physical firewall can be configured to perform as multiple virtual firewalls, helping enterprises optimize resources and security investments. Other virtualization strategies include centralized policy management, load balancing, and dynamic allocation. The use of virtualization enhances agility and improves network efficiency, reducing both capital and operational expenses.

Access Control: Authentication and Access-Layer Security

Security at the access layer is vital for protecting the campus LAN from external threats, whether inadvertent or malicious, whether presented by a user or harbored by an infected device. With the delivery of wireless and mobile access to the campus LAN, user authentication and other security measures take on increased importance. Cisco network virtualization solutions are complemented by several technologies that help ensure threats are mitigated before they can enter the campus.

One such technology is IEEE 802.1X, which is the standard for port authentication. The initial purpose of the 802.1x standard was related to authentication only. Enhancements such as VLAN assignment through RADIUS today greatly tie in with associating a user with a closer user group. Effectively, it enables an association between user IDs and their closed user groups. When a user from Group A authenticates at the network edge, 802.1X helps ensure that that port or user is connected to the Group A VLAN/VPN.

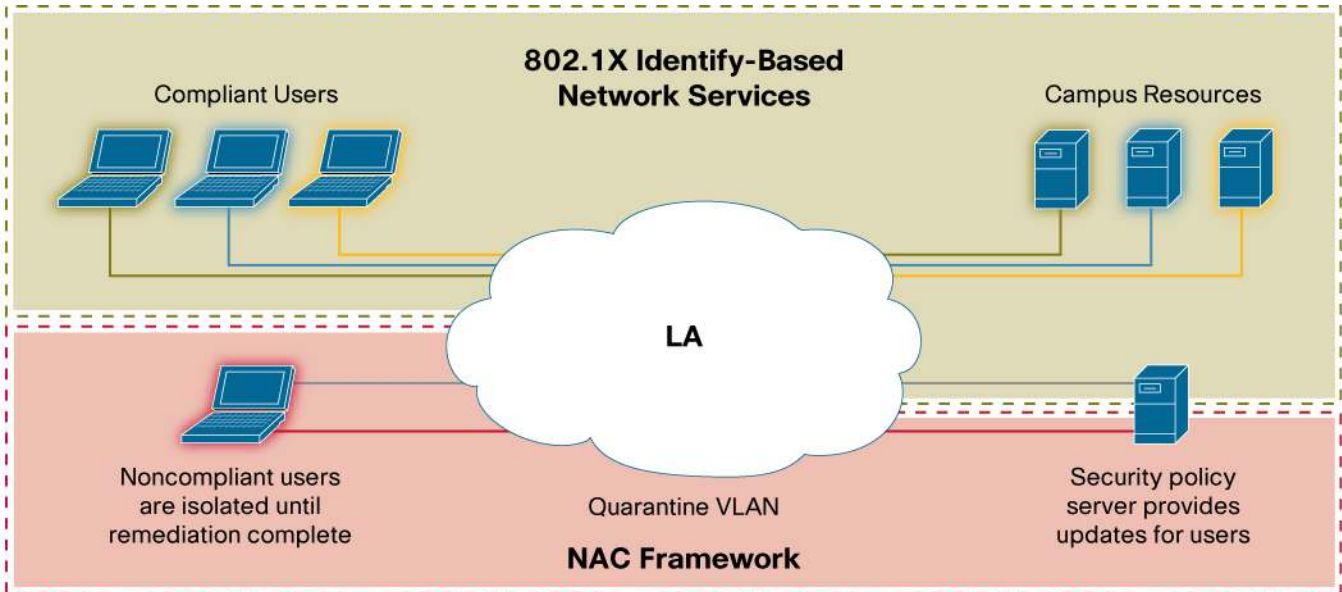
Another complementary technology is Cisco NAC. Although it is supported on all Cisco campus switch platforms and can provide benefits to the entire network, NAC is optimally deployed at the access layer. NAC's job is to mitigate threats at the edge and remove harmful traffic before it reaches other clients and servers in the network. After a user authenticates and passes through the 802.1X checkpoint, NAC comes into play to help ensure that users do not expose the campus network to any viruses, worms, or other threats.

NAC checks to see if the PC attempting access is up to date in terms of operating system and virus protection. If a user trying to gain access has missed the latest software update or is suspect for other reasons, then—after that user has authenticated with a user ID and password—the user is connected to a separate VLAN. This VLAN is connected to a closed user group for any PC not in compliance with corporate software standards. Effectively, the user is placed in quarantine, where the only network resource available is a server containing the necessary software downloads.

Before the availability of the network virtualization solution, all noncompliant PCs would either need to be placed into a campuswide VLAN or policy-based routing would need to be configured on each Layer 3. Although enterprises aim for getting rid of Layer 2 VLANs through the core of their network, this campuswide VLAN would hinder them from getting there. However, configuring policy-based routing on a per Layer 3 hop basis would represent a significant configuration overhead. Network virtualization basically allows the scaling of this solution while not relying on campuswide VLANs.

As shown in Figure 1, the users associated with the green group are fully up to date and can access resources on the campus LAN. However, the user in the red group is not up to date, and NAC sends that user to a “quarantine” VLAN until the PC has been updated with the proper software.

Figure 1. Access Control



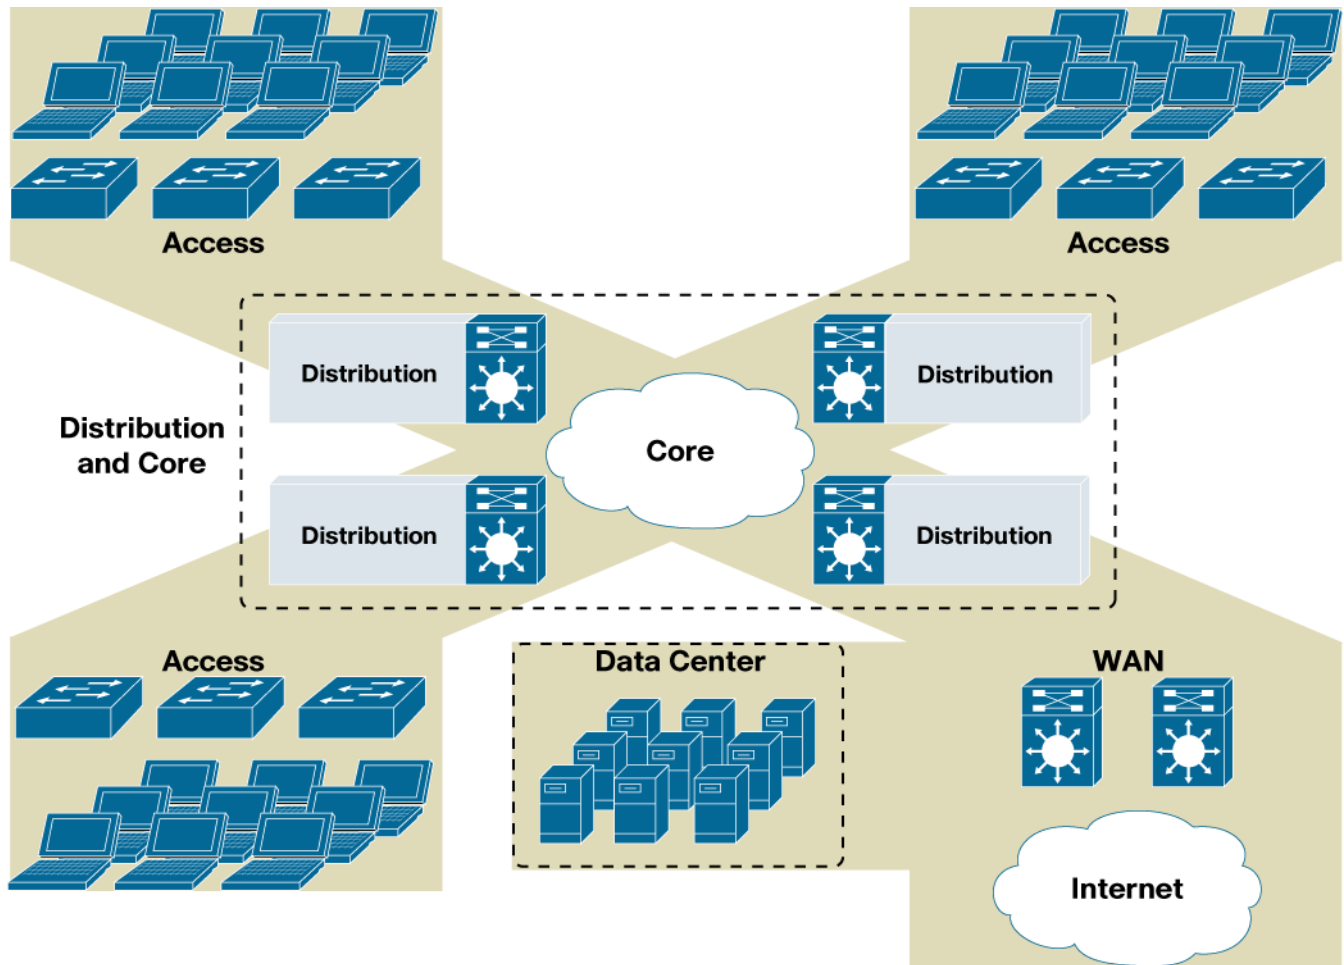
Path Isolation: Layer 3 VPNs

Typical campus network designs use a mix of switching (Layer 2) technologies at the network edge (access) and routing (Layer 3) technologies at the network core (distribution and core layers). To address network virtualization for the campus, Cisco offers three solutions that are well suited to these environments:

- GRE tunnels
- VRF-lite
- MPLS VPNs

In order for these three solutions to be effective, a scalable network design must be in place with a routed campus core and distribution layer, as shown in Figure 2. Elements of this design, such as the distribution layer, can be seen as “building blocks” and replicated many times over. With the proper building blocks in place, any or all of these solutions can be deployed in a single network.

Figure 2. Scalable Network Design



GRE Tunnels

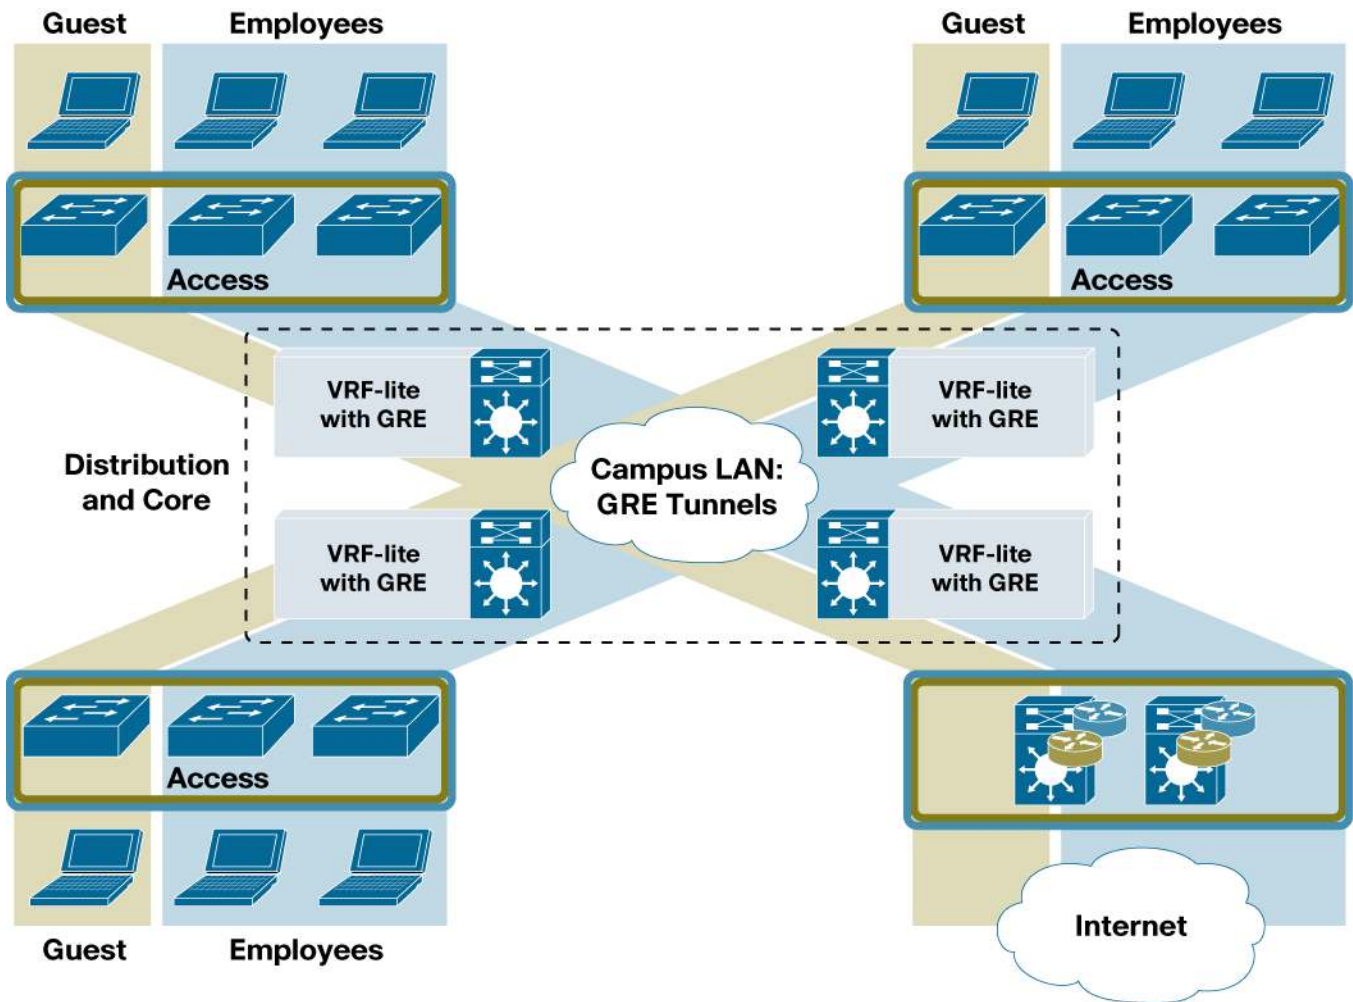
GRE tunnels represent a fairly simple approach to creating closed user groups on the campus network.

A frequent requirement for corporate IT departments is to provide access to the global Internet for onsite guests or visitors, but to prevent those users from accessing internal sites and resources. Often, the simplest solution is to extend a single “guest” VLAN across a large part of the network. However, this option has several weaknesses:

- For networks with a fully routed core, a guest VLAN cannot be extended without changing the configuration of the core.
- It does not adequately isolate guest traffic from regular corporate LAN traffic.
- It decentralizes security and quality-of-service (QoS) functions.
- It can expose the campus LAN to instability.

GRE tunnels can be used to address many network virtualization requirements, but because they are ideally suited as a guest-access solution, that application will serve as the example for this document. As shown in Figure 3, GRE tunnels used in combination with the Cisco VRF-lite feature can create a simple, easy-to-administer partitioning for guest access in any Layer 3 network where Cisco Catalyst® 6500 Series Switches are deployed with either the Cisco Catalyst 6500 Series Supervisor Engine 720 or the Cisco Catalyst 6500 Series Supervisor Engine 32.

Figure 3. GRE Tunnels Used with VRF-lite



Rather than extending a VLAN across the network to provide guest access, guest traffic is isolated to a unique VRF path at each distribution layer switch. The traffic is then transported across the corporate LAN through the GRE tunnel to a central device, such as an Internet edge router.

The advantages to this solution include:

- Can span over a typical multilayer campus network (no need for campuswide VLANs).
- Guest user traffic is isolated from the rest of the corporate LAN traffic.
- The point of ingress for all guest traffic is centralized, making security and QoS policies easier to administer.
- Can even be extended over the WAN to branches.

Each of the distribution switches in this scenario requires the configuration of an individual tunnel to the Internet edge device. A typical distribution-layer building block design calls for dual distribution-layer switches. Depending on the exact design of the building block, the gateway redundancy protocol in use, and the user's redundancy requirements, a single tunnel might suffice, or multiple tunnels might be required—one for each switch.

One consideration regarding GRE tunnels as a solution for closed user groups is that the tunnels themselves are intense to configure and manage, for which reason the solution is not advisable for more than one or two tunnels. This type of network virtualization is suitable where hub-and-spoke topologies are required.

Another consideration with this solution is IP maximum transmission unit (MTU) settings. GRE adds 24 bytes of overhead to each packet forwarded across a tunnel interface, meaning that a 1500-byte packet typically exceeds the MTU of intermediate network segments. Typically the MTU size for the guest access overlay network would be reduced by at least 24 bytes, so no changes to the existing network architecture need to be done.

Extending Guest Access Solutions

Several extensions to the basic guest-access model can offer additional functionality and flexibility:

- Deploy a firewall services module (FWSM) on the Cisco Catalyst 6500 in the Internet edge switch to further isolate the guest traffic and provide additional filtering and access control.
- Deploy other services modules on the Cisco Catalyst 6500 to provide intrusion detection or network analysis for guest traffic.
- Integrate wireless access points in the guest-access VLANs to provide wireless access for guest users.
- Use VRFs for the tunnels on the Internet edge switch to isolate guest traffic and provide limited access to defined services (such as a Web front end with corporate information or access to specific data).
- Apply security and QoS policies centrally for all guest traffic, easing the configuration burden and centralizing administration.

VRF-lite

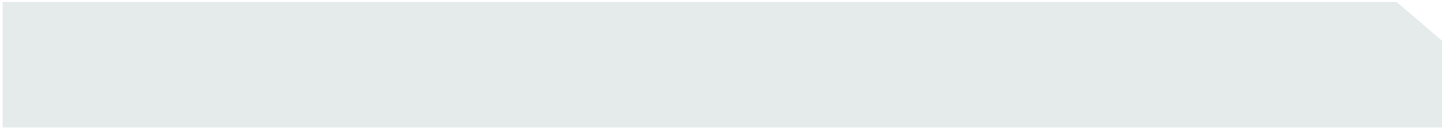
VRF-lite, a Cisco feature that also goes by the generic name of Multi-VRF Customer Edge, provides a solution for campus partitioning by enabling a single routing device to support multiple virtual routing tables. Each interface is then associated with the global or one of the virtual routing tables. VRF-lite enables support for two or more VPNs on customer edge devices, where IP addresses can be overlapped among the VPNs.

VRF-lite is effectively a lightweight version of MPLS. In case of Ethernet it works in conjunction with the IEEE 802.1Q standard for trunk ports, enabling a single physical port to support multiple VRFs. Individual groups can actually be commingled within a single building block. The router labels each packet with an 802.1Q tag, so that at each hop along the way, the routing device can identify the group to which the packet belongs.

This solution uses input interfaces to distinguish routes for different VPNs and forms virtual packet-forwarding tables by associating one or more Layer 3 interfaces with each VRF. Interfaces in a VRF can be either physical, such as Ethernet ports, or logical, such as VLAN switched virtual interfaces (SVIs)

In a model that incorporates MPLS VPNs, VRF-lite extends limited provider edge functionality to a customer edge router, enabling it to maintain separate VRF tables in order to extend the privacy and security of an MPLS VPN all the way out to the network edge. This document will explore MPLS VPNs in detail in a later section.

The customer edge router using VRF-lite can partition its LAN traffic by placing each client or organization with its own IP address space. Each IP address space can be implemented either on separate Ethernet interfaces or through one Ethernet interface segmented into multiple subinterfaces. Each subinterface contains its own IP address space to separate each different client.




As previously mentioned, each VRF maintains an independent routing domain. This characteristic provides the flexibility of using any IP address space for any given VPN, regardless of whether it overlaps or conflicts with the address spaces of other VPNs. Therefore, each group can independently use private IP addressing, as defined in RFC 1918. In this scenario, Network Address Translation (NAT) is not required. (Note: NAT might still be required in order to send traffic to the Internet.)

This addressing flexibility is beneficial in many scenarios. For example, when the networks of acquired companies are merged into a shared LAN, the acquired network can be incorporated into the infrastructure as a separate VPN. In this way, the acquired company's network can preserve its original address space without conflicting with other VPNs. Likewise, this flexibility is beneficial to enterprises that host engineering or development groups that often need to manage their own address space independently.

VRF-lite can be used either as an end-to-end solution, as shown in Figure 4, or in conjunction with another solution for closed user groups, as discussed in the next section of this document. Platform restrictions with respect to processor and memory apply; the number of VRFs supported by the customer edge router is dependent on the platform, processing power, and available memory. As a design practice, network managers should factor in the processing and memory requirements for routing processes, management, packet forwarding, and so on.

Although VRF-lite is a more scalable solution than GRE tunnels, it is best suited for networks with up to five partitions that are mostly static. One consideration with VRF-lite is that it requires manual reconfiguration of each connected network device on the campus LAN with each addition of a user group. In large networks this can be fairly labor-intensive. Subsequently, VRF-lite is usually deployed in networks with limited numbers of core and distribution switches with stable configurations. Many switches in the Cisco Catalyst Switching Series support the VRF-lite protocol, including the Cisco Catalyst 6500, 4500, 3750, 3560, and 3550 Series Switches.

Figure 4. VRF as an End-to-End Solution

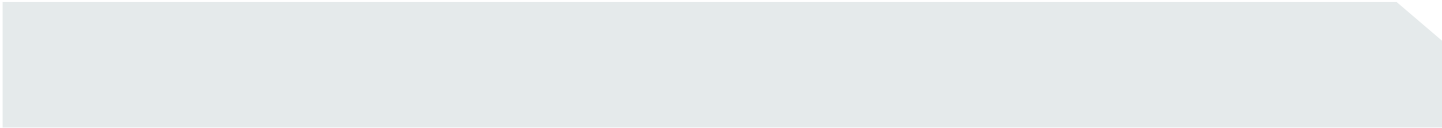


MPLS VPNs

Another way to partition a campus network for closed user groups is by overlaying MPLS-based Layer 3 VPNs onto the routed segment of the campus LAN. Like GRE tunnels and VRF-lite, MPLS VPNs provide a secure and dependable way to form logically separated networks on a common physical infrastructure.

Service providers have made use of MPLS technology for several years. In enterprise networks, though, MPLS has not been widely deployed, mainly because of the lack of support for this technology on LAN switches. But changing business requirements and, in response, new product availability, are helping MPLS emerge as a vital technology in the campus infrastructure.

As enterprises grow and their networking infrastructures evolve to support myriad converged services, IT managers have been under pressure to provide scalable and highly available networking infrastructures. More and more services (such as telephony, video, surveillance, sound systems, and so on) now rely on the enterprise's network, forcing service requirements to become more stringent and very specific. For enterprises to deliver these service levels with the appropriate security and policy customization, they now require technologies that in the past were used only by service providers.




With the introduction of MPLS VPN support on the Cisco Catalyst 6500 Series Switches in late 2003, MPLS technology became available at an affordable price point for many large enterprises. MPLS VPNs, as defined in RFC 2547, basically offer all the benefits of the other solutions discussed in this document (VRF-lite and GRE tunnels). As shown in Figure 5, closed user groups are established through VPNs that are transported independently over the core of the network, using labels for group identification. The networkwide benefit of this approach is that any VPN can be configured to connect users and resources at any location in the network, without any compromises in performance or network design. Accordingly, MPLS VPNs are the most scalable of the three solutions for network infrastructure virtualization discussed in this document.

Unlike VRF-lite, MPLS VPN as a technology automatically manages the distribution of routes and closed user groups, disseminating that information across the network. Because this process is handled dynamically, no manual reconfiguration of distribution and core links is needed when groups are added or changed, another factor that adds to its scalable nature.

As with VRF-lite, flexibility of network addressing is also a benefit of MPLS VPNs. This characteristic results from the fact that each user group is completely autonomous and each VPN makes use of its own VRF table. Therefore, addressing across VPNs is completely independent and can even overlap with addressing elsewhere in the campus infrastructure. In the event that shared or common services (for example, Domain Name System [DNS], e-mail, and Internet access) are used, NAT would be required on a per-VRF basis.

VLANs are deployed at the network edge as part of the existing campus design and can map to Layer 3 VPNs, which are overlaid onto the routed portion of the campus. In this manner, all the benefits of a hierarchical campus deployment are preserved while the solution achieves end-to-end scalable segmentation and centralized security and services in the campus LAN.

Figure 5. MPLS VPNs for Any-to-Any Connectivity



Unified Access

A crucial point regarding the three Cisco path-isolation solutions for network virtualization in the campus is that they do not limit users to any specific type of access, such as Ethernet only. Although they work within a single physical network infrastructure, these solutions can easily accommodate mobile users.

For example, a user might connect from home using an IP Security (IPsec) VPN tunnel, which would terminate in the appropriate user group within the campus LAN. Whether the solution in use is based on GRE tunnels, VRF-lite, or MPLS VPNs, users can be tied transparently to their closed user groups from wherever they have network access.

If IEEE 802.1X or Cisco NAC is deployed across the campus, configuration of access ports can be unified. Typically the number of access ports in a campus network is making up the majority of the ports in the network. These ports can be configured in a unified fashion and will be dynamically associated to the individual user's VLAN/VRF.

SERVICES EDGE: VIRTUALIZED SERVICES

Sometimes a need arises for members of different closed user groups to communicate with each other or share network resources—typically in a limited fashion. For example, traffic from the red group can go to the blue group, but it must pass through a firewall, or the communication is limited to certain hours of the day. In such cases, the network must have a central point of policy enforcement.

A highly effective way to address policy enforcement in the campus LAN is to integrate firewall services into the distribution layer.

Centralized Services

As with policy enforcement, the centralization of other services greatly simplifies and strengthens security enforcement. By helping ensure a single point of access for each VPN, centralized appliances for firewalling and intrusion detection can be shared by many VPNs. A wealth of other services that are common to the different VPNs can also be shared, and doing so can significantly reduce the capital and operational expenses of providing these services.

Some of the services that can be centralized include:

- Security policy enforcement (ACLs and firewall conduits)
- Traffic monitoring, accounting, and billing
- Shared Internet and WAN access
- Shared data centers
- Site integration when mergers or acquisitions occur across multiple companies

Centralized Security

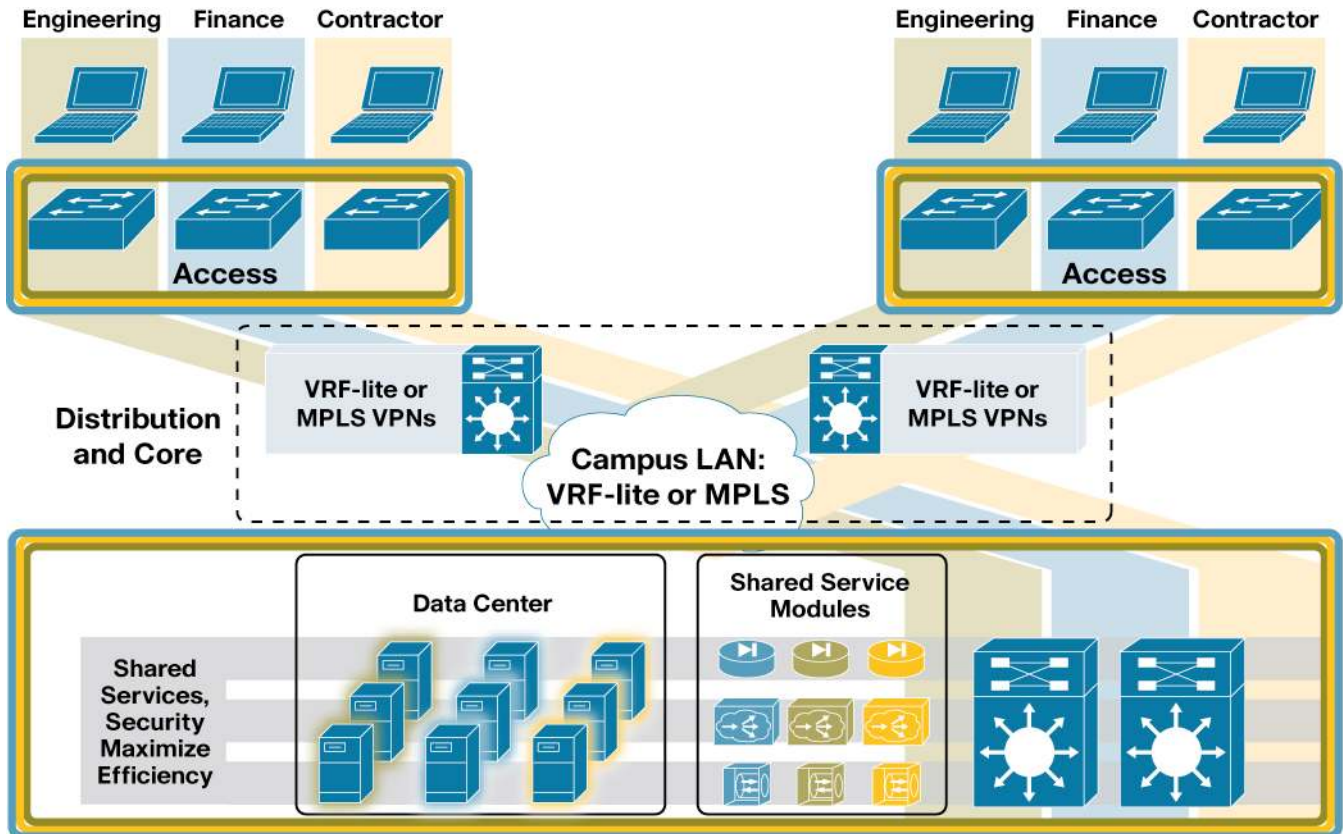
A noteworthy security benefit inherent in VPNs is that traffic cannot enter or leave a VPN unless an entry or exit route is specifically configured for it. This characteristic allows the network engineer to control the positioning and number of points of access to the VPN. It also dramatically reduces the number of ACLs that the network engineer must maintain, because the ACLs need only exist at the points of entry or exit. When these entry/exit points are placed at a central site, that site can be used as a point of centralized security enforcement, as well as a transit area between VPNs and a point of access to shared resources such as server farms, the WAN, or the Internet.

The enforcement of security policies at a central location simplifies management by reducing the number and length of ACLs that need to be maintained. ACLs that once needed to be distributed across the entire campus are now shorter and required only at the central site. The centralization of security enforcement also allows the sharing of security appliances, such as firewalls and intrusion detection systems (IDSs). Next-generation firewalls, such as the Cisco Firewall Service Module for Cisco Catalyst 6500 Series Switches, provide hundreds of virtual firewalls concurrently on a single appliance. These virtual firewalls can be tied to individual VPNs, enabling the allocation of a dedicated virtual firewall for each VPN. By using VPN-aware virtual firewalls, each group can enforce its own policies on its own virtual firewall, while the enterprise needs only to own and maintain a single firewall appliance.

The architecture also allows having multiple hand-off points between closed user groups. This typically is the case in very large networks where sending the traffic over the core to the centralized policy enforcement point is either not practical or not possible. In such cases, firewall services can be deployed throughout the distribution layer by installing Firewall Services Modules in the Catalyst 6500 Series Switches.

Figure 6 shows network virtualization in a campus LAN.

Figure 6. Network Virtualization in the Campus LAN

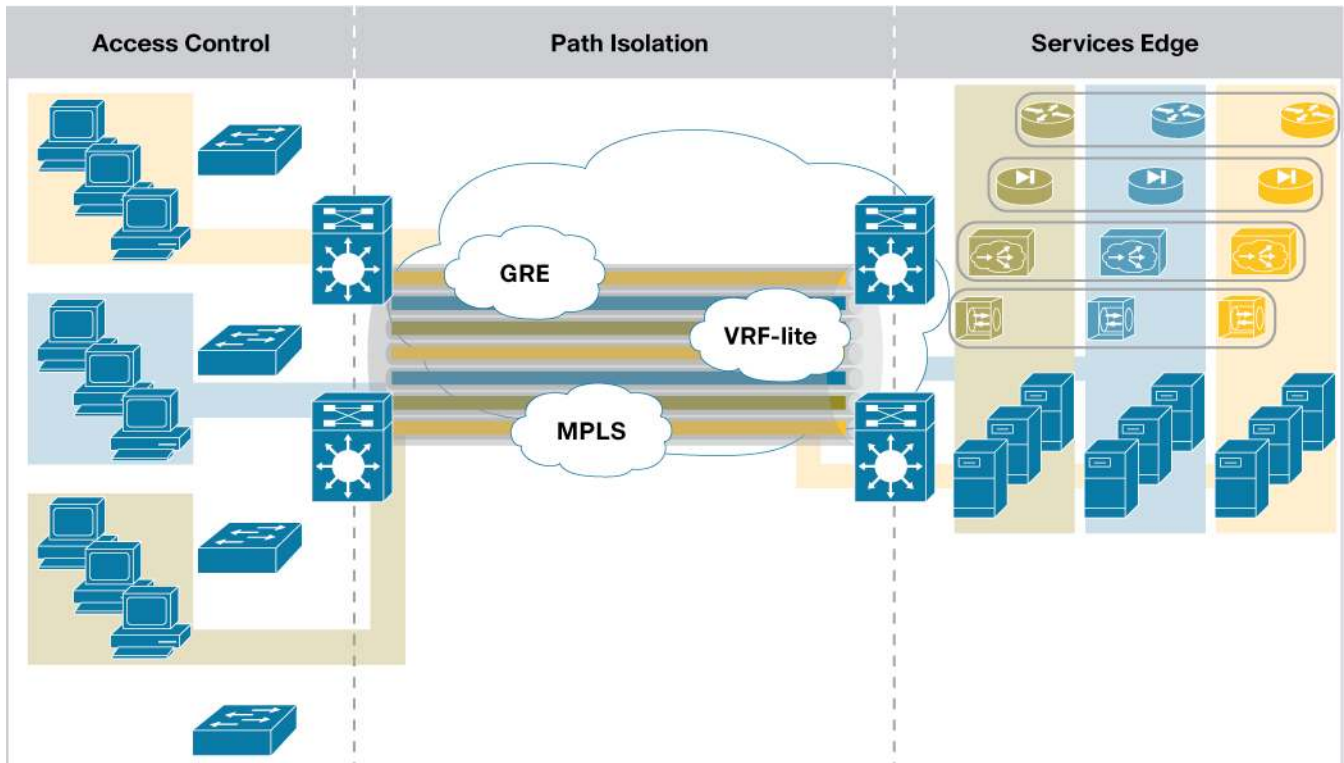


CONCLUSION

In today's evolved networking environments, typical campus network designs use a mix of switching (Layer 2) technologies at the network edge (access) and routing (Layer 3) technologies at the network core (distribution and core layers). Thus, network virtualization can be achieved at the network access layer (Layer 2) by means of VLANs and at the network core (Layer 3) by using GRE tunnels, VRF-lite, and MPLS-based Layer 3 VPNs to partition the routed domain and thus achieve scalable end-to-end virtualization.

With Cisco network virtualization solutions for the campus (Figure 7), enterprises can deploy multiple closed user groups on a single physical infrastructure, while maintaining high standards of security, scalability, manageability, and availability throughout the campus LAN. In light of their virtualized nature and their enablement of centralized services, these solutions form a crucial element of the Cisco SONA framework. A wide range of Cisco Catalyst switches enables enterprises that adopt this framework to use more of their network assets with greater efficiency, allowing them to realize cost savings even as requirements for devices, systems, services, and applications grow.

Figure 7. Complete Network Virtualization Solution



**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0601R)

