

A Basic Virtualized Enterprise

In this chapter, we define the technical requirements posed by the need to virtualize the network. Based on these requirements, we propose an architectural framework comprised of the functional areas necessary to successfully support concurrent *virtual networks* (VNs) over a shared enterprise physical network.

Networks enable users to access services and resources distributed throughout the enterprise. Some of these services and resources are public: those accessed over the Internet, and others that are private and internal to the enterprise. Every enterprise has unique security and service level policies that govern the connectivity to the different services, whether these are public or private.

One of the basic building blocks behind the virtualized network and, in fact, a key driver is security. An important element of an enterprise's security policy is the definition of a network perimeter. In general, the level of trust inside and outside of the network perimeter differs, with end stations inside the perimeter being generally trusted and any access from outside the perimeter being untrusted by default. Communications between the inside and the outside of the perimeter must happen through a checkpoint. At the checkpoint, firewalls and other security devices ensure that all traffic that enters or leaves the enterprise is tightly controlled. Therefore, we refer to the point of entry/exit to/from the enterprise network as the network perimeter.

NOTE


The network perimeter defines one layer of security and must be complemented with other security mechanisms. It is critical to incorporate mechanisms to protect the network from attacks initiated inside the perimeter. This functionality is generally provided at the network access/edge and is not impacted by the virtualization of the network.

To provide the required connectivity, create a secure perimeter and enforce the necessary policies, it is recommended that an enterprise network be based on certain functional blocks. Figure 3-1 depicts a modular enterprise network and its perimeter. The recommended functional blocks are as follows:

- The LAN/MAN transport (core and distribution)
- The LAN edge or access layer

- The Internet access module
- The data center access module
- The WAN aggregation module
- The WAN transport
- The branch

Figure 3-1 *The Modular Enterprise Network and Its Perimeter*



When a single enterprise network must service many different groups, it is often necessary to create virtual networks (VNs) so that each group can enjoy

- Private connectivity over a shared infrastructure.
- A dedicated perimeter in which independent policies can be enforced per group.
- User mobility (ubiquitous access to the appropriate virtual network regardless of the user's location).

At the risk of oversimplifying, a VN can be seen as a security zone. All devices within the security zone trust each other and communicate freely with each other. Meanwhile, any communication with other security zones, or other networks, must happen in a controlled manner over a highly secured perimeter or checkpoint. Thus, a virtualized enterprise network will simultaneously host many security zones, and their dedicated perimeters, over a shared infrastructure.

The Virtual Enterprise

A virtual enterprise network must provide each group with the same services as a traditional dedicated enterprise network would. The experience from an end-user perspective should be that of being connected to a dedicated network that provides connectivity to all the resources the user requires. The experience from the perspective of the network administrator is that they can easily create and modify virtual work environments for the different groups of users and adapt to changing business requirements in a much easier way. The latter derives from the ability to create security zones that are governed by policies enforced centrally. Because policies are centrally enforced, adding or removing users and services to or from a VN does not require any policy reconfiguration. Meanwhile, new policies affecting an entire group can be deployed centrally at the VN perimeter. To virtualize an enterprise network, the basic functional blocks of the modular enterprise must be enhanced to provide the following functionality:

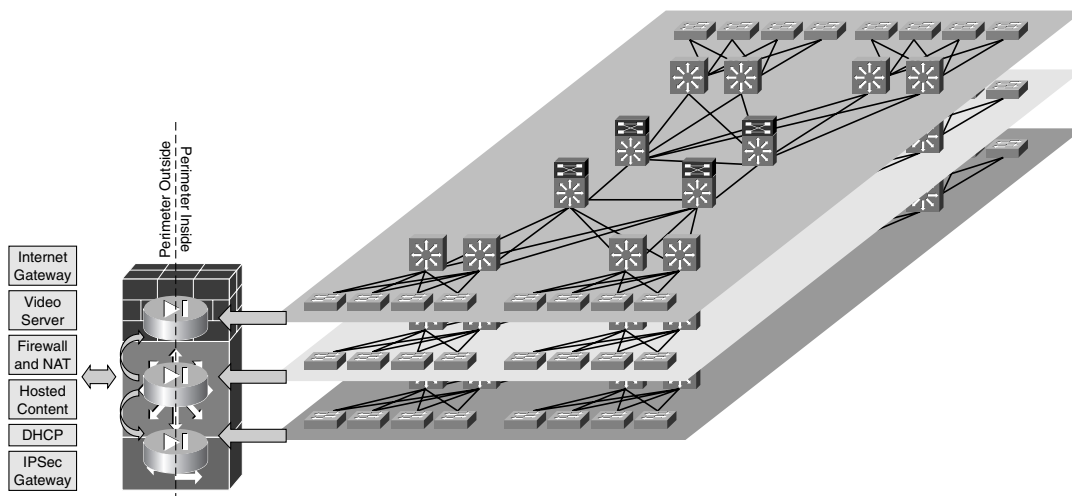
- Dynamically authenticate and authorize users into groups
- Isolate connectivity to guarantee privacy between groups
- Create well-defined and controllable ingress/egress points at the perimeter of each VN
- Enforce independent security policies for each group at the perimeter
- Centralize the enforcement of the perimeter security policies for the different VNs by
 - Allowing secure collaboration mechanisms among groups
 - Allowing secure sharing of common resources
- Provide basic networking services for the different groups, either shared or dedicated
- Provide independent routing domains and address spaces to each group

You could use many different technologies to solve the listed challenges. The technologies available and how these can be used to meet the above requirements are the topic of the remaining chapters in the book.

From an architectural perspective, the previous requirements can be addressed by segmenting the network pervasively into VNs and centralizing the application of network policies at the perimeter of each VN. These are, of course, the policies for ingress and egress to the VN or security zone. The formation of a trusted security

zone relies on traffic-isolation mechanisms rather than a distributed policy. Because traffic internal to a zone is trusted, policies are required only at the perimeter to control the access to external resources that could in many cases be shared. Figure 3-2 illustrates this concept.

Figure 3-2 *Virtual Networks with Centralized Policies at the Perimeter*



Regardless of where a user is connected, its traffic should always use the same VN and be directed through a central site of policy enforcement (VN perimeter), should it need to exit the VN. This makes users mobile and ensures that regardless of their location they will always be subject to the same policies. To ensure that users are always connected to the right VN, dynamic authentication and authorization mechanisms are required. These allow the identification of devices, users, or even applications so that these can be authorized onto the correct virtual segment and thus inherit the segment's policies.

The virtualization architecture described so far can be organized into functional areas. These functional areas provide a framework for the virtualization of networks:

- Transport virtualization
- Edge authorization
- Central services access (VN perimeter)

As you will see throughout the book, this modular framework gives the network architect a wide choice of technologies for each functional area. A key element in achieving this degree of flexibility is the definition of clear communication interfaces between the different areas.

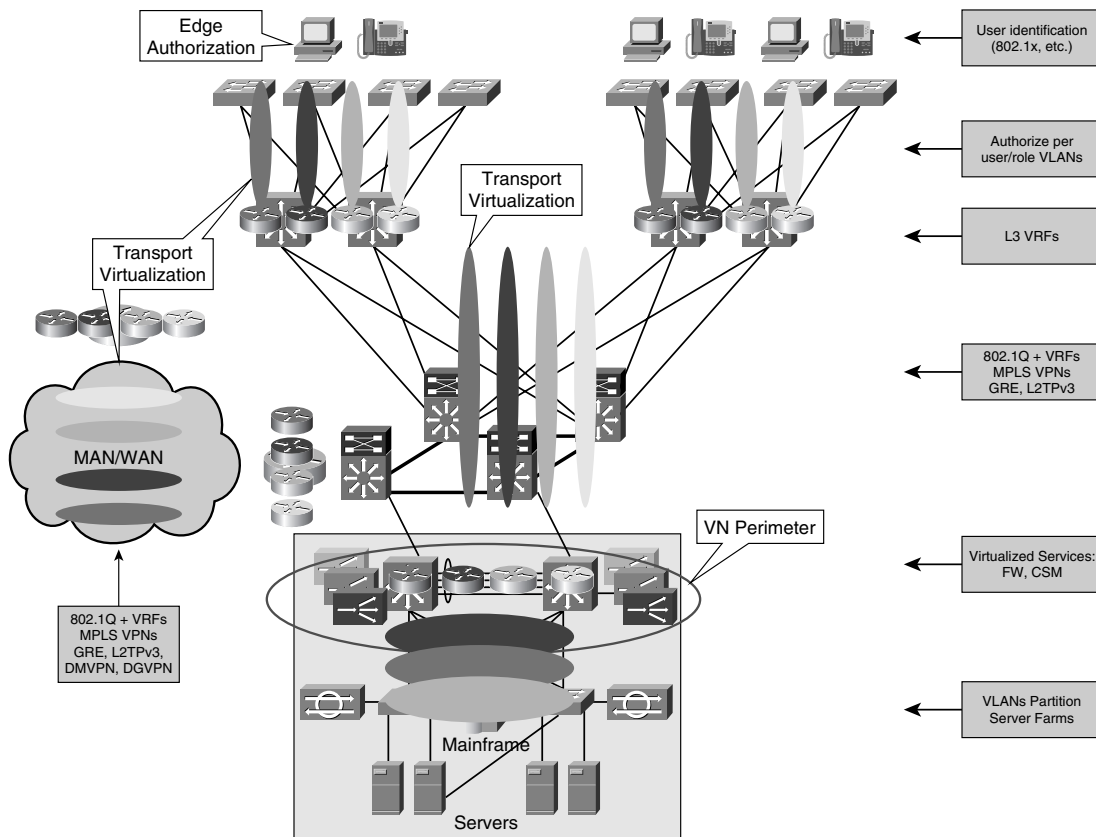
VLANs provide an example of a communication interface between functional areas. The edge authorization module assigns a user to a VLAN, and the transport module maps

that VLAN to a VN. At the destination, the transport module maps the VPN back to a VLAN. If the destination is outside the VN perimeter, the transport module hands off a VLAN to the central services access module, which maps the VLAN to the necessary virtual services. As you progress through the book, you learn that the interface between modules could very well be a label or a policy.

NOTE There are, of course, pros and cons to using different types of communication interfaces. These are analyzed as the different technologies are discussed in detail, so read on.

Figure 3-3 shows the functional areas of the virtualized enterprise. As shown, you can use a variety of technologies for each different area.

Figure 3-3 *Virtualized Enterprise Network Functional Areas*



A useful way to look at Figure 3-3 and understand the role of the different functional areas is to look at it from the top down. Starting at the top, the endpoints connected to the network are authenticated and as a result of the authentication are authorized onto a specific VLAN (edge authorization). Each VLAN maintains its traffic separate from other VLANs and is mapped to a *virtual routing and forwarding instance* (VRF).

NOTE VRFs are logical routing and forwarding tables with associated interfaces and routing processes, what could be thought of as a virtual routing instance. The section on “Control-Plane-Based Segmentation” and Chapter 4 examine the concept of a VRF in more detail.

Each VRF is connected to other VRFs in its VN and keeps its traffic separate from VRFs that belong to other VNs (transport virtualization). When traffic is destined to a resource outside the VN (for example, the data center), it is routed to the VN perimeter, where virtual services, such as firewalls and load balancers, are applied to each group (central services access—VN perimeter). Traffic destined to a subnet over the WAN is kept separate from traffic in other VNs through the virtualization of the WAN transport (transport virtualization).

Transport Virtualization—VNs

When segmenting the network pervasively, all the scalability, resiliency, and security functionality present in a nonsegmented network must be preserved and in many cases improved. As the number of groups sharing a network increases, the network devices must handle a much higher number of routes. Any technologies used to achieve virtualization must therefore provide the necessary mechanisms to preserve resiliency, enhance scalability, and improve security.

Chapter 2, “Designing Scalable Enterprise Networks,” discussed network design recommendations that provide high availability and scalability through a hierarchical and modular design. Much of the hierarchy and modularity discussed relies on the use of a routed core. Nevertheless, some areas of the network continue to benefit from the use of Layer 2 technologies, such as VLANs, ATM, or Frame Relay circuits. Thus, a hierarchical IP network is a combination of Layer 3 (routed) and Layer 2 (switched) domains. Both the Layer 2 and the Layer 3 domains must be virtualized, and the virtualized domains must be mapped to each other to create VNs.

NOTE

The term *virtual private network* is broadly used and might have different connotations to different people. To avoid confusion, we use the term *virtual network* as an implementation-independent concept. In many implementations, a VN is actually a VPN; but as you read, you might want to avoid creating a direct association between your favorite type of VPN implementation (IPsec, *Secure Sockets Layer* [SSL], IP-VPN) and the concept of a VN, which we are here introducing.

One key principle in the virtualization of the transport is that it must address the virtualization of the network devices and their interconnection. Thus, the virtualization of the transport involves two areas of focus:

- **Data-path virtualization**—Refers to the virtualization of the interconnection between devices. This could be a single-hop or multiple-hop interconnection. For example, an Ethernet link between two switches provides a single-hop interconnection that can be virtualized by means of 802.1q VLAN tags; for Frame Relay or ATM transports, separate virtual circuits provide data-path virtualization. An example of a multiple-hop interconnection would be that provided by an IP cloud between two devices. This interconnection can be virtualized through the use of multiple tunnels (*generic routing encapsulation* [GRE] for example) between the two devices.
- **Device virtualization**—Refers to the virtualization of a networking device or the creation of logical devices within the physical device. This includes the virtualization of all processes, databases, tables, and interfaces within a device.

In turn, within each networking device, there are at least two planes to virtualize:

- **Control plane**—Refers to all the protocols, databases, and tables necessary to make forwarding decisions and maintain a functional network topology free of loops or unintended blackholes. This plane could be said to draw a clear picture of the topology for the network device. A virtualized device must possess a unique picture of each VN it is to handle, hence the requirement to virtualize the control-plane components.
- **Forwarding plane**—Refers to all the processes and tables used to actually forward traffic. The forwarding plane builds forwarding tables based on the information provided by the control plane. Similar to the control plane, each VN will have a unique forwarding table that needs to be virtualized.

Furthermore, the control and forwarding planes can be virtualized at different levels, which map directly to different layers of the OSI model. For instance, a device can be VLAN aware and therefore virtualized at Layer 2, but yet have a single routing table, *Routing Information Base* (RIB), and *Forwarding Information Base* (FIB), which means it is not virtualized at Layer 3. The different levels of virtualization come in handy, depending on the technical requirements of the deployment. Sometimes Layer 2 virtualization is enough (a wiring closet, for instance). In other cases, virtualization of other layers might be necessary.

For example, providing virtual firewall services requires Layers 2, 3, and 4 virtualization, plus the ability to define independent services and management on each virtual firewall, which some may argue is Layer 7 virtualization. We delve into firewall virtualization in Chapter 4. For now, we focus on the virtualization of the transport at Layers 2 and 3.

VLANs and Scalability

Time and experience have proven the scalability benefits of limiting the size of Layer 2 domains in a network. A large amount of this experience comes from campus networks, where highly resilient topologies with redundant links are possible. This link redundancy intrinsically creates network loops that must be controlled by mechanisms such as spanning tree. The broadcast nature of a Layer 2 domain is the main reason these redundant links behave as loops rather than redundant active paths capable of load balancing. Hence, the lack of load balancing and the complexity involved in managing large and highly resilient spanning-tree domains makes a routed infrastructure much more appropriate for large-scale highly available networks. Thus, experience has taught us that meshed Layer 2 domains have their role in the network, but they must be kept small in scale. Keep in mind that we are referring to highly meshed resilient Layer 2 domains such as those you would find in a campus. This type of problem is faced less in the WAN, where point-to-point connections tend to be at the base of the architecture and are for the most part routed. Nevertheless, the introduction of technologies that extend Layer 2 domains over an IP infrastructure has brought many of the spanning-tree concerns to the table in the *metro-area network* (MAN) and WAN.

When you are virtualizing a network, it is tempting to revisit ideas such as end-to-end VLANs. After all, mapping a group of users to a specific VLAN to create an isolated workgroup was one of the original thoughts behind the creation of VLANs. Should the VLAN traverse the entire enterprise, we could say the transport has been virtualized. This type of solution will have all the scalability problems associated with large Layer 2 domains and is therefore not desirable.

Nevertheless, the use of VLANs has its place as a way of segmenting the Layer 2 portion of the network. In an enterprise campus, this is generally the mesh of links between the access and the distribution. Remember, the recommendation is to reduce the size of the broadcast domains to something manageable, not necessarily to eliminate the broadcast domains, because too much IP subnet granularity would also represent a management challenge. So, to segment the access portion of the network, VLANs are of much use.

NOTE

Later on, in the section “Policy-Based Segmentation,” you will see that there are mechanisms to achieve traffic differentiation by using code points. These techniques do not create separate broadcast domains and are effective only after entering the routed core. The use of code points will not provide separation between groups that share a broadcast domain. VLANs are required to provide Layer 2 separation at the access.

The network must preserve its hierarchy and therefore its routed core. As the periphery (access/distribution) continues to be switched (as opposed to routed), VLANs must be used for segmentation purposes. Thus, a VLAN in a wiring closet would represent the point of entry into a VN.

Because these VLANs are terminated as they reach the routed core, it is necessary to map them to segments created in the routed core. The next section looks into what is necessary in the core. From the access perspective, the VLANs must map to the corresponding segments created in the core to achieve an end-to-end VPN that spans both the switched and routed portions of the network.

We focus our analysis on a network with a routed core and a switched access. This model is widely adopted because it has been proven, optimized, and recommended by Cisco for many years.

Virtualizing the Routed Core

You can achieve the virtualization of the routed portion of the network in many ways. At the device level, the available traffic separation mechanisms can be broadly classified as follows:

- Policy-based segmentation
- Control-plane-based virtualization

Policy-Based Segmentation

Policy-based segmentation restricts the forwarding of traffic to specific destinations, based on a policy and independently of the information provided by the control plane. The policies are applied onto a single IP routing space. A classic example of this uses an *access control list* (ACL) to restrict the valid destination addresses to subnets in the VN.

Policy-based segmentation is limited by two main factors:

- Policies must be configured pervasively.
- Locally significant code points are currently used for policy selection.

The configuration of distributed policies can be a significant administrative burden, is error prone, and causes any update in the policy to have widespread impact.


The code point used for policy selection has traditionally been an IP address and therefore locally significant. Because of the diverse nature of IP addresses, and because policies must be configured pervasively, building policies based on IP addresses does not scale well. Thus, policy-based segmentation using IP addresses as code points has limited applicability. However, other code points could potentially be used. If the code point is independent of the IP addressing and globally significant (uniformly maintained throughout the network), all policies would look alike throughout the network, making their deployment and maintenance much simpler.

NOTE

An example of globally significant code points are the *differentiated services code points* (DSCPs) used for the selection of *per-hop behaviors* (PHBs) in a DiffServ *quality of service* (QoS) architecture. Different PHB policies are selected and enforced at each hop based on the traffic's DSCP label. The DSCP labels identify types of traffic through the network, regardless of source/destination subnets. DSCP is just one example of a globally significant code point; in general, any label could serve the purpose. The use of a label (code point) to identify types of traffic is a powerful concept and could be leveraged to identify traffic for policy application. Thus, if traffic is labeled appropriately, ACLs based on code points rather than IP addresses could provide a scalable alternative to policy-based segmentation.

Policy-based segmentation with the tools available today (ACLs) can address the creation of VNs with many-to-one connectivity requirements; it would be hard to provide any-to-any connectivity with such technology. This is the case for segments providing guest access to the Internet, in which many guests access a single resource in the network. This is manageable because the policies are identical everywhere in the network (allow Internet access, deny all internal access). The policies are usually applied at the edge of the Layer 3 domain. Figure 3-4 shows ACL policies applied at the distribution layer to segment a campus network.

Figure 3-4 Hub-and-Spoke Policy-Based Segmentation



As a creativity exercise, you could attempt to design an IP-based policy to provide any-to-any connectivity between guests, while keeping them separate from the rest of the users!

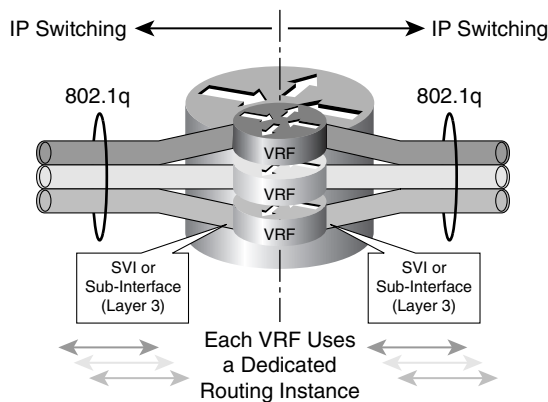
Control-Plane-Based Virtualization

Control-plane-based virtualization restricts the propagation of routing information so that only subnets that belong to a VN are included in any VN-specific routing tables and updates. Thus, this type of solution actually creates a separate IP routing space for each VN. To achieve control-plane virtualization, a device must have many control/forwarding instances, one for each VN. An example of control-plane-based device segmentation is a VRF.

A VRF could be looked at as a “virtual routing instance.” Each VRF will have its own RIB, FIB, interfaces, and routing processes. Figure 3-5 illustrates VRFs.

NOTE A VRF is not strictly a virtual router because it does not have dedicated memory, processing, or I/O resources. In Chapter 4, we discuss other levels of device virtualization, such as logical routers, and proper virtual routers. For now, we use the analogy just to help you understand what a VRF is.

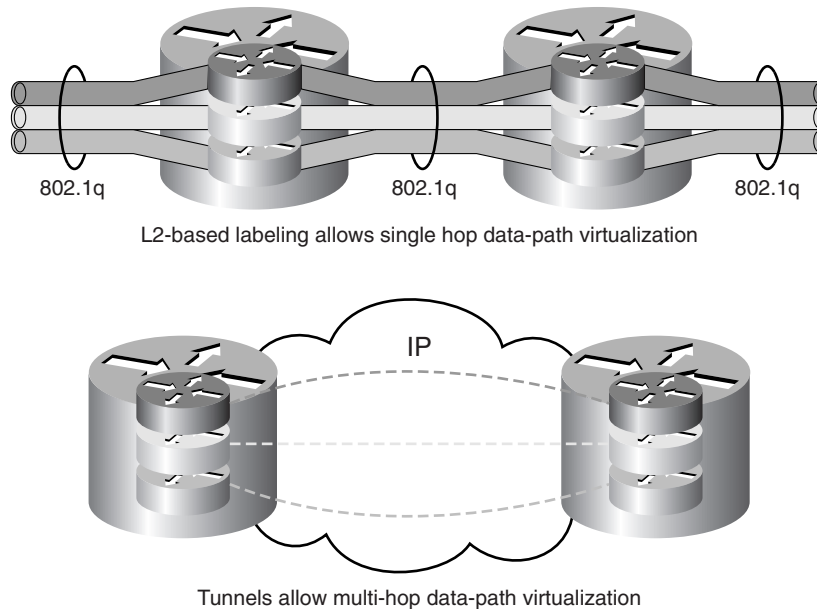
Figure 3-5 Virtual Routing and Forwarding



The VRF achieves the virtualization of the networking device at Layer 3. After the devices have been virtualized, the virtual instances in the different devices must be interconnected to form a VN. Thus, a VN is a group of interconnected VRFs. In theory, this interconnection could be achieved by using dedicated physical links for each VN (group of interconnected VRFs). In practice, this would be inefficient and costly. Hence, it is necessary to virtualize the data path between the VRFs to provide logical interconnectivity between the VRFs that participate in a VN. The type of data-path virtualization will vary depending on how far the VRFs are from each other. If the virtualized devices are directly connected to each other (single hop), link or circuit virtualization is necessary. If the virtualized devices are

connected multiple hops apart over an IP network, a tunneling mechanism is necessary. Figure 3-6 illustrates single-hop and multiple-hop data-path virtualization.

Figure 3-6 *Single- and Multiple-Hop Data-Path Virtualization*



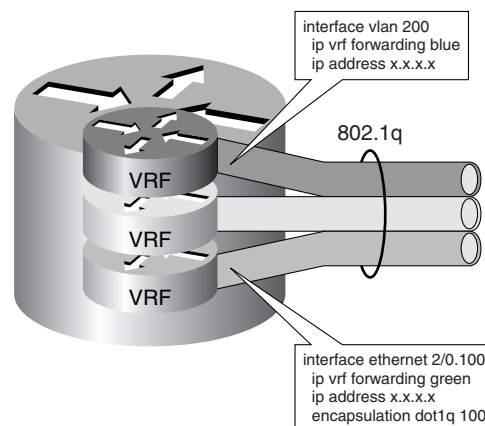
The many technologies that virtualize the data path and interconnect VRFs are discussed in Chapters 4 and 5. The different technologies have different benefits and limitations depending on the type of connectivity and services required. For instance, some technologies are good at providing hub-and-spoke connectivity, whereas others provide any-to-any connectivity. The support for encryption, multicast, and other services will also determine the choice of technologies to be used for the virtualization of the transport.

NOTE

Some technologies leverage the use of labels to “color” routing updates and/or data traffic. In theory, “coloring” allows the interconnection of virtual devices without the need for a dedicated virtual data path for each VN. For example, *multiprotocol interior Border Gateway Protocol* (MP-iBGP) uses a “coloring” mechanism to differentiate routing updates for different RFC 2547 VPNs, but the RFC 2547 forwarding plane relies on dedicated logical data paths to forward traffic (tunnels based on *label switched paths* [LSPs] or *Layer 2 Tunnel Protocol Version 3* [L2TPv3]). Other technologies such as *Multi-Topology Routing* (MTR) rely on “coloring” for both control-plane updates and forwarding, the latter implemented in a mechanism known as “class-based forwarding.” Control-plane coloring for MTR is done natively in the *interior gateway protocols* (IGPs) by labeling the routing updates, much like MP-iBGP does. Chapter 4 provides more detail about the different technologies available to virtualize devices and the data path and about “coloring” for MTR.

The VRFs must also be mapped to the appropriate VLANs at the edge of the network. This mapping provides continuous virtualization across the Layer 2 and Layer 3 portions of the network. The mapping of VLANs to VRFs is as simple as placing the corresponding VLAN interface at the distribution switch into the appropriate VRF. The same type of mapping mechanism applies to Layer 2 virtual circuits (ATM, Frame Relay) or IP tunnels, which are handled by the router as a logical interface. The mapping of VLAN logical interfaces (*switch virtual interface* [SVI]) to VRFs is illustrated in Figure 3-7.

Figure 3-7 VLAN-to-VRF Mapping



So far, we have created a virtualized transport that can keep the traffic from different groups separate from each other. The next section introduces the functionality required at the edge to place or authorize endpoints into the appropriate groups.

The LAN Edge: Authentication and Authorization

At the edge of the network, it is necessary to identify the users or devices logging on to the network so that they can be assigned to the right groups.

The process of identifying the users or devices is known as *authentication*. Two parameters affect the assignment of a user or devices: the identity of the user or device and the posture of the device. The posture of the device refers to the health of the device, measured by the level of software installed, especially operating system patches and antivirus.

When identified, the endpoints must be authorized onto the network. To this effect, the port on which an endpoint connects is activated and configured with certain characteristics and policies. This process is known as *authorization*. One example of authorization is the configuration of a port's VLAN membership based on the results of an authentication process. Another example is the dynamic configuration of port ACLs based on the authentication.

NOTE

For wireless access, the concept of a “port” is replaced by an “association” between client and access point. When authorizing a wireless device, the association is customized to reflect the policy for the user or device. This customization can take the form of the selection of a different wireless LANs, VLANs, or mobility groups depending on the wireless technology used.

In this two-phased process, authorization is the most relevant to virtualization. When an endpoint is authorized on the network, it can be associated to a specific VN. Thus, it is the authorization method that will ultimately determine the mapping of the end station to a VN. For example, when a VLAN is part of a VN, a user authorized onto that VLAN will therefore be authorized onto the VN.

The main authentication scenarios for the enterprise could be summarized as follows:

- Client-based authentication, for endpoints with client software
 - 802.1x
 - NAC
- Clientless authentication, for endpoints without any client software
 - Web-based authentication
 - MAC-based machine authentication

Regardless of the authentication method, the authorization could be done in one of the following ways:

- Assigning a port to a specific VLAN
- Uploading a policy to a port, in the form of ACLs, policy maps, or even the *modular QoS command-line interface* (MQC)

VLANs map into VRFs seamlessly and are the authorization method of choice when using a VRF-based transport virtualization approach. ACL authorization could be used to achieve policy-based transport virtualization. For a transport virtualization approach based on class-based forwarding, the ability to dynamically load a QoS policy onto the access device could prove useful.

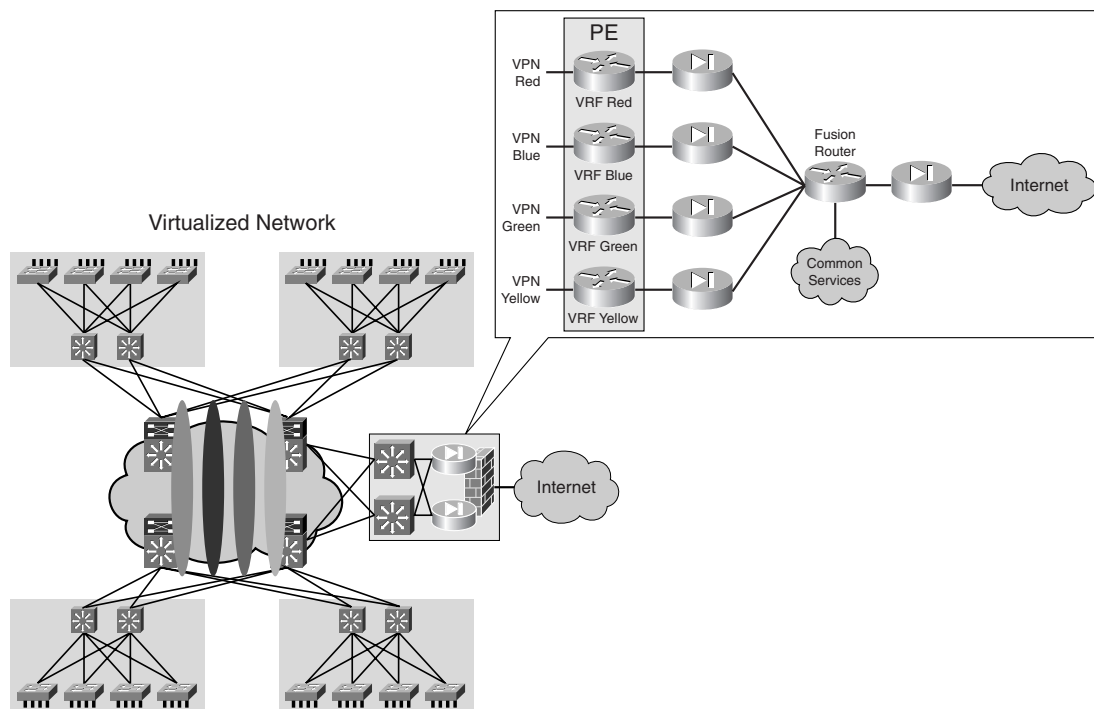
The current state of the technology provides broad support for VLAN assignment as an authorization alternative. In the cases where policy changes based on authentication are required and there is only VLAN assignment authorization available, a static assignment of a policy to a VLAN will provide the required linkage between the user authorization and the necessary policy. The policy will in effect be applied to the VLAN; as users are authorized onto different VLANs, they are subject to different policies.

Central Services Access: Virtual Network Perimeter

The default state of a VN is to be totally isolated from other VNs. In this respect, VNs could be seen as physically separate networks. However, because VNs actually belong to a common physical network, it is desirable for these VNs to share certain services such as Internet access, management stations, DHCP services, *Domain Name System* (DNS) services, or server farms. These services will usually be located outside of the different VNs or in a VN of their own. So, it is necessary for these VNs to have a gateway to connect to the “outside world.” The outside world is basically any network outside the VN such as the Internet or other VNs. Because this is the perimeter of the VN, it is also desirable for this perimeter to be protected by security devices such as firewalls and *intrusion detection systems* (IDSs). Typically, the perimeter is deployed at a common physical location for most VNs. Hence, this location is known as the central services site, and the security devices here deployed can be shared by many VNs.

The creation of VNs could be seen as the creation of security zones, each of which has a unique and controlled entry/exit point at the VN perimeter. Routing within the VNs should be configured so that traffic is steered to the common services site as required. Figure 3-8 illustrates a typical perimeter deployment for multiple VNs accessing common services. Because the services accessed through the VN perimeter are protected by firewalls, we refer to these as “protected services.”

Figure 3-8 Central Site Providing VN Perimeter Security



As shown in Figure 3-8, each VN is head ended by a dedicated firewall. This allows the creation of security policies specific to each VN and independent from each other. To access the shared services, all firewalls are connected to a “fusion” router. The fusion router can provide the VNs with connectivity to the common services, the Internet, or even inter-VN connectivity. The presence of this fusion router should raise two main concerns:

- The potential for traffic leaking between VNs
- The risk of routes from one VN being announced to another VN

The presence of dedicated per-VN firewalls prevents the leaking of traffic between VNs through the fusion router by only allowing established connections (connections initiated from “inside” the firewall) to return through the VN perimeter. It is key to configure the routing on the fusion device so that routes from one VN are not advertised to another through the fusion router. The details of the routing configuration at the central site are discussed in Chapter 8, “Traffic Steering and Service Centralization.”

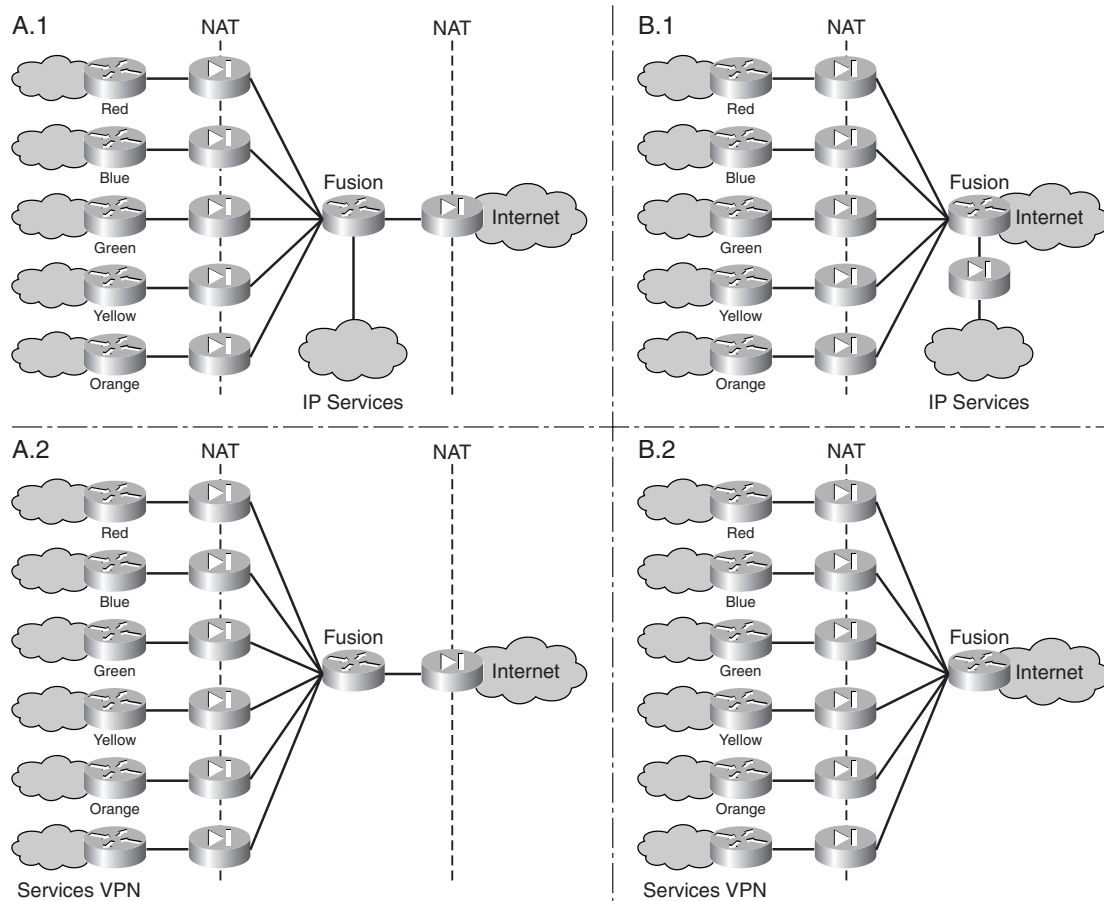
Figure 3-8 shows an additional firewall separating the fusion area from the Internet. This firewall is optional. Whether to use it or not depends on the need to keep common services or transit traffic in the fusion area protected from the Internet.

As mentioned, the common services could exist in a central location or in their own VN and therefore distributed throughout the enterprise. Depending on where the common services are located, the VN perimeter topology will vary. Figure 3-9 illustrates the different scenarios for common services positioning and the Internet firewall.

When the common services are not present, or are placed in their own VN (and therefore front-ended by a dedicated firewall context) the additional Internet firewall can be removed, as shown in scenario B.2 of Figure 3-9. If concern exists about transit traffic (between VNs or between a VN and the shared services area) being on the Internet, the firewall can be kept (see diagram A.2). The common services could be separated from the rest of the network by having their own firewall, yet not be included in a VN; this is shown in scenario B.1 in Figure 3-9.

For scenarios B.1 and B.2 in Figure 3-9, it is important to note that the fusion router is actually part of the Internet; therefore the *Network Address Translation* (NAT) pool used at the firewalls must use valid Internet addresses. The deployment of the optional Internet firewall should follow standard Internet edge design guidance, which has been extensively documented across networking literature. The reference designs proposed by Cisco and published at <http://www.cisco.com/go/srnd> are good sources of information. The “Data Center: Internet Edge Design” document contains a comprehensive discussion on the topic. We use scenario A.1 from Figure 3-9 to illustrate the relevant design and deployment considerations, but these considerations are applicable to other scenarios.

Figure 3-9 Common Services Positioning



Unprotected Services

In contrast with circuit-based technologies such as ATM or Frame Relay, most Layer 3 VPN technologies allow enough flexibility for traffic to be leaked between VNs in a controlled manner by importing and exporting routes between VNs to provide IP connectivity between the VNs. Thus, the exchange of traffic between the VNs may happen within the IP core and not have to pass through the VN perimeter firewalls at the central site. This type of inter-VN connectivity can be used to provide services, such as DHCP or DNS, that do not need to be protected by the central site firewall, or that would represent an unnecessary burden to the VN perimeter firewalls. Because of the any-to-any nature of an IP cloud, there is little

chance of controlling inter-VN traffic after the routes have been exchanged. We refer to these as “unprotected services.” This type of connectivity must be deployed carefully because it can potentially create unwanted back doors between VNs and break the concept of the VN as a “security zone” protected by a robust VN perimeter front end. Another consideration that must be made is the fact that importing and exporting routes between VNs precludes the use of overlapping address spaces between the VNs. We discuss the use of route-importing mechanisms for the creation of common services extranet VNs in detail in Chapter 8.

NOTE Although, these services are not protected by the VN perimeter firewalls, the IP segment to which they belong could potentially be head-ended by a firewall and therefore “protected.”

The deployment of protected services does not preclude the deployment of unprotected services and vice versa. In a virtualized network, a combination of protected and unprotected services is usually provided for the different VNs. For example, the DHCP and DNS services for several VNs may be shared and accessed in an unprotected manner, while all server farms and the Internet are also shared among the different VNs, but their access must be controlled by firewall policies and an IDS.

Summary

You can use many technologies to virtualize the enterprise network. Regardless of the technologies of choice, they must provide the functionality required in the three areas discussed:

- Transport virtualization
- Edge authorization
- Central services access (VN perimeter)

The network architect should be well aware of how these functional blocks interface with each other and always keep in mind that virtualizing the network must not come at the expense of important resiliency and performance characteristics in the network. However, because of the new technologies put in place, there will be an impact in the operations and processes for the maintenance of the network. In the long term, this impact is likely to be a positive one as new operational efficiencies are gained and operational costs tend to diminish.

It is also important to remember that when virtualizing a network, not everything must be migrated onto the VNs created. VN technologies are overlaid onto the existing operational network infrastructure. Therefore, the network continues to function as it did before the virtualization, but now has VNs overlaid on top of it. The endpoints using the network could belong to the original network or to a VN. This provides a clear path to a phased migration and support for groups that do not require a dedicated VN.