Network Virtualization—Guest and Partner Access Deployment Guide

This document provides deployment guidance for enterprises that want to provide Internet and limited corporate access for their guests and partners. Several solutions for guest and partner access challenges are proposed and analyzed in this document, at both the architectural and functional levels. For related information, see the following documents:

- *Network Virtualization—Access Control Design Guide* (OL-13634-01)
- *Network Virtualization—Network Admission Control Deployment Guide* (OL-13636-01)
- *Network Virtualization—Services Edge Design Guide* (OL-13637-01)
- *Network Virtualization—Path Isolation Design Guide* (OL-13638-01)

Contents

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

© 2007 Cisco Systems, Inc. All rights reserved.

Introduction

This document provides network architects with a general understanding of the various options that are available for solving the challenge of providing guest and partner access. This understanding allows the network architect to select a specific set of technologies and then to delve deeper into the specific implementation details. Challenges associated with segmenting user groups in the enterprise network are addressed in the solution framework.

Several technologies and techniques can be used to provide the functionality required in each area. Rather than attempting to provide detailed information for every available technology option within the scope of this document, a separate set of design guides provide more information on each of these technologies. This modularity provides the flexibility to combine various sections of the design guides to tailor the solution implementation documentation to better serve customer requirements and architect choices. These design guides are as follows:

- *Network Virtualization—Access Control Design Guide* (OL-13634-01)
- *Network Virtualization—Services Edge Design Guide* (OL-13637-01)
- *Network Virtualization—Path Isolation Design Guide* (OL-13638-01)

This document also provides an end-to-end analysis of guest and partner access challenges across all functional areas and all places in the network. This document provides enough information for the network architect to select a manageable subset of the available technologies, while referring to relevant sections in the individual design guides for the specific low-level implementation information. See the tables in Appendix—Design Guide Mapping, page 25 to locate the sections of the functional area design guides that are relevant for each solution.

Business Problem

Companies must currently provide network access for their customers, partners, vendors, contractors, and other guests to enable greater productivity, improved collaboration, and better service. By implementing a complete solution to offer guest and partner access service, enterprises can control network access, reduce or eliminate IT support for non-employee personnel, maintain full auditing capability, and keep their traffic securely separated from restricted internal resources.

The need for guest and partner access has evolved over the years. At one time, it was sufficient to provide guests and partners with just a chair and a phone. Today, with the availability of laptops, networked applications, and digital phone lines, visitors in enterprise facilities, at a minimum, require access to the Internet and the ability to run virtual private network (VPN) services. Partners may also need to access internal resources such as printers, web servers, and shared folders.

Guest networks are network connections provided by an enterprise to allow their guests to gain access to the Internet and to their own enterprise networks without compromising the security of the host enterprise. Authorized partners may also use guest networks to access certain parts of the corporate network.

The main technical requirements for a complete guest and partner access solution are as follows:

- Integration into the enterprise network, overcoming traditional solutions (modem or DSL ports or parallel networks)
- Seamless support for both wired and wireless users
- Logical isolation of guest traffic from the internal enterprise traffic
- Partner access to authorized corporate resources
- Accounting, filtering, content checking, and so on
- Providing guests and partners the ability to establish VPN connections to their own corporate networks
- Authentication and logging capabilities for guests

Three identity types are discussed in this document:

- Guests

  Guests are visitors to whom Internet access is extended. Guest users, once connected to the Internet, typically use VPN services to access their own corporate resources. Visitors typically require only DHCP and DNS services and need only a supported web browser (IE 6.0+, Firefox, Opera, and so on). The host organization may offer access via web authentication rather than providing dedicated modem ports or a separate physical network. The guest access profile is the default profile for unknown or unmanaged users and must also be isolated to prevent unauthorized internal access.

- Unmanaged partners

  Partners with unmanaged computers are similar to guest users but the host organization may grant them different access capabilities. These devices may or may not have an acceptable 802.1x supplicant and, if they do happen to fall into compliance, it is merely by chance and may or may not be a requirement for access. Future web authentication enhancements at the access layer will differentiate more distinctly between guests and unmanaged partner profiles and may allow the host to provide partners with access to corporate printers and internal web servers.

- Managed partners

  Partners with managed computers are allowed further access to corporate resources such as shared folders and social networking sites (wikis, blogs, and so on). Managed partner accounts require either identified credentials and approved 802.1x-capable supplicants that are properly configured or some other access technology, such as MAC authentication bypass (MAB).

Note MAB is an alternative to 802.1X that allows network access to devices (such as printers and IP phones) that do not have the 802.1X supplicant capability. MAB uses the MAC address of the connecting device to grant or deny network access.

Allowing users located in remote branch offices to gain Internet access is a typical requirement that is valid for large, medium, and small enterprises. This document presents and validates two solutions where guest and partner access occurs. The first describes a deployment where guest and partner traffic flows through the main enterprise site only. The second scenario is presented for large enterprises that span the globe and have several points-of-presence (POPs) to connect to the Internet. This is the case for the Cisco internal network, where guest and partner traffic from all Cisco offices worldwide is always backhauled to the closest geographical POP site.

Note Designs where Internet access can be provided directly at the remote branch locations (leveraging, for example, split tunneling mechanisms) are not within the scope of this document.

An example of a traditional solution to connect branch offices to the enterprise headquarters leverages a privately owned WAN, leased lines, ATM networks, and Frame Relay connections. The requirement to reduce costs has, in recent years, led to the adoption of a new type of connectivity between branch offices and headquarters. In these deployments, VPN solutions (mostly IPsec) are implemented to leverage the public Internet. Both these scenarios are addressed when describing the solutions to backhaul the guest traffic across the WAN to the main site.

End-to-End Network Virtualization Solution

To provide guests and partners with Internet access, a virtual network (also known as a segment) is created. This virtual network is logically overlaid onto the existing infrastructure and allows connectivity only to the Internet and not to any internal resources. Internet connectivity is provided through a control point that forces guests and partners to authenticate to provide legal disclaimers and maintain traffic accounting information on the various accounts. Ensuring that guests and partners (managed and unmanaged) connect only through their assigned virtual network keeps them separate from the employees and forces compliance with the guest and partner access policies of the enterprise. Per-user bandwidth contracts can be assigned for wireless clients, which can be used to restrict bandwidth to specific SSIDs and VLANs. For example, unmanaged partner accounts, while similar to guest accounts, can be provided higher bandwidth than guest accounts. Managed partner accounts are allowed further access to corporate resources such as restricted websites and shared folders.

End-to-End Overview

The goal of this solution is met with the following considerations:

- Identifying a user as a guest or partner and assigning them to the proper segment.
- Isolating guest traffic from the rest of the network while providing Internet access.
- Isolating unmanaged partner traffic from the rest of the network while providing higher bandwidth.
- Isolating managed partner traffic from the rest of the network while providing Internet, printer, and further access to approved restricted corporate resources.
- Providing network services to enterprise visitors. Network services include the following:
  - Network services—DHCP, DNS, and Internet. Access only to the Internet for unmanaged partner or guest accounts. Limited access to approved corporate resources such as printers, web servers, and shared folders for managed partner accounts.
  - Security services—Firewalls, load balancers, intrusion detection systems (IDS), SSIDs, wireless authentication mechanisms (802.11i/EAP), web authentication (wired and wireless), disclaimers, accounting, monitoring, and Auth-Fail-VLAN capability (wired).

Note IEEE 802.1x with restricted VLAN (Auth-Fail-VLAN) allows end devices that fail 802.1x authentication to be placed into a VLAN of choice.

The solution framework is divided into three functional areas (see Figure 1), each of which maps to one of the objectives:

- Access control
- Path isolation
- Services edge

Figure 1 Solution Framework—Three Functional Areas

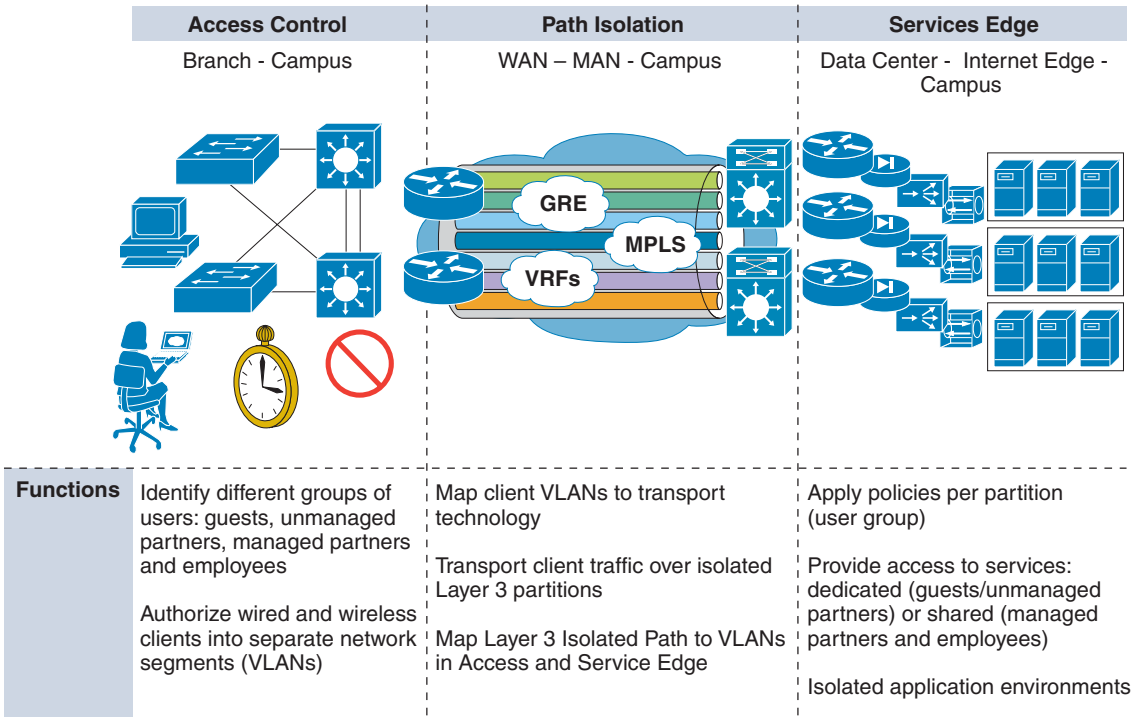

The end-to-end solution involves an optimal combination of chosen technologies available in each functional area.

Note The main goal of this deployment guide is to provide end-to-end guest and partner access solutions that are seamlessly valid for wired and wireless clients. This means that the technical elements comprising the path isolation and services edge functional areas are shared and valid for both categories of users. Cisco currently also offers a wireless-only guest access solution based on the use of WLAN Controllers. However, this solution is beyond the scope of this paper; more information can be found in the *Enterprise Mobility 3.0 Design Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns348/c649/ccmigration_09186a00807b59ca.pdf.

Access Control

Access control defines how the various users connect to the enterprise network. The goal is to provide a separate virtual environment for each group of users. For example, guest users and partners with unmanaged devices should be assigned to the guest virtual network (unmanaged partner accounts are

effectively guest accounts with higher wireless bandwidth capability), partners with managed devices should be assigned to a limited-access corporate segment, while an employee should be assigned to the employee virtual network.

Depending on the access method, employees may not be required to perform authentication. Because the various virtual networks are deployed on a common shared infrastructure, the physical access ports to the network are shared by the various groups. This implies that switch ports (for wired clients) and access points (for wireless clients) become shared network resources for internal employees, partners, and guests. A dynamic mechanism is necessary to differentiate employees from managed partners, unmanaged partners, and guests and to assign their port the appropriate policy. This policy ensures that the users of one group can access only their own virtual network while the users of other groups are assigned to their respective segments. The policy can be as simple as the assignment of the port or AP association to a specific VLAN. The VLAN belongs to a specific virtual network and therefore maps the user to the virtual network. In the case of a guest, this means recognizing that a user is a guest and confining them to the guest segment of the network and restricting their bandwidth capability. Devices in the guest segment of the network can reach only the Internet and are subject to traffic accounting controls.

Wired Media

Access control techniques for guests and partners connecting to an Ethernet port (wired access) include the following:

- Static port configuration—A port is statically assigned to a VLAN, which is in turn statically associated with a virtual network or segment. There is a guest, partner, and employee VLAN in every wiring closet. Ports are statically assigned to VLANs. Note that even deploying guest and partner VLANs in a static fashion impacts the underlying network infrastructure, so it must be planned in advance. Management of statically assigned ports can be quite time consuming if the network incurs constant changes.
- 802.1x guest VLAN—This feature dynamically assigns the port to the guest VLAN in the absence of an 802.1x supplicant. If there is a successful 802.1x authentication, the port is placed in the corresponding employee or managed partner VLAN.
- Auth_Fail_VLAN—This feature assigns the port to the Failed_Authentication VLAN when the client has a valid supplicant but not a valid login. The Failed_Authentication VLAN can be the same as the guest VLAN, therefore assigning 802.1x clients without a valid login to the guest segment. Note that when leveraging the auth-failed VLAN, it is not possible to differentiate between true guests and hackers who might try to access the network without having valid credentials. As a consequence, plan the network policies to allow only limited connectivity from clients deployed into the auth-failed VLAN, to avoid unauthorized access to the private enterprise network resources.

Note Each switch port might require both the guest and auth_failed features, because there is no easy way to determine whether a guest has an 802.1x supplicant. If the supplicant is not present, the host is considered to be a guest or unmanaged partner account. If the supplicant is present and the authentication fails, you can also consider the host to be a guest or unmanaged partner. When there is an 802.1x supplicant present and there is a successful authentication, the client is deemed a managed partner or an employee and is therefore placed in the appropriate authorized VLAN.

- MAC Authentication Bypass (MAB)—This feature can be enabled on an 802.1x port to allow the switch to use the MAC address as the client identity. In this case, the backend authentication server has a database of client MAC addresses that are allowed network access. For this solution to be applicable in a guest access scenario, you must also implement a mechanism to populate the backend

database with the MAC addresses of the clients. Cisco does not currently offer this capability, so it is up to each customer to provide it in a proprietary manner. The use of MAB is more feasible for managed partner access because the provisioning of the MAC database is easier in that scenario. However, note that MAC addresses can be easily spoofed, so appropriate security mechanisms must also be implemented, such as further segmenting the VLAN from the corporate network.

Wireless Media

The access control alternatives for guests connecting over wireless media include open authentication with dedicated guest SSID. Additional SSIDs can be established for unmanaged partners and managed partners. Alternatively, unmanaged partners can simply use the guest SSID. Also, with Maintenance Release 2 of the Airespace WLC 4.0 software (version 4.0.206.0), it is possible to configure multiple WLAN profiles with the same SSID and thus use multiple authentication types to authenticate users and place them into separate VLANs. Guest users can access the Internet via a designated open broadcast SSID to allow plug-and-play connectivity (to avoid broadcasting the SSID is not a valid form of security because a hacker can easily detect the SSID in use by simply sniffing the probe response messages in the air). Associating to the designated SSID results in guests eventually being assigned to a guest virtual network. Authenticated users are assigned to appropriate VLANs depending on whether they are managed partners or employees. The mapping between SSIDs and virtual networks can be achieved by the association of a VLAN with WLAN profile, depending on the wireless architecture deployed.

The access control functional area assigns users to the correct segment. Normally, this assignment is referred to as authorization and is linked to some means of authentication. Accounting can also be leveraged to provide the necessary information to maintain accounting records for the various users. However, in the case of guests, there is no widely adopted authentication mechanism. Therefore, Cisco has chosen to do the authorization on a guest VLAN with open authentication (no authentication). At this point, the guest users have not been authenticated; they have simply been identified as guests and assigned to a separate segment of the network. This is analogous to the previously described wired scenario where the 802.1x guest VLAN is used to provide network access to clients not equipped with 802.1x supplicants. Keep in mind that guests and partners must still be authenticated and authorized to access the Internet. This authentication is enforced at the services edge, which is covered subsequently in this guide. The access control functional area is simply assigning guests and unmanaged partners to a segment of the network in which they are not able to reach any of the enterprise resources, but are able to connect to a web authentication portal that allows or denies them access to the Internet. Access control is also used for controlling managed partner access as well. Access control can also provide the necessary information to be able to monitor the kind of traffic the user is putting on the network.

Path Isolation

After various types of users (guest, unmanaged, and managed partners) are deployed in their own segment, they should have access only to specific network resources depending on their profile. To achieve this, you can keep traffic logically isolated by using separate Layer 2 domains (VLANs or wireless domains) for guests and unmanaged partners, managed partners, and employees. To preserve end-to-end separation, those Layer 2 domains must be extended across the entire network. Extending Layer 2 domains end-to-end negates all the scalability and modularity benefits achieved by a hierarchical network design. IP routing is at the heart of the hierarchical design because of its ability to limit the size of broadcast domains and to lessen the impact of failures and changes by providing a modular structure that is capable of preventing problems from propagating and affecting the entire network. A mechanism to provide network virtualization while preserving the scalability and modularity of the routed network is necessary. Clearly, end-to-end Layer 2 extensions negate the desired scalability and modularity.

When the Layer 2 domains at the edge are connected to the routed core of the hierarchical network, the logical isolation achieved at the edge by the Layer 2 domains is lost. A mechanism to give continuity to those segments over the routed core is needed.

The following alternatives are available to maintain this logical traffic separation in the Layer 3 domain of the enterprise network:

- Distributed ACLs—ACLs can be configured at the frontier points between the edge Layer 2 domains and the routed core. These ACLs should ensure that hosts in one group can access resources only in their own group. Thus, a user in group A should be able to reach addresses of users and resources only in group A. This policy can be enforced by means of an ACL, provided that the IP prefixes belonging to a group are well-known. Keeping track of the various combinations of IP addresses that belong to a group is a cumbersome task and can reach its scale limit relatively quickly, especially when peer-to-peer connectivity is required within the segments. For certain applications, such as guest access, the requirement is for many-to-one connectivity. In this case, the use of distributed ACLs might provide a manageable mechanism for restricting guests to access only the Internet edge. The ACL should simply deny access to any internal prefix and allow access to the rest of the world (the Internet). This ACL is identical everywhere and is, therefore, relatively manageable. Distributed ACLs are presented here more as a legacy method of providing selective network access to different user groups. This method is not recommended for network virtualization projects because it lacks many of the advantages provided by true network virtualization techniques.
- Overlay of GRE tunnels interconnecting VRFs—Another mechanism to provide continuity over the routed network to the logical separation provided by VLANs at the edge is to use IP tunnel overlays. A tunnel overlay (either in a full or partial mesh) is created for each user group. Each tunnel overlay is mapped to the group VLANs at the various sites. For example, the traffic in a guest VLAN maps to the tunnel mesh created for guests, managed partner access is mapped separately as well, while all other traffic is treated normally (no tunnel overlay). Guest traffic being tunneled to specific places prevents the guests from reaching any enterprise resources not present in the guest segment. To associate the VLANs with the tunnel overlays, policy-based routing (PBR) can be used. However, this requires the use of distributed ACLs and therefore provides little added value when compared to a pure ACL approach.

  By associating the VLAN interfaces and the tunnel interfaces in a group to a VPN Routing and Forwarding instance (VRF), VLANs can be mapped to the required tunnel overlay. VRFs are considered as virtual routers (although they are not strictly that) to which different interfaces can be assigned. Assigning VLAN interfaces and tunnel interfaces to these virtual routers, or VRFs, effectively creates a virtual network that has its own links and routed hops. Thus, a virtual network built this way consists of VLANs, VRFs, and GRE tunnels, all working together to form a separate overlay topology. For the specific guest access scenario, there is an instance of a guest VLAN at every access point, a guest VRF at every distribution point, and a guest mesh of tunnels interconnecting the guest VRFs present at the distribution points. Similar considerations hold true for unmanaged and managed partner access deployments as well. A routing protocol must run between the VRFs and over the tunnel mesh to provide the necessary reachability information. The underlying infrastructure is designed according to well-known hierarchical and high resiliency principles. Therefore, the tunnel overlay enjoys these benefits. See Network Virtualization Deployment Scenarios, page 10 for a high-level deployment model. More information is provided in the *Network Virtualization—Path Isolation Design Guide*.
- VRFs at every hop interconnected with VLAN (802.1q) trunks—This approach basically creates multiple parallel networks. Each group of users has a VRF at every hop, and all the VRFs for one group are interconnected. To keep traffic from the various groups separate as they travel from hop-to-hop, dot1q trunks are used to provide logical point-to-point connections between the VRFs. For each group, this provides an end-to-end virtual network in which each routed hop is represented by a VRF and each connection is represented by an 802.1q logical link. In a traditional network, each hop is a router and each connection is a physical wire. VRFs allow you to have separate logical

routers, and 802.1q allows you to interconnect these with separate logical wires. This requires a routing protocol to run at each VRF to convey the necessary network reachability information. This model maps directly to the hierarchical model of network design and therefore enjoys the same benefits of scalability and resiliency that have become required in any network design.

- MPLS/BGP VPNs (RFC 2547)—This technique uses Multiprotocol Label Switching (MPLS) to dynamically create a tunnel mesh similar to the tunnel overlay Cisco created for the Generic Routing Encapsulation (GRE)-based solution. These dynamic tunnels are better known as label-switched paths (LSPs), which handle traffic forwarding, while Border Gateway Protocol (BGP) is used to carry routing information between the VRFs. The separation of the control plane and the data plane is the key to being able to create the LSPs dynamically. This is the most scalable technique of all of the techniques described, but it is also the most demanding in terms of platform capabilities.

Some of these techniques apply exclusively to the campus and others are better suited for the aggregation of branches over the WAN. For example, a hop-to-hop VRF technique is better suited for the LAN than for a WAN; the main reason being the need to control every hop in the network (including the core). A tunnel overlay solution is better suited for the WAN, where the tunnels allow you to segment without having control of every hop in the core of the network. Usually, these are service provider routers over which the enterprise has no control. Also, the aggregation of branches over the WAN usually follows a hub-and-spoke logical topology, which is well-suited for the implementation of a static tunnel overlay.

Whichever technique is used, it can be overlaid onto the existing infrastructure. This means that the network continues to function as usual and only traffic that is steered into the created VPNs is isolated or segmented. When providing support for guest access, the requirement is for isolation of the guests only. This can be achieved by creating an isolated path for guests and assigning all guest traffic to this isolated path. Isolating the guests without altering the behavior of employee traffic is the approach taken to provide Internet access to guests. Managed partners may need access to some corporate resources such as printers, shared folders, and web services. These resources need to be segmented from the employee network. Regular employee traffic continues to be forwarded as normal without having to create a dedicated segment.

Services Edge

When the groups (employees, partners, and guest in this scenario) have been separated, they need access to certain services. Some of these services are dedicated to each group, while others are shared among several groups. Employees require access to their data centers, network services (DHCP servers, DNS servers), and many more resources including the Internet. Partners require access to the resources mentioned above. Guests and partners require access to network services (such as DHCP, DNS, or web authentication mechanisms), as well as the Internet. The Internet represents, in this case, a resource that is very likely to be shared between all users, while other services are most likely dedicated. The services edge provides the mechanisms necessary for users from different groups to access common services without compromising the security gained by isolating the groups from each other. The services edge also provides access to services that are dedicated to each specific group. To achieve this, it provides logical connectivity and security mechanisms over shared facilities, such as firewalls, load balancers, VPN concentrators, or even intrusion detection systems. For the purposes of guest and unmanaged partner connectivity scenarios, this topic is limited to the sharing and virtualization of firewalls to provide Internet connectivity, and to the provisioning of authentication and accounting services (DHCP, DNS, web-auth, FW instances) that are dedicated to these user segments. For managed partners, network services such as DHCP, DNS, web servers, printers, and shared folders may be able to be shared with employees.

The enterprise policy can require the guest or unmanaged partner to accept a legal disclaimer before being able to access the Internet, or it might track and log user activity while browsing the web from the enterprise network. Guest and unmanaged partner machines are usually not under the administration of

the enterprise IT department, which means that they are most likely not configured in accordance with specific enterprise security policies. Also, it cannot be assumed that these machines are equipped with specific authentication software (as, for example, an 802.1x supplicant), and even when that is the case, guests most likely do not have valid credentials to successfully complete the 802.1x authentication process. For these reasons, it is assumed in this document that the only software the guest or unmanaged partner can leverage to go through an authentication and authorization process is a web browser, commonly found on any machine.

The following two Cisco products can be used to perform a web authentication process in this context:

- Cisco Clean Access (CCA, also called Cisco NAC Appliance)
- Cisco Wireless LAN Controllers (WLC) using the internal web authentication feature or an external web authentication server such as BroadHop

Note Broadhop integration with the WLC was not tested as part of this deployment guide; however, it has been tested and additional information is available in the *Enterprise Mobility 3.0 Design Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns348/c649/ccmigration_09186a00807b59ca.pdf

The deployment models that are described in this document leverage web authentication devices in-band; this implies that all the traffic originated from, and destined to, the guest subnets defined in the enterprise network is always enforced through these appliances.

Network Virtualization Deployment Scenarios

This section covers the following network virtualization deployment scenarios:

- Option A
- Option B

Option A

This end-to-end proposition includes the following components:

- Access control—802.1x Guest VLAN, 802.1x Auth-fail VLAN, and dedicated open SSID for guest/unmanaged partners. 802.1x authentication, MAB and dedicated SSID with 802.1x authentication for managed partners.
- Path isolation—Separate VRF+GRE overlays for guest/unmanaged partners and managed partners.
- Services edge—Cisco Firewall Services Module (FWSM) and an in-band web authentication appliance with dedicated services (DHCP, DNS, and so on) for guest/unmanaged partners. Shared or dedicated services (in the data center) for managed partners.

Figure 2 shows the end-to-end picture of what is proposed in the campus. All traffic is handled in the traditional way, and only guest and unmanaged partner traffic is segmented.

Figure 2 End-to-End Network Virtualization

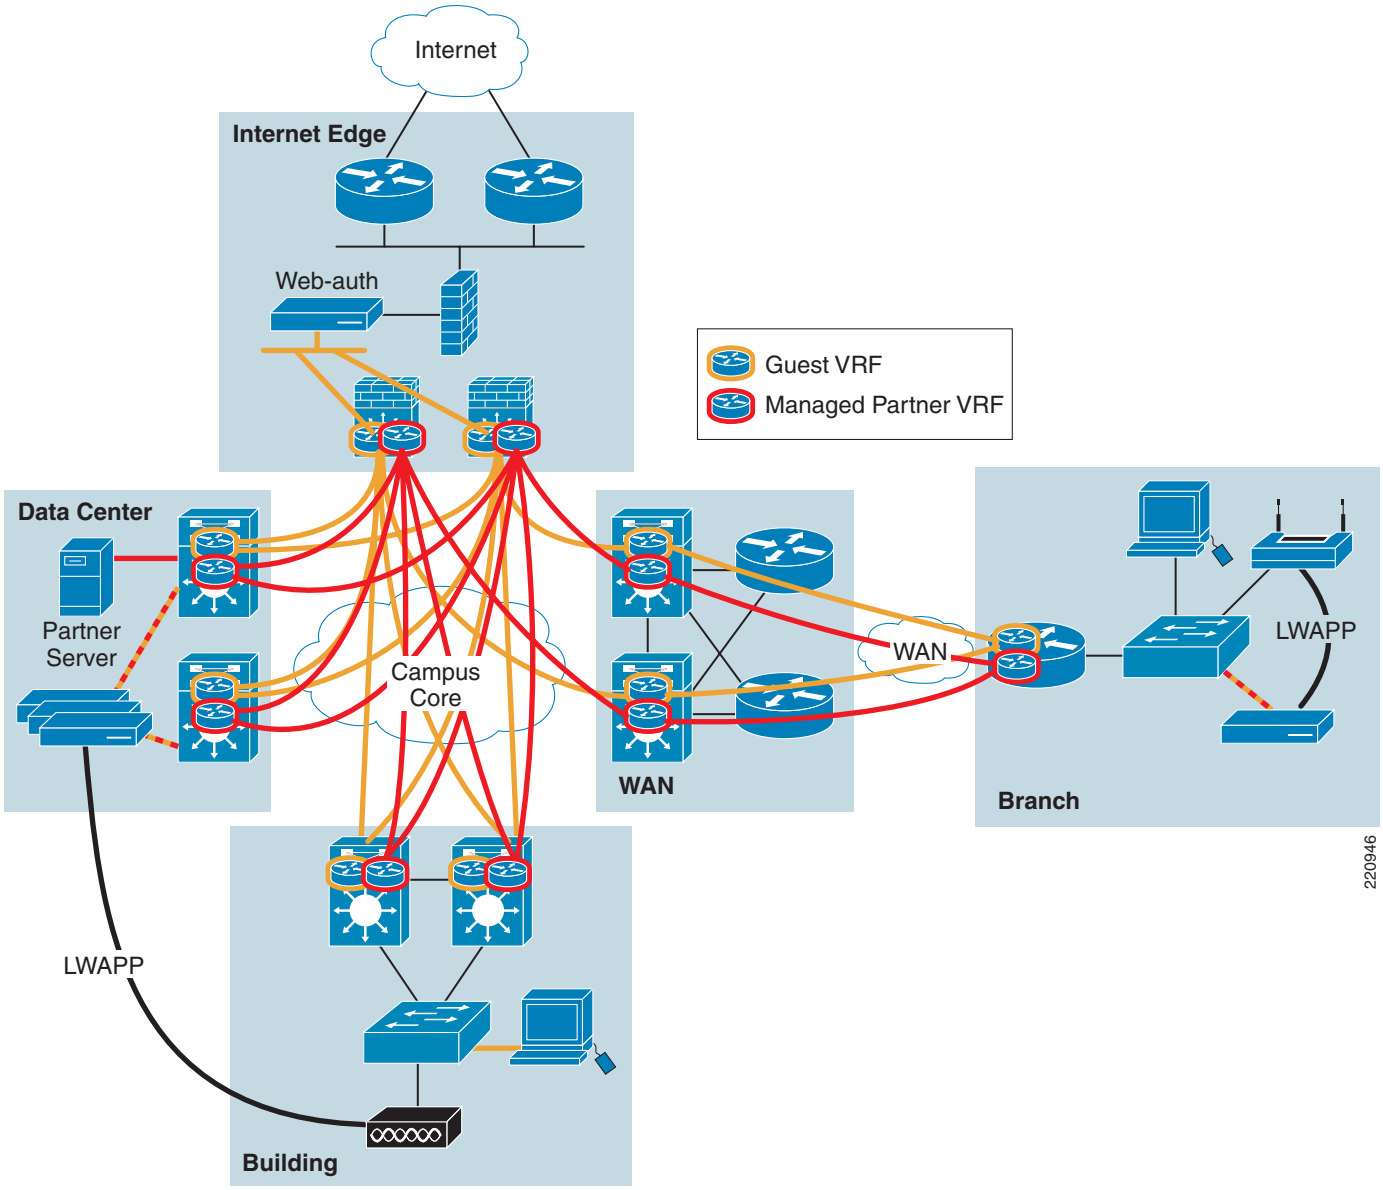

Note For wireless specific deployment information, see the *Enterprise Mobility 3.0 Design Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns348/c649/ccmigration_09186a00807b59ca.pdf

Access Control

At the edge, guests and unmanaged partners that use a wired connection are dynamically connected to the guest VLAN when they do not have an 802.1x supplicant. For guests and unmanaged partners with an 802.1x supplicant, the authentication fails and so they can be placed in the authentication-failed VLAN, which can be configured to match the same value as the guest VLAN.

Note This approach assumes that the employees in the enterprise are leveraging 802.1x to access the network. Providing guest and partner access in deployments leveraging 802.1x capabilities is not a trivial task; more considerations on this topic can be found in the *Network Virtualization—Access Control Design Guide*.

When guests and unmanaged partners use wireless access to reach the network, they are required to use a specific SSID, which eventually maps to a guest VLAN or a guest VRF in the wired network (the choice depends on the wireless architecture deployed). The VLAN or the VRF are part of the guest virtual network and allow the guest or unmanaged partner to access the web portal and, eventually, the Internet, without providing connectivity to the internal enterprise network. The following are the two main deployment models for wireless networking:

- Distributed wireless controllers
- Centralized wireless controllers

Figure 3 shows these two wireless deployment models.

Figure 3 Two Wireless Deployment Models

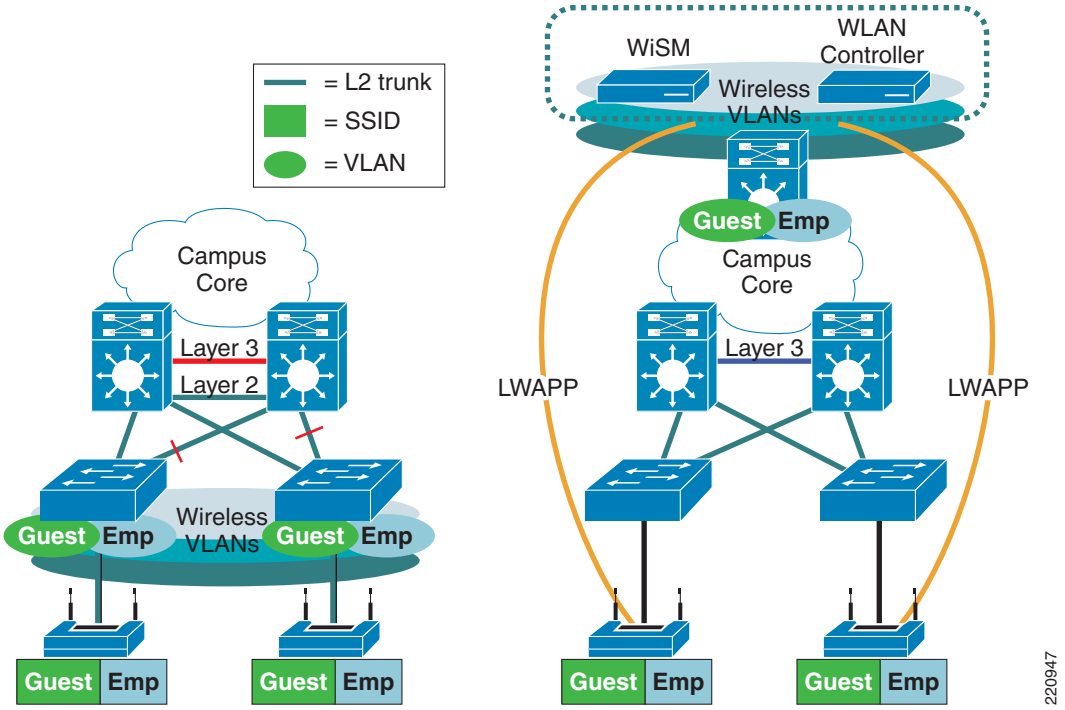

The diagrams above refer to both standalone (distributed) and controller-based (centralized) wireless deployments. Cisco recommends centralized wireless deployments moving forward, but there are many successful distributed deployments in place as well. The distributed wireless access model leverages the wired network virtualization model and maps traffic from the guest SSID into the guest VLAN at the access point. Thus, the controller trunks with the access switch and sends traffic tagged with the guest and unmanaged partner, managed partner, or employee VLAN 802.1q identifier.

The centralized wireless controller model creates its own tunnel overlay and consolidates all wireless traffic at a centralized controller. Based on the SSID being used, the wireless architecture can provide separate tunnel overlays. Traffic from each SSID travels over a separate tunnel to the central controller. The combination of SSIDs, users, and tunnel overlays is often referred to as a mobility group. Thus, the

wireless architecture is providing its own level of logical isolation. At the centralized controller, it is necessary to map the wireless traffic from each mobility group into the appropriate virtual network on the wired infrastructure. To achieve this, the mobility groups can be mapped into a guest VLAN at the central site or directly to a guest VRF at the central site.

Note Detailed information on the various wired and wireless access options is documented in the *Network Virtualization—Access Control Design Guide*.

The conceptual approach for the access control functionality is the same in the campus and in the branch. This guide provides the best practices for using the specific platforms in these two scenarios.

Path Isolation

The first stage of the path isolation solution is the creation of the various VLANs required at the network access. This is very straightforward in the case of guests and unmanaged partners because all you have to do is add a new VLAN for guests, following the same guidelines used for the deployment of other VLANs in the network access. This implies adding a new guest VLAN in each access switch. For the branch, this translates into a single new VLAN, because there is usually a single access switch. For the campus, there is a separate guest VLAN for each access layer switch. This is true when you follow the best practices for campus design at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/cdccont_0900aecd801a8a2d.pdf

Managed partner access requires a similar approach but care must be taken that only approved resources are accessible via the managed partner VLAN. Additional authorization may be needed via Active Directory or LDAP services to ensure only users with valid logins can gain access to the resources.

The best practice for campus design is to not span VLANs across multiple access switches. Depending on the access aggregation model being used, this can cause VLAN proliferation proportional to the number of access switches being aggregated. A traditional aggregation model uses a Layer 2 connection between access and distribution; in this model, every VLAN created in an access switch must be present on both distribution switches in the block.

Alternatively, a routed access can be used, in which case the VLANs created at the access do not impact the configuration at the distribution. The benefits of this simplification of the Layer 2 portion of the network must be carefully measured against the challenges of implementing Layer 3 isolation beginning at the access rather than the distribution.

Figure 4 shows both the traditional and routed access campus deployments.

Figure 4 Traditional and Routed Access Campus Deployments

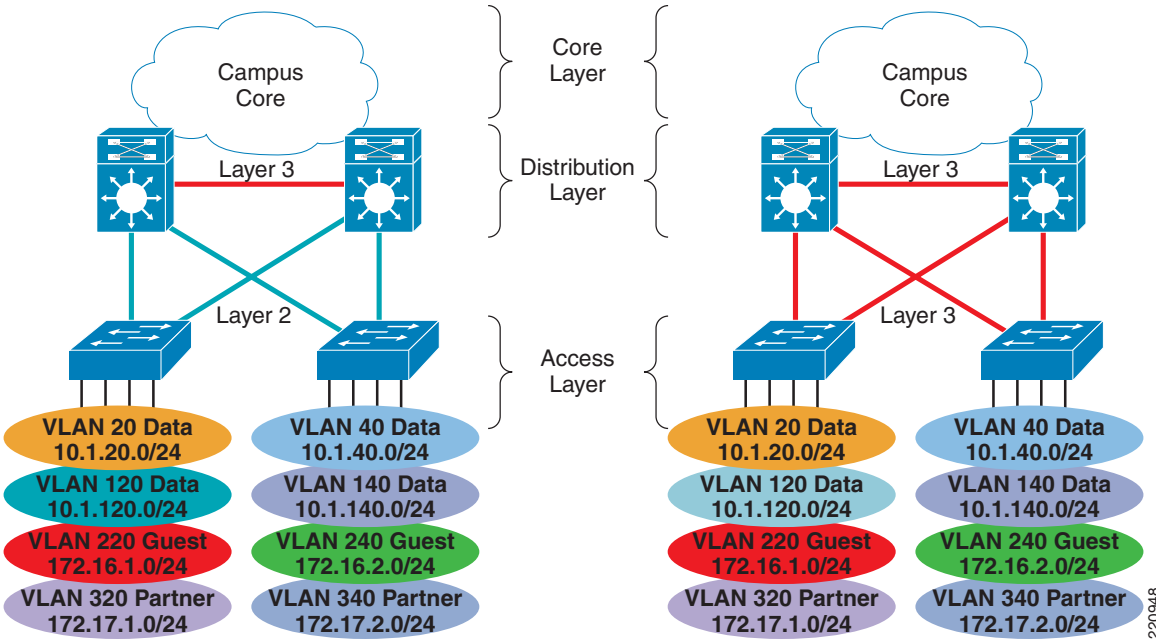

Whether the VLANs are terminated in the access or the distribution, you must map them to some kind of Layer 3 VPN. In the model used in this guide, the VLANs are terminated at the distribution in the campus. However, keep in mind that this can be generalized to include the devices at the border between the switched and routed portions of the network. In the case of a routed access, these are the access switches; in the branch, it is most likely the access router that provides the first Layer 3 hop.

Each guest and partner VLAN is kept separate from the rest of the network by mapping the VLAN to a VRF at the first Layer 3 hop. Thus, the IP default gateway for all devices in the guest VLAN is an address in the guest VRF. By making the VLAN part of the VRF, the VLAN (and its associated IP prefix) is removed from the original global network and is isolated from the rest of the network, because the VRF is an isolated routing table. The result is similar to having connected the guest VLAN to a separate physical router that is not connected to the original global network; traffic in the guest and partner VLANs does not have a route or connection to the original network, nor does the original network know how to get to the isolated VLAN or subnet.

So far, VLANs have been created and associated to isolated routing tables so that they cannot communicate with anything beyond their first hop. As described previously, a campus using a traditional Layer 2 access has several VLANs for a single user group (one per access switch). At the distribution are several VLANs, one for each access switch. These VLANs must all be associated to a single VRF. The guest and partner VLANs are simply different subnets that all belong to the same guest or partner virtual network, so all the VLANs in a distribution block can now communicate with each other, but they cannot reach any devices over the core.

Note Communication between different guest or partner subnets defined in the same campus building can be prevented by configuring ACLs on the first L3 hop devices (distribution layer or access layer switches, depending on the implemented campus design). Communication between subnets belonging to separate buildings can instead be achieved by configuring appropriate policies on the hub devices.

The next step is to interconnect the VRFs in the various distribution switches. A GRE tunnel overlay over the IP core can be used to interconnect all guest or partner VRFs at the various distribution switches. This basically creates an overlay logical topology across the core into which the VRFs at the distribution can plug. This overlay is used to carry routing updates and traffic between the VRFs. Because the VRFs are isolated and have only tunnel and VLAN interfaces associated with them, this network is effectively isolated from the rest of the infrastructure. Figure 5 shows the original global network and the overlaid virtual network sharing the same infrastructure.

Figure 5 Guest and Partner Overlay Networks in the Campus

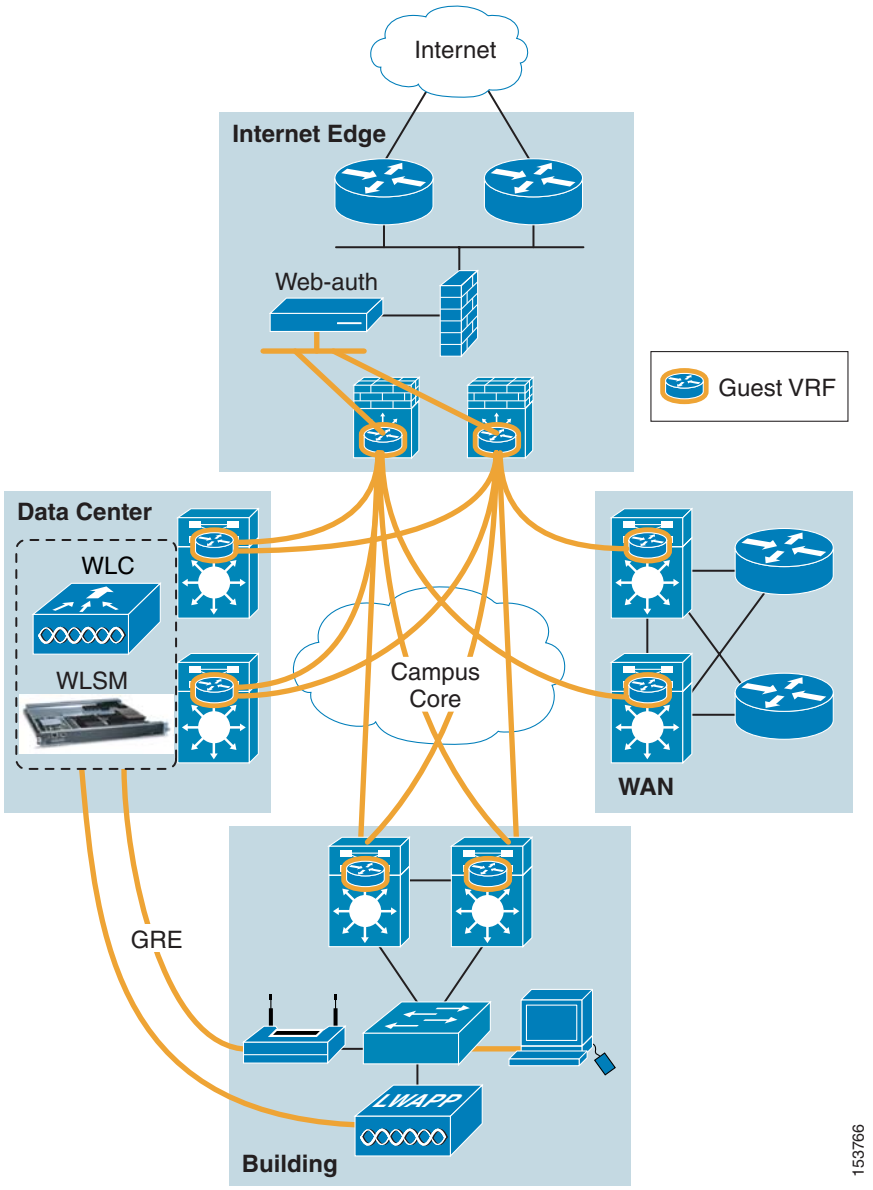

This concept can be implemented with point-to-point GRE tunnels, as well as multipoint GRE tunnels. The details on how to do this are provided in the *Network Virtualization—Path Isolation Design Guide*. For the specific guest and partner access application, keep in mind that the requirement is for many-to-one connectivity, in which many guests or partners are accessing a central services area that provides network services, AAA services, and, ultimately, Internet access.

Whether in the campus or aggregating branches over the WAN, GRE termination hubs are deployed, into which are brought all the GRE tunnels required to create the logical guest topology. Cisco recommends that separate logical topologies (tunnel overlays) be created for the campus and the WAN. This allows you to preserve the modularity and scalability provided by a WAN aggregation block, which clearly delimits the frontier between the WAN and the LAN. This modular approach also allows you to choose a variety of path isolation techniques in the LAN and WAN in an independent way. Thus, you can combine a hop-to-hop approach in the LAN with a tunnel-based approach in the WAN, or even change the WAN isolation approach in the future without affecting the LAN. Figure 6 shows the combination of LAN and WAN isolation logical topologies.

Figure 6 LAN and WAN GRE Tunnel Overlay (Private WAN)

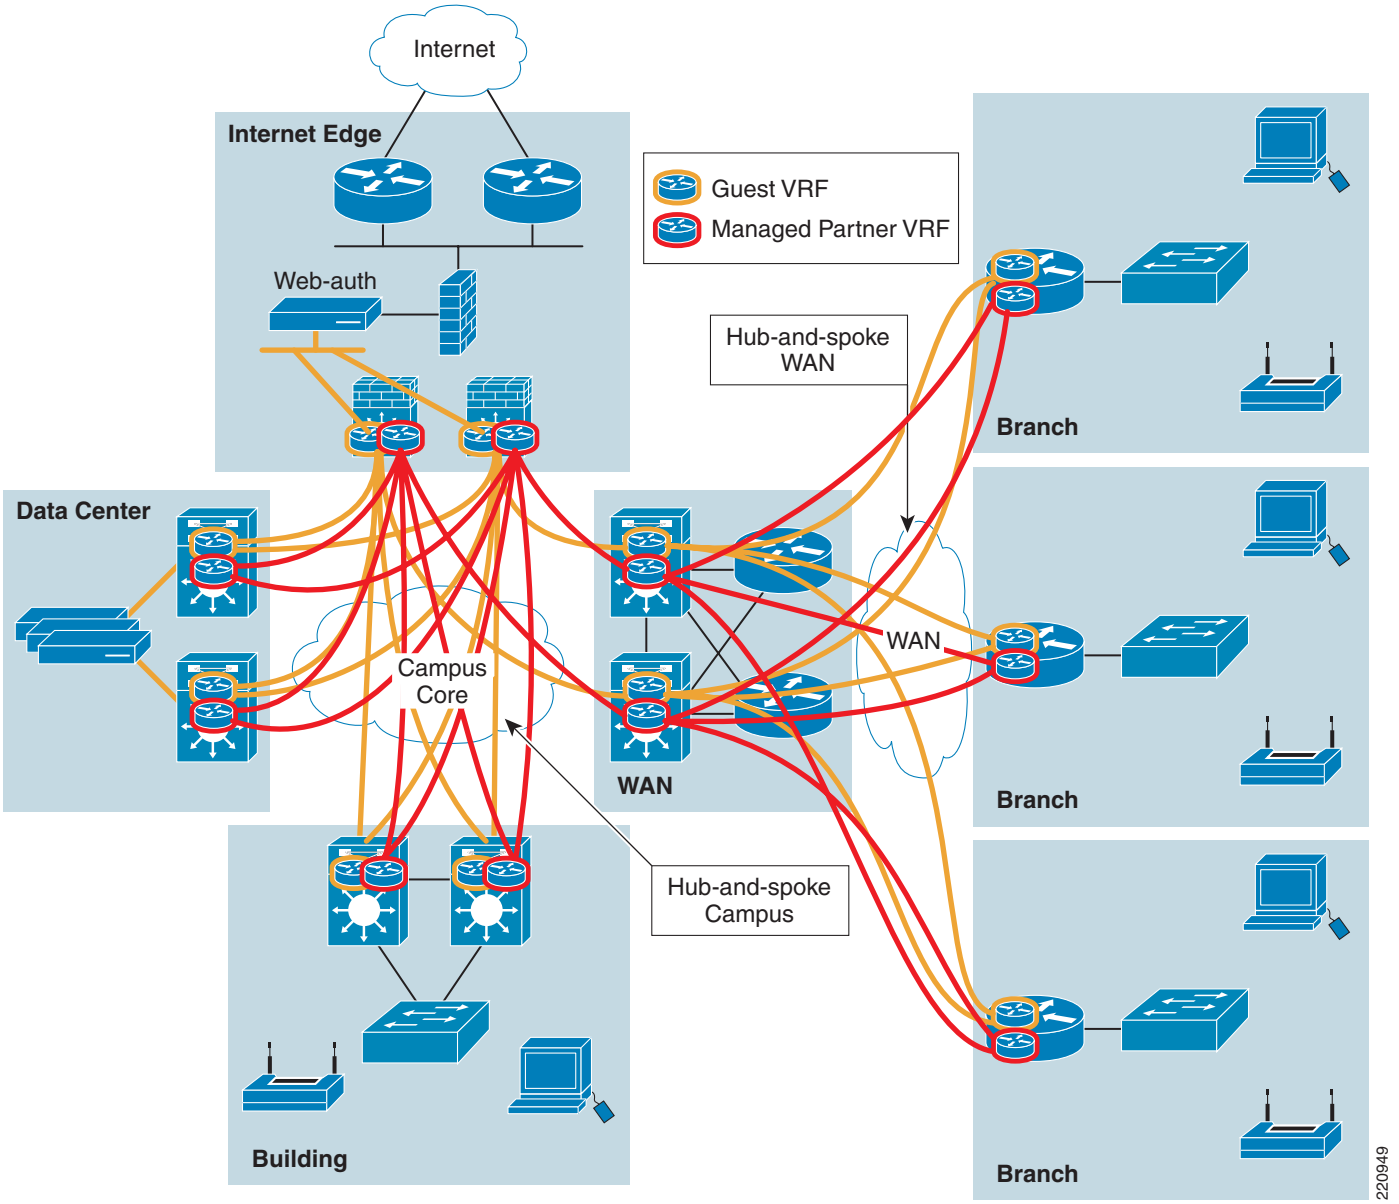

Note that the VRFs provide the point of termination for the LAN and WAN virtual networks, and also provide the mapping between the LAN and WAN segments. Thus, a VRF in the WAN aggregation module is associated to a tunnel in the WAN and a tunnel in the LAN. Figure 6 shows the scenario for a private WAN deployment in which the WAN aggregation module and the Internet edge module are separate. In this case, the LAN hub for the overlay topology is at the Internet edge module while the WAN hub for the WAN overlay topology is at the WAN aggregation module. Note that the WAN hub acts as a spoke to the LAN hub.

When providing WAN access over the Internet, the WAN aggregation and the Internet edge are in the same module. Therefore, the same set of routers serve as the WAN aggregation hub and the LAN aggregation hub simultaneously, as shown in Figure 7.

Figure 7 LAN and WAN GRE Tunnel Overlay (Internet)

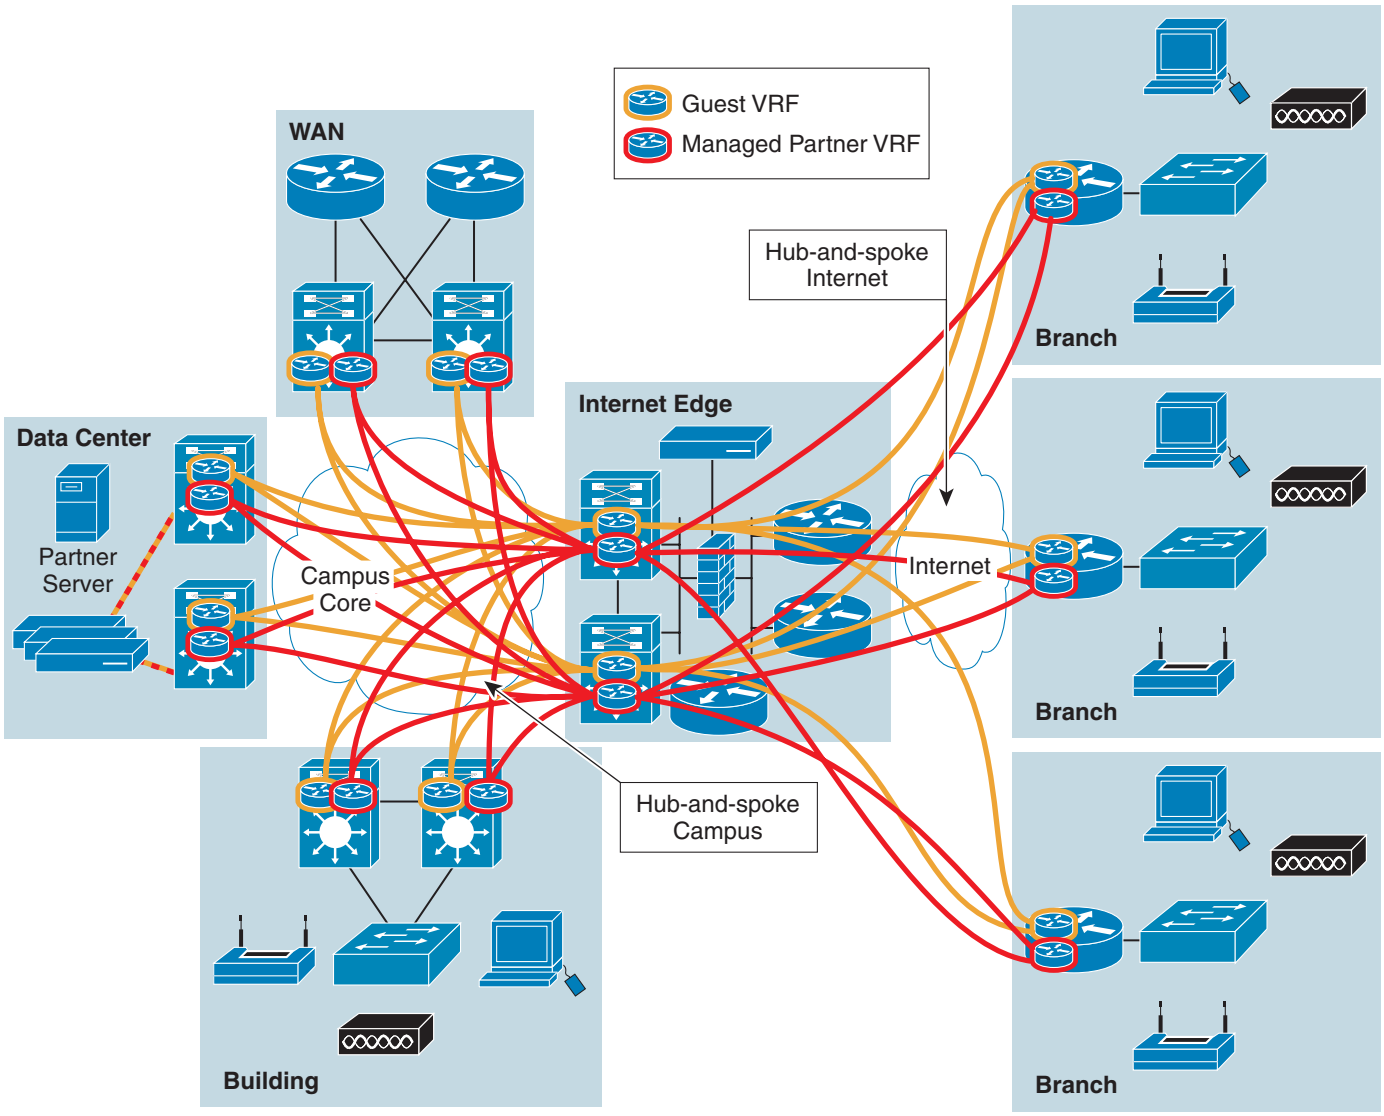

Figure 7 differs from the others in that it shows Internet traffic, as opposed to the non-Internet traffic in Figure 6. GRE tunnels over Dynamic Multipoint VPN (DMVPN) requires encryption while private WAN links may not require encryption.

One common requirement over shared IP networks is encryption. There are many ways to provide an encrypted overlay in the WAN. In the context of this document, the specific encryption solution adopted is irrelevant; the assumption is that an encrypted pipe exists between the enterprise branches and the main site, and it is leveraged as a transport to backhaul the guest traffic originated at the remote locations, as shown in Figure 8. More information on this topic can be found in the *Network Virtualization—Path Isolation Design Guide*.

Figure 8 GRE Tunnels over DMVPN

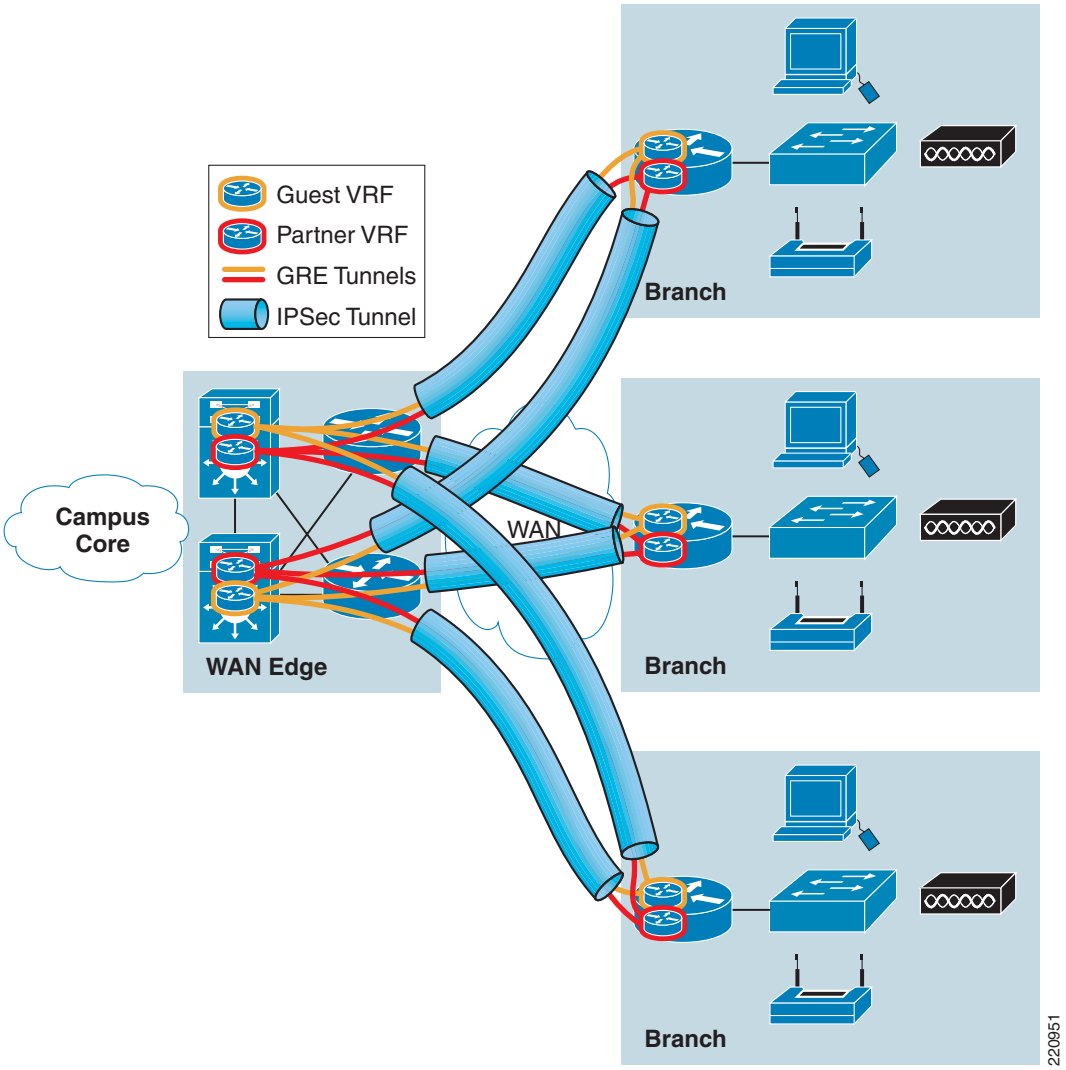

In this diagram, there is an encrypted link between the branch and the campus. Both groups are in IPsec tunnels traversing the Internet.

Services Edge

In the services edge functional area, guests, partners, and employees access the following services:

- DHCP servers
- DNS servers

- Web authentication and accounting servers
- Internet

This proposed solution assumes that each group accesses only those services to which they are allowed access, and shares limited services such as DNS, DHCP, some web servers and printers, and the Internet connection. However, it may be required to separate guests and unmanaged partners from the servers used by managed partners.

To share the Internet connection among various groups of users, the services edge should provide secure connectivity between each group and the Internet edge routers, as shown in Figure 9.

Figure 9 Shared Internet Access

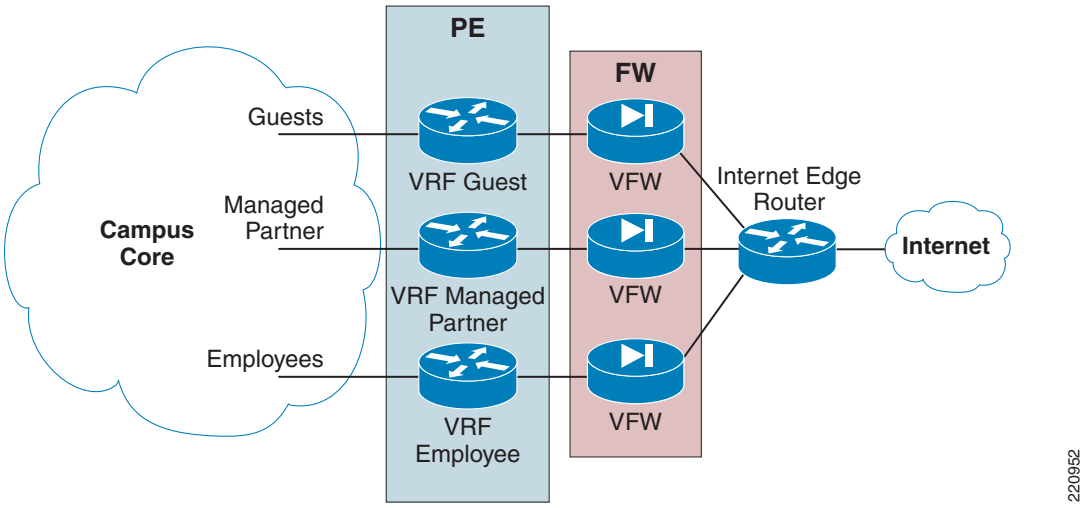

As shown in Figure 9, each group is connected to its own firewall, and all the various firewalls are connected to an Internet edge router that is shared by all VPNs. The firewall and routing configuration in this scenario is critical to preserving the isolation between the various user groups (employees, partners, and guests in this case), because this area has the potential to become a transit area between virtual networks. Detailed configuration guidelines are provided in the *Network Virtualization—Services Edge Design Guide*. This document simply states that the default state of a firewall limits any type of inter-VRF connectivity.

As the number of groups increases, so does the number of firewalls required. Cisco firewalls can be virtualized, thus allowing a single physical firewall to dedicate separate logical firewalls to each group. Furthermore, the entire topology shown in Figure 9 can be deployed inside a single device. The virtualization capabilities of the Cisco Catalyst 6500 and the integrated FWSM allow the flexibility necessary to implement this entire topology in a single device. The details on how to achieve this are documented in the *Network Virtualization—Services Edge Design Guide*.

Employee services such as DHCP and DNS servers should continue to be offered according to the best practices already adopted by the enterprise. The guest services, especially the web authentication services, require more consideration.

The goal is to force any guest or unmanaged partner attempting to access the Internet through a web authentication server. There are three main reasons for doing this:

- To prevent anyone from being able to exploit the enterprise network for Internet access
- To enforce the acceptance of a legal disclaimer before allowing access to the Internet from the enterprise network
- To enforce traffic and session monitoring and accounting

Guests and unmanaged partners are subject to the IP routing in the guest segment of the network. When attempting to access the Internet, guest traffic is routed to the bank of firewalls at the Internet edge. At that point, it is necessary to insert an inline web authentication mechanism so that the required control is enforced before the guest or unmanaged partner accesses the Internet.

Figure 10 shows how their devices are inserted in the services edge topology.

Figure 10 Services Edge with Web Authentication for Guests and Unmanaged Partners

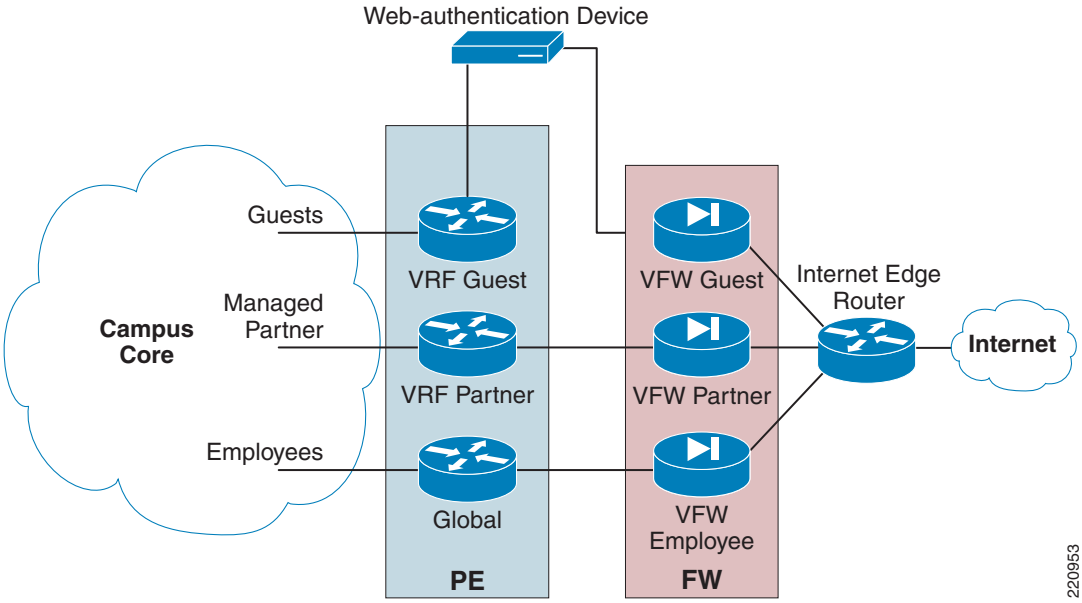

As shown in Figure 10, all guest and unmanaged partner traffic must traverse the web authentication server to reach the Internet. When the guest attempts to access the Internet for the first time, the web authentication server redirects the guest request to an authentication page. The guest then logs in with a set of credentials provided by their host. This set of credentials is associated with the specific user and must be requested in advance by the host of the user. The value of using specific credentials is that it allows the web authentication device to track and log accounting information for each session of each guest or unmanaged partner.

Cisco recommends the use of Cisco NAC Appliance for web authentication. DHCP functionality is either deployed in the web authentication device or behind the device on a separate server, as shown in Figure 11. Note that the web authentication servers do not provide DNS servers, only DNS relaying, so a separate DNS server is necessary. The implementation details are provided in the *Network Virtualization—Services Edge Design Guide*.

Figure 11 Web Authentication, DNS, and DHCP Deployment

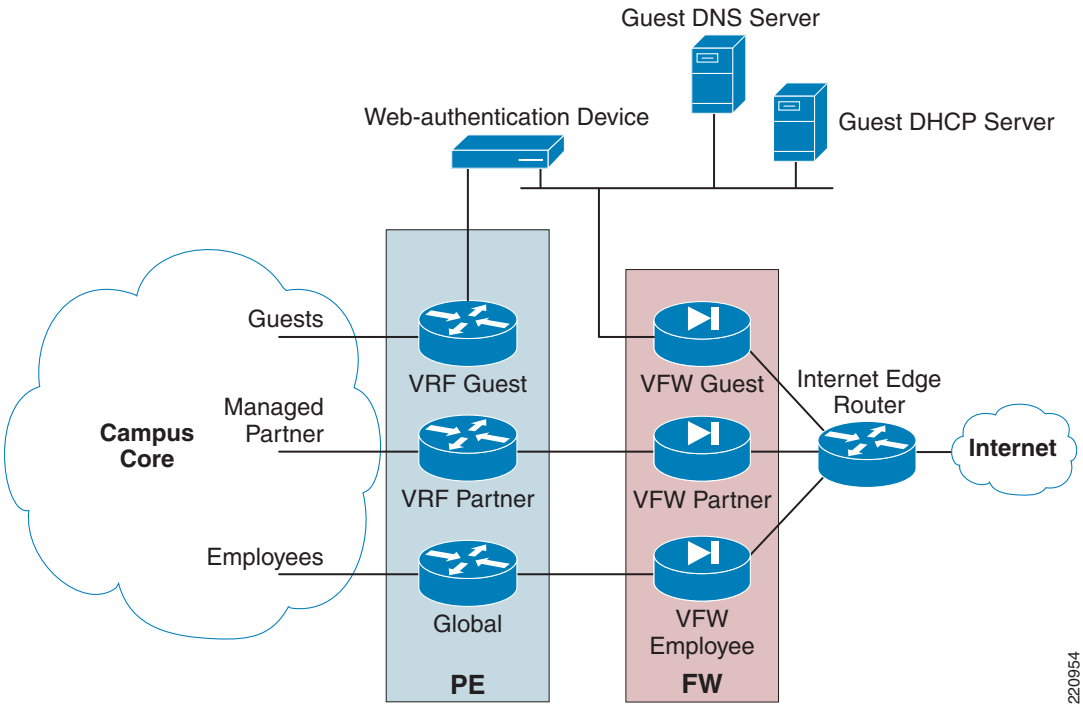

Note A concern with guests and unmanaged partners is that they are not necessarily subject to the control of the host organization regarding the traffic that they may generate. As such, providing guest and unmanaged partners with private DHCP and DNS servers, dedicated for guest use but controlled by the host organization, is the option recommended in this guide.

Option B

This end-to-end proposition includes the following components:

- Access control—802.1x Guest VLAN, 802.1x Auth-fail VLAN, and dedicated open SSID for guest/unmanaged partners. 802.1x authentication, MAB, and dedicated SSID with 802.1x authentication for managed partners.
- Path Isolation—MPLS VPN, with a dedicated VPN for guest/unmanaged partners and managed partners.
- Services edge—Cisco Firewall Services Module (FWSM) and an in-band web authentication appliance with dedicated services (DHCP, DNS, and so on) for guest/unmanaged partners. Shared or dedicated services (in the data center) for managed partners.

Access Control

Access control for Option B is handled in the same way as for Option A (see Access Control, page 11). Wired clients are assigned to the guest VLAN, either through a static configuration or, when they do not have a valid 802.1x supplicant, leveraging the 802.1x guest VLAN functionality. Wireless clients must use a specific SSID to access the guest segment.

Path Isolation

The same considerations for the creation of guest VLANs at the edge of the network that are described for Option A in Path Isolation, page 13 can be repeated here. What differs in Option B is the path isolation solution being implemented to preserve the logical traffic separation, achieved by the access VLANs as traffic traverses the routed portion of the network.

One way of keeping traffic originating in a VLAN from reaching VLANs or subnets belonging to other groups or segments is to use ACLs at the first Layer 3 hop that the VLAN encounters. These ACLs are used to restrict the prefixes that can be reached from a particular VLAN. As the number of potential destinations increases, so does the size and complexity of the ACL. ACLs cease to scale fairly quickly, so Cisco recommends that they be used to restrict reachability only for groups that require access to a centralized resource, but not for groups that have resources distributed throughout the enterprise. The simplicity of the ACLs also depends on how the IP address space of the enterprise is organized. Having an address space that allows summarization enables the creation of much simpler ACLs and is typically recommended when deploying enterprise networks. *Network Virtualization—Path Isolation Design Guide* describes the best practices for the creation of a simplified ACL to maintain isolation as guests and unmanaged partners traverse the routed enterprise network.

Services Edge

The Services Edge for Option B is the same as the Services Edge for Option A. At the services edge, guests, partners, and employees access the following services:

- DHCP servers
- DNS servers
- Web authentication and accounting servers
- Internet

These services are deployed in a dedicated manner, sharing the Internet connection only among guests, partners, and employees. The main goal is to divert guest and unmanaged partner traffic headed to the Internet to a web authentication server. The dedicated DHCP, DNS, and firewall services are located behind the web authentication server. All guest traffic must be diverted to the web authentication server, as shown in Figure 12. To divert the traffic, PBR can be used to route traffic to the web authentication server based on its source IP address. Being able to summarize the guest subnets helps to simplify the ACL necessary to achieve the anticipated PBR.

Figure 12 Use of PBR in the Internet Edge

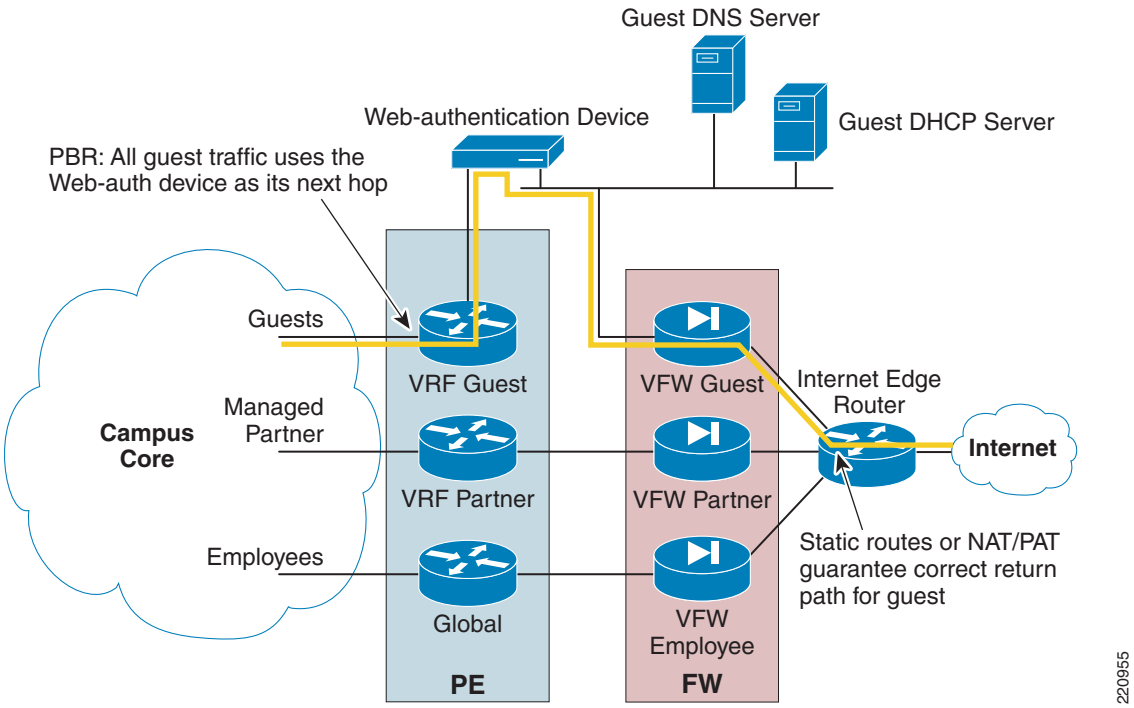

Both the DHCP and DNS services are behind the web authentication server because the web authentication server has relay functions for DHCP as well as DNS. As shown in Figure 12, guests and unmanaged partners have their own firewall instance in which guest-specific rules can be applied, independent of the managed partner or employee rules. Detailed implementation information is available in the *Network Virtualization—Services Edge Design Guide*.

The return traffic requires more specific routes to traverse the guest firewall. Rather than using PBR to divert this traffic, you can instead use static routes or leverage unique NAT address pools for guests if you are translating the guest addresses for the Internet. PBR is of little benefit here because you can base the decision on the well-known destination addresses (the guest prefixes).

Integration with Other Cisco Subsystems

IP Communications

The proposed model described in this document does not provide guests with IP communication (IPC) services. Because employee traffic continues to be treated in its traditional way and it is not migrated to a virtual network (segment), employee traffic is not affected by the addition of a guest segment to the network. The only interactions between guests and IPC occur if you want to allow network connectivity from the Ethernet port on the back of IP phones deployed in the enterprise public areas. Consideration details for this topic can be found in the section entitled “802.1x Guest VLAN” in the *Network Virtualization—Access Control Design Guide.*

QoS

One of the main concerns when allowing a guest to connect to the enterprise network is the risk that the guest machine can be the carrier of a virus or worm that might disrupt the network itself. At the same time, even in scenarios where the guest clients are not carrying any viruses, a specific enterprise policy can limit the bandwidth made available to guests. Deploying web authentication appliances in-band allows you to configure bandwidth throttling for all the guest traffic enforced through these devices. Unmanaged partner accounts may be configured to get a higher bandwidth allotment than guest accounts. This document recommends deploying the web-auth appliances in a centralized location (usually the enterprise DMZ); some additional configuration is required to provide QoS between the edge of the network (where guest clients gain access to the network) and the DMZ.

Cisco suggests two ways of accomplishing this goal:

- Statically rate-limit guest traffic
- Mark guest traffic as scavenger (less than best effort)

The first approach is straightforward. Any guest traffic entering the network is rate-limited below a certain threshold. This limits the guest traffic to a reasonable amount.

The second approach requires QoS to be implemented across the entire network and involves three phases in treating the traffic:

- Mark the guest traffic as scavenger traffic at the network edge. This is basically re-marking the guest DSCP bits to an arbitrary known value that is considered *scavenger* traffic and treated as less than best effort.
- Meter the amount of traffic flowing through the network. This is done at the aggregation points and is basically intended to identify network congestion situations.
- Aggressively drop scavenger traffic at the network aggregation when congestion is detected.

A more complete description of this topic can be found in the *Network Virtualization—Path Isolation Design Guide*.

Appendix—Design Guide Mapping

Table 1 and Table 2 provide mapping to the design guides that are relative to the solutions described in this document. Each table refers to the corresponding sections of the design guides (access control, path isolation, and services edge) where more detailed information can be found.

Table 1 Guest and Unmanaged Partner Access Option A

Functional Area/Design Guide	Technique	Design Guide Section Title
Access control	802.1x guest VLAN	Access control— "802.1x Guest VLAN"
Path isolation—Campus	GRE + VRFs	Path isolation— "Network Virtualization Using GRE and VRF"
Path isolation—WAN	GRE + VRFs	Path isolation— "Network Virtualization Using GRE and VRF"
Services edge	Internet access design	Services edge— "Shared Internet Access—Internet Edge Design"
Services edge	Web authentication	Services edge— "Centralized Web Authentication Services"

Table 2 Guest and Unmanaged Partner Access Option B

Functional Area/Design Guide	Technique	Design Guide Section Title
Access control	802.1x Guest VLAN	Access control— "802.1x Guest VLAN"
Path isolation—Campus	Distributed ACLs	Path isolation—"Network Virtualization Using Distributed ACLs"
Path isolation—WAN	Distributed ACLs	Path isolation— "Network Virtualization Using Distributed ACLs"
Services edge	Internet access design	Services edge— "Shared Internet Access: Internet Edge Design"
Services edge	Web authentication	Services edge— "Centralized Web Authentication Services"