CHAPTER 3

Campus Design Considerations

When deploying an energy management solution over a campus network, a common design practice is to view the campus as a series of interconnected modules, each with particular requirements for supporting the solution. Figure 3-1 shows an example of a campus network design from a modular perspective.

Figure 3-1 Campus Network Design Modules

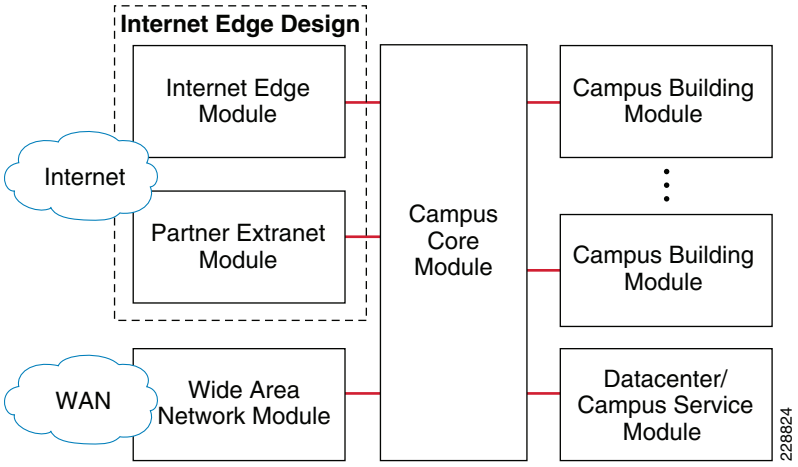

Table 3-1 provides a brief overview of the role of each module in a campus network.

Table 3-1 Campus Network Modules

Module	Description
Internet Edge	Provides centralized and secure Internet connectivity to and from the enterprise network.
Partner Extranet	Provides centralized and secure connectivity to partner networks via VPN, the Internet, or direct connections.
Wide Area Network	Provides internal connectivity between campus locations, and between campus and branch locations within the enterprise.
Campus Core	Provides a high-speed routed infrastructure between various modules within the campus network.

Table 3-1 Campus Network Modules (continued)

Campus Building	Provides both network connectivity for end-user devices (PCs, IP phones, etc.) and aggregation of those devices within each building of a campus network.
Data Center/Campus Service Module	Provides a high-speed infrastructure for the centralization of server resources within a campus network.

Depending upon the particular deployment model selected, not all of the modules are relevant for support of the energy management solution. For example, if the enterprise customer has decided to manage the energy management solution themselves, network connectivity to a partner through the Partner Extranet Module is not necessary. The reader should also note that enterprise customers may choose to collapse the functionality of several modules into a single module. Each of the modules presented in Table 3-1 are discussed in detail, either within this chapter or following chapters. This chapter focuses on the Campus Building Module, Campus Core Module, and the Wide-Area Network Module. Chapter 4, "Internet Edge Design Considerations"focuses on the Internet Edge Module and the Partner Extranet Module, which together comprise the Internet Edge Design. Chapter 5, "Data Center/Campus Service Module Design Considerations" discusses the Data Center/Campus Service Module.

Campus Building Module

In terms of the energy management solution, the function of the Campus Building Module is to provide network connectivity for the Cisco Network Building Mediators deployed within a campus environment. It also provides network connectivity and isolation for any Ethernet-based building devices which utilize protocols such as BACnet/IP and Modbus/TCP. The Campus Building Module provides strict access control to and from the management interface of the Mediators from within the enterprise network infrastructure. Finally, it provides QoS classification and marking of ingress traffic from the management interface of the Mediators, so that the data flows receive the necessary QoS service levels as they cross the enterprise network infrastructure.

A wider range of access control exists within the campus compared to the branch, due to the wider range of platforms typically deployed within the campus. Traditional Campus Building Module designs implement a hierarchical structure consisting of a distribution layer and an access layer. Typically, the distribution layer consists of a Layer-3 switch, while the access layer can be either a Layer 3 or Layer 2 switch. Campus Building Module designs with Layer-2 access switches are discussed in the Layer-2 Access Layer Switch Designs section; while Campus Building Module designs with Layer-3 access switches are discussed in the Layer-3 Access Layer Switch Designs section.

Layer-2 Access Layer Switch Designs

Figure 3-2 shows three examples of a Campus Building Module with support for the energy management solution using a Layer-2 access switch.

Figure 3-2 Layer-2 Access Switch Designs—Deployment Options

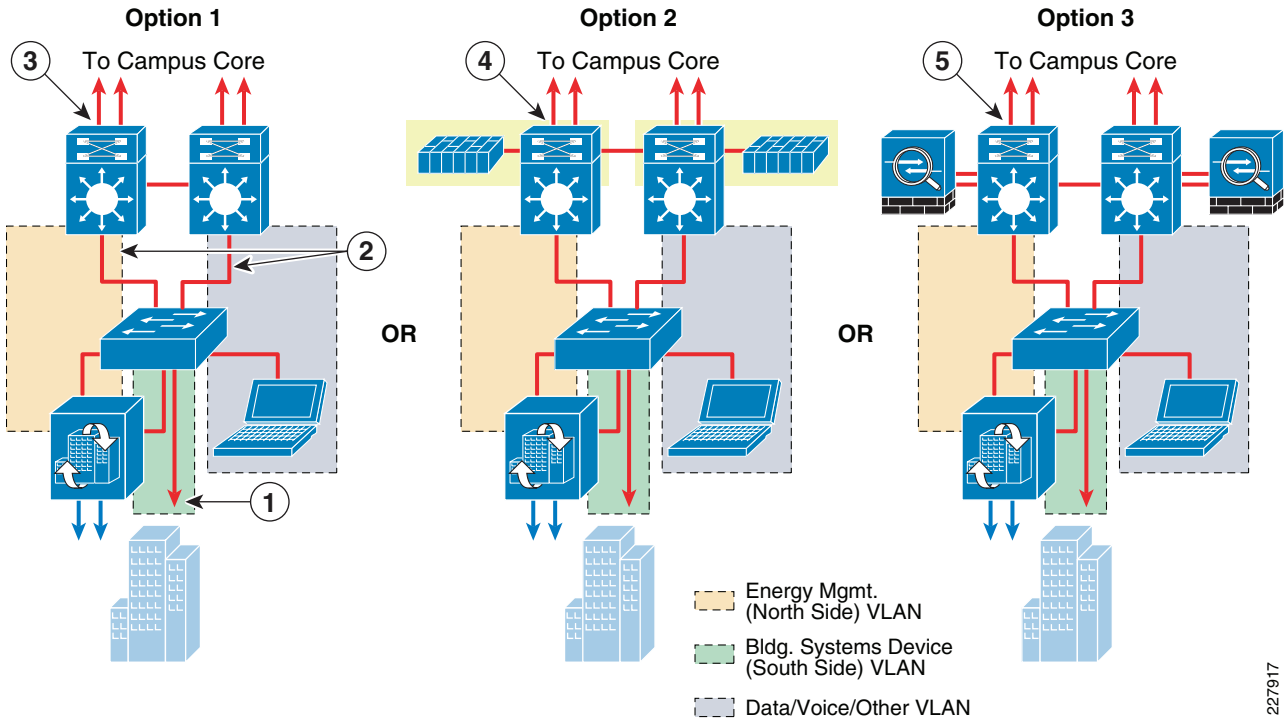

The following describes the numbers shown in Figure 3-2:

- **1**—Building systems device VLAN isolated by not trunking it to the distribution-layer switch.
- **2**—Both the energy management and the data/voice/other VLANs trunked to the distribution-layer switch.
- **3**—Layer-3 switches with ACLs at the distribution layer provide stateless access control to and from the energy management VLAN.
- **4**—Layer-3 switches with FWSM at the distribution layer provide stateful access control to and from the energy management VLAN.
- **5**—Layer-3 switches and ASA 5500 Security Appliances at the distribution layer provide stateful access control to and from the energy management VLAN.

In each of the three deployment options, a separate energy management VLAN (north side) and a building systems device VLAN (south side) is provisioned on the Layer-2 access switch. The energy management VLAN along with any data/voice/other VLANs are trunked to the Layer-3 distribution switch. However, the building systems device VLAN is not trunked, effectively isolating it within the access-layer switch. The only devices connected to the building systems VLAN are the actual building devices that use protocols such as BACnet/IP and Modbus/TCP, as well as the building systems device (south side) interface of the Mediator. With this design, all communications to the building devices occur through the Mediator.

In the first deployment option shown in Figure 3-2, access to and from the Mediator is controlled via ACLs applied to a switched virtual interface (SVI) defined for the energy management VLAN on the Layer-3 distribution switch. For redundancy purposes, a protocol such as Gateway Load Balancing Protocol (GLBP), Hot Standby Routing Protocol (HSRP), or Virtual Router Redundancy Protocol (VRRP) needs to be run between the SVI interfaces. Example 3-1 shows a partial configuration of a Catalyst 4500 switch with a basic HSRP group and access-control lists into and out of the SVI dedicated for the energy management solution.

Example 3-1 Partial Configuration of Layer 3 Distribution Switch with SVI and ACLs

```
interface Vlan192! Layer 3 SVI interface to the energy mgmt VLAN
description ENERGY MANAGEMENT (NORTH SIDE) VLAN
ip address 10.17.192.4 255.255.255.248
ip access-group 100 in ! Traffic from the energy mgmt VLAN to the network
ip access-group 101 out ! Traffic from the network to the energy mgmt VLAN
no ip redirects
standby 1 ip 10.17.192.1! HSRP group for redundancy
standby 1 priority 90
!
~
!
access-list 100 permit tcp host 10.17.192.2 host 10.16.19.2 eq 22
access-list 100 permit tcp host 10.17.192.2 host 10.16.19.2 eq ftp
access-list 100 permit tcp host 10.17.192.2 host 10.16.19.2 range 49152 49153
access-list 100 permit tcp host 10.17.192.2 host 10.16.19.2 eq www
access-list 100 permit tcp host 10.17.192.2 host 10.16.19.2 eq 8443
access-list 100 remark EXPORT DATA INITIATED FROM MEDIATOR TO MSP RA VPN HOST
!
access-list 100 permit tcp host 10.17.192.2 eq 22 host 10.16.19.2
access-list 100 permit tcp host 10.17.192.2 eq www host 10.16.19.2
access-list 100 permit tcp host 10.17.192.2 eq 81 host 10.16.19.2
access-list 100 permit tcp host 10.17.192.2 eq 443 host 10.16.19.2
access-list 100 remark RETURN MGMT TRAFFIC INITIATED FROM MSP RA VPN HOST TO MEDIATOR
!
access-list 100 permit tcp host 10.17.192.2 host 10.17.192.72 eq 22
access-list 100 permit tcp host 10.17.192.2 host 10.17.192.72 eq ftp
access-list 100 permit tcp host 10.17.192.2 host 10.17.192.72 range 49152 49153
access-list 100 permit tcp host 10.17.192.2 host 10.17.192.72 eq www
access-list 100 permit tcp host 10.17.192.2 host 10.17.192.72 eq 8443
access-list 100 remark EXPORT DATA INITIATED FROM MEDIATOR TO ENTERPRISE MGMT HOST
!
access-list 100 permit tcp host 10.17.192.2 eq 22 host 10.17.192.72
access-list 100 permit tcp host 10.17.192.2 eq www host 10.17.192.72
access-list 100 permit tcp host 10.17.192.2 eq 81 host 10.17.192.72
access-list 100 permit tcp host 10.17.192.2 eq 443 host 10.17.192.72
access-list 100 remark RETURN MGMT TRAFFIC INITIATED FROM ENTERPRISE MGMT HOST TO MEDIATOR
!
access-list 100 permit tcp host 10.17.192.2 host 10.192.2.3 eq 22
access-list 100 permit tcp host 10.17.192.2 host 10.192.2.3 eq ftp
access-list 100 permit tcp host 10.17.192.2 host 10.192.2.3 range 49152 49153
access-list 100 permit tcp host 10.17.192.2 host 10.192.2.3 eq www
access-list 100 permit tcp host 10.17.192.2 host 10.192.2.3 eq 8443
access-list 100 remark EXPORT DATA INITIATED FROM MEDIATOR TO CLOUD SERVICES INTERNET HOST
!
access-list 100 permit tcp host 10.17.192.2 host 10.16.4.10 eq 22
access-list 100 permit tcp host 10.17.192.2 host 10.16.4.10 eq ftp
access-list 100 permit tcp host 10.17.192.2 host 10.16.4.10 gt 1023
access-list 100 permit tcp host 10.17.192.2 host 10.16.4.10 eq www
access-list 100 permit tcp host 10.17.192.2 host 10.16.4.10 eq 8443
access-list 100 remark EXPORT DATA INITIATED FROM CAMPUS MEDIATOR TO CLOUD SERVICES HOST
VIA DMZ PROXY
!
access-list 100 permit tcp host 10.17.192.2 host 10.16.1.9 eq smtp
access-list 100 remark TRAFFIC FROM MEDIATOR TO EMAIL SERVER
!
access-list 100 permit udp host 10.17.192.2 eq ntp host 10.17.192.1 eq ntp
access-list 100 permit udp host 10.17.192.2 eq bootpc host 10.17.192.1 eq bootps
access-list 100 remark NTP and DHCP TRAFFIC FROM MEDIATOR TO LAYER 3 SWITCH
!
access-list 100 permit udp host 10.17.192.2 host 10.16.1.9 eq domain
access-list 100 remark TRAFFIC FROM MEDIATOR TO DNS SERVER
!
```

```
access-list 100 permit tcp host 10.17.192.2 eq 5150 host 10.17.192.70
access-list 100 remark RETURN RNA TRAFFIC INITIATED FROM DATA CENTER MEDIATOR
!
access-list 100 deny ip any any log
access-list 100 remark BLOCK and OPTIONALLY LOG ADDITIONAL ACCESS
!
!
access-list 101 permit tcp host 10.16.19.2 eq 22 host 10.17.192.2
access-list 101 permit tcp host 10.16.19.2 eq ftp host 10.17.192.2
access-list 101 permit tcp host 10.16.19.2 range 49152 49153 host 10.17.192.2
access-list 101 permit tcp host 10.16.19.2 eq www host 10.17.192.2
access-list 101 permit tcp host 10.16.19.2 eq 8443 host 10.17.192.2
access-list 101 remark RETURN SESSION DATA FROM MEDIATOR EXPORT TO MSP RA VPN HOST
!
access-list 101 permit tcp host 10.16.19.2 host 10.17.192.2 eq 22
access-list 101 permit tcp host 10.16.19.2 host 10.17.192.2 eq www
access-list 101 permit tcp host 10.16.19.2 host 10.17.192.2 eq 81
access-list 101 permit tcp host 10.16.19.2 host 10.17.192.2 eq 443
access-list 101 remark MGMT TRAFFIC INITIATED FROM MSP RA VPN HOST TO MEDIATOR
!
access-list 101 permit tcp host 10.17.192.72 eq 22 host 10.17.192.2
access-list 101 permit tcp host 10.17.192.72 eq ftp host 10.17.192.2
access-list 101 permit tcp host 10.17.192.72 range 49152 49153 host 10.17.192.2
access-list 101 permit tcp host 10.17.192.72 eq www host 10.17.192.2
access-list 101 permit tcp host 10.17.192.72 eq 8443 host 10.17.192.2
access-list 101 remark RETURN SESSION DATA FROM MEDIATOR EXPORT TO ENTERPRISE MGMT HOST
!
access-list 101 permit tcp host 10.17.192.72 host 10.17.192.2 eq 22
access-list 101 permit tcp host 10.17.192.72 host 10.17.192.2 eq www
access-list 101 permit tcp host 10.17.192.72 host 10.17.192.2 eq 81
access-list 101 permit tcp host 10.17.192.72 host 10.17.192.2 eq 443
access-list 101 remark MGMT TRAFFIC INITIATED FROM ENTERPRISE MGMT HOST TO MEDIATOR
!
access-list 101 permit tcp host 10.192.2.3 eq 22 host 10.17.192.2
access-list 101 permit tcp host 10.192.2.3 eq ftp host 10.17.192.2
access-list 101 permit tcp host 10.192.2.3 range 49152 49153 host 10.17.192.2
access-list 101 permit tcp host 10.192.2.3 eq www host 10.17.192.2
access-list 101 permit tcp host 10.192.2.3 eq 8443 host 10.17.192.2
access-list 101 remark RETURN SESSION DATA FROM MEDIATOR EXPORT TO CLOUD SERVICES INTERNET
HOST
!
access-list 101 permit tcp host 10.16.4.10 eq 22 host 10.17.192.2
access-list 101 permit tcp host 10.16.4.10 eq ftp host 10.17.192.2
access-list 101 permit tcp host 10.16.4.10 range 49152 49153 host 10.17.192.2
access-list 101 permit tcp host 10.16.4.10 eq www host 10.17.192.2
access-list 101 permit tcp host 10.16.4.10 eq 8443 host 10.17.192.2
access-list 101 remark RETURN SESSION DATA FROM MEDIATOR EXPORT TO CLOUD SERVICES HOST VIA
DMZ PROXY
!
access-list 101 permit tcp host 10.16.1.9 eq smtp host 10.17.192.2
access-list 101 remark RETURN TRAFFIC FROM EMAIL SERVER TO MEDIATOR
!
access-list 101 permit udp host 10.17.192.1 eq ntp host 10.17.192.2 eq ntp
access-list 101 permit udp host 10.17.192.1 eq bootps host 10.17.192.2 eq bootpc
access-list 101 remark RETURN NTP and DHCP TRAFFIC LAYER 3 SWITCH TO MEDIATOR
!
access-list 101 permit udp host 10.16.1.9 eq domain host 10.17.192.2
access-list 101 remark RETURN DNS SERVER TRAFFIC TO MEDIATOR
!
access-list 101 permit tcp host 10.17.192.70 host 10.17.192.2 eq 5150
access-list 101 remark RNA SESSION INITIATED FROM DATA CENTER MEDIATOR
!
access-list 101 deny ip any any log
access-list 101 remark BLOCK and OPTIONALLY LOG ADDITIONAL ACCESS
```

Example 3-1 shows both inbound (traffic from the Mediator VLAN into the SVI) and outbound (traffic from the rest of the network into the Mediator VLAN) ACLs on the SVI interface. The ACLs show connectivity allowed from the campus Mediator VLAN to the following:

- An MSP partner server via remote access VPN connection for both management and data export
- An enterprise management server for both management and data export
- A cloud services partner for data export
- A cloud services partner via a proxy server for data export
- An E-mail server for events and/or data export
- Another Mediator (located within the data center) for sharing datapoints via the RNA protocol
- The adjacent Catalyst switch for NTP and DHCP services
- A DNS server

Furthermore, all potential data export protocols are shown within the ACLs. This was done purely for illustrative purposes. In real deployments, only a subset of devices and protocols will in all likelihood need to be configured within the ACLs, making them far simpler. Finally, note that the FTP data port range has been added (and restricted to ports 49152 to 49153) since no application-layer inspection of the FTP control channel exists with simple ACLs.

For network administrators who desire or require stateful firewalling of the Mediator from the rest of the campus network, the second or third deployment options can be deployed. The second deployment option offers an integrated solution with a Firewall Services Module deployed within the Layer-3 distribution switch. In the third deployment option, a separate ASA 5500 Security Appliance can be deployed along with the Layer-3 distribution switch. Note that in either stateful firewall deployment option, only selective VLANs, such as the energy management VLAN, need to be trunked through to the firewall which provides the Layer 3 interface. Chapter 5, "Data Center/Campus Service Module Design Considerations"provides additional detail around the use of the FWSM and/or ASA 5500 Security Appliance with a Layer 3 switch.

The selection of the deployment option often depends on the hardware already implemented within the Campus Building Module. For each of the three options, since only basic Layer-2 functionality and VLAN trunking is required at the access layer, Catalyst 2900 Series switches can be implemented. Sizing of the Mediator itself depends largely on the number of points that will be monitored within the campus building. Option 1 provides the greatest flexibility in terms of the platform support at the distribution layer. The Catalyst 6500 Series, Catalyst 4500 Series, and even the Catalyst 3750 Series switch stack all support ACLs deployed across an SVI. Option 2 provides the least flexibility in terms of the platform support. At the distribution layer, only the Catalyst 6500 Series can support the Firewall Service Module (FWSM). Option 3 provides the same flexibility in terms of the platform support at the distribution layer as Option 1. The Catalyst 6500 Series, Catalyst 4500 Series, and even the Catalyst 3750 Series switch stack can be deployed at the distribution layer. However, a separate set of ASA 5500 firewalls is deployed in order to provide stateful isolation of the energy management VLAN from the rest of the data/voice/other VLANs within the campus building.

Layer-3 Access Layer Switch Designs

Figure 3-3 shows three examples of a Campus Building Module with support for the energy management solution using a Layer-3 access switch.

Figure 3-3 Layer-3 Access Switch Designs—Deployment Options

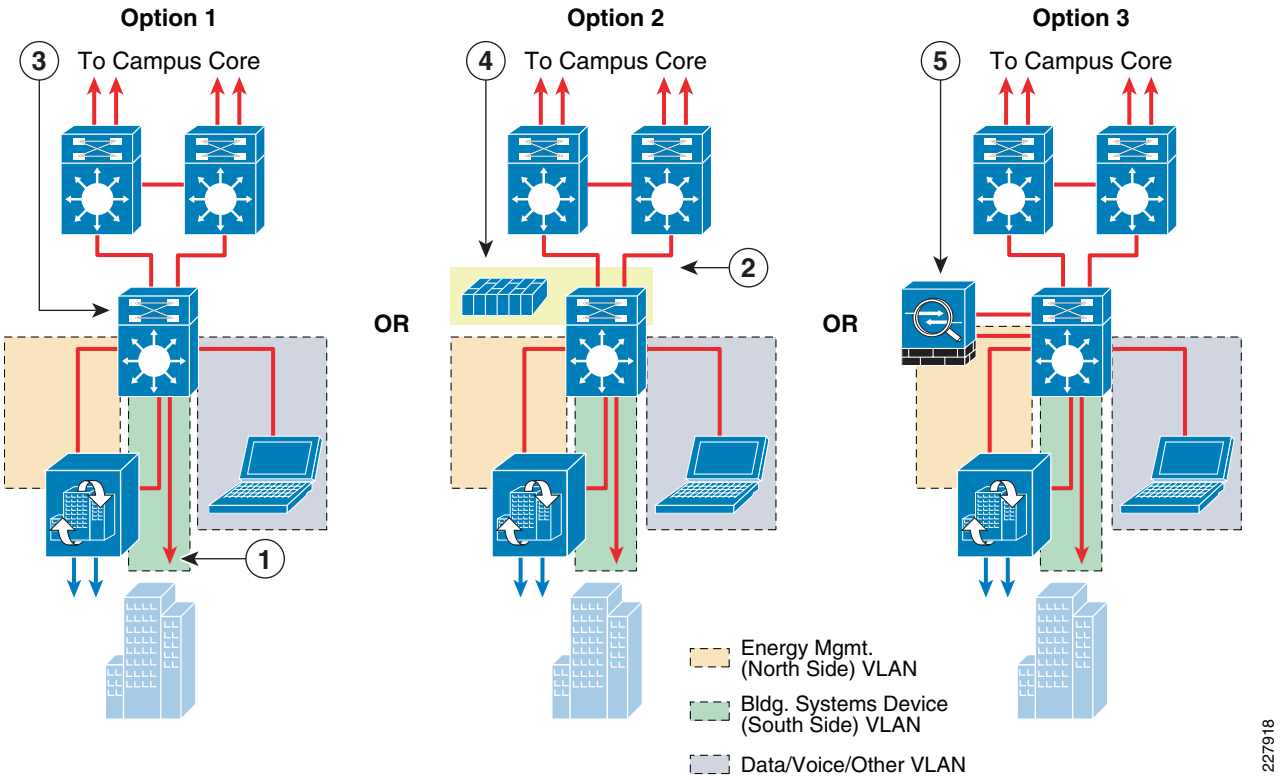

The following describes the numbers shown in Figure 3-3:

- **1**—Building systems device VLAN isolated by not configuring an SVI for the VLAN at the Layer-3 access switch.
- **2**—Routed uplinks between the access and distribution switches.
- **3**—Layer-3 access switch provides stateless access control to and from the energy management VLAN via ACLs.
- **4**—Layer-3 access switch with the FWSM provides stateful access control to and from the energy management VLAN.
- **5**—Layer-3 access switch and ASA 5500 Security Appliance provides stateful access control to and from the energy management VLAN.

When Layer-3 switches are deployed within the access layer, the access control point for traffic between VLANs is typically shifted down to the access layer. The same three deployment choices exist, but at the access layer. In each of the three deployment options, a separate energy management VLAN (north side) and a building systems device VLAN (south side) are provisioned on the Layer-3 access switch. Layer 3 interfaces (SVIs or the firewall itself) are defined for the energy management VLAN along with any data/voice/other VLANs. However, a Layer 3 interface is not defined for the building systems device VLAN, effectively isolating it within the access layer switch. Again, the only devices connected to the

building systems VLAN are the actual building devices which utilize protocols such as BACnet/IP and Modbus/TCP, as well as the building systems device (south side) interface of the Mediator. All communications to the building devices occur through the Mediator.

In the first deployment option shown in Figure 3-3, access to and from the Mediator is controlled via ACLs applied to the SVI defined for the energy management VLAN on the Layer-3 access switch. For network administrators who desire or require stateful access control to the energy management VLAN, the second or third deployment options can be deployed. The second deployment options offers an integrated solution, with a Firewall Service Module deployed within the Layer-3 access switch. In the third deployment option, a separate ASA 5500 Security Appliance can be deployed along with the Layer-3 access switch. Again, the reader should note that in either stateful firewall deployment option, only selective VLANS such as the energy management VLAN need to be trunked through to the firewall which provides the Layer-3 interface.

The selection of the deployment option often depends on the hardware already implemented within the Campus Building Module. For each of the three options, since only basic Layer-3 routing functionality is required at the distribution layer, Catalyst 6500 Series, Catalyst 4500 Series, or even Catalyst 3750 Series switch stacks can be implemented. Option 1 provides the greatest flexibility in terms of the platform support at the access layer. The Catalyst 6500 Series, Catalyst 4500 Series, Catalyst 3750 Series switch stack, and even the Catalyst 3560 Series all support ACLs deployed across an SVI. Option 2 provides the least flexibility in terms of the platform support. At the access layer, only the Catalyst 6500 Series can support the Firewall Service Module (FWSM). Option 3 provides the same flexibility in terms of the platform support at the access layer as Option 1. The Catalyst 6500 Series, Catalyst 4500 Series, Catalyst 3750 Series, and even Catalyst 3560 Series switches can be deployed at the access layer. However, a separate ASA 5500 firewall is deployed in order to provide stateful isolation of the energy management VLAN from the rest of the data/voice/other VLANs within the campus building.

Extending VRFs to the Campus Building Module

The deployment of a network virtualization for energy management systems can provide the additional advantage of path isolation of the energy management solution traffic across the campus network infrastructure. When applied to the Campus Module, the energy management VRF is extended into the Layer-3 device. An example of this when implementing an Layer 2 access switch design is shown in Figure 3-4.

Figure 3-4 Layer-2 Access Switch Campus Module Design with Energy Management VRF

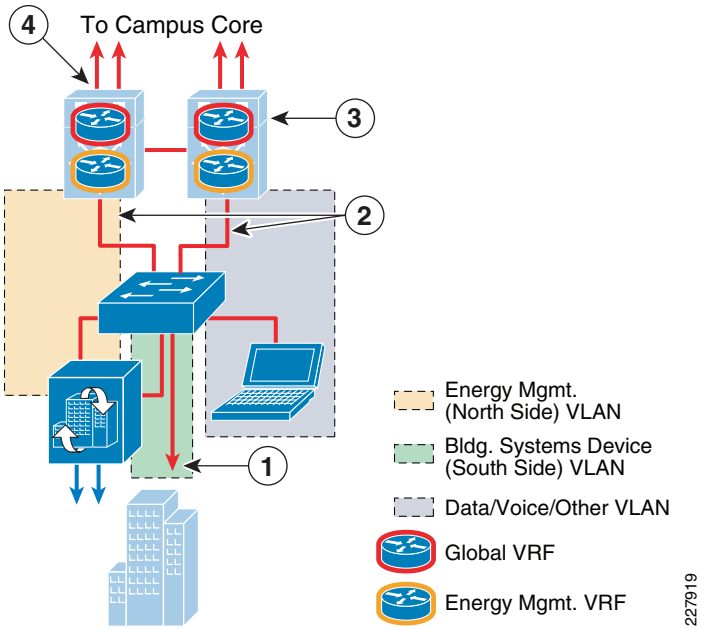

The following describes the numbers shown in Figure 3-4:

- **1**—Building systems device VLAN isolated by not trunking it to the Layer-3 distribution switch.
- **2**—Both the energy management and the data / voice / other VLANs trunked to the Layer-3 distribution switch.
- **3**—Energy management VLAN mapped to the energy management VRF, while data / voice / other VLANs mapped to the global VRF at the Layer-3 distribution switch.
- **4**—VRFs extended to the rest of the campus either via VRF-Lite end-to-end or VRF-Lite with GRE Tunnels from the Layer-3 distribution switch.

In this example, the energy management VLAN is defined on the Layer-2 access switch and trunked to the Layer-3 distribution switch where the SVI for the energy management VLAN is defined. The SVI is then mapped to an energy management VRF which is separate from the global VRF which supports the data/voice/other VLANs. Because the traffic within the energy management VRF is isolated from traffic in other VRFs, stateful firewalling is not really required within the Campus Building Module itself. This eases the administrative burden of configuring access control, and is a major advantage of deploying a VRF. However, inbound and outbound ACLs may still be applied to the SVI defined for the energy management VLAN in order to restrict access to the Mediators if desired. Example 3-2 shows a partial configuration of a Catalyst 4500 switch, this time with the SVI assigned to a Building Infrastructure Network (BIN) VRF dedicated for the energy management solution. GRE tunnels then extend the BIN VRF across the campus to the Data Center/Campus Service Module.

Example 3-2 Partial Configuration of Layer 3 Distribution Switch with VRFs using GRE Tunnels

```
!
ip vrf bin ! Defines the Building Infrastructure Network (BIN) VRF
rd 192:6
!
~
!
interface Tunnel0        ! GRE tunnel back to the Data Center Service Switch
description VRF FOR MEDIATOR NETWORK TO ME-W-DCSERV-1
```

```
ip vrf forwarding bin    ! Places the GRE tunnel within the BIN VRF
ip address 10.17.192.50 255.255.255.248
tunnel source Loopback0
tunnel destination 10.17.252.1
!
~
!
interface Loopback0
description LOOPBACK INTERFACE FOR TUNNEL TO ME-W-DCSERV-1
ip address 10.17.255.51 255.255.255.255
ip pim sparse-mode
!
~
!
interface Vlan192        ! Layer 3 SVI interface to the energy mgmt VLAN
description ENERGY MANAGEMENT (NORTH SIDE) VLAN
ip address 10.17.192.4 255.255.255.248
ip vrf forwarding bin    ! Places SVI interface within the BIN VRF
no ip redirects
standby 1 ip 10.17.192.1 ! HSRP group for redundancy
standby 1 priority 90
!
```

When Layer-3 access switches are used, both SVI for the energy management VLAN and the energy management VRF are configured on the access switch itself.

From the Campus Building Module, the energy management VRF can be extended across the campus via one of two methods. The first method (referred to as the VRF-Lite with GRE model) is to use GRE tunnels, as is shown in Example 3-2. GRE tunnels can be defined from the Campus Building Module Layer-3 switches to Layer-3 switches that support the Energy Management Operations Center (EMOC) within either a Data Center Service Module or Campus Service Module. The GRE tunnels are then mapped to the energy management VRF. Sets of GRE tunnels may also need to be defined from each branch location which supports a Mediator to the WAN Module. Another set of GRE tunnels can then be defined from the WAN Module to the Data Center Service Module. A similar set of GRE tunnels may also need to be defined from the Partner Extranet Module Layer-3 switches to the Layer-3 distribution switches within the Data Center Service Module or Campus Service Module. These tunnels support both MSP partner VPN management as well as the periodic export of logged data to the Internet via the Partner Extranet Module.

Note that with this design, routing of partner traffic goes through the Data Center Service Module or Campus Service Module before reaching the individual Mediators. This reduces the overall number of GRE tunnels required to support the energy management solution, versus defining two tunnels at each campus Mediator site—one to the Data Center Service Module and one to the Extranet Service Module. This design also facilitates the deployment of a hierarchical Mediator which is further discussed within the Chapter 5, “Data Center/Campus Service Module Design Considerations.” An visual example of the VRF-Lite with GRE deployment model is shown in Figure 3-5.

Figure 3-5 Energy Management Solution Using VRF-Lite with GRE Tunnels

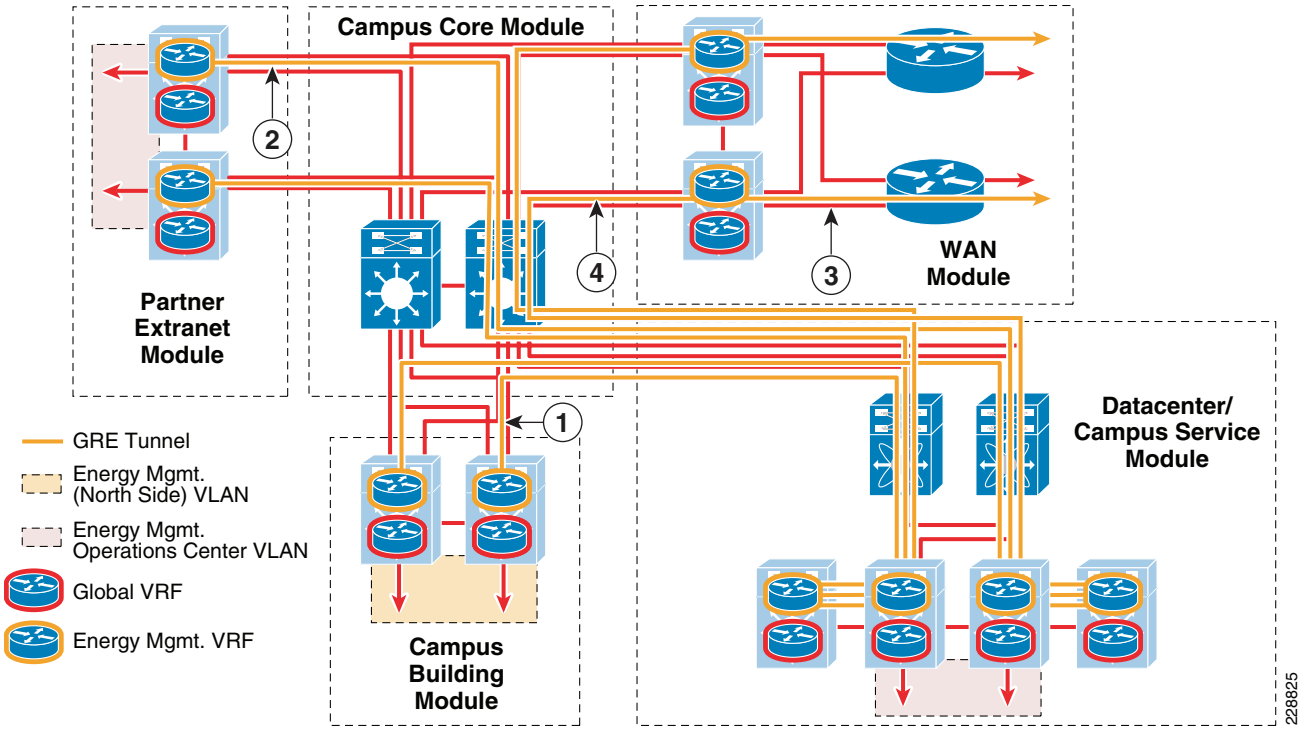

The following describes the numbers shown in Figure 3-5:

- **1**—GRE tunnels connect energy management VRF within the Campus Building module to the Datacenter / Campus Service Module.
- **2**—GRE tunnels connect energy management VRF within the Partner Extranet Module to the Data Center/Campus Service Module.
- **3**—Multiple GRE tunnels connect energy management VRFs within branch locations to the WAN Module.
- **4**—GRE Tunnels connect the energy management VRF within the WAN Module to the Data Center/Campus Service Module. Routing between GRE tunnels may occur both within the Data Center/Campus Service Module and WAN Module.

Note that in this model, the Campus Core Module switches do not need to support VRFs. The advantage of the VRF with GRE design is that only the edges of the network which participate within the energy management solution need to support VRFs. The downside to this design is that it does not allow any-to-any communications without having to backhaul all the traffic to a central point such as an Energy Management Operations Center (EMOC) within the Data Center Services Module.

The second method (referred to as VRF-Lite end-to-end) requires enabling VRFs on every Layer-3 device that supports the energy management solution. In this scenario, the energy management VRF is defined on the Campus Building Module Layer-3 switches, the Campus Core Layer-3 switches, the Layer-3 switches which support the EMOC within either a Data Center Service Module or Campus Service Module, the Layer-3 distribution switches of the Partner Extranet Module, and the Layer-3 distribution switches of the WAN Module. An example of the VRF-Lite end-to-end model is shown in Figure 3-6.

Figure 3-6 Energy Management Solution Utilizing VRF-Lite End-to-End

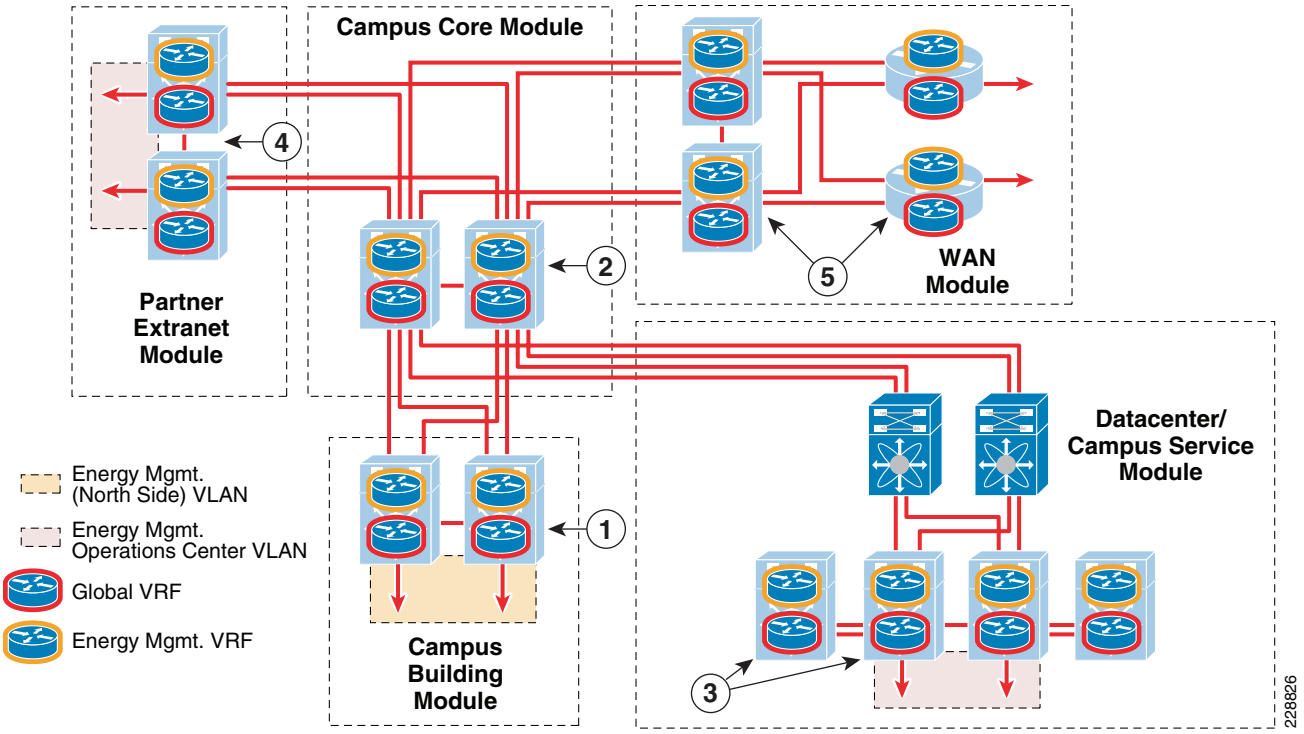

The following describes the numbers shown in Figure 3-6:

- **1**—Energy management VLAN mapped to energy management VRF at Campus Building Module Layer-3 distribution switches.
- **2**—Energy management and global VRFs extended across the Campus Core Module switches.
- **3**—Energy management and global VRFs extended to the Datacenter / Campus Service Module. Energy Management Operations Center (EMOC) VLAN mapped to energy management VRF within the Datacenter / Campus Service Module.
- **4**—Energy management and global VRFs extended to the Partner Extranet Module Layer-3 distribution switches. Energy management VLAN mapped to DMZ interface of the Partner Extranet firewall.
- **5**—Energy management and global VRFs extended to the WAN Module Layer-3 switches and/or routers for Mediators deployed within branch locations.

The advantage of this design is that it does allow any-to-any communications without having to backhaul all the traffic to a central point such as the Energy Management Operations Center, if peer-to-peer Mediator communications is needed. The downside, however, is that every Layer-3 device within the campus must support VRFs in order to implement the VRF-Lite end-to-end method.

Note The use of multipoint GRE tunnels has not been evaluated within this revision of the design guide of the energy management solution. Multipoint GRE tunnels may provide additional scalability of the VRF-Lite with GRE tunnel model. Also, the use of MPLS deployed within a campus to provide path isolation for the energy management solution has not been evaluated for this revision of the design guide. Future revisions may include discussion of such technologies.

Enterprise Client PC Access to Campus Building Module Mediators

In some scenarios, business requirements include the need for client PCs sitting on the enterprise data network to access energy usage data. Access to energy usage data can be accomplished in the following ways:

1. Client PCs access an energy scorecard website provided by a cloud services partner via the Internet.
2. Client PCs access an energy scorecard website provided internally by a partner or developed internally.
3. Client PCs directly access one or more websites deployed on the hierarchical Mediator which provides energy usage information.
4. Client PCs directly access websites deployed on Mediators deployed in branch and campus locations throughout the network infrastructure.

Option 1is discussed in the Chapter 4, "Internet Edge Design Considerations." Options 2 and 3 are discussed in the Chapter 5, "Data Center/Campus Service Module Design Considerations." This section discusses direct connectivity to Mediators within the Campus Building Module.

Direct client PC access to the Mediators poses a greater security risk than accessing an energy scorecard, since the Mediator may potentially be running logic which controls building systems. In business scenarios where access from enterprise client PCs to historical energy usage data is the only requirement, the network administrator should consider either contracting a cloud service partner to provide an energy scorecard service externally, or deploying an energy scorecard service internally (either via partner or internally developed). However, in business scenarios where access to real-time energy usage data, or access in order to change setpoints is a requirement; then direct access to the Mediators may be necessary. In such cases, the deployment of one or more hierarchical Mediator(s) within a data center Energy Management Operations Center (EMOC) LAN segment is a more secure access method, than individual access to each Mediator. When deploying a VRF, access to the overall energy management solution can be restricted to one or more strategic points within the IP infrastructure, such as the data center EMOC. Access to the hierarchical Mediator(s) can be controlled via ACLs or further tightened through the use of technologies such as an IPsec VPN deployed at a data center ASA 5500 Security Appliance. This is one major advantage of deploying path virtualization for the energy management solution. Even without a hierarchical Mediator design client access into the energy management VRF through an IPSec VPN internally deployed within the enterprise organization can be used to centrally control access and provide an audit trail for the network administrator.

When a VRF has not been implemented, access control from enterprise client PCs can be handled via ACLs or stateful firewalling within the Layer-3 distribution switch or Layer-3 access switch—depending upon the design implemented—as discussed above. In cases where an ASA 5500 Security Appliance has been deployed within the Campus Building Module, the network administrator has the option of further tightening access control to each Mediator via the deployment of an IPSec VPN on the ASA 5500 Security Appliance. However, when implementing the FWSM or basic ACLs within the Campus Building Module, access control may not be as tightly controlled. Where possible, the network administrator should try to restrict the individual PCs that have access to the Mediators via ACLs, and use the HTTPS protocol and passwords for access to the Mediators themselves.

Note Technologies that provide dynamic ACLs based on user authentication, such as Cisco's Lock-and-Key technology, have also not been evaluated as part of this design guide. Such technologies place a single dynamic ACL entry across an interface, and hence cannot provide fine-grained per-user access, as can remote-access IPSec VPN technologies. Future revisions of this design guide may investigate technologies such as Cisco's Lock-and-Key.

QoS within the Campus Building Module

A secondary function of the Campus Building Module is to provide classification and marking of Cisco Network Building Mediator traffic flows as they enter the network. Currently, the Mediator marks all traffic in the Best Effort service class (DSCP value = 0). The Mediator does not currently support VLANs either, so Layer-3 Class-of-Service (CoS) marking is not supported. In order to classify traffic flows from the Mediator in anything other than the Best Effort service class, the classification and re-marking must be performed at the ingress port of Campus Building Module access switch. The following tow different methods are discussed in this guide:

- Identifying and marking individual traffic flows from the Mediator to different service classes based upon the traffic type (FTP, HTTP, SSH, etc.) and use (periodic data export or configuration and management).
- Identifying and marking all traffic flows from the Mediator to a single service class.

Cisco recommends the deployment of a 12-class QoS model based on IETF RFC 4594 for the support of voice, video, and data across a converged IP network infrastructure, as shown in Figure 3-7.

Figure 3-7 RFC 4594-Based Enterprise 12-Class QoS Model

Application Class	PHB	Admission Control	Queueing and Dropping	Application Examples
VoIP Telephony	EF	Required	Priority Queue (PQ)	Cisco IP Phones
Broadcast Video	CS5	Required	Optional (PQ)	Cisco IP Surveillance, Cisco Enterprise TV
Realtime Interactive	CS4	Required	Optional (PQ)	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Unified Personal Communicator
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoD)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Call-Signaling	CS3		BW Queue	SCCP, SIP, H.323
OAM	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	Cisco WebEx, Cisco MeetingPlace, ERP Apps
Bulk Data	AF1		BW Queue + DSCP WRED	Email, FTP, Backup Apps, Content Distribution
Best Effort	default		Default Queue + RED	Default Class Traffic
Scavenger	CS1		Min BW Queue (Deferential)	YouTube, iTunes, BitTorrent, Xbox Live

227922

The 12-class QoS model can then be mapped to the various queueing structures of the Catalyst switch and Cisco router platforms, depending upon which platforms are deployed throughout the network infrastructure. Note that the applications listed for each of the traffic classes in Figure 3-7 are simply

suggestions. There is no strict requirement for specific traffic types to be placed into a particular traffic classes. Figure 3-8 shows a possible method of applying the 12-class QoS model specifically to the energy management solution based upon the traffic flows to and from the Cisco Network Building Mediators.

Figure 3-8 Possible Mapping of Energy Management Traffic to 12-Class QoS Model

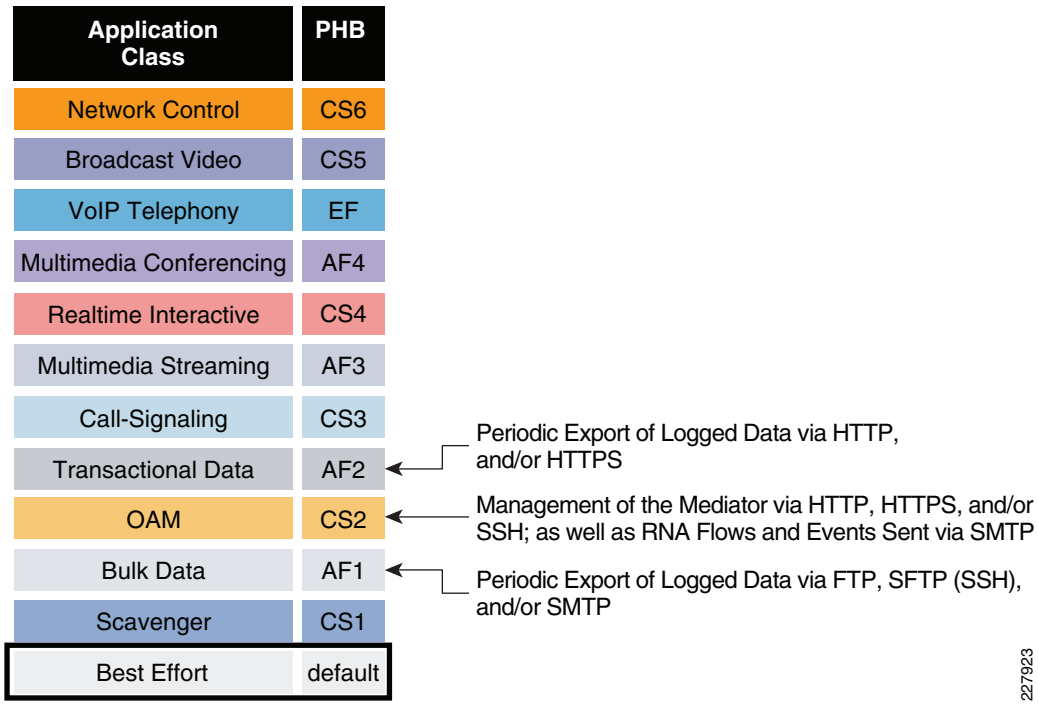

In this example, the network administrator may consider placing the periodic export of logged data into either the Bulk Data service class and marked with a DSCP value of AF1 or the Transactional Data service class and marked with a DSCP value of AF2. For example, if the logged data is exported very infrequently - perhaps ever hour-utilizing a protocol such as FTP or SFTP, then the characteristics of the traffic are typically medium to large file transfers which occur infrequently. Therefore the Bulk Data service class may be appropriate. If however, the logged data is exported very often—perhaps every few minutes—using a protocol such as an HTTP or HTTPS POST, then the characteristics of the traffic are typically small transfers which occur frequently. Therefore, the Transactional Data service class may be appropriate.

Note that in either case, the periodic export of logged data utilizes a TCP-based protocol which handles lost packets and retransmissions. Further, with the energy management solution there are no stringent time constraints in terms of end-to-end latency and/or jitter for the traffic flows, which is characteristic of voice and/or video traffic types. In terms of the actual management of the Mediators for configuration, real-time monitoring, event forwarding via E-mail (SMTP), and peer-to-peer communications between mediators via the RNA protocol, the enterprise network administrator may consider the Operations, Administration, and Maintenance (OAM) service class. This service class is often utilized for configuration and monitoring of network infrastructure devices such as routers and switches.

Classification and marking of traffic from the Mediators can be accomplished via an ingress policy-map with ACLs, applied to the access switch ports to which the Mediators are connected within the Campus Building Module. The ACLs can be configured simply to identify a particular protocol based on its TCP port number. The policy-map marks all traffic corresponding to that protocol to a particular service class. Example 3-3 shows a partial configuration from a Catalyst 2960 switch using this method.

Example 3-3 Classification and Marking via ACLs Based on Protocol

```
!
class-map match-all MGMT_TRAFFIC
match access-group name MEDIATOR_MGMT
class-map match-all DATA_EXPORT_TRAFFIC
match access-group name MEDIATOR_EXPORT
!
!
policy-map MEDIATOR_ENDPOINT
class DATA_EXPORT_TRAFFIC
set ip dscp af11
class MGMT_TRAFFIC
set ip dscp cs2
class class-default
set ip dscp default
!
~
!
interface FastEthernet0/18
description CONNECTION TO CAMPUS MEDIATOR
switchport access vlan 192
srr-queue bandwidth share 1 30 35 5
priority-queue out
service-policy input MEDIATOR_ENDPOINT
!
~
!
ip access-list extended MEDIATOR_EXPORT
permit tcp any any eq ftp
permit tcp any any range 49152 49153
ip access-list extended MEDIATOR_MGMT
permit tcp any any eq smtp
permit tcp any any eq www
permit tcp any eq www any
permit tcp any eq 81 any
permit tcp any any eq 22
permit tcp any eq 22 any
permit tcp any any eq 443
permit tcp any eq 443 any
permit udp any any eq domain
permit udp any any eq ntp
permit udp any any eq bootps
!
```

In the above example, FTP is used for the periodic export of logged data from the Mediator to a cloud services partner host. Both the control channel (TCP port 23) and a restricted data channel (TCP port range 49152 - 49153) have been set to match the class map named DATA_EXPORT_TRAFFIC. Management protocols which include SMTP (TCP port 25), HTTP (TCP port 80), TCP port 81, HTTPS (TCP port 443), SSH (TCP port 22), DNS (UDP port 53), NTP (UDP port 123), and DHCP (UDP port 68) have been set to match the class map named MGMT_TRAFFIC. Both class maps are placed into a policy map named MEDIATOR_ENDPOINT which is then applied to ingress traffic on the switch port connected to the Mediator. The policy map marks all DATA_EXPORT_TRAFFIC to AF11, and all MGMT_TRAFFIC to CS2. All other traffic is remarked to a default value.

Note that the same protocol can often be used for the periodic export of logged data as well as the configuration and real-time monitoring of the Mediators. For example, HTTPS can be used to configure the Mediators via the configTOOL application. Periodic logged data can also be exported to a cloud

services server located on the Internet via an HTTPS POST. In such cases, the only way of differentiating whether the HTTPS traffic should be classified and marked into the Transactional Data service class or the OAM service class may be the destination address to which the traffic is being sent.

Because of the complexities involved with this approach, an alternative is to simply mark all traffic inbound on the port connected to the Mediator to a particular service class. In this scenario, the access switch port can be configured to mark all ingress traffic to a single service class, such as OAM. Example 3-4 shows a partial configuration from a Catalyst 2960 switch using this method.

Example 3-4 Classification and Marking All Traffic to a Single Service Class

```
!
interface FastEthernet0/18
description CONNECTION TO CAMPUS MEDIATOR
switchport access vlan 192
mls qos cos 2
!
```

In the above example, the switch is configured to mark all ingress traffic to a class-of-service (CoS) value of 2. Based on the CoS to DSCP mapping internally within the switch, this corresponds to a DSCP value of CS2, corresponding to the OAM service class shown in Figure 3-7 above. The CoS to DSCP mapping can be viewed using the **show mls qos maps cos-dscp** command on the Catalyst 2960 switch. The output appears similar to Example 3-5.

Example 3-5 COS to DSCP Mapping on a Catalyst 2960 Switch

```
me-westrich-3#show mls qos maps cos-dscp
Cos-dscp map:
        cos: 0 1  2  3   4   5  6   7
             --------------------------------
      dscp: 0 8 16 24 32 46 48 56
```

Alternatively, all traffic inbound on the energy management VLAN itself can be remarked to the OAM service class.

The selection of which method to implement—either different service classes for different traffic types, or a single service class for all traffic types—is a matter of preference by the end customer. In either scenario, the objective is to provide a service class for the energy management traffic that is consistent with its network requirements. There is no particular need for the energy management traffic to be placed into a service class designed for real-time interactive or multimedia traffic which has tight requirements for packet loss, jitter, and end-to-end delay. However, at the same time, the network administrator may desire a service class above scavenger or best effort traffic. The network administrator should also note that as traffic from the Mediators destined for a cloud services server located on the Internet exits the enterprise network, it is likely to be remarked into the default service class as it enters the ISP network. Therefore, any marking done at the ingress edge of the Campus Building Module applies to traffic as it traverses the enterprise network only. When traffic flows cross over a VPN tunnel to a MSP network, the network administrator should work closely with their partner to ensure the desired CoS is maintained on the MSP network.

WAN Module

In terms of the energy management solution, the function of the WAN Module is to provide a redundant network infrastructure between the branch and the campus locations over which energy management traffic flows from both the enterprise EMOC and MSP partner connections (when utilizing a centralized VPN deployment model). A typical campus WAN Module consists of a set of Layer-3 Catalyst 6500 switches functioning as a distribution layer connected to one or more sets of Cisco ASR 1000 Series, Cisco 7600 Series, or Cisco 7200 series routers which terminate that actual WAN circuits. An example is shown in Figure 3-9.

Figure 3-9 Example WAN Module with VRF Design

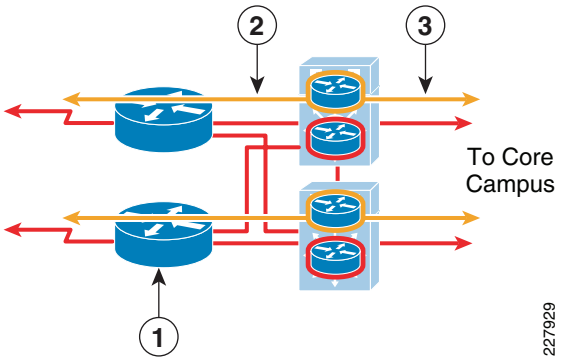

The following describes the numbers shown in Figure 3-9:

- **1**—Redundant ASR-1000 Series, Cisco 7600 Series, or Cisco 7200 Series routers provide WAN Circuit Termination from branch locations.
- **2**—GRE tunnels from branch locations terminate on Layer-3 Catalyst 6500 distribution switches.
- **3**—GRE tunnels defined from Layer-3 Catalyst 6500 distribution switches to the Energy Management Operations Center (EMOC) Layer-3 switches, or VRF-Lite End-to-End deployed across the campus.

In a non-VRF implementation, nothing specific needs to be configured on the Catalyst 6500 distribution switches in order to support the energy management solution. However, in a VRF implementation, modifications may be necessary. If the network administrator chooses to deploy the VRF-Lite with GRE model then GRE tunnels from the branch locations may be terminated on the Catalyst 6500 switches. This is because the Catalyst 6500 with Sup-720 Supervisor supports GRE in hardware. This provides a scalable platform for deploying multiple branches with Mediators. Example 6 shows a partial configuration from a Catalyst 6500 serving as the distribution layer switch of a WAN Module.

Example 3-6 Partial Configuration of a Catalyst 6500 Distribution Switch in the WAN Module

```
!
ip vrf bin      ! Defines the Building Infrastructure Network (BIN) VRF
rd 192:75
!
~
!
interface Tunnel0      ! GRE tunnel to the Branch Router
description VRF FOR MEDIATOR NETWORK TO ME-WESTRICH-1
ip vrf forwarding bin ! Places the GRE tunnel within the BIN VRF
ip address 10.17.192.26 255.255.255.248
tunnel source TenGigabitEthernet5/1
tunnel destination 10.17.252.9
```

```
!
interface Tunnel1     ! GRE tunnel back to the Data Center Service Switch
description VRF FOR MEDIATOR NETWORK TO ME-W-DCSERV-1
ip vrf forwarding bin
ip address 10.17.192.58 255.255.255.248
tunnel source Loopback2
tunnel destination 10.17.255.108
!
~
!
interface Loopback2
description LOOPBACK INTERFACE FOR TUNNEL FROM ME-W-DCSERV-1
ip address 10.17.252.3 255.255.255.255
!
~
!
interface TenGigabitEthernet5/1
description CONNECTION TO ME-WESTCORE-1 GIG5/1
ip address 10.17.100.10 255.255.255.252
ip pim sparse-mode
!
~
!
router eigrp 111
network 10.0.0.0
no auto-summary
!
address-family ipv4 vrf bin    ! Creates a routing process within the BIN VRF
network 10.17.192.0 0.0.0.255
network 10.17.252.0 0.0.0.255
network 10.17.255.0 0.0.0.255
no auto-summary
autonomous-system 99
exit-address-family
!
```

Note that other methods of supporting VRFs across the WAN, such as mapping them to an MPLS service also exist. Future revisions of this design guide may include further discussion of such technologies.

Campus Core Module

In terms of the energy management solution, the function of the Campus Core module is to provide a redundant high-speed Layer-3 infrastructure over which the energy management traffic flows as it crosses between the various campus modules. Typically two or more Catalyst 6500 switches make up the core switches of medium to large enterprise organizations, as shown in Figure 3-10.

Figure 3-10 Example Campus Core Module

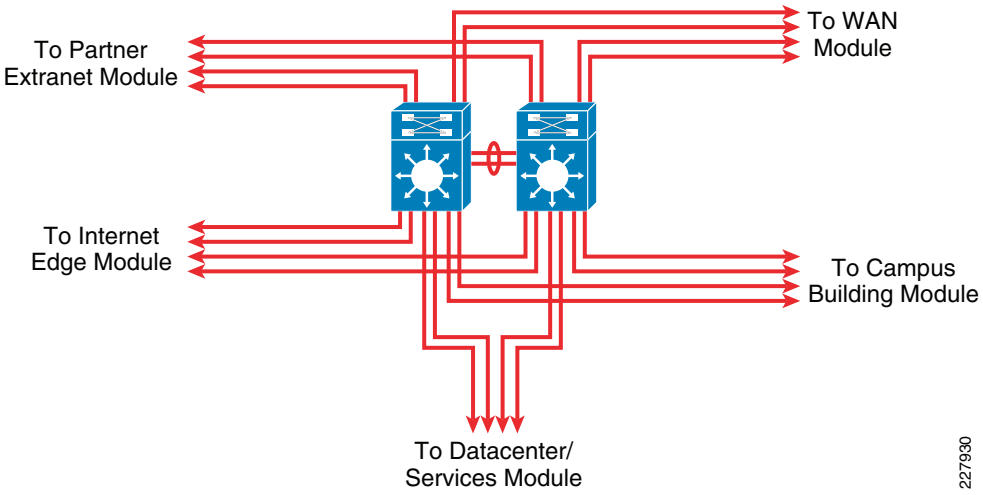

In a non-VRF implementation, nothing specific needs to be configured on the core Catalyst 6500 switches in order to support the energy management solution. However, in a VRF implementation, modifications may be necessary. If the network administrator chooses to deploy the VRF with GRE Tunnel model, then no modifications are needed to the core Catalyst 6500 switches. GRE tunnels are simply routed across the Layer-3 core switches. If the network administrator chooses to deploy the VRF-Lite end-to-end model, then the core Catalyst 6500 switches must be configured to support VRFs as well.