Mediator Design Guide

Last Updated: June 28, 2010

Solution Authors

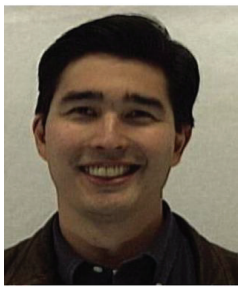

Roland Saville

Roland Saville, Technical Leader, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Roland is a Technical Leader for the Enterprise Systems Engineering team within Cisco, focused on developing best-practice design guides for enterprise network deployments. He has 14+ years of Cisco experience as a Systems Engineer, Consulting Systems Engineer, Technical Marketing Engineer, and Technical Leader. During that time, he has focused on a wide range of technology areas including the integration of voice and video onto network infrastructures, network security, and wireless LAN networking. Roland has a BS degree in Electrical Engineering from the University of Idaho and an MBA from Santa Clara University. He has co-authored the Cisco TelePresence Fundamentals book has six U.S. Patents

John Johnston

John Johnston, Technical Marketing Engineer, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

John has been with Cisco for 10 years, with previous experience as a network consulting engineer in Cisco's advanced services group. Prior to joining Cisco, he was a consulting engineer with MCI's Professional Managed Services group. John has been designing or troubleshooting enterprise networks for the past 15 years. In his spare time, he enjoys working with microprocessor-based electronic projects including wireless environmental sensors. John holds CCIE certification 5232. He holds a bachelor of science degree in electrical engineering from the University of North Carolina's Charlotte campus.

Bart McGlothin

Bart McGlothin, Vertical Solutions Architect, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Bart is a Retail Architect at Cisco Systems. With over 13 years of industry experience, Bart leads Cisco's involvement with the National Retail Federation's Association for Retail Technology Standards Committee. Prior to Cisco, Bart worked as the Network Architect at Safeway, Inc.

.

CONTENTS

CHAPTER 1

Mediator Overview

This guide discusses network design considerations when deploying an energy management solution consisting of Cisco Network Building Mediators onto a converged IP network infrastructure. More specifically, the primary focus of this design guide is on enterprise network infrastructure design in order to provide secure access to the management interface (often referred to as the north side) of the Mediator, for managed service provider (MSP) partners and internal corporate facilities management personnel.

This guide does not discuss network design and connectivity of actual building systems devices connected to what is often referred to as the south side interface of the Mediator; other than to recommend isolation of this segment from the rest of the network infrastructure. The features and functionality of the Mediator discussed within this document are primarily presented from the perspective of how to support the underlying protocols, and not to provide an in-depth understanding of the Mediator itself. This design guide discusses the deployment of the energy management solution both with and without a separate virtual routing and forwarding (VRF) instance. The VRF method tested and discussed in this document uses VRF-Lite with point-to-point GRE tunnels. This is considered scalable for small to moderate-sized deployments. Future revisions of this design guide may discuss more scalable VRF methods such as VRF-Lite end-to-end and VRFs with MPLS. Finally, this document does not discuss interoperability of the Cisco Network Building Mediator with Cisco EnergyWise technologies. Future revisions may address both designs for building systems devices (south side designs) as well as integration with Cisco EnergyWise.

Problems of Traditional Building Systems

Traditional building systems consist of siloed networks built and maintained as individual systems, such as lighting; heating, ventilating, and air conditioning (HVAC); metering; fire; uninterruptible power supplies (UPS); video surveillance; physical access; and others. The duplication of networks for each of these systems results in higher installation, commissioning, and maintenance costs. Many of the systems that consume energy within buildings implement communication protocols and formats, limiting access to important information and building functionality. Proprietary building automation systems and black boxes provide access to only a subset of the energy consuming systems within a facility. The lack of unification among all these disparate building systems and the lack of centralized monitoring and control across global operations leads to inefficiencies and increased energy consumption.

New Open Energy Management Systems

Cisco's Network Building Mediator is an open, any-to-any networked energy, facility, and sustainability platform developed specifically to connect to the wide range of existing building systems and normalize building system informational data. Since all points within the framework are identified by a unique identifier (URI) and all information can be presented in common formats, such as HTML or XML-RPC, the Mediator allows for a number of other parties to securely consume and manipulate this information. These different parties might include both operations staff performing diagnostics and executives examining customer reports via their browser. These benefits are also extended to value add service providers that specialize in specific areas, such as building systems analytics, predictive maintenance, or renewable energy solutions which rely solely on the Mediator as a systems aggregator to tenants controlling their personal environment via their VoIP phone and other intelligent machines performing automated operations. Once this data has been liberated by the Mediator and these disparate protocols represented in a uniform IP-centric fashion, all of the information from these systems, which exist in virtually every building in the world, can now be leveraged for the sole benefit of improving operations. For example, using cloud services such as Automated Demand Response (ADR), this data can be correlated across each system at a site, multiple systems at a site, and multiple sites over time. Underperforming sites can be identified and adjusted, resulting in significant energy savings and cost reductions. Through the use of controlled energy systems, it is also possible to participate in an ADR and dynamic-pricing programs from utility companies, potentially gaining additional cost savings. The Network Building Mediator will also provide critical energy usage and forecast information to Smart Grid programs as they become available.

Cisco Network Building Mediator

The Cisco Network Building Mediator is the centerpiece of the open sustainability and energy management solution. It is a hardened network appliance connecting disparate building systems of various communication protocols onto the IP network. Cisco routing platforms have connected multiprotocol networks for years; now this functionality is extended to include building systems with the Mediator. The Mediator is available in the two models shown in Table 1-1.

Table 1-1 Cisco Network Building Mediator Models

Model	Description	Building Control Protocol Licensing Options
Cisco Network Building Mediator 4800	Targeted for campus deployments, it supports up to approximately 5,000 points.	Base, intermediate, and advanced protocols
Cisco Network Building Mediator 2400	Targeted for branch deployments, it supports up to approximately 1,000 points.	Base and intermediate protocols

Note A *point* or *datapoint* is a generic term used to describe a single item of information in a building control system. Examples of points include the temperature of a room, duct pressure of an air handling unit (AHU), and chiller water flow rate.

Figure 1-1 shows a Cisco Network Building Mediator with the icon used in figures of this document.

Figure 1-1 Cisco Network Building Mediator

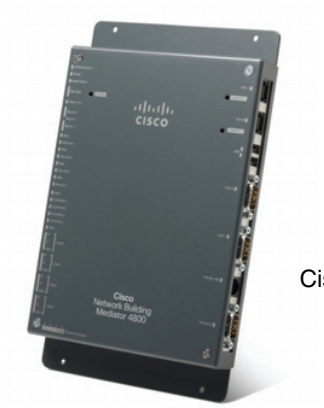

The Mediator aggregates and normalizes building systems data, making it available through an open XML interface.

Network Design Implications

When designing a converged IP network infrastructure to support both traditional IT services (for example, voice, video, and data applications) and energy management systems, the design engineer should be particularly aware of the security implications. These security requirements must be balanced against the business requirements of the energy management system itself, including its evolution over time. The Cisco Network Building Mediator contains two 10/100 Base-T Ethernet ports, one of which can be used for the management network segment, while the other can be used for the segment which houses IP-based building systems devices. These interfaces are typically referred to as north-side for the management interface and south-side for the building systems interface. In addition, the Mediator also supports a variety of communications and I/O ports, including two RS-232 ports, four RS-485 ports, four Dallas 1-Wire ports, four pulse counter inputs, and two solid-state single-pole relay outputs for connecting to building systems devices. Figure 1-2 shows a closeup of the communication and I/O ports of the Mediator.

Figure 1-2 Close-up of Mediator Ports

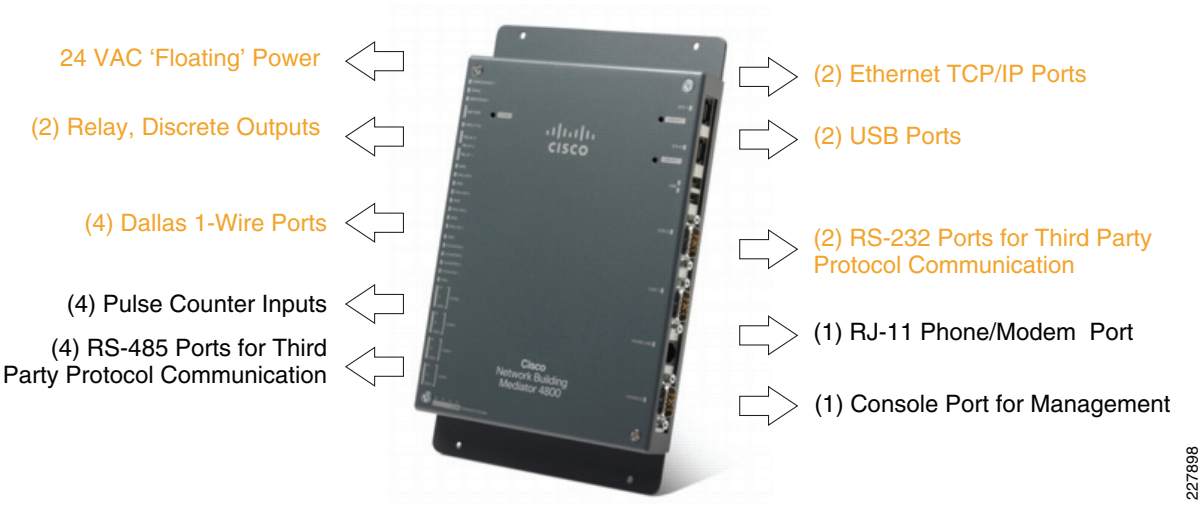

When the Mediator is integrated with critical energy and facility management systems, it is recommended to improve security by isolating the 10/100 FastEthernet network segments connected to the Mediator from the rest of the IP network infrastructure and tightly controlling access to these network segments. The management network segment (for example, a north-side segment) should be separated wherever possible from the network segment to which the building devices are connected (for example, a south-side segment), especially when using IP-based energy management systems protocols such as BACnet/IP, Modbus/TCP, etc. An example is shown in Figure 1-3.

Figure 1-3 10/100 FastEthernet Connectivity Example on the Mediator

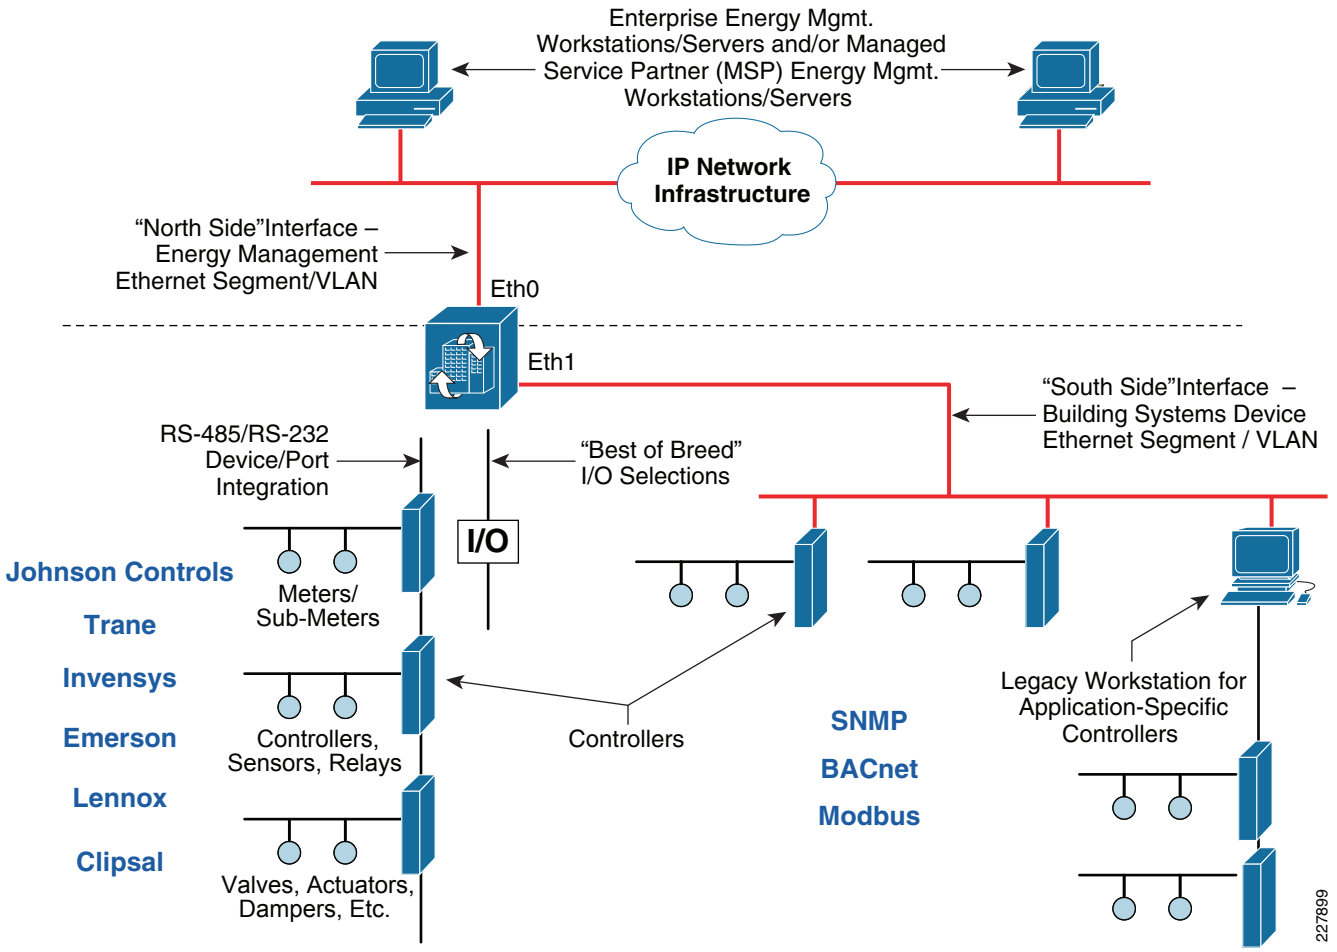

Figure 1-3 shows an example of a Mediator connected to a number of floor, zone, or room controllers by way of both traditional RS-232/RS-485 wiring and a dedicated Ethernet segment. The controllers are in turn connected to the actual building energy management devices—meters, sub-meters, valves, actuators, dampers, etc. The controllers on the Ethernet segment may be running open standards-based protocols such as Bacnet/IP or Modbus/TCP. Although many of the open standards for IP-based energy management systems protocols have security features such as encryption and authentication, actual implementations by vendors may not offer these security features. Many offerings of IP-based energy management systems protocols often use broadcast technologies, requiring the need for flat networks and/or specialized broadcast servers. Therefore, isolating these network segments is considered prudent.

The Mediator can also interface with legacy management workstations, as shown in Figure 1-3. This may be desirable in situations where application-specific controllers exist within the deployment. In such cases the Mediator serves as a Web-based thin-client monitor solution, while application changes are handled by the legacy management workstation. Alternatively, the programming for the application-specific controllers may be duplicated within the Mediator and the legacy management workstation removed.

Network isolation within the LAN infrastructure can be accomplished through several methods, including separate physical switches dedicated to energy management systems. The preferred method is to use separate logical VLAN segments provisioned off of a converged switch infrastructure. Using a

converged switch infrastructure design has the advantage of lower overall hardware and reoccurring maintenance costs. Access control to the energy management systems segments can be accomplished through the following methods:

- Dedicated firewall appliances, such as the Cisco ASA 5500 Series.
- Firewall services integrated within a router or switch platform, such as the Context-Based Access Control (CBAC) or Zone-Based Policy Firewall (ZBPF) features of Cisco ISR router platforms, or the Firewall Services Module (FWSM) of the Cisco Catalyst 6500 Series switch platforms.
- Access-control lists (ACLs) within a Layer-3 switch or a router platform.
- Site-to-site or client-based IPSec VPN connectivity.

The deployment of path virtualization technology such as virtual routing and forwarding (VRF) can also be used to isolate the energy management solution and limit access control to one or more strategic locations within the IP network infrastructure. The application of the various access control methods within different parts of the network infrastructure is discussed in detail within individual chapters of this design guide. The following chapter discusses some of the information flows and network protocols required on the energy management interface or north side of the Cisco Network Building Mediator for operation over the IP network infrastructure.

CHAPTER 2

Deployment Models and Information Flows

This chapter discusses common deployment models for an energy management solution using the Cisco Network Building Mediator. Following the deployment models, a detailed discussion is provided of the protocols necessary for the operation of the management interface of the Mediator over the IP network infrastructure.

Deployment Models

The deployment of energy management systems often follows two models. In the first model, a managed service provider (MSP) deploys or uses a Cisco partner to deploy the system for the enterprise customer. The customer or the MSP manages the system on a day-to-day basis. This deployment model implies full management and monitoring capabilities to and from the Mediators for both the MSP and the enterprise customer concurrently. The most common method for the MSP to provide this service is connectivity via IPSec VPNs. Note that other interactive data flows to entities such as a utility company may be required for automated-demand response (ADR) or dynamic pricing applications. Partner VPN connectivity may be centralized or distributed. Centralizing the partner VPN connectivity to a single entrance point, such as a campus location, provides a more scalable, cost effective, and manageable solution. However, it also requires MSP partner traffic to traverse the enterprise network infrastructure.

Centralized partner VPN connectivity is typically used for medium to large sized energy management solution deployments, where the locations are already connected via an enterprise WAN infrastructure. Distributed partner VPN connectivity is typically used for small energy management solution deployments with only a handful of independent locations, which may or may not be connected by an enterprise network infrastructure. Distributing the partner VPN connectivity typically requires Internet connectivity and VPN hardware at each site in which a Cisco Network Building Mediator is deployed. However, it does not require MSP partner traffic to traverse an enterprise network infrastructure.

With the second deployment model, the enterprise customer may use a Cisco partner to deploy and then manages the energy management system themselves. This model requires full management and monitoring capabilities to and from the Mediators and management workstations within the enterprise. However, partner VPN connectivity may not be needed. Logging data may still be exported from the Mediators to a MSP via the Internet in order to provide cloud services such as Energy Scoreboards, ADR, AFDD, and dynamic pricing applications. Alternatively, logging data may be collected within an enterprise archiving server and an internal Energy Scorecard deployed using a partner or internally developed.

In addition to the choice of whether an MSP will deploy and manage the solution, or whether the solution will be managed by the enterprise customer, the network design engineer must also decide whether traffic isolation is a requirement in order to support the energy management solution. Network virtualization refers to the creation of logically isolated network partitions overlaid on top of a common enterprise physical network infrastructure.

Network virtualization is accomplished through the deployment of Virtual Routing and Forwarding (VRF) technology. VRF technology is a path isolation technique used to restrict the propagation of routing information, so that only subnets belonging to a particular virtual network (VPN) are included in any VPN-specific routing tables. This results in the creation of independent logical traffic paths over a shared physical network infrastructure. VRFs can be used to separate and isolate energy management traffic flows from normal data traffic in order to provide an additional layer of network security for the energy management solution.

Figure 2-1 summarizes the choices for deployment of the energy management solution over an enterprise network infrastructure.

Figure 2-1 Energy Management Solution Deployment Flowchart

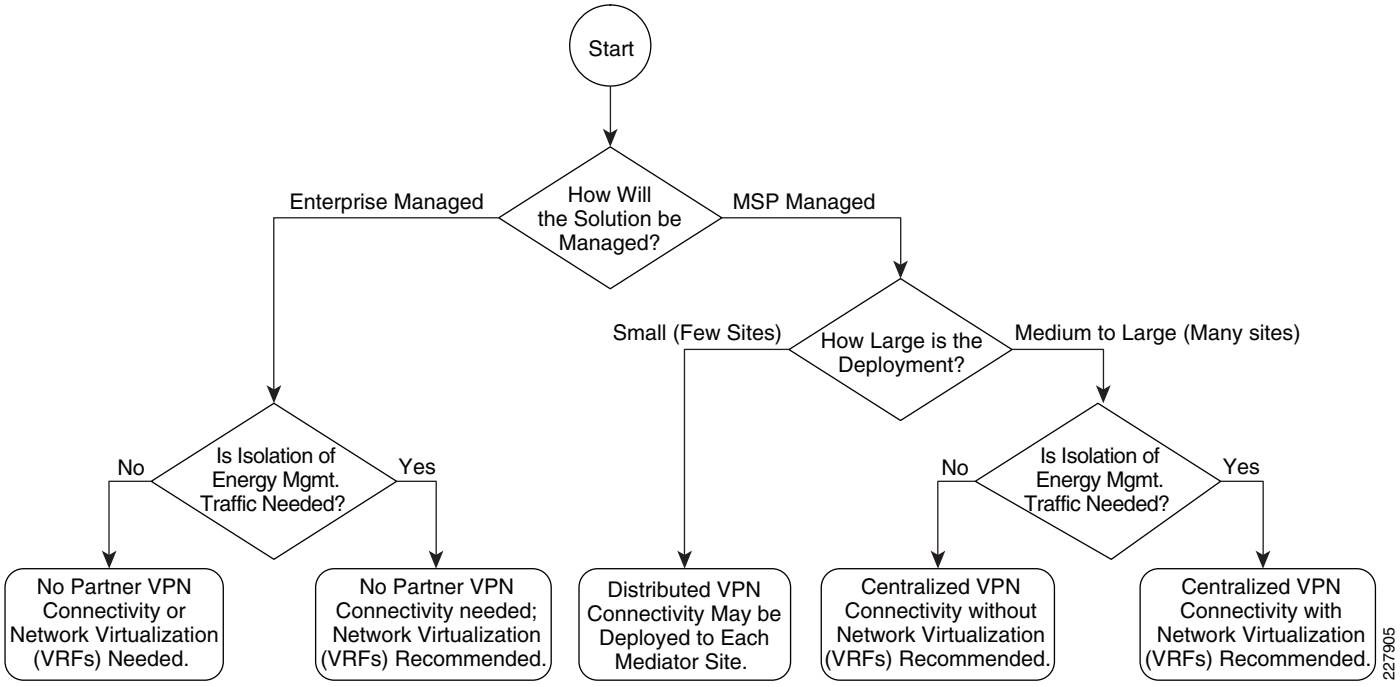

Information Flows

This section discusses some of the information flows and network protocols required on the energy management interface or north side of the Cisco Network Building Mediator for operation over the IP network infrastructure. These protocols can be separated into two broad categories---Management Services Protocols and Cloud Services Protocols.

Note The current version of this design guide does not include discussion of the information flows and protocols between the actual building devices (chillers, boilers, HVAC systems, etc) and the building systems interface or south side of the Mediator. Future revisions may include discussions of such flows as well.

Management Services Protocols

Management services include the protocols and information flows required for provisioning the Mediator onto the network infrastructure. They also include the protocols and information flows required for configuring the Mediators, creating control logic which is then deployed onto the Mediator, and for monitoring and managing the Mediators.

Device provisioning protocols can include Dynamic Host Configuration Protocol (DHCP), Domain Name Services (DNS), and Network Time Protocol (NTP). Figure 2-2 shows an example of these flows.

Figure 2-2 Management Flows—Device Provisioning

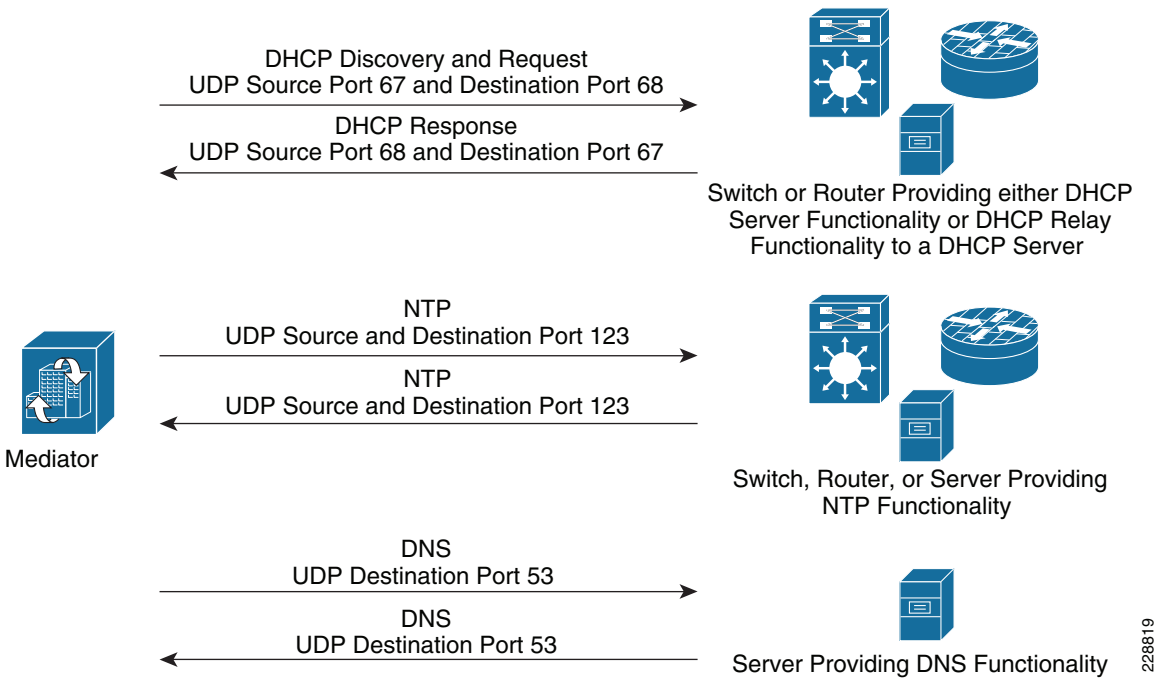

DHCP support is needed if the Mediator uses dynamic IP addressing. In such cases, either the local Catalyst switch or Cisco router connected to the Mediator Ethernet interface can be configured with DHCP server functionality to hand out IP addresses. Alternatively, the switch or router can be configured with DHCP relay functionality to relay the request to a DHCP server centrally located within a data center or campus service module. DHCP uses UDP ports 67 and 68.

Note From a security perspective, configuring access control down to specific IP addresses and protocols is discussed throughout this document. The use of DHCP to provide IP addressing to the Mediator may complicate the configuration of access control lists on firewalls, Layer 3 switches, and routers. In such

situations, the network administrator may need to either ensure the Mediator always receives the same IP address from the DHCP server, or the access control lists may need to be expanded to include the range of addresses which the Mediator may receive. Note that the second alternative presents a somewhat larger security concern since access control to the energy management solution is less specific.

DNS is required if the hostnames are configured within the Mediator. If hostnames are used, the Mediator must query a DNS server in order to translate the names to IP addresses in order to reach the destinations. For campus locations, DNS servers may be centrally located within a data center or campus service module. Additional DNS servers may be deployed within branch locations. DNS uses UDP destination port 53 for queries to the server and responses from the server.

NTP is recommended for time synchronization of devices across the network infrastructure. This is particularly important if schedules are implemented within the Mediator. Also, the periodic exporting of log data requires accurate time synchronization of Mediators in order to make sense of the logged data. Network administrators typically synchronize the clocks of network infrastructure devices, so the local Catalyst switch or Cisco router connected to the Mediator Ethernet interface can be configured with NTP functionality to synchronize the clock of the Mediator. Alternatively, a server centrally located within the data center service module or campus service module can serve as the NTP server to synchronize all Mediators. NTP uses UDP source and destination port 123.

Note that when deploying a separate VRF for the energy management solution (also referred to as the Building Infrastructure Network or BIN VRF within this document), the network administrator must carefully consider how DHCP, DNS, and NTP services are provided for the energy management solution. In deployments where centralized servers provide these functions, a separate set of servers dedicated to the energy management VRF may need to be deployed. Alternatively, connectivity between the energy management VRF and the another VRF (either the global VRF or a VRF dedicated to provide such services) may need to be provisioned. Chapter 5, “Data Center/Campus Service Module Design Considerations” discusses some design options around connectivity between the energy management VRF and other VRFs.

As mentioned above, management services also include protocols required for configuring the Mediators, creating control logic deployed onto the Mediator, and for monitoring and managing the Mediators. These services are provided through the management applications listed in Table 2-1.

Table 2-1 Mediator Configuration and Management Applications

Application Name	Description
configTOOL	ConfigTOOL is a software application that runs on an enterprise energy management workstation or managed service provider (MSP) partner workstation, which is used to configure the system settings, protocols, and services on the Mediator. The XML file created by configTOOL, which holds the Mediator configuration, is named broadway.xml.
perfectHOST	PerfectHOST is an application that runs on an enterprise energy management workstation or MSP partner workstation, which provides an intuitive graphical programming tool for creating control logic that resides on the Mediator. Logic is built by adding functional building blocks called templates to the canvas and connecting them together to create drawings. PerfectHOST comes with a pre-built library of 1,000 + templates such as I/O, logic, and protocol.
OMEGA	OMEGA is a software suite used to program and configure the Mediator. The OMEGA software tools are served to the browser of the enterprise energy management workstation or MSP partner workstation by the Mediator's internal Web server.

Specific tools within the OMEGA software suite are listed in Table 2-2.

Table 2-2 OMEGA Tools

OMEGA Tool	Description
System	OMEGA system allows the network administrator to view and modify Mediator network settings, backup and restore the Mediator, upload Mediator keys, reset OMEGA, and troubleshoot the Mediator through the Mediator message log.
EventManager	EventManager is an OMEGA tool for creating and viewing events on Mediators. It is typically used to define alarm conditions that, when met, display alarm events in EventManager and deliver e-mails. EventManager can view and acknowledge the events in multiple Mediators defined within a “cloud”.
SecurityManager	SecurityManager service allows the network administrator to define and manage access to resources on the Mediator, including authorization and restriction of the ability of users to view information, to modify settings, to add, modify, or delete files, etc.
TrendManager	TrendManager is a service that allows the network administrator to configure and manage trends. Trends are log nodes that record changes in the values of specified nodes over time. Trends can be viewed as graphs directly from the Mediator.
WebScheduler	WebScheduler allows network administrators to make customized project and system schedules with an Internet browser.
WebExpress	WebExpress is a Web page authoring tool within the Mediator which allows network administrators to create Web monitor drawings using customizable widgets, graphics, and live datapoints from the Mediator.
Web SiteBuilder	Web SiteBuilder allows network administrators to quickly create and customize the look and feel of the “Home Page” of the Mediator or default Web page.

Figure 2-3 shows the protocols required to enable the flows needed by the Mediator configuration and management applications.

Figure 2-3 Management Flows—Device Configuration, Logic Configuration, and Monitoring

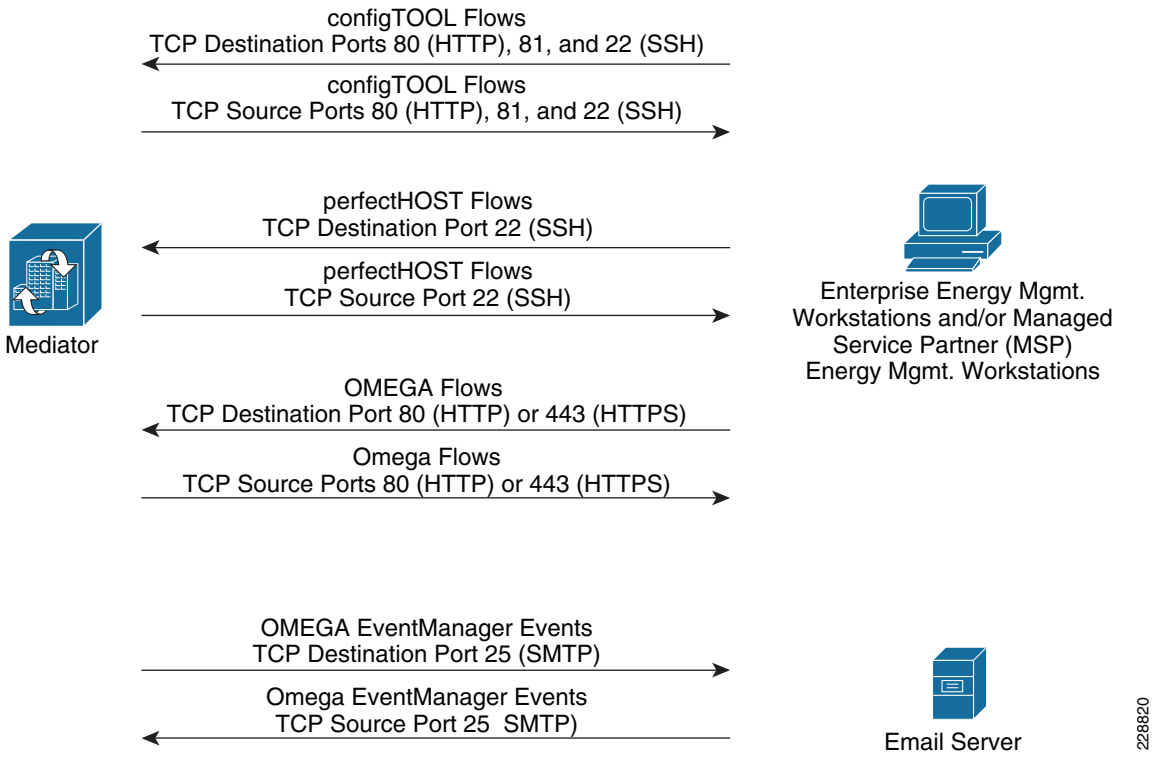

ConfigTOOL uses TCP ports 80 (HTTP), 81, and 22 (SSH) to establish connectivity to the Mediator and download the broadway XML file to configure it. PerfectHOST uses TCP port 22 (SSH) in order to establish connectivity and download the control logic it creates into the Mediators. The various applications which make up the OMEGA software suite use TCP ports 80 (HTTP) or 443 (HTTPS). Figure 2 shows that for each of these applications, the flow is initiated from the enterprise energy management or MSP partner workstation to the Mediators. However, the reverse traffic must be allowed through the network infrastructure for the session to be established as well. Note, however, that OMEGA EventManager can initiate events outbound from the Mediator to e-mail servers via TCP port 25 (SMTP).

One additional protocol that may be used by the Mediators is the Remote Node Abstraction (RNA) protocol. RNA is the protocol used to share node values (such as data points) between Mediators. RNA allows two or more Mediators on the same network to share points in a peer-to-peer manner or hierarchical manner. RNA uses TCP port 5150 as shown in Figure 2-4.

Figure 2-4 Management Flows—Remote Node Abstraction (RNA)

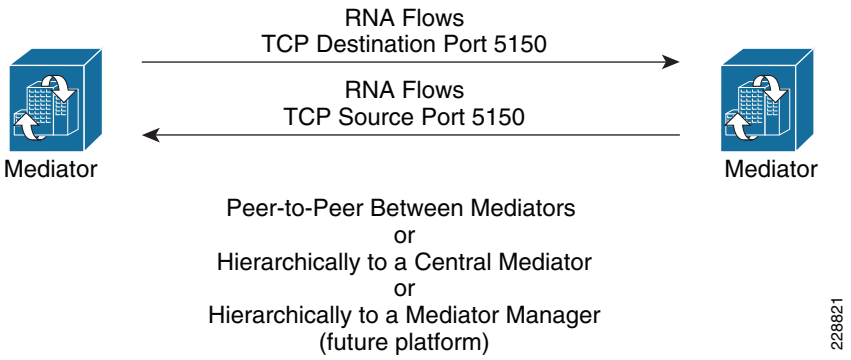

Figure 2-4 shows a scenario where the Mediator on the left has aliased a node on the Mediator on the right. Therefore the RNA flow is initiated by the Mediator on the left, with a destination TCP port of 5150. The response (return traffic) has a source TCP port of 5150. Chapter 5, "Data Center/Campus Service Module Design Considerations" discusses the hierarchical portal model further.

Cloud Services Protocols

Cloud services include protocols necessary for services such as Energy Scoreboards, Enterprise Energy Management (EEM), event reporting, ADR, dynamic pricing, and Automated Fault Detection and Diagnostics (AFDD). Data is typically logged and exported uni-directionally from the Mediator to cloud services servers located on the Internet. Datapoints can be periodically logged in intervals from 15 seconds to 1,800 seconds (30 minutes) and stored on the Mediator until they are ready to be exported. The Mediator is capable of exporting periodic logged data via File Transfer Protocol (FTP), Secure File Transfer Protocol (SFTP), Hypertext Transfer Protocol (HTTP) POST, Secure Hypertext Transfer Protocol (HTTPS) POST, and Simple Mail Transfer Protocol (SMTP). Periodic logged data can be exported in intervals from every minute to once every week (providing sufficient buffering capacity exists to hold the periodic logged data between exports). The Mediator is capable of sending the logged data in various formats, although the XML data format is most commonly used. Example 2-1 shows an example of an FSG XML format showing multiple data points.

Example 2-1 Example Mediator Data Export in FSG XML Format

```
<data info="RTP_Cisco_Systems" key="RTP000000001">
 <device info="ESE_Lab_Meter_A" key="rtp001">
  <channel name="Volts_B2C" Totalized="N" uom="Voltage" key="5" Delta="N"
meastype="Volts">
   <value timestamp="2009-09-24T13:15:00">280
   </value>
  </channel>
  <channel name="Amps_Phase_B" Totalized="N" uom="Current" key="2" Delta="N"
meastype="Amps">
   <value timestamp="2009-09-24T13:15:00">531
   </value>
  </channel>
  <channel name="Amps_Phase_C" Totalized="N" uom="Current" key="3" Delta="N"
meastype="Amps">
   <value timestamp="2009-09-24T13:15:00">516
   </value>
  </channel>
```

```
  <channel name="Amps_Phase_A" Totalized="N" uom="Current" key="1" Delta="N"
meastype="Amps">
   <value timestamp="2009-09-24T13:15:00">505
   </value>
  </channel>
  <channel name="KVA" Totalized="N" uom="Power" key="8" Delta="N" meastype="KVA">
   <value timestamp="2009-09-24T13:15:00">1261.41912
   </value>
  </channel>
  <channel name="Power_Factor" Totalized="N" uom="Ratio" key="9" Delta="N"
meastype="Ratio">
   <value timestamp="2009-09-24T13:15:00">0.95
   </value>
  </channel>
  <channel name="Volts_A2B" Totalized="N" uom="Voltage" key="4" Delta="N"
meastype="Volts">
   <value timestamp="2009-09-24T13:15:00">273
   </value>
  </channel>
  <channel name="Volts_A2C" Totalized="N" uom="Voltage" key="6" Delta="N"
meastype="Volts">
   <value timestamp="2009-09-24T13:15:00">273
   </value>
  </channel>
  <channel name="Volts_3_Phase" Totalized="N" uom="Voltage" key="7" Delta="N"
meastype="Volts">
   <value timestamp="2009-09-24T13:15:00">478
   </value>
  </channel>
 </device>
</data>
```

Figure 2-5 shows the protocols required to enable the data flows associated with the periodic export of logged data from the Mediators.

Figure 2-5 Cloud Services Flows—Periodic Export of Logged Datapoints

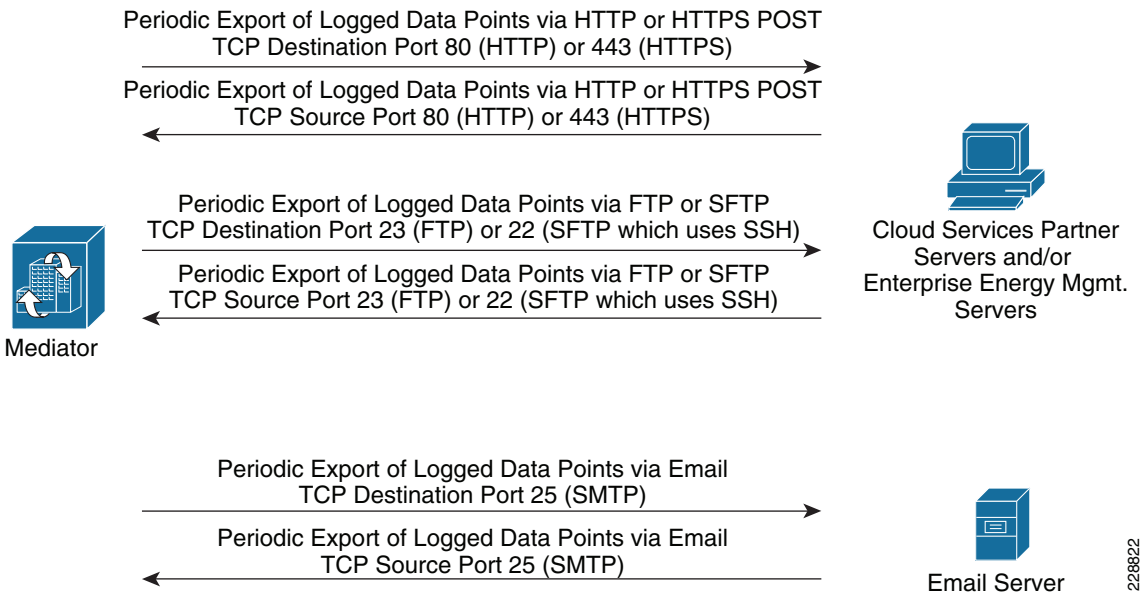

The data flows associated with the periodic export of logged data points are typically initiated from the Mediator outbound to a cloud services partner server and/or an enterprise archiving server. However, the reverse traffic must be allowed back through the network in order for the session to be established. Specific protocols required for management and cloud services should be identified and tied to unique source and destination IP addresses for firewall, ACL, or IPSec VPN connectivity. The use of secure protocols (for example, encrypted and authenticated) is highly recommended.

Note The FTP protocol dynamically opens a data channel separate from the control channel, which uses TCP port 23. Stateful firewalls with application-layer inspection open and close the data channel automatically, based upon inspection of the FTP control channel. However, when using ACLs, the network administrator may have to statically open the port range needed for the FTP control channel. Note that SFTP, which uses the SSH protocol (TCP port 22), does not use a separate dynamic data channel.

Automated Demand Response (ADR) systems often use standards-based protocols such as SOAP/WSDL over HTTP/HTTPS in order to exchange event state information between the Mediator and the energy service provider (utility companies, etc.). Figure 2-6 shows an example of some of the flows which may be involved.

Note ADR functionality has not been tested or validated with the Mediator for the current version of this design guide. The example below serves as possible example only. Future revisions may include ADR testing.

Figure 2-6 Cloud Services Flows—Example Automated Demand Response Flows

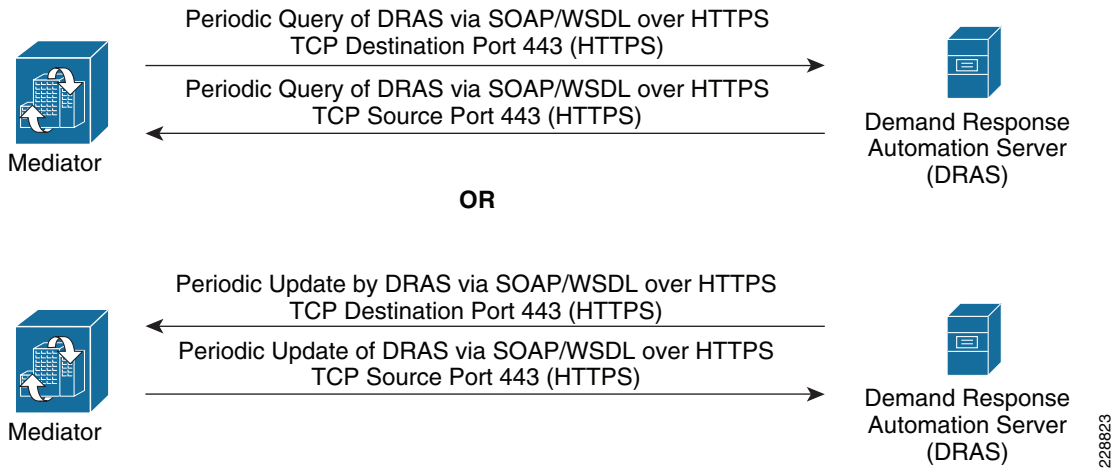

In the example shown in Figure 2-6, a server (sometimes referred to as an Demand Response Automation Server or (DRAS)) is deployed at the energy service provider location. The Mediator may function as a DRAS client, periodically polling the DRAS for event state information. Alternatively, the DRAS may periodically update the Mediator, although this requires the network administrator to allow connections initiated from the energy service provider into the enterprise network, which may be less desirable. Transport Layer Security (TLS) as well as userid and password are commonly used with ADR systems in order to ensure confidentiality and authenticate the sessions. The Mediator mat then use the event state information from the DRAS to implement pre-programmed logic, such as resetting the setpoints of thermostats, in order shed load and reduce energy consumption.

Specific design considerations for support of the energy management solution within campus and branch locations are discussed in detail in Chapter 3, "Campus Design Considerations" and Chapter 6, "Branch Design Considerations."

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

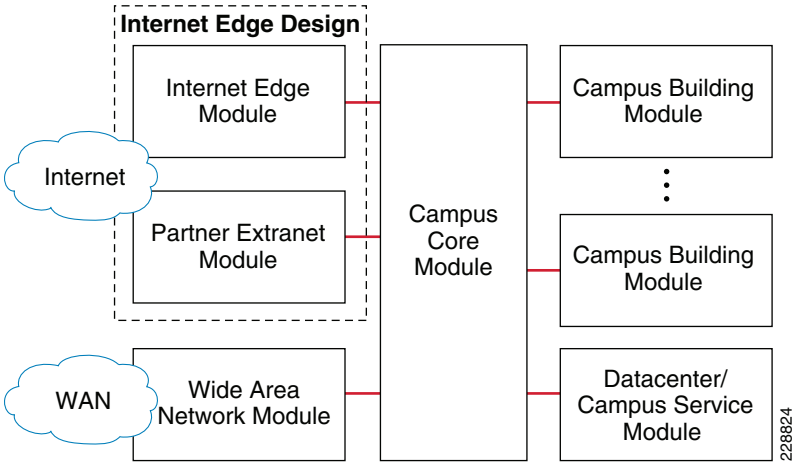

Table 3-1 provides a brief overview of the role of each module in a campus network.

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

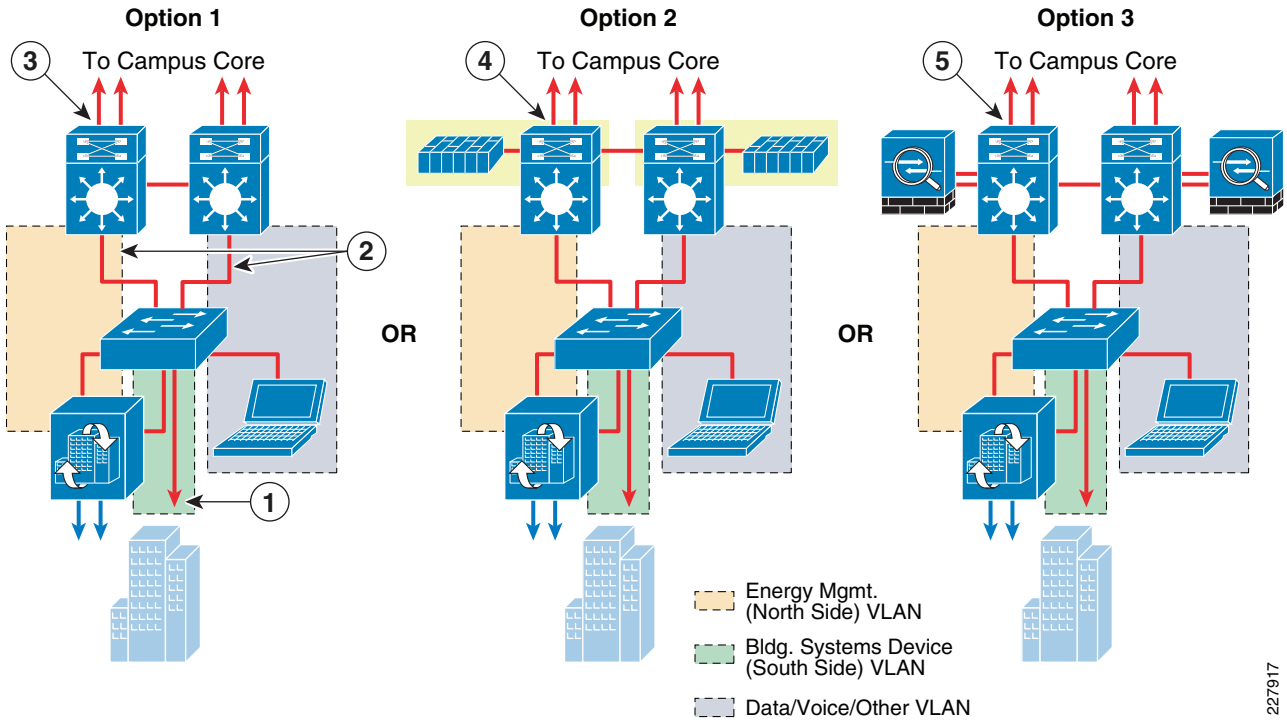

The following describes the numbers shown in Figure 3-2:

- **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

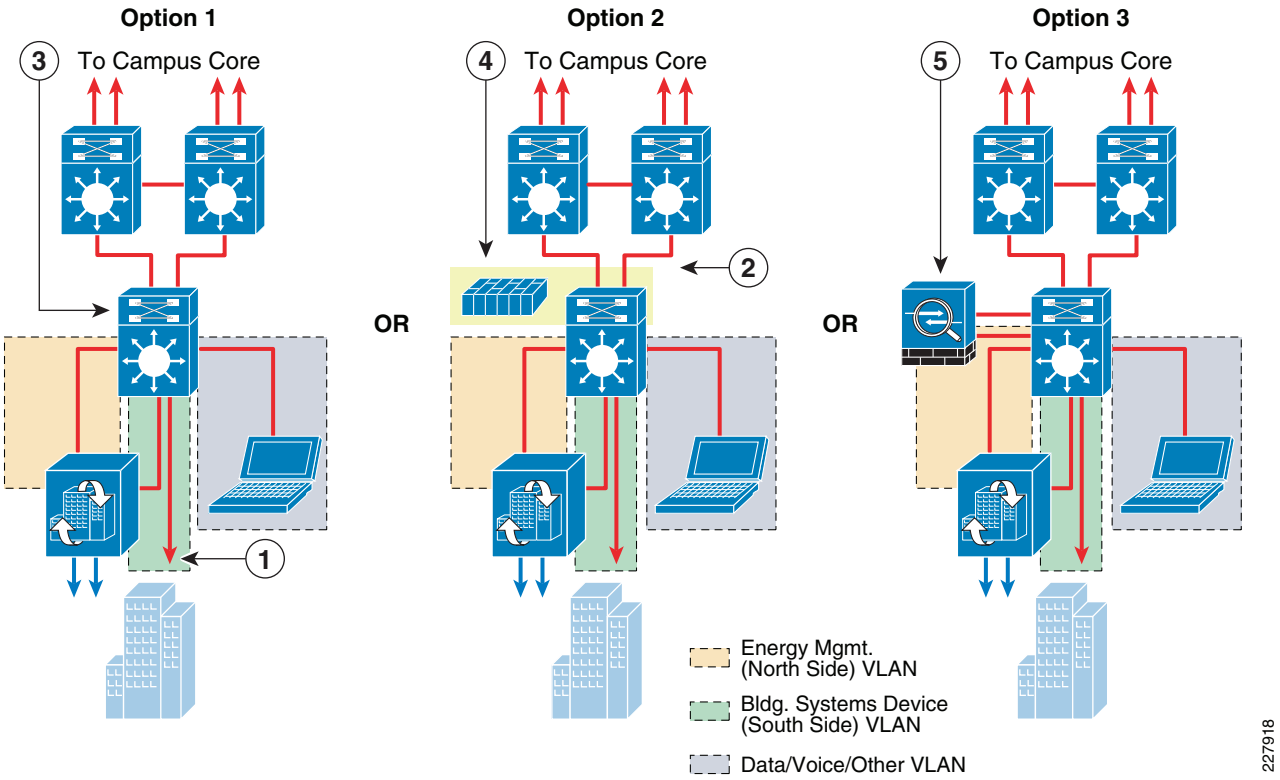

The following describes the numbers shown in Figure 3-3:

- **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

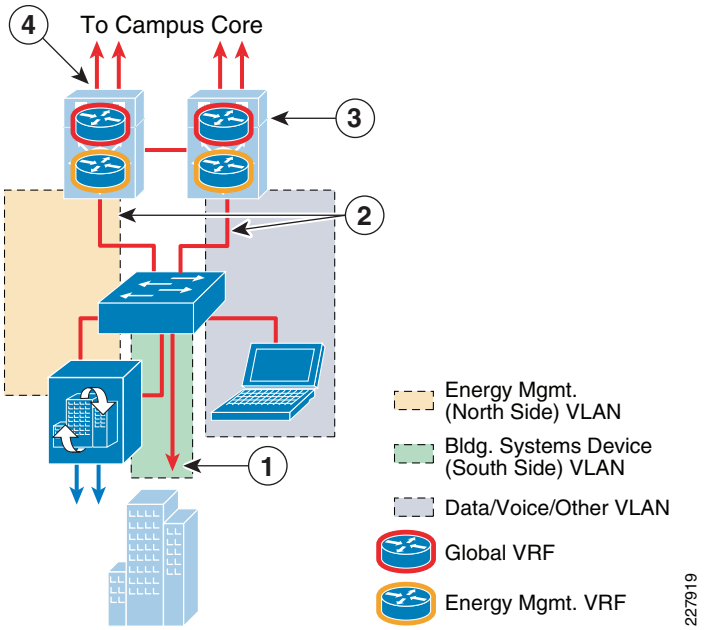

The following describes the numbers shown in Figure 3-4:

- **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

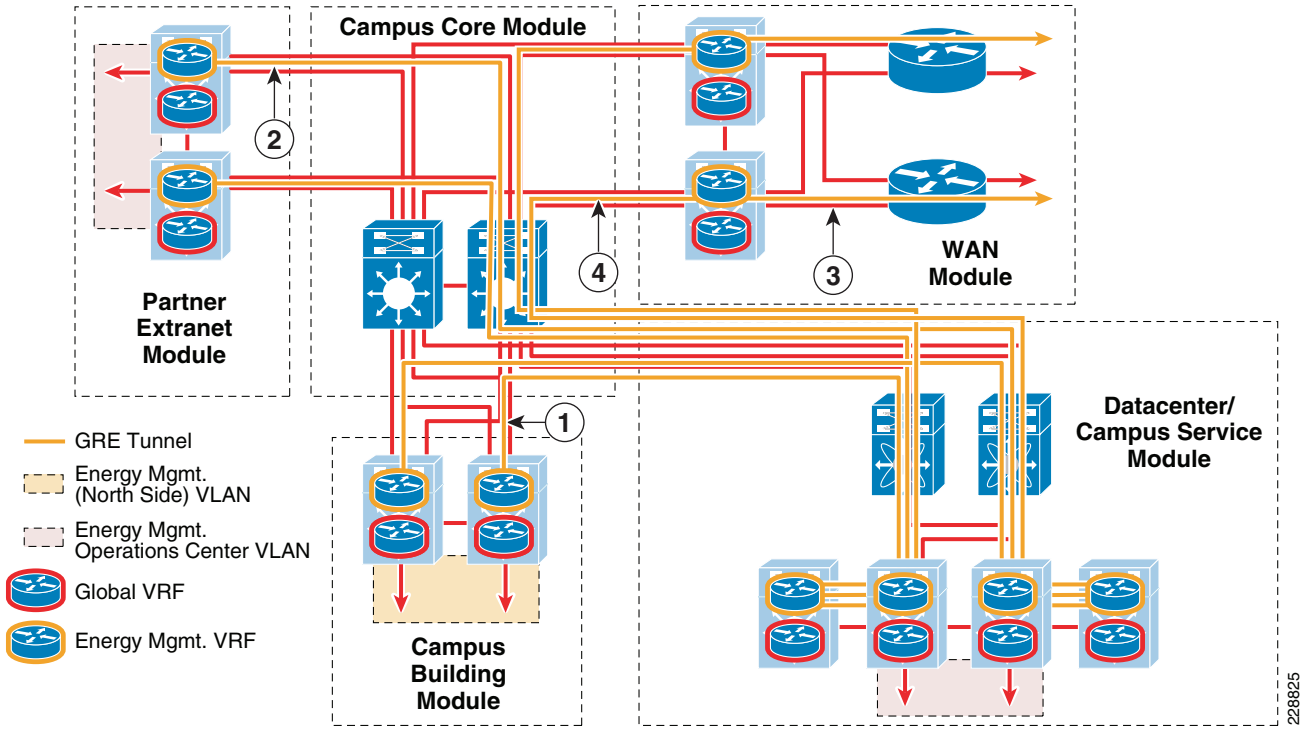

The following describes the numbers shown in Figure 3-5:

- **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

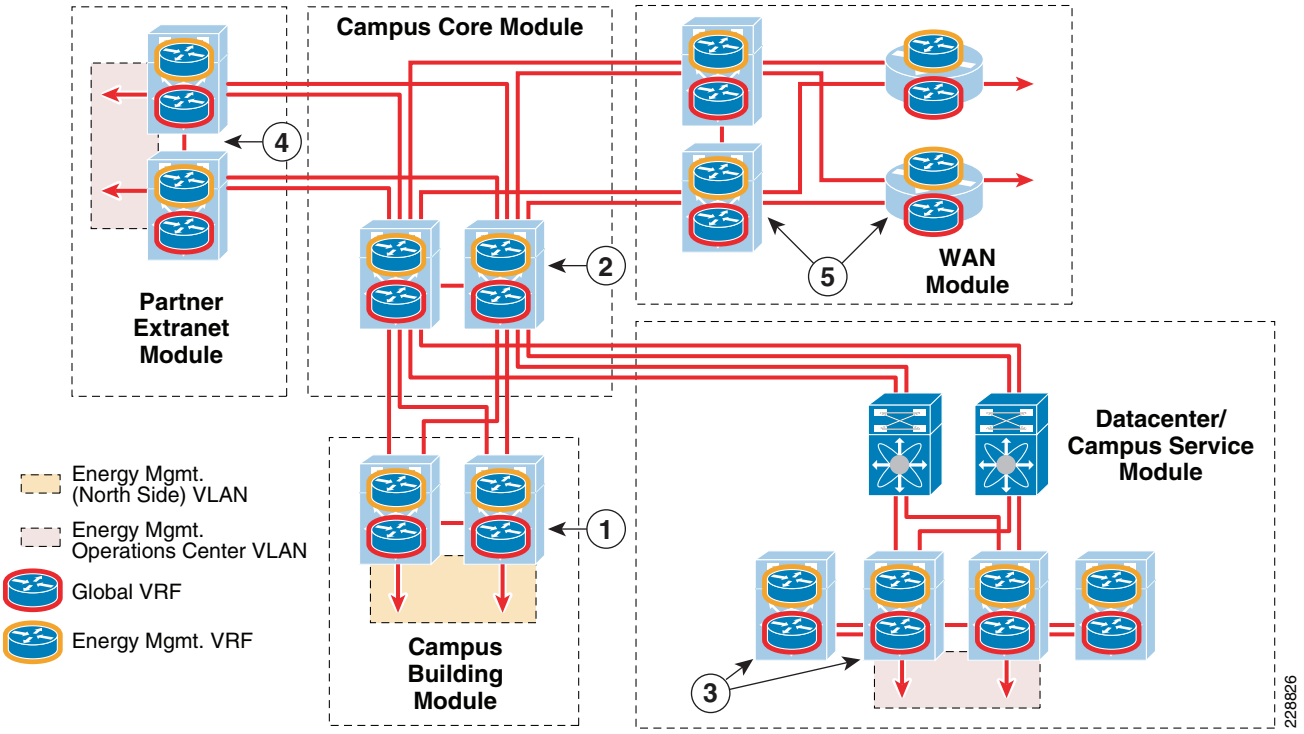

The following describes the numbers shown in Figure 3-6:

- **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

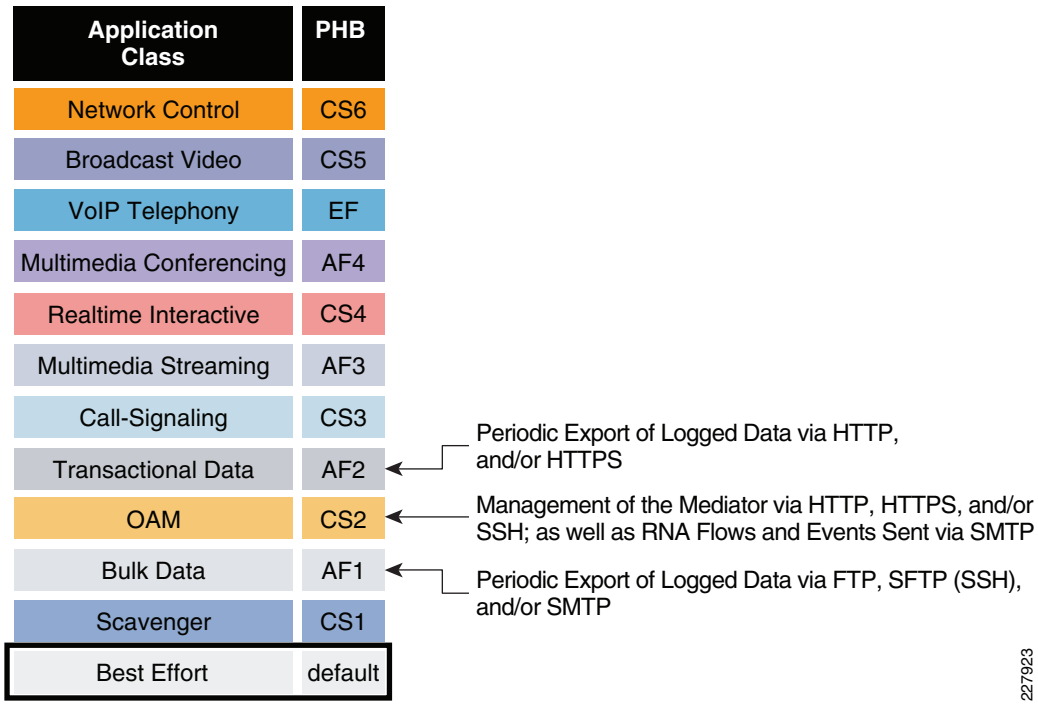

In this example, the network administrator may consider placing the periodic export of logged data into either the Bulk Data service class and marked with a DSCP value of AF1 or the Transactional Data service class and marked with a DSCP value of AF2. For example, if the logged data is exported very 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

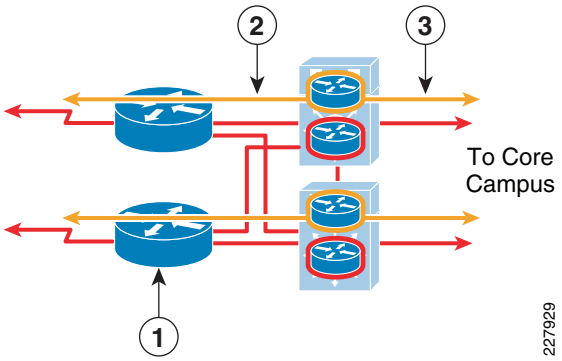

The following describes the numbers shown in Figure 3-9:

- **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

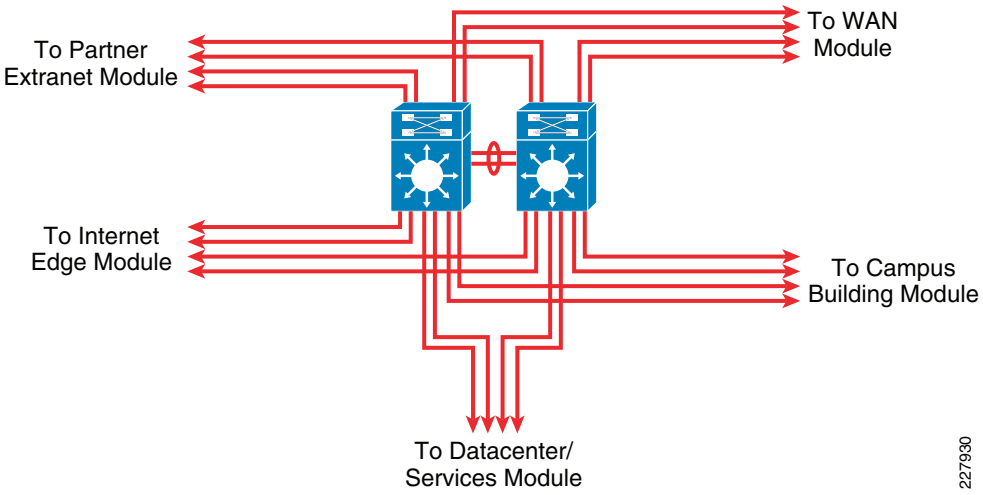

In a non-VRF implementation, nothing specific needs to be configured on the core Catalyst 6500 switches in order to support the energy management solution. However, in a VRF implementation, modifications may be necessary. If the network administrator chooses to deploy the VRF with GRE Tunnel model, then 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.

CHAPTER 4

Internet Edge Design Considerations

The Internet edge design provides three main services for the Cisco Network Building Mediator deployment:

- Internet connectivity through which logged data is exported from the Mediators to cloud services partners who provide features such as energy scorecards. Internet connectivity may also be used for interactive data exchanges to support capabilities such as automated demand response (ADR).
- VPN connectivity through which managed service provider (MSP) partners connect to the enterprise network in order to manage Mediator deployments.
- The Internet edge maintains strict access control to the Mediator deployment from devices outside the enterprise network

Internet edge designs vary considerably between organizations, and are often driven by the existing security policy of the particular business entity. Large enterprise organizations often separate employee Internet connectivity from partner connectivity through the deployment of separate network modules. Within each module, functionality may be converged onto a single device or separated onto multiple devices. This type of separation of access between employees and partners, and separation of functionality onto multiple devices, can provide more granular control of partner access to the enterprise, and is consistent with current Cisco SAFE best practices. For further information regarding network security best practices, the reader is encouraged to review the *Cisco SAFE Reference Guide* which can be found at the following URL:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html

The Internet edge designs presented in this guide follows this design paradigm; presenting an overall Internet edge design for large organizations, with a separate Internet Edge Module and Partner Extranet Module, and separate devices for each functionality. However, for smaller organizations, a collapsed Internet edge design, where a single Internet Edge Module is deployed for both employee and partner access, is also discussed. In this design, all functionality is collapsed into a single device. These two design options represent both extremes (complete separation of functionality versus complete convergence of functionality) of Internet Edge designs. Actual Internet Edge designs typically lie somewhere between the two designs presented within this document.

Internet Edge Module

With respect of the energy management solution, the function of the Internet Edge Module is to provide stateful access control for outgoing connections initiated by the Mediators to cloud services partner servers accessible via the Internet. This is typically for the periodic exporting of logged data from the Mediators. However, the network administrator should note that in certain scenarios (such as if network virtualization is implemented for the Mediator deployment), exporting of logged data may be supported

through the Partner Extranet Module. This is discussed further in the "Extending VRFs to the Partner Extranet Module" section on page 4-25. The Internet Edge Module also provides stateful access control for enterprise client PCs accessing web-based energy scorecards provided by cloud services partners reachable through the Internet. An example of a redundant Internet Edge Module design is shown in Figure 4-1. Only the highlighted components, which are relevant to the energy management solution, are discussed in this document.

Figure 4-1 Example Internet Edge Module

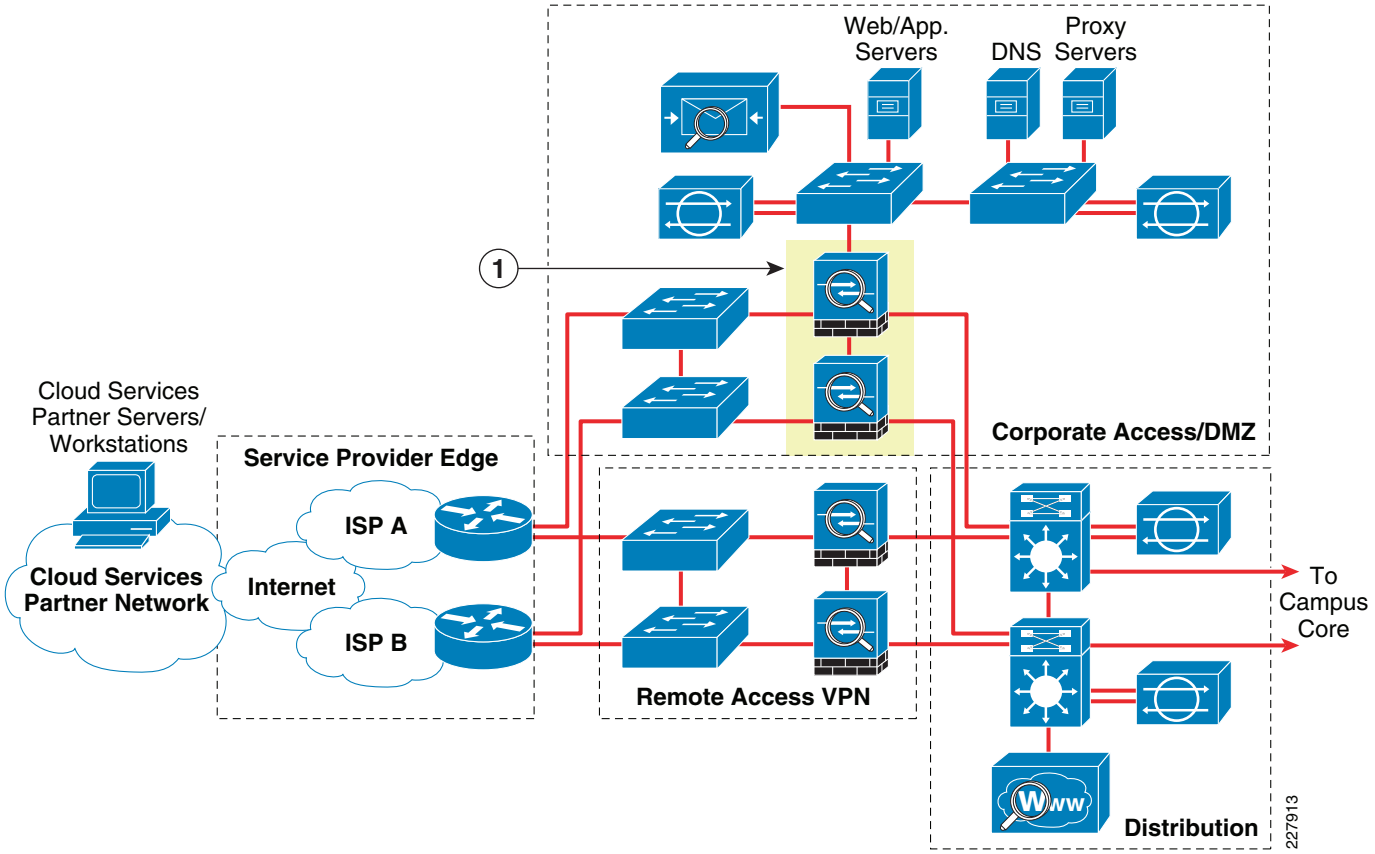

The following describes the number in Figure 4-1:

- 1—ASA 5500 Security Appliances deployed within the Corporate Access/DMZ section provides address translation and stateful access control for outgoing connections to cloud-services partners.

The internal IP addressing of the Mediators and client PCs within an enterprise network is typically hidden by translating it to Internet routable addressing via Network Address Translation (NAT) functionality within the Internet Edge Module. A redundant pair of ASA 5500 Series Security Appliances within the Corporate Access/DMZ section of the Internet Edge Module, highlighted in Figure 4-1, can provide the required address translation and stateful access control services within this module.

Note The network administrator should note that the remote access VPN section of the Internet Edge Module is intended for employee VPN access only. With this type of design, partner VPN access is instead handled within the Partner Extranet Module, discussed in the next section. This provides a clean separation of traffic between internal employees and partners accessing the enterprise network.

Examples of the connections that should be allowed through the ASA 5500 Security Appliances for energy management flows through the Internet Edge Module are shown in Figure 4-2, Figure 4-6, and Figure 4-7.

Figure 4-2 Data Export From Mediators Directly to Cloud Services Partner Servers

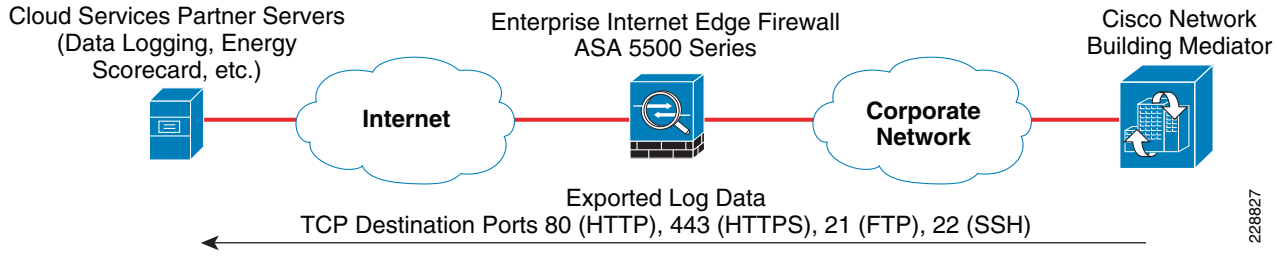

Figure 4-2 shows that the export of periodic-logged data is initiated by Mediators outbound toward the cloud-services partner servers. The Mediator is capable of exporting periodic-logged data via HTTP or HTTPS POST, FTP, or Secure FTP (SFTP) which uses the SSH protocol. Multiple data loggers and exporters can be configured within a single Mediator. However, in typical deployments, only one protocol is used for the data exports. Which protocol or protocols are used often depends on what the cloud-services partner can support, as well as what is acceptable to the enterprise organization.

Note The Mediator is also capable of exporting data via HTTP or HTTPS GET. However, this requires the cloud-services partner server to initiate an inbound connection through the Internet edge firewall directly to the Mediator. Due to the security implications of allowing an inbound connection into the enterprise network directly to the Mediator, this is not a highly recommended method of data export. The reader should also note that the Mediator shown in figures throughout this chapter can refer to remote Mediators deployed throughout branch and campus buildings; a hierarchical Mediator (as discussed within the Chapter 5, "Data Center/Campus Service Module Design Considerations") deployed within the Energy Management Operations Center (EMOC); or the future Cisco Network Building Mediator Manager appliance.

When using FTP for data export, the ASA 5500 Series Security Appliance will inspect the FTP control channel to dynamically open the data channel necessary for the actual file transfer. This is one of the benefits of implementing a stateful firewall with application-layer inspection capabilities at the Internet Edge. Application-layer inspection reduces the security holes created by having to statically open high-order port ranges in order to accommodate the randomly chosen data channel. Noted that FTP sends traffic unencrypted, including userids and passwords. For greater security, the network administrator should consider working with its cloud-services partner to implement a secure alternative to FTP such as Secure FTP (SFTP). SFTP uses the Secure Sell (SSH) protocol in order to provide data authentication and encryption of the data export. Additionally, with SFTP, the network administrator no longer has to be concerned with separate control and data channels. Likewise, when using HTTP POST for data export, the network administration should consider working with its cloud-services partner to implement HTTPS POST as an alternative, because HTTPS uses SSL/TLS to provide data authentication and encryption.

For each of the data exporters (HTTP or HTTPS POST, FTP, or SFTP), a username and password can and should be configured within the Mediator. An example of this is shown in Figure 4-3, which shows an SFTP transporter.

Figure 4-3 Example Configuration of a Username and Password for Data Export From Mediators to Cloud Services Partner Servers

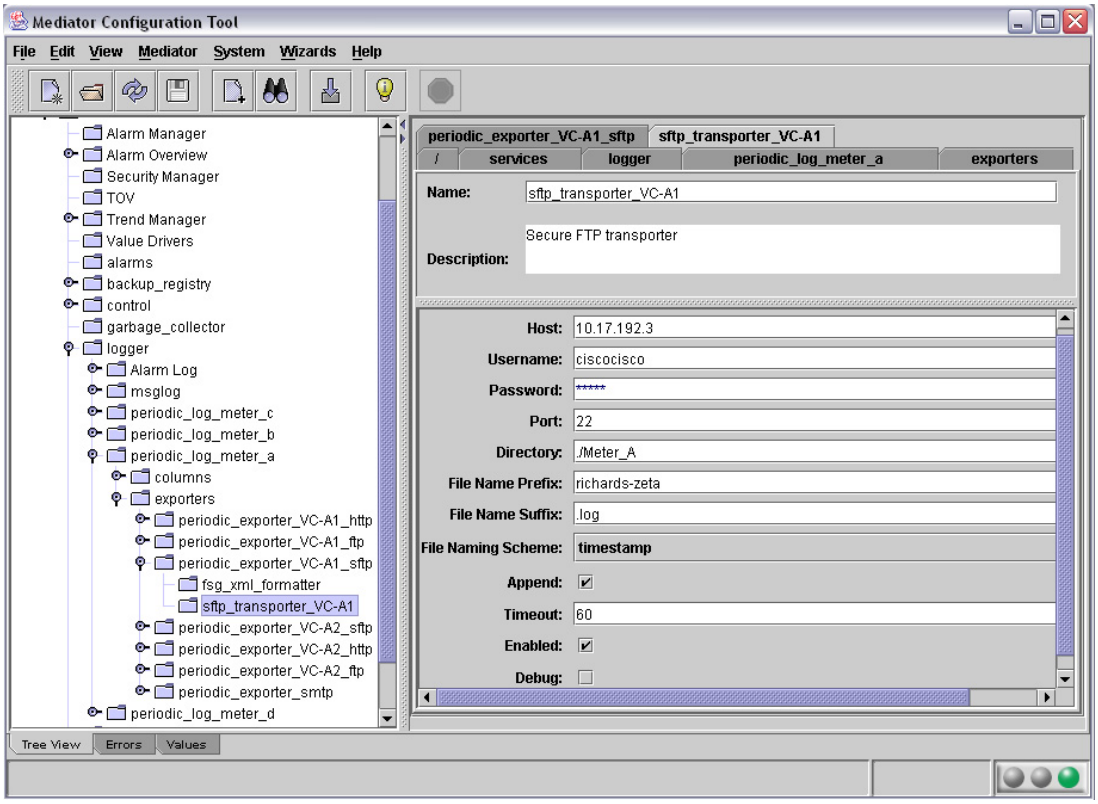

The username and password are used to authenticate the Mediator to the particular server (HTTP, FTP, or SSH server) before transferring the log files that contain datapoint information. The use of a unique username and password for the data export from each Mediator may be considered, versus using the same username and password configured for all Mediators within a deployment. The benefit to this approach is that if the username and password of a single Mediator is compromised, then only that particular username and password needs to be changed on both the Mediator and the cloud services partner server; versus having to change the usernames and passwords of all of the Mediators. The downside to this approach is that it can result in significantly more administrative overhead. Again, the network administrator should consider working with its cloud-services partner in order to determine if the additional security benefits outweigh the administrative overhead of maintaining separate usernames and passwords for data export from each Mediator.

Another alternative is to export log files from multiple Mediators through a secure proxy server located within the Corporate Access/DMZ section of the Internet Edge Module. An example of this is shown in Figure 4-4.

Figure 4-4 Data Export From Mediators to Cloud Services Partner Servers via Proxy or Drop-and-Forward Server

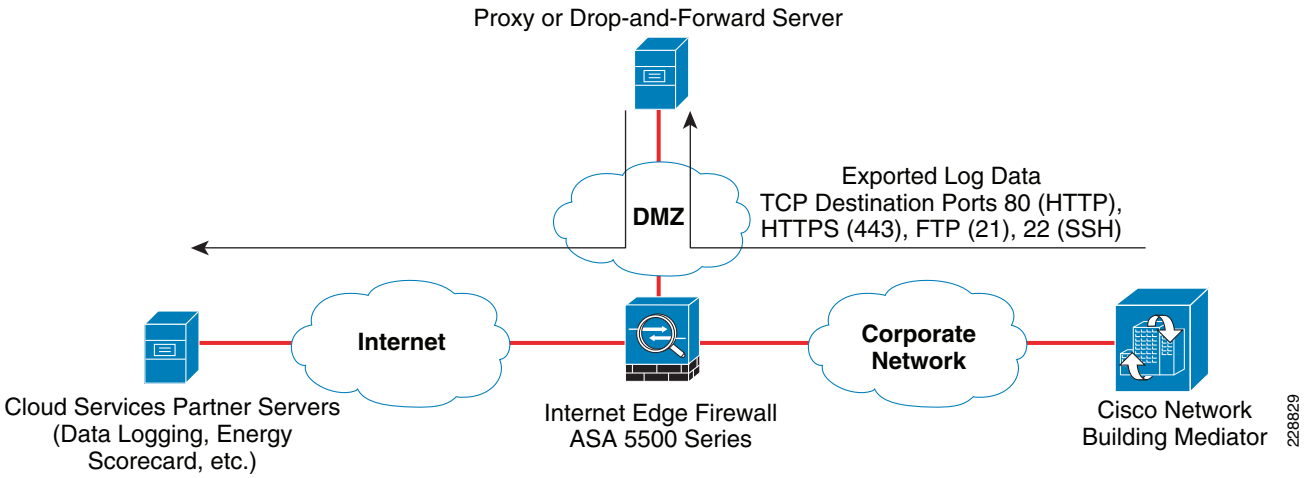

When configuring an FTP exporter within the Mediator to use a proxy server, the configuration may need to be modified, as shown in Figure 4-5.

Figure 4-5 Example Configuration for Data Export From Mediators to Cloud Services Partner Servers via FTP Proxy Server

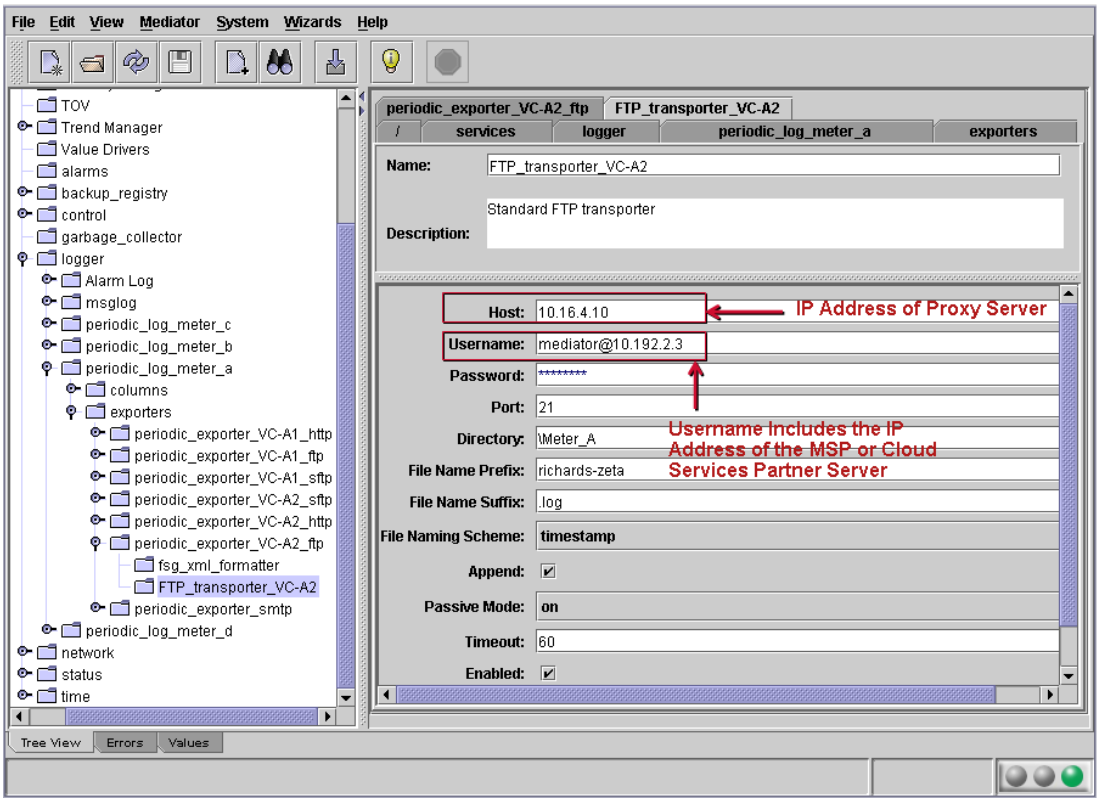

As shown inFigure 4-5 above, the host field has been modified to point to the IP address of the FTP proxy server located within the DMZ of the enterprise network. Note that a hostname can be used, if DNS resolves the hostname to an IP address. The username field has also been modified to include the username as well as specifying the IP address of the actual cloud services partner server. This is

accomplished via the use of the @ symbol. This configuration can be used to allow the FTP proxy server to automatically "forward" the exported data logs from the Mediators to the cloud-services partner servers.

Note A basic FTP proxy design was tested for this design guide using the ftp.proxy FTP proxy-server (http://www.ftpproxy.org/) running on a Fedora Linux server.

Using an FTP proxy server has the advantage in that the FTP proxy server itself is the only device that establishes an FTP connection outside the corporate network to the cloud-services partner server for data export. Firewall configurations can therefore be tightened somewhat to allow only the FTP proxy server or servers out, versus allowing every Mediator to individually establish FTP sessions out. The network administrator should note, however, that similar benefits are derived from a firewall, such as the Cisco ASA 5500 Series Security Appliance, running application-layer inspection of FTP and using NAT to hide the internal IP addressing of the Mediators. In large Mediator deployments, the scalability of the FTP proxy server or servers should be thoroughly assessed, in order to ensure that they do not result in a bottleneck for data exports from the energy management solution. The FTP proxy server itself should also be thoroughly hardened, due to its placement on the DMZ segment of the firewall.

For protocols that may be difficult to proxy, such as SFTP that uses the SSH protocol, a drop-and-forward mechanism may be implemented. In this scenario, the Mediators export to a server sitting on a the firewall DMZ segment; which in turn, periodically exports to the cloud-services partner server. An alternative method is for the network administrator to work with their cloud services partner to determine if the cloud services partner can periodically establish a secure connection inbound to the drop-and-forward server and retrieve the log files exported by the Mediators. One advantage of this scenario is that a public/private key pair can be used to authenticate the enterprise drop-and-forward server to the cloud services partner server; instead of the username and password used by the Mediator currently. The username and password configured within the Mediator are still used, but only to authenticate to the drop-and-forward server sitting on a DMZ segment within the corporate network. Therefore, the decision whether to use a single password for exports from all Mediators within the deployment, or to use individual passwords for each Mediator; now shifts completely to the network administrator. The network administrator should note, however, that the additional drop-and-forward function can add additional delays in exporting log files to the cloud services partner. For business requirements that need frequent exporting of log data for near-real-time use, this may not be a feasible design. However, for business requirements that need very infrequent exporting of log data for long-term trending use, this may be a feasible design.

Note A basic drop-and-forward design was tested for this design guide using OpenSSH (http://www.openssh.com/) running on a Fedora Linux server, along with simple shell scripts and a cron job.

The network administrator should note that both the proxy or drop-and-forward functions add another layer of complexity, additional hardware, and additional software to the energy management solution design. The network administrator should carefully evaluate the additional advantages of deploying either of these methods against the potential additional costs of supporting such a design. Finally, note that the use of proxy or drop-and-forward servers for outbound sessions is often dictated by the security policy of the enterprise organization.

Figure 4-6 shows that the Mediator can also export either periodic logged data or event data, via E-mail, using the Simple Mail Transfer Protocol (SMTP).

Figure 4-6 Data Export and/or Event Forwarding to Partners via E-mail

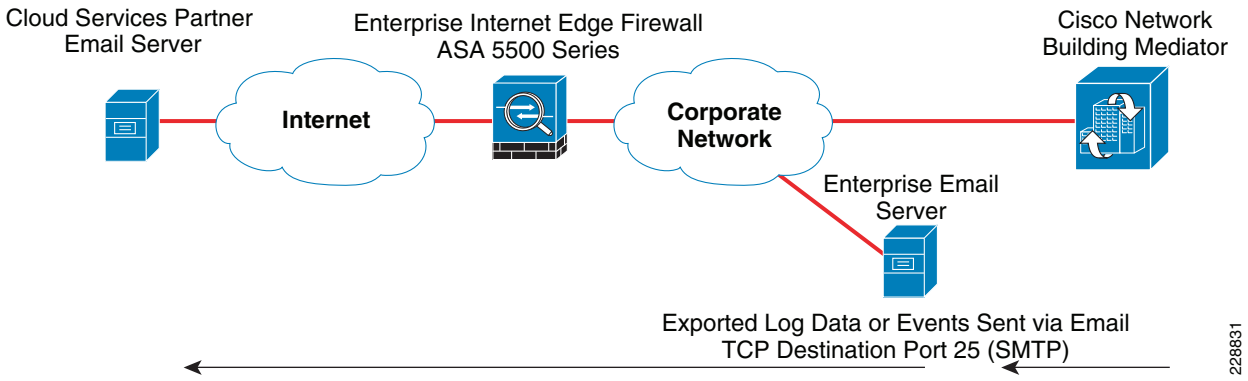

In this scenario the data originating from the Mediator is either embedded within the E-mail message or included as an attachment, and sent to the corporate E-mail servers. The logged data or event data can then be forwarded as normal E-mail from the enterprise E-mail servers to an account on the cloud services partner E-mail servers; passing through the ASA 5500 Security Appliances on the Internet edge. (Note that this example has been highly simplified from actual corporate E-mail systems.)

Figure 4-7 shows an example of an enterprise client PC accessing an energy scorecard located on a cloud services partner server, via a web browser.

Figure 4-7 Client PC Energy Scorecard Access

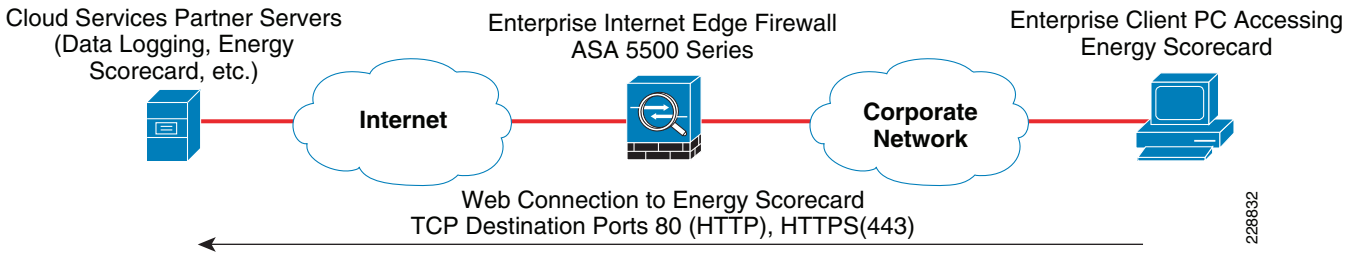

In this scenario, HTTP or HTTPS sessions initiated by client PCs need to be allowed outbound to the cloud services partner servers, through the ASA 5500 Series Security Appliance. As with the exporting of periodic logged data, the network administrator should consider working with the cloud services partner to implement HTTPS instead of HTTP in order to provide data authentication and encryption of the data flow. In addition, username and password protection of the energy scorecard site is highly recommended.

Note The port numbers for each of the service flows shown in Figure 4-2, Figure 4-6, Figure 4-7 are default values. These are configurable within the Cisco Network Building Mediator. The network administrator can choose to use different port numbers for HTTP, HTTPS, FTP, SFTP (SSH), and/or SMTP. if desired. Note that, however, changing the default ports may create issues with application-layer inspection within the Internet edge firewall, and is not highly recommended.

The actual configuration of the ASA 5500 firewall, in order to allow the exporting of periodic logged data or event data from the Mediators, as well as allow access from client PCs to the Energy Scorecard, is highly dependant upon the security policy of the organization. Security operations personnel should be involved within the discussions involving Cisco Network Building Mediator deployments. Some enterprise organizations allow all outbound connections directly to the Internet, but restrict inbound

connections. Other organizations limit direct outbound connections to certain protocols. Still other organizations may both limit outbound connections to certain protocols, and force client PCs through a proxy server for well known protocols such as HTTP and FTP. The following partial ASA 5500 firewall configuration (non-redundant) shows a very simplified example (see Example 4-1) where all direct outbound Internet access initiated by devices within the corporate network is allowed.

Example 4-1 Example of a Partial ASA 5500 Firewall Configuration for Internet Edge Access

```
!
interface GigabitEthernet0/1             ! Inside interface with higher security level.
 description CONNECTION TO ME-EASTDIST-1 G5/1
 nameif inside
 security-level 100
 ip address 10.16.2.2 255.255.255.252
!
interface GigabitEthernet0/3              ! Outside interface with lower security level.
 description CONNECTION TO OC3 INTERNET VIA ME-EASTINET-3 & 4
 nameif outside
 security-level 25
 ip address user_inet 255.255.255.248
!
~
!
access-list inside_nat_outbound extended permit ip any any
!                                          ! Causes all Inside addresses
!                                          ! to use NAT when connecting to Internet
hosts.
~
!
nat-control                                ! Enables NAT for Inside to Outside
connections.
!
!
global (outside) 1 interface               ! Uses the Outside Interface Address for the
global pool (PAT).
!
nat (inside) 1 access-list inside_nat_outbound   ! Specifies the NAT access-list for the
                                                 !Inside Interface.
!
~
!
class-map inspection_default               ! Specifies default class-map for
application-layer
 match default-inspection-traffic          ! inspection
!
policy-map global_policy                   ! Name of the inspection policy map
 class inspection_default
  inspect ftp                              ! Enables FTP inspection.
  inspect h323 h225
  inspect h323 ras
  inspect netbios
  inspect rsh
  inspect rtsp
  inspect skinny
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect tftp
  inspect sip
  inspect xdmcp
  inspect http
  inspect dns
```

```
!
service-policy global_policy global         ! Applies global_policy map to the firewall.
!
```

By default, traffic from an interface with a lower security level to an interface with a higher security level is blocked by an implicit access-list, unless specifically allowed by a user-defined access-list applied inbound on the outside interface. Therefore, all inbound access initiated by devices on the Internet is blocked. By default, traffic from an interface with a higher security level to an interface with a lower security level is allowed by an implicit access-list, unless specifically denied by a user-defined access-list applied inbound on the Inside interface. Therefore, all outbound Internet access initiated by devices within the enterprise organization is allowed. In this example, all outbound connections use the IP address of the outside interface of the firewall, in a Port Address Translation (PAT) configuration. In larger organizations, typically a pool of IP addresses is defined, and a combination of NAT and/or PAT is used to provide network address translation from internal corporate addressing to external Internet routable addresses. This not only provides a more efficient means of using the existing Internet routable address space, but can also provide an additional layer of security by hiding the real internal addressing structure of the enterprise organization. Finally, application-layer inspection of protocols such as FTP is enabled, allowing the ASA 5500 firewall to inspect FTP control channels in order to dynamically open the required data channels.

Partner Extranet Module

When managed service provider (MSP) partner access is required, separate VPN connectivity to each campus building or branch location is possible but not scalable. Instead, it is recommended to centralize the VPN access from the MSP partner, providing a more maintainable and cost effective solution. In terms of the energy management solution, the primary function of the Partner Extranet Module is to provide the centralized termination of MSP partner VPN connections by which the Mediators are subsequently managed. The Partner Extranet Module also provides stateful access control of the traffic that flows through the VPN tunnels and across the enterprise network to the Mediators. Additionally, the Partner Extranet Module can be used to as an alternative means of providing Internet connectivity through which periodic logged data is exported from the Mediators to cloud services partners; either by way of proxy servers located on a partner DMZ, or through the deployment of network virtualization within the enterprise network. An example of a redundant Partner Extranet Module design is shown in Figure 8 below.

Figure 4-8 Example Redundant Partner Extranet Module Design

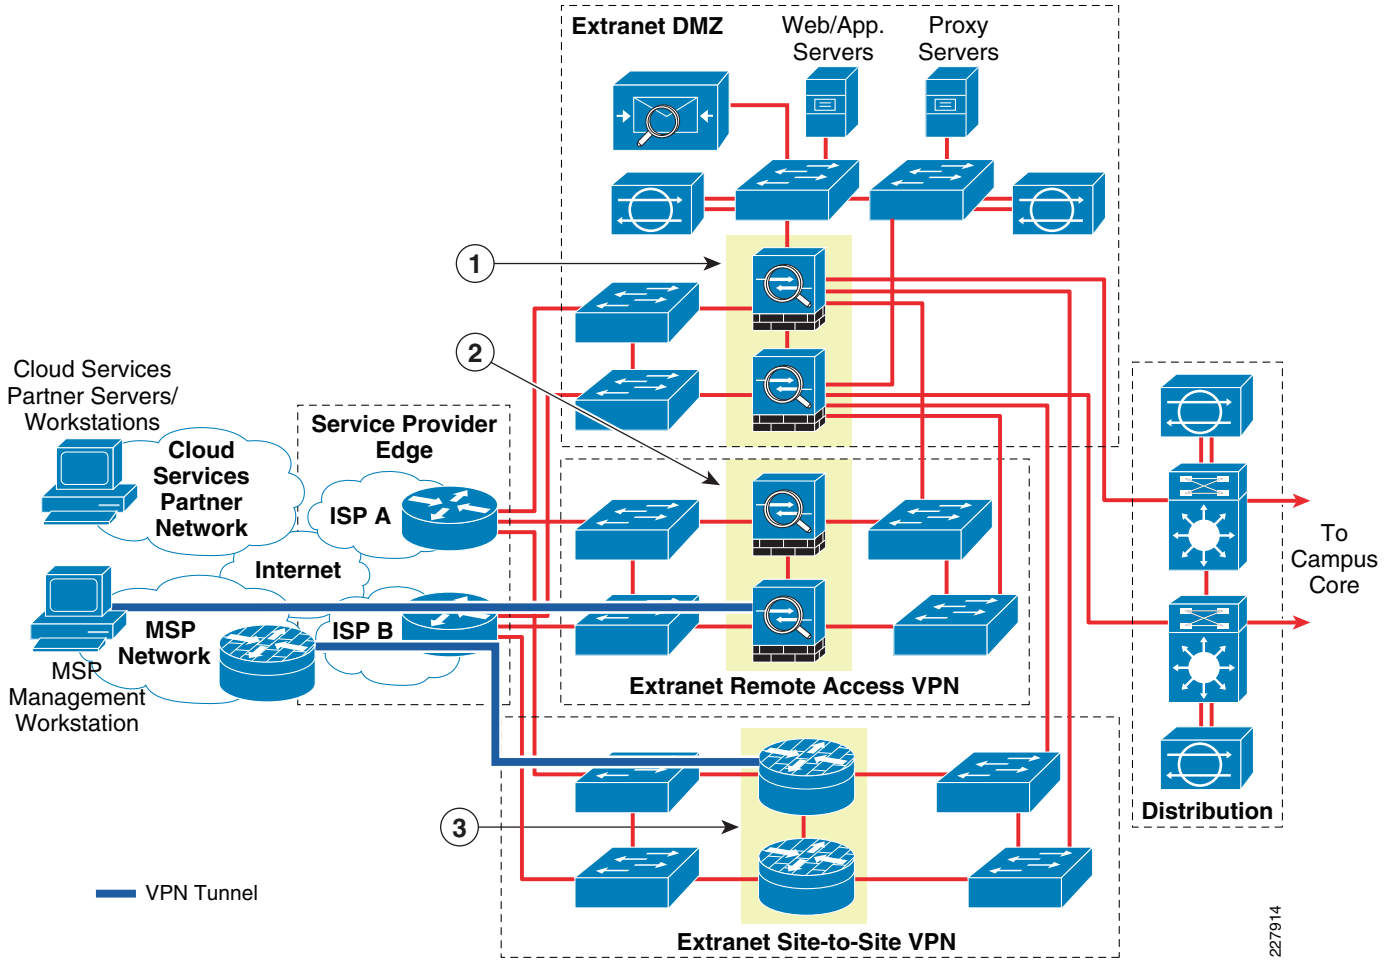

The following describes the numbers shown in Figure 4-8:

- **1**—ASA 5500 Security Appliances deployed within the Extranet DMZ section provides address translation and stateful access control for outgoing connections to cloud-services partners and incoming connections from managed service provider (MSP) partners through VPN devices.
- **2**—ASA 5500 Security Appliances deployed within the Extranet remote-access VPN section provides remote-access VPN termination, address translation, and IP address assignment for MSP partners.
- **3**—Cisco 1000 Series ASRs or Cisco 7200/7300 Series routers deployed within the Extranet site-to-site VPN section provides site-to-site VPN termination, stateful access control, and potentially address translation of managed services partner VPN connections.

Again, only the highlighted components that are relevant to the energy management solution are discussed in this document.

Note In Figure 4-8 above, separate IP subnets are shown connecting the routers within the service provider edge section to separate switches within the Extranet DMZ, Extranet remote-access VPN, and Extranet site-to-site VPN sections of the Partner Extranet Module. Since address translation is typically

performed within firewalls and VPN concentrators, the outside addresses of these devices are usually within the Internet routable addressing space. Since separate IP subnets normally require a larger amount of Internet routable IP addressing space to be provisioned from the service provider, enterprise organizations often combine these into a single IP subnet.

Examples of the connections that should be allowed through the ASA 5500 Series Security Appliances and/or VPN routers for energy management flows through the Partner Extranet Module are shown in Figure 4-9 and Figure 4-10.

Figure 4-9 Inbound Management Flows From MSP Partner Workstations to Mediators

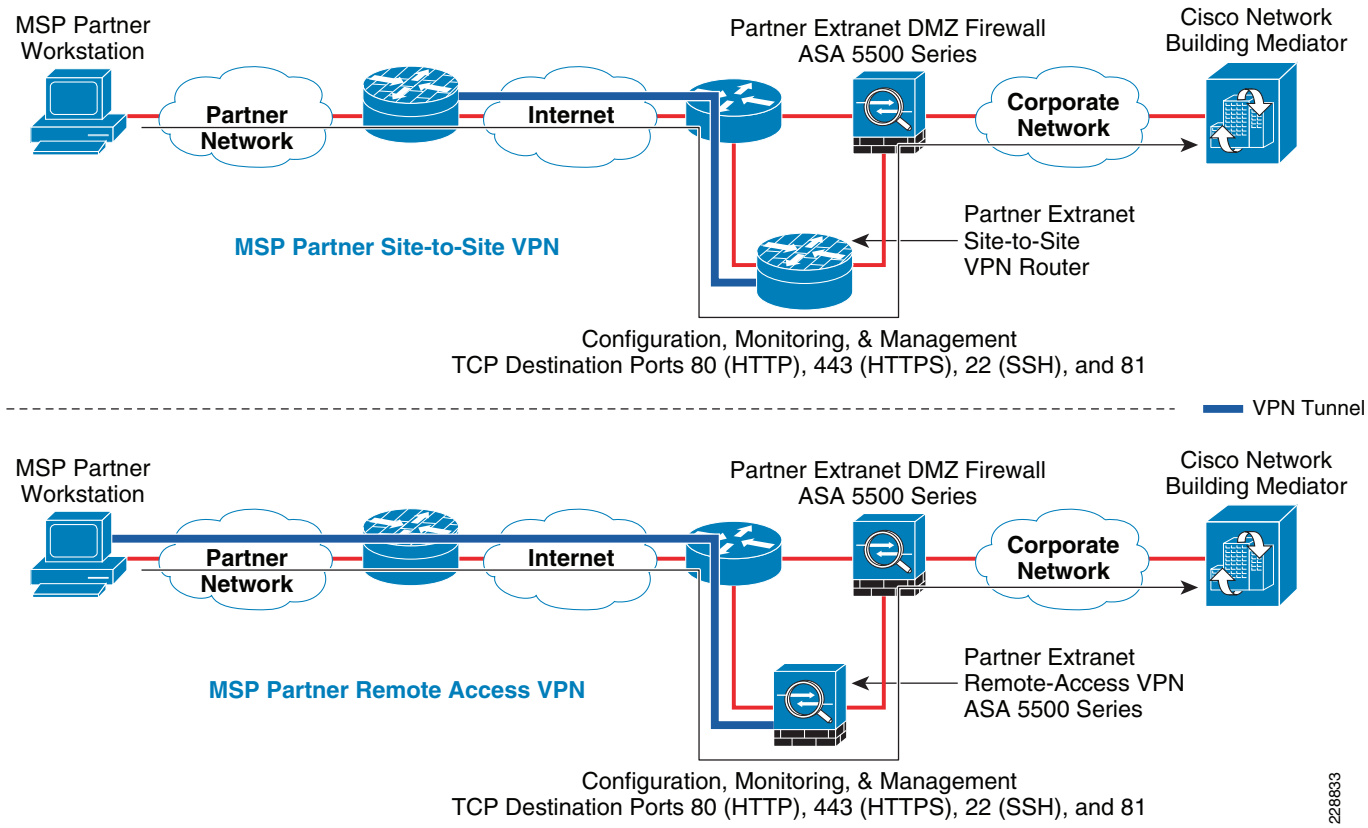

Figure 4-9 shows inbound management (configuration, monitoring, etc.) connections initiated from MSP partner workstations to the Mediators. The top half of the figure shows the flows through a site-to-site VPN tunnel, while the bottom half of the figure shows the flows through a remote-access VPN tunnel. In a site-to-site VPN configuration, the IPsec tunnel is extended from a router within the MSP partner network to a dedicated VPN router within the enterprise network. In a remote-access VPN configuration, the IPsec tunnel is extended from the MSP partner workstation to a dedicated ASA 5500 Series Security Appliance within the enterprise network. In either case, SSH (TCP port 22) is required in order for the MSP partner workstation to establish a connection to the Mediators for configuration via the ConfigTOOL application. The Omega suite of tools requires web-based access to the Mediators, either via HTTP (TCP port 80) or HTTPS (TCP port 443). Although the IPsec tunnel does protect traffic crossing the Internet between the MSP partner network and the enterprise network, the use of a secure protocol such as HTTPS can provide additional confidentiality and data integrity of the management flows as they traverse both the MSP partner network and the enterprise network. Note that the ConfigTOOL application also uses TCP port 81 to determine the state (online or offline) of the

Mediators. These protocols need to be enabled across the VPN tunnel itself; across any firewall software within the partner Extranet site-to-site VPN router or the remote-access VPN ASA 5500 Series Security Appliance; and across the partner Extranet DMZ firewall.

Figure 4-10 Data Export From Mediators to MSP Partner or Cloud Services Partner Servers via VPN

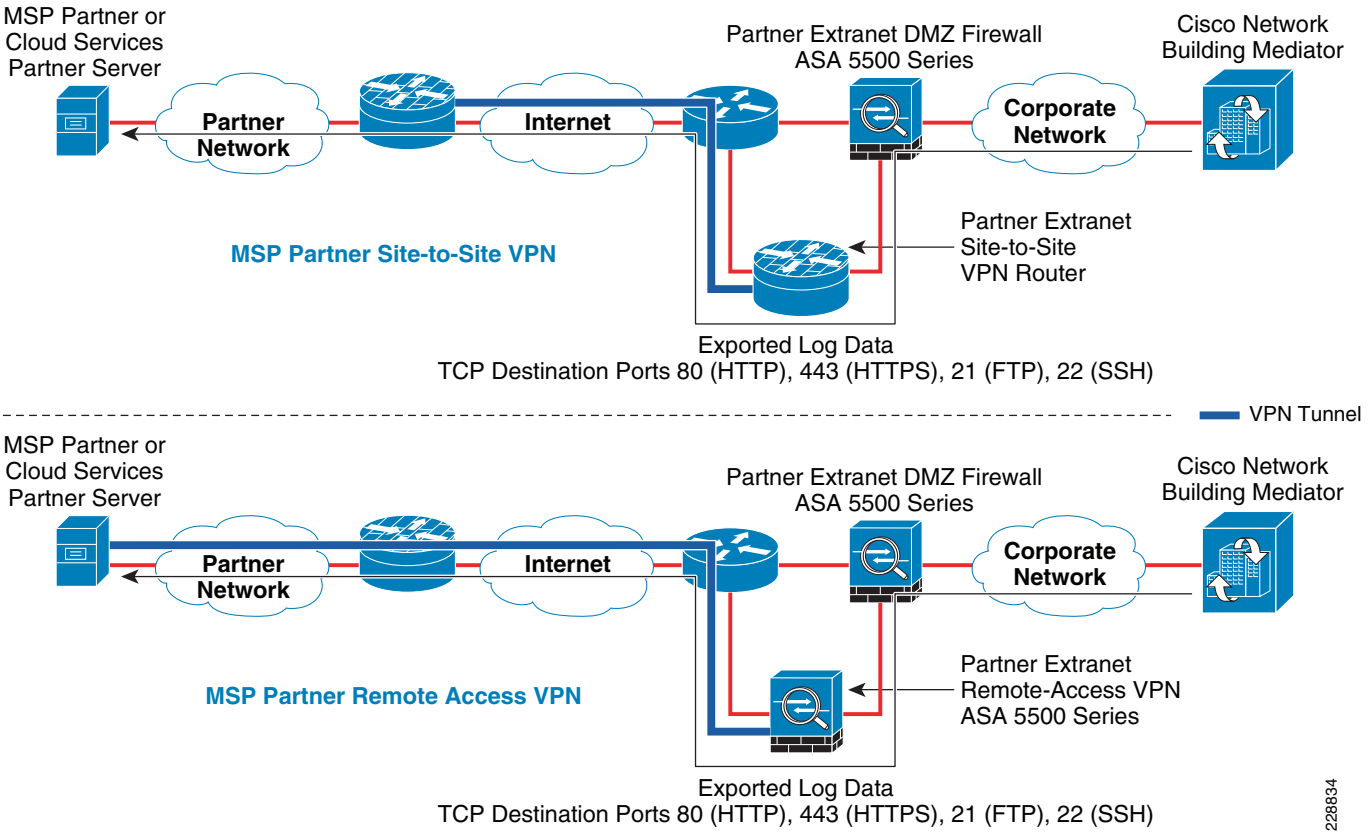

Figure 4-10 shows the export of periodic logged data initiated by the Mediators outbound toward MSP partner or cloud-services partner servers through the VPN tunnel. This scenario may be implemented when the MSP partner provides both ongoing monitoring and management of the energy management deployment, as well as providing an energy scorecard service for the enterprise customer. As discussed previously, the Mediator can periodically export logged data via HTTP, HTTPS, FTP, or Secure FTP (SSH). Typical deployments will use one of these protocols, which must be enabled across the VPN tunnel as well as through the access control of the firewalls.

When implementing a site-to-site VPN, exported log data generated from the Mediators can be used to automatically establish an IPSec tunnel, if it is not already established, with no end-user intervention. Data flows to multiple MSP partner servers (i.e., one for monitoring and management and another for energy scorecard services) is simply a matter of modifying the access-control lists that control traffic allowed across the VPN tunnels and through the various firewalls. Using a remote-access VPN for data export is somewhat more challenging. Since the Mediator cannot initiate a remote-access VPN tunnel to a MSP partner or cloud-services partner server, the remote-access VPN tunnel must already be established when the data export from the Mediator occurs. This implies the remote-access VPN tunnel is permanently established between the MSP partner or cloud-services partner server, which collects the exported log data and the enterprise network. This may require additional monitoring of the server in order to ensure the IPSec VPN tunnel is always established.

The discussion around site-to-site VPN access and remote-access VPN access in the following two sections assume a scenario in which the MSP partner is providing both ongoing monitoring and management of the Cisco Network Building Mediator deployment, as well as providing an energy scorecard service. For ease of understanding the examples, the same MSP server is providing both functions. This was done, because it demonstrates more complex site-to-site and remote-access VPN configurations. The network administrator should note that actual deployment scenarios may vary considerably. Based upon the services offered by the MSP partner, the enterprise may instead choose to export log data via the Internet (i.e., not over the VPN tunnel) to either the MSP partner or to a completely separate cloud-services partner.

Site-to-Site VPN Access

Site-to-site VPN connections are generally considered for more permanently connected requirements, meaning that a VPN tunnel is automatically established between the MSP network and the enterprise customer network based on traffic flows. This model may be useful if the MSP partner is providing ongoing monitoring and support of the energy management system, as well as collecting periodic logged data exported from the Mediators for services such as an energy scorecard. In this case, data exports initiated from the Mediators may also traverse the site-to-site VPN tunnel to the MSP partner. One advantage of this method is that the exported data and/or event data is secured by the IPSec VPN tunnel itself. Also, it is relatively easy to allow multiple MSP management workstations access through the tunnel to both manage the Mediator deployment and collect the exported data.

Site-to-site VPN connectivity can be provided with a redundant pair of Cisco 1000 Series Advanced Services Routers (ASRs), or a redundant pair of Cisco 7200 or 7300 Series routers licensed for site-to-site VPN use located within the Extranet site-to-site VPN section of the Partner Extranet Module. Example 4-2 shows an example partial configuration of a Cisco 7200 Series router (non-redundant) for site-to-site VPN connectivity.

Example 4-2 Partial Cisco 7200 Series Router Configuration for Site-to-Site VPN Access

```
!
crypto isakmp policy 10
 encr aes 256
 authentication rsa-encr
 group 2
 lifetime 180
!
!
crypto ipsec transform-set site-to-site-vpn esp-aes 256 esp-sha-hmac
!
crypto map mediator_vpn 10 ipsec-isakmp
 set peer 192.168.192.26                    ! IP address of the MSP partner VPN router.
 set transform-set site-to-site-vpn
 match address mediator_vpn_traffic         ! Controls traffic send down the VPN tunnel.
!
!
crypto key pubkey-chain rsa                 ! Public/Private key pair for authentication.
 addressed-key 192.168.192.26 encryption
  address 192.168.192.26
  key-string
   30819F30 0D06092A 864886F7 0D010101 05000381 8D003081 89028181 00E06CEC
   7B9A4D57 95D52DEC 4C55A967 2B83C85B F4F21D03 CD904B6C BECC7BB6 310115D1
   D634B04C A35F1294 886C2F40 FFEDEB34 9C515DC3 B5B7CCE5 C3E46286 950F6E30
   1983A3AE D5B03FD5 3A280657 48283817 E5384686 06DE82E9 6D49A17F 79277FAC
   C9B28AA1 55063C2E CA9F6DA1 831E7389 11FDAD02 F3A9FED7 3D83ED04 9B020301 0001
  quit
```

```
!
~
!
class-map type inspect match-all vpn-inside-tcp-81-cmap
 match access-group name mediator_mgmt_tcp_81
class-map type inspect match-all inside-vpn-ftp-cmap
 match protocol ftp
 match access-group name mediator_export_ftp_ssh_http_https
class-map type inspect match-all inside-vpn-https-cmap
 match protocol https
 match access-group name mediator_export_ftp_ssh_http_https
class-map type inspect match-all vpn-inside-https-cmap
 match protocol https
 match access-group name mediator_mgmt_web_ssh
class-map type inspect match-all inside-vpn-ssh-cmap
 match protocol ssh
 match access-group name mediator_export_ftp_ssh_http_https
class-map type inspect match-all vpn-inside-ssh-cmap
 match protocol ssh
 match access-group name mediator_mgmt_web_ssh
class-map type inspect match-ll vpn-inside-8443-cmap
 match access-group name mediator_export_https
class-map type inspect match-all inside-vpn-http-cmap
 match protocol http
 match access-group name mediator_export_ftp_ssh_http_https
class-map type inspect match-all vpn-inside-http-cmap
 match protocol http
 match access-group name mediator_mgmt_web_ssh
!
!
policy-map type inspect inside-outside-pmap  ! Policy map to inspect inside traffic going
                                             ! out the VPN tunnel.
 class type inspect inside-vpn-http-cmap
  inspect
 class type inspect inside-vpn-https-cmap
  inspect
 class type inspect inside-vpn-ftp-cmap
  inspect
 class type inspect inside-vpn-ssh-cmap
  inspect
 class type inspect vpn-inside-8443-cmap
  inspect
 class class-default
  drop
policy-map type inspect outside-inside-pmap ! Policy map to inspect VPN traffic coming in
                                              ! form the VPN tunnel.
 class type inspect vpn-inside-http-cmap
  inspect
 class type inspect vpn-inside-https-cmap
  inspect
 class type inspect vpn-inside-ssh-cmap
  inspect
 class type inspect vpn-inside-tcp-81-cmap
  inspect
 class class-default
  drop
!
zone security outside
zone security inside
zone-pair security inside-outside source inside destination outside
 service-policy type inspect inside-outside-pmap
zone-pair security outside-inside source outside destination inside
 service-policy type inspect outside-inside-pmap
!
```

```
~
!
interface GigabitEthernet0/1
 description CONNECTION TO INTERNET (VPN OUTSIDE INTERFACE)
 ip address 10.192.181.5 255.255.255.248
 ip access-group vpn_tunnel_establishment in ! Controls Internet traffic to the VPN
                                             ! router.
 zone-member security outside                ! Firewall outside zone.
 duplex auto
 speed auto
 media-type rj45
 negotiation auto
 crypto map mediator_vpn                     ! VPN tunnel originates from this interface.
!
interface GigabitEthernet0/2
 description CONNECTION TO ME-EASTFIRE-1 (VPN INSIDE)
 ip address 10.16.4.4 255.255.255.0
 zone-member security inside                 ! Firewall inside zone.
 duplex auto
 speed auto
 media-type rj45
 negotiation auto
!
~
!
ip route 10.192.2.0 255.255.255.0 10.192.181.1       ! Static route pointing to the
                                                     ! partner network.
ip route 192.168.192.24 255.255.255.252 10.192.181.1 ! Static route pointing to the
                                       ! Internet facing side of the partner VPN router.
!
~
!
ip access-list extended mediator_export_ftp_ssh_http_https
 remark Allow Mediator to Export hourly reports
 permit ip host 10.16.4.50 host 10.192.2.3
ip access-list extended mediator_export_https
 remark Allow Mediator to Export hourly reports via HTTPS
 permit tcp host 10.16.4.50 host 10.192.2.3 eq 8443
ip access-list extended mediator_mgmt_tcp_81
 remark Allow TCP 81 from partner mgmt workstations to Mediator NAT
 permit tcp host 10.192.2.3 host 10.16.4.50 eq 81
ip access-list extended mediator_mgmt_web_ssh
 remark Allow partner mgmt workstations to Mediator NAT
 permit ip host 10.192.2.3 host 10.16.4.50
ip access-list extended mediator_vpn_traffic
 remark Allow Mediator NAT to MSP managemet workstation
 permit ip host 10.16.4.50 host 10.192.2.3
ip access-list extended vpn_tunnel_establishment
 remark Allow S2S VPN from partner router
 permit esp host 192.168.192.26 host 10.192.181.5
 permit udp host 192.168.192.26 eq isakmp host 10.192.181.5
!
```

Typical site-to-site VPN connectivity uses IPSec, with AES 128-bit or higher encryption for data confidentiality and integrity. The IPSec crypto map in the example above uses AES encryption with a 256-bit key, and HMAC-SHA1 for data authentication. The crypto map also shows a self-generated public-private RSA key pair for IKE authentication. This is somewhat more secure than a shared secret, since the MSP partner VPN router must possess the private key in order to authenticate to the enterprise VPN router; without having to deploy a full public-key infrastructure (PKI) system.

Note that a classic VPN tunnel was selected for the example, as opposed to the use of virtual tunnel interface (VTI) for the VPN tunnel. Therefore, in this example, the VPN tunnel endpoint is the Internet-facing interface of the VPN router (GigabitEthernet0/1). It is recommended that the configuration should be as specific as possible, in terms of allowing access from individual MSP partner management workstations to individual Mediators (or a hierarchical mediator or Mediator Manager within the enterprise network which then allows access to individual Mediators), within the crypto ACL which controls traffic across the VPN tunnel. Access control should be specified at least down to the host level. In the example above, the access-list named "mediator_vpn_traffic" allows only traffic from IP host 10.16.4.50 to IP host 10.192.2.3 to go down the VPN tunnel. Overlapping IP address spaces between the enterprise network and the MSP partner network may often complicate site-to-site VPN deployments, requiring the deployment of NAT. IP address 10.16.4.50 is a NATed IP address of the actual Mediator sitting within the enterprise network. The NAT function could be done at the VPN router. However, for this design example NAT is done at the ASA 5500 firewall within the Extranet DMZ section of the Partner Extranet module. The crypto ACL of the VPN router could be specified down to the protocol level to allow only HTTP, HTTPS, FTP, and/or SSH traffic between the partner management workstations and individual Mediators. However, as will be discussed shortly, both the Zone-Based Policy Firewall (ZBPF) on the VPN router as well as the ASA 5500 firewall policy already restrict access down to the protocol level.

It is recommended that no active IP routing protocols be in operation between the MSP partner network and the enterprise network where possible. Instead, IP routes can be statically defined and redistributed to active routing protocols within each network. These should be restricted to only those routes necessary for the establishment of the VPN tunnel and for the partner management host subnet to reach the Mediator subnet. Note that these subnets may correspond to DMZ subnets on both the enterprise and MSP partner sides due to the use of NAT. This hides the true IP addressing of the enterprise and MSP partner networks, and simplifies the issue of overlapping IP addressing space in both the MSP partner and enterprise networks.

Access control of inbound Internet traffic to the VPN router, as well as unencrypted MSP partner traffic, can be accomplished multiple ways. Basic ACLs can be configured to restrict inbound traffic from the Internet facing interface to only IPSec and ISAKMP protocols, as is shown in the example. Technically, this may be somewhat redundant, since ZBPF is also enabled on the VPN router, and no policy exists that allows non-VPN Internet traffic from the outside security zone to the inside security zone of the VPN router. Likewise basic ACLs could be configured to restrict inbound traffic from the inside-facing interface to allow only the Mediator hosts to communicate to the MSP partner hosts. These could be specified down to the protocol level. Keep in mind, however, that FTP uses a dynamic port range for data transfer. If configuring an ACL down to the protocol level, and if FTP data export is supported over the VPN tunnel, then a range of ports may need to be opened to support the FTP exports from the Mediators. Alternatively, access control can be accomplished via ZBPF functionality on Cisco 1000 Series Advanced Services Router (ASR); or via either ZBPF functionality or Context-Based Access Control (CBAC) functionality on Cisco 7200 and 7300 Series routers.

In the example above, ZBPF functionality is enabled on the Cisco 7200 Series router, with two security zones established. The Internet-facing interface of the VPN router (GigabitEthernet0/1) is part of the outside security zone. Note that, since the VPN tunnel is established from this interface, the VPN tunnel is part of the outside security zone. The inside interface of the VPN router (GigabitEthernet0/2) is part of the inside security zone. The "inside-outside-pmap" policy map allows HTTP, HTTPS, FTP, and SSH protocols from the NATed IP address of the Mediator to the IP address of the MSP partner workstation. The "outside-inside-pmap policy" allows HTTP, HTTPS,SSH, and TCP port 81 from the IP address of the MSP partner host—reachable via the VPN tunnel—to the NATed IP address of the Mediator. These are protocols necessary for data export from the Mediator to the MSP partner server, as well as inbound management from the MSP partner server, respectively. In real deployments, a subset of protocols may be implemented, depending upon the requirements of the MSP partner. For example, the MSP partner may utilize separate IP hosts, both visible over the VPN tunnel for data export and for management. Alternatively, the data export may be sent to a separate cloud services partner on the Internet, reachable

via the ASA 5500 firewall in the Extranet DMZ section of the Partner Extranet Module. In this case, the zone-based policy firewall running on the VPN router may be configured not to allow any inbound traffic initiated from the Mediator. Finally, note that the default class within each policy map indicates that all other traffic should be dropped. The functioning of the zone-based policy firewall can be verified with the **show policy-map type inspect zone-pair** command.

Note TCP port 8443 was added in the example above to verify HTTPS functionality during the testing of the design, since the Apache Tomcat server defaults to port 8443 for HTTPS.

Besides the use of zone-based policy firewall on the VPN router, another alternative is to send the unencrypted MSP partner traffic from the site-to-site VPN router to a DMZ interface off the ASA 5500 firewall within the Extranet DMZ section of the Partner Extranet Module, as was shown in Figure 4-6 on page 4-7. This design option is in alignment with the security concept of "defense-in-depth". Alternatively, providing the stateful firewalling function within a dedicated ASA 5500 firewall and simply using ACLs on the VPN router can reduce the CPU utilization of the VPN router for additional scalability. In either case, besides a second layer of access control, this has the additional advantage in that NAT can be done at the ASA 5500 firewall, hiding the true IP addressing of the Mediators. All of the Mediators within an enterprise network appear as a series of statically NATed IP addresses off a DMZ segment within the Partner Extranet Module in this design. A partial configuration (non-redundant) from the ASA 5500 firewall for this type of design is shown in Example 4-3.

Example 4-3 Partial Configuration from the Extranet DMZ ASA 5500 Firewall

```
names
name 10.192.2.3 vc_a2_internet description Internet Address of VC-A2      ! MSP Partner
Host
name 10.16.4.50 me-westcampus-mediator-nat description NATed addresss of mediator for S2S
VPN
name 10.17.192.2 me-westcampus-mediator description West Campus Mediator
!
~
!
interface GigabitEthernet0/2
 description me-eastfire-3 g1/0/17 vlan 51
 nameif dmz
 security-level 50
 ip address 10.16.4.1 255.255.255.0
!
interface GigabitEthernet1/0
 description me-eastdist-1 g5/25 vrf bin
 nameif inside_bin
 security-level 75
 ip address 10.16.19.193 255.255.255.252
!
~
!
object-group service DM_INLINE_TCP_6 tcp     ! Inbound SSH, HTTP, & HTTPS
 port-object eq ssh
 group-object web
 port-object eq 81
object-group service DM_INLINE_TCP_3 tcp     ! Outbound FTP, SSH, HTTP, & HTTPS
 port-object eq 8443
 port-object eq ftp
 port-object eq ssh
 group-object web
object-group service web tcp
 port-object eq www
 port-object eq https
```

```
object-group network Mediators
 description Cisco Network Building Mediators
 network-object host me-westcampus-mediator
!
~
!
access-list dmz_access_in remark Allow management server access to NATed address of
Mediators.
access-list dmz_access_in extended permit tcp host vc_a2_internet host
me-westcampus-mediator-nat object-group DM_INLINE_TCP_6
access-list inside_bin_access_in remark Allow Data Export from Mediators to VC-A2 via
site-to-site VPN
access-list inside_bin_access_in extended permit tcp object-group Mediators host
vc_a2_internet object-group DM_INLINE_TCP_3
!
~
!
static (inside_bin,dmz) me-westcampus-mediator-nat me-westcampus-mediator netmask
255.255.255.255
!                                          ! Static NAT
!
~
!
access-group dmz_access_in in interface dmz
access-group inside_bin_access_in in interface inside_bin
!
~
!
route dmz 10.192.2.0 255.255.255.0 10.16.4.4 1
route inside_bin 10.17.192.0 255.255.255.248 10.16.19.194 1
!
~
!
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map global_policy
 class inspection_default
  inspect ftp                               ! Application-layer inspection of FTP traffic
  inspect h323 h225
  inspect h323 ras
  inspect netbios
  inspect rsh
  inspect rtsp
  inspect skinny
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect tftp
  inspect sip
  inspect xdmcp
  inspect http
  inspect dns
!
service-policy global_policy global
!
```

Access control configured on the ASA 5500 firewall again allows only HTTP, HTTPS, SSH, and TCP port 81 traffic from the MSP partner host to access the Mediator via the statically NATed IP address. Likewise, only HTTP, HTTPS, SSH, and FTP traffic are allowed from the Mediator to the MSP partner host. Note that TCP port 8443 was allowed to validate HTTPS functionality since Apache Tomcat defaults to TCP port 8443 for HTTPS. This provides a second line of access control on top of VPN

router, as well as a single point of entry for partner data flows into the enterprise network. Since partner connectivity requirements are typically much more well-defined than general employee access to the Internet, the Extranet DMZ firewall may be locked down much tighter with both inbound and outbound access control, versus allowing all connectivity outbound, as is often done with the Internet Edge Module firewall. This is another advantage of deploying a separate firewall within a Partner Extranet Module. Finally, note that application-layer inspection on the ASA 5500 firewall will dynamically open the port range needed for FTP transfers, versus opening a range of ports on a static ACL.

Remote-Access VPN Access

Remote-access VPN connections are generally considered for more temporarily connected requirements, although such VPN connections may be left up for extended periods of time as well, based on the particular business requirements of the energy management solution. With remote-access VPN connectivity, individual MSP partner workstations establish tunnels to a VPN concentrator in order to manage the Mediator deployment. If exporting of logged data is done through remote-access VPN connectivity, individual MSP partner or cloud service partner servers also need to establish tunnels to a VPN concentrator in order to provide a path for the data export from the Mediators to reach them.

Remote-access VPN connectivity can be provided with a redundant pair of ASA 5500 Series Security Appliances licensed for remote-access VPN use deployed within the Extranet remote-access VPN section of the Partner Extranet Module. Typical MSP remote-access VPN connections also uses IPSec with AES 128-bit or higher encryption for data confidentiality and integrity. This may require client software, such as the Cisco VPN Client, to be deployed on the MSP partner and/or cloud services partner workstations and servers. Example 4-4 shows a partial example configuration (non-redundant) of an ASA 5500 for remote-access VPN connectivity.

Example 4-4 Example of a Partial ASA 5500 Series Configuration for Remote-Access VPN Connectivity

```
!
names
name 10.192.181.6 user_inet description Firewall Outside Internet Interface
name 10.17.192.2 me-westcampus-mediator description West Campus Mediator
name 10.17.2.10 me-westserv-2 description Radius Server
name 10.16.19.2 vc-a2-vpn_alternate description MSP partner Server RAVPN IP Address
name 10.16.19.193 inside_bin_interface description Firewall Building Information Network
Interface
!
interface GigabitEthernet0/3
 description OC3 Internet Access via me-eastinet-3 & 4
 nameif outside
 security-level 25
 ip address user_inet 255.255.255.248
!
interface GigabitEthernet1/0
 description BIN VRF via me-eastdist-1 g5/25
 nameif inside_bin
 security-level 75
 ip address inside_bin_interface 255.255.255.252
!
!
object-group service web tcp
 port-object eq www
 port-object eq https
!
object-group network Mediators
 description Cisco Network Building Mediators
 network-object host me-westcampus-mediator
```

```
!
object-group service DM_INLINE_TCP_5 tcp
 port-object eq 8443
 port-object eq ftp
 port-object eq ssh
 group-object web
object-group service DM_INLINE_TCP_7 tcp
 port-object eq 8443
 port-object eq ftp
 port-object eq ssh
 group-object web
object-group service DM_INLINE_TCP_8 tcp
 port-object eq 8443
 port-object eq ftp
 port-object eq ssh
 group-object web
object-group service DM_INLINE_TCP_10 tcp
 port-object eq ftp
 port-object eq ssh
 group-object web
 port-object eq 81
object-group service DM_INLINE_TCP_11 tcp
 port-object eq ftp
 port-object eq ssh
 group-object web
 port-object eq 81
object-group network msp-ravpn-server-group
 description MSP mgmt workstation group visible via RAVPN
 network-object host vc-a2-vpn_alternate
!
access-list remote_access_user extended permit tcp object-group msp-ravpn-server-group
object-group DM_INLINE_TCP_8 object-group Mediators
access-list remote_access_user extended permit tcp object-group msp-ravpn-server-group
object-group Mediators object-group DM_INLINE_TCP_11
!
access-list remote_access_group extended permit tcp object-group msp-ravpn-server-group
object-group DM_INLINE_TCP_5 object-group Mediators
access-list remote_access_group extended permit tcp object-group msp-ravpn-server-group
object-group Mediators object-group DM_INLINE_TCP_10
!
access-list inside_bin_access_in remark Allow Data Export from Mediators to VC-A2 via
remote-access VPN
access-list inside_bin_access_in extended permit tcp object-group Mediators object-group
msp-ravpn-server-group object-group DM_INLINE_TCP_7
!
access-list inside_bin_nat0_outbound remark Allow Mediators to reach assigned IP address
for RA VPN.
access-list inside_bin_nat0_outbound extended permit ip object-group Mediators
object-group msp-ravpn-server-group
!
access-list me-nets remark east coast
access-list me-nets standard permit 10.16.0.0 255.255.0.0
access-list me-nets remark west coast
access-list me-nets standard permit 10.17.0.0 255.255.0.0
!
access-list splitTunFWIn remark Roland's office
access-list splitTunFWIn extended permit ip 64.100.160.0 255.255.255.0 any
!
ip local pool mediator_vpn 10.16.19.4-10.16.19.14 mask 255.255.255.240
!                                        ! IP address pool (not used in this example)
!
nat-control
nat (inside_bin) 0 access-list inside_bin_nat0_outbound
```

```
!                                              !Don't NAT access to the assigned address of
the RA VPN devices
nat (inside_bin) 1 0.0.0.0 0.0.0.0
!
access-group inside_bin_access_in in interface inside_bin
!                                              ! Allow Mediator export to RA VPN devices
using
!                                              ! SSH, FTP, and HTTPS
!
route outside 0.0.0.0 0.0.0.0 10.192.181.1 1
route inside_bin 10.17.192.0 255.255.255.248 10.16.19.194 1
!
dynamic-access-policy-record IPsec_RAVPN_Dynamic_Policy
 description "Dynamic Access Policy for RAVPN Clients"
 network-acl remote_access_user                ! Specifies per-user RA VPN access control
!
aaa-server radius-server protocol radius      ! Specifies the RADIUS server
aaa-server radius-server (inside) host me-westserv-2
 key ciscoese
 authentication-port 1812
 accounting-port 1813
 radius-common-pw ciscoese
!
crypto ipsec transform-set ESP-AES-256-MD5 esp-aes-256 esp-md5-hmac
crypto ipsec transform-set ESP-DES-SHA esp-des esp-sha-hmac
crypto ipsec transform-set ESP-3DES-SHA esp-3des esp-sha-hmac
crypto ipsec transform-set ESP-DES-MD5 esp-des esp-md5-hmac
crypto ipsec transform-set ESP-AES-192-MD5 esp-aes-192 esp-md5-hmac
crypto ipsec transform-set ESP-3DES-MD5 esp-3des esp-md5-hmac
crypto ipsec transform-set ESP-AES-256-SHA esp-aes-256 esp-sha-hmac
crypto ipsec transform-set ESP-AES-128-SHA esp-aes esp-sha-hmac
crypto ipsec transform-set ESP-AES-192-SHA esp-aes-192 esp-sha-hmac
crypto ipsec transform-set ESP-AES-128-MD5 esp-aes esp-md5-hmac
crypto ipsec security-association lifetime seconds 28800
crypto ipsec security-association lifetime kilobytes 4608000
crypto ipsec df-bit clear-df inside
crypto ipsec df-bit clear-df mgmt
crypto dynamic-map SYSTEM_DEFAULT_CRYPTO_MAP 65535 set transform-set ESP-AES-128-SHA
ESP-AES-128-MD5 ESP-AES-192-SHA ESP-AES-192-MD5 ESP-AES-256-SHA ESP-AES-256-MD5
ESP-3DES-SHA ESP-3DES-MD5 ESP-DES-SHA ESP-DES-MD5
crypto dynamic-map SYSTEM_DEFAULT_CRYPTO_MAP 65535 set security-association lifetime
seconds 28800
crypto dynamic-map SYSTEM_DEFAULT_CRYPTO_MAP 65535 set security-association lifetime
kilobytes 4608000
crypto map user_inet_map 65535 ipsec-isakmp dynamic SYSTEM_DEFAULT_CRYPTO_MAP
crypto map user_inet_map interface outside  ! Crypto map applied to the outside interface
crypto isakmp enable outside                ! IKE enabled on outside interface
crypto isakmp policy 5
 authentication pre-share
 encryption 3des
 hash sha
 group 2
 lifetime 86400
crypto isakmp policy 10
 authentication pre-share
 encryption des
 hash sha
 group 2
 lifetime 86400
!
group-policy DfltGrpPolicy attributes      ! Default RA VPN policy group
 banner value Unauthorized Access is Prohibited.
 vpn-filter value remote_access_group      ! Defines access control at the group level
 vpn-tunnel-protocol IPSec webvpn
```

```
 ipsec-udp enable
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value me-nets   ! Defines ACL for split tunneling at the client
 secure-unit-authentication enable
 user-authentication enable
 address-pools value mediator_vpn          ! Assigns IP address pool to the default group
policy
 client-firewall opt cisco-integrated acl-in splitTunFWIn acl-out open
!                                          !Defines client firewall ACL at the client
 webvpn
  url-list value Mediator
  svc compression none
  svc ask none default webvpn
  http-comp none
  smart-tunnel enable Mediator_Configuration
  url-entry disable
!
group-policy IPsec_VPN_Group_Policy_2 internal
!                                          ! RA VPN policy group using RADIUS user
authentication
group-policy IPsec_VPN_Group_Policy_2 attributes
 vpn-tunnel-protocol IPSec                 ! Allows only IPsec VPN
 group-lock value vc-a2-acs2
 address-pools none                        ! Overrides default group policy use of IP
address pool
!
group-policy IPsec_VPN_Group_Policy internal
           ! RA VPN policy group using local userid authentication
group-policy IPsec_VPN_Group_Policy attributes
 vpn-tunnel-protocol IPSec                 ! Allows only IPsec VPN
 group-lock value vc-a2
 address-pools none                        ! Overrides default group policy use of IP
address pool
!
username joeuser1 password aSnrkw/Y3WIwmRvT encrypted
!                                          ! Local userid definition (note that the userid
is the same as
username joeuser1 attributes               ! the IPsec connection profile in this example
configuration)
 vpn-group-policy IPsec_VPN_Group_Policy
 vpn-framed-ip-address 10.16.19.2 255.255.255.240
           ! Local assignment of IP address specific to userid
 service-type remote-access
!
tunnel-group vc-a2 type remote-access     ! IPsec connection profile for local
authentication
tunnel-group vc-a2 general-attributes
 authorization-server-group LOCAL
 default-group-policy IPsec_VPN_Group_Policy
tunnel-group vc-a2 ipsec-attributes
 pre-shared-key *
!
tunnel-group vc-a2-acs2 type remote-access ! RA VPN policy group using RADIUS
authentication
tunnel-group vc-a2-acs2 general-attributes
 authentication-server-group radius-server
 authorization-server-group radius-server
 default-group-policy IPsec_VPN_Group_Policy_2
tunnel-group vc-a2-acs2 ipsec-attributes
 pre-shared-key *
!
```

One advantage of remote-access VPN connectivity is that access control can be provisioned on a per group and/or per user basis. The configuration in Example 4-4 above shows two different remote access VPN group policies called “IPsec_VPN_Group_Policy” and “IPsec_VPN_Group_Policy2”, for illustrative purposes. Both group policies inherit some attributes from the default VPN group policy named “DfltGrpPolicy”, and also override other attributes. This hierarchical grouping allows the network administrator to define the default policy more generically, and then restrict further with policies specific to individual groups, such as a group policy specifically defined for MSP partner remote access VPN connections. In the configuration above, the default VPN group policy specifies the access-control list named “remote_access_group” as a group level filter. Both VPN group policies shown in the configuration above utilize this filter. However, both VPN group policies shown in the configuration above override their VPN tunnel protocols to use only IPSec, instead of IPSec and SSL. Likewise, both VPN group policies override the assignment of the IP address of the client. Instead of using IP address from the pool “mediator_vpn”, the IP address for the groups is assigned per user.

The Cisco Network Building Mediator periodically exports logged data to a destination based either on a hostname or an IP address. There are some inherent delays in propagating dynamic DNS updates around a large enterprise network, as well as potential issues regarding the caching of DNS entries within the Mediator itself. For this reason, if the periodic export of logged data from the Mediators is through a remote-access VPN connection, then a static IP address should be handed to the MSP partner or cloud services partner server when it establishes the remote-access VPN tunnel. Identification of the server can be done based upon the userid when the remote-access VPN tunnel is established. In other words, the assignment of the IP address can be based upon the userid. An example has been provided regarding how the IP address is specified per userid. Within the configuration in Example 4-4 above, the IPSec connection profile (tunnel-goup) called “vc-a2” has been defined to use the policy group named “IPsec_VPN_Group_Policy”. Tunnel-group “vc-a2” specifies the use of the local database for authentication of individual users. A userid of “joeuser1” has also been defined within the configuration example. Under the attributes of the username, the **vpn-framed-ip-address** command is used to statically assign an IP address to the particular userid.

However, the configuration of userids directly on the ASA 5500 Series Security Appliance is generally not considered a best practice. For any sizeable enterprise deployment, a best practice is to handle the access control decision centrally through a AAA server connected to the ASA 5500 Series Security Appliance via either the RADIUS or TACACS+ protocol. The AAA server may in turn be connect to a backend directory server. An example of a AAA server is the Cisco Secure Access Control Server (ACS), which can be deployed either within a Data Center or Campus Service Module within the campus network. Centralized control allows the network administrator to more effectively maintain the list of partner userids who have access to the enterprise network. It also allows the network administrator to add or remove partner userids quickly, without having to potentially touch multiple local databases in multiple network devices. It is essential that network administrator work closely with the MSP partner to immediately identify any employees who leave the company, so that their access to the enterprise network can be immediately revoked. Alternatives to the use of individual passwords include the use of token cards or token software installed on the MSP partner workstation. This requires the MSP employee either to have physical access to the PC or physical access to the token card in order to access the enterprise network. Within the configuration in Example 4-4, a second IPSec connection profile (tunnel-group) called “vc-a2-acs” has been defined to use the policy group named “IPsec_VPN_Group_Policy2”. Tunnel-group “vc-a2-acs” specifies the use of a RADIUS server group for authentication of individual users. Within the RADIUS server itself, specific userids are defined. The *Framed-IP-Address* and *Framed-IP-Netmask* RADIUS attributes defined for each userid can be used in order to return a specific IP address, based upon the particular userid. Per-user filtering can be accomplished through the application of a dynamic access policy which is configured to key in on a specific AAA attribute such as the userid. In the configuration above, the access-list named “remote_access_group” is used applied based on userid.

The network administrator should take note that, if the export of periodic logged data from the Mediator does not use the remote-access VPN tunnel, then assignment of IP addresses to MSP partner VPN clients can be based on an address pool. In such cases, the access-control lists need to be modified to allow for the required protocols across the range of IP addresses. Finally, as with site-to-site VPN connections, the MSP data flows which terminate on the remote-access VPN ASA 5500s can also be routed to a separate segment off of the Extranet DMZ firewalls, as shown in Figure 4-8 on page 4-10. This provides a second layer of access control and provides a single point of entry for partner traffic into the enterprise network.

Note The use of SSL VPN technology to access the Cisco Network Building Mediator has not been thoroughly tested by ESE as of the time this document was written; and has therefore not been validated to work entirely correctly.

Internet Access for Exporting of Logged Data

As mentioned previously, the Extranet Module can be used to as an alternative means of providing Internet connectivity through which data is exported from the Mediators to cloud services partners. The first method of doing so is through the use of secure proxy or drop-and-forward servers located on the Extranet DMZ section of the Extranet Module. An example of this is shown in Figure 4-11.

Figure 4-11 Example Proxy Server Deployment for Mediator Log File Export via the Extranet Module

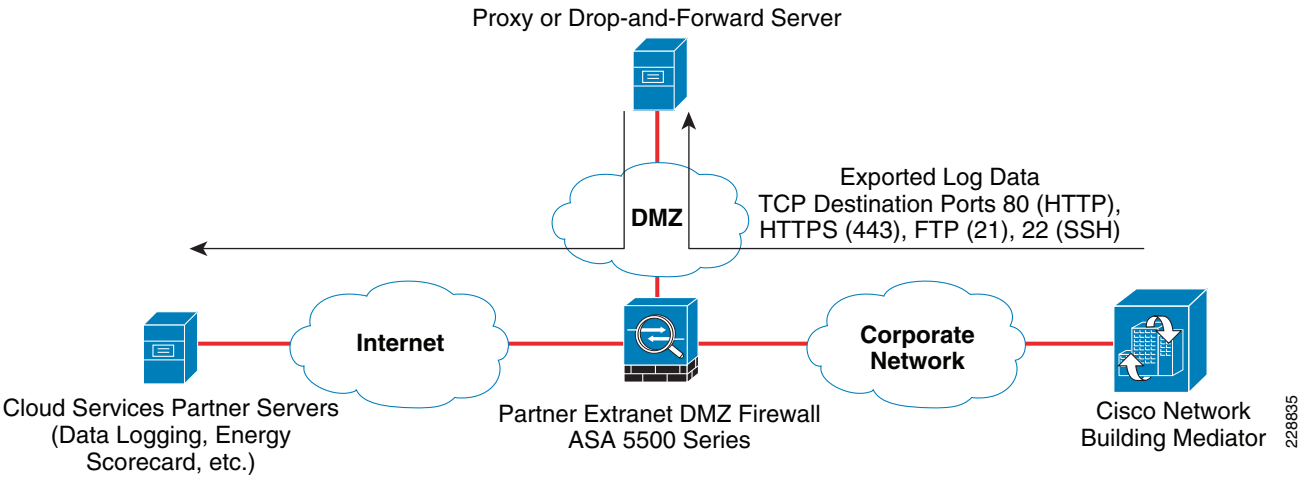

Note that Figure 4-11 is nearly identical to Figure 4-4. The only difference is that the Mediator exported logged-data passes through a secure proxy or drop-and-forward server located within the Extranet DMZ section of the Partner Extranet Module. The Partner Extranet DMZ firewall shown in Figure 4-11 is a firewall dedicated for partner traffic, separate from the firewall dedicated employee traffic deployed within the Internet Edge module. Typically, the default route for the overall enterprise network out to the Internet—therefore, to cloud-services partner servers which are reachable via the Internet—is through the Internet edge firewall. However, as discussed previously, the Mediator can be configured to use a proxy server in order to export log files via protocols such as FTP. In such cases, the Mediator only needs to be able to route to the IP address of the proxy server. The proxy server then uses a separate default route to the Internet via the Extranet DMZ firewall.

Again, the benefit of using a proxy or drop-and-forward server is that it decouples the sessions between the Mediators and the DMZ server from the sessions between the cloud services partner servers and the DMZ server. It should be noted, however, that the use of either a proxy or a drop-and-forward server adds additional administrative overhead and another possible point of failure. Therefore the additional security gained from such implementations should be weighed against the administrative overhead incurred.

A second method of using the Extranet Module as an alternative means of providing Internet connectivity through which data is exported from the Mediators to cloud services partners is through the use of virtual routing and forwarding (VRF) technology. This is discussed further in the next section.

Extending VRFs to the Partner Extranet Module

When deploying network virtualization for the energy management solution, the VRF which supports the Mediators needs to be extended through the Campus Module out to the Partner Extranet Module. Figure 4-12 shows a modified version of the Partner Extranet Module which supports virtualization for the energy management solution.

Figure 4-12 Example Redundant Partner Extranet Module Design with VRFs

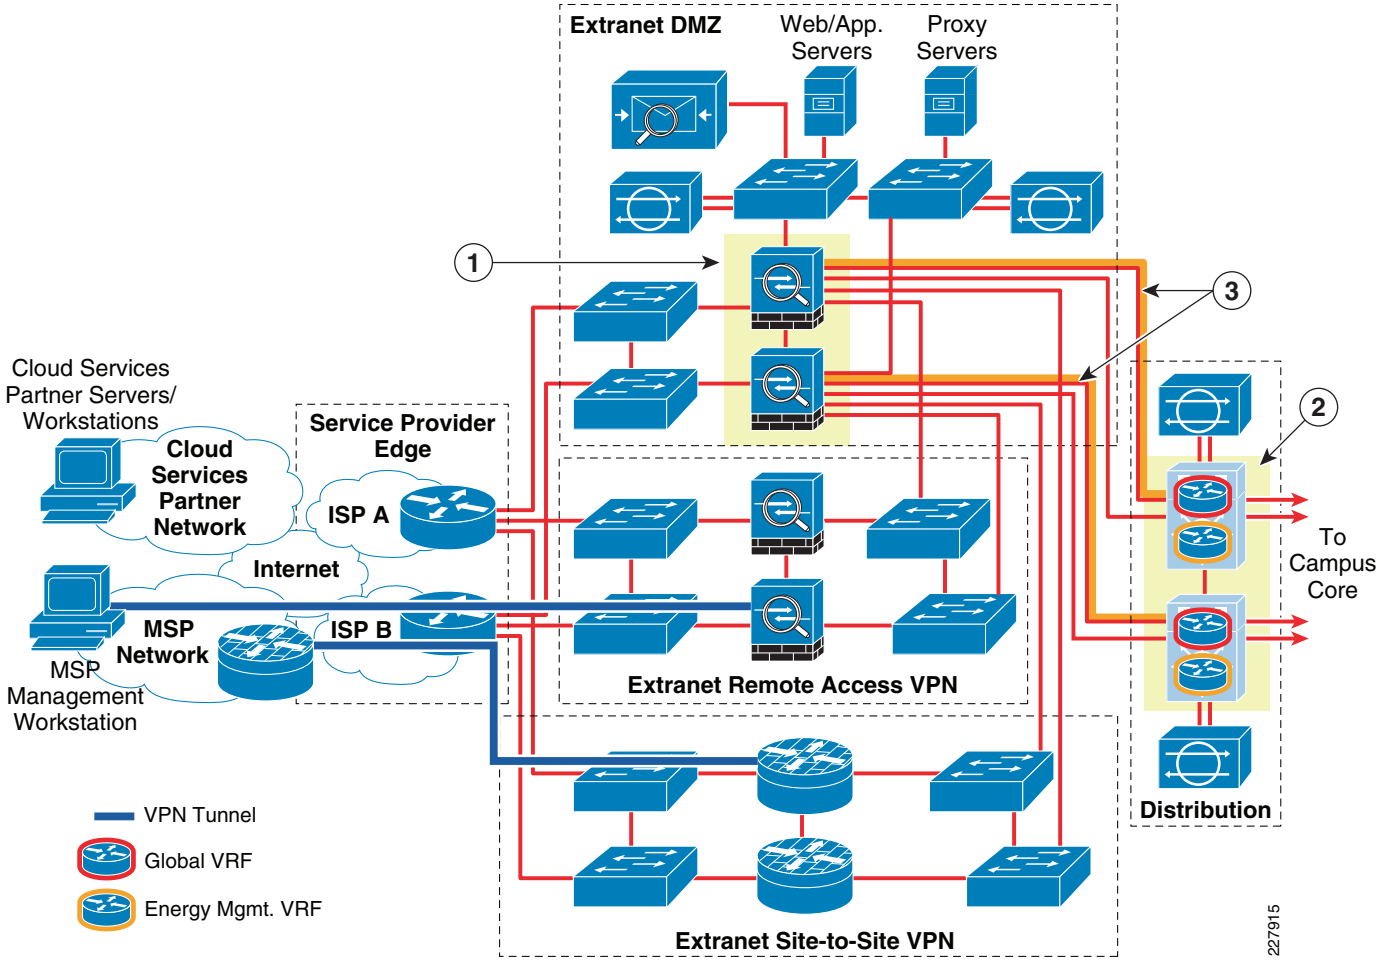

The following describes the numbers in Figure 4-12:

- **1**—ASA 5500 Security Appliances deployed within the Extranet DMZ section provides address translation and stateful access control for outgoing connections to cloud-services partners and incoming connections from MSP partners through VPN devices.
- **2**—ASA 5500 Security Appliances deployed within the Extranet remote-access VPN section provides remote-access VPN termination, address translation, and IP address assignment for MSP partners.
- **3**—Cisco 1000 Series ASRs or Cisco 7200/7300 Series routers deployed within the Extranet site-to-site VPN section provides site-to-site VPN termination, stateful access control, and potentially address translation of managed services partner VPN connections.
- **4**—Traffic destined for the energy management VRF is routed to separate interfaces on the Extranet DMZ Firewalls.

In this example, the Catalyst 6500 switches within the distribution section of the Partner Extranet Module are configured to support a separate virtual routing and forwarding (VRF) instance for the energy management solution. This is then mapped to GRE tunnels which extend the energy management solution VRF back to either a Campus Service Module or Data Center Service Module. From there, other GRE tunnels extend out to campus buildings or branch locations which house the individual Mediators. This provides a star configuration, versus provisioning separate GRE tunnels from the Partner Extranet Module to each location which houses a Mediator. Note that this configuration also facilitates a hierarchical Mediator deployment, where access to the remote mediators is accomplished via a hierarchical Mediator (or future Mediator Manager appliance) located within a Campus Service Module or Datacenter Service Module. A partial configuration example of a Catalyst 6500 switch within the Distribution section of the Partner Extranet Module is shown in Example 4-5.

Example 4-5 Partial Configuration Example of VRFs Extended to the Partner Extranet Module

```
!
ip vrf bin                                    ! Configures the Building Information Network
(BIN) VRF.
rd 251:127
!
~
!
interface Tunnel0                             ! GRE tunnel to the Datacenter Module for the
BIN VRF.
 description VRF FOR MEDIATOR NETWORK TO ME-W-DCSERV-1
 ip vrf forwarding bin
 ip address 10.17.192.42 255.255.255.248
 tunnel source Loopback0
 tunnel destination 10.17.252.2
!
interface Loopback0
 ip address 10.16.255.40 255.255.255.255
!
interface Loopback2
 description LOOPBACK INTERFACE FOR TUNNEL TO ME-W-DCSERV-1
 ip vrf forwarding bin
 ip address 10.16.255.40 255.255.255.255
!
~
!
interface GigabitEthernet1/25
! Interface which connects to the ASA 5500 firewall.
 description ASA5550 G1/0 mediator vpn
 switchport
 switchport access vlan 192
 switchport mode access
 load-interval 30
```

```
 wrr-queue bandwidth 5 25 40
 priority-queue queue-limit 30
 wrr-queue queue-limit 5 25 40
 wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 3 60 70 80 90 100 100 100 100
 wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 3 70 80 90 100 100 100 100 100
 wrr-queue cos-map 1 1 1
 wrr-queue cos-map 2 1 0
 wrr-queue cos-map 3 1 2
 wrr-queue cos-map 3 2 3
 wrr-queue cos-map 3 3 6
 wrr-queue cos-map 3 4 7
 priority-queue cos-map 1 4 5
 mls qos trust dscp
!
~
!
interface Vlan192
! SVI for the BIN VRF is part of an HSRP group
 ip vrf forwarding bin
! for high availability.
 ip address 10.16.19.195 255.255.255.240
 standby 192 ip 10.16.19.194
 standby 192 priority 110
 standby 192 track GigabitEthernet1/25 20
!
~
!
router eigrp 110
 network 10.16.0.0 0.0.255.255
 eigrp router-id 10.16.255.21
 no auto-summary
 passive-interface default
 no passive-interface GigabitEthernet1/1
 no passive-interface GigabitEthernet1/46
 no passive-interface GigabitEthernet3/0/0
 no passive-interface GigabitEthernet3/0/1
 no passive-interface TenGigabitEthernet6/1
 no passive-interface TenGigabitEthernet6/2
 redistribute static metric 1000000 100 255 1 1500 route-map FW_route_inject
 redistribute bgp 64000 metric 100000 150 255 1 15000
 !
 address-family ipv4 vrf bin                ! EIGRP autonomous system 99 provides routing
for
  autonomous-system 99                       ! the BIN VRF.
  network 10.16.19.0 0.0.0.255
  network 10.16.255.0 0.0.0.255
  network 10.17.192.0 0.0.0.255
  no auto-summary
  passive-interface Vlan192
  passive-interface Loopback2
  redistribute static metric 12100 250 255 1 1500
 exit-address-family
!
ip route vrf bin 0.0.0.0 0.0.0.0 10.16.19.193  ! Static default route for the
!                                               ! BIN VRF points to the ASA firewall.
```

In Example 4-5, the default route to the Internet for the Building Infrastructure Network (BIN) VRF is pointed toward the Extranet DMZ firewall, instead of the Internet Edge firewall. Virtual firewall contexts are also not required within the Extranet DMZ firewalls to support this configuration. Instead

a separate physical interface on the Extranet DMZ firewalls can be provisioned, dedicated for the energy management solution VRF, as long as there is no overlapping IP address spaces between the global VRF and the energy management solution VRF.

Collapsed Internet Edge Designs

Although large organizations often deploy separate Internet Edge and Partner Extranet Modules with separate components for each function (firewall, remote access VPN, site-to-site VPN, etc.), smaller organizations sometimes collapse partner connectivity and employee Internet connectivity into a single module. The benefit of this design is reduced capital expenditures for networking equipment. However the disadvantage of this design is that there is no longer a clean separation of the MSP partner traffic from employee traffic into and out of the enterprise network. Also combining multiple functions into a single device increases the operational complexity of the device and reduces the overall scalability of the solution. However, for smaller organizations the trade-off is often acceptable. Figure 4-13 shows an example of a collapsed Internet edge design, as it applies to the energy management solution.

Figure 4-13 Example of a Collapsed Internet Edge Design

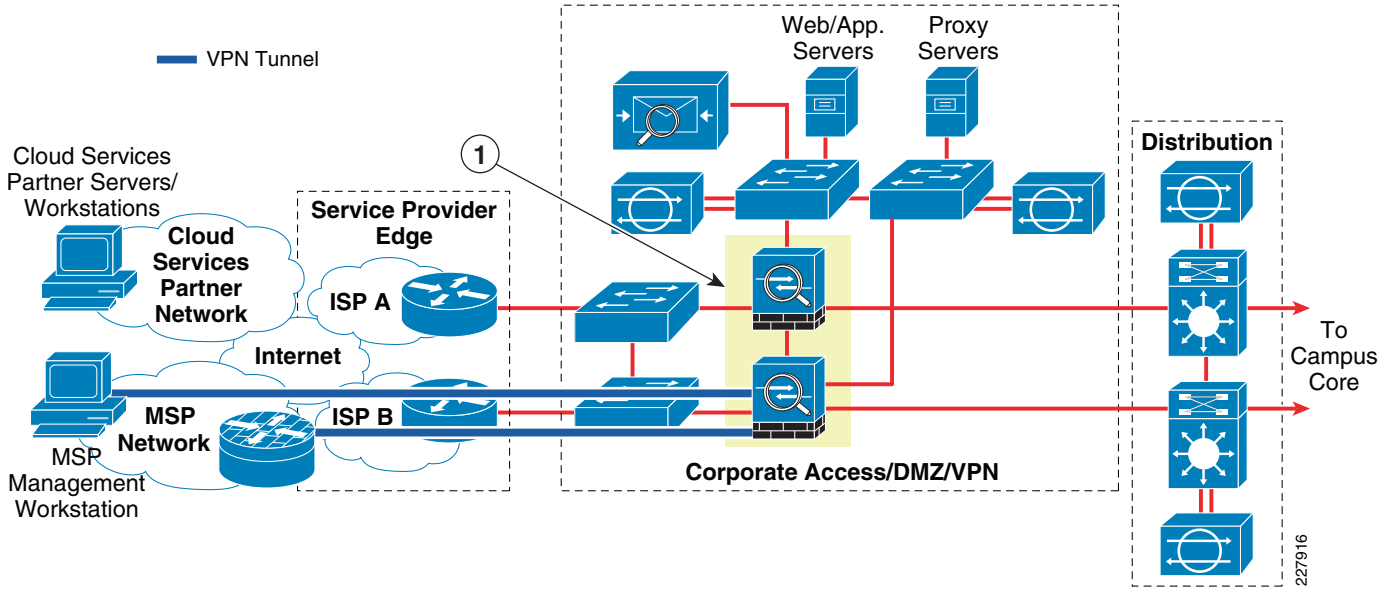

The following describes the number in Figure 4-13.

- **1**—ASA 5500 Security Appliances deployed within the corporate access/DMZ/ VPN section provides address translation and stateful access control for outgoing connections to cloud services partners. ASA 5500 Security Appliances also provide site-to-site and/or remote access VPN termination, address translation, IP address assignment, and stateful access control to MSP partner VPN connections.

With this design, a single pair of ASA 5500 Series Security Appliances can provide both site-to-site and remote access VPN connectivity for MSP partner and employee access, as well as stateful firewalling for Internet connectivity. As with the separate Partner Extranet Module design, the periodic export of data logs from the Mediators can be sent directly to cloud-services partner servers accessible through the Internet or sent via either a proxy or drop-and-forward server located on the DMZ. A separate energy

management VRF can also be extended to the ASA 5500 Series Security Appliances via a separate physical interface in order to provide path isolation for traffic to and from the Cisco Network Building Mediators, as is shown in Figure 4-14.

Figure 4-14 Example of a Collapsed Internet Edge Design with VRF

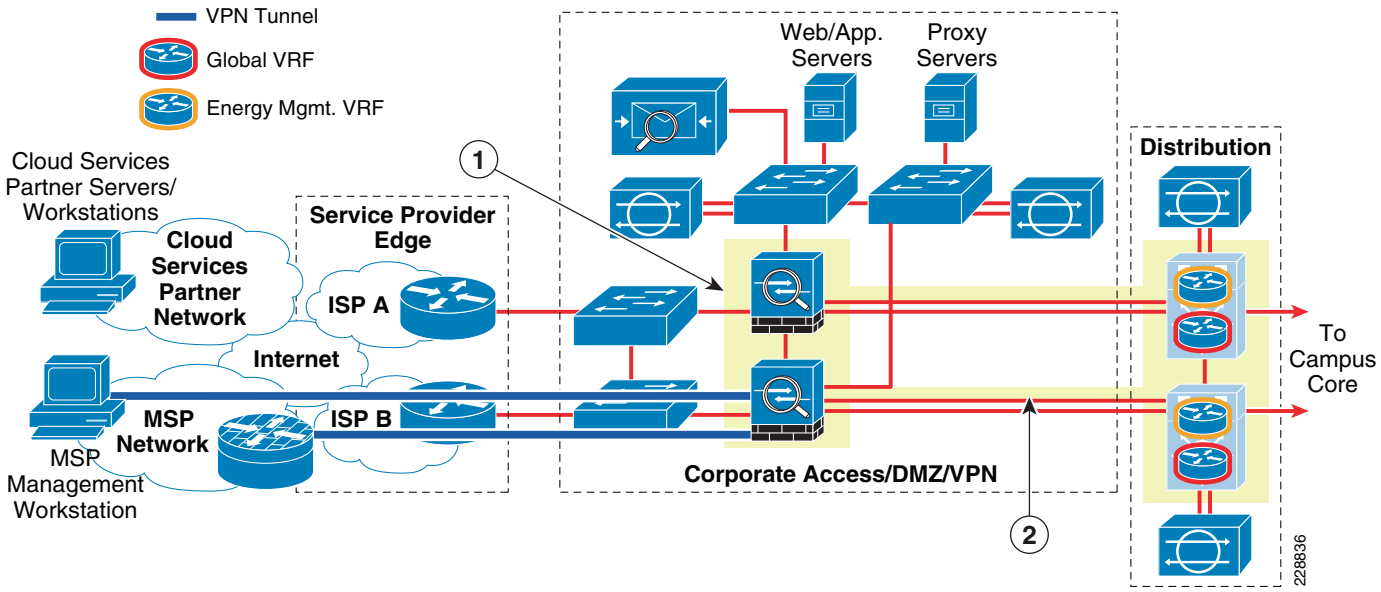

The following describes the numbers in Figure 4-14:

- **1**—ASA 5500 Security Appliances deployed within the corporate access/DMZ/VPN section provides address translation and stateful access control for outgoing connections to cloud services partners. ASA 5500 Security Appliances also provide site-to-site and/or remote access VPN termination, address translation, IP address assignment, and stateful access control to MSP partner VPN connections.
- **2**—Energy management VRF extended from the Catalyst 6500 switch within the distribution section of the collapsed Internet edge to a separate interface on the ASA 5500 Security Appliance.

CHAPTER 5

Data Center/Campus Service Module Design Considerations

Enterprise customers often handle the day-to-day operations and management of building system networks. In terms of the energy management solution, the function of the Data Center/Campus Service Module is to provide a centralized point of administration and operations, referred to as an Energy Management Operations Center (EMOC) within this document. This allows the enterprise organization to centrally manage the Mediators deployed both within the campus as well as within the branches. The components of the energy management solution that may reside within the EMOC are as follows:

- *Enterprise Energy Management Workstations*—These provide the ability to configure the Mediators via the configTOOL management application; to create and deploy graphical-based logical control applications to the Mediators via the perfectHOST application; and to monitor and manage aspects of the Mediators such as set points, events and alarms, and websites created and deployed on the Mediators, via the OMEGA suite of applications.
- *Enterprise Archiving and Energy Scorecard Servers*—These optional servers may be deployed in order to collect and archive periodically exported datapoint information from the Mediators. This information may then be used to provide an internal energy scorecard and/or historical energy usage information for both energy management personnel and business units within the enterprise organization. In business scenarios where the enterprise organization has outsourced this function to a cloud service provider, these servers may not be deployed. Note also that separate servers may be used for the datapoint archiving function and the energy scorecard function.
- *Hierarchical Mediator/Mediator Manager*—The hierarchical Mediator is an optional component which provides the ability to centrally collect and display datapoints on remote Mediators via the use of aliases and the Remote Node Abstraction (RNA) protocol. Note that any Mediator can share datapoint information with another mediator in a peer-to-peer manner. The hierarchical Mediator design is simply a design option in which one or more Mediators functions as an aggregation point for downstream Mediators. The Cisco Network Building Mediator Manager (Mediator Manager) is a future product offering, which will scale the overall energy management solution deployment by offloading the hierarchical Mediator function to a server appliance.

Note This revision of the design guide focuses primarily on hierarchical Mediator designs. Future revisions will include further discussion of the Cisco Network Building Mediator Manager.

Hierarchical Mediator Designs

A hierarchical Mediator design allows datapoints from building devices connected to remote Mediators to be defined using aliases on a hierarchical Mediator centrally located. The Mediators use the Remote Node Abstraction (RNA) protocol to share datapoint information between themselves. To the hierarchical Mediator, the datapoints appear as if they were coming from devices directly attached to it. The hierarchical mediator model is designed to provide a single point of access, via websites deployed on the hierarchical mediator, to the datapoints for the Mediators deployed within the enterprise organization; versus having to access each Mediator individually. An example of this is shown in Figure 5-1.

Figure 5-1 Mediator Access With and Without Hierarchical Mediator Functionality

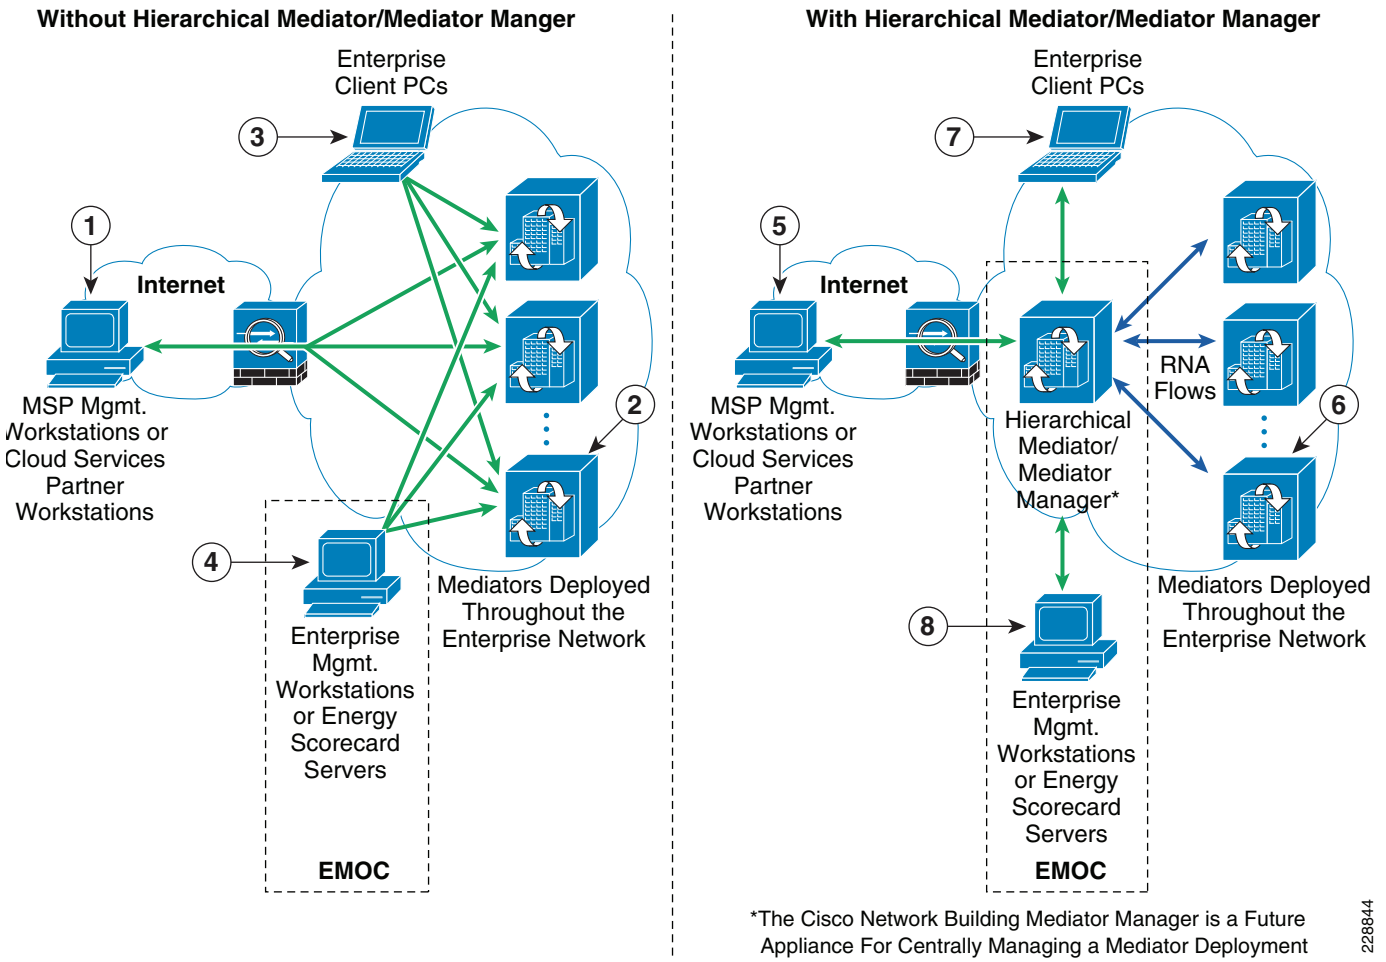

The following describes the numbers in Figure 5-1:

Without Hierarchical Mediator

- **1**—Managed service provider workstations need access to each Mediator in order to monitor and manage the deployment.

- **2**—Mediators individually export periodically logged datapoint information to cloud-services partner servers and/or enterprise energy scorecard servers. Event data may be exported by each Mediator individually, viewed and acknowledged on each Mediator individually, or viewed and acknowledged centrally on one Mediator through a cloud formation.
- **3**—Individual client PCs may need access to individual Mediators deployed throughout the enterprise organization, in order to view collected datapoint information via websites created and deployed on the Mediators.
- **4**—Enterprise energy management workstations need access to each Mediator in order to monitor and manage the deployment.

With Hierarchical Mediator

- **5**—Managed service provider workstations may need access to the hierarchical Mediator in order to view datapoint information (aggregated from remote Mediators through the RNA protocol) within websites deployed on the hierarchical Mediator. Depending on management requirements, individual access to each Mediator may still be required.
- **6**—Mediators individually export periodically logged datapoint information to cloud-services partner servers and/or enterprise energy scorecard servers. Event data may be exported by each Mediator individually, viewed and acknowledged on each Mediator individually, or viewed and acknowledged centrally on the hierarchical Mediator through a cloud formation.
- **7**—Individual enterprise client PCs may need access to the hierarchical Mediator in order to view datapoint information within websites created and deployed on the hierarchical Mediator.
- **8**—Enterprise energy management workstations may need access to the hierarchical Mediator in order to view datapoint information within websites deployed on the hierarchical Mediator. Depending upon management requirements, individual access to each Mediator may still be required.

One advantage of implementing hierarchical Mediator functionality is that it may be possible to restrict enterprise client PCs which need access to energy management information directly from the Mediators, to only the hierarchical Mediator. This has security advantages in terms of configuration of access control and monitoring to a single Mediator; versus allowing enterprise client PCs potentially spread throughout the network to access individual Mediators. It may also be possible to restrict MSP partner VPN access to only the hierarchical Mediator device deployed within the EMOC - depending upon the management requirements. Note however that current hierarchical Mediator functionality requires RNA functionality to be configured on both the hierarchical Mediator and remote Mediators. This requires initial management workstation access to both. Further, ongoing changes to the configuration of the Mediators through the configTOOL, and/or applications deployed on the Mediators via perfectHOST; may still require the MSP partner management servers and enterprise management servers to have direct access to the remote Mediators.

The current hierarchical Mediator functionality allows for the datapoints on remote Mediators to be aliased, so that they appear to be local datapoints on the hierarchical mediator, and are automatically collected from the remote Mediators via the RNA protocol. Websites created on the hierarchical Mediator can then display datapoint information on the remote Mediators as if it were local to the hierarchical Mediator. Note that the current hierarchical Mediator design has scale limitations which depend on factors such the number of remote Mediators deployed, the number of datapoints aliased from those remote Mediators, the collection interval of the datapoints, and the number websites deployed and accessed on the hierarchical Mediator. The Cisco Network Building Mediator Manager functionality is a future function targeted to be deployed on a dedicated server appliance, as opposed to a Mediator platform, to address such scalability concerns. It is targeted to scale to approximately 200 remote Mediators.

The following sections discuss the deployment of the Energy Management Operations Center (EMOC) as a service module within a larger data center design; as well as within a service module hanging off the campus core. The designs further discuss deployments in which VRFs have been used (using the VRF-Lite with GRE tunnel method) in order to provide path isolation of the energy management solution from the rest of the enterprise network; as well as non-VRF deployments.

Data Center Service Module Design

In some situations, the facilities management personnel who are managing the energy management solution are physically located within a data center of the campus. In these scenarios, a separate service module hanging off of the overall data center design can be implemented for the EMOC. The following sections discuss data center EMOC designs from both a non-VRF and VRF perspective.

Non-VRF Designs

Figure 5-2 shows an example of a Data Center Service Module design without the use of VRFs.

Figure 5-2 Non-VRF Data Center Service Module Design with Catalyst 6500 Service Switch and FWSM

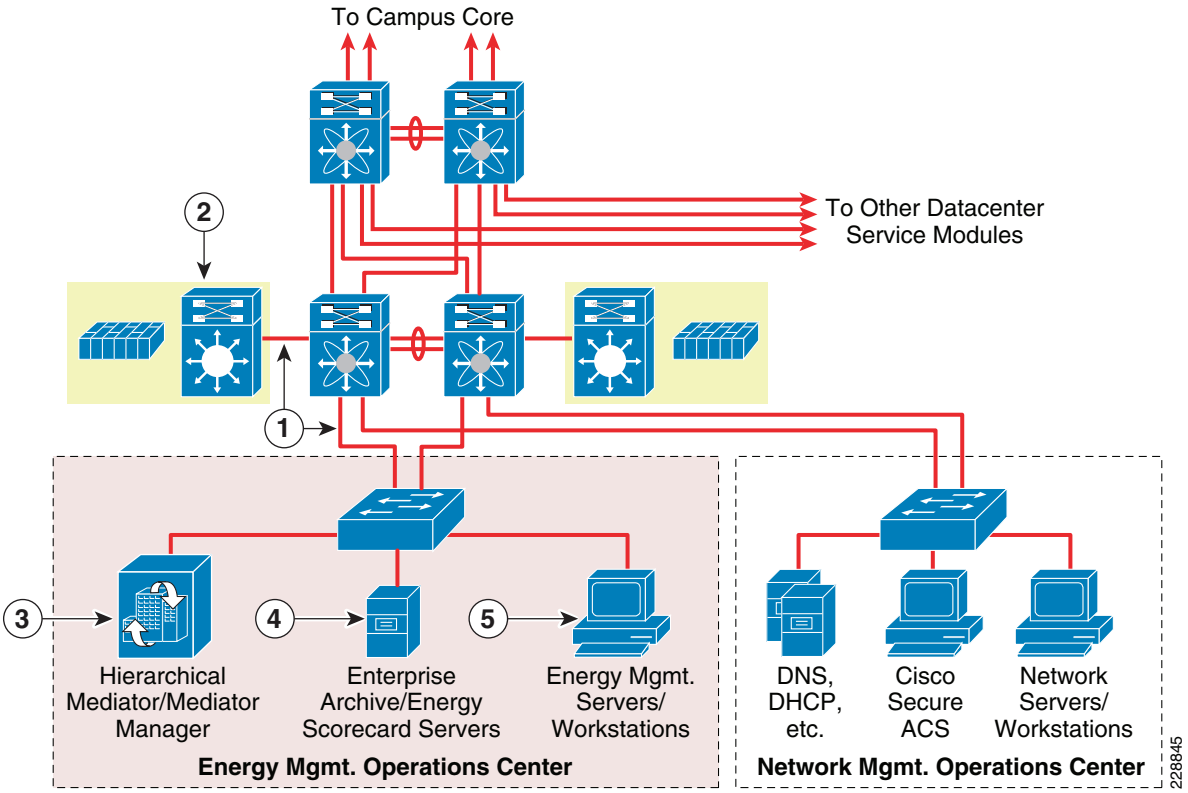

The following describes the numbers in Figure 5-2:

- **1**—Energy Management Operations Center (EMOC) VLAN trunked through Nexus 7000 Distribution Switch and the Catalyst 6500 Service Switch to the FWSM.
- **2**—FWSM in the Catalyst 6500 Service Switch provides stateful access control to and from the EMOC devices.

- **3**—Hierarchical Mediator functionality can provide a real-time single point of access to energy usage information (datapoints) from Mediators deployed throughout the enterprise network.
- **4**—Internal enterprise archive servers may collect periodically exported datapoint information from individual Mediators deployed throughout the network infrastructure. An energy scorecard service may make historical energy usage information available internally within the organization.
- **5**—Internal enterprise management workstations may be used to configure the hierarchical and/or remote Mediators; deploy logical control applications to the hierarchical and/or remote Mediators; and to create, monitor, and deploy websites to the hierarchical and/or remote Mediators.

In this example, a separate EMOC VLAN is implemented within the Data Center Service Module. The EMOC VLAN is trunked from the access switch, through the Nexus 7000 Series data center Distribution Switch, to a Layer-3 interface of the FWSM module located within the Catalyst 6500 Service Switch. The FWSM then provides stateful access control to and from the EMOC VLAN. A more detailed example is shown in Figure 5-3.

Figure 5-3 Detailed Example of Data Center Service Module Design with Catalyst 6500 Service Switch and FWSM

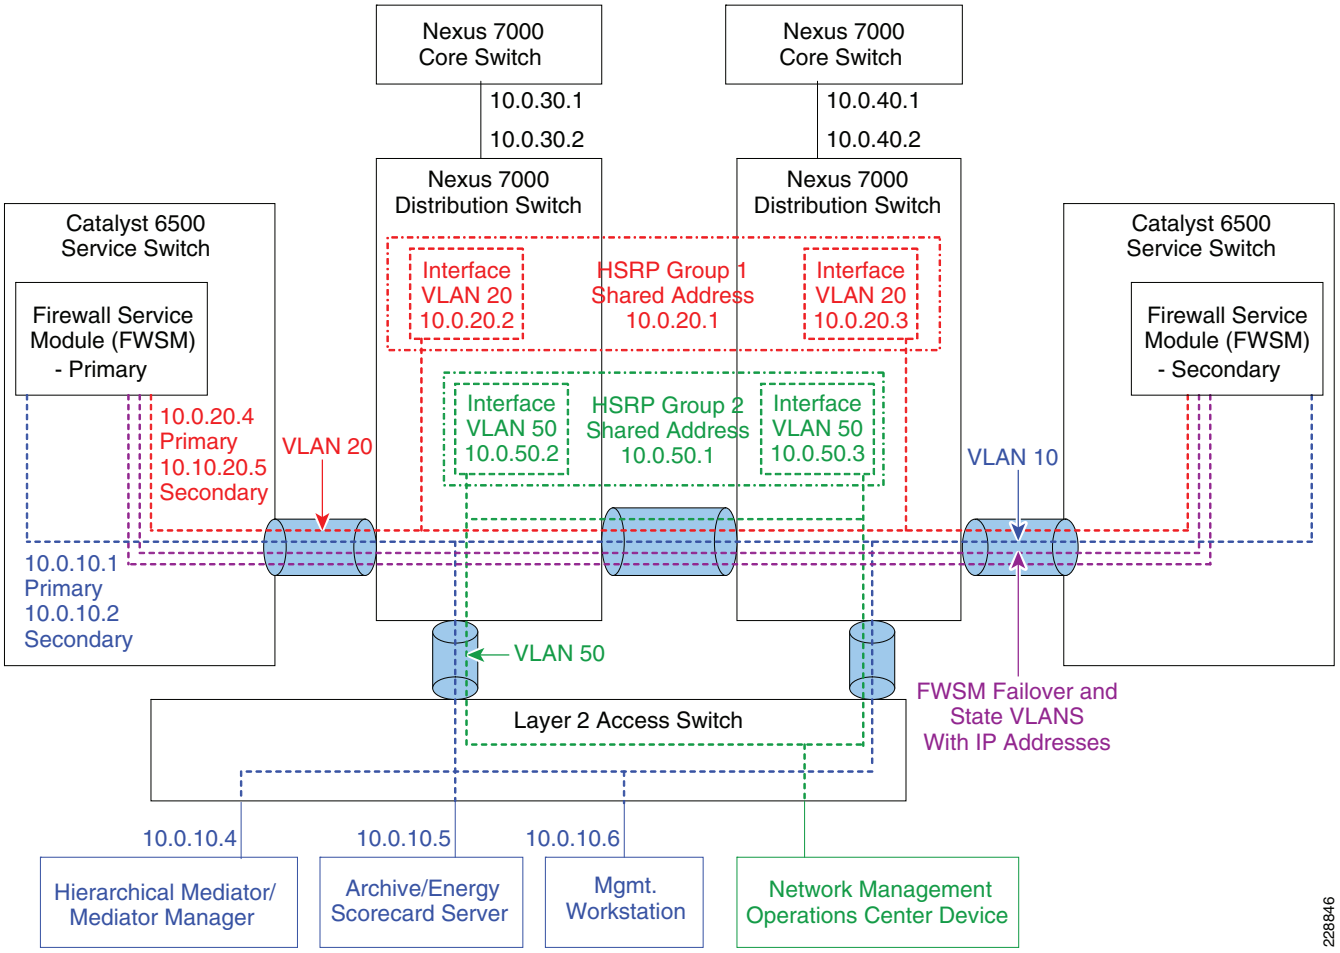

Note that there are many different ways in which firewalling of the EMOC VLAN could be achieved within the data center, because the FWSM supports both transparent (Layer 2) and routed mode (Layer 3) firewalling, high-availability through active/active or active/standby firewall configurations, and single or multiple context (virtual firewalls) mode. This example shows only one highly simplified method using a single context, routed mode firewall in an active/standby configuration. IP addressing

has been arbitrarily chosen for example purposes only. In this configuration, when the primary FWSM fails, the secondary FWSM takes over its IP and MAC addresses. From the perspective of the EMOC devices nothing has changed. However, traffic flows to the secondary FWSM located within the second Catalyst 6500 Service Switch. Stateful firewalling is done on the FWSM between VLANs 10 and 20 in the example, corresponding to IP subnets 10.0.10.0 and 10.0.20.0.

Other VLANs such as a Network Operations Center (NOC) VLAN can also be supported off the same Data Center Service Module, as shown in Figure 5-2 and Figure 5-3. The NOC VLAN can support traditional functionality such as network management servers, DNS, and DHCP, as well as the Cisco Secure Access Control Server which provides AAA services for MSP Partner VPN access to the enterprise network. For example purposes only, the design shows traffic from the Network Management Operations Center (NMOC) VLAN not passing through the FWSM. In actual deployments, this may be firewalled as well.

An alternative to the Catalyst 6500 Service Switch with FWSM design is to implement a set of ASA 5500 Security Appliances within the Data Center Service Module, as shown in Figure 5-4.

Figure 5-4 Data Center Service Module Design with ASA 5500 Security Appliances

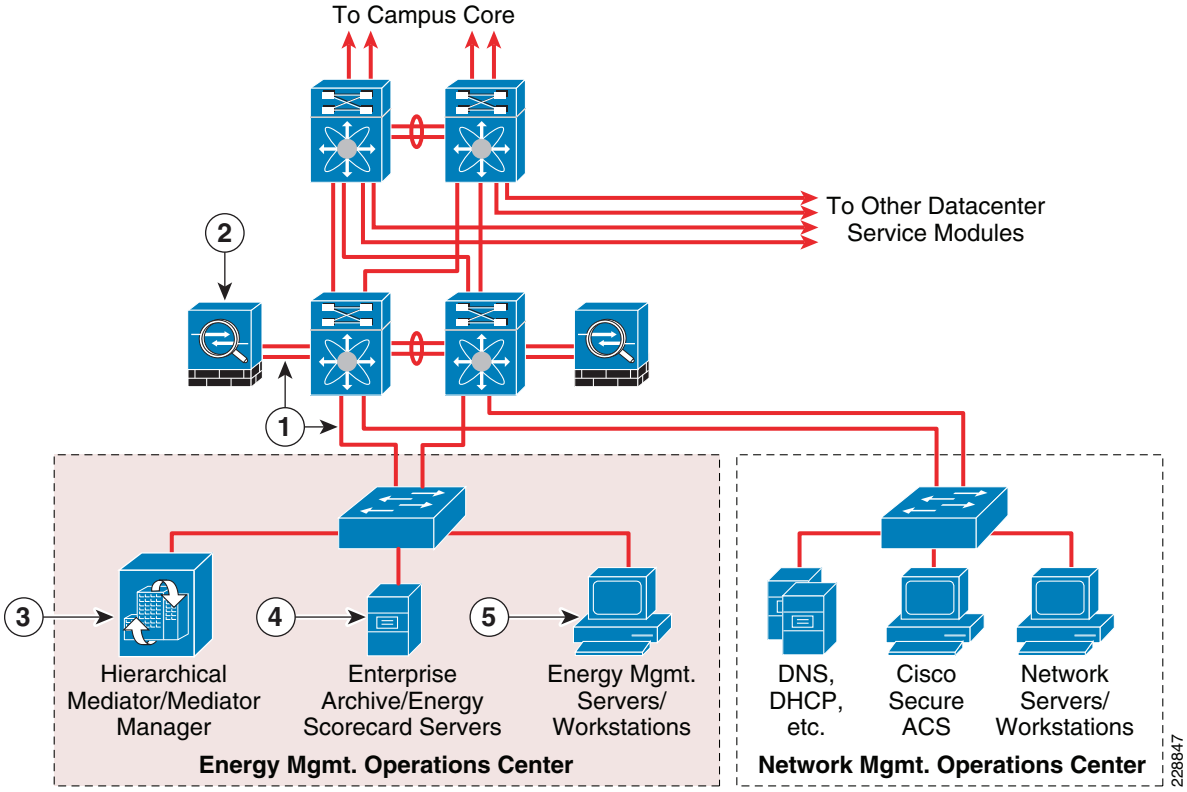

The following describes the numbers in Figure 5-4:

- **1**—EMOC VLAN trunked through Nexus 7000 Distribution Switch to ASA 5500 Security Appliances.
- **2**—ASA 5500 Security Appliances in the data center provide stateful access control to and from the EMOC devices.
- **3**—Hierarchical Mediator functionality can provide a real-time single point of access to energy usage information (datapoints) from Mediators deployed throughout the enterprise network.

- **4**—Internal enterprise archive servers may collect periodically exported datapoint information from individual Mediators deployed throughout the network infrastructure. An energy scorecard service may make historical energy usage information available internally within the organization.
- **5**—Internal enterprise management workstations may be used to configure the hierarchical and/or remote Mediators; deploy logical control applications to the hierarchical and/or remote Mediators; and to create, monitor, and deploy websites to the hierarchical and/or remote Mediators.

In this design, the EMOC VLAN is again trunked from the access switch, through the Nexus 7000 Series data center Distribution Switch, to a Layer-3 interface of the ASA 5500 Series Security Appliance. The ASA 5500 provides stateful access control to and from the EMOC VLAN. A more detailed example is shown in Figure 5-5.

Figure 5-5 Detailed Example Data Center Service Module Design with ASA 5500 Series Security Appliances

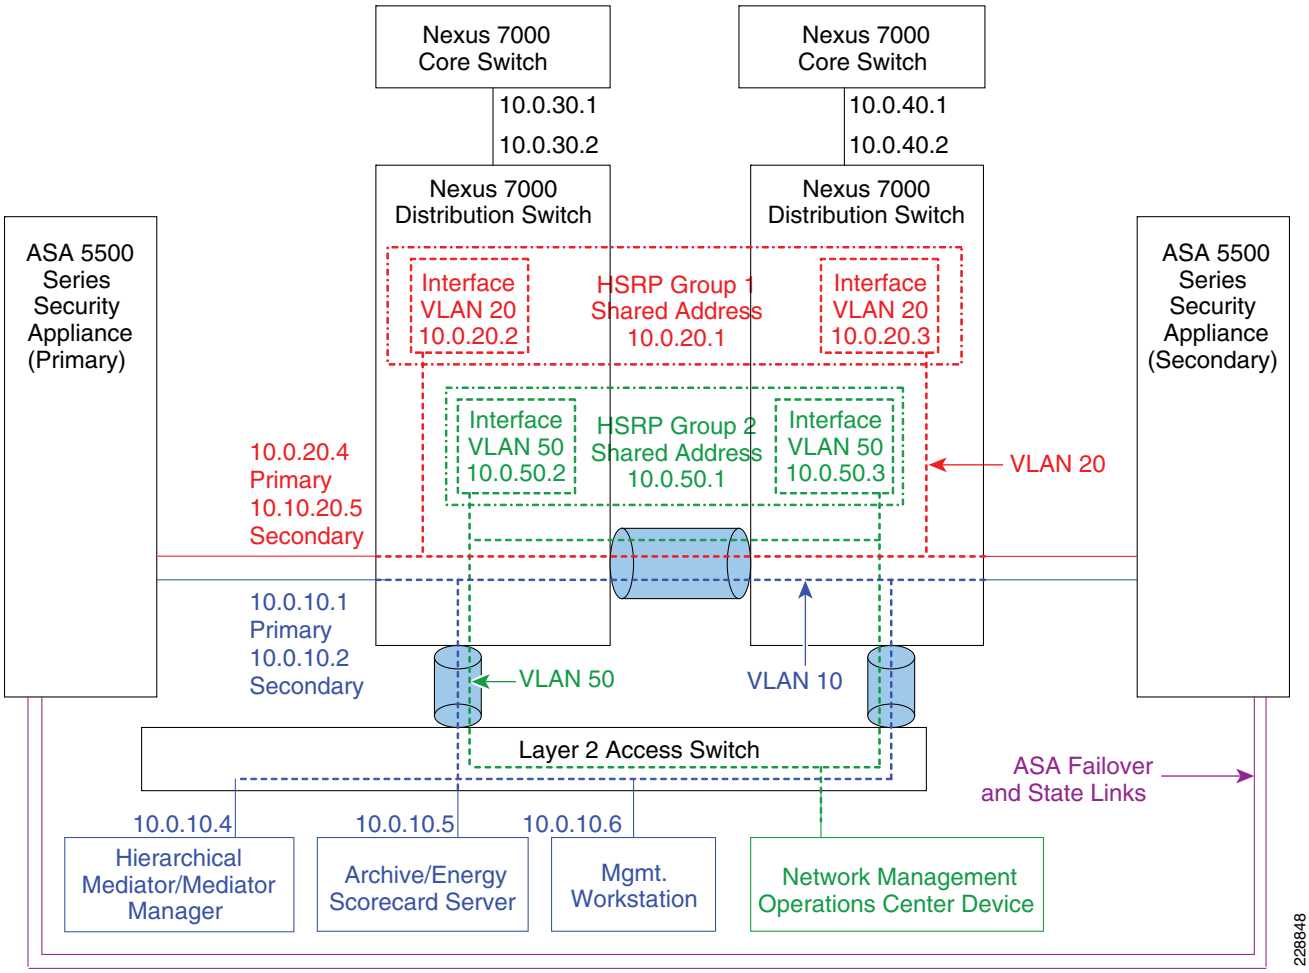

Again, note that there are many different ways in which firewalling of the EMOC VLAN could be achieved within the data center, because the ASA 5500 Series supports both transparent (Layer 2) and routed mode (Layer 3) firewalling, high-availability through active/active or active/standby firewall configurations, and single or multiple context (virtual firewalls) mode. This example shows only one highly simplified method using a single context, routed mode firewall in an active/standby configuration. IP addressing has been arbitrarily chosen for example purposes only. Individual interfaces are shown in this example, although VLAN sub-interfaces could also be used. In this configuration, if the primary ASA 5500 fails, the secondary ASA 5500 takes over its IP and MAC addresses. Again, from the

perspective of the EMOC devices nothing has changed. However, traffic flows to the secondary ASA 5500. Stateful firewalling is done on the ASA 5500 between VLANs 10 and 20 in the example, corresponding to IP subnets 10.0.10.0 and 10.0.20.0.

Since the examples in Figure 5-2 and Figure 5-4 4 do not assume a separate VRF for the Energy Management Solution, access control into the EMOC VLAN should be very tightly controlled. Figure 5-6, Figure 5-7, and Figure 5-8 show the protocols which may be allowed through the firewall.

Figure 5-6 Inbound Management and Monitoring Flows From MSP Partner

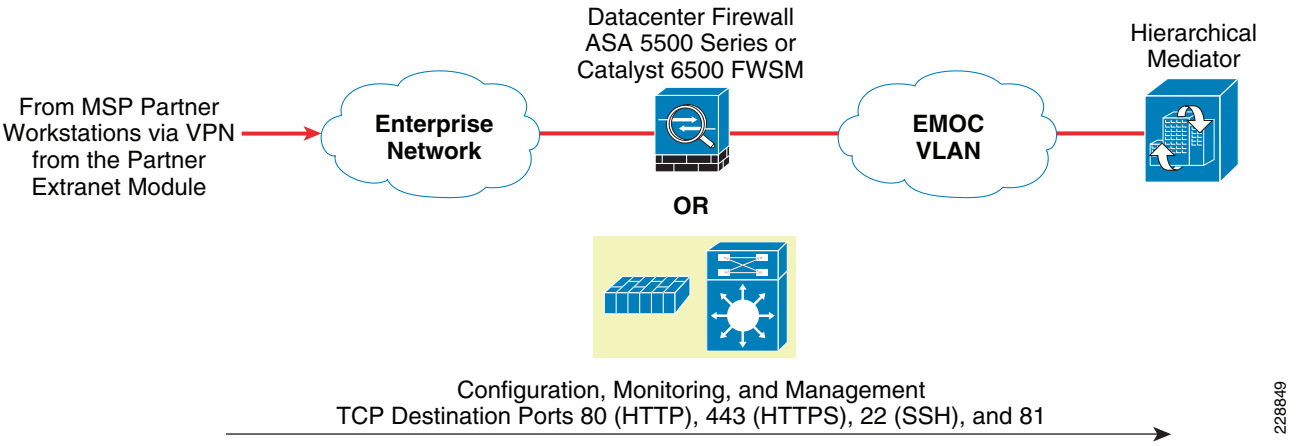

Figure 5-6 shows an example in which a MSP partner is providing management and/or monitoring of the Mediator deployment through a hierarchical Mediator deployed within the EMOC. In such cases, HTTP (TCP port 80), HTTPS (TCP port 443), SSH (TCP port 22) and TCP port 81 protocols need to be allowed inbound from the campus network through the data center firewall to the hierarchical Mediator within the EMOC. These allow inbound connectivity for the configTOOL, perfectHOST, and OMEGA suite of applications. The access control should be specified down to the IP address or addresses of the individual MSP partner workstations which connect to the enterprise network via either remote-access VPN or site-to-site VPN at the Partner Extranet Module. Note that the MSP partner workstations may still need access to each individual Mediator, but this access does not pass through the data center firewall in this design.

Figure 5-7 shows an example in which remote Mediators deployed throughout the IP network infrastructure periodically export datapoint information to an enterprise archiving server located within the EMOC. The remote Mediators may simultaneously export to a cloud services partner reachable via the Internet as well, but the flows do not pass through the data center firewall and are therefore not shown in Figure 5-7.

Figure 5-7 Inbound Exported Datapoint and Outbound Management Flows

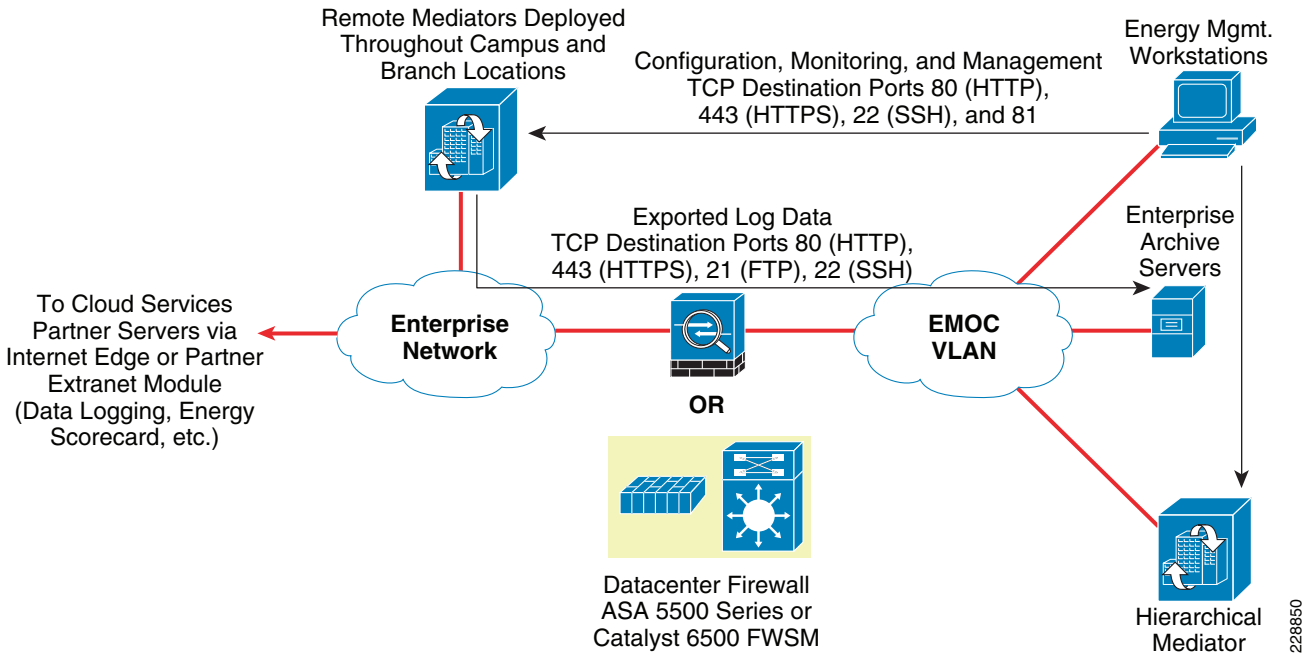

In such deployments, HTTP (TCP port 80), HTTPS (TCP port 443), FTP (TCP port 23), and/or SFTP that uses SSH (TCP port 22) needs to be allowed inbound from the remote Mediators through the data center firewall. Return traffic will automatically be allowed with a stateful firewall. In typical deployments, only a single protocol is used for the periodic exports. Note that when using FTP for periodic exporting of datapoint information, application-layer inspection of the FTP traffic is needed in order to dynamically open the data channel for the transport. Otherwise, a static range of ports may need to be opened for the data channel with FTP.

The enterprise energy management workstations also requires access each remote Mediator for configuration; deployment of graphical control applications; and the creation, deployment, and monitoring of websites on the remote Mediators. Therefore, HTTP, HTTPS, SSH, and TCP port 81 sessions initiated from the enterprise management workstations may need to be allowed from the EMOC out to the enterprise network, as is shown in Figure 5-7. No modifications to the data center firewall are needed for the enterprise management workstations to manage and monitor the hierarchical Mediator via the configTOOL, perfectHOST, and OMEGA suite of applications, provided the hierarchical Mediator is on the same VLAN segment as the enterprise management workstations.

Firewall access control should be specified down to the source IP address of the remote Mediator to the addresses of the enterprise archiving servers which receive exported datapoint information. Likewise, firewall access control should be specified down to the source IP address of the enterprise energy management workstation to the addresses of the individual Mediators deployed in branch and campus locations.

In some deployments it may be necessary for the hierarchical Mediator itself to export event data via the SMTP protocol through a corporate E-mail server. Figure 5-8 shows an example of this.

Figure 5-8 Exported Event Data via SMTP

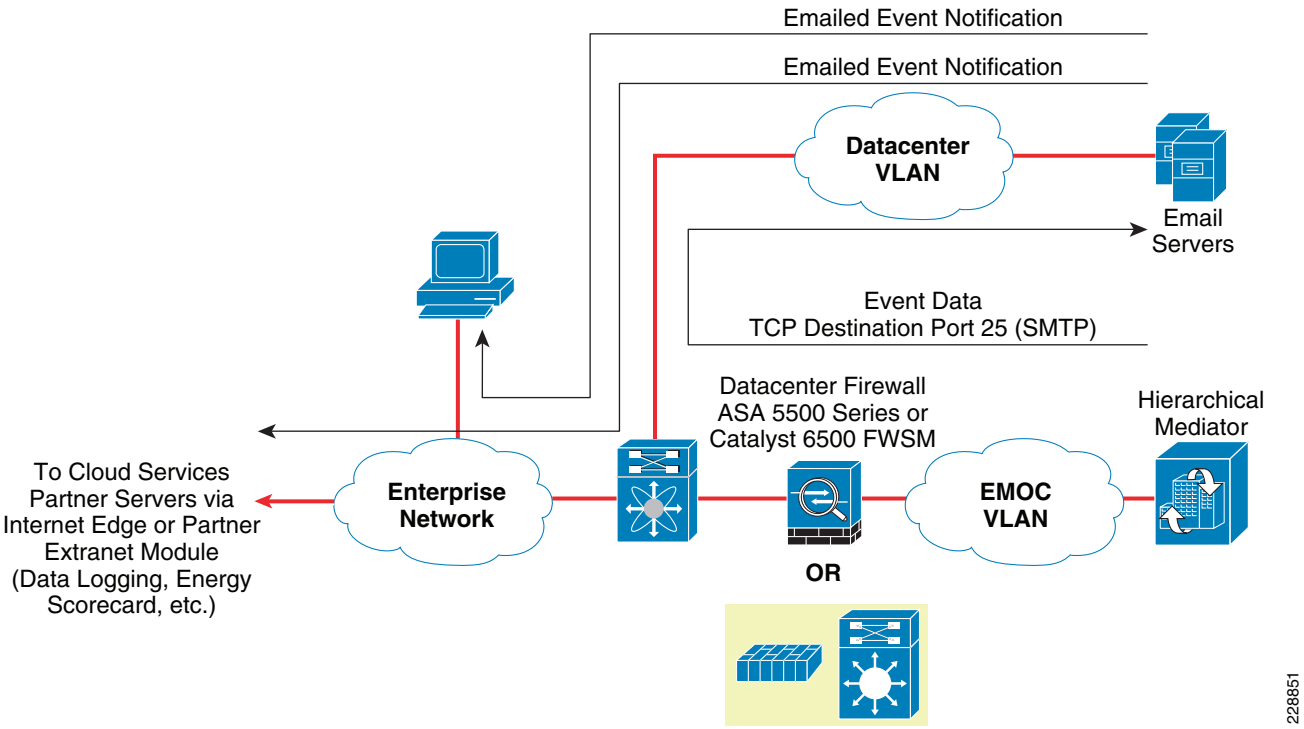

For event data, Mediators deployed throughout the network can be configured to form a cloud. Event data can be viewed and cleared on any of the Mediators. If events are exported, each Mediator can export via the SMTP protocol (TCP port 25) to an E-mail server (or RSS server). Since the hierarchical Mediator may sit within the EMOC, SMTP traffic initiated by it must be allowed outbound through the data center firewall to the corporate E-mail servers that may reside on another data center VLAN segment, possibly hanging off another Data Center Service Module. In this example, the corporate E-mail server then forwards the event information to E-mail clients on enterprise client PCs as well as possibly MSP partner E-mail addresses or E-mail aliases.

Figure 5-9 RNA Flows Between Mediators

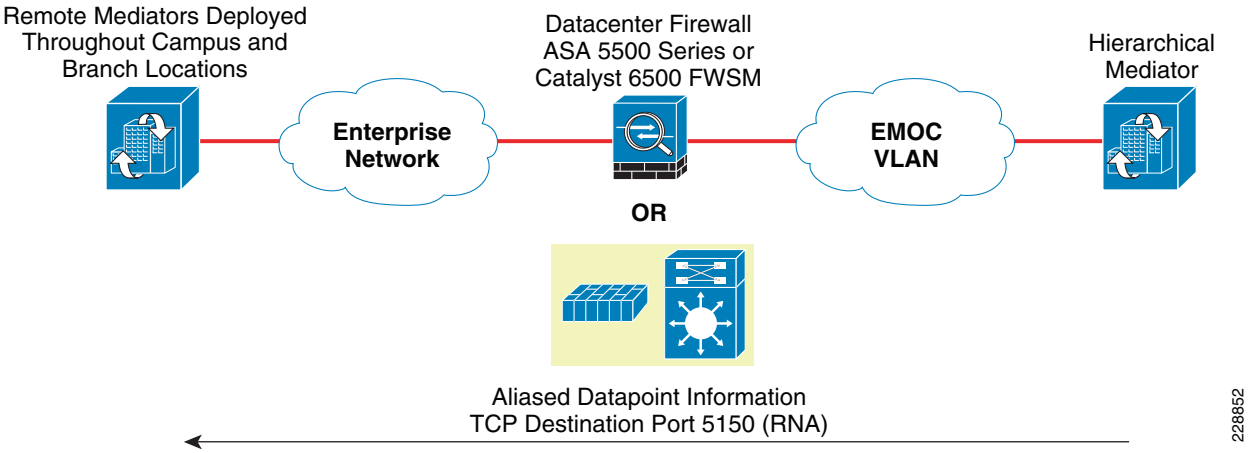

Figure 5-9 shows RNA flows (TCP port 5150) which need to be allowed outbound through the data center firewall in order for the hierarchical Mediator to automatically collect aliased datapoint information from remote Mediators deployed throughout the enterprise network. Note that in this example the hierarchical Mediator initiates the RNA flow, since the actual datapoints exist on the remote Mediators, and are aliased on the hierarchical Mediator. Access control should be specified down to the IP addresses of the individual Mediators in which datapoint information is being shared via the RNA protocol.

Figure 5-10 Data Center Flows with a Non-Hierarchical Mediator Design

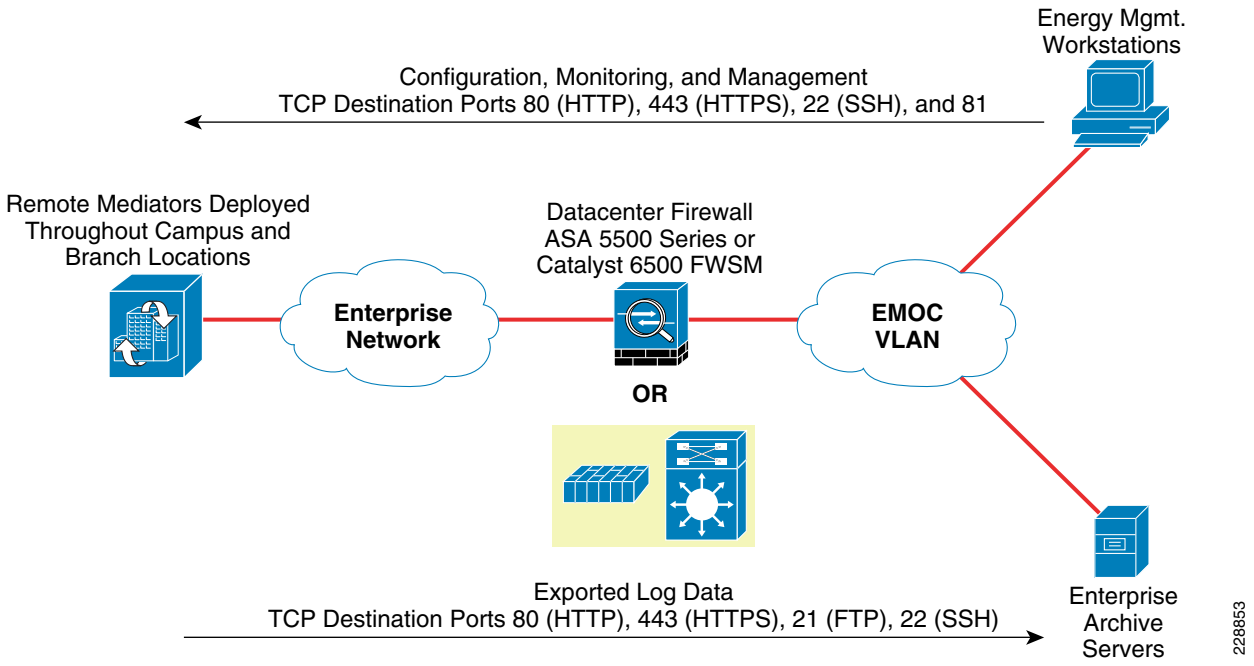

Figure 5-10 shows an example in which a hierarchical Mediator design is not deployed. In this model, the MSP partner servers individually access the Mediators deployed within the campus and branch locations. Therefore, MSP partner access into the EMOC may not be needed at all, and is therefore not shown in the figure. Sessions initiated from enterprise energy management servers located within the EMOC will still need to be allowed through the data center firewall to manage and monitor individual Mediators. Likewise, periodically exported datapoint information from each Mediator to the cloud services partner servers will not cross into the EMOC, and therefore do not need to be allowed through the data center firewall. However, periodically exported datapoint information from each Mediator to the enterprise archiving server will need to be allowed inbound through the data center firewall.

Enterprise Client PC Access to the EMOC

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.

The first option is discussed in Chapter 4, "Internet Edge Design Considerations." This chapter discusses options 2 and 3. Option 4 is discussed within both Chapter 3, "Campus Design Considerations" and Chapter 6, "Branch Design Considerations."

When deploying an internal energy scorecard server, it may be beneficial from a security standpoint to separate the archiving function which collects and stores the periodically exported datapoint information from the Mediators, from the actual web/application server interface that client PCs access. This could be accomplished by implementing two servers: an energy scorecard web/application sever and an archiving server. One method of implementation is to have the energy scorecard web/application server remotely access a SQL database on the archiving server. Alternatively, and perhaps a bit more secure, would be to replicate the database information from the archiving server to the server which provides the web/application server interface for the energy scorecard. The energy scorecard server could then sit on a DMZ segment between the EMOC VLAN and the rest of the enterprise data network. Figure 5-11 shows an example of this type of design.

Figure 5-11 Example Client PC Access to Internal Energy Scorecard

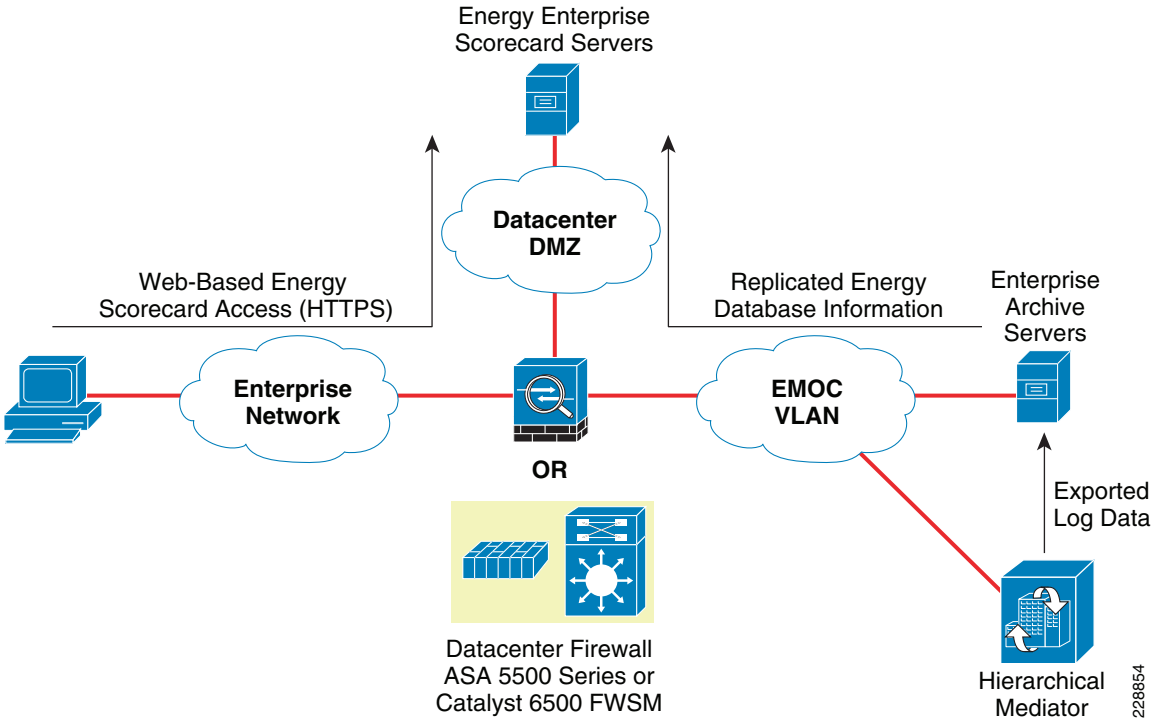

In this design, client PCs would be allowed to access only the energy scorecard server and not the actual Mediators which may be running logical control applications for building devices, or the actual archiving server which holds the historical database of logged datapoints. This continues to isolate the energy management solution from the rest of the data network, while still providing energy usage information to internal business units within the enterprise organization. The downside to this approach is that the

data may not be as “real-time” since the datapoint information is periodically exported and then replicated from the archiving server to the energy scorecard server; as opposed to directly accessing a Mediator via a website deployed directly on it.

Direct client PC access to websites deployed on the hierarchical Mediator within the EMOC is considered less secure than accessing an energy scorecard, yet still a somewhat more secure approach than allowing direct access to each remote Mediator. The hierarchical Mediator may be a dedicated device, configured only to collect aliased datapoint information from remote Mediators and display the information in custom websites developed for various internal business units within the enterprise organization; as opposed to directly interfacing with actual building systems. This approach requires direct inbound access from client PCs to the hierarchical Mediator within the EMOC; and therefore poses a higher security risk than accessing an energy scorecard. If business needs require enterprise client PCs to manage set points on the Mediators, rather than passively view datapoint information, direct access to the hierarchical Mediator may also be necessary. Where possible the use of HTTPS to access the hierarchical Mediator should be utilized. Userids and passwords on the hierarchical Mediator, restricting the particular website screens to the relevant business unit personnel, should be configured where possible. If possible, the access control from the data network to the EMOC should be restricted to a set of IP addresses which can only reach the hierarchical Mediator via the HTTPS protocol. For additional security the network administrator should consider two additional measures. First, the hierarchical Mediator itself may be placed on a DMZ segment off the data center firewall, between the enterprise client PC network and the EMOC. This is a similar design as was shown in Figure 5-11above. Client PCs would be allowed to access the hierarchical Mediator through the data center firewall via the HTTPS protocol. The hierarchical Mediator would be allowed to access the remote Mediators throughout the network via the RNA protocol. Alternatively, a second design is the deployment of IPSec remote access VPN on an ASA 5500 data center firewall. This would allow for access into the EMOC based on individual userid and password which can be controlled centrally via a AAA server. Figure 5-12 shows an example of this type of configuration.

Figure 5-12 Example Client PC Access to Hierarchical Mediator

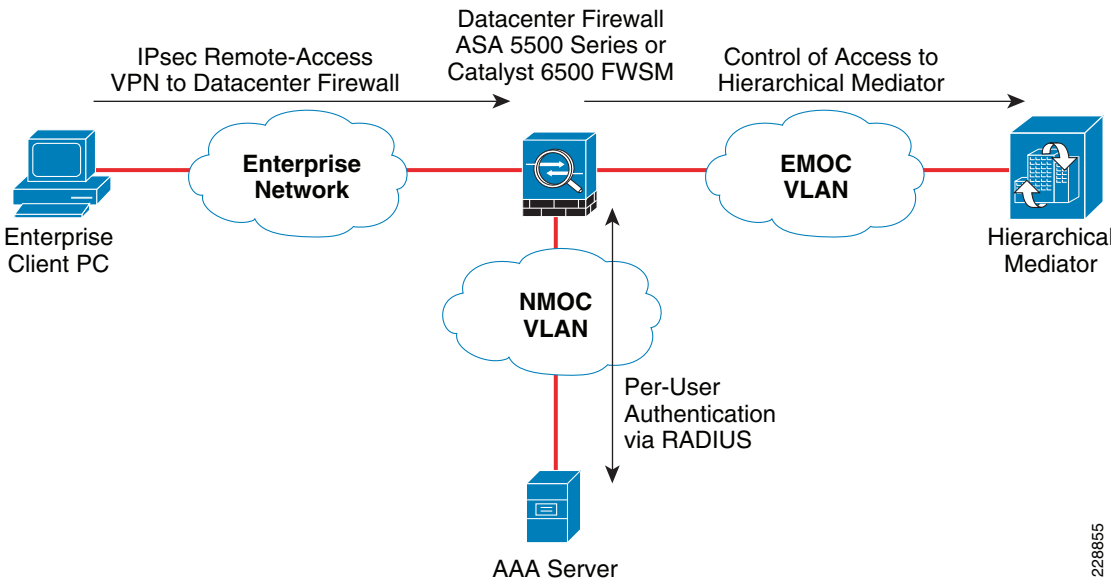

VRF Designs

The deployment of a separate virtual routing and forwarding instance (VRF) has the advantage of providing path isolation for the energy management solution. Since a separate energy management VRF already effectively isolates the traffic from the rest of the data, voice, and video traffic on the global VRF, stateful firewalling is not necessarily needed within the Campus Building Modules or within the branches. This means that controlling access between the energy management solution and the normal data network can be limited to one or more strategic points, versus being individually controlled within each campus building and each branch location. Figure 5-13 shows an example of a data center service module which makes use of VRFs to isolate the EMOC to a separate energy management VRF called the Building Infrastructure Network (BIN) VRF.

Figure 5-13 VRF Data Center Service Module Design

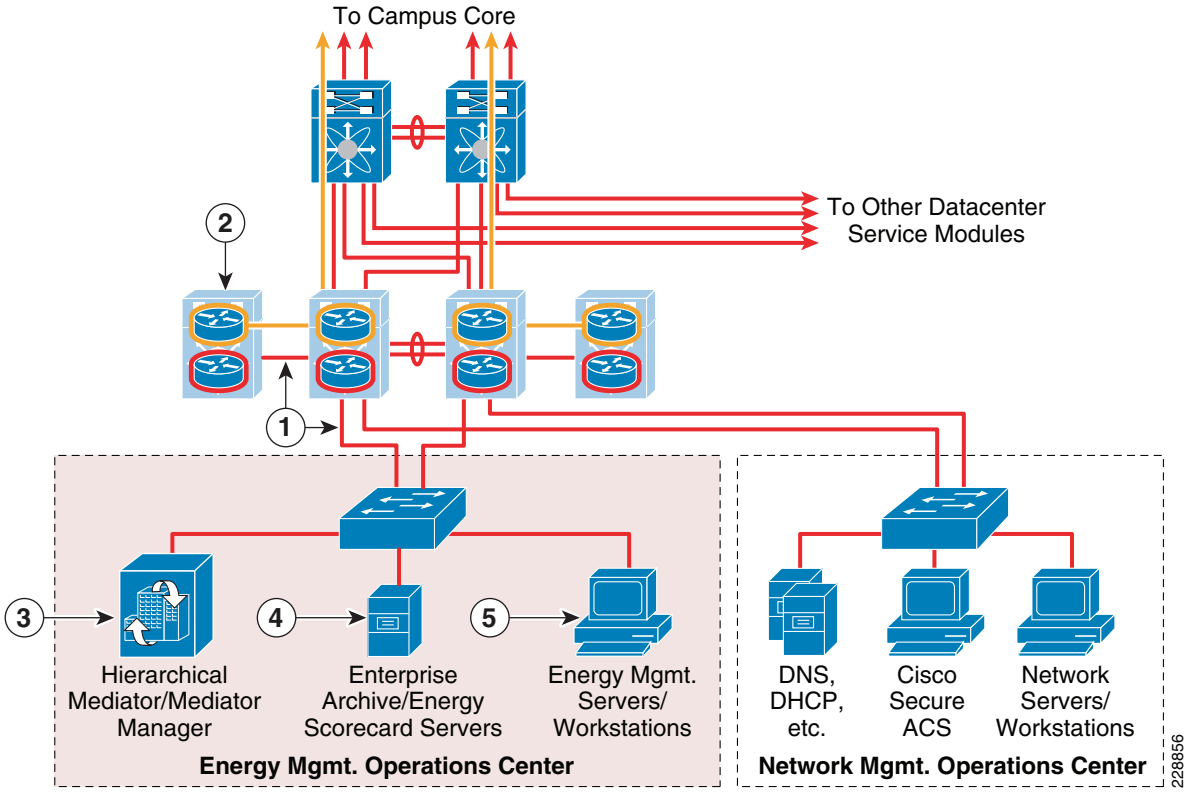

The following describes the numbers in Figure 5-13:

- **1**—EMOC VLAN trunked through Nexus 7000 Distribution Switch to the Catalyst 6500 Service Switch.
- **2**—GRE tunnels from Partner Extranet Module, Campus Building Modules, and WAN Edge Module terminate on Catalyst 6500 Service Switches; extending the BIN VRF across the campus network to these locations.
- **3**—Hierarchical Mediator functionality can provide a real-time single point of access to energy usage information (datapoints) from Mediators deployed throughout the enterprise network.
- **4**—Internal enterprise archive servers may collect periodically exported datapoint information from individual Mediators deployed throughout the network infrastructure. An energy scorecard service may make historical energy usage information available internally within the organization.

- **5**—Internal enterprise management workstations may be used to configure the hierarchical and/or remote Mediators; deploy logical control applications to the hierarchical and/or remote Mediators; and to create, monitor, and deploy websites to the hierarchical and/or remote Mediators.

In this example, a separate EMOC VLAN is again implemented within the Data Center Service Module. The EMOC VLAN is trunked from the access switch, through the Nexus 7000 Series Distribution Switch, to a Layer-3 interface defined within the Catalyst 6500 Service Switch. The Layer-3 interface sits on the BIN VRF, separate from the global VRF which houses the rest of the enterprise data network. GRE tunnels are then defined on the Catalyst 6500 Service Switch, which extend the BIN VRF out into the campus network. Figure 5-14 shows a more detailed view, without redundancy for clarity of the drawing.

Figure 5-14 Detailed Example Data Center Service Module Design Using VRF-Lite with GRE Tunnels

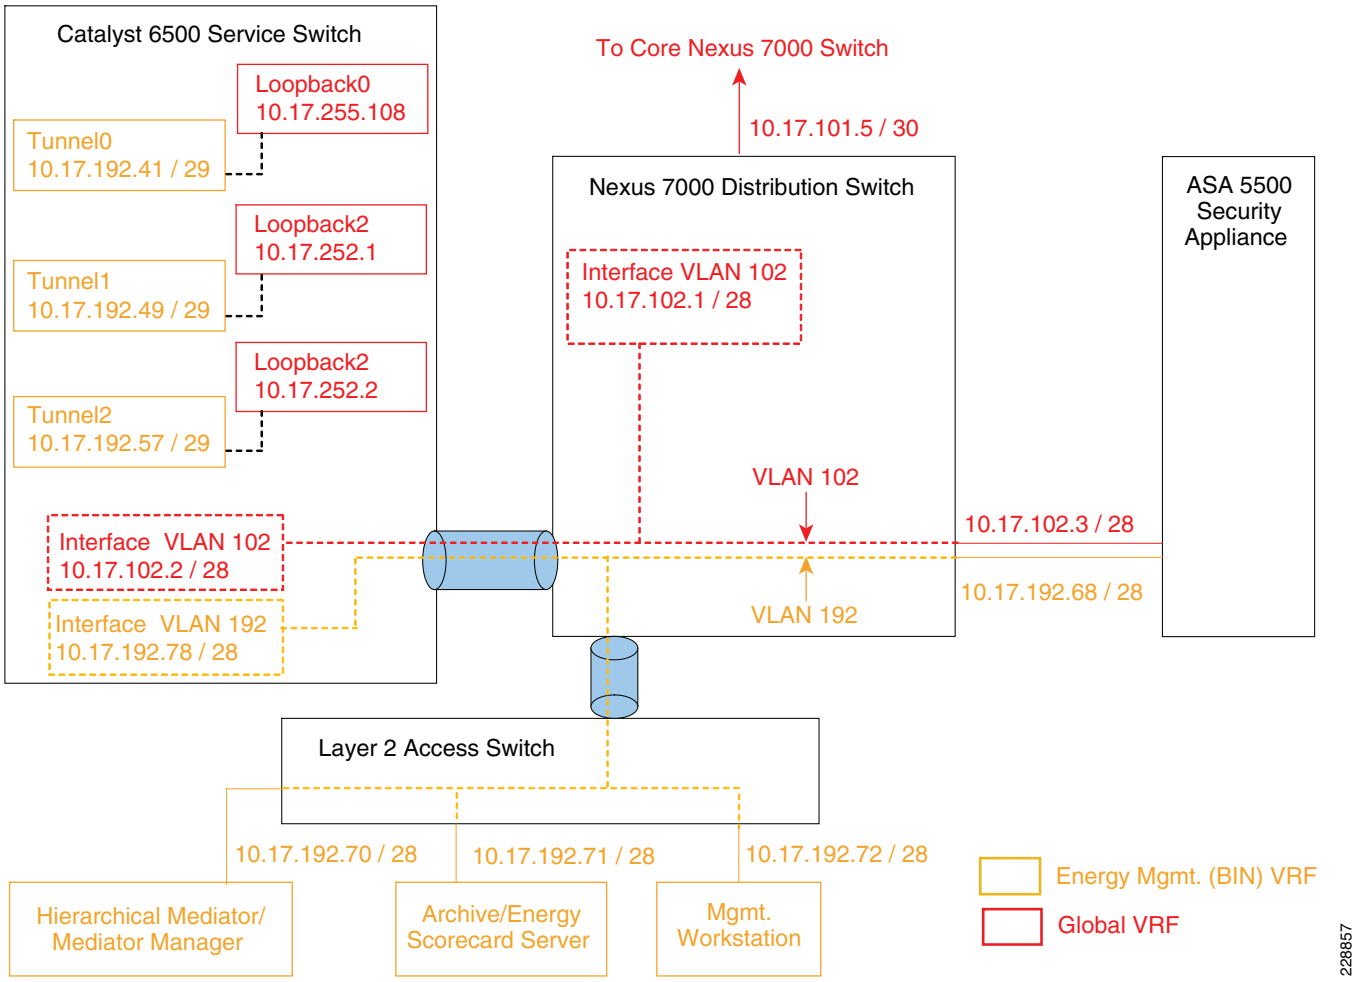

Example 5-1 provides a partial configuration from the Catalyst 6500 Service Switch.

Example 5-1 Partial Configuration Catalyst 6500 Service Switch with VRFs and GRE Tunnels

```
!
ip vrf bin
! Creates Building Infrastructure Network (BIN) VRF
rd 192:108
!
~
!
interface Tunnel0
! GRE Tunnel to Partner Extranet Module
description VRF FOR MEDIATOR NETWORK TO ME-EASTDIST-1
ip vrf forwarding bin
ip address 10.17.192.41 255.255.255.248
tunnel source Loopback4
tunnel destination 10.16.255.40
!
interface Tunnel1
! GRE Tunnel to Campus Building Module
description VRF FOR MEDIATOR NETWORK TO ME-WESTCAMP-1
ip vrf forwarding bin
ip address 10.17.192.49 255.255.255.248
tunnel source Loopback2
tunnel destination 10.17.255.51
!
interface Tunnel2
! GRE Tunnel to WAN Edge Module
description VRF FOR MEDIATOR NETWORK TO ME-WESTDIST-1
ip vrf forwarding bin
ip address 10.17.192.57 255.255.255.248
tunnel source Loopback0
tunnel destination 10.17.252.3
!
interface Loopback0
description LOOPBACK INTERFACE FOR TUNNEL FROM ME-WESTDIST-1
ip address 10.17.255.108 255.255.255.255
!
interface Loopback2
description LOOPBACK INTERFACE FOR TUNNEL FROM ME-WESTCAMP-1
ip address 10.17.252.1 255.255.255.255
!
interface Loopback4
description LOOPBACK INTERFACE FOR TUNNEL FROM ME-EASTDIST-1
ip address 10.17.252.2 255.255.255.255
!
~
!
interface TenGigabitEthernet1/2
! Trunk to Nexus 7000 Data Center Distribution Switch
description TRUNK TO ME-WESTDC7K-1
switchport
switchport trunk encapsulation dot1q
switchport mode trunk
load-interval 30
wrr-queue bandwidth 5 25 10 10 5 5 10
priority-queue queue-limit 30
wrr-queue queue-limit 5 25 10 10 5 5 10
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 4 80 100 100 100 100 100 100 100
```

```
wrr-queue random-detect min-threshold 5 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 6 80 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect 4
wrr-queue random-detect 5
wrr-queue random-detect 6
wrr-queue random-detect 7
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 2
wrr-queue cos-map 4 1 3
wrr-queue cos-map 5 1 6
wrr-queue cos-map 6 1 7
wrr-queue cos-map 7 1 4
mls qos trust dscp
!
~
!
interface Vlan102
! Extends DataCenter Global VRF to Cat6K Service Switch
ip address 10.17.102.2 255.255.255.240
!
interface Vlan192
! Extends DataCenter BIN VRF to Cat6K Service Switch
ip vrf forwarding bin
ip address 10.17.192.78 255.255.255.240
standby 1 ip 10.17.192.76
standby 1 priority 90
!
~
!
router eigrp 111
network 10.17.0.0 0.0.255.255
no auto-summary
!
address-family ipv4 vrf bin
! Creates EIGRP instance for BIN VRF
autonomous-system 99
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
exit-address-family
!
```

In the above example, separate GRE tunnels are provisioned, extending the BIN VRF to the Catalyst 6500 Series Distribution Switches within the following campus locations:

- *Partner Extranet Module*—This places inbound remote-access or site-to-site VPN traffic from the MSP partner onto the BIN VRF, effectively isolating it from the rest of the enterprise traffic.
- *WAN Edge Module*—This extends the BIN VRF to the WAN edge. A separate set of GRE tunnels then extends the BIN VRF out to each remote branch which houses a Mediator.
- *Campus Building Modules*—This extends the BIN VRF to each campus building which houses a Mediator.

The deployment of GRE tunnels which start out from the Data Center Service Module may provide a somewhat more manageable deployment for enterprise and partner access to the Mediators than provisioning tunnels directly from the Partner Extranet Module to each Mediator location within the campus or branches. An example of this is shown in Figure 5-15.

Figure 5-15 Star vs Partial-Mesh Deployment of GRE Tunnels

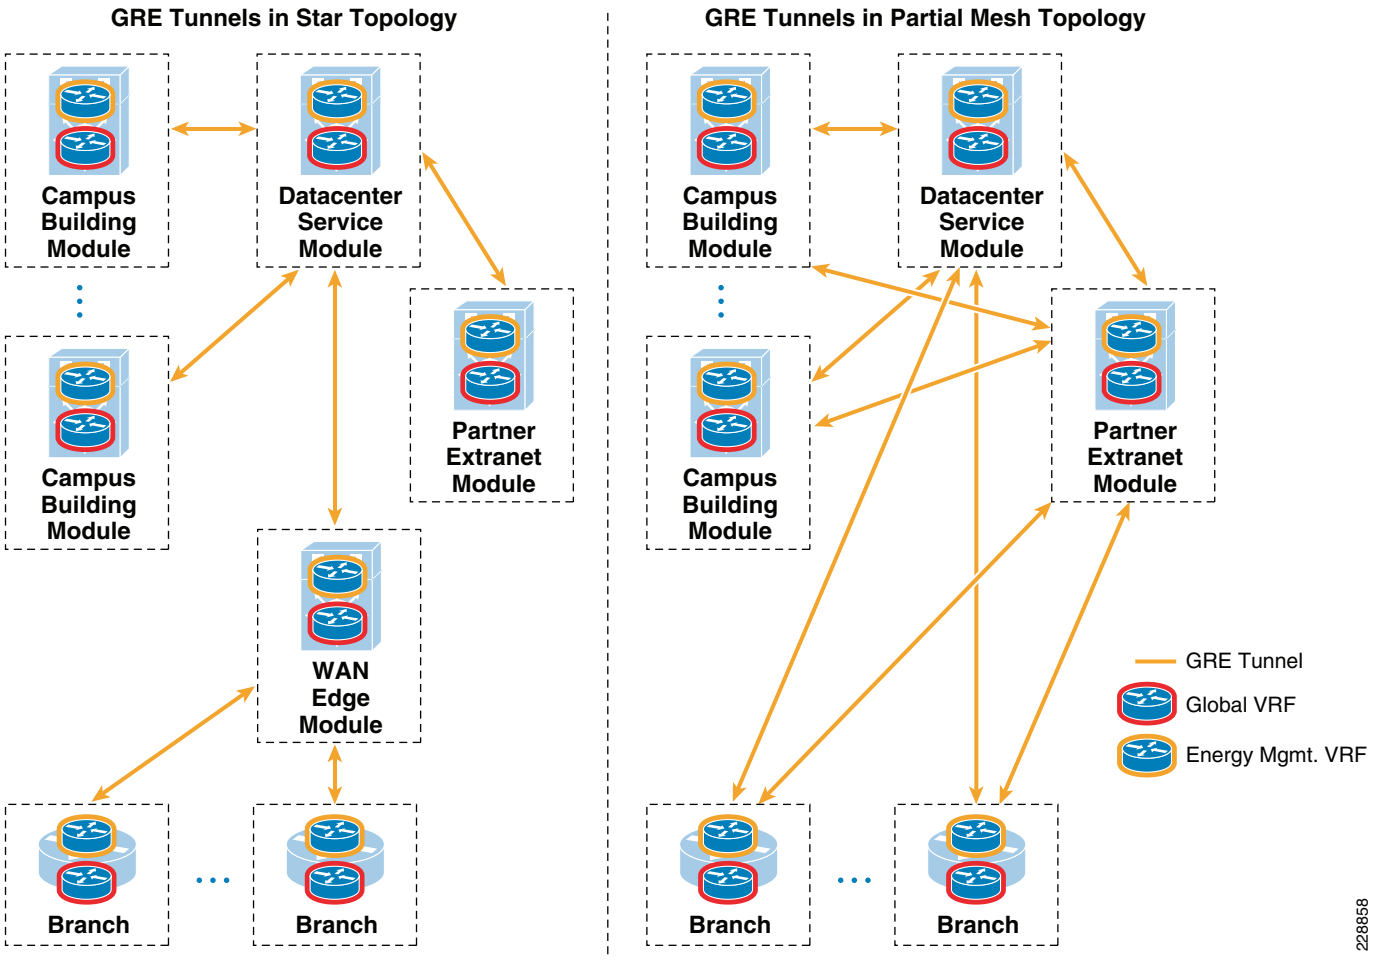

As shown in Figure 5-15, the overall number of GRE tunnels that must be provisioned is less with the star configuration versus the partial-mesh configuration. Note that Figure 5-15 does not show redundant GRE tunnels, which connect between the redundant Catalyst 6500 switch pairs. Catalyst 6500 switches handle GRE tunneling in hardware and, therefore, provide a platform sufficient for moderate-sized VRF deployments in this configuration. Further, with a star configuration, all traffic is routed back toward the Data Center Service Module before being routed out toward the remote branch and campus Mediators. This design facilitates the use of a hierarchical Mediator/Mediator Manager deployed within the data center EMOC. Since the BIN VRF is assumed to be dedicated to the energy management solution, no additional firewalling is shown within the datacenter between the EMOC and the rest of the energy management solution. Stateful access control at both the VPN concentrator and firewall within the Partner Extranet Module already restricts access from MSP partner devices. If desired, ACLs could be deployed across the GRE tunnel interfaces of the Catalyst 6500 Service Switches to further restrict access within the BIN VRF. This may be useful if the network virtualization concept is expanded to include other functionality such as video surveillance and physical access control, as well as energy management, into a single BIN VRF.

Other methods exist for providing VRF connectivity to the Data Center Service Module. These include the use of multipoint GRE tunnels that may be used to dynamically establish connectivity directly between tunnel endpoints. This may provide a more scalable deployment. Alternatively, an end-to-end

VRF-Lite implementation may be deployed in which the BIN VRF is extended throughout the Nexus 7000 data center switches as well as the campus core and Distribution Switches. Future revisions of this design guide may include designs using these technologies.

Client PC Access when Deploying a VRF for the Energy Management Solution

When deploying a VRF for the energy management solution, access control between enterprise client PCs located on the global VRF and devices on the energy management VRF should ideally be restricted to one point, or, alternatively, a small handful of points within the network infrastructure. Otherwise, the whole concept of implementing a VRF for path isolation becomes somewhat irrelevant. This also eases the administrative burden of not having to configure and manage multiple stateful firewalls deployed throughout the enterprise network when network virtualization is not implemented.

Chapter 4, "Internet Edge Design Considerations" discusses how MSP partner workstations that access the network via remote-access VPN or site-to-site VPN can gain access to the energy management VRF. This represents one access point into the energy management VRF. Therefore, one possible method of allowing enterprise client PCs access to the Mediators, when deploying a VRF, is also through the Partner Extranet Module. This could be accomplished by having the enterprise client PCs establish remote-access IPSec VPN connections to the BIN VRF via the same remote-access VPN device (in this case the ASA 5500 Series Security Appliances within the Partner Extranet Module) as MSP partner workstations. The advantage of this design is that all access into the BIN VRF, regardless of whether it is a partner workstation or enterprise client PC, is established through a single set of devices. In this case, the devices are the ASA 5500 Security Appliances deployed within the Partner Extranet Module. AAA services on the ASA 5500 Security Appliance can be used to centralize access control via RADIUS to a AAA server such as the Cisco Secure ACS server, which in turn may rely on a backend directory server or LDAP database. The disadvantage of this design is that it forces enterprise client PC traffic into the Partner Extranet Module; therefore, it violates the design paradigm of separation of partner and employee traffic at the Internet edge. Further, routing IPSec VPN traffic from enterprise client PCs into the Partner Extranet Module, where it is then decrypted and routed back into the BIN VRF, may be somewhat challenging.

Alternatively, a stateful firewall deployed within the Data Center Service Module (either a FWSM or ASA 5500 Security Appliance) can serve as a point of access control between the global VRF and the BIN VRF, as is shown in Figure 5-14 on page 5-15. The network administrator should still consider the use of remote-access IPSec VPN connections from the enterprise client PCs to the BIN VRF. This is one advantage of implementing an ASA 5500 Security Appliance over the Catalyst 6500 FWSM. This design has the advantage of maintaining the design paradigm of separating partner traffic from employee traffic at the Internet edge. Also, this design may be somewhat simpler to implement if the firewall already exists for E-mail (SMTP) traffic generated by the Mediators to be sent to corporate E-mail servers; or if the Mediators need to use corporate DNS servers. However, the disadvantage of this design is that there are effectively two access control points into the BIN VRF; one at the Internet edge for partner traffic and one within the data center for internal traffic. Note that the deployment of the hierarchical Mediator or internal energy management scorecard server on a DMZ segment can also be deployed with a VRF design, using the data center firewall.

Campus Service Module Design

In other situations, the facilities management personnel are not physically located within the data center of the campus or the campus location does not have a data center housed within it. In these scenarios, a separate Campus Service Module hanging off the Campus Core Module can be implemented for the EMOC. Again, both non-VRF and VRF designs can be implemented.

Non-VRF Campus Service Module Designs

The designs for the Campus Service Module are similar to the Data Center Service Module designs, but with some differences in the switching infrastructure. Figure 5-16 shows an example of a non-VRF Campus Service Module design with Catalyst 6500 Service Switch and FWSM

Figure 5-16 Campus Service Module Design with Catalyst 6500 Switch and FWSM

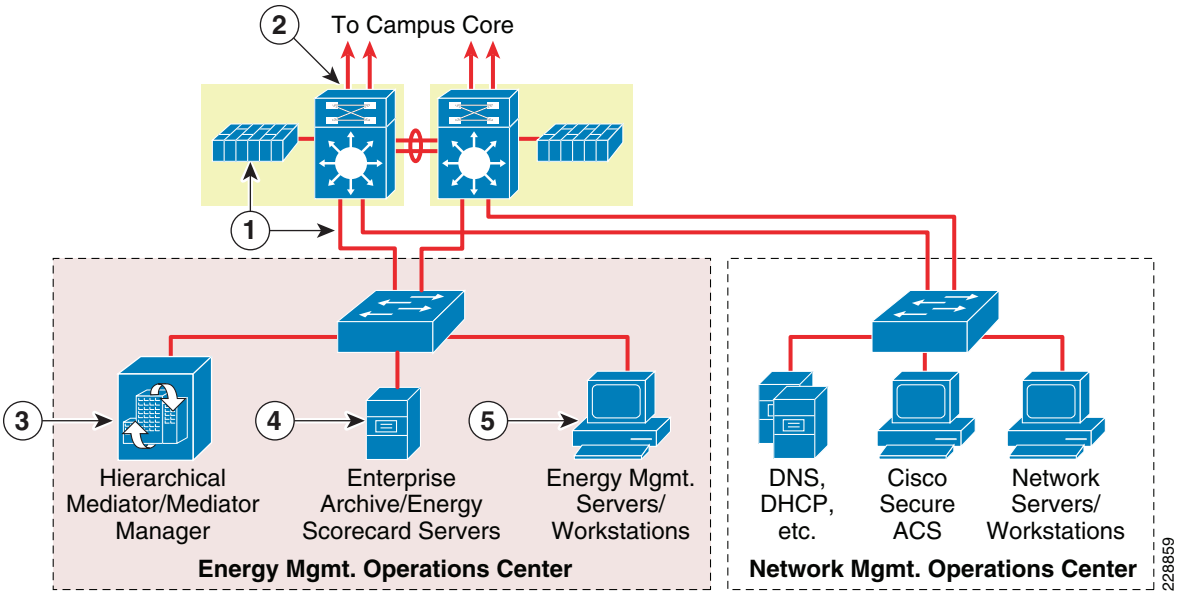

The following describes the numbers in Figure 5-16:

- **1**—EMOC VLAN trunked through Catalyst 6500 Distribution Switch to the FWSM.
- **2**—FWSM in the Catalyst 6500 Distribution Switch provides stateful access control to and from the EMOC devices.
- **3**—Hierarchical Mediator functionality can provide a real-time single point of access to energy usage information (datapoints) from Mediators deployed throughout the enterprise network.
- **4**—Internal enterprise archive servers may collect periodically exported datapoint information from individual Mediators deployed throughout the network infrastructure. An energy scorecard service may make historical energy usage information available internally within the organization.
- **5**—Internal enterprise management workstations may be used to configure the hierarchical and/or remote Mediators; deploy logical control applications to the hierarchical and/or remote Mediators; and to create, monitor, and deploy websites to the hierarchical and/or remote Mediators.

In this example, the EMOC VLAN is trunked from the access switch to a Layer-3 interface on the FWSM module. However, the FWSM is located within the Catalyst 6500 Distribution Switch within the Campus Service Module, as opposed to a separate Catalyst 6500 Service Switch. The FWSM provides stateful access control to and from the EMOC VLAN. Again, note that there are many different ways in which firewalling of the EMOC VLAN could be achieved within the Campus Service Module, because the FWSM supports both transparent (Layer 2) and routed mode (Layer 3) firewalling, high-availability through active/active or active/standby firewall configurations, and single or multiple context (virtual firewalls) mode. This example shows only one highly simplified method using a single context, routed mode firewall in an active/standby configuration.

As with the data center design, an alternative to the FWSM design is to implement a set of ASA 5500 Security Appliances within the Campus Service Module, as shown in Figure 5-17.

Figure 5-17 Non-VRF Campus Service Module Design with Catalyst Distribution Switch and ASA 5500 Security Appliance

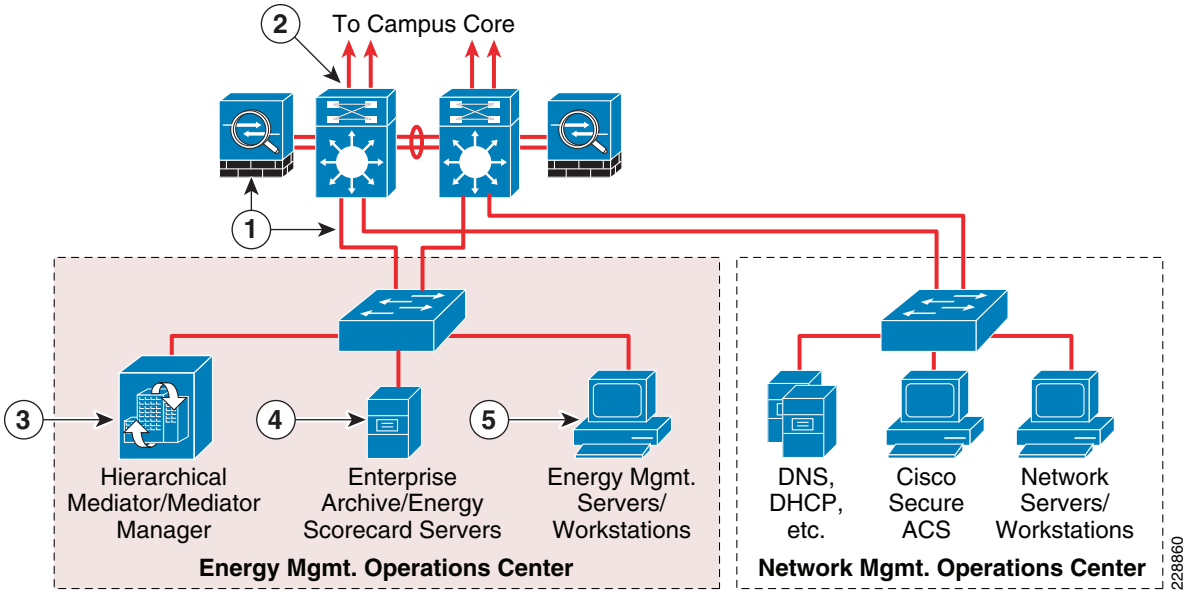

The following describes the numbers in Figure 5-17:

- **1**—EMOC VLAN trunked through Catalyst Distribution Switch to the ASA 5500 Security Appliance.
- **2**—ASA 5500 Security Appliance provides stateful access control to and from the EMOC devices.
- **3**—Hierarchical Mediator functionality can provide a real-time single point of access to energy usage information (datapoints) from Mediators deployed throughout the enterprise network.
- **4**—Internal enterprise archive servers may collect periodically exported datapoint information from individual Mediators deployed throughout the network infrastructure. An energy scorecard service may make historical energy usage information available internally within the organization.
- **5**—Internal enterprise management workstations may be used to configure the hierarchical and/or remote Mediators; deploy logical control applications to the hierarchical and/or remote Mediators; and to create, monitor, and deploy websites to the hierarchical and/or remote Mediators.

In this design, the EMOC VLAN is trunked from the access switch, through the Campus Service Module Distribution Switch, to a Layer-3 interface off of the ASA 5500 Security Appliance. One advantage of this design is that a wider range of switches (Catalyst 6500 Series, Catalyst 4500 Series, or even the Catalyst 3750 Series switch stack) can be used as the Campus Service Module switch. The ASA 5500 provides stateful access control to and from the EMOC VLAN. Again, note that there are many different ways in which firewalling of the EMOC VLAN could be achieved within the Campus Service Module, because the ASA 5500 Series supports both transparent (Layer 2) and routed mode (Layer 3) firewalling, high-availability through active/active or active/standby firewall configurations, and single or multiple context (virtual firewalls) mode. This example shows only one highly simplified method using a single context, routed mode firewall in an active/standby configuration. Individual interfaces or VLAN sub-interfaces could be used on the ASA 5500. Finally, as with the data center designs, other VLANs such as a Network Operations Center (NOC) VLAN can also be supported off of the same Campus Service Module.

VRF Campus Service Module Designs

Figure 5-18 shows an example of a Campus Service Module which makes use of VRFs to isolate the EMOC to a separate energy management VRF called the Building Infrastructure Network (BIN) VRF.

Figure 5-18 VRF Campus Service Module Design

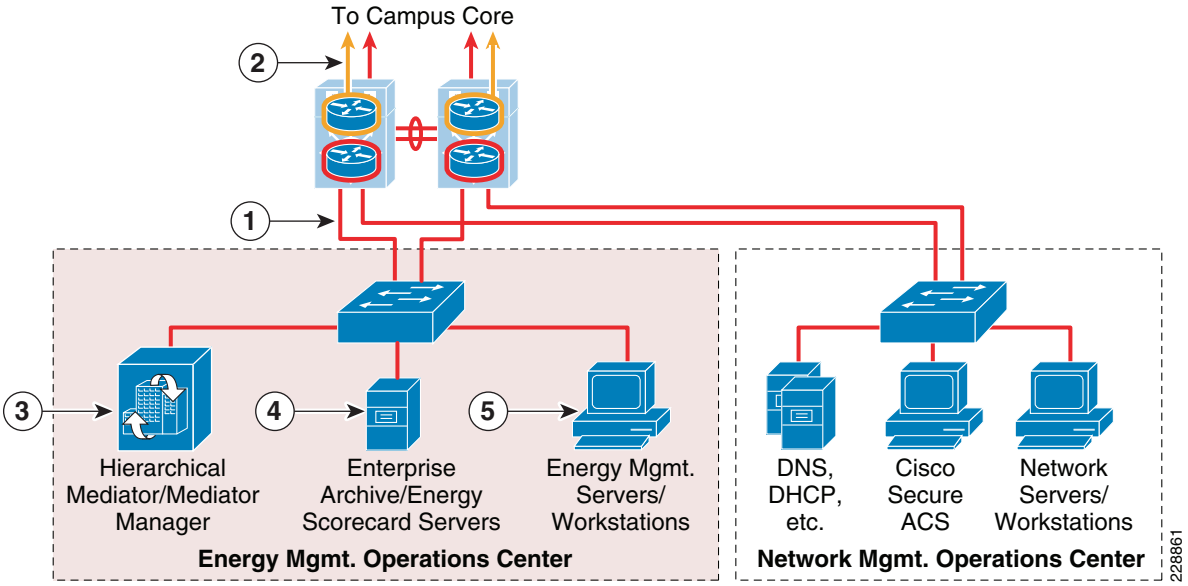

The following describes the numbers in Figure 5-18:

- **1**—EMOC VLAN trunked to the Catalyst 6500 Service Switch.
- **2**—GRE tunnels from Partner Extranet Module, Campus Building Modules, and WAN Edge Module terminate on Catalyst 6500 Campus Service Switches; extending the BIN VRF across the campus network to these locations.
- **3**—Hierarchical Mediator functionality can provide a real-time single point of access to energy usage information (datapoints) from Mediators deployed throughout the enterprise network.
- **4**—Internal enterprise archive servers may collect periodically exported datapoint information from individual Mediators deployed throughout the network infrastructure. An energy scorecard service may make historical energy usage information available internally within the organization.
- **5**—Internal enterprise management workstations may be used to configure the hierarchical and/or remote Mediators; deploy logical control applications to the hierarchical and/or remote Mediators; and to create, monitor, and deploy websites to the hierarchical and/or remote Mediators.

In this example, a separate EMOC VLAN is implemented within the Campus Service Module. The EMOC VLAN is trunked from the access switch to a Layer-3 interface defined within the Catalyst 6500 Service Switch. The Layer-3 interface sits on the BIN VRF, separate from the global VRF which houses the rest of the enterprise data network. GRE tunnels are then defined on the Catalyst 6500 Service Switch, which extend the BIN VRF out into the campus network. Figure 5-19 shows a more detailed view, without redundancy for clarity of the drawing.

Figure 5-19 Detailed Example Campus Service Module Design Using VRF-Lite with GRE Tunnels

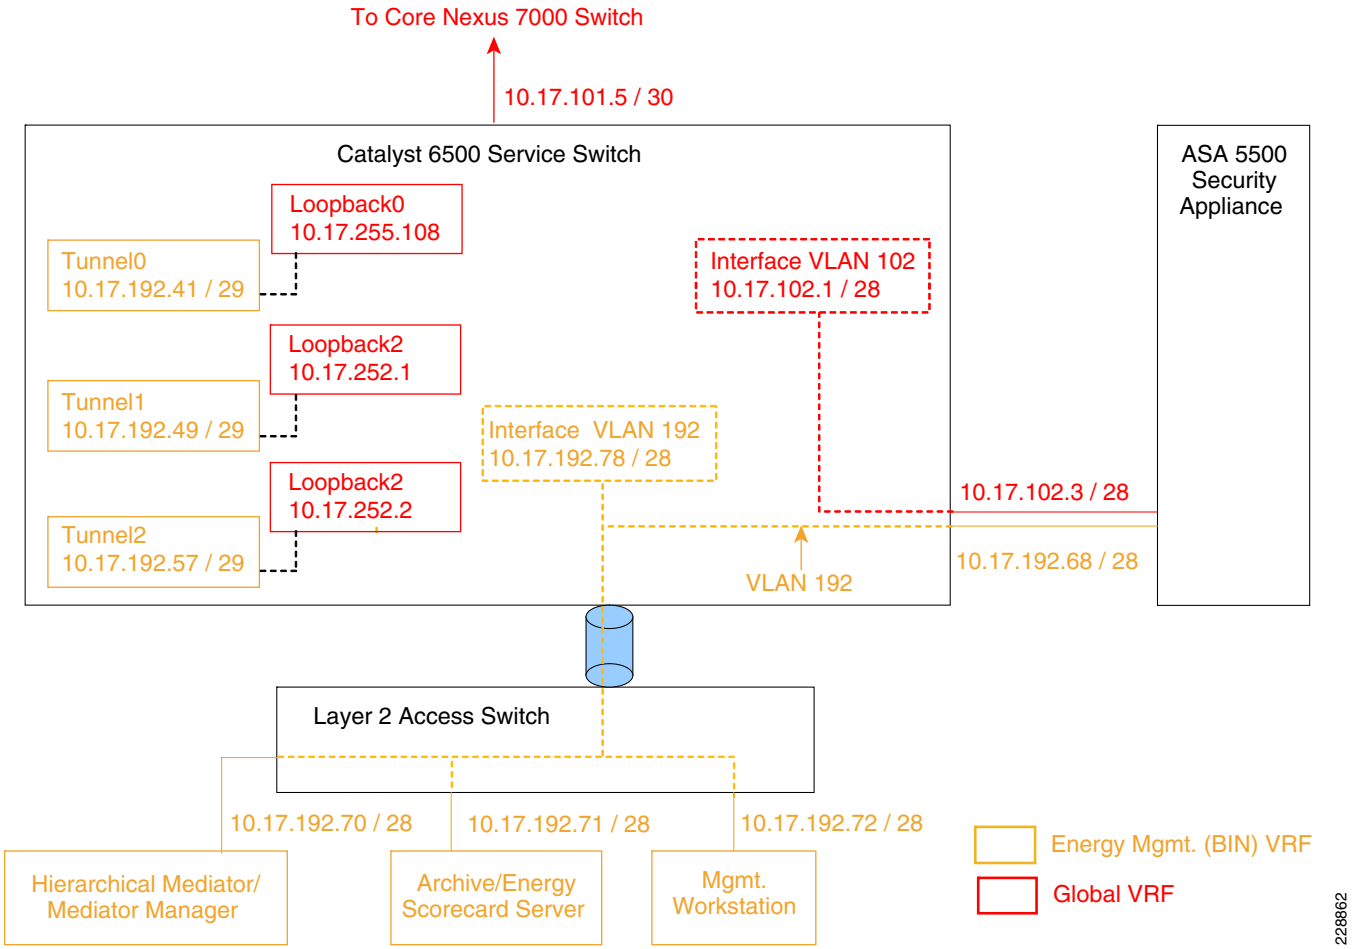

As with the Data Center Service Module VRF design, access control between the BIN VRF and global VRF can be accomplished via a ASA 5500 Security Appliance deployed within the Campus Service Module. Again, the use of remote-access IPSec VPN on the ASA 5500 can be implemented as an added security measure for enterprise client PCs who require direct access to either the hierarchical Mediator, or remote Mediators deployed throughout the network infrastructure.

CHAPTER 6

Branch Design Considerations

Branch Network Design Considerations

When designing the network to support one or more Mediators within a branch location, the design engineer must first determine the energy management solution deployment model. As mentioned previously, installations involving a managed service provider (MSP) partner often require VPN access to and from the Mediators within the branch back to the MSP network. In this deployment model, the VPN connectivity may be provisioned directly to each branch, referred to as distributed VPN connectivity. Alternatively, VPN connectivity may be centrally provisioned at a campus location and access allowed through the enterprise network to each branch. Each of these options is discussed separately.

Distributed VPN Connectivity Designs

For small energy management deployments involving only a handful of locations, provisioning separate VPN connectivity to each branch location may be acceptable. VPN access can take the form of a dedicated Catalyst switch separated from the existing branch IT network through a Cisco ASA 5500 Series Security Appliance that provides both VPN termination and stateful firewalling. An example of this design for a medium-sized branch site is shown in Figure 6-1.

Figure 6-1 Branch Design with Dedicated Switch and ASA 5500 Security Appliance

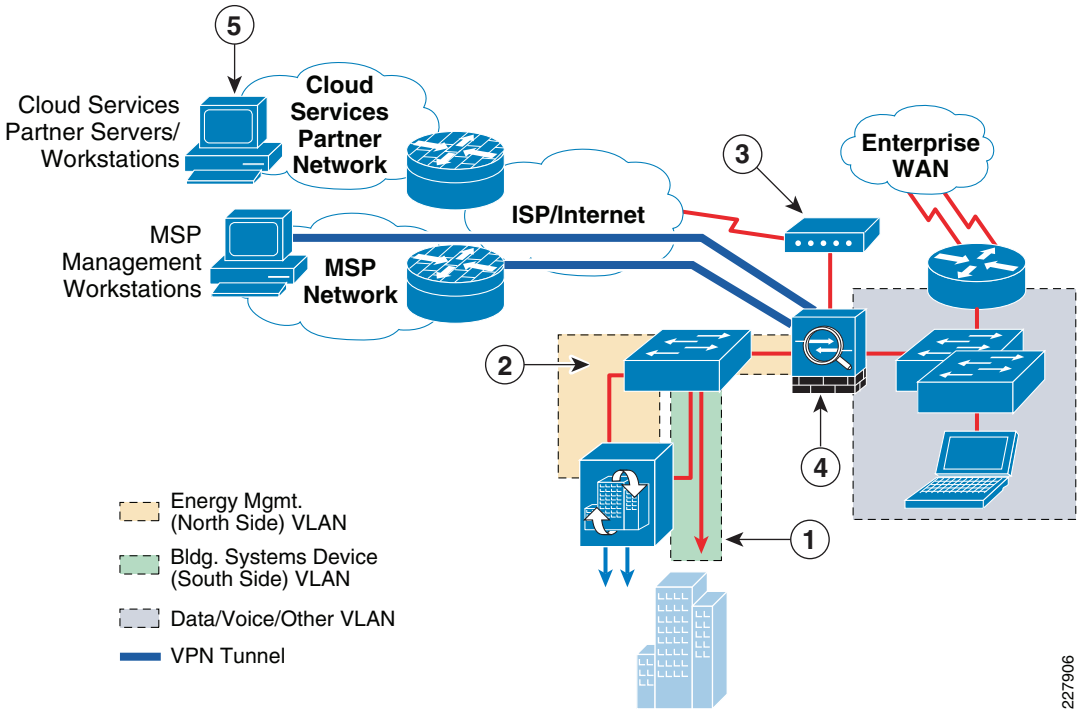

The following describes the numbers in Figure 6-1:

- **1**—Building systems device VLAN isolated by not trunking it to the dedicated ASA 5500 Security Appliance.
- **2**—Dedicated ASA 5500 Security Appliance connected to the energy management VLAN.
- **3**—Separate Internet Service (DSL, cable, etc.) provisioned for MSP partner VPN connectivity (site-to-site or remote access).
- 4—Dedicated ASA 5500 Security Appliance provides both MSP partner VPN termination and stateful access control to and from the energy management VLAN.
- **5**—Periodic export of logged data to cloud services partners, and client PCs accessing cloud services such as energy scorecards; may be provided via the Internet access at the branch or via backhauled to the campus Internet edge.

Although this type of implementation provides a high degree of isolation and access control, the duplication of the switch infrastructure, dedicated VPN router, and firewall appliance results in higher hardware and ongoing maintenance costs. The preferred approach is to provision separate VLAN segments on the existing branch Catalyst switch platform for energy management. A single Cisco ISR branch router can provide both the WAN access to the branch from the enterprise campus network, as well as the VPN access from the MSP network (with appropriate software image and licensing). An example of this type of design for a medium-sized branch site is shown in Figure 6-2.

Figure 6-2 Branch Design with VPN Access and Integrated Infrastructure

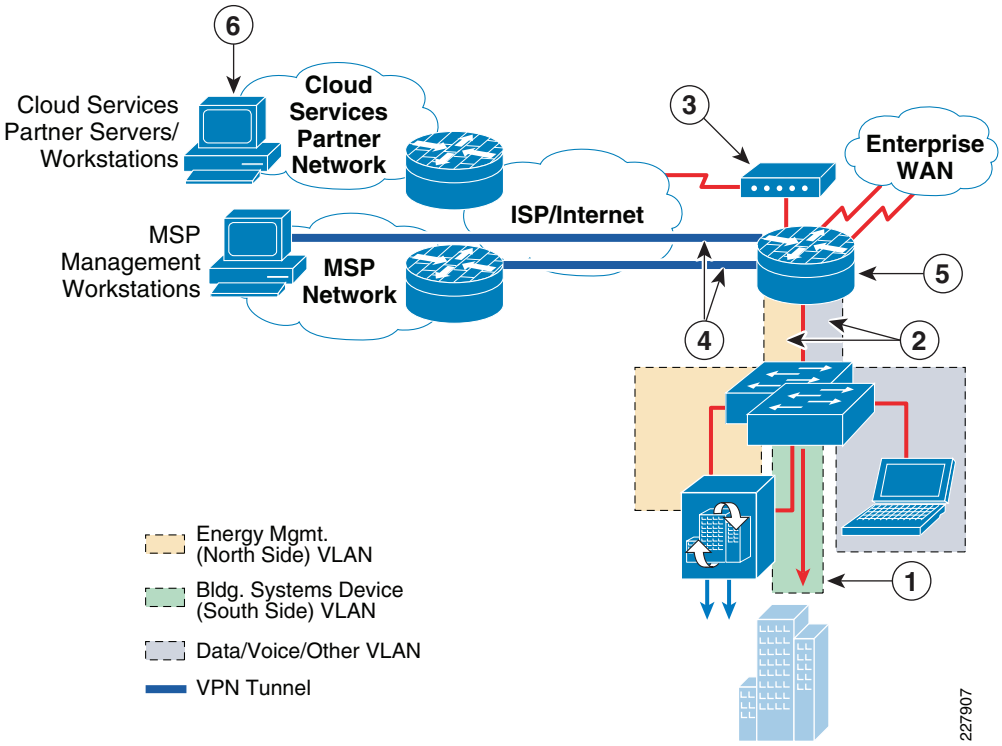

The following describes the number in Figure 6-2:

- **1**—Building systems device VLAN isolated by not trunking it to the branch router.
- **2**—The energy management VLAN and the data/voice/other VLAN trunked to the branch router.
- **3**—Separate Internet service (DSL, cable, etc.) may need to be provisioned for partner VPN connectivity.
- **4**—Branch router with IPSec VPN software provides termination of site-to-site or remote-access VPN from MSP partner.
- **5**—CBAC or ZBPF provide stateful access control to and from the energy management VLAN.
- **6**—Periodic export of logged data to cloud-services partnes, and client PCs accessing cloud services, such as energy scorecards, may be provided via the Internet access at the branch or via backhauled to the campus Internet edge.

In cases where an IPSec VPN provides the enterprise WAN connectivity, an additional VPN tunnel may be provisioned to the MSP network. In the case where private enterprise WAN connectivity is provisioned, a separate Internet connection can be provisioned on the Cisco ISR branch router, as is shown in Figure 6-2. Either CBAC or ZBPF running within the branch ISR router can be used to provide stateful access control between the energy management systems VLANs and the rest of the enterprise network and the MSP. The branch Internet connectivity could also be used to directly support the periodic export of logged data points to a cloud-services partner, as well as allow client PCs access to cloud services such as energy scorecards. However, default routing issues at the branch may limit its applicability. The network administrator may instead choose to backhaul such traffic across the enterprise WAN to the campus Internet edge.

In the designs shown in Figure 6-1 and Figure 6-2, access control should be specified down to the IP addresses and protocols required between the MSP partner and/or cloud-services partner devices and the branch Mediator. Specific protocols required are discussed Chapter 2, "Deployment Models and Information Flows." In small deployments where VPN access is provisioned to each branch, the enterprise customer may completely outsource the management of the energy management solution to the MSP. In such cases, the enterprise customer may not deploy an Energy Management Operations Center (EMOC), as discussed in Chapter 5, "Data Center/Campus Service Module Design Considerations." Therefore, no inbound access from enterprise energy management workstations may be needed. Likewise no outbound access from the mediators to enterprise archiving servers for periodic exporting of logged datapoint information may be needed with a small deployment.

Even with a small deployment, enterprise client PC access within the branch to energy usage data may still be a requirement. As discussed in Chapter 5, "Data Center/Campus Service Module Design Considerations,"access to energy usage data can be accomplished in the following ways:

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

If business needs only require non-real-time access to historical energy usage information, enterprise client PCs may only need access outbound to the Internet in order to reach a cloud services partner energy scorecard server. This can be allowed directly at the branch, or backhauled across the corporate WAN (if a corporate WAN exists within a small deployment) to the Internet edge within the campus. Note that many enterprise organizations backhaul traffic across the corporate WAN to the Internet edge simply to implement a single point of access for all employee traffic to the Internet for monitoring and control purposes, as well as ease of implementing routing tables.

In cases where real-time access to energy usage information is needed, or the enterprise client PCs may need to override setpoints, direct access to the Mediators may be required, although this poses a greater security risk, since the Mediators may be controlling building systems devices. Again, limited access through a hierarchical Mediator deployed within an EMOC provides a more secure and single point of access to the energy management solution. However, in small deployments, if no hierarchical Mediator is deployed within an EMOC, then access may need to be allowed directly between enterprise client PCs and the branch Mediator. Where possible, access should be limited to a subset of PCs, based on IP address, and restricted to the use of secure protocols such as HTTPS. Userids and passwords should be implemented on the Mediator to restrict access as well. The network administrator may consider the deployment of remote access VPN technology on the ASA 5500 Security Appliance or ISR router as an additional security mechanism for enterprise client PC access as well. Note that with the deployment of individual VPN access to each branch, a centralized AAA service is still recommended. However, with small deployments, the enterprise customer may not have a central AAA server. Therefore, individual ASA 5500 and/or ISR routers deployed within each branch may hold local databases for MSP and enterprise client PC access to the energy management solution within that branch.

Comparing both dedicated branch VPN designs shown in Figure 6-1and Figure 6-2, the design shown in Figure 6-2 results in lower hardware and ongoing maintenance costs, but the management and reoccurring costs of an additional VPN connection for each branch location may still prohibit the scaling of this implementation.

Centralized VPN Connectivity Designs

For large energy management solution implementations, centralizing the MSP VPN connectivity to a campus or data center location provides a much more scalable and manageable deployment. The headend design for this type of energy management solution is discussed in Chapter 4, "Internet Edge Design Considerations." Figure 6-3 shows an example of a large branch design for support of the energy management solution that uses the centralized VPN model.

Figure 6-3 Example Large Branch Site Design

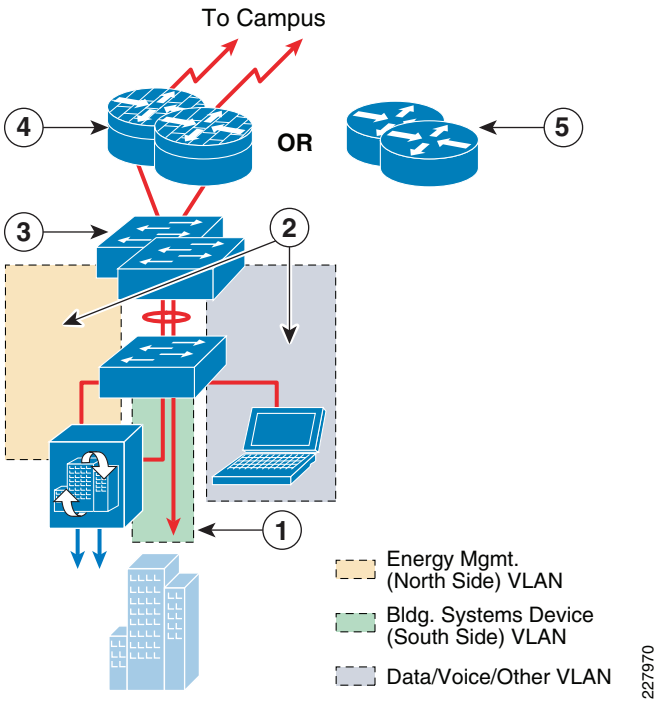

The following describes the numbers in Figure 6-3:

- **1**—The building systems device VLAN is isolated by not trunking it from the Layer-2 access switch to the Layer-3 distribution switch stack.
- **2**—The energy management VLAN and the data/voice/other VLAN are trunked to the Layer-3 distribution switch stack.
- **3**—ACLs on the Layer-3 distribution switch stack provide stateless access control to the energy management VLAN from other VLANs within the branch.
- **4**—Routers with CBAC or ZBPF provide stateful access control from the rest of the enterprise network to the energy management VLAN.
- **5**—Alternatively routers with ACLs provide stateless access control from the rest of the energy network to the energy management VLAN.

Large branch sites often implement distribution and access-layer switches for scalability, similar to a small campus building design. The distribution layer may consist of a Layer-3 Catalyst 3750 Series switch stack, while the access layer consists of Layer-2 Catalyst 2900 Series switches. In this design, a separate energy management VLAN (north side) and a separate building systems device VLAN (south side) are 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 system's

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.

Within the branch, access to and from the Mediator is controlled via ACLs applied to switched virtual interface (SVI) defined for the energy management VLAN on the Layer-3 distribution switch. Access to and from the Mediator from the devices within the MSP network, as well as the enterprise Energy Management Operations Center (EMOC) located within the campus, can further be controlled via the branch router. When stateful firewalling is desired or required, either CBAC or ZBPF can be run on the branch router. Alternatively, stateless access control can be accomplished via ACLs applied on the branch ISR router.

Figure 6-4 shows an example of a medium-sized branch site design for support of the energy management solution.

Figure 6-4 Example Medium Branch Site Design

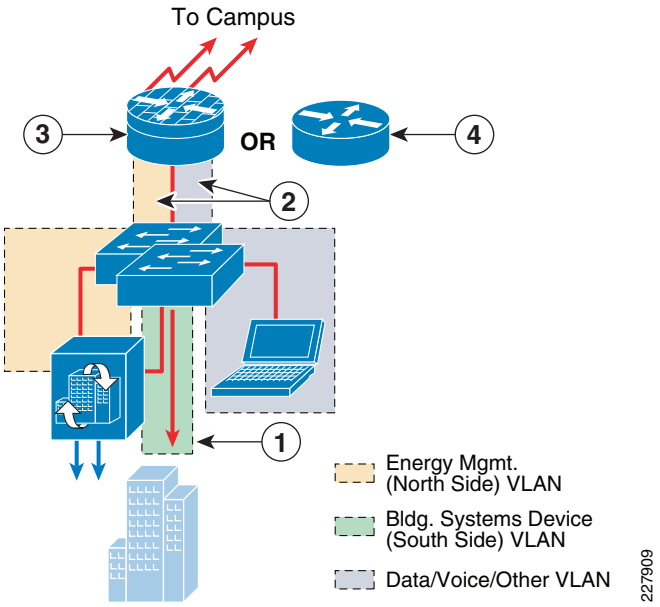

The following describes the numbers in Figure 6-4:

- **1**—Building systems device VLAN isolated by not trunking it to the branch router.
- **2**—The energy management VLAN and the data/voice/other VLAN trunked to the branch router.
- **3**—Branch router with CBAC or ZBPF provide stateful access control to and from the energy management VLAN.
- **4**—Branch router with ACLs provide stateless access control to and from the energy management VLAN.

Medium-sized branch sites often consist of just a branch ISR router and a Layer-2 Catalyst 2900 Series switch stack functioning as the access layer. In this design, a separate energy management VLAN (north side) and a separate building systems device VLAN (south side) are again provisioned on the Layer-2 access switch. The energy management VLAN, along with any data/voice/other VLANs, are trunked to a branch ISR router. However, the building systems device VLAN is not trunked, effectively isolating it within the access-layer switch stack. 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. Example 6-1 shows a partial configuration of a Catalyst 2960 switch with this design.

Example 6-1 Partial Catalyst 2960 Configuration Showing VLANs For Mediator Management and Building System Devices

```
!
interface FastEthernet0/18
description CONNECTION TO MEDIATOR MGMT INTERFACE
switchport access vlan 192                    ! Mediator management VLAN
!
interface FastEthernet0/19
description CONNECTION TO MEDIATOR BLDG SYSTEMS INTERFACE
switchport access vlan 193                    ! Mediator building device VLAN
!
~
interface GigabitEthernet0/1
description TRUNK TO ME-WESTRICH-1
switchport trunk allowed vlan 1-192,194-4094 ! VLAN 193 not trunked to router
switchport mode trunk
srr-queue bandwidth share 1 30 35 5
priority-queue out
mls qos trust
!
```

Within the branch, access to and from the Mediator is controlled via the branch ISR VLAN interfaces. When stateful firewalling is desired or required, either CBAC or ZBPF can be run on the branch ISR router. Alternatively stateless access control can be accomplished via ACLs applied on the branch ISR router energy management VLAN interface.

Figure 6-5 shows an example of a small branch site design for support of the energy management solution.

Figure 6-5 Example Small Branch Site Design

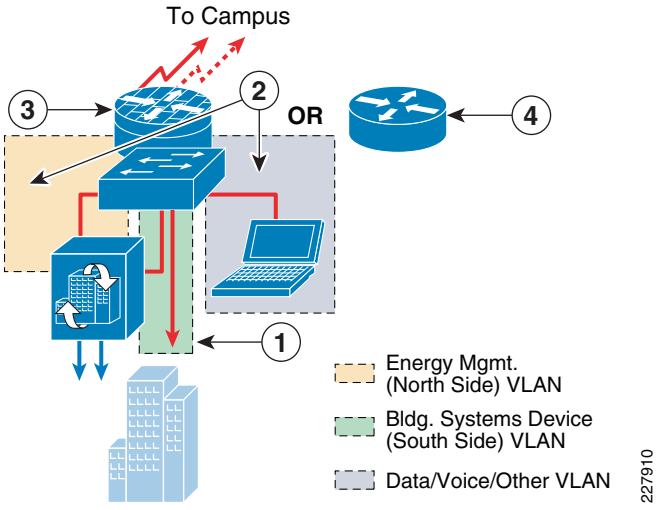

The following describes the numbers in Figure 6-5:

- **1**—Building systems device VLAN isolated by not defining a Layer-3 VLAN interface on the branch router, or by not trunking it to the branch router.
- **2**—The energy management VLAN and the data / voice / other VLAN extended to the branch router by defining Layer-3 VLAN interfaces on the branch router.
- **3**—Branch router with CBAC or ZBPF provides stateful access control to and from the energy management VLAN.
- **4**—Branch router with ACLs provide stateless access control to and from the energy management VLAN.

Small branch sites often consist of a branch ISR router in which the Layer-2 switch has been collapsed into the router as a switch module. The overall energy management design is effectively the same as if a standalone Catalyst switch were deployed.

Extending VRFs to the Branch

The deployment of network virtualization for energy management systems can provide the additional advantage of path isolation of the energy management solution traffic across the IP network infrastructure. When applied to the branch, an energy management Virtual Routing and Forwarding (VRF) instance is extended into the Layer-3 device. An example of this for the medium branch site design is shown in Figure 6-6.

Figure 6-6 Example of an Energy Management VRF Extended to the Branch

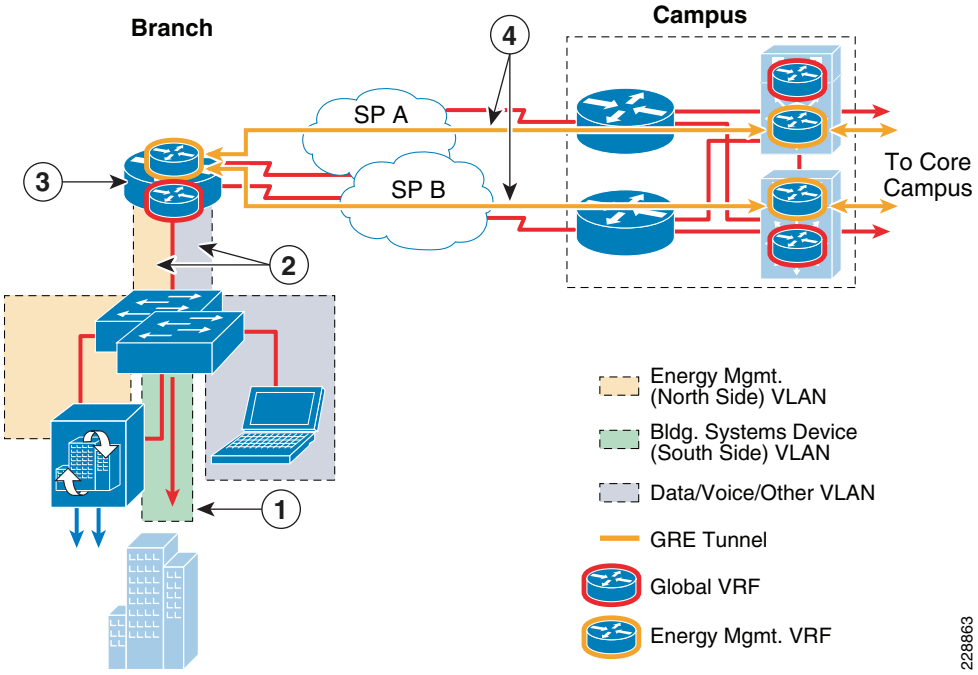

The following steps describe what occurs in Figure 6-6:

- **1**—Building systems device VLAN isolated by not trunking it to the Layer-2 distribution switch stack.
- **2**—The energy management VLAN and the data/voice/ other VLAN trunked to the ISR router.

- **3**—Energy management VLAN mapped to the energy management VRF, while data / voice / other VLANs mapped to the global VRF within the branch router.
- **4**—VRFs extended to the campus via GRE tunnels from the branch router to the Layer-3 distribution switches of the campus WAN module.

In this example, the energy management VLAN is defined on the Layer-2 access switch and trunked to the ISR router, where the Layer-3 interface for the energy management VLAN is defined. The VLAN is then mapped to an energy management VRF that is separate from the global VRF which supports the data/voice/other VLANs. Example 6-2 shows a partial configuration of a Cisco 3845 ISR router in this configuration, but with only a single GRE tunnel.

Example 6-2 Partial Branch ISR Router Configuration with VRF Implementation

```
!
ip vrf bin
! Creates Building Infrastructure Network (BIN) VRF
rd 192:118
!
~
!
interface Loopback2
description LOOPBACK INTERFACE FOR BIN GRE TUNNEL
ip address 10.17.252.9 255.255.255.252
!
interface Tunnel0                                    ! GRE tunnel extending BIN VRF to the
campus
description VRF FOR MEDIATOR NETWORK TO ME-WESTDIST-1
ip vrf forwarding bin
ip address 10.17.192.25 255.255.255.248
tunnel source Loopback2
tunnel destination 10.17.100.10
!
interface GigabitEthernet0/0
no ip address
load-interval 30
duplex auto
speed auto
media-type rj45
!
interface GigabitEthernet0/0.182                    ! VLAN sub-interface for data network
in global VRF
description Employee PCs
encapsulation dot1Q 182
ip address 10.18.2.1 255.255.255.0
!
interface GigabitEthernet0/0.192                     ! VLAN sub-interface for Mediator
mgmt network
description BRANCH MEDIATOR MGMT NETWORK
encapsulation dot1Q 192
ip vrf forwarding bin                                ! Places VLAN 192 into the BIN VRF
ip address 10.17.192.17 255.255.255.248
!
~
!
interface Serial3/0                                  ! WAN interface
ip address 192.168.64.25 255.255.255.252
load-interval 30
dsu bandwidth 44210
max-reserved-bandwidth 100
!
```

Because the traffic within the energy management VRF is isolated from traffic in other VRFs, stateful firewalling is not really required within the branch ISR itself. However, inbound and outbound ACLs may still be applied to the energy management VLAN in order to restrict access to the Mediators, if desired. The centralized VPN connectivity from the MSP to all of the enterprise energy management systems, along with the deployment of a separate energy management VRF, can provide complete path isolation of the energy management systems traffic and require only a centralized security policy.

Note The network administrator may want to consider defining a VRF not just for the energy management solution, but also to support other solutions, such as IP video surveillance and physical access control. In this scenario, a single Building Infrastructure Network (BIN) VRF may be defined. This eases the administrative burden of not having to configure and administer as many VRFs within the network infrastructure.

From the branch, the energy management VRF can be extended across the WAN via GRE tunnels. This method is referred to as the VRF-Lite with GRE Tunnel model. GRE tunnels can be defined from the branch ISR router to the Layer-3 distribution switches within the campus WAN Module (discussed in Chapter 3, "Campus Design Considerations"). The GRE tunnels are then mapped to the energy management VRF. These tunnels support both MSP partner VPN management as well as the periodic export of logged data to the Internet via the Campus Partner Extranet Module. Note that other methods of extending VRFs across the WAN exist as well; for example, the mapping of VRFs to an MPLS service. These have not been evaluated for this vision of the design guide.

When deploying a VRF for path isolation of the energy management solution, it is recommended that where possible, enterprise client PC access into the energy management solution should be centralized to one or more strategic locations, such as the EMOC or Partner Extranet Module. In other words, direct connectivity provided within the branch itself from enterprise client PCs to the Mediator should be avoided. Chapter 5, "Data Center/Campus Service Module Design Considerations" discusses how this may be achieved. This may include the use of a data center firewall between the energy management (BIN) VRF and the global VRF. Alternatively, the network administrator may consider the use of remote-access VPN technology within the data center firewall or via the Partner Extranet Module to provide access from enterprise client PCs into the energy management (BIN) VRF in order to access a hierarchical Mediator located within the EMOC or direct access to the branch Mediators. As before, if access to non-real-time energy usage information is all that is needed for enterprise client PCs, the network administrator should consider using an energy scorecard either via a cloud-services partner or internally deployed (possibly through a partner), as a more secure alternative to direct access to the Mediators.

QoS Within the Branch

A secondary function of the branch access switches is to provide classification and marking of Cisco Network Building Mediator traffic flows as they enter the network at the branch. The function of classification and marking is essentially the same at the branch as within the campus, but with different Catalyst switch platforms. Chapter 3, "Campus Design Considerations" presented the following two methods:

- Identifying and marking individual traffic flows from the Mediator to different service classes based on 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.

Using the first method above, classification and marking of traffic from the Mediators can be accomplished via an ingress policy-map which includes ACLs, applied to the access switch port to which the Mediator is connected. 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 6-3 extends the partial configuration shown in Example 6-1 to also include an example QoS configuration a Catalyst 2960 switch using this method.

Example 6-3 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                  ! QoS policy map
class DATA_EXPORT_TRAFFIC
set ip dscp af11                              ! Sets export traffic as AF11 (Bulk Data
Service Class)
class MGMT_TRAFFIC
set ip dscp cs2                               ! Sets mgmt traffic as CS2 (OAM Service
Class)
class class-default
set ip dscp default                           ! Marks all other traffic to default (best
effort)
!
~
interface FastEthernet0/18
description CONNECTION TO MEDIATOR MGMT INTERFACE
switchport access vlan 192
service-policy input MEDIATOR_ENDPOINT        ! Applies ingress QoS service policy to
Mediator Management VLAN
!
!
interface FastEthernet0/19             ! Ingress traffic set to CoS 0 (best effort) on
Building Systems VLAN
description CONNECTION TO BUILDING DEVICES VLAN
switchport access vlan 193
!
~
interface GigabitEthernet0/1
description TRUNK TO ME-WESTRICH-1
switchport trunk allowed vlan 1-192,194-4094
switchport mode trunk
srr-queue bandwidth share 1 30 35 5
priority-queue out
mls qos trust dscp
! DSCP values preserved on the trunk port
!
~
!
ip access-list extended MEDIATOR_EXPORT     ! Identifies FTP control and data
permit tcp any any eq ftp
permit tcp any any range 49152 49153
ip access-list extended MEDIATOR_MGMT       ! Identifies other management traffic
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 this example, the periodic data exports generated by the Mediator use the FTP protocol, which is identified and marked as AF11, corresponding to the bulk data service class. The remaining management protocols generated by the Mediator (which includes the return traffic from sessions initiated management workstations) are identified and marked as CS2, corresponding to the Operations, Administration, and Management (OAM) service class. Any other traffic is identified and marked as default, corresponding to the best-effort service class. The example above also shows that the Building Systems Device VLAN interface is set to not trust inbound DSCP markings, and no policy-map is applied to the interface. Therefore, all inbound building systems device traffic to the switch is classified as default or best-effort. The trunk port that connects the Catalyst 2960 switch to the ISR router is set to trust DSCP markings. This preserves traffic markings inbound across the trunk to the switch.

Note that this is just an example of using this method of identifying traffic. Actual deployments may vary in terms of the export protocol and management protocols supported. Alternatively, the network administrator may choose to mark all ingress traffic on the Mediator Mgmt VLAN port to a single DSCP value such as CS2, corresponding to the OAM service class. Note also that only part of the full QoS configuration of the switch is shown in the example above.

CHAPTER 7

Operations Energy Management

Energy Consumption Awareness

Reducing energy consumption in buildings can be accomplished in a step-wise manner. Efforts such as improving window energy efficiency, insulation, adding occupancy sensors and lighting controls can provide an initial sustainable reduction, but these changes alone will not result in a sustainable reduction such as the 30 percent reduction mandated by executive order 13423 for public sector entities.

To achieve and be able to sustain this level of energy reduction takes a systematic approach. The solution needs to achieve the following requirements:

- Accessing data from all of the energy consuming systems in the building.
- Managing those systems based upon the data to reduce energy use.
- Correlating data between systems to find new ways to lower energy use.
- Ability to integrate new Energy Star systems to reduce energy use.
- Ability to integrate alternative energy systems, if applicable, to reduce energy use.
- Ability to provide the information to accommodate new regulations to reduce carbon.
- Ability to be managed to work with Utility Company Demand Response programs to reduce energy costs.

An energy management solution that provides the proper subset of these requirements might be able to lower energy usage by 30 percent. To achieve the 30 percent goal and be able to sustain it and continue to meet new regulations and standards, the Energy Management solution must meet all of these requirements.

Cisco's Network Building Mediator provides access to all of the energy consuming systems within a building. Many of these systems have unique data protocols and unique data formats, making access alone a major hurdle. The Mediator normalizes this data and places it into standard XML format so that it is accessible by a broad range of applications that can use that data to manage energy use.

The sustainability of an energy management system is based on its ability to access energy data from all of the energy consuming devices in a building. Another dimension of sustainability is through energy visibility to the building occupants. Changing building occupants behavior is critical method to ensuring that energy reduction is sustainable. A comprehensive energy management solution is one that keeps building occupants informed about their impact on energy use and carbon footprint. See Figure 7-1.

Figure 7-1 Energy Management System

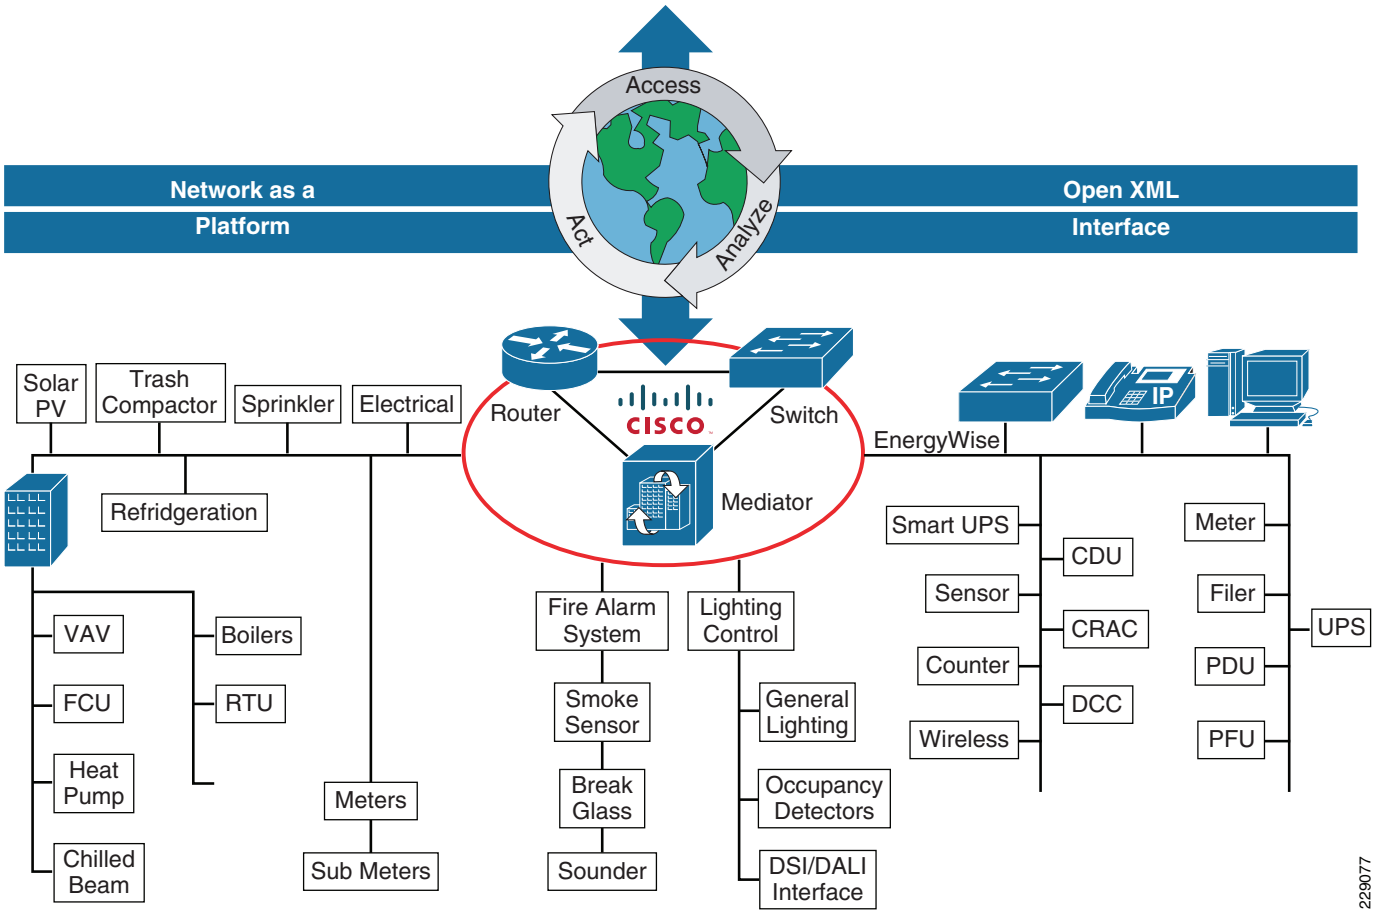

Energy Management Information Collection

Building systems and energy management applications that are able to collect energy, management, and usage information provide the foundation necessary for a comprehensive solution. The critical differentiator that enables building systems to be transformed from simple device management to an advanced level of building automation is a TCP/IP network with the ability to connect these devices to it. The Cisco Mediator connects systems and devices to the network, and facilitates the automation and data collection from these devices to energy management applications.

The Mediator

The Cisco Mediator is a system composed of both hardware and software that gathers and manipulates data from numerous sources. These sources are typically intelligent machines or sensors, found in virtually any facility, which are otherwise unable to intercommunicate. The Cisco Mediator allows communication to occur between these devices, facilitates additional processing, and provides a uniform presentation of this information to users (through web-browsers) and third-party applications such as the following:

- Connects BAS systems such as HVAC, Lighting, Electrical, Energy Management and Security with Unified Communications and other IP Applications
- Software and hardware integration platform and web server
- Embedded Linux OS
- Python Framework
- Multiple physical ports
- Software support for common open and numerous proprietary BAS protocols
- CISCO Unified Communications
- AJAX-based web page and website created by onboard tools
- Use XML-RPC

The physical connectivity functions and capabilities of the Mediator are extended by a suite of configuration and management software. The Mediator includes an extensive web-based management interface known as OMEGA. The OMEGA interface consists of operations tools and several separate applications as follows:

- Event Manager
- Security Manager
- Trend Manager
- Web Scheduler
- Web Express
- Website Builder

Two additional software applications are the Mediator configTOOL and Mediator perfectHOST. These are Windows software applications that enable graphical application programming and application monitoring. The Mediator applications communicate with the Mediator control system via two-way communication over a TCP/IP network. The Mediator collects data from multiple sources, including systems that use disparate protocols and are otherwise unable to intercommunicate, converting the data into a single, widely used format such as XML, and provides information to the end user in a uniform presentation. The Mediator enables building engineers to specify the type and quantity of information they want to receive, and omit anything they consider unessential. In addition, the Mediator is able to receive data through the perfectHost user application and convert it into commands and data that conforms to the protocols required by the facility's sensors, devices, and systems. The Mediator can run multiple services and protocols simultaneously and still have the capability to respond immediately to events generated in the network environment.

The Mediator Configuration Tool is an interactive, menu-driven software application that allows the user to specify the configuration parameters of the a mediator host system and all the connected systems including the control systems. See Figure 7-2.

Figure 7-2 Mediator Architecture

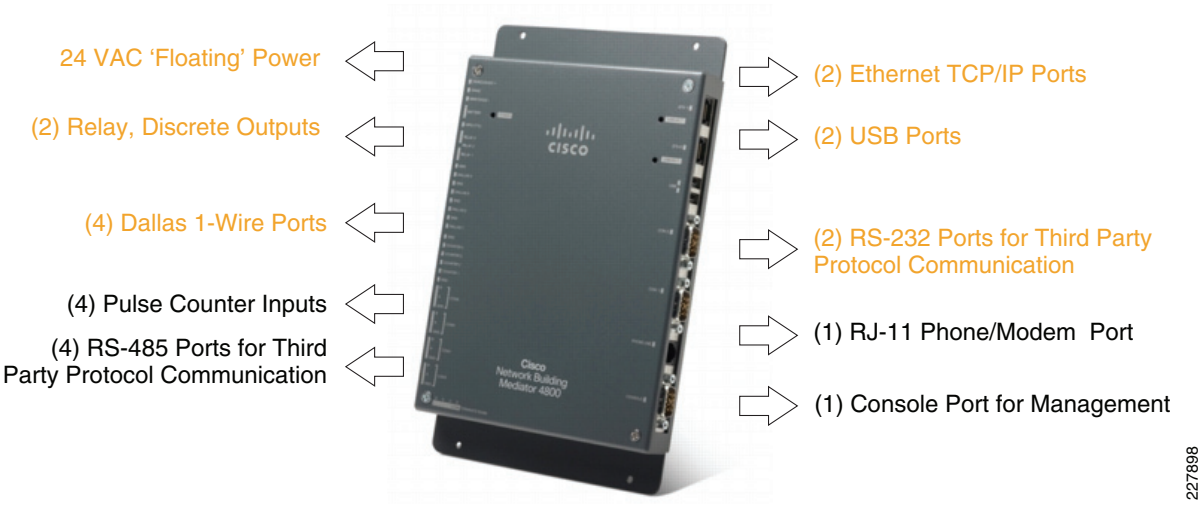

For more information on the Cisco Mediator, refer to the following URL:
http://www.cisco.com/en/US/products/ps10454/index.html

For the *Mediator End User Guide*, refer to the following URL:
http://www.cisco.com/en/US/products/ps10454/products_user_guide_list.html

Energy Management Application Partners

As collection of data for energy management operations within the Mediator itself is limited to only a few days, it is necessary to export and upload the data from individual Mediators to an external repository for the purpose of long term data collection, analysis, and trending of building systems. To collect information from many mediators covering hundreds of locations several methods are available, two of which are local databases in the enterprise and cloud-based services from Cisco partners.

Cloud-based services from Cisco partners provide added value and visibility based on their extensive experience across many different types of building systems and their respective idiosyncrasy's.

Out of the extensive list of solution partners, this validation evaluated products from Facilities Solutions Group (FSG), Noveda, and PreNova.

Facilities Solutions Group—Energy Scorecard for Operational Trend Reporting

FSG Energy's approach to the industry is to assemble "best of breed" open standard technologies, applications, and service providers. They have formed an alliance with Cisco for technology infrastructure and Sensus MI for automated HVAC Fault Detection and Diagnostics. These coupled with FSG's powerful delivery and infrastructure support allow companies to take advantage of the best combinations of technologies and applications, without being tied to one proprietary vendor.

Energy Scoreboard™

Managing energy and operations for multiple facilities requires connecting energy metering systems with building automation systems at the site/store level. By integrating these technologies a common platform is created upon the robust IT data backbone. Through centralized data aggregation and reporting, informed decisions can be made at the corporate level, keeping score on the greater energy management challenge. Energy Scoreboard™ (see Figure 7-3) is:

- An Internet-based energy management tool that is closely integrated with a control system allowing the enterprise to "call the energy plays" and immediately measure results.
- Shows how and when the enterprise uses energy through easy to view graphs, charts, tables, and reports.
- Provides the ability to compare energy use across multiple sites.
- Consists of a hardware device for data collection (the Cisco Mediator), an FTP server, database server, and web server (for web-based energy analysis).

Figure 7-3 Energy Scoreboard

For more information about FSG and their products, refer to the following URL:
http://www.fsgi.com/site/energy/products.php

Noveda—Energy Consumption Information to Building Occupants via Digital Media or Corporate Web

Noveda Technologies provides a web-based, dynamic, graphic visualization solution for real-time monitoring, diagnostics, metrics, and historical tracking of renewable energy, conventional energy, and building mechanical/environmental systems.

Their building projects include institutional, commercial, educational, governmental clients for which we provide a critical tool for dramatically improving energy efficiency; allowing substantial cost savings, reducing systems maintenance, decreasing a building's greenhouse gas emissions and carbon footprint while providing an exciting platform for communicating these efforts and achievements to customers, clients, employees, and the general public.

The Noveda systems not only measure energy use, but also analyze the information allowing building owners to better manage energy resources, improve energy and mechanical system performance, and lower energy and maintenance costs. It is understood that an immediate energy savings of 10 to 30 percent can be realized by taking simple measures. By monitoring and displaying building energy consumption on a real-time basis, opportunities for operational efficiency improvements arise as well as incentives for changes in occupant behavior in regards to energy use.

EnergyFlow Monitor™

For buildings that use conventional and/or alternative energy (electric, gas, steam, solar, wind), Noveda's real-time web-based EnergyFlow Monitor™ is the most powerful tool available to monitor a building's energy and natural resource use in order to lower energy costs and reduce the consumption of fossil fuels. It easily track the *before* and *after* effects of energy conservation efforts, such as high efficiency lighting and mechanical equipment retrofits. Receive notifications when peak demand thresholds are being reached so that billing rates can be controlled.

Studies[1] have consistently shown that the best way to save energy is to show people how much they are using. Noveda's real-time web-based EnergyFlow Monitor™ is the most powerful tool available to communicate building energy and natural resource use to effect change.

With engaging visuals and industry-leading monitoring frequency, EnergyFlow Monitor™ provides the context that is normally missing in energy consumption data. With meaningful graphics, analysis-based alerts, detailed reporting and complete data export flexibility, EnergyFlow Monitor™ gives answers, not just details.

Dynamic visualizations not only informative to building occupants, but engage them in helping achieve goals for energy and resource efficiency. Easily track the before and after effects of energy conservation efforts. EnergyFlow Monitor™ provides real-time energy in perspective, it tells occupants “How they are doing right now”, enabling timely action to be taken instead of waiting 30 to 45 days to see the impact on utility bills.

1. July 2009, European Commission's Directorate-General for Energy and Transport initiative, “*Energy Savings from Intelligent Metering and Behavioral Change (INTELLIGENT METERING)* http://www.managenergy.net/products/R1951.htm", 2009. http://www.noveda.com/resources/general/sales%20-%20energyflow-ehb093008.pdf

Figure 7-4 EnergyFlow Monitor

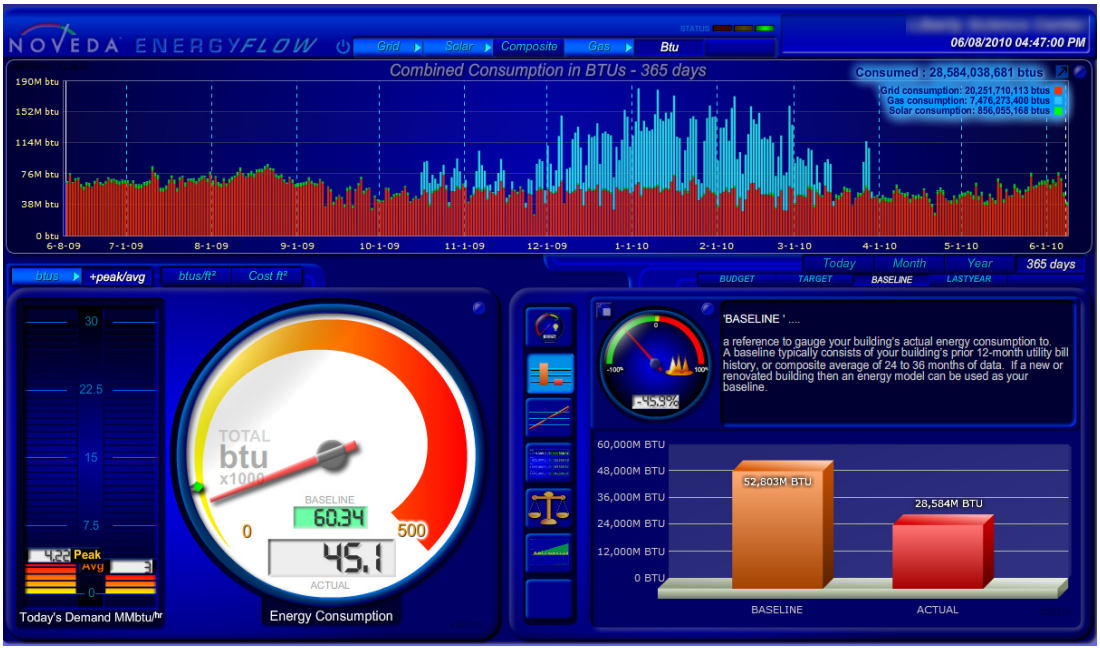

Prenova—Services Provider for Deployment, Support and Daily Operations

Prenova provides technology-based energy management services that help customers reduce utility costs and improve energy efficiency. The company works hand-in-hand with customers to develop an energy strategy that fits their unique business needs. Then they help customers implement this plan consistently across their entire organization. With a combination of services that includes Utility Management, Energy Procurement, and Remote Monitoring, Prenova drives customer savings of up to 10 to 15 percent of energy spend. From securing the best available rates in deregulated energy markets to optimizing the performance of building systems in real time, the company's solutions span the entire energy lifecycle.

Energy Management—Top Priority for Many Organizations

Energy Procurement helps customers reduce energy spend by securing lower utility prices in deregulated markets and identifying the most appropriate rates and tariffs in regulated markets. Prenova's team of experts continuously monitors global events and market trends to assess the short- and long-term impact changing conditions will have on energy supply. They also track changes in state and federal regulations that may affect rates. Extensive experience and understanding of the energy industry help Prenova develop a unique sourcing strategy for each customer.

Utility Management Services

Prenova's Utility Management solution is designed for organizations that want to free up internal resources by outsourcing the process of receiving, processing, and remitting utility invoices. This service includes a two stage audit that helps customers avoid overpaying for energy. A pre-payment review spots the most common billing mistakes, while a thorough post-payment analysis identifies billing anomalies that may indicate an error has been made by the utility. Prenova personnel investigate and resolve all billing errors, securing refunds when necessary.

Remote Monitoring

Remote Monitoring optimizes the performance of critical building systems, including HVAC, lighting, and refrigeration equipment. Prenova technicians monitor the performance of these systems from the company's remote Operations Control Center and are on hand 24/7 to respond to inbound service requests and system alarms. Using the company's advanced Asset Optimization technology, they perform sophisticated diagnostics to ensure equipment is functioning properly, within predefined operating standards. Should a problem occur, Prenova staff can often resolve the issue remotely, helping the customer avoid an expensive field service call.

Exporting Mediator Data

In order to take advantage of the applications and services of energy management partners, data from the Mediator must be exported to these partners. The Cisco Mediator stores data collected from various sensors and managed devices locally within FLASH. The flash storage available within the mediator itself is limited, and based on the number and rate of data logging typically can only store a week’s worth of information. For long term storage and trend analysis of energy management and building systems this data should be exported to a central repository on a regular basis.

Mediator Data Exporting Steps

The Mediator can be configured to export data through the following four types of transports:

- HTTP post
- HTTPS post
- SFTP
- FTP
- SMTP

There are several supported export formats, and one specifically for FSG. This example shows the steps necessary to export data to an FSG account. Out of the four transport types, FSG can only accept FTP and SFTP export options from the Mediator at this time.

Note A best practice is to assign all physical inputs and soucres via aliases. These aliases are then used in logging and are more user friendly names. If for some reason a device needed to change to a different physical interface (e.g., COM5 to COM5), the alias just needs to be pointed to the new physical interface, eliminating the tedious task of finding all other references for that resource in logs and reports.

The following are the steps to create a new Export Log, typically one is created for each device at the location:

Step 1 Under the logger right click and select **ADD** and **Periodic Log**.

Step 2 Assign an appropriate name for the log related to that location.

Step 3 Select the **Wizard is Available** button.

Step 4 Select **Next** at the bottom, then **Create New Column**.

Step 5 In the New Column wizard dialog, change the column type to *Periodic Column with Attributes of an FSG Channel* (see Figure 7-5).

Figure 7-5 - FSG Exporter —New Column Details

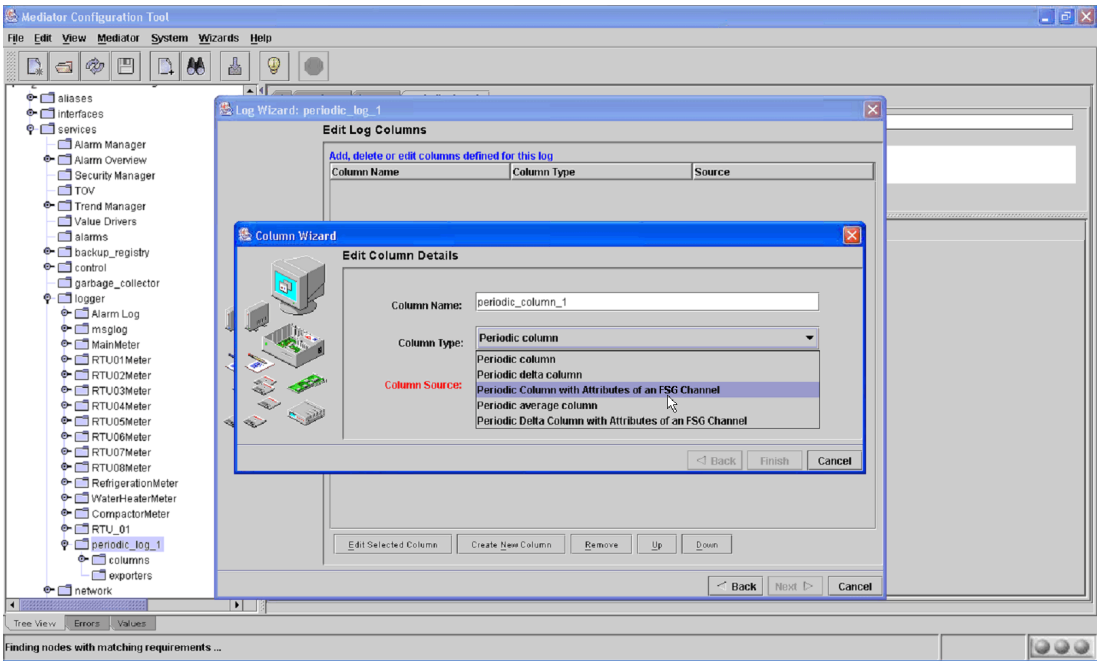

Step 6 Add the appropriate column source for this column via the Node Selector pop-up (see Figure 7-6).

Figure 7-6 FSG Exporter - New Column Node select

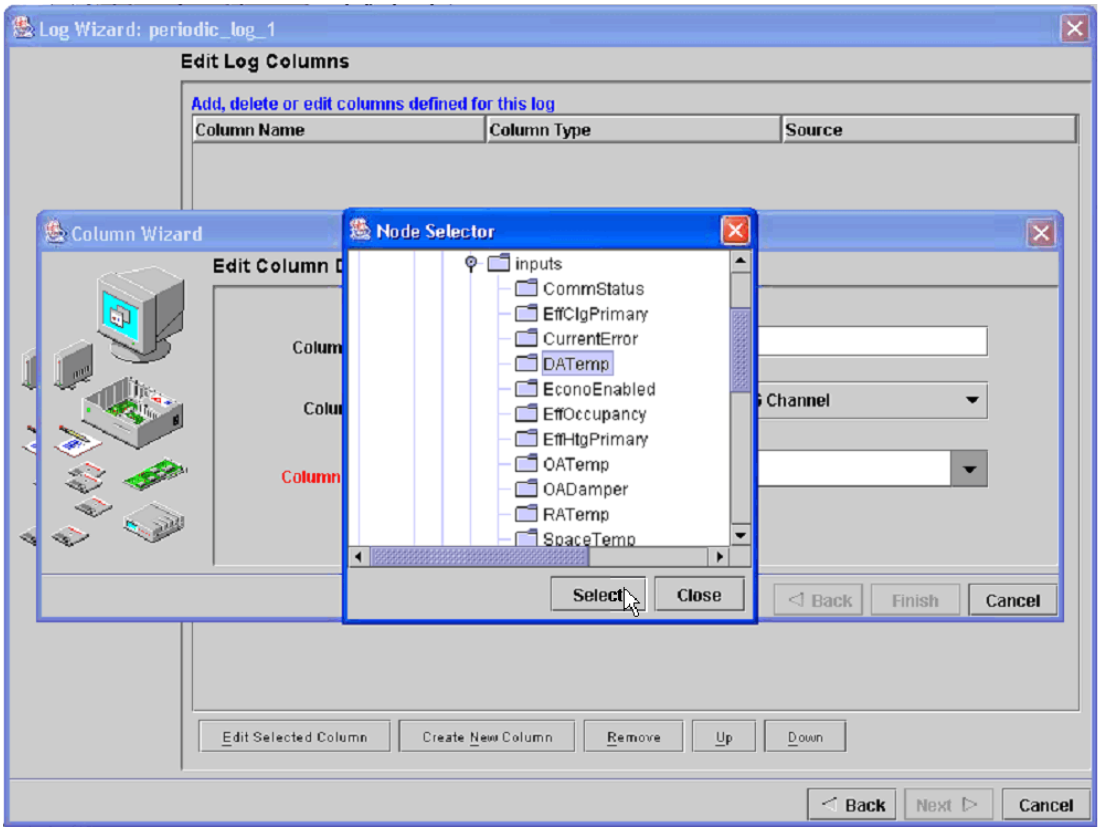

Step 7 Click **finish** and then repeat Steps 4 to 7 to add additional columns.

Step 8 Once all the desired columns have been added, click **Next**.

Step 9 On the Define Log Exporters page, select **New**.

Step 10 in the Log Exporter Wizard interface, make sure the Exporter Type is **Standard periodic data exporter** is selected (see Figure 7-7).

Figure 7-7 FSG Exporter - Log Exporter Type

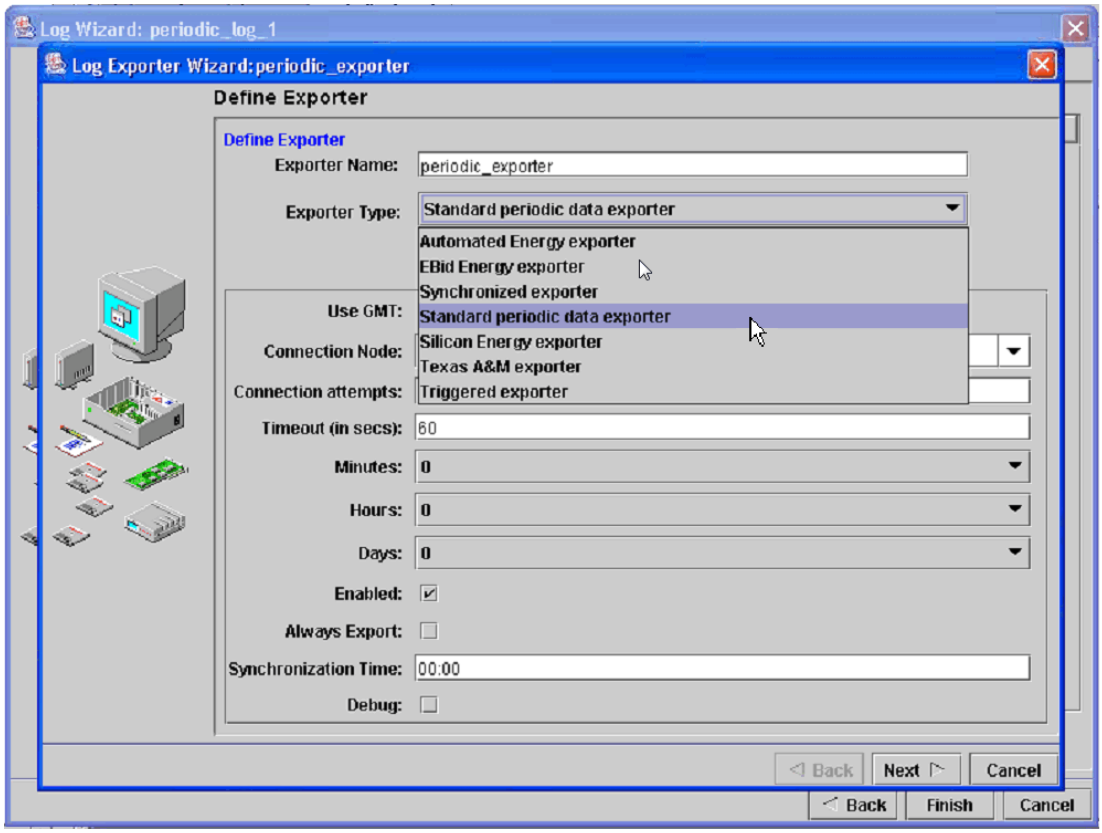

Step 11 Click **Next**.

Step 12 On the Define Formatter page, change the Formatter Type to **FSG XML formatter**.

Step 13 Click **Next**.

Step 14 Set the Transporter type to **Standard FTP** or **SFTP** as desired.

Step 15 Enter the appropriate hostname, username, and password for the FSG service account.

Step 16 Set the port as appropriate for the protocol used (21 for FTP, 22 for SFTP).

Step 17 the directory will be */home/FTPUser*.

Step 18 The file name prefix should be unique and be based on a hierarchal schema such as [*CustomerEnterpriseName*]_[*LocationID/Name*]_[*MediatorID/Name*]_[*DeviceID/Name*] (see Figure 7-8).

Figure 7-8 FSG Exporter - XML File Name

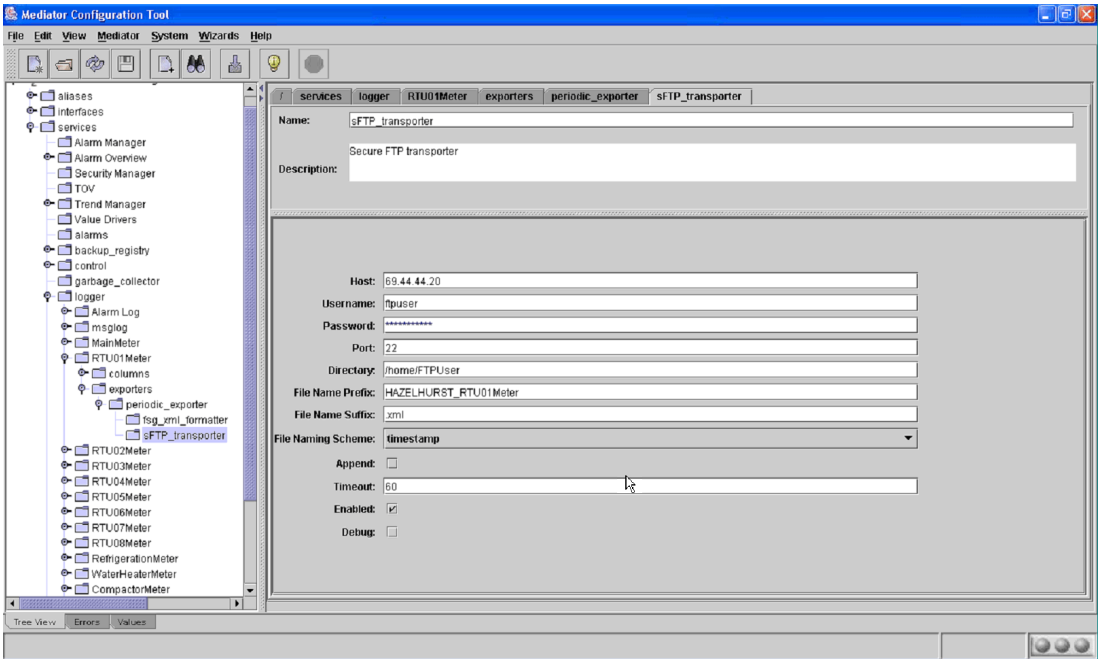

Step 19 Under the **periodic_exporter**, select the **fsg_xml_formatter** folder.

Step 20 Set the **Location Info** and **Location Key** data fields. These items must be unique to FSG and are often set to a unique store number and name relevant to the enterprise: [*StoreNumber-StoreName]*_[*MediatorID/Name*].

Step 21 Set the **Panel Info** and **Panel Key** data fields. These items must be unique to the location and are often set to the common device name alias [*StoreNumber*]_[*MediatorID/Name]*_[*DeviceID/Name]* (see Figure 7-9).

Figure 7-9 FSG Exporter - XML Location Info

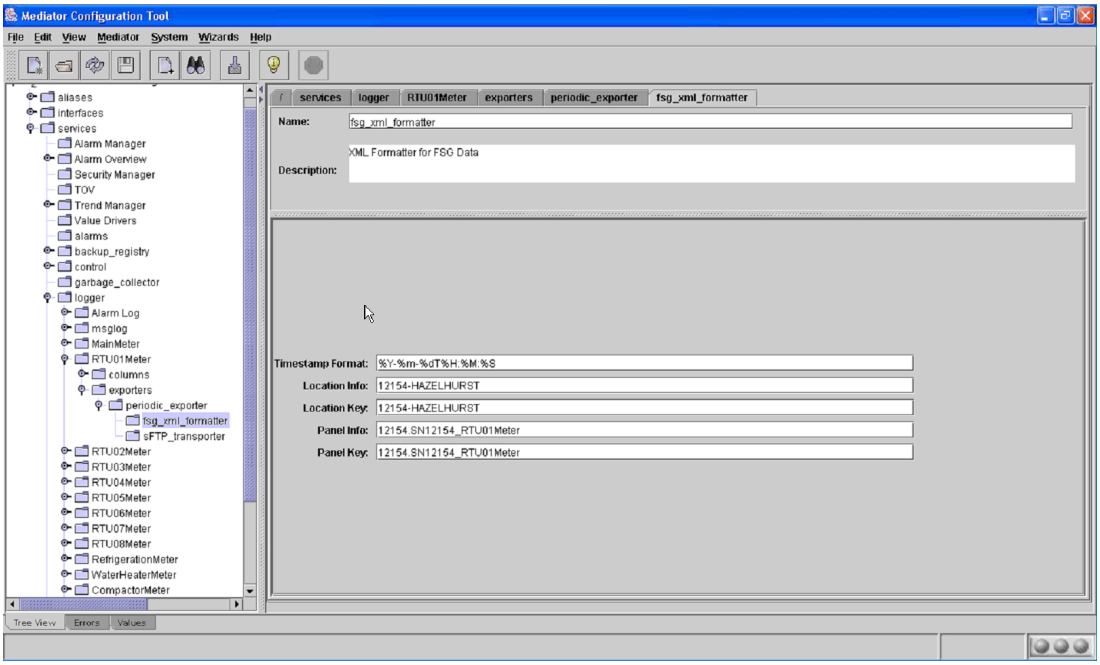

Step 22 Change the timer, if a different export rate is desired.

This completes the steps necessary for creating an XML data export to FSG.

Data that is sent to FSG is normally collected by the database system at 15-minute intervals and assimilated for use. Once assimilated it is used to update reports and status pages used by building managers, compared against trend indicators and historical system function.

Case Study of a Large Retail Merchandise Chain

Revamped Stores

A large retail merchandise chain has set its expectations to save money through energy management. Green building techniques at this retail merchandise chain involve roughly 100 different environmental projects running concurrently. The large retail merchandise chain first partnered with General Electric (GE) back in 1968 where they began developing energy-efficient lighting for stores. Fluorescent ceiling lights alone saved $5.7 million a year in energy costs. Solar energy now provides 20 percent of the electricity needs in 52 stores of the large retail merchandise chain and two distribution centers, and LEDs illuminate store refrigerators. In May, the first store of the large retail merchandise chain with a green roof opened in Chicago. The roof is planted with heat and water-absorbing plants, reducing both heating costs and water runoff. Meanwhile, other environmentally friendly policies (reduced water consumption, more recycling of construction waste, designated parking for energy-efficient cars, and bike racks for customers and workers) all contributed to this large retail merchandise chain's selection to participate in

a pilot program run by the U.S. Green Building Council to help develop environmental standards for retail construction. the large retail merchandise chain used in this case study is the only drugstore chain among 70 retailers in the program.

In its ongoing efforts to conserve energy, the large retail merchandise chain partnered with Cisco Systems and Facilities Solutions Group to automate its store building systems in a more comprehensive and coordinated fashion to monitor and control building systems globally. See Figure 7-10.

Figure 7-10 WG-1 - HVAC Cisco Mediator

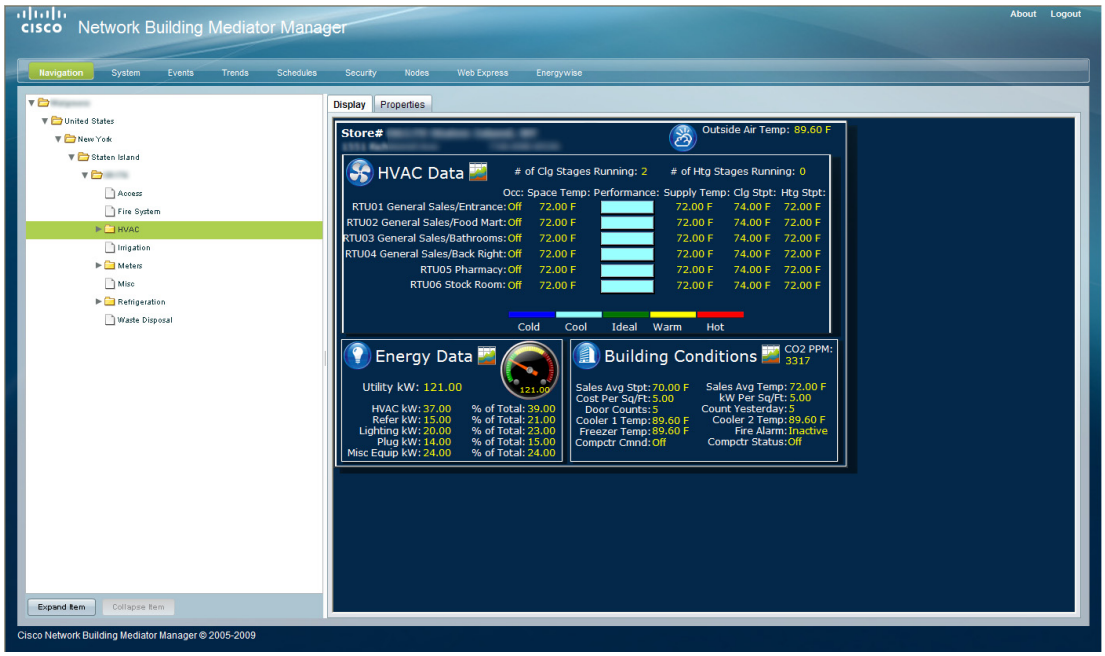

Cisco and FSG Pilot

The large retail merchandise chain pilot implemented systems to monitor and measure various systems throughout the store. Systems were connected to the Cisco Mediator as follows:

- HVAC rooftop unit (RTU) control systems and temperature sensors connected via BACnet
- Main incoming power lines via a SATEC BFM136 meter
- Refrigerators and freezers via DALLAS sensors
- Fire, sprinkler, and trash compactor systems via relays
- Lighting systems via a KMC BACnet controller

Savings Calculations for the Enterprise Energy Management System

The following items were used to calculate savings achieved in the pilot:

- FSG Energy Scoreboard usage data as shown in Figure 7-11
- Weather trend daily CDD and HDD data (from Austin Bergstrom Weather station)
- Calculation from Dr Haorong Li University of Nebraska

Figure 7-11 FSG Scoreboard Energy Collection Information

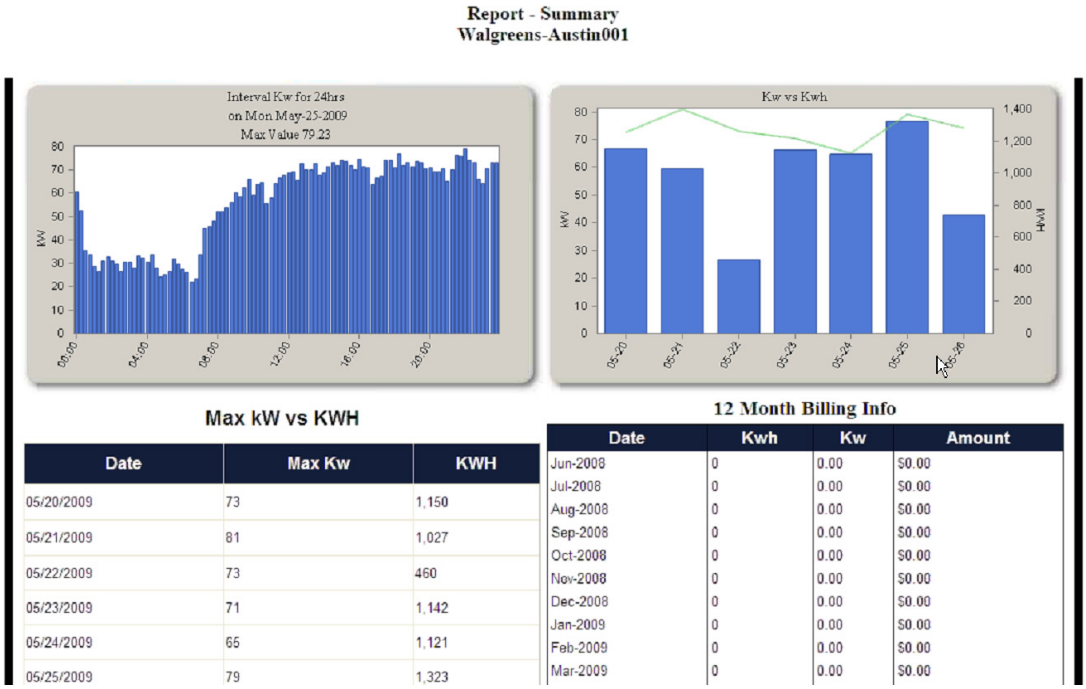

Date	Max Kw	KWH
05/20/2009	73	1,150
05/21/2009	81	1,027
05/22/2009	73	460
05/23/2009	71	1,142
05/24/2009	65	1,121
05/25/2009	79	1,323

Date	Kwh	Kw	Amount
Jun-2008	0	0.00	$0.00
Jul-2008	0	0.00	$0.00
Aug-2008	0	0.00	$0.00
Sep-2008	0	0.00	$0.00
Oct-2008	0	0.00	$0.00
Nov-2008	0	0.00	$0.00
Dec-2008	0	0.00	$0.00
Jan-2009	0	0.00	$0.00
Feb-2009	0	0.00	$0.00
Mar-2009	0	0.00	$0.00

Austin Texas—Large Retail Merchandise Chain Case Study Analysis for February

The HVAC systems were the initial focus of the analysis. They are often the largest energy loads of any building, providing the area of greatest return of effort. Each of the large retail merchandise chain stores uses six separate rooftop units (RTU) to maintain temperature through out the store. Each unit is typically controlled via a locally-zoned thermostat, independent of the other systems in the building. As the analysis in Figure 7-12 shows, these systems would often operate in opposing modes and not make use of economizer capabilities.

Figure 7-12 Analysis for February

		Off	Heating	Cooling	Economizer Potential % of Cooling	Economizer Potential Savings	Simultaneous (Hours) Heat	Simultaneous (Hours) Cool
RTU1	Cashier	53%	3%	44%	25%	56.78	0.3	43.9
RTU2	Sales	85%	13%	2%	16%	1.58	16.1	1.3
RTU3	Cosmetics	90%	0%	10%	27%	26.42	-	13.3
RTU4	Pharmacy	85%	6%	9%	24%	11.32	10.0	9.5
RTU5	Stock	73%	26%	1%	9%	0.15	34.3	0.1
RTU6	Photo Lab	98%	0%	2%	20%	1.40	-	0.9
		81%	8%	11%		97.65	19%	16%
							Savings	105.0

229087

The information gathered through comprehensive building system monitoring identified the following key problems:

- One or more RTUs were running on average 11 percent of the time
- The Cashier-zone RTU ran significantly more than the other RTUs mainly in cooling mode.

- The Stock-zone RTU ran much more in heating mode than any of the other RTUs.
- On average, the RTUs were in cooling mode 20 percent of the time where they could have been using outside air and had the compressors off, saving close to $100 for the month.
- Between 15 and 20 percent of the time, while the RTUs were running, there was another RTU operating in a different mode.
- In the main floor area covered by RTUs 1,2, and 3, RTU1 and RTU3 were often in cooling mode to balance RTU2 that was often in heating mode. RTU5 was also often running in heating mode when other RTUs were in cooling mode.

Tip Additional savings could also be realized through more advanced supervisory control implementation.

By integrating these HVAC systems into a comprehensive energy management system through the Cisco Mediator, coordinated control can be accomplished and significant savings achieved.

Total Annual Savings Projected and Measured

The following show the initial dollar cost savings achieved in the pilot:

- Total energy costs last year: $39,615.00
- Total kWh used last year: 455960 ($/kWh at .087)
- Savings for economizers and heat and cool coincidental loading: *$4,893.00/year*
- Savings for cooling loads:15 percent blended value of 306560 kWh during cooling season (March through October)
- Total HVAC savings: 45984 kWh saved; $4,000.60 saved (kWh at $.087)
- Combined annual savings: $8,893.60 ($4,000.60 + $4,893)
- Margin of error of saving overlap of staging and HVAC control is 20percent—Conservative 20 percent reduction
- Total savings is $7,114.88 (.8*$8893.60), a 17.96 percent savings annually.

Overall, the large retail merchandise chain in this case study reduced energy consumption by more than 30 percent in the Austin Texas Pilot store through monitoring and automated systems intelligence. To further validate sustainability, the large retail merchandise chain is extending its energy management solution from the pilot to a 500-store deployment.

Other retailers with Mediator deployments include the following:

- Circle K—Remote energy management deployment 300 stores
- Simon Mall—Automated meter reading of thousands of check meters for energy usages.

Other enterprise with Mediator deployments include the following:

- Google Campus —Energy control and demand response across entire campus. 25 percent reduction achieved.
- NetApp Campus—Automated demand response system from PG&E, system successfully shed 1MW in 10 minutes.
- Wipro Campus (India)—Converged systems management over IP

Summary

Reducing energy consumption in buildings can only be accomplished in a step-wise manner. The sustainability of an energy management system is based on the ability to integrate data from all of the systems in a building and intelligently control these devices together in a coordinated fashion. Cisco's Network Building Mediator provides access to all of the energy consuming systems within a building, normalizing this data and making it accessible to a broad range of applications and partners. Coordinated decisions and system programming changes can then be aggregated to efficiently operate the building as a whole. Building occupant behavior is also crucial to ensuring that energy reduction is sustainable. A comprehensive energy management solution is one that keeps building occupants informed about their impact on energy use and carbon footprint, while improving efficiencies in building operations and use.