



Solution Overview—Network Design Considerations for Cisco Mediator Deployments

Revised: 11/13/09

Contents

Problems of Traditional Building Systems	2
New Open Energy Management Systems	2
Cisco Network Building Mediator	3
Network Design Implications	4
Management Services Protocols	6
Cloud Services Protocols	10
Deployment Models	12
Branch Network Design Considerations	14
Distributed VPN Connectivity Designs	14
Centralized VPN Connectivity Designs	16
QoS Within the Branch	21
Campus Network Design Considerations	21
Internet Edge Module	22
Partner Extranet Module	23
Collapsed Internet Edge Design	27
Campus Building Module	28
Layer-2 Access Layer Switch Designs	28
Layer-3 Access Layer Switch Designs	30
Extending VRFs to the Campus Building Module	32
QoS within the Campus Building Module	35



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

Data Center/Campus Service Module	38
Data Center Service Module Design	38
Campus Service Module Design	42
Extending VRFs to the Data Center/Campus Service Module	43
QoS within the Data Center/Campus Service Module	44
WAN Module	45
Campus Core Module	46
Summary	46

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](#).

Table 1 Cisco Network Building Mediator Models

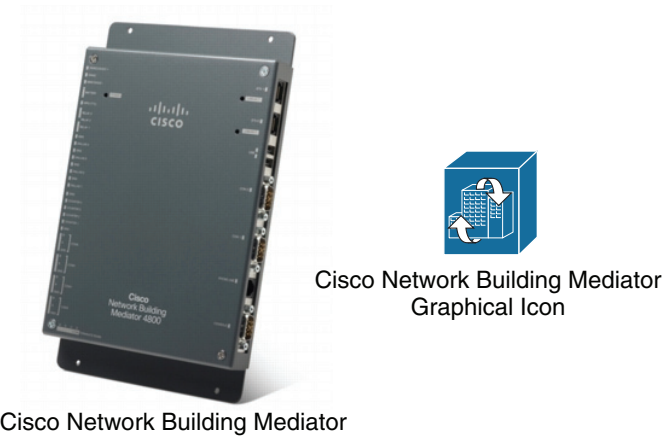
Model	Description	Building Control Protocol Licensing Options
Cisco Network Building Mediator 4800	Targeted for campus deployments, it supports up to approximately 5,000 points.	Base, Intermediate, and Advanced Protocols
Cisco Network Building Mediator 2400	Targeted for branch deployments, it supports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


Note

A point 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](#) shows a Cisco Network Building Mediator with the icon used in figures in this document.

Figure 1 Cisco Network Building Mediator

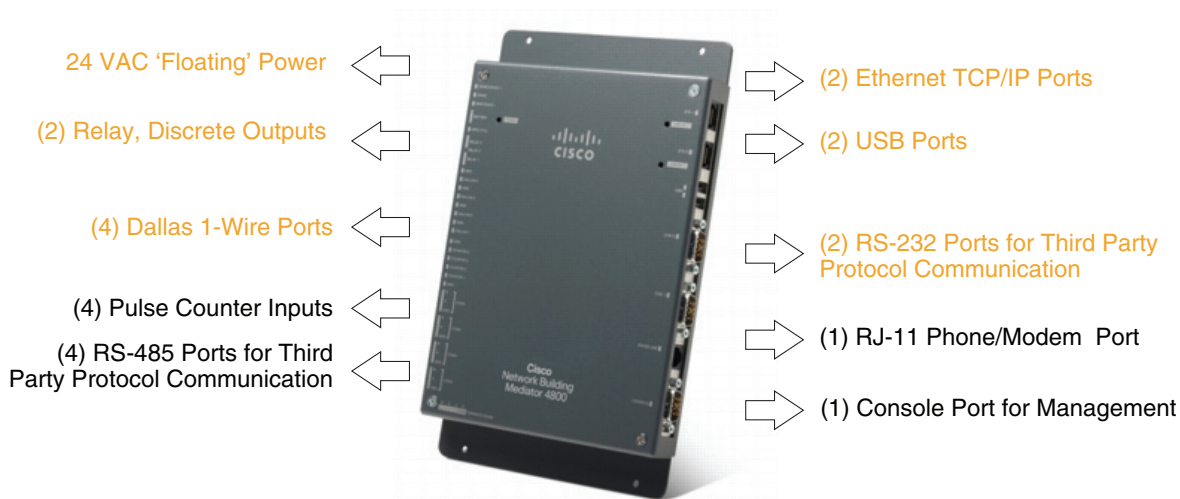


The Mediator aggregates and normalizes building systems data, making it available through an open XML interface.

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 2](#) shows a close-up of the communication and I/O ports of the Mediator.

Figure 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 3](#).

Figure 3 10/100 FastEthernet Connectivity Example on the Mediator

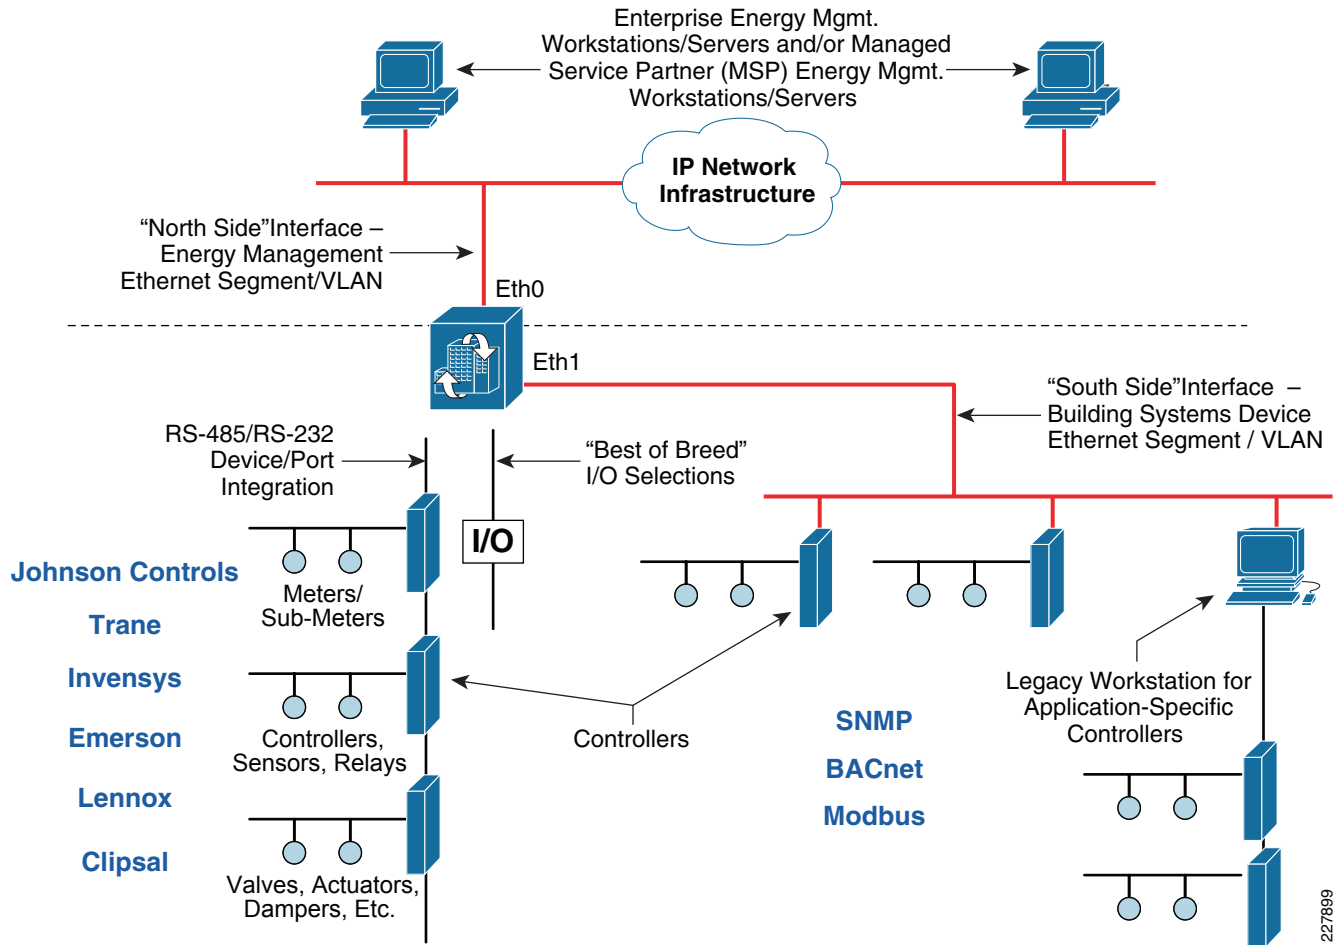


Figure 3 shows an example of a Mediator connected to a number of floor, zone, or room controllers by way of both traditional RS-232/RS-485 wiring and a dedicated Ethernet segment. The controllers are in turn connected to the actual building energy management devices—meters, sub-meters, valves, actuators, 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.



Note

The Mediator can also interface with legacy management workstations, as shown in Figure 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:

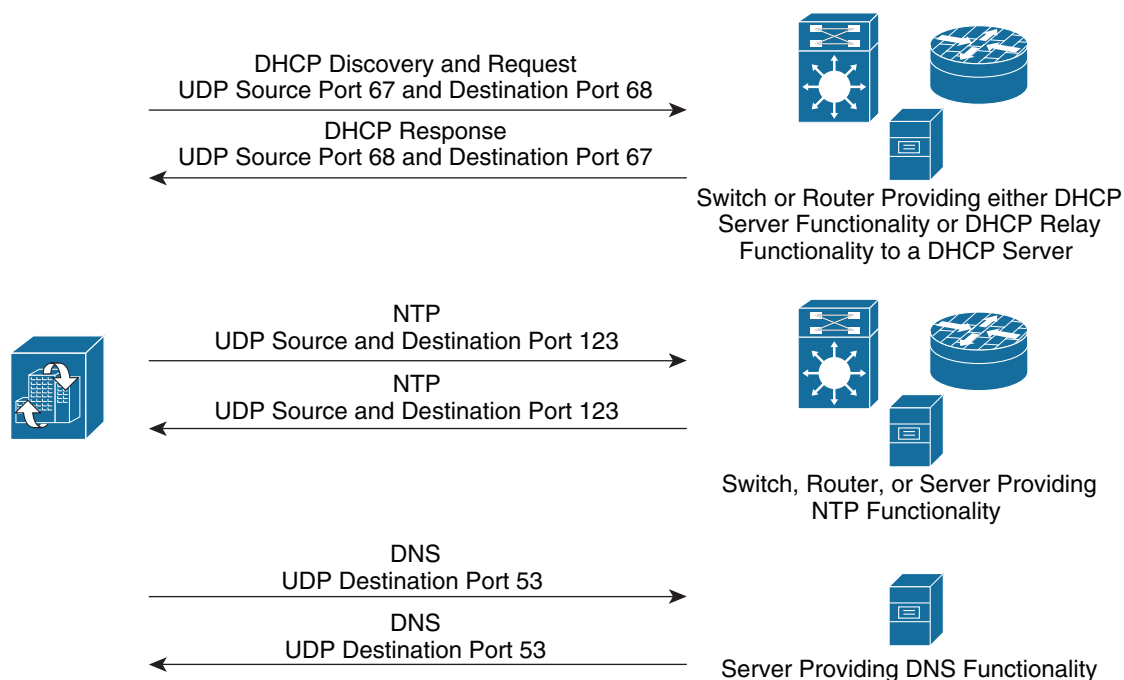
- Dedicated firewall appliances, such as the Cisco ASA 5500 Series.
- Firewall services integrated within a router or switch platform, such as the Context-Based Access Control (CBAC) or Zone-Based Policy Firewall (ZBPF) features of Cisco ISR router platforms, or the Firewall Services Module (FWSM) of 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 application of each of these access control methods within different parts of the network infrastructure is discussed in detail in the [“Deployment Models” section on page 12](#). The following sections discuss some of the network protocols required on the north side or energy management interface 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.

Management Services Protocols

Management services include the protocols required for provisioning the Mediator onto the network infrastructure. Device provisioning protocols can include Dynamic Host Configuration Protocol (DHCP), Domain Name Services (DNS), and Network Time Protocol (NTP). [Figure 4](#) shows an example of these flows.

Figure 4 Management Flows—Device Provisioning



227900

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 Service Module or Campus Service Module. DHCP uses UDP ports 67 and 68.

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 Service Module 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 with 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.

Management services also include protocols required for configuring the Mediators, creating control logic deployed onto the Mediator, and for monitoring the Mediators. These services are provided through the management applications listed in [Table 2](#).

Table 2 *Mediator Configuration and Management Applications*

Application Name	Description
configTOOL	ConfigTOOL is a software application which 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 <code>broadway.xml</code> .
perfectHOST	PerfectHOST is an application which 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.

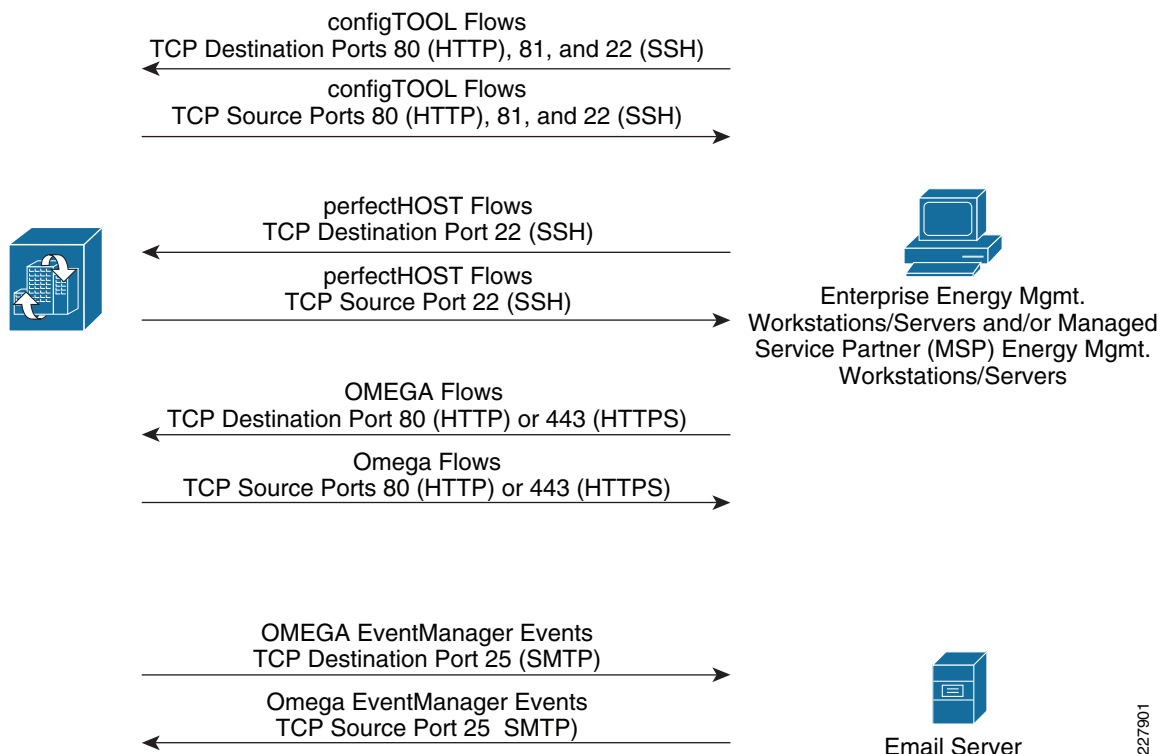
Specific tools within the OMEGA software suite are listed in [Table 3](#).

Table 3 **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 data points 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 5 shows the protocols required to enable the flows needed by the Mediator configuration and management applications.

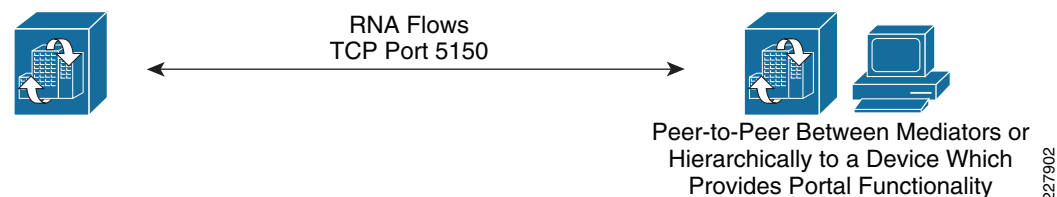
Figure 5 Management Flows—Device Configuration, Logic Configuration, and Monitoring



ConfigTOOL uses TCP ports 80 (HTTP), 81, and 22 (SSH) to establish connectivity to the Mediator and download the broadband 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 5 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 initiates events outbound from the Mediator to e-mail servers via the TCP port 25 (SMTP).

One additional protocol used by the Mediators is Remote Node Abstraction (RNA). RNA is the protocol used to share node values in between Mediators. RNA allows two or more Mediators on the same network to share points in a peer-to-peer manner or hierarchical manner in a portal configuration. RNA uses TCP port 5150 as shown in Figure 6.

Figure 6 Management Flows—Remote Node Abstraction (RNA)



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. Data points 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 1](#) shows an example of an FSG XML format showing multiple data points.

Example 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">
```

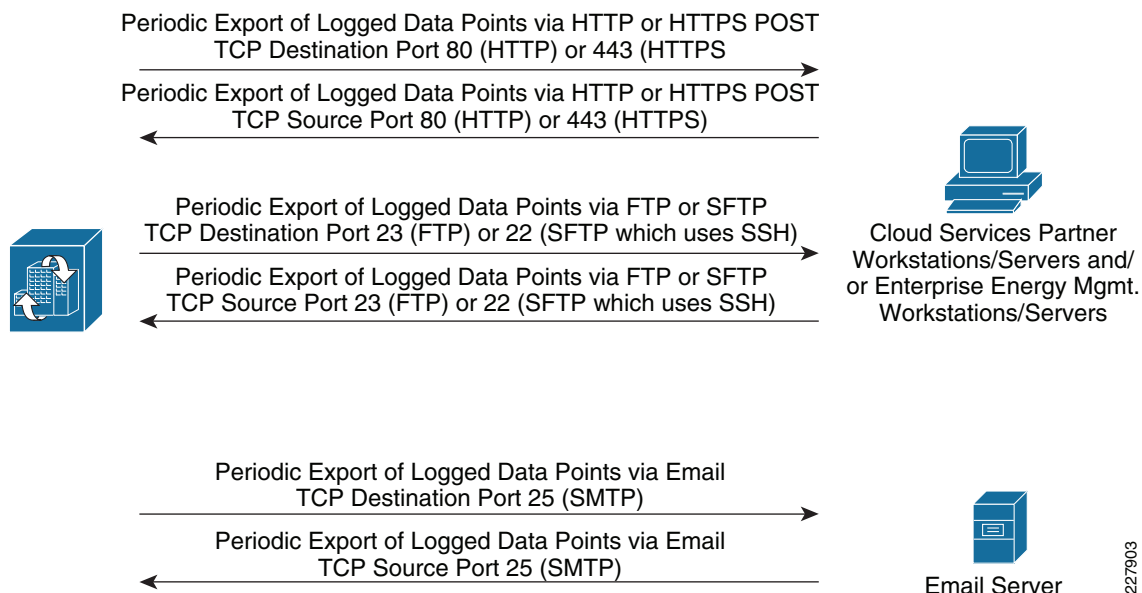
```

<value timestamp="2009-09-24T13:15:00">478
</value>
</channel>
</device>
</data>

```

Figure 7 shows the protocols required to enable the data flows associated with the periodic export of logged data from the Mediators.

Figure 7 Cloud Services Flows—Periodic Export of Logged Data Points



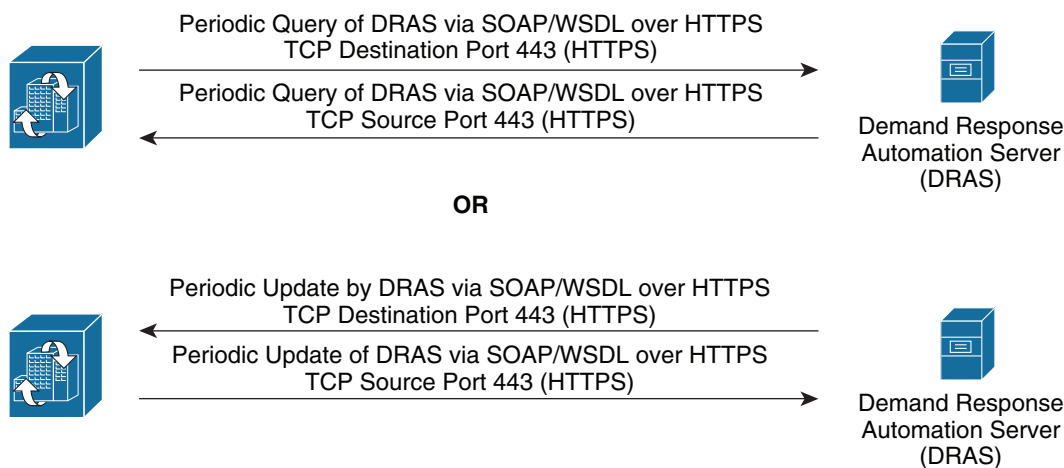
The data flows associated with the periodic export of logged data points are typically initiated from the Mediator outbound to a cloud services partner workstation or server. However, the reverse traffic must be allowed back through the network in order for the session to be established. Specific protocols 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 utilize a separate dynamic data channel.

Automated Demand Response (ADR) systems often utilize 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 8 shows an example of the flows involved.

Figure 8 Cloud Services Flows—Example Automated Demand Response Flows

In [Figure 8](#), a server—sometimes referred to as an Demand Response Automation Server or DRAS—is deployed at the energy service provider location. The Mediator functions as a DRAS client, periodically polling the DRAS for event state information. Polling times can be as often as every minute. 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 to ensure confidentiality and authenticate the sessions. The Mediator can then use the event state information from the DRAS to implement pre-programmed logic, such as re-setting the setpoints of thermostats, in order shed load and reduce energy consumption.

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.

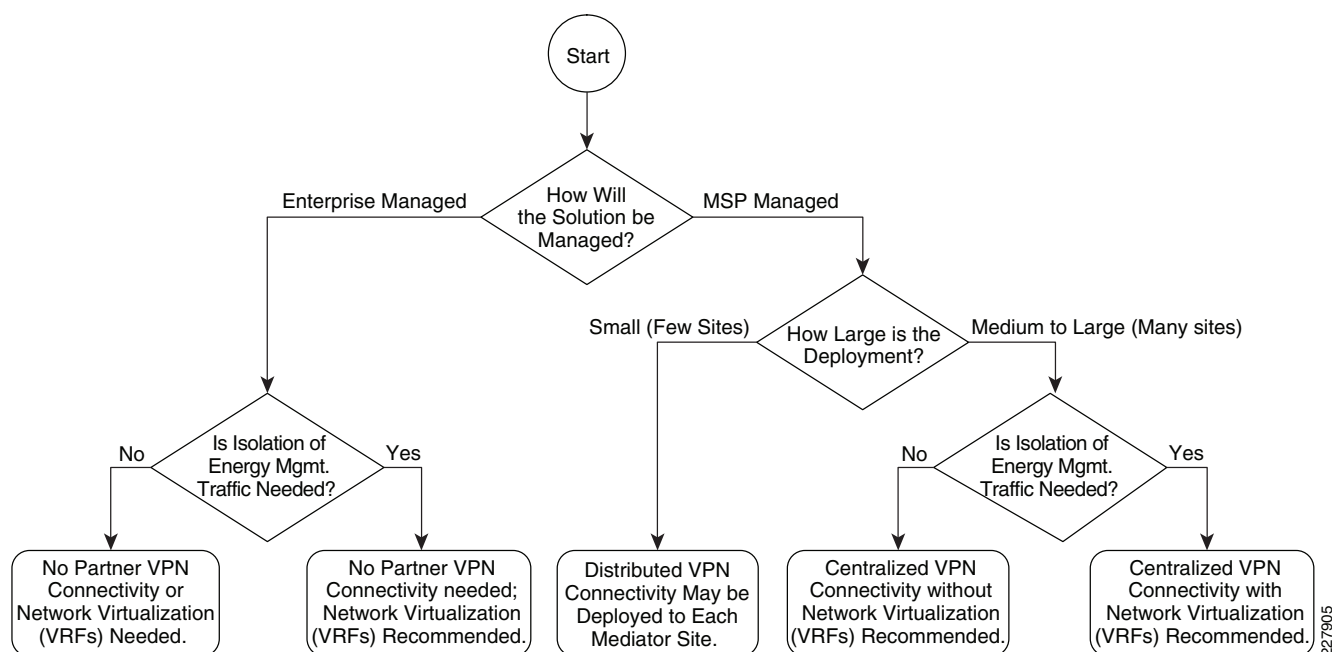
With the second deployment model, the enterprise customer may use a Cisco partner to deploy and then manages the energy management system themselves. This model requires full management and monitoring capabilities to and from the Mediators and management workstations within the enterprise.

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 an Energy Scoreboards, ADR, AFDD, and dynamic pricing applications.

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 logical 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 9 summarizes the choices for deployment of the energy management solution over an enterprise network infrastructure.

Figure 9 *Energy Management Solution Deployment Flowchart*



Specific considerations for the branch and campus Places-in-the-Network (PINs) designs are discussed in the following subsections.

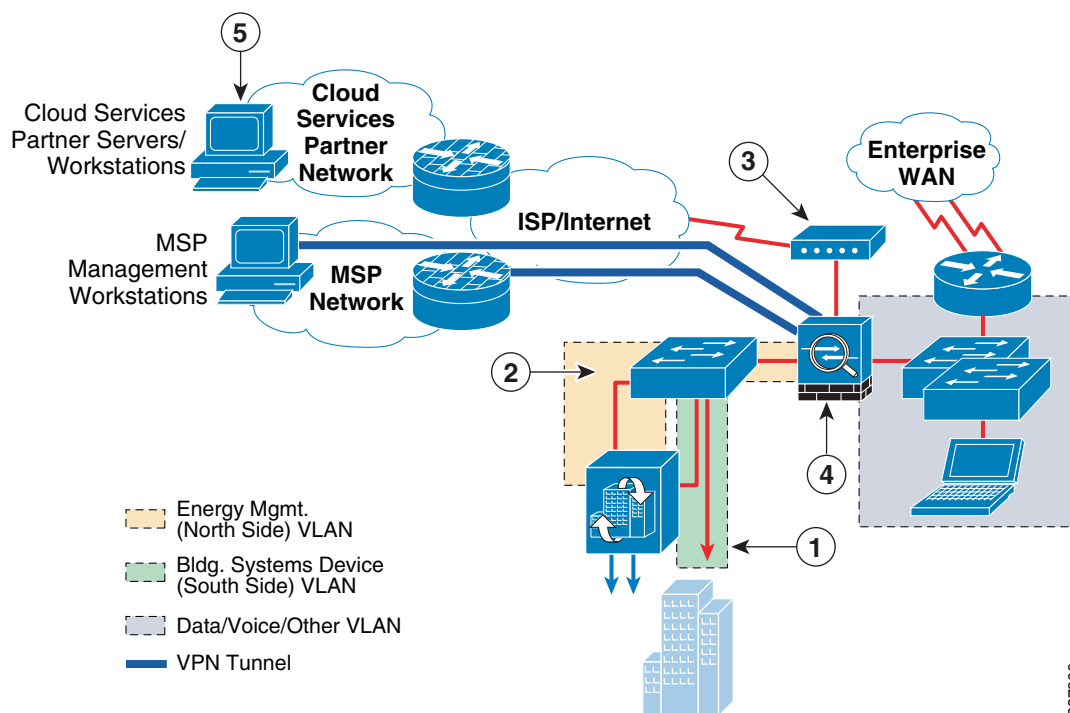
Branch Network Design Considerations

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 branch site is shown in [Figure 10](#).

Figure 10 Branch Design with Dedicated Switch and ASA 5500 Security Appliance



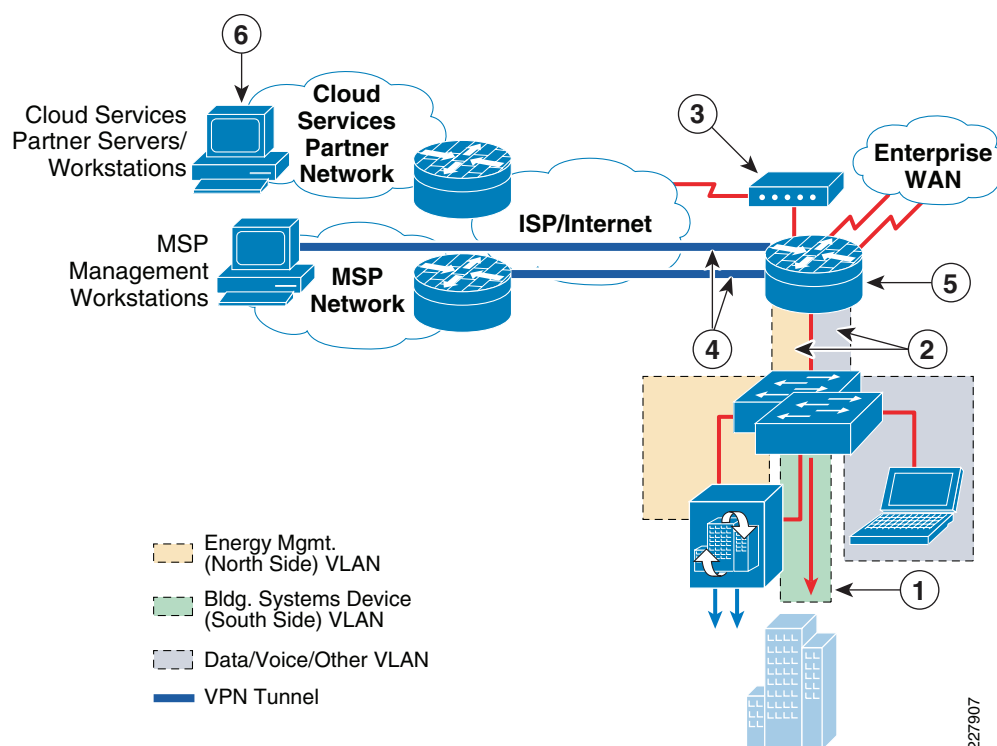
The following steps describe what occurs in [Figure 10](#):

- Step 1** Building systems device VLAN isolated by not trunking it to the dedicated ASA 5500 Security Appliance.
- Step 2** Dedicated ASA 5500 Security Appliance connected to the energy management VLAN.
- Step 3** Separate Internet Service (DSL, cable, etc.) provisioned for MSP partner VPN connectivity (site-to-site or remote access).

- Step 4** Dedicated ASA 5500 Security Appliance provides both MSP partner VPN termination and stateful access control to and from the energy management VLAN.
- Step 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 branch site is shown in [Figure 11](#).

Figure 11 Branch Design with VPN Access and Integrated Infrastructure



The following steps describe what occurs in [Figure 11](#):

- Step 1** Building systems device VLAN isolated by not trunking it to the branch router.
- Step 2** The energy management VLAN and the data/voice/other VLAN trunked to the branch router.
- Step 3** Separate Internet service (DSL, cable, etc.) may need to be provisioned for partner VPN connectivity.
- Step 4** Branch router with IPSec VPN software provides termination of site-to-site or remote-access VPN from MSP partner.
- Step 5** CBAC or ZBPF provide stateful access control to and from the energy management VLAN.

- Step 6** 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.

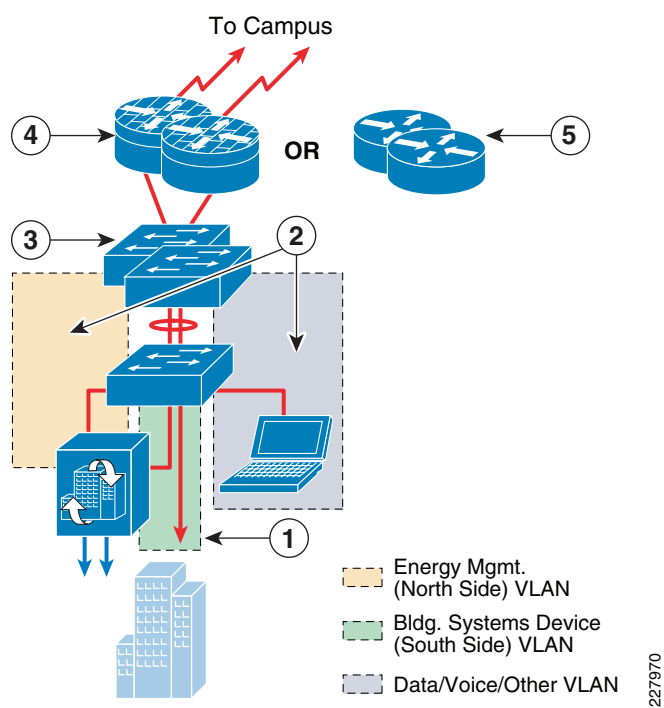
In cases where an IPSec VPN provides the enterprise WAN connectivity, an additional VPN tunnel may be provisioned to the MSP network. In the case where private enterprise WAN connectivity is provisioned, a separate Internet connection can be provisioned on the Cisco ISR branch router, as is shown in [Figure 10](#). Either CBAC or ZBPF running within the branch ISR router would 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.

Comparing both dedicated branch VPN designs shown in [Figure 10](#) and [Figure 11](#), the design shown in [Figure 11](#) 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 implementations, centralizing the MSP VPN connectivity to a campus or data center location provides a much more scalable and manageable deployment. [Figure 12](#) shows an example of a large branch site design for support of the energy management solution.

Figure 12 Example Large Branch Site Design



The following steps describe what occurs in [Figure 12](#):

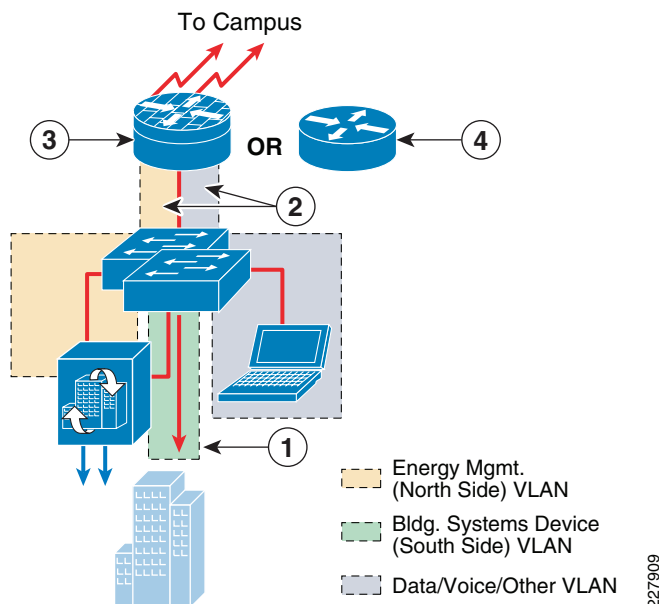
-
- | | |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | The building systems device VLAN is isolated by not trunking it from the Layer-2 access switch to the Layer-3 distribution switch stack. |
| Step 2 | The energy management VLAN and the data/voice/other VLAN are trunked to the Layer-3 distribution switch stack. |
| Step 3 | ACLs on the Layer-3 distribution switch stack provide stateless access control to the energy management VLAN from other VLANs within the branch. |
| Step 4 | Routers with CBAC or ZBPF provide stateful access control from the rest of the enterprise network to the energy management VLAN. |
| Step 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 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 which utilize 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 13 shows an example of a medium branch site design for support of the energy management solution.

Figure 13 Example Medium Branch Site Design



The following steps describe what occurs in Figure 13:

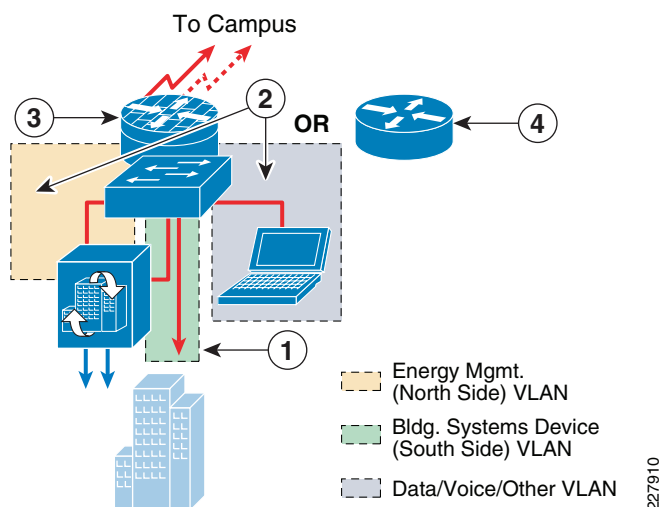
-
- Step 1** Building systems device VLAN isolated by not trunking it to the branch router.
 - Step 2** The energy management VLAN and the data/voice/other VLAN trunked to the branch router.
 - Step 3** Branch router with CBAC or ZBPF provide stateful access control to and from the energy management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

Medium branch sites often consist of just a branch ISR router and a Layer-2 Catalyst 2900 Series switch 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 which utilize 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 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 14 shows an example of a small branch site design for support of the energy management solution.

Figure 14 Example Small Branch Site Design



The following steps describe what occurs in Figure 14:

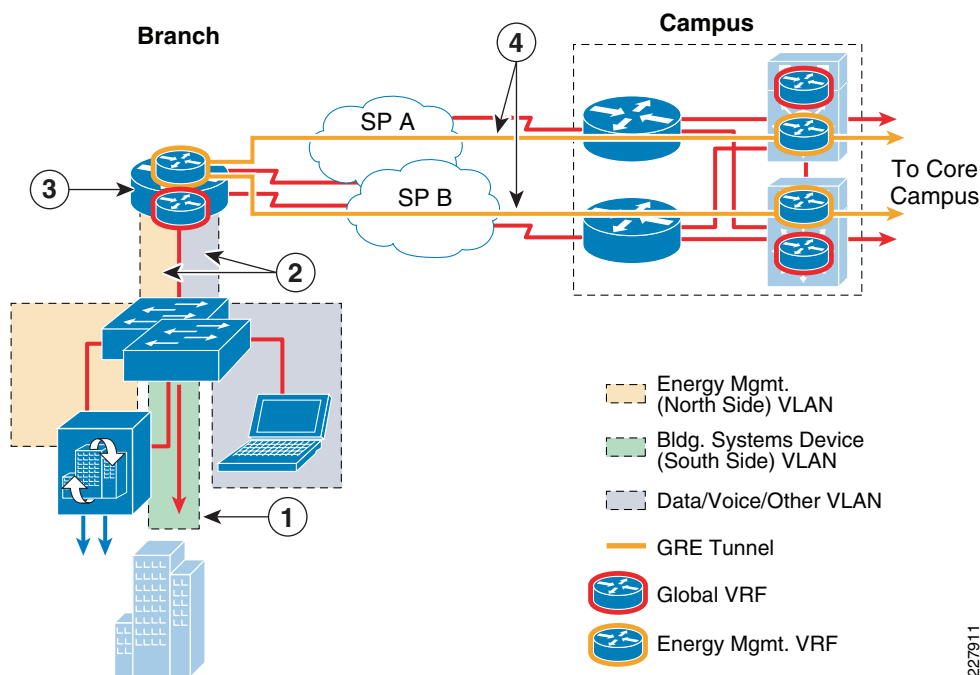
-
- Step 1** Building systems device VLAN isolated by not defining a Layer-3 VLAN interface on the branch router.
 - Step 2** The energy management VLAN and the data/voice/other VLAN extended to the branch router by defining Layer-3 VLAN interfaces on the branch router.
 - Step 3** Branch router with CBAC or ZBPF provides stateful access control to and from the energy management VLAN.
 - Step 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 15.

Figure 15 Example of an Energy Management VRF Extended to the Branch



The following steps describe what occurs in Figure 15:

- Step 1** Building systems device VLAN isolated by not trunking it to the Layer-2 distribution switch stack.
- Step 2** The energy management VLAN and the data/voice/other VLAN trunked to the ISR router.
- Step 3** Energy management VLAN mapped to the energy management VRF, while data/voice/other VLANs mapped to the global VRF within the branch router.
- Step 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 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 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. The centralized VPN connectivity from the MSP to all of the enterprise energy

227911

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 automation 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 model. GRE tunnels can be defined from the branch ISR router to the Layer-3 distribution switches within the campus WAN Module (discussed in the [“Campus Network Design Considerations” section on page 21](#)). 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. The reader should note that other methods of extending VRFs across the WAN exist as well, such as the mapping of VRFs to an MPLS service.

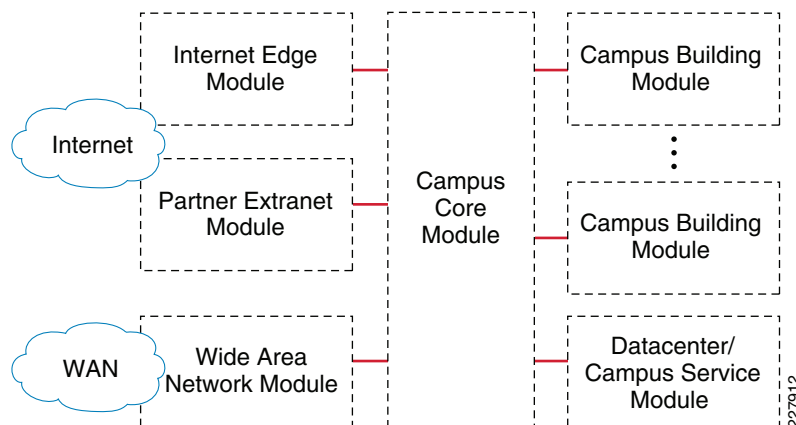
QoS Within the Branch

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. Since the function of classification and marking is the same at the branch as within the campus, but with different Catalyst switch platforms, refer to the [“QoS within the Campus Building Module” section on page 35](#) for a detailed discussion.

Campus Network 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 16](#) shows an example of a campus network design from a modular perspective.

Figure 16 *Campus Network Design Modules*



[Table 4](#) provides a brief overview of the role of each module in a campus network.

Table 4 **Campus Network Modules**

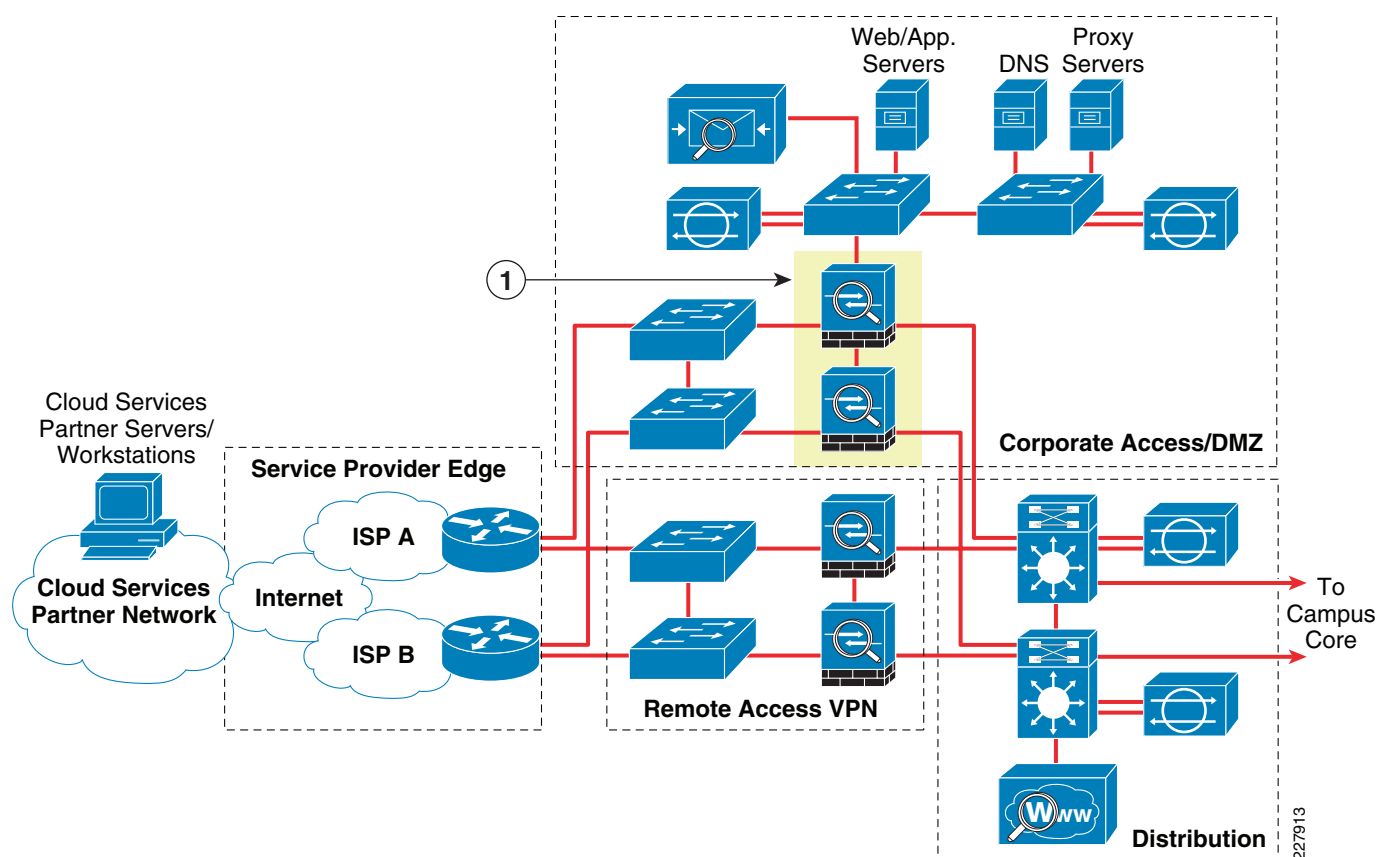
Module	Description
Internet Edge	Provides centralized and secure Internet connectivity to and from the enterprise network.
Partner Extranet	Provides centralized and secure connectivity to partner networks via VPN, the Internet, or direct connections.
Wide Area Network	Provides internal connectivity between 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.
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. An example of this is discussed in the [“Collapsed Internet Edge Design” section on page 27](#). Each of the modules presented in [Table 4](#) are discussed in detail in the following sections.

Internet Edge Module

In terms of the energy management solution, the function of the Internet Edge Module is to provide stateful access control for outgoing connections initiated by the Cisco Network Building Mediators to cloud services partner servers accessible via the Internet. This is typically for the periodic exporting of logged data from the Mediators. The Internet Edge Module also provides stateful access control for enterprise client PCs accessing Web-based energy scorecards provided by cloud services partners, also reachable through the Internet. An example of a redundant Internet Edge Module design is shown in [Figure 17](#).

Figure 17 Example Redundant Internet Edge Module Design



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 on the enterprise network is 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 Security Appliances within the Corporate Access/DMZ section of the Internet Edge Module, highlighted in Figure 16, can provide the required services within this module.

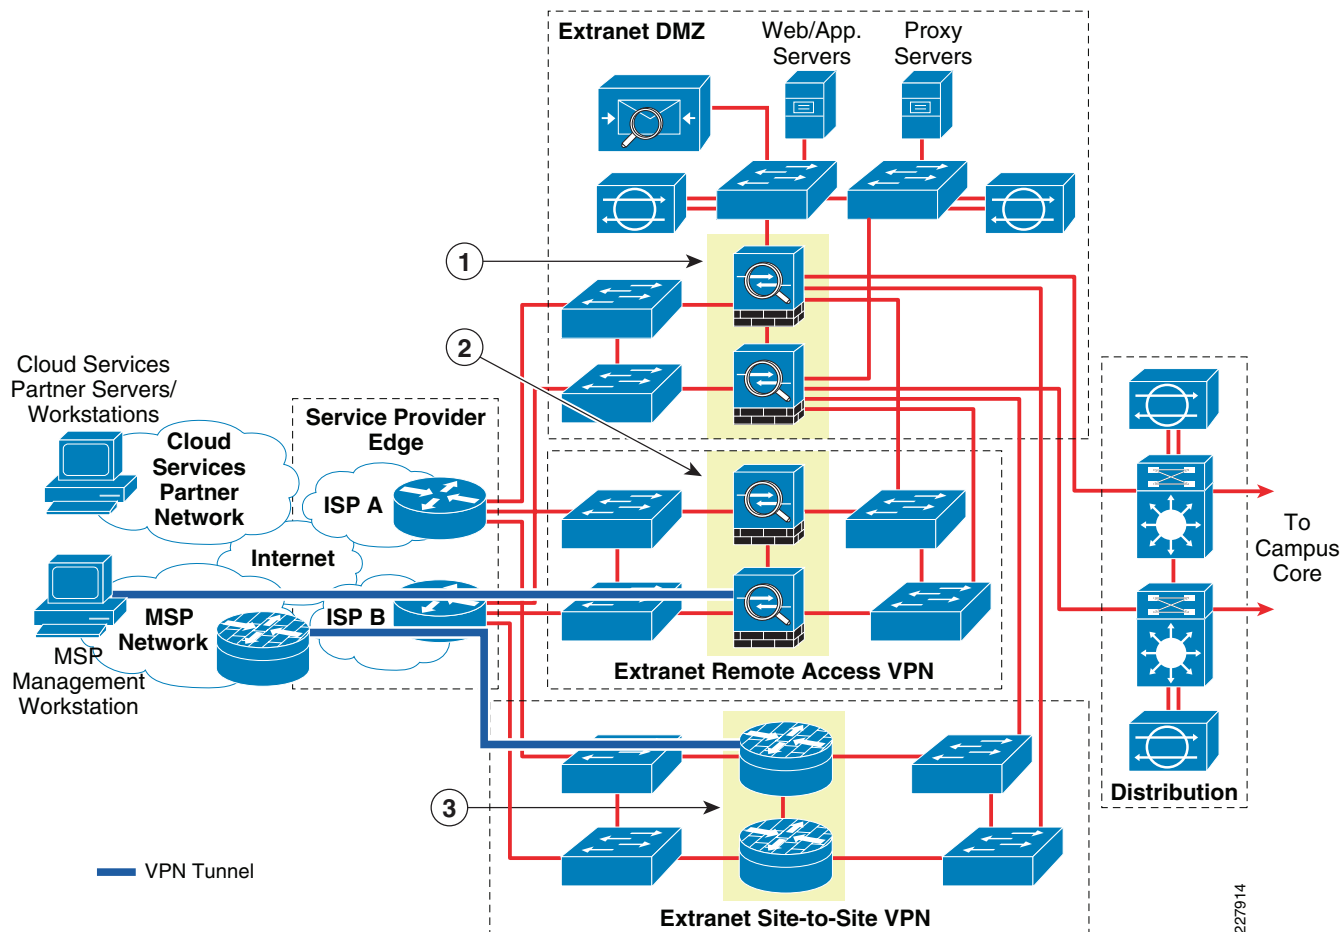
The reader should note that the remote access VPN section of the Internet Edge Module is intended for employee VPN access only. Partner VPN access is instead handled within the Partner Extranet Module, discussed in the “[Partner Extranet Module](#)” section on page 23. This provides a clean separation of traffic between internal employees and partners accessing the enterprise network.

Partner Extranet Module

When managed service provider (MSP) partner access is required, separate VPN connectivity to each campus building is possible but not scalable. Instead, it is recommended to centralize the VPN access from the MSP, providing a more maintainable and cost effective solution. In terms of the energy management solution, the primary function of the Partner Extranet Module is therefore 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 which 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 18](#).

Figure 18 Example Redundant Partner Extranet Module Design



The following steps describe what occurs in [Figure 18](#):

-
- Step 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 VPN connections from Managed Service Provider (MSP) partners.
 - Step 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.
 - Step 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, address translation, and stateful access control of managed services partner VPN connections.
-

MSP partner VPN connections can either be site-to-site or remote access. Site-to-site VPN connections are generally considered for more permanently connected requirements, meaning that a VPN tunnel is normally always established between the MSP network and the enterprise customer network. Site-to-site VPN connectivity can be provided with a pair of Cisco 1000 Series Advanced Services Routers (ASRs)

or a 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. Typical site-to-site VPN connectivity utilizes IPSec with AES 128-bit or higher encryption for data confidentiality and integrity. Multiple MSP management workstations may be allowed access through the tunnel to manage the Mediator deployment. Access control can be accomplished via Zone-Based Policy Firewall functionality on the Cisco 1000 Series Advanced Services Router (ASR) or via either Zone-Based Policy Firewall functionality or CBAC functionality on the Cisco 7200 and 7300 Series routers.

In alignment with the security concept of “defense-in-depth”, the data flows from the site-to-site VPN routers can be routed to a separate segment off of the extranet DMZ firewalls, as shown in [Figure 18](#). Access control configured on the Partner Extranet DMZ firewall allows the MSP to access only the energy management systems subnets (i.e., subnets with Mediators deployed) throughout the enterprise network. This provides a second line of access control, as well as a single point of entry for partner data flows into the enterprise network. A separate pair of ASA 5500 Series Security Appliances located within the Extranet DMZ section of the Partner Extranet Module can provide this functionality. 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 firewall. This is one advantage of deploying a separate pair of firewalls within an Partner Extranet Module.

Remote access VPN connections are generally considered for more temporarily connected requirements, although such VPN connections can be left up for extended periods of time as well. With remote access VPN connectivity, individual MSP PCs establish tunnels to a VPN concentrator device in order to manage the Mediator deployment. 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 Extranet Module. Typical MSP remote access VPN connections also utilize IPSec with AES 128-bit or higher encryption for data confidentiality and integrity. This requires client software, such as the Cisco VPN Client, to be deployed on the MSP management PCs.

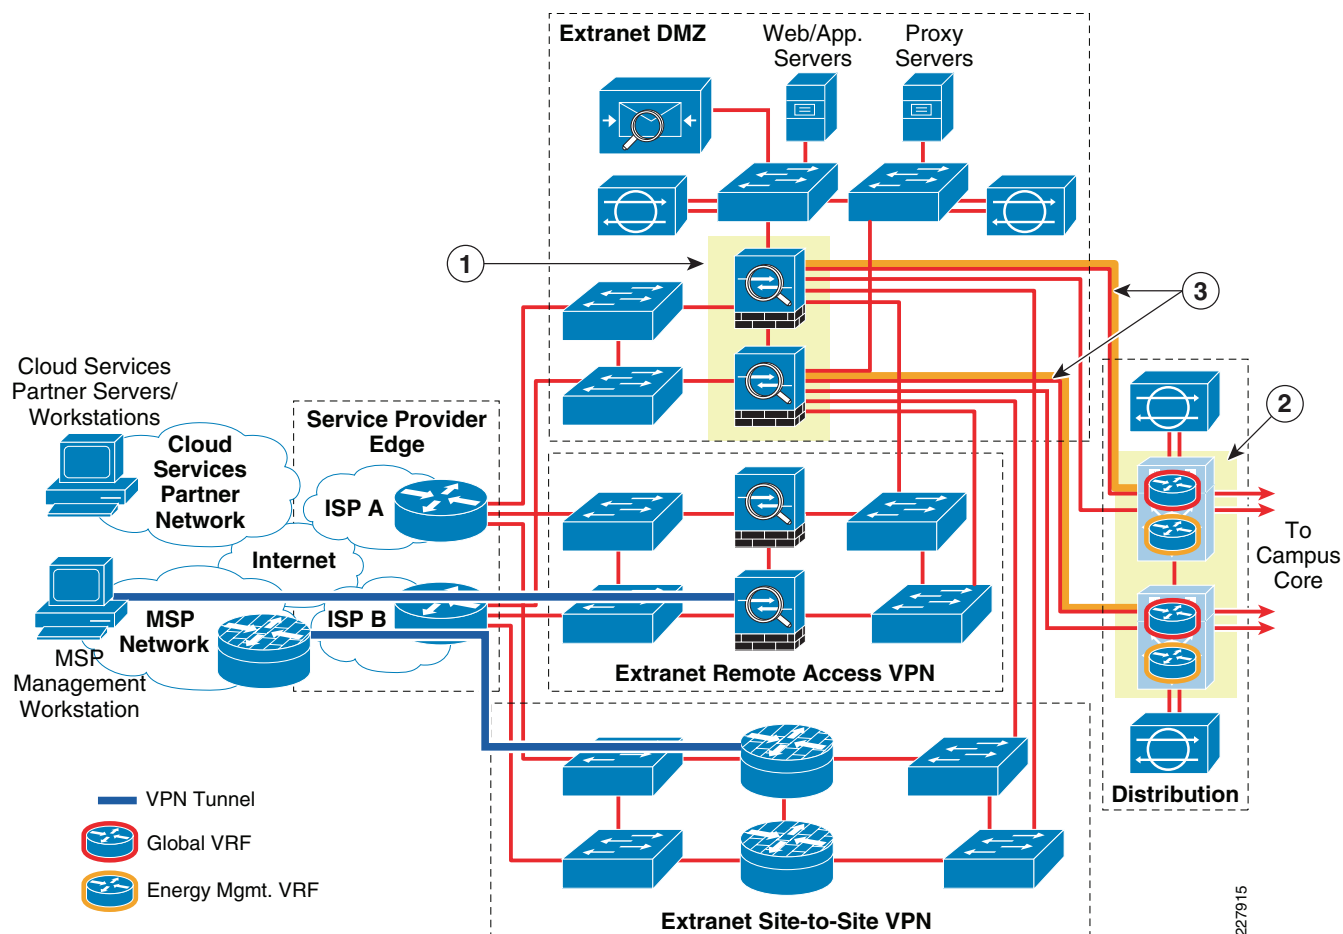
One advantage of remote access VPN connectivity is that access control can be provisioned on a per-group and/or per-user basis. The access control decision can also be centralized through a AAA server connected to the ASA 5500 Series Security Appliance via the RADIUS protocol. This can be accomplished via a Cisco Secure Access Control Server deployed either within a Data Center Service Module or a Campus Service Module within the campus network. 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 PC. 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.

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 18](#). This provides a second layer of access control and provides a single point of entry for partner traffic into the enterprise network.

Extending VRFs to the Partner Extranet Module

Figure 19 shows an example of the extension of network virtualization via VRFs to the Partner Extranet Module.

Figure 19 Energy Management VRF Extended to the Partner Extranet Module



The following steps describe what occurs in Figure 19:

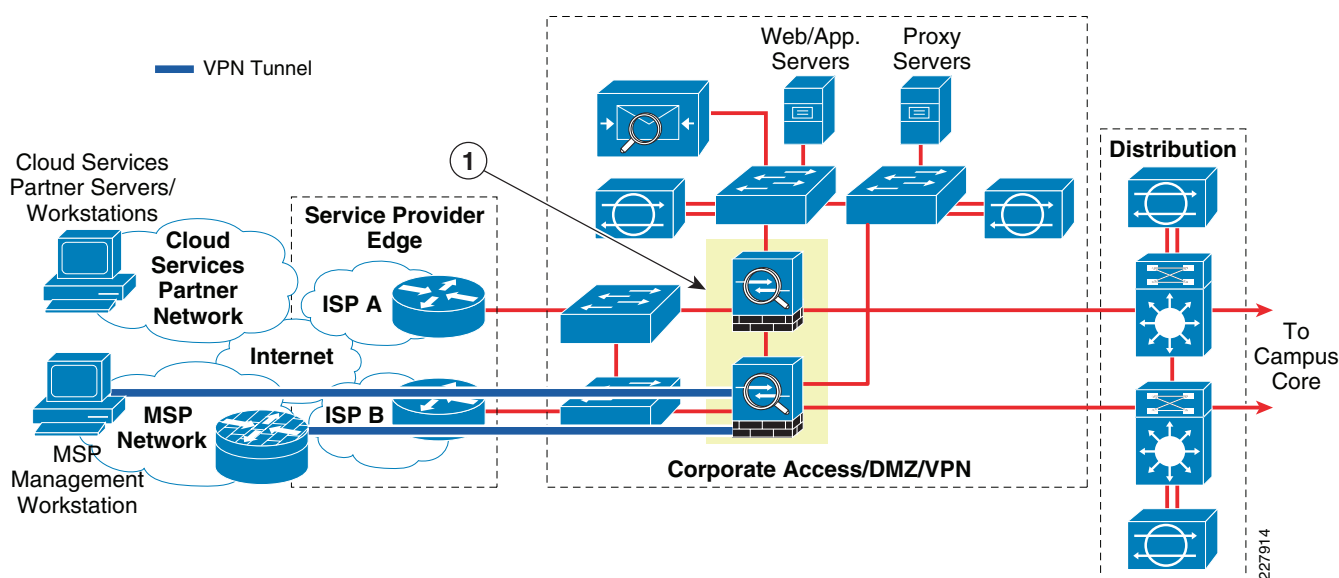
- Step 1** ASA 5500 Security Appliances deployed within the extranet DMZ section of the Extranet Module provides address translation and stateful access control for outgoing connections to cloud services partners and incoming VPN connections from managed services partners.
- Step 2** Catalyst 6500 switches within the distribution section of the Extranet Module extends the energy management Virtual Routing and Forwarding (VRF) instance from the campus.
- Step 3** Traffic destined for the energy management VRF is routed to separate interfaces on the Extranet DMZ firewalls.

By extending an separate Energy Management VRF to the Partner Extranet Module, the network administrator can ensure isolation of the MSP partner VPN traffic from the rest of the enterprise data traffic as it enters the enterprise network. The network administrator can also route the exported log data from the Cisco Network Building Mediators out the Extranet DMZ firewall (either direct or via proxy) instead of the Internet Edge firewall, which is used for the rest of the enterprise data traffic.

Collapsed Internet Edge Design

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 20 shows an example of a collapsed Internet edge design, as it applies to the energy management solution.

Figure 20 Example of a Collapsed Internet Edge Design



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 in order to provide path isolation for traffic to and from the Cisco Network Building Mediators, although this is not shown in Figure 19.

Campus Building Module

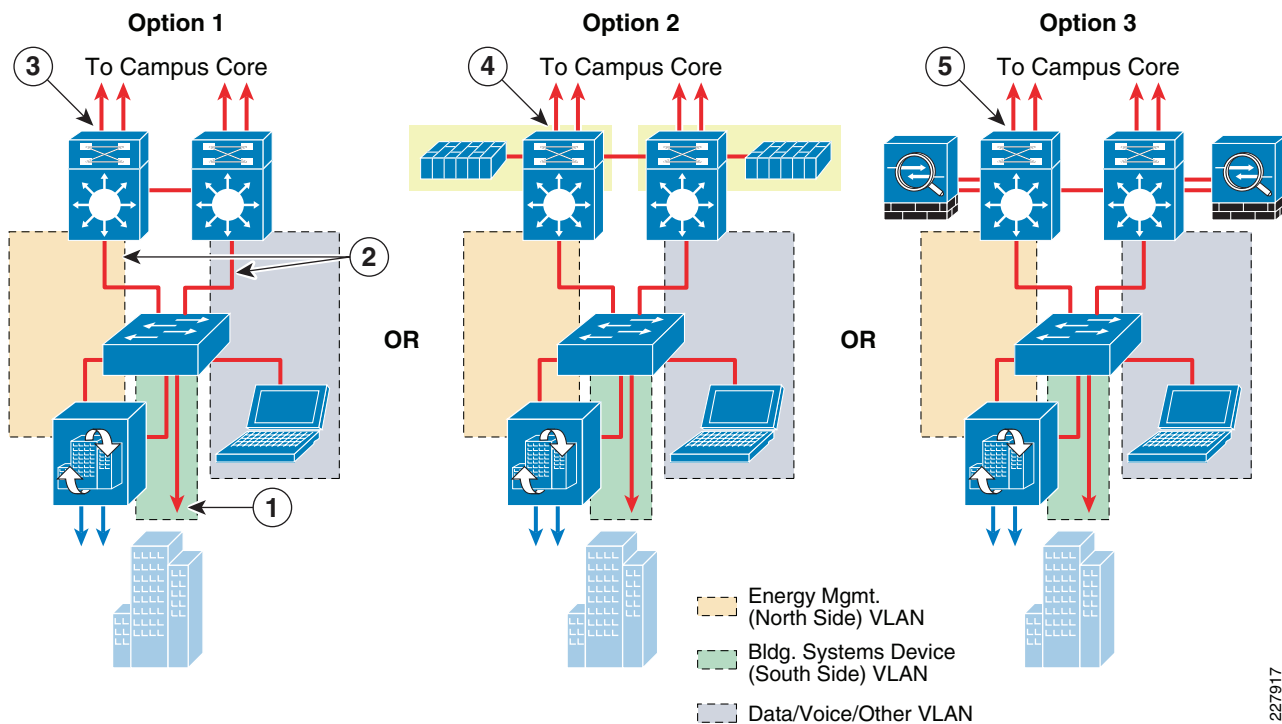
In terms of the energy management solution, the function of the Campus Building Module is to provide network connectivity for the Cisco Network Building Mediators deployed within a campus environment. It also provides network connectivity and isolation for any Ethernet-based building devices which utilize 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 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 and Layer-3 access switches are discussed in the [“Layer-2 Access Layer Switch Designs”](#) section on page 28 and the [“Layer-3 Access Layer Switch Designs”](#) section on page 30.

Layer-2 Access Layer Switch Designs

Figure 21 shows three examples of a Campus Building Module with support for the energy management solution using a Layer-2 access switch.

Figure 21 Layer-2 Access Switch Designs—Deployment Options



The following steps describe what occurs in [Figure 21](#):

-
- | | |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Building systems device VLAN isolated by not trunking it to the distribution-layer switch. |
| Step 2 | Both the energy management and the data/voice/ other VLANs trunked to the distribution-layer switch. |
| Step 3 | Layer-3 switches with ACLs at the distribution layer provide stateless access control to and from the energy management VLAN. |
| Step 4 | Layer-3 switches with FWSM at the distribution layer provide stateful access control to and from the energy management VLAN. |
| Step 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 which utilize 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.

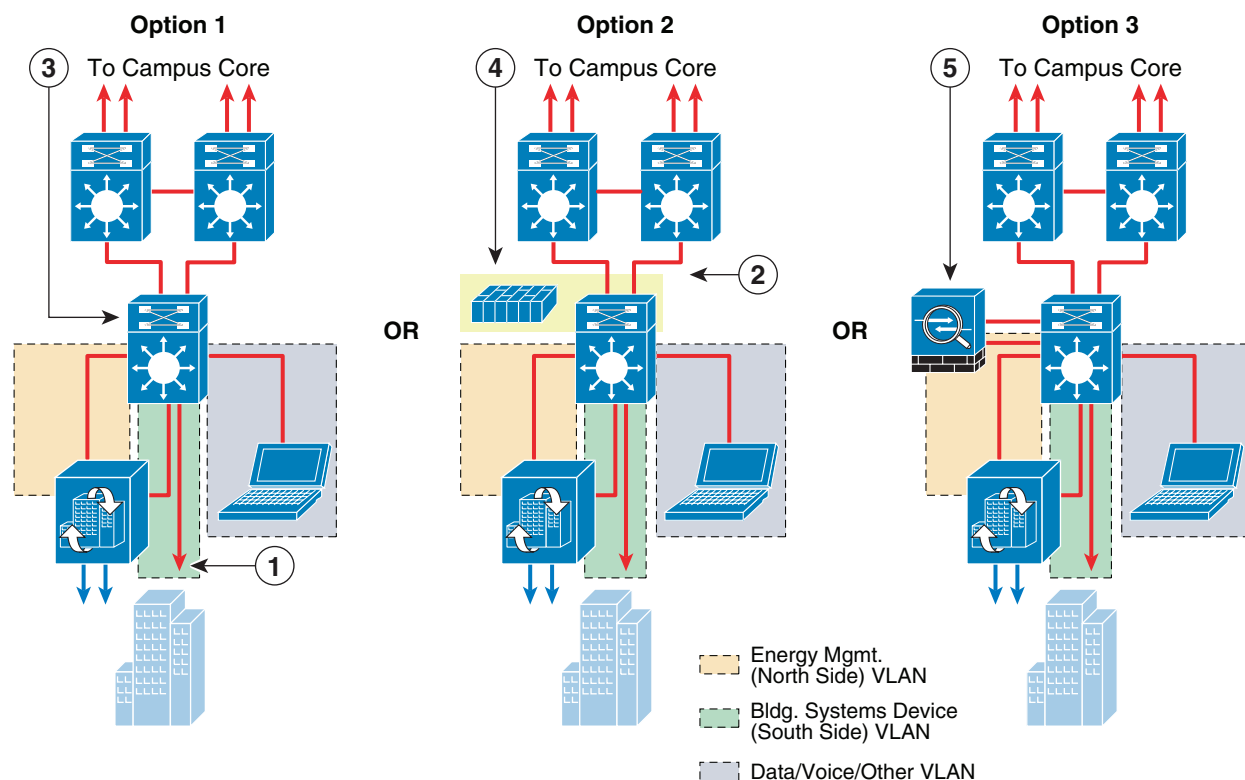
In the first deployment option shown in [Figure 6](#), access to and from the Mediator is controlled via ACLs applied to switched virtual interface (SVI) defined for the energy management VLAN on the Layer-3 distribution switch. For network administrators who desire or require stateful firewalling of 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 routed through the firewall.

The choice of the deployment option often depends upon 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 Cisco Network Building Mediator depends largely upon the number of points which 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 22 shows three examples of a Campus Building Module with support for the energy management solution using a Layer-3 access switch.

Figure 22 Layer-3 Access Switch Designs—Deployment Options



The following steps describe what occurs in Figure 22:

-
- Step 1** Building systems device VLAN isolated by not configuring an SVI for the VLAN at the Layer-3 access switch.
 - Step 2** Routed uplinks between the access and distribution switches.
 - Step 3** Layer-3 access switch provides stateless access control to and from the energy management VLAN via ACLs.
 - Step 4** Layer-3 access switch with the FWSM provides stateful access control to and from the energy management VLAN.
 - Step 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. SVIs are defined for the energy management VLAN along with any data/voice/other VLANs. However, an

SVI 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.

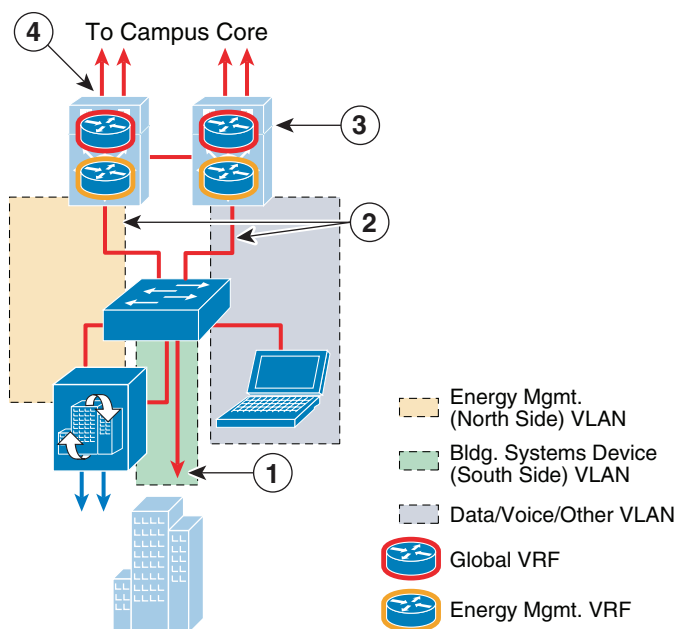
In the first deployment option shown in [Figure 22](#), access to and from the Mediator is controlled via ACLs applied to the SVI defined for the energy management VLAN on the Layer-3 access switch. For network administrators who desire or require stateful access control to the energy management VLAN, 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 routed through the firewall.

The choice of the deployment option again often depends upon 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 23](#).

Figure 23 *Layer-2 Access Switch Campus Module Design with Energy Management VRF*



The following steps describe what occurs in [Figure 23](#):

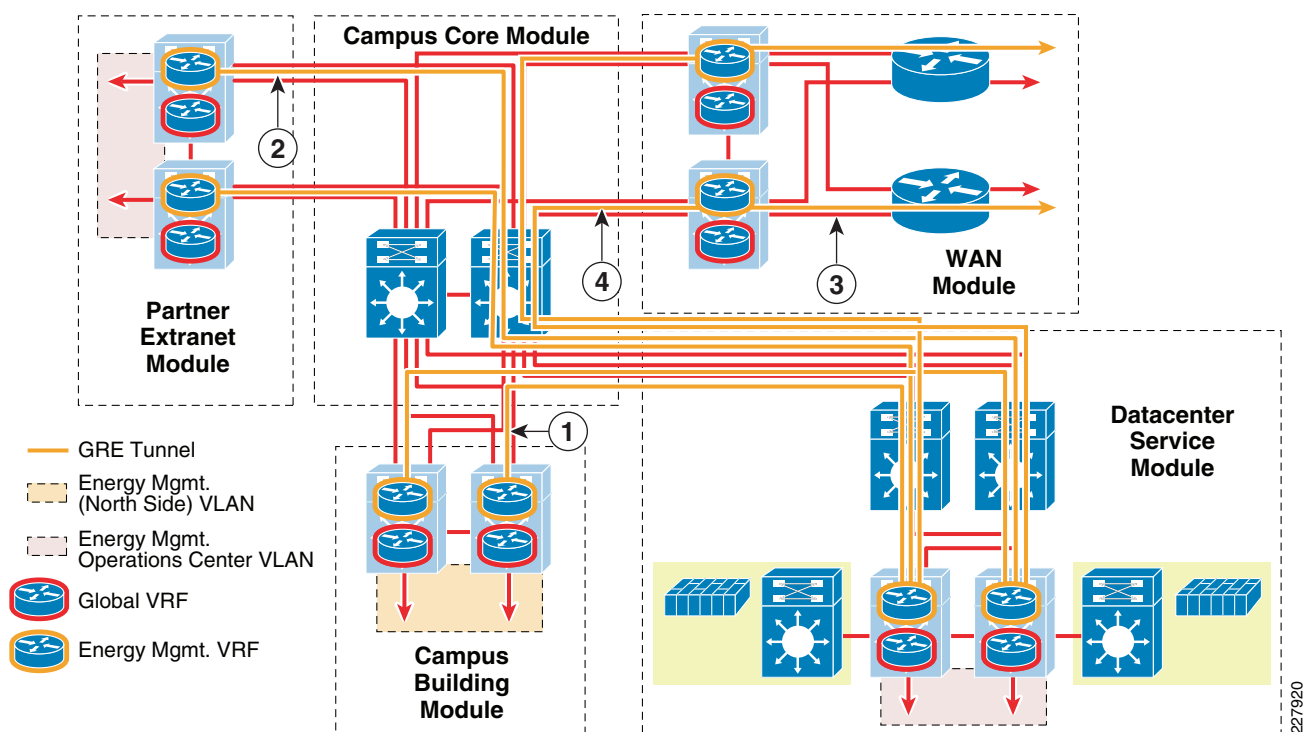
-
- Step 1** Building systems device VLAN isolated by not trunking it to the Layer-3 distribution switch.
 - Step 2** Both the energy management and the data/voice/other VLANs trunked to the Layer-3 distribution switch.
 - Step 3** Energy management VLAN mapped to the energy management VRF, while data/voice/other VLANs mapped to the global VRF at the Layer-3 distribution switch.
 - Step 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. 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. 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. GRE tunnels can be defined from the Campus Building Module Layer-3 switches to Layer-3 switches which 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 hierarchical Mediator support via portal functionality (discussed in the [“Data Center/Campus Service Module”](#) section on page 38). An example of the VRF-Lite with GRE deployment model is shown in Figure 24.

Figure 24 Energy Management Solution Utilizing VRF-Lite with GRE Tunnels



The following steps describe what occurs in Figure 24:

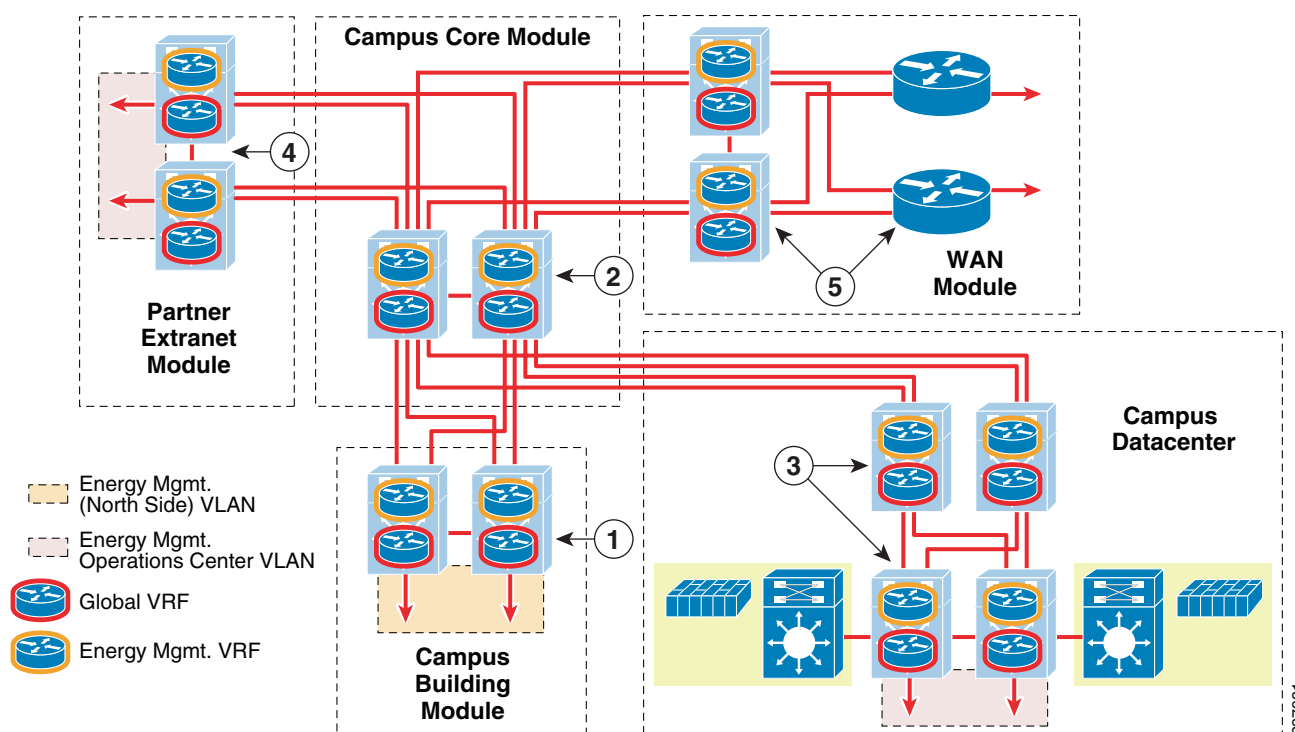
- Step 1** GRE tunnels connect energy management VRF within the campus building module to the Datacenter Service Module.
- Step 2** GRE tunnels connect energy management VRF within the partner Extranet Module to the Datacenter Service Module.
- Step 3** Multiple GRE tunnels connect energy management VRFs within branch locations to the WAN Module.

- Step 4** GRE Tunnels connect the energy management VRF within the WAN Module to the Datacenter Service Module. Routing between GRE tunnels may occur both within the Datacenter 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 the Energy Management Operations Center 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 which 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 25.

Figure 25 Energy Management Solution Utilizing VRF-Lite End-to-End



The following steps describe what occurs in Figure 25:

- Step 1** Energy management VLAN mapped to energy management VRF at campus Module Layer-3 distribution switches.
- Step 2** Energy management and global VRFs extended across the campus core module switches.

- Step 3** Energy management and global VRFs extended to the campus Datacenter Service Module. Energy Management Operations Center (EMOC) VLAN mapped to energy management VRF within the campus Datacenter Service Module.
 - Step 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.
 - Step 5** Energy management and global VRFs extended to the WAN Module Layer-3 switches 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.

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. Two different methods are discussed in this document:

- 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 26](#).

Figure 26 *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 26](#) are simply suggestions. There is no strict requirement for specific traffic types to be placed into a particular traffic classes. [Figure 27](#) 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 27 Possible Mapping of Energy Management Traffic to 12-Class QoS Model

Application Class	PHB
Network Control	CS6
Broadcast Video	CS5
VoIP Telephony	EF
Multimedia Conferencing	AF4
Realtime Interactive	CS4
Multimedia Streaming	AF3
Call-Signaling	CS3
Transactional Data	AF2
OAM	CS2
Bulk Data	AF1
Scavenger	CS1
Best Effort	default

Periodic Export of Logged Data via HTTP, and/or HTTPS

Management of the Mediator via HTTP, HTTPS, and/or SSH; as well as RNA Flows and Events Sent via SMTP

Periodic Export of Logged Data via FTP, SFTP (SSH), and/or SMTP

227923

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 every 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—utilizing 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 video traffic types. In terms of the actual management of the Cisco Network Building 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 ingress ACLs applied to the access switch ports to which the Cisco Network Building 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 and mark all traffic corresponding to that protocol to a particular service class. The reader should note that often the same protocol can 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 Cisco Network Building 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. Alternatively, all traffic inbound on the energy management VLAN itself can be remarked to the OAM service class.

The choice of which method to implement—either different service classes for different traffic types, or a single service class for all traffic types—is really a matter of preference by the end customer. In either scenario, the objective is to provide a service class for the energy management traffic which 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 Best Effort 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 class of service is maintained on the MSP network.

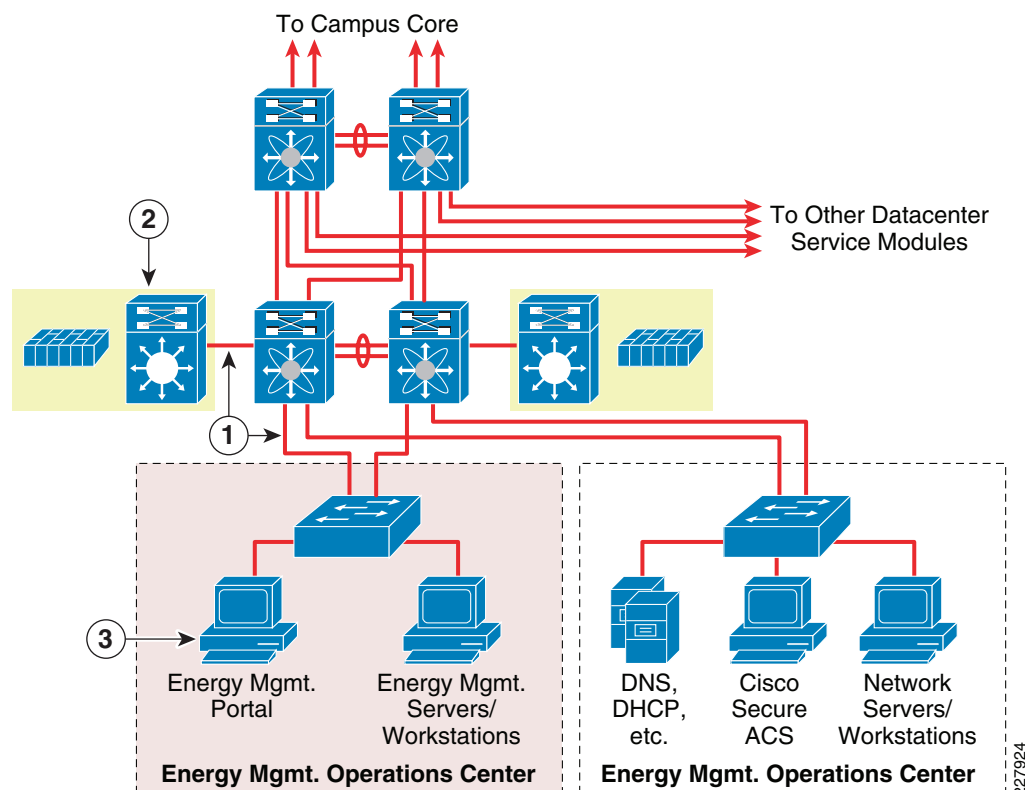
Data Center/Campus Service Module

Enterprise customers typically 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.

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. [Figure 28](#) shows an example of this design.

Figure 28 Data Center Service Module Design with Catalyst 6500 Service Switch and FWSM



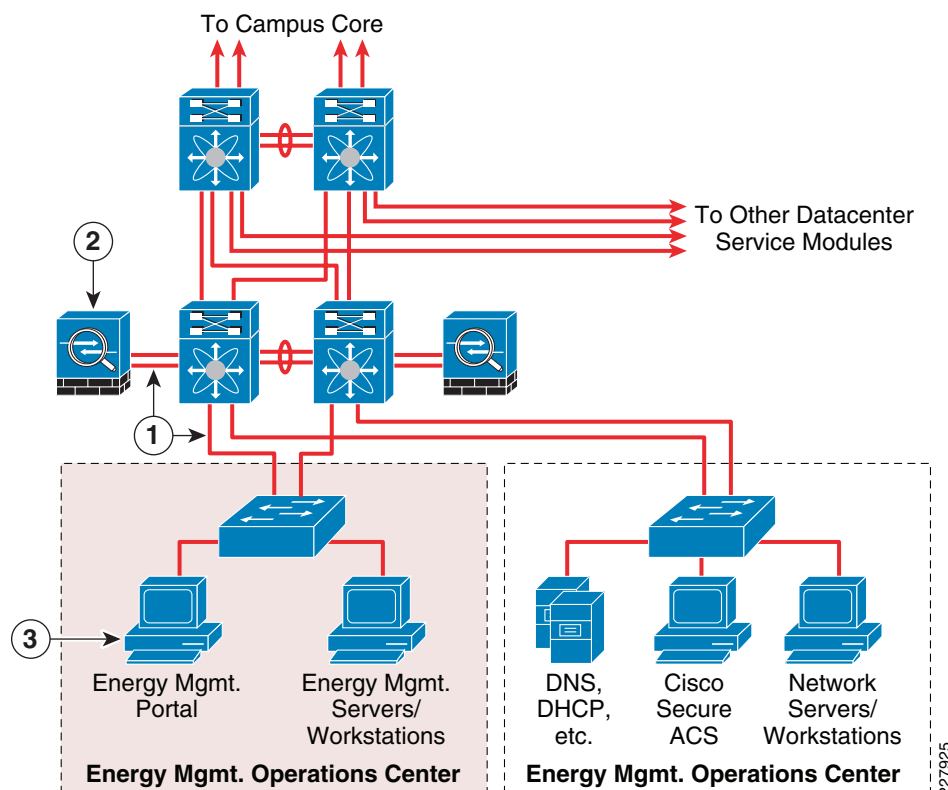
The following steps describe what occurs in [Figure 28](#):

-
- Step 1** Energy Management Operations Center VLAN trunked through Nexus 7000 distribution switch and Catalyst 6500 Service switch to the FWSM.
 - Step 2** Service switches with FWSM in the data center provide stateful access control to and from the energy Management Operations Center devices.
 - Step 3** Energy Management Portal can provide single access point to Mediators deployed throughout the enterprise network.
-

In this example, a separate EMOC VLAN is implemented within the Data Center Service Module. The EMOC VLAN is trunked from the access switches, 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 provides stateful access control to and from the EMOC VLAN.

An alternative to the Catalyst 6500 service switch design is to implement a set of ASA 5500 Security Appliances within the Data Center Service Module, as shown in [Figure 29](#).

Figure 29 Data Center Service Module Design with ASA 5500 Security Appliances

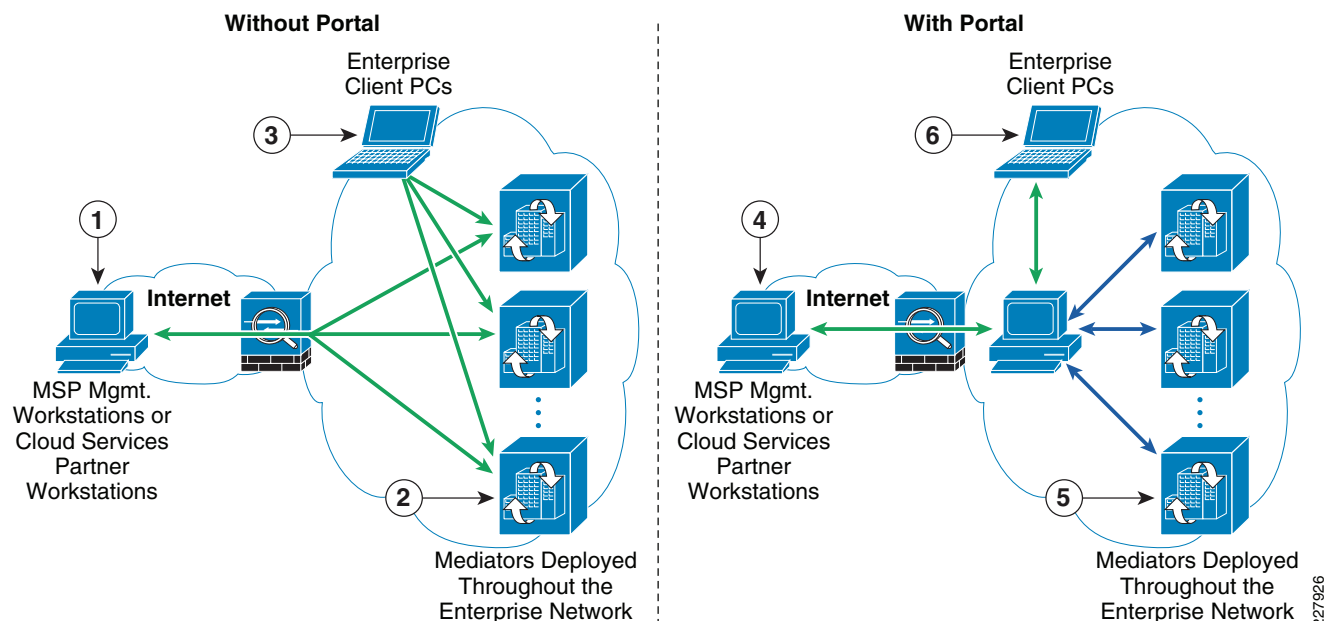


The following steps describe what occurs in [Figure 29](#):

-
- Step 1** Energy Management Operations Center VLAN trunked through Nexus 7000 distribution switch to ASA 5500 Security Appliances.
 - Step 2** ASA 5500 Security Appliances in the data center provide stateful access control to and from the Energy Management Operations Center devices.
 - Step 3** Energy Management Portal can provide single access point to Mediators deployed throughout the enterprise network.
-

In this design, the EMOC VLAN is trunked from the access switch, through the Nexus 7000 Series data center distribution switches, to a Layer-3 interface of the ASA 5500 Security Appliances. The ASA 5500 provides stateful access control to and from the EMOC VLAN. 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 28](#) and [Figure 29](#). 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 Managed Service Partner (MSP) VPN access to the enterprise network.

Portal functionality can also be deployed within the EMOC VLAN in order to provide a single hierarchical point of access for all of the Mediators deployed within the enterprise organization versus having to access each Mediator individually, as shown in [Figure 30](#).

Figure 30 Mediator Access With and Without Portal Functionality

The following steps describe what occurs in [Figure 30](#):

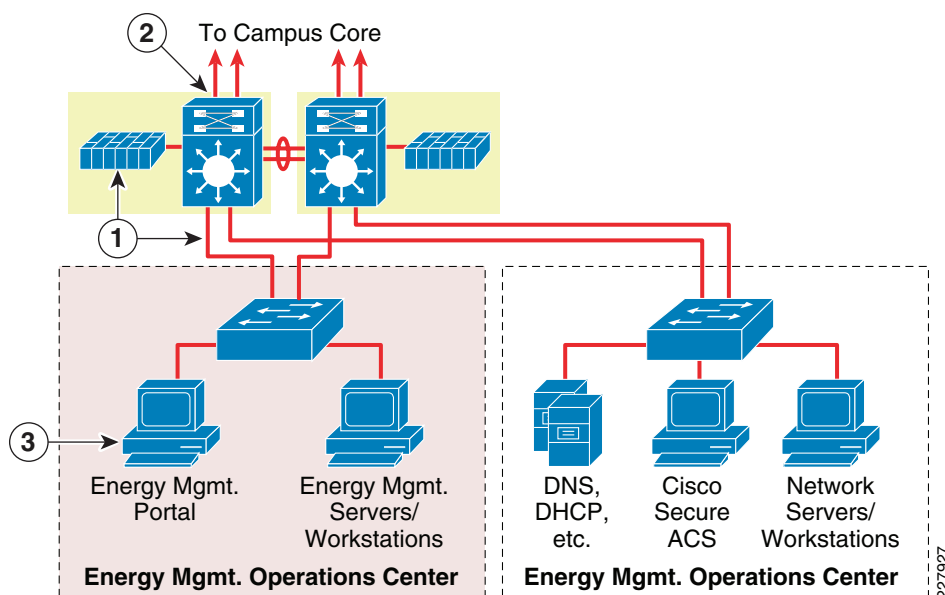
-
- Step 1** Managed Service Provider needs access to each Mediator in order to manage the deployment.
 - Step 2** Mediators individually export logged data to Cloud Services Partners.
 - Step 3** Individual enterprise client PCs may need access to individual Mediators deployed throughout the network to view information.
 - Step 4** Managed service provider needs an access to the Energy Management Portal in order to manage the deployment. Individual Mediators are accessed through the portal.
 - Step 5** Mediators send logged data to the portal, which periodically exports it to Cloud Services Partners.
 - Step 6** Individual enterprise client PCs access the portal to view information on individual Mediators deployed throughout the network.
-

One advantage of implementing portal functionality is that it may be possible to restrict MSP VPN access to only the portal device itself. The MSP would then manage the rest of the Mediators through the portal. Further, the individual Mediators can send logged data to the portal, which in turn periodically exports the data to cloud service partner servers located on the Internet. Finally, the portal functionality may also provide a centralized, and therefore more secure, mechanism for client PCs located within the enterprise network to access information from the Mediators, if this is a requirement.

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. [Figure 31](#) shows an example of this design.

Figure 31 *Campus Service Module Design with Catalyst 6500 Switch and FWSM*

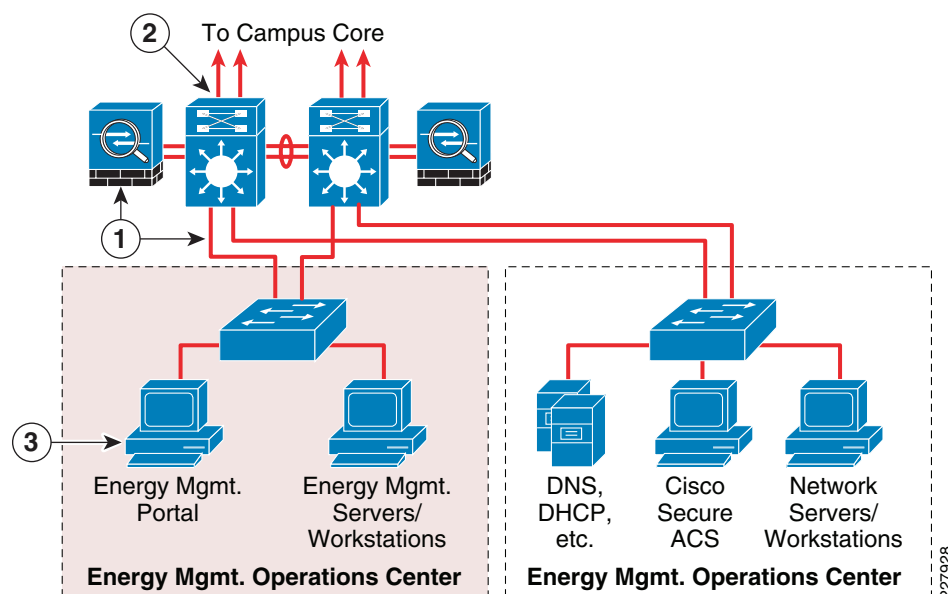


The following steps describe what occurs in [Figure 31](#):

-
- Step 1** Energy Management Operations Center VLAN trunked through Catalyst 6500 Service switch to the FWSM.
 - Step 2** FWSM in the Campus Service Module Catalyst 6500 switches provide stateful access control to and from the Energy Management Operations Center devices.
 - Step 3** Energy Management Portal can provide single access point to Mediators deployed throughout the enterprise network.
-

In this example, the EMOC VLAN is trunked from the access switch to a Layer-3 interface of the FWSM module located within the Catalyst 6500 Campus Service Module switch. The FWSM provides stateful access control to and from the EMOC VLAN. 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 32](#).

Figure 32 *Campus Service Module Design with Catalyst 6500 Switch and ASA 5500 Security Appliance*



The following steps describe what occurs in [Figure 32](#):

-
- Step 1** Energy Management Operations Center VLAN trunked through Catalyst 6500 Service switch to the ASA 5500 Security Appliance.
 - Step 2** ASA 5500 Security Appliances provide stateful access control to and from the Energy Management Operations Center devices.
 - Step 3** Energy Management Portal can provide single access point to Mediators deployed throughout the enterprise network.
-

In this design, the EMOC VLAN is trunked from the access switch, through the Campus Service Module switch, to a Layer-3 interface of the ASA 5500 Security Appliances. With this design 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. As with the data center designs, other VLANs such as a Network Operations Center (NOC) VLAN can also be supported off the same Campus Service Module. Likewise portal functionality can be deployed within the EMOC VLAN hanging off the Campus Service Module.

Extending VRFs to the Data Center/Campus Service Module

When network virtualization is implemented in order to provide path isolation for the energy management solution, the Data Center/Campus Service Module serves an additional function. Since a separate energy management VRF 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 branch. Instead the stateful firewall within the Data Center/Campus Service Module (either a FWSM or ASA 5500 Security Appliance) can serve as the single point of access control between the global VRF and the energy management VRF. This 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. ACLs may still be utilized to specifically

control traffic to and from each VLAN within the energy management VRF in order to control access to the Mediators. This is particularly useful when the network virtualization concept is expanded to include a single building automation VRF instead, which includes other functionality such as video surveillance and physical access control, as well as energy management. The reader should note that the deployment of network virtualization can also effectively isolate inbound MSP partner VPN traffic to the energy management VRF, reducing the security exposure of partner traffic mixed in with enterprise data, voice, and video traffic on the global VRF. As previously shown in [Figure 23](#) and [Figure 24](#), the choice of whether VRF with GRE or VRF-Lite End-to-End is deployed within the campus determines which Data Center/Campus Service Module switches need to support VRFs and/or GRE tunnels.

QoS within the Data Center/Campus Service Module

The access switches within the Data Center/Campus Service Module house the energy management servers and workstations which manage and possibly collect periodic log information from the Mediators deployed throughout the enterprise organization. Also, portal functionality may be housed on a device connected to the access switch with the Data Center/Campus Service Module. Typically most management workstations and servers have no capability to mark traffic to any service class other than Best Effort (DSCP value = 0). In order to classify traffic flows originating from the management workstations and/or portal device into anything other than the Best Effort service class, the classification and re-marking can be performed at the ingress port of Data Center/Campus Service Module access switch. The two different methods discussed in the [“QoS within the Campus Building Module” section on page 35](#) apply equally to the application of QoS within the Data Center/Campus Service Module. As applied to the management servers and/or portal device, the choices are:

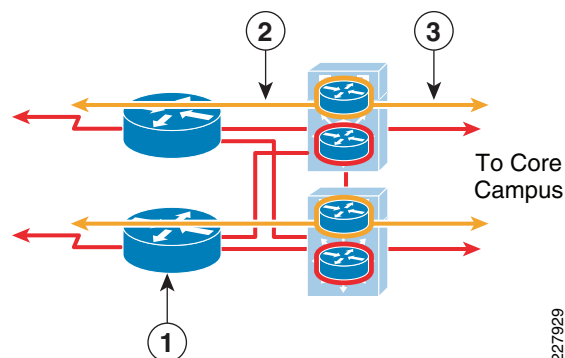
- Identifying and marking individual traffic flows from the management servers and/or portal device to different service classes based upon the traffic type (FTP, HTTP, SSH, etc.) and use (periodic data export or configuration and management).
- Identifying and marking all traffic flows from the management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

Although not necessarily a requirement, it is recommended that the network administrator implement the same QoS classification and marking method for traffic originating from the management servers and portal device as the method implemented for traffic originating from the Mediators. This ensures consistency of QoS treatment of the energy management traffic flows across the 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 the MSP (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 33](#).

Figure 33 Example WAN Module with VRF Design



The following steps describe what occurs in [Figure 33](#):

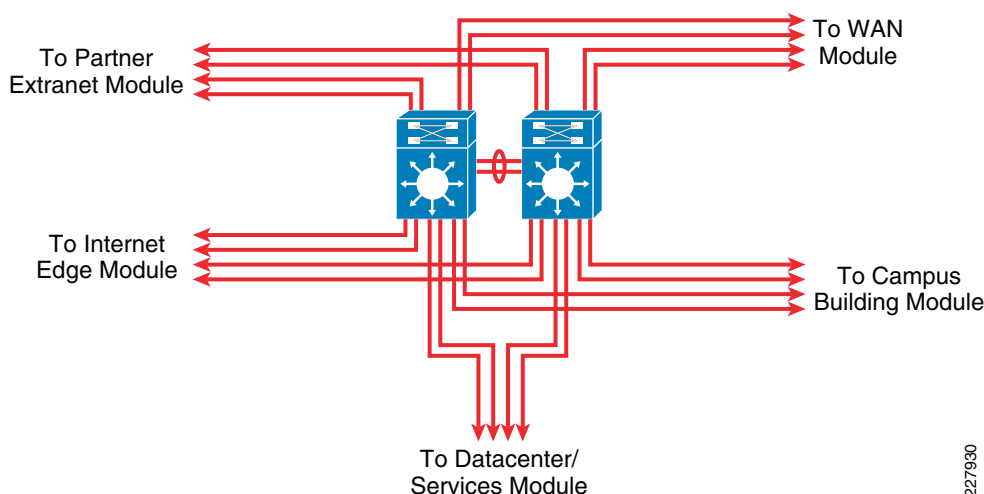
-
- Step 1** Redundant ASR-1000 Series, Cisco 7600 Series, or Cisco 7200 Series routers provide WAN Circuit Termination from branch locations.
 - Step 2** GRE tunnels from branch locations terminate on Layer-3 Catalyst 6500 distribution switches.
 - Step 3** GRE tunnels defined from Layer-3 Catalyst 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 discussed in the [“Campus Building Module”](#) section on page 28, 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. As previously mentioned in the [“Branch Network Design Considerations”](#) section on page 14, other methods of supporting VRFs across the WAN, such as mapping them to an MPLS service, also exist.

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 34](#).

Figure 34 **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 model discussed in the “[Campus Building Module](#)” section on page 28, 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, also discussed in the “[Campus Building Module](#)” section on page 28, then the core Catalyst 6500 switches must be configured to support VRFs as well.

Summary

The Cisco Network Building Mediator is the centerpiece of the open sustainability and energy management solution, aggregating and normalizing energy management systems and making them available through an open XML interface over an IP network infrastructure. As the Mediator interfaces with critical energy management systems, the design engineer must be aware of the requirements for tight access controls. Access control can be accomplished through multiple mechanisms, including IPSec VPN connectivity, firewall appliances, VRF segmentation, integrated firewall services, and ACLs within router and switch platforms.

For more information, refer to the following URL: <http://www.cisco.com/go/designzone>