Small Enterprise Design Profile Reference Guide

Last Updated: July 8, 2010

Cisco Validated Design

Building Architectures to Solve Business Problems

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit www.cisco.com/go/designzone.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2010 Cisco Systems, Inc. All rights reserved

Solution Authors

Martin Pueblas

Martin Pueblas, CCIE#2133, CISSP#40844—Technical Leader, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Martin is the lead system architect of the Cisco SAFE Security Reference Architecture. He is a network security expert with over 17 years of experience in the networking industry. He obtained his CCIE certification in 1996 and CISSP in 2004. Martin joined Cisco in 1998 and has held a variety of technical positions. Started as a Customer Support Engineer in Cisco's Technical Assistance Center (TAC) in Brussels, Belgium. In 1999 moved to the United States where soon became technical leader for the Security Team. Martin's primary job responsibilities included acting as a primary escalation resource for the team and delivering training for the support organization. At the end of 2000, he joined the Advanced Engineering Services team as a Network Design Consultant, where he provided design and security consulting services to large corporations and Service Providers. During this period, Martin has written a variety of technical documents including design guides and white papers that define Cisco's best practices for security and VPNs. Martin joined Cisco's Central Marketing Organization in late 2001, where as a Technical Marketing Engineer, he focused on security and VPN technologies. In late 2004, he joined his current position acting as a security technical leader. As part of his current responsibilities, Martin is leading the development of security solutions for enterprises.

Steve Gyurindak

Steve Gyurindak, CCIE#9057, CISSP#61046—Solutions Architect, Enterprise Solutions Engineering (ESE), Cisco Systems

Steve is a solutions architect with over 15 years of industry experience. He joined Cisco in 2000 and worked the first 8 and a half years as a Systems Engineer covering the Service Provider, North Florida/Alabama Commercial, Georgia Enterprise and US Channels sales markets. Steve has been recognized for his work with some of Cisco's most influential customers as well as for his work in South America and Europe. Steve joined ESE in 2009 to lead the development of customer-focused architectures and designs for the Education Market. Steve has a Bachelor of Science degree in Telecommunications from the State University of New York at Buffalo, and is currently pursuing a Master's of Science degree in Network Telecommunications at New York University. In addition to a CCIE in Routing and Switching, Steve holds the following certifications: CISSP, CCNP, CCDP, CCNA, CCDA, MCSE, and MCNE.

John Strika

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

John is a Technical Marketing Engineer in Cisco's Public Sector ESE team, with expertise in the areas of mobility and location-based services. He has coauthored documents on enterprise mobility and Wi-Fi location-based services. As a member of Cisco's Enterprise Architecture Board, he helps maintain Cisco's vision and architectural direction and define Cisco's roadmap for context-aware and presence solutions. Previously, John was Cisco's first mobility consulting systems engineer, responsible for architecting creative wireless solutions for large enterprise customers. His 28 years of experience spans network design and implementation, applications development, facilities planning and management, consulting, and general management. His past roles have included mission-critical telecommunications design and development at AT&T and systems programming and data communications management with Wall Street brokerages and commercial banks. Prior to joining Cisco, Strika was at Telxon Corporation (parent of Cisco's Aironet wireless acquisition) for nine years, reaching the position of Southern Division Vice President of Wireless Technologies and Services. He is a member of the IEEE and has held several Federal Communications Commission licenses in the use and modification of amateur and commercial radio. His educational background is in electrical engineering and computer applications programming from Columbia University and in finance from Fordham University's College of Business Administration, and he holds a masters of communications technology certificate from the American Institute. He was a charter Novell Certified Netware Engineer in the greater New York City area. Always seeking opportuni-

Solution Authors

ties to use his mobility and advanced communications knowledge to improve public safety as well as the safety of our public servants, John has served in volunteer search and rescue as well as a Reserve Deputy.

Rahul Kachalia

Rahul Kachalia, CCIE#11740—Technical Marketing Engineer, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Rahul is a technical marketing engineer in Cisco's Enterprise Solution Engineering group, helping to create the design guidance that will help build the 21st century school network infrastructure. Rahul has more than 14 years of broad engineering experience, primarily in service provider core and edge focused products and technologies including broadband, MPLS, VPN and managed services. He has led many assurance projects to develop solutions that can deliver design guidance and accelerate deployments from traditional WAN infrastructure to next-generation IP/MPLS managed core networks. In the Enterprise Solution Engineering group he has also worked on designing next-generation unified virtual campus networks for large enterprise customers. In addition to CCIE, Rahul holds CCNP, CCNA, MCSE, MCP, and CNE. He holds a bachelor's degree from Mumbai University, India.

Dan Hamilton

Dan Hamilton, CCIE #4080 —Technical Leader, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Dan has over 15 years experience in the networking industry. He has been with Cisco for 9 years. He joined Cisco in 2000 as a Systems Engineer supporting a large Service Provider customer. In 2004, he became a Technical Marketing Engineer in the Security Technology Group (STG) supporting IOS security features such as infrastructure security, access control and Flexible Packet Matching (FPM) on the Integrated Security Routers (ISRs), mid-range routers and the Catalyst 6500 switches. He moved to a Product Manager role in STG in 2006, driving the development of new IOS security features before joining the ESE Team in 2008. Prior to joining Cisco, Dan was a network architect for a large Service Provider, responsible for designing and developing their network managed service offerings. Dan has a Bachelor of Science degree in Electrical Engineering from the University of Florida.

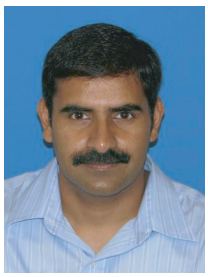

Srinivas Tenneti

Srinivas Tenneti, CCIE#10483—Technical Marketing Engineer, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Srinivas is a Technical Marketing Engineer for WAN and branch architectures in Cisco's ESE team. Prior to joining the ESE team, Srinivas worked two years in Commercial System Engineering team where he worked on producing design guides, and SE presentations for channel partners and SEs. Before that, he worked for 5 years with other Cisco engineering teams. Srinivas has been at Cisco for 8 years.

CONTENTS

CHAPTER 1 **Small Enterprise Design Profile (SEDP) Overview** 1-1

- Small Enterprise Design Profile 1-1
  - Main Site Design 1-3
  - Remote Large Site Design 1-3
  - Remote Small Site Design 1-3
  - Access Devices 1-3
- Network Foundation Design Considerations 1-3
  - LAN Design Considerations 1-4
  - High Availability Design Considerations 1-4
    - Routing Protocol Selection Criteria 1-4
  - Access Layer Design Considerations 1-5
  - LAN Foundational Services 1-5
  - WAN Design Considerations 1-6
    - WAN Transport 1-6
    - WAN Foundational Services 1-6
- Cisco SAFE Security Architecture Design Considerations 1-6
- Mobility 1-7
- Collaboration Services Design Considerations 1-7
- Unified Communications Design Considerations 1-7
  - Call Processing Considerations 1-7
  - Gateway Design Considerations 1-8
  - Dial Plan Considerations 1-8
  - Survivability Considerations 1-8
- IP Video Surveillance Design Considerations 1-8
  - Digital Media Systems 1-9
- Summary 1-10

CHAPTER 2 **Small Enterprise Design Profile(SEDP)—Network Foundation Design** 2-1

- Building Unified Small Enterprise Network Infrastructure 2-1
  - Hierarchical Network Design 2-2
  - Collapsed Core Network Design 2-4
  - Main Site Network Design 2-5
    - Resilient Distributed System 2-9

Main Site Access-Layer Edge Services 2-9
Access-Layer Network Control Services 2-11
Resilient Access-Layer Network and System 2-12
Main Site Data Center Network Design 2-13
Remote Site Network Design 2-14
Remote Site Collapsed Core Network Design 2-14
Remote Site Access-Layer Design 2-16
Deploying Network Foundation Services 2-16
Implementing EtherChannel in Network 2-16
EtherChannel Load Balancing 2-19
Deploying Core Network Layer 2-21
Routing Protocol 2-21
Routing Protocol Selection Criteria 2-21
Designing End-to-End EIGRP Routing Domain 2-22
Deploying Multi-Layer Network 2-27
Spanning-Tree in Multilayer Network 2-28
Logical Multi-Layer Network 2-29
Flat Logical Network Design 2-29
Segmented Logical Network Design 2-30
Implementing Layer 2 Trunk 2-32
Unidirectional Link Detection 2-33
Deploying Routed-Access Network 2-34
Implementing EIGRP Routing in Access-Distribution Block 2-36
Building EIGRP Network Boundary 2-37
EIGRP Adjacency Protection 2-41
Tuning EIGRP Protocol Timers 2-42
Deploying Multicast in Network 2-42
Multicast Routing Design 2-44
Deploying PIM-SM 2-46
Implementing IGMP 2-49
Multicast Security—Preventing Rogue Source 2-49
Multicast Security—Preventing Rogue RP 2-50
Deploying QoS in Network 2-51
QoS in Catalyst Fixed Configuration Switches 2-51
QoS in Cisco Modular Switches 2-52
QoS Framework 2-54
QoS Trust Boundary 2-55
Deploying QoS 2-55
Implementing QoS Trust Mode 2-56
Implementing QoS Classification 2-57

Implementing Ingress Policer 2-60
Implementing Ingress Marking 2-61
Applying Ingress Policies 2-62
Applying Ingress Queueing 2-63
Deploying Egress QoS 2-65
Deploying Network Core QoS 2-69
Deploying Main or Remote Large Site Ingress QoS 2-70
Deploying Remote Small Site Ingress QoS 2-71
Deploying Main Site Egress QoS 2-72
Deploying Remote Large Site Egress QoS 2-74
Deploying Remote Small Site Egress QoS 2-76
Building a Resilient Network 2-78
Redundant Hardware Components 2-79
Redundant Power System 2-80
Redundant Network Connectivity 2-80
Redundant Control-Plane 2-80
Operational Resiliency Strategy 2-83
Deploying High Availability in Network 2-83
Network Resiliency 2-83
Implementing IP Event Dampening 2-83
Device Resiliency 2-85
Summary 2-92

CHAPTER 3 **Small Enterprise Design Profile(SEDP)—WAN Design** 3-1
WAN Design 3-1
MPLS/VPN 3-1
Internet 3-2
Metro Ethernet 3-2
Why Metro Ethernet Service Is Needed 3-2
Types of Services Available in Metro Ethernet 3-2
WAN Service Deployed in the Small Enterprise Design 3-3
When to Consider WAN Technologies 3-4
Bandwidth Capacity 3-5
Planning 3-5
Implementation 3-7
Aggregation Structure 3-7
Routing for WAN Connections 3-8
QoS Design 3-9
QoS Policy at Main Site 3-9

QoS Policy at Remote Site 3-11

CHAPTER 4 **Small Enterprise Design Profile(SEDP)—Wireless LAN Design 4-1**

Cisco Unified Wireless Network Architecture 4-1

LWAPP Features 4-3

Small Enterprise Design Profile 4-4

Management 4-5

Connection to the Small Enterprise Design Profile Network 4-9

RF Groups and Mobility Groups 4-12

Example WLAN Configurations 4-13

Secured Employee WLAN 4-14

Secured VoWLAN 4-15

AP Deployments Considerations 4-18

AP 1250 4-18

AP 1140 4-18

Coverage and Site Surveys 4-19

Single-Band versus Dual-Band APs 4-19

WLC Discovery 4-20

WLC Failover Options 4-20

Appendix A—Devices and Software Used 4-21

CHAPTER 5 **Small Enterprise Design Profile(SEDP)—Security Design 5-1**

Small Enterprise Network Security Design 5-1

Network Foundation Protection 5-3

Internet Perimeter Protection 5-5

Internet Border Router Security Guidelines 5-6

Internet Firewall Guidelines 5-7

Intrusion Prevention Guidelines 5-10

E-mail Security Guidelines 5-14

Web Security Guidelines 5-19

Serverfarm Protection 5-23

Network Access Security and Control 5-25

Catalyst Integrated Security Features 5-26

Cisco Unified Wireless Network (CUWN) Integrated Security Features 5-26

Cisco NAC Appliance 5-27

NAC Appliance Modes and Positioning 5-29

NAC Appliance Deployment in Small Enterprise Networks 5-34

Cisco Identity-Based Network Networking Services (IBNS) 5-35

802.1X and EAP 5-36

Impacts of 802.1X on the Network 5-37
802.1X in Small Enterprise Networks 5-37
NAC 802.1X and CISF in Combination 5-38
Secure Mobility 5-38
Threats Mitigated 5-40
Network Security Deployment 5-41
Internet Border Router Deployment 5-41
Internet Firewall Deployment 5-42
Firewall Hardening and Monitoring 5-43
Network Address Translation (NAT) 5-45
Firewall Access Policies 5-45
Botnet Traffic Filter 5-47
Firewall Redundancy 5-51
Routing 5-52
Intrusion Prevention Deployment 5-55
Deploying IPS with the Cisco ASA 5-55
IPS Global Correlation Deployment 5-55
Web Security Deployment 5-60
Initial System Setup Wizard 5-61
WCCP Transparent Web Proxy 5-65
Web Access Policies 5-69
CISF Protected Ports 5-70
NAC Appliance Deployment 5-70
Adding a CAS to the CAM 5-71
Managing the CAS 5-72
Clean Access Roles 5-77
Layer 2 OOB Example 5-77
Availability Considerations 5-79
Basic Clean Access Switch Configuration 5-80
Basic Clean Access Out-of-Band Switch Configuration 5-80
NAC CAS Connection 5-80
Basic 802.1X Switch Configuration 5-81

CHAPTER 6 **Small Enterprise Design Profile (SEDP)—Context-Aware Services** 6-1
Introduction 6-1
What Are Context-Aware Services? 6-1
Why Use Context-Aware Services? 6-2
Is It Here?—Zone or Inventory Management 6-3
Where Is It?—Asset Tracking 6-3

What Is Its Condition?—Condition Tracking 6-3
What Is Its Status?—Status Monitoring 6-4
Where Is It in My Network?—Network Location Services 6-4
Cisco Context-Aware Components 6-5
Wired and Wireless Client Devices 6-5
Cisco Unified Network 6-6
Management and Applications 6-7
Context-Aware Component Interaction 6-8
Network Mobility Services Protocol (NMSP) 6-11
Context-Aware Services in Small Enterprise Designs 6-12
Component Capacities 6-15
Mobility Services Engine 6-15
WLAN Controllers 6-16
Wireless Control System (WCS) 6-16
Context-Aware Engine for Tags (AeroScout) 6-17
Integration within Small Enterprise Designs 6-18
MSE Connection to the Network 6-18
Clock Synchronization 6-18
SEDP Wireless Control System (WCS) Context-Aware Considerations 6-22
WLAN Controller and Cisco Catalyst Switch Definition and Synchronization 6-36
Wireless Client Context-Aware Considerations 6-37
RFID Asset Tag Context-Aware Considerations 6-41
Rogue Device Context-Aware Considerations 6-45
Context-Aware Considerations for Wired Device Tracking 6-48
Classification and Marking of NMSP Sessions 6-55
NMSP Traffic Flows Originating At The MSE 6-57
NMSP Traffic Generated By WLAN Controllers 6-58
NMSP Traffic Generated By Context-Aware Switches 6-59
Hardware/Software Releases 6-60
Context-Aware Services—General Best Practice References 6-61

CHAPTER 7 **Small Enterprise Design Profile(SEDP)—Digital Media and Video Surveillance Design** 7-1
Digital Media System Overview 7-1
Video Surveillance Overview 7-1
Cisco Digital Media System Architecture 7-2
DMS Solution for Small Enterprise 7-3
Desktop Video Application Overview 7-5
Desktop Video Components 7-5
Publishing Live and Video On-Demand Content 7-5

Enterprise TV Application Overview 7-6
Enterprise TV Components 7-7
Broadcasting Live TV or Video On-Demand Content 7-7
Digital Signage Application Overview 7-8
Digital Signage Components 7-9
Video Surveillance System Architecture 7-9
Cisco Video Surveillance Media Server 7-11
Cisco Video Surveillance Operations Manager 7-11
Cisco Video Surveillance IP Cameras 7-12
Deploying Digital Signage in Small Enterprise Network 7-14
Centralized Management Model 7-14
Distributed Content Storage Model 7-16
Validated Content Distribution Model 7-17
Implementing Network Services for Digital Signage 7-18
Deploying Cisco DMP in the Access Layer Network 7-19
Manual Deployment 7-19
Applying Ingress QoS Policy on Access Layer 7-21
Applying Egress QoS Policy on Access Layer 7-23
Auto Smartport Macro Deployment 7-23
Implementing Auto SmartPort Macro 7-25
Tuning Auto SmartPort Macro 7-25
Implementing Cisco Digital Media Player 7-26
Assigning IP Address to Cisco DMP 7-26
Registering Cisco DMP to Cisco DMM Database 7-27
Summary 7-36

CHAPTER 1

Small Enterprise Design Profile (SEDP) Overview

The Enterprise Design Profile delivers the foundational network design that all enterprise services, applications, and solutions use to interact and communicate with one another. The Enterprise Design Profile is constructed in a fashion that supports all the applications and services that will ride on it. Additionally, these profiles must be aware of the type of traffic traversing and treat each application or service with the correct priority based on the needs and importance of that application.

The Small Enterprise Design Profile is made up of the following four distinct components:

- Network Foundation guidance
- Cisco SAFE Security Architecture guidance
- Mobility guidance
- Collaboration Services guidance such as Unified Communications

Each of these critical foundation components have been carefully designed and tuned to allow for a secure environment that provides for business continuity, service awareness and differentiation, as well as access flexibility. See Figure 1-1.

Figure 1-1 Small Enterprise Design Profile Design Components

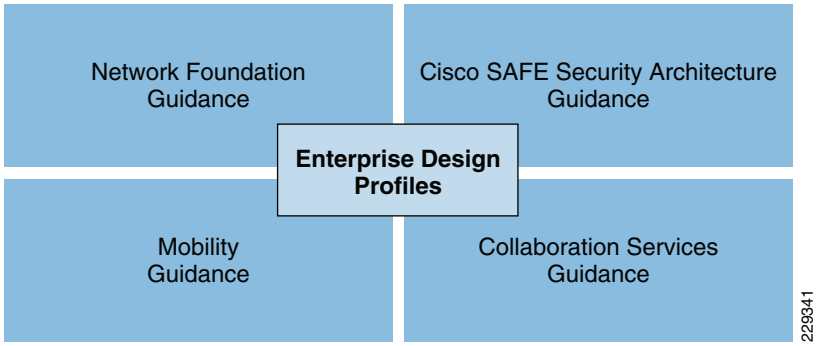

Small Enterprise Design Profile

The design used for the Small Enterprise Design Profile is intended to represent as many small-size Enterprise network environments as possible. To accomplish this, a modular design is used representing sites of varying sizes (see Figure 1-2). The Small Enterprise Design Profile is built upon a network foundation consisting of a main site, where the majority of the critical applications reside.

Connected through a Metro Ethernet WAN are remote sites of varying sizes. The remote small site is designed to support up to 100 employees. The remote large site is designed to support up to 500 employees. Each site can coexist in a small Enterprise network or can be treated as separate modules. Design guidance for remote sites of varying sizes provides flexibility, modularity, and scalability as the Enterprise grows. Additionally, it is expected that half of all network can be accessed wired and wirelessly.

Figure 1-2 Small Enterprise Design Profile Design

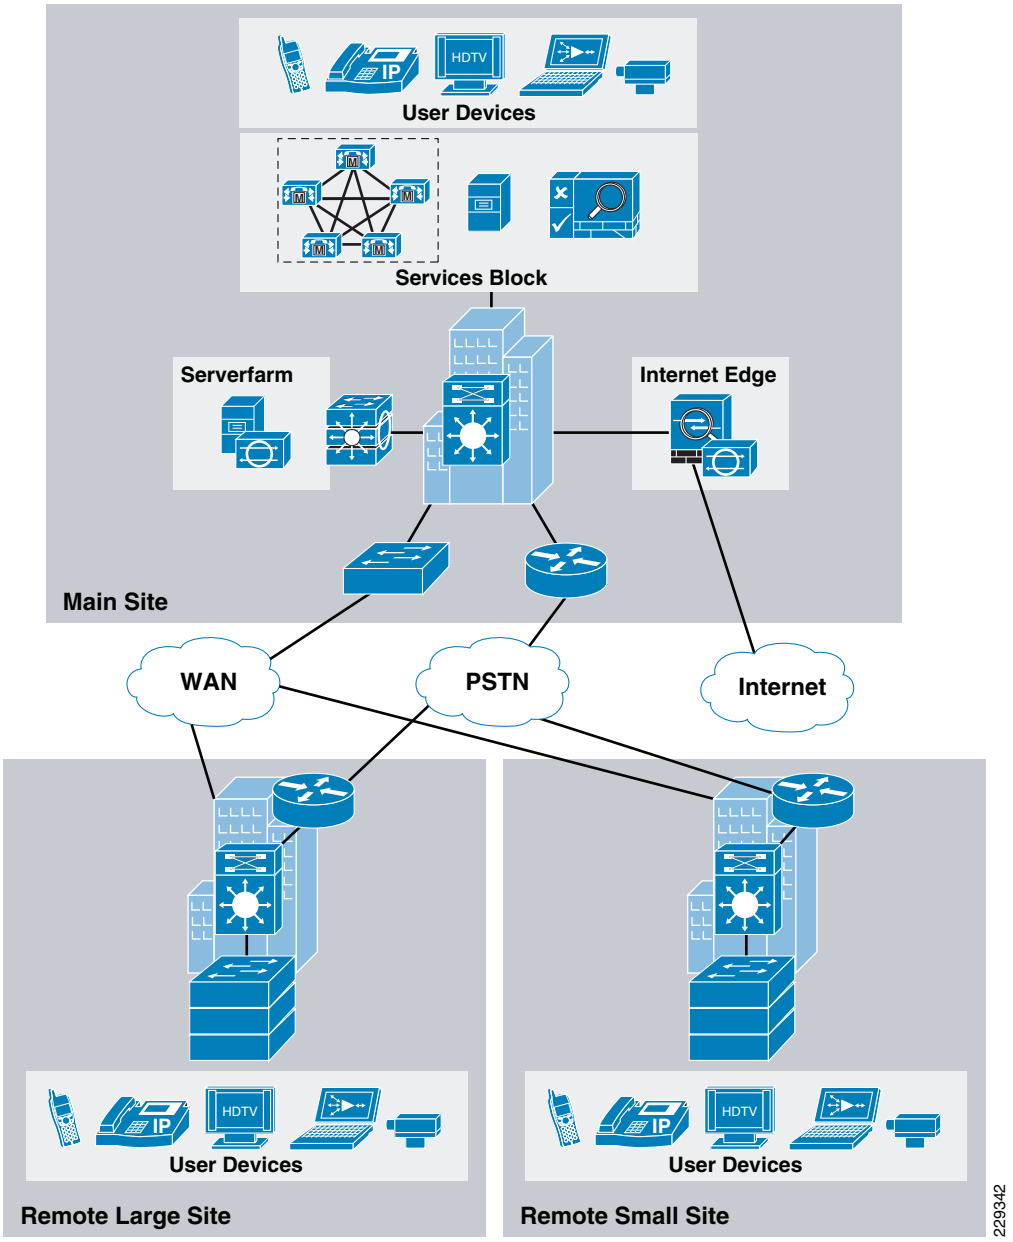

Main Site Design

The main site design represents a centralized services site that interconnects the remote sites regardless of size. The main site consists of a main building core connected to a serverfarm and service block design. The remote large site and the remote small site connect to the main site via a 100Mbps Metro Ethernet link. Because many Enterprise services and applications are centrally located within the main site rather than each remote site, high network availability must be maintained. The main site is connected to outside entities such as the partners, vendors, customers, and the Internet using the Internet edge components. All the remote sites within the small enterprise system also connect to the main site and use the main site to connect outside the Enterprise.

Remote Large Site Design

The remote large site design has been developed for Enterprise sites that have up to 500 employees. The remote large site is connected to the main site via a 100Mbps Metro Ethernet link. While this site connects to the main site where most Enterprise services and applications are centralized, this site also uses resilient application services features to maintain critical services such as Digital Media Signage, Video Surveillance, and SRST in case of a WAN failure. For resiliency, this site is designed with a dual switch and dual aggregation links.

Remote Small Site Design

The remote small site profile represents a site supporting up to 100 employees. The core and distribution layers in the remote site network are collapsed into one and use Cisco's Catalyst Stackwise switching technology for redundancy. The remote small site is connected to the main site via a 10Mbps Metro Ethernet link and like the remote large site, resilient application service features are used to maintain critical services in case of a WAN failure.

Access Devices

The devices that connect to the Small Enterprise Design Profile network include phones, cameras, displays, laptops, desktops, mobile phones, and personal devices (iPod, MP3, etc). Half of all the devices are expected to connect to the network using 802.11 ABGN wireless access.

The Small Enterprise Design Profile consists of four major components. The sections below provide a brief description of each of these components.

Network Foundation Design Considerations

The Small Enterprise Design Profile Network Foundation guidance is made up of routers and switches deployed in a three-tier hierarchical model that uses Cisco IOS to provide foundational network technologies needed to provide a highly available, application-aware network with flexible access. LAN and WAN guidance is provided under this component.

LAN Design Considerations

Hierarchical network design model components:

- *Core layer*—The site backbone consisting of a Layer-3 core network interconnecting to several distributed networks and the shared services block to access local and global information.
- *Distribution layer*—The distribution layer uses a combination of Layer-2 and Layer-3 switching to provide for the appropriate balance of policy and access controls, availability, and flexibility in subnet allocation and VLAN usage.
- *Access layer*—The Demarcation point between network infrastructure and access devices. Designed for critical network edge functionality to provide intelligent application and device aware services.

High Availability Design Considerations

To ensure business continuity and prevent catastrophic network failure during unplanned network outage, it is important to identify network fault domains and define rapid recovery plans to minimize the application impact during minor and major network outages.

The Small Enterprise Design Profile design ensures network survivability by employing three major resiliency methods that serve to mitigate most types of failures. The appropriate resiliency option should be selected given the network system tier, role, and network service type:

- *Link resiliency*—Provides redundancy during physical link failures (i.e., fiber cut, bad transceivers, incorrect cablings, etc.)
- *Device resiliency*—Protects network during abnormal node failure triggered by hardware or software (i.e., software crashes, non-responsive supervisor etc.)
- *Operational resiliency*—Enables higher level resiliency capabilities, providing complete network availability even during planned network outage conditions.

Routing Protocol Selection Criteria

Routing protocols are essential for any network, because they allow for the routing of information between buildings and sites. Selecting the right routing protocol can vary based on the end-to-end network infrastructure. The routers and switches support many different routing protocols that will work for Small enterprise network environments. Network architects must consider all the following critical design factors when selecting the right routing protocol to be implemented throughout the internal network:

- *Network design*—Proven protocol that can scale in full-mesh site network designs and can optimally function in hub-and-spoke WAN network topologies.
- *Scalability*—Routing protocol function must be network and system efficient that operates with a minimal number of updates, recomputation independent of number of routes in the network.
- *Rapid convergence*—Link state versus DUAL recomputation and synchronization. Network reconvergence also varies based on network design, configuration, and a multitude of other factors which are beyond the routing protocol.
- *Operational considerations*—Simplified network and routing protocol design that can ease the complexities of configuration, management, and troubleshooting.

Access Layer Design Considerations

The access layer represents the entry into the network, consisting of wired and wireless access from the client to the network. The switch that the client connects to will ultimately connect to the network distribution layer of and the method communication here must be considered in any design. Traditional Layer 2 connectivity is prevalent in most networks today; however, it comes at some cost in administration, configuration, and timely resiliency. The emerging method of connectivity is a Layer 3 connection, commonly referred to as *routed-access*.

Performing the routing function in the access-layer simplifies configuration, optimizes distribution performances, and allows for the use of well known end-to-end troubleshooting tools. Implementing a Layer 3 access-layer in lieu of the traditional Layer 2 access replaces the required Layer 2 trunks with a single point-to-point Layer 3 link. Pushing Layer 3 routing functionality one tier down on Layer 3 access switches changes traditional multilayer network topology and the forwarding path. Implementing a routed access layer does not require any physical or logical link reconfiguration or changes.

See Figure 1-3.

Figure 1-3 Control Function in Multi-Layer and Routed-Access Network Design

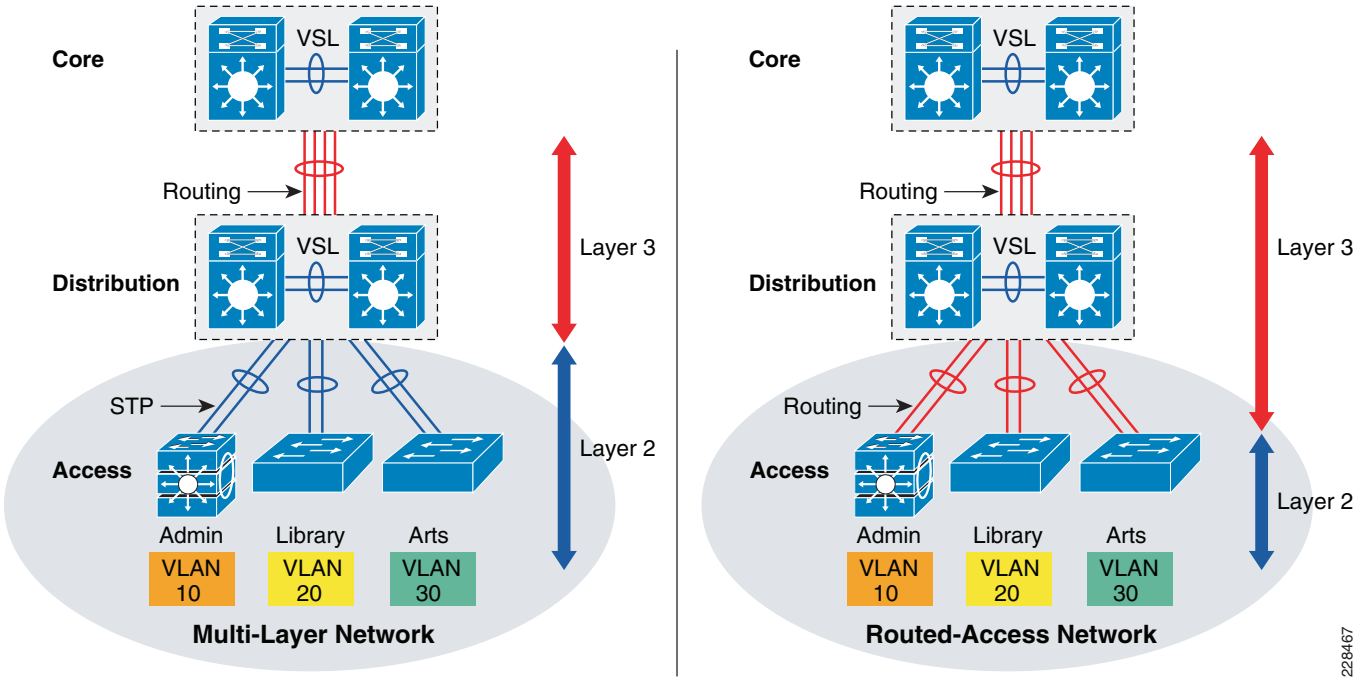

At the network edge, Layer 3 access switches provides an IP gateway function and serve as a Layer-2 demarcation point to locally connected endpoints that can be logically segmented into multiple VLANs.

LAN Foundational Services

The Small Enterprise Design Profile uses essential foundational services to efficiently disseminate information that is used by multiple clients, as well as identify and prioritize different applications traffic based on their requirements. Designing the foundational services in a manner consistent with the needs of the Small enterprise network system is paramount. Some of the key foundational services discussed include the following:

- Multicast routing protocol design considerations
- Designing QoS in the site network

WAN Design Considerations

WAN Transport

In order for sites to communicate with one another and/or to communicate outside the Enterprise network system, network traffic must traverse over a WAN. WAN transport differs greatly from LAN transport due to variables such as the type of connection used, the speed of the connection, and the distance of the connection. The service fabric design model covers the following WAN transport design considerations:

- Internet
- Metro Ethernet

WAN Foundational Services

Similar to the LAN, the WAN must employ essential foundational services to ensure the proper transport and prioritization of community college services. WAN foundation services considered are as follows:

- Routing protocol design
- Quality-of-service (QoS)
- WAN resiliency
- Multicast

Cisco SAFE Security Architecture Design Considerations

Security of Small Enterprise Design Profile is essential. Without it, Enterprise solutions, applications, and services are open to be compromised, manipulated, or shut down. The Small Enterprise Design Profile was designed with built-in security by leveraging the proven design and deployment guidelines of the Cisco SAFE security architecture. The following are the primary security design considerations:

- *Network Foundation Protection (NFP)*—Ensuring the availability and integrity of the network infrastructure, protecting the control and management planes.
- *Internet perimeter protection*— Ensuring safe connectivity to the Internet, and protecting internal resources and users from malware, viruses, and other malicious software. Protecting users from harmful content. Enforcing E-mail and web browsing policies.
- *Serverfarm protection*—Ensuring the availability and integrity of centralized applications and systems. Protecting the confidentiality and privacy of user information and records.
- *Network access security and control*—Securing the access edges. Enforcing authentication and role-based access for employees and users residing at the main and remote sites. Ensuring systems are up-to-date and in compliance with the network security policies.
- *Network endpoint protection*—Protecting servers and Enterprise-controlled from viruses, malware, botnets, and other malicious software. Enforcing E-mail and web browsing policies for users.

Mobility

Mobility is an essential part of the Small Enterprise Design Profile. Most users will connect wirelessly to site networks and other devices will also rely on the mobile network. In designing the mobility portion of the service fabric, the following design criteria were used:

- *Accessibility*—Enables employees and guests to be accessible and productive, regardless of where they meet. This design element provides for easy, secure guest access to guests such as contractors, vendors and other visitors.
- *Usability*—In addition to extremely high WLAN transmission speeds made possible by the current generation of IEEE 802.11n technology, latency sensitive applications (such as IP telephony and video-conferencing) are supported over the WLAN using appropriately applied QoS. This gives preferential treatment to real-time traffic, helping to ensure that video and audio information arrives on time.
- *Security*—Segment authorized users and block unauthorized users. Extend the services of the network safely to authorized parties. Enforce security policy compliance on all devices seeking to access network computing resources. Employees enjoy rapid and reliable authentication through IEEE 802.1x and Extensible Authentication Protocol (EAP), with all information sent and received on the WLAN being encrypted.
- *Manageability*—Network administrators must be able to easily deploy, operate, and manage hundreds of access points within multiple Enterprise network site deployments. A single, easy to understand WLAN management framework is desired to provide small and large Enterprise systems with the same level of wireless LAN management scalability, reliability and ease of deployment that is demanded by traditional enterprise business customers.
- *Reliability*—Provide adequate capability to recover from a single-layer fault of a WLAN accessibility component or controller wired link. Ensure that wireless LAN accessibility is maintained for employees and guest visitors in the event of common failures.

Collaboration Services Design Considerations

Adoption of IP technology has led to a fundamental change in designing networks. No longer are networks used solely to provide data communication between computers and servers. IP technology has extended beyond the data network and is now used extensively for Unified Communications and Video communication as well. Unified Communications, IP Video Surveillance and Digital Media systems were validated in the Small Enterprise Design Profile.

Unified Communications Design Considerations

Call Processing Considerations

How calls are processed in the Small Enterprise Design Profile environment is an important design consideration. Guidance in designing scalable and resilient call processing systems is essential for the successful deployment of a unified communications system. Considerations include the following:

- *Scale*—The number of users, locations, gateways, applications, and so forth
- *Performance*—The call rate

- *Resilience*—The amount of redundancy

Gateway Design Considerations

Gateways provide a number of methods for connecting an IP telephony network to the Public Switched Telephone Network (PSTN). Design considerations for gateways include the following:

- PSTN trunk sizing
- Traffic patterns
- Interoperability with the call processing system

Dial Plan Considerations

Dial plan is one of the key elements of a unified communications system, and an integral part of all call processing agents. Generally, the dial plan is responsible for instructing the call processing agent on how to route calls. Specifically, the dial plan performs the following main functions:

- Endpoint addressing
- Path selection
- Calling privileges
- Digit manipulation
- Call coverage

Survivability Considerations

Voice communications are a critical service that must be maintained in the event of a network outage and for this reason the service fabric must take survivability into consideration. Guidance on how the Small Enterprise Design Profile is equipped and designed to keep voice communications active in the event of an outage is provided.

IP Video Surveillance Design Considerations

Video surveillance systems have proven their value in a wide range of applications. Video documentation of critical incidents enhances employee safety and better protects valuable assets. However, traditional analog Closed-circuit TeleVision (CCTV) surveillance systems have many limitations—they are unable to store recorded video in local and remote locations or provide video access to mobile or remote users.

Network-centric video surveillance components include the following:

- *Cisco Video Surveillance Manager*—Enables IT administrators and security personnel to view, manage and record video locally and remotely using the IP network and a standard Internet browser. Video can be securely accessed anywhere, at any time, enabling faster response, investigation and resolution of incidents. Video can be recorded and stored locally off and at the main site allowing it to be managed and aggregated with video from multiple locations.
- *Cisco Video Surveillance Media Server*—A highly scalable and reliable video management platform that manages, replicates, distributes and archives video systems.

- *Cisco Video Surveillance Operations Manager*—A web-based user interface that authenticates and manages access to video feeds. It is a centralized administration tool for the management of Media Server hosts, Virtual Matrix hosts, cameras, encoders, and viewers.
- *Cisco Video Surveillance Media Virtual Matrix*—Monitors video feeds in command center and other 24-hour monitoring environments. It allows operators to control the video being displayed on multiple local and remote digital monitors.

Digital Media Systems

- *The Cisco Digital Media Suite*—A comprehensive portfolio of digital signage, desktop video, and enterprise TV components and applications that can be centrally managed. The Cisco Digital Media Suite is comprised of three distinct subsystems.
- *Cisco Digital Signs*—Provides scalable centralized management and publishing of compelling digital media to networked, on-premise digital signage displays. It enables the disseminations news and emergency information to large screens connected to the Enterprise existing network. The same content to all signs in the network can be delivered.
- *Cisco Cast*—Uses the same hardware as Cisco Digital Signs, but has different usage models. With Cast, all control switches to the end user via a remote control—and with the latest release, IP phones, smartphones, and touch screens—can be used to control what content comes to the screen. Cast has three user interfaces, one to access VoDs, one for scrolling through live channels, and a channel guide,
- *Cisco Show and Share*—Enables employees to create, capture, and receive live and pre-recorded video on their desktop computers.Digital media can be browsed, searched, and viewed over the network through a unique, easy-to-use Cisco video portal experience—anywhere, anytime.

The components of the Cisco Digital Media Suite include the following:

- *Cisco Digital Media Manager (DMM)*—The central management application for all Cisco Digital Media Suite products. It is used to manage, schedule, and publish compelling digital media for Cisco Digital Signs, Cisco Cast and Cisco Show and Share. As an integrated part of the Cisco Digital Media Suite, this web-based media management application enables content owners to easily upload, catalogue, edit, package, and publish digital media content for live or on-demand playback.
- *Scientific Atlanta Encoder*—Encodes live video input into a MPEG-2 or MPEG-4 multicast stream for Cisco Digital Signs and Cisco Cast.
- *Digital Media Encoders (DME) for Cisco Show and Share*—Register with the DMM and broadcast live video to the Streaming Server using RTSP. Recorded video may be archived to the Content Repository.
- *Streaming Server for Cisco Show and Share*—Provides stream splitting capabilities, allowing many clients to view a single live stream from a DME or pre-recorded source (live rebroadcast).
- *Show and Share User Interface*—Provides the web-based interface for clients. All navigation and authentication is completed through this server.
- *Web Server/Content Repository*—Holds all VoDs referenced by the Show and Share server. All VoD streaming requests to the Show and Share server are redirected to this server. The web server/content repository is also the component that holds all VoDs referenced by the DMM for Cisco Digital Signs and Cisco Cast. For Digital Signs and Cast, all VoD streaming requests issued to the DMP are serviced from this server.
- *Digital Media Player (DMP) for Cisco Digital Signs and Cisco Cast*—Decodes and displays unicast (prerecorded content streamed over http or RTSP) and multicast streamed video as well as flash content (live content from the Scientific Atlanta encoder or other multicast encoder).

- *Cisco LCD Professional Series Displays*—An integral part of the Digital Media System (DMS) suite of products and are used to display information, Cisco LCD displays are available in different sizes and models and offer full 1080p resolution.

Summary

The Small Enterprise Design Profile delivers validated network design and deployment best practices to evolve small enterprise networks into Borderless networks. Enterprise Design Profile applies the best practices from a collection of Cisco Validated Designs (CVD) and integrates these best practices into a customer-based design profile.

For the existing network, the Small Enterprise Design Profile provides guidance for the evolution of the Enterprise network to a Borderless Network. For new networks, the Small Enterprise Design profile steps you from planning your Enterprise network to using technology to enable and solve your business needs.

CHAPTER 2

Small Enterprise Design Profile(SEDP)—Network Foundation Design

This chapter describes the Small Enterprise Design Profile network design, which is a well designed and validated network architecture that is flexible, and cost effective to support a wide range of network foundational services. Key features of this network design include the following:

- High availability
- Single fabric, multi services
- Differentiated services
- Layer 2 and Layer 3 access

This chapter provides design guidance to build a highly resilient, manageable, and cost-effective small enterprise network that provides a solid foundation for seamless integration and operation of applications and network services. The network has been specifically designed to meet the challenges of the small enterprise environment.

Building Unified Small Enterprise Network Infrastructure

Cisco has years of experience developing high performance, highly available, multi service networks. The key to developing a robust design is applying a proven methodology. The following design principles were applied to develop the Small Enterprise Network Design architecture:

- Hierarchy
  - Clarifies the role of each device in each tier
  - Simpler to deploy, operate, and manage the network
  - Reduces fault domains at every tier
- Modularity
  - Enables growing the network on demand basis
- Resiliency
  - Meet users expectation of network always being available.
- Flexibility
  - Allows intelligent traffic load-sharing by using all network resources

The Unified Campus network is designed to be highly available, and cost effective, while delivering capabilities necessary to enable advanced services, such as IP telephony, video, security, wireless LANs. The network design includes the following key features;

- Hierarchical design with collapsed Core
- Quality-of-service (QoS) to ensure real-time data (telephony, video) are given higher priority
- Application of resilient design principles
- Multi cast
- Routed access
- Redundancy

Hierarchical Network Design

The three-tier hierarchical model (see Figure 2-1) is the approach typically employed to achieve a high performance, highly available, scalable network design. This design employs the four key design principles of hierarchy, modularity, resiliency and flexibility.

Figure 2-1 Three-Tier Hierarchical Model

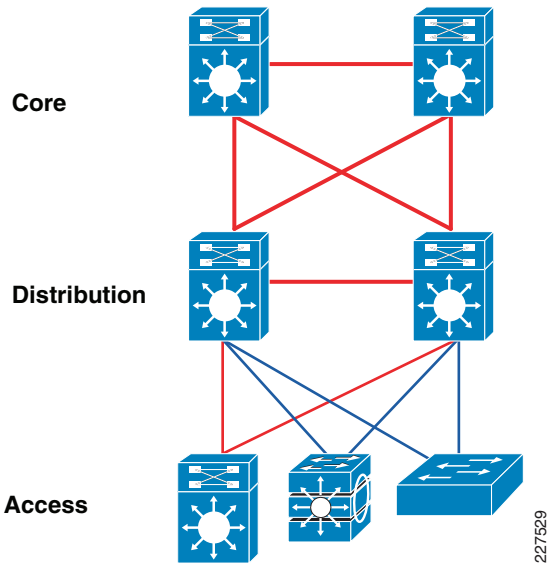

Each layer in the three-tier hierarchical model has a unique role to perform:

- *Access Layer*—The primary function of an access-layer is to provide network access to the end user. This layer often performs OSI Layer-2 bridge function that interconnects logical Layer-2 broadcast domains and provides isolation to groups of users, applications, and other endpoints. The access-layer interconnects to the distribution layer.
- *Distribution Layer*—Multi-purpose system that interfaces between access layer and core layer. Some of the key function for a distribution layer include the following:
  - Aggregate and terminate Layer-2 broadcast domains
  - Provide intelligent switching, routing, and network access policy function to access the rest of the network.

  - Redundant distribution layer switches provides high availability to the end-user and equal-cost paths to the core. It can provide differentiated services to various class-of-service applications at the edge of network.

- *Core Layer*—The core-layer provides high-speed, scalable, reliable and low-latency connectivity. The core layer aggregates several distribution switches that may be in different buildings. Backbone core routers are a central hub-point that provides transit function to access the internal and external network.

Table 2-1 lists the key functions of each layer.

Table 2-1 Key Functions of Hierarchical Network Layer Devices

Key Function	Access	Distribution	Core
Network Transit	Rest of the network.	Internal and External network	
Intelligent Services	PoE, IEEE 802.1AD, Mobility, AutoQoS, Auto-SmartPort Macro(ASP)	Route optimization Network and System Virtualization Layer-2 Interconnect	
Forwarding Decision	Layer 2/Layer 3		Layer 3
Security Services	CISF, 802.1x, NAC, ACL etc.	CISF, ACL, Route Filter, CoPP etc.	ACL, Route Filter, CoPP etc.
QoS Services	Classification, Marking, Policer and Queueing	Classification, Marking, and Queueing	

To learn more about typical network designs, refer to the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html

Figure 2-2 illustrates a sample network diagram for a multi-building small enterprise network design.

Figure 2-2 Multi Building Small Enterprise Network Design

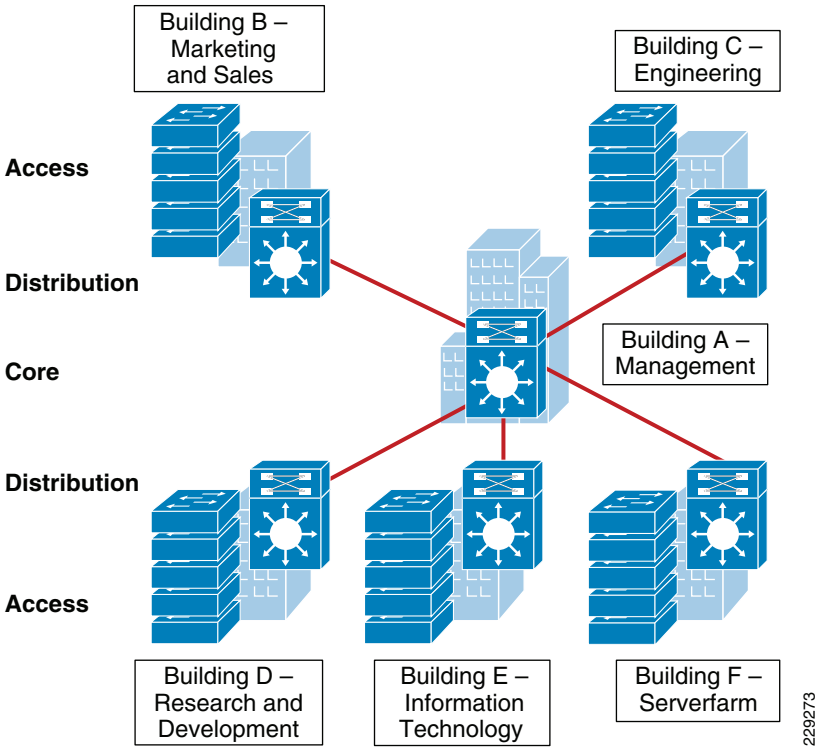

Collapsed Core Network Design

The three-tier hierarchical design maximizes performance, network availability, and the ability to scale the network design. Most small enterprise campus' do not grow significantly larger over time, and most small enterprise campus are small enough to be well served by a two-tier hierarchical design, where the core and distribution layers are collapsed into one layer. The primary motivation for the collapsed core design is reducing network cost, while maintaining most of the benefits of the three-tier hierarchical model.

Deploying a collapsed core network results in the distribution layer and core layer functions being implemented in a single device. The collapsed core/distribution device must provide the following:

- High speed physical and logical paths connecting to the network
- Layer-2 aggregation and demarcation point
- Define routing and network access policies
- Intelligent network services—QoS, Network virtualization, etc.

Note If the main site or a remote site campus has multiple buildings, and is expected to grow over time, then implementing the three-tier hierarchical model is a better choice.

Figure 2-3 illustrates a sample network diagram for a single main site building.

Figure 2-3 Main Site—Collapsed Core Network Design

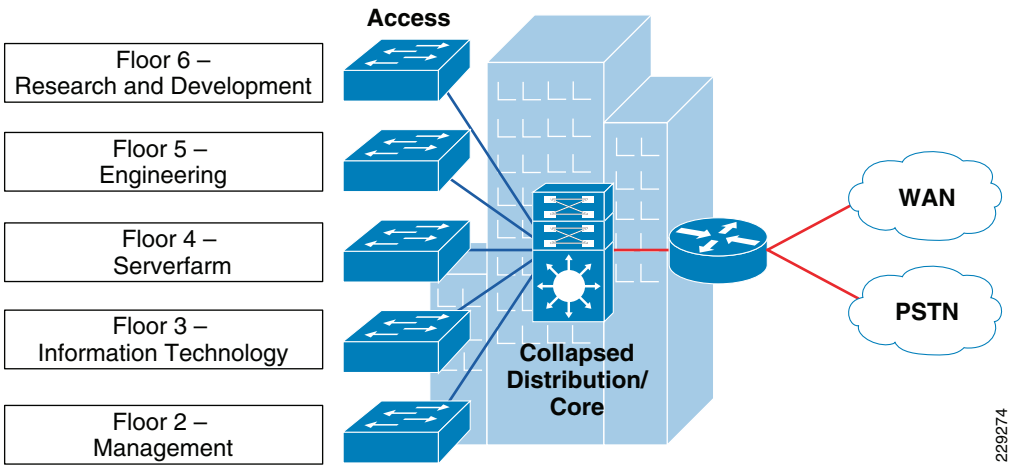

Main Site Network Design

If the main site has multiple buildings and it is expected to grow significantly over time, then implementing the three-tier hierarchical model is a good choice. For small size main sites that are unlikely to grow significantly, the collapsed core model is more cost effective. The Small Enterprise Design Profile uses the collapsed core network design in the main site.

The collapsed core network (see Figure 2-4) may be deployed with redundant core/distribution router, or consolidated core/distribution router.

Figure 2-4 Small Enterprise Design —Collapsed Core Network Models

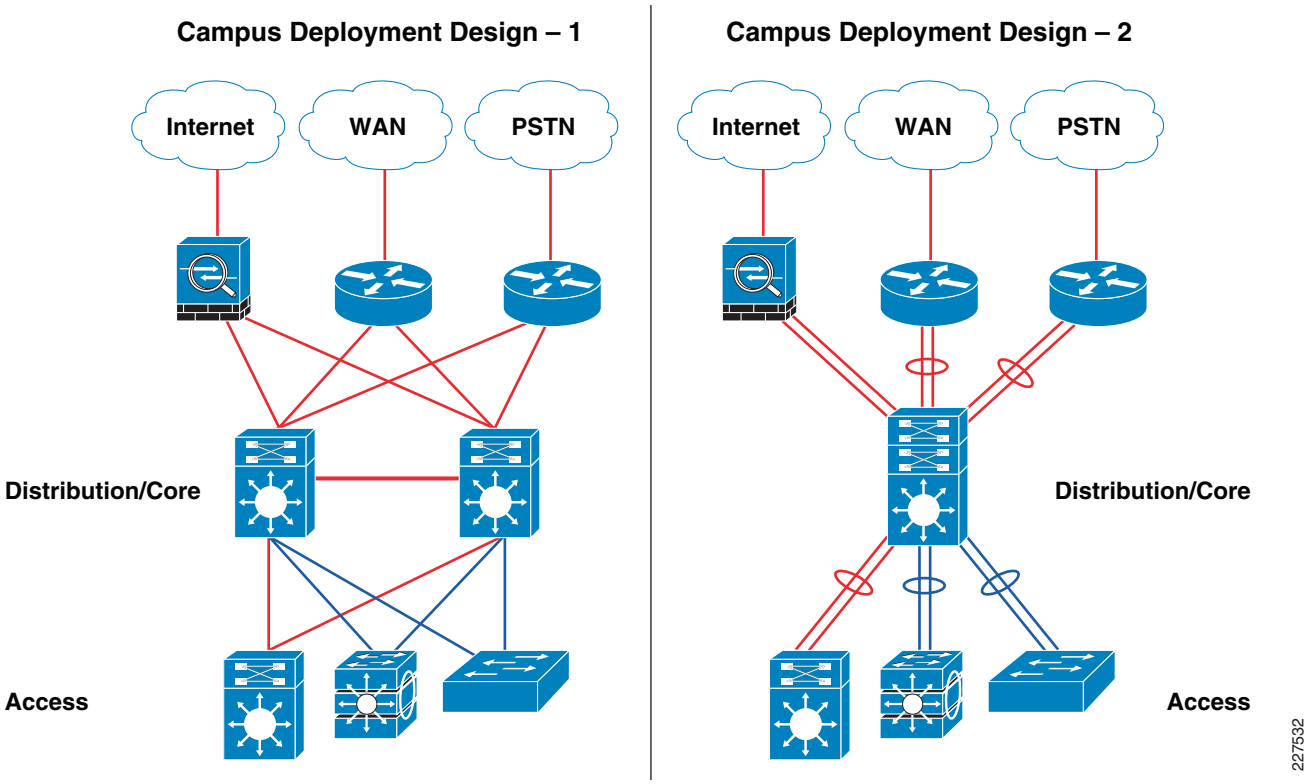

The redundant design is more complex, because all of the core/distribution functions must be implemented on two routers in a complimentary fashion. To learn more about the redundant designs, refer to *High Availability Campus Recovery Analysis Design Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/HA_recovery_DG/campusRecovery.html

The main site is designed with a consolidated core/distribution router to maximize performance, while keeping costs affordable (design 2). The consolidated collapsed core model has the following benefits:

- Simplifies network protocols (eases network operations)
- Enables symmetric forwarding paths
- Delivers deterministic network recovery performance

With this design, the default behavior of Layer-2 and Layer-3 network control protocols is to create a redundant view between two systems. The core router builds a ECMP routing topology which results in symmetric forwarding paths beyond the main site.

Default Layer-2 configuration eliminates the need for FHRP, automatically eliminating the asymmetric forwarding behavior which causes unicast flooding in the network. This simplifies the network operation, since there is no need to configure or tune FHRP protocols.

The disadvantage of this Layer-2 network design is that the network is under-utilized. This is due to the way Layer-2 protocols are designed to build loop-free network topologies. When two Layer-2 bridges are directly connected, the STP protocol will block low-priority STP physical port in the forwarding table. Figure 2-5 illustrates the control-plane, and the forwarding-plane for this design.

Figure 2-5 Design Model 2 – Developing Control and Forwarding Paths

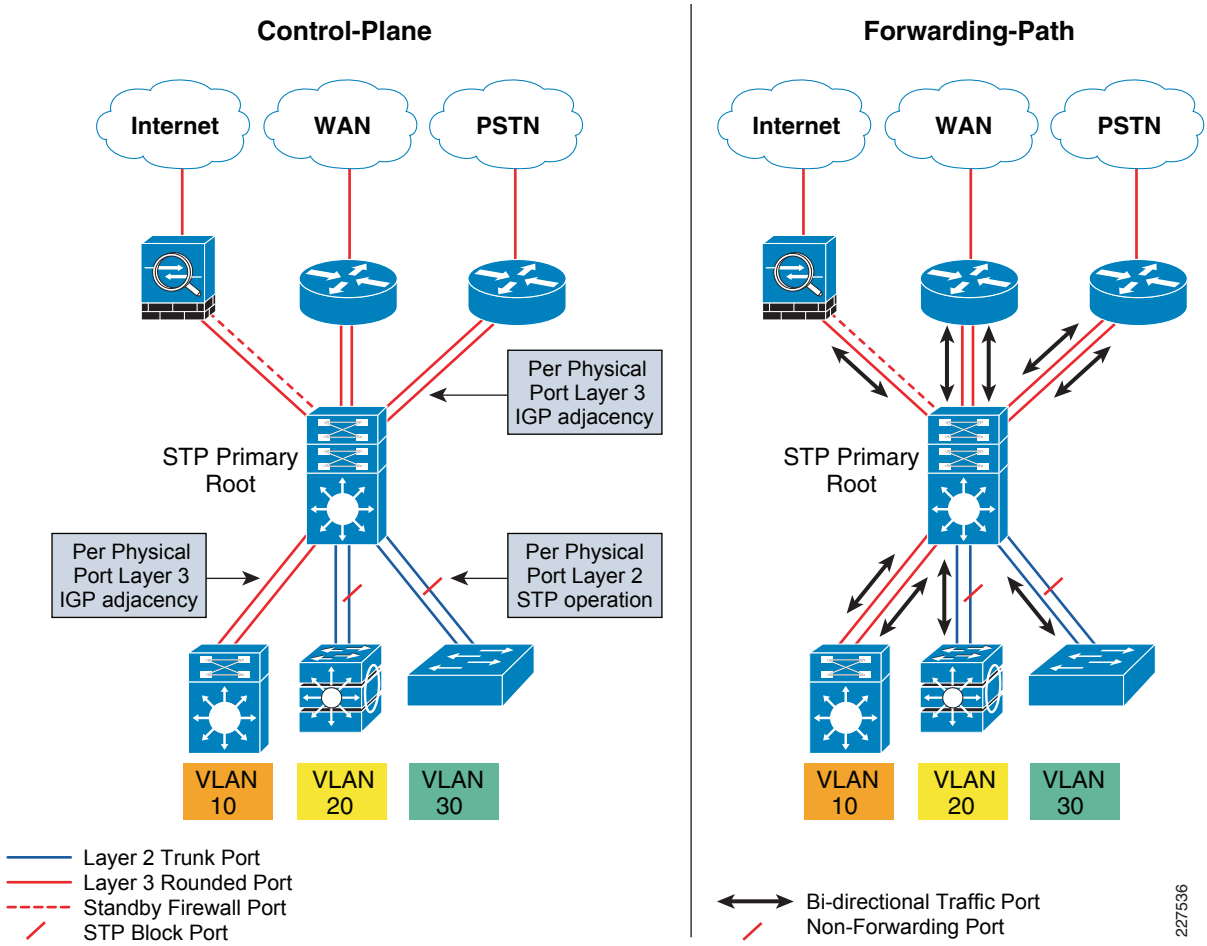

This design suffers from two challenges:

- Multiple-routing adjacencies between two Layer-3 systems. This configuration doubles the control-plane load between each of the Layer-3 devices. It also uses more system resources like CPU and memory to store redundant dynamic-routing information with different Layer-3 next-hop addresses connected to same router.
- As depicted in Figure 2-5, STP protocol blocks one of the physical ports in the Layer-2 network. Since this design employs point-to-point links between the collapsed core and peer devices, the solution is to tune the network to enable a single control plane, to improve forwarding efficiency and resource utilization. The recommendation is to aggregate all physical ports into a single logical channel-group. This logical aggregated Ethernet bundle interface is known as *EtherChannel*.

EtherChannel Fundamentals

EtherChannel provides inverse-multiplexing of multiple ports into a single logical port to a single neighbor. This technique increases bandwidth, link efficiency, and resiliency. EtherChannel technology operates on the MAC layer. Upper layer protocols require a single instance to operate over the logical interface. EtherChannel provides efficient network operation and graceful recovery to higher layer protocols during bundle port failure and restoration.

The control-plane depicted in Figure 2-5 builds redundant Layer- 2 or Layer-3 network information over each physical links. Each device builds common network prefix entries with a different next-hop path pointing to same next hop device. Implementing EtherChannel results in a network topology with a single destination entry for single next-hops, via the egress logical EtherChannel port. EtherChannel reduces storing redundant network entries in the database and forwarding tables, which automatically improves network convergence times and system resources utilization.

EtherChannel helps improve the overall network stability and availability. Failure of individual physical link will cause network topology recomputation, restoration, and may be rerouted. Such process requires CPU interruption that could impact the overall application performance. EtherChannel significantly simplifies the network response to a individual link failure. If an individual link in EtherChannel fails, the interface will not trigger any network topology changes. All underlying hardware changes remain transparent to higher-layer protocols, thus minimizing impact to network and application performance, and improving network convergence. Figure 2-6 illustrates how enabling EtherChannel in Layer-2 and Layer-3 network simplifies control-plane and forwarding-plane.

Figure 2-6 Design Model 2 – Optimized Control and Forwarding Paths with EtherChannel

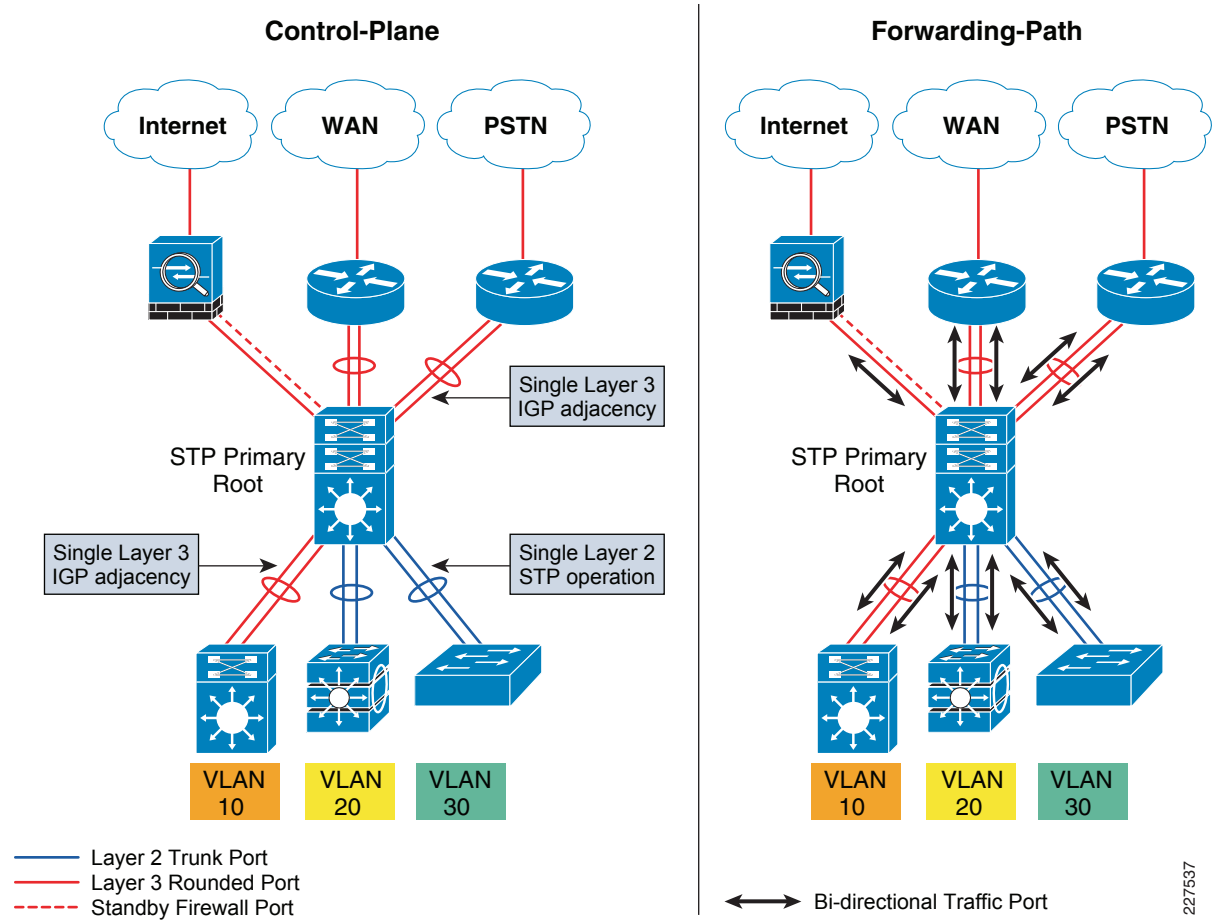

Resilient Distributed System

The Small Enterprise Design Profile uses the Cisco Catalyst 4500 with next-generation Supervisor-6E in the consolidated core/distribution layer. It is chosen for its price performance, and the high availability features within the device. The Cisco Catalyst 4500 switch supports redundant supervisor engines and provides Stateful Switchover (SSO) and Non-Stop Forwarding (NSF) capabilities. SSO ensures the Layer-2 and Layer-3 protocol state-machines and network forwarding entries on the standby supervisor engine are maintained, and can quickly assume control-plane responsibilities and gracefully restore the control-plane in the event of a primary supervisor failure. While the control-plane is gracefully recovering, the NSF function continues to switch traffic in hardware.

The Cisco Catalyst 6500 platform is an enterprise-class system providing integrated network services for large scale and high-speed networks. For large, multi building sites, or in situations where future scalability is important, the Catalyst 6500 is a better choice for core/distribution layer switch. The design principles remain the same when deploying a Catalyst 6500.

Main Site Access-Layer Edge Services

The access layer is the first tier or edge of the network. It is the layer where end-devices (PCs, printers, cameras, etc.) attach to the small enterprise network. It is also the layer where devices that extend the network out one more level are attached; IP phones and wireless access points (APs) are examples of devices that extend the connectivity out from the access switch. The wide variety of devices that can connect and the various services and dynamic configuration mechanisms required, make the access layer the most feature-rich layer of the small enterprise network. Figure 2-7 illustrates a main site network deployment with various types of trusted and untrusted endpoints.

Figure 2-7 Access-Layer Trust Boundary and Network Control Services

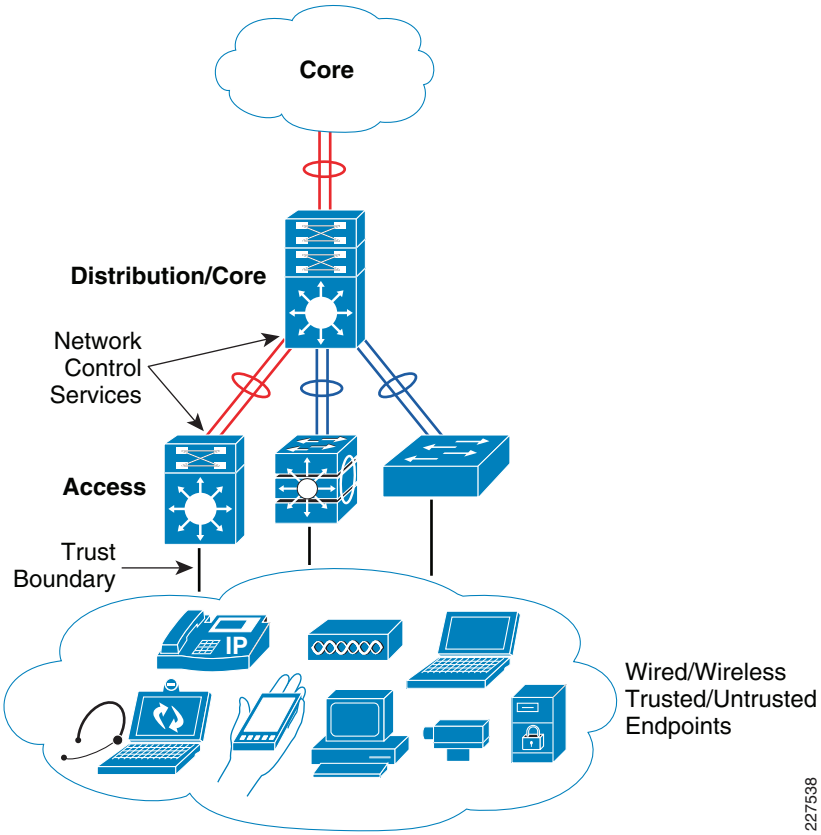

Table 2-2 examples of the types of services and capabilities that need to be defined and supported in the access layer of the network.

Table 2-2 Access-layer Services and Capabilities

Service Requirements	Service Features
Discovery and Configuration Services	802.1AF, CDP, LLDP, LLDP-MED
Integrated Security Services	IBNS (802.1X), CISF – Port-Security, DHCP Snooping, DAI and IPSG
Network Identity and Access	802.1X, MAB, Web-Auth
Application Recognition Services	QoS marking, policing, queueing, deep packet inspection NBAR
Intelligent Network Control Services	PVST+, Rapid PVST+, EIGRP, OSPF, DTP, PAgP/LACP, UDLD, FlexLink, Portfast, UplinkFast, BackboneFast, LoopGuard, BPDUGuard, Port Security, RootGuard
Energy Efficient Services	Power over Ethernet, EnergyWise, Energy efficient systems
Management Services	Auto-SmartPort Macro, Cisco Network Assistant

The access layer provides the intelligent demarcation between the network infrastructure and the computing devices that use the infrastructure. It provides network edge security, QoS, and policy trust boundary. It is the first point of negotiation between the network infrastructure and the end devices seeking access to the network.

A flexible network design, and the demand for mobility are two requirements which drive the access layer design. A flexible network design allows any legitimate device to be connected anywhere in the network (eg IP Phone, printer, video surveillance camera, digital signage, etc). Network users expect to be able to move around their devices (laptops, PDAs, printers, etc) and gain network access wherever necessary.

In order to allow devices to be moved within the network and ensure they associate with the correct network policies and services; the following access services are integrated into the small enterprise architecture:

- Ability to physically attach to the network and be associated with or negotiate the correct Layer-1 and Layer-2 network services—PoE, link speed and duplex, subnet (VLAN or SSID)
- Ability to provide device identification and, where needed, perform network access authentication
- Ability for the network to apply the desired QoS policies for the specific user, device or traffic flow (such as RTP streams)
- Ability for the network to apply the desired security policies for the specific user or device
- Ability for the network and device to determine and then register the location of the attaching device
- Ability for the device to negotiate and register the correct end station parameters (such as DHCP), as well as register for any other necessary network services (such as register for Unified Communications presence and call agent services)

The basic steps for deploying edge access switch features are as follows:

Step 1 Configure the baseline switching foundation

Step 2 Protect the network infrastructure

Step 3 Protect the end devices and their application data flows

Step 4 Apply the necessary network policies (QoS) to provide for the required service levels.

Step 5 Create the final template macro to allow for simplified configuration

Access-Layer Network Control Services

Properly designing the distribution block ensures the stability of the overall architecture. In the collapsed core model, the access-distribution block includes the access and distribution layers. Each of these layers has specific service and feature requirements. The network control plane choice (i.e., routing or spanning tree protocols) are central to determining how the distribution block fits within the overall architecture. The Small Enterprise Design Profile includes two designs for configuring the access-distribution block: multi-layer and routed-access. See Figure 2-8.

Figure 2-8 Access-Distribution Deployment Model

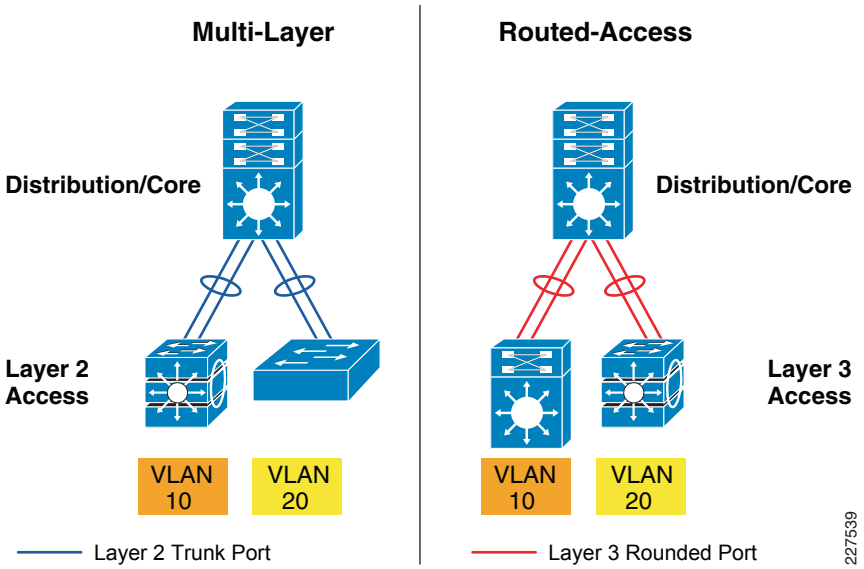

While both of these designs use the same basic physical topology and cabling plant, there are several key differences:

- Where the Layer-2 and Layer-3 boundaries exist
- How the network redundancy is implemented
- How load-balancing works

A complete configuration description of each access-distribution block model is provided in "Logical Multi-Layer Network" section on page 2-29 and the "Deploying Routed-Access Network" section on page 2-34 of this document.

Resilient Access-Layer Network and System

The access-layer provides endpoint connectivity to the rest of the network. Typical access switches like the Cisco Catalyst 2900 Series and Cisco Catalyst 3500 Series switches becomes single-point-of-failure (SPOF), if the hardware fails or if there is a software upgrade. Disrupting communication to mission critical endpoints (e.g., physical security camera) may be unacceptable.

The Small Enterprise Design Profile is designed with 2 to 4 uplink ports for each access switch, providing link-failure protection. For mission critical endpoints, this design employs the Cisco StackWise Plus and FlexStack solution in the access. It is designed to physically stack and interconnect multiple Layer-2 or Layer-3 switches using special cables. Stacking multiple switches into a logical ring creates a single unified and resilient access-layer network topology (see Figure 2-9). The next-generation Cisco Catalyst 2960-S FlexStack can be deployed in Layer-2 network domain and the Cisco Catalyst 3750-X StackWise Plus is deployed for routed access implementations.

Figure 2-9 Resilient, Scalable and Efficient Access-Layer Network Design

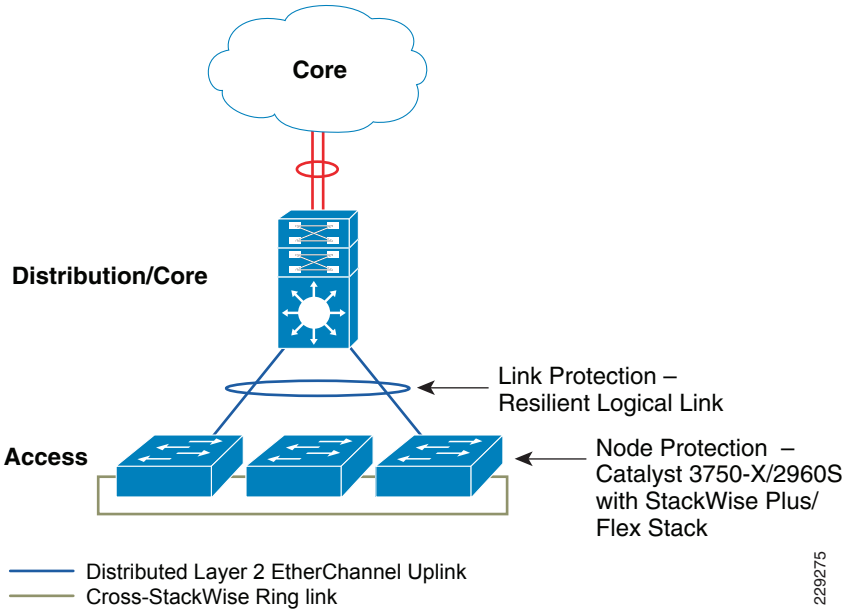

Main Site Data Center Network Design

The serverfarm is a central location which houses servers and storage. These resources must be available to users throughout the small enterprise network. The serverfarm may be collocated at the small enterprise network, or in a nearby site. Typically, small enterprise network are unable to afford high-speed redundant WAN links between the serverfarm and the remote sites. This makes the design vulnerable to service outage at the remote sites, in the event of WAN link failure. The Small Enterprise Design Profile recommends which services should be placed in the centralized serverfarm, and which services should be distributed (i.e., hosted at each remote site). The key criteria to consider when making this decision include:

- *Scalability*—The compute capacity and storage capacity of the centralized serverfarm must be sufficient to handle peak loads.
- *Network Load*—The overall network design, from serverfarm to remote site must have enough capacity to carry the anticipated traffic (data and control traffic) to ensure good application performance.
- *Redundancy*—A WAN link failure will result in service outage between serverfarm and the remote site. This can impact network data services and can expose security issue.
- *Synchronization*—Some applications which are hosted locally will require content synchronization with a centralized server.This can often be scheduled for off hours, to avoid adding traffic to WAN links during normal working hours.

Table 2-3 provides a sample list of centralized and distributed servers.

Table 2-3 Sample List of Centralized and Distribution Servers

Data Center Model	Server Function	Deployment Location
Centralized	Database server(i.e., Oracle, Sybase, etc)	Main Site Data Center
	Cisco Unified Call Manager	
	Cisco Presence Server	
	Cisco Digital Media Manager (DMM)	
	E-mail Messaging Server	
Distributed	Hosted services – Web, FTP, DHCP, DNS, NTP	Main and Remote Site Data Center
	Access-Control – Cisco Access Control Server	
	Cisco Video Surveillance Operation Management	
	Media Storage Server	

Remote Site Network Design

The Small Enterprise Design Profile includes two remote site designs. One for larger remote site, and another for medium to smaller remote sites. The typical remote site is a single building with a limited population which makes the collapsed core network design a suitable choice.

Remote Site Collapsed Core Network Design

The key criteria to consider when designing a remote site network are the network size, bandwidth capacity and high-availability requirements. The Small Enterprise Design Profile includes two models: one for smaller remote site and another for larger remote site. Both designs offer high capacity, performance, and availability. Figure 2-10 illustrates the two remote site network design models.

Figure 2-10 Remote Site—Collapsed Core Network Models

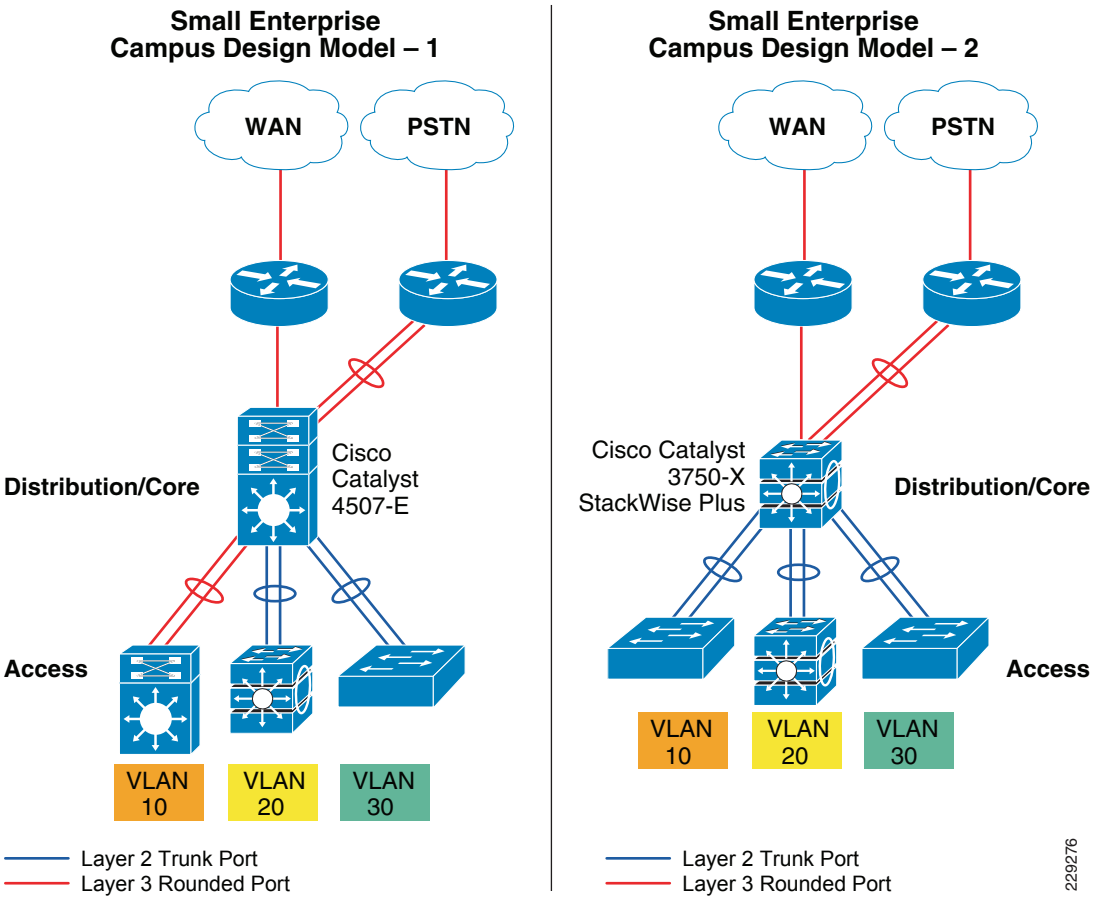

Design Model-1 is for a larger remote site. The network design is the same as the main site network design, with the same performance capabilities, scalability options, and high availability features.

Design Model-2 is for a medium to small remote site. The primary difference is the use of the Cisco Catalyst 3750-X Stack Wise Plus switch in the collapsed core/distribution layer. The 3750-X Stack Wise Plus deploys up to nine 3750-X switches in a ring topology as a single virtual switch. Each chassis replicates the control functions, and provides packet forwarding. If the master switch fails, another switch will quickly assume the control plane 'master' function. This results in a cost effective, high performance, scalable solution, with built in resiliency.

- *Performance*—Provides wire-rate network connection to access switches
- *Scalable*—May deploy up to 9 switches in a stack to aggregate a reasonable number of access switches
- *High Availability*—Stack provides a virtual switch, with distributed control plane, delivering subsecond system failure recovery times

The Cisco 3750-X StackWise Plus delivers high performance routing and switching capability and robust IOS feature support. The control-plane and forwarding paths functions for the Cisco 3750-X StackWise Plus in the collapsed core network design remain the same. However, the switching architecture of the Cisco 3750-X StackWise differs significantly from the high-end distributed and modular switch platforms like Cisco Catalyst 4500-E and 6500-E Series switches.

For more information about the Cisco 3750-X StackWise-Plus architecture, refer to the following URL:

http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps5023/prod_white_paper09186a00801b096a_ps7077_Products_White_Paper.html

Remote Site Access-Layer Design

The access-layer network designate the remote site is the same as at the main site. The same devices are available, and the same design choices may be deployed to achieve a high performance, secure and resilient access layer. To simplify the overall system design, and network operations, it is recommended to use consistent design and platform selections in the access-layer role, at the main and remote sites. This will allow a common configuration template and simplify operations and troubleshooting procedures.

Deploying Network Foundation Services

The two-tier hierarchical design delivers a *reliable and resilient, scalable*, and *manageable* foundation network design. This subsection provides design and deployment guidelines for the small enterprise campus core layer, and access-distribution block.

The access-distribution block, as described in the "Main Site Data Center Network Design" section on page 2-13, uses a combination of Layer-2 and Layer-3 switching to provide a balance of policy and access controls, availability, and flexibility in subnet allocation and VLAN usage. Deployment guidelines are provided to implement multi-layer, and routed access designs in the access-distribution block.

Implementing EtherChannel in Network

Etherchannel is used throughout the network design, and the implementation guidelines are the same for multi-layer, and routed-access models, and in the WAN edge design. As recommended in the "EtherChannel Fundamentals" section on page 2-7, there should be single logical point-to-point EtherChannel deployed between collapsed core and access-layer. The EtherChannel configuration on each end of the link in the access-distribution block must be consistent to prevent a link bundling problem. EtherChannels use link bundling protocols to dynamically bundle physical interfaces into a logical interface.

The following are the benefits of building EtherChannel in dynamic mode:

- Ensure link aggregation parameters consistency and compatibility between switches.
- Ensure compliance with aggregation requirements.
- Dynamically react to runtime changes and failures on local and remote Etherchannel systems
- Detect and remove unidirectional links and multi-drop connections from the Etherchannel bundle.

EtherChannels can be deployed in dynamic or static modes. Both EtherChannel modes can coexist in a single system; however, the protocols (PagP, LACP) can not interoperate with each other.

- *Cisco proprietary link aggregation*—Cisco's implementation of Port Aggregation group Protocol (PAgP) is supported on all the Cisco Catalyst platforms. The PAgP protocol is not supported when the Cisco Catalyst 2960-S or 3750-X Series switches are deployed in FlexStack or StackWise mode. The PAgP protocol can operate in the different channel-group modes shown in Table 2-4 to initialize link bundling process.

Table 2-4 Cisco's Proprietary PAgP Channel-Group Mode

	Distribution	Access-switch and WAN Edge	EtherChannel State
channel-group mode	auto	auto	Non-Operational
	desirable (recommended)	desirable (recommended)	Operational State
	desirable	auto	
	auto	desirable	

- *IEEE 802.3ad link aggregation*—Link Aggregation Control Protocol (LACP) is based on IEEE 802.3ad specification to operate in vendor-independent network environment. LACP link bundling protocol is developed with same goal as Cisco's PAgP. Cisco Catalyst switches in FlexStack and StackWise Plus mode must use LACP to dynamically bundle. LACP can operate in the following different channel-group modes to initialize the link bundling process. See Table 2-5.

Table 2-5 IEEE 802.3ad LACP channel-group Mode

	Distribution	Access-switch and WAN Edge	EtherChannel State
channel-group mode	passive	passive	Non-Operational
	active (recommended)	active (recommended)	Operational State
	active	passive	
	passive	active	

- *Static Mode*—Each system statically bundles selected physical ports into a logical port-channel. In static mode, Etherchannel consistency check is not performed between two switches, which may lead to network protocol instability or network outage due to mis-configuration. This mode is not recommended and should only be considered when EtherChannel is required but side of the link does not support PAgP or LACP link aggregation protocol.

Figure 2-11 Implementing EtherChannel in Main Site Network

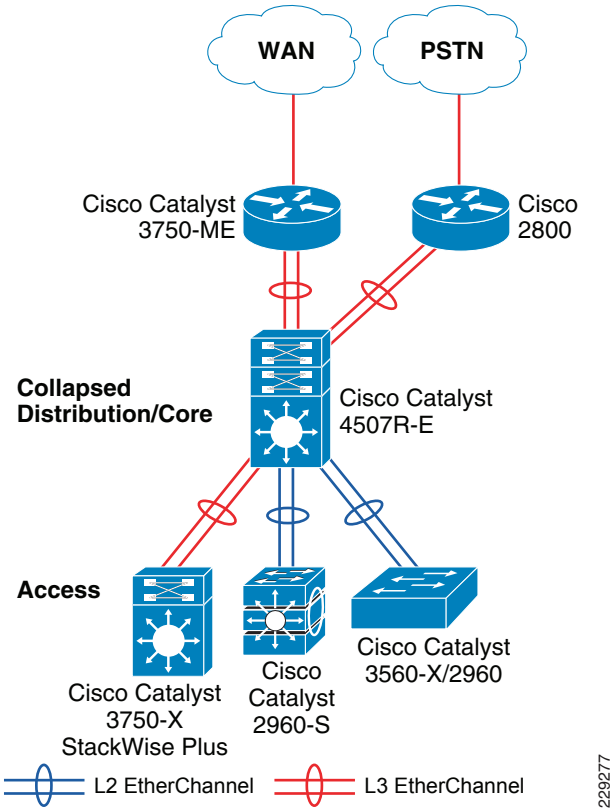

The following sample configuration shows how to build Layer-2 and Layer-3 EtherChannel configuration and bundling physical ports into appropriate logical EtherChannel-group:

```
cr24-4507-1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
cr24-4507-1(config)#interface Port-channel1
cr24-4507-1(config-if)# description Connected to cr24-3750ME-1
cr24-4507-1(config-if)#
cr24-4507-1(config-if)#interface Port-channel11
cr24-4507-1(config-if)# description Connected to cr24-2960-1
cr24-4507-1(config-if)# switchport
cr24-4507-1(config-if)#
cr24-4507-1(config-if)#interface Port-channel16
cr24-4507-1(config-if)# description Connected to cr25-3750s-1
cr24-4507-1(config-if)# switchport
cr24-4507-1(config-if)#
cr24-4507-1(config-if)#interface range Gig 3/3 , Gig 4/3
cr24-4507-1(config-if-range)# description Connected to cr24-3750ME-1
cr24-4507-1(config-if-range)# channel-protocol pagp
cr24-4507-1(config-if-range)# channel-group 1 mode desirable
cr24-4507-1(config-if-range)#
cr24-4507-1(config-if-range)#interface range Gig 1/1 , Gig 2/1
cr24-4507-1(config-if-range)# description Connected to cr24-2960-1
cr24-4507-1(config-if-range)# channel-protocol pagp
cr24-4507-1(config-if-range)# channel-group 11 mode desirable
cr24-4507-1(config-if-range)#
cr24-4507-1(config-if-range)#interface range Gig 1/6 , Gig 2/6
cr24-4507-1(config-if-range)#description Connected to cr26-3750s-1
cr24-4507-1(config-if-range)# channel-protocol lacp
cr24-4507-1(config-if-range)# channel-group 16 mode active
```

Enabling EtherChannel on each switch endpoint will automatically form a logical connection and can be verified using following CLI command:

```
cr24-4507-1#show etherchannel summary | inc Po
Group  Port-channel  Protocol    Ports
1      Po1(RU)         PAgP      Gi3/3(P)    Gi4/3(P)
11     Po11(SU)        PAgP      Gi1/1(P)    Gi2/1(P)
16     Po16(SU)        LACP      Gi1/6(P)    Gi2/6(P)
```

EtherChannel Load Balancing

EtherChannel load-sharing is based on a polymorphic algorithm. On per protocol basis, load sharing is done based on source XOR destination address or port from Layer 2 to 4 header and ports. For higher granularity and optimal utilization of each member-link port, an EtherChannel can intelligently load-share egress traffic using different algorithms. EtherChannel load balancing method support varies on Cisco Catalyst platforms. Table 2-6 summarizes the currently supported EtherChannel load-balancing methods.

Table 2-6 EtherChannel Load Balancing Support Matrix

Packet Type	Classification Layer	Load Balancing Mechanic	Supported Cisco Catalyst Platform
Non-IP	Layer 2	src-dst-mac	29xx, 35xx, 3750, 4500
IP		src-mac	
		dst-mac	
		src-dst-mac	
IP	Layer 3	src-ip	
		dst-ip	
		src-dst-ip	
IP	Layer 4	src-port	4500
		dst-port	
		src-dst-port	

EtherChannel load-balancing mechanisms function on a per-system basis. By default, EtherChannel will use the hash computation algorithm. The network administrator can globally configure the load balancing mechanism. In Cisco Catalyst platforms, EtherChannel load balancing is performed in hardware and it cannot perform per-packet-based load balancing among different member links within EtherChannel. Bandwidth utilization of each member-link may not be equal in default load balancing mode. The Ether Channel load balancing method should be changed to source and destination IP address-based throughout the main and remote site network for the following reasons:

- One cannot optimize load balancing using hash tuning in a general network deployment model. This is due to variations in application deployment and usage patterns.
- EtherChannel does not take into account the bandwidth of each flow. Instead, it relies on the statistical probability that the load is equally distributed across the links of the port-channel group, given a large number of flows of relatively equal bandwidths. However, this may not always be true.

Tuning the load-balancing to source-and-destination IP address allows for statistically-superior load-distribution. When loads are balanced in this manner, packets belonging to a single flow will retain their packet order. See Figure 2-12.

Figure 2-12 EtherChannel Load-Balance Method

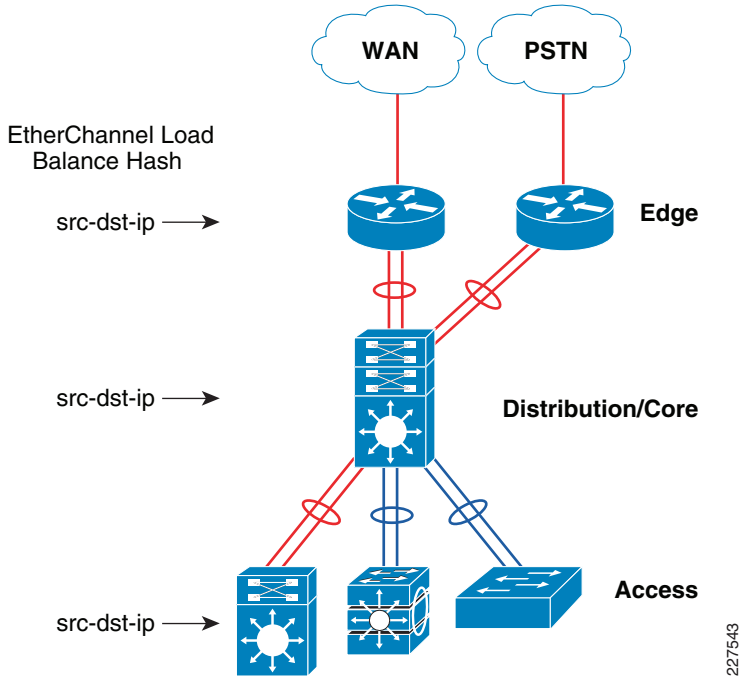

The following output provides sample configuration guideline for changing the default **port-channel load-balance** setting to **source-destination-ip** based. Aside from Layer-2 or Layer-3 EtherChannel mode, similar configuration must be applied on each system in the access-distribution block and WAN edge.

```
cr24-4507-1(config)#port-channel load-balance src-dst-ip
cr24-4507-1#show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
        src-dst-ip
EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
  IPv4: Source XOR Destination IP address
  IPv6: Source XOR Destination IP address
```

The following additional EtherChannel design and configuration must be taken into consideration for an optimal EtherChannel design:

- Enable single EtherChannel between access-layer and distribution system. Enabling more than a single Ether Channel in a collapsed core network design imposes the same limitations as discussed in non-EtherChannel scenario in Figure 2-5.
- For optimal load sharing and hashing computation, it is recommended to bundle the number of physical ports in powers of 2 (i.e., 2, 4, and 8).
- EtherChannel is a logical interface in Cisco Catalyst platform. EtherChannel scalability in collapsed core and distribution must be taken into account. The Cisco Catalyst 4500 can support up to 64 EtherChannels, whereas the Cisco Catalyst 3750 StackWise can support up to 48 EtherChannels per-system.

Deploying Core Network Layer

This section provides implementation and best practice guidelines for deploying the core-layer in both the main and remote enterprise sites. Proper design of the core network layer ensures reachability, transparency and availability. This section focuses on building a unicast routing topology.

Routing Protocol

Enabling routing in the small enterprise network is a simple task. However, the network physical layout must be carefully planned and designed to ensure flexible, stable and efficient routing. Developing a hierarchical network addressing scheme enables a stable, efficient and scalable design.

- *Hierarchical network addressing*—Structured IP network addressing in small enterprise LAN/WAN network is a must to make network scalable, stable.
- *Routing protocol*—Cisco IOS supports wide range of Interior Gateway Protocol (IGP). It is recommended to deploy a single choice of routing protocol across the network infrastructure. This solution guide does not recommended any particular IGP to deploy in the small enterprise network architecture as it significantly varies based on different network infrastructure. However it will provide some key points to be considered when selecting unicast routing protocol.
- *Hierarchical routing domain*—Routing protocols must be designed in a hierarchical model that allows network to scale and operate with greater stability. Building routing boundaries and summarizing the network addresses minimizes topology size and synchronization procedure, which improves the overall network resource utilization and reconvergence.

Routing Protocol Selection Criteria

- *Efficient address allocation*—Hierarchical addressing enables efficient use of address space, since groups are contiguous.
- *Improves routing efficiency*—Using contiguous ip addresses enables efficient route summarization. Route summarization simplifies the routing database, and computations during topology changes. This reduces the network bandwidth used by the routing protocol, and improves routing protocol performance by reducing network convergence time.
- *Improves system performance*—Hierarchical, contiguous ip addressing reduces router memory usage by eliminating dis-contiguous and non-summarized route entries. It saves on CPU cycles needed to compute the routing database during topology changes. This contributes to a more stable routing network, and simplifies the task of network operations and management.

Cisco IOS supports many Interior Gateway Protocols (IGP), including EIGRP and OSPF, either of which are suitable for large network deployments. While OSPF is capable of greater scale, it is also more complex, and hence more difficult to configure, operate and manage. The Small Enterprise Design Profile is designed and validated using EIGRP, since it is a stable, high performance, efficient protocol, which is simple to implement and manage. The same design principles apply whether using EIGRP or OSPF.

Table 2-7 lists some of the EIGRP and OSPF side-by-side feature comparison information.

Table 2-7 EIGRP and OSPF feature comparison chart

Feature	EIGRP	OSPF
Classless Routing	Both routing protocols support classless routing that allows to partition networks into VLSM.	
Loop Prevention	Built-in mechanic to prevent routing loop in network.	
Robust metric	Aggregated link Bandwidth + Delay	Aggregated link Bandwidth
Efficient routing	Partial update	
Multi-access routing adjacency	Full-mesh	Hub-n-spoke
Hierarchical Routing	No. All routers considered in backbone. Non-backbone or routers in non-transit path can be deployed in Stub role.	OSPF area is divided in multiple routing domains. Backbone area maintains complete summarized network topology; non-backbone area can be transit or non-transit OSPF routers.
Network convergence	Both routing protocol offers rapid network recovery during link failure.	
Graceful-Restart Support	Yes	Yes. Cisco and IETF based
Route Summarization	Flexibility to manual summarized on any routing node.	Can only be performed on ABR or ASBR
Load-Balancing	Support equal and un-equal cost load balancing	Equal-cost path only.
Standard	Cisco proprietary	IETF standard

Designing End-to-End EIGRP Routing Domain

EIGRP is a balanced hybrid routing protocol that builds neighbor adjacency and a flat routing topology on a per-autonomous-system (AS) basis. The LAN/WAN infrastructure of Small Enterprise Design Profile should be deployed in a single EIGRP AS to prevent route redistribution, loops, and other problems that may occur due to misconfiguration. See Figure 2-13.

Figure 2-13 End-to-End EIGRP Routing Design in Small Enterprise Architecture

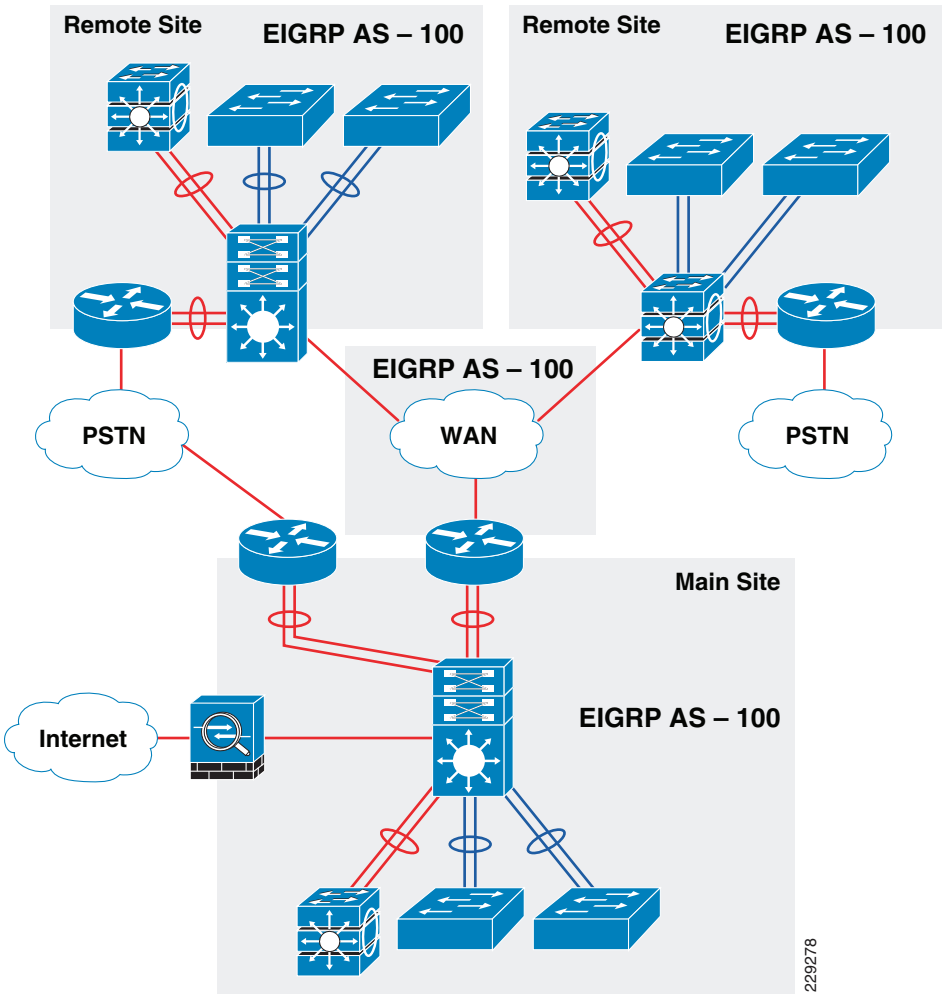

Implementing EIGRP Routing

The main site is the central hub in the network. Each remote site is connected to the main site over the WAN infrastructure. The main site network includes the Internet gateway and provides access to the central data-center. Since main and remote sites use the collapsed core design, the routing configuration of the core routers is the same.

The following is a sample configuration to enable EIGRP routing process at the edge of the main site collapsed core network. EIGRP is enabled in the remote sites network with the same configuration:

```
cr24-4507-1(config)#interface Loopback0
cr24-4507-1(config-if)# ip address 10.125.100.1 255.255.255.255

cr24-4507-1(config-if)#interface Port-channel1
cr24-4507-1(config-if)# description Connected to cr24-3750ME-1
cr24-4507-1(config-if)#no switchport
cr24-4507-1(config-if)# ip address 10.125.32.4 255.255.255.254

cr24-4507-1(config-if)#interface Port-channel2
cr24-4507-1(config-if)# description Connected to cr24-2851-1
cr24-4507-1(config-if)#no switchport
cr24-4507-1(config-if)# ip address 10.125.32.6 255.255.255.254
```

```
cr24-4507-1(config)#interface Vlan200
cr24-4507-1(config-if)# description Connected to cr24_ASA_Inside_Port
cr24-4507-1(config-if)# ip address 10.125.33.9 255.255.255.0
```

```
cr24-4507-1(config)#router eigrp 100
cr24-4507-1(config-router)# no auto-summary
cr24-4507-1(config-router)# eigrp router-id 10.125.100.1
cr24-4507-1(config-router)# network 10.125.0.0 0.0.255.255

cr24-4507-1#show ip eigrp neighbor port-channel 13
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface     Hold Uptime   SRTT   RTO   Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.125.33.10Vl200111d00h       1   200  0  171
0   10.125.32.7Po2161d02h       1   200  0  304
2   10.125.32.5Po1141d02h       2   200  0  25038
```

EIGRP Adjacency Protection

Implementing summarization in the EIGRP routing process automatically enables EIGRP routing process on each interface that is summarized. By default, the router transmits and accept EIGRP hello messages from remote device to form an adjacency on all EIGRP enabled interfaces. This behavior needs to be modified to ensure a secure, efficient and stable routing design:

- *System efficiency*—There is no need to send EIGRP hellos on an interface where there is no trusted EIGRP neighbor. In a large network, sending EIGRP hello messages periodically to such interfaces consumes unnecessary CPU resource. EIGRP route processing should only be enabled on interfaces where trusted network devices are connected. All other interfaces can be suppressed in passive mode. The following configuration shows how to automatically disable EIGRP processing on all the Layer-3 interfaces and only enable on the trusted interface. This design principle must be applied on each EIGRP router, including distribution and core routers:

```
cr24-4507-1(config)#router eigrp 100
cr24-4507-1(config-router)# network 10.125.0.0 0.0.255.255
cr24-4507-1(config-router)# passive-interface default
cr24-4507-1(config-router)# no passive-interface Port-channel1
cr24-4507-1(config-router)# no passive-interface Port-channel2
cr24-4507-1(config-router)# no passive-interface Vlan200

cr24-3560r-1#show ip eigrp interface
EIGRP-IPv4:(100) interfaces for process 100

                    Xmit           Queue   Mean   Pacing Time   Multicast Pending
Interface    Peers      Un/Reliable  SRTT   Un/Reliable   Flow Timer  Routes
Vl2001       0/01   0/1     50           0
Po1 1        0/02   0/1     50           0
Po2 1        0/04   0/1     50           0
```

```
cr24-4507-1#show ip protocols | inc Passive|Vlan
  Passive Interface(s):
    Vlan1
    Vlan101
    Vlan102
```

```
Vlan103
Vlan104
```

- *Network Security*—Sending unnecessary EIGRP Hello messages opens a security vulnerability in two ways. An attacker can detect EIGRP operation and send flood of EIGRP hello messages to destabilize the network. Or an attacker could establish a "fake" EIGRP adjacency and advertise a best metric default-route into the network to black hole and compromise all critical traffic. Each EIGRP system should implement MD5 authentication, and each EIGRP neighbor should validate MD5 authentication is enabled on adjacent systems. This provides a secure method of transmitting and receiving routing information between devices in the network. Following is a sample configuration to enable EIGRP neighbor authentication using MD5:
  - Distribution

```
cr24-4507-1(config)#key chain eigrp-key
cr24-4507-1(config-keychain)# key 1
cr24-4507-1(config-keychain-key)#   key-string <password>

cr24-4507-1(config)#interface Port-channel1
cr24-4507-1(config-if)# description Connected to cr24-3750ME-1
cr24-4507-1(config-if)# ip authentication mode eigrp 100 md5
cr24-4507-1(config-if)# ip authentication key-chain eigrp 100 eigrp-key
```

  - WAN Aggregation

```
cr24-3750ME-1(config)#key chain eigrp-key
cr24-3750ME -1(config-keychain)# key 1
cr24-3750ME -1(config-keychain-key)#   key-string <password>

cr24-3750ME -1(config)#interface Port-channel1
cr24-3750ME -1(config-if)# description Connected to cr24-4507-1
cr24-3750ME -1(config-if)# ip authentication mode eigrp 100 md5
cr24-3750ME -1(config-if)# ip authentication key-chain eigrp 100 eigrp-key
```

- *System Stability*—As mentioned in Table 8, EIGRP allows network administrator to summarize multiple individual and contiguous networks into a single summarized network before advertising to neighbors. Route summarization improves performance, stability, and convergence times, and it makes the network easier to manage operate and troubleshoot.

EIGRP provides the flexibility to summarize at any point in the network. Proper design requires determining which routers will serve as Aggregators, and advertise summarized network information to peers. Routers which connect multiple access devices, or connect to the WAN edge should be made Aggregators. Figure 2-14 provides an example small enterprise network with route aggregator devices identified with the direction of route summarization illustrated.

Figure 2-14 Route Aggregator and Summary Route Advertisement Direction

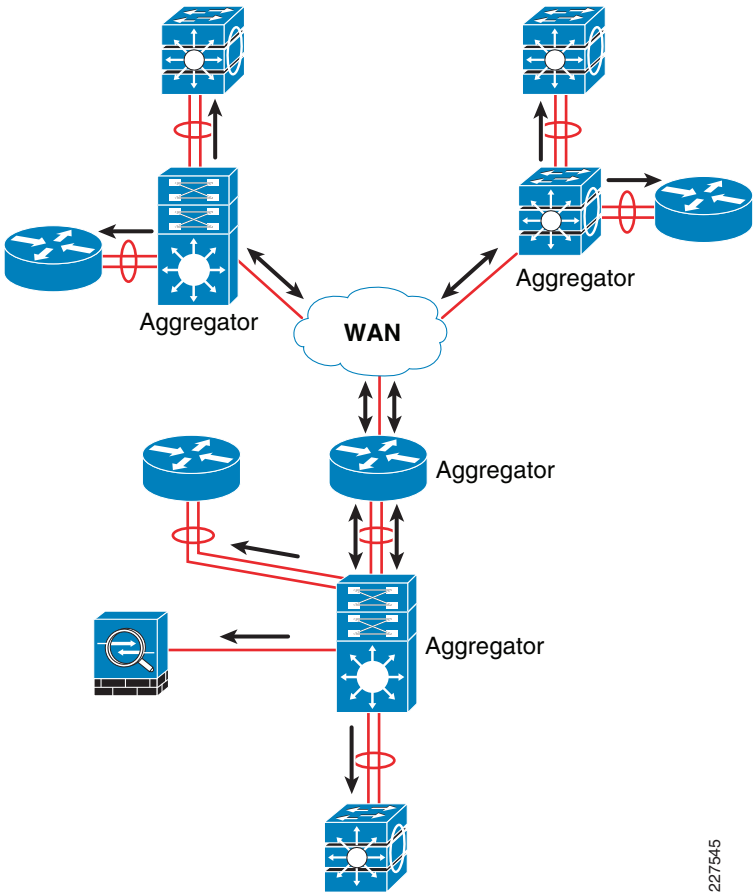

The following sample configuration shows EIGRP route summarization. In this example, the entire access-layer network is summarized into a single classless network and advertised to the WAN edge, the ASA firewall and the PSTN gateway:

- Distribution

```
cr24-4507-1(config)#interface Port-channel1
cr24-4507-1(config-if)# description Connected to cr24-3750ME-1
cr24-4507-1(config-if)# ip summary-address eigrp 100 10.125.0.0 255.255.0.0

cr24-4507-1(config-if)#interface Port-channel2
cr24-4507-1(config-if)# description Connected to cr24-2851-1
cr24-4507-1(config-if)# ip summary-address eigrp 100 10.125.0.0 255.255.0.0

cr24-4507-1(config-if)#interface Vlan200
cr24-4507-1(config-if)# description Connected to cr24_ASA_Inside_Port
cr24-4507-1(config-if)# ip summary-address eigrp 100 10.125.0.0 255.255.0.0

cr24-4507-1#show ip protocols | inc Address|10.125.0.0
Address Family Protocol EIGRP-IPv4:(100)
  Address Summarization:
    10.125.0.0/16 for Port-channel1, Vlan200, Port-channel2
```

- WAN Aggregation

  Verifying main site EIGRP summarized route status at WAN aggregation layer as follows:

```
cr24-3750ME-1#show ip route 10.125.0.0 255.255.0.0
Routing entry for 10.125.0.0/16
  Known via "eigrp 100", distance 90, metric 1792, type internal
  Redistributing via eigrp 100
  Last update from 10.125.32.4 on Port-channel1, 1d04h ago
  Routing Descriptor Blocks:
  * 10.125.32.4, from 10.125.32.4, 1d04h ago, via Port-channel1
      Route metric is 1792, traffic share count is 1
      Total delay is 20 microseconds, minimum bandwidth is 2000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Tuning EIGRP Protocol Timers

EIGRP uses Hello messages to form adjacencies and determine if neighbors are alive. EIGRP adjacency is declared down if it fails to receive Hello messages within the Hold down timer interval. All the prefixes discovered from a dead neighbor are removed from the routing table. By default, EIGRP transmits a Hello message every 5 seconds to notify neighbors that it is still alive. The EIGRP hold-down timer gets reset each time the router receives a EIGRP Hello message. Default EIGRP adjacency hold-down timer is 15 seconds.

Lowering EIGRP hello and hold-down timer intervals improves network convergence times (i.e. time to detect and respond to an outage). For small enterprise network design it is recommended to use the default EIGRP Hello and Hold timer values for the following reasons:

- *EtherChannel benefits*—EIGRP operates over the Layer-3 EtherChannel. In the event of a single member-link failure condition, layer 2 will respond more quickly than the routing protocol, and switchover traffic from the impacted link to an alternate member link. EIGRP routing is not impacted by individual link member and no change in the routing table is required. Thus reducing the EIGRP timers will not result in quicker convergence, and may adversely impact system stability.
- *High availability*—The Cisco Catalyst 4507R-E and 3750-X Stack Wise Plus layer 3 switches support graceful-restart protocol extensions which enables a redundant module or member switch to gracefully assume the active role while maintaining adjacency with neighbors, during a active supervisor failure condition. The backup supervisor requires sufficient time to detect a failure and initiate graceful recovery with neighbors. Implementing aggressive timers may abruptly terminate adjacency and cause network outage before a stateful switch over is accomplished. Thus, default EIGRP Hello and Hold timers are recommended on Cisco Catalyst 4507R-E and 3750-X Stackwise Plus Series Layer-3 platforms.

Deploying Multi-Layer Network

Multilayer design is one of the two access-distribution block designs included in the Small Enterprise Design Profile. This section provides implementation and best practices guidelines the multi-layer design. The deployment and configuration guidelines for the multi-layer access-distribution block are the same for both main and remote site networks.

Spanning-Tree in Multilayer Network

Spanning Tree (STP) is a Layer-2 protocol that prevents logical loops in switched networks with redundant links. The Small Enterprise Design Profile uses Etherchannel (point-to-point logical Layer-2 bundle) connection between access-layer and distribution switch which inherently simplifies the STP topology and operation. In this design, the STP operation is done on a logical port, therefore, it will be assigned automatically in forwarding state.

Over the years, the STP protocols have evolved into the following versions:

- *Per-VLAN Spanning Tree Plus (PVST+)*—Provides a separate 802.1D STP for each active VLAN in the network.
- *IEEE 802.1w – Rapid PVST+*—Provides an instance of RSTP (802.1w) per VLAN. It is easy to implement, proven in large scale networks that support up to 3000 logical ports and greatly improves network restoration time.
- *IEEE 802.1s – MST*—Provides up to 16 instances of RSTP (802.1w) and combines many VLANs with the same physical and logical topology into a common RSTP instance.

Following is the example configuration to enable STP protocol in multi-layer network:

Distribution

```
cr24-4507-1(config)#spanning-tree mode rapid-pvst

cr24-4507-1#show spanning-tree summary | inc mode
Switch is in rapid-pvst mode
```

Access-Layer Switch

```
cr24-2960-1(config)#spanning-tree mode rapid-pvst
```

Default STP parameters optimize the network for packet forwarding. Best practice design includes hardening STP parameters in the access and distribution switch to protect against STP misconfiguration, or malicious user by deploying spanning-tree toolkit in the access-distribution block. See Figure 2-15.

Figure 2-15 Hardening Spanning-Tree Toolkit in Multi-Layer Network

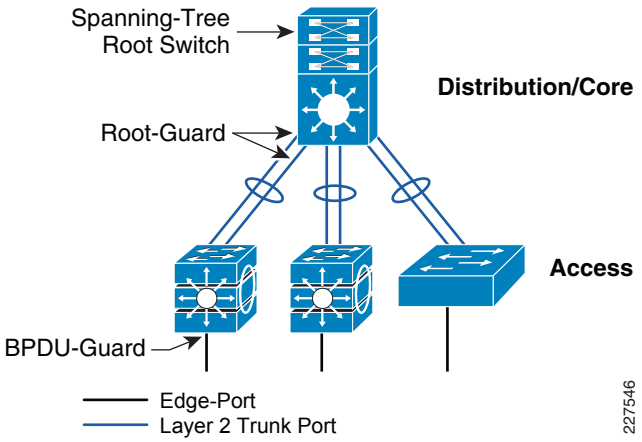

The following is the configuration deploys spanning-tree toolkit in the access-distribution block:

Distribution

```
cr24-4507-1(config)#spanning-tree vlan 1-4094 root primary
cr24-4507-1(config)#interface range Gig 1/1 - 2 , Gig 2/1 - 2
cr24-4507-1(config)#spanning-tree guard root
```

Access

```
cr26-2960S-1(config)#interface GigabitEthernet1/0/1
cr26-2960S-1(config-if)#description CONNECTED TO UNTRUSTED-PC
cr26-2960S-1(config-if)#spanning-tree bpduguard enable
```

Other STP Toolkit Consideration

When the access-distribution block multi-layer design is deployed using the recommended best practices, it automatically minimizes the need for deploying the following additional spanning-tree toolkit technologies:

- *UplinkFast*—Improves the network convergence time by providing direct access to the root switch link failure. UplinkFast is not necessary in this design, because there is no alternate STP path and RSTP protocol natively includes rapid recovery mechanism.
- *BackBone Fast*—Provides rapid convergence from indirect Layer-2 link failures in a redundant distribution switch configuration. This is feature is not necessary for the same reason as stated for UplinkFast.
- *LoopGuard*—Protects Layer-2 networks from loops that occur due to any malfunction that prevents normal BPDU forwarding. A STP loop is created when a blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports in a physically redundant topology (not necessarily the blocking port) stopped receiving BPDUs. Because there is single point-to-point STP forwarding port in this design, enabling Loopguard does not provide any additional benefit. UDLD protocol must be implemented to prevent STP loop that may occur in the network due to network malfunction, mis-wiring, etc.

Logical Multi-Layer Network

VLAN assignment can have a significant impact on network performance and stability. There are three basic ways to assign VLANs within the access-distribution block.

Flat Logical Network Design

Spanning a single VLAN across multiple access-layer switches is much simpler with a single collapsed core-distribution device versus a design with redundant distribution devices. The flat multi-layer design has a single VLAN across multiple access devices, as shown in See Figure 2-16.

Figure 2-16 Multi-Layer Flat Network Design

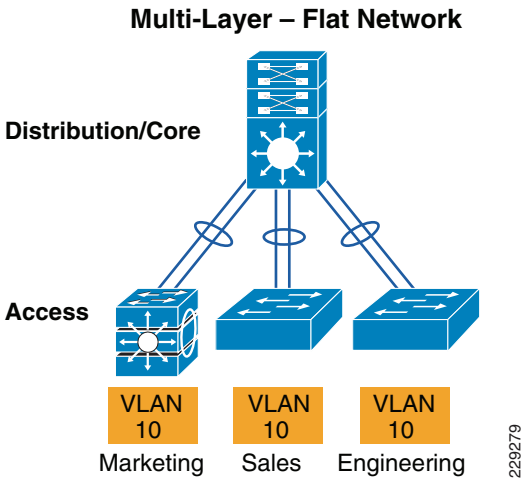

A flat multi-layer network deployment introduces the following challenges:

- *Scalability*—Spanning the same VLAN in different access-layer switches will create a large Layer-2 broadcast domain that dynamically discovers and populates MAC address entries for endpoints that may not need to communicate. In a large network, this may become a scalability issue (i.e. memory required to hold large CAM table).
- *Performance*—In a large network, spanning a large number of broadcast domains will impact the performance of all network devices in the access-distribution block, because the switch will have to process many more broadcast packets such as ARP.
- *Security*—The flat muli0layer design widens the fault domain which increases possible attacks to a larger number of users.The number of users is not necessarily due to the number switches spanned and applications during DoS or viruses attack.

Segmented Logical Network Design

Best practice design includes identifying meaningful groups within the user community, and assigning a unique VLAN to each group. These groups may be departments, user groups, or any other logical grouping of users. Enabling a unique VLAN for each group will segment the network and build a logical network structure. All network communication between groups will pass through the routing and forwarding policies defined at the distribution layer. See Figure 2-17.

Figure 2-17 Multi-Layer Segmented Network Design

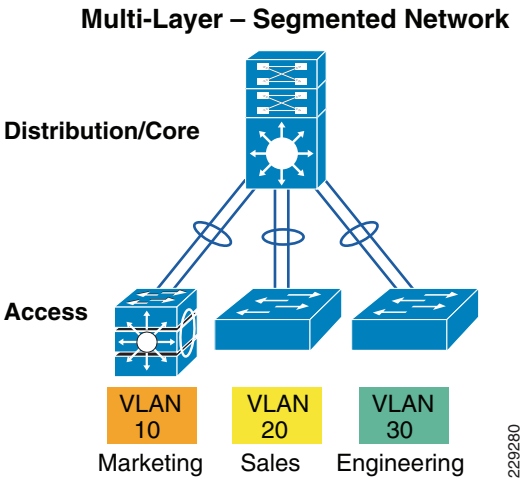

A segmented VLAN design is the solution to the challenges described in the flat network design. VLAN segmentation improves the scalability, performance, and security of the network.

Hybrid Logical Network Design

The segmented logical network design improves scalability, performance and security, and addresses the challenges of a flat network design. In real world deployments, there is usually a need for some users or applications to communicate with all users (eg system administrator). The hybrid network design is the segmented design, with the addition of a exceptional VLAN which spans the entire access-distribution block. See Figure 2-18.

Figure 2-18 Multi-Layer Hybrid Network Design

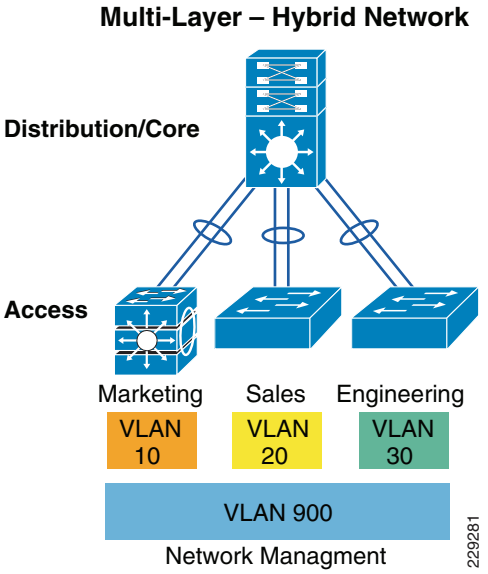

Cisco recommends the segmented VLAN network design and optionally hybrid network for centralized users or applications that requires distributed function across the access-layer network.

Following are the sample VLAN configuration steps in the access and the distribution layer switches.

Distribution

VLAN Trunking Protocol (VTP) is a Cisco proprietary Layer -messaging protocol that manages the addition, deletion, and renaming of VLANs on a network-wide basis. Cisco's VTP simplifies administration in a switched network. VTP can be configured in three modes: server, client, and transparent. Set the VTP domain name and change the mode to the transparent mode as follows:

```
cr24-4507-1(config)#vtp domain campus
cr24-4507-1(config)#vtp mode transparent

cr24-4507-1(config)#vlan 10
cr24-4507-1(config-vlan)#name cr24-3750-Mktg-Dept
cr24-4507-1(config-vlan)#vlan 20
cr24-4507-1(config-vlan)#name cr24-3560-Sales-Dept
cr24-4507-1(config-vlan)#vlan 30
cr24-4507-1(config-vlan)#name cr24-2960-Engg-Dept
```

Access

Set VTP domain name and change the mode to the transparent mode as follows:

```
cr24-3750-1(config)#vtp domain campus
cr24-3750-1(config)#vtp mode transparent

cr24-3750-1(config)#vlan 10
cr24-3750-1(config-vlan)#name cr24-3750-Sales-Dept
```

Implementing Layer 2 Trunk

In a typical network design, a single access switch will have more than one VLAN, for example a Data VLAN and a Voice VLAN. The network connection between Distribution and Access device is a trunk. VLAN's tag their traffic to maintain separation between VLANS across the trunk. By default on Cisco Catalyst switches, the native VLAN on each layer 2 trunk port is VLAN 1, and cannot be disabled or removed from VLAN database. The native VLAN remains active on all access switches layer 2 ports.

There are two choices for encapsulating the tagged VLAN traffic on the trunk: IEEE 802.1Q or Cisco ISL. It is recommended to implement trunk encapsulation in static mode instead of negotiating mode, to improve the rapid link bring-up performance. Not all Cisco Catalyst platforms support ISL encapsulation; therefore IEEE 802.1Q is recommended, and validated in the access and distribution switches.

Enabling the Layer-2 trunk on a port-channel, automatically enables communication for all of the active VLANs between the access and distribution. This means an access-switch which has implemented, for example, VLANs 10 to 15, will receive flood traffic destined for VLANs 20 to 25, which are implemented on another access switch. RPVST +, using logical ports, operates on a per-VLAN basis to load balance traffic. In a large network, it is important to limit traffic on Layer-2 trunk ports to only the assigned VLANS, to ensure efficient and secure network performance. Allowing only assigned VLANs on a trunk port automatically filters rest.

The default native VLAN must be properly configured to avoid several security risks—Attack, worm and virus or data theft. Any malicious traffic originated in VLAN 1 will span across the access-layer network. With a VLAN-hopping attack it is possible to attack a system which does not reside in VLAN 1. Best practice to mitigate this security risk is to implement a unused and unique VLAN ID as a native VLAN on the Layer-2 trunk between the access and distribution switch. For example, configure VLAN

802 in the access-switch and in the distribution switch. Then change the default native VLAN setting in both the switches. Thereafter, VLAN 802 must not be used anywhere for any purpose in the same access-distribution block.

Following is the configuration example to implement Layer-2 trunk, filter VLAN list and configure the native-VLAN to prevent attacks on port channel interface. When the following configurations are applied on port-channel interface (i.e., Port-Channel 11), they are automatically inherited on each bundled member-link (i.e., Gig1/1 and Gig2/1):

Distribution

```
cr24-4507-1(config)#vlan 802
cr24-4507-1(config-vlan)#name Admin-Hopping-VLAN

cr24-4507-1(config)#interface Port-channel 11
cr24-4507-1(config-if)# description Connected to cr24-3750-1
cr24-4507-1(config-if)# switchport
cr24-4507-1(config-if)# switchport mode trunk
cr24-4507-1(config-if)# switchport trunk allowed vlan 101-110,900
cr24-4507-1(config-if)# switchport trunk native vlan 802
```

```
cr24-4507-1#show interface port-channel 11 trunk

Port        Mode             Encapsulation  Status        Native vlan
Po11        on               802.1q         trunking      802

Port        Vlans allowed on trunk
Po11        101-110,900

Port        Vlans allowed and active in management domain
Po11        101-110,900

Port        Vlans in spanning tree forwarding state and not pruned
Po11        101-110,900
```

Access-switch

```
cr24-3750-1(config)#vlan 802
cr24-3750-1(config-vlan)#name Admin-Hopping-VLAN

cr24-3750-1(config)#interface Port-channel 1
cr24-3750-1(config-if)# description Connected to cr24-4507-1
cr24-3750-1(config-if)# switchport
cr24-3750-1(config-if)# switchport mode trunk
cr24-3750-1(config-if)# switchport trunk allowed vlan 101-110,900
cr24-3750-1(config-if)# switchport trunk native vlan 802
```

Unidirectional Link Detection

UDLD is a Layer 2 protocol that works with the Layer 1 features to determine the physical status of a link. At Layer 1, auto-negotiation takes care of physical signaling and fault detection. UDLD performs tasks that auto-negotiation cannot perform, such as detecting the identity of neighbors and shutting down

misconnected ports. When both auto-negotiation and UDLD are enabled, Layer 1 and Layer 2 detection works together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

Copper media ports use Ethernet link pulse as a link monitoring tool and are not susceptible to unidirectional link problems. Because one-way communication is possible in fiber-optic environments, mismatched transmit/receive pairs can cause a link up/up condition even though bidirectional upper-layer protocol communication has not been established. When such physical connection errors occur, it can cause loops or traffic black holes. UDLD functions transparently on Layer-2 or Layer-3 physical ports. UDLD operates in one of two modes:

- *Normal mode*—If bidirectional UDLD protocol state information times out; it is assumed there is no-fault in the network, and no further action is taken. The port state for UDLD is marked as undetermined. The port behaves according to its STP state.
- *Aggressive mode*—If bidirectional UDLD protocol state information times out, UDLD will attempt to reestablish the state of the port, if it detects the link on the port is operational. Failure to reestablish communication with UDLD neighbor will force the port into the err-disable state. That must be manually recovered by user or the switch can be configured for auto recovery within specified interval of time.

Following is the configuration example to implement UDLD protocol:

Distribution

```
cr24-4507-1(config)#interface range Gig 1/2 , Gig 2/2
cr24-4507-1(config-int)#udld port

cr24-4507-1#show udld neighbor
PortDevice Name   Device ID     Port IDNeighbor State
------     ----------------    ------------ -------- --------------
Gi1/2    FOC1318Y06V    1            Gi1/0/49Bidirectional
Gi2/2    FOC1318Y06J     1            Gi3/0/49Bidirectional
```

Access

```
cr26-2975-1(config)#interface Gig 1/0/49 , Gig 3/0/49
cr26-2975-1(config-if)#description Connected to cr24-4507-1
cr26-2975-1(config-if)#udld port

cr26-2975-1#show udld neighbor
PortDevice Name   Device ID     Port IDNeighbor State
------     ----------------    ------------ -------- --------------
Gi1/0/49 FOX1216G8LT     1            Gi1/2      Bidirectional
Gi3/0/49 FOX1216G8LT     1            Gi2/2      Bidirectional
```

Deploying Routed-Access Network

This section provides implementation and best practices guidelines to deploy routed-access in the access-distribution block. The routed access design moves the boundary between Layer 2 and Layer 3 from the distribution layer to the access layer as seen in Figure 2-19.

Figure 2-19 Control Function in Multi-Layer and Routed-Access Network Design

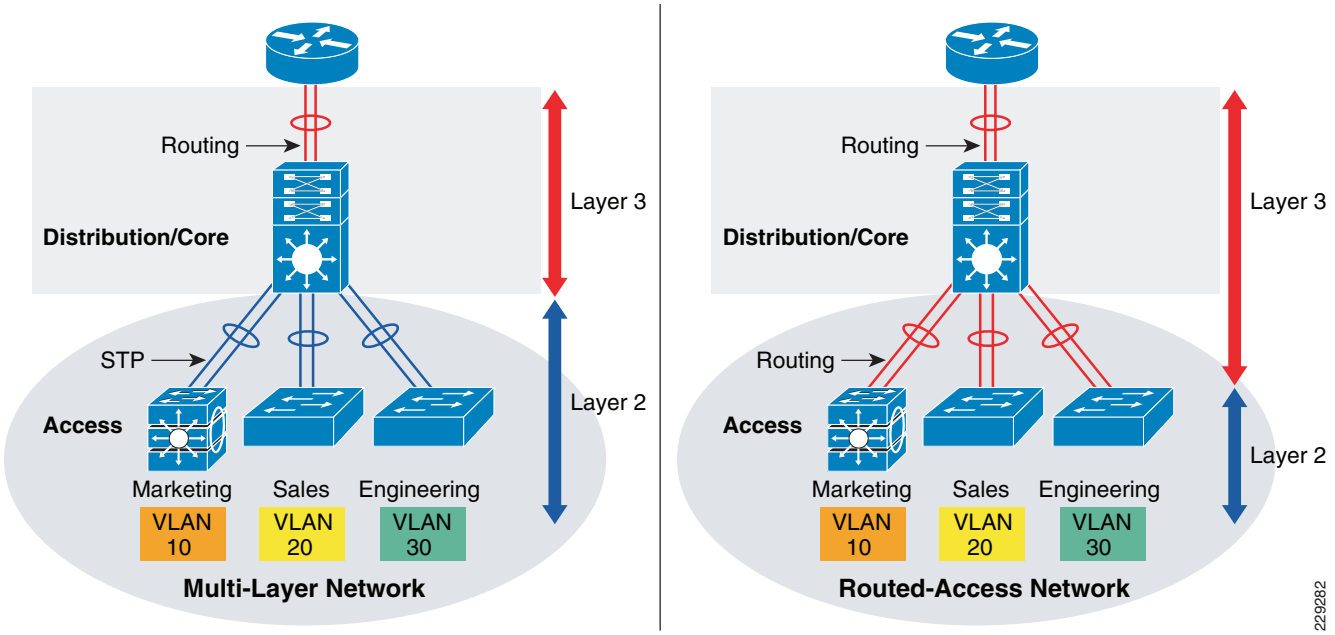

Routing in the access-layer simplifies configuration, optimizes distribution performance, and improves end-to-end troubleshooting tools. Implementing routing in the access-layer replaces Layer-2 trunk configuration with single point-to-point Layer-3 interface in distribution layer. Placing Layer-3 function one tier down on access-switches, changes the multilayer network topology and forwarding path. Implementing Layer-3 function in the access-switch does not require a physical or logical link reconfiguration; the same EtherChannel in access-distribution block can be used.

At the network edge, Layer-3 access-switches provides an IP gateway and become the Layer-2 demarcation point to locally connected endpoints that could be logically segmented into multiple VLANs. Following are the benefits of implementing routed-access in the access-distribution block:

- Eliminates the need to implement STP and the STP toolkit in the distribution layer. As a best practice, STP toolkit must be hardened at the access-layer.
- Shrinks the Layer-2 fault domain, which minimizes the number of endpoints affected by a DoS/DDoS attack.
- Improves Layer-3 uplink bandwidth efficiency by suppressing Layer-2 broadcasts at the access edge port.
- Improves performance by reducing resource utilization in collapsed core-distribution layer. In a large multilayer network, the aggregation layer may consume more CPU cycles due to the large number of MAC and ARP discovery and processing and storing required for each end-station. Routed-access reduces the load of this Layer-2 processing and storage in the distribution layer, by moving the load to layer-3 access-switches. Figure 2-20 illustrates where Layer-2 and Layer-3 forwarding entry processing and storage takes place when access-distribution block is implemented as multi-layer versus routed-access network.

Figure 2-20 Forwarding Entry Development in Multi-tier Network

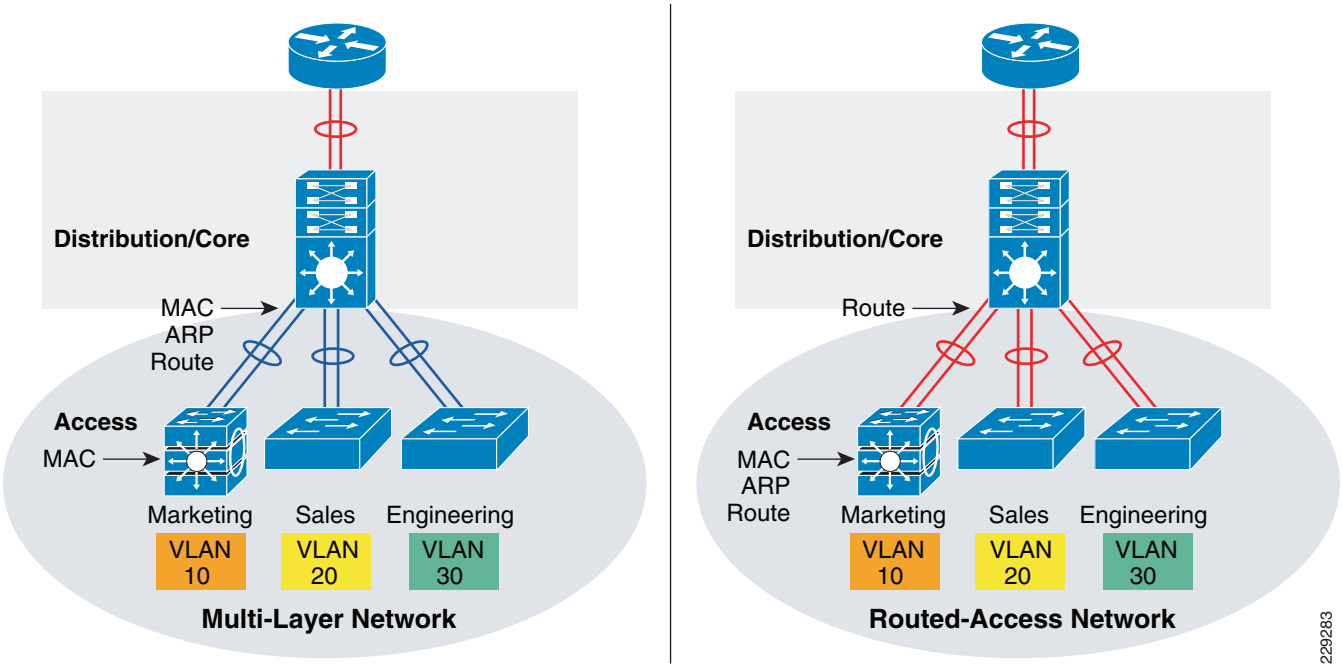

While the routed access design is appropriate for many small enterprise networks it is not suitable for all environments. Routed access does not allow a VLAN to span multiple access switches. Refer to following URL for detailed design guidance for the routed access distribution block design:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/routed-ex.html

Implementing EIGRP Routing in Access-Distribution Block

The Small Enterprise Design Profile uses EIGRP routing protocol, and all the devices in the LAN and WAN sub-networks are deployed in a single AS. This subsection focuses on implementing EIGRP in the access-distribution block. All the deployment and configuration guidelines in this section are the same for deploying in the main or remote site network.

Following is the example configuration to enable basic EIGRP routing in the distribution layer and in the access layer:

Distribution

```
cr24-4507-1(config)#interface Port-channel13
cr24-4507-1(config-if)# description Connected to cr24-3560r-1
cr24-4507-1(config-if)#no switchport
cr24-4507-1(config-if)# ip address 10.125.32.0 255.255.255.254

cr24-4507-1(config)#router eigrp 100
cr24-4507-1(config-router)# no auto-summary
cr24-4507-1(config-router)# eigrp router-id 10.125.100.1
cr24-4507-1(config-router)# network 10.125.0.0 0.0.255.255

cr24-4507-1#show ip eigrp neighbor port-channel 13
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO   Q  Seq
                                            (sec)         (ms)           Cnt Num
3   10.125.32.1             Po13              14 00:02:14    2         200   0  385
```

Access

```
cr24-3560r-1(config)#interface Loopback0
cr24-3560r-1(config-if)# ip address 10.125.100.4 255.255.255.255
cr24-3560r-1(config-if)#
cr24-3560r-1(config-if)#interface Port-channel1
cr24-3560r-1(config-if)# description Connected to cr24-4507-1
cr24-3560r-1(config-if)# no switchport
cr24-3560r-1(config-if)# ip address 10.125.32.1 255.255.255.254


cr24-3560r-1(config)#ip routing

cr24-3560r-1(config)#router eigrp 100
cr24-3560r-1(config-router)# no auto-summary
cr24-3560r-1(config-router)# eigrp router-id 10.125.100.4
cr24-3560r-1(config-router)# network 10.125.0.0 0.0.255.255

cr24-3560r-1#show ip eigrp neighbor port-channel 1
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface       Hold   Uptime   SRTT   RTO  Q  Seq
                                            (sec)            (ms)             Cnt Num
0   10.125.32.0             Po1               13     00:10:00    1       200   0  176
```

Building EIGRP Network Boundary

EIGRP creates and maintains a single flat routing network topology between EIGRP peers. Building a single routing domain enables complete network visibility and reach ability between all of the elements within the network.(access, distribution, core, serverfarm, WAN, etc)

In a tiered design, the access layer always has a single physical or logical forwarding path to the distribution layer. The access switch will build a forwarding topology pointing to same distribution switch as a single Layer-3 next-hop. Since the distribution switch provides a gateway function to the access switch, the routing design can be optimized with the following two techniques to improve performance and network convergence in the access-distribution block:

- Deploy Layer 3 access-switch in EIGRP stub mode
- Summarize network view to Layer-3 access-switch for intelligent routing function

Deploy Layer 3 Access-Switch in EIGRP Stub Mode

The Layer-3 access switch can be deployed to announce itself as a stub router that acts as a non-transit router and does not connect any other Layer-3 stub or non-stub routers. Announcing itself as a non-transit stub Layer-3 router is one way to notify the distribution router that it should not include the Layer-3 access switch in the EIGRP topology recomputation process. This optimized recomputation process will prevent unnecessary EIGRP network queries, which reduces network traffic, and simplifies the route computation.

As illustrated in Figure 2-21, implementing EIGRP stub function in the access switches, greatly reduces the number of EIGRP network queries.

Figure 2-21 EIGRP Query Path with and without Stub Implementation

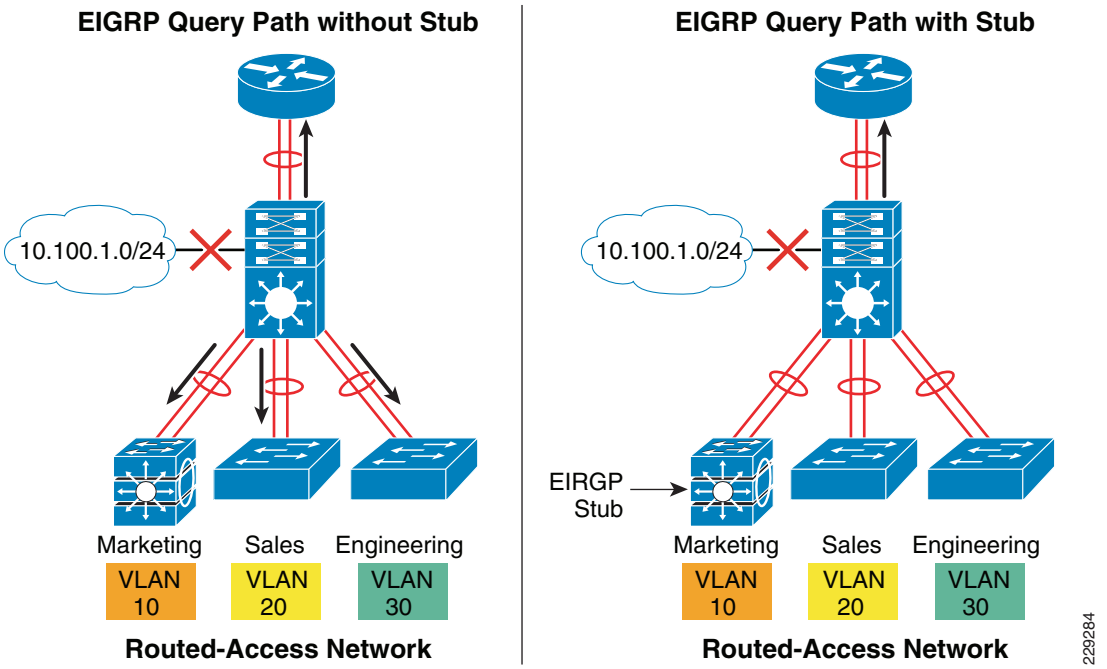

EIGRP stub router in Layer-3 access-switch can announce routes to a distribution-layer router with great flexibility.

EIGRP stub router can be deployed to announce routes dynamically discovered or statically configured. Best practice design is to deploy EIGRP stub router to announce locally learned routes to aggregation layer.

Following is the example configuration to enable EIGRP stub routing in the Layer-3 access-switch, no configuration changes are required in distribution system:

Access

```
cr24-3560r-1(config)#router eigrp 100
cr24-3560r-1(config-router)#eigrp stub connected

cr24-3560r-1#show eigrp protocols detailed

Address Family Protocol EIGRP-IPv4:(100)
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  EIGRP NSF-aware route hold timer is 240
  EIGRP stub, connected
  Topologies : 0(base)
```

Distribution

```
cr24-4507-1#show ip eigrp neighbors detail port-channel 13
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.125.32.1             Po13              13 00:19:19   16   200  0  410
   Version 12.2/3.0, Retrans: 0, Retries: 0, Prefixes: 11
   Topology-ids from peer - 0
```

```
Stub Peer Advertising ( CONNECTED ) Routes
Suppressing queries
```

Summarizing Stub Routed-Access Network

Enabling the EIGRP stub function on the access switch does not change the distribution router behavior of forwarding the full EIGRP topology table. The Distribution router must be configured to advertise summarized routes that do not compromise end-to-end reach ability, and help access switches maintain minimal routing information. In a network with a well designed IP addressing scheme, the aggregation system can advertise summarized routes in a classless address configuration, that reduce individual network advertisements, improve network scalability and network convergence. The distribution router must have full network topology information to ensure efficient reachability paths. Therefore, it is recommended to summarize at the distribution router, and not summarize at the access-layer.

Route summarization must be implemented on the distribution layer of main and each remote site network. This includes devices such as the WAN aggregation in the main site. The distribution router must advertise the following summarized network information to Layer 3 access-switch:

- *Local Network*—Distribution router can be implemented in hybrid access-distribution configuration that interconnects several multi-layer or routed-access enabled access-layer switches. Independent of route origination source (connected or dynamic route) and network size within the access-distribution block, the distribution router in main and remote site network must advertise a single, concise and summarized Layer 3 network to each Layer 3 access-switch and to core devices.
- *Remote Network*—Summarized network will be propagated dynamically across the network. Single summarization of all remote networks may be advertised to local Layer 3 access-switches, since it improves bandwidth efficiency. During a network outage, Layer 3 access-switch may drop traffic at the network edge instead of transmitting it to the distribution router to black hole traffic.
- *WAN Network*—Announcing a single summarized WAN network provides flexibility to troubleshoot and verify network availability.
- *Default Network*—When Layer 3 access-switch receives un-known destination traffic from the edge that does not match any of the above mentioned summarized networks, then it is sent to the distribution router to make a forwarding decision. The distribution router performs a forwarding table lookup and may forward to appropriate path or black hole the traffic. In a typical small enterprise network environment, a default route is announced by an Internet edge system, to forward all internet traffic. Distribution router must propagate this default route to the Layer 3 access-switch.

Figure 2-22 illustrates a summarized EIGRP network advertisement, by route aggregation system, that provides end-to-end internal and external network reachability.

Figure 2-22 End-to-End Routed-Access Network

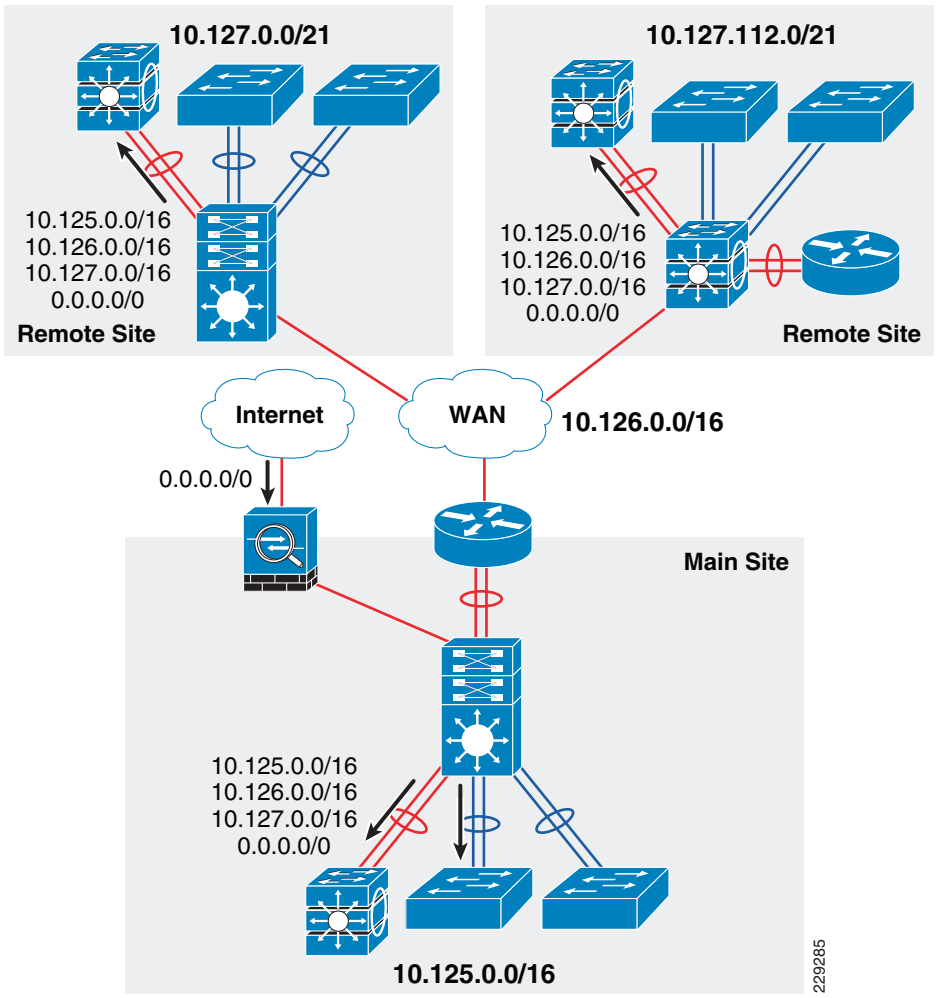

Following is configuration example to deploy summarized and filtered Layer-3 network information to Layer-3 access-switch.

Distribution

```
interface Port-channel13
 description Connected to cr24-3560r-1
 dampening
 ip address 10.125.32.0 255.255.255.254
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-key
 ip summary-address eigrp 100 10.125.0.0 255.255.0.0 5
 load-interval 30
 carrier-delay msec 0
!
!configure ACL and route-map to allow summarized route advertisement to Layer 3 access-
switch
!
access-list 1 permit 0.0.0.0
access-list 1 permit 10.126.0.0
access-list 1 permit 10.127.0.0
access-list 1 permit 10.125.0.0
```

```
!
route-map EIGRP_STUB_ROUTES permit 10
 match ip address 1
!
router eigrp 100
 distribute-list route-map EIGRP_STUB_ROUTES out Port-channel13


cr24-4507-1#show ip protocols | inc Outgoing|filtered
  Outgoing update filter list for all interfaces is not set
    Port-channel13 filtered by
```

Access

```
cr24-3560r-1#show ip route eigrp
     10.0.0.0/8 is variably subnetted, 15 subnets, 4 masks
D       10.126.0.0/16 [90/3328] via 10.125.32.0, 01:37:21, Port-channel1
D       10.127.0.0/16 [90/3584] via 10.125.32.0, 01:37:21, Port-channel1
D       10.125.0.0/16 [90/1792] via 10.125.32.0, 01:34:29, Port-channel1
D*EX 0.0.0.0/0 [170/515072] via 10.125.32.0, 00:03:15, Port-channel1
cr24-3560r-1#
```

EIGRP Adjacency Protection

EIGRP adjacency protection guidelines discussed earlier for the core network, apply equally to routed access in the access-distribution block. The two challenges, system efficiency, and network security also apply equally to the routed access design, and the same solution is applied:

- *System efficiency*—EIGRP hello transmission must be blocked on an interface where there are no trusted EIGRP neighbors, to reduce CPU utilization and prevent network attacks. EIGRP routing process should only be enabled on interfaces where trusted enterprise devices are connected. All other interfaces can be suppressed in passive mode.

  Following is the example configuration on Layer-3 access-switch that advertises networks enabled on SVI interfaces; however, keeps them in passive mode and explicitly allows EIGRP function on uplink port-channel to distribution router. Same configuration principle must be applied on each EIGRP router including distribution and core routers:

```
cr24-3560r-1(config)#router eigrp 100
cr24-3560r-1(config-router)# network 10.125.0.0 0.0.255.255
cr24-3560r-1(config-router)# passive-interface default
cr24-3560r-1(config-router)# no passive-interface Port-channel1
cr24-3560r-1#show ip eigrp interface
EIGRP-IPv4:(100) interfaces for process 100

                      Xmit           Queue   Mean   Pacing Time   Multicast Pending
Interface        Peers      Un/Reliable  SRTT   Un/Reliable   Flow Timer  Routes
Po1                  1        0/0         1       0/1           50           0

cr24-3560r-1#show ip protocols | inc Passive|Vlan
  Passive Interface(s):
    Vlan1
    Vlan11
    Vlan12
    Vlan13
    Vlan14
```

- Network Security—EIGRP adjacency between distribution and Layer-3 access-switch must be secured. Following is the example configuration to enable EIGRP neighbor authentication using MD5:

Distribution

```
cr24-4507-1(config)#key chain eigrp-key
cr24-4507-1(config-keychain)# key 1
cr24-4507-1(config-keychain-key)#   key-string <password>

cr24-4507-1(config)#interface Port-channel13
cr24-4507-1(config-if)# description Connected to cr24-3560r-1
cr24-4507-1(config-if)# ip authentication mode eigrp 100 md5
cr24-4507-1(config-if)# ip authentication key-chain eigrp 100 eigrp-key
```

Access

```
cr24-3560r-1(config)#key chain eigrp-key
cr24-3560r-1(config-keychain)# key 1
cr24-3560r-1(config-keychain-key)#   key-string <password>

cr24-3560r-1(config)#interface Port-channel1
cr24-3560r-1(config-if)# description Connected to cr24-4507-1
cr24-3560r-1(config-if)# ip authentication mode eigrp 100 md5
cr24-3560r-1(config-if)# ip authentication key-chain eigrp 100 eigrp-key
```

Tuning EIGRP Protocol Timers

EIGRP protocol functions the same in routed-access as it does in the core network. It is highly recommended to retain default EIGRP hello and hold timers on distribution and Layer 3 access-switch and rely on EtherChannel and SSO-based recovery mechanisms, that offers sub-second network convergence, during individual link or supervisor failure scenarios.

Deploying Multicast in Network

Communications in a IP network can be:

- *Unicast*—One source sends a message to one destination
- *Broadcast*—One source sends a message to all destinations
- *Multicast*—One source sends a message to a subset of destinations

IP multicast allows a source to transmit a message as a group transmission to a subset of hosts on the network. Many collaboration applications, such as video conferencing, distance learning, software distribution, utilize multicast techniques. IP multicast improves network bandwidth utilization, by reducing un necessary duplicate traffic. Multicast improves efficiency by reducing data processing on the source server, and sending a single flow into the network. Multicast packets are replicated in the network where paths diverge, by Protocol Independent Multicast (PIM)-enabled routers, and other supporting multicast protocols. See Figure 2-23.

Figure 2-23 Unicast versus Multicast Communication in Network

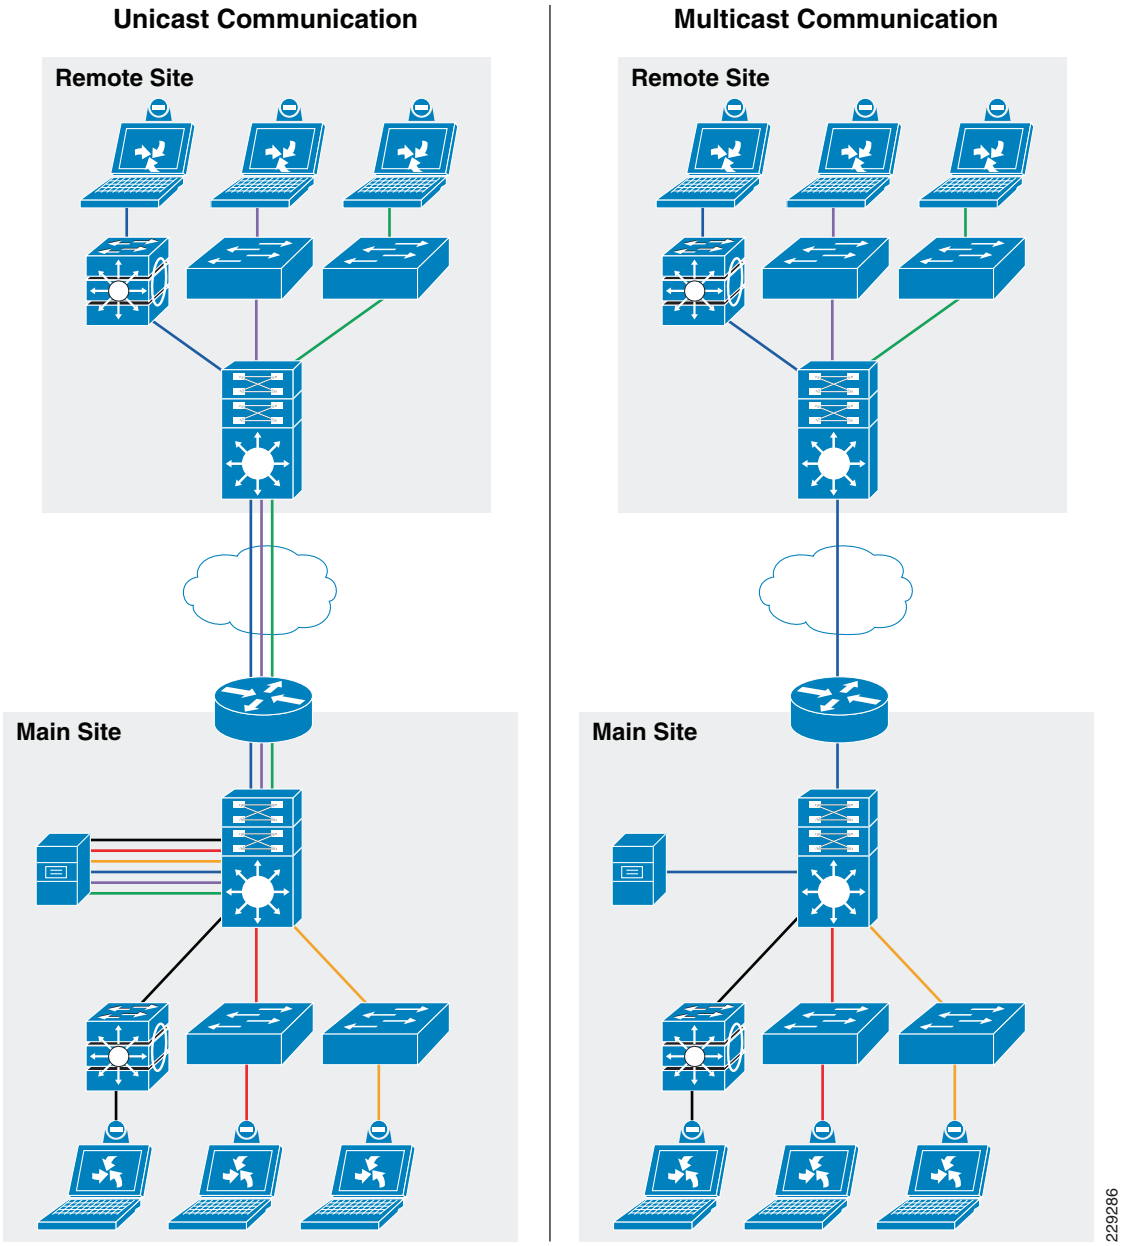

Multicast IP Addressing

The Internet Assigned Numbers Authority (IANA) controls the assignment of IP multicast addresses. A range of class D address space is assigned for IP multicast applications. All multicast group addresses fall in the range of 224.0.0.0 through 239.255.255.255. In IP multicast packets, the destination IP address is in the multicast group range, while the source IP address is always in the unicast address range. The multicast IP address space is further divided into several pools for well-known multicast network protocols, and inter-domain multicast communications as shown in Table 2-8.

Table 2-8 Multicast Address Range Assignments

Application	Address Range
Reserved – Link Local Network Protocols	224.0.0.0/24
Globally Scope – Group communication between organization and Internet	224.0.1.0 – 238.255.255.255
Source Specific Multicast (SSM) – PIM extension for one-to-many unidirectional multicast communication	232.0.0.0/8
GLOP – Inter-domain Multicast group assignment with reserved global Autonomous System (AS)	233.0.0.0/8
Limited Scope – Administratively scope address that remains constrained within local organization or AS. Commonly deployed in enterprise and other organization.	239.0.0.0/8

For the Schools SRA network design, the multicast IP addresses must be selected from the Limited Scope pool (239.0.0.0/8).

Multicast Routing Design

Each device between a multicast source and receiver must enable dynamic multicast. The technique for creating a multicast forwarding table is different than unicast routing and switching techniques. Multicast requires Multicast Routing Protocol (MRP) and Dynamic Group Membership (DGM) to enable communication.

Multicast Routing Protocol

IP multicast delivers source traffic to multiple receivers using the least amount of network resources, without placing additional burden on the source or the receivers. Multicast packet replication in the network is performed by Cisco routers and switches enabled with Protocol Independent Multicast (PIM) and other multicast routing protocols.

The network must build a packet distribution tree that specifies a unique forwarding path between the source subnet and each multicast group members subnet. A primary goal for the tree is to ensure that only one copy of each packet is forwarded on each branch of the tree. The two basic types of multicast distribution trees are source trees and shared trees:

- *Source trees*—The simplest form of a multicast distribution tree is a source tree, with the source at the root and the receivers at the branches. Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT).
- *Shared trees*—A shared tree uses a single common root placed at a chosen point in the network. This shared root is called a Rendezvous Point (RP).

PIM protocol has two modes which support both types of multicast distribution trees:

- *Dense Mode*—This mode assumes that most routers in the network will distribute multicast traffic to each multicast group. PIM-DM builds distribution trees by initially flooding the entire network and then pruning back the small number of paths without receivers.
- *Sparse Mode*—This mode assumes that relatively few routers in the network will be involved in each multicast group. The hosts belonging to the group are usually widely dispersed, as would be the case for most multicast over the WAN. PIM-SM begins with an empty distribution tree and adds branches only as the result of explicit IGMP requests to join.

It is recommended to deploy multicast in PIM-SM in the small enterprise network design. All the recommended platforms in this design support PIM-SM mode on physical or logical (SVI and EtherChannel) interfaces.

Dynamic Group Membership

Multicast receiver registration and deletion is done via Internet Group Management Protocol (IGMP) signaling. IGMP operates between a multicast receiver in the access-layer and a collapsed core router at the distribution layer in the main or the remote site.

In a multi-layer design, the layer 3 boundary is at the distribution switch. Multi-layer access-switches do not run PIM, and therefor flood the traffic on all ports. This multi-layer access-switch limitation is solved by using IGMP snooping feature, which is enabled by default. Best practice is to not disable IGMP snooping feature.

In a routed-access network design, the Layer-3 boundary is at the access-layer and IGMP communication is between receiver and access-switch. Along with unicast routing protocol, PIM-SM must be enabled on the Layer 3 access-switch to communicate with RP in the network.

Figure 2-24 demonstrates multicast source and receiver registration procedure and how shared-tree is dynamically developed for multicast data delivery.

Figure 2-24 Multicast Source and Receiver Registration Procedure

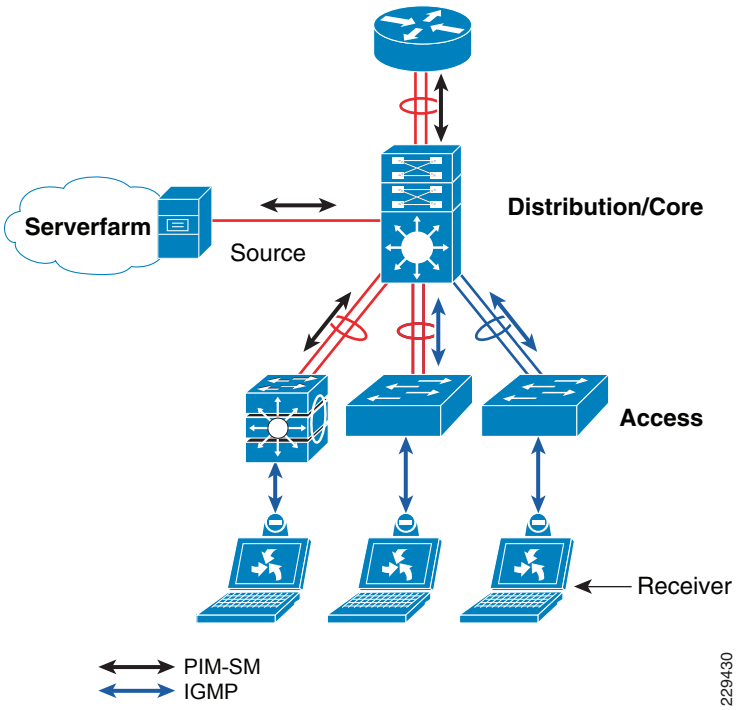

Deploying PIM-SM

PIM-SM distributes information about active sources by forwarding data packets on the shared tree. Because PIM-SM uses shared trees initially, it requires the use of a RP. It is recommended to deploy the RP close to the multicast source (collapsed core-distribution router in the main site is a good choice). Multicast sources centrally deployed in main site will register themselves with the RP and then data is forwarded down the shared tree to the receivers that could be located anywhere in the network.

PIM-SM Rendezvous Point

PIM-SM distributes information about active sources by forwarding data packets on the shared tree. Because PIM-SM uses shared trees initially, it requires the use of a RP. It is recommended to deploy the RP close to the multicast source (collapsed core-distribution router in the district office is a good choice). Multicast sources centrally deployed in district office will register themselves with the RP and then data is forwarded down the shared tree to the receivers that could be located anywhere in the network.

PIM-SM supports RP deployment in the following three different modes in the network:

- *Static*—As the name implies, RP must be statically identified and configured on each PIM router in the network. RP load-balancing and redundancy can be achieved using Anycast RP.
- *Auto-RP*—Dynamic method to discover and announce RP in the network. Auto-RP implementation is beneficial when there are multiple RPs and groups that often change in the network. To prevent network reconfiguration during change, RP mapping agent router must be designated in the network to receive RP group announcements and arbitrate conflicts. This capability is part of PIM version 1 specification.
- *BootStrap Router (BSR)*—Performs same task as Auto-RP but different mechanism. This capability is pare of PIM version 2 specification. Auto-RP and BSR cannot coexist or interoperate in the same network.

In a small to mid-size multicast network, static RP configuration is best overall, due primarily to the amount of administrative overhead that Auto-RP or BSR introduce. Static RP implementation offers same RP redundancy and load sharing and a simple ACL can be applied to deploy RP without compromising multicast network security. See Figure 2-25.

Figure 2-25 PIM-SM Network Design in Network Infrastructure

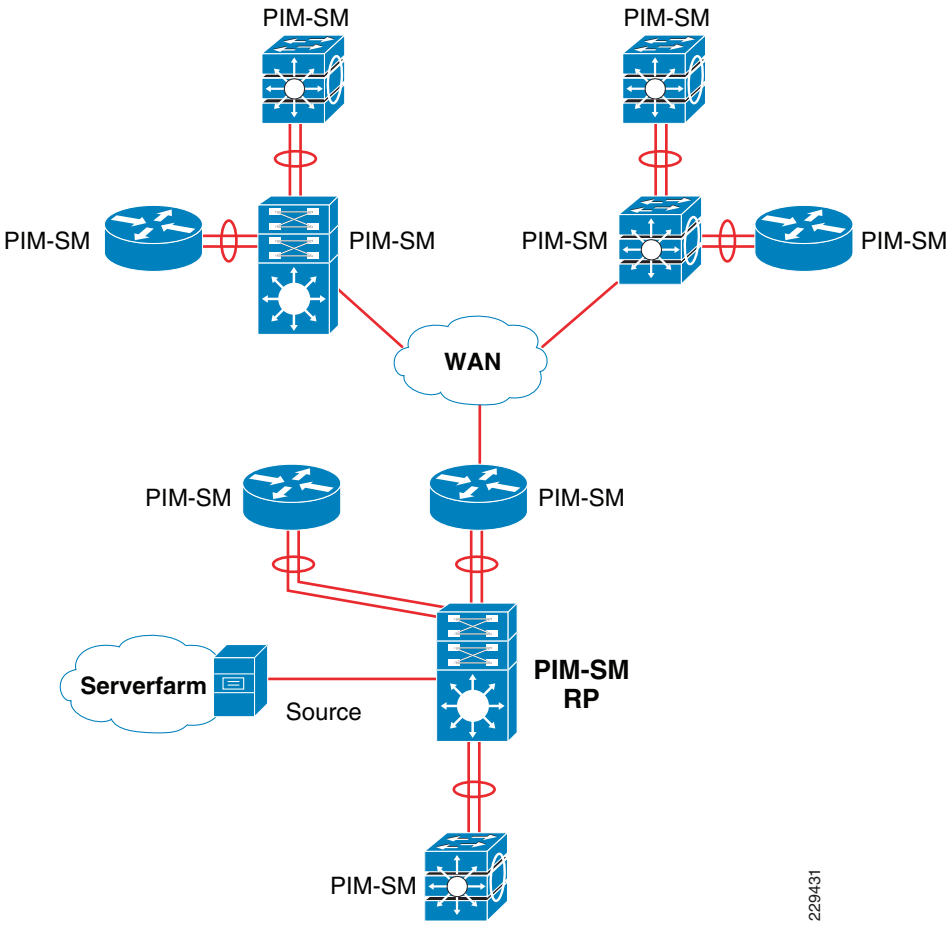

Following is an example configuration to deploy PIM-SM RP in the small enterprise network. Similar static PIM-SM configuration must be enabled on each Layer-3 PIM router or an access-switch in the remote sites:

Distribution - RP

```
cr24-4507-1(config)#interface Loopback1
cr24-4507-1(config-if)# description RP
cr24-4507-1(config-if)# ip address 10.125.100.100 255.255.255.255

cr24-4507-1(config)#ip multicast-routing
cr24-4507-1(config)#ip pim rp-address 10.125.100.100
```

Layer 3 Access

```
cr24-3560r-1(config)#ip multicast-routing distributed
cr24-3560r-1(config)#ip pim rp-address 10.125.100.100
```

Remote Site Core

```
cr36-3750s-1(config)#ip multicast-routing distributed
cr36-3750s-1(config)#ip pim rp-address 10.125.100.100
```

Upon successful PIM-SM RP implementation throughout the enterprise network, PIM-SM must be enabled on Layer-3 edge and core network-facing ports. The following sample configuration provides a simple PIM-SM implementation guideline to be implemented on every intermediate Layer-3 systems between receiver and source:

Distribution - RP

```
! Main Site - Access Network
cr24-4507-1(config)#interface range Vlan101 - 140
cr24-4507-1(config-if-range)# ip pim sparse-mode

! Main Site - Data Center Network
cr24-4507-1(config)#interface range Vlan141 - 150
cr24-4507-1(config-if-range)# ip pim sparse-mode

! Layer 3 Core and Routed-Access Port-Channel
cr24-4507-1(config)#interface range Port-channel 1, Port-channel 13, Port-channel 15
cr24-4507-1(config-if-range)# ip pim sparse-mode


cr24-4507-1#show ip pim interface
Address          Interface Ver/   Nbr    Query  DR     DR
                                   Mode   Count  Intvl  Prior
10.125.32.4     Port-channel1v2/S   1      30     1    10.125.32.4
<omitted>
10.125.1.1       Vlan101         v2/S 0      30     1      10.125.1.1

cr24-4507-1#show ip mroute sparse
(*, 239.192.51.8), 02:33:37/00:03:12, RP 10.125.100.100, flags: SJC
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan111, Forward/Sparse, 02:04:33/00:02:44, H
    Vlan101, Forward/Sparse, 02:04:59/00:02:58, H
    Port-channel15, Forward/Sparse, 02:04:59/00:02:47, H
    Vlan131, Forward/Sparse, 02:04:59/00:02:32, H
    Port-channel13, Forward/Sparse, 02:04:59/00:03:12, H
    Vlan121, Forward/Sparse, 02:04:59/00:02:14, H
    Vlan146, Forward/Sparse, 02:21:26/00:02:01, H
```

Layer 3 Access

```
! Main Site - Layer 3 Access Network
cr24-3560r-1(config)#interface range Vlan11 - 20
cr24-3560r-1(config-if-range)# ip pim sparse-mode

! Routed-Access Port-Channel
cr24-4507-1(config)#interface Port-channel 1
cr24-4507-1(config-if)# ip pim sparse-mode

cr24-3560r-1#show ip pim interface
Address          Interface Ver/   Nbr    Query  DR     DR
                                   Mode   Count  Intvl  Prior
10.125.32.1Port-channel1v2/S   1      30     1      10.125.32.0
10.125.11.1      Vlan11v2/S   0      30     1      10.125.11.1
```

Implementing IGMP

By default the Layer-2 access-switch will dynamically detect IGMP hosts and multicast-capable Layer-3 routers in the network. The IGMP snooping and multicast router detection functions on a per VLAN basis, and is globally enabled by default for all the VLANs. The IGMP configuration can be validated using the show command on the Layer-2 access-switch:

```
cr24-2960-1#show ip igmp snooping
Global IGMP Snooping configuration:
-------------------------------------------------
IGMP snooping               : Enabled
IGMPv3 snooping (minimal)    : Enabled
Report suppression           : Enabled
TCN solicit query            : Disabled
TCN flood query count        : 2
Robustness variable          : 2
Last member query count      : 2
Last member query interval   : 1000

cr24-2960-1#show ip igmp snooping mrouter
Vlan    ports
------    ------
 101    Po1(dynamic)
 102    Po1(dynamic)

cr24-2960-1#show ip igmp snooping group
Vlan      GroupType        Version     Port List
------------------------------------------------------------------------------
101       239.192.51.1igmp        v2     Fa0/1, Po1
101       239.192.51.2igmp        v2     Fa0/2, Po1
```

Multicast routing function changes when the access-switch is deployed in routed-access mode. PIM operation is performed at the access layer, therefore multicast router detection process is eliminated. The following output from a Layer-3 switch verifies that the local multicast ports are in router mode, and provide a snooped Layer-2 uplink port-channel which is connected to the collapsed core router, for multicast routing:

```
cr24-3560r-1#show ip igmp snooping mrouter
Vlan    ports
------    ------
  11    Router
  12    Router

cr24-3560r-1#show ip igmp membership | inc Channel|Vl
 Channel/Group                 Reporter        Uptime       Exp.FlagsInterface
 *,239.192.51.8                10.125.11.2000:17:52  02:45 2A     Vl11
 *,239.192.51.9                10.125.11.13100:17:52  02:43 2A     Vl12
```

Multicast Security—Preventing Rogue Source

This section provides basic multicast security configuration guidelines to prevent an unauthorized host in the network from acting like a rogue source in the network and sending multicast traffic.

In a PIM-SM network, an unwanted traffic source can be controlled with the pim accept-register command. When the source traffic hits the first-hop router, the first-hop router (DR) creates (S,G) state and sends a PIM Source Register message to the RP. If the source is not listed in the accept-register filter list (configured on the RP), then the RP rejects the Register and sends back an immediate Register-Stop

message to the DR. The drawback with this method of source-filtering is that the pim accept-register command on the RP, PIM-SM (S,G) state is still created on the source's first-hop router. This can result in traffic reaching receivers local to the source and located between the source and the RP. Furthermore, the pim accept-register command works on the control plane of the RP, which could be used to overload the RP with “fake” register messages, and possibly cause a DoS condition.

Best practice is to apply the pim accept-register command on the RP in addition to other edge-filtering methods, such as simple data plane ACLs on all DRs and on all ingress points into the network. While ingress ACLs on the DR are sufficient in a perfectly configured and operated network, best practice includes configuring the pim accept-register command on the RP in the main site as a secondary security mechanism in case of misconfiguration on the edge routers.

Following is the sample configuration with a simple ACL which has been applied to the RP to filter only on the source address. It is also possible to filter the source and the group with the use of an extended ACL on the RP:

Distribution-RP

```
cr24-4507-1(config)#ip access-list extended PERMIT-SOURCES
cr24-4507-1(config-ext-nacl)# permit ip 10.125.31.80 0.0.0.15 239.192.0.0 0.0.255.255

cr24-4507-1(config)#ip pim accept-register list PERMIT-SOURCES
```

Multicast Security—Preventing Rogue RP

Any router can be misconfigured or maliciously advertise itself as a multicast RP in the network. with the valid multicast group address. With a static RP configuration, each PIM-enabled router in the network can be configured to use the static RP for the multicast source and ignore any Auto-RP or BSR multicast router announcement.

Following is the sample configuration that must be applied to each PIM-enabled router in the main and remote sites, to accept PIM announcements only from the static RP and ignore dynamic multicast group announcement from any other RP:

Distribution-RP

```
cr24-4507-1(config)#ip access-list standard Allowed_MCAST_Groups
cr24-4507-1(config-std-nacl)# permit 224.0.1.39
cr24-4507-1(config-std-nacl)# permit 224.0.1.40
cr24-4507-1(config-std-nacl)# permit 239.192.0.0 0.0.255.255

cr24-4507-1(config)#ip pim rp-address 10.125.100.100 Allowed_MCAST_Groups override

cr24-4507-1#show ip pim rp mapping
PIM Group-to-RP Mappings
Acl: Allowed_MCAST_Groups, Static-Override
    RP: 10.125.100.100 (?)
```

Deploying QoS in Network

IP networks forward traffic on a best-effort basis by default. The routing protocol forwards packets over the best path, but offers no guarantee of delivery. This model works well for TCP-based data applications that adapt gracefully to variations in latency, jitter, and loss. The Small Enterprise Design Profile is a multi-service network design which supports voice and video as well as data traffic on a single network. Real time applications (such as voice, video) require packets delivered with in specified loss, delay and jitter parameters. Quality-of-Service (QoS) is a collection of features which allows the network to dedicate network resources for higher priority real time applications, while reserving sufficient network resources to service lower priority traffic. QoS accomplishes this by providing differentiated services, depending on the traffic type. For a detailed discussion of QoS, refer to the Enterprise QoS SRND at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html

While design principles are common, QoS implementation varies between fixed-configuration switches and the modular switching platforms like the Cisco Catalyst 4500-E/6500-E. This section discusses the internal switching architecture and the differentiated QoS structure on a per-hop-basis.

QoS in Catalyst Fixed Configuration Switches

The QoS implementation in Cisco Catalyst 2960, 2960S, 3560-X and 3750-X Series switches is similar. There is no difference in ingress or egress packet classification, marking, queuing and scheduling implementation among these Catalyst platforms. The Cisco Catalyst switches allow users to create a policy-map by classifying incoming traffic (Layer 2 to Layer 4). Catalyst switches allow attaching the policy-map to an individual physical port or to logical interfaces (SVI or port-channel). This creates a common QoS policy which may be used in multiple networks. To prevent switch fabric and egress physical port congestion, the ingress QoS policing structure can strictly filter excessive traffic at the network edge. All ingress traffic from edge ports passes through the switch fabric and congestion may occur at the egress ports. Congestion in access-layer switch can be prevented by tuning queuing scheduler and Weighted Tail Drop (WTD) drop parameters. See Figure 2-26.

Figure 2-26 Fixed Configuration Catalyst QoS Architecture

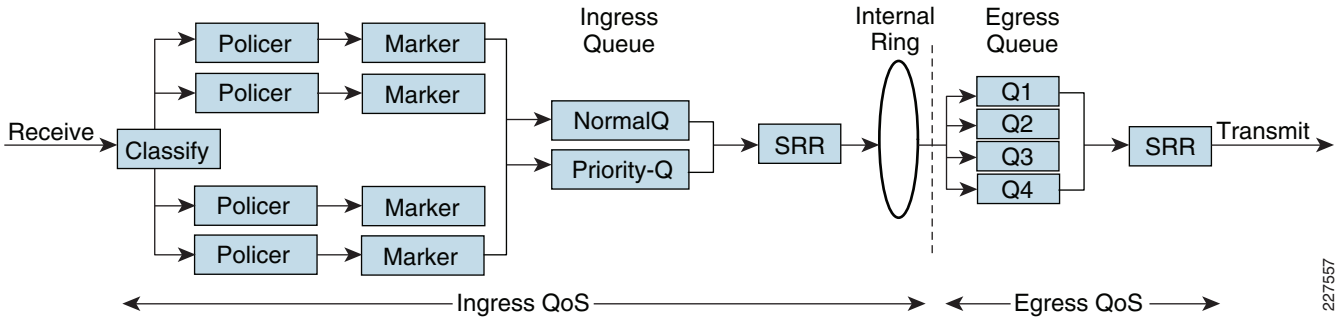

The main difference between these platforms is the switching capacity which ranges from 1G to 10G. The switching architecture and some of the internal QoS structure differs between these switches also. Following are some important differences to consider when selecting the access switch:

- The Cisco Catalyst 2960 does not support multilayer switching and does not support per-VLAN or per-port/per-VLAN policies.

- The Cisco Catalyst 2960 can police to a minimum rate of 1 Mbps; all other switches including next-generation Cisco Catalyst 2960-S Series within this product family can police to a minimum rate of 8 kbps.
- Only the Cisco Catalyst 3560-X and 3750-X support IPv6 QoS.
- Only the Cisco Catalyst 3560-X and 3750-X support policing on 10-Gigabit Ethernet interfaces.
- Only the Cisco Catalyst 3560-X and 3750-X support SRR shaping weights on 10-Gigabit Ethernet interfaces.

The next-generation Cisco Catalyst 2960-S Series platform introduces modified QoS architecture. To reduce the latency and improve application performance, the new Cisco 2960-S platform does not support ingress queueing and buffer function in harware. All other ingress and egress queuing, buffer and bandwidth sharing function remain consistent as the Cisco Catalyst 2960 platform. Each physical ports, including StackPort, have 2 MB buffer capacity to prevent traffic drop during congestion. This buffer allocation is static and cannot be modified by the user. However, when the Cisco Catalyst 2960-S is deployed in FlexStack configuration mode, there is a flexibility to assign different buffer size on egress queue of StackPort. Figure 2-27 illustrates QoS architecture on Catalyst 2960-S Series platform.

Figure 2-27 QoS Implementation in Catalyst 2960-S Switches

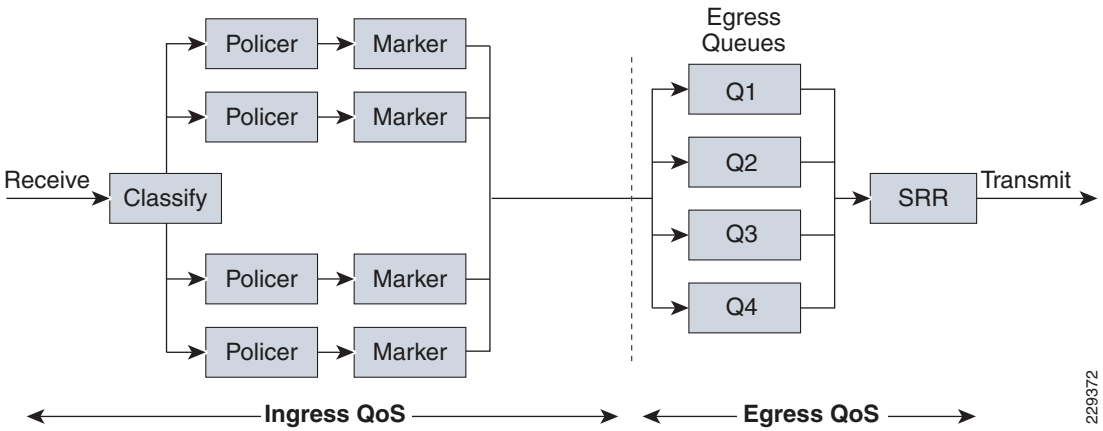

QoS in Cisco Modular Switches

Cisco Catalyst 4507R-E and 6500-E are high density, resilient switches for large scale networks. The Small Enterprise Design Profile uses the Cisco Catalyst 4507R-E in the main and larger remote site designs, hence all the QoS recommendations in this section will be based on 4500-E architecture. Cisco Catalyst 4500-E Series platform are widely deployed with classic and next-generation supervisors.

The classification function in the classic supervisor module is based on incoming DSCP or CoS setting in the pack, which was assigned by the access-layer switch. Catalyst 4500-E with classic supervisor performs ingress and egress QoS function based on internal mapping table that performs DSCP, ToS, or CoS interworking. Classic supervisor relies on trust model configuration; redirection of ingress traffic to an appropriate queue is based on the trust model defined on the edge port. See Figure 2-28.

Figure 2-28 Catalyst 4500-E– Classic Supervisor QoS Architecture

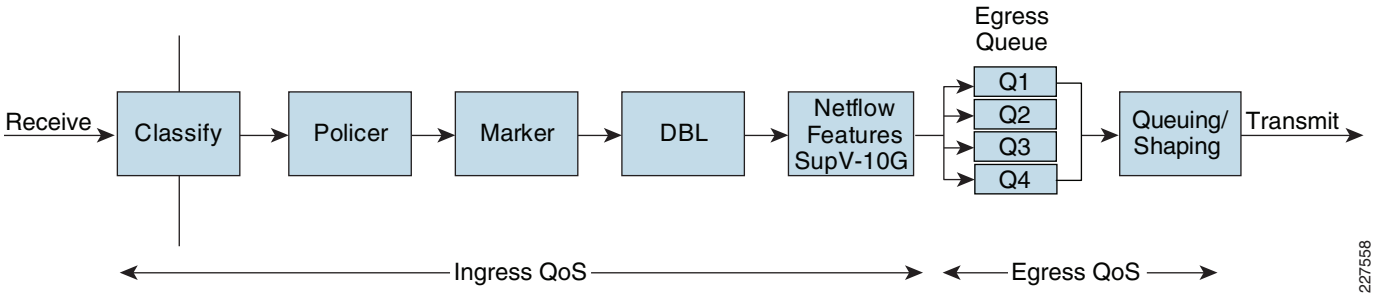

The Cisco Catalyst 4500-E with next generation Sup-6E (see Figure 2-29) is designed to offer better differentiated and preferential QoS services for various class-of-service traffic. New QoS capabilities in the Sup-6E enable administrators to take advantage of hardware-based intelligent classification and take action to optimize application performance and network availability. The QoS implementation in Sup-6E supports Modular QoS CLI (MQC) as implemented in IOS-based routers that overall enhances QoS capabilities and eases implementation and operations. Following are some of the key QoS features which differentiate the Sup-6E versus classic supervisors:

- *Trust and Table-Map*—MQC based QoS implementation offers a number of implementation and operational benefits over classic supervisors that rely on Trust model and internal Table-map as a tool to classify and mark ingress traffic.
- *Internal DSCP*—The queue placement in Sup-6E is simplified by leveraging the MQC capabilities to explicitly map DSCP or CoS traffic in hard-coded egress Queue structure,. For example, DSCP 46 can be classified with ACL and can be matched in PQ class-map of an MQC in Sup-6E.
- *Sequential vs Parallel Classification*—With MQC-based QoS classification, the Sup6-E provides sequential classification rather than parallel. Sequential classification method allows the network administrator to classify traffic at egress based on the ingress markings.

Figure 2-29 Catalyst 4500-E—Supervisor 6-E QoS Architecture

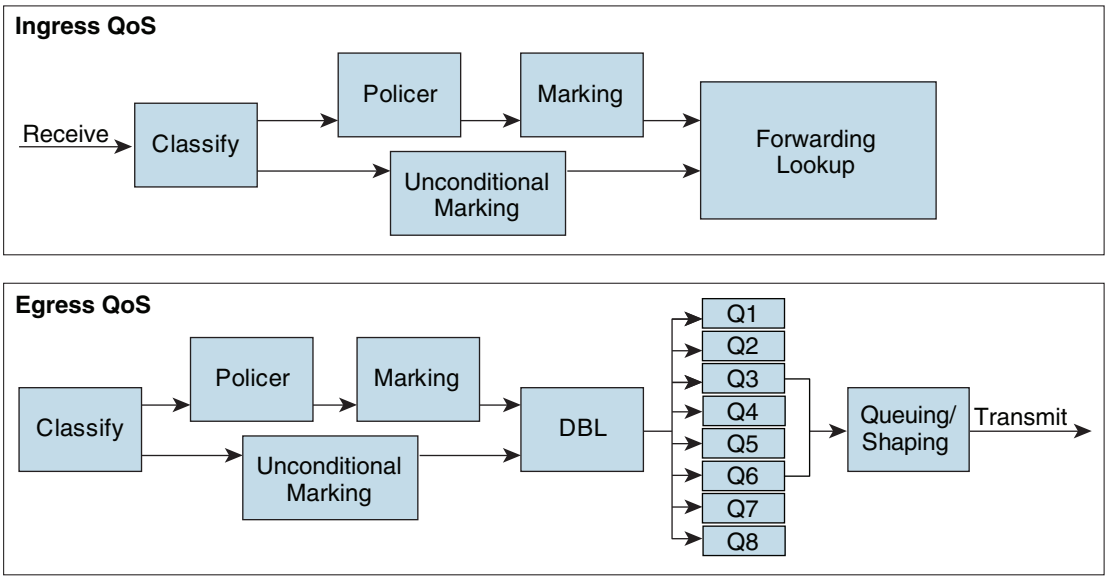

QoS Framework

QoS needs to be designed and implemented considering the entire network. This includes defining trust points, and determining which policies to enforce at each device within the network. Developing the trust model, guides policy implementations for each device.

Figure 2-30 depicts QoS trust model that guides QoS policy implementation in the main and remote site networks.

Figure 2-30 Small Enterprise Network QoS Framework

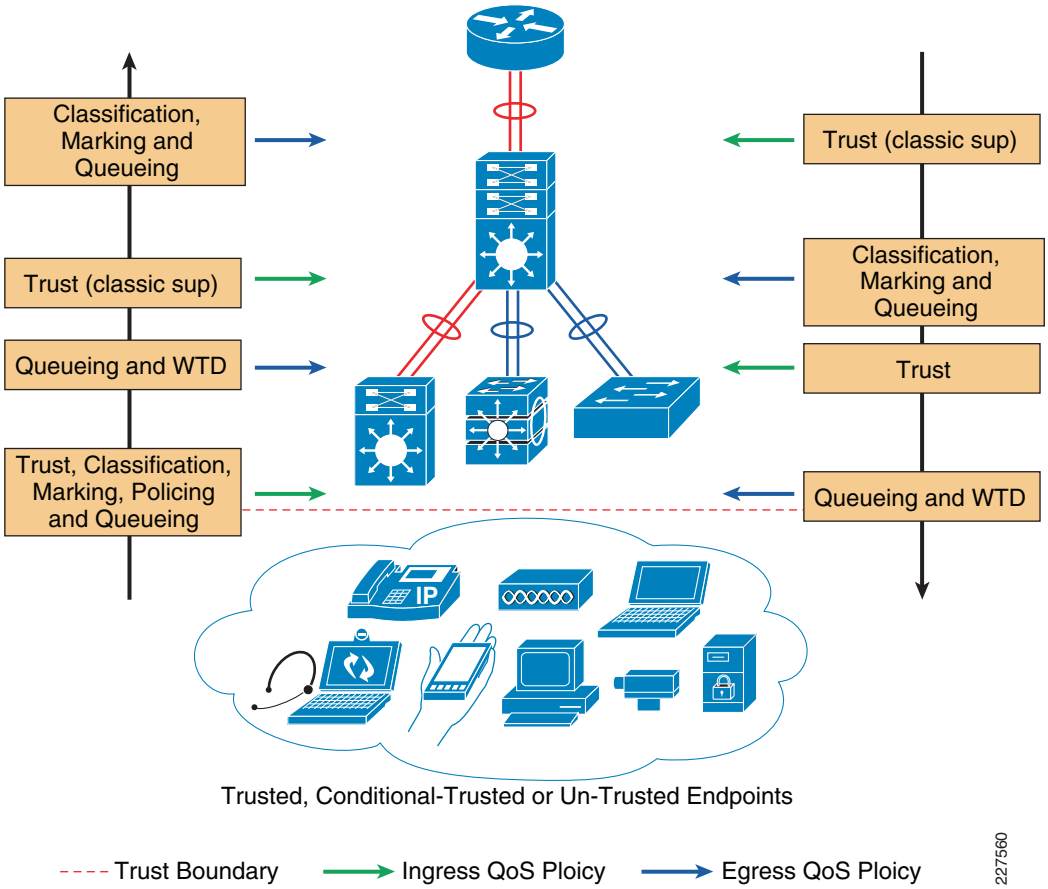

The devices (routers, switches) within the internal network are managed by the system administrator, and hence are classified as trusted devices. Access-layer switches communicate with devices that are beyond the network boundary and within the internal network domain. QoS trust boundary at the access-layer communicates with various devices that could be deployed in different trust models (Trusted, Conditional-Trusted, or Un-Trusted). This section discusses the QoS policies for the traffic that traverses access-switch QoS trust boundary. The QoS function is unidirectional; it provides flexibility to set different QoS polices for traffic entering the network versus traffic that is exiting the network. See Figure 2-31.

Figure 2-31 Small Enterprise Network Edge QoS Boundary

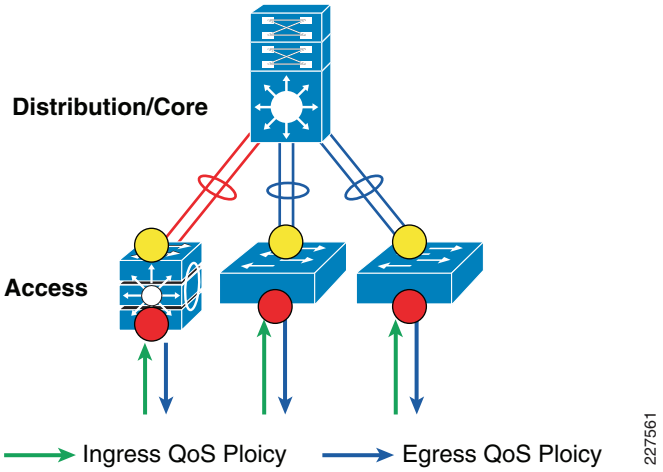

QoS Trust Boundary

The access-switch provides the entry point to the network for end devices. The access-switch must decide whether to accept the QoS markings from each endpoint, or whether to change them. This is determined by the QoS policies, and the trust model with which the endpoint is deployed.

End devices are classified into one of three different trust models; each with its own unique security and QoS policies to access the network:

- *Untrusted*—An unmanaged device that does not pass through the network security policies. For example, employee-owned PC or network printer. Packets with 802.1p or DSCP marking set by untrusted endpoints are reset to default by the access-layer switch at the edge. Otherwise, it is possible for an unsecured user to take away network bandwidth that may impact network availability and security for other users.
- *Trusted*—Devices that passes through network access security policies and are managed by network administrator. For example, secure PC or IP endpoints (i.e., servers, cameras, DMP, wireless access points, VoIP/video conferencing gateways, etc). Even when these devices are network administrator maintained and secured, QoS policies must still be enforced to classify traffic and assign it to the appropriate queue to provide bandwidth assurance and proper treatment during network congestion.
- *Conditionally-Trusted*—A single physical connection with one trusted endpoint and a indirect untrusted endpoint must be deployed as conditionally-trusted model. The trusted endpoints are still managed by the network administrator, but it is possible that the untrusted user behind the endpoint may or may not be secure. For example, Cisco Unified IP Phone + PC. These deployment scenarios require hybrid QoS policy that intelligently distinguishes and applies different QoS policy to the trusted and untrusted endpoints that are connected to the same port.

Deploying QoS

The ingress QoS policy at the access-switches needs to be established, since this is the trust boundary, where traffic enters the network. The following ingress QoS techniques are applied to provide appropriate service treatment and prevent network congestion:

- *Trust*—After classifying the endpoint the trust settings must be explicitly set by a network administrator. By default, Catalyst switches set each port in untrusted mode when QoS is enabled.

- *Classification*—IETF standard has defined a set of application classes and provides recommended DSCP settings. This classification determines the priority the traffic will receive in the network. Using the IETF standard, simplifies the classification process and improves application and network performance.
- *Policing*—To prevent network congestion, the access-layer switch limits the amount of inbound traffic up to its maximum setting. Additional policing can be applied for known applications, to ensure the bandwidth of an egress queue is not completely consumed by one application.
- *Marking*—Based on trust model, classification, and policer settings the QoS marking is set at the edge before approved traffic enters through the access-layer switching fabric. Marking traffic with the appropriate DSCP value is important to ensure traffic is mapped to the appropriate internal queue, and treated with the appropriate priority.
- *Queueing*—To provide differentiated services internally in the Catalyst switching fabric, all approved traffic is queued into priority or non-priority ingress queue. Ingress queueing architecture assures real-time applications, like VoIP traffic, are given appropriate priority (eg transmitted before data traffic).

Implementing QoS Trust Mode

By default, QoS is disabled on all Catalyst switches and must be explicitly enabled in global configuration mode. The QoS configuration is the same for a multilayer or routed-access deployment. The following sample QoS configuration must be enabled on all the access-layer switches deployed in main and remote sites.

```
cr24-2960-1(config)#mls qos
cr24-2960-1#show mls qos
QoS is enabled
QoS ip packet dscp rewrite is enabled
```

Upon enabling QoS in the Catalyst switches, all physical ports are assigned untrusted mode. The network administrator must explicitly enable the trust settings on the physical port where trusted or conditionally trusted endpoints are connected. The Catalyst switches can trust the ingress packets based on 802.1P (CoS-based), ToS (ip-prec-based) or DSCP (DSCP-based) values. Best practice is to deploy DSCP-based trust mode on all the trusted and conditionally-trusted endpoints. This offers a higher level of classification and marking granularity than other methods. The following sample DSCP-based trust configuration must be enabled on the access-switch ports connecting to trusted or conditionally-trusted endpoints.

Access (Multilayer or Routed-Access)

Trusted Port

```
cr24-2960-1(config)#interface FastEthernet0/5
cr24-2960-1(config-if)# description CONNECTED TO IPVS 2500 - CAMERA
cr24-2960-1(config-if)# mls qos trust dscp

cr24-2960-1#show mls qos interface f0/5
FastEthernet0/5
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based
```

Conditionally-Trusted Port

```
cr24-2960-1(config)#interface FastEthernet0/3
cr24-2960-1(config-if)# description CONNECTED TO PHONE
cr24-2960-1(config-if)# mls qos trust device cisco-phone
cr24-2960-1(config-if)# mls qos trust dscp

cr24-2960-1#show mls qos interface f0/3
FastEthernet0/3
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: cisco-phone
qos mode: port-based
```

UnTrusted Port

As described earlier, the default trust mode is *untrusted* when globally enabling QoS function. Without explicit trust configuration on Fas0/1 port, the following **show** command verifies current trust state and mode:

```
cr24-2960-1#show mls qos interface f0/1
FastEthernet0/1
trust state: not trusted
trust mode: not trusted
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based
```

Implementing QoS Classification

When creating QoS classification policies, the network administrator needs to consider what applications are present at the access edge (in the ingress direction) and whether these applications are sourced from trusted or untrusted endpoints. If PC endpoints are secured and centrally administered, then endpoint PCs may be considered trusted endpoints. In most deployments, this is not the case, thus PCs are considered untrusted endpoints for the remainder of this document.

Not every application class, as defined in the Cisco-modified RFC 4594-based model, is present in the ingress direction at the access edge; therefore, it is not necessary to provision the following application classes at the access-layer:

- *Network Control*—It is assumed that access-layer switch will not transmit or receive network control traffic from endpoints; hence this class is not implemented.
- *Broadcast Video*—Broadcast video and multimedia streaming server are centrally deployed at the main and multicast traffic is originated from trusted serverfarm servers and is unidirectional to remote site endpoints (and should not be sourced from remote site endpoints).
- *Operation, Administration and Management*—Primarily generated by network devices (routers, switches) and collected by management stations which are typically deployed in the trusted serverfarm network, or a network control center.

All applications present at the access edge need to be assigned a classification, as shown in Figure 34. Voice traffic is primarily sourced from Cisco IP telephony devices residing in the voice VLAN (VVLAN). These are trusted devices, or conditionally trusted, if users also attach PC's, etc to the same port. Voice communication may also be sourced from PC's with soft-phone applications, like Cisco Unified Personal Communicator (CUPC). Since such applications share the same UDP port range as multimedia conferencing traffic (UDP/RTP ports 16384-32767) this soft-phone VoIP traffic is indistinguishable, and should be classified with multimedia conferencing streams. See Figure 2-32.

Figure 2-32 QoS Classes

Application	PHB	Application Examples	Present at Campus Access-Edge (Ingress)?	Trust Boundary
Network Control	CS6	EIGRP, OSPF, HSRP, IKE		
VoIP	EF	Cisco IP Phone	Yes	Trusted
Broadcast Video		Cisco IPVS, Enterprise TV		
Realtime Interactive	CS4	Cisco TelePresence	Yes	Trusted
Multimedia Conferencing	AF4	Cisco CUPC, WebEx	Yes	Untrusted
Multimedia Streaming	AF3	Cisco DMS, IP/TV		
Signaling	CS3	SCCP, SIP, H.323	Yes	Trusted
Transactional Data	AF2	ERP Apps, CRM Apps	Yes	Untrusted
OAM	CS2	SNMP, SSH, Syslog		
Bulk Data	AF1	Email, FTP, Backup	Yes	Untrusted
Best Effort	DF	Default Class	Yes	Untrusted
Scavenger	CS1	YouTube, Gaming, P2P	Yes	Untrusted

227562

MQC offers scalability and flexibility in configuring QoS to classify all 8 application classes by using match statements or an extended access-list to match the exact value or range of Layer-4 known ports that each application uses to communicate on the network. The following sample configuration creates an extended access-list for each application and then applies it under class-map configuration mode.

```
cr24-3560r-1(config)#ip access-list extended MULTIMEDIA-CONFERENCING
cr24-3560r-1(config-ext-nacl)# remark RTP
cr24-3560r-1(config-ext-nacl)# permit udp any any range 16384 32767
cr24-3560r-1(config-ext-nacl)#!
cr24-3560r-1(config-ext-nacl)#ip access-list extended SIGNALING
cr24-3560r-1(config-ext-nacl)# remark SCCP
cr24-3560r-1(config-ext-nacl)# permit tcp any any range 2000 2002
cr24-3560r-1(config-ext-nacl)# remark SIP
cr24-3560r-1(config-ext-nacl)# permit tcp any any range 5060 5061
cr24-3560r-1(config-ext-nacl)# permit udp any any range 5060 5061
cr24-3560r-1(config-ext-nacl)#
cr24-3560r-1(config-ext-nacl)#ip access-list extended TRANSACTIONAL-DATA
cr24-3560r-1(config-ext-nacl)# remark HTTPS
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 443
cr24-3560r-1(config-ext-nacl)# remark ORACLE-SQL*NET
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 1521
```

```
cr24-3560r-1(config-ext-nacl)# permit udp any any eq 1521
cr24-3560r-1(config-ext-nacl)# remark ORACLE
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 1526
cr24-3560r-1(config-ext-nacl)# permit udp any any eq 1526
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 1575
cr24-3560r-1(config-ext-nacl)# permit udp any any eq 1575
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 1630
cr24-3560r-1(config-ext-nacl)#
cr24-3560r-1(config-ext-nacl)#ip access-list extended BULK-DATA
cr24-3560r-1(config-ext-nacl)# remark FTP
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq ftp
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq ftp-data
cr24-3560r-1(config-ext-nacl)# remark SSH/SFTP
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 22
cr24-3560r-1(config-ext-nacl)# remark SMTP/SECURE SMTP
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq smtp
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 465
cr24-3560r-1(config-ext-nacl)# remark IMAP/SECURE IMAP
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 143
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 993
cr24-3560r-1(config-ext-nacl)# remark POP3/SECURE POP3
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq pop3
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 995
cr24-3560r-1(config-ext-nacl)# remark CONNECTED PC BACKUP
cr24-3560r-1(config-ext-nacl)# permit tcp any eq 1914 any
cr24-3560r-1(config-ext-nacl)#
cr24-3560r-1(config-ext-nacl)#ip access-list extended DEFAULT
cr24-3560r-1(config-ext-nacl)# remark EXPLICIT CLASS-DEFAULT
cr24-3560r-1(config-ext-nacl)# permit ip any any
cr24-3560r-1(config-ext-nacl)#
cr24-3560r-1(config-ext-nacl)#ip access-list extended SCAVENGER
cr24-3560r-1(config-ext-nacl)# remark KAZAA
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 1214
cr24-3560r-1(config-ext-nacl)# permit udp any any eq 1214
cr24-3560r-1(config-ext-nacl)# remark MICROSOFT DIRECT X GAMING
cr24-3560r-1(config-ext-nacl)# permit tcp any any range 2300 2400
cr24-3560r-1(config-ext-nacl)# permit udp any any range 2300 2400
cr24-3560r-1(config-ext-nacl)# remark APPLE ITUNES MUSIC SHARING
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 3689
cr24-3560r-1(config-ext-nacl)# permit udp any any eq 3689
cr24-3560r-1(config-ext-nacl)# remark BITTORRENT
cr24-3560r-1(config-ext-nacl)# permit tcp any any range 6881 6999
cr24-3560r-1(config-ext-nacl)# remark YAHOO GAMES
cr24-3560r-1(config-ext-nacl)# permit tcp any any eq 11999
cr24-3560r-1(config-ext-nacl)# remark MSN GAMING ZONE
cr24-3560r-1(config-ext-nacl)# permit tcp any any range 28800 29100
cr24-3560r-1(config-ext-nacl)#
```

Creating class-map for each application services and applying match statement:

```
cr24-3560r-1(config)#class-map match-all VVLAN-SIGNALING
cr24-3560r-1(config-cmap)# match ip dscp cs3
cr24-3560r-1(config-cmap)#
cr24-3560r-1(config-cmap)#class-map match-all VVLAN-VOIP
cr24-3560r-1(config-cmap)# match ip dscp ef
cr24-3560r-1(config-cmap)#
cr24-3560r-1(config-cmap)#class-map match-all MULTIMEDIA-CONFERENCING
cr24-3560r-1(config-cmap)# match access-group name MULTIMEDIA-CONFERENCING
cr24-3560r-1(config-cmap)#
cr24-3560r-1(config-cmap)#class-map match-all SIGNALING
cr24-3560r-1(config-cmap)# match access-group name SIGNALING
cr24-3560r-1(config-cmap)#
cr24-3560r-1(config-cmap)#class-map match-all TRANSACTIONAL-DATA
```

```
cr24-3560r-1(config-cmap)# match access-group name TRANSACTIONAL-DATA

cr24-3560r-1(config-cmap)#
cr24-3560r-1(config-cmap)#class-map match-all BULK-DATA
cr24-3560r-1(config-cmap)# match access-group name BULK-DATA
cr24-3560r-1(config-cmap)#
cr24-3560r-1(config-cmap)#class-map match-all DEFAULT
cr24-3560r-1(config-cmap)# match access-group name DEFAULT
cr24-3560r-1(config-cmap)#
cr24-3560r-1(config-cmap)#class-map match-all SCAVENGER
cr24-3560r-1(config-cmap)# match access-group name SCAVENGER
```

Implementing Ingress Policer

It is important to limit how much bandwidth each class may use at the ingress to the access-layer for two primary reasons:

- *Bandwidth Bottleneck*—To prevent network congestion, each physical port at trust boundary must be rate-limited. The rate-limit value may differ based on several factors—end-to-end network bandwidth capacity, end-station, and application performance capacities, etc.
- *Bandwidth Security*—Well-known applications like Cisco IP telephony, use a fixed amount of bandwidth per device, based on codec. It is important to police high-priority application traffic which is assigned to the high-priority queue, otherwise it could consume too much overall network bandwidth and impact other application performance.

In addition to policing, the rate-limit function also provides the ability to take different actions on the excess incoming traffic which exceeds the established limits. The exceed-action for each class must be carefully designed based on the nature of application to provide best effort service based on network bandwidth availability. Table 2-9 provides best practice policing guidelines for different classes to be implemented for trusted and conditional-trusted endpoints at the network edge.

Table 2-9 Best Practice Policing Guidelines

Application	Policing Rate	Conform-Action	Exceed-Action
VoIP Signaling	<32 kbps	Pass	Drop
VoIP Bearer	<128 kbps	Pass	Drop
Multimedia Conferencing	<5Mbps[1]	Pass	Drop
Signaling	<32 kbps	Pass	Drop
Transactional Data	<10 Mbps [1]	Pass	Remark to CS1
Bulk Data	<10 Mbps [1]	Pass	Remark to CS1
Best Effort	<10 Mbps [1]	Pass	Remark to CS1
Scavenger	<10 Mbps [1]	Pass	Drop

1. Rate varies based on several factors as defined earlier. This table depicts sample rate-limiting values.

As described in the "QoS in Catalyst Fixed Configuration Switches" section on page 2-51, the policer capabilities differ in Cisco Catalyst switching platforms. When deploying policer policies on the access-layer switches the following platform limitations must be taken into consideration:

- The Catalyst 2960 can only police to a minimum rate of 1 Mbps; all other platforms, including next-generation Cisco Catalyst 2960-S Series, within this switch-product family can police to a minimum rate of 8 kbps.

- Only the Cisco Catalyst 3560-X and 3750-X support policing on 10 Gigabit Ethernet interfaces.

The following sample configuration shows how to deploy policing for multiple classes on trusted and conditionally-trusted ingress ports in access-layer switches.

Trusted or Conditionally-Trusted Port

```
cr24-3560r-1(config)#policy-map Phone+PC-Policy
cr24-3560r-1(config-pmap)# class VVLAN-VOIP
cr24-3560r-1(config-pmap-c)# police 128000 8000 exceed-action drop
cr24-3560r-1(config-pmap-c)# class VVLAN-SIGNALING
cr24-3560r-1(config-pmap-c)# police 32000 8000 exceed-action drop
cr24-3560r-1(config-pmap-c)# class MULTIMEDIA-CONFERENCING
cr24-3560r-1(config-pmap-c)# police 5000000 8000 exceed-action drop
cr24-3560r-1(config-pmap-c)# class SIGNALING
cr24-3560r-1(config-pmap-c)# police 32000 8000 exceed-action drop
cr24-3560r-1(config-pmap-c)# class TRANSACTIONAL-DATA
cr24-3560r-1(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
cr24-3560r-1(config-pmap-c)# class BULK-DATA
cr24-3560r-1(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
cr24-3560r-1(config-pmap-c)# class SCAVENGER
cr24-3560r-1(config-pmap-c)# police 10000000 8000 exceed-action drop
cr24-3560r-1(config-pmap-c)# class DEFAULT
cr24-3560r-1(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
```

All ingress traffic (default class) from untrusted endpoint be must be policed without explicit classification that requires differentiated services. The following sample configuration shows how to deploy policing on untrusted ingress ports in access-layer switches:

UnTrusted Port

```
cr24-3560r-1(config)#policy-map UnTrusted-PC-Policy
cr24-3560r-1(config-pmap)# class class-default
cr24-3560r-1(config-pmap-c)#  police 10000000 8000 exceed-action drop
```

Implementing Ingress Marking

Accurate DSCP marking of ingress traffic at the access-layer switch is critical to ensure proper QoS service treatment as traffic traverses through the network. All classified and policed traffic must be explicitly marked using the policy-map configuration based on an 8-class QoS model as shown in Figure 2-32.

Best practice is to use a explicit marking command (**set dscp**) even for trusted application classes (like VVLAN-VOIP and VVLAN-SIGNALING), rather than a trust policy-map action. A trust statement in a policy map requires multiple hardware entries, while the use of an explicit (seemingly redundant) marking command, improves the hardware efficiency.

The following sample configuration shows how to implement explicit marking for multiple classes on trusted and conditionally-trusted ingress ports in access-layer switches:

Trusted or Conditionally-Trusted Port

```
cr24-3560r-1(config)#policy-map Phone+PC-Policy
cr24-3560r-1(config-pmap)# class VVLAN-VOIP
cr24-3560r-1(config-pmap-c)# set dscp ef
cr24-3560r-1(config-pmap-c)# class VVLAN-SIGNALING
cr24-3560r-1(config-pmap-c)# set dscp cs3
cr24-3560r-1(config-pmap-c)# class MULTIMEDIA-CONFERENCING
cr24-3560r-1(config-pmap-c)# set dscp af41
```

```
cr24-3560r-1(config-pmap-c)# class SIGNALING
cr24-3560r-1(config-pmap-c)# set dscp cs3
cr24-3560r-1(config-pmap-c)# class TRANSACTIONAL-DATA
cr24-3560r-1(config-pmap-c)# set dscp af21
cr24-3560r-1(config-pmap-c)# class BULK-DATA
cr24-3560r-1(config-pmap-c)# set dscp af11
cr24-3560r-1(config-pmap-c)# class SCAVENGER
cr24-3560r-1(config-pmap-c)# set dscp cs1
cr24-3560r-1(config-pmap-c)# class DEFAULT
cr24-3560r-1(config-pmap-c)# set dscp default
```

All ingress traffic (default class) from an untrusted endpoint must be marked without a explicit classification. The following sample configuration shows how to implement explicit DSCP marking:

Untrusted Port

```
cr24-3560r-1(config)#policy-map UnTrusted-PC-Policy
cr24-3560r-1(config-pmap)# class class-default
cr24-3560r-1(config-pmap-c)#  set dscp default
```

Applying Ingress Policies

After creating a complete policy-map with all the QoS policies defined, the service-policy must be applied on the edge interface of the access-layer to enforce the QoS configuration. Cisco Catalyst switches offer three simplified methods to apply service-policies. Depending on the deployment model, any of these methods may be used:

- *Port-based QoS*—Applying service-policy on a per physical port basis will force traffic to pass-through the QoS policies before entering the network. Port-based QoS functions on a per-physical port basis even if the port is associated with a logical VLAN.
- *VLAN-based QoS*—Applying service-policy on per VLAN basis requires the policy-map to be attached to a logical Layer-3 SVI interface. Every physical port associated with the VLAN will require an extra configuration to enforce the QoS policies defined on a logical interface.
- *Per-Port/Per-VLAN-based QoS*—Not supported on all the Catalyst platforms and the configuration commands are platform-specific. Per-port/per-VLAN-based QoS creates a nested hierarchical policy-map that operates on a trunk interface. A different policy-map can be applied on each logical SVI interface that is associated to a single physical port.

Figure 2-33 Depicts All Three QoS Implementation Method

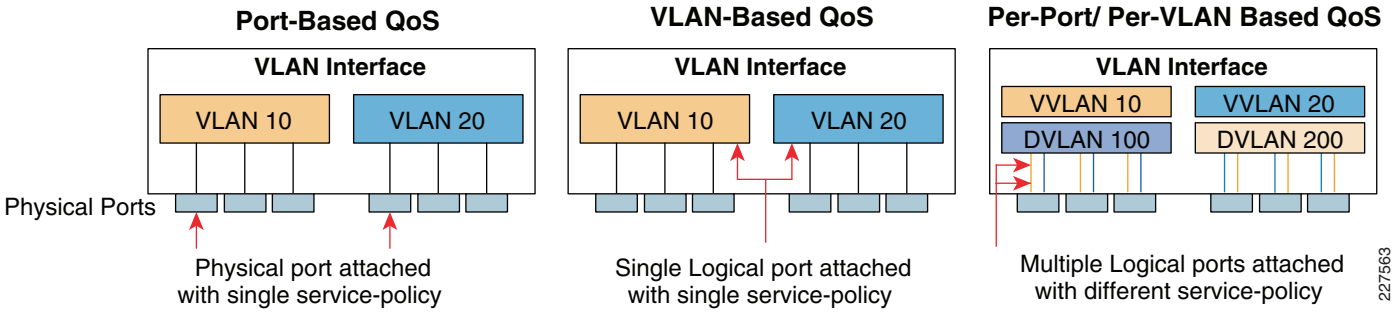

The following sample configuration shows how to deploy port-based QoS on the access-layer switches:

```
cr24-3560r-1(config)#interface fastethernet0/4
cr24-3560r-1(config-if)# description CONNECTED TO PHONE+PC
```

```
    cr24-3560r-1(config-if)# service-policy input Phone+PC-Policy

cr24-3560r-1#show policy-map interface f0/4 | inc Service|Class
Service-policy input: Phone+PC-Policy
Class-map: VVLAN-VOIP (match-all)
Class-map: VVLAN-SIGNALING (match-all)
Class-map: MULTIMEDIA-CONFERENCING (match-all)
Class-map: SIGNALING (match-all)
Class-map: TRANSACTIONAL-DATA (match-all)
Class-map: BULK-DATA (match-all)
Class-map: SCAVENGER (match-all)
 Class-map: DEFAULT (match-all)
    Class-map: class-default (match-any)
```

Applying Ingress Queueing

Fixed configuration Cisco Catalyst switches (2960 and 3xxx) not only offer differentiated services on the network ports but also internally on the switching fabric. Note, Cisco Catalyst 2960-S Series platform do not support ingress queueing and buffer allocation. After enabling QoS and attaching inbound policies on the physical ports, all the packets that meet the specified policy are forwarded to the switching fabric for egress switching. The aggregate bandwidth from all edge ports may exceed switching fabric bandwidth and cause internal congestion.

Cisco Catalyst 2960 and 3xxx platforms support two internal ingress queues: normal queue and priority queue. The ingress queue inspects the DSCP value on each incoming frame and assigns it to either the normal or priority queue. High priority traffic, like DSCP EF marked packets, are placed in the priority queue and switched before processing the normal queue.

The Catalyst 3750-X family of switches supports the weighted tail drop (WTD) congestion avoidance mechanism. WTD is implemented on queues to manage the queue length. WTD drops packets from the queue, based on dscp value, and the associated threshold. If the threshold is exceeded for a given internal DSCP value, the switch drops the packet. Each queue has three threshold values. The internal DSCP determines which of the three threshold values is applied to the frame. Two of the three thresholds are configurable (explicit) and one is not (implicit). This last threshold corresponds to the tail of the queue (100 percent limit).

Figure 2-34 depicts how different class-of-service applications are mapped to the Ingress Queue structure (1P1Q3T) and how each queue is assigned a different WTD threshold.

Figure 2-34 Catalyst 2960 and 3xxx Ingress Queueing Model

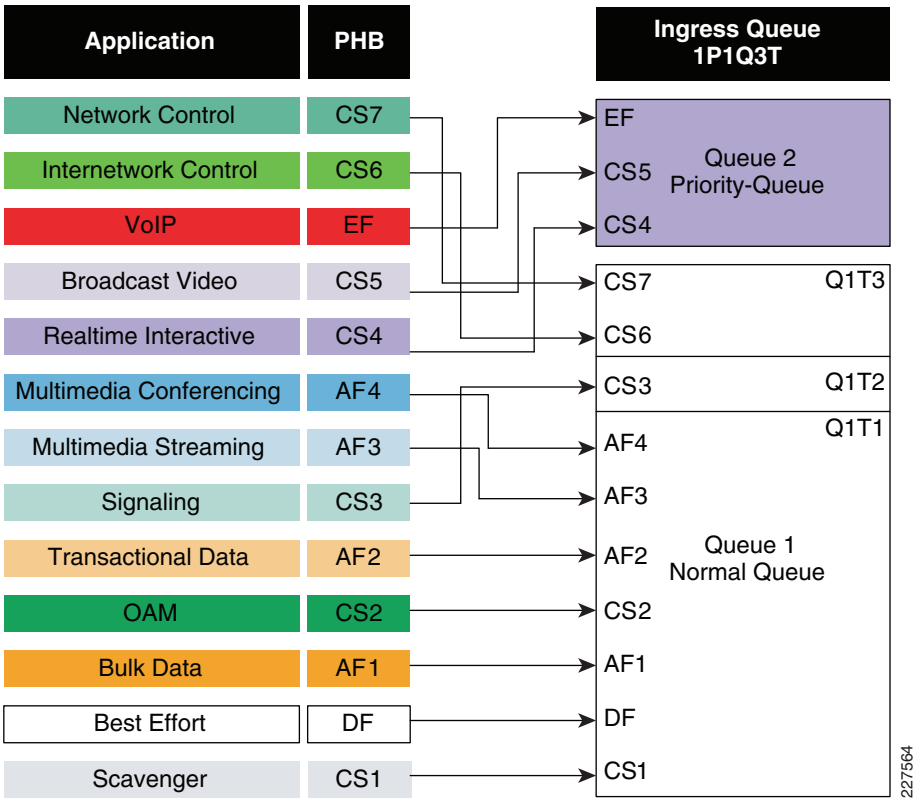

The DSCP marked packets in the policy-map must be assigned to the appropriate queue and each queue must be configured with the recommended WTD threshold as defined in Figure 2-34. The following ingress queue configuration must be enabled in global configuration mode on every access-layer switch.

```
cr25-3750-1(config)#mls qos srr-queue input priority-queue 2 bandwidth 30
 ! Q2 is enabled as a strict-priority ingress queue with 30% BW

cr25-3750-1(config)#mls qos srr-queue input bandwidth 70 30
 ! Q1 is assigned 70% BW via SRR shared weights
 ! Q1 SRR shared weight is ignored (as it has been configured as a PQ)

cr25-3750-1(config)#mls qos srr-queue input threshold 1 80 90
 ! Q1 thresholds are configured at 80% (Q1T1) and 90% (Q1T2)
 ! Q1T3 is implicitly set at 100% (the tail of the queue)
 ! Q2 thresholds are all set (by default) to 100% (the tail of Q2)

! This section configures ingress DSCP-to-Queue Mappings
cr25-3750-1(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 0 8 10 12 14
 ! DSCP DF, CS1 and AF1 are mapped to ingress Q1T1
cr25-3750-1(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 16 18 20 22
 ! DSCP CS2 and AF2 are mapped to ingress Q1T1
cr25-3750-1(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 26 28 30 34 36 38
 ! DSCP AF3 and AF4 are mapped to ingress Q1T1
cr25-3750-1(config)#mls qos srr-queue input dscp-map queue 1 threshold 2 24
 ! DSCP CS3 is mapped to ingress Q1T2
cr25-3750-1(config)#mls qos srr-queue input dscp-map queue 1 threshold 3 48 56
 ! DSCP CS6 and CS7 are mapped to ingress Q1T3 (the tail of Q1)
cr25-3750-1(config)#mls qos srr-queue input dscp-map queue 2 threshold 3 32 40 46
 ! DSCP CS4, CS5 and EF are mapped to ingress Q2T3 (the tail of the PQ)
```

```
cr25-3750-1#show mls qos input-queue
Queue:       12
-----------------------------------
buffers   :9010
bandwidth :7030
priority  :030
threshold1:80100
threshold2:90100

cr25-3750-1#show mls qos maps dscp-input-q
   Dscp-inputq-threshold map:
     d1 :d2    0        1         2        3          4         5        6          7
8          9

-------------------------------------------------------------------------------------
      0 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01
      1 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01
      2 :    01-01 01-01 01-01 01-01 01-02 01-01 01-01 01-01 01-01 01-01
      3 :    01-01 01-01 02-03 01-01 01-01 01-01 01-01 01-01 01-01 01-01
      4 :    02-03 02-01 02-01 02-01 02-01 02-01 02-03 02-01 01-03 01-01
      5 :    01-01 01-01 01-01 01-01 01-01 01-01 01-03 01-01 01-01 01-01
      6 :    01-01 01-01 01-01 01-01
```

Deploying Egress QoS

The QoS implementation for egress traffic toward the network edge on access-layer switches is much simpler than the ingress traffic QoS. The egress QoS implementation provides optimal queueing policies for each class and sets the drop thresholds to prevent network congestion and application performance impact. Cisco Catalyst switches support 4 hardware queues that are assigned the following policies:

- Real-time queue (to support a RFC 3246 EF PHB service)
- Guaranteed bandwidth queue (to support RFC 2597 AF PHB services)
- Default queue (to support a RFC 2474 DF service)
- Bandwidth constrained queue (to support a RFC 3662 scavenger service)

As a best practice each physical or logical link must diversify bandwidth assignment to map with hardware queues:

- Real-time queue should not exceed 33 percent of the link's bandwidth.
- Default queue should be at least 25 percent of the link's bandwidth.
- Bulk/scavenger queue should not exceed 5 percent of the link's bandwidth.

Figure 2-35 shows best practice egress queue bandwidth allocation for each class.

Figure 2-35 Engress QoS

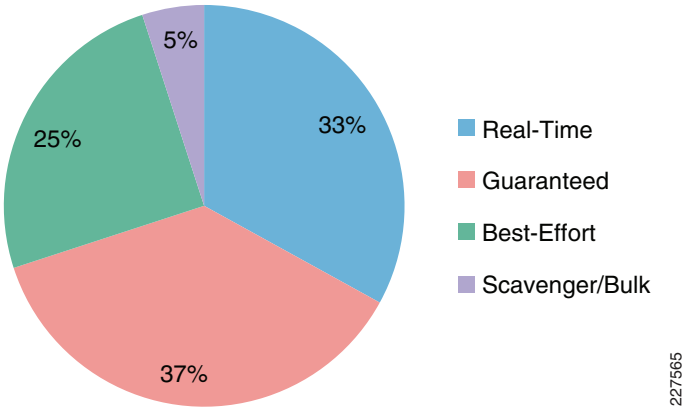

Given these minimum queuing requirements and bandwidth allocation recommendations, the following application classes can be mapped to the respective queues:

- *Realtime Queue*—Voice, broadcast video, and realtime interactive may be mapped to the realtime queue (per RFC 4594).
- *Guaranteed Queue*—Network/internetwork control, signaling, network management, multimedia conferencing, multimedia streaming, and transactional data can be mapped to the guaranteed bandwidth queue. Congestion avoidance mechanisms (i.e., selective dropping tools), such as WRED, can be enabled on this class. If configurable drop thresholds are supported on the platform, these may be enabled to provide intra-queue QoS to these application classes, in the respective order they are listed (such that control plane protocols receive the highest level of QoS within a given queue).
- *Scavenger/Bulk Queue*—Bulk data and scavenger traffic can be mapped to the bandwidth-constrained queue and congestion avoidance mechanisms can be enabled on this class. If configurable drop thresholds are supported on the platform, these may be enabled to provide inter-queue QoS to drop scavenger traffic ahead of bulk data.
- *Default Queue*—Best effort traffic can be mapped to the default queue; congestion avoidance mechanisms can be enabled on this class.

The egress queueing is designed to map traffic, based on DSCP value, to four egress queues. as shown above. The egress QoS model for a platform that supports DSCP-to-queue mapping with a 1P3Q8T queuing structure is depicted in Figure 2-36.

Figure 2-36 Access-Layer 1P3Q3T Egress Queue Model

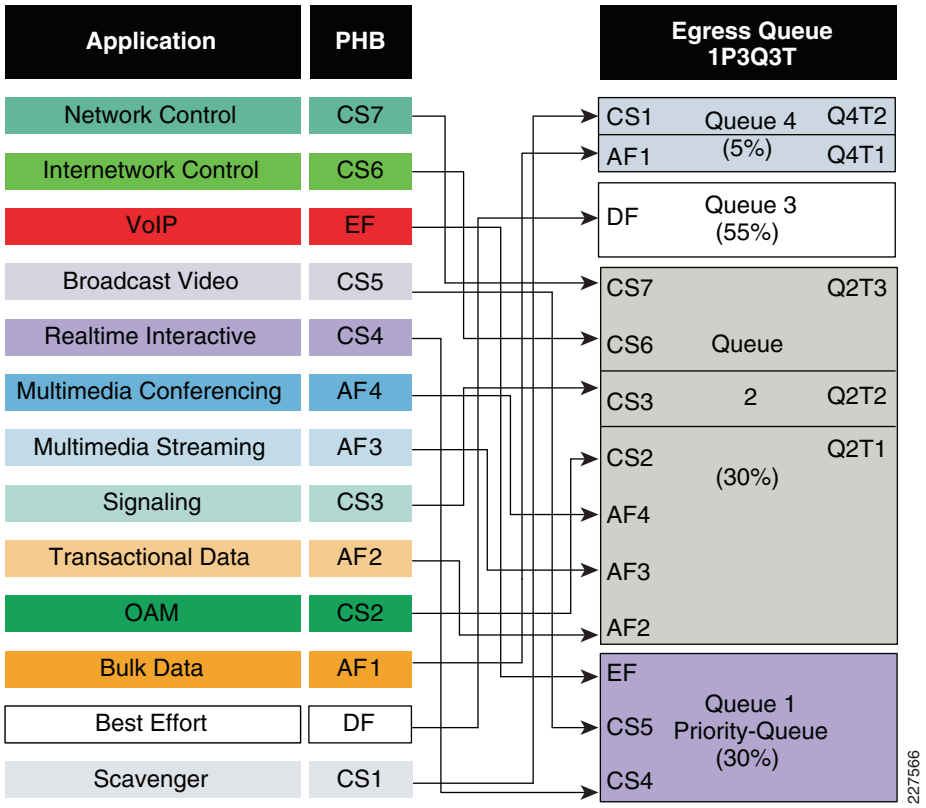

DSCP marked packets are assigned to the appropriate queue and each queue is configured with appropriate WTD threshold as defined in Figure 2-36. Egress queueing is the same on network edge port as well as on uplink connected to internal network, and it is independent of trust mode. The following egress queue configuration in global configuration mode, must be enabled on every access-layer switch in the network.

```
cr25-3750-1(config)#mls qos queue-set output 1 buffers 15 30 35 20
! Queue buffers are allocated
cr25-3750-1(config)#mls qos queue-set output 1 threshold 1 100 100 100 100
! All Q1 (PQ) Thresholds are set to 100%
cr25-3750-1(config)#mls qos queue-set output 1 threshold 2 80 90 100 400
! Q2T1 is set to 80%; Q2T2 is set to 90%;
 ! Q2 Reserve Threshold is set to 100%;
 ! Q2 Maximum (Overflow) Threshold is set to 400%
cr25-3750-1(config)#mls qos queue-set output 1 threshold 3 100 100 100 400
! Q3T1 is set to 100%, as all packets are marked the same weight in Q3
 ! Q3 Reserve Threshold is set to 100%;
 ! Q3 Maximum (Overflow) Threshold is set to 400%
cr25-3750-1(config)#mls qos queue-set output 1 threshold 4 60 100 100 400
! Q4T1 is set to 60%; Q4T2 is set to 100%
 ! Q4 Reserve Threshold is set to 100%;
 ! Q4 Maximum (Overflow) Threshold is set to 400%


 ! This section configures egress DSCP-to-Queue mappings
cr25-3750-1(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 32 40 46
 ! DSCP CS4, CS5 and EF are mapped to egress Q1T3 (tail of the PQ)
cr25-3750-1(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
 ! DSCP CS2 and AF2 are mapped to egress Q2T1
```

```
cr25-3750-1(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
 ! DSCP AF3 and AF4 are mapped to egress Q2T1
cr25-3750-1(config)#mls qos srr-queue output dscp-map queue 2 threshold 2 24
 ! DSCP CS3 is mapped to egress Q2T2
cr25-3750-1(config)#mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
 ! DSCP CS6 and CS7 are mapped to egress Q2T3

cr25-3750-1(config)#mls qos srr-queue output dscp-map queue 3 threshold 3 0
 ! DSCP DF is mapped to egress Q3T3 (tail of the best effort queue)
cr25-3750-1(config)#mls qos srr-queue output dscp-map queue 4 threshold 1 8
 ! DSCP CS1 is mapped to egress Q4T1
cr25-3750-1(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
 ! DSCP AF1 is mapped to Q4T2 (tail of the less-than-best-effort queue)

 ! This section configures interface egress queuing parameters
cr25-3750-1(config)#interface range GigabitEthernet1/0/1-48
cr25-3750-1(config-if-range)# queue-set 1
 ! The interface(s) is assigned to queue-set 1
cr25-3750-1(config-if-range)# srr-queue bandwidth share 1 30 35 5
 ! The SRR sharing weights are set to allocate 30% BW to Q2
 ! 35% BW to Q3 and 5% BW to Q4
 ! Q1 SRR sharing weight is ignored, as it will be configured as a PQ
cr25-3750-1(config-if-range)# priority-queue out
 ! Q1 is enabled as a strict priority queue

cr25-3750-1#show mls qos interface GigabitEthernet1/0/27 queueing
GigabitEthernet1/0/27
Egress Priority Queue : enabled
Shaped queue weights (absolute) :  25 0 0 0
Shared queue weights  :  1 30 35 5
The port bandwidth limit : 100  (Operational Bandwidth:100.0)
    The port is mapped to qset : 1
```

Table 2-10 and Table 2-11 summarize the ingress and egress QoS policies at the access-layer for several types of validated endpoints.

Table 2-10 Summarized Network Edge Ingress QoS Deployment Guidelines

endpoint	Trust Model	DSCP Trust	Classification	Marking	Policing	Ingress Queueing
Unmanaged devices, printers etc	UnTrusted	Don't Trust. Default.	None	None	Yes	Yes
Managed secured devices, Servers etc	Trusted	Trust	8 Class Model	Yes	Yes	Yes
Phone	Trusted	Trust	Yes	Yes	Yes	Yes
Phone + Mobile PC	Conditionally-Trusted	Trust	Yes	Yes	Yes	Yes
IP Video surveillance Camera	Trusted	Trust	No	No	No	Yes
Digital Media Player	Trusted	Trust	No	No	No	Yes
Core facing Uplinks	Trusted	Trust	No	No	No	Yes

Table 2-11 Summarized Network Edge Egress QoS Deployment Guidelines

Endpoint	Trust Model	Classification / Marking / Policing	Egress Queueing	Bandwidth Share
Unmanaged devices, printers etc	UnTrusted	None	Yes	Yes
Managed secured devices, Servers etc	Trusted	None	Yes	Yes
Phone	Trusted	None	Yes	Yes
Phone + Mobile PC	Conditionally-Trusted	None	Yes	Yes
IP Video surveillance Camera	Trusted	None	Yes	Yes
Digital Media Player	Trusted	None	Yes	Yes
Core facing Uplinks	Trusted	None	Yes	Yes

Deploying Network Core QoS

All connections between internal network devices that are deployed within the network domain boundary are classified as trusted devices and follow the same QoS best practices recommended in the previous section. Ingress and egress core QoS policies are simpler than those applied at the network edge, See Figure 2-37.

Figure 2-37 Network Core QoS Boundary

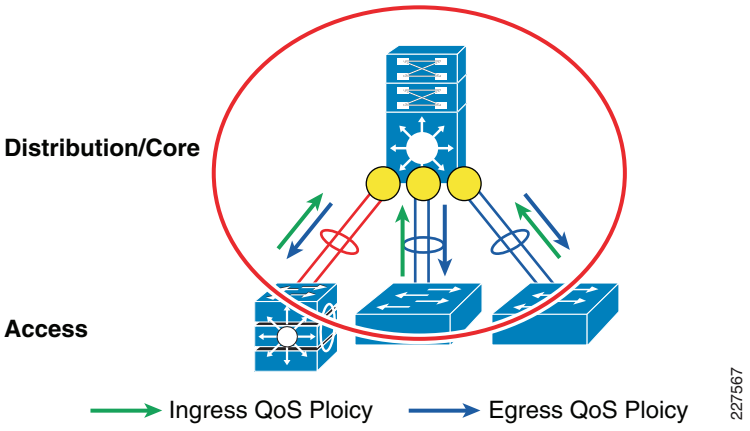

The core network devices are considered trusted and rely on the access-switch to properly mark DSCP values. The core network is deployed to ensure consistent differentiated QoS service across the network. This ensures there is no service quality degradation for high-priority traffic, such as IP telephony or video.

The QoS implementation at the main and remote large site differ from the remote small site, due to different platforms used as the collapsed core router (Catalyst 4500-E vs Catalyst 3750-X StackWise Plus).

Deploying Main or Remote Large Site Ingress QoS

The collapsed core at main site is deployed with Cisco Catalyst 450R-E with Supervisor-6E, whereas the remote large site collapsed core is deployed with Cisco Catalyst 4507R-E with either Supervisor-6E, Supervisor-6LE or Supervisor-V. The next-generation Sup-6E and Sup-6LE module has a redesigned QoS implementation which matches Cisco IOS routers. No ingress QoS configuration is required, since QoS is enabled by default, and all ports are considered trusted.

The Cisco Catalyst 4507R-E with Supervisor-V requires ingress QoS configuration similar to trusted endpoints in the access-layer. Following is a sample configuration which enables QoS in the Catalyst 4507R-E with Supervisor-V:

```
cr35-4507-1(config)#qos
! Enables QoS function in the switch

cr35-4507-1#show qos
QoS is enabled globally
IP header DSCP rewrite is enabled
```

After QoS is globally enabled, all interfaces are in the untrusted mode by default. QoS trust settings must be set on each Layer 2 or Layer 3 port that is physically connected to another device within the network trust boundary. When Cisco Catalyst 4500 is deployed in EtherChannel mode, the QoS trust settings must be applied to every physical member-link and logical port-channel interface. Best practice is to enable trust DSCP settings on each physical and logical interface that connects to another internal trusted device (e.g., access-layer switches in wiring closet or data-center, a router, wireless LAN controller (WLC)).

```
cr35-4507-1(config)#interface range Po11 , Gi1/2 , Gi2/2
cr35-4507-1(config-if-range)#description Connected to cr35-2960-1
cr35-4507-1(config-if-range)#qos trust dscp

cr35-4507-1#show qos interface Port-channel 11
QoS is enabled globally
Port QoS is enabled
Administrative Port Trust State: 'dscp'
Operational Port Trust State: 'dscp'
Trust device: none
Default DSCP: 0 Default CoS: 0
```

Additional ingress QoS techniques (such as classification, marking, and policing) are not required at the collapsed core layer since these functions are already performed by the access-layer switches. The architecture of Catalyst 4500-E with classic or next-generation Supervisor do not need ingress queueing since all of the forwarding decisions are made centrally on the supervisor. There are no additional QoS configurations required at the collapsed core-layer system.

Deploying Remote Small Site Ingress QoS

The remote small site is deployed using Cisco Catalyst 3750-X StackWise Plus as the collapsed core switch. The QoS implementation remains the same whether deployed as 3750-X StackWise or as a standalone switch. By default, QoS is disabled on the 3750-X switch. Following is a sample configuration to enable QoS in global configuration mode:

```
cr36-3750s-1(config)#mls qos
! Enables QoS function in the switch

cr36-3750s-1#show mls qos

QoS is enabled
QoS ip packet dscp rewrite is enabled
```

After QoS is globally enabled, all interfaces are in the untrusted mode by default. QoS trust settings must be set on each Layer 2 or Layer 3 port that is physically connected to another device within the network trust boundary. When Cisco Catalyst 3750-E StackWise Plus is deployed in EtherChannel mode, the QoS trust settings must be applied to every physical member-link. Best practice is to enable trust DSCP settings on each physical and logical interface that connects to another internal trusted device (e.g., access-layer switches in wiring closet or data-center, a router, wireless LAN controller (WLC)).

```
cr36-3750s-1(config)#int range gi1/0/49 , gi3/0/49
cr36-3750s-1(config-if-range)# description Connected to cr36-2960-1
cr36-3750s-1(config-if-range)#mls qos trust dscp

cr36-3750s-1#show mls qos interface Gi1/0/49
GigabitEthernet1/0/49
trust state: trust dscp
trust mode: trust dscp
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based
```

Additional ingress QoS techniques (such as classification, marking, and policing) are not required at the collapsed core layer since these functions are already performed by the access-layer switches. The ingress queueing and DSCP-Ingress-Queue function in 3750-X StackWise Plus must be enabled to allow differentiation between normal versus high-priority traffic. The ingress queuing configuration is consistent with the implementation at the access-edge. Following is a sample configuration for the ingress queues of the Catalyst 3750-X StackWise collapsed core switch:

```
cr36-3750-1(config)#mls qos srr-queue input priority-queue 2 bandwidth 30
 ! Q2 is enabled as a strict-priority ingress queue with 30% BW

cr36-3750-1(config)#mls qos srr-queue input bandwidth 70 30
 ! Q1 is assigned 70% BW via SRR shared weights
 ! Q1 SRR shared weight is ignored (as it has been configured as a PQ)

cr36-3750-1(config)#mls qos srr-queue input threshold 1 80 90
 ! Q1 thresholds are configured at 80% (Q1T1) and 90% (Q1T2)
 ! Q1T3 is implicitly set at 100% (the tail of the queue)
 ! Q2 thresholds are all set (by default) to 100% (the tail of Q2)

! This section configures ingress DSCP-to-Queue Mappings
cr36-3750-1(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 0 8 10 12 14
 ! DSCP DF, CS1 and AF1 are mapped to ingress Q1T1
cr36-3750-1(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 16 18 20 22
```

```
 ! DSCP CS2 and AF2 are mapped to ingress Q1T1
cr36-3750-1(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 26 28 30 34 36 38
 ! DSCP AF3 and AF4 are mapped to ingress Q1T1
cr36-3750-1(config)#mls qos srr-queue input dscp-map queue 1 threshold 2 24
 ! DSCP CS3 is mapped to ingress Q1T2
cr36-3750-1(config)# mls qos srr-queue input dscp-map queue 1 threshold 3 48 56
 ! DSCP CS6 and CS7 are mapped to ingress Q1T3 (the tail of Q1)
cr36-3750-1(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 32 40 46
 ! DSCP CS4, CS5 and EF are mapped to ingress Q2T3 (the tail of the PQ)

cr36-3750s-1#show mls qos input-queue
Queue     :       1       2

-------------------------------------
buffers   :      90      10
bandwidth :      70      30
priority  :       0      30
threshold1:      80     100
threshold2:      90     100

cr36-3750s-1#show mls qos maps dscp-input-q
   Dscp-inputq-threshold map:
     d1 :d2    0        1         2        3          4         5        6          7
8          9

-----------------------------------------------------------------------------------------

      0 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01
      1 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01
      2 :    01-01 01-01 01-01 01-01 01-02 01-01 01-01 01-01 01-01 01-01
      3 :    01-01 01-01 02-03 01-01 01-01 01-01 01-01 01-01 01-01 01-01
      4 :    02-03 02-01 02-01 02-01 02-01 02-01 02-03 02-01 01-03 01-01
      5 :    01-01 01-01 01-01 01-01 01-01 01-01 01-03 01-01 01-01 01-01
      6 :    01-01 01-01 01-01 01-01
```

Deploying Main Site Egress QoS

The main site is deployed with Cisco Catalyst 4507R-E with Supervisor-6E as the collapsed core router. Egress QoS from the collapsed core router provides optimized queueing and drop thresholds to drop excess low-priority traffic and protect high-priority traffic.

The Supervisor-6E supports up to 8 traffic classes for QoS mapping. It also supports a platform-specific congestion avoidance algorithm to provide Active Queue Management (AQM) with Dynamic Buffer Limiting (DBL). DBL tracks the queue length for each traffic flow in the switch. When the queue length of a flow exceeds its limit, DBL drops packets or sets the Explicit Congestion Notification (ECN) bit in the TCP packet header. With 8 egress (1P7Q1T) queues and DBL capability in the Sup-6E, the bandwidth distribution for each class changes, as shown in Figure 2-38.

Figure 2-38 Small Enterprise Main Site Network

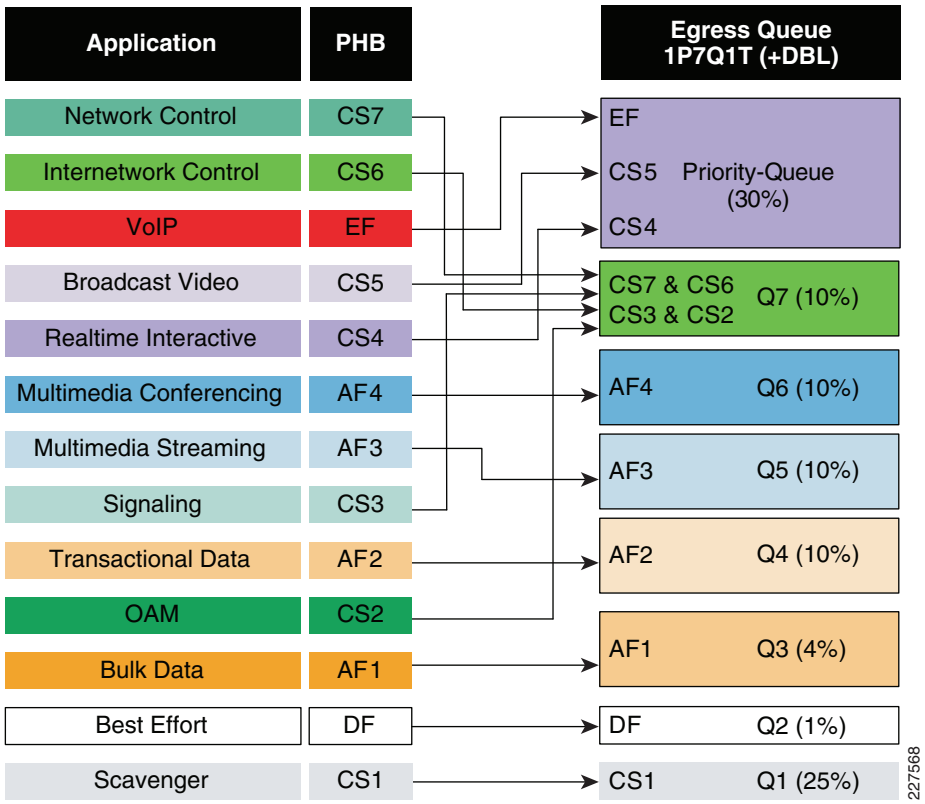

Implementing QoS policies on Sup-6E-based Catalyst 4500 platform follows IOS (MQC)-model. The egress QoS implementation bundles the queueing and policing functions on EtherChannel based networks. To provide low-latency for high priority traffic, all lower priority traffic must wait until the priority-queue is empty. Best practice includes implementing a policer along with the priority-queue to provide more fair treatment for all traffic.

The following sample configuration shows how to create an 8-class egress queueing model and protect from high-priority traffic consuming more bandwidth than global policies allow. The egress QoS service-policy must be applied to all the physical EtherChannel member-links connected to different service-blocks (i.e., WAN edge, serverfarm, access-layer switches, etc).

```
! Creating class-map for each classes using match dscp statement as marked by edge systems
cr24-4507-1(config)#class-map match-all PRIORITY-QUEUE
cr24-4507-1(config-cmap)# match  dscp ef
cr24-4507-1(config-cmap)# match  dscp cs5
cr24-4507-1(config-cmap)# match  dscp cs4
cr24-4507-1(config-cmap)#class-map match-all CONTROL-MGMT-QUEUE
cr24-4507-1(config-cmap)# match  dscp cs7
cr24-4507-1(config-cmap)# match  dscp cs6
cr24-4507-1(config-cmap)# match  dscp cs3
cr24-4507-1(config-cmap)# match  dscp cs2
cr24-4507-1(config-cmap)#class-map match-all MULTIMEDIA-CONFERENCING-QUEUE
cr24-4507-1(config-cmap)# match  dscp af41  af42  af43
cr24-4507-1(config-cmap)#class-map match-all MULTIMEDIA-STREAMING-QUEUE
cr24-4507-1(config-cmap)# match  dscp af31  af32  af33
cr24-4507-1(config-cmap)#class-map match-all TRANSACTIONAL-DATA-QUEUE
cr24-4507-1(config-cmap)# match  dscp af21  af22  af23
cr24-4507-1(config-cmap)#class-map match-all BULK-DATA-QUEUE
cr24-4507-1(config-cmap)# match  dscp af11  af12  af13
cr24-4507-1(config-cmap)#class-map match-all SCAVENGER-QUEUE
```

```
cr24-4507-1(config-cmap)# match  dscp cs1
```

```
! Creating policy-map and configure queueing for class-of-service
cr24-4507-1(config)#policy-map EGRESS-POLICY
cr24-4507-1(config-pmap)# class PRIORITY-QUEUE
cr24-4507-1(config-pmap-c)# priority
cr24-4507-1(config-pmap-c)# class CONTROL-MGMT-QUEUE
cr24-4507-1(config-pmap-c)# bandwidth remaining percent 10
cr24-4507-1(config-pmap-c)# class MULTIMEDIA-CONFERENCING-QUEUE
cr24-4507-1(config-pmap-c)# bandwidth remaining percent 10
cr24-4507-1(config-pmap-c)# class MULTIMEDIA-STREAMING-QUEUE
cr24-4507-1(config-pmap-c)# bandwidth remaining percent 10
cr24-4507-1(config-pmap-c)# class TRANSACTIONAL-DATA-QUEUE
cr24-4507-1(config-pmap-c)# bandwidth remaining percent 10
cr24-4507-1(config-pmap-c)# dbl
cr24-4507-1(config-pmap-c)# class BULK-DATA-QUEUE
cr24-4507-1(config-pmap-c)# bandwidth remaining percent 4
cr24-4507-1(config-pmap-c)# dbl
cr24-4507-1(config-pmap-c)# class SCAVENGER-QUEUE
cr24-4507-1(config-pmap-c)# bandwidth remaining percent 1
cr24-4507-1(config-pmap-c)# class class-default
cr24-4507-1(config-pmap-c)# bandwidth remaining percent 25
cr24-4507-1(config-pmap-c)# dbl

! Attaching egress service-policy on all physical member-link ports
cr24-4507-1(config)#int range Gi1/1 - 6 , Gi2/1 - 6
cr24-4507-1(config-if-range)# service-policy output EGRESS-POLICY
```

EtherChannel is an aggregated logical bundle interface that does not perform queueing and relies on individual member-links to queue egress traffic. The policer to rate-limit priority class traffic must be implemented on EtherChannel and not on individual member-links since it governs the aggregate egress traffic limits. The following additional policy-map must be created to classify priority-queue class traffic and rate-limit the traffic to 30 percent of egress link capacity:

```
cr24-4507-1(config)#class-map match-any PRIORITY-QUEUE
cr24-4507-1(config-cmap)# match  dscp ef
cr24-4507-1(config-cmap)# match  dscp cs5
cr24-4507-1(config-cmap)# match  dscp cs4

cr24-4507-1(config)#policy-map PQ-POLICER
cr24-4507-1(config-pmap)# class PRIORITY-QUEUE
cr24-4507-1(config-pmap-c)# police cir 300 m conform-action transmit exceed-action drop

cr24-4507-1(config)#interface range Port-Channel 1 , Port-channel 11 - 17
cr24-4507-1(config-if-range)#service-policy output PQ-POLICER
```

Deploying Remote Large Site Egress QoS

The remote large site is deployed with Cisco Catalyst 450R-E and either Supervisor-6E, Supervisor-6LE or Supervisor-V as the collapsed core router. If the remote large site network is deployed with Sup-6E or Sup-6LE, then the configuration is the same as described in the previous section.

The QoS deployment and implementation guidelines differ when the Cisco Catalyst 4500-E is deployed with the classic Supervisor-V module. The SupV supervisor can have up to four egress queues like the Cisco Catalyst 2960, 2960-S and 35xx/37xx-X Series switches. Before forwarding egress traffic, each packet must be internally classified and placed in the appropriate egress-queue. Placing traffic into different class-of-service queues, will offer traffic prioritization and guaranteed bandwidth to the network. The following sample configuration shows how to implement egress QoS on the Catalyst 4500-E with Supervisor-V:

```
cr35-4507-1(config)#qos dbl
 ! DBL is globally enabled
cr35-4507-1(config)#no qos dbl dscp-based 32
cr35-4507-1(config)#no qos dbl dscp-based 40
cr35-4507-1(config)#no qos dbl dscp-based 46
 ! DBL is explicitly disabled on DSCP CS4, CS5 and EF
 ! as these DSCP values are assigned to the PQ
 ! and as such should never experience congestion avoidance drops
cr35-4507-1(config)#qos dbl exceed-action ecn
 ! DBL will mark IP ECN bits in the event of congestion

! This section configures the DBL policy-map
cr35-4507-1(config)#policy-map DBL
cr35-4507-1(config-pmap)# class class-default
cr35-4507-1(config-pmap-c)# dbl
 ! DBL is enabled on all flows
 ! (with the exception of DSCP CS4, CS5 and EF)
 ! This section configures the DSCP-to-Queue mappings


cr35-4507-1(config)#qos map dscp 8 10 12 14 to tx-queue 1
 ! DSCP CS1 and AF1 are mapped to Q1 (the less than best effort queue)
cr35-4507-1(config)#qos map dscp 0 to tx-queue 2
 ! DSCP DF is mapped to Q2 (the best effort/default queue)
cr35-4507-1(config)#qos map dscp 32 40 46 to tx-queue 3
 ! DSCP CS4, CS5 and EF are mapped to Q3 (the PQ)
cr35-4507-1(config)#qos map dscp 16 18 20 22 to tx-queue 4
 ! DSCP CS2 and AF2 are mapped to Q4 (guaranteed BW queue)
cr35-4507-1(config)#qos map dscp 24 26 28 30 to tx-queue 4
 ! DSCP CS3 and AF3 are mapped to Q4 (guaranteed BW queue)
cr35-4507-1(config)#qos map dscp 34 36 38 to tx-queue 4
 ! DSCP AF4 is mapped to Q4 (guaranteed BW queue)
cr35-4507-1(config)#qos map dscp 48 56 to tx-queue 4
 ! DSCP CS6 and CS7 are mapped to Q4 (guaranteed BW queue)

 ! This section configures all the EtherChannel member-link for egress queuing
cr35-4507-1(config)#interface range Gig1/1 - 6 , Gig2/1 - 6
cr35-4507-1(config-if-range)# tx-queue 1
cr35-4507-1(config-if-tx-queue)# bandwidth percent 5
 ! Q1 (less than best effort queue) is assigned 5% BW
cr35-4507-1(config-if-tx-queue)# tx-queue 2
cr35-4507-1(config-if-tx-queue)# bandwidth percent 35
 ! Q2 (default/best effort queue) is assigned 35% BW
cr35-4507-1(config-if-tx-queue)# tx-queue 3
cr35-4507-1(config-if-tx-queue)# priority high
cr35-4507-1(config-if-tx-queue)# bandwidth percent 30
 ! Q3 is enabled as a PQ and assigned 30% BW
cr35-4507-1(config-if-tx-queue)# tx-queue 4
cr35-4507-1(config-if-tx-queue)# bandwidth percent 30
 ! Q4 (guaranteed BW queue) is assigned 30% BW
cr35-4507-1(config-if-range)# service-policy output DBL
 ! DBL policy-map is attached to the interface(s)

cr35-4507-1#show qos dbl
QOS is enabled globally
```

```
DBL is enabled globally on DSCP values:
     0-31,33-39,41-45,47-63
DBL flow includes vlan
DBL flow includes layer4-ports
DBL uses ecn to indicate congestion
DBL exceed-action probability: 15%
DBL max credits: 15
DBL aggressive credit limit: 10
DBL aggressive buffer limit: 2 packets
```

```
cr35-4507-1#show qos maps dscp tx-queue
DSCP-TxQueue Mapping Table (dscp = d1d2)
d1 : d2  0   1   2    3   4   5   6    7   8   9
------------------------------------------------
 0 :    02 01 01 01 01 01 01 01 01 01
 1 :    01 01 01 01 01 01 04 02 04 02
 2 :    04 02 04 02 04 02 04 02 04 02
 3 :    04 02 03 03 04 03 04 03 04 03
 4 :    03 03 03 03 03 03 03 03 04 04
 5 :    04 04 04 04 04 04 04 04 04 04
 6 :    04 04 04 04
```

```
cr35-4507-1#show qos interface Gig1/2
QoS is enabled globally
Port QoS is enabled
Administrative Port Trust State: 'dscp'
Operational Port Trust State: 'dscp'
Trust device: none
Default DSCP: 0 Default CoS: 0
Appliance trust: none
Tx-Queue   Bandwidth   ShapeRate   Priority   QueueSize
             (bps)       (bps)                (packets)
  1        50000000     disabled    N/A        2080
  2        350000000     disabled    N/A2080
  3        300000000disabled    high2080
  4        300000000disabledN/A2080
```

Deploying Remote Small Site Egress QoS

Collapsed Core— Catalyst 3750-X StackWise Plus

The remote small site is deployed with Cisco Catalyst 3750-X StackWise Plus as the collapsed core router. The Catalyst 3750-X can have up to four egress queues. Before forwarding egress traffic, each packet is placed in the appropriate egress-queue as shown in Figure 2-36. The Catalyst 3750-E switch supports Shaped Round Robin (SRR) packet schedule service which can be deployed in two different modes:

- *Shaped*—To provide guaranteed bandwidth, the shaped egress queue reserves some of the bandwidth of the port for each queue. Traffic load exceeding the shape parameter gets dropped. The queue cannot take advantage of excess bandwidth capacity when other queues are not using their bandwidth allocations.
- *Shared*—Shared mode also provides guaranteed bandwidth for each queue; however, it allows the flexibility of using excess bandwidth when there is any available.

The following sample configuration shows how to implement egress QoS on the Catalyst 3750-X:

```
cr36-3750s-1(config)#mls qos queue-set output 1 buffers 15 30 35 20
! Queue buffers are allocated
cr36-3750-1(config)#mls qos queue-set output 1 threshold 1 100 100 100 100
! All Q1 (PQ) Thresholds are set to 100%
cr36-3750s-1(config)#mls qos queue-set output 1 threshold 2 80 90 100 400
! Q2T1 is set to 80%; Q2T2 is set to 90%;
 ! Q2 Reserve Threshold is set to 100%;
 ! Q2 Maximum (Overflow) Threshold is set to 400%
cr36-3750s-1(config)#mls qos queue-set output 1 threshold 3 100 100 100 400
! Q3T1 is set to 100%, as all packets are marked the same weight in Q3
 ! Q3 Reserve Threshold is set to 100%;
 ! Q3 Maximum (Overflow) Threshold is set to 400%
cr36-3750s-1(config)#mls qos queue-set output 1 threshold 4 60 100 100 400
! Q4T1 is set to 60%; Q4T2 is set to 100%
 ! Q4 Reserve Threshold is set to 100%;
 ! Q4 Maximum (Overflow) Threshold is set to 400%

 ! This section configures egress DSCP-to-Queue mappings
cr36-3750s-1(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 32 40 46
 ! DSCP CS4, CS5 and EF are mapped to egress Q1T3 (tail of the PQ)
cr36-3750s-1(config)# mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22
 ! DSCP CS2 and AF2 are mapped to egress Q2T1
cr36-3750s-1(config)#mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30 34 36 38
 ! DSCP AF3 and AF4 are mapped to egress Q2T1
cr36-3750s-1(config)#mls qos srr-queue output dscp-map queue 2 threshold 2 24
 ! DSCP CS3 is mapped to egress Q2T2
cr36-3750s-1(config)#mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
 ! DSCP CS6 and CS7 are mapped to egress Q2T3
cr36-3750s-1(config)#mls qos srr-queue output dscp-map queue 3 threshold 3 0
 ! DSCP DF is mapped to egress Q3T3 (tail of the best effort queue)
cr36-3750s-1(config)#mls qos srr-queue output dscp-map queue 4 threshold 1 8
 ! DSCP CS1 is mapped to egress Q4T1
cr36-3750s-1(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
 ! DSCP AF1 is mapped to Q4T2 (tail of the less-than-best-effort queue)

 ! This section configures interface egress queuing parameters
cr36-3750s-1(config)#interface range GigabitEthernet1/0/1-48
cr36-3750s-1(config-if-range)# queue-set 1
 ! The interface(s) is assigned to queue-set 1
cr36-3750s-1(config-if-range)# srr-queue bandwidth share 1 30 35 5
 ! The SRR sharing weights are set to allocate 30% BW to Q2
 ! 35% BW to Q3 and 5% BW to Q4
 ! Q1 SRR sharing weight is ignored, as it will be configured as a PQ
cr36-3750s-1(config-if-range)# priority-queue out
 ! Q1 is enabled as a strict priority queue

cr36-3750s-1#show mls qos interface GigabitEthernet1/0/49 queueing
GigabitEthernet1/0/49
Egress Priority Queue : enabled
Shaped queue weights (absolute) :  25 0 0 0
Shared queue weights  :  1 30 35 5
The port bandwidth limit : 100  (Operational Bandwidth:100.0)
    The port is mapped to qset : 1
```

Building a Resilient Network

The Small Enterprise Design Profile is a high performance, resilient and scalable network design. A network outage may be caused by the system, human error, or natural disaster. The small enterprise network is designed to minimize the impact of a failure regardless of the cause. Network outages may be either planned or unplanned.

- *Planned Outage*—Planned network outage occurs when a portion of the network is taken out of service as part of a scheduled event (e.g., a software upgrade).
- *Unplanned Outage*—Any unscheduled network outage is considered an unplanned outage. Such outages may be caused by internal faults in the network, or devices due to hardware or software malfunctions.

The network is designed to recover from most un planned outages in less than a second (milliseconds). In many situations, the user will not even notice the outage occurred. If the outage lasts longer (several seconds), then the user will notice the lack of application responsiveness. The network is designed to minimize the overall impact of a unplanned network outage, and gracefully adjust and recover from many outage conditions. Figure 2-39 shows an example of a real-time VoIP application and user impact depending on duration of outage event.

Figure 2-39 VoIP User Impact for Minor and Major Network Outage

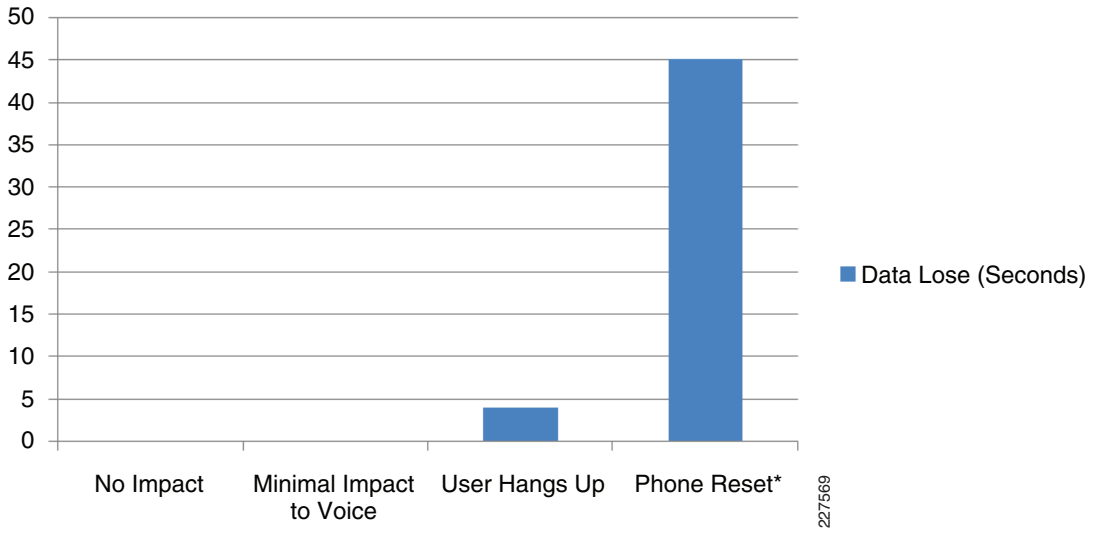

Several techniques are used to make the network design more resilient. Deploying redundant devices and redundant connections between devices, enables the network to recover from fault conditions. Identifying critical versus non critical applications, and network resources optimizes cost performance, by focusing on the most important elements of the network design. The resiliency of a system design is often categorized as follows:

- *Network Resiliency*—Provides redundancy during physical link outages (e.g., fiber cut, bad transceivers, incorrect cabling, etc).
- *Device Resiliency*—Protects network during device outage triggered by hardware or software (e.g. software crash, non-responsive supervisor, etc).

- *Operational Resiliency*—Capabilities which provide network availability even during planned network outage conditions (e.g., ISSU features which enable software upgrades while device is operating).

Figure 2-40 Resiliency Deployment Strategy

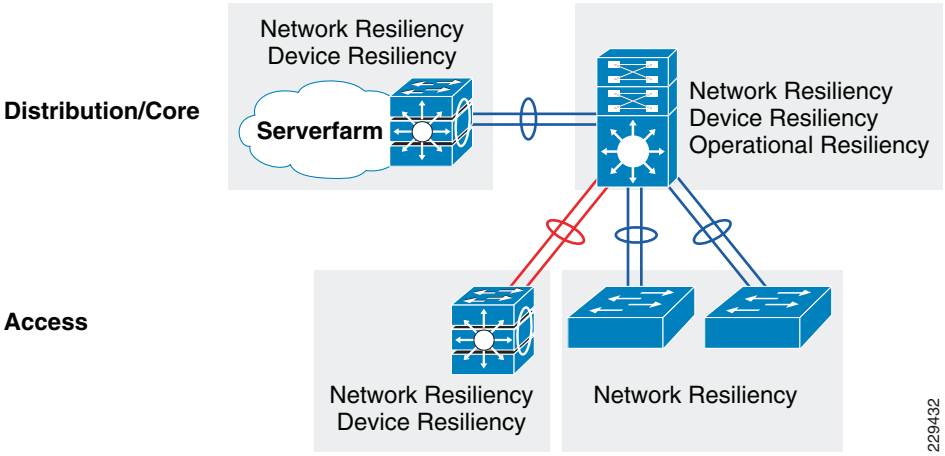

The high availability framework is based upon the three resiliency categories described in the previous section. Figure 2-41 shows which technologies are implemented to achieve each category of resiliency.

Figure 2-41 High-Availability Categories and Technologies

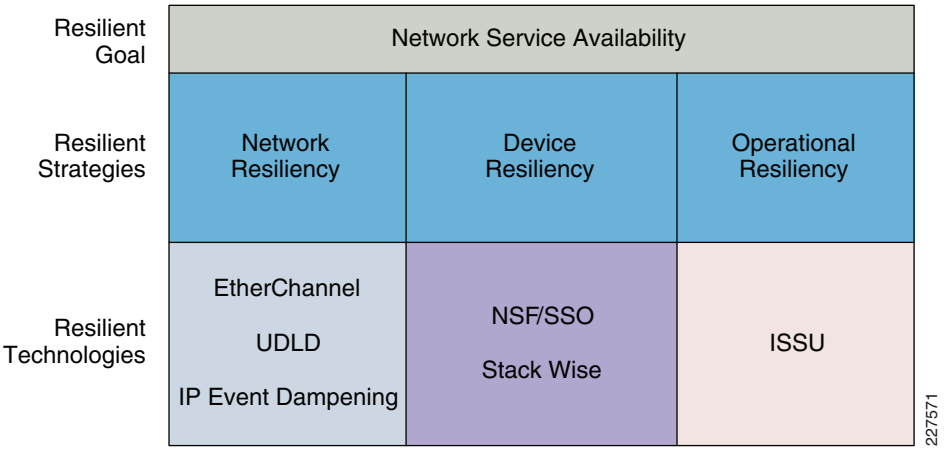

Redundant Hardware Components

Redundant hardware implementations vary between fixed configuration and modular Cisco Catalyst switches. Selective deployment of redundant hardware is an important element of the Small Enterprise Design Profile which delivers device resiliency.

Redundant hardware component for device resiliency varies between fixed configuration and modular Cisco Catalyst switches. To protect against common network faults or resets, all critical main and remote site campus network devices must be deployed with similar device resiliency configuration. This subsection provides a basic redundant hardware deployment guideline at the access-layer and collapsed core switching platforms in the campus network.

Redundant Power System

Redundant power supplies protect the device from power outage or power supply failure. Protecting the power is not only important for the network device, but also the endpoints that rely on power delivery over the Ethernet network. Redundant power supplies are deployed differently depending on the switch type:

- *Modular Switch*—Dual power supplies can be deployed in the modular switching platforms like the Cisco Catalyst 4500-E. By default, the Cisco Catalyst 4500 power supply operates in 1+1 redundant mode (both power supplies are active).
- *Fixed Configuration Switch*—Fixed configuration switches are deployed with internal power supplies and they may also use Cisco RPS 2300 external power supply. A single Cisco RPS 2300 power supply has modular power supplies and fans to deliver power to multiple switches. Deploying internal and external power-supplies provides a redundant power solution for fixed configuration switches.

Redundant Network Connectivity

Redundant network connections protect the system from failure due to cable or transceiver faults. Redundant network connections attached to a single fixed configuration switch or network module in the Cisco Catalyst 4500 switch do not protect against internal device hardware or software fault.

Best practice design is to deploy redundant network modules within the Catalyst 4500 switch and the Cisco 3750-X StackWise Plus solution in the small remote site collapsed core network. Deploying the 3750-X StackWise Plus in critical access-layer switches in the serverfarm network and in the main site is also best practice. Connecting redundant paths to different hardware elements provides both network and device resiliency.

Figure 2-42 Redundant Network Connectivity

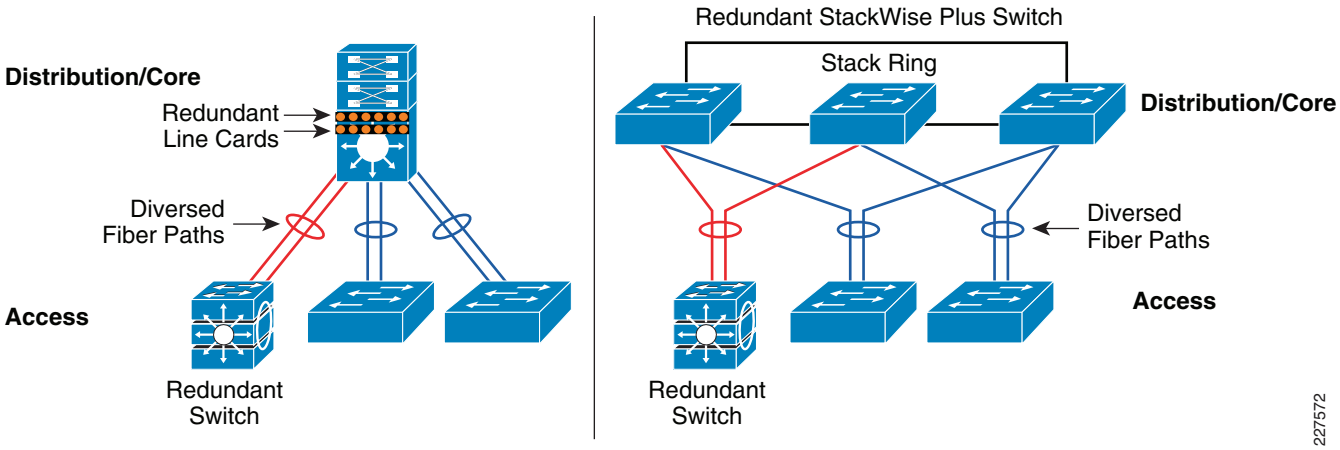

Redundant Control-Plane

The processing software operation is different in standalone or StackWise fixed configuration switches, and on a supervisor module of a modular switch. Network communication and forwarding operations can be disrupted when the processing unit fails, causing a network outage. Network recovery techniques vary based on the different platforms.

The standalone and non-stackable fixed configuration switches like the Cisco Catalyst 2960 or 3560-E feature power redundancy and network resiliency support; however they do not protect against a processing unit failure. During a processing unit failure event, all endpoints attached to the switch are impacted and network recovery time is undeterministic.

Device resiliency in Cisco StackWise and modular switching platforms provides 1+1 redundancy with enterprise-class high availability and deterministic network recovery time.

Cisco StackWise Plus

Cisco Catalyst 3750-X switches can be deployed in StackWise mode using a special stack cable. Up to nine switches can be integrated into a single stack that delivers distributed forwarding architecture and unified single control and management plane. Device level redundancy in StackWise mode is achieved via stacking multiple switches using the Cisco StackWise technology. One switch from the stack is selected automatically to serve as the master, which manages the centralized control-plane process. Cisco StackWise solution provides 1:N redundancy. In the event of a active master-switch outage, a new master is selected automatically. See Figure 2-43.

Figure 2-43 Cisco Stack Wise Switching Architecture

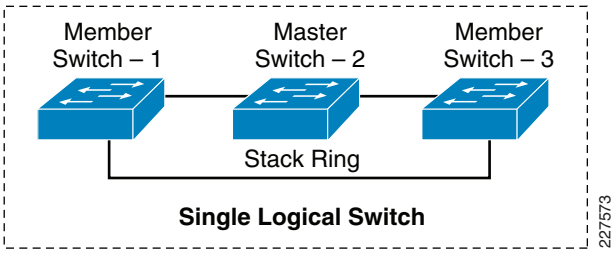

Since Cisco StackWise enables up to 9 switches to appear as one logical switch, it has centralized management and control functions. Most Layer 2 and Layer 3 functions are centrally performed, however Layer-2 topology development is distributed (i.e., each switch performs the function independently). Table 2-12 lists network protocol functions and identifies which are centralized and which are distributed.

Table 2-12 Cisco StackWise Centralized and Distributed Control-Plane

Protocols		Function
Layer 2 Protocols	MAC Table	Distributed
	Spanning-Tree Protocol	Distributed
	CDP	Centralized
	VLAN Database	Centralized
	EtherChannel - LACP	Centralized
Layer 3 Protocols	Layer 3 Management	Centralized
	Layer 3 Routing	Centralized

Cisco StackWise Plus solution offers network and device resiliency with distributed forwarding. In the event of a master switch outage, Non-Stop Forwarding (NSF) enables packet forwarding to continue based on current state information, while a new master switch is selected. New master switch selection

is accomplished in the range of 700 to 1000 milliseconds; the amount of time to reestablish the control-plane and develop distributed forwarding will vary depending on the size and complexity of the network.

Following is a best practice to reduce Layer-3 disruption in the event of a master switch outage: Determine the master switch with the higher switch priority, and isolate the uplink Layer-3 EtherChannel bundle path by using physical ports from member switches (i.e. don't use the master switches ports for Etherchannel uplinks). With NSF capabilities enabled, this design decreases network downtime during a master-switch outage.

An understanding of SSO and StackWise components and failover events associated with NSF provides significant insight in designing a network that enables supervisor redundancy. The following subsection uses the above concepts and principles to identify the design parameters, and applies them to develop a best-practice hierarchical network with the highest availability.

Cisco FlexStack

The next-generation Catalyst 2960-S Series Layer 2 access-switch introduces high-speed, low-latency stacking capability based on “pay as you grow” model. Following the Catalyst 3750-X StackWise Plus success, the Catalyst 2960-S model offers high availability, increased port-density with unified single control-plane and management to reduce the cost for small enterprise network. However, the architecture of FlexStack on Catalyst 2960-S Series platform differs from StackWise Plus. The Cisco FlexStack is comprised with hardware module and software capabilities. The FlexStack module must be installed on each Catalyst 2960-S switches that are intended to be deployed in stack-group. Cisco FlexStack module is hot-swappable module providing flexibility to deploy FlexStack without impacting business network operation.

Cisco FlexStack allows up to four Catalyst 2960-S Series switches into a single stacking group; it is recommended to deploy each switch member with dual FlexStack cable to provide increased 20G bidirectional stack bandwidth capacity and FlexStack redundancy. The FlexStack protocol dynamically detects switch member and allows it to join the stack group if all stacking criteria is met. The unique data forwarding architecture and FlexStack QoS is on per-hop basis, the unknown unicast, broadcast, and multicast traffic will be flooded between stack group switch members. The FlexStack protocol detects and breaks the loop between the FlexStack group switches. Once the destination switch member is determined, Catalyst 2960-S use shortest egress stack port path to forward traffic. Any packet traverses across FlexStack is encapsulated with 32 bytes of FlexStack header carrying unique information to provide centralized control-plane and distributed forwarding design.

Cisco Modular Switch

The Cisco Catalyst 4500-E modular switch supports redundant supervisors, and Stateful Switch Over (SSO). When deployed along with NSF, the 4500-E provides a enterprise-class highly available system with network and device resiliency.

SSO is a Cisco IOS service used to synchronize critical forwarding and protocol state information between redundant supervisors configured in a single chassis. With SSO enabled, one supervisor in the system assumes the role of *active* and the other supervisor becomes the *hot-standby*. Each is ready to backup the other, thus providing 1:1 hot redundancy to protect from a control-plane outage. Since both supervisors are active, the system benefits by using the physical ports from both supervisors during normal operation. SSO synchronizes system services such as DHCP snooping, Switched Port Analyzer (SPAN), security access control lists (ACLs), and QoS policies so ensure the switch provides the same level of protection and service after a supervisor failover event.

NSF enables packets to continue to be forwarded using existing routing table information, during switchover. NSF also provides graceful restart to the routing protocol such that during the failover, the routing protocol remains aware of the change and does not react by resetting its adjacency. If the routing protocol were to react to the failure event, and alter routing path information, the effectiveness of stateful switch over would be diminished.

Operational Resiliency Strategy

Designing the network to recover from unplanned outages is important. It is also important to consider how to minimize the disruption caused by planned outages. These planned outages can be due to standard operational processes, configuration changes, software and hardware upgrades, etc.

The same redundant components which mitigate the impact of unplanned outages can also be used to minimize the disruption caused by planned outages. The ability to upgrade individual devices without taking them out of service is enabled by having internal component redundancy (such as with power supplies, and supervisors) complemented with the system software capabilities. Two primary mechanisms exist to upgrade software in a live network:

- Full-image In-Service Software Upgrade (ISSU) on the Cisco Catalyst 4500-E leverages dual supervisors to allow for a full, in-place Cisco IOS upgrade. This leverages the NSF/SSO capabilities of the switch and provides for less than 200 msec of traffic loss during a full Cisco IOS upgrade.
- Network and device level redundancy, along with the necessary software control mechanisms, guarantee controlled and fast recovery of all data flows following a fault condition, and provide the ability to manage the fault tolerant infrastructure during planned outage events.

Validating operational resiliency is beyond the scope of this design guide, refer to CCO documentation for deployment guidelines.

Deploying High Availability in Network

Many of the design features of the Small Enterprise Design Profile which were described in "Deploying Network Foundation Services" section on page 2-16, contribute to the network high availability capabilities. This section focuses on how to implement additional features which complete the high availability design in small enterprise network design.

Network Resiliency

Etherchannel and UDLD are two design features which are included in the network foundation services, which contribute to network resiliency.

Implementing IP Event Dampening

Poor signaling or a loose connection may cause continuous port-flap (port alternates between active state and inactive state). A single interface flapping can impact the stability and availability of the network. Route summarization is one technique which mitigates the impact of a flapping port. Summarization isolates the fault domain with a new metric announcement by the aggregator and thereby hides the local networks fault within the domain.

A best practice to mitigate local network domain instability due to port-flap, is implementing IP Event Dampening on all layer 3 interfaces. Each time the Layer-3 interface flaps the IP dampening tracks and records the flap event. Upon multiple flaps, a logical penalty is assigned to the port and suppresses link status notification to IP routing until the port becomes stable. IP event dampening is a local function and

does not have a signaling mechanism to communicate with a remote system. It can be implemented on each individual physical or logical Layer-3 interface: physical ports, SVI or port-channels. Following is a example configuration to implement IP Event Damenping:

Distribution/Core

```
cr24-4507-1(config)#int range Port1 , Gig5/6 , Gig6/6 , Vlan 101 - 110
cr24-4507-1(config-if-range)#dampening

    cr24-4507-1#show interface dampening | be Port
    Port-channel1 Connected to cr24-3750ME-1
    Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
      0       0   FALSE       0           5       1000       2000      20
16000       0
```

The following output illustrates how the IP event dampening keeps track of port flaps and makes a decision to notify IP routing process based on interface suppression status:

```
cr24-4507-1#debug dampening interface
cr24-4507-1#show logging | inc EvD|IF-EvD


12:32:03.274: EvD(GigabitEthernet5/6): charge penalty 1000, new accum. penalty 1000, flap
count 2
12:32:03.274: EvD(GigabitEthernet5/6): accum. penalty 1000, not suppressed
12:32:03.274: IF-EvD(GigabitEthernet5/6): update IP Routing state to DOWN, interface is
not suppressed

cr24-4507-1#show interface dampening | be 5/6
Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
2       0   FALSE       0           5       1000       2000      20           16000
0
```

In a multilayer access-distribution design, the Layer-2 and Layer-3 demarcation is at the collapsed core-distribution device. IP event dampening is enabled on per-logical VLAN (SVI) interface basis on the collapsed core device. IP event dampening becomes more effective when each access-layer switch is deployed with a unique set of Layer-2 VLANs.

Assigning unique VLANs on each access-layer switch also helps IP event dampening to isolate the problem and prevent network faults triggered in a multilayer network. The following output illustrates how IP event dampening keeps track of individual logical VLAN networks associated to same Layer-2 physical trunk ports. When a Layer-2 trunk port flaps, the state of SVI also flaps, and forces dampening to track and penalize unstable interfaces:

```
12:58:41.332: EvD(Vlan101): charge penalty 1000, new accum. penalty 2627, flap count 3
12:58:41.332: EvD(Vlan101): accum. penalty 2627, now suppressed with a reuse intervals of
7
12:58:41.332: IF-EvD(Vlan101): update IP Routing state to DOWN, interface is suppressed

    cr24-4507-1#show interface dampening
    Vlan101 Connected to cr24_2960_Dept_1_VLAN
    Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
    3       71   FALSE       0           5       1000       2000      20           16000
0
```

Device Resiliency

As described earlier, redundant hardware is an important technique for achieving device resiliency. The small enterprise network design applies hardware redundancy considering the cost/performance tradeoffs.

Implementing Redundant Power Supply

Redundant power supplies can prevent a system outage due to power outage, power supply or fan hardware failure. All Cisco Catalyst switching platforms supports robust 1+1 redundant power capabilities that can be deployed with internal or external power source management. See Figure 2-44.

Figure 2-44 Cisco Catalyst Internal and External Power Redundancy Option

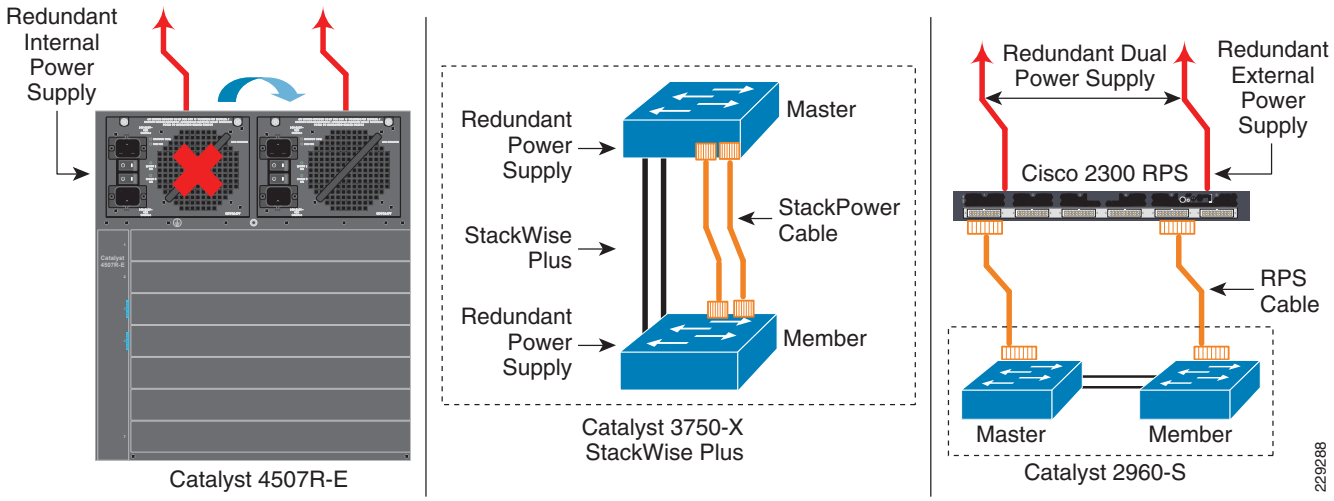

Catalyst 4500-E—Redundant Internal Power Supply

The Cisco Catalyst 4500-E provides power to internal hardware components and external devices like IP phones. All the power is provided by the internal power supply. Dual-power supplies in the Catalyst 4500-E can operate in one of two different modes:

- *Redundant Mode*—By default, Catalyst 4500-E power supply operates in redundant mode offering 1+1 redundant option. The system determines power capacity and number of power supplies required based on the power required for all internal and external power components. Both power supplies must have sufficient power to support all the installed modules and operate in 1+1 redundant mode.

```
cr24-4507-1(config)#power redundancy-mode redundant

cr24-4507-1#show power supplies
Power supplies needed by system    :1
Power supplies currently available :2
```

- *Combined Mode*—If the system power requirement exceeds the capacity of a single power supply, then both power supplies can be combined to increase the capacity. In this mode, the power system does not provide 1+1 power redundancy. The following global configuration will enable power supplies to operate in combined mode:

```
cr24-4507-1(config)#power redundancy-mode combined
```

```
cr24-4507-1#show power supplies
Power supplies needed by system:2
Power supplies currently available:2
```

Catalyst 3750-X— Cisco StackPower Redundancy

The next-generation Catalyst 3750-X Series platform introduces innovative Cisco StackPower technology to provide power redundancy solution for fixed configuration switches. Cisco StackPower unifies the individual power supplies installed in the switches and creates a pool of power, directing that power where it is needed. Up to four switches can be configured in a StackPower stack with the special Cisco proprietary StackPower cable. The StackPower cable is different than the StackWise data cables and is available on all Cisco Catalyst 3750-X models.

During individual power supply fault from the stack can regain power from global power pool to provide seamless operation in the network. With the modular power supply design in Catalyst 3750-X series platform the defective power supply can be swapped without disrupting network operation. The Cisco StackPower can be deployed in following two modes:

- *Sharing Mode*—All input power is available to be used for power loads. The total aggregated available power in all switches in the power stack (up to four) is treated as a single large power supply. All switches in stack can share power with available power to all powered devices connected to PoE ports. In this mode, the total available power is used for power budgeting decisions and no power is reserved to accommodate power-supply failures. If a power supply fails, powered devices and switches could be shut down. Default mode.
- *Redundant Mode*—The power from the largest power supply in the system is subtracted from the power budget, which reduces the total available power, but provides backup power in case of a power-supply failure. Although there is less available power in the pool for switches and powered devices to draw from, the possibility of having to shut down switches or powered devices in case of a power failure or extreme power load is reduced. It is recommended to budget the required power and deploy each Catalyst 3750-X switch in stack with dual power supply to meet the need. Enabling redundant mode will offer power redundancy as a backup during one of the power supply unit failure event.

Since Cisco StackWise Plus can group up to nine 3750-X Series switches in the stack-ring, the Cisco StackPower must be deployed with two power stack group to accommodate up to four switches. Following sample configuration demonstrate deploying Cisco StackPower redundancy mode and grouping the stack-member into power stack group, to make new power configuration effective, it is important that network administrator must plan downtime as all the switches in the stack ring must be reloaded:

```
cr36-3750X-xSB(config)# stack-power stack PowerStack
cr36-3750X-xSB(config-stackpower)#mode redundant

cr36-3750X-xSB(config)#stack-power switch 1
cr36-3750X-xSB(config-switch-stackpower)#stack-id PowerStack
%The change may not take effect until the entire data stack is reloaded

cr36-3750X-xSB(config)#stack-power switch 2
cr36-3750X-xSB(config-switch-stackpower)#stack-id PowerStack
%The change may not take effect until the entire data stack is reloaded
```

Catalyst 2960 (External Power Redundancy)

The Cisco Redundant Power Supply (RPS) 2300 can support up to six RPS ports to provide seamless power backup to critical access-layer switches in the campus network. Upto two devices can be protected by Cisco RPS 2300 against power or power supply failure. Additional power resiliency on Cisco RPS 2300 can be added by deploying dual power supply to backup to two devices simultaneously. Note that external power redundancy requires special RPS cable and specific 2960 models currently do not support external power redundancy. Deploying external power redundancy on Cisco Catalyst 2960 and 2960-S with Cisco RPS 2300 is performed automatically and do not require any extra configuration. Cisco Catalyst 3560-X and 3750-X switches can be used if RPS 2300 configuration is required.

Implementing Redundant Control Plane System

The collapsed core device in the main and remote sites (Catalyst 4500-E or 3750-X StackWise Plus) is deployed with redundant supervisor, or StackWise Plus to enable graceful recovery from switch hardware outage. Any access-switch which is deemed critical may be deployed as StackWise Plus and FlexStack to improve device resiliency. The implementation for each switch is different, and is discussed separately in the sections which follow.

Resilient Cisco FlexStack and StackWise Plus

Starting in Cisco IOS Release 12.2(53)SE1, Cisco Catalyst 2960-S supports FlexStack and it can be used when higher port-density, increased uplink bandwidth capacity, and the resiliency at Layer-2 access is required. Starting in Cisco IOS Release 12.2(53)SE2, Cisco Catalyst 3750-X supports StackWise Plus, and is used when a resilient Layer 2 or Layer 3 access-switch is required. Cisco 3750-X StackWise Plus is deployed for the collapsed core in the small remote site network.

Cisco Catalyst 3750-X and 2960-S switches are provisioning dynamically in the stack group by the StackWise or FlexStack protocols. Cisco IOS automatically adjusts the interface addressing and its associated configuration based on the number of provisioned switches in the stack.

```
cr26-2960s-1#show run | inc provision

switch 1 provision ws-c2960s-48ts-s
switch 2 provision ws-c2960s-48ts-s
```

Master Switch Election

The centralized control-plane and management plane is managed by the master switch in the stack. By default, the master switch selection within the ring is performed dynamically by negotiating several parameters and capabilities between each switch within the stack. Each StackWise-capable switch is by default configured with priority 1.

```
cr26-C2960S-1# show switch
Switch/Stack Mac Address : 0022.bdc4.1d80
                                            H/W   Current
Switch#  Role   Mac Address     Priority Version  State
-----------------------------------------------------------
*1       Master 0022.bdc4.1d80     1       1       Ready
 2       Member 0026.0ac1.3e00     1       1       Ready
```

As described in previous section, the Cisco StackWise Plus architecture is not SSO-capable. This means that all the centralized Layer-3 functions must be reestablished with the neighbor switch during a master-switch outage. To minimize the control-plane impact and improve network convergence, the Layer-3 uplinks should be diverse, originating from member switches instead of the master switch. The

default switch priority must be increased manually after identifying the master switch and switch number. The new switch priority becomes effective after switch reset. It is recommended to modify default switch priority on the Catalyst 2960-S FlexStack group as Catalyst 3750-X StackWise.

```
cr26-3750r-1(config)#switch 1 priority 15
Changing the Switch Priority of Switch Number 1 to 15
cr26-3750r-1(config)#switch 2 priority 14
Changing the Switch Priority of Switch Number 2 to 15

cr26-3750r-1#show switch
Switch/Stack Mac Address : 0023.eb7b.e580
                                           H/W   Current
Switch# RoleMac AddressPriorityVersionState
------------------------------------------------------------
  1         Member0023.eb7b.e580150       Ready
* 2         Master0026.5284.ec80140       Ready
```

StackWise Layer-3 MAC Management

To provide a single unified logical network view in the network, the MAC addresses of Layer-3 interfaces on the StackWise (physical, logical, SVIs, port channel) are derived from the Ethernet MAC address pool of the master switch in the stack. All the Layer-3 communication from the StackWise switch to the endpoints (like IP phone, PC, servers and core network system) is based on the MAC address pool of the master switch.

```
cr26-3750r-1#show switch
Switch/Stack Mac Address : 0026.5284.ec80
                                           H/W   Current
Switch# Role   Mac AddressPriority     Version       State
-----------------------------------------------------------------------------------------
    1Member0023.eb7b.e580      1                 0            Ready

* 2 Master  0026.5284.ec80      1              0           Ready

cr26-3750r-1#show version
.  .  .
Base ethernet MAC Address       : 00:26:52:84:EC:80
.  .  .
```

After a master-switch outage, the new master switch in the stack assigns new MAC addresses to all Layer-3 interfaces, from the local MAC address pool. Once the new MAC address is assigned, it will force the switch to generate a gratuitous ARP in the network to make sure no other system is using the same MAC address. The default timer to retain the MAC address from the failed master switch is four minutes. While the new MAC address is not assigned on Layer-3 interface and not being propagated and updated in the network, the traffic will blackhole in the network.

```
cr26-3750r-1#reload slot 2
Proceed with reload? [confirm]
Switch 2 reloading...

cr26-3750r-1#show switch
Switch/Stack Mac Address : 0023.eb7b.e580
                                           H/W   Current
Switch# Role   Mac AddressPriority     Version       State
-----------------------------------------------------------------------------------------
* 1 Master  0023.eb7b.e580      1                 1            Ready
   2Member000.0000.0000      0              1           Removed
```

To prevent this network instability, the old MAC address assignments on Layer-3 interfaces can be retained even after the master switch fails. The new active master switch can continue to use the MAC addresses assigned by the old master switch, which prevents ARP and routing outages in the network. The default stack-mac timer settings must be changed in Cisco Catalyst 2960-S FlexStack and 3750-E StackWise Plus switch mode using the global configuration CLI mode as shown below:

```
cr26-3750r-1(config)#stack-mac persistent timer 0

cr26-3750r-1#show switch
Switch/Stack Mac Address : 0023.eb7b.e580
Mac persistency wait time: Indefinite

                                             H/W   Current
Switch# Role   Mac AddressPriority     Version       State
-----------------------------------------------------------------------------------------
    1Master0023.eb7b.e580      1                 0            Ready
* 2 Master  0026.5284.ec80      1              0           Ready
```

Non-Stop Forwarding (NSF)

The Cisco Catalyst 3750-X switch in StackWise Plus mode is not SSO-capable. When the master switch fails, the new master switch is required to reform the Layer-3 adjacencies with the neighbors in the network. The forwarding architecture in StackWise switch is designed to provide non-stop forwarding during the master switch outage using NSF technology. Each 3750-X switch in the stack maintains distributed Layer-3 FIB from the old master switch and continues to forward upstream traffic, until they are updated by the new master in the stack ring.

To enable NSF capability, explicit configuration must be enabled under the routing process. NSF-aware feature is enabled by default on all Layer-3 Ethernet switches to function in helper mode to perform graceful recovery during NSF-capable Cisco 3750-X master switch outage. NSF-capable system can also operate in NSF aware role:

NSF Capable Layer-3 Switch

```
cr36-3750s-1(config)#router eigrp 100
cr36-3750s-1(config-router)#nsf




cr36-3750s-1#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
  EIGRP NSF enabled
     NSF signal timer is 20s
     NSF converge timer is 120s
```

NSF-Aware Layer-3 Switch

```
cr24-3560r-1#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
```

NSF Timers

As depicted in the above **show** commands, the default NSF-aware system hold timer is 240 seconds. Lowering the timer value may abruptly terminate graceful recovery, causing network instability. Best practice is to use the default NSF hold timer, unless it is observed that NSF recovery takes longer than 240 seconds.

600 seconds after a graceful-recovery starts on a NSF-aware system, NSF clears the route stale marking and resumes using the synchronized routing database.

```
! NSF Aware received graceful-restart message from new master switch
11:56:15.365: %DUAL-5-NBRCHANGE: EIGRP-IPv4:(100) 100: Neighbor 10.125.32.3
(Port-channel15) is resync: peer graceful-restart
11:56:15.365: EIGRP: NSF: AS100, NSF or GR initiated by 10.125.32.3 at 00:00:00, flags 0x4

! NSF route hold timer expires and searches and removes all stale route entries
received graceful-restart message from new master switch
12:00:15.392: EIGRP: NSF: AS100. route hold timer expiry
12:00:15.392: DUAL: Search for outdated routes from 10.125.32.3
```

Resilient Cisco Catalyst 4500-E

A modular switching platform like the Cisco Catalyst 4507R-E is fully NSF/SSO-capable, providing 1+1 control plane redundancy. In the Catalyst 4507R-E, all the intelligent Layer-2 and Layer-3 functions are performed centrally on the supervisor module. Deploying redundant supervisor in SSO mode in same system will allow the primary supervisor to fully synchronize the adjacencies, forwarding, configuration, counters, and more information on redundant hot-standby supervisor.

The Cisco Catalyst 4507R-E ports are independent of the supervisor state. Because of this hardware design, during a supervisor switchover, the ports connected to the failed supervisor do not go down. Because paths and ports are not down, hardware keeps forwarding the packet to a valid next-hop while supervisor switchover is occurring.

The configuration and implementation guidelines for implementing NSF/SSO on the Cisco Catalyst 4507R-E are the same for main and remote site network designs.

Increasing Supervisor Uplink Port Availability

There are restrictions on which supervisor uplink ports can be actively configured. Multiple ports can be simultaneously active on the supervisor. However Cisco IOS Release 12.2(25)SG or later is required for concurrent use of both 10G and 1G. The Small Enterprise Design Profile uses the 1G interface to connect to the Cisco 3750-MetroE WAN aggregation switch. To use 10G port in 1G mode with redundancy, the following configuration must be applied on collapsed core Catalyst 4507R-E switch:

```
cr24-4507-1(config)#hw-module uplink mode shared-backplane
cr24-4507-1(config)#hw-module module 3 port-group 1 select gigabitethernet
cr24-4507-1(config)#hw-module module 4 port-group 1 select gigabitethernet

cr24-4507-1#show hw-module uplink
 Active uplink mode configuration is Shared-backplane
```

Stateful Switchover (SSO)

SSO redundancy mode in the Cisco Catalyst 4507R-E supervisor is turned on by default starting with Cisco IOS Release 12.2(20)EWA. To provide 1+1 redundancy, all the technical specifications between active and standby supervisor must be identical. Also note that the Cisco Catalyst 4507R-E and 4510R-E are the only models that support supervisor redundancy. SSO is supported on all supervisors running IOS except Sup II-Plus-TS. The NSF-awareness feature is supported by all the supervisors supporting EIGRP, OSPF, IS-IS, and BGP routing protocols, while the NSF-capable feature is supported only on supervisors IV, V, and V-10G. NSF/SSO on Catalyst 4500-E requires a minimum boot ROM version and must be the same on both supervisors. For additional details on hardware requirements, refer to the Release Notes at the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/53SG/configuration/NSFwSSO.html#wp1135767

```
cr24-4507-1(config)#redundancy
cr24-4507-1(config-red)# mode sso
cr24-4507-1(config-red)# main-cpu
cr24-4507-1(config-r-mc)# auto-sync standard

cr24-4507-1#show module | inc Chassis | 6-E | SSO
Chassis Type : WS-C4507R-E


 3     6  Sup 6-E 10GE (X2), 1000BaseX (SFP)     WS-X45-SUP6-E      JAE1132SXQ3
 4     6  Sup 6-E 10GE (X2), 1000BaseX (SFP)     WS-X45-SUP6-E      JAE1132SXRQ

 3   Active Supervisor   SSOActive
 4   Standby Supervisor  SSOStandby hot
```

The active supervisor dynamically detects the secondary supervisor installed in the same chassis and initiates several SSO dependency configuration checks. If the SSO dependency check fails, then the standby supervisor falls back into RPR mode. For example, IOS release mismatch between two supervisors may not allow SSO to synchronize.

If the SSO dependency configuration checks successfully pass, then SSO communication between both supervisors goes through several synchronization states before it transitions to hot-standby state as illustrated in the following output:

```
cr24-4507-1#show redundancy states
my state = 13 -ACTIVE
peer state = 8  -STANDBY HOT
. . .

Redundancy Mode (Operational) = Stateful Switchover
Redundancy Mode (Configured)  = Stateful Switchover
Redundancy State              = Stateful Switchover
    Maintenance Mode = Disabled
    Manual Swact = enabled
    Communications = Up
    . . .
```

All the state-machines and dynamic information of SSO-capable protocols are automatically synchronized to the standby supervisor module without any additional operational requirement. The hot-standby supervisor takes over the ownership of control-plane process when the active supervisor outage or removal from the chassis is detected.

Non-Stop Forwarding (NSF)

All the state-machines and dynamic information of SSO-capable protocols are automatically synchronized to the standby supervisor module. The hot-standby supervisor takes over the ownership of control-plane process if the active supervisor suffers an outage or is removed from the chassis.

NSF-Capable Layer 3 Switch

```
cr24-4507-1(config)#router eigrp 100
cr24-4507-1 (config-router)#nsf

cr24-4507-1#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
  EIGRP NSF enabled
     NSF signal timer is 20s
     NSF converge timer is 120s
```

NSF-Aware Layer 3 Switch

```
cr24-3560r-1#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
```

Summary

Designing the LAN network aspects for the small enterprise network design establishes the foundation for all other aspects within the design profile including WAN, Security, and Mobility.

This chapter reviews LAN design models recommended by Cisco, as well as where to apply these models within the various locations of a small enterprise network. Each of the layers is discussed and design guidance is provided on where to place and how to deploy these layers. Finally, key network foundation services such as routing, switching, QoS, multicast, and high availability best practices are given for the entire small enterprise design.

CHAPTER 3

Small Enterprise Design Profile(SEDP)—WAN Design

This chapter discusses how to design and deploy WAN architecture for Small Enterprise Design Profile. The primary components of the WAN architecture are as follows:

- WAN technology
- Bandwidth capacity planning
- WAN IP addressing structure
- Routing
- QoS

WAN Design

The success of Small Enterprise Design depends on having an effective and better collaborative environment between employees located at different geographic locations. Some of the technologies that enhance this environment are interactive-video, on-demand video, voice and web collaboration, video to mobile devices, and TelePresence. To enable these technologies that foster the collaborative environment, WAN design is a critical consideration. There are several WAN technologies available today to provide WAN services, such as MPLS/VPN, Internet, and Metro Ethernet. See Figure 3-1.

Figure 3-1 WAN Technology Representation

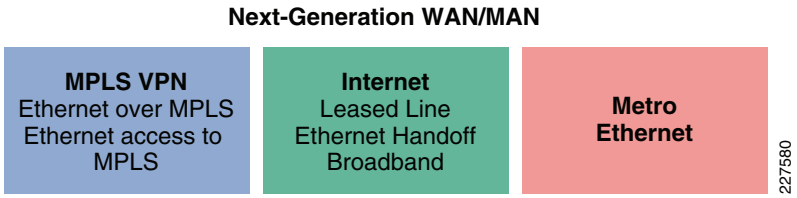

MPLS/VPN

The MPLS/VPN provides Layer-2 or Layer-3 VPN services. It provides the capability to an IP network infrastructure that delivers private network services over a shared network infrastructure. For more information about deploying MPLS/VPN, refer to the following URL: www.cisco.com/go/srnd

Internet

Out of the three WAN technologies, Internet is the cheapest and easiest to deploy. However, to deploy a VPN service over Internet requires an overlay VPN network such as DMVPN to provide secure VPN service.

Metro Ethernet

Metro Ethernet is one of the fastest growing transport technologies in the telecommunications industry. In the current WAN design, it is recommended to use Metro Ethernet Service as WAN transport between remote sites and main office. The following subsection describes the advantages of Metro service.

Why Metro Ethernet Service Is Needed

The deployment of carrier Ethernet services had raced to an astounding $7 billion by 2007* with predictions of more than $30 billion by 2012. This is driven by the following benefits to end users (IT, network, and applications departments) and service providers alike:

- *Scalability, ubiquity, and reachability*
  - Global availability of standardized services independent of physical access type dramatically reduce complexity and cost
- *Performance, QoS, and suitability for convergence*
  - Inherently, Ethernet networks require less processing to operate and manage at higher bandwidth than other technologies
  - Low latency and delay variation make it the best solution for video, voice, and data
- *Cost savings*
  - Carrier Ethernet brings the cost model of Ethernet to the wide area network WAN)
- *Control, simplicity, familiarity*
  - IT departments manage Ethernet networks every day and now are in control of their IP routed networks, worldwide
- *Expediting and enabling new applications*
  - Accelerates implementations with reduced resources for overburdened IT departments
  - Enables new applications requiring high bandwidth and low latency that were previously not possible or prohibited by high cost

Types of Services Available in Metro Ethernet

The following are two popular methods of Metro service:

- E-line, also known as Ethernet Virtual Private Line (EVPL), provides a point-to-point service.
- E-LAN provides multipoint or any to any connectivity.

EVPL, like Frame Relay, provides multiplexing multiple point-to-point connections over a single physical link. In the case of Frame Relay, the access link is a serial interface to a Frame Relay switch with individual data-link connection identifiers (DLCIs) identifying the multiple virtual circuits or connections. In the case of EVPL, the physical link is Ethernet, typically FastEthernet or Gigabit Ethernet, and the multiple circuits are identified as VLANs by way of an 802.1q trunk.

One of the major advantage E-LAN provides is any-to-any connectivity within the Metro area, which allows flexibility. It passes 802.q trunks across the SP network known as Q-in-Q. Figure 3-2 shows the difference between E-line and E-LAN.

Figure 3-2 Different Services Available

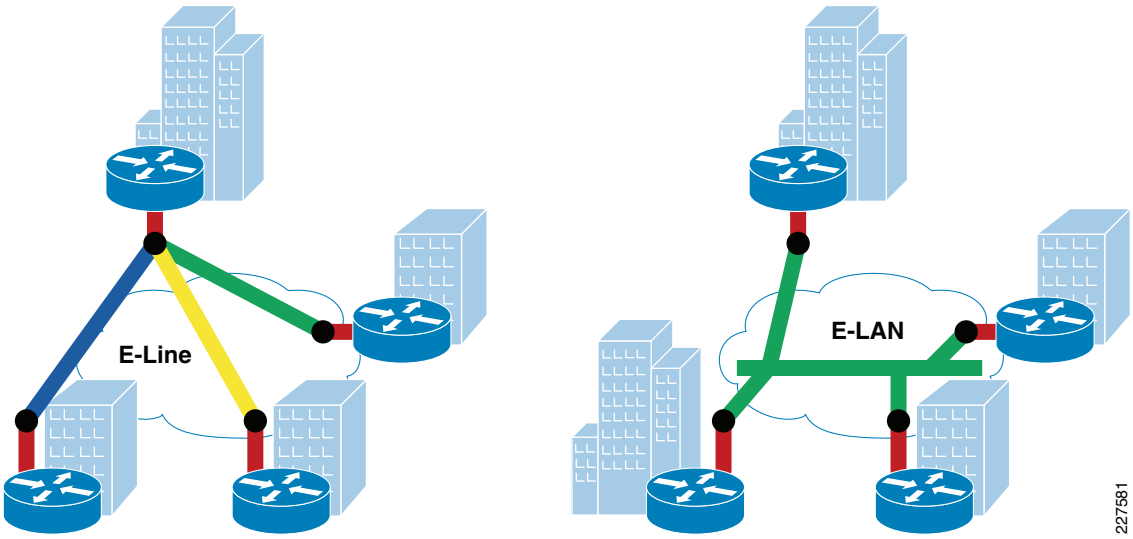

WAN Service Deployed in the Small Enterprise Design

In this version of the guide, E-line with point-to-point service is chosen as the model for connecting remote sites to main site. All the remote sites have a point-to-point connection to the main site. Each remote site is assumed to have 100Mpbs of Metro service to the service provider. As mentioned in the previous section, each circuit is represented by a VLAN using dot1q trunk. Figure 3-3 shows how this is implemented in this design.

Figure 3-3 EVPN Service Used in Remote WAN Architecture

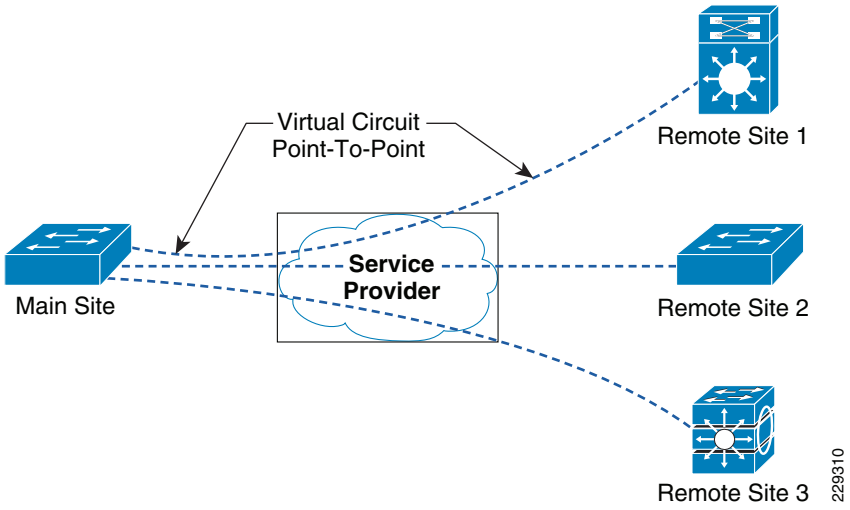

The following is the configuration of the WAN interface at the main site:

```
interface GigabitEthernet1/1/1
 description Connected to SP-MPLS-Core-cr24-6500-1
 switchport trunk native vlan 801
 switchport trunk allowed vlan 501-550
 switchport mode trunk
 logging event trunk-status
 load-interval 30
 carrier-delay msec 0
 priority-queue out
 mls qos trust dscp
 spanning-tree portfast trunk
 spanning-tree bpdufilter enable
 spanning-tree guard root
 service-policy output Remote-1to50-Parent-Policy-Map
 hold-queue 2000 in
 hold-queue 2000 out
```

As shown in the above configuration, the link is carrying 50 VLANs, which are connected to 50 small office sites. The solution shown above is one of the ways to deploy WAN service; there are other ways to deploy the WAN service. The next section discusses when to look at alternative to Metro service.

When to Consider WAN Technologies

In this design, all the remote sites are connected to the main site using point-to-point links that provide high bandwidth, symmetric traffic flows, and rich media services to all the sites. Since it is a point-to-point circuit, when the number of circuits increases to 1K and beyond, then we need not only to have 1K point-to-point circuits at the main site but also have 1K routing adjacency relationships, which will increase the CPU utilization. Therefore, when the scalability numbers increase more than 1K, an alternative Layer-3 WAN technology such as MPLS/VPN should be considered. MPLS/VPN Layer-3 technology deployments will not require different hardware, but only changing the WAN service. Figure 3-4 shows the routing behavior difference between Layer 2 (Metro service) and Layer 3 (MPLS/VPN service) from the routing protocol perspective.

Figure 3-4 Difference in Routing Adjacency Between Layer 2 and Layer 3 Service

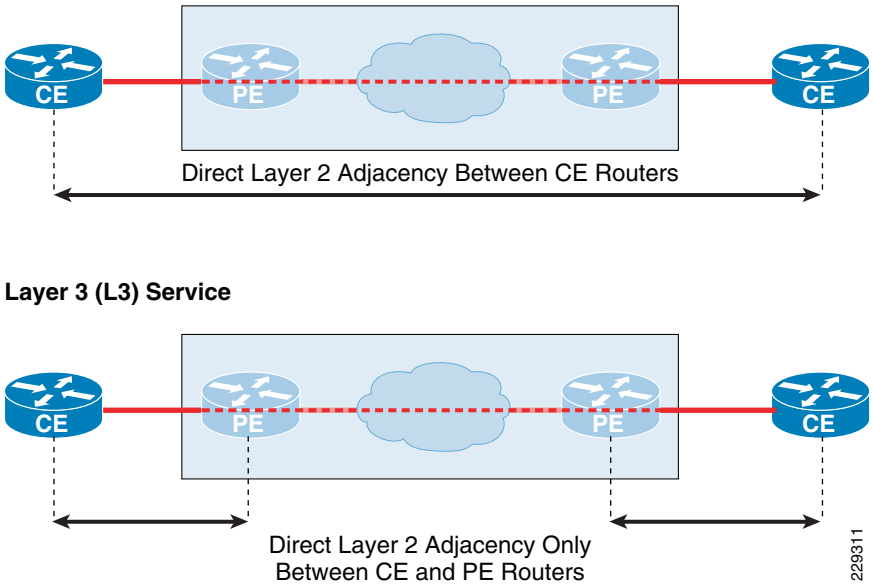

Bandwidth Capacity

This is critical component of the overall WAN design architecture because the application performance largely depends upon guaranteed level of bandwidth at remote sites and the main site. This section primarily discusses the general WAN bandwidth capacity planning and implementation steps.

Planning

To start with, each site must purchase the Metro Ethernet service from local service provider (SP) with bandwidth capacity that can meet the current minimum network load and provides enough flexibility to scale higher in the future, when additional applications need to be added. This bandwidth must be shared based on 4 class QoS model as described in the "Deploying QoS in Network" section on page 2-51 for optimal delivery. Logical bandwidth assignment to each circuit must be symmetric at the remote small office sites and at the main site. A mismatch in bandwidth capacity will force the traffic drop in the SP core due to inconsistent bandwidth service-level agreement.

This design requirement remains consistent for all the remote sites. Depending on large, medium, or small size, the WAN bandwidth capacity may also vary; however, the bandwidth sharing, routing, QoS, multicast etc. principles remains consistent.

Note Main site WAN router is an aggregator that connects multiple small office sites logically over a single media. This media type *must* be GigE, because the platform is 3750ME and the WAN egress port is a non-negotiated fiber port.

Calculating the optimum guaranteed bandwidth required at each remote site and the main site must be done by considering the following factors:

- Number of remote sites
- Bandwidth required at each location
- Platforms scalability limits, mainly, at main office

To explain how much bandwidth is required at each remote site, two terminologies (CIR and interface bandwidth) are used:

- *CIR*—The committed bandwidth that is guaranteed from the service provider, based on the logical connection.
- *Interface speed*—The interface speed is actual Ethernet handoff, which is 1Gbps for both the remote sites and the main site.

For example, consider a scenario where there are three remote sites and one main site. The CIR required at each remote site is 100Mbps and, because there are three remote sites, the main office needs three virtual circuits each having a CIR of 100Mbs. This also means that the aggregate bandwidth at main office is 300Mbps. Figure 3-5 illustrates this example.

Figure 3-5 Bandwidth Capacity Planning for Three Remote Sites

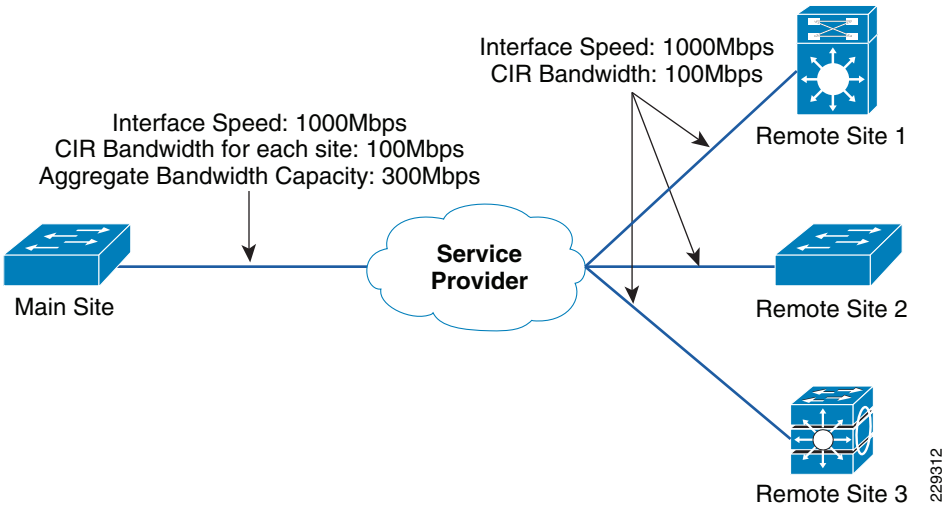

Similarly, if the CIR required at each remote site is 100Mbps and if there are 100 remote sites, then the bandwidth required at main office is 10Gbps. To support an aggregate CIR bandwidth of 10Gbps at main office, we need to think about scalability limits on the WAN aggregation switch. If the aggregated logical connection speed to the remote sites exceed the media capacity on 3750-ME, which is WAN aggregation switch for this design, then there are two approaches to mitigate the problem:

- *Migrate to a Modular switching platform*—Migrating to a modular switching platform for the WAN aggregation tier will enable wan capability to scale higher, reduce operational and management complexities and may offer better performance and resiliency depending on modular platform capabilities. When planning to migrate to a modular switching platform, please make sure that it supports the same QoS features of 3750ME, along with redundancy.
- *Deploy another 3750-ME in same tier*—Alternatively, it is possible to deploy secondary 3750-ME system to connect another set of remote networks. Deploying another set of 3750-ME does not change any WAN principles, In fact, except VLAN and IP address, all the configurations can be replicated to the secondary system.

Implementation

This section discusses how WAN bandwidth was implemented and validated at remote sites and main office of this design. The following are the considerations used to validate this model:

- The main office has dual WAN connections, and each connection is 1Gbps.
- On each WAN connection at the main office, there are 50 virtual circuits, each with CIR of 20Mbps.
- The remote sites have 1Gbps connection, and CIR value of 20Mbs on each circuit.

Figure 3-6 shows how this design is validated with 100 remote sites.

Figure 3-6 Bandwidth Capacity Planning for 100 Remote Sites

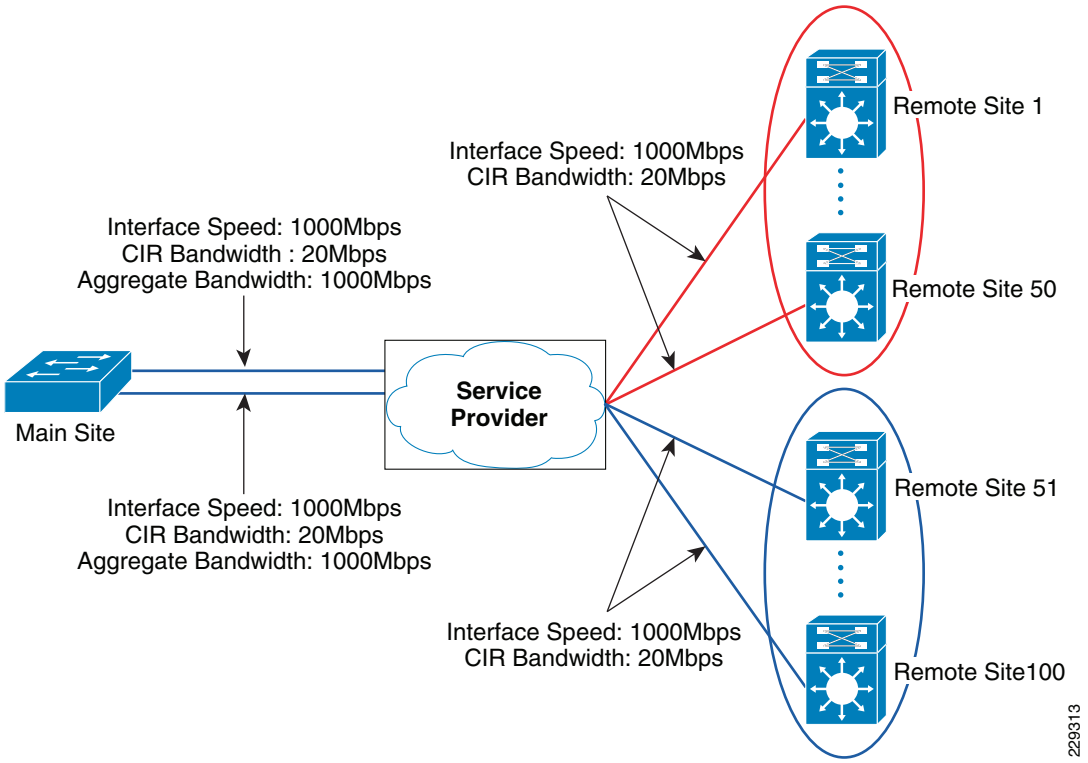

The next section discusses how to implement an IP addressing structure for the WAN interfaces.

Aggregation Structure

This section describes how to aggregate IP address space at remote sites and the main site on the WAN interface.

For designing IP address spaces for the LAN, refer to "Multicast IP Addressing" section on page 2-43.

As explained in the previous section, all the remote sites are connected to the main site using point-to-point links, which means that every remote site should be on a different subnet. One common method that some customers deploy is using /30 subnets at both ends. However, this would mean that every subnet uses up to four addresses. The recommendation is to use /31 subnets, this way only two addresses are used in every link. The following is a sample configuration for deploying this link at the remote site:

Remote Site	Main Office
interface Vlan501 description Connected to cr24-3750ME-DO dampening **ip address 10.127.0.1 255.255.255.254** no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.127.0.0 255.255.248.0 5 load-interval 30	interface Vlan501 description Connected to cr35-4507-SS1 dampening **ip address 10.127.0.0 255.255.255.254** no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.124.0.0 255.252.0.0 5 load-interval 30 hold-queue 2000 in hold-queue 2000 out

After planning an IP addressing scheme, the next part is to complete a routing design for the WAN connections, which is discussed in the following section.

Routing for WAN Connections

This section discusses how to implement routing on WAN interfaces. The key consideration while designing routing scheme is as follows:

- Summarization on network boundaries—This is very important aspect of the design as it prevents unnecessary routing updates to flow across the WAN interface when there is a link state change in the network. For example, let us consider the main site where there are subnets in the following range:

```
10.127.0.0/26
10.127.0.64/26
10.127.0.128/26
          .
          .
          .
10.127.7.64/26
```

Since the above subnets belong to a particular remote site (the main site pair, they could be summarized as 10.127.0.0/21. The following configuration shows how to perform summarization on the WAN interface, which is Vlan501 in this example:

Remote Site	Main Office
interface Vlan501 description Connected to cr24-3750ME-DO dampening ip address 10.1270.1 255.255.255.254 no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.127.0.0 255.255.248.0 5 load-interval 30	interface Vlan501 description Connected to cr35-4507-SS1 dampening ip address 10.1270.0 255.255.255.254 no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.124.0.0 255.252.0.0 5 load-interval 30 hold-queue 2000 in hold-queue 2000 out

The next most important component is QoS design, which the following section discusses.

QoS Design

This section discusses how QoS is implemented at the main office and remote sites. With Ethernet defined as the WAN medium, QoS design becomes most important because the router or switch on the WAN edge may believe that it has the complete line rate available for it to transmit. If so, the service provider would start dropping packets due to policers defined at its end. Figure 3-7 shows what may happen without a proper QoS design.

Figure 3-7 Policing at Service Provider Due to Lack of Proper QoS at Main Office and Remote Sites

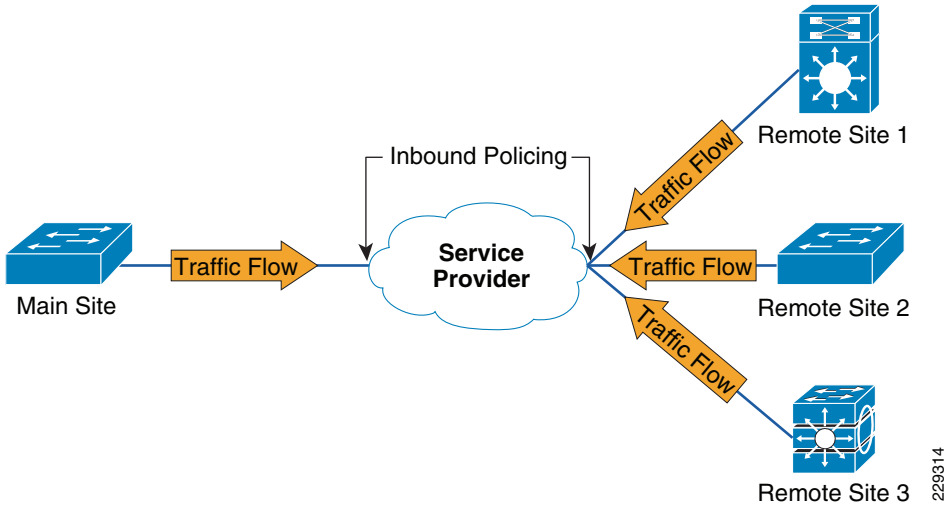

To prevent packets getting dropped at the service provider network, it is very important to have proper and consistent QoS policies at the remote sites and main site. Having a proper QoS policy would ensure that all the traffic is queued and serviced as per the bandwidth defined. This in turn ensures that packets will not be dropped by the service provider. The following subsections discuss how QoS policies are implemented at the remote and main sites.

QoS Policy at Main Site

As mentioned in the "When to Consider WAN Technologies" section on page 3-4, the main office has several point-to-point circuits to the remote sites, and the number (point-to-point circuits) depends upon the number of remote sites. Therefore, the QoS policy at the main site has the following objectives:

- The aggregate traffic going out to the remote site does not exceed 20Mbs.
- All the traffic going out is put in four classes.

To accomplish the above objectives, an hierarchical CBWFQ (HCBWFQ) is implemented. For more information about HCBWFQ, refer to "Chapter 4, Medianet QoS Design Considerations" of the *Medianet Reference Guide* at following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/Medianet_Ref_Gd/medianet_ref_gd.html

To implement HCWFQ, the following two policies are needed:

1. Parent policy that defines the aggregate shape rate.
2. Child policy that enables queuing within the shaped rate.

Figure 3-8 shows the representation of hierarchical policy.

Figure 3-8 Hierarchical Policy Implementation at Main Office

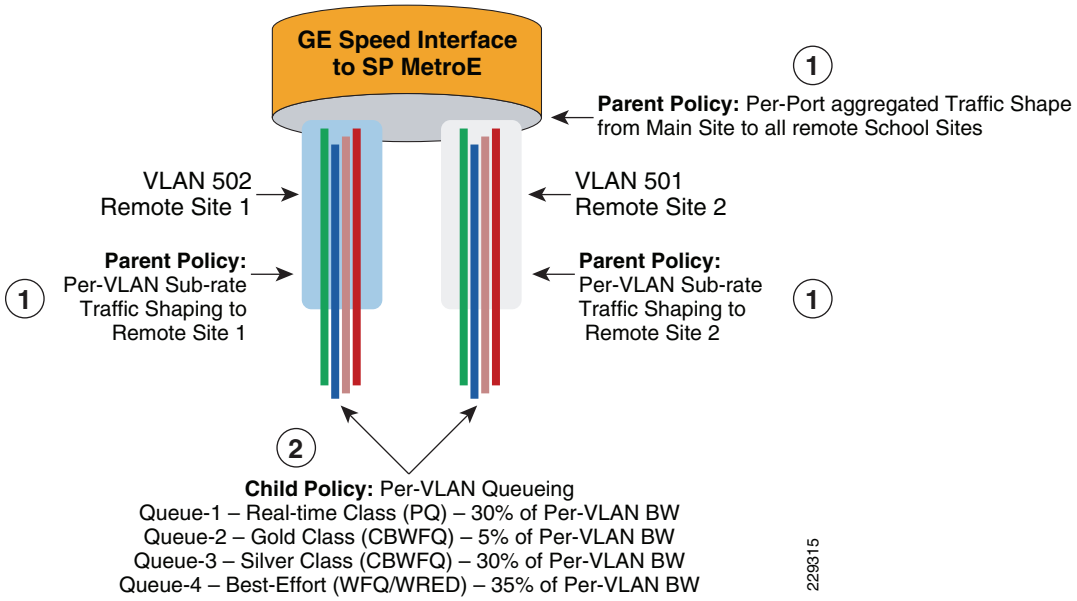

Table 3-1 shows how these classes are defined.

Table 3-1 QoS Policies

Class	Queuing type	Bandwidth Allocation
REAL_TIME	LLQ	30%
GOLD	CBWFQ	5%
SILVER	CBWFQ	30%
DEFAULT	CBWFQ	35%

The following is the configuration of the QoS policy at the main office. For brevity, only configuration for one VLAN is shown.

```
class-map match-all Remote_Site1 ← This class map would match the traffic going to a
Remote site
  description cr2-4507-SS1
 match vlan  501

policy-map Remote-Child-Policy-Map  ← This is child policy
 class REAL_TIME
    priority
   police cir percent 30 conform-action set-cos-transmit 5 exceed-action drop
violate-action drop
 class GOLD
    bandwidth percent 5
  set cos 3
```

```
 class SILVER
    bandwidth percent 30
  set cos 2
 class class-default
    bandwidth percent 35
  set cos 0
!
```

The following configuration is of the parent policy, which in this design shapes to 20Mbps for each remote site:

```
Policy Map Remote-1to50-Parent-Policy-Map ← This is parent policy map
    Class Remote_Site1
      shape average 20000000 (bits/sec)
      service-policy Remote-Child-Policy-Map ← This is child policy map
```

After defining the policies, they are applied to WAN interfaces. The following example shows the configuration of the Metro switch on its WAN interface:

```
interface GigabitEthernet1/1/1
 description Connected to SP-MPLS-Core-cr24-6500-1
 switchport trunk native vlan 801
 switchport trunk allowed vlan 501-550
 switchport mode trunk
 logging event trunk-status
 load-interval 30
 carrier-delay msec 0
 priority-queue out
 mls qos trust dscp
 spanning-tree portfast trunk
 spanning-tree bpdufilter enable
 spanning-tree guard root
 max-reserved-bandwidth 100
 service-policy output Remote-1to50-Parent-Policy-Map ← The policy-map
 hold-queue 2000 in
 hold-queue 2000 out
```

After completing the QoS policy at main office, we need to define the QoS policy at remote sites. The following subsection describes how to define QoS policy at the remote site.

QoS Policy at Remote Site

At the remote sites, the objectives are similar to the one at main site, which is to ensure that 20Mbps is the aggregate traffic leaving out the switch and the ingress traffic is queued in 4 classes. However, the implementation is different from the main office due to lack of hierarchical QoS support on the particular switch deployed at the remote site. It is however possible to have a line card with the capability to do HCBWFQ, if it is needed. The following configuration shows how to do QoS policy when the device does not support HCBWFQ.

The first step is to queue the ingress traffic in the four queues. Table 3-2 shows the queues and the bandwidth allocated.

Table 3-2 QoS Policies for Remote Site

Class	Queuing type	Bandwidth Allocation
REAL_TIME	Per-class	6mbps
GOLD	Per-class	1mbps

Table 3-2 QoS Policies for Remote Site

SILVER	Per-class	7mpbs
DEFAULT	Per-class	6mbps

The following is configuration of egress interface on the remote site:

```
interface GigabitEthernet1/1
 description Connected to MetroE-Core-cr25-6500-1
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 801
 switchport trunk allowed vlan 501
 switchport mode trunk
 logging event link-status
 load-interval 30
 carrier-delay msec 0
 qos trust dscp
 udld port disable
 tx-queue 1
   bandwidth 1 mbps
 tx-queue 2
   bandwidth 7 mbps
 tx-queue 3
   bandwidth 6 mbps
   priority high
 tx-queue 4
   bandwidth 6 mbps
 no cdp enable
 spanning-tree portfast trunk
 spanning-tree bpdufilter enable
 spanning-tree guard root
 service-policy output WAN-EGRESS-PARENT
```

The above configuration ensures that each class of the traffic is queued as shown in Table 3-2. However, the bandwidth only ensures the minimum amount of bandwidth available. It does not control the upper threshold, which needs to be 20Mbs in our solution. Therefore, to ensure that the traffic on the egress interface does not exceed 20Mps, WAN-EGRESS-PARENT policy is used to police the traffic to 20Mbps. The following is configuration of the WAN egress policy:

```
cr35-4507-SS1#show policy-map WAN-EGRESS-PARENT
  Policy Map WAN-EGRESS-PARENT
    Class class-default
      police 20 mbps 1000 byte conform-action transmit exceed-action drop
      service-policy WAN-EGRESS-CHILD
cr35-4507-SS1#
```

CHAPTER 4

Small Enterprise Design Profile(SEDP)—Wireless LAN Design

Cisco Unified Wireless Network Architecture

WLANs have emerged as one of the most effective means for connecting to a network, given the mobility of users. The Cisco Unified Wireless Network (CUWN) is a unified wired and wireless network solution that addresses the wireless network security, deployment, management, and control aspects of deploying a wireless network. It combines the best elements of wireless and wired networking to deliver secure, scalable wireless networks with a low total cost of ownership.

Figure 4-1 shows a high-level topology of the CUWN architecture, which includes Lightweight Access Point Protocol (LWAPP) access points (APs), mesh LWAPP APs (MAPs), the Wireless Control System (WCS), and the Wireless LAN Controller (WLC). Alternate WLC platforms include the Wireless LAN Controller Module (WLCM) or Wireless Services Module (WiSM). The Cisco Access Control Server (ACS) and its Authentication, Authorization, and Accounting (AAA) features complete the solution by providing RADIUS services in support of wireless user authentication and authorization.

Figure 4-1 Cisco Unified Wireless Network Architecture Overview

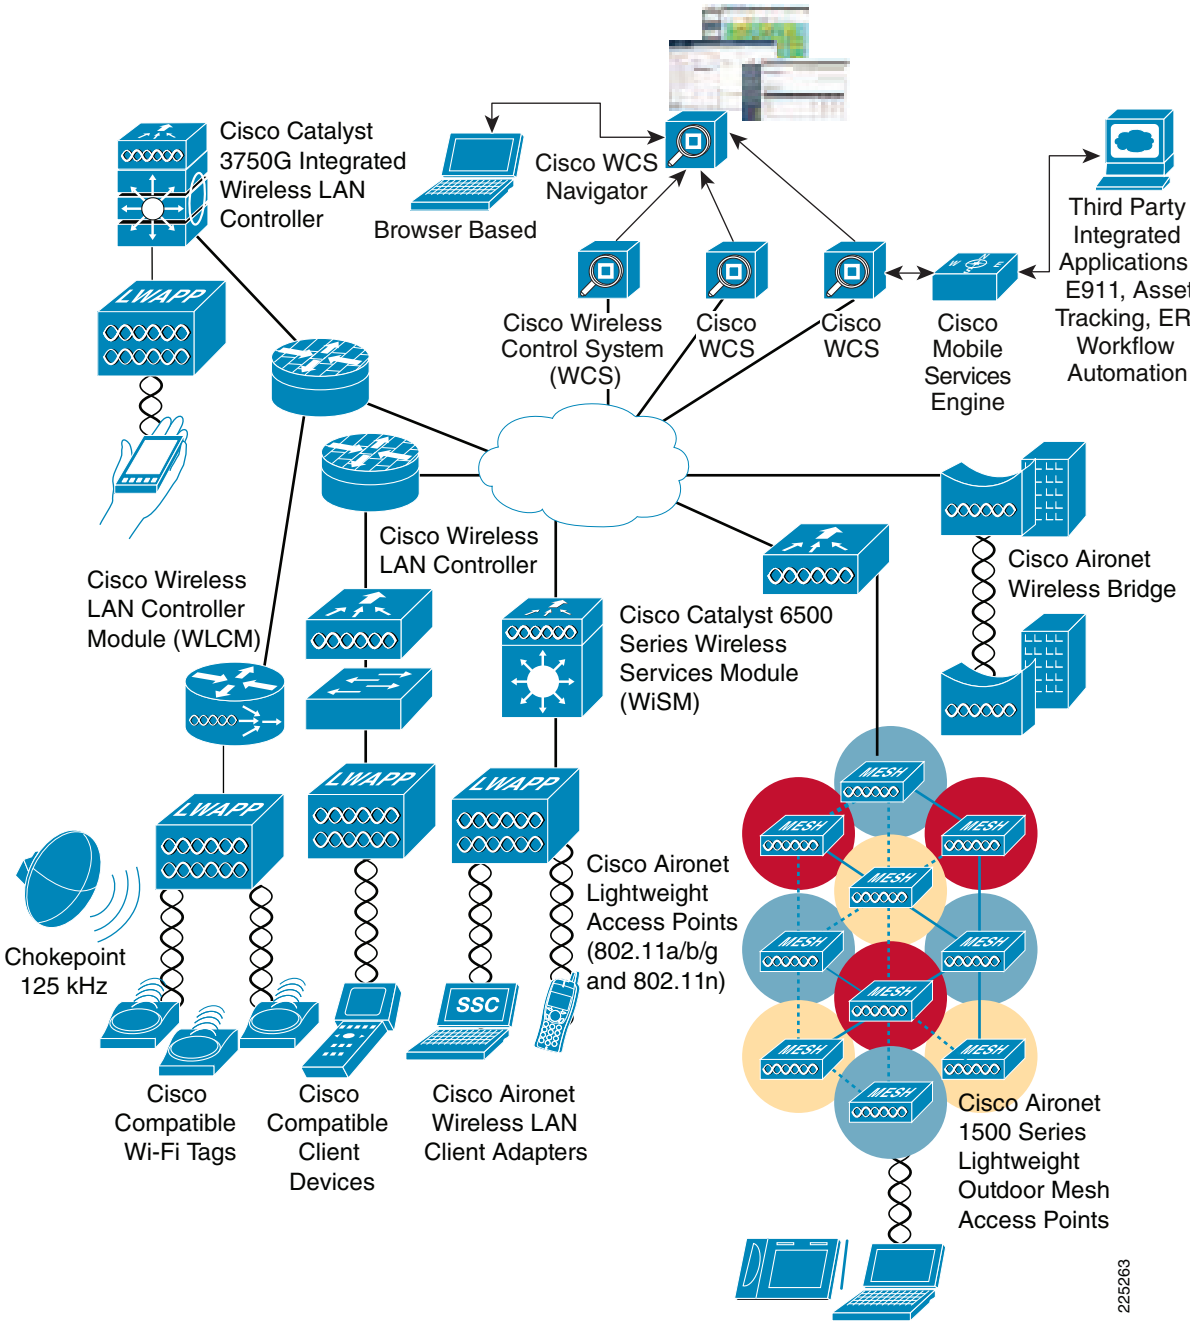

The Cisco Unified Wireless Network is composed of two key elements: Wireless LAN Controllers and Access Points (APs). These form the core of the Wireless LAN system, where the APs provide the radio connection between wireless clients and the network, and the WLCs provide network.

Note Figure 4-2 illustrates one of the primary features of the architecture: how LWAPP or Control and Provisioning of Wireless Access Points (CAPWAP) access points use the LWAPP/CAPWAP protocol to communicate with and tunnel traffic to a WLC.CUWN is migrating from the LWAPP protocol to CAPWAP, and the WLC software version in the Small Enterprise Design Profile uses CAPWAP. The

fundamentals of the architecture and operation are the same. Documents discussing the LWAPP architecture operation and behavior are still valid for CAPWAP, apart from the UDP port numbers. For the purposes of this document and other documents referring to LWAPP, the Cisco CAPWAP implementation can be considered as a superset of LWAPP features and behavior.

Figure 4-2 LAP and WLC Connection

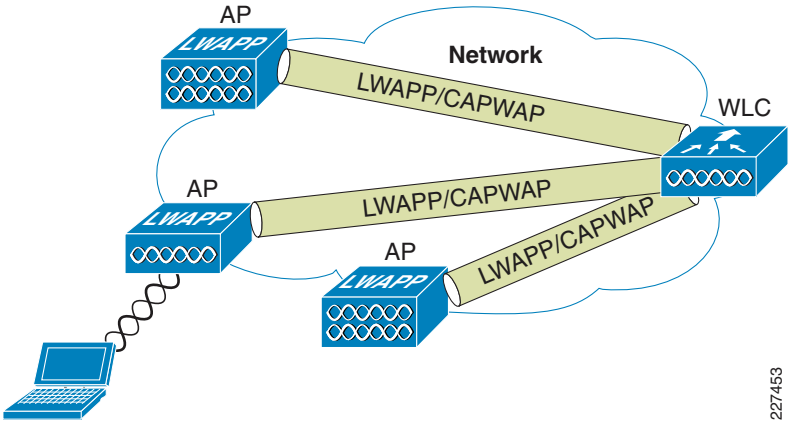

LWAPP/CAPWAP has three primary functions:

- Control and management of the LAP
- Tunneling of WLAN client traffic to the WLC
- Collection of 802.11 data for the management of the Cisco Unified Wireless System

LWAPP Features

The easier a system is to deploy and manage, the easier it will be to manage the security associated with that system. Early implementers of WLAN systems that used “fat” APs (autonomous or intelligent APs) found that the implementation and configuration of such APs was the equivalent of deploying and managing hundreds of individual firewalls, each requiring constant attention to ensure correct firmware, configuration, and safeguarding. Even worse, APs are often deployed in physically unsecured areas where theft of an AP could result in someone accessing its configuration to gain information to aid in some other form of malicious activity.

LWAPP addresses deployment, configuration, and physical security issues by doing the following:

- Removing direct user interaction and management of the AP. Instead, the AP is managed by the WLC through its LWAPP connection. This moves the configuration and firmware functions to the WLC, which can be further centralized through the use of the WCS.
- Having the AP download its configuration from the WLC, and be automatically updated when configuration changes occur on the WLC.
- Having the AP synchronize its firmware with its WLC, ensuring that the AP is always running the correct software version
- Storing sensitive configuration data at the WLC, and storing only IP address information on the AP. In this way, if the AP is physically compromised, there is no configuration information resident in NVRAM that can be used to perform further malicious activity.
- Mutually authenticating LAPs to WLCs, and AES encrypting the LWAPP control channel.

In addition to the improvements in physical security, firmware, and configuration management offered by LWAPP, the tunneling of WLAN traffic in an LWAPP-based architecture improves the ease of deployment without compromising the overall security of the solution. LAPs that support multiple WLAN VLANs can be deployed on access-layer switches without requiring dot1q trunking or additional client subnets at the access switches. All WLAN client traffic is tunneled to centralized locations (where the WLC resides), making it simpler to implement enterprise-wide WLAN access and security policies.

Small Enterprise Design Profile

Figure 4-3 shows a simple schematic of the CUWN integration into the small enterprise design profile. The key features of the CUWN integration is the use of a WLC at each location, with the management function (WCS) located at the main site. If context-aware services are implemented, the Cisco Mobility Services Engine (MSE) may be placed at the remote site; for smaller remote sites, an MSE at the main site may provide a centralized service.

The standalone WLCs used in this design support AP capacities from 12 to 250 APs per WLC, and multiple WLCs may be deployed at the same site if more than 250 APs are required or if a load sharing or higher availability WLAN solution is required. An alternate higher availability solution is to use a WLC at the main site as a backup WLC for the remote site's WLCs. This is known as an N+1 solution, where a main site WLC maintains sufficient capacity to support the APs of any individual remote site.

A similar principle to N+1 is used to provide high availability for the AAA service provided by the Cisco ACS server. Each remote site will have a local ACS server to provide AAA services, and use the main site ACS server as its secondary AAA server.

Figure 4-3 High level view of the CUWN Integration

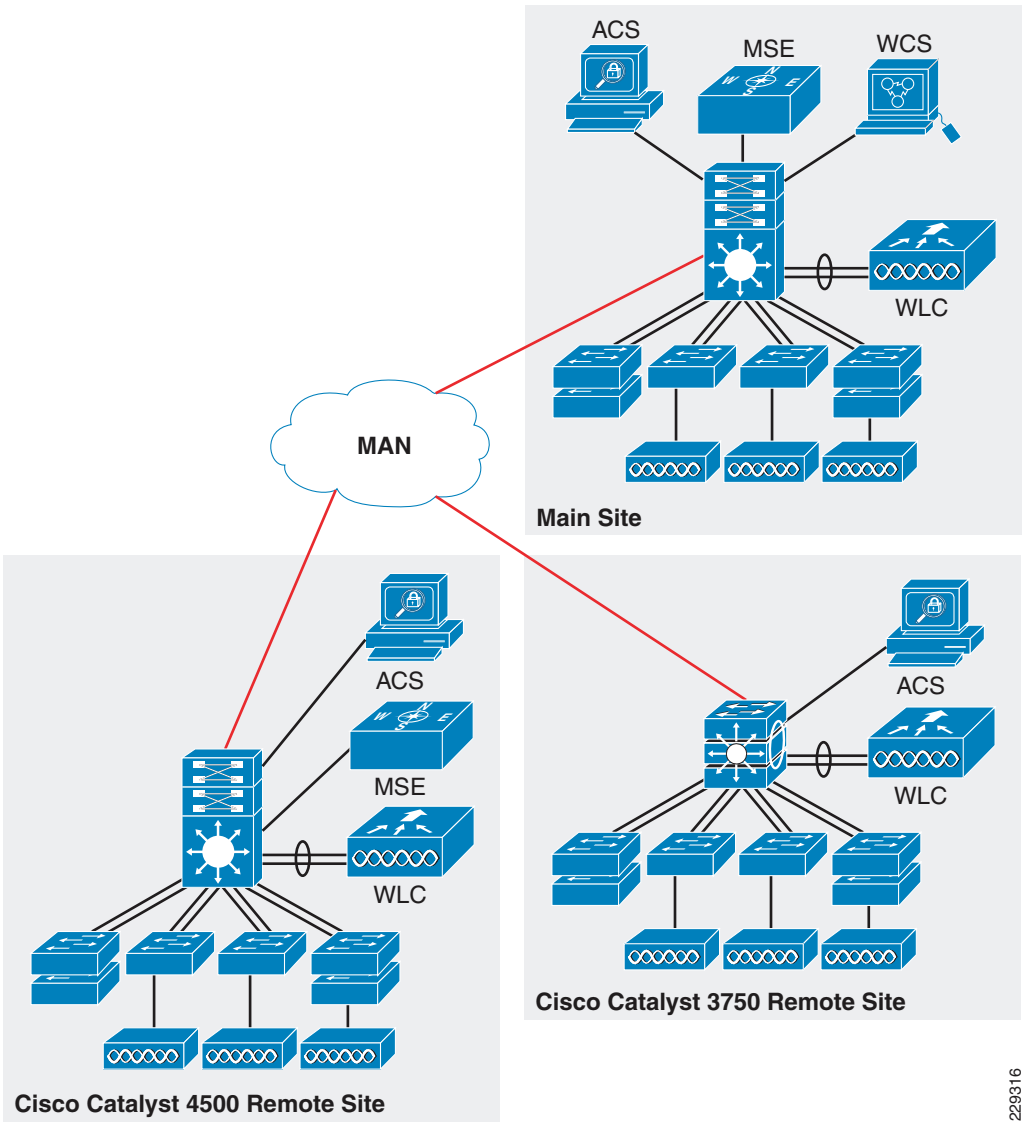

Management

Each WLCs has both a CLI and web interface to provide WLAN configuration and management features, but for a complete lifecycle management solution, the Cisco Wireless Control System (WCS) is needed. The WCS supports the delivery of high-performance applications and mission-critical solutions that simplify business operations and improve productivity. This comprehensive platform scales to meet the needs of small-, mid-, and large-scale wireless LANs across local, remote, national, and international locations. The WCS provides IT managers immediate access to the tools they need, when they need them, to more efficiently implement and maintain new or expanding WLANs—all from a centralized location requiring minimal IT staffing. Operational costs are significantly reduced through the Cisco WCS's intuitive GUI, simplified ease-of-use, and built-in tools that deliver improved IT efficiency,

lowered IT training costs, and minimized IT staffing requirements, even as the network grows. Cisco WCS lowers operational costs by incorporating the full breadth of management requirements, from radio frequency to controllers services, into a single unified platform.

The Cisco WCS scales to manage hundreds of Cisco wireless LAN controllers, which in turn can manage thousands of Cisco Aironet® access points, including the next-generation Cisco Aironet 1140 and 1250 Series 802.11n access points. For large-scale indoor and outdoor deployments, Cisco WCS Navigator can be included to simultaneously support up to 20 Cisco WCS platforms and 30,000 Cisco access points. Adding mobility services such as context-aware software and adaptive wireless intrusion prevention systems (wIPS) is simplified through Cisco WCS integration with the Cisco MSE.

Designing a wireless LAN that effectively supports business-critical data, voice, and video services is simplified with the Cisco WCS suite of built-in planning and design tools. Figure 4-4 shows an example of the simplified Wireless LAN Planning and Design Cisco WCS planning and design tools, simplify the process of defining access-point placement and determining access-point coverage areas for standard and irregularly shaped buildings. These tools give IT administrators clear visibility into the radio frequency (RF) environment. They make it easier to visualize the ideal RF environment, anticipate future coverage needs, and assess wireless LAN behavior. They help IT administrators reduce, and in many cases eliminate, improper RF designs and coverage problems that can lead to end-user trouble tickets. Specialized Cisco WCS planning tools enable real-time assessment of the WLAN's readiness to support voice-over-WLAN (VoWLAN) and context-aware (location) services. VoWLAN services support single and dual-mode Wi-Fi-enabled phones. Context-aware services use Cisco's patent pending “RF fingerprinting” technology to locate, track, and manage Wi-Fi-enabled devices and their contextual information in conjunction with Cisco MSE.

Figure 4-4 WCS Planning Tools

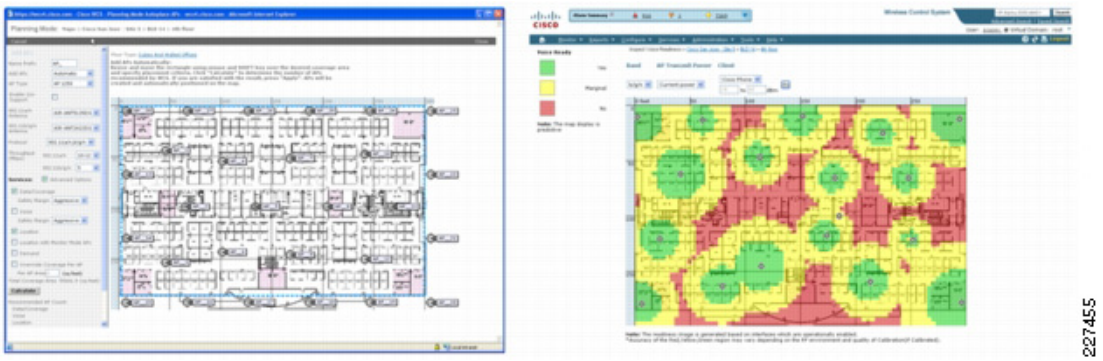

Getting the WLAN up and running quickly and cost-effectively to meet end-user needs is streamlined with the broad array of Cisco WCS integrated configuration templates. These easy-to-use templates and deployment tools help IT managers provision and configure the wireless LAN to expressly deliver the services that their business requires. Figure 4-5 shows an example of the Flexible Deployment Tools and Configuration Templates available through an easy-to-use interface, make it simple to apply common configurations across one or more wireless LAN controllers, regardless of their location in the network—whether on the same LAN as Cisco WCS, on separate routed subnets, or across a wide-area connection. At the click of a button, IT administrators can streamline even the most complex controller configurations, updates, and scheduling across the entire wireless network. Auto-provisioning access points is just as simple, with easy-to-use templates that support customized configuration of single or multiple access points.

Figure 4-5 WCS Deployment Templates

Cisco WCS is the ideal management platform for monitoring the entire WLAN to maintain robust performance and deliver an optimal wireless experience to mobile end users. Cisco WCS centralized interface makes it easy to access information where it is needed, when it is needed, on-demand or as scheduled. Figure 4-6 shows an example of the customizable dashboard and easy-to-use web-based interface. The Cisco WCS easy-to-use graphical displays contained within Cisco WCS serve as a starting point for maintenance, security, troubleshooting, and future capacity planning activities. Quick access to actionable data about healthy and unhealthy events occurring on the network is available from a variety of entry points, making Cisco WCS vital to ongoing network operations. The ever-present alarm summary in the Cisco WCS simplifies access to critical information, faults, and alarms based on their severity. Detecting, locating, and containing unauthorized (rogue) devices is fully supported when location services are enabled. Figure 4-7 shows an example of the ever-present alarm summary and simplified rogue device detection and location capabilities found within Cisco WCS.

Figure 4-6 WCS Monitoring Dashboard

Figure 4-7 WCS Alarm Panels

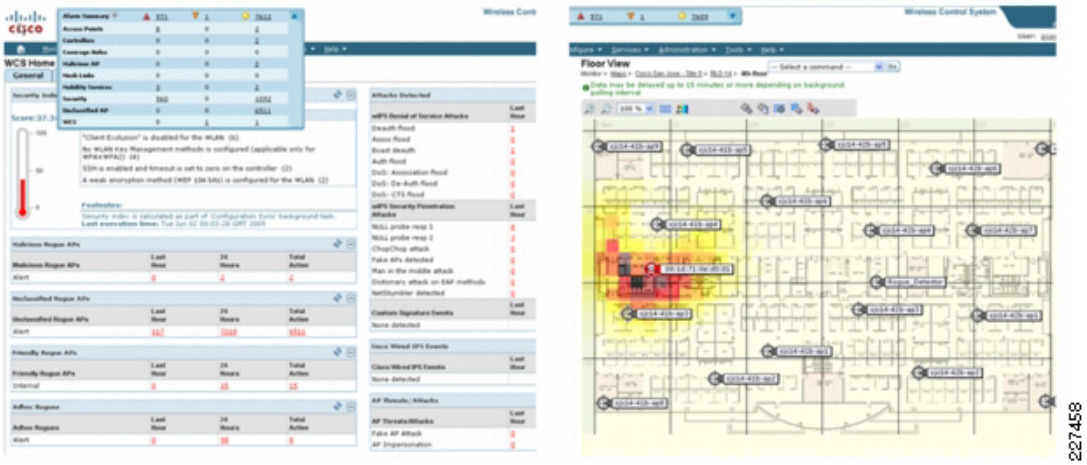

The integrated workflow and expansive array of troubleshooting tools in the Cisco WCS help IT administrators quickly identify, isolate, and resolve problems across all components of the Cisco Unified Wireless Network. Cisco WCS supports rapid troubleshooting of any size WLAN with minimal IT staffing. Figure 4-8 shows an example of the integrated workflows and troubleshooting tools found in Cisco WCS. Cisco WCS makes it easy to quickly assess service disruptions, receive notices about performance degradation, research resolutions, and take action to remedy non-optimal situations. Integrated workflows support seamless linkage between all tools, alarms, alerts, searches, and reports for all infrastructure components and client devices. A variety of tools work together to help IT administrators understand the operational nuances occurring on the WLAN and discover non-optimal events occurring outside baseline parameters such as client connection or roaming problems. The ever-present search tool in Cisco WCS facilitates cross-network access to real-time and historic information about devices and assets located anywhere in the wireless network. A built-in client troubleshooting tool provides a step-by-step method to analyze problems for all client devices. Cisco CleanAir supports finding, classifying, and correlating sources of interference from Wi-Fi and non-Wi-Fi sources such as Bluetooth devices and cordless phones.

Figure 4-8 WCS Troubleshooting Tools

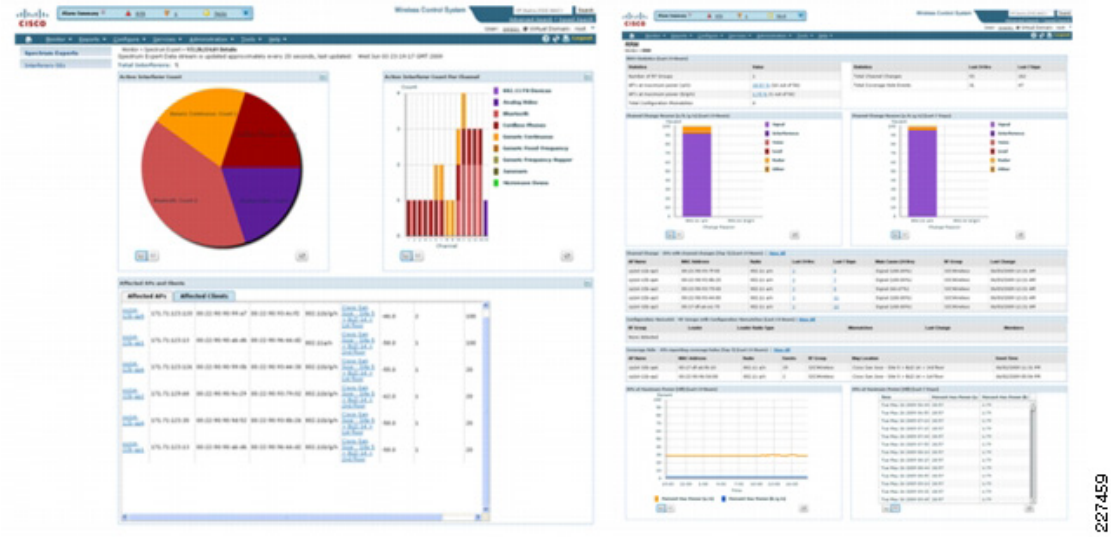

Cisco WCS includes customizable reporting that assists IT teams in more effectively managing, maintaining, and evolving the wireless LAN to meet ongoing business and operations requirements. Flexible reports provide access to the right data, at the right time, in a format to meet any requirement, as illustrated in Figure 4-9. An extensive variety of reports is available to help IT managers stay on top of network trends, maintain network control, audit operations, and quickly address changing business and end-user requirements. Reports are customizable based on user-defined parameters. Detailed analysis of what is going on, where and when in the network, as well as capacity planning, is simplified by collecting data from several reports and analyzing trends to understand how the WLAN has changed over time. Understanding WLAN trends makes it easier to plan for future enhancements and growth.

Figure 4-9 WCS Customizable Reports

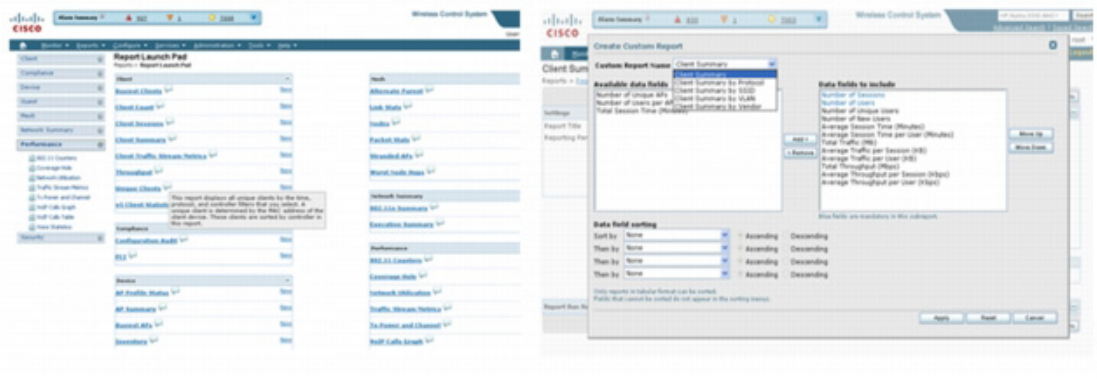

Connection to the Small Enterprise Design Profile Network

Figure 4-10 and Figure 4-11 show the remote site switch to WLC physical connection in more detail. A key feature of the WLC interface is its direct connection to the core distribution switch via a port-channel interface. This uses multiple Gigabit Ethernet connections from the WLC to the core/distribution switch. These Gigabit Ethernet connections are different line cards on switches or line card to ensure that a single switch or line card failure does not result in the loss of the WLC connection to the remote site network. The switch feature to achieve this is the same switch feature used for the Ether Channel connections between switches in the Small Enterprise Design Profile. The WLC feature is called *link aggregation* (LAG). LAG is disabled by default on the WLC and requires a WLC reboot to be enabled. This allows the WLC to use the same port channel configuration as the access switches when connecting to the core/distribution switch.

Figure 4-10 4500 Site Switch WLC Physical Connection

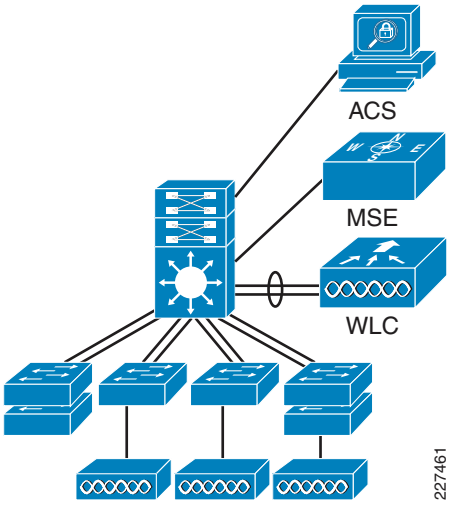

Figure 4-11 3750 Site Switch WLC Physical Connection

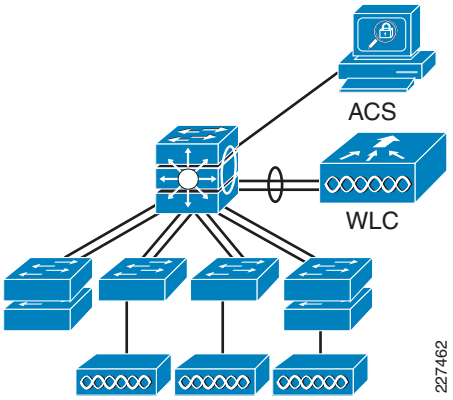

The WLC connects to the switch via a 802.1Q trunk connection, as shown in Figure 4-12. Multiple SVIs need to be configured on the switch to support the CUWN implementation. The key SVIs are an SVI for the management and AP manager interface of the WLC, and the SVIs for each of the different WLANs implemented on the WLC. There is not always a one-to-one relationship between SVIs and WLANs, but in most simple WLAN deployments this is the case.

Figure 4-12 Switch WLC Layer-2 Connection

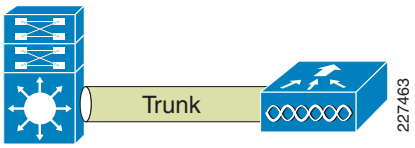

Figure 4-13 shows an example of the interface configuration summary on the remote site WLC. The key interfaces of interest are *ap-manager*, *manager*, and *wlan data1*, *wlan data2*, and *wlan voice1* interfaces.

The server port is an out-of-band management interface not used in this design guide. The virtual interface and its interface address are used to assist in the provisioning of seamless mobility. The virtual interface is assigned an address during the initial configuration of the WLC and this address is typically 1.1.1.1 for all controllers.

Figure 4-13 WLC Interface Example

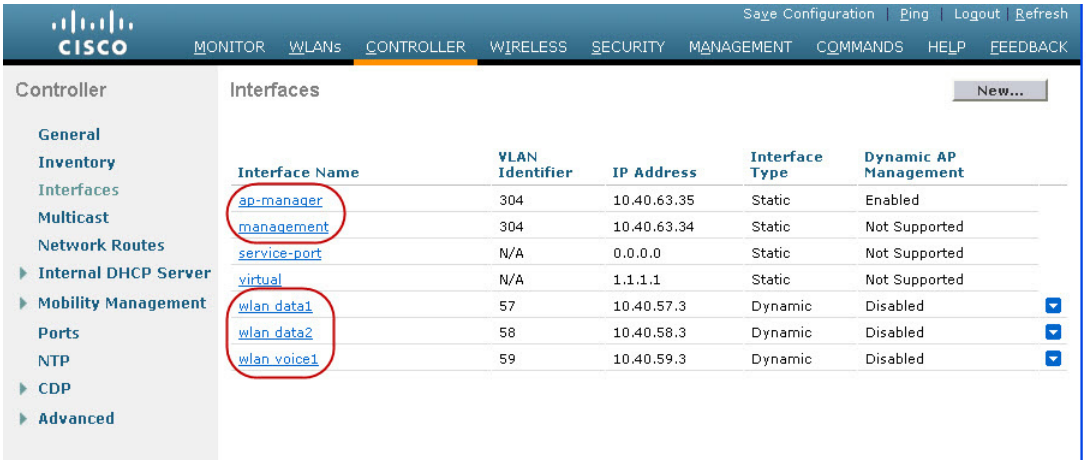

Figure 4-14 shows the mapping of a particular WLAN SSID to a defined interface. A WLAN can be mapped to the management interface (this is normally not recommended), or any dynamic interface.

Figure 4-14 WLAN Example

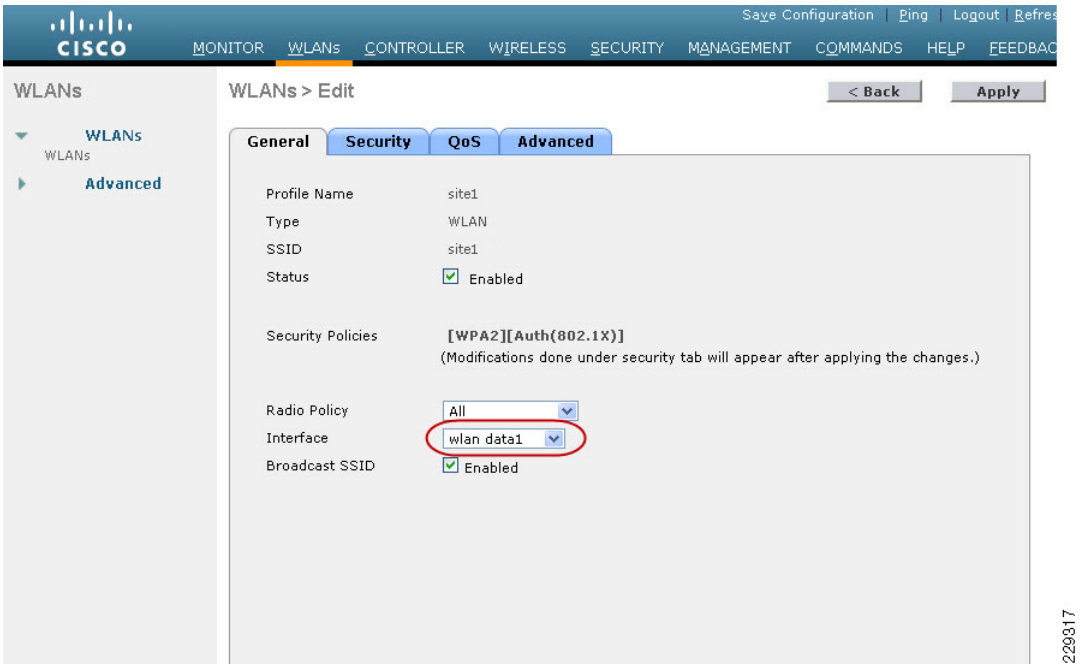

RF Groups and Mobility Groups

Part of a WLCs role is to manage the RF network in its area, and to provide mobility services to WLCs in its network. To define the area of the RF network that you are interested in managing, use an RF group name. To define the mobility services domain, use a mobility group. The details of RF groups and mobility groups are beyond the scope of this design guide, but the key point for the design is that the RF network area and the mobility services domain will typically be a single remote site, and only WLCs that are at the same site should have the same RF group name or mobility group name. Figure 4-15 shows an example of the RF and mobility group configuration on the controllers. Each remote site can be given a different RF group and mobility groups, since the WLCs are in different remote sites and are not expected to be in the same RF group or mobility group.

Figure 4-15 Mobility Groups and RF Groups Example

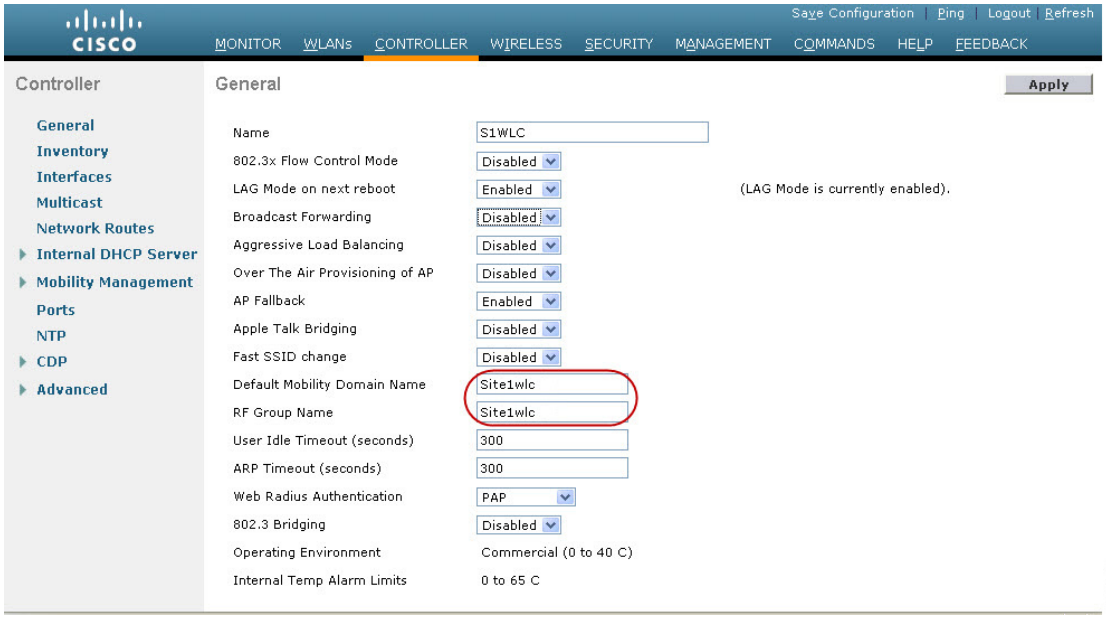

A remote site with only one WLC will have a mobility group with only its own details in the mobility group. If there is more than one WLC at the remote site, then the mobility group configuration will contain both WLCs.

Figure 4-16 shows a single WLAN example and Figure 4-17 shows a multiple WLC example. If there is only one WLC, the mobility group information is automatically populated. Additional WLCs must have the MAC address and management IP address added manually.

Figure 4-16 Mobility Groups for a Single WLC

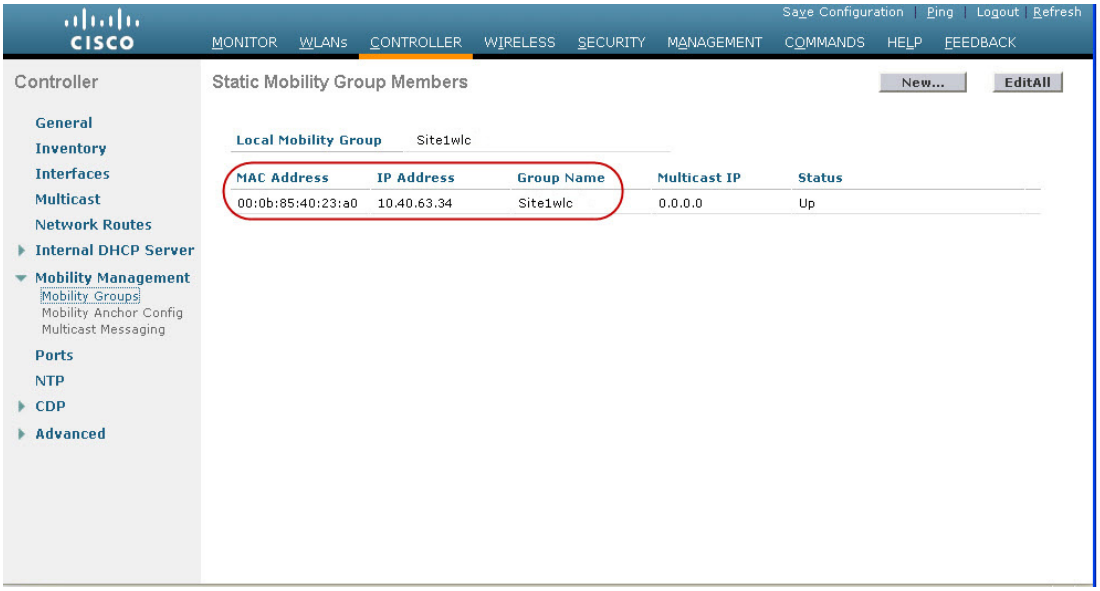

Figure 4-17 Mobility Groups for a Multiple WLCs

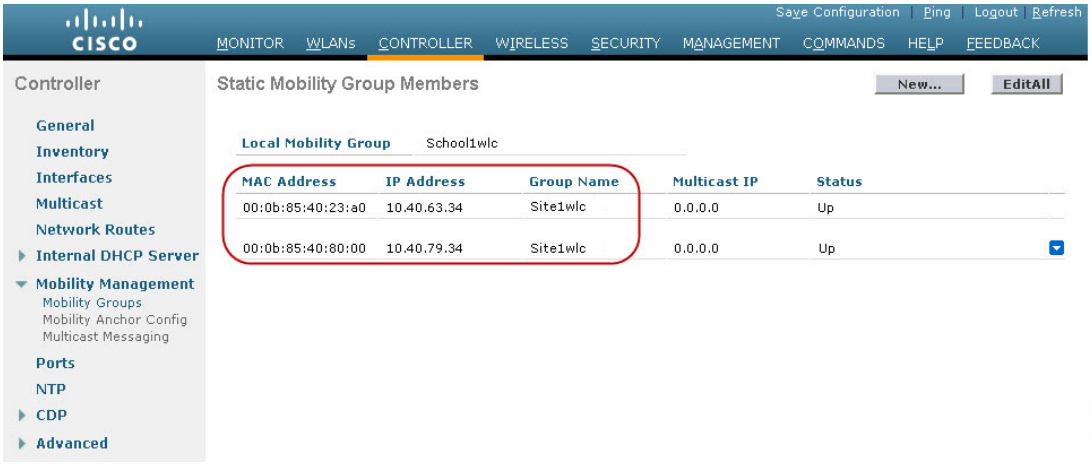

Example WLAN Configurations

In a typical remote site WLAN environment, it is expected that there be multiple WLANs (SSIDs) serving different purposes and different client groups. This section addresses the examples of what would be considered typical WLAN examples.

- A secured data WLAN network that uses 802.1X/EAP to provide AAA functionality and dynamically generated per-user, per-session encryption key.
- A secured VoWLAN network that also uses 802.1X/EAP to provide AAA functionality and optimized for voice.

- An open unencrypted WLAN for access to a WLAN network for unmanaged clients such as laptops, iPods, and iPhones.

For ease of administration and support for users who visit multiple sites, the WLAN SSIDs should be the same for each site in the enterprise. In addition, the SSIDs should be broadcast and have meaningful names.

Secured Employee WLAN

Figure 4-18 shows the general WLAN configuration tab for the secured data WLAN network. The key point shown are the security policy that has been set under the security tab and the WLC interface that the WLAN has been mapped to. The security configuration recommended is to use WPA2 with 802.1X+CCKM. Most WLANs should now support WPA2, and CCKM has been added to 802.1X as it provides a faster roaming for WLAN clients. This is for clients that support CCKM, while using the AAA features of 802.1X/AP to secure the WLAN connection.

Figure 4-18 General Configuration for Secured WLAN

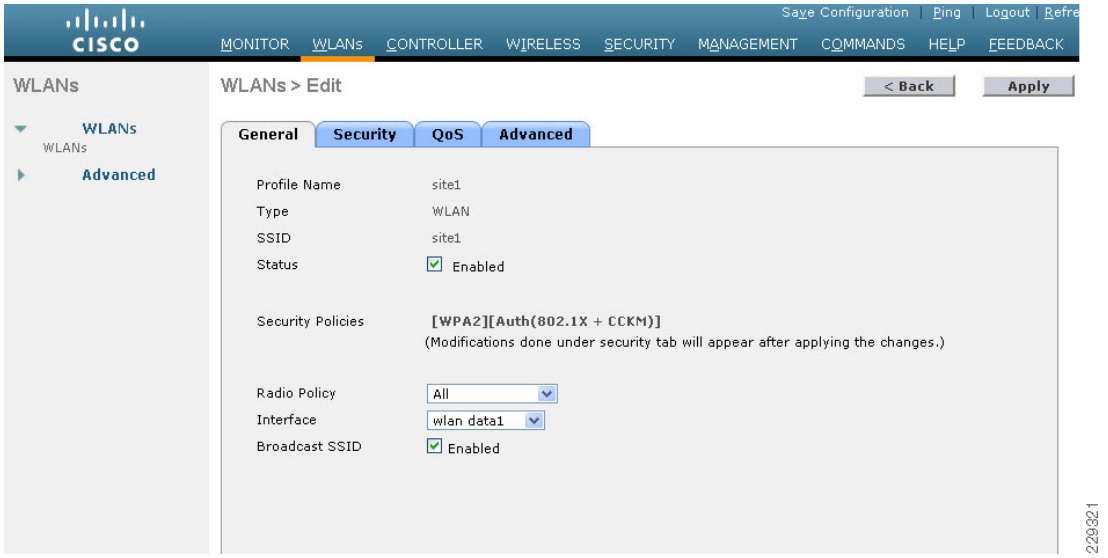

Figure 4-18 shows the QoS configuration for the secured data WLAN; in this case, the QoS profile is set to *Silver*, which is best effort setting. The WMM policy is set to disabled, as disabled WMM is the equivalent of best effort. The primary role of WMM is to give higher priority to voice and video traffic over the WLAN. Unless the site is planning to deliver interactive voice and video applications to their WLAN data clients, WMM can remain disabled.

Note 802.11n standard requires WMM be enabled and, therefore, WMM must be enabled on all WLANs in the 802.11n deployments. In this case, the WMM policy would be set to allowed.

Figure 4-19 Secured Employee WLAN QoS

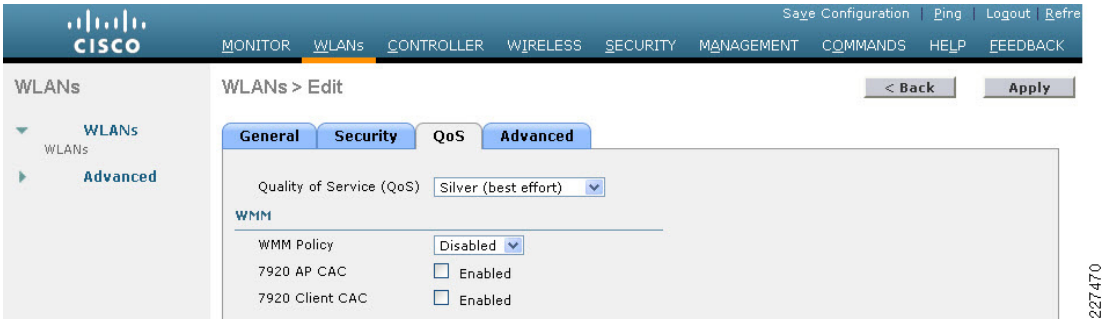

Figure 4-20 shows the secured data WLAN advanced configuration. The only change from the default settings on the tab is enabling the DHCP address assignment required feature. Typically, WLAN mobile clients use DHCP, and any statically configured client runs the risk of introducing an address duplication issue.

Figure 4-20 Secured Employee Advanced Configuration

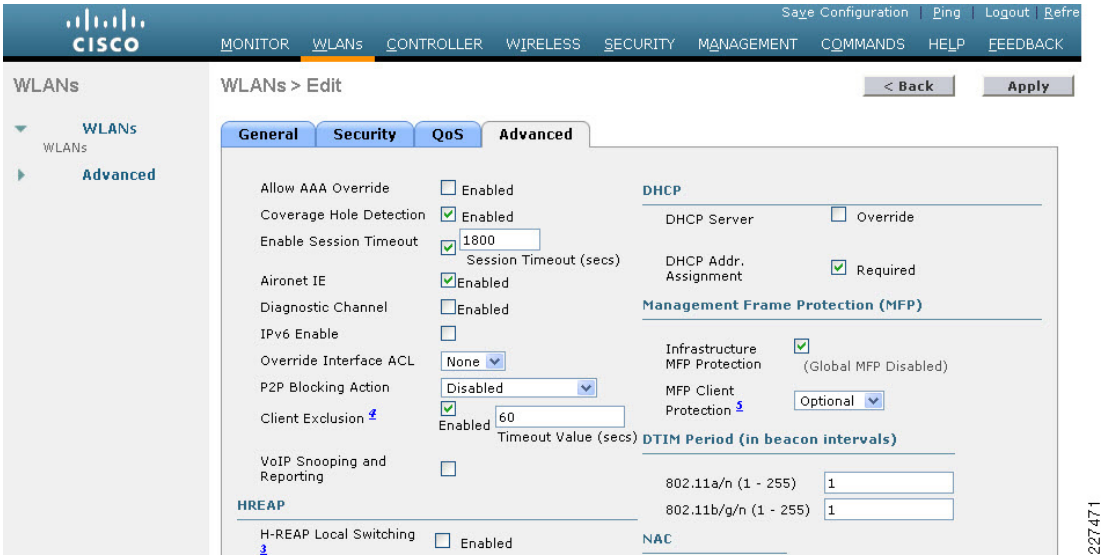

Secured VoWLAN

Figure 4-21 shows the General Tab of the voice over WLAN (VoWLAN). The primary difference between this WLAN and the secured data WLAN is that the security policy is WPA with CCKM, because this is the optimum security configuration for the Cisco 7921G and 7925G. The other difference is that the radio policy has been set for 802.11a only.

The use of 802.11a for the VoWLAN will depend on a number of factors, but the Cisco 7921G and 7925G are dual-band phones, and can use both bands but do not roam between bands. This means that once the handset associates with a network in one band, it will not leave that band while call quality is maintained. Keeping the VoWLAN handsets in the 802.11a band will ensure that the 2.4GHz band remains available for other client devices. Whether this is a viable option for a site depends on the required call capacity of the site's WLAN and the type of AP network that has been deployed.

Figure 4-21 VoWLAN General Configuration

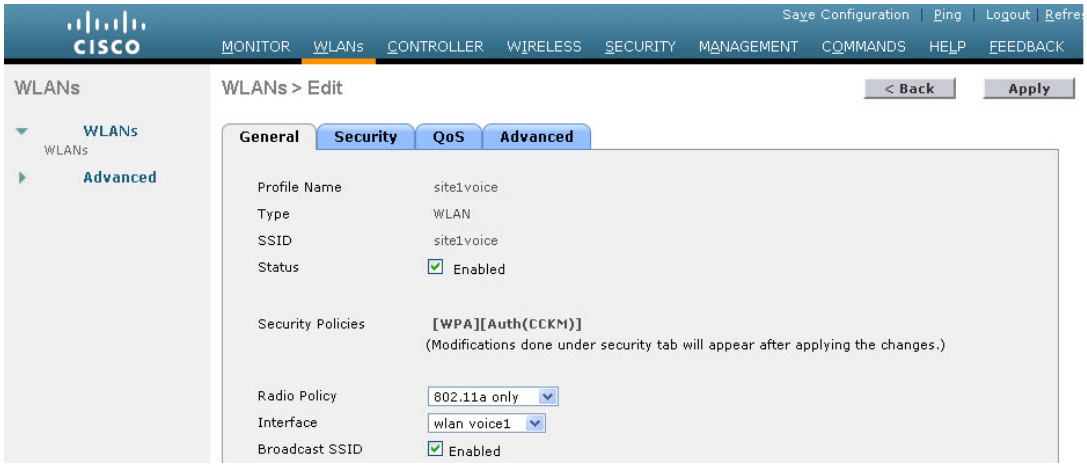

Figure 4-22 shows the QoS Tab for the VoWLAN. In this WLAN configuration, WMM is required. Both the 7921G and 7921G support WMM, and WMM will give voice traffic priority over other WLAN traffic on the network. The QoS profile is set to *Platinum* to ensure that the QoS classification is appropriate for voice. The QoS profile controls the maximum classification value for both the WLAN frames and LWAPP packets.

Figure 4-22 VoWLAN QoS Configuration

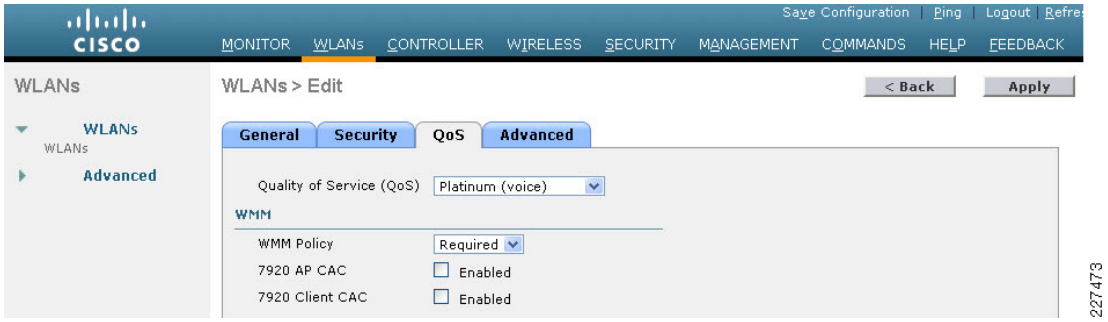

The Advanced Tab for the VoWLAN is the secured data WLAN. There is an option for VoIP snooping and reporting, but this option pertains only to a particular type of SIP and is not applicable to the Cisco 7921G and 7925G handsets.

To protect VoIP call quality, the WLC can perform call admission control (CAC) to prevent VoWLAN calls being added to an access point that cannot take any additional VoWLAN calls without compromising call quality. An example of the CAC configuration page is shown in Figure 4-23.

Note There is a separate CAC page for each RF band.

Figure 4-23 VoWLAN Call Admission Control

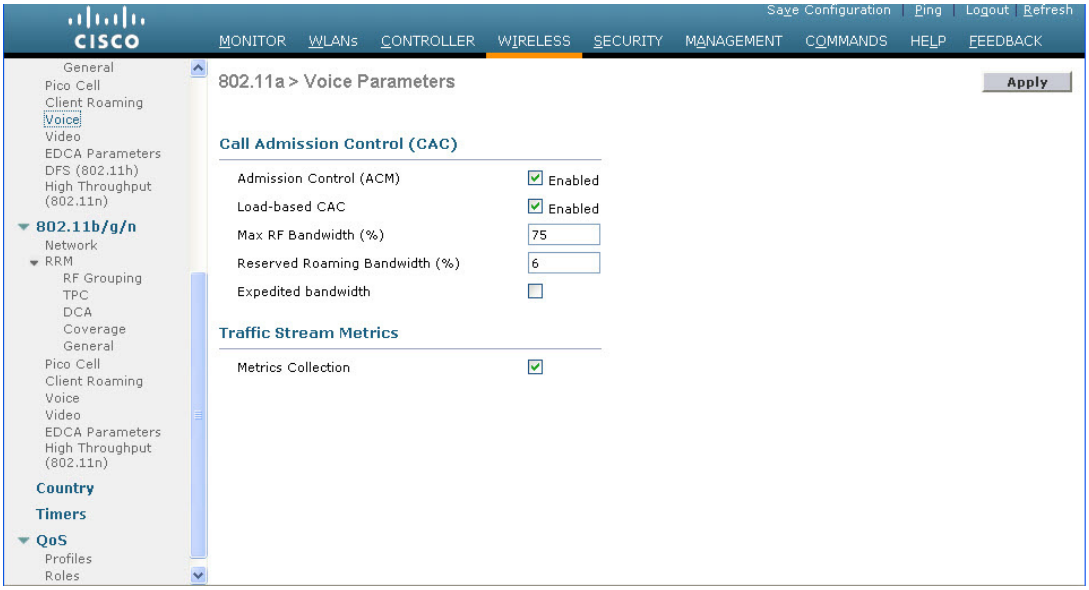

The CUWN prioritizes traffic based upon the QoS profiles applied to each WLAN, but it does not change the IP QoS classification (DSCP) of the client traffic carried by the CUWN. This means that client traffic that leaves the CUWN may need to be reclassified based upon the network policy. There are two ways of achieving this.

1. Applying policy at each of the network SVIs that connect the WLC to the network.
2. Learning the QoS policy that was applied within the CUWN as this should be in alignment with the network policy.

The second method is preferable as it requires less configuration and maintenance of the policy; the policy only needs to be maintained upon WLCs, and not open the WLCs and the connected switch. To achieve this, the Wired Protocol in the QoS profiles (Platinum, Gold, Sliver, and Bronze) must be set to 802.1p and all other settings may remain as default. This configures the WLC to set the 802.1p marking of the frames sent from the WLC to reflect QoS policy on that WLAN. For example, the IP packet was from a Platinum WLAN and had a DSCP value of EF, the WLC would use a CoS value of 5 in the frame header. If the same packet had been on a Silver WLAN, the CoS value would be 0. Therefore, if the WLC is connected to switch network that is configured to trust CoS and maintains a translation table between CoS and DSCP for its network, the translation between CUWN policy and network policy will occur automatically. See Figure 4-24.

For a further information on WLAN QoS, refer to the *Voice over WLAN Design Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html

Figure 4-24 Controller QoS Profiles

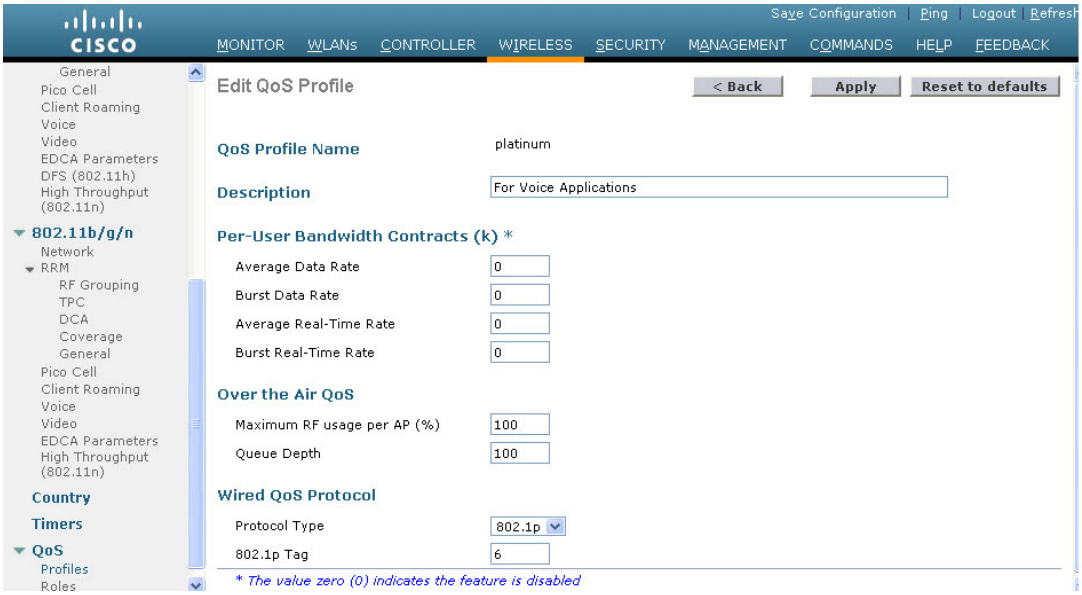

AP Deployments Considerations

As with any other WLAN deployment, the key design decisions are as follows: which areas require coverage and what level of performance is required in those areas with WLAN coverage. The general guidance for enterprise AP deployments is 15 to 20 active clients per AP. The number of APs required depends on many factors, including the number of clients, the type of applications, and the expected performance.

AP 1250

The Cisco 1250 Series is a rugged indoor access point designed for challenging RF environments that require the versatility associated with external antennas, a rugged metal enclosure, and a broad operating temperature range. The combined data rates of up to 600 Mbps to provide users with mobile access to high-bandwidth data, voice, and video applications. 802.11n provides reliable and predictable WLAN coverage to improve the end-user experience for both existing 802.11a/b/g clients and new 802.11n clients.

AP 1140

The Cisco 1140 Series Access Point is a business-ready, 802.11n access point designed for simple deployment and energy efficiency. The high-performance platform, which offers at least six times the throughput of existing 802.11a/g networks, prepares the business for the next wave of mobile devices and applications. Designed for sustainability, the Cisco 1140 Series delivers high performance from standard 802.3af PoE while decreasing waste with multi-unit eco-packs and Energy Star certified power supplies. As part of the CUWN, the Cisco 1140 Series provides the industry's lowest total cost of ownership and investment protection by integrating seamlessly with the existing network.

Coverage and Site Surveys

The WLAN coverage requirements can be expected to vary from enterprise to enterprise depending upon their goals and their budget. If the enterprise is simply to try to provide wireless network connectivity in selected areas, then simple tactical placement of APs in the selected rooms is likely to be sufficient. If the enterprise is planning to leverage the productivity associated with mobile applications and mobile access, then a more strategic approach is required.

If the enterprise is planning to implement a mobility solution, they need to examine the expected workflow and movement of the users of these applications to determine the range of coverage required and perform a site survey based on these coverage requirements. If the customer is considering WLAN location-based services as a possibility for future deployments, this should also be taken into account during the site survey process as the density and placement of APs can be substantially different when providing a suitable WLAN platform for location-based services.

Single-Band versus Dual-Band APs

There are both single-band and dual-band APs available for remote site solutions. The single-band APs support the 2.4GHz band and the dual-band APs support both the 2.4GHz and 5GHz band. It is a general recommendation that a dual-band solution be deployed.

Number of APs Per Room, Coverage in the Remote site

Single band APs vs Dual Band APs

If your goal is to simply provide WLAN coverage without trying to optimize capacity and performance, then a single-band AP is an appropriate choice; however, in most cases, a dual-band AP is a better long term choice.

The longevity of a WLAN deployment is fundamentally determined by its capacity. A quick look at the dual-band deployment shows that it has twice the capacity of a single-band solution, but a deeper look will reveal that the advantage of a dual-band solution is much greater than an additional radio.

The additional 5GHz radio, of a dual-band AP, is able to support a much higher capacity WLAN network, as it has access to approximately 7 times the number of non-overlapping channels as does a the 2.4GHz AP radio. In almost all 2.4GHz deployments, APs reusing the three non-overlapping channels interfere with each other and prevent the WLAN deployment from delivering a full WLAN capacity increase when the number of APs is increased, realizing its full theoretical capacity. A 5GHz AP is 7 times more likely to be able to delivery additional capacity for the addition of an AP.

Another consideration in the single-band versus dual-band AP discussion is 802.11n performance. 802.11n uses two primary mechanism to provide data rate improvements over the existing 802.11g and 802.11a standards. The first mechanism changes in the modulation, and error correction that can provide a data rate of up to 150Mbps, and the second mechanism is channel binding that combines non-overlapping channels to deliver data rates that are multiples of what a single channel could achieve. Channel binding is only available for the 5GHz band, as there is not sufficient channel capacity to support it in an enterprise 2.4GHz deployment.

Deploying a dual-band WLAN system is not a matter of simply replacing the APs in place, the 5GHz band has different power constraints, and has different propagation properties that need to be considered when deciding on AP density and placement. If fiscally possible, a dual-band AP solution should be planned and deployed initially. This will save an expensive rework layer.

For further discussion on 2.4GHz vs 5GHz capacity, refer to the *Voice over WLAN Design Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html

Client Considerations

One additional consideration in the single-band versus dual-band AP decision is the client devices that the WLAN network is going to support. Many earlier laptops and mobile devices only supported the 2.4GHz band, and this is still true for many consumer WLAN clients. To take advantage of a dual-band solution, a concerted effort needs to be made to ensure that as many clients as possible are also dual-band. For cases where the remote site is purchasing WLAN clients, they should favor dual-band devices. When recommending WLAN client devices, they should point out that the dual-band client devices will have access to a higher performance network. Of course, the first step is to have the dual band network in place, in order to for client devices to take advantage of their investment in a higher performance clients.

WLC Discovery

CUWN provides auto-discovery functionality for its APs, where an AP upon connection to an appropriately connected network can automatically find and connect to a WLC. The WLC will ensure that the AP is running the appropriate software version, apply the appropriate configuration to that AP, and adjust the radio settings to optimize the AP for its current environment.

Multiple auto-discovery options are available in the CUWN:

- Over the air: The APs learns the IP address of WLCs from APs in the area which are currently attached to those WLCs
- DHCP: The APs learns the IP address(es) of the WLCs as part of its DHCP address assignment
- DNS: The APs learn the IP address(es) of the WLCs by querying and well known DNS name CISCO-LWAPP-CONTROLLER.<*localdomain.com*>
- Staging: Have the AP join a WLC prior to them being deployed, and the APs will attempt to rejoin this WLC when reconnected to the network
- Static Configuration: The APs can be manually configured with the WLC IP address prior to being connected to the networks

Given that the small enterprise design profile uses a local DNS server for sites to ensure survivability, the use of the DNS discovery provides the simplest WLC discovery mechanism.

For details about how to configure DHCP discovery, refer to the DHCP OPTION 43 for Lightweight Cisco Aironet Access Points Configuration Example at the following URL:

http://www.cisco.com/en/US/tech/tk722/tk809/technologies_configuration_example09186a00808714fe.shtml

WLC Failover Options

CUWN provides multiple failover options, allowing APs to make a choice between WLCs based on configured priorities. When an AP goes through its discovery process, it learns about all of the WLCs in the mobility group, and can prioritize based on its high availability (HA) configuration or choose an WLC based on current loads.

In network architectures, such as the small enterprise design profile, where there is a high-speed WAN/MAN that makes AP failover to a remote WLC—such as the main site WLC—feasible, APs can be configured to failover to a WLC outside their mobility group. In this scenario, the remote WLC would not be in the Mobility Group that is learned during the AP discovery process, and the IP address of the remote WLC needs to be provided in the HA configuration.

This feature allows the main site to become a backup WLC for the remote sites in an event of an WLC outage at the remote site. For this to be effective, a common WLAN SSID naming policy for key WLANs needs to be implemented within the enterprise to ensure that WLAN clients do not have to be reconfigured in the event of an AP failover to the main site WLC. This type of HA configuration is call N+1 where a single main site WLC is able to provide HA at a much lower cost than a traditional 1+1 design which would require additional WLCs at each remote site. See Figure 4-25.

Figure 4-25 AP High Availability Configuration Example

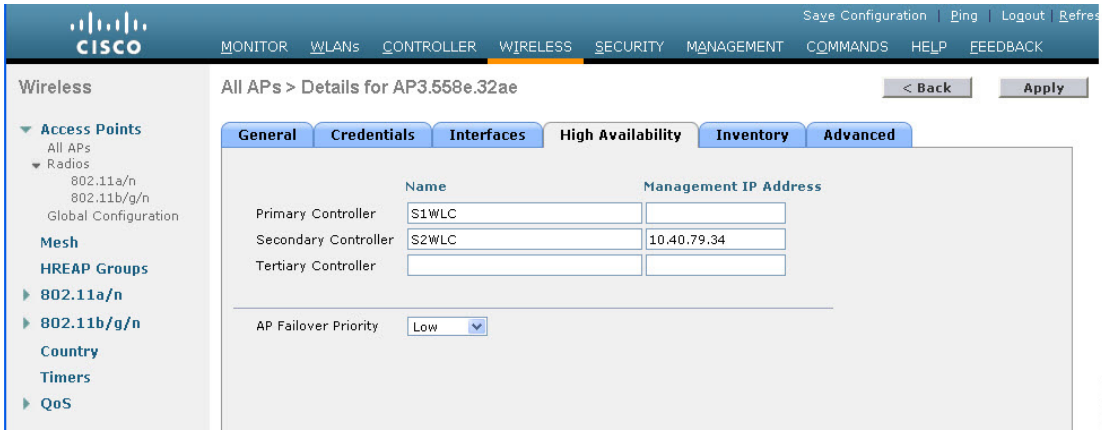

Appendix A—Devices and Software Used

Table 4-1 lists the devices and software used for the CUWN in this design guide.

Table 4-1 WLAN Devices and Software

Name	Version
WCS	6.0.132
WLC 4402	6.0.182.0
WLC 4404	6.0.182.0
AP1252	AIR-LAP1252AG-A-K9
AP1142	AIR-LAP1142N-A-K9

CHAPTER 5

Small Enterprise Design Profile(SEDP)—Security Design

This chapter describes how the Small Enterprise Design Profile sets the foundation for safe and secure enterprise networks by leveraging the proven design and deployment guidelines of the Cisco SAFE security architecture. The Small Enterprise Design Profile is a well-designed and validated network architecture that enables the small enterprise to deliver all of the services required for an enhanced business environment. Cisco SAFE is a security reference architecture that provides detailed design and implementation guidelines for organizations looking to build highly secure and reliable networks.

Small Enterprise Network Security Design

The Small Enterprise Design Profile was designed with built-in security to protect the infrastructure and to provide a safe environment for business. Following the proven guidelines of the Cisco SAFE security architecture, a series of network security technologies and products are strategically deployed throughout the network to protect employees and company assets, to guarantee the confidentiality of sensitive data, and to ensure the availability and integrity of systems and data.

Small enterprises are not immune to violence, theft, vandalism, and other threats. The adoption of new network collaboration and Internet-based technologies also opens the possibility for a number of cyber threats. Understanding the nature and diversity of these threats, and how they may evolve over time, is the first step towards a successful enterprise security strategy. The following are some of the common threats to enterprise environments:

- *Service disruption*—Disruption to the infrastructure, applications and other business resources caused by botnets, worms, malware, adware, spyware, viruses, denial-of-service (DoS) attacks, and Layer 2 attacks.
- *Network abuse*—Use of non-approved applications by employees; peer-to-peer file sharing and instant messaging abuse; and access to non-business related content.
- *Unauthorized access*—Intrusions, unauthorized users, escalation of privileges, and unauthorized access to restricted resources.
- *Data loss*—Theft or leakage of private and confidential data from servers, endpoints, while in transit, or as a result of spyware, malware, key-loggers, viruses, etc.
- *Identity theft and fraud*—Theft of personnel identity or fraud on servers and end users through phishing and E-mail spam.

The Small Enterprise Design Profile is based on a validated network architecture designed around both business operations and technical considerations. Recognizing the fact cost is a common limiting factor to small enterprise network designs, the architecture topologies and platforms were carefully selected to increase productivity while reducing overall costs.

The Small Enterprise Design Profile accommodates a main site and multiple remote sites of various sizes, all interconnected over a Metro Ethernet core. The architecture design is illustrated in Figure 5-1. At the heart of the architecture is a robust routing and switching network. Operating on top of this network are the all the services used within the enterprise. This often includes the majority of the business most critical applications such as databases, payroll, accounting and customer relationship management (CRM) applications. The core of those services are deployed and managed at the main site, allowing the enterprise to reduce the need for separate services to be operated and maintained at the various remote locations. Centralized systems and applications are served by a main site serverfarm.

Figure 5-1 Small Enterprise Network Design

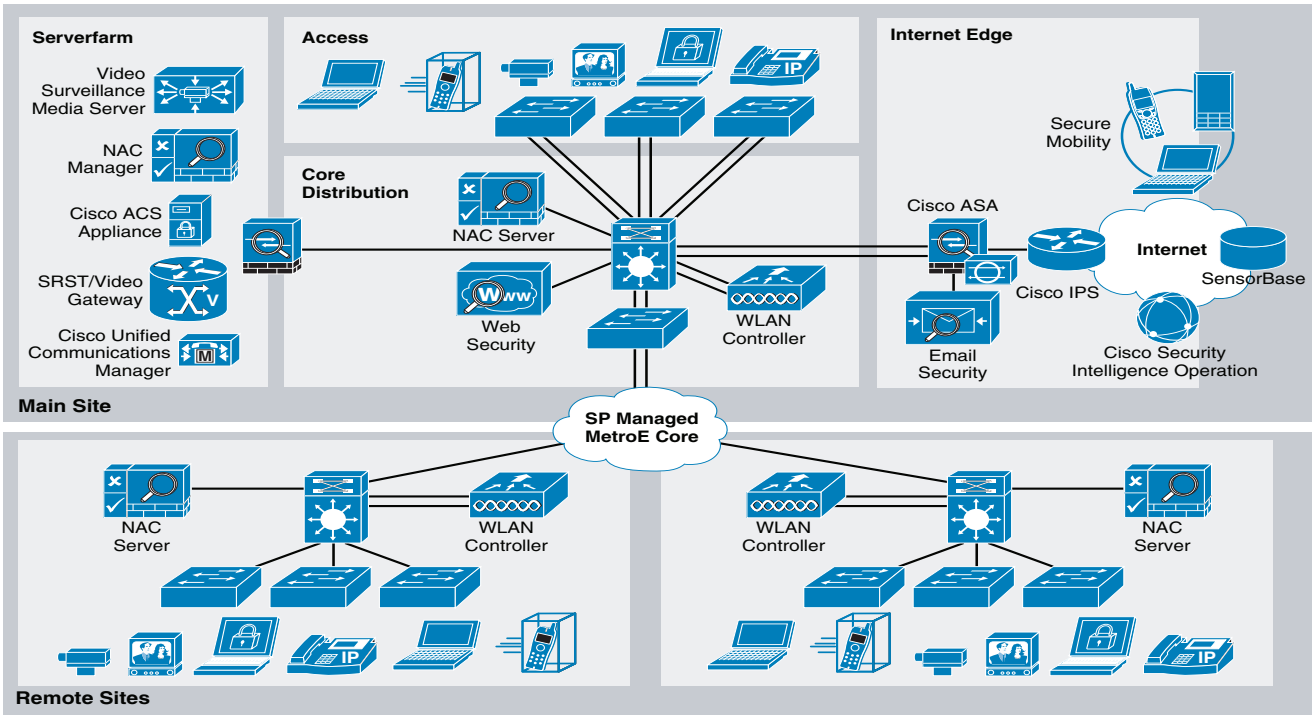

As shown in Figure 5-1, the small enterprise network design follows a defense-in-depth approach, where multiple layers of protection are built into the architecture. The different security tools are combined together for enhanced visibility and control.

The security design of the architecture focuses on the following key areas:

- *Network Foundation Protection (NFP)*
  - Ensuring the availability and integrity of the network infrastructure, protecting the control and management planes.

- *Internet Perimeter Protection*
  - Ensuring safe Internet connectivity, and protecting internal resources and users from malware, viruses, and other malicious software.
  - Protecting personnel from harmful and inappropriate content.
  - Enforcing E-mail and web browsing policies.
- *Serverfarm Protection*
  - Ensuring the availability and integrity of centralized applications and systems.
  - Protecting the confidentiality and privacy of sensitive data.
- *Network Access Security and Control*
  - Securing the access edges. Enforcing authentication and role-based access for users residing at the main site and remote offices.
  - Ensuring systems are up-to-date and in compliance with the enterprise network security policies.
- *Secure Mobility*
  - Providing secure, persistent connectivity to all mobile employees on laptops, smartphones and other mobile platforms. Enforcing encryption, authentication and role-based access to all mobile users.
  - Delivering consistent protection to all mobile employees from viruses, malware, botnets, and other malicious software.
  - Ensuring a persistent enforcement of enterprise network security policies to all users. Making sure systems comply with corporate policies and have up-to-date security.

The following sections discuss the key areas of the small enterprise network security design.

Network Foundation Protection

Small enterprise networks are built with routers, switches, and other network devices that keep the applications and services running. Therefore, properly securing these network devices is critical for continued operation.

The Small Enterprise Design Profile protects the network infrastructure by implementing the Cisco SAFE best practices for the following areas:

- *Infrastructure device access*
  - Restrict management device access to authorized parties and for the authorized ports and protocols.
  - Enforce Authentication, Authorization and Accounting (AAA) with Terminal Access Controller Access Control System (TACACS+) or Remote Authentication Dial-In User Service (RADIUS) to authenticate access, authorize actions and log all administrative access.
  - Display legal notification banners.
  - Ensure confidentiality by using secure protocols like Secure Shell (SSH) and HTTPS.
  - Enforce idle and session timeouts. Disable unused access lines.
- *Routing infrastructure*
  - Restrict routing protocol membership by enabling Message-Digest 5 (MD5) neighbor authentication and disabling default interface membership.

  - Enforce route filters to ensure that only legitimate networks are advertised; and networks that are not supposed to be propagated are never advertised.
  - Log status changes of neighbor sessions to identify connectivity problems and DoS attempts on routers.
- *Device resiliency and survivability*
  - Disable unnecessary services, implement control plane policing (CoPP).
  - Enable traffic storm control.
  - Implement topological, system and module redundancy for the resiliency and survivability of routers and switches and to ensure network availability.
  - Keep local device statistics.
- *Network telemetry*
  - Enable Network Time Protocol (NTP) time synchronization.
  - Collect system status and event information with Simple Network Management Protocol (SNMP), Syslog, TACACS+/RADIUS accounting.
  - Monitor CPU and memory usage on critical systems.
  - Enable NetFlow to monitor traffic patterns and flows.
- *Network policy enforcement*
  - Implement access edge filtering.
  - Enforce IP spoofing protection with access control lists (ACLs), Unicast Reverse Path Forwarding (uRPF), and IP Source Guard.
- *Switching infrastructure*
  - Implement a hierarchical design, segmenting the LAN into multiple IP subnets or virtual LANs (VLANs) to reduce the size of broadcast domains.
  - Protect the Spanning Tree Protocol (STP) domain with BPDU Guard and STP Root Guard.
  - Use per-VLAN Spanning Tree to reduce the scope of possible damage.
  - Disable VLAN dynamic trunk negotiation on user ports.
  - Disable unused ports and put them into an unused VLAN.
  - Implement Catalyst Infrastructure Security Features (CISF) including port security, dynamic ARP inspection, DHCP snooping.
  - Use a dedicated VLAN ID for all trunk ports.
  - Explicitly configure trunking on infrastructure ports.
  - Use all tagged mode for the native VLAN on trunks and drop untagged frames.
- *Network management*
  - Ensure the secure management of all devices and hosts within the enterprise network.
  - Authenticate, authorize, and keep record of all administrative access.
  - If possible, implement a separate out-of-band (OOB) management network (hardware or VLAN based) to manage systems at the main site.
  - Secure the OOB by enforcing access controls, using dedicated management interfaces or virtual routing and forwarding (VRF) tables.

- Provide secure in-band management access for systems residing at the remote sites by deploying firewalls and ACLs to enforce access controls, using Network Address Translation (NAT) to hide management addresses, and using secure protocols like SSH and HTTPS.
- Ensure time synchronization by using NTP.
- Secure servers and other endpoint with endpoint protection software and operating system (OS) hardening best practices.

For more detailed information on the Network Foundation Protection best practices, refer the "Chapter 2, Network Foundation Protection" of the *Cisco SAFE Reference Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap2.html

Internet Perimeter Protection

The Small Enterprise Design Profile assumes the existence of a centralized Internet connection at the headquarters or main site serving users at all locations. Common services include E-mail and web browsing for employees, the hosting of a company's website accessible to clients and partners over the Internet, and secure remote access for mobile users and remote workers. Other services may also be provided using the same infrastructure.

The network infrastructure that provides Internet connectivity is defined as the Internet perimeter, illustrated in Figure 5-2.

Figure 5-2 Internet Perimeter

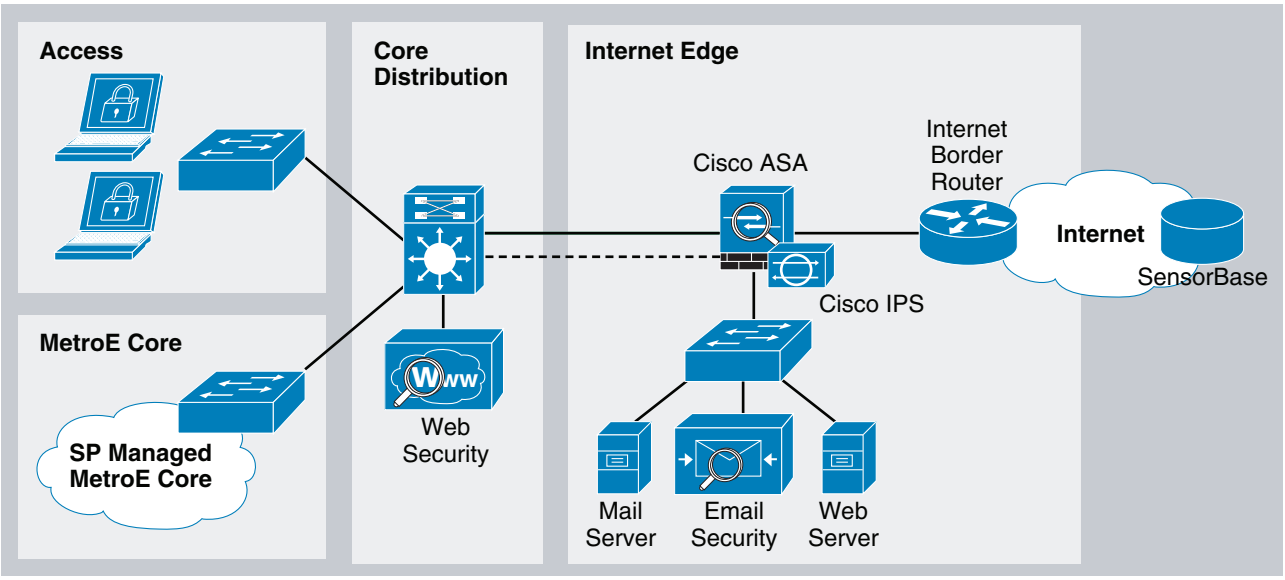

The primary functions of the Internet perimeter is to allow for safe and secure access for users residing at all locations, and to provide public services without compromising the confidentiality, integrity and availability of the company's resources and data. To that end, the Internet perimeter incorporates the following security functions:

- *Internet Border Router*—The Internet border router is the Internet gateway responsible for routing traffic between the enterprise network and the Internet. The Internet border router may be administered by the company's personnel or may be managed by the Internet service provider (ISP). The router provides the first line of protection against external threats and should be hardened following the Network Foundation Protection (NFP) best practices.
- *Internet Firewall*—A Cisco ASA provides stateful access control and deep packet inspection to protect company resources and data from unauthorized access and disclosure. The security appliance is configured to prevent incoming access from the Internet, to protect the enterprise Internet public services, and to control user traffic bound to the Internet. In addition, the Cisco ASA Botnet Traffic Filter feature can be enabled to defend the enterprise against botnet threats. Once enabled, the Botnet Traffic Filter feature monitors network traffic across all ports and protocols for rogue activity and to prevent infected internal endpoints from sending command and control traffic back to external hosts on the Internet. The security appliance may also provide secure remote access to employees with the Cisco AnyConnect Secure Mobility client.
- *Intrusion Prevention*—An Advanced Inspection and Prevention Security Service Module (AIP SSM) on the Cisco ASA or a separate IPS appliance can be implemented for enhanced threat detection and mitigation. The IPS module or appliance is responsible for identifying and blocking anomalous traffic and malicious packets recognized as well-known attacks. IPS can be deployed either in inline or promiscuous mode. The module or appliance may be configured to participate in Cisco IPS Global Correlation, allowing the IPS to gain visibility on global threats as they emerge in the Internet, and to quickly react to contain them. IPS may also be configured to help block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and more.
- *Public Services DMZ*— The company's Internet website, mail server and other public-facing services may be placed on a demilitarized zone (DMZ) for security and control purposes. The DMZ acts as a middle stage between the Internet and company's private resources, preventing external users from directly accessing any internal servers and data. The Internet firewall is responsible for restricting incoming access to the public services and by limiting outbound access from DMZ resources out to the Internet. Systems residing on the DMZ are hardened with endpoint protection software and operating system (OS) hardening best practices.
- *E-mail Security*—A Cisco Ironport C Series E-mail Security Appliance (ESA) is deployed at the DMZ to inspect incoming and outgoing E-mails and eliminate threats such as E-mail spam, viruses and worms. The ESA appliance also offers E-mail encryption to ensure the confidentiality of messages, and data loss prevention (DLP) to detect the inappropriate transport of sensitive information.
- *Web Security*—A Cisco IronPort S Series Web Security Appliance (WSA) is deployed at the distribution switches to inspect HTTP and HTTPS traffic bound to the Internet. This system enforces URL filtering policies to block access to websites containing non-business related content or that are known to be the source of spyware, botnets or other type of malware. The WSA may also be configured to block certain Internet applications such as AOL Messenger, BitTorrent, Skype, etc.

Following are the design guidelines for implementing the security functions.

Internet Border Router Security Guidelines

The Internet border router provides connectivity to the Internet through one or more Internet service providers. The router act as the first line-of-defense against unauthorized access, DDoS, and other external threats. Access control lists (ACLs), uRPF, and other filtering mechanisms may be implemented for anti-spoofing and to block invalid packets. NetFlow, Syslog, SNMP may be used to gain visibility on traffic flows, network activity and system status. In addition, the Internet border router should be secured

following the practices explained in the "Network Foundation Protection" section on page 5-3. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information.

The "Internet Border Router Deployment" section on page 5-41provides an example of Internet edge ACL. For more information on how to configure the Internet border router, refer to the "Chapter 6, Enterprise Internet Edge" of the *Cisco SAFE Reference Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap6.html

Internet Firewall Guidelines

The Cisco ASA deployed at the Internet perimeter is responsible for protecting the enterprise internal resources and data from external threats by preventing incoming access from the Internet; protecting public resources served by the DMZ by restricting incoming access to the public services and by limiting outbound access from DMZ resources out to the Internet; and controlling user's Internet-bound traffic.

To that end, the security appliance is configured to enforce access policies, keep track of connection status, and inspect packet payloads following these guidelines:

- Deny any connection attempts originating from the Internet to internal resources and subnets.
- Allow outbound Internet access for users residing at any of the enterprise locations and for the protocols permitted by the organization's policies (i.e., HTTP and HTTPS).
- Allow outbound Internet SSL access for administrative updates, SensorBase, IPS signature updates, etc.
- Allow users access to DMZ services such as company's website, E-mail, and domain name resolution (HTTP, SMTP, POP, IMAP, and DNS).
- Restrict inbound Internet access to the DMZ for the necessary protocols and servers (HTTP to web server, SMTP to the mail transfer agent, DNS to DNS server, etc.).
- Restrict connections initiated from DMZ to the only necessary protocols and sources (DNS from DNS server, SMTP from the mail server, HTTP/SSL from Cisco IronPort ESA).
- Enable stateful inspection for the used protocols to ensure returning traffic is dynamically allowed by the firewall.
- Implement Network Address Translation (NAT) and Port Address Translation (PAT) to shield the internal address space from the Internet.

Figure 5-3 illustrates the protocols and ports explicitly allowed by the Cisco ASA.

Figure 5-3 Allowed Protocols and Ports

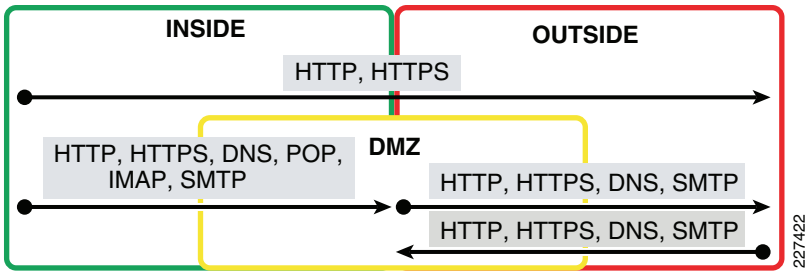

Note Figure 5-3 does not include any management traffic destined to the firewall. Whenever available, a dedicated management interface should be used. In case the firewall is managed in-band, identify the protocols and ports required prior to configuring the firewall ACLs.

It is also important to remember that the Cisco ASA should be hardened following the Network Foundation Protection best practices. This includes restricting and controlling administrative access, securing the dynamic exchange of routing information with MD5 authentication, and enabling firewall network telemetry with SNMP, Syslog, and NetFlow.

In the Small Enterprise Design Profile, high availability is achieved by using redundant physical interfaces. This represents the most cost-effective solution for high availability. As an alternative, a pair of firewall appliances could be deployed in stateful failover, as discussed in the "Internet Firewall Deployment" section on page 5-42.

Cisco ASA Botnet Traffic Filter

The Cisco ASA Botnet Traffic Filter feature can be enabled to monitor network ports for rogue activity and to prevent infected internal endpoints from sending command and control traffic back to an external host on the Internet. The Botnet Traffic Filter on the ASA provides reputation-based control for an IP address or domain name, similar to the control that Cisco IronPort SensorBase provides for E-mail and web servers.

The Cisco Botnet Traffic Filter is integrated into all Cisco ASA appliances, and inspects traffic traversing the appliance to detect rogue traffic in the network. When internal clients are infected with malware and attempt to phone home to an external host on the Internet, the Botnet Traffic Filter alerts the system administrator of this through the regular logging process and can be automatically blocked. This is an effective way to combat botnets and other malware that share the same phone-home communications pattern.

The Botnet Traffic Filter monitors all ports and performs a real-time lookup in its database of known botnet IP addresses and domain names. Based on this investigation, the Botnet Traffic Filter determines whether a connection attempt is benign and should be allowed, or is a risk and should be blocked.

The Cisco ASA Botnet Traffic Filter has three main components:

- *Dynamic and administrator blacklist data*—The Botnet Traffic Filter uses a database of malicious domain names and IP addresses that is provided by Cisco Security Intelligence Operations. This database is maintained by Cisco Security Intelligence Operations and is downloaded dynamically from an update server on the SensorBase network. Administrators can also configure their own local blacklists and whitelists.
- *Traffic classification and reporting*—Botnet Traffic Filter traffic classification is configured through the **dynamic-filter** command on the ASA. The dynamic filter compares the source and destination addresses of traffic against the IP addresses that have been discovered for the various lists available (dynamic black, local white, local black), and logs and reports the hits against these lists accordingly.
- *Domain Name System (DNS) snooping*—To map IP addresses to domain names that are contained in the dynamic database or local lists, the Botnet Traffic Filter uses DNS snooping in conjunction with DNS inspection. Dynamic Filter DNS snooping looks at User Datagram Protocol (UDP) DNS replies and builds a DNS reverse cache (DNSRC), which maps the IP addresses in those replies to the domain names they match. DNS snooping is configured via the Modular Policy Framework (MPF) policies

The Botnet Traffic Filter uses two databases for known addresses. Both databases can be used together, or the dynamic database can be disabled and the static database can be used alone. When using the dynamic database, the Botnet Traffic Filter receives periodic updates from the Cisco update server on the Cisco IronPort SensorBase network. This database lists thousands of known bad domain names and IP addresses.

Note The SensorBase network is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware, and other abnormal behavior. The network is composed of the Cisco IronPort appliances, Cisco ASA, and Cisco IPS appliances and modules installed in more than 100,000 organizations worldwide, providing a large and diverse sample of Internet traffic patterns.

The Cisco ASA uses the dynamic database as follows:

1. When the domain name in a DNS reply matches a name in the dynamic database, the Botnet Traffic Filter adds the name and IP address to the DNS reverse lookup cache.
2. When the infected host starts a connection to the IP address of the malware site, the Cisco ASA sends a syslog message reporting the suspicious activity and optionally drops the traffic if the Cisco ASA is configured to do so.
3. In some cases, the IP address itself is supplied in the dynamic database, and the Botnet Traffic Filter logs or drops any traffic to that IP address without having to inspect DNS requests.

The database files are stored in running memory rather than Flash memory. The database can be deleted by disabling and purging the database through the configuration.

Note To use the database, be sure to configure a domain name server for the Cisco ASA so that it can access the URL of the update server. To use the domain names in the dynamic database, DNS packet inspection with Botnet Traffic Filter snooping needs to be enabled; the Cisco ASA looks inside the DNS packets for the domain name and associated IP address.

In addition to the dynamic database, a static database can be used by manually entering domain names or IP addresses (host or subnet) that you want to tag as bad names in a blacklist. Static blacklist entries are always designated with a Very High threat level. Domain names or IP addresses can also be entered in a whitelist,

When a domain name is added to the static database, the Cisco ASA waits one minute, and then sends a DNS request for that domain name and adds the domain name/IP address pairing to the DNS host cache. This action is a background process, and does not affect your ability to continue configuring the Cisco ASA. Cisco also recommends that DNS packet inspection be enabled with Botnet Traffic Filter snooping. When enabled, the Cisco ASA uses Botnet Traffic Filter snooping instead of the regular DNS lookup to resolve static blacklist domain names in the following cases:

- The Cisco ASA DNS server is unavailable.
- A connection is initiated during the one minute waiting period before the Cisco ASA sends the regular DNS request.

If DNS snooping is used, when an infected host sends a DNS request for a name on the static database, the Cisco ASA looks inside the DNS packets for the domain name and associated IP address and adds the name and IP address to the DNS reverse lookup cache.

If Botnet Traffic Filter snooping is not enabled, and one of the above circumstances occurs, that traffic is not monitored by the Botnet Traffic Filter.

Note It is important to realize that a comprehensive security deployment should include Cisco Intrusion Prevention Systems (IPS) with its reputation-based Global Correlation service and IPS signatures in conjunction with the security services provided by the Cisco ASA security appliance such as Botnet Traffic Filter.

For more information on the Cisco ASA Botnet Traffic Filter feature, see the following URL: http://www.cisco.com/en/US/prod/vpndevc/ps6032/ps6094/ps6120/botnet_index.html

Intrusion Prevention Guidelines

IPS is responsible for identifying and blocking anomalous traffic and packets recognized as well-known attacks. An AIP SSM module on the Cisco ASA Internet firewall or a separate IPS appliance can be implemented in the Internet perimeter for enhanced threat detection and mitigation. IPS may also be configured to help block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and so on.

Integrating IPS on a Cisco ASA appliance using an AIP SSM provides a cost-effective solution for small enterprise networks. The AIP SSM is supported on Cisco ASA 5510 and higher platforms. The AIP SSM runs advanced IPS software providing proactive, full-featured intrusion prevention services to stop malicious traffic before it can affect the enterprise network.

The AIP SSM may be deployed in inline or promiscuous mode:

- *Inline mode*—The AIP SSM is placed directly in the traffic flow (see the left side of Figure 5-4). Traffic identified for IPS inspection cannot continue through the ASA without first passing through and being inspected by the AIP SSM. This mode is the most secure because every packet that has been identified for inspection is analyzed before being allowed through. Also, the AIP SSM can implement a blocking policy on a packet-by-packet basis. This mode, however, can affect throughput if not designed or sized appropriately.
- *Promiscuous mode*—A duplicate stream of traffic is sent to the AIP SSM. This mode is less secure, but has little impact on traffic throughput. Unlike inline mode, in promiscuous mode the AIP SSM can block traffic only by instructing the ASA to shun the traffic or by resetting a connection on the ASA. Also, while the AIP SSM is analyzing the traffic, a small amount of traffic might pass through the ASA before the AIP SSM can shun it. The right side of Figure 5-4 shows the AIP SSM in promiscuous mode.

Figure 5-4 IPS Inline and Promiscuous Modes

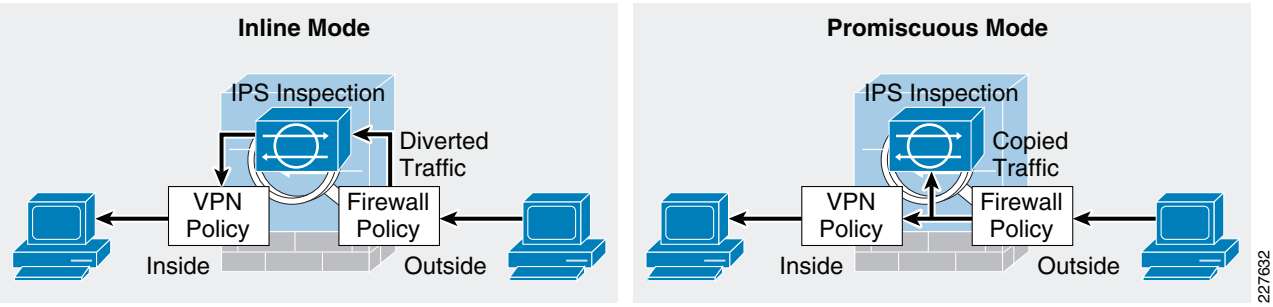

The recommended IPS deployment mode depends on the goals and policies of the enterprise. IPS inline mode is more secure because of its ability to stop malicious traffic in real-time; however, it may impact traffic throughput if not properly designed or sized. Conversely, IPS promiscuous mode has less impact on traffic throughput but is less secure because there may be a delay in reacting to the malicious traffic.

Although the AIP SSM runs as a separate application within the Cisco ASA, it is integrated into the traffic flow. The AIP SSM contains no external interfaces itself, except for the management interface on the SSM itself. When traffic is identified for IPS inspection on the ASA, traffic flows through the ASA and the AIP SSM in the following sequence:

Step 1 Traffic enters the adaptive security appliance.

Step 2 Firewall policies are applied.

Step 3 Traffic is sent to the AIP SSM over the backplane.

Step 4 The AIP SSM applies its security policy to the traffic and takes appropriate actions.

Step 5 Valid traffic (for inline mode only) is sent back to the adaptive security appliance over the backplane; the AIP SSM might block some traffic according to its security policy and that traffic is not passed on.

Step 6 VPN policies are applied (if configured).

Step 7 Traffic exits the adaptive security appliance.

The AIP SSM card may be configured to fail open or close when the module becomes unavailable. When configured to fail open, the ASA allows all traffic through, uninspected, if the AIP SSM becomes unavailable. Conversely, when configured to fail close, the ASA blocks all traffic in case of an AIP SSM failure.

Cisco IPS Global Correlation

If desired, the AIP SSM module on the Cisco ASA (or IPS appliance) may participate in Cisco IPS Global Correlation for further threat visibility and control. Once enabled, the participating IPS sensor receives threat updates from the Cisco SensorBase Network at regular intervals. The Cisco SensorBase Network contains detailed information about known threats on the Internet, including serial attackers, botnet harvesters, malware outbreaks, and dark nets. The IPS uses this information to filter out the worst attackers before they have a chance to attack critical assets. It then incorporates the global threat data into its system to detect and prevent malicious activity even earlier.

IPS Global Correlation is an important improvement in the basic functions of IPS because it enables the system to understand the world in which it operates: an understanding of who the attacker is and whether the attacker has a record of bad behavior. With Global Correlation, the sensor does not have to rely on just the data in the packet or connection to make a decision about the intent of the activity and determine whether the activity is malicious. Now, the sensor can look at a ping sweep and know that the source of the ping sweep does not have a negative reputation, but later can look at another ping sweep and see that the source is a known malicious site with a history of web attacks, and the sensor can block access to and from that site. Global Correlation gives users greater confidence in the actions the sensor takes because these actions are applied to attackers that have shown a predisposition for malicious behavior.

Global Correlation provides a process through which security data is collected for IP addresses and a reputation score is developed for each IP address globally by Cisco. Cisco IPS 7.0 uses this reputation data in two ways: for its reputation filters and for Global Correlation inspection.

- Reputation filters are used to block a subset of IP networks that are owned wholly by malicious groups or were unused and have been hijacked. This first line of defense helps prevent malicious contact ranging from spam to intelligence gathering in preparation for directed attacks. Reputation filters also prevent attempts by botnets to phone home if the botnet controller machine resides in one of these networks.
- Global Correlation inspection uses reputation scores for normal IP addresses to increase the percentage of attacks that the sensor can block. First, the sensor must detect some sort of malicious activity and fire an event as a result. When an event is triggered, that event is processed to determine whether the attacker's IP address has a negative reputation and to what degree. If the event is sourced from an attacker with a negative reputation, the sensor will add risk to the event, raising its risk rating and making it more likely that the sensor will deny the event. This enables the sensor to deny packets and attackers based on the fact that the event has a negative reputation in addition to a high risk rating calculated on the sensor.

Once Global Correlation is configured, the IPS works in the following manner:

Step 1 When a packet enters the sensor, the first check is against the preprocessor, which performs Layer-2 packet normalization and atomic signature checks. Here, the packet header is processed to help ensure that the packet is an IP packet, that the header is not incorrectly formed, and that the packet does not match any atomic signatures.

Step 2 The packet is sent through the Cisco IPS reputation filters.

a. Packets that match are discarded immediately, assuming that the reputation filters are enabled and not in Audit mode.

b. Packets that do not match go to the inspection engines, starting with the Layer 3 and 4 normalization engine, then all the signature engines, and then anomaly detection.

Step 3 If any events are triggered, alerts are sent to the Global Correlation inspection processor, where the source IP address for any alert is checked for negative reputation, and the risk rating is modified and actions are added as appropriate.

Step 4 Any actions assigned to alerts are processed and acted upon, including event action overrides to add new actions and event action filters to remove actions.

Reputation Filters

Cisco IPS reputation filters use a list of hundreds of networks that can be safely blocked because they are not owned by any legitimate source. The reputation of the networks on this list can be considered to be –10. This list includes only IP networks consisting entirely of stolen, "zombie" address blocks and address blocks controlled entirely by malicious organizations. Individual IP addresses are never found on this list. Because there is no way that a legitimate IP address can go from a positive or neutral reputation and then, because of malicious activity, earn a place on the Cisco IPS reputation filter list, users can confidently block all activity to and from networks on this list.

The primary purpose of the IPS reputation filters is to provide protection from direct scanning, botnet harvesting, spamming, and distributed denial-of-service (DDoS) attacks originating from these malicious address blocks, and from connections being attempted back to these networks from systems already infected.

Note There is currently no capability to view the networks on this list, but the networks that are being blocked get logged by the sensor in the Statistics section of Cisco IPS Manager Express (IME).

The only user configuration required for reputation filters is enabling or disabling them and specifying whether Global Correlation is set to Audit mode (a global configuration setting for the entire sensor). In Audit mode, the sensor will report potential deny actions due to reputation filters instead of actually denying the activity.

Global Correlation Inspection

The primary activity of an IPS sensor is detection of malicious behavior. After the packet goes through the IPS reputation filter process, the signature inspection occurs. This involves inspection of packets flowing through the sensor by the various engines looking for the various types of malicious behavior. Alerts that are created are passed to the Global Correlation inspection process for reputation lookups.

When an event occurs, the Global Correlation inspection process performs a local lookup of the source (attacker) IP address of the event in its reputation database. This lookup process returns a value ranging from –.1 to –10; the more negative the value, the more negative the reputation of the source IP address. This reputation score is calculated for Cisco IPS sensors using the data in Cisco SensorBase and is sent to the sensor as a reputation update. If an IP address returns no value for reputation, than it is considered to be neutral. Cisco IPS, unlike E-mail and web security reputation applications, has no concept of positive reputation. When an event is checked for reputation, this checking occurs entirely on the sensor using data downloaded previously from Cisco SensorBase. Unlike other devices, the sensor will not send a live request for information about an IP address that it has just seen. It will look in the data that it has, and if it finds the address, it will use that data; otherwise, the sensor will assume that the address has a neutral reputation.

Global Correlation inspection has three modes of primary operation: Permissive, Standard (default), and Aggressive; you can also select Off:

- Permissive mode tells the sensor to adjust the risk rating of an event, but not to assign separate reputation only actions to the event.
- Standard mode tells the sensor to adjust the risk rating and to add a Deny Packet action due to reputation if the risk rating is greater than or equal to 86. It also adds a Deny Attacker action due to reputation if the risk rating is greater than or equal to 100.
- Aggressive mode also adjusts the risk rating due to reputation, adds a Deny Packet action due to reputation if the risk rating is greater than or equal 83, and adds a Deny Attacker action due to reputation if the risk rating is greater than or equal to 95.
- Selecting Off in the Global Correlation Inspection window prevents the sensor from using updates from Cisco SensorBase to adjust reputation.

If Global Correlation inspection is enabled and an event is generated by an attacker with a negative reputation, the risk rating for the event will be elevated by a certain amount that is determined by a statistical formula. The amount by which the risk rating is raised depends on the original risk rating of the event and the reputation of the attacker.

Network Participation and Correlation Updates

The IPS sensor pulls reputation information for addresses on the global Internet from Cisco SensorBase. When the sensor is configured initially, a DNS server needs to be configured for the sensor to use to connect to Cisco SensorBase or an HTTP or HTTPS proxy (that has DNS configured) needs to be configured. After the sensor has this information, the sensor will make an outbound connection to check for the latest updates from Cisco SensorBase. It will initiate an HTTPS request to Cisco SensorBase update servers and download a manifest that contains the latest versions of the files related to Global Correlation. The sensor will check Cisco SensorBase every 5 minutes for updates. If changes are needed, the sensor will perform a DNS lookup of the server name returned in the initial request. This lookup will return the location of the server nearest to the sensor. The sensor will then initiates an HTTP connection

that will actually transfer the data. The size of a full update is about 2 MB; incremental updates average about 100 KB. If a sensor loses connection to Cisco SensorBase, Global Correlation information will begin to timeout within days, and sensor health will change accordingly.

The other component of Global Correlation is network participation. This feature sends data from events that the sensor fires back to Cisco SensorBase to adjust the reputation of IP addresses; this information is then packaged in future reputation data downloads from Cisco SensorBase. The sensor passes this information back to Cisco SensorBase according to the sensor configuration. The possible configuration options are Off, Partial, and Full.

- With the Off (default) setting, the sensor will not send back any data. The sensor will still receive reputation data, and this setting does not affect its use of that data except that the reputations of addresses attacking the network being protected will not be influenced by their generation on the sensor.
- With the Partial setting, the sensor will send back alert information. This information consists of protocol attributes such as the TCP maximum segment size and TCP options string, the signature ID and risk rating of the event, the attacker IP address and port, and Cisco IPS performance and deployment mode information.
- The Full setting adds victim IP address and port information to the information reported with the Partial setting.

Note No actual packet content information is sent to Cisco. In addition, events having RFC 1918 addresses, because they are not unique, are not considered interesting. So all events reported to Cisco SensorBase will have any such IP address information stripped from the reported data.

The mechanism used to update Cisco SensorBase with new attack information is fairly straightforward. The sensor takes event information, parses out the important pieces of data, and buffers this data for transmission back to Cisco SensorBase. It sends this data in the form of an HTTPS connection that it initiates on average every 10 minutes. The average size of an update is 2 to 4 KB, with weekly averages of about 0.5 to 1 MB. Some higher-volume sensors have average update sizes of about 50 KB, with weekly totals in the 45-MB range. Sensors with very high alert volumes can have average update sizes of about 850 KB, with weekly totals of up to 900 MB; these sensors, though, are at the extreme end of the range.

For more information on IPS Global Correlation including configuration information, see the following URL:
http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/cli/cli_collaboration.html.

E-mail Security Guidelines

The Small Enterprise Design Profile implements a Cisco Ironport C Series E-mail Security Appliance (ESA) at the DMZ with the purpose of inspecting E-mails and eliminating threats such as E-mail spam, viruses, and worms. The ESA can be described as a firewall and threat monitoring system for Simple Mail Transfer Protocol (SMTP) traffic (TCP port 25). Logically speaking, the ESA acts as a Mail Transfer Agent (MTA) within the E-mail delivery chain, as illustrated in Figure 5-5. Upon reception, E-mails are evaluated using a reputation score mechanism based on the SensorBase network. The SensorBase network is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware and other and abnormal behavior. The network is composed of Cisco IronPort appliances, Cisco ASA and Cisco IPS appliances and modules installed in more than 100,000 organizations worldwide, providing a large and diverse sample of Internet traffic patterns. By leveraging the SensorBase Network, messages originating from domain names or servers known to be the source of spam or malware, and therefore with a low reputation score, are automatically dropped or quarantined

by preconfigured reputation filters. Optionally, the enterprise may choose to implement some of the other functions offered by the ESA appliance, including anti-virus protection with virus outbreak filters and embedded anti-virus engines (Sophos and McAfee), encryption to ensure the confidentiality of messages, and data loss prevention (DLP) for E-mail to detect the inappropriate transport of sensitive information.

Note Alternatively, Cisco offers managed hosted and hybrid hosted E-mail security services. These services are provided through a dedicated E-mail infrastructure hosted in a network of Cisco data centers. For more information, refer to http://www.cisco.com/go/designzone.

Figure 5-5 E-mail Delivery Chain

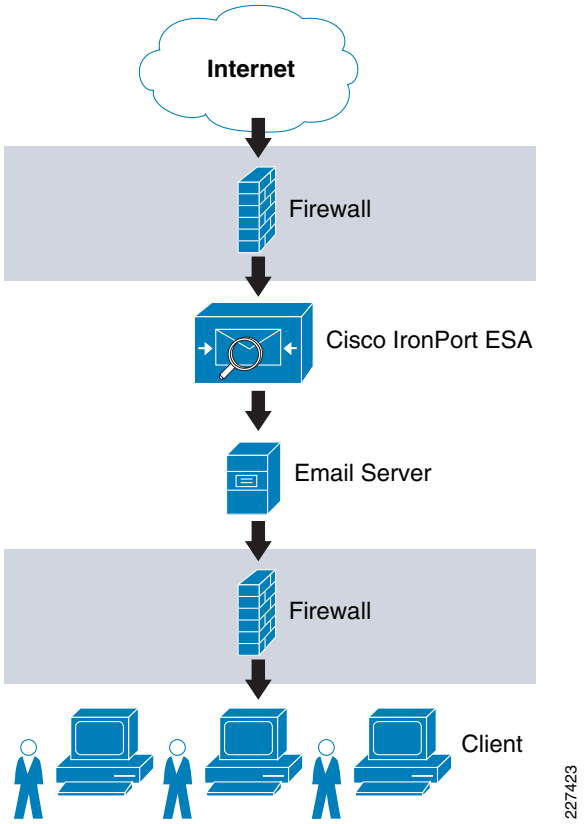

Note Figure 5-5 shows a logical implementation of a DMZ hosting the E-mail server and ESA appliance. This can be implemented physically by either using a single firewall or two firewalls in sandwich.

There are multiple deployment approaches for the security appliance depending on the number of interfaces used (see Figure 5-6):

- *Dual-armed configuration*—Two physical interfaces used to serve a public mail listener and a private mail listener, each one configured with a separate logical IP address. The public listener receives E-mail from the Internet and directs messages to the internal E-mail servers; while the private listener receives E-mail from the internal servers and directs messages to the Internet. The public listener interface may connect to the DMZ, while the private listener interface may connect to the inside of the firewall.

- *One-armed configuration*—A single ESA interface configured with a single IP address and used for both incoming and outgoing E-mail. A public mail listener is configured to receive and relay E-mail on that interface. The best practice is to connect the ESA interface to the DMZ where the E-mail server resides.

For simplicity, the Small Enterprise Design Profile implements the ESA appliance with a single interface. In addition, using a single interface leaves other data interfaces available for redundancy.

Figure 5-6 Common ESA Deployments

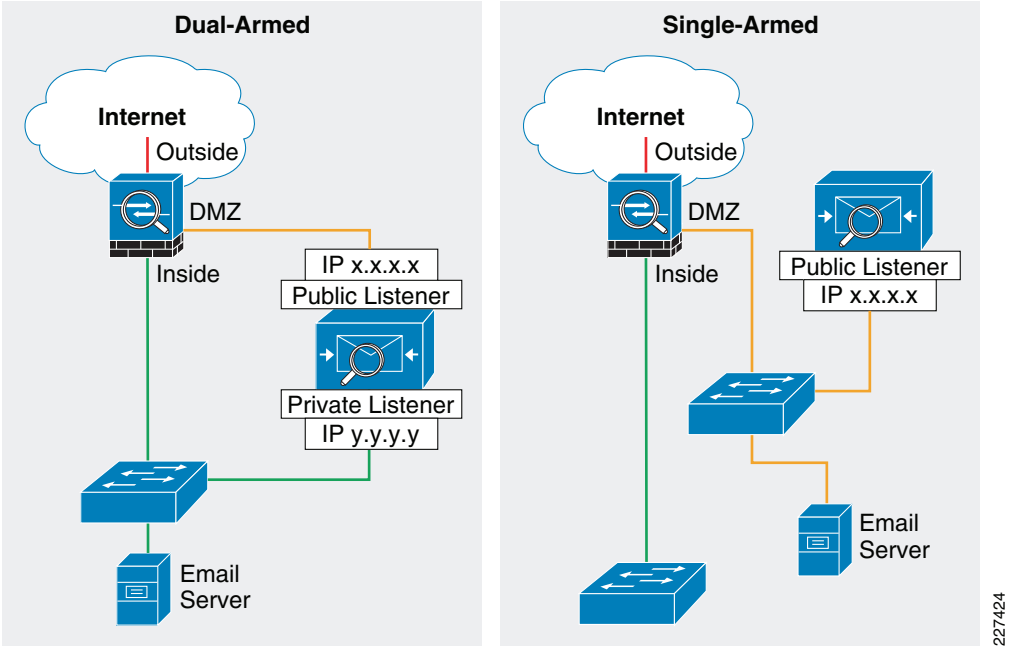

Figure 5-7 illustrates the logical location of the ESA within the E-mail flow chain.

Figure 5-7 Typical Data Flow for Inbound E-mail Traffic

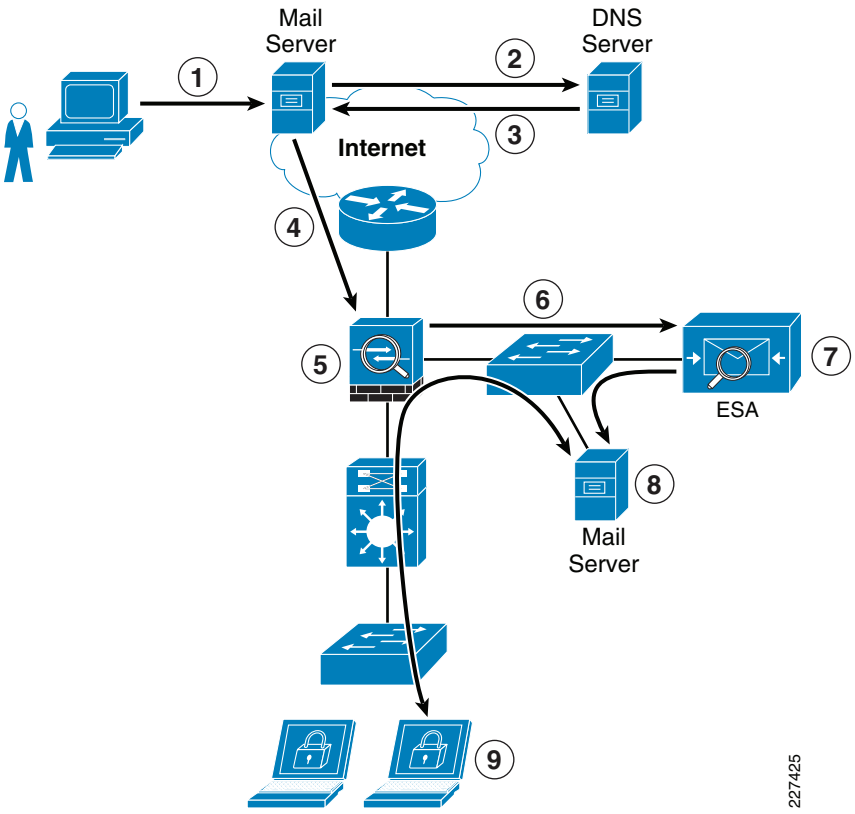

The following steps explain what is taking place in Figure 7:

Step 1 Sender sends an E-mail to xyz@domain X.

Step 2 What's the IP address of domain X?

Step 3 It's a.b.c.d (public IP address of ESA).

Step 4 E-mail server sends message to a.b.c.d using SMTP.

Step 5 Firewall permits incoming SMTP connection to the ESA, and translates its public IP address.

Step 6 ESA performs a DNS query on sender domain and checks the received IP address in its reputation database, and drops, quarantines E-mail based on policy.

Step 7 ESA forwards E-mail to preconfigured inbound E-mail server.

Step 8 E-mail server stores E-mail for retrieval by receiver.

Step 9 Receiver retrieves E-mail from server using POP or IMAP.

The Cisco IronPort ESA appliance functions as a SMTP gateway, also known as a mail exchange (MX). The following are the key deployment guidelines:

- Ensure that the ESA appliance is both accessible via the public Internet and is the first hop in the E-mail infrastructure. If you allow another MTA to sit at your network's perimeter and handle all external connections, then the ESA appliance will not be able to determine the sender's IP address.

The sender's IP address is needed to identify and distinguish senders in the Mail Flow Monitor, to query the SensorBase Reputation Service for the sender's SensorBase Reputation Service Score (SBRS), and to improve the efficacy of the anti-spam and virus outbreak filters features.

- Features like Cisco IronPort Anti-Spam, Virus Outbreak Filters, McAfee Antivirus and Sophos Anti-Virus require the ESA appliance to be registered in DNS. To that end, create an A record that maps the appliance's hostname to its public IP address, and an MX record that maps the public domain to the appliance's hostname. Specify a priority for the MX record to advertise the ESA appliance as the primary (or backup during testing) MTA for the domain. A static address translation entry needs to be defined for the ESA public IP address on the Internet firewall if NAT is configured.
- Add to the Recipient Access Table (RAT) all the local domains for which the ESA appliance will accept mail. Inbound E-mail destined to domains not listed in RAT will be rejected. External E-mail servers connect directly to the ESA appliance to transmit E-mail for the local domains, and the ESA appliance relays the mail to the appropriate groupware servers (for example, Exchange™, Groupwise™, and Domino™) via SMTP routes.
- For each private listener, configure the Host Access Table (HAT) to indicate the hosts that will be allowed to send E-mails. The ESA appliance accepts outbound E-mail based on the settings of the HAT table. Configuration includes the definition of Sender Groups associating groups or users, upon which mail policies can be applied. Policies include Mail Flow Policies and Reputation Filtering. Mail Flow Policies are a way of expressing a group of HAT parameters (access rule, followed by rate limit parameters and custom SMTP codes and responses). Reputation Filtering allows the classification of E-mail senders and to restrict E-mail access based on sender's trustworthiness as determined by the IronPort SensorBase Reputation Service.
- Define SMTP routes to direct E-mail to appropriate internal mail servers.
- If an out-of-band (OOB) management network is available, use a separate interface for administration.

A failure on the ESA appliance may cause service outage; therefore, a redundant design is recommended. There are multiple ways to implement redundancy:

- *IronPort NIC Pairing*—Redundancy at the network interface card level by teaming two of the Ethernet interfaces on the ESA appliance. If the primary interface fails, the IP addresses and MAC address are assumed by the secondary.
- *Multiple MTAs*—Consists of adding a second ESA appliance or MTA with an equal cost secondary MX record.
- *Load Balancer*—A load balancer such as Cisco ACE Application Control Engine (ACE) load-balances traffic across multiple ESA appliances.

IronPort NIC pairing is the most cost-effective solution (see Figure 5-8), because it does not require the implementation of multiple ESA appliances and other hardware. It does not, however, provide redundancy in case of chassis failure.

Figure 5-8 Cisco IronPort ESA NIC Pairing

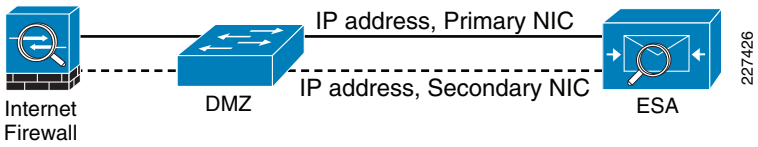

Finally, the Internet firewall should be configured to accommodate traffic to and from the Cisco IronPort ESA deployed at the DMZ. The protocols and ports to be allowed vary depending on the services configured on the appliance. For details, refer to the *Cisco IronPort User's Guide* at the following URL: http://www.ironport.com/support/

The following are some of the most common services required:

- Outbound SMTP (TCP/25) from ESA to any Internet destination
- Inbound SMTP (TCP/25) to ESA from any Internet destination
- Outbound HTTP (TCP/80) from ESA to downloads.ironport.com and updates.ironport.com
- Outbound SSL (TCP/443) from ESA to updates-static.ironport.com and phonehome.senderbase.org
- Inbound and Outbound DNS (TCP and UDP port 53)
- Inbound IMAP (TCP/143), POP (TCP/110), SMTP (TCP/25) to E-mail server from any internal client

Also remember that if the ESA is managed in-band, appropriate firewall rules need to be configured to allow traffic such as SSH, NTP, and syslog.

For more information on how to configure the Cisco Ironport ESA, see the following guides:

- *Cisco SAFE Reference Guide*—http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html
- *Cisco IronPort ESA User Guide*—http://www.ironport.com/support

Web Security Guidelines

The Small Enterprise Design Profile implements a Cisco IronPort S Series Web Security Appliance (WSA) to block access to sites with non-business related content, and to protect the enterprise from web-based malware and spyware.

Cisco IronPort WSA's protection relies in two independent services:

- *Web Proxy*—This provides URL filtering, web reputation filters, and optionally anti-malware services. The URL filtering capability defines the handling of each web transaction based on the URL category of the HTTP requests. Leveraging the SensorBase network, the web reputation filters analyze the web server behavior and characteristics to identify suspicious activity and protect against URL-based malware. The anti-malware service leverages anti-malware scanning engines such as Webroot and McAfee to monitor for malware activity.
- *Layer 4 Traffic Monitoring (L4TM)*—Service configured to monitor all Layer-4 traffic for rogue activity and to detect infected clients.

Note The SensorBase network is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware, and other abnormal behavior. The network is composed of the Cisco IronPort appliances, Cisco ASA, and Cisco IPS appliances and modules installed in more than 100,000 organizations worldwide, providing a large and diverse sample of Internet traffic patterns.

As the small enterprise network design assumes a centralized Internet connection, the WSA is implemented at the core/distribution layer of the main site network. This allows the inspection and enforcement of web access policies to all users residing at any of the enterprise locations. Logically, the WSA sits in the path between web users and the Internet, as shown in Figure 5-9.

Figure 5-9 Cisco IronPort WSA

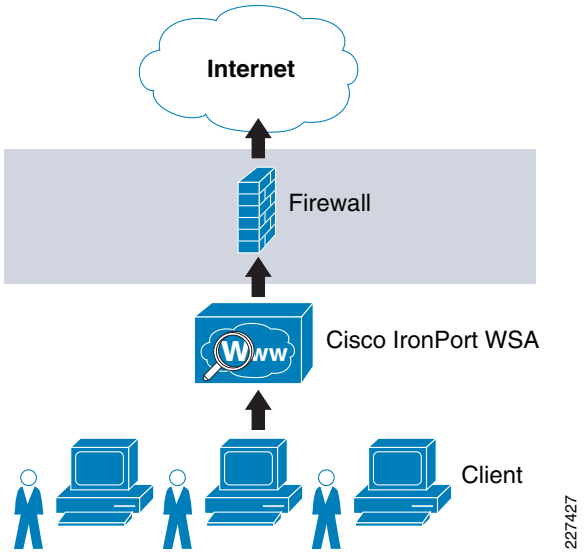

There are two deployment modes for the Web Proxy service:

- *Explicit Forward Proxy*—Client applications, such as web browsers, are aware of the Web Proxy and must be configured to point to the WSA. The web browsers can be either configured manually or by using Proxy Auto Configuration (PAC) files. The manual configuration does not allow for redundancy, while the use of PAC files allows the definition of multiple WSAs for redundancy and load balancing. If supported by the browser, the Web Proxy Autodiscovery Protocol (WPAD) can be used to automate the deployment of PAC files. WPAD allows the browser to determine the location of the PAC file using DHCP and DNS lookups.
- *Transparent Proxy*—Client applications are unaware of the Web Proxy and do not have to be configured to connect to the proxy. This mode requires the implementation of a Web Cache Communications Protocol (WCCP) enabled device or a Layer-4 load balancer in order to intercept and redirect traffic to the WSA. Both deployment options provide for redundancy and load balancing.

Explicit forward proxy mode requires the enterprise to have control over the configuration of the endpoints, which may not be always possible. For example, the enterprise may allow the use of personal laptops, smart-phones and other devices outside the company's administration. Transparent proxy mode, on the other hand, provides a transparent integration of WSA without requiring any configuration control over the endpoints. It also eliminates the possibility of users reconfiguring their web browsers to bypass the appliance without knowledge of the administrators. For these reasons, the Small Enterprise Design Profile implements transparent proxy with WCCP. In this configuration, the Cisco ASA at the Internet perimeter is leveraged as a WCCP server while the WSA act as a WCCP Traffic Processing Entity.

The Cisco ASA uses WCCP version 2, which has a built-in failover and load balancing mechanism. Per WCCPv2 specification, multiple appliances (up to 32 entities) can be configured as part of the same service group. HTTP and HTTPS traffic is load-balanced across the active appliances based on source and destination IP addresses. The server (Cisco ASA) monitors the availability of each appliance in the group, and can identify appliance failures within 30 seconds. After failure, traffic is redirected across the remaining active appliances. In the case where no appliances are active, WCCP takes the entire service group offline and subsequent requests bypass redirection. In addition, WCCPv2 supports MD5 authentication for the communications between WCCP server and WSA appliances.

Note In the event the entire service group fails, WCCP automatically bypasses redirection, allowing users to browse the Internet without the Web controls. In case it is desired to handle a group failure by blocking all traffic, an outbound ACL may be configured on the Cisco ASA outside interface to permit HTTP/HTTPS traffic originated from the WSA appliance itself and to block any direct requests from clients. The ACL may also have to be configured to permit HTTP/HTTPS access from IPS and other systems requiring such access.

WCCPv2 supports Generic Route Encapsulation (GRE) and Layer-2-based redirection; however, the Cisco ASA only supports GRE. In addition, WCCP is supported only on the ingress of an interface. The only topology supported is one where both clients and WSA are reachable from the same interface, and where the WSA can directly communicate with the clients without going through the Cisco ASA. For these reasons, the WSA appliance is deployed at the inside segment of the Cisco ASA.

Figure 5-10 illustrates the how WCCP redirection works in conjunction with Cisco ASA.

Figure 5-10 WCCP Redirection

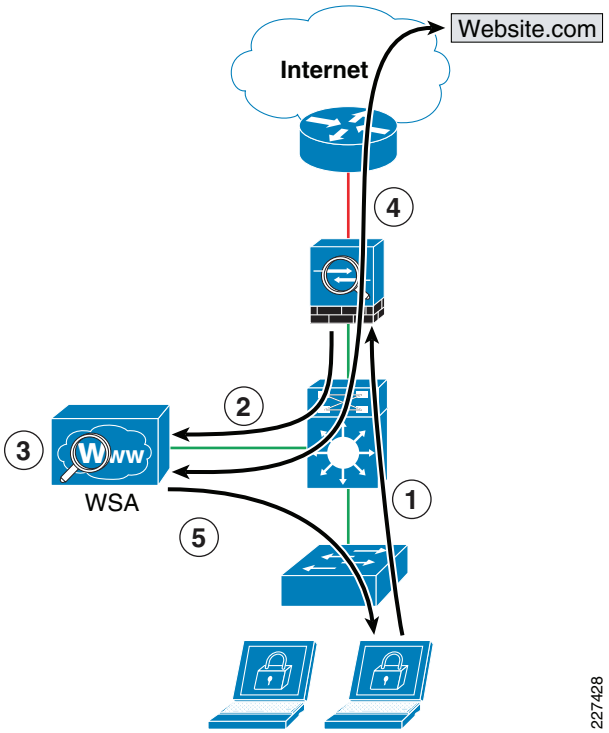

The following steps describe what takes place in Figure 5-10:

Step 1 Client's browser requests connection to http://website.com.

Step 2 Cisco ASA intercepts and redirects HTTP requests over GRE.

Step 3 If content not present in local cache, WSA performs a DNS query on destination domain and checks the received IP address against URL and reputation rules, and allows/denies request accordantly.

Step 4 WSA fetches content from destination web site.

Step 5 Content is inspected and then delivered directly to the requesting client.

The WSA appliance may also be configured to control and block peer-to-peer file-sharing and Internet applications such as AOL Messenger, BitTorrent, Skype, Kazaa, etc. The way WSA handles these applications depends on the TCP port used for transport:

- *Port 80*—Applications that use HTTP tunneling on port 80 can be handled by enforcing access policies within the web proxy configuration. Application access may be restricted based on applications, URL categories, and objects. Applications are recognized and blocked based on their user agent pattern, and by the use of regular expressions. The user may also specify categories of URL to block, including the predefined *chat* and *peer-to-peer* categories. Custom URL categories may also be defined. Peer-to-peer access may also be filtered based on object and MIME Multipurpose Internet Mail Extensions (MIME) types.
- *Ports other than 80*—Applications using ports other than 80 can be handled with the L4TM feature. L4TM blocks access to a specific application by preventing access to the server or block of IP addresses to which the client application must connect.

Note The Cisco IPS appliances and modules, and the Cisco ASA (using the modular policy framework), may also be used to block peer-to-peer file sharing and Internet applications.

The following are the guidelines for implementing a Cisco IronPort WSA appliance with WCCP on a Cisco ASA:

- Deploy WSA on the inside of the firewall so that the WSA can communicate with the clients without going through the firewall.
- Implement MD5 authentication to protect the communications between the Cisco ASA and the WSA(s).
- Configure a redirect-list on the firewall to indicate what traffic needs to be redirected. Make sure the WSA is always excluded from redirection.
- Ingress ACL on the firewall takes precedence over WCCP redirection, so make sure the ingress ACL is configured to allow HTTP and HTTPS traffic from clients and the WSA itself.
- In an existing proxy environment, deploy the WSA downstream from the existing proxy servers (closer to the clients).
- Cisco ASA does not support WCCP IP source address spoofing, therefore any upstream authentication or access controls based on client IP addresses are not supported. Without IP address spoofing, requests originating from a client are sourced with the IP address of the Web Proxy, and not the one of the client.
- TCP intercept, authorization, URL filtering, inspect engines, and IPS features do not apply to redirected flows of traffic served by the WSA cache. Content requested by the WSA is still subject to all the configured features on the firewall.
- Configure WSA access policies to block access to applications (AOL Messenger, Yahoo Messenger, BitTorrent, Kazaa, etc.) and URL categories not allowed by the enterprise Internet access policies.
- If an out-of-band (OOB) management network is available, use a separate interface for administration.

Note WCCP, firewall, and other stateful features usually require traffic symmetry, whereby traffic in both directions should flow through the same stateful device. The Small Enterprise Design Profile is designed with a single Internet path ensuring traffic symmetry. Care should be taken when implementing active-active firewall pairs as they may introduce asymmetric paths.

The Layer-4 Traffic Monitor (L4TM) service is deployed independently from the Web Proxy functionality, and its mission is to monitor network traffic for rogue activity and for any attempts to bypass port 80. L4TM works by listening to all UDP and TCP traffic and by matching domain names and IP addresses against entries in its own database tables to determine whether to allow incoming and outgoing traffic. The L4TM internal database is continuously updated with matched results for IP addresses and domain names. Additionally, the database table receives periodic updates from the IronPort update server (https://update-manifests.ironport.com).

For more information on how to configure the L4TM feature on Cisco Ironport WSA, see the following guides:

- *Cisco SAFE Reference Guide*—http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html
- *Cisco IronPort WSA User Guide*—http://www.ironport.com/support

Configuration steps and examples are included in the "Web Security Deployment" section on page 5-60.

Serverfarm Protection

Small enterprise networks typically include a serverfarm at the main site that hosts the systems that serve business applications and store the data accessible to internal users. The infrastructure supporting it may include application servers, the storage media, routers, switches, load balancers, off-loaders, application acceleration devices and other systems. In addition, they may also host foundational services as part of the enterprise network such as identity and security services, unified communication services, mobility services, video services, partner applications, and other services.

Depending on the size of the enterprise network, the serverfarm may be constructed following different design models. Figure 5-11 illustrates a collapsed design, and the less common for small enterprise networks, multi-tier design. In the collapsed design all services are hosted in a shared physical server-farm, and high availability is achieved by using redundant processors and interfaces. Large enterprises may implement a more scalable multi-tier design serverfarm with chassis redundancy.

Figure 5-11 Serverfarm Designs

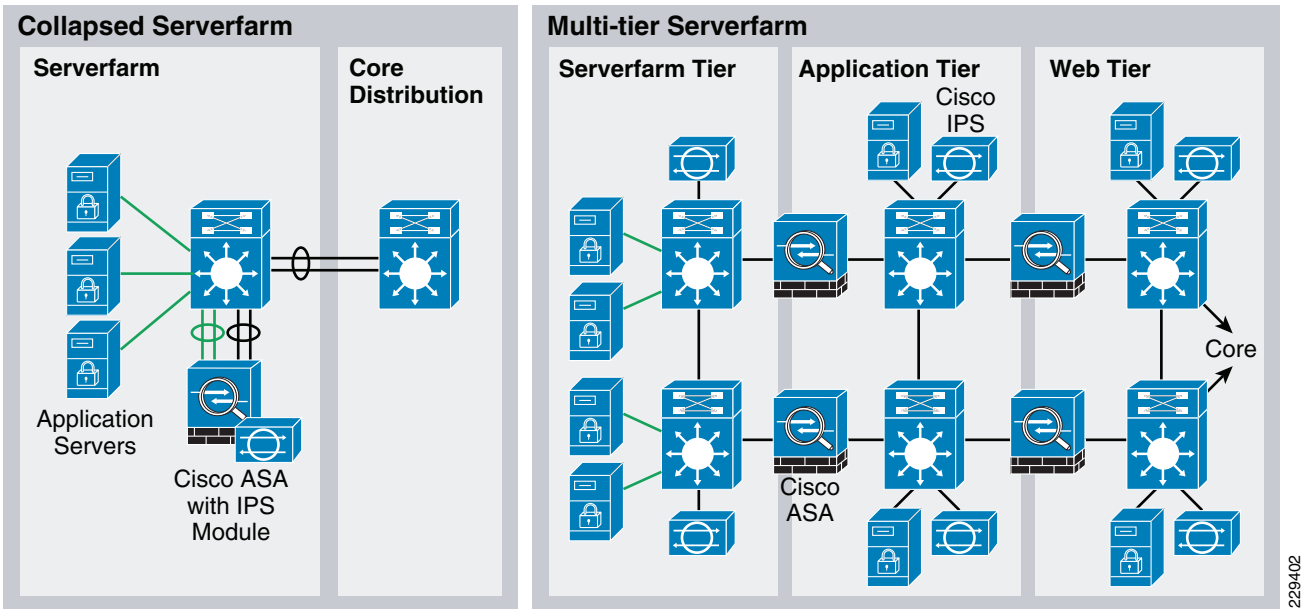

Independently from the design model adopted by the enterprise, the following are the primary security guidelines for the serverfarm design:

- *Network Foundation Protection*— All infrastructure equipment should be protected following the Network Foundation Protection best practices described earlier in this document. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the switching and routing planes.
- *Firewall*—A stateful firewall may be deployed to limit access to only the necessary applications and services, and for the intended users. The firewall should be configured to control and inspect both traffic entering and leaving the server farm segments. The firewall may also be leveraged to ensure the appropriate segregation between application layers or groups. In addition, the firewall's deep packet inspection may be used to mitigate DoS attacks and enforce protocol compliance.
- *Intrusion Prevention*— An IPS module on the Cisco ASA or a separate IPS appliance may be implemented for enhanced threat detection and mitigation. The IPS is responsible for identifying and blocking anomalous traffic and packets recognized as well-known attacks. The Cisco IPS may be configured either in inline or promiscuous mode. When deployed in inline mode, the Cisco IPS is placed in the traffic path and is capable of stopping malicious traffic before it reaches the intended target.
- *Service Isolation*— Services and applications serving different group of users or under different security requirements should be properly isolated. Isolation helps prevent data leakage and contain possible compromises from spreading across different server farm groups. Logical isolation may be achieved by separating applications and services in different VLANs and by assigning them into different firewall interfaces (physical or logical). This is illustrated in Figure 5-12.
- *Switch Security*—Private VLANs, port security, storm control and other switch security features may be leveraged to mitigate spoofing, man-in-the-middle, denial-of-service and other network-based attacks directed to the serverfarm applications and the switching infrastructure.
- *Endpoint Protection*— Servers residing at the different layers should be protected with host-based IPS or other endpoint security software.

Figure 5-12 Service Isolation

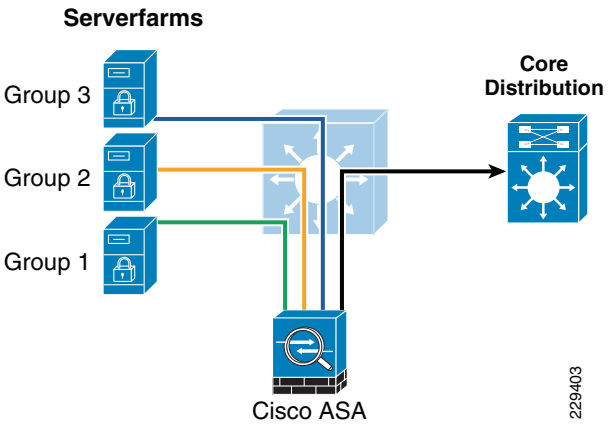

SSL termination and inspection, Web Application Firewall (WAF), Application Control Engine (ACE), and other solutions may be leveraged to complement the guidelines described above. For a more detailed discussion of serverfarm security, refer to "Chapter 4, Intranet Data Center" of the *Cisco SAFE Reference Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap4.html

Network Access Security and Control

Some of the most vulnerable points of the network are the access edges where users connect to the network. With the proliferation of wireless networks, increased use of laptops and smart mobile devices, the enterprise cannot simply rely on physical controls hoping to prevent unauthorized systems from being plugged into the ports of the access switches. Protection should rather be embedded into the network infrastructure, leveraging the native security features available in switches, routers, and WLAN systems. Furthermore, the network infrastructure should also provide dynamic identity or role-based access controls for all systems attempting to gain access.

Implementing role-based access controls for users and devices helps reduce the potential loss of sensitive information by enabling enterprises to verify a user or device identity, privilege level, and security policy compliance before granting network access. Security policy compliance could consist of requiring antivirus software, OS updates or patches. Unauthorized or noncompliant devices can be placed in a quarantine area where remediation can occur prior to gaining access to the network.

The Small Enterprise Design Profile achieves access security and control by leveraging the following technologies:

- Catalyst Integrated Security Features (CISF), wired
- Cisco Unified Wireless Network (CUWN) Integrated Security Features, wireless
- Cisco NAC Appliance, wired and wireless
- Cisco Identity-Based Network Networking Services (IBNS), wired and wireless

Catalyst Integrated Security Features

Catalyst Integrated Security Features (CISF) is a set of native security features available on Cisco Catalyst Switches and designed to protect the access infrastructure and users from spoofing, man-in-the-middle, DoS and other network-based attacks. CISF includes features such as private VLANs, port security, DHCP snooping, IP Source Guard, secure Address Resolution Protocol (ARP) detection, and Dynamic ARP Inspection (DAI). CISF features are considered to be part of a security baseline and should be deployed on all access ports.

- *Port Security*—Mitigates MAC flooding and other Layer 2 CAM overflow attacks by restricting the MAC addresses that are allowed to send traffic on a particular port. After Port Security is enabled on a port, only packets with a permitted source MAC address are allowed to pass through the port. A permitted MAC address is referred to as a secure MAC address.
- *DHCP snooping*—Inspects and filters DHCP messages on a port to ensure DHCP server messages come only from a trusted interface. Additionally, it builds and maintains a DHCP snooping binding table that contains the MAC address, IP address, lease time, binding type, VLAN number, and interface information corresponding to the local untrusted interfaces of a switch. This binding table is used by the other CISF features.
- *IP Source Guard*—Restricts IP traffic on a port based on DHCP or static IP address MAC bindings to prevent IP spoofing attacks. IP address bindings are validated using information in the DHCP Snooping binding table.
- *Dynamic ARP inspection*—Validates that the source MAC and IP address in an ARP packet received on an untrusted interface matches the source MAC and IP address registered on that interface (using the DHCP snooping binding table) to prevent ARP spoofing and MITM attacks.
- *ARP rate limiting*—Where an excessive rate of ARP request (which must be processed by network hosts CPUs), and the switch responds with access restriction if this rate is exceeded.
- *Storm Control*—Prevents broadcast and multicast storms by monitoring packets passing from an interface to the switching bus and determines whether the packet is unicast, multicast, or broadcast. The switch counts the number of packets of a specified type received within the 1-second time interval and compares the measurement with a predefined suppression-level threshold. When the suppression-level threshold is reached, the port blocks traffic until the traffic falls below the threshold level.

Cisco Unified Wireless Network (CUWN) Integrated Security Features

The Cisco Unified Wireless Network adds to the 802.11 security standards by providing additional security features. Some of these are the WLAN equivalent of CiSF features such as, Dynamic Host Configuration Protocol (DHCP) and Address Resolution Protocol (ARP) protection, peer-to-peer blocking, and access control list and firewall features. Additionally other more WLAN specific features are provided, including Enhanced WLAN security options, wireless intrusion detection system (IDS), client exclusion, rogue AP detection, management frame protection, dynamic radio frequency management, and network IDS integration.

The Cisco Unified Wireless Network solutions are discussed in the *Wireless and Network Security Integration Solution Design Guide* at the following URL:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns820/landing_sec_wireless.html

Cisco NAC Appliance

The Cisco Network Admission Control (NAC) Appliance (formerly known as Cisco Clean Access) uses the network infrastructure to enforce security policy compliance on all devices seeking to access network computing resources. With Cisco NAC Appliance, network administrators can authenticate, authorize, evaluate, and remediate wired, wireless, and remote users and their machines before network access. The NAC Appliance identifies whether networked devices such as laptops, IP phones, or game consoles are compliant with your network security policies, and can repair any vulnerability before permitting access to the network.

The Cisco NAC Appliance provides the following benefits:

- Recognizes users, their devices, and their roles in the network. This first step occurs at the point of authentication, before malicious code can cause damage.
- Evaluates whether machines are compliant with security policies. Security policies can include specific anti-virus or anti-spyware software, OS updates, or patches. Cisco NAC Appliance supports policies that vary by user type, device type, or operating system.
- Enforces security policies by blocking, isolating, and repairing non-compliant machines.
- Non-compliant machines are redirected to a quarantine network, where remediation occurs at the discretion of the administrator.

The NAC solution provides the following four functions, as shown in Figure 5-13:

- Authenticates and authorizes
- Scans and evaluates
- Quarantines and enforces
- Updates and remediates

Figure 5-13 The Four Functions of the NAC Framework

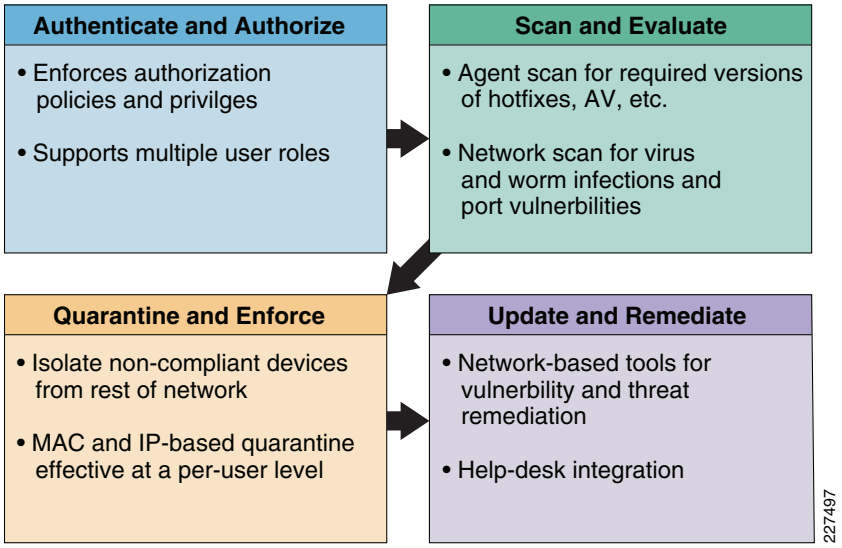

For more details of the NAC Appliance solution, see the following URL: http://www.cisco.com/go/nacappliance.

Cisco NAC Appliance is a network-centric integrated solution administered from the Cisco Clean Access Manager (NAC Manager) web console and enforced through the Clean Access Server (NAC Server) and (optionally) the Clean Access Agent or Cisco NAC Web Agent. Cisco NAC Appliance checks client systems, enforces network requirements, distributes patches and antivirus software, and quarantines vulnerable or infected clients for remediation before clients access the network. Cisco NAC Appliance consists of the components shown in Figure 5-14.

Figure 5-14 NAC Appliance Components

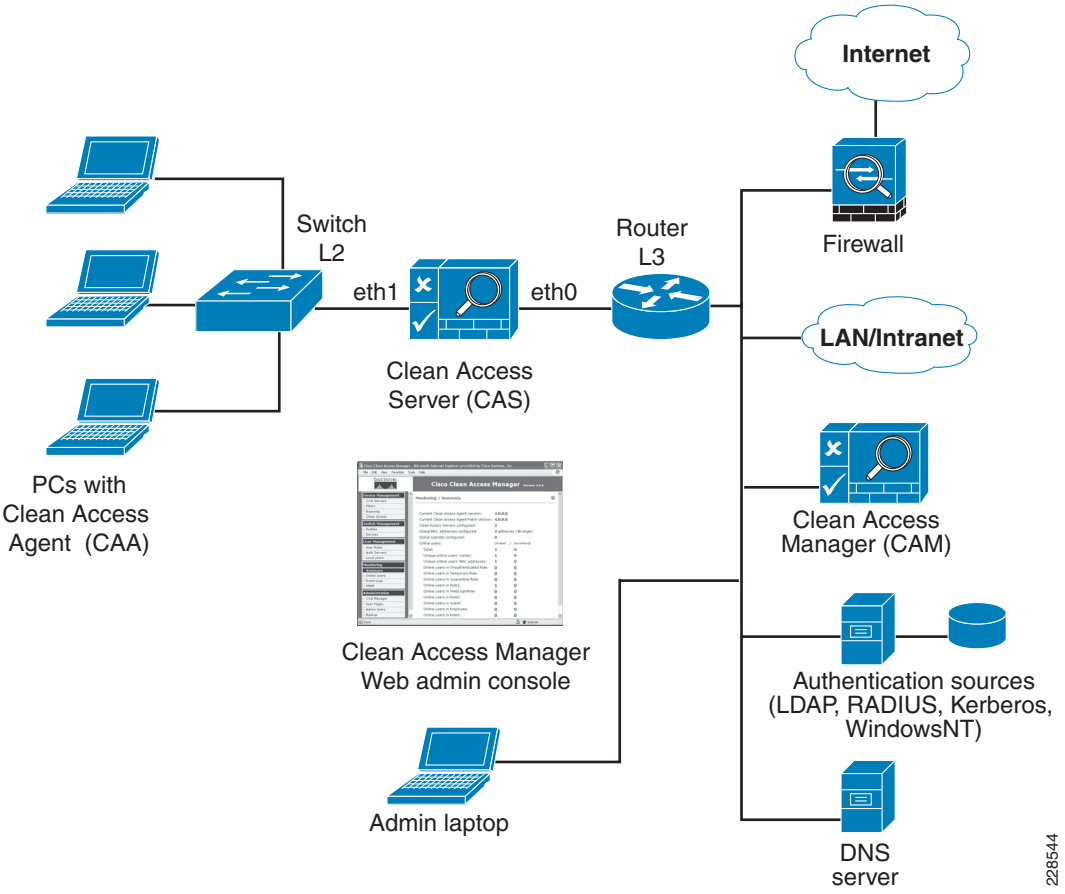

Clean Access Manager (CAM)

The Cisco CAM (also known as NAC Manager) is the administration server for Clean Access deployment. The secure web console of the Clean Access Manager is the single point of management for up to 20 Clean Access Servers in a deployment (or 40 CASs if installing a SuperCAM). For Out-of-Band (OOB) deployment, the web admin console allows you to control switches and VLAN assignment of user ports through the use of SNMP. In the Small Enterprise Design Profile, the CAM would be located at the main site.

Clean Access Server (CAS)

The Cisco CAS (a.k.a NAC Server) is the enforcement server between the untrusted (managed) network and the trusted network. The CAS enforces the policies you have defined in the CAM web admin console, including network access privileges, authentication requirements, bandwidth restrictions, and

Clean Access system requirements. You can install a CAS either as a standalone appliance (like the Cisco NAC-3300 Series) or as a network module (Cisco NME-NAC-K9) in a Cisco ISR chassis and deploy it in-band (always inline with user traffic) or OOB (inline with user traffic only during authentication/posture assessment).

The CAS can also be deployed in Layer 2 mode (users are Layer-2-adjacent to the CAS) or Layer 3 mode (users are multiple Layer-3 hops away from the CAS). You can also deploy several CASs of varying size/capacity to fit the needs of varying network segments. You can install Cisco NAC-3300 Series appliances in your company headquarters core, for example to handle thousands of users and simultaneously install one or more Cisco NAC network modules in ISR platforms to accommodate smaller groups of users at a satellite office, for example.

In the Small Enterprise Design Profile, the CAS would be located at the main site and the remote locations, and it would be used to provide Layer-2 or Layer-3 OOB authentication/posture assessment.

Clean Access Agent (CAA)

CAA is optional read-only agent that resides on Windows clients. It checks applications, files, services, or registry keys to ensure that clients meets your specified network and software requirements prior to gaining access to the network.

Note There is no client firewall restriction with CAA posture assessment. The agent can check the client registry, services, and applications even if a personal firewall is installed and running.

If NAC is implemented as part of the Small Enterprise Design Profile it is recommended that the CAA be used.

Cisco NAC Web Agent

The Cisco NAC Web Agent provides temporal posture assessment for client machines. Users launch the Cisco NAC Web Agent executable, which installs the Web Agent files in a temporary directory on the client machine via ActiveX control or Java applet. When the user terminates the Web Agent session, the Web Agent logs the user off of the network and their user ID disappears from the Online Users list.

Clean Access Policy Updates

Regular updates of prepackaged policies/rules that can be used to check the up-to-date status of operating systems, antivirus (AV), antispyware (AS), and other client software. It provides built-in support for 24 AV vendors and 17 AS vendors.

NAC Appliance Modes and Positioning

NAC Appliance allows multiple deployment options and may be placed at different points in the network. The modes of operation can be generally defined as follows:

- Out-of-band (OOB) virtual gateway
- OOB IP real gateway
- In-band (IB) virtual gateway
- IB real IP gateway

Out-of-Band Modes

Out-of-Band (OOB) deployments require user traffic to traverse through the NAC Appliance only during authentication, posture assessment, and remediation. When a user is authenticated and passes all policy checks, their traffic is switched normally through the network and bypasses the appliance. See Figure 5-15.

Figure 5-15 Layer-2 OOB Topology

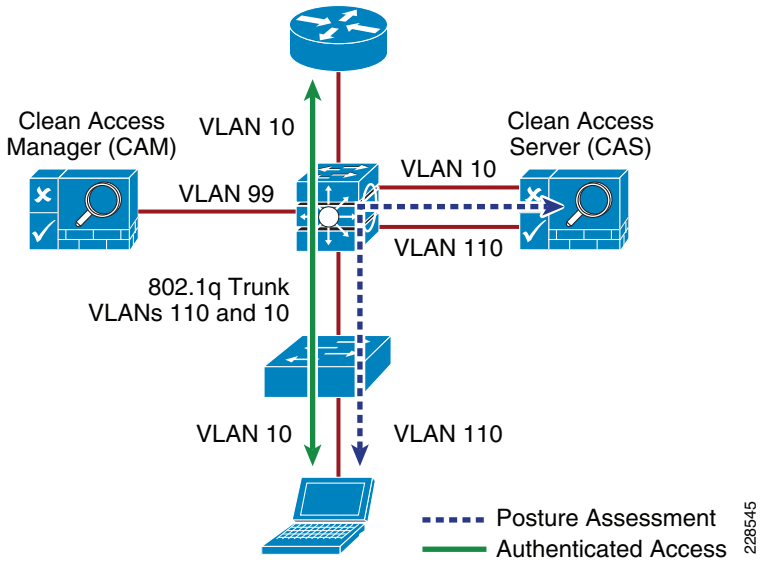

To deploy the NAC Appliance in this manner, the client device must be directly connected to the network via a Catalyst switch port. After the user is authenticated and passes posture assessment, the Clean Access Manager (CAM) instructs the switch to map the user port from an unauthenticated VLAN (which switches or routes user traffic to the NAC) to an authenticated (authorized) VLAN that offers full access privileges. For example, as shown in Figure 5-15, the client PC is connected through VLAN 110 to the NAC Clean Access Server for the authentication and posture assessment, and is moved to VLAN 10 once it successfully completes the authentication and authorization, scan, and evaluation phases of the NAC solution.

In-Band Modes

When the NAC Appliance is deployed in-band, all user traffic, both unauthenticated and authenticated, passes through the NAC Appliance, which may be positioned logically or physically between end users and the network(s) being protected. See Figure 5-16 for a logical in-band topology example and Figure 5-17 for a physical in-band topology example.

Figure 5-16 In-Band Virtual Gateway Topology

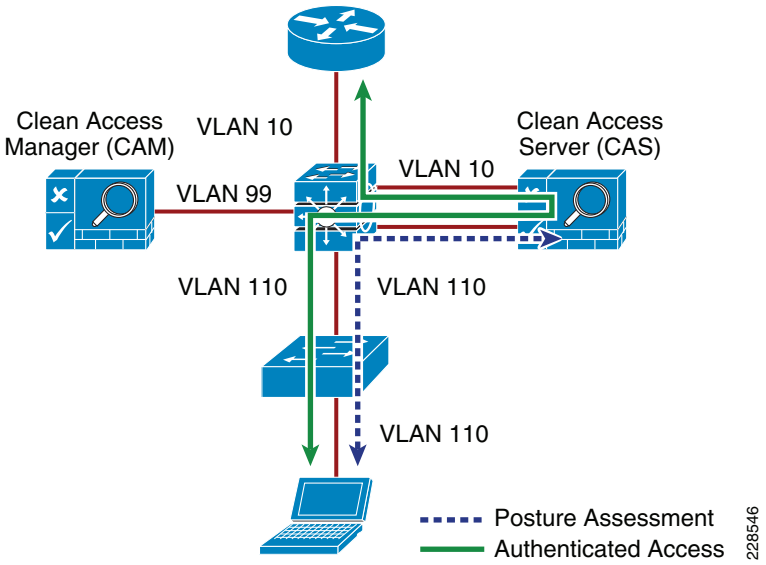

Figure 5-17 Physical In-Band Topology

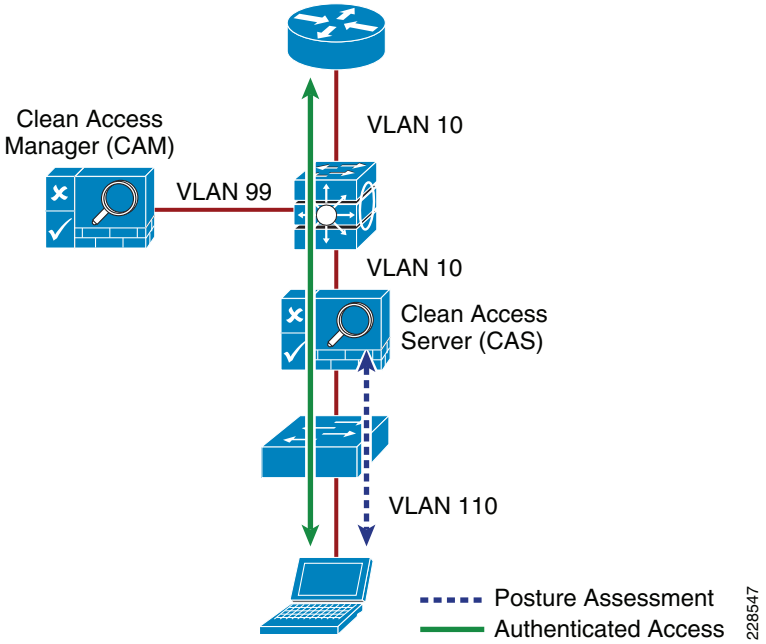

In-Band Virtual Gateway

When the NAC Appliance is configured as a virtual gateway, it acts as a bridge between end users and the default gateway (router) for the client subnet being managed. The following two bridging options are supported by the NAC Appliance:

- *Transparent*—For a given client VLAN, the NAC Appliance bridges traffic from its untrusted interface to its trusted interface. Because the appliance is aware of "upper layer protocols", by default it blocks all traffic except for Bridge Protocol Data Unit (BPDU) frames (spanning tree) and those protocols explicitly permitted in the "unauthorized" role; for example, DNS and DHCP. In other words, it permits those protocols that are necessary for a client to connect to the network, authenticate, undergo posture assessment, and remediation. This option is viable when the NAC Appliance is positioned physically in-band between end users and the upstream network(s) being protected, as shown in Figure 5-17.
- *VLAN mapping*—This is similar in behavior to the transparent method except that rather than bridging the same VLAN from the untrusted side to the trusted side of the appliance, two VLANs are used. For example, Client VLAN 131 is defined for the untrusted interface of the NAC Appliance. There is no routed interface or switched virtual interface (SVI) associated with VLAN 131. VLAN 31 is configured between the trusted interface of the NAC Appliance and the next-hop router interface/SVI for the client subnet. A mapping rule is made in the NAC Appliance that forwards packets arriving on VLAN 131 and forwards them out VLAN 31 by swapping VLAN tag information. The process is reversed for packets returning to the client. Note that in this mode, BPDUs are not passed from the untrusted-side VLANs to their trusted-side counterparts.

The VLAN mapping option is usually selected when the NAC Appliance is positioned logically in-band between clients and the networks being protected. This is the bridging option that should be used if the NAC Appliance is going to be deployed in the virtual gateway mode.

In-Band Real IP Gateway

When the NAC Appliance is configured as a "real" IP gateway, it behaves like a router and forwards packets between its interfaces. In this scenario, one or more client VLAN/subnets reside behind the untrusted interface. The NAC Appliance acts as a default gateway for all clients residing on those networks. Conversely, a single VLAN/subnet is defined on the trusted interface, which represents the path to the protected upstream network(s).

After successful client authentication and posture assessment, the NAC Appliance by default routes traffic from the untrusted networks to the trusted interface, where it is then forwarded based on the routing topology of the network.

The NAC Appliance is not currently able to support dynamic routing protocols. As such, static routes must be configured within the trusted side of the Layer 3 network for each client subnet terminating on or residing behind the untrusted interface. These static routes should reference, as a next hop, the IP address of the trusted interface of the NAC.

If one or more Layer-3 hops exist between the untrusted NAC interface and the end-client subnets, static routes to the client networks must be configured in the NAC Appliance. Likewise, a static default route (0/0) is required within the downstream Layer 3 network (referencing the IP address of the untrusted NAC interface) to facilitate default routing behavior from the client networks to the NAC Appliance.

Depending on the topology, multiple options exist to facilitate routing to and from the NAC Appliance, including static routes, VRF-Lite, MPLS VPN, and other segmentation techniques. It is beyond the scope of this design guide to examine all possible methods.

In-Band Versus Out-of-Band

Table 5-1 summarizes different characteristics of each type of deployment.

Table 5-1 In-Band Versus Out-of-Band Deployment Characteristics

In-Band Deployment Characteristics	Out-of-Band Deployment Characteristics
The CAS is always inline with user traffic (both before and following authentication, posture assessment and remediation). Enforcement is achieved through being inline with traffic.	The CAS is inline with user traffic only during the process of authentication, assessment and remediation. Following that, user traffic does not come to the CAS. Enforcement is achieved through the use of SNMP to control switches and VLAN assignments to ports.
The CAS can be used to securely control authenticated and unauthenticated user traffic by using traffic policies (based on port, protocol, subnet), bandwidth policies, and so on.	The CAS can control user traffic during the authentication, assessment and remediation phase, but cannot do so post-remediation since the traffic is out-of-band.
Does not provide switch port level control.	Provides port-level control by assigning ports to specific VLANs as necessary using SNMP.
In-band deployment is supported for wired and wireless clients.	OOB deployments support wired and wireless clients. Wireless OOB requires a specific network topology.[1]
Cisco NAC Appliance In-Band deployment with supported Cisco switches is compatible with 802.1x	Cisco does not recommend using 802.1x in an OOB deployment, as conflicts will likely exist between Cisco NAC Appliance OOB and 802.1x to set the VLAN on the switch interfaces/ports.

1. OOB NAC deployments for wireless require the NAC server to be deployed in Layer 2 OOB virtual gateway (bridge) mode, and the NAC server must be placed Layer 2-adjacent to the wireless LAN controller (WLC).

Out-of-Band Requirements

OOB implementation of Cisco NAC Appliance requires the switches and Wireless LAN Controllers to be supported by the Cisco NAC Appliance software. All the switches tested as part of the development of the Small Enterprise Design Profile, apart from the Cisco Catalyst 2975, are supported by the Cisco NAC OOB, and the Wireless LAN Controllers are also supported by the NAC Appliance software used in this design guide. If the Catalyst 2975 is to be used as an access switch with the Cisco NAC Appliance, the NAC solution must be an in-band solution.

Note To obtain the latest list of supported devices, check the latest version of the *Cisco NAC Appliance-Clean Access Manager Installation and Administration Guide* at the following URL: http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/45/cam/45cam-book.html

Out-Of-Band, Layer 2 and Layer 3

The proposed design for the Small Enterprise Design Profile is an OOB design, in order to get the highest possible performance and scalability for traffic that has passed through the authentication, posture assessment, and remediation stages of NAC. The Small Enterprise Design Profile offers two different access layer options, a Layer-2 access layer for smaller sites and a hybrid Layer-2/Layer-3 access layer for larger sites. This means that either a Layer-2 OOB solution or a Layer-3 OOB NAC solution may be deployed.

NAC Appliance Deployment in Small Enterprise Networks

Within the Small Enterprise Design Profile, a Cisco NAC Appliance is deployed at each of the site types, the headquarters or main site, and remote sites. A centralized CAM is deployed at the main site, likely within the serverfarm. A CAS is deployed at the main site and each remote site and is directly connected to the core/distribution layer at each of the locations.

The simple topology used in each site type means that a VLAN from an access layer to the untrusted interface of the NAC Appliance is always available as a standard component of the design, and untrusted traffic should never need to be tunneled to the CAS. This allows a common network configuration, to support NAC at any of the enterprise sites, regardless of whether the client devices are using a Layer-2 or Layer-3 access model. As the client can use a Layer-2 connection to the untrusted interface of the NAS in either Layer 2 or Layer 3 access mode (this requires a trunk between the Layer-3 access switch and the core/distribution. One VLAN of the trunk would carry the untrusted VLAN, and the other VLAN the IP routing for all other traffic), and the VLAN used once the client is trusted will be either be a Layer-2 access VLAN from the core/distribution switch or a Layer-3 access switch VLAN depending upon the site requirements.

This is illustrated in Figure 5-18 and Figure 5-19. In Figure 5-19, there is a simple Layer-2 NAC OOB connection where a client device upon initial connection to the network is given VLAN 264, which connects them directly to the untrusted interface of the NAS. VLAN 264 on the untrusted interface is mapped to VLAN 64 on the trusted interface within the NAC appliance, which allows the client to obtain an IP address that belongs on VLAN 64. Upon successful completion of the NAC authentication and validation functions, the access switch is instructed, via SNMP from the CAM, to change the client VLAN to VLAN 64. Even though the client has changed Layer-2 VLANs, its Layer-3 network connections are unchanged and the traffic from the client no longer passes through the NAC appliance.

Figure 5-18 Layer 2 OOB Topology

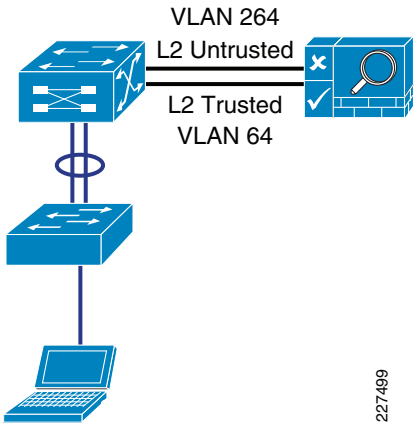

In Figure 5-19, the same processes are followed when the client is untrusted, but once the client has successfully completed its NAC functions the access switch is instructed via SNMP to change the client VLAN to VLAN 67—a subnet local to the access switch. As the Layer-3 information for the client has changed, the switch is also instructed to “bounce” the client switch port to initiate a new DHCP request for an IP address appropriate to VLAN 67.

Figure 5-19 Layer 3 OOB Topology

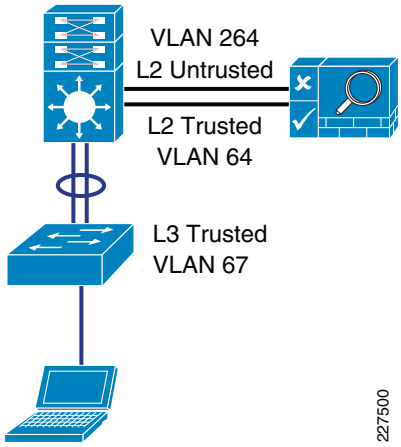

Cisco Identity-Based Network Networking Services (IBNS)

The best and most secure solution to vulnerability at the access edge is to leverage the intelligence of the network. The Cisco Identity-Based Network Networking Services solution (IBNS) is a set of Cisco IOS software services designed to enable secure user and host access to enterprise networks powered by Cisco Catalyst switches and wireless LANs. It provides standards-based network access control at the access layer by using the 802.1X protocol to secure the physical ports where end users connect. 802.1X is an IEEE standard for media-level (Layer 2) access control, offering the capability to permit or deny network connectivity based on the identity of the end user or device. 802.1X is well-known as a way to secure wireless network access. It is equally essential in securing wired network access.

IEEE 802.1X Protocol

The IEEE 802.1X protocol allows Cisco Catalyst switches to offer network access control at the port level. Every port on the switch is individually enabled or disabled based on the identity of the user or device connecting to it. When 802.1X is first enabled on a port, the switch automatically drops all traffic received on that port. There is one exception to this rule. The only traffic a switch will accept is a request to start 802.1X authentication. Only after the 802.1X authentication has successfully completed will the switch accept any other kind of traffic on the port.

The high-level message exchange in Figure 5-20 illustrates how port-based access control works within an identity-based system.

Figure 5-20 Port-Based Access Control

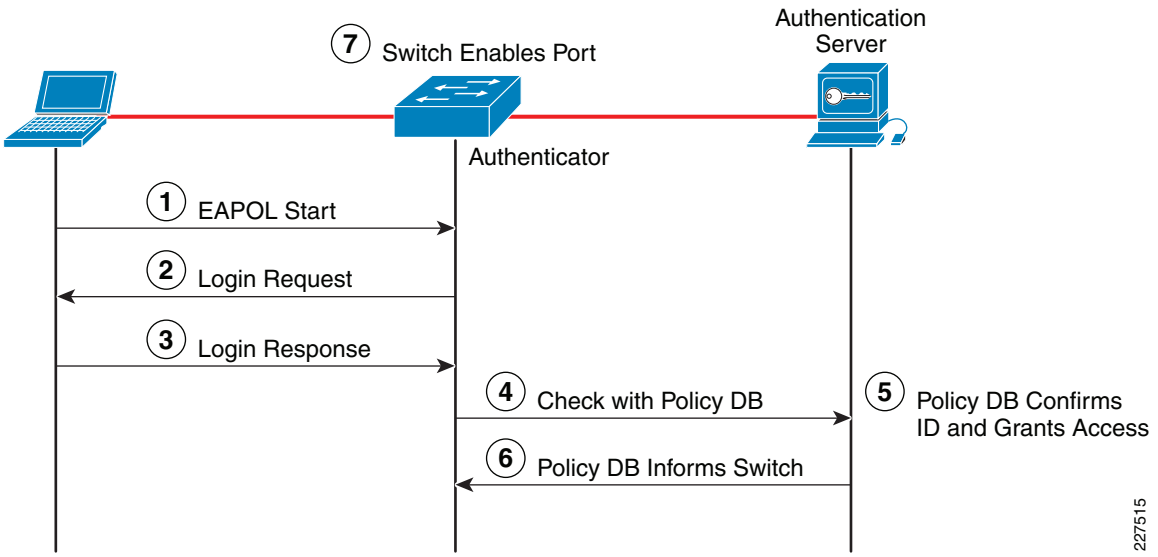

The following steps describe the port-based access control flow shown in Figure 5-20:

Step 1 A client, such as a laptop with an 802.1X supplicant, connects to an IEEE 802.1X-enabled network and sends a start message to the LAN switch (the authenticator).

Step 2 When the start message is received, the LAN switch sends a login request to the client.

Step 3 The client replies with a login response.

Step 4 The switch forwards the response to the policy database (authentication server).

Step 5 The authentication server authenticates the user.

Step 6 After the user identity is confirmed, the policy database authorizes network access for the user and informs the LAN switch.

Step 7 The LAN switch then enables the port connected to the client.

The user and device credentials are processed by an AAA server. The AAA server is able to reference user or device policy profile information either internally, using the integrated user database, or externally, using database sources such as Microsoft Active Directory, LDAP, Novell NDS or Oracle databases. This enables the integration of the system into exiting user management structures and schemes, thereby simplifying overall deployment.

802.1X and EAP

When authenticating users for the purposes of network access control, the system must provide user and/or device identification using strong authentication technologies known to be secure and reliable. IEEE 802.1X does not by itself dictate how this is achieved. Rather, the 802.1X protocol defines an encapsulation for the transport of the Extensible Authentication Protocol (EAP) from the client to the switch. The 802.1X encapsulation is sometimes referred to as EAP over LAN (EAPoL). The switch in turn relays the EAP information to the authentication server using the RADIUS protocol (EAP over RADIUS).

EAP, which is defined by RFC 3748, is itself a framework---not a specific authentication method. EAP provides a way for the client and the authentication server to negotiate an authentication method that they both support. There are many EAP methods but the ones used more frequently for 802.1X wired authentication include EAP-TLS, EAP-PEAP, and EAP-FAST.

Impacts of 802.1X on the Network

Before enabling 802.1X in the network, it is essential to review the default security posture of a port enabled for 802.1X authentication: all traffic is dropped except 802.1X EAPoL packets. This is a fundamental change from the traditional model in which the port is enabled and all traffic is allowed from the moment that a device plugs into the port. Ports that were traditionally open will now be closed by default. This is one of the cornerstones of the strong security and network access control provided by 802.1X. However, this change in the default network access model can have a profound impact on network devices and applications. Understanding and providing for the impacts of this change will make for a smooth deployment of 802.1X network access control.

Non-802.1X-Enabled Devices

802.1X must be enabled on both the host device and on the switch to which the device connects. If a device without an 802.1X supplicant attempts to connect to a port that is enabled for 802.1X, it will be subjected to the default security policy. The default security policy says that 802.1X authentication must succeed before access to the network is granted. Therefore, by default, non-802.1X-capable devices cannot get access to an 802.1X-protected network.

Although many devices increasingly support 802.1X, there will always be devices that require network connectivity but do not and/or cannot support 802.1X. Examples of such devices include network printers, badge readers, legacy servers, and PXE boot machines. Some provision must be made for these devices.

Cisco provides two features to accommodate non-802.1X devices. These are MAC Authentication Bypass (MAB) and the Guest VLAN. These features provide fallback mechanisms when there is no 802.1X supplicant. After 802.1X times out on a port, the port can move to an open state if MAB succeeds or if the Guest VLAN is configured. Judicious application of either or both of these features will be required for a successful 802.1X deployment.

Note Network-specific testing will be required to determine the optimal values for 802.1X timers to accommodate the various non-802.1X-capable devices on your network.

802.1X in Small Enterprise Networks

As mentioned above, one of the requirements for 802.1X authentication is the requirement for a supplicant. This has typically been a challenge in enterprise environments with a wide range of devices and limited or no management of many of these devices. In many enterprises this is still the case, and this makes a company-wide 802.1X very challenging. At the same time there are pockets of an enterprise network where 802.1X may be a good choice.

For example, 802.1X protected ports may be a good choice for the network ports in the company's headquarters or main site, because these locations are more likely to have managed PCs.

Other locations in the enterprise network still need protection, but user network access may be better served by a NAC Appliance solution. For networks requiring role-based access control using posture assessments to ensure security compliance, Cisco NAC Appliance should be considered. In addition, network access ports in open areas such as lobbies and meeting rooms may use 802.1X or Cisco Clean Access NAC to protect these ports.

When considering the 802.1X deployment, the following four main 802.1X authentication options need to be considered:

- *Basic 802.1X Authentication*—An 802.1X controlled port with an 802.1X client directly connected
- *IP Phone Ports*—An IP Phone and an 802.1X controlled port with an 802.1X client connected to the phone
- *MAC Auth By-Pass*—Using the MAC address of the client to provide authentication and bypass the 802.1X authentication process. Printer and legacy device support are typical applications
- *Web Auth*—Allowing a user to authenticate by entering username and passwords in a web page. Legacy device support and guest access are typical deployment applications

For more information on the Cisco IBNS 802.1X network access solution, see the following URL: http://www.cisco.com/go/ibns.

NAC 802.1X and CISF in Combination

The three key access security features discussed above have been discussed in isolation, but can be combined. In particular, the CISF features should be considered “baseline” features that are applied on all access ports, and either NAC or 802.1X maybe overlaid on top of the CISF configuration.

The Cisco Clean Access and 802.1X configuration are also compatible (although they are not often combined in wired networks), the key consideration in combining the two is how to give the appearance of a SSO for the end user. Both 802.1X and NAC require authentication, as 802.1X authenticates the client initially, a mechanism of communicating the 802.1X authentication result to the Cisco Clean Access system is required.

If the authenticating clients join an Windows Active Directory network, the Cisco Clean Access Active Directory SSO feature allows the clients to authenticate to active directory once they have performed their 802.1X authentication. The CAM, when a client is detected, checks Active Directory to see if the client has authenticated; this allows a SSO experience for client devices that are using 802.1X and NAC.

Secure Mobility

Today workers use laptops, smartphones and other smart mobile devices to access information and applications at anytime and from anywhere there is an Internet connection. While embracing a mobile workforce clearly boosts productivity and makes the small enterprise more competitive, there are a number of challenges that arise from the use of mobile technologies. To start with, workers often use the same devices to access both business and personal information. Devices used outside the enterprise onsite controls may potentially introduce viruses, worms, spyware and other type of malware as mobile workers connect back to the corporate network. Confidential and proprietary information may also be lost or stolen while mobile users connect outside the company premises. At the other hand, the great variety in hardware types, operating systems, and applications represents a clear challenge to the enforcement of security controls and policies. In order to continue to foster innovation, enable productivity, and meet the needs of the mobile workforce, companies must adapt to the changing trends in mobility. A viable solution is one that enables access for mobile workers while ensuring that the

corporate data, assets and network remain secure. At the other hand, users want the flexibility of choosing how, when, and where to access both personal and professional information to be productive without being inconvenienced by security checks.

The Small Enterprise Design Profile provides persistent and secure mobile access by implementing the Cisco AnyConnect Secure Mobility solution (see Figure 5-21). This solution delivers secure, persistent connectivity to all mobile employees independently from the type of mobile device used. The Cisco AnyConnect Secure Mobility solution also ensures a consistent enforcement of the network security policies to all users, no matter when, where and how they connect to the network.

The Cisco AnyConnect Secure Mobility is a collection of features across multiple Cisco products that extends control and security into borderless networks. The products that work together to provide AnyConnect Secure Mobility are as follows:

- Cisco IronPort Web Security appliance (WSA)
- Cisco ASA 5500 series adaptive security appliance (ASA)
- Cisco AnyConnect client

Cisco AnyConnect Secure Mobility addresses the challenges of a mobile workforce by offering the following features:

- *Secure, persistent connectivity*—Cisco AnyConnect (with the Cisco ASA at the headend) provides the remote access connectivity portion of AnyConnect Secure Mobility. The connection is secure because both the user and device must be authenticated and validated prior to being provided access to the network. The connection is persistent because Cisco AnyConnect is typically configured to be always-on even when roaming between networks. Although Cisco AnyConnect is always-on, it is also flexible enough to apply different policies based on location, allowing users access to the Internet in a “captive portal” situation, when users must accept terms of agreement before accessing the Internet.
- *Persistent security and policy enforcement*—The Web Security appliance applies context-aware policies, including enforcing acceptable use policies and protection from malware for all users, including mobile (remote) users. The WSA also accepts user authentication information from the AnyConnect client, providing an automatic authentication step for the user to access web content.

Figure 5-21 Cisco AnyConnect Secure Mobility Solution

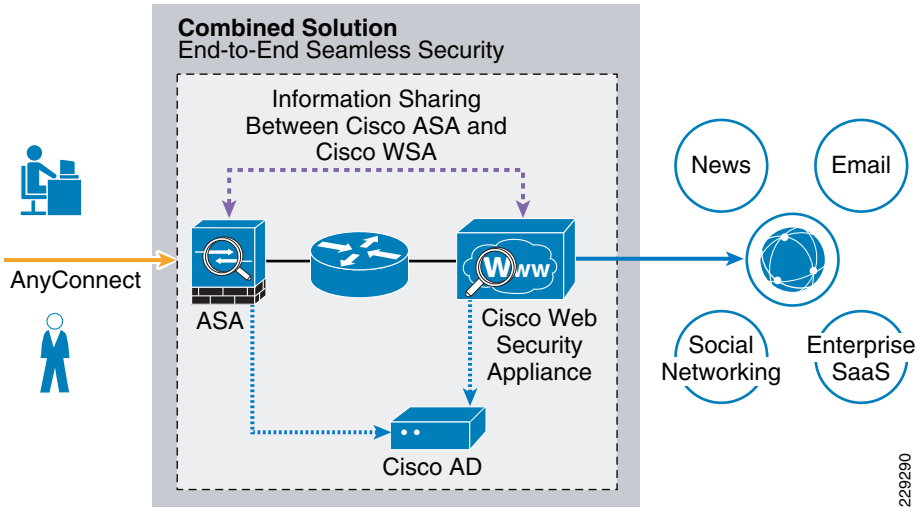

Figure 5-21 illustrates the relationship between the various elements of the Cisco AnyConnect Secure Mobility solution. Remote and mobile users use the Cisco AnyConnect Secure VPN client to establish VPN sessions with the Cisco ASA appliance. The Cisco ASA sends web traffic to the WSA appliance along with information identifying the user by IP address and user name. The WSA scans the traffic, enforces acceptable use policies, and protects the user from security threats. The Cisco ASA returns all traffic deemed safe and acceptable to the user.

All Internet traffic scanning is done by the WSA, not the client on the mobile device. This improves overall performance by not burdening the mobile device, some of which have limited processing power. In addition, by scanning Internet traffic on the network, the enterprise can more easily and quickly update security updates and acceptable use policies since the enterprise does not have to wait days, weeks, or even months to push the updates to the client. The WSA tracks the requests it receives and applies policies configured for remote users to traffic received from remote users.

For complete details about the Cisco AnyConnect Secure Mobility solution, refer to the documentation available at the following URL:

http://www.cisco.com/en/US/netsol/ns1049/index.html

Threats Mitigated

The success of the security tools and measures in place ultimately depends on the degree they enhance visibility and control. Simply put, security can be defined as a function of visibility and control. Without any visibility, it is difficult to enforce any control, and without any control it is hard to achieve an adequate level of security. Therefore, the security tools selected in the enterprise network design were carefully chosen not only to mitigate certain threats but also to increase the overall visibility and control.

Table 5-2 summarizes how the security tools and measures used in the Small Enterprise Design Profile help mitigate certain threats, and how they contribute to increasing visibility and control. Please note the table is provided for illustration purposes and it is not intended to include all possible security tools, and threats.

Table 5-2 Security Measures of the Small Enterprise Design Profile

	Service Disruption	Harmful Content	Network Abuse	Unauthorized Access	Data Loss	Visibility	Control
Network Foundation Protection	Yes			Yes	Yes	Yes	Yes
Stateful Firewall	Yes		Yes	Yes		Yes	Yes
IPS	Yes	Yes	Yes	Yes		Yes	Yes
Mobile Security	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Web Security		Yes	Yes	Yes	Yes	Yes	Yes
E-mail Security		Yes	Yes		Yes	Yes	Yes
Access Security and Control			Yes	Yes		Yes	Yes

Network Security Deployment

This section describes the deployment best practices for the key security platforms and features used in the Small Enterprise Design Profile. This includes deployment and setting guidelines, and configuration examples.

Internet Border Router Deployment

Whether the Internet border router is managed by the enterprise or the ISP, it must be hardened following the best practices listed in the "Network Foundation Protection" section on page 5-3. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information. In addition, the Internet border router may be leveraged as the first layer of protection against outside threats. To that end, edge ACLs, uRPF, and other filtering mechanisms may be implemented for anti-spoofing and to block invalid packets.

The following configuration snippet illustrates the structure of an edge ACL applied to the upstream interface of the Internet border router. The ACL is designed to block invalid packets and to protect the infrastructure IP addresses from the Internet. The configuration assumes the enterprise is assigned the 198.133.219.0/24 address block for its public-facing services, and that the upstream link is configured in the 64.104.10.0/24 subnet.

```
! !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 1: Anti-spoofing Denies
!--- These ACEs deny fragments, RFC 1918 space,
!--- invalid source addresses, and spoofs of
!--- internal space (space as an external source).
!
!--- Deny fragments.
access-list 110 deny tcp any 198.133.219.0 0.0.0.255 fragments
access-list 110 deny udp any 198.133.219.0 0.0.0.255 fragments
access-list 110 deny icmp any 198.133.219.0 0.0.0.255 fragments
!--- Deny special-use address sources.
!--- See RFC 3330 for additional special-use addresses.
access-list 110 deny ip host 0.0.0.0 any
access-list 110 deny ip 127.0.0.0 0.255.255.255 any
access-list 110 deny ip 192.0.2.0 0.0.0.255 any
access-list 110 deny ip 224.0.0.0 31.255.255.255 any
!--- Filter RFC 1918 space.
access-list 110 deny ip 10.0.0.0 0.255.255.255 any
access-list 110 deny ip 172.16.0.0 0.15.255.255 any
access-list 110 deny ip 192.168.0.0 0.0.255.255 any
!--- Deny packets spoofing the enterprise public addresses
access-list 110 deny ip 198.133.219.0 0.0.0.255 any
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 2:  Explicit Permit
!--- Permit only applications/protocols whose destination
!--- address is part of the infrastructure IP block.
!--- The source of the traffic should be known and authorized.
!
!--- Permit external BGP to peer 64.104.10.113
access-list 110 permit tcp host 64.104.10.114 host 64.104.10.113 eq bgp
access-list 110 permit tcp host 64.104.10.114 eq bgp host 64.104.10.113
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 3:  Explicit Deny to Protect Infrastructure
access-list 110 deny ip 64.104.10.0 0.0.0.255 any
!
```

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!--- Module 4:  Explicit Permit for Traffic to Company's Public
!--- Subnet.
access-list 110 permit ip any 198.133.219.0 0.0.0.255
!
```

Note The 64.104.0.0/16 and 198.133.219.0/24 address blocks used in the examples in this chapter are reserved for the exclusive use of Cisco Systems, Inc.

Internet Firewall Deployment

The mission of the Internet firewall is to protect the company's internal resources and data from external threats, secure the public services provided by the DMZ, and to control user's traffic to the Internet. The Small Enterprise Design Profile uses a Cisco ASA appliance as illustrated in Figure 5-22.

Figure 5-22 Internet Edge Firewall

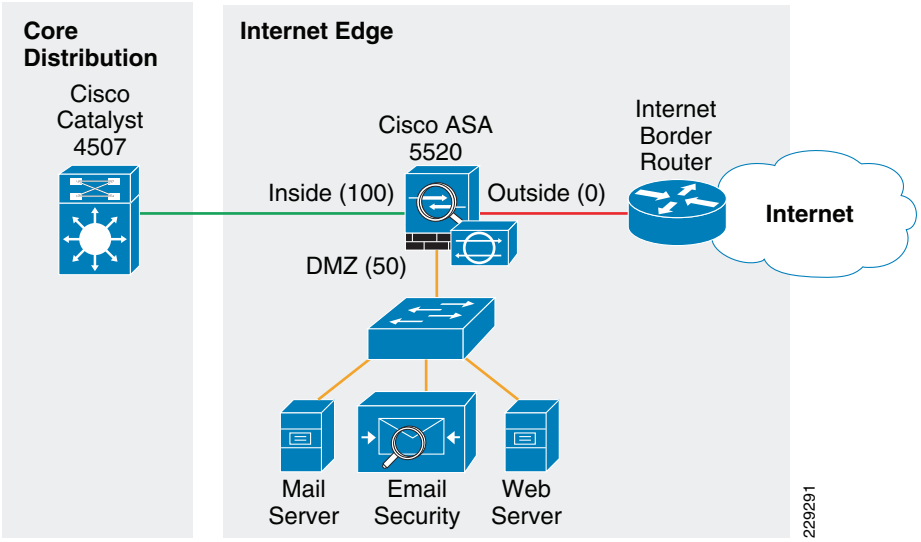

The Cisco ASA is implemented with three interface groups, each one representing a distinct security domain:

- *Inside*—The interface connecting to the core/distribution switch that faces the interior of the network where internal users and resources reside. The inside interface connects to the internal trusted networks, therefore it is given the highest security level, 100.
- *Outside*—Interface connecting to the Internet border router. The router may be managed either by the enterprise or a service provider. The outside interface connects to the Internet, hence it's given the lowest security level, 0.
- *Demilitarized Zone (DMZ)*—The DMZ hosts services that are accessible over the Internet. These services may include the company's website and E-mail services. The DMZ serves a medium level security segment, therefore should be given any security value between the ones defined for the inside and the outside interfaces, for example, 50.

The Internet firewall acts as the primary gateway to the Internet; therefore, its deployment should be carefully planned. The following are key aspects to be considered when implementing the firewall:

- Firewall hardening and monitoring
- Network Address Translation (NAT)
- Firewall access policies
- Botnet Traffic Filter
- Firewall redundancy
- Routing

Firewall Hardening and Monitoring

The Cisco ASA should be hardened in a similar fashion as the infrastructure routers and switches. According to the Cisco SAFE security best practices, the following is a summary of the measures to be taken:

- Implement dedicated management interfaces to the OOB management network.
- Present legal notification for all access attempts.
- Use HTTPS and SSH for device access. Limit access to known IP addresses used for administrative access.
- Configure AAA for role-based access control and logging. Use a local fallback account in case AAA server is unreachable.
- Use NTP to synchronize the time.
- Use syslog or SNMP to keep track of system status, traffic statistics, and device access information.
- Authenticate routing neighbors and log neighbor changes.
- Implement firewall access policies (explained in Firewall Access Policies).

The Cisco ASA 5510 and higher appliance models come with a dedicated management interface that should be used whenever possible. Using a dedicated management interface keeps the management plane of the firewall isolated from threats originating from the data plane. The management interface should connect to the OOB management network, if one is available.

The following is an example of the configuration of a dedicated management interface.

```
interface Management0/0
 nameif management
 security-level 100
 ip address 172.26.160.225 255.255.252.0
 management-only
!
```

Note Any physical interface or logical sub-interface can be configured as a management-only interface using the **management-only** command.

It is recommended that a legal notification banner is presented on all interactive sessions to ensure that users are notified of the security policy being enforced and to which they are subject. The notification banner should be written in consultation with your legal advisors.

The following example displays the banner after the user logs in:

```
banner motd UNAUTHORIZED ACCESS TO THIS DEVICE IS PROHIBITED.
banner motd You must have explicit, authorized permission to access or configure this
device.
banner motd Unauthorized attempts and actions to access or use this system may result in
civil and/or criminal penalties.
banner motd All activities performed on this device are logged and monitored.
```

Management access to the firewall should be restricted to SSH and HTTPS. SSH is needed for CLI access and HTTPS is needed for the firewall GUI-based management tools such as CSM and ADSM. Additionally, this access should only be permitted for users authorized to access the firewalls for management purposes.

The following ASA configuration fragment illustrates the configuration needed to generate a 768 RSA key pair and enabling SSH and HTTPS access for devices located in the management subnet.

```
! Generate RSA key pair with a key modulus of 768 bits
crypto key generate rsa modulus 768
! Save the RSA keys to persistent flash memory
write memory
! enable HTTPS
http server enable
! restrict HTTPS access to the firewall to permitted management stations
http <CSM/ADSM-IP-address> 255.255.255.255 management
! restrict SSH access to the firewall to well-known administrative systems
ssh <admin-host-IP-address> 255.255.255.255 management
! Configure a timeout value for SSH access to 5 minutes
ssh timeout 5
```

Administrative users accessing the firewalls for management must be authenticated, authorized, and access should be logged using AAA. The following ASA configuration fragment illustrates the AAA configurations needed to authenticate, authorize, and log user access to the firewall:

```
aaa-server tacacs-servers protocol tacacs+
 reactivation-mode timed
aaa-server tacacs-servers host <ACS-Server>
 key <secure-key>
aaa authentication ssh console tacacs-servers LOCAL
aaa authentication serial console tacacs-servers LOCAL
aaa authentication enable console tacacs-servers LOCAL
aaa authentication http console tacacs-servers LOCAL
aaa authorization command tacacs-servers LOCAL
aaa accounting ssh console tacacs-servers
aaa accounting serial console tacacs-servers
aaa accounting command tacacs-servers
aaa accounting enable console tacacs-servers
aaa authorization exec authentication-server
! define local username and password for local authentication fallback
username admin password <secure-password> encrypted privilege 15
```

As with the other infrastructure devices in the network, it is important to synchronize the time on the firewall protecting the management module using NTP.

The following configuration fragment illustrates the NTP configuration needed on an ASA to enable NTP to an NTP server located in the management network:

```
ntp authentication-key 10 md5 *
ntp authenticate
ntp trusted-key 10
```

```
ntp server <NTP-Server-address> source management
```

Syslog and SNMP can be used to keep track of system status, device access and session activity. NetFlow Security Event Logging (NSEL), now supported on all Cisco ASA models, may also be used for the monitoring and reporting of session activity. The following configuration fragment illustrates the configuration of Syslog.

```
logging trap informational
logging host management <Syslog-Server-address>
logging enable
```

The routing protocol running between the Internet firewall and the distribution/core should be secured. The following ASA configuration fragment illustrates the use of EIGRP MD5 authentication to authenticate the peering session between the inside firewall interface and the core/distribution switch:

```
interface Redundant1
 nameif inside
 security-level 100
 ip address 10.125.33.10 255.255.255.0
 authentication key eigrp 100 <removed> key-id 1
 authentication mode eigrp 100 md5
```

Network Address Translation (NAT)

NAT is required because the enterprise typically gets a limited number of public IP addresses. In addition, NAT helps shield the company's internal address space from reconnaissance and another malicious activity.

The following illustrates the NAT configuration:

```
! Static translation for servers residing at DMZ
static (dmz,outside) 198.133.219.10 10.25.34.10 netmask 255.255.255.255
static (dmz,outside) 198.133.219.11 10.25.34.11 netmask 255.255.255.255
static (dmz,outside) 198.133.219.12 10.25.34.12 netmask 255.255.255.255
static (dmz,outside) 198.133.219.13 10.25.34.13 netmask 255.255.255.255
!
! Dynamic Port Address Translation (PAT) for inside hosts going to the Internet
global (outside) 10 interface
nat (inside) 10 10.0.0.0 255.0.0.0
!
! Static translation for inside hosts going to the DMZ and vice-versa. The inside IP
addresses are visible to the DMZ.
static (inside,dmz) 10.0.0.0 10.0.0.0 netmask 255.0.0.0
```

Firewall Access Policies

As previously explained, the Internet firewall should be configured to:

- Protect the company's internal resources and data from external threats by preventing incoming access from the Internet.
- Protect public resources served by the DMZ by restricting incoming access to the public services and by limiting outbound access from DMZ resources out to the Internet.
- Control user's Internet-bound traffic.

Enforcing such policies requires the configuration of the appropriate interface security levels and the deployment of ACLs governing what traffic is allowed or prevented from transiting between interfaces.

By default, the Cisco ASA appliance allows traffic from higher to lower security level interfaces (i.e., from inside to outside). However, due to the sensitivity of enterprise environments, the security administrators are recommended to override the default rules with more stringent rules indicating exactly what ports and protocols are permitted.

In our configuration example the inside, DMZ and outside interfaces were configured with the security levels of 100, 50 and 0 respectively. With this, by default any traffic originating from the inside to the DMZ and outside, and from the DMZ to the outside interface will be allowed freely. At the same time, any traffic originating from the outside to the DMZ and inside, and from the DMZ to the inside interface will be blocked. While this may satisfy the basic access control requirements of the organization, it is always a good idea to reinforce the policies by enforcing granular ACLs.

Before designing the ACLs, it should also be noted that, as the Cisco ASA inspects traffic it is able to recognize packets belonging to already established sessions. The stateful inspection engine of the firewall dynamically allows the returning traffic. Therefore, the firewall ACLs should be constructed to match traffic in the direction in which it is being initiated. In our sample configurations ACLs are applied in the ingress direction.

The following are the guidelines and configuration examples of ACLs controlling access and traffic flows:

- Ingress Inside

  Allow Internet access to users residing at all enterprise sites for the allowed ports and protocols. This typically includes HTTP and HTTPS access.

```
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 any eq http
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 any eq https
```

Allow users access to DMZ services such as the company's website, E-mail, and domain name resolution (HTTP, HTTPS, SMTP, POP, IMAP, and DNS). Note that the previous entries in the ACL already permit HTTP and HTTPS traffic.

```
! Allow DNS queries to DNS server
access-list Outbound extended permit udp 10.0.0.0 255.0.0.0 host 10.25.34.13 eq domain
! Allow SMTP, POP3 and IMAP access to DMZ mail server
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 host 10.25.34.12 eq smtp
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 host 10.25.34.12 eq pop3
access-list Outbound extended permit tcp 10.0.0.0 255.0.0.0 host 10.25.34.12 eq imap4
! Apply ACL to inside interface
access-group Outbound in interface inside
```

- Ingress DMZ

  Restrict connections initiated from DMZ only to the necessary protocols and sources. This typically includes DNS queries and zone transfer from DNS server, SMTP from E-mail server, HTTP/SSL access from the Cisco IronPort ESA for updates, Sensorbase, etc.

```
! Allow DNS queries and zone transfer from DNS server
access-list DMZ extended permit udp host 10.25.34.13 any eq domain
access-list DMZ extended permit tcp host 10.25.34.13 any eq domain
!
! Allow SMTP from Cisco IronPort ESA
access-list DMZ extended permit tcp host 10.25.34.11 any eq smtp
!
! Allow update and SensorBase access to Cisco IronPort ESA
access-list DMZ extended permit tcp host 10.25.34.11 any eq http
access-list DMZ extended permit tcp host 10.25.34.11 any eq https
!
! Apply ACL to DMZ interface
access-group DMZ in interface dmz
```

- Ingress Outside

  Inbound Internet access should be restricted to the public services provided at the DMZ such as SMTP, web, and DNS. Any connection attempts to internal resources and subnets from the Internet should be blocked. ACLs should be constructed using the servers' global IP addresses.

```
! Allow DNS queries and zone transfer to DNS server
access-list Inbound extended permit udp any host 198.133.219.13 eq domain
access-list Inbound extended permit tcp any host 198.133.219.13 eq domain
!
! Allow SMTP to Cisco IronPort ESA
access-list Inbound extended permit tcp any host 198.133.219.11 eq smtp
!
! Allow HTTP/HTTPS access to the company's public web portal
access-list Inbound extended permit tcp any host 198.133.219.10 eq http
access-list Inbound extended permit tcp any host 198.133.219.10 eq https
!
! Apply ACL to outside interface
access-group Inbound in interface outside
```

Botnet Traffic Filter

The Small Enterprise Design Profile uses the Cisco ASA Botnet Traffic Filter on the Internet Firewall to detect malware that attempts network activity such as sending private data (passwords, credit card numbers, key strokes, or other proprietary data) when the malware starts a connection to a known bad IP address. The Botnet Traffic Filter checks incoming and outgoing connections against a dynamic database of known bad domain names and IP addresses (the blacklist), and then logs or blocks any suspicious activity.

Configuring the Botnet Traffic Filter requires the following steps:

Step 1 Configure DNS Server

Step 2 Enable Use of the Dynamic Database

Step 3 Enable DNS Snooping

Step 4 Enable Traffic Classification and Actions for the Botnet Traffic Filter

Step 5 Verify and Monitor Botnet Traffic Filter Operation

The following sections provides configuration examples for each of these steps.

Configure DNS Server

The Botnet Traffic Filter requires a DNS Server to access Cisco's dynamic database update server and to resolve entries in the static database. The following configuration illustrates this configuration.

```
! Enable DNS requests to a DNS Server out the outside interface
dns domain-lookup outside
! Specify the DNS Server Group and the DNS Servers
dns server-group DefaultDNS
 name-server 68.238.112.12
 name-server 68.238.96.12
 domain-name cisco.com
```

Enable Use of the Dynamic Database

The Botnet Traffic Filter can receive periodic updates for the dynamic database from the Cisco update server. This database lists thousands of known bad domain names and IP addresses. The following configuration enables database updates, and also enables use of the downloaded dynamic database by the adaptive security appliance.

```
! enable downloading of the dynamic database from the Cisco Update server
dynamic-filter updater-client enable
! enable use of the dynamic database
dynamic-filter use-database
```

Enable DNS Snooping

DNS Snooping enables inspection of DNS packets and enables Botnet Traffic Filter Snooping which compares the domain name with those in the dynamic or static databases, and adds the name and IP address to the DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

It is recommended that DNS snooping is only enabled on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the Cisco ASA. For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface.

The following configuration example illustrates enabling DNS Snooping on the outside interface:

```
! create a class map to identify the traffic you want to inspect DNS
class-map dynamic-filter-snoop-class
 match port udp eq domain
! create a policy map to enable DNS inspection with Botnet Traffic Filtering snooping for
the class map
policy-map dynamic-filter-snoop-policy
 class dynamic-filter-snoop-class
  inspect dns preset_dns_map dynamic-filter-snoop
! activate the policy map on the outside interface
service-policy dynamic-filter-snoop-policy interface outside
```

Enable Traffic Classification and Actions for the Botnet Traffic Filter

The Botnet Traffic Filter compares the source and destination IP address in each initial connection packet to the IP addresses in the dynamic database, static database, DNS reverse lookup cache, and DNS host cache, and sends a syslog message or drops any matching traffic. When an address matches, the Cisco ASA sends a syslog message and can optionally be configured to drop the connection. You can enable Botnet Traffic filter on a subset of traffic or for all traffic by enabling an access list to classify traffic.

The following configuration example enables the Botnet Traffic Filter feature on all traffic and additionally enables dropping of connections going to IP addresses with a severity of moderate and higher.

```
! identify the traffic that you want to monitor or drop.
access-list btf-filter-acl extended permit ip any any
! enable Botnet Traffic Filter on the outside interface for traffic classified by the
btf-filter-acl access list
dynamic-filter enable interface outside classify-list btf-filter-acl
! enable automatic dropping of traffic with threat level moderate or higher
dynamic-filter drop blacklist interface outside action-classify-list btf-filter-acl
threat-level range moderate very-high
```

Botnet Traffic Filter Verification

To monitor and verify the operation of the Botnet Traffic Filter feature, the following commands can be used:

- **show dynamic-filter updater-client**—Shows information about the updater server, including the server IP address, the next time the adaptive security appliance will connect with the server, and the database version last installed.

```
cr12-asa-1-ie# show dynamic-filter updater-client
Dynamic Filter updater client is enabled
Updater server URL is https://update-manifests.ironport.com
Application name: threatcast, version: 1.0
Encrypted UDI:
0bb93985f42d941e50dc8f022350d1a8a8c5097dc1d252b9cff62d26d4ec58e202883d704fc62b85bf8629
fa757fe36b
Last update attempted at 15:14:11 UTC Apr 7 2010,
  with result: Downloaded file successfully
Next update is in 00:52:14
Database file version is '1270651144' fetched at 15:14:11 UTC Apr 7 2010, size:
2097152
cr12-asa-1-ie#
```

- **show dynamic-filter data**—Shows information about the updater server, including the server IP address, the next time the adaptive security appliance will connect with the server, and the database version last installed.

```
cr12-asa-1-ie# show dynamic-filter data
Dynamic Filter is using downloaded database version '1270651144'
Fetched at 15:14:11 UTC Apr 7 2010, size: 2097152
Sample contents from downloaded database:
  win-antimalware2.com  firstlook.com  red-devil-sport-club.gymdb.com
mswindowsupdate.info
  zardoz.wizardz.com  exchange.bg  bisexual-photo.com  lookmbox.com
Sample meta data from downloaded database:
  threat-level: very-high,      category: Malware,
  description: "These are sources that use various exploits to deliver adware, spyware
and other malware to victim computers.  Some of these are associated with rogue online
vendors and distributors of dialers which deceptively call premium-rate phone
numbers."
  threat-level: high,   category: Bot and Threat Networks,
  description: "These are rogue systems that control infected computers.  They are
either systems hosted on threat networks or systems that are part of the botnet
itself."
  threat-level: moderate,       category: Spyware,
  description: "These are sources that distribute spyware, adware, greyware, and other
potentially unwanted advertising software.  Some of these also run exploits to install
such software."
  threat-level: low,    category: Ads,
  description: "These are advertising networks that deliver banner ads, interstitials,
rich media ads, pop-ups, and pop-unders for websites, spyware and adware.  Some of
these networks send ad-oriented HTML emails and email verification services."
Total entries in Dynamic Filter database:
  Dynamic data: 82119 domain names , 2565 IPv4 addresses
  Local data: 0 domain names , 0 IPv4 addresses
Active rules in Dynamic Filter asp table:
  Dynamic data: 0 domain names , 2565 IPv4 addresses
  Local data: 0 domain names , 0 IPv4 addresses
cr12-asa-1-ie#
```

- **show dynamic-filter statistics detail**—Shows how many connections were monitored and dropped with the Botnet Traffic Filter, and how many of those connections match the whitelist, blacklist, and greylist. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.) The detail keyword shows how many packets at each threat level were classified or dropped.

```
cr12-asa-1-ie# show dynamic-filter statistics detail
Enabled on interface outside using classify-list btf-filter-acl
 Total conns classified 35, ingress 0, egress 35
 Total whitelist classified 0, ingress 0, egress 0
 Total greylist classified 16, dropped 0, ingress 0, egress 16
  Threat-level  very-high: classified        0, dropped        0, ingress        0,
egress        0
  Threat-level       high: classified        0, dropped        0, ingress        0,
egress        0
  Threat-level   moderate: classified        0, dropped        0, ingress        0,
egress        0
  Threat-level        low: classified       16, dropped        0, ingress        0,
egress       16
  Threat-level   very-low: classified        0, dropped        0, ingress        0,
egress        0
 Total blacklist classified 19, dropped 0, ingress 0, egress 19
  Threat-level  very-high: classified        9, dropped        0, ingress        0,
egress        9
  Threat-level       high: classified        0, dropped        0, ingress        0,
egress        0
  Threat-level   moderate: classified        0, dropped        0, ingress        0,
egress        0
  Threat-level        low: classified       10, dropped        0, ingress        0,
egress       10
  Threat-level   very-low: classified        0, dropped        0, ingress        0,
egress        0
cr12-asa-1-ie#
```

Note To clear the statistics, enter the clear dynamic-filter statistics [interface name] command.

Other commands that are useful for monitoring the Botnet Traffic filter include the following:

- **show dynamic-filter reports top [malware-sites | malware-ports | infected-hosts]**—Generates reports of the top 10 malware sites, ports, and infected hosts monitored. The top 10 malware-sites report includes the number of connections dropped, and the threat level and category of each site. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected.
- **show dynamic-filter reports infected-hosts {max-connections | latest-active | highest-threat | subnet ip_address netmask | all}**—Generates reports about infected hosts. These reports contain detailed history about infected hosts, showing the correlation between infected hosts, visited malware sites, and malware ports. The max-connections keyword shows the 20 infected hosts with the most number of connections. The latest-active keyword shows the 20 hosts with the most recent activity. The highest-threat keyword shows the 20 hosts that connected to the malware sites with the highest threat level. The subnet keyword shows up to 20 hosts within the specified subnet. The all keyword shows all buffered infected-hosts information. This display might include thousands of entries. You might want to use ASDM to generate a PDF file instead of using the CLI.
- **show dynamic-filter dns-snoop [detail]**—Shows the Botnet Traffic Filter DNS snooping summary, or with the detail keyword, the actual IP addresses and names. All inspected DNS data is included in this output, and not just matching names in the blacklist. DNS data from static entries are not included.

- **show asp table dynamic-filter [hits]**—Shows the Botnet Traffic Filter rules that are installed in the accelerated security path.

Firewall Redundancy

The Internet perimeter of the Small Enterprise Design Profile uses a single Cisco ASA appliance configured with redundant interfaces. The use of redundant interfaces makes the design resilient to link level failures, representing an affordable option for high availability. In cases where chassis redundancy is desirable, the enterprise may consider deploying a pair of Cisco ASA appliances configured for stateful failover. Both active/active and active/standby failover modes are supported. While stateful failover protects against chassis failures, it requires the deployment of two identical Cisco ASA appliances and the adjustment of the topologies around the firewalls, so its deployment should be carefully planned.

This guide explains the use of redundant interfaces. For information on how to configure stateful failover, refer to the *Cisco ASA 5500 Series Adaptive Security Appliances Configuration Guides* at the following URL:

http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.htm

A Cisco ASA redundant interface is a logical interface that pairs two physical interfaces, called active and standby interfaces. Under normal operation the active interface is the only one passing traffic. The active interface uses the IP address defined at the redundant interface, and the MAC address of the first physical interface associated with the redundant interface. When the active interface fails, the standby interface becomes active and starts passing traffic. The same IP address and MAC address are maintained so that traffic is not disrupted. Figure 5-23 illustrates the concept of redundant interface.

Figure 5-23 Cisco ASA Redundant Interface

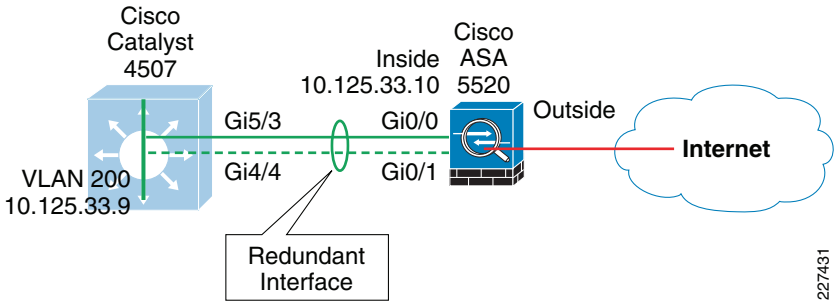

The configuration of a redundant interface consists in the configuration of the physical interface parameters and the logical redundant interface. Physical parameters such as media type, duplex, and speed are still configured within the physical interface. IP address, interface name, routing protocols, security level are configured as part of the redundant interface. The following configuration example corresponds to Figure 5-23.

```
! Physical interface and Ethernet parameters
interface GigabitEthernet0/0
 description Connected to cr24-4507-DO
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/1
 description backup to cr24-4507-DO
 no nameif
```

```
 no security-level
 no ip address
!
! Defines logical redundant interface associated with physical
! interfaces. Configures IP and logical interface parameters.
interface Redundant1
 description Connected to cr24-4507-DO
 member-interface GigabitEthernet0/0
 member-interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 10.125.33.10 255.255.255.0
 authentication key eigrp 100 <removed> key-id 1
 authentication mode eigrp 100 md5
!
```

Routing

An interior gateway protocol, EIGRP in our configuration examples, is used for dynamic routing. The Internet firewall may participate in routing by learning the internal routes and by injecting a default route pointing to the Internet. The default route should be removed dynamically if the Internet connection becomes unavailable.

As part of the Small Enterprise Design Profile, two different approaches were validated for the injection of the default route:

- *OSPF*—The Cisco ASA appliance learns the default route from the Internet border router using OSPF. The default route is then redistributed into EIGRP, and from there propagated into the rest of the internal network.
- *Static Route*—The Cisco ASA appliance is configured with a static default route pointing to the Internet gateway. Object tracking is configured to dynamically remove the default route when the Internet connection becomes unavailable. The default route is redistributed into EIGRP, and from there propagated into the rest of the internal network.

Injecting a default route with OSPF requires the configuration of an OSPF process between the Cisco ASA and the Internet border router, as illustrated in Figure 5-24. If the router is managed by the ISP, the configuration will require coordination with the service provider. This scenario also requires the default route to be propagated over OSPF. The actual default route may originate from the Internet border router itself or somewhere in the ISP network

Figure 5-24 OSPF Default Route Injection

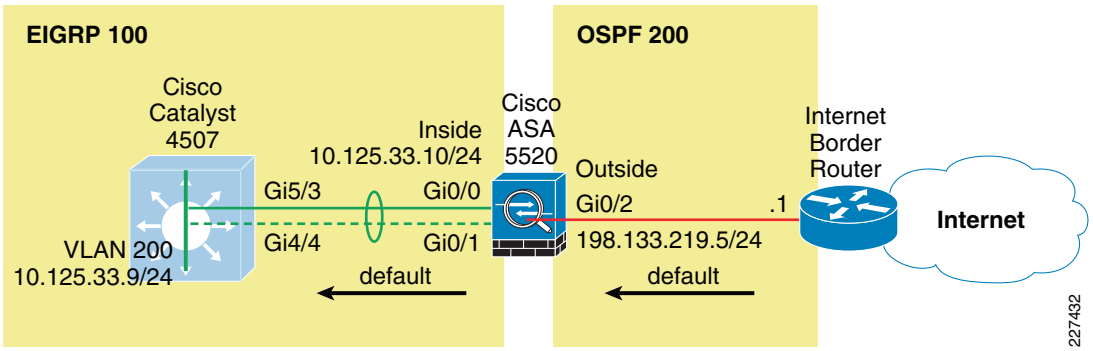

The following are the guidelines for using OSPF for the injection of a default route:

- Whenever possible, use MD5 authentication to secure the routing session between the Cisco ASA and the Internet border router.
- Since NAT is configured on the Cisco ASA and the inside address space is not visible outside the firewall, there is no need to redistribute routes from the internal EIGRP into OSPF.
- Route redistribution from OSPF into the internal EIGRP should be limited to the default route only. No other routes should be propagated into EIGRP.

The following configuration snippet illustrates the routing configuration of the Cisco ASA appliance. The configuration includes the route redistribution from OSPF into EIGRP with the enforcement of a route-map allowing only the injection of the default route. MD5 authentication is used for OSPF, and the logging of neighbor status changes is enabled.

```
! Permit default only
access-list Inbound-Routes standard permit host 0.0.0.0
!
interface GigabitEthernet0/2
 ospf message-digest-key 1 md5 <removed>
 ospf authentication message-digest
!
route-map Inbound-EIGRP permit 10
 match ip address Inbound-Routes
!
router eigrp 100
 no auto-summary
 network 10.125.33.0 255.255.255.0
 passive-interface default
 no passive-interface inside
 redistribute ospf 200 metric 1000000 2000 255 1 1500 route-map Inbound-EIGRP
!
router ospf 200
 network 198.133.219.0 255.255.255.0 area 100
 area 100 authentication message-digest
 log-adj-changes
!
```

Note The **hello-interval** and **dead-interval** OSPF timers can be adjusted to detect topological changes faster.

The other validated alternative for the default route injection is the definition of a static default route, which then can be redistributed into the internal EIGRP process. This is shown in Figure 5-25. This option does not require the configuration of the Internet border router.

Figure 5-25 Static Default Route with Object Tracking

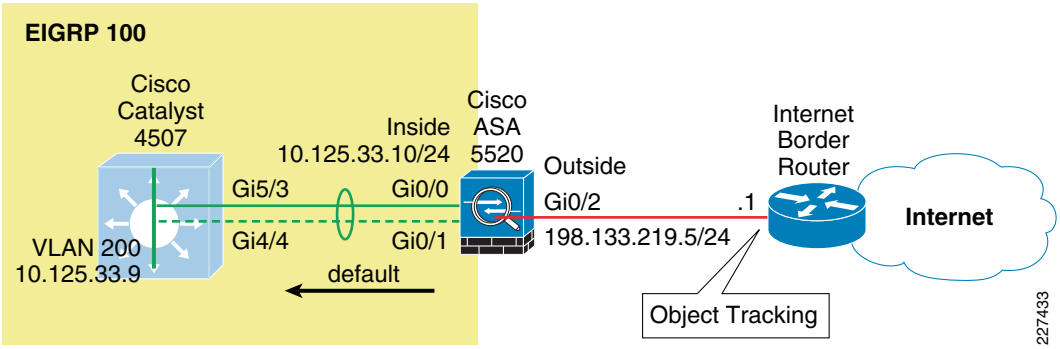

It is highly recommended to use object tracking so the default route is removed when the Internet connection becomes unavailable. Without object tracking, the default route will be removed only if the outside interface of the appliance goes down. So there is a possibility that the default route may remain in the routing table even if the Internet border router becomes unavailable. To avoid that problem, the static default route can be configured with object tracking. This consists in associating the default route with a monitoring target. The Cisco ASA appliance monitors the target using ICMP echo requests. If an echo reply is not received within a specified time period, the object is considered down and the associated default route is removed from the routing table.

The monitoring target needs to be carefully selected. First, pick one that can receive and respond to ICMP echo requests sent by the Cisco ASA. Second, it is better to use a persistent network object. In the configuration example below the Cisco ASA monitors the IP address of the next hop gateway, which helps identifying if the Internet gateway goes down, but it will not help if the connection is lost upstream. If available, you may want to monitor a persistent network object located somewhere in the ISP network. Static route tracking can also be configured for default routes obtained through DHCP or PPPoE.

In the following configuration the IP address of the next hop gateway (198.133.219.1) is used as the monitoring target. The static default route is then redistributed into EIGRP.

```
router eigrp 100
 no auto-summary
 network 10.125.33.0 255.255.255.0
 passive-interface default
 no passive-interface inside
 redistribute static metric 1000000 2000 255 1 1500
!
route outside 0.0.0.0 0.0.0.0 198.133.219.1 1 track 10
!
sla monitor 1
 type echo protocol ipIcmpEcho 198.133.219.1 interface outside
sla monitor schedule 1 life forever start-time now
!
track 10 rtr 1 reachability
```

Note The *frequency* and *timeout* parameters of object tracking can be adjusted to detect topological changes faster.

Intrusion Prevention Deployment

The Small Enterprise Design Profile implements Intrusion Prevention using an Advanced Inspection and Prevention Security Services Module (AIP SSM) on the Cisco ASA appliance deployed at the Internet perimeter. This section describes the best practices for integrating and configuring the IPS module for maximum threat control and visibility, as well as the deployment of the IPS Global Correlation feature.

Deploying IPS with the Cisco ASA

The Advanced Inspection and Prevention Security Services Module (AIP SSM) is supported on Cisco ASA 5510 and higher platforms. The AIP SSM runs advanced IPS software that provides proactive, full-featured intrusion prevention services to stop malicious traffic, including worms and network viruses, before they can affect your network.

As described in the "Intrusion Prevention Guidelines" section on page 5-10, the AIP SSM may be deployed in inline or promiscuous mode. In inline mode the AIP SSM is placed directly in the traffic flow, while in promiscuous mode the Cisco ASA sends a duplicate stream of traffic to the AIP SSM.

When deploying the AIP SSM in inline mode it is particularly important to determine how traffic will be treated in case of a module failure. The AIP SSM card may be configured to fail open or close when the module becomes unavailable. When configured to fail open, the Cisco ASA appliance allows all traffic through, uninspected, if the AIP SSM becomes unavailable. Per contrary, when configured to fail close, the adaptive security appliance blocks all traffic in case of an AIP SSM failure.

The following example illustrates how a Cisco ASA can be configured to divert all IP traffic to the AIP SSM in inline mode, and to block all IP traffic if the AIP SSM card fails for any reason:

```
access-list IPS permit ip any any
class-map my-ips-class
  match access-list IPS
policy-map my-ips-policy
  class my-ips-class
    ips inline fail-close
service-policy my-ips-policy global
```

IPS Global Correlation Deployment

There are a number of aspects that need to be understood prior to the deployment of the IPS Global Correlation feature. To start with, before configuration be sure that you are using Cisco IPS Version 7.0 with the latest patch and signature updates and that Cisco IPS is configured for network connectivity in either IDS or IPS mode.

The configuration of Global Correlation can be performed using the command-line interface (CLI), Cisco IDS Device Manager (IDM), Cisco IME, or the Cisco Security Manager. The figures here provided are screenshots from Cisco IME that illustrate the basic steps in the configuration of IPS Global Correlation.

The first step in configuring the IPS sensor (module or appliance) to use Global Correlation is to add either a DNS address and/or the proxy server setup. This step, illustrated in Figure 5-26, enables connection to Cisco SensorBase. After you configure the DNS and proxy settings, these settings will go into effect as soon as the sensor has downloaded the latest Global Correlation updates.

Figure 5-26 DNS and HTTP Proxy Within the Network Setting Configuration Screen

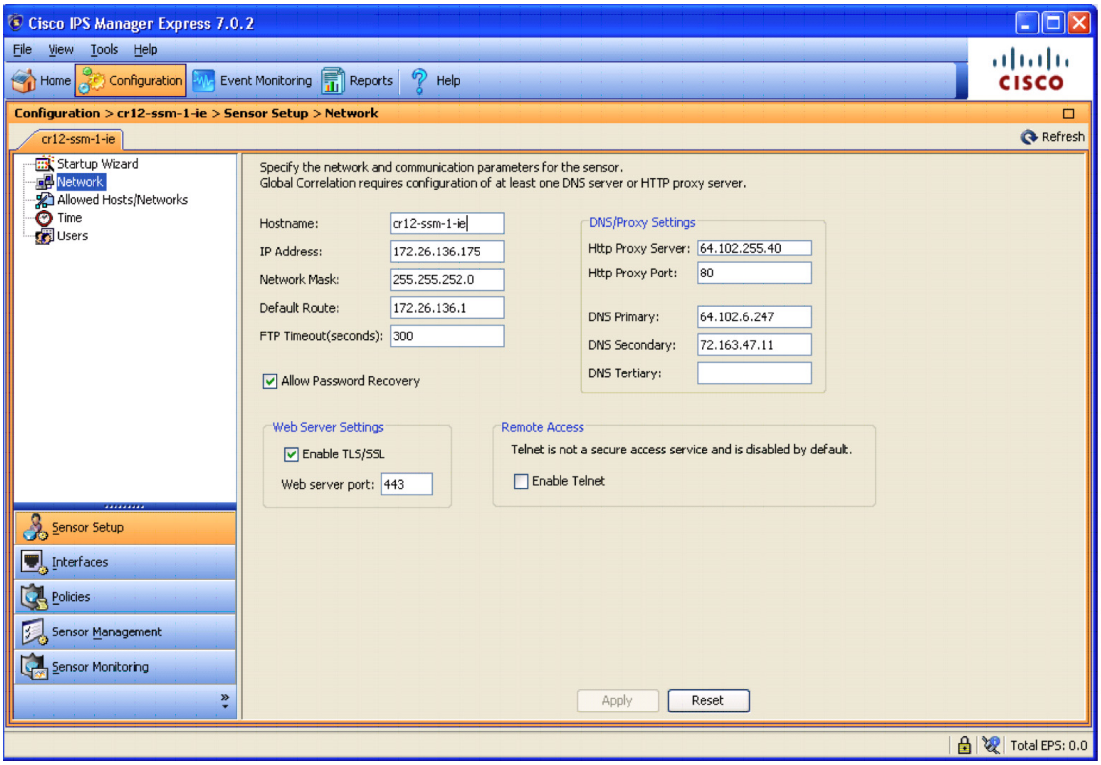

By default, a sensor runs Global Correlation Inspection in Standard mode, and enables the reputation filters, both illustrated in Figure 5-27. A good practice is to configure Global Correlation Inspection initially in permissive mode while monitoring the effects, and later change the configuration into Standard or Aggressive mode.

Figure 5-27 Global Correlation Inspection Settings

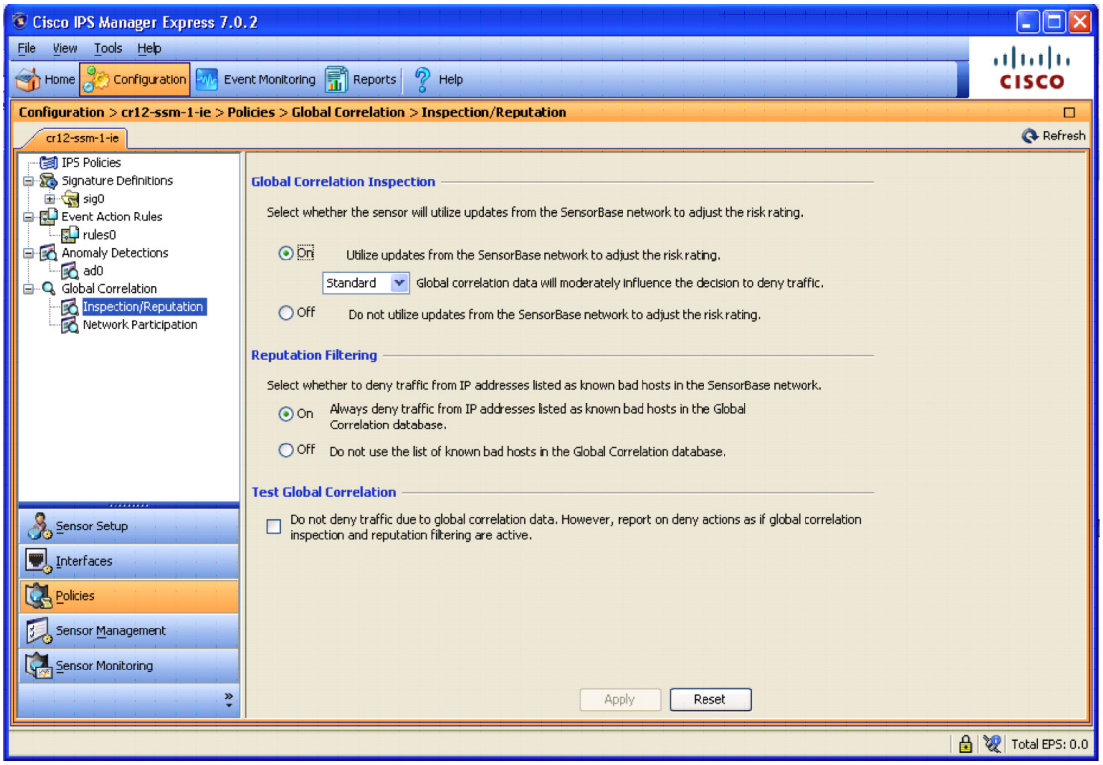

By default network participation is disabled, which means the sensor does not share any event data back to the Cisco SensorBase network. The event data provided by all devices participating in the Cisco SensorBase network is a key element that provides real-time and worldwide visibility into the threat activity, accelerating the identification and mitigation of threats propagating throughout the Internet. For this reason, Cisco recommends to configure the IPS sensors with partial or full network participation. See Figure 5-28.

Figure 5-28 Network Participation Settings (Off by Default)

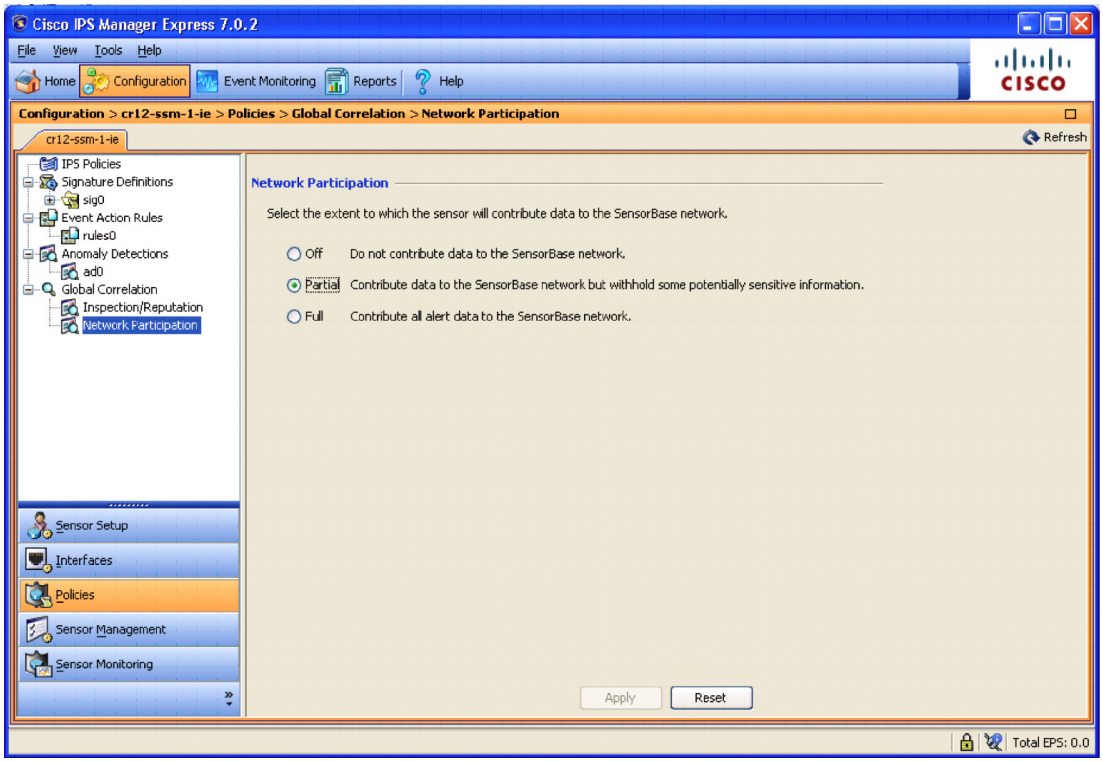

Event Monitoring with Global Correlation

Event monitoring with Global Correlation is similar to event monitoring with signature-only IPS. The primary difference is the potential addition of reputation scores representing the Global Correlation data. Figure 5-29 shows Cisco IPS events with reputation scores in Cisco IME.

Figure 5-29 Event Monitoring with Global Correlation in Cisco IME

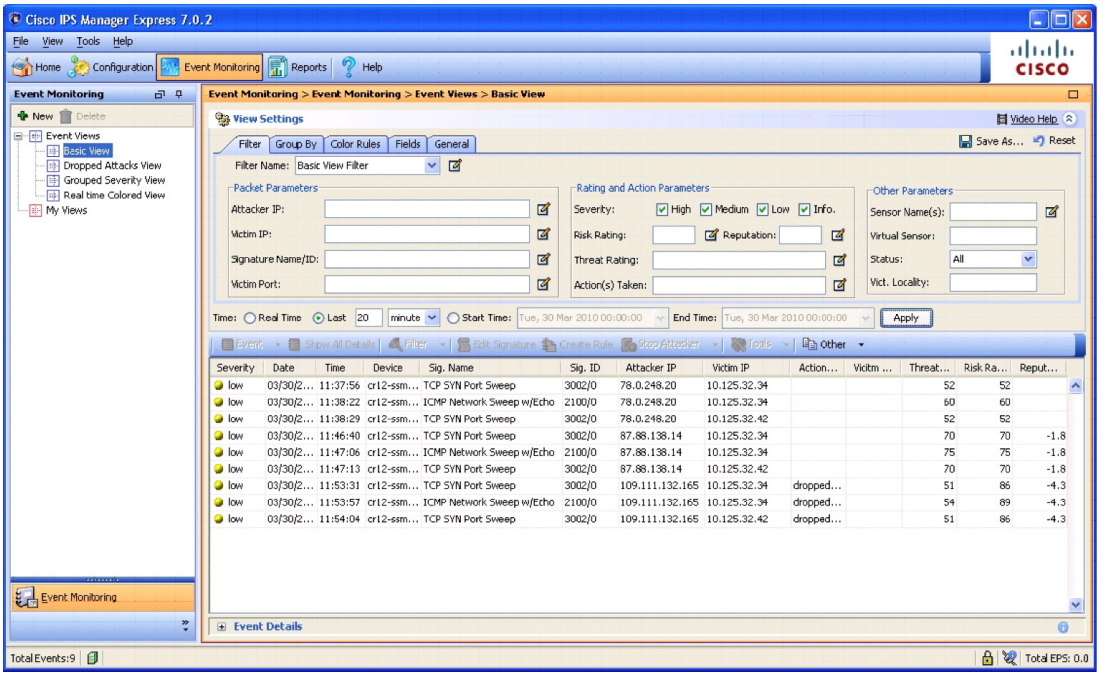

Figure 5-29 shows several TCP SYN Port Sweep and ICMP Network Sweep attacks that were seen by the sensor. The first three events had no reputation, and the event's risk ratings were 52 and 60 which did not meet the threshold for the packets to be dropped. The next three events were identical except that the attacker had a negative reputation of -1.8 which elevated the risk ratings to 70 and 75 which still did not meet the thresholds to be dropped in Standard Mode. The last events were also identical except this time the attacker has a negative reputation if -4.3 which elevated the risk ratings to 86 and 89. This time the risk rating was high enough for the packets to be dropped.

Figure 5-30 illustrates the detail view of the TCP SYN Port Sweep event coming from the attacker with a negative reputation of -4.3.

Figure 5-30 Detailed view of a TCP SYN Port Sweep from an Attacker with a Negative Reputation Score

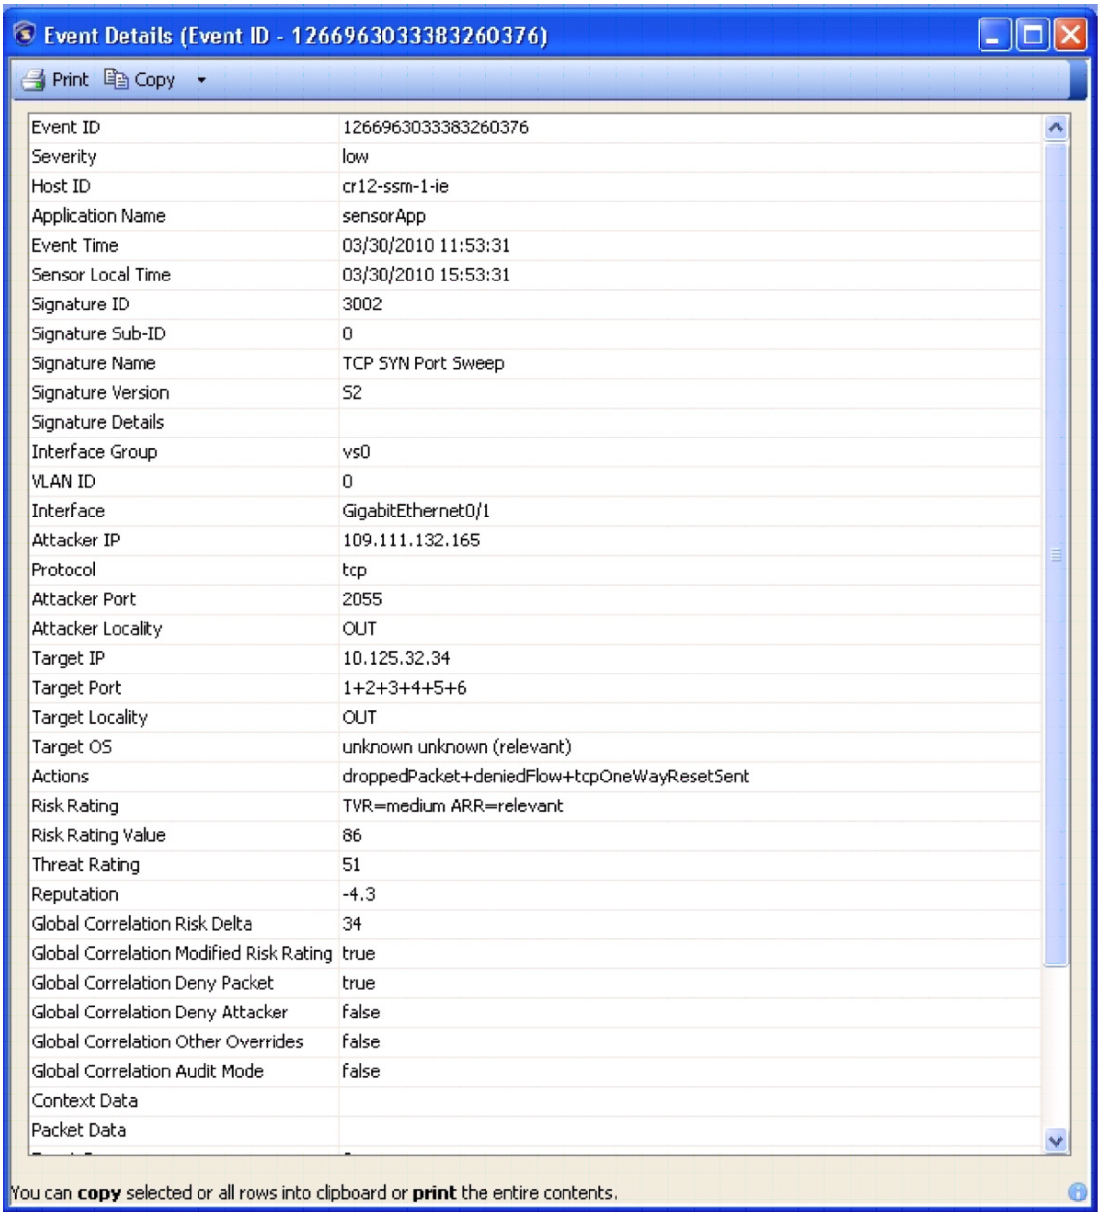

Web Security Deployment

The Small Enterprise Design Profile implements a Cisco IronPort WSA at the core/distribution layer of the main site, as illustrated in Figure 5-31. The WSA is located at the inside of the Cisco ASA acting as the Internet firewall. That ensures that clients and WSA are reachable over the same inside interface of the firewall, and that the WSA can communicate with them without going through the firewall. At the same time, deploying the WSA at the core/distribution layer gives complete visibility to the WSA on the traffic before getting out to the Internet through the firewall.

Figure 5-31 Cisco IronPort WSA deployment

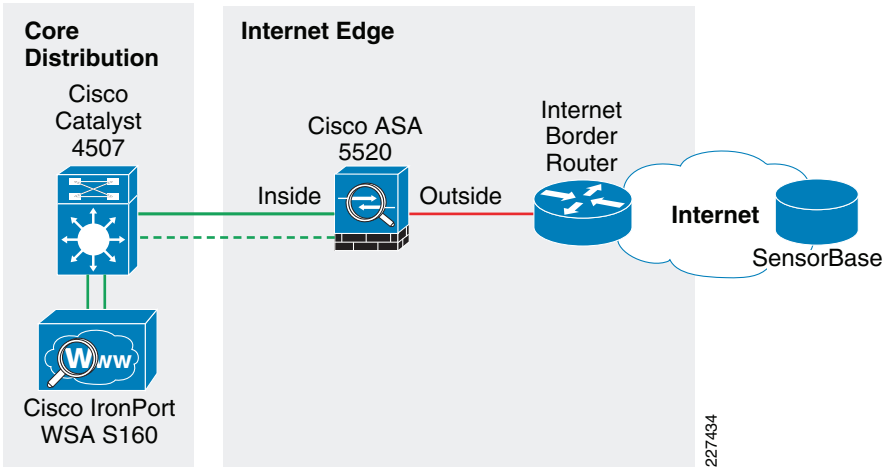

Following are the guidelines for the WSA configuration and deployment.

Initial System Setup Wizard

The WSA provides a browser-based system setup wizard that must be executed the first time the appliance is installed. The System Setup Wizard guides the user through initial system configuration such as network and security settings. It should be noted that running the initial System Setup Wizard completely reconfigures the WSA appliance and resets the administrator password. Only use the System Setup Wizard the first time you install the appliance, or if you want to completely overwrite the existing configuration.

The following are some of the default settings when running the System Setup Wizard:

- Web Proxy is deployed in transparent mode.
- The L4 Traffic Monitor is active and set to monitor traffic on all ports.

All these settings can be changed any time after the initial configuration by running the WSA web-based configuration GUI.

Interface and Network Configuration

The following need to be completed as part of the initial setup of the WSA appliance:

Step 1 Configuring network interfaces

Step 2 Adding routes

Step 3 Configuring DNS

Step 4 Setting time

Step 5 Working with upstream proxy (if present)

Configuring Network Interfaces

Independently from the model, all Cisco IronPort WSA appliances are equipped with six Ethernet interfaces as shown in Figure 5-32.

Figure 5-32 WSA Interfaces

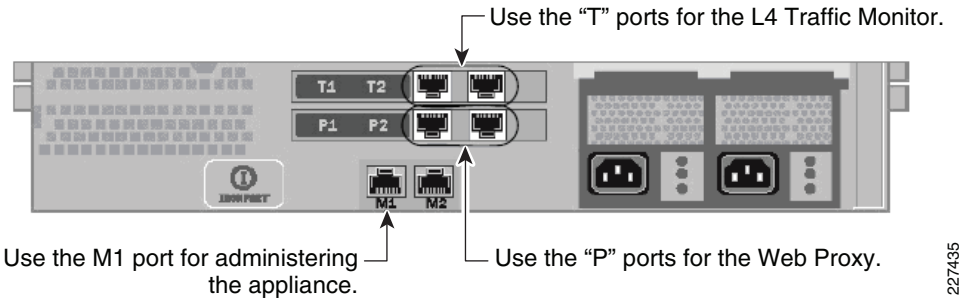

The WSA interfaces are grouped for the following functions:

- *Management*—Interfaces M1 and M2 are out-of-band (OOB) management interfaces. However, only M1 is enabled. In the Small Enterprise Design Profile, interface M1 connects to the out-of-band management network. Interface M1 can optionally be used to handle data traffic in case the enterprise does not have an out-band management network.
- *Web Proxy*—Interfaces P1 and P2 are Web Proxy interfaces used for data traffic. Only the P1 interface is used in the Enterprise Design Profile for Small Enterprise Networks. P1 connects to the inside subnet of the firewall.
- *L4 Traffic Monitor (L4TM)*—T1 and T2 are the L4TM interfaces. L4TM is not used in the Small Enterprise Design Profile, Cisco IPS Global Correlation and the Cisco ASA Botnet Filter features are used instead.

Figure 5-33 illustrates the network topology around the WSA used in the Cisco validation lab.

Figure 5-33 WSA Network Topology

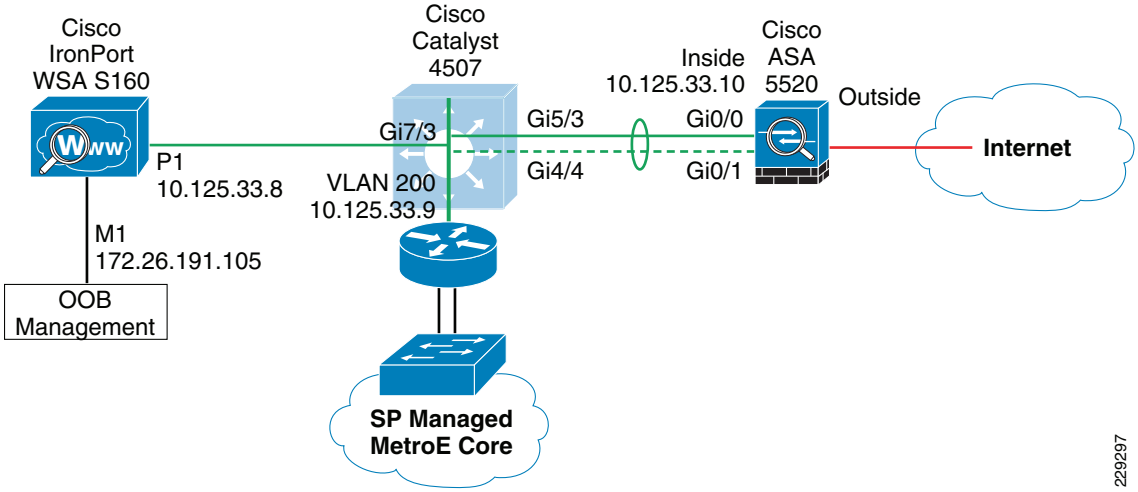

Figure 5-34 illustrates the IP address and hostname configurations for the interfaces used. In this case, an out-of-band management network is used; therefore the M1 port is configured with an IP address in the management subnet. In addition, the WSA is configured to maintain a separate routing instance for the M1 management interface. This allows the definition of a default route for management traffic separate from the default route used for data traffic.

Figure 5-34 WSA Interface Configuration

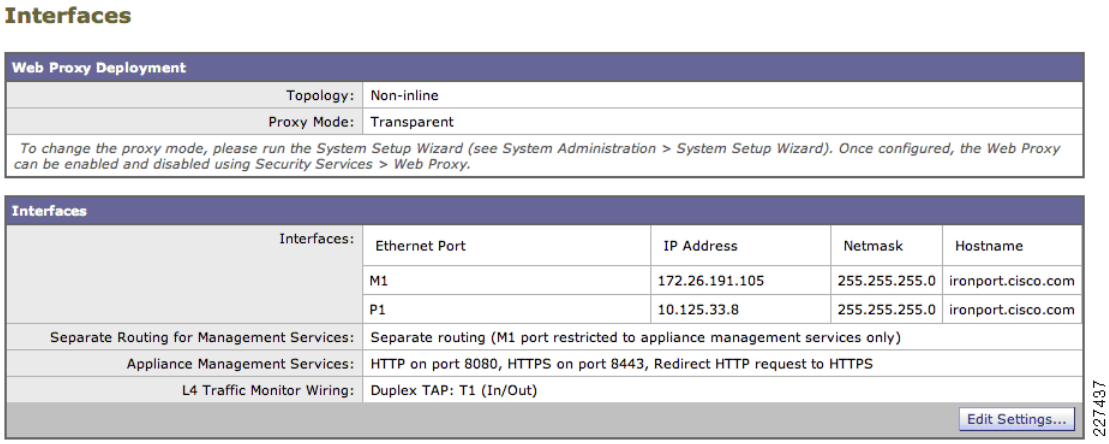

Adding Routes

A default route is defined for management traffic pointing to the OOB management default gateway (172.26.191.1). A separate default route is defined for the data traffic pointing to the inside IP address of the firewall (10.125.33.10). As all internal networks are reachable throughout the core/distribution switch, a route to 10.0.0.0/8 is defined pointing to the switch IP address (10.125.33.9) to allow the WSA to communicate with the clients directly without having to go to the firewall first. These settings are illustrated in Figure 5-35.

Figure 5-35 WSA Route Configuration

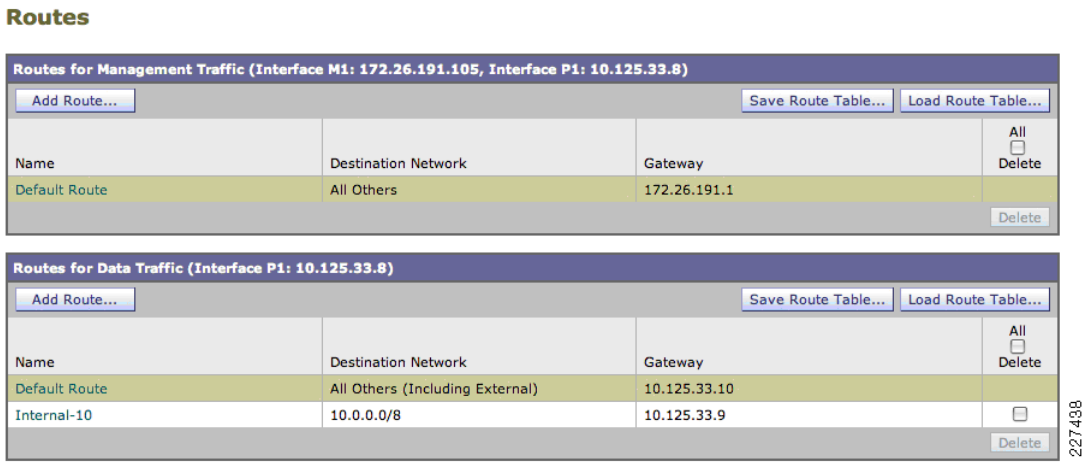

Configuring DNS

The initial setup requires the configuration of a hostname for the WSA appliance, and listing the DNS servers. Figure 5-36 shows the DNS configuration.

Figure 5-36 WSA DNS Configuration

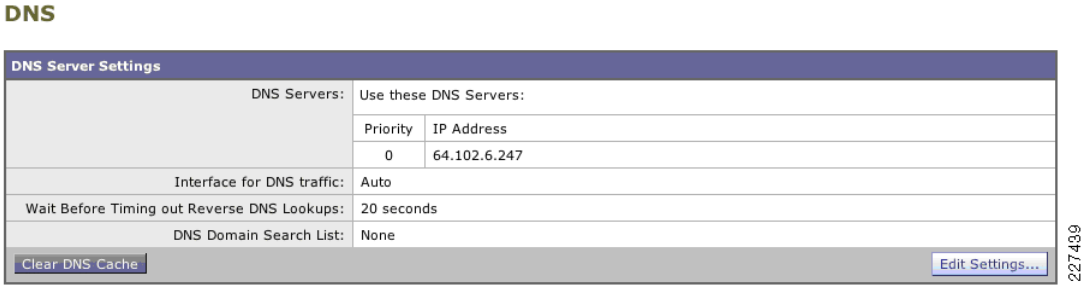

Setting Time

Time synchronization is critical for forensic analysis and troubleshooting, therefore enabling NTP is highly recommended. Figure 5-37 shows how the WSA is configured to synchronize its clock with an NTP server located on the OOB management network.

Figure 5-37 WSA NTP Configuration

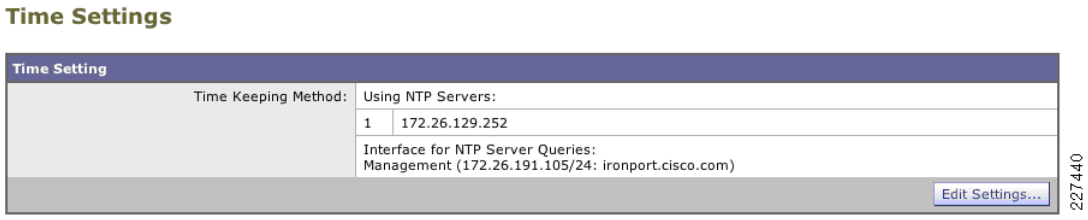

Working with Upstream Proxies

If Internet access is provided by an upstream proxy, then the WSA must be configured to use the proxy for component updates and system upgrades. This is illustrated in Figure 5-38 and Figure 5-39.

Figure 5-38 WSA Upgrade Settings

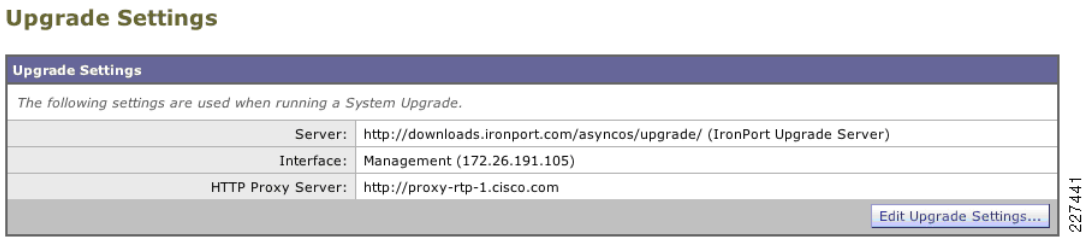

Figure 5-39 WSA Component Updates

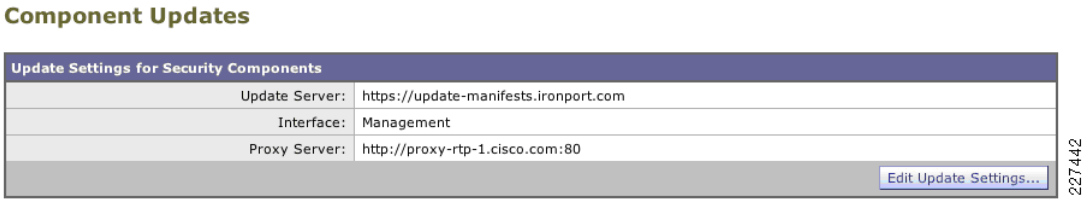

WCCP Transparent Web Proxy

The configuration of the WCCP Transparent Web Proxy includes the following:

Step 1 Defining WSA WCCP Service Group

Step 2 Enabling WSA Transparent Redirection

Step 3 Enabling WCCP redirection on the Cisco ASA

Step 4 Enabling WSA HTTPS scanning

Step 5 Working with upstream proxy (if present)

Defining WSA WCCP Service Group

Web Proxy settings are configured as part of an initial setup using the System Setup Wizard and can be later modified with the WSA Web-based GUI. The Web Proxy setting include the following:

- *HTTP Ports to Proxy*—List the ports to be proxied. Default is 80 and 3128.
- *Caching*—Defines whether or not the WSA should cache response and requests. Caching helps reduce latency and the load on the Internet links. Default is enabled.
- *IP Spoofing* — Defines whether or not the Web Proxy should spoof IP addresses when forwarding requests to upstream proxies and servers. The Cisco ASA does not support source address spoofing.

Figure 5-40 illustrates the Web Proxy settings.

Figure 5-40 WSA Proxy Settings

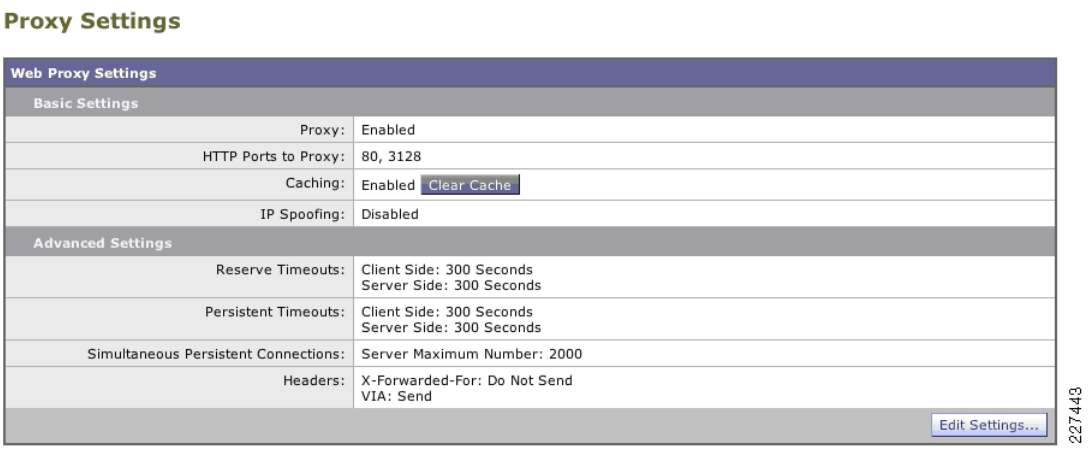

Enabling WSA Transparent Redirection

Configuring WCCP Transparent Redirection requires the definition of a WCCP service profile in the WSA. If redirecting HTTP and HTTPS, define a dynamic service ID to be used with the Cisco ASA. Use MD5 authentication to protect the WCCP communication between the WSA and Cisco ASA. Figure 5-41 shows an example.

Figure 5-41 WSA Transparent Proxy

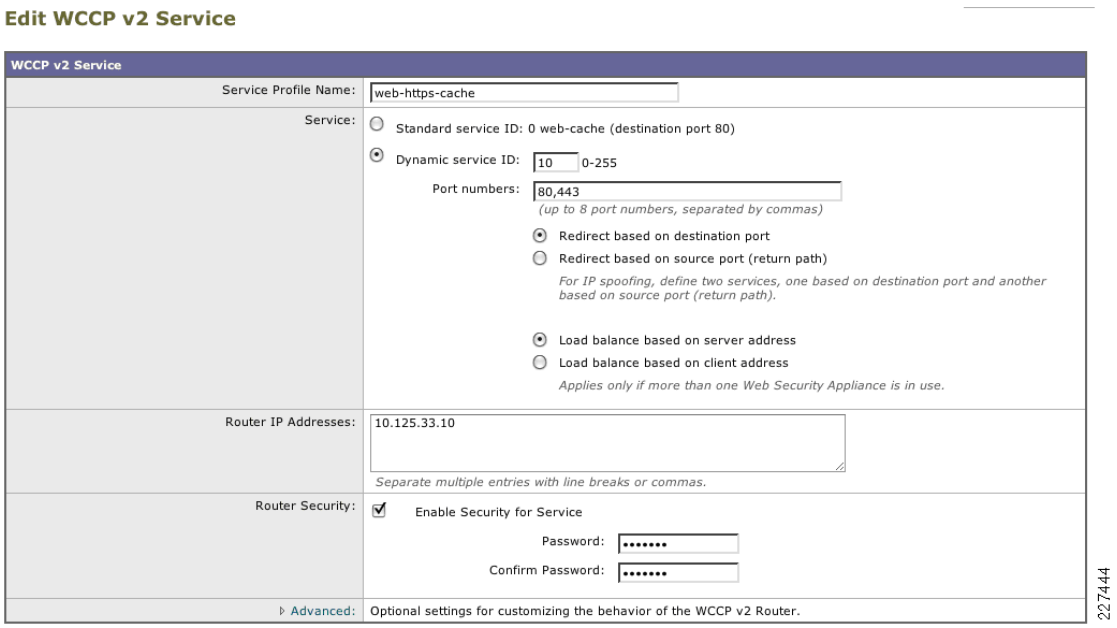

Enabling WCCP Redirection on Cisco ASA

The configuration of WCCP on the Cisco ASA appliance requires:

- A group-list indicating the IP addresses of the appliances member of the service group. In the example provided below the group-list is called **wsa-farm**.
- A redirect-list indicating the ports and subnets of traffic to be redirected. In the example, the ACL named proxylist is configured to redirect any HTTP and HTTPS traffic coming from the 10.0.0.0/8 subnet. It is critical to ensure traffic from the WSA(s) bypasses redirection. To that end, add an entry to the redirect-list explicitly denying traffic sourced from the WSA(s).
- WCCP service indicating the service ID. Make sure you use the same ID as defined on the WSAs. Use a password for MD5 authentication.
- Enabling WCCP redirection on an interface. Apply the WCCP service on the inside interface of the Cisco ASA.

The following is a Cisco ASA WCCP configuration example:

```
! Group-list defining the IP addresses of all WSAs
access-list wsa-farm extended permit ip host 10.125.33.8 any
!
! Redirect-list defining what ports and hosts/subnets should be redirected
access-list proxylist extended deny ip host 10.125.33.8 any
access-list proxylist extended permit tcp 10.0.0.0 255.0.0.0 any eq www
access-list proxylist extended permit tcp 10.0.0.0 255.0.0.0 any eq https
!
! WCCP service
wccp 10 redirect-list proxylist group-list wsa-farm password cisco
!
! Applies WCCP on an interface
wccp interface inside 10 redirect in
```

The WCCP connection status and configuration can be monitored on the Cisco ASA with the **show wccp** command, as shown below:

```
cr26-asa5520-do# show wccp

Global WCCP information:
    Router information:
    Router Identifier:                   198.133.219.5
    Protocol Version:                    2.0

    Service Identifier: 10
    Number of Cache Engines:             1
    Number of routers:                   1
    Total Packets Redirected:            428617
    Redirect access-list:                proxylist
    Total Connections Denied Redirect:   0
    Total Packets Unassigned:            4
    Group access-list:                   wsa-farm
    Total Messages Denied to Group:      0
    Total Authentication failures:       0
    Total Bypassed Packets Received:     0
cr26-asa5520-do#
```

Enabling WSA HTTPS Scanning

To monitor and decrypt HTTPS traffic, you must enable HTTPS scanning on the WSA. The HTTPS Proxy configuration is illustrated in Figure 5-42.

Figure 5-42 WSA HTTPS Proxy

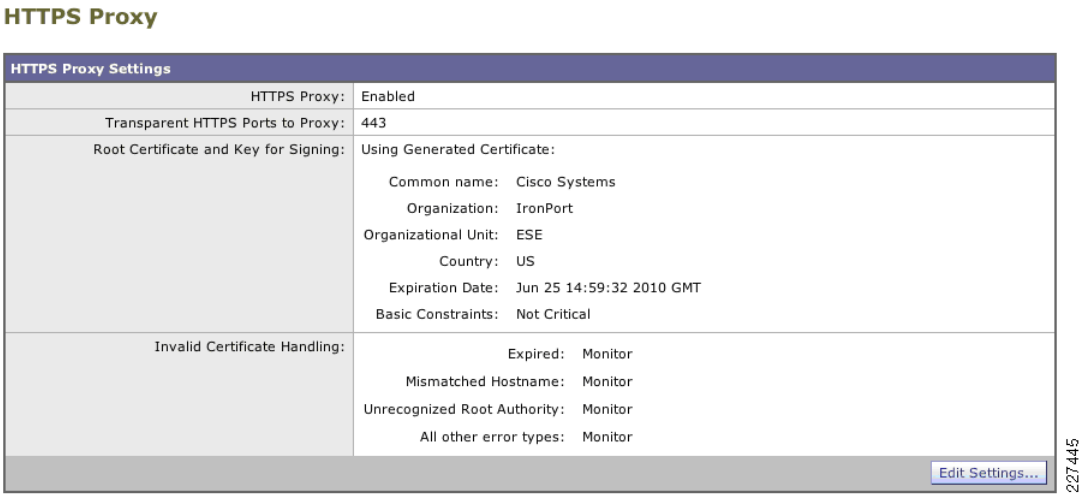

Working with Upstream Proxies

In case Internet traffic is handled by one or more upstream proxies, follow these guidelines:

- Add an Upstream Proxy Group
- Define a routing policy to direct traffic to the upstream proxies

The Upstream Proxy Group lists the IP addresses or domain names of the proxies to be used for traffic sent to the Internet. When multiple proxies are available, the WSA can be configured for failover or load balancing.

The following are the options available:

- *None (failover)*—The first proxy in the list is used. If one proxy cannot be reached, the Web Proxy attempts to connect to the next one in the list.
- *Fewest connections*—Transactions are directed to the proxy servicing the fewest number of connections.
- *Hash-based*—Requests are distributed using a hash function. The function uses the proxy ID and URL as inputs so that requests for the same URL are always directed to the same upstream proxy.
- *Least recently used*—Transactions are directed to the proxy that least recently received a transaction if all proxies are currently active.
- *Round robin*—The Web Proxy cycles transactions equally among all proxies in the group in the listed order.

Figure 5-43 illustrates the upstream proxy group configuration. Two upstream proxies are used, and transactions are forwarded to the proxy servicing the fewest number of connections.

Figure 5-43 WSA Upstream Proxy Group

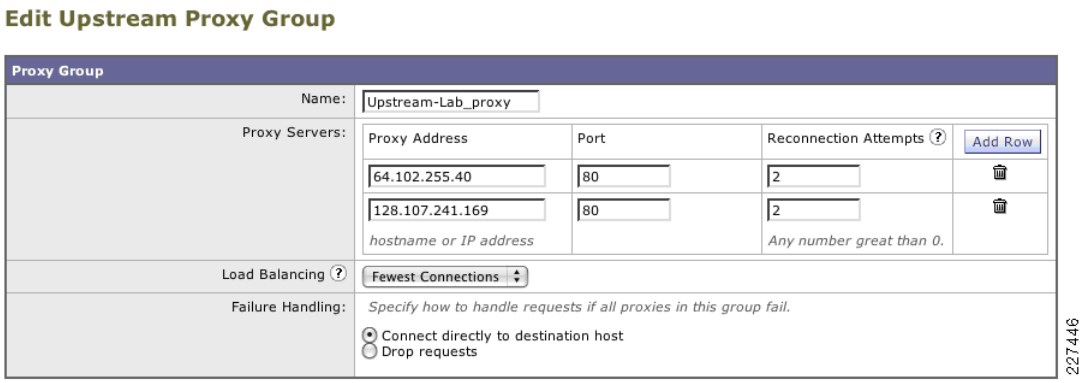

Next, a routing rule needs to be defined to indicate when and how to direct transactions to the upstream proxy group. Use the Global Routing Policy if all traffic is to be handled by the upstream proxies. If no proxies are present, then leave the routing destination of the Global Routing Policy configured as **Direct Connection**. Figure 5-44 presents an example where all traffic is directed to the proxies in the Upstream-Lab_proxy group.

Figure 5-44 WSA Routing Policies

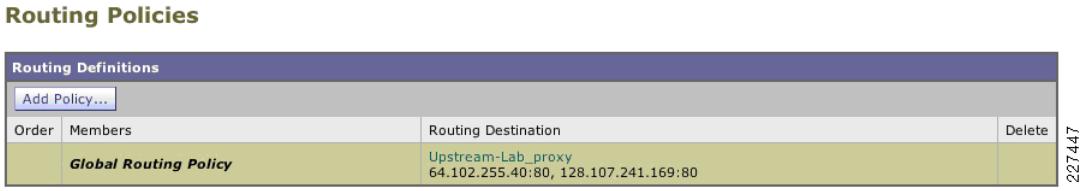

Web Access Policies

The access policies define how the Web Proxy handles HTTP requests and decrypted HTTPS connections for network users. By configuring access policies the enterprise can control what Internet applications (instant messaging clients, peer-to-peer file-sharing, web browsers, Internet phone services, etc.) and URL categories users may access. In addition, access policies can be used to block file downloads based on file characteristics, such as file size and file type.

The WSA comes with a default Global Policy that applies to all users. However, multiple policies can be defined when different policies need to be applied to different group of users. Figure 5-45 shows the global policy.

Figure 5-45 Global Access Policy

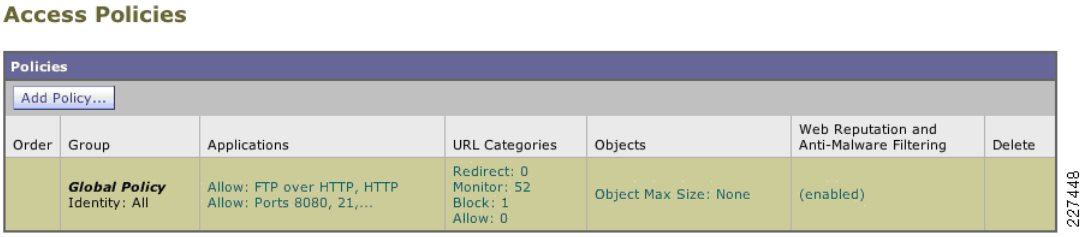

URL categories corresponding to non-business related content should be blocked in compliance with the company's Internet access policies. Figure 5-46 provides an example on how the "Adult/Sexually Explicit" category is blocked.

Figure 5-46 URL Categories

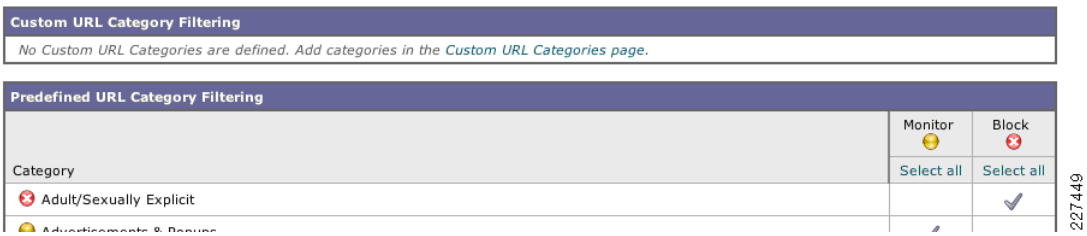

CISF Protected Ports

Catalyst Integrated Security Features (CISF) is a set of native security features available on Cisco Catalyst Switches that protect the network against attacks such as man-in-the-middle, spoofing, and infrastructure denial-of-services (DoS) attacks. CISF includes the following:

- Port Security
- DHCP snooping
- IP Source Guard
- Dynamic ARP inspection
- ARP rate limiting
- Storm Control

The following is a sample CISF port configuration:

```
! configure port security parameters
switchport port-security maximum 2
switchport port-security
switchport port-security aging time 2
switchport port-security violation restrict
switchport port-security aging type inactivity
! configure arp inspection rate limiting
ip arp inspection limit rate 100
! configure dhcp snooping rate limiting
ip dhcp snooping limit rate 100
! configure storm control parameters
storm-control broadcast level 20.00 10.00
storm-control multicast level 50.00 30.00
```

NAC Appliance Deployment

In the Small Enterprise Design Profile, a NAC Appliance solution is deployed at all locations, the main site and each of the remote offices. To that end, a centralized CAM is deployed at the main site, likely within the serverfarm. And a CAS is deployed at the main site and each remote site, and each of which directly connects to the core/distribution layer at each of the locations.

Deploying the NAC appliance solution requires the configuration of the CAM and CAS components. The initial CAM and CAS configuration is performed via console access, and this is described in the installation guide, *Cisco NAC Appliance Hardware Installation.* During the configuration stage the multiple steps must be followed in configuring the NAC Appliance with IP addresses, VLANs. passwords, etc. The installation guide contains worksheets that assist in the gathering and preparation of this information for both CAM and CAS.

After performing the initial setup through the console, the rest of the configuration of the CAM and CAS is performed using the CAM web-based GUI. The first task on the CAM before beginning configuration is the installation the licenses for the solution. A license must be installed for the CAM, and the CAS servers controlled by the CAM. The *Cisco NAC Appliance Ordering Guide*, provides information on the ordering options.

Licenses can be entered via the CAM web-based GUI, as shown in Figure 5-47.

Figure 5-47 NAC Appliance Licensing

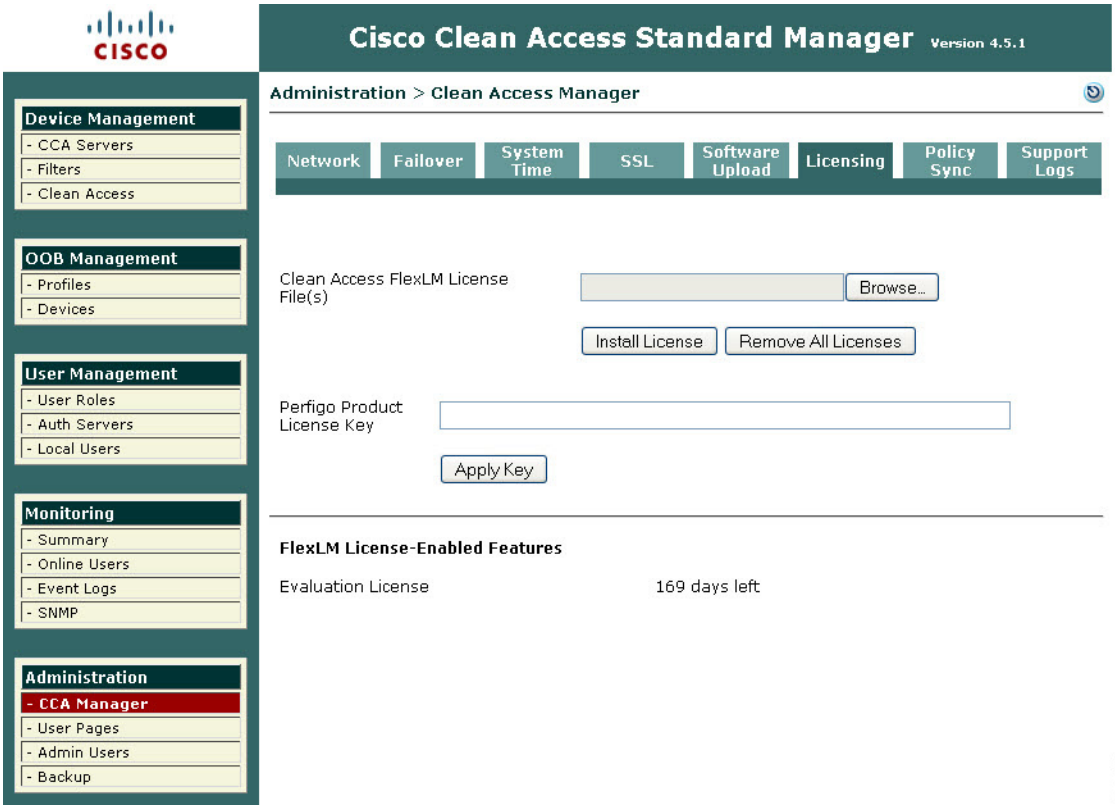

Adding a CAS to the CAM

For a CAS server to be managed by the CAM, it must be first added to the list of managed servers on the CAM. To do this the CAM needs to know the IP address of the CAS, and the Server Type (its role in the network) of the CAS. In addition to this, the CAS and the CAM must have the same shared secret. The shared secret is configured during the server installation. An example of adding a CAS to the CAM is shown in Figure 5-48.

Figure 5-48 Adding a new CAS to the CAM

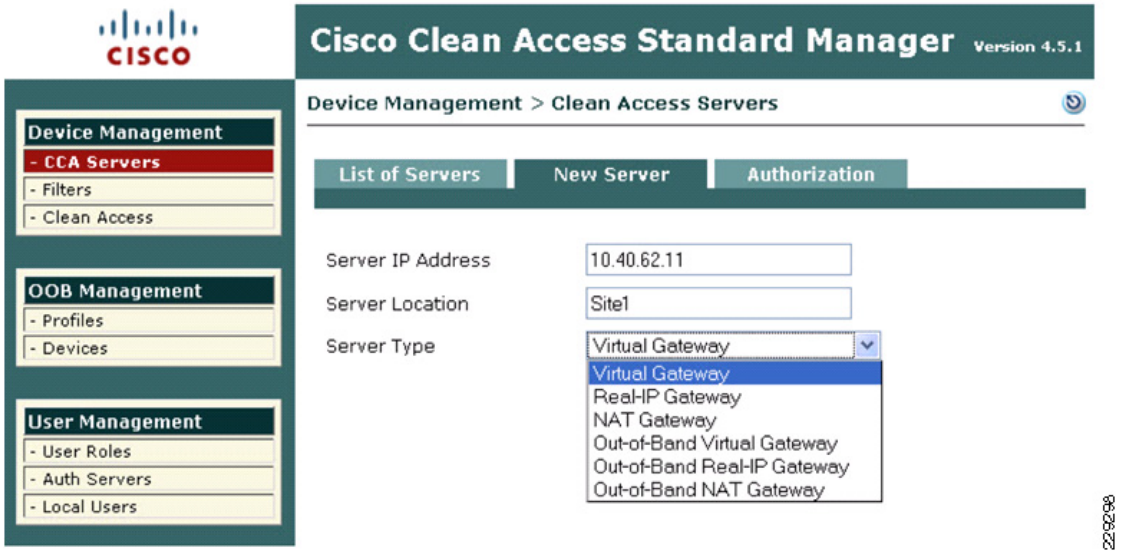

Once the CAS has been added to the CAM, it appears in the list of servers on the CAM. From this point, it can be managed directly from the CAM for almost all tasks. An example of a list of servers is shown in Figure 5-49.

Figure 5-49 List of CAS Servers

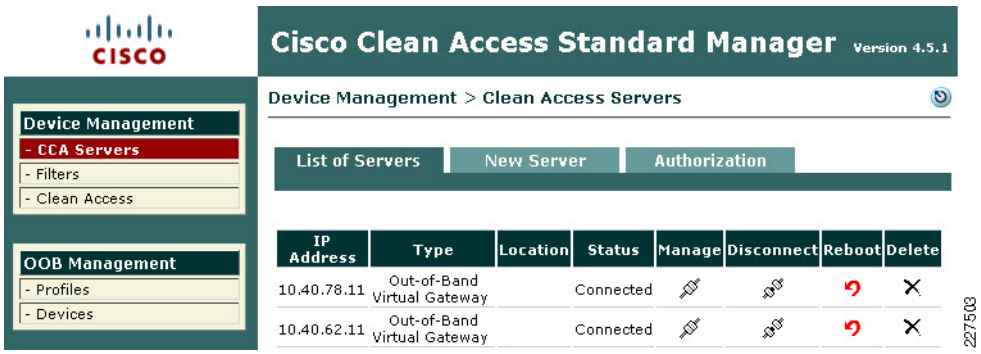

Managing the CAS

Once the CAS is in the list of servers managed by the CAM, it can be configured further for its role in the network. To manage the server click the icon under the Manage heading in the server list, this will connect you to the CAS server and present you with the summary menu shown in Figure 5-50.

Figure 5-50 AS Management Menu

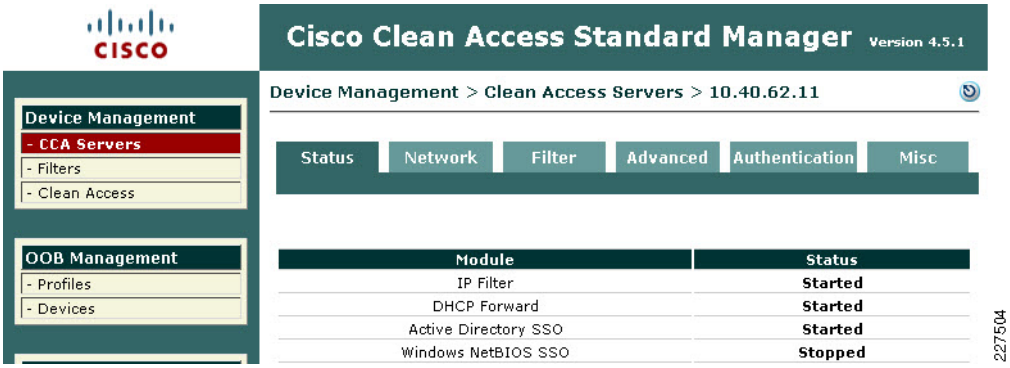

Under the CAS Network setting tab, shown in Figure 5-51, the basic network settings for the CAS can be seen and altered, if needed. In this example, we are keeping the network configuration performed during the server installation. The primary dialog under the Network Tab is the IP dialog, shown in Figure 5-51, the other dialogs allow the DHCP options to the configured—our example uses the default of DHCP passthrough— and the DNS options where host name, domain name, and DNS server information is added, as shown in Figure 5-52.

Figure 5-51 CAS Network Settings

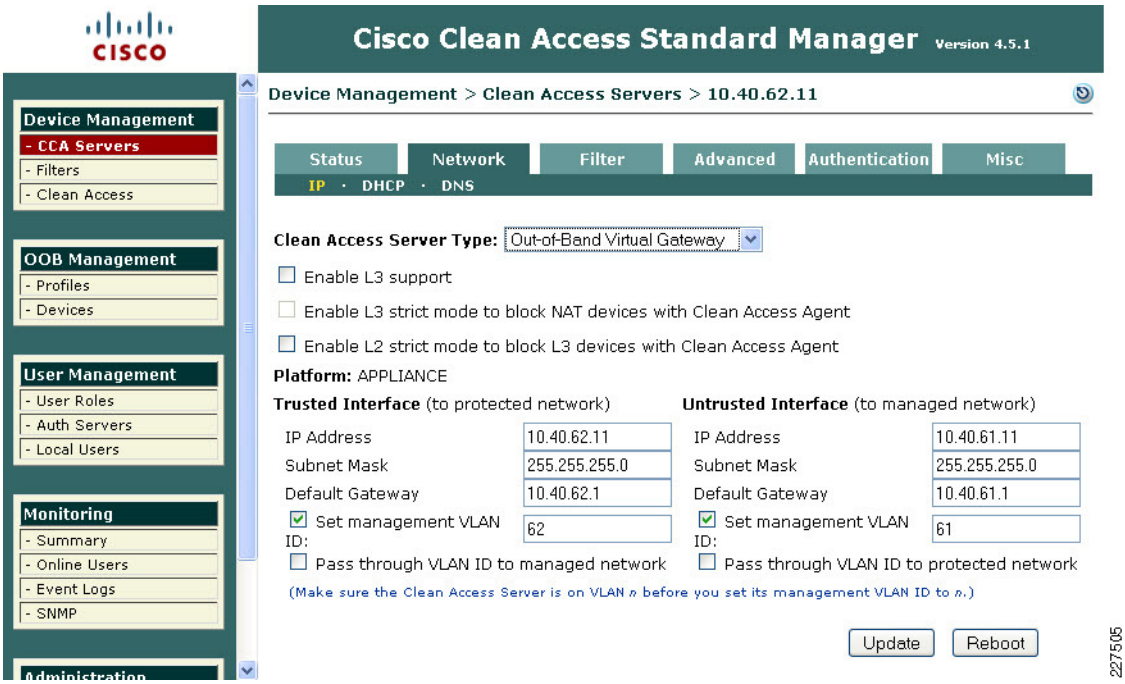

Figure 5-52 CAS DNS Settings

The next tab in the CAS configuration is the Filter tab (see Figure 5-53), for the purposes of our example the important dialog is the Roles where network traffic filters may be applied to different user Roles. The Role of interest at this moment is the default Unauthenticated Role. By default the Unauthenticated Role blocks all traffic. In this example we are allowing the Unauthenticated Role to pass Active Directory client authentication traffic to pass to the Active Directory Sever. This will allow a windows client to join the active Directory Domain, and windows users to authenticate to the domain although they have not been through the NAC process. This is often important to allow printer and drive mapping information to be sent to the winders users. As the user has already authenticated to the Active Directory Domain the user authentication information maybe learned from Active Directory, and the user does not have to reauthenticate for the NAC server.

Note The creation of Roles and their associated filters is performed in the CAM **User Management -> User Roles** menu.

Figure 5-53 CAS Filter Settings

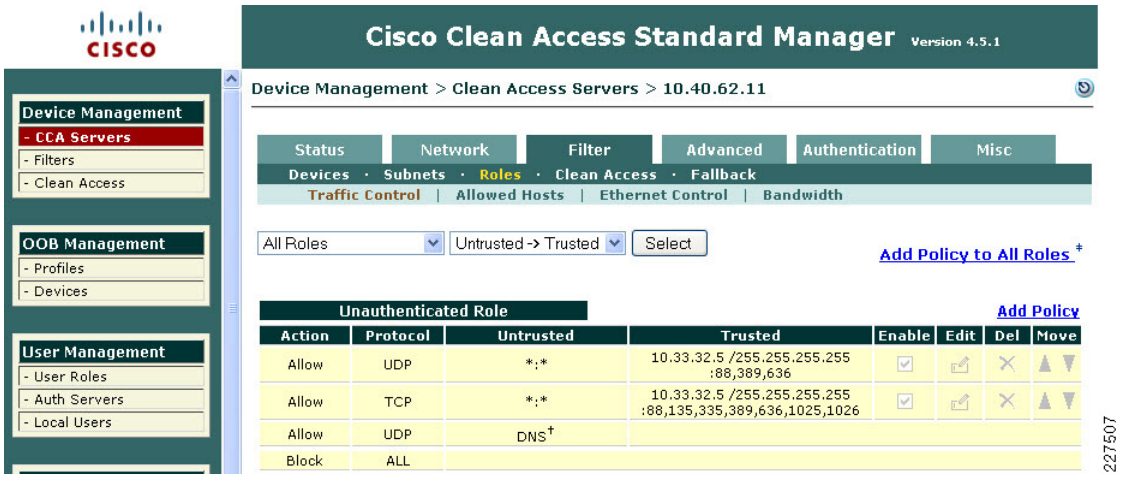

The next tab that requires configuration is the Advanced tab. This has multiple dialogs that require configuration. The first of these is the Managed Subnet dialog, where each of the trusted VLAN subnets is added to the CAS for management. An example of this shown in Figure 5-54.

Figure 5-54 CAS Managed Subnet

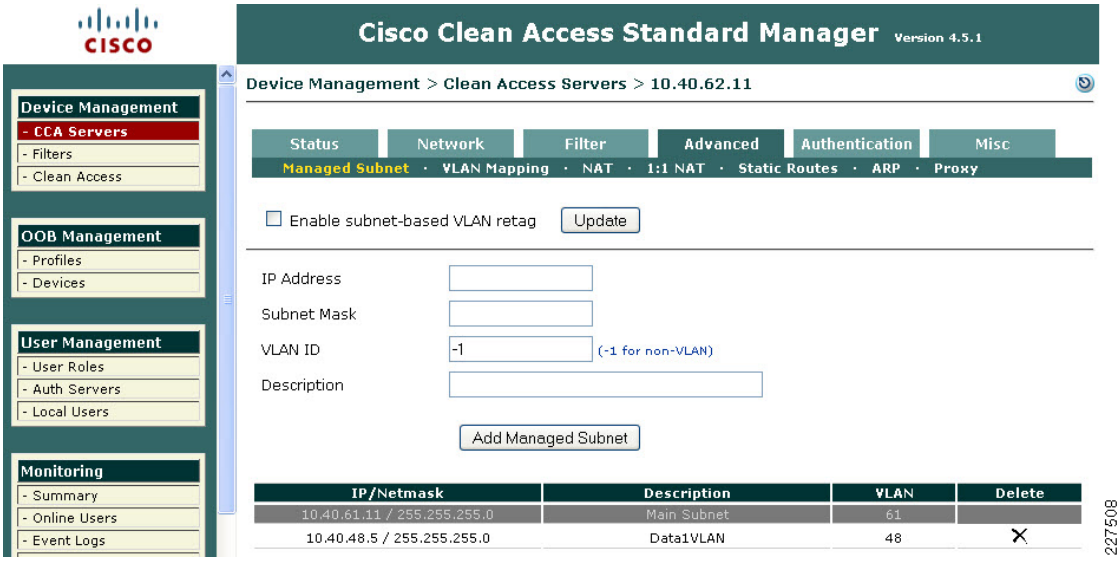

The next dialog of the Advanced Tab is the VLAN Mapping dialog, which tells the CAS which trusted VLAN to be mapped to an untrusted VLAN, an example of this is shown in Figure 5-55. In our example, VLAN Prunning and VLAN Mapping is also enabled.

Figure 5-55 VLAN Mapping

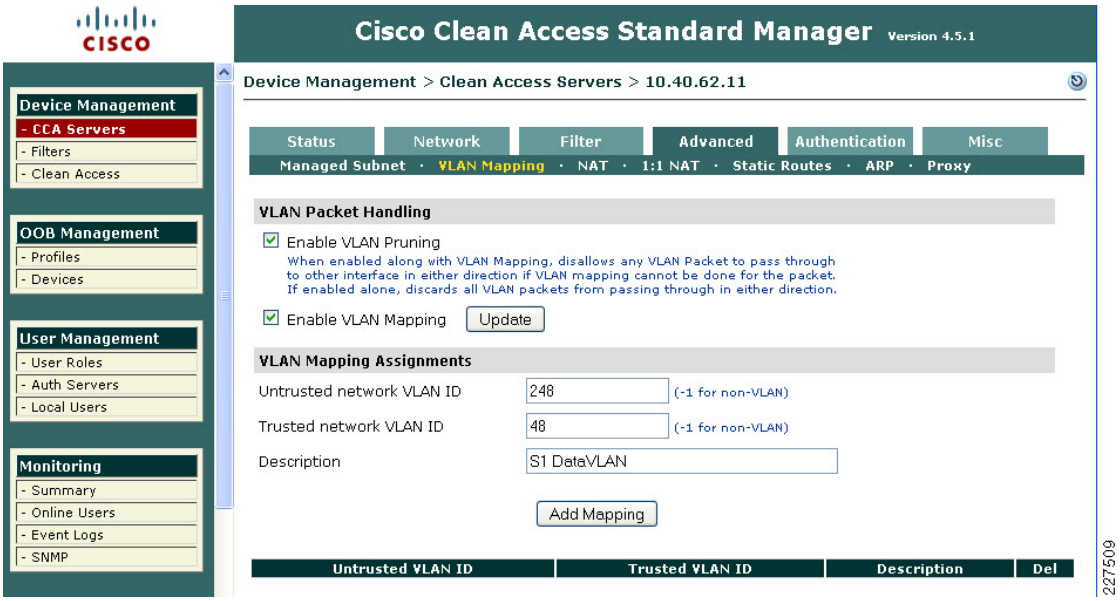

The next tab of interest is the Authentication Tab (see Figure 5-56), this tab has multiple dialogs for configuring different authentication options. The first dialog is the Login Page Dialog. This allows the configuration of different web login pages depending upon the untrusted subnet being used for authenticating client.

Figure 5-56 Authentication Login Page

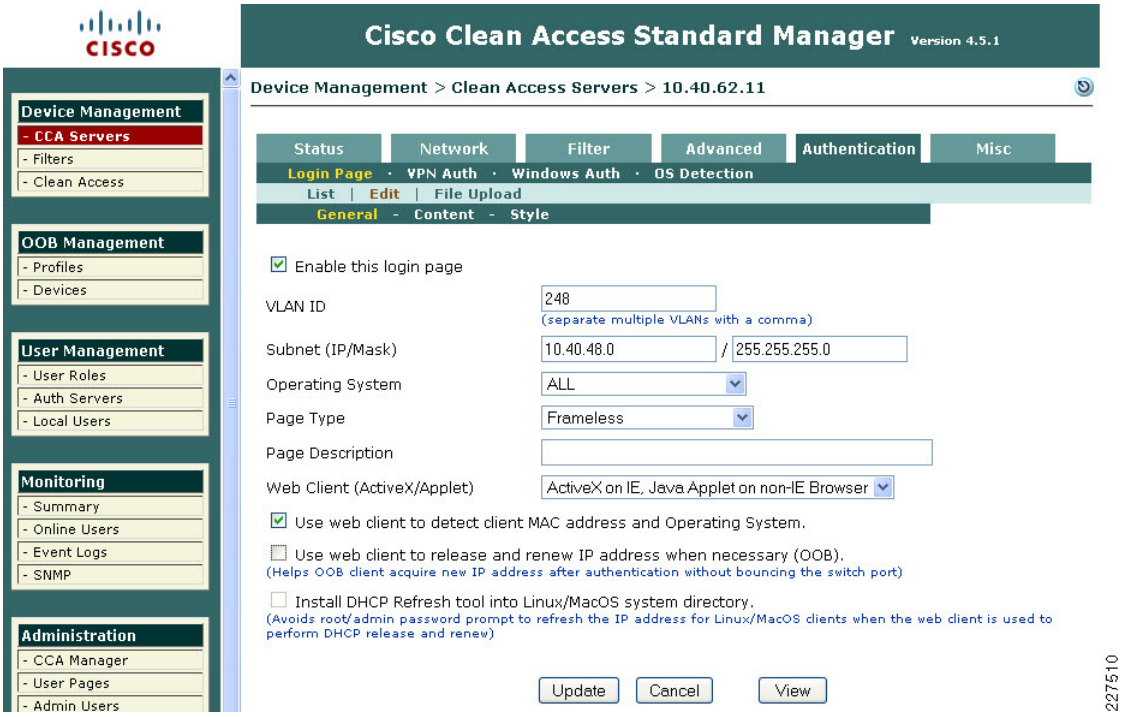

The other Authentication dialog of interest in this example is the Windows Auth dialog, as Windows Single Sign On (SSO) is used in this example. To perform Windows SSO the CAS needs to be able to communicate with Active Directory to determine the authentication state of the windows user. If Active Directory confirms that the user has authenticated to Active Directory the user doesn't need to perform additional authentication to the CAS. An example of this configuration is shown in Figure 5-57. There are a number of steps required to configure Active Directory SSO, as these are described in the *Cisco NAC Appliance —Clean Access Server Installation and Configuration Guide*. The key components in this configuration are:

- The creation of a Active Directory client account for the CAS
- Using the KTPass Application on Active Directory to convert the account encryption to DES encryption

Figure 5-57 CAS Windows Authentication

Clean Access Roles

The unauthenticated role is common to all clients, but once the client has been authenticated a different role may be applied based upon the identity of the client, different roles may be assigned for admin staff, users, etc.

User roles allow you to aggregate various policies into a user role. These policies include:

- Traffic policies
- Bandwidth policies

Note If bandwidth policies are to be enforced by the Clean Access Server it must be operating in band.

- VLAN ID retagging
- Clean Access network port scanning plugins
- Clean Access Agent/Cisco NAC Web Agent client system requirements

For example, the Admin and User roles could each have different traffic polices and VLANs, in addition the User role may enforce bandwidth policies by keeping the user traffic in band.

For more information on roles, refer to the *Cisco NAC Appliance - Clean Access Manager Installation and Configuration Guide* at the following URL:

http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/45/cam/45cam-book.html

Layer 2 OOB Example

Figure 5-58 shows an example of a Layer-2 OOB deployment. Where a wired client connected to an access switch is originally on the untrusted VLAN 264 and is switched to a trusted VLAN 64 once it has completed the NAC functions. The first NAC function is the authentication and authorization function, and this is the first design decision in implementing the NAC solution. That is, how will authentication and authorization be achieved, and what will the user experience be.

This example is focused upon the virtual gateway example, as virtual gateway provides the simplest deployment. In the virtual gateway example the original IP addressing, interfaces, and VLANs are maintained, and normal traffic flows are maintained. The only changes are the addition of the untrusted VLANs that carry client traffic during the NAC Authentication and Authorization, Scanning and evaluation, remediation, and quarantine modes.

Figure 5-58 Layer 2 OOB Example

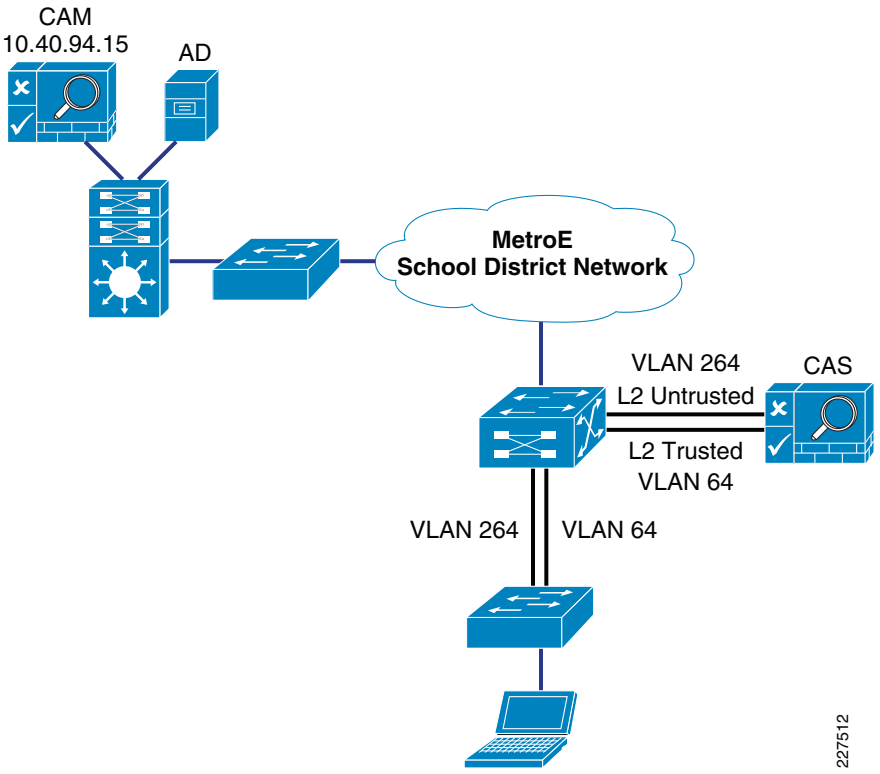

NAC Authentication Options

The authentication option in the NAC solution can be broadly categorized as NAC Authentication or NAC Single Sign On

- *NAC Authentication*—NAC authentication gives the NAC system the role of authenticating users, a user database, either local to the NAC system or a separate system such as RADIUS, or LDAP
- *NAC Single Sign On*—NAC SSO, addresses systems that already perform authentication as part of their normal operation. For example 802.1X, VPN access, or Active Directory. NAC SSO learns the authentication state of clients through RADIUS accounting, or Active Director and therefore doesn't require the user to reenter authentication.

Topology Considerations

The Layer-2 OOB solution relies upon there being a Layer-2 network connection available between the client devices and the Cisco CAS. In Figure 5-58 a trunk connects the access switch to the core/distribution switch. The Cisco CAS is connected to the core/distribution switch through two

interfaces—trusted and untrusted. In such a simple network it is relatively easy to provide a Layer-2 connection between the client and the Cisco CAS, for larger networks Layer-3 OOB may be a better choice.

The roles of the untrusted and trusted interfaces:

- *Untrusted Interface*—The untrusted interface connects the client to the to the Cisco CAS during the NAC Authentication and Authorization, Scanning and Evaluation, Remediation, and Quarantine modes
- *Trusted Interface*—The trusted interface connects the NAC CAS to the "normal" network interface. This makes a network connection available to the CAS while it is sitting between the client and the network, thus allowing client access to services such as DHCP and DNS -and user defined services. Once a client has successfully completed its authentication and scanning phases, the CAM uses SNMP to change the client VLAN, on the access switch, from the untrusted VLAN to the trusted VLAN. Thus providing a direct connection to the network that was on the other side of the CAS (the trusted network).

Availability Considerations

Both the CAS and CAM are both highly involved in client network access, and consideration must be given to the impact on clients if either a CAS or CAM should fail or need to be taken out of service for a period of time.

The CAS is inline with client devices during the authentication, authorization, and posture assessment phases of NAC, and if "In Band NAC" is being used it may be inline at all times. A CAS outage in an OOB deployment would not impact already connected clients but would prevent network access for new clients. A CAS outage for "In Line" clients prevents access for all clients.

In situations where availability of a CAS is critical an HA CAS solution may be implemented where a pair of CAS servers are installed using a primary CAS, and a secondary in hot standby. For more information refer to the *Cisco NAC Appliance - Clean Access Server Installation and Configuration Guide*.

The CAM is also a critical part of the authentication, authorization, and posture assessment phases of NAC. Although the CAM doesn't pass client traffic, the impact of its availability needs to be considered in the network design as well. Like the CAS the CAM has a HA solution that provides for a primary server and a hot standby secondary server. In addition, each CAS may be configured with a "fallback" option (as shown in Figure 5-59) that defines how it will manage client traffic in a situation where the CAM is unavailable.

In both HA CAM, and HA CAS, HA licenses are use that address the HA role of the server.

The use of the HA features will be dependent upon the company unique requirements, but CAS fallback should always configured to ensure that critical network services are available in even of a network outage.

Figure 5-59 CAS Fallback

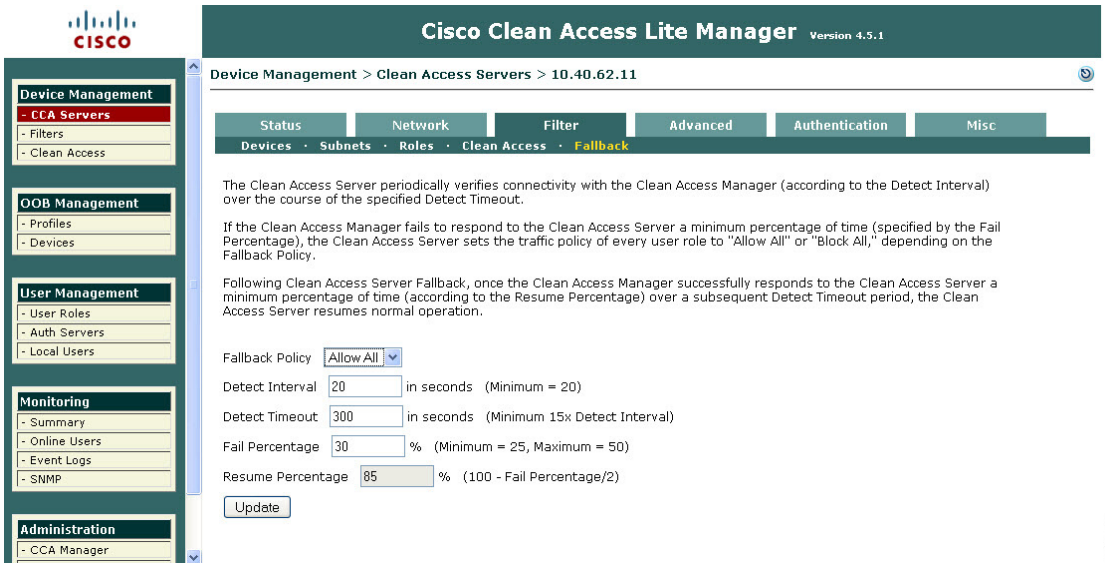

Basic Clean Access Switch Configuration

For OOB-based Clean Access some simple configuration must be performed on the switches implementing NAC This configuration is primarily to enable SNMP communication between the switches and the CAM. The table below shows a simple SNMP v1 configuration (SNMPv2c and SNMPv3 are supported).

In addition to the switch SNMP configuration, the required trusted and untrusted VLANs must exist and be operational on the switch, as illustrated in the configuration shown in Table 4. If a switch has more than one IP address the **snmp-server** source interface must be specified, as the CAM must be configured with the source IP address that OOB SNMP messages will originate from, alternatively all IP addresses of interfaces on the switch can be added to the CAM. If SNMP access filtering is applied on the switch (as recommended as a best practice) the CAM must be added as a trusted address.

Basic Clean Access Out-of-Band Switch Configuration

Switch Port Configuration	Global Switch Configuration

`snmp trap mac-notification change added`	`snmp-server enable traps mac-notification`
	` snmp-server enable traps snmp linkup linkdown`
	` mac-address-table aging-time 3600`
	` snmp-server host 172.16.1.61 traps version 1 cam_v1 udp-port 162 mac-notification snmp`

NAC CAS Connection

The NAC CAS connects to the core/distribution switch using two switch ports (which are not configured in EtherChannel). One is the untrusted port that serves the VLANs used by the clients prior to their authentication and authorization. The second is the trusted port that connects to the VLANs used once

clients have successfully completed the NAC process. Both trusted and untrusted ports are required, even if OOB NAC is used, as the CAS requires access to the trusted VLANs during the NAC process. The following is an example NAC CAS connection configuration.

```
interface GigabitEthernet1/0/4
description NAC Trusted Eth0
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 48,57,62
switchport mode trunk
spanning-tree portfast trunk
!
interface GigabitEthernet1/0/8
description NAC Untrusted Eth1
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 61,248,257
switchport mode trunk
spanning-tree portfast trunk
```

Basic 802.1X Switch Configuration

The basic 802.1X configuration controls access to an access VLAN depending upon the success or failure of the 802.1X authentication. If the 802.1X authentication is successful, there are three basic options:

- Access to the VLAN configured on the switch port
- Access to the VLAN configured on the switch port an controlled by a access list downloaded from the AAA server
- Access to a VLAN passed to the switch by the AAA server

The following is an example 802.1X configurations.

Example 3750 802.1X PC Port Configuration	Example 3750 Global Configuration
authentication port-control auto authentication periodic dot1x pae authenticator	aaa new-model aaa authentication dot1x default group radius dot1x system-auth-control ip radius source-interface Vlan300 radius-server host 10.40.62.9 auth-port 1812 acct-port 1813 key *cisco* radius-server host 10.40.94.9 auth-port 1812 acct-port 1813 key *cisco*

For more information about the 3750 802.1X configuration, refer to the following documents:

- *Catalyst 3750-E and 3560-E Switch Software Configuration Guide*, 12.2(50)SE ->Configuring IEEE 802.1x Port-Based Authentication

  http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_50_se/configuration/guide/sw8021x.html

- *Catalyst 2960 Switch Software Configuration Guide*, Rel. 12.2(50)SE Configuring IEEE 802.1x Port-Based Authentication

  http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/sw8021x.htm

CHAPTER 6

Small Enterprise Design Profile (SEDP)—Context-Aware Services

This chapter focuses on the application of general design best practices for Cisco Context-Aware Services (CAS) and the Cisco Mobility Services Engine (MSE) as they relate to small enterprise designs. Note that while this chapter attempts to be comprehensive, it is not intended to be a standalone technical guide on Cisco CAS, RFID technology, or the Cisco MSE. For comprehensive configuration and deployment information, the reader should refer to the in-depth configuration and deployment guides mentioned throughout this chapter.

Introduction

What Are Context-Aware Services?

Context-Aware Services provides the ability to dynamically capture and use contextual information about assets to optimize existing communications flows and organizational processes or facilitate the establishment of new ones. Contextual information can be collected for assets involved in almost any activity or process and this includes not just network endpoint devices (such as laptops and VoIP Phones) and products (such as video projectors), but in some cases also the users that are associated with such devices.

In environments where the Cisco Unified Wireless Network has been deployed, Context-Aware Services makes use of embedded 802.11 wireless network interface adapters (radios) in wireless client devices to accumulate contextual information about those assets or the user associated with the asset. For example, the location of a wireless laptop can be calculated via several different approaches or the user name associated with the laptop's user may be collected.

For assets that do not possess embedded wireless interfaces, external active Radio Frequency Identification (RFID) tags and sensors can be used to provide location input and monitor ambient environmental characteristics. Sensor capabilities can be directly embedded into active RFID tags in order to link the data captured (for instance, whether the asset is currently in motion) with the location of the asset. The algorithms used to determine location vary depending on the Radio Frequency (RF) environment and the accuracy required for a specific application.

In some cases, it may be necessary to track an asset with a high degree of accuracy throughout a small enterprise, such as when it is necessary to determine where a missing valuable asset is currently located). On the other hand, some applications using context-aware services may only require general indication of whether an asset is in or out of a permissible zone (such as the confines of an equipment storage area, for example).

Context-Aware Services can also provide location and other contextual information for wired devices attached to Cisco Catalyst LAN switches, such as the 2960G, 3560E, 3750E, 3750G, 4500, and 4900 series.

Note For the most current information regarding whether any specific Cisco Catalyst switch supports context-aware services, please refer to your Cisco Catalyst switch documentation.

Catalyst switches that are context-aware can provide civic location details for wired devices to the Cisco Mobility Services Engine based on pre-configured information specified for each switch port. This information can then be presented to users in a tabular format combined with other contextual information such as user name, device serial number and emergency location identifier numbers (ELINs).

Note A civic location specifies the civic address and postal information for a physical location using fields such as the number, street or road name, community, and county assigned to residential, commercial, institutional, and industrial buildings (e.g., 31 Main Street, AnyTown, AnyState 90004). An emergency location identifier number (ELIN) is a number that can be used by the local public safety answering point (PSAP) to look up the geographic location of the caller in a master database known as the automatic location information (ALI) database. The ELIN also allows the PSAP to call back the emergency caller directly in the event the phone call is disconnected.

For a more detailed overview of Context-Aware Services, refer to the following URL: http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns788/solution_overview_c22-475173.html.

Why Use Context-Aware Services?

The information that can be provided by Context-Aware Services across its application API can generally be classified into five functional categories, as shown in Figure 6-1.

Figure 6-1 Five Functional Categories of Context-Aware Services

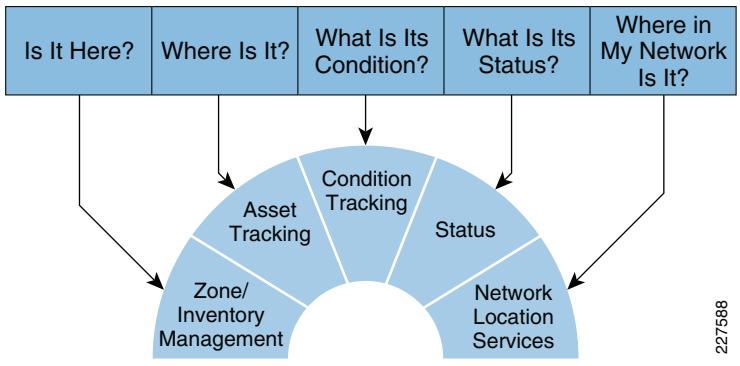

Is It Here?—Zone or Inventory Management

Zone or inventory management applications that utilize Cisco Context-Aware Services can define specific zones in which they monitor mobile assets that possess embedded wireless interfaces or have been outfitted with Cisco Compatible Extensions compliant RFID tags. These devices can be tracked and monitored when they enter and exit both permissible and non-permissible areas. Notifications can be generated when monitored assets stray into areas where they are not allowed. Examples of how zone or inventory management may be used in the small enterprise environment can include the following:

- Triggering the issuance of notifications to administrators or security personnel when monitored assets are moved out of authorized areas
- Alerting appropriate parties when persons equipped with RFID-enabled ID badges enter unauthorized areas
- Providing indication of proper staffing levels in critical environments. The number of doctors or nurses currently located within a hospital emergency room or the confines of an intensive care unit (ICU) can be carefully monitored. If the level of staff drops unacceptably, an alert can be issued to the appropriate recipients so that minimum staffing requirements can be maintained.

Further details about zone or inventory management can be found at the following URL: http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns788/solution_overview_c22-475178.html.

Where Is It?—Asset Tracking

Asset tracking applications that incorporate Context-Aware Services can help locate assets within the enterprise, whether they are connected to wired or wireless infrastructure. In this way, Cisco Context-Aware Services can provide great value to administrators, staff, security personnel, or anyone who must quickly and effectively locate and recover missing assets. Examples of how asset tracking may be used in the small enterprise environment include:

- Locating wired and wireless portable assets (such as a portable video projector or flat panel display) for meetings, customer presentations or other activities
- Quickly locating personnel possessing wireless VoWLAN phones, RFID-enabled badges or other trackable devices in both emergency and non-emergency situations
- Identifying the past pattern of movement associated with an asset by reviewing its past location history, in both tabular and graphical formats. Such audit trails can be especially useful when incorporated into security applications that can combine this information with other information sources, such as video surveillance.

Further detail regarding asset tracking can be found at the following URL: http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns788/solution_overview_c22-475177.html.

What Is Its Condition?—Condition Tracking

Condition tracking applications utilizing Context-Aware Services can monitor select characteristics of an asset's internal or external environment, such as variations in temperature, humidity, pressure, quantity, fluid volume, etc. Any change in these characteristics beyond set thresholds can trigger alerts, notifications, or other application-dependent actions. Examples of how condition tracking may be used in the small enterprise environment include:

- Temperature and humidity telemetry passed by RFID tag external sensors can be utilized in food service applications to monitor the temperature of food refrigeration units, guarding against costly spoilage and premature replacement.
- External fluid level sensors placed in combination with compatible RFID tags can be used to monitor critical fluid levels in building maintenance applications, such as fuel oil levels for emergency power generators and remote fuel storage for building heating.
- Indication of excessive or insufficient pressure in heating and cooling and other critical systems, which can be passed as telemetry data using properly equipped RFID tag sensors.

What Is Its Status?—Status Monitoring

Applications that monitor changes in user and asset status can use Context-Aware Services to detect status transitions, such as a change from a normal state to one indicating that an extraordinary event has occurred. Examples of how this may apply to the small enterprise environment include:

- RFID tags with user signalling buttons could be used to covertly pass indication of situations where assistance is needed, along with the location of the tag at the time of activation.
- Attempted asset tampering, such as the removal of RFID asset tags themselves or jostling of any type, can trigger status monitoring applications to generate alerts.
- The introduction of new assets into any small enterprise site, or any changes in the motion status of existing assets above a certain threshold, can serve as preliminary indication to building energy management systems regarding potential building environmental modifications. Such modifications may include changes to lighting, zone heating, or zone cooling settings for optimal efficiency.

Where Is It in My Network?—Network Location Services

Network location applications interfacing with the Context-Aware Services can help optimize the distribution of both wired and wireless network resources, reduce troubleshooting time, help eliminate waste due to use of network resources by unauthorized “rogue” devices, and help lower the overall total cost of network operation. Examples of how this may apply to the small enterprise environment include:

- Location and removal of unauthorized 802.11 wireless devices operating within a site, which can help reduce the enterprise’s exposure to various risks, such as maliciously introduced software and illegal file sharing.
- Ongoing tuning of the wireless network by identifying areas where wireless users routinely congregate. This may prompt adjustments in wireless network parameters, the number of access points deployed, or the placement of access points.

Further detail regarding network location services can be found at the following URL: http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns788/solution_overview_c02-474514.html.

When combined with other applications constructed to take advantage of the Cisco Context-Aware Services API, Context-Aware Services can serve as an enabler for entirely new application functionality. For example, a context-aware application might provide the ability for an employee to locate other employee team members with specialized skills (such as an employee who is also a certified emergency medical technician or emergency response team member). Once such specialists are located, contact can be initiated with the nearest such specialist that is available to deliver assistance. Context-aware services enhance the experience of users while at the same time improving their overall efficiency and productivity.

Cisco Context-Aware Components

The components of the Cisco Context-Aware Services are shown in Figure 6-2.

Figure 6-2 Cisco Context-Aware Services Components

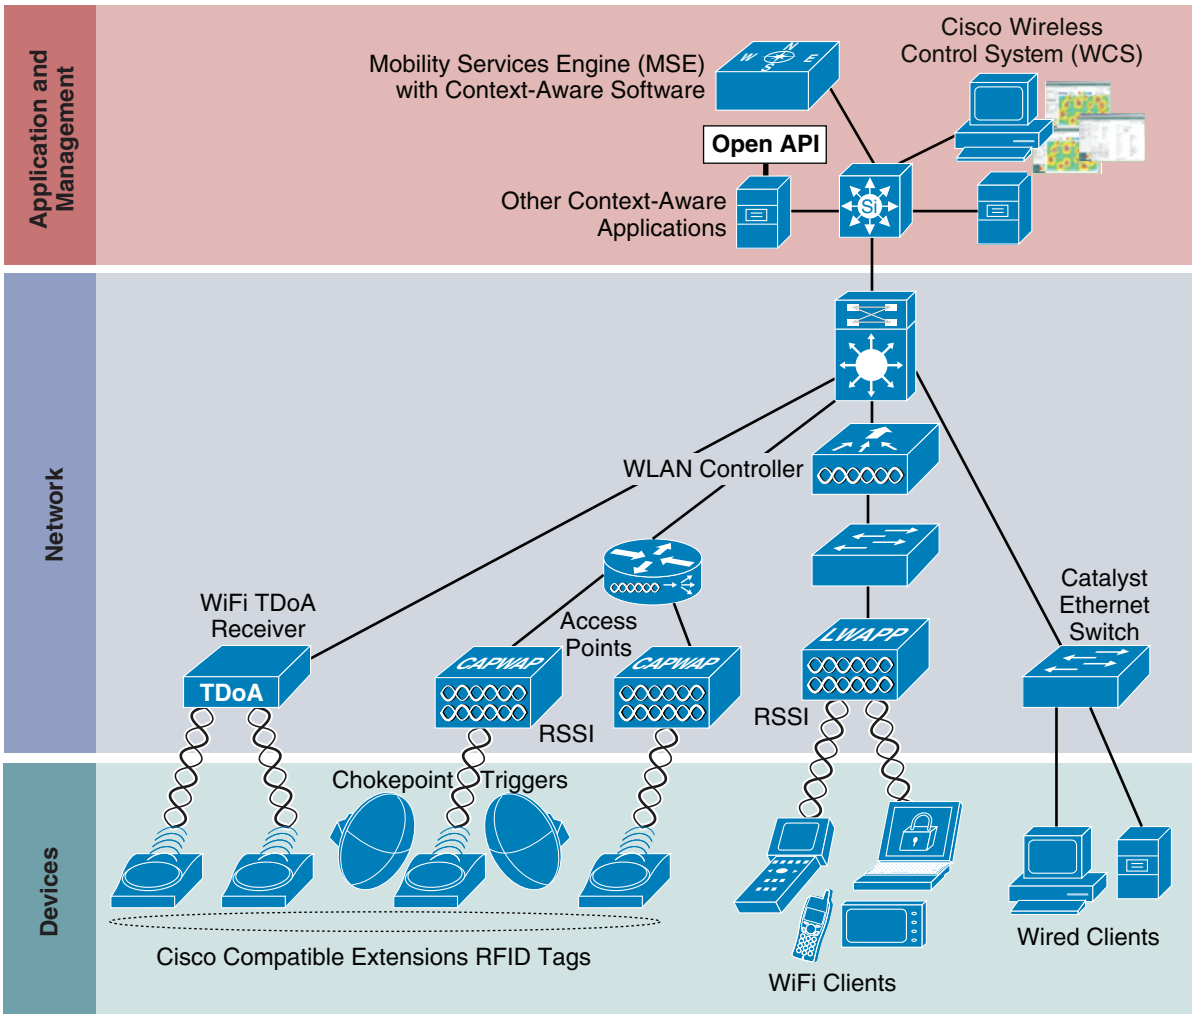

Wired and Wireless Client Devices

- Wired or wireless (Wi-Fi) devices—Mobile wireless devices (asset tags, WiFi equipped computers, mobile stations, etc.) that interact with the network and whose location and other contextual parameters can be monitored by Context-Aware Services. Wired devices are generally equipped with an Ethernet interface which is attached to a Cisco Catalyst switch that supports context-aware services. In addition, some devices may possess both wired and wireless interfaces. Without context-aware services deployed, a wired Cisco IP phone originally deployed in a conference room that is subsequently moved to an adjacent conference room across the hall, would require any location information associated with it to be manually updated. With Context-Aware Services and a context-aware Cisco Catalyst switch, the location of the Cisco IP can be dynamically updated to

reflect its new location within seconds after it has been plugged into its new location. Via the Mobility Services Engine's API, this information can be provided to various context-aware applications, including the Cisco Wireless Control System (WCS).

- Cisco Compatible Extensions RFID Tags—These RFID tags can be physically attached to assets (regardless of whether the asset itself contains a wired or wireless network interface adapter) and can pass some contextual information on behalf of the asset to Cisco Context-Aware Services. Compliance with the Cisco Compatible Extensions program helps ensure that RFID tags adhere to predefined transmission formats such that the contextual information they capture can be made readily available to other Context-Aware Services components, including safety and security applications from Cisco partners. Sensor capabilities can be externally mounted or directly embedded into tags in order to link the data captured (for instance, motion, or temperature data) with the location of the mobile asset. Externally mounted sensors can also be placed in fixed locations, like a refrigerator or a storage room.

  For more information about the Cisco Compatible Extensions for RFID Tags program, refer to the following URL: http://www.cisco.com/web/partners/pr46/pr147/ccx_wifi_tags.html.

- Chokepoint triggers—Chokepoint triggers (sometimes referred to as Exciters) are an optional component that can greatly enhance asset tag functionality, providing for finer tag location granularity and improved accuracy by localizing tagged assets within multi-floor structures, in presence detection scenarios, or when passing through chokepoint areas, such as entrances and exits. RFID asset tags that enter the proximity of a chokepoint trigger can change their normal behavior based on a set of pre-programmed instructions. RFID tags that have been stimulated by chokepoint triggers in this fashion send notifications and contextual information via the wireless infrastructure to the Cisco Mobility Services Engine.

Cisco Unified Network

This multipurpose network contains the wired and wireless infrastructure required to address converged data, voice, and video requirements, as well as providing the foundation for use of context-aware services.

- *Context-Aware Cisco Catalyst switches*—Such as the 2960G, 3560E, 3750E, 3750G, and 4500 series switches that support the specification of civic and emergency location identification number (ELIN) location information and the transmission of this information to Cisco Context-Aware Services. This functionality allows for contextual information associated with wired devices to be tracked using Context-Aware software on the Mobility Services Engine. Switches transmit relevant contextual information to the MSE for all of the devices attached to them. This information may include the physical mailing or street address location associated with the attached device (the civic address) as well as other information such as the IP address, MAC address, port, VLAN, and user name. Typically, this information is obtained using switch features such as IEEE 802.1x, Dynamic Host Configuration Protocol (DHCP) snooping, Dynamic Address Resolution Protocol (ARP) Inspection (DAI), and IP Source Guard. Additionally, if the end device runs the Cisco Discovery Protocol or Link Layer Discovery Protocol for Media Endpoint Devices (LLDP-MED), additional information, such as the version number and serial number, can also be sent to the MSE.

Note At this time, serial numbers of attached devices are reported to the MSE Context-Aware Service only if the device supports LLDP-MED.

- *WLAN controllers*—WLAN controllers (and the embedded software residing within them) provide for the aggregation and transfer of device tracking and statistics information for RFID tags, mobile wireless clients, and any rogue devices detected.

- Access points—In addition to their fundamental role of providing access for wireless clients, Cisco Aironet access points provide measurements of received signal strength from both wireless client devices and RFID tags and subsequently forward this information to the Mobility Services Engine via their registered WLAN controller.
- Received Signal Strength Indication (RSSI)—This is a mechanism used to determine device location by carefully considering the measured strength of a radio signal at several points in an indoor environment. Used by the Cisco Mobility Services Engine for WLAN clients, RFID tags, and rogue devices, this algorithm is based on the signal sent from the mobile asset to different access points deployed within the small enterprise site. RSSI is usually preferred for indoor or low ceiling environments, both of which can result in high degrees of signal reflection.
- Wi-Fi Time Difference of Arrival (TDoA) Receiver—Wi-Fi TDoA receivers are optional components used in very large, open environments to locate assets equipped with RFID tags with greater accuracy and precision than is possible using other techniques.

Note Although useful in extending Context-Aware Services for RFID tagged assets in outdoor venues, the use of TDoA receivers were not included in the SEDP Mobility design.

For a much more detailed explanation of both RSSI and other wireless location algorithms, refer to the *Wi-Fi Location-Based Services Design Guide 4.1* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/wifich2.html#wp1049520.

Management and Applications

- *Cisco Mobility Services Engine (MSE) with Context Aware Services Softwar*e—The Cisco Mobility Services platform can host multiple independent services possessing high-level capabilities that can enhance both wireless and wired network infrastructures. One of these is Cisco Context-Aware Services which can capture, store, and analyze contextual information from multiple wired and wireless networks simultaneously. When Context-Aware Services is deployed in accordance with generally accepted best practices, both wired and wireless network infrastructure devices (controllers and switches) may send raw location measurement data, device attachment, and other contextual information to the MSE regarding the presence of any wired clients, wireless clients, RFID tags, or rogue devices. Both wired and wireless network infrastructures communicate with the MSE using the Cisco Network Management Services Protocol (NMSP), which is a Cisco-defined protocol used for secure communication between the MSE and other context-aware network infrastructure components. The MSE sits out of the data path of the wireless LAN and receives data from WLAN controllers and context-aware switches via the use of NMSP.

  The location of WLAN clients and RFID tags on the Cisco Mobility Service Engine is calculated by one of two software service modules:

  - Cisco Context-Aware Engine for Clients, which handles all context-aware operations involving RSSI location of Wi-Fi clients, rogue clients, and rogue access points. This engine also handles context-aware operations for wired clients.
  - Cisco Context-Aware Engine for Tags, which handles all context-aware operations involving TDoA and RSSI location of Cisco Compatible Extensions compliant RFID tags.

  Context-Aware Services software, when operating alone on the Cisco MSE, is capable of servicing up to a maximum of 18,000 simultaneously tracked devices per single MSE-3350 appliance and 2,000 simultaneously tracked devices per single MSE-3310 appliance.

Refer to the following data sheet for more information regarding the Cisco 3300 Series Mobility Services Engines
http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c78-475378.html

- *Wireless Control System (WCS)*—The Cisco Wireless Control System is a management platform that also contains a context-aware client application that interacts with the Mobility Services Engine. The primary role of the context-aware client application is to provide access to the contextual information contained on the MSE using the MSE's application programming interface (API). The application can then either present this information to the user directly (such as is seen in a graphical location map or a table of location values) or enable other processes to accomplish relevant tasks using this information that would otherwise be difficult to achieve. The Cisco WCS can also serve in a special secondary role as a control client that possesses the ability to configure MSE operational parameters.

  For more detailed information on the Cisco Wireless Control System management server and its capabilities, including its ability to serve as a context-aware application client, refer to the documentation at the following URL:

  http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html.

- *Other Context-Aware Applications*—Other context-aware client applications from third party Cisco Technology Development Partners may also access the MSE via its open API, which is based on Simple Object Access Protocol (SOAP) and XML protocol. Access to this API is available to any Cisco technology partner. Context-aware applications developed by Cisco Partners often deliver specifically targeted application functionality that is often not available from other sources.

  For more information on the Cisco Context-Aware Services API, refer to the following URL: http://developer.cisco.com/web/contextaware/home.

Context-Aware Component Interaction

Figure 6-3 provides a more detailed illustration of the protocol interaction between the various Context-Aware Service components.

Figure 6-3 Protocol Interaction Between Context-Aware Components

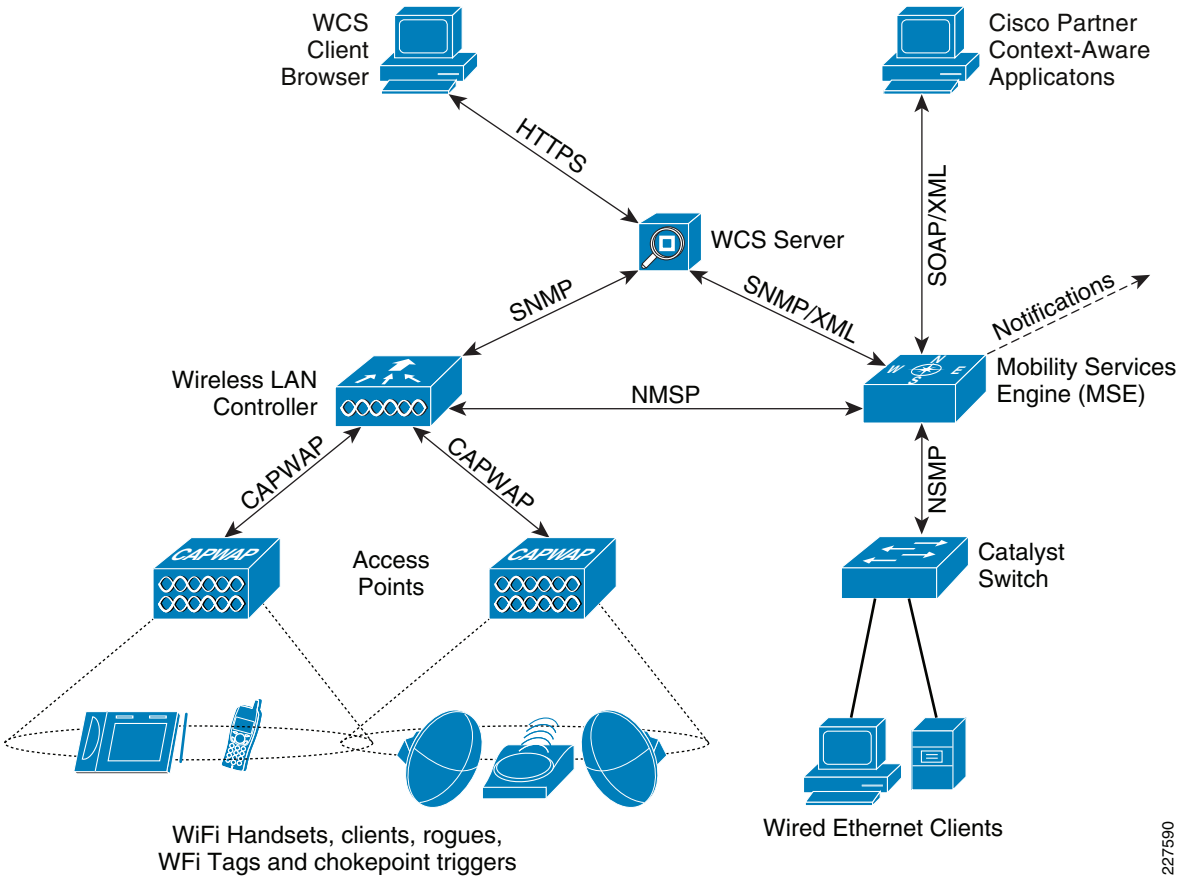

For wireless clients, tags, and rogues, Cisco Aironet access points use the Control and Provisioning of Wireless Access Points (CAPWAP) protocol to forward the RSSI of detected clients, tags, and rogues to the WLAN controller to which they are currently registered. The wireless LAN controller aggregates this information on a per device basis from all registered access points detecting the wireless device's signal. This information is then forwarded to the MSE using the NMSP protocol via an authenticated and encrypted session. The appropriate Context Aware software engine on the MSE then uses the RSSI data received from one or more WLAN controllers to determine the location of the wireless device

Note A rogue access point is any access point that is determined not to be a member of the same mobility group as the WLAN controller to which the detecting access points belong. A rogue client is any client that is currently associated to a rogue access point.

For wireless clients, rogue access points, and rogue clients, the Context-Aware Engine for Clients is used to process the received RSSI information. For operations involving RFID tags, however, the Context-Aware Engine for Tags is used instead. Using Context-Aware Services and the Mobility Services Engine, the Cisco Unified Network can readily detect 802.11 Wi-Fi active RFID tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification (such as those from AeroScout, WhereNet, G2 Microsystems, and others). Through the MSE and WCS, the location of these RFID tags can then be displayed on WCS floor maps using a yellow tag icon.

Note The Context-Aware Engine for Tags in Cisco Context-Aware Services Release 6.0 only tracks RFID tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification.

Cisco Compatible Extensions compliant active RFID tags are detected on a Wi-Fi network based on periodic frames[1] that are sent by the tag using a Layer-2 multicast. The delay between these periodic frames can be programmed based on the specific application use case. In most cases, tags are configured to transmit periodic frames every three to five minutes in order to strike an equitable balance between location accuracy and good tag battery life.

RFID tags can also pass tag telemetry information upstream to the MSE as part of the tag message payload. This contextual information (battery status, motion, temperature, pressure, humidity, etc.) is received by access points and collected by WLAN controllers in a similar fashion as that described earlier in this section. WLAN controllers will aggregate telemetry traffic from multiple tags and eliminate any duplicate tag telemetry values that might be received. After the telemetry has been distilled and cleansed of any duplicate information, the WLAN controller passes it to the MSE. The MSE updates its internal databases with this information and in turn makes this information available to application programs.

Properly equipped RFID tags can also indicate the occurrence of a priority event, such as one that might result from the triggering of a tag tamper sensor or the depression of a tag call button. RFID tags indicate that these types of events have occurred via additional information embedded in the tag messages that are sent to the WLAN controller. This information is in turn passed northbound from the WLAN controller to the MSE and the MSE can make this information available to application programs.

Several manufacturer's RFID tags include a secondary on-board magnetic signaling receiver, typically set up to respond to the magnetic field component of a 125 kHz RF carrier. This secondary receiver provides for additional tag functionality when tags enter into areas that are within close proximity to a magnetic signaling transmitter or chokepoint trigger. Chokepoint triggers are proximity communication devices that trigger asset tags to alter their configuration or behavior when the tag enters the chokepoint trigger's area of operation or stimulation zone. This alteration could be as simple as causing the asset tag to transmit its unique MAC address identifier. It could be significantly more complex, including causing the tag to change its internal configuration and status. One of the prime functions of a chokepoint trigger is to stimulate the asset tag such that it provides indication to the system that tag has entered (or exited) the confines of a constricted physical area known as a chokepoint. Typical chokepoints include entrances, exits or other types of physical constrictions that provide passage between connected regions of a facility (such as a corridor or hallway).

Note While chokepoint triggers are electronic devices that are typically deployed within physical chokepoints, it is not unusual to hear the term “chokepoint” used rather loosely to refer to both the physically constricted area and the associated electronic device.

Chokepoint triggers (including a very popular model manufactured by AeroScout Ltd. known as an Exciter) may be connected to the wired infrastructure and are configured using each vendor's configuration software. Once they have been configured, they can either remain connected to the infrastructure full-time for management purposes or can be disconnected and operate in a standalone mode, requiring a source of electrical power but no actual connectivity to the network.

1. These periodic frames are also sometimes referred to as “beacons”. They should not be confused with the 802.11 beacons that are sent by access points.

The chokepoint trigger address information contained in the tag packet provides the MSE with enough information to temporarily override any RSSI or TDoA localization currently in place for the tag and set the current location of the RFID tag to the location of the chokepoint trigger. The size of a chokepoint trigger's stimulation zone, or range, can extend from a radius one foot or less to over twenty feet, dependent upon the vendor and the capabilities of the particular model.

Catalyst switches supporting context-aware services also make use of NMSP to interact with the MSE similar to the manner described earlier for WLAN controllers. A major difference between how WLAN controllers and Catalyst switches interact with context-aware services lies with the method used to determine the location of switch attached wired devices. As explained earlier, localization of wireless devices is performed by the MSE using a signal or time based technique, whereas the location of wired devices is based on information sent to the MSE that originates in the switch configuration. The information recorded by the MSE for the wired device includes the device MAC address, switch MAC address, slot or port, IP address, and user name (if available). This information is sent whenever a device link changes state. Context-aware Cisco Catalyst switches provide the MSE with the latest relevant civic location and emergency location identification number (ELIN) information for all attached IP endpoints. These endpoints may include IP phones, PCs, access points and other devices.

In Release 6.0, all civic and ELIN location information is configured locally at the switch, and shortly after location changes are made using the CLI, they are propagated to the MSE. NMSP is used between the switches and the MSE to maintain synchronization, and alert the switch as to the connection or disconnection of devices.

Additional information regarding civic address location is available from the IETF in the following RFCs:

- DHCP Option for Civic Addresses Configuration Information
  http://www.rfc-editor.org/rfc/rfc4776.txt
- Revised Civic Location Format for Presence Information Data Format Location Object
  http://www.rfc-editor.org/rfc/rfc5139.txt

Note Proper validation of certificates between context-aware service components requires the participants to possess sane clocks (clocks whose configured time does not differ from one another by large amounts). In order to facilitate this, it is highly recommended that the clocks in the MSE, WCS, WLAN Controllers, and any context-aware Cisco Catalyst switches be synchronized to a common time base using the Network Time Protocol (NTP). The lack of clock sanity amongst context-aware components in the network can cause NMSP sessions to fail if the configured date and time fall outside of the certificate validity period and cause certificate validation to fail.

Network Mobility Services Protocol (NMSP)

The Network Management Service Protocol (NMSP) was designed to define intercommunication between Mobility Service Engines and network access controllers over a switched or routed IP network. An access controller can provide network access for either wired or wireless endpoints. Within the scope of the small enterprise design discussed in this document, access controllers are represented by WLAN controllers and context-aware Cisco Catalyst switches.

NMSP is a two-way protocol that can be run over a connection-oriented or a connectionless transport. WLAN controllers and context-aware switches can use NMSP to communicate with one or more MSEs. NMSP is based upon a bidirectional system of requests and responses between the MSE and the access controllers.

MSP also provides for a keep alive feature that allows either partner in a NMSP session to determine if the adjacent partner is still active and responsive. Should an MSE fail, a WLAN controller or a context-aware Cisco Catalyst switch will try to contact another MSE with which to communicate. If the WLAN controller or context-aware Cisco Catalyst switch fails, all context-aware services being provided to that WLAN controller or context-aware Cisco Catalyst switch are disabled until that WLAN controller or context-aware Cisco Catalyst switch once again becomes active and re-establishes its NMSP session.

Note It is important to understand that the failure of an NMSP session has no direct impact on the ability of a WLAN controller or context-aware capable Cisco Catalyst switch to pass normal client session traffic to applications on the network. In other words, a failed NMSP session to a WLAN controller may affect the ability of the MSE to provide updated contextual information on that controller and its resources. But it does not affect the ability of the WLAN clients using that WLAN controller to logon to applications residing on the network. This also applies to wired clients and context-aware Cisco Catalyst switches.

NMSP uses Transport Layer Security (TLS) and TCP port 16113 on the WLAN controller or context-aware Cisco Catalyst switch. The MSE will initiate the connection to the WLAN controller or context-aware Cisco Catalyst switch, although once a secure session has been established between MSE and its session partner, messages may be initiated in either direction. The TCP port (16113) that the controller and mobility services engine communicate over must be open on any firewall that exists between the controller and mobility services engine.

The MSE and the WLAN controller or context-aware Cisco Catalyst switch use Echo Request and Echo Response control messages to maintain an active channel of communication so that the data messages can be sent. The Echo Request message is a keep alive mechanism that allows either NMSP session partner to determine if the other partner remains active and responsive. Echo Requests are sent periodically (upon expiration of a heartbeat timer) by the MSE or its session partner to determine the state of the NMSP session. When the Echo Request is sent, a NeighborDeadInterval timer is started. The NeighborDeadInterval timer specifies the minimum time a session partner must wait without having received Echo Responses to its Echo Requests, before the other session partner can be considered non-responsive and the NMSP session is placed in an idle state.

Context-Aware Services in Small Enterprise Designs

Figure 6-4 provides a high level illustration of the integration of Cisco Context-Aware Services into the Small Enterprise Design Profile. The key points of this integration into a Metropolitan Area Network deployment are:

- The presence of a centralized management entity (the wireless control system (WCS) at the main site. In the case of Context-Aware Services, WCS also serves as a context-aware application client. Main and remote small enterprise site context-aware users can log into WCS and query the contextual characteristics (such as location) associated with wireless and wired client devices, rogues and asset tags. In some cases, third-party context-aware application servers may also contain context-aware applications located in the main site.
- The presence of a per-site local Mobility Services Engine may be deployed to provide Context-Aware Services to larger sites, where the anticipated number of total tracked devices is significantly higher (e.g., greater than 500 and most likely 1000 or more). A locally deployed MSE may also be justified if context-aware services are being utilized for applications and tasks that are considered mission-critical to the function of the small enterprise or the safety and security of employees, guests and visitors.

- The option of a centralized Mobility Services Engine with Context Aware software at the main site used to provide the Context-Aware Services to smaller remote sites where the anticipated number of total tracked devices per small enterprise is relatively low (e.g., less than 500).

Figure 6-4 High-Level View of the SEDP With Hybrid Context-Aware Services

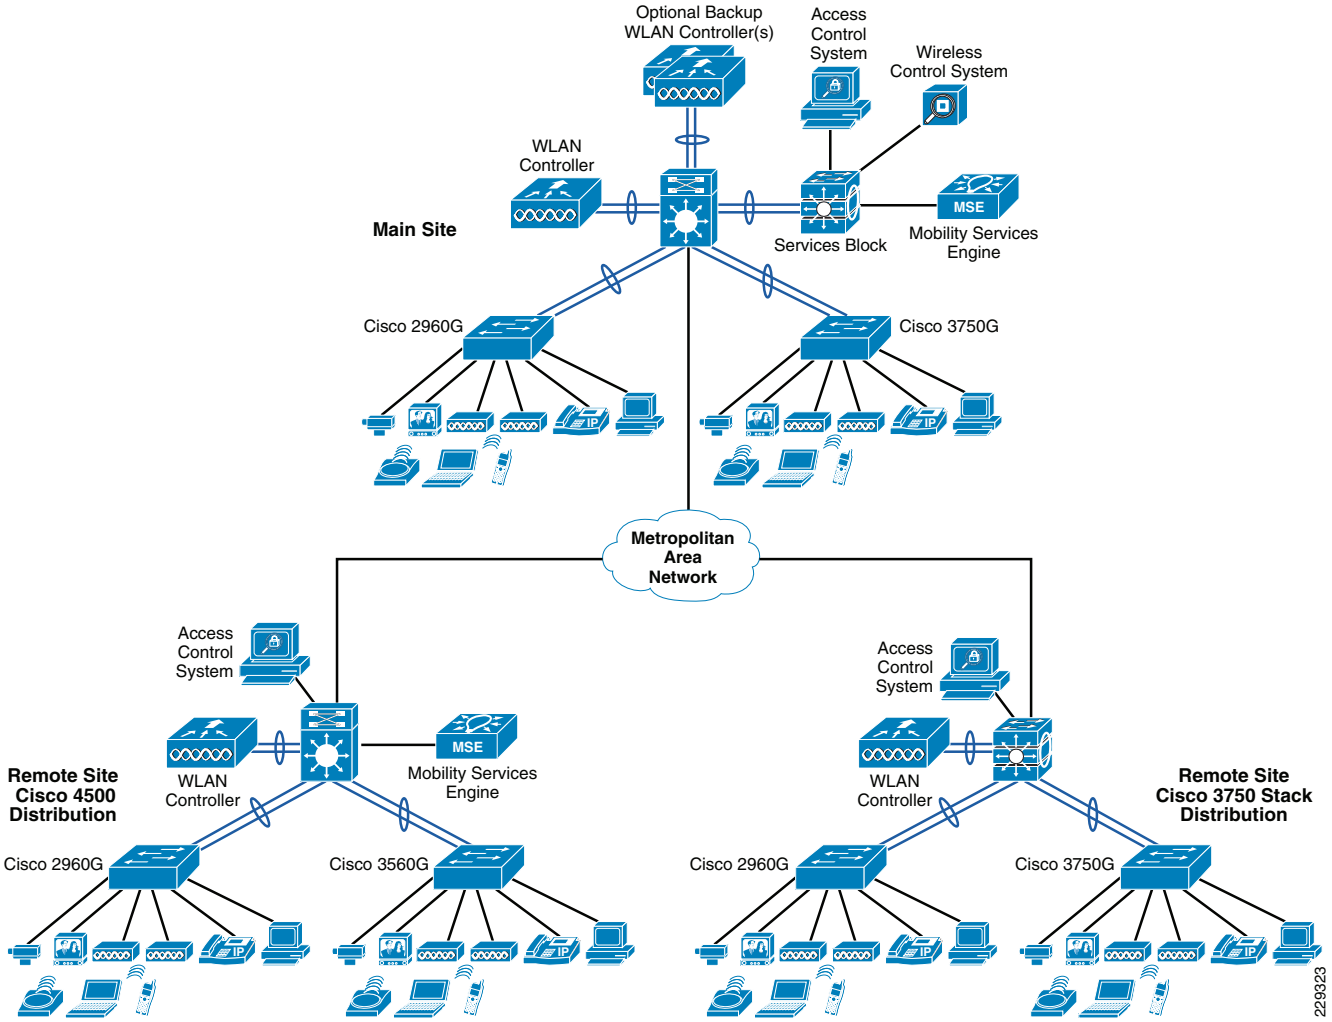

It is our understanding that the enterprises addressed by the Small Enterprise Design Profile will contain a mix of site sizes, with the main or headquarters site typically being larger than the remote sites. We anticipate that many of the smaller remote sites may be well served at the distribution layer by the 3750 switch stack, whereas the larger main site may be outfitted instead with the 4500E distribution switch. In situations such as this, the deployment model shown in Figure 6-4 can be used to provide context-aware services to both types of locations.

Historically, context-aware services and the Cisco Mobility Services Engine (as well as its predecessor, the Cisco Wireless Location Appliance) were designed to be deployed within modern switched LAN environments. In such deployments, FastEthernet speeds and capacity (or better) are typically assumed to be present throughout the local area network. Due to the lower speeds associated with traditional wide-area networking technologies (such as frame relay, T-1, and so on), context-aware components have not been recommended for deployment across traditional WANs. In fact, if context-aware services are to be deployed in remote sites possessing only traditional WAN connectivity, Cisco Systems has typically always recommended that the MSE be deployed locally along with WLAN controllers and any

other components establishing NMSP sessions to the MSE. In addition, designers may wish to break up large WCS network designs into smaller designs to avoid time outs that can occur between WCS and the MSE during synchronization of very large network designs when using low-speed links[2]. While the cost of local MSE deployment can be very applicable to very large remote sites whose device population can justify it, this is typically not true in the case of smaller sites with lower device populations.

A key advantage of the SEDP design is that the use of a modern high-speed metropolitan area network (MAN) to interconnect the main and remote sites offers far more bandwidth to each site than would be seen with a traditional WAN deployment. In this case, with modern LAN-like speeds available across the metropolitan area network, our approach begins to take on the look and feel of a local LAN deployment, and thus the idea of deploying an MSE remotely from the other context-aware components becomes much more feasible. Note however, this assumes sufficient bandwidth and infrastructure is in place to assure that FastEthernet-like speeds are available to each site and that proper network protocol identification, classification and QoS are all applied properly to manage congestion in the network.

In the context-aware services model shown in Figure 6-4, we take advantage of high-speed connectivity across the MAN allow for a centralized MSE to provide context-aware services for smaller remote sites in the enterprise. In this case, we assume that a small site might have a maximum of 250 to 500 simultaneously tracked devices. An exception clearly would be made for sites where context-aware services are used for mission critical applications, such as enterprise safety and security applications. An example of such a safety and security application might one that is used to ascertain the location of all employees and staff during an emergency event, such as a fire, flood or crime evacuation. This type of application obviously must be available at all times, including during any potential network outages, hence the use of a centralized MSE would not be an option. Any other supporting applications that are required for such mission-critical deployments of context-aware services should also be deployed locally in this case.

Note Based on our analysis of MAN capacity, traffic flows, classification and QoS, we believe the centralized deployment of an MSE across a modern high speed metropolitan network is a viable concept. Although a great deal of intensive functional testing was performed during the preparation of this chapter, time constraints did not allow us to complete validation of centralized MSE deployments across metropolitan area networks.

Larger enterprise sites using a 4500 series Catalyst switch for distribution are assumed to possess at least 500 or, more likely, 1000 or more simultaneously tracked wired or wireless devices. While it may be possible to service these larger sites using a centralized MSE, the larger number of clients and the increased amount of traffic placed onto the MAN between controllers, switches, and the MSE in this case can pose more of a challenge, especially when there are large device populations that move frequently and generate location updates on a regular basis. Careful analysis of data traffic and the judicious application of QoS in the network becomes especially important.

At the current time, the most reliable and robust solution in the case of sites with large tracked device populations is to deploy an MSE locally. Once again, in cases where context-aware services are regarded as mission critical for the small enterprise site, other context-aware components (such as third-party context aware application servers or in some cases the WCS as well) should also be deployed locally in order to ensure that the context-aware solution is functional even in the rare case of a prolonged MAN failure.

In small enterprises where there are many sites with either large tracked device populations or mission-critical context-aware applications, it is important to keep in mind that at this time Cisco officially supports the management of up to five (5) MSE platforms from a single WCS system. While this is not a “hard” limitation on the number of MSE platforms that can be assigned to a single WCS

2. Or use a locally deployed WCS at each remote site that could optionally be managed by WCS-Navigator at a central site.

system, it is the limit to which testing has been performed. Therefore, in designs where there may be greater than five sites equipped with locally deployed MSE platforms, you may wish to consider using additional WCS systems as necessary. In this case, the use of WCS-Navigator (not shown in Figure 6-4) should be considered at the main site to provide a single interface portal to as many as twenty (20) WCS management systems and their associated Mobility Services Engines. WCS-Navigator is a management aggregation platform that delivers enhanced scalability, manageability, and visibility of large-scale implementations of Cisco WCS and the Cisco Unified Network. WCS-Navigator provides straightforward access to information from multiple Cisco WCS management platforms. A single WCS-Navigator management aggregator can support up to twenty (20) WCS management systems and 30,000 access points.

Note Due to time constraints, scalability testing of WCS-Navigator in the Small Enterprise Design Profile beyond four WCS systems was not able to be completed. Further information on WCS-Navigator can be found at http://www.cisco.com/en/US/products/ps7305/index.html.

Component Capacities

Mobility Services Engine

Each Cisco Mobility Services Engine has a maximum device tracking capacity, defined as the maximum number of active wired and wireless (clients, rogues, and RFID tags) that can be tracked by a single Mobility Services Engine. This is a “hard” limit that is dictated by the licensing purchased for the Context-Aware software as well as the presence of any other applications on the MSE. Once a Mobility Services Engine has reached its maximum tracking capacity, any new devices that the MSE becomes aware of beyond that limit are simply not tracked. It is important to note that while this section discusses the maximum device tracking limits for the MSE, MSE licenses can be purchased supporting device limits significantly lower than those shown here. Refer to the *MSE Licensing and Ordering Guide* (http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html) for more information regarding the various combination of client and RFID tag tracking capacities available for the MSE.

For Release 6.0, the maximum device tracking limits when using only the Context-Aware Services software on the MSE are shown in the following table:

Mobility Service Engine	Maximum Tracked Device Capacity
MSE-3350	18,000
MSE-3310	2,000

If you intend to use the MSE-3350 to deliver other services in addition to Context-Aware, the maximum capacity shown above in Table 2 will likely be reduced. See the *MSE Licensing and Ordering Guide* (http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html) for information on Context-Aware maximum tracked device limits with other co-resident MSE services.

When working within these maximum capacities, it is important to note that further category-specific limits can be initiated at the designer's discretion via the MSE configuration. This allows, for example, a maximum capacity of 2,000 tracked devices on a MSE-3310 to be further limited as 1,000 wired and wireless clients, 500 RFID tags, and 500 rogue access points and clients. Partitioning the maximum tracking capacity of the context-aware software in this manner prevents any single device category from consuming more than its authorized share of the maximum tracking capacity of the system.

In the high-level diagram shown in Figure 6-4, the MSE-3310 might be a good design choice for a locally deployed MSE at our 4500-based main site. Its maximum tracked device capacity of 2000 devices should scale well in this type of application, where we are assuming an estimated 1000-1250 total tracked devices. This would leave ample MSE capacity in reserve for future growth at this location. Of course, you could deploy with less capacity in reserve should you choose to, and elect to address future tracked device growth at a later date via a hardware addition or upgrade. 4500-based sites that possess or anticipate near-term tracked device needs exceeding 2000 devices should consider the MSE-3350 instead.

Except for very small enterprises, the MSE-3350 would typically be the best overall choice when considering a centralized deployment using a high-speed metropolitan area network. In this way, it can provide context-aware services for several smaller sites, each of which might possess an estimated 500 or fewer tracked devices. In very small enterprises (e.g., enterprises containing up to a total of four small sites, for example) a centralized MSE-3310 may prove to be even more cost-effective.

WLAN Controllers

WLAN controllers also possess limitations on the maximum number of devices for which the controller will track and aggregate contextual information. In Release 6.0, these limits are shown in Table 6-1.

Table 6-1 Maximum Device Limitation

WLAN Controller	Clients	Tags	Rogue Access Points	Rogue Clients
4404	5000	2500	625	500
4402	2500	1250	625	500
2106	500	256	125	100

Note that these are indeed “hard” limits. In other words, once these limits have been achieved on a WLAN controller, contextual tracking information for any new clients, RFID tags, rogue access points, or rogue clients beyond these limits will be dropped until such time that older entries are pruned from the controller's internal database. The client and tag limits are quite high and should not be easily exceeded within most small enterprise sites.

The Context-Aware System Performance chapter of the *Mobility Services Engine Context Aware Deployment Guide* (http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#casysperf) also points out that a single MSE can support up to 500 total NMSP connections. Keep in mind this includes not only the NMSP sessions to WLAN controllers, but NMSP sessions to any context-aware Cisco Catalyst switches conducted from that MSE.

Note Although a single MSE can technically support up to 500 NMSP sessions, scalability testing constraints have only allowed for testing of 100 simulated NMSP connections to a single MSE at this time.

Wireless Control System (WCS)

With regard to the MSE, WCS interacts as a context-aware client and does not track devices itself when used in conjunction with the MSE. Thus, there are no direct constraints on the maximum number of tracked devices imposed by WCS itself.

In addition to established WCS sizing and capacity guidelines for the number of supported controllers and access points (listed in the *System Requirements section of the Wireless Control System Configuration Guide*, Release 6.0, http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0wst.html#wp1061082), there are a few indirect constraints relating to Context-Aware services that you should be aware of:

- To maintain clarity and the speed of its graphical user interface, WCS only displays the first 250 wireless clients, RFID tags, rogue clients, or access points on a single floor map. To view graphical location displays for any of these device categories beyond this limit, filtering (based on MAC address, asset name, asset group, asset category, or controller) should be used to limit the number of devices displayed at once (see the Floor Settings section of the *Wireless Control System Configuration Guide*, Release 6.0, http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0maps.html#wp1210969).
- In Release 6.0, a single WCS can manage Context-Aware Services on up to five (5) Mobility Service Engines. While defining more than five MSEs to a single WCS is possible, Cisco Systems has not validated this level of operation.
- In Release 6.0, Context-Aware Services on the Mobility Services Engine can be managed by only one Wireless Control System.
- In Release 6.0, WCS supports the creation of up to 124 WCS virtual domains.
- Although there is no “hard” limit on the number of simultaneous WCS users supported per WCS server in Release 6.0, Cisco has not validated more than five (5) simultaneous users per WCS server at this time.

Context-Aware Engine for Tags (AeroScout)

The Context-Aware Engine for Tags used in version 6.0.85.0 of the MSE Context-Aware Services software supports network designs containing up to 255 floor maps. In small enterprise designs, a network design might be used to describe the overall small enterprise from the point of view of context-aware services. An MSE network design typically consists of campus, buildings, and floor maps. The limitation on the number of floor maps then could be interpreted as saying that a small enterprise should not contain more than 255 floor maps (regardless of the number of buildings these floor maps are spread amongst). If more than 255 floor maps are required, it is necessary to divide the small enterprise up into two or more network designs, with each network design containing a subset of the total number of floor maps.

Obviously, the impact of this restriction will vary with the number of floors in each building, multiplied by the number of buildings contained within the entire small enterprise. For example, if all buildings in the enterprise contain only a single floor, then up to 255 buildings could be defined in a single network design before the conditions of this restriction are encountered. If each building each contained three floors however, then only 85 site buildings could be defined before this limitation would take effect.

Integration within Small Enterprise Designs

MSE Connection to the Network

Figure 6-5 is an illustration of the rear panel of both the MSE-3350 and the MSE-3310. The Cisco Mobility Services Engine is equipped with two 10/100/1000BASE-T Gigabit Ethernet ports (shown in Figure 6-5 by the solid arrows) that can be used to directly connect the MSE to two different IP networks (dual-homed). This makes it a simple affair, for example, to configure a MSE for service on network A while affording it the capability to be managed out-of-band on network B if the need arises.

In the Small Enterprise Design Profile, we attach the MSE to the network via a single connection to the NIC0 interface.

Figure 6-5 Rear Panel Illustration of MSE-3350

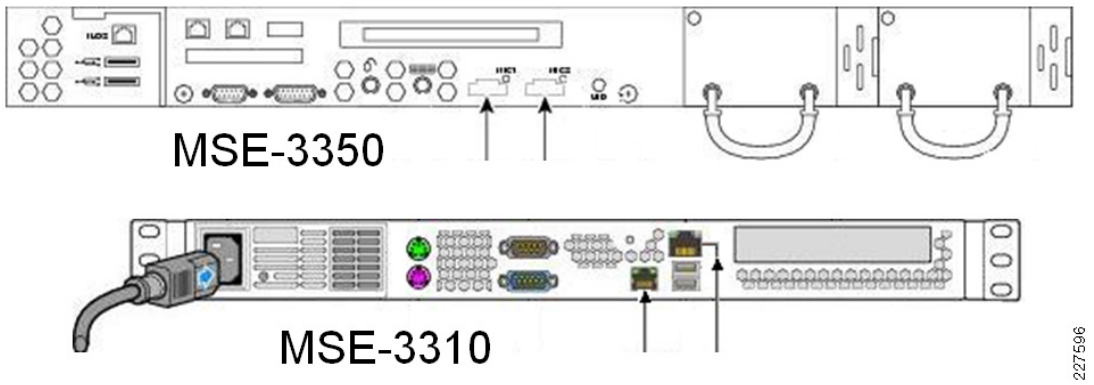

Note The dual on-board Ethernet controllers on the MSE are not intended for redundant or simultaneous connections to the same IP network. Any attempt to manually configure the MSE in order to try and establish parallel, load balancing, or redundant Ethernet connections to the same IP network is not recommended or supported at this time.

Clock Synchronization

The Mobility Services Engine, WCS, WLAN controllers, and any switches that support context-aware services use Coordinated Universal Time (UTC)[3] in their interaction. Because proper certificate authentication relies on time base consistency between participating components, it is important to ensure that these components are synchronized to a common time base throughout our design. In addition, having components synchronized to a common time source makes troubleshooting much easier, especially when having to look at events occurring within the logs of different network components. And the output of coordinated information by a central source, such as WCS, makes much more sense when the time stamps of all information displayed follow a logical flow and make sense to the user.

Once time and date in each network component has been set initially, time synchronization should be maintained using the Network Time Protocol (NTP). In the Small Enterprise Design Profile, these components should be synchronized to the ISR router found at the site.

3. For applications such those anticipated in the small enterprise designs discussed in this document, Universal Coordinated Time may be considered as equivalent to Greenwich Mean Time (GMT).

NTP Configuration of the Mobility Services Engine

Configuration of the NTP server addresses used by the MSE is handled during installation and the execution of the MSE automatic configuration script. An excerpt of that script is shown below. Detailed information regarding the automatic configuration script can be found in the Automatic Installation Script section (http://www.cisco.com/en/US/docs/wireless/mse/3350/quick/guide/mse_qsgmain.html#wp1057105) of the Mobility Services Engine Getting Started Guide, http://www.cisco.com/en/US/docs/wireless/mse/3350/quick/guide/mse_qsgmain.html#wp1057105, and in Appendix A (http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#appena) of the Mobility Services Engine Context Aware Deployment Guide, http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#appena.

```
If you choose to enable NTP, the system time will be configured from NTP servers that you
select. Otherwise, you will be prompted to enter the current date and time.

NTP is currently disabled.

Configure NTP related parameters? (Y)es/(S)kip/(U)se default (S)kip: Y

Enter whether or not you would like to set up the Network Time Protocol (NTP) for this
machine.

If you choose to enable NTP, the system time will be configured from NTP servers that you
select. Otherwise, you will be prompted to enter the current date and time.

Enable NTP (yes/no) no : yes
Enter NTP server name or address: <IP address or DNS name of NTP server>
Enter another NTP server IP address (or none) none: none
```

NTP Configuration of WLAN Controllers

Configuration of the internal clock and the specification of which NTP servers to use for periodic time synchronization can be performed on the WLAN controller using either the web GUI interface or the command line interface.

If you did not configure the system date and time through the configuration wizard when the controller was initially configured, or if you want to change your configuration, you can follow the instructions located in the section entitled Managing the System Date and Time (http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/c60intf.html) in the WLAN Controller Configuration Guide 6.0 (http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/c60intf.html#wp1144340) in order to configure the controller to obtain the date and time from a Network Time Protocol (NTP) server.

NTP Configuration of the Wireless Control System (WCS) Server

Configuration of the internal clock and the specification of which NTP servers to use for periodic time synchronization must be performed on the WCS server using the time and date capabilities of the WCS host operating system in use (either Windows or Linux).

RHEL-Based WCS Server

For a Redhat Linux-based WCS server, login to the host OS as root and use the following procedure to synchronize the internal software clock to the NTP server, synchronize the software clock to the server's hardware clock, and then maintain synchronization by starting the ntpd client daemon:

1. **clock**—Displays the current setting of the software clock.
2. **/etc/init.d stop**—Stops the ntpd client if it is already running.
3. **ntpdate <*ntp server name or address*>**—Synchronizes the system software clock with the NTP server.
4. **setup**—Brings up a setup utility that allows you to choose to set the time zone (shown in Figure 6-6).
5. **hwclock-systohc**—Writes the software clock settings to the hardware clock.
6. **/etc/init.d/ntpd start**—Starts the ntpd daemon to keep the clock synchronized going forward.[4]

Figure 6-6 RHEL Setup Utility

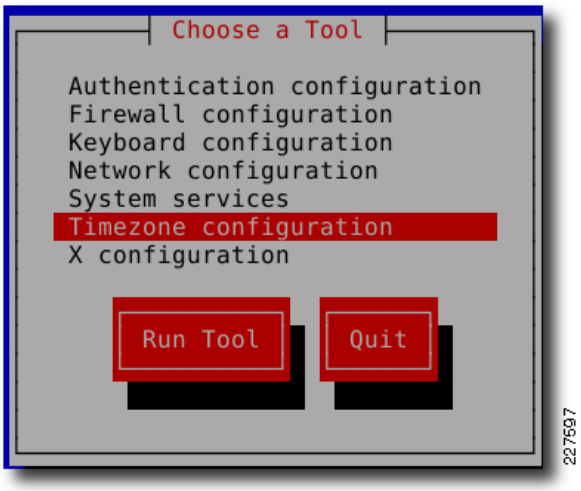

Note There are various other approaches that can be used to set the time zone on a Linux system. The reader is encouraged to consult the Redhat documentation for methods involving the use of the TZ variable or symbolic links to the localtime file or a particular time zone file in the system's time zone directory.

Windows 2003-Based WCS Server

For a WCS server based on the Microsoft Windows 2003 Server OS, use the following procedure to synchronize and maintain the correct system time via the Windows Time service (see Figure 6-7):

1. Check **Settings>Control Panel >Administrative Tools>Services** for the Windows Time service and ensure that it has been started.
2. Right click on the **Task Bar** clock and select **Adjust Date/Time**.
3. Under the **Date & Time** tab, set the current date and clock time to the approximate time of your NTP server.
4. Set the Time Zone and Daylight Savings time selections appropriately.

4. If ntpd does not start as part of your system boot script, you might want to add it using the command **chkconfig --add ntpd**.

5. Select the Internet Time tab, check the box to Automatically Synchronize With An Internet Time Server, type in the DNS name or address of your NTP server, and then Apply.

Figure 6-7 Setting Time and NTP Server on Windows 2003

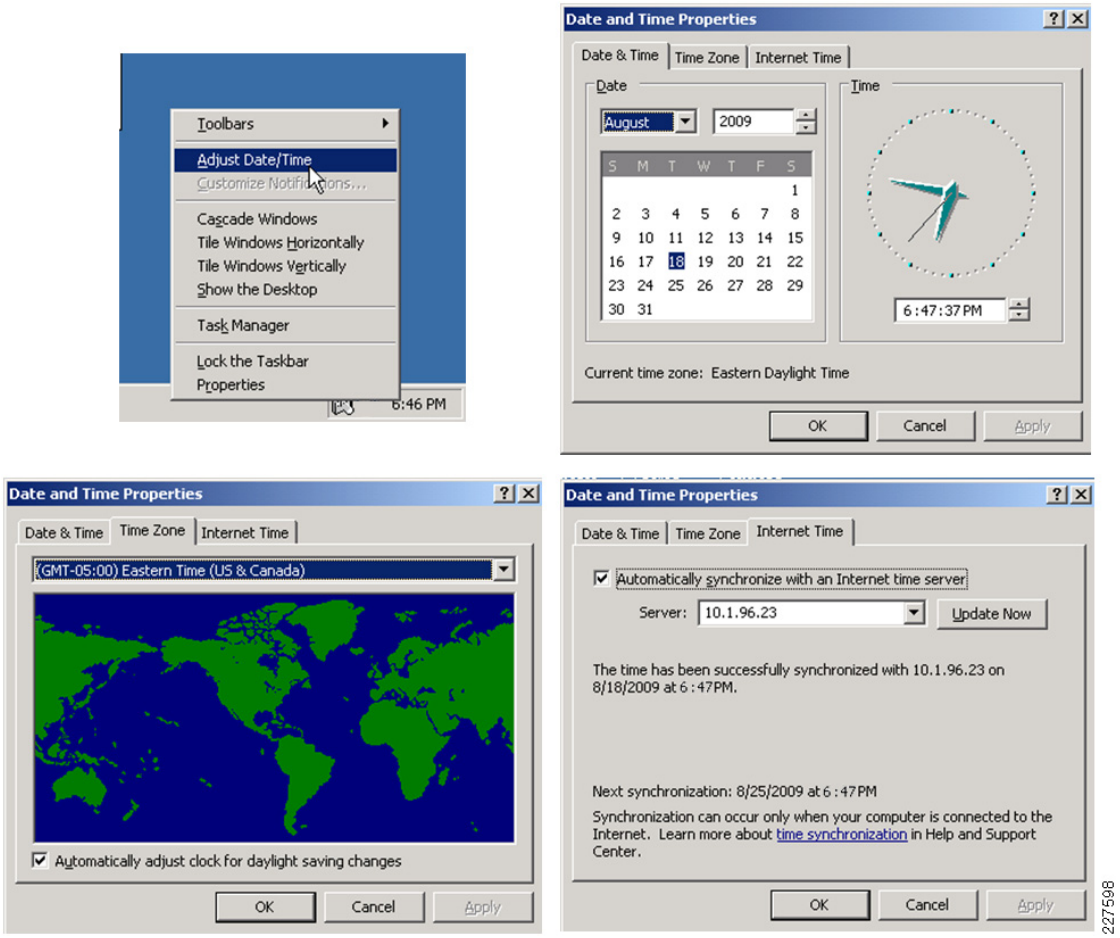

NTP Configuration of Context-Aware Catalyst Switches

In order to prevent any issues with authentication and NMSP session initiation, context-aware Cisco Catalyst switches should be configured to utilize NTP in order to keep their clocks in synchronization with other context-aware components. NTP is configured similarly amongst the various switch models discussed in this chapter, and the most comprehensive information on how to configure a Catalyst switch as an NTP client can usually be found in the configuration guide for the particular switch model. For example, for the Catalyst 2960G NTP configuration is documented in Configuring NTP section of the Catalyst 2960 Switch Software Configuration Guide (http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/swadmin.html#wp1053923).

It is good general best practice to ensure time synchronization of all network components when possible. However, from the perspective of context-aware services in the small enterprise designs discussed here, only those switches that are actually participating in an NMSP session with the MSE require clock synchronization.

SEDP Wireless Control System (WCS) Context-Aware Considerations

The Wireless Control System is used for several important configuration tasks relating to context-aware services in small enterprise designs:

- Creation of a network design at either the enterprise or site level, which may include campus, building, floor, and outdoor level maps.
- Definition of WCS User Groups, which are used to define what management actions context-aware users are authorized to perform.
- Definition of Virtual Domains, which can be used to restrict which network resources users will have the ability to manage via WCS.
- Configuration of Mobility Service Engine operating parameters. This represents the next level of MSE setup beyond that performed by the MSE automatic configuration script discussed in section "NTP Configuration of the Mobility Services Engine" section on page 6-19.
- Definition of Context-Aware Conditional Notifications, which defines how applications and parties external to the site may receive notification of specific events pertaining to changes in contextual characteristics associated with clients, tags or rogue devices.

In this section, we discuss only those areas where, in our testing we made use of significant WCS features relevant to the integration of context-aware services, or where important configuration changes were made that significantly differ from the defaults. This is not meant to serve as a comprehensive configuration guide to all aspects of the WCS and MSE. Readers should refer to the WCS and MSE configuration documents already cited throughout this document (including Context-Aware Services General Best Practice References) for additional information regarding configuration parameters and procedures that, while not discussed in detail here, must still be configured or performed properly.

Creation of a Network Design

Once access points have been installed and have registered with a controller, WCS can be configured to manage the controllers and a network design can be set up. A network design is a representation within WCS of the physical placement of access points and other context-aware components throughout a facility or group of facilities. A hierarchy consisting of a single campus, the buildings that compose that campus, the floors of each building, and any outdoor areas constitutes a single network design.

In small enterprise designs, the choice of whether to configure the entire enterprise or each individual site at the campus layer depends to a large part on the on whether the sites each contain a single building, or multiple buildings. If each site is comprised of one building and one building only, then the campus layer of the network design can be the entire enterprise. On the campus map, each site would be represented by a single building with one or more floors per building.

In other cases, each enterprise site might be composed of multiple buildings. In these scenarios, it makes more sense to define each site as a campus unto itself.

A step-by-step set of configuration instructions regarding how to configure network designs consisting of campus, building, and floor maps can be found in Chapter 5 of the *Cisco Wireless Control System Configuration Guide*, Release 6.0, http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0maps.html#wp1203275.

Figure 6-8, Figure 6-9, and Figure 6-10 give an example of what a campus, building, and floor level network design might look like for a small enterprise where all small sites are assumed to be comprised of single buildings. In Figure 6-8, we use satellite imagery of the small enterprise's physical location as the backdrop for the campus map. Clicking on any of the icons takes us to the building map for the small enterprise site. In this case, we select site number 1278, which then brings us to Figure 6-9.

Figure 6-8 Campus Level View of the Small Enterprise

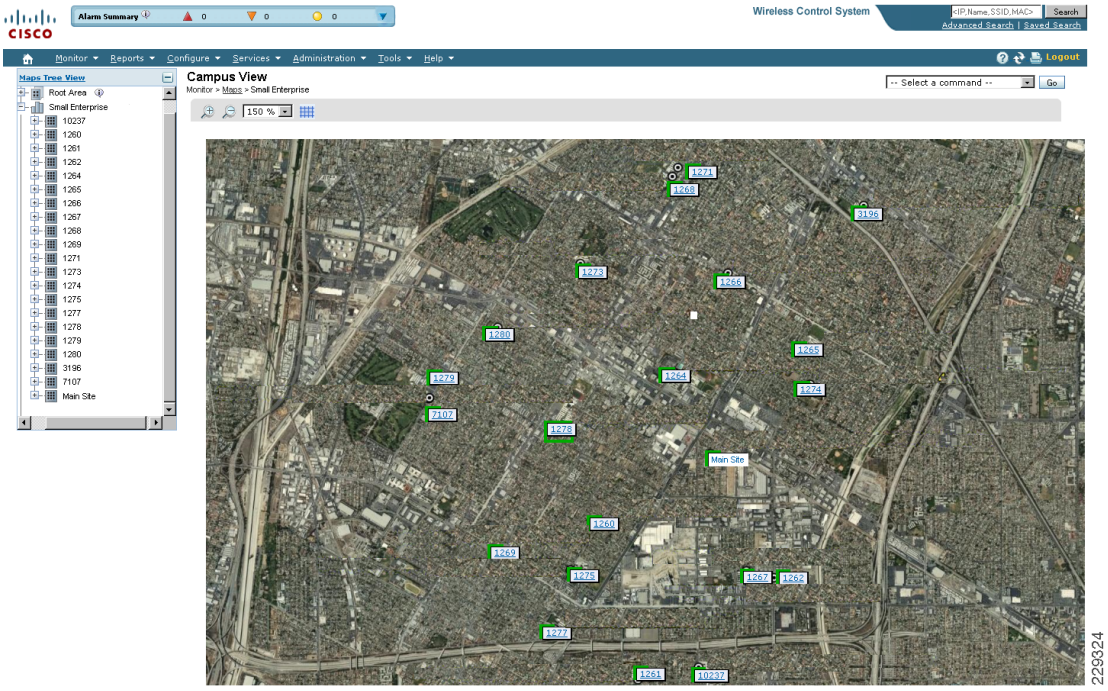

The building map shown in Figure 6-9 indicates that this building contains a single floor. Clicking directly on the building map takes us to the floor definition shown in Figure 6-10. This is where we would actually see the location of wireless clients, active RFID tags and rogues displayed. Wired devices that are attached to context-aware Cisco Catalyst switches are not displayed on floor maps in Release 6.0 of Context-Aware Services.

Figure 6-9 Building Level View for Site #1278

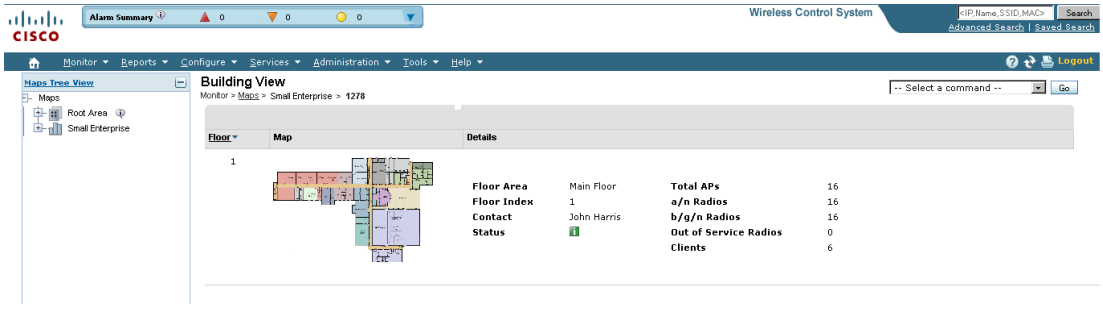

Figure 6-10 Floor Level for Site #1278

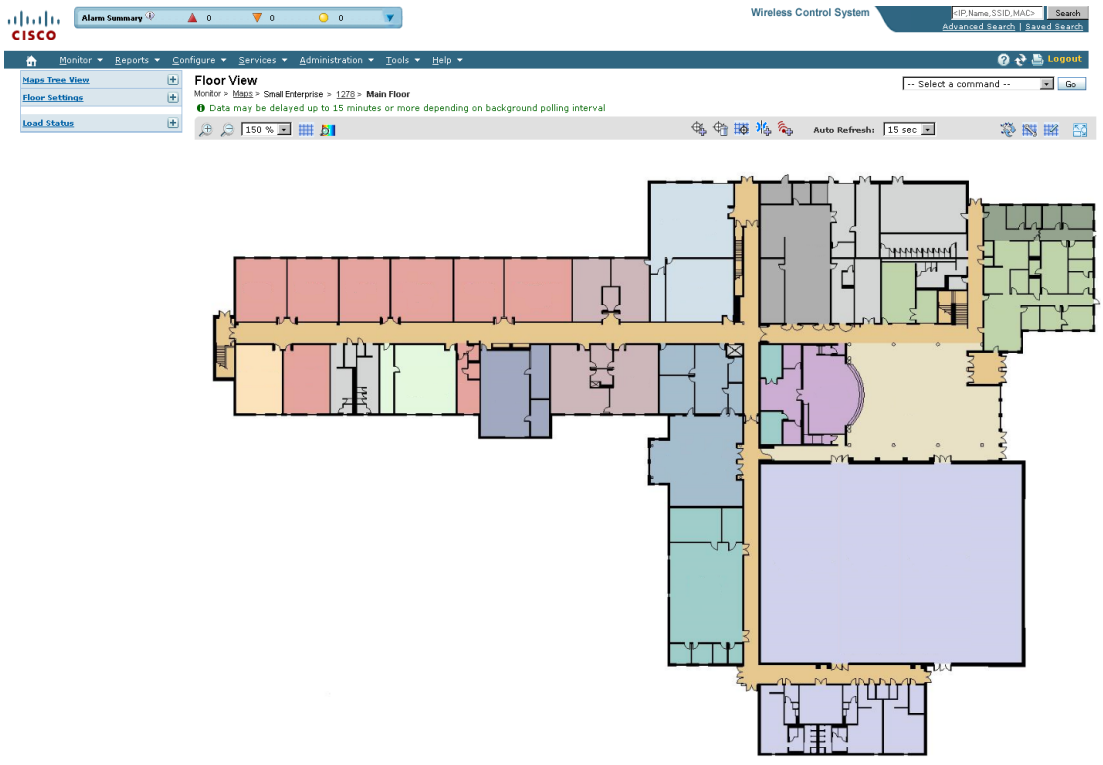

Note Network designs are created in WCS, but they are not actually used for device tracking until they are transmitted to the MSE via a process known as network design synchronization. Only after network design synchronization has successfully occurred between WCS and its associated MSE will the network design actually be used by the Context-Aware Engine for Clients and the Context-Aware Engine for Tags. Synchronization of network designs and other components with *the MSE is discussed in detail in the chapter entitled "Synchronizing Mobility Services Engines" in the Context-Aware Service Configuration Guide* 6.0, http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch3_CAS.html.

WCS Users, User Groups, and Virtual Domains

When installed, WCS provides for a single root user, which will have access to all WCS functions. The password for this root user should be protected and only known by those personnel at the main site data center with a true need to know (e.g. those personnel responsible for the installation, maintenance, and detailed administration of WCS). Instead of using the root user password for routine access to WCS, you should create other users and grant them administrative access with privileges assigned as necessary via the use of WCS user groups. Chapter 7 of the *Cisco Wireless Control System Configuration Guide*, Release 6.0 (http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0manag.html) provides comprehensive instructions with regard to the proper procedure for configuring users and group privileges on your WCS server. This chapter also contains a complete listing of the user groups available in WCS as well as the privileges contained in each group.

Common sense should be applied in the assignment of user privileges in small enterprise designs. For example, while only a very small set of key technical personnel should have access to the actual WCS root user ID and password, you may wish to assign the ability to make WCS configuration changes to a

somewhat larger audience. This larger group can be assigned as WCS “admin” users or assigned to the “superuser” group. Most small enterprise users will likely require only the capability to view or monitor network services that pertain to them or their workgroups, as opposed to administering the network in WCS. For these users, the privileges accorded them by the WCS System Monitoring or Monitor Lite user groups may be all that is required, depending upon the specific WCS monitoring functions you wish to grant those users.

In our small enterprise design validation, the custom user-defined group shown in Figure 6-11 was found to be very useful in limiting users to only monitoring context-aware information, as well as some basic WCS alerts and events. For example, we may wish to allow a research specialist in a small enterprise document center access to the context-aware functions listed under the WCS “Maps” function or search for a device by name. But this same research specialist probably has no need to monitor network security compliance reports, thus we have not enabled access to those reports for the research specialist’s account. Keep in mind that the requirements of small enterprise users in your environment may be different, so it may make sense for you to develop a custom WCS user group that closely fits your needs.

Figure 6-11 Custom User Group to Allow Context-Aware Monitoring

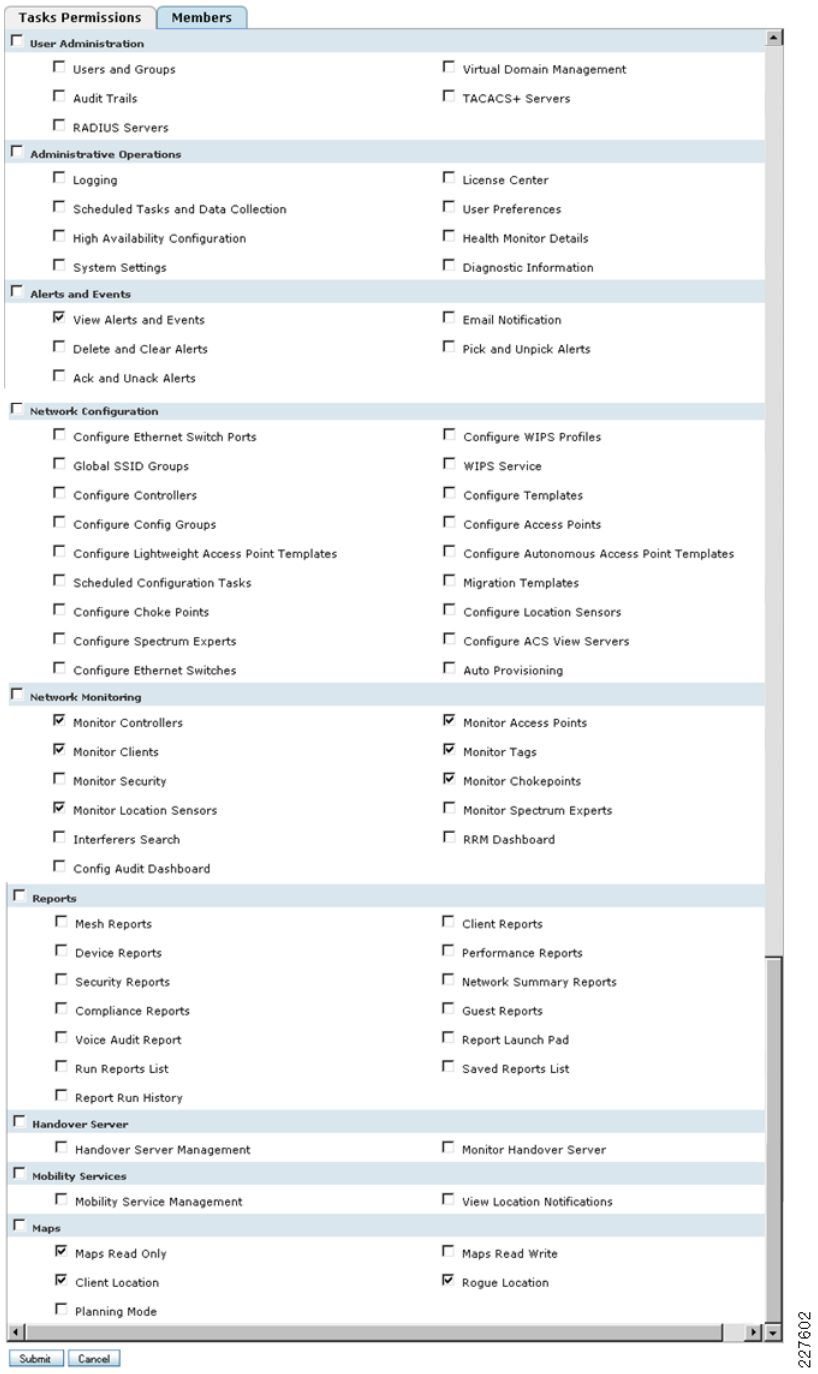

While WCS user groups define the WCS functionality users have been granted, WCS virtual domains allow the network administrator logically partition the WCS management domain and limit management access. In this way, the group of resources that the WCS functionality assigned to a user group may be exercised against is restricted. A WCS virtual domain consists of a set of assigned devices and maps, and restricts a user's scope to only information that is relevant to those devices and maps. Through a assigned virtual domain, users are only able to utilize WCS functionality against a pre-defined subset of the devices managed by WCS.

Users can be assigned one or more virtual domains, however only one assigned virtual domain may be active for a user at WCS login. The user can change the current virtual domain in use by selecting a different permitted virtual domain using the WCS Virtual Domain drop-down menu.

The WCS virtual domain can be used to limit the user's ability to even view certain resources inside WCS that are not contained in their active assigned virtual domain. For example, the site manager of small enterprise site A have the ability to view the location and other context-aware characteristics of wireless assets due to his WCS user account being assigned to an appropriate user group permitting this level of WCS functionality. But the virtual domain that this site manager is assigned may only allow such functionality to be exercised against these assets if they are located within the small enterprise he is assigned to manage. Thus, if our site manager attempted to discover the quantity and location of RFID-tagged equipment in small enterprise site "B", his assigned virtual domain prevent him from being able to view site B resources.

Administrative personnel with enterprise-wide responsibilities, on the other hand, would be assigned a virtual domain that includes all resources in the enterprise, across all sites, and could exercise the functionality assigned to them by their user group against any of these resources. The virtual domain assignment also helps prevent unnecessary inter-site WCS traffic, especially traffic whose nature might be based more upon curiosity rather than actual need.

Note WCS user groups assign what actions a user can take against a resource, whereas WCS virtual domains determine what resources those user group actions can be applied towards.

There are two basic steps necessary to enable the use of virtual domains within WCS:

1. A virtual domain must be created, and the resources that we wish to include assigned to the virtual domain. The process for creating and assigning network resources to the virtual domain is detailed in Chapter 20 Virtual Domains of the WCS Configuration Guide 6.0 (http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0virtual.html#wp1040002).
2. The virtual domain must be assigned to the user. The process for assigning a virtual domain to a user is detailed in Chapter 7 Managing WCS User Accounts (http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0manag.html#wp1097733).

Note It is important to note that in release 6.0, non-root WCS virtual domain users cannot access WCS functions listed under the **Services > Mobility Services** main menu heading. This includes wired switch and device location. Refer to Understanding Virtual Domains as a User, WCS Configuration Guide 6.0 (http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0virtual.html#wp1120787) for a complete list of WCS functions that are not available in non-root virtual domains.

In Release 6.0, since wired devices attached to context-aware Cisco Catalyst switches are displayed using **Services > Mobility Services > Context Aware Service > Wired > Wired Clients**, only users that are assigned to the root virtual domain are able to display context-aware information for these devices.

Figure 6-12, Figure 6-13, and Figure 6-14 demonstrate the effectiveness of WCS virtual domains (note that the current virtual domain in use by the logged-in WCS user is highlighted in each figure by the red oval). In Figure 6-12, we can see that the "Small Enterprise" virtual domain user can see the entire set of sites comprising the small enterprise, and is capable of applying any of the WCS functionality accorded to them by their WCS user group assignment. This virtual domain setting might be appropriate, for example, in the case of a person requiring the ability to view and potentially take action upon all resources in the network. An example of personnel in a small enterprise that might require such

capability would be network administration staff. However, keep in mind the immediately preceding note regarding the viewing or administration of WCS resources appearing under the Services > Mobility Services main menu heading.

Figure 6-12 Virtual Domain for Entire Small Enterprise

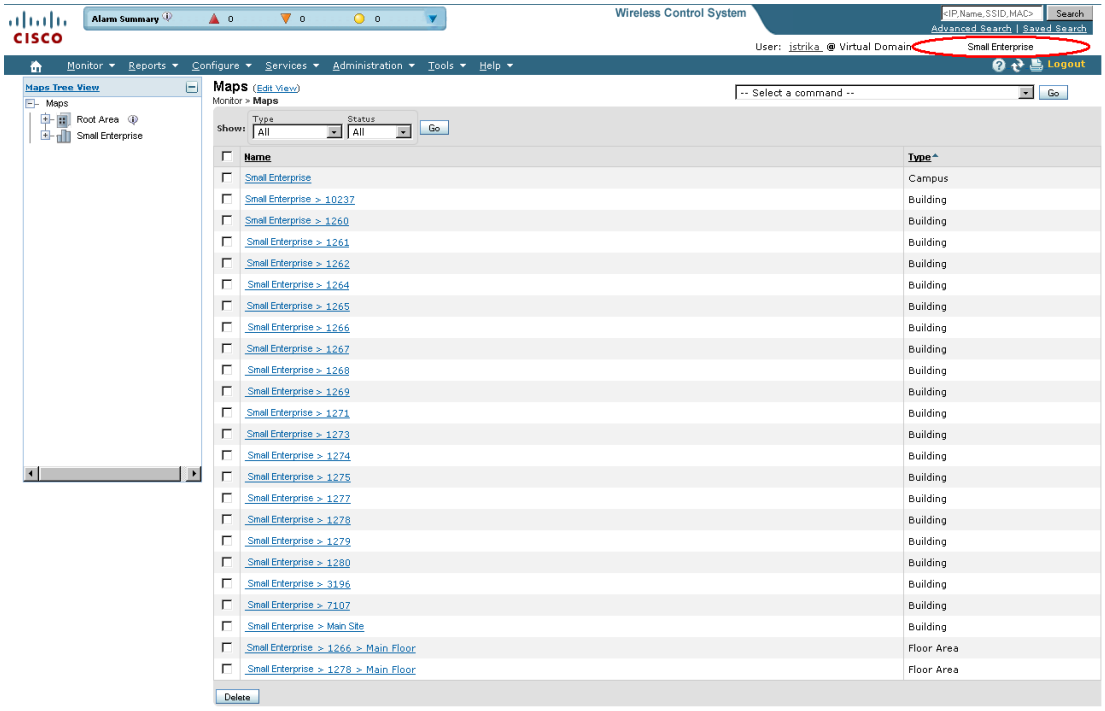

In contrast, Figure 6-13 and Figure 6-14 each illustrate how the user's view of the small enterprise can be severely curtailed when a WCS virtual domain is applied. A virtual domain setting such as this might be appropriate for most of the personnel within a site that might only need to work with resources located in their site only. Figure 6-13 illustrates what a user in site #1278 would see if they were assigned a WCS virtual domain that limited their resource visibility to only those resources associated with site #1278.

Figure 6-13 Virtual Domain Limited to Site #1278

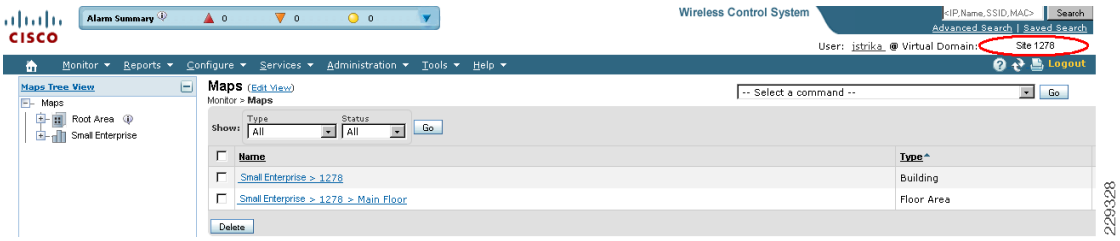

In Figure 6-14, we see the results of a WCS virtual domain for a user in site #1266. Note that a user that is assigned a virtual domain for site #1266 does not have visibility to any resources associated with other sites within the enterprise. All that is visible to the site #1266 user in this case are the buildings and floor maps that are associated with site #1266 and site #1266 only. We can see that plainly in Figure 6-14, as the resources associated with site #1278 or any other site are simply not displayed to the site 1266 virtual domain user.

Figure 6-14 Virtual Domain Limited to Site #1266

Figure 6-15 illustrates the typical result that occurs when a user attempts to view information for resources outside the scope of their assigned virtual domain. In this case, we see the result of a user in site #1266 attempting to access resources that are located in site #1278.

Figure 6-15 Virtual Domain Permission Error

Mobility Services NMSP Parameters

WCS 6.0 provides us with several NMSP parameters available that affect various NMSP protocol timing characteristics between the MSE and its session partners. These parameters can be found at **Services > Mobility Services > System > NMSP Parameters**, and apply globally to all NMSP sessions between the selected MSE and any of its WLAN controller or Cisco Catalyst switch session partners. Complete configuration information for configuring NMSP session timing parameters, as well as the default values for these parameters, can be found at Configuring Mobility Services Engine Properties section of the *Context-Aware Service Configuration Guide* (http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch4_CAS.html#wp1014368).

When deploying an MSE locally in the small enterprise, it is unlikely that these parameters would require modification. This would be more the norm in large enterprise sites, where there might be several thousand tracked devices, and a large quantity of wireless devices moving about on a regular basis. However, in a centralized deployment, there is more of a chance that network congestion or other factors may cause delays that could cause NMSP session time outs. While this should be minimized by the appropriate identification and classification of NMSP data flows in the network along with properly defined network QoS, there may be instances where adjustments to NMSP timing is required. In these cases, the NMSP echo interval, neighbor dead interval, and response time out values can be increased to limit the number of failed echo acknowledgments that may occur, especially in a centralized MSE deployment.

Note Readers are advised that a tremendous amount of functional validation was performed in association with the content contained in this chapter. However, time constraints limited the degree of performance validation that could be completed for Context-Aware Services across a simulated Metropolitan Area Network (MAN). The deployment of Mobility Services Engines in a centralized fashion across a MAN was not able to be fully validated due to these time constraints.

To aid in determining whether echo packets are being dropped, you can use the WCS function **Services > Mobility Services > System > Status > NMSP Connection Status** as shown in Figure 6-16. This WCS menu panel displays all the NMSP session partners for this MSE, along with echo request and response counts. For example, in the figure for the Mobility Services Engine with the hostname MSE1, we see that there are currently two NMSP sessions to two different WLAN controllers.

Figure 6-16 NMSP Connection Status

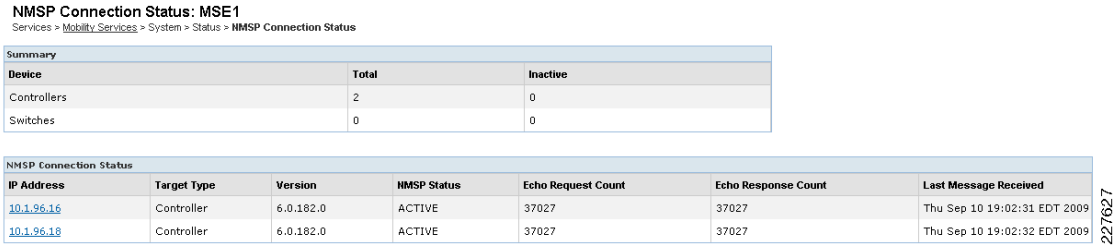

NMSP Connection Status: MSE1

Services > Mobility Services > System > Status > **NMSP Connection Status**

Summary

Device	Total	Inactive
Controllers	2	0
Switches	0	0

NMSP Connection Status

IP Address	Target Type	Version	NMSP Status	Echo Request Count	Echo Response Count	Last Message Received
10.1.96.16	Controller	6.0.182.0	ACTIVE	37027	37027	Thu Sep 10 19:02:31 EDT 2009
10.1.96.18	Controller	6.0.182.0	ACTIVE	37027	37027	Thu Sep 10 19:02:32 EDT 2009

227627

We can see in Figure 6-16 that both of the NMSP WLAN controller sessions appear to be functioning properly, as the number of Echo Requests issued is seen to be exactly equal to the number of Echo Responses received. This might not always be the case and small static differences over a long period of time do not necessarily indicate a serious problem. However, sluggish performance combined with a regularly increasing discrepancy in the delta between the number of requests issued and responses received could be indicative of the NMSP session timing out. In a centralized deployment, this may be due to unforeseen levels of congestion or other resource constraint. In this case, raising the response time-out may assist in alleviate the time-outs. Keep in mind, however, that a successful centralized deployment assumes that there is sufficient MAN capacity available to each site and that QoS has been applied appropriately.

Figure 6-17 gives an example of the NMSP Parameter screen that is located at **Services > Mobility Services > System > NMSP Parameters**, which we can use to change the system defaults. In Figure 6-17, the network administrator or network technician has instituted the following changes from the defaults: the echo interval has been raised to 30 seconds, the neighbor dead interval has been raised to 60 seconds, and the response time-out has been raised to 5 seconds.

Figure 6-17 Example of NMSP Parameter Modification

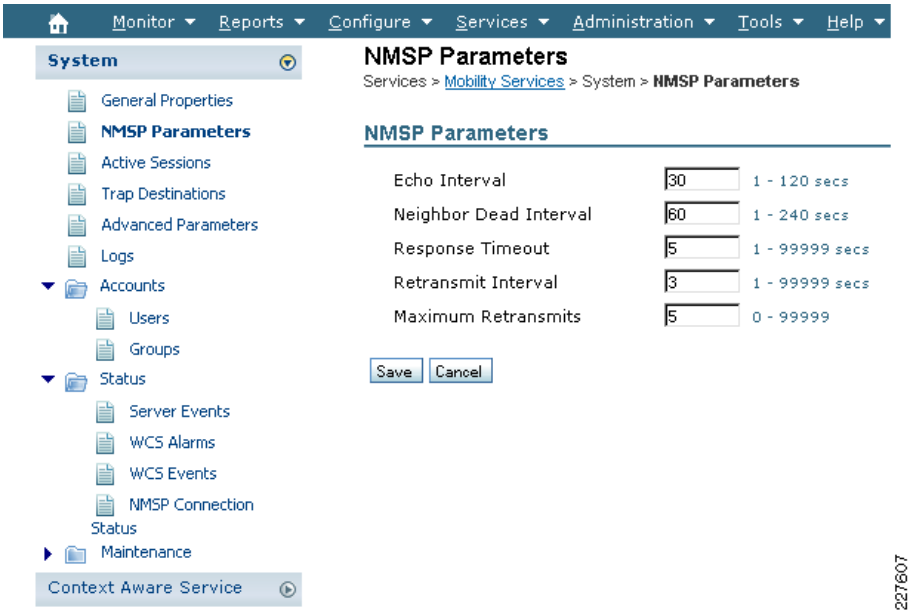

Note Although the NMSP configuration worked well for us in our lab testing, we cannot predict and simulate each and every condition that might occur in a production deployment. Therefore, it is important that you take the time to understand the function of these parameters and especially the fact that they can be adjusted beyond the values illustrated here in order to promote improved NMSP session stability and network performance.

Further information on these NMSP parameters can be found in the Configuring Mobility Services Engine Properties section of the *Context-Aware Service Configuration Guide* (http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch4_CAS.html#wp1014368).

Context-Aware Service Parameters—Tracking

As mentioned earlier, Context-Aware Services can track up to a maximum of 18,000 licensed devices when using the MSE-3350 hardware platform, and up to a maximum of 2,000 licensed devices when using the MSE-3310 platform. The absolute limit on the number of clients or tags that can be tracked is determined by the hardware platform used, the presence of any other applications co-residing on the MSE, and the level of licensing purchased. The WCS tracking parameters configuration panel (located at **Services > Mobility Services > Context Aware Service > Administration> Tracking Parameters**) allows the administrator to pre-determine just how much of the MSE's maximum licensed tracking capacity will be allocated towards the tracking of specific device categories. This is useful in the small enterprise environment in order to allow the tracking of device categories such as nearby rogue access points and rogue clients, but also limit these categories such that an uncontrolled introduction of rogues is not allowed to consume all of the remaining context-aware tracking capacity on the MSE.

We can use the Context-Aware Service Tracking configuration to:

- Entirely enable or disable the tracking of wired and wireless client stations, asset tags, rogue access points, and rogue clients.

- Set limits on how much MSE tracked device capacity will be allocated to certain device categories. Figure 6-18 provides us with an example of how this can be achieved, where the maximum number of tracked clients and rogue clients/APs are capped at 4,000 devices each. No limit is placed on the number of RFID tags tracked, which in effect means that the maximum number of tags tracked will be allowed to rise until the tag licensing limit is reached (3,000 tags).

Note that any devices that are detected but excluded from tracking due to the enforcement of a tracking limit will be reflected in the "Not Tracked" device count column shown on the right side of the display.

Figure 6-18 Context-Aware Service Tracking Parameters

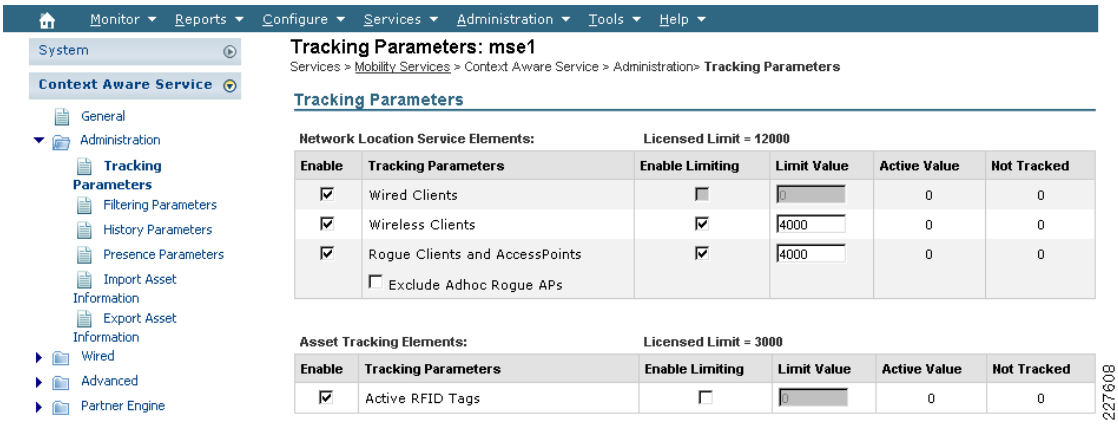

Note In Release 6.0, wired client tracking can be enabled or disabled, but imposing a limit on the number of wired clients that are tracked simultaneously is not supported at this time. As a work around, use switch CLI commands such as the global nmsp disable or the interface nmsp suppress attachment commands to limit the number of tracked wired clients that are presented to the MSE.

Context-Aware Service Parameters—History

The MSE records and maintains historical location and statistics information for wireless clients, tags, rogue access points, and rogue clients. This information is available for viewing through WCS or via third-party context-aware application clients, and can be very valuable in helping establish patterns of movement for tracked assets and rogues. Historical information can be used for location trending, asset loss investigation, RF capacity management, and facilitation of network problem resolution. Contextual information such as whether an emergency button was depressed or whether an asset tag has moved into close proximity to a chokepoint trigger is also tracked in the history data.

The collection of historical information must be explicitly enabled for each desired category of device (as shown in Figure 6-19). By default, 30 days of historical data are stored in the MSE.

Figure 6-19 MSE History Parameters

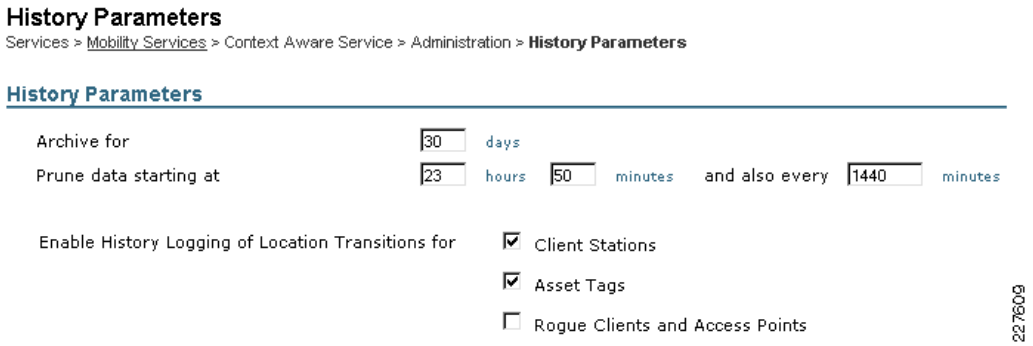

There are several variables that can affect how much historical information can be stored by the MSE for tracked assets. Among these variables are the average number of elements that move, average distance covered every time there is a movement, information transitions, telemetry information from tags, and so on. Depending on these variables as well as the number of items for which you are tracking history information, you may wish to decrease the number of days of historical data to a value below 30 days.

Changes to the default history archive period should be done with careful consideration, since longer history periods typically increase the amount of space consumed by the history database. Users unfamiliar with the way in which Context-Aware Services on the Cisco MSE archives historical information for tracked devices may wish to consult with your Cisco field technical representative or the Cisco Technical Assistance Center

Figure 6-20 illustrates what you can expect to see when recalling the history for a tracked device. Here, we recall the history for a specific WLAN client. As shown in Figure 6-20, the focus of the display is the list of past locations recorded for the client. Setting the “change selection time” parameter on the screen and then clicking play displays the various client locations on the small floor map at the right side of the image (this can be enlarged for easier viewing). In this way, you can step back through all of the stored locations for the client within the history database. Obviously, this information could be very useful to administrators, WLAN engineers, as well as enterprise security officers and other law enforcement officials that may be looking for information useful in recreating a pattern of potentially suspicious past movement.

Figure 6-20 Display of Location History for a WLAN Client

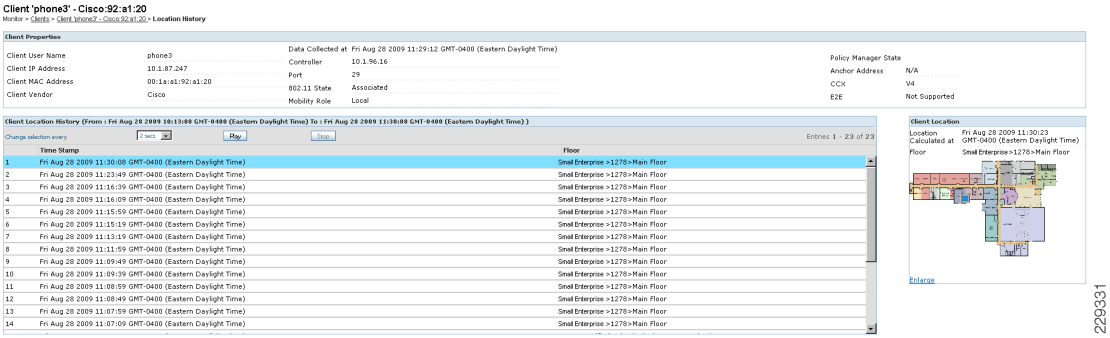

Further information on the procedure to follow when making adjustments to history parameters can be found in the following documents:

- *Context Aware Solution Deployment Guide*
  http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#rfid

- *Context Aware Service Configuration Guide 6.0*
  http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1128896

Context Aware Service Parameters—Notifications

Cisco WCS allows you to define certain conditions for tags, WLAN clients, and rogues that cause the MSE to send notifications to application programs that are monitoring specific ports. You can use Cisco WCS to define and enable both conditional notifications and northbound notifications.

Conditional notifications are those notifications that the mobility services engine sends to Cisco WCS and other applications that can receive short, relatively simple messages via SOAP/XML (either HTTP or HTTPS), SMTP, UDP Syslog, or as an SNMP trap. Conditional notifications can be triggered by WLAN clients, tags, or rogue devices. The conditions available include:

- *Missing*—The MSE generates a Missing Asset conditional notification if it has not located the asset for more than a specified number of minutes.
- *In/Out*—The MSE generates an In/Out conditional notification if the asset is found to be inside of, or outside of, a selected area.
- *Distance From Marker*—The MSE generates a Distance From Marker conditional notification if the asset is found to be beyond a specified distance from a designated marker.
- *Battery Level*—The MSE generates a Battery Level conditional notification if the battery level reported by an asset tag is equal to a selected value.
- *Location Change*—The MSE generates a Location Change conditional notification if the asset experiences a change in location.
- *Emergency*—The MSE generates an Emergency conditional notification if a tag button, tamper or detached event is detected.
- *Chokepoint*—The MSE generates a Chokepoint Conditional notification if a tag enters into the proximity of a chokepoint trigger.

Northbound notifications are a special category of notification that is specific to RFID tags only. They define which tag notifications the MSE will send to third-party applications using SOAP/XML. Northbound notifications can include chokepoint, telemetry, emergency, battery, and tag vendor data. Optionally, the tag's location coordinates can be included within the northbound notifications as well. An important difference between conditional notifications and northbound notifications is that northbound notifications can contain considerably more information. The information sent in the northbound notification is sent in a pre-defined data format. Details regarding the data format for northbound notifications are available on the Cisco developers support portal at http://developer.cisco.com/web/contextaware.

Note In Release 6.0, neither conditional nor northbound notifications can be applied to tracked wired devices.

Complete and detailed information regarding how to configure Context-Aware Notifications can be found in the following documents:

- "Configuring Event Notifications" chapter *of the Context-Aware Configuration Guide* 6.0 (http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch6_CAS.html).

- "Enabling Notifications and Configuring Notification Parameters" section of the *Context-Aware Configuration Guide* 6.0 (http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1129909).
- "Context-Aware System Performance" section of the *Context-Aware Solution Deployment Guide* (http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#caasperf).

Via the menu panel located at **Services > Mobility Services > Context Aware Service > Advanced> Notification Parameters** (Figure 6-21),WCS allows the user to modify several advanced timing parameters pertaining to conditional notifications and northbound notifications. For example, you can limit the rate at which the MSE generates notifications, set a maximum queue size for notifications, and set a retry limit for notifications with in a certain period.

Figure 6-21 Context Aware Advanced Notification Parameters

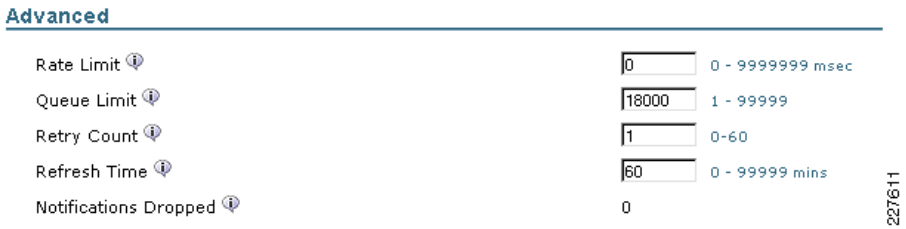

Note Modify advanced notification parameters only when you can reasonably expect the Mobility Services Engine to transmit a large number of notifications or when it is noticed that notifications are being dropped.

- *Rate Limit*—(Pertains to northbound notifications only.) This is the rate in milliseconds at which the MSE generates northbound notifications. A value of 0 (default) means that the Mobility Services Engine generates event notifications as fast as possible. If you are using the northbound notifications system to communicate to a third-party application (such as AeroScout MobileView, for example) in a small enterprise design, it is recommended that you consider how often notifications are expected to be generated and how many total notifications will be requested of the MSE per minute. For example, an exception based notification (such as an emergency notification) is not a normal event and thus it would be extremely unusual to see a great deal of emergency events generated from many sites at the same time. In contrast, more routine events may be seen to generate substantially greater traffic. In cases where a large number of notifications might be a normal and routine event, it is recommended to increase the rate limit so as to allow the MSE to pace the transmission of notifications (for example, a rate limit value of 200 would slow the rate of notification transmission down to one every 200 msec). The correct notification delay will vary from deployment to deployment depending on the amount of traffic generated.
- *Queue Limit*—Specifies the size of the output notification queue of the MSE. Default queue limit value for the MSE-3350 is 18000 and for the MSE-3310 it is 5,000. The MSE drops any outbound notifications above this limit if the output notification queue size is exceeded. In small enterprise designs, if you are using context-aware notifications, it is recommended that you use the Queue Limit parameter in conjunction with the Notifications Dropped counter to avoid any notifications from being dropped. If you notice that the Notifications Dropped counter is greater than zero, you should consider increasing the Queue Limit parameter to avoid any future increases. Given the relatively large default sizes for this parameter however, it is unlikely that adjustment will be required except in rare cases where very many notifications are generated.

- *Retry Count*—For each matching condition, the retry count specifies the number of times to generate an event notification before the refresh timer expires. The default value is 1. The total number of event notifications transmitted between Refresh Time periods is equal to one plus the value specified for Retry Count. After the value of one plus the Retry Count has been reached, the location appliance skips firing any further northbound notifications for this condition and device for the time period specified by the Refresh Time. Once the Refresh Time has expired, this cycle repeats unless the event has been cleared. Retry count is intended to help ensure (to a limited extent) that notifications reach their intended destination. If notifications are not indicated as being dropped by the Notifications Dropped counter, but are not reliably reaching their destination application, you may wish to increase the number of notifications sent between Refresh Time periods by raising the Retry Count judiciously.

Note The Mobility Service Engine transmits notifications using a "Fire and Forget" technique. Notifications are not retained in any database within the MSE after they are transmitted.

- *Refresh Time*—The wait time in minutes that must pass before an event notification is resent. The default is 60 minutes. Refresh Time and Retry Count are used cooperatively to help limit the number of notifications repeatedly generated for events that have not been cleared. Retry Count limits the number of notifications that are sent by the MSE, while Refresh Time imposes a "waiting period" during which time no further notifications are sent for this event condition and device. If you are noticing that repeated notifications are being generated for the same event, it may be due to the condition not clearing within the refresh time interval. In this case, you may wish to investigate why the triggered condition does not clear and possibly extend the refresh time.
- *Notifications Dropped*—The number of event notifications dropped from the queue since startup. The Notifications Dropped counter should be used in conjunction with the Queue Limit parameter to reduce the number of total dropped notifications.

WLAN Controller and Cisco Catalyst Switch Definition and Synchronization

To allow for proper tracking of the devices that may be registered or attached to them, WLAN controllers and context-aware Cisco Catalyst switches must be defined to WCS and then synchronized with the Mobility Services Engine.

Detailed information regarding how to add WLAN controller definitions to WCS using the **WCS Configure > Controllers** menu panel can be found in the section titled "Adding Controllers" in the *WCS Configuration Guide* 6.0: (http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0ctrlcfg.html#wp1041451) .

Detailed information regarding how to add Cisco Catalyst switch definitions to WCS using the **WCS Configure > Add Ethernet Switches** menu panel can be found in the *WCS Configuration Guide* 6.0: (http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0ctrlcfg.html#wp1089752) .

Detailed information regarding how to synchronize the MSE with WLAN controllers and context-aware Cisco Catalyst switches using the **WCS Services >Mobility Services > Synchronize WCS and MSE(s)** menu panel can be found in the "Synchronizing Mobility Services Engines" chapter of the *Context-Aware Service Configuration Guide* 6.0: (http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch3_CAS.html#wp998995).

Note Always ensure that the MSE is synchronized with the primary WLAN controllers it is providing Context-Aware Services for, as well as any backup WLAN controllers. If the MSE is not kept in synchronization with your backup WLAN controllers, Context-Aware Services may not function properly or may not be available at all for WLAN clients and RFID tags in the event of a primary WLAN controller failure.

Wireless Client Context-Aware Considerations

The ability to track WLAN client location using Cisco Context-Aware Services can be useful in a small enterprise to locate WLAN clients residing on the network (such as authorized Cisco 7921G and 7925 VoWLAN phones, laptops, wireless desktops, etc.). Generally speaking, provided that a 802.11 wireless client device has its 802.11 wireless network interface adapter powered on and is sending periodic transmissions (probe requests), these WLAN client devices can be located by Context-Aware Services. Since devices such as portable VoWLAN phones are very often powered on for the entire day and carried on the belt, purse, or pocket of the user, these devices become useful in helping locate not only the device itself but the user to which the device is assigned. As the device becomes larger in size and heavier in weight, we find that the chances of the device being with the assigned user all the time diminishes and thus its usefulness as a way to determine the location of the user diminishes as well.

Tracking WLAN client devices by their wireless network interface adapter works well if the goal is to track the location of the device when it is in operation and use this information to enhance the operation of the device on the network, or its interaction with an application. For example, using WLAN client tracking in the small enterprise to perform one or more of the following tasks are just some examples of how this functionality can be put to use:

- Determining where wireless users and their portable devices may congregate, allowing for the placement of access points to be further optimized to enhance overall coverage and performance.
- Investigate where a missing device is currently located within the main or remote site. A good example might be when equipment is “borrowed” from one office or conference room to another (or even from one site to another) without permission.
- Determining the location of users that are known to be visitors to one or more enterprise sites, and helping ensure that they do not stray into areas that are off-limits to them.
- Using the location history of a wireless device to establish a pattern of usage and location in order to clarify past actions, such as those that might relate to safety and security concerns.
- Use the location of wireless LAN clients as a parameter for network troubleshooting and security audits.
- Use the location of wireless printers (or other output devices) as input to an application designed to determine which device is available and most convenient for a particular wireless user depending on the user's location.

Figure 6-22 illustrates how Cisco Context-Aware Services can be used with Cisco WCS to display the current location of employees equipped with Cisco 7925G IP phones, laptops, and PDAs. Note that we have chosen to assign and display the user name associated with each device, rather than the device MAC address. Clicking on any of the blue WLAN client icons shown in Figure 6-22 displays a plethora of information about the client device, including its client properties, association history, connection troubleshooting information, as well as its event history.

Figure 6-22 WLAN Client Tracking

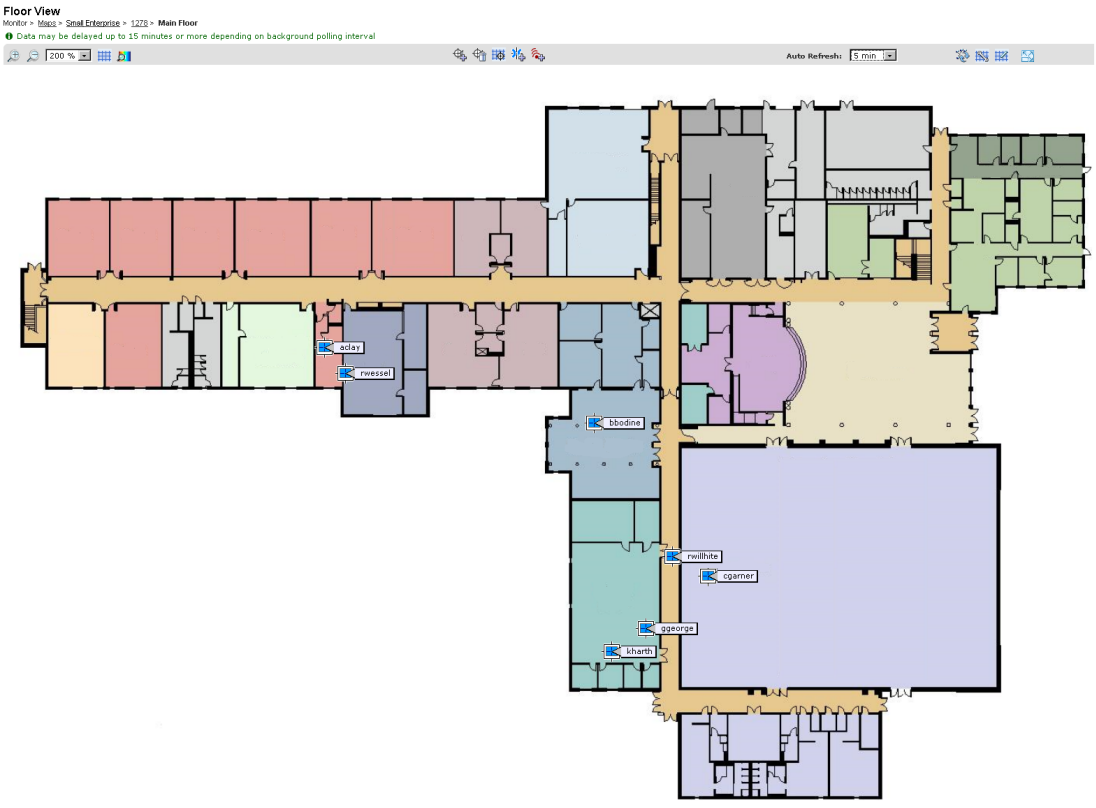

If the intention is to try and track a wireless device for inventory or loss prevention purposes, it is important to ascertain whether the embedded network interface remains active for any devices that may be in a powered-down or standby state. Even so, if the device remains active but accesses the network only very infrequently in order to conserve power, the degree of location fidelity obtained will likely be less than optimal, especially if the device can be expected to move frequently while in this state. A higher accuracy solution for this type of full-time location tracking application would be to attach or embed a RFID asset tag (see the "RFID Asset Tag Context-Aware Considerations" section on page 6-41) to the device. RFID tags are independently powered and operate without concern for the power state of the WLAN client device, and can be stimulated to transmit immediately when passing in proximity of properly equipped doorways and exits.

In general, best results are obtained in context-aware designs when all WLAN clients are compliant with the Cisco Compatible Extensions for WLAN Clients specification v2 at minimum, and preferably with the Cisco Compatible Extensions for WLAN Clients specification v4 or v5[5].

Additional information on the Cisco Compatible Extensions program for Wi-Fi Client Devices is available at http://www.cisco.com/web/partners/pr46/pr147/partners_pgm_concept_home.html.

Note that this document does not detail the steps involved with important procedures such as calibration of the Context-Aware Engine for WLAN Clients, or the definition of location inclusion and exclusion regions and other deployment procedures. For information on these and other procedures that should be understood prior to deployment, refer to the *MSE Context Aware Service Deployment Guide* (http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml).

5. Readers interested in the technical details concerning compatibility with the Cisco Compatible Extensions for WLAN Clients specification v2 are referred to the section titled "Tracking Clients, Assets and Rogue Devices" in the *Wi-Fi Location-Based Services Design Guide* 4.1, http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/wifich3.html#wp1049277

Intel® Wi-Fi Clients and ProSet Client Supplicant Context-Aware Considerations

When using clients equipped with the Intel® Wireless WiFi Link 4965AGN, Intel® PRO/Wireless 3945ABG Network Connection, or the Intel® PRO/Wireless 2915ABG Network Connection adapter, it is important to note that the default "Personal Security" settings of the Intel® ProSet Configuration Utility does not include compatibility with the Cisco Compatible Extensions specification. In order to enable compatibility with the Cisco Compatible Extensions specification, the Intel ProSet client supplicant must be used to reconfigure the client for "Enterprise Security" and to enable Cisco Compatible Extensions using the "Cisco Options" (Figure 6-23).

Figure 6-23 Enabling Cisco Compatible Extensions on Intel® ProSet Clients

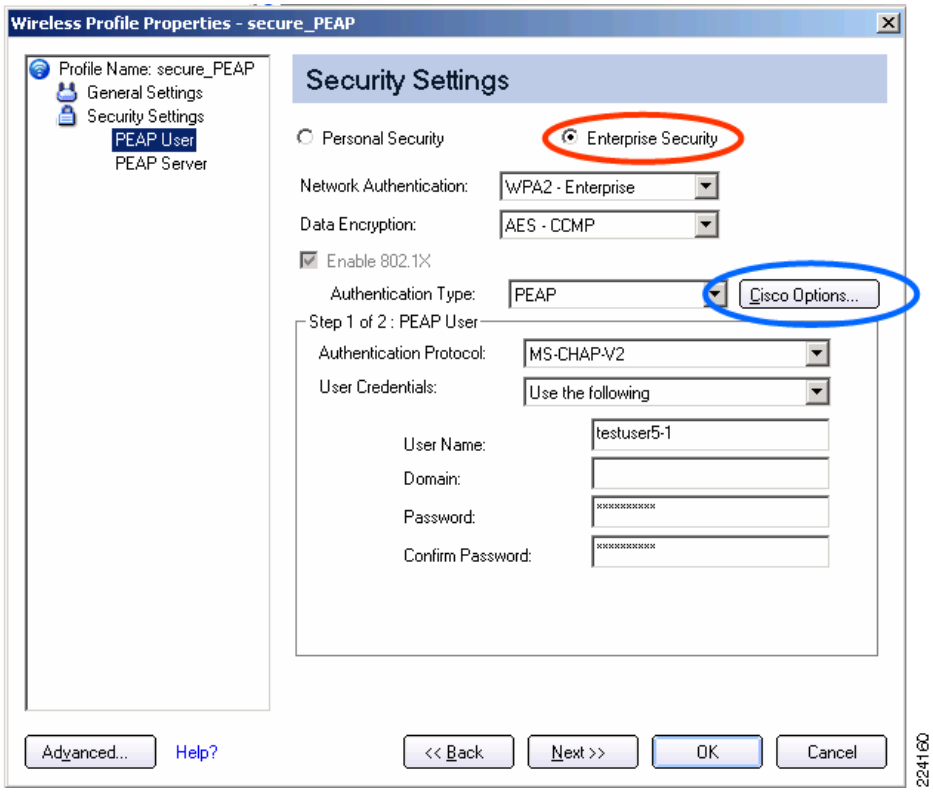

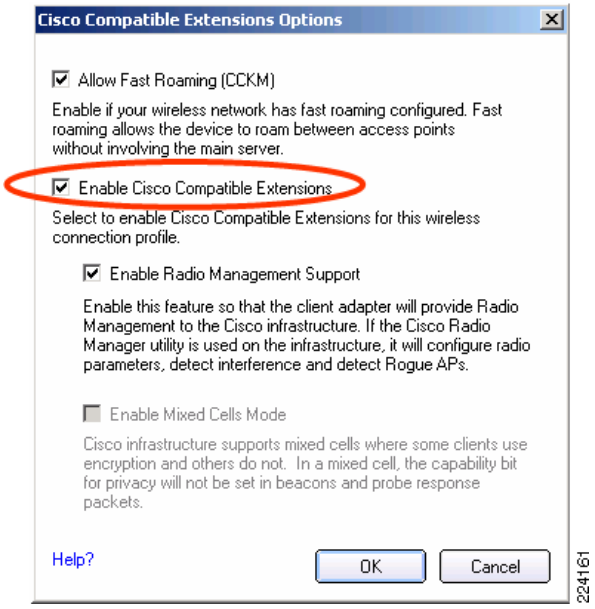

Cisco 7921 and 7925G Context-Aware Considerations

7921G and 7925G Unified IP Wireless Phone users should note that phones that are idle and not currently participating in an active call may not transmit 802.11 Probe Requests with sufficient frequency to ensure that changes in the actual location of the phone user are promptly reflected in the calculated location coordinates provided to the contest-aware services software in the MSE. In some cases this can be of concern, especially if the 7921G or 7925G user remains within the primary coverage area of the same access point for long periods of time. In cases where an access point might have a large coverage footprint, a roaming event may not occur very often despite a significant change in user location.

If you experience situations where the location fidelity of a 7921G or 7925G Unified IP Wireless Phone appears to be much better for those users actively participating on a call versus those that are on hook and idle, you may wish to consider changing the scan mode parameter associated with the 7921G or 7925G device in the Cisco Unified Communications Manager. In some cases, improved location fidelity can be achieved by enabling the “continuous” scan mode on the Device=>Phone configuration page of Cisco Unified Communications Manager Administration (shown in Figure 6-24). Note that scan mode options listed are “auto”, “continuous”, and “single AP”, where auto is the default. “Continuous” scan mode causes the wireless IP phone to issue a probe request approximately every two seconds, whereas “auto” scan mode causes the device to issue probe requests primarily only when the device is engaged on an active call, when roaming, or when preparing to roam. “Single AP” is used in installations where the wireless IP phone is only used in the vicinity of a single access point at all times, as probe requests are issued only when the wireless IP phone is first powered on. “Single AP” mode is not applicable and should not be used in the small enterprise designs discussed here.

Figure 6-24 Scan Mode Option in Cisco Unified Communications Manager

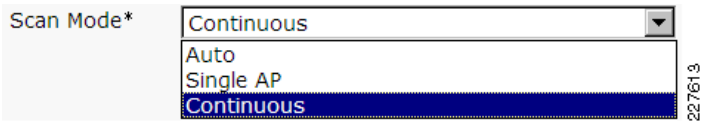

Note It is recommended that continuous scan mode be used only in situations where the anomaly described here is actually witnessed. This is because a trade-off associated with any increase in the frequency of probe requests transmitted is the potential for reduced 7921G or 7925G battery life. If you do not notice the anomaly described in this section, it is recommended that you leave the CUCM scan mode setting at the "auto" default.

RFID Asset Tag Context-Aware Considerations

Radio Frequency Identification (RFID) has many potential safety and security applications in the enterprise arena. Already used to improve efficiency and productivity across many business sectors, there are a wide variety of applicable use cases that embrace RFID in the enterprise in order to address a plethora of challenges.

Figure 6-25 Floor Map Showing Various Uses for RFID Tags and Cisco Context-Aware Services

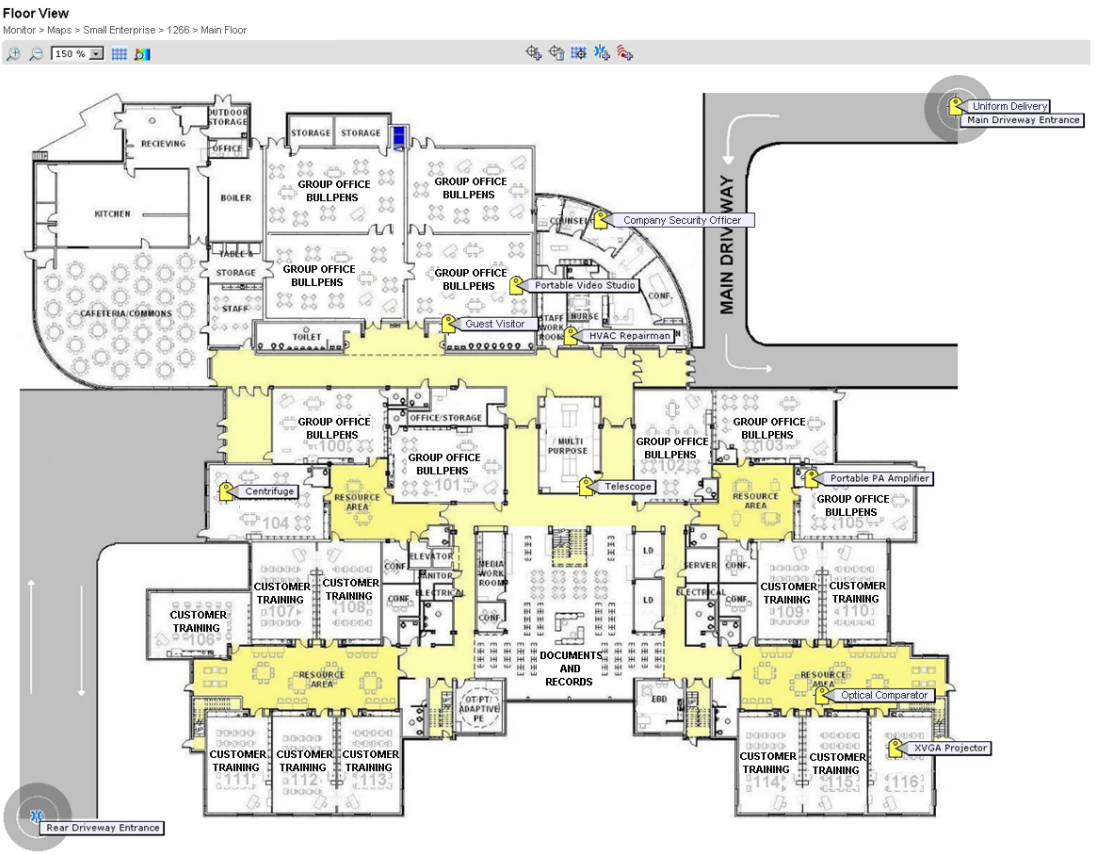

The illustration in Figure 6-25 provides us with a visual representation of just some of the ways RFID can be used in an enterprise in conjunction with Context-Aware Services:

- High value assets (such as projectors, audio/video equipment, etc.) can be kept safe and secure, since the application of RFID tags to these assets provides small enterprise administrators and security staff with the ability to locate the assets quickly and efficiently. Figure 6-25 illustrates how the present location of assets equipped in this fashion can be quickly ascertained by a quick check of the small enterprise floor map. Here, we can see the last known location of XVGA portable projectors and other equipment belonging to different departments at this enterprise site, such as sales, marketing, R&D, engineering, customer service, and technical support. For buildings containing more than a single floor, database search techniques can be used to search for the desired asset by name or attached RFID tag MAC address.

Figure 6-26 ID Badge with Embedded Active RFID

- Employees, visitors, guests and administrators can wear specially manufactured badges (see Figure 6-26) that combine a traditional identification card with active RFID technology, allowing them to be located quickly in the event of an emergency. These same devices can also transmit special notifications using a push button sequence, which can be interpreted in various ways by Context-Aware Services, including as a signal that an emergency event is in progress. Figure 6-25 illustrates how a person can be located if they are carrying an RFID-enabled ID card, as seen in the displayed location of the site's security officer, whose location is currently indicated as being inside of the site manager's office. Being able to physically locate the site security officer using a tool such as this could help in saving precious seconds in the event of an emergency.
- The whereabouts of enterprise site visitors and guests (such as maintenance or repair contractors) can be monitored and alerts triggered to security officers or others if these visitors stray into areas that they do not have authorization for. For example, in Figure 6-25, we see that we have an HVAC repair contractor as well as a guest visitor that has come in for an employment interview. In addition to making us aware as to the presence of these visitors, our system can alert us if these personnel access areas that we do not want them to by comparing their current locations to location boundaries we have defined as notification rules. Third-party applications that can access context-aware information from the MSE may perform more advanced tasks as well.
- RFID tag technology can be combined with chokepoint triggers (Exciters) to enable the use of proximity applications. In this fashion, chokepoint triggers can serve many purposes, including stimulating asset tags affixed to assets that might be in the process of being removed from the site without authorization. This makes it possible to notify officials or security officers of such action, so as to act quickly and determine whether such movement is legitimate.

In addition to simply displaying the location of assets, Cisco Context-Aware Services makes other contextual characteristics of the asset and its environment available to applications accessing the MSE via its SOAP/XML API. For instance, if the asset tags used contain on-board temperature sensors, we can also read the current temperature surrounding the asset tag via the MSE. This can be seen in Figure 6-27, where we see from WCS that the ambient temperature surrounding our HVAC repairman is a comfortable 20 degrees Celsius (68 degrees Fahrenheit). From this information, it certainly appears that the site air conditioning system is doing its job! It is important to note that while WCS is used to view this information in Figure 6-27, any authorized context-aware application that is written to use the MSE SOAP/XML API could have accessed this data as well.

Figure 6-27 Rogue AP Details

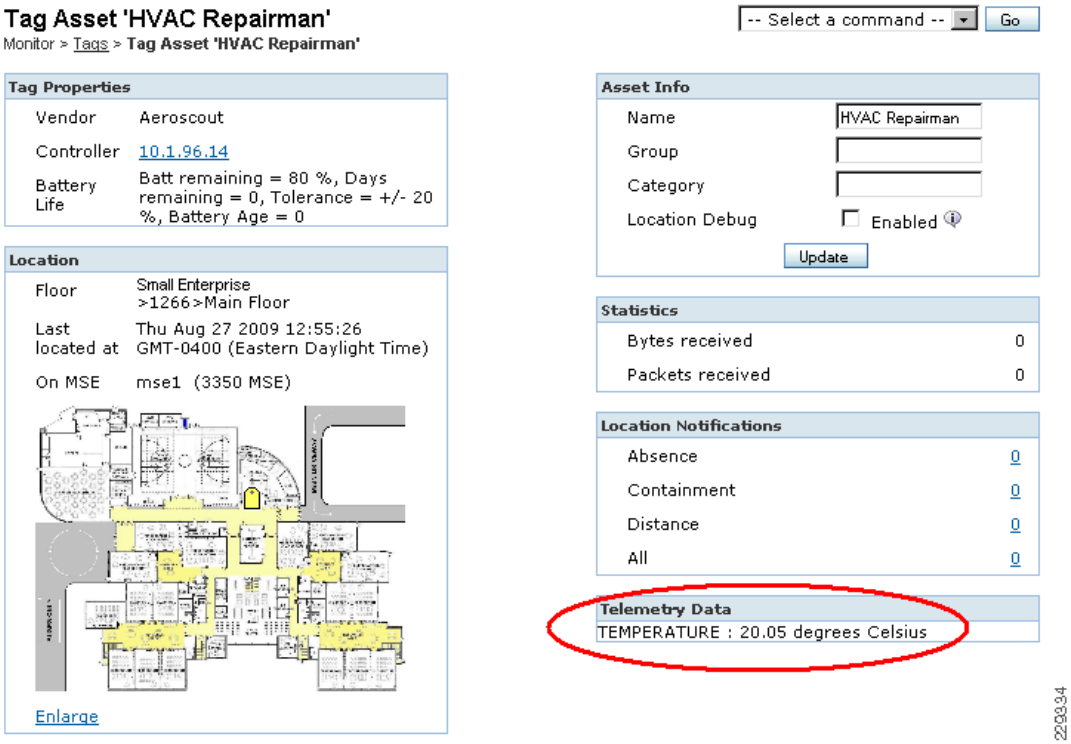

Simply put, RFID technology and Cisco Context-Aware Services can help small enterprises improve their overall safety, security, and efficiency.

Cisco Context-Aware Software is designed to function with active RFID asset tags from vendors compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. A list of current Cisco Compatible Extensions for Wi-Fi Tags compliant vendors can be found at http://www.cisco.com/web/partners/pr46/pr147/ccx_wifi_tags.html. Although all vendor RFID tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification share a great deal of functionality in common, the parameter names and means used to configure each brand of tags can differ from vendor to vendor, therefore no set prescribed configuration parameter list would apply to all. That being said, there are several general configuration functions that tag vendors share and although the exact parameter names may differ, it is important that you understand them and use this knowledge accordingly when configuring the particular tags of choice for your installation. More information regarding the configuration procedure for AeroScout tags can be found in the section entitled RFID Tag

and WLC Configuration/Tuning in the Cisco MSE Context-Aware Deployment Guide (http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#rfid-wlc)[6].

- *RF Channel Configuration*—It is recommended that tags be configured for the standard set of 2.4 Ghz non-overlapping channels, which is typically channels 1, 6, and 11 (this may vary depending on your international regulatory domain).
- *Stationary Transmission Interval*—This is the time between periodic tag transmissions that are normally generated when the tag is stationary and not in motion. It is recommended that this be configured for values between 3 and 5 minutes.
- *Motion Transmission Interval*—(Applies only to tags with motion sensors.) This is the time between periodic tag transmissions generated when the tag and asset is in motion. A recommended initial value is 15 seconds.
- *Number of Tag Message Repetitions*—Some popular RFID tags default to transmitting a single transmission on all defined channels. Tag parameters that control the number of tag message repetitions specify the number of times each transmitted message is repeated, per channel. It is generally recommended that this parameter be set to a value of three. Doing this helps protect against lost tag transmissions due to congestion or interference, which is a primary cause of poor tag location accuracy.
- *Message Repetitions Interval*—The delay between subsequent message repetitions on the same channel. This is often defaulted to 512msec, although in lab testing we have seen some evidence of improved location accuracy when a message repetition interval of 256 msec is used with the default controller value of 2 seconds for NMSP notification interval. This parameter is discussed in more detail the section entitled RFID Tag and WLC Configuration/Tuning of the Cisco MSE Context-Aware Deployment Guide (http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#rfid-wlc).

This chapter does not detail the steps involved with procedures such as calibration of the Context-Aware Engine for Tags and other deployment procedures. For information on these and other procedures that should be understood prior to deployment, refer to the MSE Context Aware Service Deployment Guide (http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml). In addition, the AeroScout Context-Aware Engine for Tags for the Cisco MSE Users Guide, version 3.2, available from your AeroScout representative or https://support.aeroscout.com, is highly recommended.

RFID Tag Chokepoint Trigger Considerations

Chokepoint triggers are proximity communication devices that trigger RFID asset tags to alter their behavior when the tag enters into close range (otherwise known as the stimulation zone) of the chokepoint trigger. This behavioral modification may, for example, cause the RFID tag to immediately transmit its unique identifier (MAC address) or cause the tag to change its internal configuration, depending on how the tag is programmed. A very popular use of the chokepoint trigger is to stimulate the asset tag such that it provides indication to the MSE that the tag has entered or exited a given area, known as a chokepoint. Chokepoints are entry or exit points that provide passage between connected regions. Common chokepoints are entrances and exits such as doorways, hallways, and stairwells.

Note An Exciter is a registered trademark of AeroScout Ltd., and represent a popular example of a chokepoint trigger.

6. Additional information can also be found in the *Wi-Fi Location-Based Services Design Guide 4.1* in the section titled “Configuring Asset Tags”, located at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/wifich6.html#wp1077248.

Chokepoint triggers are useful in causing RFID tags to react quickly when assets are moved past certain points. This could be a server, for example, with an affixed RFID tag moving past a chokepoint trigger located near a facility exit, an employee with an RFID badge entering the site through the front entrance at 7:00 AM, or it could be a delivery vehicle that has a RFID tag on the front visor coming past a chokepoint trigger located at the site's property entrance. In all these cases, the chokepoint trigger causes the tag to change its behavior, most likely to immediately transmit its MAC address (as well as the MAC address of the chokepoint trigger that stimulated it) to the MSE via one or more access points that can receive the tag's transmissions. The net result is to cause the MSE to indicate that the current location of the asset and its asset tag is within a known, pre-determined proximity of the chokepoint trigger.

In order to use chokepoint triggers with Cisco Context-Aware Services, they must be properly configured using the appropriate vendor-supplied software utility, defined to WCS, placed on floor maps, and synchronized as part of an updated network design to the MSE. After all of this is complete, the MSE is able to recognize the transmissions generated by asset tags that have been stimulated by specific chokepoint trigger MAC addresses. Based on this information, the MSE can attempt to localize the tag to the proximity of the chokepoint trigger. Applications such as WCS (or third party context-aware applications) may then display the asset tag's location at the chokepoint icon associated with the chokepoint trigger's MAC address.

Various chokepoint trigger specific parameters such as transmission range, IP address, transmission interval, transmission repetitions, and so on are set using vendor-specific utilities. Note that each vendor maintains their set of software tools necessary for configuration of their chokepoint triggers. These software configuration tools are not interoperable between vendors (for example, AeroScout software configuration tools cannot be used to configure WhereNet chokepoint triggers or vice-versa).

The individual configuration of each vendor's chokepoint trigger is beyond the scope of this chapter. Complete and detailed configuration information relating to the specific configuration of each vendor's chokepoint trigger can be found in the appropriate vendor's documentation, which can be obtained from your tag vendor representative.

AeroScout EX-3200 User Guide https://support.aeroscout.com

AeroScout Exciter EX-2000 User Guide https://support.aeroscout.com

AeroScout Context-Aware Engine for Tags for the Cisco MSE Users Guide, version 3.2 https://support.aeroscout.com

Technical documentation for WhereNet WherePort chokepoint triggers and the necessary software and hardware for configuration of WherePorts is available from WhereNet Corporation (http://www.wherenet.com) or via your WhereNet account representative.

Rogue Device Context-Aware Considerations

The use of context-aware services for locating clients and RFID tags that we have defined and authorized in the small enterprise environment is often what first comes to mind for many of us when we consider this solution. However, another equally important and useful function of context aware services is the ability to detect the location associated with those wireless clients and access points that we have not authorized to operate within our domain. In other words, context-aware services in our small enterprise can help us in locating rogue access points or rogue clients that may have been installed by employees, contractors, visitors, vendors, or even administration members without authorization. Even if an unauthorized access point is innocently installed by an enterprise user that is otherwise authorized to use the network, the unauthorized and potentially insecure portal provided by such an access point can unnecessarily expose our otherwise secure small enterprise network to outside intruders.

The Cisco Wireless Control System can use the location capabilities provided by the Cisco Mobility Services Engine and the Cisco Context-Aware Engine for Clients to define the location of unauthorized access points and the wireless clients that may be using these access points, as shown in Figure 6-28. In

Figure 6-28, we can see icons for both rogue access points as well as rogue clients displayed over their predicted positions on a floor map of a site located within our small enterprise. This capability is very useful both to the local site administrator, as well as the main site administrator and their technical teams. It allows them to determine the location of wireless equipment that may have been brought into the enterprise site and used in an attempt to gain unauthorized access to enterprise site computing resources. In Figure 6-28, the round icon with the "skull and crossbones" logo represents the rogue access point and the rectangular icon with the same logo represents a rogue client that is associated to this rogue access point.

Figure 6-28 Using Context-Aware Services to Determine Location of Rogue Access Points and Clients

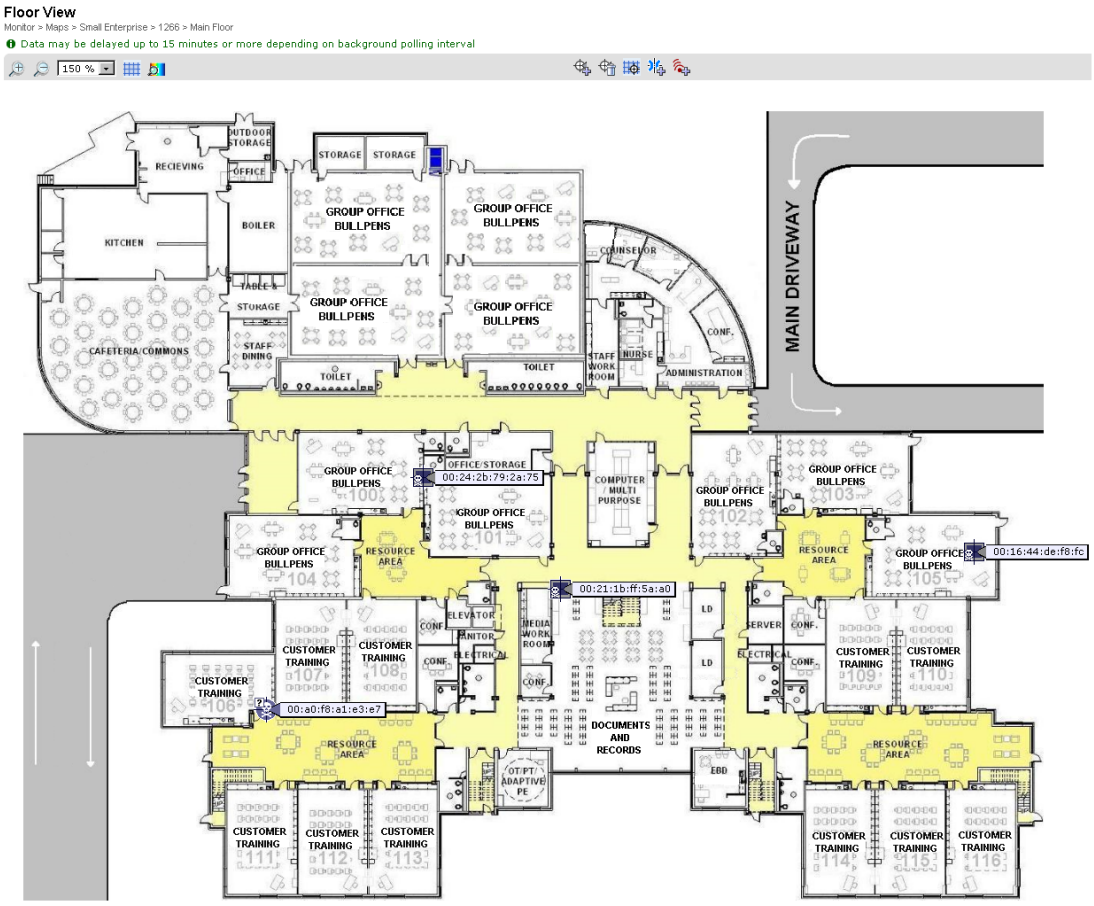

Note In comparison to WLAN clients and RFID tags, it is normal to experience a reduced accuracy when localizing rogue access points and rogue clients. The very nature of "rogue" devices establishes that these devices are not under the control of site administration. Therefore, the configuration of such these rogue devices may not facilitate optimal context-aware location accuracy (for example, they may not be Cisco Compatible Extensions devices, etc.).

As shown in Figure 6-29, clicking on either the rogue access point icon or the rogue client icon reveals additional information and capabilities that can help in further identifying (and even isolating) these devices. This includes the ability to perform switch port tracing with the latest versions of WCS, which

allows tracing of rogue access points that have been connected to the switch infrastructure (for more information on switch port tracing, refer to the *WCS Configuration Guide* 6.0 at the following URL: (http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0ctrlcfg.html#wp1089752)

Figure 6-29 Rogue AP Details

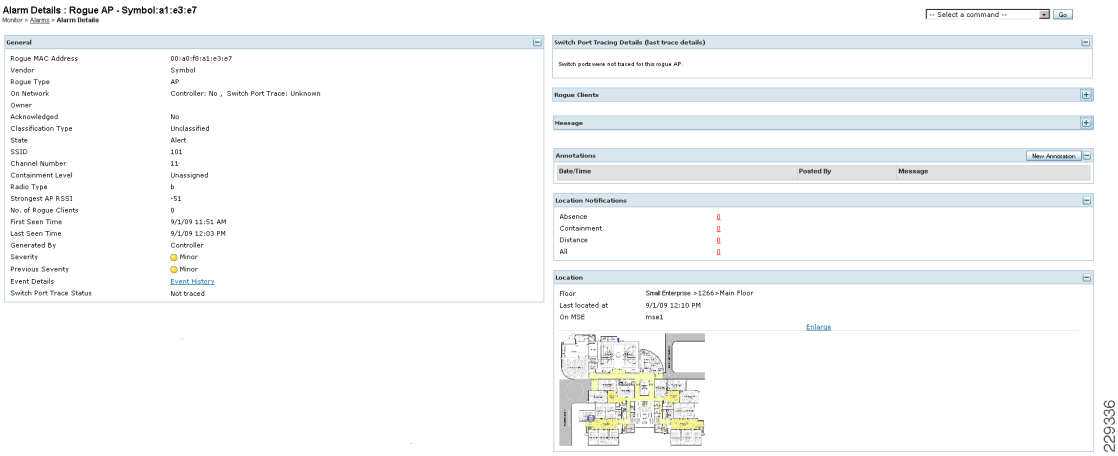

The collection of rogue AP and rogue client information is enabled by default on WLAN controllers. In order to enable the tracking of rogue access points and rogue clients by the context-aware service, it is important to ensure that the collection of rogue information has also been enabled for context-aware services on the MSE. Refer back to Figure 6-18 in this chapter and ensure that the "Rogue Client and Access Points" tracking parameter check box has been enabled under **Mobility Services > Context-Aware Services > Administration >Tracking Parameters**.

Note Be advised that depending on the environment in which your enterprise site is located, as well as the number of rogue devices present, the number of rogue devices detected can rise very quickly. Since the size of the rogue device population is typically not under the direct control of enterprise IT staff or local site administration, it is highly advisable that you enable limiting for rogue clients and access points and set a limiting value. This is so that any unforeseen increase in the number of rogue devices detected does not consume all the remaining tracked device capacity on the MSE, thereby depriving the MSE of capacity that might be of more significance to enterprise site administration and employees. This is especially important if the MSE is used to service multiple sites within your small enterprise.

In addition to enabling the collection of rogue access point and rogue client information by the context-aware services on the MSE, you must also enable the display rogue device location when displaying floor maps using WCS. This is accomplished via the WCS "Floor Settings" submenu that is displayed on every WCS floor map, as shown in Figure 6-30. Information on how to use the Floor Settings sub-menu for all the categories of device shown here can be found in Chapter 5 of the *WCS Configuration Guide* 6.0, http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0maps.html#wp1210969.

Figure 6-30 WCS Floor Settings Sub-Menu

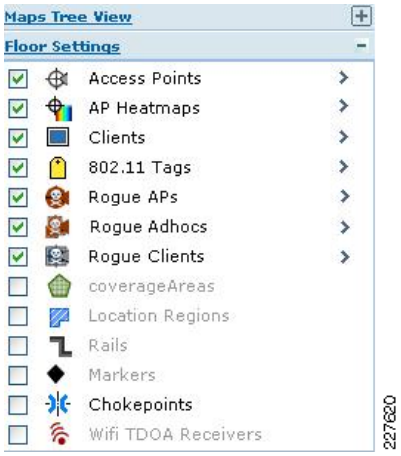

For further information regarding rogue access points and clients, refer to the following documents:

- *Context Aware Service Configuration Guide,* Rlease 6.0
  http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/CAS_60.html
- *Cisco Wireless Control System Configuration Guide*, Release 6.0
  http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/WCS60cg.html
- *Cisco MSE Context-Aware Deployment Guide*
  http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml

Context-Aware Considerations for Wired Device Tracking

As described previously in this chapter, beginning with release 6.0 Cisco Context-Aware Services provides the capability to determine the civic location and emergency line identifiers of devices connected to Cisco Catalyst switches, such as the 2960G, 3560E, 3750E, 3750G, 4500, and 4900 series. As participants in context-aware services, switches are configured with and provide the relevant contextual information for all the IP endpoints attached to them. These endpoints may include IP phones, PCs, host servers, access points, etc. The NMSP protocol is used between the switches and MSE to deliver this contextual information to the MSE. Location information may include the physical location address (also known as the civic address) as well as other information about endpoints such as the IP address, MAC address, port, VLAN, and username. If the end device makes use of the Cisco Discovery Protocol or Link Layer Discovery Protocol (LLDP), additional information, such as the version number and serial number, can also be sent to the MSE.

Note The use of Context-Aware Services for wired device location is entirely optional. In the Small Enterprise Design Profile, Context-Aware Services may be deployed for wireless devices, wired devices, or for both.

The Small Enterprise Design Profile provides that in the event of a failure of a switch line card in a 4500 series context-aware Catalyst switch, or a stack member in the 3750 switch stack, NMSP sessions recover from the failure without intervention from the user. NMSP session recovery time from isolated switch stack member or line card failure was observed to be very quick during lab testing, and in most cases the recovery time was almost unnoticeable from the perspective of the Mobility Services Engine. Careful examination of the NMSP session status during simulated failures indicated that sessions remained up, intact, and passing NMSP data while stack member or line card hand-off occurred.

While still a relatively new context-aware capability, the inclusion of wired device tracking provides new visibility into the location of wired IP endpoints in your small enterprise network. For example, some of the ways that this new exciting capability can be used in small enterprise designs include:

- Determining the whereabouts of missing IP phones, PCs, and peripheral devices such as printers and network scanners that have disconnected and moved from their originally installed locations to other locations within a site (or moved to another site altogether). Once reconnected to the network, the MSE would be updated with the wired device's new attachment information and any location and/or ELIN information defined for that switch port.
- Keeping track of the location of host computers that are physically present at the main site data center or in any other site. The wired location capabilities contained in release 6.0 of Context-Aware Services make it possible to specify the location of a device down to the room, cubicle, seat, or even rack/slot location (see Figure 6-31). This can be important in verifying that equipment to be de-commissioned is actually removed from service and sanitized of all corporate and employee information.

Figure 6-31 Displaying the Location of a Wired Host

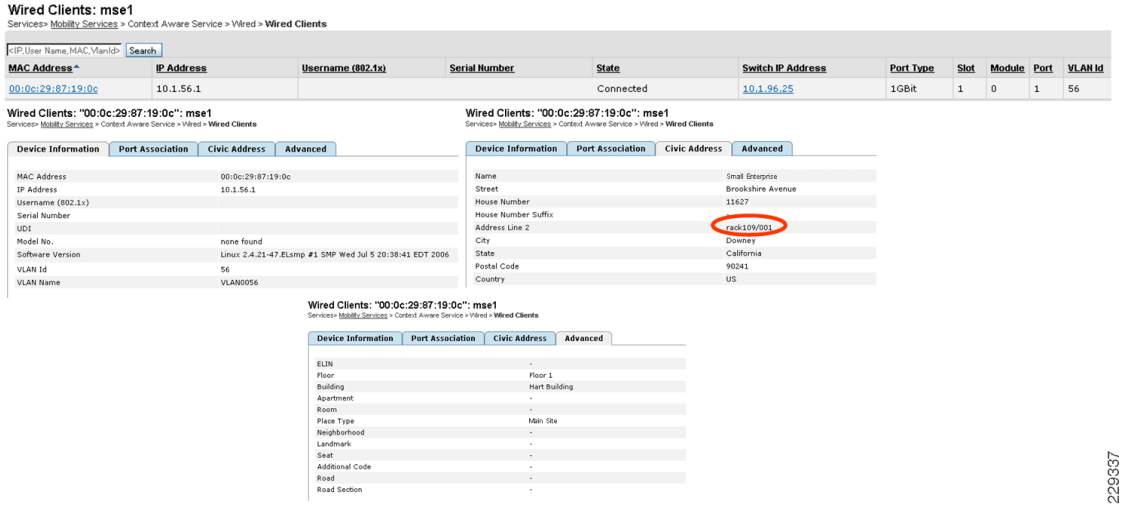

- Determining the civic location of users based on their IP address or the user name that was specified during 802.1x / EAP login (refer to Figure 6-32). Context-Aware Services for wired devices makes it possible to quickly determine the civic or emergency line identifier information associated with the switch port. Figure 6-32 illustrates how we can search for the wired device by the known username.

Figure 6-32 Searching for Wired Device by Username

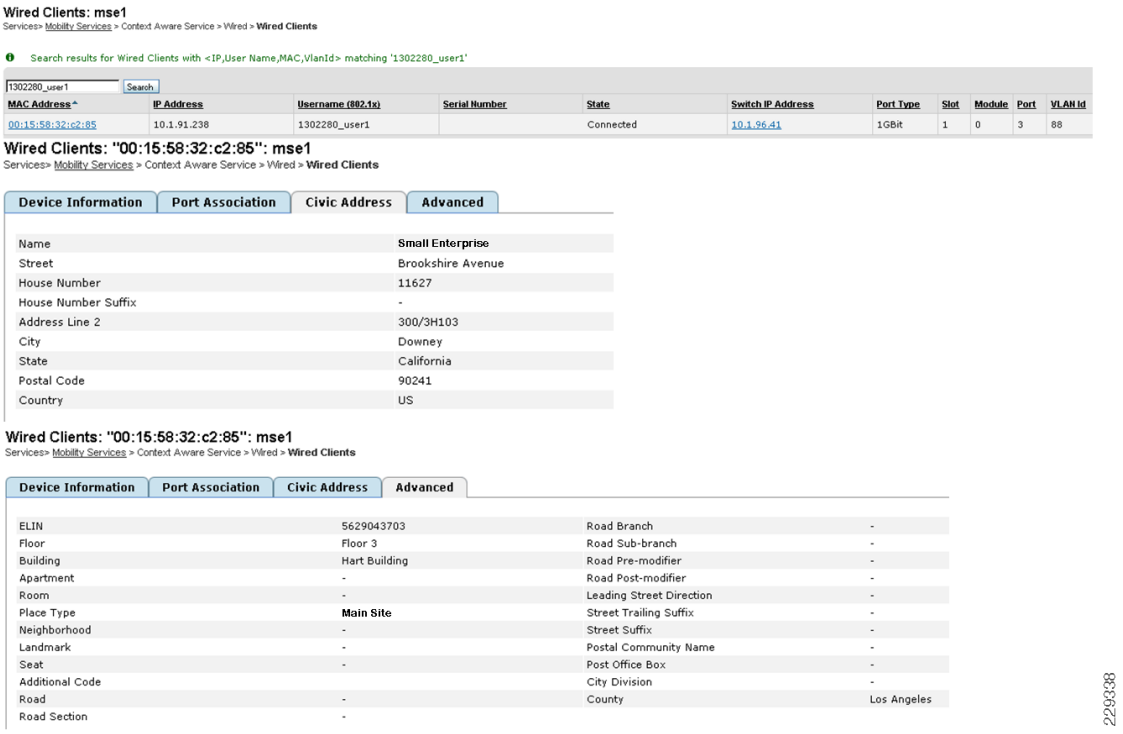

- For devices that support it, impromptu inventory checks of devices across the main site by examining wired device listings by serial number (see Figure 6-33) and comparing this information to deployment records. This can help enterprise and site administrators better determine whether assets have been moved between enterprise sites without authorization.

Figure 6-33 Examining Device Serial Numbers Via Wired Device Tracking

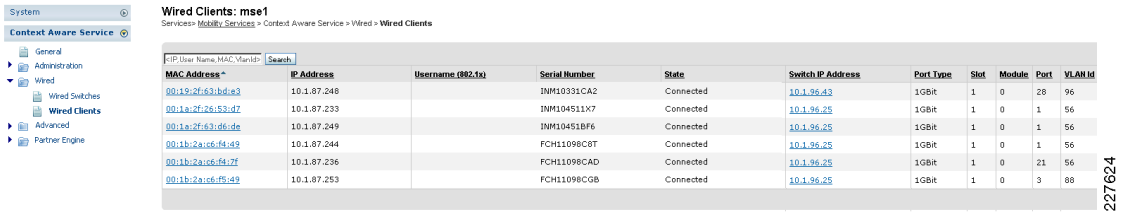

In contrast to the manner in which searches for wireless clients and tags is handled via the **WCS Monitor > Maps**, in Release 6.0 of Context-Aware Services, all wired client searches are handled via the **Services > Mobility Services >Context Aware Services > Wired > Wired Clients** menu panel. Searches using the wired client WCS menu panel are specific to the particular MSE that you have selected.

Hardware and Software Requirements for Wired Device Tracking

As mentioned previously, at the current time wired device tracking is only performed on Catalyst switch hardware supporting context-aware services.

Note The images used for testing of wired device tracking during the production of this "Hardware/Software Releases" section on page 6-60. Readers should note that cryptography-enabled (k9) switch images are required in order to enable wired device tracking in Catalyst switches.

- You should be aware that including wired device tracking in your design will require additional tracked device licenses on the MSE over and above that required for wireless device tracking alone. This is because wired tracked devices are included in the maximum number of simultaneous devices that can are licensed for tracking by an MSE. For example, a small enterprise that would otherwise possess a maximum licensed tracked device requirement of 500 for wireless LAN clients, RFID tags, and rogues might require 750 or more when wired device tracking is considered, depending on the number of switches deployed and whether wired device tracking is enabled for all of them. Be sure to plan appropriately for the number of wired devices that you intend to track when purchasing MSE client tracking licenses for the Cisco Context-Aware Engine for Clients.

In addition to planning appropriately for an increase in MSE tracked device licensing due to the use of wired device tracking, keep in mind that although a single MSE can technically support up to 500 NMSP sessions, scalability testing limits have only allowed Cisco to test up to 100 simulated NMSP connections to a single MSE at this time. Each context-aware switch that is enabled for wired device tracking in your network establishes one NMSP session to the MSE and counts against this limit. Therefore, when enabling many switches for context-aware wired device tracking, it is recommended that you plan for the total number of MSEs that may be required to support the total number of NMSP sessions in your network.

Enabling Context-Aware Wired Device Tracking

In order to track wired devices on Catalyst switch ports, each switch whose devices we wish to track must be configured to enable NMSP and other important parameters, and to contain the appropriate civic address and ELIN location information. In addition, WCS must be configured to be aware of the context-aware switches in the network and to be able to communicate with them. WCS is also used to transmit information about the switches to the Mobility Services Engine via the synchronization process.

A complete, step by step guide to configuring Catalyst switches and WCS for wired device tracking can be found in the *Context-Aware Service Configuration Guide* 6.0 http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1224011.

In addition, the following chapters and documents provide valuable and detailed background information concerning the wired device tracking capability of Catalyst switches:

- "Configuring LLDP, LLDP-MED, and Wired Location Service" in the *Catalyst 3750 Switch Software Configuration Guide*, 12.2(50)SE http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2_50_se/configuration/guide/swlldp.html
- "Configuring LLDP, LLDP-MED, and Wired Location Service" in the *Catalyst 2960 Switch Software Configuration Guide*, 12.2(50)SE http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/swlldp.html
- "Configuring LLDP and LLDP-MED" in the *Catalyst 4500 Series Switch Software Configuration Guide*, 12.2(53)SG http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/53SG/configuration/swlldp.html#wp1097119

Readers are reminded about the following:

- NMSP is disabled on switches by default and must be explicitly enabled via the nmsp enable global configuration command.
- IP device tracking must be enabled on the switch in order for Context-Aware wired device tracking to function properly. It can be enabled by issuing the ip device tracking command in global configuration mode on the context-aware switch.
- The civic location identifier in the LLDP-MED TLV is limited to 250 bytes or less. To avoid receiving error messages regarding available buffer space during switch configuration, the total length of all civic location information specified for each civic-location identifier must not exceed 250 bytes.
- In Release 6.0, all wired device client and switch tracking is available only to the root WCS virtual domain user. Because of this, you may wish to limit the use of context-aware wired device tracking to only those users with whom you are comfortable assigning WCS root virtual domain privileges.

NMSP Attachment Notification Interval

After an NMSP session is established between the MSE and a context-aware Catalyst switch, the MSE transmits an Echo Response packet to the switch every echo interval period, which is specified on the MSE (for all NMSP session partners) using the WCS menu entitled Services > Mobility Services > System > NMSP Parameters. In addition to Echo Responses, the switch will periodically send attachment notifications to the MSE via the NMSP session. Any link-up or link-down events that are detected by the switch are aggregated during a configurable time interval and sent to the MSE via an attachment notification at the conclusion of that time interval.

This interval is known as the nmsp attachment notification interval and is configurable on the switch via the **nmsp notification interval attachment** *interval-seconds* command. The range of values for interval-seconds is from 1 to 30 seconds, with 30 seconds being the default. In large networks where there are many NMSP sessions active across the MAN to an MSE and the number of users connecting and disconnecting from the switch is high, configuring **nmsp attachment notification** *interval* to a very short interval can increase the amount of NMSP traffic between switches and the MSE and is not recommended[7] without carefully understanding the nature of the traffic present in your network.

Civic Address Configuration

The information contained in the *Context-Aware Service Configuration Guide* 6.0 (http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1224011) provides the information necessary to configure context-aware switches and the WCS for wired device tracking.

As can be seen in the Context Aware Service Configuration Guide, in release 6.0 of Context-Aware Services all configuration of civic and ELIN location information is performed on each context-aware switch using the switch CLI. Once a switch is configured with the desired civic and ELIN location information, the switch will share all of the configured port information with the MSE when the NMSP session is initially established and will periodically update the MSE if any location updates are performed.

Note Readers should find the following IETF RFC documents helpful in better understanding the types of values that should be specified for the various civic location fields: RFC 4776 (http://www.ietf.org/rfc/rfc4776.txt), RFC 4589 (http://www.ietf.org/rfc/rfc4589.txt), and RFC 5139 (http://www.ietf.org/rfc/rfc5139.txt).

7. Except for switches that are local to the Mobility Services Engine and need not traverse the MAN.

During the course of our lab testing, we discovered other useful facets of information regarding civic location and ELIN configuration in Catalyst switch IOS releases 12.2-52(SE) and 12.2-53(SG):

- *Global Scope of Civic and ELIN location*—Civic address or ELIN information must be configured at a global level and then assigned to each switch interface using the appropriate civic or ELIN location identifier. Globally defined civic address parameters (such as the building, county, postal code, floor, and so on) cannot be individually over-ridden at the interface level in release 6.0. If more than one switch port shares the same civic location or ELIN, then the same globally specified civic and ELIN location identifiers can be used on each switch port interface. However, if all ports possess unique civic address characteristics, then uniquely specified global civic address parameters for each port must be used. This can be seen in the following example where three unique civic location identifiers are applied to three different ports (Gi1/0/3 - 1/0/5). The test application server that is being tested on port Gi 1/0/6 is in the same physical location as the device connected to port Gi1/0/5, hence they share civic-location identifier 3.

```
location civic-location identifier 1
additional-code "Small Enterprise"
building "Hart Building"
city Downey
country US
county "Los Angeles"
floor "Floor 3"
name "Main Site"
postal-code 90241
state California
street-group "Brookshire Avenue"
number 11627
room 300
seat 3H103
type-of-place "Main Site"
!
location civic-location identifier 2
additional-code "Small Enterprise"
building "Hart Building"
city Downey
country US
county "Los Angeles"
floor "Floor 3"
name "Main Site"
postal-code 90241
state California
street-group "Brookshire Avenue"
number 11627
room 300
seat 3H104
type-of-place "Main Site"
!
location civic-location identifier 3
additional-code "Small Enterprise"
building "Hart Building"
city Downey
country US
county "Los Angeles"
floor "Floor 3"
name "Main Site"
postal-code 90241
state California
street-group "Brookshire Avenue"
number 11627
room 300
seat 3H105
type-of-place "Main Site"
```

```
!
interface GigabitEthernet0/3
description 802.1x data access only
location civic-location-id 1
switchport access vlan 88
switchport mode access
authentication port-control auto
dot1x pae authenticator
!
interface GigabitEthernet0/4
description 802.1x data access only
location civic-location-id 2
switchport access vlan 88
switchport mode access
authentication port-control auto
dot1x pae authenticator
!
interface GigabitEthernet0/5
description 802.1x data access only
location civic-location-id 3
switchport access vlan 88
switchport mode access
authentication port-control auto
dot1x pae authenticator
!
interface GigabitEthernet0/6
description local RFID gate application server
location civic-location-id 3
switchport access vlan 56
switchport mode access
!
```

- *Civic location* ***additional-location-information*** *subcommand*—During the course of our testing we found the **additional-location-information** subcommand useful in adding miscellaneous information about devices that can be displayed on the civic address tab of the WCS wired clients display under “Address Line 2". In our testing, we made use of this facility to label the rack and position location of various servers and other devices that are deployed in a common location. Figure 6-34 illustrates how the information entered using **additional-location-information** is then displayed on the civic address tab of the WCS wired clients display for the device (indicated by the red arrow). Below the configuration for the switch port, we can see the results of displaying the civic location for the switch interface using the l**ocation civic-location interface** *interface* command.

Figure 6-34 Additional-Location-Information

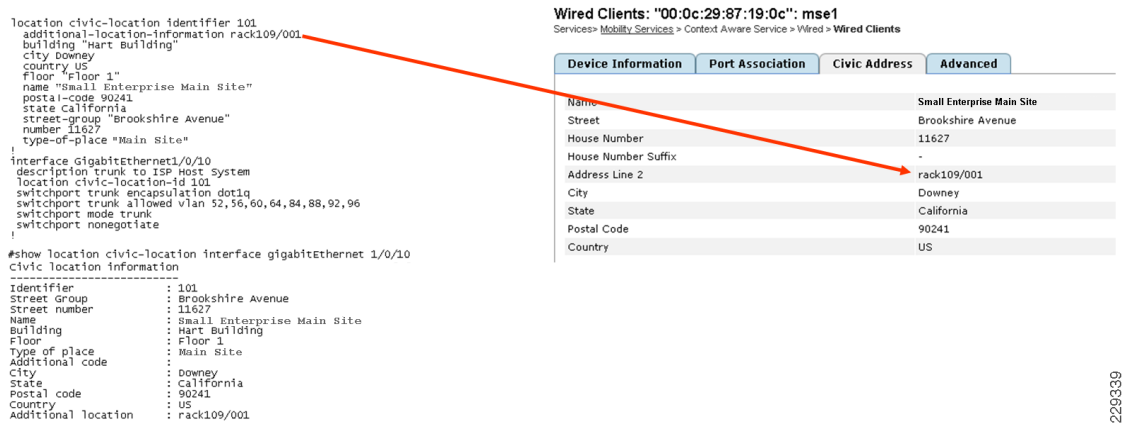

- *Civic Location* ***street-group*** *subcommand*—In order to display a value under the street name component on the civic address tab of the WCS wired clients display, we found it was necessary to use the civic location **street-group** subcommand in the switch configuration. If we look again at the left hand portion of Figure 6-34, we can see the street-group is specified as "Brookshire Avenue". On the right hand of Figure 6-34, we see that this value appears under the civic address tab of the WCS wired clients display under the label "Street".

Excluding Device Tracking on Select Switch Ports

In the majority of cases, when using context-aware wired device tracking, it is usually acceptable to simply enable NMSP on the switch and allow the attachment status of all ports to be reported to the MSE. In this way, the attachment status and any location or ELIN information specified for each port in the switch is reported and can be accessed from the MSE. However, in some cases, it might be desirable to enable wired device tracking on a switch, but exclude selected switch ports from reporting device attachments to the MSE. A reason for doing this might be to help reduce the number of MSE tracked device licenses required by eliminating the NMSP reporting of device attachments on select ports where not much chance of device migration is expected. Recall that in release 6.0, while it is possible to enable or disable wired device tracking entirely on a per MSE basis, it is not possible to limit the number of wired devices that are tracked in each MSE at this time (i.e., wired device tracking is either on or off). Therefore, any such limiting must be done manually by either disabling NMSP sessions with selected switches entirely (**no nmsp enable**) or disabling only the tracking of select switch ports on a context-aware switch that is otherwise reported device attachments normally.

To disable device tracking for select switch ports on a switch where NMSP has been enabled, the switch interface configuration **nmsp attachment suppress** command should be specified on each switch interface where device tracking is not desired. The nmsp attachment suppress interface command is used to configure the interface to not send any attachment notifications to a Cisco Mobility Services Engine (MSE).

If you are using Location MAC Filtering (**Services > Mobility Services > Context Aware Service > Administration> Filtering Parameters**) to specifically limit or block (by MAC address) tracked wireless clients and tags, be advised that these address filters also apply to wired device clients as well. Make sure that any filtering specifications that you set using Location MAC Filtering are flexible enough to allow tracking of not only your wireless clients and tags, but wired devices as well. Any devices that have been blocked from location tracking as a result of a defined filter will be viewable under the "Blocked MACs" listing on the Filtering Parameters page. Detailed information regarding how to configure Location MAC filtering can be found in Modifying Filtering Parameters section of the Cisco Context-Aware Service Configuration Guide 6.0 (http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1100062).

Classification and Marking of NMSP Sessions

A vital component in assuring NMSP session stability and acceptable CAS performance during periods of network congestion is the application of QoS to NMSP data flows between the MSE and any WLAN controllers or context-aware Cisco Catalyst switches in the Small Enterprise Design Profile. Classification, marking and QoS should be applied ideally in both cases of locally deployed as well as centralized MSE implementations, however, it is especially important when using a centralized MSE at the main site and remote WLAN controllers and context-aware Cisco Catalyst switches in the remote sites.

In order to ensure that QoS prioritization can occur properly for NMSP sessions the NMSP data flows must be properly identified, classified and marked as close as possible to their points of origin. This section explains where such marking of NMSP data flows should occur in the network, and how it should be performed.

Figure 6-35 provides an example of where identification, classification and marking of NMSP data flows should occur in the case of the small enterprise design containing:

- A main site with adjoining data center
- A remote site that is based on the Catalyst 4500 for distribution
- A remote site that is based on the Catalyst 3750 switch stack for distribution.

In Figure 6-35, we have identified several points where classification and remarking needs to occur in order to properly mark NMSP traffic inbound from WLAN controllers and context aware switches, and outbound from Mobility Services Engines. These points are indicated by the yellow and red numbered circles.

Figure 6-35 Points of NMSP Classification and Marking

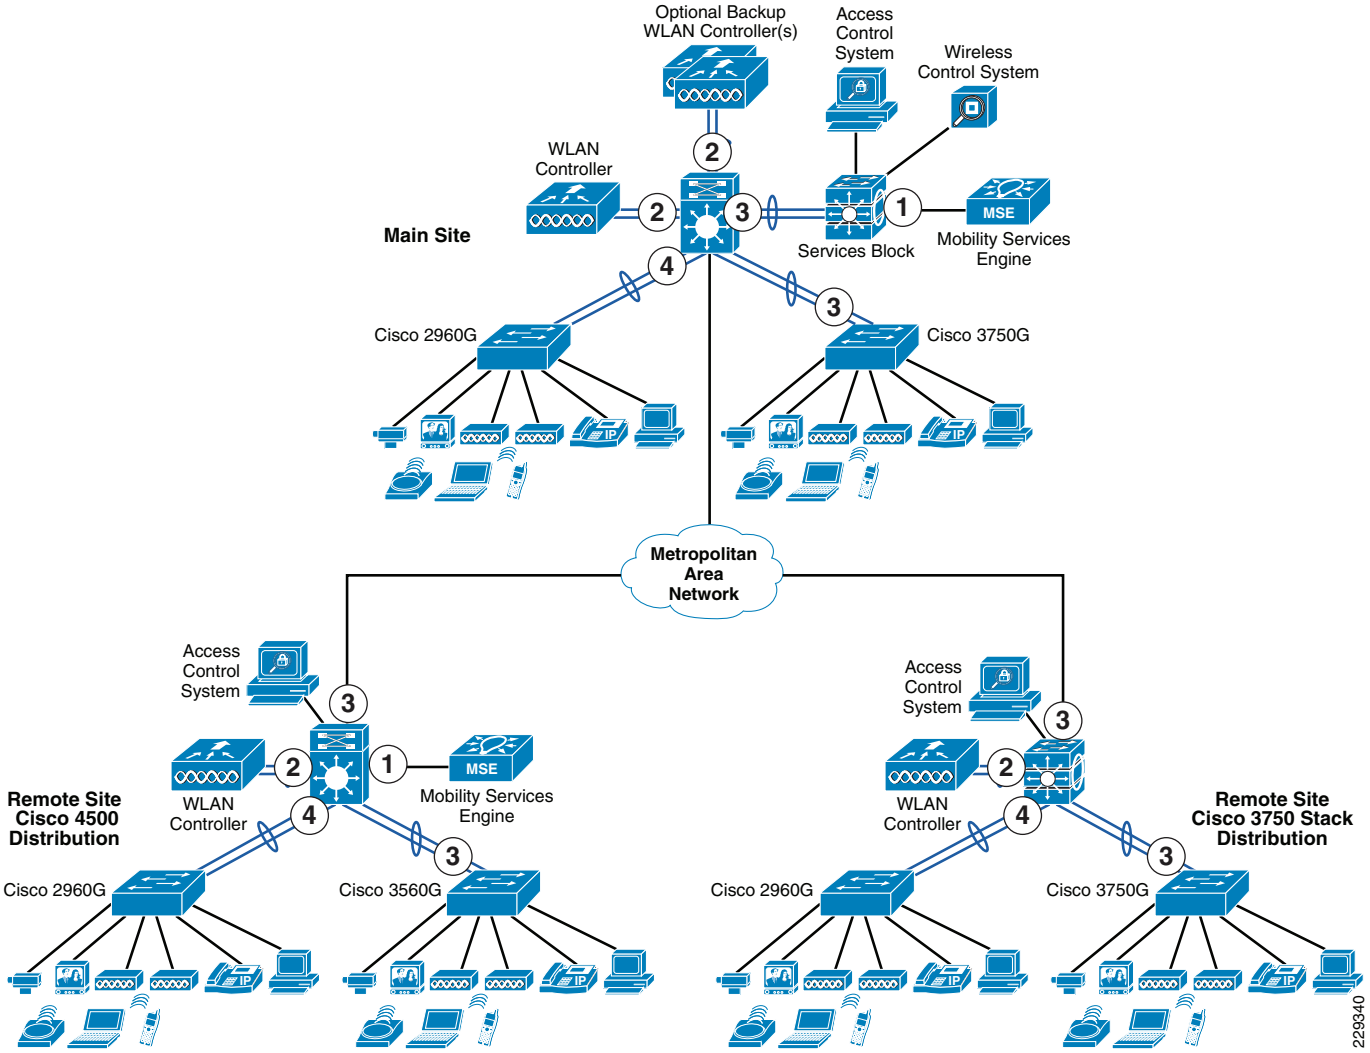

NMSP Traffic Flows Originating At The MSE

NMSP sessions between the MSE, WLAN controllers and context-aware switches are bi-directional in nature. Here, we are referring simply to that portion of any flow whose source address is that belonging to the Mobility Services Engine. Since the Mobility Services Engine does not mark the DSCP for the NMSP traffic it transmits into the network, all such traffic originating at the MSE will contain the default DSCP marking of 0x00. Left unchanged, when congestion is encountered, all NMSP traffic marked in this fashion will be treated with the lowest priority. This increases the probability that NMSP session data will be dropped in the network during periods of congestion. Left unchecked, this can result in NMSP session stability issues, and poor context-aware performance. Clearly, this is not desirable. We can avoid this by classifying and remarking NMSP data appropriately, as described in this section.

Since the MSE currently does not support classification and marking the DSCP value assigned to its traffic, we must make use of the classification and marking capabilities available to us in the Cisco Catalyst switch to which the MSE is attached. Thus, where **1** appears in Figure 6-35, we will:

- Define the criteria to select our NMSP traffic
- Define a class-map that will filter NMSP traffic using this criteria
- Define a service-policy to assign our desired DSCP value to the filtered traffic using the class-map
- Apply the service-policy to the port to which the MSE is attached.

Since we know that NMSP traffic will involve TCP port 16113, we can make use of this to identify NMSP traffic and proceed to mark it appropriately.

The following example defines a policy map that will mark NMSP traffic inbound to the network from the MSE as DSCP 0x12 (also referred to as DSCP 18, Assured Forwarding 21), police down to 10 Mbps per 8k burst, and mark down any NMSP traffic exceeding this accordingly using the QoS map.

```
mls qos
!
class-map match-all CAS
 match access-group name NMSP
!
policy-map MSE-Policy
 class CAS
 set dscp af21
 police 10000000 8000 exceed-action policed-dscp-transmit
!
ip access-list extended NMSP
 remark Identify NMSP traffic
 permit tcp any any eq 16113
 permit tcp any eq 16113 any
```

On the switch interface where the MSE is attached, it is imperative that service-policy statement appears to assign the policy map to the interface. For example:

```
interface GigabitEthernet1/0/2
description Mobility Services Engine MSE1
    .
      .
service-policy input MSE-Policy
```

NMSP Traffic Generated By WLAN Controllers

As was the case with the MSE, NMSP traffic entering the network originating at the WLC is also marked with the default DSCP value of 0x00. Left unchanged, when congestion is encountered all NMSP traffic marked in this fashion will be treated with the lowest priority. Once again, this is not desirable and we shall address it in this section.

Like the MSE, the WLAN controller does not provide us with the ability to mark the NMSP traffic as we see fit, thus we must instead classify and mark this traffic inside the network. In accordance with general best practices, this operation is performed as close in the network as possible to the WLAN controller. Therefore, in Figure 6-35, the points at which this should occur are shown by **2**. Note that traffic coming from both normally active as well as any backup WLAN controllers located in the data center services switch block must be classified and marked.

The method used to accomplish this classification and marking for WLAN controllers is very similar to that presented in the previous section for the MSE. However, since the WLAN controller is attached to the Cisco Catalyst switch using an Etherchannel port-channel group, there will be some relevant differences relating to whether the port-channel attachment is to a Catalyst 3750 switch stack or Catalyst 4500 distribution switch.

For example, when using the Catalyst 3750 switch stack in distribution, the following configuration would apply:

```
mls qos
!
class-map match-all CAS
 match access-group name NMSP
!
policy-map CAS-Policy
 class CAS
 set dscp af21
 police 10000000 8000 exceed-action policed-dscp-transmit
!
ip access-list extended NMSP
 remark Identify NMSP traffic
 permit tcp any any eq 16113
 permit tcp any eq 16113 any
```

The 3750 switch stack would also have a port-channel definition that would refer to the two physical Ethernet interfaces comprising the port-channel group.

```
interface Port-channel4
 description EC trunk to site 1266, WLC, interfaces gig1/0/28, 2/0/28
```

It is important to note that the 3750 switch stack does not support the use of a policy-map statement on a port-channel definition. In order to assure that NMSP traffic originating at the WLAN Controller in-bound to the network is properly classified and marked, ensure that a service-policy statement appears on each of the two physical interfaces that comprise the WLAN controller port-channel group in the 3750 switch stack. For example:

```
interface GigabitEthernet1/0/28
 description trunk to site 1266 WLC, port-channel4
    .
    .
 mls qos trust dscp
 channel-group 4 mode on
 service-policy input CAS-Policy
!
interface GigabitEthernet2/0/28
 description trunk to site 1266 WLC, port-channel4
```

```
    .
    .
 mls qos trust dscp
 channel-group 4 mode on
 service-policy input CAS-Policy
```

When using the Catalyst 4500 as the distribution switch, the scenario is a bit different. The Catalyst 4500 requires the service-policy to be assigned to the port-channel group used for the WLAN controller, and not the physical interfaces. Thus, our recommended configuration when using a Catalyst 4500 in distribution would be as follows:

```
!
class-map match-all CAS
 match access-group name NMSP
!
policy-map CAS-Policy
 class CAS
  set dscp af21
    police cir 10000000
      conform-action transmit
      exceed-action set-dscp-transmit default
!
ip access-list extended NMSP
 remark Identify NMSP traffic
 permit tcp any any eq 16113
 permit tcp any eq 16113 any
!
interface Port-channel6
 description trunk to main site, interfaces gig2/11, gig3/11
    .
    .
 service-policy input CAS-Policy
!
!
interface GigabitEthernet2/11
 description trunk to main site WLC, port-channel6
    .
    .
 channel-group 6 mode on
!
interface GigabitEthernet3/11
 description trunk to main site WLC, port-channel6
    .
    .
 channel-group 6 mode on
!
```

NMSP Traffic Generated By Context-Aware Switches

As you will recall, specific models of Catalyst switches described earlier in this document (such as the 2960G, 3560, 3750, 4500, 4900 and others) can provide context-aware information to the MSE relating to IP-based attached devices. When this capability is enabled in a context-aware Cisco Catalyst switch, the switch itself participates in an NMSP session with the MSE. This NMSP session is separate and independent of any NMSP sessions that may pass through the switch to or from attached devices (such as the WLAN controllers described earlier, or other downstream context-aware Cisco Catalyst switches).

In order to ensure that the NMSP traffic originating at these switches is treated appropriately in the network during times of congestion, it is important that we properly classify the NMSP data originating at the context-aware switch and destined in-bound to the network for the MSE. Depending on the type of switch used, the exact method we shall use to apply this classification and marking will vary.

Layer-Three Context-Aware Switches

In the case of context-aware Layer-3 switches such as the 3750, 3560 and 4500, we can make use of local policy routing to assign an IP precedence 2 (DSCP 0x10) to NMSP traffic originating from the switch itself. Local policy routing in this fashion is performed internal to the context-aware L3 switch originating the NMSP traffic, and the traffic is marked prior to being introduced into the network. At the locations in Figure 6-35 marked with **3**, we can perform the required classification using a local policy route-map as follows:

```
ip local policy route-map switch-NMSP
!
route-map switch-NMSP permit 10
 match ip address NMSP
 set ip precedence 2
 set ip tos max-throughput
!
ip access-list extended NMSP
 permit tcp any any eq 16113
 permit tcp any eq 16113 any
!
```

Layer Two Context-Aware Switches

However, the use of the Cisco 2960G context-aware Layer 2 (L2) switch in the access layer presents a challenge. Local policy routing is not supported by the L2-only 2960G. Therefore, as a work around, we must classify and mark the NMSP traffic originating from the 2960G at the next switch upstream to it in the network. In the small enterprise designs discussed in this document, this upstream switch would be the 3750 switch stack or 4500 distribution switch. In Figure 6-35, the points at where this must be performed are indicated by **4**.

Note that in the Small Enterprise Design Profile, the 2960G access layer switches are attached to the distribution switches using an Etherchannel port-group, in a fashion very similar to that of the WLAN controllers. Therefore, similar techniques can be used to classify and mark the NMSP traffic originating at the 2960G and coming across the Etherchannel link.

Thus, for a 3750 switch stack used in distribution, a **class-map** and **policy-map** can be applied as described previously, and the **service-policy input CAS-Policy** statement applied to the physical interfaces in the 3750 switch stack that comprise the port-channel group to the 2960G switch. Just as for an Etherchannel-attached WLAN controller, a service-policy cannot be applied to the port-channel group definition used for the 2960G.

When using a Catalyst 4500 as the distribution switch with a context-aware 2960G at the access layer, apply the **service-policy input CAS-Policy** statement to the port-channel group definition for the 2960G. This would be done in a similar fashion to that described earlier for a Etherchannel-attached WLAN controller.

Hardware/Software Releases

Table 6-2 Hardware/Software Releases

Component	Version	Comments
Wireless Control System	6.0.132.0	For licensing and part number information, see http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd804b4646.html

Table 6-2 Hardware/Software Releases

Mobility Services Engine 3350	6.0.85.0	For client and tag licensing information, see http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html
Mobility Services Engine 3310	6.0.85.0	For client and tag licensing information, see http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html
WLAN Controller 4404	6.0.182.0	Standalone Cisco Wireless LAN Controller
WLAN Controller 4402	6.0.182.0	Standalone Cisco Wireless LAN Controller
Catalyst 4500	12.2.53-SG	Large site distribution switch; must use crypto (K9) image if CAS for wired devices is desired
Catalyst 3750E Switch Stack	12.2(52)SE	Small site distribution switch; must use crypto (K9) image if CAS for wired devices is desired
Catalyst 2960G	12.2(52)SE	Access switch; must use crypto (K9) image if CAS for wired devices is desired
Catalyst 3750E	12.2(52)SE	Access switch; must use crypto (K9) image if CAS for wired devices is desired
LAP1252 Access Point	6.0.182.0	Cisco Aironet 1250 Series Wireless Access Point
LAP1142 Access Point	6.0.182.0	Cisco Aironet 1140 Series Wireless Access Point
WCS Navigator	1.5.132.0	Optional component; for licensing information, see http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps7305/product_data_sheet0900aecd8065bd19.html
AeroScout T2 RFID Asset Tag	4.33	Available from http://www.aeroscout.com/; RFID tags are required only if RFID tracking is desired
AeroScout T3 RFID Asset Tag	6.05	Available from http://www.aeroscout.com/; RFID tags are required only if RFID tracking is desired
AeroScout EX-3200 Exciter	33007/60007	Chokepoint trigger, optional RFID enhancement; EX-2000 model recommended if outdoor placement is necessary

Context-Aware Services—General Best Practice References

The following are recommended references with regard to general best practice deployment recommendations for Cisco Unified Networks making use of Context-Aware Services release 6.0:

- A cornerstone of a successful design is knowledge of established best practices. Thus, it is highly recommended that you become familiar with the material presented in the following documents:
  - *Context-Aware Solution Deployment Guide* http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml
  - *VoWLAN Design Guide* 4.1 http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html
  - *Context-Aware Services Configuration Guide*, Release 6.0 http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/CAS_60.html

  - *Wireless LAN Controller Configuration Guide*, Release 6.0 http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.html
  - *Wireless Control System Configuration Guide*, Release 6.0 http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/WCS60cg.html

- If you intend to make use of RFID tags in your Context-Aware solution, it is also recommended that you become familiar with the following document which explains the operation of the Cisco Context-Aware Engine for Tags:
  - *AeroScout Context-Aware Engine for Tags for Cisco MSE User Guide*, version 3.2 http://support.aeroscout.com
- If you intend to track the location and status of wired devices attached to Cisco Catalyst switches, it is recommended that you familiarize yourself with the appropriate configuration guide for this feature in the switches you will be using. For example:
  - *Catalyst 2960 Switch Software Configuration Guide*, 12.2(50)SE http://www.cisco.com/en/US/docs/switches/lan/catalyst2960/software/release/12.2_50_se/configuration/guide/scg.html
  - *Catalyst 3750 Switch Software Configuration Guide*, 12.2(50)SE http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2_50_se/configuration/guide/scg.html
  - *Catalyst 4500 Series Switch Software Configuration Guide*, 12.2(53)SG http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/53SG/configuration/config.html

During any deployment of Context-Aware Services and Cisco Unified Wireless Networks, detailed site surveys should be performed by a Cisco Wireless LAN Specialized Partner with expertise in voice, high speed data, and Context-Aware wireless network deployment. Cisco Systems also offers a complete package of bundled design and deployment services via the Cisco Advanced Services team. Cisco and our Wireless LAN Specialized Partners offer Context-Aware Design and Implementation Services to help you successfully deploy enterprise-class wireless connectivity. These services include the installation and configuration of crucial components such as the Mobility Services Engine (MSE), helping you to take full advantage of the strong security, management, and investment protection features that are built into Cisco Context-Aware components. In addition to planning, design, and implementation, we also offer services based on proven methodologies for operating and optimizing the performance of a Context-Aware Mobility solution, along with its associated technologies and strategies.

Note The importance of a properly performed wireless site survey of your facility cannot be over-emphasized. For more information on Cisco bundled planning, design, and deployment services, refer to http://www.cisco.com/en/US/services/ps2961/ps6899/ps8306/services_overview_context_aware.pdf. To locate a Cisco Wireless LAN Specialized Partner, refer to http://tools.cisco.com/WWChannels/LOCATR/openAdvanceSearch.do.

CHAPTER 7

Small Enterprise Design Profile(SEDP)—Digital Media and Video Surveillance Design

Digital Media System Overview

The Small Enterprise Design network architecture in combination with the Cisco Digital Media System (DMS) creates an environment that streamlines and automates information flow and process throughout enterprise network.

The evolution of digital communication in the 21st century enterprise network environment is transforming processes and empowering educators to develop and deploy “eye-catching”, compelling, and integrated video content.

- Digital media communication in small enterprise systems is an extremely effective tool to deliver dynamic and cost effective subject-matter to staff and customers.
- Using digital media systems enables enterprise organization to keep their customers connected, increases awareness, and integrates with safety and security system notifications.
- Integrating digital media technologies assists corporate information, business events and internal professional development.
- Comprehensive digital video systems transform video delivery processes in individual business function; network administrator can create customized, on-demand educational video content or even relay live video feed during national events or local emergencies.

Video Surveillance Overview

Video surveillance has been a key component of safety and security groups for many organizations. As an application, video surveillance has demonstrated its value and benefits countless times by providing real-time monitoring of a facility’s environment, people, and assets as well as by recording events for subsequent investigation, proof-of-compliance, and audit.

For small enterprise network systems that need to visually monitor or record events video, surveillance has become more important as the number of security risks increases. In addition to video analytics, the value of video surveillance has grown significantly with the introduction of motion, heat, and environmental sensors.

In a typical enterprise network environment, several systems are deployed for disparate applications, such as physical access control, fire and smoke detection, and video surveillance. These applications typically do not communicate with each other and require different management and support personnel.

As a result, owners and operators suffer from a lack of operational consistency, interoperability, and capabilities that translate into higher capital and operational costs and limit the return on their investment.

Cisco's solution offers software and hardware to support video transmission, monitoring, recording, and management. Cisco video surveillance solutions work in unison with the advanced features and functions of the IP network infrastructure—switches, routers, and other network security devices—to enable secure, policy-based access to live and recorded video.

Cisco Digital Media System Architecture

Cisco has developed a comprehensive, scalable, and network-centric DMS architecture that is built on three major digital application components. Each application is specifically designed to address key challenges, regardless of how or where the digital content is designed and developed. The multifunction management and adaptive media system is common to these three applications:

- *Desktop video*—Interactive education training application that allow employees to watch instructional and corporate event videos on-demand or via live streaming. With the right privilege, employees can navigate a Cisco Video Portal database to securely access relevant training video content.
- *Enterprise TV*—While the targeted users for desktop video applications are individuals or group of employees, enterprise TV expands the same video capability to a larger audience. Live or on-demand pre-recorded training video can be broadcast in a enterprise network independent of geographical location. In addition to internally developed video material, main and remote site network administration can also integrate live TV programming feed from the local content provider.
- *Digital signage*—Enables innovative ways to publish content and information that improves the user experience, allows dynamic updates, and increases campus safety and security. Some of the common digital signage use cases in enterprise networks include announcing main and remote site news, major events, corporate initiatives, board meetings, etc.

Figure 7-1 shows a Cisco DMS solution suite that is a set of product and technologies developed to create an end-to-end digital media network. The products are divided into three major functions, create, manage, and access.

Figure 7-1 Cisco DMS Solution Suite

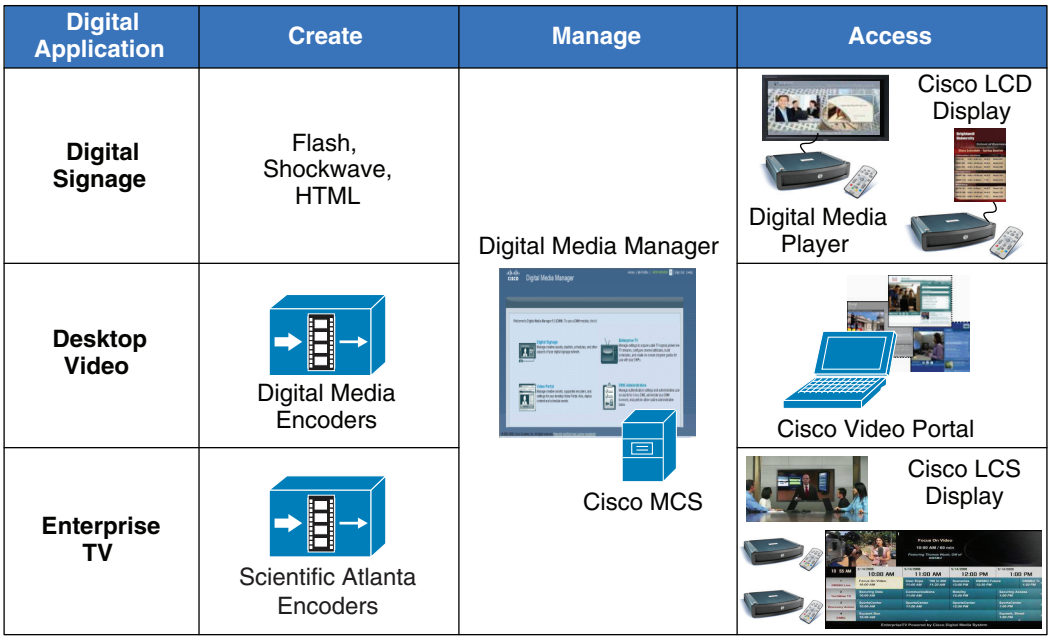

The digital media components in the Cisco DMS solution suite assist small enterprise network in deploying digital media solutions at their own pace; e.g., initially deploy digital signage with simple development applications and deploy interactive video solutions in subsequent phases.

The Cisco Digital Media Manager (DMM) is a Web-based application that simplifies deploying all three digital media applications in small enterprise network.

DMS Solution for Small Enterprise

DMS solutions in small enterprise have proven valuable in overcoming key challenges in the development of next-generation education delivery processes. DMS provides the flexibility to develop training content that can be accessed by employees anytime, and anywhere. Cisco DMS relies on a resilient, scalable, and reliable network infrastructure for seamless end-to-end content delivery.

Figure 7-2 shows an end-to-end digital media reference model with all three applications enabling a unified digital network service for the main site of small enterprise network.

The following are some of the key benefits of this DMS design model:

- Centralized management at the main site providing consistent publishing policies, security, scalability and reduced operational and maintenance cost.
- Distributed storage and media access points enabling main and remote site to use centrally developed content with reduced bandwidth capacity and increased availability during network instability
- In large-scale video networks, the Cisco Application and Content Network Services (ACNS) or WAN optimization appliances can be deployed to increase media performance and reduce expensive WAN bandwidth requirements.

Figure 7-2 End-to-End Reference DMS Solution in Small Enterprise Network Architecture

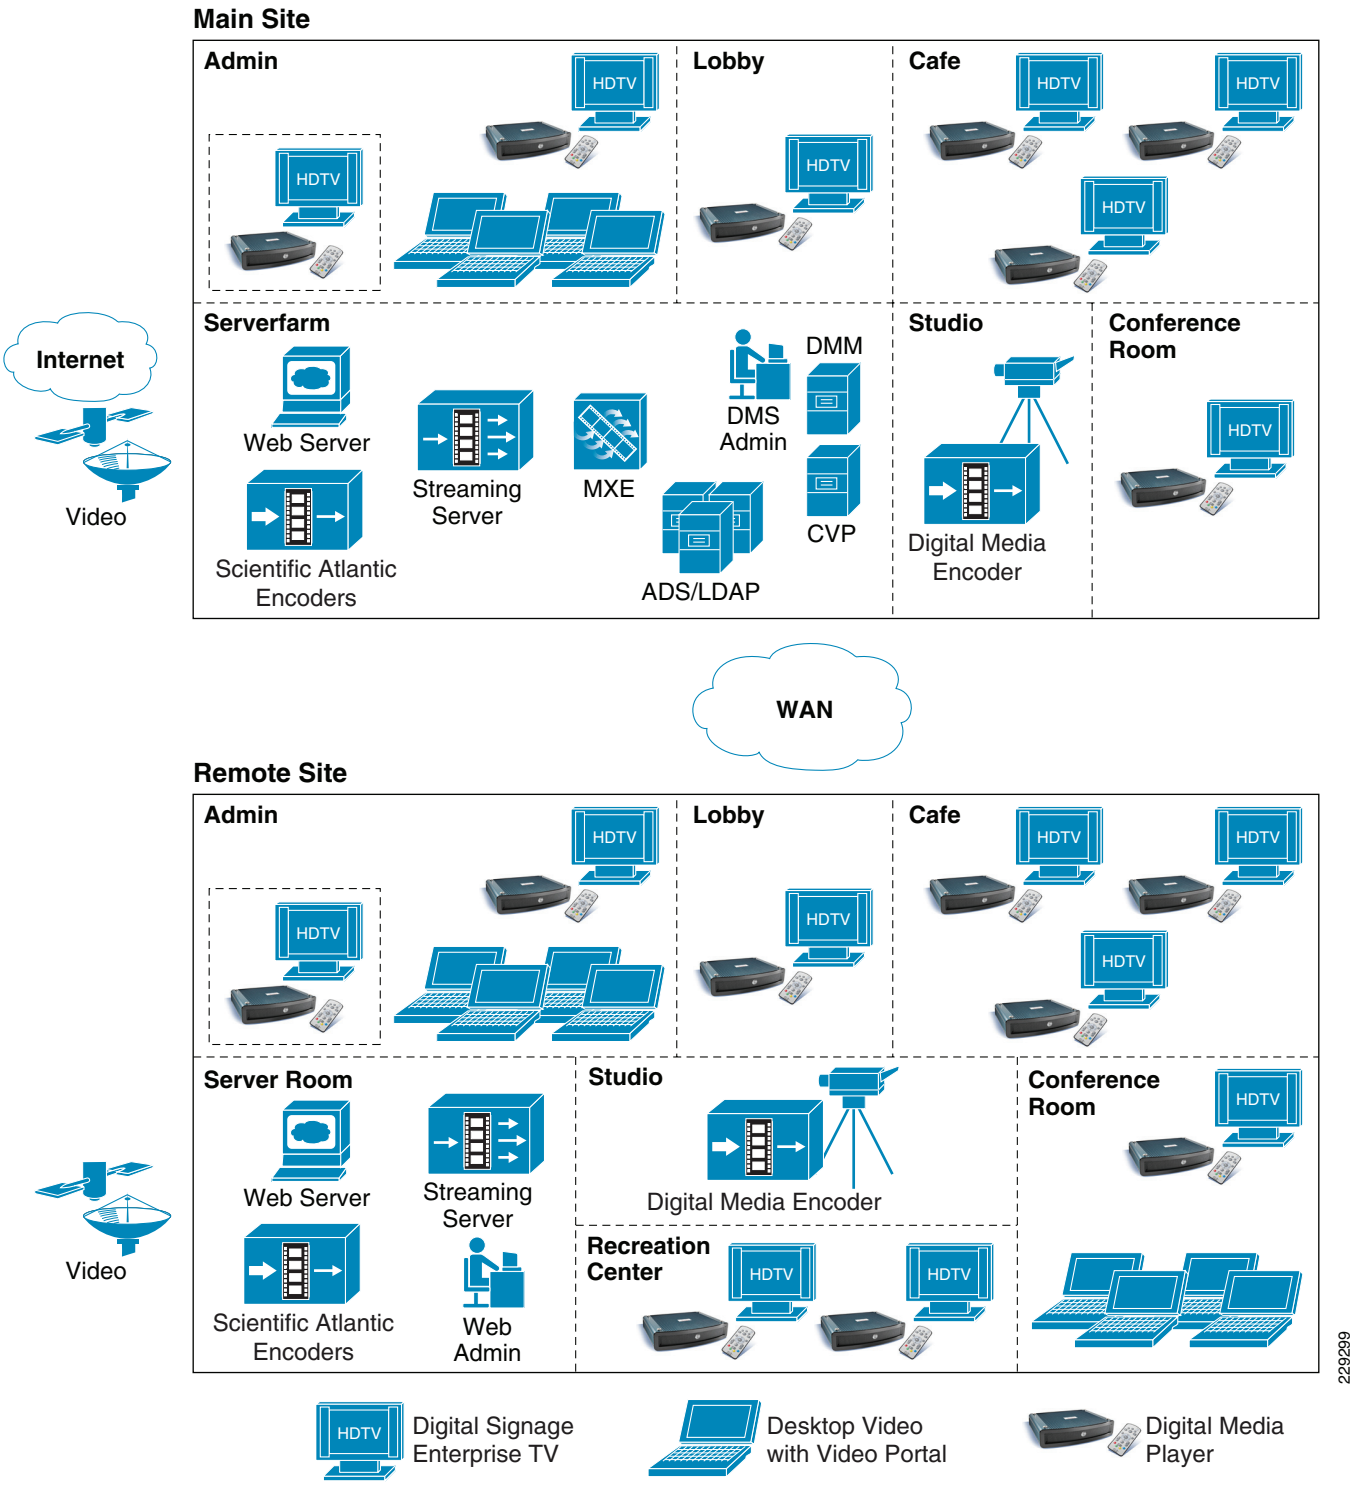

The following section provides an overview of various deployment scenarios, device components, communication, and network requirements for digital media applications in small enterprise network. For a detailed digital media design and implementation guide, refer to the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/DMS_DG/DMS_dg.html

Desktop Video Application Overview

Employees and administrative staff can watch live or pre-recorded video events from their personal computers at any location and at any time. The Cisco DMS Digital Video application empowers employees to extend the audience reach ability to remote site locations by broadcasting live or recording training sessions available as video on-demand (VoD).

Cisco Desktop Video applications offer several benefits:

- Customizable interface with program guide and search window.
- Employees can create a personalized video playlist.
- Questions and comments can be made during live video broadcast events.
- Restrict video content access based on Active Directory or LDAP authentication and privilege.
- Wide format support—Adobe Flash, Windows Media, H.264, QuickTime, etc.
- Player Controls—Synchronized slides, advanced video and controls, etc.

Desktop Video Components

As one of the integrated components of the Cisco DMS solution suite, the digital video application uses common video development, management, and publishing components. In combination, external authentication servers can provide secure, on-demand video content and live video broadcast services to the desktop or on Cisco LCD TVs deployed in different physical locations.

The Cisco DMS solution suite consists of:

- *Digital Media Manager (DMM)*—Centralized management appliance in main site governs the content and communicates with local or remotely deployed critical desktop video components, e.g., CVP, HTTP server etc.
- *Digital Media Encoder (DME)*—Single or multi-channel media encoder receives live analog/digital feed from cameras or television service providers and transports over the IP network to the streaming server.
- *Streaming Server*—Provides stream splitting capabilities, allowing many clients to view a single live stream from DME or pre-recorded source (replay)
- *Cisco Video Portal (CVP)*—Web-based video navigation engine provides access to users after successful authentication with Active Directory or LDAP server.
- *Web Server/Content Repository*—Stores all VoDs referenced by Video Portal server. User triggers VoD request to access video content through CVP and the request gets redirected to pull video file from the content repository to the requested user.

Publishing Live and Video On-Demand Content

A critical feature of the Cisco Digital Media System for Cisco Desktop Video is its ability to simplify the publishing of live and on-demand digital media files to the Cisco Video Portal (CVP). On-demand video content can be uploaded from the developer's computer directly to the DMM server for staging and previewing prior to deployment. This staging capability includes the addition of an approval process within the content workflow to help ensure that school branding, publishing policies, and messaging are properly incorporated in the content. Post approval process, the content can be moved or deployed to the Cisco Video Portal using secure file transmission.

The Cisco DMM works in conjunction with Cisco Digital Media Encoders (DME) to create and deploy live content to the Cisco Video Portal. The Cisco DMM first manages the Cisco DME to set up their encoding profiles, defining the bit rate, format, and media type. The Cisco DMM also defines the port that the Cisco Digital Media Encoders will stream from, so that the streaming servers can pull the stream to their live publishing points. These publishing points are then deployed to the Cisco Video Portal through the Cisco Digital Media Manager deployment process. The same workflow defined for the on-demand digital media content is applied to live events, providing a consistent, easy-to-use process for all types of deployments. Figure 7-3 shows a schematic of Cisco DMM video management.

Figure 7-3 VoD and Live Video Broadcast Using Digital Video Application

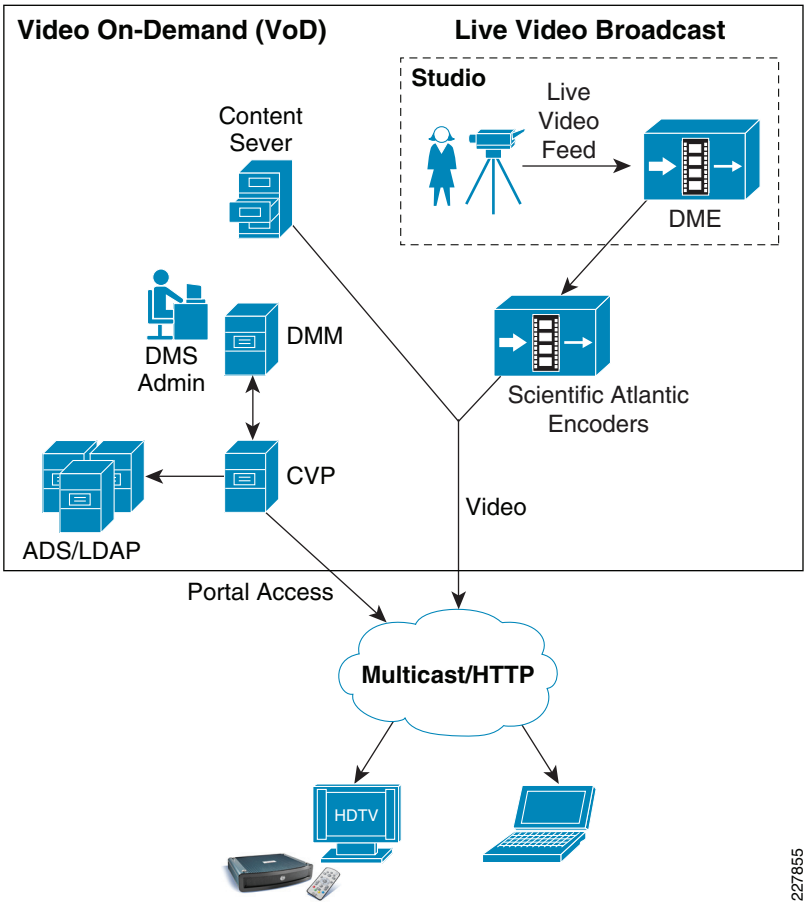

Enterprise TV Application Overview

The Enterprise TV (ETV) application brings standard or high-definition television network channels into IP-based networks. Deploying Scientific Atlanta encoders in a video head-end role performs the interworking function that transforms video source from television service provider to an IP based video delivery within the campus network. When the ETV module is enabled in Cisco DMM appliance, the network administrator can program the channel guide information to be broadcast in different physical locations, e.g., channel number, name, port number, etc. To improve the user experience in navigating video channels, ETV Electronic Program Guide (EPG) can be programmed to provide information on channel lineup, and current and future programming information, which is similar to television service

provided programming guide. To watch the live video channels, users can use Cisco DMP and remote control to navigate and access the channel. Video delivery over IP networks can be unicast or multicast, depending on how IP/multicast is designed in campus network.

With larger displays in key physical locations, the ETV application becomes the primary communications interface in the main site targeting large audiences. For example, live or VoD broadcast for professional training, demonstrations, meetings, etc., targeting all the employees in a personal office, conference rooms, or live news, etc.

Enterprise TV Components

To broadcast developed VoD or live video through ETV, the core Cisco DMS components used in the desktop video application can be leveraged to integrate ETV in the network. The primary difference between ETV and desktop video in this case would be Digital Media Player (DMP) instead of a PC as an access end point for a large screen display targeting a larger audience. To broadcast television network channels in the campus network, the Scientific Atlanta encoder must be integrated along with other ETV components. It is recommended to deploy distributed television service in local campus and not forward non-critical video traffic over the WAN infrastructure.

- *Scientific Atlanta Encoder*—An encoder system that provides interworking function between analog or digital television service provider and IP network. Encodes live video input and transforms into MPEG-2 or MPEG-4 multicast stream.
- *Cisco Digital Media Player (DMP)*—Key media access end point that connects to Cisco LCD TV for large size displays. DMP provides capabilities to decode multi-format graphics and stream video content received over unicast or multicast IP network.

Broadcasting Live TV or Video On-Demand Content

The communication flow between the DMP and the DMM and Video Portal function is similar to desktop video applications. Proper planning, technologies, and equipment must be deployed in campus network for successful live television video delivery. When designing the playlist in Cisco Enterprise TV module, the network administrator must understand that it can support up to 99 live and on-demand video channels broadcast in the campus network. The DMM administrator in the main site can use the DMM-ETV software module to create customized TV navigation interfaces, such as adding corporate logo and skins, programming video channel assignments, and configuring specific video channel assignment to DMP deployed in specific campus location. Figure 7-4 shows a schematic of the Enterprise TV video architecture.

Main and remote site architects must understand the codec type required for publishing video in the campus network. Deployed digital encoders must follow the MPEG2 standard specification to stream the video. It is recommended to deploy Scientific Atlanta 9032SD or 9050HD encoder to stream live video stream to DMP for Enterprise TV application.

Figure 7-4 Live Video Broadcast and VoD Using Enterprise TV Application

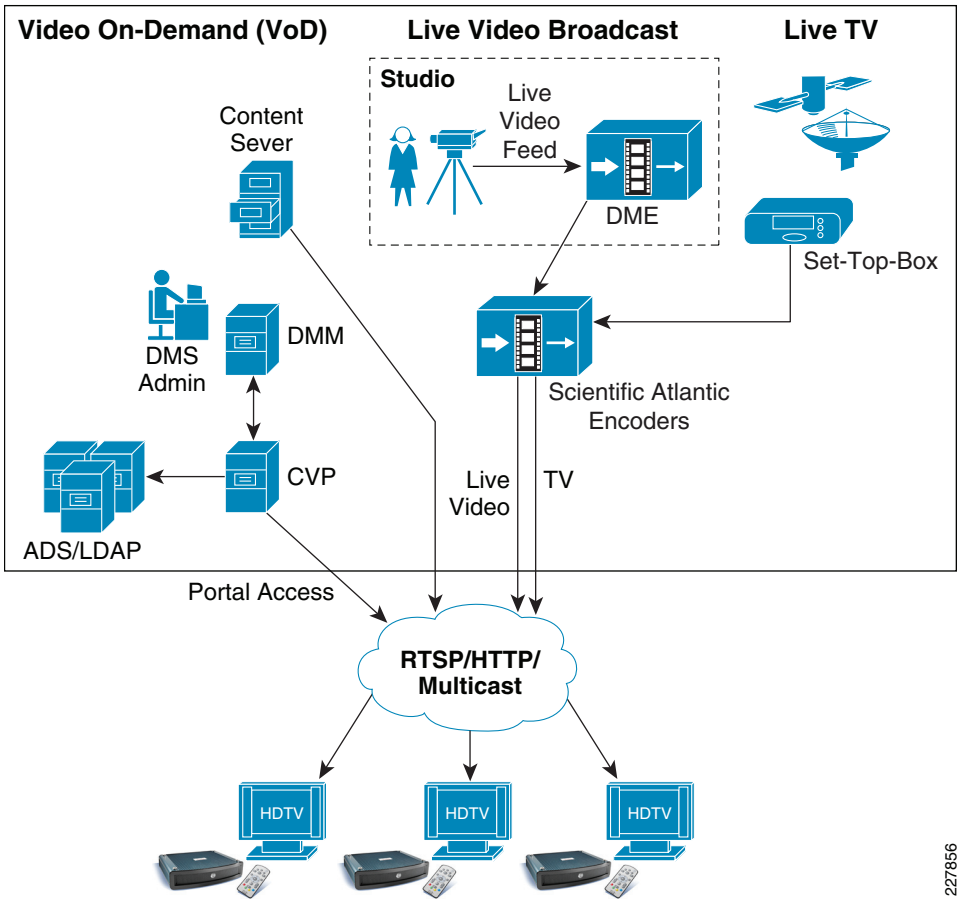

Digital Signage Application Overview

Cisco's Digital Signage solution is a comprehensive solution for the publishing of dynamic and on-demand signage using digital media displays deployed locally or regionally in small enterprise network over an IP network. The key benefits of digital signage over traditional static signs are that the digital content can be exchanged and updated more dynamically, using digital media tools to make the content more relevant and interactive. Publishing corporate messages, local announcements, or emergency alerts through Cisco digital signage becomes more effective and with better investment return compared to traditional models.

The Cisco digital signage application is a Web-based media management and publishing application that creates a playlist with a set of content that is required to be published to a single or group of DMPs in a network. The digital signage application use standard HTTP protocols to communicate with centrally deployed DMM in the main site serverfarm and single or distributed Web servers to pull and publish the real-time information to display on large Cisco LCD TVs. With flexible user administration, the DMM administrator is empowered to create groups of users with different privileges who can develop, publish, and manage the signage content; e.g., user group like Web admin, IT admin, and security admin can create content assets, control display properties, etc.

Digital Signage Components

The requirements of the integrated digital media components depend on which content needs to be published through a signage module. The Cisco digital signage application provides the flexibility for the DMM administrator or Web administrator to develop multi-functional, integrated content that can play key messages, stream VoD files from a content server, and keep connected with external information. For example, administrator or employees can publish a single Adobe flash file that is composed of static text information, embedded with education short VoD stream and live news information with RSS feed. The basic digital signage components are:

- *Digital Media Manager (DMM)*—Centralized management appliance in main site governs the content and communicate with local or remotely deployed critical desktop video components, e.g., CVP, HTTP server, etc.
- *Cisco Digital Media Player (DMP)*—The Cisco DMP is a highly reliable IP-based hardware endpoint for video decoding and playback of digital media content—including high-definition live broadcasts and VoD, Flash animations, text tickers, and other Web content—across digital displays. DMP is a critical component of the digital signage and ETV applications allow for the networking of digital displays and the broadcasting of live and on-demand media. The current DMP portfolio includes the Cisco DMP 4305G for standard signage and ETV and the Cisco DMP 4400G for high-end signage and ETV.
- *Cisco LCD Professional Series*—For an end-to-end Cisco digital media solution, the Cisco LCD professional series displays is a high definition LCD display that can be centrally managed through DMM.
- *Web Server/Content Repository*—Stores all HTML, VoDs, flash files referenced by HTTP server or Video Portal server. Multiple files can be played on same DMP; based on Web application design, the program triggers the content request and it is pulled by DMP from a source server.

Video Surveillance System Architecture

The Cisco Video Surveillance solution relies on an IP network infrastructure to link all components. The design of a highly available hierarchical network has been proven and tested for many years and allows applications to converge on an intelligent and resilient infrastructure.

Figure 7-5 shows the main components of the Cisco Physical Security solution, including video surveillance, physical access control, incident response and integration with third-party systems.

Figure 7-5 Cisco Physical Security Components

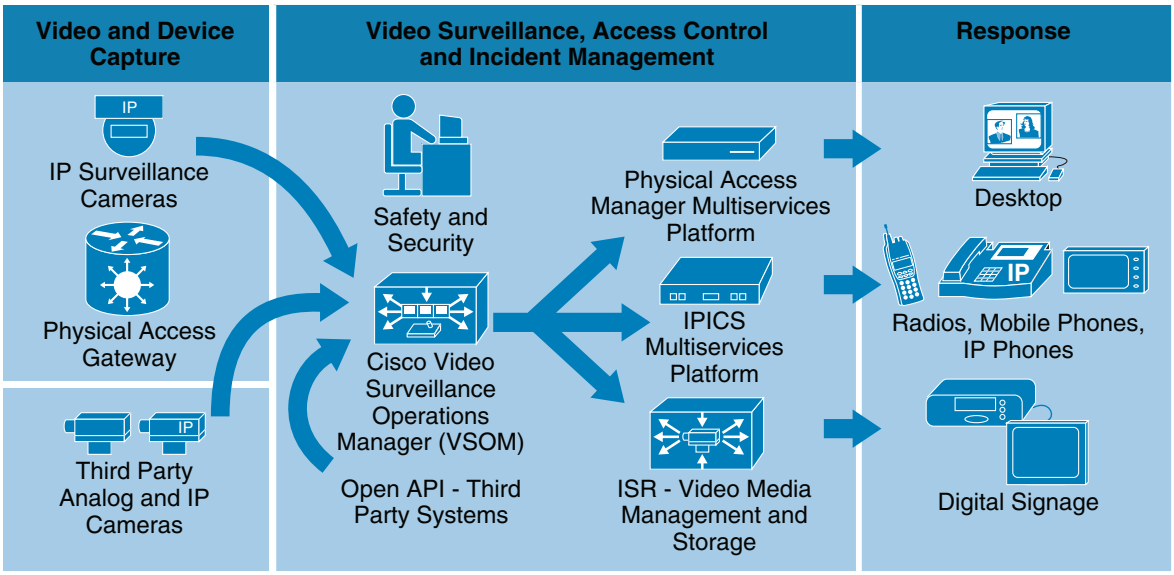

Some of the benefits of Cisco's Video Surveillance solution include the following:

- Access to video at any time from any network location, enabling real-time incident response and investigation.
- Transfer of control and monitoring to any other point in the network in an emergency situation.
- Ability to manage devices and alarms from a centralized location.
- Ability for products from various vendors to interoperate in the same network.
- An open, standards-based infrastructure that enables the deployment and control of new security applications.

The main components of the Cisco Video Surveillance solution include the following:

- *Cisco Video Surveillance Media Server*—The core component of the network-centric Video Surveillance Manager solution. This software manages, stores, and delivers video from a wide range of cameras and encoders over an IP network
- *Cisco Video Surveillance Operations Manager*—The Operations Manager authenticates and manages access to video feeds. It is a centralized administration tool for management of Media Servers, Virtual Matrixes, cameras, encoders, and viewers and for viewing network-based video.
- *Cisco Video Surveillance IP Cameras*—The high-resolution digital cameras are designed for superior performance in a wide variety of environments.
- *Cisco Video Surveillance Virtual Matrix*—The Virtual Matrix monitors video feeds in command center and other 24-hour monitoring environments. It allows operators to control the video being displayed on multiple local and remote monitors.
- *Cisco Video Surveillance Encoding Server*—This all-in-one appliance encodes, distributes, manages, and archives digital video feeds for analog cameras. Each server encodes up to 64 channels and provides up to 12 TB of storage.
- *Cisco Video Surveillance Storage System*—This complementary component allows the Media Server's internal storage to be expanded with direct attached storage (DAS) and storage area networks (SANs). The Storage System allows video to be secured and accessed locally or remotely.

The following subsections describe the components used for this solution.

Cisco Video Surveillance Media Server

The Cisco Video Surveillance Media Server (VSMS) is the core component in the Cisco Video Surveillance Manager solution and performs the following networked video surveillance system functions:

- Collection and routing of video from a wide range of third-party cameras and video encoders over an IP network
- Event-tagging and recording of video for review and archival purposes
- Secure local, remote, and redundant video archive capabilities

In Figure 7-6, the Media Server is responsible for receiving video streams from different IP cameras and encoders and replicating them as necessary to different viewers.

Figure 7-6 Video Surveillance Media Server (VSMS)

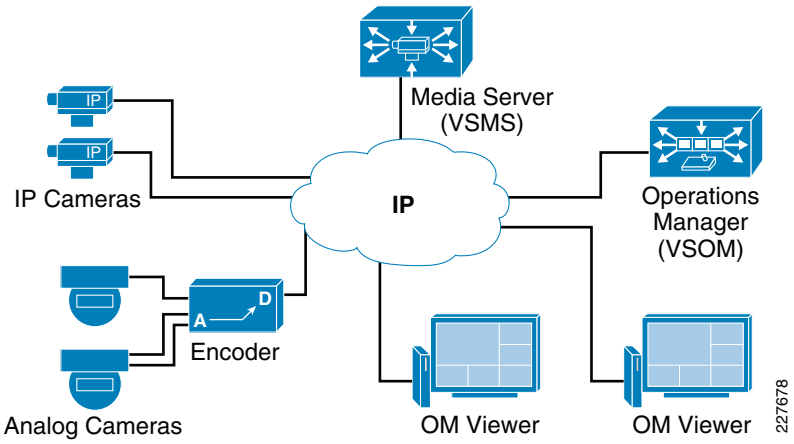

By using the power and advanced capabilities of today's IP networks, the Media Server software allows third-party applications, additional users, cameras, and storage to be added over time. This system flexibility and scalability supports the following:

- Hundreds of simultaneous users viewing live or recorded video
- Standard video compression algorithms such as MJPEG, MPEG-2, MPEG-4, and H.264 simultaneously via a single Media Server
- Conservation of storage using events and loop-based archival options
- Integration with other security applications

Cisco Video Surveillance Operations Manager

Working in conjunction with the Cisco Video Surveillance Media Server, the Cisco Video Surveillance Operations Manager (VSOM) enables organizations to quickly and effectively configure, manage, and view video streams throughout the enterprise. Figure 7-7 shows the Operations Manager main screen, which is accessed through a Web browser.

Figure 7-7 Video Surveillance Operations Manager

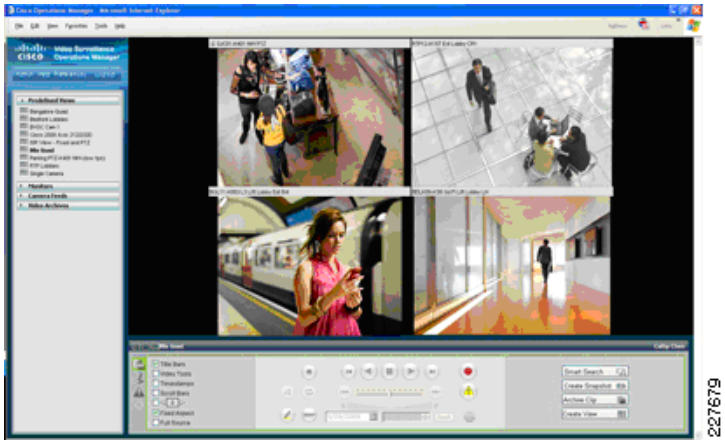

The Operations Manager meets the diverse needs of administrators, systems integrators, and operators by providing the following:

- Multiple Web-based consoles to configure, manage, display, and control video throughout a customer's IP network.
- The ability to manage a large number of Cisco Video Surveillance Media Servers, Cisco Video Surveillance Virtual Matrixes, cameras and users.
- Customizable interface, ideal for branded application delivery.
- Encoder and camera administration.
- Scheduled and event-based video recording.
- Interface to Media Server and Virtual Matrix software for pushing predefined views to multiple monitors.
- User and role management.
- Live and archived video views.
- Friendly user interface for PTZ controls and presets, digital zoom, and instant replay.
- Event setup and event notifications.
- *"Record Now"* feature while viewing live video

Cisco Video Surveillance IP Cameras

Cisco 2500 Series Video Surveillance IP Camera

The Cisco 2500 Series Video Surveillance IP camera is a high resolution standard-definition, feature-rich digital camera designed for secure performance in a wide variety of environments. The camera supports MPEG-4 and MJPEG compressions with up to 30 frames per second.

Contact closure and two-way audio allow integration with microphones, speakers, and access control systems. By providing wired and wireless models, the Cisco 2500 IP camera provides an ideal platform for integration and operation as an independent device or as part of the Cisco Video Surveillance network. Figure 7-8 shows both the wired and wireless models of the 2500 IP Camera.

Figure 7-8 Cisco 2500 Series IP Cameras

The 2500 Series IP camera provides the following features:

- The camera employs powerful digital imaging technology, allowing it to capture high-quality images in a wide variety of indoor and outdoor lighting conditions. It uses a progressive scan image-sensor with global electronic shuttering to ensure natural color rendition, and minimal motion blurring.
- The wireless IP camera model supports 1X2 Multiple Input Multiple Output (MIMO) communication, which provides better data throughput and higher link range than single antenna designs. The wireless IP camera offers strong wireless security using Wi-Fi Protected Access (WPA)/WPA2 and supports various network protocols for 802.1x authentication.
- Power over Ethernet (PoE) 802.3af or DC power through an optional external power supply.
- Support for the Cisco Media API, an open, standards-based interface that allows integration with compatible video surveillance management systems.
- Support for 802.1x authentication on both the wired and wireless models.

Cisco 4000 Series Video Surveillance IP Camera

The Cisco Video Surveillance 4000 Series IP Cameras employ true high-definition (HD) video and H.264 compression, streaming up to 30 frames per second at 1080p (1920 x 1080) resolution. The Cisco 4000 IP Camera series also supports contact closure and two-way audio allow integration with microphones, speakers, and access control systems.

The Cisco 4000 Series includes two models: the CIVS-IPC-4300 and CIVS-IPC-4500. These cameras have identical feature sets, with the exception of the additional digital signal processor capabilities specifically designed to support real-time video analytics at the edge on the CIVS-IPC-4500. On this model, applications and end users have the option to run multiple analytics packages without compromising video streaming performance on the camera.

Figure 7-9 shows a Cisco 4000 IP Camera with an optional DC Auto Iris Lens.

Figure 7-9 Cisco 4000 Series IP Camera

The 4000 Series IP camera provides the following features:

- True high-definition video—The camera streams crisp and clear 1080p (1920 x 1080) video at 30 frames per second while maintaining surprisingly low network bandwidth.
- Progressive scan video—The camera captures each frame at its entire resolution using progressive scan rather than interlaced video capture, which captures each field of video.
- Embedded security and networking—The camera provides hardware-based Advanced Encryption Standard (AES).
- IP Multicast for enhanced bandwidth management.
- Event notification—The camera can examine designated areas for activity and notify users or other applications when it detects activity that exceeds a predefined sensitivity and threshold.
- True day/night functionality that includes an infrared (IR) filter that automatically switches to night mode in low light scenes.
- The camera supports Power-over-Ethernet (PoE) 802.3af, 12 VDC or 24 VAC power through an optional external power supply.
- The camera can be installed with a fixed mount or with an optional external pan/tilt mount and motorized zoom lens.

For more information, refer to the *IP Video Surveillance Deployment Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/IPVS/IPVS_DG/IPVS_DG.pdf

Deploying Digital Signage in Small Enterprise Network

To ensure a successful digital media deployment, network, display, and management planning must be done prior to deploying digital signage in the small enterprise network. A well-planned digital signage network design provides flexibility to incrementally deploy desktop video and ETV digital media applications without making major infrastructure changes. As described earlier, digital signage uses standard HTTP protocols to pull and publish the signage content on displays. Network bandwidth consumption for digital signs varies widely as it depends on playlist and content types. This document provides the best practices to design and configure the digital signage with content developed with rich text, flash, and animation and located on distributed Web servers to increase network efficiency.

Centralized Management Model

Cisco DMM is highly scalable appliance server that can be deployed centrally in main site location to manage up to 1000 DMPs deployed in local and regional small enterprise campus network. Cisco DMP deployed in main or remote sites can communicate with centralized DMM over LAN and WAN network

using standard HTTP as the control protocol to receive Web or content server re-direction information to display content. Deploying DMM in a centralized location allows the DMM administrators in main site to manage all registered DMP in various ways:

- Add and archive digital content and assign metadata and keywords.
- Create and manage play lists, ticker alerts, messages, closed captions, and promotional interstitials.
- Preview digital signage content and manage approval workflow.
- Ability to pre-configure the playlist and schedule for instant and future deployments.
- Take WAN optimization solution advantage and provide tight integration with Cisco ACNS and Cisco Content Engines.
- Manage user administrator accounts and permissions.

Centralizing DMP management and publishing signage content centrally at a main site allows the DMM administrator to advertise consistent information and messaging throughout the network. To minimize WAN network utilization, the remote site network administrator must leverage internal storage or their local Web server to store and advertise the local, regional, and department news and information. However the main site and Internet news must be communicated over the WAN.

The network architect must consider integrating enterprise-class Cisco Application and Content Networking System (ACNS) that uses caching technology and offers higher scalability and reliable video delivery solution in the network to improve end user experience and application response time. When integrated with the Cisco Wide-Area Application Services (WAAS) solution, it helps in optimizing WAN bandwidth utilization significantly with local redirection and high data compression over the WAN.

Figure 7-10 is a validated design to integrate digital signage with centralized management in the main site with distributed DMP in the network.

Figure 7-10 Centralized DMM with Distributed DMP in Small Enterprise Network

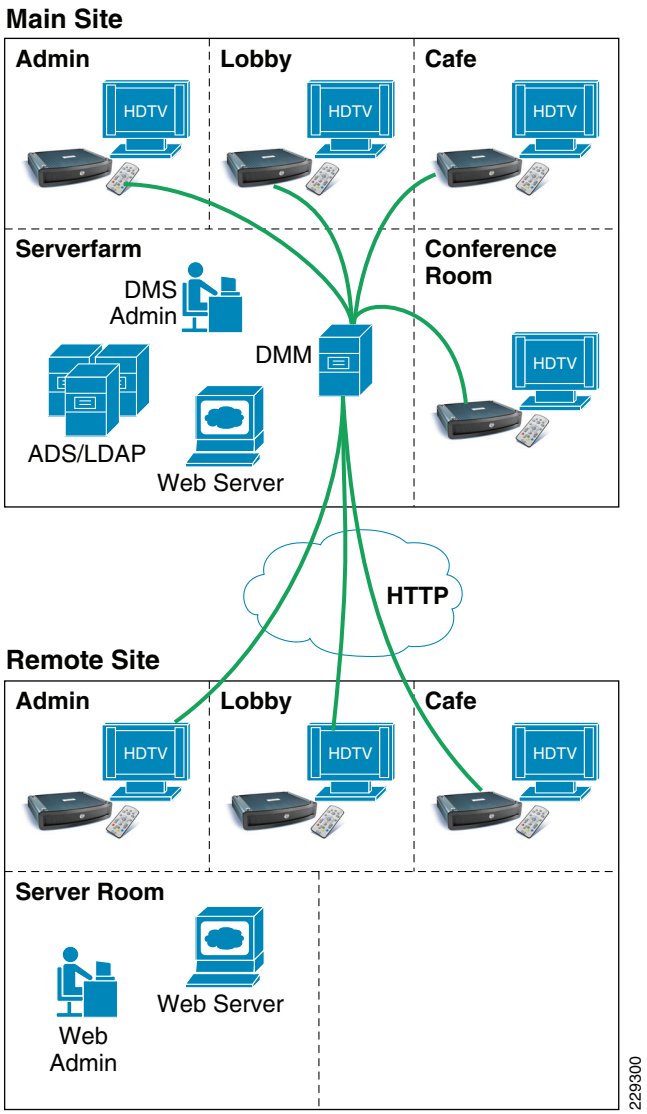

Distributed Content Storage Model

Digital signage content is typically wrapped with HTML or Adobe Flash applications that provide greater flexibility for a Web administrator to display more types of content from various sources on a single page. When deploying large numbers of DMPs with rich and static digital signage content, the unicast communication between DMP and the distributed content server may waste network bandwidth by retrieving the same content for continuous display. Hence it becomes important for the network architect to understand the content distribution and network level requirements to optimally deploy signage application in a campus network.

Depending on DMP scalability, overall network capacity, and the bandwidth allocation for digital signage application, Cisco DMS offers the following three distributed content storage solutions for Cisco DMP to pull and display the static content from the local network instead of downloading all of it through the WAN network:

- Cisco DMS-Content Distribution (CD)—Is an ideal solution for a small-scale enterprise network with few Cisco DMPs. Cisco DMM can push the static HTML or flash content via FTP or SFTP protocol to Cisco DMM on internal storage or to external storage device like USB drive and redirect Cisco DMP to access internal storage. This solution helps to minimize WAN bandwidth utilization.
- Local Content Server (Web or CIFS)—Storing the digital media content on a single local content server, like HTTP or CIFS sever, gives the network administrator more flexibility and management of the content distribution solution. The Web administrator can dynamically add, modify, and store the updated HTML or Flash file on a centralized server for Cisco DMP to retrieve and display. On the next HTTP request from DMP, the refreshed copy is displayed automatically. This solution provides more flexibility compared to Cisco DMS-CD solution, as it can dynamically update content without updating and managing content on each individual Cisco DMP.
- Cisco ACNS—Highly scalable and intelligent large size video content distribution to remote locations. Hierarchical content distribution system at main and remote sites distributes single copy of pre-recorded video to ACNS edge at the remote sites. To increase WAN network efficiency, Cisco ACNS leverages the caching technology and provides unicast VoD delivery in local LAN networks to end users instead of downloading one copy for each user over the WAN network.

Validated Content Distribution Model

To provide a simplified, scalable, and cost-effective content distribution and management solution in small enterprise network architecture, it is recommended to leverage local Web or CISF servers in the main and remote sites to store and publish local digital signage content. Cisco DMM can be programmed to redirect local DMPs to a local Web server and remote DMPs to pull the content from a local Web server. Such distributed content storage design minimizes the critical WAN bandwidth usage to publish local information. However, the WAN network may still be utilized to access global signage information, such as county or state level education and emergency news that can be broadcast by Web server from main site, and similarly real-time news ticker from the Internet can be embedded in major content that provides constant world-wide news updates.

Small enterprise network architects and Web administrators must perform pre-deployment exercise to assess the type of local versus distributed content (text/graphics/VoD/RSS) embedded in signage and the number of DMPs to be deployed in remote sites. This assessment provides WAN bandwidth guidelines to integrate digital signage in the network. As described earlier, Cisco WAN optimization solution like ACNS and WAAS must be integrated in the network if it demands higher WAN bandwidth. Figure 7-11 depicts the unicast communication flow between Cisco DMP deployed across the network and Web servers located in intranet and Internet domains.

Figure 7-11 Distributed Content Server Communication

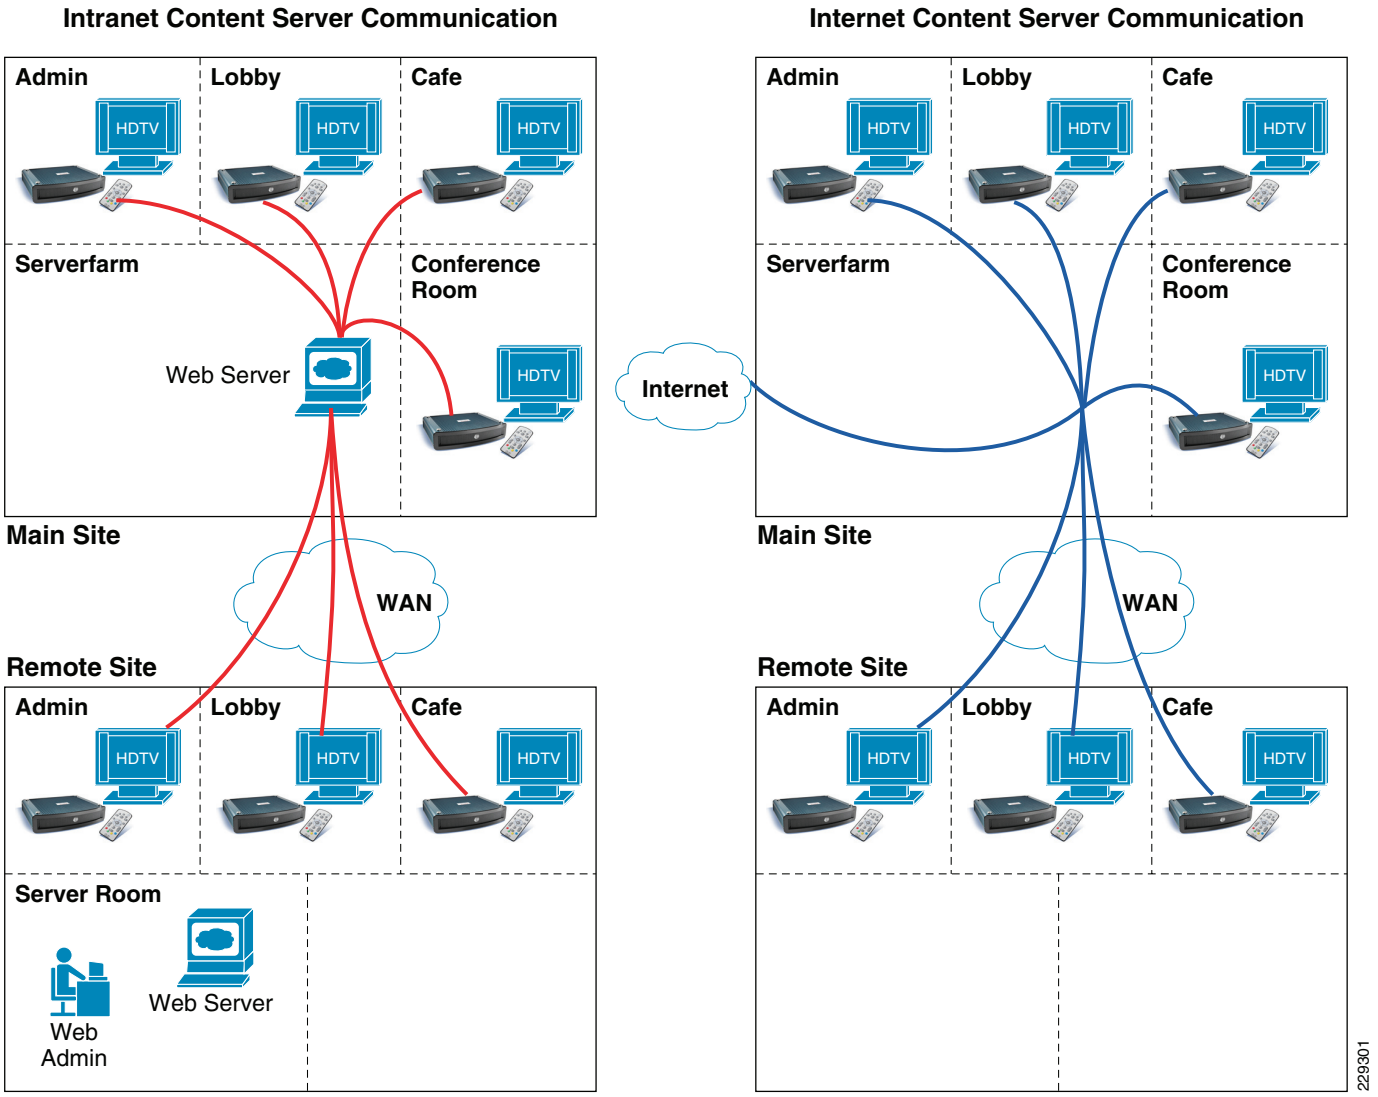

Implementing Network Services for Digital Signage

Prior to integrating digital signage applications, the network architect must make network services ready with the best practices for resilient and seamless operation and integration. Building the network as a highly-available platform is a foundational requirement for applications like IP telephony and digital media solutions as they demand constant bandwidth and network availability. Deploying centralized DMM with a distributed content server spans the digital communication beyond the campus boundary; hence it is recommended to deploy digital media solutions based on these network design principles:

- *Low latency*—Deploy the high-speed campus network that offers lower latency for real-time applications like voice and video.
- *High availability*—To increase network resiliency it is recommended to deploy the network with redundant modules, systems, and power supplies offering non-stop communication and sub-second network recovery during minor or major network failure events.

- QoS—Enhance user experience and content quality with robust QoS policies at the campus network edge.
- Confidentiality—Protect digital media end points, appliances, and data with centralized authentication and encryption.

This section provides access layer design and configuration guidelines to deploy Cisco DMP at the campus network edge and DMM in a centralized serverfarm in a main site. For more information on design and implementation guideline for building a strong and resilient core and foundational campus network, refer to Chapter 2, "Small Enterprise Design Profile(SEDP)—Network Foundation Design."

Deploying Cisco DMP in the Access Layer Network

Cisco DMP is an enterprise managed and trusted end point in the campus access layer, hence the network administrator must apply the common security and QoS policy for DMP as defined for other trusted endpoints like IP phones. In a typical deployment scenario, a single access layer switch may be connected to several other trusted and un-trusted endpoints, hence it becomes an important task for administrator to provide secure and suitable network services.

Cisco access layer switches provides the flexibility to deploy Cisco DMP in two different modes, manual deployment and plug-and-play Auto Smartport (ASP) macro.

Manual Deployment

Network administrator must manually implement the following three major network services to successfully integrate DMP in the network:

- Assigning unique Layer 2 VLAN
- Implement network edge security
- Implement network edge QoS

Assigning Unique Layer 2 VLAN

To provide secured and simplified digital signage communication, the DMP must be assigned a unique broadcast domain. De-coupling DMP with other trusted and un-trusted end points makes DMP more secure during any attacks and is easier to manage and troubleshoot. When single access layer connects to multiple DMPs, then all the DMPs can be assigned on the same Layer 2 VLAN. Like any other logical network partition design, it is recommended to use unique Layer 2 VLAN for DMP that are physically deployed on different Cisco access layer switches. Figure 7-12 provides recommended Cisco DMP Layer 2 segmentation guidelines in the access-layer:

Figure 7-12 Cisco DMP-Layer 2 VLAN Segmentation

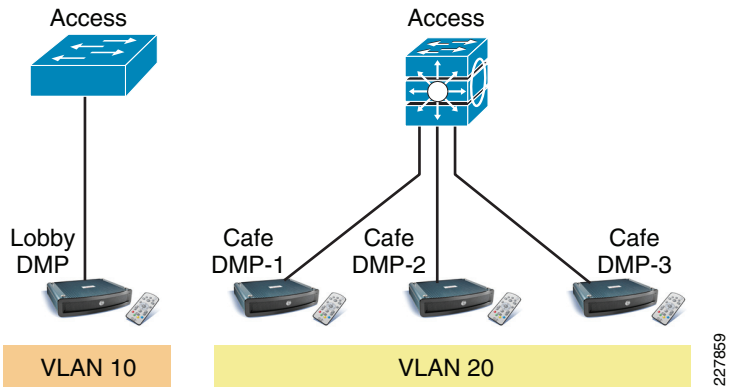

Cisco DMP cannot transmit or receive 802.1Q tagged frames, hence it is recommended to change default switchport mode from dynamic to access mode. The following is a sample configuration to enable VLAN in the database and apply VLAN on the DMP physical port:

2960-S

```
cr24-2960-DO(config-if)#interface FastEthernet0/7
cr24-2960-DO(config-if)# description CONNECTED TO LOBBY DMP
cr24-2960-DO(config-if)# switchport mode access
cr24-2960-DO(config-if)# spanning-tree portfast
cr24-2960-DO(config-if)# switchport access vlan 10
```

3750-X

```
cr25-3750-DO(config-if)#interface range GigabitEthernet 1/0/1 - 3
cr25-3750-DO(config-if-range)# switchport mode access
cr25-3750-DO(config-if-range)# spanning-tree portfast
cr25-3750-DO(config-if-range)# switchport access vlan 20
```

For flexible and scalable DMP deployment, it is recommended that the DMP edge port be in Layer 2 mode even when the access layer switch is deployed in a multilayer or routed access network design.

Implement Network Edge Security

Cisco DMP player requires communication with certain critical systems in the network, such as an IP gateway, Cisco DMM, Web servers, SNMP, and NTP. To display the digital signage content and synchronize with management servers, Cisco DMP receives more data from the network than transmitting to the network.

Cisco Catalyst integrated security feature must be deployed on the physical port to protect DMP from being attacked by viruses or unauthorized hosts. Based on the protocol and data communication characteristics of Cisco DMP, it is recommended to apply the following set of security configurations to protect the network and the DMP from unknown traffic floods and attacks:

Access

```
interface FastEthernet0/7
! Block transmitting all unknown unicast traffic
   switchport block unicast
! Enable port-security on this port
   switchport port-security
! Default, allow single-host to access this port
   switchport port-security maximum 1
! Block receiving BPDU from this port
   spanning-tree bpduguard enable
```

Implement Network Edge QoS

It is important to implement differential service treatment for digital media applications over non-critical network traffic in the network. Depending on the digital media applications and the distributed content, appropriate QoS services must be implemented at the network edge that connects media endpoints and in the serverfarm where typically centralized management and content servers are deployed. As described earlier, the Cisco DMP primarily use standard HTTP protocol to communicate with centralized DMM management server and the distributed Web server to publish the digital signage content.

By default, HTTP packets between digital media endpoints are set with default DSCP values and rely on intermediate network devices to classify the traffic and provide advanced QoS techniques to protect the digital media communication between DMP and other back end systems. Since the communication and publish content is delivered using HTTP protocol, it becomes challenging to distinguish between HTTP control traffic versus digital content in the campus network. Following RFC 4594 QoS deployment guidelines, the unicast control plane communication between Cisco DMM and DMP system can be classified as signaling traffic and must be marked an appropriate DSCP value and assigned a proper queue. Figure 7-13 provides QoS references to deploy digital media application in a campus network:

Figure 7-13 Digital Signage QoS Reference Chart

Application Class	PHB	Admission Control	Congestion Managment and Congestion Avoidance	Cisco Video Applications
VoIP Telephony	EF	Required	Priority Queue (PQ)	
Broadcast Video	CS5	Required	Optional (PQ)	Cisco DMS (Live Streams)/Enterprise TV/IPVS
Realtime Interactive	CS4	Required	Optional (PQ)	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco CUPC/CUVA/CI IP Phone 7985G
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco DMS (VoDs)
Network Control	CS6		BW Queue	
Call-Signaling	CS3*		BW Queue	DMM/DMP Control Traffic
OAM	CS2		BW Queue	
Transactional Data	AF2		BW Queue + DSCP WRED	Cisco WebEx/CU MeetingPlace
Bulk Data	AF1		BW Queue + DSCP WRED	
Best Effort	DF		Default Queue + RED	
Scavenger	CS1		Min BW Queue (Deferential)	YouTube/Xbox Live/iTunes/BitTorent/etc.

227860

*Cisco classification method slightly differs from RFC 4594. CS3 and CS5 definition are interchanged.

Applying Ingress QoS Policy on Access Layer

Network QoS policies must be set at the campus access edge to mark recommended DSCP bit for control or management traffic between DMM and DMP. HTML or Flash based digital signage content can remain in same best-effort class. The control traffic can be identified from digital signage content based on TCP flow between Cisco DMM and DMP. Figure 7-14 provides QoS marking guidelines between Cisco DMP, Cisco DMM appliance, and content server:

Figure 7-14 QoS Marking Between Digital Signage Components

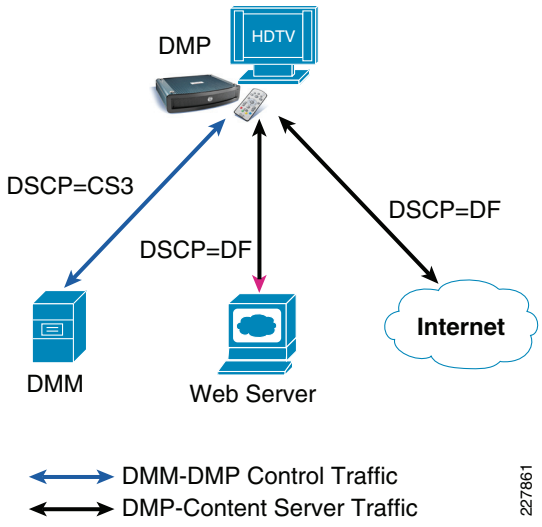

Based on TCP and static Cisco DMM and DMP player information, the following configuration guideline must be implement QoS policy on access layer switches that connect to Cisco DMP and Cisco DMM in a centralized serverfarm:

```
! Classify DMP and DMM HTTP traffic with extended ACL
ip access-list extended DMS-SIGNALING
remark DMM-DMP-MGMT
permit tcp host <DMP-IP-Address> host <DMM-IP-Address>
permit tcp host <DMM-IP-Address> host <DMP-IP-Address>
!
class-map DMS-SIGNALING
match access-group name DMS-SIGNALING
!
policy-map DMS-Policy
 class DMS-SIGNALING
 set dscp cs3  ! Explicit mark DSCP CS3
!
interface FastEthernet0/7
description CONNECTED TO LOBBY DMP
mls qos trust dscp
service-policy input DMS-Policy!Apply ingress service-policy
!

interface FastEthernet0/10
description CONNECTED TO Cisco DMM Appliance
mls qos trust dscp
service-policy input DMS-Policy!Apply ingress service-policy

cr24-3560r-DO#show mls qos interface fas0/7 | inc policy-map|dscp
Attached policy-map for Ingress: DMS-Policy
trust state: trust dscp
trust mode: trust dscp
```

Additional ingress QoS policies, such as policers, can be implemented on access switches if the network administrator is concerned about securing the restricting ports to consume higher bandwidth.

Applying Egress QoS Policy on Access Layer

Ingress QoS policy helps the network to distinguish between HTTP control and digital media content traffic within the campus backbone. Similar QoS techniques are required to provide differential services between control and digital media content traffic exiting the port connected to Cisco DMP and DMM appliance on the access layer switches. For global egress policy for trusted and un-trusted device, it is recommended to share the egress bandwidth to each hardware queue and enable prioritization for the low-latency traffic:

```
cr24-3560r-DO(config)#interface FastEthernet0/7
cr24-3560r-DO(config-if)# srr-queue bandwidth share 1 30 35 5 ! Enable BW share
cr24-3560r-DO(config)#priority-queue out! Enable Priority-Queueing

cr24-3560r-DO#show mls qos interface fast0/7 queueing
FastEthernet0/7
Egress Priority Queue : enabled
Shaped queue weights (absolute) :  25 0 0 0
Shared queue weights  :  1 30 35 5
The port bandwidth limit : 100  (Operational Bandwidth:100.0)
The port is mapped to qset : 1
```

To deploy QoS in a small enterprise campus network design, refer to the "Deploying QoS in Network" section on page 2-51.

Auto Smartport Macro Deployment

Cisco access layer switches provide a zero-touch or plug-n-play type network provisioning solution by dynamically detecting connected endpoints and automatically applying best practices and recommended configurations. Cisco Auto Smartport macro helps network architects to reduce the number of challenges in implementation when deploying complex configurations. Cisco validated design comprehensively validates multiple set of tools and technologies to solve the critical business problems. Cisco Auto Smartport leverages validated and recommended network configuration and parameters that dynamically provision the network without any user intervention. Implementing recommended and validated configuration parameters helps enterprise network administrators to automatically provide network and device security and optimize application performance.

Cisco Auto Smartport leverages several Layer-2 protocol techniques to dynamically detect the endpoint platform that intelligently triggers the function and applies the configuration from pre-defined recommended templates. To further increase operational efficiency, Cisco Auto Smartport removes all dynamically applied configurations when the device is un-plugged or removed from the network. The following is the list of Layer 2 technologies and endpoint types that Cisco Auto Smartport macros use to dynamically apply configurations:

- Layer 2 Technologies
  - Cisco Discovery Protocol (CDP)
  - IEEE AB - LLDP
  - 802.1x
  - MAC-Authentication Bypass (MAB)
  - Layer 2 Source MAC address
  - Ethernet OUI
- Supported Endpoint Platforms
  - IP Phones—Cisco and Avaya IP Phone

- Wireless Access-Points—Cisco AP 11xx series
- IP Video Surveillance—Cisco IPVS 25xx and 4xxx series camera
- Digital Media Player—Cisco DMP 4x00 series players

This section focuses on providing guidelines for the basic configuration needed to enable Cisco Auto Smartport to dynamically provision Cisco DMP in the network.

Cisco Auto Smartport macro leverages a simple shell function to execute the pre-defined configuration template embedded in the switch for each type of supported endpoint. Cisco DMP configuration gets executed dynamically based on the port event triggers. The following output provides the Auto Smartport event trigger and dynamic configuration guideline when it detects Cisco DMP on the physical port:

```
cr26-3750#show shell functions CISCO_DMP_AUTO_SMARTPORT

function CISCO_DMP_AUTO_SMARTPORT () {
!!Provision this configuration when Link up event is triggered
!!and Cisco DMP is detected:
    if [[ $LINKUP -eq YES ]]; then
        conf t
           interface  $INTERFACE
               macro description $TRIGGER
               switchport access vlan $ACCESS_VLAN
               switchport mode access
                switchport block unicast
                mls qos trust dscp
                spanning-tree portfast
                switchport port-security
                switchport port-security maximum 1
                switchport port-security violation shutdown
                spanning-tree bpduguard enable
                priority-queue out
            exit
        end
    fi
!!Remove dynamic configuration when Link Down event is triggered.
    if [[ $LINKUP -eq NO ]]; then
        conf t
            interface  $INTERFACE
               no macro description
               no switchport access vlan $ACCESS_VLAN
                no switchport block unicast
                no switchport port-security
                no switchport port-security maximum 1
                no switchport port-security violation shutdown
                no mls qos trust dscp
                no spanning-tree portfast
                no spanning-tree bpduguard enable
                no priority-queue out
                if [[ $AUTH_ENABLED -eq NO ]]; then
                    no switchport mode access
                fi
            exit
        end
    fi
}
```

Comparing the Auto Smartport macro configuration with a recommended manual configuration, it can be seen that the majority of the recommended manual configuration is in the macro template. Some of the advanced QoS parameters, such as MQC-based DSCP marking, may have to be manually configured.

Implementing Auto SmartPort Macro

Enterprise network administrators must enable Cisco Auto SmartPort Macro function along with basic network parameters on access layer switches to dynamically detect and provision the configuration for various types of endpoints. Enabling Cisco Auto Smartport macro function globally enables all the physical ports and provisions and un-provisions the network configuration for the Cisco DMP based on shell triggers and functions. It also provides the flexibility to disable the Auto Smartport function on a per-port basis where the static configuration is required. The following is the simple global configuration to enable Cisco Auto Smartport processing for all the physical ports:

3750-X

```
cr26-3750(config)#macro auto global processing

cr26-3750#show macro auto interface | inc Auto
 Global Auto Smart Port Status
 Auto Smart Ports Enabled
```

As described earlier, Cisco Auto Smartport can leverage multiple Layer 2 technologies to detect the endpoints, by default Auto SmartPort use the pre-defined MAC address-group range to dynamically detect the Cisco DMP based on Ethernet OUI address. To deploy Cisco DMP based on Ethernet OUI, no additional configuration is required. To enable secure access-control solution, Cisco Secure ACS and MAB can be integrated to authenticate Cisco DMP based on registered MAC address in the ACS database.

```
cr26-3750#show macro auto address-group CISCO_DMP_EVENT
MAC Address Group Configuration:
Group Name                              OUI               MAC ADDRESS
------------------------------------------------------------------------------
CISCO_DMP_EVENT  0023.AC
```

Because some of the network parameters like VLAN IDs are unique in the network, the network administrator needs to determine the common VLAN ID to deploy for a common set of end points. For example, when Cisco DMP is detected on Switch1, then it must dynamically detect the media player and assign it to appropriate VLAN and execute the network configuration template. By default, when Cisco Auto Smartport detects the Cisco DMP device on the port it configures the port in access-mode and assigned it a default VLAN ID = 1. After applying following single global configuration, the Auto Smartport automatically assigns all the Cisco DMP in to user-defined VLAN.

```
cr26-3750(config)#macro auto device media-player ACCESS_VLAN=58

cr26-3750#show macro auto device media-player | inc Device|ACCESS
Device:media-player
Configurable Parameters:ACCESS_VLAN
Defaults Parameters:ACCESS_VLAN=1
Current Parameters:ACCESS_VLAN=58
```

Tuning Auto SmartPort Macro

Cisco Auto Smartport is optimized in detecting the end points and provisioning the configuration for rapid and error-free deployments. The shell function that performs the configuration provisioning and un-provisioning task is based on physical link up and down events. In initial deployment, the endpoints can be detected using different Layer 2 techniques and all dynamically provisioned configurations can be saved in configuration files. Due to its nature the configuration is removed and then the same configuration is re-applied when the link goes down temporarily for any common reason, e.g., link flap, endpoint is power cycled, etc.

To make configuration persistent during such a link flap, the following single global configuration can be applied to retain the dynamic configuration during a link flap:

```
cr26-3750(config)#macro auto sticky
```

Implementing Cisco Digital Media Player

Once the recommended network edge configuration on the access layer is implemented, the Cisco DMP is ready to be deployed. Since Cisco DMP uses an embedded Linux OS which is not accessible directly to end users, the basic network parameters must be provisioned using Web-based Cisco DMP-Device Manager (DM). The Cisco DMP-DM is divided into three major configuration modes:

- Settings (Network/Browser/Storage)
- Display
- DMP Administration

The network administrator must configure the basic network parameters to deploy Cisco DMP in production network; the DMM administrator from the main site can apply the global display and management parameters to DMP without intervening network administrator for any advanced configuration task.

This chapter provides the following deployment guidelines to successfully deploy Cisco DMP and Cisco DMM appliance server communication in the network:

- Assigning IP address to Cisco DMP
- Registering Cisco DMP to Cisco DMM database

Assigning IP Address to Cisco DMP

The default IP setting on Cisco DMP is to dynamically acquire IP address and gateway information from DHCP server. It is recommended to assign a unique static IP address to DMP as it provides flexibility to the network administrator to provide secured DMP-DM GUI access with ACLs and the ability to provide distinguished QoS treatment for control traffic between Cisco DMP and DMM appliance server. Figure 7-15 is a simple IP configuration task that the network administrator must perform to change the default IP address method from DHCP to static IP address mode:

Figure 7-15 Assigning DMP Static IP Address

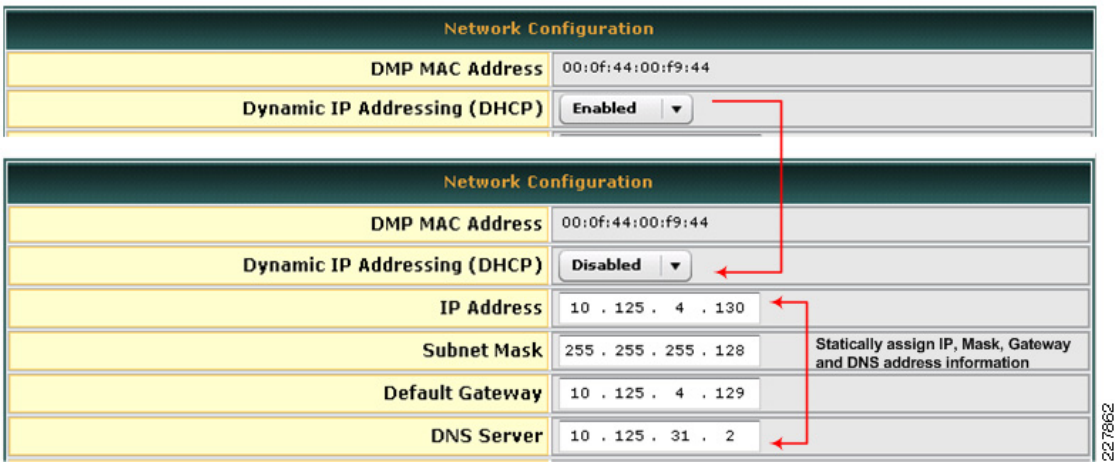

Registering Cisco DMP to Cisco DMM Database

The first step to enable communication between digital signage components is to register the Cisco DMP into the centralized Cisco DMM appliance database. Cisco DMM appliance requires basic information from Cisco DMM to register-MAC and IP address firmware version and storage information. Centralized DMM appliance can register large numbers of DMPs across the enterprise network and it may become an operational challenge to manage each DMP player. DMM administrator can create a logical DMP group along with the range of IP subnet; DMM-DSM automatically groups the registered DMPs based on assigned IP address. The following are some of the advantages in deploying DMP group in small enterprise network architecture:

- Organize registered DMPs into a single logical group.
- Instead of managing each individual DMP, the DMP group allows the DMM administrator to manage a group of DMPs collectively.
- Accelerate digital signage content deployment and instruction to the DMP group instead of individual players.
- Like content management, common display attributes to all DMPs in the DMP group.

The Cisco DMS solution provides flexibility to the DMM administrator for manual or automatic Cisco DMP registration into the database. Depending on the number of Cisco DMP deployed across the network, either of the registration process can be selected.

- *Manual DMP Registration*—The DMM administrator must manually enter Cisco DMP MAC and IP address information into the DMM database. For better operational management and troubleshooting, the DMP must be assigned to an logical DMP groups.
- *Auto DMP Registration*—Is a highly scalable, simplified, and error-free DMP registration solution for large deployments. To auto-register and auto-group the DMP in user-defined groups, the range of IP subnets or CIDR ranges must be specified to scan the Cisco DMP players in the network. Multiple IP subnets can be configured for single DMP group. The DMM appliance transmits TCP based unicast packet with pre-defined port number 7777 to scan Cisco DMP, which cannot be modified by the user. Hence for successful auto-registration process, it is recommended to make sure TCP port 7777 is not filtered anywhere in the network. The DMM admin can control the DMM appliance to trigger on-demand or schedule scan to locate DMP in the network for auto registration.

The DMM admin must complete these three steps sequentially to successfully deploy DMM groups and auto DMP registration in the DMM-DSM:

Step 1 Configure DMM Group and assign an IP subnet.

Figure 7-16 Configuring DMP Group Using DMM-DSM

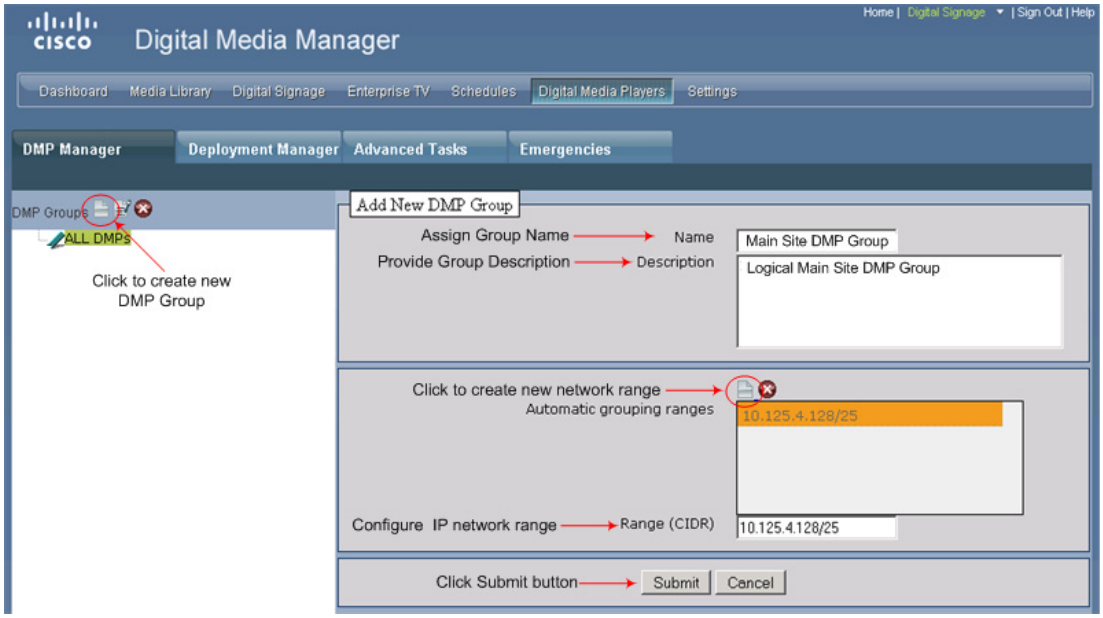

Step 2 Configure DMP Discovery Application and assign an IP subnet.

Figure 7-17 Configuring DMP Discovery Application

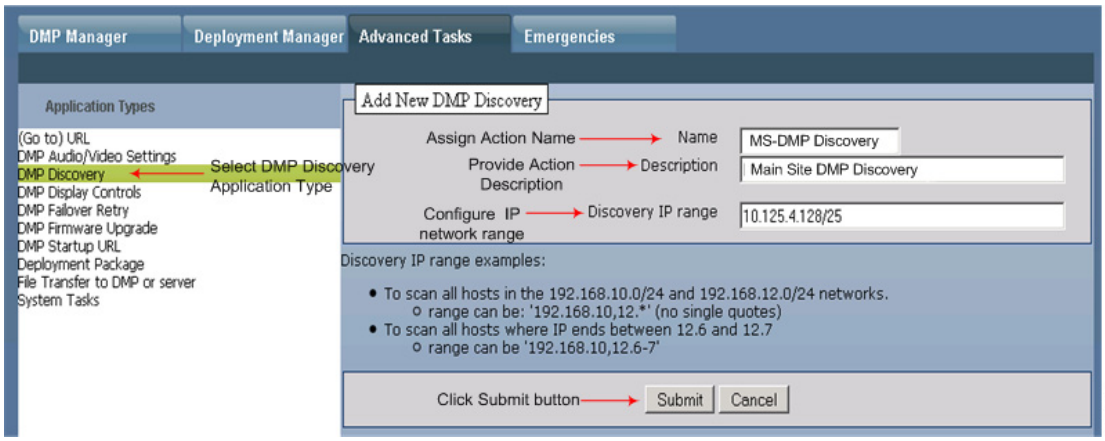

Step 3 Trigger the DMP discovery with manual action or schedule to discover in future.

Manual DMP Discovery Trigger

Figure 7-18 Triggering DMP Discovery Manually

Refreshing the window in few seconds will reflect the dynamically discovered Cisco DMP that gets automatically registered and grouped as depicted in Figure 7-19. The default name of the auto-registered DMP is the same as their MAC address; the DMM admin must change to reflect with proper name or location name.

Figure 7-19 Dynamically Discovered DMP Discovery Triggered Manually

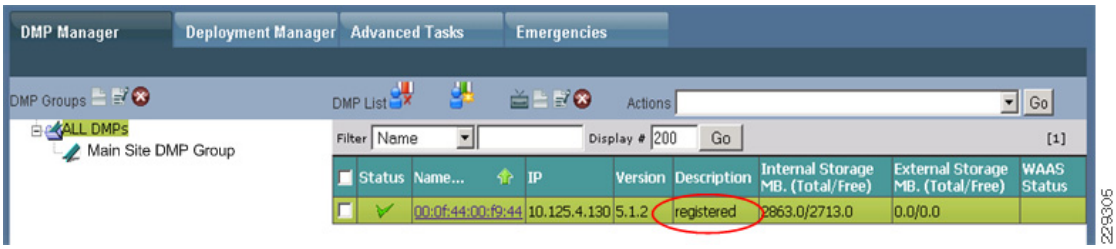

Scheduling DMP Discovery

Depending on the number of DMP groups, network ranges, and DMP players in the network, the DMP discovery may take some time. In large deployments, it is recommended to schedule Cisco DMP discovery during non-business hours. Scheduling DMP discovery in the network is identical to scheduling the digital signage publishing time. To schedule DMP discovery using Cisco DMM-DSM:

1. Click on **Schedules ->Play in Future** and select discovery month and date
2. Select the DMP group to be discovered and click on **Add an Event** button.
3. Select DMP group from Select Group Tab and click **OK**.
4. Select Advanced Task from Task Type Tab and click on Select Advanced Tasks Tab.
5. Select DMP Discovery from Types and the Action Name and click **OK**.
6. Configure Start and Stop Action run time and optionally configure repeat value to dynamically discover Cisco DMP based on schedule.

Implementing Digital Signage

Upon successfully discovering and registering the Cisco DMP in the DMM appliance database, the DMM administrator can start preparing to implement digital signage in the network. The provision checklist must include exact content path and schedule to display the content on individual or group

DMP in the network. As described earlier, the content must be stored on a Web or CISF server that is physically located on the same campus network as a Cisco DMP. Hence when creating the playlists, it is recommended to specify the URL path where the content is stored.

Publishing digital signage requires three simple configuration step on DMM-DSM as depicted in Figure 7-20.

Figure 7-20 Digital Signage Deployment Steps

Executing each step populates content information in DMM in the common repository, provides flexibility to compile playlist with distributed content, and schedules the digital signage publishing time by mapping the playlist to individual or group DMPs. The network topology in Figure 7-21 is used as a reference point to configure each deployment step.

Figure 7-21 Digital Signage Network Topology

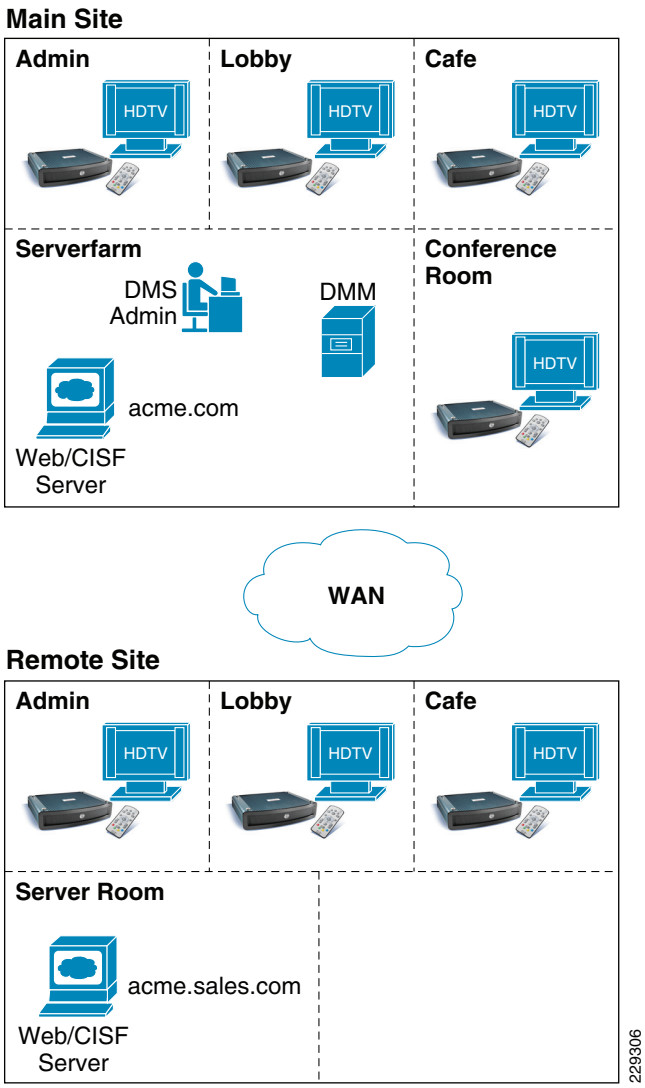

Step 1 Adding Asset in Media Library.

Cisco DMM-DSM builds an asset of digital media content in a common Media Library database. Figure 7-22 provides information about the variety of digital media content types supported and can be stored in two major locations—remote Web/CISF content server or it must be uploaded to local DMM appliance server. As described earlier, the recommended content distribution model is to keep content distributed on remote servers that reside on the same LAN as Cisco DMPs. Cisco DMM appliance must not be used as a content server. Each asset type must be entered one at a time in the media library. Each asset is executed serially within the playlist; the Cisco DMM-DSM also provides flexibility to publish digital signage content in random order.

Figure 7-22 Supported Asset Types

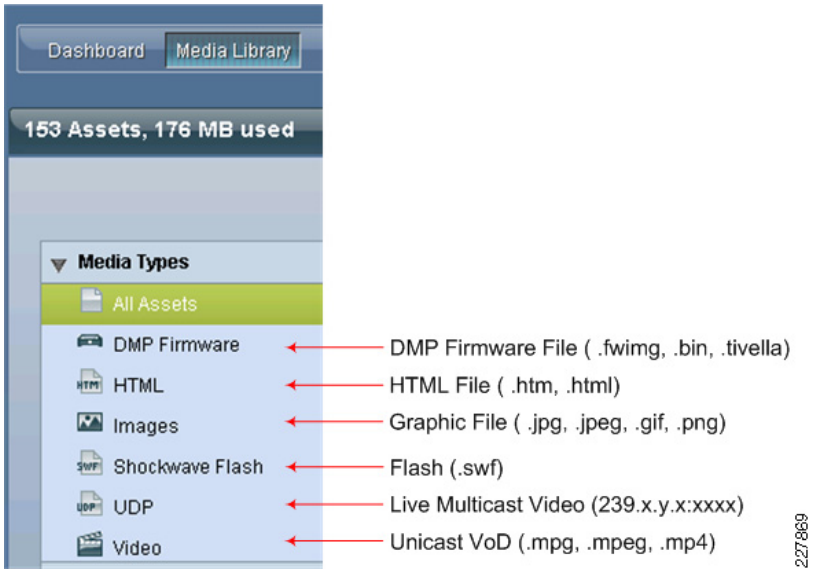

Execute the following steps to add distributed digital signage (non-video) content asset into Media Library:

1. Click on **Add Media Asset** button.
2. Click on **Single** tab.
3. Click URL in radio button as a Source and type the exact URL to access content (e.g., http://creek.rosenelemetary.com/default.html). Prior to deploying, the content must be tested and verified by applying same URL from local internet browser.
4. Do not click on download check box. This will prevent downloading provided HTTP URL content to the local DMM hard-drive.
5. Select **File Type** from drop down window based on URL extension.
6. Enter estimated or planned playback time for this asset.
7. Select the media category from Category window in which this URL fits in.
8. Optionally, provide description and content owner/developer information.
9. Click on **Save** button to review the entered information.
10. Click the **Close** button to exit the window.

Step 2 Creating Playlist.

Playlist is a user-defined logical group that is packaged with the compiled list of distributed digital signage assets which are being added in the shared Media Library database. For example, a playlist can include the intranet home page of the main and remote site and corporate document catalog developed in Adobe Flash.

Cisco DMM-DSM provides the flexibility to develop playlists in two different modes—Standalone and Cisco Digital Media Designer (DMD). When the digital signage content is designed and developed outside the DMM-DMD, then DMM administrator must create a standalone playlist.

Cisco Digital Media Designer (DMD) is a Java-based, powerful, drag and drop design tool to create customized digital signage content. Cisco DMM offers pre-designed assets that can be leverage to create personalized design. DMD also offers flexibility to customized horizontal and vertical screen display orientation.

Figure 7-23 Digital Signage Playlist Options

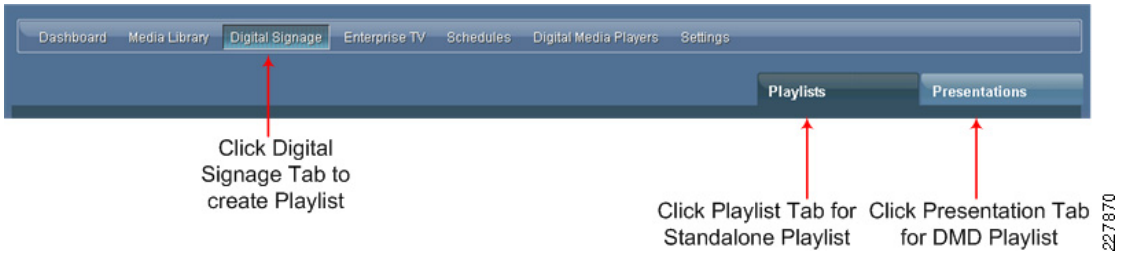

This section provides guideline to implement a standalone playlist. Multiple playlists can be created and associated to an individual or group of DMPs. When designing the asset in the playlist, it is important to remember that the content publishing time and mappings to DMP groups are applied based on per-playlist and not on per-asset basis. The order of playlist execution is done serially based on specified time. Figure 7-24 depicts the playlist order and duration time on per-cycle basis.

Figure 7-24 Digital Signage Display Order and Duration

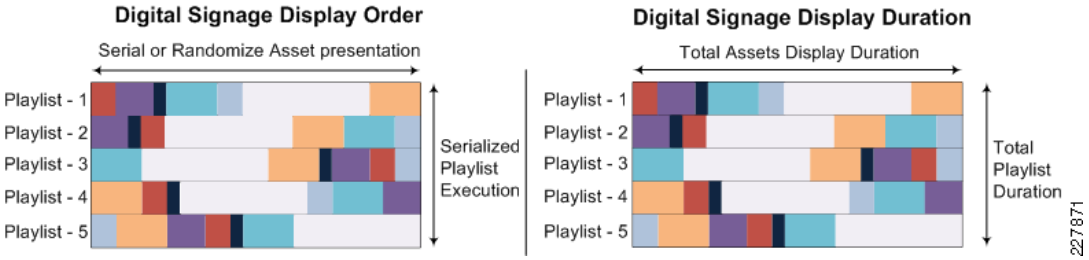

In a best practice, the playlist can be created based on individual main or remote site departments that can provide flexibility to announce the news or any other department specific content at specific time without impacting the playlist for other DMP groups deployed in different departments. Execute the following steps to compile the distributed digital signage asset and estimated or planned run time for each individual asset.

Click on **Create Playlist** button and follow step-by-step instruction provided in Figure 7-25 to create a complied and distributed digital signage content in playlist.

Figure 7-25 Compiling Assets with Standalone Playlist

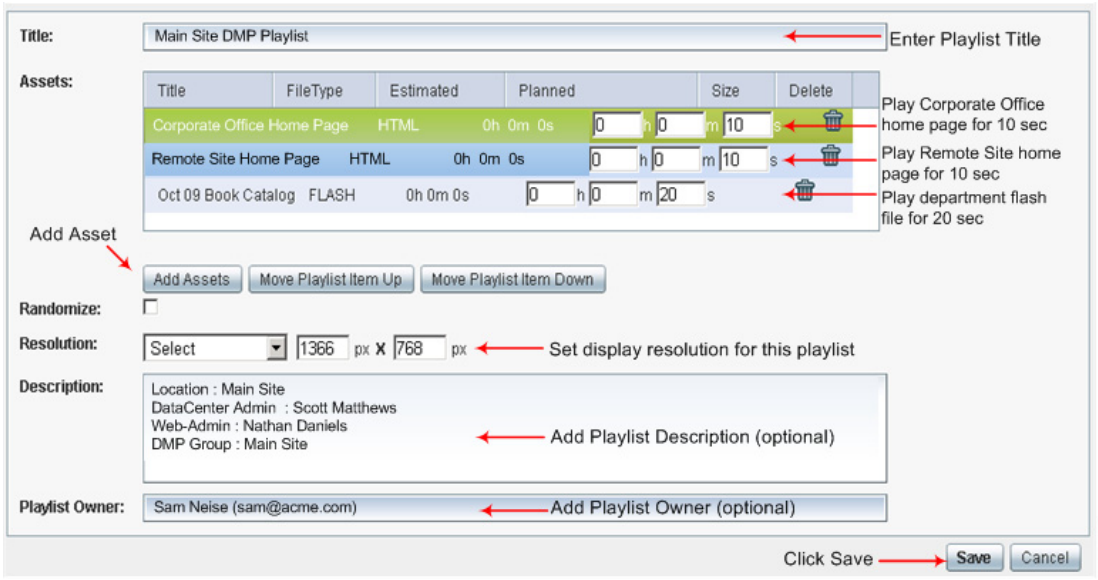

Step 3 Scheduling Playlist.

After successfully completing the above two steps the URL of distributed digital signage content is added in media library database and the playlist is compiled with the list of signage content that needs to be played for DMP group. The DMM administrator can send the playback command in two different modes—Instant Play and Future Play.

Instant Play or "Play Now" sends the playback command to selected DMP groups and all the DMPs can start displaying digital signage content immediately. Instant Play option provides an option for immediate signage content deployment; it can also be used to publish the newly added or updated signage asset in an already playing playlist. When Cisco DMP receives the new and updated playlist command from centralized DMM appliance, it aborts playing the previous playlist commands and immediately starts the display based on new information.

Future Play or "Play in Future" gives flexibility to the DMM administrator to pre-deploy digital signage in the DMM appliance database and schedule to play the playlist on a pre-determined month, date, or specific time. For example, you could have next month's café lunch menu automatically published in the last week of the current month.

Figure 7-26 depicts tab selection to schedule the playlist in both deployments modes.

Figure 7-26 Playlist Scheduling Modes

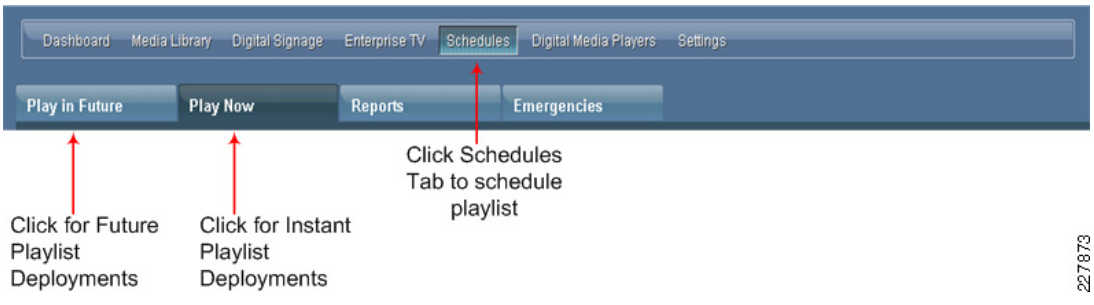

Implementing Instant Play Mode

Execute the step-by-step procedure as depicted in Figure 7-27 to deploy digital signage instantly in the network. Cisco DMM-DSM provides the flexibility to select single, multiple, or all the DMPs using shift-key to instantaneously publish digital signage in large deployments.

Figure 7-27 Publishing Instant Digital Signage

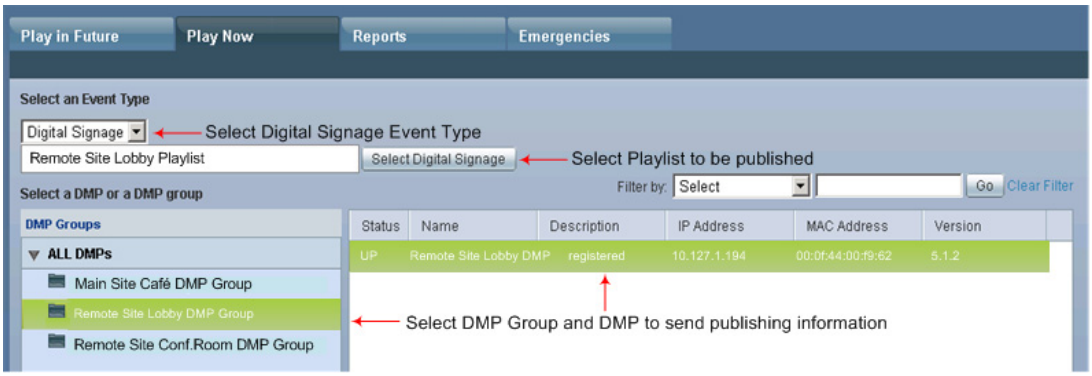

Implementing Future Play Mode

Click on **Play** in Future tab to schedule a future digital signage content publishing time. Click on **Add an Event** button from the bottom window and follow the step-by-step configuration guideline as displayed in Figure 7-28 to schedule future signage deployment.

Figure 7-28 Scheduling Signage Deployment

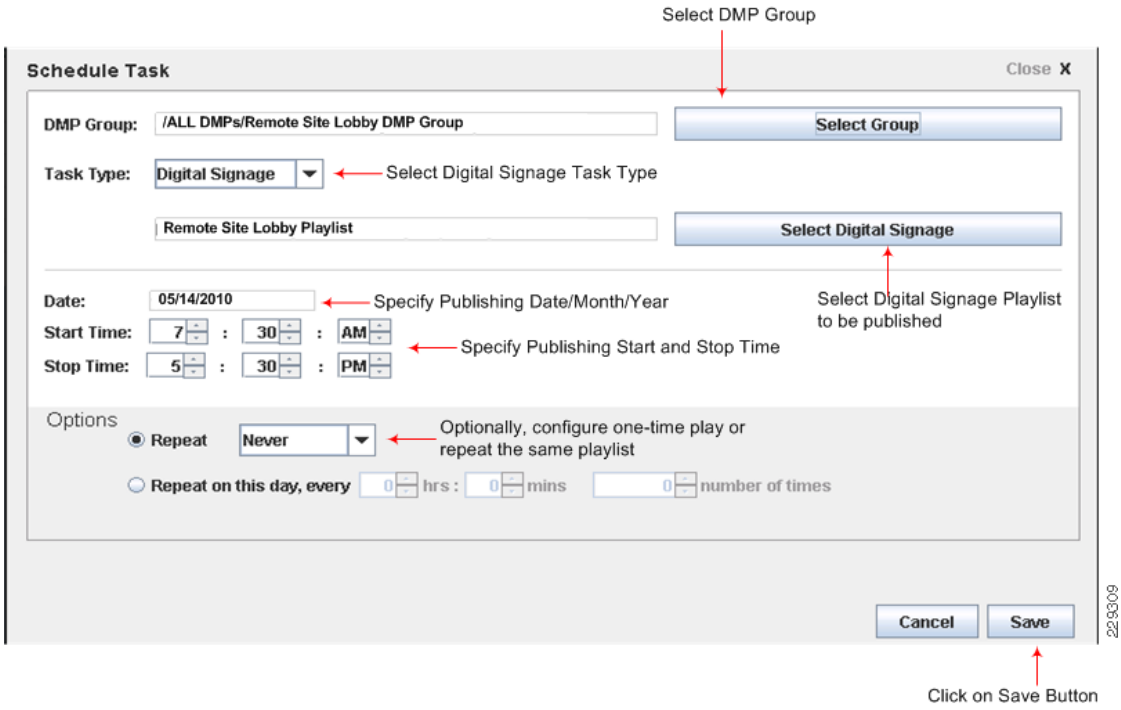

Summary

The Cisco Digital Media Suite (DMS) is a comprehensive, scalable, and network-centric solution that offers next-generation enterprise network for enabling business communication through social video, digital signage, and IPTV systems. The emerging video technology can help transform the organization to learn, grow, communicate, and collaborate between employees, partners, and customers. Cisco DMS modules help increase customer experience, communication, and training to improve productivity at reduced cost. Cisco DMS enables Web 2.0 technologies to users to share innovative ideas and expertise through wide range of video collaboration techniques.