



# Medium 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](http://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

#### **Medium Enterprise Design Profile (MEDP)—Service Fabric Design Considerations 1-1**

Service Fabric Design	1-1
Main and Large Site Design	1-2
Medium Site Design	1-3
Small Site Design	1-3
Building Profiles	1-3
Large Building Design	1-3
Medium Building Design	1-3
Small Building Design	1-3
Extra Small Building Design	1-4
Access Devices	1-4
LAN/WAN Design Considerations	1-4
LAN Design Considerations	1-4
Routing Protocol Selection Criteria	1-4
High Availability Design Considerations	1-5
Access Layer Design Considerations	1-5
LAN Service Fabric Foundational Services	1-6
WAN Design Considerations	1-6
WAN Transport	1-6
WAN Service Fabric Foundational Services	1-7
Security Design Considerations	1-7
Mobility	1-7
Unified Communications	1-8
Call Processing Considerations	1-8
Gateway Design Considerations	1-8
Dial Plan Considerations	1-9
Survivability Considerations	1-9

---

### CHAPTER 2

#### **Medium Enterprise Design Profile (MEDP)—LAN Design 2-1**

LAN Design	2-1
LAN Design Principles	2-4
Medium Enterprise LAN Design Models	2-7
Main Site Network Design	2-9

Remote Large Campus Site Design	2-10
Remote Medium Campus Site Design	2-11
Remote Small Campus Network Design	2-12
Multi-Tier LAN Design Models for Medium Enterprise	2-13
Campus Core Layer Network Design	2-13
Core Layer Design Option 1—Cisco Catalyst 6500-E-Based Core Network	2-14
Core Layer Design Option 2—Cisco Catalyst 4500-E-Based Campus Core Network	2-15
Core Layer Design Option 3—Cisco Catalyst 4500-E-Based Collapsed Core Campus Network	2-17
Campus Distribution Layer Network Design	2-18
Distribution Layer Design Option 1—Cisco Catalyst 6500-E Based Distribution Network	2-19
Distribution Layer Design Option 2—Cisco Catalyst 4500-E-Based Distribution Network	2-21
Distribution Layer Design Option 3—Cisco Catalyst 3750-X StackWise-Based Distribution Network	2-22
Campus Access Layer Network Design	2-23
Access Layer Design Option 1—Modular/StackWise Plus/FlexStack Access Layer Network	2-24
Access Layer Design Option 2—Fixed Configuration Access Layer Network	2-24
Deploying Medium Enterprise Network Foundation Services	2-25
Implementing LAN Network Infrastructure	2-25
Deploying Cisco Catalyst 6500-E in VSS Mode	2-26
Deploying Cisco Catalyst 4500-E	2-34
Deploying Cisco Catalyst 3750-X StackWise Plus	2-38
Deploying Cisco Catalyst 3560-X and 2960-S FlexStack	2-41
Designing EtherChannel Network	2-41
Network Addressing Hierarchy	2-49
Network Foundational Technologies for LAN Design	2-50
Designing the Core Layer Network	2-50
Designing the Campus Distribution Layer Network	2-56
Designing the Multilayer Network	2-56
Spanning-Tree in Multilayer Network	2-59
Designing the Routed Access Network	2-60
Multicast for Application Delivery	2-64
Multicast Addressing Design	2-64
Multicast Routing Design	2-65
Designing PIM Rendezvous Point	2-66
Dynamic Group Membership	2-73
Designing Multicast Security	2-74
QoS for Application Performance Optimization	2-75
Medium Enterprise LAN QoS Framework	2-76

Designing Medium Enterprise LAN QoS Trust Boundary and Policies	2-79
Medium Enterprise LAN QoS Overview	2-80
Deploying QoS in Campus LAN Network	2-84
QoS in Catalyst Fixed Configuration Switches	2-84
QoS in Cisco Modular Switches	2-85
Deploying Access-Layer QoS	2-87
Deploying Network-Layer QoS	2-105
High-Availability in LAN Network Design	2-119
Medium Enterprise High-Availability Framework	2-119
Baselining Campus High Availability	2-120
Network Resiliency Overview	2-121
Device Resiliency Overview	2-122
Operational Resiliency Overview	2-124
Design Strategies for Network Survivability	2-126
Implementing Network Resiliency	2-127
Implementing Device Resiliency	2-130
Implementing Operational Resiliency	2-141
Summary	2-150

## CHAPTER 3

### Medium Enterprise Design Profile (MEDP)—WAN Design 3-1

WAN Design	3-1
WAN Transport	3-3
Private WAN Service	3-3
Internet Service	3-4
Metro Service	3-4
Leased-Line Service	3-7
WAN Aggregation Platform Selection in the Medium Enterprise Design Profile	3-7
Main Site WAN Aggregation Platform Selection	3-8
Large Remote Site WAN Aggregation Platform Selection	3-10
Medium Remote Site WAN Aggregation Platform Selection	3-10
Small Remote Site WAN Aggregation Platform Selection	3-11
Implementation of WAN Reference Design	3-11
WAN Infrastructure Design	3-11
Leased-Line Service	3-12
Routing Design	3-13
QoS	3-19
QoS Implementation	3-21
QoS Implementation at WAN Aggregation Router 1	3-22
Implementation Steps for QoS Policy at WAN Aggregation Router 1	3-24

QoS Policy Implementation for WAN Aggregation Router 2	3-26
QoS Policy Between the Main Site and Large Remote Site	3-29
QoS Policy Between the Main Site and Medium Remote Site Location	3-30
QoS Policy Between Main Site and Small Remote Site Location	3-32
QoS Policy Implementation Between the Main Site and Core	3-33
QoS Policy Between Large Remote Site and Main Site Location	3-34
QoS Policy Between Remote Medium Site and Main Site Location	3-36
QoS Policy Implementation Between Small Remote Site and Main Site Location	3-37
Redundancy	3-38
Multicast	3-43
Summary	3-46

## CHAPTER 4

### Medium Enterprise Design Profile (MEDP)— Mobility Design 4-1

Mobility Design	4-1
Accessibility	4-5
WLAN Controller Location	4-7
WLAN Controller Connectivity	4-8
Controller Connectivity to the Wired Network	4-8
Controller Connectivity to Wireless Devices	4-10
Access Points	4-19
Usability	4-26
Quality-of-Service	4-26
Guest Access	4-27
Manageability	4-32
Reliability	4-35
Controller Link Aggregation	4-35
Controller Redundancy	4-38
AP Controller Failover	4-40
Wireless LAN Controller Configuration	4-40
WLAN Controller and Wired Network Connections	4-41
Remote Site	4-43
Mobility Groups	4-43
WLAN Configuration	4-45
Staff Data WLAN	4-45
Staff Voice WLAN	4-46
Guest Access WLAN	4-46
WLAN QoS	4-49
Access Point Configuration	4-50
AP 1520 Configuration	4-51



Adding the AP1520 MAC Address to the WLAN Controller	4-52
Configuring the AP1520 as a Root Access Point (RAP)	4-52
WCS Configuration	4-54
WCS Users and User Groups	4-54
WCS Virtual Domains	4-55
Reference Documents	4-57

## CHAPTER 5

### Medium Enterprise Design Profile (MEDP)—Network Security Design 5-1

Security Design	5-1
Network Foundation Protection	5-6
Internet Perimeter Protection	5-8
Internet Border Router Security	5-10
Internet Firewall	5-10
Cisco ASA Botnet Traffic Filter	5-11
Intrusion Prevention	5-13
Cisco IPS Global Correlation	5-14
E-Mail Security Guidelines	5-17
Web Security Guidelines	5-22
Data Center Protection	5-26
Network Access Security and Control	5-28
Cisco Catalyst Integrated Security Features	5-29
Cisco Unified Wireless Network (CUWN) Integrated Security Features	5-29
Cisco Identity-Based Network Services (IBNS)	5-30
IEEE 802.1X Protocol	5-30
802.1X and EAP	5-31
Impacts of 802.1X on the Network	5-31
802.1X in Medium Enterprise Networks	5-32
Cisco NAC Appliance	5-32
NAC Appliance Components	5-33
NAC Appliance Modes and Positioning	5-35
NAC Deployment in the Medium Enterprise Design Profile	5-39
Secure Mobility	5-44
Threats Mitigated	5-46
Medium Enterprise Network Security Deployment Guidelines	5-47
Internet Border Router Edge ACL Deployment	5-47
Module 1—Implement Anti-spoofing Denies	5-47
Module 2—Implement Explicit Permits	5-48
Module 3—Implement Explicit Deny to Protect Infrastructure	5-48
Module 4—Implement Explicit Permit for Traffic to the Enterprise Public Subnet	5-48

Internet Firewall Deployment	5-48
Firewall Hardening and Monitoring	5-50
Network Address Translation (NAT)	5-52
Firewall Access Policies	5-52
Firewall Redundancy	5-55
Routing	5-56
Botnet Traffic Filter	5-57
Intrusion Prevention Deployment	5-61
Deploying IPS with the Cisco ASA	5-61
IPS Global Correlation Deployment	5-61
Web Security Deployment	5-66
Initial System Setup Wizard	5-67
Interface and Network Configuration	5-67
WCCP Transparent Web Proxy	5-71
Web Access Policies	5-74
Catalyst Integrated Security Features Deployment	5-75
NAC Appliance Deployment	5-76
NAC Deployment for Wired Clients	5-76
NAC Deployment for Wireless Clients	5-89
Additional Information	5-98

**APPENDIX A**

**Reference Documents A-1**



# CHAPTER 1

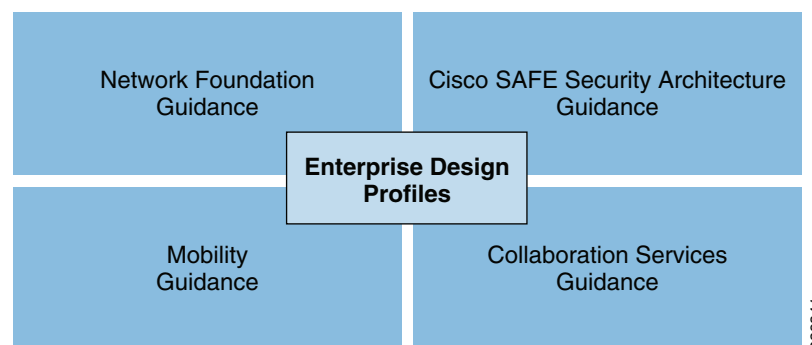
## Medium Enterprise Design Profile (MEDP)—Service Fabric Design Considerations

The service fabric is the foundational network that all enterprise services, applications, and solutions use to interact and communicate with one another. The service fabric is the most important component of the Medium Enterprise Design Profile. If it fails, all applications, solutions, and technologies deployed in the Medium Enterprise Design Profile will also fail. Like the foundation of a house, the service fabric must be constructed in a fashion that supports all the applications and services that will ride on it. Additionally, it must be aware of what is type of traffic is transversing and treat each application or service with the right priority based on the needs and importance of that application.

The service fabric is made up of four distinct components local and wide area network (LAN/WAN), security, mobility, and unified communications. Each of these critical foundation components must be carefully designed and tuned to allow for a secure environment that provides business continuity, service awareness and differentiation, as well as access flexibility.

See [Figure 1-1](#).

**Figure 1-1**      **Service Fabric Foundation Network**



## Service Fabric Design

The model used for the Medium Enterprise Design Profile service fabric is based around the desire to represent as many medium enterprise environments as possible. To do that a modular design is used, represented by sites and buildings of varying sizes (see [Figure 1-2](#)). The sites are made up of one or more building, depending on the site size profile; buildings are also sized with the determining factor being the number of users or connections to the network in that building as well as physical size. When

**Figure 1-2** *Medium Enterprise Design Profile Overview*



The main and large site designs are meant to represent significantly sized sites containing the largest user populations. The profile of the main/large site is made up of six buildings, the buildings range in size from large to extra small. The buildings will connect back to the resilient core via multiple 10Gb Ethernet links. The core will also connect to a serverfarm design and service block. The large site will connect to the main site via a 1Gb Metro Ethernet link. The main site and large site are almost identical, with the exception that the main site is connected to outside entities such as the Internet using the Internet edge components, and will also have all other sites within the enterprise connecting to it.

## Medium Site Design

The medium site design is targeted at enterprise sites that have approximately 3 buildings ranging in size from medium to small. The buildings will connect to the medium site core via multiple 10Gb links, and the core will also connect to a small serverfarm and service block. The medium site is connected to the main site via a 100mb Metro Ethernet link. This link interconnects the medium site to the other sites as well as external networks such as the Internet.

## Small Site Design

The small site profile represents a site made up of just one building; in this case, the core and distribution networks are collapsed into one. The small site is connected to the main site via a fractional DS3 with a 20mb bandwidth rating. This link interconnects the small site to the other sites as well as external networks such as the Internet.

## Building Profiles

There are four building profiles: large, medium, small, and extra small. All buildings have access switches that connect users. The buildings also have distribution switches that connect the access switches together as well as connect the building itself to the core network.

### Large Building Design

The large building is designed for 1600 Ethernet access ports ranging in bandwidth from 100mb to 1Gb. The ports are distributed over four different floors, each floor having 400 access ports. There are 80 wireless access points using the IEEE 802.11 ABGN standards, there are 20 access points per floor; additionally, there are 6 outdoor mesh access points to cover the outdoor skirt of the building. The large building designed for 160 phones.

### Medium Building Design

The medium building was designed for 800 Ethernet access ports ranging in bandwidth from 100mb to 1Gb. The ports are distributed over two different floors, each floor having 400 access ports. There are 40 wireless access points using the IEEE 802.11 ABGN standards, there are 20 access points per floor; additionally, there are four outdoor mesh access points to cover the outdoor skirt of the building. The medium building is made up of designed for 80 phones.

### Small Building Design

The small building is designed for 200 Ethernet access ports ranging in bandwidth from 100mb to 1Gb. The ports are all located on one floor. There are 10 wireless access points using the IEEE 802.11 ABGN standards; additionally, there are 2 outdoor mesh access points to cover the outdoor skirt of the building. The small building is designed for 30 phones.

## Extra Small Building Design

The extra small building is designed for 48 100mb Ethernet access ports. The ports are all located on one floor. There are 3 wireless access points using the IEEE 802.11 ABGN standards; additionally, there is 1 outdoor mesh access point to cover the outdoor skirt of the building. The extra small building designed for up of 10 phones.

## Access Devices

The devices that connect to the Medium 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 service fabric consists of four major components. The sections below provide a brief description of each of these components.

## LAN/WAN Design Considerations

The service fabric LAN/WAN is made up of routers and switches deployed in a three-tier hierarchical model that use Cisco IOS to provide foundational network technologies needed to provide a highly available, application-aware network with flexible access.

## 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*—Demarcation point between network infrastructure and access devices. Designed for critical network edge functionality to provide intelligent application and device aware services.

## 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 service fabric routers and switches support many different routing protocols that will work medium enterprise 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.

## 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 service fabric design must ensure network survivability by following three major resiliency methods pertaining to most types of failures. Depending on the network system tier, role, and network service type the appropriate resiliency option should be deployed:

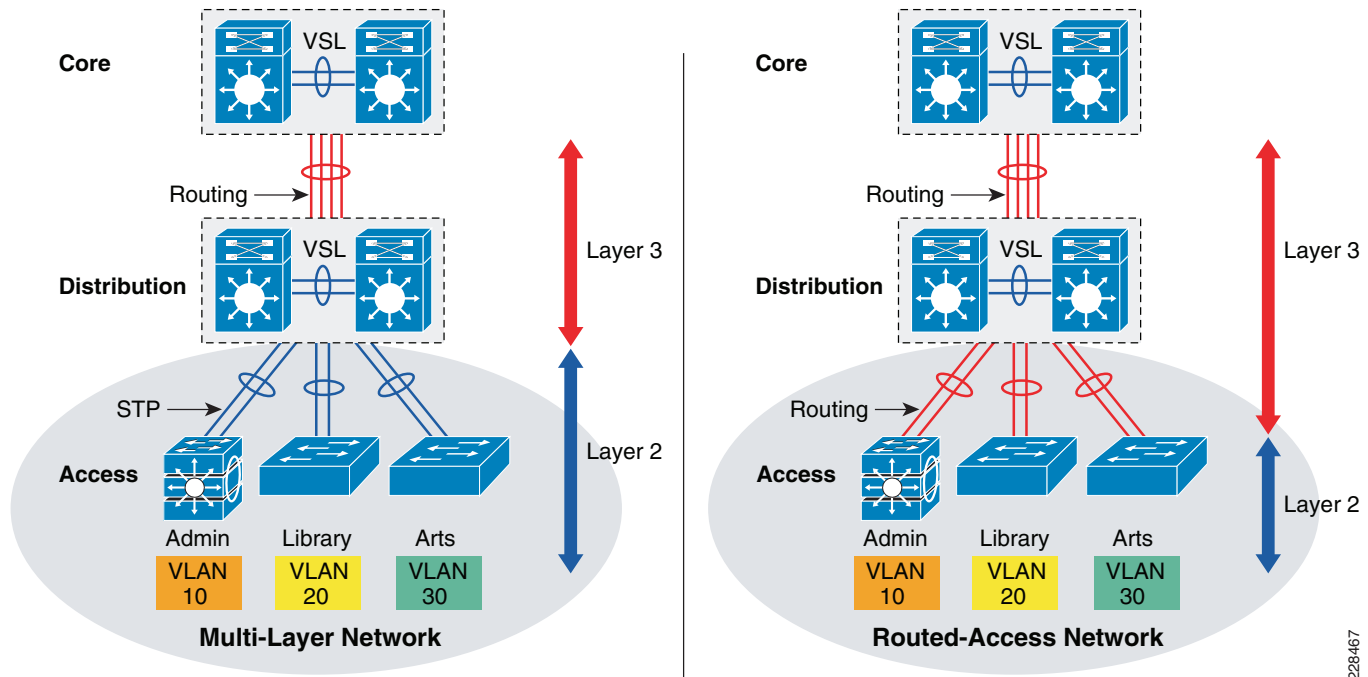
- *Link resiliency*—Provides redundancy during physical link failures (i.e., fiber cut, bad transceivers, incorrect cablings, etc.)
- *Device resiliency*—Protects network during abnormal node failure triggered by hardware or software (i.e., software crashes, non-responsive supervisor etc.)
- *Operational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

## 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 up to the network distribution, and the layer of communication used 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 function one tier down on Layer 3 access switches changes traditional multilayer network topology and the forwarding path. The implementing of a Layer 3 access 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 becomes a Layer-2 demarcation point to locally connected endpoints that could be logically segmented in multiple VLANs.

## LAN Service Fabric Foundational Services

The service fabric uses essential foundational services to efficiently disseminate information that are 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 medium enterprise is paramount. Some of the key foundational services discussed include the following:

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

## WAN Design Considerations

### WAN Transport

In order for sites to communicate with one another and/or to communicate outside the medium enterprise network, the network traffic must traverse over a WAN. WAN transport differs greatly from LAN transport due to the 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:

- MPLS/VPN
- Internet
- Metro Ethernet

## WAN Service Fabric Foundational Services

Similar to the LAN, the WAN must deploy essential foundational services to ensure the proper transport and prioritization of medium enterprise services, the WAN Service Fabric Foundation Services considered are as follows:

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

## Security Design Considerations

Security of the Medium Enterprise Design Profile service fabric is essential. Without it, medium enterprise solutions, applications, and services are open to be compromised, manipulated, or shut down. The service fabric was developed with the following 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 external (extranets) networks 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.
- *Data center protection*—Ensuring the availability and integrity of centralized applications and systems. Protecting the confidentiality and privacy of users.
- *Network access security and control*—Securing the access edges. Enforcing authentication and role-based access for users residing at the main and remote sites. Ensuring systems are up-to-date and in compliance with the medium enterprises' network security policies.
- *Network endpoint protection*—Protecting servers and enterprise-controlled systems (desktops, laptops, etc.) from viruses, malware, botnets, and other malicious software. Enforcing E-mail and web browsing policies for enterprise users.

Each of these security design considerations are discussed in further detail in [Chapter 5, “Medium Enterprise Design Profile \(MEDP\)—Network Security Design.”](#)

## Mobility

Mobility is an essential part of the enterprise environment. Most users will connect wirelessly to site networks. Additionally, 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 enterprise users and guests to be accessible and productive, regardless of whether they are meeting in a conference room, at lunch with colleagues in the site cafeteria, or simply enjoying a breath of fresh air outside a site building. Provide easy, secure guest access to guests such as temporary workers, visiting colleagues, 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. Enterprise users 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*—Enterprise network administrators must be able to easily deploy, operate, and manage hundreds of access points within multiple enterprise site deployments. A single, easy to understand WLAN management framework is desired to provide small, medium and large sites within the enterprise with the same level of wireless LAN management scalability, reliability and ease of deployment that is demanded by very large 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 users and visitors in the event of common failures.

## Unified Communications

### Call Processing Considerations

How calls are processed in the medium enterprise environment is an important design consideration, guidance on designing scalable and resilient call processing systems is essential for deploying a unified communications system. Some of the 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). Several considerations for gateways include the following:

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

## Dial Plan Considerations

The dial plan is one of the key elements of an 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 for this reason the service fabric must take survivability into consideration.







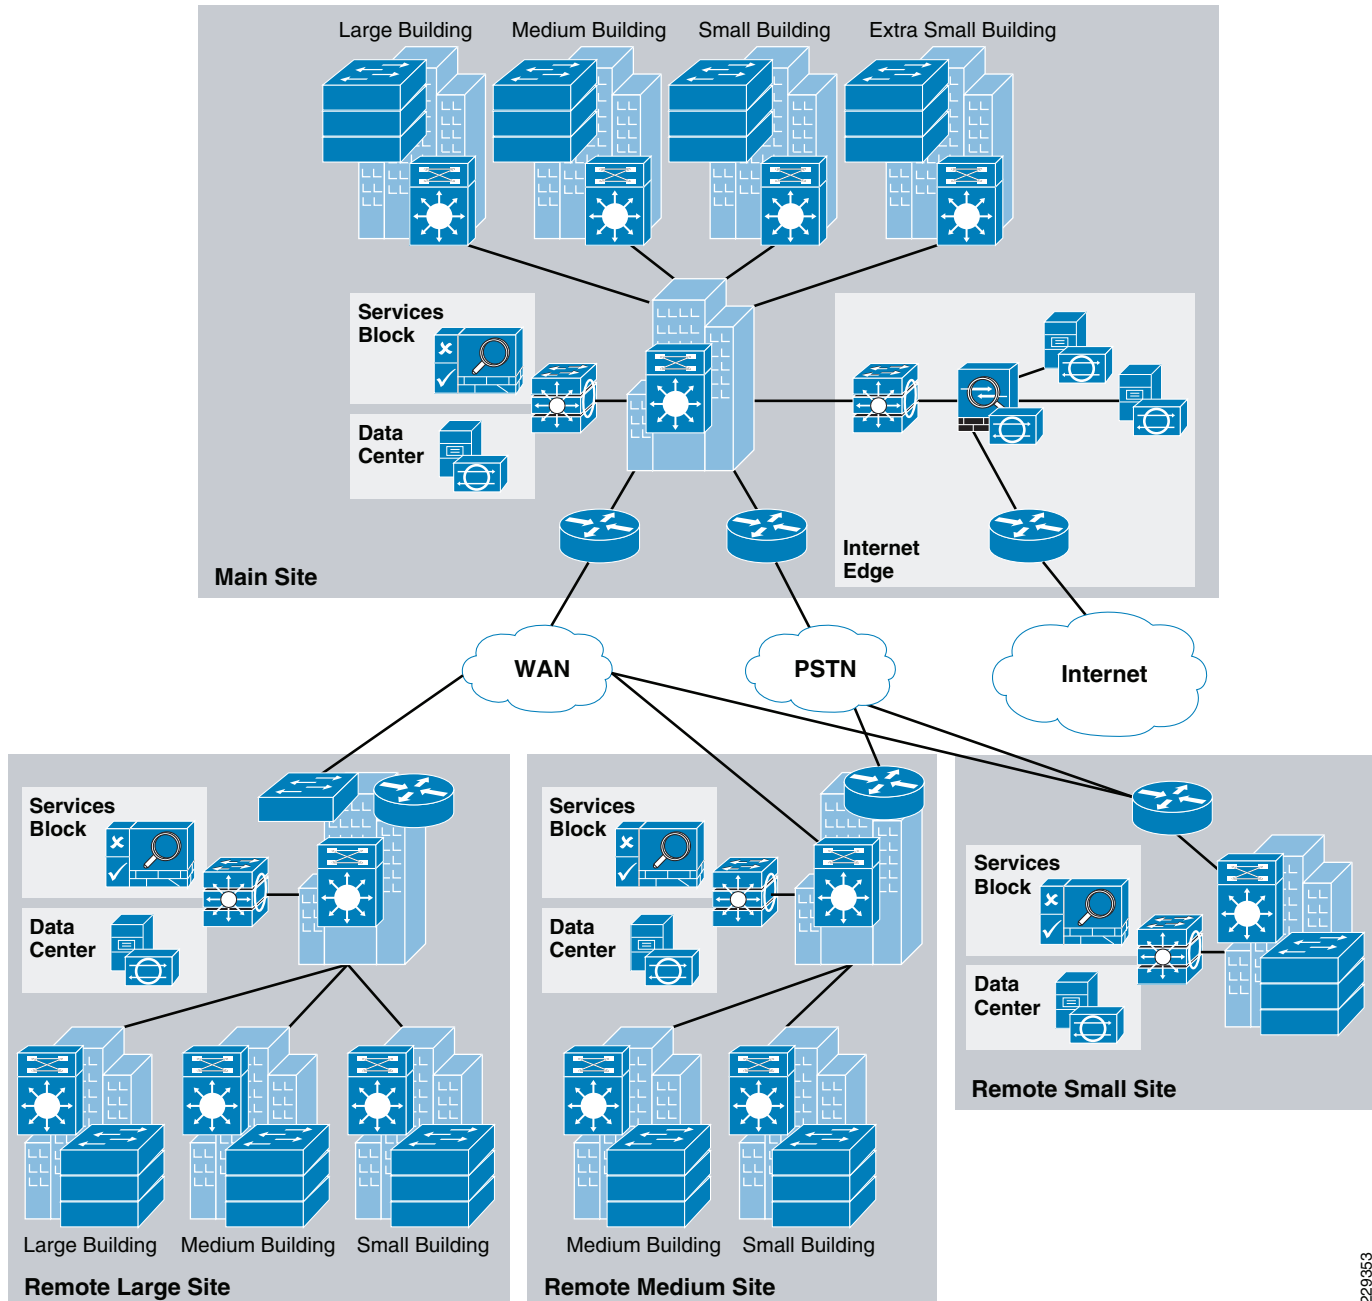
## CHAPTER 2

# Medium Enterprise Design Profile (MEDP)—LAN Design

---


## LAN Design

The Medium Enterprise LAN design is a multi-campus design, where a campus consists of multiple buildings and services at each location, as shown in [Figure 2-1](#).

**Figure 2-1** Medium Enterprise LAN Design

229353

Figure 2-2 shows the service fabric design model used in the medium enterprise LAN design.

**Figure 2-2 Medium Enterprise LAN Design**

This chapter focuses on the LAN component of the overall design. The LAN component consists of the LAN framework and network foundation technologies that provide baseline routing and switching guidelines. The LAN design interconnects several other components, such as endpoints, data center, WAN, and so on, to provide a foundation on which mobility, security, and unified communications (UC) can be integrated into the overall design.

This LAN design provides guidance on building the next-generation medium enterprise network, which becomes a common framework along with critical network technologies to deliver the foundation for the service fabric design. This chapter is divided into following sections:

- *LAN design principles*—Provides proven design choices to build various types of LANs.
- *LAN design model for the medium enterprise*—Leverages the design principles of the tiered network design to facilitate a geographically dispersed enterprise campus network made up of various elements, including networking role, size, capacity, and infrastructure demands.
- *Considerations of a multi-tier LAN design model for medium enterprises*—Provides guidance for the enterprise campus LAN network as a platform with a wide range of next-generation products and technologies to integrate applications and solutions seamlessly.
- *Designing network foundation services for LAN designs in medium enterprise*—Provides guidance on deploying various types of Cisco IOS technologies to build a simplified and highly available network design to provide continuous network operation. This section also provides guidance on designing network-differentiated services that can be used to customize the allocation of network resources to improve user experience and application performance, and to protect the network against unmanaged devices and applications.

# LAN Design Principles


Any successful design or system is based on a foundation of solid design theory and principles. Designing the LAN component of the overall medium enterprise LAN service fabric design model is no different than designing any large networking system. The use of a guiding set of fundamental engineering design principles serves to ensure that the LAN design provides for the balance of availability, security, flexibility, and manageability required to meet current and future advanced and emerging technology needs. This chapter provides design guidelines that are built upon the following principles to allow a medium enterprise network architect to build enterprise campuses that are located in different geographical locations:

- *Hierarchical*
  - Facilitates understanding the role of each device at every tier
  - Simplifies deployment, operation, and management
  - Reduces fault domains at every tier
- *Modularity*—Allows the network to grow on an on-demand basis
- *Resiliency*—Satisfies user expectations for keeping network always on
- *Flexibility*—Allows intelligent traffic load sharing by using all network resources

These are not independent principles. The successful design and implementation of a campus network requires an understanding of how each of these principles applies to the overall design. In addition, understanding how each principle fits in the context of the others is critical in delivering a hierarchical, modular, resilient, and flexible network required by medium enterprises today.

Designing the medium enterprise LAN building blocks in a hierarchical fashion creates a flexible and resilient network foundation that allows network architects to overlay the security, mobility, and UC features essential to the service fabric design model, as well as providing an interconnect point for the WAN aspect of the network. The two proven, time-tested hierarchical design frameworks for LAN networks are the three-tier layer and the two-tier layer models, as shown in [Figure 2-3](#).

**Figure 2-3** Three-Tier and Two-Tier LAN Design Models



228470

The key layers are access, distribution and core. Each layer can be seen as a well-defined structured module with specific roles and functions in the LAN network. Introducing modularity in the LAN hierarchical design further ensures that the LAN network remains resilient and flexible to provide critical network services as well as to allow for growth and changes that may occur in a medium enterprise.

- *Access layer*

The access layer represents the network edge, where traffic enters or exits the campus network. Traditionally, the primary function of an access layer switch is to provide network access to the user. Access layer switches connect to the distribution layer switches to perform network foundation technologies such as routing, quality of service (QoS), and security.

To meet network application and end-user demands, the next-generation Cisco Catalyst switching platforms no longer simply switch packets, but now provide intelligent services to various types of endpoints at the network edge. Building intelligence into access layer switches allows them to operate more efficiently, optimally, and securely.

- *Distribution layer*

The distribution layer interfaces between the access layer and the core layer to provide many key functions, such as the following:

- Aggregating and terminating Layer 2 broadcast domains
- Aggregating Layer 3 routing boundaries
- Providing intelligent switching, routing, and network access policy functions to access the rest of the network
- Providing high availability through redundant distribution layer switches to the end-user and equal cost paths to the core, as well as providing differentiated services to various classes of service applications at the edge of network


- *Core layer*

The core layer is the network backbone that connects all the layers of the LAN design, providing for connectivity between end devices, computing and data storage services located within the data center and other areas, and services within the network. The core layer serves as the aggregator for all the other campus blocks, and ties the campus together with the rest of the network.

**Note**

For more information on each of these layers, see the enterprise class network framework at the following URL: <http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html>.

Figure 2-4 shows a sample three-tier LAN network design for medium enterprises where the access, distribution, and core are all separate layers. To build a simplified, cost-effective, and efficient physical cable layout design, Cisco recommends building an extended-star physical network topology from a centralized building location to all other buildings on the same campus.

**Figure 2-4 Three-Tier LAN Network Design Example**


229354

The primary purpose of the core layer is to provide fault isolation and backbone connectivity. Isolating the distribution and core into separate layers creates a clean delineation for change control between activities affecting end stations (laptops, phones, and printers) and those that affect the data center, WAN, or other parts of the network. A core layer also provides for flexibility in adapting the campus design to meet physical cabling and geographical challenges. If necessary, a separate core layer can use a different transport technology, routing protocols, or switching hardware than the rest of the campus, providing for more flexible design options when needed.

In some cases, because of either physical or network scalability, having separate distribution and core layers is not required. In smaller locations where there are less users accessing the network or in campus sites consisting of a single building, separate core and distribution layers are not needed. In this scenario, Cisco recommends the two-tier LAN network design, also known as the collapsed core network design.

Figure 2-5 shows a two-tier LAN network design example for a medium enterprise LAN where the distribution and core layers are collapsed into a single layer.



**Figure 2-5 Two-Tier Network Design Example**

If using the small-scale collapsed campus core design, the enterprise network architect must understand the network and application demands so that this design ensures a hierarchical, modular, resilient, and flexible LAN network.

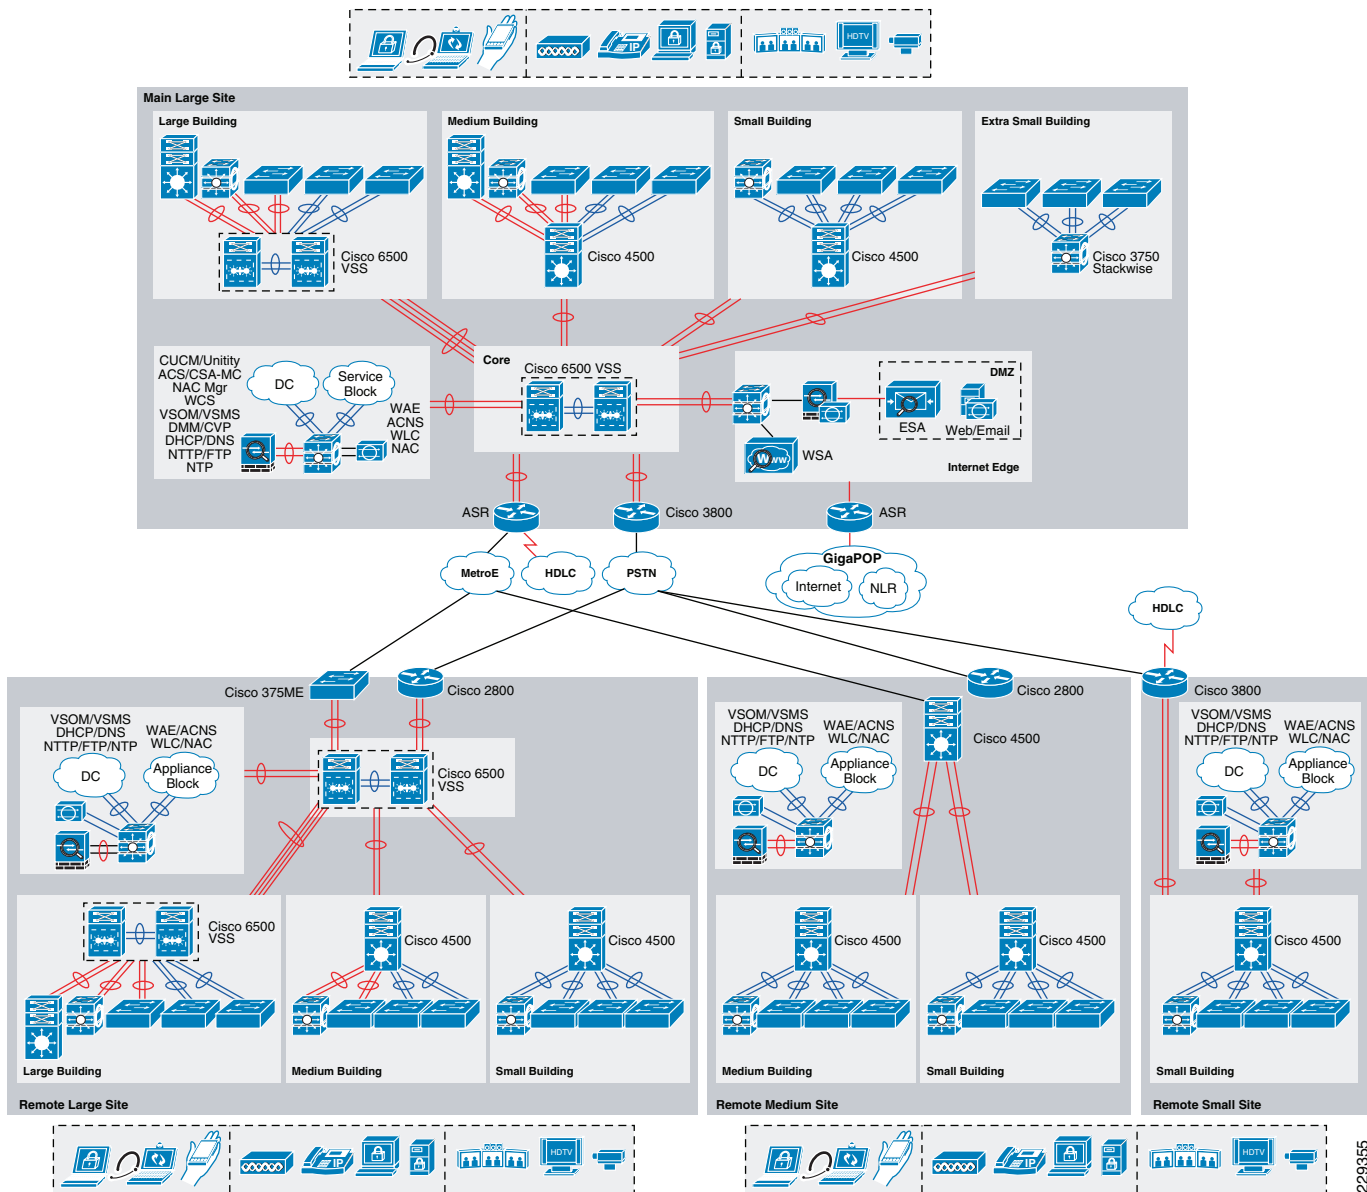
## Medium Enterprise LAN Design Models

Both LAN design models (three-tier and two-tier) have been developed with the following considerations:

- *Scalability*—Based on Cisco enterprise-class high-speed 10G core switching platforms for seamless integration of next-generation applications required for medium enterprises. Platforms chosen are cost-effective and provide investment protection to upgrade network as demand increases.
- *Simplicity*—Reduced operational and troubleshooting cost via the use of network-wide configuration, operation, and management.
- *Resilient*—Sub-second network recovery during abnormal network failures or even network upgrades.
- *Cost-effectiveness*—Integrated specific network components that fit budgets without compromising performance.

As shown in [Figure 2-6](#), multiple campuses can co-exist within a single medium enterprise system that offers various academic programs.

Figure 2-6 Medium Enterprise LAN Design Model



Depending on the remote campus office facility, the number of employees and the networked devices in remote campuses may be equal to or less than the main site. Campus network designs for the remote campus may require adjusting based on overall campus capacity.

Using high-speed WAN technology, all the remote medium enterprise campuses interconnect to a centralized main site that provides shared services to all the employees independent of their physical location. The WAN design is discussed in greater detail in the next chapter, but it is worth mentioning in the LAN section because some remote sites may integrate LAN and WAN functionality into a single platform. Collapsing the LAN and WAN functionality into a single Cisco platform can provide all the needed requirements for a particular remote site as well as provide reduced cost to the overall design, as discussed in more detail in the following section.

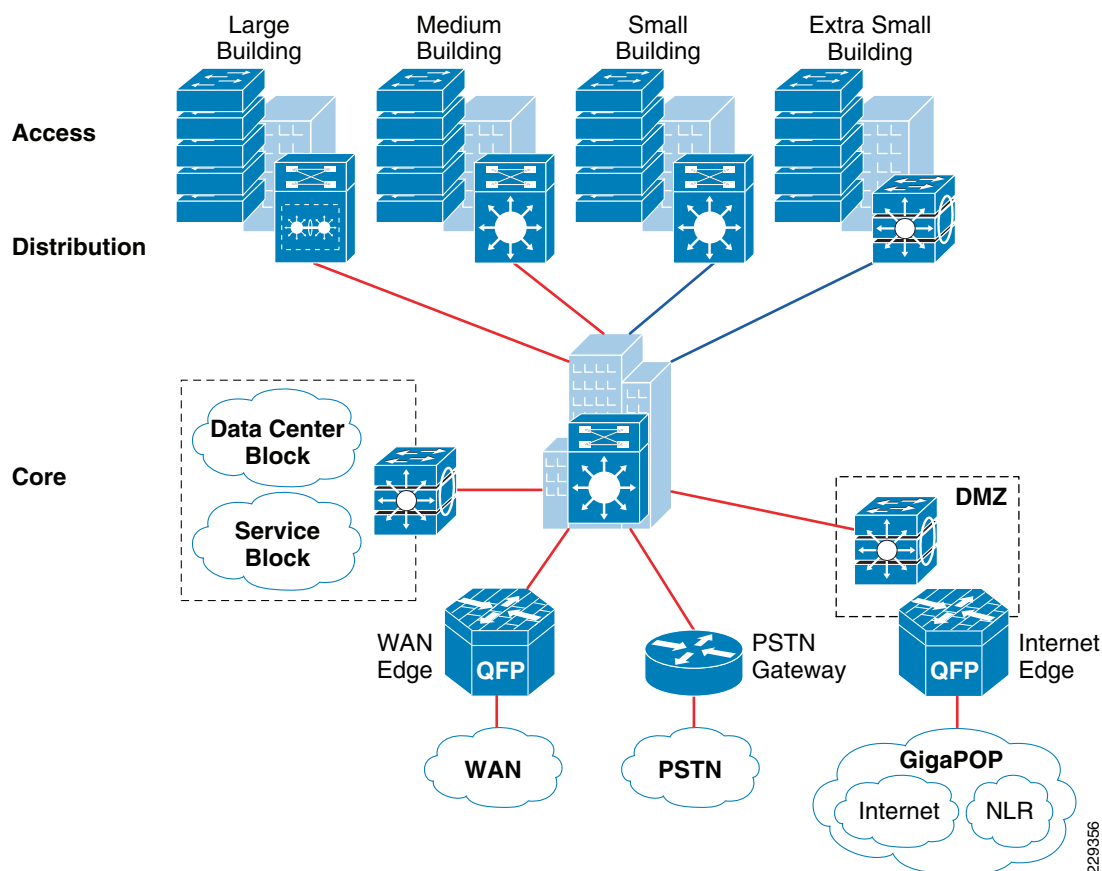
Table 2-1 shows a summary of the LAN design models as they are applied in the overall medium enterprise network design.

**Table 2-1** Medium Enterprise Recommended LAN Design Model

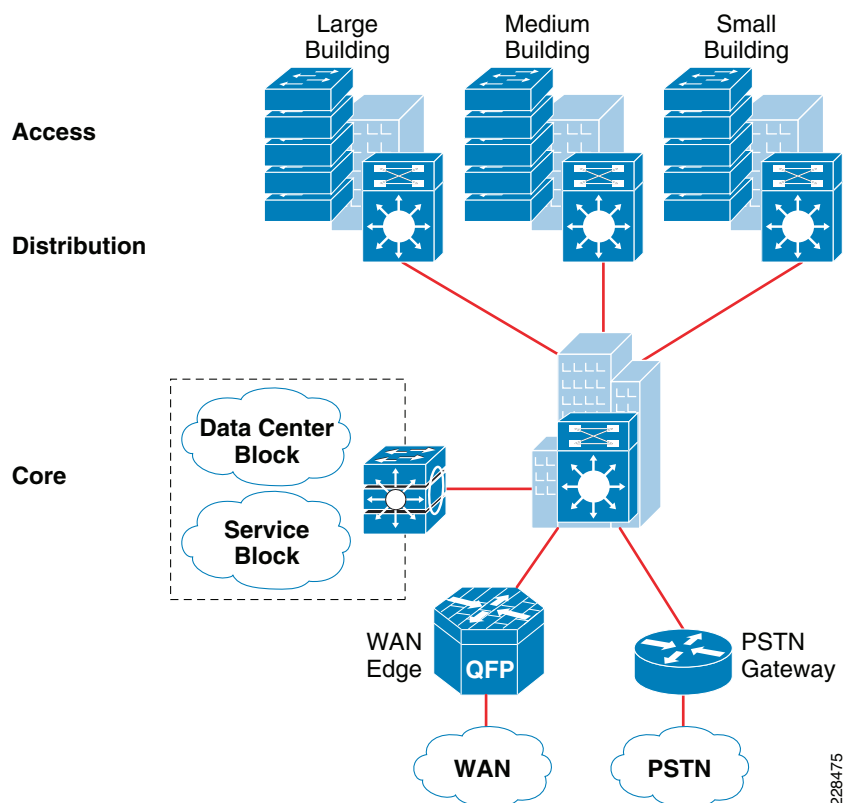
Medium Enterprise Location	Recommended LAN Design Model
Main campus	Three-tier
Remote large campus	Three-tier
Remote medium campus	Three-tier with collapsed WAN edge
Remote small campus	Two-tier

## Main Site Network Design

The main site in the medium enterprise design consists of a centralized hub campus location that interconnects several sizes of remote campuses to provide end-to-end shared network access and services, as shown in [Figure 2-7](#).

**Figure 2-7** Main Site Reference Design

The main site typically consists of various sizes of building facilities and various organization department groups. The network scale factor in the main site is higher than the remote campus site, and includes end users, IP-enabled endpoints, servers, and security and network edge devices. Multiple buildings of various sizes exist in one location, as shown in [Figure 2-8](#).

**Figure 2-8 Main Site Reference Design**

The three-tier LAN design model for the main site meets all key technical aspects to provide a well-structured and strong network foundation. The modularity and flexibility in a three-tier LAN design model allows easier expansion and integration in the main site network, and keeps all network elements protected and available.

To enforce external network access policy for each end user, the three-tier model also provides external gateway services to the employees for accessing the Internet.

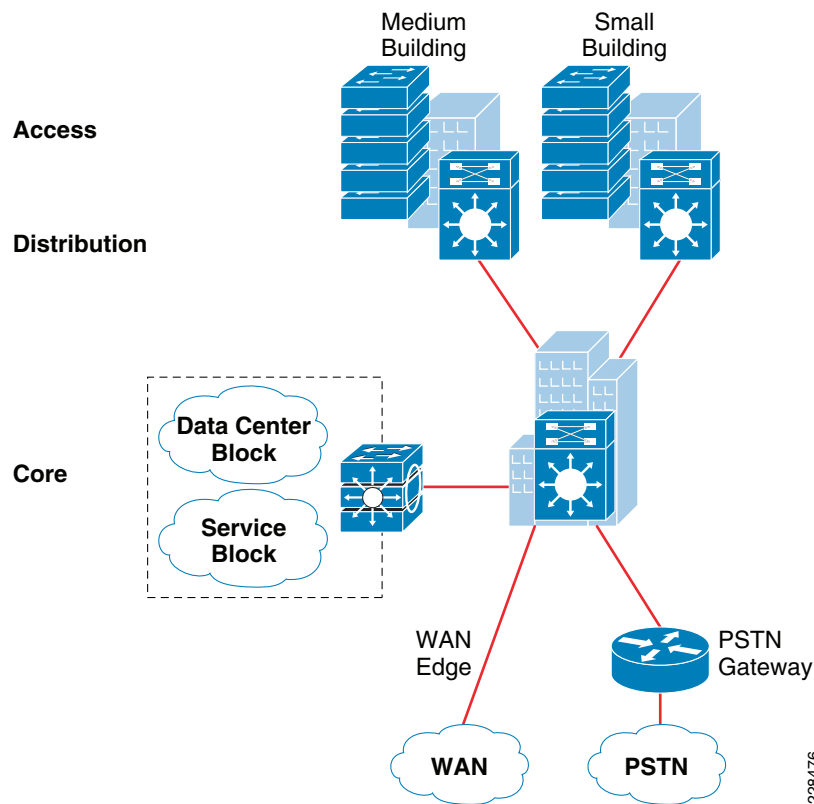
**Note**

The WAN design is a separate element in this location, because it requires a separate WAN device that connects to the three-tier LAN model. WAN design is discussed in more detail in [Chapter 3, “Medium Enterprise Design Profile \(MEDP\)—WAN Design.”](#)

## Remote Large Campus Site Design


From the location size and network scale perspective, the remote large site is not much different from the main site. Geographically, it can be distant from the main campus site and requires a high-speed WAN circuit to interconnect both campuses. The remote large site can also be considered as an alternate campus to the main campus site, with the same common types of applications, endpoints, users, and network services. Similar to the main site, separate WAN devices are recommended to provide application delivery and access to the main site, given the size and number of employees at this location.

Similar to the main site, Cisco recommends the three-tier LAN design model for the remote large site campus, as shown in [Figure 2-9](#).

**Figure 2-9 Remote Large Campus Site Reference Design**

## Remote Medium Campus Site Design

Remote medium campus locations differ from a main or remote large site campus in that there are less buildings with distributed organization departments. A remote medium campus may have a fewer number of network users and endpoints, thereby reducing the need to build a similar campus network to that recommended for main and large campuses. Because there are fewer employees and networked devices at this site as compared to the main or remote large site campus sites, the need for a separate WAN device may not be necessary. A remote medium campus network is designed similarly to a three-tier large campus LAN design. All the LAN benefits are achieved in a three-tier design model as in the main and remote large site campus, and in addition, the platform chosen in the core layer also serves as the WAN edge, thus collapsing the WAN and core LAN functionality into a single platform. [Figure 2-10](#) shows the remote medium campus in more detail.

**Figure 2-10 Remote Medium Campus Site Reference Design**


## Remote Small Campus Network Design

The remote small campus is typically confined to a single building that spans across multiple floors with different academic departments. The network scale factor in this design is reduced compared to other large campuses. However, the application and services demands are still consistent across the medium enterprise locations.

In such smaller scale campus network deployments, the distribution and core layer functions can collapse into the two-tier LAN model without compromising basic network demands. Before deploying a collapsed core and distribution layer in the remote small campus network, considering all the scale and expansion factors prevents physical network re-design, and improves overall network efficiency and manageability.

WAN bandwidth requirements must be assessed appropriately for this remote small campus network design. Although the network scale factor is reduced compared to other larger campus locations, sufficient WAN link capacity is needed to deliver consistent network services to employees. Similar to the remote medium campus location, the WAN functionality is also collapsed into the LAN functionality. A single Cisco platform can provide collapsed core and distribution LAN layers. This design model is recommended only in smaller locations, and WAN traffic and application needs must be considered.

[Figure 2-11](#) shows the remote small campus in more detail.

**Figure 2-11 Remote Small Campus Site Reference Design**

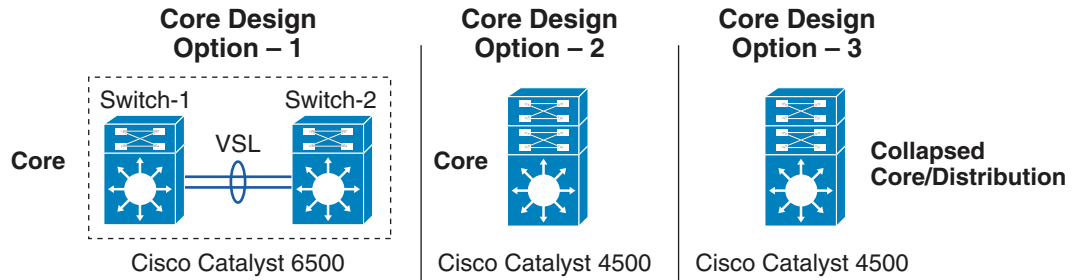
## Multi-Tier LAN Design Models for Medium Enterprise

The previous section discussed the recommended LAN design model for each medium enterprise location. This section provides more detailed design guidance for each tier in the LAN design model. Each design recommendation is optimized to keep the network simplified and cost-effective without compromising network scalability, security, and resiliency. Each LAN design model for a medium enterprise location is based on the key LAN layers of core, distribution, and access.

### Campus Core Layer Network Design

As discussed in the previous section, the core layer becomes a high-speed intermediate transit point between distribution blocks in different premises and other devices that interconnect to the data center, WAN, and Internet edge.

Similarly to choosing a LAN design model based on a location within the medium enterprise design, choosing a core layer design also depends on the size and location within the design. Three core layer design models are available, each of which is based on either the Cisco Catalyst 6500-E Series or the Cisco Catalyst 4500-E Series Switches. [Figure 2-12](#) shows the three core layer design models.


**Figure 2-12 Core Layer Design Models for Medium Enterprises**

Each design model offers consistent network services, high availability, expansion flexibility, and network scalability. The following sections provide detailed design and deployment guidance for each model as well as where they fit within the various locations of the medium enterprise design.

## Core Layer Design Option 1—Cisco Catalyst 6500-E-Based Core Network

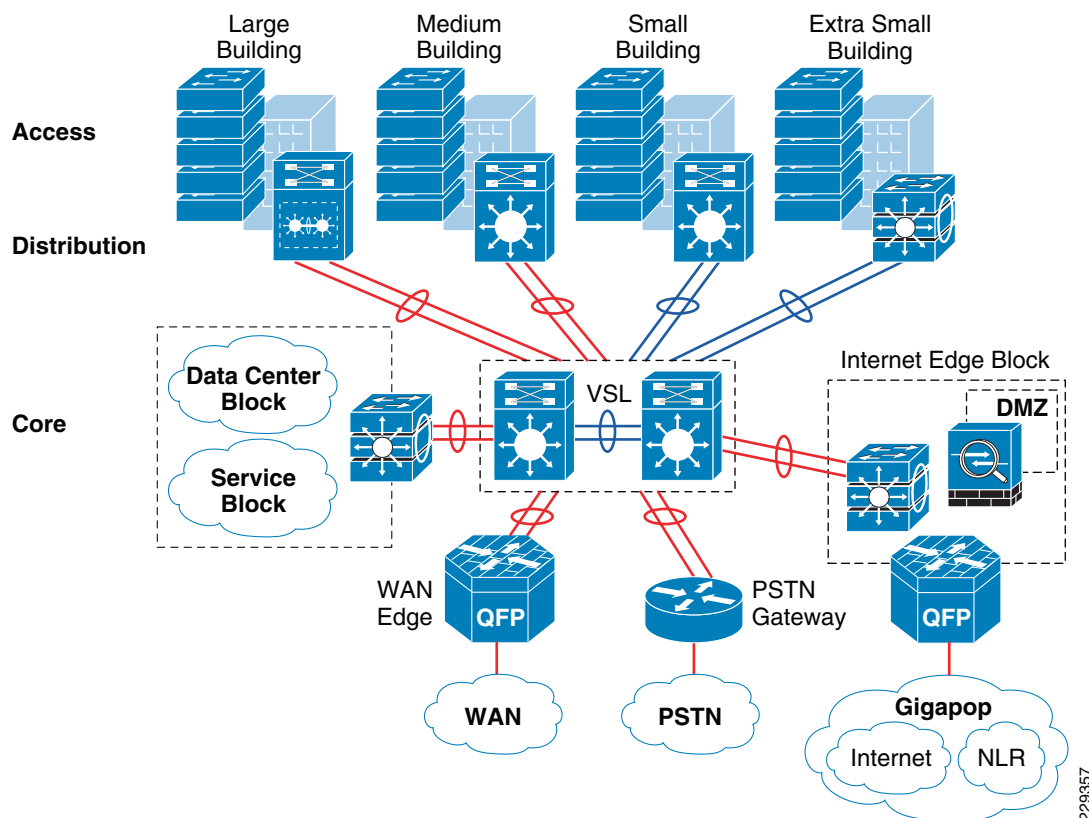
Core layer design option 1 is specifically intended for the main and remote large site campus locations. It is assumed that the number of network users, high-speed and low-latency applications (such as Cisco TelePresence), and the overall network scale capacity is common in both sites and thus, similar core design principles are required.

Core layer design option 1 is based on Cisco Catalyst 6500 Series switches using the Cisco Virtual Switching System (VSS), which is a software technology that builds a single logical core system by clustering two redundant core systems in the same tier. Building a VSS-based network changes network design, operation, cost, and management dramatically. [Figure 2-13](#) shows the physical and operational view of VSS.

**Figure 2-13 VSS Physical and Operational View**

To provide end-to-end network access, the core layer interconnects several other network systems that are implemented in different roles and service blocks. Using VSS to virtualize the core layer into a single logical system remains transparent to each network device that interconnects to the VSS-enabled core. The single logical connection between core and the peer network devices builds a reliable, point-to-point connection that develops a simplified network topology and builds distributed forwarding tables to fully use all resources. [Figure 2-14](#) shows a reference VSS-enabled core network design for the main campus site.



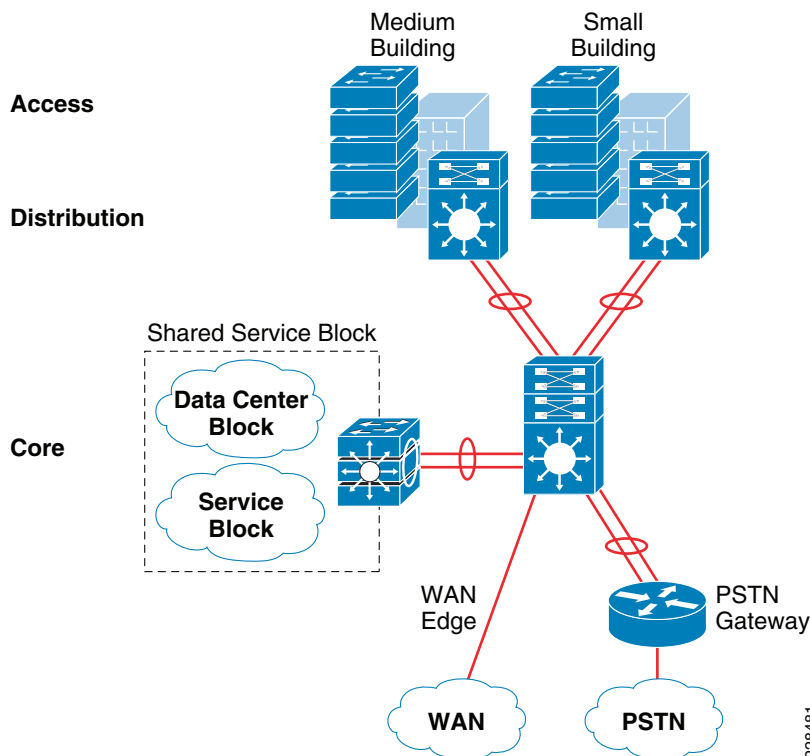
**Figure 2-14 VSS-Enabled Core Network Design****Note**

For more detailed VSS design guidance, see the *Campus 3.0 Virtual Switching System Design Guide* at the following URL:


[http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVSS\\_DG.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVSS_DG.html).

## Core Layer Design Option 2—Cisco Catalyst 4500-E-Based Campus Core Network

Core layer design option 2 is intended for a remote medium-sized campus and is built on the same principles as for the main and remote large site campus locations. The size of this remote site may not be large, and it is assumed that this location contains distributed building premises within the remote medium campus design. Because this site is smaller in comparison to the main and remote large site campus locations, a fully redundant, VSS-based core layer design may not be necessary. Therefore, core layer design option 2 was developed to provide a cost-effective alternative while providing the same functionality as core layer design option 1. [Figure 2-15](#) shows the remote medium campus core design option in more detail.

**Figure 2-15 Remote Medium Campus Core Network Design**

The cost of implementing and managing redundant systems in each tier may introduce complications in selecting the three-tier model, especially when network scale factor is not too high. This cost-effective core network design provides protection against various types of hardware and software failure and offers sub-second network recovery. Instead of a redundant node in the same tier, a single Cisco Catalyst 4500-E Series Switch can be deployed in the core role and bundled with 1+1 redundant in-chassis network components. The Cisco Catalyst 4500-E Series modular platform is a one-size platform that helps enable the high-speed core backbone to provide uninterrupted network access within a single chassis. Although a fully redundant, two-chassis design using VSS as described in core layer option 1 provides the greatest redundancy for large-scale locations, the redundant supervisors and line cards of the Cisco Catalyst 4500-E provide adequate redundancy for smaller locations within a single platform. Figure 2-16 shows the redundancy of the Cisco Catalyst 4500-E Series in more detail.

**Figure 2-16 Highly Redundant Single Core Design Using the Cisco Catalyst 4500-E Platform**

This core network design builds a network topology that has similar common design principles to the VSS-based campus core in core layer design option 1. The future expansion from a single core to a dual VSS-based core system becomes easier to deploy, and helps retain the original network topology and the management operation. This cost-effective single resilient core system for a medium-size enterprise network meets the following four key goals:

- *Scalability*—The modular Cisco Catalyst 4500-E chassis enables flexibility for core network expansion with high throughput modules and port scalability without compromising network performance.
- *Resiliency*—Because hardware or software failure conditions may create catastrophic results in the network, the single core system must be equipped with redundant system components such as supervisor, line card, and power supplies. Implementing redundant components increases the core network resiliency during various types of failure conditions using Non-Stop Forwarding/Stateful Switch Over (NSF/SSO) and EtherChannel technology.
- *Simplicity*—The core network can be simplified with redundant network modules and diverse fiber connections between the core and other network devices. The Layer 3 network ports must be bundled into a single point-to-point logical EtherChannel to simplify the network, such as the VSS-enabled campus design. An EtherChannel-based campus network offers similar benefits to an Multi-chassis EtherChannel (MEC)- based network.
- *Cost-effectiveness*—A single core system in the core layer helps reduce capital, operational, and management cost for the medium-sized campus network design.

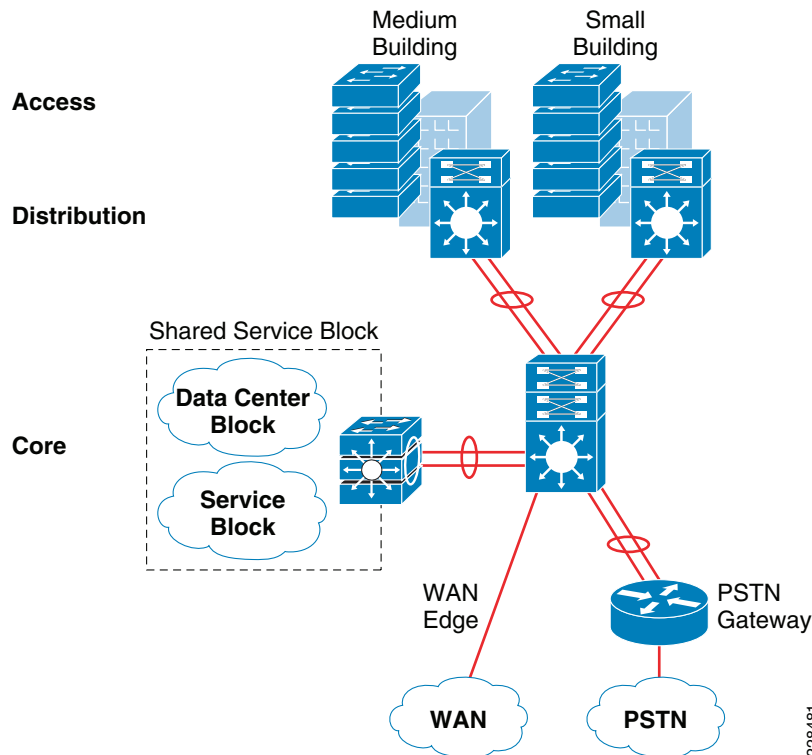
### Core Layer Design Option 3—Cisco Catalyst 4500-E-Based Collapsed Core Campus Network

Core layer design option 3 is intended for the remote small campus network that has consistent network services and applications service-level requirements but at reduced network scale. The remote small campus is considered to be confined within a single multi-story building that may span academic departments across different floors. To provide consistent services and optimal network performance, scalability, resiliency, simplification, and cost-effectiveness in the small campus network design must not be compromised.

As discussed in the previous section, the remote small campus has a two-tier LAN design model, so the role of the core system is merged with the distribution layer. Remote small campus locations have consistent design guidance and best practices defined for main, remote large site, and remote medium-sized campus cores. However, for platform selection, the remote medium campus core layer design must be leveraged to build this two-tier campus core.

Single highly resilient Cisco Catalyst 4500-E switches with a Cisco Sup6L-E supervisor must be deployed in a centralized collapsed core and distribution role that interconnects to wiring closet switches, a shared service block, and a WAN edge router. The cost-effective supervisor version supports key technologies such as robust QoS, high availability, security, and much more at a lower scale, making it an ideal solution for small-scale network designs. [Figure 2-17](#) shows the remote small campus core design in more detail.

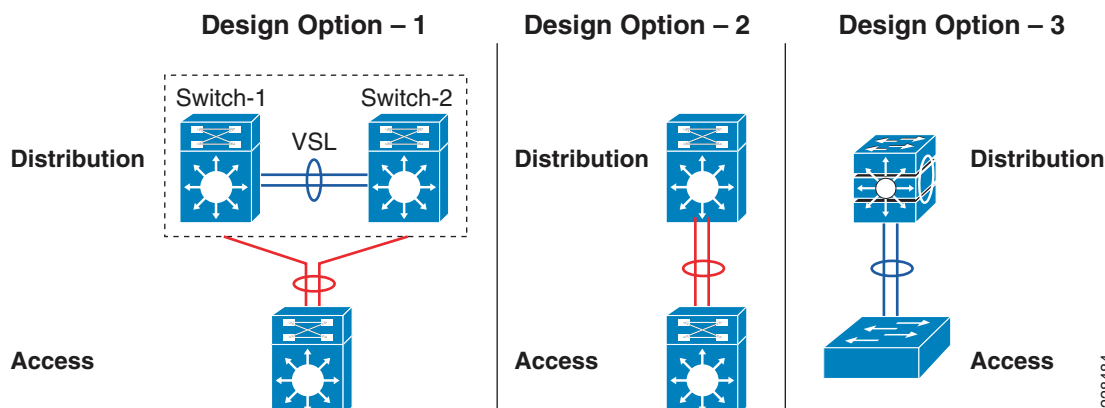
**Figure 2-17** Core Layer Option 3 Collapsed Core/Distribution Network Design in Remote Small Campus Location



## Campus Distribution Layer Network Design


The distribution or aggregation layer is the network demarcation boundary between wiring-closet switches and the campus core network. The framework of the distribution layer system in the medium enterprise design is based on best practices that reduce network complexities and accelerate reliability and performance. To build a strong campus network foundation with the three-tier model, the distribution layer has a vital role in consolidating networks and enforcing network edge policies.

Following the core layer design options in different campus locations, the distribution layer design provides consistent network operation and configuration tools to enable various network services. Three simplified distribution layer design options can be deployed in main or remote campus locations, depending on network scale, application demands, and cost, as shown in [Figure 2-18](#). Each design model offers consistent network services, high availability, expansion flexibility, and network scalability.

**Figure 2-18 Distribution Layer Design Model Options**

## Distribution Layer Design Option 1—Cisco Catalyst 6500-E Based Distribution Network

Distribution layer design option 1 is intended for main campus and remote large site campus locations, and is based on Cisco Catalyst 6500-E Series switches using the Cisco VSS, as shown in [Figure 2-19](#).

**Figure 2-19 VSS-Enabled Distribution Layer Network Design**

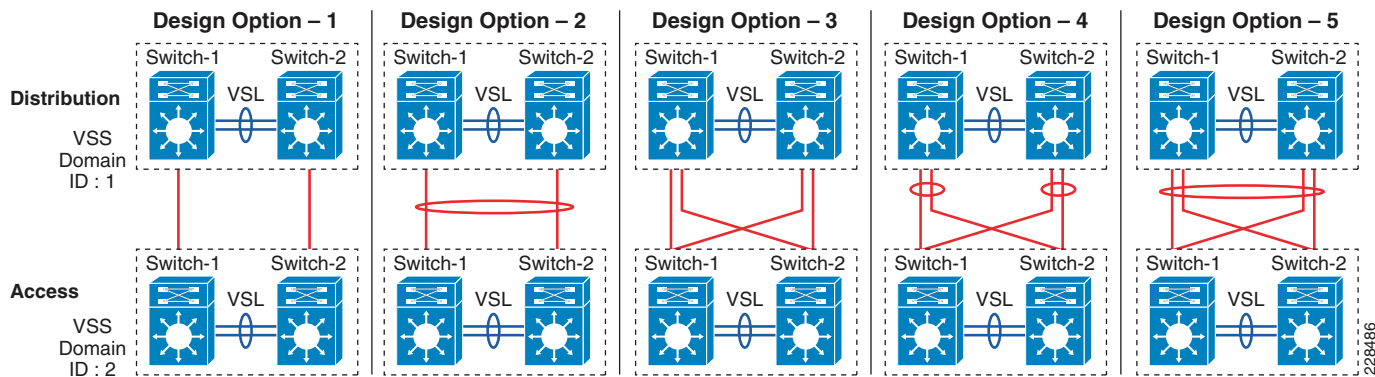
The distribution block and core network operation changes significantly when redundant Cisco Catalyst 6500-E Series switches are deployed in VSS mode in both the distribution and core layers. Clustering redundant distribution switches into a single logical system with VSS introduces the following technical benefits:

- A single logical system reduces operational, maintenance, and ownership cost.
- A single logical IP gateway develops a unified point-to-point network topology in the distribution block, which eliminates traditional protocol limitations and enables the network to operate at full capacity.
- Implementing the distribution layer in VSS mode eliminates or reduces several deployment barriers, such as spanning-tree loop, Hot Standby Routing Protocol (HSRP)/Gateway Load Balancing Protocol (GLBP)/Virtual Router Redundancy Protocol (VRRP), and control plane overhead.

- Cisco VSS introduces unique inter-chassis traffic engineering to develop a fully-distributed forwarding design that helps in increased bandwidth, load balancing, predictable network recovery, and network stability.

Deploying VSS mode in both the distribution layer switch and core layer switch provides numerous technology deployment options that are not available when not using VSS. Designing a common core and distribution layer option using VSS provides greater redundancy and is able to handle the amount of traffic typically present in the main and remote large site campus locations. Figure 2-20 shows five unique VSS domain interconnect options. Each variation builds a unique network topology that has a direct impact on steering traffic and network recovery.

**Figure 2-20 Core/Distribution Layer Interconnection Design Considerations**



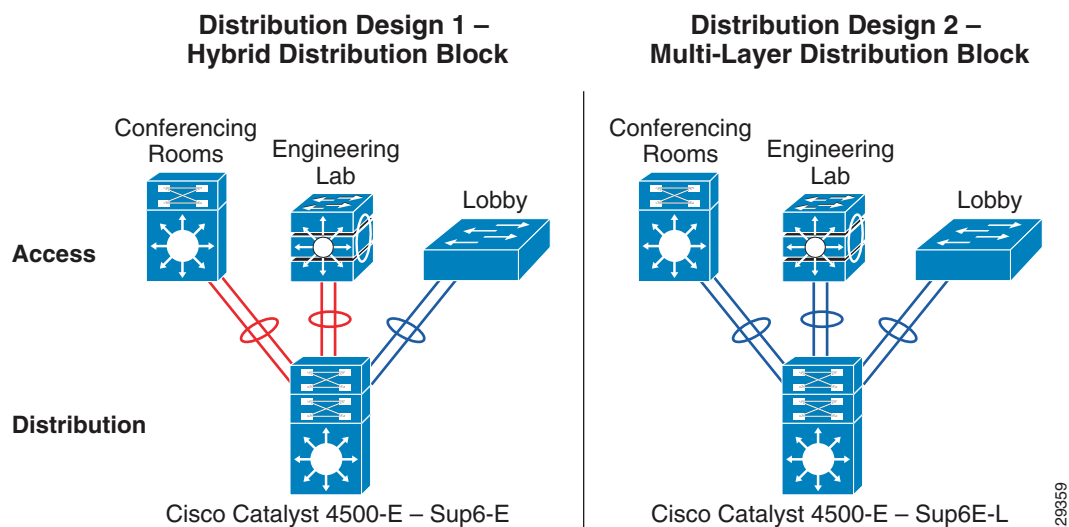
The various core/distribution layer interconnects offer the following:

- *Core/distribution layer interconnection option 1*—A single physical link between each core switch with the corresponding distribution switch.
- *Core/distribution layer interconnection option 2*—A single physical link between each core switch with the corresponding distribution switch, but each link is logically grouped to appear as one single link between the core and distribution layers.
- *Core/distribution layer interconnection option 3*—Two physical links between each core switch with the corresponding distribution switch. This design creates four equal cost multi-path (ECMP) with multiple control plane adjacency and redundant path information. Multiple links provide greater redundancy in case of link failover.
- *Core/distribution layer interconnection option 4*—Two physical links between each core switch with the corresponding distribution switch. There is one link direction between each switch as well as one link connecting to the other distribution switch. The additional link provides greater redundancy in case of link failover. Also these links are logically grouped to appear like option 1 but with greater redundancy.
- *Core/distribution layer interconnection option 5*—This provides the most redundancy between the VSS-enabled core and distribution switches as well as the most simplified configuration, because it appears as if there is only one logical link between the core and the distribution. Cisco recommends deploying this option because it provides higher redundancy and simplicity compared to any other deployment option.

## Distribution Layer Design Option 2—Cisco Catalyst 4500-E-Based Distribution Network

Two cost-effective distribution layer models have been designed for the medium-sized and small-sized buildings within each campus location that interconnect to the centralized core layer design option and distributed wiring closet access layer switches. Both models are based on a common physical LAN network infrastructure and can be chosen based on overall network capacity and distribution block design. Both distribution layer design options use a cost-effective single and highly resilient Cisco Catalyst 4500-E as an aggregation layer system that offers consistent network operation like a VSS-enabled distribution layer switch. The Cisco Catalyst 4500-E Series provides the same technical benefits of VSS for a smaller network capacity within a single Cisco platform. The two Cisco Catalyst 4500-E-based distribution layer options are shown in [Figure 2-21](#).


**Figure 2-21** Two Cisco Catalyst 4500-E-Based Distribution Layer Options



The hybrid distribution block must be deployed with the next-generation supervisor Sup6-E module. Implementing redundant Sup6-Es in the distribution layer can interconnect access layer switches and core layer switches using a single point-to-point logical connection. This cost-effective and resilient distribution design option leverages core layer design option 2 to take advantage of all the operational consistency and architectural benefits.

Alternatively, the multilayer distribution block option requires the Cisco Catalyst 4500-E Series Switch with next-generation supervisor Sup6L-E deployed. The Sup6L-E supervisor is a cost-effective distribution layer solution that meets all network foundation requirements and can operate at moderate capacity, which can handle a medium-sized enterprise distribution block.

This distribution layer network design provides protection against various types of hardware and software failure, and can deliver consistent sub-second network recovery. A single Catalyst 4500-E with multiple redundant system components can be deployed to offer 1+1 in-chassis redundancy, as shown in [Figure 2-22](#).

**Figure 2-22 Highly Redundant Single Distribution Design**


Distribution layer design option 2 is intended for the remote medium-sized campus locations, and is based on the Cisco Catalyst 4500-E Series switches. Although the remote medium and the main and remote large site campus locations share similar design principles, the remote medium campus location is smaller and may not need a VSS-based redundant design. Fortunately, network upgrades and expansion become easier to deploy using distribution layer option 2, which helps retain the original network topology and the management operation. Distribution layer design option 2 meets the following goals:

- *Scalability*—The modular Cisco Catalyst 4500-E chassis provides the flexibility for distribution block expansion with high throughput modules and port scalability without compromising network performance.
- *Resiliency*—The single distribution system must be equipped with redundant system components, such as supervisor, line card, and power supplies. Implementing redundant components increases network resiliency during various types of failure conditions using NSF/SSO and EtherChannel technology.
- *Simplicity*—This cost-effective design simplifies the distribution block similarly to a VSS-enabled distribution system. The single IP gateway design develops a unified point-to-point network topology in the distribution block to eliminate traditional protocol limitations, enabling the network to operate at full capacity.
- *Cost-effectiveness*—The single distribution system in the core layer helps reduce capital, operational, and ownership cost for the medium-sized campus network design.

### Distribution Layer Design Option 3—Cisco Catalyst 3750-X StackWise-Based Distribution Network

Distribution layer design option 3 is intended for a very small building with a limited number of wiring closet switches in the access layer that connects remote classrooms or and office network with a centralized core, as shown in [Figure 2-23](#).



**Figure 2-23 Cisco StackWise Plus-enabled Distribution Layer Network Design**

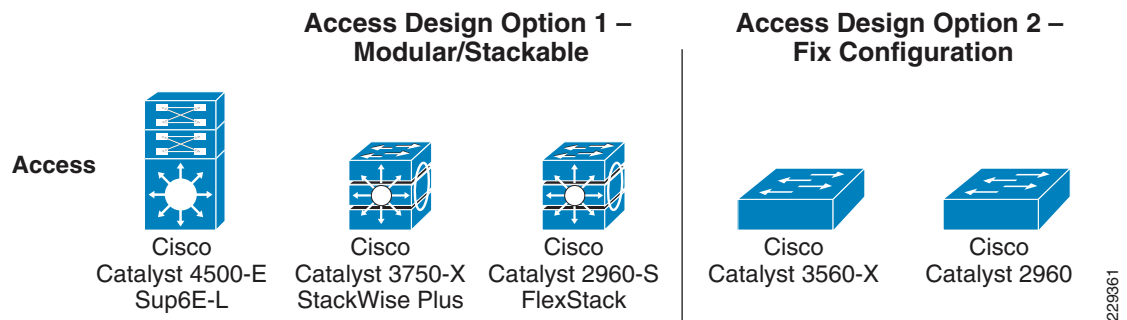
While providing consistent network services throughout the campus, a number of network users and IT-managed remote endpoints can be limited in this building. This distribution layer design option recommends using the Cisco Catalyst 3750-X StackWise Plus Series platform for the distribution layer switch.

The fixed-configuration Cisco Catalyst 3750-X Series switch is a multilayer platform that supports Cisco StackWise Plus technology to simplify the network and offers flexibility to expand the network as it grows. With Cisco StackWise Plus technology, multiple Catalyst 3750-X can be stacked into a high-speed backplane stack ring to logically build as a single large distribution system. Cisco StackWise Plus supports up to nine switches into single stack ring for incremental network upgrades, and increases effective throughput capacity up to 64 Gbps. The chassis redundancy is achieved via stacking, in which member chassis replicate the control functions with each member providing distributed packet forwarding. This is achieved by stacked group members acting as a single virtual Catalyst 3750-X switch. The logical switch is represented as one switch by having one stack member act as the master switch. Thus, when failover occurs, any member of the stack can take over as a master and continue the same services. It is a 1:N form of redundancy where any member can become the master. This distribution layer design option is ideal for the remote small campus location.

## Campus Access Layer Network Design

The access layer is the first tier or edge of the campus, where end devices such as PCs, printers, cameras, Cisco TelePresence, and so on attach to the wired portion of the campus network. It is also the place where devices that extend the network out one more level, such as IP phones and wireless access points (APs), are attached. The wide variety of possible types of devices that can connect and the various services and dynamic configuration mechanisms that are necessary, make the access layer one of the most feature-rich parts of the campus network. Not only does the access layer switch allow users to access the network, the access layer switch must provide network protection so that unauthorized users or applications do not enter the network. The challenge for the network architect is determining how to implement a design that meets this wide variety of requirements, the need for various levels of mobility, the need for a cost-effective and flexible operations environment, while being able to provide the appropriate balance of security and availability expected in more traditional, fixed-configuration environments. The next-generation Cisco Catalyst switching portfolio includes a wide range of fixed and modular switching platforms, each designed with unique hardware and software capability to function in a specific role.

Enterprise campuses may deploy a wide range of network endpoints. The campus network infrastructure resources operate in shared service mode, and include IT-managed devices such as Cisco TelePresence and non-IT-managed devices such as employee laptops. Based on several endpoint factors such as function and network demands and capabilities, two access layer design options can be deployed with campus network edge platforms, as shown in [Figure 2-24](#).

**Figure 2-24 Access Layer Design Models**

## Access Layer Design Option 1—Modular/StackWise Plus/FlexStack Access Layer Network

Access layer design option 1 is intended to address the network scalability and availability for the IT-managed critical voice and video communication network edge devices. To accelerate user experience and campus physical security protection, these devices require low latency, high performance, and a constant network availability switching infrastructure. Implementing a modular, Cisco StackWise Plus and latest Cisco's innovation FlexStack-capable platform provides flexibility to increase network scale in the densely populated campus network edge.

The Cisco Catalyst 4500-E with supervisor Sup6E-L can be deployed to protect devices against access layer network failure. Cisco Catalyst 4500-E Series platforms offer consistent and predictable sub-second network recovery using NSF/SSO technology to minimize the impact of outages on enterprise business and IT operation.

The Cisco Catalyst 3750-X Series is the alternate Cisco switching platform in this design option. Cisco StackWise Plus technology provides flexibility and availability by clustering multiple Cisco Catalyst 3750-X Series Switches into a single high-speed stack ring that simplifies operation and allows incremental access layer network expansion. The Cisco Catalyst 3750-X Series leverages EtherChannel technology for protection during member link or stack member switch failure.

The Catalyst 2960-S with FlexStack technology is Cisco's latest innovation in access-layer tier. Based on StackWise Plus architecture, the FlexStack design is currently supported on Layer-2 Catalyst 2960-S Series switches. Following to 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.

## Access Layer Design Option 2—Fixed Configuration Access Layer Network

This entry-level access layer design option is widely chosen for enterprise environments. The fixed configuration Cisco Catalyst switching portfolio supports a wide range of access layer technologies that allow seamless service integration and enable intelligent network management at the edge.

The next-generation fixed configuration Cisco Catalyst 3560-X and Catalyst 2960 Series is a commonly deployed platform for wired network access that can be in a mixed configuration with critical devices such as Cisco IP Phones and non-mission critical endpoints such as library PCs, printers, and so on. For non-stop network operation during power outages, the Catalyst 3560-X must be deployed with an internal or external redundant power supply solution using the Cisco RPS 2300. Increasing aggregated

power capacity allows flexibility to scale with enhanced power-over-Ethernet (PoE+) on a per-port basis. With its wire-speed 10G uplink forwarding capacity, this design reduces network congestion and latency to significantly improve application performance.

For a campus network, the Cisco Catalyst 3560-X is an alternate switching solution for the multilayer distribution block design option discussed in the previous section. The Cisco Catalyst 3560-X Series Switches offer limited software feature support that can function only in a traditional Layer 2 network design. To provide a consistent end-to-end enhanced user experience, the Cisco Catalyst 2960-S supports critical network control services to secure the network edge, intelligently provide differentiated services to various class-of-service traffic, as well as simplified management. The Cisco Catalyst must leverage the 1G dual uplink ports to interconnect the distribution system for increased bandwidth capacity and network availability.

Both design options offer consistent network services at the campus edge to provide differentiated, intelligent, and secured network access to trusted and untrusted endpoints. The distribution options recommended in the previous section can accommodate both access layer design options.

## Deploying Medium Enterprise Network Foundation Services

After each tier in the model has been designed, the next step for the medium enterprise design is to establish key network foundation services. Regardless of the application function and requirements that medium enterprises demand, the network must be designed to provide a consistent user experience independent of the geographical location of the application. The following network foundation design principles or services must be deployed in each campus location to provide resiliency and availability for all users to obtain and use the applications the medium enterprise offers:

- Implementing LAN network infrastructure
- Network addressing hierarchy
- Network foundation technologies for LAN designs
- Multicast for applications delivery
- QoS for application performance optimization
- High availability to ensure user experience even with a network failure

Design guidance for each of these six network foundation services are discussed in the following sections, including where they are deployed in each tier of the LAN design model, the campus location, and capacity.

### Implementing LAN Network Infrastructure

The preceding sections provided various design options for deploying the Cisco Catalyst platform in multi-tier centralized main campus and remote campus locations. The Medium Enterprise Reference network is designed with consistency to build simplified network topology for easier operation, management, and troubleshooting independent of campus location. Depending on network size, scalability, and reliability requirements, the Medium Enterprise Reference design applies the following common set of Cisco Catalyst platforms in different campus network layers:

- Cisco Catalyst 6500-E in VSS mode
- Cisco Catalyst 4500-E
- Cisco Catalyst 3750-X Stackwise and Catalyst 2960-S FlexStack
- Cisco Catalyst 3560-X and 2960

This subsection focuses on building the initial LAN network infrastructure setup to bring the network up to the stage to start establishing network protocol communication with the peer devices. The deployment and configuration guidelines remain consistent for each recommended Catalyst platform independent of their network role. Advanced network services implementation and deployment guidelines will be explained in subsequent section.

## Deploying Cisco Catalyst 6500-E in VSS Mode

All the VSS design principles and foundational technologies defined in this subsection remains consistent when the Cisco Catalyst 6500-E is deployed in VSS mode at campus core or distribution layer.

Prior to enabling the Cisco Catalyst 6500-E in VSS mode, enterprise network administrator must adhere to Cisco recommended best practices to take complete advantage of virtualized system and minimize the network operation downtime when migration is required in a production network. Migrating VSS from the standalone Catalyst 6500-E system requires multiple pre and post-migration steps to deploy virtual-system that includes building virtual-system itself and migrating the existing standalone network configuration to operate in virtual-system environment. Refer to the following document for step-by-step migration procedure:

[http://www.cisco.com/en/US/products/ps9336/products\\_tech\\_note09186a0080a7c74c.shtml](http://www.cisco.com/en/US/products/ps9336/products_tech_note09186a0080a7c74c.shtml)

This subsection is divided into the following categories that provide guidance for deploying mandatory steps and procedure in implementing VSS and its components in campus distribution and core.

- VSS Identifiers
- Virtual Switch Link
- Unified Control-Plane
- Multi-Chassis EtherChannel
- VSL Dual-Active Detection and Recovery

### VSS Identifiers

This is the first premigration step to be implemented on two standalone Cisco Catalyst 6500-E in the same campus tier that are planned to be clustered into a single logical entity. Cisco VSS defines the following two types of physical node identifiers to distinguish remote node within the logical entity as well as to set logical VSS domain identity to uniquely identify beyond the single VSS domain boundary.

#### Domain ID


Defining the domain identifier (ID) is the initial step in creating a VSS with two physical chassis. The domain ID value ranges from 1 to 255. Virtual Switch Domain (VSD) is comprised with two physical switches and they must be configured with common domain ID. When implementing VSS in multi-tier campus network design, the unique domain ID between different VSS pair will prevent network protocol conflicts and allow simplified network operation, troubleshooting, and management.

#### Switch ID

In current software version, each VSD supports up to two physical switches to build a logical virtual switch. The switch ID value is 1 or 2. Within VSD, each physical chassis must be uniquely configure switch-ID to successfully deploy VSS. Post VSS migration when two physical chassis is clustered, from the control-plane and management plane perspective, it will create single large system; therefore, all the distributed physical interfaces between two chassis are automatically appended with the switch ID (i.e.,

<switch-id>/<slot#>/<port#> or TenGigabitEthernet 1/1/1. The significance of the switch ID remains within VSD and all the interfaces ID associated to the switch ID will be retained independent of control-plane ownership. See Figure 2-25.

**Figure 2-25 VSS Domain and Switch ID**



The following simple configuration shows how to configure VSS domain ID and switch ID:

Standalone Switch 1:

```
VSS-SW1(config)# switch virtual domain 20
VSS-SW1(config-vs-domain)# switch 1
```

Standalone Switch 2:

```
VSS-SW2(config)# switch virtual domain 20
VSS-SW2(config-vs-domain)# switch 2
```

### Switch Priority

During both virtual-switch bootup processes, the switch priority is negotiated between both virtual switches to determine the control-plane ownership. Virtual-switch configured with high priority takes the control-plane ownership while the low priority switch boots up in redundant mode. The default switch priority is 100, the lower switch ID is a tie-breaker when both virtual-switch node are deployed with default settings.

Cisco recommends deploying both virtual-switch nodes with identical hardware and software to take full advantage of distributed forwarding architecture with centralized control and management plane. The control-plane operation is identical on either of the virtual-switch nodes. Modifying the default switch priority is an optional setting since either of the virtual-switch can provide transparent operation to network and the user.

## Virtual Switch Link

To cluster two physical chassis into single a logical entity, the Cisco VSS technology enables the capability to extend various types of single-chassis internal system components to multi-chassis level. Each virtual-switch must be deployed with the direct physical links and extend the backplane communication boundary over the special links known as Virtual-Switch Link (VSL).

VSL can be considered as Layer 1 physical links between two virtual-switch nodes and is designed to not operate any network control protocols. Therefore, the VSL links cannot establish network protocol adjacencies and are excluded when building the network topology tables. With the customized traffic engineering on VSL, it is tailored to carry the following major traffic categories:

- Inter-Switch Control Traffic
  - Inter-Chassis Ethernet Out Band Channel (EOBC) traffic— Serial Communication Protocol (SCP), IPC, and ICC.
  - Virtual Switch Link Protocol (VSLP) —LMP and RRP control-link packets.
- Network Control Traffic
  - Layer 2 Protocols —STP BPDU, PagP+, LACP, CDP, UDLD, LLDP, 802.1x, DTP, etc.
  - Layer 3 Protocols—ICMP, EIGRP, OSPF, BGP, MPLS LDP, PIM, IGMP, BFD, etc.
- Data Traffic
  - End-user data application traffic in single-home network designs.
  - Integrated service-module with centralized forwarding architecture (i.e., FWSM)
  - Remote SPAN

Using EtherChannel technology, the VSS software design provides the flexibility to increase on-demand VSL bandwidth capacity and to protect the network stability during the VSL link failure or malfunction.

The following sample configuration shows how to configure VSL EtherChannel:

Standalone Switch 1:

```
VSS-SW1(config)# interface Port-Channel 1
VSS-SW1(config-if)# switch virtual link 1

VSS-SW1(config)# interface range Ten 1/1 , Ten 5/4
VSS-SW1(config-if)# channel-group 1 mode on
```

Standalone Switch 2:

```
VSS-SW2(config)# interface Port-Channel 2
VSS-SW2(config-if)# switch virtual link 2

VSS-SW2(config)# interface range Ten 1/1 , Ten 5/4
VSS-SW2(config-if)# channel-group 2 mode on
```

## VSL Design Consideration

Implementing VSL EtherChannel is a simple task; however, the VSL design may require proper design with high reliability, availability, and optimized. Deploying VSL requires careful planning to keep system virtualization intact during VSS system component failure on either virtual-switch node. The strategy for reliable VSL design requires the following three categories of planning:

- VSL Links Diversification
- VSL Bandwidth Capacity
- VSL QoS

### VSL Links Diversification


Complete VSL link failure may break the system virtualization and create network instability during VSL link failure. Designing VSL link redundancy through diverse physical paths on both systems prevents network instability, reduces single point of failure conditions and optimizes bootup process.

All the traffic traverses over the VSL are encoded with special encapsulation header, hence VSL protocols is not designed to operate all Catalyst 6500-E supported linecard module. The next-generation specialized Catalyst 6500-E 10G based supervisor and linecard modules are fully capable and equipped with modern hardware ASICs to build VSL communication. VSL EtherChannel can bundle 10G member-links with any of following next-generate hardware modules:

- Sup720-10G
- WS-X6708
- WS-X6716 (must be deployed in performance mode to enable VSL capability)

Figure 2-26 shows an example of how to build VSL EtherChannel with multiple diverse physical fiber paths from supervisor 10G uplink ports and the VSL-capable 10G hardware modules.

**Figure 2-26 Recommended VSL Links Design**



Deploying VSL with multiple diversified VSL-link design offers the following benefits:

- Leverage 10G port from supervisor and use remaining available ports for other network connectivity.
- Use 10G ports from VSL-capable WS-X6708 or WS-X6716 linecard module to protect against any abnormal failure on supervisor uplink port (i.e., GBIC failure).
- Reduces the single point-of-failure chances as triggering multiple hardware faults on diversified cables, GBIC and hardware modules are rare conditions.
- VSL-enabled 10G module boot up rapidly than other installed modules in system. This software design is required to initialize VSL protocols and communication during bootup process. If the same 10G module is shared to connect other network devices, then depending on the network module type and slot bootup order, it is possible to minimize traffic losses during system initialization process.
- Use 4 class built-in QoS model on each VSL member-links to optimize inter-chassis communication traffic, network control, and user data traffic.

### VSL Bandwidth Capacity

From each virtual-switch node, VSL EtherChannel can bundle up to 8 physical member-links. Therefore, VSL can be bundled up to 80G of bandwidth capacity, the requirement on exact capacity may truly depend on number of the following factors:

- Aggregated network uplink bandwidth capacity on per virtual-switch node basis. For example, 2 x 10GE diversified to same remote peer system.
- Designing the network with single-homed devices connectivity (no MEC) will force at least half of the downstream traffic to flow over the VSL link. This type of connectivity is highly discouraged.

- Remote SPAN from one switch member to other. The SPANed traffic is considered as a single flow, thus the traffic hashes only over a single VSL link that can lead to oversubscription of a particular link. The only way to improve the probability of traffic distribution is to have an additional VSL link. Adding a link increases the chance of distributing the normal traffic that was hashed on the same link carrying the SPAN traffic, which may then be sent over a different link.
- If the VSS is carrying the services hardware (such as FWSM, WiSM, etc.), then depending on the service module forwarding design, it may be carried over the VSL. Capacity planning for each of the supported services blades is beyond the scope of this design guide.

For an optimal traffic load-sharing between VSL member-links, it is recommended to bundle VSL member-link in the power of 2 (i.e., 2, 4, and 8).

### VSL QoS

The network infrastructure and the application demands of next-generation enterprise networks have tremendous amount of dependencies on the strong and resilient network for constant network availability and on-demand bandwidth allocation to provide services compromising performance. Cisco VSS is designed with application intelligence and automatically enables QoS on VSL interface to provide bandwidth and resource allocation for different class-of-service traffic.

The QoS implementation on VSL EtherChannel operates in restricted mode as it carries critical inter-chassis backplane traffic. Independent of global QoS settings, the VSL member-links are automatically configured with system generated QoS settings to protect different class of applications. To retain system stability, the inter-switch VSLP protocols the QoS settings are fine tuned to protect high priority traffic with different thresholds even during VSL link congestion.

To deploy VSL in non-blocking mode and increase the queue depth, the Sup720-10G uplink ports can be configured in one of the following two QoS modes:

- *Default (Non-10G-only mode)*—In this mode, all ports must follow a single queuing mode. If any 10-Gbps port is used for the VSL link, the remaining ports (10 Gbps or 1Gbps) follow the same CoS-mode of queuing for any other non-VSL connectivity because VSL only allows class of service (CoS)-based queuing.
- *Non-blocking (10G-only mode)*—In this mode, all 1-Gbps ports are disabled, as the entire supervisor module operates in a non-blocking mode. Even if only one 10G port used as VSL link, still both 10-Gbps ports are restricted to CoS-based trust model.

Implementing 10G mode may assist in increasing the number of transmit and receive queue depth level; however, restricted VSL QoS prevents reassigning different class-of-service traffic in different queues. Primary benefit in implementing 10G-only mode is to deploy VSL port in non-blocking mode to dedicate complete 10G bandwidth on port. Deploying VSS network based on Cisco's recommendation significantly reduces VSL link utilization, thus minimizing the need to implement 10G-only mode and using all 1G ports for other network connectivities (i.e., out-of-band network management port).

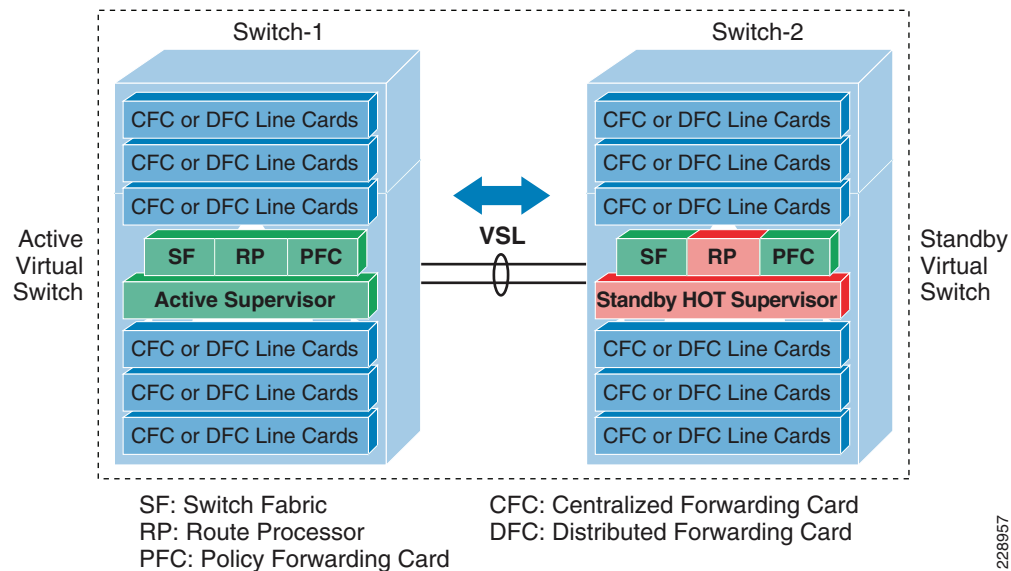
## Unified Control-Plane

Deploying redundant supervisor with common hardware and software components into single standalone Cisco Catalyst 6500-E platform automatically enables the Stateful Switch Over (SSO) capability to provide in-chassis supervisor redundancy in highly redundant network environment. The SSO operation on active supervisor holds control-plane ownership and communicates with remote Layer 2 and Layer 3 neighbors to build distributed forwarding information. SSO-enabled active supervisor is tightly synchronized with standby supervisor with several components (protocol state-machine, configuration, forwarding information, etc.). As a result, if an active supervisor fails, a hot-standby supervisor takes over control-plane ownership and initializes protocol graceful-recovery with peer devices. During network protocol graceful-recovery process the forwarding information remains non-disrupted to continue nonstop packet switching in hardware.



Leveraging the same SSO and NSF technology, the Cisco VSS supports inter-chassis SSO redundancy by extending the supervisor redundancy capability from single-chassis to multi-chassis level. Cisco VSS uses VSL EtherChannel as a backplane path to establish SSO communication between active and hot-standby supervisor deployed in separate physical chassis. Entire virtual-switch node gets reset during abnormal active or hot-standby virtual-switch node failure. See [Figure 2-27](#).

**Figure 2-27 Inter-Chassis SSO Operation in VSS**



To successfully establish SSO communication between two virtual-switch nodes, the following criteria must match between both virtual-switch node:

- Identical software version
- Consistent VSD and VSL interface configuration
- Power mode and VSL-enabled module power settings
- Global PFC Mode
- SSO and NSF-enabled


During the bootup process, the SSO synchronization checks all the above criteria with remote virtual-system. If any of the criteria fails to match, it will force the virtual-switch node to boot in RPR or cold-standby state that cannot synchronize protocol and forwarding information.

### VSL Dual-Active Detection and Recovery

The preceding section described VSL EtherChannel functions as extended backplane link that enables system virtualization by transporting inter-chassis control traffic, network control plane and user data traffic. The state machine of the unified control-plane protocols and distributed forwarding entries gets dynamically synchronized between the two virtual-switch nodes. Any fault triggered on VSL component leads to a catastrophic instability in VSS domain and beyond. The virtual-switch member that assumes the role of hot-standby keeps constant communication with the active switch. The role of the hot-standby switch is to assume the active role as soon as it detects a loss of communication with its peer via all VSL links without the operational state information of the remote active peer node. Such network condition is known as *dual-active*, where both virtual switches get split with common configuration and takes

control plane ownership. The network protocols detect inconsistency and instability when VSS peering devices detect two split systems claiming the same addressing and identifications. Figure 2-28 depicts the state of campus topology in a single active-state and during dual-active state.


**Figure 2-28 Single Active and Dual-Active Campus Topology**



The system virtualization gets impacted during the dual-active network state and splits the single virtual system into two identical Layer 2/3 system. This condition that can destabilize the campus network communication with two split system advertising duplicate information. To prevent such network instability, Cisco VSS introduces the following two methods to rapidly detect dual-active condition and recover the situation by isolating the old active virtual-switch from network operation before the network gets destabilized:

- **Direct Detection Method**—This method requires extra physical connection between both virtual-switch nodes. Dual-Active Fast-Hello (Fast-Hello) and Bidirectional Forwarding Decision (BFD) protocols are specifically designed to detect the dual-active condition and protect network malfunction. All VSS supported Ethernet media and module can be used to deploy this methods. For additional redundancy, VSS allows configuring up to four dual-active fast-hello links between virtual-switch nodes. Cisco recommends deploying Fast-Hello in lieu of BFD for the following reasons:
  - Fast-Hello can rapidly detects dual-active condition and trigger recovery procedure. Independent of routing protocols and network topology, Fast-Hello offers faster network recovery.
  - Fast-Hello enables the ability to implement dual active detection in multi-vendor campus or data-center network environments.
  - Fast-Hello optimize protocol communication procedure without reserving higher system CPU and link overheads.
  - Fast-Hello supersedes BFD-based detection mechanism.
- **Indirect Detection Method**—This method relies on intermediate trusted L2/L3 MEC Cisco Catalyst remote platform to detect the failure and notify to old-active switch about the dual-active detection. Cisco extended the capability of PAGP protocol with extra TLVs to signal the dual-active condition and initiate recovery procedure. Most of the Cisco Catalyst switching platforms can be used as trusted PAGP+ partner to deploy indirect detection method.

All dual-active detection protocol and methods can be implemented in parallel. As depicted in Figure 2-29, in a VSS network deployment peering with Cisco Catalyst platforms, Cisco recommends deploying Fast-Hello and PAGP+ methods for rapid detection, to minimize network topology instability, and to retain application performance intact.

**Figure 2-29 Recommended Dual-Active Detection Method**

The following sample configuration illustrates implementing both methods:

- Dual-Active Fast-Hello

```
cr23-VSS-Core(config)#interface range Gig1/5/1 , Gig2/5/1
cr23-VSS-Core(config-if-range)# dual-active fast-hello

! Following logs confirms fast-hello adjacency is established on
! both virtual-switch nodes.
%VSDA-SW1_SP-5-LINK_UP: Interface Gi1/5/1 is now dual-active detection capable
%VSDA-SW2_SPSTBY-5-LINK_UP: Interface Gi2/5/1 is now dual-active detection capable

cr23-VSS-Core#show switch virtual dual-active fast-hello
Fast-hello dual-active detection enabled: Yes
Fast-hello dual-active interfaces:
Port          Local StatePeer Port    Remote State
-----
Gi1/5/1      Link up      Gi2/5/1      Link up
```

- PAgP+

Enabling or disabling dual-active trusted mode on L2/L3 MEC requires MEC to be in administration shutdown state. Prior to implementing trust settings, network administrator must plan for downtime to provision PAgP+-based dual-active configuration settings:

```
cr23-VSS-Core(config)#int range Port-Channel 101 - 102
cr23-VSS-Core(config-if-range)#shutdown

cr23-VSS-Core(config)#switch virtual domain 20
cr23-VSS-Core(config-vs-domain)#dual-active detection pagp trust channel-group 101
cr23-VSS-Core(config-vs-domain)#dual-active detection pagp trust channel-group 102

cr23-VSS-Core(config)#int range Port-Channel 101 - 102
cr23-VSS-Core(config-if-range)#no shutdown

cr23-VSS-Core#show switch virtual dual-active pagp
PAgP dual-active detection enabled: Yes
PAgP dual-active version: 1.1
```

Channel group 101 dual-active detect capability w/nbrs

Dual-Active trusted group: Yes

Port	Dual-Active Detect Capable	Partner Name	Partner Port	Partner Version
Te1/1/2	Yes	cr22-6500-LB	Te2/1/2	1.1
Te1/3/2	Yes	cr22-6500-LB	Te2/1/4	1.1
Te2/1/2	Yes	cr22-6500-LB	Te1/1/2	1.1
Te2/3/2	Yes	cr22-6500-LB	Te1/1/4	1.1

Channel group 102 dual-active detect capability w/nbrs

Dual-Active trusted group: Yes

Port	Dual-Active Detect Capable	Partner Name	Partner Port	Partner Version
Te1/1/3	Yes	cr24-4507e-MB	Te4/2	1.1
Te1/3/3	Yes	cr24-4507e-MB	Te3/1	1.1
Te2/1/3	Yes	cr24-4507e-MB	Te4/1	1.1
Te2/3/3	Yes	cr24-4507e-MB	Te3/2	1.1

## Virtual Routed MAC

The MAC address allocation for the interfaces does not change during a switchover event when the hot-standby switch takes over as the active switch. This avoids gratuitous ARP updates (MAC address changed for the same IP address) from devices connected to VSS. However, if both chassis are rebooted at the same time and the order of the active switch changes (the old hot-standby switch comes up first and becomes active), then the entire VSS domain will use that switch's MAC address pool. This means that the interface will inherit a new MAC address, which will trigger gratuitous ARP updates to all Layer-2 and Layer-3 interfaces. Any networking device connected one hop away from the VSS (and any networking device that does not support gratuitous ARP), will experience traffic disruption until the MAC address of the default gateway/interface is refreshed or timed out. To avoid such a disruption, Cisco recommends using the configuration option provided with the VSS in which the MAC address for Layer-2 and Layer-3 interfaces is derived from the reserved pool. This takes advantage of the virtual-switch domain identifier to form the MAC address. The MAC addresses of the VSS domain remain consistent with the usage of virtual MAC addresses, regardless of the boot order.

The following configuration illustrates how to configure virtual routed MAC address for Layer 3 interface under switch-virtual configuration mode:

```
cr23-VSS-Core(config)#switch virtual domain 20
cr23-VSS-Core(config-vs-domain)#mac-address use-virtual
```

## Deploying Cisco Catalyst 4500-E

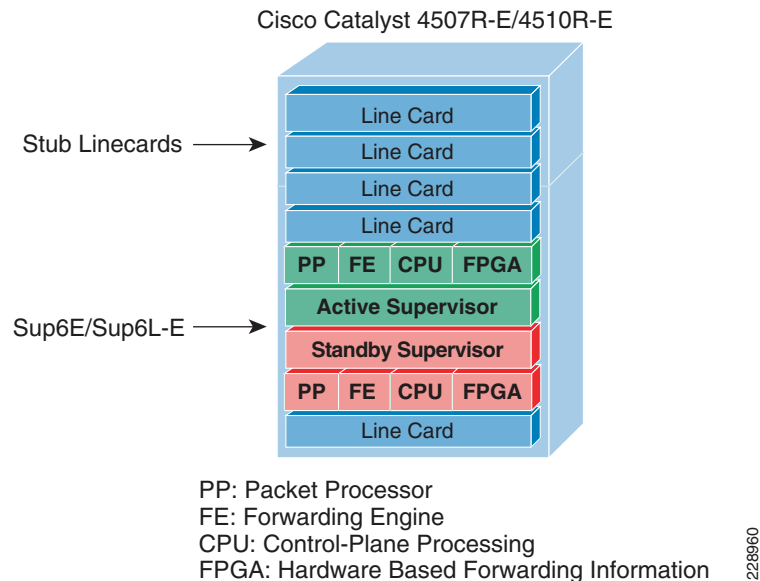
In a mid-size medium enterprise campus network, it is recommended to deploy single highly redundant Cisco Catalyst 4500-E Series platform in the different campus network tiers-access, distribution, core. Cisco Catalyst 4500-E Series switches is a multi-slots modular and scalable and high-speed resilient platform. Single Catalyst 4500-E Series platform in medium enterprise design is build with multiple redundant hardware components to develop consistent network topology as Catalyst 6500-E VSS based large network design. For Catalyst 4500-E in-chassis supervisor redundancy, the network administrators must consider Catalyst 4507R-E or 4510R-E slot chassis to accommodate redundant supervisors and use remaining for LAN network modules.

Cisco Catalyst 4500-E Series supports wide-range of supervisor modules designed for high-performance Layer 2 and Layer 3 network. This reference design recommends deploying next-generation Sup6E and Sup6L-E that supports next-generation hardware switching capabilities, scalability, and performance for various types application and services deployed in campus network.

## Implementing Redundant Supervisor

Cisco Catalyst 4507R-E supports intra-chassis or single-chassis supervisor redundancy with dual-supervisor support. Implementing single Catalyst 4507R-E in highly resilient mode at various campus layer with multiple redundant hardware components will protect against different types of abnormal failures. This reference design guide recommends deploying redundant Sup6E or Sup6L-E supervisor module to deploy full high-availability feature parity. Mid-size core or distribution layer Cisco Catalyst 4507R-E Series platform currently do not support inter-chassis supervisor and node redundancy with VSS technology. Therefore, implementing intra-chassis supervisor redundancy and initial network infrastructure setup will be simplified for medium and small size campus network. [Figure 2-30](#) illustrates Cisco Catalyst 4500-E-based intra-chassis SSO and NSF capability.

**Figure 2-30 Intra-Chassis SSO Operation**



During bootup process, the SSO synchronization checks various criteria to assure both supervisors can provide consistent and transparent network services during failure event. If any of the criteria fails to match, it forces the standby supervisor to boot in RPR or cold-standby state which cannot synchronize protocol and forwarding information from active supervisor. The following sample configuration illustrates how to implement SSO mode on Catalyst 4507R-E and 4510R-E chassis deployed with Sup6E and Sup6L-E redundant supervisors:

```
cr24-4507e-MB#config t
cr24-4507e-MB (config)#redundancy
cr24-4507e-MB (config-red)#mode sso

cr24-4507e-MB#show redundancy states
my state = 13 - ACTIVE
peer state = 8 - STANDBY HOT
< snippet >
```

## Sup6L-E Enhancement

Starting in IOS Release 12.2(53)SG, Cisco introduced new Catalyst 4500 – Sup6L-E supervisor module that is designed and built on the next-generation supervisor Sup6E architecture. As a cost-effective solution, the Sup6L-E supervisor is built with reduced system resources, but also addresses several types of key business and technical challenges for mid- to small-scale size Layer-2 network design.

Initial IP-based IOS Release for Sup6L-E supports SSO capability for multiple types of Layer 2 protocols. To extend its high availability and enterprise-class Layer 3 feature-parity support on Sup6L-E supervisor, it is recommended to deploy IOS Release 12.2(53)SG2 software version with Enterprise license.



### Note

This validated design guide provides the Sup6L-E supervisor deployment guidance and validated test results based on the above recommended software version.

## Deploying Supervisor Uplinks

Every supported supervisor module in Catalyst 4500-E supports different types of uplink ports for core network connectivity. Each Sup6E and Sup6L-E supervisor module supports up to two 10G or can be deployed as four different 1G uplinks using Twin-Gigabit converters. To build high speed low-latency campus backbone network, it is recommended to leverage and deploy 10G uplinks to accommodate various types of bandwidth demanding network application operating in the network.

Cisco Catalyst 4500-E Series supervisors are designed with unique architecture to provide constant network availability and reliability during supervisor reset. Even during supervisor switchover or administrative reset events, the state-machines of all deployed uplinks remain in operation and with centralized forwarding architecture it continues to switch packets without impacting any time-sensitive application like Cisco TelePresence. Such unique architecture protects bandwidth capacity while administrative supervisor switchover is to upgrade IOS software or during abnormal software triggers supervisor reset.

## Sup6E Uplink Port Design

### Non-Redundant Mode

In non-redundant mode, there is a single supervisor module deployed in Catalyst 4500-E chassis. In non-redundant mode, by default both uplink physical ports can be deployed in 10G or 1G with Twin-Gigabit converters. Each port operates in non-blocking state and can switch traffic at the wire-rate performance.

### Redundant Mode


In recommended redundant mode, Catalyst 4507R-E chassis is deployed with dual supervisor. To provide wire-rate switching performance, by default port-group 1 from active and hot-standby supervisor are in active mode and port-group 2 is in the in-active state. The default configuration can be modified by changing Catalyst 4500-E backplane settings to sharing mode. The shared backplane mode enables operation of port-group 2 of both supervisors. Note that sharing the 10G backplane ASIC between two 10G ports does not increase switching capacity; it creates 2:1 oversubscription. If the upstream device is deployed with chassis-redundancy (i.e., Catalyst 6500-E VSS), then it is highly recommended to deploy all four uplink ports for the following reasons:

- Helps developing full-mesh or V shape physical network topology from each supervisor module.
- Increases high availability in the network during individual link, supervisor, or other hardware component failure event.

- Reduces latency and network congestion during rerouting traffic through non-optimal path.

Figure 2-31 summarizes the uplink port support on Sup6E model depends on non-redundant and redundant deployment scenario.

**Figure 2-31 Catalyst 4500-E Sup6E Uplink Mode**



The following sample configuration provides guideline to modify default backplane settings on Catalyst 4507R-E platform deployed with Sup6E supervisors in redundant mode. The new backplane settings will be effective only after complete chassis gets reset; therefore, it is important to plan the downtime during this implementation:

```
cr24-4507e-MB#config t
cr24-4507e-MB(config)#hw-module uplink mode shared-backplane

!A 'redundancy reload shelf' or power-cycle of chassis is required
! to apply the new configuration

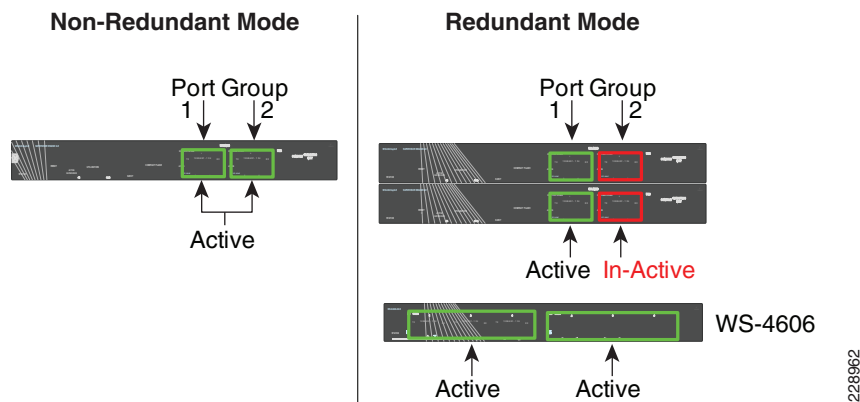
cr24-4507e-MB#show hw-module uplink
Active uplink mode configuration is Shared-backplane

cr24-4507e-MB#show hw-module mod 3 port-group
Module Port-group ActiveInactive
-----
3      1      Te3/1-2Gi3/3-6

cr24-4507e-MB#show hw-module mod 4 port-group
Module Port-group ActiveInactive
-----
4      1      Te4/1-2Gi4/3-6
```


## Sup6L-E Uplink Port Design

The Sup6L-E uplink port function same as Sup6E in non-redundant mode. However, in redundant mode the hardware design of Sup6L-E differs from Sup6E—currently does not support shared backplane mode that allow using all uplink ports actively. The Catalyst 4507R-E deployed with Sup6L-E may use 10G uplink of port group 1 from active and standby supervisor when the upstream device is a single, highly redundant Catalyst 4507R-E chassis. If the upstream device is deployed with chassis-redundancy, (i.e., Cisco VSS), then it is recommended to build full-mesh network design between each supervisor and virtual-switch node. For such design, the network administrator must leverage the existing WS-4606 Series 10G linecard to build full-mesh uplink. Figure 2-32 illustrates the deployment guideline for highly resilient Catalyst 4507R-E-based Sup6L-E uplink.

**Figure 2-32 Catalyst 4500-E Sup6L-E Uplink Mode**

## Deploying Cisco Catalyst 3750-X StackWise Plus

The next-generation Cisco Catalyst 3750-X switches can be deployed in StackWise mode using special stack cable that develops bidirectional physical ring topology. Up to nine switches can be integrated into a single stack ring that offers robust 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 Plus technology. Single switch from the stack ring is selected in master role that manages centralized control-plane process while keeping all member switches in member role. Cisco StackWise Plus solution is designed based on 1:N redundancy option. Master switch election in stack ring is determined based on internal protocol negotiation. During the active master switch failure, the new master is selected based on reelection process that takes place internally through the stack ring. See [Figure 2-33](#).

**Figure 2-33 Cisco StackWise Plus Switching Architecture**

Since Cisco StackWise Plus solution is developed with high redundancy, it offers unique centralized control and management plane with forwarding architecture design. To logically appear as a single virtual switch, the master switch manages complete management-plane and Layer-3 control-plane operations (i.e., IP Routing, CEF, PBR, etc.). Depending on the implemented network protocols, the master switch communicates with rest of the Layer 3 network through stack ring and dynamically develops the best path global routing and updates local hardware with forwarding information.

Unlike centralized Layer-3 management function on master switch, the Layer-2 network topology development is completely based on distributed design. Each member switch in the stack ring dynamically discovers MAC entry from the local port and use internal stack ring network to synchronize MAC address table on each member switch in the stack ring. [Table 2-2](#) lists the network protocols that are designed to operate in centralized versus distributed model in Cisco StackWise Plus architecture.



**Table 2-2 Cisco StackWise Plus 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

Using stack ring as a backplane communication path, master switch updates the Layer-3 forwarding information base (FIB) to each member-switch in the stack ring. Synchronizing common FIB in member switch will develop distributed forwarding architecture. Each member switch performs local forwarding physical path lookup to transmit the frame instead of having master switch performing forwarding path lookup, which may cause traffic hair-pinning problem.

### SSO Operation in 3750-EX StackWise Plus

Cisco StackWise Plus solution offers network and device resiliency with distributed forwarding, but the control plane is not designed like 1+1 redundant design. This is because Cisco Catalyst 3750-X StackWise switch is not an SSO-capable platform that can synchronize control-plane state-machines to a standby switch in the ring. However, it can be configured in NSF-capable mode to gracefully recover from the network during master switch failure. Therefore, when the master switch failure occurs, all the Layer 3 function that is primarily deployed on the uplink ports may get disrupted until new master election occurs and reforms Layer 3 adjacency. Although the new master switch in the stack ring identification is done in range of 0.7 to 1 second, the amount of time for rebuilding the network and forwarding topology depends on the protocol function and scalability.

To prevent Layer 3 disruption in the network caused by master switch failure, the determined master switch with the higher switch priority can be isolated from the uplink Layer 3 EtherChannel bundle path and use physical ports from switches in member role. With the Non-Stop Forwarding (NSF) capabilities in the Cisco StackWise Plus architecture, this network design helps to decrease major network downtime during master switch failure.

### Implementing StackWise Mode

As described earlier, Cisco Catalyst 3750-E switch dynamically detects and provision member-switches in the stack ring without any extra configuration. For planned deployment, network administrator can pre-provision the switch in the ring with the following configuration in global configuration mode:

```
cr36-3750x-xSB(config)#switch 3 provision WS-C3750E-48PD

cr36-3750x-xSB#show running-config | include interface GigabitEthernet3/
interface GigabitEthernet3/0/1
interface GigabitEthernet3/0/2
```

## Switch Priority

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 member-switch is by default configured with switch priority 1.

```
cr36-3750x-xSB#show switch
Switch/Stack Mac Address : 0023.eb7b.e580
```

Switch#	Role	Mac Address	Priority	Version	State	H/W	Current
* 1	Master	0023.eb7b.e580	10		Ready		
2	Member	0026.5284.ec80		1	0		Ready

As described in previous section, the Cisco StackWise architecture is not SSO-capable. This means 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.

```
cr36-3750x-xSB (config)#switch 1 priority 15
Changing the Switch Priority of Switch Number 1 to 15
cr36-3750x-xSB (config)#switch 2 priority 14
Changing the Switch Priority of Switch Number 2 to 14
```

```
cr36-3750x-xSB # show switch
Switch/Stack Mac Address : 0023.eb7b.e580
```

Switch#	Role	Mac Address	Priority	Version	State	H/W	Current
1	Master	0023.eb7b.e580	15	0	Ready		
* 2	Member	0026.5284.ec80	14	0	Ready		

## Stack-MAC Address

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.

```
cr36-3750x-xSB#show switch
Switch/Stack Mac Address : 0023.eb7b.e580
```

Switch#	Role	Mac Address	Priority	Version	State	H/W	Current
1	Master	0023.eb7b.e580	15	0	Ready		
* 2	Member	0026.5284.ec80	14	0	Ready		

```
cr36-3750s-xSB #show version
. . .
Base ethernet MAC Address      : 00:23:EB:7B:E5:80
. . .
```

To prevent 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 Catalyst 3750-X StackWise switch mode using the global configuration CLI mode as shown below:

```
cr36-3750x-xSB (config)#stack-mac persistent timer 0
cr36-3750x-xSB #show switch
Switch/Stack Mac Address : 0026.5284.ec80
Mac persistency wait time: Indefinite
```

Switch#	Role	Mac Address	Priority	H/W Version	Current State
1	Master	0023.eb7b.e580	150	Ready	
* 2	Member	0026.5284.ec80	140	Ready	

## Deploying Cisco Catalyst 3560-X and 2960-S FlexStack

The Medium Enterprise Reference design recommends deploying fixed configuration Cisco Catalyst 3560-X and 2960 Series platform at the campus network edge. The hardware architecture of access-layer fixed configuration is standalone and non-modular in design. These switches are designed to go above traditional access-layer switching function to provide robust next-generation network services (i.e., edge security, PoE+ EnergyWise, etc.).

Cisco Catalyst 3560-X and 2960 Series platform do not support StackWise technology, therefore, these platforms are ready to deploy with a wide-range of network services at the access-layer. All recommended access-layer features and configuration will be explained in following relevant sections.

The access-layer Cisco Catalyst 2960-S Series switches can be stacked using Cisco FlexStack technology that allows stacking up to four switches into single stack ring using special proprietary cable. Cisco FlexStack leverages several architecture components from Cisco Catalyst 3750-X StackWise Plus. However it offers flexibility to upgrade hardware capability in standalone Cisco Catalyst 2960-S series platform to support FlexStack with hot-swappable FlexStack module. The FlexStack module supports dual on-board StackPort each design to support upto 10G switching capacity. The StackPorts on FlexStack module is not a network ports hence it does not run any Layer 2 network protocols, i.e. STP, to develop virtual-switch environment each participating Cisco Catalyst 2960-S in stack-ring runs FlexStack protocol to keep protocols, ports and forwarding information synchronized within the ring. The port configuration and QoS configuration StackPorts are preset and cannot be modified by user, it is design to minimize the network impact due to misconfiguration. From an operational perspective Cisco Catalyst 2960-S FlexStack technology is identical as Cisco Catalyst 3750-X StackWise Plus. Therefore, all the deployment guidelines and best practices defined in [“Deploying Cisco Catalyst 3750-X StackWise Plus” section on page 2-38](#) must be leverage to deploy Cisco Catalyst 2960-S FlexStack in the campus access-layer.

## Designing EtherChannel Network

In this reference design, multiple parallel physical paths are recommended to build highly scalable and resilient medium enterprise network design. Without optimizing the network configuration, by default each interfaces requires network configuration, protocol adjacencies and forwarding information to load-share traffic and provide network redundancy.


The reference architecture of medium enterprise network is design is built upon small- to mid-size enterprise-class network. Depending on the network applications, scalability, and performance requirement, it offers wide-range of campus network designs, platform and technology deployment options in different campus locations and building premises. Each campus network design offers the following set of operation benefits:

- Common network topologies and configuration (all campus network design)
- Simplifies network protocols (eases network operations)
- Increase network bandwidth capacity with symmetric forwarding paths
- Delivers deterministic network recovery performance

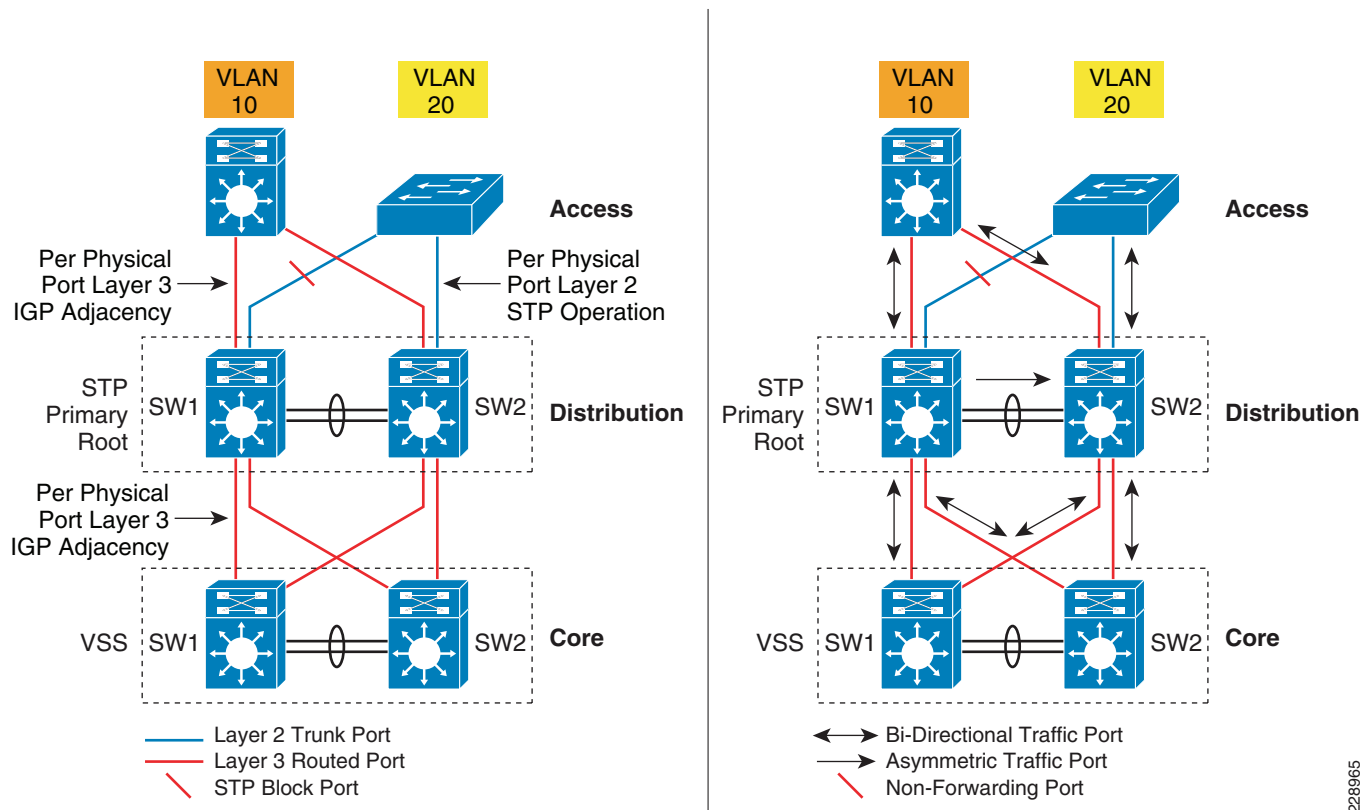
### Diversified EtherChannel Physical Design

As a general best practice to build resilient network designs, it is highly recommended to interconnect all network systems with full-mesh diverse physical paths. Such network design automatically creates multiple parallel paths to provide load-sharing capabilities and path redundancy during network fault events. Deploying single physical connection from a standalone single system to separate redundant upstream systems creates a “V” shape physical network design instead non-recommended partial-mesh “square” network design.

Cisco recommends building full-mesh fiber path between each Layer 2 or Layer 3 operating in standalone, redundant (dual-supervisor) or virtual systems (Cisco VSS and StackWise Plus. Independent of network tier and platform role, this design principle is applicable to all systems across campus network. [Figure 2-34](#) demonstrates recommended deployment physical network design model for various Catalyst platforms.

**Figure 2-34** Designing Diverse Full-mesh Network Topology

Deploying diverse physical network design with redundant mode standalone or the virtual-system running single control-plane will require extra network design tuning to gain all EtherChannel benefits. Without designing the campus network with EtherChannel technology, the individual redundant parallel paths will create network operation state depicted in [Figure 2-35](#). Such network design cannot leverage distributed forwarding architecture and increase operational and troubleshooting complexities. [Figure 2-35](#) demonstrates the default network design with redundant and complex control-plane operation with under-utilized forwarding plane design.

**Figure 2-35 Non-optimized Campus Network Design**

228965

The design in [Figure 2-35](#) suffers from the following challenges for different network modes:

- Layer 3**—Multiple routing adjacencies between two Layer-3 systems. This configuration doubles or quadruples 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. It develops Equal Cost Multi Path (ECMP) symmetric forwarding paths between same Layer 3 peers and offers network scale-dependent Cisco CEF-based network recovery.
- Layer 2**—Multiple parallel Layer-2 paths between STP Root (distribution) and the access switch will build the network loop. To build loop-free network topology, the STP blocks the non-preferred individual link path from forwarding state. With the single STP root virtual-switch, such network topologies cannot fully use all the network resources as well as it creates non-optimal and asymmetric traffic forwarding design.
- VSL Link Utilization**—In a Cisco VSS-based distribution network, it is highly recommended to prevent the condition where it creates hardware or network protocol-driven asymmetric forwarding design (i.e., single-home connection or STP block port). As described in [“Deploying Cisco Catalyst 4500-E” section on page 2-34](#), VSL is not regular network port; it is a special inter-chassis backplane connection used to build virtual system and the network must be designed to switch traffic across VSL-only as a last-resort.

Implementing campus wide MEC or EtherChannel across all the network platforms is the solution for all of the above challenges. Bundling multiple parallel paths into single logical connection builds single loop-free, point-to-point topology that helps to eliminate all protocol-driven forwarding restrictions and program hardware for distributed forwarding to fully use all network resources.

## EtherChannel Fundamentals

In a standalone EtherChannel mode, multiple and diversified member-links are physically connected in parallel between two same physical systems. All the key network devices in the Medium Enterprise Reference design support EtherChannel technology. Independent of campus location and the network layer—campus, data center, WAN/Internet edge, all the EtherChannel fundamentals and configuration guideline described in this section remain consistent.

### Multi-Chassis EtherChannel Fundamentals

Cisco's Multi-Chassis EtherChannel (MEC) technology is a breakthrough innovation that lifts up barrier to create logical point-to-point EtherChannel by distributing physical connection to each highly resilient virtual-switch node in the VSS domain. Deploying Layer 2 or Layer 3 MEC with VSS introduces the following benefits:

- In addition to all EtherChannel benefits, the distributed forwarding architecture in MEC helps increasing network bandwidth capacity.
- Increases network reliability by eliminating single point-of-failure limitation compare to traditional EtherChannel technology.
- Simplifies network control-plane, topology, and system resources with single logical bundled interface instead multiple individual parallel physical paths.
- Independent of network scalability, MEC provides deterministic hardware-based subsecond network recovery.
- MEC technology which remains transparent operation to remote peer devices.

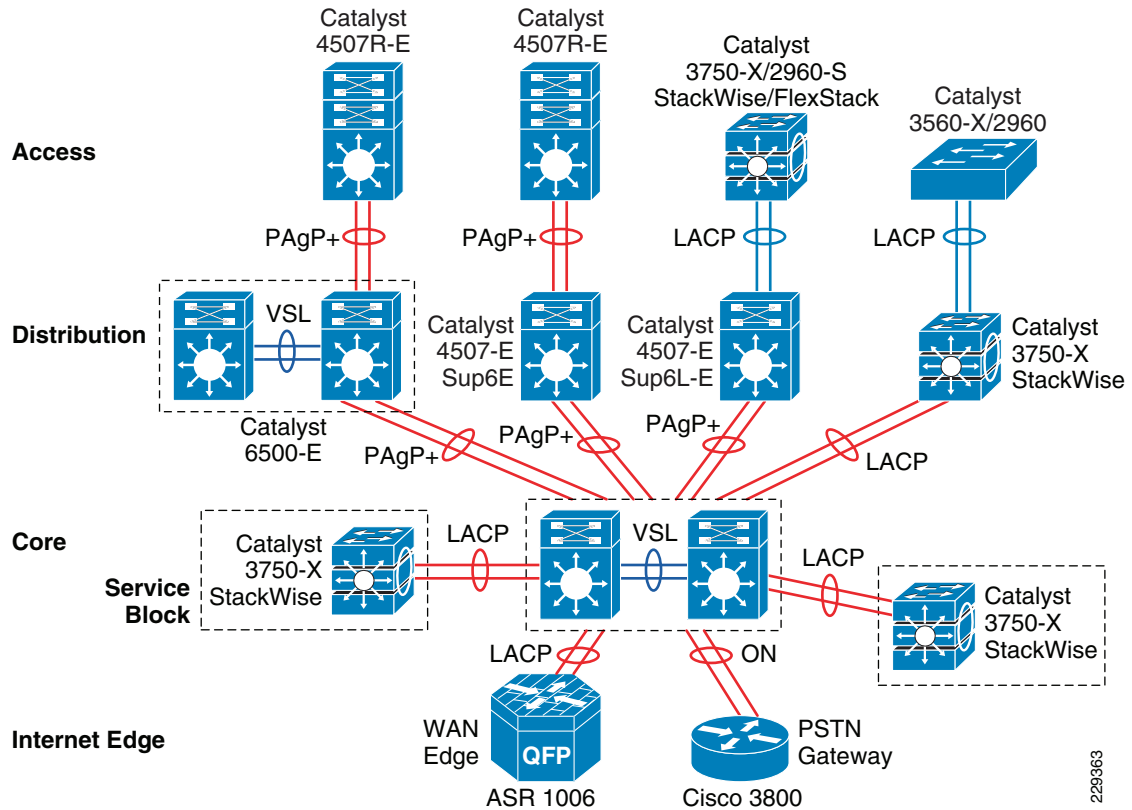
## Implementing EtherChannel

In a standalone EtherChannel mode, multiple and diversified member-links are physically connected in parallel between two same physical systems. All the key network devices in the medium enterprise network design support EtherChannel technology. Independent of campus location and the network layer—campus, data center, WAN/Internet edge, all the EtherChannel fundamentals and configuration guideline described in this section remain consistent.

## Port-Aggregation Protocols

The member-links of EtherChannel must join the port-channel interface using Cisco PAgP+ or industry standard LACP port-aggregation protocols. Both protocols are designed to provide identical benefits. Implementing these protocols provides the following additional benefits:

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

**Figure 2-36 Network-Wide Port-Aggregation Protocol Deployment Guidelines**

Port-aggregation protocol support varies on various types of Cisco platforms; therefore, depending on each end of EtherChannel device types, Cisco recommends deploying the port-channel settings specified in [Table 2-3](#).

**Table 2-3 MEC Port-Aggregation Protocol Recommendation**

Port-Agg Protocol	Local Node	Remote Node	Bundle State
PAgP+	Desirable	Desirable	Operational
LACP	Active	Active	Operational
None <sup>1</sup>	ON	ON	Operational

1. None or Static Mode EtherChannel configuration must be deployed in exceptional cases when remote node do not support either of the port-aggregation protocols. To prevent network instability, network administrator must implement static mode port-channel with special attention that assures no configuration in-compatibility between bundling member-link ports.

The implementation guidelines to deploy EtherChannel and MEC in Layer 2 or Layer 3 mode are simple and consistent. The following sample configuration provides a guidance to implement single point-to-point Layer-3 MEC from diverse physical ports in different module slots that physically resides in two virtual-switch chassis to a single redundant mode, standalone Catalyst 4507R-E system:

- MEC—VSS-Core

```
cr23-VSS-Core(config)#interface Port-channel 102
cr23-VSS-Core(config-if)# ip address 10.125.0.14 255.255.255.254
! Bundling single MEC diversified physical ports and module on per node basis.
```



```

cr23-VSS-Core(config)#interface range Ten1/1/3 , Ten1/3/3 , Ten2/1/3 , Ten2/3/3
cr23-VSS-Core(config-if-range)#channel-protocol pagp
cr23-VSS-Core(config-if-range)#channel-group 102 mode desirable

cr23-VSS-Core#show etherchannel 102 summary | inc Te
102      Po102(RU)      PAgP      Te1/1/3(P)      Te1/3/3(P)      Te2/1/3(P)      Te2/3/3(P)
cr23-VSS-Core#show pagp 102 neighbor | inc Te
Te1/1/3      cr24-4507e-MB      0021.d8f5.45c0      Te4/2      27s SC      10001
Te1/3/3      cr24-4507e-MB      0021.d8f5.45c0      Te3/1      28s SC      10001
Te2/1/3      cr24-4507e-MB      0021.d8f5.45c0      Te4/1      11s SC      10001
Te2/3/3      cr24-4507e-MB      0021.d8f5.45c0      Te3/2      11s SC      10001

```

- EtherChannel—Catalyst 4507R-E Distribution

```

cr24-4507e-MB (config)#interface Port-channel 1
cr24-4507e-MB (config-if)# ip address 10.125.0.15 255.255.255.254
! Bundling single EtherChannel diversified on per physical ports and per supervisor
basis.
cr24-4507e-MB (config)#interface range Ten3/1 - 2 , Ten4/1 - 2
cr24-4507e-MB (config-if-range)#channel-protocol pagp
cr24-4507e-MB (config-if-range)#channel-group 1 mode desirable

cr24-4507e-MB #show etherchannel 101 summary | inc Te
1      Po1 (RU)      PAgP      Te3/1(P)      Te3/2(P)      Te4/1(P)      Te4/2(P)

cr24-4507e-MB#show pagp 1 neighbor | inc Te
Te3/1      cr23-VSS-Core      0200.0000.0014      Te1/3/3      26s SC      660001
Te3/2      cr23-VSS-Core      0200.0000.0014      Te2/3/3      15s SC      660001
Te4/1      cr23-VSS-Core      0200.0000.0014      Te2/1/3      25s SC      660001
Te4/2      cr23-VSS-Core      0200.0000.0014      Te1/1/3      11s SC      660001

```

### EtherChannel Load-Sharing

The numbers of applications and their function in campus network design becomes highly variable, especially when the network is provided as a common platform for business operation, campus security and open accessibility to the users. It becomes important for the network to become more intelligence-aware with deep packet-inspection and load-share the traffic by fully using all network resources.

Fine tuning EtherChannel and MEC add an extra computing intelligence in the network to make protocol-aware egress forwarding decision between multiple local member-links paths. For each traffic flow, such tuning optimizes the egress path-selection procedure with multiple levels of variable information that are originated by the source host (i.e., Layer 2 to Layer 4). EtherChannel load-balancing method supports varies on Cisco Catalyst platforms. [Table 2-4](#) summarizes the currently supported EtherChannel load-balancing methods.

**Table 2-4 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, 6500
		src-mac	
		dst-mac	
		src-dst-mac	
IP	Layer 3	src-ip	
		dst-ip	
		src-dst-ip (recommended)	
IP	Layer 4	src-port	4500, 6500
		dst-port	
		src-dst-port	
IP	XOR L3 and L4	src-dst-mixed-ip-port (recommended)	6500

## Implementing EtherChannel Load-Sharing

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 the higher granularity and optimal utilization of each member-link port, an EtherChannel can intelligently load-share egress traffic using different algorithms.

All Cisco Catalyst 29xx-S, 3xxx-X, and 4500-E switching must be tuned with optimal EtherChannel load-sharing capabilities similar to the following sample configuration:

```
cr24-4507e-MB(config)#port-channel load-balance src-dst-ip
cr24-4507e-MB#show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
    src-dst-ip
```

## Implementing MEC Load-Sharing

The next-generation Catalyst 6500-E Sup720-10G supervisor introduces more intelligence and flexibility to load-share traffic with upto 13 different traffic patterns. Independent of virtual-switch role, each node in VSD uses same polymorphic algorithm to load-share egress Layer 2 or Layer 3 traffic across different member-links from local chassis. When computing the load-sharing hash, each virtual-switch node includes local physical ports of MEC instead remote switch ports; this customized load-sharing is design to prevent traffic reroute over the VSL. It is recommended to implement the following MEC load-sharing configuration in the global configuration mode:

```
cr23-VSS-Core(config)#port-channel load-balance src-dst-mixed-ip-port

cr23-VSS-Core#show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
    src-dst-mixed-ip-port vlan included
```

**Note**

MEC load-sharing becomes effective only when each virtual-switch node have more than one physical path in same bundle interface.

## MEC Hash Algorithm

Like MEC load sharing, the hash algorithm is computed independently by each virtual-switch to perform load share via its local physical ports. Traffic-load share is defined based on number of internal bits allocated to each local member-link ports. Cisco Catalyst 6500-E system in VSS mode assigns 8 bits to every MEC, 8-bit can be represented as 100 percent switching load. Depending on number of local member-link ports in an MEC bundle, the 8-bit hash is computed and allocated to each port for optimal load-sharing result. Like standalone network design, VSS supports the following EtherChannel hash algorithms:

- *Fixed*—Default setting. Keep it default if each virtual-switch node has single local member-link port bundled in same L2/L3 MEC (total 2 ports in MEC).
- *Adaptive*—Best practice is to modify to adaptive hash method if each virtual-switch node has greater than or equal to two physical ports in the same L2/L3 MEC.

When deploying full-mesh V-shape network VSS-enabled campus core network, it is recommended to modify default MEC hash algorithm from default settings as shown in the following sample configuration:

```
cr23-VSS-Core(config)#port-channel hash-distribution adaptive
```

Modifying MEC hash algorithm to adaptive mode requires the system to internally reprogram hash result on each MEC. Therefore, plan for additional downtime to make new configuration effective.

```
cr23-VSS-Core(config)#interface Port-channel 101
cr23-VSS-Core(config-if)#shutdown
cr23-VSS-Core(config-if)#no shutdown

cr23-VSS-Core#show etherchannel 101 detail | inc Hash
Last applied Hash Distribution Algorithm: Adaptive
```

## Network Addressing Hierarchy

Developing a structured and hierarchical IP address plan is as important as any other design aspect of the medium enterprise network to create an efficient, scalable, and stable network design. Identifying an IP addressing strategy for the network for the entire medium enterprise network design is essential.

**Note**

This section does not explain the fundamentals of TCP/IP addressing; for more details, see the many Cisco Press publications that cover this topic.

The following are key benefits of using hierarchical IP addressing:

- *Efficient address allocation*
  - Hierarchical addressing provides the advantage of grouping all possible addresses contiguously.
  - In non-contiguous addressing, a network can create addressing conflicts and overlapping problems, which may not allow the network administrator to use the complete address block.
- *Improved routing efficiencies*

- Building centralized main and remote campus site networks with contiguous IP addresses provides an efficient way to advertise summarized routes to neighbors.
- Route summarization simplifies the routing database and computation during topology change events.
- Reduces network bandwidth utilization used by routing protocols.
- Improves overall routing protocol performance by flooding less messages and improves network convergence time.
- *Improved system performance*
  - Reduces the memory needed to hold large-scale discontinuous and non-summarized route entries.
  - Reduce higher CPU power to re-compute large-scale routing databases during topology change events.
  - Becomes easier to manage and troubleshoot.
  - Helps in overall network and system stability.

## Network Foundational Technologies for LAN Design

In addition to a hierarchical IP addressing scheme, it is also essential to determine which areas of the medium enterprise design are Layer 2 or Layer 3 to determine whether routing or switching fundamentals need to be applied. The following applies to the three layers in a LAN design model:

- *Core layer*—Because this is a Layer 3 network that interconnects several remote locations and shared devices across the network, choosing a routing protocol is essential at this layer.
- *Distribution layer*—The distribution block 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. Both routing and switching fundamentals need to be applied.
- *Access layer*—This layer is the demarcation point between network infrastructure and computing devices. This is designed for critical network edge functions to provide intelligent application and device-aware services, to set the trust boundary to distinguish applications, provide identity-based network access to protected data and resources, provide physical infrastructure services to reduce greenhouse emission, and more. This subsection provides design guidance to enable various types of Layer 1 to 3 intelligent services, and to optimize and secure network edge ports.

The recommended routing or switching scheme of each layer is discussed in the following sections.

### Designing the Core Layer Network

Because the core layer is a Layer 3 network, routing principles must be applied. Choosing a routing protocol is essential, and routing design principles and routing protocol selection criteria are discussed in the following subsections.

## Routing Design Principles

Although enabling routing functions in the core is a simple task, the routing blueprint must be well understood and designed before implementation, because it provides the end-to-end reachability path of the enterprise network. For an optimized routing design, the following three routing components must be identified and designed to allow more network growth and provide a stable network, independent of scale:

- *Hierarchical network addressing*—Structured IP network addressing in the medium enterprise LAN and/or WAN design is required to make the network scalable, optimal, and resilient.
- *Routing protocol*—Cisco IOS supports a wide range of Interior Gateway Protocols (IGPs). Cisco recommends deploying a single routing protocol across the medium enterprise network infrastructure.
- *Hierarchical routing domain*—Routing protocols must be designed in a hierarchical model that allows the network to scale and operate with greater stability. Building a routing boundary and summarizing the network minimizes the topology size and synchronization procedure, which improves overall network resource use and re-convergence.

## Routing Protocol Selection Criteria

The criteria for choosing the right protocol vary based on the end-to-end network infrastructure. Although all the routing protocols that Cisco IOS currently supports can provide a viable solution, 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*—Requires a proven protocol that can scale in full-mesh campus network designs and can optimally function in hub-and-spoke WAN network topologies.
- *Scalability*—The routing protocol function must be network- and system-efficient and operate with a minimal number of updates and re-computation, independent of the number of routes in the network.
- *Rapid convergence*—Link-state versus DUAL re-computation and synchronization. Network re-convergence also varies based on network design, configuration, and a multitude of other factors that may be more than a specific routing protocol can handle. The best convergence time can be achieved from a routing protocol if the network is designed to the strengths of the protocol.
- *Operational*—A simplified routing protocol that can provide ease of configuration, management, and troubleshooting.

Cisco IOS supports a wide range of routing protocols, such as Routing Information Protocol (RIP) v1/2, Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Intermediate System-to-Intermediate System (IS-IS). However, Cisco recommends using EIGRP or OSPF for this network design. EIGRP is a popular version of an Interior Gateway Protocol (IGP) because it has all the capabilities needed for small to large-scale networks, offers rapid network convergence, and above all is simple to operate and manage. OSPF is popular link-state protocol for large-scale enterprise and service provider networks. OSPF enforces hierarchical routing domains in two tiers by implementing backbone and non-backbone areas. The OSPF area function depends on the network connectivity model and the role of each OSPF router in the domain. OSPF can scale higher but the operation, configuration, and management might become too complex for the medium enterprise LAN network infrastructure.

Other technical factors must be considered when implementing OSPF in the network, such as OSPF router type, link type, maximum transmission unit (MTU) considerations, designated router (DR)/backup designated router (BDR) priority, and so on. This document provides design guidance for using simplified EIGRP in the medium enterprise campus and WAN network infrastructure.

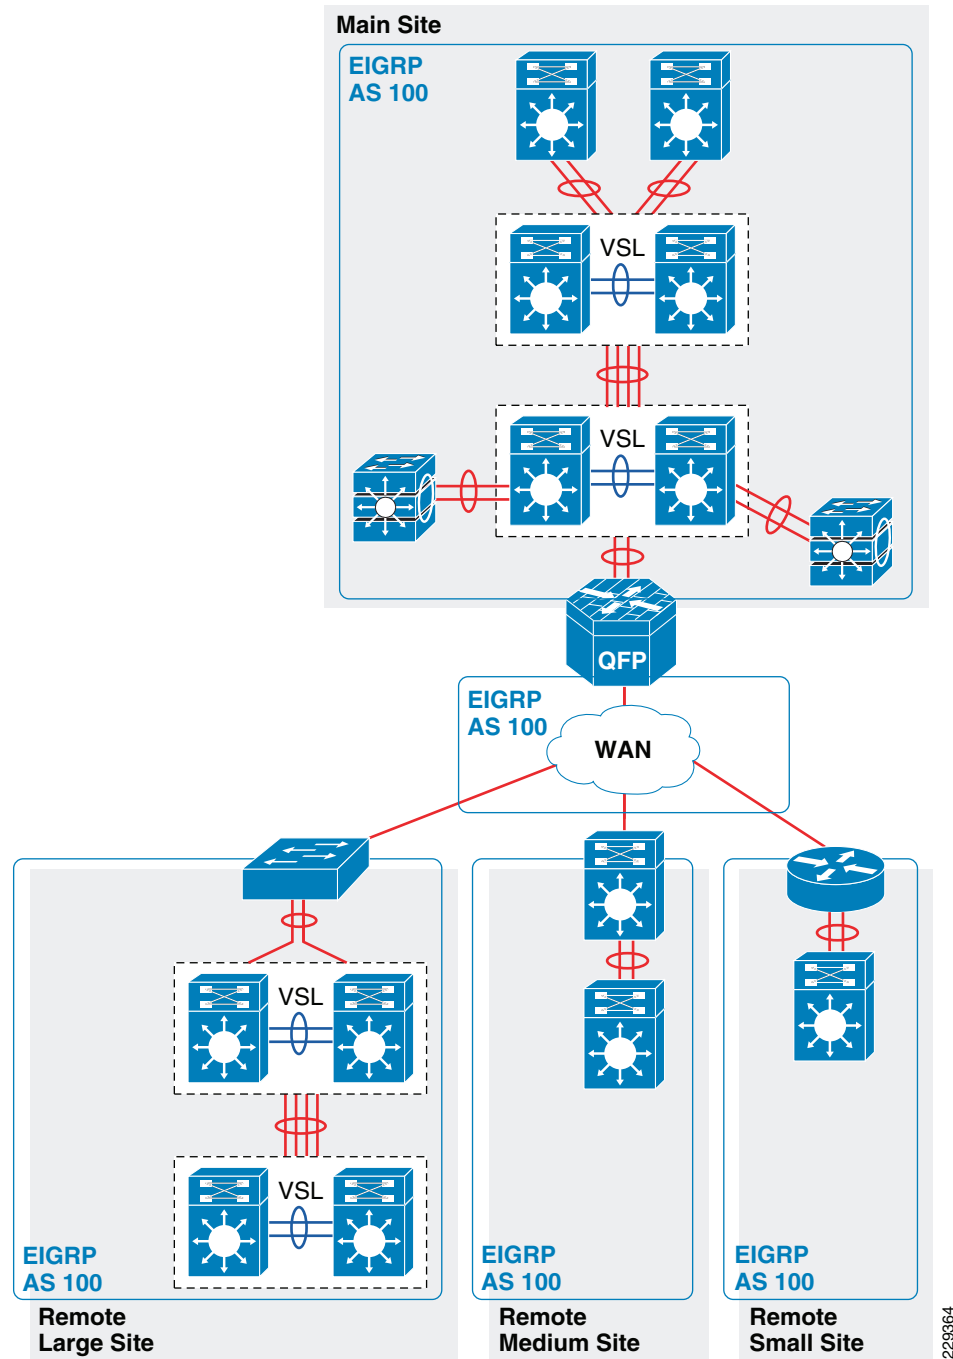
**Note**

For detailed information on EIGRP and OSPF, see the following URL:  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/routed-ex.html>.

## Designing an End-to-End EIGRP Routing Network

EIGRP is a balanced hybrid routing protocol that builds neighbor adjacency and flat routing topology on a per autonomous system (AS) basis. Cisco recommends considering the following three critical design tasks before implementing EIGRP in the medium enterprise LAN core layer network:

- *EIGRP autonomous system*—The Layer 3 LAN and WAN infrastructure of the medium enterprise design must be deployed in a single EIGRP AS, as shown in [Figure 2-37](#). A single EIGRP AS reduces operational tasks and prevents route redistribution, loops, and other problems that may occur because of misconfiguration. [Figure 2-37](#) illustrates end-to-end single EIGRP Autonomous network design in medium enterprise network.

**Figure 2-37** Sample End-to-End EIGRP Routing Design in Medium Enterprise LAN Network

229364

## Implementing EIGRP Routing Protocol

The following sample configuration provides deployment guideline for implement EIGRP routing protocol on all Layer-3 network devices into a single Autonomous System (AS):

```
cr23-VSS-Core(config)#router eigrp 100
cr23-VSS-Core(config-router)# network 10.0.0.0
```

```

cr23-VSS-Core(config-router)# eigrp router-id 10.125.200.254
cr23-VSS-Core(config-router)# no auto-summary

cr23-VSS-Core#show ip eigrp neighbors
EIGRP-IPv4 neighbors for process 100
H   Address                Interface      Hold    Uptime    SRTT    RTO    Q    Seq
                               (sec)        (ms)          Cnt  Num
7   10.125.0.13             Po101         12      3d16h     1       200    0    62
0   10.125.0.15             Po102         10      3d16h     1       200    0    503
1   10.125.0.17             Po103         11      3d16h     1       200    0    52
...

cr23-VSS-Core#show ip route eigrp | inc /16|/20|0.0.0.0
10.0.0.0/8 is variably subnetted, 41 subnets, 5 masks
D    10.126.0.0/16 [90/3072] via 10.125.0.23, 08:33:16, Port-channel106
D    10.125.128.0/20 [90/3072] via 10.125.0.17, 08:33:15, Port-channel103
D    10.125.96.0/20 [90/3072] via 10.125.0.13, 08:33:18, Port-channel101
D    10.125.0.0/16 is a summary, 08:41:12, Null0
...
D*EX 0.0.0.0/0 [170/515072] via 10.125.0.27, 08:33:20, Port-channel108

```

- **EIGRP adjacency protection**—This increases network infrastructure efficiency and protection by securing the EIGRP adjacencies with internal systems. This task involves two subset implementation tasks on each EIGRP-enabled network devices:
  - **Increases system efficiency**—Blocks EIGRP processing with passive-mode configuration on physical or logical interfaces connected to non- EIGRP devices in the network, such as PCs. The best practice helps reduce CPU utilization and secures the network with unprotected EIGRP adjacencies with untrusted devices. The following sample configuration provide guidelines to enable EIGRP protocol communication on trusted interface and block on all system interfaces. This recommended best practice must be enabled on all the EIGRP Layer 3 systems in the network:
 

```

cr23-VSS-Core(config)#router eigrp 100
cr23-VSS-Core(config-router)# passive-interface default
cr23-VSS-Core(config-router)# no passive-interface Port-channel101
cr23-VSS-Core(config-router)# no passive-interface Port-channel102
<snippet>

```
  - **Network security**—Each EIGRP neighbor in the LAN/WAN network must be trusted by implementing and validating the Message-Digest algorithm 5 (MD5) authentication method on each EIGRP-enabled system in the network. Following recommended EIGRP MD5 adjacency authentication configuration must on each non-passive EIGRP interface to establish secure communication with remote neighbors. This recommended best practice must be enabled on all the EIGRP Layer 3 systems in the network:
 

```

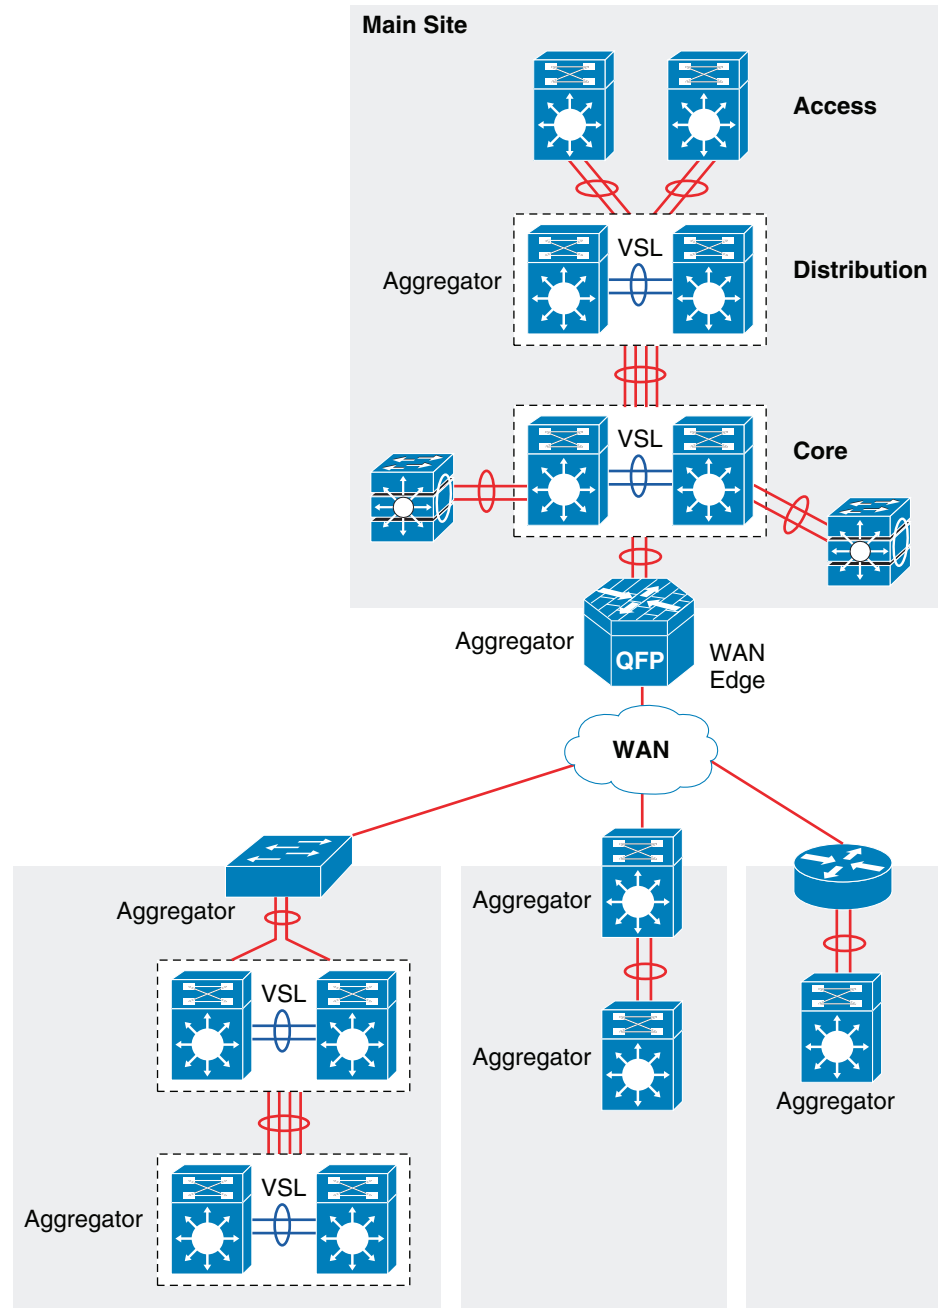
cr23-VSS-Core(config)#key chain eigrp-key
cr23-VSS-Core(config-keychain)# key 1
cr23-VSS-Core(config-keychain-key)#key-string <password>

cr23-VSS-Core(config)#interface range Port-Channel 101 - 108
cr23-VSS-Core(config-if-range)# ip authentication mode eigrp 100 md5
cr23-VSS-Core(config-if-range)# ip authentication key-chain eigrp 100 eigrp-key

```

- **Optimizing EIGRP topology**—EIGRP allows network administrators to summarize multiple individual and contiguous networks into a single summary network before advertising to the neighbor. Route summarization helps improve network performance, stability, and convergence by hiding the fault of an individual network that requires each router in the network to synchronize the routing topology. Each aggregating device must summarize a large number of networks into a single summary route. [Figure 2-38](#) shows an example of the EIGRP topology for the medium enterprise LAN design.



**Figure 2-38 EIGRP Route Aggregator Design**

The following configuration must be applied on each EIGRP route aggregator system as depicted in Figure 2-38. EIGRP route summarization must be implemented on upstream logical port-channel interface to announce single prefix from each block.

```
cr22-6500-LB(config)#interface Port-channel100
cr22-6500-LB(config-if)# ip summary-address eigrp 100 10.125.96.0 255.255.240.0
```

```

cr22-6500-LB#show ip protocols
...
  Address Summarization:
    10.125.96.0/20 for Port-channel100
<snippet>

cr22-6500-LB#s ip route | inc Null0
D      10.125.96.0/20 is a summary, 3d16h, Null0

```

- *EIGRP Timers*—By default, EIGRP speakers transmit Hello packets every 5 seconds, and terminates EIGRP adjacency if the neighbor fails to receive it within 15 seconds of hold-down time. In this network design, Cisco recommends retaining default EIGRP Hello and Hold timers on all EIGRP-enabled platforms.

## Designing the Campus Distribution Layer Network

This section provides design guidelines for deploying various types of Layer 2 and Layer 3 technology in the distribution layer. Independent of which implemented distribution layer design model is deployed, the deployment guidelines remain consistent in all designs.

Because the distribution layer can be deployed with both Layer 2 and Layer 3 technologies, the following two network designs are recommended:

- Multilayer
- Routed access

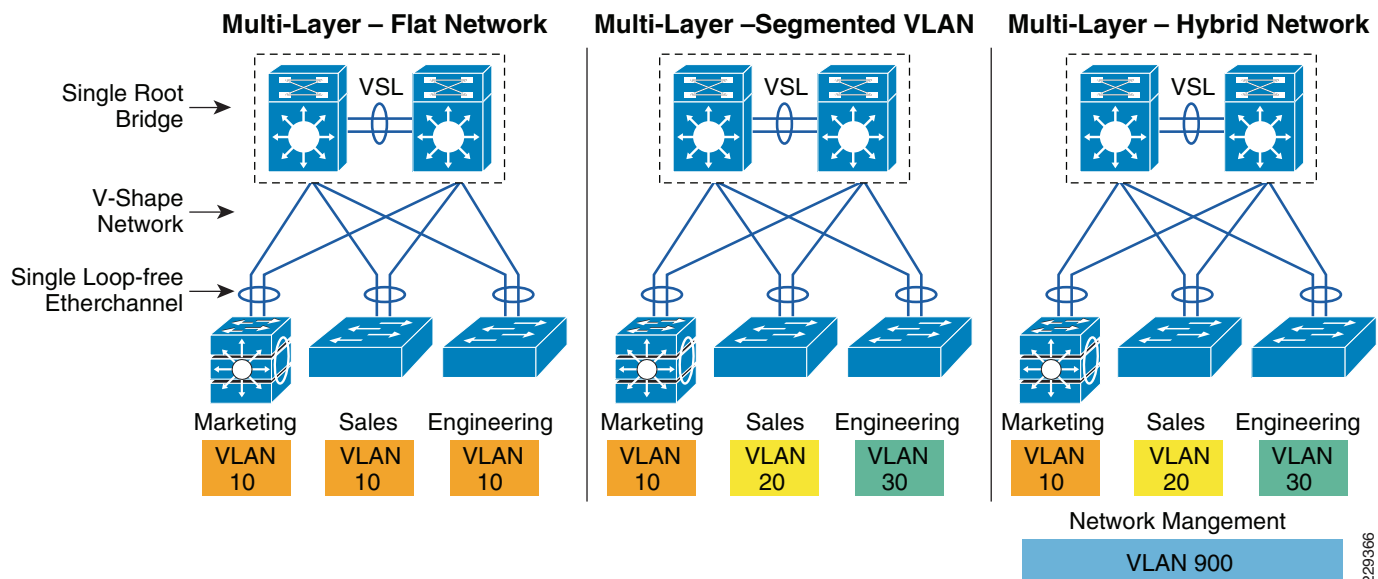
## Designing the Multilayer Network

A multilayer network is a traditional, simple, and widely deployed scenario, regardless of network scale. The access layer switches in the campus network edge interface with various types of endpoints and provide intelligent Layer 1/2 services. The access layer switches interconnect to distribution switches with the Layer 2 trunk, and rely on the distribution layer aggregation switch to perform intelligent Layer 3 forwarding and to set policies and access control.

There are the following three design variations to build a multilayer network; all variations must be deployed in a V-shape physical network design and must be built to provide a loop-free topology:

- *Flat*—Certain applications and user access requires that the broadcast domain design span more than a single wiring closet switch. The multilayer network design provides the flexibility to build a single large broadcast domain with an extended star topology. Such flexibility introduces scalability, performance, and security challenges, and may require extra attention to protect the network against misconfiguration and miswiring that can create spanning-tree loops and de-stabilize the network.
- *Segmented*—Provides a unique VLAN for different organization divisions and enterprise business function segments to build a per-department logical network. All network communication between various enterprise and administrative groups passes through the routing and forwarding policies defined at the distribution layer.
- *Hybrid*—A hybrid logical network design segments VLAN workgroups that do not span different access layer switches, and allows certain VLANs (for example, that net management VLAN) to span across the access-distribution block. The hybrid network design enables flat Layer 2 communication without impacting the network, and also helps reduce the number of subnets used.

Figure 2-39 shows the three design variations for the multilayer network.

**Figure 2-39 Multilayer Design Variations**

Cisco recommends that the hybrid multilayer access-distribution block design use a loop-free network topology, and span a few VLANs that require such flexibility, such as the management VLAN.

The following sample configuration provides guideline to deploy several types of multilayer network components for hybrid multilayer access-distribution block. All the configuration and best practices remains consistent and can deployed independent of Layer 2 platform type and campus location:

## VTP

VLAN Trunking Protocol (VTP) is a Cisco proprietary Layer 2-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. It is recommended to deploy VTP in transparent mode, set the VTP domain name and change the mode to the transparent mode as follows:

```
cr22-3750-LB(config)#vtp domain CCVE-LB
cr22-3750-LB(config)#vtp mode transparent
cr22-3750-LB(config)#vtp version 2

cr22-3750-LB#show vtp status
VTP Version capable:1 to 3
VTP version running:2
VTP Domain Name:CCVE-LB
```

## VLAN

```
cr22-3750-LB(config)#vlan 101
cr22-3750-LB(config-vlan)#name Untrusted_PC_VLAN
cr22-3750-LB(config)#vlan 102
cr22-3750-LB(config-vlan)#name Lobby_IP_Phone_VLAN
cr22-3750-LB(config)#vlan 900
cr22-3750-LB(config-vlan)#name Mgmt_VLAN

cr22-3750-LB#show vlan | inc 101|102|900
101 Untrusted_PC_VLANactive Gi1/0/1
```

```

102 Lobby_IP_Phone_VLANactive    Gi1/0/2
900 Mgmt_VLANactive

```

## Implementing Layer 2 Trunk

In a typical campus network design, a single access switch will be deployed with more than single VLAN, for example a Data VLAN and a Voice VLAN. The Layer-2 network connection between the distribution and access device is a trunk interface. VLAN tag is added to maintain logical separation between VLANs across the trunk. It is recommended to implement 802.1Q trunk encapsulation in static mode instead of negotiating mode, to improve the rapid link bring-up performance.

Enabling the Layer-2 trunk on a port-channel automatically enables communication for all of the active VLANs between the access and distribution. This may create an adverse impact in the large scale network, the access-layer switch may receive traffic flood destined to another access switch. Hence it is important to limit traffic on Layer-2 trunk ports by statically allowing the active VLANs to ensure efficient and secure network performance. Allowing only assigned VLANs on a trunk port automatically filters rest.

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. 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 801 in the access-switch and in the distribution switch. Then change the default native VLAN setting in both the switches. Thereafter, VLAN 801 must not be used anywhere for any purpose in the same access-distribution block.

The following is the configuration example to implement Layer-2 trunk, filter VLAN list and configure the native-VLAN to prevent attacks and optimize port channel interface. When the following configurations are applied on port-channel interface (i.e., Port-Channel 1), they are automatically inherited on each bundled member-link (i.e., Gig1/0/49 and Gig1/0/50):

### Access-Layer

```

cr22-3750-LB(config)#vlan 801
cr22-3750-LB(config-vlan)#name Hopping_VLAN

cr22-3750-LB(config)#interface Port-channel1
cr22-3750-LB(config-if)#description Connected to cr22-6500-LB
cr22-3750-LB(config-if)#switchport
cr22-3750-LB(config-if)#switchport trunk encapsulation dot1q
cr22-3750-LB(config-if)#switchport trunk native vlan 801
cr22-3750-LB(config-if)#switchport trunk allowed vlan 101-110,900
cr22-3750-LB(config-if)#switchport mode trunk

cr22-3750-LB#show interface port-channel 1 trunk

```

Port	Mode	Encapsulation	Status	Native vlan
Po1	on	802.1q	trunking	801

Port	Vlans allowed on trunk
Po1	101-110,900

Port	Vlans allowed and active in management domain
Po1	<b>101-110,900</b>
Port	Vlans in spanning tree forwarding state and not pruned
Po1	101-110,900

## Spanning-Tree in Multilayer Network

Spanning Tree (STP) is a Layer-2 protocol that prevents logical loops in switched networks with redundant links. The medium enterprise LAN network design uses Etherchannel or MEC (point-to-point logical Layer-2 bundle) connection between access-layer and distribution switch which inherently simplifies the STP topology and operation. In this point-to-point network 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.

The following is the example configuration for enabling STP in multilayer network:

### Distribution-Layer

```
cr22-6500-LB(config)#spanning-tree mode rapid-pvst

cr22-6500-LB #show spanning-tree summary | inc mode

!Switch is in rapid-pvst mode
```

### Access-Layer


```
cr22-3750-LB(config)#spanning-tree mode rapid-pvst
```

## Hardening Spanning-Tree Toolkit

Ensuring a loop-free topology is critical in a multilayer network design. Spanning-Tree Protocol (STP) dynamically develops a loop-free multilayer network topology that can compute the best forwarding path and provide redundancy. Although STP behavior is deterministic, it is not optimally designed to mitigate network instability caused by hardware miswiring or software misconfiguration. Cisco has developed several STP extensions to protect against network malfunctions, and to increase stability and availability. All Cisco Catalyst LAN switching platforms support the complete STP toolkit suite that must be enabled globally on individual logical and physical ports of the distribution and access layer switches.

Figure 2-40 shows an example of enabling various STP extensions on distribution and access layer switches in all campus sites.

**Figure 2-40** Protecting Multilayer Network with Cisco STP Toolkit

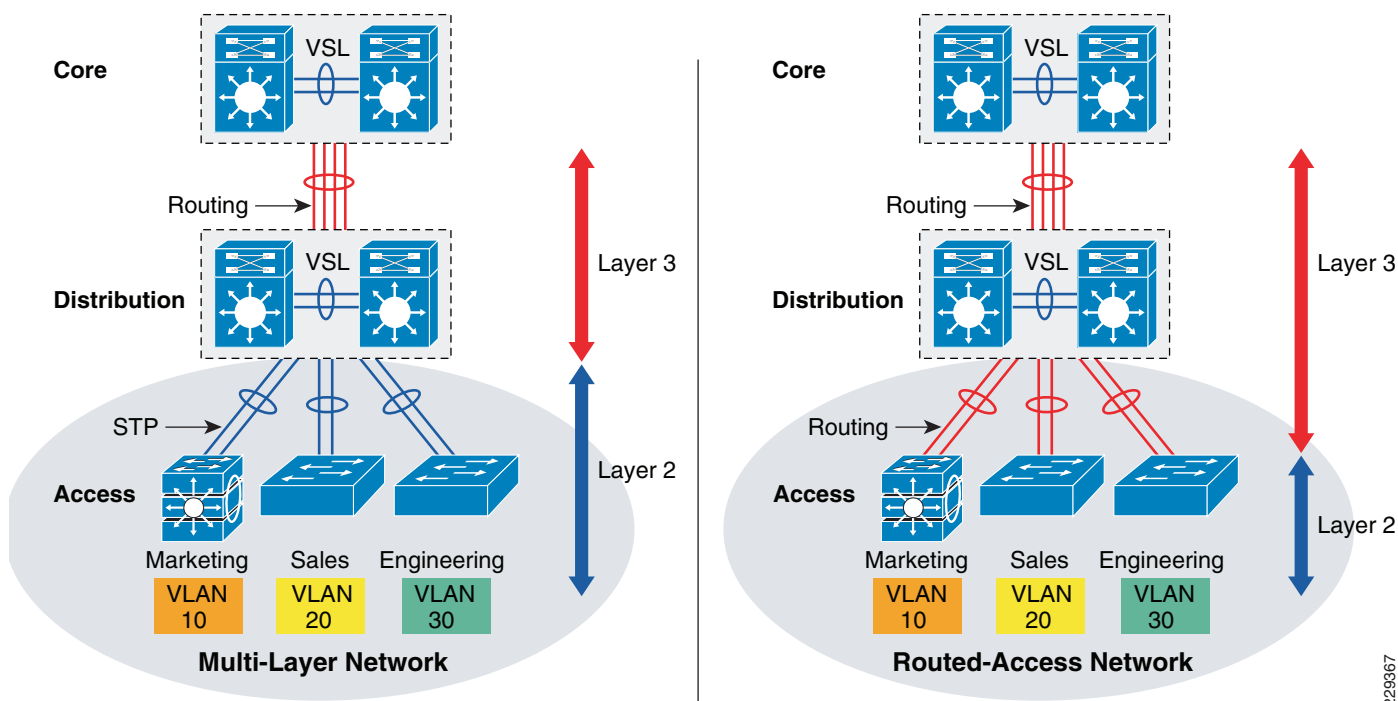
**Note**

For additional STP information, see the following URL:

[http://www.cisco.com/en/US/tech/tk389/tk621/tsd\\_technology\\_support\\_troubleshooting\\_technotes\\_list.html](http://www.cisco.com/en/US/tech/tk389/tk621/tsd_technology_support_troubleshooting_technotes_list.html).

## Designing the Routed Access Network

Routing functions in the access layer network simplify configuration, optimize distribution performances, and provide end-to-end troubleshooting tools. Implementing Layer 3 functions in the access layer replaces Layer 2 trunk configuration to a single point-to-point Layer 3 interface with a collapsed core system in the aggregation layer. Pushing Layer 3 functions one tier down on Layer 3 access switches changes the traditional multilayer network topology and forwarding development path. Implementing Layer 3 functions in the access switch does not require any physical or logical link reconfiguration; the access-distribution block can be used, and is as resilient as in the multilayer network design. [Figure 2-41](#) shows the differences between the multilayer and routed access network designs, as well as where the Layer 2 and Layer 3 boundaries exist in each network design.

**Figure 2-41 Layer 2 and Layer 3 Boundaries for Multilayer and Routed Access Network Design**

Routed-access network design enables Layer 3 access switches to perform Layer 2 demarcation point and provide Inter-VLAN routing and gateway function to the endpoints. The Layer 3 access switches makes more intelligent, multi-function and policy-based routing and switching decision like distribution-layer switches.

Although Cisco VSS and a single redundant distribution design are simplified with a single point-to-point EtherChannel, the benefits in implementing the routed access design in medium enterprises are as follows:

- Eliminates the need for implementing STP and the STP toolkit on the distribution system. As a best practice, the STP toolkit must be hardened at the access layer.
- Shrinks the Layer 2 fault domain, thus minimizing the number of denial-of-service (DoS)/distributed denial-of-service (DDoS) attacks.
- Bandwidth efficiency—Improves Layer 3 uplink network bandwidth efficiency by suppressing Layer 2 broadcasts at the edge port.
- Improves overall collapsed core and distribution resource utilization.

Enabling Layer 3 functions in the access-distribution block must follow the same core network designs as mentioned in previous sections to provide network security as well as optimize the network topology and system resource utilization:

- *EIGRP autonomous system*—Layer 3 access switches must be deployed in the same EIGRP AS as the distribution and core layer systems.
- *EIGRP adjacency protection*—EIGRP processing must be enabled on uplink Layer 3 EtherChannels, and must block remaining Layer 3 ports by default in passive mode. Access switches must establish secured EIGRP adjacency using the MD5 hash algorithm with the aggregation system.

- *EIGRP network boundary*—All EIGRP neighbors must be in a single AS to build a common network topology. The Layer 3 access switches must be deployed in EIGRP stub mode for a concise network view.

## Implementing Routed Access in Access-Distribution Block

Cisco IOS configuration to implement Layer 3 routing function on the Catalyst access-layer switch remains consistent. Refer to EIGRP routing configuration and best practices defined in Designing End-to-End EIGRP Network section to routing function in access-layer switches.

EIGRP creates and maintains a single flat routing topology network between EIGRP peers. Building a single routing domain in a large-scale campus core design allows for complete network visibility and reachability that may interconnect multiple campus components, such as distribution blocks, services blocks, the data center, the WAN edge, and so on.

In the three- or two-tier deployment models, the Layer 3 access switch must always have single physical or logical forwarding to a distribution switch. The Layer 3 access switch dynamically develops the forwarding topology pointing to a single distribution switch as a single Layer 3 next hop. Because the distribution switch provides a gateway function to rest of the network, the routing design on the Layer 3 access switch can be optimized with the following two techniques to improve performance and network reconvergence in the access-distribution block, as shown in [Figure 2-42](#):

- Deploying the Layer 3 access switch in EIGRP stub mode

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

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

- Access layer

```
cr22-4507-LB(config)#router eigrp 100
cr22-4507-LB(config-router)# eigrp stub connected

cr22-4507-LB#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 NSF enabled
    NSF signal timer is 20s
    NSF converge timer is 120s
    Time since last restart is 2w2d
EIGRP stub, connected
  Topologies : 0(base)
```

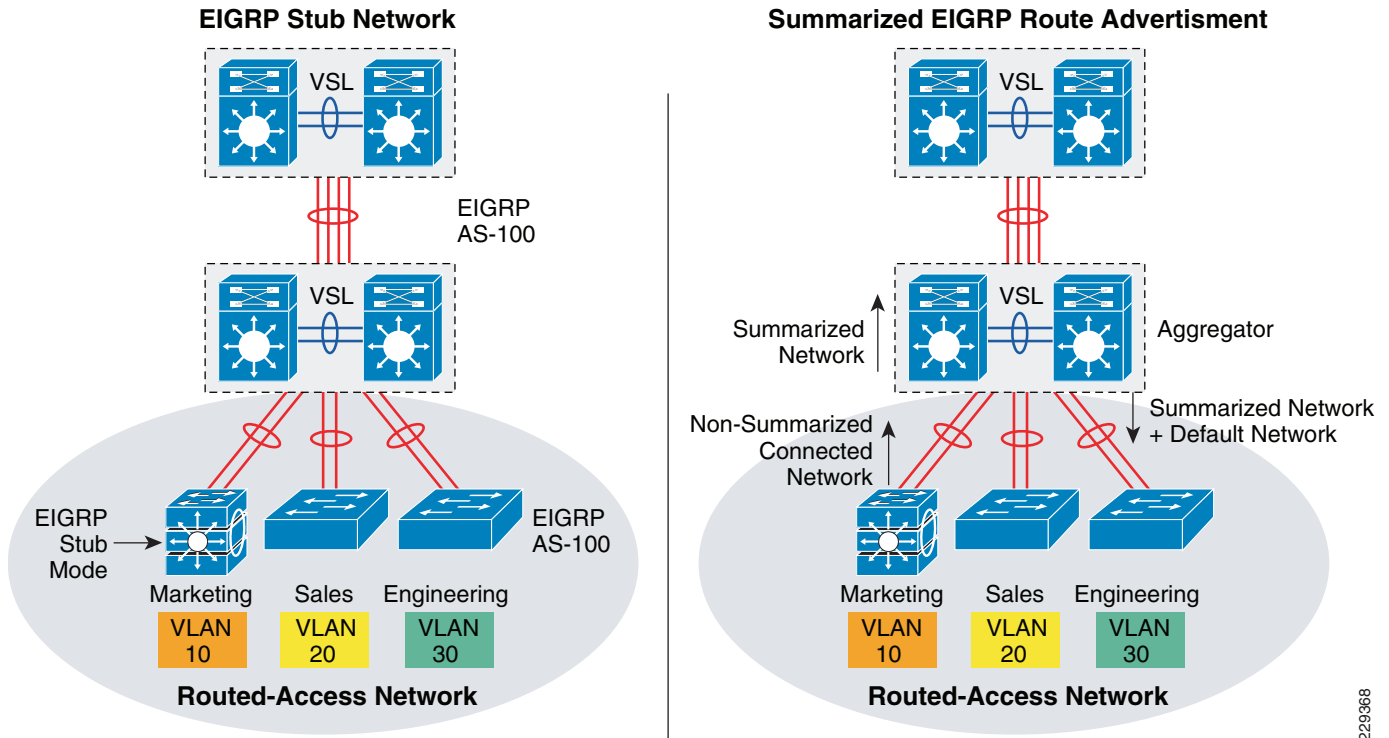
- Distribution layer

```
cr22-6500-LB#show ip eigrp neighbors detail port-channel 101
EIGRP-IPv4 neighbors for process 100
H   Address                Interface          Hold UptimeSRTT   RTO  Q Seq
                               (sec)              (ms)   Cnt Num
2   10.125.0.1              Po101              13 3d18h         4   2000 98
Version 4.0/3.0, Retrans: 0, Retries: 0, Prefixes: 6
Topology-ids from peer - 0
Stub Peer Advertising ( CONNECTED ) Routes
Suppressing queries
```



- Summarizing the network view with a default route to the Layer 3 access switch for intelligent routing functions

**Figure 2-42** *Designing and Optimizing EIGRP Network Boundary for the Access Layer*



The following sample configuration demonstrate the procedure to implement route filtering at the distribution layer that allows summarized and default-route advertisement to build concise network topology at the access layer:

- Distribution layer

```
cr22-6500-LB(config)# ip prefix-list EIGRP_STUB_ROUTES seq 5 permit 0.0.0.0/0
cr22-6500-LB(config)# ip prefix-list EIGRP_STUB_ROUTES seq 10 permit 10.122.0.0/16
cr22-6500-LB(config)# ip prefix-list EIGRP_STUB_ROUTES seq 15 permit 10.123.0.0/16
cr22-6500-LB(config)# ip prefix-list EIGRP_STUB_ROUTES seq 20 permit 10.124.0.0/16
cr22-6500-LB(config)# ip prefix-list EIGRP_STUB_ROUTES seq 25 permit 10.125.0.0/16
cr22-6500-LB(config)# ip prefix-list EIGRP_STUB_ROUTES seq 30 permit 10.126.0.0/16
```

```
cr22-6500-LB(config)#router eigrp 100
cr22-6500-LB(config-router)#distribute-list route-map EIGRP_STUB_ROUTES out
Port-channel101
cr22-6500-LB(config-router)#distribute-list route-map EIGRP_STUB_ROUTES out
Port-channel102
cr22-6500-LB(config-router)#distribute-list route-map EIGRP_STUB_ROUTES out
Port-channel103
```

```
cr22-6500-LB#show ip protocols
Outgoing update filter list for all interfaces is not set
Port-channel101 filtered by
Port-channel102 filtered by
```

**Port-channel103 filtered by**

- Access layer

```
cr22-4507-LB#show ip route eigrp
  10.0.0.0/8 is variably subnetted, 12 subnets, 4 masks
D       10.122.0.0/16 [90/3840] via 10.125.0.0, 07:49:11, Port-channel1
D       10.123.0.0/16 [90/3840] via 10.125.0.0, 01:42:22, Port-channel1
D       10.126.0.0/16 [90/3584] via 10.125.0.0, 07:49:11, Port-channel1
D       10.124.0.0/16 [90/64000] via 10.125.0.0, 07:49:11, Port-channel1
D       10.125.0.0/16 [90/768] via 10.125.0.0, 07:49:13, Port-channel1
D *EX 0.0.0.0/0 [170/515584] via 10.125.0.0, 07:49:13, Port-channel1
```

## Multicast for Application Delivery

Because unicast communication is based on the one-to-one forwarding model, it becomes easier in routing and switching decisions to perform destination address lookup, determine the egress path by scanning forwarding tables, and to switch traffic. In the unicast routing and switching technologies discussed in the previous section, the network may need to be made more efficient by allowing certain applications where the same content or application must be replicated to multiple users. IP multicast delivers source traffic to multiple receivers using the least amount of network resources as possible without placing an additional burden on the source or the receivers. Multicast packet replication in the network is done by Cisco routers and switches enabled with Protocol Independent Multicast (PIM) as well as other multicast routing protocols.

Similar to the unicast methods, multicast requires the following design guidelines:

- Choosing a multicast addressing design
- Choosing a multicast routing protocol
- Providing multicast security regardless of the location within the medium enterprise design

## Multicast Addressing Design

The Internet Assigned Numbers Authority (IANA) controls the assignment of IP multicast addresses. A range of class D address space is assigned to be used for IP multicast applications. All multicast group addresses fall in the range from 224.0.0.0 through 239.255.255.255. Layer 3 addresses in multicast communications operate differently; while the destination address of IP multicast traffic is in the multicast group range, the source IP address is always in the unicast address range. Multicast addresses are assigned in various pools for well-known multicast-based network protocols or inter-domain multicast communications, as listed in [Table 2-5](#).

**Table 2-5 Multicast Address Range Assignments**

Application	Address Range
Reserved—Link local network protocols.	224.0.0.0/24
Global scope—Group communication between an organization and the 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

**Table 2-5 Multicast Address Range Assignments (continued)**

GLOP—Inter-domain multicast group assignment with reserved global AS.	233.0.0.0/8
Limited scope—Administratively scoped address that remains constrained within a local organization or AS. Commonly deployed in enterprise, education, and other organizations.	239.0.0.0/8

During the multicast network design phase, medium enterprise network architects must select a range of multicast sources from the limited scope pool (239/8).

## Multicast Routing Design

To enable end-to-end dynamic multicast operation in the network, each intermediate system between the multicast receiver and source must support the multicast feature. Multicast develops the forwarding table differently than the unicast routing and switching model. To enable communication, multicast requires specific multicast routing protocols and dynamic group membership.

The medium enterprise LAN design must be able to build packet distribution trees that specify a unique forwarding path between the subnet of the source and each subnet containing members of the multicast group. A primary goal in distribution trees construction is to ensure that no more than one copy of each packet is forwarded on each branch of the tree. The two basic types of multicast distribution trees are as follows:

- *Source trees*—The simplest form of a multicast distribution tree is a source tree, with its root at the source and branches forming a tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT).
- *Shared trees*—Unlike source trees that have their root at the source, shared trees use a single common root placed at a selected point in the network. This shared root is called a rendezvous point (RP).

The PIM protocol is divided into the following two modes to support both types of multicast distribution trees:

- *Dense mode (DM)*—Assumes that almost all routers in the network need to distribute multicast traffic for each multicast group (for example, almost all hosts on the network belong to each multicast group). PIM in DM mode builds distribution trees by initially flooding the entire network and then pruning back the small number of paths without receivers.
- *Sparse mode (SM)*—Assumes that relatively few routers in the network are involved in each multicast. The hosts belonging to the group are widely dispersed, as might be the case for most multicasts over the WAN. Therefore, PIM-SM begins with an empty distribution tree and adds branches only as the result of explicit Internet Group Management Protocol (IGMP) requests to join the distribution. PIM-SM mode is ideal for a network without dense receivers and multicast transport over WAN environments, and it adjusts its behavior to match the characteristics of each receiver group.

Selecting the PIM mode depends on the multicast applications that use various mechanisms to build multicast distribution trees. Based on the multicast scale factor and centralized source deployment design for one-to-many multicast communication in medium enterprise LAN infrastructures, Cisco recommends deploying PIM-SM because it is efficient and intelligent in building multicast distribution tree. All the recommended platforms in this design support PIM-SM mode on physical or logical (switched virtual interface [SVI] and EtherChannel) interfaces.

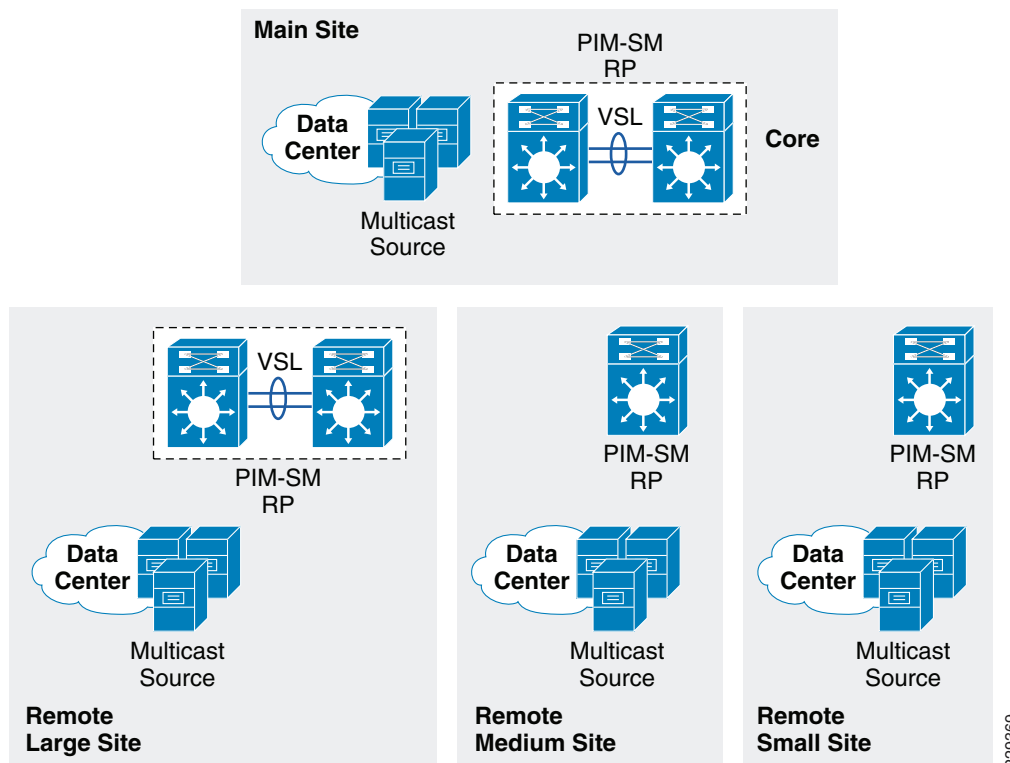
## Designing PIM Rendezvous Point

The following sections discuss best practices in designing and deploying the PIM-SM Rendezvous Point.

### PIM-SM RP Placement

It is assumed that each medium enterprise site has a wide range of local multicast sources in the data center for distributed medium enterprise IT-managed media and employee research and development applications. In such a distributed multicast network design, Cisco recommends deploying PIM RP on each site for wired or wireless multicast receivers and sources to join and register at the closest RP. The Medium Enterprise Reference design recommends PIM-SM RP placement on a VSS-enabled and single resilient core system in the three-tier campus design, and on the collapsed core/distribution system in the two-tier campus design model. See [Figure 2-43](#).

**Figure 2-43** Distributed PIM-SM RP Placement



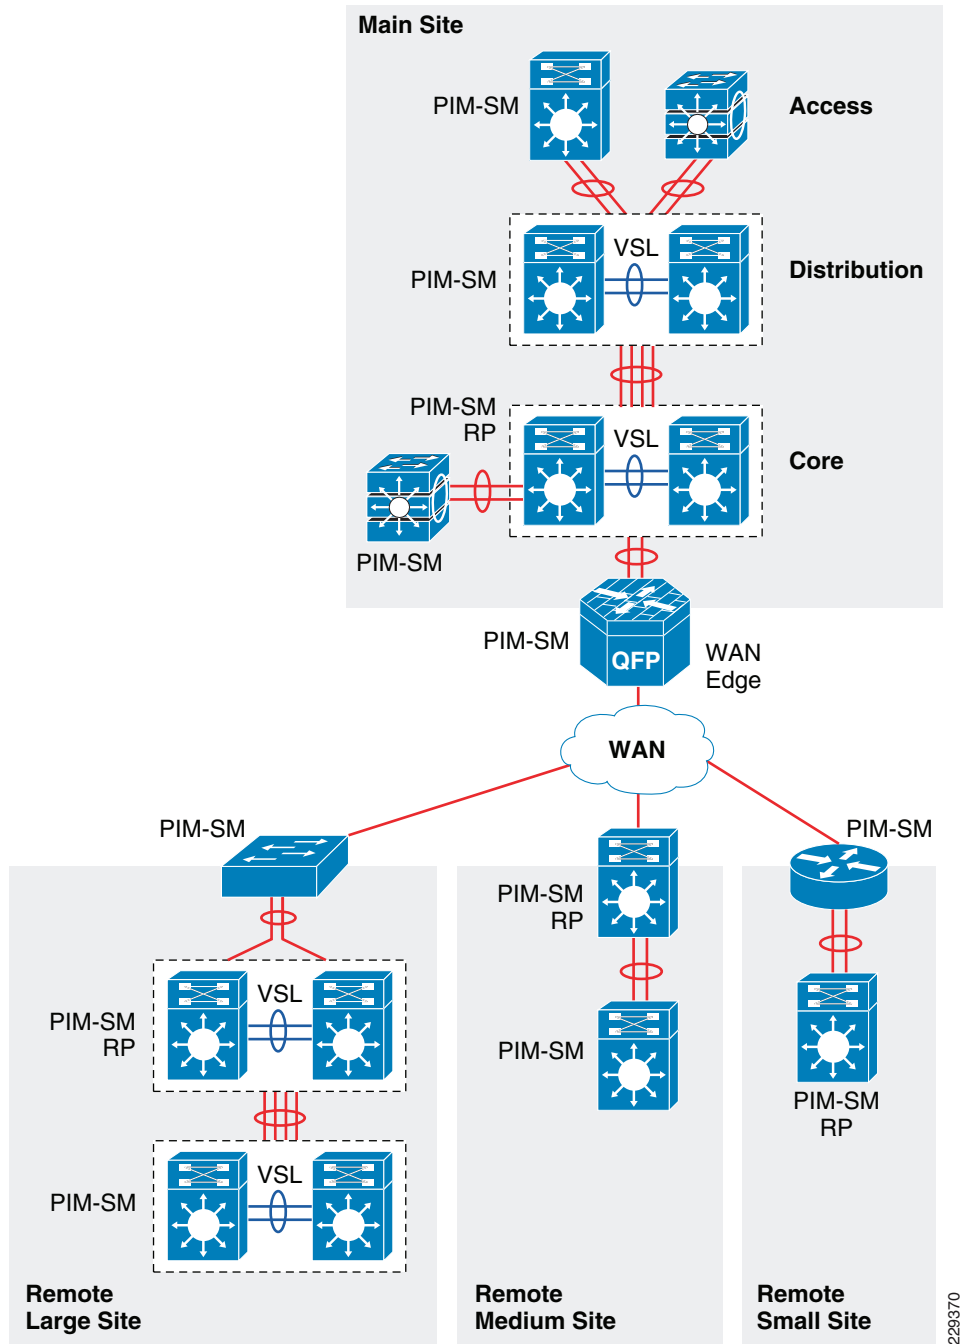
### PIM-SM RP Mode

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

- *Static*—In this mode, 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*—This mode is a dynamic method for discovering and announcing the 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 a change, the RP mapping agent router must be designated in the network to receive RP group announcements and to arbitrate conflicts, as part of the PIM version 1 specification.

- *Bootstrap Router (BSR)*—This mode performs the same tasks as Auto-RP but in a different way, and is part of the PIM version 2 specification. Auto-RP and BSR cannot co-exist or interoperate in the same network.

In a small- to mid-sized multicast network, static RP configuration is recommended over the other modes. Static RP implementation offers RP redundancy and load sharing, and an additional simple access control list (ACL) can be applied to deploy RP without compromising multicast network security. Cisco recommends designing the medium enterprise LAN multicast network using the static PIM-SM mode configuration. See [Figure 2-44](#).

**Figure 2-44 PIM-SM Network Design in Medium Enterprise Network**

The following is an example configuration to deploy PIM-SM RP on all PIM-SM running systems. To provide transparent PIM-SM redundancy, static PIM-SM RP configuration must be identical across the campus LAN network and on each PIM-SM RP routers.

- Core layer

```
cr23-VSS-Core(config)#ip multicast-routing
```

```
cr23-VSS-Core(config)#interface Loopback100
```

```

cr23-VSS-Core(config-if)#description Anycast RP Loopback
cr23-VSS-Core(config-if)#ip address 10.100.100.100 255.255.255.255

cr23-VSS-Core(config)#ip pim rp-address 10.100.100.100

cr23-VSS-Core#show ip pim rp

Group: 239.192.51.1, RP: 10.100.100.100, next RP-reachable in 00:00:34
Group: 239.192.51.2, RP: 10.100.100.100, next RP-reachable in 00:00:34
Group: 239.192.51.3, RP: 10.100.100.100, next RP-reachable in 00:00:34

cr23-VSS-Core#show ip pim interface

Address          Interface          Ver/  Nbr   Query  DR    DR
                  Mode    Count  Intvl Prior
10.125.0.12      Port-channel101    v2/S   1      30     1
10.125.0.13
10.125.0.14      Port-channel102    v2/S   1      30     1
10.125.0.15
...

cr23-VSS-Core#show ip mroute sparse
(*, 239.192.51.8), 3d22h/00:03:20, RP 10.100.100.100, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Port-channel105, Forward/Sparse, 00:16:54/00:02:54
    Port-channel101, Forward/Sparse, 00:16:56/00:03:20

(10.125.31.147, 239.192.51.8), 00:16:54/00:02:35, flags: A
  Incoming interface: Port-channel105, RPF nbr 10.125.0.21
  Outgoing interface list:
    Port-channel101, Forward/Sparse, 00:16:54/00:03:20

cr23-VSS-Core#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps

Group: 239.192.51.1, (?)
  Source: 10.125.31.153 (?)
    Rate: 2500 pps/4240 kbps(1sec), 4239 kbps(last 30 secs), 12 kbps(life avg)

```

- Distribution layer

```

cr23-6500-LB(config)#ip multicast-routing
cr23-6500-LB(config)#ip pim rp-address 10.100.100.100

cr23-6500-LB(config)#interface range Port-channel 100 - 103
cr22-6500-LB(config-if-range)#ip pim sparse-mode

cr23-6500-LB(config)#interface range Vlan 101 - 120
cr22-6500-LB(config-if-range)#ip pim sparse-mode

cr22-6500-LB#show ip pim rp
Group: 239.192.51.1, RP: 10.100.100.100, uptime 00:10:42, expires never
Group: 239.192.51.2, RP: 10.100.100.100, uptime 00:10:42, expires never
Group: 239.192.51.3, RP: 10.100.100.100, uptime 00:10:41, expires never
Group: 224.0.1.40, RP: 10.100.100.100, uptime 3d22h, expires never

cr22-6500-LB#show ip pim interface

Address          Interface          Ver/  Nbr   Query  DR    DR
                  Mode    Count  Intvl Prior
10.125.0.13Port-channel100v2/S   1      30     1      10.125.0.13

```

```

10.125.0.0Port-channel101v2/S 1 30 1 10.125.0.1
...
10.125.103.129Vlan101v2/S 0 30 1 10.125.103.129
...

cr22-6500-LB#show ip mroute sparse
(*, 239.192.51.1), 00:14:23/00:03:21, RP 10.100.100.100, flags: SC
Incoming interface: Port-channel100, RPF nbr 10.125.0.12, RPF-MFD
Outgoing interface list:
Port-channel102, Forward/Sparse, 00:13:27/00:03:06, H
Vlan120, Forward/Sparse, 00:14:02/00:02:13, H
Port-channel101, Forward/Sparse, 00:14:20/00:02:55, H
Port-channel103, Forward/Sparse, 00:14:23/00:03:10, H
Vlan110, Forward/Sparse, 00:14:23/00:02:17, H

cr22-6500-LB#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps

Group: 239.192.51.1, (?)
RP-tree:
Rate: 2500 pps/4240 kbps(1sec), 4240 kbps(last 10 secs), 4011 kbps(life avg)

```

- Access layer

```

cr23-3560-LB(config)#ip multicast-routing distributed
cr23-3560-LB(config)#ip pim rp-address 10.100.100.100

cr23-3560-LB(config)#interface range Vlan 101 - 110
cr22-3560-LB(config-if-range)#ip pim sparse-mode

cr22-3560-LB#show ip pim rp
Group: 239.192.51.1, RP: 10.100.100.100, uptime 00:01:36, expires never
Group: 239.192.51.2, RP: 10.100.100.100, uptime 00:01:36, expires never
Group: 239.192.51.3, RP: 10.100.100.100, uptime 00:01:36, expires never
Group: 224.0.1.40, RP: 10.100.100.100, uptime 5w5d, expires never
cr22-3560-LB#show ip pim interface

Address          Interface          Ver/  Nbr   Query  DR    DR
                  Mode    Count  Intvl Prior
10.125.0.5        Port-channel1      v2/S   1     30     1    10.125.0.5
10.125.101.1      Vlan101            v2/S   0     30     1     0.0.0.0
...
10.125.103.65    Vlan110            v2/S   0     30     1    10.125.103.65

cr22-3560-LB#show ip mroute sparse
(*, 239.192.51.1), 00:06:06/00:02:59, RP 10.100.100.100, flags: SC
Incoming interface: Port-channel1, RPF nbr 10.125.0.4
Outgoing interface list:
Vlan101, Forward/Sparse, 00:06:08/00:02:09
Vlan110, Forward/Sparse, 00:06:06/00:02:05

```

- WAN edge layer

```

cr11-asr-we(config)#ip multicast-routing distributed

cr11-asr-we(config)#ip pim rp-address 10.100.100.100

cr11-asr-we(config)#interface range Port-channel1 , Gig0/2/0 , Gig0/2/1.102
cr11-asr-we(config-if-range)#ip pim sparse-mode
cr11-asr-we(config)#interface Ser0/3/0
cr11-asr-we(config-if)#ip pim sparse-mode

```



```

cr11-asr-we#show ip pim rp
Group: 239.192.57.1, RP: 10.100.100.100, uptime 00:23:16, expires never
Group: 239.192.57.2, RP: 10.100.100.100, uptime 00:23:16, expires never
Group: 239.192.57.3, RP: 10.100.100.100, uptime 00:23:16, expires never

cr11-asr-we#show ip mroute sparse

(*, 239.192.57.1), 00:24:08/stopped, RP 10.100.100.100, flags: SP
  Incoming interface: Port-channel11, RPF nbr 10.125.0.22
  Outgoing interface list: Null

(10.125.31.156, 239.192.57.1), 00:24:08/00:03:07, flags: T
  Incoming interface: Port-channel11, RPF nbr 10.125.0.22
  Outgoing interface list:
    Serial0/3/0, Forward/Sparse, 00:24:08/00:02:55

cr11-asr-we#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps

Group: 239.192.57.1, (?)
  Source: 10.125.31.156 (?)
    Rate: 625 pps/1130 kbps(1sec), 1130 kbps(last 40 secs), 872 kbps(life avg)

```

## PIM-SM RP Redundancy

PIM-SM RP redundancy and load sharing becomes imperative in the medium enterprise LAN design, because each recommended core layer design model provides resiliency and simplicity. In the Cisco Catalyst 6500 VSS-enabled core layer, the dynamically discovered group-to-RP entries are fully synchronized to the standby switch. Combining NSF/SSO capabilities with IPv4 multicast reduces the network recovery time and retains the user and application performance at an optimal level. In the non-VSS-enabled network design, PIM-SM uses Anycast RP and Multicast Source Discovery Protocol (MSDP) for node failure protection. PIM-SM redundancy and load sharing is simplified with the Cisco VSS-enabled core. Because VSS is logically a single system and provides node protection, there is no need to implement Anycast RP and MSDP on a VSS-enabled PIM-SM RP.

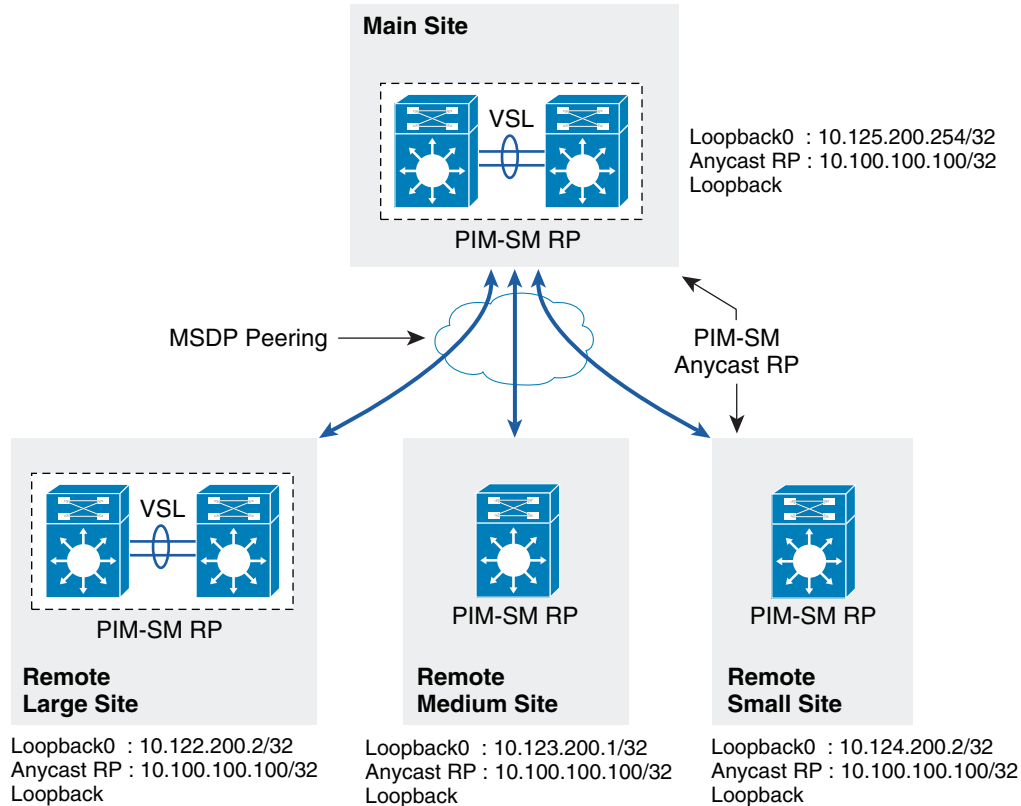
## Inter-Site PIM Anycast RP

MSDP allows PIM RPs to share information about the active sources. PIM-SM RPs discover local receivers through PIM join messages, while the multicast source can be in a local or remote network domain. MSDP allows each multicast domain to maintain an independent RP that does not rely on other multicast domains, but does enable RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.

Anycast RP is a useful application of MSDP. Originally developed for interdomain multicast applications, MSDP used with Anycast RP is an intradomain feature that provides redundancy and load sharing capabilities. Large networks typically use Anycast RP for configuring a PIM-SM network to meet fault tolerance requirements within a single multicast domain.

The medium enterprise LAN multicast network must be designed with Anycast RP. PIM-SM RP at the main or the centralized core must establish an MSDP session with RP on each remote site to exchange distributed multicast source information and allow RPs to join SPT to active sources as needed.

Figure 2-45 shows an example of a medium enterprise LAN multicast network design.

**Figure 2-45 Medium Enterprise Inter-Site Multicast Network Design**

229371

## Implementing MSDP Anycast RP

### Main Campus

```
cr23-VSS-Core(config)#ip msdp peer 10.122.200.2 connect-source Loopback0
cr23-VSS-Core(config)#ip msdp description 10.122.200.2 ANYCAST-PEER-6k-RemoteLrgCampus
cr23-VSS-Core(config)#ip msdp peer 10.123.200.1 connect-source Loopback0
cr23-VSS-Core(config)#ip msdp description 10.123.200.1 ANYCAST-PEER-4k-RemoteMedCampus
cr23-VSS-Core(config)#ip msdp peer 10.124.200.2 connect-source Loopback0
cr23-VSS-Core(config)#ip msdp description 10.124.200.2 ANYCAST-PEER-4k-RemoteSmlCampus
cr23-VSS-Core(config)#ip msdp cache-sa-state
cr23-VSS-Core(config)#ip msdp originator-id Loopback0

cr23-VSS-Core#show ip msdp peer | inc MSDP Peer|State
MSDP Peer 10.122.200.2 (?), AS ?
    State: Up, Resets: 0, Connection source: Loopback0 (10.125.200.254)
MSDP Peer 10.123.200.1 (?), AS ?
    State: Up, Resets: 0, Connection source: Loopback0 (10.125.200.254)
MSDP Peer 10.124.200.2 (?), AS ?
    State: Up, Resets: 0, Connection source: Loopback0 (10.125.200.254)
```

### Remote Large Campus

```
cr14-6500-RLC(config)#ip msdp peer 10.125.200.254 connect-source Loopback0
cr14-6500-RLC(config)#ip msdp description 10.125.200.254 ANYCAST-PEER-6k-MainCampus
cr14-6500-RLC(config)#ip msdp cache-sa-state
```

```
cr14-6500-RLC(config)#ip msdp originator-id Loopback0

cr14-6500-RLC#show ip msdp peer | inc MSDP Peer|State|SAs learned
MSDP Peer 10.125.200.254 (?), AS ?
State: Up, Resets: 0, Connection source: Loopback0 (10.122.200.2)
SAs learned from this peer: 94
```

### Remote Medium Campus

```
cr11-4507-RMC(config)#ip msdp peer 10.125.200.254 connect-source Loopback0
cr11-4507-RMC(config)#ip msdp description 10.125.200.254 ANYCAST-PEER-6k-MainCampus
cr11-4507-RMC(config)#ip msdp cache-sa-state
cr11-4507-RMC(config)#ip msdp originator-id Loopback0

cr11-4507-RMC#show ip msdp peer | inc MSDP Peer|State|SAs learned
MSDP Peer 10.125.200.254 (?), AS ?
State: Up, Resets: 0, Connection source: Loopback0 (10.123.200.1)
SAs learned from this peer: 94
```

### Remote Small Campus

```
cr14-4507-RSC(config)#ip msdp peer 10.125.200.254 connect-source Loopback0
cr14-4507-RSC(config)#ip msdp description 10.125.200.254 ANYCAST-PEER-6k-MainCampus
cr14-4507-RSC(config)#ip msdp cache-sa-state
cr14-4507-RSC(config)#ip msdp originator-id Loopback0

cr14-4507-RSC#show ip msdp peer | inc MSDP Peer|State|SAs learned
MSDP Peer 10.125.200.254 (?), AS ?
State: Up, Resets: 0, Connection source: Loopback0 (10.124.200.2)
SAs learned from this peer: 94
```

## Dynamic Group Membership

Multicast receiver registration is done via IGMP protocol signaling. IGMP is an integrated component of an IP multicast framework that allows the receiver hosts and transmitting sources to be dynamically added to and removed from the network. Without IGMP, the network is forced to flood rather than multicast the transmissions for each group. IGMP operates between a multicast receiver host in the access layer and the Layer 3 router at the distribution layer.

The multicast system role changes when the access layer is deployed in the multilayer and routed access models. Because multilayer access switches do not run PIM, it becomes complex to make forwarding decisions out of the receiver port. In such a situation, Layer 2 access switches flood the traffic on all ports. This multilayer limitation in access switches is solved by using the IGMP snooping feature, which is enabled by default and is recommended to not be disabled.

IGMP is still required when a Layer 3 access layer switch is deployed in the routed access network design. Because the Layer 3 boundary is pushed down to the access layer, IGMP communication is limited between a receiver host and the Layer 3 access switch. In addition to the unicast routing protocol, PIM-SM must be enabled at the Layer 3 access switch to communicate with RPs in the network.

### Implementing IGMP

By default, the Layer-2 access-switch dynamically detects IGMP hosts and multicast-capable Layer-3 PIM 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.

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:

The IGMP configuration can be validated using the following **show** command on the Layer-2 and Layer-3 access-switch:

#### Layer 2 Access

```
cr22-3750-LB#show ip igmp snooping groups
Vlan      Group                Type      Version  Port List
-----
110       239.192.51.1          igmp      v2       Gi1/0/20, Po1
110       239.192.51.2          igmp      v2       Gi1/0/20, Po1
110       239.192.51.3          igmp      v2       Gi1/0/20, Po1

cr22-3750-LB#show ip igmp snooping mrouter
Vlan      ports
-----
110       Po1 (dynamic)
```

#### Layer 3 Access

```
cr22-3560-LB#show ip igmp membership
Channel/Group      Reporter      Uptime  Exp.  Flags  Interface
*,239.192.51.1      10.125.103.106 00:52:36 02:09 2A     Vl1110
*,239.192.51.2      10.125.103.107 00:52:36 02:12 2A     Vl1110
*,239.192.51.3      10.125.103.109 00:52:35 02:16 2A     Vl1110
*,224.0.1.40        10.125.0.4     3d22h   02:04 2A     Po1
*,224.0.1.40        10.125.101.129 4w4d    02:33 2LA    Vl1103

cr22-3560-LB#show ip igmp snooping mrouter
Vlan      ports
-----
103       Router
106       Router
110       Router
```

## Designing Multicast Security

When designing multicast security in the medium enterprise LAN design, two key concerns are preventing a rogue source and preventing a rogue PIM-RP.

### Preventing Rogue Source

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 the (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), 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 with the **pim accept-register** command on the RP, the PIM-SM (S,G) state is still created on the first-hop router of the source. This can result in traffic reaching receivers local to the source and located between the source and the RP. Furthermore, because the **pim accept-register** command works on the control plane of the RP, this can be used to overload the RP with fake register messages and possibly cause a DoS condition.

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

```
cr23-VSS-Core(config)#ip access-list extended PERMIT-SOURCES
cr23-VSS-Core(config-ext-nacl)# permit ip 10.120.31.0 0.7.0.255 239.192.0.0 0.0.255.255
cr23-VSS-Core(config-ext-nacl)# deny ip any any

cr23-VSS-Core(config)#ip pim accept-register list PERMIT-SOURCES
```

## Preventing Rogue PIM-RP

Like the multicast source, any router can be misconfigured or can 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 static RP for the multicast source and override any other Auto-RP or BSR multicast router announcement from the network.

The following is the sample configuration that must be applied to each PIM-enabled router in the campus network, to accept PIM announcements only from the static RP and ignore dynamic multicast group announcement from any other RP:

```
cr23-VSS-Core(config)#ip access-list standard Allowed_MCAST_Groups
cr23-VSS-Core(config-std-nacl)# permit 224.0.1.39
cr23-VSS-Core(config-std-nacl)# permit 224.0.1.40
cr23-VSS-Core(config-std-nacl)# permit 239.192.0.0 0.0.255.255
cr23-VSS-Core(config-std-nacl)# deny any

cr23-VSS-Core(config)#ip pim rp-address 10.100.100.100 Allowed_MCAST_Groups override
```

## QoS for Application Performance Optimization

The function and guaranteed low latency bandwidth expectation of network users and endpoints has evolved significantly over the past few years. Application and device awareness has become a key tool in providing differentiated service treatment at the campus LAN edge. Media applications, and particularly video-oriented media applications, are evolving as the enterprise networks enters the digital era of doing business, as well as the increased campus network and asset security requirements. Integrating video applications in the medium enterprise LAN network exponentially increases bandwidth utilization and fundamentally shifts traffic patterns. Business drivers behind this media application growth include remote learning, as well as leveraging the network as a platform to build an energy-efficient network to minimize cost and go "green". High-definition media is transitioning from the desktop to conference rooms, and social networking phenomena are crossing over into enterprise settings. Besides internal and enterprise research applications, media applications are fueling a new wave of IP convergence, requiring the ongoing development of converged network designs.

Converging media applications onto an IP network is much more complex than converging voice over IP (VoIP) alone. Media applications are generally bandwidth-intensive and bursty (as compared to VoIP), and many different types of media applications exist; in addition to IP telephony, applications can include live and on-demand streaming media applications, digital signage applications, high-definition room-based conferencing applications, as well as an infinite array of data-oriented applications. By

embracing media applications as the next cycle of convergence, medium enterprise IT departments can think holistically about their network design and its readiness to support the coming tidal wave of media applications, and develop a network-wide strategy to ensure high quality end-user experiences.

The medium enterprise LAN infrastructure must set the administrative policies to provide differentiated forwarding services to the network applications, users and endpoints to prevent contention. The characteristic of network services and applications must be well understood, so that policies can be defined that allow network resources to be used for internal applications, to provide best-effort services for external traffic, and to keep the network protected from threats.

The policy for providing network resources to an internal application is further complicated when interactive video and real-time VoIP applications are converged over the same network that is switching mid-to-low priority data traffic. Deploying QoS technologies in the campus allows different types of traffic to contend inequitably for network resources. Real-time applications such as voice, interactive, and physical security video can be given priority or preferential services over generic data applications, but not to the point that data applications are starving for bandwidth.

## Medium Enterprise LAN QoS Framework

Each group of managed and un-managed applications with unique traffic patterns and service level requirements requires a dedicated QoS class to provision and guarantee these service level requirements. The medium enterprise LAN network architect may need to determine the number of classes for various applications, as well as how should these individual classes should be implemented to deliver differentiated services consistently in main and remote campus sites. Cisco recommends following relevant industry standards and guidelines whenever possible, to extend the effectiveness of your QoS policies beyond your direct administrative control.

With minor changes, the medium enterprise LAN QoS framework is developed based on RFC4594 that follows industry standard and guidelines to function consistently in heterogeneous network environment. These guidelines are to be viewed as industry best-practice recommendations. Enterprise organizations and service providers are encouraged to adopt these marking and provisioning recommendations, with the aim of improving QoS consistency, compatibility, and interoperability. However, because these guidelines are not standards, modifications can be made to these recommendations as specific needs or constraints require. To this end, to meet specific business requirements, Cisco has made a minor modification to its adoption of RFC 4594, namely the switching of call-signaling and broadcast video markings (to CS3 and CS5, respectively).

RFC 4594 outlines twelve classes of media applications that have unique service level requirements, as shown in [Figure 2-46](#).

**Figure 2-46 Campus 12-Class QoS Policy Recommendation**

Application Class	Media Application Examples	PHB	Admission Control	Queuing and Dropping
VoIP Telephony	Cisco IP Phone	EF	Required	Priority Queue (PQ)
Broadcast Video	Cisco IPVS, Enterprise TV	CS5	Required	(Optional) PQ
Real-Time Interactive	Cisco TelePresence	CS4	Required	(Optional) PQ
Multimedia Conferencing	Cisco CUPC, WebEx	AF4	Required	BW Queue + DSCP WRED
Multimedia Streaming	Cisco DMS, IP/TV	AF3	Recommended	BW Queue + DSCP WRED
Network Control	EIGRP, OSPF, HSRP, IKE	CS6		BW Queue
Call-Signaling	SCCP, SIP, H.323	CS3		BW Queue
Ops/Admin/Mgmt (OAM)	SNMP, SSH, Syslog	CS2		BW Queue
Transactional Data	ERP Apps, CRM Apps	AF2		BW Queue + DSCP WRED
Bulk Data	E-mail, FTP, Backup	AF1		BW Queue + DSCP WRED
Best Effort	Default Class	DF		Default Queue + RED
Scavenger	YouTube, Gaming, P2P	CS1		Min BW Queue

228497

The twelve classes are as follows:

- *VoIP telephony*—This service class is intended for VoIP telephony (bearer-only) traffic (VoIP signaling traffic is assigned to the call-signaling class). Traffic assigned to this class should be marked EF. This class is provisioned with expedited forwarding (EF) per-hop behavior (PHB). The EF PHB-defined in RFC 3246 is a strict-priority queuing service and, as such, admission to this class should be controlled (admission control is discussed in the following section). Examples of this type of traffic include G.711 and G.729a.
- *Broadcast video*—This service class is intended for broadcast TV, live events, video surveillance flows, and similar *inelastic* streaming video flows, which are highly drop sensitive and have no retransmission and/or flow control capabilities. Traffic in this class should be marked class selector 5 (CS5) and may be provisioned with an EF PHB; as such, admission to this class should be controlled. Examples of this traffic include live Cisco Digital Media System (DMS) streams to desktops or to Cisco Digital Media Players (DMPs), live Cisco Enterprise TV (ETV) streams, and Cisco IP Video Surveillance.
- *Real-time interactive*—This service class is intended for (inelastic) room-based, high-definition interactive video applications and is intended primarily for voice and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the transactional data traffic class. Traffic in this class should be marked CS4 and may be provisioned with an EF PHB; as such, admission to this class should be controlled. A sample application is Cisco TelePresence.
- *Multimedia conferencing*—This service class is intended for desktop software multimedia collaboration applications and is intended primarily for voice and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the transactional data traffic class. Traffic in this class should be marked assured forwarding (AF) Class 4 (AF4) and should be provisioned with a guaranteed bandwidth queue with Differentiated Services Code Point (DSCP)-based Weighted Random Early Detection (WRED) enabled. Admission to this class should be controlled;

additionally, traffic in this class may be subject to policing and re-marking. Sample applications include Cisco Unified Personal Communicator, Cisco Unified Video Advantage, and the Cisco Unified IP Phone 7985G.

- *Multimedia streaming*—This service class is intended for video-on-demand (VoD) streaming video flows, which, in general, are more elastic than broadcast/live streaming flows. Traffic in this class should be marked AF Class 3 (AF31) and should be provisioned with a guaranteed bandwidth queue with DSCP-based WRED enabled. Admission control is recommended on this traffic class (though not strictly required) and this class may be subject to policing and re-marking. Sample applications include Cisco Digital Media System VoD streams.
- *Network control*—This service class is intended for network control plane traffic, which is required for reliable operation of the enterprise network. Traffic in this class should be marked CS6 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because network control traffic should not be dropped (if this class is experiencing drops, the bandwidth allocated to it should be re-provisioned). Sample traffic includes EIGRP, OSPF, Border Gateway Protocol (BGP), HSRP, Internet Key Exchange (IKE), and so on.
- *Call-signaling*—This service class is intended for signaling traffic that supports IP voice and video telephony. Traffic in this class should be marked CS3 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because call-signaling traffic should not be dropped (if this class is experiencing drops, the bandwidth allocated to it should be re-provisioned). Sample traffic includes Skinny Call Control Protocol (SCCP), Session Initiation Protocol (SIP), H.323, and so on.
- *Operations/administration/management (OAM)*—This service class is intended for network operations, administration, and management traffic. This class is critical to the ongoing maintenance and support of the network. Traffic in this class should be marked CS2 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because OAM traffic should not be dropped (if this class is experiencing drops, the bandwidth allocated to it should be re-provisioned). Sample traffic includes Secure Shell (SSH), Simple Network Management Protocol (SNMP), Syslog, and so on.
- *Transactional data (or low-latency data)*—This service class is intended for interactive, “foreground” data applications (foreground refers to applications from which users are expecting a response via the network to continue with their tasks; excessive latency directly impacts user productivity). Traffic in this class should be marked AF Class 2 (AF21) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Sample applications include data components of multimedia collaboration applications, Enterprise Resource Planning (ERP) applications, Customer Relationship Management (CRM) applications, database applications, and so on.
- *Bulk data (or high-throughput data)*—This service class is intended for non-interactive “background” data applications (background refers to applications from which users are not awaiting a response via the network to continue with their tasks; excessive latency in response times of background applications does not directly impact user productivity). Traffic in this class should be marked AF Class 1 (AF11) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Sample applications include E-mail, backup operations, FTP/SFTP transfers, video and content distribution, and so on.
- *Best effort (or default class)*—This service class is the default class. The vast majority of applications will continue to default to this best-effort service class; as such, this default class should be adequately provisioned. Traffic in this class is marked default forwarding (DF or DSCP 0) and should be provisioned with a dedicated queue. WRED is recommended to be enabled on this class.



- *Scavenger (or low-priority data)*—This service class is intended for non-business-related traffic flows, such as data or video applications that are entertainment and/or gaming-oriented. The approach of a less-than Best-Effort service class for non-business applications (as opposed to shutting these down entirely) has proven to be a popular, political compromise. These applications are permitted on enterprise networks, as long as resources are always available for business-critical voice, video, and data applications. However, as soon as the network experiences congestion, this class is the first to be penalized and aggressively dropped. Traffic in this class should be marked CS1 and should be provisioned with a minimal bandwidth queue that is the first to starve should network congestion occur. Sample traffic includes YouTube, Xbox Live/360 movies, iTunes, BitTorrent, and so on.

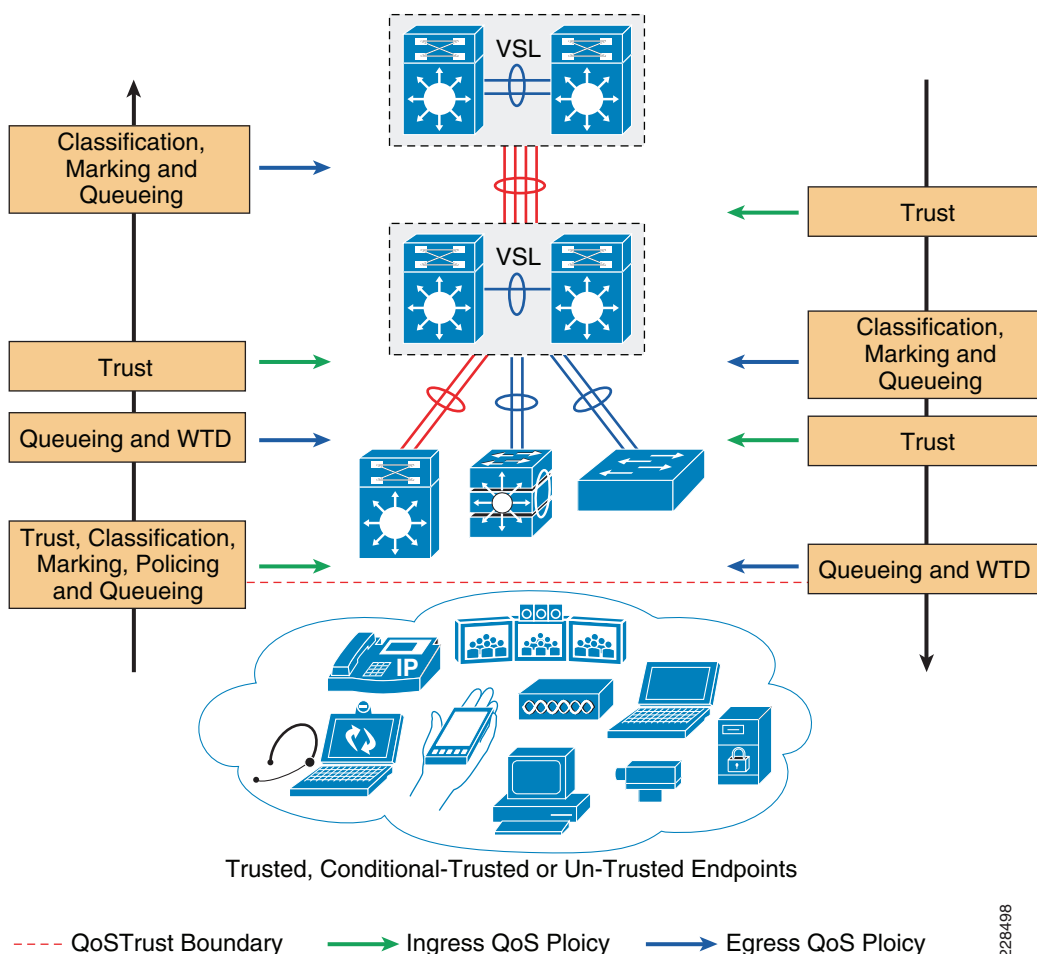
## Designing Medium Enterprise LAN QoS Trust Boundary and Policies

To build an end-to-end QoS framework that offers transparent and consistent QoS service without compromising performance, it is important to create a blueprint of the network, classifying a set of trusted applications, devices, and forwarding paths; and then define common QoS policy settings independent of how QoS is implemented within the system.

QoS settings applied at the LAN network edge sets the ingress rule based on deep packet classification and marks the traffic before it is forwarded inside the campus core. To retain the marking set by access layer switches, it is important that other LAN network devices in the campus trust the marking and apply the same policy to retain the QoS settings and offer symmetric treatment. Bi-directional network communication between applications, endpoints, or other network devices requires the same treatment when traffic enters or leaves the network, and must be taken into account when designing the trust model between network endpoints and core and edge campus devices.

The trust or un-trust model simplifies the rules for defining bi-directional QoS policy settings.

[Figure 2-47](#) shows the QoS trust model setting that sets the QoS implementation guidelines in medium enterprise campus networks.

**Figure 2-47 Campus LAN QoS Trust and Policies**

228498

## Medium Enterprise LAN QoS Overview

With an overall application strategy in place, end-to-end QoS policies can be designed for each device and interface, as determined by their roles in the network infrastructure. However, because the Cisco QoS toolset provides many QoS design and deployment options, a few succinct design principles can help simplify strategic QoS deployments, as discussed in the following sections.

### Hardware versus Software QoS

A fundamental QoS design principle is to always enable QoS policies in hardware rather than software whenever possible. Cisco IOS routers perform QoS in software, which places incremental loads on the CPU, depending on the complexity and functionality of the policy. Cisco Catalyst switches, on the other hand, perform QoS in dedicated hardware application-specific integrated circuits (ASICs) on Ethernet-based ports, and as such do not tax their main CPUs to administer QoS policies. This allows complex policies to be applied at line rates even up to Gigabit or 10-Gigabit speeds.

## Classification and Marking

When classifying and marking traffic, a recommended design principle is to classify and mark applications as close to their sources as technically and administratively feasible. This principle promotes end-to-end differentiated services and PHBs.

In general, it is not recommended to trust markings that can be set by users on their PCs or other similar devices, because users can easily abuse provisioned QoS policies if permitted to mark their own traffic. For example, if an EF PHB has been provisioned over the network, a PC user can easily configure all their traffic to be marked to EF, thus hijacking network priority queues to service non-realtime traffic. Such abuse can easily ruin the service quality of realtime applications throughout the campus. On the other hand, if medium enterprise network administrator controls are in place that centrally administer PC QoS markings, it may be possible and advantageous to trust these.

Following this rule, it is recommended to use DSCP markings whenever possible, because these are end-to-end, more granular, and more extensible than Layer 2 markings. Layer 2 markings are lost when the media changes (such as a LAN-to-WAN/VPN edge). There is also less marking granularity at Layer 2. For example, 802.1P supports only three bits (values 0-7), as does Multiprotocol Label Switching Experimental (MPLS EXP). Therefore, only up to eight classes of traffic can be supported at Layer 2, and inter-class relative priority (such as RFC 2597 Assured Forwarding Drop Preference markdown) is not supported. Layer 3-based DSCP markings allow for up to 64 classes of traffic, which provides more flexibility and is adequate in large-scale deployments and for future requirements.

As the network border blurs between enterprise network and service providers, the need for interoperability and complementary QoS markings is critical. Cisco recommends following the IETF standards-based DSCP PHB markings to ensure interoperability and future expansion. Because the medium enterprise voice, video, and data applications marking recommendations are standards-based, as previously discussed, medium enterprises can easily adopt these markings to interface with service provider classes of service.

## Policing and Markdown

There is little reason to forward unwanted traffic that gets policed and drop by a subsequent tier node, especially when unwanted traffic is the result of DoS or worm attacks in the enterprise network. Excessive volume attack traffic can destabilize network systems, which can result in outages. Cisco recommends policing traffic flows as close to their sources as possible. This principle applies also to legitimate flows, because worm-generated traffic can masquerade under legitimate, well-known TCP/UDP ports and cause extreme amounts of traffic to be poured into the network infrastructure. Such excesses should be monitored at the source and marked down appropriately.

Whenever supported, markdown should be done according to standards-based rules, such as RFC 2597 (AF PHB). For example, excess traffic marked to AFx1 should be marked down to AFx2 (or AFx3 whenever dual-rate policing such as defined in RFC 2698 is supported). Following such markdowns, congestion management policies, such as DSCP-based WRED, should be configured to drop AFx3 more aggressively than AFx2, which in turn should be dropped more aggressively than AFx1.

## Queuing and Dropping

Critical media applications require uncompromised performance and service guarantees regardless of network conditions. Enabling outbound queuing in each network tier provides end-to-end service guarantees during potential network congestion. This common principle applies to campus-to-WAN/Internet edges, where speed mismatches are most pronounced; and campus interswitch links, where oversubscription ratios create the greater potential for network congestion.

Because each application class has unique service level requirements, each should be assigned optimally a dedicated queue. A wide range of platforms in varying roles exist in medium enterprise networks, so each must be bounded by a limited number of hardware or service provider queues. No fewer than four queues are required to support QoS policies for various types of applications, specifically as follows:

- Realtime 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)

Additional queuing recommendations for these classes are discussed next.

## Strict-Priority Queuing

The realtime or strict priority class corresponds to the RFC 3246 EF PHB. The amount of bandwidth assigned to the realtime queuing class is variable. However, if the majority of bandwidth is provisioned with strict priority queuing (which is effectively a FIFO queue), the overall effect is a dampening of QoS functionality, both for latency- and jitter-sensitive realtime applications (contending with each other within the FIFO priority queue), and also for non-realtime applications (because these may periodically receive significant bandwidth allocation fluctuations, depending on the instantaneous amount of traffic being serviced by the priority queue). Remember that the goal of convergence is to enable voice, video, and data applications to transparently co-exist on a single medium enterprise network infrastructure. When realtime applications dominate a link, non-realtime applications fluctuate significantly in their response times, destroying the transparency of the converged network.

For example, consider a 45 Mbps DS3 link configured to support two Cisco TelePresence CTS-3000 calls with an EF PHB service. Assuming that both systems are configured to support full high definition, each such call requires 15 Mbps of strict-priority queuing. Before the TelePresence calls are placed, non-realtime applications have access to 100 percent of the bandwidth on the link; to simplify the example, assume there are no other realtime applications on this link. However, after these TelePresence calls are established, all non-realtime applications are suddenly contending for less than 33 percent of the link. TCP windowing takes effect and many applications hang, timeout, or become stuck in a non-responsive state, which usually translates into users calling the IT help desk to complain about the network (which happens to be functioning properly, albeit in a poorly-configured manner).



### Note

As previously discussed, Cisco IOS software allows the abstraction (and thus configuration) of multiple strict priority LLQs. In such a multiple LLQ context, this design principle applies to the sum of all LLQs to be within one-third of link capacity.

It is vitally important to understand that this strict priority queuing rule is simply a best practice design recommendation and is not a mandate. There may be cases where specific business objectives cannot be met while holding to this recommendation. In such cases, the medium enterprise network administrator must provision according to their detailed requirements and constraints. However, it is important to recognize the tradeoffs involved with over-provisioning strict priority traffic and its negative performance impact, both on other realtime flows and also on non-realtime-application response times.

And finally, any traffic assigned to a strict-priority queue should be governed by an admission control mechanism.

## Best Effort Queuing


The best effort class is the default class for all traffic that has not been explicitly assigned to another application-class queue. Only if an application has been selected for preferential/differential treatment is it removed from the default class. Because most medium enterprises may have several types of applications running in networks, adequate bandwidth must be provisioned for this class as a whole to handle the number and volume of applications that default to it. Therefore, Cisco recommends reserving at least 25 percent of link bandwidth for the default best effort class.

## Scavenger Class Queuing

Whenever the scavenger queuing class is enabled, it should be assigned a minimal amount of link bandwidth capacity, such as 1 percent, or whatever the minimal bandwidth allocation that the platform supports. On some platforms, queuing distinctions between bulk data and scavenger traffic flows cannot be made, either because queuing assignments are determined by class of service (CoS) values (and both of these application classes share the same CoS value of 1), or because only a limited amount of hardware queues exist, precluding the use of separate dedicated queues for each of these two classes. In such cases, the scavenger/bulk queue can be assigned a moderate amount of bandwidth, such as 5 percent.

These queuing rules are summarized in [Figure 2-48](#), where the inner pie chart represents a hardware or service provider queuing model that is limited to four queues and the outer pie chart represents a corresponding, more granular queuing model that is not bound by such constraints.

**Figure 2-48** Compatible 4-Class and 12-Class Queuing Models



228499

## Deploying QoS in Campus LAN Network

All Layer 2 and Layer 3 systems in IP-based networks forward traffic based on a best-effort, providing no differentiated services between different class-of-service network applications. The routing protocol forwards packets over the best low-metric or delay 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 medium enterprise LAN and WAN is a multi-service network designed to support a wide-range of low-latency voice and high bandwidth video with critical and non-critical data traffic over a single network infrastructure. For an optimal user-experience the real time applications (such as voice, video) require packets delivered within specified loss, delay and jitter parameters. Cisco quality-of-service (QoS) is a collection of features and hardware capabilities that allow the network to intelligently dedicate the network resources for higher priority real-time applications, while reserving sufficient network resources to service medium to lower non-real-time traffic. QoS accomplishes this by creating a more application-aware Layer 2 and Layer 3 network to provide differentiated services to network applications and traffic. For a detailed discussion of QoS, refer to the *Enterprise QoS Design Guide* at the following URL:

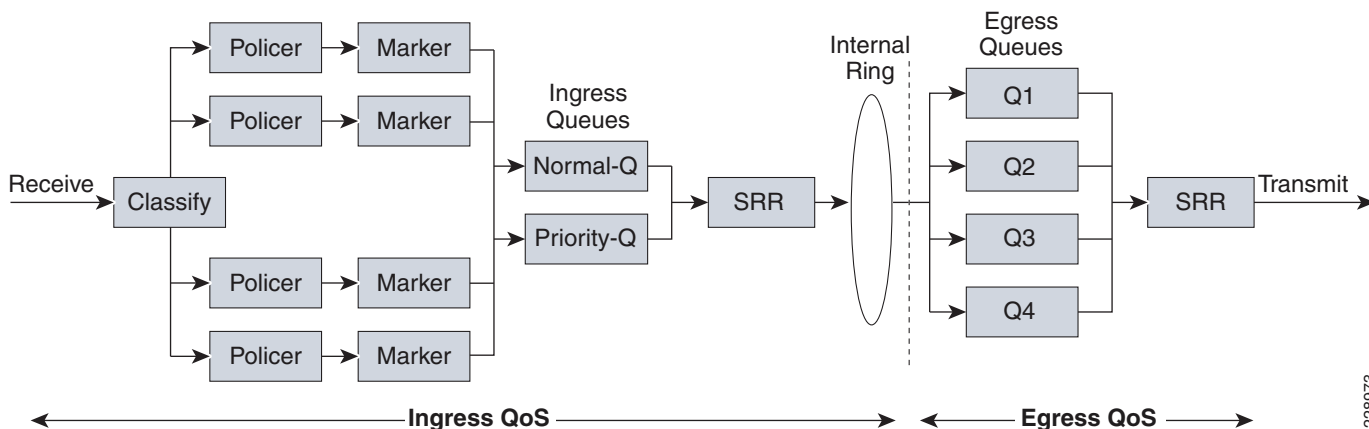
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND/QoS-SRND-Book.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html)

While the QoS design principles across the network are common, the QoS implementation in hardware and software-based switching platforms vary due to internal system design. 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, 3560-X, and 3750-X Series switches are similar to one another. There is no difference in the ingress or egress packet classification, marking, queuing and scheduling implementation among these Catalyst platforms. The Cisco Catalyst switches allow users to create policy-maps by classifying incoming traffic (Layer 2 to Layer 4), and then attaching the policy-map to an individual physical port or to logical interfaces (SVI or port-channel). This creates a common QoS policy that 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 move to the egress ports, where congestion may occur. Congestion in access-layer switches can be prevented by tuning queuing scheduler and Weighted Tail Drop (WTD) drop parameters. See [Figure 2-49](#).

**Figure 2-49** QoS Implementation in Cisco Catalyst Switches




228973

The main difference between these platforms is the switching capacity that ranges from 1G to 10G. The switching architecture and some of the internal QoS structure also differs between these switches. The following are some important differences to consider when selecting an 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 2960-S platform do not support ingress queueing and buffer function in hardware. All other ingress and egress queueing, buffer and bandwidth sharing function remain consistent as 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 Catalyst 2960-S is deployed in FlexStack configuration mode, there is a flexibility to assign different buffer size on egress queue of StackPort. [Figure 2-50](#) illustrates QoS architecture on Catalyst 2960-S Series platform

**Figure 2-50 QoS Implementation in Catalyst 2960-S Switches**



229372

## QoS in Cisco Modular Switches

The Cisco Catalyst 4500-E and 6500-E are high-density, resilient switches for large scale networks. The medium enterprise LAN network design uses both platforms across the network; therefore, all the QoS recommendations in this section for these platforms will remain consistent. Both Catalyst platforms are modular in design; however, there are significant internal hardware architecture differences between the two platforms that impact the QoS implementation model.

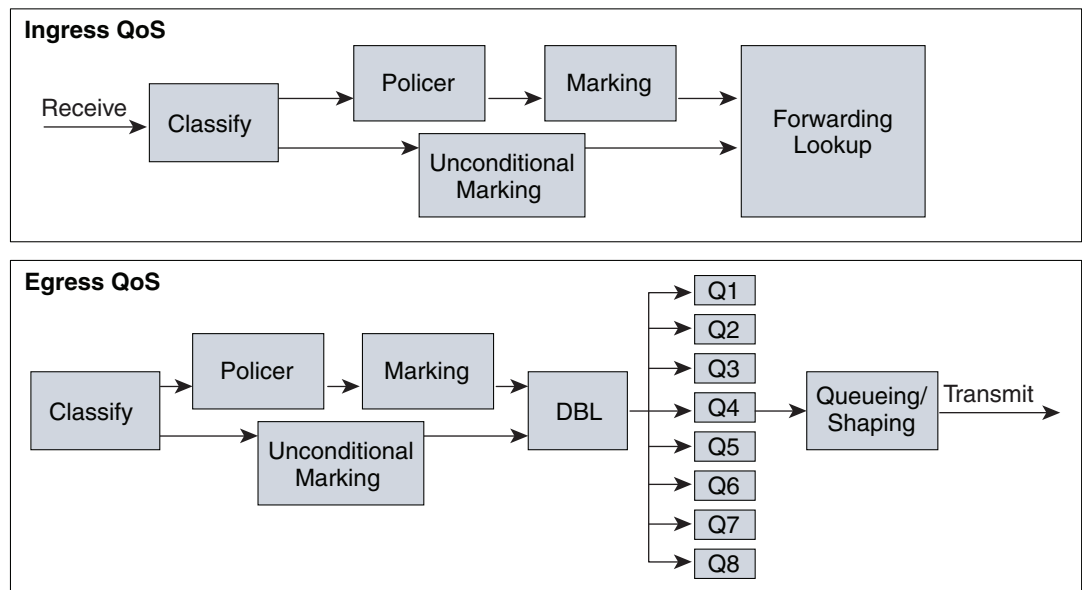
## Catalyst 4500-E QoS

The Cisco Catalyst 4500-E Series platform are widely deployed with classic and next-generation supervisors. This design guide recommends deploying the next-generation supervisor Sup6E and Sup6L-E that offers a number of technical benefits that are beyond QoS.

The Cisco Catalyst 4500 with next generation Sup-6E and Sup6L-E (see [Figure 2-51](#)) are designed to offer better differentiated and preferential QoS services for various class-of-service traffic. New QoS capabilities in the Sup-6E and Sup6L-E 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 and Sup6L-E supports the Modular QoS CLI (MQC) as implemented in IOS-based routers that enhances QoS capabilities and eases implementation and operations. The following are some of the key QoS features that 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 the Trust model and internal Table-map as a tool to classify and mark ingress traffic.
- **Internal DSCP**—The queue placement in Sup-6E and Sup6L-E is simplified by leveraging the MQC capabilities to explicitly map DSCP or CoS traffic in a 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 and Sup6L-E.
- **Sequential vs Parallel Classification**—With MQC-based QoS classification, the Sup6-E and Sup6L-E provides sequential classification rather than parallel. The sequential classification method allows network administrators to classify traffic at the egress based on the ingress markings.

**Figure 2-51 Catalyst 4500—Supervisor 6-E and 6L-E QoS Architecture**



## Catalyst 6500-E QoS

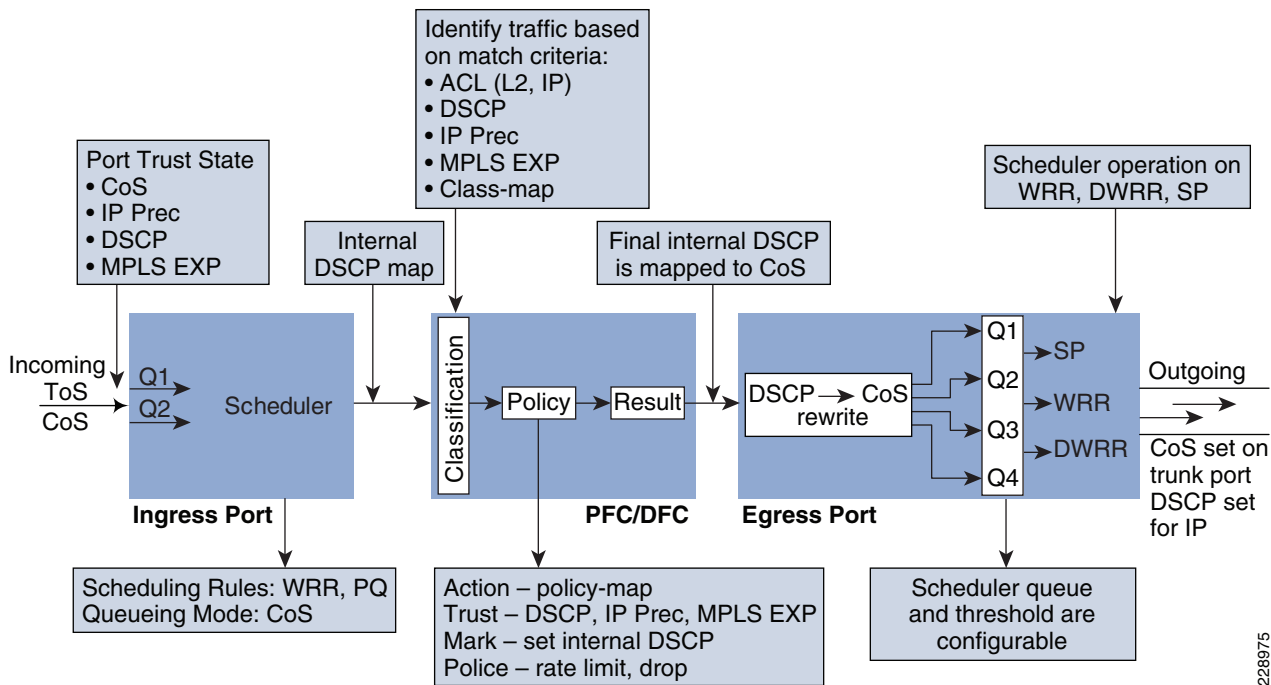
The Cisco Catalyst 6500-E Series are enterprise-class switches, with next-generation hardware and software capabilities designed to deliver innovative, secure, converged network services regardless of its place in the network. The Cisco Catalyst 6500-E can be deployed as a service-node in the campus



network to offer a high performance, robust, intelligent application and network awareness services. The Catalyst 6500-E provides leading-edge Layer 2-Layer 7 services, including rich high availability, manageability, virtualization, security, and QoS feature sets, as well as integrated Power-over-Ethernet (PoE), allowing for maximum flexibility in virtually any role within the campus.

Depending on the network services and application demands of the Cisco Catalyst 6500-E, the platform can be deployed with different types of Supervisor modules—Sup720-10GE, Sup720 and Sup32. This design guide uses the Sup720-10GE supervisor, which is built with next-generation hardware allowing administrators to build virtual-network-systems in the enterprise LAN network. These supervisors leverage various featured daughter cards, including the Multilayer Switch Feature Card (MSFC) that serves as the routing engine, the Policy Feature Card (PFC) that serves as the primary QoS engine, as well as various Distributed Feature Cards (DFCs) that serve to scale policies and processing. Specifically relating to QoS, the PFC sends a copy of the QoS policies to the DFC to provide local support for the QoS policies, which enables the DFCs to support the same QoS features that the PFC supports. Since Cisco VSS is designed with a distributed forwarding architecture, the PFC and DFC functions are enabled and active on active and hot-standby virtual-switch nodes. Figure 2-52 provides internal PFC based QoS architecture.

**Figure 2-52 Cisco Catalyst 6500-E PFC QoS Architecture**



## Deploying Access-Layer QoS

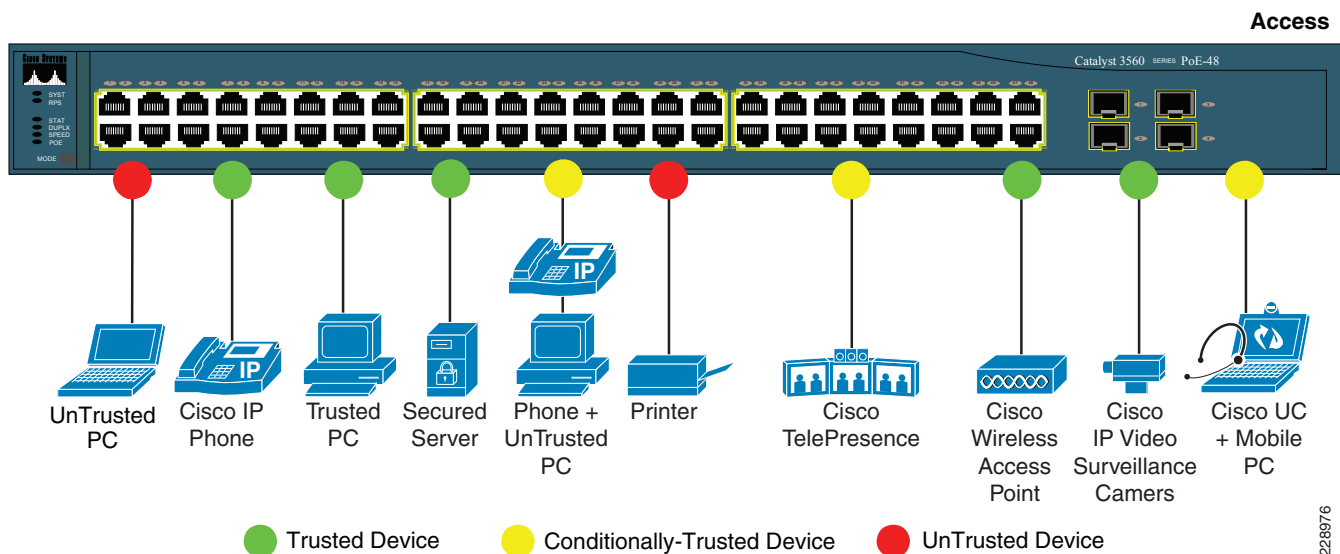
The campus access switches provide the entry point to the network for various types of end devices managed by medium enterprise IT department or employee's personal devices (i.e., laptop etc.). The access switch must decide whether to accept the QoS markings from each endpoint, or whether to change them. This is determined by the QoS policies, and the trust model with which the endpoint is deployed.

## QoS Trust Boundary

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.

The devices (routers, switches, WLC) within the internal network boundary 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 untrusted). Figure 2-53 illustrates several types of devices in the network edge.

**Figure 2-53** Campus LAN QoS Trust Boundary



Enterprise network administrator must identify and classify each of this device type 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 pass through network access security policies and are managed by network administrator. 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 an 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.

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.
- *Queuing*—To provide differentiated services internally in the Catalyst 29xx and 3xxx switching fabric, all approved traffic is queued into priority or non-priority ingress queue. Ingress queuing architecture assures real-time applications, like VoIP traffic, are given appropriate priority (eg transmitted before data traffic).

## Enabling QoS

By default, QoS is disabled on all Catalyst 29xx and 3xxx Series 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 campus network LAN network.

### Access-Layer 29xx and 3xxx (Multilayer or Routed Access)

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



**Note** QoS function on Catalyst 4500-E with Sup6E and Sup6L-E is enabled with the policy-map attached to the port and do not require any additional global configuration.

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.

## QoS Trust Mode (Multilayer or Routed-Access)

### Trusted Port

- 29xx and 3xxx (Multilayer or Routed Access)

```
cr22-3560-LB(config)#interface GigabitEthernet0/5
cr22-3560-LB(config-if)# description CONNECTED TO IPVS 2500 - CAMERA
cr22-3560-LB(config-if)# mls qos trust dscp
cr22-3560-LB#show mls qos interface Gi0/5
GigabitEthernet0/5
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
```

- 4500-E-Sup6LE (Multilayer or Routed Access)

By default all the Sup6E and Sup6L-E ports are in trusted mode, such configuration leverages internal DSCP mapping table to automatically classify QoS bit settings from incoming traffic and place it to appropriate to queue based on mapping table. To appropriate network policy the default settings must be modified by implementing ingress QoS policy-map. Refer to the [“Implementing Ingress QoS Policing”](#) section on page 2-94 for further details.

### Conditionally-Trusted Port

```
cr22-3560-LB(config)#interface Gi0/4
cr22-3560-LB(config-if)# description CONNECTED TO PHONE+PC
cr22-3560-LB(config-if)# mls qos trust device cisco-phone
cr22-3560-LB(config-if)# mls qos trust dscp

cr22-3560-LB#show mls qos interface Gi0/4
GigabitEthernet0/4
trust state: not trusted
trust mode: trust dscp
trust enabled flag: dis
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: cisco-phone
qos mode: port-based
```

- 4500-E-Sup6LE (Multilayer or Routed Access)

```
cr22-4507-LB(config)#interface GigabitEthernet3/3
cr22-4507-LB(config-if)# qos trust device cisco-phone

cr22-4507-LB#show qos interface Gig3/3
Operational Port Trust State: Trusted
Trust device: cisco-phone
Default DSCP: 0 Default CoS: 0
Appliance trust: none
```

### UnTrusted Port

As described earlier, the default trust mode is untrusted when globally enabling QoS function. Without explicit trust configuration on Gi0/1 port, the following show command verifies current trust state and mode:

- 29xx and 3xxx (Multilayer or Routed Access)

```
cr22-3560-LB#show mls qos interface Gi0/1
GigabitEthernet0/1
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
```

- 4500-E-Sup6LE (Multilayer or Routed Access)

QoS trust function on Cisco Catalyst 4500-E with Sup6E and Sup6L-E is enabled by default and must be modified with the policy-map attached to the port.

```
cr22-4507-LB#show qos interface GigabitEthernet3/1
Operational Port Trust State: Trusted
Trust device: none
Default DSCP: 0 Default CoS: 0
Appliance trust: none
```

## Implementing Ingress 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 can be distributed across the campus network which may be broadcasting live video feed using multicast streams must be originated from trusted distributed data center servers.
- *Operation, Administration and Management*—Primarily generated by network devices (routers, switches) and collected by management stations which are typically deployed in the trusted data center network, or a network control center.

All applications present at the access edge need to be assigned a classification, as shown in [Figure 2-54](#). Voice traffic is primarily sourced from Cisco IP telephony devices residing in the voice VLAN (VLAN). These are trusted devices, or conditionally trusted (if users also attach PCs, etc.) to the same port. Voice communication may also be sourced from PCs 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 2-54](#).

**Figure 2-54 Ingress QoS Application Model**

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, Backups	Yes	Untrusted
Best Effort	DF	Default Class	Yes	Untrusted
Scavenger	CS1	YouTube, Gaming, P2P	Yes	Untrusted

228977

Modular QoS 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.

- Catalyst 29xx, 3xxx and 4500-E (MultiLayer and Routed Access)

```

cr22-4507-LB(config)#ip access-list extended MULTIMEDIA-CONFERENCING
cr22-4507-LB(config-ext-nacl)# remark RTP
cr22-4507-LB(config-ext-nacl)# permit udp any any range 16384 32767

cr22-4507-LB(config-ext-nacl)#ip access-list extended SIGNALING
cr22-4507-LB(config-ext-nacl)# remark SCCP
cr22-4507-LB(config-ext-nacl)# permit tcp any any range 2000 2002
cr22-4507-LB(config-ext-nacl)# remark SIP
cr22-4507-LB(config-ext-nacl)# permit tcp any any range 5060 5061
cr22-4507-LB(config-ext-nacl)# permit udp any any range 5060 5061

cr22-4507-LB(config-ext-nacl)#ip access-list extended TRANSACTIONAL-DATA
cr22-4507-LB(config-ext-nacl)# remark HTTPS
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 443
cr22-4507-LB(config-ext-nacl)# remark ORACLE-SQL*NET
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 1521
cr22-4507-LB(config-ext-nacl)# permit udp any any eq 1521
cr22-4507-LB(config-ext-nacl)# remark ORACLE
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 1526
cr22-4507-LB(config-ext-nacl)# permit udp any any eq 1526
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 1575
cr22-4507-LB(config-ext-nacl)# permit udp any any eq 1575
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 1630

```

```

cr22-4507-LB(config-ext-nacl)#ip access-list extended BULK-DATA
cr22-4507-LB(config-ext-nacl)# remark FTP
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq ftp
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq ftp-data
cr22-4507-LB(config-ext-nacl)# remark SSH/SFTP
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 22
cr22-4507-LB(config-ext-nacl)# remark SMTP/SECURE SMTP
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq smtp
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 465
cr22-4507-LB(config-ext-nacl)# remark IMAP/SECURE IMAP
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 143
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 993
cr22-4507-LB(config-ext-nacl)# remark POP3/SECURE POP3
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq pop3
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 995
cr22-4507-LB(config-ext-nacl)# remark CONNECTED PC BACKUP
cr22-4507-LB(config-ext-nacl)# permit tcp any eq 1914 any

cr22-4507-LB(config-ext-nacl)#ip access-list extended DEFAULT
cr22-4507-LB(config-ext-nacl)# remark EXPLICIT CLASS-DEFAULT
cr22-4507-LB(config-ext-nacl)# permit ip any any

cr22-4507-LB(config-ext-nacl)#ip access-list extended SCAVENGER
cr22-4507-LB(config-ext-nacl)# remark KAZAA
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 1214
cr22-4507-LB(config-ext-nacl)# permit udp any any eq 1214
cr22-4507-LB(config-ext-nacl)# remark MICROSOFT DIRECT X GAMING
cr22-4507-LB(config-ext-nacl)# permit tcp any any range 2300 2400
cr22-4507-LB(config-ext-nacl)# permit udp any any range 2300 2400
cr22-4507-LB(config-ext-nacl)# remark APPLE ITUNES MUSIC SHARING
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 3689
cr22-4507-LB(config-ext-nacl)# permit udp any any eq 3689
cr22-4507-LB(config-ext-nacl)# remark BITTORRENT
cr22-4507-LB(config-ext-nacl)# permit tcp any any range 6881 6999
cr22-4507-LB(config-ext-nacl)# remark YAHOO GAMES
cr22-4507-LB(config-ext-nacl)# permit tcp any any eq 11999
cr22-4507-LB(config-ext-nacl)# remark MSN GAMING ZONE
cr22-4507-LB(config-ext-nacl)# permit tcp any any range 28800 29100

```

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

```

cr22-4507-LB(config)#class-map match-all VVLAN-SIGNALING
cr22-4507-LB(config-cmap)# match ip dscp cs3

cr22-4507-LB(config-cmap)#class-map match-all VVLAN-VOIP
cr22-4507-LB(config-cmap)# match ip dscp ef

cr22-4507-LB(config-cmap)#class-map match-all MULTIMEDIA-CONFERENCING
cr22-4507-LB(config-cmap)# match access-group name MULTIMEDIA-CONFERENCING

cr22-4507-LB(config-cmap)#class-map match-all SIGNALING
cr22-4507-LB(config-cmap)# match access-group name SIGNALING

cr22-4507-LB(config-cmap)#class-map match-all TRANSACTIONAL-DATA
cr22-4507-LB(config-cmap)# match access-group name TRANSACTIONAL-DATA

cr22-4507-LB(config-cmap)#class-map match-all BULK-DATA
cr22-4507-LB(config-cmap)# match access-group name BULK-DATA

cr22-4507-LB(config-cmap)#class-map match-all DEFAULT
cr22-4507-LB(config-cmap)# match access-group name DEFAULT

```

```
cr22-4507-LB(config-cmap)#class-map match-all SCAVENGER
cr22-4507-LB(config-cmap)# match access-group name SCAVENGER
```

## Implementing Ingress QoS Policing

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 the 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 2-6](#) provides best practice policing guidelines for different classes to be implemented for trusted and conditional-trusted endpoints at the network edge.

**Table 2-6 Access-Layer Ingress Policing Guidelines**

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

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

### Catalyst 29xx

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

- Trusted or Conditionally-Trusted Port Policer

```
cr22-2960-LB(config)#policy-map Phone+PC-Policy
cr22-2960-LB(config-pmap)# class VVLAN-VOIP
cr22-2960-LB(config-pmap-c)# police 1000000 8000 exceed-action drop
cr22-2960-LB(config-pmap-c)# class VVLAN-SIGNALING
cr22-2960-LB(config-pmap-c)# police 1000000 8000 exceed-action drop
cr22-2960-LB(config-pmap-c)# class MULTIMEDIA-CONFERENCING
cr22-2960-LB(config-pmap-c)# police 5000000 8000 exceed-action drop
cr22-2960-LB(config-pmap-c)# class SIGNALING
cr22-2960-LB(config-pmap-c)# police 1000000 8000 exceed-action drop
cr22-2960-LB(config-pmap-c)# class TRANSACTIONAL-DATA
```



```

cr22-2960-LB(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
cr22-2960-LB(config-pmap-c)# class BULK-DATA
cr22-2960-LB(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
cr22-2960-LB(config-pmap-c)# class SCAVENGER
cr22-2960-LB(config-pmap-c)# police 10000000 8000 exceed-action drop
cr22-2960-LB(config-pmap-c)# class DEFAULT
cr22-2960-LB(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit

```

### Catalyst 2960-S, 3xxx and 4500-E (Multilayer and Routed-Access)

- Trusted or Conditionally-Trusted Port Policer

```

cr22-4507-LB(config)#policy-map Phone+PC-Policy
cr22-4507-LB(config-pmap)# class VVLAN-VOIP
cr22-4507-LB(config-pmap-c)# police 128000 8000 exceed-action drop
cr22-4507-LB(config-pmap-c)# class VVLAN-SIGNALING
cr22-4507-LB(config-pmap-c)# police 32000 8000 exceed-action drop
cr22-4507-LB(config-pmap-c)# class MULTIMEDIA-CONFERENCING
cr22-4507-LB(config-pmap-c)# police 5000000 8000 exceed-action drop
cr22-4507-LB(config-pmap-c)# class SIGNALING
cr22-4507-LB(config-pmap-c)# police 32000 8000 exceed-action drop
cr22-4507-LB(config-pmap-c)# class TRANSACTIONAL-DATA
cr22-4507-LB(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
cr22-4507-LB(config-pmap-c)# class BULK-DATA
cr22-4507-LB(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
cr22-4507-LB(config-pmap-c)# class SCAVENGER
cr22-4507-LB(config-pmap-c)# police 10000000 8000 exceed-action drop
cr22-4507-LB(config-pmap-c)# class DEFAULT
cr22-4507-LB(config-pmap-c)# police 10000000 8000 exceed-action policed-dscp-transmit
Catalyst 29xx, 3xxx and 4500-E (Multilayer and Routed-Access)

```

- UnTrusted Port Policer

All ingress traffic (default class) from untrusted endpoint 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:

```

cr22-2960-LB(config)#policy-map UnTrusted-PC-Policy
cr22-2960-LB(config-pmap)# class class-default
cr22-2960-LB(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 2-59](#).

The best practice is to use an 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, with the use of an explicit (seemingly redundant) marking command, and 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

- Catalyst 29xx, 3xxx and 4500-E (Multilayer and Routed-Access)

```

cr22-3750-LB(config)#policy-map Phone+PC-Policy

```

```

cr22-3750-LB(config-pmap)# class VVLAN-VOIP
cr22-3750-LB(config-pmap-c)# set dscp ef
cr22-3750-LB(config-pmap-c)# class VVLAN-SIGNALING
cr22-3750-LB(config-pmap-c)# set dscp cs3
cr22-3750-LB(config-pmap-c)# class MULTIMEDIA-CONFERENCING
cr22-3750-LB(config-pmap-c)# set dscp af41
cr22-3750-LB(config-pmap-c)# class SIGNALING
cr22-3750-LB(config-pmap-c)# set dscp cs3
cr22-3750-LB(config-pmap-c)# class TRANSACTIONAL-DATA
cr22-3750-LB(config-pmap-c)# set dscp af21
cr22-3750-LB(config-pmap-c)# class BULK-DATA
cr22-3750-LB(config-pmap-c)# set dscp af11
cr22-3750-LB(config-pmap-c)# class SCAVENGER
cr22-3750-LB(config-pmap-c)# set dscp cs1
cr22-3750-LB(config-pmap-c)# class DEFAULT
cr22-3750-LB(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

- Catalyst 29xx, 3xxx and 4500-E (Multilayer and Routed-Access)

```

cr22-3750-LB(config)#policy-map UnTrusted-PC-Policy
cr22-3750-LB(config-pmap)# class class-default
cr22-3750-LB(config-pmap-c)# set dscp default

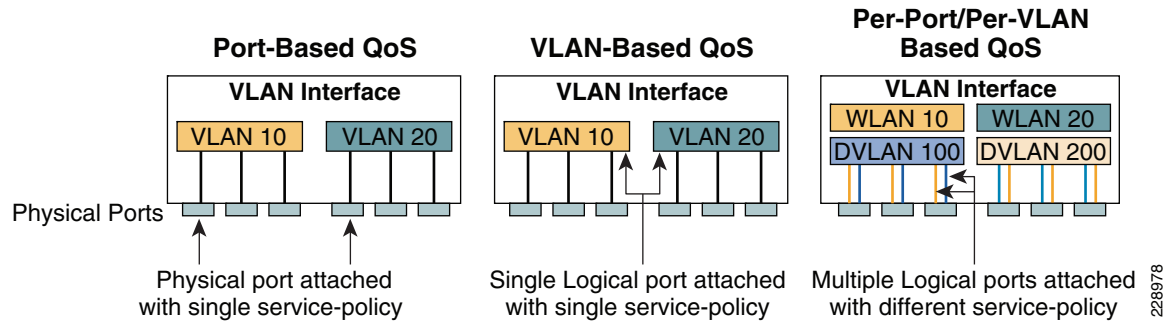
```

### Applying Ingress Policies

After creating complete a policy-map on all the Layer 2 and Layer 3 access-switches with 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 offers three simplified methods to apply service-policies; depending on the deployment model either of the methods can be implemented:

- *Port-Based QoS*—Applying the service-policy on per physical port basis will force traffic to pass-through the QoS policies before entering in to the campus network. Port-Based QoS discretely functions on a per-physical port basis even if it is associated with a logical VLAN which is applied on multiple physical ports.
- *VLAN-Based QoS*—Applying the service-policy on a per VLAN has requires the policy-map to be attached to a logical Layer 3 SVI interface. Every physical port associated to VLAN requires an extra configuration to ensure all traffic to passes through the QoS policies defined on an logical interface.
- *Per-Port / Per-VLAN-Based QoS*—This is not supported on all the Catalyst platforms and the configuration commands are platform-specific. Per-Port/Per-VLAN-based QoS create 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 same physical port.

See [Figure 2-55](#).

**Figure 2-55** *Depicts all three QoS implementation method*

The following sample configuration provides guideline to deploy port-based QoS on the access-layer switches in campus network:

- Catalyst 29xx, 3xxx and 4500-E (Multilayer and Routed-Access)

```
cr22-2960-LB(config)#interface FastEthernet0/1
cr22-2960-LB(config-if)# service-policy input UnTrusted-PC-Policy

cr22-2960-LB#show mls qos interface FastEthernet0/1
FastEthernet0/1
Attached policy-map for Ingress: UnTrusted-PC-Policy
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
```

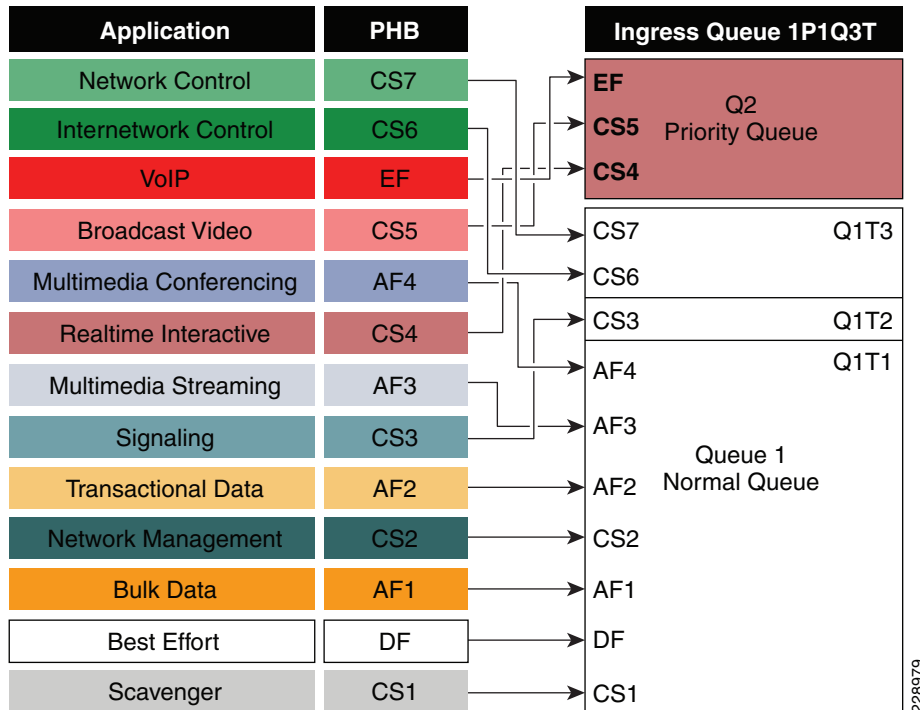
## Applying Ingress Queuing

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 the 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 2-56 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 2-56 Catalyst 2960 and 3xxx Ingress Queuing Model**

- **Catalyst 2960 and 3xxx (Multilayer and Routed-Access)**

```

cr22-3750-LB(config)#mls qos srr-queue input priority-queue 2 bandwidth 30
! Q2 is enabled as a strict-priority ingress queue with 30% BW

cr22-3750-LB (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)

cr22-3750-LB (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
cr22-3750-LB (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
cr22-3750-LB (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
cr22-3750-LB (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
cr22-3750-LB (config)#mls qos srr-queue input dscp-map queue 1 threshold 2 24
! DSCP CS3 is mapped to ingress Q1T2

cr22-3750-LB(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)
cr22-3750-LB(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)

cr22-3750-LB#show mls qos input-queue

```

```

Queue:      12
-----
buffers    :9010
bandwidth  :7030
priority   :030
threshold1:80100
threshold2:90100

cr22-3750-LB#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

```

**Note**

The ingress queuing function on Catalyst 4500-E Sup6E and Sup6L-E is not supported as described in [Figure 2-51](#).

## Implementing Access-Layer Egress QoS

The QoS implementation of egress traffic towards network edge devices on access-layer switches are much simplified compared to ingress traffic which requires stringent QoS policies to provide differentiated services and network bandwidth protection. Unlike the Ingress QoS model, the egress QoS model must provide optimal queuing policies for each class and set the drop thresholds to prevent network congestion and prevent an application performance impact. With egress queuing in DSCP mode, the Cisco Catalyst switching platforms are bounded by a limited number of hardware queues.

### Catalyst 2960 and 3xxx Egress QoS


Cisco Catalyst 29xx and 3xxx Series platform supports four egress queues that are required to support the variable class QoS policies for the medium enterprise campus LAN network; specifically, the following queues would be considered a minimum:

- Realtime 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 interfaces must be deployed with IETF recommended bandwidth allocations for different class-of-service applications:

- The real-time queue should not exceed 33 percent of the link's bandwidth.
- The default queue should be at least 25 percent of the link's bandwidth.
- The bulk/scavenger queue should not exceed 5 percent of the link's bandwidth.

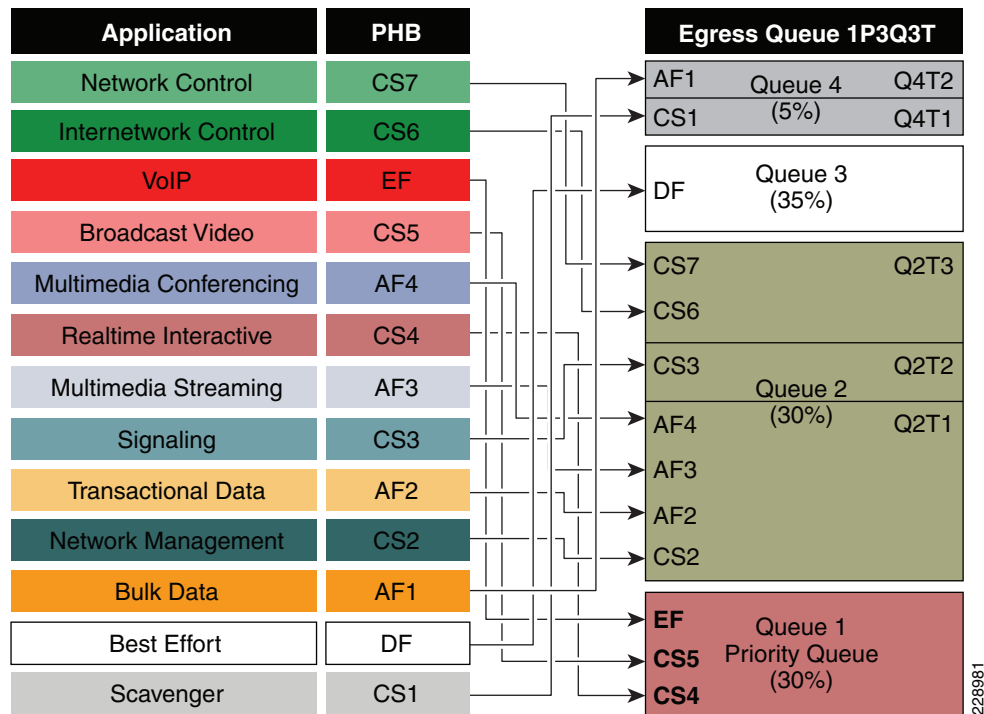
[Figure 2-57](#) illustrates the egress bandwidth allocation best practices design for different classes.

**Figure 2-57 Class-of-Service Egress Bandwidth Allocations**

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; furthermore, 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.

Like the ingress queuing structure that maps various applications based on DSCP value into two ingress queues, the egress queuing must be similar designed to map with four egress queues. The DSCP-to-queue mapping for egress queuing must be mapped to each egress queues as stated above which allows better queuing-policy granularity. A campus egress QoS model example for a platform that supports DSCP-to-queue mapping with a 1P3Q8T queuing structure is depicted in [Figure 2-58](#).

**Figure 2-58 1P3Q3T Egress QoS Model on Catalyst 29xx and 3xxx platforms**

DSCP marked packets are assigned to the appropriate queue and each queue is configured with appropriate WTD threshold as defined in Figure 2-58. Egress queuing settings are common between all the trust-independent network edge ports as well as on the Layer 2 or Layer 3 uplink connected to internal network. The following egress queue configuration entered in global configuration mode must be enabled on every access-layer switch in the network.

- Catalyst 2960, 2960-S and 3xxx (Multilayer and Routed-Access)

```
cr22-3750-LB(config)#mls qos queue-set output 1 buffers 15 30 35 20
! Queue buffers are allocated
cr22-3750-LB (config)#mls qos queue-set output 1 threshold 1 100 100 100 100
! All Q1 (PQ) Thresholds are set to 100%
cr22-3750-LB (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%
cr22-3750-LB (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%
cr22-3750-LB (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%

cr22-3750-LB(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)
cr22-3750-LB(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
cr22-3750-LB(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
```

```

cr22-3750-LB(config)#mls qos srr-queue output dscp-map queue 2 threshold 2 24
! DSCP CS3 is mapped to egress Q2T2
cr22-3750-LB(config)#mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
! DSCP CS6 and CS7 are mapped to egress Q2T3
cr22-3750-LB(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)
cr22-3750-LB(config)#mls qos srr-queue output dscp-map queue 4 threshold 1 8
! DSCP CS1 is mapped to egress Q4T1
cr22-3750-LB(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 edge and uplink port interface with common egress queuing
parameters
cr22-3750-LB(config)#interface range GigabitEthernet1/0/1-48
cr22-3750-LB(config-if-range)# queue-set 1
! The interface(s) is assigned to queue-set 1
cr22-3750-LB(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
cr22-3750-LB(config-if-range)# priority-queue out
! Q1 is enabled as a strict priority queue


cr22-3750-LB#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

```

- Catalyst 4500-E Sup6E and Sup6L-E Egress QoS

The enterprise-class 4500-E switch with next-generation supervisor hardware architecture are designed to offers better egress QoS techniques, capabilities, and flexibilities to provide for a well diverse queuing structure for multiple class-of-service traffic types. Deploying the next-generation Sup-6E and Sup6L-E in the campus network provides more QoS granularity to map the 8-class traffic types to hardware-based egress-queues as illustrated in [Figure 2-59](#).

**Figure 2-59 8 Class-of-Service Egress Bandwidth Allocations**

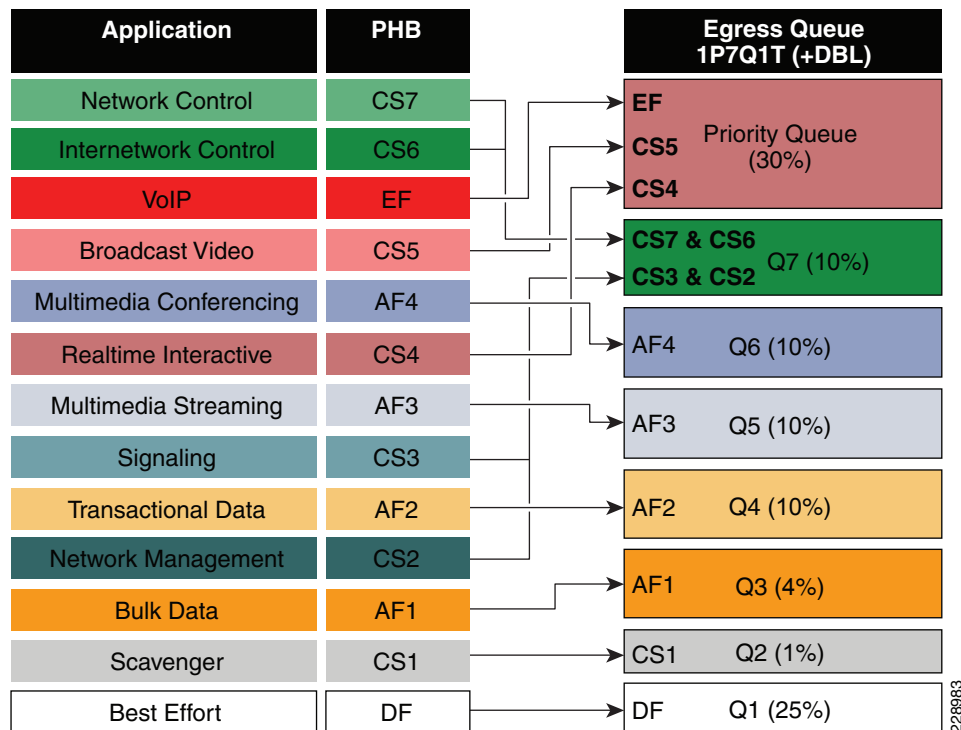


The Cisco Catalyst 4500-E Sup-6E and Sup6L-E supervisor supports platform-specific congestion avoidance algorithms to provide Active Queue Management (AQM), namely Dynamic Buffer Limiting (DBL). DBL tracks the queue length for each traffic flow in the switch. When the queue length of a flow exceeds its limit, DBL drops packets or sets the Explicit Congestion Notification (ECN) bits in the TCP



packet headers. With 8 egress (1P7Q1T) queues and DBL capability in the Sup-6E-based supervisor, the bandwidth distribution for different classes change. Figure 2-60 provides the new recommended bandwidth allocation.

**Figure 2-60 1P7Q1T Egress QoS Model on Catalyst 4500-E with Sup6E and Sup6L-E**



The QoS architecture and implementation procedure are identical between Sup-6E and Sup6L-E modules. Implementing QoS policies on Sup-6E-based Catalyst 4500 platform follows the IOS (MQC) based configuration model instead of the Catalyst OS-based QoS model. To take advantage of hardware-based QoS egress, the queuing function using MQC must be applied on per member-link of the EtherChannel interface. Therefore, load-sharing egress per-flow traffic across EtherChannel links offers the advantage to optimally use distributed hardware resources.

Recommended DSCP markings for each traffic class can be classified in a different class-map for egress QoS functions. Based on Figure 2-60, the following configuration use the new egress policy-map with queuing and DBL function implemented on the Catalyst 4500-E deployed with a Sup6E and SupL-E supervisor module. All network edge port and core-facing uplink ports must use a common egress policy-map.

- Catalyst 4500 Sup-6E and SupL-E (MultiLayer and Routed-Access)

```
! Creating class-map for each classes using match dscp statement as marked by edge systems
cr22-4507-LB(config)#class-map match-all PRIORITY-QUEUE
cr22-4507-LB(config-cmap)# match dscp ef
cr22-4507-LB(config-cmap)# match dscp cs5
cr22-4507-LB(config-cmap)# match dscp cs4
cr22-4507-LB(config-cmap)#class-map match-all CONTROL-MGMT-QUEUE
cr22-4507-LB(config-cmap)# match dscp cs7

cr24-4507-LB(config-cmap)# match dscp cs6
cr24-4507-LB(config-cmap)# match dscp cs3
cr24-4507-LB(config-cmap)# match dscp cs2
```

```

cr24-4507-LB(config-cmap)#class-map match-all MULTIMEDIA-CONFERENCING-QUEUE
cr24-4507-LB(config-cmap)# match dscp af41 af42 af43
cr24-4507-LB(config-cmap)#class-map match-all MULTIMEDIA-STREAMING-QUEUE
cr24-4507-LB(config-cmap)# match dscp af31 af32 af33
cr24-4507-LB(config-cmap)#class-map match-all TRANSACTIONAL-DATA-QUEUE
cr24-4507-LB(config-cmap)# match dscp af21 af22 af23
cr24-4507-LB(config-cmap)#class-map match-all BULK-DATA-QUEUE
cr24-4507-LB(config-cmap)# match dscp af11 af12 af13
cr24-4507-LB(config-cmap)#class-map match-all SCAVENGER-QUEUE
cr24-4507-LB(config-cmap)# match dscp cs1

! Creating policy-map and configure queueing for class-of-service
cr22-4507-LB(config)#policy-map EGRESS-POLICY
cr22-4507-LB(config-pmap)# class PRIORITY-QUEUE
cr22-4507-LB(config-pmap-c)# priority
cr22-4507-LB(config-pmap-c)# class CONTROL-MGMT-QUEUE
cr22-4507-LB(config-pmap-c)# bandwidth remaining percent 10
cr22-4507-LB(config-pmap-c)# class MULTIMEDIA-CONFERENCING-QUEUE
cr22-4507-LB(config-pmap-c)# bandwidth remaining percent 10
cr22-4507-LB(config-pmap-c)# class MULTIMEDIA-STREAMING-QUEUE
cr22-4507-LB(config-pmap-c)# bandwidth remaining percent 10
cr22-4507-LB(config-pmap-c)# class TRANSACTIONAL-DATA-QUEUE
cr22-4507-LB(config-pmap-c)# bandwidth remaining percent 10
cr22-4507-LB(config-pmap-c)# dbl
cr22-4507-LB(config-pmap-c)# class BULK-DATA-QUEUE
cr22-4507-LB(config-pmap-c)# bandwidth remaining percent 4
cr22-4507-LB(config-pmap-c)# dbl
cr22-4507-LB(config-pmap-c)# class SCAVENGER-QUEUE
cr22-4507-LB(config-pmap-c)# bandwidth remaining percent 1
cr22-4507-LB(config-pmap-c)# class class-default
cr22-4507-LB(config-pmap-c)# bandwidth remaining percent 25
cr22-4507-LB(config-pmap-c)# dbl

! Attaching egress service-policy on all physical member-link ports
cr24-4507-DO(config)#int range Ten3/1 , Te4/1 , Ten5/1 , Ten5/4, Ten Gi1/1 - 6
cr24-4507-DO(config-if-range)# service-policy output EGRESS-POLICY

```

## Policing Priority-Queue

EtherChannel is an aggregated logical bundle of interfaces that do not perform queuing and rely on individual member-links to queue egress traffic by using hardware-based queuing. The hardware-based priority-queue implementation on the Catalyst 4500-E does not support a built-in policer to restrict traffic during network congestion. To mitigate this challenge, it is recommended to implement an additional policy-map to rate-limit the priority class traffic and must be attached on the EtherChannel to govern the aggregated egress traffic limits. The following additional policy-map must be created to classify priority-queue class traffic and rate-limit up to 30 percent egress link capacity:

```

cr22-4507-LB(config)#class-map match-any PRIORITY-QUEUE
cr22-4507-LB (config-cmap)# match dscp ef
cr22-4507-LB (config-cmap)# match dscp cs5
cr22-4507-LB (config-cmap)# match dscp cs4

cr22-4507-LB (config)#policy-map PQ-POLICER
cr22-4507-LB (config-pmap)# class PRIORITY-QUEUE
cr22-4507-LB (config-pmap-c)# police cir 300 m conform-action transmit exceed-action drop

cr22-4507-LB (config)#interface range Port-Channel 1
cr22-4507-LB (config-if-range)#service-policy output PQ-POLICER

```

**Table 2-7 Summarized Access-Layer Ingress QoS Deployment Guidelines**

End-Point	Trust Model	DSCP Trust	Classification	Marking	Policing	Ingress Queuing <sup>1</sup>
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

1. Catalyst 29xx and 3xxx only

**Table 2-8 Summarized Access-Layer Egress QoS Deployment Guidelines**

End-Point	Trust Model	Classification / Marking / Policing	Egress Queuing	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	Yes (PQ Policer)	Yes	Yes

## Deploying Network-Layer QoS

Campus network systems at the main site and remote campus are managed and maintained by the enterprise IT administration to provide key network foundation services such as routing, switching, QoS, and virtualization. In a best practice network environment, these systems must be implemented with the recommended configuration to provide differentiated network services on per-hop basis. To allow for consistent application delivery through the network, it is recommended to implement bidirectional QoS policies on distribution and core layer systems.

## QoS Trust Boundary

All medium enterprise IT managed campus LAN and WAN network systems can be classified as trusted device and must follow same QoS best practices recommended in previous subsection. It is recommended to avoid deploying trusted or untrusted endpoints directly to the campus distribution and core layer systems.

Based on global network QoS policy each class-of-service applications get common treatment. Independent of enterprise network tier—LAN/WAN, platform type and their capabilities—each devices in the network will protect service quality and enable communication across the network without degrading the application performance.

## Implementing Network-Layer Ingress QoS

As described earlier, the internal campus core network must be considered to be trusted. The next-generation Cisco Catalyst access-layer platform must be deployed with more application-aware and intelligence at the network edge. The campus core and distribution network devices should rely on the access-layer switches to implement QoS classification and marking based on a wide-range of applications and IP-based devices deployed at the network edge.

To provide consistent and differentiated QoS services on per-hop basis across the network, the distribution and core network must be deployed to trust incoming pre-marked DSCP traffic from the downstream Layer 2 or Layer 3 network device. This medium enterprise LAN network design recommends deploying a broad-range of Layer-3 Catalyst switching platforms in the campus distribution and core layer. As mentioned in the previous section, the hardware architecture of each switching platform is different, based on the platform capabilities and resources. This will change how each various class-of-service traffic will be handled in different directions: ingress, switching fabric, and egress.

Cisco Catalyst access-layer switches must classify the application and device type to marks DSCP value based on the trust model with deep packet inspection using access-lists (ACL) or protocol-based device discovery; therefore, there is no need to reclassify the same class-of-service at the campus distribution and core layer. The campus distribution and core layers can trust DSCP markings from access-layer and provide QoS transparency without modifying the original parameters unless the network is congested.

Based on the simplified internal network trust model, the ingress QoS configuration also becomes more simplified and manageable. This subsection provides common ingress QoS deployment guidelines for the campus distribution and core for all locations:

### QoS Trust Mode

As described earlier, the Catalyst 4500-E deployed with either a Sup6E or Sup6L-E supervisor module in the distribution or core layer will automatically sets the physical ports in the trust mode. The Catalyst 4500-E by default will perform DSCP-CoS or CoS-DSCP mappings to transmit traffic transparently without any QoS bits rewrites. However the default QoS function on campus distribution or core platforms like the Catalyst 3750-X and 6500-E Series switches is disabled.

The network administrator must manually enable QoS globally on the switch and explicitly enable DSCP trust mode on each logical EtherChannel and each member-link interface connected to upstream and downstream devices. The distribution layer QoS trust configuration is the same for a multilayer or routed-access deployment. The following sample QoS configuration must be enabled on all the distribution and core layer switches deployed in campus LAN network.

### Distribution-Layer Catalyst 3750-X and 6500-E

- 3750-X and 6500-E (Multilayer or Routed Access)

```
cr22-6500-LB(config)#mls qos
cr22-6500-LB#show mls qos
```

```
QoS is enabled globally
...
```

### Implement DSCP Trust Mode

- Catalyst 6500-E (Multilayer or Routed Access)

```
cr22-6500-LB(config)#interface Port-channel100
cr22-6500-LB(config-if)# description Connected to cr22-4507-LB
cr22-6500-LB(config-if)# mls qos trust dscp
```

Catalyst 6500-E will automatically replicate “mls qos trust dscp” command from port-channel interface to each bundled member-links.

```
cr22-6500-LB#show queueing interface Ten1/1/2 | inc QoS|Trust
Port QoS is enabled
Trust boundary disabled
Trust state: trust DSCP
```

### Catalyst 3750-X (Multilayer or Routed Access)

Catalyst 3750-X does not support **mls qos trust dscp** command on port-channel interface; therefore, network administrator must apply this command on each bundled member-links.

```
cr36-3750x-xSB(config)#interface range Ten1/0/1 - 2 , Ten2/0/1 - 2
cr36-3750x-xSB(config-if-range)# description Connected to cr23-VSS-Core
cr36-3750x-xSB(config-if-range)# mls qos trust dscp
```

```
cr36-3750x-xSB#show mls qos interface Ten1/0/1
TenGigabitEthernet1/0/1
trust state: trust dscp
trust mode: trust dscp
...
```

## Applying Ingress Queuing

When Cisco Catalyst 3750-X and 6500-E switching platforms receive various class-of-service requests from different physical ports, then depending on the DSCP and CoS markings it can queue the traffic prior sending it to the switching fabric in a FIFO manner. Both Catalyst platforms support up to two ingress queues but how they are implemented differs. The Cisco Catalyst 4500-E deployed with a Sup6E or a Sup6L-E supervisor module does not support ingress queuing.

### Implementing Catalyst 3750-X Ingress Queuing

The ingress queuing function in the distribution-layer Catalyst 3750-X StackWise Plus must be deployed to differentiate and place the normal versus high-priority class traffic in separate ingress queue before forwarding it to the switching fabric.

For consistent QoS within the campus network, the core and access layers should map DSCP-marked traffic into ingress queues the same way. Refer to the [“Applying Ingress Queuing” section on page 2-97](#) for implementation detail.

### Implementing Catalyst 6500-E Ingress Queuing

There are two main considerations relevant to ingress queuing design on the Catalyst 6500/6500-E:

- The degree of oversubscription (if any) of the linecard
- Whether the linecard requires trust-CoS to be enabled to engage ingress queuing

Some linecards may be designed to support a degree of oversubscription that theoretically offers more traffic to the linecard than the sum of all GE/10GE switch ports than can collectively access the switching backplane at once. Since such a scenario is extremely unlikely, it is often more cost-effective to use linecards that have a degree of oversubscription within the campus network. However, if this design choice has been made, it is important for network administrators to recognize the potential for drops due to oversubscribed linecard architectures. To manage application-class service levels during such extreme scenarios, ingress queuing models may be enabled.

While the presence of oversubscribed linecard architectures may be viewed as the sole consideration as to enabling ingress queuing or not, a second important consideration that many Catalyst 6500-E linecards only support CoS-based ingress queuing models that reduces classification and marking granularity—limiting the administrator to an 8-class 802.1Q/p model. Once CoS is trusted, DSCP values are overwritten (via the CoS-to-DSCP mapping table) and application classes sharing the same CoS values are longer distinguishable from one another. Therefore, given this classification and marking limitation and the fact that the value of enabling ingress queuing is only achieved in extremely rare scenarios, it is not recommended to enable CoS-based ingress queuing on the Catalyst 6500-E; rather, limit such linecards and deploy either non-oversubscribed linecards and/or linecards supporting DSCP-based queuing at the distribution and core layers of the campus network.

Table 2-9 summarizes recommended linecards consideration by listing and oversubscription ratios and whether the ingress queuing models are CoS or DSCP-based.

**Table 2-9 Catalyst 6500-E Switch Module Ingress Queuing Architecture**

Switch Module	Maximum Input	Maximum Output (To Backplane)	Oversubscription Ratio	Ingress Queuing Structure	CoS / DSCP Based	Ingress Queuing Recommendations
WS-6724-SFP	24 Gbps (24 x GE ports)	40 Gbps (2 x 20 Gbps)	-	1P3Q8T	CoS based	Not Required
WS-6704-10GE	40 Gbps (4 x 10GE ports)		-	8Q8T	CoS or DSCP based	Not Required
WS-6708-10GE	80 Gbps (8 x 10GE ports)		2:1	8Q4T	CoS or DSCP based	Use DSCP-based 8Q4T ingress queuing
WS-6716-10GE	160 Gbps (16 x 10GE ports)		4:1	8Q4T / 1P7Q2T*	CoS or DSCP based	Use DSCP-based 1P7Q2T ingress queuing



**Note**

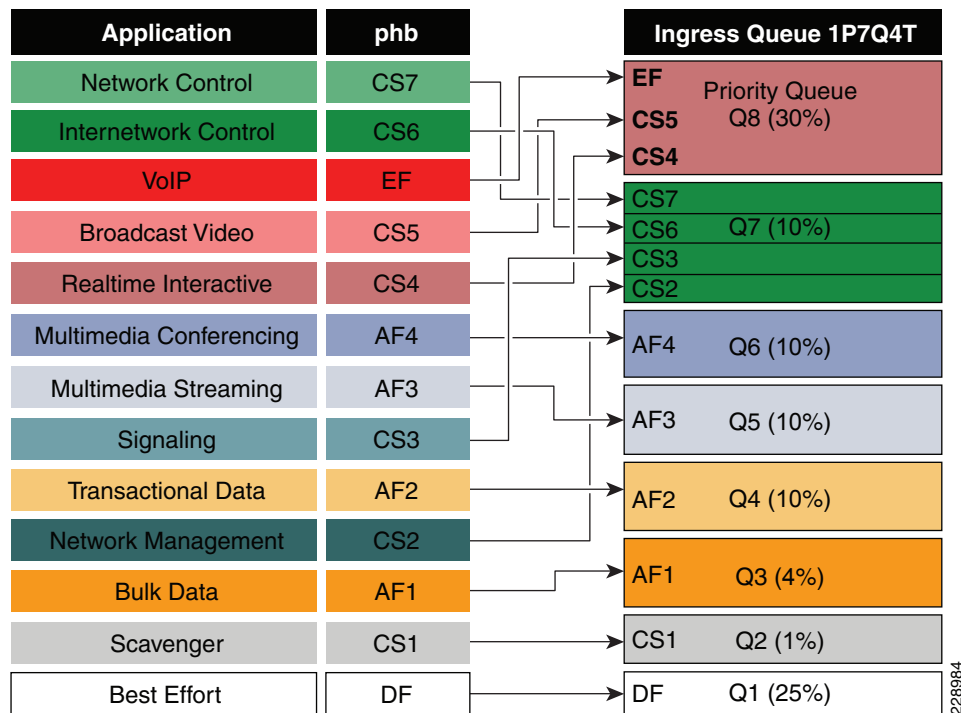
The Catalyst WS-X6716-10GE can be configured to operate in Performance Mode (with an 8Q4T ingress queuing structure) or in Oversubscription Mode (with a 1P7Q2T ingress queuing structure). In Performance mode, only one port in every group of four is operational (while the rest are administratively shut down), which eliminates any oversubscription on this linecard and as such ingress queuing is not required (as only 4 x 10GE ports are active in this mode and the backplane access rate is also at 40 Gbps). In Oversubscription Mode (the default mode), all ports are operational and the maximum oversubscription ratio is 4:1. Therefore, it is recommended to enable 1P7Q2T DSCP-based ingress queuing on this linecard in Oversubscription Mode.

Additional details on these WS-X6716-10GE operational modes can be found at the following URL:  
[http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/qa\\_cisco\\_catalyst\\_6500\\_series\\_16port\\_10gigabit\\_ethernet\\_module.html](http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/qa_cisco_catalyst_6500_series_16port_10gigabit_ethernet_module.html)

If 6708 and 6716 linecards (with the latter operating in oversubscription mode) are used in the distribution and core layers of the campus network, then 8Q4T DSCP-based ingress queuing and 1P7Q2T DSCP-based ingress queuing (respectively) are recommended to be enabled. These queuing models are detailed in the following sections.

Figure 2-61 depicts how different class-of-service applications are mapped to the Ingress Queue structure (8Q4T) and how each queue is assigned a different WTD threshold.

**Figure 2-61 Catalyst 6500-E Ingress Queuing Model**



The corresponding configuration for 8Q8T (DSCP-to-Queue) ingress queuing on a Catalyst 6500-E VSS in distribution and core layer is shown below. PFC function is active on active and hot-standby virtual-switch nodes; therefore, ingress queuing must be configured on each distributed member-links of Layer 2 or Layer 3 MEC.

- Distribution and Core-Layer Catalyst 6500-E in VSS mode

```
! This section configures the port for DSCP-based Ingress queuing
cr22-vss-core(config)#interface range TenGigabitEthernet 1/1/2 - 8 , 2/1/2-8
cr22-vss-core(config-if-range)# mls qos queue-mode mode-dscp
! Enables DSCP-to-Queue mapping
```

```
! This section configures the receive queues BW and limits
cr22-vss-core(config-if-range)# rcv-queue queue-limit 10 25 10 10 10 10 10 15
! Allocates 10% to Q1, 25% to Q2, 10% to Q3, 10% to Q4,
! Allocates 10% to Q5, 10% to Q6, 10% to Q7 and 15% to Q8
cr22-vss-core(config-if-range)# rcv-queue bandwidth 1 25 4 10 10 10 10 30
! Allocates 1% BW to Q1, 25% BW to Q2, 4% BW to Q3, 10% BW to Q4,
! Allocates 10% BW to Q5, 10% BW to Q6, 10% BW to Q7 & 30% BW to Q8
```

```
! This section enables WRED on all queues except Q8
cr22-vss-core(config-if-range)# rcv-queue random-detect 1
! Enables WRED on Q1
cr22-vss-core(config-if-range)# rcv-queue random-detect 2
```

```

! Enables WRED on Q2
cr22-vss-core(config-if-range)# rcv-queue random-detect 3
! Enables WRED on Q3
cr22-vss-core(config-if-range)# rcv-queue random-detect 4
! Enables WRED on Q4
cr22-vss-core(config-if-range)# rcv-queue random-detect 5
! Enables WRED on Q5
cr22-vss-core(config-if-range)# rcv-queue random-detect 6
! Enables WRED on Q6
cr22-vss-core(config-if-range)# rcv-queue random-detect 7
! Enables WRED on Q7
cr22-vss-core(config-if-range)# no rcv-queue random-detect 8
! Disables WRED on Q8

! This section configures WRED thresholds for Queues 1 through 7
cr22-vss-core(config-if-range)# rcv-queue random-detect max-threshold 1 100 100 100 100
! Sets all WRED max thresholds on Q1 to 100%
cr22-vss-core(config-if-range)# rcv-queue random-detect min-threshold 1 80 100 100 100
! Sets Q1T1 min WRED threshold to 80%
cr22-vss-core(config-if-range)# rcv-queue random-detect min-threshold 2 80 100 100 100
! Sets Q2T1 min WRED threshold to 80%
cr22-vss-core(config-if-range)# rcv-queue random-detect max-threshold 2 100 100 100 100
! Sets all WRED max thresholds on Q2 to 100%

cr22-vss-core(config-if-range)# rcv-queue random-detect min-threshold 3 70 80 90 100
! Sets WRED min thresholds for Q3T1, Q3T2, Q3T3 to 70 %, 80% and 90%
cr22-vss-core(config-if-range)# rcv-queue random-detect max-threshold 3 80 90 100 100
! Sets WRED max thresholds for Q3T1, Q3T2, Q3T3 to 80%, 90% and 100%
cr22-vss-core(config-if-range)# rcv-queue random-detect min-threshold 4 70 80 90 100
! Sets WRED min thresholds for Q4T1, Q4T2, Q4T3 to 70 %, 80% and 90%
cr22-vss-core(config-if-range)# rcv-queue random-detect max-threshold 4 80 90 100 100
! Sets WRED max thresholds for Q4T1, Q4T2, Q4T3 to 80%, 90% and 100%
cr22-vss-core(config-if-range)# rcv-queue random-detect min-threshold 5 70 80 90 100
! Sets WRED min thresholds for Q5T1, Q5T2, Q5T3 to 70 %, 80% and 90%
cr22-vss-core(config-if-range)# rcv-queue random-detect max-threshold 5 80 90 100 100
! Sets WRED max thresholds for Q5T1, Q5T2, Q5T3 to 80%, 90% and 100%
cr22-vss-core(config-if-range)# rcv-queue random-detect min-threshold 6 70 80 90 100
! Sets WRED min thresholds for Q6T1, Q6T2, Q6T3 to 70 %, 80% and 90%
cr22-vss-core(config-if-range)# rcv-queue random-detect max-threshold 6 80 90 100 100
! Sets WRED max thresholds for Q6T1, Q6T2, Q6T3 to 80%, 90% and 100%
cr22-vss-core(config-if-range)# rcv-queue random-detect min-threshold 7 60 70 80 90
! Sets WRED min thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 60%, 70%, 80% and 90%, respectively
cr22-vss-core(config-if-range)# rcv-queue random-detect max-threshold 7 70 80 90 100
! Sets WRED max thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 70%, 80%, 90% and 100%, respectively

! This section configures the DSCP-to-Receive-Queue mappings
cr22-vss-core(config-if-range)# rcv-queue dscp-map 1 1 8
! Maps CS1 (Scavenger) to Q1T1
cr22-vss-core(config-if-range)# rcv-queue dscp-map 2 1 0
! Maps DF (Best Effort) to Q2T1
cr22-vss-core(config-if-range)# rcv-queue dscp-map 3 1 14
! Maps AF13 (Bulk Data-Drop Precedence 3) to Q3T1
cr22-vss-core(config-if-range)# rcv-queue dscp-map 3 2 12
! Maps AF12 (Bulk Data-Drop Precedence 2) to Q3T2
cr22-vss-core(config-if-range)# rcv-queue dscp-map 3 3 10
! Maps AF11 (Bulk Data-Drop Precedence 1) to Q3T3
cr22-vss-core(config-if-range)# rcv-queue dscp-map 4 1 22
! Maps AF23 (Transactional Data-Drop Precedence 3) to Q4T1
cr22-vss-core(config-if-range)# rcv-queue dscp-map 4 2 20
! Maps AF22 (Transactional Data-Drop Precedence 2) to Q4T2
cr22-vss-core(config-if-range)# rcv-queue dscp-map 4 3 18
! Maps AF21 (Transactional Data-Drop Precedence 1) to Q4T3

```



```

cr22-vss-core(config-if-range)# rcv-queue dscp-map 5 1 30
! Maps AF33 (Multimedia Streaming-Drop Precedence 3) to Q5T1
cr22-vss-core(config-if-range)# rcv-queue dscp-map 5 2 28
! Maps AF32 (Multimedia Streaming-Drop Precedence 2) to Q5T2
cr22-vss-core(config-if-range)# rcv-queue dscp-map 5 3 26
! Maps AF31 (Multimedia Streaming-Drop Precedence 1) to Q5T3
cr22-vss-core(config-if-range)# rcv-queue dscp-map 6 1 38
! Maps AF43 (Multimedia Conferencing-Drop Precedence 3) to Q6T1
cr22-vss-core(config-if-range)# rcv-queue dscp-map 6 2 36
! Maps AF42 (Multimedia Conferencing-Drop Precedence 2) to Q6T2
cr22-vss-core(config-if-range)# rcv-queue dscp-map 6 3 34
! Maps AF41 (Multimedia Conferencing-Drop Precedence 1) to Q6T3
cr22-vss-core(config-if-range)# rcv-queue dscp-map 7 1 16
! Maps CS2 (Network Management) to Q7T1
cr22-vss-core(config-if-range)# rcv-queue dscp-map 7 2 24
! Maps CS3 (Signaling) to Q7T2
cr22-vss-core(config-if-range)# rcv-queue dscp-map 7 3 48
! Maps CS6 (Internetwork Control) to Q7T3
cr22-vss-core(config-if-range)# rcv-queue dscp-map 7 4 56
! Maps CS7 (Network Control) to Q7T4
cr22-vss-core(config-if-range)# rcv-queue dscp-map 8 4 32 40 46
! Maps CS4 (Realtime Interactive), CS5 (Broadcast Video),
! and EF (VoIP) to Q8

cr23-VSS-Core#show queueing interface Ten1/1/2 | begin Rx
Queueing Mode In Rx direction: mode-dscp
Receive queues [type = 8q4t]:
Queue Id      Scheduling  Num of thresholds
-----
      01          WRR              04
      02          WRR              04
      03          WRR              04
      04          WRR              04
      05          WRR              04
      06          WRR              04
      07          WRR              04
      08          WRR              04

WRR bandwidth ratios:   1[queue 1]  25[queue 2]   4[queue 3]  10[queue 4]  10[queue
5]  10[queue 6]  10[queue 7]  30[queue 8]
queue-limit ratios:    10[queue 1]  25[queue 2]  10[queue 3]  10[queue 4]  10[queue
5]  10[queue 6]  10[queue 7]  15[queue 8]

queue tail-drop-thresholds
-----
1      70[1] 80[2] 90[3] 100[4]
2      100[1] 100[2] 100[3] 100[4]
3      100[1] 100[2] 100[3] 100[4]
4      100[1] 100[2] 100[3] 100[4]
5      100[1] 100[2] 100[3] 100[4]
6      100[1] 100[2] 100[3] 100[4]
7      100[1] 100[2] 100[3] 100[4]
8      100[1] 100[2] 100[3] 100[4]

queue random-detect-min-thresholds
-----
1      80[1] 100[2] 100[3] 100[4]
2      80[1] 100[2] 100[3] 100[4]
3      70[1] 80[2] 90[3] 100[4]
4      70[1] 80[2] 90[3] 100[4]
5      70[1] 80[2] 90[3] 100[4]
6      70[1] 80[2] 90[3] 100[4]
7      60[1] 70[2] 80[3] 90[4]
8      100[1] 100[2] 100[3] 100[4]

```

```

queue random-detect-max-thresholds
-----
 1    100[1] 100[2] 100[3] 100[4]
 2    100[1] 100[2] 100[3] 100[4]
 3    80[1] 90[2] 100[3] 100[4]
 4    80[1] 90[2] 100[3] 100[4]
 5    80[1] 90[2] 100[3] 100[4]
 6    80[1] 90[2] 100[3] 100[4]
 7    70[1] 80[2] 90[3] 100[4]
 8    100[1] 100[2] 100[3] 100[4]

WRED disabled queues:      8

...
queue thresh dscp-map
-----
47 1      1      1 2 3 4 5 6 7 8 9 11 13 15 17 19 21 23 25 27 29 31 33 39 41 42 43 44 45
   1      2
   1      3
   1      4
   2      1      0
   2      2
   2      3
   2      4
   3      1      14
   3      2      12
   3      3      10
   3      4
   4      1      22
   4      2      20
   4      3      18
   4      4
   5      1      30 35 37
   5      2      28
   5      3      26
   5      4
   6      1      38 49 50 51 52 53 54 55 57 58 59 60 61 62 63
   6      2      36
   6      3      34
   6      4
   7      1      16
   7      2      24
   7      3      48
   7      4      56
   8      1
   8      2
   8      3
   8      4      32 40 46

...
Packets dropped on Receive:
  BPDUs packets:  0

queue                dropped  [dscp-map]
-----
41 1                  0  [1 2 3 4 5 6 7 8 9 11 13 15 17 19 21 23 25 27 29 31 33 39
42 2                  0  [0 ]
43 3                  0  [14 12 10 ]
44 4                  0  [22 20 18 ]
45 5                  0  [30 35 37 28 26 ]
46 6                  0  [38 49 50 51 52 53 54 55 57 58 59 60 61 62 63 36 34 ]
47 7                  0  [16 24 48 56 ]
   8                  0  [32 40 46 ]

```

## Implementing Network Core Egress QoS

The QoS implementation of egress traffic towards network edge devices on access-layer switches are much simplified compared to ingress traffic which requires stringent QoS policies to provide differentiated services and network bandwidth protection. Unlike the Ingress QoS model, the egress QoS model must provide optimal queuing policies for each class and sets the drop thresholds to prevent network congestion and an application performance impact. With egress queuing in DSCP mode, the Cisco Catalyst switching platforms and linecards are bounded by a limited number of egress hardware queues.

### Catalyst 3750-X and 4500-E

The configuration and implementation guideline for egress QoS on Catalyst 3750-X StackWise and Catalyst 4500-E with Sup6E and Sup6L-E in distribution and access-layer roles remains consistent. All conformed traffic marked with DSCP values must be manually assigned to each egress queue based on a four class-of-service QoS model. Refer to the [“Implementing Access-Layer Egress QoS”](#) section on page 2-99 for the deployment details.

### Catalyst 6500-E – VSS

The Cisco Catalyst 6500-E in VSS mode operates in a centralized management mode but uses a distributed forwarding architecture. The Policy Feature Card (PFC) on active and hot-standby is functional on both nodes and is independent of the virtual-switch role. Like ingress queuing, the network administrator must implement egress queuing on each of the member-links of the Layer 2 or Layer 3 MEC. The egress queuing model on the Catalyst 6500-E is based on linecard type and its capabilities, when deploying Catalyst 6500-E in VSS mode only the WS-67xx series 1G/10G linecard with daughter card – CFC or DFC3/DFC3CXL is supported.

[Table 2-10](#) describes the deployment guidelines for the Catalyst 6500-E Series linecard module in the campus distribution and core layer network. In the solutions lab, the WS-6724-SFP and WS-6708-10GE was validated in the campus distribution and core layers. Both modules support different egress queuing models, this sub-section will provide deployment guidelines for both module types.

**Table 2-10 Catalyst 6500-E Switch Module Egress Queuing Architecture**

Switch Module	Daughter Card	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Egress Buffer Size
WS-6724-SFP	CFC or DFC3	1P3Q8T	DWRR	1.3 MB	1.2 MB
WS-6704-10GE	CFC	1P7Q8T	DWRR	16 MB	14 MB
	DFC3				
WS-6708-10GE	DFC3	1P7Q4T	DWRR	198 MB	90 MB
WS-6716-10GE	DFC3	1P7Q8T (Oversubscription and Perf. Mode)	SRR	198 MB <sup>1</sup>	90 MB <sup>1</sup>
				91 MB <sup>2</sup>	1 MB <sup>2</sup>

1. Per Port Capacity in Performance Mode

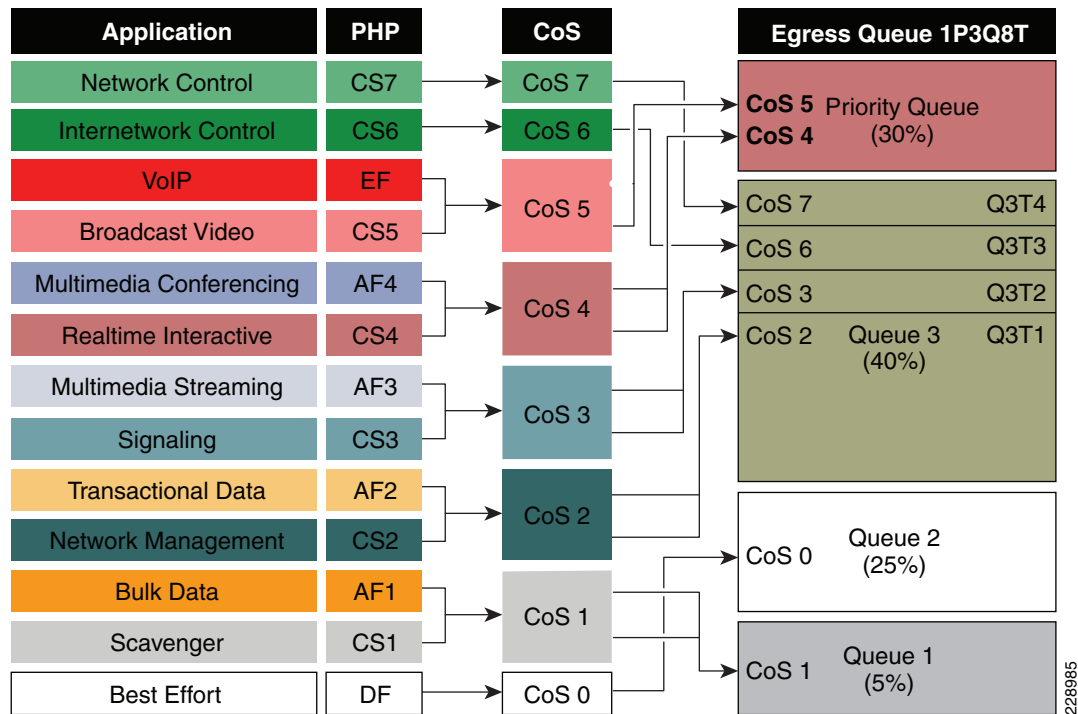
2. Per Port Capacity in Oversubscription Mode

### WS-6724-SFP – 1P3Q8T Egress Queuing Model

On the WS-6724-SFP module the egress queuing functions on per physical port basis and independent of link-layer and above protocols settings, these functions remain consistent when the physical port is deployed in standalone or bundled into an EtherChannel. Each 1G physical port support 4 egress queues

with default CoS based on the transmit side. This module is a cost-effective 1G non-blocking high speed network module but does not provide deep application granularity based on different DSCP markings. It does not have the flexibility to use various class-of-service egress queue for applications. Campus LAN QoS consolidation to a 4 class model occurs on the physical paths that connects to the WAN or Internet Edge routers, which forwards traffic across a private WAN or the Internet. Deploying the WS-6724-SFP module in 4 class model would be recommended in that design. Figure 2-62 illustrates 1P3Q8T egress queuing model to be applied on Catalyst 6500-E – WS-6724-SF module.

**Figure 2-62 1P3Q8T Egress Queuing Model**



The following corresponding 1P3Q8T egress queuing configuration must be applied on each member-links of MEC.

- Catalyst 6500-E VSS (Distribution and Core)

```
cr23-vss-core(config)#interface range GigabitEthernet 1/2/1-24 , Gi2/2/1 - 24
cr23-vss-core(config-if-range)# wrr-queue queue-limit 20 25 40
! Allocates 20% of the buffers to Q1, 25% to Q2 and 40% to Q3
cr23-vss-core(config-if-range)# priority-queue queue-limit 15
! Allocates 15% of the buffers to the PQ
cr23-vss-core(config-if-range)# wrr-queue bandwidth 5 25 40
! Allocates 5% BW to Q1, 25% BW to Q2 and 30% BW to Q3

! This section enables WRED on Queues 1 through 3
cr23-vss-core(config-if-range)# wrr-queue random-detect 1
! Enables WRED on Q1
cr23-vss-core(config-if-range)# wrr-queue random-detect 2
! Enables WRED on Q2
cr23-vss-core(config-if-range)# wrr-queue random-detect 3
! Enables WRED on Q3

! This section configures WRED thresholds for Queues 1 through 3
```

```

cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 1 100 100 100 100
100 100 100 100
! Sets all WRED max thresholds on Q1 to 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 1 80 100 100 100 100
100 100 100
! Sets Q1T1 min WRED threshold to 80%; all others set to 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 2 100 100 100 100
100 100 100 100
! Sets all WRED max thresholds on Q2 to 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 2 80 100 100 100 100
100 100 100
! Sets Q2T1 min WRED threshold to 80%; all others set to 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 3 70 80 90 100 100
100 100 100
! Sets Q3T1 max WRED threshold to 70%; Q3T2 max WRED threshold to 80%;
! Sets Q3T3 max WRED threshold to 90%; Q3T4 max WRED threshold to 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 3 60 70 80 90 100
100 100 100
! Sets Q3T1 min WRED threshold to 60%; Q3T2 min WRED threshold to 70%;
! Sets Q3T3 min WRED threshold to 80%; Q3T4 min WRED threshold to 90%

! This section configures the CoS-to-Queue/Threshold mappings
cr23-vss-core(config-if-range)# wrr-queue cos-map 1 1 1
! Maps CoS 1 (Scavenger and Bulk Data) to Q1T1
cr23-vss-core(config-if-range)# wrr-queue cos-map 2 1 0
! Maps CoS 0 (Best Effort) to Q2T1
cr23-vss-core(config-if-range)# wrr-queue cos-map 3 1 2
! Maps CoS 2 (Network Management and Transactional Data) to Q3T1
cr23-vss-core(config-if-range)# wrr-queue cos-map 3 2 3
! Maps CoS 3 (Signaling and Multimedia Streaming) to Q3T2
cr23-vss-core(config-if-range)# wrr-queue cos-map 3 3 6
! Maps CoS 6 (Internetwork Control) to Q3T3
cr23-vss-core(config-if-range)# wrr-queue cos-map 3 4 7
! Maps CoS 7 (Network Control) to Q3T4
cr23-vss-core(config-if-range)# priority-queue cos-map 1 4 5
! Maps CoS 4 (Realtime Interactive and Multimedia Conferencing) to PQ
! Maps CoS 5 (VoIP and Broadcast Video) to the PQ

cr23-VSS-Core#show queueing interface GigabitEthernet 1/2/1
Interface GigabitEthernet1/2/1 queueing strategy: Weighted Round-Robin
Port QoS is enabled
Trust boundary disabled

Trust state: trust DSCP
Extend trust state: not trusted [COS = 0]
Default COS is 0
Queueing Mode In Tx direction: mode-cos
Transmit queues [type = lp3q8t]:
Queue Id      Scheduling  Num of thresholds
-----
    01         WRR             08
    02         WRR             08
    03         WRR             08
    04         Priority          01

WRR bandwidth ratios:      5[queue 1]  25[queue 2]  40[queue 3]
queue-limit ratios:       20[queue 1]  25[queue 2]  40[queue 3]  15[Pri Queue]

queue tail-drop-thresholds
-----
1      70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2      70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]

```

```

queue random-detect-min-thresholds
-----
 1    80[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
 2    80[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
 3    60[1] 70[2] 80[3] 90[4] 100[5] 100[6] 100[7] 100[8]

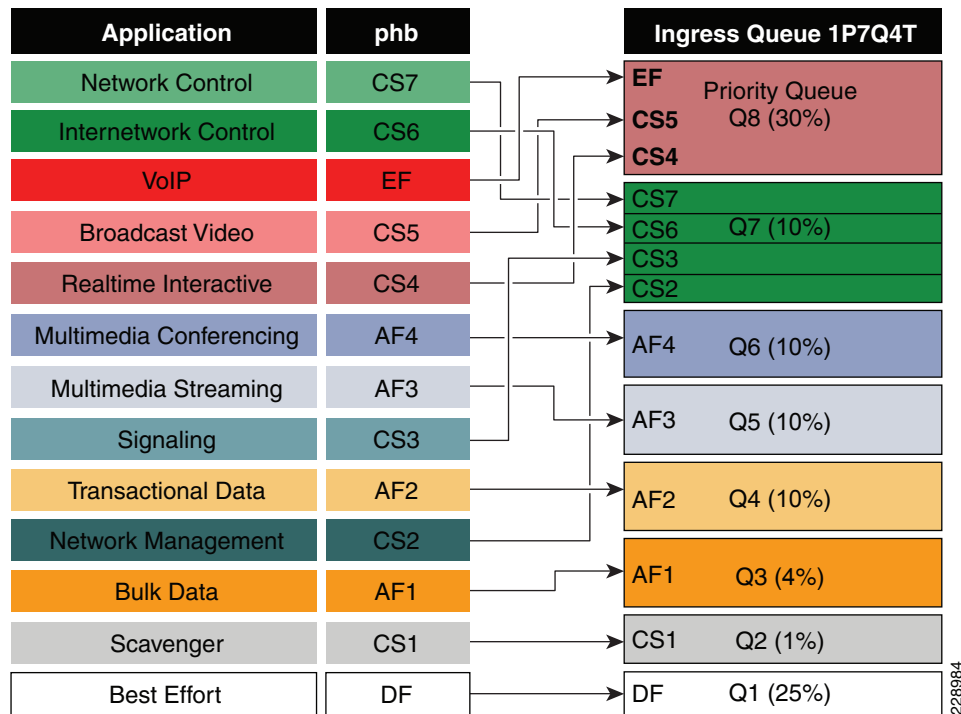
queue random-detect-max-thresholds
-----
 1   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
 2   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
 3   70[1] 80[2] 90[3] 100[4] 100[5] 100[6] 100[7] 100[8]

WRED disabled queues:
queue thresh cos-map
-----
 1      1      1
 1      2
 1      3
 1      4
 1      5
 1      6
 1      7
 1      8
 2      1      0
 2      2
 2      3
 2      4
 2      5
 2      6
 2      7
 2      8
 3      1      2
 3      2      3
 3      3      6
 3      4      7
 3      5
 3      6
 3      7
 3      8
 4      1      4 5
...

```

### WS-6708-10GE and WS-6716-10GE – 1P7Q4T Egress Queuing Model

The hardware design of the next-generation 10G linecards are designed with advanced ASICs and higher capacity to ensure the campus backbone of large enterprise networks are ready for future. Both modules support DSCP based on the 8 queue model to deploy flexible and scalable QoS in the campus core. With 8-egress queue support the WS-6708-10G and WS-6716-10G modules increased application granularity based on various DSCP markings are done at the network edge. [Figure 2-63](#) illustrates DSCP-based 1P7Q4T egress queuing model.

**Figure 2-63 P7Q4T Egress Queuing Model**

The following corresponding 1P7Q4T egress queuing configuration must be applied on each member-links of MEC.

- Catalyst 6500-E VSS (Distribution and Core)

```
cr23-vss-core(config)#interface range TenGigabitEthernet 1/1/2 - 8 , 2/1/2 - 8
cr23-vss-core(config-if-range)# wrr-queue queue-limit 10 25 10 10 10 10 10
! Allocates 10% of the buffers to Q1, 25% to Q2, 10% to Q3, 10% to Q4,
! Allocates 10% to Q5, 10% to Q6 and 10% to Q7
cr23-vss-core(config-if-range)# wrr-queue bandwidth 1 25 4 10 10 10 10
! Allocates 1% BW to Q1, 25% BW to Q2, 4% BW to Q3, 10% BW to Q4,
! Allocates 10% BW to Q5, 10% BW to Q6 and 10% BW to Q7
cr23-vss-core(config-if-range)# priority-queue queue-limit 15
! Allocates 15% of the buffers to the PQ

! This section enables WRED on Queues 1 through 7
cr23-vss-core(config-if-range)# wrr-queue random-detect 1
! Enables WRED on Q1
cr23-vss-core(config-if-range)# wrr-queue random-detect 2
! Enables WRED on Q2
cr23-vss-core(config-if-range)# wrr-queue random-detect 3
! Enables WRED on Q3
cr23-vss-core(config-if-range)# wrr-queue random-detect 4
! Enables WRED on Q4
cr23-vss-core(config-if-range)# wrr-queue random-detect 5
! Enables WRED on Q5
cr23-vss-core(config-if-range)# wrr-queue random-detect 6
! Enables WRED on Q6
cr23-vss-core(config-if-range)# wrr-queue random-detect 7
! Enables WRED on Q7

! This section configures WRED thresholds for Queues 1 through 7
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 1 100 100 100 100
```

```

! Sets all WRED max thresholds on Q1 to 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 1 80 100 100 100
! Sets Q1T1 min WRED threshold to 80%
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 2 100 100 100 100
! Sets all WRED max thresholds on Q2 to 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 2 80 100 100 100
! Sets Q2T1 min WRED threshold to 80%
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 3 80 90 100 100
! Sets WRED max thresholds for Q3T1, Q3T2, Q3T3 to 80%, 90% and 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 3 70 80 90 100
! Sets WRED min thresholds for Q3T1, Q3T2, Q3T3 to 70 %, 80% and 90%

cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 4 70 80 90 100
! Sets WRED min thresholds for Q4T1, Q4T2, Q4T3 to 70 %, 80% and 90%
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 4 80 90 100 100
! Sets WRED max thresholds for Q4T1, Q4T2, Q4T3 to 80%, 90% and 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 5 70 80 90 100
! Sets WRED min thresholds for Q5T1, Q5T2, Q5T3 to 70 %, 80% and 90%
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 5 80 90 100 100
! Sets WRED max thresholds for Q5T1, Q5T2, Q5T3 to 80%, 90% and 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 6 70 80 90 100
! Sets WRED min thresholds for Q6T1, Q6T2, Q6T3 to 70 %, 80% and 90%
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 6 80 90 100 100
! Sets WRED max thresholds for Q6T1, Q6T2, Q6T3 to 80%, 90% and 100%
cr23-vss-core(config-if-range)# wrr-queue random-detect min-threshold 7 60 70 80 90
! Sets WRED min thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 60%, 70%, 80% and 90%, respectively
cr23-vss-core(config-if-range)# wrr-queue random-detect max-threshold 7 70 80 90 100
! Sets WRED max thresholds for Q7T1, Q7T2, Q7T3 and Q7T4
! to 70%, 80%, 90% and 100%, respectively

! This section configures the DSCP-to-Queue/Threshold mappings
cr23-vss-core(config-if-range)# wrr-queue dscp-map 1 1 8
! Maps CS1 (Scavenger) to Q1T1
cr23-vss-core(config-if-range)# wrr-queue dscp-map 2 1 0
! Maps DF (Best Effort) to Q2T1
cr23-vss-core(config-if-range)# wrr-queue dscp-map 3 1 14
! Maps AF13 (Bulk Data-Drop Precedence 3) to Q3T1
cr23-vss-core(config-if-range)# wrr-queue dscp-map 3 2 12
! Maps AF12 (Bulk Data-Drop Precedence 2) to Q3T2
cr23-vss-core(config-if-range)# wrr-queue dscp-map 3 3 10
! Maps AF11 (Bulk Data-Drop Precedence 1) to Q3T3
cr23-vss-core(config-if-range)# wrr-queue dscp-map 4 1 22
! Maps AF23 (Transactional Data-Drop Precedence 3) to Q4T1
cr23-vss-core(config-if-range)# wrr-queue dscp-map 4 2 20
! Maps AF22 (Transactional Data-Drop Precedence 2) to Q4T2
cr23-vss-core(config-if-range)# wrr-queue dscp-map 4 3 18
! Maps AF21 (Transactional Data-Drop Precedence 1) to Q4T3
cr23-vss-core(config-if-range)# wrr-queue dscp-map 5 1 30
! Maps AF33 (Multimedia Streaming-Drop Precedence 3) to Q5T1
cr23-vss-core(config-if-range)# wrr-queue dscp-map 5 2 28
! Maps AF32 (Multimedia Streaming-Drop Precedence 2) to Q5T2
cr23-vss-core(config-if-range)# wrr-queue dscp-map 5 3 26
! Maps AF31 (Multimedia Streaming-Drop Precedence 1) to Q5T3
cr23-vss-core(config-if-range)# wrr-queue dscp-map 6 1 38
! Maps AF43 (Multimedia Conferencing-Drop Precedence 3) to Q6T1
cr23-vss-core(config-if-range)# wrr-queue dscp-map 6 2 36
! Maps AF42 (Multimedia Conferencing-Drop Precedence 2) to Q6T2
cr23-vss-core(config-if-range)# wrr-queue dscp-map 6 3 34
! Maps AF41 (Multimedia Conferencing-Drop Precedence 1) to Q6T3
cr23-vss-core(config-if-range)# wrr-queue dscp-map 7 1 16
! Maps CS2 (Network Management) to Q7T1
cr23-vss-core(config-if-range)# wrr-queue dscp-map 7 2 24

```



```

! Maps CS3 (Signaling) to Q7T2
cr23-vss-core(config-if-range)# wrr-queue dscp-map 7 3 48
! Maps CS6 (Internetwork Control) to Q7T3
cr23-vss-core(config-if-range)# wrr-queue dscp-map 7 4 56
! Maps CS7 (Network Control) to Q7T4
cr23-vss-core(config-if-range)# priority-queue dscp-map 1 32 40 46
! Maps CS4 (Realtime Interactive), CS5 (Broadcast Video),
! and EF (VoIP) to the PQ

```

**Note**

Due to the default WRED threshold settings, at times the maximum threshold needs to be configured before the minimum (as is the case on queues 1 through 3 in the example above); at other times, the minimum threshold needs to be configured before the maximum (as is the case on queues 4 through 7 in the example above).

## High-Availability in LAN Network Design

Network reliability and availability is not a new demand, but is well planned during the early network design phase. To prevent a catastrophic network failure during an unplanned network outage event, it is important to identify network fault domains and define rapid recovery plans to minimize the application impact during minor and major network outage conditions.

Because every tier of the LAN network design can be classified as a fault domain, deploying redundant systems can be effective. However, this introduces a new set of challenges, such as higher cost and the added complexity of managing more systems. Network reliability and availability can be simplified using several Cisco high availability technologies that offer complete failure transparency to the end users and applications during planned or unplanned network outages.

Cisco high availability technologies can be deployed based on critical versus non-critical platform roles in the network. Some of the high availability techniques can be achieved with the LAN network design inherent within the medium enterprise network design, without making major network changes. However, the critical network systems that are deployed in the main campus that provide global connectivity may require additional hardware and software components to provide non-stop communications. The following three major resiliency requirements encompass most of the common types of failure conditions; depending on the LAN design tier, the resiliency option appropriate to the role and network service type must be deployed:

- *Network resiliency*—Provides redundancy during physical link failures, such as fiber cut, bad transceivers, incorrect cabling, and so on.
- *Device resiliency*—Protects the network during abnormal node failure triggered by hardware or software, such as software crashes, a non-responsive supervisor, and so on.
- *Operational resiliency*—Enables resiliency capabilities to the next level, providing complete network availability even during planned network outage conditions, using In Service Software Upgrade (ISSU) features.

## Medium Enterprise High-Availability Framework

Independent of the business function, the network architects builds strong, scalable, and resilient next-generation IP network. Networks that are built on these three fundamentals, offers high availability to use network as a core platform that enables flexibility to overlay advanced and emerging technologies and provide non-stop network communications. The medium enterprise campus network must be build based on same fundamentals that can provide constant “on” network service for uninterrupted business operations and protects campus physical security and assets.


Network faults domains in this reference architecture are identifiable but the failure conditions within the domains are un-predicted. Improper network design or non-resilient network systems can experience higher number of faults that not only degrades user experience but may severely impact application performance and may not capture the critical physical security video information. For example failure of 1-Gigabit Ethernet backbone connection for 10 seconds can the drop network information for more than 1Gig, which may include critical medium enterprise data or video surveillance captured data. The fault levels can range from network interruption to disaster, which can be triggered by system, human, or even by nature. Network failures can be classified in one of the following two ways:

- *Planned Failure*—Planned network outage occurs when any network systems is administratively planned to disable in the network for scheduled event (i.e., software upgrade etc.).
- *Unplanned Failure*—Any unforeseen failures of network elements can be considered as unplanned failure. Such failures can include internal faults in the network device caused by hardware or software malfunctions which includes software crash, linecard, or link transceiver failures conditions.

## Baselining Campus High Availability

Typical application response time is in milliseconds when the campus network is build with high speed backbone connection and is in fully-operational state. When constantly working in deterministic network response time environment the learning and work practice of end-users is rapid; however, during abnormal network failure causing traffic loss, congestion and application retries will impact the performance and alerts the user about the network faults. During the major network fault event, user determines network connection problem based on routine experience even before an application protocols determines connection problem (i.e., slow internet browsing response time). Protocol-based delayed failure detection are intentional, they are designed to minimize overall productivity impact and allows network to gracefully adjust and recover during minor failure conditions. Every protocol operation is different in the network; while the retries for non-critical data traffic is acceptable the applications running in real-time may not. Figure 2-64 provides a sample real-time VoIP application in campus network and sequence of user experience in different phases during minor and major unplanned network outage:

**Figure 2-64** VoIP Impact During Minor and Major Network Outage



This high availability framework is based on the three major resiliency strategies to solve a wide-range of planned and unplanned network outage types described in the previous section. Several high availability technologies must be deployed at each layer to provide higher network availability and

rapid recovery during failure conditions, to prevent communication failure or degraded network-wide application performance. (See [Figure 2-65](#).)

**Figure 2-65 High-Availability Goals, Strategy, and Technologies**

Resilient Goal	Network Service Availability		
Resilient Strategies	Network Resiliency	Device Resiliency	Operational Resiliency
Resilient Technologies	EtherChannel/MEC UDLD IP Event Dampening	NSF/SSO Stack Wise	ISSU eFSU

228500

## Network Resiliency Overview

The most common network fault occurrence in the LAN network is a link failure between two systems. Link failures can be caused by issues such as a fiber cut, miswiring, linecard module failure and so on. In the modular platform design the redundant parallel physical links between distributed modules between two systems reduces fault probabilistic and can increase network availability. It is important to remember how multiple parallel paths between two systems also changes overall higher layer protocols construct the adjacency and loop-free forwarding topology.

Deploying redundant parallel paths in the recommended medium enterprise LAN design by default develops a non-optimal topology that keeps the network underutilized and requires protocol-based network recovery. In the same network design, the routed access model eliminates such limitations and enables the full load balancing capabilities to increase bandwidth capacity and minimize the application impact during a single path failure. To develop a consistent network resiliency service in the centralized main and remote campus sites, the following basic principles apply:

- Deploying redundant parallel paths are the basic requirement to employ network resiliency at any tier. It is critical to simplify the control plane and forwarding plane operation by bundling all physical paths into a single logical bundled interface (EtherChannel). Implement a defense-in-depth approach to failure detection and recovery mechanisms. An example of this is configuring the UniDirectional Link Detection (UDLD) protocol, which uses a Layer 2 keep-alive to test that the switch-to-switch links are connected and operating correctly, and acts as a backup to the native Layer 1 unidirectional link detection capabilities provided by 802.3z and 802.3ae standards. UDLD is not an EtherChannel function; it operates independently over each individual physical port at Layer 2 and remains transparent to the rest of the port configuration. Therefore, UDLD can be deployed on ports implemented in Layer 2 or Layer 3 modes.
- Ensure that the network design is self-stabilizing. Hardware or software errors may cause ports to flap, which creates false alarms and destabilizes the network topology. Implementing route summarization advertises a concise topology view to the network, which prevents core network instability. However, within the summarized boundary, the flood may not be protected. Deploy IP event dampening as an tool to prevent the control and forwarding plane impact caused by physical topology instability.

These principles are intended to be a complementary part of the overall structured modular design approach to the campus design, and serve primarily to reinforce good resilient design practices.

## Device Resiliency Overview

Another major component of an overall campus high availability framework is providing device or node level protection that can be triggered during any type of abnormal internal hardware or software process within the system. Some of the common internal failures are a software-triggered crash, power outages, line card failures, and so on. LAN network devices can be considered as a single-point-of-failure and are considered to be major failure condition because the recovery type may require a network administrator to mitigate the failure and recover the system. The network recovery time can remain undeterministic, causing complete or partial network outage, depending on the network design.

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

### Redundant Power System

Redundant power supplies for network systems protect against power outages, power supply failures, and so on. It is important not only to protect the internal network system but also the endpoints that rely on power delivery over the Ethernet network. Redundant power systems can be deployed in the two following configuration modes:

- *Modular switch*—Dual power supplies can be deployed in modular switching platforms such as the Cisco Catalyst 6500-E and 4500-E Series platforms. By default, the power supply operates in redundant mode, offering the 1+1 redundant option. Overall power capacity planning must be done to dynamically allow for network growth. Lower power supplies can be combined to allocate power to all internal and external resources, but may not be able to offer power redundancy.
- *Fixed configuration switch*—Depending on the Catalyst switch capability the fixed configuration switches offers wide range of power redundancy options includes the latest innovation Cisco StackPower in Catalyst 3750-X series platform. To prevent network outage on fixed configuration the Catalyst switches they must be deployed with Cisco StackPower technology, an internal redundant power supplies on Catalyst 3560-X and use Cisco RPS 2300 external power supplies solution on Catalyst 2960-S Series switches. A single Cisco RPS 2300 power supply uses a modular power supply and fan for flexibility, and can deliver power to multiple switches. Deploying an internal and external power supply solution protects critical access layer switches during power outages, and provides completes fault transparency and constant network availability.

### Redundant Control Plane

Device or node resiliency in modular Cisco Catalyst 6500-E/4500-E platforms and Cisco StackWise provides a 1+1 redundancy option with enterprise-class high availability and deterministic network recovery time. The following subsections provide high availability design details, as well as graceful network recovery techniques that do not impact the control plane and provide constant forwarding capabilities during failure events.

## Stateful Switchover

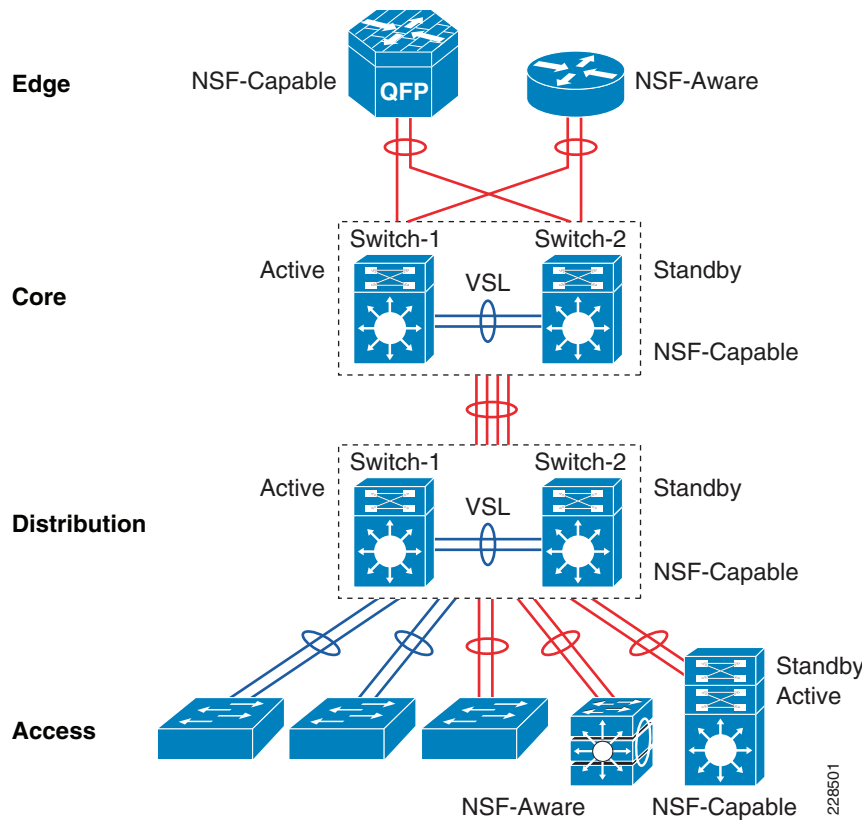
The stateful switchover (SSO) capability in modular switching platforms such as the Cisco Catalyst 4500 and 6500 provides complete carrier-class high availability in the campus network. Cisco recommends distribution and core layer design model be the center point of the entire enterprise communication network. Deploying redundant supervisors in the mission-critical distribution and core system provides non-stop communication throughout the network. To provide 99.999 percent service availability in the access layer, the Catalyst 4500 must be equipped with redundant supervisors to critical endpoints, such as Cisco TelePresence.

Cisco StackWise is an low-cost solution to provide device-level high availability. Cisco StackWise is designed with unique hardware and software capabilities that distribute, synchronize, and protect common forwarding information across all member switches in a stack ring. During master switch failure, the new master switch re-election remains transparent to the network devices and endpoints. Deploying Cisco StackWise according to the recommended guidelines protects against network interruption, and recovers the network in sub-seconds during master switch re-election.

Bundling SSO with NSF capability and the awareness function allows the network to operate without errors during a primary supervisor module failure. Users of realtime applications such as VoIP do not hang up the phone, and IP video surveillance cameras do not freeze.

## Non-Stop Forwarding

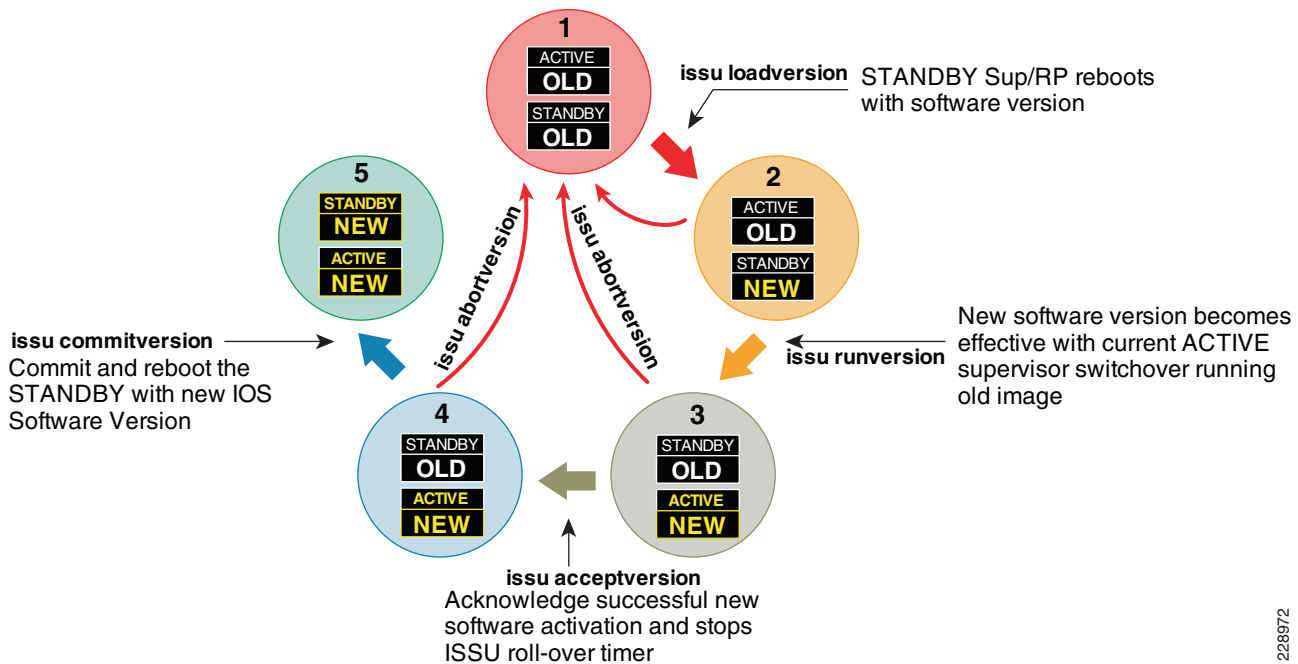
Cisco VSS and the single highly resilient-based campus system provides uninterrupted network availability using non-stop forwarding (NSF) without impacting end-to-end application performance. The Cisco VSS and redundant supervisor system is an NSF-capable platform; thus, every network device that connects to VSS or the redundant supervisor system must be NSF-aware to provide optimal resiliency. By default, most Cisco Layer 3 network devices are NSF-aware systems that operate in NSF helper mode for graceful network recovery. (See [Figure 2-66](#).)

**Figure 2-66 Medium Enterprise NSF/SSO Capable and Aware Systems**

## Operational Resiliency Overview

Designing the network to recover from failure events is only one aspect of the overall campus non-stop design. Converged network environments are continuing to move toward requiring true 7x24x365 availability. The medium enterprise LAN network is part of the backbone of the enterprise network and must be designed to enable standard operational processes, configuration changes, and software and hardware upgrades without disrupting network services.

The ability to make changes, upgrade software, and replace or upgrade hardware becomes challenging without a redundant system in the campus core. Upgrading individual devices without taking them out of service is similarly based on having internal component redundancy (such as with power supplies and supervisors), complemented with the system software capabilities. The Cisco Catalyst 4500-E, 6500-E and ASR 1000 series platform support realtime upgrade software in the campus. The Cisco In-Service Software Upgrade (ISSU) and Enhanced Fast Software Upgrade (eFSU) leverages NSF/SSO technology to provide continuous network availability while upgrading the critical systems that eliminates network services downtime planning and maintenance window. [Figure 2-67](#) demonstrates platform-independent Cisco IOS software upgrade flow process using ISSU technology.

**Figure 2-67 Cisco ISSU Software Process Cycle****Catalyst 4500—ISSU**

Full-image ISSU on the Cisco Catalyst 4500-E leverages dual redundant supervisors to allow for a full, in-place Cisco IOS upgrade, such as moving from IOS Release 12.2(53)SG to 12.2(53)SG1 for example. This leverages the NSF/SSO capabilities and unique uplink port capability to keep in operational and forwarding state even when supervisor module gets reset, such design helps in retaining bandwidth capacity while upgrading both supervisor modules at the cost of less than sub-second of traffic loss during a full Cisco IOS upgrade.

Having the ability to operate the campus as a non-stop system depends on the appropriate capabilities being designed-in from the start. Network and device level redundancy, along with the necessary software control mechanisms, guarantee controlled and fast recovery of all data flows following any network failure, while concurrently providing the ability to proactively manage the non-stop infrastructure.

**Catalyst 6500 VSS—eFSU**

A network upgrade requires planned network and system downtime. VSS offers unmatched network availability to the core. With the Enhanced Fast Software Upgrade (eFSU) feature, the VSS can continue to provide network services during the upgrade. With the eFSU feature, the VSS network upgrade remains transparent and hitless to the applications and end users. Because eFSU works in conjunction with NSF/SSO technology, the network devices can gracefully restore control and forwarding information during the upgrade process, while the bandwidth capacity operates at 50 percent and the data plane can converge within sub-seconds.

For a hitless software update, the ISSU process requires three sequential upgrade events for error-free software install on both virtual switch systems. Each upgrade event causes traffic to be re-routed to a redundant MEC path, causing sub-second traffic loss that does not impact realtime network applications, such as VoIP.

## Design Strategies for Network Survivability

The network reliability and availability is not a new demand, it is one of the critical integrated component that gets well planned during early network design phase. To prevent catastrophic network failure during un-planned network outage event, it is important to identify network fault domains and define rapid recovery plans to minimize the application impact during minor and major network outage conditions.

Each network tier can be classified as a fault domains, deploying redundant components and systems increases redundancy and load sharing capabilities. However, it introduces the new set of challenges – higher cost and complexities to manage more number of systems. Network reliability and availability can be simplified using several Cisco high-availability and virtual-system technologies like VSS offers complete failure transparency to the end-users and applications during planned or un-planned network outage conditions. Minor and major network failure are the broad terms that's includes several types of network faults that must be taken into consideration and implement the rapid recovery solution.

Cisco high-availability technologies can be deployed based on critical versus non-critical platform role in the network. Some of the high-availability techniques can be achieved with inherent campus network design without making major network changes; however, the critical network systems that is deployed in the center of the network to provide global connectivity may require additional hardware and software component to offer non-stop communication. The network survivability strategy can categorized in following three major resiliency requirements that can encompass most of the common types of failure conditions. Depending on the network system tier, role and network service type appropriate resilient option must be deployed. See [Table 2-11](#).

**Table 2-11 Medium Enterprise Network High Availability Strategy**

Platform	Role	Network Resiliency	Device Resiliency	Operational Efficiency
Catalyst 2960-S FlexStack	Access	EtherChannel <sup>1</sup> UDLD Dampening	RPS 2300 NSF-Aware	Cisco FlexStack
Catalyst 3560-X			Redundant Power Supplies	None. Standalone systems
Catalyst 3750-X				
Catalyst 3750ME	WAN Edge			
Catalyst 3750-X StackWise	Access		Cisco StackPower	Stackwise Plus
	Distribution		NSF-Capable and Aware	
Catalyst 4500-E	Access		Red. Power Supplies <sup>2</sup>	ISSU
	Distribution		Red. Linecard modules <sup>2</sup>	
	Core		Red. Supervisor modules <sup>3</sup>	
Catalyst 6500-E	Distribution		SSO/NSF Capable & Aware <sup>2</sup>	VSS
	Core			eFSU



**Table 2-11 Medium Enterprise Network High Availability Strategy (continued)**

ASR 1006	WAN Edge	EtherChannel Dampening	Red. Power Supplies Red. ESP modules Red. Route Processors SSO/NSF Capable & Aware	ISSU
ASR 1004	Internet Edge		Red. Power Supplies SSO/NSF Capable & Aware <sup>4</sup>	ISSU
Cisco ISR	PSTN Gateway		-	None. Standalone system

1. Redundant uplinks from each 3750-E member switch in Stack ring and 6500-E virtual-switch in VSS domain
2. Redundant power and hardware components from each 3750-E member switch in Stack ring and 6500-E virtual-switch in VSS domain
3. Redundant supervisor per VSS Domain (One per virtual-switch node basis). Starting 12.2(33)SX14 it is recommended to deploy redundant supervisor on each virtual-switch in a VSS domain.
4. Software based SSO redundancy

## Implementing Network Resiliency

The medium enterprise design guide recommends deploying a mix of hardware and software resiliency designed to address the most common campus LAN network faults and instabilities. It is important to analyze the network and the application impacts from a top-down level to adapt and implement the appropriate high availability solution for creating a resilient network. Implementing a resilient hardware and software design increases network resiliency and maintains the availability of all upper layer network services that are deployed in a medium enterprise campus network design.

### EtherChannel / Multi-Chassis EtherChannel

In a non-EtherChannel network environment, the network protocol requires fault detection, topology synchronization, and best-path recomputation to reroute traffic which requires variable time to restart the forwarding traffic. Conversely, EtherChannel or MEC network environments provide significant benefits in such conditions, as network protocol remains unaware of the topology changes and allows the hardware to self-recover from faults. Re-routing traffic over an alternate member-link of EtherChannel or MEC is based on minor system internal EtherChannel hash re-computations instead of an entire network topology re-computation. Hence an EtherChannel and MEC based network provides deterministic sub-second network recovery of minor to major network faults.

The design and implementation considerations for deploying diverse physical connectivity across redundant standalone systems and virtual-systems to create a single point-to-point logical EtherChannel is explained in the [“Designing EtherChannel Network” section on page 2-41](#).

### EtherChannel/MEC Network Recovery Analysis

The network recovery with EtherChannel and MEC is platform and diverse physical path dependent instead of Layer 2 or Layer 3 network protocol dependent. The medium enterprise campus LAN network design deploys EtherChannel and MEC throughout the network to develop a simplified single point-to-point network topology which does not build any parallel routing paths between any devices at any network tiers.


During individual member-link failures, the Layer 2 and Layer 3 protocols dynamically adjusts the metrics of the aggregated port-channel interfaces. Spanning-Tree updates the port-cost and Layer 3 routing protocols like EIGRP updates the composite metric or OSPF may change the interface cost. In

such events, the metric change will require minor update messages in the network and do not require end-to-end topology recomputation that impacts the overall network recovery process. Since the network topology remains intact during individual link failures, the re-computation to select alternate member-links in EtherChannel and MEC becomes locally significant on each end of the impacted EtherChannel neighbors. EtherChannel re-computation requires recreating new logical hash table and re-programming the hardware to re-route the traffic over the remaining available paths in the bundled interface. The Layer 2 or Layer 3 EtherChannel and MEC re-computation is rapid and network scale independent.

### Catalyst 6500-E VSS MEC Link Recovery Analysis

Several types of network faults can trigger link failures in the network (i.e., fiber pullout, GBIC failure, etc.). The network recovery remains consistent and deterministic in all network fault conditions. In standalone or non-virtual systems like Catalyst 2960-S or 4500-E, the EtherChannel recomputation is fairly easy as the alternate member-link resides within the system. However, with the distributed forwarding architecture in virtual-systems like Catalyst 6500-E VSS and Catalyst 3750-X StackWise Plus may require extra computation to select alternate member-link paths through its inter-chassis backplane interface—VSL or StackRing. Such designs still provides deterministic recovery, but with an additional delay to recompute a new forwarding path through the remote virtual-switch node. The link failure analysis chart with inter-chassis reroute in [Figure 2-68](#) summarizes several types of faults induced in large scale Cisco lab during developing this validated design guide.


**Figure 2-68 Catalyst 6500-E VSS Inter-Chassis MEC Link Recovery Analysis**



The medium enterprise LAN can be designed optimally for deterministic and bidirectional symmetric network recovery for unicast and multicast traffic. Refer to the [“Redundant Linecard Network Recovery Analysis”](#) section on [page 2-134](#) for intra-chassis recovery analysis with the same network faults tested in inter-chassis scenarios.

### Catalyst 4507R-E EtherChannel Link Recovery Analysis

In the medium enterprise campus reference design, a single Catalyst 4507R-E with redundant hardware components is deployed in the different campus LAN network tiers. A Cisco Catalyst 4507R-E can only be deployed in standalone mode with in-chassis supervisor and module redundancy. However, the traffic load balancing and rerouting across different EtherChannel member-links occurs within the local chassis. The centralized forwarding architecture in Catalyst 4500-Es can rapidly detect link failures and reprogram the hardware with new EtherChannel hash results. The test results in [Figure 2-69](#) confirm the deterministic and consistent network recovery during individual Layer 2/3 EtherChannel member-link failures.

**Figure 2-69 Catalyst 4507R-E EtherChannel Link Recovery Analysis**

228988

### Unidirectional Link Detection (UDLD)

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 auto-negotiation and UDLD are enabled together, the Layer 1 and Layer 2 detection methods work together to prevent physical and logical unidirectional connections and prevent malfunctioning of other protocols.

Copper media ports use Ethernet link pulses as a link monitoring tool and are not susceptible to unidirectional link problems. However, 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 (Recommended)*—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 and 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 the user or the switch can be configured for auto recovery within a specified interval of time.

The following illustrates a configuration example to implement the UDLD protocol:

```
cr22-6500-LB#config t
cr22-6500-LB(config)#interface range gi1/2/3 , gi2/2/3
cr22-6500-LB(config-if-range)#udld port
```

```
cr22-6500-LB#show udld neighbors
```

Port	Device Name	Device ID	Port ID	Neighbor State
Gi1/2/3	FD01328R0E2	1	Gi1/0/49	Bidirectional
Gi2/2/3	FD01328R0E2	1	Gi1/0/50	Bidirectional

## IP Event Dampening

Unstable physical network connectivity with poor signaling or loose connection may cause continuous port-flaps. When the medium enterprise network is not deployed using best practice guidelines to summarize the network boundaries at the aggregation layer, a single interface flap can severely impact stability and availability of the entire campus network. Route summarization is one technique used to isolate the fault domain and contain local network faults within the domain.

To ensure local network domain stability during to port-flaps, all Layer 3 interfaces can be implemented with IP Event Dampening. It uses the same fundamental principles as BGP dampening. Each time the Layer 3 interface flaps, IP dampening tracks and records the flap events. On multiple flaps, a logical penalty is assigned to the port and suppresses link status notifications to IP routing until the port becomes stable.

IP Event Dampening is a local specific function and does not have any signaling mechanism to communicate with remote systems. It can be implemented on each individual physical or logical Layer 3 interface—physical ports, SVI, or port-channels:

- Layer 3 Port-Channel

```
cr24-4507e-MB(config)#interface Port-Channel 1
cr24-4507e-MB(config-if)#no switchport
cr24-4507e-MB(config-if)#dampening
```

- Layer 2 Port-Channel

```
cr24-4507e-MB(config)#interface Port-Channel 15
cr24-4507e-MB(config-if)#switchport
cr24-4507e-MB(config-if)#dampening
```

- SVI Interface

```
cr24-4507e-MB(config)#interface range Vlan101 - 120
cr24-4507e-MB(config-if-range)#dampening
```

```
cr24-4507e-MB#show interface dampening
```

```
Vlan101
  Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV   MaxSTm   MaxP Restart
      3         0  FALSE      0       5    1000    2000     20  16000     0
...
TenGigabitEthernet3/1 Connected to cr23-VSS-Core
  Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV   MaxSTm   MaxP Restart
     10         0  FALSE      0       5    1000    2000     20  16000     0
...
Port-channel11 Connected to cr23-VSS-Core
  Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV   MaxSTm   MaxP Restart
      3         0  FALSE      0       5    1000    2000     20  16000     0
Port-channel15 Connected to cr24-2960-S-MB
  Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV   MaxSTm   MaxP Restart
      3         0  FALSE      0       5    1000    2000     20  16000     0
```

## Implementing Device Resiliency


Each device in the medium enterprise LAN and WAN network design is connected to a critical system or end-point to provide network connectivity and services for business operations. Like network resiliency, the device resiliency solves the problem by integrating redundant hardware components and software based solutions into single standalone or virtual systems. Depending on the platform architecture of the Cisco router or switch deployed in the campus network design, the device redundancy is divided into four major categories—Redundant Power Supplies, Redundant Line cards, Redundant Supervisor/RP, and Non-Stop Forwarding (NSF) with Stateful Switchover (SSO).

## Redundant Power

To provide non-stop network communication during power outages, critical network devices must be deployed with redundant power supplies. Network administrators must identify the network systems that provide network connectivity and services to mission critical servers. This would also include Layer 1 services like PoE to boot IP Phone and IP Video Surveillance Cameras for campus physical security and communications.

Depending on the Cisco platform design, the in-chassis power redundancy option allows flexibility to deploy dual power supplies into a single system. The next-generation borderless network ready Cisco Catalyst 3750-X introduces latest Cisco StackPower innovation that creates a global pool of power that can provide power load sharing and redundancy option. While the Cisco Catalyst 3560-X Series switches are designed to increase device resiliency with dual redundant power supplies and fans. The Catalyst platforms like the 2960 and 2960-S can be deployed with Cisco RPS 2300 for external power redundancy solution. Figure 2-70 provides complete power redundancy design and solution on various Cisco Catalyst switching platforms:

**Figure 2-70 Power Supply Redundancy Design**



The following configuration examples provide guidelines to deploy in-chassis and external power redundancy in the Catalyst switching platforms.

### 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. This is the 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 6 RPS ports to provide seamless power backup to critical access-layer switches in the campus network. Additional power resiliency can be added by deploying dual power supply to backup to two devices simultaneously. Cisco RPS 2300 can be provisioned for the 3750-E or 3560-E series switches through CLI:

### Catalyst 4500-E and 6500-E (In-Chassis Power Redundancy)

The Cisco Catalyst 4500-E and 6500-E Series modular platforms allocate power to several internal hardware components and external power devices like IP Phones, Wireless Access Points, etc. All the power allocation is assigned from the internal power supply. Dual power supplies in these systems can operate in two different modes as listed below:

- *Redundant Mode*—By default, power supplies operate in redundant mode offering a 1+1 redundant option. The system determines power capacity and the number of power supplies required based on the allocated power to all internal and external power components. Both power supplies must have sufficient power to allocate power to all the installed modules in order to operate in 1+1 redundant mode.

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

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

cr22-vss-core(config)#power redundancy-mode redundant switch 1
cr22-vss-core(config)#power redundancy-mode redundant switch 2

cr2-6500-vss#show power switch 1 | inc Switch|mode
Switch Number: 1
system power redundancy mode = redundant
```

```
cr2-6500-vss#show power switch 2 | inc Switch|mode
Switch Number: 2
system power redundancy mode = redundant
```

- **Combined mode**—If the system power requirement exceeds the single power supply capacity, then the network administrator can utilize both power supplies in combined mode to increase capacity. However it may not offer 1+1 power redundancy during a primary power supply failure event. The following global configuration will enable power redundancy mode to operate in combined mode:

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

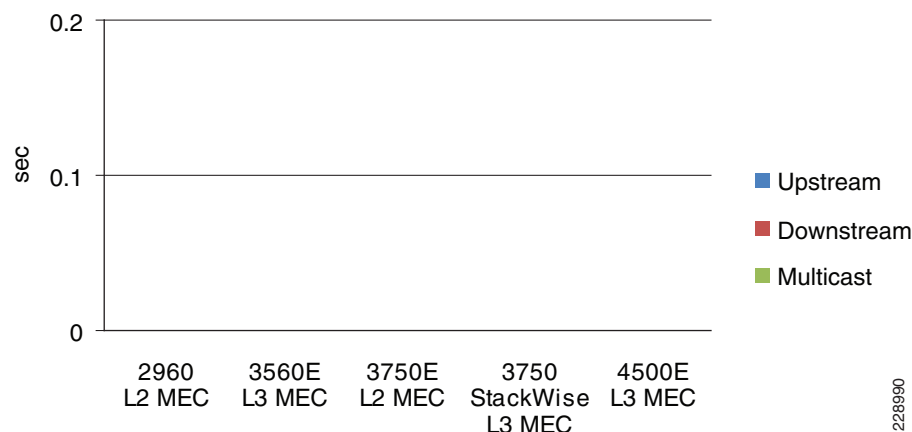
cr24-4507-LB#show power supplies
Power supplies needed by system:2
Power supplies currently available:2
```

### Network Recovery Analysis with Power Redundancy

Each campus LAN router and switch providing critical network services must be protected with either the in-chassis or external redundant power supply system. This best practice is also applicable to the standalone or virtual-systems devices. Each physical Catalyst 6500-E chassis in VSS mode at the campus distribution and core layer must be deployed with a redundant in-chassis power supply. The Catalyst 3750-X StackWise Plus must be deployed following the same rule, the master and member-switches in the stack ring must be deployed with the external redundant power system. Protecting virtual-systems with redundant power supplies will prevent reducing network bandwidth capacity, topology changes, and poor application performance.

Several power failures on power redundant systems were conducted to characterize overall network and application impact. The lab test results shown in [Figure 2-71](#) performed on all power redundant campus systems confirms zero-packet loss during individual power supply failure. Note that the network administrator must analyze the required power capacity that will be drawn by different hardware components (i.e., Network modules, PoE+ etc.).

**Figure 2-71 Redundant Power Analysis**




### Redundant Linecard Modules

Modular Catalyst platforms support a wide range of linecards for network connectivity to the network core and edge. The high speed core design linecards are equipped with special hardware components to build the campus backbone whereas the network edge linecards are developed with more intelligence and application awareness. Using internal system protocols, each line card communicates with the

centralized control-plane processing supervisor module through the internal backplane. Any type of internal communication failure or protocol malfunction may disrupt the communication between the linecard and the supervisor, which may lead to the linecard and all the physical ports associated with it to forcibly reset to resynchronize with the supervisor.

When the distribution and core layer Catalyst 4500-E and 6500-E systems are deployed with multiple redundant line cards, the network administrator must design the network by diversifying the physical cables across multiple linecard modules. A per system “V”-shaped, full-mesh physical design must have quad paths to address multiple types of faults. Deploying redundant linecards and diversifying paths across the modules will allow for inter-chassis re-route and, more importantly, the Cisco VSS traffic-engineering will prevent VSL reroute which may cause network congestion if there is not sufficient bandwidth to accommodate the rerouted traffic. [Figure 2-72](#) demonstrates inter-chassis reroute (without linecard redundancy) and intra-chassis re-route (with linecard redundancy).

**Figure 2-72 Intra-Chassis versus Inter-Chassis Traffic Re-route**



The single standalone Catalyst 4500-E in distribution or core layer must be deployed with linecard redundancy. The campus LAN network may face a complete network outage during linecard failures without deploying linecard redundancy as it can be considered a single point-of-failure.


## Redundant Linecard Network Recovery Analysis

### Catalyst 6500-E VSS Linecard module Recovery Analysis

The distributed forwarding architecture in Catalyst 6500-Es operating in VSS mode is designed with unique traffic-engineering capabilities. The centralized control-plane design on the active virtual-switch node builds Layer 2/3 peerings with the neighboring devices. However with MEC, both virtual-switch nodes program their local linecard modules to switch egress data plane traffic. This design minimizes data traffic re-routing across VSL links. Data traffic traverses the VSL links as a “last-resort” in hardware if either of the virtual-switch nodes lose a local member-link from the MEC link due to a fiber cut or linecard failure. The impact on traffic could be in the sub-second to seconds range and may create congestion on the VSL Etherchannel link if rerouting traffic exceeds overall VSL bandwidth capacity.

At the critical large campus LAN core and distribution layer, traffic loss can be minimized and consistent bi-directional sub-second network recovery can be achieved by deploying redundant network modules on a per virtual-switch node basis. Additionally, proper Cisco VSS traffic-engineering will prevent traffic routing over the VSL which may cause network congestion during individual link or entire high-speed network module failure. [Figure 2-72](#) provides an example of asymmetric traffic-loss statistics when traffic is rerouted via remote virtual-switch node across VSL links. [Figure 2-73](#) illustrates intra-chassis network recovery analysis showing symmetric sub-second traffic loss during individual member-links and the entire linecard module at the campus core and distribution-layer.




**Figure 2-73 Catalyst 6500-E VSS Intra-Chassis Link and Linecard Module Recovery Analysis**

228992

**Catalyst 4507R-E Linecard module Recovery Analysis**

The centralized forwarding architecture in a Catalyst 4507R-E programs all the forwarding information on the active and standby supervisor Sup6E or Sup6L-E modules. All the redundant linecards in the chassis are stub and maintains low level information to handle ingress and egress forwarding information. During a link or linecard module failure, the new forwarding information gets rapidly reprogrammed on both supervisors in the chassis. However, deploying the EtherChannel utilizing diversified fibers across different linecard modules will provide consistent sub-second network recovery during abnormal failure or the removal of a linecard from the Catalyst 4507R-E chassis. The chart in [Figure 2-74](#) provides test results conducted by removing a linecard from the Catalyst 4507R-E chassis deployed in campus network in various roles.

**Figure 2-74 Catalyst 4507R-E Linecard Recovery Analysis**

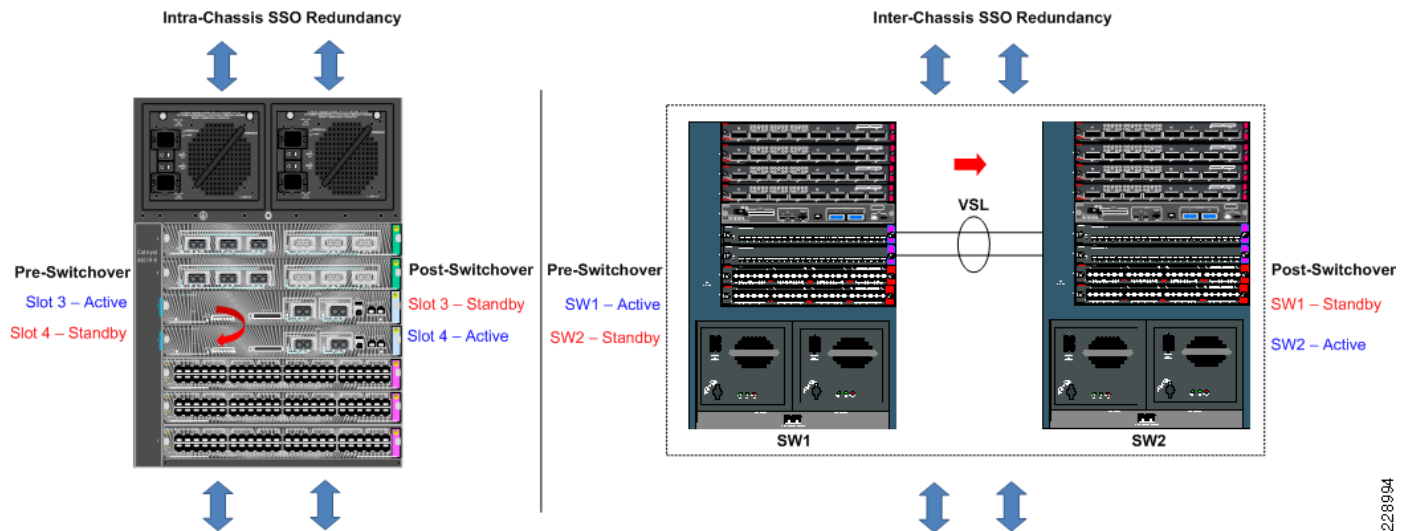
228993

**Redundant Supervisor**

Enterprise-class modular Cisco Catalyst 4500-E and 6500-E platforms support dual-redundant supervisor modules to prevent disrupting the network control-plane and topology during abnormal supervisor module failures or when forced by the admin reset. The Cisco Catalyst 4507R-E and 4510R-E Series platforms and all current generation Catalyst 6500-E Series chassis and supervisors support in-chassis redundant supervisor modules. However, with Cisco's latest Virtual-Switching System (VSS)

innovation and the next-generation Sup720-10GE supervisor module, supervisor redundancy can be extended across dual chassis by logically clustering them into one single large virtual-switch. See Figure 2-75.

**Figure 2-75 Intra-Chassis versus Inter-Chassis SSO Redundancy**



### Intra-Chassis SSO Redundancy

Intra-Chassis SSO redundancy in the Catalyst 4500-E switch provides continuous network availability across all the installed modules and the uplinks ports from active and standby supervisor modules. The uplink port remains in operation and forwarding state during an active supervisor switchover condition. Thus, it provides full network capacity even during SSO switchover. Cisco Catalyst 6500-E deployed in standalone mode also synchronizes all the hardware and software state-machine info in order to provide constant network availability during intra-chassis supervisor switchover.

- **Inter-Chassis SSO Redundancy**

The Cisco VSS solution extends supervisor redundancy by synchronizing SSO and all system internal communication over the special VSL EtherChannel interface between the paired virtual systems. Note VSS does not currently support intra-chassis supervisor redundancy on each individual virtual nodes. The virtual-switch node running in active supervisor mode will be forced to reset during the switchover. This may disrupt the network topology if not deployed with the best practices defined in this design guide. The “V”-shaped, distributed, full-mesh fiber paths combined with single point-to-point EtherChannel or MEC links play a vital role during such type of network events. During the failure, the new active virtual-switch node will perform a Layer 3 protocol graceful recovery with its neighbors in order to provide constant network availability over the local interfaces.

- **Implementing SSO Redundancy**

To deploy supervisor redundancy, it is important to remember that both supervisor modules must be identical in type and all the internal hardware components like memory and bootflash must be the same to provide complete operational transparency during failure.

The default redundancy mode on Catalyst 4500-E and Catalyst 6500-E series platforms is SSO. Hence it does not require any additional configuration to enable SSO redundancy. The following sample configuration illustrates how to implement VSS in SSO mode:

```
cr23-VSS-Core#config t
```

228994

```

cr23-VSS-Core(config)#redundancy
cr23-VSS-Core(config-red)#mode sso

cr23-VSS-Core#show switch virtual redundancy
My Switch Id = 1
Peer Switch Id = 2
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso

Switch 1 Slot 5 Processor Information :
-----
Current Software state = ACTIVE
<snippet>
Fabric State = ACTIVE
Control Plane State = ACTIVE

Switch 2 Slot 5 Processor Information :
-----
Current Software state = STANDBY HOT (switchover target)
<snippet>
Fabric State = ACTIVE
Control Plane State = STANDBY

```

## Non-Stop Forwarding (NSF)

When implementing NSF technology in SSO redundancy mode systems, the network disruption remains transparent and provides seamless availability to the campus users and applications remains during control-plane processing module (Supervisor/Route-Processor) gets reset. During a failure, the underlying Layer 3 NSF capable protocols perform graceful network topology re-synchronization and the preset forwarding information in hardware on the redundant processor or distributed linecards remain intact in order to continue switching network packets. This service availability significantly lowers the Mean Time To Repair (MTTR) and increases the Mean Time Between Failure (MTBF) to achieve highest level of network availability.

NSF is an integral part of a routing protocol and depends on the following fundamental principles of Layer 3 packet forwarding:

- *Cisco Express Forwarding (CEF)*—CEF is the primary mechanism used to program the network path into the hardware for packet forwarding. NSF relies on the separation of the control plane update and the forwarding plane information. The control plane is the routing protocol graceful restart, and the forwarding plane switches packets using hardware acceleration where available. CEF enables this separation by programming hardware with FIB entries in all Catalyst switches. This ability plays a critical role in NSF/SSO failover.
- *Routing protocol*—The motivation behind NSF is route convergence avoidance. From the protocol operation perspective, this requires the adjacent routers to support a routing protocol with special intelligence that allows a neighbor to be aware that NSF-capable routers can undergo switchover so that its peer can continue to forward packets, but may bring its adjacency to hold-down (NSF recovery mode) for a brief period, and requests routing protocol information to be resynchronized.

A router that has the capability for continuous forwarding during a switchover is *NSF-capable*. Devices that support the routing protocol extensions to the extent that they continue to forward traffic to a restarting router are *NSF-aware*. A Cisco device that is NSF-capable is also NSF-aware. The NSF capability must be manually enabled on each redundant system on a per routing protocol basis. The NSF aware function is enabled by default on all Layer 3 platforms. [Table 2-11](#) describes the Layer 3 NSF-capable and aware platforms deployed in the campus network environment.

The following configuration illustrates how to enable the NSF capability within EIGRP on each Layer 3 campus LAN/WAN systems deployed with redundant supervisor, route-processors or in virtual-switching modes (i.e., Cisco VSS and StackWise Plus):

```

cr23-vss-core(config)#router eigrp 100
cr23-vss-core (config-router)#nsf
cr23-vss-core #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

cr23-vss-core #show ip protocols | inc NSF
*** IP Routing is NSF aware ***
    EIGRP NSF-aware route hold timer is 240

```

### Graceful Restart Example

The following example demonstrates how the EIGRP protocol will gracefully recover when active supervisor/chassis switchover on a Cisco VSS core system is forced by a reset:

- NSF Capable System

```

cr23-VSS-Core#redundancy force-switchover
This will reload the active unit and force switchover to standby[confirm]y

NSF Aware/Helper System

! VSS active system reset will force all linecards and ports to go down
!the following logs confirms connectivity loss to core system
%LINK-3-UPDOWN: Interface TenGigabitEthernet2/1/2, changed state to down
%LINK-3-UPDOWN: Interface TenGigabitEthernet2/1/4, changed state to down

! Downed interfaces are automatically removed from EtherChannel/MEC,
! however additional interface to new active chassis retains port-channel in up/up
state
%EC-SW1_SP-5-UNBUNDLE: Interface TenGigabitEthernet2/1/2 left the port-channel
Port-channel100
%EC-SW1_SP-5-UNBUNDLE: Interface TenGigabitEthernet2/1/4 left the port-channel
Port-channel100

! EIGRP protocol completes graceful recovery with new active virtual-switch.
%DUAL-5-NBRCHANGE: EIGRP-IPv4:(613) 100: Neighbor 10.125.0.12 (Port-channel100) is
resync: peer graceful-restart

```

### NSF Timers

As depicted in the above show commands, up to 240 seconds NSF aware system can hold the routing information until routing protocol do not gracefully synchronize routing database. Lowering the timer values may abruptly terminate graceful recovery causing network instability. The default timer setting is well tuned for a well structured and concise campus LAN network topology. It is recommended to retain the default route hold timers in the network unless it is observed that NSF recovery takes more than 240 seconds.

600 seconds after the protocol graceful-recovery starts, the NSF route hold-timer expires on the NSF aware system and clears the stale NSF route marking and continues to use the synchronized routing database.

### NSF/SSO Recovery Analysis

As described in the previous section, the NSF/SSO implementation and its recovery process differs on Catalyst 4507R-E (Intra-Chassis) and Catalyst 6500-E VSS (Inter-Chassis) in the medium enterprise campus LAN network design. In both deployment scenarios, Cisco validated the network recovery and

application performance by inducing several types of active supervisor faults that trigger Layer 3 protocol graceful recovery. During each test, the switches continued to provide network accessibility during the recovery stage.


During the SSO switchover process, the Cisco Catalyst 4507R-E deployed with redundant Sup6E or Sup6L-E will retain the operational and forwarding state of the uplink ports and linecard modules in the chassis.

The inter-chassis SSO implementation in Catalyst 6500-E VSS differs from the single-chassis redundant implementation, in that during active virtual-switch node failure the entire chassis and all the linecards installed will reset. However, with Layer 2/3 MEC links, the network protocols and forwarding information remains protected via the remote virtual-switch node that can provide seamless network availability.

#### Catalyst 4507R-E NSF/SSO Recovery Analysis

Figure 2-76 illustrates intra-chassis NSF/SSO recovery analysis for the Catalyst 4507R-E chassis deployed with Sup6E or Sup6L-E in redundant mode. With EIGRP NSF/SSO capability the unicast traffic recovers consistently within 200 msec or less. However, Catalyst 4507R-E does not currently support redundancy for Layer 3 multicast routing and forwarding information. Therefore, there may be around 2 second multicast traffic loss since the switch has to re-establish all the multicast routing information and forwarding information during the Sup6E or Sup6L-E switchover event.

**Figure 2-76 Catalyst 4507R-E NSF/SSO Recovery Analysis**



In the remote medium campus, the Catalyst 4507R-E is also deployed as the PIM-SM RP with MSDP Anycast-RP peering to the Cisco VSS core in the main campus location. If a user from the remote medium campus location joins the multicast source from the main campus location then during Sup6E switchover there could be around a 3 second multicast packet loss. However unicast recovery will still remain in the 200 msec or less range in the same scenario.

#### Catalyst 4507R-E Standby Supervisor Failure and Recovery Analysis


The standby Sup6E or Sup6L-E supervisor remains in redundant mode while the active supervisor is in the operational state. If the standby supervisor gets reset or gets re-inserted, this event will not trigger protocol graceful recovery or any network topology change. The uplink port of the standby supervisor remains in operational and forwarding state and the network bandwidth capacity remains intact during a standby supervisor removal or insertion event.

### Catalyst 6500-E VSS NSF/SSO Recovery Analysis

As described earlier, the entire chassis and all linecard modules installed gets reset during an active virtual-switch switchover event. With a diverse full-mesh fiber network design, the Layer 2/3 remote device perceives this event as a loss of a member-link since the alternate link to the standby switch is in an operational and forwarding state. The standby virtual-switch detects the loss of the VSL Etherchannel and transitions in active role and initializes Layer 3 protocol graceful recovery with the remote devices. Since there is no major network topology changes and there are member-links still in an operational state, the NSF/SSO recovery in Catalyst 6500-E VSS system is identical as losing individual links.

Additionally, the Cisco Catalyst 6500-E supports Multicast Multilayer Switching (MMLS) NSF with SSO enabling the system to maintain the multicast forwarding state in PFC3 and DFC3 based hardware during an active virtual-switch reset. The new active virtual-switch reestablishes PIM adjacency while continuing to switch multicast traffic based on pre-switchover programmed information. See [Figure 2-77](#).

**Figure 2-77 Catalyst 6500-E VSS NSF/SSO Recovery Analysis**



### Catalyst 6500-E VSS Standby Failure and Recovery Analysis

The network impact during a VSS standby failure is similar to a failure of a VSS active virtual-switch node. The primary difference with a standby virtual-switch failure is that it will not trigger a Layer 3 protocol graceful recovery since the active virtual-switch is in an operational state. Each MEC neighbors will lose their physical path to standby switch and re-route traffic to the remaining MEC member-links connected to the active virtual-switch node. The VSS standby virtual-switch failure will trigger a bidirectional subsecond loss as illustrated in [Figure 2-77](#).

Since VSS is developed with the distributed forwarding architecture it can create certain race conditions during a standby re-initialization state since the virtual-switch receives traffic from the network while it is not fully ready to switch the traffic. The amount and the direction of traffic loss depend on multiple factors – VSL interface, ingress and egress module type, boot up ordering etc.

When the upstream device is a Catalyst 6500-E and it is deployed in standalone mode, then Cisco recommends configuring the **port-channel load-defer** command under the port-channel to prevent the traffic loss during the standby initialization state. It is possible to configure the same command line under the MEC interface when the upstream device is Catalyst 6500-E and it is deployed in VSS mode instead of standalone.

Cisco recommends not configuring the **port-channel load-defer** command under the MEC as it will create an adverse impact to the downstream unicast and multicast traffic:

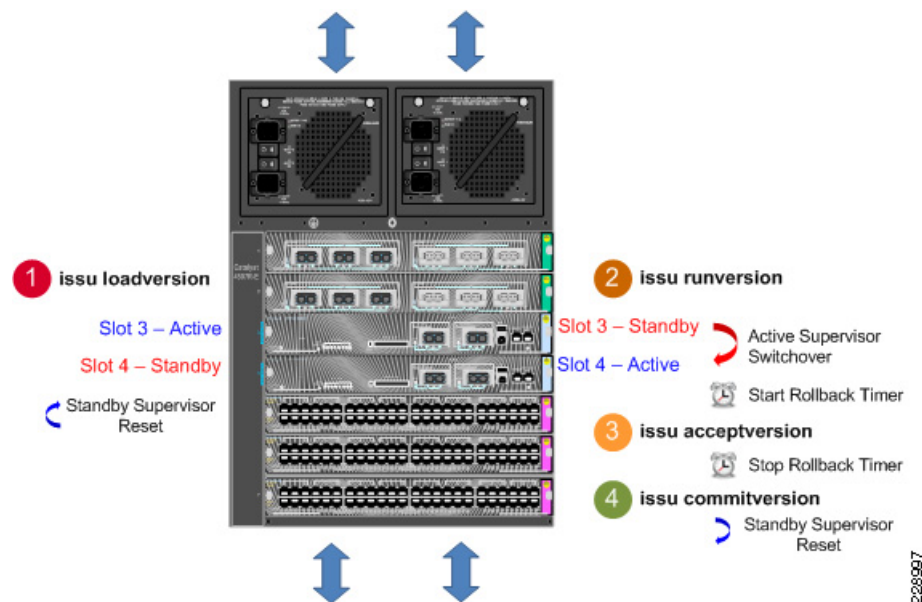
- The **port-channel load-defer** command is primarily developed for Catalyst 6500-E based standalone systems and does not have much effect when the campus upstream device type is Catalyst 6500-E deployed in VSS mode.
- There is no software restriction on turning on the feature on VSS systems. However, it may create an adverse impact on downstream multicast traffic. With the default multicast replication configuration, the MEC may drop multicast traffic until the defer timer expires (120 second default timer). Therefore, the user may experience traffic loss for a long period of time.
- Modifying the default (egress) multicast mode to the ingress replication mode may resolve the multicast traffic loss problem. However, depending on the network scale size, it may degrade performance and scalability.

## Implementing Operational Resiliency

Path redundancy often is used to facilitate access during periods of maintenance activity, but the single standalone systems are single points of failure sometimes exist or the network design simply does not allow for access if a critical node is taken out of service. Leveraging enterprise-class high availability features like NSF/SSO in the distribution and core layer Catalyst 4500-E and 6500-E Series platforms supports ISSU to enable real-time network upgrade capability. Using ISSU and eFSU technology, the network administrator can upgrade the Cisco IOS software to implement new features, software bug fixes or critical security fixes in real time.

### Catalyst 4500-E ISSU Software Design and Upgrade Process

**Figure 2-78** Catalyst 4500-E ISSU Software Upgrade Process



## ISSU Software Upgrade Pre-Requirement

### ISSU Compatibility Matrix

When a redundant Catalyst 4500-E system is brought up with a different Cisco IOS software version, the ISSU stored compatibility matrix information is analyzed internally to determine interoperability between the software running on the active and standby supervisors. ISSU provides SSO compatibility



between several versions of software releases shipped during a 18 month period. Prior to upgrading the software, the network administrator must verify ISSU software compatibility with the following command. Incompatible software may cause the standby supervisor to boot in RPR mode which may result in a network outage:

```
cr24-4507e-MB#show issu comp-matrix stored
Number of Matrices in Table = 1
      My Image ver:  12.2(53)SG
      Peer Version   Compatibility
      -----
      12.2(44)SGBase(2)
      12.2(46)SG           Base(2)
      12.2(44)SG1          Base(2)
      ...
```

### Managing System Parameters

#### Software

Prior to starting the software upgrade process, it is recommended to copy the old and new Cisco IOS software on Catalyst 4500-E active and standby supervisor into local file systems—Bootflash or Compact Flash.

```
cr24-4507e-MB#dir slot0:
Directory of slot0:/
 1  -rw- 25442405 Nov 23 2009 17:53:48 -05:00 cat4500e-entservicesk9-mz.122-53.SG1  ← new image
 2  -rw- 25443451 Aug 22 2009 13:26:52 -04:00 cat4500e-entservicesk9-mz.122-53.SG  ← old image

cr24-4507e-MB#dir slaveslot0:
Directory of slaveslot0:/
 1  -rw- 25443451 Aug 22 2009 13:22:00 -04:00 cat4500e-entservicesk9-mz.122-53.SG  ← old image
 2  -rw- 25442405 Nov 23 2009 17:56:46 -05:00 cat4500e-entservicesk9-mz.122-53.SG1  ← new image
```

### Configuration

It is recommended to save the running configuration to NVRAM and other local or remote locations such as bootflash or TFTP server prior upgrading IOS software.

#### Boot Variable and String

The system default boot variable is to boot from the local file system. Make sure the default setting is not changed and the configuration register is set to 0x2102.

Modify the boot string to point to the new image to boot from new IOS software version after the next reset triggered during ISSU upgrade process. Refer to following URL for additional ISSU pre-requisites:

<http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst4500/12.2/53SG/configuration/issu.html#wp1072849>

## Catalyst 4500-E ISSU Software Upgrade Procedure

This subsection provides the realtime software upgrade procedure for a Catalyst 4500-E deployed in the medium enterprise campus LAN network design in several different roles—access, distribution, core, collapsed core, and Metro Ethernet WAN edge. ISSU is supported on Catalyst 4500-E Sup6E and Sup6L-E supervisor running Cisco IOS Enterprise feature set.

In the following sample output, the Sup6E supervisor is installed in Slot3 and Slot4 respectively. The Slot3 supervisor is in the SSO Active role and the Slot4 supervisor is in Standby role. Both supervisors are running identical 12.2(53)SG Cisco IOS software version and is fully synchronized with SSO.

```
cr24-4507e-MB#show module | inc Chassis|Sup|12.2
```



```

Chassis Type : WS-C4507R-E
!Common Supervisor Module Type
 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
!Common operating system version
 3 0021.d8f5.45c0 to 0021.d8f5.45c5 0.4 12.2(33r)SG ( 12.2(53)SG Ok
 4 0021.d8f5.45c6 to 0021.d8f5.45cb 0.4 12.2(33r)SG ( 12.2(53)SG Ok
!SSO Synchronized
 3 Active Supervisor SSO Active
 4 Standby Supervisor SSO Standby hot

```

The following provides the step-by-step procedure to upgrade the Cisco IOS Release 12.2(53)SG to 12.2(53)SG1 Cisco IOS release without causing network topology and forwarding disruption. Each upgrade steps can be aborted at any stage by issuing the **issu abortversion** command if software detects any failure.

- **ISSU loadversion**—This first step will direct the active supervisor to initialize the ISSU software upgrade process.

```

cr24-4507e-MB#issu loadversion 3 slot0:cat4500e-entservicesk9-mz.122-53.SG1 4 slaveslot0:
cat4500e-entservicesk9-mz.122-53.SG1

```

After issuing the above command, the active supervisor ensures the new IOS software is downloaded on both supervisors file system and performs several additional checks on the standby supervisor for the graceful software upgrade process. ISSU changes the boot variable with the new IOS software version if no errors are found and resets the standby supervisor module.

```
%RF-5-RF_RELOAD: Peer reload. Reason: ISSU Loadversion
```



#### Note

Resetting the standby supervisor will not trigger a network protocol graceful recovery and all standby supervisor uplink ports will remain in operational and forwarding state for the transparent upgrade process.

With the broad range of ISSU version compatibility to form SSO communication the standby supervisor will successfully bootup again in its original standby state, see the following output.

```

cr24-4507e-MB#show module | inc Chassis|Sup|12.2
Chassis Type : WS-C4507R-E
! Common Supervisor Module Type
 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
! Mismatch operating system version
 3 0021.d8f5.45c0 to 0021.d8f5.45c5 0.4 12.2(33r)SG( 12.2(53)SG Ok
 4 0021.d8f5.45c6 to 0021.d8f5.45cb 0.4 12.2(33r)SG( 12.2(53)SG1 Ok
!SSO Synchronized
 3 Active Supervisor SSO Active
 4 Standby Supervisor SSO Standby hot

```

This bootup process will force the active supervisor to re-synchronize all SSO redundancy and checkpoints, VLAN database and forwarding information with the standby supervisor and will notify the user to proceed with the next ISSU step.

```

%C4K_REDUNDANCY-5-CONFIGSYNC: The config-reg has been successfully synchronized to the
standby supervisor
%C4K_REDUNDANCY-5-CONFIGSYNC: The startup-config has been successfully synchronized to the
standby supervisor
%C4K_REDUNDANCY-5-CONFIGSYNC: The private-config has been successfully synchronized to the
standby supervisor
%C4K_REDUNDANCY-5-CONFIGSYNC_RATELIMIT: The vlan database has been successfully
synchronized to the standby supervisor

%ISSU_PROCESS-7-DEBUG: Peer state is [ STANDBY HOT ]; Please issue the runversion command

```

- *ISSU runversion*—After performing several steps to assure the new loaded software is stable on the standby supervisor, the network administrator must proceed to the second step.

```
cr24-4507e-MB#issu runversion 4
This command will reload the Active unit. Proceed ? [confirm]y
%RF-5-RF_RELOAD: Self reload. Reason: Admin ISSU runversion CLI
%SYS-5-RELOAD: Reload requested by console. Reload reason: Admin ISSU runversion
```

This step will force the current active supervisor to reset itself which will trigger network protocol graceful recovery with peer devices, however the uplink ports of the active supervisor remains intact and the data plane will remain un-impacted during the switchover process. From the overall network perspective, the active supervisor reset caused by the **issu runversion** command will be no different than similar switchover procedures (i.e., administrator-forced switchover or supervisor online insertion and removal). During the entire software upgrade procedure; this is the only step that performs SSO-based network graceful recovery. The following syslog on various Layer 3 systems confirm stable and EIGRP graceful recovery with the new supervisor running the new Cisco IOS software version.

- NSF-Aware Core

```
cr23-VSS-Core#
%DUAL-5-NBRCHANGE: EIGRP-IPv4:(415) 100: Neighbor 10.125.0.15 (Port-channel102) is
resync: peer graceful-restart
```

- NSF-Aware Layer 3 Access

```
cr24-3560-MB#
%DUAL-5-NBRCHANGE: EIGRP-IPv4:(100) 100: Neighbor 10.125.0.10 (Port-channel1) is
resync: peer graceful-restart
```

The previously active supervisor module will boot up in the standby role with the older IOS software version instead the new IOS software version.

```
cr24-4507e-MB#show module | inc Chassis|Sup|12.2
Chassis Type : WS-C4507R-E
! Common Supervisor Module Type
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 JAE1132SXQ
! Mismatch operating system version
3 0021.d8f5.45c0 to 0021.d8f5.45c5 0.4 12.2(33r)SG( 12.2(53)SG Ok
4 0021.d8f5.45c6 to 0021.d8f5.45cb 0.4 12.2(33r)SG( 12.2(53)SG1 Ok
!SSO Synchronized
3 Active Supervisor SSO Standby hot
4 Standby Supervisor SSO Active
```

This safeguarded software design provides an opportunity to roll back to the previous IOS software if the system upgrade causes any type of network abnormalities. At this stage, ISSU automatically starts internal rollback timers to re-install old IOS image. The default rollback timer is up to 45 minutes which provides a network administrator an opportunity to perform several sanity checks. In small to mid size network designs, the default timer may be sufficient. However, for large networks, network administrators may want to adjust the timer up to 2 hours:

```
cr24-4507e-MB#show issu rollback-timer
Rollback Process State = In progress
Configured Rollback Time = 45:00
Automatic Rollback Time = 19:51
```

The system will notify the network administrator with the following syslog to instruct them to move to the next ISSU upgrade step if no stability issues are observed and all the network services are operating as expected.

```
%ISSU_PROCESS-7-DEBUG: Peer state is [ STANDBY HOT ]; Please issue the acceptversion
command
```

- *ISSU acceptversion*—This step provides confirmation from the network administrator that the system and network is stable after the IOS install and they are ready to accept the new IOS software on the standby supervisor. This step stops the rollback timer and instructs the network administrator to issue the final commit command. However, it does not perform any additional steps to install the new software on standby supervisor.

```
cr24-4507e-MB#issu acceptversion 4
% Rollback timer stopped. Please issue the commitversion command.
```

```
cr24-4507e-MB#show issu rollback-timer
Rollback Process State = Not in progress
Configured Rollback Time = 45:00
```

```
cr24-4507e-MB#show module | inc Chassis|Sup|12.2
Chassis Type : WS-C4507R-E
! Common Supervisor Module Type
 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
! Mismatch operating system version
 3      0021.d8f5.45c0 to 0021.d8f5.45c5 0.4 12.2(33r)SG( 12.2(53)SG      Ok
 4      0021.d8f5.45c6 to 0021.d8f5.45cb 0.4 12.2(33r)SG( 12.2(53)SG1      Ok
!SSO Synchronized
 3      Active Supervisor      SSO Standby hot
 4      Standby Supervisor      SSO Active
```

- *ISSU commitversion*—This final ISSU step forces the active supervisor to synchronize its configuration with the standby supervisor and force it to reboot with the new IOS software. This stage concludes the ISSU upgrade procedure and the new IOS version is permanently committed on both supervisor modules. If for some reason the network administrator wants to rollback to the older image, then it is recommended to perform an ISSU-based downgrade procedure to retain the network operational state without any downtime planning.

```
cr24-4507e-MB#issu commitversion 3
Building configuration...
Compressed configuration from 24970 bytes to 10848 bytes[OK]
%C4K_REDUNDANCY-5-CONFIGSYNC: The private-config has been successfully synchronized to the
standby supervisor
%RF-5-RF_RELOAD: Peer reload. Reason: ISSU Commitversion
```

```
cr24-4507e-MB#show module | inc Chassis|Sup|12.2
Chassis Type : WS-C4507R-E
! Common Supervisor Module Type
 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
! Common new operating system version
 3      0021.d8f5.45c0 to 0021.d8f5.45c5 0.4 12.2(33r)SG( 12.2(53)SG1      Ok
 4      0021.d8f5.45c6 to 0021.d8f5.45cb 0.4 12.2(33r)SG( 12.2(53)SG1      Ok
!SSO Synchronized
 3      Active Supervisor      SSO Standby hot
 4      Standby Supervisor      SSO Active
```


## Catalyst 6500-E VSS eFSU Software Design and Upgrade Process

Cisco Catalyst VSS was introduced in the initial IOS Release 12.2(33)SXH that supported Fast Software Upgrade (FSU). In the initial introduction, it had limited high-availability consideration to upgrade the IOS software release. The ISSU mismatched software version compatibility was not supported by the FSU infrastructure which could cause network down time. This may not be a desirable solution when deploying Catalyst 6500-E in the critical aggregation or core network tier.

Starting with the IOS Release 12.2(33)SXI, the Catalyst 6500-E supports true hitless IOS software upgrade in standalone and virtual-switch network designs. Enhanced Fast Software Upgrade (eFSU) made it completely ISSU infrastructure compliant and enhances the software and hardware design to retain its functional state during the graceful upgrade process.

### Catalyst 6500-E VSS eFSU Software Design and Upgrade Process

**Figure 2-79 Catalyst 6500-E VSS eFSU Software Upgrade Process**



Since eFSU in the Catalyst 6500-E system is built on the ISSU infrastructure, most of the eFSU pre-requisites and IOS upgrade procedures remain consistent as explained in previous sub-section. As described earlier, the Cisco VSS technology enables inter-chassis SSO communication between two virtual-switch nodes. However, while the software upgrade procedure for inter-chassis eFSU upgrades is similar, the network operation slightly differs compared to ISSU implemented on intra-chassis based SSO design.

## Catalyst 6500-E eFSU Software Upgrade Procedure

This subsection provides the software upgrade procedure for Catalyst 6500-Es deployed in VSS mode in the medium enterprise campus LAN network design. eFSU is supported on the Catalyst 6500-E Sup720-10GE supervisor module running Cisco IOS release with the Enterprise feature set.

In the following sample output, a VSS capable Sup720-10G supervisor module is installed in Slot5 of virtual-switch SW1 and SW2 respectively. The virtual-Switch SW1 supervisor is in the SSO Active role and the SW2 supervisor is in the Standby hot role. In addition, with MEC and the distributed forwarding architecture, the forwarding plane is in an active state on both virtual-switch nodes. Both supervisor are running identical the Cisco IOS Release 12.2(33)SX12a software version and is fully synchronized with SSO.

```
cr23-VSS-Core#show switch virtual redundancy | inc Mode|Switch|Image|Control
! VSS switch node with control-plane ownership
```

```

My Switch Id = 1
Peer Switch Id = 2
! SSO Synchronized
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
! Common operating system version
Switch 1 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI2a
Control Plane State = ACTIVE
Switch 2 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI2a
Control Plane State = STANDBY

```

The following provides a step-by-step procedure to upgrade from Cisco IOS Release 12.2(33)SXI2a to 12.2(33)SXI3 without causing network topology and forwarding disruption. Each upgrade step can be aborted at any stage by issuing the **issu abortversion** command if the software detects any failures.

- **ISSU loadversion**—This first step will direct the active virtual-switch node to initialize the ISSU software upgrade process.

```

cr23-VSS-Core#issu loadversion 1/5 disk0: s72033-adventerprisek9_wan-mz.122-33.SXI3 2/54
slavedisk0: s72033-adventerprisek9_wan-mz.122-33.SXI3

```

After issuing the above command, the active virtual-switch ensures the new IOS software is downloaded on both supervisors file system and performs several additional checks on the standby supervisor on the remote virtual-switch for the graceful software upgrade process. ISSU changes the boot variable to the new IOS software version if no error is found and resets the standby virtual-switch and installed modules.

```

%RF-SW1_SP-5-RF_RELOAD: Peer reload. Reason: ISSU Loadversion
%SYS-SW2_SPSTBY-5-RELOAD: Reload requested - From Active Switch (Reload peer unit).

```



#### Note

Resetting standby virtual-switch node will not trigger the network protocol graceful recovery process and will not reset the linecards on the active virtual-switch. It will remain in operational and forwarding state for the transparent upgrade process.

With the broad range of ISSU version compatibility to form SSO communication the standby supervisor will successfully bootup again in its original standby state, see the following output.

```

cr23-VSS-Core#show switch virtual redundancy | inc Mode|Switch|Image|Control
! VSS switch node with control-plane ownership
My Switch Id = 1
Peer Switch Id = 2
! SSO Synchronized
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
! Mismatch operating system version
Switch 1 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI2a, RELEASE SOFTWARE (fc2)
Control Plane State = ACTIVE
Switch 2 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI3, RELEASE SOFTWARE (fc2)
Control Plane State = STANDBY

```

To rejoin the virtual-switch domain, both nodes will reestablish the VSL EtherChannel communication and force the active supervisor to resynchronize all SSO redundancy and checkpoints, VLAN database and forwarding information with the standby virtual-switch and the network administrator is notified to proceed with the next ISSU step.

```
%HA_CONFIG_SYNC-6-BULK_CFGSYNC_SUCCEED: Bulk Sync succeeded
%PFREDUN-SW2_SPSTBY-6-STANDBY: Ready for SSO mode
```

```
%ISSU_PROCESS-SW1_SP-7-DEBUG: Peer state is [ STANDBY HOT ]; Please issue the runversion
command
```

- *ISSU runversion*—After performing several steps to assure the new loaded software is stable on the standby virtual-switch, the network administrator is now ready to proceed to the runversion step.

```
cr23-VSS-Core#issu runversion 2/5
This command will reload the Active unit. Proceed ? [confirm]
%issu runversion initiated successfully

%RF-SW1_SP-5-RF_RELOAD: Self reload. Reason: Admin ISSU runversion CLI
```

This step will force the current active virtual-switch (SW1) to reset itself which will trigger network protocol graceful recovery with peer devices; however the linecard on the current standby virtual-switch (SW2) will remain intact and the data plane traffic will continue get switched during the switchover process. From the network perspective, the affects of the active supervisor resetting during the ISSU runversion step will be no different than the normal switchover procedure (i.e., administration-forced switchover or supervisor online insertion and removal). In the entire eFSU software upgrade procedure, this is the only time that the systems will perform an SSO-based network graceful recovery. The following syslogs confirm stable and EIGRP graceful recovery on the virtual-switch running the new Cisco IOS software version.

### NSF-Aware Distribution

```
cr24-4507e-MB#
%DUAL-5-NBRCHANGE: EIGRP-IPv4:(100) 100: Neighbor 10.125.0.14 (Port-channel1) is resync:
peer graceful-restart
```

After re-negotiating and establishing the VSL EtherChannel link and going through the VSLP protocol negotiation process, the rebooted virtual-switch module boots up in the standby role with the older IOS software version instead the new IOS software version.

```
cr23-VSS-Core#show switch virtual redundancy | inc Mode|Switch|Image|Control
! VSS switch node with control-plane ownership changed to SW2
My Switch Id = 2
Peer Switch Id = 1
! SSO Synchronized
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
! Mismatch operating system version
Switch 2 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI3, RELEASE SOFTWARE (fc2)
Control Plane State = ACTIVE
Switch 1 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI2a, RELEASE SOFTWARE (fc2)
Control Plane State = STANDBY
```

Like intra-chassis ISSU implementation, eFSU also provides a safeguarded software design for additional network stability and opportunity to roll back to the previous IOS software if the system upgrade causes any type of network abnormalities. At this stage, ISSU automatically starts internal rollback timers to re-install old IOS image if there are any problems. The default rollback timer is up to 45 minutes which provides the network administrator an opportunity to perform several sanity checks. In small to mid size network designs, the default timer may be sufficient. However for large networks, the network administrator may want to adjust the timer up to 2 hours:

```
cr23-VSS-Core#show issu rollback-timer
Rollback Process State = In progress
```

```
Configured Rollback Time = 00:45:00
Automatic Rollback Time = 00:36:08
```

The system will notify the network administrator with following syslog to continue to the next ISSU upgrade step if no stability issues are observed and all the network services are operating as expected.

```
%ISSU_PROCESS-SW2_SP-7-DEBUG: Peer state is [ STANDBY HOT ]; Please issue the
acceptversion command
```

- *ISSU acceptversion*—This eFSU step provides confirmation from the network administrator regarding the system and network stability after installing the new software and confirms they are ready to accept the new IOS software on the standby supervisor. This step stops the rollback timer and instructs the network administrator to continue to the final commit state. However, it does not perform any additional steps to install the new software on standby supervisor.

```
cr23-VSS-Core#issu acceptversion 2/5
% Rollback timer stopped. Please issue the commitversion command.
cr23-VSS-Core#show issu rollback-timer
Rollback Process State = Not in progress
Configured Rollback Time = 00:45:00
```

```
cr23-VSS-Core#show switch virtual redundancy | inc Mode|Switch|Image|Control
! VSS switch node with control-plane ownership changed to SW2
My Switch Id = 2
Peer Switch Id = 1
! SSO Synchronized
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
! Mismatch operating system version
Switch 2 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI3, RELEASE SOFTWARE (fc2)
Control Plane State = ACTIVE
Switch 1 Slot 5 Processor Information :
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M), Version
12.2(33)SXI2a, RELEASE SOFTWARE (fc2)
Control Plane State = STANDBY
```

- *ISSU commitversion*—The final eFSU step forces the active virtual-switch to synchronize the configuration with the standby supervisor and force it to reboot with the new IOS software. This stage concludes the eFSU upgrade procedure and the new IOS version is permanently committed on both virtual-switches. If for some reason the network administrator needs to rollback to the older image, then it is recommended to perform the eFSU-based downgrade procedure to maintain the network operational state without any downtime planning.

```
cr23-VSS-Core#issu commitversion 1/5
Building configuration...
[OK]
%RF-SW2_SP-5-RF_RELOAD: Peer reload. Reason: Proxy request to reload peer
%SYS-SW1_SPSTBY-5-RELOAD: Reload requested - From Active Switch (Reload peer unit).
%issu commitversion executed successfully
```

```
cr23-VSS-Core#show switch virtual redundancy | inc Mode|Switch|Image|Control
! VSS switch node with control-plane ownership
My Switch Id = 2
Peer Switch Id = 1
! SSO Synchronized
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
! Common operating system version
Switch 2 Slot 5 Processor Information :
```



```
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M) ,  
Version 12.2(33)SXI3, RELEASE SOFTWARE (fc2)  
Control Plane State = ACTIVE  
Switch 1 Slot 5 Processor Information :  
Image Version = Cisco IOS Software, s72033_rp Software (s72033_rp-ADVENTERPRISEK9_WAN-M) ,  
Version 12.2(33)SXI3, RELEASE SOFTWARE (fc2)  
Control Plane State = STANDBY
```

## Summary

Designing the LAN network aspects for the medium enterprise network design establishes the foundation for all other aspects within the service fabric (WAN, security, mobility, and UC) as well as laying the foundation to provide safety and security, operational efficiencies, virtual learning environments, and secure classrooms.

This chapter reviews the two LAN design models recommended by Cisco, as well as where to apply these models within the various locations of a medium 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 medium enterprise design.





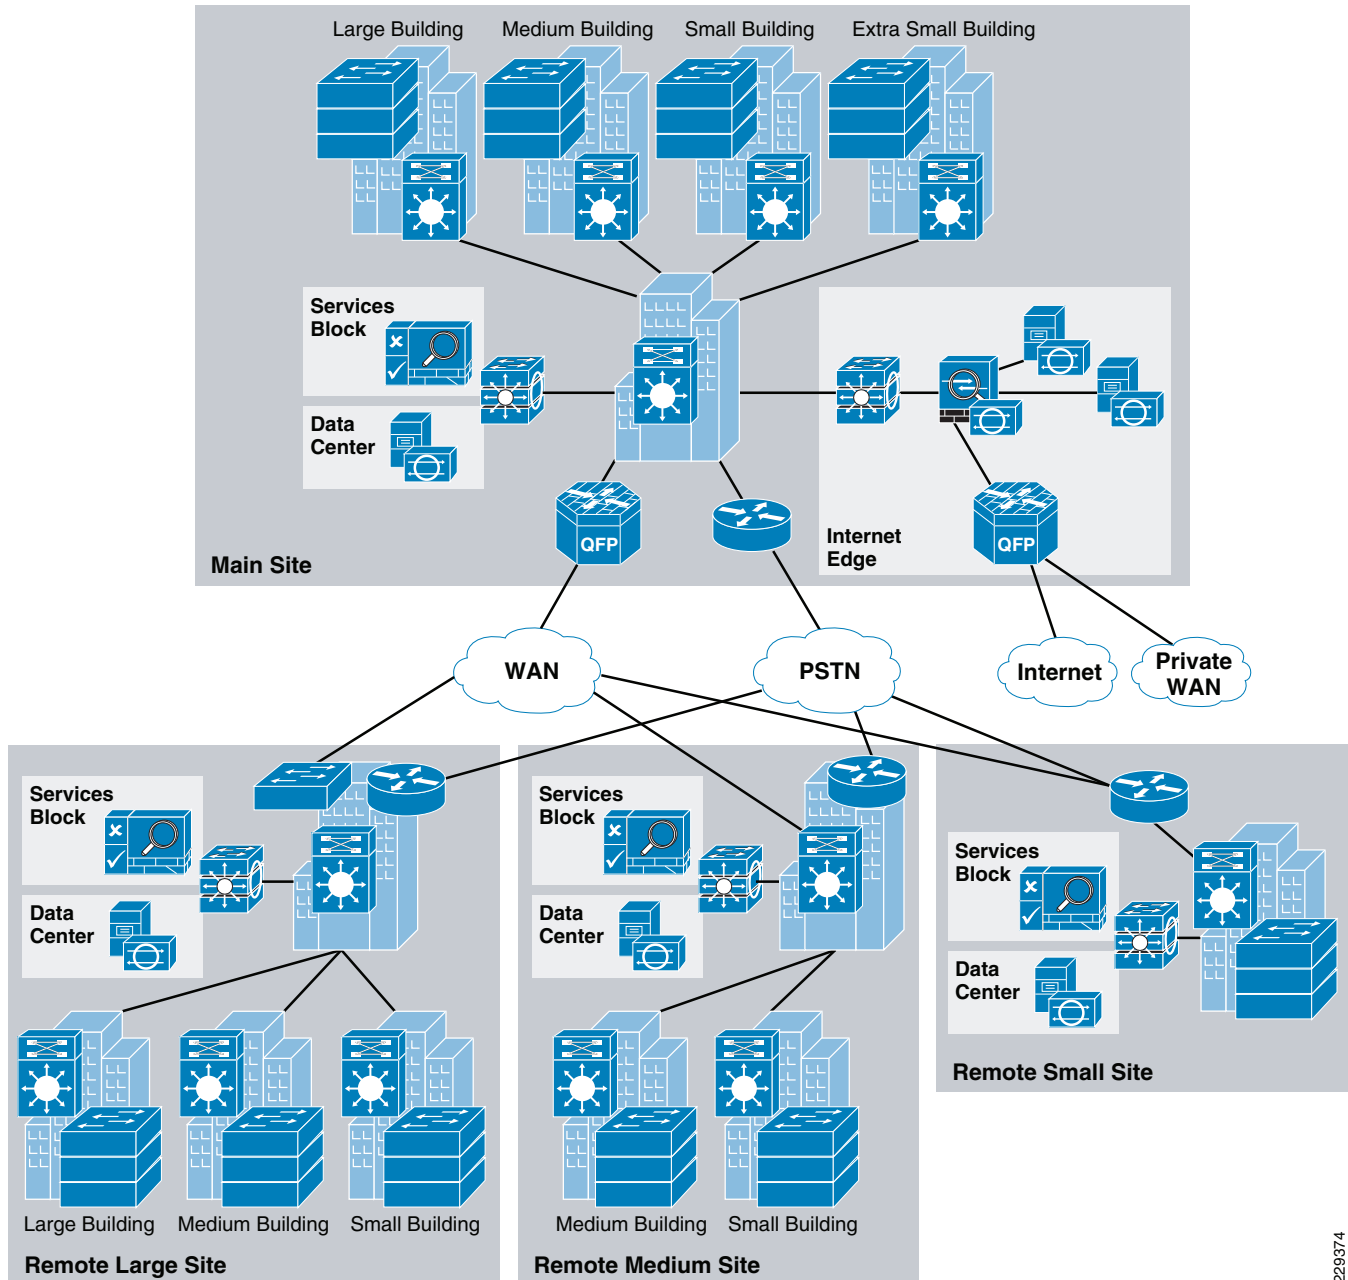
## CHAPTER 3

# Medium Enterprise Design Profile (MEDP)—WAN Design

---

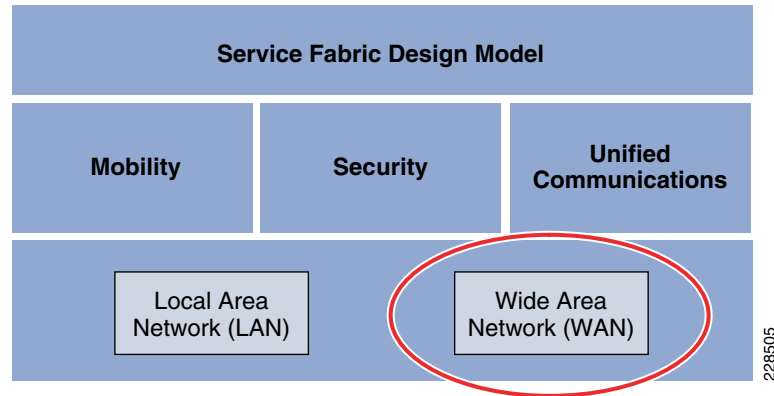
## WAN Design

The Medium Enterprise WAN Design Profile is a multi-site design where a site consists of multiple buildings and services. The sites are interconnected through various WAN transports as shown in [Figure 3-1](#).

**Figure 3-1** Medium Enterprise WAN Design Diagram

229374

Within the Medium Enterprise Design Profile, the service fabric network provides the foundation on which all the solutions and services are built upon to solve the business challenges. This service fabric consists of four distinct components as shown in [Figure 3-2](#).

**Figure 3-2 The Service Fabric Design Model**

This chapter discusses the WAN design component of the Medium Enterprise Design Profile. This section discusses how the WAN design is planned for medium enterprises, the assumptions made, the platforms chosen, and the justification for choosing a platform. The WAN design is highly critical to provide network access for remote sites to the main site, as well as connectivity to other networks, and general Internet access for the entire enterprise. The WAN design should not be viewed merely for providing access, but mainly to see how the business requirements can be met. Therefore, it is important for communication to exist between the employees, customers, and partners. This communication could be with voice, video, or data applications. Moreover, the video applications, may possess, flavors ranging from desktop video to real-time video. To provide this collaborative environment, highly resilient and, highly performing WAN designs are required.

The main components of Medium Enterprise Design Profile for WAN architecture are as follows:

- WAN transport
- WAN devices
- Network Foundation services—Routing, QoS, and multicast

## WAN Transport


This section discusses the different WAN transports present in the Medium Enterprise Design Profile.

### Private WAN Service

The Medium Enterprise Design Profile consists of several locations. These locations have similar architecture as the main site. However, these sites need to collaborate with each other to meet the business objectives. Therefore, a WAN network that can support the following requirements is needed:


- High performance
- Support different classes of traffic
- Native routing
- Multicast capability
- Security

To support these requirements enterprises need to have a private WAN service to provide connectivity between remote sites, and main site. See [Figure 3-3](#).

**Figure 3-3 Medium Enterprise Connectivity to Other Remote Sites Using Private WAN**

## Internet Service

The physical connection for reaching the Internet and the private WAN network is same; however, both circuits are logically separated using different subinterfaces. Therefore, it is similar to a situation where a customer is connected to different service providers. See [Figure 3-4](#).

**Figure 3-4 Medium Enterprise Internet Service**

## Metro Service

Metro Ethernet is one of the fastest growing WAN transport technologies in the telecommunications industry. The advantages of using this WAN transport are as follows:

- Scalability and reachability
  - The services offered would scale from 1Mbps to 10Gbps, and beyond in granular increments, which makes this transport highly scalable.
  - Service providers worldwide are migrating their networks to provide metro services; thereby, it is available at large number of places.
- Performance, QoS, and suitability for convergence
  - Inherently Ethernet networks require less processing to operate and manage and operate at higher bandwidth than other technologies.
  - The granular options in bandwidth, ability to provide different SLAs based on voice, video, and data applications that provide QoS service to customers.

- Low latency and delay variation make it the best solution for video, voice, and data.
- Cost savings
  - Metro Ethernet brings the cost model of Ethernet to the WAN.
- 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.


There are two popular methods of service for Metro Ethernet:

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

EVPL, like Frame Relay, provides for 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.

E-LAN, also known as Virtual Private LAN Services (VPLS), provides 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.

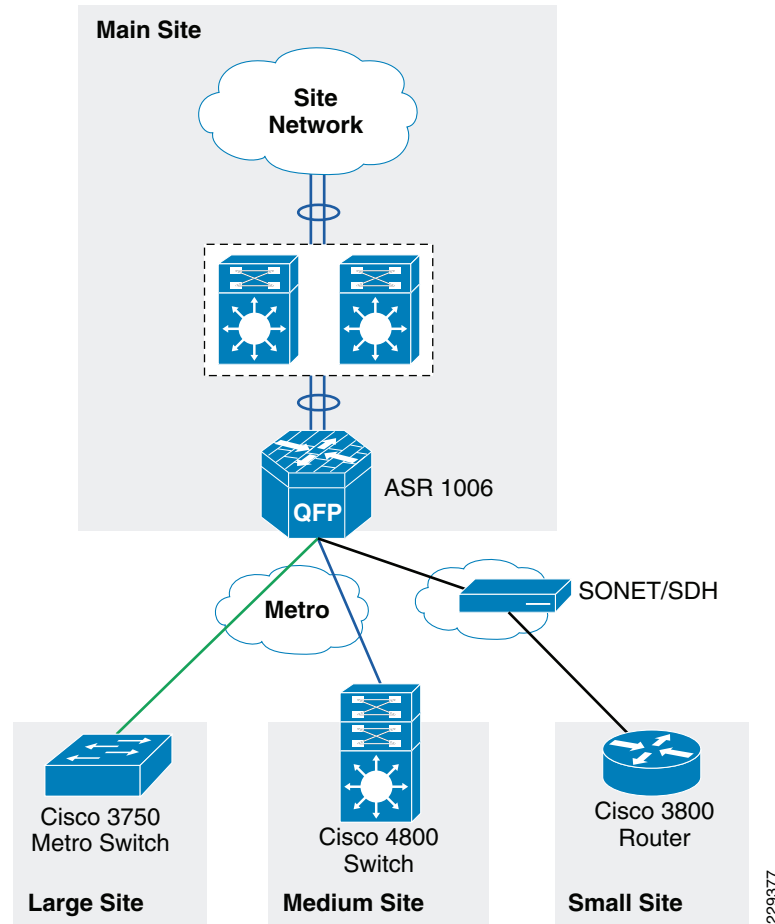
Figure 3-5 shows the difference between these services.

**Figure 3-5** *Different Services Available*

This section discusses how the Metro service is designed in the Medium Enterprise Design Profile. The Metro service is used to provide connectivity between the remote sites to the main site. The key reasons for recommending Metro service for Medium Enterprise are as follows:

- *Centralized administration and management*—E-line service provides point-to-point connectivity, whereas, E-LAN provides point-to-multipoint connectivity. Having a point-to-point connectivity mandates that all the remote site sites need to traverse the main site to reach the other, making the centralized administration applicable.
- *Performance*—Since all the application services are centrally located at main site, the WAN bandwidth required for remote sites to main site should be at least 100 Mbps. The Metro transport can provide 100Mbps, and more if needed in the future.

Therefore, in this design, it is recommended that the remote large and medium remote site locations use E-line service to connect to the main site. [Figure 3-6](#) shows how the remote site locations are connected to main site using Metro service.

**Figure 3-6** The Metro Transport Deployment in Medium Enterprise WAN Design

## Leased-Line Service

The WAN bandwidth requirement for a small remote site is assumed to be 20Mbps. Cisco recommends that the small remote site connect to the main site using a private leased-line service. The leased-line service is more readily available for these type of locations and the bandwidth is sufficient for the small remote site application requirements.

## WAN Aggregation Platform Selection in the Medium Enterprise Design Profile

In addition to selecting the WAN service for connectivity between remote site locations and access to the Internet, choosing the appropriate WAN aggregation router is essential. For each location in the Medium Enterprise Design Profile various WAN aggregation platforms are selected based on the requirements.

## Main Site WAN Aggregation Platform Selection

A WAN aggregation router aggregates all the incoming WAN circuits from various locations in the network as well as the Internet and also provides the proper QoS required for application delivery. Cisco recommends the Cisco ASR family of routers as the WAN aggregation platform for the main site location.

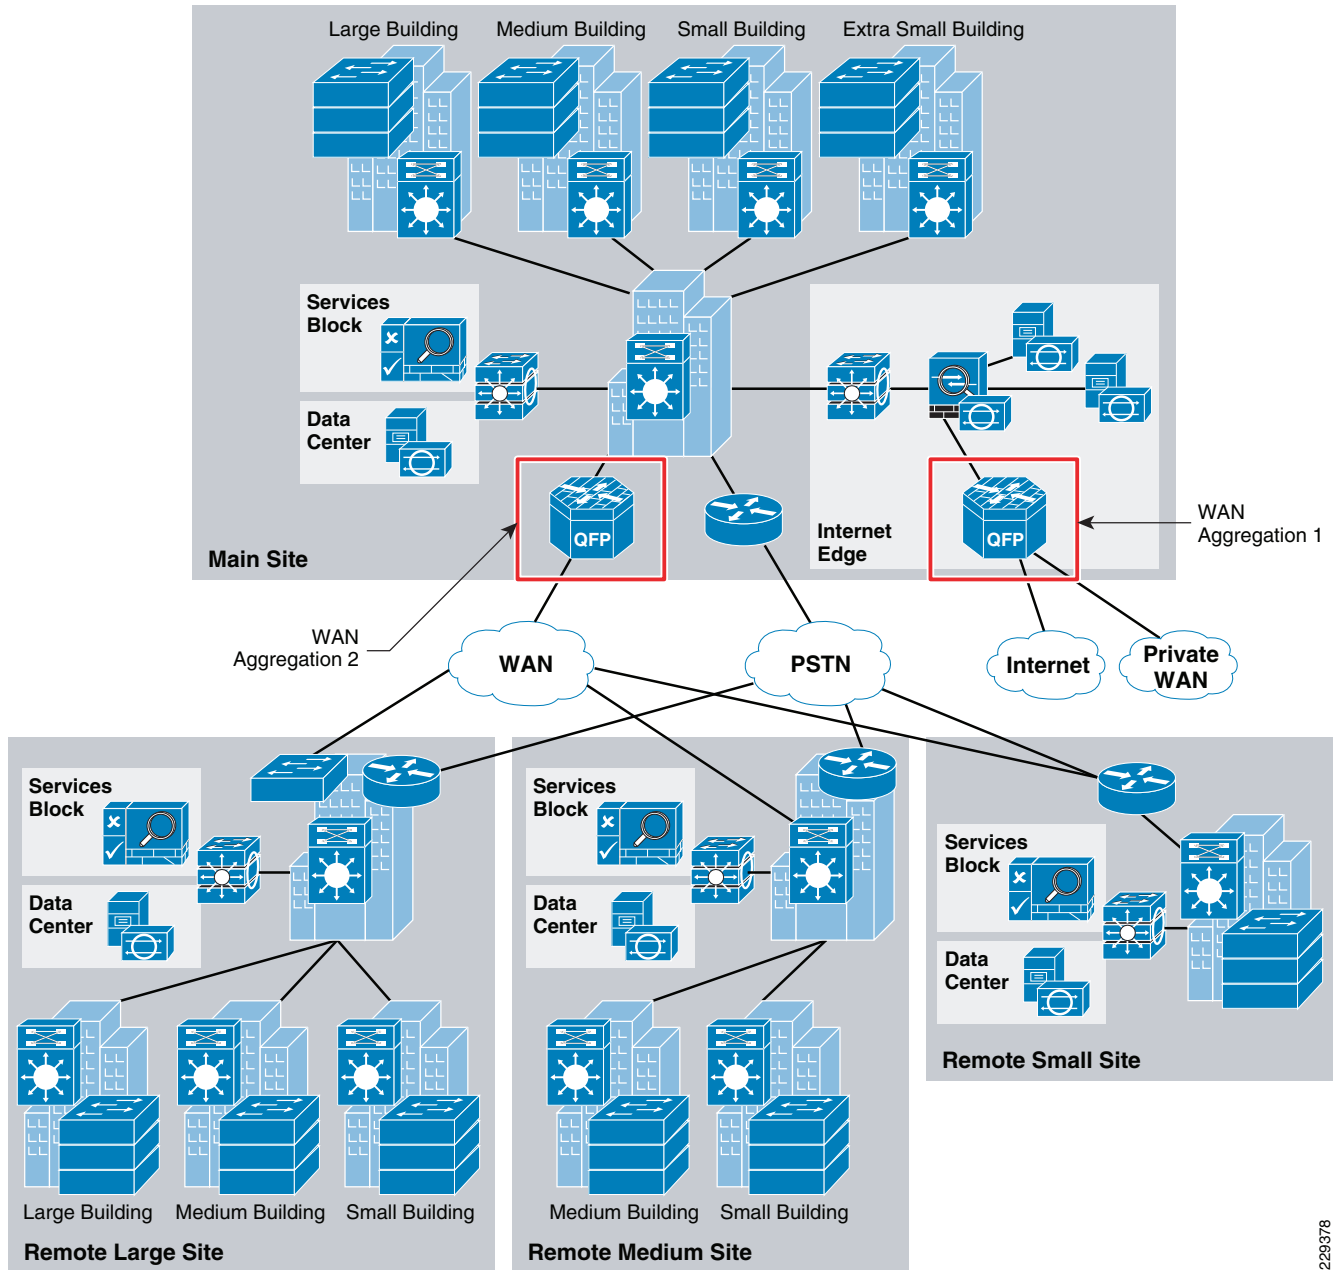
The Cisco ASR 1000 Series Router family consists of three different models:

- The Cisco ASR 1002 Router is a 3-SPA, 2-rack-unit (RU) chassis with one Embedded Services Processor (ESP) slot that comes with an integrated Router Processor (RP), integrated Cisco ASR 1000 Series Shared Port Adapter Interface Processor (SIP), and integrated four Gigabit Ethernet ports.
- The Cisco ASR 1004 Router is an 8-SPA, 4-RU chassis with one ESP slot, one RP slot, and two SIP slots.
- The Cisco ASR 1006 Router is a 12-SPA, 6-RU, hardware redundant chassis with two ESP slots, two RP slots and three SIP slots.

In Medium Enterprise Design Profile, there are two places where the WAN aggregation occurs in the main site location. The first place is where the main site location connects to outside world using private WAN and Internet networks. The second place is where all the remote sites connect to the main site.

[Figure 3-7](#) shows the two different WAN aggregation devices.



**Figure 3-7** The WAN Aggregation Points in Medium Enterprise WAN Design

229378

### WAN Aggregation 1

A Cisco ASR 1004 Series router is recommended as the WAN aggregation platform for private WAN/Internet connectivity. This choice was made considering the cost and required features—performance, QoS, routing, and resiliency—that are essential requirements for WAN aggregation router. Moreover, this platform contains built-in resiliency capabilities such as ISSU and IOS-based redundancy.

## WAN Aggregation 2

The second WAN aggregation device provides connectivity to the large and medium remote sites to the main site. To perform this aggregation, the Cisco ASR 1006 router with redundant route processors and redundant ESP's has been recommended for the following reasons:

- *Performance*—Up to 20 Gbps throughput
- *Port density*—Up to 12 shared port adapters (SPAs), the highest port density solution of the three Cisco ASR 1000 routers
- *Resiliency*—Cisco ASR 1006 router supports hardware redundancy and in-service software upgrades (ISSU). This chassis would support dual route processors, and dual ESP modules to support the hardware redundancy. Moreover, this router would also support EtherChannel load balancing feature.

## Large Remote Site WAN Aggregation Platform Selection

The WAN connectivity between the large remote site to the main site is fairly simpler because of the lack of requirements of advanced encryption technologies. Therefore, the main purpose is to reduce the cost and try to consolidate the WAN functionality into the distribution device at the large site. However, at the large remote site, as per the site LAN design document, VSS has been chosen as technology on the distribution switch, and VSS does not support WAN functionality. Therefore, a dedicated WAN aggregation device is needed to perform that functionality, which can be an ASR, 7200, or 3750ME switches. Out of these choices, considering the cost/performance criteria, the Cisco 3750ME switch was selected to perform the WAN aggregation. The Cisco 3750 Metro switch has the following features/capabilities to adequately meet the requirements:

- Hierarchical QoS
- Routing support: OSPF, EIGRP, and BGP
- Multicast support: PIM
- Redundant power supply

## Medium Remote Site WAN Aggregation Platform Selection

As discussed in [Chapter 2, “Medium Enterprise Design Profile \(MEDP\)—LAN Design,”](#) the medium remote site collapses the WAN edge and core-layer LAN functionality into a single switch to provide cost effectiveness to meet the budget needs for this size location. The remote medium site is connected to the main site location through Metro service. At the remote medium site, the WAN and LAN aggregation platform is the Cisco Catalyst 4507 switch. This switch has the necessary features to perform as WAN router. However, if there is the need for advanced WAN features such as MPLS, the Cisco Catalyst 3750 ME, Cisco ISR Series router or upgrading to the Cisco Catalyst 6500 series could be explored as an option. For this design, the Cisco Catalyst 4500 Series switches has been chosen to perform the dual functionality as WAN router, in addition to its role as core-layer LAN switch.

## Small Remote Site WAN Aggregation Platform Selection

The small remote site is connected to main site using a private leased-line service. The WAN speed between the small remote site and the main site location is assumed to be around 20Mbps, and this service is provided by a traditional leased line. Since it is a leased-line circuit, WAN devices such as Cisco 3750 Metro or 4507 switch cannot be used. Therefore, an integrated services router is needed to meet the requirement. For this reason, the Cisco 3845 Series router is chosen as the WAN platform for the small remote site. The main advantages of using the Cisco 3845 Series router are as follows:

- Enhanced Network Module Slot
- Support for over 90 existing and new modules
- Voice Features: Analog and digital voice call support and optional voice mail support
- Support for majority of existing AIMs, NMs, WICs, VWICs, and VICs
- Integrated GE ports with copper and fiber support

## Implementation of WAN Reference Design

The following section discusses the implementation details for the Medium Enterprise Design Profile. The major components of the implementation are the following:

- WAN infrastructure design
- Routing
- QoS
- Resiliency
- Multicast

## WAN Infrastructure Design

As explained in the design considerations (**where??? in chapter 1??**), the Medium Enterprise Design Profile uses two different services to connect the remote site locations to the main site location. The large remote site and medium remote site would connect to main site using Metro Ethernet services. The small remote site uses a leased-line service to connect to the main site location. The large remote site, due to its size, is recommended to have 1Gbps Metro service to the main site where as the small remote site location is recommended to have at least 20Mbps of bandwidth to main site. The following section provides the configuration details of all the WAN devices needed to establish the WAN connectivity.

### Configuration of WAN Interfaces at WAN Aggregation Router 2

The following is configuration of WAN interfaces on WAN aggregation router 2, which aggregates all the connections from the remote site locations to the main site:

```
interface GigabitEthernet0/2/0
  description Connected to cr11-3750ME-RLC
  ip address 10.126.0.1 255.255.255.254
!
interface GigabitEthernet0/2/1
  description Connected to cr11-4507-RMC
  dampening
  no ip address
  load-interval 30
  carrier-delay msec 0
```

```
negotiation auto
cdp enable
service-policy output PARENT_POLICY
hold-queue 2000 in
hold-queue 2000 out
!
interface GigabitEthernet0/2/1.102
encapsulation dot1Q 102
ip address 10.126.0.3 255.255.255.254
!
!
```

### Configuration of WAN Interface at 3750 Large Remote Site

The following is configuration of WAN interface at the 3750 large remote site switch, which is connected to main site:

```
interface GigabitEthernet1/1/1
description Connected to cr11-ASR-WE
no switchport
dampening
ip address 10.126.0.0 255.255.255.254
```

### Configuration of WAN interface at 4500 Medium Remote Site

The following is the configuration of WAN interface at the medium remote site connected to the main site:

```
interface GigabitEthernet4/1
description link connected to cr13-6500-pe2 gi3/2
switchport trunk native vlan 802
switchport trunk allowed vlan 102
switchport mode trunk
logging event link-status
load-interval 30
carrier-delay msec 0
no cdp enable
spanning-tree portfast trunk
spanning-tree guard root
!
interface Vlan102
description Connected to cr11-ASR-WE
dampening
ip address 10.126.0.2 255.255.255.254
load-interval 30
carrier-delay msec 0
```

## Leased-Line Service

The WAN bandwidth requirement for a small remote site is assumed to be 20Mbps. Cisco recommends that the small remote site connect to the main site using a private leased-line service. The leased-line service is more readily available for these type of locations and the bandwidth is sufficient for the small remote site application requirements. To implement this design, a serial SPA is needed on the ASR 1006 WAN aggregation router at the main site and this SPA needs to be enabled for T3 interface type. The configuration below illustrates how to enable and configure the T3 interface.

The following configuration steps are needed to build the lease-line service between the main site and small remote site:

**Step 1** Enable the T3 interface on the SPA on ASR1006:

```
card type t3 0 3
```

**Step 2** Configure the WAN interface:

```
interface Serial0/3/0
dampening
ip address 10.126.0.5 255.255.255.254
```

## Configuration of WAN Interface at Small Remote Site Location

The following is configuration of WAN interface at the small remote site location:

```
interface Serial2/0
dampening
ip address 10.126.0.4 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
ip pim sparse-mode
service-policy output RSC_PARENT_POLICY
ip summary-address eigrp 100 10.124.0.0 255.255.0.0 5
load-interval 30
carrier-delay msec 0
dsu bandwidth 44210
```

## Routing Design

This section discusses how routing is designed and implemented in the Medium Enterprise Design Profile. As indicated in the WAN transport design, the Medium Enterprise Design Profile has multiple transports—Private WAN, Internet, Metro Service, and leased-line services. The private network would provide access to reach other remote sites globally. Internet service would help the medium enterprise to reach Internet. Metro/leased-line service would help to connect remote site locations to the main site. To provide connectivity using these transport services we have designed two distinct routing domains—external and internal. The external routing domain is where the medium enterprise would connect with external autonomous system, and the internal routing domain is where the entire routing domain is within single autonomous system. The following section discusses about the external routing domain design, and the internal routing domain design.

### External Routing Domain

As indicated above, the external routing domain would connect with different service providers, Private WAN, and the Internet service. This is applicable only to the WAN aggregation router 1, which interfaces with both Private WAN, and the Internet service, because it is the only router which interfaces with the external domain.

The main design considerations for routing for the Internet/private WAN edge router are as follows:

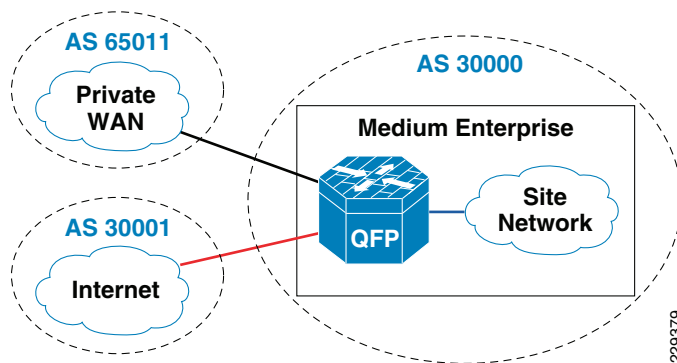
- Scale up to large number of routes
- Support for multi-homing—connection to different service providers
- Ability to implement complex policies—Have separate policies for incoming and outgoing traffic

To meet the above requirements, BGP has been chosen as the routing protocol because of the following reasons:

- *Scalability*—BGP is far superior when routing table entries are quite large.
- *Complex policies*—IGP protocol is better in environments where the neighbors are trusted, whereas when dealing with different service providers' complex policies are needed to deal with incoming and outgoing entries. BGP supports having different policies for incoming and outgoing prefixes.

Figure 3-8 shows the BGP design.

**Figure 3-8** BGP Design in Medium Enterprise



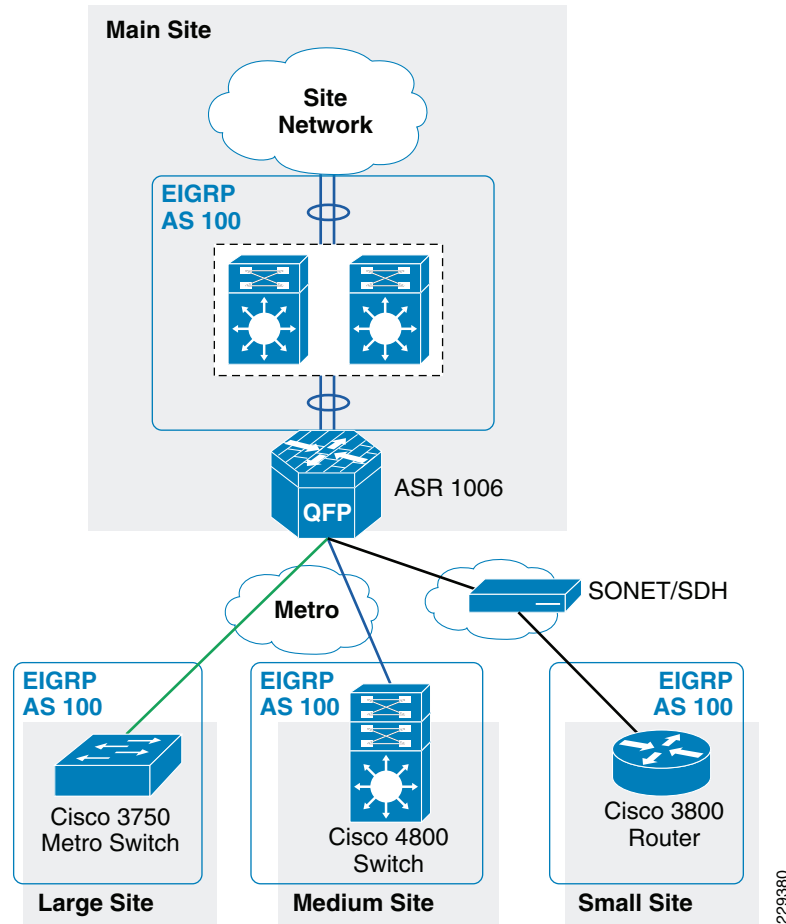
For more information on designing and configuring BGP on the Internet border router, refer to the *SAFE Reference Design* at the following link:

<http://www.cisco.com/en/US/netsol/ns954/index.html#~five>

## Internal Routing Domain

EIGRP is chosen as the routing protocol for designing the internal routing domain, which is basically connecting all the devices in the site network. EIGRP is a balanced hybrid routing protocol that builds neighbor adjacency and flat routing topology on per autonomous-system (AS)-basis. It is important to design EIGRP routing domain in site infrastructure with all the design principles defined earlier in this section. The Medium Enterprise Design Profile network infrastructure must be deployed in recommended EIGRP protocol design to secure, simplify, and optimize the network performance.

Figure 3-9 depicts the design of EIGRP for internal network.

**Figure 3-9 EIGRP Design Diagram****EIGRP Configuration on WAN Aggregation Router2 –ASR1006**

The EIGRP is used on the following links:

1. Port-channel link, which is link between the ASR1006 router and the core
2. The 1Gbps Metro link to large remote site location
3. The 100Mbps Metro link to medium remote site location
4. 20Mbps leased-line service to small remote Site location

**Step 1** Configure the neighbor authentication on interface links:

```
interface Port-channel1
 ip address 10.125.0.23 255.255.255.254
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-key
!
interface GigabitEthernet0/2/0
 description Connected to cr11-3750ME-RLC
 ip address 10.126.0.1 255.255.255.254
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-key
!
```

```

interface GigabitEthernet0/2/1
description Connected to cr11-4507-RMC
dampening
no ip address
load-interval 30
carrier-delay msec 0
negotiation auto
cdp enable
hold-queue 2000 in
hold-queue 2000 out
!
interface GigabitEthernet0/2/1.102
encapsulation dot1Q 102
ip address 10.126.0.3 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
!
interface Serial0/3/0
dampening
ip address 10.126.0.5 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key

```

**Step 2** Configure the summarization on the member links:

```

interface Port-channel1
ip address 10.125.0.23 255.255.255.254
ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5
!
interface GigabitEthernet0/2/0
description Connected to cr11-3750ME-RLC
ip address 10.126.0.1 255.255.255.254
ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5
!
interface GigabitEthernet0/2/1
description Connected to cr11-4507-RMC
!
interface GigabitEthernet0/2/1.102
encapsulation dot1Q 102
ip address 10.126.0.3 255.255.255.254
ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5
!
interface Serial0/3/0
ip address 10.126.0.5 255.255.255.254
ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5

```

**Step 3** Configure EIGRP routing process:

```

router eigrp 100
network 10.0.0.0
eigrp router-id 10.125.200.24
no auto-summary
passive-interface default
no passive-interface GigabitEthernet0/2/0
no passive-interface GigabitEthernet0/2/1.102
no passive-interface Serial0/3/0
no passive-interface Port-channel1
nsf

```

The ASR1006 router is enabled with nonstop forwarding feature. The following command is used to verify the status:

```
cr11-asr-we#show ip protocols
```



```

*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  EIGRP NSF-aware route hold timer is 240s
  EIGRP NSF enabled
    NSF signal timer is 20s
    NSF converge timer is 120s
    Time since last restart is 2w1d
  Automatic network summarization is not in effect
  Address Summarization:
    10.126.0.0/16 for Port-channel1, GigabitEthernet0/2/0, GigabitEthernet0/2/1.102
    Serial0/3/0
    Summarizing with metric 2816
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Passive Interface(s):
    GigabitEthernet0/2/1
    GigabitEthernet0/2/2
    GigabitEthernet0/2/3
    GigabitEthernet0/2/4
    Serial0/3/1
    Group-Async0
    Loopback0
    Tunnel0
  Routing Information Sources:
    Gateway         Distance      Last Update
    (this router)    90           2w1d
    10.125.0.22      90           1d17h
    10.126.0.4       90           1d17h
    10.126.0.0       90           1d17h
    10.126.0.2       90           1d17h
  Distance: internal 90 external 170

cr11-asr-we#

```

## EIGRP Configuration on 3750 Large Remote Site Switch

The EIGRP configuration at the 3750 large remote site also has similar steps compared to main site.

### Step 1 Enable authentication on the link:

```

interface GigabitEthernet1/1/1
  description Connected to cr11-ASR-WE
  no switchport
  dampening
  ip address 10.126.0.0 255.255.255.254
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 eigrp-key
  router eigrp 100
  network 10.0.0.0
  passive-interface default
  no passive-interface Port-channel1

```

```
no passive-interface GigabitEthernet1/1/1
eigrp router-id 10.122.200.1
```

**Step 2** Configure summarization on the link:

```
interface GigabitEthernet1/1/1
description Connected to cr11-ASR-WE
no switchport
dampening
ip address 10.126.0.0 255.255.255.254
ip summary-address eigrp 100 10.122.0.0 255.255.0.0
```

**Step 3** Configure EIGRP routing process:

```
router eigrp 100
network 10.0.0.0
passive-interface default
no passive-interface Port-channel1
no passive-interface GigabitEthernet1/1/1
eigrp router-id 10.122.200.1
!
```

**EIGRP Configuration at 4750 Medium Site Switch****Step 1** Enable authentication on the WAN link:

```
interface Vlan102
description Connected to cr11-ASR-WE
dampening
ip address 10.126.0.2 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
Step2) Enable summarization on the WAN links
interface Vlan102
ip summary-address eigrp 100 10.123.0.0 255.255.0.0 5
load-interval 30
carrier-delay msec 0
```

**Step 2** Enable EIGRP routing process:

```
router eigrp 100
passive-interface default
no passive-interface Vlan102
no auto-summary
eigrp router-id 10.123.200.1
network 10.98.0.1 0.0.0.0
network 10.123.0.0 0.0.255.255
network 10.126.0.0 0.0.255.255
nsf
!
```

**EIGRP Configuration at 3800 Small Remote Site Router****Step 1** Configure link authentication:

```
interface Serial2/0
dampening
ip address 10.126.0.4 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
```

```

Step2) Configure Summarization
interface Serial2/0
dampening
ip summary-address eigrp 100 10.124.0.0 255.255.0.0 5
load-interval 30
carrier-delay msec 0
dsu bandwidth 44210

```

## Step 2 Configure EIGRP process:

```

router eigrp 100
network 10.0.0.0
no auto-summary
eigrp router-id 10.124.200.1
!

```

To obtain more information about EIGRP design, refer to the [“Designing an End-to-End EIGRP Routing Network”](#) section on page 2-52.

## QoS


QoS is a part of foundation services, which is very critical to the application performance. The traditional applications such as voice, video, and data together with newer applications such as broadcast video, real-time video, video surveillance, and many other applications have all converged into IP networks. Moreover, each of these applications require different performance characteristics on the network. For example, data applications may need only high throughput, but are tolerant to delay and loss. Similarly, voice applications need constant low bandwidth and low delay performance. To cater to these performance characteristics, Cisco IOS has several robust QoS tools such as classification and marking, queuing, WRED, policing, shaping, and many other tools to effect the traffic characteristics. Before discussing the QoS design, the following subsection provides a brief introduction on these characteristics.

### Traffic Characteristics

The main traffic characteristics are bandwidth, delay, loss, and jitter.

- **Bandwidth**—Lack of proper bandwidth can cause applications from performing poorly. This problem would be exacerbated if there were more centralized applications. The bandwidth constraint occurs because of the difference between the bandwidth available at LAN and the WAN. As shown in [Figure 3-10](#), the bandwidth of the WAN transport dictates the amount of traffic received at each remote site. Applications are constrained by the amount of WAN bandwidth.

**Figure 3-10 Bandwidth Constraint Due to Difference in Speeds**



- **Jitter**—Occurs when there are bandwidth mismatches between the sender and receiver, which could result in poor performance of delay sensitive applications like voice and video.

- *Loss*—occurs when the queues become full, and there is not enough bandwidth to send the packets.
- *Delay*—Is an important characteristic, which plays a large role in determining the performance of the applications. For a properly designed voice network, the one-way delay must be less than 150 msec.

## QoS Design for WAN Devices

For any application regardless of whether it is video, voice, or data, the traffic characteristics discussed above need to be fully understood before making any decisions on WAN transport or the platforms needed to deploy these services. Cisco QoS tools help to optimize these characteristics so that voice, video, and data applications performance is optimized. The voice and video applications are highly delay-and drop-sensitive, but the difference lies in the bandwidth requirement. The voice applications have a constant and low bandwidth requirement, but the video applications have variable bandwidth requirements. Therefore, it is important to have a good QoS policy to accommodate these applications.

Regardless of the WAN transport chosen, QoS design is the most significant factor in determining the success of network deployment. There are many benefits in deploying a consistent, coherent QoS scheme across all network layers. It helps not only in optimizing the network performance, it helps to mitigate network attacks and manage the control plane traffic. Therefore, when the platforms are selected at each network layer, QoS must always be considered in the design choice.

In the WAN links, the congestion can occur when there are speed mismatches. This may occur because there is significant difference between LAN speeds and WAN speeds. To prevent that from occurring, the following two major tools can be used:

- Low-Latency Queuing (LLQ), which is used for highest-priority traffic (voice/ video).
- Class-based Weighted-Fair Queuing (CBWFQ), which can be used for guaranteeing bandwidth to data applications.

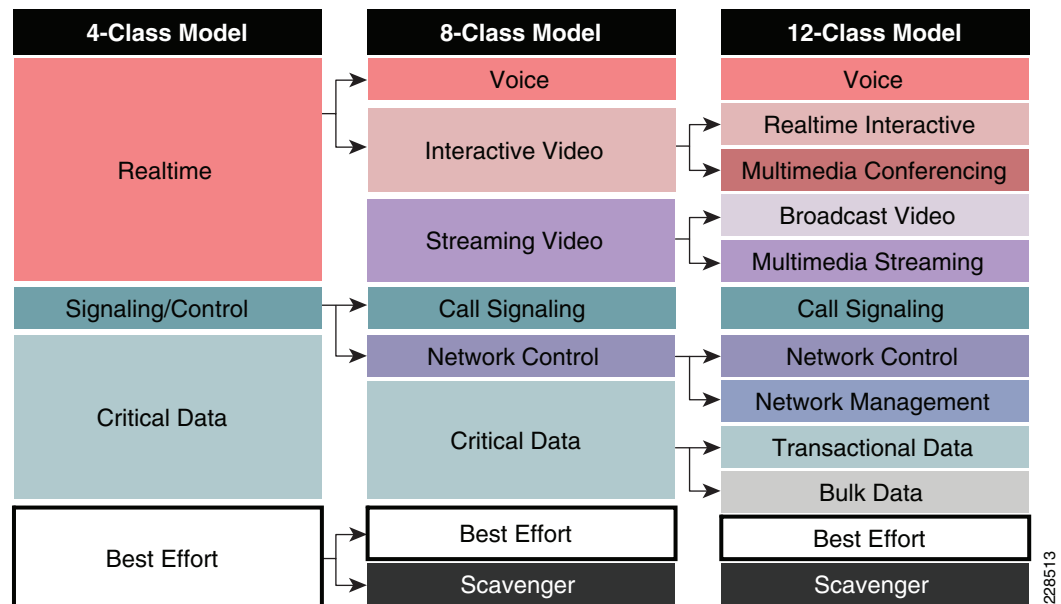
The general guidelines for deploying the WAN edge device considerations are as follows:

- For WAN speeds between 1Mbps to 100Mbps, use hierarchical policies for sub-line-rate Ethernet connections to provide shaping and CBWFQ/LLQ.
- For WAN speeds between 100Mbps to 10Gbps, use ASR1000 with QFP or hardware queuing via Cisco Catalyst 3750-Metro and 6500/7600 WAN modules.

When designing the QoS for WAN architecture, there are two main considerations to start with:

- Whether the service provider will provide four classes of traffic
- Whether the service provider will only provide one class of traffic

This document assumes that the service provider will support at least 4 classes of traffic such as REAL\_TIME, GOLD, SILVER, and DEFAULT. The Medium Enterprise site LAN supports 12 classes of traffic, which will be mapped to 4 classes of traffic on the WAN side. [Figure 3-11](#) illustrates the recommended markings for different application traffic.

**Figure 3-11 Mapping of 12-Class Model to 4-Classes**

Once the QoS policy is designed, the next pertinent question is the appropriate allocation of bandwidth for the 4 classes of traffic. [Table 3-1](#) describes the different classes, the percentage, and actual bandwidth allocated for each class of traffic.

**Table 3-1 Classes of Traffic**

Class of Traffic	4-class SP Model	Bandwidth Allocated	Actual Bandwidth
Voice, Broadcast Video, Real Time Interactive	SP- Real-Time	30%	33 Mbps
Network Control Signaling Transactional Data	SP-Critical 1	20%	36 Mbps
Multi-media Conferencing Multimedia streaming OAM	SP-Critical 2	20%	25 Mbps
Bulk data Scavenger Best Effort	SP-Best Effort	30%	6 Mbps

## QoS Implementation

This section discusses how QoS is implemented in Medium Enterprise Design Profile. As explained in the QoS design considerations, the main objective of the QoS implementation is to ensure that the 12 classes of LAN traffic is mapped into 4 classes of WAN traffic. Each class should receive the adequate bandwidth, and during congestion, each class must received the guaranteed minimum bandwidth. To accomplish this objective, the following methods are used to implement QoS policy:

- *Three-layer hierarchical design*—This is needed when multiple sites need to share a common bandwidth, and each site needs dedicated bandwidth, and queuing within the reserved policy.
- *Two-layer hierarchical design*—This design is needed when the interface bandwidth is higher than the SLA bandwidth allocated by the service provider. For example, if the physical link is 100Mbps, but the service provider has only allocated 50 Mbps. In this scenario we need two policies. The first policy, which is parent policy would shape the entire traffic to 50Mbps then the child policy would queue and allocated bandwidth for each class.
- *Single-layer design*—If the interface bandwidth, and the SLA bandwidth of the provider are equal then we can use a single QoS policy to share the bandwidth among the classes of traffic, which is four in our design.

This section describes detailed implementation of QoS policies at various parts of the network. The devices that need QoS design are as follows:

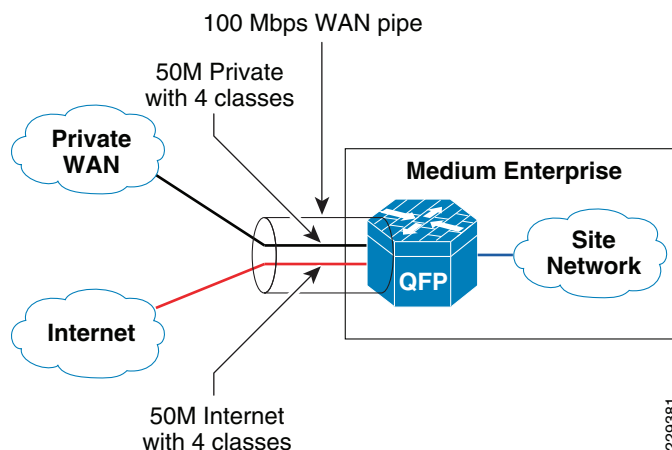
- WAN aggregation router 1 for connection to the Internet and PRIVATE WAN network
- WAN aggregation router 2 for connection to remote site
- Cisco 3750 Metro switch at the large remote site
- Cisco 4500 switch at the medium remote site
- Cisco 3800 router at the small remote site

## QoS Implementation at WAN Aggregation Router 1

The WAN aggregation router 1 connects to two different providers: private WAN network and Internet. It is assumed that the aggregate bandwidth is 100Mbps that should be shared between both services—50Mbps is dedicated for private WAN network and 50Mbps is dedicated for Internet traffic. As explained in the previous section, to implement this granular policy, a three-layer hierarchical QoS design needs to be used.


Figure 3-12 depicts the bandwidth allocation at the WAN aggregation router 1.

**Figure 3-12 The Bandwidth Allocation at WAN Aggregation Router 1**



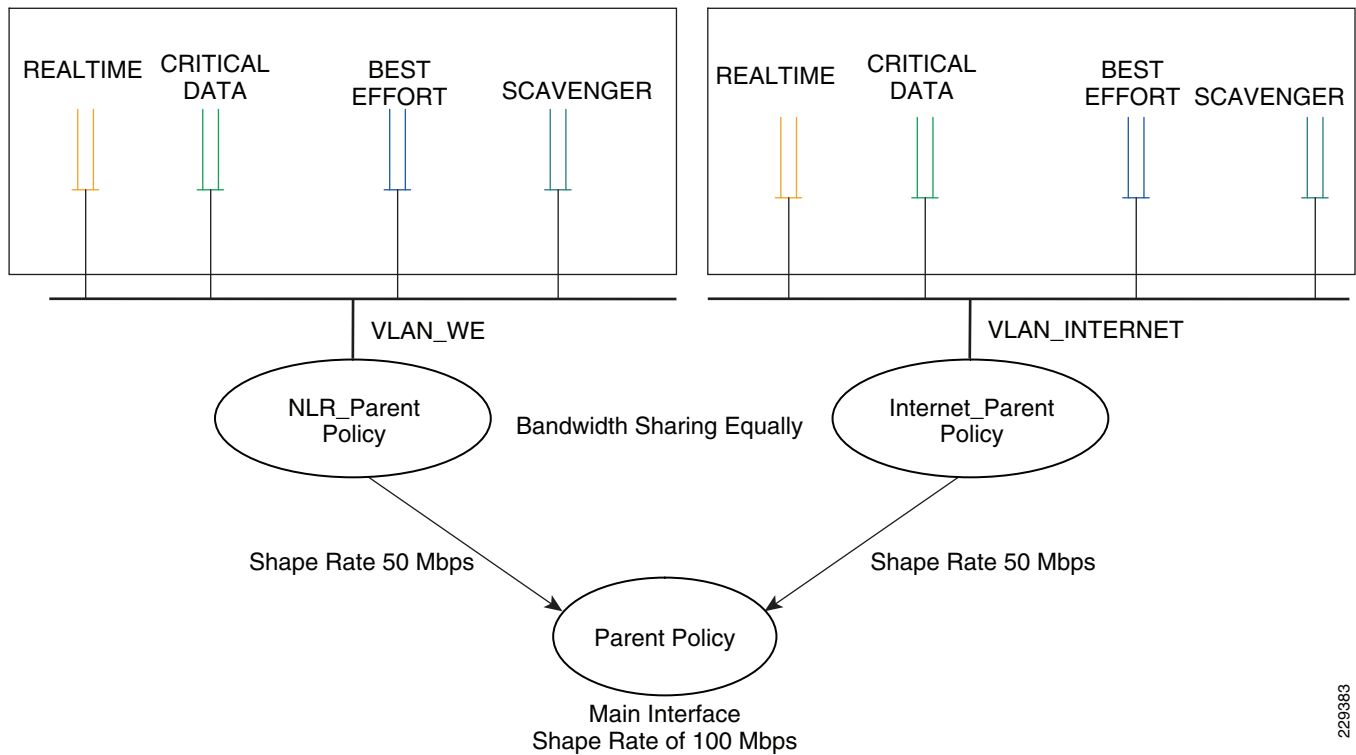
To implement a three-layer hierarchical QoS policy on the WAN aggregation1 router, a higher-level parent policy is defined that would shape the aggregate WAN speed to 100Mbps, then subparent policies are defined, which would further shape it to 50Mbps. Within each of the subparent policies, there are four defined classes: REALTIME, CRITICAL\_DATA, BEST\_EFFORT, and SCAVENGER classes. Figure 3-13 depicts this hierarchical QoS design.

**Figure 3-13 Hierarchical QoS Design**



The hierarchical three-layer QoS policy is implemented in three steps as follows:

- 
- Step 1** Define parent policy—Enforces the aggregate bandwidth policy for the entire interface. This is like a grandfather of policy.
  - Step 2** Define the individual subparent policies—These would be specific to each service type. For example, PRIVATE WAN\_PARENT is a policy dedicated for PRIVATE WAN traffic, and PRIVATE WAN\_Internet is specific to Internet traffic.
  - Step 3** Define the child policies—Classifies, queues, and allocate bandwidth within each subparent policy. For example, PRIVATE WAN\_PARENT would have a PRIVATE WAN\_Child policy that would classify, queue, and allocate the bandwidth within each allocated bandwidth. The following diagram shows the hierarchical allocation.



229383

## Implementation Steps for QoS Policy at WAN Aggregation Router 1

This section would describes the detailed steps needed to implement the three-layer QoS policy in the WAN\_Aggregation\_router1.

### Step 1 Define the class-maps.

```
class-map match-all REALTIME
match ip dscp cs4 af41 cs5 ef

class-map match-all CRITICAL_DATA
match ip dscp af11 af21 cs3 cs6

class-map match-all BEST_EFFORT
match ip dscp default

class-map match-all SCAVENGER
match ip dscp cs2
```

228926

### Step 2 Define the child policy maps.



```

policy-map IE_CHILD_POLICY
class REALTIME
  priority percent 33
class CRITICAL_DATA
  bandwidth remaining ratio 6
class SCAVENGER
  bandwidth remaining ratio 1
class BEST_EFFORT
  bandwidth remaining ratio 4

policy-map NLR_CHILD_POLICY
class REALTIME
  priority percent 33
class CRITICAL_DATA
  bandwidth remaining ratio 6
class BEST_EFFORT
  bandwidth remaining ratio 4
class SCAVENGER
  bandwidth remaining ratio 1

```

228927

**Step 3** Define the parent policy maps.

```

class-map match-all dummy
!
policy-map PARENT_POLICY
class dummy service-fragment share
  shape average 10000000

policy-map NLR_PARENT_POLICY
class class-default fragment share
  shape average 50000000
  service-policy NLR_CHILD_POLICY

policy-map IE_PARENT_POLICY
class class-default fragment share
  shape average 50000000
  service-policy IE_CHILD_POLICY
!

```

class-map match-all dummy → Dummy class does not classify anything  
 policy-map PARENT\_POLICY  
 class dummy service-fragment share → Defining service-fragment would allow other policies to point for share of bandwidth.  
 shape average 10000000 → The parent policy would shape to 100 Mbps.  
 policy-map NLR\_PARENT\_POLICY  
 class class-default fragment share  
 shape average 50000000 → Parent policy allocates 50% of bandwidth  
 service-policy NLR\_CHILD\_POLICY → Child policy gets attached to parent policy  
 policy-map IE\_PARENT\_POLICY  
 class class-default fragment share  
 shape average 50000000  
 service-policy IE\_CHILD\_POLICY  
 !

228928

**Step 4** Apply the policy maps created in Steps 1 to 3.

```

interface GigabitEthernet1/0/0
  dampening
  no ip address
  load-interval 30
  carrier-delay msec 0
  negotiation auto
  service-policy output PARENT_POLICY
  hold-queue 2000 in
  hold-queue 2000 out
!
interface GigabitEthernet1/0/0.65
  description link to 6500
  encapsulation dot 1Q.65
  ip address 64.104.10.113 255.255.255.252
  service-policy output IE_PARENT_POLICY
!
interface GigabitEthernet1/0/0.75
  description link to 6500
  encapsulation dot 1Q.75
  ip address 64.104.10.125 255.255.255.252
  service-policy output NLR_PARENT_POLICY
!

```

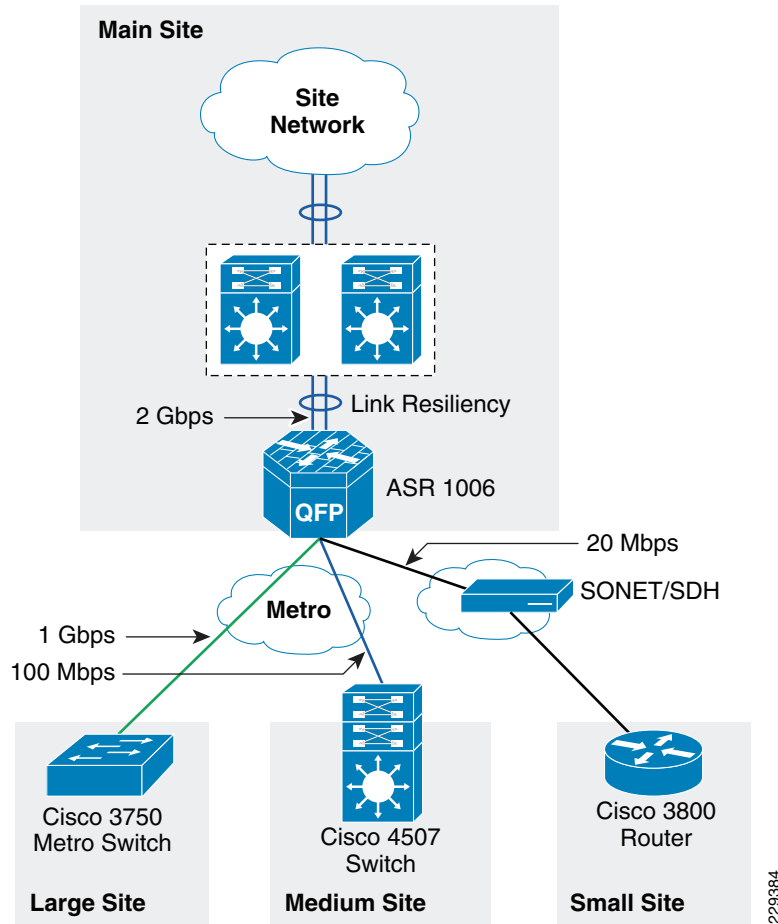
Aggregate policy (grand-father) applied on main interface

The parent policy applied on sub-interface

228929

## QoS Policy Implementation for WAN Aggregation Router 2

QoS configuration at WAN aggregation router 2 is more complex than the QoS configuration of WAN aggregation router 1 because of different speeds connected to the router. [Figure 3-14](#) depicts the different types of WAN speeds

**Figure 3-14** WAN Link Speeds at WAN Aggregation Router 2 Device

The requirements of the QoS design at the WAN aggregation router 2 are as follows:

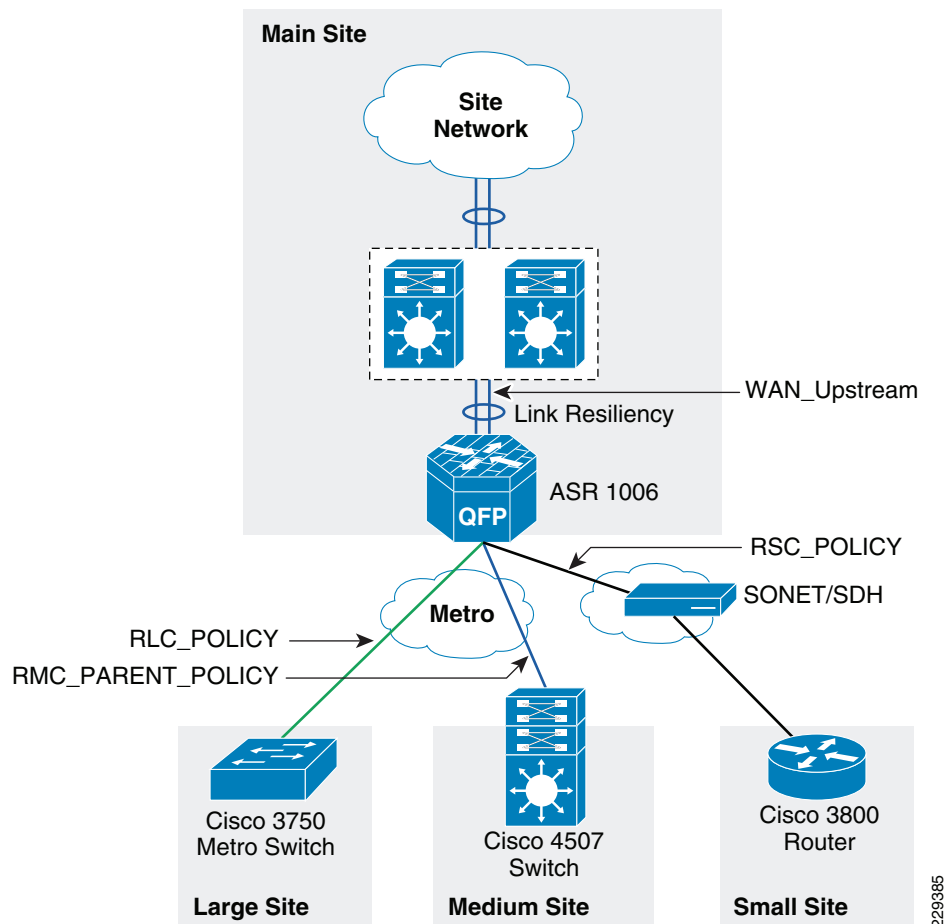
- The link speed between the main site and large site is 1Gbps. Therefore, a single-layer QoS policy can be defined on the link.
- The SLA between the main site and remote medium site is assumed to be 100Mbps; however, the link speed is assumed to be 1Gbps. In addition, there is an assumption that there could be more than one remote medium site present in this design. Therefore, each medium remote site would connect to the main site using these 100Mbps links, requiring a three-layer hierarchical QoS policy is needed. The link between the main site and small remote site is 20Mbps. The physical link speed is 44Mbps, requiring a two-level hierarchical QoS policy is needed.
- The EtherChannel link between the ASR router and the core is 2Gbps, which contains two links of 1Gbps link speeds. Since the physical link speed and the actual WAN speed is 1Gbps, a single-level QoS policy can be applied on each of the links.

Table 3-2 describes the different QoS policy names applied at the WAN aggregation router 2.

**Table 3-2 QoS Policy for WAN Aggregation Route 2**

QoS Policy Name	Description	WAN Speed
RLC_POLICY	Applied on link between Main Site, and Large Remote Site	1Gbps
PARENT_POLICY RMC_PARENT_POLICY RMC_CHILD_POLICY	Hierarchical Qos Policy between the Main Site, and Medium Remote Site location.	100 Mbps
WAN_Upstream	Applied on link between Main Site, and core	2Gbps
RSC_PARENT_POLICY RSC_POLICY	Applied on link between Main Site and small site	20Mbps

Figure 3-15 depicts the various points where QoS policies are applied.

**Figure 3-15 The allocation of QoS Policy at Different Places on WAN Aggregation Router 2**

## QoS Policy Between the Main Site and Large Remote Site

The WAN physical link speed is 1Gbs. Also, the actual SLA between the main site and the large remote site is assumed to be 1Gbps. Therefore, a single-layer QoS policy is implemented in this scenario.

**Step 1** Define the class-maps.

```
class-map match-all REALTIME
 match ip dscp cs4 af41 cs5 ef
class-map match-all CRITICAL_DATA
 match ip dscp af11 af21 cs3 af31 cs6
class-map match-all BEST_EFFORT
 match ip dscp default
class-map match-all SCAVENGER
 match ip dscp cs2
```

228932

**Step 2** Define the policy map.

```
policy-map RLC_POLICY
 class REALTIME
  priority percent 33
  set cos 5
 class CRITICAL_DATA
  bandwidth remaining ratio 6
  set cos 3
 class SCAVENGER
  bandwidth remaining ratio 1
  set cos 0
 class BEST_EFFORT
  bandwidth remaining ratio 4
  set cos 2
!
```


228933

**Step 3** Apply the class-maps and policy map defined in Steps 1 and 2 on the interface connected between the main site to the large site.

```

interface GigabitEthernet0/2/0
description Connected to cr11-3750ME-RLC
ip address 10.126.0.1 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
ip pim sparse-mode
ip summary-address eigrp 100 10.126.0.0 255.255.0.0
logging event link-status
load-interval 30
negotiation auto
service-policy output RLC_POLICY
!

```

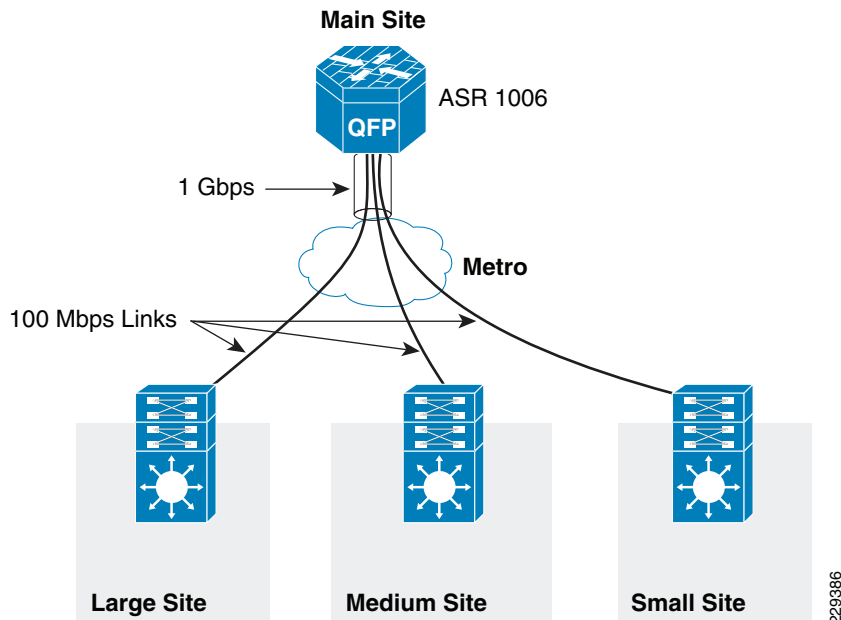


228934

## QoS Policy Between the Main Site and Medium Remote Site Location

A three-layer QoS design is needed between the main site and large remote medium site location, because there could be a couple of remote medium site locations connected on a single metro link to the main site. [Figure 3-16](#) shows how this design looks like when there are more than one remote medium site.

**Figure 3-16** The WAN Link Design for Connectivity Between Main Site and Medium Remote Site



229386

Here, the implementation details are provided for only a single medium site location; however, more medium site locations could be added, if desired.

The following are implementation steps for this QoS policy:

**Step 1** Define the child policy maps.

```

policy-map RMC_CHILD_POLICY
class REALTIME
  priority percent 33
  set cos 5
class CRITICAL_DATA
  bandwidth remaining ratio 6
  set cos 3
class SCAVENGER
  bandwidth remaining ratio 1
  set cos 0
class BEST_EFFORT
  set cos 2
  
```

228936

**Step 2** Define the parent policy maps.

```

class-map match-all dummy
!
policy-map PARENT_POLICY
class dummy service-frgment share
  shape average 10000000
  
```

→ Sets the total bandwidth to 1G

```

policy-map RMC_PARENT_POLICY
class class-default fragment share
  shape average 10000000
  service-policy RMC_CHILD_POLICY
  
```

→ Sets the bandwidth for single medium campus to 100Mbps

228937

**Step 3** Apply the policy maps.

```

interface GigabitEthernet0/2/1
description Connected to cr11-4507-RMC
dampening
no ip address
load-interval 30
carrier-delay msec 0
negotiation auto
cdp enable
service-policy output PARENT_POLICY
hold-queue 2000 in
hold-queue 2000 out
!

```

```

interface GigabitEthernet0/2/1.102
encapsulation dot1Q 102
ip address 10.126.0.3 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
ip pim sparse-mode
ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5
service-policy output RMC_PARENT_POLICY

```

```

policy-map RMC_PARENT_POLICY
class class-default fragment share
shape average 100000000
service-policy RMC_CHILD_POLICY

```

First level policy applied  
to main interface

Second level policy applied  
to sub-interface

Third level policy applied  
to main parent policy

Large Campus



Metro



Medium Campus

## QoS Policy Between Main Site and Small Remote Site Location

The following is the QoS policy implementation steps between main site and small remote site location. The actual WAN speed is 44Mbps; however, the SLA is assumed to be 20Mbps. Therefore, a two-layer hierarchical QoS design is needed to implement the above policy.

### Step 1 Define the policy map.

```

policy-map RSC_POLICY
class REALTIME
priority percent 33
class CRITICAL_DATA
bandwidth remaining ratio 6
class SCAVENGER
bandwidth remaining ratio 1
class BEST_EFFORT
bandwidth remaining ratio 4
!

```

```

policy-map RSC_PARENT_POLICY
class class-default
shape average 20000000
service-policy RSC_POLICY

```

### Step 2 Apply the policy map to the interface.




```

interface Serial0/3/0
  dampening
  ip address 10.126.0.5 255.255.255.254
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 eigrp-key
  ip pim sparse-mode
  ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5
  logging event link-status
  load-interval 30
  carrier-delay msec 0
  dsu bandwidth 44210
  framing c-bit
  cablelength 10
  service-policy output RSC_PARENT_POLICY
end

cr11-asr-we#

```

Apply parent policy  
to interface



## QoS Policy Implementation Between the Main Site and Core

The following is the QoS policy implementation between main site and core. There are two links between the ASR 1006 and core, which is VSS. QoS policy needs to be configured on both links.

**Step 1** Define the policy-map.

```

policy-map WAN_Upstream
  class REALTIME
    priority percent 33
  class CRITICAL_DATA
    bandwidth remaining ratio 6
  class SCAVENGER
    bandwidth remaining ratio 1
  class BEST_EFFORT
    bandwidth remaining ratio 4

```

228941

**Step 2** Apply the policy-map on both interfaces going up to the core.

```

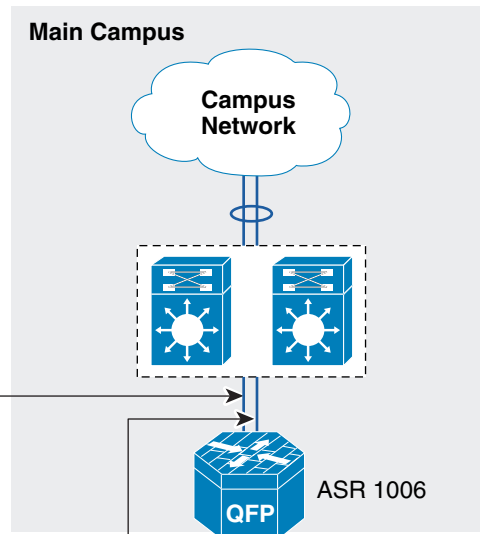
policy-map WAN_Upstream
class REALTIME
  priority percent 33
class CRITICAL_DATA
  bandwidth remaining ratio 6
class SCAVENGER
  bandwidth remaining ratio 1
class BEST_EFFORT
  bandwidth remaining ratio 4

```

```

interface GigabitEthernet0/2/3
dampening
no ip address
load-interval 30
carrier-delay msec 0
negotiation auto
cdp enable
service-policy output WAN_Upstream
channel-group 1 mode active
hold-queue 2000 in
hold-queue 2000 out
!
interface GigabitEthernet0/2/4
dampening
no ip address
load-interval 30
carrier-delay msec 0
negotiation auto
cdp enable
service-policy output WAN_Upstream
channel-group 1 mode active
hold-queue 2000 in
hold-queue 2000 out

```



228942

## QoS Policy Between Large Remote Site and Main Site Location

The WAN interface between the large remote site and main site is 1 Gbps, which is also equal to the link speed; therefore, a single-layer QoS policy map can be created.

### Step 1 Define the class-maps.

```

class-map match-all REALTIME
match ip dscp cs4 af41 cs5 ef
class-map match-all CRITICAL_DATA
match ip dscp af11 cs2 af21 cs3 af31 cs6
class-map match-all BEST_EFFORT
match ip dscp default
class-map match-all SCAVENGER
match ip dscp cs1

```

228943

### Step 2 Define the policy-map.

```

policy-map ME_POLICY
class REALTIME
  priority
  police 220000000 8000 exceed-action drop —> The realtime traffic get 330 Mbps
  set cos 5
class CRITICAL_DATA
  bandwidth remaining ratio 40
  set cos 3
class BEST_EFFORT
  bandwidth remaining ratio 35
  set cos 2
class SCAVENGER
  bandwidth remaining ratio 25
  set cos 0
!
!

```


228944

**Step 3** Apply the QoS policy-map to the WAN interface.

```

interface GigabitEthernet1/1/1
description Connected to cr11-ASR-WE
no switchport
dampening
ip address 10.126.0.0 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
ip pim sparse-mode
ip summary-address eigrp 100 10.122.0.0 255.255.0.0
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 30 35 5
priority-queue out
mls qos trust dscp
service-policy output ME_POLICY
hold-queue 2000 in
hold-queue 2000 out
!

```



228945

## QoS Policy Between Remote Medium Site and Main Site Location

The remote medium site location uses 4500 as WAN device, which uses 4500-E supervisor. The physical link speed is 100Mbps and the actual SLA is also 100Mbps. Therefore, a single-layer QoS policy meets the requirement.

### Step 1 Define the class-maps.

```
class-map match-all REALTIME
 match ip dscp cs4 af41 cs5 ef
class-map match-all CRITICAL_DATA
 match ip dscp af11 cs2 af21 cs3 af31 cs6
class-map match-all BEST_EFFORT
 match ip dscp default
class-map match-all SCAVENGER
 match ip dscp cs1
```

228946

### Step 2 Define the policy-maps.

```
policy-map RMC_POLICY
class REALTIME
 priority
 police cir 33000000
 conform-action transmit
 exceed-action drop
 set cos 5
class CRITICAL_DATA
 set cos 3
 bandwidth percent 36
class SCAVENGER
 bandwidth percent 5
 set cos 0
class BEST_EFFORT
 set cos 2
 bandwidth percent 25
!
```


228947

### Step 3 Apply the defined class and policy maps to the interface.

```

interface GigabitEthernet4/1
description link connected to cr13-6500-pe2 gi3/2
switchport trunk native vlan 802
switchport trunk allowed vlan 102
switchport mode trunk
logging event link-status
load-interval 30
carrier-delay msec 0
no cdp enable
spanning-tree portfast trunk
spanning-tree guard root
service-policy output RMC_POLICY
!

```



## QoS Policy Implementation Between Small Remote Site and Main Site Location

This section describes the QoS policy implementation between the small remote site location and the main site. The physical link speed is T3, which is 45Mbps, but the SLA is 20 Mbps. Therefore, a hierarchical two-layer QoS policy is implemented. The parent policy shapes the link speed to 20Mbps and the child policy would queue and allocate the bandwidth within the 20Mbps.

### Step 1 Define the class-maps.

```

class-map match-all REALTIME
match ip dscp cs4 af41 cs5 ef
class-map match-all CRITICAL_DATA
match ip dscp af11 af21 cs3 af31 cs6
class-map match-all BEST_EFFORT
match ip dscp default
class-map match-all SCAVENGER
match ip dscp cs2

```

228949

### Step 2 Define the child policy map.

```

policy-map RSC_POLICY
class REALTIME
priority percent 33
class CRITICAL_DATA
bandwidth remaining percent 40
class SCAVENGER
bandwidth remaining percent 25
class BEST_EFFORT
bandwidth remaining percent 35

```

228950

### Step 3 Define the parent policy map.

```

policy-map RSC_PARENT_POLICY
class class-default
  shape average 20000000
  service-policy RSC_POLICY

```


228951

**Step 4** Apply the policy map to interface.

```

interface Serial2/0
  dampening
  ip address 10.126.0.4 255.255.255.254
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 eigrp-key
  ip pim sparse-mode
  service-policy output RSC_PARENT_POLICY
  ip summary-address eigrp 100 10.124.0.0 255.255.0.0 5
  load-interval 30
  carrier-delay msec 0
  dsu bandwidth 44210

```



## Redundancy

Redundancy must be factored into the WAN design for a number of reasons. Since the WAN may span across several service provider networks, it is likely that network will be subjected to different kinds of failures occurring all the time. Some of the following failures can occur over a period of time: route flaps, brownouts, fibers being cut, and device failures. The probability of these occurring over a short period of time is low, but the occurrence is highly likely over a long period of time. To meet these challenges, different kind of redundancy should be planned. The following are some of the ways to support redundancy:

- NSF/SSO—For networks to obtain 99.9999% of availability, technologies such as NSF/SSO are needed. The NSF would route packets until route convergence is complete, whereas SSO allows standby RP to take immediate control and maintain connectivity protocols.
- Service Software Upgrade (ISSU) allows software to be updated or modified, while packet forwarding continues with minimal interruption.
- Ether channel load balancing—Enabling this feature provides link resiliency and load balancing of traffic. This feature is enabled on the WAN aggregation 2 device. [Figure 3-17](#) shows where this feature is enabled.

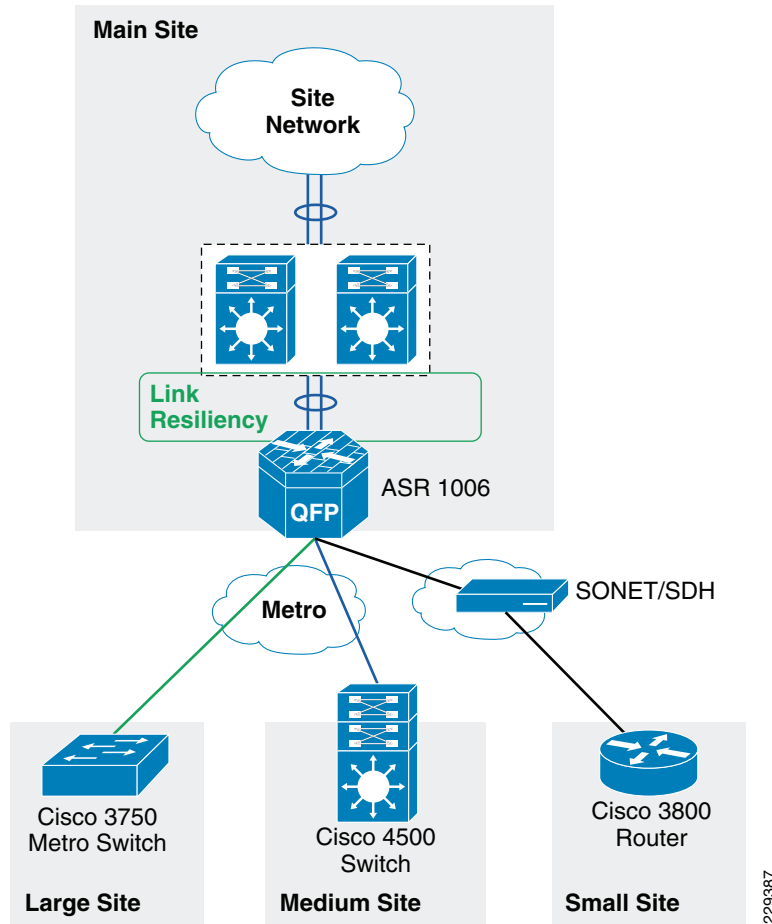
**Figure 3-17 Link Resiliency**

Table 3-3 shows the various WAN devices that are designed for resiliency.

**Table 3-3 WAN Devices**

Device	WAN Transport	Resiliency Feature
WAN aggregation 1	Private WAN/Internet	ISSU, IOS based redundancy
WAN aggregation 2	Metro	Redundant ESP, RP

This section discusses how to incorporate the resiliency principle in Cisco Medium Enterprise Design Profile for the WAN design. Enabling resiliency adds cost and complexity to the design. Therefore, resiliency has been added at certain places where it is absolutely critical to the network architecture rather than designing redundancy at every place of the network.

In the Cisco Medium Enterprise Design Profile, the redundancy is planned at both WAN aggregation router1 and WAN aggregation router 2 in the main site location. As explained in the [“WAN Aggregation Platform Selection in the Medium Enterprise Design Profile”](#) section on page 3-7, ASR routers have been selected at both WAN aggregation locations. However, there are different models at both WAN aggregation locations. When the ASR router interfaces with the private WAN, Internet networks the ASR 1004 with IOS-based redundancy. Similarly, for the ASR router that interfaces with Metro connections, the ASR 1006 with dual RP and dual ESP has been chosen to provide for hardware-based redundancy.

Both of these models support In Service Software Upgrade (ISSU) capabilities to allow a user to upgrade Cisco IOS XE Software while the system remains in service. To obtain more information on ASR resiliency capabilities, see the ASR page at following URL: <http://www.cisco.com/go/asr1000>

## Implementing IOS-based Redundancy at WAN Aggregation Router 1

The key requirement for implementing software-based redundancy on the ASR1004 is that it must have 4GB DRAM on ASR1004. The following are steps for implementing the IOS-based redundancy:

### Step 1 Check the memory on ASR 1004 router.

```
CR11-ASR-IE#show version
Cisco IOS Software, IOS-XE Software (PPC_LINUX_IOSD-ADVENTERPRISE-M), Version
12.2(33)XND3, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2010 by Cisco Systems, Inc.
Compiled Tue 02-Mar-10 09:51 by mcpre

Cisco IOS-XE software, Copyright (c) 2005-2010 by cisco Systems, Inc.
All rights reserved. Certain components of Cisco IOS-XE software are
licensed under the GNU General Public License ("GPL") Version 2.0. The
software code licensed under GPL Version 2.0 is free software that comes
with ABSOLUTELY NO WARRANTY. You can redistribute and/or modify such
GPL code under the terms of GPL Version 2.0. For more details, see the
documentation or "License Notice" file accompanying the IOS-XE software,
or the applicable URL provided on the flyer accompanying the IOS-XE
software.

ROM: IOS-XE ROMMON

CR11-ASR-IE uptime is 3 weeks, 6 days, 2 hours, 4 minutes
Uptime for this control processor is 3 weeks, 6 days, 2 hours, 6 minutes
System returned to ROM by SSO Switchover at 14:41:38 UTC Thu Mar 18 2010
System image file is "bootflash:asr1000rp1-adventerprise.02.04.03.122-33.XND3.bin"
Last reload reason: redundancy force-switchover

cisco ASR1004 (RP1) processor with 736840K/6147K bytes of memory.
5 Gigabit Ethernet interfaces
32768K bytes of non-volatile configuration memory.
4194304K bytes of physical memory.
937983K bytes of eUSB flash at bootflash:.
39004543K bytes of SATA hard disk at harddisk:.
15641929K bytes of USB flash at usb1:.

Configuration register is 0x2102

CR11-ASR-IE#
```

### Step 2 Enable the redundancy:

```
redundancy
 mode sso
!
```

### Step 3 Verify that redundancy is enabled:



```

CR11-ASR-IE#show redun
CR11-ASR-IE#show redundancy
Redundant System Information :
-----
    Available system uptime = 3 weeks, 6 days, 2 hours, 11 minutes
Switchovers system experienced = 3
    Standby failures = 0
    Last switchover reason = active unit removed

    Hardware Mode = Duplex
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
    Maintenance Mode = Disabled
    Communications = Up

Current Processor Information :
-----
    Active Location = slot 7
    Current Software state = ACTIVE
    Uptime in current state = 3 weeks, 6 days, 2 hours, 0 minutes
        Image Version = Cisco IOS Software, IOS-XE Software
(PPC_LINUX_IOSD-ADVENTERPRISE-M), Version 12.2(33)XND3, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2010 by Cisco Systems, Inc.
Compiled Tue 02-Mar-10 09:51 by mcpre
        BOOT =
bootflash:asr1000rp1-adventerprise.02.04.03.122-33.XND3.bin,1;
        CONFIG_FILE =
    Configuration register = 0x2102

Peer Processor Information :
-----
    Standby Location = slot 6
    Current Software state = STANDBY HOT
    Uptime in current state = 3 weeks, 6 days, 1 hour, 59 minutes
        Image Version = Cisco IOS Software, IOS-XE Software
(PPC_LINUX_IOSD-ADVENTERPRISE-M), Version 12.2(33)XND3, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2010 by Cisco Systems, Inc.
Compiled Tue 02-Mar-10 09:51 by mcpre
        BOOT =
bootflash:asr1000rp1-adventerprise.02.04.03.122-33.XND3.bin,1;
        CONFIG_FILE =
    Configuration register = 0x2102

CR11-ASR-IE#

```

## Implementation of Hardware-based Redundancy at WAN Aggregation Router 2

As explained in the design considerations documents, the WAN aggregation router 2 has redundant RPs and redundant ESPs. Therefore, with this configuration, we nonstop forwarding of data can be achieved even when there are failures with either ESP or RPs. The following steps are needed to enable hardware redundancy on WAN aggregation router 2:

### Step 1 Configuration of SSO redundancy:

```

redundancy
mode sso

```

**Step 2** Verify the redundancy information:

```

cr11-asr-we#show redundancy
Redundant System Information :
-----
    Available system uptime = 3 weeks, 6 days, 3 hours, 32 minutes
Switchovers system experienced = 4
    Standby failures = 0
    Last switchover reason = active unit removed

    Hardware Mode = Duplex
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
    Maintenance Mode = Disabled
    Communications = Up

Current Processor Information :
-----
    Active Location = slot 6
    Current Software state = ACTIVE
    Uptime in current state = 2 weeks, 1 day, 19 hours, 3 minutes
    Image Version = Cisco IOS Software, IOS-XE Software
(PPC_LINUX_IOSD-ADVENTERPRISEK9-M), Version 12.2(33)XND2, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled Wed 04-Nov-09 18:53 by mcpre
    BOOT =
    CONFIG_FILE =
    Configuration register = 0x2102

Peer Processor Information :
-----
    Standby Location = slot 7
    Current Software state = STANDBY HOT
    Uptime in current state = 2 weeks, 1 day, 18 hours, 52 minutes
    Image Version = Cisco IOS Software, IOS-XE Software
(PPC_LINUX_IOSD-ADVENTERPRISEK9-M), Version 12.2(33)XND2, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled Wed 04-Nov-09 18:53 by mcpre
    BOOT =
    CONFIG_FILE =
    Configuration register = 0x2102

cr11-asr-we#

```

**Implementation of Link Resiliency Between the WAN Aggregation Router 2 and VSS Core**

The following are the implementation steps to deploy link resiliency:

**Step 1** Configure the EtherChannel between the ASR1006 and the VSS core:

```

interface GigabitEthernet0/2/3
  dampening
  no ip address
  load-interval 30
  carrier-delay msec 0
  negotiation auto
  cdp enable
  service-policy output WAN_Upstream

```

```

channel-group 1 mode active
hold-queue 2000 in
hold-queue 2000 out
!
interface GigabitEthernet0/2/4
dampening
no ip address
load-interval 30
carrier-delay msec 0
negotiation auto
cdp enable
service-policy output WAN_Upstream
channel-group 1 mode active
hold-queue 2000 in
hold-queue 2000 out
!
Step 2) Configure the port-channel interface
interface Port-channel1
ip address 10.125.0.23 255.255.255.254
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
ip pim sparse-mode
ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5
logging event link-status
load-interval 30
carrier-delay msec 0
negotiation auto
!

```

## Multicast

The main design considerations for multicast are as follows:

- The number of groups supported by the WAN edge device. This is scalability factor of the WAN edge device. The platform chosen must support the number of required groups.
- The placement of the RP—There are couple of options available with RP placement, which include Anycast with Static, Anycast with Auto-RP, or Anycast with BSR
- Multicast protocols—PIM-Sparse mode, IGMP
- QoS policy must be configured for multicast traffic, so that this traffic does not affect the unicast traffic

In the Medium Enterprise Design Profile, it is assumed that multicast traffic would be present only within the site, and not external enterprise/WAN networks. Therefore, the multicast design looks at only between the main site and small remote site locations. The implementation section in the document shows how to enable multicast on the WAN device only. Therefore, to obtain more information about multicast design for site, refer to the [“Multicast for Application Delivery” section on page 2-64](#).

### Multicast Configuration on WAN Aggregation Router 2

This section shows how to enable multicast routing, and what interfaces to be enabled with PIM-Sparse mode on the WAN aggregation router 2 that connects to different remote sites.

- 
- Step 1** Enable multicast routing:
- ```
ip multicast-routing distributed
```

**Step 2** Enable PIM-Sparse mode on the following WAN interfaces:

- Port-channel—Connects to the VSS core
- Gi0/2/0—Connects to Large Remote Site site
- Gi0/2/1—Connects to Medium Remote Site site
- S0/3/0—Connects to Small Remote Site site

```

interface Port-channel1
 ip address 10.125.0.23 255.255.255.254
 ip pim sparse-mode
 negotiation auto
!
interface GigabitEthernet0/2/0
 description Connected to cr11-3750ME-RLC
 ip address 10.126.0.1 255.255.255.254
 ip pim sparse-mode
 logging event link-status
 load-interval 30
 negotiation auto
!
interface GigabitEthernet0/2/1
 description Connected to cr11-4507-RMC
 dampening
 no ip address
 load-interval 30
 carrier-delay msec 0
 negotiation auto
 cdp enable
 hold-queue 2000 in
 hold-queue 2000 out
!
interface GigabitEthernet0/2/1.102
 encapsulation dot1Q 102
 ip address 10.126.0.3 255.255.255.254
 ip pim sparse-mode
!
!
interface Serial0/3/0
 dampening
 ip address 10.126.0.5 255.255.255.254
 ip pim sparse-mode
 load-interval 30
 carrier-delay msec 0
 dsu bandwidth 44210
 framing c-bit
 cablelength 10
!
Step 3) Configure the RP location
 ip pim rp-address 10.100.100.100

```

**Configuration of Multicast on Large Remote Site**

This section discusses how to implement multicast on large remote site. The following are implementation steps:

**Step 1** Enable multicast routing:

```
ip multicast-routing distributed
```

- Step 2** Enable PIM-Sparse mode on the WAN interface that connects to main site.

```
interface GigabitEthernet1/1/1
description Connected to cr11-ASR-WE
no switchport
dampening
ip address 10.126.0.0 255.255.255.254
ip pim sparse-mode
hold-queue 2000 in
hold-queue 2000 out
!
```

### Configuration of Multicast on Medium Remote Site

This section discusses on how to implement multicast on medium remote site.

- Step 1** Enable multicast routing:

```
ip multicast-routing
```

- Step 2** Enable PIM-Spare mode on the WAN interface:

```
interface Vlan102
description Connected to cr11-ASR-WE
dampening
ip address 10.126.0.2 255.255.255.254
ip pim sparse-mode
load-interval 30
carrier-delay msec 0
```

### Configuration of Multicast on Small Remote Site

This section discusses on how to implement multicast on small remote site.

- Step 1** Enable multicast routing:

```
ip multicast-routing
```

- Step 2** Enable PIM -pare mode on the WAN interface:

```
interface Serial2/0
dampening
ip address 10.126.0.4 255.255.255.254
ip pim sparse-mode
load-interval 30
carrier-delay msec 0
dsu bandwidth 44210
```

- Step 3** Configure the RP location:

```
ip pim rp-address 10.100.100.100 Allowed_MCAST_Groups override
```

- Step 4** Configure the multicast security:

```
ip pim spt-threshold infinity
ip pim accept-register list PERMIT-SOURCES
!
ip access-list standard Allowed_MCAST_Groups
```

```
permit 224.0.1.39
permit 224.0.1.40
permit 239.192.0.0 0.0.255.255
deny any
ip access-list standard Deny_PIM_DM_Fallback
deny 224.0.1.39
deny 224.0.1.40
permit any
!
ip access-list extended PERMIT-SOURCES
permit ip 10.125.31.0 0.0.0.255 239.192.0.0 0.0.255.255
deny ip any any
!
```

## Summary

Designing the WAN network aspects for the Cisco Medium Enterprise Design Profile interconnects the various LAN locations as well as lays the foundation to provide safety and security, operational efficiencies, virtual learning environments, and secure classrooms.

This chapter reviewed the WAN design models recommended by Cisco and where to apply these models within the various locations within a medium enterprise network. Key WAN design principles such as WAN aggregation platform selection, QoS, multicast, and redundancy best practices are discussed for the entire Medium Enterprise Design Profile. Designing the WAN network of a medium enterprise using these recommendations and best practices will establish a network that is resilient in case of failure, scalable for future growth, simplified to deploy and manage, and cost efficient to meet the budget needs of a medium enterprise.



## CHAPTER 4

# Medium Enterprise Design Profile (MEDP)— Mobility Design

---

## Mobility Design

The Cisco Medium Enterprise Design Profile is intended to assist enterprises in the design and deployment of advanced network-based solutions within twenty-first century business environments.

At the heart of the Medium Enterprise Design Profile is the network service fabric, which is a collection of products, features, and technologies that provide a robust routing and switching foundation upon which all solutions and services are built. Operating on top of the network service fabric are all the services used within the medium enterprise network to solve business problems.

Today's enterprise worker is dynamic, mobile, and technology-savvy. When at the enterprise site, they move about while equipped with an array of mobility-enabled devices including PDAs, phones, and laptops. Business professionals tend to use state of the art applications and the enterprise network for many aspects of their lives, demanding connectivity, performance and network flexibility wherever they may be located. This connected generation of professionals is untethered from wired networks and typically assume that high-performance, reliable wireless LANs (WLANs) are present at all medium enterprise environments.

The mobility design implemented by a medium enterprise must meet the needs of these mobile workers while also addressing the requirements of guests and visitors. The challenge facing a medium enterprise is to create a robust, end-to-end, mobility-enabled network that supports their requirements at a cost that makes good business sense. Medium enterprises should be equipped with mobility solutions that support the following:

- Secure communications between local and remote sites to support employees, guests and visitors, using mobility-enabled devices and mobile applications
- A scalable design model that can easily accommodate the addition of new local and remote buildings as well as modifications to existing buildings
- Support for bandwidth-intensive, high-speed multimedia applications
- Simplified management tools to facilitate system-wide mobility maintenance
- The use of tools and applications for mobile conferencing, collaboration, and operations
- Effective communication and inter operation with public safety first responders in the event of an emergency.

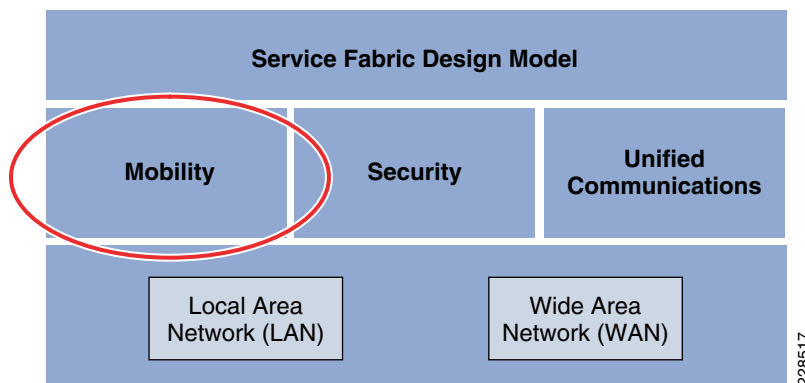
Medium enterprises must remain competitive and must differentiate themselves from their peers, both for competitive customer marketing purposes as well as to attract and retain the best employee talent. Prospective employees want to be part of medium enterprises that provide services relevant to the way

they live, work, and spend their free time. They want to take full advantage of what the medium enterprise has to offer, in ways that serve to enhance both their quality of life and their individual success potential. A medium enterprise with a pervasive, high-speed wireless network not only provides technological leadership and innovation, but enables the deployment of innovative applications that streamline operations, enhance collaboration, and improve productivity.

This mobile enterprise lifestyle helps to drive the need for careful wireless capacity and coverage planning. Keep in mind that traditional offices and conference rooms are by no means the only environments seen within medium enterprises any longer. In fact, high performance, secure wireless technologies can enable “virtual offices” even in non-traditional settings such as leased space in professional buildings, temporary office spaces, and even in employee homes. Administrators need secure access to tools, records, and resources, as well as access to mobile voice capabilities throughout medium enterprise sites. Secure, reliable, and high-performance wireless guest access for contractors, vendors, and other guests of the medium enterprise has become a standard and expected part of modern-day mobile business environments.

To meet these needs, medium enterprises must evolve into mobility-enabled local and remote sites and twenty-first century business centers. In support of this, this chapter discusses design considerations surrounding the requirements, expectations and trade-offs that must be taken into account when integrating mobility into the Cisco Medium Enterprise Design Profile. These design considerations form a critical part of the overall service fabric design model, as shown in [Figure 4-1](#).

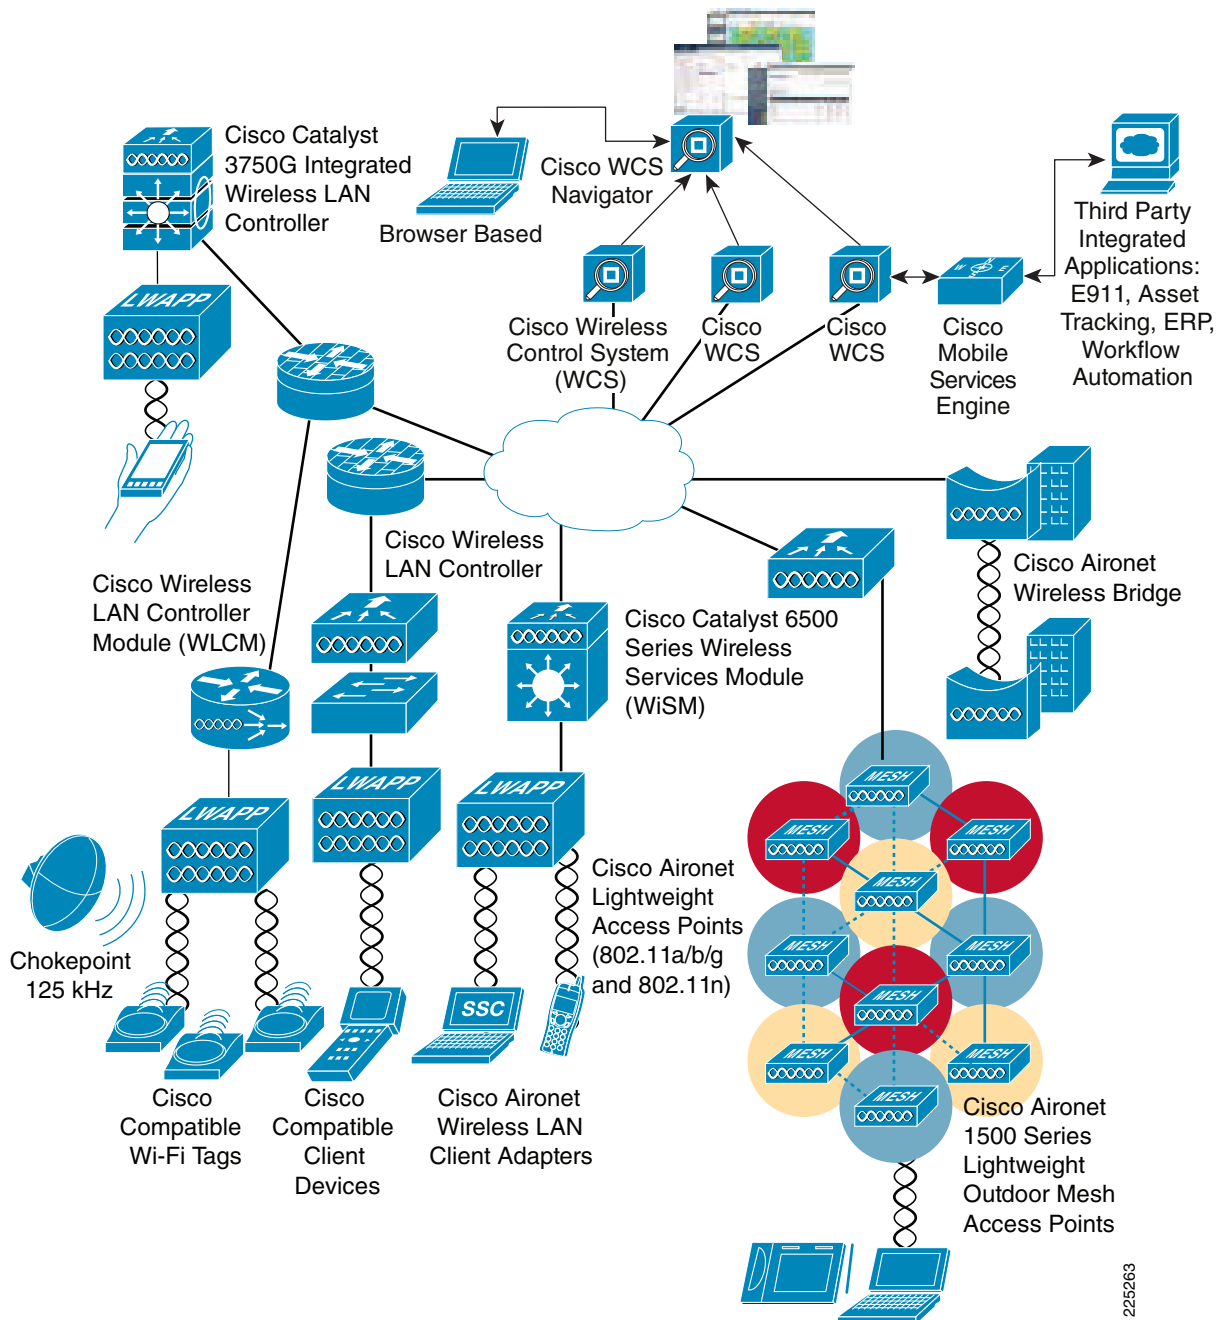
**Figure 4-1 Service Fabric Design Model**



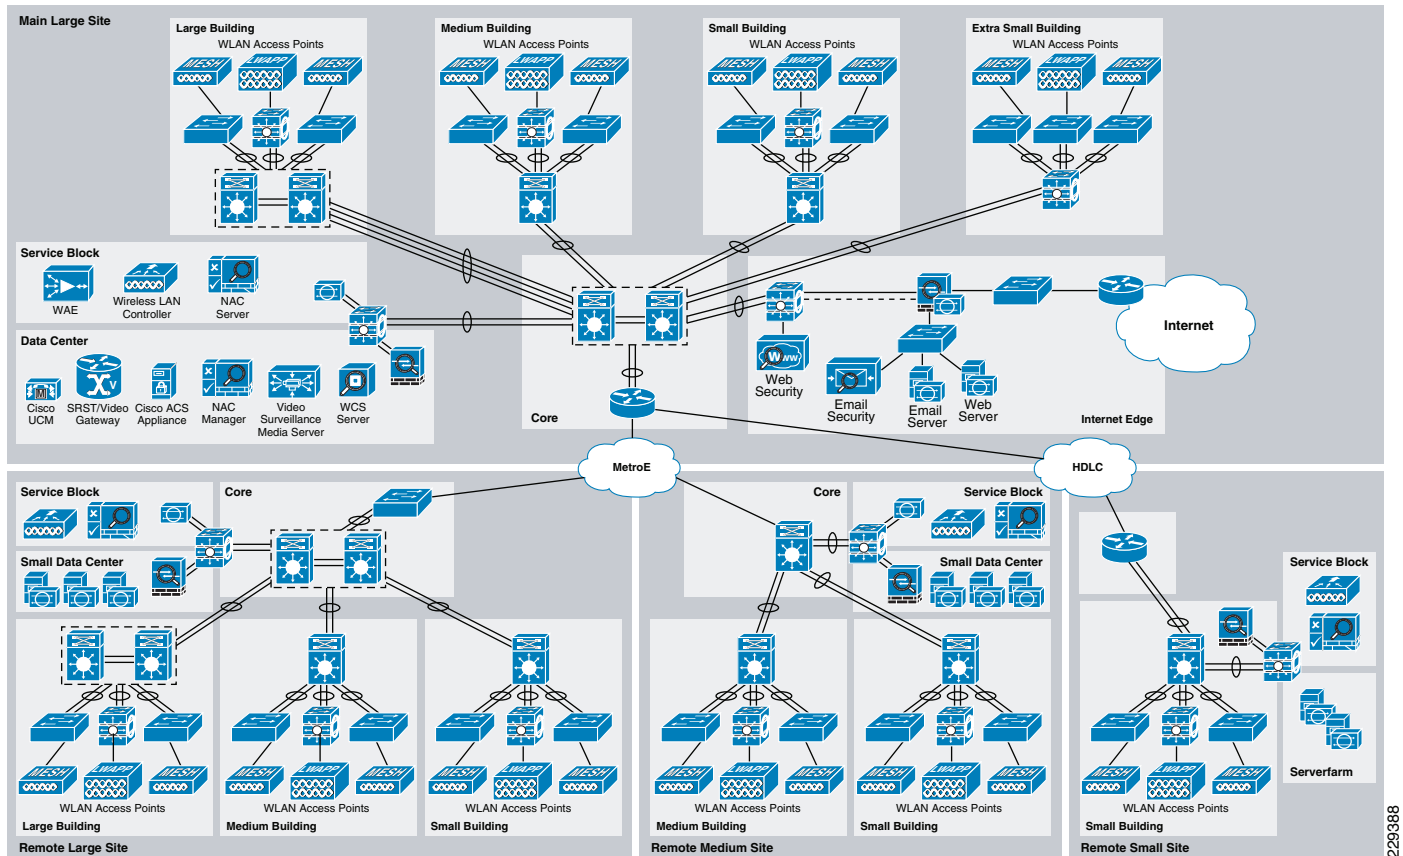
Given the mobility requirements of medium enterprise professionals, guests, and visitors, wireless LANs have emerged as one of the most effective and high performance means for these mobile users to access the medium enterprise network. The Cisco Unified Wireless Network (Cisco UWN) is a unified solution 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 4-2](#) shows a high-level topology of the Cisco Unified Network, which includes access points that use the Control and Provisioning of Lightweight Access Points (CAPWAP) protocol, the Cisco Wireless Control System (WCS), and the Cisco Wireless LAN Controller (WLC). In addition to the traditional standalone WLAN controller, alternate hardware platforms include the Cisco ISR router Wireless LAN Controller Module (WLCM) or the Cisco Catalyst 6500 Wireless Services Module (WiSM). The Cisco Access Control Server (ACS) and its Authentication, Authorization, and Accounting (AAA) features complete the solution by providing Remote Authentication Dial-In User Service (RADIUS) services in support of user authentication and authorization.



**Figure 4-2 Cisco Unified Wireless Network Overview**

The Cisco Medium Enterprise Design Profile accommodates a main site and one or more remote sites interconnected over a metro Ethernet or managed WAN service. Each of these sites may contain one or more buildings of various sizes, as shown in [Figure 4-3](#).

**Figure 4-3** Medium Enterprise Design Profile Overview

Operating on top of this network are all the services used within the medium enterprise environment such as safety and security systems, voice communications, video surveillance equipment, and so on. The core of these services are deployed and managed at the main (or headquarters) site building, allowing each remote site to reduce the need for separate services to be operated and maintained. These centralized systems and applications are served by a data center at the main site.

As Figure 4-3 shows, the Cisco Medium Enterprise Design Profile uses a centralized approach in which key resources are centrally deployed. The key feature of this integration is the use of one or more WLAN controllers at each site, with the overall WLAN management function (the Cisco WCS) located at the main site. This approach simplifies the deployment and operation of the network, helping to ensure smooth performance, enhance security, enhance network maintainability, maximize network availability, and reduce overall operating costs.

The Cisco Medium Enterprise Design Profile takes into account that cost and limited network administrative resources can, in some cases, be limiting factors for medium enterprises. The topologies and platforms are carefully selected to increase productivity while minimizing the overall cost and complexity of operation. In certain instances, trade-offs are necessary to reach these goals, and this document helps to point out and clarify some of these trade-offs.

The Cisco mobility approach within the Cisco Medium Enterprise Design Profile focuses on the following key areas:

- *Accessibility*

- Enabling mobile professionals, administrators, guests and visitors to be accessible and productive on the network, regardless of whether they are in a traditional office setting, collaborating in a conference room, having lunch with colleagues within enterprise site dining areas, or simply enjoying a breath of fresh air outside on-site buildings
- Enabling easy, secure guest access to guests such as prospective customers, future employees, 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 quality-of-service (QoS) classification. This gives preferential treatment to real-time traffic, helping to ensure that video and audio information arrives on time.
- *Security*
  - Segmenting authorized users and blocking unauthorized users
  - Extending the services of the network safely to authorized parties
  - Enforcing security policy compliance on all devices seeking to access network computing resources. Staff 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.

**Note**

For information on how security design is addressed within the Cisco Medium Enterprise Design Profile, see [Chapter 5, “Medium Enterprise Design Profile \(MEDP\)—Network Security Design.”](#)

- *Manageability*

A relatively small team of network administrators should be able to easily deploy, operate, and manage hundreds of access points that may reside within a multisite medium enterprise. A single, easy-to-understand WLAN management framework provides small, medium, and large sites with the level of WLAN management scalability, reliability, and ease of deployment required in the medium enterprise domain.
- *Reliability*
  - Providing adequate capability to recover from a single-layer fault of a WLAN access component or controller wired link.
  - Ensuring that WLAN accessibility is maintained for employees, administrators, staff, guests, and visitors, in the event of common failures.

## Accessibility

This section provides a brief introduction to the fundamental protocol used for communication between access points and WLAN controllers, followed by a discussion of mobility design considerations pertaining to those aspects of the Cisco Medium Enterprise Design Profile relevant to accessibility, such as the following:


- WLAN controller location
- WLAN controller connectivity

- Access points

The basic mobility components involved with providing WLAN access in the Cisco Medium Enterprise Design Profile consists of WLAN controllers and access points that communicate with each other using the IETF standard CAPWAP protocol. In this arrangement, access points provide the radio connection to wireless clients, and WLAN controllers manage the access points and provide connectivity to the wired network.

Figure 4-4 shows the use of CAPWAP by access points to communicate with and tunnel traffic to a WLAN controller.

**Figure 4-4 CAPWAP Access Point to WLC Communication**



CAPWAP enables the controller to manage a collection of wireless access points, and has the following three primary functions in the mobility design:

- Control and management of the access point
- Tunneling of WLAN client traffic to the WLAN controller
- Collection of 802.11 data for overall WLAN system management

CAPWAP is also intended to provide WLAN controllers with a standardized mechanism with which to manage radio-frequency ID (RFID) readers and similar devices, as well as enable controllers to interoperate with third-party access points in the future.

In controller software Release 5.2 or later, Cisco lightweight access points use CAPWAP to communicate between the controller and other lightweight access points on the network. Controller software releases before Release 5.2 use the Lightweight Access Point Protocol (LWAPP) for these communications. Note that most CAPWAP-enabled access points are also compatible with the preceding LWAPP protocol. An exception is that the Cisco Aironet 1140 Series Access Point supports only CAPWAP.

The mobility approach in the Cisco Medium Enterprise Design Profile is based on the feature set available in Cisco Wireless LAN Controller software Release 6.0, which uses CAPWAP.

For detailed CAPWAP protocol information, see the following URL:  
<http://www.ietf.org/rfc/rfc5415.txt>.

## WLAN Controller Location

WLAN deployments are typically categorized into two main categories, *distributed* and *centralized*:

- *Distributed controller*—In this model, WLAN controllers are located throughout the medium enterprise network, typically on a per-building basis, and are responsible for managing the access points resident in a given building. This technique is commonly used to connect controllers to the medium enterprise network using distribution routers located within each building. In the distributed deployment model, the CAPWAP tunnels formed between access points and WLAN controllers are typically fully contained within the confines of the building.
- *Centralized controller*—In this model, WLAN controllers are placed at a centralized location within the enterprise. Because centralized WLAN controllers are typically not located in the same building as the access points they manage, the CAPWAP tunnels formed between them must traverse the site backbone network.

The Cisco Medium Enterprise Design Profile is based on the centralization of WLAN controllers, on a per-site basis, and follows established best practices, such as those contained in Chapter 2 of the *Enterprise Mobility 4.1 Design Guide* at the following URL:

[http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns820/landing\\_ent\\_mob\\_design.html](http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns820/landing_ent_mob_design.html).

Figure 4-3 shows the planned deployment of WLAN controllers within distinct per-site service blocks, each associated with the main (headquarters), large remote, medium remote, and small remote sites respectively. Service blocks tend to be deployed at locations in the network where high availability routing, switching, and power is present. In addition, these areas tend to be locally or remotely managed by network staff possessing higher skill sets.

Some of the advantages underlying the decision to centralize the deployment of WLAN controllers on a per-site basis include the following:

- *Reduced acquisition and maintenance costs*—By servicing the needs of all wireless users from a central point, the number of WLAN controller hardware platforms deployed can be reduced compared to that required for a distributed, per-building design. Similarly, incremental software licensing costs associated with WLAN controllers are reduced as well. These economies of scale typically increase with the size of the enterprise WLAN.
- *Reduced administrative requirements*—By minimizing the total number of WLAN controllers deployed, the controller management burden imposed on site network administrators is minimized.
- *Cost-effective capacity management*—The use of a centralized WLAN controller model allows the designer the ability to centrally service access points located in multiple building locations and efficiently manage controller capacity.
- *Simplified network management and high availability*—Centralized WLAN controller designs simplify overall network management of controllers, as well as facilitate cost-effective controller high availability approaches. This can protect sites from a loss of WLAN access in the rare event of a controller failure, without the expense of 1:1 controller duplication.
- *Reduced component interaction points*—Centralizing WLAN controllers minimizes the number of integration points that must be managed when interfacing the controller with other devices. When integrating the WLAN controller with the Network Admission Control (NAC) appliance on any given site, for example, only one integration point must be administered.
- *Increased performance and reliability*—Centralized WLAN controller deployments usually lead to highly efficient inter-controller mobility. For large sites, there is also an incremental economy of scale that occurs as the network grows larger. By centralizing WLAN controllers on a per-site basis, CAPWAP tunneling between access points and WLAN controllers is not normally required to traverse WAN links (except during controller fail over), thereby conserving WAN bandwidth and improving performance overall.

**Note**

For additional information on inter-controller mobility and roaming, see the following URL: [http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2\\_Arch.html#wp1028197](http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2_Arch.html#wp1028197).

The choice of WLAN controller for the Cisco Medium Enterprise Design Profile is the Cisco 5508 Wireless Controller, as shown in [Figure 4-5](#).

**Figure 4-5** Cisco 5508 Wireless Controller



The Cisco 5508 Wireless Controller is a highly scalable and flexible platform that enables system-wide services for mission-critical wireless in medium to large-sized enterprise environments. Designed for 802.11n performance and maximum scalability, the Cisco 5508 Wireless Controller offers the ability to simultaneously manage from 12 to a maximum of 250 access points per controller. Base access point controller licensing provides the flexibility to purchase only the number of access point licenses required, with the ability to add additional access point licenses in the future when medium enterprise site growth occurs. In sites requiring more than 250 total access points, or load sharing/high availability is required, multiple controllers can be deployed as necessary.

More information on the Cisco 5508 Wireless Controller can be found at the following URL: [http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/data\\_sheet\\_c78-521631.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/data_sheet_c78-521631.html).

## WLAN Controller Connectivity

This section discusses WLAN controller connectivity, including the following:

- Controller connectivity to the wired network
- Controller connectivity to the wireless devices
- Defining WLANs and Service Set Identifiers (SSIDs)
- WLAN controller mobility groups
- WLAN controller access point groups
- WLAN controller RF groups

### Controller Connectivity to the Wired Network

WLAN controllers possess physical entities known as *ports* that connect the controller to its neighboring switch (the Cisco 5508 Wireless Controller supports up to eight Gigabit Ethernet Small Form-Factor Pluggable [SFP] ports). Each physical port on the controller supports, by default, an 802.1Q VLAN trunk, with fixed trunking characteristics.

**Note**

For more information concerning the various types of ports present on Cisco WLAN controllers, see the *Cisco Wireless LAN Controller Configuration Guide, Release 6.0* at the following URL:  
<http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.html>.

*Interfaces* are logical entities found on the controller. An interface may have multiple parameters associated with it, including an IP address, default gateway, primary physical port, optional secondary physical port, VLAN identifier, and Dynamic Host Configuration Protocol (DHCP) server. Each interface is mapped to at least one primary port, and multiple interfaces can be mapped to a single controller port.

**Note**

For more information concerning the various types of interfaces present on Cisco WLAN controllers, see the *Cisco Wireless LAN Controller Configuration Guide, Release 6.0* at the following URL:  
<http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.html>.

A special type of controller interface is known as the *AP manager interface*. A controller has one or more AP manager interfaces, which are used for all Layer 3 communications between the controller and its joined access points. The IP address of the AP manager interface is used as the tunnel source for CAPWAP packets from the controller to the access point, and as the destination for CAPWAP packets from the access point to the controller. The AP manager interface communicates through a distribution system port by listening across the Layer 3 network for CAPWAP “join” messages generated by access points seeking to communicate with and “join” the controller.


*Link aggregation (LAG)* is a partial implementation of the 802.3ad port aggregation standard. It bundles all of the controller distribution system ports into a single 802.3ad port channel, thereby reducing the number of IP addresses needed to configure the ports on your controller. When LAG is enabled, the system dynamically manages port redundancy and load balances traffic transparently to the user. LAG bundles all the enabled distribution ports on the WLAN controller into a single EtherChannel interface.

Currently published best practices specify either multiple AP manager interfaces (with individual Ethernet links to one or more switches) or link aggregation (with all links destined for the same switch or switch stack) as the recommended methods of interconnecting WLAN controllers with wired network infrastructure. For more information, see the following URL:

<http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/c60mint.html#wp1277659>.

In the Cisco Medium Enterprise Design Profile, the Cisco 5508 Wireless Controllers are interconnected with the modular switches or switch stacks found in the services block using link aggregation and EtherChannel exclusively, as shown in [Figure 4-6](#).

**Figure 4-6** WLAN Controller Link Aggregation to Services Block



In this way, one or more centralized WLAN controllers are connected via the services block to the site core. This design can make use of up to eight Gigabit Ethernet connections from the Cisco 5508 Wireless Controller to the services block. These Gigabit Ethernet connections should be distributed among different modular line cards or switch stack members as much as possible, so as to ensure that the failure of a single line card or switch stack failure does not result in total failure of the



WLAN controller connection to the site network. The switch features required to implement this connectivity between the WLAN controller and the services block are the same switch features that would otherwise be used for EtherChannel connectivity between switches in general.

Further discussion of the advantages of using controller link aggregation, as well as the considerations concerning its implementation in the Cisco Medium Enterprise Design Profile can be found in [Controller Link Aggregation, page 4-35](#).

The key advantage of using link aggregation in this fashion instead of multiple AP manager interfaces is design performance, reliability, and simplicity:

- With the Ethernet bundle comprising up to eight Gigabit Ethernet links, link aggregation provides very high traffic bandwidth between the controller and the site network.
- With link aggregation, if any of the controller ports fail, traffic is automatically migrated to one of the other controller ports. As long as at least one controller port is functioning, the system continues to operate, access points remain connected to the network, and wireless clients continue to send and receive data. Terminating on different modules within a single Catalyst modular switch, or different switch stack members (as shown in [Figure 4-6](#)), provides redundancy and ensures that connectivity between the services block switch and the controller is maintained in the rare event of a failure.
- Link aggregation also offers simplicity in controller configuration; for example, configuring primary and secondary ports for each interface is not required.

## Controller Connectivity to Wireless Devices

This section deals with the design considerations that involve provisioning wireless access for the various user groups that reside within the medium enterprise, such as the administrators, employees, and guests. These considerations include the WLAN controllers deployed in the services blocks, as well as the access points that are located in buildings.

### Defining WLANs and SSIDs

In most medium enterprises, various user groups will likely require access to the WLAN for a variety of different purposes. Although usage peaks may occur, it is safe to assume that a large portion of these groups will likely want access to the WLAN at more or less the same time. Thus, in designing for mobility in the Cisco Medium Enterprise Design Profile, the wireless infrastructure must support logical segmentation in such a fashion that a reasonable proportion of all users can be serviced simultaneously and with an appropriate degree of security and performance.

One of the basic building blocks used in the WLAN controller to address this need is the ability to provision logical WLANs, each of which are mapped to different wired network interfaces by the WLAN controller. These WLANs are configured and assigned a unique SSID, which is a sequence of characters that uniquely names a WLAN. For this reason, an SSID is also sometimes referred to simply as a *network name*.



#### Note

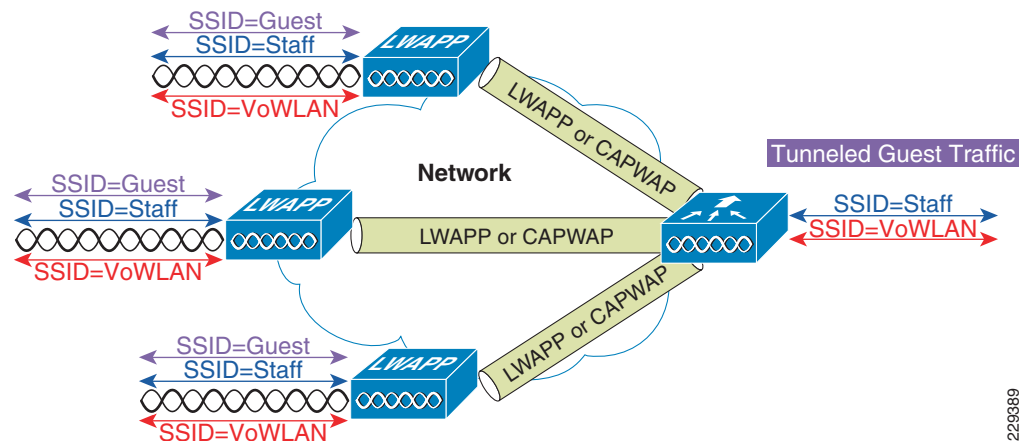
Each set of wireless devices communicating directly with each other is called a basic service set (BSS). Several BSSs can be joined together to form one logical WLAN segment, referred to as an extended service set (ESS). An SSID is simply the 1–32 byte alphanumeric name given to each ESS.

To promote ease of administration, the value chosen for the SSID should bear some direct relationship to the intended purpose of the WLAN.



Figure 4-7 provides a high-level illustration of the three logical WLANs that provide mobility within the Cisco Medium Enterprise Design Profile, and how they are mapped to WLAN controller network interfaces or tunneled to another controller. For ease of administration and the support of employees, administrators, guests and visitors that frequent multiple sites, the names chosen for the WLAN SSIDs should be consistent within each site in the medium enterprise system. For example, in Figure 5-7 employee wireless access is made available anywhere there is WLAN RF coverage using the SSID titled “staff”.

**Figure 4-7 WLAN SSIDs**



In the Medium Enterprise Design Profile, the set of WLAN SSIDs provide access to the following WLANs:

- A *secured staff* WLAN network with dynamically generated per-user, per-session encryption keys.

This WLAN would be used by enterprise employees, administrators and other staff members using managed client devices, such as laptops, PDAs, and so on. The secured staff WLAN is designed to provide secure access and good performance for devices supported by the medium enterprise network administration staff. Devices that are used on the secured staff WLAN are usually procured and deployed by (or with the knowledge and cooperation of) the medium enterprise network administration staff on behalf of full-time and temporary employees. These employees are typically prohibited from bringing their own personal PDAs, laptops, or voice over WLAN (VoWLAN) phones to use on the secured staff WLAN. This allows, for example, a uniform baseline level of authentication and encryption to be deployed for the secured staff WLAN across all such devices. An underlying assumption made here is that only devices supporting compatible authentication and encryption would be considered for deployment at all.

The characteristics of this WLAN include the following:

- Wi-Fi Protected Access 2 (WPA2) encryption with 802.1x/EAP authentication, and Cisco Centralized Key Management (Cisco CKM, also referred to as CCKM) for enhanced roaming.

Most modern WLAN client devices being produced today support this level of authentication and encryption. The addition of Cisco CKM in this case provides for faster roaming by enabling Cisco CKM-equipped clients to securely roam from one access point to another without the need to re-authenticate after the roam completes.

- Broadcast SSID enabled. Enabling this helps to avoid potential connectivity difficulties with some clients. There is no real disadvantage to enabling broadcast SSID.
- QoS profile setting of *silver* (best effort delivery).



**Note**

For more details on WLAN QoS, see the references contained at the end of [Quality-of-Service, page 4-26](#).

- Wi-Fi Multimedia (WMM) policy of allowed. This allows devices and applications that can support 802.1e enhanced QoS prioritization to do so. Enabling the use of WMM in this way is also in compliance with the 802.11n.
- Mandatory IP address assignment via DHCP. Eliminating the configuration of static IP addresses helps to mitigate the risk of IP address duplication.
- Radio policy set to allow clients to use either 2.4 GHz or 5 GHz to access this WLAN. This allows clients that can take advantage of benefits of 5 GHz operation (such as increased capacity and reduced interference) to do so.



**Note**

The 802.11b and 802.11g physical layers (PHYs) are applied in the unlicensed 2.4 GHz industrial, scientific, and medical (ISM) frequency band, whereas the 802.11a PHY is applied in the unlicensed 5 GHz ISM band. “Dual-band” 802.11a/bg clients are capable of operating in either 2.4 or 5 GHz frequency bands because they are capable of using any of the three PHYs. Selection between PHYs is typically achieved via software configuration.

Clients using the very high speed 802.11n PHY may be designed to operate in a single band, or they may be 802.11n “dual-band” clients. Unlike the 802.11b, 802.11g, and 802.11a PHYs, simply stating that a client is 802.11n does not precisely indicate what frequency bands the client is capable of operating within.

For more information about the 802.11n PHY and its application to the 2.4 and 5 GHz frequency bands, see the following URL:

[http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/white\\_paper\\_80211n\\_design\\_and\\_deployment\\_guidelines.html](http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/white_paper_80211n_design_and_deployment_guidelines.html).

- A *secured VoWLAN* network that is optimized for VoWLAN usage by employee staff and administrators using managed VoWLAN client devices.

As was the case with the secured staff WLAN, this WLAN is designed to provide secure access and good performance when used with VoWLAN devices (such as the Cisco Unified Wireless IP Phone 7925G) that are usually procured, deployed, and managed by (or with the knowledge and cooperation of) the medium enterprise network administration staff. Such procurement is usually conducted on behalf of full-time and temporary employee staff users. To assure proper security and promote effective device management, employee staff users are typically prohibited from bringing their own personal VoWLAN phones and using them on this WLAN. This allows, for example, a baseline level of authentication and encryption to be deployed for this WLAN with the knowledge that the devices using this WLAN will support that level of security. The key differences between this WLAN and the secured staff WLAN include the following:

- The security policy on this WLAN is WPA with Cisco CKM, which is recommended as a best practice for the Cisco 7921G and 7925G VoWLAN phones.
- WLAN controller QoS profile setting of *platinum*, which assigns the highest prioritization to voice traffic.
- WMM policy is *required* (this precludes the use of clients that do not support WMM).

- Load-based Call Admission Control (CAC) should be specified for this WLAN. This prevents VoWLAN calls from being added to an access point that is unable to accept them without compromising call quality.
- The radio policy should be set to allow clients to access only this WLAN using 5 GHz. This helps to ensure that all secured voice devices take full advantage of the robust call capacity and reduced co-channel interference characteristics associated with 5 GHz.

For further information on best practices for voice applications, see the *Voice over Wireless LAN 4.1 Design Guide* at the following URL:

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

- A *guest access* WLAN that uses web authentication for guest users of the enterprise network.

Traffic to and from this guest access WLAN is tunneled to the DMZ transparently, with no visibility by, or interaction with, other traffic in the enterprise. The Cisco Medium Enterprise Design Profile uses the Cisco Unified Wireless Network to provide a flexible, easy-to-implement method for deploying wireless guest access by using Ethernet in IP (RFC3378). Ethernet in IP is used to create a tunnel across a Layer 3 topology between two WLAN controller endpoints (known as the *foreign* and *anchor* controllers). The foreign controller is the controller resident in the respective site services block described earlier, whereas the anchor controller is resident within the network DMZ. The benefit of this approach is that no additional protocols or segmentation techniques must be implemented to isolate guest traffic travelling within the tunnel from all other enterprise traffic.

See [Guest Access, page 4-27](#) for further information regarding considerations surrounding the products and techniques used to provide guest access when designing for mobility in the Cisco Medium Enterprise Design Profile.

For technical information on Guest Access best practices in wireless networks, see the Guest Access section in the *Enterprise Mobility 4.1 Design Guide* at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch10GuAc.html>.

The guest access WLAN must be designed to accommodate the needs of enterprise guests (such as customers, vendors, contractors, prospective employee candidates, and so on) as well as the wide variety of WLAN guest client devices they may bring into the enterprise. Although their numbers will likely be much less compared to that of employees, the WLAN clients brought into the enterprise environment by guest users are typically not managed or directly supported by medium enterprise network administrative staff. Because of the lack of control over the type of device used, mandating the use of 802.1x authentication and WPA or WPA2 encryption does not usually facilitate a practical guest access solution.

Characteristics of the guest access WLAN include the following:

- To provide access control and an audit trail, the guest access WLAN authenticates the user via a web portal (web authentication) where all network access, apart from DHCP and Domain Name Service (DNS), is blocked until the user enters a correct user name and password into an authentication web page.
- The guest access WLAN user is re-directed to a web authentication web page whenever the user attempts to open any web page before successful authentication via the web portal. This authentication web page is provided by an internal WLAN controller web server in the Cisco Medium Enterprise Design Profile. However, there is an option of using a non-controller-based web authentication server, such as the Cisco NAC Appliance. User names and passwords for authentication can reside on a RADIUS AAA server (Cisco ACS).
- Broadcast SSID is enabled.
- The guest access WLAN uses a QoS profile setting of *bronze* (less than best effort).
- WMM policy is set to *allowed*.

- Radio policy should be set such that client access is allowed to use either 2.4 GHz or 5 GHz.

Additional information about the definition of controller WLANs and SSIDs can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>.

## WLAN Controller Mobility Groups

A *mobility group* is a group of WLAN controllers that behave as a single virtual WLAN controller, sharing essential end client, access point, and RF information. A given WLAN controller is able to make decisions based on data received from other members of the mobility group, rather than relying solely on the information learned from its own directly connected access points and clients. The WLAN controllers in a mobility group form a mesh of authenticated tunnels between themselves, affording any member controller the ability to efficiently communicate with any other member controller within the group.

Mobility groups are used to help facilitate seamless client roaming between access points that are joined to different WLAN controllers. The primary purpose of a mobility group is to create a virtual WLAN domain (across multiple WLAN controllers) to provide a comprehensive view of a wireless coverage area. Typically, two WLAN controllers should be placed in the same mobility group when an inter-controller roam is possible between access points. If the possibility of a roaming event does not exist, it may not make sense to put the WLAN controllers in the same mobility group.

For example, consider the scenario illustrated in Figure 4-8. Here we see a large and a medium building located within the same medium enterprise site. The buildings are in relatively close proximity to one another, with a small building located on a remote site some distance away from the main site. Assume for the purposes of this example that the access points of each building are joined to a different WLAN controller, with the controllers servicing the large and medium building being located within the main service block at the main site, and the WLAN controller servicing the smaller building located in the remote site. The circular and oval patterns surrounding each building are intended to represent a very simplistic view of hypothetical outdoor RF coverage.

**Figure 4-8** Roaming




Figure 4-8 shows that there is overlapping RF coverage between the large and medium buildings, but not between the small building and any other building. This is because users must leave the main site and traverse through a part of the town to get to the smaller remote site, and vice versa. Because roaming is

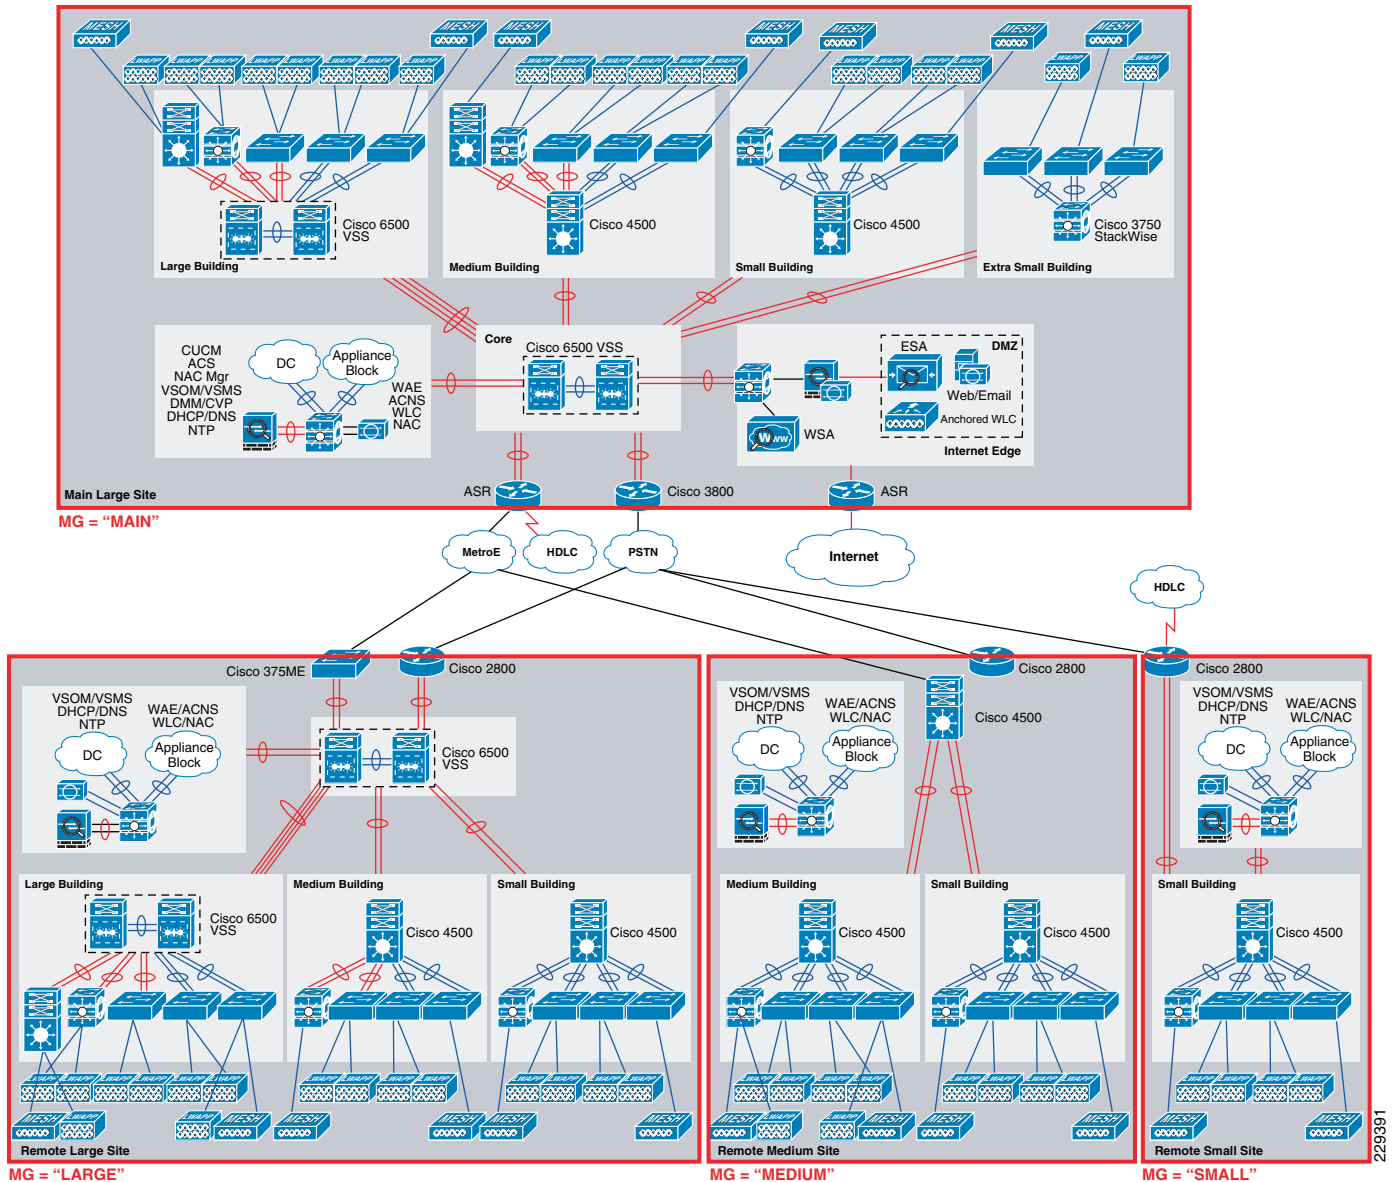
clearly possible between the medium and large building, but not between the small building and any other building on any site, only the WLAN controllers servicing the medium and large building are required to be in the same mobility group. The WLAN controller servicing the small building may be configured to be a member of the same mobility group, but it is not mandatory that this be done in this case.

In applying the concept of mobility groups to the Cisco Medium Enterprise Design Profile, consider the following:

- Within a medium enterprise comprised of one or more sites, it is assumed that intra-site roaming is possible between all buildings resident within the same site. Keep in mind that in reality this may not always be the case, as some sites may have collocated buildings with coverage voids between them. However, assuming that intra-site roaming is possible between all buildings allows us to make a design assumption that is generally applicable to both situations. Thus, in our Medium Enterprise Design Profile, all WLAN controllers serving access points deployed on the same site are placed within the same mobility group.
- It is also assumed that in the vast majority of cases, remote sites are sufficiently distant from the main site (as well as from one another) to render inter-site roaming impractical. Allowing of course for the rare exception that two sites may be adjacent to one another, we assume that roaming between buildings located on different sites is very unlikely.

[Figure 4-9](#) provides a high-level illustration of how mobility group assignment can be handled in the Medium Enterprise Design Profile. Note that *MG* refers to the mobility group name assigned for the site.

Figure 4-9 Medium Enterprise Mobility Groups



The following are some of the key design considerations concerning mobility groups:

- The controllers present at each site are defined as members of a mobility group unique to that site. Each controller in the same mobility group is defined as a peer in the mobility list of all controllers for that mobility group.
- If inter-site roaming between multiple sites is possible, the controllers at each sites should be assigned into the same mobility group and defined as peers in the mobility list of all controllers for that mobility group.
- Because of high-speed WAN/MAN connectivity between sites, access point fail over to a remote backup controller resident at the main site becomes possible. To support this, access points can be configured to fail over to a WLAN controller outside of their mobility group. This is discussed further in [Controller Redundancy](#), page 4-38 and [AP Controller Failover](#), page 4-40.

- A single mobility group can contain a maximum of 72 WLAN controllers. The number of access points supported in a mobility group is bound by the number of controllers and the access point capacity of each controller. Thus, for the Cisco 5508 Wireless Controller, a mobility group can have up to  $72 * 250$ , or 18,000 access points.

The advantage of this mobility group approach is clarity and simplicity in deployment and administration. This is a key point when dealing with medium enterprises that may have limited network administrative staff. By using mobility groups as indicated in [Figure 4-9](#), design simplicity is maintained. Given the large capacity of the Cisco 5508 Wireless Controller, the limitation on the maximum number of controllers per mobility group is not considered to be a significant trade-off.

Additional information about WLAN controller mobility groups, including best practice information, can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL:

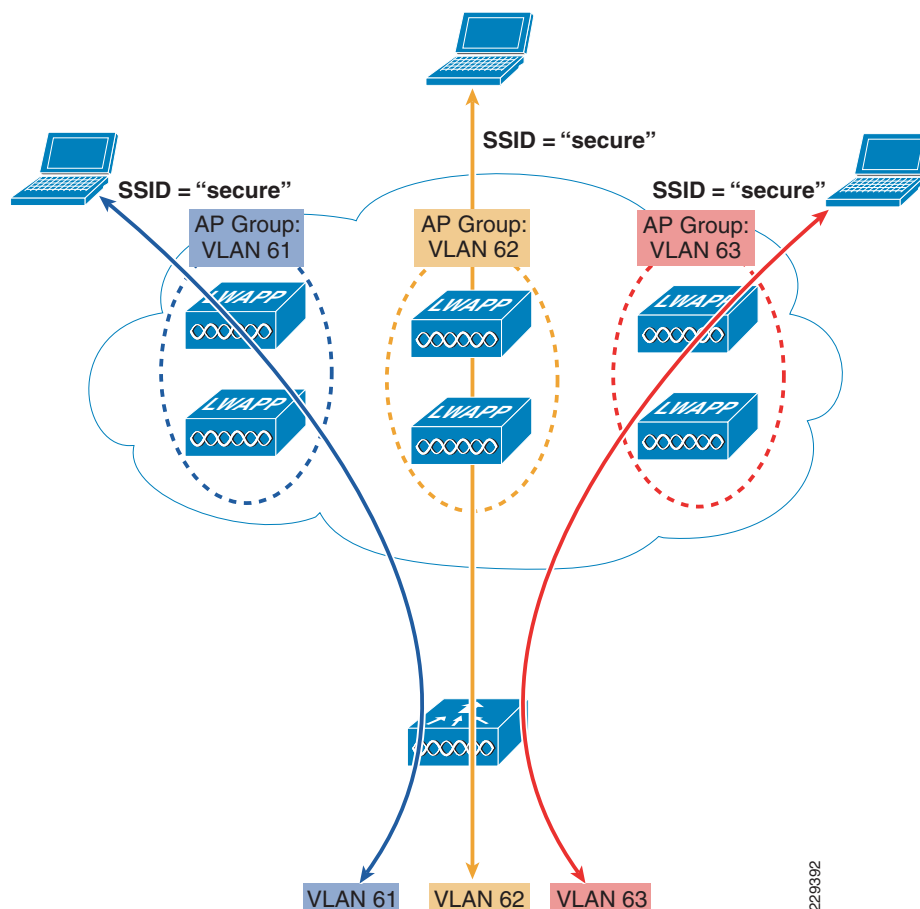
[http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2\\_Arch.html#wp1028143](http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2_Arch.html#wp1028143).

### WLAN Controller Access Point Groups

Typically, each WLAN defined on the controller is mapped to a single dynamic interface (as shown earlier for the secure staff, VoWLAN, and guest access WLANs). Consider the case however, where the Cisco 5508 Wireless Controller is deployed and licensed for 250 access points. Assume also that there are 10 users associated to each access point, using the same WLAN and SSID. This would result in 2500 users sharing the single VLAN to which the WLAN is mapped. A potential issue with this approach is that, depending on the particular overall network design, the use of subnets large enough to support 2500 users may not be possible.

To address this issue, the WLAN can be divided into multiple segments using the AP grouping capability of the WLAN controller. AP grouping allows a single WLAN to be supported across multiple dynamic VLAN interfaces on the controller. This is done by assigning a group of access points to an access point group at the WLAN controller, and then mapping the group to a specific dynamic interface. In this way, access points can be grouped logically, such as by building or set of buildings. [Figure 4-10](#) shows the use of AP grouping based on site-specific VLANs.



**Figure 4-10 Access Point (AP) Groups**

As shown in [Figure 4-10](#), three dynamic interfaces are configured, each mapping to a site-specific VLAN: VLANs 61, 62, and 63. Each site-specific VLAN is mapped to a group of access points that uses the same WLAN/SSID. Each of these access point groups are denoted by an AP Group Name (AP Group VLAN61, VLAN62 or VLAN63). Thus, a staff member associating to the WLAN using an access point that is part of AP group VLAN61 is assigned an IP address from the VLAN 61 IP subnet. Likewise, a staff member associating to the WLAN using an access point that is part of AP group VLAN62 is assigned an IP address from the VLAN 62 IP subnet, and so on. Roaming between the site-specific VLANs is then handled internally by the WLAN controller as a Layer 3 roaming event. As such, the WLAN client maintains its original IP address.

Cisco 5508 Wireless Controllers can contain up to 192 access point group definitions, with up to 16 WLANs defined in each group. Each access point advertises only the enabled WLANs that belong to its access point group. Access points do not advertise disabled WLANs that are contained within its access point group, or WLANs belonging to another access point group.

In implementations of the Cisco Medium Enterprise Design Profile where addressing limitations are present, the use of access point grouping to allow a single WLAN to be supported across multiple dynamic VLAN interfaces on the controller can be extremely beneficial.



## WLAN Controller RF Groups

The strategy behind how *RF groups*, otherwise known as *RF domains*, are deployed within the Cisco Medium Enterprise Design Profile represents another important deployment consideration that can affect overall accessibility. An RF group is a cluster of WLAN controllers that collectively coordinate and calculate their dynamic radio resource management (RRM) settings. Grouping WLAN controllers into RF groups in this way allows the dynamic RRM algorithms used by the Cisco Unified Wireless Network to scale beyond a single WLAN controller. In this way, the benefits of Cisco RRM for a given RF group can be extended between floors, buildings, and even across sites.

**Note**

Complete information regarding Cisco Radio Resource Management can be found in the *Cisco Radio Resource Management under Unified Wireless Networks* at the following URL:  
[http://www.cisco.com/en/US/tech/tk722/tk809/technologies\\_tech\\_note09186a008072c759.shtml](http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml).

If there is any possibility that an access point joined to one WLAN controller may receive RF transmissions from an access point joined to a different WLAN controller, the implementation of system-wide RRM is recommended, to include both controllers and their access points. In this way, RRM can be used to optimize configuration settings to avoid 802.11 interference and contention as much as possible. In this case, both WLAN controllers should be configured with the same RF group name.

In general, simplicity is preferred in the configuration of RF groups within the mobility design. Thus, all WLAN controllers in the Medium Enterprise Design Profile are configured with the same RF group name. Although it is true that geographically disparate WLAN controllers have very little chance of experiencing RF interaction, and thus need not be contained in the same RF domain, for most medium enterprise deployments there is no real disadvantage to doing so. An exception to this would be in extremely large deployments, as the maximum number of controllers that can be defined in a single mobility group is twenty. A clear advantage to this approach is simplicity of configuration and better support of N+1 controller redundancy (see [Controller Redundancy](#), page 4-38 for further details).

A more detailed discussion as well as best practice recommendations regarding the use of RF groups can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL:  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2\\_Arch.html#wp1028184](http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2_Arch.html#wp1028184).

## Access Points

In the Cisco Medium Enterprise Design Profile, it is anticipated that each building requiring WLAN access will be outfitted with dual-band 802.11n access points providing RF coverage in both the 2.4 and 5 GHz bands. It is generally assumed that users will require WLAN access in most building interior areas, plus a 50–75 yard outdoor perimeter area surrounding each building. Of course, it is important to consider that most buildings will almost certainly contain some areas not intended for human entry or occupancy at any time. Similarly, some buildings may possess areas within the aforementioned outdoor perimeter that simply may not be accessible to any users at any time. During your initial mobility design, these vacant areas may not be identified. Therefore, the precise subset of interior and exterior areas requiring WLAN access will likely be better determined instead during the site survey planning process, which is typically an integral part of any wireless network deployment.

**Note**

For more information on site survey planning, see the *Cisco 802.11n Design and Deployment Guidelines* at the following URL:  
[http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/white\\_paper\\_80211n\\_design\\_and\\_deployment\\_guidelines.html](http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/white_paper_80211n_design_and_deployment_guidelines.html).

In most medium enterprises, the vast majority of interior building WLAN access can be provided by the Cisco Aironet 1140 Series 802.11n access point (see [Figure 4-11](#)), which delivers pervasive wireless connectivity while blending in seamlessly with the aesthetics of modern-day enterprise environments.

**Figure 4-11** Cisco Aironet 1140 Series 802.11n Access Point (AIR-LAP1142N)



To deliver the right mix of style and performance, the Cisco Aironet 1140 Series 802.11n access point contains six integrated omni-directional antenna elements that incorporate the use of three hidden discrete elements for each frequency band. Ideal for indoor environments such as offices, conference rooms, corridors, and so on, the Cisco Aironet 1140 Series 802.11n access point has a visually pleasing metal housing covered by a white plastic shell that blends with the most elegant environments. The Aironet 1140 series 802.11n access point provides the ability to be powered directly from 802.3af power-over-Ethernet (PoE) while sustaining full-performance 802.11n connections on both of its radios simultaneously. In the Cisco Medium Enterprise Design Profile, the model of the Cisco 1140 Series 802.11n access point recommended for most interior building locations is the AIR-LAP1142N.



**Note**

Complete information (including country-specific ordering information) regarding the Cisco Aironet 1140 series 802.11n Access Point can be found at the following URL:  
<http://www.cisco.com/en/US/products/ps10092/index.html>

Although the Cisco Aironet 1140 Series 802.11n access point is capable of servicing the bulk of all medium enterprise interior wireless needs, there are some trade-offs to consider in specialized situations. For example, in situations where the results of pre-site survey planning indicate that the use of external antennas are required to best meet specific RF coverage requirements, an access point that provides antenna connectors is necessary. This might include situations where a focused directional antenna pattern is required, or simply one where aesthetic requirements demand that the access point be completely hidden, with only a small antenna footprint exposed to public view. In other cases, perhaps one or more access points will need to be deployed in refrigerated storage, research or even product testing environments where the anticipated operating temperature extremes are not within common norms. Here, extended operating temperature tolerances beyond that of the Cisco Aironet 1140 Series 802.11n access point may be required.

To assist in addressing these and other rare but still significant deployment challenges that may be encountered within medium enterprise sites, the Cisco Aironet 1250 Series 802.11n access point is recommended (see [Figure 4-12](#)).

**Figure 4-12 Cisco Aironet 1250 Series 802.11n Access Point (AIR-LAP1252AG)**

Designed with a next-generation ruggedized modular form factor, the Cisco Aironet 1250 Series 802.11n access point is intended for no-compromise performance in combination with the inherent expand-ability required to address challenging deployment situations. With robust modularized construction and six RP-TNC antenna jacks that allow antennas to be positioned independently of the access point itself, the Cisco Aironet 1250 Series 802.11n access point can be used to address situations requiring focused directional coverage patterns, extended operating temperature capabilities or minimal-footprint installations where it is highly preferable that the access point chassis is totally hidden from view. In the Cisco Medium Enterprise Design Profile, the AIR-LAP1252AG model of the Cisco 1250 Series of access points is recommended for those and other types of demanding deployments.

**Note**

To help discourage theft and vandalism, both the Cisco 1140 as well as 1250 Series 802.11n access points are manufactured with a security slot machined into the access point casing. You can secure either model access point by installing a standard security cable (such as the Kensington Notebook MicroSaver, model number 64068) into the access point security cable slot.

Complete information regarding the Cisco Aironet 1250 series 802.11n access point can be found at the following URL: <http://www.cisco.com/en/US/products/ps8382/index.html>. Additional information concerning the antenna options available for the Cisco Aironet 1250 Series 802.11n access point can be found at the following URL:

[http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/at\\_a\\_glance\\_c45-513837.pdf](http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/at_a_glance_c45-513837.pdf)

Note that Cisco Aironet 1140 Series 802.11n access points can power both 802.11n radios, at full transmit power running two spatial streams with encryption, while drawing only 15.4 watts of power from an 802.3af PoE Catalyst switch. A tradeoff associated with the use of Cisco Aironet 1250 Series 802.11n access points is that the AP-1250 Series requires slightly more power to reach its peak levels of performance, approximately 18.5 to 20 watts of power from a switch capable of providing enhanced-PoE (ePoE). Keep in mind, however, that if the full performance capability of the Cisco Aironet 1250 series access point is not necessary in your particular deployment, or you wish to support only a single RF band (i.e., either 2.4 GHz or 5 GHz) the Cisco Aironet 1250 Series 802.11n access point can also operate with 15.4 watts from a 802.3af PoE Catalyst switch.

To provide the Cisco Aironet 1250 Series 802.11n access point with 20 watts of input power, Cisco recommends the following power options:

- An ePoE Cisco Catalyst switch or switch blade module (such as the 3560-E, 3750-E, 4500E and 6500E Series).
- The use of a mid-span ePoE injectors (Cisco part number AIR-PWRINJ4). This option allows the Cisco Aironet 1250 series 802.11n access point to deliver full 802.11n performance while connected to any Cisco Catalyst switch. Power is injected directly onto the wire by the AIR-PWRINJ4 mid-span injector without reliance on the power output level of the switch itself.

Although its deployment flexibility is unparalleled within the marketplace, in most medium enterprise installations, the Cisco Aironet 1250 series 802.11n access point is typically only deployed only in those locations where they are necessary to address challenging situations. Other trade offs include a higher total cost per access point because of the added cost of external antennas, a larger footprint, and a heavier mounting weight as compared to the Cisco Aironet 1140 series 802.11n access point.



**Note**

For the Cisco Aironet 1250 Series 802.11n access point, Cisco recommends performing your site survey using the same levels of PoE input power as you expect to use in your final deployment. For example, if you plan to deploy Cisco Aironet 1250 Series 802.11n access points with 15.4 watts of PoE, it is recommended for consistency and accuracy that perform your site survey using the same PoE input power levels.

The following design considerations regarding dual-band access points should be kept in mind when designing networks for dense user environments (for example, large cubicle areas, cafeterias, and employee auditoriums within site buildings):

- *Use the 5 GHz band whenever possible*

In general, this applies for both 802.11n as well as pre-802.11n wireless clients. The characteristics of 5 GHz operation make it advantageous for most users, and especially 802.11n users, for the following reasons:

- Despite the maturity of 802.11 wireless LAN technology, the installed base of 5 GHz 802.11a clients generally is not nearly as widespread as 2.4 GHz 802.11b and 802.11g clients. A smaller installed base of users translates into less contention with existing clients and better operation at higher throughput rates.
- The number of non-802.11 interferers (such as cordless phones and wireless personal networks) operating in the 5 GHz band is still just a fraction of the number found within the 2.4 GHz band.
- The amount of available bandwidth found in the 5 GHz band is much greater than that of the 2.4 GHz band. In the United States, there are twenty-one 5 GHz non-overlapping channels that can be deployed. This translates into the ability to deploy with density and capacity in mind, and allow background resources such as Cisco RRM to handle channel and power output requirements accordingly.

- *Design and survey for capacity, not just maximum coverage*

It is a natural tendency to try to squeeze the most coverage from each access point deployed, thereby servicing as much of the site as possible with the lowest total access point investment. When designing networks for high-speed applications, attempting to design for maximum coverage at maximum transmitter output power can be counter-productive, as the maximum coverage footprint is typically attained using lower data rates and degraded signal-to-noise ratios. In addition, such false economies often sacrifice the ability to effectively make use of advanced tools such as Cisco RRM to address anomalies such as “coverage holes” and other deficiencies. Instead, the successful designer should design for capacity and generally aim to have access points installed closer together at lower power output settings. This approach allows for access point transmitter power to be dynamically managed via Cisco RRM. It also allows the practical use of higher data rates, provides RRM with the necessary transmission power “headroom” to allow for the ability to compensate for environmental changes, and facilitates the use of advanced capabilities such as location-based context-aware services.

- *Mount access points or antennas on the ceiling when possible*

Cisco Aironet AP-1140 Series 802.11n access points should be mounted on ceilings only. Ceiling mounting is recommended in general for the types of indoor environments found within medium enterprises, especially for voice applications. In the majority of carpeted indoor environments, ceiling-mounted antennas typically have better signal paths to handheld phones, taking into consideration signal loss because of attenuation of the human head and other obstacles.

Ceiling mounting locations are usually readily available, and more importantly, they place the radiating portion of the antenna in open space, which usually allows for the most efficient signal propagation and reception. Cisco Aironet 1250 Series 802.11n access points can be mounted as deemed necessary during pre-site survey planning or during the actual site survey process. However, ceiling mounting of Cisco Aironet 1250 Series access point antennas is highly recommended, especially when using omni-directional antennas.

- *Avoid mounting on surfaces that are highly reflective to RF*

Cisco recommends that all antennas be placed one to two wavelengths from surfaces that are highly reflective to RF, such as metal. The separation of one or more wavelengths between the antenna and reflective surfaces allows the access point radio a better opportunity to receive a transmission, and reduces the creation of nulls when the radio transmits. Based on this recommendation, a good general rule of thumb then is to ensure that all access point antennas are mounted at least five to six inches away from any large metal reflective surfaces. Note that although recent technological advances have helped greatly in mitigating problems with reflections, nulls, and multipath, a sensible antenna placement strategy still is very important to ensure a superior deployment.

- *Disable legacy and low speed data rates*

Clients operating at low data rates (for example, 1, 2, and 5.5 Mbps) consume more airtime when compared to clients transmitting the same data payloads at higher data rates such as 36 Mbps and 54 Mbps. Overall system performance in any given access point cell drops significantly when a large percentage of low data rate frames tend to consume available airtime. By designing for capacity and disabling lower data rates, aggregate system capacity can be increased.

Unless you are aware of specific reasons why one of the data rates described below are required in your deployment (such as the presence of clients that can transmit or receive *only* at these rates), the following actions are recommended:

- For 2.4 GHz, disable the 1, 2, 5.5, 6, and 9 Mbps rates.
- For 5 GHz, disable at a minimum the 6 and 9 Mbps rates.

A common question concerning 2.4 GHz is why not disable 802.11b entirely? In other words, why not disable the 1, 2, 5.5, and 11 Mbps 2.4 GHz rates altogether? Although this certainly may offer advantages relating to better performance for 802.11g users, this approach may not be entirely practical, especially on guest access WLANs where a visitor might attempt to gain access using a device with embedded legacy radio technology that may not support 802.11g. Because of this, depending on the mix of clients in the environment, it may be wiser to simply disable only the three 802.11b data rates below 11 Mbps. Only if you completely confident that the situation just described is entirely not applicable in your environment should you consider completely disabling all 802.11b data rates.

Additional best practice guidelines for access point and antenna deployments can be found in the following reference documents:

- *Enterprise Mobility 4.1 Design Guide*—  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>
- *Voice Over Wireless LAN 4.1 Design Guide*—  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan41dg/vowlan41dg-book.html>

To provide outdoor WLAN access around the immediate perimeter area of each building, the Cisco Aironet 1520 Series Lightweight Outdoor Access Point is recommended (see [Figure 4-13](#)).

**Figure 4-13** Cisco Aironet 1520 Series Lightweight Outdoor Access Point



As part of the Cisco Medium Enterprise Design Profile, the Cisco Aironet 1520 Series Lightweight Outdoor Access Point provides an outdoor extension to the enterprise wireless network, with central management provided through WLAN controllers and the Cisco Wireless Control System. A very rugged enclosure allows for deployment outdoors without the need to purchase additional housings or third-party National Electrical Manufacturers Association (NEMA) enclosures to provide protection from extreme weather. The robust, weatherized housing of the Cisco Aironet 1520 Series Lightweight Outdoor Access Point can be painted to adapt to local codes and aesthetics.

Although the Cisco Aironet 1520 Series Lightweight Outdoor Access Point is part of the outdoor mesh series of Cisco access point products, a full outdoor mesh infrastructure is beyond the scope of the Cisco Medium Enterprise Design Profile at this time. Rather, in this design Cisco Aironet 1520 Series Lightweight Outdoor Access Points are deployed only as root access points (RAPs), located outdoors on each building in such a manner that a satisfactory outdoor perimeter area is established. The precise location of these outdoor access points, as well as antenna choices, depends on the characteristics associated with the required coverage area and other particulars, and should be determined during pre-site survey planning.

For readers who wish to augment the recommendations made in this design guide and deploy a full site outdoor mesh configuration, see the *Cisco Aironet 1520, 1130, 1240 Series Wireless Mesh Access Points, Design and Deployment Guide*, Release 6.0 at the following URL:

[http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP\\_60.html](http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP_60.html).

In choosing among the various models of Cisco Aironet 1520 Lightweight Outdoor Access Points, readers may also wish to consider whether local, municipal, state or other public safety agencies are currently using or otherwise plan to deploy compatible 4.9 GHz public safety equipment (see note below) in emergency response vehicles. If this is the case, it may be wise to plan ahead in conjunction with on-site and local public safety agencies to accommodate the use of this licensed band for connectivity from properly equipped first responders and emergency vehicles to your WLAN. In the event of a site emergency, the ability to connect to and monitor in-building events, or access key safety and security applications, can significantly enhance the ability of law enforcement and other agencies to locate and combat threats.



**Note**

In 2003, the U.S. Federal Communications Commission (FCC) allocated 50 MHz of spectrum in the 4.9 GHz band to public safety services. Public safety agencies can use this 4.9 GHz band to implement wireless networks with advanced services for the transmission of mission-critical information. Because of the limited number of transmitters and the requirement for licensing, interference on the 4.9 GHz band tends to be below that of other bands, such as 2.4 GHz and 5 GHz. Communications using the 4.9 GHz

public safety band must be related to the protection of life, health, or property. Examples include WLANs for incident scene management, mobile data, video surveillance, VoWLAN, fixed point-to-point, and so on.

Even if 4.9 GHz access is not available at your site, public safety agencies may still be able to access the WLAN using standard 2.4 GHz or 5 GHz unlicensed bands. This depends on whether the emergency response vehicles of the agencies in question are equipped to do so, as well as the configuration of their equipment. Keep in mind that when public safety users access WLANs using unlicensed 2.4 GHz and 5 GHz frequencies, they must also contend for access with other unlicensed users of these frequencies, as well as deal with any interference from other sources located within those bands.

With this in mind, the particular model of outdoor access point recommended for outdoor perimeter building coverage, depending on the inclusion of 4.9 GHz as follows:

- The Cisco Aironet 1524PS (Public Safety) Lightweight Outdoor Access Point includes 4.9 GHz capability and provides flexible and secure outdoor WLAN coverage for both public safety and mobility services. The Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point is a multiple-radio access point that complies with the IEEE 802.11a and 802.11b/g standards, as well as 4.9 GHz public safety licensed operation parameters. This access point can support independent data exchanges across all three radios simultaneously. The main trade-off with the Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point is the added purchase and deployment cost. However, in environments where public safety agencies are already equipped with compatible 4.9 GHz clients, the added benefits and advantages afforded by the 1524PS are often considered worthwhile. The model of Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point recommended in the Cisco Medium Enterprise Design Profile is the AIR-LAP1524PS.
- The Cisco Aironet 1522 Outdoor Lightweight Access Point is a dual-radio, dual-band product that is compliant with IEEE 802.11a (5-GHz) and 802.11b/g standards (2.4-GHz). Designed for demanding environments, the Cisco Aironet 1522 provides high performance device access through improved radio sensitivity and range performance. The trade offs of deploying this model are the lack of 4.9 GHz licensed public safety support in environments where 4.9 GHz is in use among public safety agencies. The model of Cisco Aironet 1522 Lightweight Outdoor Access Point recommended in the Cisco Medium Enterprise Design Profile for deployments without 4.9GHz is the AIR-LAP1522AG.

Cisco offers a wide array of antenna options for the entire range of Cisco Aironet 1520 Series Lightweight Outdoor Access Points. Information on these antenna options can be found in the *Cisco Aironet 1520 Series Lightweight Outdoor Access Point Ordering Guide* at the following URL: [http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps8368/product\\_data\\_sheet0900aecd8066a157.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps8368/product_data_sheet0900aecd8066a157.html).

All models of the Cisco Aironet 1520 Series Lightweight Outdoor Access Point can be powered from a multitude of sources, including PoE, direct DC, or direct AC. The entire range of power input options is described in the *Cisco Aironet 1520 Series Lightweight Outdoor Access Point Ordering Guide*.

**Note**

Although the Cisco Aironet 1520 Series Lightweight Outdoor Access Point can be conveniently powered via PoE, a power injector (Cisco AIR-PWRINJ1500-2) specific to this product line must be used. Do not use any other power injector or Ethernet switch PoE capability (including enhanced PoE switches) in an attempt to directly provide PoE to Cisco Aironet 1520 Series Lightweight Outdoor Access Points. The Cisco Aironet 1520 Series Lightweight Outdoor Access Point is approved for use only with the Cisco AIR-PWRINJ1500-2 power injector. Keep in mind that although the Cisco Aironet 1520 Series Lightweight Outdoor Access Point is intended to be installed exposed to outdoor weather elements, the AIR-PWRINJ1500-2 power injector is approved for indoor installation only.



Some Cisco partners and customers may choose instead to integrate a standard access point into their own weatherproof outdoor enclosure. In this case, it is highly recommended that the Cisco Aironet 1250 Series 802.11n access point be used as the basis for that integration, as its external antenna capabilities would facilitate connection to external antennas via bulkhead connectors. However, integrating a standard indoor access point into a weatherproof outdoor enclosure in this manner has the disadvantage of lacking 4.9 GHz support in areas where public safety agencies are so equipped.

## Usability

This section discusses the mobility design considerations pertaining to those aspects of the Cisco Medium Enterprise Design Profile that are relevant to overall usability, such as the following:

- Quality-of-service (QoS)
- Guest access

## Quality-of-Service

The WLAN controller should be configured to set the 802.1p marking of frames received and forwarded onto the wired VLAN to reflect the QoS policy used on this WLAN. Therefore, if the WLAN controller is connected to a switch that is configured to trust the class-of-service (CoS) and maintain a translation table between CoS and Differentiated Services Code Point (DSCP), the translation between wireless QoS policy and wired network QoS policy occurs automatically.

In the Cisco Medium Enterprise Design Profile, WLAN traffic is prioritized based on the QoS profiles (platinum, silver, bronze, and so on) applied to each WLAN. However, this does not change the IP QoS classification (DSCP) of the client traffic carried, which means that client traffic leaving WLAN controllers may need to be reclassified based on network policy.

This may be achieved via one of following approaches:

- Applying policy at each of the switch virtual interfaces (SVIs) connecting the WLAN controller to the wired network
- Learning the QoS policy that has already been applied by the wireless networking components, because this should already be in alignment with the overall network policy

In the Cisco Medium Enterprise Design Profile, the plan is to use the latter approach, because it provides both the advantage of initial configuration simplicity as well as ongoing ease of maintenance. This technique requires only that the QoS profiles be maintained on the WLAN controllers themselves, without the need to configure explicit policies on adjacent switches. Switches need to be configured to trust only the QoS of frames forwarded to them by the WLAN controller.

To implement this approach, the WLAN controller should be configured to set the 802.1p marking of packets forwarded onto wired VLANs to reflect the QoS policy used on the specific WLAN from which they were received. Therefore, if the WLAN controller is connected to a switch that is configured to trust CoS and maintain a translation table between CoS and DSCP, the translation between wireless and wired network QoS policy occurs automatically.

For example, assume a packet received originates from a WLAN to which a platinum QoS profile has been assigned. This translates to a DSCP value of EF; therefore, the WLAN controller assigns a CoS value of 5 in the header of the frame that carries this data to the wired switch. Similarly, if the same packet originates from a WLAN assigned a QoS profile of silver, the translated CoS value is 0.

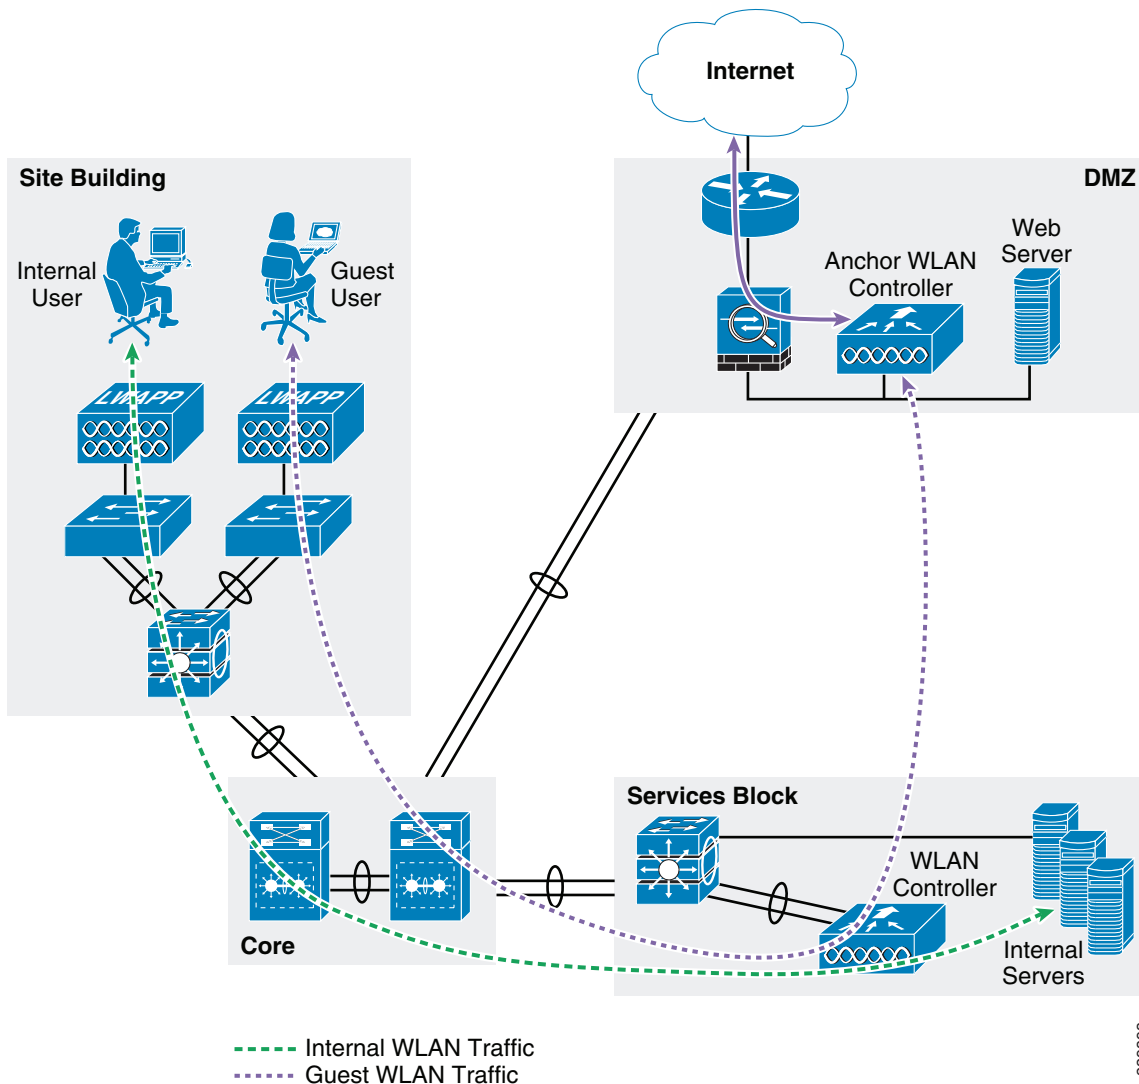
For more information on WLAN QoS, see the following URLs:



- *Voice over Wireless LAN 4.1 Design Guide 4.1*—  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>.
- *Enterprise Mobility 4.1 Design Guide*—  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch5\\_QoS.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch5_QoS.html)

## Guest Access

The Cisco Medium Enterprise Design Profile uses the Cisco Unified Wireless LAN Guest Access option to offer a flexible, easy-to-implement method for deploying wireless guest access via Ethernet over IP (EoIP), as described in RFC3378. EoIP tunneling is used between two WLAN controller endpoints in the centralized network design. The benefit of this approach is that there are no additional protocols or segmentation techniques necessary to achieve guest traffic isolation in relation to other internal traffic. [Figure 4-14](#) shows a high-level view of guest access using this technique with a centralized WLAN controller design.

**Figure 4-14 Guest Access Solution High-Level Overview**

As shown in [Figure 4-14](#), a WLAN controller with a specific purpose is located in the main site DMZ, where it is referred to as an anchor controller. The anchor controller is responsible for terminating EoIP tunnels originating from centralized site WLAN controllers, and interfacing the traffic from these controllers to a firewall or border router. As described in earlier sections of this document, the centralized site WLAN controllers are responsible for termination, management, and standard operation of the various WLANs provisioned throughout the enterprise, including one or more guest WLANs. Instead of being switched locally to a corresponding VLAN on the site controller, guest WLANs are instead transported via the EoIP tunnel to the anchor controller in the DMZ.

When an access point receives information from a WLAN client via the guest access WLAN/SSID, these frames are encapsulated using CAPWAP from the access point to the site WLAN controller. When received at the WLAN controller, they are encapsulated in EoIP from there to the anchor controller. After reaching the anchor controller, these frames are de-encapsulated and passed to a firewall or border router via the guest VLAN. The use of EoIP and an anchor WLAN controller in the DMZ allows guest user traffic to be transported and forwarded to the Internet transparently, with no visibility by, or interaction with, other traffic in the enterprise.

Because the anchor controller is responsible for termination of guest WLAN traffic and is positioned within the Internet DMZ, firewall rules must be established to limit communication between the anchor controller and only those controllers authorized to establish EoIP tunnels to them. Such rules might include filtering on source or destination controller addresses, UDP port 16666 for inter-WLAN controller communication, and IP protocol ID 97 (Ethernet over IP) for client traffic. Other rules that might be needed include the following:

- TCP 161 and 162 for SNMP
- UDP 69 for TFTP
- TCP 80 or 443 for HTTP, or HTTPS for GUI access
- TCP 23 or 22 for Telnet, or SSH for command-line interface (CLI) access

The following are other important considerations to keep in mind regarding the use of this guest access solution:

- For the best possible performance, Cisco strongly recommends that the anchor controller be dedicated to supporting EoIP guest access tunneling only. In other words, do not use the anchor controller for any other purpose but EoIP guest access tunneling. In particular, in addition to its guest access role, the anchor controller should not be used to control and manage other access points in the enterprise.
- When deploying a Cisco 5508 Wireless Controller as an anchor controller, keep in mind that because the anchor controller is not going to be used to manage access points, it can be licensed to support only a minimal number of access points. For example, a Cisco CT5508-12 (12 access point-licensed capacity) can function quite well as an anchor controller in the Cisco Medium Enterprise Design Profile, even in networks where hundreds or thousands of access points may be joined to other Cisco 5508 Wireless Controllers.
- Multicast traffic is not supported over guest tunnels, even if multicast is enabled on wireless controllers.
- The mobility group name of the anchor controller should differ from that configured for site controllers. This is done to keep the anchor controllers logically separate from the mobility groups associated with the general wireless deployment.
- The mobility group name for every WLAN controller that establishes EoIP tunnels with the anchor controller must be configured as a mobility group member in the anchor controller configuration.

Finally, although the focus for the Cisco Medium Enterprise Design Profile is on the pure controller-based guest access solution, note that other, equally functional solutions are available that combine what is discussed in this section with the use of an access control platform external to the WLAN controller. For example, the guest access solution topology described in this section can be integrated with the Cisco NAC Appliance. This might be the case, for example, if the medium enterprise has already deployed the Cisco NAC Appliance within their Internet DMZ to support wired guest access services. As shown in [Figure 4-15](#), the wireless guest access topology remains the same, except that in this scenario, the guest VLAN interface on the anchor controller connects to an inside interface on the NAC Appliance, instead of to a firewall or border router.

Figure 4-15 Cisco UWN Guest Access with Anchor WLC and NAC Appliance

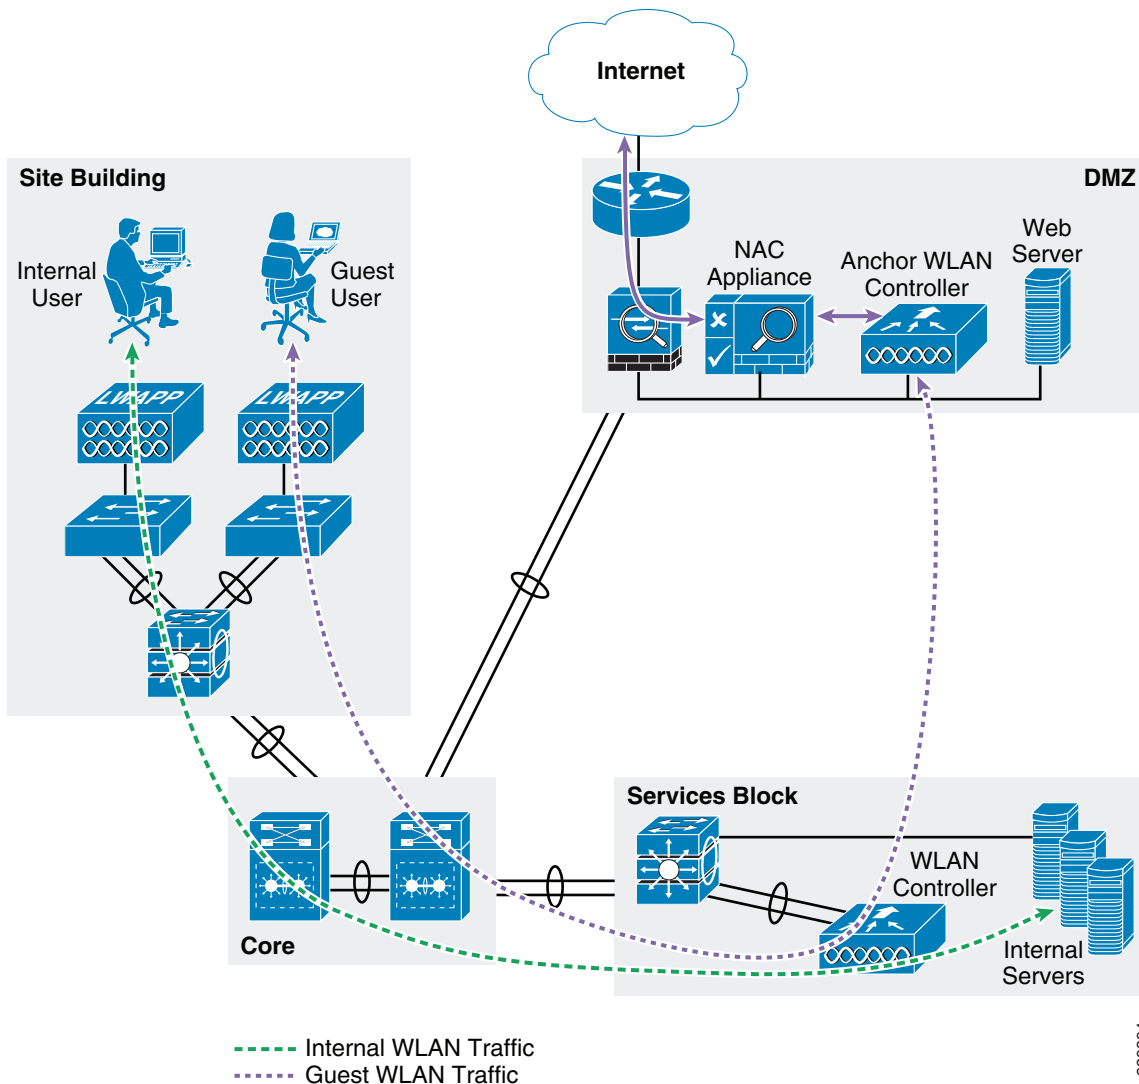


Figure 4-15 shows that the NAC Appliance is responsible for redirection, web authentication, and subsequent access to the Internet. The site and anchor controllers are used only to tunnel guest WLAN traffic across the enterprise into the DMZ, where the NAC appliance is used to actually control guest access. The trade-off here is the added cost of the external access control solution, versus the benefits it affords in relation to your particular deployment.

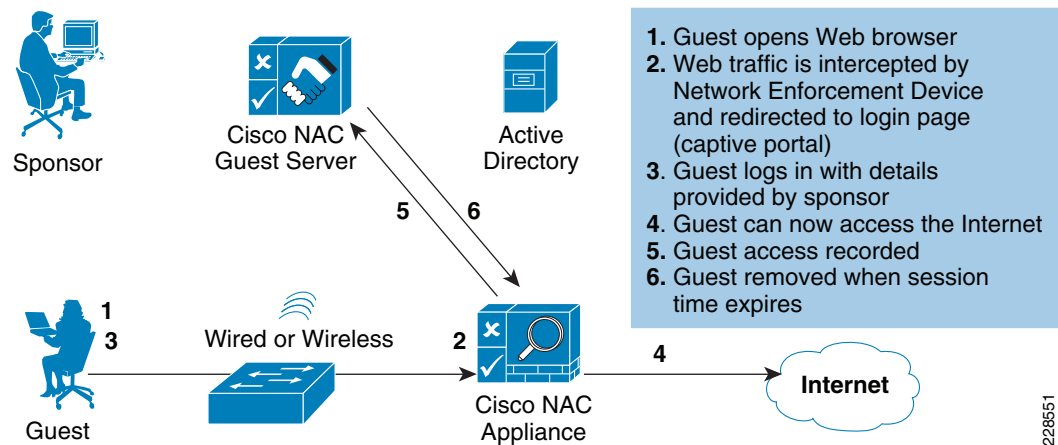
**Note**

Additional information concerning the design and deployment of the Cisco Unified Wireless Network guest access solution can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL: <http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch10GuAc.html#wp999659>.

The Cisco NAC Guest Access Server is another member of the Cisco Network Admission Control solution family that can further enhance the utility of your design by assisting network administrators in the provisioning of guest access user accounts. The NAC Guest Access Server facilitates the creation of guest accounts for temporary network access by permitting provisioning by authorized personnel in a

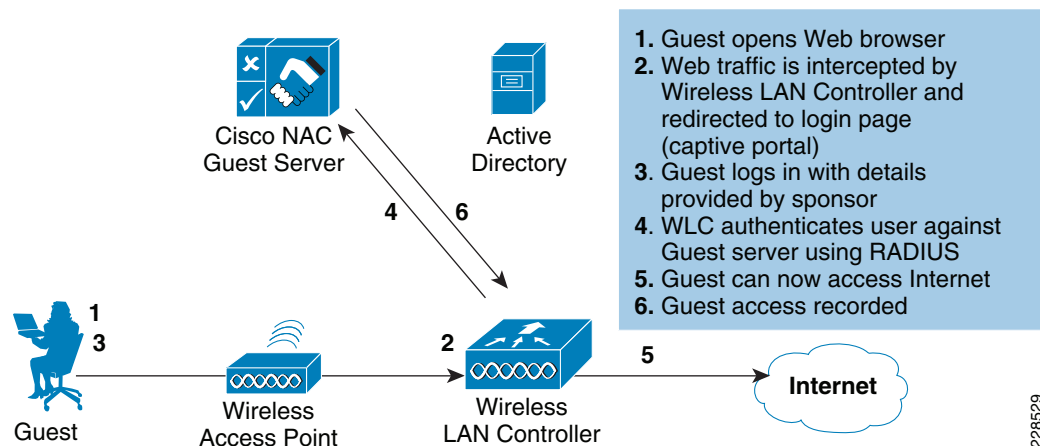
simple and secure manner. In addition, the whole process is recorded in a single place and stored for later reporting, including details of the network access activity. Cisco NAC Guest Server integrates with Cisco NAC Appliance through an application programming interface (API), allowing for guest accounts to be controlled via the Guest Server user interface, including creation, editing, suspension, and deletion of accounts. The Cisco NAC Guest Server then controls these accounts on the Cisco NAC Appliance through the API (shown in Figure 4-16). In addition, the Guest Server receives accounting information from the NAC Appliance to enable full reporting.

**Figure 4-16 NAC Guest Server with NAC Appliance and WLAN Controller**



Cisco NAC Guest Server can also integrate directly with Cisco WLAN controllers through the RADIUS protocol, allowing for guest accounts to be controlled via the Guest Server user interface, including the creation, editing, and deletion of guest accounts. In this case, the WLAN controller makes use of the NAC Guest Server to authenticate guest users (shown in Figure 4-17). In addition, the Guest Server receives accounting information from the WLAN controller to enable full reporting.

**Figure 4-17 NAC Guest Server with WLAN Controller Alone**



**Note**

For more information on the Cisco NAC Guest Server, see the following URL:

[http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5707/ps8418/ps6128/product\\_data\\_sheet0900aecd806e98c9.html](http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5707/ps8418/ps6128/product_data_sheet0900aecd806e98c9.html).

# Manageability

As mentioned earlier, each WLAN controller in the Cisco Medium Enterprise Design Profile provides both a CLI as well as a graphical web user interface, which are primarily used for controller configuration and management. These user interfaces provide ready access to the network administrator. However, for a full-featured, centralized complete life cycle mobility management solution that enables network administrators to successfully plan, configure, deploy, monitor, troubleshoot, and report on indoor and outdoor wireless networks, the use of the Cisco Wireless Control System (WCS) is highly recommended (see Figure 4-18).

**Figure 4-18 Cisco Wireless Control System**



The Cisco Wireless Control System allows very effective management of wireless networks supporting high-performance applications and mission-critical solutions. Effective management of these networks helps to simplify network operation and improve the productivity of employees, administrators, guests and site visitors. The comprehensive Cisco WCS platform scales to meet the needs of small, midsize, and large-scale WLANs across local and remote sites. Cisco WCS gives network administrators immediate access to the tools they need when they need them, wherever they may be located within the enterprise.

Operational costs are significantly reduced through a simplified and intuitive GUI, with built-in tools delivering improved efficiency and helping to reduce training costs, even as the site network grows incrementally larger. Cisco WCS lowers operational costs by addressing the whole range of mobility management requirements (radio frequency, access points, controllers, mobility services, and so on) using a single unified management platform deployed in a centralized location, and with minimal impact on staffing requirements.

Cisco WCS can scale to manage hundreds of Cisco WLAN controllers, which in turn can manage thousands of Cisco Aironet access points. For installations where network management capabilities are considered mission-critical, WCS also supports a software-based high availability option that provides failover from a primary (active) WCS server to a secondary (standby). Adding mobility services such as context-aware software and adaptive wireless intrusion prevention systems (wIPS) is simplified through Cisco WCS integration with the Cisco Mobility Services Engine (MSE).




## Note

A detailed description of each management feature and benefit available in the Cisco Wireless Control System is beyond the scope of this chapter, but the information can be found at the following URL: [http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product\\_data\\_sheet0900aec802570d0.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aec802570d0.html).

In the Cisco Medium Enterprise Design Profile, a centralized WCS management server located in the data center block within the main site is used. The data center block was initially shown in Figure 4-3. Figure 4-19 provides greater detail and magnification.

**Figure 4-19 WCS Within the Data Center Block**



The current upper limit for scaling WCS on a high-end server is up to 3000 Cisco Aironet CAPWAP-based access points, and up to 750 Cisco WLAN controllers. As such, most implementations of the Cisco Medium Enterprise Design Profile are well served by a mobility design using a WCS management server located on the main site.



**Note**

For further information on WCS hardware platforms and requirements, see the following URL:  
[http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6\\_0wst.html#wp1061082](http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0wst.html#wp1061082).

The planning, configuration, deployment, monitoring, reporting, auditing, and troubleshooting made available by WCS are accessible to any authorized medium enterprise network administrator via standard secured web browser access.

Generally speaking, it is anticipated that access to WCS will be restricted to network administrators and their staff located at the main and remote sites, as well as local site administrators for sites where network administrators are not present. However, these groups will most likely not have equivalent resource and functionality access. It is anticipated that resource access will be limited further, based on administrative level and assigned site or sites.

In this design, the ability to query and manage site mobility resources is regulated using the virtual domain feature of WCS, in conjunction with the appropriate assignment of WCS user rights. Thus, although key members of the main site central network administration staff may possess the authority to manage any and all mobility resources located on any site throughout the enterprise, remote site administrators may be limited by the following:

- *Site resource management visibility policy*—This is performed by assigning the network mobility infrastructure components associated with each site to a WCS virtual domain, and assigning the virtual domains to appropriate network administrators. Key members of the central administrative staff are assigned to the WCS root domain, granting them overall authority to view and configure all mobility infrastructure resources, on any site, via their WCS management consoles. However, personnel responsible for local site network administration are restricted to the discrete mobility infrastructure components associated with the virtual domain representing their local site. These infrastructure components include WLAN controllers, access points, configuration templates, WCS events, reports, alarms, WLAN clients, and so on.
- *Site resource management access policy*—Although the visibility of a resource is determined by WCS virtual domain assignment, the subset of acceptable actions that are allowed against any visible resources are further regulated by the assignment of appropriate WCS user and group rights, which allow policies to be applied that further limit what actions each may be allowed against any visible resources.



Via the WCS GUI interface, virtual domains (as well as WCS user rights) can be assigned at the WCS server or using an external security manager such as Cisco Secure ACS.

**Note**

Further information regarding how WCS virtual domains may be used to limit individual site network administrator access to segments of the mobility network outside of their scope of responsibility, while still providing for overall “root” administrator control of the entire wireless network, may be found at the following URL:

[http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/brochure\\_c02-474335.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/brochure_c02-474335.html).

Guest access credentials can be created and managed centrally using the Cisco WCS. A network administrator can create a limited privilege account within WCS that permits “lobby ambassador” access for the purpose of creating guest credentials. With such an account, the only function a lobby ambassador is permitted is to create and assign guest user credentials to controllers that have web-policy configured WLANs. In the rare event that a centralized WCS management system is not available because of a server failure, a network administrator can establish a local administrator account on the anchor WLAN controller, with lobby ambassador privileges, as a backup means of managing the guest access solution.

The use of a centralized WCS management server in the Cisco Medium Enterprise Design Profile provides key advantages such as reduced initial deployment cost and ease of maintaining server resources in a centralized location, coupled with good performance across modern high-speed LANs and WANs. Of course, as with any design choice, certain trade offs exist, such as the following:

- *WCS server failure*

In the Cisco Medium Enterprise Design Profile, the centralized mobility network management services provided by WCS are not regarded as being mission-critical for the majority of medium enterprise deployments. Thus, in the rare event of a WCS server failure, it is assumed that direct WLAN controller management work a rounds (such as that described earlier for guest access management) are an acceptable cost compromise. Any downtime realized because of a WCS server failure, although undoubtedly very inconvenient, would in most cases not be viewed as entirely catastrophic. This being the case, the Cisco Medium Enterprise Design Profile does not at this time provide for the added cost of a secondary WCS management server in an N+1 software-based high-availability arrangement. However, deployments where WCS management services are critical to the mission of the medium enterprise should instead consider modifying the design to include the services of a secondary WCS management platform configured for N+1 software-based high-availability.

**Note**

For more information on WCS high availability configurations, see the following URL:

[http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6\\_0admin.html#wp1132580](http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0admin.html#wp1132580).

- *Unrecoverable WAN failure*

A catastrophic, unrecoverable WAN failure can interrupt management traffic between WCS and the WLAN controllers that are located on remote sites. One way to protect against this is to distribute the WCS management server function out further into the network, and centralize WCS management on a per-site basis. However, this increases the cost of WCS deployment significantly, requiring one WCS management server per site, and preferably a Cisco WCS Navigator management aggregation platform located at the main site. Because it is believed that the centralized mobility network management services provided by WCS are not regarded as mission-critical to the majority of medium enterprises, these decentralized management options are not included in the



Cisco Medium Enterprise Design Profile at this time. Instead, it is assumed that in this type of a rare occurrence, the aforementioned ability to minimally manage WLAN controllers directly will suffice, should any network management intervention be required in such circumstances.

**Note**

For more information on WCS Navigator, see the following URL:  
<http://www.cisco.com/en/US/products/ps7305/index.html>.

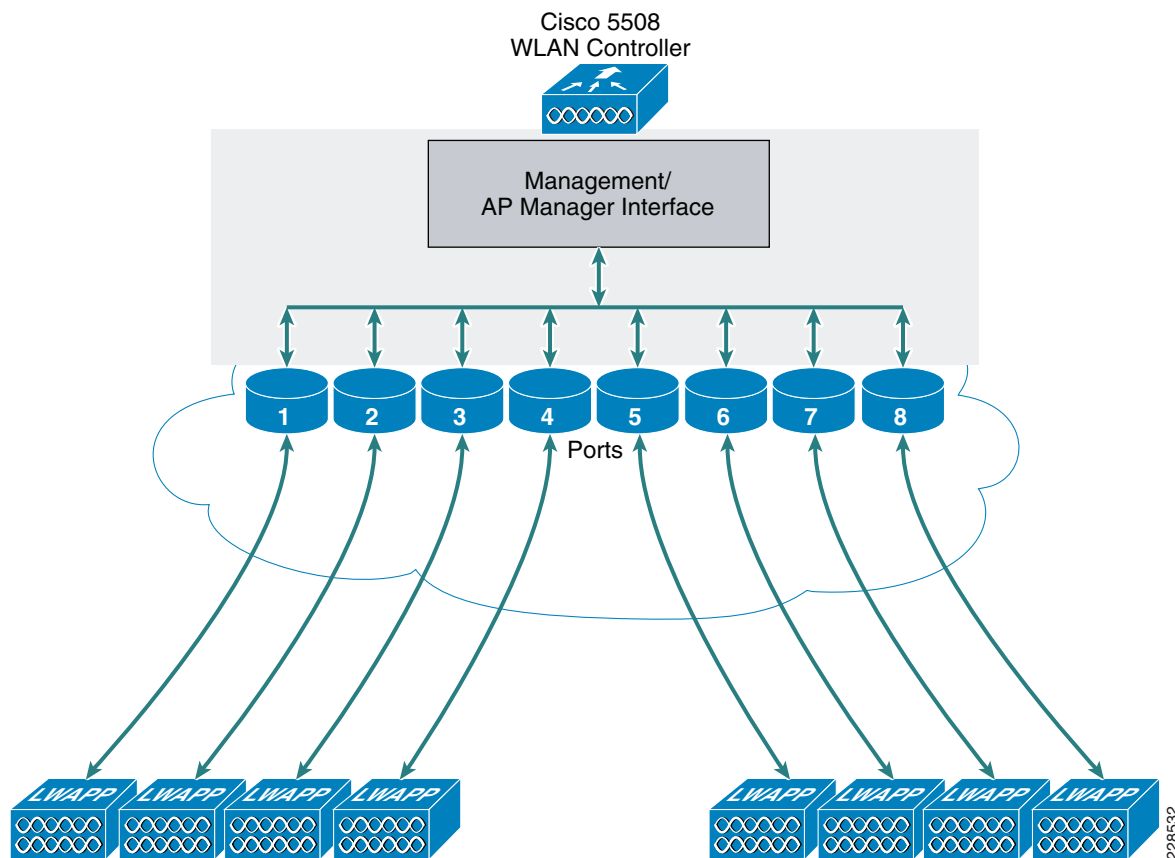
## Reliability

This section discusses the mobility design considerations pertaining to those aspects of the Cisco Medium Enterprise Design Profile relevant to overall reliability, and includes the following:

- Controller link aggregation
- Controller redundancy
- AP controller failover

### Controller Link Aggregation

An important capability used to enhance the reliability of WLAN controller interconnection to the wired network is *link aggregation (LAG)*. As mentioned earlier, LAG is a partial implementation of the 802.3ad port aggregation standard. It bundles all the controller distribution system ports into a single 802.3ad port channel, thereby reducing the number of IP addresses needed to make use of all controller wired ports. When LAG is enabled, the system dynamically manages port redundancy and load balances access points across each port, without interaction from the network administrator. With the Cisco 5508 Wireless Controller and the release 6.0 software used in the Cisco Medium Enterprise Design Profile, all eight ports can be bundled together into a single Gigabit EtherChannel interface. LAG is effective in distributing access point traffic across all controller ports, as shown in [Figure 4-20](#). This can be especially important with high capacity controllers licensed for many access points, such as the Cisco CT5508-250.

**Figure 4-20** LAG in the Cisco 5508 WLC

LAG simplifies controller configuration and improves the overall solution reliability. If any of the controller ports fail, traffic is automatically migrated to one of the remaining ports. As long as at least one controller port is functioning, the system continues to operate, access points remain connected to the network, and wireless clients continue to send and receive data.

The Gigabit Ethernet connections comprising the LAG (up to eight on the Cisco 5508 Wireless Controller) should be distributed among different modular line cards or switch stack members in the services block to the greatest degree possible. This is done to ensure that the failure of a single line card or switch stack member does not result in total failure of the WLAN controller interconnection to the network.

For example, if there are four switch stack members in the services block and LAG is configured using all eight WLAN controller interfaces, the Gigabit Ethernet links from the services switch block to the WLAN controller should be distributed two per services block switch stack member. In this way, if any switch stack member fails, six other Gigabit Ethernet links to the WLAN controller remain ready, active, and available to pass data.

The switch features required to implement this connectivity between the WLAN controller and the services block are the same switch features that are otherwise generally used for EtherChannel connectivity between switches.

When using a Cisco 5508 Wireless Controller with link aggregation enabled, it is important to keep the following considerations in mind:

- When the port channel is configured as “on” at both ends of the link, it does not matter if the Cisco Catalyst switch is configured for either Link Aggregation Control Protocol (LACP) or Cisco proprietary Port Aggregation Protocol (PAgP), because no channel negotiation occurs between the controller and the switch.

The recommended load balancing method for Cisco Catalyst switches is by use of the CLI command **src-dest-ip**.

- You cannot configure the controller ports into separate link aggregation groups. Only one link aggregation group is supported per controller. Therefore, you can connect a controller in link aggregation mode to only one neighbor switch device (note that this can be a switch stack with multiple member switches).
- When you enable link aggregation or make any changes to the link aggregation configuration, you must immediately reboot the controller.
- When you enable link aggregation, only one AP manager interface is needed because only one logical port is needed. The in-band management interface of the Cisco 5508 Wireless Controller can also serve as the AP manager interface.
- When you enable link aggregation, all Cisco 5508 Wireless Controller distribution ports participate in link aggregation by default. Therefore, you must configure link aggregation for all the connected ports in the neighbor switch that have been outfitted with small form-factor plug-in (SFP) modules.
- When you enable link aggregation, only one functional physical distribution port is needed for the controller to pass client traffic. Although Cisco 5508 Wireless Controllers have no restrictions on the number of access points per port, Cisco recommends that if more than 100 access points are connected to the controller, make sure that at least two or more Gigabit Ethernet interfaces are used to connect the controller to the services block.
- As mentioned previously, there are eight SFP interfaces on the Cisco 5508 Wireless Controller. These may be fully deployed to take full advantage of multi layer site design guidelines regarding the oversubscription of access layer uplinks. By doing so, it is relatively straightforward to design a solution that delivers access layer uplinks from the WLAN controller with an oversubscription rate of between 8:1 and 20:1 (Note that these oversubscription rates are not unique to wireless products and are equivalent with what is typically seen in wired networks as well.)

Table 4-1 provides information for the Cisco 5508 Wireless Controller deployed with its maximum complement of 250 access points.

**Table 4-1 Cisco 5508 Wireless Controller Oversubscription Rates**

| Throughput per AP (Mbps) | Cisco 5508 Wireless Controller Oversubscription Rate (8 Gbps) |
|--------------------------|---------------------------------------------------------------|
| 25                       | 1:1                                                           |
| 50                       | 2:1                                                           |
| 100                      | 4:1                                                           |
| 150                      | 5:1                                                           |
| 200                      | 7:1                                                           |
| 250                      | 8:1                                                           |

Table 4-1 shows that even if designing for peak 802.11n throughput of 250 Mbps per access point, oversubscription is not expected to exceed site design guidelines of 8:1 when using all the available controller interfaces with LAG.

**Note**

For more information concerning WLAN controller link aggregation, see *Deploying Cisco 440X Series Wireless LAN Controllers* at the following URL:

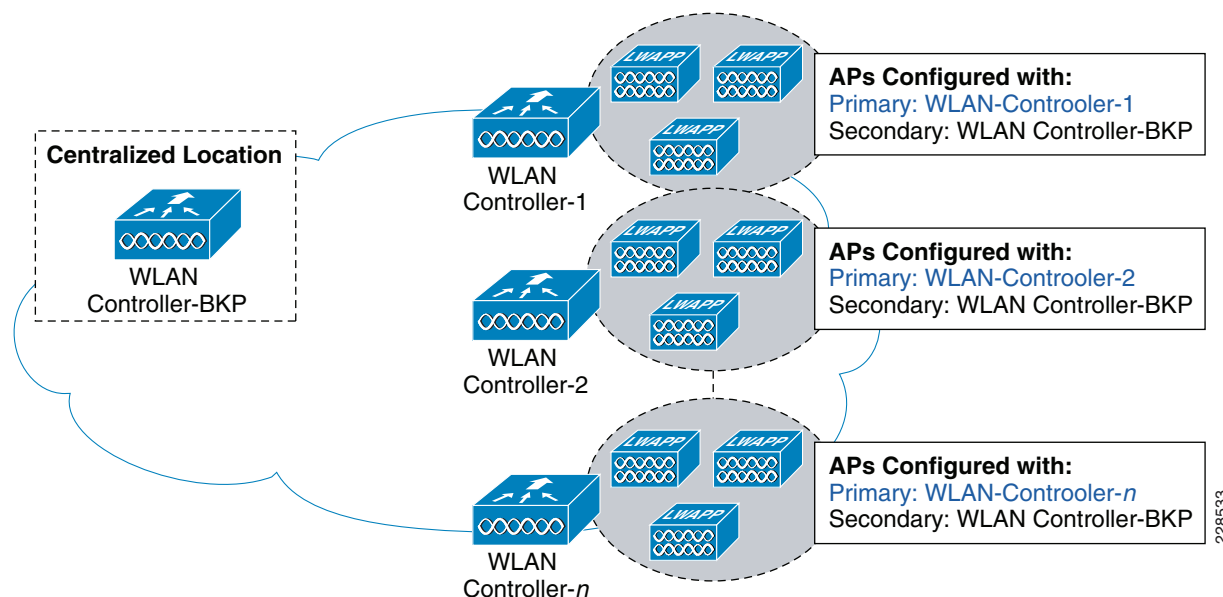
<http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html#wp1062211>.

## Controller Redundancy

The ability of the solution to recover from a reasonable degree of component failure is important in ensuring the reliability of any WLAN networking solution. This is especially important when there are many users that may rely on a centralized component, such as a WLAN controller, for access into the network. An easy solution is to have a “hot” standby secondary controller always at the ready for each primary controller in active service (otherwise known as 1:1 controller redundancy). Although this offers the highest degree of protection from any number of failed primary controllers, it is also the most costly approach.

In the Cisco Medium Enterprise Design Profile, unforeseen controller failures are avoided using an “N+1” controller redundancy model, in which the redundant WLAN controller is placed in a central location and acts as a backup for multiple active WLAN controllers. Each access point is configured with the name or IP address of its primary WLAN controller, but is also configured with the name or IP address of the redundant controller as its secondary WLAN controller. The N+1 controller redundancy approach is based on the assumption that the probability of more than one primary WLAN controller failure occurring simultaneously is very low. Thus, by allowing one centralized redundant controller to serve as the backup for many primary controllers, high availability controller redundancy can be provided at a much lower cost than in a traditional 1:1 redundancy arrangement. Figure 4-21 provides a general illustration of the principle of N+1 controller redundancy.

**Figure 4-21 General N+1 WLAN Controller Redundancy**



The main tradeoff associated with the N+1 redundancy approach is that the redundant controller may become oversubscribed if multiple primary controllers fail simultaneously. In reality, experience indicates that the probability of multiple controller failures is low, especially at geographically separate

site locations. However, when designing an N+1 redundant controller solution, you should assess the risk of multiple controller failures in your environment as well as the potential consequences of an oversubscribed backup controller. In situations where there is reluctance to assume even this generally small degree of risk, other controller redundancy approaches are available that can provide increasingly greater degrees of protection, albeit with associated increases in complexity and equipment investment.

**Note**

For more details on controller redundancy, see *Deploying Cisco 440X Series Wireless LAN Controllers* at the following URL:

<http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html#wp1060810>.

The configuration of N+1 redundancy in any mobility design depends greatly on the licensed capacity of the controllers used and the number of access points involved. In some cases, configuration is rather straightforward, emulating what is shown in [Figure 4-21](#) by having the access points of the main site as well as all remote sites address a common redundant controller located in the main site services block. In other cases, there may be sufficient capacity on the primary controllers located on the main site themselves to accommodate the access point and user load of a single failed controller on any of the remote sites. This approach requires that main site controllers be licensed for a greater number of access points than necessary for the support of the main site alone. Additional licensing of existing controllers is performed in place of providing a dedicated additional controller platform at the main site for system-wide redundancy. In this case, the available capacity of the primary main site WLAN controllers allow them to act as the secondary destination for the access associated with the largest remote site. Thus, in this particular case, the need to deploy hardware at the main site explicitly for the purposes of controller redundancy may be avoided.

For example, assume that the main site shown in [Figure 4-3](#) contains a total of 250 combined access points across all main site buildings, and the largest of the remote sites also contains 250 combined access points across all remote site buildings. In this case, if the main site services block is equipped with two Cisco CT5508-250 WLAN controllers (the “-250” signifies that this particular Cisco 5508 Wireless Controller is licensed for 250 access points), the access point load of the main site alone can be split equally between the two controllers (125 access points on each controller). This leaves ample capacity in the main site for one of the following scenarios to occur:

- Either of the main site controllers may fail and allow up to 125 joined access points to migrate (fail over) to the other controller in the pair. This results in the remaining functional controller bearing the full load of 250 access points.
- Any remote site controller may fail and allow its joined access points to migrate (fail over) to the main site controllers. In the case of a failure of the largest remote site, this results in each of the main site controllers operating at their full licensed capacity.

Further information regarding WLAN controller redundancy may be found in the following documents:

- *Deploying Cisco 440X Series Wireless LAN Controllers*—  
<http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html#wp1060810>
- *Enterprise Mobility 4.1 Design Guide*—  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>

## AP Controller Failover

The Cisco Unified Wireless Network provides for multiple failover options that can allow access points to determine which WLAN controller to migrate in the event of a controller failure, based on pre-configured priorities. When an access point goes through its discovery process, it learns about all the WLAN controllers in its mobility group. The access point can prioritize which controller it attempts to join based on its high availability configuration, or choose a WLAN controller based on loading.

In the Cisco Medium Enterprise Design Profile, a high-speed WAN/MAN is present between sites, thus making access point fail over to a remote WLAN controller feasible, as described in the previous section. To accomplish this in the Cisco Medium Enterprise Design Profile, access points can be configured to fail over to a WLAN controller that is outside their mobility group. In this scenario, the remote WLAN controller is not in the mobility group that is learned during the AP discovery process, and the IP address of the remote WLAN controller must be provided in the HA configuration.

For this to be effective, however, a common WLAN SSID naming policy for key WLANs must be implemented to ensure that WLAN clients do not have to be re configured in the event of an access point fail over to the main site backup controller.

Best practice considerations regarding to AP controller failover include the following:

- After access points initially discover a WLAN controller, access points should be manually assigned to primary and secondary controllers. By doing this, AP assignment and WLAN redundancy behavior is deterministic.
- A common WLAN SSID naming policy is necessary to ensure that WLAN clients do not have to be re configured in the event of an access point fail over to a central backup controller. The SSID used to access a particular WLAN throughout the multisite medium enterprise should be the same, regardless of the controller.
- WLAN controllers have a configurable parameter known as *AP Fallback* that causes access points to return to their primary controllers after a failover event, after the primary controller comes back online. This feature is enabled by default. However, leaving this parameter at the default value can have some unintended consequences. When an access point “falls back” to its primary controller, there is a brief window of time, usually approximately 30 seconds or so, during which service to wireless clients is interrupted because the access points are busy re-joining the primary controller. In addition, if connectivity to the primary WLAN controller becomes unstable for some reason, the access point might “flap” between the primary controller and the backup. For this reason, it is preferable to disable AP Fallback and, in the rare event of a controller failure, move the access points back to the primary controller in a controlled fashion during a scheduled service window.

**Note**

For more information and best practices regarding AP controller failover, see the *Enterprise Mobility 4.1 Design Guide* at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>.

## Wireless LAN Controller Configuration

The core component of the Cisco Unified Wireless architecture is the WLAN Controller (WLAN controller) that provides the interface between the “split-MAC” wireless network and the wired network. That is, the WLAN controller is the Layer-2 connection point between WLAN client traffic and the wired network, making the WLAN controller an aggregation and control point for WLAN traffic. In addition, the WLAN controller is the primary control point of AP and RF management.


The reference design used for testing WLAN services in the Medium Enterprise Design Profile mobility design uses the following four WLAN controllers:

- Two WLAN controllers (cr23-5508-1, cr23-5508-2) for the main site
- One WLAN controller (cr14-5508-1) for a remote site
- One anchor WLAN controller (cr11-5508-wlc) for guest services

## WLAN Controller and Wired Network Connections

The WLAN controllers in the main site are centralized for that site and connected to a 3750E stack in services block connected to the site core, as shown in [Figure 4-22](#). These WLAN controllers provide WLAN services for the entire site, as well as fail over support for APs in remote sites, in the event of WLAN controller outage at that location. The number of WLAN controllers for the main site is driven by the number of APs deployed and the type of fail over support required. In this example, two WLAN controllers are used to illustrate the basic configuration requirements.

**Figure 4-22 Services Block WLAN Controller Connection**



The two main site WLAN controllers share the VLAN and subnet configuration, differing only in their IP addressing. [Figure 4-23](#) and [Figure 4-24](#) show the interface summary on the two WLAN controllers. The two key interfaces are highlighted, that is the management and virtual interfaces. The management interface is used as the interface for in-band communication with the WLAN controller, including CAPWAP tunnel termination (there is no AP-manager interface), and the virtual interface is used to support mobility.



### Note

Although the 1.1.1.1 address has been used in example mobility configurations, the 1.0.0.0/8 address range has now been assigned, and it is best that customers use a private address that would not be a valid address within their own network.

**Figure 4-23 cr23-5508-1 Interfaces**

CISCO

MONITORWLANsCONTROLLERWIRELESSSECURITYMANAGEMENTCOMMANDSHELPFEEDBACK

Controller

GeneralInventoryInterfacesMulticastNetwork RoutesInternal DHCP ServerMobility ManagementPortsNTP

Interfaces

New...

| Interface Name               | VLAN Identifier | IP Address     | Interface Type | Dynamic AP Management |
|------------------------------|-----------------|----------------|----------------|-----------------------|
| <a href="#">staff_voice</a>  | 112             | 10.125.30.18   | Dynamic        | Disabled              |
| <a href="#">staff_data</a>   | 113             | 10.125.30.34   | Dynamic        | Disabled              |
| <a href="#">management</a>   | 111             | 10.125.30.2    | Static         | Enabled               |
| <a href="#">nac</a>          | 117             | 10.125.30.99   | Dynamic        | Disabled              |
| <a href="#">service-port</a> | N/A             | 172.26.158.243 | Static         | Not Supported         |
| <a href="#">virtual</a>      | N/A             | 1.1.1.1        | Static         | Not Supported         |

29395



**Figure 4-24** *cr23-5508-2 Interfaces*

| Interface Name               | VLAN Identifier | IP Address     | Interface Type | Dynamic AP Management |
|------------------------------|-----------------|----------------|----------------|-----------------------|
| <a href="#">staff_voice</a>  | 112             | 10.125.30.19   | Dynamic        | Disabled              |
| <a href="#">staff_data</a>   | 113             | 10.125.30.35   | Dynamic        | Disabled              |
| <a href="#">management</a>   | 111             | 10.125.30.3    | Static         | Enabled               |
| <a href="#">nac</a>          | 117             | 10.125.30.100  | Dynamic        | Disabled              |
| <a href="#">service-port</a> | N/A             | 172.26.158.244 | Static         | Not Supported         |
| <a href="#">virtual</a>      | N/A             | 1.1.1.1        | Static         | Not Supported         |

Figure 4-25 shows the management interface of WLAN controller cr23-5508-1. Note that Link Aggregation (LAG) is enabled on all the WLAN controllers used in the Medium Enterprise Design Profile mobility design.

**Figure 4-25** *cr23-5508-1 Management Interface*

**General Information**

Interface Name: management  
MAC Address: 00:24:97:cf:3f:a0

**Configuration**

Quarantine: ☐  
Quarantine Vlan Id: 0

**NAT Address**

Enable NAT Address: ☐

**Interface Address**

VLAN Identifier: 111  
IP Address: 10.125.30.2  
Netmask: 255.255.255.240  
Gateway: 10.125.30.1

**Physical Information**

The interface is attached to a LAG.  
Enable Dynamic AP Management: ☒

Example 4-1 and Example 4-2 show examples of the switch configuration for the 3750 stack switch connecting the main WLAN controllers to the wired network.

#### **Example 4-1** *Example of WLAN Controller 3750 Stack Port Channel Configuration*

```
interface Port-channel11
```



```
description cr23-5508-1
switchport trunk encapsulation dot1q
switchport trunk native vlan 801
switchport trunk allowed vlan 111-115,117,313,317
switchport mode trunk
switchport nonegotiate
load-interval 30
carrier-delay msec 0
hold-queue 2000 in
hold-queue 2000 out
end
```

#### Example 4-2 Example of WLAN Controller 3750 Stack Interface Configuration

```
interface GigabitEthernet1/0/10
description Connected to cr23-5508-1 port Gi0/0/1 via CG#11
switchport trunk encapsulation dot1q
switchport trunk native vlan 801
switchport trunk allowed vlan 111-115,117,313,317
switchport mode trunk
switchport nonegotiate
load-interval 30
carrier-delay msec 0
udld port
mls qos trust cos
channel-group 11 mode on
hold-queue 2000 in
hold-queue 2000 out
end

interface GigabitEthernet2/0/10
description Connected to cr23-5508-1 port Gi0/0/2 via CG#11
switchport trunk encapsulation dot1q
switchport trunk native vlan 801
switchport trunk allowed vlan 111-115,117,313,317
switchport mode trunk
switchport nonegotiate
load-interval 30
carrier-delay msec 0
udld port
mls qos trust cos
channel-group 11 mode on
hold-queue 2000 in
hold-queue 2000 out
end
```

## Remote Site

The remote site WLAN controller and wired network connection is the same as that used in the main site. In other words, WLAN controller is connected to a 3750E stack that acts as a services block for the remote site. The configuration is the same; therefore, the details are not duplicated here.

## Mobility Groups

The primary purpose of a Mobility Group in the Cisco Unified Wireless Network (CUWN) is to share client information between WLAN controllers. This helps to ensure seamless mobility when clients roam between APs that are connected to WLAN controllers within the same Mobility Group. The default Mobility Group Name is created in the Controller General configuration page, as shown in [Figure 4-26](#).

**Figure 4-26** *cr23-5508-1 Mobility Group Definition*

The screenshot shows the Cisco Wireless LAN Controller configuration interface for controller **cr23-5508-1**. The **General** tab is selected under the **Mobility Management** section. The **Default Mobility Domain Name** is set to **MAIN**, which is highlighted with a red circle. Other settings include Name (cr23-5508-1), 802.3x Flow Control Mode (Disabled), LAG Mode on next reboot (Enabled), Broadcast Forwarding (Disabled), AP Multicast Mode (Multicast), AP Fallback (Enabled), Fast SSID change (Disabled), RF Group Name (Enterprise), and User Idle Timeout (300 seconds).

| Parameter                    | Value       |
|------------------------------|-------------|
| Name                         | cr23-5508-1 |
| 802.3x Flow Control Mode     | Disabled    |
| LAG Mode on next reboot      | Enabled     |
| Broadcast Forwarding         | Disabled    |
| AP Multicast Mode            | Multicast   |
| AP Fallback                  | Enabled     |
| Fast SSID change             | Disabled    |
| Default Mobility Domain Name | MAIN        |
| RF Group Name                | Enterprise  |
| User Idle Timeout (seconds)  | 300         |

The default Mobility Domain Name is automatically entered in the Mobility Group membership for that controller, along with the necessary IP address and MAC address information for that controller. The IP address and MAC address information of other controllers in that Mobility Group must be entered manually.

Figure 4-27 and Figure 4-28 show the Mobility Group membership information for both main site WLAN controllers. It can be seen that the Mobility Group membership has two main members for the two WLAN controllers that are providing WLAN access within the main site. These WLAN controllers are also members of another Mobility Group GUEST\_ACCESS. This Mobility Group has been configured to provide guest access tunneling and is discussed later in this chapter.

The remote site WLAN controller Mobility Group membership configuration uses a different mobility group name, and does not include either of the main site WLAN controllers. The reason for it not including either of the main site WLAN controllers is because it is not expecting to support seamless roaming between the remote site and main site. There is no point of providing seamless roaming between controllers when there is no seamless WLAN coverage between APs connected to those controllers. Because this design includes supporting guest access tunneling for users at the remote site, the GUEST\_ACCESS mobility group-member information also appears on the remote site WLAN controller.

**Figure 4-27** *cr23-5508-1 Mobility Group Members*

The screenshot shows the **Static Mobility Group Members** page for controller **cr23-5508-1**. The table lists the following members:

| Local Mobility Group: MAIN |              |              |              |        |
|----------------------------|--------------|--------------|--------------|--------|
| MAC Address                | IP Address   | Group Name   | Multicast IP | Status |
| 00:24:97:cf:3f:a0          | 10.125.30.2  | MAIN         | 0.0.0.0      | Up     |
| 00:24:97:cf:3e:a0          | 10.125.32.34 | GUEST_ACCESS | 0.0.0.0      | Up     |
| 00:24:97:cf:48:60          | 10.125.30.3  | MAIN         | 0.0.0.0      | Up     |

**Figure 4-28** *cr23-5508-2 Mobility Group Members*

| Local Mobility Group | MAC Address       | IP Address   | Group Name   | Multicast IP | Status |
|----------------------|-------------------|--------------|--------------|--------------|--------|
| MAIN                 | 00:24:97:cf:48:60 | 10.125.30.3  | MAIN         | 0.0.0.0      | Up     |
|                      | 00:24:97:cf:3e:a0 | 10.125.32.34 | GUEST_ACCESS | 0.0.0.0      | Up     |
|                      | 00:24:97:cf:3f:a0 | 10.125.30.2  | MAIN         | 0.0.0.0      | Up     |

229004

## WLAN Configuration

### Staff Data WLAN

Figure 4-29 shows the general WLAN configuration for the staff data WLAN network. The key point to note on this tab is the security policy that has been set under the security tab, and the WLAN controller interface that the WLAN has been mapped to. The security configuration recommended is to use WPA2 with 802.1X+CCKM. Most WLAN clients should now support WPA2, and CCKM has been added to 802.1X as it provides faster roaming for WLAN clients that support CCKM, while using the AAA features of 802.1X to secure the WLAN connection.

**Figure 4-29** *Staff Data WLAN*

229398

Apart from setting DHCP as required in the advanced settings, the remainder of the WLAN configuration uses default settings. Unless static IP address are needed, obtaining IP addresses using DHCP is recommended as a best practice.

## Staff Voice WLAN

Figure 4-30 shows the general WLAN configuration for the Staff VoWLAN network. The key point to note on this tab is the security policy that has been set under the security tab, and the WLAN controller interface that the WLAN has been mapped to. The security configuration recommended is to use WPA with CCKM. The VoWLAN clients (7921 and 7925) support WPA and CCKM. CCKM provides optimal roaming performance for voice calls, and the level of security provided by Enterprise WPA is sufficient for VoWLAN traffic. The radio policy of this WLAN is to use the 5GHz (802.11a) band for VoWLAN support, in order to ensure optimal VoWLAN capacity and performance.

The QoS requirements for the WLAN are that it be set for the platinum profile and that WMM be required. Apart from the QoS differences, the remainder of the WLAN configuration is the same as the “Staff Data WLAN” section on page 4-45.

**Figure 4-30 Staff Voice WLAN**

The screenshot displays the Cisco WLAN configuration interface. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh. The main menu on the left shows WLANs, with the Advanced tab selected. The central panel is titled 'WLANs > Edit' and features tabs for General, Security, QoS, and Advanced. The Security tab is active, showing the following configuration:

- Profile Name: Staff VoWLAN
- Type: WLAN
- SSID: vowlan
- Status: ☒ Enabled
- Security Policies: [WPA][Auth(CCKM)] (Modifications done under security tab will appear after applying the changes.)
- Radio Policy: 802.11a only
- Interface: staff voice
- Broadcast SSID: ☒ Enabled

Red circles highlight the Security Policies and the Radio Policy/Interface settings. The bottom right corner of the interface shows the number 229399.

## Guest Access WLAN

Although the configuration for the Guest WLAN indicates that it has been assigned to the management interface, the true interface used by the Guest WLAN is on the anchor WLAN controller that is located in the DMZ. The WLAN client traffic from the Guest WLAN is tunneled by the WLAN controller to the anchor WLAN.

**Figure 4-31 Guest WLAN**

WLANs > Edit

General Security QoS Advanced

Profile Name: Guest\_Access  
 Type: WLAN  
 SSID: Guest\_Access  
 Status: ☒ Enabled

Security Policies: **Web-Auth**  
 (Modifications done under security tab will appear after applying the changes.)

Radio Policy: All  
 Interface: management  
 Broadcast SSID: ☒ Enabled

Figure 4-32 and Figure 4-33 show the first steps in configuring Guest Access Tunneling for the WLAN, namely, the creation of a mobility anchor for the Guest WLAN. The address chosen for the mobility anchor is the management address of the anchor WLAN controller that is located in the DMZ.

**Figure 4-32 WLAN Mobility Anchor Selection**

WLANs > Edit

General Security QoS Advanced

Profile Name: Guest\_Access  
 Type: WLAN  
 SSID: Guest\_Access  
 Status: ☒ Enabled

Security Policies: **Web-Auth**  
 (Modifications done under security tab will appear after applying the changes.)

Radio Policy: All  
 Interface: management  
 Broadcast SSID: ☒ Enabled

**Figure 4-33** Mobility Anchor Selection

The screenshot shows the 'Mobility Anchors' configuration page in the Cisco WLC GUI. The 'WLAN SSID' is 'Guest\_Access'. The 'Switch IP Address (Anchor)' is highlighted with a red circle and contains the value '10.125.32.34'. The 'Data Path' and 'Control Path' are both set to 'up'. There is a 'Mobility Anchor Create' button and a dropdown for 'Switch IP Address (Anchor)' set to '(local)'.

229012

Figure 4-34 shows the DMZ anchor Guest WLAN configuration. The WLAN configuration must be exactly the same as the home controller, except that it has a real local interface, and shown in Figure 4-34 and Figure 4-35.

**Figure 4-34** Anchor Guest WLAN

The screenshot shows the 'WLANs > Edit' configuration page for the 'Guest\_Access' profile. The 'Web-Auth' checkbox is highlighted with a red circle. The 'Radio Policy' is set to 'All' and the 'Interface' is set to 'guest-vlan'. The 'Broadcast SSID' checkbox is also checked.

229013

**Figure 4-35** Anchor WLAN Controller Interfaces

The screenshot shows the 'Controller' configuration page with the 'Interfaces' tab selected. A table lists the configured interfaces. The 'quest-vlan' interface is highlighted with a red circle.

| Interface Name               | VLAN Identifier | IP Address    | Interface Type | Dynamic AP Management |
|------------------------------|-----------------|---------------|----------------|-----------------------|
| <a href="#">quest-vlan</a>   | 104             | 10.125.32.66  | Dynamic        | Disabled              |
| <a href="#">management</a>   | 102             | 10.125.32.34  | Static         | Enabled               |
| <a href="#">service-port</a> | N/A             | 172.26.137.82 | Static         | Not Supported         |
| <a href="#">virtual</a>      | N/A             | 1.1.1.1       | Static         | Not Supported         |

229014

The WLAN on the DMZ anchor WLAN controller must also be configured with a mobility anchor, but in this case the Mobility Anchor is its own local management address, as shown in Figure 4-36.

**Figure 4-36** Anchor Guest WLAN Mobility Anchor

Save Configuration | Ping | Logout | Refresh

WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Mobility Anchors < Back

WLAN SSID Guest\_Access

| Switch IP Address (Anchor) | Data Path | Control Path |
|----------------------------|-----------|--------------|
| local                      | up        | up           |

Mobility Anchor Create

Switch IP Address (Anchor) 10.124.2.66

229015

## WLAN QoS

The Cisco Unified Wireless Network (CUWN) prioritizes traffic based on the QoS profiles applied to each WLAN, but it does not change the IP QoS classification (DSCP) of the client traffic. This means that client traffic leaving the CUWN may need to be reclassified based on the network QoS policy. There are two ways to achieve this reclassification:

1. Applying policy at each of the network SVIs that connect the WLAN controller to the network.
2. Learning the QoS policy that was applied within the CUWN, because this should be aligned with the network policy.

The latter method is preferable as it requires less configuration and less policy maintenance (the policy only needs to be maintained on WLAN controllers and not on the connected switches as well). To achieve this, each of the four QoS profiles (platinum, gold, silver and bronze) on the WLAN controller must have its Wired QoS Protocol Type set to 802.1p. All other QoS profile settings can remain at the defaults (an example is shown in Figure 4-37). This procedure configures the WLAN controller to set the 802.1p marking of the frames sent from the WLAN controller to reflect QoS policy for that WLAN. For example, if the IP packet was from a platinum WLAN and had a DSCP value of EF, the WLAN controller would use a CoS of 5 in the frame header. If the same packet had been on a silver WLAN, the CoS value assigned would be 0. Therefore, as long as the WLAN controller is connected to a switch network that is configured to trust CoS and maintain a translation table between CoS and DSCP for its network, the translation between CUWN policy and network policy will occur automatically.

For more 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 4-37** Wired QoS Protocol Configuration

The screenshot displays the Cisco Unified Wireless Network configuration interface. The left sidebar shows a tree view with categories like General, Network, and QoS. The main content area is titled 'Edit QoS Profile' and contains several sections: 'QoS Profile Name' (platinum), 'Description' (For Voice Applications), 'Per-User Bandwidth Contracts (k) \*' (with fields for Average Data Rate, Burst Data Rate, Average Real-Time Rate, and Burst Real-Time Rate, all set to 0), 'Over the Air QoS' (with fields for Maximum RF usage per AP (%) and Queue Depth, both set to 100), and 'Wired QoS Protocol' (highlighted with a red box). The 'Wired QoS Protocol' section includes 'Protocol Type' (802.1p) and '802.1p Tag' (6). At the bottom, a note states: '\* The value zero (0) indicates the feature is disabled'.

## Access Point Configuration

The configuration and software management of Cisco Unified Wireless Network access points is determined by the WLAN controller they ultimately join. Therefore, establishing the connection between APs and the correct WLAN controller is a key component of the design.

The CUWN provides many different options to allow APs to discover the correct WLAN controller (DHCP, DNS, over the air, or static configuration). These are detailed in the *Deploying Cisco 440X Series Wireless LAN Controllers* document at the following URL:

<http://www.cisco.com/en/US/partner/docs/wireless/technology/controller/deployment/guide/dep.html>

For the purposes of testing in this design, the APs used DHCP to discover a WLAN controller appropriate for their location. The configuration of DHCP for APs is discussed in the *DHCP OPTION 43 for Lightweight Cisco Aironet Access Points Configuration Example* document at the following URL:

[http://www.cisco.com/en/US/partner/tech/tk722/tk809/technologies\\_configuration\\_example09186a00808714fe.shtml](http://www.cisco.com/en/US/partner/tech/tk722/tk809/technologies_configuration_example09186a00808714fe.shtml)

Once an AP is in communication with a WLAN controller that has been defined using a discovery mechanism, it learns about all of the WLAN controllers in the default mobility group of the discovered WLAN controller. An AP can be configured for preferred primary, secondary, and tertiary WLAN controllers within that mobility group. Figure 4-38 shows an example of this where the AP is configured with its preferred WLAN controller (primary controller), and its preferred fail over WLAN controller (secondary controller).



**Figure 4-38 AP Controller Preferences**

The screenshot shows the Cisco Wireless Controller configuration interface. The left sidebar lists 'Wireless' > 'Access Points' > 'All APs' > 'Radios' > '802.11a/n' > '802.11b/g/n' > 'Global Configuration'. The main content area is titled 'All APs > Details for cr22-1142-1-LB'. The 'General' tab is selected, showing a table with columns 'Name' and 'Management IP Address'. The 'Primary Controller' is 'cr23-5508-1' and the 'Secondary Controller' is 'cr23-5508-2'. A red oval highlights these two rows. The 'Tertiary Controller' field is empty. The 'Apply' button is visible in the top right.

|                      | Name        | Management IP Address |
|----------------------|-------------|-----------------------|
| Primary Controller   | cr23-5508-1 |                       |
| Secondary Controller | cr23-5508-2 |                       |
| Tertiary Controller  |             |                       |

The configuration of access point WLAN controller preference will determine the fail over models for the WLAN deployment. For example, all the APs on the site could be configured to prefer one WLAN controller as primary, with the other WLAN controller used solely as a back-up controller. An alternative configuration would be to spread the AP load across both WLAN controllers, on a per building basis, thereby ensuring that all controllers are actively engaged in passing traffic. The advantage of this approach is that a developing controller failure would potentially be discovered more readily if both controllers were always actively carrying some degree of traffic load, rather than with one of them sitting idle.

In situations where the APs are expected to fail over to a WLAN controller outside of its primary WLAN controllers mobility group, the AP must be configured with the IP address and name of that fail over WLAN controller, rather than just the WLAN controller name. An example of this configuration, from the remote site, is shown in Figure 4-39.

**Figure 4-39 AP Failover to a WLAN Controller Outside the Mobility Group**

The screenshot shows the Cisco Wireless Controller configuration interface for a remote site. The left sidebar lists 'Wireless' > 'Access Points' > 'All APs' > 'Radios' > '802.11a/n' > '802.11b/g/n' > 'Global Configuration'. The main content area is titled 'All APs > Details for cr14-1252-1-RSC'. The 'General' tab is selected, showing a table with columns 'Name' and 'Management IP Address'. The 'Primary Controller' is 'cr14-5508-1' and the 'Secondary Controller' is 'cr23-5508-2' with a Management IP Address of '10.125.30.3'. A red oval highlights these two rows. The 'Tertiary Controller' field is empty. The 'Apply' button is visible in the top right.

|                      | Name        | Management IP Address |
|----------------------|-------------|-----------------------|
| Primary Controller   | cr14-5508-1 |                       |
| Secondary Controller | cr23-5508-2 | 10.125.30.3           |
| Tertiary Controller  |             |                       |

## AP 1520 Configuration

AP1520 access points require somewhat further configuration over and above what has been shown in the preceding paragraphs. By default, AP1520 access points are configured for outdoor mesh operation, and in order to use these access points to provide outdoor coverage as root access points, some basic configuration changes must be implemented.

## Adding the AP1520 MAC Address to the WLAN Controller

AP1520 series access points will not join a WLAN controller unless the MAC address of the access point has been defined to the WLAN controller. This can be done by adding the BVI MAC of the access point (this is the MAC address printed on a label on the outside of the access point) via the **Security > AAA > MAC Filtering** GUI panel, as shown in Figure 4-40.

**Figure 4-40 Adding the AP1520 MAC Address to the WLAN Controller**

| MAC Address       | Profile Name | Interface  | IP Address | Description             |
|-------------------|--------------|------------|------------|-------------------------|
| 00:24:50:36:9a:00 | Any WLAN     | management | unknown    | AP1522 on cr24-3750-MB  |
| 00:24:50:36:b6:00 | Any WLAN     | management | unknown    | AP1522 on cr22-4507-LB  |
| 00:24:50:36:b9:00 | Any WLAN     | management | unknown    | AP1522 on cr14-4507-RSC |
| 00:24:50:36:c2:00 | Any WLAN     | management | unknown    | AP1522 on cr14-3750s-SB |

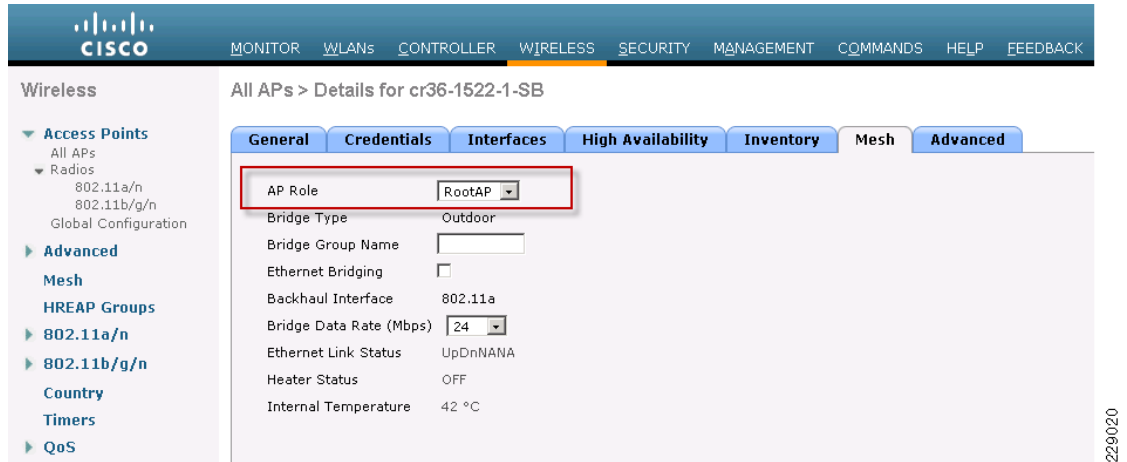
Note that MAC addresses must be defined to all WLAN controllers that an AP1520 access point may join. This includes not only WLAN controllers defined as primary controllers, but any WLAN controllers that are defined as secondary or tertiary as well.

You can also validate the MAC addresses of AP1520 access points externally using Cisco ACS. For complete details on how to do this, refer to the *Cisco Wireless Mesh Access Points Design and Deployment Guide*, Release 6.0 at the following URL:  
[http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP\\_60.html#wp1194149](http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP_60.html#wp1194149)

## Configuring the AP1520 as a Root Access Point (RAP)

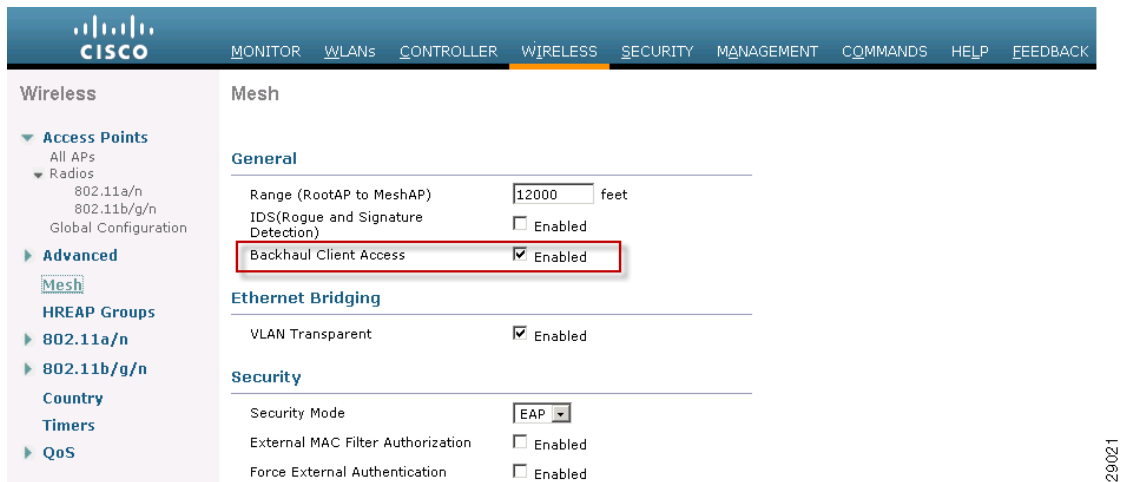
AP1520 series access points are shipped with a default outdoor Mesh Access Point (MAP) configuration. In the Medium Enterprise Design Profile mobility design, the AP1520 series access point is used as an outdoor root access point (RAP)<sup>1</sup>. In order to reconfigure the AP1520 to be a RAP, once the access point has joined the controller, the AP role is changed to “RootAP” in the **Wireless > Access Points > All APs > Details > Mesh** configuration panel on the WLAN controller, as shown in Figure 4-41. None of the other parameters need to be changed on this screen.

1. MAPs and RAPs are explained in much more detail in the *Cisco Wireless Mesh Access Points Design and Deployment Guide*, Release 6.0 at the following URL:  
[http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP\\_60.html#wp1194149](http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP_60.html#wp1194149)

**Figure 4-41**      **Setting the AP Role**

## 5 GHz Backhaul Client Access

By default, the 5 GHz radio interface on the AP1522 is enabled only as a back haul interface, and will not allow any 5 GHz clients to associate. In order to enable the use of this interface for 5 GHz client traffic, it must be enabled using the Back haul Client Access check box on the WLAN controller's **Wireless > Mesh** configuration panel, as shown in [Figure 4-42](#). Enabling this once on the WLAN controller enables back haul client access for all AP1520 series access points that join this controller.

**Figure 4-42**      **Enabling Backhaul Client Access**

## Primary Backhaul Scanning

Under normal circumstances, an AP1520 configured as a root AP (RAP) communicates with the WLAN controller via its wired Ethernet interface. However, if the Ethernet port is “down” on a RAP, or a RAP fails to connect to a controller when its Ethernet port is “up”, the AP1520 will attempt to use the 5 GHz radio interface as the primary backhaul for 15 minutes. Failing to find another AP1520 neighbor or failing to connect to a WLAN controller via the 5 GHz interface causes the AP1520 to begin to scan for reassignment of the primary backhaul, beginning with the Ethernet interface.

In most cases we did not find this behavior to cause any issues in our validation and we recommend that it be left as is. We found this behavior beneficial in that should a switch port for an AP 1520 series access point go down, the connected AP1520 can establish a connection to another AP1520 in the same building or at an adjacent building using the 5 GHz backhaul. This can be especially useful if the neighbor AP1520 is attached to the wired network via a different Ethernet switch. Within 15 minutes of the failed Ethernet port being repaired, the AP1520 should revert back to operation over the Ethernet connection.

If you do not wish to allow primary backhaul scanning, you may either:

- Disable the use of 5 GHz entirely on the AP1520 series access point. In this case, backhaul operation will not occur over any wireless medium (2.4 GHz is never used for backhaul purposes by the AP1520). This is an acceptable alternative if there is no need to support 5 GHz clients within the outdoor perimeter of the buildings where AP1520s are installed.
- Use AP 1250 access points installed within traditional weatherproof outdoor NEMA-rated enclosures (supplied by Cisco partners) to provide outdoor coverage.

## WCS Configuration

Configuring WCS to allow basic management of WLAN controllers for each site in the Medium Enterprise Design Profile mobility design is a relatively straightforward process. After installing WCS in the main site, each WLAN controller must be added to WCS, as described in the *Cisco Wireless Control System Configuration Guide* at the following URL:

[http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6\\_0ctrlcfg.html#wp1041451](http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0ctrlcfg.html#wp1041451)

Once the WLAN controllers are properly defined and reachable from WCS, the network administrator can begin to use the multitude of configuration, monitoring, and reporting options available under the WCS to begin to manage not only the WLAN controllers themselves, but the access points and devices that connect through them. These capabilities are far too numerous to be described here, but a comprehensive description of these capabilities and how to enable them can be found in the *Cisco Wireless Control System Configuration Guide*, at the above URL.

## WCS Users and User Groups

By default, WCS provides for a single root user, which allows access to all WCS functions. The password for this root user should be protected and only known by those who are responsible for the overall Medium Enterprise Design Profile mobility design and with a real need to know (for example, those personnel responsible for the actual installation, maintenance, and detailed administration of WCS). For these users and others who require routine administrative access to WCS, alternate user credentials should be created, with administrative access granted and privileges assigned as necessary via the use of appropriate WCS user groups settings. 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](http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0manag.html)) provides comprehensive instructions for configuring users and group privileges on the 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 used when assigning user privileges. 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 users who are only interested in viewing the information available to them on WCS will not need more than the ability

to simply monitor network activity in WCS. For these users, the privileges accorded to 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.

## WCS Virtual Domains

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 an assigned virtual domain, users are only able to use WCS functionality against a predefined 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 the WCS that are not contained in their active assigned virtual domain. For example, the site manager for a medium enterprise may have the ability to view and report on certain characteristics of wireless assets for his site due to his WCS user account being assigned to an appropriate user group permitting this level of WCS functionality. However, the virtual domain that this site manager is assigned to may only allow such functionality to be exercised against these assets if they are located within his site. Thus, if the site manager for site “A” attempted to use WCS to discover or manage wireless infrastructure located in site “B”, his assigned virtual domain might not allow the ability to manage or even view resources on site “B”.

Administrative personnel with system-wide responsibilities, on the other hand, could be assigned a virtual domain that includes all resources in the system (i.e., all sites), and could exercise the functionality assigned to them by their WCS user group against any of these resources. In this way, WCS virtual domain assignment can be useful in 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 we must assign the resources we wish to include in the virtual domain. [Figure 4-43](#) provides an illustration of how controller resources were assigned during lab testing for the “main site” virtual domain.

**Figure 4-43** Assigning WLAN Controller Resources to the Main Site Virtual Domain

**Virtual Domains**  
Administration > Virtual Domains

Virtual Domain Name: main site  
Description: main site

**Summary** | **Maps** | **Controllers** | **Access Points**

**Available Controllers**

- cr14-5508-1

**Selected Controllers**

- cr23-5508-1
- cr11-5508-wlc
- cr23-5508-2

Add > < Remove

Submit Cancel

**Footnotes**

1. Manage each controller from only one Virtual Domain at a time. If a controller's configuration is modified by multiple Virtual Domains, complications may arise.
2. Adding a controller to a Virtual Domain adds all the associated APs to that Virtual Domain automatically.
3. Adding a map to a Virtual Domain adds all the associated APs to that Virtual Domain automatically.
4. Associate each Virtual Domain to users by going to Administration->AAA->Users. [click here to view Users](#)

229400

The process for creating and assigning network resources to the virtual domain is detailed in “Chapter 20, “Virtual Domains” of the *WCS Configuration Guide*, Release 6.0, found at the following URL:

[http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6\\_0virtual.html#wp1040002](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 the main site virtual domain to the “main1” user is shown in [Figure 4-44](#). This process is detailed in a step-by-step fashion in “Chapter 7, Managing WCS User Accounts” at the following URL:

[http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6\\_0manag.html#wp1097733](http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0manag.html#wp1097733)

**Figure 4-44** Assigning the Virtual Domain to a User

**User Detail :main1**  
Administration > AAA > Users > User Detail

**General** | **Virtual Domains**

**Available Virtual Domains**

- root
- remote
- small site

**Selected Virtual Domains**

- main site

Add > < Remove

Submit Cancel

**Footnotes:**

1. Click [here](#) for current password policy.
2. If user belongs to 'LobbyAmbassador' or 'Monitor Lite' or 'North Bound API' or 'Users Assistant' group then he cannot belong to any other group.
3. Root group is only assignable to 'root' user and that assignment cannot be changed.
4. 'root' Virtual Domain cannot be removed from Selected Virtual Domains for 'root' user.

229401

**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. 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](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.

Additional information on creating WCS users, user groups, and virtual domains can be found in the “Context-Aware Service Design” chapter of the *Cisco Service Ready Architecture for Schools Design Guide* at the following URL:

[http://cisco.com/en/US/docs/solutions/Enterprise/Education/SchoolsSRA\\_DG/SchoolsSRA\\_chap6.html#wp1054537](http://cisco.com/en/US/docs/solutions/Enterprise/Education/SchoolsSRA_DG/SchoolsSRA_chap6.html#wp1054537)

## Reference Documents

A cornerstone of a successful design relies on the knowledge of established best practices. Thus, it is highly recommended that you become familiar with the following general best practice deployment recommendations for Cisco Unified Wireless Networks:

- *Enterprise Mobility Design Guide 4.1*  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>
- *Cisco 802.11n Design and Deployment Guidelines*  
[http://www.cisco.com/en/US/docs/solutions/collateral/ns340/ns394/ns348/ns767/white\\_paper\\_80211n\\_design\\_and\\_deployment\\_guidelines.html](http://www.cisco.com/en/US/docs/solutions/collateral/ns340/ns394/ns348/ns767/white_paper_80211n_design_and_deployment_guidelines.html)
- *Voice over Wireless LAN 4.1 Design Guide*  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>
- *Cisco Radio Resource Management*  
[http://www.cisco.com/en/US/tech/tk722/tk809/technologies\\_tech\\_note09186a008072c759.shtml](http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml)
- *Cisco Wireless Mesh Access Point Design and Deployment Guide*, Release 6.0  
[http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP\\_60.html](http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP_60.html)

A successful deployment also involves strong knowledge of how to set key infrastructure configuration procedures. The following documents provide comprehensive configuration guidance and should be referenced as needed:

- *Cisco Wireless LAN Controller Configuration Guide*, Release 6.0  
<http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.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>

Additional product information on the Cisco wireless infrastructure discussed in this chapter can be found at the following locations:

- *Cisco 5508 Wireless Controller*  
[http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/data\\_sheet\\_c78-521631.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/data_sheet_c78-521631.html)
- *Cisco 1140 Series 802.11n Access Point*  
<http://www.cisco.com/en/US/products/ps10092/index.html>
- *Cisco 1250 Series 802.11n Access Point*  
<http://www.cisco.com/en/US/products/ps8382/index.html>
- *Cisco 1250 Series Antenna Options*  
[http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/at\\_a\\_glance\\_c45-513837.pdf](http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/at_a_glance_c45-513837.pdf)
- *Cisco Aironet 1520 Lightweight Outdoor Access Point Ordering Guide*  
[http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps8368/product\\_data\\_sheet0900aecd8066a157.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps8368/product_data_sheet0900aecd8066a157.html)
- *Cisco Wireless Control System (WCS)*  
[http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product\\_data\\_sheet0900aecd802570d0.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html)
- *Cisco Wireless Control System Virtual Domains*  
[http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/brochure\\_c02-474335.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/brochure_c02-474335.html)
- *Cisco Wireless Control System Navigator*  
<http://www.cisco.com/en/US/products/ps7305/index.html>





## CHAPTER 5

# Medium Enterprise Design Profile (MEDP)—Network Security Design

## Security Design

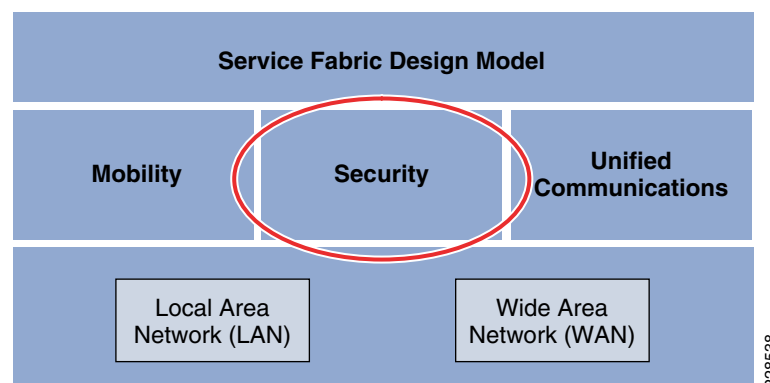
As medium enterprises embrace new communication and collaboration tools, transitioning to more Internet-based, media-rich applications, a whole new set of network security challenges arise. Medium enterprise network infrastructures must be adequately secured to protect employees from harmful content, to guarantee confidentiality of private data, and to ensure the availability and integrity of the systems and data. Providing a safe and secure network environment is a top responsibility for enterprise administrators.

This chapter describes how the Medium 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 Medium Enterprise Design Profile is a well-designed and validated network reference design that enables medium enterprises to deliver all of the services required for an enhanced business environment.

Within the Cisco Medium Enterprise Design Profile, the service fabric network provides the foundation on which all solutions and services are built to solve the business challenges facing medium enterprises. These business challenges include borderless access, secure mobile workforce, operational efficiencies, and user experience.

The network fabric consists of four distinct components: LAN/WAN, security, mobility, and unified communications, as shown in [Figure 5-1](#).

**Figure 5-1**      **Service Fabric Design Model**

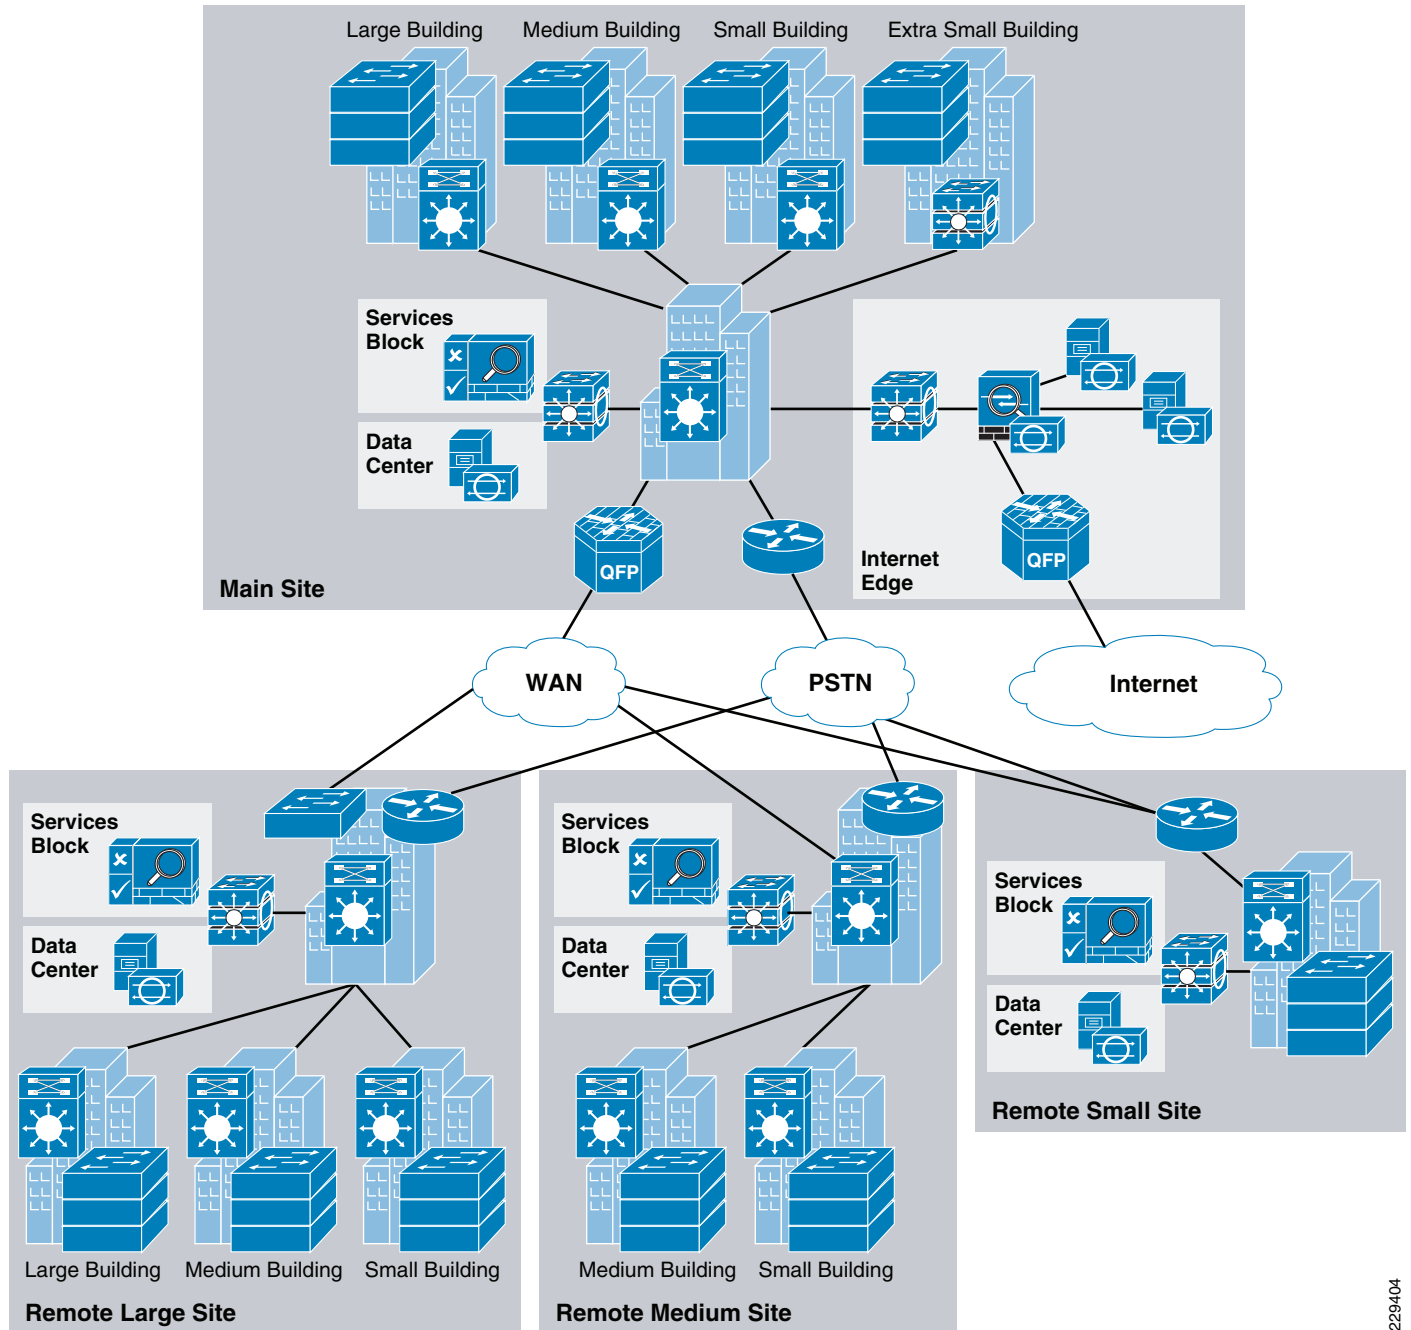


The Medium Enterprise Design Profile includes security to protect the infrastructure and its services to provide a safe and secure online business environment. This design profile leverages the proven design and deployment guidelines of the Cisco SAFE Security Reference Architecture to secure the service fabric by deploying security technologies throughout the entire network solution to protect employees from harmful non-business content, to guarantee the confidentiality of the enterprise and employees private data, and to ensure the availability and integrity of the systems and data.

Protecting the infrastructure and its services requires implementation of security controls capable of mitigating both well-known and new forms of threats. Common threats to enterprise environments include the following:

- 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, IP spoofing, and unauthorized access to restricted resources
- Data loss—Loss or leakage of private data from servers and endpoints while in transit or as a result of spyware, malware, key-loggers, viruses, and so on
- Identity theft and fraud—Theft of personal identity or fraud on servers and end users through phishing and E-mail spam

The Medium Enterprise Design Profile accommodates a main site and one or more remote sites of various sizes, interconnected over a metro Ethernet or managed WAN service. Each of these sites may contain one or more buildings of varying sizes, as shown in [Figure 5-2](#).

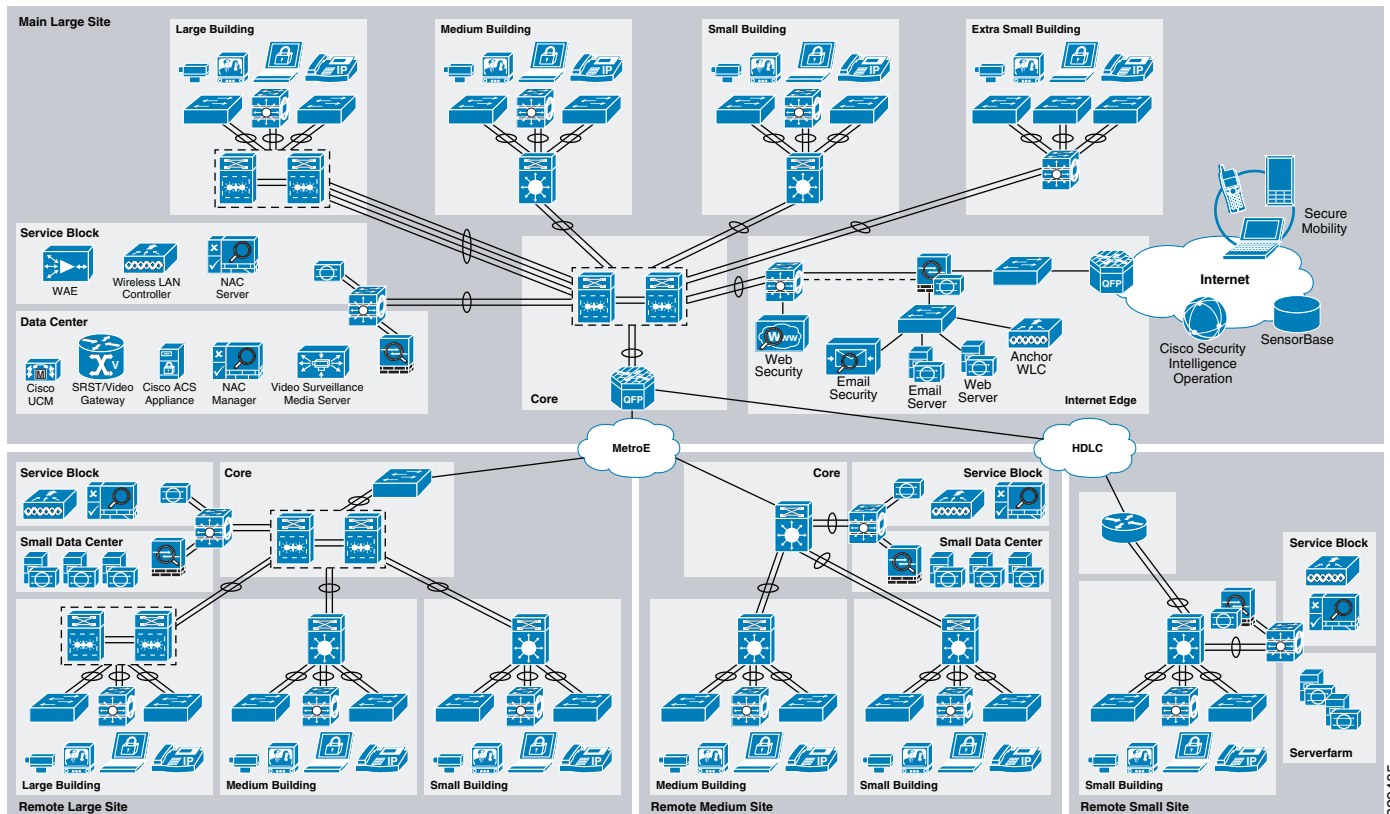
**Figure 5-2** Medium Enterprise Design Profile Network Overview

229404

Operating on top of this network are all the services used within the enterprise environment, such as safety and security systems, voice communications, business databases, ordering systems, payroll, accounting and customer relationship management (CRM) applications, and so on. The core of these 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 various remote locations. These centralized systems and applications are served by a data center at the main site.

The security design uses a defense-in-depth approach where multiple layers of security protection are integrated into the architecture. Various security products and technologies are combined to provide enhanced security visibility and control, as shown in Figure 5-3.

**Figure 5-3** Medium Enterprise Design Profile Network Security Design Overview



The following security elements may be included in the Medium Enterprise Security design as shown in Figure 5-3:

- *Endpoint security*—Desktop and server endpoint protection for day-zero attack protection, data loss prevention, and signature-based antivirus
- *Network foundation protection (NFP)*—Device hardening, control plane, and management plane protection throughout the entire infrastructure to maximize availability and resiliency
- *Catalyst Integrated Security Features*—Access layer protection provided by port security, Dynamic ARP inspection, IP Source Guard, and DHCP Snooping
- *Threat detection and mitigation*—Intrusion prevention and infrastructure based network telemetry to identify and mitigate threats
- *Internet access*—E-mail and web security, stateful firewall inspection, intrusion prevention and global correlation, and granular Internet access control
- *Cisco Video Surveillance*—Monitor activities throughout the enterprise environment to prevent and deter safety incidents
- *Enhanced availability and resiliency*—Hardened devices and high availability design to ensure optimal service availability, and system and interface-based redundancy

- *Unified communications*—Security and emergency services, enhanced 911 support, conferencing and collaboration for planning and emergency response
- *Network access control*—Authentication and policy enforcement via Cisco Identity-Based Networking Services (IBNS), and role-based access control and device security compliance via Cisco Network Admission Control (NAC) Appliance
- Secure mobility
  - Always-on VPN protection for PC-based and smartphone mobile users
  - Persistent and consistent policy enforcement independently from user location
  - Enforcement of client firewall policies
  - Optimal gateway selection to ensure best connectivity
  - Integration with web security and malware threat defense systems deployed at the enterprise premises
  - Consolidated SaaS access control

The Medium Enterprise Design Profile recognizes that cost and limited resources may be common limiting factors. Therefore, architecture topologies and platforms are carefully selected to increase productivity while minimizing the overall cost and operational complexities. In certain cases, tradeoffs are made to simplify operations and reduce costs where needed.

The Medium Enterprise Design Profile focuses on the following key areas for securing the service fabric within medium enterprises:

- Network foundation protection (NFP)—Ensuring the availability and integrity of the network infrastructure by protecting the control and management planes to prevent service disruptions network abuse, unauthorized access, and data loss
- Internet perimeter protection
  - Ensuring safe connectivity to the Internet
  - Protecting internal resources and users from botnets, malware, viruses, and other malicious software
  - Protecting employees from harmful content
  - Enforcing E-mail and web browsing policies to prevent identity theft and fraud
  - Blocking command and control traffic from infected internal bots to external hosts
- Data center protection
  - Ensuring the availability and integrity of centralized applications and systems
  - Protecting the confidentiality and privacy of proprietary and 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 sites
  - Ensuring that systems are up-to-date and in compliance with the enterprise's 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

Together, these key security areas create a defense-in-depth solution for protecting medium enterprises from common security threats such as service disruption, network abuse, unauthorized access, data loss, and identity theft and fraud. The design guidelines and best practices for each of the above security focus areas are detailed in the following sections. For more detailed information on each of these areas, see the *Cisco SAFE Reference Guide* at the following URL: <http://www.cisco.com/go/safe>.

## Network Foundation Protection

Medium enterprise networks are built with routers, switches, and other infrastructure network devices that keep the applications and services running. These infrastructure devices must be properly hardened and secured to maintain continued operation and access to these services.

To ensure the availability of the medium enterprise network infrastructure, the Medium Enterprise Design Profile leverages the Network Foundation Protection best practices for the following areas:

- Infrastructure device access
  - Restrict management device access to authorized parties and via only 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 such as 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, and 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 (PVST) 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 Cisco Catalyst Infrastructure Security Features (CISF) including port security, Dynamic ARP Inspection, DHCP Snooping, and IP Source Guard.
  - 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 infrastructure.
  - Authenticate, authorize, and keep records of all administrative access.
  - If possible, implement a separate out-of-band (OOB) management network (hardware- or VLAN-based) to manage systems local to the main site.
  - Secure the OOB management access 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 such as SSH and HTTPS.
  - Ensure time synchronization by using NTP.
  - Secure management servers and endpoints with endpoint protection software and operating system (OS) hardening best practices.

For more detailed information on the NFP best practices including configuration examples, see “Chapter 2, Network Foundation Protection” in the *Cisco SAFE Reference Guide* at the following URL: [http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE\\_RG/chap2.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap2.html).

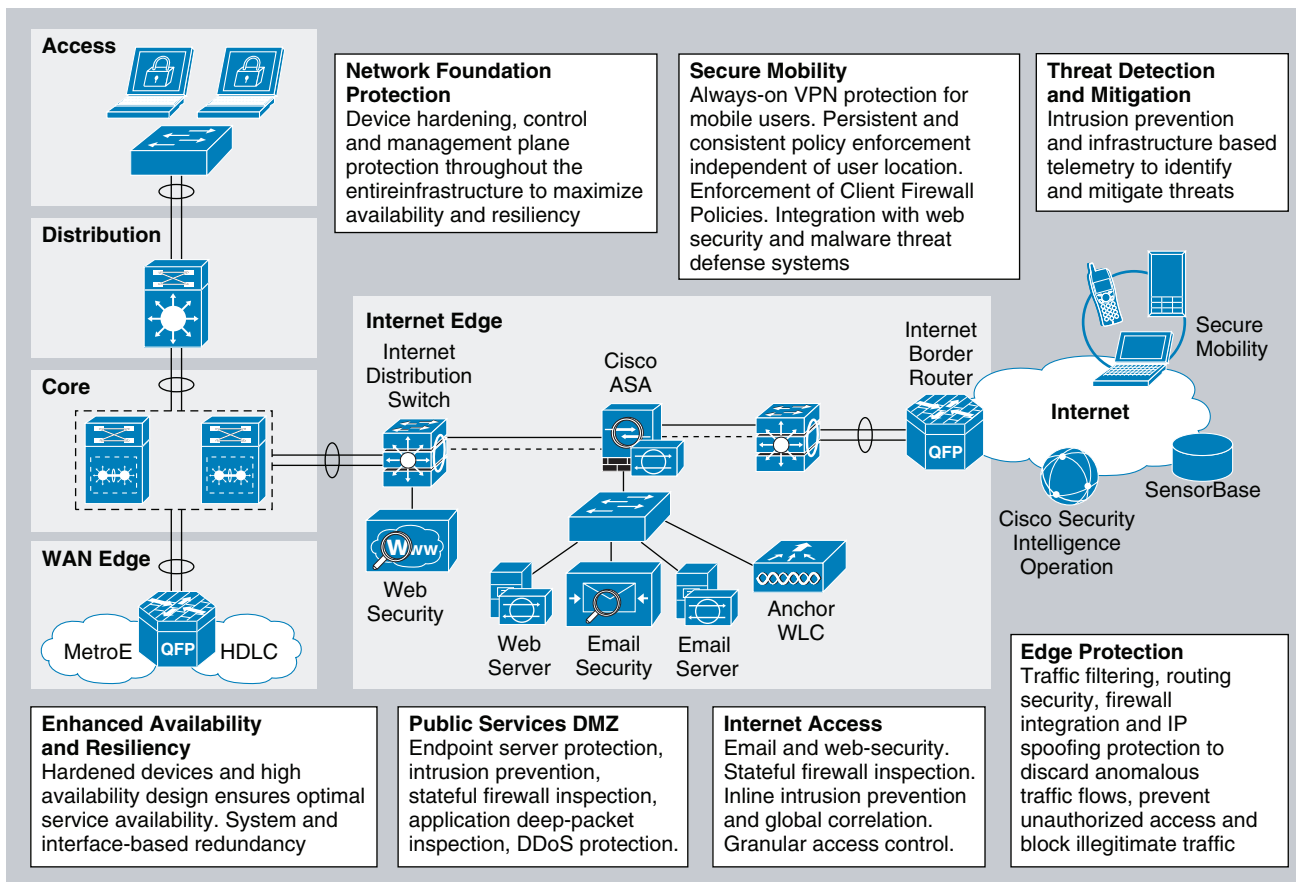
# Internet Perimeter Protection

The Medium Enterprise Design Profile assumes a centralized connection to the Internet at the headquarters or main site. This connection serves users at the main site as well as all remote sites or offices. Common services typically provided by this connection include the following:

- E-mail and Internet web browsing for employees
- Hosting of a company's web portal accessible to clients and partners over the Internet
- Secure remote access to the enterprise network for mobile users and remote workers
- Other services may also be provided using the same infrastructure

The part of the network infrastructure that provides connectivity to the Internet is defined as the Internet perimeter, as shown in Figure 5-4.

**Figure 5-4** Internet Perimeter



The Internet perimeter provides safe and secure centralized access to the Internet for employees and guest users residing at all locations. It also provides access to public services such as the company's web portal and public services without compromising the confidentiality, integrity, and availability of the resources and data of the enterprise.



To provide secure access, the Internet perimeter should incorporate the following security functions:

- *Internet border router*—The Internet border router is the gateway responsible for routing traffic between the enterprise and the Internet. It may be administered by the company's IT staff or may be managed by the Internet service provider. This router provides the first line of protection against external threats and should be hardened according to the NFP best practices.
- *Internet firewall*—A Cisco Adaptive Security Appliance (ASA) provides stateful access control and deep packet inspection to protect enterprise resources and data from unauthorized access and disclosure. 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 ports for rogue activity and detects and blocks traffic from infected internal endpoints, sending command and control traffic back to a host on the Internet. The Cisco ASA is configured to control or prevent incoming access from the Internet, to protect the enterprise web portal and other Internet public services, and to control user traffic bound towards 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 so on.
- *Public services DMZ*—The company's external Internet web portal, mail server, and other public facing servers and services are placed on a demilitarized zone (DMZ) for security and control purposes. The DMZ acts as a middle stage between the Internet and enterprise 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 in the DMZ, and controls outbound access from DMZ resources to the Internet. Systems residing within the DMZ should be hardened with endpoint protection software and OS hardening best practices.
- *E-mail security*—A Cisco IronPort C Series E-Mail Security Appliance (ESA) is deployed in 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. The WSA enforces URL filtering policies to block access to websites containing non-business-related content or that are known sources of spyware, adware, botnets, or other types of malware. The WSA may also be configured to block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and so on.
- *Guest access wireless LAN controller*—The Cisco Unified Wireless LAN Guest Access option offers a flexible, easy-to-implement method for deploying wireless guest access via Ethernet over IP (RFC3378). Ethernet over IP (EoIP) tunneling is used between two wireless LAN controller (WLC) endpoints in the centralized network design. A WLC is located in the Internet perimeter DMZ, where it is referred to as an *anchor controller*. The anchor controller is responsible for terminating EoIP tunnels originating from centralized campus WLCs located in the services block, and interfacing the traffic from these controllers to a firewall or border router. Traffic to and from this guest access WLAN is tunneled to the DMZ transparently, with no visibility by, or interaction with, other traffic in the enterprise internal network. For more information on the wireless guest access solution, see the *Medium Enterprise Design Profile Mobility Design Guide*.

The following subsections describe the design guidelines for implementing the above security functions.

## Internet Border Router Security

The Internet border router provides connectivity to the Internet through one or more Internet service providers. The router acts as the first line of defense against unauthorized access, distributed DoS (DDoS), and other external threats. ACLs, uRPF, and other filtering mechanisms should be implemented for anti-spoofing and to block invalid packets. NetFlow, Syslog, and SNMP should be used to gain visibility on traffic flows, network activity, and system status. In addition, the Internet border router should be hardened and secured following the best practices explained in [Network Foundation Protection, page 5-6](#). 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 Edge ACL Deployment section provides a sample configuration of an Internet Edge ACL. For more information on how to secure the Internet border router, see “Chapter 6, Enterprise Internet Edge” in the *Cisco SAFE Reference Guide* at the following URL: [http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE\\_RG/chap6.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap6.html).

## Internet Firewall

A Cisco ASA firewall should be deployed at the Internet perimeter to protect the enterprise internal resources and data from external threats, and is responsible for the following:

- Preventing incoming access from the Internet
- Protecting public resources deployed in the DMZ by restricting incoming access to the public services and by limiting outbound access from DMZ resources out to the Internet
- Controlling user Internet-bound traffic
- Monitoring network ports for rogue activity and preventing infected internal endpoints from communicating with botnet masters or command and control hosts on the Internet

The Cisco ASA should be configured to enforce access policies, keep track of connection status, and inspect packet payloads. Examples of the needed access policies include the following:

- Deny or control 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; that is, HTTP and HTTPS, and so on.
- Allow outbound SSL access to the Internet for devices requiring administrative updates such as SensorBase, IPS signature updates, and so on.
- Allow user access to DMZ services such as the company’s web portal, E-mail, and domain name resolution (HTTP, HTTPS, Simple Mail Transfer Protocol (SMTP), point-of-presence [POP], Internet Message Access Protocol (IMAP), Domain Name Service [DNS]).
- Restrict inbound Internet access to the DMZ for the necessary protocols and servers (HTTP to web server, SMTP to Mail Transfer Agent, DNS to DNS servers, and so on).
- Restrict connections initiated from the DMZ to only necessary protocols and sources (DNS from DNS server, SMTP from mail server, HTTP/SSL from Cisco IronPort ESA).
- Enable stateful inspection for the outbound protocols being used to ensure returning traffic is dynamically allowed by the firewall.

- Prevent access to the anchor WLC deployed in the DMZ for guest access except for tunneled traffic coming from the centralized campus WLCs (UDP port 16666 and IP protocol ID 97) and traffic needed to manage it (SNMP, TFTP, HTTP, HTTPS, SSH).
- Implement NAT and Port Address Translation (PAT) to shield the internal address space from the Internet.

**Note**

Whenever available, a dedicated management interface should be used. However, in cases where the firewall is managed in-band, identify the protocols and ports required before configuring the firewall ACLs.

When deploying the Internet firewall, it is important to understand the traffic and policy requirements when selecting a firewall. An appropriately-sized Cisco ASA model should be chosen so that it does not become a bottleneck. The Cisco ASA should also be hardened following the NFP best practices as described in [Network Foundation Protection, page 5-6](#). This includes restricting and controlling administrative access, securing dynamic exchange of routing information with MD5 authentication, and enabling firewall network telemetry with SNMP, Syslog, and NetFlow.

In the Medium 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 can be deployed in stateful failover using separate boxes at a higher cost.

## 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 master host on the Internet. The Botnet Traffic Filter on the Cisco 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 Cisco 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 this 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 circumstances:

- 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](http://www.cisco.com/en/US/prod/vpndevc/ps6032/ps6094/ps6120/botnet_index.html).


## Intrusion Prevention

IPS is responsible for identifying and blocking anomalous traffic and packets recognized as well-known attacks. An AIP SSM IPS 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 medium 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 5-5](#)). Traffic identified for IPS inspection cannot continue through the Cisco 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 Cisco ASA to shun the traffic or by resetting a connection on the Cisco ASA. Also, while the AIP SSM is analyzing the traffic, a small amount of traffic might pass through the Cisco ASA before the AIP SSM can shun it. The right side of [Figure 5-5](#) shows the AIP SSM in promiscuous mode.

**Figure 5-5** *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 Cisco ASA, traffic flows through the Cisco ASA and the AIP SSM in the following sequence:

1. Traffic enters the ASA.
2. Firewall policies are applied.
3. Traffic is sent to the AIP SSM over the backplane.
4. The AIP SSM applies its security policy to the traffic and takes appropriate actions.
5. (Inline mode only) Valid traffic is sent back to the ASA over the backplane; the AIP SSM might block some traffic according to its security policy, and that traffic is not passed on.
6. Remote access VPN policies are applied (if configured).
7. Traffic exits the Cisco ASA.

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 allows all traffic through, uninspected, if the AIP SSM becomes unavailable. Conversely, when configured to fail close, the Cisco ASA blocks all traffic in case of an AIP SSM failure.

## Cisco IPS Global Correlation

The AIP SSM module on the Cisco ASA (or IPS appliance) may also participate in Cisco 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. IPS incorporates the global threat data into its system to detect and prevent malicious activity even earlier, allowing IPS to filter out the worst attackers before they have a chance to attack critical assets.

IPS Global Correlation is an important improvement in the basic functions of IPS because it enables it 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 provides 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 adds 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.

After Global Correlation is configured, the IPS works in the following manner. 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. Next, the packet is sent through the Cisco IPS reputation filters. Packets that match are discarded immediately, assuming that the reputation filters are enabled and not in Audit mode. 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. 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. Finally, 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. Packets that match the IPS reputation filters, are dropped before signature inspection.



### 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 reports potential deny actions because of 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, 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 does not send a live request for information about an IP address that it has just seen. It inspects the data that it has, and if it finds the address, it uses that data; otherwise, the sensor assumes 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 is 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 makes an outbound connection to check for the latest updates from Cisco SensorBase. It initiates an HTTPS request to Cisco SensorBase update servers and downloads a manifest that contains the latest versions of the files related to Global Correlation. The sensor checks Cisco SensorBase every five minutes for updates. If changes are needed, the sensor performs a DNS lookup of the server name returned in the initial request. This lookup returns the location of the server nearest to the sensor. The sensor then initiates an HTTP connection that



actually transfers the data. The size of a full update is approximately 2 MB; incremental updates average about 100 KB. If a sensor loses connection to Cisco SensorBase, Global Correlation information begins to time out within days, and sensor health changes 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 as follows:

- *Off (default)*—The sensor does not send back any data. The sensor still receives reputation data, and this setting does not affect its use of that data except that the reputations of addresses attacking the network being protected are not influenced by their generation on the sensor.
- *Partial*—The sensor sends 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.
- *Full*—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. Therefore, all events reported to Cisco SensorBase have any such IP address information stripped from the reported data.

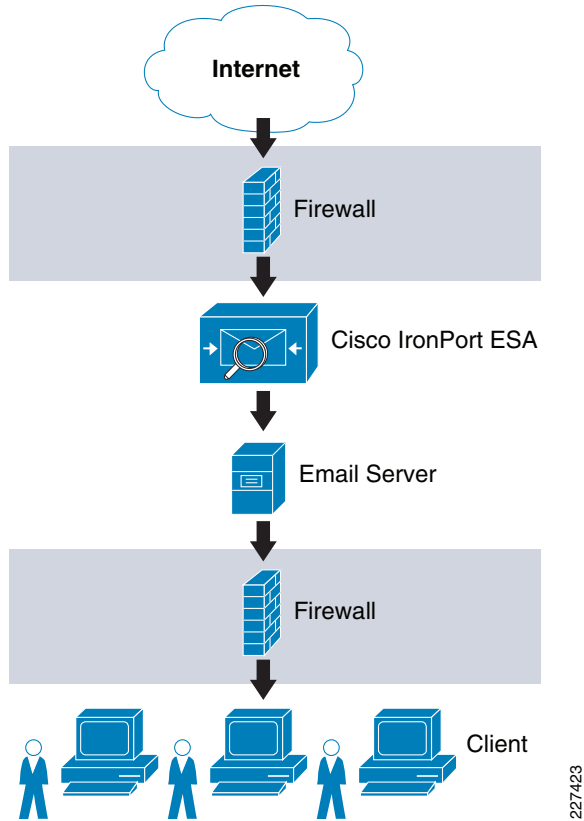
The mechanism used to update Cisco SensorBase with new attack information is fairly straightforward. The sensor takes event information, parses out the important pieces of data, and buffers this data for transmission back to Cisco SensorBase. It sends this data in the form of an HTTPS connection that it initiates on average every 10 minutes. The average size of an update is 2 to 4 KB, with weekly averages of approximately 0.5 to 1 MB. Some higher-volume sensors have average update sizes of approximately 50 KB, with weekly totals in the 45-MB range. Sensors with very high alert volumes can have average update sizes of approximately 850 KB, with weekly totals of up to 900 MB; these sensors, however, 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](http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/cli/cli_collaboration.html).

## E-Mail Security Guidelines

The Medium Enterprise Design Profile implements a Cisco IronPort C Series E-Mail Security Appliance (ESA) in the DMZ to inspect E-mails and prevent threats such as E-mail spam, viruses, and worms. The ESA acts as a firewall and threat monitoring system for SMTP traffic (TCP port 25). Logically, the ESA acts as a Mail Transfer Agent (MTA) within the E-mail delivery chain, as shown in [Figure 5-6](#).

**Figure 5-6 Logical E-Mail Delivery Chain**

227423

**Note**

Figure 5-6 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 a “sandwich” configuration.

When the ESA receives the E-mails, they are evaluated using a reputation score mechanism based on the SensorBase network, which is an extensive network that monitors global E-mail and web traffic for anomalies, viruses, malware, and other abnormal behavior. The SensorBase network consists of Cisco IronPort appliances, Cisco ASA, and IPS appliances installed in more than 100,000 organizations worldwide. This provides a large and diverse sample of Internet traffic patterns. By leveraging the information in 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.

In addition, an enterprise may optionally 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**

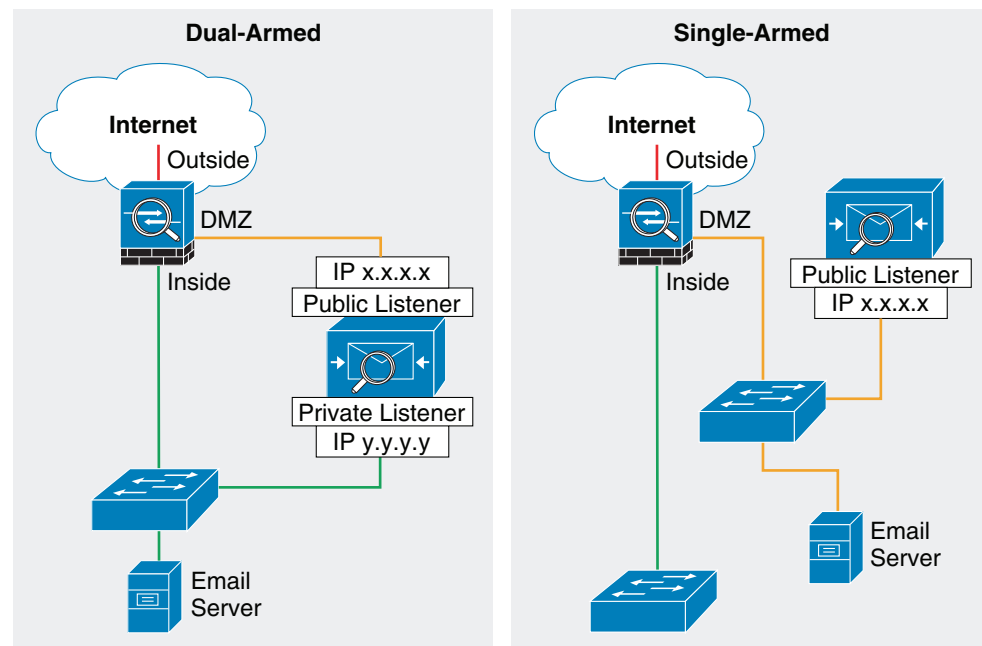
Alternatively, Cisco offers managed hosted and hybrid hosted E-mail security services. These services are provided through a dedicated E-mail infrastructure hosted in a network of Cisco data centers. For more information, see the following URL: <http://www.cisco.com/go/designzone>.

There are two options for deploying the ESA appliance, depending on the number of interfaces used:

- *Dual-armed configuration*—Two physical interfaces are used to serve as a public mail listener and a private mail listener where each interface is configured with a separate logical IP address. The public listener receives E-mail from the Internet and directs messages to the internal mail servers. The private listener receives E-mail from the internal servers and directs messages to the Internet. The public listener interface would connect to the DMZ and the private listener interface can connect to the inside of the firewall closer to the mail server.
- *One-armed configuration*—A single interface is configured on the ESA 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.

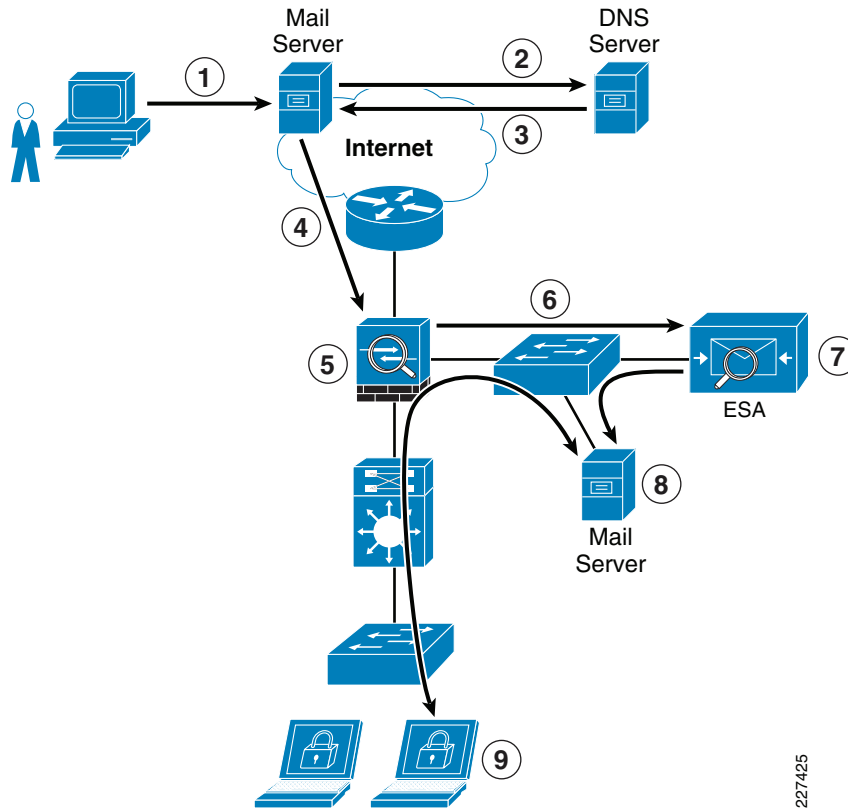
Figure 5-7 shows both configurations.

**Figure 5-7 Common ESA Deployments**



For simplicity, the Medium Enterprise Design Profiles implements the ESA with a single interface in a single-armed configuration. This also leaves the other data interfaces available for redundancy.

Figure 5-8 shows the logical location of the ESA within the E-mail flow chain and the typical data flow for inbound E-mail traffic.

**Figure 5-8 Typical Data Flow for Inbound E-Mail Traffic**

227425

The following steps explain what is taking place in [Figure 5-8](#):

- 
- |               |                                                                                                                                                         |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Sender sends an E-mail to xyz@domain X.                                                                                                                 |
| <b>Step 2</b> | What's the IP address of domain X?                                                                                                                      |
| <b>Step 3</b> | It is a.b.c.d (public IP address of ESA).                                                                                                               |
| <b>Step 4</b> | The E-mail server sends message to a.b.c.d using SMTP.                                                                                                  |
| <b>Step 5</b> | The firewall permits incoming SMTP connection to the ESA, and translates its public IP address.                                                         |
| <b>Step 6</b> | 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. |
| <b>Step 7</b> | ESA forwards E-mail to preconfigured inbound E-mail server.                                                                                             |
| <b>Step 8</b> | The E-mail server stores E-mail for retrieval by receiver.                                                                                              |
| <b>Step 9</b> | The receiver retrieves E-mail from server using POP or IMAP.                                                                                            |
- 

The Cisco IronPort ESA appliance functions as an SMTP gateway, also known as a mail exchange (MX). The following outlines some of the key deployment guidelines for the ESA within the Medium Enterprise Design Profile:


- The ESA appliance needs to be accessible via the public Internet and is the first hop in the E-mail infrastructure. If another MTA sits at your network's perimeter and handles all external connections, the ESA is not able to determine the sender's IP address. The IP address of the sender is needed to

identify and distinguish the senders in the Mail Flow Monitor to query the SensorBase Reputation Service for the SensorBase Reputation Service Score (SBRS) of the sender. Therefore, a separate MTA should not be deployed at the network perimeter to handle the external connections.

- The ESA needs to be registered in DNS for features such as IronPort Anti-Spam, Virus Outbreak Filters, MacAfee Antivirus, and Sophos Antivirus. A DNS “A” record should be created to map the appliance hostname to its public IP address, and an MX record that maps the public domain to the appliance hostname. A priority is specified for the MX record to advertise the ESA appliance as the primary MTA for the domain.
- A static IP address translation entry on the Internet firewall should be defined to map the public IP address of the ESA to its private internal address if NAT is configured on the Internet firewall.
- All the local domains for which the ESA appliance accepts mail needs to be added to the Recipient Access Table (RAT). Inbound E-mail destined to domains not listed in the RAT are 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, Domino) via SMTP routes.
- For each private listener, the host access table (HAT) must be configured to indicate the hosts that are 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, and on which mail policies can be applied. Policies include the following:
  - Mail flow policies—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 restricting E-mail access based on sender trustworthiness as determined by the IronPort SensorBase Reputation Service.
- Define SMTP routes to direct E-mail to the appropriate internal mail servers.
- If an OOB management network is available, a separate interface for administration should be used.

Because a failure on the ESA appliance may cause a service outage, a redundant design is recommended. One way to implement redundancy is to use IronPort NIC pairing, as shown in [Figure 5-9](#).

**Figure 5-9 Cisco IronPort ESA NIC Pairing**



IronPort NIC pairing provides redundancy at the network interface card level by teaming two of the Ethernet interfaces in the ESA appliance. If the primary interface fails, the IP address and MAC address are moved to the secondary interface. IronPort NIC pairing is the most cost-effective solution because it does not require the deployment of multiple ESA appliances and other hardware. However, it does not provide redundancy in case of chassis failure.

Alternative redundant designs include the following:

- *Multiple MTAs*—Adding a second ESA appliance or MTA and using a secondary MX record with an equal cost to load balance between the MTAs.
- *Load balancer*—Using a load balancer such as the Cisco Application Control Engine (ACE) to load balance traffic across multiple ESA appliances.

To accommodate traffic to and from the IronPort ESA provisioned in the DMZ, the Internet firewall needs to be configured to allow this communication. Protocols and ports to be allowed vary depending on the services configured on the ESA.

The following are some of the common services required to be allowed through the Internet firewall:

- 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, if the ESA is managed in-band, the appropriate firewall rules need to be configured to allow traffic such as SSH, NTP, and SYSLOG to/from the ESA.

For more information on how to configure the ESA, see the following guides:

- Cisco SAFE Reference Guide—[http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE\\_RG/SAFE\\_rg.html](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 Medium Enterprise Design Profile implements a Cisco IronPort S Series Web Security Appliance (WSA) to block HTTP and HTTPS access to sites on the Internet with non-business-related content and to protect the enterprise network from web-based malware and spyware.

The Cisco IronPort WSA relies on two independent services to protect the network:

- *Web proxy*—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)*—Monitors all Layer 4 traffic for rogue activity, and to detect infected clients.

The Medium Enterprise Design Profile assumes a centralized Internet connection implemented at the main site. The WSA should be implemented at the distribution layer in the Internet perimeter network. This allows for the inspection and enforcement of web access policies to all users located at any of the enterprise sites. Logically, the WSA sits in the path between web users and the Internet, as shown in [Figure 5-10](#).

**Figure 5-10 Cisco IronPort WSA**

There are two deployment modes for enabling the Cisco IronPort WSA 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 as its proxy. The web browsers can be configured either manually or by using proxy auto configuration (PAC) files. 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 Auto-discovery 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 to intercept and redirect traffic to the WSA before going to the Internet. Both WCCP and Layer 4 load balancer 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, or other devices outside the company's administration. Conversely, transparent proxy mode provides transparent integration of WSA without requiring any configuration control over the endpoints. In addition, transparent proxy also eliminates the possibility of users reconfiguring their web browsers to bypass the WSA appliance without the knowledge of the administrators. For these reasons, the Medium Enterprise Design Profile implements transparent proxy mode with WCCP. In the Medium Enterprise Design Profile, the Cisco Catalyst 3750 Stackwise distribution switches deployed in the Internet perimeter can be leveraged as the WCCP server while the WSA acts as a WCCP traffic processing entity.

The Cisco Catalyst 3750 switch uses WCCP version 2, which has a built-in failover and load balancing mechanism. Per the WCCPv2 specifications, 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 WSA appliances based on source and destination IP addresses. The WCCP server (Cisco Catalyst 3750 switch) monitors the availability of each appliance in the group and can identify the appliance failures in 30 seconds. After failure, the 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 communication between the WCCP server and the WSA appliances.

Figure 5-11 shows how WCCP redirection works in conjunction with the Cisco Catalyst 3750 StackWise distribution switches.

**Figure 5-11 WCCP Redirection**



As shown in Figure 5-11, the following steps take place:

1. The client browser requests a connection to `http://website.com`.
2. The Cisco Catalyst 3750 Internet perimeter distribution switch intercepts and redirects HTTP/HTTPS requests to WSA via Layer 2 redirection.
3. If the content is not present in the local cache, WSA performs a DNS query on the destination site and checks the received IP address against URL and reputation rules, and allows/denies the request accordingly.
4. If allowed, WSA fetches the content from the destination website.
5. The content is inspected and then delivered to the requesting client.



**Note**

In the event that the entire service group fails, WCCP automatically bypasses redirection, allowing users to browse the Internet without the web controls. If it is desired to handle a group failure by blocking all traffic, an inbound ACL may be configured on the Cisco ASA inside interface to permit only 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 direct access to the Internet without going through the WSA proxy.

WCCPv2 supports Generic Route Encapsulation (GRE) and Layer 2-based redirection. The Cisco Catalyst 6500 and 3750 switches support Layer 2-based redirection, and the redirection is supported in hardware. Therefore, the WSA must be directly connected to the switch running WCCP. In addition, WCCP is supported only on the ingress of an interface. For these reasons, WSA should connect directly to the Internet perimeter distribution switch using a VLAN that is different than the VLAN from where the client traffic is coming.

**Note**

The Cisco Catalyst 4500 does not provide the ability to create WCCP traffic redirect exception lists, which is an important component of the design. If a Cisco Catalyst 4500 is implemented as the distribution layer switch, another device, such as the Cisco ASA, should be used as the WCCP server.

The following describes some of the design considerations and caveats for implementing a Cisco IronPort WSA with WCCP on a Cisco Catalyst 3750 switch:

- The WSA must be Layer 2-adjacent to the Cisco Catalyst 3750 switch.
- The WSA and switches in the same service group must be in the same subnet directly connected to the switch that has WCCP enabled.
- Configure the switch interfaces that are facing the downstream web clients, the WSA(s), and the web servers as Layer 3 interfaces (routed ports or switch virtual interfaces [SVIs]).
- Use inbound redirection only.
- WCCP is not compatible with VRF-Lite. WCCP does not have visibility into traffic that is being used by the virtual routing tables with VRFs.
- WCCP and policy-based routing (PBR) on the same switch interface are not supported.
- WCCP GRE forwarding method for packet redirection is not supported.
- Use MD5 authentication to protect the communication between the Cisco Catalyst 3750 switches and the WSA(s).
- Use redirect-lists to specifically control what hosts/subnets should be redirected.
- Cisco Catalyst 3750 switches support switching in hardware only at Layer 2; therefore, no counters increment when the command **show ip wccp** is issued on the switch.
- In an existing proxy environment, deploy the WSA downstream from the existing proxy servers (closer to the clients).
- If an OOB management network is available, use a separate interface for WSA administration.

For more information on WCCP in relation to the Cisco Catalyst 3750 switch, see the following URL: [http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e\\_3560e/software/release/12.2\\_46\\_se/configuration/guide/swwccp.html](http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_46_se/configuration/guide/swwccp.html).

**Note**

WCCP, firewall, and other stateful features typically require traffic symmetry where traffic in both directions should flow through the same stateful device. Care should be taken when implementing active-active firewall pairs because they may introduce asymmetric paths.

The WSA appliance may also be configured to control or block peer-to-peer file sharing and instant messaging applications such as AOL Messenger, BitTorrent, Skype, Kazaa, and so on. Depending on the port used for transport, the WSA handles these applications as follows:

- *Port 80*—Applications that use HTTP tunneling on port 80 can be handled by enforcing access policies within the web proxy configuration. Applications access can be controlled based on applications, URL categories, and objects. Applications are matched based on their user agent pattern, and the use of regular expressions. URLs can be blocked based on specific categories, such as predefined chat and peer-to-peer categories, or custom categories defined by the administrator. Peer-to-peer access can also be filtered based on object and Multipurpose Internet Mail Extensions (MIME) types.
- *Ports other than 80*—Applications using ports other than 80 can be handled with the L4TM feature. L4TM can block access to specific applications by preventing access to the server or IP address blocks to which the client application must connect.

In the medium enterprise design, the Internet perimeter firewall can be configured to allow only web traffic (HTTP and HTTPS) outbound to the Internet from only the WSA. This prevents users from bypassing the WSA to browse the Internet.

**Note**

Peer-to-peer file sharing and Internet instant messaging applications can also be blocked using Cisco IPS appliances and modules and the Cisco ASA firewall (using modular policy framework).

The WSA L4TM service is deployed independently from the web proxy functionality. L4TM monitors network traffic for rogue activity and for any attempts to bypass port 80. It 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 periodic updates from the Cisco 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](http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html).
- *IronPort WSA User Guide*—<http://www.ironport.com/support>.

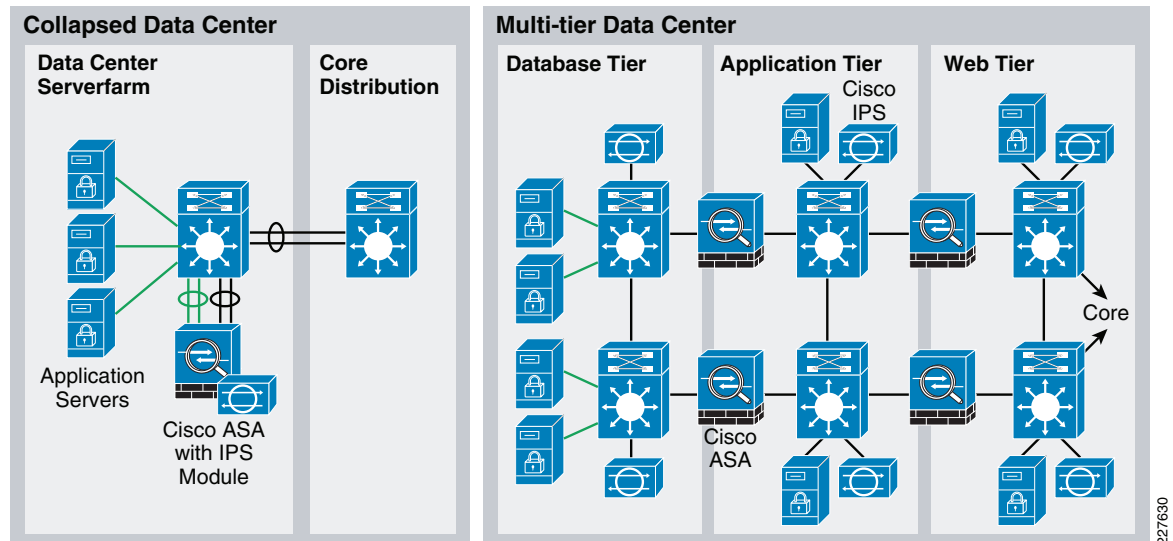
## Data Center Protection

Medium enterprise networks typically include a data center at the main site that hosts the systems that serve business applications and store the data accessible to internal users. The infrastructure supporting them 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 need and the size of the enterprise network, a single data center may be deployed at the main site. If needed, smaller data centers or server farms may also be deployed in remote sites.

The data center may be constructed following various design models. Figure 5-12 illustrates a collapsed and multi-tier data center 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 data center with chassis redundancy.

**Figure 5-12 Collapsed and Multi-tier Data Center Designs**




Independent from the design model adopted by the enterprise, the following are the primary security guidelines for the data center design:

- *Network Foundation Protection*—All infrastructure equipment should be protected following the NFP 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 groups 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 5-13.
- *Switch security*—Private VLANs, port security, storm control, and other switch security features may be leveraged to mitigate spoofing, man-in-the-middle, DoS, and other network-based attacks directed to the data center 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 5-13 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 data center security, see “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](http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap4.html).

## Network Access Security and Control

One of the most vulnerable points of a network is at the access edge. The access layer is where end users connect to the network. In the past, network administrators have largely relied on physical security to protect this part of the network. Unauthorized users were not allowed to enter secure buildings where they could plug into the network. Today, with the proliferation of wireless networks, and the increased use of laptops and smart mobile devices, enterprises cannot simply rely on physical controls to prevent these unauthorized devices from plugging into ports of the access switches. Contractors and consultants regularly have access to secure areas, and there is nothing preventing them from plugging into a wall jack in a cubicle or conference room to gain access to the enterprise network. When connected to the network, everyone has access to all resources on the network.

Protection needs to be embedded into the network infrastructure, leveraging the native security features available in switches and routers. In addition, the network infrastructure should also provide dynamic identity or role-based access controls for all devices attempting to gain access. Implementing role-based access controls for users and devices helps reduce the potential loss of sensitive information by enabling administrators to verify a user or device identity, privilege level, and security policy compliance before granting access to the network. Security policy compliance can consist of requiring anti-virus software, OS updates, or patches. Unauthorized or noncompliant devices can be placed in a quarantine area where remediation can occur before network access.

The Medium Enterprise Design Profile achieves access security and control by leveraging the following technologies:

- Cisco Catalyst Integrated Security Features (CISF), wired users
- Cisco Unified Wireless Network (CUWN) Integrated Security Features, wireless users

- Cisco Identity-Based Network Services (IBNS), wired and wireless users
- Cisco Network Admission Control (NAC) Appliance, wired and wireless users

## Cisco Catalyst Integrated Security Features

Cisco CISF is a set of native security features available on Cisco Catalyst switches designed to protect the access layer infrastructure and users from spoofing, man-in-the-middle (MITM), DoS, and other network-based attacks. CISF should be considered part of the security baseline of any network and should be deployed on all ports on the access switches within the enterprise network architecture.

CISF includes the following features:

- *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.
- *Dynamic ARP inspection (DAI)*—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 man-in-the-middle attacks.
- *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.
- *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](http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns820/landing_sec_wireless.html).

## Cisco Identity-Based Network Services (IBNS)

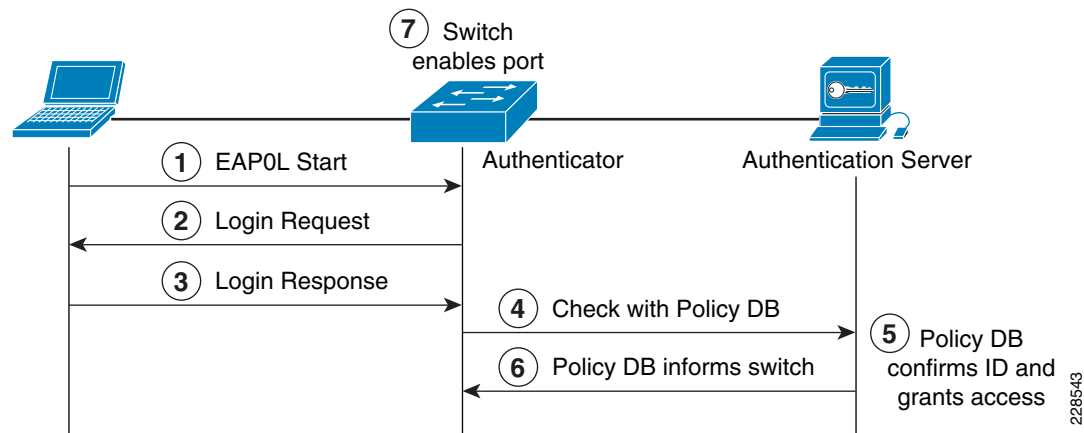
The Cisco IBNS solution is a set of Cisco IOS software services that provide secure user and host access to enterprise networks powered by Cisco Catalyst switches and wireless LANs. It provides standards-based network access control at the access layer by using the 802.1X protocol to secure the physical 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 a well-known way to secure wireless network access and is also capable of 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 the port except the request to start 802.1X authentication. After the 802.1X authentication successfully completes, the switch starts accepting other kinds of traffic on the port.

The high-level message exchange shown in [Figure 5-14](#) illustrates how port-based access control works within an identity-based system.

**Figure 5-14 Port-Based Access Control**



The following steps describe the port-based access control flow shown in [Figure 5-14](#):

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).
2. When the start message is received, the LAN switch sends a login request to the client.
3. The client replies with a login response.
4. The switch forwards the response to the policy database (authentication server).
5. The authentication server authenticates the user.
6. After the user identity is confirmed, the policy database authorizes network access for the user and informs the LAN switch.
7. The LAN switch then enables the port connected to the client.

The user or device credentials are processed by a AAA server. The AAA server is able to reference user or device profile information either internally, using the integrated user database, or externally using database sources such as Microsoft Active Directory, Lightweight Directory Access Protocol (LDAP), Novelle Directory, or Oracle databases. This enables the IBNS solution to be integrated into existing user management structures and schemes, which simplifies overall deployment.

## 802.1X and EAP

When authenticating users for network access, the client must provide user and/or device identification using strong authentication technologies. IEEE 802.1X does not dictate how this is achieved. Instead, 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 is defined by RFC 3748. EAP is a framework and not a specific authentication method. It 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

When 802.1X is enabled on a port, the default security posture is to drop all traffic except 802.1X EAPoL packets. This is a fundamental change from the traditional model, where traffic is allowed from the moment a port is enabled and a device is plugged into the port. Ports that were traditionally open are now closed by default. This is one of the key elements 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 accommodating for this change in access behavior facilitates 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 it connects. If a device without an 802.1X supplicant attempts to connect to a port that is enabled for 802.1X, it is subjected to the default security posture. The default security posture 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 a 802.1X-protected network.

Although an increasing number of devices 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 Preboot Execution Environment (PXE) boot machines. Some provisions must be made for these devices.

The Cisco IBNS solution provides two features to accommodate non 802.1X devices. These are MAC Authentication Bypass (MAB) and 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 a Guest VLAN is configured. Application of either or both of these features is required for a successful 802.1X deployment.



#### Note

Network-specific testing is required to determine the optimal values for the 802.1X timers to accommodate the various non-802.1X-capable devices on your network.



## 802.1X in Medium Enterprise Networks

As mentioned in the previous sections, 802.1X authentication requires a supplicant on the host device. This typically has been a challenge in enterprise environments that have a wide range of devices and limited or no management of many of these devices. In many enterprise environments, this is still the case, which makes a company-wide 802.1X deployment very challenging. However, there may be 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 (discussed in the next section). 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 NAC Appliance to protect these ports.

When considering an 802.1X deployment, there are four main 802.1X authentication options to consider:

- *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 bypass*—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>.

## Cisco NAC Appliance

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.

When deployed, Cisco NAC Appliance provides the following benefits:

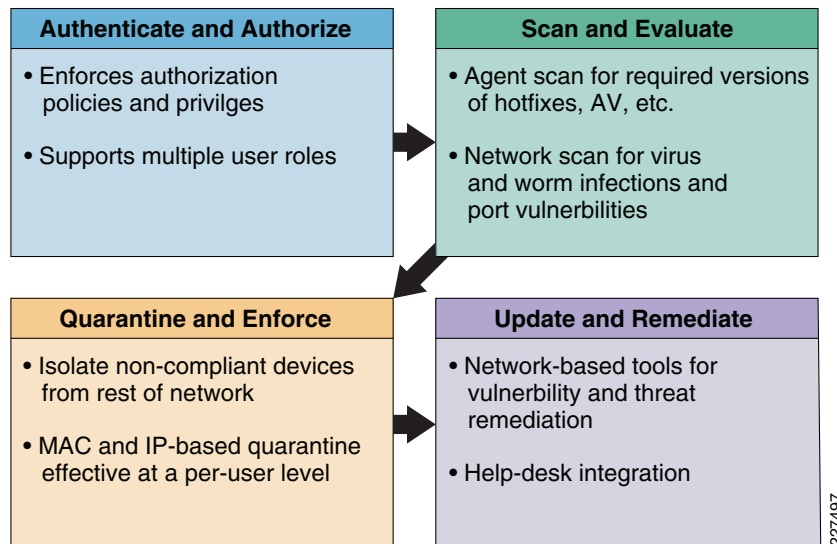
- Recognizes users, their devices, and their roles in the network. This first step occurs at the point of authentication, before malicious code can cause damage.
- Evaluates whether machines are compliant with security policies. Security policies can include requiring specific anti-virus or anti-spyware software, OS updates, or patches. Cisco NAC Appliance supports policies that vary by user type, device types, 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 5-15:

- Authenticates and authorizes
- Scans and evaluates
- Quarantines and enforces
- Updates and remediates

**Figure 5-15 Four Functions of the NAC Solution**

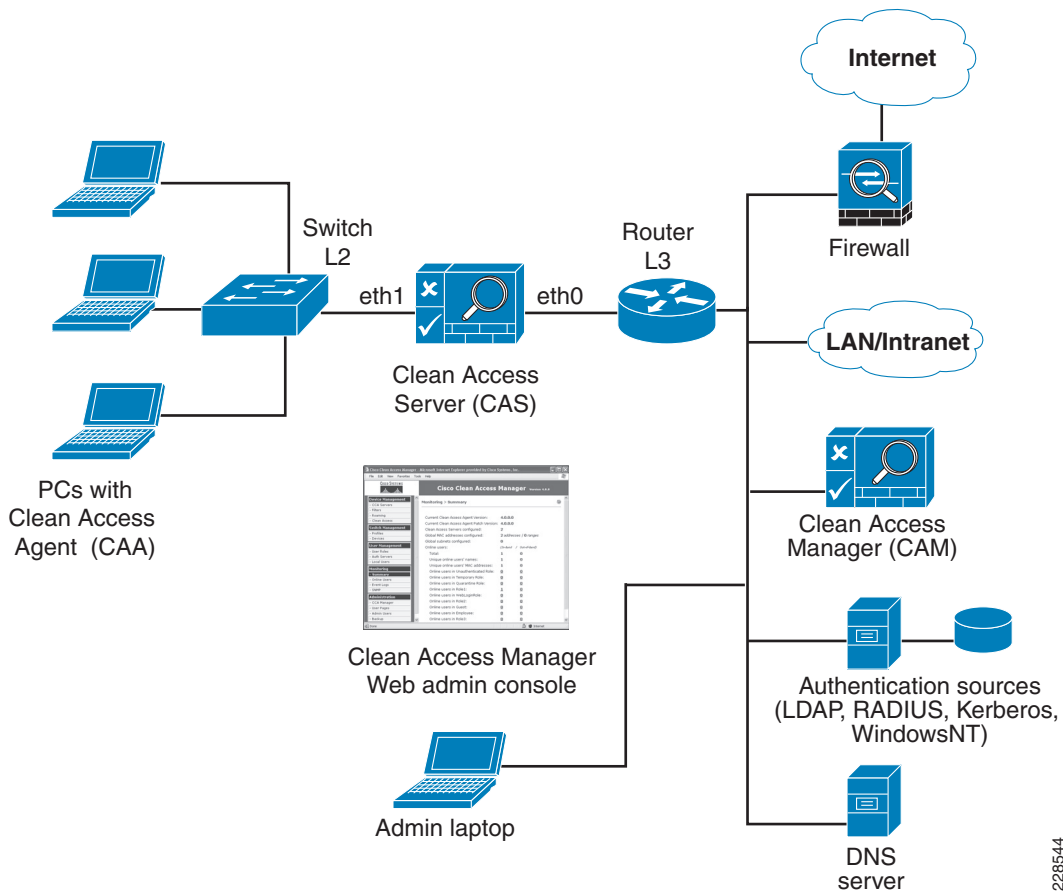


For more details of the NAC Appliance Solution, see the following URL:  
<http://www.cisco.com/go/nacappliance>.

## NAC Appliance Components

Cisco NAC Appliance is a network-centric, integrated solution administered from the Cisco Clean Access Manager (CAM) web console and enforced through the Cisco Clean Access Server (CAS) and (optionally) the Clean Access Agent (CAA) or NAC Web Agent. Cisco NAC Appliance checks client systems, enforces network requirements, distributes patches and anti-virus software, and quarantines vulnerable or infected clients for remediation before clients access the network.

Figure 5-16 shows Cisco NAC Appliance components.

**Figure 5-16 NAC Appliance Components**

### Cisco Clean Access Manager

The Cisco CAM is the administration server for NAC Appliance deployments. The secure web console of the CAM is the single point of management for up to 20 Clean Access Servers in a deployment (or 40 CASs if using a SuperCAM). For OOB deployments, the web administration console controls the switches and VLAN assignment of user ports through the use of SNMP. In the Medium Enterprise Design Profile, the CAM is located in the data center at the main campus site.

### Cisco Clean Access Server (CAS)

The Cisco CAS is the enforcement server between the untrusted network and the trusted network. The CAS enforces the policies defined by the CAM web administration console. Policies can include network access privileges, authentication requirements, bandwidth restrictions, and system requirements. The CAS can be installed as either a standalone appliance (such as the Cisco NAC-3300 Series) or as a network module (Cisco NME-NAC-K9) in a Cisco ISR chassis. The CAS can be deployed in in-band (always inline with user traffic) or OOB (inline with user traffic only during authentication and posture assessment).

Additionally, the CAS can 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). Multiple CASs of varying size/capacity can be deployed to fit the needs of various network segments. For example,

Cisco NAC-3300 Series appliances can be installed in a main site core to handle thousands of users, and one or more Cisco NAC network modules can be simultaneously installed in ISR platforms to accommodate smaller groups of users in a satellite office.

In the Medium Enterprise Design Profile, the CAS would be located at the main site and the remote campus sites, and deployed in Layer 2 OOB (for wireless clients) and Layer 3 OOB (for wired clients) modes for authentication and posture assessments.

### Cisco Clean Access Agent (CAA)

The Cisco CAA is an optional read-only agent that resides on Windows clients. It checks applications, files, services, or registry keys to ensure that clients meet the specified network and software requirements before 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 Medium 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. Using a Web browser, users launch the Cisco Web Agent executable file, 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 the network and their user ID disappears from the online users list.

In the Medium Enterprise Design Profile, the NAC Web Agent can be used for unmanaged clients such as guest users and contractors.

### Clean Access Policy Updates

Regular updates of prepackaged policies/rules can be used to check the up-to-date status of operating systems, anti-virus (AV), anti-spyware (AS), and other client software. Built-in support is provided for 24 AV and 17 AS vendors.


## NAC Appliance Modes and Positioning

The NAC Appliance can be deployed in multiple deployment options and placed at various locations in the network. The modes of operation can be generally defined as follows:

- Out-of-band (OOB) virtual gateway
- OOB real IP gateway
- In-band (IB) virtual gateway
- IB real IP gateway

### OOB Modes



OOB deployments require user traffic to traverse through the NAC Appliance only during authentication, posture assessment, and remediation (see [Figure 5-17](#)). When a user is authenticated and passes all policy checks, their traffic is switched normally through the network and bypasses the NAC Appliance.

**Figure 5-17 Layer 2 OOB Topology**

To deploy the NAC Appliance in OOB mode, the client device must be directly connected to the network via a Cisco Catalyst switch port. After the user is authenticated and passes posture assessment, the CAM instructs the switch to map the user port from an unauthenticated VLAN (which switches or routes user traffic to the CAS) to an authenticated (authorized) VLAN that offers full access privileges. For example, as shown in [Figure 5-17](#), the client PC is connected through VLAN 110 to the NAC CAS for the authentication and posture assessment, and is moved to VLAN 10 after it successfully completes the authentication/authorization and scan/evaluation phases of the NAC Appliance solution.

## In-Band Modes

When the NAC Appliance is deployed in-band, all user traffic, both unauthenticated and authenticated, passes through the NAC Appliance. The CAS may be positioned logically or physically between the end users and the networks being protected. [Figure 5-18](#) shows a logical in-band topology example, and [Figure 5-19](#) shows a physical in-band topology example.

**Figure 5-18 In-Band Virtual Gateway Topology****Figure 5-19 Physical In-Band Topology****In-Band Virtual Gateway**

When the NAC Appliance is configured as a virtual gateway, it acts as a bridge between the end users and the default gateway (router or switch) for the client subnet being managed. The following two bridging options are supported by the NAC server:

- *Transparent*—For a given client VLAN, the NAC server bridges traffic from its untrusted interface to its trusted interface. The NAC server is aware of “upper layer” protocols and is able to permit those protocols that are necessary for a client to connect to the network, authenticate, and undergo

posture assessment and remediation. 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, such as DNS and DHCP. This option is viable when the NAC server is positioned physically in-band between the end users and the upstream network(s) being protected, as shown in [Figure 5-19](#).

- *VLAN mapping*—This is similar in behavior to the transparent option except that rather than bridging the same VLAN from the untrusted side to the trusted side of the NAC server, two separate VLANs are used. For example, client VLAN 110 is defined for the untrusted interface of the NAC server. There is no routed interface or SVI associated with VLAN 110. VLAN 10 is configured between the trusted interface of the NAC server and the next-hop router interface (or SVI) for the client subnet. A mapping rule is made in the NAC server that forwards packets arriving on VLAN 110 and forwards them out VLAN 10 by swapping VLAN tag information. The process is reversed for packets returning to the client. Also, in this mode, BPDUs are not passed from the untrusted-side VLANs to their trusted-side counterparts.

The VLAN mapping option is typically used when the NAC server is positioned logically in-band between clients and the network(s) being protected, as shown in [Figure 5-18](#). This is the bridging option that should be used if the NAC Appliance is deployed in virtual gateway mode.

### In-Band Real IP Gateway

When the NAC server is configured as a “real” IP gateway, it behaves like a router and routes packets between its interfaces. In this scenario, one or more client VLAN/subnets resides behind the untrusted interface. The NAC server 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 server 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 server is not currently able to support dynamic routing protocols. Therefore, 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 the IP address of the NAC server trusted interface as its next hop.

If one or more Layer 3 hops exist between the untrusted NAC interface and the end-client subnets, static routes must be configured in the NAC server. In addition, a static default route is required within the downstream Layer 3 network (referencing the IP address of the untrusted NAC server interface) to facilitate default routing behavior from the client networks to the NAC server.

Depending on the topology, multiple options exist to facilitate routing clients to and from the NAC server, including ACLs, static routes, PBR, VRF-Lite, Multiprotocol Label Switching (MPLS) VPN, and other segmentation techniques. These options are discussed in later sections.

## In-Band Versus Out-of-Band

[Table 5-1](#) summarizes various characteristics of the two deployment types.

**Table 5-1** *In-Band versus Out-of-Band Characteristics*

| In-Band Deployment Characteristics                                                                                                                                                 | Out-of-Band Deployment Characteristics                                                                                                                                                                                                                                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The CAS is always inline with user traffic (both before and after authentication, posture assessment, and remediation). Enforcement is achieved through being inline with traffic. | The CAS is inline with the user traffic only during the process of authentication, posture assessment, and remediation. After that, user traffic does not go to the CAS. Enforcement is achieved through the use of SNMP to control switches and VLAN assignments to end-user ports. |

**Table 5-1 In-Band versus Out-of-Band Characteristics (continued)**

|                                                                                                                                                                   |                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The CAS can be used to securely control authenticated and unauthenticated user traffic policies (based on port, protocol, subnet), bandwidth policies, and so on. | The CAS can control user traffic during the authentication, posture assessment, and remediation phases but cannot do so post remediation because 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. <sup>1</sup>                                                                            |
| Cisco NAC in-Band deployment with supported Cisco switches is compatible with 802.1X.                                                                             | Cisco does not recommend using 802.1X in an OOB deployment, because conflicts will likely exist between Cisco NAC Appliance OOB and 802.1X in setting 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 access switches and WLCs to be supported by the NAC Appliance software for the NAC Manager to make the necessary changes to the switch ports and WLCs during the authentication, assessment, and remediation process. If access switches are to be used that are not supported, the NAC Solution must be deployed in in-band mode.

To obtain the latest list of supported devices, see 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/47/cam/47cam-book.html](http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/47/cam/47cam-book.html).

## Layer 2 and Layer 3 Out-of-Band

The proposed deployment option for the Medium Enterprise Design Profile is an OOB design. This provides the highest possible performance and scalability for traffic that has completed the authentication, posture assessment, and remediation stages of NAC. For wireless clients, a Layer 2 OOB solution should be deployed and for wired users, a Layer 2 OOB or Layer 3 OOB solution can be deployed, depending on the topology of your network.

## NAC Deployment in the Medium Enterprise Design Profile

Within the Medium Enterprise Design Profile, a Cisco NAC Appliance solution is deployed at each of the site types; main or headquarters site, remote large site, remote medium site, and remote small site. A centralized CAM is deployed at the main site and is deployed within the data center at that site. A CAS is deployed at each of the sites (main and remote sites) and is connected within the service block connecting to the core switches at each of the sites.

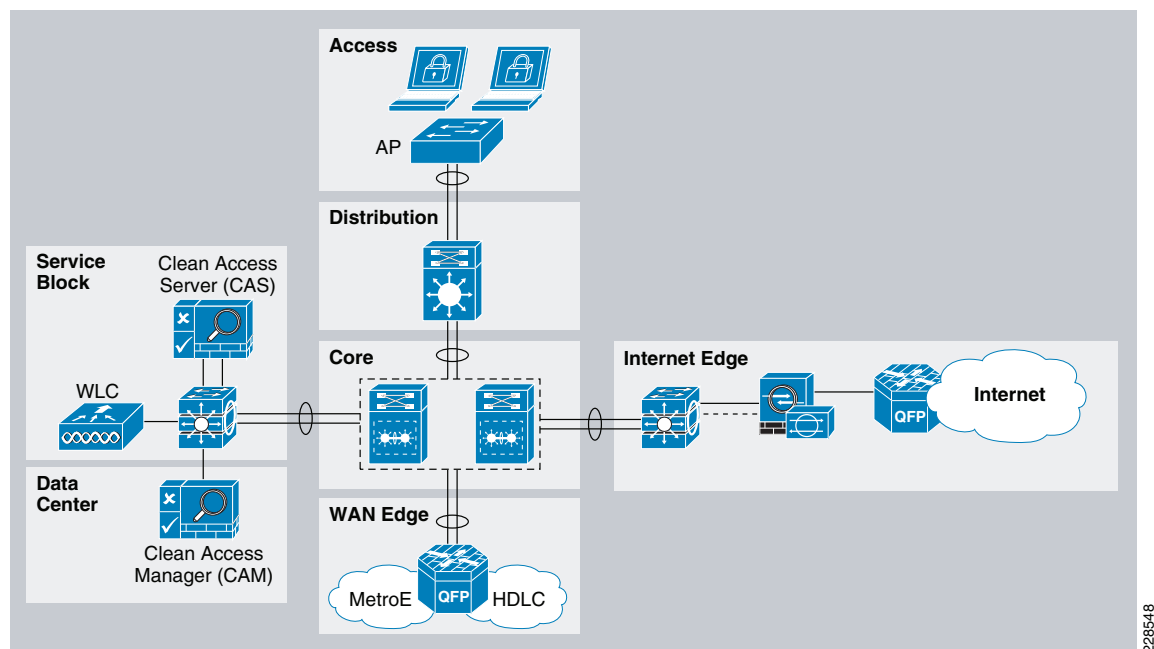
The Medium Enterprise Design Profile accommodates host network connectivity using wired and wireless technologies. As such, the NAC Appliance solution must provide a solution for both connectivity options. For wireless clients, a Layer 2 OOB NAC solution is deployed, and for wired clients, a Layer 2 OOB or a Layer 3 OOB NAC solution may be deployed.

## NAC Deployment for Wireless Clients

To provide network access control for wireless clients within the Medium Enterprise Design Profile, the recommended design is the virtual gateway (bridge mode) and central OOB deployment. In this design, the NAC server must be placed Layer 2-adjacent to the WLC. In the Medium Enterprise Design Profile, the WLCs are centrally deployed at each site and are implemented in the service block off the core switches, as detailed in [Chapter 4, “Medium Enterprise Design Profile \(MEDP\)—Mobility Design.”](#)

Therefore, the NAC server must also be implemented in the service block. The NAC Manager is implemented in the data center block, as shown in [Figure 5-20](#).


**Figure 5-20 NAC OOB Deployment for Wireless Clients**



The WLC connects to the service block switch using a trunk port carrying the unauthenticated quarantine VLAN and authenticated access VLAN (VLAN 20 and 120). On the switch, the quarantine VLAN is trunked to the untrusted interface on the NAC server (CAS), and the access VLAN is trunked directly to the Layer 3 switch interface. Traffic that reaches the quarantine VLAN on the CAS is mapped to the access VLAN based on a static mapping configuration within the CAS.

When a wireless client associates to the WLC, it initially maps the WLAN/SSID to the quarantine VLAN interface and the client traffic flows in the quarantine VLAN (VLAN 120), which is trunked to the CAS untrusted interface. When NAC authentication, posture assessment, and remediation stages are complete and the user is certified, the NAC Manager sends an SNMP set message to the WLC that updates the VLAN ID from the quarantine VLAN to the access VLAN. After this occurs, the traffic then bypasses the NAC server and goes directly to the network. See [Figure 5-21](#).



**Figure 5-21 Wireless NAC OOB Traffic Flow**

When implementing the NAC OOB wireless solution, it is recommended to enable RADIUS single sign-on (SSO), which is an option that does not require user intervention and is relatively easy to implement. This option makes use of the VPN SSO capability of the NAC solution, coupled with the Clean Access Agent software that runs on the client PC. VPN SSO uses RADIUS accounting records to notify the NAC Appliance about authenticated remote access users that connect to the network. In the same way, this feature can be used in conjunction with the WLAN controller to automatically inform the NAC server about authenticated wireless clients that connect to the network.

The most transparent method to facilitate wireless user authentication is to enable VPN SSO authentication on the NAC server and configure the WLCs to forward RADIUS accounting to the NAC server. In the event that accounting records need to be forwarded to a RADIUS server upstream in the network, the NAC server can be configured to forward the accounting packet to the RADIUS server.

**Note**

If VPN SSO authentication is enabled without the Clean Access Agent installed on the client PC, users are still automatically authenticated. However, they are not automatically connected through the NAC Appliance until their web browser is opened and a connection attempt is made. In this case, when users open their web browser, they are momentarily redirected (without a logon prompt) within the agentless phase. When the SSO process is complete, they are connected to their originally requested URL.

For more information on deploying NAC OOB for wireless environments, see the *NAC Out-Of-Band (OOB) Wireless Configuration Example* at the following URL:  
[http://www.cisco.com/en/US/products/ps6128/products\\_configuration\\_example09186a0080a138cc.shtml](http://www.cisco.com/en/US/products/ps6128/products_configuration_example09186a0080a138cc.shtml).

## NAC Deployment for Wired Clients

For wired clients, the Medium Enterprise Design Profile also uses a central OOB NAC deployment with a NAC server implemented at each of the sites deployed in the service block off the core switch. Depending on the type of network topology deployed, a Layer 3 OOB or Layer 2 OOB solution can be deployed. If the Layer 2 OOB solution is used, the same NAC server can be leveraged for both wired and wireless clients. However, if the Layer 3 OOB solution is deployed, separate NAC servers must be deployed for wired and wireless users.

### Layer 3 Out-of-Band Deployment

Layer 3 (L3) OOB is best suited for routed access designs and has rapidly become one of the most popular deployment methodologies for NAC. By deploying NAC in an L3 OOB methodology, a single NAC Appliance can scale to accommodate more users. This deployment also allows NAC Appliances to be centrally located rather than distributed across the site or organization. Thus, L3 OOB deployments are much more cost-effective, both from a capital and operational expense standpoint.

For the main, large, and medium remote site locations, an L3 OOB NAC deployment is recommended, given the 3-tier hierarchical design. In the L3 OOB NAC solution, when a user connects to the access switch before being certified by the NAC server, the user is placed in the authentication VLAN (also called “dirty” VLAN). The user should not have access to any part of the network from the authentication VLAN except for the NAC server and the remediation servers in the quarantine segment. After users are certified by the NAC server, they are placed in the authenticated access VLAN, where their traffic is switched normally through the network and bypasses the NAC server.

The following are three widely deployed techniques for redirecting client traffic from the dirty VLAN to the NAC server for authentication, posture assessment, and remediation purposes:

- *Access control lists*—Use ACLs on the edge access switches to allow traffic from the unauthenticated VLAN only to the NAC server untrusted interface and specific infrastructure resources needed to get on the network for authentication purposes such as DHCP, DNS, and remediation servers. All other traffic from the dirty VLAN must be blocked.
- *VRFs/GRE/MPLS*—Use VRFs to route unauthenticated traffic to the CAS. Traffic policies configured on the NAC server (CAS) are used for enforcement on the dirty network. This approach has two sub-approaches. In the first approach, VRFs are pervasive throughout the infrastructure, in which case all Layer 3 devices participate in the tag switching. The second approach uses VRF-Lite and GRE tunnels to tunnel the VRFs through the Layer 3 devices that do not understand tag switching. The benefit to the second approach is that minimal configuration changes are required to your core infrastructure. For more information on this approach, see the following URL: [http://www.cisco.com/en/US/products/ps6128/products\\_configuration\\_example09186a0080a3a8a7.shtml](http://www.cisco.com/en/US/products/ps6128/products_configuration_example09186a0080a3a8a7.shtml).
- *Policy-based routing*—Use PBR to redirect all traffic in the dirty VLAN to the NAC server. PBR needs to be configured on every Layer 3 hop between the dirty VLAN and the NAC server to ensure that traffic is appropriately redirected.

The most common approach used for isolating the dirty VLAN traffic is to use ACLs. The ACLs on the Layer 3 edge access switches act as the enforcement point to ensure segregation between the “clean” and “dirty” networks. When clients first attach to the network, they are placed in a quarantine or dirty VLAN on the access switches. ACLs should be applied on the SVIs for the dirty VLAN. This ACL should block all access from the dirty VLAN going to the internal networks and allow traffic only to the untrusted interface on the NAC server, the needed remediation servers, and a few infrastructure devices needed for network access such as the DNS, DHCP, and Active Directory servers.

The clients need to communicate with the NAC server untrusted interface for the certification process. The ACLs on the access switches act as the enforcement point for path isolation for the dirty VLAN traffic. Methods for getting the dirty VLAN traffic to the untrusted interface vary, depending on whether the NAC Client Agent is used.

When the NAC agent is used, the NAC Agent communicates with the NAC server untrusted interface to initiate the login process. The NAC Agent tries to discover the NAC server based on the known discovery host value. The discovery host value in the NAC Agent points to the untrusted interface of the NAC server. In the Medium Enterprise Design Profile, the NAC Agent can be used for managed PCs.

Web login may also be required for devices that are not managed. With the ACL isolation technique, the NAC server untrusted interface is not directly in the path of the data traffic; therefore, the user is not automatically redirected to the login page when first opening the browser. The following two options can enable the end host to get the login page:

- *Option 1*—Have a guest login URL known to the users (for example, *guest.nac.local*). In this case, the guest must open a browser and manually enter this URL, which redirects them to the login page.
- *Option 2*—Create a dummy DNS server for the unauthenticated user subnet. This dummy DNS server resolves every URL to the untrusted interface of the NAC server. When guests open a browser, regardless of which URL they are trying to reach, they are redirected to the login page. When users are then moved to the respective Role/VLAN, they get a new DNS address assignment when performing IP release/renew on a successful login.


### Layer 2 Out-of-Band Deployment

For the small remote sites, a two-tier, collapsed core/distribution LAN design is recommended, as explained in [Chapter 2, “Medium Enterprise Design Profile \(MEDP\)—LAN Design.”](#)

In a collapsed core/distribution design, the CAS should be deployed in the services block connected to the core/distribution switch. In this simple topology, a Layer 2 Out-of-Band (L2 OOB) NAC deployment can be used.

In the L2 OOB NAC design for the small remote site, the unauthenticated and authenticated VLANs on the access switch (VLANs 30 and 130) are extended to the core/distribution switch using a trunk connection, as shown in [Figure 5-22](#).

**Figure 5-22 Layer 2 OOB Topology**



When a client device initially connects to the access switch before authentication, it is placed in the unauthenticated VLAN (VLAN 130), which connects the client directly to the untrusted interface of the CAS. The CAS maps VLAN 130 to the VLAN 30 trusted interface, allowing the client to obtain an IP address that belongs on VLAN 30. After the client is authenticated and passes the posture assessment, the access switch is instructed, via SNMP from the CAM, to change the client VLAN to the authenticated VLAN (VLAN 30), where the traffic now bypasses the CAS to access the rest of the network. Although the client has changed Layer 2 VLANs, its Layer 3 network connections are unchanged.

## NAC Availability Considerations

Both the CAS and CAM are highly involved in client network access. Consideration must be given to the impact on clients if either a CAS or CAM fails or needs 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 NAC is deployed in in-band mode, it is inline even after authentication and certification. A CAS outage for inline clients prevents access for all clients. However, if NAC is deployed in OOB mode, a CAS outage does not affect already connected clients but does prevent network access for new clients.

In situations where availability of a CAS is critical, a high availability (HA) CAS solution can be implemented where a pair of CAS servers are installed using a primary CAS, and a secondary in hot standby. For more information, see the *Cisco NAC Appliance - Clean Access Server Installation and Configuration Guide* at the following URL:  
[http://www.cisco.com/en/US/docs/security/nac/appliance/configuration\\_guide/461/cas/461cas-book.html](http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/461/cas/461cas-book.html).

The CAM is also a critical part of the authentication, authorization, and posture assessment phases of NAC. Although it does not pass client traffic, the impact of its availability needs to be considered in the network design as well. Like the CAS, the CAM has an 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 that defines how it manages client traffic in a situation where the CAM is unavailable.

The use of the CAM and CAS HA features depends on the requirements of the enterprise. However, CAS fallback should always be configured to ensure that critical network services are available, even during a network outage.

# Secure Mobility

Today's 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 medium enterprise more competitive, a number of challenges arise from the use of mobile technologies. 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 types 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. In addition, the great variety in hardware types, operating systems, and applications represents a clear challenge to the enforcement of security controls and policies.

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. Additionally, 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 Medium Enterprise Design Profile provides persistent and secure mobile access by implementing the Cisco AnyConnect Secure Mobility solution. 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:

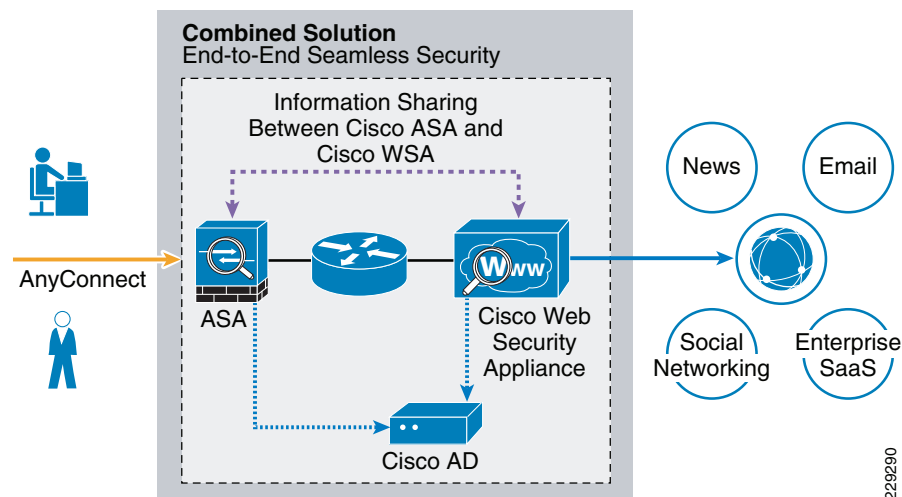
- Cisco IronPort Web Security Appliance (WSA)
- Cisco ASA 5500 Series Adaptive Security Appliance (ASA)
- Cisco AnyConnect client

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 before 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 5-23 shows the relationship between the various elements of the Cisco AnyConnect Secure Mobility solution.

**Figure 5-23 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 because 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, see 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 Medium Enterprise Design Profile help mitigate certain threats, and how they contribute to increasing visibility and control. Note that 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 Enterprise Design Profile for Small Enterprise Networks**

|                                      | <b>Service Disruption</b> | <b>Harmful Content</b> | <b>Network Abuse</b> | <b>Unauthorized Access</b> | <b>Data Loss</b> | <b>Visibility</b> | <b>Control</b> |
|--------------------------------------|---------------------------|------------------------|----------------------|----------------------------|------------------|-------------------|----------------|
| <b>Network Foundation Protection</b> | Yes                       |                        |                      | Yes                        | Yes              | Yes               | Yes            |
| <b>Stateful Firewall</b>             | Yes                       |                        | Yes                  | Yes                        |                  | Yes               | Yes            |
| <b>IPS</b>                           | Yes                       | Yes                    | Yes                  | Yes                        |                  | Yes               | Yes            |
| <b>Security Mobility</b>             | Yes                       | Yes                    | Yes                  | Yes                        | Yes              | Yes               | Yes            |
| <b>Web Security</b>                  |                           | Yes                    | Yes                  | Yes                        | Yes              | Yes               | Yes            |
| <b>E-mail Security</b>               |                           | Yes                    | Yes                  |                            | Yes              | Yes               | Yes            |
| <b>Access Security and Control</b>   |                           |                        | Yes                  | Yes                        |                  | Yes               | Yes            |

# Medium Enterprise Network Security Deployment Guidelines

The previous sections of this document provide design guidelines and considerations for deploying security within a medium enterprise network environment. The sections that follow provide deployment and configuration examples and guidelines for deploying some of these features. Security features and devices covered include the following:

- Internet Border Router Edge ACL Deployment
- Internet Firewall Deployment
- IPS Global Correlation Deployment
- Web Security Deployment
- Catalyst Integrated Security Features Deployment
- NAC Deployment for Wired and Wireless Clients

## Internet Border Router Edge ACL Deployment

Whether the Internet border router is managed by the enterprise or the ISP, it must be hardened following the best practices discussed in the [“Network Foundation Protection” section on page 5-6](#). 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 snippets illustrate 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—Implement Anti-spoofing Denies

These ACEs deny fragments, RFC 1918 space, invalid source addresses, and spoofs of the internal address space.

- 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—Implement Explicit Permits

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—Implement Explicit Deny to Protect Infrastructure

```
access-list 110 deny ip 64.104.10.0 0.0.0.255 any
```

## Module 4—Implement Explicit Permit for Traffic to the Enterprise 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 document are reserved for the exclusive use of Cisco Systems, Inc.

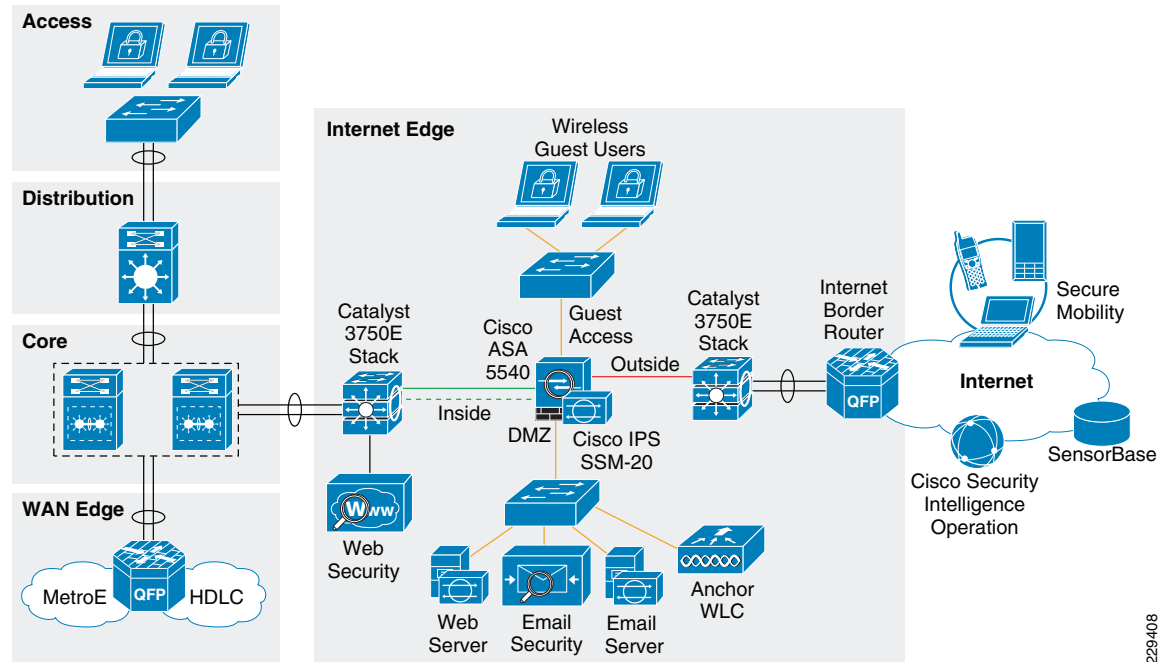
For more information and configuration examples on how to secure the Internet border router using the other Network Foundation Protection features, see “Chapter 6, Enterprise Internet Edge” in the *Cisco SAFE Reference Guide* at the following URL:

[http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE\\_RG/chap6.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap6.html).

## Internet Firewall Deployment

The Internet firewall is responsible for protecting the enterprise's internal resources and data from external threats, securing the public services provided by the DMZ, and to control user's traffic to the Internet. The Medium Enterprise Design Profile uses a Cisco ASA appliance for the Internet Firewall, as illustrated in [Figure 5-24](#).



**Figure 5-24 Internet Edge Firewall**

229408

The Cisco ASA is implemented with four interface groups with each group representing a distinct security domain:

- **Inside**—The interface connecting to the 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; therefore, it is given the lowest security level, 0.
- **Demilitarized zone (DMZ)**—The DMZ hosts the enterprise's public-facing 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, and therefore should be given any security value between the ones defined for the inside and the outside interfaces; for example, 50.
- **Guest access**—The interface connecting to the LAN segment that wireless guest access in which clients will be placed by the anchor WLC in the DMZ. Wireless guest access clients should have access only to the Internet and not to any internal resources. Therefore, they should be given a value lower than the internal and DMZ interfaces; for example, 10.

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
- Firewall redundancy
- Routing
- Botnet Traffic Filter

## 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 the 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” section on page 5-52](#)).

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 172.26.136.170 255.255.254.0
 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 ASDM. Additionally, this access should only be permitted for users authorized to access the firewalls for management purposes.

The following Cisco ASA configuration fragment illustrates the configuration needed to generate a 768 RSA key pair and enable 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-subnet> 255.255.255.0 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 Cisco 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 a Cisco 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 switch should be secured. The following Cisco ASA configuration fragment illustrates the use of EIGRP MD5 authentication to authenticate the peering session between the inside firewall interface and the Internet edge distribution switch:

```
interface Redundant1
  description connection to CR12-3750s-IE distribution switch
  nameif inside
  security-level 100
  ip address 10.125.32.18 255.255.255.240
```

```
authentication key eigrp 100 <removed> key-id 1
authentication mode eigrp 100 md5
```

## Network Address Translation (NAT)

NAT is required because enterprises typically get a limited number of public IP addresses. In addition, NAT helps shield the company's internal address space from reconnaissance and other malicious activity. The following illustrates the NAT configuration:

```
! Static translation for servers residing at DMZ
static (dmz,outside) 198.133.219.35 10.125.32.35 netmask 255.255.255.255
static (dmz,outside) 198.133.219.36 10.125.32.36 netmask 255.255.255.255
static (dmz,outside) 198.133.219.40 10.125.32.40 netmask 255.255.255.255
static (dmz,outside) 198.133.219.41 10.125.32.41 netmask 255.255.255.255
!
! Dynamic Port Address Translation (PAT) for inside hosts and wireless guest access
! clients going to the Internet
global (outside) 10 interface
nat (inside) 10 10.0.0.0 255.0.0.0
nat (guestaccess) 10 10.125.32.64 255.255.255.240
!
! 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

The Internet firewall should be configured with access policies to do the following:

- Protect the enterprise'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
- Preventing wireless guest access users from accessing internal resources

Enforcing such policies requires configuration of the appropriate interface security levels and the deployment of ACLs to control 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 (that is, from inside to outside). However, depending on the sensitivity of an enterprise environment, the security administration is recommended to override the default rules with more stringent rules indicating exactly what ports and protocols are permitted.

In this configuration example, the inside, DMZ, guestaccess, and outside interfaces were configured with the security levels of 100, 50, 10, and 0, respectively. With this, by default any traffic originating from the inside to the DMZ, guestaccess, and outside, from the DMZ to the guestaccess and outside interface, and from the guestaccess to the outside is allowed freely. At the same time, any traffic originating from the outside to the DMZ, guestaccess, and inside, and from the DMZ to the guestaccess and inside interfaces is blocked. Although 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.

Note also 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 associated with those sessions. Therefore, the firewall ACLs should be constructed to match traffic in the direction in which it is being initiated. In the following sample configurations, ACLs are applied in the ingress direction. The following are guidelines and configuration examples for the ACLs controlling access and traffic flows:

- Ingress inside

Allow users residing at all enterprise sites to access the Internet for the allowed ports and protocols. Depending on the policy of the enterprise, this may allow only HTTP and HTTPS access, or may be less restrictive to allow additional protocols and ports. The following example allows only HTTP and HTTPS access to the Internet:

```
access-list outbound extended permit tcp 10.0.0.0 255.0.0.0 any eq www
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 web portal, E-mail, and domain name resolution. This can include HTTP, HTTPS, SMTP, POP, IMAP, and DNS protocols. Permit tunneled control and user traffic from internal WLCs to the Anchor WLC in the DMZ for wireless guest access (UDP 16666, UDP 16667, IP Protocol 97). Permit management traffic from the management segment to the Anchor WLC in the DMZ (SNMP, SSH, and HTTPS). Allow WSA access to the IronPort SensorBase network (HTTPS) for updates. Note that the previous entries in the ACL already permit HTTP and HTTPS traffic.

```
! Allow DNS queries to DNS server located in DMZ
access-list outbound extended permit udp 10.0.0.0 255.0.0.0 host 10.125.32.35 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 10.125.32.40 eq
smtp
access-list outbound extended permit tcp 10.0.0.0 255.0.0.0 host 10.125.32.40 eq
pop3
access-list outbound extended permit tcp 10.0.0.0 255.0.0.0 host 10.125.32.40 eq
imap4
! Allow access to the Anchor WLC on the DMZ from the internal WLCs for wireless
Guest access
access-list outbound extended permit udp host 10.125.30.2 host 10.125.32.34 eq
16666
access-list outbound extended permit udp host 10.125.30.3 host 10.125.32.34 eq
16666
access-list outbound extended permit udp host 10.124.2.66 host 10.125.32.34 eq
16666
access-list outbound extended permit udp host 10.125.30.2 host 10.125.32.34 eq
16667
access-list outbound extended permit udp host 10.125.30.3 host 10.125.32.34 eq
16667
access-list outbound extended permit udp host 10.124.2.66 host 10.125.32.34 eq
16667
access-list outbound extended permit 97 host 10.125.30.2 host 10.125.32.34
access-list outbound extended permit 97 host 10.125.30.3 host 10.125.32.34
access-list outbound extended permit 97 host 10.124.2.66 host 10.125.32.34
! Allow management access to the Anchor WLC on the DMZ
access-list outbound extended permit udp 10.125.31.0 255.255.255.0 host
10.125.32.34 eq snmp
access-list outbound extended permit udp 10.125.31.0 255.255.255.0 host
10.125.32.34 eq snmptrap
access-list outbound extended permit tcp 10.125.31.0 255.255.255.0 host
10.125.32.34 eq ssh
!
! 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, and so on.

```
! Allow DNS queries and zone transfer from DNS server
access-list dmz-acl extended permit udp host 10.125.32.35 any eq domain
access-list dmz-acl extended permit tcp host 10.125.32.35 any eq domain
!
! Allow SMTP from Cisco IronPort ESA
access-list dmz-acl extended permit tcp host 10.125.32.36 any eq smtp
!
! Allow update and SensorBase access to Cisco IronPort ESA
access-list dmz-acl extended permit tcp host 10.125.32.36 any eq www
access-list dmz-acl extended permit tcp host 10.125.32.36 any eq https
!
! Apply ACL to DMZ interface
access-group dmz-acl in interface dmz
```

- Ingress outside

Inbound traffic from the Internet 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 198.133.219.35 eq domain
access-list inbound extended permit tcp any host 198.133.219.35 eq domain
!
! Allow SMTP to Cisco IronPort ESA
access-list inbound extended permit tcp any host 198.133.219.36 eq smtp
!
! Allow HTTP/HTTPS access to the company's public web portal
access-list inbound extended permit tcp any host 198.133.219.41 eq www
access-list inbound extended permit tcp any host 198.133.219.41 eq https
!
! Apply ACL to outside interface
access-group inbound in interface outside
```

- Ingress guest access

Wireless guest access users should be restricted to having access only to the Internet. Access to the internal enterprise network should not be allowed. Because the security level of the Guest Access interface is lower than the internal and DMZ interfaces, traffic coming from the Guest Access interface going to the internal and DMZ segments is automatically blocked. In addition, because the security level for the Guest Access interface is also higher than the outside interface, traffic is permitted to the Internet. If the guest wireless clients need to access the DMZ servers, an ACL must be configured and applied to allow this access. However, if desired to reinforce this policy, granular ACLs may also be applied to the guestaccess interface.

```
! Deny access to internal networks
access-list guestaccess-acl extended deny ip any 10.0.0.0 255.0.0.0
access-list guestaccess-acl extended deny ip any 192.168.0.0 255.255.0.0
access-list guestaccess-acl extended deny ip any 172.16.0.0 255.240.0.0
! Permit all other access to the Internet
access-list guestaccess-acl extended permit ip any any
!
! Apply ACL to guest-access interface
access-group guestaccess-acl in interface guestaccess
```

## Firewall Redundancy

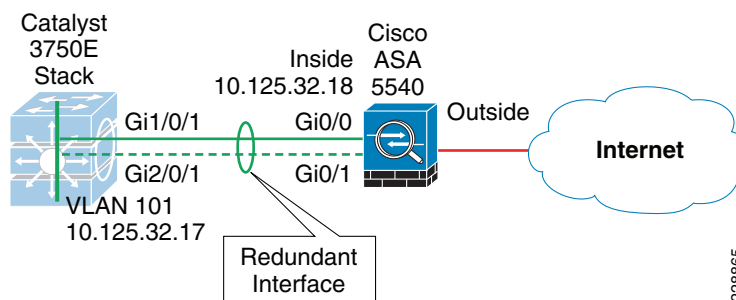
The Internet perimeter of the Medium Enterprise Design Profile uses a single Cisco ASA appliance configured with redundant interfaces. The use of redundant interfaces makes the design resilient to link-level failures, representing an affordable option for high availability. In cases where chassis redundancy 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 section explains the use of redundant interfaces. For information on how to configure stateful failover using multiple platforms, 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](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 5-25 illustrates the concept of redundant interface.

**Figure 5-25 Cisco ASA Redundant Interface**



The configuration of a redundant interface requires 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 logical redundant interface. The following configuration example corresponds to Figure 5-25.

```
! Physical interface and Ethernet parameters
interface GigabitEthernet0/0
description Connection to CR12-3750s-IE port Gig1/0/1
no nameif
no security-level
no ip address
!
interface GigabitEthernet0/1
description backup connection to CR12-3750s-IE port Gig2/0/1
no nameif
no security-level
no ip address
!
! Define logical redundant interface and associate with physical interfaces.
! Configures IP and logical interface parameters.
```

```

interface Redundant1
  description connected to CR12-3750s-IE
  member-interface GigabitEthernet0/0
  member-interface GigabitEthernet0/1
  nameif inside
  security-level 100
  ip address 10.125.32.18 255.255.255.240
  authentication key eigrp 100 ***** key-id 1
  authentication mode eigrp 100 md5
!

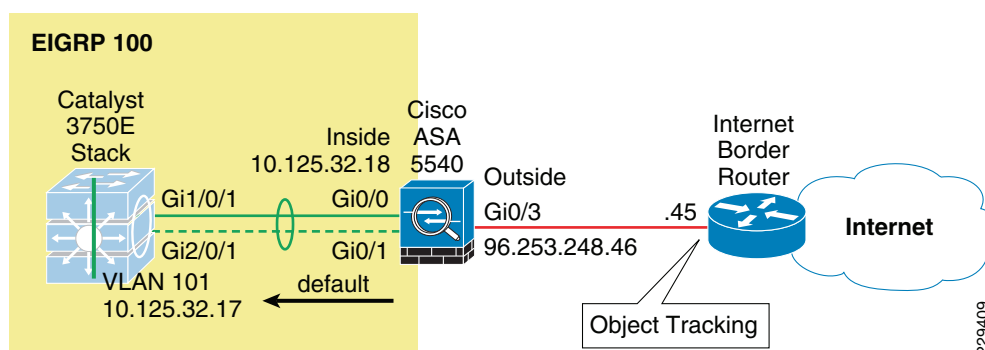
```

## Routing

Within the Medium Enterprise Design Profile, an interior gateway protocol, EIGRP, is used for dynamic routing. The Internet firewall may participate in routing by learning the internal routes within the enterprise network and by injecting a default route pointing to the Internet. The default route should be removed dynamically if the Internet connection becomes unavailable.

Within the Medium Enterprise Design Profile, 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, as shown in [Figure 5-26](#).

**Figure 5-26 Cisco ASA Static Route**



It is highly recommended to use object tracking so that the default route is removed when the Internet connection becomes unavailable. Without object tracking, the default route is removed only if the outside interface of the appliance goes down. Therefore, there is a possibility that the default route may remain in the routing table even if the Internet border router becomes unavailable. To avoid this problem, the static default route can be configured with object tracking. This is accomplished by 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. Pick one that can receive and respond to ICMP echo requests sent by the Cisco ASA. 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 does 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 snippet, the IP address of the next hop gateway (96.253.248.45) is used as the monitoring target. The static default route is then redistributed into EIGRP.

```

router eigrp 100

```



```

network 10.125.32.0 255.255.255.0
passive-interface default
no passive-interface dmz
no passive-interface inside
redistribute static metric 1000000 2000 255 1 1500
!
route outside 0.0.0.0 0.0.0.0 96.253.248.45 1 track 10
!
sla monitor 1
type echo protocol ipIcmpEcho 96.253.248.45 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.

Another option for dynamically controlling the injection and removal of a default route in the enterprise routing table is to use OSPF, where 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. Injecting a default route with OSPF requires the configuration of an OSPF process between the Cisco ASA and the Internet border router. If the Internet border router is managed by the ISP, the configuration requires 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.

## Botnet Traffic Filter

The Medium Enterprise Design Profile uses the 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:

1. Configure DNS server.
2. Enable use of the dynamic database.
3. Enable DNS snooping.
4. Enable traffic classification and actions for the Botnet Traffic Filter.
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 enabled only 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 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 and/or drops any matching traffic. When an address matches, the 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 lookmbbox.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 btbf-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.

## Intrusion Prevention Deployment

The Medium 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 some of the best practices for integrating and configuring the IPS service 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 AIP SSM is supported in the 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.

The AIP SSM may be deployed in inline or promiscuous mode as described in [Intrusion Prevention, page 5-13](#). In inline mode, the AIP SSM is placed directly in the traffic flow; in promiscuous mode, the Cisco ASA sends a duplicate stream of traffic to the AIP SSM for inspection.

When deploying the AIP SSM in inline mode, it is important to determine how traffic should be treated in case of module failure. The AIP SSM 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, leaving your network unprotected. Conversely, when configured to fail close, the Cisco ASA blocks all traffic in case of an AIP SSM failure. This is more secure but impacts traffic during a 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 in the AIP SSM card fails for any reason:

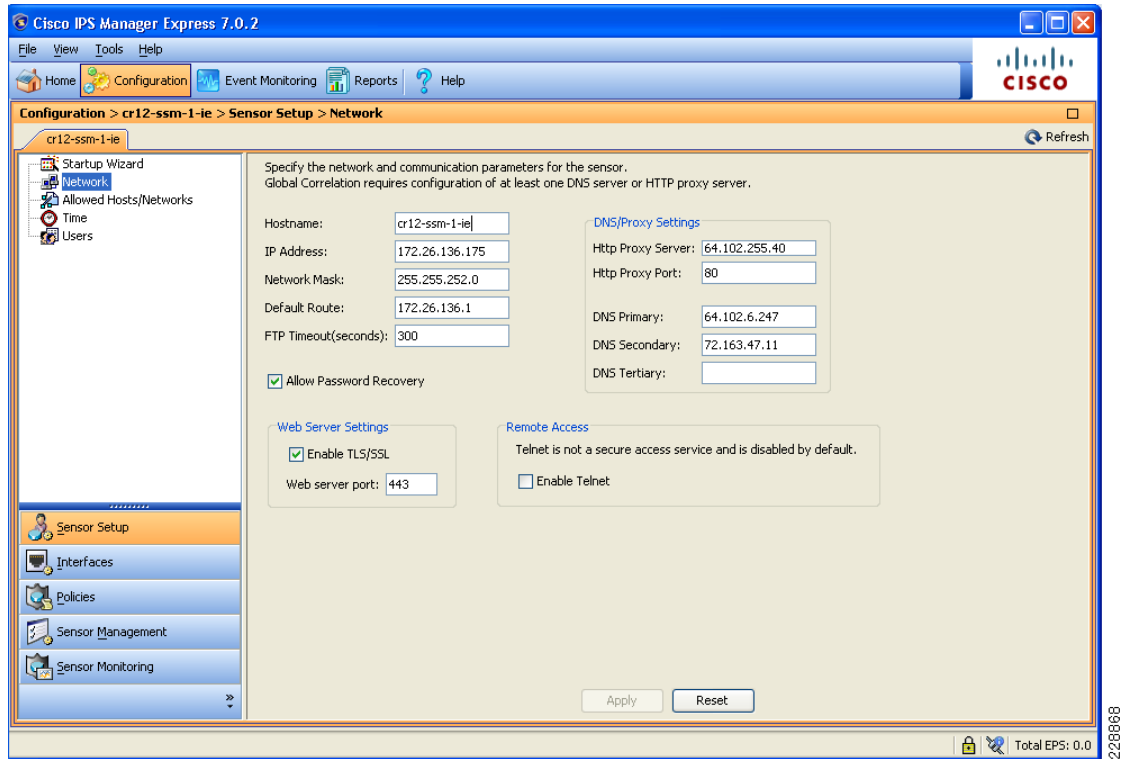
```
access-list IPS extended permit ip any any
class-map ips_class
  match access-list IPS
policy-map ips_policy
  class ips_class
    ips inline fail-close
service-policy ips_policy global
```

### IPS Global Correlation Deployment

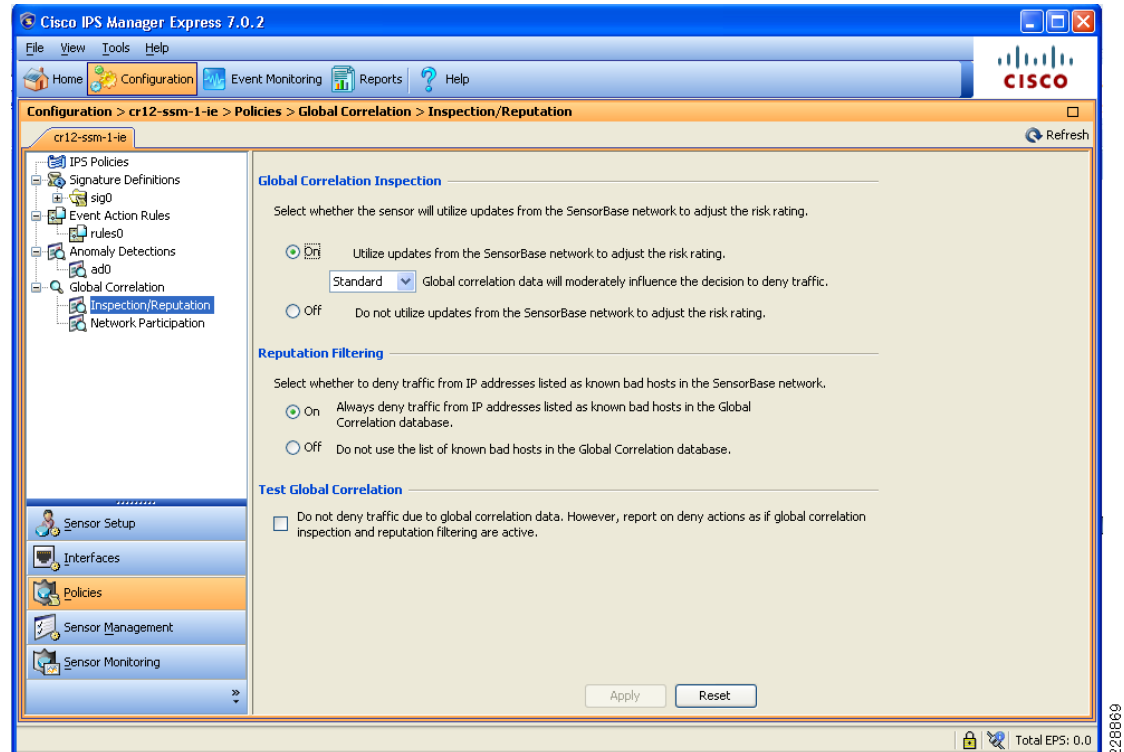
Before configuring IPS Global Correlation, be sure that you are using Cisco IPS Version 7.0 with the latest patch and signature updates and that Cisco IPS is configured for network connectivity in either IDS or IPS mode.

The configuration of Global Correlation can be performed using the command-line interface (CLI), Cisco IDS Device Manager (IDM), Cisco IME, or Cisco Security Manager. The following screenshots from Cisco IME illustrate the basic steps in the configuration of the IPS Global Correlation.

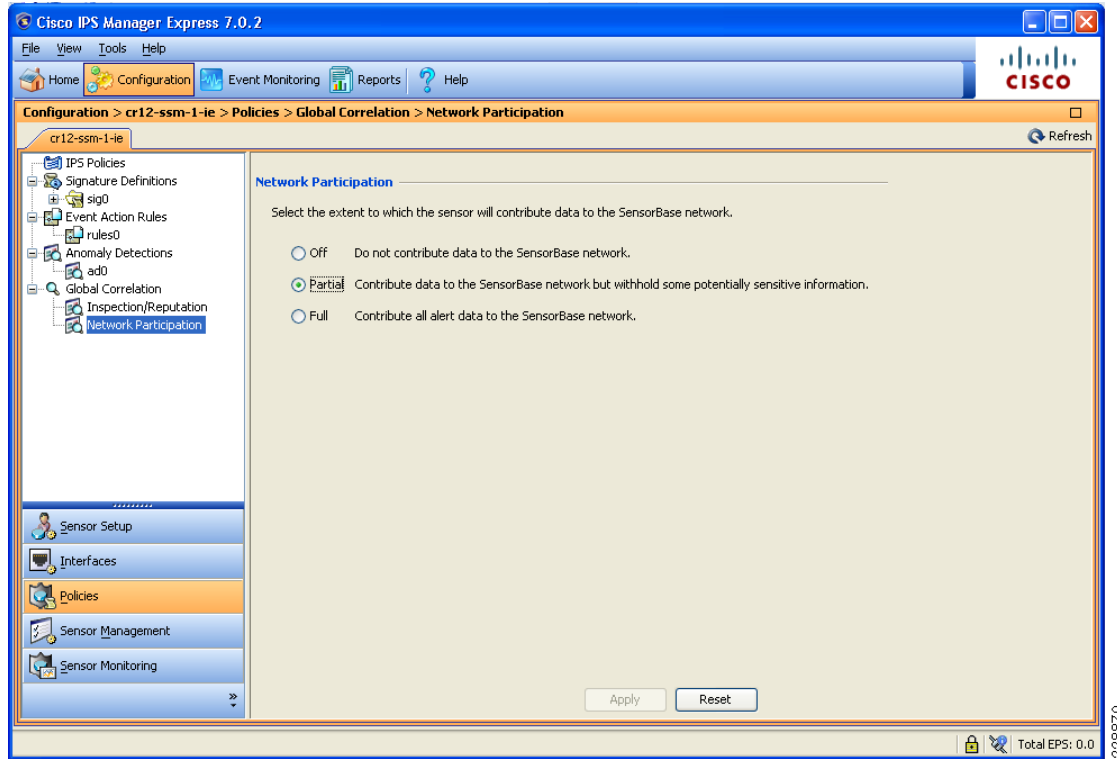
The first step in configuring the IPS sensor (module or appliance) to use Global Correlation is to add either a DNS address and/or the proxy server setup. This step enables a connection to Cisco SensorBase and is illustrated in [Figure 5-27](#). After you configure the DNS and proxy settings, the Global Correlation settings goes into effect as soon as the sensor has downloaded the latest Global Correlation updates.

**Figure 5-27 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, as illustrated in [Figure 5-28](#). A good practice is to configure Global Correlation Inspection initially in permissive mode while monitoring the effects, and then later change the configuration to Standard or Aggressive mode as desired.

**Figure 5-28 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 realtime and worldwide visibility into threat activity, which accelerates the identification and mitigation of threats propagating throughout the Internet. For this reason, it is recommended to configure the IPS sensors with partial or full network participation. Network participation configuration is illustrated in [Figure 5-29](#).

**Figure 5-29 Network Participation Settings (Off by Default)**

## Event Monitoring with Global Correlation

Event monitoring with IPS 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 5-30](#) shows Cisco IPS events with reputation scores in Cisco IME.



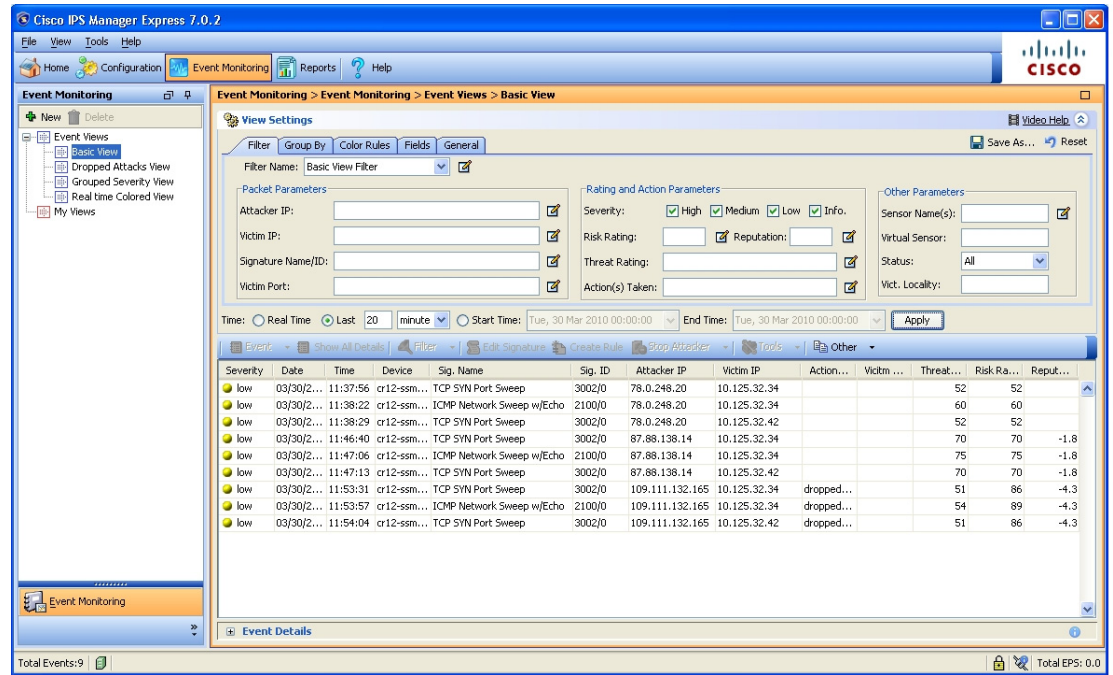
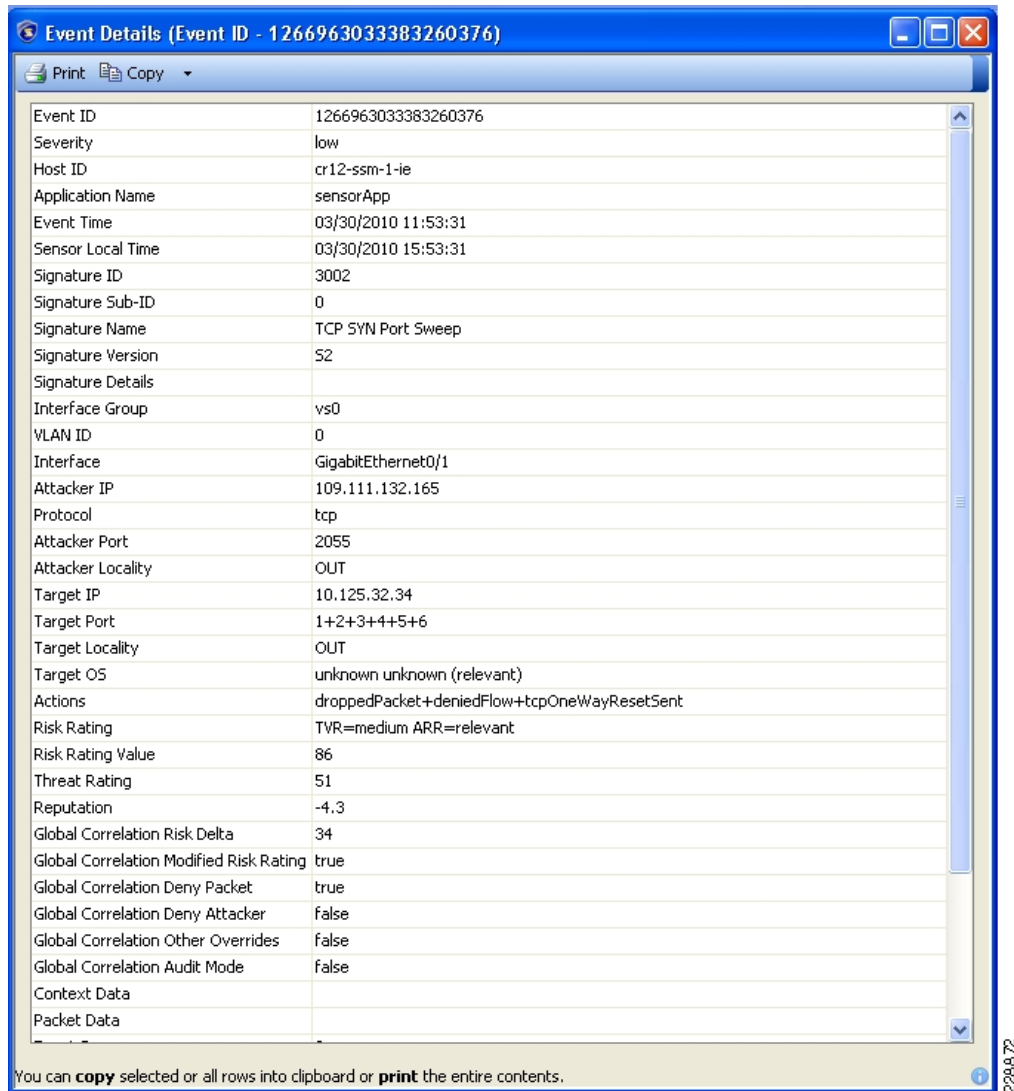
**Figure 5-30 Event Monitoring with Global Correlation in Cisco IME**

Figure 5-30 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, elevating 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 of -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 5-31 illustrates the detail view of the TCP SYN Port Sweep event coming from the attacker with a negative reputation of -4.3.

**Figure 5-31** Detailed View of a TCP SYN Port Sweep from an Attacker with a Negative Reputation Score

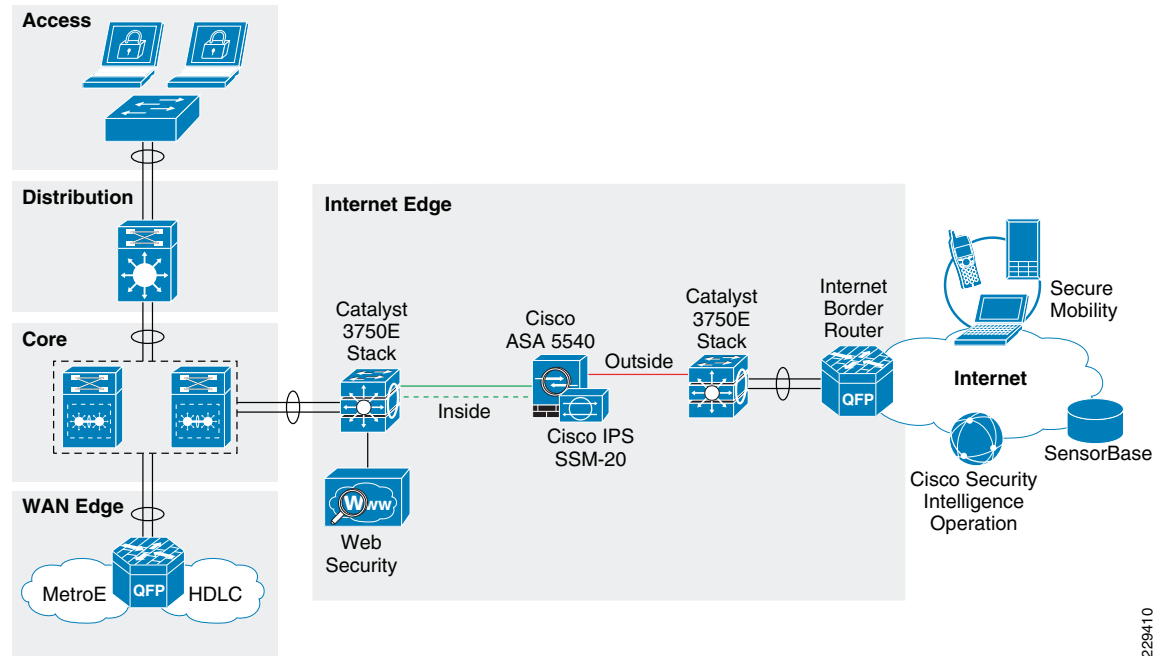


| Event Details (Event ID - 1266963033383260376) |                                             |
|------------------------------------------------|---------------------------------------------|
| Event ID                                       | 1266963033383260376                         |
| Severity                                       | low                                         |
| Host ID                                        | cr12-ssm-1-ie                               |
| Application Name                               | sensorApp                                   |
| Event Time                                     | 03/30/2010 11:53:31                         |
| Sensor Local Time                              | 03/30/2010 15:53:31                         |
| Signature ID                                   | 3002                                        |
| Signature Sub-ID                               | 0                                           |
| Signature Name                                 | TCP SYN Port Sweep                          |
| Signature Version                              | S2                                          |
| Signature Details                              |                                             |
| Interface Group                                | vs0                                         |
| VLAN ID                                        | 0                                           |
| Interface                                      | GigabitEthernet0/1                          |
| Attacker IP                                    | 109.111.132.165                             |
| Protocol                                       | tcp                                         |
| Attacker Port                                  | 2055                                        |
| Attacker Locality                              | OUT                                         |
| Target IP                                      | 10.125.32.34                                |
| Target Port                                    | 1+2+3+4+5+6                                 |
| Target Locality                                | OUT                                         |
| Target OS                                      | unknown unknown (relevant)                  |
| Actions                                        | droppedPacket+deniedFlow+tcpOneWayResetSent |
| Risk Rating                                    | TVR=medium ARR=relevant                     |
| Risk Rating Value                              | 86                                          |
| Threat Rating                                  | 51                                          |
| Reputation                                     | -4.3                                        |
| Global Correlation Risk Delta                  | 34                                          |
| Global Correlation Modified Risk Rating        | true                                        |
| Global Correlation Deny Packet                 | true                                        |
| Global Correlation Deny Attacker               | false                                       |
| Global Correlation Other Overrides             | false                                       |
| Global Correlation Audit Mode                  | false                                       |
| Context Data                                   |                                             |
| Packet Data                                    |                                             |

You can **copy** selected or all rows into clipboard or **print** the entire contents.

## Web Security Deployment

The Medium Enterprise Design Profile implements a Cisco IronPort WSA at the Internet edge distribution layer at the main site, as illustrated in [Figure 5-32](#). The WSA is located at the inside of the Cisco ASA acting as the Internet firewall. Deploying the WSA at the Internet edge distribution layer gives the WSA complete visibility on the traffic before getting out to the Internet through the firewall.

**Figure 5-32 WSA Deployment**

229410

The following subsections provides 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 the initial system configuration such as network and security settings. Note that running the initial System Setup Wizard completely reconfigures the WSA appliance and resets the administrator password. Use the System Setup Wizard only 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.

## Interface and Network Configuration

As part of the initial setup of the WSA, the following steps need to be completed:


1. Configuring network interfaces
2. Adding routes
3. Configuring DNS
4. Setting time

These settings are configured as part of an initial setup using the system setup wizard, but can be later modified using the WSA web-based GUI.

## Configuring Network Interfaces

Independent of the model, all Cisco IronPort WSA appliances are equipped with six Ethernet interfaces as shown in [Figure 5-33](#).

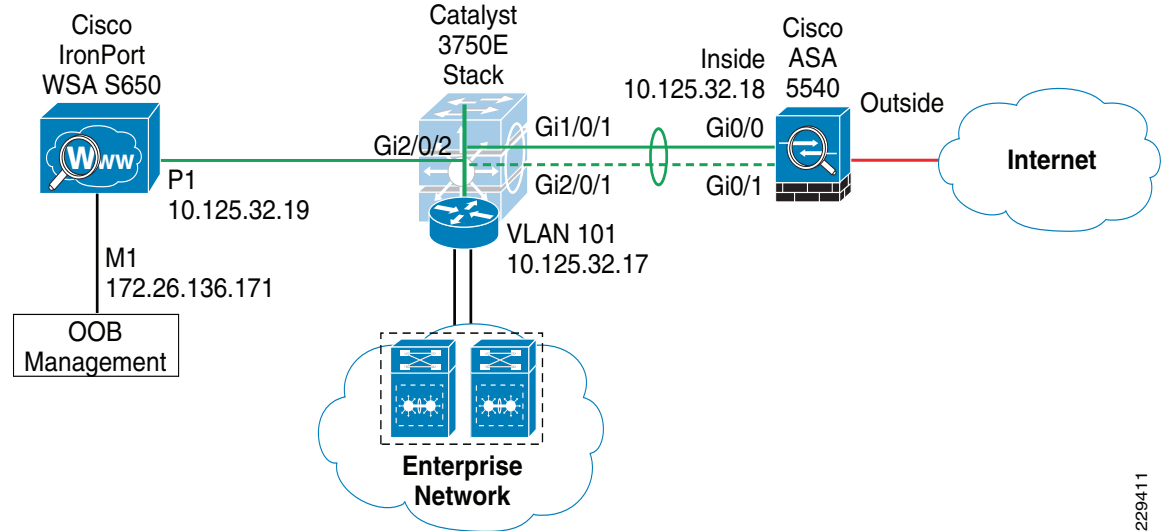
**Figure 5-33 WSA Interfaces**



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 Medium 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 Medium Enterprise Design Profile. P1 connects to the inside subnet of the firewall.
- *L4 Traffic Monitor (L4TM)*—T1 and T2 are the L4TM interfaces. These ports are used to capture traffic for inspection using either SPAN on a switch or a network tap. L4TM was not validated as part of the Medium Enterprise Design Profile, Cisco IPS Global Correlation and the Cisco ASA Botnet Traffic Filter features were used instead. For more information on L4TM, refer to the WSA configuration guides.

[Figure 5-34](#) illustrates the network topology for the WSA design in the validation lab.

**Figure 5-34 WSA Network Topology**

229411

Figure 5-35 shows the IP address and hostname configurations for the interfaces used within the WSA web-based GUI. In this case, an out-of-band management network is used where 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 5-35 WSA Interface Configuration****Interfaces**

| Interfaces                                |                                                                             |                |                 |                    |
|-------------------------------------------|-----------------------------------------------------------------------------|----------------|-----------------|--------------------|
| Interfaces:                               | Ethernet Port                                                               | IP Address     | Netmask         | Hostname           |
|                                           | M1                                                                          | 172.26.136.171 | 255.255.252.0   | ironport.cisco.com |
|                                           | P1                                                                          | 10.125.32.19   | 255.255.255.240 | ironport.cisco.com |
| Separate Routing for Management Services: | Separate routing (M1 port restricted to appliance management services only) |                |                 |                    |
| Appliance Management Services:            | HTTP on port 8080, HTTPS on port 8443, Redirect HTTP request to HTTPS       |                |                 |                    |
| L4 Traffic Monitor Wiring:                | Duplex TAP: T1 (In/Out)                                                     |                |                 |                    |
|                                           |                                                                             |                |                 | Edit Settings...   |

228875

**Adding Routes**

A default route is defined for management traffic pointing to the OOB management default gateway (172.26.136.1). A separate default route is defined for the data traffic pointing to the inside IP address of the firewall (10.125.32.18). As all internal networks are reachable through the Internet edge distribution switch, a route to 10.0.0.0/8 is defined pointing to the switch IP address (10.125.32.17) to allow the WSA to communicate with the clients directly. These settings are illustrated in Figure 5-36.

**Figure 5-36 WSA Route Configuration****Routes**

**Routes for Management Traffic (Interface M1: 172.26.136.171, Interface P1: 10.125.32.19)**

Add Route... Save Route Table... Load Route Table...

| Name          | Destination Network | Gateway      | All<br>Delete |
|---------------|---------------------|--------------|---------------|
| Default Route | All Others          | 172.26.136.1 |               |

Delete

**Routes for Data Traffic (Interface P1: 10.125.32.19)**

Add Route... Save Route Table... Load Route Table...

| Name          | Destination Network             | Gateway      | All<br>Delete |
|---------------|---------------------------------|--------------|---------------|
| Default Route | All Others (Including External) | 10.125.32.18 |               |
| Internal-10   | 10.0.0.0/8                      | 10.125.32.17 |               |

Delete

228876

**Configuring DNS**

The initial setup requires the configuration of a host name for the WSA appliance, and defining the DNS servers. [Figure 5-37](#) shows the DNS configuration.

**Figure 5-37 WSA DNS Configuration****DNS**

**DNS Server Settings**

|                                             |                        |               |
|---------------------------------------------|------------------------|---------------|
| DNS Servers:                                | Use these DNS Servers: |               |
|                                             | Priority               | IP Address    |
|                                             | 0                      | 10.125.31.2   |
|                                             | 0                      | 68.238.112.12 |
| Routing Table for DNS traffic:              | Data                   |               |
| Wait Before Timing out Reverse DNS Lookups: | 20 seconds             |               |
| DNS Domain Search List:                     | None                   |               |

Clear DNS Cache Edit Settings...

228877

**Setting Time**

Time synchronization is critical for forensic analysis and troubleshooting, therefore enabling NTP is highly recommended. [Figure 5-38](#) shows how the WSA is configured to synchronize its clock with an NTP server located on the OOB management network.

**Figure 5-38 WSA NTP Configuration****Time Settings**

| Time Keeping Method: Using NTP Servers: |                |
|-----------------------------------------|----------------|
| 1                                       | 172.26.129.252 |

Routing Table for NTP Server Queries: Management

Edit Settings...

**Note**

If Internet access is provided by an upstream proxy, the WSA must be configured to use the proxy for component updates and system upgrades. For information on configuring upstream proxies for upgrades, see the WSA configuration guides located at the following URL: [http://www.cisco.com/en/US/docs/security/wsa/wsa6.3/user\\_guide/WSA\\_6.3.0\\_GA\\_UserGuide.pdf](http://www.cisco.com/en/US/docs/security/wsa/wsa6.3/user_guide/WSA_6.3.0_GA_UserGuide.pdf).

## WCCP Transparent Web Proxy

The configuration of the WCCP Transparent Web Proxy includes the following steps:

1. Defining WSA WCCP Service Group
2. Enabling WSA Transparent Redirection
3. Enabling WCCP redirection on the Cisco ASA
4. Enabling WSA HTTPS scanning

### 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 settings include the following:

- *HTTP Ports to Proxy*—Lists 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.

Figure 5-39 illustrates the Web Proxy settings.

**Figure 5-39 WSA Proxy Settings****Proxy Settings**

| Web Proxy Settings                   |                                                      |
|--------------------------------------|------------------------------------------------------|
| <b>Basic Settings</b>                |                                                      |
| Proxy:                               | Enabled                                              |
| HTTP Ports to Proxy:                 | 80, 3128                                             |
| Caching:                             | Enabled <a href="#">Clear Cache</a>                  |
| Proxy Mode:                          | Transparent                                          |
| IP Spoofing:                         | Not Enabled                                          |
| <b>Advanced Settings</b>             |                                                      |
| Persistent Connection Timeout:       | Client Side: 300 Seconds<br>Server Side: 300 Seconds |
| In-Use Connection Timeout:           | Client Side: 300 Seconds<br>Server Side: 300 Seconds |
| Simultaneous Persistent Connections: | Server Maximum Number: 2000                          |
| Headers:                             | X-Forwarded-For: Do Not Send<br>VIA: Send            |
| <a href="#">Edit Settings...</a>     |                                                      |

228879

**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 Catalyst Internet Edge Distribution Switch. Use MD5 authentication to protect the WCCP communication between the WSA and Cisco Catalyst Switch. [Figure 5-40](#) shows an example.

**Figure 5-40 WSA Transparent Proxy**

| WCCP v2 Service                                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service Profile Name:                                                                           | <input type="text" value="web-https-cache"/>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Service:                                                                                        | <input type="radio"/> Standard service ID: 0 web-cache (destination port 80)<br><input checked="" type="radio"/> Dynamic service ID: <input type="text" value="10"/> 0-255<br>Port numbers: <input type="text" value="80,443"/><br><i>(up to 8 port numbers, separated by commas)</i><br><input checked="" type="radio"/> Redirect based on destination port<br><input type="radio"/> Redirect based on source port (return path)<br><i>For IP spoofing, define two services, one based on destination port and another based on source port (return path).</i><br><input checked="" type="radio"/> Load balance based on server address<br><input type="radio"/> Load balance based on client address<br><i>Applies only if more than one Web Security Appliance is in use.</i> |
| Router IP Addresses:                                                                            | <input type="text" value="10.125.32.17"/><br><i>Separate multiple entries with line breaks or commas.</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Router Security:                                                                                | <input checked="" type="checkbox"/> Enable Security for Service<br>Password: <input type="password" value="....."/><br>Confirm Password: <input type="password" value="....."/>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <a href="#">Advanced:</a> Optional settings for customizing the behavior of the WCCP v2 Router. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

228880



## Enabling WCCP Redirection on Catalyst 3750E Distribution Switch

The configuration of WCCP on the Cisco Catalyst 3750 switch requires the following components:

- A group-list indicating the IP addresses of the WSA appliances that are members 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.
- WCCP service indicating the service ID. The service ID defined on the Catalyst switch must be the same as the service ID defined on the WSAs. Use a password for MD5 authentication.
- Enabling WCCP redirection on an interface. Apply the WCCP service on the Internet Edge distribution switch interface facing the Core switch.

The following is a Cisco Catalyst switch configuration example:

```
! Group-list defining the IP addresses of all WSAs
ip access-list standard wsa-farm
 permit 10.125.32.19
!
! Redirect-list defining what ports and hosts/subnets should be redirected
ip access-list extended proxylist
 permit tcp 10.0.0.0 0.255.255.255 any eq www
 permit tcp 10.0.0.0 0.255.255.255 any eq 443
!
! Configure WCCP service
ip wccp 10 redirect-list proxylist group-list wsa-farm password <MD5-password>
!
! Apply WCCP on an interface
interface Port-channel1
ip wccp 10 redirect in
!
```

The WCCP connection status and configuration can be monitored on the Cisco Catalyst 3750 switch with the **show ip wccp** command. An example is provided below:

```
cr12-3750s-ie# show ip wccp
Global WCCP information:
  Router information:
    Router Identifier:          10.125.200.23
    Protocol Version:          2.0

  Service Identifier: 10
    Number of Service Group Clients: 1
    Number of Service Group Routers: 1
    Total Packets s/w Redirected: 0
      Process: 0
      CEF: 0
    Redirect access-list: proxylist
    Total Packets Denied Redirect: 0
    Total Packets Unassigned: 5
    Group access-list: wsa-farm
    Total Messages Denied to Group: 0
    Total Authentication failures: 0
    Total Bypassed Packets Received: 0

cr12-3750s-ie#
```

**Note**

Cisco Catalyst 3750 switches support switching in hardware only at Layer 2; therefore, no counters increment when the **show ip wccp** command is issued on the switch.

## 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 5-41.

**Figure 5-41 WSA HTTPS Proxy**

### HTTPS Proxy

| HTTPS Proxy Settings                  |                                                                                                                                                                                                                 |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HTTPS Proxy:                          | Enabled                                                                                                                                                                                                         |
| Transparent HTTPS Ports to Proxy:     | 443                                                                                                                                                                                                             |
| Root Certificate and Key for Signing: | Using Generated Certificate:<br>Common name: Cisco Systems, Inc<br>Organization: CMO<br>Organizational Unit: ESE<br>Country: US<br>Expiration Date: Oct 14 16:30:47 2010 GMT<br>Basic Constraints: Not Critical |
| Invalid Certificate Handling:         | Expired: Monitor<br>Mismatched Hostname: Monitor<br>Unrecognized Root Authority: Monitor<br>All other error types: Monitor                                                                                      |

2298861

[Edit Settings...](#)

**Note**

In cases where Internet access is handled by upstream proxies, you must configure the WSA to route through the upstream proxies. For information regarding the configuration of upstream proxies, see the Cisco IronPort WSA configuration guide located at the following URL:  
[http://www.cisco.com/en/US/docs/security/wsa/wsa6.3/user\\_guide/WSA\\_6.3.0\\_GA\\_UserGuide.pdf](http://www.cisco.com/en/US/docs/security/wsa/wsa6.3/user_guide/WSA_6.3.0_GA_UserGuide.pdf)

## 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, and so on) and URL categories users can 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 5-42 shows the global policy.

**Figure 5-42 Global Access Policy****Access Policies**

| Policies      |                                       |                                                                                           |                                                                                |                                                               |                                           |        |
|---------------|---------------------------------------|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|---------------------------------------------------------------|-------------------------------------------|--------|
| Add Policy... |                                       |                                                                                           |                                                                                |                                                               |                                           |        |
| Order         | Group                                 | Applications                                                                              | URL Categories                                                                 | Objects                                                       | Web Reputation and Anti-Malware Filtering | Delete |
|               | <b>Global Policy</b><br>Identity: All | Allow: FTP over HTTP, HTTP, Native FTP<br>Block: User Agents<br>Allow: Ports 8080, 21,... | Redirect: 0<br>Allow: 0<br>Monitor: 51<br>Warn: 0<br>Block: 3<br>Time-Based: 0 | HTTP/HTTPS Object Max Size: None<br>FTP Object Max Size: None | (enabled)                                 |        |

228882

URL categories corresponding to non-business related content should be blocked in compliance with the company's Internet access policies. Figure 5-43 provides an example on how the Adult/Sexually Explicit and Chat categories are blocked.

**Figure 5-43 URL Categories****Access Policies: URL Categories: Global Policy**

| Custom URL Category Filtering                                                                                           |              |            |            |                 |
|-------------------------------------------------------------------------------------------------------------------------|--------------|------------|------------|-----------------|
| No Custom URL Categories are defined. Add categories in the Custom URL Categories page.                                 |              |            |            |                 |
| Predefined URL Category Filtering                                                                                       |              |            |            |                 |
| These URL Categories are defined as group membership criteria. All other categories are not applicable for this policy. |              |            |            |                 |
| Category                                                                                                                | Monitor<br>🔍 | Warn<br>⚠️ | Block<br>🚫 | Time-Based<br>🕒 |
| Adult/Sexually Explicit                                                                                                 | Select all   | Select all | Select all | (Unavailable)   |
| Advertisements & Popups                                                                                                 | ✓            |            |            | —               |
| Alcohol & Tobacco                                                                                                       | ✓            |            |            | —               |
| Arts                                                                                                                    | ✓            |            |            | —               |
| Blogs & Forums                                                                                                          | ✓            |            |            | —               |
| Business                                                                                                                | ✓            |            |            | —               |
| Chat                                                                                                                    |              |            | ✓          | —               |
| Computing & Internet                                                                                                    | ✓            |            |            | —               |

228883

## Catalyst Integrated Security Features Deployment

Within the Medium Enterprise Design Profile, the Cisco Catalyst Integrated Security Features (CISF) were implemented in the access layer switches. CISF is a set of native security features available on the Cisco Catalyst switches that protect the infrastructure and users from spoofing, man-in-the-middle (MITM), DoS, and other access layer attacks. The following configuration illustrates an example of the CISF configurations used on a Cisco 3750E switch in the Medium Enterprise Design Profile in the access layer.

```
! configure dhcp snooping on the access VLANs in global configuration mode
ip dhcp snooping vlan 101-113
no ip dhcp snooping information option
ip dhcp snooping
!
! configure arp inspection on the access VLANs in global configuration mode
ip arp inspection vlan 101-113
ip arp inspection validate src-mac dst-mac ip allow zeros
!
```

```

! configure the port recovery parameters for ports being disabled by dhcp snooping,
arp-inspection, or storm control
errdisable recovery cause dhcp-rate-limit
errdisable recovery cause storm-control
errdisable recovery cause arp-inspection
errdisable recovery interval 120
!
! configure port specific parameters on access ports
interface GigabitEthernet1/0/1
! configure port security parameters
switchport port-security
switchport port-security aging time 5
switchport port-security violation restrict
switchport port-security aging type inactivity
! configure arp inspection rate limiting
ip arp inspection limit rate 100
! configure storm control parameters
storm-control broadcast level pps 1k
storm-control multicast level pps 2k
storm-control action trap
! configure IP Source Guard parameters
ip verify source

```

**Note**

When deploying Cisco Catalyst 3000 switches in the access layer in a routed Layer 3 deployment, configuring IP Source Guard causes edge router ACLs and VLAN ACLs to be ineffective for blocking traffic. When IP Source Guard is enabled, it creates a port-based ACL to permit traffic only from IP addresses that were assigned via the DHCP server. On Cisco Catalyst 3000 switches, port-based ACLs overrides router and VLAN ACLs resulting in all traffic being permitted to all destinations.

## NAC Appliance Deployment

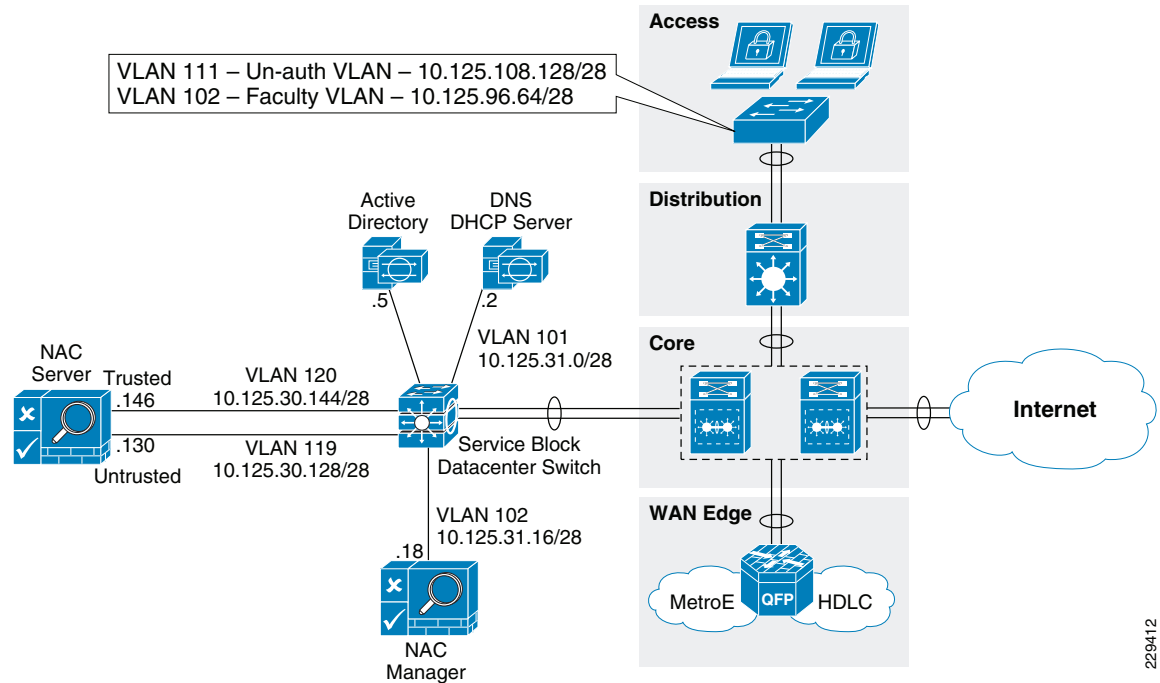
In the Medium Enterprise Design Profile, a NAC Appliance solution is deployed at all locations; the main site and each of the remote offices. A centralized NAC Manager is deployed at the main site and is deployed within the data center at that site. A NAC server is deployed at the main site and each of the remote site, and is connected within the service block connecting to the core switches at each of the sites.

The Medium Enterprise Design Profile provides host network connectivity using wired and wireless technologies. As such, the NAC Appliance solution must provide a solution for both connectivity options. For wireless clients, a Layer 2 OOB NAC solution is deployed; and for wired clients, a Layer-2 OOB or a Layer-3 OOB NAC solution may be deployed.

The following subsections provide configuration steps for configuring a Layer-3 OOB NAC solution for wired clients and a Layer-2 OOB NAC solution for wireless clients.

### NAC Deployment for Wired Clients

Within the Medium Enterprise Design Profile, a NAC Layer-3 OOB deployment using ACLs was used for the wired clients. [Figure 5-44](#) shows the L3 OOB logical network diagram that was used to validate NAC for wired clients in the Medium Enterprise Design Profile.

**Figure 5-44 NAC L3 OOB Logical Topology Diagram**

The following subsections illustrate the needed steps to configure a L3 Real-IP OOB NAC Deployment using ACLs.

### Configuring the Edge Access Switch for Enforcement

VLANs and edge ACLs are used on the access switches to restrict access to the network based on the NAC assigned user roles. The following configuration snippets provide sample configurations for two VLANs (Unauthenticated and Employee) and the associated edge ACLs. Edge ACLs and VLANs should be configured on all access switches that users are connecting to.

- *Unauthenticated role*—VLAN 111 and ACL Name: *nac-unauth-acl*

```
! create NAC unauthenticated VLAN
vlan 111
  name nac-unauth-vlan
! create SVI for unauthenticated VLAN
interface Vlan111
  ip address 10.125.108.129 255.255.255.192
  ip helper-address 10.125.31.2
!
! configure ACL for the unauthenticated role
ip access-list extended nac-unauth-acl
! allow Discovery packets from the NAC Agent to the NAC Server
  permit udp any host 10.125.30.130 eq 8906
! allow Discovery packets from the NAC Agent to the NAC Server for ADSSO
  permit udp any host 10.125.30.130 eq 8910
! allow web traffic from the PC to the NAC Server
  permit tcp any host 10.125.30.130 eq www
! allow SSL traffic from the PC to the NAC Server
  permit tcp any host 10.125.30.130 eq 443
! allow DHCP traffic to the DHCP server
  permit udp any host 255.255.255.255 eq bootps
  permit udp any host 10.125.31.2 eq bootps
! allow DNS traffic to the DNS Server
```

```

permit udp any host 10.125.31.2 eq domain
permit tcp any host 10.125.31.2 eq domain
! allow traffic to the remediation servers
permit tcp any host 12.120.79.206 eq www
permit tcp any host 12.120.10.243 eq www
permit tcp any host 12.120.11.243 eq www
permit tcp any host 12.120.78.208 eq www
permit tcp any host 216.151.177.81 eq ftp
!
! apply ACL to the Unauthenticated VLAN
interface Vlan111
ip access-group nac-unauth-acl in

```

- *Employee role*—VLAN 102 and ACL name: *employee-acl*

After the client is moved to this VLAN, if the native NAC Agent is used, it still attempts to discover the NAC Server. This NAC Agent behavior is by design. If the Agent is able to reach the NAC Server, the Agent pops up trying to perform the login process again, even though the client is already granted access. To prevent this, an ACL entry needs to be added to the ACL on the employee VLAN to prevent UDP 8906 Discovery packets originating from the Agent are dropped once the client is authenticated. The below configuration snippet illustrates the ACL entry needed to drop these discovery packets on the authenticated employee VLAN.

```

! create NAC employee VLAN
vlan 102
 name employee-vlan
! create SVI for employee VLAN
interface Vlan102
 ip address 10.125.96.65 255.255.255.192
 ip helper-address 10.125.31.2
! configure ACL for the employee role to prevent NAC Discovery packets from
! reaching NAC Server
ip access-list extended employee-acl
 deny udp any host 10.125.30.130 eq 8906
 permit ip any any
!
! apply ACL to the employee VLAN
interface Vlan102
 ip access-group employee-acl in

```

## NAC Manager and NAC Servers Initial Setup

The initial installation and configuration of the NAC Manager and NAC Server is performed via console access, and the install utility guides you through the initial configuration for both NAC Manager and NAC Server. See the following link to perform initial setup:

[http://www.cisco.com/en/US/docs/security/nac/appliance/installation\\_guide/hardware/47/hi\\_instal.html](http://www.cisco.com/en/US/docs/security/nac/appliance/installation_guide/hardware/47/hi_instal.html)

## Apply License to the NAC Manager

After performing the initial setup through the console, the rest of the configuration of the NAC Manager and Server is performed using the NAC Manager GUI. The first step is to upload the NAC Manager and Server licenses that came with the appliances. See the following URL for more detail on uploading the licenses:

[http://www.cisco.com/en/US/docs/security/nac/appliance/installation\\_guide/hardware/47/hi\\_instal.html#wp1113597](http://www.cisco.com/en/US/docs/security/nac/appliance/installation_guide/hardware/47/hi_instal.html#wp1113597)

## Update Policies from Cisco.com on the NAC Manager

The NAC Manager needs to be configured to retrieve periodic updates from the central update server located at Cisco. The Cisco NAC Appliance Supported AV/AS Product List is a versioned XML file distributed from a centralized update server that provides the most current matrix of supported anti-virus (AV) and antispyware (AS) vendors and product versions used to configure AV or AS Rules and AV or AS Definition Update requirements for posture assessment/remediation. This list is updated regularly for the AV/AS products and versions supported in each Agent release and include new products for new Agent versions. The list provides version information only. When the CAM downloads the Supported AV/AS Product List, it is downloading the information about what the latest versions are for AV/AS products; it is not downloading actual patch files or virus definition files. Based on this information, the agent can then trigger the native AV/AS application to perform updates. See the following URL for details on setting this up:

[http://www.cisco.com/en/US/docs/security/nac/appliance/configuration\\_guide/47/cam/m\\_agntd.html#wp1351880](http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/47/cam/m_agntd.html#wp1351880)

## Installing Certificates from Third-Party Certificate Authority (CA)

During installation, the initial configuration utility script for both the NAC Manager and NAC Server requires you to generate a temporary SSL certificate. For a lab environment, you may continue to use the self-signed CERTs. However, the self-signed CERTs are not recommended for a production network. For more information on installing certificates on the NAC Manager from a third-party CA, see the following URL:

[http://www.cisco.com/en/US/docs/security/nac/appliance/configuration\\_guide/47/cam/m\\_admin.html#wp1078189](http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/47/cam/m_admin.html#wp1078189)

For more information on installing certificates on the NAC Server from a third-party CA, see the following URL:

[http://www.cisco.com/en/US/docs/security/nac/appliance/configuration\\_guide/47/cas/s\\_admin.html#wp1040111](http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/47/cas/s_admin.html#wp1040111)



### Note

If you are using the self-signed certificates in the lab environment, the NAC Manager and NAC Server need to trust the certificate of each other, which requires you to upload each other's certificates as a **Trusted Certificate Authority** under **SSL > Trusted Certificate Authorities**.

## Adding the NAC Server to the NAC Manager

- Step 1** To add the NAC Server to the NAC Manager, from within the NAC Manager GUI, click **CCA Servers > New Server**.
- Step 2** Add the IP address of the NAC Server's **Trusted** interface, select **Out-of-Band Real-IP-Gateway** from the *Server Type* dropdown list, and click **Add Clean Access Server**. See [Figure 5-45](#).

**Figure 5-45** Adding the NAC Server to the NAC Manager

The screenshot shows the Cisco Clean Access Standard Manager interface (Version 4.7.2). The left sidebar contains navigation menus for Device Management (CCA Servers, Filters, Clean Access), OOB Management (Profiles, Devices), and User Management (User Roles, Auth Servers, Local Users). The main content area is titled 'Device Management > Clean Access Servers' and has three tabs: 'List of Servers', 'New Server', and 'Authorization'. The 'New Server' tab is active, displaying a form to add a new Clean Access Server. The form includes fields for 'Server IP Address' (10.125.30.146), 'Server Location' (Main Site DC), and 'Server Type' (Virtual Gateway). An 'Add Clean Access Server' button is at the bottom.

Once added, the NAC Server appears in the list.

**Note**

The NAC Manager and NAC Server have to trust each other's certificate authority (CA) for NAC Manager to successfully add the NAC server.

**Configure the NAC Server**

- Step 3** Click the **Manage** icon for the NAC Server to continue the configuration. See [Figure 5-46](#).

**Figure 5-46** NAC Server Managed by NAC Manager

The screenshot shows the Cisco Clean Access Standard Manager interface (Version 4.7.2) with the 'List of Servers' tab active. It displays a table of configured NAC servers.

| IP Address    | Type                        | Location        | Status    | Manage | Disconnect | Reboot | Delete |
|---------------|-----------------------------|-----------------|-----------|--------|------------|--------|--------|
| 10.125.30.146 | Out-of-Band Real-IP Gateway | Main Site DC    | Connected |        |            |        |        |
| 10.125.30.114 | Out-of-Band Virtual Gateway | Wireless L2 OOB | Connected |        |            |        |        |

- Step 4** After clicking the **Manage** icon, click the **Network** tab.

**Layer 3 Support**

- Step 5** To enable Layer 3 support for L3 OOB, check (enable) the options for the following:
- Enable L3 Support
  - Enable L3 strict mode to block NAT devices with Clean Access Agent
- Step 6** Click **Update** and reboot the NAC Server as instructed. [Figure 5-47](#).



Figure 5-47 NAC Server Network Details

The screenshot shows the Cisco Clean Access Standard Manager web interface, Version 4.7.2. The breadcrumb navigation is "Device Management > Clean Access Servers > 10.125.30.146". The left sidebar contains a tree view with categories: Device Management (CCA Servers, Filters, Clean Access), OOB Management (Profiles, Devices), User Management (User Roles, Auth Servers, Local Users), Monitoring (Summary, Online Users, Event Logs, SNMP), and Administration (CCA Manager). The main content area has tabs for Status, Network (selected), Filter, Advanced, Authentication, and Misc. Under the Network tab, there are sub-tabs for IP, DHCP, and DNS. The "Clean Access Server Type" is set to "Out-of-Band Real-IP Gateway". Checkboxes are present for "Enable L3 support", "Enable L3 strict mode to block NAT devices with NAC Agent", and "Enable L2 strict mode to block L3 devices with NAC Agent". The "Platform" is set to "APPLIANCE". Below this, there are two sections: "Trusted Interface (to protected network)" and "Untrusted Interface (to managed network)". Each section has fields for IP Address, Subnet Mask, and Default Gateway. The Trusted Interface fields are: IP Address (10.125.30.146), Subnet Mask (255.255.255.240), and Default Gateway (10.125.30.145). The Untrusted Interface fields are: IP Address (10.125.30.130), Subnet Mask (255.255.255.240), and Default Gateway (10.125.30.129). There are also checkboxes for "Set management VLAN ID" and "Pass through VLAN ID to managed network" for both interfaces. A note at the bottom states: "(Make sure the Clean Access Server is on VLAN n before you set its management VLAN ID to n.)". At the bottom right, there are "Update" and "Reboot" buttons. A vertical text "226867" is visible on the far right edge of the interface.

**Note**

Always generate the certificate for the NAC Server with the IP address of its *untrusted* interface. For name-based certificate, the name should resolve to the untrusted interface IP address. When the endpoint communicates with the untrusted interface of the NAC Server to begin the NAC process, the NAC Server will redirect the user to the certificate hostname or IP. If the certificate points to the trusted interface, the login process will not function correctly.

**Static Routes**

Once the NAC Server reboots, return to managing the NAC Server and continue the configuration. The NAC Server will need to communicate with endpoints on the unauthenticated VLAN with the untrusted interface.

- Step 1** Go to **Advanced > Static Routes** to add routes to the unauthenticated VLAN. Fill in the appropriate subnets for the unauthenticated VLANs and click **Add Route**. Be sure to select **untrusted interface [eth1]** for these routes. See [Figure 5-48](#).

**Figure 5-48** Adding Static Route to Reach the Unauthenticated User Subnet

The screenshot shows the Cisco Clean Access Standard Manager interface. The left sidebar contains navigation menus for Device Management, OOB Management, User Management, and Monitoring. The main content area is titled "Cisco Clean Access Standard Manager Version 4.7.2" and shows the path "Device Management > Clean Access Servers > 10.125.30.146". Below this, there are tabs for Status, Network, Filter, Advanced, Authentication, and Misc. The "Static Routes" tab is selected. The configuration fields are as follows:

| Field                     | Value                                                                                                                                       |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| Dest. Subnet Address/Mask | 10.125.108.128 / 26                                                                                                                         |
| Gateway (optional)        | 10.125.30.129<br><small>(gateway should be the address of an external gateway for the dest. subnet, not of the Clean Access Server)</small> |
| Link                      | Untrusted [eth1]                                                                                                                            |
| Description               | static route for unauthenticated users on cr22-4                                                                                            |

An "Add Route" button is located at the bottom right of the configuration area.

## Setup Profiles for Managed Switches in the NAC Manager

- Step 2** Each switch will be associated with a profile. Add a profile for each type of edge switch the NAC Manager will manage by going to **Profiles** and clicking on the **Device** tab. In the example shown in Figure 5-49, a Cisco Catalyst 4507 switch is added.

**Figure 5-49** SNMP Profile Used to Manage a Cisco Catalyst 4507 Switch

The screenshot shows the Cisco Clean Access Standard Manager interface. The left sidebar contains navigation menus for Device Management, OOB Management, User Management, and Monitoring. The main content area is titled "Cisco Clean Access Standard Manager Version 4.7.2" and shows the path "OOB Management > Profiles". Below this, there are tabs for Group, Device, Port, VLAN, and SNMP Receiver. The "Device" tab is selected. The configuration fields are as follows:

| Field                      | Value                             |
|----------------------------|-----------------------------------|
| Profile Name               | CR22_4507_LB                      |
| Device Model               | Cisco Catalyst 4000/4500 series   |
| SNMP Port                  | 161                               |
| Description                | Cisco 4507 L3 Access Switch in Ma |
| <b>SNMP Read Settings</b>  |                                   |
| SNMP Version               | SNMP V2C                          |
| Community String           | cisco123                          |
| <b>SNMP Write Settings</b> |                                   |
| SNMP Version               | SNMP V2C                          |
| Community String           | enterprise                        |

"Update" and "Reset" buttons are located at the bottom right of the configuration area.

## Switch Configuration for SNMP

- Step 3** The edge access switches should be configured for SNMP read/write community strings which are the same as those configured on the NAC Manager.

```
snmp-server community enterprise RW
snmp-server community cisco123 RO
```

## Configuring Port Profiles

**Step 4** For individual port control, configure a port profile under **OOB Management > Profiles > Ports** that includes the default unauthenticated VLAN and default access VLAN.

**Step 5** In the access VLAN section, specify the User Role VLAN. The NAC Manager changes the unauthenticated VLAN to the access VLAN based on the VLAN defined in the role where the user belongs.

The next step is to define the port profile to control the port's VLAN based upon User Roles and VLANs implemented.

In the example shown in [Figure 5-50](#), the Auth VLAN is the unauthenticated VLAN to which unauthenticated devices are initially assigned. The default access VLAN is the employee VLAN. This is used if the authenticated user does not have a role-based VLAN defined.

**Step 6** For the access VLAN, select **User Role VLAN** to map users to the VLAN configured in the user's role. The **Access VLAN** can override the default VLAN to a user role VLAN, which is defined under the **User Role**.

**Figure 5-50** Port Profile to Manage the Switch Port

The screenshot shows the Cisco Clean Access Standard Manager web interface. The left sidebar contains navigation menus for Device Management, OOB Management (with 'Profiles' highlighted), User Management, and Monitoring. The main content area is titled 'Cisco Clean Access Standard Manager Version 4.7.2' and shows the 'OOB Management > Profiles' configuration page. A breadcrumb trail indicates 'List > New'. The configuration form includes fields for Profile Name (Employee\_Port), Description (Access Switches in Large), and a checked box for 'Manage this port'. Under the 'VLAN Settings' section, there is a note about the supported VLAN Name format. The form contains four rows of settings: Auth VLAN (VLAN ID 111), Default Access VLAN (VLAN ID 102), Access VLAN (User Role VLAN), and VLAN Profile (Default).



### Note

You can also define VLAN names instead of IDs. This offers the flexibility of having different VLAN IDs on different switches across the site, but the same VLAN name attached to a particular Role.

Additional options are available under the port profile for IP release/renew options. If the user is behind an IP phone, then uncheck the option for bouncing the port, which will likely reboot the IP Phone when the port is bounced. See [Figure 5-51](#).

**Figure 5-51** Various Options Available under Port Profile

**Cisco Clean Access Standard Manager** Version 4.7.2

**Options: Device Connected to Port**

The CAM discovers the device connected to the switch port when it receives SNMP mac-notification or linkup traps for the device. The CAM then instructs the switch to assign the **Auth VLAN** to the port if the device is not certified, or **Access VLAN** if the device is certified and user is authenticated. You can additionally configure the following options:

- ☒ Change VLAN according to global device filter list (device must be in list).  
When set, the VLAN of the port will be assigned by global device filter settings (ALLOW=Default Access VLAN, DENY=Auth VLAN, ROLE/CHECK=User Role VLAN, IGNORE=ignore SNMP traps from managed switches (IP Phones)).
- ☒ Change to Auth VLAN if the device is certified but not in the out-of-band user list.  
Select the VLAN to assign when device is certified and user is reconnecting to network.
- ☐ Bounce the port after VLAN is changed.  
Check this box to help clients update their IP settings for non-Virtual Gateways. You can leave this field unchecked for Virtual Gateways.
- ☒ Bounce the port based on role settings after VLAN is changed.
- ☒ Generate event logs when there are multiple MAC addresses detected on the same switch port.

**Options: Device Disconnected from Port**

The device is considered disconnected after: SNMP linkdown trap received or admin removal of user. Additional configuration options are:

- ☒ Remove out-of-band online user when SNMP linkdown trap is received, and then change to Auth VLAN.  
Ensure Access VLAN client is removed from OOB online user list if disconnecting/reconnecting to same port.
- ☒ Remove other out-of-band online users on the switch port when a new user is detected on the same port.  
Ensure only one valid user is allowed on one switch port at the same time.
- ☐ Remove out-of-band online user without bouncing the port.  
This prevents port bouncing for IP phone connected users.

228891

## SNMP Receiver Setting

In addition to setting up the SNMP community string for Read/Write, you also need to configure the NAC Manager to receive SNMP traps from the switch. These traps are sent when the user connects and disconnects from the port. When the NAC Server sends the MAC/IP address information of a particular end point to the NAC Manager, the Manager is able to build a mapping table internally for MAC/IP and switch port. See [Figure 5-52](#).

**Figure 5-52** NAC Manager SNMP Receiver Setting to Collect SNMP Traps/Informs

**Cisco Clean Access Standard Manager** Version 4.7.2

OOB Management > Profiles

**SNMP Trap** · Advanced Settings

(Configure the SNMP daemon running on the Clean Access Manager. The device setup must match these settings to be able to send traps to the Clean Access Manager)

Trap Port on Clean Access Manager: 162

**SNMP V1 Settings**

Community String:

**SNMP V2c Settings**

Community String: NacTraps

**SNMP V3 Settings**

Security Method (Auth/Priv): No Auth No Priv

User Name:

User Auth:

User Priv:

228892

The switch needs to be configured to enable SNMP traps to be sent to the NAC Manager. In addition, it is recommended to increase the default switch CAM table entry flush timer to 1 hour per Cisco best practice recommendations for NAC OOB. This reduces the frequency of MAC notifications that are sent out from already connected devices to the NAC Manager. Having a source trap command ensures a consistent source address will be used to send out the traps.

```
! global applicable SNMP configurations
snmp-server trap-source Loopback0
snmp-server enable traps mac-notification change move threshold
snmp-server host 10.125.31.18 version 2c NacTraps
! interface specific configurations
mac-address-table aging-time 3600
```

You can optionally configure Linkup/Linkdown traps to send to the NAC Manager. They are only used in a deployment scenario where the end hosts are *not* connected behind an IP phone.

## Adding Switches as Devices in the NAC Manager

The switch profile created in the previous section will be used to add the managed switches.

- Step 1** Under the **Device Profile**, use the profile you created, but do not change the default port profile value when adding the switch. See [Figure 5-53](#).

**Figure 5-53 Adding Edge Switch in the NAC Manager to Control via SNMP**

The screenshot displays the Cisco Clean Access Standard Manager web interface, Version 4.7.2. The left sidebar shows a navigation menu with sections: Device Management (CCA Servers, Filters, Clean Access), OOB Management (Profiles, **Devices**), User Management (User Roles, Auth Servers, Local Users), and Monitoring (Summary, Online Users, Event Logs, SNMP). The main content area is titled 'OOB Management > Devices' and has tabs for 'Devices' and 'Discovered Clients'. Under the 'Devices' tab, there are links for 'List', 'New', and 'Search'. The 'New' form includes the following fields: Device Profile (dropdown menu showing 'CR22\_4507\_LB'), Device Group (dropdown menu showing 'default'), Default Port Profile (dropdown menu showing 'uncontrolled'), IP Addresses (text input field with '10.125.200.1'), and Description (text input field with 'Cisco 4507 L3 Access Sw'). At the bottom of the form are 'Add' and 'Reset' buttons. A vertical text '2288893' is visible on the right edge of the interface.

## Configuring Switch Ports for the Devices to be Managed by NAC

- Step 2** Once the switch is added to the NAC Manager, you can select the ports that you want to manage. See [Figure 5-54](#).

**Figure 5-54** Port Control Selection available for a Managed Switch

**Cisco Clean Access Standard Manager** Version 4.7.2

OOB Management > Devices > Switch[10.125.200.1]

**Config** **Ports**

**List** **Manage**

For trunk ports (blue background), the VLAN value refers to **trunk native VLAN**.  
For Private VLAN ports (green background), the VLAN value refers to **private secondary VLAN**.

Update Refresh Advanced >>

Search For: -- Select Field -- starts with Show

Ports/Page: 12 Ports 1-12 of 63 | First | Previous | Next | Last

| Name  | Index | Description           | Status | Bounce | Initial VLAN | Current VLAN | MAC Notif. | Client MAC | Profile                | Note |
|-------|-------|-----------------------|--------|--------|--------------|--------------|------------|------------|------------------------|------|
| Te2/2 | 4     | TenGigabitEthernet2/2 | ●      | ⚙      | 1            | 1            | X          | 🔍          | Default [uncontrolled] |      |
| Gi2/3 | 5     | GigabitEthernet2/3    | ●      | ⚙      | 1            | 111          | X          | 🔍          | Employee_Port          |      |
| Gi2/4 | 6     | GigabitEthernet2/4    | ●      | ⚙      | 1            | 111          | X          | 🔍          | Employee_Port          |      |
| Gi2/5 | 7     | GigabitEthernet2/5    | ●      | ⚙      | 1            | 111          | X          | 🔍          | Employee_Port          |      |
| Gi2/6 | 8     | GigabitEthernet2/6    | ●      | ⚙      | 1            | 111          | X          | 🔍          | Employee_Port          |      |
| Gi3/1 | 9     | GigabitEthernet3/1    | ●      | ⚙      | N/A          | 102          | ✓          | 🔍          | Employee_Port          |      |
| Gi3/2 | 10    | GigabitEthernet3/2    | ●      | ⚙      | 1            | 111          | X          | 🔍          | Employee_Port          |      |
| Gi3/3 | 11    | GigabitEthernet3/3    | ●      | ⚙      | 1            | 111          | X          | 🔍          | Employee_Port          |      |
| Gi3/4 | 12    | GigabitEthernet3/4    | ●      | ⚙      | 1            | 104          | X          | 🔍          | Default [uncontrolled] |      |

## Configuring User Roles

- Step 3** The next step is to configure the user roles and map the appropriate VLANs to these roles. The screenshots in [Figure 5-55](#) show the creation of the employee role for the employee clients. The VLANs were already created in the edge access switches that correspond to each role. Additional roles and VLAN can be created for more granular access control if desired.

**Figure 5-55** Creating Employee Role and Mapping it to VLAN 102

## Adding Users and Assigning to Appropriate User Role

For user authentication, a local user database can be defined on the NAC Manager. However, in environments where there is a large user base or pre-existing authentication servers, integrating NAC with external authentication servers using RADIUS, LDAP, Kerberos, and so on, is typically preferred. When using external authentication servers, users are mapped to a particular role via RADIUS or LDAP attributes. For information on configuring external authentication servers with NAC, see the following URL:

[http://www.cisco.com/en/US/docs/security/nac/appliance/configuration\\_guide/47/cam/m\\_auth.html](http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/47/cam/m_auth.html)

## Customizing User Login Page for Web Login

A default login page is already created in the NAC Manager. However, the login page can be customized to change the appearance of the web portal. For a NAC L3 OOB solution, it is important to download the ActiveX or Java component to the end client. This is done to perform the following:

- Fetch the MAC address of the client machine
- Perform IP address release/renew

To do this, Go to **Administration > User Pages**. Edit the page to make sure these options are enabled as shown in [Figure 5-56](#).

**Figure 5-56** User Page Settings for Web Login

The screenshot shows the Cisco Clean Access Standard Manager interface, Version 4.7.2. The left sidebar contains navigation menus for Device Management, OOB Management, User Management, Monitoring, and Administration. The main content area is titled 'Administration > User Pages' and shows the 'Login Page' configuration. The 'General' tab is selected, displaying settings for enabling the login page, VLAN ID, Subnet (IP/Mask), Operating System, Page Type, Page Description, and Web Client. Checkboxes are present for enabling the login page, using a web client to detect client MAC address and Operating System, and using a web client to release and renew IP address when necessary (OOB). An unchecked checkbox is for installing the DHCP Refresh tool into the Linux/macOS system directory. Buttons for 'Update', 'Cancel', and 'View' are at the bottom right.

**Cisco Clean Access Standard Manager** Version 4.7.2

Administration > User Pages

Login Page | File Upload | Guest Registration Page

List · Add · Edit

General | Content | Style

☒ Enable this login page

VLAN ID   
(separate multiple VLANs with a comma)

Subnet (IP/Mask)  /

Operating System

Page Type

Page Description

Web Client (ActiveX/Applet)

☒ Use web client to detect client MAC address and Operating System.

☒ Use web client to release and renew IP address when necessary (OOB).  
(Helps OOB client acquire new IP address after authentication without bouncing the switch port.)

☐ Install DHCP Refresh tool into Linux/macOS system directory.  
(Avoids root/admin password prompt to refresh the IP address for Linux/macOS clients when the web client is used to perform DHCP release and renew.)

Update Cancel View

## Customizing the Agent for the User Roles

The NAC Manager can be configured to make the Agent mandatory for any user role. The agent should be made mandatory for any role that you want to perform posture assessment prior to granting them access to the network. In the example in [Figure 5-57](#), the NAC Web Agent is made mandatory for the employee role.

**Figure 5-57** Agent Login Required for Employee Role

The screenshot shows the Cisco Clean Access Standard Manager interface, Version 4.7.2. The left sidebar contains navigation menus for Device Management, OOB Management, User Management, Monitoring, and Administration. The main content area is titled 'Device Management > Clean Access' and shows the 'Agent Login' configuration. The 'Agent Login' tab is selected, displaying settings for User Role, Operating System, and checkboxes for requiring the use of the Agent and the Cisco NAC Web Agent. Text areas for Agent Download Page Message and Cisco NAC Web Agent Launch Page Message are also visible. Buttons for 'Update', 'Cancel', and 'View' are at the bottom right.

**Cisco Clean Access Standard Manager** Version 4.7.2

Device Management > Clean Access

Certified Devices | General Setup | Network Scanner | Clean Access Agent | Updates

Web Login · Agent Login

User Role

Operating System   
(By default, 'ALL' settings apply to all client operating systems if no OS-specific settings are specified.)

☐ Require use of Agent (for Windows & Macintosh OSX only)  
Agent Download Page Message (or URL):

☒ Require use of Cisco NAC Web Agent (for Windows 7/2000/XP/Vista only)  
Cisco NAC Web Agent Launch Page Message (or URL):

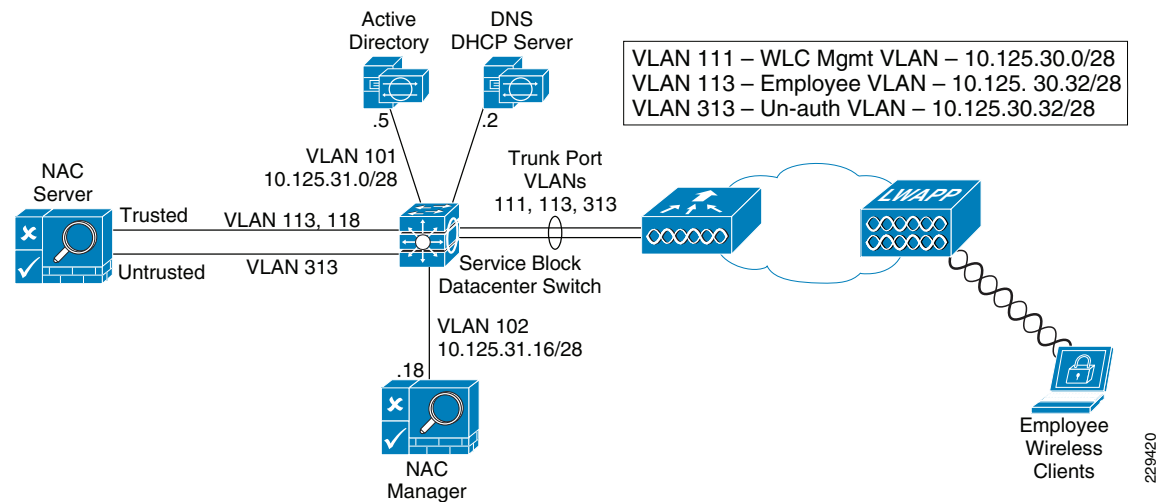
Update Cancel View



## NAC Deployment for Wireless Clients

Within the Medium Enterprise Design Profile, a NAC Layer-2 OOB deployment was used for wireless clients. Figure 5-58 shows the L2 OOB logical network diagram that was used to validate the NAC L2 OOB deployment in the Medium Enterprise Design Profile. Figure 5-58 shows the specifics for the employee wireless clients.

**Figure 5-58 Layer-2 OOB NAC Deployment Topology for Employee Wireless Clients**



As illustrated in Figure 5-58, the WLC is connected to a trunk port that carries the quarantine VLAN and access VLAN for the employee clients (VLANs 113 and 313). On the switch, the quarantine VLAN traffic is trunked to the NAC appliance, and the access VLAN traffic is trunked directly to the Layer 3 switch. Traffic that reaches the quarantine VLAN on the NAC appliance is mapped to access the VLAN based on static mapping configuration. When client associates and complete the L2 Auth, it checks whether the quarantine interface is associated; if yes, the data is sent on the quarantine interface. The client traffic that flows in the quarantine VLAN is trunked to the NAC appliance. After posture validation is done, the NAC server (CAS) sends an SNMP set message that updates the access VLAN ID to the controller, and the data traffic starts to switch from the WLC directly to the network without going through the NAC server.

The following subsections illustrate the configurations needed for deploying L2 OOB NAC for the Employee clients. Similar steps would be taken to enable NAC for additional clients/roles as needed.

### Catalyst Switch Configuration

The following Cisco Catalyst 3750E configuration example illustrates the configurations used on the service block switch in the Medium Enterprise Design Profile for the NAC wireless deployment:

```
interface GigabitEthernet2/0/9
  description Connected to cr25-nac-mgr-1
  switchport access vlan 102
  switchport mode access
  spanning-tree portfast
!
interface GigabitEthernet2/0/19
  description NAC Server trusted interface - Ethernet 0
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 802
  switchport trunk allowed vlan 113,118
  switchport mode trunk
```

```

!
interface GigabitEthernet2/0/20
description NAC Server untrusted interface - ethernet 1
switchport trunk encapsulation dot1q
switchport trunk native vlan 803
switchport trunk allowed vlan 313
switchport mode trunk
!
interface Port-channel11
description Connection to WLC cr23-5508-1
switchport trunk encapsulation dot1q
switchport trunk native vlan 801
switchport trunk allowed vlan 111,113,313
switchport mode trunk
switchport nonegotiate
!
interface Vlan111
description WLC Management VLAN
ip address 10.125.30.1 255.255.255.240
!
interface Vlan113
description Employee Client Subnet Access VLAN
ip address 10.125.30.33 255.255.255.240
!

```

## NAC OOB Configuration Steps on the WLC and NAC Manager

Perform the following steps to configure the WLC and the NAC Manager for a NAC L2 OOB deployment.

- Step 1** Enable SNMPv2 mode on the controller. See [Figure 5-59](#).

**Figure 5-59** Enabling SNMPv2

The screenshot shows the Cisco NAC Manager GUI. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh. The main menu includes MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT (highlighted), COMMANDS, HELP, and FEEDBACK. The left sidebar shows the Management section with a tree view including Summary, SNMP (expanded), General, SNMP V3 Users, Communities, Trap Receivers, Trap Controls, Trap Logs, HTTP, Telnet-SSH, Serial Port, Local Management Users, User Sessions, Logs, Mgmt Via Wireless, Software Activation, and Tech Support. The main content area displays the SNMP System Summary configuration page. It includes an Apply button and the following fields:

|                    |                      |
|--------------------|----------------------|
| Name               | cr23-5508-1          |
| Location           |                      |
| Contact            |                      |
| System Description | Cisco Controller     |
| System Object ID   | 1.3.6.1.4.1.9.1.1069 |
| SNMP Port Number   | 161                  |
| Trap Port Number   | 162                  |
| SNMP v1 Mode       | Disable              |
| SNMP v2c Mode      | Enable               |
| SNMP v3 Mode       | Disable              |

- Step 2** Create a profile for WLC on the NAC Manager. Click **OOB Management Profile > Device > New** from within the NAC Manager GUI. See [Figure 5-60](#).

**Figure 5-60** Creating Profile for WLC

**Cisco Clean Access Standard Manager** Version 4.7.2

OOB Management > Profiles

Group Device Port VLAN SNMP Receiver

List · New · Edit

(These settings must match the device setup to ensure that the Clean Access Manager can read/write to the device correctly)

Profile Name: WLC\_Main\_Site

Device Model: Cisco Wireless LAN Controllers

SNMP Port: 161

Description: Main Site WLC

**SNMP Read Settings**

SNMP Version: SNMP V2C

Community String: cisco123

**SNMP Write Settings**

SNMP Version: SNMP V2C

Community String: enterprise

Update Reset

229421

**Step 3** After the profile is created in the NAC Manager, add the WLC in the profile; go to **OOB Management > Devices > New** and enter the management IP address of WLC. See [Figure 5-61](#).

**Figure 5-61** Adding WLC in Profile

**Cisco Clean Access Standard Manager** Version 4.7.2

OOB Management > Devices

Devices Discovered Clients

List · New · Search

Device Profile: WLC\_Main\_Site

Device Group: default

IP Addresses: 10.125.30.2

Description: WLC 1 at Main Site

Add Reset

229422

**Step 4** Add the NAC Manager as the SNMP trap receiver in the WLC. Use the exact name of the trap receiver in the NAC Manager as the SNMP receiver. See [Figure 5-62](#).

**Figure 5-62 Adding MAC Manager as the SNMP Trap Receiver**

The screenshot shows the Cisco WLC Management interface. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh. The main menu has tabs for MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT (selected), COMMANDS, HELP, and FEEDBACK. On the left, the 'Management' sidebar shows a tree view with 'Summary' expanded, and 'SNMP' selected under it. The main content area is titled 'SNMP Trap Receiver > New'. It contains three fields: 'Community Name' with the value 'NacTraps', 'IP Address' with the value '10.125.31.18', and 'Status' with a dropdown menu set to 'Enable'. There are '< Back' and 'Apply' buttons at the top right of the form.

228903

- Step 5** Configure the SNMP trap receiver in the NAC Manager with the same name that was specified in the WLC controller; click **OOB Management > Profiles > SNMP Receiver**.

**Figure 5-63 Configure the SNMP Trap Receiver in the NAC Manager**

The screenshot shows the Cisco Clean Access Standard Manager interface, Version 4.7.2. The left sidebar has a tree view with categories: Device Management (CCA Servers, Filters, Clean Access), OOB Management (Profiles selected, Devices), User Management (User Roles, Auth Servers, Local Users), Monitoring (Summary, Online Users, Event Logs, SNMP), and Administration. The main content area is titled 'OOB Management > Profiles'. It has a tabbed interface with 'Group', 'Device', 'Port', 'VLAN', and 'SNMP Receiver' (selected). Under 'SNMP Receiver', there is a sub-tab 'Advanced Settings'. A note states: '(Configure the SNMP daemon running on the Clean Access Manager. The device setup must match these settings to be able to send traps to the Clean Access Manager)'. The configuration fields include: 'Trap Port on Clean Access Manager' (162), 'SNMP V1 Settings' (Community String), 'SNMP V2c Settings' (Community String: NacTraps), and 'SNMP V3 Settings' (Security Method: No Auth, No Priv; User Name, User Auth, User Priv). An 'Update' button is at the bottom right.

228904

At this stage, the WLC and the NAC Manager can talk to each other for client posture validation and access/quarantine state updates.

- Step 6** In the controller, create a dynamic interface with access and quarantine VLAN mapped to it. See [Figure 5-64](#).

**Figure 5-64** Creating Dynamic Interface in the Controller

The screenshot shows the Cisco WLC Controller configuration page for a dynamic interface. The left sidebar lists various configuration categories, with 'Advanced' selected. The main content area is divided into several sections:

- General Information:** Interface Name: staff data, MAC Address: 00:24:97:cf:3f:af.
- Configuration:** Guest Lan: ☐, Quarantine: ☒, Quarantine Vlan Id: 313.
- Physical Information:** The interface is attached to a LAG. Enable Dynamic AP Management: ☐.
- Interface Address:** VLAN Identifier: 113, IP Address: 10.125.30.34, Netmask: 255.255.255.240, Gateway: 10.125.30.33.
- DHCP Information:** Primary DHCP Server: 10.125.31.2.

Navigation links at the top include: Save Configuration, Ping, Logout, Refresh, MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, HELP, and FEEDBACK.

**Step 7** Create the WLAN and associate it with the dynamic interface. See [Figure 5-65](#).

**Figure 5-65** Creating the WLAN

The screenshot shows the Cisco WLC Controller configuration page for a new WLAN. The left sidebar lists various configuration categories, with 'Advanced' selected. The main content area is divided into several sections:

- General:** Profile Name: Staff Data, Type: WLAN, SSID: data, Status: ☒ Enabled.
- Security Policies:** [WPA2][Auth(802.1X + CCKM)] (Modifications done under security tab will appear after applying the changes.)
- Radio Policy:** All (dropdown), Interface: staff data (dropdown), Broadcast SSID: ☒ Enabled.

Navigation links at the top include: Save Configuration, Ping, Logout, Refresh, MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, HELP, and FEEDBACK.

**Step 8** Enable NAC in the WLAN on the WLC Controller. See [Figure 5-66](#).

**Figure 5-66 Enabling NAC in the WLAN on the WLC Controller**

The screenshot shows the Cisco WLC Controller configuration page for WLANs. The 'WLANs > Edit' page is open, and the 'Advanced' tab is selected. The 'NAC' section is expanded, showing the following settings:

- Allow AAA Override:** ☐ Enabled
- Coverage Hole Detection:** ☒ Enabled
- Enable Session Timeout:** ☒ 1800 (Session Timeout (secs))
- Aironet IE:** ☒ Enabled
- Diagnostic Channel:** ☐ Enabled
- IPv6 Enable:** ☐
- Override Interface ACL:** ☐ None
- P2P Blocking Action:** ☐ Disabled
- Client Exclusion:** ☒ Enabled 60 (Timeout Value (secs))
- Media Session Snooping:** ☐
- DHCP:**
  - DHCP Server:** ☐ Override
  - DHCP Addr. Assignment:** ☐ Required
- Management Frame Protection (MFP):**
  - Infrastructure MFP Protection:** ☒ (Global MFP Disabled)
  - MFP Client Protection:** ☐ Optional
- DTIM Period (in beacon intervals):**
  - 802.11a/n (1 - 255): 1
  - 802.11b/g/n (1 - 255): 1
- NAC:**
  - State:** ☒ Enabled

Buttons at the bottom include '< Back' and 'Apply'.

- Step 9** Add the client subnet in the CAS server as the managed subnet by clicking **CAS server > Select your CAS server > Manage > Advanced > Managed Subnets**. Add an unused IP address from the client subnet and put the quarantine VLAN (untrusted VLAN) for the managed subnet. See Figure 5-67.

**Figure 5-67 Adding the Client Subnet in the CAS Server**

The screenshot shows the Cisco Clean Access Standard Manager configuration page. The 'Device Management > Clean Access Servers > 10.125.30.114' page is open, and the 'Advanced' tab is selected. The 'Managed Subnet' section is expanded, showing the following settings:

- Enable subnet-based VLAN retag:** ☐ Update
- IP Address:** 10.125.30.36
- Subnet Mask:** 255.255.255.240
- VLAN ID:** 313 (-1 for non-VLAN)
- Description:** Staff Wireless Subnet (PEAP)

Buttons at the bottom include 'Add Managed Subnet'.

- Step 10** Create VLAN mappings on the CAS. Click **CAS server > Select your CAS server > Manage > Advanced > VLAN Mapping**. Add the access VLAN as trusted and the quarantine VLAN as untrusted. See Figure 5-68.

Figure 5-68 Creating VLAN Mappings

**Cisco Clean Access Standard Manager** Version 4.7.2

Device Management > Clean Access Servers > 10.125.30.114

Status Network Filter Advanced Authentication Misc

Managed Subnet · **VLAN Mapping** · Static Routes · ARP · Proxy

**VLAN Packet Handling**

☒ Enable VLAN Pruning  
When enabled along with VLAN Mapping, disallows any VLAN Packet to pass through to other interface in either direction if VLAN mapping cannot be done for the packet. If enabled alone, discards all VLAN packets from passing through in either direction.

☒ Enable VLAN Mapping

**VLAN Mapping Assignments**

Untrusted network VLAN ID  (-1 for non-VLAN)

Trusted network VLAN ID  (-1 for non-VLAN)

Description

| Untrusted VLAN ID | Trusted VLAN ID | Description                     | Del |
|-------------------|-----------------|---------------------------------|-----|
| 313               | 113             | Wireless Staff Users 313 -> 113 | X   |

229426

## Configuring Single SignOn (SSO) with the OOB Wireless Solution

The following is required to enable VPN SSO for a wireless NAC OOB deployment:

- Enabling VPN authentication on the NAC server with the WLC defined as the VPN concentrator in the NAC appliance.
- Enabling RADIUS accounting on the WLC controller. The WLC that is defined in the NAC appliance must be configured to send RADIUS accounting records to the NAC appliance for each 802.1x/EAP WLAN that is a managed subnet in the NAC.

The following steps outline the needed configuration on the NAC Manager to enable SSO.

- Step 1** From the NAC Manager GUI, click **CAS server > Select your CAS server > Manage > Authentication > VPN Auth.** See [Figure 5-69](#).

Figure 5-69 NAC Manager Configuration—Enabling SSO

The screenshot shows the Cisco Clean Access Standard Manager interface. The left sidebar contains navigation menus for Device Management, OOB Management, User Management, and Monitoring. The main content area is titled "Cisco Clean Access Standard Manager Version 4.7.2" and shows the configuration path: Device Management > Clean Access Servers > 10.125.30.114. The "Authentication" tab is selected, showing the "VPN Auth" sub-tab. The configuration fields are as follows:

| Field                      | Value                               |
|----------------------------|-------------------------------------|
| Single Sign-On:            | <input checked="" type="checkbox"/> |
| Agent VPN Detection Delay: | 0 seconds (0 means no delay)        |
| Auto Logout:               | <input checked="" type="checkbox"/> |
| RADIUS Accounting Port:    | 1813                                |

An "Update" button is located at the bottom right of the configuration fields.

- Step 2** Select the **VPN Concentrators** tab to add a new entry for the WLC. Populate the entry fields for the WLC Management IP address and shared secret you want to use between the WLC and NAC server. See [Figure 5-70](#).

Figure 5-70 Adding New Entry for WLC

The screenshot shows the Cisco Clean Access Standard Manager interface. The left sidebar contains navigation menus for Device Management, OOB Management, User Management, and Monitoring. The main content area is titled "Cisco Clean Access Standard Manager Version 4.7.2" and shows the configuration path: Device Management > Clean Access Servers > 10.125.30.114. The "Authentication" tab is selected, showing the "VPN Concentrators" sub-tab. The configuration fields are as follows:

| Field                  | Value |
|------------------------|-------|
| Name:                  |       |
| IP Address:            |       |
| Shared Secret:         |       |
| Confirm Shared Secret: |       |
| Description:           |       |

An "Add VPN Concentrator" button is located below the fields. Below the button is a table showing the existing entries:

| VPN Concentrator | IP Address  | Description        | Del |
|------------------|-------------|--------------------|-----|
| cr23-5508        | 10.125.30.2 | WLC-1 at Main Site | X   |

- Step 3** For role mapping, add the new authentication server with type **vpn sso** under **User Management > Auth Servers**. See [Figure 5-71](#).



**Figure 5-71 Adding New Authentication Server for Role Mapping**

**Cisco Clean Access Standard Manager** Version 4.7.2

User Management > Auth Servers

Auth Servers | Lookup Servers | Mapping Rules | Auth Test | Accounting

List • New

Authentication Cache Timeout (seconds): 120

| Provider Name | Authentication Type | Description                | Mapping | Edit | Delete |
|---------------|---------------------|----------------------------|---------|------|--------|
| Local DB      | local               | Cisco local authentication |         |      |        |
| Cisco VPN     | vpn sso             | Single Sign-on             |         |      |        |

**Step 4** Click the **Mapping** icon and then add **Mapping Rule**. The mapping varies depending on the class attribute 25 value that WLC sends in the accounting packet. This attribute value is configured in the RADIUS server and varies based on the user authorization. In this example, the attribute value is *employee*, and it is placed in the *Wireless\_Employee* role. See [Figure 5-72](#).

**Figure 5-72 Mapping Class Attribute from WLC to User Roles**

**Cisco Clean Access Standard Manager** Version 4.7.2

User Management -> Auth Servers

Auth Servers | Lookup Servers | Mapping Rules | Auth Test | Accounting

Configure one or more conditions first using the Add/Save Condition form, then add or save the mapping rule to the selected Role using the Add/Save Mapping form. Note that if the mapping is not added or saved, conditions are not preserved.

Provider Name: Cisco VPN Priority: 1

Role Name: Wireless\_Employee Description: Wireless Employee mapping

Rule Expression: ( 0,25 equals employee )

Condition Type: VLAN ID Operator: equals

Property Name: VLAN ID Property Value:

VLAN IDs may not be available for mapping if there are multiple hops between the CAS and the VPN concentrator.

| # | Type      | Left Operand | Operator | Right Operand | Edit | Del |
|---|-----------|--------------|----------|---------------|------|-----|
| 1 | Attribute | 0,25         | equals   | employee      |      |     |

**Step 5** To configure VPN SSO on the Wireless LAN Controller, RADIUS accounting needs to be enabled and sent to the NAC Server. See [Figure 5-73](#).

**Figure 5-73 Enabling RADIUS Accounting for VPN SSO**

The screenshot shows the Cisco WLAN configuration interface. The left sidebar has a tree view with 'WLANs' expanded and 'Advanced' selected. The main content area is titled 'WLANs > Edit' and has tabs for 'General', 'Security', 'QoS', and 'Advanced'. The 'Advanced' tab is active, and within it, the 'AAA Servers' sub-tab is selected. A message states: 'Select AAA servers below to override use of default servers on this WLAN'. Below this, there are sections for 'Radius Servers', 'Local EAP Authentication', and 'Authentication priority order for web-auth user'. The 'Radius Servers' section contains two columns: 'Authentication Servers' and 'Accounting Servers'. Both columns have a 'Server 1' entry with 'IP:10.125.31.66, Port:1812' and 'IP:10.125.30.114, Port:1813' respectively, and 'Server 2' and 'Server 3' entries set to 'None'. The 'Local EAP Authentication' section has a 'Local EAP Authentication' checkbox that is unchecked. The 'Authentication priority order for web-auth user' section has a 'Not Used' button and an 'Order Used For Authentication' button. The top of the interface includes a navigation bar with 'MONITOR', 'WLANs', 'CONTROLLER', 'WIRELESS', 'SECURITY', 'MANAGEMENT', 'COMMANDS', 'HELP', and 'FEEDBACK'. The bottom right corner of the interface shows the number '228914'.

**Note**

When deploying a wireless NAC solution that requires single sign-on for some WLANs and non-single sign-on for other WLANs, RADIUS accounting must be disabled for the WLANs not requiring SSO. Otherwise, NAC mistakenly authenticates the non-single sign-on clients without prompting them.

## Additional Information

The configuration guidelines and examples in the previous sections were based on the features, devices, and designs that were used for validating the Medium Enterprise Design Profile. For more information on all of these features, refer to [Appendix A, “Reference Documents.”](#)



# APPENDIX A

## Reference Documents

| Document Title                                                                                                     | URL                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>CCVE SRA Solution Overview</i>                                                                                  | <a href="http://www.cisco.com/en/US/docs/solutions/Verticals/Education/srajrcollegesoverview.html">http://www.cisco.com/en/US/docs/solutions/Verticals/Education/srajrcollegesoverview.html</a>                                                                                 |
| <i>Cisco SAFE Reference Guide</i>                                                                                  | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html</a>                                                                                             |
| <i>Enterprise Mobility Design Guide 4.1</i>                                                                        | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html</a>                                                                         |
| <i>Cisco Wireless LAN Controller Configuration Guide, Release 6.0</i>                                              | <a href="http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.html">http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.html</a>                                                                   |
| <i>Cisco 802.11n Design and Deployment Guidelines</i>                                                              | <a href="http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/white_paper_80211n_design_and_deployment_guidelines.html">http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/white_paper_80211n_design_and_deployment_guidelines.html</a> |
| <i>Cisco 5500 Series Wireless Controllers Data Sheet</i>                                                           | <a href="http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/data_sheet_c78-521631.html">http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/data_sheet_c78-521631.html</a>                                                         |
| <i>RFC 5415 Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification</i>                 | <a href="http://www.ietf.org/rfc/rfc5415.txt">http://www.ietf.org/rfc/rfc5415.txt</a>                                                                                                                                                                                           |
| <i>Voice over Wireless LAN 4.1 Design Guide</i>                                                                    | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html</a>                                                                     |
| <i>Cisco Aironet 1520, 1130, 1240 Series Wireless Mesh Access Points, Design and Deployment Guide, Release 6.0</i> | <a href="http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP_60.html">http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP_60.html</a>                                                                                         |
| <i>Cisco Aironet 1520 Series Lightweight Outdoor Access Point Ordering Guide</i>                                   | <a href="http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps8368/product_data_sheet0900aecd8066a157.html">http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps8368/product_data_sheet0900aecd8066a157.html</a>                                               |
| <i>Cisco NAC Guest Server Overview</i>                                                                             | <a href="http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5707/ps8418/ps6128/product_data_sheet0900aecd806e98c9.html">http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5707/ps8418/ps6128/product_data_sheet0900aecd806e98c9.html</a>                                   |
| <i>Cisco Wireless Control System (WCS) Overview</i>                                                                | <a href="http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html">http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html</a>                                 |
| <i>Cisco Wireless Control System Configuration Guide, Release 6.0</i>                                              | <a href="http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/WCS60cg.html">http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/WCS60cg.html</a>                                                                                               |

| Document Title                                                                         | URL                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Deploying Cisco 440X Series Wireless LAN Controllers</i>                            | <a href="http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html">http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html</a>                                                                         |
| <i>Cisco ASA Botnet Traffic Filter</i>                                                 | <a href="http://www.cisco.com/en/US/prod/vpndevc/ps6032/ps6094/ps6120/botnet_index.html">http://www.cisco.com/en/US/prod/vpndevc/ps6032/ps6094/ps6120/botnet_index.html</a>                                                                                             |
| <i>Configuring Global Correlation</i>                                                  | <a href="http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/cli/cli_collaboration.html">http://www.cisco.com/en/US/docs/security/ips/7.0/configuration/guide/cli/cli_collaboration.html</a>                                                           |
| <i>Cisco IronPort Support</i>                                                          | <a href="http://www.ironport.com/support/">http://www.ironport.com/support/</a>                                                                                                                                                                                         |
| <i>WCCP Configuration Guide</i>                                                        | <a href="http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_46_se/configuration/guide/swwccp.html">http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_46_se/configuration/guide/swwccp.html</a> |
| <i>Identity Based Networking Services</i>                                              | <a href="http://www.cisco.com/en/US/products/ps6638/products_ios_protocol_group_home.html">http://www.cisco.com/en/US/products/ps6638/products_ios_protocol_group_home.html</a>                                                                                         |
| <i>NAC Appliance Support</i>                                                           | <a href="http://www.cisco.com/go/nacappliance">http://www.cisco.com/go/nacappliance</a>                                                                                                                                                                                 |
| <i>Clean Access Manager Configuration Guide</i>                                        | <a href="http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/47/cam/47cam-book.html">http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/47/cam/47cam-book.html</a>                                                       |
| <i>Clean Access Server Configuration Guide</i>                                         | <a href="http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/461/cas/461cas-book.html">http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/461/cas/461cas-book.html</a>                                                   |
| <i>NAC Out-Of-Band Wireless Configuration Example</i>                                  | <a href="http://www.cisco.com/en/US/products/ps6128/products_configuration_example09186a0080a138cc.shtml">http://www.cisco.com/en/US/products/ps6128/products_configuration_example09186a0080a138cc.shtml</a>                                                           |
| <i>Overall Campus Design</i>                                                           | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html</a>                                                                                                       |
| <i>Campus Network for High Availability Design Guide</i>                               | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html</a>                                                                         |
| <i>High Availability Campus Network Design-Routed Access Layer using EIGRP or OSPF</i> | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/routed-ex.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/routed-ex.html</a>                                                                                                     |
| <i>High Availability Campus Recovery Analysis Design Guide</i>                         | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html</a>                                                           |
| <i>Campus Virtual Switching System Design Guide</i>                                    | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVS_S_DG.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVS_S_DG.html</a>                                                                             |
| <i>Nonstop Forwarding with Stateful Switchover on the Cisco Catalyst 6500</i>          | <a href="http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps708/prod_white_paper0900aecd801c5cd7.html">http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps708/prod_white_paper0900aecd801c5cd7.html</a>                             |
| <i>Cisco Catalyst 4500 E-Series High Availability</i>                                  | <a href="http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps4324/prod_white_paper0900aecd806f0663.html">http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps4324/prod_white_paper0900aecd806f0663.html</a>                           |
| <i>Cisco StackWise Technology White Paper</i>                                          | <a href="http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps5023/prod_white_paper09186a00801b096a.html">http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps5023/prod_white_paper09186a00801b096a.html</a>                           |
| <i>Campus QoS Design 4.0</i>                                                           | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html</a>                                                           |
| <i>Service Ready Architecture for Schools Design Guide</i>                             | <a href="http://www.cisco.com/en/US/docs/solutions/Enterprise/Education/SchoolsSRA_DG/SchoolsSRA-DG.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Education/SchoolsSRA_DG/SchoolsSRA-DG.html</a>                                                           |