




# Community College and Vocational Education (CCVE) Deployment Guide

Last Updated: July 20, 2010



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



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.



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 Busi-

## Solution Authors




Rahul Kachalia

ness 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 opportunities 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, 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.





## CONTENTS

---

### CHAPTER 1

<b>Community College Reference Design Solution Overview</b>	<b>1-1</b>
Executive Summary	1-1
The Community College Environment	1-1
External Influences that Impact Community College Education	1-2
Vision for 21st Century Learning in Community College Education	1-2
Community College Challenges	1-3
Cisco Community College Reference Design	1-4
Community College Reference Design Service Fabric	1-4
High Availability	1-4
Differentiated Services	1-5
Access Layer Flexibility	1-7
Security	1-9
Network Security	1-9
Security Management	1-10
Secure Access Control	1-10
Mobility	1-11
Unified Communication	1-12
Network Management	1-13
Unified Communications Management	1-13
TelePresence Network Management	1-14
Performance Assurance	1-14
Routing and Switching Management	1-14
Identity Management	1-14
Video, Cable, and Content Delivery Management	1-14
Virtual Learning Environment	1-14
Secure Remote Access for Faculty and Students	1-15
Virtual Classroom	1-16
Online Collaborative Classroom Using WebEx Training Center	1-17
Review Streaming and Stored Video Using Video Portal	1-18
Operational Efficiencies	1-18
Network as a Platform	1-18
Data Center Optimization and Design	1-19
Facilities Management	1-20

- Secure Connected Classroom 1-23
  - Classroom Connectivity to the Network 1-23
  - Network Admission Control for Guests and Students 1-23
  - Application and Network Control 1-24
- Campus Safety and Security 1-25
  - IP-Based Video Surveillance 1-26
  - Communicate Campus Events and Emergencies with Digital Signage 1-26
  - IPICS for Emergency Collaboration 1-27
- Conclusion 1-29

**CHAPTER 2**

**Community College Reference Design—Service Fabric Design Considerations 2-1**

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



**CHAPTER 3****Community College LAN Design 3-1**

LAN Design	3-1
LAN Design Principles	3-4
Community College LAN Design Models	3-7
Main College Campus Network Design	3-9
Remote Large College Campus Site Design	3-10
Remote Medium College Campus Site Design	3-11
Remote Small College Campus Network Design	3-12
Multi-Tier LAN Design Models for Community College	3-13
Campus Core Layer Network Design	3-13
Core Layer Design Option 1—Cisco Catalyst 6500-Based Core Network	3-14
Core Layer Design Option 2—Cisco Catalyst 4500-Based Campus Core Network	3-15
Core Layer Design Option 3—Cisco Catalyst 4500-Based Collapsed Core Campus Network	3-17
Campus Distribution Layer Network Design	3-18
Distribution Layer Design Option 1—Cisco Catalyst 6500-E Based Distribution Network	3-19
Distribution Layer Design Option 2—Cisco Catalyst 4500-E-Based Distribution Network	3-21
Distribution Layer Design Option 3—Cisco Catalyst 3750-E StackWise-Based Distribution Network	3-22
Campus Access Layer Network Design	3-23
Access Layer Design Option 1—Modular/StackWise Plus Access Layer Network	3-24
Access Layer Design Option 2—Fixed Configuration Access Layer Network	3-24
Deploying Community College Network Foundation Services	3-25
Implementing LAN Network Infrastructure	3-25
Deploying Cisco Catalyst 6500-E in VSS Mode	3-26
Deploying Cisco Catalyst 4500-E	3-34
Deploying Cisco Catalyst 3750-E StackWise Plus	3-38
Deploying Cisco Catalyst 3560-E and 2960	3-41
Designing EtherChannel Network	3-41
Network Addressing Hierarchy	3-48
Network Foundational Technologies for LAN Design	3-49
Designing the Core Layer Network	3-49
Designing the Campus Distribution Layer Network	3-55
Designing the Multilayer Network	3-55
Spanning-Tree in Multilayer Network	3-58
Designing the Routed Access Network	3-59
Multicast for Application Delivery	3-63
Multicast Addressing Design	3-63
Multicast Routing Design	3-64

- Designing PIM Rendezvous Point 3-65
- Dynamic Group Membership 3-72
- Designing Multicast Security 3-73
- QoS for Application Performance Optimization 3-74
  - Community College LAN QoS Framework 3-75
  - Designing Community College LAN QoS Trust Boundary and Policies 3-78
  - Community College LAN QoS Overview 3-79
- Deploying QoS in College Campus LAN Network 3-83
  - QoS in Catalyst Fixed Configuration Switches 3-83
  - QoS in Cisco Modular Switches 3-84
  - Deploying Access-Layer QoS 3-86
  - Deploying Network-Layer QoS 3-104
- High-Availability in LAN Network Design 3-118
  - Community College LAN Design High-Availability Framework 3-119
  - Baselining Campus High Availability 3-119
  - Network Resiliency Overview 3-120
  - Device Resiliency Overview 3-121
  - Operational Resiliency Overview 3-123
  - Design Strategies for Network Survivability 3-125
  - Implementing Network Resiliency 3-126
  - Implementing Device Resiliency 3-129
  - Implementing Operational Resiliency 3-140
- Summary 3-150

**CHAPTER 4**

**Community College WAN Design 4-1**

- WAN Design 4-1
  - WAN Transport 4-3
    - Private WAN Service 4-3
    - Internet Service 4-4
    - Metro Service 4-5
    - Leased-Line Service 4-7
  - WAN Aggregation Platform Selection in the Community College Reference Design 4-7
    - Main Campus WAN Aggregation Platform Selection 4-8
    - Remote Large Campus WAN Aggregation Platform Selection 4-10
    - Remote Medium Campus WAN Aggregation Platform Selection 4-10
    - Remote Small Campus WAN Aggregation Platform Selection 4-10
  - Implementation of Community College WAN Reference Design 4-11
    - WAN Infrastructure Design 4-11
    - Leased-Line Service 4-12

Routing Design	4-13
QoS	4-19
QoS Implementation	4-22
QoS Implementation at WAN Aggregation Router 1	4-22
Implementation Steps for QoS Policy at WAN Aggregation Router 1	4-24
QoS Policy Implementation for WAN Aggregation Router 2	4-26
QoS Policy Between the Main Campus and Large Campus	4-29
QoS Policy Between the Main Campus and Medium Campus Location	4-30
QoS Policy Between Main Campus and Remote Small Campus Location	4-32
QoS Policy Implementation Between the Main Campus and Core	4-33
QoS Policy Between Large Campus and Main Campus Location	4-34
QoS Policy Between Medium Campus and Main Campus Location	4-35
QoS Policy Implementation Between Remote Small Campus and Main Campus Location	4-36
Redundancy	4-37
Multicast	4-42
Summary	4-45

**CHAPTER 5****Community College Mobility Design 5-1**

Mobility Design	5-1
Accessibility	5-5
WLAN Controller Location	5-7
WLAN Controller Connectivity	5-8
Controller Connectivity to the Wired Network	5-9
Controller Connectivity to Wireless Devices	5-10
Access Points	5-20
Usability	5-27
Quality-of-Service	5-27
Guest Access	5-28
Traffic and Performance	5-33
Manageability	5-33
Reliability	5-37
Controller Link Aggregation	5-37
Controller Redundancy	5-40
AP Controller Failover	5-42
Community College Mission Relevancy	5-43
Safety and Security	5-43
Virtual Learning	5-44
Secure Connected Classrooms	5-45
Operational Efficiencies	5-45

- Wireless LAN Controller Configuration 5-47
  - WLC and Wired Network Connections 5-47
    - Remote Campus 5-50
    - Mobility Groups 5-50
  - WLAN Configuration 5-52
    - Faculty and Staff Data WLAN 5-52
    - Faculty and Staff Voice WLAN 5-52
    - Student WLAN 5-53
    - Guest Access WLAN 5-55
  - WLAN QoS 5-58
- Access Point Configuration 5-59
  - AP 1520 Configuration 5-60
    - Adding the AP1520 MAC Address to the WLC 5-60
    - Configuring the AP1520 as a Root Access Point (RAP) 5-61
- WCS Configuration 5-63
  - WCS Users and User Groups 5-63
  - WCS Virtual Domains 5-63
  - Reference Documents 5-66

**CHAPTER 6**

**Community College Security Design 6-1**


- Security Design 6-1
- Network Foundation Protection 6-5
- Internet Perimeter Protection 6-7
  - Internet Border Router Security 6-10
  - Internet Firewall 6-10
  - Intrusion Prevention 6-13
  - E-Mail Security 6-14
  - Web Security 6-19
- Data Center Protection 6-24
- Network Access Security and Control 6-25
  - Cisco Catalyst Integrated Security Features 6-25
  - Cisco Identity-Based Network Services 6-26
    - IEEE 802.1X Protocol 6-26
    - 802.1X and EAP 6-27
    - Impacts of 802.1X on the Network 6-27
    - 802.1X in Community Colleges 6-28
  - Cisco NAC Appliance 6-28
    - NAC Appliance Components 6-29
    - NAC Appliance Modes and Positioning 6-31

- NAC Deployment in the Community College Reference Design 6-35
- Endpoint Protection 6-41
- Community College Mission Relevancy 6-41
  - Virtual Learning Environments 6-41
  - Secure Connected Classrooms 6-42
  - Safety and Security 6-42
  - Operational Efficiencies 6-43
- Community College Security Configuration Guidelines 6-44
  - Internet Border Router Edge ACL Deployment 6-44
    - Module 1: Implement Anti-spoofing Denies 6-44
    - Module 2: Implement Explicit Permits 6-45
    - Module 3: Implement Explicit Deny to Protect Infrastructure 6-45
    - Module 4: Explicit Permit for Traffic to Community College's Public Subnet 6-45
  - Internet Firewall Deployment 6-45
    - Firewall Hardening and Monitoring 6-47
    - Network Address Translation (NAT) 6-49
    - Firewall Access Policies 6-49
    - Firewall Redundancy 6-52
    - Routing 6-53
    - Botnet Traffic Filter 6-54
  - IPS Global Correlation Deployment 6-58
    - How IPS with Global Correlation Works 6-59
  - Web Security Deployment 6-66
    - Initial System Setup Wizard 6-67
    - Interface and Network Configuration 6-68
    - WCCP Transparent Web Proxy 6-71
    - Web Access Policies 6-74
  - Catalyst Integrated Security Features Deployment 6-75
  - NAC Deployment 6-76
    - NAC Deployment for Wired Clients 6-76
    - NAC Deployment for Wireless Clients 6-89
- Further Information 6-98

**APPENDIX A**

**Reference Documents A-1**





# CHAPTER 1

## Community College Reference Design Solution Overview

---

### Executive Summary

The Cisco Community College reference design is a framework designed to assist Community Colleges in designing and implementing a network for the 21st century learning environment. The design is created around solving complex business challenges that these institutions face. At its foundation is the network service fabric, which is a collection of features and technologies that serve to provide a highly available network that understands and adapts to the different services that it facilitates. The Cisco Community College reference design supports business solutions that utilize the service fabric were created to help these institutions:

- Create a 21st century virtual learning environment to enable highly interactive and collaborative learning and teaching experiences while delivering any content, anytime, anywhere, to any device.
- Increase operational efficiencies by using the network as a platform and optimizing data center design to extend cost reduction, improve utilization of under-used network capacity, and add flexibility to organizations through business process improvements.
- Design and implement secure connected classrooms that serve the educational needs of students and faculty by leveraging network and application control.
- Provide safety and security on campus by utilizing a platform architecture that proactively protects students, faculty, and staff.
- Allow for facilities management to interact with building controls, measure power consumption, and control energy output to reduce energy cost and carbon footprint, creating greener and more energy efficient campuses.

### The Community College Environment

In the United States, community colleges, sometimes called junior colleges, technical colleges, or city colleges, are primarily two-year public institutions providing higher education and lower-level tertiary education, granting certificates, diplomas, and associate degrees. Traditionally, after graduating from a community college, some students transfer to a four-year liberal arts college or university for two to three years to complete a bachelor's degree.

## External Influences that Impact Community College Education

Current economic conditions, a rise in continuing education enrollment, and the addition of courses for traditional secondary schools have all led to a substantial increase in enrollment at community colleges.

The worldwide recession has led to a significant reduction in funding for institutions of higher learning, which in turn have tightened budgets and increased tuition rates. Increased tuition costs, as well as the reduction of available income due to the high unemployment rate worldwide, have led students that may have attended a traditional institution of higher learning to turn to community colleges as a lower-cost educational option. This allows the student to begin a post-secondary education at a community college to attain a two-year Associate's degree, while having the option to continue on to an institution of higher learning to earn a traditional undergraduate degree.

Continuing education for adults has also been on the rise. Adults who have lost their jobs have been attending community colleges to augment their existing skill set or to take workforce development programs to retrain for another profession. The addition of distance learning as an option for working adults has also contributed to the rise in enrollment for continuing education students.

Secondary school students have also been a factor in the increased enrollment of community colleges. As the children of the baby boomers enter their college years, the competition to get into top-rated institutions has increased. One tool that secondary school students use to stand out from the crowd of applicants is to demonstrate their academic prowess at a college level by attending and passing community college courses while in secondary school.

## Vision for 21st Century Learning in Community College Education


The 21st century learning environment will be an environment where anyone, from anywhere, at anytime can access community college resources. The traditional classroom will be extended by the use of online communities of learning. Students will be able to access their course work online, receive instruction by attending class either in person or remotely, and be able to retrieve the instruction at a later time through audio and video recordings augmented by online instructor and class notes.

This style of learning requires a collaborative environment in which instructors and students are not bound by geographic distances; students will be able to work together remotely to seamlessly complete projects and course work.

21st century learning will also be ecologically friendly by reducing the need to expand brick and mortar schools and commuting to campus to attend class. In addition, buildings and infrastructure optimized to reduce energy usage will all help reduce green house gasses and lead to greener and more energy efficient campuses.

Security, whether physical or logical, will become increasingly important in the 21st century learning environment. Security elements will be ubiquitous across traditional and virtual campuses and will work in concert to provide a safer environment for all students, faculty, and staff. See [Figure 1-1](#).



**Figure 1-1** *Characteristics of a 21st Century School*

227410

## Community College Challenges

In the United States from 2000 to 2006, there was a 10 percent growth in overall enrollment at two-year institutions, according to the most recent figures from the Department of Education. During the 2006-2007 academic years, 6.2 million students were enrolled in the country's 1,045 community colleges, 35 percent of all postsecondary pupils that year, according to a new National Center for Education Statistics study. Though full national figures for the 2007-2008 academic year are not yet available and most colleges only have estimates for their enrollments this fall, many colleges are projecting increases of around 10 percent over last fall.

This increased enrollment presents new challenges for delivering educational course work. The demand for instruction is increasing at a pace that does not allow the brick and mortar campus to expand quickly enough to handle growth. Community colleges have turned to online learning as the predominant method of handling this growth. While online learning has helped to handle the increase demand, it has been criticized for lacking the face-to-face experience that traditional learning provides, as well as lower than traditional passing rates for students. Community colleges are faced with the task of delivering a true virtual learning environment that delivers experiences that are comparable to the traditional environment. They must also allow for secure remote working environment for faculty and staff.

While community colleges are growing, their funding are flat or decreasing similar to institutions of higher learning, so they have to do more with less. Operational efficiencies are being streamlined to allow community colleges to produce the same quality of education with fewer resources.

The rapid and expansive adoption of technology by students has led community colleges to offer connected classrooms and laboratories. Allowing the student to be connected to the network from the classroom or lab while receiving lecture has the benefits of mutual use of online resources, but it also requires community colleges to ensure that their networks are protected and that only authorized users are allowed access. Additionally, they must be able to control the use of applications and resources that reside on the network.

Since the incidents at Columbine and Virginia Tech, campus safety and security have become paramount to all educational institutions. Creating a safe campus is a major challenge for all community. They must employ the right tools to ensure the safety of the students, faculty, and staff. The safety and security systems in place must allow campus safety personnel to respond immediately and effectively in the case of an incident. A safe campus environment is a key differentiator for student and faculty recruitment and is an integral part of the community that welcomes local citizens and contributes positively to the area in which it resides.

As the world changes and becomes "more green", educational institutions are put in the position of leading that cause. Students are overwhelmingly concerned about greenhouse gasses as well as energy usage. The facility managers of community colleges must be able to strike the right balance between conducting business and optimizing energy use.

Budget reductions, increased enrollment, and limited staff are business constraints that impact the ability of community colleges to effectively address these challenges. A well thought out plan that optimizes resources, minimizes costs, and allows for flexibility in implementation is needed to allow community colleges to achieve the vision of 21st century learning.

## Cisco Community College Reference Design

The Cisco Community College reference design is a framework designed around the vision, challenges, and constraints that community colleges face. Cisco has employed a business down approach in this design. The first step in creating the Community College reference design was to understand the vision that these institutions have for 21st century learning. Next, we identified the challenges that these institutions are facing. After understanding the vision and challenges, we recognized the business constraints that shape these institutions' ability to adopt solutions to address the challenges. Finally, we selected the best technologies, features, and equipment that allow these institutions to solve these challenges within the business constraints.

The Community College reference design is composed of the following solutions:

- Community College Reference Design Service Fabric
- Virtual Learning Environment
- Operational Efficiencies
  - Data Center Design
  - Facilities Management
- Secure Connected Classroom
- Campus Safety and Security

## Community College Reference Design Service Fabric

The service fabric is the foundational network on which all solutions and services build. It comprises local and wide area networking equipment, security appliances, unified communications hardware as well as network, security, and mobility services that all work in concert to provide the fundamental network building block that all solutions and services use.

### High Availability


The high availability technologies used in the Cisco Community College reference design allow network equipment to eliminate the effects of any unplanned link or network failures by understanding the typology of the infrastructure and using that information to immediately re-route network traffic without the need to relearn (reconverge) the network. The use of these technologies allows critical service communications to remain unaffected by network outages.

The service fabric is designed to provide nonstop communications with resiliency throughout all the layers of the network. Many elements of the network must be correctly designed and implemented to ensure a highly available network.

Network resiliency is achieved by the careful design and implementation of network paths, devices, and power:

- *Path resiliency*—End-to-end resilient paths are required (Figure 1-2).
- *Device resiliency*—Resilient devices are usually preferred over resilient components within a single device. While resilient components within a single device are valuable, the best availability is usually achieved with completely separate devices (and paths).
- *Power resiliency*—Power diversity is another area that must be addressed because resilient devices attached to a single power source are vulnerable to simultaneous failure. For example, resilient core switches should have at least two unique power sources. Otherwise, a single power failure will bring down both core switches. Alternatively, backup power could be implemented. These types of mundane issues are very important when creating a highly available service fabric.

**Figure 1-2 End-to-End Resilient Paths**



## Differentiated Services

Certain network services demand more from the network than others. For example, voice communications do not work if parts of the conversation drop out. Video conferencing is not useful if the picture keeps freezing. Additionally, a professor's use of the network to conduct class should take precedence over a student surfing the Web. Finally, if there are more traffic demands than the network can handle, the network should be able to make decisions as to which traffic is most important. The ability to understand, mark, shape, and limit traffic is embedded into the Cisco Community College reference design using Cisco's extensive array of quality-of-service (QoS) technologies.

There is some debate in the networking industry about the need to deploy QoS in campus architectures due to ample amounts of bandwidth and the rarity of congestion. However, during network attacks or a partial outage, this situation can change dramatically. It has been shown that QoS can serve as a vital tool to maintain the performance of priority applications and traffic during a degraded network condition.

The following are some reasons why QoS is important in the campus portion of the network:


- The introduction of 10Gbps (and higher) link speeds is creating greater mismatches between high-speed and low-speed links in the campus. This increases the need to buffer and prioritize traffic.
- Well-known applications ports, like HTTP, are being used by a large number of applications. There is a need to distinguish between high-priority and low-priority traffic using the same port numbers to make sure priority traffic is transmitted.
- Prioritized traffic, like voice and video, must continue to flow even during a network attack or during a partial failure in the network. Attack traffic often masquerades as legitimate traffic using well-known port numbers. There is a need to distinguish between legitimate and bogus traffic by inspecting data packets more deeply.

The following principles should guide QoS deployments:

- Classify and mark traffic as close to the network edge as possible. This is called creating a *trust boundary*. Traffic crossing the trust boundary is considered “trusted” and the QoS markings are adhered to in the rest of the network.
- Police/rate-limit traffic as close to the source as possible. It is most efficient to drop unwanted traffic as close to the source as possible, rather than transmitting it further into the network before dropping it.
- Perform QoS functions in hardware rather than software. Software-based QoS functions can easily overwhelm the CPU of networking devices. High-speed networks require hardware-based QoS functions.

Figure 1-3 summarizes key QoS functions and where they should be performed.

Figure 1-3 QoS Functions



## Access Layer Flexibility


In a hierarchical network design, the core and distribution layers can reconverge in less than 1 second after most types of failures. The access layer typically has longer convergence times due to the inherent deficiencies of a flat Layer-2 architecture. Bridging loops, broadcast storms, and slow reconvergence are examples of access layer problems that reduce end-to-end availability. Spanning tree typically takes up to 1 minute to recover from a link or system outage, which is far too long to support real-time mission critical applications or provide 99.999 percent availability. There are several design changes and software features that can be implemented to improve availability in the access layer.

Currently, there are three different ways to design the access-layer control plane. Although all three of them use the same physical layout, however they differ in performance and availability.

The traditional multi-tier network is designed where all access switches run in Layer 2 mode between the access and the distribution, while they run in Layer 3 mode between distribution and the core. Cross-connects between distribution switches are usually Layer 2 links. When not optimized, this model is dependent on spanning tree, with all its inherent limitation, to detect and recover from network failures. As mentioned, load balancing of resilient uplinks is not possible because spanning tree usually blocks one uplink. HSRP, VRRP, or GLBP must be used to provide First Hop Routing Protocol (FHRP) resiliency. While deficiencies are evident in the traditional multi-tier approach, design changes and feature enhancements are available to greatly enhance availability and performance.

The current multi-tier best practice is to create unique VLANs on each access switch as shown in [Figure 1-4](#). The best practice design offers several benefits. First, a loop-free topology is created. This means spanning tree does not impact reconvergence times. Traffic is load balanced across two active uplinks, achieving maximum throughput and minimum failover times. This loop-free topology also reduces the risk of broadcast storms and unicast flooding.

**Figure 1-4 Best Practice Multi-Tier Has Unique VLANs on each Access Switch**



One disadvantage of the best-practice multi-tier design is the requirement to redesign the VLAN and IP addressing scheme: unique IP subnet(s)/VLAN(s) per switch. This can be a significant challenge in large mature networks. The routed access model discussed below has this same drawback.

The routed access layer design is an improvement over the traditional multi-tier, as the name implies this design pushes routing into the access layer switches and creates an end-to-end routed infrastructure. Several important benefits are gained:

- Spanning tree issues are virtually eliminated.
- Reconvergence times for the end-to-end network can be reduced to 1 second or less.
- Reconvergence times become more predictable with the elimination of spanning tree.
- Resilient uplinks can be fully utilized.
- HSRP/VRRP is no longer needed to provide host resiliency. This simplifies configuration, management, and troubleshooting.
- Troubleshooting is accomplished using well-known Layer 3 tools, such as traceroute, ping, etc.
- Network layout, naming, and VLAN numbering can become standardized across buildings and campuses.

A drawback to the routed access model is the requirement to have separate IP subnets and VLANs on every access switch. This is in contrast to the traditional multi-tier model where a user VLAN can span several switches. However, the convergence times of the routed access layer are much less than that of flat Layer 2 networks.

Employing a hybrid access-layer design allows the network administrator to leverage their existing Layer 2 network while giving them the flexibility to implement and slowly migrate their existing network to a routed access layer design model. Advantages of a routed access design include the following:

- Prevention of loops without the need of multiple complex Layer 2 technologies such as spanning tree protocol.
- High availability and ease of network troubleshooting and management by leveraging well-known Layer 3 troubleshooting tools and technologies.

# Security

Building a secure Community College reference design is paramount to the community college environment. Community Colleges have to balance network access given to students, guests, faculty and staff with protecting critical data and personal information of students and staff. The Community College reference design approaches security as described in the following subsections.

## Network Security

Build a network security infrastructure that inherently detects and blocks invasive software attacks and intruder access.

### Firewalls

- Combines firewall, VPN, and optional content security and intrusion prevention to distribute network security across your operations
- Provides threat defense and highly secure communications services to stop attacks before they affect business continuity
- Reduces deployment and operational costs while delivering comprehensive network security for networks of all sizes

### Intrusion Prevention

- Identifies, classifies, and stops malicious traffic, including worms, spyware, adware, viruses, and application abuse
- Delivers high-performance, intelligent threat detection and protection over a range of deployment options
- Uses reputation filtering and global inspection to give businesses actionable intelligence and prevent threats with confidence
- Promotes business continuity and helps businesses meet compliance needs

### E-mail and Web Security

Reduce costly downtime associated with E-mail-based spam, viruses, and web threats.

### E-mail Security Appliances

- Fights spam, viruses, and blended threats to protect organizations of all sizes with industry-leading security capabilities
- Prevents data leaks, enforces compliance, and protects reputation and brand assets
- Reduces downtime, simplifies administration of community college mail systems, and eases the technical support burden

### Web Security Appliances

- Integrates industry-leading web-usage controls, reputation filtering, malware filtering, and data security
- Takes advantage of Cisco Security Intelligence Operations (SIO) and global threat correlation technology to help optimize threat detection and mitigation

- Combines multiple layers of web security technology to combat complex and sophisticated web-based threats
- Supports built-in management capabilities to simplify administration and provide visibility into threat-related activity

## Security Management

Simplify the configuration, monitoring, and management of your Cisco security capabilities.

### Cisco IronPort Security Management Appliances

- Simplifies security management across Cisco IronPort E-mail and web security products
- Delivers centralized reporting, message tracking, and spam quarantine for the E-mail security appliances
- Provides centralized web policy management for web security appliances
- Allows for delegated administration of web access policies and custom URL categories

### Cisco Security Manager

- Facilitates the configuration and management of Cisco firewalls, VPNs, IPS sensors, and integrated security services
- Ideal for controlling large or complex deployments of Cisco network and security devices
- Supports role-based access control and an approval framework for proposing and integrating changes
- Delivers flexible device management options, including policy-based management and methods for deploying configuration changes

### Cisco Security Monitoring, Analysis and Response System

- Identifies threats by learning the topology, configuration, and behavior of the network environment
- Facilitates troubleshooting and identifying attacks or vulnerabilities for a wide range of enterprise networks
- Visually characterizes an attack path, identifies the threat source, and makes precise recommendations for threat mitigation
- Simplifies incident management and response through integration with Cisco Security Management software

## Secure Access Control

Enforce network security policies; help secure user and host access control, and control network access based on dynamic conditions and attributes.

### Network Admission Control Appliance

- Enforces network security policies on all devices by allowing access only to compliant and trusted devices



- Blocks access by noncompliant devices and limits the potential damage from emerging security threats and risks
- Reduces virus, worm, and unwanted access threats by promoting efficiency and integrating with other Cisco products

### Cisco Secure Access Control System

- Controls network access based on dynamic conditions and attributes through an easy-to-use management interface
- Meets evolving access requirements with rule-based policies for flexibility and manageability
- Simplifies management and increases compliance with integrated monitoring, reporting, and troubleshooting capabilities
- Adopts an access policy that takes advantage of built-in integration capabilities and distributed deployment

## Mobility

Cisco Mobility and Wireless Solutions for Community Colleges give students and staff the freedom to be anywhere on campus and still perform all the tasks they would normally do in a classroom's, or an office's wired network. The solutions enable new network connections to PCs, laptops, PDAs, printers, video cameras, videoconferencing units, IP phones, and other devices, making school resources more widely available and improving collaboration among students, Faculty and Staff

Mobility products include the following:

- Cisco Aironet Access Points connect Wi-Fi devices to networks in a variety of wireless environments. Cisco Next-Generation Wireless solutions use 802.11n technology to deliver unprecedented reliability and up to nine times the throughput of 802.11a/b/g networks. Wi-Fi certified for interoperability with a variety of client devices, these access points support robust connectivity for both indoor and outdoor environments.
- Wireless LAN controllers simplify the deployment and operation of wireless networks, helping to ensure smooth performance, enhanced security, and maximum network availability. Cisco wireless LAN controllers communicate with Cisco Aironet access points over any Layer 2 or Layer 3 infrastructure to support systemwide wireless LAN (WLAN) functions such as the following:
  - Enhanced security with WLAN policy monitoring and intrusion detection
  - Intelligent radio frequency (RF) management
  - Centralized management
  - Quality of service (QoS)
  - Mobility services such as guest access, voice over Wi-Fi and location services

Cisco wireless LAN controllers support 802.11a/b/g and the IEEE 802.11n draft 2.0 standard, so you can deploy the solution that meets your individual school requirements. From voice and data services to location tracking, Cisco wireless LAN controller products provide the control, scalability, security, and reliability you need to build highly secure, district-wide wireless networks.

- Cisco Wireless Location Appliance allows school districts to simultaneously track thousands of devices from within the WLAN infrastructure, bringing the power of a cost-effective, high-resolution location solution to critical applications such as the following:
  - High-value asset tracking

- IT management
- Location-based security

This easy-to-deploy solution smoothly integrates with Cisco WLAN Controllers and Cisco lightweight access points to track the physical location of wireless devices to within a few meters. This appliance also records historical location information that can be used for location trending, rapid problem resolution, and RF capacity management.

## Unified Communication

Cisco Unified Communications solutions provide many solutions for community colleges that wish to take advantage of media-rich unified communications functionality. Each aspect of the total unified communications architecture provides opportunities for enhancing links within the higher education community. Functionality includes IP telephony, unified client software, presence, instant messaging, unified messaging, rich-media conferencing, mobility solutions, and application development.

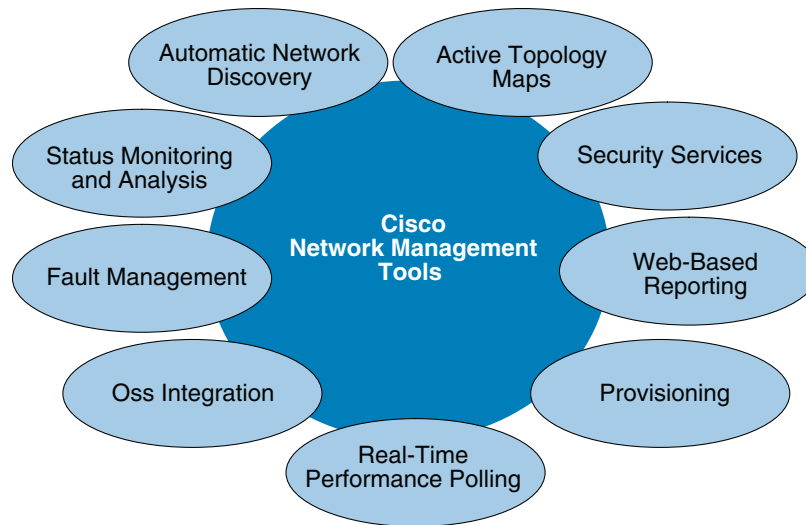
- *IP telephony*—At the foundation of the Cisco Unified Communications solution is its proven, industry-leading call processing system, Cisco Unified Communications Manager. This highly available, enterprise-class system delivers call processing, video, mobility, and presence services to IP phones, media processing devices, VoIP gateways, mobile devices, and multimedia applications. The system can scale to one million users across 1000 sites or more or 60,000 users within a single clustered system. Built-in resiliency keeps service reliable. Cisco also offers several unified communications platforms for small community colleges. All of these standards-based systems work with an array of third-party phones and dual-mode devices. The systems also provide integration with existing desktop applications such as calendar solutions, E-mail, enterprise resource planning (ERP) systems, and customer relationship management (CRM) software. Cisco unified communications capabilities can also be extended to a variety of mobile phones, including those that run Symbian, Blackberry, and Windows Mobile operating systems.
- *Unified client software*—Cisco offers several rich-media client applications that improve productivity and simplify processes. Available on Microsoft Windows and Mac OS environments as well as mobile operating systems, these clients support a range of applications, including voice, presence/messaging, unified messaging, video, and conferencing. Communications functionality has also been unified with applications from industry partners. For example, call control and presence can be launched and managed from within Microsoft Outlook through a Cisco Unified Personal Communicator widget or toolbar.
- *Presence and instant messaging*—Cisco presence solutions based on Session Initiation Protocol (SIP) or (SIMPLE) provide SIP presence and proxy services to deliver IM and click-to-call features. Through the presentation of dynamic presence information, presence solutions allow users to check the availability of colleagues in real time, reducing “phone tag” and improving productivity. Cisco presence and instant messaging solutions work in conjunction with Cisco Unified Communications Manager and support Cisco Unified Personal Communicator, Cisco IP phones, Cisco IP Phone Messenger, WebEX Connect, IBM Sametime clients, and Microsoft clients.
- *Unified messaging*—Cisco unified messaging solutions easily integrate with existing environments and provide flexible deployment options to meet each organization’s individual needs. The broad range of easy-to-manage solutions includes products tailored for small, medium-sized, and very large organizations, with feature-rich functionality aligned intelligently with business requirements.
- *Rich-media conferencing*—Cisco conferencing solutions help remote workers and teams communicate more effectively to save time and reduce costs. Integrated voice, video, and Web conferences can be set up and attended in a single step from IP phones, instant messaging clients, Web browsers, and Microsoft Outlook and IBM Lotus Notes calendars.

- *Mobility solutions*—Cisco Unified Communications extends rich call control and collaboration services to facilitate easy collaboration among mobile workers on campus or on the move. By anchoring communications in the network, Cisco Mobile Unified Communications solutions connect different mobile worker types and workspaces, provide a consistent collaboration experience regardless of location, maintain business continuity and compliance, and take advantage of least-cost routing of mobile communications over the education network. Cisco Mobile Unified Communications solutions support a wide range of popular handheld platforms, enabling workers to communicate quickly and easily using their familiar mobile equipment.
- *Application development*—Community colleges may operate in unique educational environments that require specialized applications. To meet these needs, Cisco provides a versatile service creation platform, enabling institutions and partners to rapidly and easily develop and deliver innovative media-rich and Web-rich applications. The platform also allows organizations to easily blend unified communications capabilities with existing business process systems.

## Network Management

As community colleges implement more services and their networks become more instrumental as the platform for 21st century learning, the need to understand how the network is operating, what issues it is experiencing, and how those issues are impacting students, faculty, and staff become critical. Network management tools (see [Figure 1-5](#)) have been developed to help the IT staff understand the status and operation of service fabric and the services that are in operation in the network. This section discusses some of the specific network management options available to community colleges.

**Figure 1-5 Cisco Network Management Tools**



## Unified Communications Management

The broad range of products in the Cisco Unified Communications portfolio provide enormous flexibility for applications, rich media collaboration, call control and messaging, and IP communications. Networks that deliver data, voice, video and rich media applications require unified, system-level management.

## TelePresence Network Management

Cisco TelePresence integrates advanced audio, high-definition video and interactive elements with the power of the underlying network to deliver an immersive, face-to-face experience for collaboration. Cisco TelePresence network management is essential to the Cisco TelePresence experience.

## Performance Assurance

Cisco network management products can help network administrators effectively manage network resources, plan for changes in resource usage, and resolve problems before they affect users. Quick access to configuration menus and easy-to-read performance reports on data, voice, and video traffic helps network operators to monitor trends, plan capacity, and optimize performance.

## Routing and Switching Management

Cisco network management products support more than 400 types of Cisco devices with detailed reporting, monitoring, and configuration. They can save network administrators time and effort with improved inventory and configuration management, rapid software deployment, and simplified troubleshooting.

## Identity Management

The ever-increasing number of methods for accessing networks makes security breaches and uncontrolled user access a primary concern. Network operators can use Cisco network management products with identity management features to protect systems and information through internal trust and identity policies, access control, and compliance features. The result is security assurance and protection of company profits and assets.

## Video, Cable, and Content Delivery Management

Designed to be ready for advanced applications, Cisco network management products help ensure high performance and high availability, leading to higher subscriber satisfaction. With Cisco network management products, subscribers can access next-generation services such as IP telephony, video on demand, and interactive gaming.


## Virtual Learning Environment

One of the key challenges that face community colleges is extending their learning environments beyond the campus to allow for online/distance learning, professor collaboration, and anytime/anywhere access for students to obtain course and educational materials. This virtual learning environment is key in allowing community to continue to grow at current rates and enhance the learning experience.

Cisco has several offerings for the virtual learning environment:

- Secure remote access
- Virtual classroom
- WebEx training center
- Video portal

See [Figure 1-6](#).

**Figure 1-6 21st Century Learning Environment**

227412

## Secure Remote Access for Faculty and Students

Secure remote access is a way for community colleges to extend the network using secure remote access to anyone, anytime, anywhere, with virtually any device, in order to increase productivity and reduce costs. Secure remote access allows you to deliver network access safely and easily to a wide range of users and devices.

Cisco Secure Remote Access is a comprehensive and versatile remote access solution that supports the widest range of connectivity options, endpoints, and platforms to meet the changing and diverse remote access needs of community colleges.

The Secure Remote Access solution gives IT administrators a single point of control to assign granular access based on both user and device. It provides both full and controlled client-based network access to Web-based applications and network resources for a highly secure, flexible, remote access deployment.

Benefits include the following:

- Web-based access without preinstalled desktop software:
  - Facilitates customized remote access based on user and security requirements
  - Reduces desktop interaction and support costs
- Threat-protected Virtual Private Network (VPN) access:
  - Protects against viruses, worms, spyware, and hackers by integrating network and endpoint security in the Cisco Secure Sockets Layer (SSL) VPN platform
  - Eliminates the need for additional security equipment and management infrastructure
- Multiple VPN support from a single platform:
  - Supports both IP Security (IPSec) and SSL connectivity
  - Supports unified management of remote access and site-to-site VPN services to help reduce costs and management complexity

## Virtual Classroom

Communication, collaboration, and learning are the fundamental building blocks of higher education. Students expect to use the latest technologies and many prefer dynamic online content to static printed materials. Distance learning and e-learning enable community colleges to deliver more engaging content to both on-campus and remote students, creating a new and potentially significant revenue stream.

Enhancing education through video and rich media elements can:

- Provide anywhere, anytime learning experiences not traditionally available to all students
- Offer a better way to present abstract ideas, making them easier to understand
- Eliminate the barriers of time, distance, and resources
- Permit faculty, staff, and students worldwide to function as if they were in the same room

Cisco's Virtual Classroom solution is an integrated learning and administrative environment that enables academic excellence and administrative efficiencies. The virtual classroom strategy focuses on the implementation of a network platform that can enable highly interactive and collaborative learning and teaching learning experiences while delivering any content, anytime, anywhere, to any device. As a result, Cisco's goal is to provide a scalable a solution that provides educational institutions with the necessary technology to solve business problems and address important issues, such as increasing student participation and graduation rates, in a cost effective and successful manner.

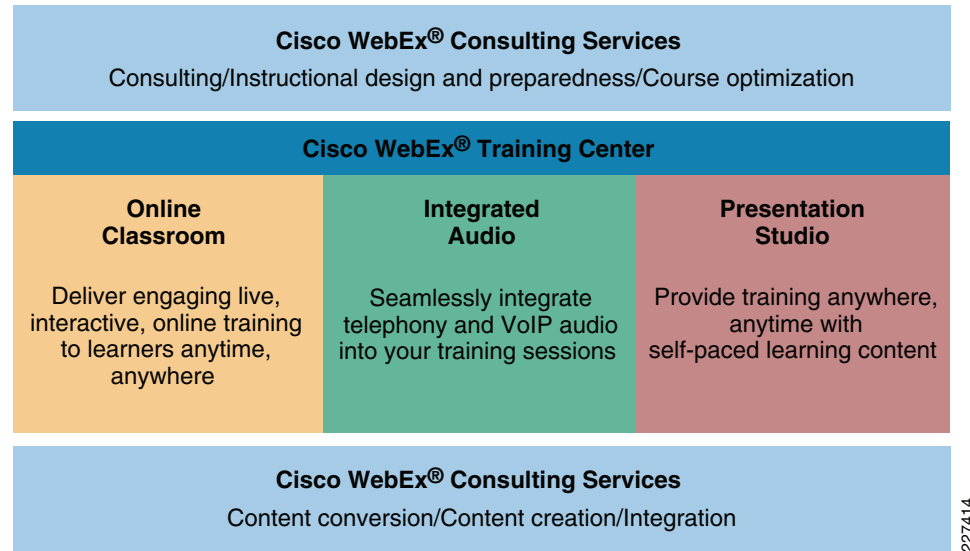
The Virtual Classroom solution is composed of campus-hosted technologies:

- Unified Communications
- TelePresence
- Video over IP technologies
- Wide area application services

The Virtual Classroom solution is designed to allow educational institutions to expand the reach of their offerings, both in a geographic and time-based manner. By using Web, video, and audio collaboration and scalable content delivery technologies, educational institutions can now reach students that are unable to physically attend class. Additionally, instructors and professors can record and edit pre-existing content into the recorded sessions and then post those content objects for the students to download to various devices, such as a mobile device. As a result, the schools can now scale their assets beyond their physical presence, such as subject matter expert or sign language teacher, to other classrooms and locations. Lastly, the ability to record the classroom events allows the student to also refer back to earlier classes for review or playback a class they missed.

## Online Collaborative Classroom Using WebEx Training Center

**Figure 1-7 Online Collaborative Classroom Using WebEx Training Center**



WebEx Training Center is a Cisco web-hosted solution designed to facilitate online instruction for anywhere, anytime learning experiences. Features include the following:

- The ability to capture each student’s attention with live, interactive instruction:
  - Share presentations, stream multimedia, and live video.
  - Connect online learners with remote computers, applications, and simulations before, during, or after live training sessions.
  - Pass control to attendees to demo applications.
- Encourage, improve, and track interaction:
  - Enhance and test retention with features like polling, testing, and breakout sessions.
  - Extend the reach of your educational facilities to students across the globe.
  - Simplify session registration and track attendance.
  - Record sessions and offer them on demand.
- Extend the reach of your institution while reducing costs:
  - Connect with more learners more often, while you eliminate travel and venue costs.
  - Charge for classes and online certification programs to turn your training center into a profit center.
  - Manage costs and pay as you go for an affordable, predictable monthly fee.

WebEx solutions are software delivered as a service (SaaS). Therefore community colleges do not need to worry about providing servers, maintenance, or support. Those items are handled as part of the subscription service.

Some advantages of SaaS:

- Performance and reliability for your critical communications.
- Keep sessions as private and safe as you need with exceptional security.

- No need to handle maintenance and upgrades.

## Review Streaming and Stored Video Using Video Portal

Students can browse, search, and view digital media content interactively at the desktop with the Cisco Video Portal. An integrated component of the Cisco Digital Media System for Cisco Desktop Video, the Cisco Video Portal is a sophisticated video playback portal that uses standard Web technologies to deliver compelling live Webcasts and on-demand video to your audiences. Platform independent, the Cisco Video Portal fits easily into the existing network and infrastructure of community colleges.

The Cisco Video Portal features include:

- Customizable interface, program guide, and keyword search
- Personalized and featured playlists
- Advanced player controls—Full-screen video playback, fast forward, rewind
- Slide synchronization with video
- Submission and management of questions during live Webcasts
- Video sharing
- Secure log in and access to user-specific content based on Active Directory/LDAP
- Support for major video formats—Windows Media, Flash, MPEG-4/H.264, QuickTime
- Detailed content and user access reporting—Who, what, when, and how often

With the Cisco Digital Media Manager, the look and feel of the Cisco Video Portal can be customized to reflect the image of the educational institution.

## Operational Efficiencies


Community colleges are faced with the daunting task of doing more with less, facing explosive growth as budgets are reduced due to funding cuts. The Cisco Community College reference design leverages the use of the network as a platform to deliver an expanded array of education services and data center optimization as a means for creating operational efficiencies to reduce costs and capitalize on under-utilized network capacity.

## Network as a Platform

The concept of using the network as a platform is the next phase in the evolution of network convergence. In the past, there was an effort to consolidate voice, video, and data networks onto a single IP network to allow organizations to reduce the cost of communication and take advantage of under-used network capacity. The network as a platform extends that concept beyond voice, video, and data services to allow for any IP-based service to use the network, wired or wireless, to extend cost reduction, improve utilization of under-used network capacity, and add flexibility to organizations through business process improvements. See [Figure 1-8](#).



Figure 1-8 Network as a Platform



The network infrastructure or fabric must be able to understand the requirements of these non-traditional services and remain flexible and adaptable to their needs (discussed in detail later in this document). While the concept of adding non-traditional services like building controls and contextual awareness on top of an existing network seems like an easy task, the reality is that the underlying service fabric must be designed to accommodate and differentiate those services, especially as those services travel alongside others.

## Data Center Optimization and Design

Cisco data center networking best practices give customers guidance and assistance in developing the data center network architecture most appropriate to meet changing IT requirements. These best practices augment the Cisco data center network architecture technologies and solutions to help IT architects and data center professionals take a phased approach to building and operating a comprehensive network platform for their next-generation data centers. By taking advantage of Cisco data center networking best practices, IT professionals can build a data center-class network, deploy solutions more quickly with lower risk, facilitate technology evolution and upgrades, and help ensure that IT staff are equipped with the right skills and expertise.

The benefits of data center optimization and design include the following:

- *Build and maintain a data center-class network*—Use validated and documented data center network solution designs to plan and implement networks that can achieve the stability and scalability required for mission critical data centers. By using proven best practices, community colleges can minimize downtime and accelerate recovery from disruptions. These designs also provide a robust foundation that customers or Cisco Advanced Services can use to make customizations to meet specific requirements.
- *Deploy solutions more quickly, with less risk and complexity*—Use Cisco data center best practices and designs to reduce the time, cost, and investment required for pre-production testing. Tried and tested designs help avoid the risks associated with technology disruptions, security exposure, non-scalable designs, and inappropriate software selection.

- *Facilitate technology evolution and upgrades*—The data center network is evolving to meet the challenges associated with cost, business alignment, resilience, and facilities concerns such as power and cooling. Cisco data center network best practices are constantly updated to incorporate these changes, so that customers can adopt them in a timely manner, with minimal risk.
- *Accelerate knowledge transfer*—The expertise and skills required to design and maintain increasingly sophisticated integrated data center networks are provided through constant training and knowledge transfer programs and infrastructure services. These programs include specialized Cisco CCIE® training such as the storage specialization, data center training labs, Cisco Press® books, Cisco Networkers, and executive briefing sessions.

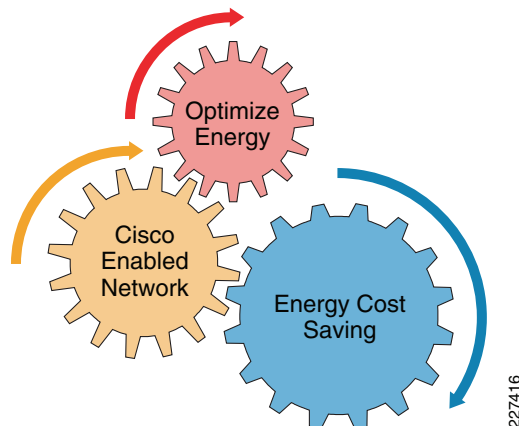
## Facilities Management

In response to energy costs, environmental concerns, and government directives, there is an increased need for sustainable and “green” IT operations at community colleges. Methods to measure power consumption and control energy output are now the focus of businesses worldwide, with all customers looking for a method to reduce energy costs and implement increased efficient operation.

Cisco EnergyWise is a new energy management architecture that allows IT operations and facilities to measure and fine-tune power usage to realize significant cost savings. Cisco EnergyWise focuses on reducing power utilization on all devices connected to a Cisco network ranging from Power-over-Ethernet (PoE) devices such as IP phones and wireless access points to IP-enabled building and lighting controllers. It uses an intelligent network-based approach, allowing IT and building facilities operations to understand, optimize, and control power across an entire campus infrastructure, potentially affecting any powered device.

This section illustrates how community colleges can use Cisco EnergyWise with a network enabled by Cisco to better understand the power footprint of their organization and optimize to reduce energy costs (see [Figure 1-9](#)).


**Figure 1-9 Cisco EnergyWise Optimize and Cost Saving**



Cisco EnergyWise is an energy management architecture designed to measure power consumption and optimize power usage, resulting in effective delivery of power across the campus. Community college IT professionals can quickly optimize the power consumed in a building and the result is immediate cost saving with a clear return on investment.

Cisco EnergyWise measures current power consumption, can automate and take actions to optimize power levels, and can advise how much power is being consumed to demonstrate cost saving. After power consumption is understood, regulation using Cisco EnergyWise network protocols provides command and control of power usage. Energy consumed per location can easily be found with a realistic view of power consumed per wiring closet, building floor, or campus building (see [Figure 1-10](#)).

**Figure 1-10 Cisco EnergyWise Optimized Power Delivery and Verification**

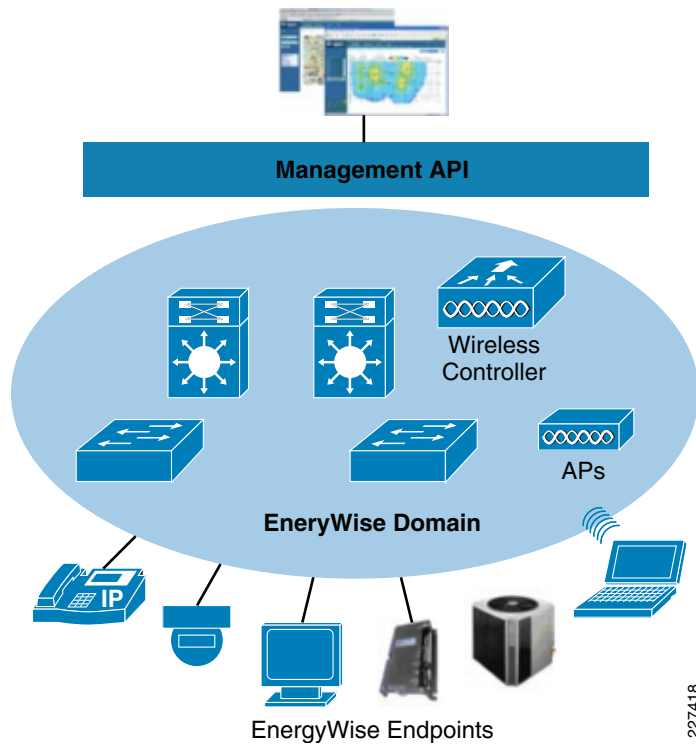


The Cisco EnergyWise network is used to intelligently and proactively manage power consumption and consistently enforce policies to provide lower energy consumption. Cisco EnergyWise has the ability to monitor, manage, and reduce energy use by creating visibility to how electricity is consumed and create the ability to turn devices from always on to always available based on business needs. Cisco EnergyWise offers orchestration and coordinated power management utilizing the Cisco network for scalability and communication. For example, when a staff member enters a building, a series of events can take place enhancing efficient building operation. An employee's badge access might trigger the office phone to power up, wireless access point coverage to be assured, computers to boot up, and temperature of the office to be brought to a proper value. As a result, the user of Cisco EnergyWise is saving energy by powering off components when they are not needed.

In many cases, individual management systems are dedicated to each type of device in a building, with management systems for building controls, another for phones, and another for access points. Today a large number of systems need to be integrated together to perform orchestration of events for power management. Disparate system integration is difficult to achieve and not always used. Cisco EnergyWise network wide policies can control device power management, eliminating the need for a myriad of systems to integrate and coordinate with each other. Orchestration is a primary benefit for the above scenarios and it is the Cisco network acting as a proxy of information that allows systems to communicate in a synchronized fashion that reduces complexity and costs, assuring power saving.

[Figure 1-11](#) depicts a typical Cisco network enabled by Cisco EnergyWise, including the management layer and endpoints.

Figure 1-11 Network Enabled by Cisco EnergyWise



The cost savings realized by using Cisco EnergyWise are significant. In many countries the government mandates saving energy for the business and proof of saving energy can provide financial incentives. As compared to today's typical campus building or branch, the savings realized by just controlling IT power devices is significant.

The Cisco Network Building Mediator ("Mediator") is the industry's first solution that extends the network as a platform to transform the way buildings are built, operated, and experienced. The Mediator:

- Enables energy reduction across global operations
- Takes advantage of Cisco's expertise in collaboration, convergence, and security to foster sustainable energy use
- Provides flexible integration of new technologies that deliver energy efficiency, clean energy, and environmental stewardship

The Mediator collects data from the building, IT, energy supply, and energy demand systems, which use different protocols. The Mediator then normalizes the data into a common data representation. This enables the Mediator to perform any-to-any protocol translation and to provide information to the end user in a uniform presentation.

This network-based framework creates a common, standards-based, open platform that allows campus applications, cloud services, and building/IT systems to communicate. The Mediator is protocol-agnostic and extends the network to serve as an effective foundation for sustainability management. The Mediator provides the following benefits:

- Reduced total cost of ownership (TCO)
- Simplified management of energy and facilities
- Flexible integration of building, IT, and clean technology systems
- Enhanced uptime and resiliency with networking technology

- Secure, high-quality delivery of concurrent building and IT services
- Future proofed investment with third-party applications and cloud services

The Mediator provides a network-based framework that interconnects four key systems: building, IT, energy supply, and energy demand. The integration of these disparate systems onto an IP network leads to a truly converged, energy-efficient building.

The Mediator's strategy is built on:

- *Any-to-any connectivity*—Building, IT, and “green” technologies
- *End-to-end management*—Efficiency, conservation, and decarbonization
- *Extensible platform*—Third-party applications and cloud services

## Secure Connected Classroom

### Classroom Connectivity to the Network

Providing connectivity to students while attending class is the foundation of 21st century learning, however it also poses many problems for community colleges. They must ensure that the person accessing the network should be allowed on the network and that the computer connecting to the network is free of viruses and other ailments that might adversely impact the network or others users. Secondly, while connectivity is provided, all steps should be taken to ensure the person connected is using the network for educational purposes and not illegal activities, such as sharing copyrighted material. Some community colleges chose to restrict the student to only access certain network resources while in class.

The density of wireless users in one location can also be problematic. Wireless designs must take into consideration the number of users, radio interference, and network utilization. The Community College reference design addresses these challenges in a variety of ways.


### Network Admission Control for Guests and Students

Network admission control allows community colleges to stop unauthorized or noncompliant devices and users from propagating threats into the network. Cisco Network Admission Control (NAC) enforces your institution's security policies and posture on all devices and users seeking network access.

Current business mechanisms such as Web 2.0, social networking, and cloud computing increase the likelihood of sensitive data residing outside of controlled devices. Traditional security products designed to protect closed environments with well-defined security boundaries are not effective in the new Web 2.0 environment.

Cisco NAC prevents loss of sensitive information by giving institutions a powerful, role-based method of allowing only compliant and authorized access and improving network resiliency. With Cisco NAC, only compliant and trusted endpoints—from PCs to printers, IP phones, and PDAs—are allowed onto the network, thereby limiting the potential damage from emerging security threats and risks.

Figure 1-12 Cisco Network Admission Control



## Application and Network Control

Cisco NAC helps reduce the potential loss of sensitive information by enabling organizations to verify a user's privilege level before granting network access. When that access is granted, the user is placed into a "role." Using role-based access control, community colleges can define security policies based on the role of the person using the network. For example, if a student connects to the network in a classroom, they can be put in a "student" role, which can then control where they can go and what they can use on the network, internally or externally.

As students, faculty, and staff carry their laptops to external locations, it is critical that the security protection on each endpoint device is up to date. The security policy is applied when an endpoint device attempts to connect to the internal network. Cisco NAC provides comprehensive policy enforcement and support. Cisco NAC integrates with a wide range of endpoint security applications. It supports built-in policies for more than 350 applications from leading antivirus and other security and management software solution providers. Many user-friendly capabilities, such as silent remediation and auto-remediation, help bring devices into compliance without causing user impact.

Cisco NAC helps community colleges provide secured guest access and assigns internal user access based on a user's role in the organization. Secure guest access allows visitors and guests to utilize the network without sacrificing the network security of the community college.

Cisco NAC provides full integration with wireless, VPN, and 802.1X and can be implemented in a single-sign-on (SSO) manner to maximize security benefits and minimize user impact.

Controlling peer-to-peer and instant messaging applications present several challenges, especially in the community college environment. Peer-to-peer applications, such as Gnutella and BitTorrent, are often used to share copyrighted material, such as music and movies, and instant messaging applications, like yahoo IM or AIM, can be used in the education environment as a way to pass notes in class. Both can be a challenge to control as often they will use common application ports such as port 80, which is also used to connect to Web pages, so just turning off the port is not an option. Educational institutions need to be able to look deeper inside the packet that is going across the network to ensure that these ports are not being used to circumvent security policies. Cisco has several ways of inspecting this traffic to ensure security compliance.

## Campus Safety and Security

Cisco physical security solutions provide broad network-centric capabilities in video surveillance, IP cameras, electronic access control, and ground breaking technology that converges voice, data, and physical security in one modular appliance. Our connected physical security solution enables community colleges use the IP network as an open platform to build more collaborative and integrated physical security systems while preserving their existing investments in analog-based technology. As customers converge their physical security infrastructures and operations and begin using the IP network as the platform, they can gain significant value through rapid access to relevant information and interoperability between other IP-centric systems. This creates a higher level of situational awareness and allows intelligent decisions to be made more quickly.

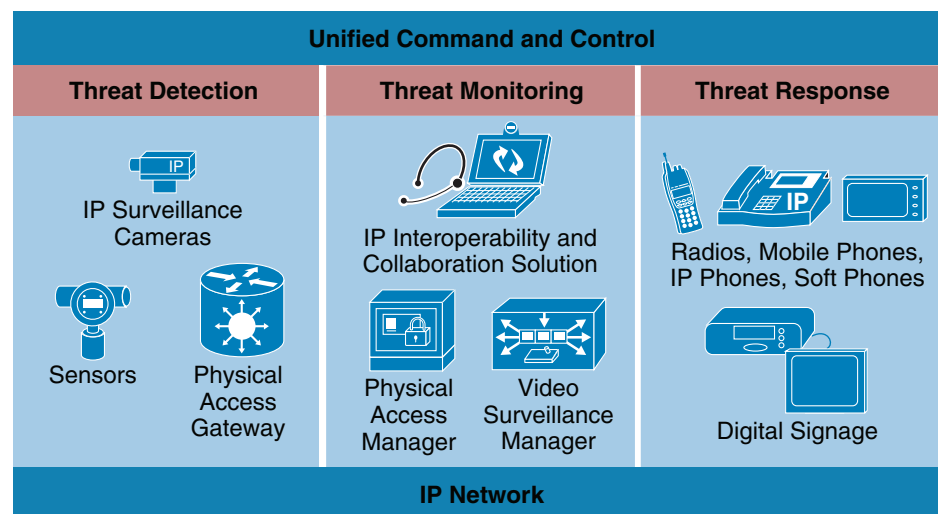
Cisco enables customers to build cost-effective, modular physical security solutions that are both best in class and interoperable. Cisco physical security products support the company's vision of a single unified security product suite that enables integration with all security operations within the IP network and with many non-security applications. Using the network as an open, scalable platform for integrating security provides community colleges with several benefits, such as operational flexibility, greater protection capabilities, lower cost of ownership, and reduced risk.

The Cisco Open Platform for Safety and Security is a platform architecture that proactively protects students, faculty, and staff through a scalable, tested network design. The architecture provides a more complete common operating picture, improves decision and response cycle times, and takes advantage of the network to expand the range and effectiveness of your emergency operations teams.

The platform takes advantage of a converged, IP network and provides the following benefits:

- Increases student, faculty, and staff safety and security through emergency notification and early warning
- Improves risk mitigation by facilitating continuity of operations (COOP), crisis management, all-hazards incident response, as well as facilities and critical infrastructure protection
- Reduces cost of operations
- Overcomes interoperability issues

**Figure 1-13** Unified Command and Control



## IP-Based Video Surveillance

Every day, you strive to make your schools as safe as possible. You develop plans, deploy systems, and train your staff on how to prevent, deter, detect, and respond to safety incidents. And you are doing a great job. Statistics show that community colleges continue to reduce the number of safety incidents.

For many decades, video surveillance has been a key component of the safety and security groups of community colleges. As an application, video surveillance has demonstrated its value and benefits countless times by:

- Providing real-time monitoring of a facility's environment, people, and assets
- Recording events for subsequent investigation, proof of compliance, and audit purposes

As security risks increase, the need to visually monitor and record events in an institution's environment has become even more important. Moreover, the value of video surveillance has grown significantly with the introduction of motion, heat, and sound detection sensors as well as sophisticated video analytics. Video surveillance can be integrated with and complement access control policies, providing video corroboration of access credential use.

These systems are realized through an open, standards-based, IP-network-centric functional and management architecture. As a network-centric company, Cisco has enabled the migration of many applications and systems onto a converged infrastructure. As a global enterprise organization, Cisco has developed and adopted a network-centric system architecture that meets the extensive requirements for a world-class video surveillance system.

The Cisco video surveillance architecture provides several benefits:

- Increased reliability and availability
- Greater utility (any camera to any monitoring or recording device for any application)
- Increased accessibility and mobility
- Multivendor video surveillance system “best of breed” interoperability
- The ability to enhance other building management system capabilities through improved interoperability

## Communicate Campus Events and Emergencies with Digital Signage

Traditionally, campuses have advertised events on posters tacked to bulletin boards around campus. The drawbacks of paper-based communications include clutter, out-of-date information, the time needed to constantly put up and take down posters, and paper waste.

Cisco Digital Signage provides more timely and eye-catching communications that can be scheduled to appear in different parts of the campus. Install the networked digital signs in high-traffic areas such as the entrances to buildings, student union, and faculty lounge areas, then display information about campus events and up-to-date emergency alerts and instructions. Assign any staff person, not necessarily an IT staff member, to use the interface to schedule content. You can even deliver different content to different signs—for example, promoting plays in the Theater Department building and advertising specials in the book store.

Popular uses of digital signage in community colleges include:

- Emergency notifications and instructions
- Event announcements, such as sports, guest speakers, registration/drop deadlines, etc.
- Classroom changes
- Student and staff group training



- Advertising in bookstores and stadiums
- Way finding
- Information for major events, such as graduation or donor recognition receptions
- Room scheduling

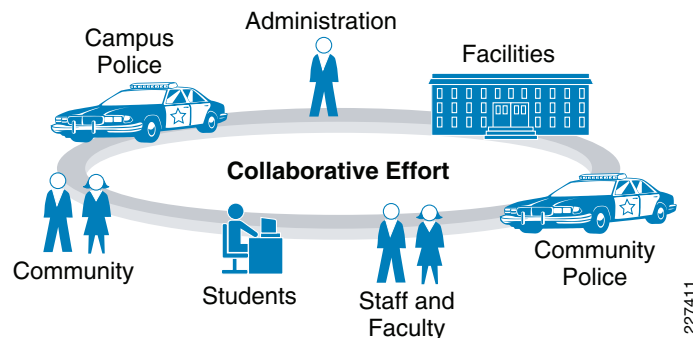
## IPICS for Emergency Collaboration

An emergency by definition is a chaotic event. Whether the emergency is a motor vehicle accident, a crime in progress, or a natural disaster that strikes a wide area, those who are responsible for responding require real-time, accurate information in order to effectively manage the event. Responding agencies—traditional first responders (police, fire, and emergency medical services), allied agencies (such as power utilities or other enterprises), or nongovernmental organizations such as the Red Cross and Red Crescent—need to work efficiently together to mitigate the effects of the incident.

Push-to-talk (PTT) Land Mobile Radio (LMR) systems have been the backbone of emergency response for decades. Unfortunately, one of the problems of LMR has been a legacy of incompatibility. Radios that do not use the same frequencies, LMR vendor-proprietary enhancements to established standards, and high infrastructure costs have led to a fractured LMR landscape that prevents effective coordination. Agencies that may have to work together may not be able to talk to each other. According to a report prepared by COMCARE, the United States alone has more than 100,000 emergency response agencies, most of which cannot easily communicate with each other or the public.

Another challenge is that responders now need to communicate with devices other than LMR systems, including Sprint/Nextel Push-To-Talk (PTT) phones, IP phones, and PCs. Technology is no longer an optional or a luxury item for emergency response. In an increasing number of cases, technology is vital to the situational awareness, span of control, scalability, and efficiency of incident response. However, incompatible communications technologies also build barriers that complicate interagency collaboration. Organizations must be able to break down these communications silos to realize the full benefit of their technology investments and to operate efficiently.

**Figure 1-14** Campus Safety



Cisco IPICS provides simple, scalable, comprehensive communications interoperability that encompasses radio networks, IP and non-IP networks, telephones, cell phones, and PC clients. Benefits of the Cisco IPICS solution include:

- *PTT everywhere*—By extending PTT and voice services from the LMR networks to IP networks, Cisco IPICS provides communications interoperability between wired and wireless networks.

- *Flexible and efficient operations and incident management*—Cisco IPICS provides an easy-to-use, Web-based interface for managing users, user groups, and radio channels across multiple networks and operational domains. Resources can be quickly added and then removed when no longer necessary, allowing graceful escalation and de-escalation based on the incident scope.
- *One-click activation of predefined policies*—Cisco IPICS Policy Engine, new in Cisco IPICS, enables administrators to create policies that define standard operating procedures—including talk group establishment and user notification—and then activate those policies with a single click. Notification methods can include radio, cell phone, public switched telephone network (PSTN) phone, Cisco Unified IP phone, Cisco IPICS Push-to-Talk Management Center (PMC) Client, pager, E-mail, or Short Message Service (SMS) text message. (Some methods require a Simple Mail Transfer Protocol [SMTP] gateway.) The agency defines policies using an intuitive, Web-based interface.
- *Customization*—Cisco IPICS can be customized to meet organizations' individual requirements. As an organization's needs change over time, Cisco IPICS can adapt with them.
- *Low cost and investment protection*—Cisco IPICS enables comprehensive communications interoperability at a fraction of the cost of replacing existing radio systems. By capitalizing on existing communications networks and devices, Cisco IPICS avoids the expense of unnecessary upgrades to existing radio networks. Furthermore, by enabling a graceful migration to IP networks and services, Cisco IPICS protects what can be a significant investment in traditional radio networks and devices. Agencies can also eliminate the expense of purchasing radios for office personnel by using the Cisco IPICS PMC Client for PCs and laptops or the Cisco IPICS Phone Client for IP phones.
- *Unified command and control*—Dispatchers and incident commanders can manage operations from one or more locations using the Web-based Cisco IPICS Administration Console.
- *Standards compliance*—Cisco IPICS takes advantage of industry-standard hardware and a proven IP architecture to create a framework for interoperable voice, video, and data communications. Organizations that currently use multiple wireless devices, including PTT, cellular, and wireless LAN (WLAN), can smoothly migrate to Cisco IPICS, which provides the infrastructure and feature set needed to achieve wide-ranging business and service goals. A standards-based solution also gives organizations the flexibility to add communications devices from any vendor.

Controlling physical access into buildings, rooms, and labs traditionally meant the use of an independent security network. The Cisco Physical Access Control solution is scalable and flexible, able to manage from one to several thousand doors. With this solution, institutions can combine modules to customize solutions and to manage the entire system remotely. In addition, this physical access solution easily integrates with Cisco's Video Surveillance solution and can use IP network services.

The Cisco Physical Access Gateway is an intelligent, distributed processing networking edge device module that connects door hardware, such as locks and readers, to the network. Accessory modules are available to handle additional doors and input/outputs.


The Cisco Physical Access Manager is the management application is used to configure hardware, monitor activity, enroll users, and integrate with IT applications and data stores. The data it collects can easily be shared with other security devices using the Cisco Open Platform for Safety and Security to create a holistic security view of the campus.

# Conclusion

The Cisco Community College reference design is built upon a highly resilient and flexible service fabric to provide community colleges with design solutions to solve business problems. It provides solutions that enable a 21st century learning environment, allowing for highly interactive and collaborative learning and teaching experiences while delivering any content, anytime, anywhere to any device.

To learn more about the Cisco Community College reference design, refer to the following URL:  
<http://www.cisco.com/go/education>






## CHAPTER 2

# Community College Reference Design—Service Fabric Design Considerations

The service fabric is the foundational network which all Community College services, applications, and solutions use to interact and communicate with one another. Service fabric is the most important component of the Community College reference design. If it fails, all applications, solutions, and technologies employed in the Community College reference design 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 2-1](#).


**Figure 2-1** Service Fabric Foundation Network



# Service Fabric Design Model

The model used for the Community College reference design service fabric is based around the desire to represent as many community college environments as possible. To do that a modular design is used, represented by campuses and buildings of varying sizes (see Figure 2-2). The campuses are made up of one or more building, depending on the campus 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 representing a classroom, an average size of 35 students per classroom or lab is used. Additionally, it is expected that half of all network can be accessed via wireless. This approach allows the network architect to essentially build their own community college environment by mixing the different campus and building profiles provided.

Figure 2-2 Community College Reference Design Overview



## Main and Large Campus Design

The main and large campus designs are meant to represent significantly sized campuses containing the largest student, faculty, and staff populations. The profile of the main/large campus 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 data center design and service block. The large campus will connect to the main campus via a 1Gb Metro Ethernet link. The main campus and large campus are almost identical, with the exception that the main campus is

connected to outside entities such as the Internet, Internet2 (I2), regional networks, and the National Lambda Rail using the Internet edge components, and will also have all other campuses within its community college system connecting to it.

## Medium Campus Design

The medium campus design is targeted at community colleges campuses that have approximately 3 buildings ranging in size from medium to small. The buildings will connect to the medium campus core via multiple 10Gb links, and the core will also connect to a small data center and service block. The medium campus is connected to the main campus via a 100mb Metro Ethernet link. This link interconnects the medium campus to the other campuses as well as external networks such as the Internet and I2.

## Small Campus Design

The small campus profile represents a campus made up of just one building; in this case, the core and distribution networks are collapsed into one. The small campus is connected to the main campus via a fractional DS3 with a 20mb bandwidth rating. This link interconnects the small campus to the other campuses as well as external networks such as the Internet and I2.

## 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 is made up of 80 classrooms, 30 professor offices, 10 administrative offices, and 40 college professionals collectively this represents 160 phones for the large building.

### 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 40 classrooms, 15 professor offices, 5 administrative offices, and 20 college professionals collectively this represents 80 phones for the medium building.

## 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 made up of 10 classrooms, 8 professor offices, 2 administrative offices, and 10 college professionals collectively this represents 30 phones for the small building.

## 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 is made up of 3 classrooms and 7 other phones, totaling 10 phones for the extra small building.

## Access Devices

The devices that connect to the Cisco Community College reference design 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 campus 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 campuses. 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 for community college 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 campus 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:

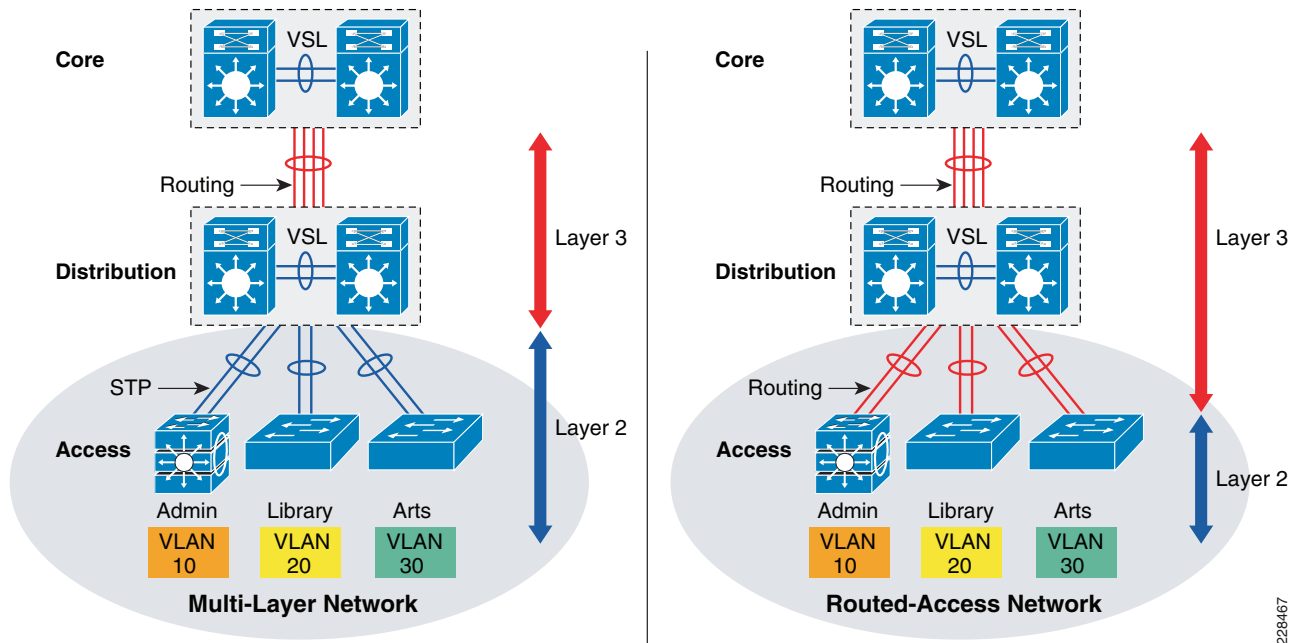
- *Link resiliency*—Provides redundancy during physical link failures (i.e., fiber cut, bad transceivers, incorrect cabling, etc.)
- *Device resiliency*—Protects network during abnormal node failure triggered by hardware or software (i.e., software crashes, non-responsive supervisor etc.)
- *Operational 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 2-2](#).

Figure 2-3 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 community college system is paramount. Some of the key foundational services discussed include the following:

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

## WAN Design Considerations

### WAN Transport

In order for campuses to communicate with one another and/or to communicate outside the community college system, 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 community college 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 Community College reference design service fabric is essential. Without it, community college 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, Internet2 and National LambdaRail (NLR) networks and protecting internal resources and users from malware, viruses, and other malicious software. Protecting students, staff and faculty 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 student, staff and faculty records.
- *Network access security and control*—Securing the access edges. Enforcing authentication and role-based access for students, staff and faculty residing at the main and remote campuses. Ensuring systems are up-to-date and in compliance with the CCVE institution's network security policies.
- *Network endpoint protection*—Protecting servers and school-controlled systems (computer labs, school-provided laptops, etc.) from viruses, malware, botnets, and other malicious software. Enforcing E-mail and web browsing policies for staff and faculty.

Each of these security design considerations are discussed in further detail in [Chapter 6, “Community College Security Design.”](#)

## Mobility

Mobility is an essential part of the community college environment. Most students will connect wirelessly to campus 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 students, staff and guests to be accessible and productive, regardless of whether they are meeting in a study hall, at lunch with colleagues in the campus cafeteria, or simply enjoying a breath of fresh air outside a campus building. Provide easy, secure guest access to college guests such as alumni, prospective students, 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. Faculty and other 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.
- *Manageability*—College network administrators must be able to easily deploy, operate, and manage hundreds of access points within multiple community college campus deployments. A single, easy to understand WLAN management framework is desired to provide small, medium and large community college systems with the same level of wireless LAN management scalability, reliability and ease of deployment that is demanded by traditional enterprise business customers.
- *Reliability*—Provide adequate capability to recover from a single-layer fault of a WLAN accessibility component or controller wired link. Ensure that wireless LAN accessibility is maintained for students, faculty, staff and visitors in the event of common failures.

## Unified Communications

### Call Processing Considerations

How calls are processed in the community college 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 a unified communications system, and an integral part of all call processing agents. Generally, the dial plan is responsible for instructing the call processing agent on how to route calls. Specifically, the dial plan performs the following main functions:

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

## Survivability Considerations

Voice communications are a critical service that must be maintained in the event of a network outage for this reason the service fabric must take survivability into consideration.





## CHAPTER 3

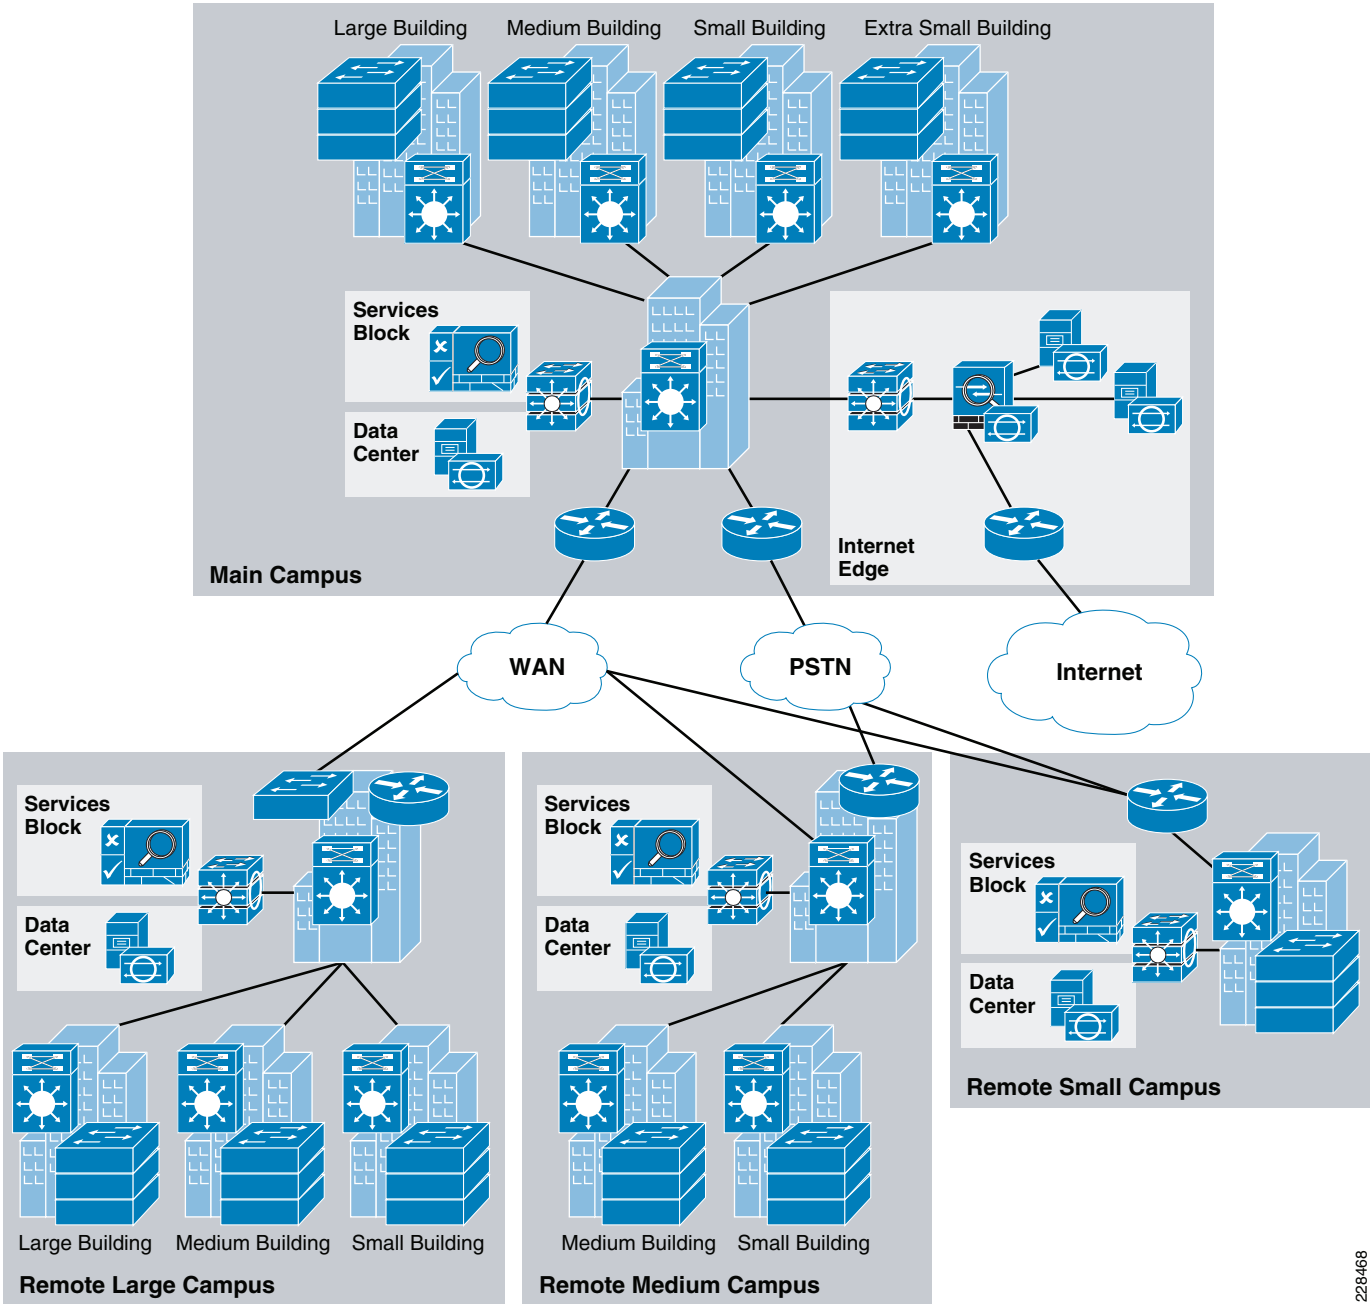
# Community College LAN Design

---

## LAN Design

The community college LAN design is a multi-campus design, where a campus consists of multiple buildings and services at each location, as shown in [Figure 3-1](#).

Figure 3-1 Community College LAN Design

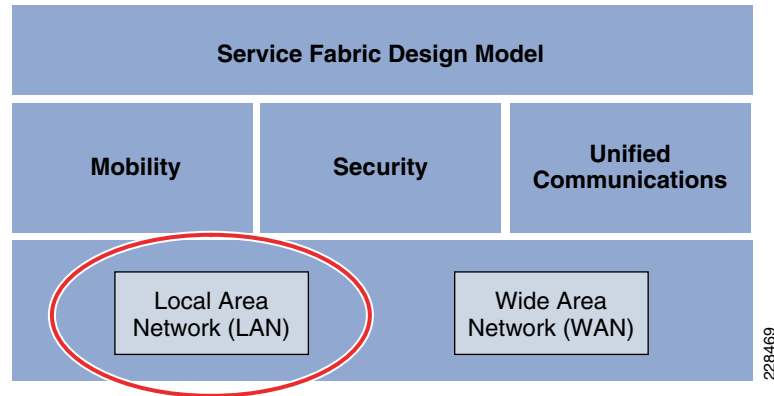


228468

Figure 3-2 shows the service fabric design model used in the community college LAN design.



**Figure 3-2 Community College 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 community college 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 community college*—Leverages the design principles of the tiered network design to facilitate a geographically dispersed college campus network made up of various elements, including networking role, size, capacity, and infrastructure demands.
- *Considerations of a multi-tier LAN design model for community colleges*—Provides guidance for the college 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 community colleges*—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

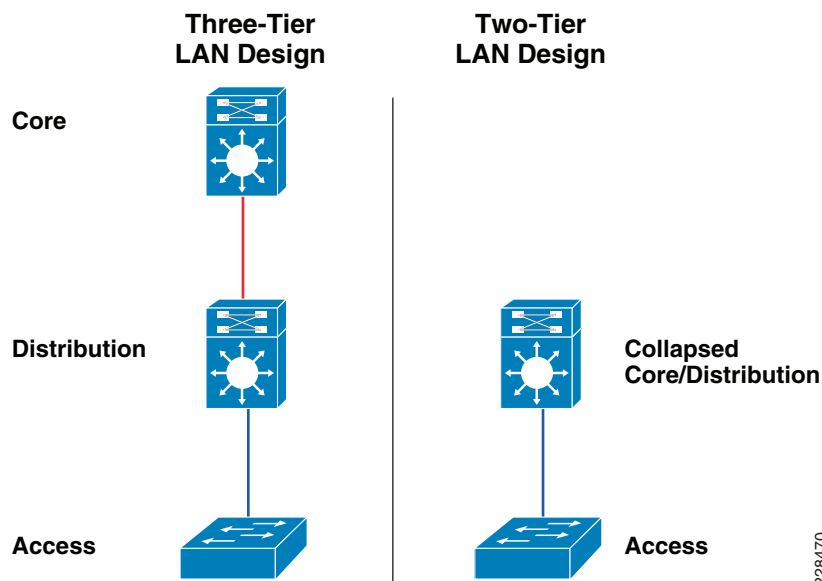
Any successful design or system is based on a foundation of solid design theory and principles. Designing the LAN component of the overall community college LAN service fabric design model is no different than designing any large networking system. The use of a guiding set of fundamental engineering 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 college and technology needs. This chapter provides design guidelines that are built upon the following principles to allow a community college network architect to build college 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 college 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 community colleges today.

Designing the community college 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 3-3](#).

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



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 community college.

- *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*

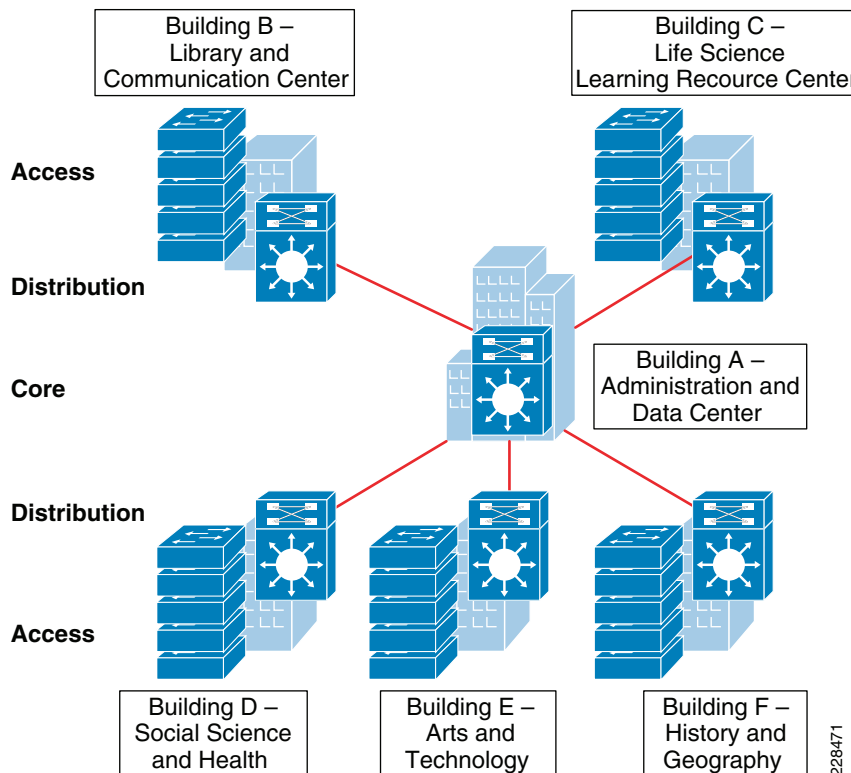
The core layer is the network backbone that connects all the layers of the LAN design, providing for connectivity between end devices, computing and data storage services located within the data center and other areas, and services within the network. The core layer serves as the aggregator for all the 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 3-4 shows a sample three-tier LAN network design for community colleges 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 3-4 Three-Tier LAN Network Design Example**

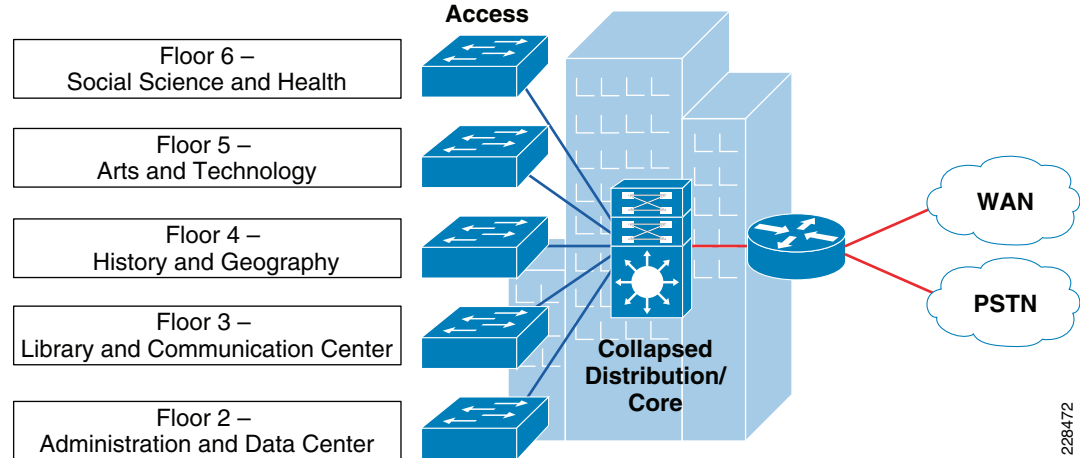


The primary purpose of the core layer is to provide fault isolation and backbone connectivity. Isolating the distribution and core into separate layers creates a clean delineation for change control between activities affecting end stations (laptops, phones, and printers) and those that affect the data 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 college 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 3-5 shows a two-tier LAN network design example for a community college LAN where the distribution and core layers are collapsed into a single layer.

**Figure 3-5 Two-Tier Network Design Example**



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

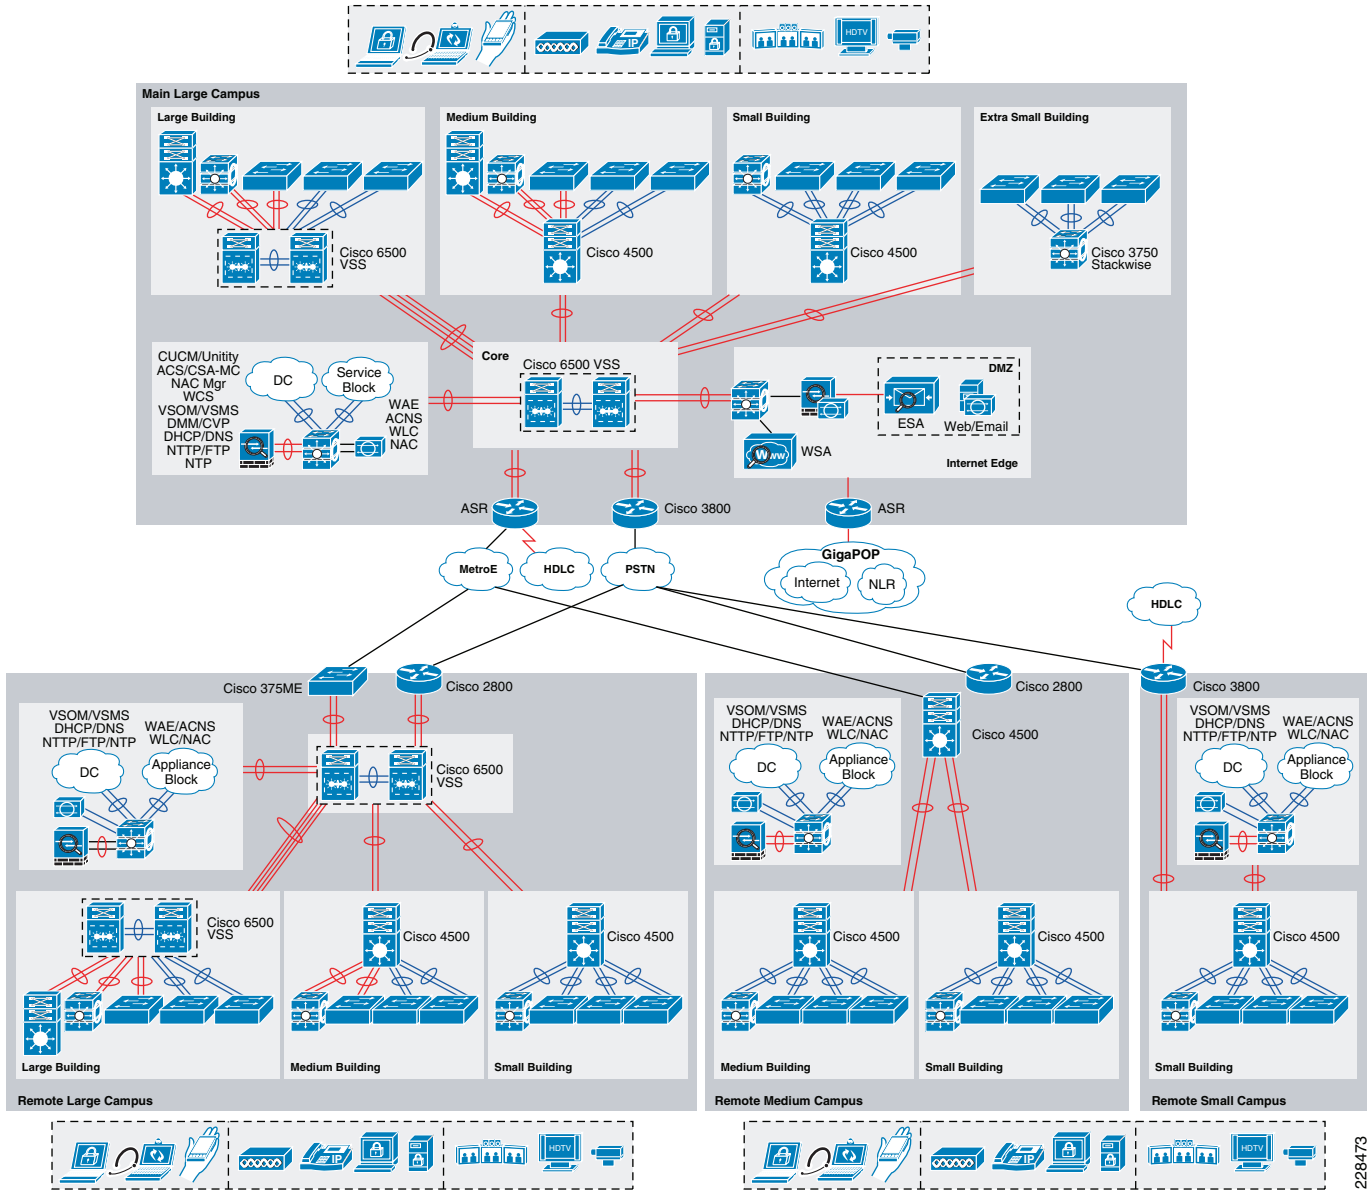
## Community College LAN Design Models

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 community colleges. 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 3-6](#), multiple campuses can co-exist within a single community college system that offers various academic programs.

Figure 3-6 Community College LAN Design Model



Depending on the number of available academic programs in a remote campus, the student, faculty, and staff population in remote campuses may be equal to or less than the main college campus site. Campus network designs for the remote campus may require adjusting based on overall college campus capacity.

Using high-speed WAN technology, all the remote community college campuses interconnect to a centralized main college campus that provides shared services to all the students, faculty, and staff, 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 3-1 shows a summary of the LAN design models as they are applied in the overall community college network design.

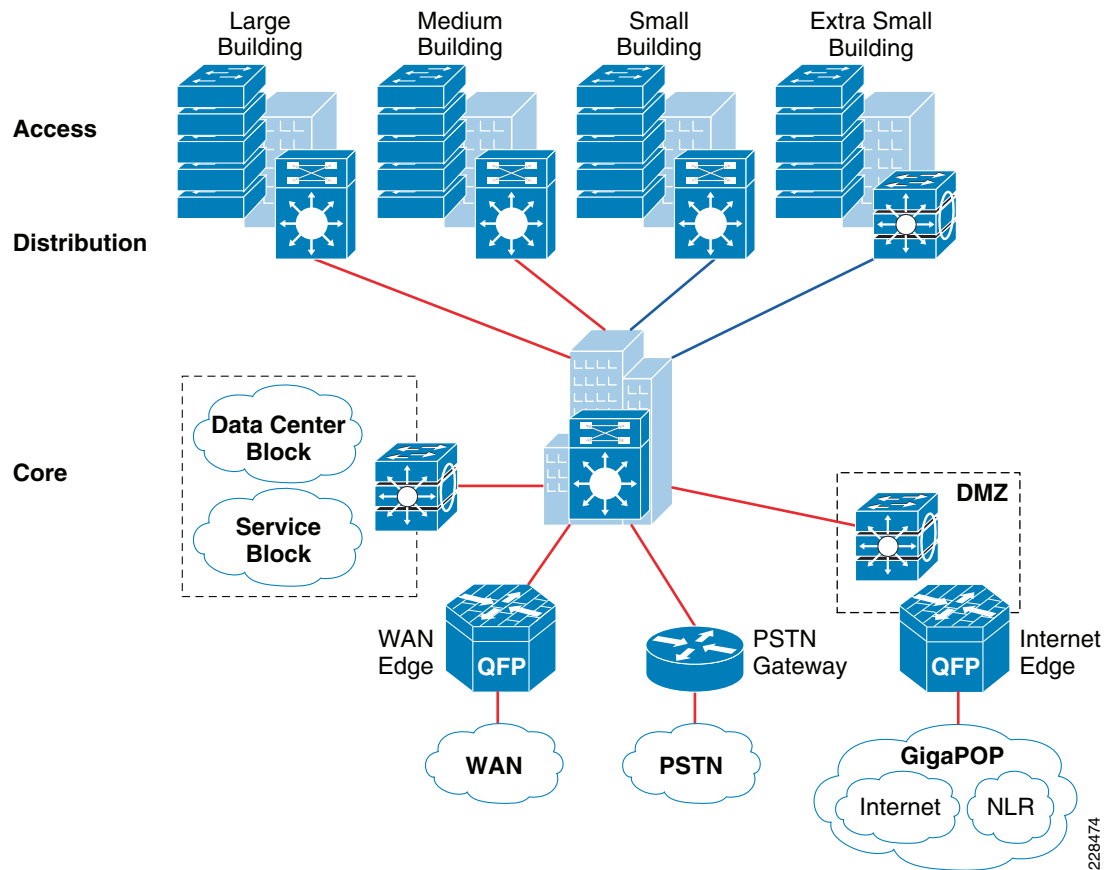
228473

**Table 3-1** Community College Recommended LAN Design Model

Community College 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 College Campus Network Design

The main college campus in the community college 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 3-7](#).

**Figure 3-7** Main College Campus Site Reference Design

The main college campus typically consists of various sizes of building facilities and various education department groups. The network scale factor in the main college campus site is higher than the remote college 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 3-8](#).

**Figure 3-8 Main College Campus Site Reference Design**



The three-tier LAN design model for the main college campus 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 college 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 students and staff for accessing the Internet as well as private education and research networks.



**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 4, “Community College WAN Design.”](#)

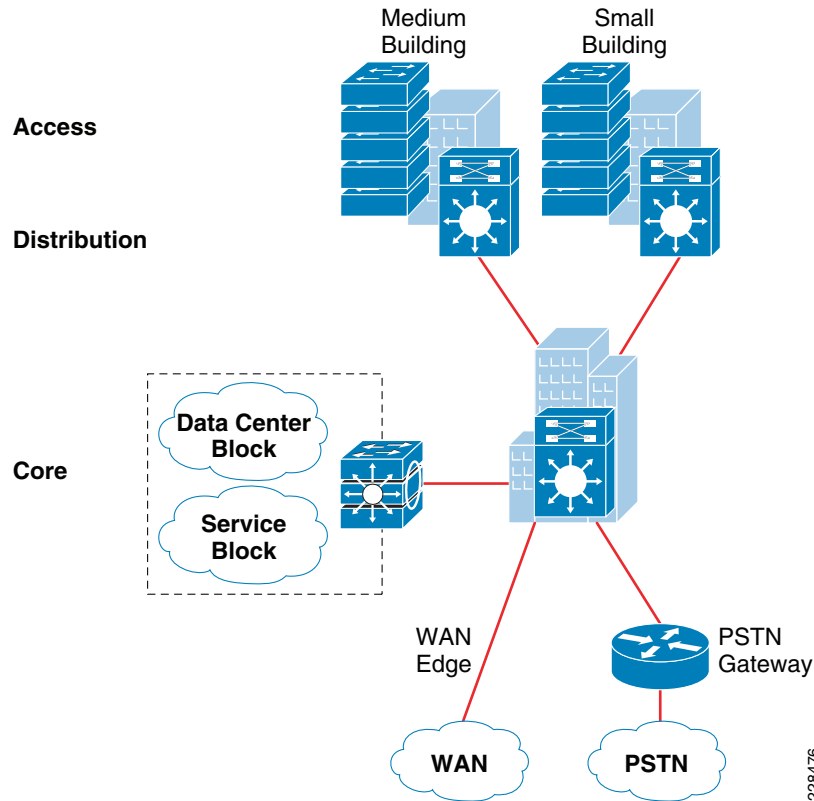
## Remote Large College Campus Site Design

From the location size and network scale perspective, the remote large college is not much different from the main college campus 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 college can also be considered as an alternate college campus to the main campus site, with the same common types of applications, endpoints, users, and network services. Similar to the main college campus, separate WAN devices are recommended to provide application delivery and access to the main college campus, given the size and number of students at this location.



Similar to the main college campus, Cisco recommends the three-tier LAN design model for the remote large college campus, as shown in Figure 3-9.


**Figure 3-9 Remote Large College Campus Site Reference Design**



## Remote Medium College Campus Site Design

Remote medium college campus locations differ from a main or remote large campus in that there are less buildings with distributed education departments. A remote medium college 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 college campuses. Because there are fewer students, faculty, and end users at this site as compared to the main or remote large campus sites, the need for a separate WAN device may not be necessary. A remote medium college 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 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 3-10 shows the remote medium campus in more detail.

**Figure 3-10 Remote Medium College Campus Site Reference Design**



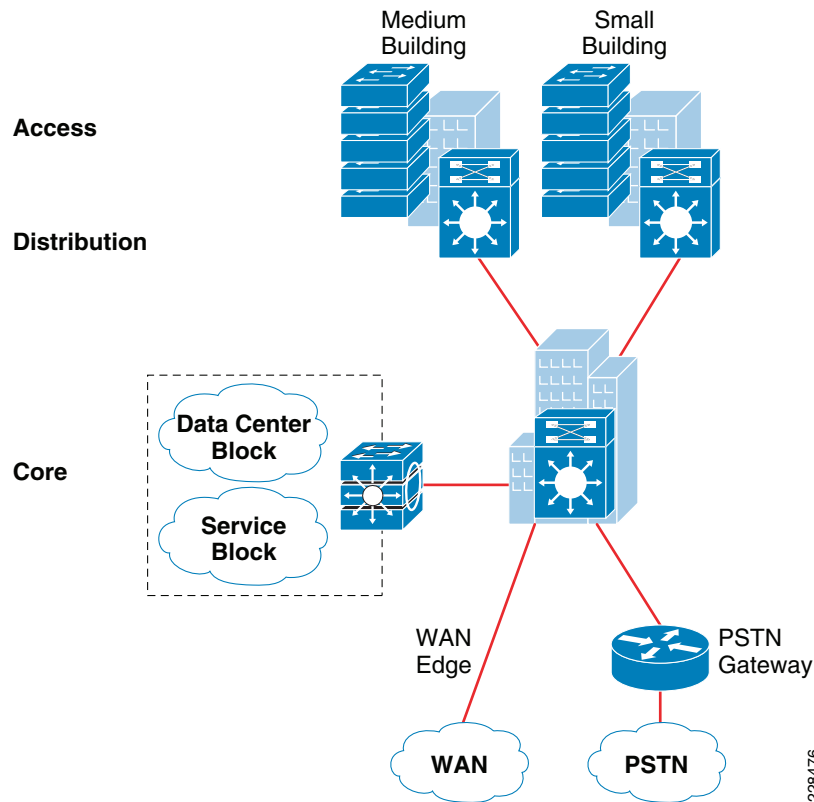
## Remote Small College Campus Network Design

The remote small college 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 college campuses. However, the application and services demands are still consistent across the community college 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 college campus locations, sufficient WAN link capacity is needed to deliver consistent network services to student, faculty, and staff. 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 3-11](#) shows the remote small campus in more detail.

**Figure 3-11 Remote Small College Campus Site Reference Design**



## Multi-Tier LAN Design Models for Community College

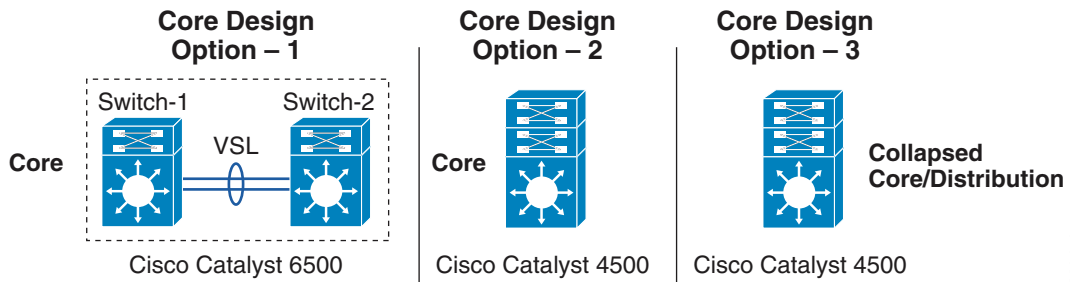
The previous section discussed the recommended LAN design model for each community college location. This section provides more detailed design guidance for each tier in the LAN design model. Each design recommendation is optimized to keep the network simplified and cost-effective without compromising network scalability, security, and resiliency. Each LAN design model for a community college 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 community college 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 Series or the Cisco Catalyst 4500 Series Switches. [Figure 3-12](#) shows the three core layer design models.

**Figure 3-12 Core Layer Design Models for Community Colleges**




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 community college design.

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

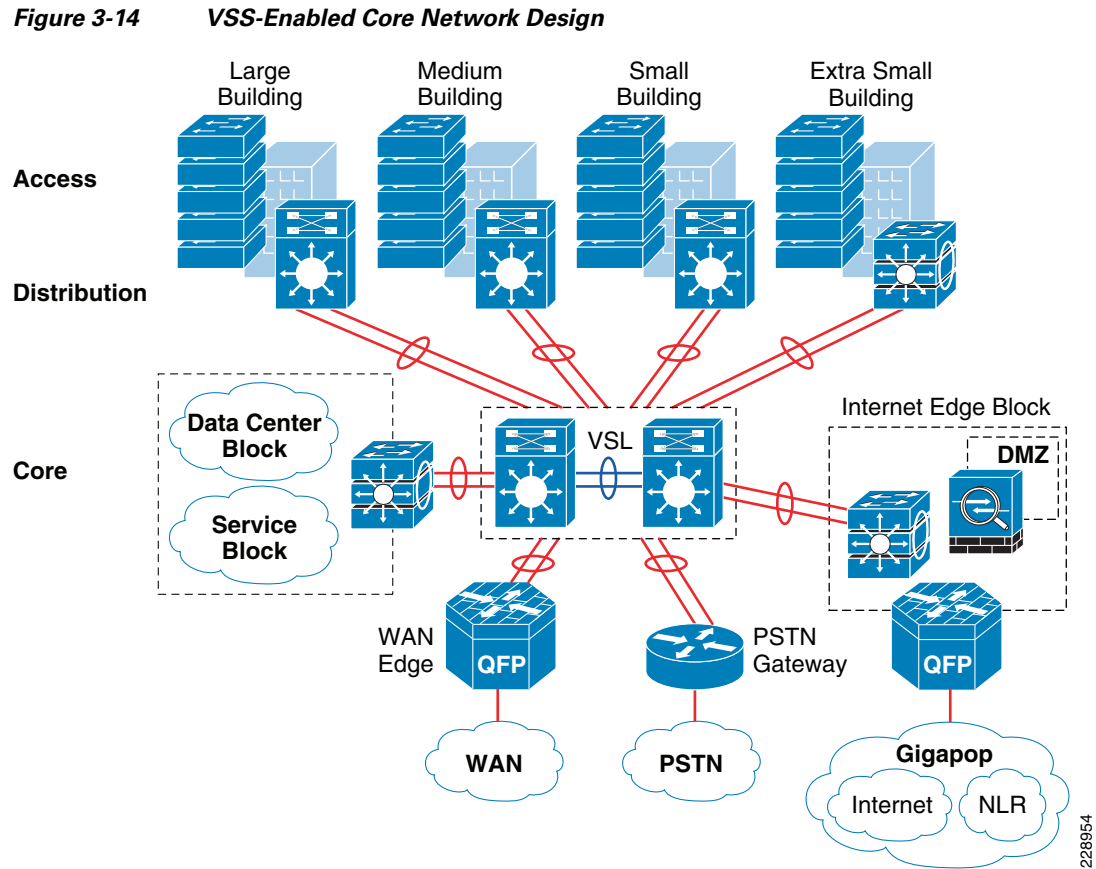
Core layer design option 1 is specifically intended for the main and remote large campus locations. It is assumed that the number of network users, high-speed and low-latency applications (such as Cisco TelePresence), and the overall network scale capacity is common in both sites and thus, similar core 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 3-13 shows the physical and operational view of VSS.

**Figure 3-13 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 3-14 shows a reference VSS-enabled core network design for the main campus site.

**Note**

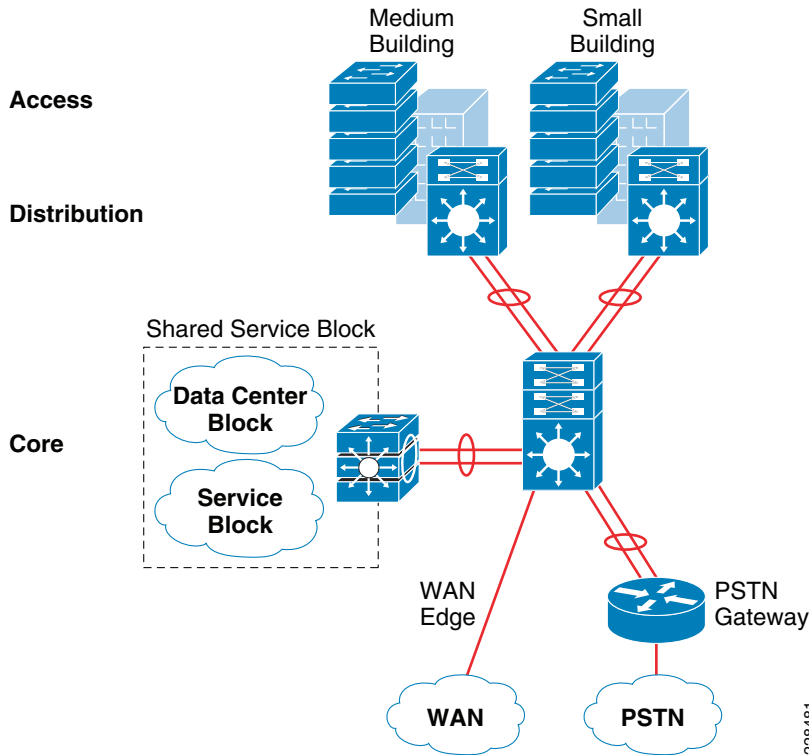
For more detailed VSS design guidance, see the *Campus 3.0 Virtual Switching System Design Guide* at the following URL:

[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-Based Campus Core Network


Core layer design option 2 is intended for a remote medium-sized college campus and is built on the same principles as for the main and remote large 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 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 3-15 shows the remote medium campus core design option in more detail.

Figure 3-15 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 3-16 shows the redundancy of the Cisco Catalyst 4500-E Series in more detail.

Figure 3-16 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 college network meets the following four key goals:

- *Scalability*—The modular Cisco Catalyst 4500 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.

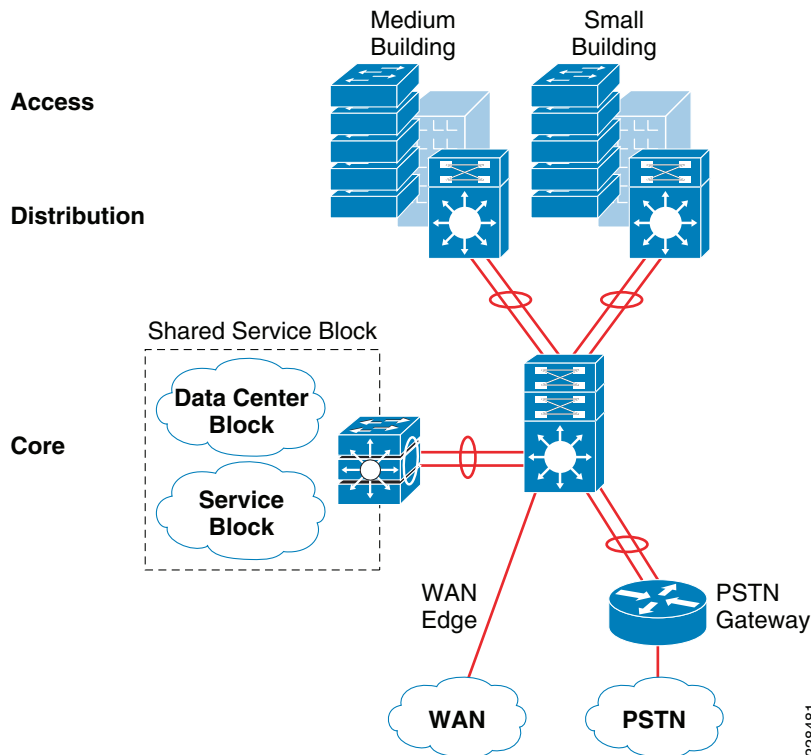
### Core Layer Design Option 3—Cisco Catalyst 4500-Based Collapsed Core Campus Network

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, 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 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 3-17](#) shows the remote small campus core design in more detail.

**Figure 3-17** 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 community college 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 college campus locations, depending on network scale, application demands, and cost, as shown in Figure 3-18. Each design model offers consistent network services, high availability, expansion flexibility, and network scalability.




Figure 3-18 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 campus locations, and is based on Cisco Catalyst 6500 Series switches using the Cisco VSS, as shown in Figure 3-19.

Figure 3-19 VSS-Enabled Distribution Layer Network Design



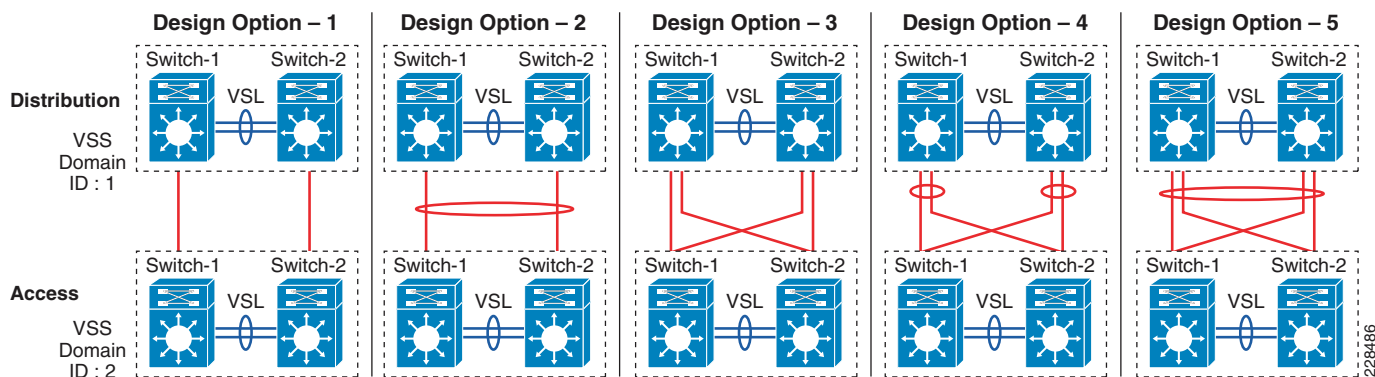
The distribution block and core network operation changes significantly when redundant Cisco Catalyst 6500-E Series switches are deployed in VSS mode in both the distribution and core layers. Clustering redundant distribution switches into a single logical system with VSS introduces the following technical 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 campus locations. Figure 3-20 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 3-20 Core/Distribution Layer Interconnection Design Considerations**



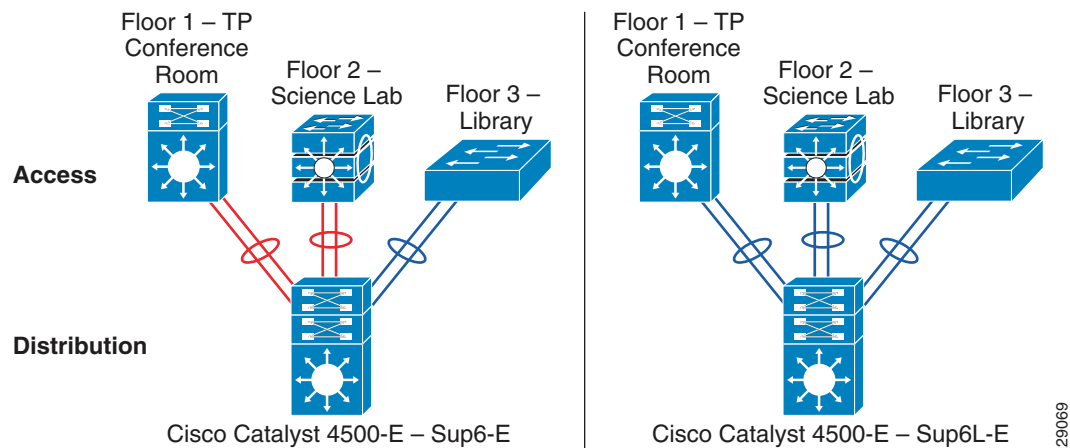
The various core/distribution layer interconnects offer the following:

- *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 as an aggregation layer system that offers consistent network operation like a VSS-enabled distribution layer switch. The Cisco Catalyst 4500 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 3-21](#).

**Figure 3-21** 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 college 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 3-22](#).

**Figure 3-22 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 Series switches. Although the remote medium and the main and remote large 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 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-E 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 3-23](#).

**Figure 3-23 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-E StackWise Plus Series platform for the distribution layer switch.

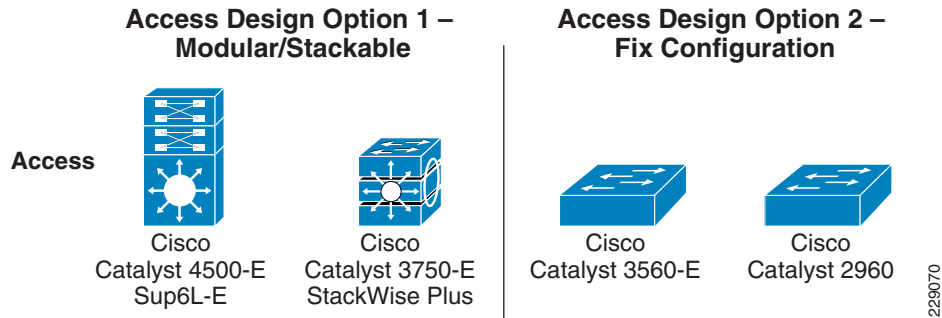
The fixed-configuration Cisco Catalyst 3750-E Series switch is a multilayer platform that supports Cisco StackWise Plus technology to simplify the network and offers flexibility to expand the network as it grows. With Cisco StackWise Plus technology, multiple Catalyst 3750-E can be stacked into a high-speed 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-E 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.

Community college 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 student laptops. Based on several endpoint factors such as function and network demands and capabilities, two access layer design options can be deployed with college campus network edge platforms, as shown in [Figure 3-24](#).

Figure 3-24 Access Layer Design Models



## Access Layer Design Option 1—Modular/StackWise Plus 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 college campus physical security protection, these devices require low latency, high performance, and a constant network availability switching infrastructure. Implementing a modular and Cisco StackWise Plus-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 college business and IT operation.

The Cisco Catalyst 3750-E 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-E Series Switches into a single high-speed stack ring that simplifies operation and allows incremental access layer network expansion. The Cisco Catalyst 3750-E Series leverages EtherChannel technology for protection during member link or stack member switch failure.

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

This entry-level access layer design option is widely chosen for educational 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 fixed configuration Cisco Catalyst 3560-E 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-E 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 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 college campus network, the Cisco Catalyst 3560-E is an alternate switching solution for the multilayer distribution block design option discussed in the previous section. The Cisco Catalyst 3560-E 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 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 Community College Network Foundation Services

After each tier in the model has been designed, the next step for the community college design is to establish key network foundation services. Regardless of the application function and requirements that community colleges 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 community college 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 college main campus and remote college campus locations. The Community College 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 Community College 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-E Stackwise
- Cisco Catalyst 3560-E 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, college 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.

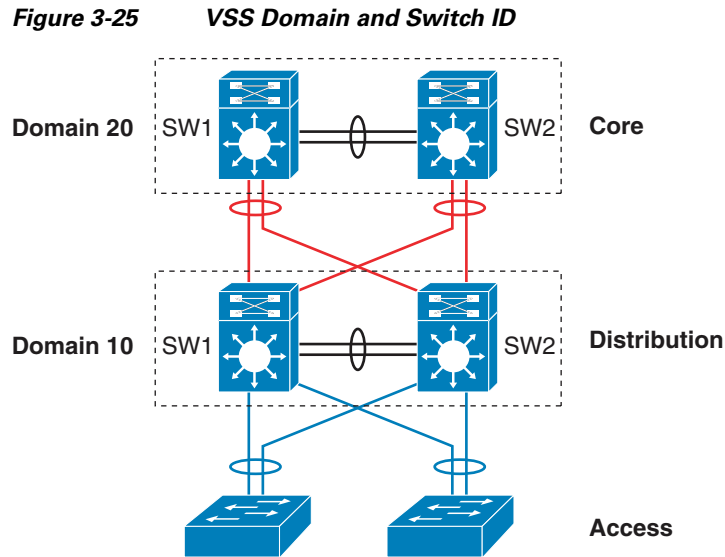
#### Domain ID

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 3-25](#).





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 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 3-26 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 3-26 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 21<sup>st</sup> century education communities 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.

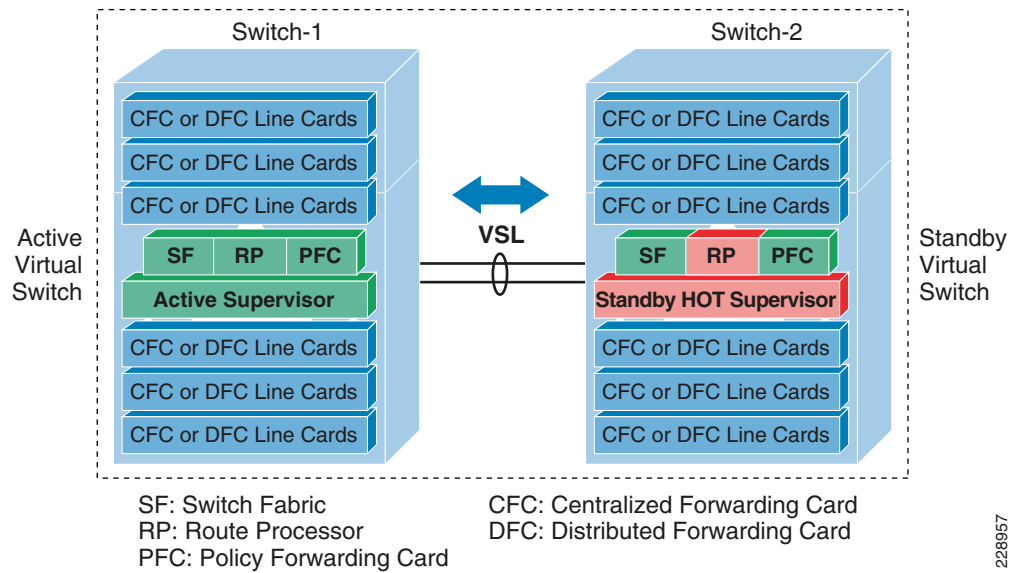
Implementing 10G mode may assist in increasing the number of transmit and receive queue depth level; however, restricted VSL QoS prevents reassigning different class-of-service traffic in different queues. Primary benefit in implementing 10G-only mode is to deploy VSL port in non-blocking mode to dedicate 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 3-27](#).

Figure 3-27 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 3-28 depicts the state of campus topology in a single active-state and during dual-active state.

Figure 3-28 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 3-29](#), 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 3-29 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 community college 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 community college 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 college 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 college campus network. Figure 3-30 illustrates Cisco Catalyst 4500-E-based intra-chassis SSO and NSF capability.

**Figure 3-30 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 college 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 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 3-31 summarizes the uplink port support on Sup6E model depends on non-redundant and redundant deployment scenario.

**Figure 3-31 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

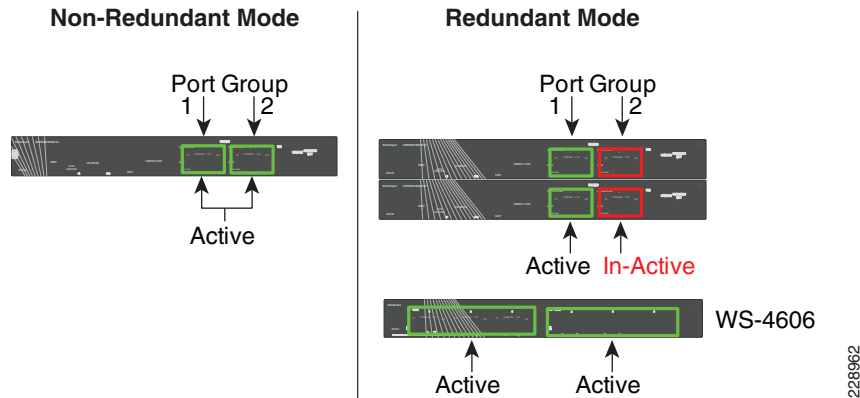
cr24-4507e-MB#show hw-module mod 3 port-group
Module Port-group ActiveInactive
-----
3      1      Te3/1-2Gi3/3-6

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 3-32 illustrates the deployment guideline for highly resilient Catalyst 4507R-E based Sup6L-E uplink.


Figure 3-32 Catalyst 4500-E Sup6L-E Uplink Mode



## Deploying Cisco Catalyst 3750-E StackWise Plus

Cisco Catalyst 3750-E 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 3-33](#).

Figure 3-33 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 3-2](#) lists the network protocols that are designed to operate in centralized versus distributed model in Cisco StackWise Plus architecture.

**Table 3-2 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-E 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-E 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:

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

cr36-3750s-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.

```
cr36-3750s-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.

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

```
cr36-3750s-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 150 Ready
* 2   Member 0026.5284.ec80 140 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 150 Ready
* 2   Member 0026.5284.ec80 140 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-E StackWise switch mode using the global configuration CLI mode as shown below:

```
cr36-3750s-xSB (config)#stack-mac persistent timer 0
cr36-3750s-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-E and 2960

The Community College Reference design recommends deploying fixed configuration Cisco Catalyst 3560-E 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-E 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.

## Designing EtherChannel Network

In this reference design, multiple parallel physical paths are recommended to build highly scalable and resilient community college 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 community college 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:

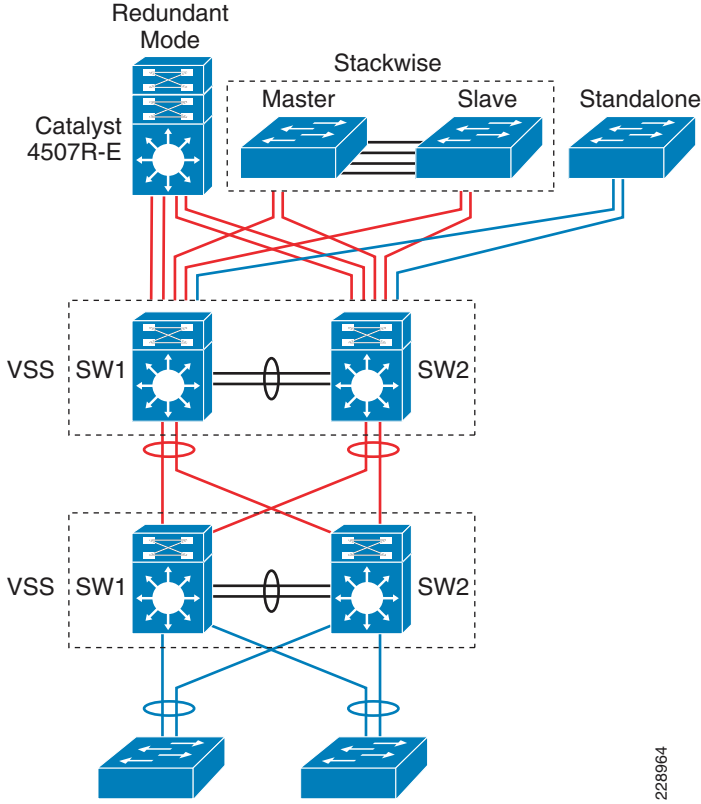
- Common network topologies and configuration (all campus network design)
- Simplifies network protocols (eases network operations)
- Increase network bandwidth capacity with symmetric forwarding paths
- Delivers deterministic network recovery performance

### 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). Independent of network tier and platform role, this design principle is applicable to all systems across college campus network. [Figure 3-34](#) demonstrates recommended deployment physical network design model for various Catalyst platforms.

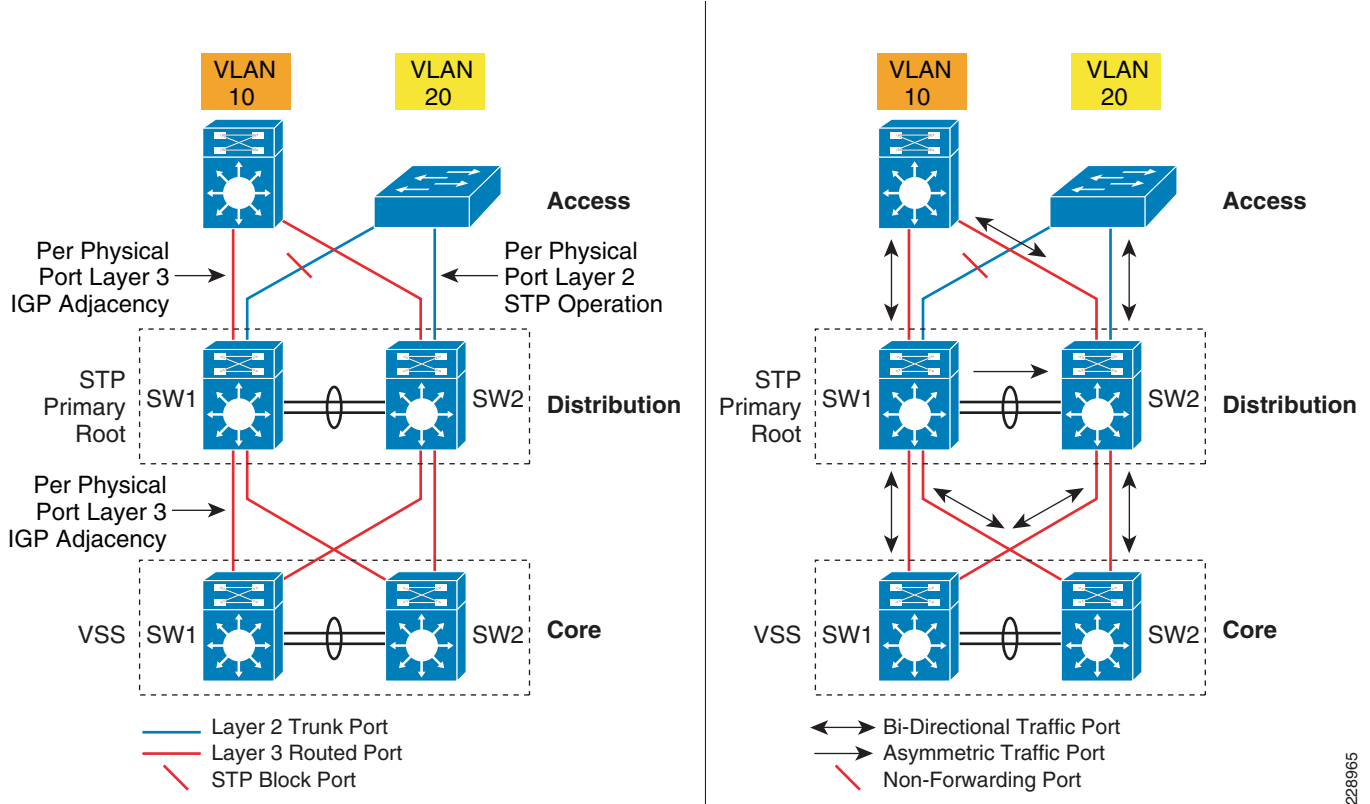
**Figure 3-34** Designing Diverse Full-mesh Network Topology



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 3-35](#). Such network design cannot leverage distributed forwarding architecture and increase operational and troubleshooting complexities. [Figure 3-35](#) demonstrates the default network design with redundant and complex control-plane operation with under-utilized forwarding plane design.



Figure 3-35 Non-optimized Campus Network Design



228965

The design in [Figure 3-35](#) 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 3-34](#), 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

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

### 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.

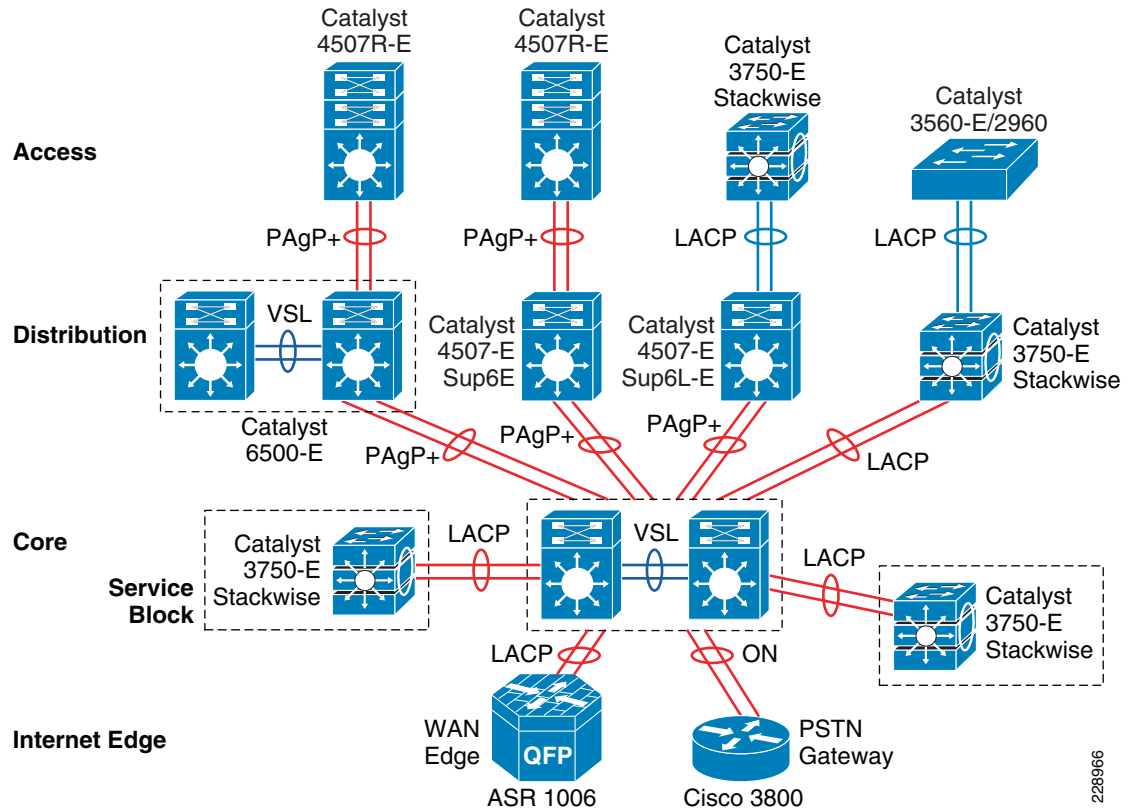
## Implementing EtherChannel

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 Community College Reference design support EtherChannel technology. Independent of college 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 3-36 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 3-3.

**Table 3-3 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

- 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 diversified 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 college 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 tunings 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 3-4](#) summarizes the currently supported EtherChannel load-balancing methods.

**Table 3-4 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-E, 3xxx-E, 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 community college network to create an efficient, scalable, and stable network design. Identifying an IP addressing strategy for the network for the entire community college 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 college 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 community college 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 college 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 community college 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 community college 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 community college 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 community college 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>.

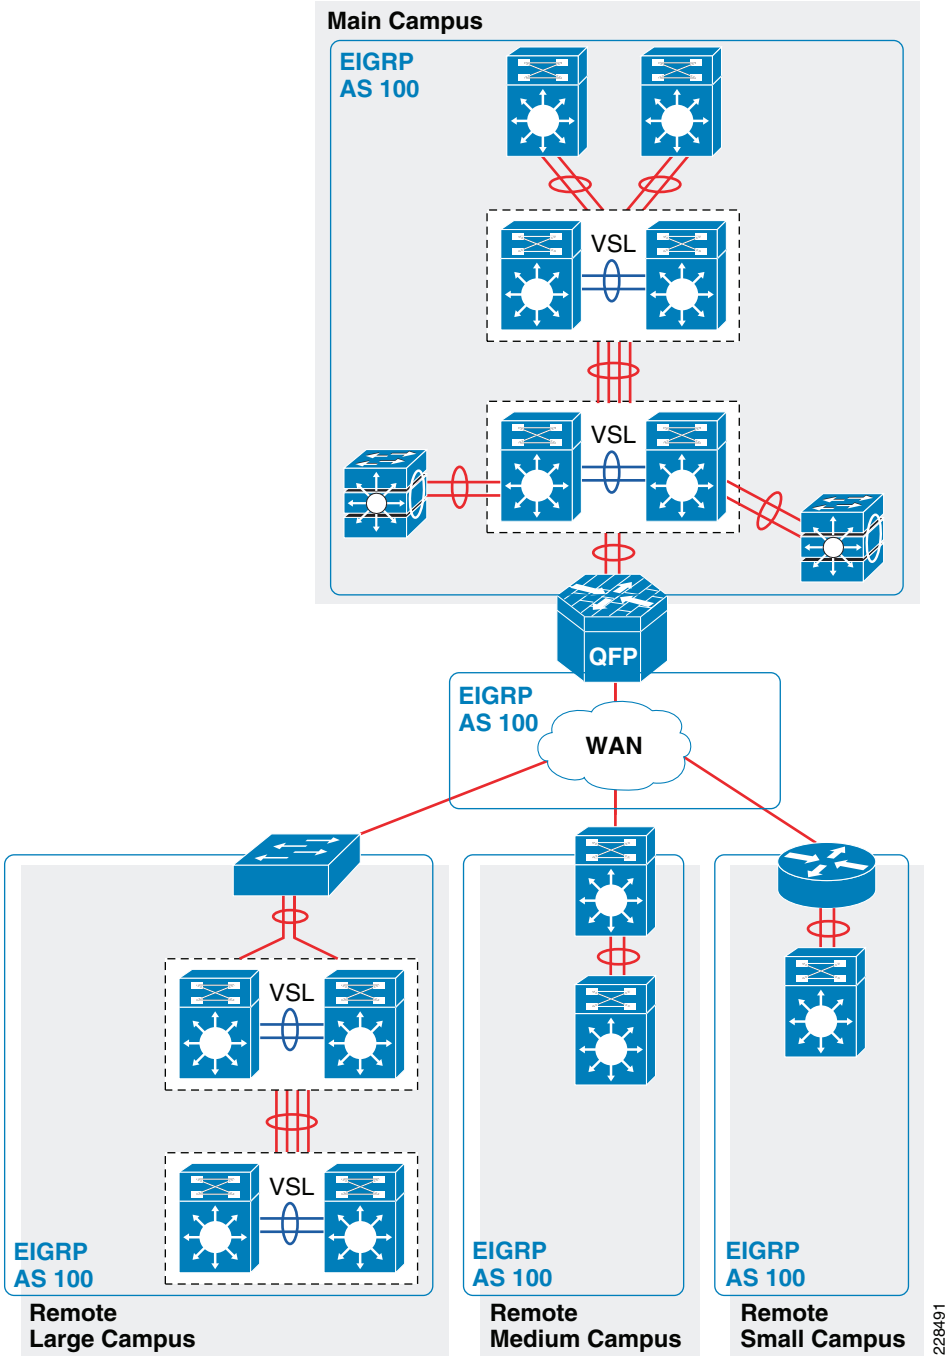


## Designing an End-to-End EIGRP Routing Network

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 community college LAN core layer network:

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

Figure 3-37 Sample End-to-End EIGRP Routing Design in Community College LAN Network



**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
```

228491

```

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)          (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:

```

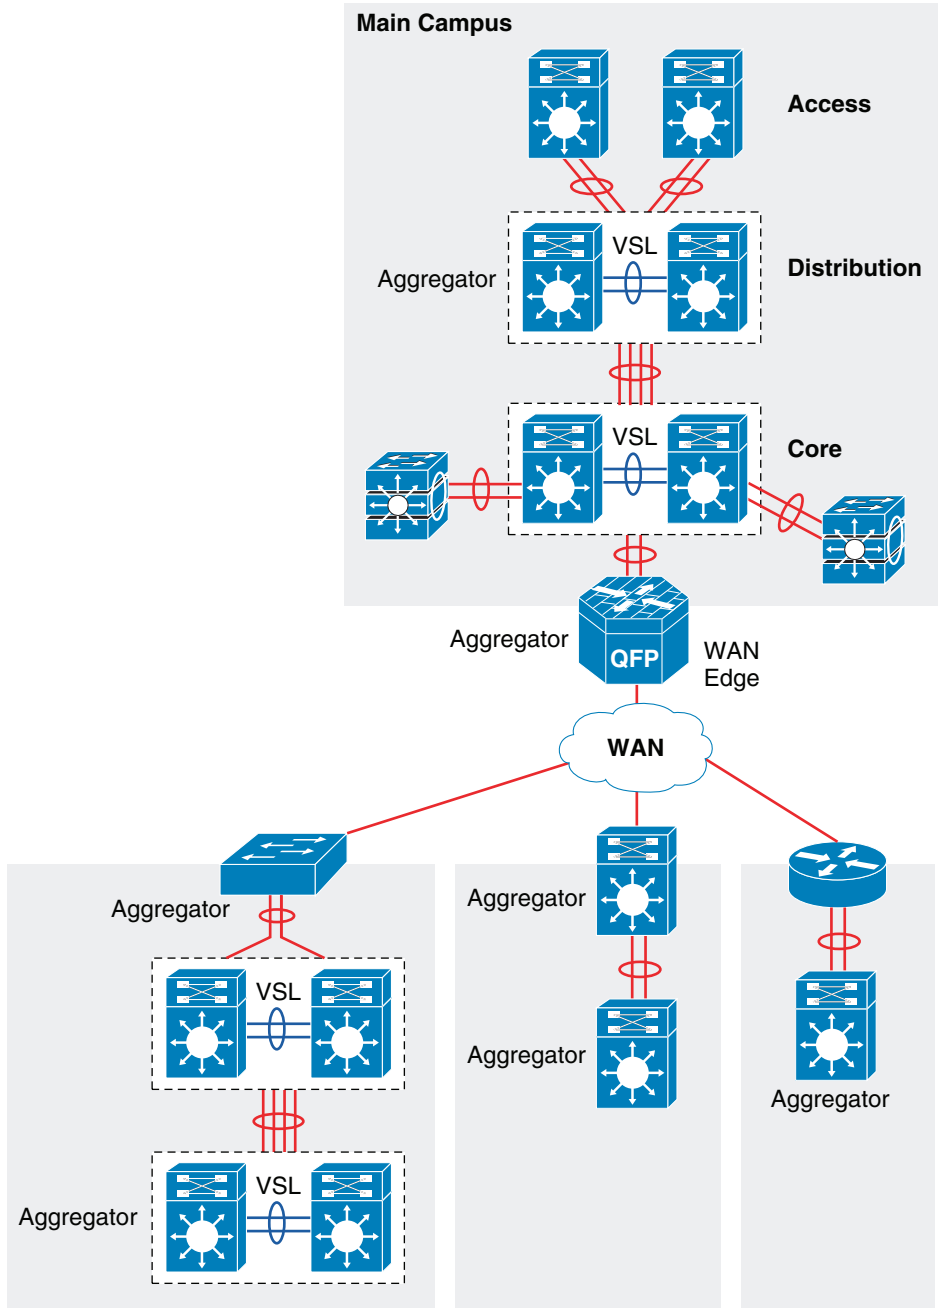
cr23-VSS-Core(config)#key chain eigrp-key
cr23-VSS-Core(config-keychain)# key 1
cr23-VSS-Core(config-keychain-key)#key-string <password>

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 3-38](#) shows an example of the EIGRP topology for the community college LAN design.

Figure 3-38 EIGRP Route Aggregator Design



The following configuration must be applied on each EIGRP route aggregator system as depicted in Figure 3-38. 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

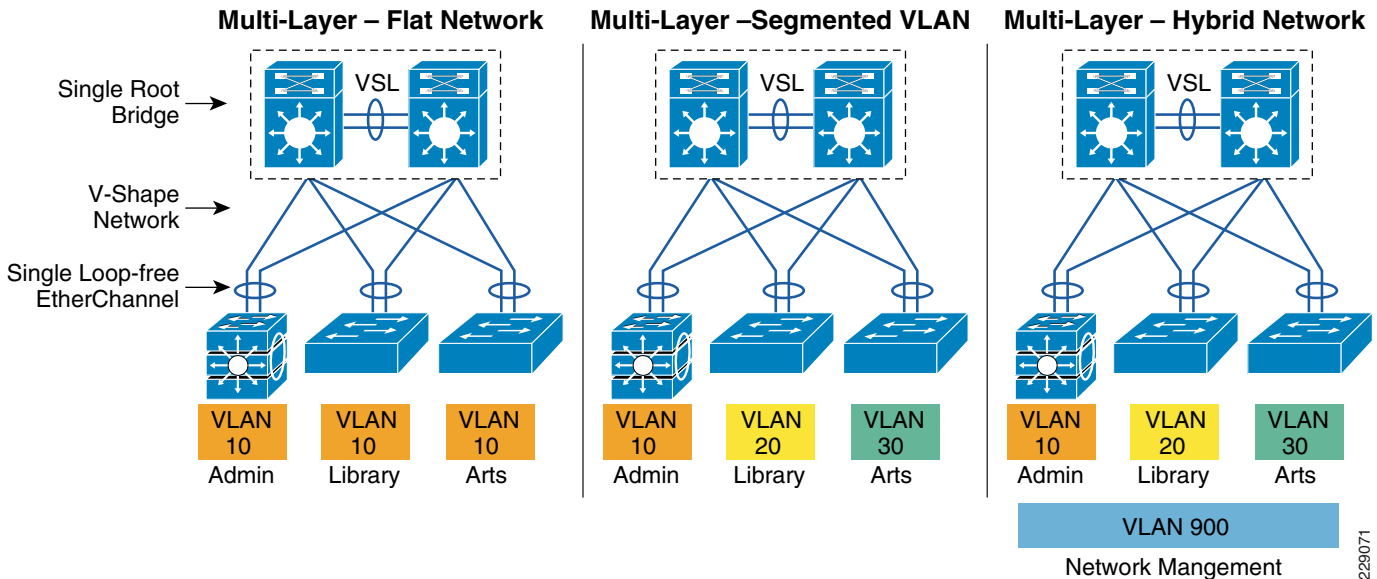
A multilayer network is a traditional, simple, and widely deployed scenario, regardless of network scale. The access layer switches in the campus network edge interface with various types of endpoints and provide intelligent Layer 1/2 services. The access layer switches interconnect to distribution switches 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 education divisions and college business function segments to build a per-department logical network. All network communication between education 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 3-39 shows the three design variations for the multilayer network.

Figure 3-39 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 be 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
Pol	on	802.1q	trunking	801

```

Port          Vlans allowed on trunk

```

```

Po1          101-110,900

Port         Vlans allowed and active in management domain
Po1          101-110,900

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 community college 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 3-40 shows an example of enabling various STP extensions on distribution and access layer switches in all campus sites.



Figure 3-40 Protecting Multilayer Network with Cisco STP Toolkit

**Note**

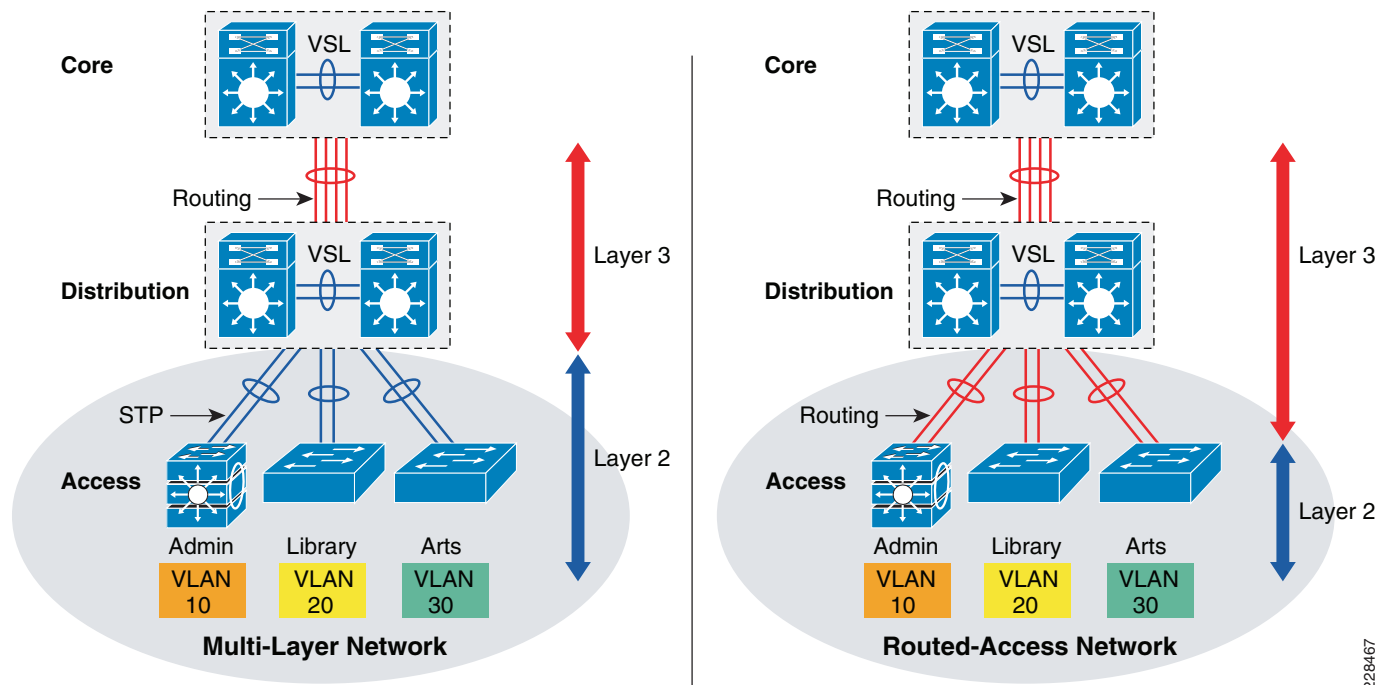
For additional STP information, see the following URL:

[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 3-41 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 3-41 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 community colleges 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.

228467

- *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 3-42](#):

- 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

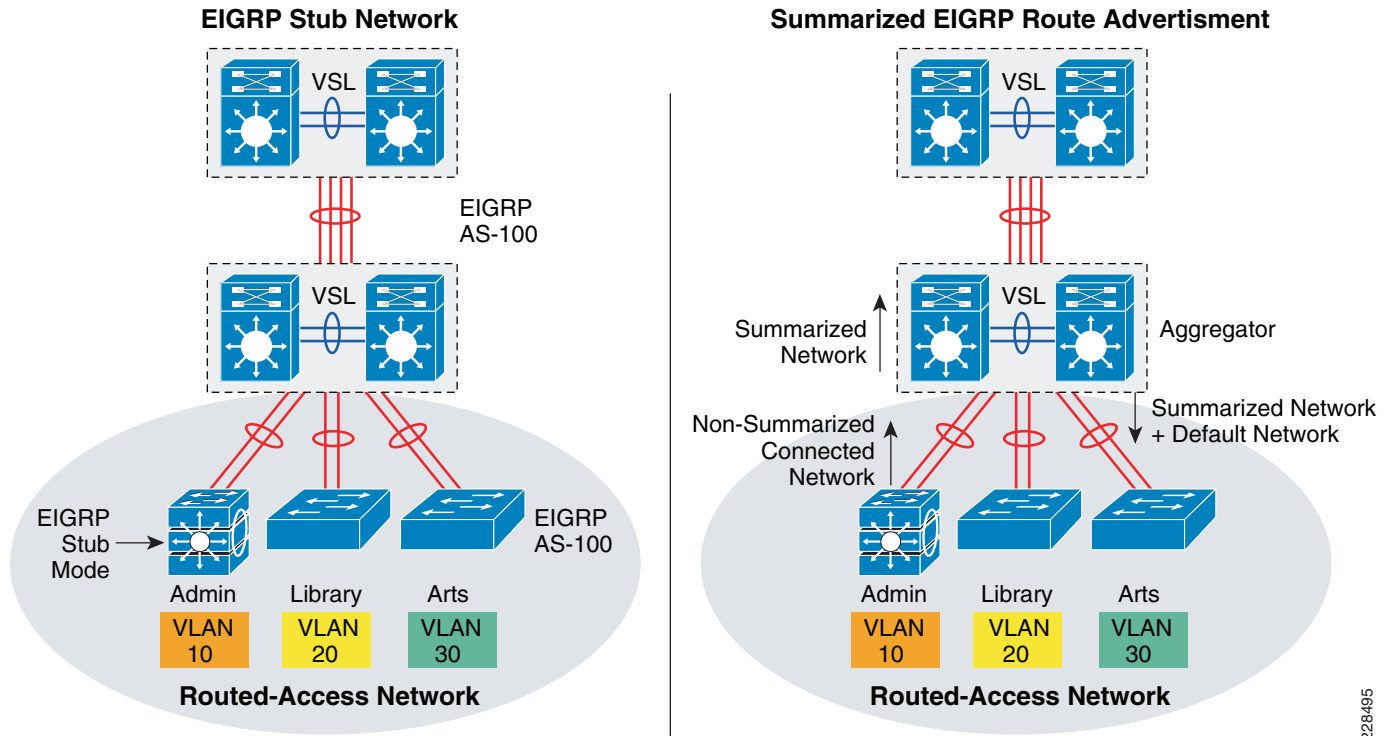
Address Family Protocol EIGRP-IPv4:(100)
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  EIGRP NSF-aware route hold timer is 240
  EIGRP NSF enabled
    NSF signal timer is 20s
    NSF converge timer is 120s
    Time since last restart 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 3-42 Designing and Optimizing EIGRP Network Boundary for the Access Layer



228-495

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 community college 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 3-5](#).

**Table 3-5** 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 3-5 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, community college 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 community college 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 community college 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.

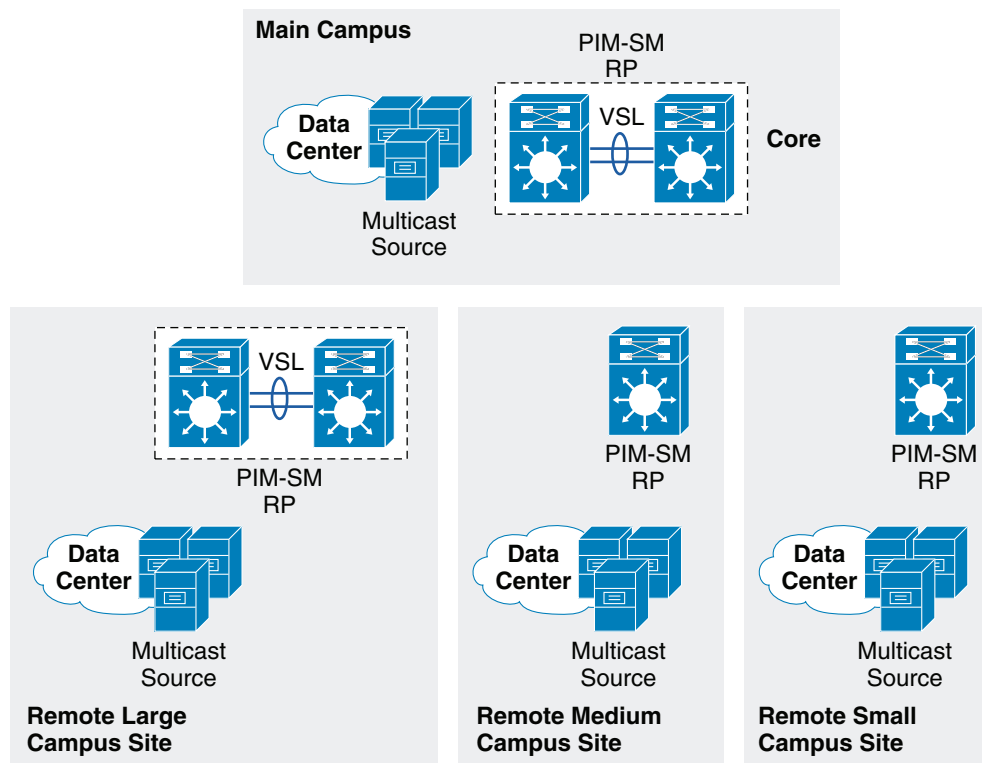
## Designing PIM Rendezvous Point

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

### PIM-SM RP Placement

It is assumed that each community college site has a wide range of local multicast sources in the data center for distributed community college IT-managed media and student 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 Community College 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 3-43](#).

**Figure 3-43** Distributed PIM-SM RP Placement



### PIM-SM RP Mode

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

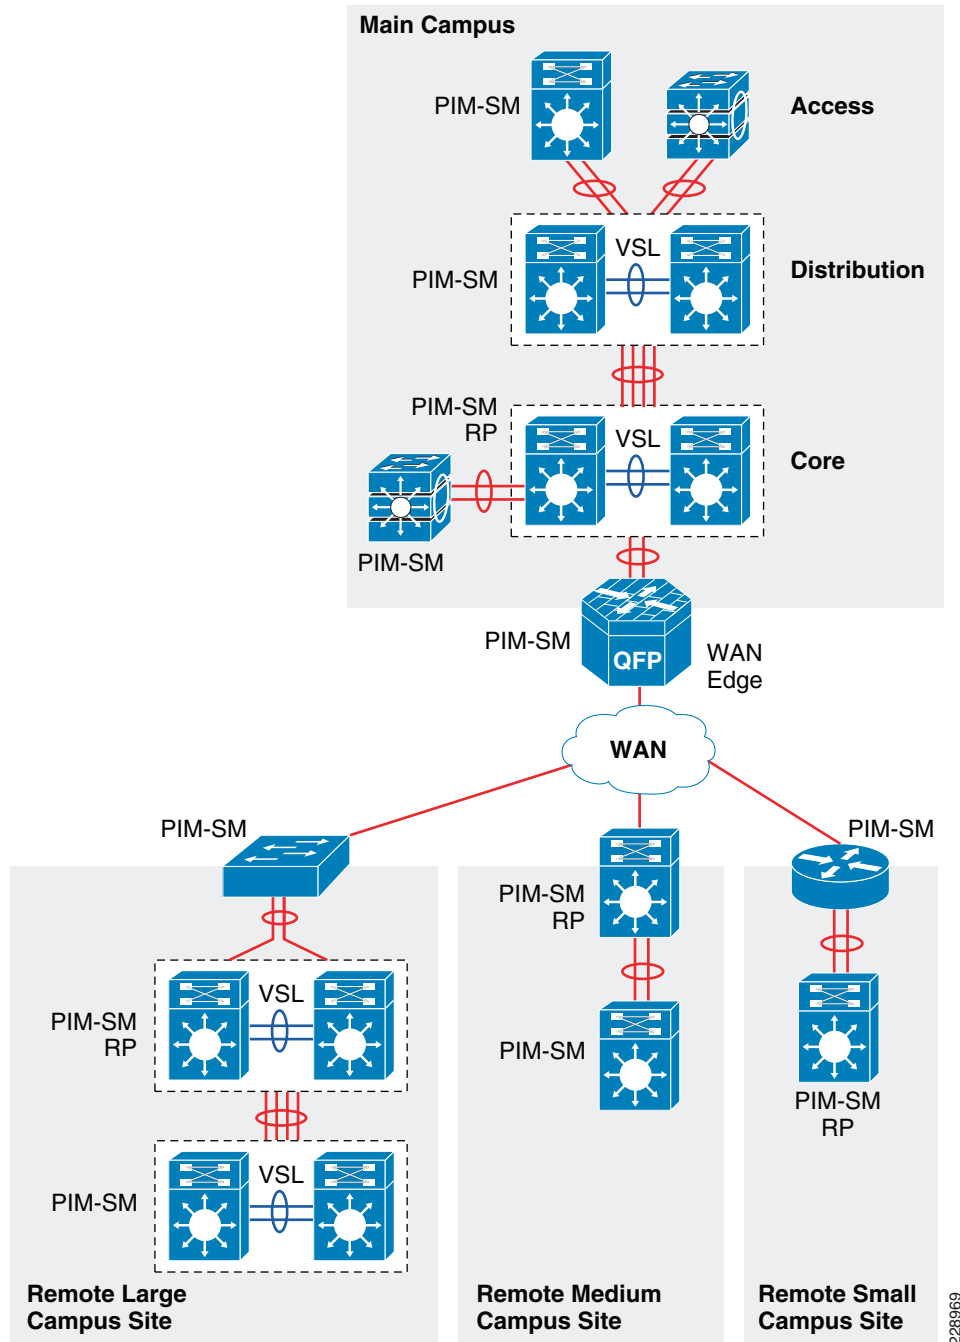
- *Static*—In this mode, RP must be statically identified and configured on each PIM router in the network. RP load balancing and redundancy can be achieved using anycast RP.
- *Auto-RP*—This mode is a dynamic method for discovering and announcing the RP in the network. Auto-RP implementation is beneficial 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 community college LAN multicast network using the static PIM-SM mode configuration. See [Figure 3-44](#).



Figure 3-44 PIM-SM Network Design in Community College 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 college 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   QueryDR  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-channell1, 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-channell1, 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 community college 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

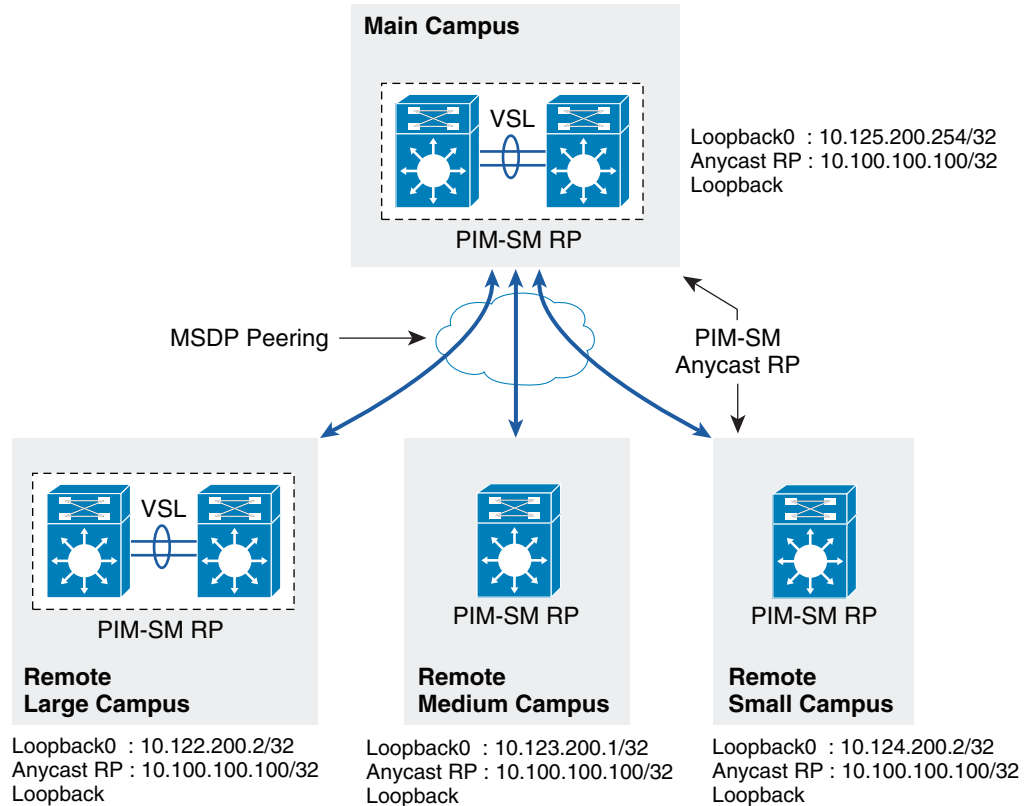
MSDP allows PIM RPs to share information about the active sources. PIM-SM RPs discover local receivers through PIM join messages, while the multicast source can be in a local or remote network domain. MSDP allows each multicast domain to maintain an independent RP that does not rely on other multicast 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 community college 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 3-45](#) shows an example of a community college LAN multicast network design.

Figure 3-45 Community College LAN Inter-Site Multicast Network Design



## 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 community college 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 college 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 education community enters the digital era of delivering education, as well as the increased campus network and asset security requirements. Integrating video applications in the community college 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 educational settings. Besides internal and college 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, community college 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 community college 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.

## Community College 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 community college 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 college 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 community college 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. Community college 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 3-46](#).

Figure 3-46 Community College LAN 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.

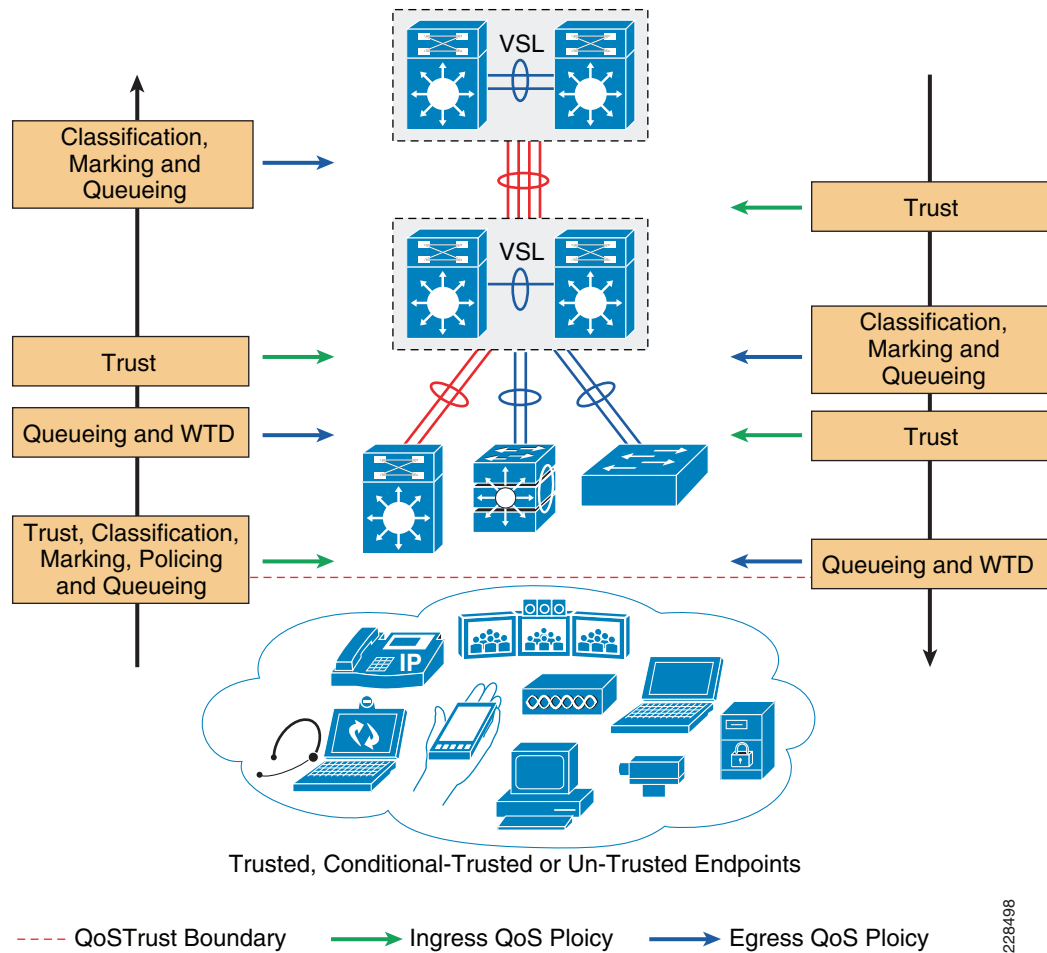
## Designing Community College LAN QoS Trust Boundary and Policies

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 college 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 3-47](#) shows the QoS trust model setting that sets the QoS implementation guidelines in community college campus networks.

Figure 3-47 Campus QoS Trust and Policies



228498

## Community College 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 college campus. On the other hand, if community college 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 and education community 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 community college voice, video, and data applications marking recommendations are standards-based, as previously discussed, community colleges 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 college 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 community college 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 community college 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 community college 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

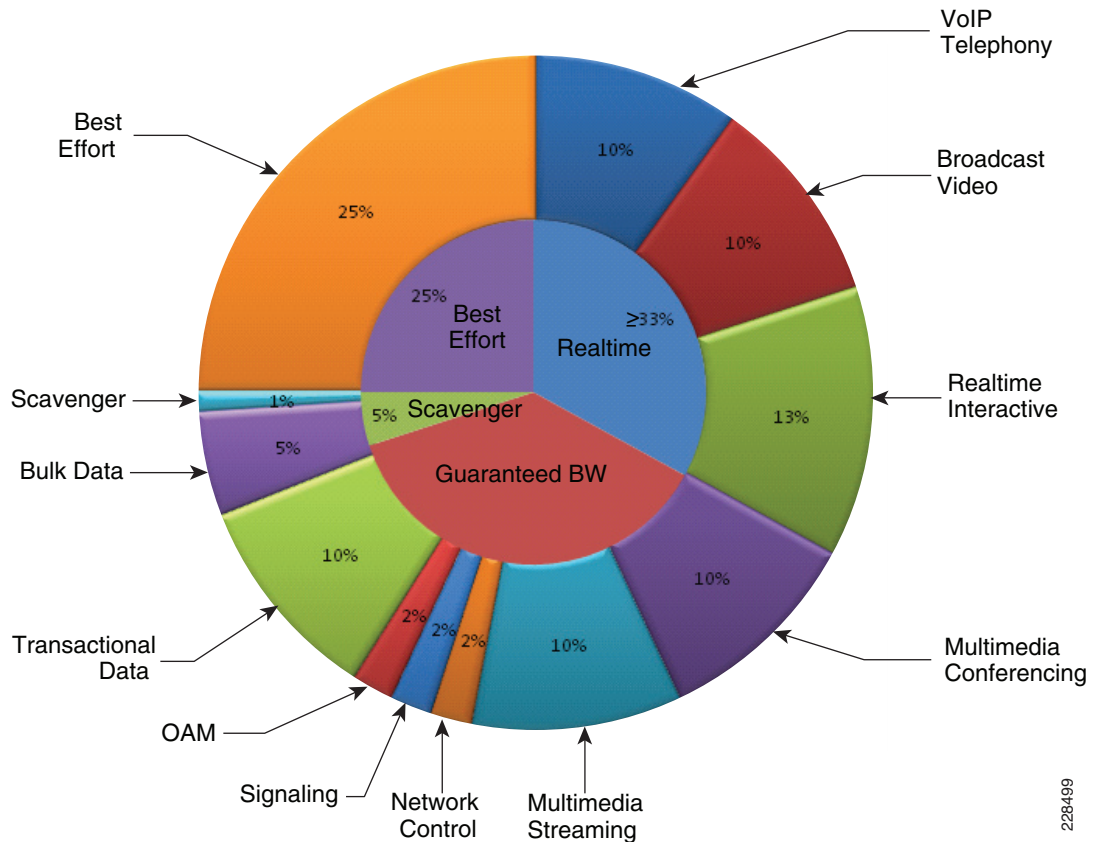
The best effort class is the default class for all traffic that has not been explicitly assigned to another application-class queue. Only if an application has been selected for preferential/differential treatment is it removed from the default class. Because most community colleges may have several 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 3-48, 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 3-48 Compatible 4-Class and 12-Class Queuing Models



228499



## Deploying QoS in College 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 community college 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:

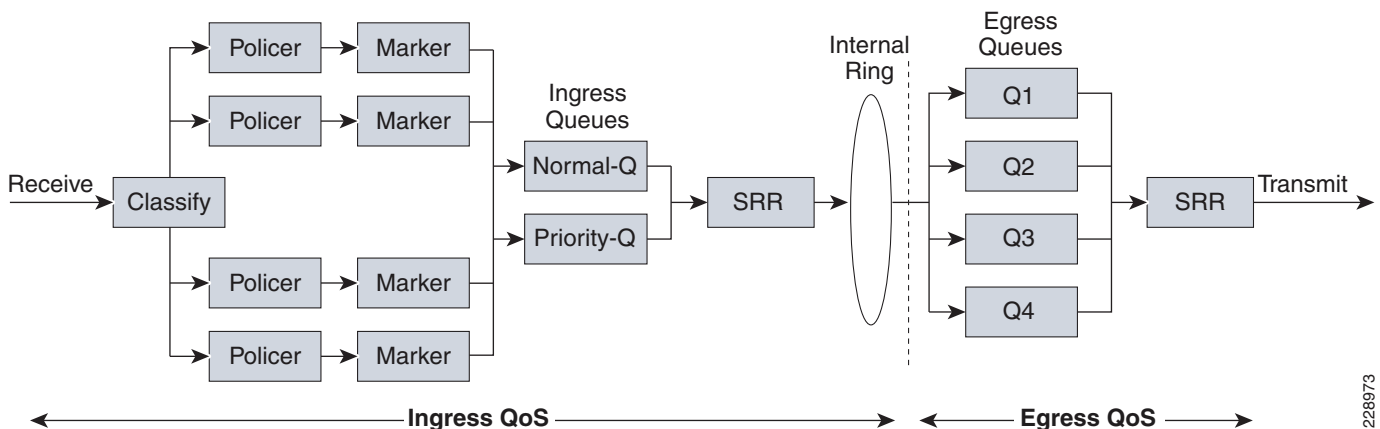
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND/QoS-SRND-Book.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html)

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-E, and 3750-E 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 3-49](#).

**Figure 3-49** QoS Implementation in Cisco Catalyst Switches



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 within this product family can police to a minimum rate of 8 kbps.
- Only the Cisco Catalyst 3560-E and 3750-E support IPv6 QoS.
- Only the Cisco Catalyst 3560-E and 3750-E support policing on 10-Gigabit Ethernet interfaces.
- Only the Cisco Catalyst 3560-E and 3750-E support SRR shaping weights on 10-Gigabit Ethernet interfaces.

## QoS in Cisco Modular Switches

The Cisco Catalyst 4500-E and 6500-E are high-density, resilient switches for large scale networks. The community college 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.

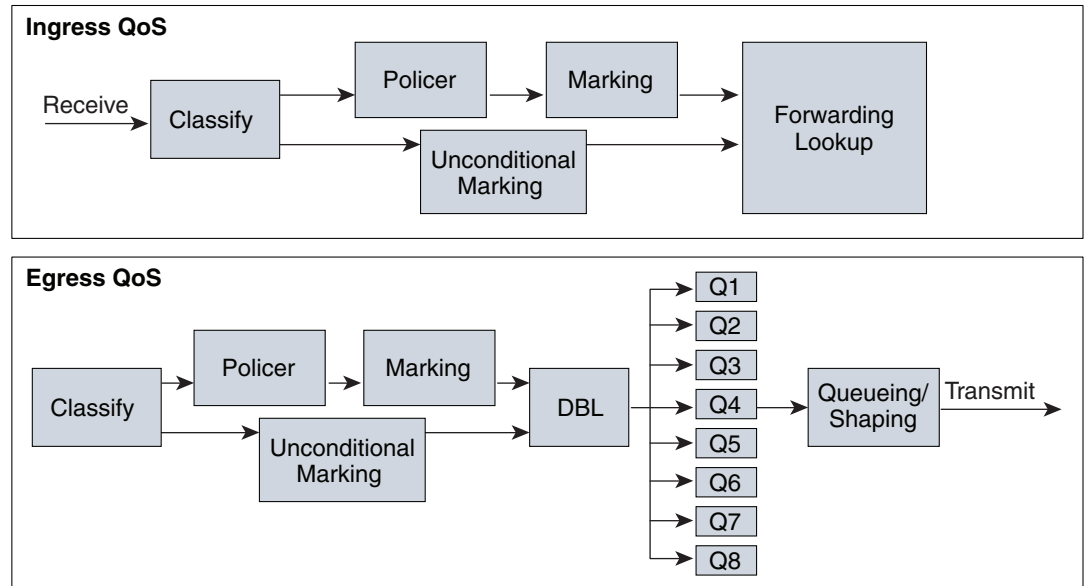
### Catalyst 4500-E QoS

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 3-50](#)) 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 3-50 Catalyst 4500 - Supervisor 6-E and 6L-E QoS Architecture**

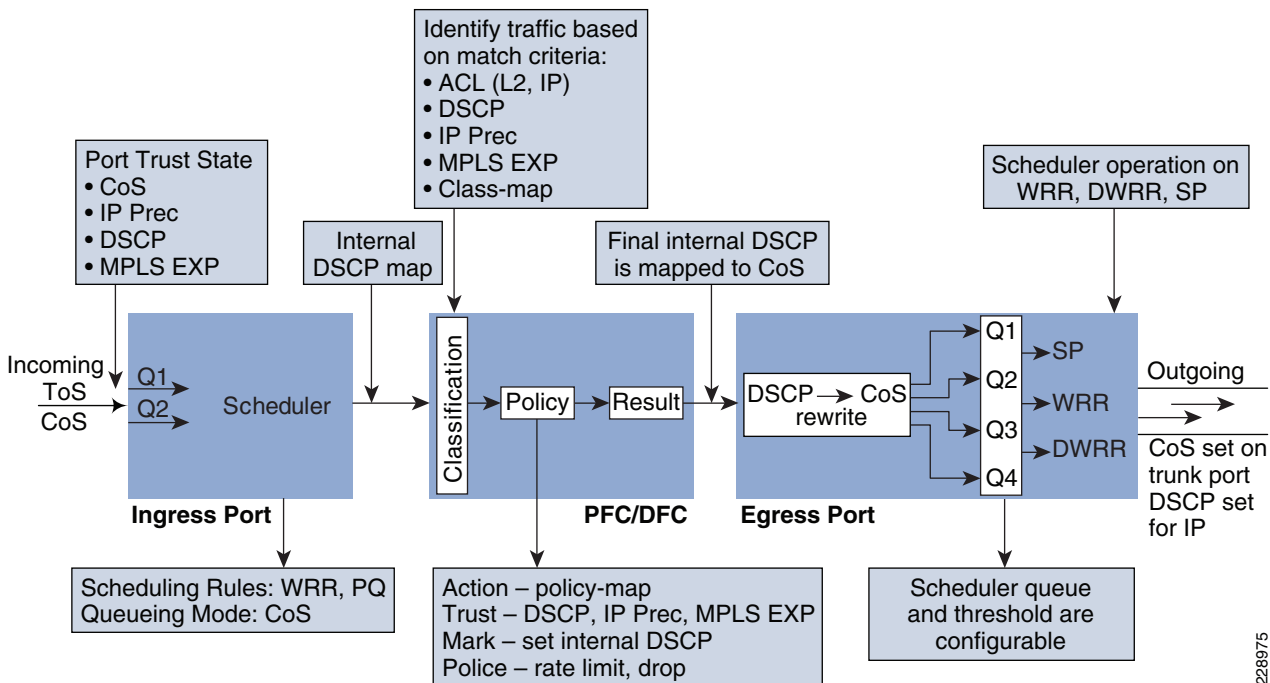


### Catalyst 6500-E QoS

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 college 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 3-51](#) provides internal PFC based QoS architecture.

Figure 3-51 Cisco Catalyst 6500-E PFC QoS Architecture



## Deploying Access-Layer QoS

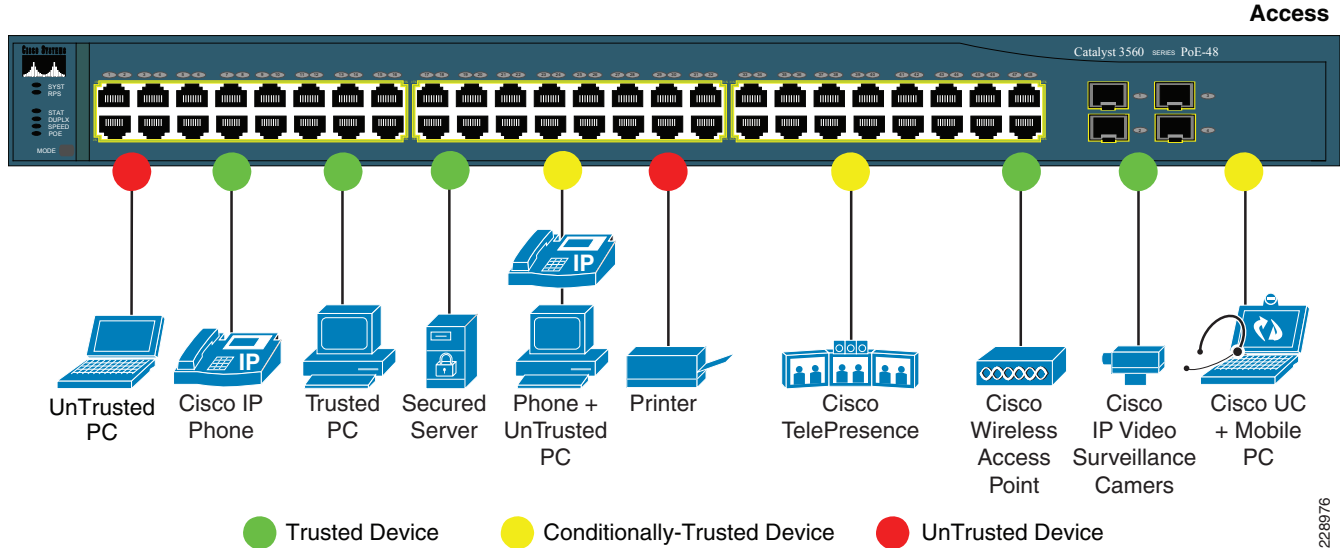
The campus access switches provide the entry point to the network for various types of end devices managed by community college IT department or student's personal devices (i.e., laptop etc.). The access switch must decide whether to accept the QoS markings from each endpoint, or whether to change them. 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 3-52 illustrates several types of devices in the network edge.

Figure 3-52 Campus LAN QoS Trust Boundary



Community college 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, student-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 college campus network LAN network.

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

```
cr24-2960-LB(config)#mls qos
cr24-2960-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 3-93 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 college 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 3-53](#). Voice traffic is primarily sourced from Cisco IP telephony devices residing in the voice VLAN (VVLAN). These are trusted devices, or conditionally trusted (if users also attach 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 3-53](#).



Figure 3-53 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 3-6](#) provides best practice policing guidelines for different classes to be implemented for trusted and conditional-trusted endpoints at the network edge.

**Table 3-6 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 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 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 3-58](#).

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 VLAN-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

```

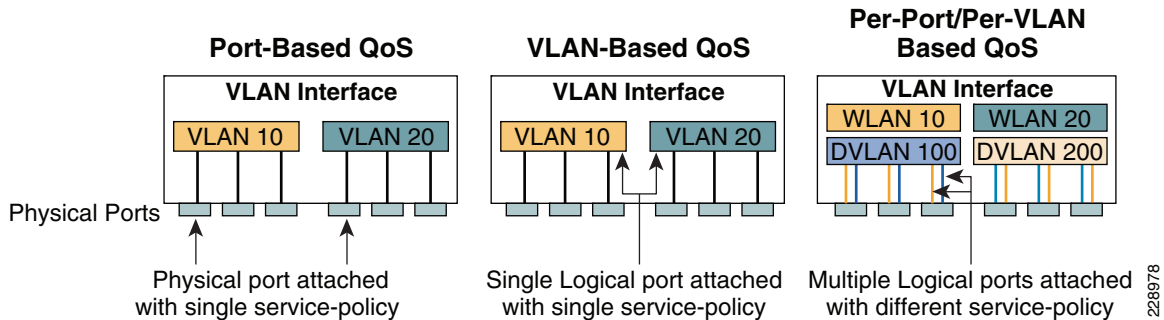
### Applying Ingress Policies

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 bas 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 3-54](#).

**Figure 3-54** 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

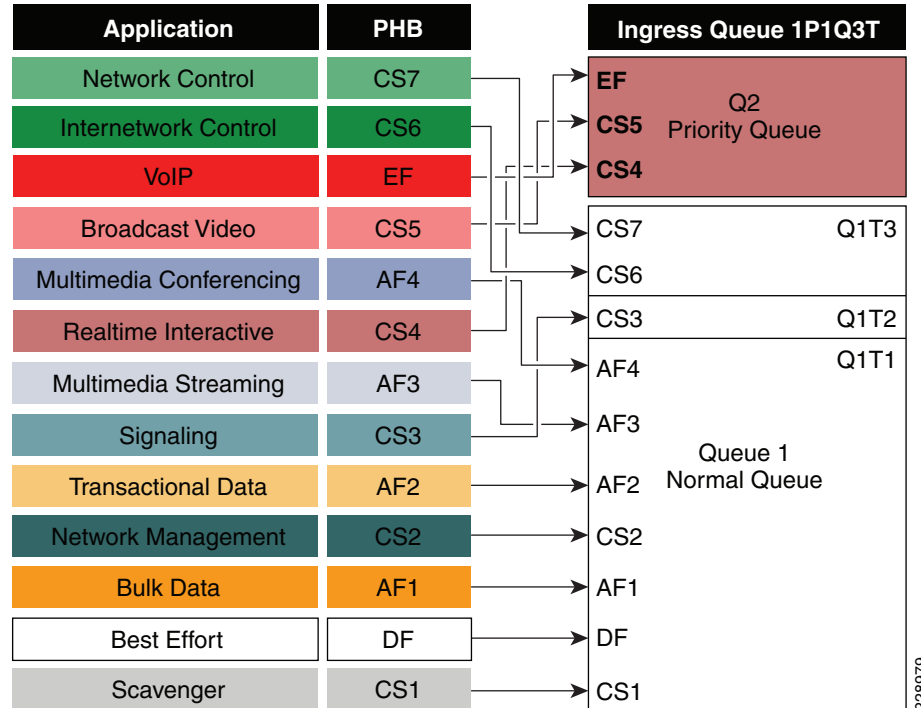
Fixed configuration Cisco Catalyst switches (29xx and 3xxx) not only offer differentiated services on the network ports, but also internally on the switching fabric. After enabling QoS and attaching inbound policies on the physical ports, all the packets that meet the specified policy are forwarded to the switching fabric for egress switching. The aggregate bandwidth from all edge ports may exceed the switching fabric bandwidth and cause internal congestion.

These 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-E 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 3-55 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 3-55 Catalyst 29xx and 3xxx Ingress Queuing Model



- Catalyst 29xx 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 02-01 02-03 02-01 01-03 01-01
5 :   01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-03 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 3-50](#).

## 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 29xx 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 community college 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 3-56](#) illustrates the egress bandwidth allocation best practices design for different classes.



**Figure 3-56 Class-of-Service Egress Bandwidth Allocations**

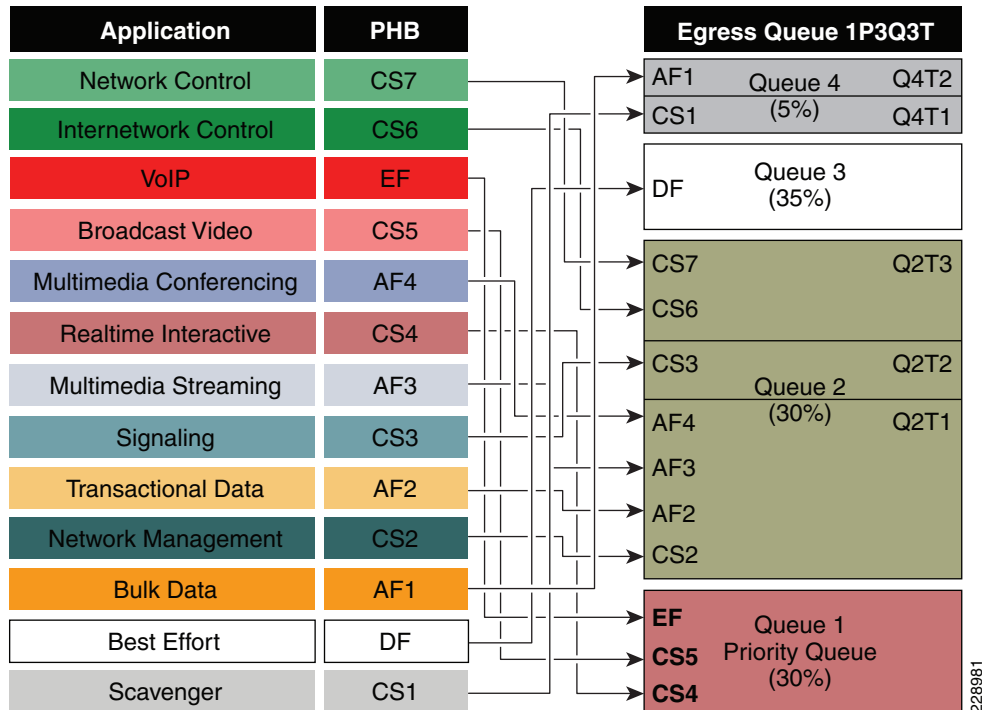


Given these minimum queuing requirements and bandwidth allocation recommendations, the following application classes can be mapped to the respective queues:

- *Realtime Queue*—Voice, broadcast video, and realtime interactive may be mapped to the realtime queue (per RFC 4594).
- *Guaranteed Queue*—Network/internet 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 3-57](#).

Figure 3-57 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 3-57. 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 29xx 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 3-58](#).

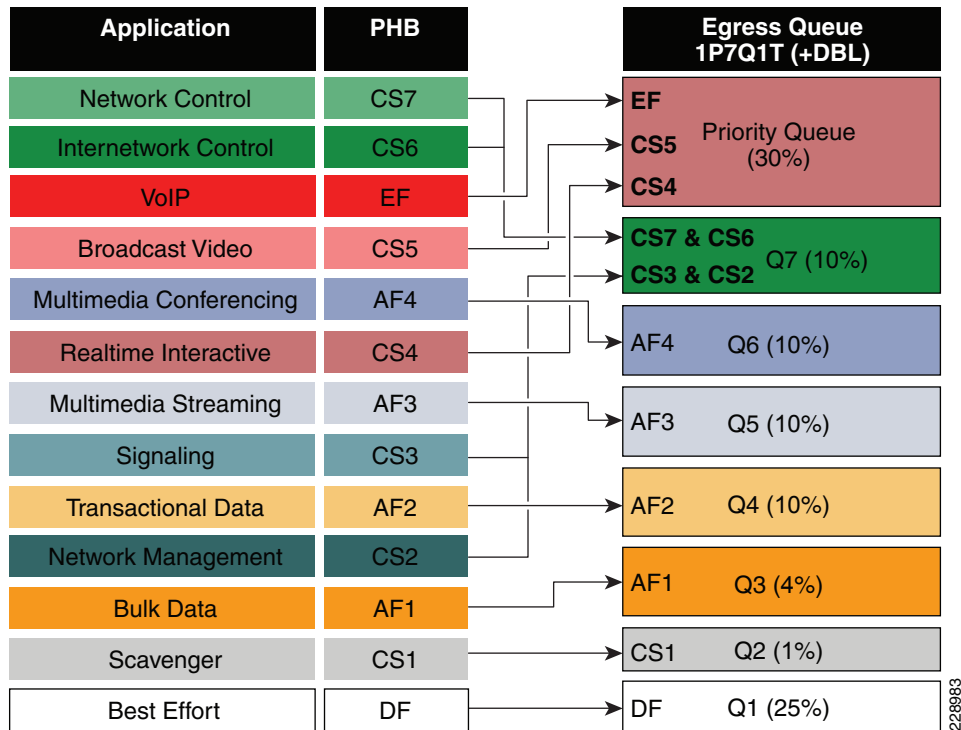
**Figure 3-58 8 Class-of-Service Egress Bandwidth Allocations**



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 3-59 provides the new recommended bandwidth allocation.

**Figure 3-59** 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 3-59, 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 (multi-layer 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 3-7 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 3-8 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 college campus and remote campus are managed and maintained by the college 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 community college 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 college 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 college 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 college 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 community college 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 college 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-E 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 college campus LAN network.

**Distribution-Layer Catalyst 3750-E and 6500-E**

- 3750-E 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-E (Multilayer or Routed Access)**

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

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

```
cr36-3750s-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-E 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-E Ingress Queuing**

The ingress queuing function in the distribution-layer Catalyst 3750-E 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 3-96](#) 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 college campus network.

Table 3-9 summarizes recommended linecards consideration by listing and oversubscription ratios and whether the ingress queuing models are CoS or DSCP-based.

**Table 3-9 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

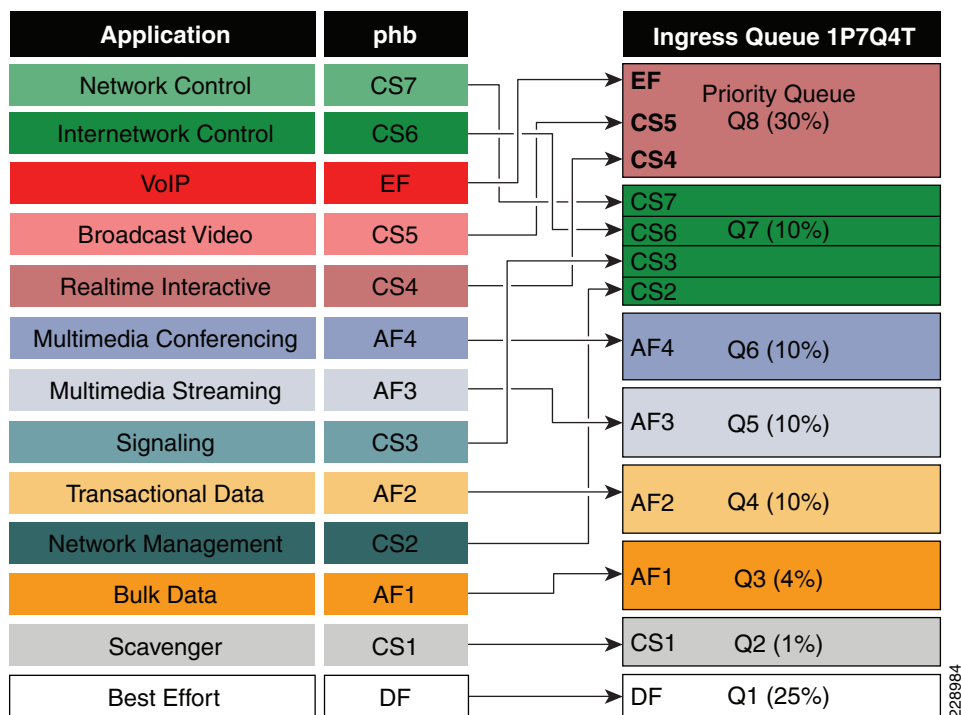
also at 40 Gbps). In Oversubscription Mode (the default mode), all ports are operational and the maximum oversubscription ratio is 4:1. Therefore, it is recommended to enable 1P7Q2T DSCP-based ingress queuing on this linecard in Oversubscription Mode.

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 college 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 3-60 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 3-60 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 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: 0

queue                dropped [dscp-map]
-----

```

```

1          0 [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 47 ]
2          0 [0 ]
3          0 [14 12 10 ]
4          0 [22 20 18 ]
5          0 [30 35 37 28 26 ]
6          0 [38 49 50 51 52 53 54 55 57 58 59 60 61 62 63 36 34 ]
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-E and 4500-E

The configuration and implementation guideline for egress QoS on Catalyst 3750-E 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 3-98 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 DFC3C/DFC3CXL is supported.

[Table 3-10](#) describes the deployment guidelines for the Catalyst 6500-E Series linecard module in the college 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 3-10 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				

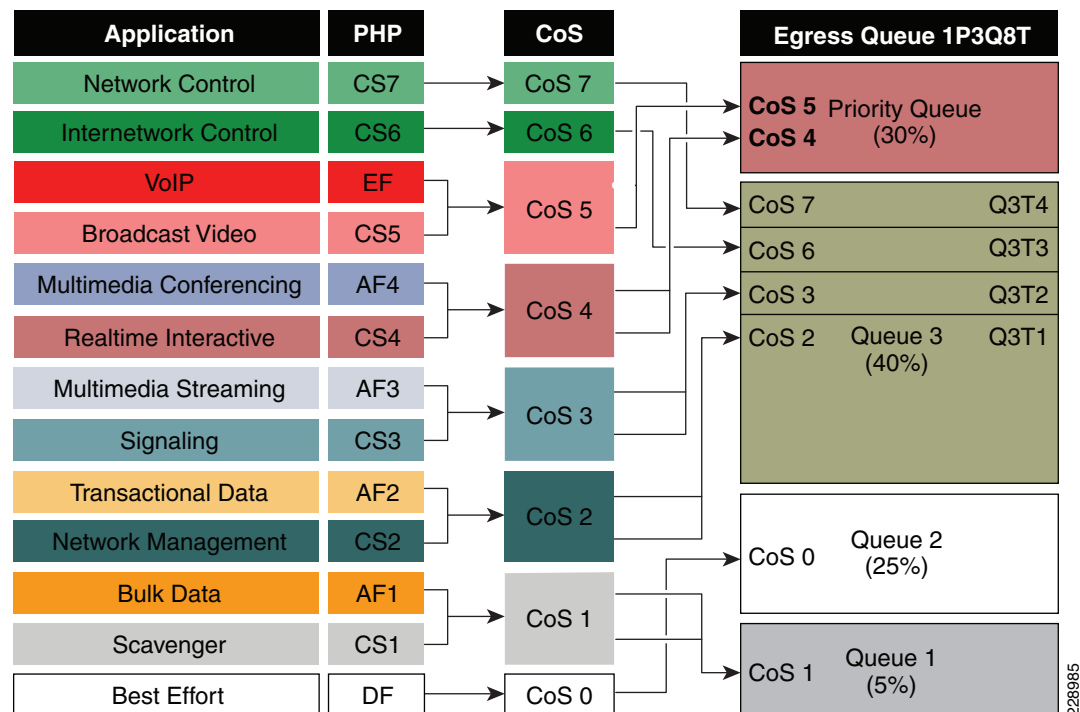
**Table 3-10 Catalyst 6500-E Switch Module Egress Queuing Architecture (continued)**

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> 91 MB <sup>2</sup>	90 MB <sup>1</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 3-61 illustrates 1P3Q8T egress queuing model to be applied on Catalyst 6500-E – WS-6724-SF module.

**Figure 3-61 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 = 1p3q8t]:
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]

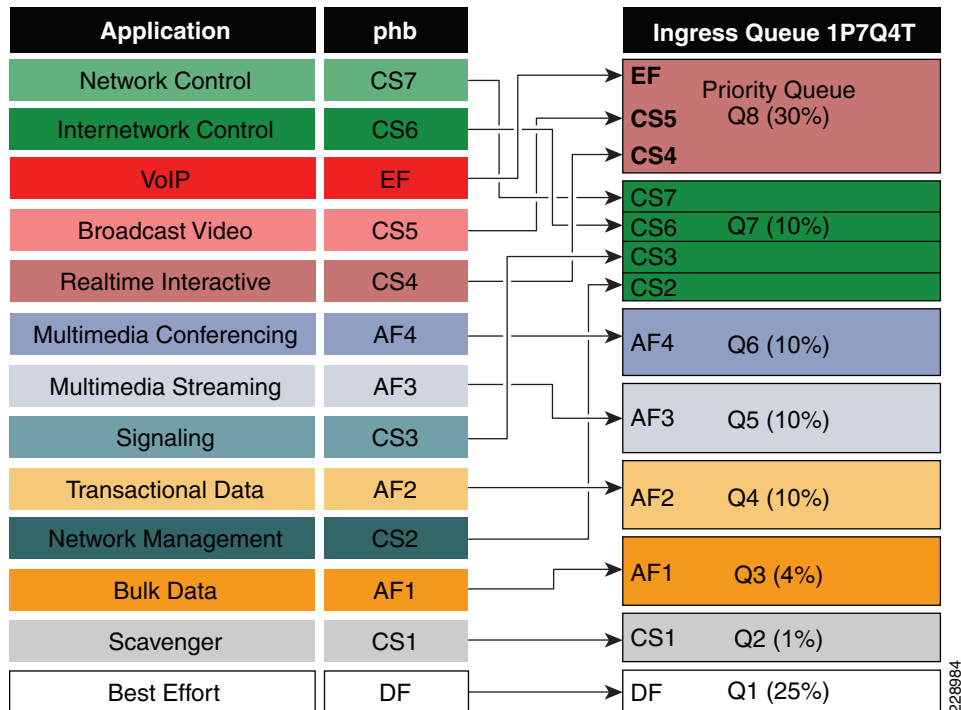
WRED disabled queues:
queue thresh cos-map
-----
 1   1       1
 1   2
 1   3
 1   4
 1   5
 1   6
 1   7
 1   8
 2   1       0
 2   2
 2   3
 2   4
 2   5
 2   6
 2   7
 2   8
 3   1       2
 3   2       3
 3   3       6
 3   4       7
 3   5
 3   6
 3   7
 3   8
 4   1       4 5
...

```

### 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 3-62 illustrates DSCP-based 1P7Q4T egress queuing model.

Figure 3-62 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 community college 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.

## Community College LAN Design High-Availability Framework

Independent of the business function, the network architect 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 community college 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 community college 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 inthe 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 3-63](#) 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 3-63 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 3-64.)

Figure 3-64 Community College LAN Design 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

### 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 community college 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 college 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 a 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 community college 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*—The power supply in fixed configuration switches can be internal or use Cisco RPS 2300 external power supplies. 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 complete fault transparency and constant network availability.

## Redundant Control Plane

Device or node resiliency in modular Cisco Catalyst 6500/4500 platforms and Cisco StackWise provides a 1+1 redundancy option with enterprise-class high availability and deterministic network recovery time. The following sub-sections 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 college 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 a 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.

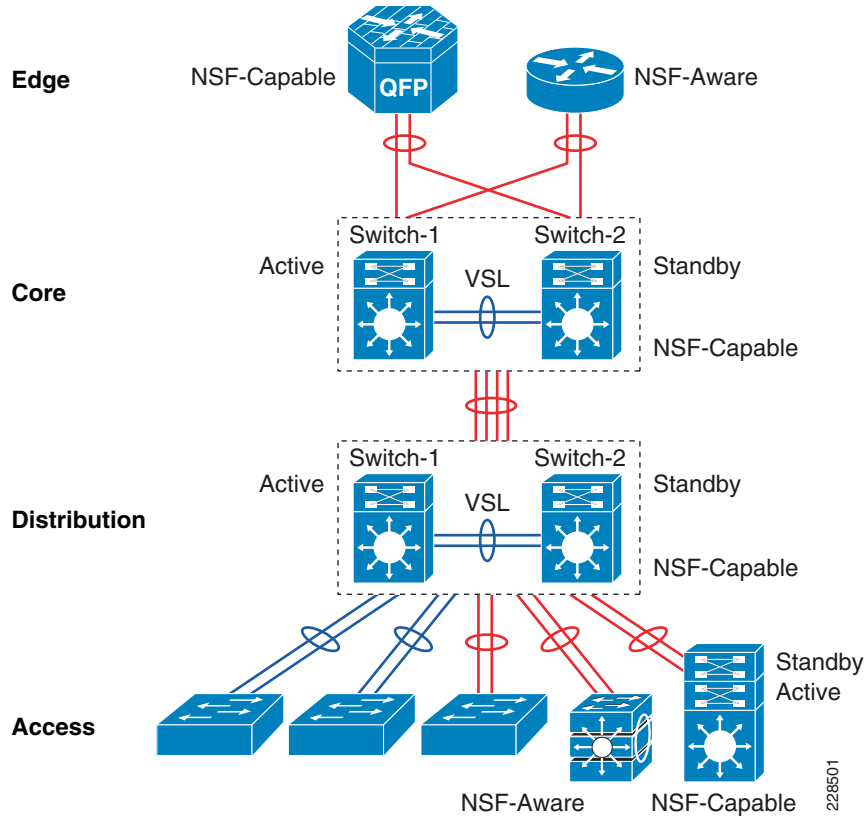
Bundling SSO with NSF capability and the awareness function allows the network to operate without errors during a primary supervisor module failure. Users of realtime applications such as VoIP do not hang up the phone, and IP video surveillance cameras do not freeze.

## 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 3-65](#).)



**Figure 3-65 Community College LAN Design NSF/SSO Capable and Aware Systems**

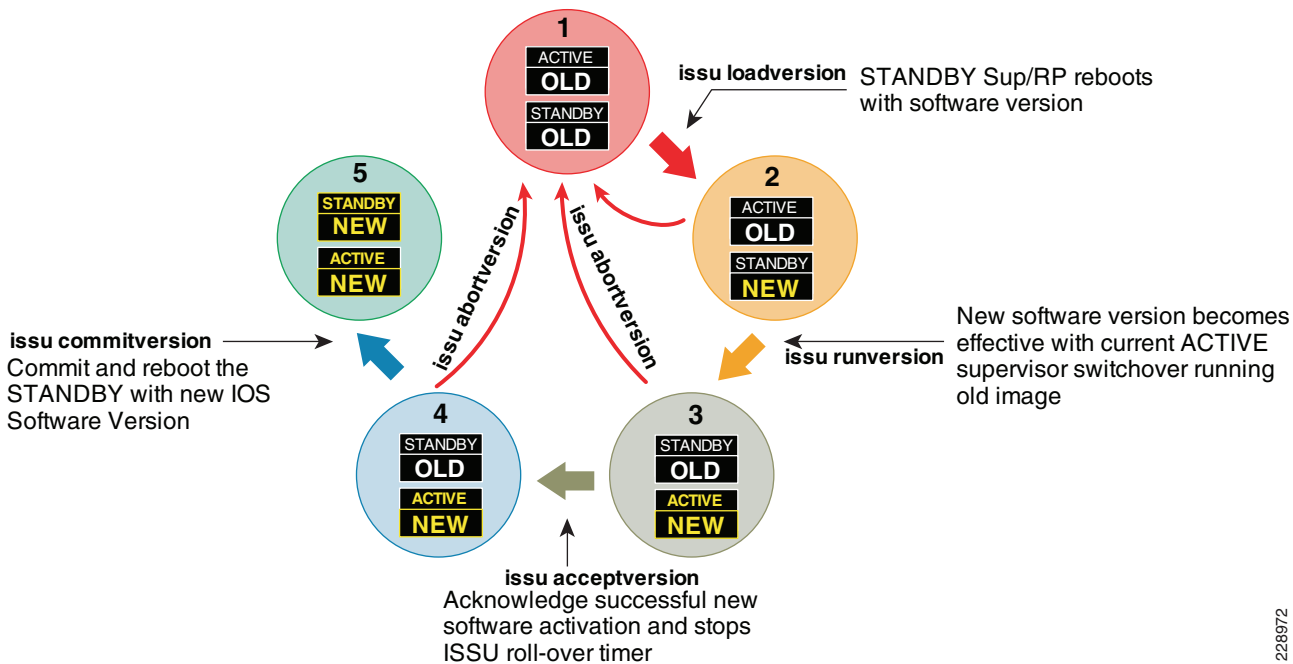


## Operational Resiliency Overview

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 community college LAN network is part of the backbone of the college 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 3-66](#) demonstrates platform-independent Cisco IOS software upgrade flow process using ISSU technology.

Figure 3-66 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 3-11](#).

**Table 3-11 Community College Network High Availability Strategy**

Platform	Role	Network Resiliency	Device Resiliency	Operational Efficiency
Catalyst 2960	Access	EtherChannel <sup>1</sup> UDLD Dampening	Red. Power Supplies RPS 2300 NSF Aware	None. Standalone systems
Catalyst 3560-E				
Catalyst 3750-E				
Catalyst 3750ME	WAN Edge			
Catalyst 3750-E StackWise	Access		Red. Power Supplies RPS 2300 NSF Capable & Aware	Stackwise Plus
	Distribution			
Catalyst 4500-E	Access		Red. Power Supplies <sup>2</sup> Red. Linecard modules <sup>2</sup> Red. Supervisor modules <sup>3</sup>	ISSU
	Distribution			
	Core			
Catalyst 6500-E	Distribution		SSO/NSF Capable & Aware <sup>2</sup>	VSS eFSU
	Core			

**Table 3-11 Community College 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>	
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 community college 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 community college 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 3-41.

### 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 community college 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.

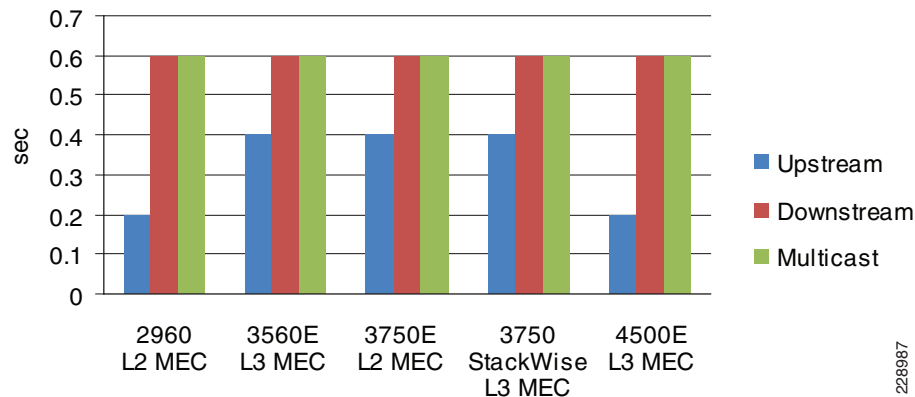
During individual member-link failures, the Layer 2 and Layer 3 protocols dynamically adjusts the metrics of the aggregated port-channel interfaces. Spanning-Tree updates the port-cost and Layer 3 routing protocols like EIGRP updates the composite metric or OSPF may change the interface cost. In

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 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-E 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 3-67](#) summarizes several types of faults induced in large scale Cisco lab during developing this validated design guide.


**Figure 3-67 Catalyst 6500-E VSS Inter-Chassis MEC Link Recovery Analysis**



The Community College 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 3-133](#) for intra-chassis recovery analysis with the same network faults tested in inter-chassis scenarios.

### Catalyst 4507R-E EtherChannel Link Recovery Analysis

In the community college 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 3-68](#) confirm the deterministic and consistent network recovery during individual Layer 2/3 EtherChannel member-link failures.

**Figure 3-68 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	FDO1328R0E2	1	Gi1/0/49	Bidirectional
Gi2/2/3	FDO1328R0E2	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 community college 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-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 community college 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 college 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 college 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. For some Catalyst platforms like the 3750-E/3560-E and 2960, a Cisco RPS 2300 can be used for external power redundancy.

The next-generation borderless network ready Cisco Catalyst 3560-X and 3750-X Series switches are designed to increase device resiliency with dual redundant power supplies and fans. These latest released Catalyst switching platforms were not shipping during the community college solution development process; therefore, they are not covered in this document. See following URL for further product details:

[http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps6406/data\\_sheet\\_c78-584733\\_ps10744\\_Products\\_Data\\_Sheet.html](http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps6406/data_sheet_c78-584733_ps10744_Products_Data_Sheet.html)

**Figure 3-69 Power Supply Redundancy Design**



The following configuration examples provide guidelines to deploy in-chassis and external power redundancy in the Catalyst switching platforms.

**Catalyst 29xx and 3xxx (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:



```
cr36-3750s-LB#power rps <switch id> name CiscoRPS
cr36-3750s-LB#power rps <switch id> port <rps port id> active
```

```
cr36-3750s-LB#show env rps
SW      Status  RPS Name      RPS Serial#    RPS Port#
-----
1       Active  CiscoRPS      FD01246SG3L    1
2       Active  CiscoRPS      FD01246SG3L    3
3       Active  CiscoRPS      FD01246SG3L    5
```

```
RPS Name: CiscoRPS
State: Active?
PID: PWR-RPS2300
Serial#: FD01246SG3L
Fan: Good
Temperature: Green
```

```
RPS Power Supply A : Present
PID                : C3K-PWR-1150WAC
Serial#            : DTM124000XX
System Power       : Good
PoE Power          : Good
Watts              : 300/800 (System/PoE)
```

**Redundant RPS  
Power Supply**

```
RPS Power Supply B : Present
PID                : C3K-PWR-1150WAC
Serial#            : DTM124000WW
System Power       : Good
PoE Power          : Good
Watts              : 300/800 (System/PoE)
```

DCOut	State	Connected	Priority	BackingUp	WillBackup	Portname	SW#
1	Active	Yes	6	No	Yes	cr36-3750s-LB	1
2	Active	Yes	6	No	Yes	<>	<>
3	Active	Yes	6	No	Yes	cr36-3750s_LB	2
4	Active	Yes	6	No	Yes	<>	<>
5	Active	Yes	6	No	Yes	cr36-3750s_LB	3
6	Active	Yes	6	No	Yes	<>	<>

### 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 college 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-E 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 3-70](#) 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 3-70 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 re-route which may cause network congestion if there is not sufficient bandwidth to accommodate the re-routed traffic. [Figure 3-71](#) demonstrates inter-chassis reroute (without linecard redundancy) and intra-chassis re-route (with linecard redundancy).

**Figure 3-71 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 college 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.

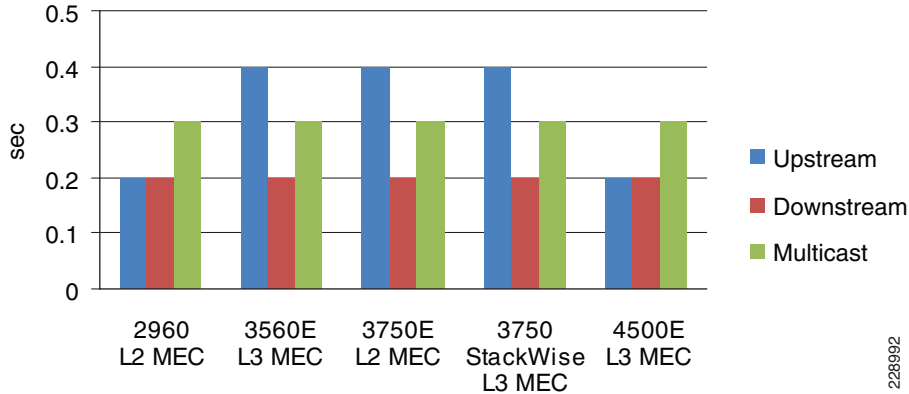
## Redundant Linecard Network Recovery Analysis

### 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 3-71](#) provides an example of asymmetric traffic-loss statistics when traffic is rerouted via remote virtual-switch node across VSL links. [Figure 3-72](#) 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 3-72 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 3-73 provides test results conducted by removing a linecard from the Catalyst 4507R-E chassis deployed in college campus network in various roles.

Figure 3-73 Catalyst 4507R-E Linecard Recovery Analysis



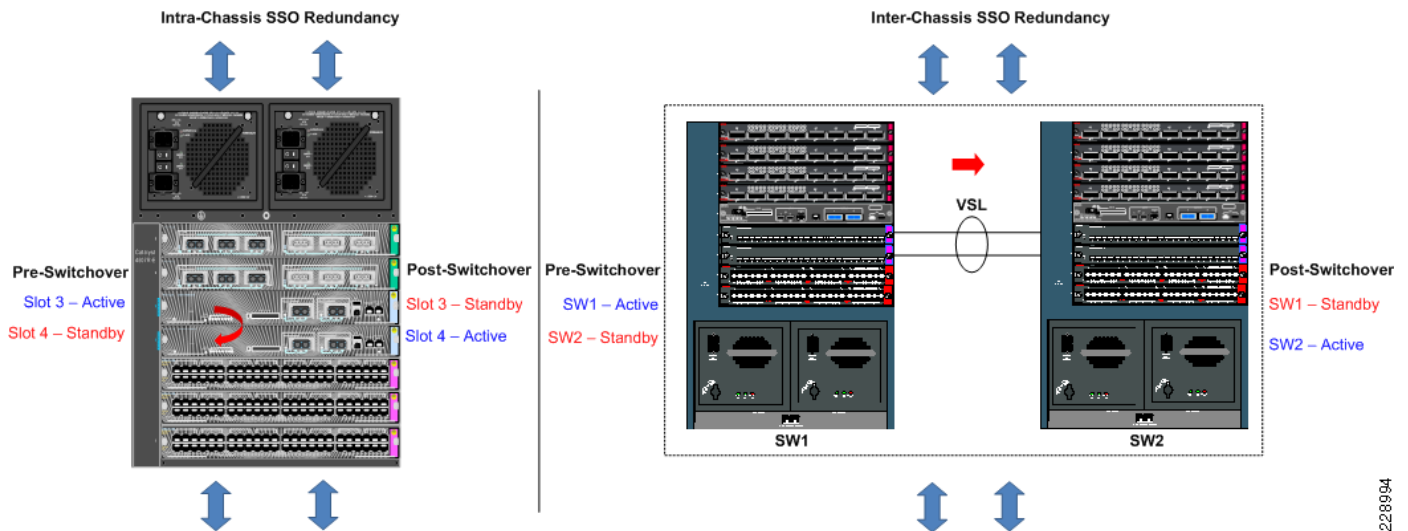
228993

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 3-74.

**Figure 3-74 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
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 college 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 3-11 describes the Layer 3 NSF-capable and aware platforms deployed in the college 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 college 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 community college 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.

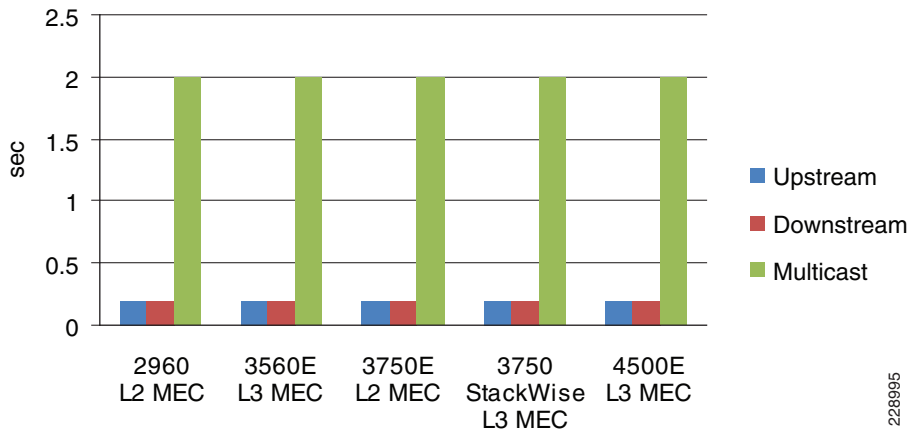
During the SSO switchover process, the Cisco Catalyst 4507R-E deployed with redundant Sup6E or Sup6L-E will retain the operational and forwarding state of the uplink ports and linecard modules in the chassis.

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 3-75 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 3-75 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 college 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 3-76](#).

**Figure 3-76 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 3-76](#).

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:

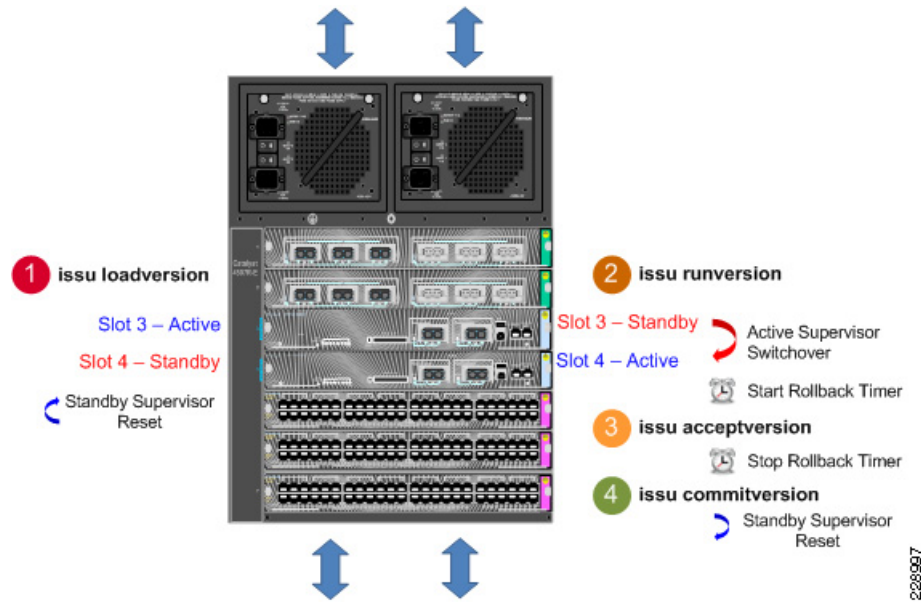
- The **port-channel load-defer** command is primarily developed for Catalyst 6500-E based standalone systems and does not have much effect when the campus upstream device type is Catalyst 6500-E deployed in VSS mode.
- There is no software restriction on turning on the feature on VSS systems. However, 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 3-77 Catalyst 4500-E ISSU Software Upgrade Process



## ISSU Software Upgrade Pre-Requisite

### 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 community college 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   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
```

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 3-78 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 community college 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)SX12a
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)SX12a
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 community college 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 community college 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 community college design.



## CHAPTER 4


# Community College WAN Design

---

## WAN Design

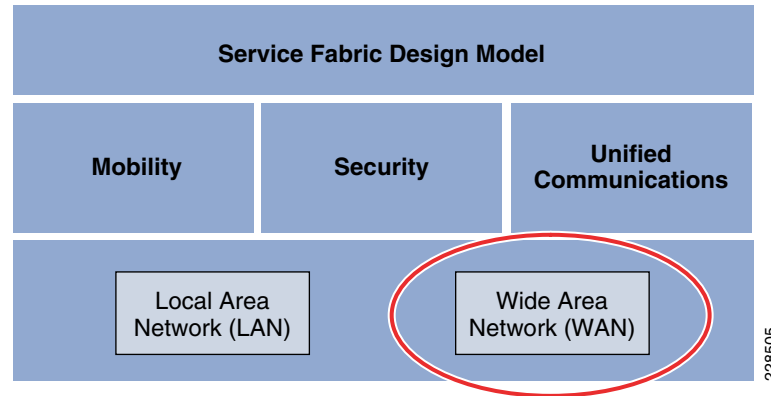
The Cisco Community College reference design is a multi-campus design where a campus consists of multiple buildings and services. The campuses are interconnected through various WAN transports as shown in [Figure 4-1](#).

Figure 4-1 Community College WAN Design Diagram



228504

Within the Community College reference design, the service fabric network provides the foundation on which all the solutions and services are built upon to solve the business challenges facing community colleges. These challenges include virtual learning, secure connected classrooms, and safety and security. This service fabric consists of four distinct components as shown in Figure 4-2.

**Figure 4-2 The Service Fabric Design Model**

This chapter discusses the WAN design component of the community college service fabric design. This section discusses how the WAN design is planned for community colleges, 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 campus locations to the main campus site, as well as connectivity between community colleges, and general Internet access for the entire college. The WAN design should not be viewed merely for providing access, but mainly to see how the business requirements can be met. In today's collaborative learning environment, it is important for communication to exist between students and teachers. 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 Community College WAN design 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 community college.

### Private WAN Service

One of the main requirements for community colleges is the ability to collaborate with other colleges within North America and globally. To achieve the inter connectivity between the colleges, the network should be connected to certain providers, such as Lambda rail, Internet2. The community colleges need to connect to Gigapops—regional networks, which provide access to these private WAN networks. The following sections provide a brief description on these two network types:

Internet2 is a not-for-profit advanced networking consortium comprising more than 200 U.S. universities in cooperation with 70 leading corporations, 45 government agencies, laboratories, and other institutions of higher learning as well over 50 international partner organizations. Internet2 provides its members both leading-edge network capabilities and unique partnership opportunities that together facilitate the development, deployment and use of revolutionary Internet technologies. The Internet2 network's physical implementation is comprised of an advanced IP network, virtual circuit network and core optical network. It provides the necessary scalability for member institutions to efficiently provision

resources to address bandwidth-intensive requirements of their campuses such as, collaborative applications, distributed research experiments, grid-based data analysis and social networking. For more information on the Internet2 network, refer to the following URL:

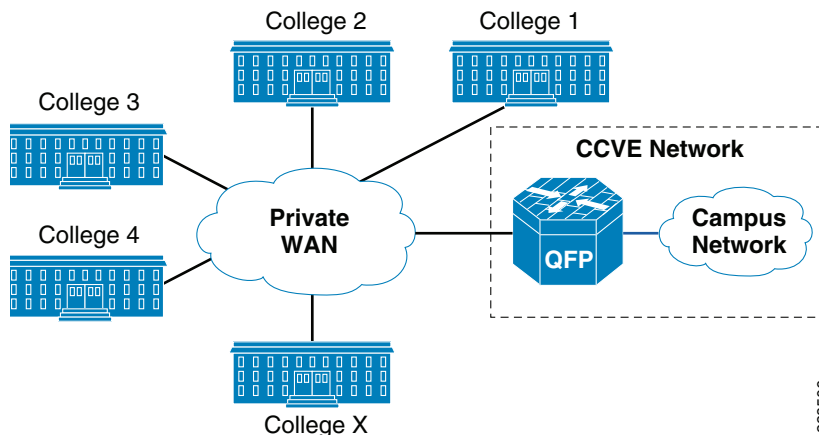
<http://www.internet2.edu/network/>

National LambdaRail (NLR) is a high-speed, fiber-optic network infrastructure linking over 30 cities in 21 states. It is owned by the U.S. research and education community and is dedicated to serving the needs of researchers and research groups. NLR's high-performance network backbone offers unrestricted usage and bandwidth, a choice of cutting-edge network services and applications, and customized service for individual researchers and projects. NLR services include high-capacity 10Gb Ethernet LAN-PHY or OC-192 lambdas, point-to-point or multipoint Ethernet-based transport, routed IP-based services, and TelePresence video-conference services. For more information on the NLR network and its services, refer to the following URL:

<http://www.nlr.net/>

This design assumes that community colleges are connected to one of these networks using either Layer 2 or Layer 3 networks for WAN connectivity, using WAN speeds of 100Mbps. The physical connection is assumed to be one connection to the service provider, but there will be two logical connections—one for accessing private networks, and the second one for Internet access. Figure 4-3 depicts how community college would connect to different colleges, universities, and research networks using either NLR or Internet2 service.

**Figure 4-3** Community College Connection to Other Colleges 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 sub-interfaces. Therefore, it is similar to a situation where a customer is connected to different service providers. See Figure 4-4.



**Figure 4-4 Community College 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 SLA 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 4-5 shows the difference between these services.

Figure 4-5 Different Services Available

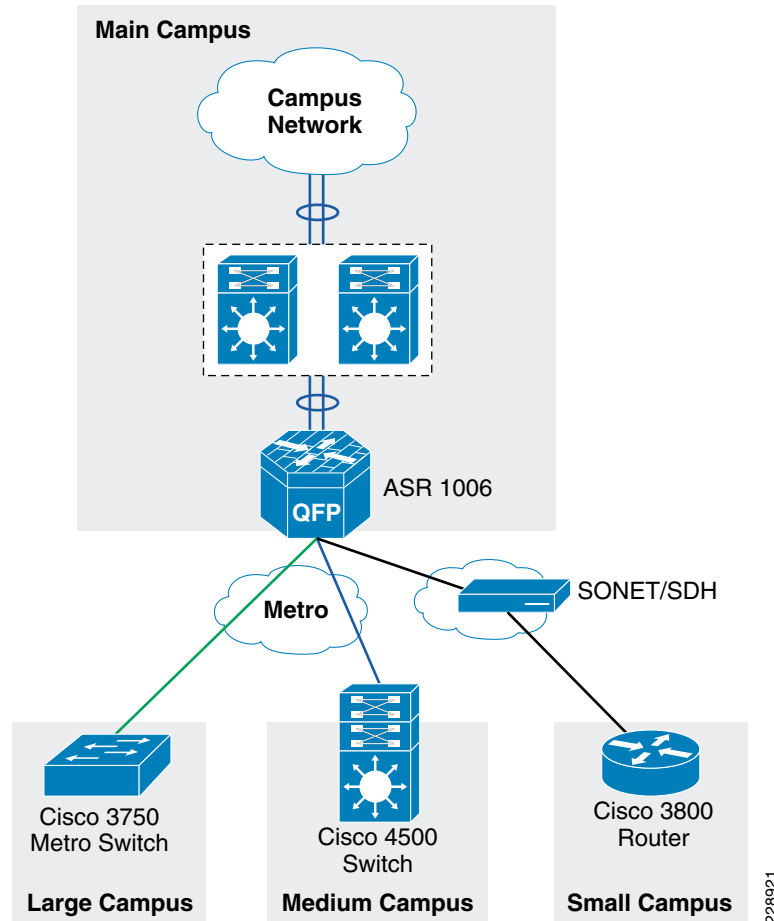


This section discusses how the Metro service is designed in the Community College reference design. The Metro service is used to provide connectivity between the remote campuses to the main campus site. The key reasons for recommending Metro service for community college are as follows:

- *Centralized administration and management*—E-line service provides point-to-point connectivity, where as, E-LAN provides point-to-multipoint connectivity. Having a point-to-point connectivity mandates that all the remote campus sites need to traverse the main campus site to reach the other, making the centralized administration applicable.
- *Performance*—Since all the application services are centrally located at main campus site, the WAN bandwidth required for remote campus sites to main campus 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 remote medium campus locations use E-line service to connect to the main campus site. Figure 4-6 shows how the remote campus locations are connected to main campus site using Metro service.

**Figure 4-6** The Metro Transport Deployment in Community College WAN Design



## Leased-Line Service

The WAN bandwidth requirement for a small remote site is assumed to be 20Mbps. Cisco recommends that the remote small campus site connect to the main campus 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 remote small campus application requirements.

## WAN Aggregation Platform Selection in the Community College Reference Design

In addition to selecting the WAN service for connectivity between college campus locations and access to the Internet, choosing the appropriate WAN aggregation router is essential. For each location in the Community College reference design, various WAN aggregation platforms are selected based on the requirements.

## Main Campus WAN Aggregation Platform Selection

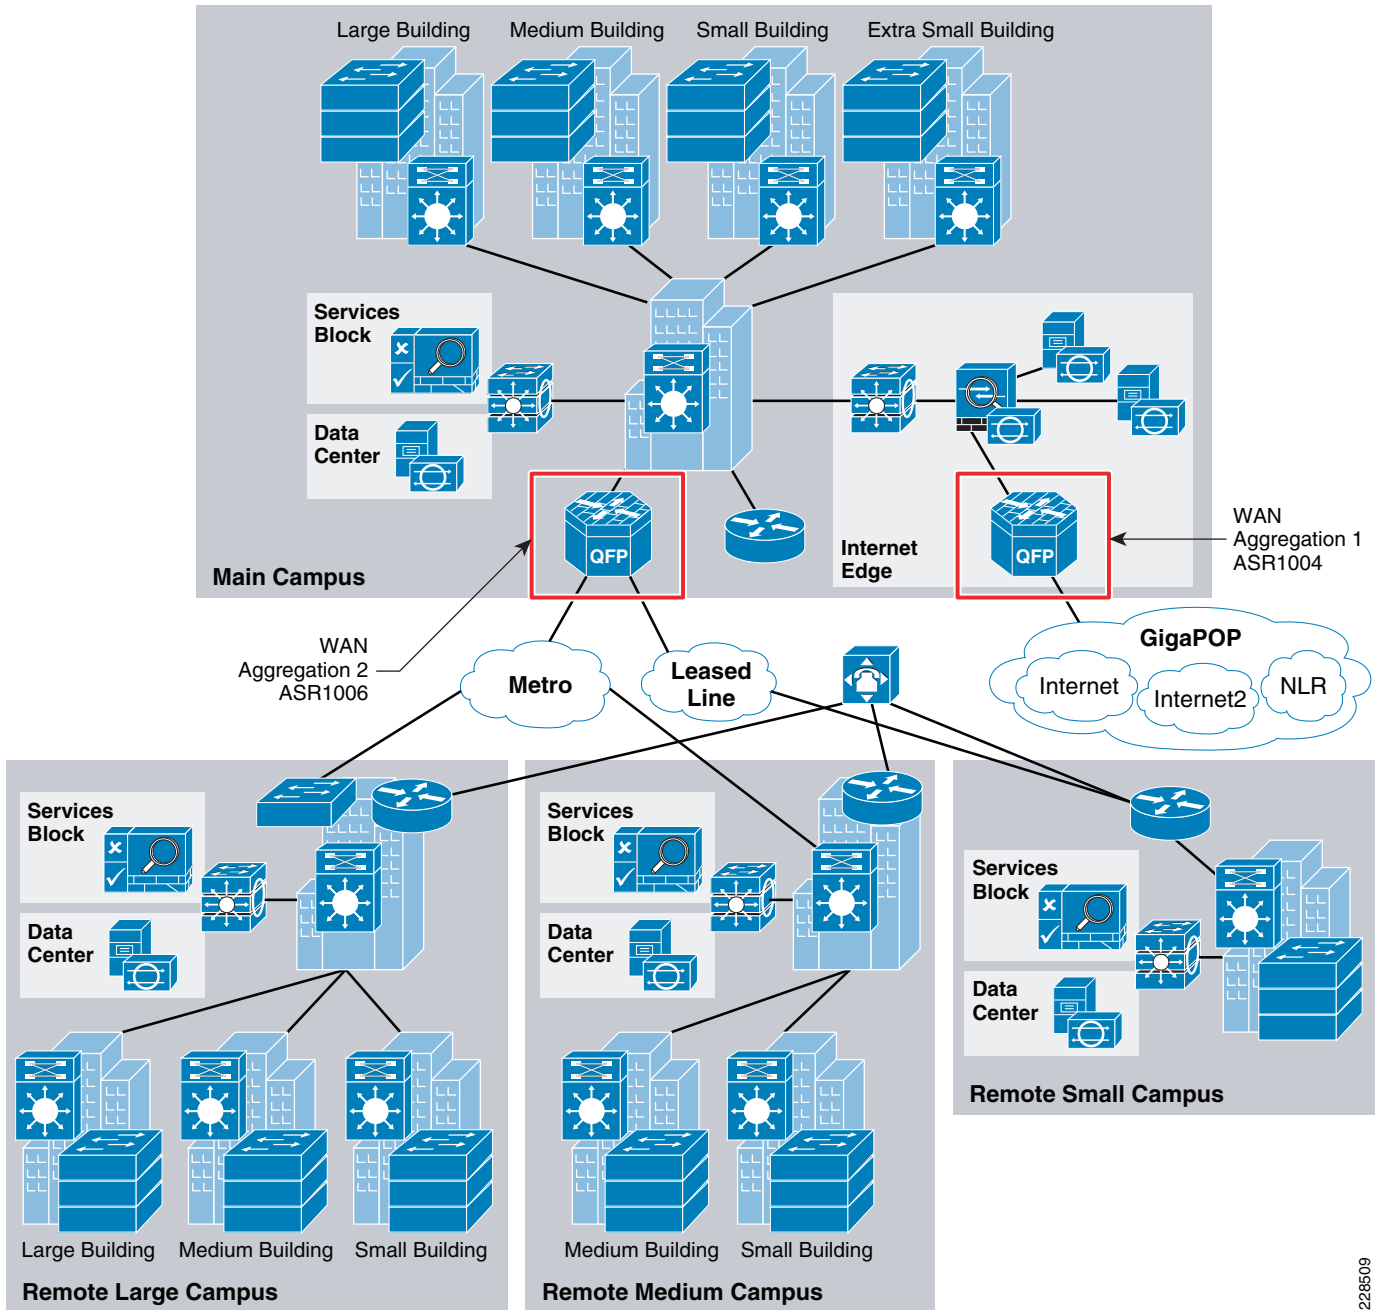
A WAN aggregation router aggregates all the incoming WAN circuits from various locations in the network as well as the Internet and also provides the proper QoS required for application delivery. Cisco recommends the Cisco ASR family of routers as the WAN aggregation platform for the main campus location.

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 community college WAN design, there are two places where the WAN aggregation occurs in the main campus location. The first place is where the main campus location connects to outside world using private WAN and Internet networks. The second place is where all the remote campus locations connect to the main campus sites. [Figure 4-7](#) shows the two different WAN aggregation devices.

Figure 4-7 The WAN Aggregation Points in Community College



228509

### WAN Aggregation 1

Cisco ASR 1004 Series router is recommended as WAN aggregation platform for private WAN/Internet connectivity. This choice was made considering the cost and required features—performance, QoS, routing, and resiliency, which 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 community college campuses. 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.

## Remote Large Campus WAN Aggregation Platform Selection

The WAN connectivity between the large remote campus sites to the main campus site is fairly simpler because of the lack of requirements of advanced encryption technologies. Therefore, the main idea is to reduce the cost and try to consolidate the WAN functionality into the distribution device at the large campus site. However, at the large campus site, as per the campus LAN design document VSS has been chosen as distribution switch, and it does not support WAN functionality. Therefore, a dedicated WAN aggregation device needed to perform that functionality, and the choice 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, BGP
- Multicast support: PIM
- Redundant power supply

## Remote Medium Campus WAN Aggregation Platform Selection

As discussed in [Chapter 3, “Community College LAN Design,”](#) the remote medium campus 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 campus location is connected to the main campus location through Metro service. At the remote medium campus location, the WAN and LAN aggregation platform is the Cisco Catalyst 4507 switch. This switch has 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 or 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.

## Remote Small Campus WAN Aggregation Platform Selection

The remote small campus is connected to main campus using a private leased-line service. The WAN speed between the remote small campus and the main campus 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 can not be used. Therefore, an integrated services router is needed to meet the requirement. For this reason, the Cisco 3845 Series router is chosen as WAN platform for remote small campus. 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 Community College WAN Reference Design

The following section would discuss on the implementation details for Community College WAN Reference design. 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, the Community College WAN design uses two different services to connect remote campus locations to main campus location. The remote large campus site, and remote medium campus sites would connect to main campus site using Metro services. The remote small campus site uses leased-line service to connect to the main campus location. The remote large campus site, due to its size, is recommended to have 1Gbps Metro service to the main campus site where as the remote small campus location is recommended to have at least 20Mbps of bandwidth to main campus 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 campus locations to main campus 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 Remote Large Campus

The following is configuration of WAN interface at 3750 remote large campus switch, which is connected to main campus 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 Remote Medium Campus

The following is the configuration of WAN interface at remote medium campus connected to main campus 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 remote small campus site connect to the main campus 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 remote small campus application requirements. To implement this design, a serial SPA is needed on the ASR 1006 WAN aggregation router at the main campus site and this SPA needs to be enabled for T3 interface type. The following configuration illustrates how to enable and configure the T3 interface:

The following configuration steps are needed to build the lease-line service between the main campus and remote small campus:

---

**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 Remote Small Campus Location

The following is configuration of WAN interface at remote small campus 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 Community College reference WAN design. As indicated in the WAN transport design, the Community College reference design has multiple transports—NLR or I2 networks, Internet, Metro Service, and leased-line services. The NLR or I2 networks would provide access to reach other community colleges, universities, and research networks globally. Internet service would help the Community College to reach Internet. Metro/leased-line service would help to connect remote campus locations to the main campus. To provide connectivity using these transport services we have designed two distinct routing domains – external and internal. The external routing domain is where the Community College 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 would discuss 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, NLR or I2, and the Internet service. This is applicable only to the WAN aggregation router 1, which interfaces with both NLR or I2, and the Internet service, because it the only router which interfaces with the external domain.

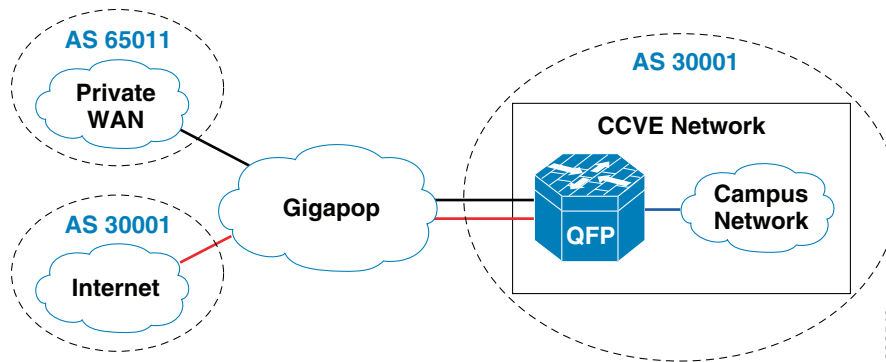
The main design considerations for routing for the Internet/private WAN edge router are as follows:

- Scale up to large number of routes
- Support for multi-homing—connection to different service providers
- Ability to implement complex polices—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 is 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 entries, and outgoing entries. BGP supports having different policies for incoming and outgoing prefixes. [Figure 4-8](#) shows the BGP design.

**Figure 4-8 BGP Design in Community College**



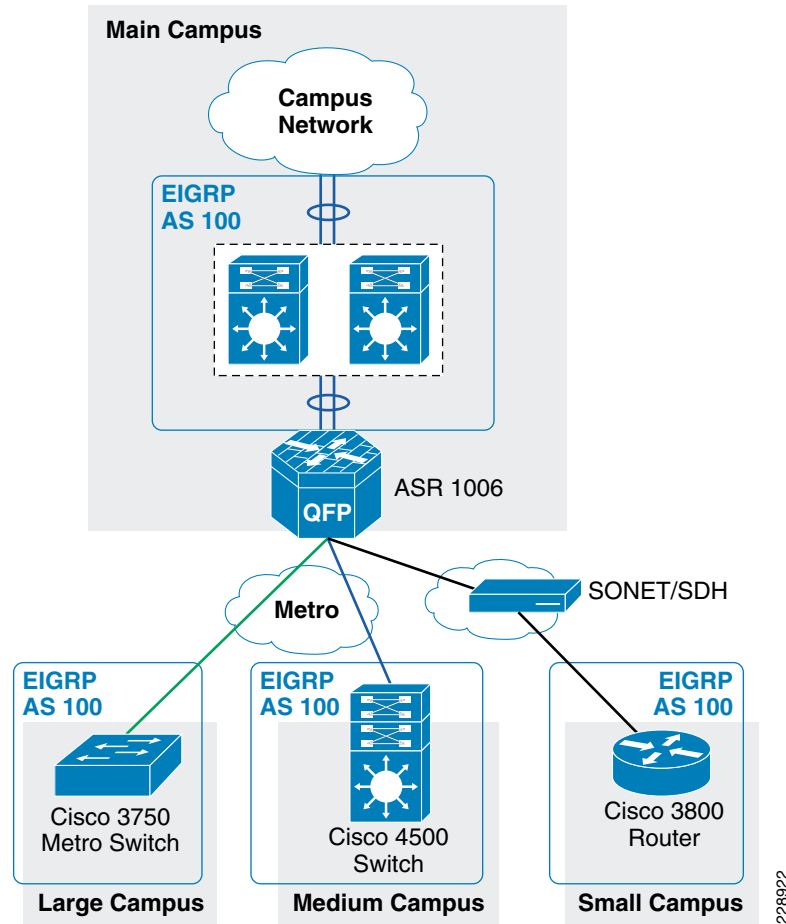
For more information on designing and configuring BGP on the Internet border router, refer to the *SAFE Reference Design* at the following link:

<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 campus 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 college infrastructure with all the design principles defined earlier in this section. CCVE SRA network infrastructure must be deployed in recommended EIGRP protocol design to secure, simplify, and optimize the network performance. [Figure 4-9](#) depicts the design of EIGRP for internal network.

Figure 4-9 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 remote large campus location.
3. The 100Mbps Metro link to remote medium campus location.
4. 20Mbps leased-line service to remote small campus 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-channell
  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-channell
  nsf

```

The ASR1006 router is enabled with non-stop 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 Remote Large Campus Switch

The EIGRP configuration at 3750 remote large campus site also has similar steps compared to Main campus 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 Campus 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 Remote Small Campus 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 “[Deploying Community College Network Foundation Services](#)” section on page 3-25.

## QoS


QoS is a part of foundation services, which is very critical to the application performance. Today’s networks are rapidly converging into IP network. The traditional applications, which used the networks, were voice, video, and data. However, 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 rich 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 4-10](#), 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 4-10** 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 just mentioned 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 number of 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 also 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.

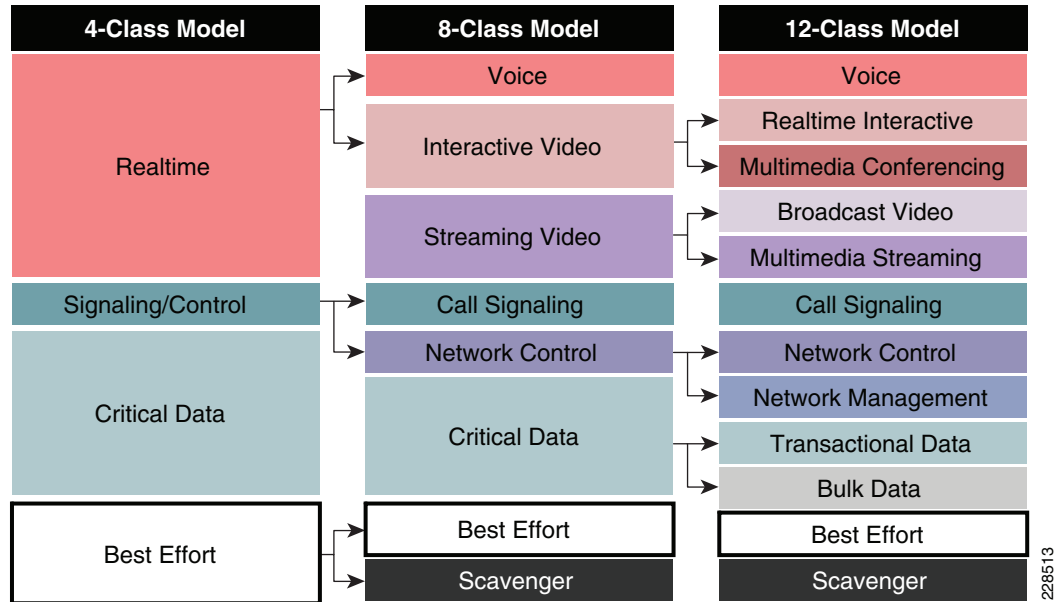
When designing the QoS for WAN architecture, there are two main considerations to start with:

- Whether the service provider will provide four classes of traffic.
- 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 community college campus LAN supports 12 classes of traffic, which will be mapped to 4 classes of traffic on the WAN side. [Figure 4-11](#) illustrates the recommended markings for different application traffic.



**Figure 4-11 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 4-1 describes the different classes, and the percentage, and actual bandwidth allocated for each class of traffic.

**Table 4-1 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 community college WAN design network. 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:

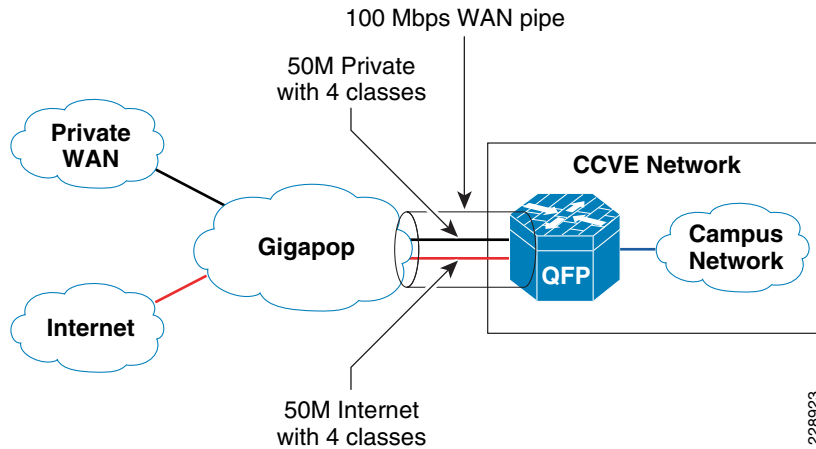
- WAN aggregation router 1 for connection to the Internet and NLR network
- WAN aggregation router 2 for connection to remote campus sites
- Cisco 3750 Metro switch at the remote large campus
- Cisco 4500 switch at the remote medium campus
- Cisco 3800 router at the remote small campus

## QoS Implementation at WAN Aggregation Router 1

The WAN aggregation router1 connects to two different providers: NLR network and Internet. It is assumed that the aggregate bandwidth is 100Mbps that should be shared between both services—50Mbps is dedicated for NLR 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 implemented.

[Figure 4-12](#) depicts the bandwidth allocation at the WAN aggregation router 1.

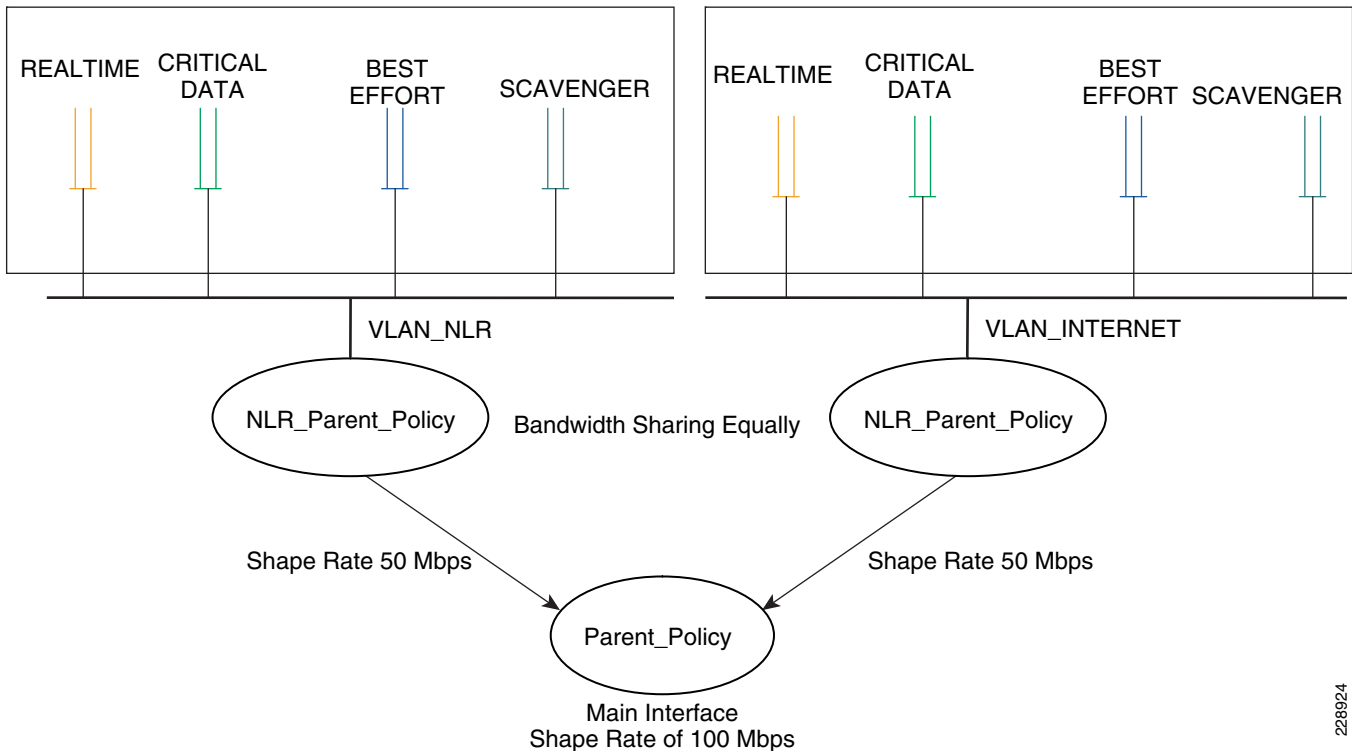
**Figure 4-12 The Bandwidth Allocation at WAN Aggregation Router 1**



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 sub-parent policies are defined, which would further shape it to 50Mbps. Within each of the sub-parent policies, there are four defined classes: REALTIME, CRITICAL\_DATA, BEST\_EFFORT, and SCAVENGER classes.

Figure 4-13 depicts this hierarchical QoS design.

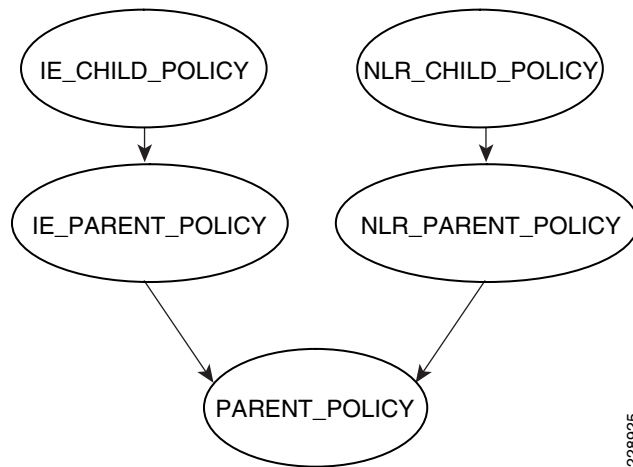
**Figure 4-13 Hierarchical QoS Design**



228924

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 sub-parent policies—These would be specific to each service type. For example, NLR\_PARENT is a policy dedicated for NLR traffic, and NLR\_Internet is specific to Internet traffic.
  - Step 3** Define the child policies—Classifies, queues, and allocate bandwidth within each sub-parent policy. For example, NLR\_PARENT would have a NLR\_Child policy that would classify, queue, and allocate the bandwidth within each allocated bandwidth. The following diagram shows the hierarchical allocation.



## 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 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 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 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
!
    
```

Dummy class does not classify anything

Defining service-fragment would allow other policies to point for share of bandwidth.

The parent policy would shape to 100 Mbps.

Parent policy allocates 50% of bandwidth  
Child policy gets attached to parent 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

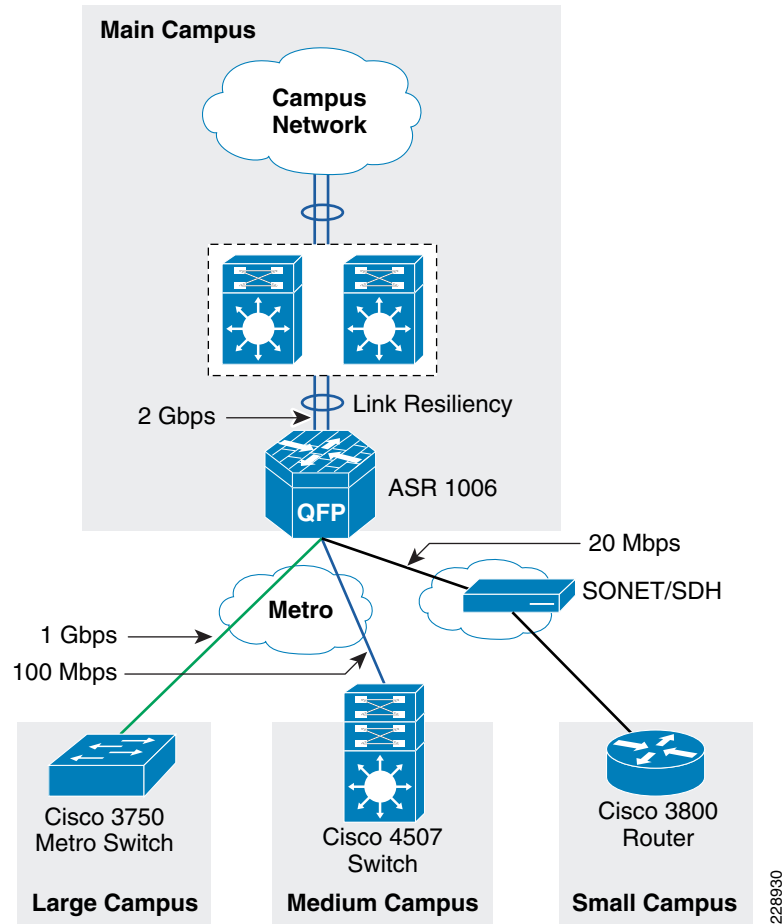
The parent policy applied on sub-interface

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 4-14](#) depicts the different types of WAN speeds

Figure 4-14 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 campus, and large campus is 1Gbps. Therefore; a single-layer QoS policy can be defined on the link.
- The SLA between the main campus, and medium campus is assumed to be 100Mbps; however, link speed is assumed to be 1Gbps. In addition, there is an assumption that there could be more than one medium campus present in this design. Therefore, each remote medium campus would connect to the main campus using these 100Mbps links, requiring a three-layer hierarchical QoS policy is needed. The link between the main campus and remote small campus is 20Mbps. The physical link speed is 44Mbps, requiring a two-level hierarchical QoS policy is needed.
- The Ether channel 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.

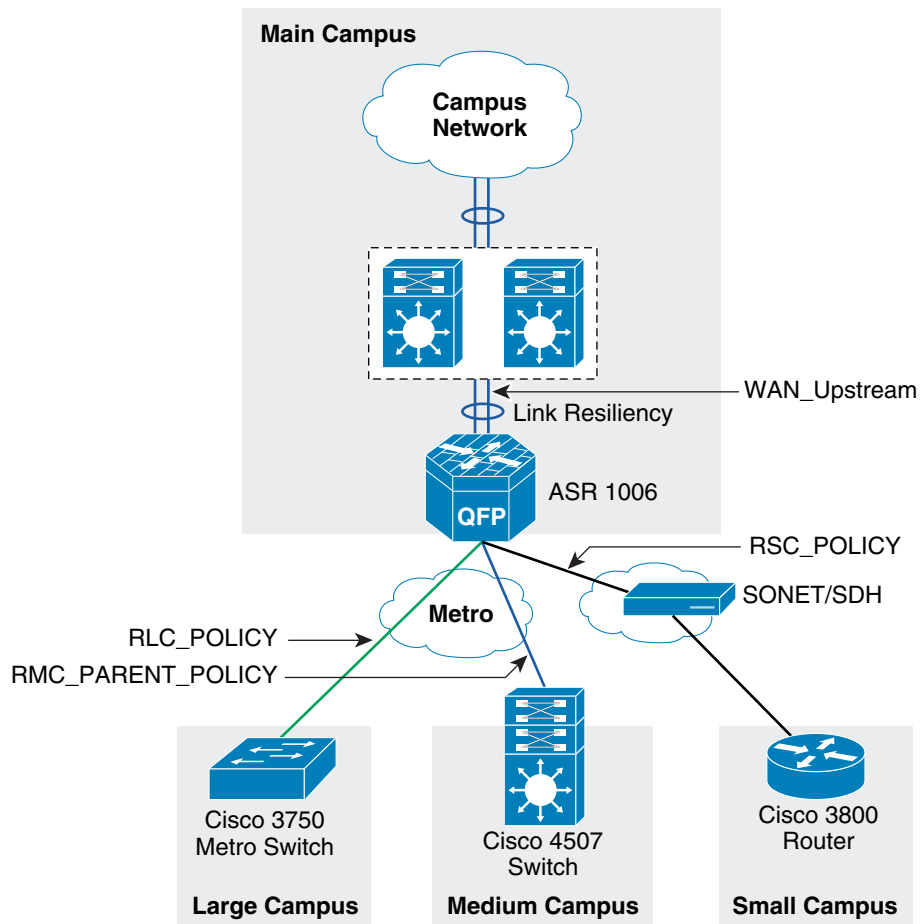
Table 4-2 describes the different QoS policy names applied at the WAN aggregation router 2.

**Table 4-2 QoS Policy for WAN Aggregation Route 2**

QoS Policy Name	Description	WAN Speed
RLC_POLICY	Applied on link between main campus, and remote large campus	1Gbps
PARENT_POLICY RMC_PARENT_POLICY RMC_CHILD_POLICY	Hierarchical QoS Policy between the main campus, and remote medium campus location.	100 Mbps
WAN_Upstream	Applied on link between main campus, and core	2Gbps
RSC_PARENT_POLICY RSC_POLICY	Applied on link between main campus and small campus	20Mbps

Figure 4-15 depicts the various points where QoS policies are applied.

**Figure 4-15 The allocation of QoS Policy at Different Places on WAN Aggregation Router 2**



228931



## QoS Policy Between the Main Campus and Large Campus

The WAN physical link speed is 1Gbps. Also, the actual SLA between the main campus and large campus 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 main campus to the large campus 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 Campus and Medium Campus Location

A three-layer QoS design is needed between the main campus site and large medium campus location because there could be couple of medium campus locations connected on a single metro link to the main campus site. [Figure 4-16](#) shows how this design looks like when there are more than one medium campus site.

**Figure 4-16** *The WAN Link Design for Connectivity Between Main Campus and Remote Medium Campus*



228935

Here, the implementation details are provided for only a single medium campus location; however, more medium campus 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-fragment 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

```

228938

**QoS Policy Between Main Campus and Remote Small Campus Location**

The following is the QoS policy implementation steps between main campus and remote small campus 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

```

228939

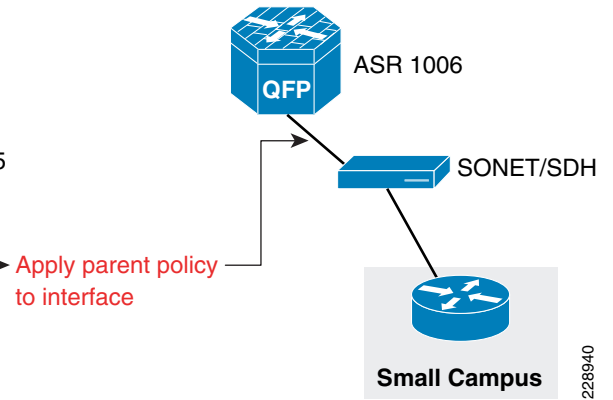
**Step 2** Apply the policy map to the interface.

```

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#

```



## QoS Policy Implementation Between the Main Campus and Core

The following is the QoS policy implementation between main campus 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 of 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

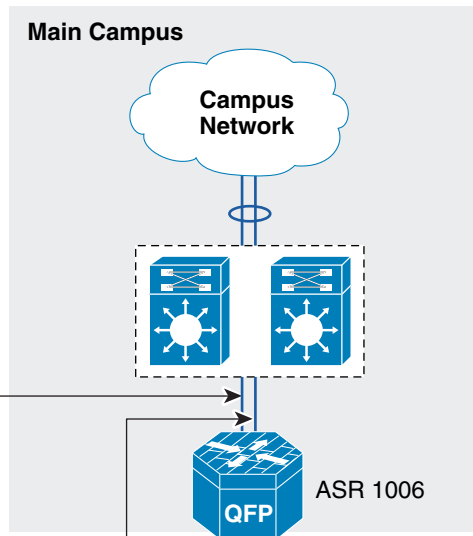
**Step 2** Apply the policy-map on both interfaces going up to the core.

```

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 Campus and Main Campus Location

The WAN interface between the large campus and main campus 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 Medium Campus and Main Campus Location

The medium campus 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
!
```



228948

## QoS Policy Implementation Between Remote Small Campus and Main Campus Location

The following section describes the QoS policy implementation between the remote small campus location and main campus. 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
```



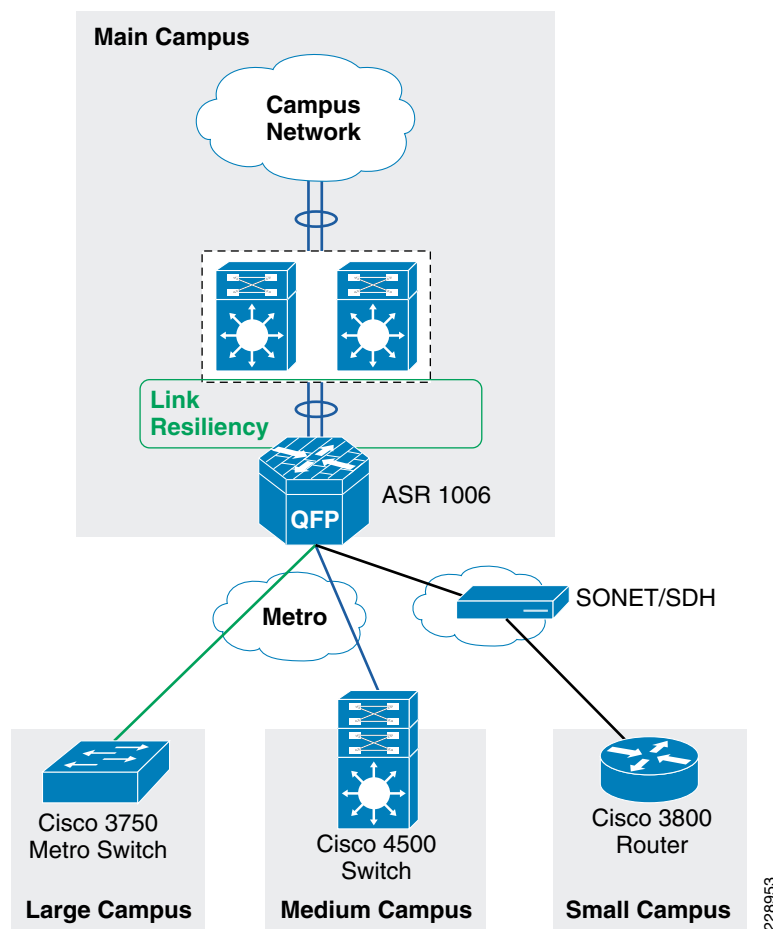
## Redundancy

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. One 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 the 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, where as 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 4-17](#) shows where this feature is enabled.

**Figure 4-17 Link Resiliency**



[Table 4-3](#) shows the various WAN devices that are designed for resiliency.

**Table 4-3 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 Community College reference design for the WAN design. To enable 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 Community College reference design the redundancy is planned at both WAN aggregation router1, and WAN aggregation router 2 in the main campus location. As explained in the WAN aggregation platform selection for the main campus location discussion ASR routers have been selected at both WAN aggregation locations places. However, we have different models, at both WAN aggregation places. When the ASR router interfaces with the private WAN, Internet networks the ASR 1004 with IOS-based redundancy has been chosen. Similarly, for the ASR router that interfaces with Metro connections, the ASR 1006 with dual RP, and dual ESP to provide for hardware-based redundancy has been chosen. 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 you 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 can achieve non-stop forwarding of data even when there 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 implementation steps to deploy link resiliency:

### Step 1 Configure the ether channel 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-channell
  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 Community College Reference design, we are assuming that multicast traffic would be present only within the campus, and not between the community colleges. Therefore, the multicast design looks at only between the main campus, and remote small campus 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 campus, refer to “[Multicast for Application Delivery](#)” section on page 3-63.

## 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 router2 that connects to different remote campus sites.

### Step 1 Enable multicast routing:

```
ip multicast-routing distributed
```

### Step 2 Enable PIM-Spare mode on the following WAN interfaces:

- Port-channel—Connects to the VSS core
- Gi0/2/0—Connects to remote large campus site
- Gi0/2/1—Connects to remote medium campus site
- S0/3/0—Connects to remote small campus 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 Remote Large campus

This section discusses how to implement multicast on remote large campus 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 campus 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 Remote Medium Campus

This section discusses on how to implement multicast on remote medium campus 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 Remote Small Campus Site

This section discusses on how to implement multicast on remote small campus site.

---

**Step 1** Enable multicast routing:

```
ip multicast-routing
```

**Step 2** Enable PIM Spare 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 Community College reference design 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 community college network. Key WAN design principles such as WAN aggregation platform selection, QoS, multicast and redundancy best practices were discussed for the entire community college design. Designing the WAN network of a community college 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 community college.





## CHAPTER 5

# Community College Mobility Design

---

## Mobility Design

The Cisco Community College reference design is intended to assist community colleges in the design and deployment of advanced network-based solutions within twenty-first century learning environments.

The reference design addresses the business challenges currently facing community colleges. At the heart of the reference design 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 community college network to solve business problems, which include the following:

- Safety and security
- Virtual learning
- Secure connected classrooms
- Operational efficiencies

Community college students are dynamic, mobile, and technology-savvy. When on campus, they move about while equipped with an array of mobility-enabled devices including PDAs, phones, and laptops. In contrast to the typical enterprise business environment, community colleges consist of a large student population that typically experiences a complete turnover every few years. Typical community college students tend to use new applications and the network for many aspects of their lives, demanding connectivity wherever they are. This connected generation is untethered from wired access connectivity and assumes the presence of a high-performance, reliable wireless LAN (WLAN) in all major campus areas.

The mobility design implemented by a community college must meet the needs of this mobile generation while also addressing the requirements of faculty, staff, administrators, and visitors. The challenge for community colleges is to create a robust, end-to-end, mobility-enabled network that supports their requirements at a cost that is within their often constrained budgets. Community colleges should be equipped with a mobility solution that supports the following:

- Secure communications between local and remote campus sites to support students, faculty, staff, administrators, and visitors, using the new generation of mobility-enabled devices and applications in the current marketplace
- A scalable design model that can easily accommodate the addition of new campus buildings as well as existing building modifications
- Built-in support for bandwidth-intensive, high-speed multimedia applications
- Simplified management tools to facilitate maintenance of the system-wide mobility solution
- The use of new tools and applications for mobile learning, collaboration, and campus operations

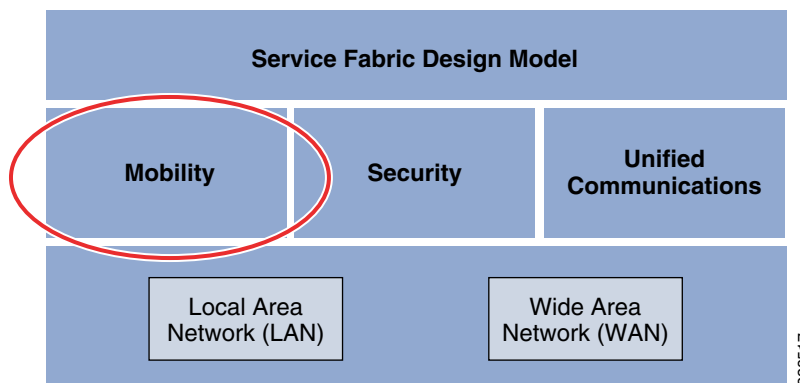
- Effective communication with public safety first responders in the event of an emergency

In addition, each community college must remain competitive, differentiating itself from its peer institutions so as to attract and retain the best students and faculty. Students want to attend quality community colleges that provide technology services relevant to the way they live, work, and learn. They want to take full advantage of community college capabilities to facilitate their success while they are students, as well as when they are pursuing post-graduation placement. A community college with a pervasive, high-speed wireless network not only demonstrates technological leadership and innovation, but enables the deployment of innovative applications that improve learning, the streamlining of operations, collaboration enhancements, and productivity improvements.

This mobile campus lifestyle helps to drive the need for careful wireless capacity and coverage planning. Keep in mind that the traditional scenario of a mass of students filing into a large lecture hall within a monolithic campus building is no longer the only learning environment seen within higher educational institutions. High performance, secure wireless technologies can enable “virtual classrooms” even in non-traditional settings, such as leased space in shopping malls, retail plazas, and even from homes and offices. School administrators need secure access to tools, records, and resources, as well as access to mobile voice capabilities throughout the campus. In addition, the expectation for secure, reliable, high-performance guest access by contractors, vendors, and other guests of the community college establishment has become a standard and expected component of doing business.

To meet these and other student, faculty, and guest needs, community colleges must evolve into mobility-enabled campuses and twenty-first century learning centers. The primary objectives of this document are the design considerations surrounding the requirements and expectations that must be considered when integrating mobility into the Cisco Community College reference design, as well as the tradeoffs required to facilitate the four service requirements stated previously. These design considerations form a critical part of the overall service fabric design model, shown in [Figure 5-1](#).

**Figure 5-1 Service Fabric Design Model**

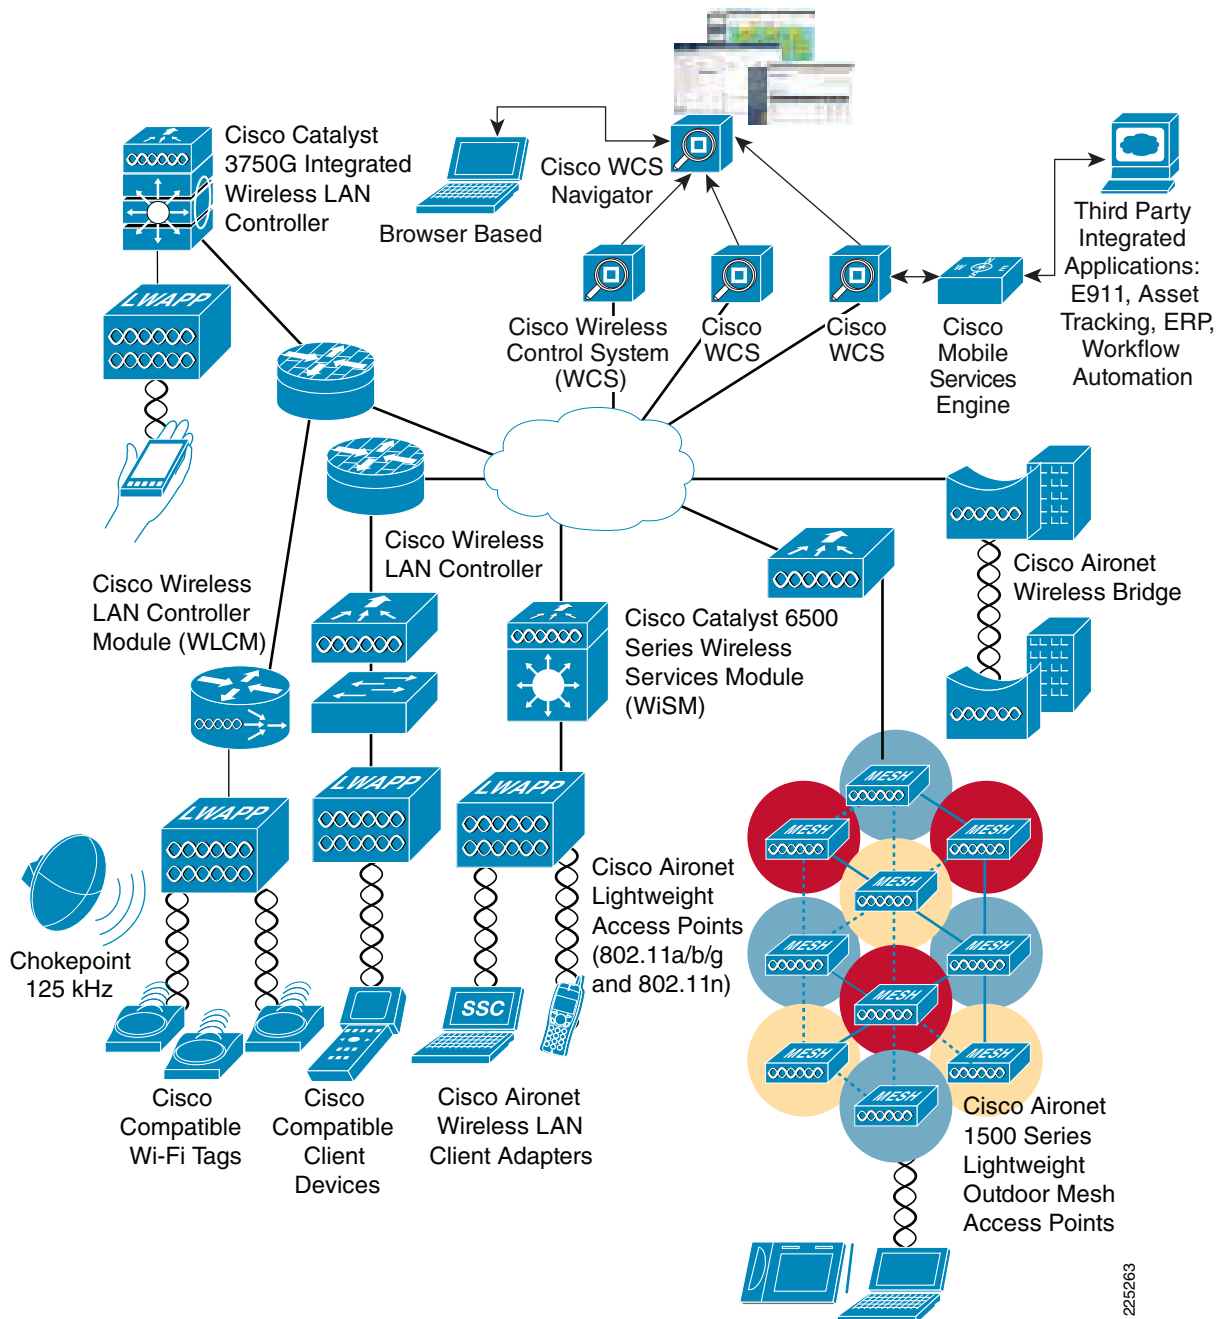


Given the mobility of students, staff, and visitors, wireless LANs have emerged as one of the most effective and high performance means for these mobile users to access the campus network. The Cisco Unified Wireless Network (Cisco UWN) is a unified solution that addresses the wireless network security, 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 5-2](#) 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

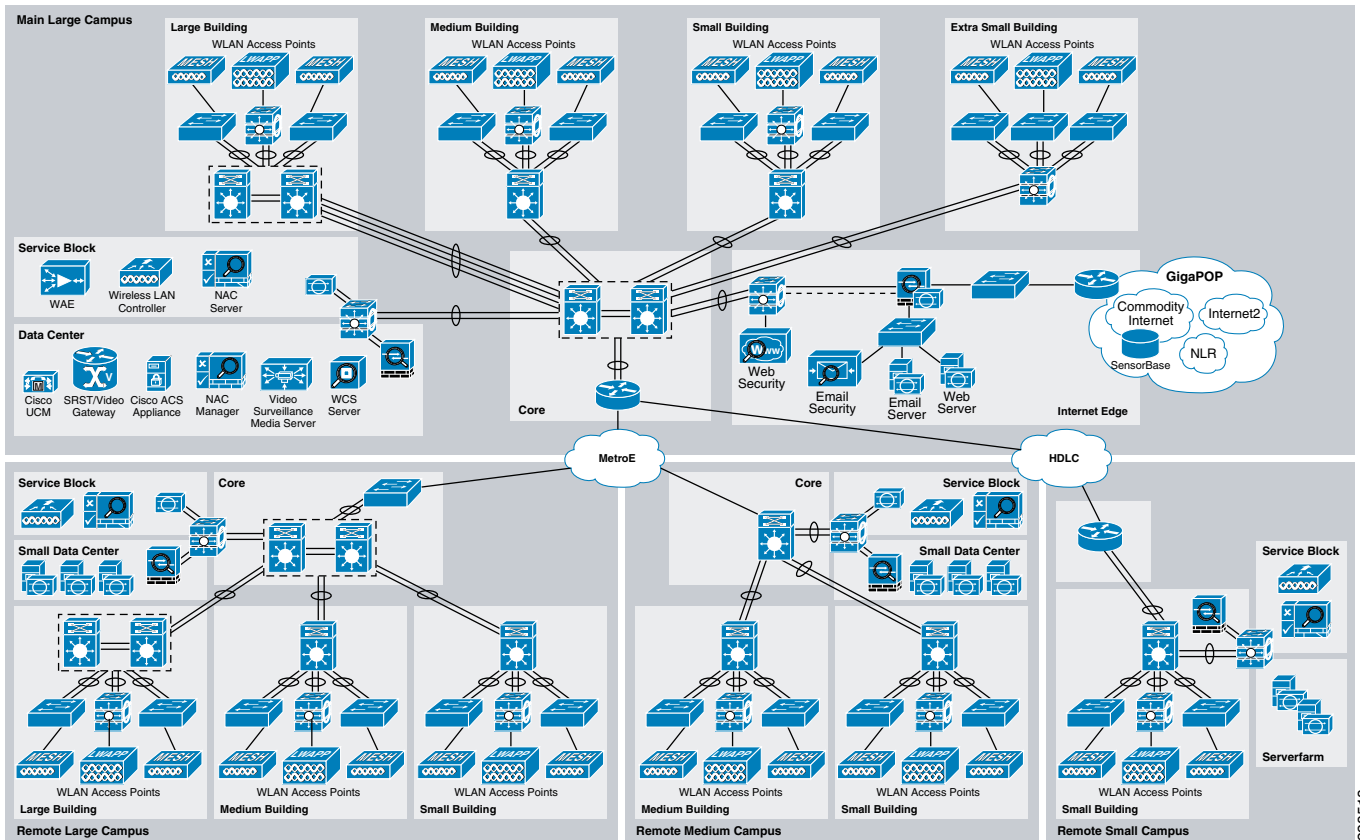
Cisco Access Control Server (ACS) and its Authentication, Authorization, and Accounting (AAA) features complete the solution by providing Remote Authentication Dial-In User Service (RADIUS) services in support of user authentication and authorization.

Figure 5-2 Cisco Unified Wireless Network Overview



The Cisco Community College reference design accommodates a main campus and one or more remote smaller campuses interconnected over a metro Ethernet or managed WAN service. Each of these campuses may contain one or more buildings of varying sizes, as shown in Figure 5-3.

Figure 5-3 Community College Reference Design Overview



Operating on top of this network are all the services used within the community college 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 campus building, allowing each remote campus to reduce the need for separate services to be operated and maintained by community college IT personnel. These centralized systems and applications are served by a data center in the main campus.

As Figure 5-3 shows, the Cisco Community College reference design uses a centralized approach in which key resources are centrally deployed at either the campus or college level. The key feature of this integration is the use of one or more WLAN controllers at each campus, with the overall WLAN management function (the Cisco WCS) located at the main campus. 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 Community College reference design takes into account that cost and limited network administrative resources are common limiting factors for most community colleges. The topologies and platforms are carefully selected to increase productivity while minimizing the overall cost and complexity of operation. In certain instances, tradeoffs are necessary to reach these goals, and this document points out such areas.

The Cisco mobility approach within the Cisco Community College reference design focuses on the following key areas:

- *Accessibility*

- Enabling students, staff, and guests to be accessible and productive on the network, regardless of whether they are in a traditional classroom setting, collaborating in a study hall, having lunch with colleagues within campus eating areas, or simply enjoying a breath of fresh air outside a campus building
- Enabling easy, secure guest access to college guests such as alumni, prospective students, 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. Faculty and other 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 Community College reference design, see [Chapter 6, “Community College Security Design.”](#)

---

- *Manageability*

A relatively small team of college network administrators must be able to easily deploy, operate, and manage hundreds of access points that may reside within a multi-campus community college. A single, easy-to-understand WLAN management framework provides small, medium, and large community colleges with the same level of WLAN management scalability, reliability, and ease of deployment demanded by traditional enterprise business customers.
- *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 students, faculty, staffs 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 Community College reference design 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 Community College reference design 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 5-4 shows the use of CAPWAP by access points to communicate with and tunnel traffic to a WLAN controller.

**Figure 5-4 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 Community College reference design 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 campus deployments are typically categorized into two main categories, *distributed* and *centralized*:

- *Distributed controller*—In this model, WLAN controllers are located throughout the campus 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 campus 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 in the network. 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 campus backbone network.

The Cisco Community College reference design is based on the centralization of WLAN controllers, on a per-campus 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 5-3 shows the planned deployment of WLAN controllers within distinct per-campus service blocks, each associated with the main, large remote, medium remote, and small remote campus 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-campus basis include the following:

- *Reduced acquisition and maintenance costs*—By servicing the needs of all campus 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 campus WLAN.
- *Reduced administrative requirements*—By minimizing the total number of WLAN controllers deployed, the controller management burden imposed on community college campus 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 the campus 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 campus, 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 campuses, there is also an incremental economy of scale that occurs as the network grows larger. By centralizing WLAN controllers on a per-campus

basis, CAPWAP tunneling between access points and WLAN controllers is not normally required to traverse WAN links (except during controller failover), 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 Community College reference design is the Cisco 5508 Wireless Controller, as shown in [Figure 5-5](#).

**Figure 5-5** 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 enterprises and campus 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 community college campus growth occurs. In campuses 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 Community College reference design, 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 5-6](#).

**Figure 5-6 WLAN Controller Link Aggregation to Services Block**

In this way, one or more centralized WLAN controllers are connected via the services block to the campus 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 campus 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 Community College reference design can be found in [Controller Link Aggregation, page 5-37](#).

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 campus 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 5-6](#)), 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 community college campus system, such as the faculty, administrators, students, and guests. These considerations include the WLAN controllers deployed in the campus services blocks, as well as the access points that are located in the campus buildings.

### Defining WLANs and SSIDs

In most community colleges, various campus user groups likely require access to the WLAN for a variety of purposes. Although peaks in usage may occur at different times, it is safe to assume that a large portion of these groups will likely want access to the WLAN at the same time. Thus, in designing for mobility within the Cisco Community College reference design, the physical campus wireless infrastructure needs to 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 5-7 provides a high-level illustration of the four logical WLANs that provide mobility within the Cisco Community College reference design, and how they are mapped to WLAN controller network interfaces or tunneled to another controller. For ease of administration and the support of students, faculty, and guests that frequent multiple campuses, the names chosen for the WLAN SSIDs should be consistent within each campus in the community college system. For example, student wireless access should be available anywhere there is WLAN RF coverage within this particular community college system using the SSID entitled *student*.

**Figure 5-7** WLAN SSIDs



228521

In the Community College reference design, 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 college faculty, staff, and administration 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 controlled by the community college network administration staff. Unlike the student and guest access WLANs, devices that are used on the secured staff WLAN are usually procured and deployed by (or with the knowledge and cooperation of) the community college network administration staff on behalf of faculty and other university staff users. Faculty and staff users 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 baseline level of authentication and encryption to be deployed for the secured staff WLAN without concern for whether or not the devices using the secured staff WLAN can support this level of authentication and encryption.

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 5-27](#).

---

- 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 college faculty, staff, and administration using managed 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 community college network administration staff. Such procurement is usually conducted on behalf of faculty and other university staff users. To assure proper security and promote effective device management, faculty and 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 can 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 *student WLAN* that uses web authentication for wireless access to the network using unmanaged and privately owned clients such as laptops, PDAs, iPod Touch, iPhones, and so on.

This method of access is normally simple enough for all WLAN users and all platforms, regardless of manufacturer or model. A key challenge in managing wireless access for any large population of users possessing the freedom to choose their wireless clients is how to provide ubiquitous access while still providing an acceptable level of security. Because the ratio of students to network administrative staff is so heavily skewed in favor of the number of students, any student access WLAN solution should require virtually “zero touch” from campus community college network staff, while allowing the vast majority of devices on the marketplace to successfully connect to the network. Characteristics of the student WLAN include the following:

- 802.1x /EAP authentication is not used. For simplicity of configuration across all devices, encryption is not configured on the student WLAN. Transport-level or application-layer encryption may be used if deemed applicable.
- To provide access control and an audit trail, the student 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 username and password into an authentication web page.
- The student WLAN client device user is re-directed to the web authentication web page whenever the client attempts to open any web page before successful web authentication. This authentication web page can be provided either by an internal WLAN controller web server or the NAC appliance in the Cisco Community College reference design. Usernames and passwords for authentication can reside on a RADIUS AAA server (such as Cisco ACS).
- Broadcast SSID is enabled.
- QoS profile setting of *silver* (best effort delivery).
- WMM policy is set to *allowed*.

- Radio policy should be set such that client access is allowed using either 2.4 GHz or 5 GHz.
- A *guest access* WLAN that uses web authentication for guest users of the campus 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 Community College reference design 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) within the Cisco Community College reference design. 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 campus 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 5-28](#) for further information regarding considerations surrounding the products and techniques used to provide guest access when designing for mobility in the Cisco Community College reference design.

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>.

Similar to the requirements stated earlier for the student access WLAN, the guest access WLAN must also be designed to accommodate campus guests (such as alumni, vendors, contractors, prospective students, parents, and so on) as well as the wide variety of WLAN guest clients they may bring onto the campus. Although their numbers will likely be much less compared to that of students, the WLAN clients brought onto campus by guest users are typically not managed or directly supported by community college campus 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 is usually not practical for guest access.

Characteristics of the guest access WLAN include the following:

- The guest access WLAN uses web authentication in a fashion similar to what was described in the student access WLAN, in order to provide access control and an audit trail.
- 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 Community College reference design. However, there is an option of using a non-controller-based web authentication server, such as the Cisco NAC Appliance. Usernames 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 5-8](#). Here we see a large and a medium building located on the same campus, in relatively close proximity to one another, with a small building located on a remote campus some distance away. 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 located within the main campus service block, and the WLAN controller servicing the smaller building located on the remote campus. The circular and oval patterns surrounding each building are intended to represent a very simplistic view of hypothetical outdoor RF coverage.

**Figure 5-8** Campus Roaming



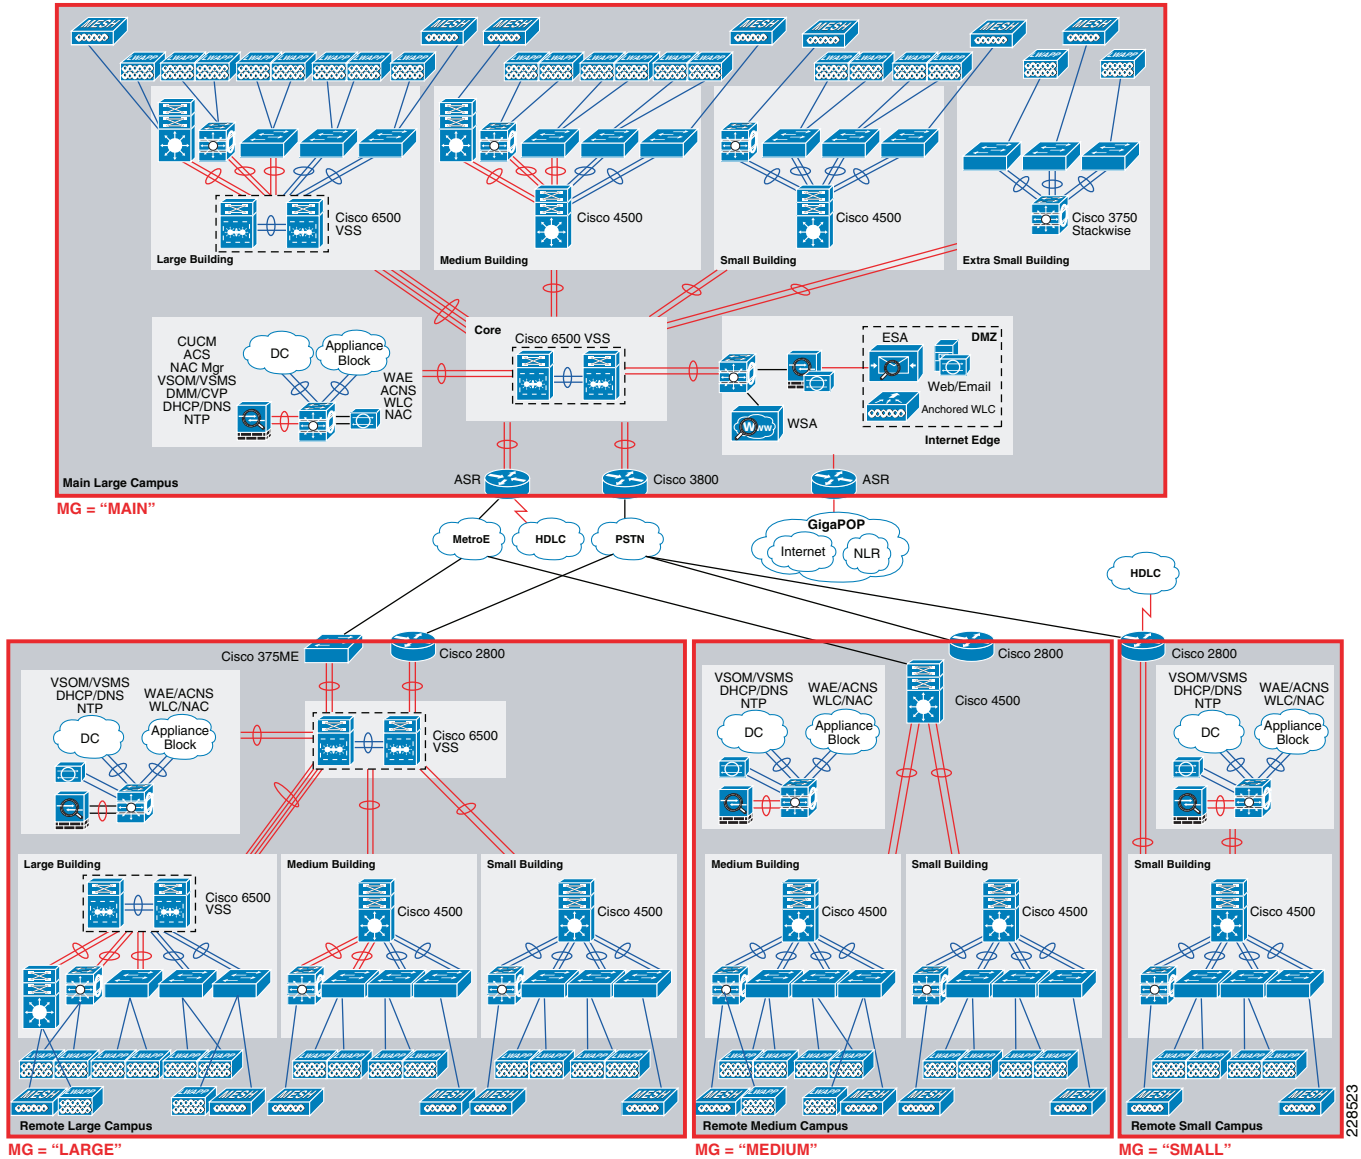
[Figure 5-8](#) shows that there is overlapping coverage between the large and medium buildings, but not between the small building and any other building. This is because users must leave the main campus and traverse through a part of the town to get to the smaller remote campus, and vice versa. Because roaming is clearly possible between the medium and large building, but not between the small building and any other building, 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 in this case.

In applying the concept of mobility groups to the Cisco Community College reference design, consider the following:

- Within a community college or community college system comprised of one or more campuses, it is assumed that intra-campus roaming is possible between all buildings resident on the same campus. This may not actually be the case in all campuses, as some may have buildings co-located on the same campus where areas of non-coverage exist between them. However, assuming that intra-campus roaming is possible between all buildings allows us to make a design assumption that is generally applicable to both situations. Thus, in our Community College reference design, all WLAN controllers serving access points deployed on the same campus are placed within the same mobility group.
- It is also assumed that in the vast majority of cases, remote campuses are sufficiently distant from the main campus (as well as from one another) to render inter-campus roaming impractical. Allowing for the rare exception that two campuses may be adjacent or otherwise overlap one another, for the most part it is assumed that roaming between buildings located on different campuses is very unlikely.

Figure 5-9 provides a high-level illustration of how mobility group assignment can be handled in the Community College reference design. Note that *MG* refers to the mobility group name assigned for the campus.

Figure 5-9 Community College Mobility Groups



The following are some of the key design considerations concerning mobility groups:

- The controllers present at each campus are defined as members of a mobility group unique to that campus. 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-campus roaming between two campuses is possible, the controllers at both campuses 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 campuses, access point failover to a remote backup controller resident at the main campus becomes feasible. To support this, access points can be configured to failover to a WLAN controller outside of their mobility group. This is discussed further in [Controller Redundancy, page 5-40](#) and [AP Controller Failover, page 5-42](#).

- 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 times 250, or 18,000 access points.

The advantage of this approach to mobility group use is clarity and simplicity in deployment and administration. This is a key point when keeping in mind that the typical community college has a limited network administrative staff that is usually resource-constrained and very busy. By dividing the community college system into mobility groups as indicated in [Figure 5-9](#), 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 a significant tradeoff.

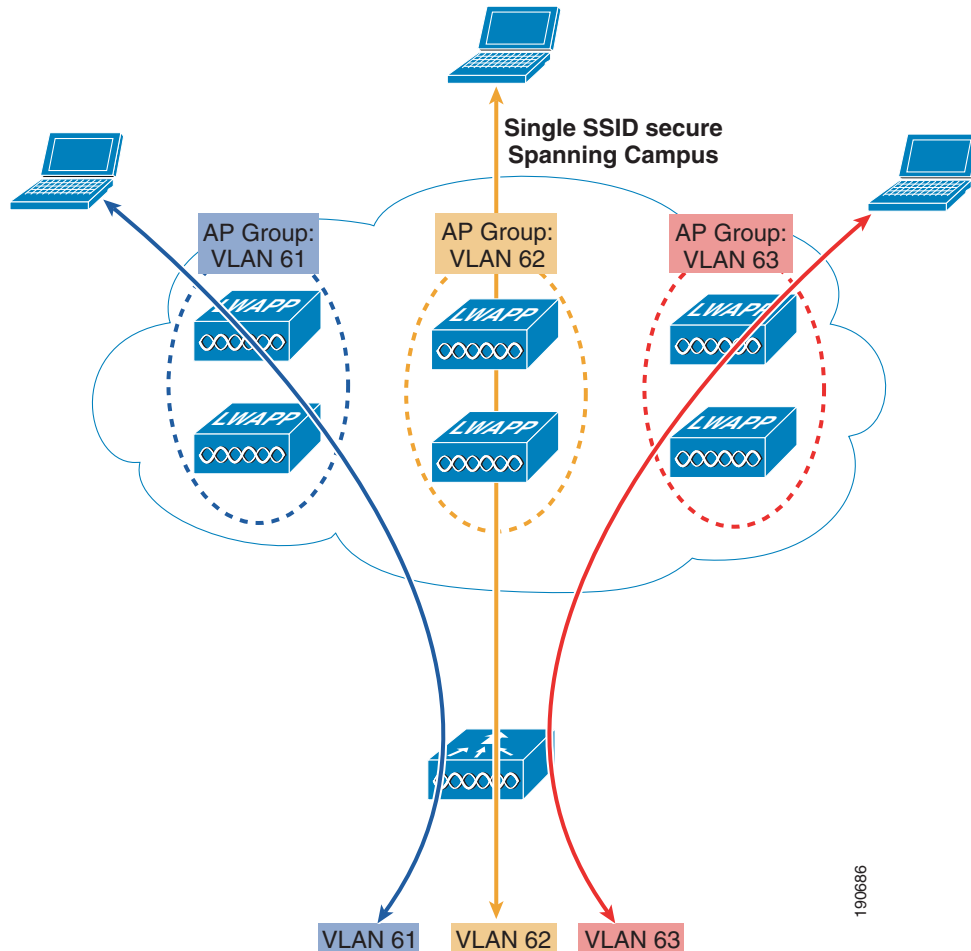
Additional information about WLAN controller mobility groups, including best practice information, can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL:  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2\\_Arch.html#wp1028143](http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2_Arch.html#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 student 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 5-10](#) shows the use of AP grouping based on site-specific VLANs.

Figure 5-10 Access Point (AP) Groups



As shown in [Figure 5-10](#), 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 (AP groups one, two, and three). Thus, a faculty member associating to the WLAN using an access point that is part of AP group one is assigned an IP address from the VLAN 61 IP subnet. Likewise, a faculty member associating to the WLAN using an access point that is part of AP group two 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 Community College reference design 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 Community College reference design 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 campuses.



### 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, Cisco prefers simplicity in the configuration of RF groups within the mobility design. Thus, all WLAN controllers in the Community College reference design 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 community college deployments there is no 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 5-40 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#wp102814](http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2_Arch.html#wp102814).

## Access Points

In the Cisco Community College reference design, it is anticipated that each campus 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 campus 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 campus users at any time. During your initial mobility design, these vacant areas may not be identified, so the precise subset of interior and exterior areas requiring WLAN access will likely be better determined during the site survey planning process that is 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 community colleges, the vast majority of interior building WLAN access can be provided by the Cisco Aironet 1140 Series 802.11n access point (see [Figure 5-11](#)), which delivers pervasive wireless connectivity while blending in seamlessly with the aesthetics of most modern campus learning environments.

**Figure 5-11** 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 classrooms, corridors, libraries, faculty offices, 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 learning 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 Community College reference design, the model of the Cisco 1140 Series 802.11n access point recommended for most interior campus 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 community college interior wireless access needs, there are some tradeoffs 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 providing external antenna connectors will be necessary. This can be a situation 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 laboratory 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 on the community college campus, the Cisco Aironet 1250 Series 802.11n access point is recommended (see [Figure 5-12](#)).

**Figure 5-12 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 expandability and customizability 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 Community College reference design, 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 community college installation cases, 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 tradeoffs 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, interior classrooms and lecture halls within community college campus 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 campus 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 community colleges, 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 for omni-directional style 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*

Globally disable any unnecessary low speed 802.11a/b/g 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/vowlan/41dg/vowlan41dg-book.html>

To provide outdoor WLAN access around the immediate perimeter area of each campus building, the Cisco Aironet 1520 Series Lightweight Outdoor Access Point is recommended (see [Figure 5-13](#)).

**Figure 5-13 Cisco Aironet 1520 Series Lightweight Outdoor Access Point**



As part of the Cisco Community College reference design, the Cisco Aironet 1520 Series Lightweight Outdoor Access Point provides an outdoor extension to the campus 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 campus infrastructure is beyond the scope of the Cisco Community College reference design 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 campus 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 campus, 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 campus 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 campus WLAN. In the event of a campus 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 on campus, public safety agencies may still be able to access the campus 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 campus WLANs using unlicensed 2.4 GHz and 5 GHz frequencies during crisis events, 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 tradeoff 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 Community College reference design 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 tradeoffs 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 Community College reference design 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, include 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 Community College reference design that are relevant to overall usability, such as the following:

- Quality-of-service (QoS)
- Guest access
- Traffic and performance

## 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 Community College reference design, 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 Community College reference design, 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.

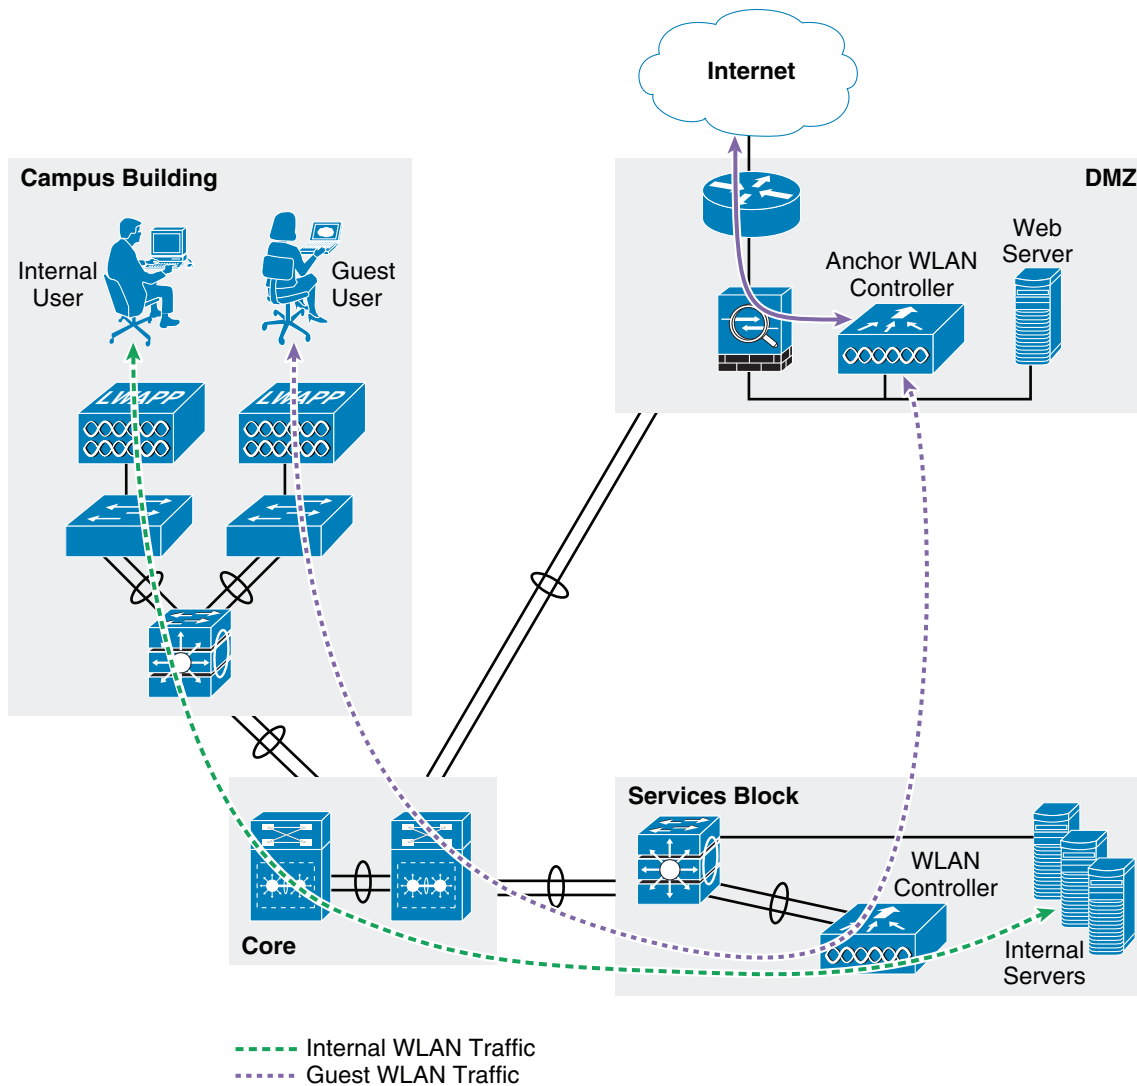
For more information on WLAN QoS, see the following URLs:

- *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 Community College reference design 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 5-14](#) shows a high-level view of guest access using this technique with a centralized WLAN controller design.

Figure 5-14 Guest Access Solution High-Level Overview



As shown in [Figure 5-14](#), a WLAN controller with a specific purpose is located in the main campus DMZ, where it is referred to as an *anchor controller*. The anchor controller is responsible for terminating EoIP tunnels originating from centralized campus 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 campus 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 campus 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 campus 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 Community College reference design, even in networks where hundreds or thousands of access points may be joined to other campus 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 campus controllers. This is done to keep the anchor controllers logically separate from the mobility groups associated with the general campus wireless deployment.
- The mobility group name for every campus 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 Community College reference design 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 community college has already deployed the Cisco NAC Appliance within their Internet DMZ to support wired guest access services. As shown in [Figure 5-15](#), 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 5-15 Cisco UWN Guest Access with Anchor WLC and NAC Appliance

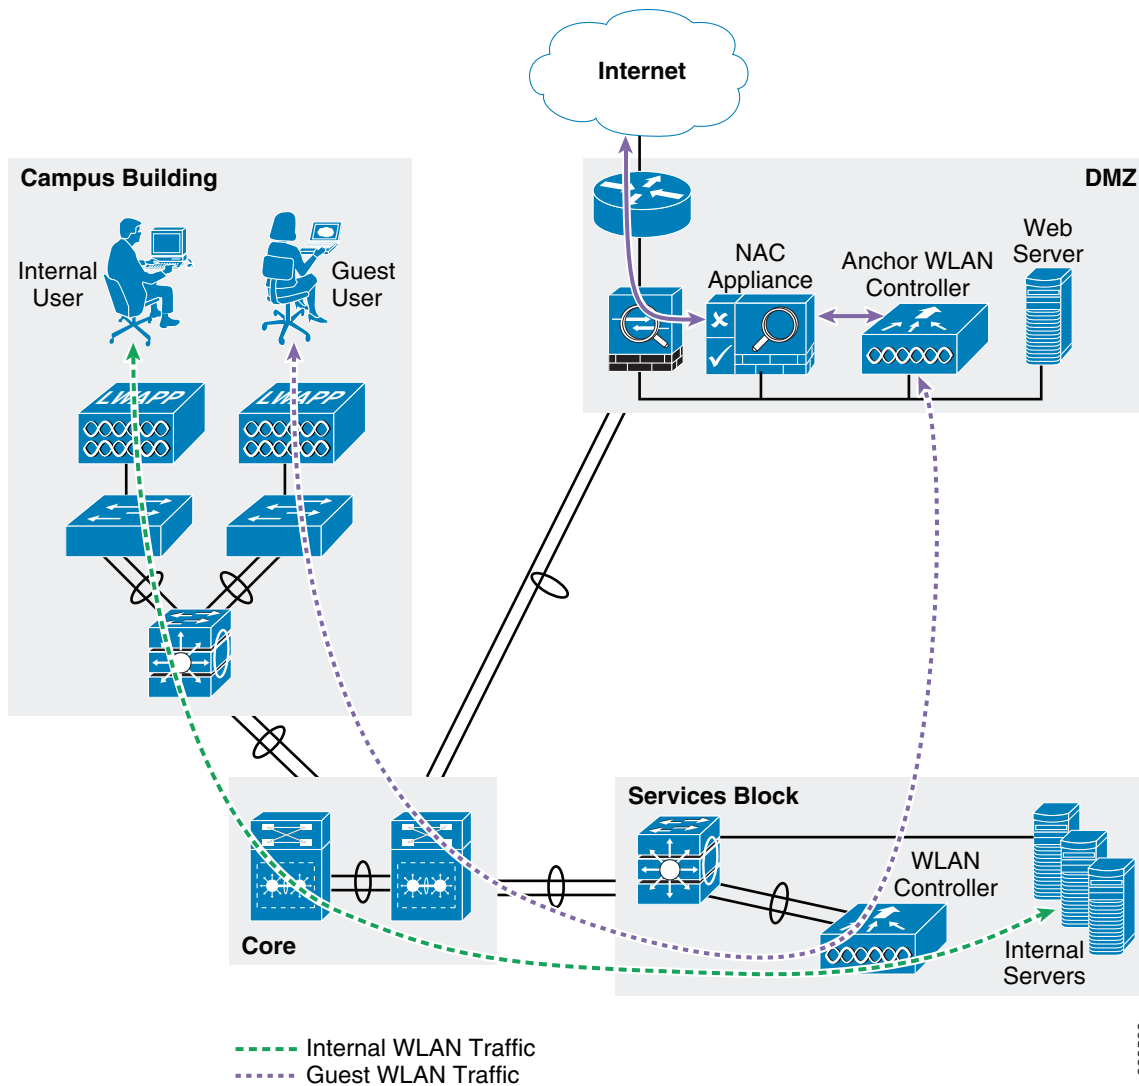


Figure 5-15 shows that the NAC Appliance is responsible for redirection, web authentication, and subsequent access to the Internet. The campus and anchor controllers are used only to tunnel guest WLAN traffic across the enterprise into the DMZ, where the NAC appliance is used to actually control guest access. The tradeoff here is the added cost of the external access control solution, versus the benefits it affords in relation to your particular deployment.

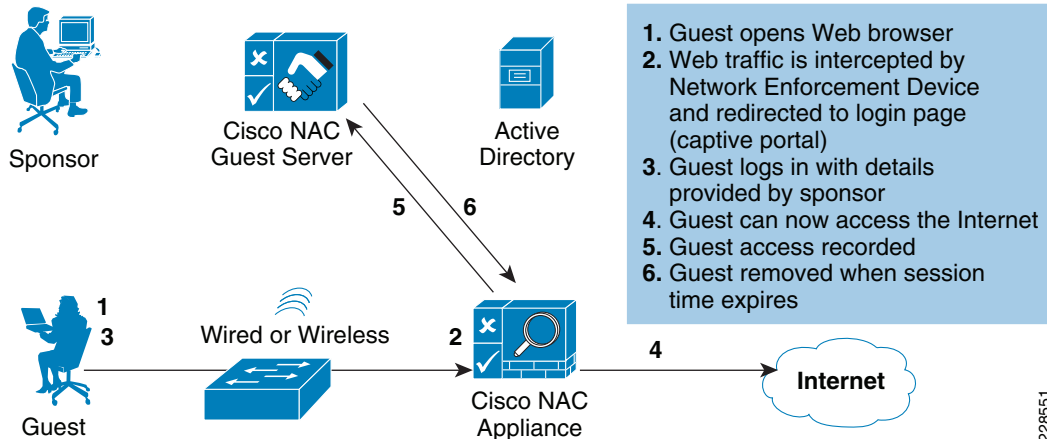
**Note**

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

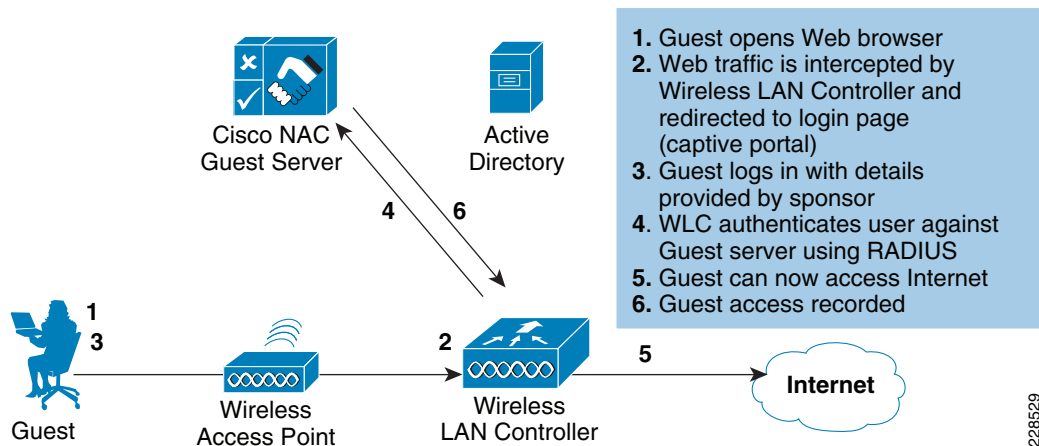
simple and secure manner. In addition, the whole process is recorded in a single place and stored for later reporting, including details of the network access activity. Cisco NAC Guest Server integrates with Cisco NAC Appliance through an application programming interface (API), allowing for guest accounts 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 5-16). In addition, the Guest Server receives accounting information from the NAC Appliance to enable full reporting.

**Figure 5-16** 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 5-17). In addition, the Guest Server receives accounting information from the WLAN controller to enable full reporting.

**Figure 5-17** 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\\_sheet0900aec806e98c9.html](http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5707/ps8418/ps6128/product_data_sheet0900aec806e98c9.html)

## Traffic and Performance

When designing mobility solutions incorporating the tunneling of CAPWAP traffic across campus infrastructure, questions often arise concerning the impact of such tunneling on network performance. In examining the impact of CAPWAP traffic in relation to overall network traffic volume, the following three points should be considered:

- *CAPWAP control traffic volume*—CAPWAP control traffic volume can vary considerably depending on the current activity state of the network. For example, this type of traffic volume usually reaches a zenith during a software upgrade or WLAN controller reboot. In most campuses, however, this degree of sporadic loading is considered negligible, and is of no consequence when considering the merits of a centralized deployment model over other options.
- *Tunneling overhead*—A Layer 3 CAPWAP tunnel adds a relatively negligible amount of overhead to a typical IP packet traversing to and from a WLAN client.

**Note**

---

A previous examination of the Light Weight Access Point Protocol (LWAPP), the predecessor to CAPWAP and similar in many ways, indicates that this overhead is approximately 44 bytes. With that said, traffic studies have concluded that the average load LWAPP control traffic places on the network is approximately 0.35 Kb/sec. Given that average packets sizes found on large scale network deployments are approximately 300 bytes, this represents an overhead of approximately 15 percent.

---

Once again, this is generally viewed as resulting in little to no consequence, especially in light of the considerable merits associated with a centralized deployment versus other options.

- *Traffic engineering*—WLAN traffic tunneled to a centralized controller is typically routed from the location of the WLAN controller to its final destination in the network. In the case of the Cisco Community College reference design, established best practices are followed concerning the placement of WLAN controllers within each per-campus centralized services block. With that said, the longer tunnels and traffic flows associated with a centralized deployment model can be mitigated by positioning the WLAN controllers in that part of the network where a large portion of the client traffic is already destined. In the Cisco Community College reference design, client-to-host/server traffic is typically destined for a local campus or main campus data center. This being the case, the overhead associated with any inefficiencies introduced because of centralized placement is not seen as adding significant delay or overhead.

## Manageability

As mentioned earlier, each WLAN controller in the Cisco Community College reference design 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 lifecycle mobility management solution that enables community college 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 5-18](#)).

**Figure 5-18 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 college network operation and improve the productivity of administrators, staff, and faculty. The comprehensive Cisco WCS platform scales to meet the needs of small, midsize, and large-scale WLANs across local and remote campuses. Cisco WCS gives college network administrators immediate access to the tools they need when they need them, wherever they may be located within the community college.

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 campus 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\\_sheet0900aecd802570d0.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html).

In the Cisco Community College reference design, a centralized WCS management server located in the data center block within the main campus is used. The data center block was initially shown in [Figure 5-3](#). [Figure 5-19](#) provides greater detail and magnification.

**Figure 5-19 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 Community College reference design are well served by a mobility design using a WCS management server located on the main campus.

**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 community college 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 campuses, as well as faculty administrators and staff. However, these groups will not all have equivalent resource and functionality access. It is anticipated that resource access will be limited further, based on administrative level and assigned campus or campuses. With few exceptions, it is not anticipated that most students will be required nor authorized to use the majority of services offered by WCS.

In this design, the ability to query and manage campus 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 campus central network administration staff may possess the authority to manage any and all mobility resources located on any campus throughout the college system, remote campus administrators may be limited by the following:

- *Campus resource management visibility policy*—This is performed by assigning the network mobility infrastructure components associated with each campus 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 campus, via their WCS management consoles. However, personnel responsible for local campus network administration are restricted to the discrete mobility infrastructure components associated with the virtual domain representing their local campus. These infrastructure components include WLAN controllers, access points, configuration templates, WCS events, reports, alarms, WLAN clients, and so on.
- *Campus 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 campus 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 Community College reference design 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 tradeoffs exist, such as the following:

- *WCS server failure*

In the Cisco Community College reference design, the centralized mobility network management services provided by WCS are not regarded as being mission-critical for the majority of community college deployments. Thus, in the rare event of a WCS server failure, and given the cost constraints of most community college environments, it is assumed that direct WLAN controller management workarounds (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 Community College reference design 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 community college 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 campuses. 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-campus basis. However, this increases the cost of WCS deployment significantly, requiring one WCS management server per campus, and preferably a Cisco WCS Navigator management aggregation platform located at the main campus 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 community colleges, these decentralized management options are not included in the Cisco Community College reference design 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

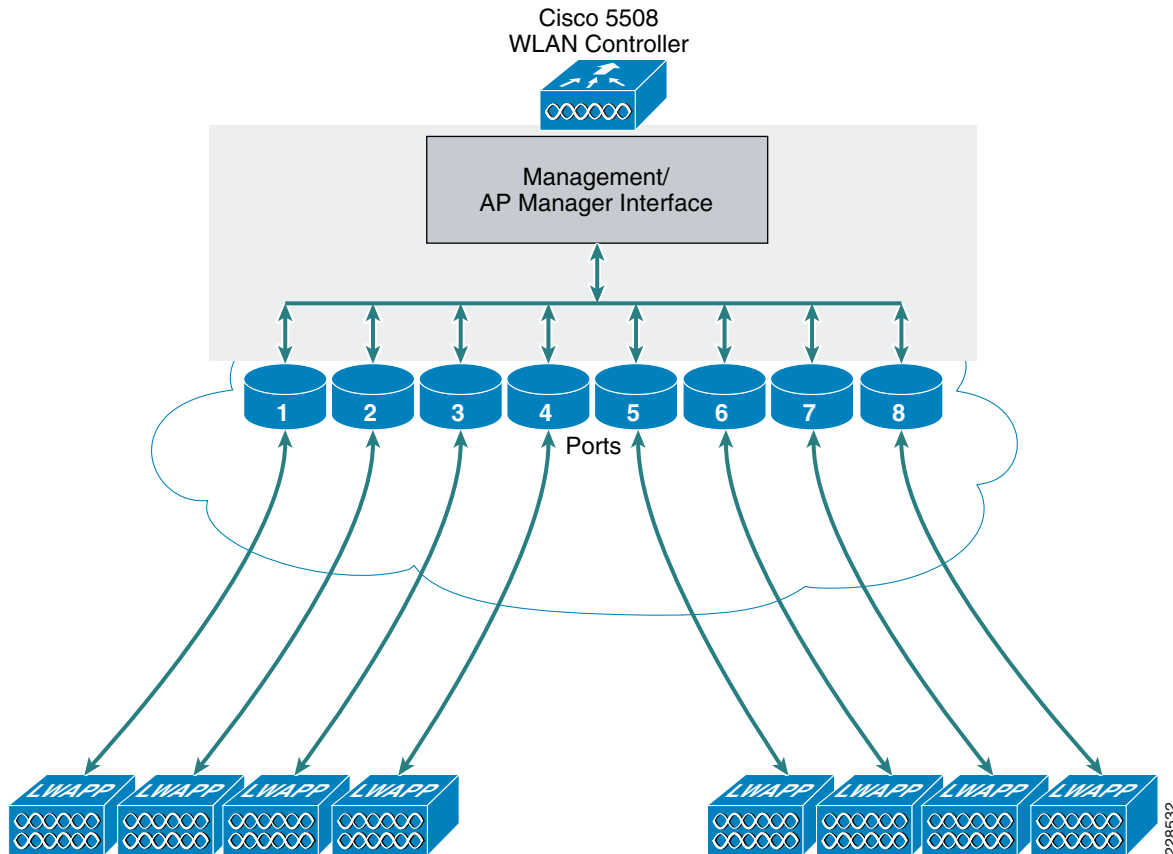
This section discusses the mobility design considerations pertaining to those aspects of the Cisco Community College reference design relevant to overall reliability, and includes the following:

- 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 Community College reference design, 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 5-20](#). This can be especially important with high capacity controllers licensed for many access points, such as the Cisco CT5508-250.

Figure 5-20 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 campus 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 multilayer campus 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 5-1](#) provides information for the Cisco 5508 Wireless Controller deployed with its maximum complement of 250 access points.

**Table 5-1 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 5-1](#) shows that even if designing for peak 802.11n throughput of 250 Mbps per access point, oversubscription is not expected to exceed campus 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:

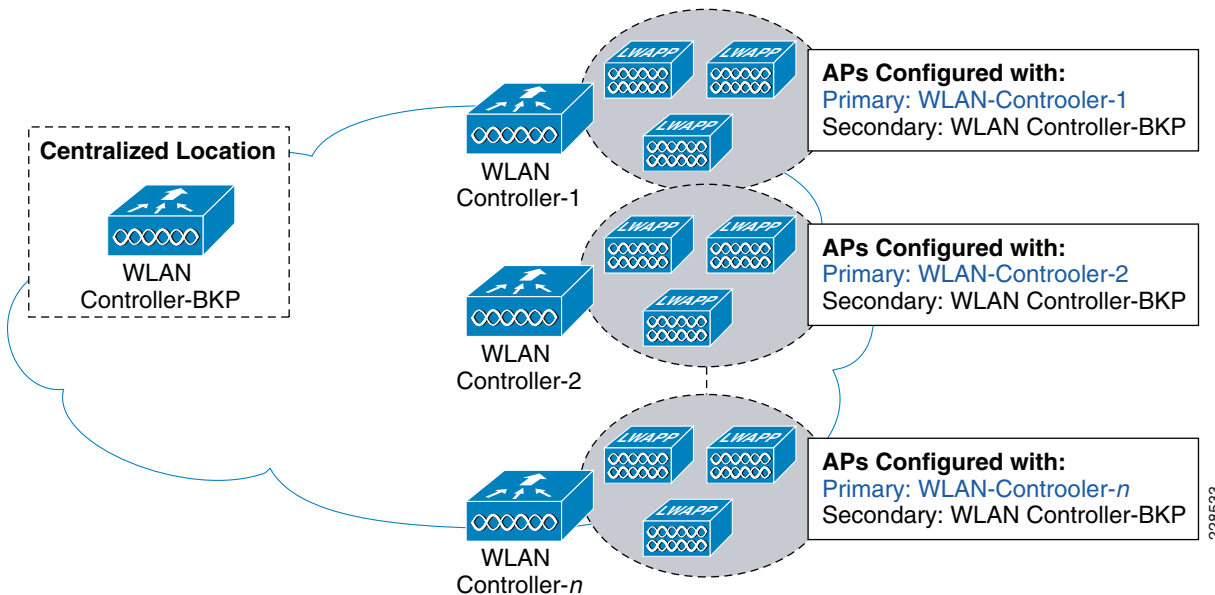
<http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html#wp1062211>.

## 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 Community College reference design, 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 5-21 provides a general illustration of the principle of N+1 controller redundancy.

**Figure 5-21** 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 5-21](#) by having the access points of the main campus as well as all remote campuses address a common redundant controller located in the main campus services block. In other cases, there may be sufficient capacity on the primary controllers located on the main campus themselves to accommodate the access point and user load of a single failed controller on any of the remote campuses. This approach requires that main campus controllers be licensed for a greater number of access points than necessary for the support of the main campus alone. Additional licensing of existing controllers is performed in place of providing a dedicated additional controller platform at the main campus for system-wide redundancy. In this case, the available capacity of the primary main campus WLAN controllers allow them to act as the secondary destination for the access associated with the largest remote campus. Thus, in this particular case, the need to deploy hardware at the main campus site explicitly for the purposes of controller redundancy may be avoided.

For example, assume that the main campus shown in [Figure 5-3](#) contains a total of 250 combined access points across all main campus buildings, and the largest of the remote campuses also contains 250 combined access points across all remote campus buildings. In this case, if the main campus 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 campus alone can be split equally between the two controllers (125 access points on each controller). This leaves ample capacity in the main campus for one of the following scenarios to occur:

- Either of the main campus controllers may fail and allow up to 125 joined access points to migrate (failover) to the other controller in the pair. This results in the remaining functional controller bearing the full load of 250 access points.
- Any remote campus controller may fail and allow its joined access points to migrate (failover) to the main campus controllers. In the case of a failure of the largest remote campus, this results in each of the main campus 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 Community College reference design, a high-speed WAN/MAN is present between campuses, thus making access point failover to a remote WLAN controller feasible, as described in the previous section. To accomplish this in the Cisco Community College reference design, access points can be configured to failover 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 reconfigured in the event of an access point failover to the main campus 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 reconfigured in the event of an access point failover to a central backup controller. The SSID used to access a particular WLAN throughout the multi-campus community college 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>.

# Community College Mission Relevancy

This document attempts to present the mobility design considerations that comprise an important part of a successful implementation of the Cisco Community College reference design. The goal is to provide stakeholders with a reference design that assists in solving the complex business challenges that community colleges must face in the 21st century.

This closing section steps back from the technical intricacies of system design to examine how these design considerations relate to the foundation services described in the opening paragraphs of this document.

## Safety and Security

The mission of the Cisco Community College reference design in this area is to enhance safety and security on campus by using a design model that proactively protects students, faculty, and staff. Maintaining safe buildings and grounds while keeping the network secure for today's community colleges. The Cisco Community College reference design helps to facilitate and enhance the effectiveness of physical campus security, track assets, protect the network, and prevent unauthorized network access.

The mobility aspects of Cisco Safety and Security Solutions includes the following three solution sets:

- *Campus physical safety and security*—Is the physical campus protected and safe? The Cisco Community College reference design helps enable community colleges to maintain safe buildings and grounds by the following:
  - Supporting the monitoring of unauthorized behavior and delivering alerts about detected events. Real-time monitoring helps campus security staff to prevent, deter, detect, and respond more quickly to incidents.
  - Providing reliable, secure, and high-performance WLAN communications throughout building interiors and outside buildings to students, faculty, administrators, and community college guests. This level of reliable wireless connectivity can be the key to ensuring rapid notification of campus personnel in the event of a safety incident.
  - Real-time tie-in to wired and wireless video surveillance systems as well as portable security devices and third-party campus safety systems, to ensure that unfolding events are detected quickly and monitored by the right personnel in the right location.
  - Offering WLAN connectivity from strategic campus locations to public safety emergency professionals using licensed 4.9 GHz frequencies as well as traditional 2.4 GHz and 5 GHz unlicensed frequencies. During periods of crisis, indoor and outdoor WLANs can provide first responders with vital tactical information about what is happening within the campus. The 4.9 GHz band that is available on the Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point provides access via radio frequencies that are reserved by the FCC for public safety usage only.

Other Cisco products and solutions that work in collaboration with the Cisco Community College reference design to enable these and other capabilities include the Cisco Mobility Services Engine (MSE), Cisco Context-Aware Mobility Solution, Cisco Unified Communications, the Cisco Unified Wireless IP Phone 7925G, and Cisco Video Surveillance products.

- *Network and data security*—Is the wireless network secure? The Cisco Community College reference design addresses this issue by the following:

- Protecting confidential data and transmissions by using the highest level of authentication and encryption applicable to the tasks at hand, helping to ensure that wireless transmissions remain secure and protected.
- Helping to prevent misguided students or malicious intruders from hacking into restricted servers or issuing attacks against the wireless network via the inclusion of the optional Mobility Services Engine with Wireless Intrusion Protection System (wIPS). It also helps quickly locate rogue access points anywhere on campus.
- Providing an economical guest access solution that furnishes safe and secure guest access for campus guests.
- *Context-aware mobility*—Where is an asset located on campus and what is its status? The Cisco Unified Wireless Network, in conjunction with Cisco Context-Aware Mobility solutions, supports the ability to do the following:
  - Capture and integrate into community college application and administrative processes, detailed contextual information about an asset such as its location, movement, status, and state. This solution helps community colleges automatically collect information about mobile assets, analyze it, and use it to reduce errors, improve asset security, prevent delays, improve scalability beyond manual processes, and enhance learning functions.
    - Any asset that is emitting a Wi-Fi signal can be monitored, tracked, and found with this solution. A Wi-Fi signal can be generated from a built-in wireless card or an attached Wi-Fi tag from third-party vendors including AeroScout, WhereNet, and others.
    - Expensive items such as projectors, televisions, portable plants, lab equipment, tools, laptops, or any asset that moves can be easily tracked.
  - Alerts can be issued about the movement of a device in or out of an area. Costs for misplaced items, loss, and theft can be reduced.
  - Faculty and staff can use context-aware mobility in conjunction with third-party applications to automatically send announcements, assignments, room change notifications, campus event updates, and emergency alerts to students as they roam on campus.
  - Security personnel can use this solution to receive silent alerts and notifications about asset movement and rogue devices, track the areas of the campus they have inspected or secured, and quickly learn the location and of emergency-triggered events.
  - Administrators can use this solution to quickly locate students, faculty, or staff anywhere on campus.

## Virtual Learning

The traditional scenario of a mass of students filing into a large lecture hall within a large, monolithic campus building is by no means the only such model available to today's modern-day community college student. High performance, secure wireless technologies can enable “virtual classrooms” even in non-traditional settings, such as leased space in shopping malls, retail plazas and even from homes and offices.

School administrators need secure access to tools, records, and resources, as well as ubiquitous access to mobile voice capabilities throughout the campus.

Using the solutions and technologies presented within the Cisco Community College reference design, state-of-the art instructional sites can be deployed in such non-traditional settings within urban, suburban, and rural venues. These types of facilities can help bring much-needed skills to areas that may not be within convenient reach of conventional community college campuses. For example, a community college location at a shopping mall may operate as a science, technology, engineering, and math learning

center. Such centers may range in size from one or two classroom sites to larger-scale deployments with ten or twelve classrooms, a hundred or more student computers, a science lab, two auditoriums, and even testing, conference, and office space.

## Secure Connected Classrooms

Providing connectivity to students while attending class is the foundation of twenty-first century learning. However, it also presents several challenges for community colleges. For example, the density of wireless users in one location can be problematic. Wireless designs must take into consideration the number of users, radio interference, and network utilization.

The Cisco Community College reference design addresses these challenges in a variety of ways, including the following:

- High-performance dual-band access points that provide options to migrate users to 802.11n and better performing bands (5 GHz) that offer increased data rates with less interference.
- Advanced radio resource management algorithms and techniques that can automate the fine-tuning of transmit power and other parameters to best accommodate high-density user populations.
- Comprehensive wireless network management systems (WCS) that can assist in the identification of interference sources and rogue access points, including their location.
- Detailed reporting mechanisms that can enable administrators to better understand the points of congestion in the network and how best to address them.
- A high-performance controller platform, optimized for use with high-performance 802.11n access points, that offers aggregate wired interface bandwidth of up to 8 Gbps.

## Operational Efficiencies

Delivering quick and cost-effective broadband access anywhere on campus extends learning beyond the classroom and improves campus operations, collaboration, and productivity. The Cisco Community College reference design supports secure, easy wireless network access to voice, video, and data applications for students, administrators, faculty, staff, and visitors as they roam about the campus.

The operational efficiencies enabled by the Cisco Community College reference design encompass the following solution sets:

- *Pervasive wireless on campus*—Is the WLAN available ubiquitously in all required indoor and outdoor areas? As a key component of the Cisco Community College reference design, the Cisco Unified Wireless Network delivers broadband access quickly and cost-effectively to all the required indoor and outdoor areas in the typical community college. The benefits of this are as follows:
  - When wireless access is available pervasively on campus, users do not need to hunt for wired ports because they can gain access to network resources using their wireless connection.
  - Users can stay connected to their applications as they roam, without having to re-log onto the network while they are in motion.
  - As long as an area is covered by the wireless infrastructure, faculty, students, and guests can work, share resources, collaborate, and communicate.
  - With a pervasive wireless network, instruction is no longer limited to the classroom.
  - Faculty can teach inside or outside the classroom, accessing the Internet and applications while on the move.

- Access to resources is improved because faculty and administrators do not have to return to their desk to perform online administration tasks, access research information, or check E-mail.
- Student satisfaction is increased and trouble calls are decreased because wireless access is predictable and consistent.
- With a pervasive Cisco WLAN, community colleges can deliver network access to locations where hardwiring is too expensive, too difficult, or implausible. Examples are refurbished buildings, older buildings with environmental concerns such as asbestos remediation, or sites with protected-building restrictions such as historical landmarks.
- Costs for cabling temporary spaces or for providing network access to new faculty or staff can be reduced or eliminated.

In fact, you may find that it is more cost-effective to provide wireless network access pervasively on campus than it is to install individual wired ports over the same geographic area.

- *High-speed wireless access*—Are bandwidth-intensive applications supported on the WLAN? The Cisco Unified Wireless Network facilitates the creation of solutions that accelerate the delivery of bandwidth-intensive applications and provides a better end-user experience.
  - The Cisco high-speed wireless network, based on the 802.11n standard, delivers unprecedented reliability, greater performance, and extended reach for pervasive wireless connectivity. It excels at supporting bandwidth-intensive applications that are used for research, learning, virtual environments, and social networking. This solution also delivers predictable and continuous WLAN coverage for areas with dense wireless usage such as lecture halls, auditoriums, open spaces, and social areas.
  - Community colleges that deploy 802.11n are demonstrating a commitment to technology innovation and leadership. They are building a solid technology foundation to attract new students and remain competitive in the ever-evolving global community college education marketplace.
- *Secure guest access*—Can visitors easily access the network? The Cisco Community College reference design supports secure wireless guest access that cost-effectively simplifies the process of providing temporary Internet access to visitors such as prospective students, alumni, parents, visiting lecturers, and temporary personnel. Wireless guest access eliminates the frustration that visitors experience when they are limited to wired-only ports in small areas on campus. It also eliminates the costs that community colleges might incur from wiring and maintaining wired ports to accommodate visitors. With the Cisco secure guest access solution, community colleges can do the following:
  - Enhance the community college experience for prospective students
  - Provide Internet access to guests attending campus events
  - Easily support network access for conference attendees and guest lecturers
- *Campus automation*—Are managing and tracking campus resources automated? The solutions enabled by the Cisco Community College reference design can help community colleges reduce costs by supporting Wi-Fi-enabled services that automatically manage, track, and maintain campus resources and assets.
  - The wireless network can assist with better management of real estate components to support green initiatives, improve energy efficiency, and create smart buildings.
  - Alarms, bells, and clocks can be wirelessly enabled to reduce the labor costs associated with managing them. Wi-Fi tags can be placed on assets to automatically track their movement and help reduce costs for misplaced items, loss, and theft.



- *Facilities management*—The Cisco Community College reference design includes adaptive power management capabilities that are built into the Cisco Unified Wireless Network through its Cisco Wireless Control System (WCS) management platform and software release 6.0. Cisco WCS adaptive power management allows community colleges to shrink their carbon footprint immediately through measurable reductions in energy usage and operational expenses.

By using Cisco WCS adaptive power management to turn access point radios on or off at scheduled intervals (hour, day, and week), power requirements and operating expenses can be reduced almost immediately. The power savings gained vary based on the Cisco Aironet access point model deployed. Using this feature can help organizations create a sustainable culture and gain momentum for “Green IT” initiatives.

## Wireless LAN Controller Configuration

The core component of the Cisco Unified Wireless architecture is the Wireless LAN Controller (WLC) that provides the interface between the “split-MAC” wireless network and the wired network. That is, the WLC is the Layer-2 connection point between WLAN client traffic and the wired network, making the WLC an aggregation and control point for WLAN traffic. In addition, the WLC is the primary control point of AP and RF management.


The reference design used for testing WLAN services in the CCVE network uses four WLCs:

- Two WLCs (cr23-5508-1, cr23-5508-2) for the main campus
- One WLC (cr14-5508-1) for a remote campus
- One anchor WLC (cr11-5508-wlc) for guest services

## WLC and Wired Network Connections

The WLCs in the main campus are centralized for that campus and connected to a 3750E stack in services block connected to the campus core, as shown in [Figure 5-22](#). These WLCs provide WLAN services for the entire campus, as well as failover support for APs in remote campuses, in the event of WLC outage at that location. The number of WLCs for the main campus is driven by the number of APs deployed and the type of failover support required. In this example, two WLCs are used to illustrate the basic configuration requirements.

**Figure 5-22 Services Block WLC Connection**



The two main campus WLCs share the VLAN and subnet configuration, differing only in their IP addressing. [Figure 5-23](#) and [Figure 5-24](#) show the interface summary on the two WLCs. 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 WLC, 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 5-23 cr23-5508-1 Interfaces**

Interface Name	VLAN Identifier	IP Address	Interface Type	Dynamic AP Management
faculty & staff voice	112	10.125.30.18	Dynamic	Disabled
faculty_and_staff_data	113	10.125.30.34	Dynamic	Disabled
management	111	10.125.30.2	Static	Enabled
nac	117	10.125.30.99	Dynamic	Disabled
service-port	N/A	172.26.158.243	Static	Not Supported
student_access	114	10.125.30.50	Dynamic	Disabled
virtual	N/A	1.1.1.1	Static	Not Supported

**Figure 5-24 cr23-5508-2 Interfaces**

Interface Name	VLAN Identifier	IP Address	Interface Type	Dynamic AP Management
faculty & staff voice	112	10.125.30.19	Dynamic	Disabled
faculty_and_staff_data	113	10.125.30.35	Dynamic	Disabled
management	111	10.125.30.3	Static	Enabled
nac	117	10.125.30.100	Dynamic	Disabled
service-port	N/A	172.26.158.244	Static	Not Supported
student_access	114	10.125.30.51	Dynamic	Disabled
virtual	N/A	1.1.1.1	Static	Not Supported

Figure 5-25 shows the management interface of WLC cr23-5508-1. Note that Link Aggregation (LAG) is enabled on all the WLCs used in the CCVE design.

Figure 5-25 cr23-5508-1 Management Interface

The screenshot shows the Cisco WLC Management Interface for controller cr23-5508-1. The navigation menu on the left includes: Controller, General, Inventory, Interfaces, Multicast, Network Routes, Internal DHCP Server, Mobility Management, Ports, NTP, CDP, and Advanced. The main content area is titled 'Interfaces > Edit' and contains the following configuration sections:

- 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:

Navigation buttons '< Back' and 'Apply' are visible at the top right of the configuration area.

Example 5-1 and Example 5-2 show examples of the switch configuration for the 3750 stack switch connecting the main WLCs to the wired network.

#### Example 5-1 Example WLC 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 5-2 Example WLC 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
```

229/001

```

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 Campus

The remote campus WLC and wired network connection is the same as that used in the main campus. In other words, WLC is connected to a 3750E stack that acts as a services block for the remote campus. 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 WLCs. This helps to ensure seamless mobility when clients roam between APs that are connected to WLCs within the same Mobility Group. The default Mobility Group Name is created in the Controller General configuration page, as shown in [Figure 5-26](#).

**Figure 5-26** cr23-5508-1 Mobility Group Definition

The screenshot shows the Cisco Unified Wireless Network (CUWN) configuration interface for Controller cr23-5508-1. The 'General' configuration tab is active, and the 'Default Mobility Domain Name' field is highlighted with a red circle, showing the value 'MAIN'. Other configuration options 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 (CCVE), and User Idle Timeout (seconds) (300). The 'Multicast Group Address' is set to 239.255.1.57. The interface includes a navigation menu on the left with options like General, Inventory, Interfaces, Multicast, Network Routes, Internal DHCP Server, Mobility Management, Ports, NTP, CDP, and Advanced. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh.

229002

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 5-27 and Figure 5-28 show the Mobility Group membership information for both main campus WLCs. It can be seen that the Mobility Group membership has two main members for the two WLCs that are providing WLAN access within the main campus. These WLCs 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 campus WLC Mobility Group membership configuration uses a different mobility group name, and does not include either of the main campus WLCs. The reason for it not including either of the main campus WLCs is because it is not expecting to support seamless roaming between the remote campus and main campus. 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 campus, the GUEST\_ACCESS mobility group-member information also appears on the remote campus WLC.

Figure 5-27 cr23-5508-1 Mobility Group 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

229003

Figure 5-28 cr23-5508-2 Mobility Group Members

Local Mobility Group	MAIN			
MAC Address	IP Address	Group Name	Multicast IP	Status
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

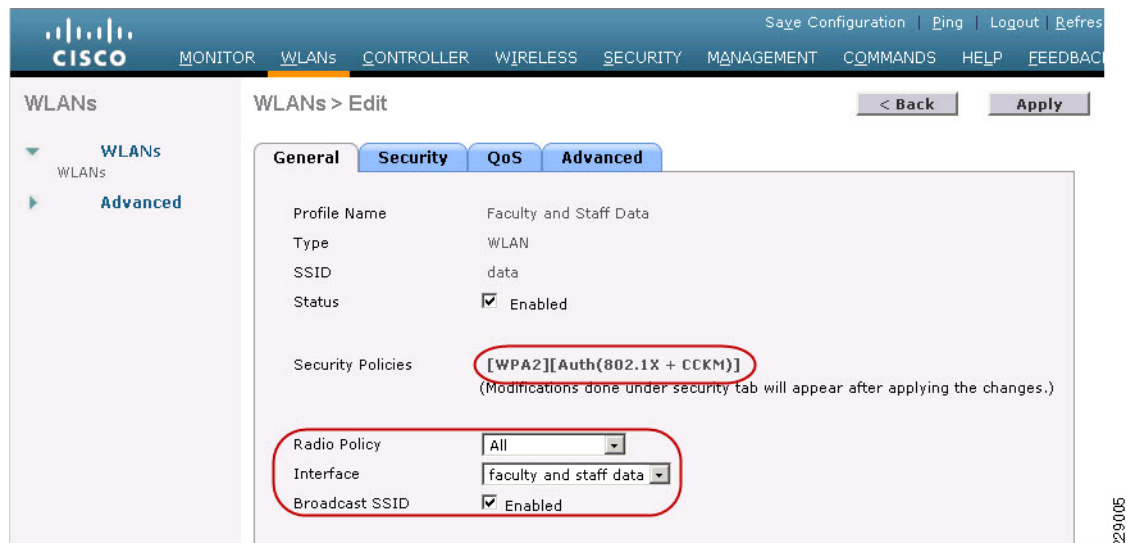
229004

## WLAN Configuration

### Faculty and Staff Data WLAN

Figure 5-29 shows the general WLAN configuration for the faculty and 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 WLC 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 5-29 Faculty and Staff Data WLAN



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.

### Faculty and Staff Voice WLAN

Figure 5-30 shows the general WLAN configuration for the Faculty and Staff Voice WLAN network. The key point to note on this tab is the security policy that has been set under the security tab, and the WLC interface that the WLAN has been mapped to. The security configuration recommended is to use 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 “Faculty and Staff Data WLAN” section on page 5-52.

Figure 5-30 Faculty and Staff Voice WLAN

The screenshot shows the Cisco Wireless LAN Controller configuration interface for the 'Faculty and Staff Voice WLAN'. The 'Security' tab is active, displaying the following configuration details:

- Profile Name: Faculty and 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: faculty & staff voice
- Broadcast SSID:  Enabled

Navigation options include '< Back' and 'Apply' buttons. The interface also shows a sidebar with 'WLANs' and 'Advanced' tabs, and a top navigation bar with 'MONITOR', 'WLANs', 'CONTROLLER', 'WIRELESS', 'SECURITY', 'MANAGEMENT', 'COMMANDS', 'HELP', and 'FEEDBACK'.

229006

## Student WLAN

Figure 5-31 shows the Student Access WLAN. This WLAN is configured for open authentication, which allows students to join the WLAN without providing security credentials. However, authentication and posture assessment of all devices using the Student WLAN is performed by the NAC system that is Layer-2 adjacent to the WLC.

Figure 5-31 Student WLAN

The screenshot shows the Cisco Wireless LAN Controller configuration interface for the 'Student Access WLAN'. The 'Security' tab is active, displaying the following configuration details:

- Profile Name: Student Access
- Type: WLAN
- SSID: student
- Status:  Enabled
- Security Policies: **None** (Modifications done under security tab will appear after applying the changes.)
- Radio Policy: All
- Interface: student access
- Broadcast SSID:  Enabled

Navigation options include '< Back' and 'Apply' buttons. The interface also shows a sidebar with 'WLANs' and 'Advanced' tabs, and a top navigation bar with 'MONITOR', 'WLANs', 'CONTROLLER', 'WIRELESS', 'SECURITY', 'MANAGEMENT', 'COMMANDS', 'HELP', and 'FEEDBACK'.

229007

To enable NAC on the WLAN, the NAC state option under **WLAN > Advanced** must be selected (see [Figure 5-32](#)).

**Figure 5-32 Student WLAN Advanced Options**

The screenshot shows the Cisco Wireless LAN Controller configuration page for a WLAN. The 'Advanced' tab is selected, and the 'NAC' section is highlighted with a red circle, indicating that the 'State' is set to 'Enabled'. Other sections visible include 'DHCP' (with 'DHCP Addr. Assignment' set to 'Required'), 'Management Frame Protection (MFP)', and 'DTIM Period (in beacon intervals)'. The 'Client Exclusion' section is also visible with 'Enabled' checked and a timeout value of 60 seconds.

Section	Option	Value
General	Allow AAA Override	<input type="checkbox"/> Enabled
	Coverage Hole Detection	<input checked="" type="checkbox"/> Enabled
	Enable Session Timeout	<input checked="" type="checkbox"/> 1800 (Session Timeout (secs))
	Aironet IE	<input checked="" type="checkbox"/> Enabled
	Diagnostic Channel	<input type="checkbox"/> Enabled
	IPv6 Enable	<input type="checkbox"/>
	Override Interface ACL	None
	P2P Blocking Action	Disabled
	Client Exclusion	<input checked="" type="checkbox"/> Enabled 60 (Timeout Value (secs))
	DHCP	DHCP Server
DHCP Addr. Assignment		<input checked="" type="checkbox"/> Required
Management Frame Protection (MFP)	Infrastructure MFP Protection	<input checked="" type="checkbox"/> (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	<input checked="" type="checkbox"/> Enabled

A quarantine VLAN must also be configured on the WLAN interface, under the **Controller > Interfaces > Edit** menu, as shown in [Figure 5-33](#).



Figure 5-33 Student Access Interface with Quarantine VLAN

The screenshot shows the Cisco Wireless LAN Controller configuration page for the 'student access' interface. The interface is configured with the following settings:

Section	Field	Value
General Information	Interface Name	student access
	MAC Address	00:24:97:cf:3f:af
Configuration	Quarantine	<input checked="" type="checkbox"/>
	Quarantine Vlan Id	314
Interface Address	VLAN Identifier	114
	IP Address	10.125.30.50
	Netmask	255.255.255.240
	Gateway	10.125.30.49

For a complete description of NAC and NAC WLAN integration, refer to [Chapter 6, “Community College Security Design.”](#)

## Guest Access WLAN

The Guest Access WLAN configuration has much in common with the Student Access WLAN configuration, with the major differences being only how the traffic is handled once it leaves the WLAN. [Figure 5-34](#) shows the Guest WLAN using the same security configuration as the Student WLAN, but its interface configuration is significantly different. 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 WLC that is located in the DMZ. The WLAN client traffic from the Guest WLAN is tunneled by the WLC to the anchor WLAN.

229009

Figure 5-34 Guest WLAN

The screenshot shows the Cisco WLC configuration interface for a WLAN named 'Guest\_Access'. The 'Security' tab is active, displaying the following configuration:

- Profile Name: Guest\_Access
- Type: WLAN
- SSID: Guest\_Access
- Status:  Enabled
- Security Policies: **Web-Auth** (circled in red)
- Radio Policy: All (dropdown menu)
- Interface: management (dropdown menu)
- Broadcast SSID:  Enabled

Buttons for '< Back' and 'Apply' are visible at the top right. The Cisco logo and navigation tabs (MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, HELP, FEEDBACK) are at the top.

229010

Figure 5-35 and Figure 5-36 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 WLC that is located in the DMZ.

Figure 5-35 WLAN Mobility Anchor Selection

The screenshot shows the 'WLANs' configuration page with a table of WLANs. The 'Web-Auth' entry is selected, and the 'Remove Mobility Anchors' option is circled in red.

WLAN ID	Type	Profile Name	WLAN SSID	Admin Status	Security Policies
1	WLAN	Faculty and Staff VoWLAN	vowlan	Enabled	[WPA][Auth(CCKM)]
2	WLAN	Faculty and Staff Data	data	Enabled	[WPA2][Auth(802.1X + CCKM)]
3	WLAN	Student Access	student	Disabled	Web-Passthrou
5	WLAN	Guest_Access	Guest_Access	Enabled	Web-Auth

Buttons for 'Remove' and 'Mobility Anchors' are visible next to the selected entry. The Cisco logo and navigation tabs are at the top.

229011

Figure 5-36 Mobility Anchor Selection

The screenshot shows the 'Mobility Anchors' configuration page. The 'Switch IP Address (Anchor)' field is circled in red, showing the value 10.125.32.34. The 'Data Path' and 'Control Path' are both set to 'up'.

Buttons for 'Mobility Anchor Create' and 'Switch IP Address (Anchor)' are visible. The Cisco logo and navigation tabs are at the top.

229012

Figure 5-37 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 5-37 and Figure 5-38.

**Figure 5-37** Anchor Guest WLAN

The screenshot shows the Cisco WLC configuration interface for the 'Guest\_Access' WLAN. The 'Security' tab is active, displaying the following configuration:

- Profile Name: Guest\_Access
- Type: WLAN
- SSID: Guest\_Access
- Status:  Enabled
- Security Policies: **Web-Auth** (circled in red)
- (Modifications done under security tab will appear after applying the changes.)
- Radio Policy: All
- Interface: **guest-vlan** (circled in red)
- Broadcast SSID:  Enabled

229013

The WLAN on the DMZ anchor WLC 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 5-39.

**Figure 5-38** Anchor WLC interfaces

The screenshot shows the Cisco WLC configuration interface for the 'Interfaces' section. The following table lists the configured interfaces:

Interface Name	VLAN Identifier	IP Address	Interface Type	Dynamic AP Management
<b>quest-vlan</b> (circled in red)	104	10.125.32.66	Dynamic	<b>Disabled</b> (circled in red)
management	102	10.125.32.34	Static	Enabled
service-port	N/A	172.26.137.82	Static	Not Supported
virtual	N/A	1.1.1.1	Static	Not Supported

229014

Figure 5-39 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 WLC 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 WLCs and not on the WLCs as well as on the connected switch). 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 5-40). This procedure configures the WLC to set the 802.1p marking of the frames sent from the WLC 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 WLC 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 WLC 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 5-40 Wired QoS Protocol Configuration

The screenshot shows the Cisco configuration interface for 'Edit QoS Profile'. The profile name is 'platinum' and the description is 'For Voice Applications'. Under 'Per-User Bandwidth Contracts (k)\*', all values are set to 0. Under 'Over the Air QoS', Maximum RF usage per AP (%) and Queue Depth are both set to 100. The 'Wired QoS Protocol' section is highlighted with a red box, showing 'Protocol Type' set to '802.1p' and '802.1p Tag' set to '6'. A note at the bottom 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 WLC they ultimately join. Therefore, establishing the connection between APs and the correct WLC is a key component of the design.

The CUWN provides many different options to allow APs to discover the correct WLC (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 WLC 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 WLC that has been defined using a discovery mechanism, it learns about all of the WLCs in the default mobility group of the discovered WLC. An AP can be configured for preferred primary, secondary, and tertiary WLCs within that mobility group. Figure 5-41 shows an example of this where the AP is configured with its preferred WLC (primary controller), and its preferred failover WLC (secondary controller).

Figure 5-41 AP Controller Preferences

	Name	Management IP Address
Primary Controller	cr23-5508-1	
Secondary Controller	cr23-5508-2	
Tertiary Controller		

229017

The configuration of access point WLC preference will determine the failover models for the WLAN deployment. For example, all the APs on the campus could be configured to prefer one WLC as primary, with the other WLC used solely as a back-up controller. An alternative configuration would be to spread the AP load across both WLCs, 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 failover to a WLC outside of its primary WLCs mobility group, the AP must be configured with the IP address and name of that failover WLC, rather than just the WLC name. An example of this configuration, from the remote campus, is shown in Figure 5-42.

Figure 5-42 AP Failover to a WLC Outside the Mobility Group

	Name	Management IP Address
Primary Controller	cr14-5508-1	
Secondary Controller	cr23-5508-2	10.125.30.3
Tertiary Controller		

229018

## 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 WLC

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 5-43.

Figure 5-43 Adding the AP1520 MAC Address to the WLC

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 CCVE 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 5-44. None of the other parameters need to be changed on this screen.

Figure 5-44 Setting the AP role to Root AP

All APs > Details for cr36-1522-1-SB

AP Role: RootAP

Bridge Type: Outdoor

Bridge Group Name: [ ]

Ethernet Bridging:

Backhaul Interface: 802.11a

Bridge Data Rate (Mbps): 24

Ethernet Link Status: UpDnNANA

Heater Status: OFF

Internal Temperature: 42 °C

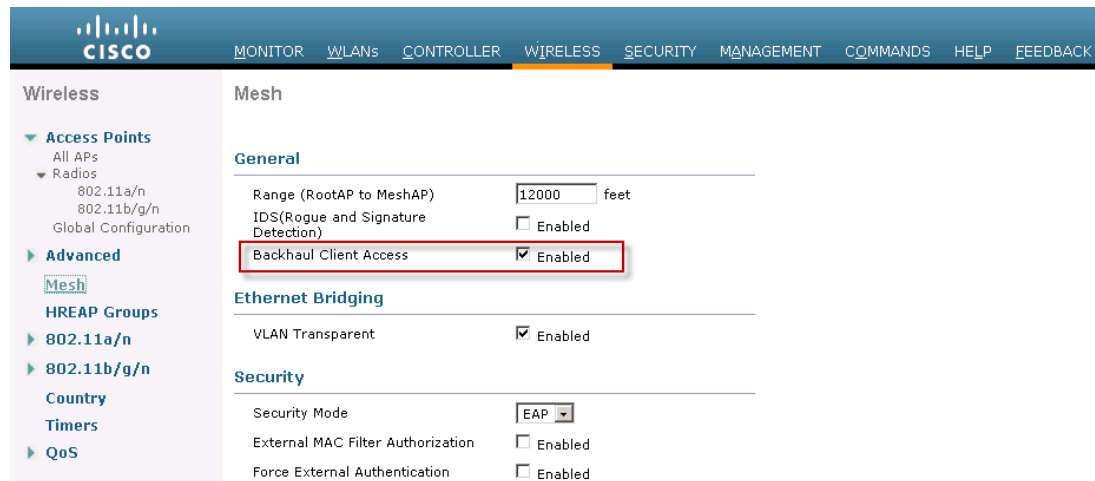
1. MAPs and RAPs are explained in much more detail in the *Cisco Wireless Mesh Access Points Design and Deployment Guide*, Release 6.0 at the following URL:

[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)

## 5 GHz Backhaul Client Access

By default, the 5 GHz radio interface on the AP1522 is enabled only as a backhaul 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 **Backhaul Client Access** checkbox on the WLAN controller's **Wireless > Mesh** configuration panel, as shown in [Figure 5-45](#). Enabling this once on the WLAN controller enables backhaul client access for all AP1520 series access points that join this controller.

**Figure 5-45** Enabling Backhaul Client Access



229021

## 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 on each campus in the CCVE design is a relatively straightforward process. After installing WCS in the main campus, 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 CCVE system 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 CCVE 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 department chairman of a community college may have the ability to view and report on certain characteristics of wireless assets for his college campus 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 department chairman is assigned to may only allow such functionality to be exercised against these assets if they are located within his college campus. Thus, if the department chairman for campus “A” attempted to use WCS to discover or manage wireless infrastructure located in campus “B”, his assigned virtual domain might not allow the ability to manage or even view resources on campus “B”.

Administrative personnel with college system-wide responsibilities, on the other hand, could be assigned a virtual domain that includes all resources in the system (i.e., all campuses), 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-campus WCS traffic, especially traffic whose nature might be based more upon curiosity rather than actual need.



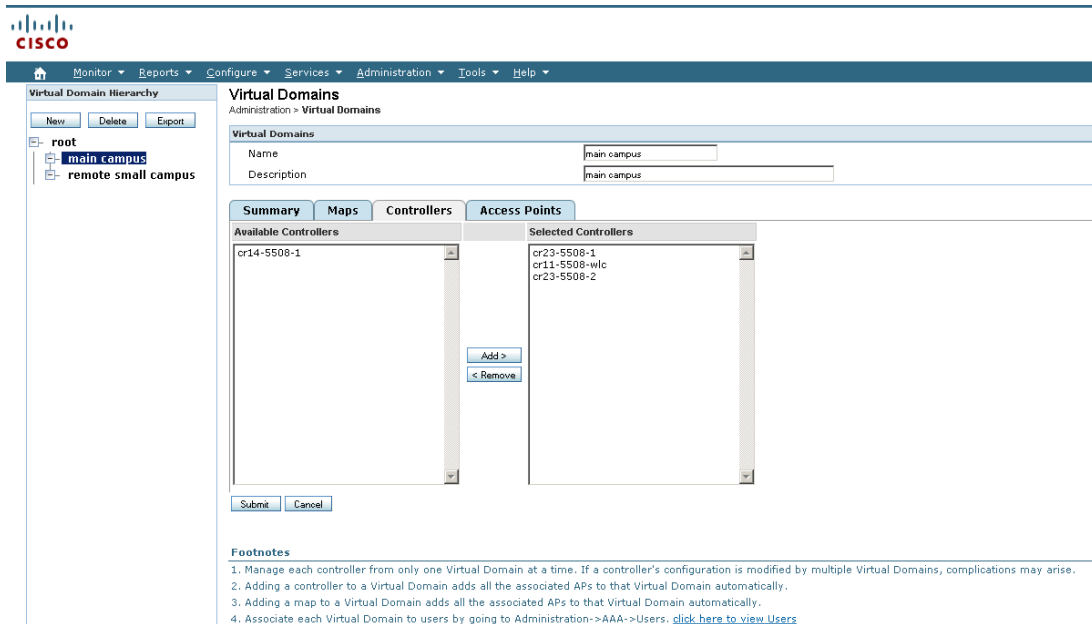
**Note**

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 5-46](#) provides an illustration of how controller resources were assigned during lab testing for the “main campus” virtual domain.

**Figure 5-46** Assigning WLC resources to the Main Campus Virtual Domain



229051

The process for creating and assigning network resources to the virtual domain is detailed in Chapter 20, “Virtual Domains” of the *WCS Configuration Guide 6.0*, 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)

- The virtual domain must be assigned to the user. The process for assigning the main campus virtual domain to the “main1” user is shown in Figure 5-47. 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 5-47 Assigning the Virtual Domain to a User**

**User Detail :main1**  
Administration > AAA > Users > User Detail

**Virtual Domains**

Available Virtual Domains	Selected Virtual Domains
root	main campus
remote small campus	

Buttons: Add > < Remove

Submit Cancel

**Footnotes:**

- Click [here](#) for current password policy.
- If user belongs to 'LobbyAmbassador' or 'Monitor Lite' or 'North Bound API' or 'Users Assistant' group then he cannot belong to any other group.
- Root group is only assignable to 'root' user and that assignment cannot be changed.
- 'root' Virtual Domain cannot be removed from Selected Virtual Domains for 'root' user.



**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/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)
- 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 6

## Community College Security Design

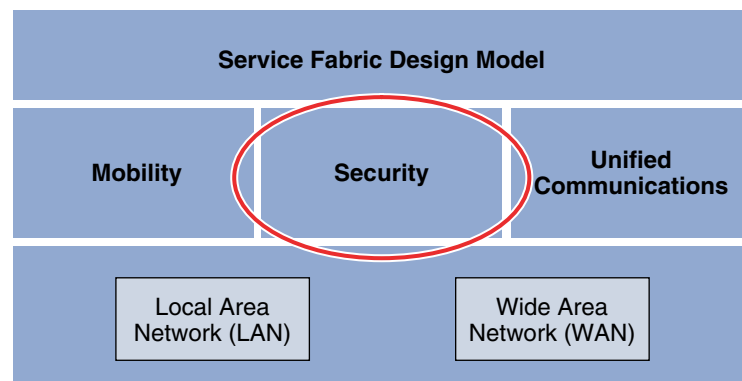
As community colleges embrace new communication and collaboration tools, transitioning from traditional classroom teaching into an Internet-based, media-rich education and learning environment, a whole new set of network security challenges arise. Community college network infrastructures must be adequately secured to protect students, staff, and faculty 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 community college administrators and community leaders.

### Security Design

Within the Cisco Community College reference design, the service fabric network provides the foundation on which all solutions and services are built to solve the business challenges facing community colleges. These business challenges include building a virtual learning environment, providing secure connected classrooms, ensuring safety and security, and operational efficiencies.

The service fabric consists of four distinct components: LAN/WAN, security, mobility, and unified communications, as shown in [Figure 6-1](#).

**Figure 6-1** Service Fabric Design Model



The Community College reference design includes security to protect the infrastructure and its services to provide a safe and secure online environment for teaching and learning. This design 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 solution to protect students and faculty from harmful content; to guarantee the confidentiality of student, staff, and faculty private data; and to ensure the availability and integrity of the systems and data.

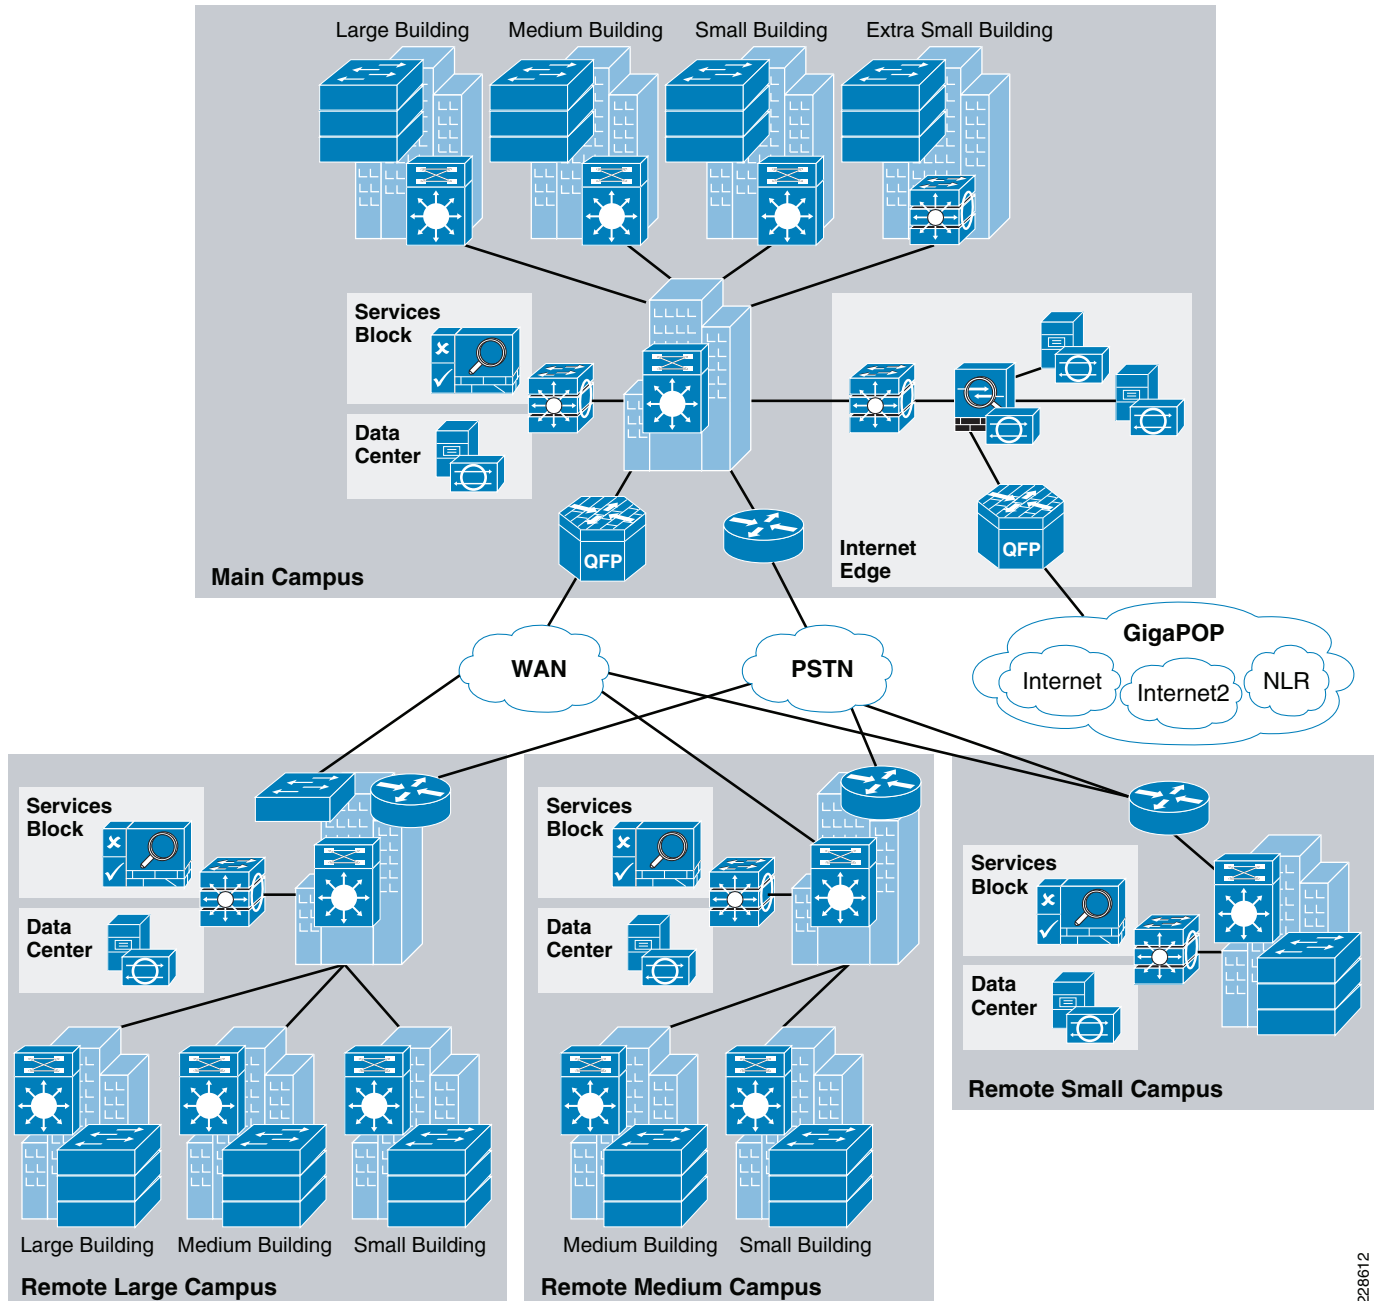
Protecting the infrastructure and its services requires implementation of security controls capable of mitigating both well-known and new forms of threats. Common threats to community college environments include the following:

- *Service disruption*—Disruption to the infrastructure and learning resources such as computer labs 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 students, faculty, and staff; peer-to-peer file sharing and instant messaging abuse; and access to forbidden content
- *Unauthorized access*—Intrusions, unauthorized users, escalation of privileges, IP spoofing, and unauthorized access to restricted learning and administrative resources
- *Data loss*—Loss or leakage of student, staff, and faculty private data from servers and user endpoints
- *Identity theft and fraud*—Theft of student, staff, and faculty identity or fraud on servers and end users through phishing and E-mail spam

The Community College reference design accommodates a main campus and one or more remote smaller campuses interconnected over a metro Ethernet or managed WAN service. Each of these campuses may contain one or more buildings of varying sizes, as shown in [Figure 6-2](#).



Figure 6-2 Community College Reference Design Overview

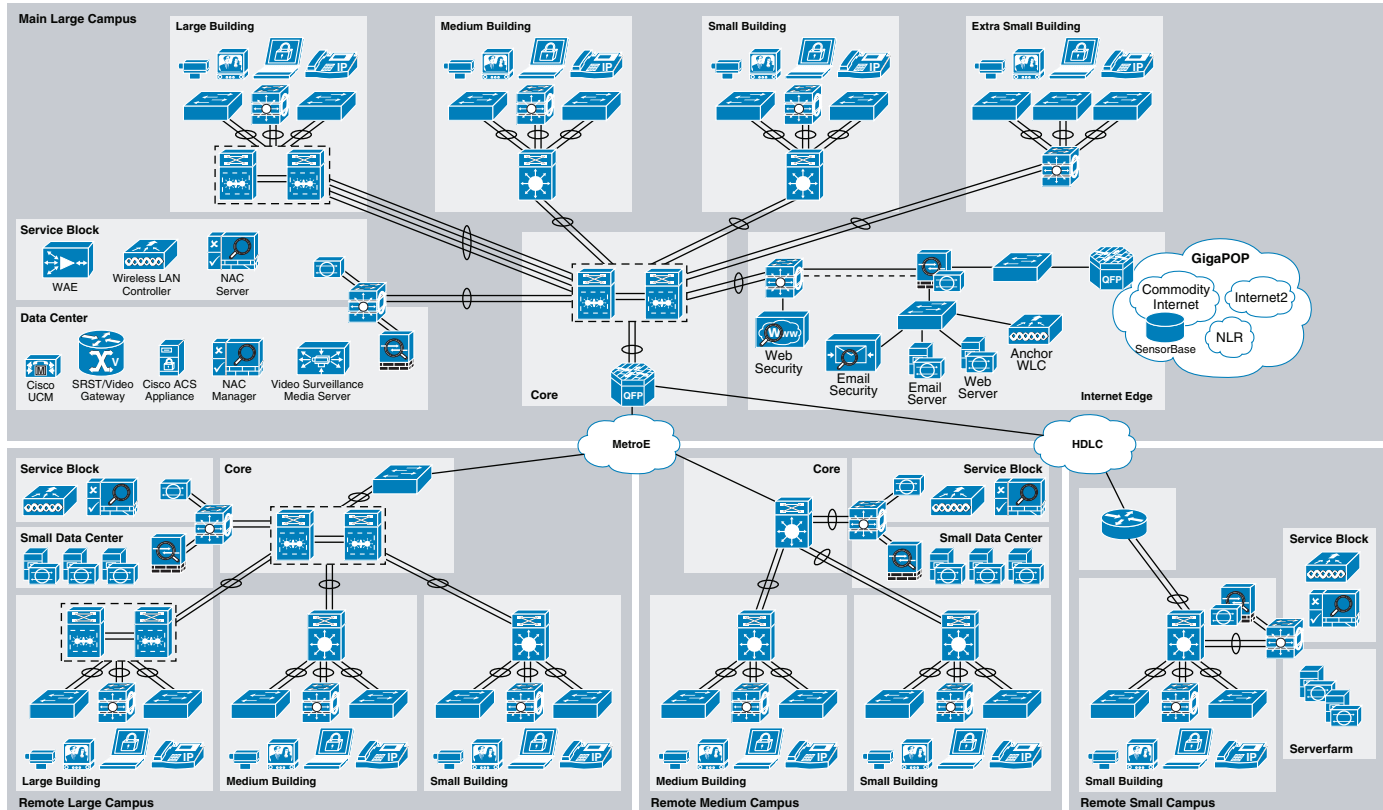


228612

Operating on top of this network are all the services used within the community college 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 campus building, allowing each remote campus to reduce the need for separate services to be operated and maintained by community college IT personnel. These centralized systems and applications are served by a data center in the main campus.

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 6-3.

Figure 6-3 Community College Network Security Design Overview



The following security elements should be included in the Community College Security design depicted in Figure 6-3:

- *Endpoint Security*—Desktop endpoint protection for day-zero attack protection, data loss prevention, and signature-based antivirus.
- *Network Foundation Protection*—Device hardening, control, 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. Granular access control.
- *Cisco Video Surveillance*—Monitor activities throughout the school environment to prevent and deter safety incidents.
- *Enhanced Availability and Resiliency*—Hardened devices and high availability design ensure optimal service availability. 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). Role-Based access control and device security compliance via Cisco Network Admission Control (NAC) Appliance.

The Community College reference design recognizes that cost and limited resources are 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 security design for the community college service fabric focuses on the following key areas.

- *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, Internet2, and National LambdaRail (NLR) networks
  - Protecting internal resources and users from botnets, malware, viruses, and other malicious software
  - Protecting students, staff, and faculty 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 student, staff, and faculty records
- *Network access security and control*
  - Securing the access edges
  - Enforcing authentication and role-based access for students, staff, and faculty residing at the main and remote campuses
  - Ensuring that systems are up-to-date and in compliance with the community college’s network security policies
- *Network endpoint protection*
  - Protecting servers and school-controlled systems (computer labs, school-provided laptops, and so on) from viruses, malware, botnets, and other malicious software
  - Enforcing E-mail and web browsing policies for staff and faculty

Together, these key security areas create a defense-in-depth solution for protecting community colleges 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

The community college network is 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 community college network infrastructure, the NFP best practices should be implemented 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 community college 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 campus.
  - 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 remote campus 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 Community College reference design assumes a centralized connection to the Internet, Internet2, and National LambdaRail (NLR) networks at the main campus site. This connection serves students, staff, and faculty at the main campus as well as all remote campus sites. Common services typically provided by this connection include the following:

- E-mail for staff and faculty
- Internet browsing for everyone
- Community college web portal accessible over the Internet
- Connectivity to other educational institutions over the Internet2 and NLR network
- Remote access to the community college network

The Internet2 network is a not-for-profit advanced networking consortium comprised of more than 200 U.S. universities in cooperation with 70 leading corporations, 45 government agencies, laboratories, and other institutions of higher learning as well as over 50 international partner organizations. Internet2 provides its members both leading-edge network capabilities and unique partnership opportunities that together facilitate the development, deployment, and use of revolutionary Internet technologies. The physical implementation of Internet2 network consists of an advanced IP network, virtual circuit network, and core optical network. The Internet2 network provides the necessary scalability for member institutions to efficiently provision resources to address the bandwidth-intensive requirements of their

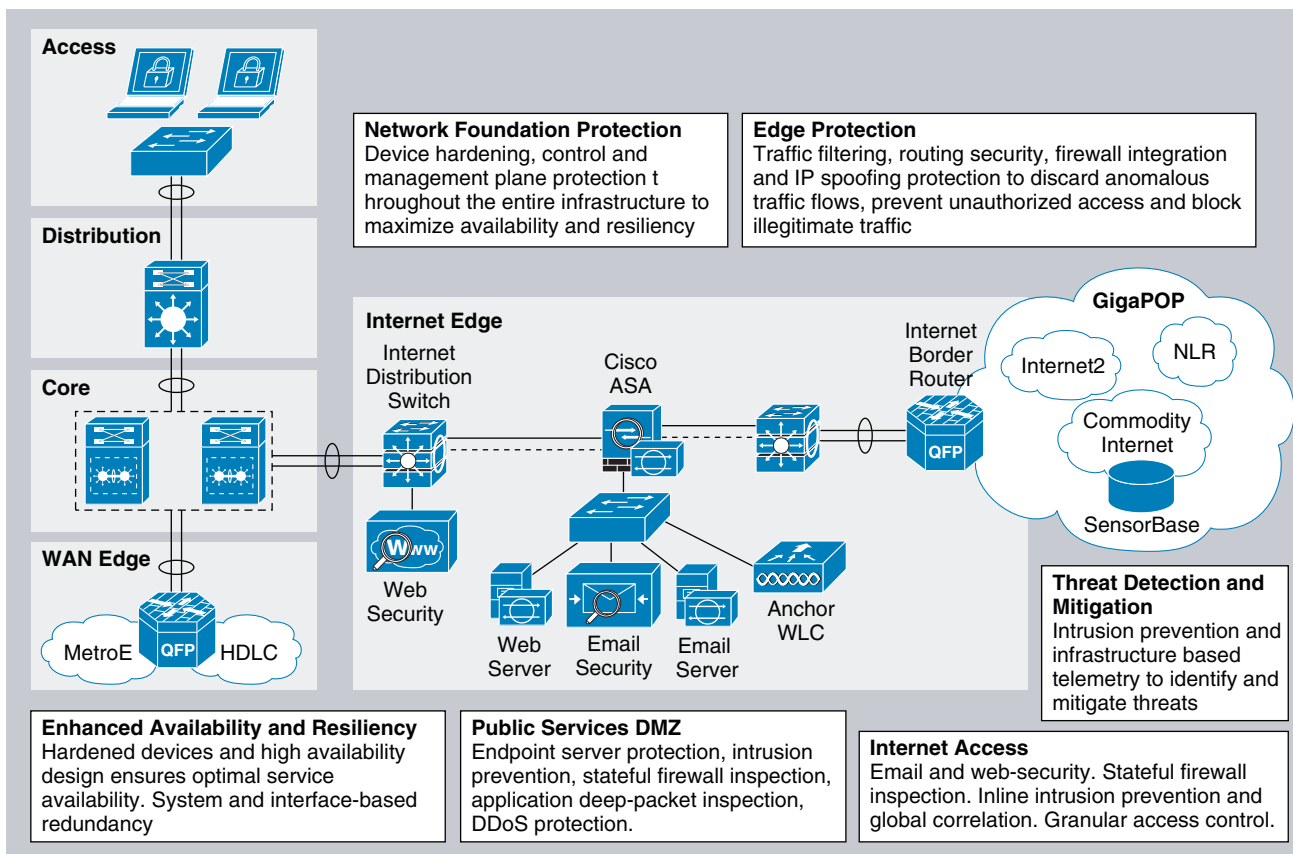
campuses, such as collaborative applications, distributed research experiments, grid-based data analysis, and social networking. For more information on the Internet2 network, see the following URL: <http://www.internet2.edu/network/>.

The National LambdaRail (NLR) network is a high-speed fiber optic network infrastructure linking over 30 cities in 21 states. It is owned by the U.S. research and education community and is dedicated to serving the needs of researchers and research groups. The NLR high-performance network backbone offers unrestricted usage and bandwidth, a choice of cutting-edge network services and applications, and customized service for individual researchers and projects. NLR services include high-capacity 10 Gigabit Ethernet LAN-PHY or OC-192 lambdas, point-to-point or multi-point Ethernet-based transport, routed IP-based services, and telepresence video conferencing services. For more information on the NLR network and its services, see the following URL: <http://www.nlr.net/>.

Community colleges typically connect to a local Gigabit point-of-presence (GigaPOP) or regional network service provider to gain access to the Internet, Internet2, and NLR networks. The same security controls are applicable regardless of whether they connect to a GigaPOP or regional network. For details on how community colleges connect to these networks, see Chapter 4, “Community College WAN Design.”

The part of the network infrastructure that provides connectivity to the Internet, Internet2, and NLR is defined as the Internet perimeter, as shown in Figure 6-4.

Figure 6-4 Internet Perimeter



The Internet perimeter provides safe and secure access to the Internet, Internet2, and NLR networks for students, staff, and faculty. It also provides access to public services such as the community college web portal without compromising the confidentiality, integrity, and availability of the resources and data of the educational institution. 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 community college and the Internet, Internet2, and NLR networks. It may be administered by the community college IT staff or may be managed by the Internet, Internet2, or NLR 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 community college resources and data from unauthorized access and disclosure. The ASA 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 ASA is configured to control or prevent incoming and outgoing access for the Internet, Internet2, and NLR networks; to protect the community college web portal and other Internet public services; and to control student, staff, and faculty traffic bound towards the Internet. The security appliance may also provide secure remote access to faculty, staff, and students in the form of a Secure Socket Layer (SSL) or IPsec virtual private network (VPN).
- *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 configured either in inline or promiscuous mode. IPS may also be configured to help block certain Internet applications such as AOL Messenger, BitTorrent, Skype, and so on.
- *Public services DMZ*—The community college 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 community college 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 (such as Cisco Security Agent) 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 content that may be harmful for students, staff, and faculty such as sites known to be 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, and for monitoring Layer 4 traffic for rogue activity and infected systems.
- *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 community college. For more information on the wireless guest access solution, see [Chapter 5, “Community College Mobility Design.”](#)

The following subsections describe the design guidelines for implementing the above security functions.

## Internet Border Router Security

The Internet border router connects to a local GigaPOP and provides connectivity to the Internet, Internet2, and NLR networks for the community college. 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 6-5](#). This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information.

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 community college internal resources and data from external threats by doing the following:

- Preventing incoming access from the Internet, Internet2, and NLR networks
- 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-, Internet2- and NLR-bound traffic
- Monitoring network ports for rogue activity and preventing infected internal endpoints from sending command and control traffic back to a host on the Internet

The 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, Internet2, and NLR to internal resources and subnets.
- Allow outbound Internet HTTP/HTTPS access for students, staff, and faculty residing at any of the community college campuses.
- Allow outbound SSL access to the Internet for devices requiring administrative updates such as SensorBase, IPS signature updates, and so on.
- Deny or control access between the community college internal network and the external Internet, Internet2, or NLR networks.



- Allow students, staff, and faculty access to DMZ services such as the community college 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.

In addition, the Cisco ASA Botnet Traffic Filter feature can be enabled to monitor network ports for rogue activity and to prevent infected internal endpoints from sending command and control traffic back to an external host on the Internet. The Botnet Traffic Filter on the ASA provides reputation-based control for an IP address or domain name, similar to the control that Cisco IronPort SenderBase provides for E-mail and web servers.

The Cisco Botnet Traffic Filter is integrated into all Cisco ASA appliances, and inspects traffic traversing the appliance to detect rogue traffic in the network. When internal clients are infected with malware and attempt to phone home to an external host on the Internet, the Botnet Traffic Filter alerts the system administrator of this through the regular logging process and can be automatically blocked. This is an effective way to combat botnets and other malware that share the same phone-home communications pattern.

The Botnet Traffic Filter monitors all ports and performs a real-time lookup in its database of known botnet IP addresses and domain names. Based on this investigation, the Botnet Traffic Filter determines whether a connection attempt is benign and should be allowed, or is a risk and should be blocked.

The Cisco ASA Botnet Traffic Filter has three main components:

- *Dynamic and administrator blacklist data*—The Botnet Traffic Filter uses a database of malicious domain names and IP addresses that is provided by Cisco Security Intelligence Operations. This database is maintained by Cisco Security Intelligence Operations and is downloaded dynamically from an update server on the SensorBase network. Administrators can also configure their own local blacklists and whitelists.
- *Traffic classification and reporting*—Botnet Traffic Filter traffic classification is configured through the **dynamic-filter** command on the ASA. The dynamic filter compares the source and destination addresses of traffic against the IP addresses that have been discovered for the various lists available (dynamic black, local white, local black), and logs and reports the hits against these lists accordingly.
- *Domain Name System (DNS) snooping*—To map IP addresses to domain names that are contained in the dynamic database or local lists, the Botnet Traffic Filter uses DNS snooping in conjunction with DNS inspection. Dynamic Filter DNS snooping looks at User Datagram Protocol (UDP) DNS replies and builds a DNS reverse cache (DNSRC), which maps the IP addresses in those replies to the domain names they match. DNS snooping is configured via the Modular Policy Framework (MPF) policies

The Botnet Traffic Filter uses two databases for known addresses. Both databases can be used together, or the dynamic database can be disabled and the static database can be used alone. When using the dynamic database, the Botnet Traffic Filter receives periodic updates from the Cisco update server on the Cisco IronPort SensorBase network. This database lists thousands of known bad domain names and IP addresses.

The 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 ASA sends a syslog message reporting the suspicious activity and optionally drops the traffic if the 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 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 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 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 ASA. Cisco also recommends that DNS packet inspection be enabled with Botnet Traffic Filter snooping. When enabled, the ASA uses Botnet Traffic Filter snooping instead of the regular DNS lookup to resolve static blacklist domain names in the following circumstances:

- The ASA DNS server is unavailable.
- A connection is initiated during the one minute waiting period before the 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 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 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).

When deploying the Internet firewall, it is important to understand the traffic and policy requirements when selecting a firewall. An appropriately sized 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 6-5](#). 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.

Given budget and resource constraints for community colleges, high availability is achieved by using redundant physical interfaces, which provides a cost-effective solution. As an alternative, a pair of firewall appliances can be deployed in stateful failover using separate boxes at a higher cost.

## Intrusion Prevention

IPS is responsible for identifying and blocking anomalous traffic and packets recognized as well-known attacks. An 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 community colleges. 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 community college network.

The AIP SSM module may also participate in Cisco Global Correlation for further threat visibility and control. If 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. It then incorporates the global threat data into its system to detect and prevent malicious activity even earlier. The IPS uses this information to filter out the worst attackers before they have a chance to attack critical assets.

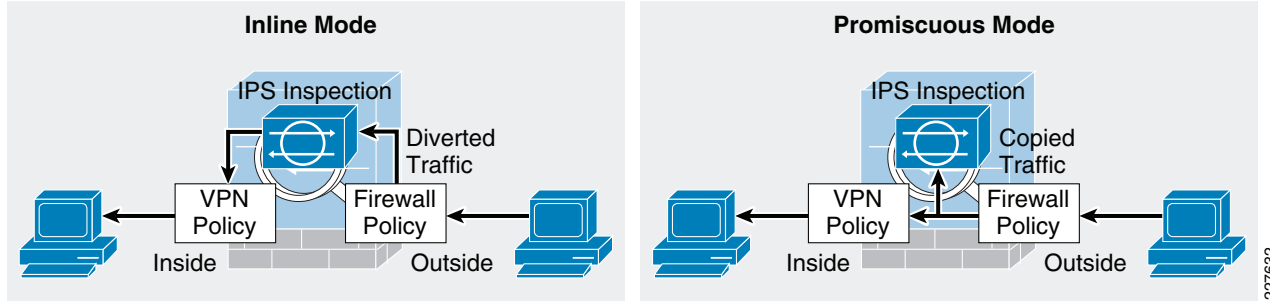
For more information on IPS Global Correlation including configuration information, see the following URL:

[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).

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 6-5](#)). Traffic identified for IPS inspection cannot continue through the ASA without first passing through and being inspected by the AIP SSM. This mode is the most secure because every packet that has been identified for inspection is analyzed before being allowed through. Also, the AIP SSM can implement a blocking policy on a packet-by-packet basis. This mode, however, can affect throughput if not designed or sized appropriately.
- *Promiscuous mode*—A duplicate stream of traffic is sent to the AIP SSM. This mode is less secure, but has little impact on traffic throughput. Unlike inline mode, in promiscuous mode the AIP SSM can block traffic only by instructing the ASA to shun the traffic or by resetting a connection on the ASA. Also, while the AIP SSM is analyzing the traffic, a small amount of traffic might pass through the ASA before the AIP SSM can shun it. The right side of [Figure 6-5](#) shows the AIP SSM in promiscuous mode.

Figure 6-5 IPS Inline and Promiscuous Modes



The recommended IPS deployment mode depends on the goals and policies of the community college. IPS inline mode is more secure because of its ability to stop malicious traffic in real-time; however, it may impact traffic throughput if not properly designed or sized. Conversely, IPS promiscuous mode has less impact on traffic throughput but is less secure because there may be a delay in reacting to the malicious traffic.

Although the AIP SSM runs as a separate application within the Cisco ASA, it is integrated into the traffic flow. The AIP SSM contains no external interfaces itself, except for the management interface on the SSM itself. When traffic is identified for IPS inspection on the ASA, traffic flows through the ASA and the AIP SSM in the following sequence:


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 ASA.

The AIP SSM card may be configured to fail open or close when the module becomes unavailable. When configured to fail open, the ASA allows all traffic through, uninspected, if the AIP SSM becomes unavailable. Conversely, when configured to fail close, the ASA blocks all traffic in case of an AIP SSM failure.

## E-Mail Security

Cisco recommends that the Cisco IronPort C Series E-Mail Security Appliance (ESA) be deployed 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 6-6](#).

Figure 6-6 Logical E-Mail Delivery Chain



227423

**Note**

Figure 6-6 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, a community college 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.

There are two options for deploying the ESA appliance, depending on the number of interfaces used:

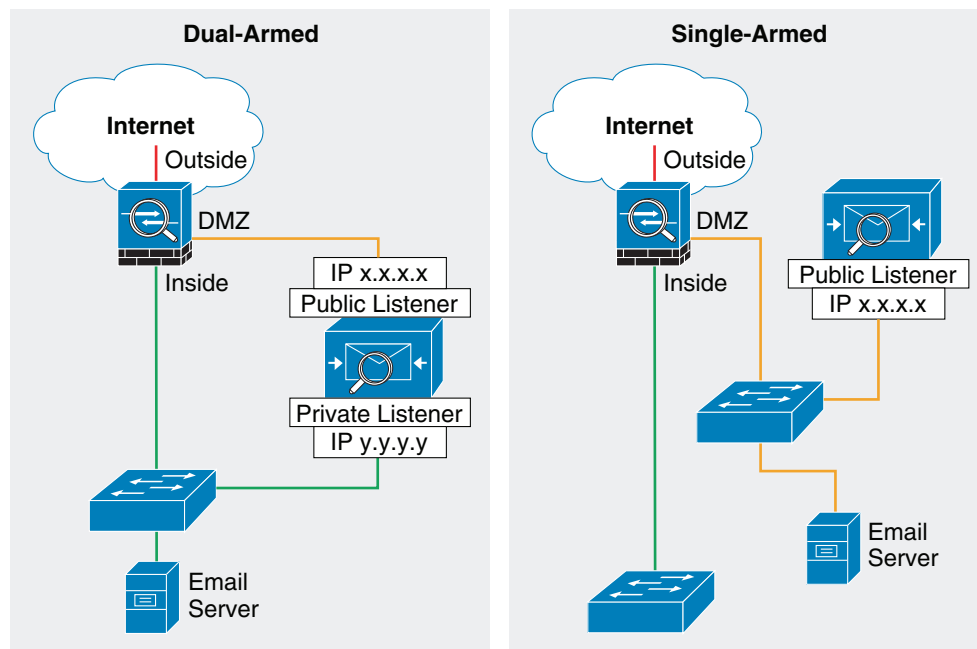
- *Dual-armed configuration*—Two physical interfaces are used to serve as a public mail listener and a private mail listener where each interface is configured with a separate logical IP address. The public listener receives E-mail from the Internet and directs messages to the internal mail servers.

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 6-7 shows both configurations.

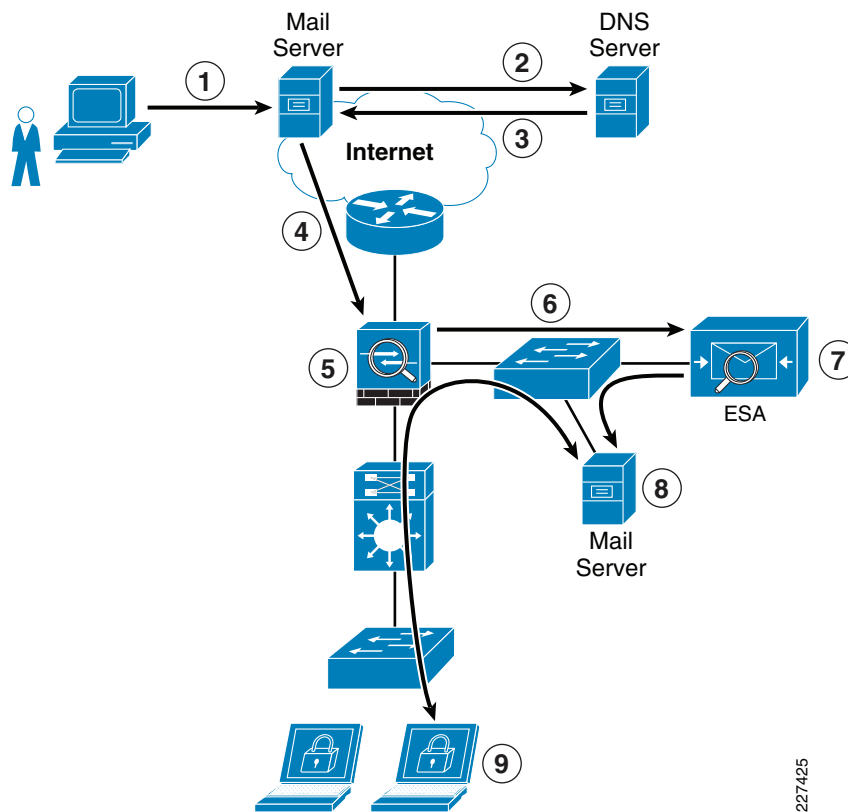
**Figure 6-7 Common ESA Deployments**



For simplicity, Cisco recommends that the community college network implement the ESA with a single interface in a single-armed configuration. This also leaves the other data interfaces available for redundancy.

Figure 6-8 shows the logical location of the ESA within the E-mail flow chain and the typical data flow for inbound E-mail traffic.

Figure 6-8 Typical Data Flow for Inbound E-Mail Traffic



227425


The following steps explain what is taking place in [Figure 6-8](#):

- 
- Step 1** Sender sends an E-mail to xyz@domain X.
  - Step 2** What's the IP address of domain X?
  - Step 3** It is a.b.c.d (public IP address of ESA).
  - Step 4** E-mail server sends message to a.b.c.d using SMTP.
  - Step 5** Firewall permits incoming SMTP connection to the ESA, and translates its public IP address.
  - Step 6** ESA performs a DNS query on sender domain and checks the received IP address in its reputation database, and drops, quarantines E-mail based on policy.
  - Step 7** ESA forwards E-mail to preconfigured inbound E-mail server.
  - Step 8** E-mail server stores E-mail for retrieval by receiver.
  - Step 9** Receiver retrieves E-mail from server using POP or IMAP.
- 

The ESA appliance functions as an SMTP gateway, also known as a mail exchange (MX). The following outlines some of the deployment design points for the ESA within the community college design:

- The ESA appliance needs to be accessible via the public Internet and is the first hop in the E-mail infrastructure. 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.
- All the local domains for which the ESA appliances accept mail need to be added to the recipient access table (RAT). Inbound E-mail destined to domains not listed in the RAT is 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.
- SMTP routes are defined 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 6-9](#).



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

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


The Community College reference design implements a Cisco IronPort S Series Web Security Appliance (WSA) to block HTTP and HTTPS access to sites on the Internet with content that may be harmful, and to protect the community college network from web-based malware and spyware.

The following services may be enabled on the WSA:

- *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 community college design assumes a centralized Internet connection implemented at the main campus 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 students, staff, and faculty located at any of the community college campuses. Logically, the WSA sits in the path between web users and the Internet, as shown in [Figure 6-10](#).

Figure 6-10 Cisco IronPort WSA



There are the following two deployment modes when 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 administrators to have control over the configuration of the endpoints, which is often not the case in community college environments. For example, community colleges may allow students, guests, or visiting professors to use personal laptops, smart phones, or other devices outside the administration of the institution. 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, Cisco recommends that community colleges implement transparent proxy mode with WCCP. In the Community College reference design, 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 switches support 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 6-11 shows how WCCP redirection works in conjunction with the Cisco Catalyst 3750 StackWise distribution switches.

**Figure 6-11 WCCP Redirection**



As shown in Figure 6-11, the following steps take place:

- 
- Step 1** The client browser requests a connection to `http://website.com`.
  - Step 2** The Cisco Catalyst 3750 Internet perimeter distribution switch intercepts and redirects HTTP/HTTPS requests to WSA via Layer 2 redirection.
  - Step 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.
  - Step 4** If allowed, WSA fetches the content from the destination website.

**Step 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 community college 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).

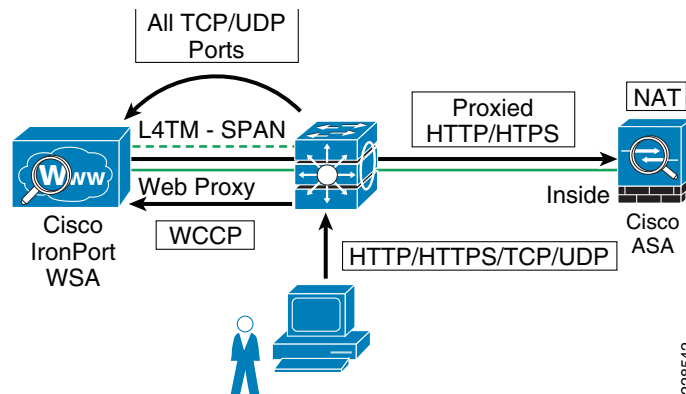
The WSA L4TM service is deployed independently from the web proxy functionality. L4TM monitors network traffic for rogue activity and for any attempts to bypass port 80. It works by listening to all UDP and TCP traffic and by matching domain names and IP addresses against entries in its own database 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>).

The following are some of the key guidelines when deploying L4TM:

- **Physical connection**—L4TM requires a copy of the traffic to be redirected to the WSA for monitoring. This can be done by connecting a physical network tap, configuring Switch Port Analyzer (SPAN) port mirroring on a Cisco Catalyst switch, or using a hub. Network TAPs forward packets in hardware, while SPAN port mirroring is generally done in software. However, SPAN port mirroring can be easily reconfigured, providing more flexibility.
- **Location**—L4TM should be deployed in the network where it can see as much traffic as possible before going out to the Internet through the firewall. Monitoring should occur before any device that performs NAT on client IP addresses.
- **Action setting**—The default action setting for L4TM is to passively monitor only. Optionally, you can configure L4TM to monitor and actively block suspicious traffic. TCP connections are reset by the generation of TCP resets, and UDP sessions are torn down using ICMP unreachables. The use of L4TM blocking requires that the L4TM and web proxy services are placed on the same network so that all clients are accessible on routes that are configured for data traffic.

In the community college design, L4TM can be deployed by configuring a SPAN session on the Internet perimeter distribution switch to monitor all TCP and UDP traffic on the links connecting to the core distribution switches. Using SPAN provides greater flexibility. Monitoring the distribution switches link to the core switches ensures that all client traffic is inspected before NAT and before traffic is sent to the Internet. Figure 6-12 shows this L4TM deployment option.

**Figure 6-12 L4TM Deployment**



If the Internet perimeter firewall is configured to block all traffic bound to the Internet except HTTP and HTTPS traffic from the WSA, or the ASA Botnet Traffic Filter feature is enabled, L4TM may not provide any additional benefit. Also, if active mitigation is required, the Cisco ASA Botnet Traffic Filter Feature or a Cisco IPS appliance or module deployed in inline mode is recommended. Both the ASA Botnet Traffic Filter and inline IPS provide better mitigation by blocking traffic automatically inline, stopping malicious traffic before it reaches the intended target.

For more information on how to configure the 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

Community colleges typically implement a data center that hosts the systems that serve the administrative and educational applications and store the data accessible to internal users. The infrastructure supporting them may include application servers, 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 Community College reference design 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 community college, a single data center may be deployed in the main campus. Smaller data centers or server farms may also be deployed in remote campuses as required.

Securing the data center is beyond the scope of this document. For more information on the best practices for securing a data center, 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 such as students, staff, and faculty 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, and students did not carry computers with them. Today, with the proliferation of wireless networks, increased use of laptops and smart mobile devices, the community college IT department 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 students carrying laptops are common. There is nothing preventing a contractor or student from plugging into a wall jack in a classroom, lab, or conference room to gain access to the community college 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 help 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 Community College reference design achieves access security and control by leveraging the following technologies:

- Cisco Catalyst Integrated Security Features (CISF)
- Cisco Network Admission Control (NAC) Appliance
- Cisco Identity-Based Network Services (IBNS)

## Cisco Catalyst Integrated Security Features

Cisco CISF is a set of 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 access switches and ports within the community college 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 MITM 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 Identity-Based Network Services

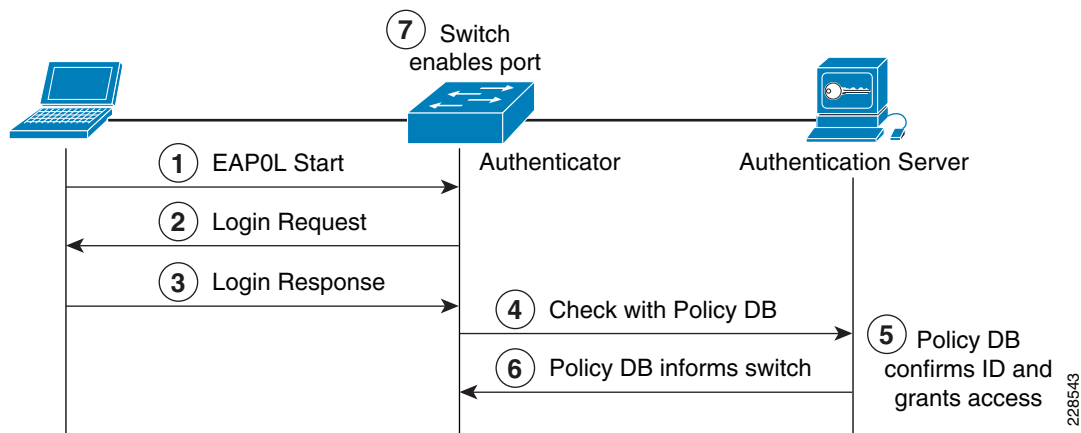
The Cisco IBNS solution is a set of Cisco IOS software services that provide secure user and host access to enterprise networks powered by Cisco Catalyst switches and wireless LANs. It provides standards-based network access control at the access layer by using the 802.1X protocol to secure the physical 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 6-13](#) illustrates how port-based access control works within an identity-based system.

**Figure 6-13** Port-Based Access Control



The following steps describe the port-based access control flow shown in [Figure 6-13](#):



- 
- Step 1** A client, such as a laptop with an 802.1X supplicant, connects to an IEEE 802.1X-enabled network and sends a start message to the LAN switch (the authenticator).
  - Step 2** When the start message is received, the LAN switch sends a login request to the client.
  - Step 3** The client replies with a login response.
  - Step 4** The switch forwards the response to the policy database (authentication server).
  - Step 5** The authentication server authenticates the user.
  - Step 6** After the user identity is confirmed, the policy database authorizes network access for the user and informs the LAN switch.
  - Step 7** The LAN switch then enables the port connected to the client.
- 

The user 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. 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 are always 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 Community Colleges

As mentioned in the previous sections, 802.1X authentications require a supplicant on the host device. This typically poses a problem in community college environments that have a wide range of host devices and limited or no management of many of these devices. This makes a community college-wide 802.1X deployment very challenging. However, there may be pockets of a community college 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 school administration office or shared labs where PCs are managed. Other locations in the community college network still need protection, but student and faculty network access may be better served by a NAC Appliance Solution (discussed in the next section). In addition, for networks requiring role-based access control using posture assessments to ensure security compliance, Cisco NAC Appliance should be considered.

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.

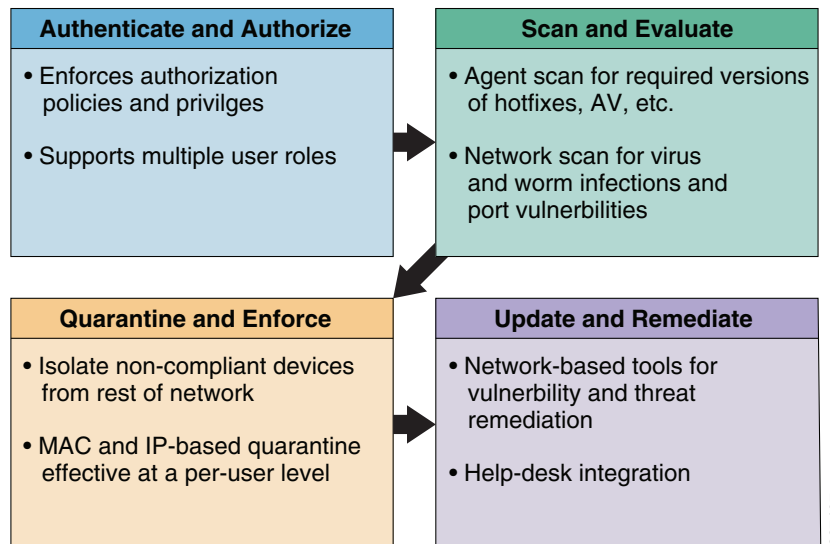
When deployed, Cisco NAC Appliance provides the following benefits:

- 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 6-14:

- Authenticates and authorizes
- Scans and evaluates
- Quarantines and enforces
- Updates and remediates

**Figure 6-14** 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 antivirus software, and quarantines vulnerable or infected clients for remediation before clients access the network.

Figure 6-15 shows Cisco NAC Appliance components.

Figure 6-15 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 Community College reference design, the CAM is located in the data center at the main campus site.

## Cisco Clean Access Server

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 (like 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 campus 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 Community College reference design, the CAS would be located at the main campus 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

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.

---

In the community college, Cisco recommends that the CAA be used for the managed PCs, such as those for administrators and faculty.

### 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 Community College reference design, the NAC Web Agent is used for unmanaged clients such as student laptops and guest professors.

### 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 6-16](#)). When a user is authenticated and passes all policy checks, their traffic is switched normally through the network and bypasses the NAC Appliance.

Figure 6-16 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 6-16, 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 6-17 shows a logical in-band topology example, and Figure 6-18 shows a physical in-band topology example.

**Figure 6-17 In-Band Virtual Gateway Topology**



**Figure 6-18 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 6-18](#).

- *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 6-17](#). 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 6-1](#) summarizes various characteristics of the two deployment types.

**Table 6-1** *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 6-1 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 recommended deployment option for the Community College reference design 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 Community College Reference Design

Within the Community College reference design, a NAC Appliance solution is deployed at each of the site types; main campus, remote large campus, remote medium campus, and remote small campus. A centralized CAM is deployed at the main campus and is deployed within the data center at that site. A CAS is deployed at each of the campus sites (main and remote sites) and is connected within the service block connecting to the core switches at each of the sites.

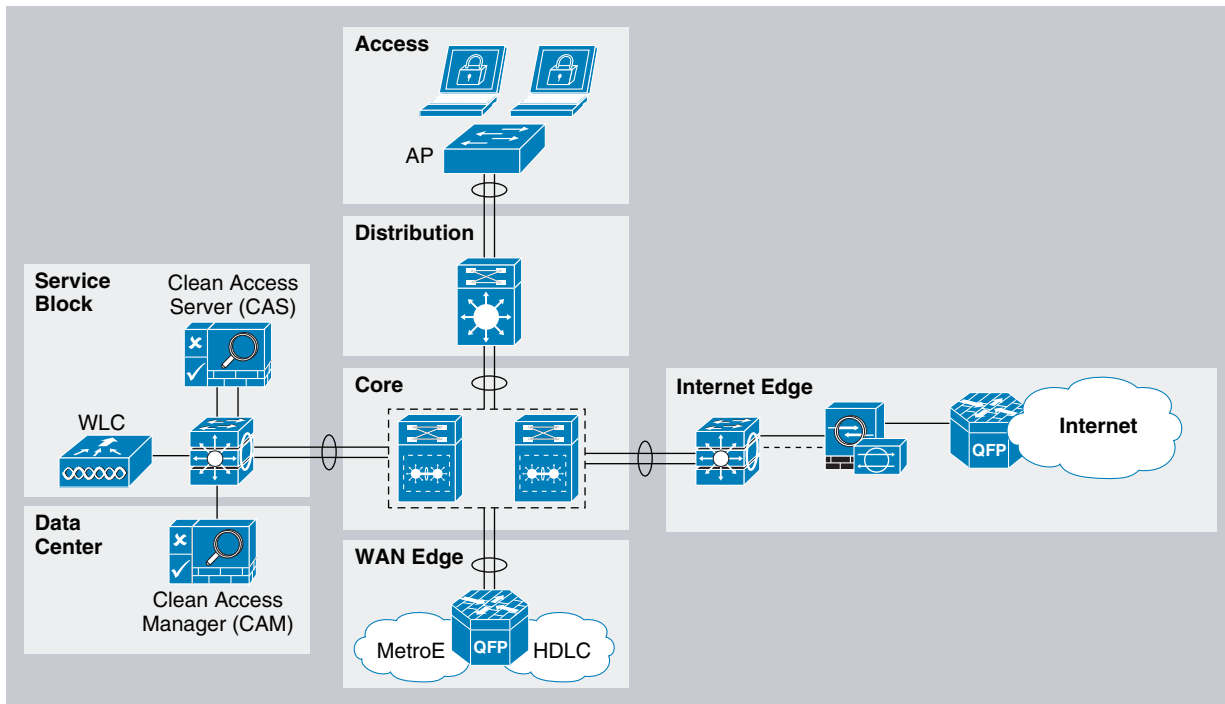
The Community College reference design provides host network connectivity using wired and wireless technologies. As such, the NAC Appliance solution must provide a solution for both connectivity options. For wireless clients, a Layer 2 OOB NAC solution is deployed, and for wired clients, a Layer 2 OOB or a Layer 3 OOB NAC solution may be deployed.

## NAC Deployment for Wireless Clients

To provide network access control for wireless clients within the Community College reference design, the recommended design is the virtual gateway (bridge mode) and central deployment OOB solution. In this design, the NAC server must be placed Layer 2-adjacent to the WLC. In the Community College reference design, the WLCs are centrally deployed at each campus and are implemented in the service block off the core switches, as detailed in [Chapter 5, “Community College 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 6-19.](#)

**Figure 6-19** NAC OOB Deployment for Wireless Clients




228548

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 6-20.](#))

Figure 6-20 Wireless NAC OOB Traffic Flow



When implementing the NAC OOB wireless solution, Cisco also recommends enabling 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 Community College reference design also uses a central OOB NAC deployment with a NAC server implemented at each of the campus 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 campus 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 campus 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 the 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 Community College reference design, the NAC Agent can be used for managed PCs such as administrative staff and faculty.

Web login is typically required for student login sessions because student laptops are typically 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 campus locations, a two-tier, collapsed core/distribution LAN design is recommended, as explained in [Chapter 3, “Community College 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 campus, 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 6-21](#).

Figure 6-21 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 community college. However, CAS fallback should always be configured to ensure that critical network services are available, even during a network outage.

# Endpoint Protection

Servers, desktop computers, laptops, printers, and IP phones are examples of the diverse network endpoints found in community college environments. Properly securing the endpoints requires not only adoption of the appropriate technical controls but also end-user awareness. The community college security strategy must include security awareness campaigns and programs. Students, staff, and faculty must be continuously educated in current threats, Internet-use best practices, and the security measures needed for keeping endpoints up-to-date with the latest updates, patches, and fixes.

The Community College reference design implements a range of security controls designed to protect the endpoints, which include Cisco host-based IPS, network-based IPS, and web and E-mail traffic security.

For host-based IPS, the Community College reference design leverages the Cisco Security Agent on managed end-user workstations and servers. Cisco Security Agent uses behavior-based security to take a proactive approach to preventing malicious activity on the hosts. When an application attempts an operation on the host, Cisco Security Agent checks the operation against the security policy of the application, and makes a real-time decision to allow or deny the operation along with determining whether to log the operation request. Security policies are assigned by IT or security administrators individually, per department, or organization-wide.

Cisco Security Agents are centrally managed with the Cisco Security Agent Management Center, which is placed in a secure segment in the data center. Cisco Security Agent Management Center also provides centralized reporting and global correlation.

## Community College Mission Relevancy

The service fabric provides the network foundation for the Community College reference design. The network service fabric is a collection of products, features, and technologies that provide a robust routing and switching foundation on which all solutions and services are built to solve the business challenges facing community colleges. These business challenges include the following:

- Virtual learning environments
- Secure connected classrooms
- Safety and security
- Operational efficiencies

The previous sections of this chapter focused on the specific design considerations for securing the community college service fabric network. This section discusses how these security design considerations relate to these business challenges.

## Virtual Learning Environments

One of the key challenges that face community colleges is extending their learning environments beyond brick and mortar campuses to allow online/distance learning, professor collaboration, and anytime, anywhere access for students to obtain course materials. Maintaining a secure virtual learning environment is critical for community colleges. The community college security design helps establish the foundation for providing a secure virtual learning environment in the following areas:

- *Secure remote access*

- Allows community colleges to extend their network to anyone, anytime, anywhere by providing a secure client-based or web-based remote access solution.
- Provides granular and encrypted access to learning resources based on user or security requirements.
- *Internet perimeter*—Protects and controls access to the community college network infrastructure from remote students, faculty, staff, and guest professors by properly securing the Internet perimeter by using firewalls, intrusion prevention systems, and a DMZ to protect against unauthorized access and malicious attacks.
- *Securing Video Portal*—Secures the network and data center to prevent misdirected students or malicious intruders from hacking into restricted servers or issuing attacks on the video portal learning infrastructure.
- *Network telemetry and monitoring*—Provides visibility into attacks or malicious activity on the network by monitoring the network using NetFlow, Syslog, and SNMP.

## Secure Connected Classrooms

Although providing connectivity to students while attending class is the foundation of twenty-first century learning, this poses many problems for community colleges. They must ensure that the person accessing the network should be allowed on the network, and that the computer accessing the network is free from viruses and other malware that might adversely affect the network and other users. In addition, while connectivity is provided, steps should be taken to prevent unauthorized access to restricted resources and protect against inadvertent or deliberate network attacks. The security design within the community college service fabric helps to address these challenges in the following ways:

- *Network access control*—Implementing role-based network access controls for wired, wireless, and remote users and devices to help 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.
- *Access layer security*—Enabling Cisco CISF on the access layer switches to protect the access layer infrastructure and users from spoofing, man-in-the-middle, DoS and other network-based attacks. CISF includes features such as Port Security, DHCP Snooping, Dynamic ARP Inspection, and IP Source Guard.
- *Web security*—Deploying a Cisco IronPort Web Security Appliance (WSA) to block access to sites with content that may be harmful or inappropriate, and to protect community colleges from web-based malware or spyware.
- *Network and data security*—Securing the network and data center to prevent misdirected students or malicious intruders from hacking into restricted servers or issuing attacks on the network infrastructure.

## Safety and Security

Providing a safe and secure environment is a top responsibility for community college administrators and community leaders. Without adequate protection, schools may be threatened by harmful or inappropriate content that can put the well-being of the students at risk, the theft of student records and private data, the loss of school network and service availability, as well as the abuse of internal applications and network resources. A safe community college is one that successfully uses the right



tools to ensure the safety and security of students, staff, and faculty, and guarantees an immediate and effective response to security and safety incidents. The most effective strategy is one that combines physical and network controls, not in isolation but rather in collaboration and with a common purpose.

Because many of the physical security components such as video surveillance and unified communication services rely on the IP infrastructure, it is critical to ensure the availability and integrity of this infrastructure. The security design within the Community College reference design helps to ensure the availability and the integrity of the network infrastructure that the physical security components rely on by focusing on the following key areas:

- *Network Foundation Protection (NFP)*—Ensuring the availability and integrity of the network infrastructure, protecting the control and management planes.
- *Internet perimeter protection*
  - Ensuring safe connectivity to the Internet, Internet2, and NLR networks and protecting internal resources and users from malware, viruses, and other malicious software.
  - Protecting students, staff, and faculty 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 student, staff, and faculty records.
- *Network access security and control*
  - Securing the access edges.
  - Enforcing authentication and role-based access for students, staff, and faculty residing at the main and remote campuses.
  - Ensuring systems are up-to-date and in compliance with the community colleges network security policies.
- *Network endpoint protection*
  - Protecting servers and school-controlled systems (computer labs, school-provided laptops, and so on) from viruses, malware, botnets, and other malicious software.
  - Enforcing E-mail and web browsing policies for staff and faculty.

## Operational Efficiencies

Community colleges are faced with the daunting task of doing more with less, facing explosive growth as budgets and resources are reduced. The Community College reference design leverages the network as a platform to deliver expanded educational services and data center optimizations to create operational efficiencies to reduce costs and capitalize on under-used network capacity. The network as a platform goes beyond merely consolidating voice, video, and data services on a single converged network; rather it consolidates all IP-based services to use the network (wired or wireless) to extend cost reduction, improve utilization on under-used networks, and add flexibility to community colleges through business process improvements.

With these critical services relying heavily on the network infrastructure, it is imperative that the IP infrastructure remains operational at all times, and it is critical that security be implemented throughout the network infrastructure to ensure the availability and the integrity of the network.

# Community College Security Configuration Guidelines

The previous sections of this chapter provides design guidelines and considerations for deploying security within a community college network environment. The sections that follow provide 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 community college or the ISP, it must be hardened following the best practices discussed in the “[Network Foundation Protection](#)” section on page 6-5. 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 Community College 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 school's 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: Explicit Permit for Traffic to Community College's Public Subnet

```
access-list 110 permit ip any 198.133.219.0 0.0.0.255
```

**Note**

---

The 64.104.0.0/16 and 198.133.219.0/24 address blocks used in the examples in this document are reserved for the exclusive use of Cisco Systems, Inc.

---

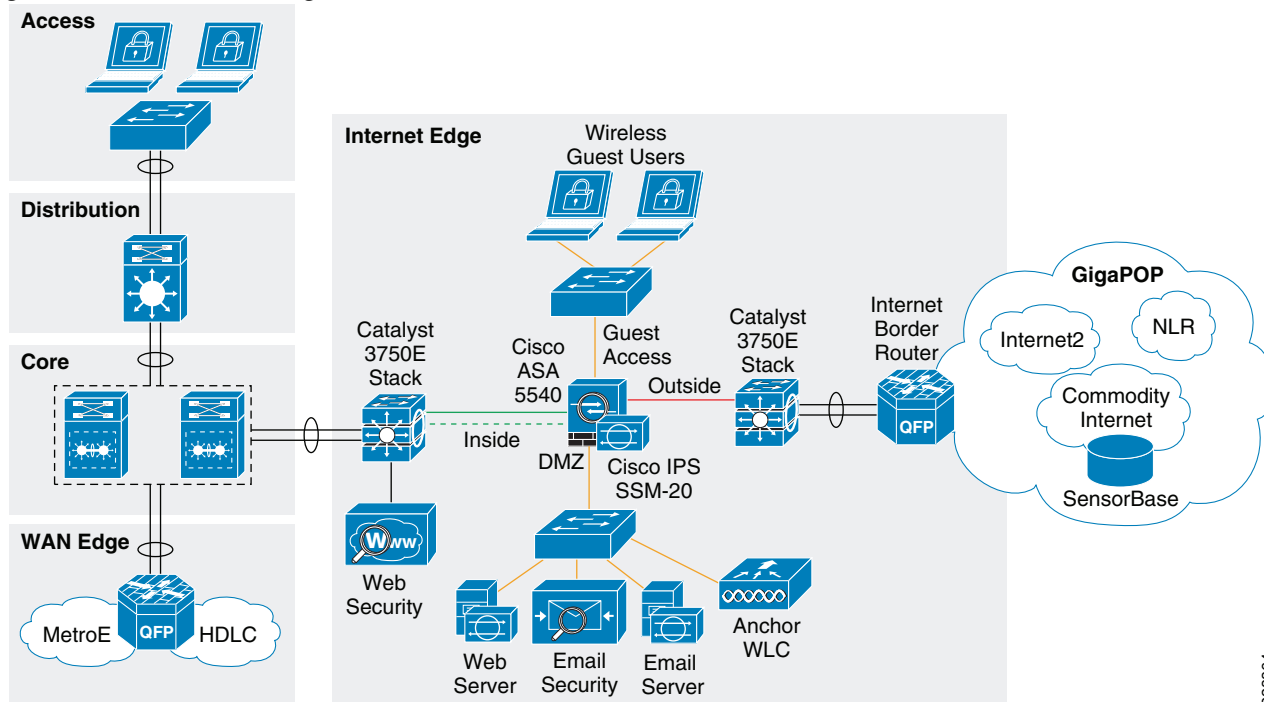
For more information and configuration examples on how to secure the Internet border router using the other Network Foundation Protection features, see “Chapter 6, Enterprise Internet Edge” in the *Cisco SAFE Reference Guide* at the following URL:

[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 community college'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 community college design uses a Cisco ASA appliance for the Internet Firewall as illustrated in [Figure 6-22](#).

Figure 6-22 Internet Edge Firewall



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.
- *Outside*—Interface connecting to the Internet border router. The router may be managed either by the community college or a service provider.
- *Demilitarized Zone (DMZ)*—The DMZ hosts the community college's public facing services that are accessible over the Internet, NLR or Internet2. These services may include a web portal and E-mail services.
- *Guest Access*—The interface connecting to the LAN segment that wireless guest access clients will be placed in by the Anchor WLC in the DMZ. Wireless guest access clients should only have access to the Internet, NLR and Internet-2 networks.

The Internet firewall acts as the primary gateway to the Internet, NLR and Internet2 networks; 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 6-49).

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 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 ASA configuration fragment illustrates the AAA configurations needed to authenticate, authorize, and log user access to the firewall:

```
aaa-server tacacs-servers protocol tacacs+
 reactivation-mode timed
aaa-server tacacs-servers host <ACS-Server>
 key <secure-key>
aaa authentication ssh console tacacs-servers LOCAL
aaa authentication serial console tacacs-servers LOCAL
aaa authentication enable console tacacs-servers LOCAL
aaa authentication http console tacacs-servers LOCAL
aaa authorization command tacacs-servers LOCAL
aaa accounting ssh console tacacs-servers
aaa accounting serial console tacacs-servers
aaa accounting command tacacs-servers
aaa accounting enable console tacacs-servers
aaa authorization exec authentication-server
! define local username and password for local authentication fallback
username admin password <secure-password> encrypted privilege 15
```

As with the other infrastructure devices in the network, it is important to synchronize the time on the firewall protecting the management module using NTP.

The following configuration fragment illustrates the NTP configuration needed on an ASA to enable NTP to an NTP server located in the management network:

```
ntp authentication-key 10 md5 *
ntp authenticate
ntp trusted-key 10
ntp server <NTP-Server-address> source management
```

Syslog and SNMP can be used to keep track of system status, device access and session activity. NetFlow Security Event Logging (NSEL), now supported on all Cisco ASA models, may also be used for the monitoring and reporting of session activity. The following configuration fragment illustrates the configuration of Syslog.

```
logging trap informational
logging host management <Syslog-Server-address>
logging enable
```

The routing protocol running between the Internet firewall and the distribution should be secured. The following ASA configuration fragment illustrates the use of EIGRP MD5 authentication to authenticate the peering session between the inside firewall interface and the 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 the community college typically gets a limited number of public IP addresses. In addition, NAT helps shield the community college'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, NLR, and Internet2 networks
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:

- Protect community colleges internal resources and data from external threats by preventing incoming access from the Internet, Internet2 and NLR networks.
- 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-, Internet2- and NLR-bound traffic.
- Prevent wireless guest access users from accessing internal resources.

Enforcing such policies requires 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 (i.e., from inside to outside). However, due to the sensitivity of community college environments, the community college administration may opt to override the default rules with more stringent rules indicating exactly what ports and protocols are permitted.

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 students, staff, and faculty residing at all campus sites to access the Internet, Internet2, and NLR networks for the allowed ports and protocols. Depending on the policy of the community college, this may only allow HTTP and HTTPS access or may be less restrictive to allow additional protocols and ports. The following example only allows HTTP and HTTPS access:

```
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 students, staff, and faculty access to DMZ services such as the community college web portal, E-mail, and domain name resolution. This could include HTTP, HTTPS, SMTP, POP, IMAP, and DNS protocols. Permit tunneled control and user traffic (UDP 16666, UDP 16667, IP Protocol 97) from internal WLCs to Anchor WLC in DMZ for wireless guest access. Permit management traffic (SNMP, SSH, and HTTPS) from management segment to Anchor WLC in DMZ. 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, etc.

```
! 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 Internet access should be restricted to the public services provided at the DMZ such as SMTP, web, and DNS. Any connection attempts to internal resources and subnets from the Internet should be blocked. ACLs should be constructed using the servers' global IP addresses.

```
! Allow DNS queries and zone transfer to DNS server
access-list inbound extended permit udp any host 198.133.219.35 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 school 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 only having access to the Internet, Internet2, and NLR networks. Access to the internal community college network should not be allowed.

```
! Deny access to internal networks
access-list guestaccess-acl extended deny ip 10.125.32.64 255.255.255.240 10.0.0.0
255.0.0.0
access-list guestaccess-acl extended deny ip 10.125.32.64 255.255.255.240
192.168.0.0 255.255.0.0
access-list guestaccess-acl extended deny ip 10.125.32.64 255.255.255.240
172.16.0.0 255.240.0.0
! Permit all other access to the Internet, Internet2 and NLR network
access-list guestaccess-acl extended permit ip 10.125.32.64 255.255.255.240 any
!
! Apply ACL to guest-access interface
access-group guestaccess-acl in interface guestaccess
```

## Firewall Redundancy

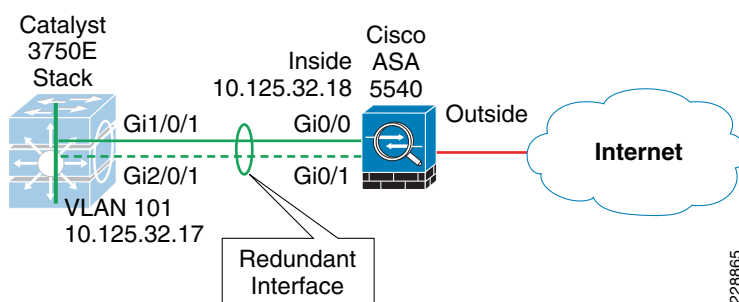
The Internet perimeter of the community college design uses a single Cisco ASA appliance configured with redundant interfaces. The use of redundant interfaces makes the design resilient to link-level failures, representing an affordable option for high availability. In cases where chassis redundancy is desirable, the community college 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, 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 6-23 illustrates the concept of redundant interface.

**Figure 6-23 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 redundant interface. The following configuration example corresponds to Figure 6-23.

```
! 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 back 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.
```

228865

```

! Configure 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 community college network design, an interior gateway protocol, EIGRP, is used for dynamic routing. The Internet firewall may participate in routing by learning the internal routes within the community college 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 community college design, 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 6-24](#).

**Figure 6-24 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 will be 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 will not help if the connection is lost upstream. If available, you may want to monitor a persistent network object located somewhere in the ISP network. Static route tracking can also be configured for default routes obtained through DHCP or PPPoE.

In the following configuration, the IP address of the next hop gateway (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 community college 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 will require coordination with the service provider. This scenario also requires the default route to be propagated over OSPF. The actual default route may originate from the Internet border router itself or somewhere in the ISP network.

## Botnet Traffic Filter

The Community College design 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:

- 
- Step 1** Configure DNS Server.
  - Step 2** Enable Use of the Dynamic Database.
  - Step 3** Enable DNS Snooping.
  - Step 4** Enable Traffic Classification and Actions for the Botnet Traffic Filter.
  - Step 5** Verify and Monitor Botnet Traffic Filter Operation
- 

The following sections provides configuration examples for each of these steps.

## Configure DNS Server

The Botnet Traffic Filter requires a DNS server to access Cisco's dynamic database update server and to resolve entries in the static database. The following configuration illustrates this configuration:

```
! Enable DNS requests to a DNS Server out the outside interface
dns domain-lookup outside
! Specify the DNS Server Group and the DNS Servers
dns server-group DefaultDNS
  name-server 68.238.112.12
  name-server 68.238.96.12
  domain-name cisco.com
```

## Enable Use of the Dynamic Database

The Botnet Traffic Filter can receive periodic updates for the dynamic database from the Cisco update server. This database lists thousands of known bad domain names and IP addresses. The following configuration enables database updates, and also enables use of the downloaded dynamic database by the adaptive security appliance.

```
! enable downloading of the dynamic database from the Cisco Update server
dynamic-filter updater-client enable
! enable use of the dynamic database
dynamic-filter use-database
```

## Enable DNS Snooping

DNS Snooping enables inspection of DNS packets and enables Botnet Traffic Filter Snooping, which compares the domain name with those in the dynamic or static databases and adds the name and IP address to the DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

It is recommended that DNS Snooping is only enabled on interfaces where external DNS requests are going. Enabling DNS Snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the 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 lookmbox.com
Sample meta data from downloaded database:
threat-level: very-high, category: Malware,
```

```

description: "These are sources that use various exploits to deliver adware, spyware
and other malware to victim computers. Some of these are associated with rogue online
vendors and distributors of dialers which deceptively call premium-rate phone
numbers."
threat-level: high, category: Bot and Threat Networks,
description: "These are rogue systems that control infected computers. They are
either systems hosted on threat networks or systems that are part of the botnet
itself."
threat-level: moderate, category: Spyware,
description: "These are sources that distribute spyware, adware, greyware, and other
potentially unwanted advertising software. Some of these also run exploits to install
such software."
threat-level: low, category: Ads,
description: "These are advertising networks that deliver banner ads, interstitials,
rich media ads, pop-ups, and pop-unders for websites, spyware and adware. Some of
these networks send ad-oriented HTML emails and email verification services."
Total entries in Dynamic Filter database:
Dynamic data: 82119 domain names , 2565 IPv4 addresses
Local data: 0 domain names , 0 IPv4 addresses
Active rules in Dynamic Filter asp table:
Dynamic data: 0 domain names , 2565 IPv4 addresses
Local data: 0 domain names , 0 IPv4 addresses
cr12-asa-1-ie#

```

- **show dynamic-filter statistics detail**—Shows how many connections were monitored and dropped with the Botnet Traffic Filter, and how many of those connections match the whitelist, blacklist, and greylist. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.) The detail keyword shows how many packets at each threat level were classified or dropped.

```

cr12-asa-1-ie# show dynamic-filter statistics detail
Enabled on interface outside using classify-list btf-filter-acl
Total conns classified 35, ingress 0, egress 35
Total whitelist classified 0, ingress 0, egress 0
Total greylist classified 16, dropped 0, ingress 0, egress 16
Threat-level very-high: classified 0, dropped 0, ingress 0,
egress 0
Threat-level high: classified 0, dropped 0, ingress 0,
egress 0
Threat-level moderate: classified 0, dropped 0, ingress 0,
egress 0
Threat-level low: classified 16, dropped 0, ingress 0,
egress 16
Threat-level very-low: classified 0, dropped 0, ingress 0,
egress 0
Total blacklist classified 19, dropped 0, ingress 0, egress 19
Threat-level very-high: classified 9, dropped 0, ingress 0,
egress 9
Threat-level high: classified 0, dropped 0, ingress 0,
egress 0
Threat-level moderate: classified 0, dropped 0, ingress 0,
egress 0
Threat-level low: classified 10, dropped 0, ingress 0,
egress 10
Threat-level very-low: classified 0, dropped 0, ingress 0,
egress 0
cr12-asa-1-ie#

```

**Note**

To clear the statistics, enter the **clear dynamic-filter statistics** [*interface name*] command.

Other commands that are useful for monitoring the Botnet Traffic Filter include the following:

- **show dynamic-filter reports top [malware-sites | malware-ports | infected-hosts]**—Generates reports of the top 10 malware sites, ports, and infected hosts monitored. The top 10 malware-sites report includes the number of connections dropped, and the threat level and category of each site. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected.
- **show dynamic-filter reports infected-hosts {max-connections | latest-active | highest-threat | subnet ip\_address netmask | all}**—Generates reports about infected hosts. These reports contain detailed history about infected hosts, showing the correlation between infected hosts, visited malware sites, and malware ports. The max-connections keyword shows the 20 infected hosts with the most number of connections. The latest-active keyword shows the 20 hosts with the most recent activity. The highest-threat keyword shows the 20 hosts that connected to the malware sites with the highest threat level. The subnet keyword shows up to 20 hosts within the specified subnet. The all keyword shows all buffered infected-hosts information. This display might include thousands of entries. You might want to use ASDM to generate a PDF file instead of using the CLI.
- **show dynamic-filter dns-snoop [detail]**—Shows the Botnet Traffic Filter DNS Snooping summary, or with the detail keyword, the actual IP addresses and names. All inspected DNS data is included in this output, and not just matching names in the blacklist. DNS data from static entries are not included.
- **show asp table dynamic-filter [hits]**—Shows the Botnet Traffic Filter rules that are installed in the accelerated security path.

## IPS Global Correlation Deployment

Within the Community College design, IPS is implemented within the Internet edge using a Security Service Module in the Cisco ASA Internet Firewall. IPS is implemented using the IPS Version 7.0 Global Correlation feature. 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 will add risk to the event, raising its risk



rating and making it more likely that the sensor will deny the event. This enables the sensor to deny packets and attackers based on the fact that the event has a negative reputation in addition to a high risk rating calculated on the sensor.

## How IPS with Global Correlation Works

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.

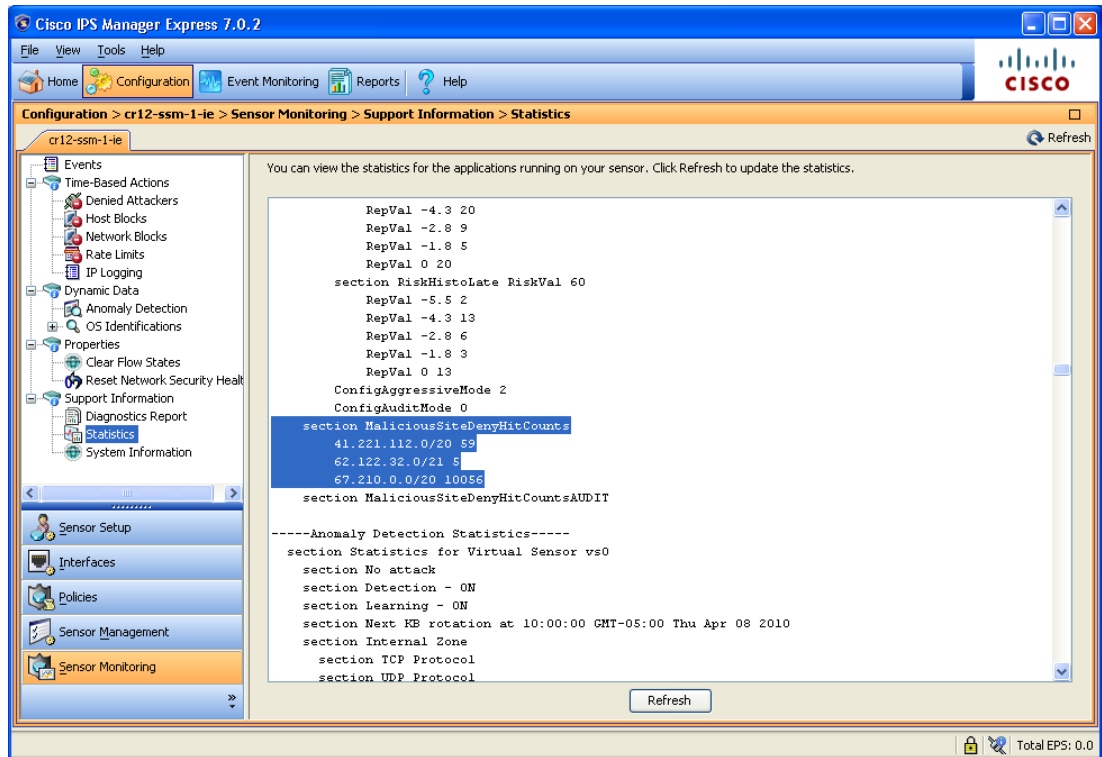
### Reputation Filters

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 prior to signature inspection.

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), as shown in [Figure 6-25](#).

**Figure 6-25 Cisco IME Statistics Section Showing Networks Currently Being Logged by IPS Reputation Filters**



The only user configuration required for reputation filters is enabling or disabling them and specifying whether Global Correlation is set to Audit mode (a global configuration setting for the entire sensor). In Audit mode, the sensor will report potential deny actions due to reputation filters instead of actually denying the activity.

## Global Correlation Inspection

The primary activity of a 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, then it is considered to be neutral. Cisco IPS, unlike E-mail and web security reputation applications, has no concept of positive reputation. When an event is checked for reputation, this checking occurs entirely on the sensor using data downloaded previously from Cisco SensorBase. Unlike other devices, the sensor will not send a live request for information about an IP address that it has just seen. It will look in the data that it has, and if it finds the address, it will use that data; otherwise, the sensor will assume that the address has a neutral reputation.

Global Correlation inspection has three modes of primary operation: permissive, standard (default), and aggressive; you can also select Off:

- Permissive mode tells the sensor to adjust the risk rating of an event, but not to assign separate reputation-only actions to the event.
- Standard mode tells the sensor to adjust the risk rating and to add a Deny Packet action due to reputation if the risk rating is greater than or equal to 86. It also adds a Deny Attacker action due to reputation if the risk rating is greater than or equal to 100.
- Aggressive mode also adjusts the risk rating due to reputation, adds a Deny Packet action due to reputation if the risk rating is greater than or equal to 83, and adds a Deny Attacker action due to reputation if the risk rating is greater than or equal to 95.
- Selecting Off in the Global Correlation Inspection window prevents the sensor from using updates from Cisco SensorBase to adjust reputation.

If Global Correlation inspection is enabled and an event is generated by an attacker with a negative reputation, the risk rating for the event will be elevated by a certain amount that is determined by a statistical formula. The amount by which the risk rating is raised depends on the original risk rating of the event and the reputation of the attacker.

### Network Participation and Correlation Updates

The sensor pulls reputation information for addresses on the global Internet from Cisco SensorBase. When the sensor is configured initially, a DNS server needs to be configured for the sensor to use to connect to Cisco SensorBase or an HTTP or HTTPS proxy (that has DNS configured) needs to be configured. After the sensor has this information, the sensor will make an outbound connection to check for the latest updates from Cisco SensorBase. It will initiate an HTTPS request to Cisco SensorBase update servers and download a manifest that contains the latest versions of the files related to Global Correlation. The sensor will check Cisco SensorBase every 5 minutes for updates. If changes are needed, the sensor will perform a DNS lookup of the server name returned in the initial request. This lookup will return the location of the server nearest to the sensor. The sensor will then initiate an HTTP connection that will actually transfer the data. The size of a full update is about 2 MB; incremental updates average about 100 KB. If a sensor loses connection to Cisco SensorBase, Global Correlation information will begin to time out within days, and sensor health will change accordingly.

The other component of Global Correlation is network participation. This feature sends data from events that the sensor fires back to Cisco SensorBase to adjust the reputation of IP addresses; this information is then packaged in future reputation data downloads from Cisco SensorBase. The sensor passes this information back to Cisco SensorBase according to the sensor configuration. The possible configuration options are *Off*, *Partial*, and *Full*.

- With the *Off* (default) setting, the sensor will not send back any data. The sensor will still receive reputation data, and this setting does not affect its use of that data except that the reputations of addresses attacking the network being protected will not be influenced by their generation on the sensor.
- With the *Partial* setting, the sensor will send back alert information. This information consists of protocol attributes such as the TCP maximum segment size and TCP options string, the signature ID and risk rating of the event, the attacker IP address and port, and Cisco IPS performance and deployment mode information.
- The *Full* setting adds victim IP address and port information to the information reported with the *Partial* setting.



#### Note

No actual packet content information is sent to Cisco. In addition, events having RFC 1918 addresses, because they are not unique, are not considered interesting. So all events reported to Cisco SensorBase will have any such IP address information stripped from the reported data.

The mechanism used to update Cisco SensorBase with new attack information is fairly straightforward. The sensor takes event information, parses out the important pieces of data, and buffers this data for transmission back to Cisco SensorBase. It sends this data in the form of an HTTPS connection that it initiates on average every 10 minutes. The average size of an update is 2 to 4 KB, with weekly averages of about 0.5 to 1 MB. Some higher-volume sensors have average update sizes of about 50 KB, with weekly totals in the 45-MB range. Sensors with very high alert volumes can have average update sizes of about 850 KB, with weekly totals of up to 900 MB; these sensors, though, are at the extreme end of the range.

## IPS Global Configuration Overview

Before configuring 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 first step in configuring the sensor to use Global Correlation is to add either a DNS address and/or the proxy server setup. This step enables connection to Cisco SensorBase. By default, a sensor runs Global Inspection in standard mode, enables the reputation filters, and does not allow network participation. After you configure the DNS and proxy settings, these settings will go into effect as soon as the sensor has downloaded the latest Global Correlation updates.

The Configuration of Global Correlation can be performed in various ways using the command-line interface (CLI), Cisco IDS Device Manager (IDM), Cisco IME, or Cisco Security Manager. [Figure 6-26](#), [Figure 6-27](#), and [Figure 6-28](#) illustrate screenshots from Cisco IME to configure IPS Global Correlation.

**Figure 6-26** DNS and HTTP Proxy Within the Network Setting Configuration Screen

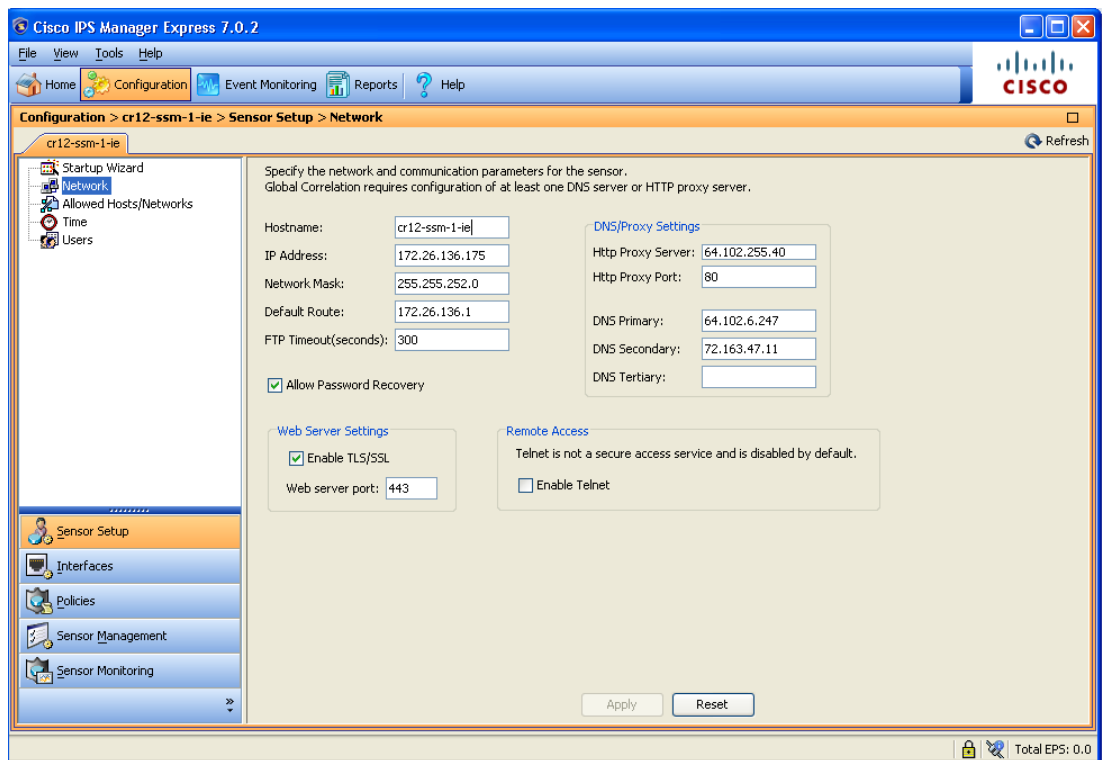


Figure 6-27 Global Correlation Inspection Settings

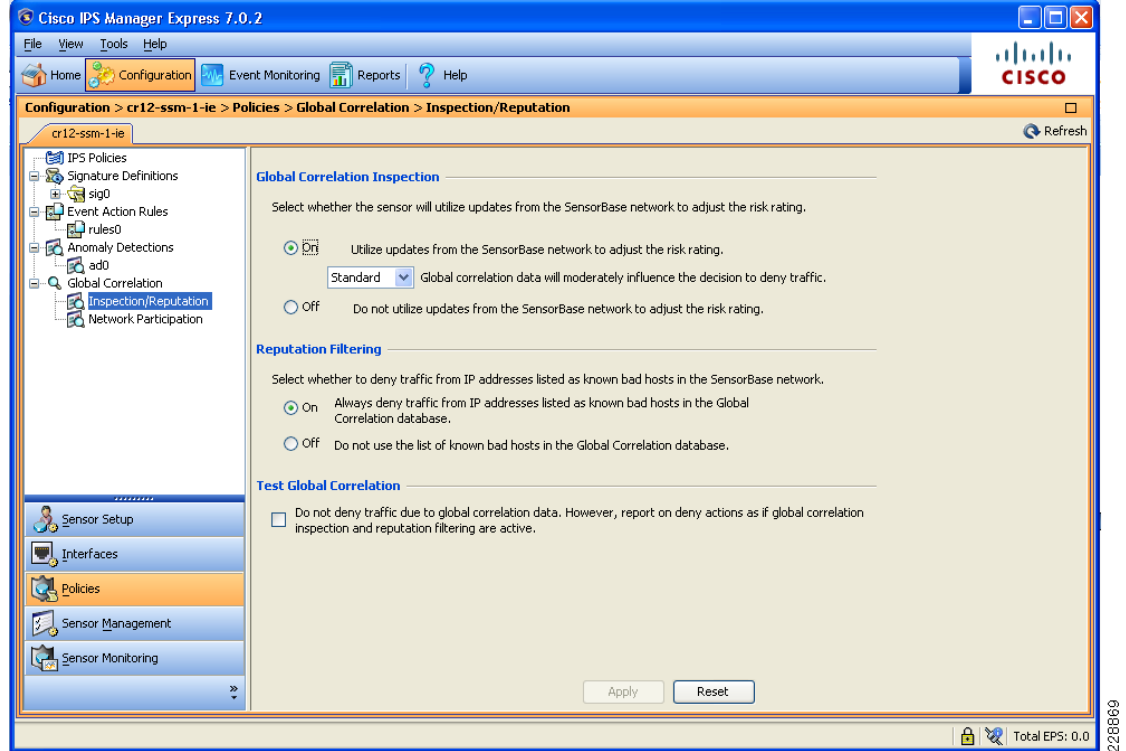
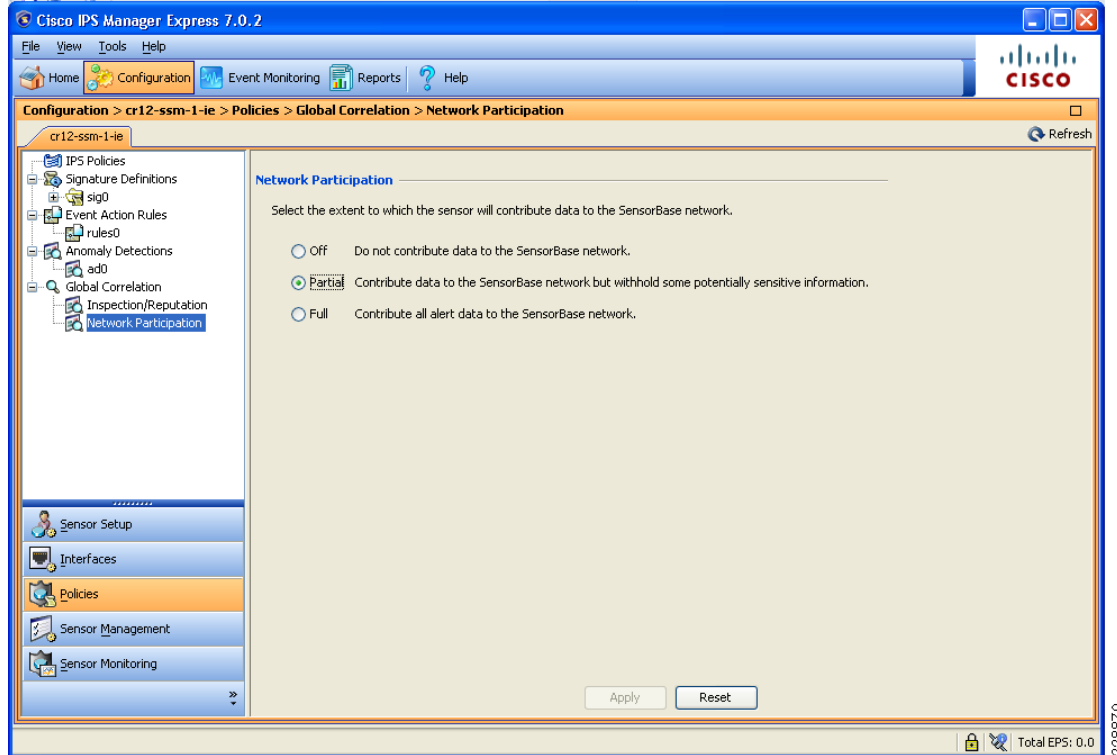


Figure 6-28 Network Participation Settings (Off by Default)



## Event Monitoring with Global Correlation

Event monitoring with Global Correlation is similar to event monitoring with signature-only IPS. The primary difference is the potential addition of reputation scores representing the Global Correlation data. Figure 6-29 shows Cisco IPS events with reputation scores in Cisco IME.

Figure 6-29 Event Monitoring with Global Correlation in Cisco IME

The screenshot displays the Cisco IPS Manager Express 7.0.2 Event Monitoring interface. The main window shows a table of events with the following columns: Severity, Date, Time, Device, Sig. Name, Sig. ID, Attacker IP, Victim IP, Action, Threat, Risk Rating, and Reputation. The table contains several rows of events, including TCP SYN Port Sweep and ICMP Network Sweep attacks. The events are sorted by date and time, showing a sequence of attacks from 11:37:56 to 11:54:04. The risk ratings and reputations vary, with some events having negative reputations (-1.8 and -4.3) and higher risk ratings (70, 75, 86, 89) compared to others (52, 60).

Severity	Date	Time	Device	Sig. Name	Sig. ID	Attacker IP	Victim IP	Action	Threat	Risk Ra...	Reput...
low	03/30/2...	11:37:56	cr12-ssm...	TCP SYN Port Sweep	3002/0	78.0.248.20	10.125.32.34			52	52
low	03/30/2...	11:38:22	cr12-ssm...	ICMP Network Sweep w/Echo	2100/0	78.0.248.20	10.125.32.34			60	60
low	03/30/2...	11:38:29	cr12-ssm...	TCP SYN Port Sweep	3002/0	78.0.248.20	10.125.32.42			52	52
low	03/30/2...	11:46:40	cr12-ssm...	TCP SYN Port Sweep	3002/0	87.88.138.14	10.125.32.34			70	70
low	03/30/2...	11:47:06	cr12-ssm...	ICMP Network Sweep w/Echo	2100/0	87.88.138.14	10.125.32.34			75	75
low	03/30/2...	11:47:13	cr12-ssm...	TCP SYN Port Sweep	3002/0	87.88.138.14	10.125.32.42			70	70
low	03/30/2...	11:53:31	cr12-ssm...	TCP SYN Port Sweep	3002/0	109.111.132.165	10.125.32.34	dropped...		51	86
low	03/30/2...	11:53:57	cr12-ssm...	ICMP Network Sweep w/Echo	2100/0	109.111.132.165	10.125.32.34	dropped...		54	89
low	03/30/2...	11:54:04	cr12-ssm...	TCP SYN Port Sweep	3002/0	109.111.132.165	10.125.32.42	dropped...		51	86

Figure 6-29 shows several TCP SYN Port Sweep and ICMP Network Sweep attacks that were seen by the sensor. The first three events had no reputation, and the event's risk ratings were 52 and 60 which did not meet the threshold for the packets to be dropped. The next three events were identical except that the attacker had a negative reputation of -1.8, which elevated the risk ratings to 70 and 75 which still did not meet the thresholds to be dropped in Standard Mode. The last events were also identical except this time the attacker has a negative reputation 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 6-30 illustrates the detail view of the TCP SYN Port Sweep event coming from the attacker with a negative reputation of -4.3.

**Figure 6-30** Detail View of a TCP SYN Port Sweep from an Attacker with a Negative Reputation Score

The screenshot shows a window titled "Event Details (Event ID - 1266963033383260376)". The window contains a table of event details. At the bottom of the window, there is a note: "You can **copy** selected or all rows into clipboard or **print** the entire contents."

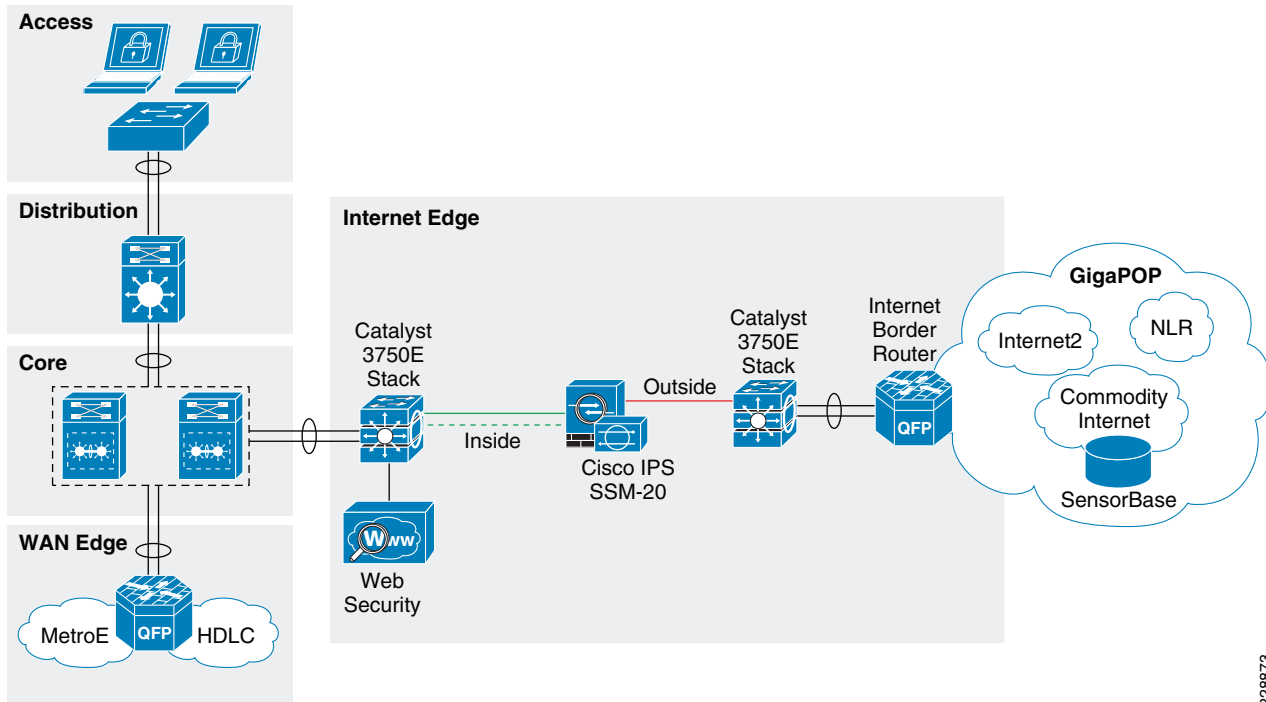
Event ID	1266963033383260376
Severity	low
Host ID	cr12-ssm-1-ie
Application Name	sensorApp
Event Time	03/30/2010 11:53:31
Sensor Local Time	03/30/2010 15:53:31
Signature ID	3002
Signature Sub-ID	0
Signature Name	TCP SYN Port Sweep
Signature Version	52
Signature Details	
Interface Group	vs0
VLAN ID	0
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	

## Web Security Deployment

The Community College design implements a Cisco IronPort WSA at the Internet edge distribution layer at the main campus, as illustrated in [Figure 6-31](#). 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 6-31 WSA Deployment



228873

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 some of the initial settings cannot be changed afterwards without resetting the appliance's configuration to its factory defaults. Care should be taken in choosing the right configuration options. Plan not only for the features to be implemented immediately, but also for what might be required in the future.

Some guidelines when running the system setup wizard include the following:

- *Deployment Options*—Step 2 of the wizard gives the user the options to enable only Layer-4 (L4) Traffic Monitoring, enable only Secure Web Proxy, or enable both functions. Select enable both **Secure Web Proxy** and **L4 Traffic Monitor**, if you plan to use both functions.
- *Proxy Mode*—If the Secure Web Proxy function has been enabled, Step 2 of the wizard requires the user to choose between Forward and Transparent modes. Note that if a WSA appliance is initially configured in Transparent mode, it can still be configured as a Forward Web Proxy. However, if the WSA is initially configured in Forward mode, the Transparent Web Proxy function is not available without having to reset the WSA to factory defaults. Therefore, select Forward mode only if you are certain that the Transparent mode will never be required.



### Note

The deployment and proxy mode options cannot be changed after the initial configuration without resetting the WSA appliance to its factory defaults. Plan your configuration carefully.

## Interface and Network Configuration

As part of the initial setup of the WSA, the following steps need to be completed:


- 
- Step 1** Configuring network interfaces.
  - Step 2** Adding routes.
  - Step 3** Configuring DNS.
  - Step 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

Independently from the model, all Cisco IronPort WSA appliances are equipped with six Ethernet interfaces as shown in [Figure 6-32](#).

**Figure 6-32** WSA Interfaces

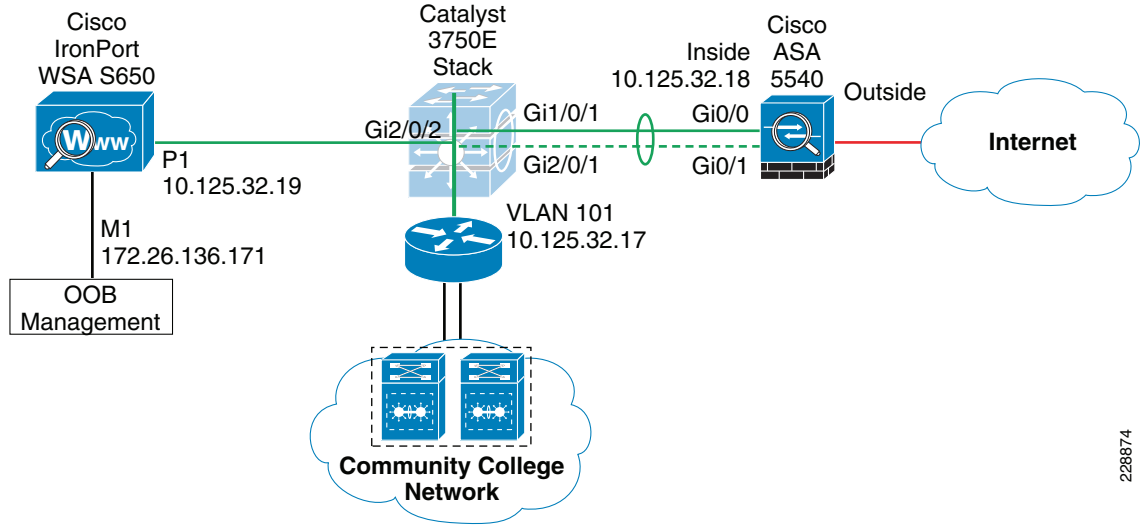


The WSA interfaces are grouped for the following functions:

- *Management*—Interfaces M1 and M2 are out-of-band (OOB) management interfaces. However, only M1 is enabled. In the community college design, interface M1 connects to the out-of-band management network. Interface M1 can optionally be used to handle data traffic in case the community college does not have an out-of-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 community college design. 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 community college design. For more information, please refer to the WSA configuration guides.

[Figure 6-33](#) illustrates the network topology for the WSA design in the validation lab.

**Figure 6-33 WSA Network Topology**



228874

Figure 6-34 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 6-34 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)			

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 6-35.

**Figure 6-35 WSA Route Configuration****Routes**

The screenshot displays two route configuration windows. The top window is titled "Routes for Management Traffic (Interface M1: 172.26.136.171, Interface P1: 10.125.32.19)". It contains a table with the following data:

Name	Destination Network	Gateway	All <input type="checkbox"/> Delete
Default Route	All Others	172.26.136.1	<input type="checkbox"/>

The bottom window is titled "Routes for Data Traffic (Interface P1: 10.125.32.19)". It contains a table with the following data:

Name	Destination Network	Gateway	All <input type="checkbox"/> Delete
Default Route	All Others (Including External)	10.125.32.18	<input type="checkbox"/>
Internal-10	10.0.0.0/8	10.125.32.17	<input type="checkbox"/>

228876

**Configuring DNS**

The initial setup requires the configuration of a host name for the WSA appliance, and defining the DNS servers. [Figure 6-36](#) shows the DNS configuration.

**Figure 6-36 WSA DNS Configuration****DNS**

The screenshot shows the "DNS Server Settings" window. It includes the following configuration details:

- 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

Buttons for "Clear DNS Cache" and "Edit Settings..." are visible at the bottom.

228877

**Setting Time**

Time synchronization is critical for forensic analysis and troubleshooting, therefore enabling NTP is highly recommended. [Figure 6-37](#) shows how the WSA is configured to synchronize its clock with an NTP server located on the OOB management network.

**Figure 6-37 WSA NTP Configuration****Time Settings**

Time Setting	
Time Keeping Method:	Using NTP Servers:
1	172.26.129.252
Routing Table for NTP Server Queries: Management	
<a href="#">Edit Settings...</a>	

**Note**

If Internet access is provided by an upstream proxy, then the WSA must be configured to use the proxy for component updates and system upgrades. For information on configuring upstream proxies for upgrades, refer to 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:

- Step 1** Defining WSA WCCP Service Group.
- Step 2** Enabling WSA Transparent Redirection.
- Step 3** Enabling WCCP redirection on the Cisco ASA.
- Step 4** Enabling WSA HTTPS scanning.

### 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*—List the ports to be proxied. Default is 80 and 3128.
- *Caching*—Defines whether or not the WSA should cache response and requests. Caching helps reduce latency and the load on the Internet links. Default is enabled.
- *IP Spoofing*—Defines whether or not the Web Proxy should spoof IP addresses when forwarding requests to upstream proxies and servers.

Figure 6-38 illustrates the Web Proxy settings.

Figure 6-38 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 Server Side: 300 Seconds
In-Use Connection Timeout:	Client Side: 300 Seconds Server Side: 300 Seconds
Simultaneous Persistent Connections:	Server Maximum Number: 2000
Headers:	X-Forwarded-For: Do Not Send 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 Switch. Use MD5 authentication to protect the WCCP communication between the WSA and Cisco Catalyst Switch. Figure 6-39 shows an example.

Figure 6-39 WSA Transparent Proxy

WCCP v2 Service	
Service Profile Name:	web-https-cache
Service:	<input type="radio"/> Standard service ID: 0 web-cache (destination port 80) <input checked="" type="radio"/> Dynamic service ID: <input type="text" value="10"/> 0-255 Port numbers: <input type="text" value="80,443"/> <small>(up to 8 port numbers, separated by commas)</small> <input checked="" type="radio"/> Redirect based on destination port <input type="radio"/> Redirect based on source port (return path) <small>For IP spoofing, define two services, one based on destination port and another based on source port (return path).</small> <input checked="" type="radio"/> Load balance based on server address <input type="radio"/> Load balance based on client address <small>Applies only if more than one Web Security Appliance is in use.</small>
Router IP Addresses:	10.125.32.17 <small>Separate multiple entries with line breaks or commas.</small>
Router Security:	<input checked="" type="checkbox"/> Enable Security for Service Password: <input type="password" value="....."/> Confirm Password: <input type="password" value="....."/>
▶ Advanced:	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.

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 6-40](#).

**Figure 6-40 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: Common name: Cisco Systems, Inc Organization: CMO Organizational Unit: ESE Country: US Expiration Date: Oct 14 16:30:47 2010 GMT Basic Constraints: Not Critical
Invalid Certificate Handling:	Expired: Monitor Mismatched Hostname: Monitor Unrecognized Root Authority: Monitor All other error types: Monitor

228881

**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, refer to 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 community college can control what Internet applications (instant messaging clients, peer-to-peer file-sharing, web browsers, Internet phone services, etc.) and URL categories students, staff and faculty 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 6-41](#) shows the global policy.



**Figure 6-41 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> Identity: All	Allow: FTP over HTTP, HTTP, Native FTP Block: User Agents Allow: Ports 8080, 21,...	Redirect: 0 Allow: 0 Monitor: 51 Warn: 0 Block: 3 Time-Based: 0	HTTP/HTTPS Object Max Size: None FTP Object Max Size: None	(enabled)	

228882

URL categories corresponding to inappropriate content should be blocked in compliance with the community college's Internet access policies. [Figure 6-42](#) provides an example on how the **Adult/Sexually Explicit** and **Chat** categories are blocked.

**Figure 6-42 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 🟡	Warn ⚠️	Block 🛑	Time-Based 🕒	
🛑 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 Community College design, the Cisco Catalyst Security features were implemented in the access layer switches to protect the infrastructure and users from spoofing, man-in-the-middle (MITM), DoS, and other access layer attacks. The following configurations illustrates an example of the CISF configurations used on a Cisco 3750 switch in the Community College design.

```
! 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 parameters
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 Catalyst 3K switches in the access layer in a routed Layer 3 deployment, configuring IP Source Guard will cause 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 only permit traffic from IP addresses that were assigned via the DHCP server. On Catalyst 3K switches, port-based ACLs overrides router and VLAN ACLs resulting in all traffic being permitted to all destinations.

---

## NAC Deployment

Within the Community College design, a NAC Appliance solution is deployed at each of the site types; main campus, remote large campus, remote medium campus, and remote small campus. A centralized NAC Manager is deployed at the main campus and is deployed within the data center at that site. A NAC server is deployed at each of the campus sites (main and remote sites) and is connected within the service block connecting to the core switches at each of the sites.

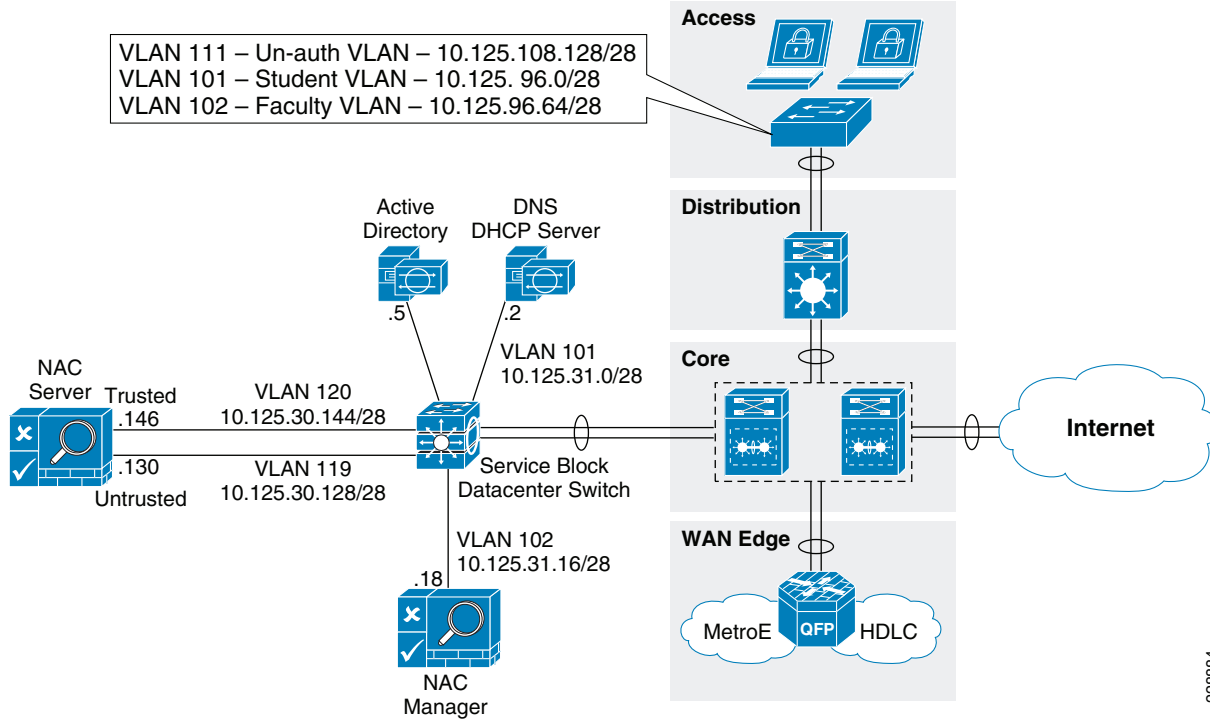
The Community College reference design provides host network connectivity using wired and wireless technologies. As such, the NAC Appliance solution must provide a solution for both connectivity options. For wireless clients, a Layer 2 OOB NAC solution is deployed, and for wired clients, a Layer-2 OOB 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 Community College design, a NAC Layer-3 OOB deployment using ACLs was used for the wired clients. [Figure 6-43](#) shows the L3 OOB logical network diagram which was used to validate NAC for wireless clients in the Community College design.

Figure 6-43 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 below configuration snippets provides sample configurations for three VLANs (Unauthenticated, Student, and Faculty) 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

```

- Student Role: VLAN 101 and ACL name: student-acl

```

! create NAC student VLAN
vlan 101
name student-vlan
! create SVI for student VLAN
interface Vlan101
ip address 10.125.96.1 255.255.255.192
ip helper-address 10.125.31.2
! configure ACL for the student role
ip access-list extended nac-student-acl
! 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
! allow web and SSL traffic to the DMZ subnet
permit tcp any 10.125.32.32 0.0.0.15 eq www
permit tcp any 10.125.32.32 0.0.0.15 eq 443
! allow web and SSL traffic to select internal student accessible resources
permit tcp any 10.125.31.113 0.0.0.15 eq www
permit tcp any 10.125.31.145 0.0.0.15 eq www
permit tcp any 10.125.31.113 0.0.0.15 eq 443
permit tcp any 10.125.31.145 0.0.0.15 eq 443
! deny access to rest of internal resources
deny ip any 10.0.0.0 0.255.255.255 log
deny ip any 192.168.0.0 0.0.255.255
deny ip any 172.16.0.0 0.15.255.255
! allow access to the Internet
permit ip any any
!
! apply ACL to the Student VLAN
interface Vlan101
ip access-group nac-student-acl in

```

- Faculty Role: VLAN 102 and ACL name: faculty-acl

The existing production VLAN and ACL can be used for the faculty NAC role. Once 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 faculty 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 faculty VLAN.

```
! create NAC faculty VLAN
vlan 102
 name faculty-vlan
! create SVI for faculty VLAN
interface Vlan10
 ip address 10.125.96.65 255.255.255.192
 ip helper-address 10.125.31.2
! configure ACL for the faculty role to prevent NAC Discovery packets from
reaching NAC Server
ip access-list extended faculty-acl
 deny udp any host 10.125.30.130 eq 8906
 permit ip any any
!
! apply ACL to the Faculty VLAN
interface Vlan101
 ip access-group faculty-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. Refer to 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. Refer to 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 antivirus (AV) and antispymware (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. Refer to the following URI 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, refer to 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, refer to 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, then 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

To add the NAC Server to the NAC Manager, from within the NAC Manager GUI click on **CCA Servers > New Server**.

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 on **Add Clean Access Server**. See [Figure 6-44](#).

**Figure 6-44** Adding the NAC Server to the NAC Manager

Once added, the NAC Server will appear in the list.

**Note**

The NAC Manager and NAC Server have to trust each other's certificate authority (CA) in order for NAC Manager to successfully add the NAC server.

## Configure the NAC Server

Click the **Manage** icon for the NAC Server to continue the configuration. See [Figure 6-45](#).

**Figure 6-45** NAC Server Managed by NAC Manager

IP Address	Type	Location	Status	Manage	Disconnect	Reboot	Delete
10.125.30.146	Out-of-Band Real-IP Gateway	CCVE Main Campus DC	Connected				
10.125.30.114	Out-of-Band Virtual Gateway	Wireless L2 OOB	Connected				

After clicking the **Manage** icon, click the **Network** tab.

## Layer 3 Support

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

Click **Update** and reboot the NAC Server as instructed. [Figure 6-46](#).

**Figure 6-46** NAC Server Network Details

**Clean Access Server Type:** Out-of-Band Real-IP Gateway

Enable L3 support

Enable L3 strict mode to block NAT devices with NAC Agent

Enable L2 strict mode to block L3 devices with NAC Agent

**Platform:** APPLIANCE

**Trusted Interface (to protected network)**

IP Address: 10.125.30.146  
Subnet Mask: 255.255.255.240  
Default Gateway: 10.125.30.145

Set management VLAN ID: 0

Pass through VLAN ID to managed network

**Untrusted Interface (to managed network)**

IP Address: 10.125.30.130  
Subnet Mask: 255.255.255.240  
Default Gateway: 10.125.30.129

Set management VLAN ID: 0

Pass through VLAN ID to protected network

(Make sure the Clean Access Server is on VLAN n before you set its management VLAN ID to n.)

**Update** **Reboot**

**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.

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 6-47](#).

**Figure 6-47** Adding Static Route to Reach the Unauthenticated User Subnet

The screenshot shows the Cisco Clean Access Standard Manager web interface. The breadcrumb navigation is "Device Management > Clean Access Servers > 10.125.30.146". The "Static Routes" tab is selected under the "Advanced" section. The configuration form is as follows:

Status	Network	Filter	Advanced	Authentication	Misc
Managed Subnet	VLAN Mapping	<b>Static Routes</b>	ARP	Proxy	
Dest. Subnet Address/Mask	10.125.108.128 / 26		Gateway (optional)	10.125.30.129 <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				
<input type="button" value="Add Route"/>					

On the left side of the interface, there is 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.

**Setup Profiles for Managed Switches in the NAC Manager**

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 6-48](#), a Catalyst 4507 switch is added.



Figure 6-48 SNMP Profile used to Manage a Catalyst 4507 Switch

**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: CR22\_4507\_LB

Device Model: Cisco Catalyst 4000/4500 series

SNMP Port: 161

Description: Cisco 4507 L3 Access Switch in Ma

**SNMP Read Settings**

SNMP Version: SNMP V2C

Community String: cisco123

**SNMP Write Settings**

SNMP Version: SNMP V2C

Community String: ccve

Update Reset

2268669

## Switch Configuration for SNMP

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 ccve RW
snmp-server community cisco123 RO
```

## Configuring Port Profiles

For individual port control, configure a port profile under **OOB Management > Profiles > Ports** that includes the default unauthenticated VLAN and default access VLAN. In the access VLAN section, specify the User Role VLAN. The NAC Manager will change 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 6-49, the Auth VLAN is the unauthenticated VLAN to which unauthenticated devices are initially assigned. The default access VLAN is the student VLAN. This is used if the authenticated user does not have a role-based VLAN defined.

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 6-49 Port Profile to Manage the Switch Port

The screenshot shows the Cisco Clean Access Standard Manager interface. The left sidebar contains navigation menus for Device Management, OOB Management (with 'Profiles' selected), User Management, and Monitoring. The main content area is titled 'OOB Management > Profiles' and shows a table with columns for Group, Device, Port, VLAN, and SNMP Receiver. Below the table, the configuration for the 'Student\_Port' profile is displayed. The 'Manage this port' checkbox is checked. Under 'VLAN Settings', the 'Auth VLAN' is set to VLAN ID 111, the 'Default Access VLAN' is set to VLAN ID 101, and the 'Access VLAN' is set to 'User Role VLAN'. The 'VLAN Profile' is set to '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 campus, 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 6-50.

Figure 6-50 Various Options Available under Port Profile

The screenshot shows the Cisco Clean Access Standard Manager interface with the 'Options: Device Connected to Port' section expanded. The text explains that the CAM discovers the device and assigns the Auth VLAN or Access VLAN based on certification. Below this, several options are listed with checkboxes:
 

- 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.

 The 'Options: Device Disconnected from Port' section is also visible, with the following options:
 

- 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.

 An 'Update' button is located at the bottom of the configuration area.

## 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 6-51](#).

**Figure 6-51** NAC Manager SNMP Receiver Setting to Collect SNMP Traps/Informs

The screenshot shows the Cisco Clean Access Standard Manager web interface. The left sidebar contains navigation menus for Device Management, OOB Management, User Management, Monitoring, and Administration. 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 'SNMP Trap > Advanced Settings'. The configuration fields include:

- Trap Port on Clean Access Manager: 162
- SNMP V1 Settings: Community String (empty)
- SNMP V2c Settings: Community String: NacTraps
- SNMP V3 Settings: Security Method (Auth/Priv): No Auth, No Priv; User Name, User Auth, and User Priv (all empty)

An 'Update' button is located at the bottom of the configuration area. A vertical ID number '228692' is visible on the right edge of the interface.

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.

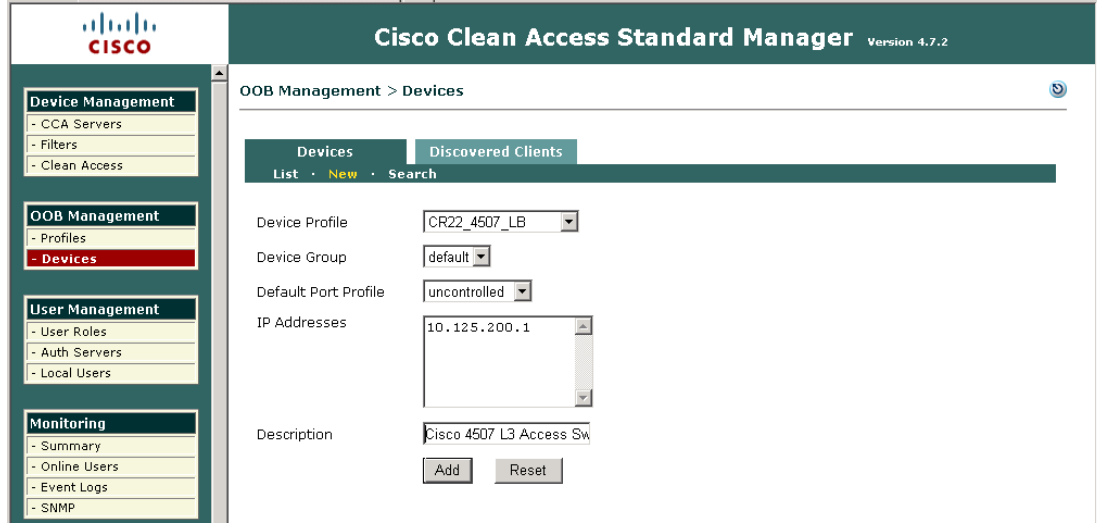
```
! global applicable SNMP configurations
snmp-server trap-source Loopback0
snmp-server enable traps mac-notification change move threshold
snmp-server host 10.125.31.18 version 2c NacTraps
! interface specific configurations
mac-address-table aging-time 3600
```

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.

## Add Switches as Devices in the NAC Manager

The switch profile created in the previous section will be used to add the managed switches. Under the **Device Profile**, use the profile you created, but do not change the default port profile value when adding the switch. See [Figure 6-52](#).

Figure 6-52 Adding Edge Switch in the NAC Manager to Control via SNMP

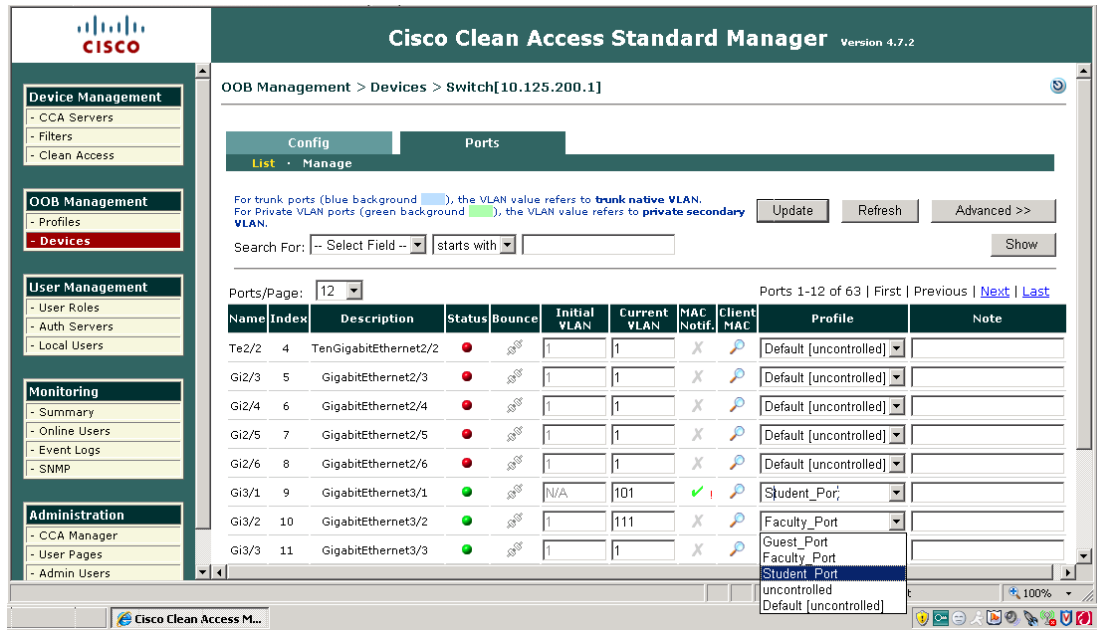


228893

## Configure Switch Ports for the Devices to be Managed by NAC

Once the switch is added to the NAC Manager, you can select the ports that you want to manage. See [Figure 6-53](#).

Figure 6-53 Port Control Selection available for a Managed Switch



228894

## Configure User Roles

The next step is to configure the user roles and map the appropriate VLANs to these roles. The screenshots in [Figure 6-54](#) and [Figure 6-55](#) depict the creation of two additional roles for the student and faculty clients. The VLANs were already created in the edge access switches which correspond to each role.

**Figure 6-54** Creating Student Role and Mapping to the Limited Access VLAN 101

The screenshot shows the Cisco Clean Access Standard Manager interface (Version 4.7.2) in the 'User Management > User Roles' section. The 'Edit Role' tab is active. The configuration for the 'Student' role is as follows:

- Disable this role
- Role Name: Student
- Role Description: Student Role
- Role Type: Normal Login Role
- \*Max Sessions per User Account (  Case-Insensitive ): 0 (1 - 255; 0 for unlimited)
- Retag Trusted-side Egress Traffic with VLAN (In-Band): (0 - 4095, or leave it blank)(\*This option has been deprecated, and it will be removed in upcoming releases)
- \*Out-of-Band User Role VLAN: VLAN ID 101 (if left blank, it will default to the default access vlan settings in the Port Profile)
- \*Bounce Switch Port After Login (OOB):  Enable  Disable (This option is effective only when port profile is set to use it)
- \*Refresh IP After Login (OOB):  Enable  Disable (This option only applies to L2 OOB Virtual Gateway with Role VLAN as Access VLAN and switch port is NOT bounced after VLAN change)
- \*After Successful Login Redirect to:  previously requested URL  this URL:

**Figure 6-55** Creating Faculty Role and Mapping it to VLAN 102

The screenshot shows the Cisco Clean Access Standard Manager interface (Version 4.7.2) in the 'User Management > User Roles' section. The 'Edit Role' tab is active. The configuration for the 'Faculty' role is as follows:

- Disable this role
- Role Name: Faculty
- Role Description: Faculty Role
- Role Type: Normal Login Role
- \*Max Sessions per User Account (  Case-Insensitive ): 0 (1 - 255; 0 for unlimited)
- Retag Trusted-side Egress Traffic with VLAN (In-Band): (0 - 4095, or leave it blank)(\*This option has been deprecated, and it will be removed in upcoming releases)
- \*Out-of-Band User Role VLAN: VLAN ID 102 (if left blank, it will default to the default access vlan settings in the Port Profile)
- \*Bounce Switch Port After Login (OOB):  Enable  Disable (This option is effective only when port profile is set to use it)
- \*Refresh IP After Login (OOB):  Enable  Disable (This option only applies to L2 OOB Virtual Gateway with Role VLAN as Access VLAN and switch port is NOT bounced after VLAN change)
- \*After Successful Login Redirect to:  previously requested URL  this URL:

## Add Users and Assign 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, etc. 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, refer to 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)

## Customize 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 6-56](#).

**Figure 6-56** User Page Settings for Web Login

The screenshot shows the Cisco Clean Access Standard Manager interface. 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 configuration for a 'Login Page'. The 'General' tab is selected, and the following settings are visible:

- Enable this login page
- VLAN ID: \* [ ] (separate multiple VLANs with a comma)
- Subnet (IP/Mask): \* [ ] / \* [ ]
- Operating System: ALL
- Page Type: Frameless
- Page Description: [ ]
- Web Client (ActiveX/Applet): ActiveX on IE, Java Applet on non-IE Browser
- 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)

Buttons for 'Update', 'Cancel', and 'View' are located at the bottom of the configuration area.

## Customize 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 6-57](#), the NAC Web Agent is made mandatory for the student role.

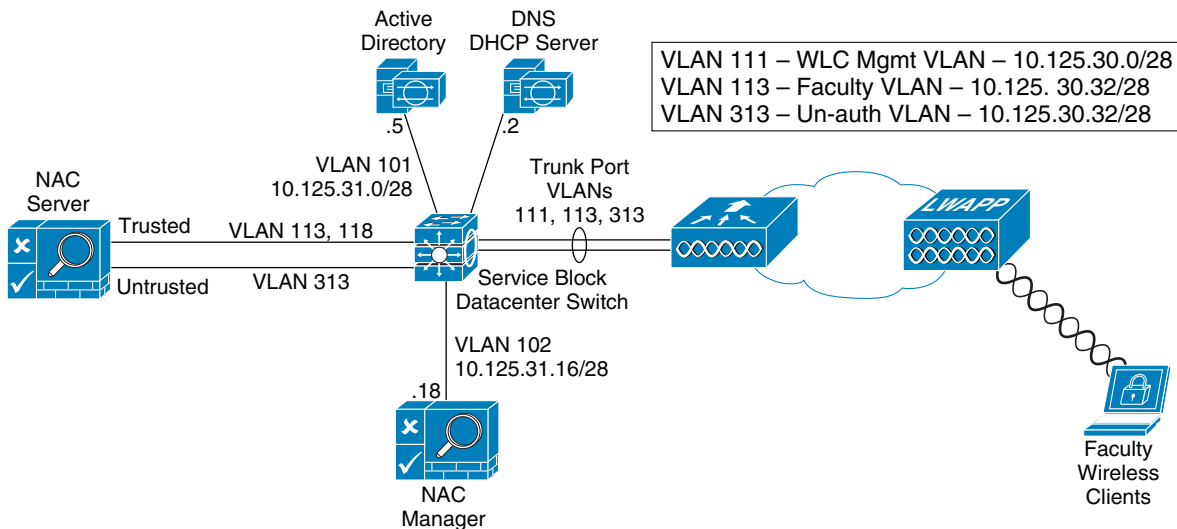
Figure 6-57 Agent Login Required for Student Role



## NAC Deployment for Wireless Clients

Within the Community College design, a NAC Layer-2 OOB deployment was used for wireless clients. Figure 6-58 shows the L2 OOB logical network diagram which was used to validate the NAC L2 OOB deployment in the Community College design. Figure 6-58 depicts the specifics for the faculty wireless clients. The deployment for student wireless clients would be similar.

Figure 6-58 Layer-2 OOB NAC Deployment Topology for Faculty Wireless Clients



As illustrated in Figure 6-58, the WLC is connected to a trunk port that carries the quarantine VLAN and access VLAN for the faculty 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 if 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. Once 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 Faculty clients. Similar steps would be taken to enable NAC for the Student clients.

### Catalyst Switch Configuration

The following Catalyst 3750 configuration example illustrates the configurations used on the appliance block switch in the Community College design for the NAC 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 Faculty Client Subnet Access VLAN
ip address 10.125.30.33 255.255.255.240
!
```

### NAC OOB Configuration Steps on the WLC and NAC Manager

The following outlines the steps needed to configure the WLC and the NAC Manager for a NAC L2 OOB deployment:

- 
- Step 1** Enable SNMPv2 mode on the controller. See [Figure 6-59](#).



Figure 6-59 Enabling SNMPv2

The screenshot shows the Cisco NAC Manager interface for configuring SNMP. The left sidebar contains a 'Management' menu with options like Summary, SNMP, HTTP, Telnet-SSH, Serial Port, Local Management Users, User Sessions, Logs, Mgmt Via Wireless, Software Activation, and Tech Support. The main area is titled 'SNMP System Summary' and includes an 'Apply' button. The configuration fields are as follows:

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

228900

- 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 6-60](#).

Figure 6-60 Creating Profile for WLC

The screenshot shows the Cisco Clean Access Standard Manager interface for creating a profile. The left sidebar contains a 'Device Management' menu with options like 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 area is titled 'Cisco Clean Access Standard Manager Version 4.7.2' and 'OOB Management > Profiles'. The configuration fields are as follows:

Profile Name	WLC_Main_Campus
Device Model	Cisco Wireless LAN Controllers
SNMP Port	161
Description	Main Campus WLC
SNMP Read Settings - Version	SNMP V2C
SNMP Read Settings - Community String	cisco123
SNMP Write Settings - Version	SNMP V2C
SNMP Write Settings - Community String	ccve

228901

- Step 3** Once 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 6-61](#).

Figure 6-61 Adding WLC in Profile

The screenshot shows the Cisco Clean Access Standard Manager interface. The left sidebar contains navigation menus for Device Management, OOB Management (with 'Devices' selected), User Management, and Monitoring. The main content area is titled 'OOB Management > Devices' and includes a 'New' button. The form fields are as follows:

- Device Profile: WLC\_Main\_Campus
- Device Group: default
- IP Addresses: 10.125.30.2
- Description: WLC 1 at Main Campus

Buttons for 'Add' and 'Reset' are located below the form.

- 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 6-62](#).

Figure 6-62 Adding MAC Manager as the SNMP Trap Receiver

The screenshot shows the Cisco Clean Access Standard Manager interface. The top navigation bar includes 'MONITOR', 'WLANS', 'CONTROLLER', 'WIRELESS', 'SECURITY', 'MANAGEMENT' (selected), 'COMMANDS', 'HELP', and 'FEEDBACK'. The left sidebar shows 'Management' with 'SNMP' expanded. The main content area is titled 'SNMP Trap Receiver > New' and includes a 'Back' button and an 'Apply' button. The form fields are as follows:

- Community Name: NacTraps
- IP Address: 10.125.31.18
- Status: Enable

- 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 6-63 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 contains navigation menus for Device Management, OOB Management (with Profiles selected), User Management, Monitoring, and Administration. The main content area is titled 'OOB Management > Profiles' and shows the configuration for an 'SNMP Trap' receiver under 'Advanced Settings'. The configuration includes:

- Trap Port on Clean Access Manager: 162
- SNMP V1 Settings: Community String (empty)
- SNMP V2c Settings: Community String: NacTraps
- SNMP V3 Settings: Security Method (Auth/Priv): No Auth, No Priv; User Name, User Auth, and User Priv (all empty)

An 'Update' button is located at the bottom right of the configuration area. A vertical ID '228904' is visible on the right edge of the screenshot.

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. [Figure 6-64](#).

Figure 6-64 Creating Dynamic Interface in the Controller

The screenshot shows the Cisco Controller configuration page for a dynamic interface. The top navigation bar includes 'MONITOR', 'WLANs', 'CONTROLLER' (selected), 'WIRELESS', 'SECURITY', 'MANAGEMENT', 'COMMANDS', 'HELP', and 'FEEDBACK'. The left sidebar shows a navigation menu for the Controller, with 'Advanced' selected. The main content area is titled 'General Information' and shows the configuration for the 'faculty and staff data' interface:

- Interface Name: faculty and 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

A vertical ID '228905' is visible on the right edge of the screenshot.

**Step 7** Create the WLAN and associate it with the dynamic interface. See [Figure 6-65](#).

Figure 6-65 Creating the WLAN

The screenshot shows the Cisco WLAN configuration interface. The 'WLANs > Edit' page is displayed with the 'Security' tab selected. The configuration details are as follows:

Profile Name	Faculty and Staff Data
Type	WLAN
SSID	data
Status	<input checked="" type="checkbox"/> Enabled
Security Policies	[WPA2][Auth(802.1X + CCKM)] (Modifications done under security tab will appear after applying the changes.)
Radio Policy	All
Interface	faculty and staff data
Broadcast SSID	<input checked="" type="checkbox"/> Enabled

**Step 8** Enable NAC in the WLAN on the WLC Controller.

Figure 6-66 Enabling NAC in the WLAN on the WLC Controller

The screenshot shows the Cisco WLAN configuration interface with the 'Advanced' tab selected. The configuration details are as follows:

Allow AAA Override	<input type="checkbox"/> Enabled
Coverage Hole Detection	<input checked="" type="checkbox"/> Enabled
Enable Session Timeout	<input checked="" type="checkbox"/> 1800 Session Timeout (secs)
Aironet IE	<input checked="" type="checkbox"/> Enabled
Diagnostic Channel	<input type="checkbox"/> Enabled
IPv6 Enable	<input type="checkbox"/>
Override Interface ACL	None
P2P Blocking Action	Disabled
Client Exclusion	<input checked="" type="checkbox"/> Enabled 60 Timeout Value (secs)
Media Session Snooping	<input type="checkbox"/>
Off Channel Scanning Defer	
DHCP Server	<input type="checkbox"/> Override
DHCP Addr. Assignment	<input type="checkbox"/> Required
Management Frame Protection (MFP)	
Infrastructure MFP Protection	<input checked="" type="checkbox"/> (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	<input checked="" type="checkbox"/> Enabled
Load Balancing and Band Select	

**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 6-67.

Figure 6-67 Adding the client subnet in the CAS server

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 for a specific CAS server (10.125.30.114). The "Advanced" tab is selected, and the "Managed Subnet" sub-tab is active. The configuration includes:

- 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: Faculty Wireless subnet (PEAP)
- Add Managed Subnet button

- 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.

Figure 6-68 Creating VLAN Mappings

The screenshot shows the Cisco Clean Access Standard Manager interface, specifically the "VLAN Mapping" configuration page. The left sidebar is the same as in Figure 6-67. The main content area is titled "Cisco Clean Access Standard Manager Version 4.7.2" and shows the configuration for a specific CAS server (10.125.30.114). The "Advanced" tab is selected, and the "VLAN Mapping" sub-tab is active. The configuration includes:

- 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 (Update)
- VLAN Mapping Assignments**
- Untrusted network VLAN ID: (input field) (-1 for non-VLAN)
- Trusted network VLAN ID: (input field) (-1 for non-VLAN)
- Description: (input field)
- Add Mapping button

At the bottom, there is a table showing the mapping:

Untrusted VLAN ID	Trusted VLAN ID	Description	Del
313	113	Wireless Faculty Users 313 -> 113	X

## Configuring Single SignOn (SSO) with the OOB Wireless Solution

The following is required to enable VPN SSO for a wireless NAC OOB deployment:

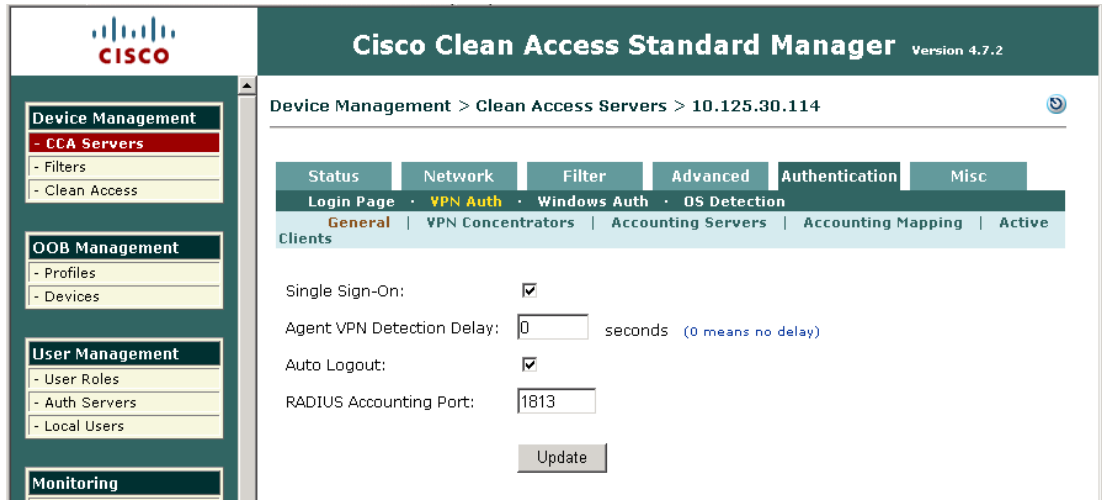
- Enable VPN authentication on the NAC server with the WLC defined as the VPN concentrator in the NAC appliance.

- Enable 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.

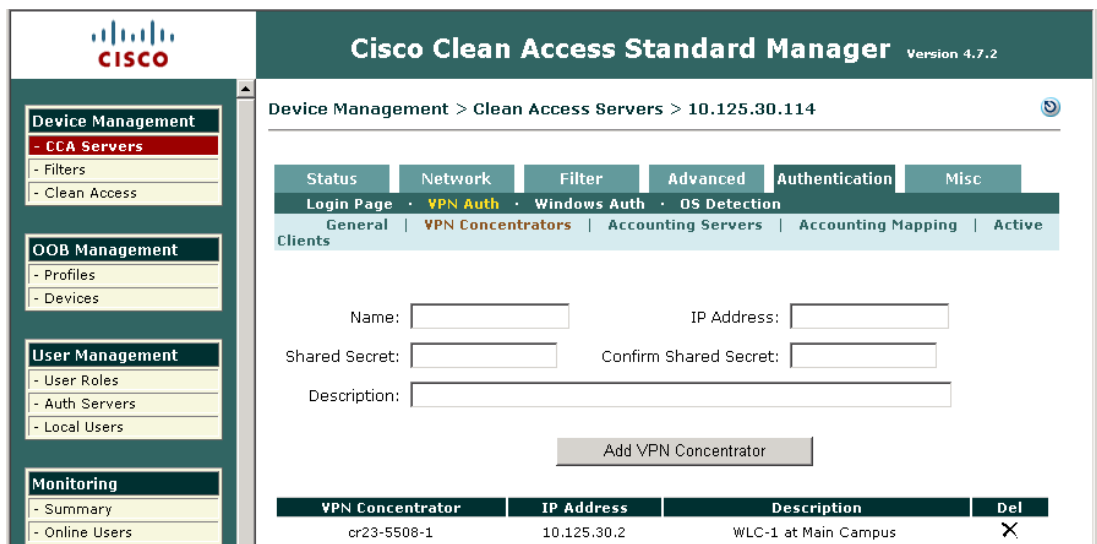
**Step 1** From the NAC Manager GUI, click **CAS server > Select your CAS server > Manage > Authentication > VPN Auth**. See [Figure 6-69](#).

**Figure 6-69 NAC Manager Configuration--Enabling SSO**



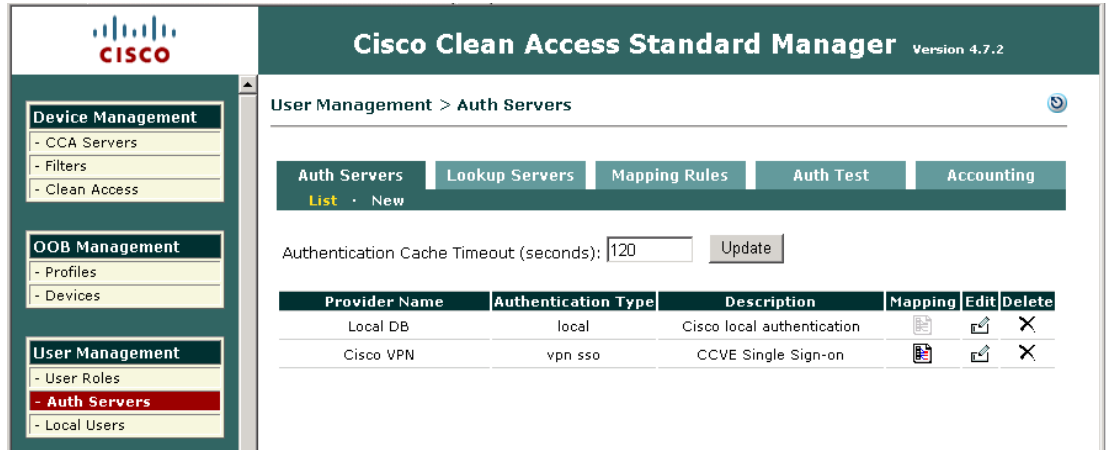
**Step 2** Select the **VPN Concentrators** tab to add a new entry for the WLC. Populate the entry fields for the WLC Management IP address and shared secret you want to use between the WLC and NAC server. See [Figure 6-70](#).

**Figure 6-70 Adding New Entry WLC**



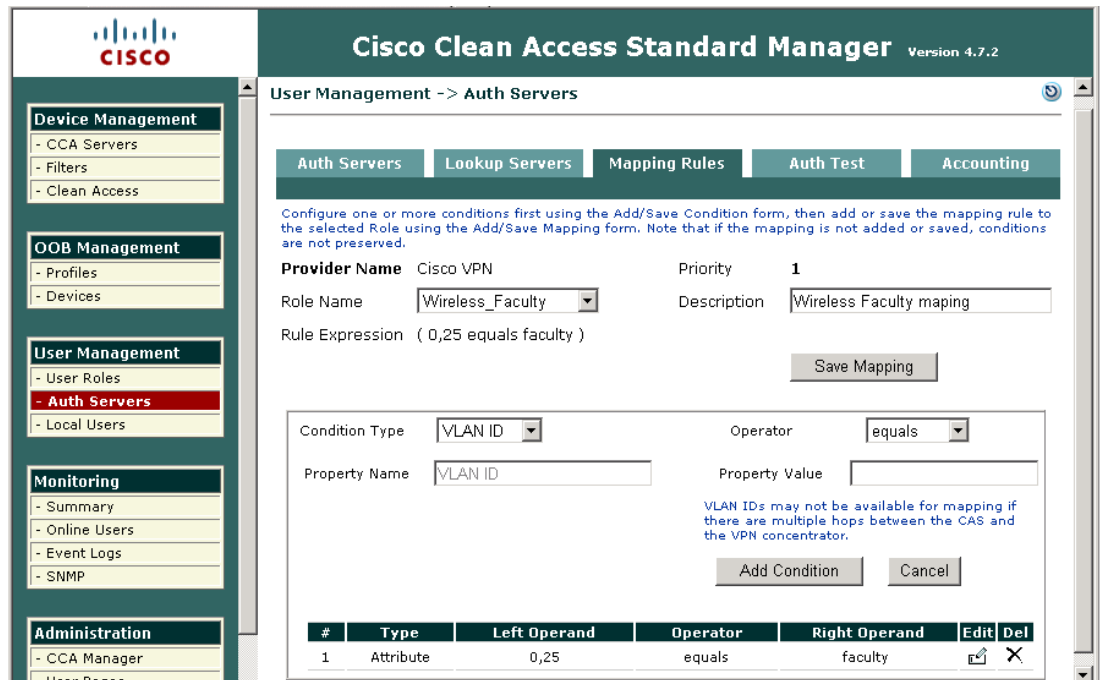
**Step 3** For role mapping, add the new authentication server with type **vpn sso** under **User Management > Auth Servers**. See [Figure 6-71](#).

**Figure 6-71 Role Mapping**



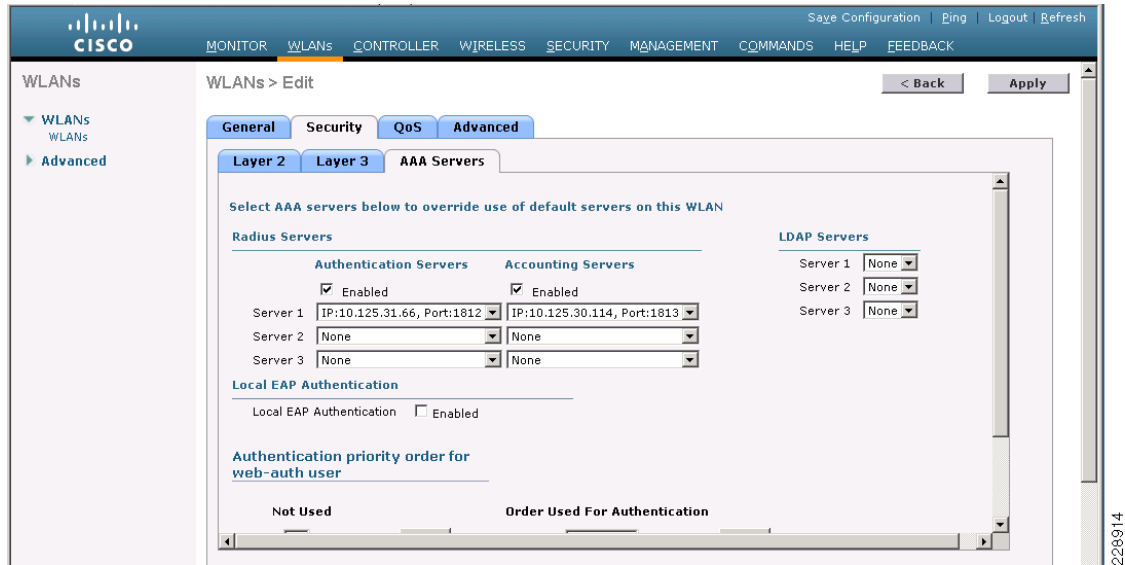
**Step 4** Click the **Mapping** icon and then add **Mapping Rule**. The mapping varies dependent upon the class attribute 25 value that WLC sends in the accounting packet. This attribute value is configured in the RADIUS Server and varies based upon the user authorization. In this example, the attribute value is **faculty**, and it is placed in the **Wireless\_Faculty** role.

**Figure 6-72 Mapping Icons to Rules**



To configure VPN SSO on the Wireless LAN Controller, RADIUS accounting needs to be enabled and sent to the NAC Server.

Figure 6-73 Enabling RADIUS Accounting for VPN SSO

**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 will mistakenly authenticate the non-single sign-on clients without prompting them.

## Further Information

The configuration guidelines and examples in the previous sections were based on the features, devices, and designs that were used for validating the Community College Reference design. For more information on all of these features, refer to the list of links [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/campus_VS_DG.html">http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campus_VS_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>