

Cisco Bring Your Own Device (BYOD) CVD Release 2.5

Revised: September 27, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors



Zeb Hallock

Zeb Hallock, Technical Marketing Engineer, Systems Development Unit, Cisco Systems

Zeb Hallock is in the Enterprise Systems Engineering group of Cisco, focusing on digital media systems. He is also pursuing creation and development of future-based collaboration systems, holding two patents in the field. He has been with Cisco for 10 years working on enterprise system testing, system design and testing of H.323 based video conferencing, and network infrastructure. He has also been a specialist working on Cisco Unified IP Contact Center Cisco Unified MeetingPlace. Before Cisco he worked as a consultant designing and implementing local and wide area networks.



John Johnston

John Johnston, Technical Marketing Engineer, Systems Development Unit, Cisco Systems

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



Fernando Macias

Fernando Macias, Solutions Architect, Systems Development Unit, Cisco Systems

Fernando Macias is a Solutions Architect in the Systems Development Unit (SDU), focusing on solutions for the enterprise market segment. Fernando has been with Cisco for over 13 years and has developed networking solutions that include Video, Security, and Content Delivery products. Fernando was also a member of Cisco's Advanced Services, where he provided network design support to large enterprise companies and a Systems Engineer for Cisco's commercial region. Prior to joining SDU, Fernando focused on Cloud Computing and Security, with the goal of leading Public Sector customers in their path to Cloud Computing.

In addition to Masters Degrees in Technology Management and Software Engineering, Fernando holds his Cisco CCIE certification in Routing and Switching.



Roland Saville

Roland Saville, Technical Leader, Systems Development Unit, Cisco Systems

Roland Saville is a technical leader for the Systems Development Unit (SDU) at Cisco, focused on developing best-practice design guides for enterprise network deployments. He has more than 15 years of experience at Cisco as a systems engineer, consulting systems engineer, technical marketing engineer, and technical leader. During that time, he has focused on a wide range of technology areas including the integration of voice and video onto network infrastructures, network security, wireless LAN networking, RFID, and energy management. He has also spent time focusing on the retail market segment. Prior to Cisco, he spent eight years as a communications analyst for Chevron Corporation. Saville holds a bachelor of science degree in electrical engineering from the University of Idaho and a master of business administration degree from Santa Clara University. He co-authored the book "Cisco TelePresence Fundamentals" and has eight U.S. patents.



Srinivas Tenneti

Srinivas Tenneti, Technical Marketing Engineer, Systems Development Unit, Cisco Systems

Srinivas Tenneti is a Technical Marketing Engineer in the Systems Development Unit (SDU) at Cisco, responsible for the design and architecture of security components for the BYOD project. He has more than 11 years of experience at Cisco where has worked in development, System Testing, and in Enterprise design and architecture. In the last 5 years Srinivas has worked on the design and architecture of GETVPN, DMVPN, Branch Architecture, Internet Edge, and PKI as a service for VPN protocols. Srinivas has co-authored the book PKI Uncovered: Certificate-Based Security Solutions for Next Generation Networks.



Mike Jessup

Mike Jessup, Technical Lead Engineer/Architect, Systems Development Unit, Cisco Systems

Mike Jessup joined Cisco in 1996 and is a Technical Lead Engineer/Architect within Cisco's Systems Development Unit (SDU). Mike's primary focus is identifying and qualifying BYOD system architectures and solutions as well as future trends and working with solutions engineers in incorporating, validating, and documenting these systems in Cisco Validated Designs. Prior to joining SDU in November of 2011, Mike worked in the US Sales Organization as an Enterprise System Engineer specializing in Routing and Switching, Data Center, Application Optimization, and Network Management. Prior to joining Cisco, Mike worked in positions as a Field Engineer and Systems Engineer in various Information Systems companies and Partner organizations since 1982.



Suyog Deshpande

Suyog Deshpande, Technical Marketing Engineer, Systems Development Unit, Cisco Systems

Suyog Deshpande is a Technical Marketing Engineer within Systems Development Unit (SDU) at Cisco Systems, focussing on converged wired and wireless access products. Suyog has been with Cisco for 9 years with previous experience in the Wireless Networking Business Unit, where he focussed on RF systems and converged access products. Before Cisco, Suyog held various positions designing and deploying large scale wireless networks and working with spectrum intelligence and analyzer products.



Tim Szigeti

Tim Szigeti, Technical Leader, Systems Development Unit, Cisco Systems

Tim Szigeti is a technical leader in the Systems Design Unit (SDU) at Cisco Systems, where his role is to design network architectures for BYOD solutions. He has also specialized in Quality of Service technologies for over a dozen years, during which he has authored many technical papers, design guides, as well as the Cisco Press Books: End-to-End QoS Network Design and Cisco TelePresence Fundamentals. Szigeti holds CCIE certification 9794 and holds a bachelor of commerce degree with a specialization in management information systems from the University of British Columbia.

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 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.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at http://www.cisco.com/go/trademarks. Third party trademarks mentioned 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. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Bring Your Own Device (BYOD) CVD Release 2.5

© 2013 Cisco Systems, Inc. All rights reserved.

PART 1 Introduction

CHAPTER 1 BYOD Solution Overview 1-1

Introduction 1-1

BYOD Market Landscape 1-1

BYOD and the Enterprise Network 1-2

Role of MDM in the Enterprise Network 1-3

User Need and Role of IT 1-3

Protecting Data and Loss Prevention 1-4

Challenges for End Users 1-4

Key Advantages of the Cisco BYOD Solution 1-4

PART 2 BYOD Design Overview

CHAPTER 2 Summary of Design Overview 2-1

CHAPTER 3 Cisco BYOD Solution Components 3-1

Cisco Access Points 3-3

Cisco Wireless Controller 3-4

Cisco Identity Services Engine 3-4

Cisco Adaptive Security Appliance 3-4

Cisco AnyConnect Client 3-4

Cisco Integrated Services Routers 3-5

Cisco Aggregation Services Routers 3-5

Cisco Catalyst Switches 3-5

Cisco Converged Access Switches 3-5

Cisco Nexus Series Switches 3-6

Cisco Prime Infrastructure 3-6

Secure Access to the Corporate Network 3-6

Certificate Enrollment and Mobile Device Provisioning 3-7

CHAPTER 4 BYOD Use Cases 4-1

Enhanced Access—Personal and Corporate Devices 4-2

Limited Access—Corporate Devices 4-3

Advanced Access—MDM Posture 4-4

Basic Access—Guest-Like 4-5

CHAPTER 5 **Campus and Branch Network Design for BYOD** 5-1

Campus Network Design 5-1
Centralized (Local Mode) Wireless Design 5-1
Security Group Tag Overview 5-3
ACL Complexity and Considerations 5-4
Security Group Tag 5-4
SGT Deployment Scenarios in this CVD 5-5
Campus Wired Design 5-6
Converged Access Campus Design 5-7
Campus Migration Path 5-10
Initial Overlay Model 5-10
Centralized/Local Mode Only 5-11
Hybrid Converged Access and Local Mode 5-13
Full Converged Access 5-15
Wireless LAN Controller High Availability 5-16
Branch Wide Area Network Design 5-17
Branch WAN Infrastructure 5-17
Branch WAN Bandwidth Requirements 5-18
Encryption Requirement 5-18
Transport 5-19
Branch LAN Network Design 5-19
FlexConnect Wireless Design 5-20
Branch Wired Design 5-21
Converged Access Branch Design 5-22

CHAPTER 6 **Mobile Device Managers for BYOD** 6-1

Cisco ISE 1.2 with MDM API Integration 6-1
MDM Deployment Options and Considerations 6-2
On-Premise 6-4
Cloud-Based 6-5
Enterprise Integration Considerations 6-6
Corporate Directory Services Integration 6-6
Corporate Certificates Authority and Public Key Infrastructure Integration 6-6
Email Integration 6-6
Content Repository Integration 6-6
Management Integration 6-7
Integration Servers 6-7

CHAPTER 7 **Application Considerations and License Requirements for BYOD** 7-1

Quality of Service 7-1
Rate Limiting 7-4
Application Visibility and Control (AVC) 7-4
Cisco Jabber 7-11
Cisco Jabber Clients and the Cisco BYOD Infrastructure 7-12
Use Case Impact on Jabber 7-12
Other Jabber Design Considerations 7-12
Cisco Virtual Workspace (VXI) Smart Solution 7-13
Cisco Virtual Workspace (VXI) Architecture 7-13
Cisco Virtual Workspace (VXI) Application Virtualization and Citrix 7-14
License Requirements for BYOD Solution 7-14

PART 3 **Configuring the Infrastructure**

CHAPTER 8 **Summary of Configuring the Infrastructure** 8-1

CHAPTER 9 **BYOD Wireless Infrastructure Design** 9-1

Campus—Unified Wireless LAN Design 9-1
Centralized Campus—Dual SSID Design 9-2
Centralized Campus—Single SSID Design 9-7
Centralized Campus—Policy Enforcement using TrustSec 9-8
Branch—Unified Wireless LAN Design 9-8
FlexConnect Wireless LAN Design 9-8
Branch Wireless IP Address Design 9-11
FlexConnect Branch—Dual SSID Design 9-12
Dual SSID—Central Switching Provisioning 9-16
Dual SSID—Local Switching Provisioning 9-18
FlexConnect Branch—Single SSID Design 9-22
FlexConnect Access Point Configuration 9-24
FlexConnect Groups 9-26
FlexConnect VLAN Override 9-30
Campus—Converged Access Design 9-30
Campus Converged Access—Dual SSID Design 9-31
Campus Converged Access—Single SSID Design 9-35
Campus Converged Access—Mobility 9-36
Branch—Converged Access Design 9-39
Branch Converged Access—Dual SSID Design 9-39

Branch Converged Access—Single SSID Design 9-41

CHAPTER 10 **Identity Services Engine for BYOD** 10-1

Identity Certificate for ISE 10-1

Network Device Definition within ISE 10-2

ISE Integration with Active Directory 10-3

Guest and Self-Registration Portals 10-4

ISE Using Certificates as an Identity Store 10-6

Identity Source Sequences 10-6

SCEP Profile Configuration on ISE 10-7

Authentication Policies 10-8

Authentication Policy for Wireless 10-10

Client Provisioning 10-11

Client Provisioning Resources—Apple iOS and Android 10-12

Client Provisioning Policy—Apple iOS and Android Devices 10-13

Client Provisioning Resources—Mac OS 10-13

Client Provisioning Policy for Mac OS Devices—Wireless 10-15

Client Provisioning Policy for Windows Devices—Wireless/Wired 10-17

Profiling 10-18

Enabling the DHCP and RADIUS Probes 10-19

Profiling Android Devices 10-21

Logical Profiles 10-22

Authorization Policies and Profiles 10-23

Authorization Profiles 10-24

Wireless CWA Authorization Profile for Dual SSID Provisioning 10-24

Dual SSID Provisioning Authorization Rule 10-27

Wireless NSP Authorization Profile for Single SSID Provisioning 10-28

Single SSID Provisioning Authorization Rule 10-29

Certificate Authority Server 10-30

NDES Server Configuration for SCEP 10-31

Certificate Template 10-31

CHAPTER 11 **BYOD Wired Infrastructure Design** 11-1

Campus Wired Design 11-1

VLAN Design for Wired Switches at Campus 11-2

IP Address Design for Campus Wired Infrastructure 11-2

Policy Enforcement in the Campus for Wired Devices 11-2

ACL Design for Campus Access Layer Switches 11-3

Provisioning ACL 11-4
802.1X and AAA Configuration for Campus Switches 11-5
Branch Wired Design—Non-Converged Access 11-7
VLAN Design at Branch Locations 11-8
IP Address Allocation at Branch Location 11-8
Policy Enforcement in the Branch for Wired Devices 11-9
ACL Design at Branch Location 11-10
Provisioning ACL 11-11
802.1X and AAA Configuration for Branch Switches 11-12
Branch Wired Design—Converged Access 11-12
VLAN Design at Branch Locations 11-13
IP Address Allocation at Branch Location 11-13
Policy Enforcement at the Branch Using Converged Access Switches 11-14
ACL Design at Branch Location for Converged Access Switches 11-14
802.1X and AAA Configuration for Branch Switches 11-16
MAB Devices at Branch or at Campus Location 11-16

CHAPTER 12 **Security Group Access for BYOD** 12-1
Unified Infrastructure Design to Support SGA 12-1
Policy Configuration for SGACLs in Scenario 1 12-2
Policy Configuration in Scenario 2 12-5
TrustSec Summary 12-7

CHAPTER 13 **Mobile Device Manager Integration for BYOD** 13-1
Establishing IP Connectivity for an On-Premise MDM 13-1
Establishing IP Connectivity for a Cloud-Based MDM 13-2
Configure ISE to Authenticate the MDM API 13-3
Creating the MDM API User Account 13-4
Setting Up the MDM Connection 13-4
Verifying the MDM Connection 13-5
Configuring the MDM 13-7

PART 4 **BYOD Use Cases**

CHAPTER 14 **Summary of BYOD Use Cases** 14-1

CHAPTER 15 **BYOD Enhanced Use Case—Personal and Corporate Devices** 15-1
Active Directory Groups 15-3

Distributing Digital Certificates 15-5
Mobile Device On-boarding 15-6
Provisioning Flows 15-7
Provisioning Apple iOS Devices 15-7
Provisioning Android Devices 15-8
Provisioning Wired Devices 15-10
Key and Certificate Storage 15-12
Network Device Groups 15-13
Policy Enforcement for Security Group Access 15-14
Security Group Access Tags 15-14
Security Group ACL 15-15
Authorization Polices for SGT 15-16
Personal Wireless Devices—Full Access 15-16
Simple and Compound Conditions 15-17
Permissions 15-20
Centralized Campus with SGT 15-20
Centralized Campus with ACLs 15-21
Branch with FlexConnect 15-21
Converged Access Branch or Campus 15-22
Personal Wireless Devices—Partial Access 15-23
Permissions 15-25
Centralized Campus with SGT 15-25
Centralized Campus with ACLs 15-26
Branches with FlexConnect 15-28
Converged Access Branch or Campus 15-30
Personal Wireless Devices—Internet Only Access 15-32
Permissions 15-34
Centralized Campus with SGT or ACLs 15-34
Branches with FlexConnect 15-36
Converged Access Branch or Campus 15-39
Personal Wired Devices—Full Access 15-41
Wired Simple and Compound Conditions 15-43
Permission 15-43
Campus Wired 15-44
Branch Wired 15-45
Converged Access Branch or Campus 15-46
Personal Wired Devices—Partial Access 15-47
Permissions 15-49
Campus Wired 15-49

Branch Wired 15-50
Converged Access Branch and Campus 15-51
Personal Wired Devices—Internet Only Access 15-52
Permissions 15-54
Wired Campus 15-54
Branch Wired 15-56
Converged Access Branch and Campus 15-57
Android Devices—Deny Access 15-57
ISE Authorization Policy 15-58

CHAPTER 16 **BYOD Limited Use Case—Corporate Devices** 16-1
Whitelist Identity Group 16-2
Policy Enforcement for Security Group Access 16-3
Security Group Access Tags 16-4
Security Group ACL 16-4
Authorization Polices for SGT 16-5
Corporate Devices—Full Access 16-5
Simple and Compound Conditions 16-8
Permissions for Wireless Users 16-10
Centralized Campus with SGTs 16-11
Centralized Campus with ACLs 16-11
Branch with FlexConnect 16-12
Converged Access Branch or Campus 16-13
Corporate Wired Devices—Full Access 16-14
Wired Simple and Compound Conditions 16-16
Permissions for Wired Users 16-17
Campus Wired 16-17
Branch Wired 16-19
Converged Access Branch or Campus 16-19
ISE Authorization Policy 16-20

CHAPTER 17 **BYOD Advanced Use Case—Mobile Device Manager Integration** 17-1
Supported MDM Functions 17-3
Integration Process Flow 17-5
ISE Compliance Check 17-6
MDM Compliance Check 17-7
ISE Configuration 17-7
Logical Profile 17-7

Authorization Policy 17-8
MDM Enrollment Rule 17-9
Remediate Non-ISE Compliant Rule 17-13
Remediate Non-MDM Compliant Rule 17-17
MDM Reports 17-21

CHAPTER 18 **BYOD Basic Access Use Case** 18-1
Extending Guest Wireless Access to Employee Personal Devices 18-1
Allowing Employees to Provision Guest Credentials for Themselves 18-2
Extending Web Auth to Use Microsoft AD when Authenticating Employees with Personal Devices 18-4
Deploying Guest-Like Wireless Access for Employee Personal Devices 18-5
Dedicated SSID for Employee Personal Devices 18-7
Wireless Controller Configuration 18-9
Campus Controller Configuration 18-9
Cisco ISE Policy Configuration 18-14
ASA Firewall Configuration 18-17
Differentiated Quality of Service Treatment 18-19
Accessing Corporate Resources 18-20
Securing Mirror Sites for Personal Devices 18-20
DNS Support 18-23
Outlook Web Access for Employee Devices 18-23
ActiveSync Support 18-24
VPN Client 18-24
Virtual Desktop Client 18-26
Summary 18-27

CHAPTER 19 **User Experience—How To On-board a BYOD Device** 19-1
Apple iOS Devices 19-1
Android Devices 19-6
Windows Devices 19-11
Windows Wireless Devices 19-12
Windows Wired Devices 19-14
Mac OS/X Devices 19-16

PART 5 **BYOD Operations and Services**

CHAPTER 20 **Summary of Operations and Services** 20-1

CHAPTER 21 **BYOD Guest Wireless Access** 21-1

Overview 21-1
IP Addressing and DNS 21-3
Authentication and Authorization 21-6
Designing Guest Access for Campus and Branch locations 21-7
WLC Configuration 21-7
Campus Controller 21-8
Guest Controller 21-13
Cisco ISE Policy Configuration 21-19
Cisco ISE Sponsor Portal 21-23
Configuring the Cisco ISE Sponsor Portal 21-24
Cisco ISE Guest Portal 21-27
Configuring the Cisco ISE Guest Portal 21-27
ASA Firewall Configuration 21-30
Additional Considerations 21-31
Wireless Guest Access at the Branch 21-32
Rate-Limiting Guest Wireless Access 21-34
Multiple Guest SSIDs and AP Groups 21-37
Managing the Downstream Load 21-38

CHAPTER 22 **Managing a Lost or Stolen Device** 22-1

Blacklist Identity Group 22-2
Blacklisting Wireless Devices 22-2
Blacklisting Wired Devices 22-7
Creating a Downloadable ACL on ISE 22-8
Creating the URL REDIRECT ACL on the Switch 22-8
Configuring the Authorization Profile 22-9
Creating a Rule in the Authorization Policy 22-10
Employees' My Devices Portal 22-10
Reporting a Device as Lost 22-11
Reinstating a Device 22-13
PIN Lock and Device Wipes 22-14
Administrators—Blacklisting a Device 22-14
MDM Actions 22-16
Endpoint Protection Services (EPS) 22-17
Revocation of Digital Certificate 22-21

Disable the RSA SecurID Token 22-23

CHAPTER 23 **BYOD Policy Enforcement Using Security Group Access 23-1**

Security Group Tag Overview 23-1
ACL Complexity and Considerations 23-1
Security Group Tag 23-2
Security Group Access Domain Infrastructure 23-3
TrustSec using 802.1X for Link Encryption 23-3
TrustSec in Manual Mode for Link Encryption 23-3
Security Group Tag Distribution and Forwarding Mechanisms 23-5
User Policies with SGT 23-7
SGT Assignment for Data Center Servers 23-8
SGT Deployment Scenarios in this CVD 23-10
Configuring the Infrastructure for TrustSec 23-12
Unified Infrastructure Design to Support TrustSec 23-12
TrustSec Policy Configuration for SGACLs in Infrastructure Deployment Scenario 1 23-14
Configuring ISE to Support TrustSec 23-17
Configuring ISE for Network Access Device Authentication 23-27
Configuring Network Devices with a Device SGT 23-31
Configuring Network Access Devices for Authentication at ISE 23-34
RADIUS Server Configuration on the Wireless Controller 23-35
RADIUS Server Configuration on the Catalyst 6500 23-36
RADIUS Server Configuration on the Nexus 7000 23-37
Catalyst 6500 Platform Specific Considerations 23-38
Configuring Switching Infrastructure to Support TrustSec with 802.1ae MACsec Encryption 23-38
Catalyst 6500 Commands 23-40
Nexus 7000 Commands 23-41
Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers 23-41
Wireless Controller Configuration 23-42
Catalyst 6500 SXP Configuration 23-43
Configuring Static IP/SGT Bindings on Nexus Switches 23-44
TrustSec Policy Configuration Using the ASA and Security Group Firewall in the Data Center—Deployment Scenario 2 23-46
Configuring ISE to Support TrustSec 23-48
Configuring ISE for Network Access Device Authentication 23-51
Wireless Controller Configuration 23-52
Nexus 7000 Configuration 23-54
ASA Firewall Configuration 23-56
Configuring the Network Devices for Integration with ISE 23-59

RADIUS Server Configuration on the Wireless Controller 23-60
RADIUS Server Configuration on the ASA Firewall 23-61
RADIUS Server Configuration on the Nexus 7000 23-67
Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers 23-67
Wireless Controller Configuration 23-68
Catalyst 6500 SXP Configuration 23-69
Nexus 7000 SXP Configuration 23-70
ASA SXP Configuration 23-71
Configuring Static IP/SGT Bindings on Nexus Switches 23-74
Configuring SG-FW Role-Based Policies at ASA 23-75
Device On-boarding, Provisioning, Authentication, and Authorization Policies for TrustSec in ISE 23-82

CHAPTER 24 **Mobile Traffic Engineering with Application Visibility and Control (AVC)** 24-1
Executive Summary 24-1
Macro Trends and Business Requirements 24-2
Cisco Application Visibility and Control (AVC) for Wireless LAN Controllers 24-3
Challenges and Solutions for Managing Application Quality over Wireless Media 24-5
IEEE 802.11 Distributed Coordination Function (DCF) 24-5
Collision Sense Multiple Access/Collision Avoidance (CSMA/CA) 24-6
Short Interframe Space (SIFS) 24-7
DCF Interframe Space (DIFS) 24-7
Contention Window (CW) 24-8
DCF Operation 24-9
IEEE 802.11e/WMM 24-10
802.11e User Priorities 24-11
Access Categories (AC) 24-12
Enhanced Distributed Coordination Function (EDCF) 24-13
Arbitration Interframe Spacing (AIFS) 24-13
Contention Window Enhancements 24-14
Transmission Opportunity (TXOP) 24-15
Call Admission Control (TSpec) 24-15
EDCF Operation 24-15
Cisco's Strategic Approach to Application Quality Management 24-16
Cisco's Strategic Application-Class QoS Recommendations 24-17
Application-Class Expansion 24-20
Application-Class Mapping to WMM 24-20
Cisco 802.11e/802.1p/DSCP Mappings 24-21
Default DSCP-to-CoS/UP Mappings 24-22
Cisco WLC/AP QoS Translation Table 24-22

Residual DSCP-to-WMM Mappings 24-23
WLC AVC/QoS Profile-to-DSCP Mappings 24-23
Cisco Wired-Wireless Mapping Points 24-24
Configuring Downstream QoS Policies for Mobile Applications 24-25
Wireless Controller QoS Profile GUI Configuration 24-26
WLC QoS Profile CLI Configuration 24-28
Wireless Controller AVC QoS Profile Configuration 24-29
AVC Applications By Business Use Case 24-31
WLC AVC Profile GUI Configuration 24-33
WLC AVC Profile CLI Configuration 24-37
AP Access Switch Downstream QoS Policies 24-40
Four-Class Enterprise to WMM Mapping Model 24-41
Eight-Class Enterprise to WMM Mapping Model 24-46
Twelve-Class Enterprise to WMM Mapping Model 24-51
Configuring Upstream QoS Policies for Mobile Applications 24-53
Wireless Device Marking and WMM Policies 24-54
AP Access Switch Upstream QoS Policies 24-56
Catalyst 3750 Configuration Example 24-56
Catalyst 4500 Configuration Example 24-57
Catalyst 6500 Configuration Example 24-57
Application Visibility and Management 24-58
Global Application Visibility 24-58
Per-WLAN Application Visibility 24-59
Per-Client Application Visibility 24-60
NetFlow 24-61
NetFlow Configuration-GUI 24-61
NetFlow Configuration-CLI 24-62
Summary 24-62
Additional Reading 24-63

CHAPTER 25 **Managing Bonjour Services for BYOD** 25-1
Executive Summary 25-1
Why Bonjour? 25-1
Bonjour Overview 25-2
Bonjour Addressing 25-2
Bonjour Naming 25-3
Bonjour Naming Rules 25-3
Bonjour Service Discovery 25-5
Bonjour Optimization 25-5

Caching 25-6

Suppression of Duplicate Responses 25-6

Exponential Back-off and Service Announcement 25-6

Cisco Bonjour Gateway Solution in WLC 7.4+ 25-6

Bonjour Gateway Service Policy Deployment Options 25-10

Bonjour Gateway BYOD Use Cases and Configuration Examples 25-12

Use Case 1—Wireless-to-Wired Bonjour Gateway Service Policy—BYOD Employee AirPrint Example 25-13

Step 1—Enable mDNS Global Snooping 25-14

Step 2—Editing the Default mDNS Profile 25-15

Step 3—Apply the Default mDNS Profile an Interface (or Interface-Group) 25-18

Use Case 2—Wireless-to-Wireless Bonjour Gateway Service Policy—BYOD Guest AirPlay Example 25-21

Step 1—Creating a New mDNS Profile 25-22

Step 2—Adding Bonjour Services to the New mDNS Profile 25-23

Step 3—Enable mDNS Snooping and the New mDNS Profile on the WLAN 25-25

Verifying Bonjour Gateway Operation 25-27

Advanced Bonjour Gateway Scenario Operation 25-30

Guest Anchoring 25-31

Layer 3 Roaming 25-31

FlexConnect 25-32

Summary 25-32

References 25-33

CHAPTER 26 **BYOD Remote Device Access** 26-1

Solution Components 26-2

RSA SecurID 26-2

ISE Integration with RSA 26-3

VPN Design Considerations 26-3

ASA Configuration 26-5

ASA's Identity Certificate 26-5

ASA Trust Point to Authenticate Remote Users 26-7

Creating Groups for Different Types of Users 26-7

Connection Profile Configuration 26-8

Enabling AnyConnect VPN on the ASA 26-8

Provisioning a Windows Device to Rmotely Connect to the Network 26-8

Verifying What ISE Policy Rules Are Applied 26-12

Provisioning an Apple iOS Device to Remotely Connect to the Network 26-14

CHAPTER 27 **BYOD Network Management** 27-1

Key Acronyms and Terminology 27-1

Cisco Prime Infrastructure Overview 27-2

Prime Infrastructure and Supporting Components 27-2

BYOD User and Device Tracking 27-3

Components 27-4

Locating Users and Devices 27-5

Context Aware Dashboard Search 27-6

Standard and Advanced Search 27-8

BYOD Template-Based Configuration 27-14

Consistent Configuration of WLCs 27-15

Multiple Templates for Variations of Deployment 27-15

Rapid Deployment of New or Replacement Components 27-16

Staged Rollout of Configuration Changes with Rapid Rollback 27-16

Template Creation and Implementation 27-16

Template Creation 27-16

Detailed Steps for Ensuring WLAN ID Consistency 27-24

PART 6 **Appendices**

APPENDIX A **BYOD Converged Access Configurations** A-1

Converged Access—Campus A-1

Converged Access—Branch A-8

APPENDIX B **References** B-1

APPENDIX C **Software Versions** C-1

APPENDIX D **BYOD System Release Notes** D-1

IOS XE 3.2.2 SE Software Release—Catalyst 3850 Series Switches and Cisco 5760 Wireless Controllers D-1

Additional Observations D-2



PART 1

Introduction



Preface

Revised: September 27, 2013

This document is a Cisco Validated Design (CVD) for Cisco Bring Your Own Device (BYOD) Solutions. It presents system-level requirements, recommendations, guidelines, and best practices for deploying personal, corporate, and guest devices onto a network to fit your business needs. As Cisco continues to develop and enhance the technologies required to implement a BYOD solution, this CVD will continue to evolve and be updated to provide the latest guidelines, recommendations, and best practices for designing and deploying a BYOD solution.

How to Use this Document

This document is organized into five main parts after the initial Chapter 1, "BYOD Solution Overview."

BYOD Design Overview

The chapters in this part of the document describe the main components of Cisco BYOD solution and explain how these components work together to form a complete end-to-end solution:

- Chapter 3, "Cisco BYOD Solution Components"—Highlights the different network components used in the design guide.
- Chapter 4, "BYOD Use Cases"—Addresses four different use cases based on the type of network access allowed by an organization.
- Chapter 5, "Campus and Branch Network Design for BYOD"—Describes different campus and branch designs used to support BYOD, including WAN infrastructure, FlexConnect, and Converged Access.
- Chapter 6, "Mobile Device Managers for BYOD"—Introduces the ISE integration with different third-party Mobile Device Managers and explores different deployment models.
- Chapter 7, "Application Considerations and License Requirements for BYOD"—Highlights different requirements that need to be present to provide the proper network service to applications.

Configuring the Infrastructure

The chapters in this part of the document describe the network infrastructure design and configuration foundations to deploy a BYOD solution in a customer environment:

- Chapter 9, "BYOD Wireless Infrastructure Design"—Presents different network designs used to support BYOD, including Campus and Branch designs.
- Chapter 10, "Identity Services Engine for BYOD"—Focuses on digital certificates, authentication and authorization policies, and device profiling and different ways to on-board devices with either single or dual SSID configurations.
- Chapter 11, "BYOD Wired Infrastructure Design"—Highlights how to on-board wired devices and how to enforce BYOD policies and network access for wired devices.
- Chapter 12, "Security Group Access for BYOD"—Presents two different deployment scenarios that rely on Security Group Tags to enforce BYOD policies.
- Chapter 13, "Mobile Device Manager Integration for BYOD"—Focuses on how to configure ISE to integrate with third-party MDM products through an XML-based API.

BYOD Use Cases

The chapters in this part of the document describe the following four BYOD use case examples of access requirements an organization may enforce as well as a user interaction with ISE during on-boarding:

- Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices"—Provides network access for personal devices, as well as corporate-issued devices. It allows a business to build a policy that enables granular role-based application access and extends the security framework on and off-premises.
- Chapter 16, "BYOD Limited Use Case—Corporate Devices"—Enables access exclusively to corporate-issued devices.
- Chapter 17, "BYOD Advanced Use Case—Mobile Device Manager Integration"—This comprehensive use case also provides network access for personal and corporate-issued devices. However it includes posture of the device into the network access control decision through integration with third-party Mobile Device Managers (MDM).
- Chapter 18, "BYOD Basic Access Use Case"—This use case is an extension of traditional wireless guest access. It represents an alternative where the business policy is to not on-board/register employee wireless personal devices, but still provides Internet-only or partial access to the network.
- Chapter 19, "User Experience—How To On-board a BYOD Device"—Captures a typical user interaction with ISE during the on-boarding process.

BYOD Operations and Services

The chapters in this part of the document describe various services in addition to the use cases described in the previous section:

- Chapter 21, "BYOD Guest Wireless Access"—Describes a traditional wireless guest access solution where users do not have to on-board or register a device on the network, but only Internet-only access is provided to users.
- Chapter 22, "Managing a Lost or Stolen Device"—Describes how to deny access to a device that is reported lost or stolen to prevent unauthorized access to the network.
- Chapter 23, "BYOD Policy Enforcement Using Security Group Access"—Describes in depth how to use Security Group Access as an alternative policy enforcement method to access control lists.

- Chapter 24, "Mobile Traffic Engineering with Application Visibility and Control (AVC)"—Describes in depth how to use application visibility and control techniques to ensure a seamless user experience.
- Chapter 25, "Managing Bonjour Services for BYOD"—Describes a use case where users need to connect to Bonjour devices.
- Chapter 26, "BYOD Remote Device Access"—Describes how to accommodate devices that attempt to connect remotely to access internal resources.
- Chapter 27, "BYOD Network Management"—Describes how to configure and deploy Cisco Prime Infrastructure management suite to manage the BYOD solution.

Appendices

The appendices contain useful information that is not covered in the main chapters of this CVD:

- Appendix A, "BYOD Converged Access Configurations"—Provides the configurations required to deploy Cisco's Converged Access solution presented in the BYOD CVD design.
- Appendix B, "References"—Provides links to references that supplement this design guide.
- Appendix C, "Software Versions"—Provides the software versions and devices leveraged in this design.
- Appendix D, "BYOD System Release Notes"—Provides important information you should be aware of when designing and implementing this release of the BYOD CVD.

For Experienced Users

Readers who are familiar with previous versions of this CVD or who are experienced at designing a BYOD solution can use this document as a reference source. Rather than reading every page or every chapter, this document has been broken into modules that can be easily searched for a particular topic. Updates to the topics in this CVD will be published periodically.

For New Users

This document is long and contains an extensive amount of complex technical information. It can seem intimidating, particularly, if you are a first time reader of this document or do not have much experiencing a BYOD solution.

To orient yourself to the document, we recommend you begin with Chapter 2, "Summary of Design Overview," which provides an overview of the major components required to deploy a BYOD solution and typical access control use cases. From this section, you can then determine if you need particular design guidance around the infrastructure, the uses cases, or a set operation such as Remote Access.

Where to Find Additional Information

Because the document covers a wide spectrum of Cisco Network Infrastructure, Security, and Mobility products and possible solution designs, it cannot provide all the details of individual products, features, or configurations. For that type of detailed information, refer to the specific product documentation available at: http://www.cisco.com.

This document provides general guidance on how to design your own BYOD solution. Cisco has developed, tested, and documented specific solutions for certain applications and has made those solutions available for customers to copy and deploy. They are part of the Cisco Validated Design program described and documented at: http://www.cisco.com/go/designzone.

Revision History

This document may be updated at any time without notice. You can obtain the latest version of this document online at:
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/own_device.html.

Visit this website periodically and check for documentation updates by comparing the revision date of your copy with the revision date of the online document.

Table 1 lists the revision history for this document.

Table 1 Revision History

Revision Date	Comments
August 7, 2013	Initial version of this BYOD CVD.
September 27, 2013	Added note on potential incompatibilities introduced by Apple iOS 7 to chapters 10, 19, and 21. Corrected table 21-2.

Command Syntax Conventions

Table 2 describes the syntax used with the commands in this document.

Table 2 Command Syntax Guide

Convention	Description
boldface	Commands and keywords.
italic	Command input that is supplied by you.
[ ]	Keywords or arguments that appear within square brackets are optional.
{ x | x | x }	A choice of keywords (represented by x) appears in braces separated by vertical bars. You must select one.
^ or Ctrl	Represent the key labeled *Control*. For example, when you read ^*D* or *Ctrl-D*, you should hold down the Control key while you press the D key.
`screen font`	Examples of information displayed on the screen.
boldface screen font	Examples of information that you must enter.

Table 2 Command Syntax Guide

Convention	Description
< >	Nonprinting characters, such as passwords, appear in angled brackets.
[ ]	Default responses to system prompts appear in square brackets.



BYOD Solution Overview

Revised: August 7, 2013

Introduction

Bring Your Own Device (BYOD) continues to be one of the most influential trends reshaping the landscape of the mobile enterprise and the evolution of IT organizations. The influx of powerful mobile devices into the workplace is changing how users access and consume enterprise resources. IT managers are establishing policies with BYOD access as the norm rather than the exception due to increasing demands from employees and executives who embrace this megatrend. Enterprises are beginning to see BYOD as an opportunity rather than a challenge. There is no longer any doubt that enterprise IT departments are adapting to mobile devices (smartphones, tablets, laptops, etc.) in the corporate workplace to meet user expectations and leverage new technologies to boost worker productivity.

IT needs to balance productivity with security and coordinate business justification with the various line of business (LOB) owners to implement BYOD programs within an enterprise. On one hand, employees are demanding access from devices not only within the corporation, but also beyond the firewall. On the other hand, risk management dictates that corporate data must remain protected

The Cisco BYOD Smart Solution delivers a unified workspace that increases workforce productivity with high quality collaboration on any device, anywhere. Cisco BYOD Smart Solution is a complete, yet flexible and secure BYOD solution that one can easily tailor to meet an enterprise's needs. Cisco customers get proven solution designs that are fully system tested and documented in a Cisco Validated Design (CVD). The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments.

When BYOD is properly implemented, it delivers an uncompromising, work-your-way user experience and enables organizations to secure data with unified policies and essential controls.

BYOD Market Landscape

Mobile devices are turning into general computing devices with faster components and better integration and support in connecting to enterprise systems. They are being used by employees to perform their regular job functions. In many cases, each employee owns multiple devices which they use for personal as well as business work.

The BYOD market was initially centered on allowing employees, partners, and guests to connect to the corporate networks and be able to perform a limited number of basic functions, such as access the Internet, access corporate email, calendar, and contacts, etc. As more and more personal devices become

prevalent in the enterprise, workers are using these devices to access enterprise resources and applications to do their daily jobs. With employees using personal devices for mission critical job functions, mobile device managers (MDM) are becoming increasingly important. Functions such as ensuring that a device can be locked and wiped remotely in case it gets lost or stolen or when the employee is terminated are becoming a necessity.

An enterprise's BYOD strategy should build momentum towards meaningful infrastructure, security, and wireless innovation and provide a solid rationale for BYOD investments.

To understand the benefits and the challenges BYOD poses, it is helpful to understand the business trends that are hindering or driving BYOD adoption.

BYOD and the Enterprise Network

The network is at the heart of all business functions and is a key enabler for service delivery. With the proliferation of mobile devices, protecting enterprise networks is getting more complex. Inconsistent management tools and policies across the wired and wireless segments of the network can increase the burden for network managers and drive up management costs and complexity. An architectural approach and design can help businesses realize greater levels of manageability, enhance user/customer experience, provide greater levels of flexibility and performance, and achieve improvements in security and policy.

Today enterprises have diverse wired and wireless LAN infrastructure implementations. The traditional workforce model involves a worker going to an office and performing their job functions while tethered to an IT-provided computer and an IP phone. Hence the focus has been on ensuring that the wired enterprise infrastructure is robust and secure. The wired network is built for high availability and performance, with adequate capacity and intelligent features such as QoS to make phone calls, exchange data and video, etc., within and outside of the enterprise. Communication within the premises is based on a trust model, protected by firewalls at the perimeter and other security tools.

In contrast, the enterprise wireless infrastructure was built for convenience, at least until the BYOD phenomenon emerged. With BYOD, users will have three or more mobile/WLAN devices (laptop, tablet, Smartphone) all connecting to this infrastructure. The devices themselves could be user-owned or corporate-owned. What is the trust level of these devices that are accessing secure enterprise resources? How are they accounted for at all times? With all these questions demanding answers, one can expect a lot more focus on the wireless network. The WLAN needs to become as robust, secure, scalable, and predictable as the wired network to support BYOD.

BYOD Security and Policy-Many enterprise CIOs and architects cite security as a major inhibitor in the deployment, adoption, and acceleration of BYOD within their enterprise. It can be overwhelming to manage security for various applications for a huge set of new BYOD devices with diverse operating systems and users. Based on each customer's environment and business policy, there are many options for implementing and configuring personal and network firewalls, intrusion prevention, anti-spyware, data security, device control for smartphones, desktops, and tablets for various user roles such as employees, guests, partners, and remote workers.

Policy management is critical to a successful BYOD strategy which enables endpoint security for multiple devices and diverse user roles. Application management and security are on the rise in a BYOD environment as enterprises seek to manage and secure applications rather than devices, with more granular policies applied based on worker or application type. More companies want enterprise application stores that allow for a central deployment method. The security needs of a mobile workforce include:

- Ensuring that the devices used to access the corporate network are safe and are not jail-broken or rooted. They should not have threatening malware, spam, or applications that can compromise the corporate network or data.

- Making sure users and devices that are accessing the corporate network can be identified and allowed connectivity only if they are authorized and meet company policy.
- Securing access with client-based or clientless access to ensure data loss prevention with encryption or containerization with VPN optimized for efficient application delivery and capabilities.
- Enforcing device-level security functionality such as remote wipe/lock with integrated Network Access Control (NAC), thus ensuring that an action can be taken on non-compliant devices at any time (not just during access).
- Having visibility into users, devices, and the applications they are running on the corporate network.

Mobility—The hottest segment of the enterprise networking market is wireless LAN equipment to support enterprise requirements for user mobility and wireless devices. As smartphones and tablets become more powerful, they are increasingly used by employees to connect to the WLAN and other enterprise resources to do their normal job functions. It then becomes important for the WLAN to have same or similar intelligence and feature functionality, such as high availability, QoS, local switching, application visibility and control, ease of collaboration, etc., that users have become accustomed to on a wired infrastructure.

Role of MDM in the Enterprise Network

Mobile Device Managers are designed to help an enterprise rapidly and securely deploy mobile devices and applications with policy, compliance, configuration, and application management to minimize risk. Large enterprises are extremely interested in delivering general-purpose as well as custom-built corporate applications to their workforce (for example, mobile sales force automation applications).

BYOD security and device management are the foundations of an enterprise BYOD strategy which must consider all mobile worker types and functions before deploying solutions. Organizations need to consider solutions across the security sub-segments that secure endpoints, provide protection for the corporate network, and protect data as it moves over their infrastructure.

The ability to integrate MDM functionality, coupled with a policy-based network access, ensures that legitimate devices and users can access the network and attempted violations can be controlled by prohibiting or limiting access to the network or network resources.

User Need and Role of IT

BYOD connectivity may look like a simple extension of enterprise mobile services, however broad user expectations and the diversity of devices create unique infrastructure demands and challenges for IT operations with end-to-end and network optimization and lifecycle management to support BYOD.

Device choice does not mean sacrificing security. IT must establish the minimum security baseline that any device must meet to be used on the corporate network, including WiFi security, VPN access, and perhaps add-on software to protect against malware.

In addition, due to the wide range of devices, it is critical to be able to identify each device connecting to the network and authenticate both the device and the person using it.

Protecting Data and Loss Prevention

One of the largest challenges with any BYOD implementation is ensuring protection of corporate data. If a corporate asset, such as a laptop, is used to access business applications and data, typically that asset is tightly controlled by IT and likely subject to more restrictive usage policies.

Some industries need to comply with confidentiality regulations like HIPAA, security compliance regulations like PCI, or more general security practice regulations like Sarbanes-Oxley and others. Companies need to show compliance is possible with BYOD adoption, which can be more challenging than with a corporate-owned and managed device.

An employee-owned mobile device is likely being routinely used for personal access and business applications. Cloud-based file sharing and storage services are convenient for personal data, but can be potential sources of leakage for confidential corporate data.

IT must have a strategy for protecting business data on all devices whether corporate managed or employee self-supported and managed. This may include a secure business partition on the device which acts as a container of corporate data that can be tightly controlled and may also include the need for a Virtual Desktop Infrastructure (VDI) application to allow access to sensitive or confidential data without storing the data on the device.

At some point in the lifecycle of a device or employee, it may become necessary to terminate access to the device or the device's access to the network. This could be due to a lost or stolen device, an employee termination, or even an employee changing roles within the company. IT needs the ability to quickly revoke access granted to any device and possibly remotely wipe some or all of the data (and applications) on the device.

Challenges for End Users

BYOD solutions and technologies are quickly evolving, however one of the largest challenges is how to make it simple for people to get connected to and use corporate resources and applications to do their work. The number of device possibilities, the range of connection types and locations, and the lack of widely adopted approaches can translate to difficulties for users.

Security precautions and steps may also vary depending upon how and where the user is trying to connect. For example, the corporate WiFi may require credentials, whereas connecting through a public WiFi hotspot may require credentials, a virtual private network (VPN), and other security steps. If such security measures are too intrusive, they could erase productivity gains of BYOD.

Key Advantages of the Cisco BYOD Solution

- Cisco networks are integrated into one control panel, resulting in greater security and ease of management.
- The Cisco Solution relies on a centralized policy server that integrates tightly with an organization's Active Directory and PKI infrastructure.
- Fully integrated, centralized, single point of visibility and control of users, devices, location, network, and applications resulting in greater security and ease of management.
- Services across end-to-end partner ecosystem—Pre-integrated and tested solutions with world-class partners.

The Cisco BYOD solution integrates the Cisco products, third-party products, and devices discussed previously into a comprehensive BYOD approach which is tightly integrated across the network infrastructure. This offers a unique set of advantages such as flexibility to allow for multiple diverse user groups, including deskbound workers, mobile workers, customers, guests, etc.



PART 2

BYOD Design Overview



Summary of Design Overview

Revised: August 7, 2013

This part of the CVD describes design considerations to implement a successful BYOD solution and different deployment models to address diverse business cases. Other parts of the CVD provide more details on how to implement unique use cases.

There are numerous ways to enable a BYOD solution based on the unique business requirements of a specific organization. While some organizations may take a more open approach and rely on basic authentication, other organizations will prefer more secure ways to identify, authenticate, and authorize devices. A robust network infrastructure with the capabilities to manage and enforce these policies is critical to a successful BYOD deployment.

The Cisco BYOD solution builds on the Cisco Borderless Network Architecture and assumes best practices are followed in network infrastructure designs for campus, branch offices, Internet edge, and converged access implementations. The solution showcases the critical components to allow secure access for any device, ease of accessing the network, and centralized enforcement of company usage policies. This robust architecture supports a multitude of wired or wireless devices, both employee-owned and corporate-owned, accessing the network locally or from remote locations, as well as on-premise guest users.

This part of the CVD includes the following chapters:

- Cisco BYOD Solution Components—This section highlights the different network components used in the design guide. These components provide a solid network infrastructure required as the enforcement point for BYOD policies. Because of the reliance on digital certificates, a discussion regarding the secure on-boarding of devices is included in this section.
- BYOD Use Cases—This CVD addresses four different use cases based on the type of network access allowed by an organization. These use cases vary from personal, corporate-owned, and guest access. Permissions are enforced using Active Directory credentials, digital certificates, ISE identity groups, and other unique identifiers.
- Campus and Branch Network Design for BYOD—Policy enforcement is effective if and only if there is a well-designed network infrastructure in place. This section describes different campus and branch designs used to support BYOD, including WAN infrastructure, FlexConnect, and Converged Access.
- Mobile Device Managers for BYOD—The section introduces the ISE integration with different third-party Mobile Device Managers and explores different deployment models.

- Application Considerations and License Requirements for BYOD—This section highlights different requirements that need to be present to provide the proper network service to applications. These include features such as Quality of Service, Rate Limiting, Application Visibility and Control (AVC), and others. The chapter also highlights Cisco Jabber and Virtual Workspace (VXI) architecture.



Cisco BYOD Solution Components

Revised: August 7, 2013

Cisco provides a comprehensive BYOD solution architecture, combining elements across the network for a unified approach to secure device access, visibility, and policy control. To solve the many challenges described earlier, a BYOD implementation is not a single product, but should be integrated into an intelligent network.

The following figures show a high-level illustration of the Cisco BYOD solution architecture. The architecture has been separated into campus and branch diagrams simply for ease of viewing. These infrastructure components are explained in detail in the following sections.

Figure 3-1 HIgh-Level BYOD Solution Architecture—Campus View

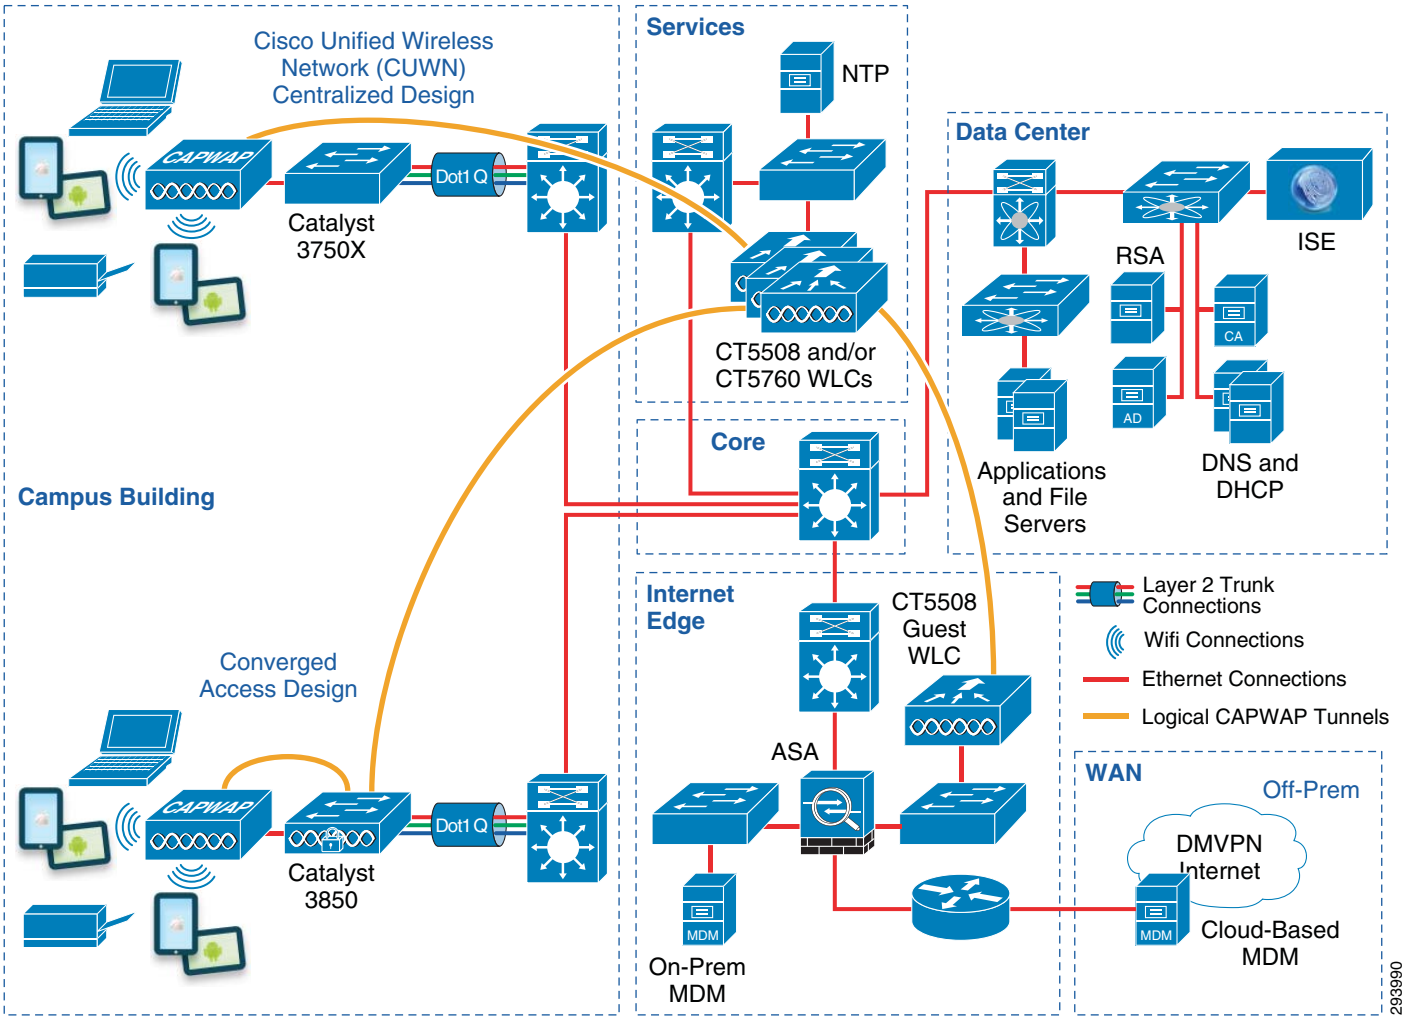
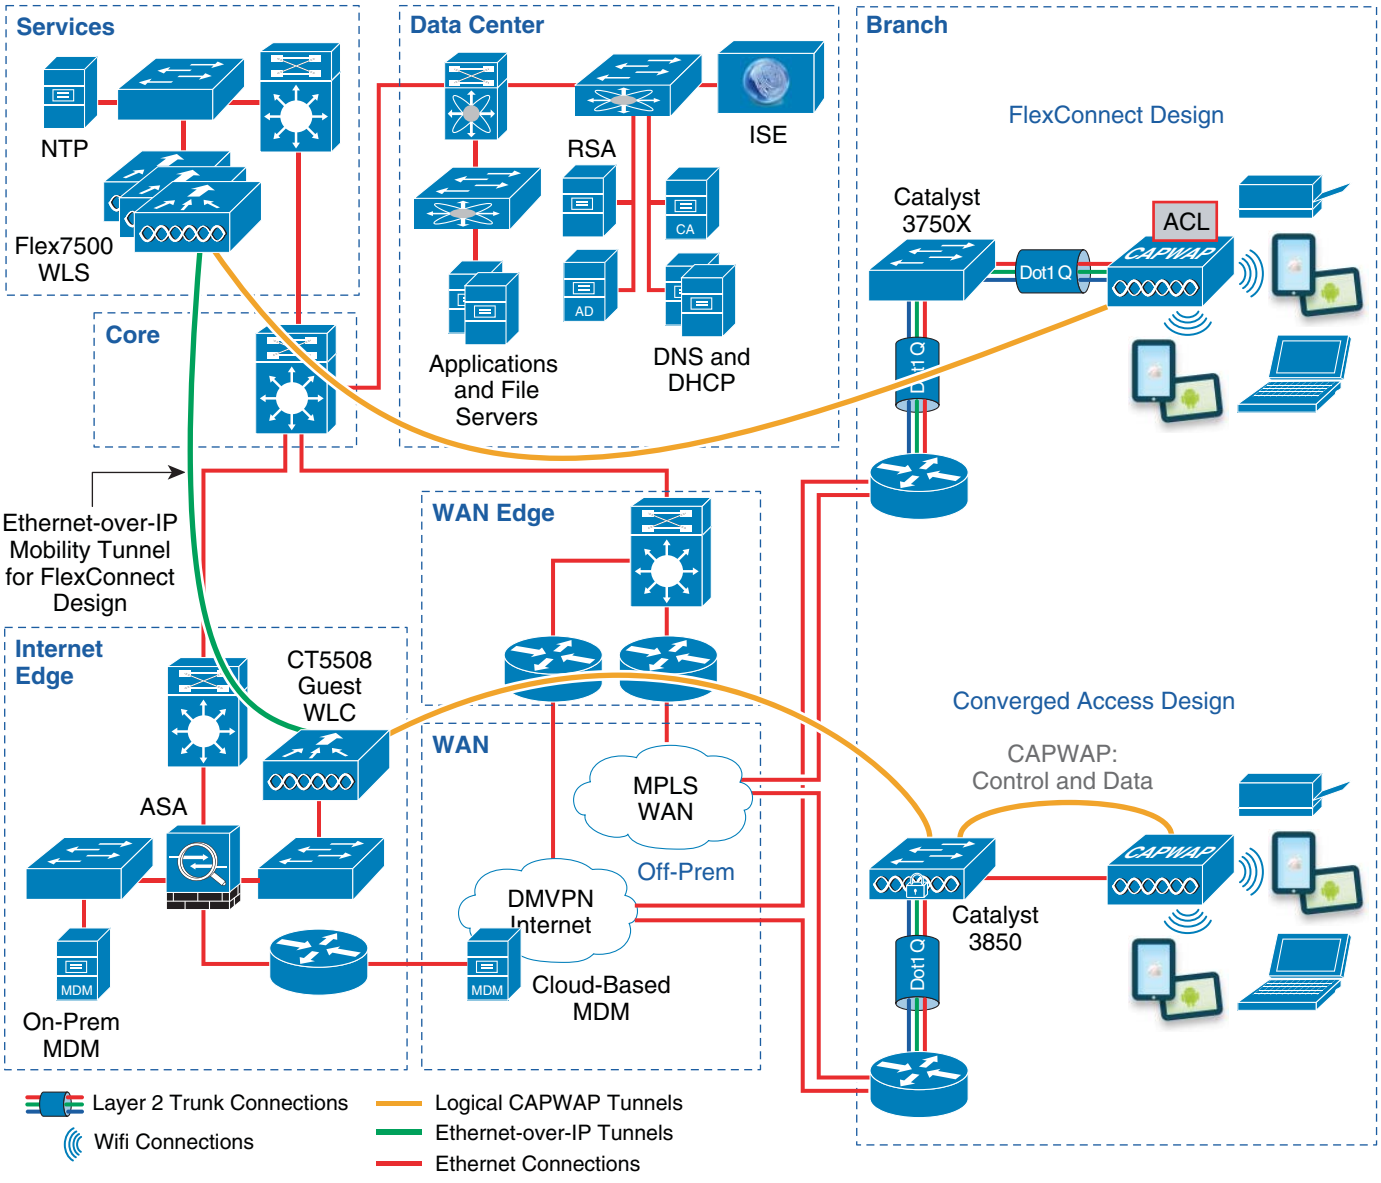


Figure 3-2 High-Level BYOD Branch Solution Architecture—Branch View



Cisco Access Points

Cisco Access Points provide WiFi connectivity for the corporate network and handle authentication requests to the network via 802.1X. In addition, the Cisco Access Points at the branch location can either tunnel all the traffic to the campus or switch traffic locally based on the configuration.

Cisco Wireless Controller

Cisco Wireless LAN Controller (WLC) is used to automate wireless configuration and management functions and to provide the visibility and control of the WLAN. The WLC extends the same access policy and security from the wired network core to the wireless edge while providing a centralized access point configuration. The WLC interacts with the Cisco Identity Services Engine (ISE) to enforce authentication and authorization policies across device endpoints. Multiple WLCs may be managed and monitored by Cisco Prime Infrastructure. Wireless LAN Controller functionality can be within standalone appliances, integrated within Catalyst switch products, or run virtually on Cisco Unified Computing System (UCS). Integrated controller functionality is discussed in Converged Access Campus Design in Chapter 5, "Campus and Branch Network Design for BYOD."

Cisco Identity Services Engine

Cisco Identity Services Engine (ISE) is a core component of the Cisco BYOD solution architecture. It delivers the necessary services required by enterprise networks, such as Authentication, Authorization, and Accounting (AAA), profiling, posture, and guest management on a common platform. The ISE provides a unified policy platform that ties organizational security policies to business components.

The ISE also empowers the user to be in charge of on-boarding their device through a self-registration portal in line with BYOD policies defined by IT. Users have more flexibility to bring their devices to their network with features such as sponsor-driven guest access, device classification, BYOD on-boarding, and device registration.

The ISE is able to integrate with third-party Mobile Device Managers (MDM) to enforce more granular policies based on device posture received from the MDM compliance rules.

Cisco Adaptive Security Appliance

Cisco Adaptive Security Appliance (ASA) provides traditional edge security functions, including firewall and Intrusion Prevention System (IPS), as well as providing the critical secure VPN (AnyConnect) termination point for mobile devices connecting over the Internet, including home offices, public WiFi hotspots, and 3G/4G mobile networks. The ASA delivers solutions to suit connectivity and mobility requirements for corporate-owned devices as well as employee-owned laptops, tablets, or mobile devices.

Cisco AnyConnect Client

Cisco AnyConnect™ client provides 802.1X supplicant capability on trusted networks and VPN connectivity for devices that access the corporate network from un-trusted networks, including public Internet, public WiFi hotspots, and 3G/4G mobile networks. Deploying and managing a single supplicant client has operational advantages as well as provides a common look, feel, and procedure for users.

In addition, the AnyConnect client can be leveraged to provide device posture assessment of the BYOD device, as well as a degree of policy enforcement and enforcing usage policies.

The AnyConnect client can be provisioned centrally with use of a third-party MDM. This enhances the user experience and reduces the support costs. MDM policy can be configured to manage who is entitled to use AnyConnect.

Cisco Integrated Services Routers

Cisco Integrated Services Routers (ISR), including the ISR 2900 and ISR 3900 families, provide WAN and LAN connectivity for branch and home offices. The LAN includes both wired and wireless access. In addition, ISRs may provide direct connectivity to the Internet and cloud services, application and WAN optimization services, and may also serve as termination points for VPN connections by mobile devices.

Cisco Aggregation Services Routers

Cisco Aggregation Services Routers (ASR), available in various configurations, provide aggregate WAN connectivity at the campus WAN edge. In addition, ASRs may provide direct connectivity to the Internet and cloud services and may also serve as a firewall. The ASR runs Cisco IOS XE software and offers Flexible Packet Matching (FPM) and Application Visibility and Control (AVC).

Cisco Catalyst Switches

Cisco Catalyst® switches, including the Catalyst 3000, Catalyst 4000, and Catalyst 6000 families, provide wired access to the network and handle authentication requests to the network via 802.1X. In addition, when deployed as access switches, they provide power-over-Ethernet (PoE) for devices such as VDI workstations, IP phones, and access points.

Cisco Converged Access Switches

Cisco Catalyst 3850 Series switches provide converged wired and wireless network access for devices. As a switch, the Catalyst 3850 provides wired access to the network and handles authentication requests to the network via 802.1X. In addition, the Catalyst 3850 contains wireless LAN controller functionality integrated within the platform. As a wireless controller, it allows for the termination of wireless traffic from access points directly attached to the Catalyst 3850 switch, rather than backhauling wireless traffic to a centralized wireless controller. This can provide greater scalability for wireless traffic, as well as provide increased visibility of wireless traffic on the switch. The Catalyst 3850 Series switch interacts with Cisco ISE to enforce authentication and authorization policies across device endpoints, providing a single point of policy enforcement for wired and wireless devices.

When deployed at the access-layer within a branch location, the Catalyst 3850 can be configured to function as both a Mobility Controller (MC) and a Mobility Agent (MA), providing full wireless controller functionality. When deployed within a large campus, the Catalyst 3850 can be configured to function as a Mobility Agent (MA), which allows for the termination of wireless traffic directly on the switch itself. For increased scalability, the Mobility Controller (MC) function, which handles Radio Resource Management (RRM), Cisco CleanAir, and roaming functions, among other things, can be moved to a dedicated CT5760 or CT5508 wireless controller. Both the Catalyst 3850 and the CT5760 wireless controller run IOS XE software, allowing for the full feature richness of Cisco IOS platforms.

Appendix C, “Software Versions” discusses the feature sets and licensing required for wireless controller functionality on the Catalyst 3850 Series platform.

Cisco Nexus Series Switches

Cisco Nexus switches, including the Nexus 7000 and 5000 families, serve as the data center switches within the CVD. The Nexus 7000 switches provide 10GE Layer 3 connectivity between the Campus Core, Data Center Core, and Aggregation Layers and 10GE Layer 2 connectivity, utilizing VPC, for the Nexus 5000 switches in the Data Center Access Layer to which all servers are attached.

Cisco Prime Infrastructure

Cisco Prime Infrastructure (PI) is an exciting new offering from Cisco aimed at managing wireless and wired infrastructure while consolidating information from multiple components into one place. While allowing management of the infrastructure, Prime Infrastructure gives a single point to discover who is on the network, what devices they are using, where they are, and when they accessed the network.

Cisco Prime Infrastructure 1.2 is the evolution of Cisco Prime Network Control System 1.1 (NCS). It provides additional infrastructure and wired device management and configuration capabilities while improving on existing capabilities in NCS 1.1.

Cisco Prime Infrastructure interacts with many other components to be a central management and monitoring portal. Prime Infrastructure has integration directly with two other appliance-based Cisco products, the Cisco Mobility Services Engine (MSE) and Identity Services Engine (ISE) for information consolidation. Prime Infrastructure controls, configures, and monitors all Cisco Wireless LAN Controllers (WLCs), and by extension, all Cisco access points (APs) on the network. Prime Infrastructure also configures and monitors Cisco Catalyst switches and Cisco routers.

Secure Access to the Corporate Network

On-boarding for new devices (certificate enrollment and profile provisioning) should be easy for end users with minimal intervention by IT, especially for employee owned devices. Device choice does not mean giving up security. IT needs to establish the minimum security baseline that any device must meet to access the corporate network. This baseline should include WiFi security, VPN access, and add-on software to protect against malware. Proper device authentication is critical to ensure secure on-boarding of new devices and to ensure secure access to other devices on the network. Hence, proper device authentication protects the entire network infrastructure.

Who is accessing the network, *what* device they are using, and *where* they are located need to be considered before implementing a BYOD solution. The user can initiate the provisioning process from a campus or a branch location. This design allows the user to provision and access resources from either location. In the past, a username/password was all that was needed as most employees accessed the network from a wired workstation. Often a simple server was used to collect and authenticate user credentials. As organizations implemented wireless into their network, a unique SSID (Wireless Network name) with a username and password was also needed.

Today, digital certificates and two-factor authentication provide a more secure method to access the network. Typically the end user must download client software to request a certificate and/or provide a secure token for access. One of the challenges with deploying digital certificates to client endpoints is that the user and endpoint may need to access the company's certification authority (CA) server directly (after being authenticated to the corporate network) to manually install the client certificate. This method requires the end user manually install the client certificate and ensuring it is installed in the proper certificate store on the local endpoint.

Deploying digital certificates on non-PC based devices requires a different process as many of these devices do not natively support all the features and functionality needed to create/download and install digital client certificates. As users become more and more mobile, authenticating users and devices accessing the network is an important aspect of BYOD.

Certificate Enrollment and Mobile Device Provisioning

Deploying digital certificates to endpoint devices requires a network infrastructure that provides the security and flexibility to enforce different security policies, regardless of where the connection originates. This solution focuses on providing digital certificate enrollment and provisioning while enforcing different permission levels. This design guide covers Android™ and Apple® iOS™ mobile devices, in addition to Windows 7 and Mac OS X.

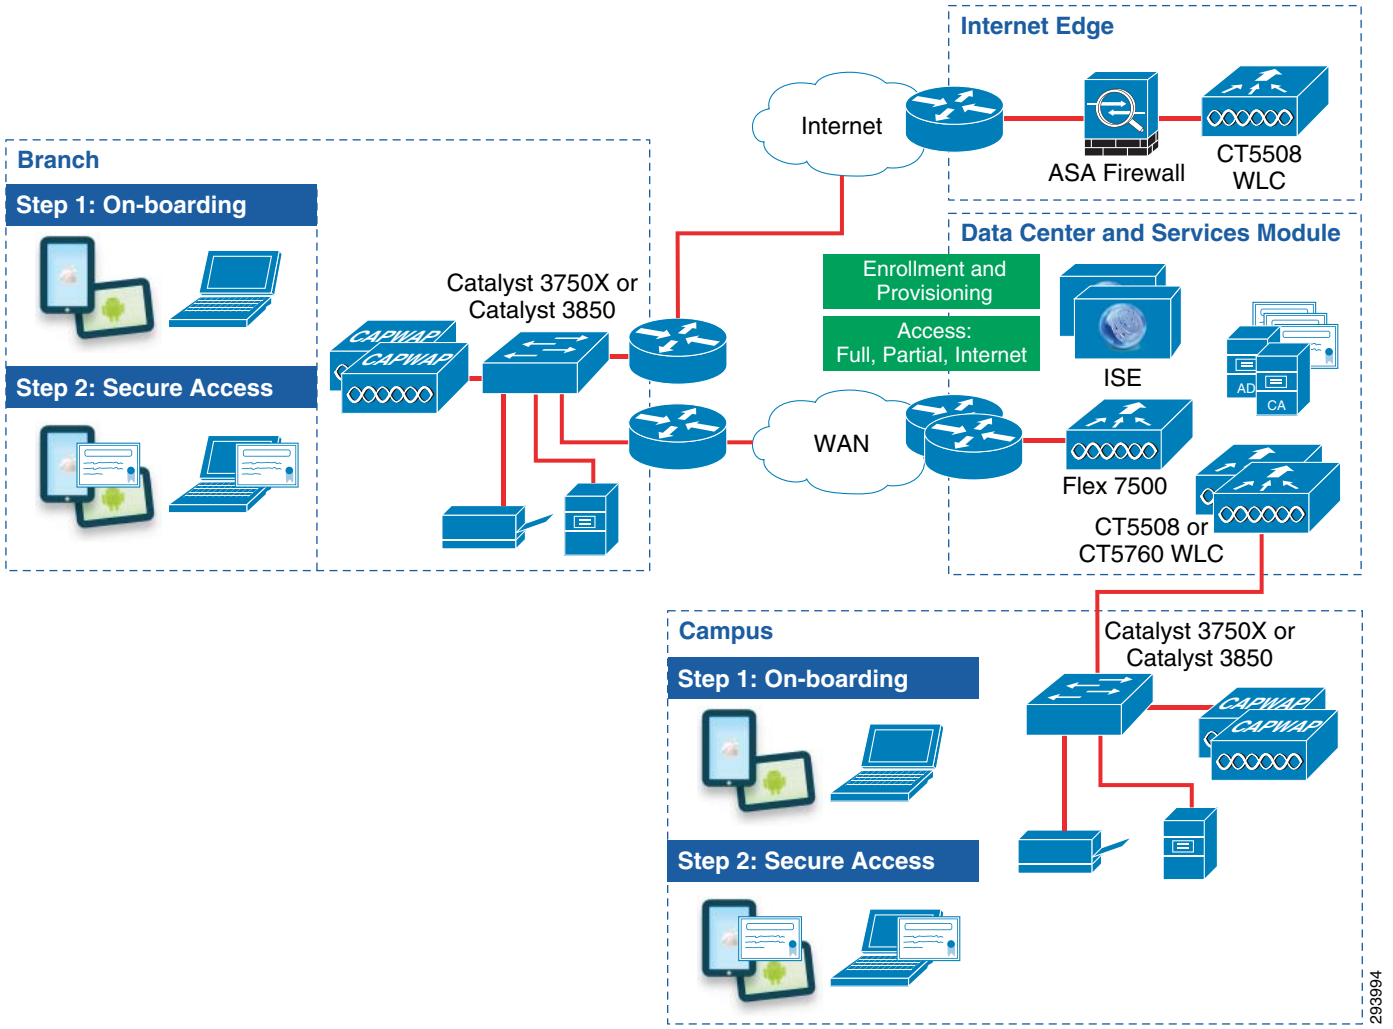
Figure 3-3 highlights the general steps that are followed for this solution when a mobile device connects to the network:

1. A new device connects to a provisioning SSID, referred to as the BYOD_Provisioning SSID. This SSID (open or secured with PEAP) is configured to redirect the user to a guest registration portal.
2. The certificate enrollment and profile provisioning begins after the user is properly authenticated.
3. The provisioning service acquires information about the mobile device and provisions the configuration profile, which includes a WiFi profile with the parameters to connect to a secure SSID, called the BYOD_Employee SSID.
4. For subsequent connections, the device uses the BYOD_Employee SSID and is granted access to network resources based on different ISE authorization rules.

The design guide also covers a single SSID environment, where the same SSID is used for both provisioning and secure access.

Employee devices that do not go through the provisioning process simply connect to a guest SSID, a or dedicated guest-like SSID; which may be configured to provide Internet-only or limited access for guests or employees.

Figure 3-3 Enrollment and Provisioning for Mobile Devices





BYOD Use Cases

Revised: August 7, 2013

An organization's business policies will dictate the network access requirements which their BYOD solution must enforce. The following four use cases are examples of access requirements an organization may enforce:

- Enhanced Access—This use case provides network access for personal devices, as well as corporate issued devices. It allows a business to build a policy that enables granular role-based application access and extends the security framework on and off-premises.
- Limited Access—This use case enables access exclusively to corporate-issued devices.
- Advanced Access—This comprehensive use case also provides network access and for personal and corporate issued devices. However it includes the posture of the device into the network access control decision through integration with third party Mobile Device Managers (MDMs).
- Basic Access—This use case is an extension of traditional wireless guest access. It represents an alternative where the business policy is to not on-board/register employee wireless personal devices, but still provides Internet-only or partial access to the network.

ISE evaluates digital certificates, Active Directory group membership, device type, etc. to determine which network access permission level to apply. ISE provides a flexible toolset to identify devices and enforce unique access based on user credentials and other conditions.

Figure 4-1 shows the different permission levels configured in this design guide. These access levels may be enforced using access lists in the wireless controller or Catalyst switches, assigning Security Group Tags (SGTs) to the device traffic or by relying on dynamic virtual LAN (VLAN) assignment. The design guide shows different ways to enforce the desired permissions.

Figure 4-1 Permission Levels

	Permission	Access
	Full Access	Internet plus all corporate resources
	Partial Access	Internet plus some corporate applications
	Internet Only	Internet Only
	Deny Access	Explicitly deny network access

292599

Enhanced Access—Personal and Corporate Devices

This use case builds on the Limited Access use case and provides the infrastructure to on-board personal devices onto the network by enrolling digital certificates and provisioning configuration files. The use case focuses on how to provide different access levels to personal devices based on authentication and authorization rules.

Employees that have registered their devices using the self-registration portal and have received a digital certificate are granted unique access based on their Active Directory group membership:

- Full Access—If the employee belongs to the BYOD_Full_Access Active Directory group.
- Partial Access—If the employee belongs to the BYOD_Partial_Access Active Directory group.
- Internet Access—If the employee belongs to the Domain Users Active Directory group.

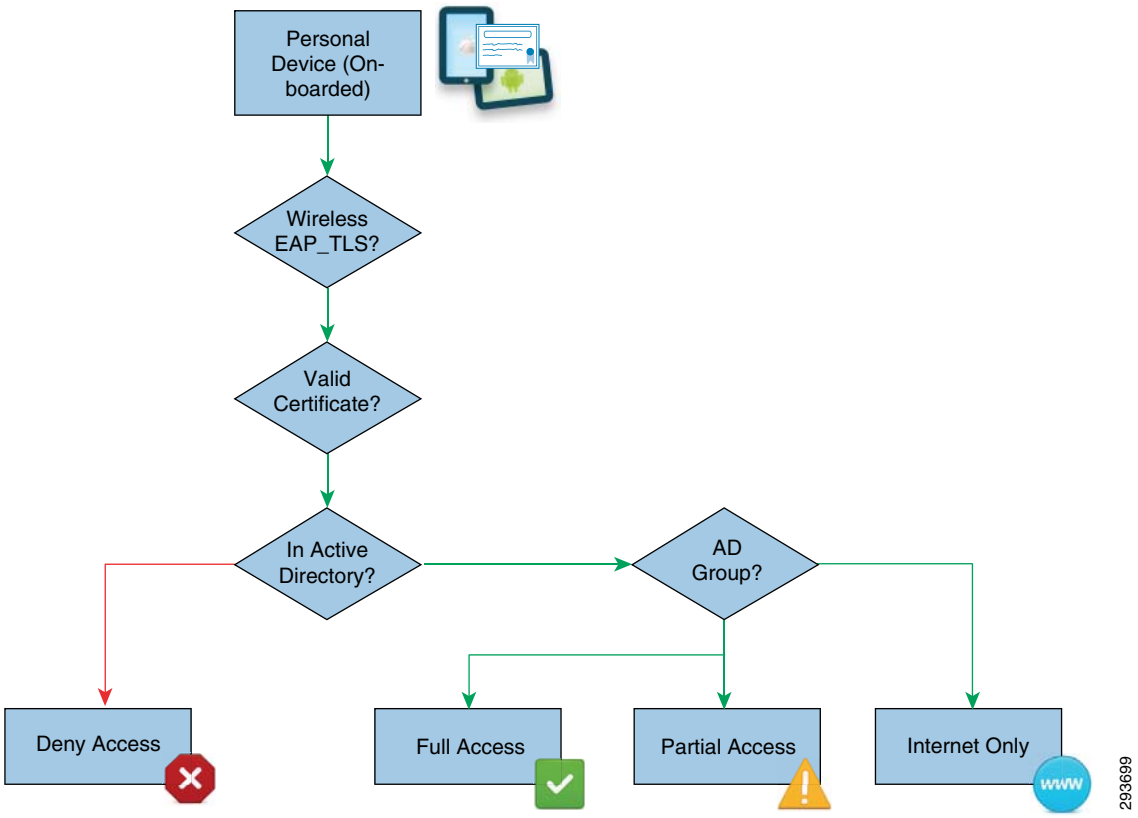
Corporate owned devices are granted full access in this use case.

The use case also explains how to prevent personal owned devices, for example Android devices, from accessing the network. Some organizations may not be ready to allow employees to connect their personal devices into the network and may decide to block their access until business or legal requirements are met. Cisco ISE provides the capability of identifying (profiling) the device type and preventing those devices from connecting to the network. As an example, this use case includes device profiling in ISE to deny access to Android devices.

The use of Security Group Tags will be used as an alternative to ACLs for enforcing role-based policies for campus wireless users and devices. Security Group Tags provide a complimentary technology offering a scalable approach to enforcing policy and traffic restrictions with minimal and in some cases, little or no ACLs at all if TCP/UDP port level granularity is not required.

Figure 4-2 highlights the connectivity flow for personal devices.

Figure 4-2 Personal Devices BYOD Access



This use case provides an effective way for organizations to embrace a BYOD environment for their employees and provide differentiated access to network resources.

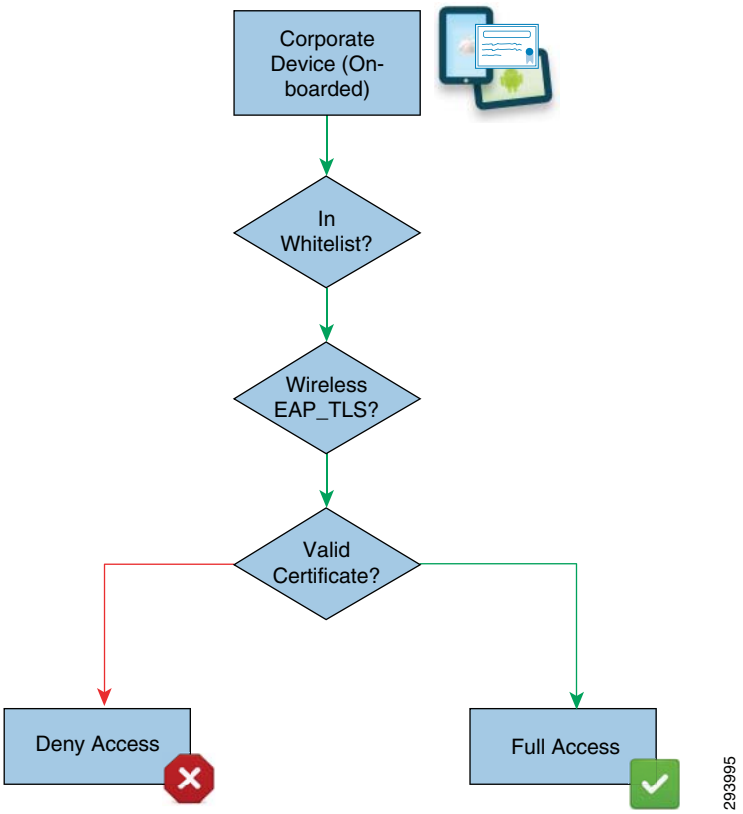
Limited Access—Corporate Devices

This use case applies to organizations that decide to enforce a more restrictive policy that allows only devices owned or managed by the corporation to access the network and denies access to employee personal devices.

ISE grants devices full access to the network based on the device's certificate and inclusion in the Whitelist identity group. This use case introduces the use of a Whitelist, a list of corporate devices maintained by the Cisco ISE that is evaluated during the authorization phase.

Figure 4-3 shows connectivity flow for corporate devices.

Figure 4-3 Corporate Device BYOD Access



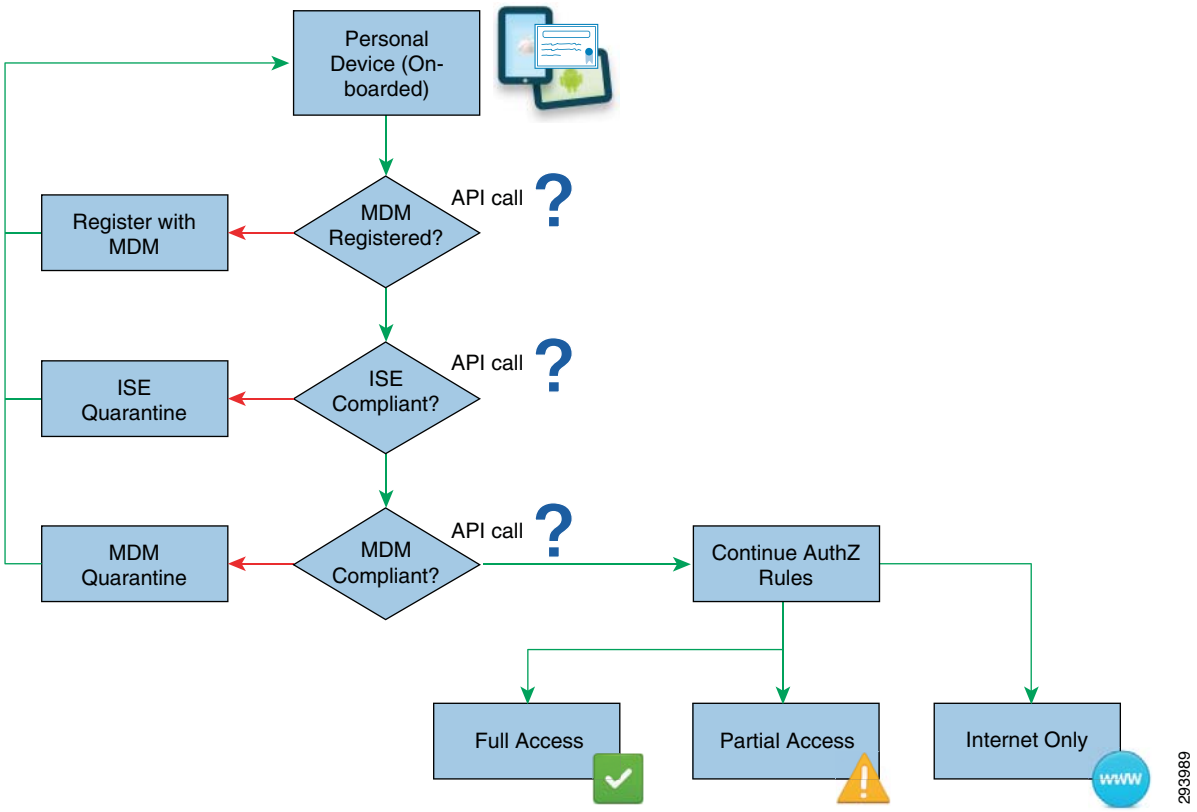
Advanced Access—MDM Posture

This use case applies to organizations that have invested in a Mobile Device Manager (MDM) to manage and secure mobile endpoints. While MDMs are not able to enforce Network Access Control policies, they provide unique device posture information not available on the ISE. Combining ISE policies with additional MDM information, a robust security policy may be enforced on mobile endpoints.

The integration between ISE and third-party MDMs is through a REST API, allowing the ISE to query the MDM for additional compliance and posture attributes.

Figure 4-4 shows the connectivity flow to obtain MDM compliance information and network access.

Figure 4-4 MDM Compliance



Basic Access—Guest-Like

Some organizations may implement a business policy which does not on-board wireless employee personal devices, yet provides some access to corporate services and the Internet for such devices. Some of the possible reasons include:

- The organization does not have the desire or the ability to deploy digital certificates on employees' personal devices.
- The employees may be unwilling to allow the organization to "manage" their personal device.
- The organization does not wish to manage and maintain separate lists of registered devices or manage a user's network access level when using personal devices.

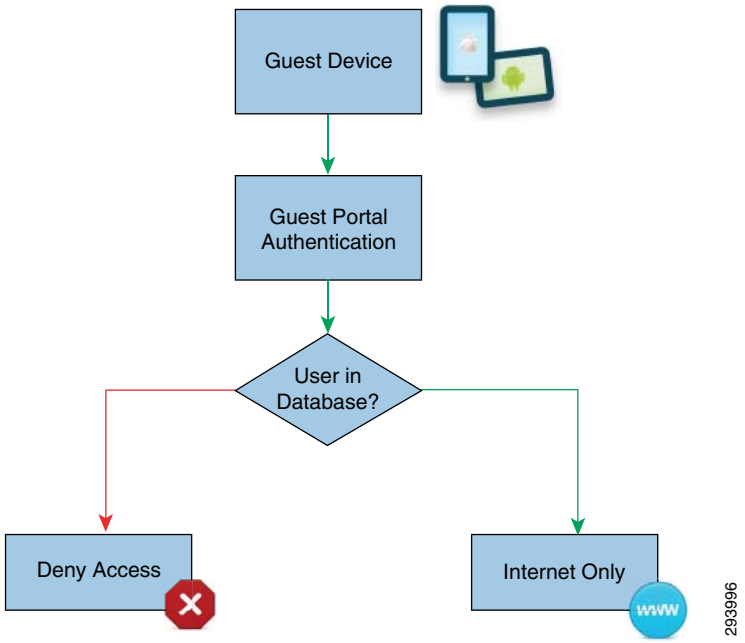
The design for this use case is based around extending traditional guest wireless access and providing similar guest-like wireless access for employee personal devices. The design guide focuses on two methods for extending guest wireless access to allow employee personal devices access to the guest network:

- Allowing employees to provision guest credentials for themselves.
- Extending guest web authentication (Web Auth) to also utilize the Microsoft Active Directory (AD) database when authenticating guests or employees using personal devices.

In addition, the design guide discusses another option in which a second guest-like wireless SSID is provisioned for employee personal devices.

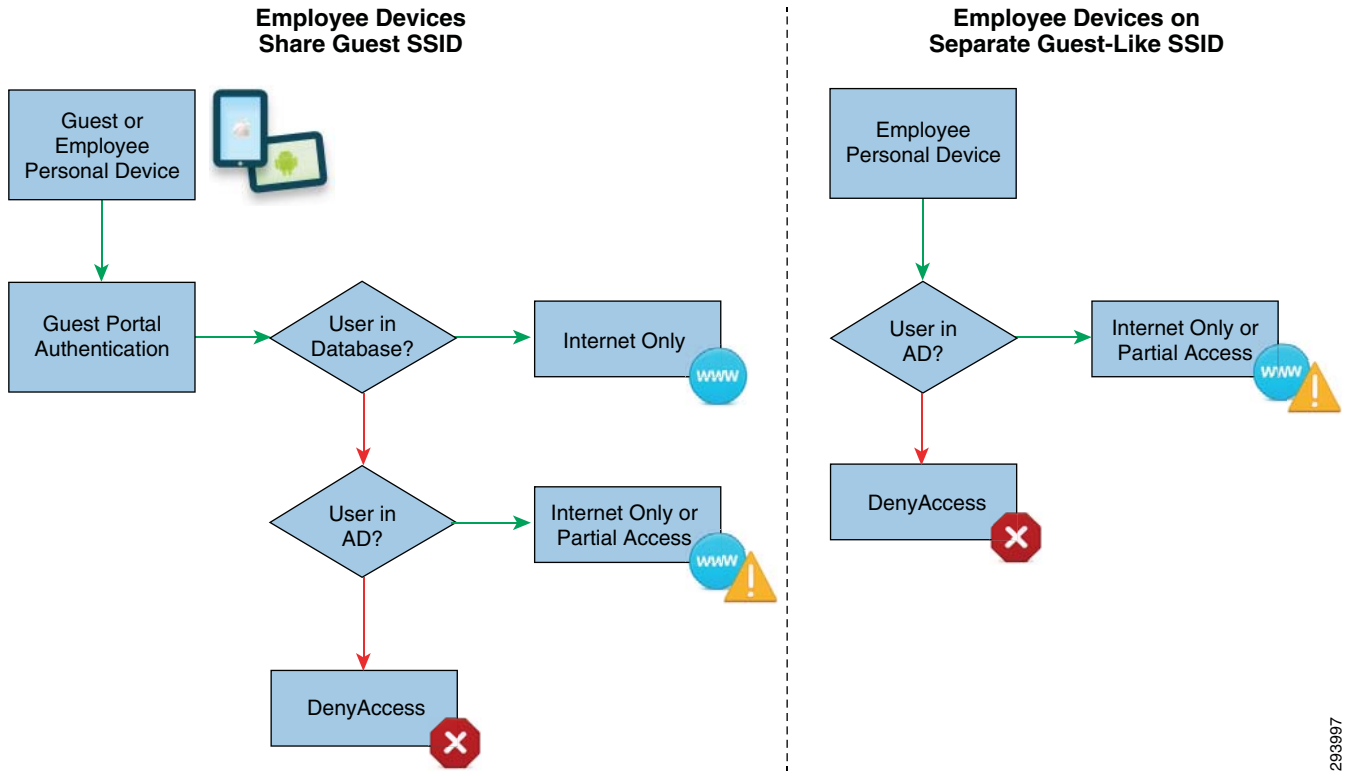
The Basic Access use case builds on traditional wireless guest access. Figure 4-5 shows the typical method for authenticating a device connecting to the guest wireless network.

Figure 4-5 Guest Wireless Access



This design guide discusses two approaches for modifying an existing guest wireless access implementation to enable Basic Access for employee personal devices, as shown in Figure 4-6.

Figure 4-6 Basic Access





CHAPTER 5

Campus and Branch Network Design for BYOD

Revised: August 7, 2013

Campus Network Design

As with the branch design, policy enforcement is effective if and only if there is a well-designed campus network infrastructure in place. This section discusses the high-level key design elements of campus LAN design.

The two wireless LAN designs for the campus which will be discussed within this design guide are Centralized (Local Mode) and Converged Access designs.

Centralized (Local Mode) Wireless Design

Cisco Unified Wireless Network (CUWN) Local Mode designs, refer to wireless LAN designs in which all data and control traffic is backhauled from the access point to a wireless controller before being terminated and placed on the Ethernet network. This type of design is also referred to as a centralized wireless design or centralized wireless overlay network. A typical recommended design within a large campus is to place all of the wireless controllers into a separate services module connected to the campus core.

The potential advantages of this design are:

- Centralized access control of all wireless traffic from a single point within the campus network.
- Less complexity for wireless roaming, since the wireless controllers can share larger IP address pool for wireless clients.

The potential disadvantages of this design are:

- Potential for scalability bottlenecks at the wireless controllers or the network infrastructure connecting to the wireless controllers. This is because all wireless traffic is backhauled to a central point within the campus network where the wireless controllers are deployed, before being terminated on the Ethernet network. Note however, that this may be alleviated by deploying additional centralized wireless controllers, by upgrading to newer platforms such as the Cisco CT5760 wireless controller, and/or by moving wireless controllers out to the building distribution modules.
- Less visibility of wireless traffic, since the wireless traffic is encapsulated within a CAPWAP tunnel as it crosses the campus network infrastructure.

With a Local Mode design, access points that are connected to the access-layer switches within the building distribution modules are configured and controlled via one or more centralized Wireless LAN Controllers. In the case of this design guide, these controllers are a set of Cisco CT5508 wireless controllers—dedicated for the campus—since they provide greater scalability for supporting Local Mode access points than Cisco Flex 7500 wireless controllers. As mentioned previously, all data and control traffic is backhauled from the access points to wireless controllers before being terminated and placed onto the Ethernet network. Guest wireless traffic is backhauled across the campus infrastructure to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus.

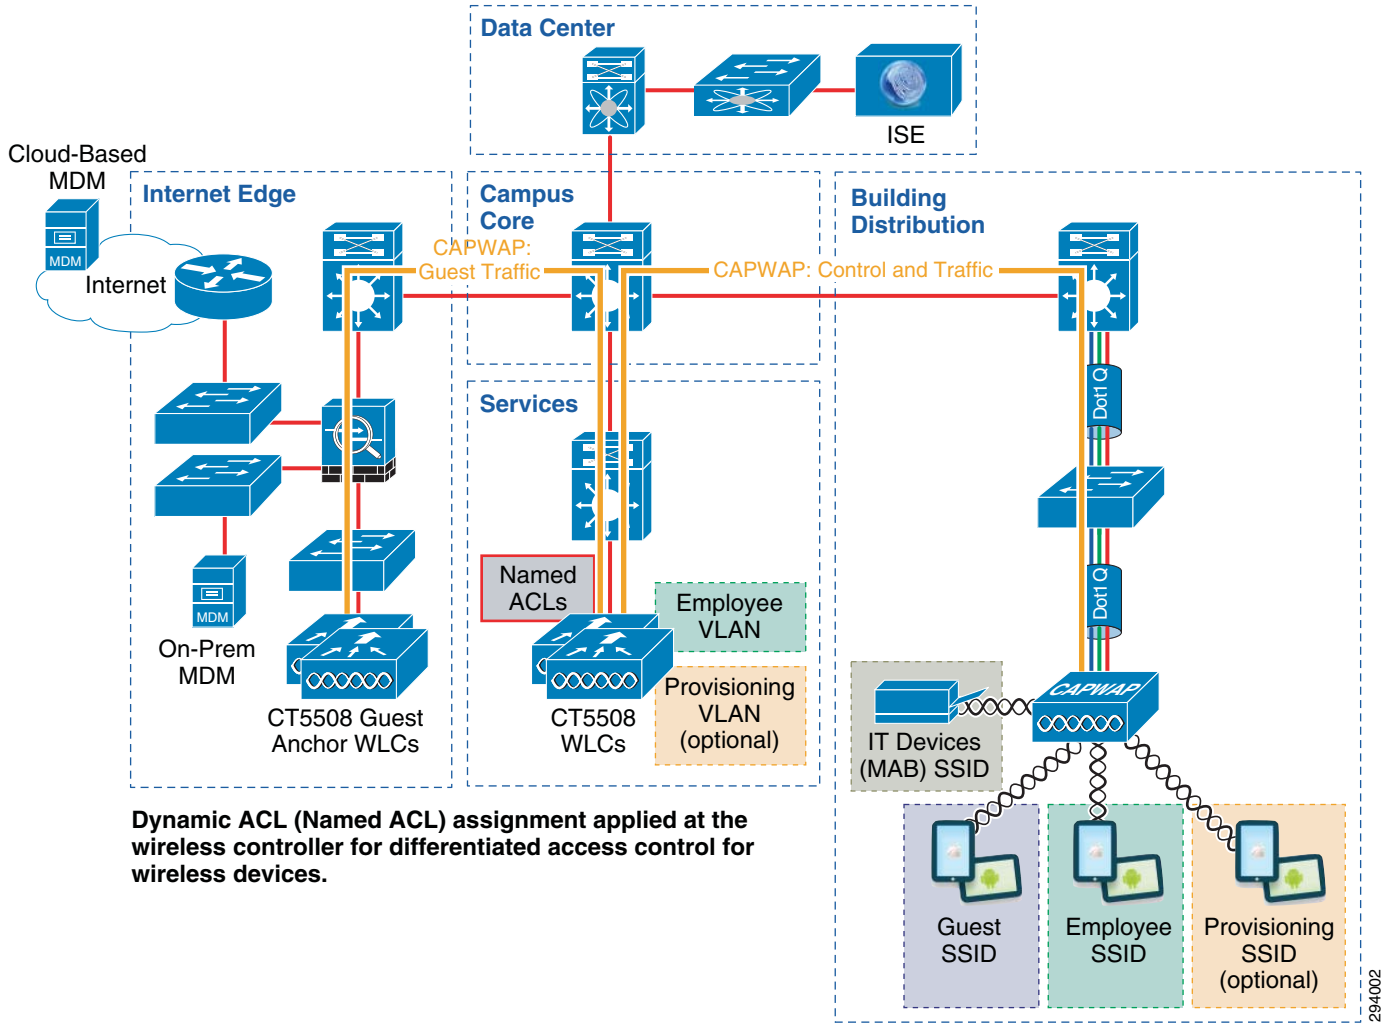
In order to implement the BYOD use cases, two separate methods of providing differentiated access control for campuses utilizing a Local Mode wireless design are examined. These methods are:

- Applying the appropriate dynamic ACL after the device is authenticated and authorized.
- Applying the appropriate Security Group Tag (SGT) to the device after it is authenticated and authorized.

When implementing access control via dynamic ACLs, the particular form of dynamic ACL chosen for the design guide are RADIUS specified local ACLs, otherwise known as named ACLs. These named ACLs must be configured on each CT5508 wireless controller. For example, a personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different named ACLs are applied to each device, granting different access to the network.

Figure 5-1 shows at a high level how a centralized (Local Mode) wireless BYOD design using named ACLs for access control is implemented in the campus.

Figure 5-1 High-Level View of the Centralized (Local Mode) Wireless Campus BYOD Design



Dynamic ACL (Named ACL) assignment applied at the wireless controller for differentiated access control for wireless devices.

When implementing access control via Security Group Association (SGA), various source and destination Security Group Tags (SGTs) must be configured within Cisco ISE. A personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different SGTs are applied to each device, thereby granting different access to the network.

Security Group Tag Overview

Throughout all versions of the BYOD CVD, policy enforcement has been accomplished through the use of Access Control Lists and VLANs to restrict user traffic as appropriate upon successful authentication and subsequent authorization. The use of ACLs can become a daunting administrative burden when factoring the number of devices upon which they are applied and the continual maintenance required to securely control network access.

BYOD v2.5 uses a complimentary technology known as TrustSec and the use of Security Group Tags (SGT). Security Group Tags offer a streamlined and alternative approach to enforcing role-based policies with minimal and in some cases, little or no ACLs at all if TCP/UDP port level granularity is not required.

The use of Security Group Tags are used as an alternative to ACLs for Campus wireless users and devices where the Cisco Wireless Controllers have been centrally deployed in a shared services block and configured for operation in local mode.

ACL Complexity and Considerations

To date, variations of named ACLs on wireless controllers, static and downloadable ACLs on various routing and switching platforms, as well as FlexACLs for FlexConnect wireless traffic in the branch have been used as a means of enforcing traffic restrictions and policies. In order to configure and deploy these ACLs, a combination of either command line (CLI) access to each device via Telnet/SSH or network management such as Prime Infrastructure have been required and used for statically configured ACLS while the Cisco Identity Services Engine (ISE) has been used to centrally define and push downloadable ACLs (DACL) to switching platforms.

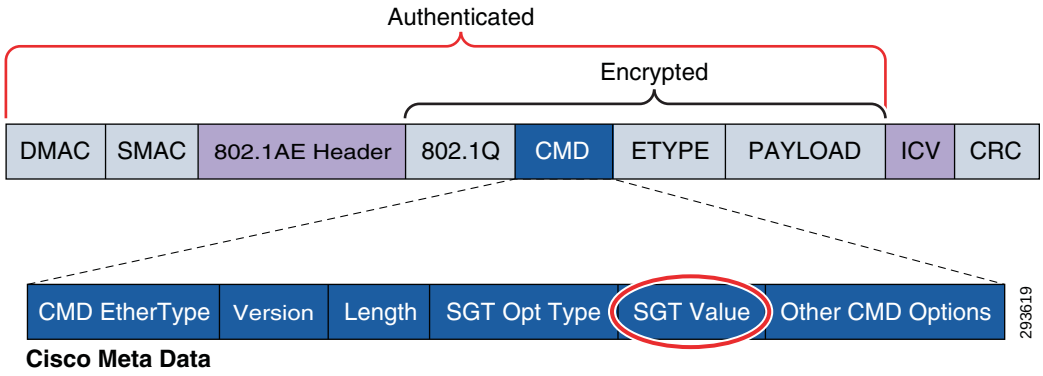
- Unique ACLs may be required for different locations such as branches or regional facilities, where user permissions may need to be enforced for local resources such as printers, servers, etc.
- The operational complexity of ACLs may be impacted by changes in business policies.
- The risk of security breaches increases with potential device misconfigurations.
- ACL definitions become more complex when policy enforcement is based on IP addresses.
- Platform capabilities, such as processor memory, scalability, or TCAM resources may be impacted by complex ACLs.

Cisco's TrustSec provides a scalable and centralized model for policy enforcement by implementing Cisco's Security Group Access architecture and the use of Security Group Tags.

Security Group Tag

Security Group Tags, or SGT as they are known, allow for the abstraction of a host's IP Address through the arbitrary assignment to a Closed User Group represented by an arbitrarily defined SGT. These tags are centrally created, managed, and administered by the ISE. The Security Group Tag is a 16-bit value that is transmitted in the Cisco Meta Data field of a Layer 2 Frame as depicted in Figure 5-2.

Figure 5-2 Layer 2 SGT Frame Format



The Security Group Tags are defined by an administrator at Cisco ISE and are represented by an arbitrary name and a decimal value between 1 and 65,535 where 0 is reserved for "Unknown". Security Group Tags allow an organization to create policies based on a user's or device's role in the network providing a layer of abstraction in security policies based on a Security Group Tag as opposed to IP Addresses in ACLs.

For a complete overview of the Security Group Access architecture and Security Group Tags and how it will be incorporated within the CVD, refer to Chapter 23, "BYOD Policy Enforcement Using Security Group Access."

SGT Deployment Scenarios in this CVD

Specifically, SGT will be used as a means of policy enforcement for the Enhanced Access Use Case where a Campus Wireless user/device terminated centrally at a Wireless Controller in Local Mode is granted either full or partial access to the network. Different classes of servers will be defined to which those users may or may not have access to. The CVD also defines a class that has access to the Internet only through the use of an ACL on the wireless controller to deny access to all internal addresses. The Converged Access products such as the Catalyst 3850 and CT-5760 will not be addressed relative to SGT in this CVD as Security Group Tags and SXP are currently not supported. More about SGT and the Enhanced Use Case will be discussed in the ensuing sections discussing the actual authorization policies.

Two deployment scenarios will be depicted within this CVD. The first will make use of Security Group ACLs (SGACLs) to enforce policies at the Nexus 7000 Data Center switches, whereas the second scenario will enforce policies configured at a Cisco ASA configured as a Security Group Firewall (SGFW). SGACLs are role-based policies enforced on Catalyst switching platforms and specifically define whether traffic is permitted or denied based on source and destination SGT values. Again, these deployment scenarios are not mutually exclusive and can be used together. This first scenario can be seen in Figure 5-3 and the second scenario in Figure 5-4.

Figure 5-3 Policy Enforcement Using SGACL

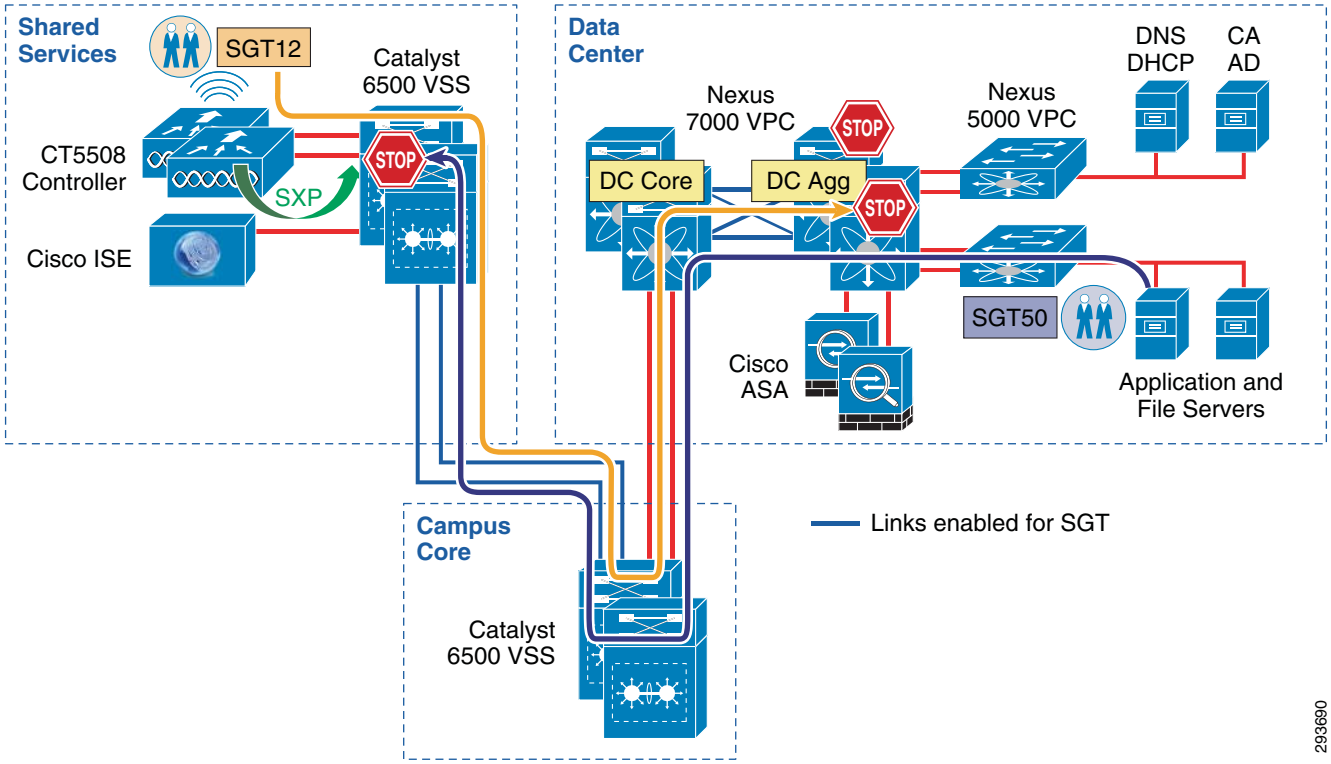
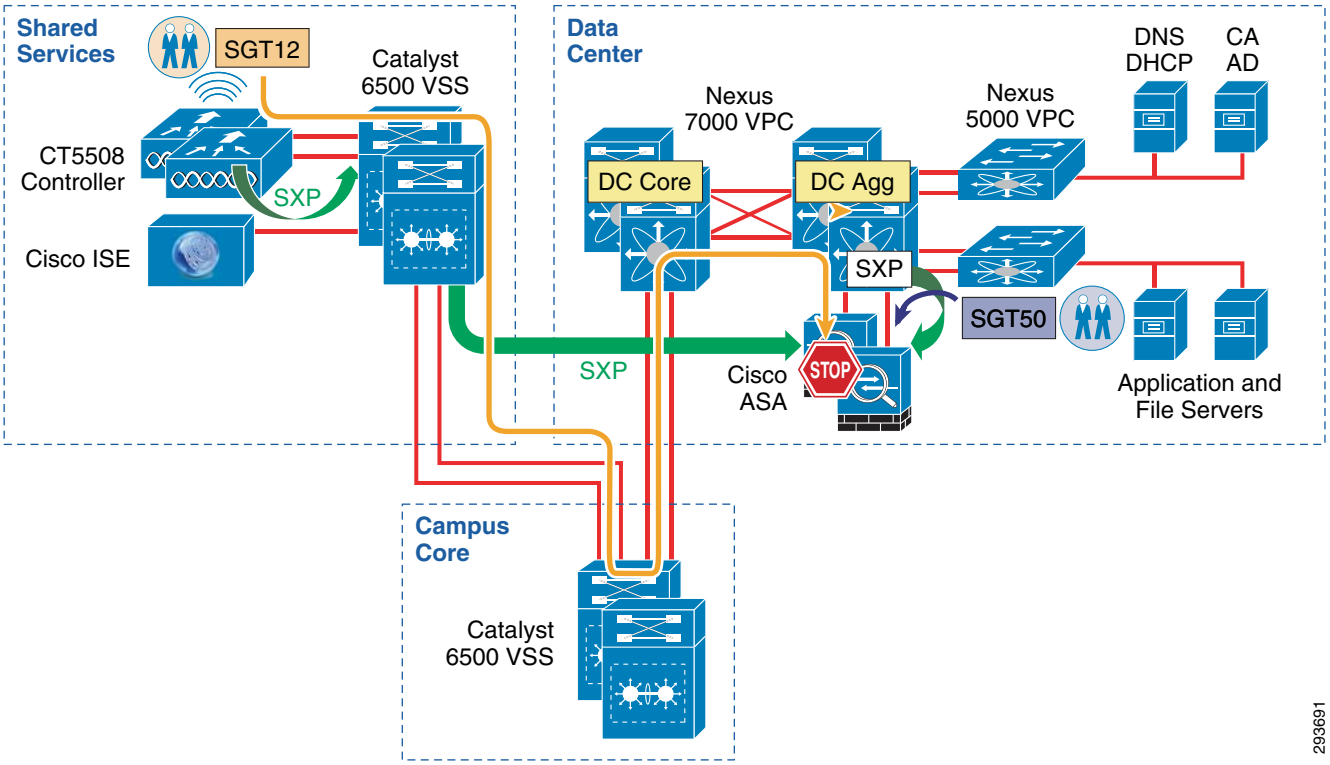


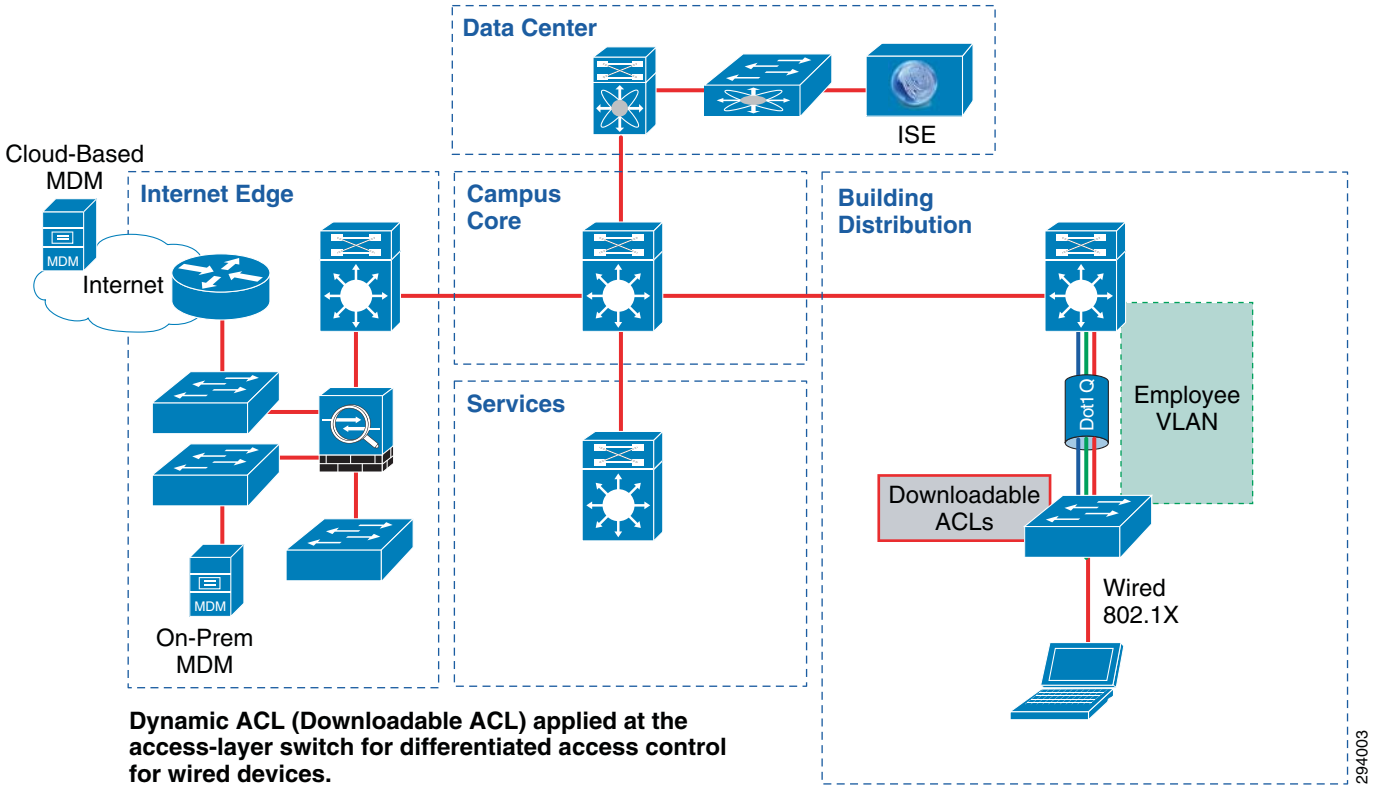
Figure 5-4 Policy Enforcement Using SG-FW



Campus Wired Design

Figure 5-5 shows the wired design for a campus which does not implement Converged Access Catalyst 3850 Series switches. In other words, this is the wired design for a campus which implements switches such as the Catalyst 3750X and 4500 series at the access-layer of building distribution modules, along with a centralized (Local Mode) wireless design.

Figure 5-5 High-Level View of Non-Converged Access Wired Campus Design



Dynamic ACL (Downloadable ACL) applied at the access-layer switch for differentiated access control for wired devices.

This design guide assumes Catalyst switches deployed as Layer 2 devices within the access-layer of the campus building modules. Wired devices authenticate using 802.1X against the ISE server located within the campus data center. For this design, wired devices are all statically assigned to a single VLAN, the Employee VLAN. Differentiated access control for wired devices is provided by different RADIUS downloadable ACLs applied to the access-layer switch, which override a pre-configured static ACL on each Catalyst switch port.

Converged Access Campus Design

The Converged Access campus BYOD design highlights multiple Catalyst 3850 Series switches or switch stacks deployed at the access layer of each building distribution module of a large sized campus. Switch stacks form Switch Peer Groups (SPGs) in which all switches contain the Mobility Agent (MA) function. Roaming within a SPG is handled through a full mesh of mobility tunnels between MAs within the SPG. Multiple SPGs exist within the large sized campus.

This design guide will assume Catalyst 3850 Series switches deployed as Layer 2 access switches within the campus location. Layer 3 connectivity within each campus building distribution module is provided by Catalyst 6500 distribution switches. In keeping with campus design best practices for minimizing spanning-tree issues, VLANs are assumed not to span multiple Catalyst 3850 Series switch stacks deployed in separate wiring closets.Future design guidance may address Catalyst 3850 Series switches deployed as Layer 3 switches within the branch location.

Cisco CT5760 wireless controllers deployed within a centralized service module within the campus contains the Mobility Controller (MC) function. Multiple SPGs connecting to a single MC form a Mobility Sub-Domain. Multiple Mobility Sub-Domains exist within the large sized campus. Roaming between SPGs within a Mobility Sub-Domain is done through the Cisco CT5760 wireless controller. The CT5760 wireless controllers also manage Radio Resource Management (RRM), WIPs, etc.

Multiple Cisco CT5760 wireless controllers form a Mobility Group. Hence a Mobility Group also consists of multiple Mobility Sub-Domains. Roaming between Mobility Sub-domains is done through the Cisco CT5760 wireless controllers within the Mobility Group. The design within this design guide assumes a single Mobility Group and hence a single Mobility Domain extends across and is entirely contained within the large campus.

Note Cisco CT5508 wireless controllers can also implement the Mobility Controller (MC) function within the Converged Access campus design. However the CT5508, being an older platform has less overall throughput than the newer CT5760 platform. This version of the design guide only discusses the CT5760 wireless controller functioning as the Mobility Controller within a Converged Access campus deployment. Future versions of this design guide may include the CT5508 wireless controller deployed in this manner.

Access points within the campus building distribution modules are configured and controlled via the wireless controller Mobility Agent (MA) functionality integrated within the Catalyst 3850 Series switch. Guest wireless traffic is still backhauled to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus. Provisioning traffic (i.e., traffic from devices attempting to on-board with ISE) is terminated locally on the Catalyst 3850 Series switch with the Converged Access campus design. When implementing a dual-SSID design, provisioning traffic is terminated on a separate VLAN. All on-boarded devices terminate on a single VLAN with this design.

Note This design guide only discusses wireless guest access. Wired guest access may be discussed within future revisions of this design guide.

The potential advantages of this design are as follows:

- Increased scalability of the wireless deployment, since wireless traffic is terminated on every access-layer Catalyst 3850 Series switch within the campus, instead of being backhauled to one or more centralized wireless controllers.
- Increased visibility of the wireless traffic, since wireless traffic is terminated on every access-layer Catalyst 3850 Series switch within the campus.

The potential disadvantages of this design are as follows:

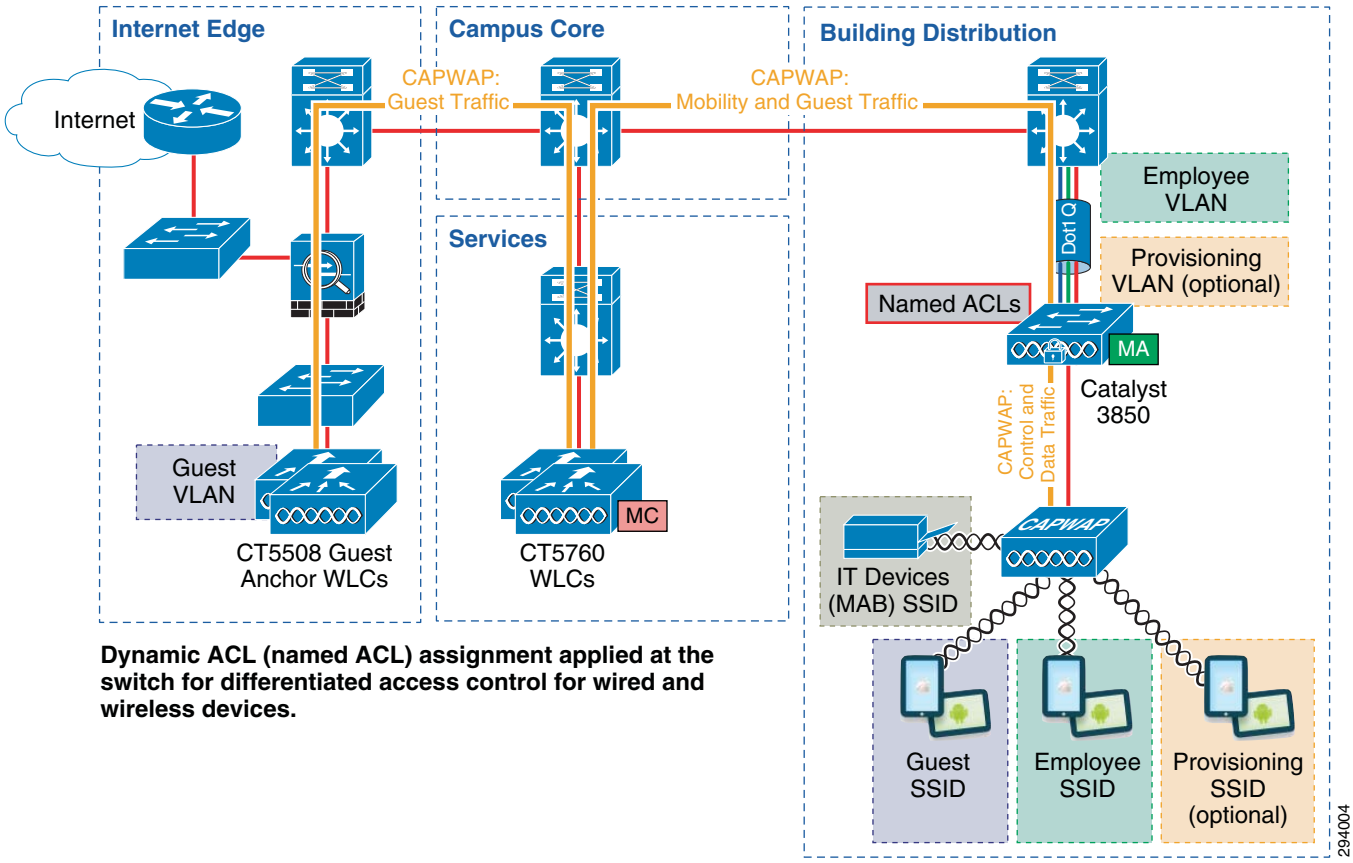
- Less centralized access control of wireless traffic from a single point within the campus network. Access control is spread out to each Catalyst 3850 Series access switch. Note however, that with Converged Access designs, traffic from a particular WLAN can still be backhauled to a centralized CT5760 wireless controller and switched centrally. This is touched upon in Campus Migration Path.
- Increased potential for more complexity for wireless roaming, since each Catalyst 3850 Series switch implements the Mobility Agent (MA) functionality, effectively functioning as a wireless controller.

In order to implement the BYOD use cases, the method adopted in this design guide for a campus utilizing a Converged Access design is to apply the appropriate named ACL after the device is authenticated and authorized. This applies to both wired and wireless devices. These named ACLs, which must be configured on each Catalyst 3850 Series switch, provide differentiated access control.

For example, a personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different named ACLs are applied to each device, granting different access to the network.

Figure 5-6 shows at a high level a simplified Converged Access BYOD design with a single Catalyst 3850 Series switch functioning as a Mobility Agent (MA) and a single CT5760 wireless controller functioning as a Mobility Controller (MC) in the campus.

Figure 5-6 High-Level View of the Converged Access Campus BYOD Design



Dynamic ACL (named ACL) assignment applied at the switch for differentiated access control for wired and wireless devices.

Note The Converged Access campus BYOD design may also be referred to as the External Controller Large Campus BYOD design within this document. Future versions of this design guide may address small campus and/or large branch Converged Access designs, in which multiple Catalyst 3850 switch stacks implement both the Mobility Controller (MC) and Mobility Agent (MA) functionality. In such a design, referred to as the Integrated Controller Small Campus / Large Branch design, no external CT5760 wireless controllers are needed.

Note that in the case of this design guide, on-boarded wired devices are also statically assigned to the same VLAN as wireless devices. Hence on-boarded wired and wireless devices will share the same VLAN, and hence the same IP subnet addressing space. It is recognized that customers may implement separate subnets for wired and wireless devices due to issues such as additional security compliance requirements for wireless devices. This is not addressed within this version of the design guidance. Dynamically assigned named ACLs provide differentiated network access for wired devices.

Assuming all campus switches implement the same set of ACLs for access control, RADIUS downloadable ACLs may alternatively be deployed within the campus. The benefit of implementing a downloadable ACL within the campus is that changes to the access control entries only have to be configured once within the Cisco ISE server versus having to touch all campus Catalyst 3850 Series switches. However this option also requires separate ISE policy rules for campus and branch Converged Access deployments, assuming named ACLs are still deployed within branch locations.

Implementing downloadable ACLs within branch locations presents scaling issues if access to local branch servers is required within the ACL. In such scenarios, each branch would require a separate downloadable ACL and, therefore, a separate Cisco ISE policy rule to identify that ACL for that branch. This becomes administratively un-scalable as the number of deployed branches increases.

Hence this design guide only discusses the use of named ACLs for access control of on-boarded devices both within the Converged Access branch and campus designs. Because named ACLs are used for both designs, the same Cisco ISE policies rules can be used for both Converged Access campus and branch deployments. Hence one set of policy rules can be used for Converged Access designs regardless of where the device is located. This reduces the administrative complexity of the Cisco ISE policy; albeit it at the expense of increased complexity of having to configure and maintain ACLs at each campus Catalyst 3850 Series switch.

Note Management applications such as Cisco Prime Infrastructure may ease the burden of ACL administration by providing a point of central configuration and deployment of named ACLs for the Converged Access BYOD branch and campus designs.

Campus Migration Path

For large campus designs, a migration path from a traditional CUWN centralized (Local Mode) wireless overlay network design to a Converged Access design is necessary. It is considered unfeasible for a customer to simply “flash cut” a large campus over to a Converged Access design. There are many potential migration paths from a traditional CUWN centralized design to a Converged Access design. This section discusses one possible migration path. The steps of the migration path from the initial overlay model are as follows:

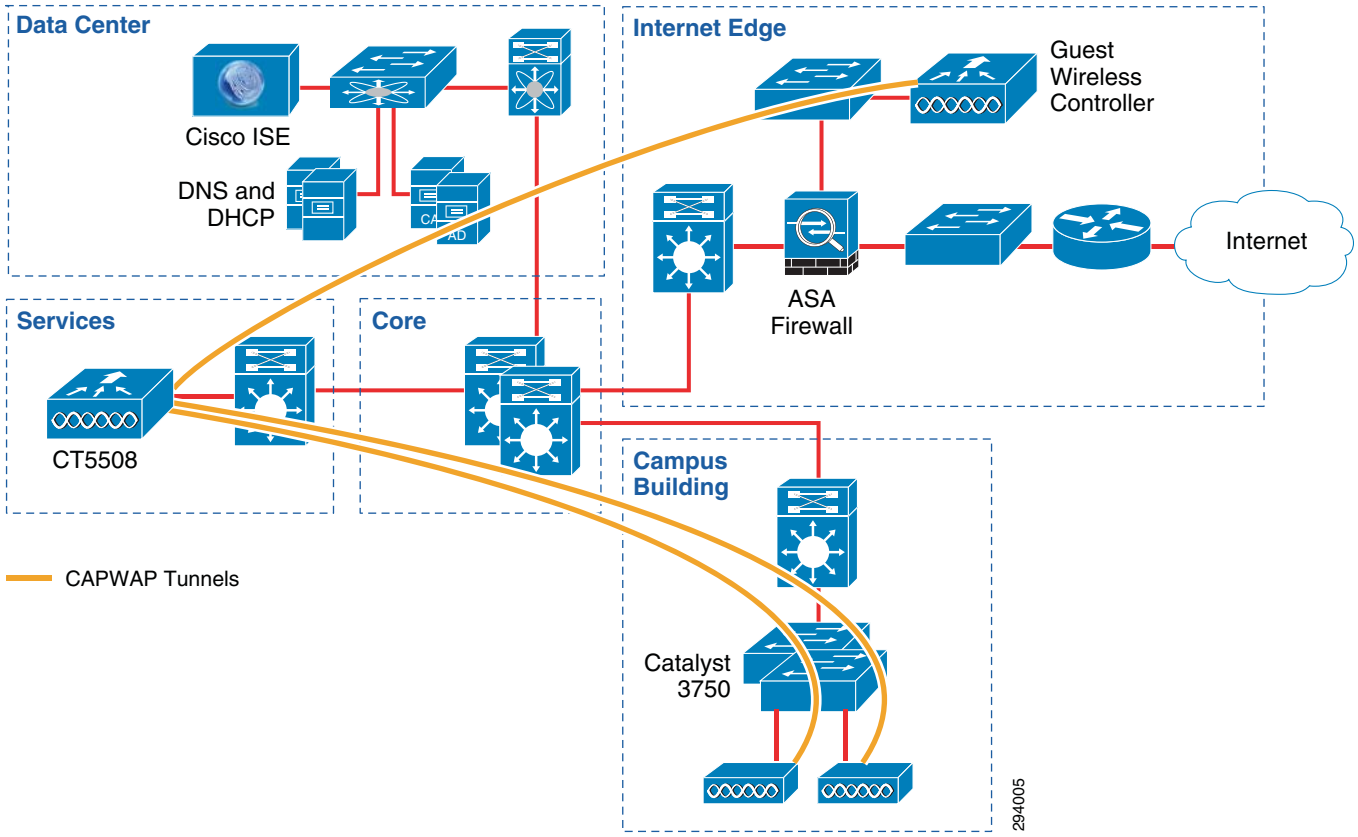
1. Local/Centralized Mode Only
2. Hybrid Converged Access and Centralized
3. Full Converged Access

Each is discussed in the following sections.

Initial Overlay Model

Figure 5-7 shows the logical components for the initial state in the migration path - the Initial Overlay Model.

Figure 5-7 Initial State in the Migration Path—Initial Overlay Model



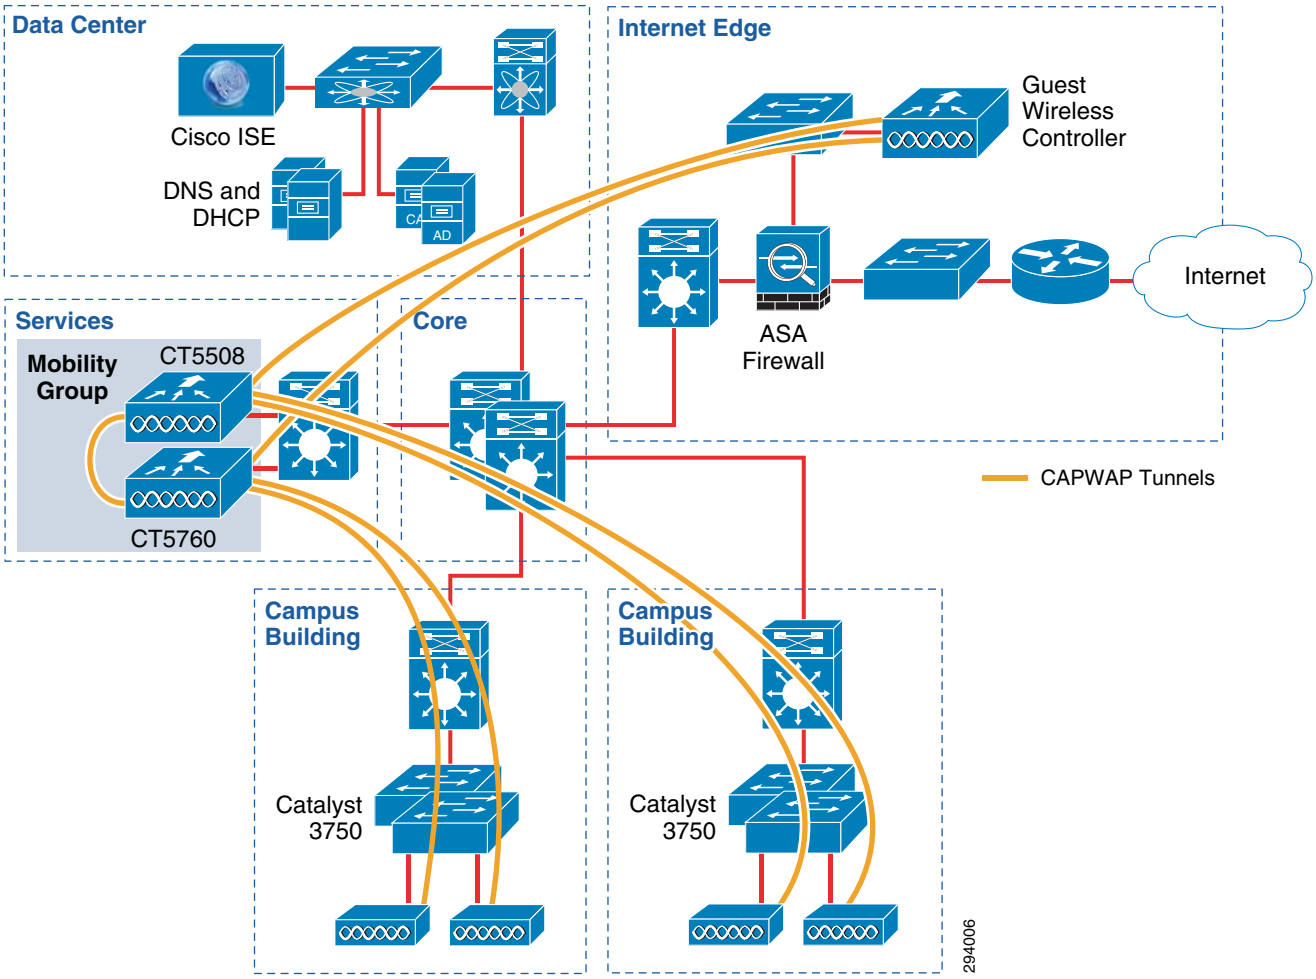
The initial overlay model consists of access points, operating in Local Mode, connected to Catalyst 3750-X series switches at the access-layer of individual building modules within the campus. The access points are controlled by a CT5508 wireless controller located within a services module within the campus. CAPWAP tunnels extend from individual access points to the CT5508 wireless controller. A second CT5508 wireless controller on a DMZ segment within the Internet edge module functions as a dedicated wireless guest anchor controller. A mobility tunnel (Ethernet-over-IP or CAPWAP, depending on CUWN software version) extends from the campus (foreign) CT5508 wireless controller to the guest (anchor) CT5508 wireless controller.

This is the campus BYOD design which is discussed in Centralized (Local Mode) Wireless Design.

Centralized/Local Mode Only

Figure 5-8 shows the logical components for the first step in the migration path—Centralized/Local Mode Only.

Figure 5-8 First Step in Migration Path—Centralized/Local Mode Only



Note Note that the term "Local Mode" is used with CUWN controllers, while the term "Centralized Mode" is used with Converged Access controllers within Cisco documentation. Both refer to the same model with a centralized data and control plane for wireless traffic. In other words, all traffic is backhauled to the wireless controller before being placed on the Ethernet network.

In this step of the migration path, the customer simply adds more wireless controller capacity. Since the CT5760 is a newer platform and offers higher aggregate throughput, the customer may decide to begin transitioning to this platform by adding them to the existing campus wireless overlay design. The CT5760 supports up to 1,000 access points and up to 12,000 clients with up to 60 Gbps throughput per wireless controller.

Note The wireless capabilities of the CT5760 are roughly equivalent to Cisco Unified Wireless Network software version 7.0 with some features from later software versions. The network administrator must ensure that all the necessary features exist in the CT5670 before migrating access points from existing CT5508 wireless controllers to CT5760 wireless controllers. For a list of supported features, refer to the

CT5760 Controller Deployment Guide at: http://www.cisco.com/en/US/docs/wireless/technology/5760_deploy/CT5760_Controller_Deployment_Guide.html.

At this point, it is assumed that the access-layer switches within the building module wiring closets have not reached their replacement cycle. Hence the access points, operating in local mode, are still connected to Catalyst 3750-X series switches at the access-layer of individual building modules within the campus. The access points are controlled by either the CT5508 or the CT5760 wireless controller located within a services module within the campus. Both are members of the same Mobility Group. CAPWAP tunnels extend from individual access points to either the CT5508 or CT5760 wireless controller. A mobility tunnel extends between the CT5508 and CT5760.

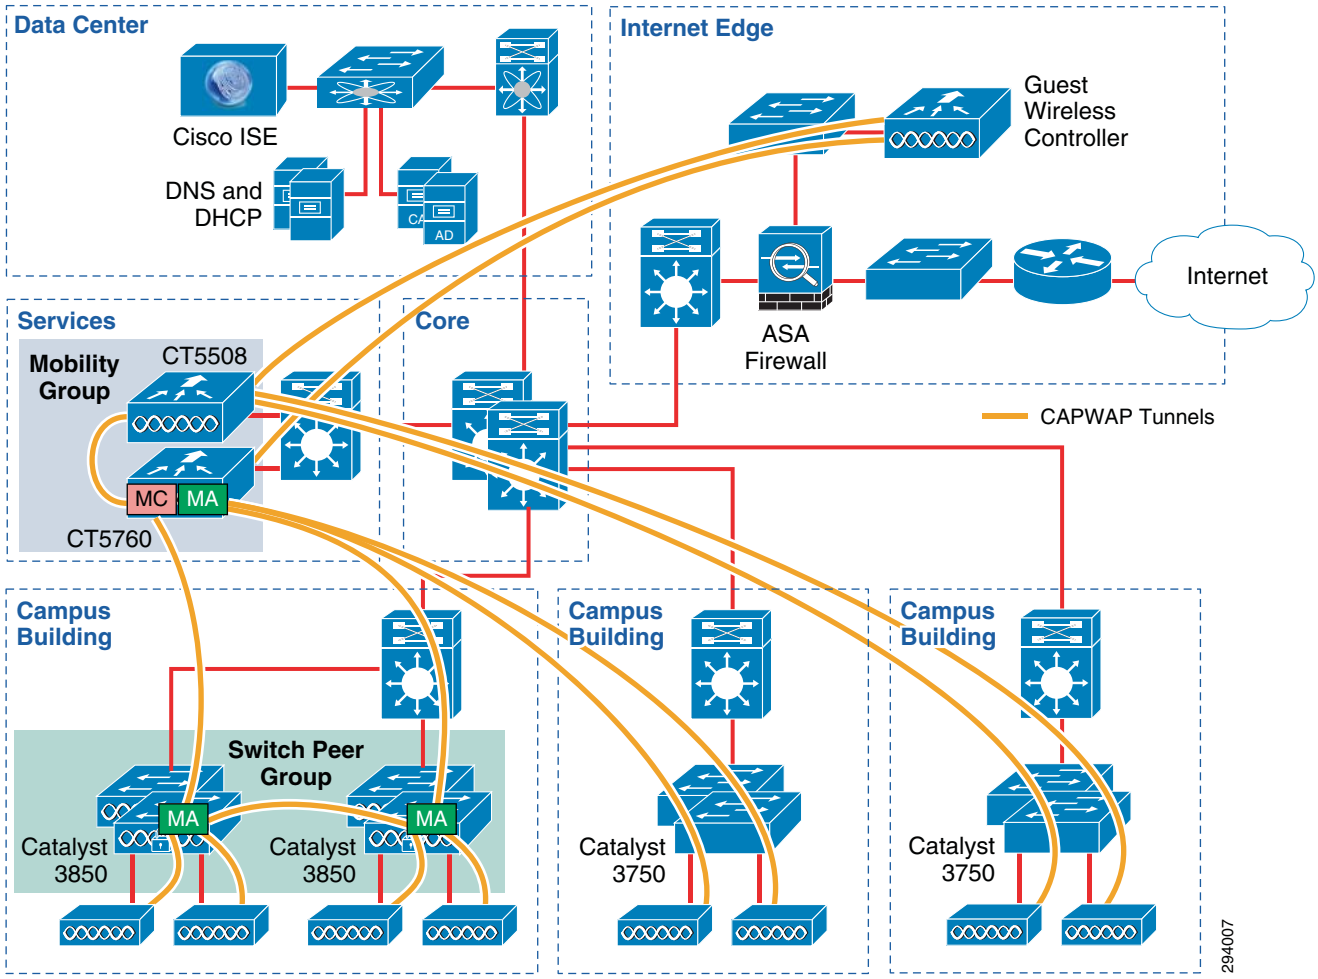
A logical choice for migration to the CT5760 wireless controller would initially be at the building level. In other words, one building of a campus could be migrated—potentially floor by floor—from an existing CT5508 to a CT5760 wireless controller.

In order to maintain mobility across the campus, the existing CT5508 wireless controllers need to be upgraded to CUWN software version 7.5. CUWN software version 7.5 supports the new mobility tunneling method, which uses CAPWAP instead of Ethernet-over-IP, which is compatible with IOS XE 3.2.2 software running on CT5760 wireless controllers. Note that this includes upgrading the CT5508 wireless controller dedicated for wireless guest access. Mobility tunnels (in this case CAPWAP tunnels) extend from the foreign CT5508 and CT5760 wireless controllers to the anchor CT5508 wireless controller.

Hybrid Converged Access and Local Mode

Figure 5-9 shows the logical components for the second step in the migration path—a Hybrid Converged Access and Local Mode model.

Figure 5-9 Second Step in Migration Path—Hybrid Converged Access and Local Mode



At this point in the migration path, it is assumed that the access-layer switches within the building module wiring closets have begun to reach their replacement cycle. In this scenario, the customer has chosen to deploy Catalyst 3850 Series switches at the access-layer of their building modules and begin migrating to a converged access model. Again, a logical choice for migration would be at the building level. In other words, one building of a campus would be migrated—potentially floor by floor—from access points operating in centralized mode connected to a Catalyst 3750-X Series switch and controlled by the CT5760, to access points operating in converged mode connected to and controlled by a Catalyst 3850 Series switch.

With this design, the Catalyst 3850 Series switches function as the Mobility Agent (MA), while the CT5760 wireless controller functions as the Mobility Controller (MC) and possibly the Mobility Oracle (MO). However during the migration of floors, the CT5760 wireless controller will still have to function in centralized mode as well for access points still connected to Catalyst 3750-X series switches. Hence the design is a “hybrid” of centralized and converged access designs.

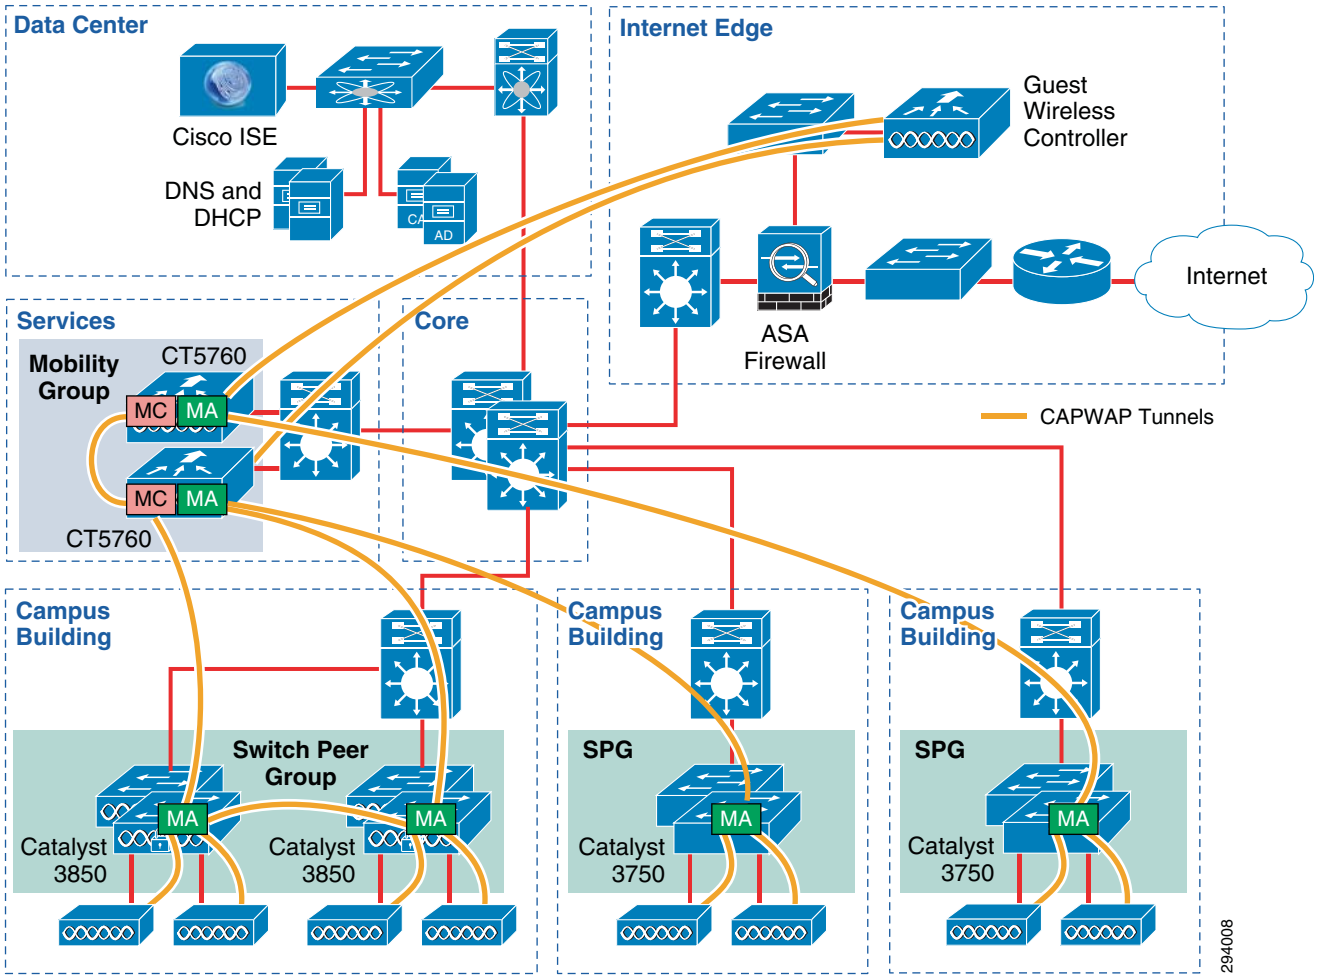
CAPWAP tunnels extend from individual access points which are connected to Catalyst 3750-X Series switches to either the CT5508 or CT5760 wireless controller. CAPWAP tunnels also extend from individual access points which are connected to Catalyst 3850 Series switches to the Catalyst 3850 Series switches. CAPWAP mobility tunnels extend from the MA within the Catalyst 3850 Series switches to the MC within the CT5760 wireless controller. Finally, CAPWAP mobility tunnels extend between MAs within the Catalyst 3850 Switches which are part of a Switch Peer Group (SPG). SPGs

offload mobility traffic for groups of switches in which a large amount of mobility is expected. When roaming between access points connected to Catalyst 3850 Series switches which are part of the same SPG, the MC located within the CT5760 is not involved in the roam. A SPG may extend across part of a floor within a building, the entire floor, or in some cases multiple floors. A CAPWAP mobility tunnel also extends between the CT5508 and CT5760. Finally CAPWAP mobility tunnels extend from the foreign CT5508 and CT5760 back to the anchor CT5508 for wireless guest access.

Full Converged Access

Figure 5-10 shows the logical components for the third step in the migration path—the Full Converged Access model.

Figure 5-10 Third Step in Migration Path—Full Converge Access



This design assumes the customer has retired existing CT5508 wireless controllers operating in Local Mode and moved to a converged access design with CT5670 wireless controllers. At this point in the migration path, it is assumed that the access-layer switches within the building module wiring closets have completed their replacement cycle. In this scenario, the customer has chosen to deploy only Catalyst 3850 Series switches at the access-layer of their building modules and completely migrate to a converged access model.

Note We realize that some customers may never fully migrate to a full Converged Access model, while others may take years to reach a full Converged Access deployment.

With this design, the Catalyst 3850 Series switches function as the Mobility Agent (MA), while the CT5760 wireless controller functions as the Mobility Controller (MC) and possibly the Mobility Oracle (MO).

CAPWAP tunnels extend from individual access points which are connected to Catalyst 3850 Series switches to the Catalyst 3850 Series switches. CAPWAP mobility tunnels extend from the MA within the Catalyst 3850 Series switches to the MC within the CT5760 wireless controller. CAPWAP mobility tunnels extend between MAs within the Catalyst 3850 Switches which are part of a Switch Peer Group (SPG). A CAPWAP mobility tunnel also extends between the two CT5760 wireless controllers. Finally CAPWAP mobility tunnels extend from the foreign CT5760 wireless controllers back to the anchor CT5508 for wireless guest access.

Note Roaming between sub-domains (i.e., roaming between two CT5760 wireless controllers functioning as MCs) has not been validated with this version of the design guide.

Wireless LAN Controller High Availability

High availability of the wireless network is becoming increasingly important as more devices with critical functions move to the wireless medium. Real-time audio, video, and text communication relies on the corporate wireless network and the expectation of zero downtime is becoming the norm. The negative impacts of wireless network outages are just as impactful as outages of the wired network.

With Cisco Unified Wireless Network (CUWN) Software Release 7.3 and above, the ability to have an active and hot-standby wireless controller has been introduced, allowing the access points (APs) to perform a rapid stateful switchover (SSO). This capability allows all the AP sessions to statefully switch over to the hot-standby WLC with an identical configuration to the primary WLC. All unique configuration parameters and groupings specific to individual APs and AP groups are retained. An example of retained configuration is Flex-Connect grouping, which applies different restrictions and settings to sub-sets of APs based on branch location. Clients will be disassociated when a failover occurs. However, clients should automatically re-associate after the stateful switchover of the access point occurs.

The active and standby WLCs use a dedicated redundant interface to send keep-alives every 100 milliseconds, as well as sending configuration, operational data synchronization, and role negotiation information between them. The redundancy interface is a dedicated port that is directly connected between WLCs by an Ethernet cable. For the WiSM2, a dedicated redundancy VLAN is used in place of the redundancy port. Failovers are triggered by loss of keep-alives as well as network faults. The active and standby WLCs share the same management IP address, with only the active being up until a failure occurs.

For more information, refer to the WLC High Availability Deployment Guide: http://www.cisco.com/en/US/products/ps10315/products_tech_note09186a0080bd3504.shtml.

Note The Cisco CT5760 wireless controller is an IOS XE based controller. IOS XE 3.2.2 does not support access point stateful switchover (AP SSO). Instead, access points controlled by the CT5760 support the configuration of a primary, secondary, and tertiary wireless controller for high availability. A Catalyst 3850 Series switch stack with wireless controller functionality enabled supports high availability

through the switch stack itself. High availability of both of these platforms is not covered within this design guide. Future versions of this design guide will address high availability on the CT5760 and Catalyst 3850.

Branch Wide Area Network Design

Many network administrators will re-examine the wide area network (WAN) prior to deploying a BYOD solution at the branch. Guest networks in particular have the ability to increase loads to a rate that can consume WAN bandwidth and compromise corporate traffic. While wired rates have increased from 10 Mbps to 1 Gbps and cellular networks have increase bandwidth from 30 Kbps for GPRS to around 20 Mbps for LTE, traditional branch WAN bandwidths have not experienced the same increase in performance. Employees and guests expect bandwidth, delay, and jitter on the corporate network to be at least as good as they experience at home or on the cellular network.

Furthermore, because WiFi access is typically free for corporate users and because most hand held devices will prefer WiFi over cellular, corporate users will likely continue using the guest or corporate SSID for Internet access, even when the LTE network offers faster speeds. This is forcing network administrators to explore new WAN transport mechanisms such as Metro Ethernet and VPN-over-Cable to meet user expectations. Another approach is to offload guest Internet traffic at the branch in an effort to preserve WAN bandwidth for corporate traffic. Corporate Security Policy will need to be considered, however, before providing direct Internet access from the branch. As a result, the WAN is experiencing increased loads. While there are no new WAN requirements for branch BYOD services, some areas such as transport technology, access speeds, and encryption should be reviewed.

Branch WAN Infrastructure

The branch WAN infrastructure within this design includes Cisco ASR 1006s as the head-end routers. Two different WAN connections are terminated on these devices; the first router is configured as a service provider MPLS circuit and the second router is configured with an Internet connection. These head-end routers are both placed in a “WAN edge” block that exists off of the campus core. The ASR that terminates the Internet connection also makes use of IOS Zone-Based Firewall (ZBFW) and only tunneled traffic towards the branch is permitted.

Within the branch, two different designs have been validated. The first design consists of two Cisco 2921 ISR-G2 routers. One of the two routers terminates the SP MPLS circuit, while the second router terminates the Internet connection which can be utilized as a branch backup exclusively or as an alternate path for corporate traffic. The second design consists of a single Cisco 2921 ISR-G2 router that terminates both circuits.

In both deployment modes, the Cisco IOS Zone-Based Firewall (ZBFW) has been implemented to protect the branch's connection to the Internet. Although entirely feasible, local Internet access from the branch is not permitted. For this purpose as well as for corporate data, DMVPN has been implemented and only tunnel access granted for secure connectivity back to the campus head-end routers. This provides for access to the data center. Internet access is available through the corporate firewall/gateway. DMVPN is additionally used to secure traffic across the service provider's MPLS circuit.

It is beyond the scope of this document to provide configuration information and design guidance around DMVPN, ZBFW configuration, QOS, and other aspects of the WAN infrastructure.

For detailed reference information around Next Generation Enterprise WAN (NGEW) design, refer to the documentation on Design Zone:
http://www.cisco.com/en/US/netsol/ns816/networking_solutions_program_home.html.

For additional QOS Design Guidance, refer to the *Medianet Design Guide* at: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns819/landing_vid_medianet.html.

Branch WAN Bandwidth Requirements

This design guide presents two branch wireless LAN designs—FlexConnect and Converged Access. In FlexConnect designs, branch access points are managed by a wireless LAN controller in the campus data center or services module. A CAPWAP tunnel is established between the wireless controller and the access points within the branch locations. This CAPWAP tunnel is used for control traffic and possibly data traffic during the on-boarding process in some designs. This traffic is transported over the WAN. Even though devices may use a FlexConnect design to locally terminate traffic onto local VLANs within the branch, a large percentage of traffic will continue to flow over the WAN to the corporate data center.

In Converged Access designs, branch access points are managed by the integrated wireless LAN controller functionality within the Catalyst 3850 Series switch. A CAPWAP tunnel is established between the Catalyst 3850 Series switch and the access points within the branch locations. This CAPWAP tunnel is used for all wireless control and data traffic. However, even though devices may use a Converged Access design to locally terminate traffic onto local VLANs within the branch, a large percentage of traffic will again continue to flow over the WAN to the corporate data center.

Since both branch wireless LAN designs presented in this document utilizes a centralized AAA server (such as Cisco ISE), there may be an increase in authentication and authorization traffic as more employee managed devices are on-boarded. These new endpoints may also generate additional new traffic. Further, guest Internet access is carried back to an anchor controller in the campus DMZ with both branch wireless LAN designs. All of this may result in increased loads on the WAN circuit as a result of the BYOD deployment.

It may be difficult to forecast the additional amount of traffic loading because the level of participation may not be well known prior to deploying BYOD. Wireless guest traffic in particular can be difficult to budget and may vary substantially depending on local events. A reasonable design goal is to provision a minimum of 1.5 Mbps at each branch that offers BYOD. The head-end WAN aggregation circuits should be provisioned to follow traditional oversubscription ratios (OSR) for data. This will allow adequate bandwidth for smaller deployments. Larger branch locations will likely need additional bandwidth, especially if the guest users are likely to expect the use of high bandwidth applications such as streaming video traffic. The WAN architecture should offer enough flexibility to adjust service levels to meet demand. Sub-rate MPLS access circuits or a dedicated WAN router with incremental bandwidth capabilities can accomplish this. Address space adequate for each branch should also be considered because both FlexConnect and Converged Access designs can allow wireless DHCP clients to pull from local scopes. Additional information concerning bandwidth management techniques such as rate-limiting is discussed in Chapter 21, “BYOD Guest Wireless Access.”

Encryption Requirement

Another component of both BYOD enabled branch wireless LAN designs is local termination of branch wireless traffic. This allows branch wireless devices to directly access resources located on the branch LAN without the need to traverse a CAPWAP tunnel to a centralized wireless controller. This reduces the amount of traffic that needs to be carried by the WAN by eliminating the hair-pinning of traffic from the branch location, back to the wireless controller within the campus, and then back to the branch server. The effect reduces load in both directions-upstream within a CAPWAP tunnel and downstream outside of the CAPWAP tunnel. The benefits are realized when a wireless branch device is connecting to a server located in the same branch. If the traffic is destined for the data center, it still transits the WAN but outside of a CAPWAP tunnel, benefiting from the same level of security and performance as wired

traffic. Depending on the application, it may not be encrypted so additional WAN security might be needed. If the branch is using a broadband connection as either the primary or backup path, then obviously encryption technologies such as DMVPN should be deployed. However, even if an MPLS VPN service is being used, the enterprise may still want to consider encrypting any traffic that passes off premise.

Transport

With both the FlexConnect and Converged Access designs, not all wireless traffic is terminated locally. In this design guide guest traffic is still tunneled within a CAPWAP tunnel to a central controller at a campus location. Also, depending upon the on-boarding design implemented (single SSID versus dual SSID), traffic from devices which are in the process of being on-boarded may also remain in the CAPWAP tunnel to the central controller with the FlexConnect design. This traffic may compete for bandwidth with the corporate traffic also using the WAN link, but not inside a CAPWAP tunnel. These concerns are addressed with a mix of traditional QoS services and wireless rate-limiting. In some situations, the transport will determine what is appropriate.

If Layer 2 MPLS tunnels are in place, destination routing can be used to place CAPWAP traffic on a dedicated path to the wireless controllers. This may be useful as an approach to isolate guest traffic from the branch towards the campus since FlexConnect with local termination will pass most corporate traffic outside of a CAPWAP tunnel directly to its destination. Return traffic from the campus towards the branch is more difficult to manage without more complex route policies, but may be possible with careful planning.

Figure 3-2 illustrates at a high level a typical WAN architecture.

Branch LAN Network Design

The anywhere, any device requirement of BYOD implies that employees can use either corporate or personal devices at either campus or branch locations. When they do, the pertinent component of the BYOD architecture is the ability to enforce policies on these devices at either the branch or at the campus location. Policy enforcement is effective if and only if there is a well-designed branch network infrastructure in place. This branch network infrastructure can be categorized into WAN and LAN components. This section discusses the high level key design elements of branch LAN design.

Cisco access points can currently operate in one of two implementation modes in the Cisco Unified Wireless Network (CUWN) architecture:

- Local mode (also referred to as a centralized controller design)
- FlexConnect mode

In addition, Cisco has recently integrated wireless LAN controller functionality directly into the latest generation access-layer switches—the Catalyst 3850. Hence there is a now a third implementation choice:

- Converged Access

FlexConnect is a wireless design which primarily applies to branch locations and is discussed in this section. Local mode is a wireless design which primarily applies to campus locations within this design guide and is discussed in Campus Network Design. Converged Access designs apply to both wired and wireless designs within both the branch and campus and hence are discussed in both sections of this chapter.

Note Local mode can be deployed within branches which are large enough to justify the requirement for wireless controllers deployed within the branch itself. In such cases, the BYOD design for the large branch is similar to the campus design.

FlexConnect Wireless Design

FlexConnect is an innovative Cisco technology which provides more flexibility in deploying a wireless LAN. For example, the wireless LAN may be configured to authenticate users using a centralized AAA server, but once the user is authenticated the traffic is switched locally on the access point Ethernet interface. Alternatively, the traffic may be backhauled and terminated on the wireless controller Ethernet interface if desired. The local switching functionality provided by FlexConnect eliminates the need for data traffic to go all the way back to the wireless controller when access to local resources at the branch is a requirement. This may reduce the Round Trip Time (RTT) delay for access to applications on local branch servers, increasing application performance. It can also reduce unnecessary hair-pinning of traffic when accessing resources local to the branch.

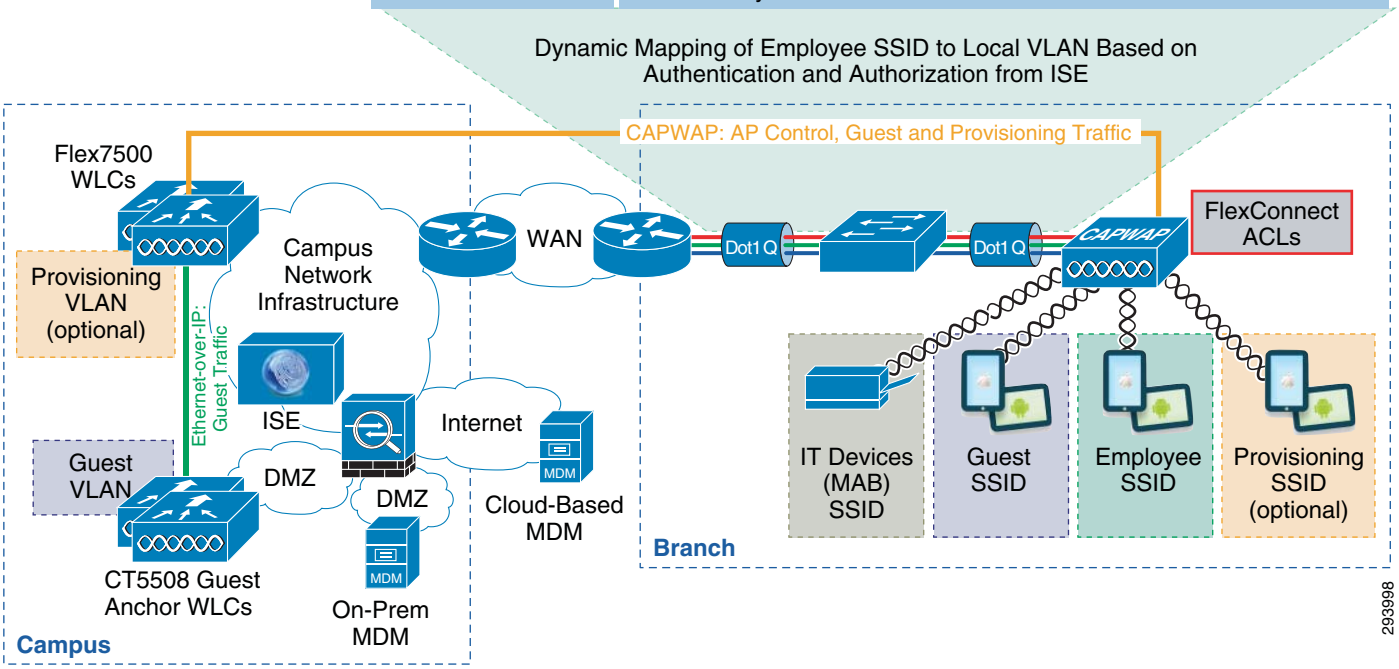
Access points connected to the access-layer switches within branch locations are still configured and controlled via one or more centralized wireless LAN controllers. In the case of this design guide, these controllers are a set of Cisco Flex 7500 wireless controllers—dedicated for branches—since they provide greater scalability for supporting access points in FlexConnect mode than Cisco CT5508 wireless controllers. Note also that with this design, guest wireless traffic is backhauled across the WAN to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus. Provisioning traffic (i.e. traffic from devices attempting to on-board with ISE) may also be backhauled across the WAN to the Flex7500 wireless controllers located within the campus.

Figure 5-11 shows at a high level how FlexConnect is implemented in the branch design.

Figure 5-11 High-Level View of the FlexConnect Wireless Branch Design

Dynamic VLAN assignment with a FlexConnect ACL applied at the wireless access point for differentiated access control.

VLAN Name	Description
Wireless_Full	Full Internal and Internet Access for On-Boarded Wireless Devices
Wireless_Partial	Partial Internal Access and Internet Access for On-Boarded Wireless Devices
Wireless_Internet	Internet Only Access for On-Boarded Wireless Devices

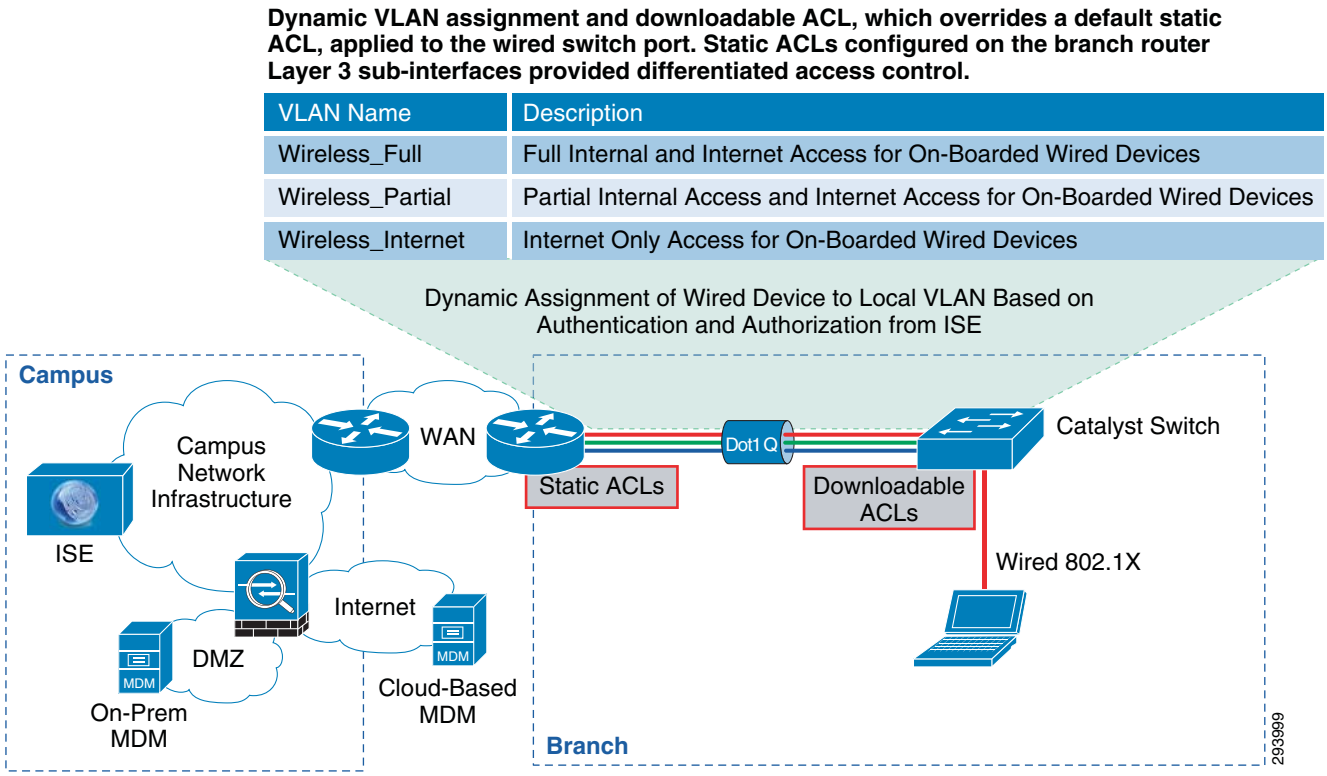


To implement the BYOD use cases for on-boarded devices, the method presented in this design guide for branch locations utilizing a FlexConnect wireless design is to place the device into an appropriate VLAN after it is authenticated and authorized. Statically configured FlexConnect ACLs applied per access point (or access point group) and per VLAN, provide differentiated access control for wireless devices. For example, a personal device which needs full access to the network is placed into a VLAN in which a FlexConnect ACL is configured on the access point with the right permissions. Personal devices that are granted partial access are placed in a different VLAN which has a different FlexConnect ACL.

Branch Wired Design

Figure 5-12 shows the wired design for a branch which does not implement Converged Access Catalyst 3850 Series switches. In other words, this is the wired design for a branch which implements switches such as the Catalyst 3750X, along with a FlexConnect wireless design.

Figure 5-12 High-Level View of Non-Converged Access Wired Branch Design

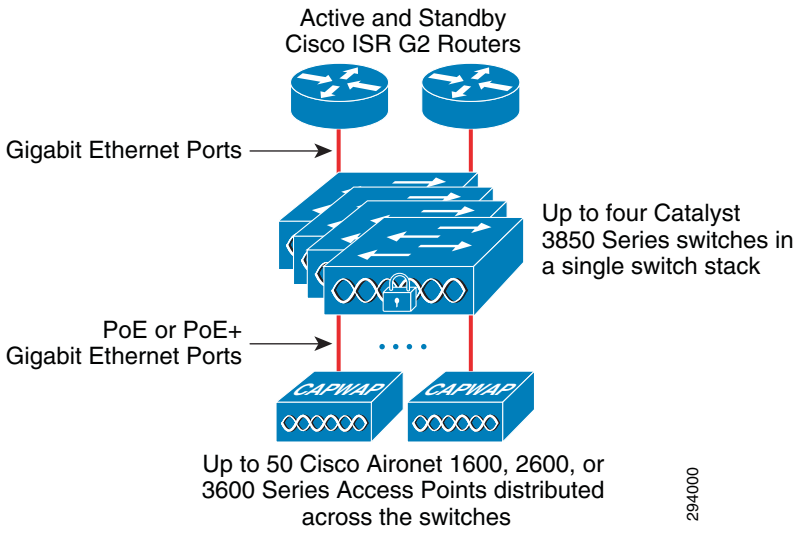


This design guide assumes that Catalyst switches are deployed as Layer 2 devices within the branch location. Wired devices authenticate using 802.1X against the ISE server centrally located within the campus. For this design, wired devices are also dynamically assigned to separate VLANs based on their access control requirements. A RADIUS downloadable ACL applied to the Catalyst 3750X Series switch overrides a pre-configured static ACL on each Catalyst switch port. Differentiated access control for the wired devices is provided by statically configured ACLs applied to the Cisco ISR G2 router Layer 3 sub-interfaces.

Converged Access Branch Design

The Converged Access branch BYOD design assumes a single Catalyst 3850 Series switch or switch stack deployed within a branch location. Hence this design applies to small to mid-sized branches only. This is shown in Figure 5-13.

Figure 5-13 Converged Access Branch Design Hardware



Up to four Catalyst 3850 Series switches may be deployed within a switch stack. The maximum number of access points supported per switch stack is 50, with up to a maximum of 2,000 wireless clients. The Catalyst 3850 Series supports up to 40 Gbps wireless throughput per switch (48-port models). Note that wireless performance requirements and physical distance limitations will often dictate the actual number of wireless access points and clients which can be deployed with this design. When a switch stack is implemented, APs should be deployed across the switches for wireless resilience purposes. This design guide will assume Catalyst 3850 Series switches deployed as Layer 2 switches within the branch location. Layer 3 connectivity within the branch is provided by the ISR routers which also serve as the WAN connectivity point for the branch. Future design guidance may address Catalyst 3850 Series switches deployed as Layer 3 switches within the branch location.

Note The Converged Access branch BYOD design may also be referred to as the Integrated Controller Branch BYOD design within this document.

As mentioned previously, Cisco has integrated wireless LAN controller functionality directly in the Catalyst 3850 Series switch. When access to local resources at the branch is a requirement, this allows for the termination of wireless traffic on the Catalyst 3850 switch itself, rather than backhauling traffic to a centralized wireless controller. As with FlexConnect designs, Converged Access designs can reduce Round Trip Time (RTT) delay, increase application performance, and reduce unnecessary hair-pinning of traffic when accessing resources local to the branch.

For the Converged Access branch BYOD design, the single Catalyst 3850 Series switch stack will implement the following wireless controller functionality:

- Mobility Agent (MA)—Terminates the CAPWAP tunnels from the access points (APs), and maintains the wireless client database.
- Mobility Controller (MC)—Manages mobility within and across sub-domains. Also manages radio resource management (RRM), WIPS, etc.

Since there is only a single switch stack, there is only a single Switch Peer Group (SPG). The Mobility Group, Mobility Sub-Domain, and Mobility Domain are entirely contained within the branch. No additional centralized wireless controllers are needed at the campus location, except for the Cisco CT5508 wireless controllers which function as the dedicated anchor controllers for wireless guest traffic. The access points within the branch locations are configured and controlled via the wireless LAN

controller functionality integrated within the Catalyst 3850 Series switch. Guest wireless traffic is still backhauled to a dedicated CT5508 guest anchor controller located on a DMZ segment within the campus. Provisioning traffic (i.e., traffic from devices attempting to on-board with ISE) is terminated locally on the Catalyst 3850 Series switch, with the Converged Access branch design. When implementing a dual-SSID design, provisioning traffic is terminated on a separate VLAN. All on-boarded devices terminate on a single VLAN with this design.

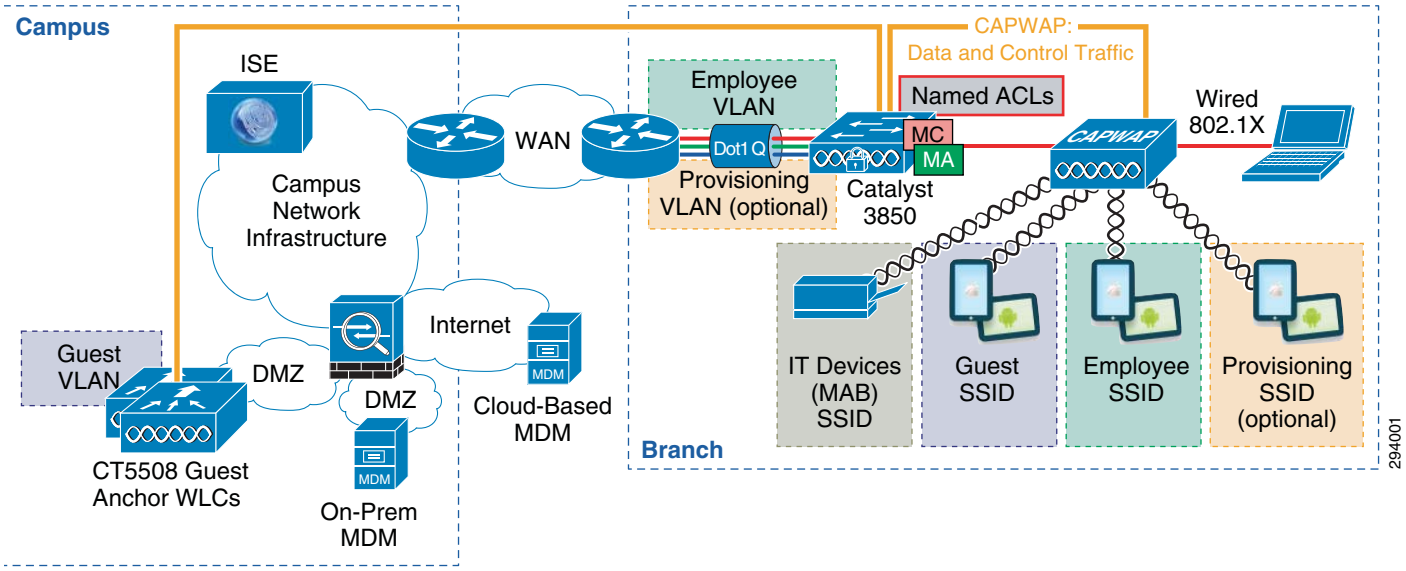
Note When deploying converged access wireless designs in which the Catalyst 3850 Series switch functions as the Mobility Controller (MC) and Mobility Agent (MA), it should be noted that the mobility tunnel for wireless guest access initiates from the Catalyst 3850 to the Guest anchor controller located within the DMZ. Hence, each branch will initiate a mobility tunnel for wireless guest access with this design. The maximum number of mobility controllers within a mobility domain is 72 for the CT5508 wireless controller. Therefore the maximum number of mobility anchor tunnels is limited to 71 for the CT5508 wireless controller. Therefore the network administrator may need to deploy additional CT5508 guest anchor controllers. Alternatively, the network administrator may look at providing direct Internet access from the branch for guest access. Future versions of this design guide may address such designs.

In order to implement the BYOD use cases, the method adopted in this design guide for branch locations utilizing a Converged Access design is to apply the appropriate dynamic ACL after the device is authenticated and authorized. This applies to both wired and wireless devices. The particular form of dynamic ACL is a RADIUS specified local ACL, otherwise known as a named ACL. These named ACLs, which must be configured on each Catalyst 3850 Series switch, provide differentiated access control. For example, a personal device which is granted full access to the network is statically assigned to the same VLAN as a personal device which is granted partial access. However different named ACLs are applied to each device, granting different access to the network. Since the named ACL is configured on the Catalyst 3850 switch specific to the particular branch, a single Cisco ISE policy can be implemented across multiple branches. However the Access Control Entries (ACEs) within the ACL for each branch can be unique to the IP addressing of the branch. This reduces the administrative complexity of the Cisco ISE policy, albeit at the expense of increased complexity of having to configure and maintain ACLs at each branch Catalyst 3850 Series switch.

Figure 5-14 shows at a high level how a Converged Access BYOD design is implemented in the branch.

Figure 5-14 High-Level View of the Converged Access Branch BYOD Design

Dynamic ACL (Named ACL) assignment applied at the switch for differentiated access control for wired and wireless devices.



Note that in the case of this design guide, on-boarded wired devices are also statically assigned to the same VLAN as wireless devices. Hence on-boarded wired and wireless devices will share the same VLAN, and hence the same IP subnet addressing space. It is recognized that customers may implement separate subnets for wired and wireless devices due to issues such as additional security compliance requirements for wireless devices. This will not be addressed within this version of the design guidance. Dynamically assigned named ACLs provide differentiated network access for wired devices.

The reason for the two methods of providing differentiated access between the FlexConnect and Converged Access branch designs is that prior to CUWN software version 7.5, FlexConnect did not allow the dynamic assignment of an ACL to an access point. It only allowed the dynamic assignment of a VLAN. The FlexConnect wireless design in this design guide is carried forward from the previous version of the design guide and continues to require a separate VLAN for each separate level of access control. This can increase the administrative burden of managing the branch network configuration. Converged access designs are more consistent with the campus wireless designs, requiring a single VLAN for multiple levels of access control.



Mobile Device Managers for BYOD

Revised: August 7, 2013

Mobile Device Managers (MDMs) secure, monitor, and manage mobile devices, including both corporate-owned devices as well as employee-owned BYOD devices. MDM functionality typically includes Over-the-Air (OTA) distribution of policies and profiles, digital certificates, applications, data and configuration settings for all types of devices. MDM-supported and managed devices include not only handheld devices, such as smartphones and tablets, but increasingly laptop and desktop computing devices as well.

Critical MDM functions include-but are not limited to:

- PIN enforcement—Enforcing a PIN lock is the first and most effective step in preventing unauthorized access to a device; furthermore, strong password policies can also be enforced by an MDM, reducing the likelihood of brute-force attacks.
- Jailbreak/Root Detection—Jailbreaking (on Apple iOS devices) and rooting (on Android devices) are means to bypass the management of a device and remove SP control. MDMs can detect such bypasses and immediately restrict a device's access to the network or other corporate assets.
- Data Encryption—Most devices have built-in encryption capabilities-both at the device and file level. MDMs can ensure that only devices that support data encryption and have it enabled can access the network and corporate content.
- Data Wipe—Lost or stolen devices can be remotely full- or partial-wiped, either by the user or by an administrator via the MDM.
- Data Loss Prevention (DLP)—While data protection functions (like PIN locking, data encryption and remote data wiping) prevent unauthorized users from accessing data, DLP prevents authorized users from doing careless or malicious things with critical data.
- Application Tunnels—Secure connections to corporate networks are often a mandatory requirement for mobile devices.

Cisco ISE 1.2 with MDM API Integration

While Cisco ISE provides critical policy functionality to enable the BYOD solution, it has limited awareness of device posture. For example, ISE has no awareness of whether a device has a PIN lock enforced or whether the device has been jailbroken or whether a device is encrypting data, etc. On the other hand, MDMs have such device posture awareness, but are quite limited as to network policy enforcement capacity.

Therefore, to complement the strengths of both ISE and MDMs, ISE 1.2 includes support of an MDM integration API which allows it to both:

- Pull various informational elements from MDM servers in order to make granular network access policy decisions that include device-details and/or device-posture.
- Push administrative actions to the managed devices (such as remote-wiping) via the MDM.

As of the publication date of this CVD, ISE 1.2 supports an API for MDM integration with the following third-party MDM vendors:

- AirWatch
- MobileIron
- Good Technologies
- XenMobile
- SAP Afaria
- FiberLink Maas360

The following MDM API pull/push capabilities are supported in ISE 1.2 for all third-party MDM systems:

- PIN lock Check
- Jailbroken Check
- Data Encryption Check
- Device Augmentation Information Check
- Registration Status Check
- Compliance Status Check
- Periodic Compliance Status Check
- MDM Reachability Check
- (Full/Partial) Remote Wipe
- Remote PIN lock

MDM Deployment Options and Considerations

With MDM solutions, there are two main deployment models:

- On-Premise—In this model, MDM software is installed on servers in the corporate DMZ or data center, which are supported and maintained by the enterprise IT staff.
- Cloud-based—In this model-also known as a MDM Software-as-a-Service (SaaS) model-MDM software is hosted, supported and maintained by a provider at a remote Network Operation Center (NOC); customers subscribe on a monthly or yearly basis and are granted access to all MDM hardware/software via the Internet.

Before deploying a MDM, businesses must make the pivotal decision of whether their MDM solutions should be on premise (on-prem) or cloud-based. Several business and technical factors are involved in this decision, including:

- Cost—Cloud-based MDM solutions often are more cost-effective than on-prem; this is because these eliminate the need for incremental and ongoing hardware, operating system, database and networking costs associated with a dedicated MDM server. Also avoided is any additional training

that may be required by IT staff to support these servers. From a cloud-provider's perspective: since these fixed infrastructure costs have already been invested, there are very little marginal cost to provisioning custom-tailored virtual-instances to enterprise subscribers, and as such, these can be priced attractively.

- Control—On-prem models offer enterprises the greatest degree of control, of not only the MDM solution, but also the enterprise systems that these integrate with (such as the corporate directory, certificate authority, email infrastructure, content repositories and management systems-all of which will be discussed in additional detail below). This is because an on-prem model requires no transmission or storage of corporate data offsite. Conversely, a cloud-based service requires giving up a level of control over the overall solution, as confidential information, data and documents will be required to be transmitted to the provider, and (depending on the details of the service) may also be stored offsite. Cloud providers may also update the software on the servers without following the enterprise change control protocol.
- Security—On-prem MDM models are often perceived as being more secure than cloud-based models; however, this perceived difference in security-levels may be lessening, especially when considering that over $14B of business was securely conducted via SaaS in 2012 alone. Ultimately, the security of a system will principally depend, not only on the technologies deployed, but also on the processes in place to keep the hardware and software updated and managed properly.
- Intellectual Property—Most MDMs support secure isolation of corporate data on the devices they manage; however, these systems typically require corporate data to be passed through the MDM in order to be transmitted OTA to the device's secure and encrypted compartment. This process may represent an additional security concern in a cloud-based model, as now the enterprise is called on to trust the MDM SaaS provider with not only device management, but also with intellectual property and confidential data.
- Regulatory Compliance—Regulatory compliance can dictate where and how financial, healthcare and government (and other) organizations can store their data. Such regulations include PCI, HIPAA, HITECH, Sarbanes-Oxley, and even the US Patriot Act. Such regulations may preclude storing sensitive information in the cloud, forcing the choice of an on-prem MDM model.
- Scalability—Cloud-based models offer better scalability than on-prem models, as these can accommodate either small or large deployments (and anything in-between) without any increased infrastructure costs to the subscriber. Conversely, on-prem models may have difficulty in cost-effectively accommodating small deployments. For example, consider the cost of deploying an MDM server that can support 100,000 devices being deployed to support only 100. Additionally, on-prem models will incrementally require more hardware and infrastructure as the number of devices increases.
- Speed of Deployment—Cloud-based solutions are typically faster to deploy (and can often be enabled the same-day as these are ordered), whereas on-prem solutions often take a couple of weeks (or more) to plan out, install and deploy.
- Flexibility—Cloud-based MDM solutions typically have day-one support for new releases of device hardware and software; alternatively, on-prem solutions will require an upgrade to the MDM software for each new device/software supported.
- Ease of Management—With on-prem models, the IT department must ensure the MDM has all the latest updates; in a cloud-based system, this responsibility rests with the provider.



Note Cisco is not advocating the use of one MDM deployment model over another, nor does Cisco recommend any specific third-party MDM solution. These business and technical considerations are included simply to help draw attention to the many factors that an IT architect may find helpful in reviewing when evaluating which MDM solution works best to meet their specific business needs.

On-Premise

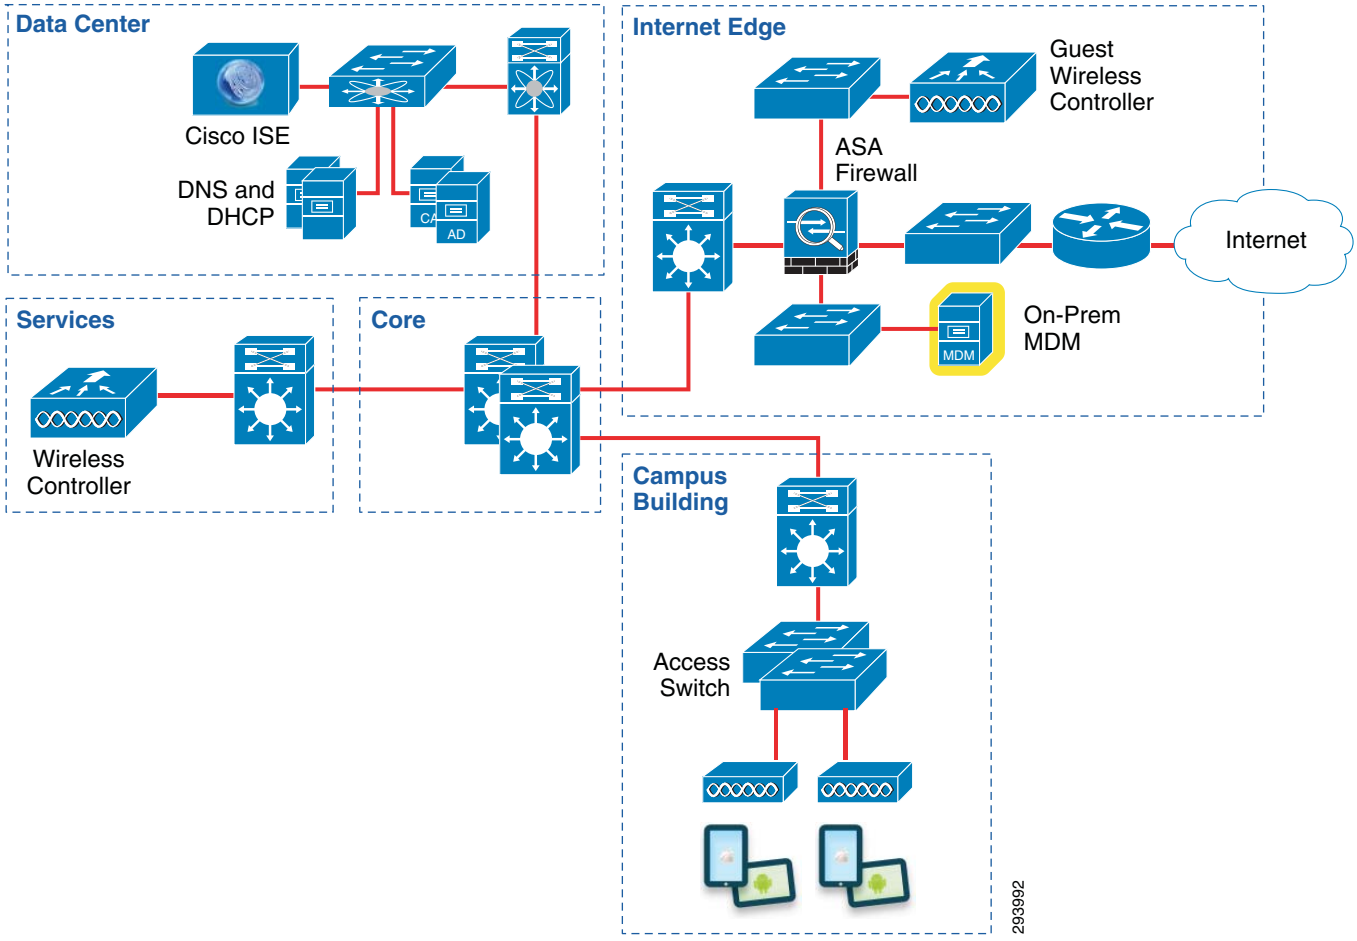
In the on-premise MDM deployment model, the MDM software resides on premises on a dedicated server (or servers), typically within the Internet Edge or DMZ.

This model is generally better suited to IT staff that have a higher-level of technical expertise (such that they can configure, periodically-update and manage such a server) or to enterprises that may have stricter security/confidentiality requirements (which may preclude the management of their devices by a cloud-based service).

The on-premise model may also present moderate performance benefits to some operational flows (due to its relative proximity to the devices, as opposed to a cloud-based service). For example, if a network access policy included the "MDM Reachability" check, this test would likely be much more responsive in an on-premise MDM deployment model versus a cloud-based model.

The network topology for a campus BYOD network utilizing an On-Prem MDM deployment model is illustrated in Figure 6-1.

Figure 6-1 Campus BYOD Network with On-Prem MDM (at the Internet Edge)



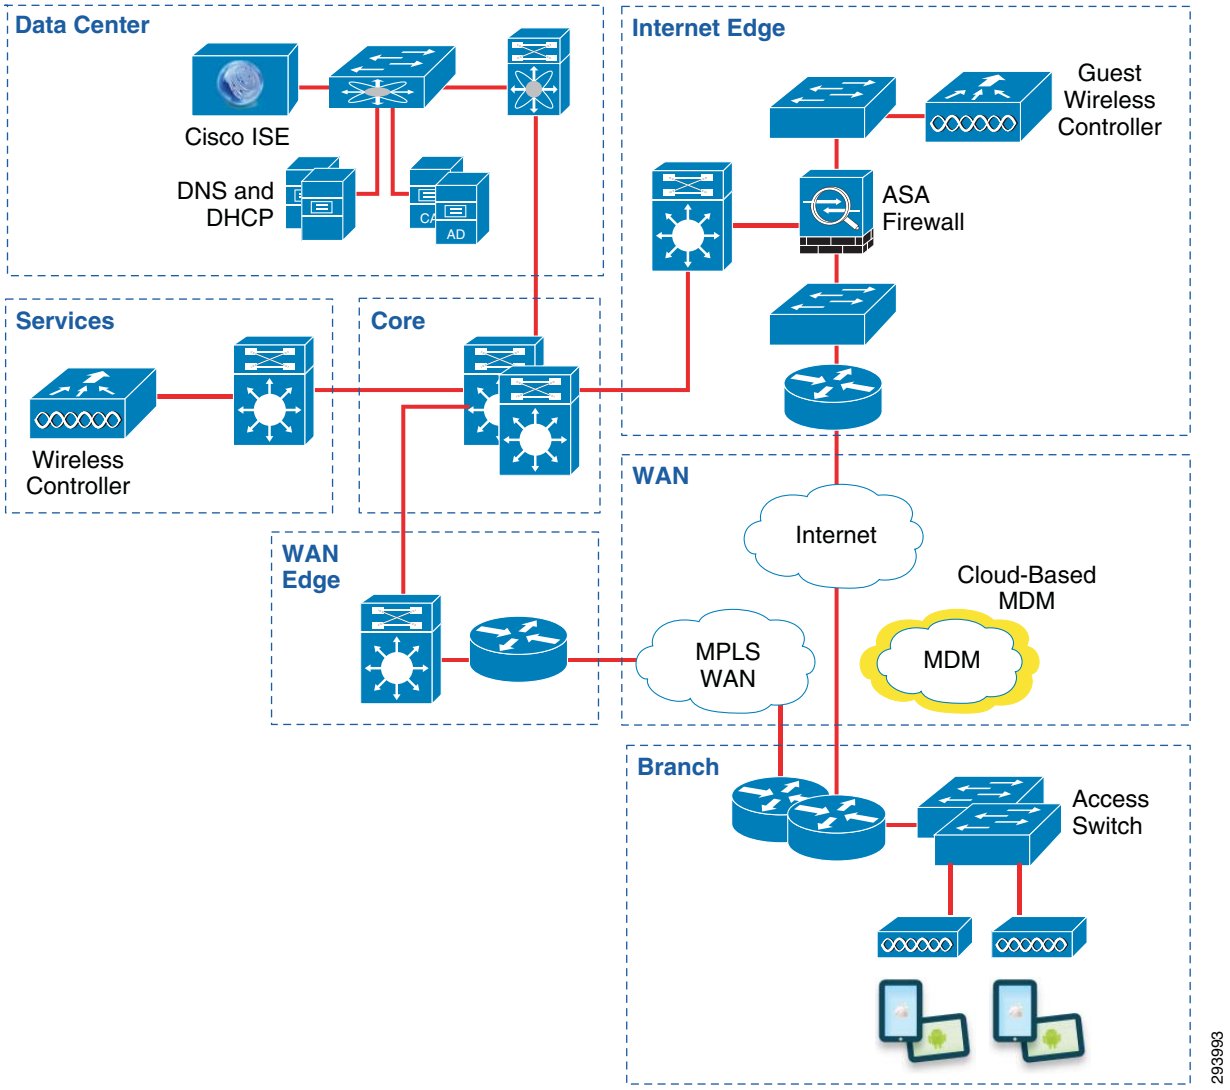
Cloud-Based

In the Cloud-Based MDM deployment model, MDM functionality is delivered to customers in a SaaS manner: the software resides wholly within the MDM vendor's cloud, with a custom-tailored virtual instance provided for each customer.

From a customer's perspective, this model is greatly simplified (as now they do not have to configure, update, maintain and manage the MDM software); however, as a trade-off, they relinquish a degree of control over all their devices (and also some of the data on these devices) to the third-party MDM SaaS provider, which may pose security concerns. As such, this model may be better suited to small- or medium-sized businesses that have moderate IT technical expertise and unexceptional security requirements.

The network topology for a branch BYOD network utilizing a cloud-based MDM deployment model is illustrated in Figure 6-2.

Figure 6-2 Branch BYOD Network with Cloud-Based MDM



Enterprise Integration Considerations

In addition to the integrating the corporate network with the MDM—which is discussed in great detail in this document—other enterprise services and resources are also important to integrate with the MDM system, including:

- Corporate Directory Services
- Certificate Authority (CA) and Public Key Infrastructure (PKI)
- Email Infrastructure
- Content Repositories
- Management Systems

Corporate Directory Services Integration

Corporate directory services (such as LDAP-based directory, Active Directory, etc.) can be leveraged by MDMs to efficiently organize and manage user access. Administrators can assign device profiles, apps, and content to users based on their directory-group memberships. Additionally, some MDMs can detect directory changes and automatically update device-policies. For example, if a user is deactivated in a directory system, then the MDM can remove device-based corporate network access and selectively wipe the device.

Corporate Certificates Authority and Public Key Infrastructure Integration

Certificate Authorities (such as Microsoft CA) or SCEP certificate services providers (such as MSCEP and VeriSign) can be leveraged by MDMs to assign and verify certificates for advanced user authentication and to secure access to corporate systems. CA integration ensures message integrity, authenticity and confidentiality. Additionally CA integration enables client authentication, encryption and message signatures.

Furthermore, MDMs can also integrate with Public Key Infrastructure (PKI) or third-party providers to configure certificates and distribute these to devices without user interaction.

Email Integration

The corporate email infrastructure can be integrated with the MDM solution to provide security, visibility and control in managing mobile email. This enables employees to access corporate email on their mobile devices without sacrificing security. Additionally such integration facilitates the management of mobile email (such as configuring email settings over-the-air, blocking unmanaged devices from receiving email, enforcing device encryption, etc.) The MDMs approach to email management varies among MDM providers and is feature differentiator. Email policy information is not available to ISE via the API.

Content Repository Integration

Integrating MDM systems with content repositories enables administrators to deliver secure mobile access to corporate documents while managing document distribution and access permissions (including the ability to view, view-offline, email, or print on a document). This ensures the right content gets to

the right employees without sacrificing the security of the documents themselves, which are distributed to mobile devices over encrypted connections. Furthermore, files and documents can be synchronized with corporate file systems and share points, so that the latest version of a document is automatically updated on employee mobile devices. To ensure security, users can be authenticated with a username, password, and certificate before they can access corporate content. Additionally, document metadata (including author, keywords, version, and dates created or modified) can be restricted on a per-user basis.

Management Integration

MDM systems can be integrated with enterprise management systems for enhanced logging, recording and reporting of device and console events. Event logging settings can be configured based on severity levels, with the ability to send specific levels to external systems via Syslog integration. Events can include login events, failed login attempts, changes to system settings and configurations, changes to profiles, apps and content, etc. Such management systems integration ensures security and compliance with regulations and corporate policies.

Integration Servers

The integration of these enterprise systems with MDMs in on-premise deployment models is relatively straightforward, as it is largely a matter of ensuring the proper protocols are configured correctly and the necessary ports are opened in any firewalls within the paths. However, in cloud-based deployment models, such integration requires secure transport protocols (such as over HTTPS) from the customer to the MDM service provider and/or a specialized MDM integration servers (or similar proxy-servers) located within the client's DMZ.



Application Considerations and License Requirements for BYOD

Revised: August 7, 2013

When implementing a BYOD solution, the applications that run on employee-owned devices need to be considered before selecting which of the particular BYOD use cases discussed above to deploy. The application requirements for these devices determine the level of network connectivity needed. The network connectivity requirements in turn influence the choice of the BYOD use case to apply.

Quality of Service

In addition to network connectivity, quality of service (QoS) is an important consideration for applications, especially those delivering real-time media. Device specific hardware, such as dedicated IP phones which send only voice traffic, allowed for the configuration of dedicated voice wireless networks. However, with the widespread use of smartphones and tablets which support collaboration software (such as Cisco's Jabber client), devices are capable of sending voice, video, and data traffic simultaneously. Hence, QoS is necessary to provide the necessary per-hop behavior as such traffic traverses the network infrastructure.

QoS can be categorized into the following broad functions:

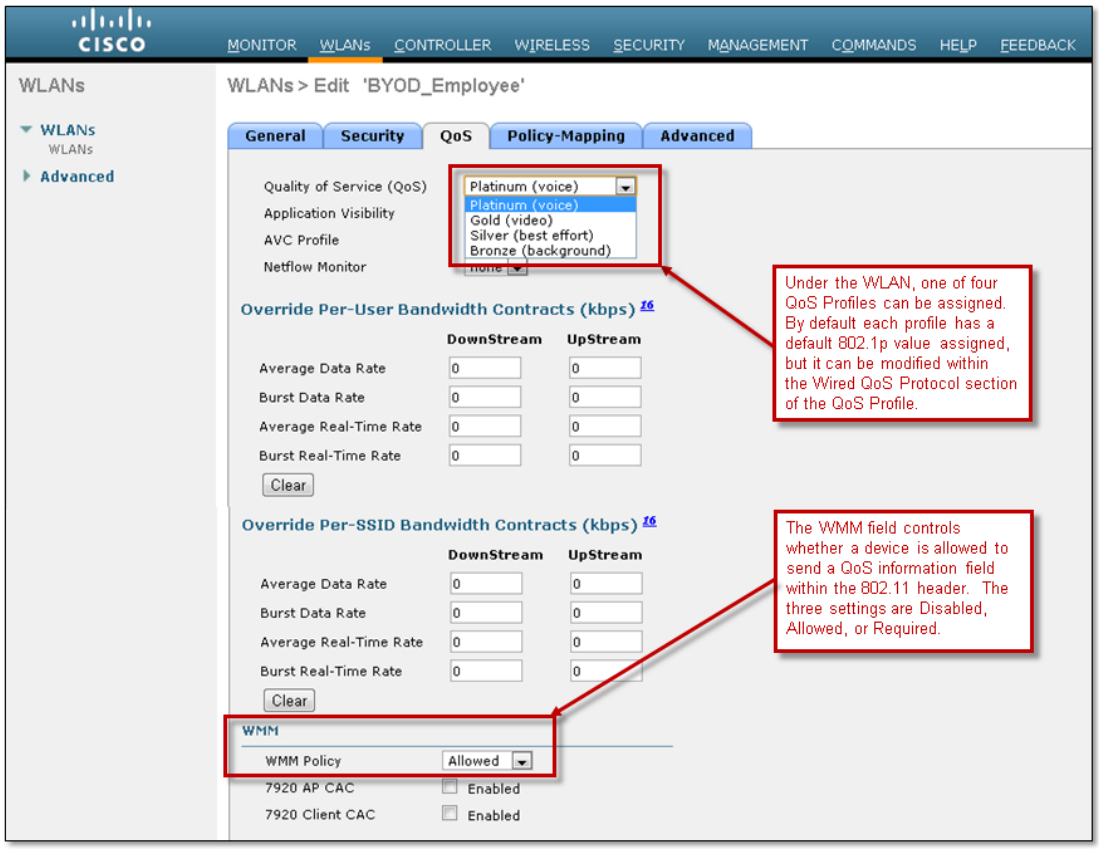
- Classification and Marking-including Application Visibility and Control (AVC)
- Bandwidth Allocation/Rate Limiting (Shaping and/or Policing)
- Trust Boundary Establishment
- Queueing

For a discussion regarding implementing wired QoS, refer to Medianet Campus QoS Design 4.0 at: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html.

The following sections discuss various aspects of wireless QoS.

As of Cisco Unified Wireless Network (CUWN) software release 7.3 and above, wireless QoS is configured by applying one of four QoS Profiles—Platinum, Gold, Silver, or Bronze—to the WLAN to which a particular client device is associated. An example of the configuration is shown in Figure 7-1.

Figure 7-1 Application of a QoS Profile to a WLAN



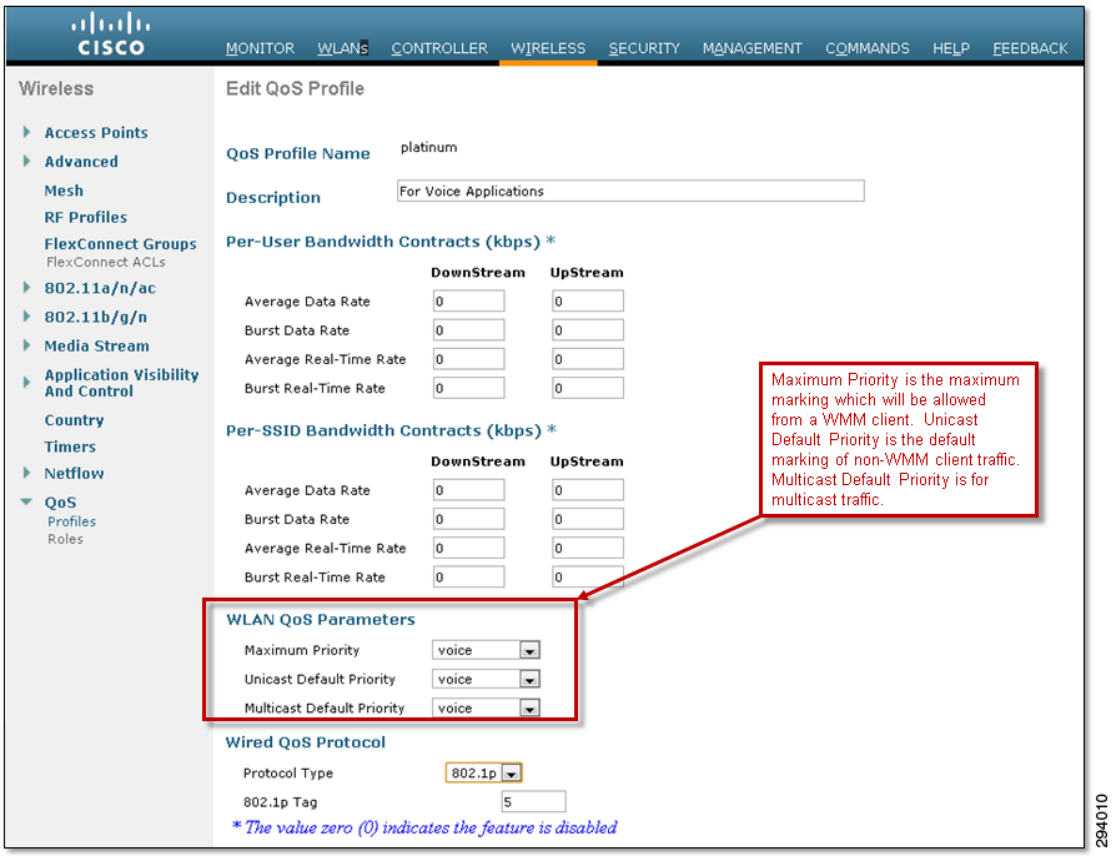
Note that the QoS settings for the profile can be overridden on a per-WLAN basis from within the QoS tab of the WLAN configuration.

The DSCP marking of client traffic, as it traverses the network within a CAPWAP tunnel, is controlled by three fields within the WLAN QoS Parameters field within the QoS Profile:

- Maximum Priority—This is the maximum 802.11 User Priority (UP) value of a packet sent by a Wi-Fi Multimedia (WMM)-enabled client which will be allowed by the access point. The User Priority maps to a DSCP value within the outer header of the CAPWAP tunnel as the packet traverses the network infrastructure. If the WMM-enabled client sends an 802.11 packet with a User Priority higher than allowed, the access point marks the packet down to the maximum allowed User Priority. This in turn maps to the DSCP value of the external CAPWAP header as the packet is sent over the network infrastructure.
- Unicast Default Priority—This is the default 802.11 User Priority (UP) to which a unicast packet sent by a non-WMM-enabled client is assigned. This User Priority also maps to the DSCP value within the outer header of the CAPWAP tunnel as the packet traverses the network infrastructure.
- Multicast Default Priority—This is the default 802.11 User Priority (UP) for multicast traffic. This User Priority maps to a DSCP value within the outer header of the CAPWAP tunnel as the packet traverses the network infrastructure.

An example of the configuration of the WLAN QoS Parameters is shown in Figure 7-2.

Figure 7-2 Controlling the Marking of Wireless Packets



It should be noted that these settings apply primarily to Local Mode (centralized wireless controller) designs and FlexConnect designs with central termination of traffic, since the WLAN QoS Parameters field results in the mapping of the 802.11 User Priority to the DSCP value within the outer header of the CAPWAP tunnel.

The original DSCP marking of the packet sent by the wireless client is always preserved and applied as the packet is placed onto the Ethernet segment, whether that is at the wireless controller for centralized wireless controller designs or at the access point for FlexConnect designs with local termination.

The wireless trust boundary is established via the configuration of the WMM Policy within the QoS tab of the WLAN configuration. An example was shown in Figure 7-1. The three possible settings for WMM Policy are:

- Disabled—The access point will not allow the use of QoS headers within 802.11 packets from WMM-enabled wireless clients on the WLAN.
- Allowed—The access point will allow the use of QoS headers within 802.11 packets from wireless clients on the WLAN. However the access point will still allow non-WMM wireless clients (which do not include QoS headers) to associate to the access point for that particular WLAN.
- Required—The access point requires the use of QoS headers within 802.11 packets from wireless clients on the WLAN. Hence, any non-WMM-enabled clients (which do not include QoS headers) will not be allowed to associate to the access point for that particular WLAN.

Note Where possible, it is advisable to configure WMM policy to Required. Some mobile devices may incorrectly mark traffic from collaboration applications when the WMM policy is set to Allowed versus Required. Note however that this requires all devices on the WLAN to support WMM before being allowed onto the WLAN. Before changing the WMM policy to Required, the network administrator should verify that all devices which utilize the WLAN are WMM-enabled. Otherwise, non-WMM-enabled devices will not be able to access the WLAN.

The configuration of the WMM Policy, along with the WLAN QoS Parameters, together create the wireless QoS trust boundary and determine the marking of wireless traffic within the CAPWAP tunnel as it traverses the network infrastructure.

Note The Cisco CT5760 wireless controller and the Catalyst 3850 Series switch both run IOS XE software. QoS configuration uses the Modular QoS based CLI (MQC) which is in alignment other platforms such as Catalyst 4500E Series switches. This version of the design guide does not address QoS on the Cisco CT5760 wireless controller and Catalyst 3850 Series switch. Future versions may address QoS on these platforms.

Rate Limiting

One additional option to prevent the wireless medium from becoming saturated, causing excessive latency and loss of traffic, is rate limiting. Rate limiting may be implemented per device or per SSID to prevent individual devices from using too much bandwidth and negatively impacting other devices and applications. Rate limiting is particularly useful for guest access implementations and is discussed in detail in Chapter 21, “BYOD Guest Wireless Access.”

Application Visibility and Control (AVC)

Beginning with Cisco Unified Wireless Network (CUWN) software release 7.4, the Application Visibility and Control set of features—already supported on Cisco routing platforms such as ASR 1000s and ISR G2s—became available on WLC platforms, including the Cisco 2500, 5500, 7500, 8500 WLCs, and WiSM2 controllers on Local and FlexConnect Modes (for WLANs configured for central switching only in 7.4 release).

The AVC feature set increases the efficiency, productivity, and manageability of the wireless network. Additionally, the support of AVC embedded within the WLAN infrastructure extends Cisco’s application-based QoS solutions end-to-end.

Business use-cases for AVC policies include:

- Guaranteeing voice quality from wireless applications meets enterprise VoIP requirements.
- Ensuring video applications—both interactive and streaming—are delivered to/from wireless devices with a high Quality of Experience, so that users can communicate and collaborate more efficiently and effectively-regardless of their location or device.
- Provisioning preferred services for business-critical applications running on wireless devices, such as Virtual Desktop applications, sales applications, customer relationship management (CRM) applications, and enterprise resource planning (ERP) applications, etc.

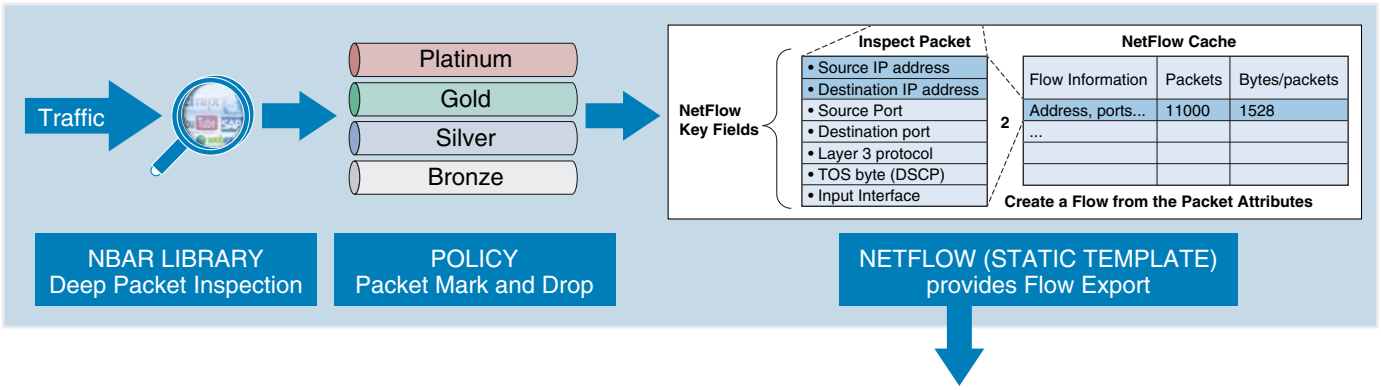
- De-prioritizing “background” application traffic (i.e., applications that send data to/from servers, rather than directly to other users and which do not directly impact user-productivity), such as email, file-transfers, content distribution, backup operations, software updates, etc.
- Identifying and de-prioritizing (or dropping) non-business applications, which can include social networking applications, peer-to-peer file-sharing applications, and type of entertainment and/or gaming applications so that network resources are always available for business-oriented applications.

AVC includes these components:

- Next-generation Deep Packet Inspection (DPI) technology called Network Based Application Recognition (NBAR2), which allows for identification and classification of applications. NBAR is a deep-packet inspection technology available on Cisco IOS based platforms, which includes support of stateful L4-L7 classification.
- QoS—Ability to remark applications using DiffServ, which can then be leveraged to prioritize or de-prioritize applications over both the wired and wireless networks.
- A template for Cisco NetFlow v9 to select and export data of interest to Cisco Prime or a third-party NetFlow collector to collect, analyze, and save reports for troubleshooting, capacity planning, and compliance purposes.

These AVC components are shown in Figure 7-3.

Figure 7-3 Cisco AVC Components



AVC on the WLC inherits NBAR2 from Cisco IOS that provides deep packet inspection technology to classify stateful L4-L7 application classification. This is critical technology for application management, as it is no longer a straightforward matter of configuring an access list based on the TCP or UDP port number(s) to positively identify an application. In fact, as applications have matured—particularly over the past decade—an ever increasing number of applications have become opaque to such identification. For example, HTTP protocol (TCP port 80) can carry thousands of potential applications within it and in today’s networks seems to function more as a transport protocol rather than as the OSI application-layer protocol that it was originally designed as. Therefore to identify applications accurately, Deep Packet Inspection technologies—such as NBAR2—are critical.

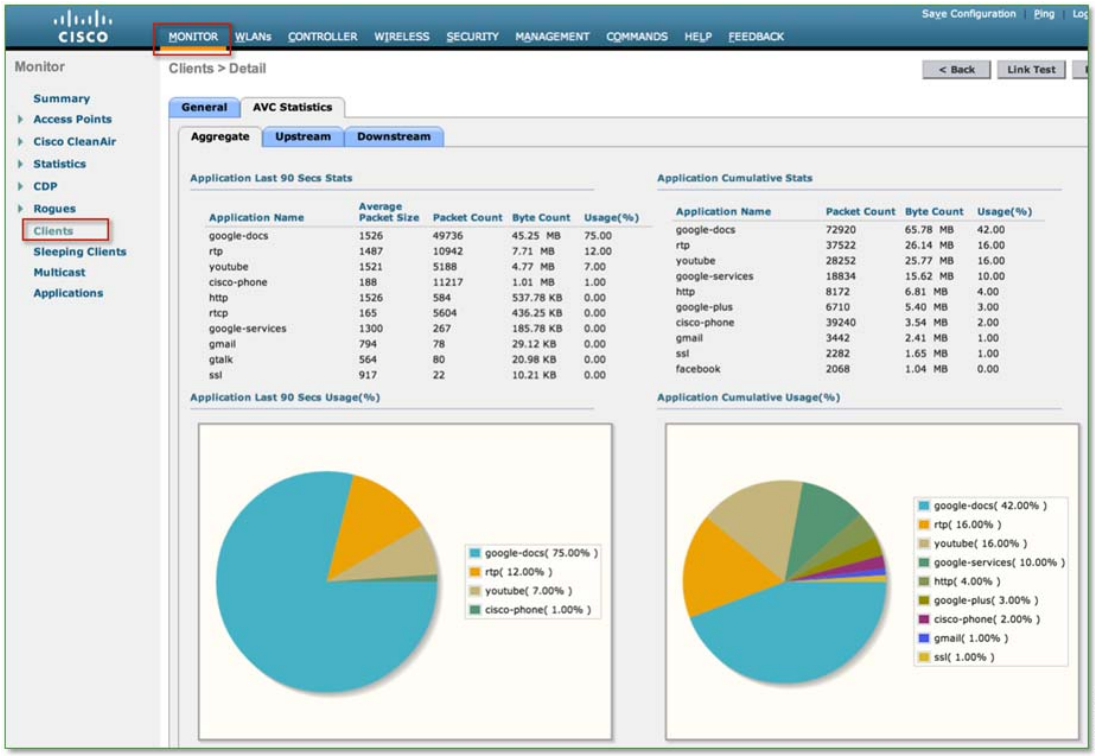
Once applications are recognized by the NBAR engine by their discrete protocol signatures, it registers this information in a Common Flow Table so that other WLC features can leverage this classification result. Such features include Quality of Service (QoS), NetFlow, and firewall features, all of which can take action based on this detailed classification.

Thus AVC provides:

- Application Visibility on the Cisco WLC by enabling Application Visibility for any WLAN configured. Once Application Visibility is turned on, the NBAR engine classifies applications on that particular WLAN. Application Visibility on the WLC can be viewed at an overall network level, per WLAN, or per client. An example of a per-client application visibility report is illustrated in Figure 7-4.
- Application Control on the Cisco WLC by creating an AVC profile (or policy) and attaching it to a WLAN. The AVC Profile supports QoS rules per application and provides the following actions to be taken on each classified application: Mark (with DSCP), Permit (and transmit unchanged), or Drop. An example of an AVC profile is shown in Figure 7-5, Figure 7-6, and Figure 7-7.

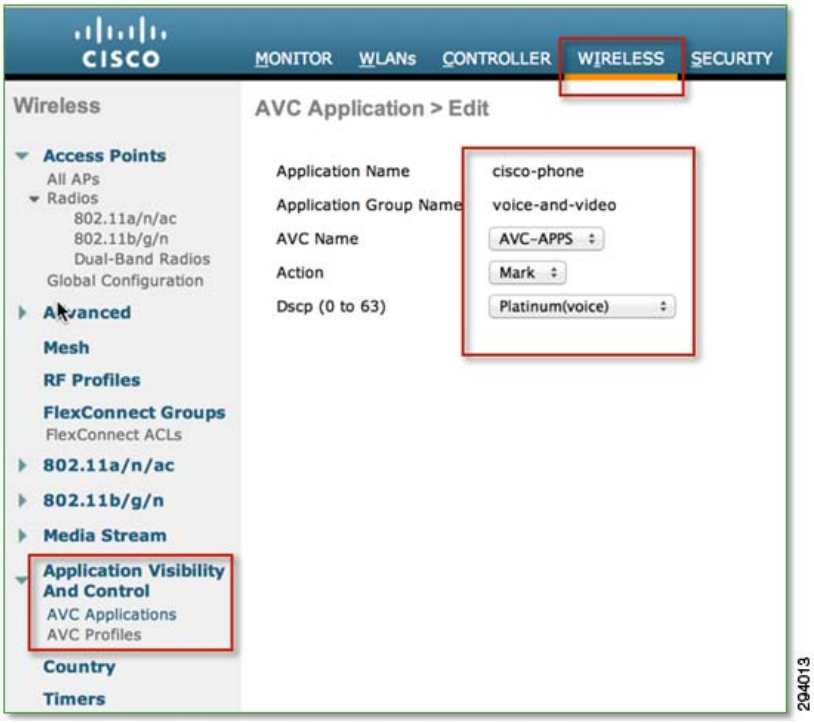
A client-based AVC report—such as shown in Figure 7-4—can show the top applications by device. AVC reports can also be compiled by WLAN or at the overall network level.

Figure 7-4 Cisco AVC Application Visibility Reports



An AVC profile—a collection of individual application policy rules—can be configured via the WLC GUI or CLI. In Figure 7-5 an AVC application rule is being configured for voice traffic sourced-from or destined-to Cisco wireless devices. This traffic is identified via an NBAR2 signature named **cisco-phone** and is marked as DSCP 46 (EF) and assigned to the Platinum Wireless Multi-Media (WMM) access-category for the highest level of service over the air.

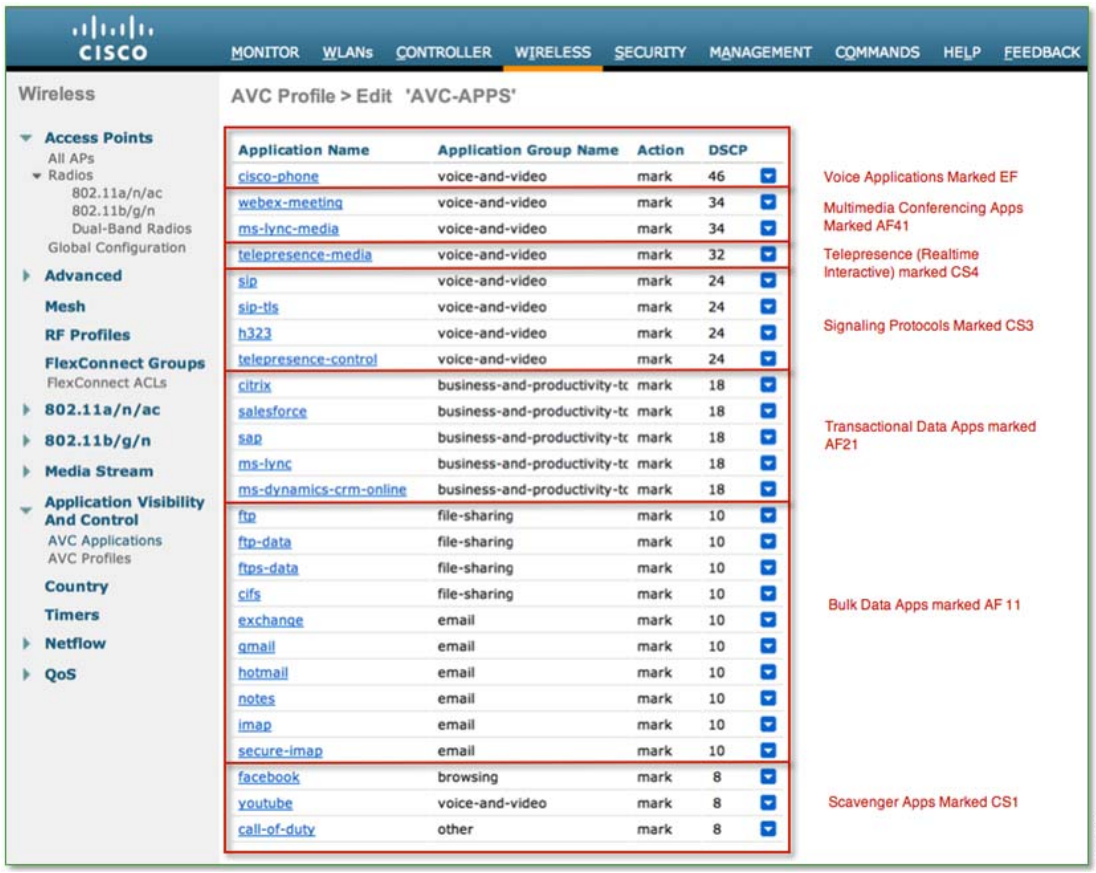
Figure 7-5 Cisco AVC Profile Example 1—Creating an AVC Policy Rule



An AVC profile can contain up to 32 individual application rules, as is shown in Figure 7-6, containing recommended policies for the following classes of application traffic (as based on RFC 4594):

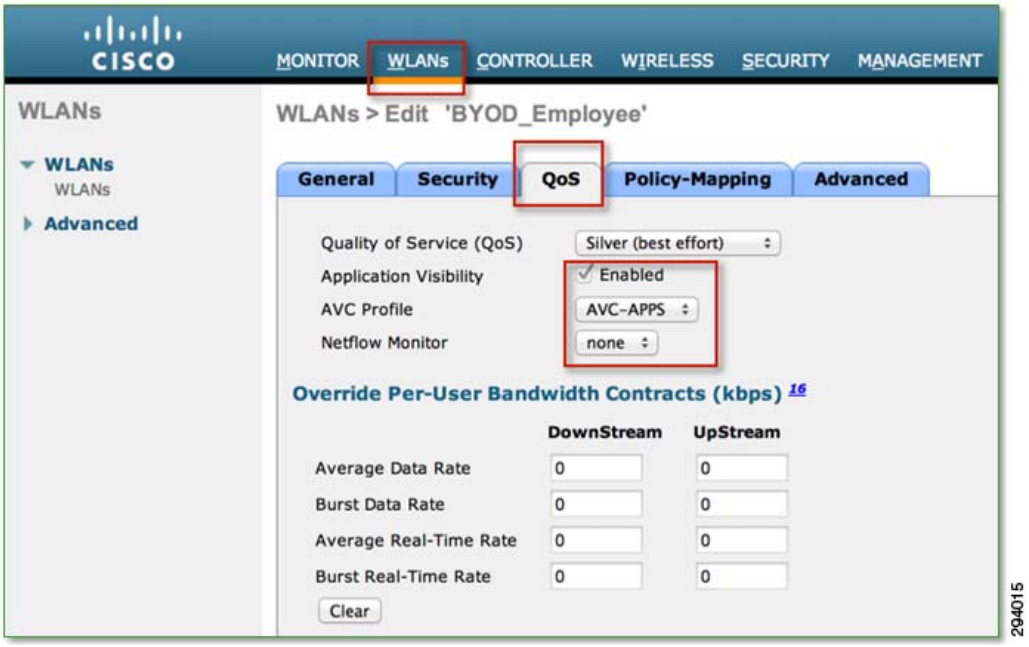
- Voice
- Video
- Multimedia Conferencing
- Multimedia Streaming
- Transactional Data
- Bulk Data
- Scavenger applications

Figure 7-6 Cisco AVC Profile Example 2—Displaying a Comprehensive AVC Policy



Once an AVC profile has been assembled, it can be applied to a WLAN(s), as shown in Figure 7-7. AVC policies are applied bi-directionally—that is, in the upstream and downstream directions simultaneously.

Figure 7-7 Cisco AVC Profile Example 3—Applying an AVC Profile to a WLAN



AVC supports over 1000 applications in its initial release for WLCs. Some of these applications-grouped by business case-are:

To ensure voice quality for wireless devices, the **cisco-phone** application would typically be assigned to the Platinum (Voice) WMM access category via AVC. However, additional VoIP applications may include:

- **aol-messenger-audio**
- **audio-over-http**
- **fring-voip**
- **gtalk-voip**
- **yahoo-voip-messenger**
- **yahoo-voip-over-sip**

Similarly, to protect video and multimedia applications, the following applications might be assigned to the Gold (Video) WMM access-category via AVC:

- **cisco-ip-camera**
- **telepresence-media**
- **webex-meeting**
- **ms-lync-media**
- **aol-messenger-video**
- **fring-video**
- **gtalk-video**
- **livemeeting**
- **msn-messenger-video**
- **rhapsody**

- **skype**
- **video-over-http**

Note It may be that some of these video conferencing applications may be considered non-business in nature (such as Skype and gtalk-video), in which case these may be provisioned into the Bronze (Background) WMM access category.

To deploy AVC policies to protect the signaling protocols relating to these voice and video applications, the following applications might be marked to the Call-Signaling marking of CS3 (DSCP 24) via AVC:

- **sip**
- **sip-tls**
- **skinny**
- **telepresence-control**
- **h323**
- **rtcp**

To deploy policies to protect business-critical applications, the following applications might be marked AF21 (DSCP 18) via AVC:

- **citrix**
- **ms-lync**
- **ms-dynamics-crm-online**
- **salesforce**
- **sap**
- **oraclenames**
- **perforce**
- **phonebook**
- **semantix**
- **synergy**

On the other hand, some business applications would be best serviced in the background by assigning these to the Bronze (Background) WMM access category via AVC:

- **ftp/ftp-data/ftps-data**
- **cifs**
- **exchange**
- **notes**
- **smtp**
- **imap/secure imap**
- **pop3/secure pop3**
- **gmail**
- **hotmail**
- **yahoo-mail**

And finally, many non-business applications can be controlled by either being assigned to the Bronze (Background) WMM access category or dropped via AVC policies:

- **youtube**
- **netflix**
- **facebook**
- **twitter**
- **bittorrent**
- **hulu**
- **itunes**
- **picasa**
- **call-of-duty**
- **doom**
- **directplay8**



Note It is important to note that these are only example applications and do not represent an exhaustive list of applications by class. With over a thousand applications to choose from, these lists are simplified for the sake of brevity and serve only to illustrate AVC policy options and concepts.

For comprehensive design guidance on using the AVC feature for WLCs, see: Chapter 24, "Mobile Traffic Engineering with Application Visibility and Control (AVC)."

Cisco Jabber

Cisco's Jabber clients are unified communications (UC) applications that are available for Android and Apple mobile devices as well as Microsoft Windows and Apple Mac computers. These client applications provide instant messaging (IM), presence, voice, video, and visual voicemail features. These features require that the employee-owned device is allowed to establish call signaling flows between the device itself and the corporate Cisco Unified Communications Manager (Unified CM) server, typically deployed within the campus data center. Note that the Basic Access use case discussed above terminates employee-owned devices on a DMZ segment off of the Internet Edge firewall. Cisco Jabber requires only Internet access to access WebEx cloud-based services like IM, meetings, and point-to-point voice and video calls. However, to deliver these same services with on-premise corporate assets such as Unified CM and other back-end UC applications, connectivity through the firewall is required for Jabber features to function. In addition to signaling, media flows also need to be allowed between the Jabber client and other IP voice and video endpoints, such as corporate IP phones deployed throughout the corporate network. This requires the network administrator to allow a range of addresses and ports inbound from the DMZ segment through the Internet Edge firewall. Given these connectivity considerations for real time communications and collaboration, the network administrator may instead decide to implement the Enhanced Access use case discussed above. With this BYOD model, the employee-owned devices are on-boarded (registered with the Cisco ISE server and provisioned with digital certificates) and terminated on the inside of the corporate network. This requires no modifications to the Internet Edge firewall, and potentially fewer security concerns.

Cisco Jabber Clients and the Cisco BYOD Infrastructure

Cisco Jabber, a Cisco mobile client application, provides core Unified Communications and collaboration capabilities, including voice, video, and instant messaging to users of mobile devices such as Android and Apple iOS smartphones and tablets. When a Cisco Jabber client device is attached to the corporate wireless LAN, the client can be deployed within the Cisco Bring Your Own Device (BYOD) infrastructure.

Because Cisco Jabber clients rely on enterprise wireless LAN connectivity or remote secure attachment through VPN, they can be deployed within the Cisco Unified Access network and can utilize the identification, security, and policy features and functions delivered by the BYOD infrastructure.

The Cisco BYOD infrastructure provides a range of access use cases or scenarios to address various device ownership and access requirements. The following high-level access use case models should be considered:

- Enhanced Access—This comprehensive use case provides network access for corporate, personal, and contractor/partner devices. It allows a business to build a policy that enables granular role-based application access and extends the security framework on and off-premises.
- Advanced Access —This use case introduces MDM integration with Enhanced Access.
- Limited Access—Enables access exclusively to corporate issued devices.
- Basic Access—This use case is an extension of traditional wireless guest access. It represents an alternative where the business policy is to not on-board/register employee wireless personal devices, but still provides Internet-only or partial access to the network.

Use Case Impact on Jabber

The Enhanced use case allows the simplest path for implementing a Cisco Jabber solution. Cisco Jabber clients, whether running on corporate or personal devices, require access to numerous back-end, on-premise enterprise application components for full functionality. The Enhanced Access use case will allow access from corporate devices with the option of allowing access from personal devices for Jabber back-end applications.

The Limited Access use case will allow Jabber use only from corporate devices.

Basic Access adds a significant layer of complexity for personal devices, requiring them to have access to back-end on-premise Jabber applications from the DMZ. Various signal, control, and media paths must be allowed through the firewall for full functionality.

In the case of cloud-based collaboration services, Cisco mobile clients and devices connect directly to the cloud through the Internet without the need for VPN or full enterprise network attachment. In these scenarios, user and mobile devices can be deployed using the Basic Access model because these use cases require only Internet access.

Other Jabber Design Considerations

When deploying Cisco Jabber clients within the Cisco BYOD infrastructure, consider the following high-level design and deployment guidelines:

- The network administrator should strongly consider allowing voice- and video-capable clients to attach to the enterprise network in the background (after initial provisioning), without user intervention, to ensure maximum use of the enterprises telephony infrastructure. Specifically, use of certificate-based identity and authentication helps facilitate an excellent user experience by minimizing network connection and authentication delay.
- In scenarios where Cisco Jabber clients are able to connect remotely to the enterprise network through a secure VPN:
  - The network administrator should weigh the corporate security policy against the need for seamless secure connectivity without user intervention to maximize utilization of the enterprise telephony infrastructure. The use of certificate-based authentication and enforcement of a device PIN lock policy provides seamless attachment without user intervention and functionality similar to two-factor authentication because the end user must possess the device and know the PIN lock to access the network. If two-factor authentication is mandated, then user intervention will be required in order for the device to attach remotely to the enterprise.
  - It is important for the infrastructure firewall configuration to allow all required client application network traffic to access the enterprise network. Failure to open access to appropriate ports and protocols at the corporate firewall could result in an inability of Cisco Jabber clients to register to on-premises Cisco call control for voice and video telephony services and/or the loss of other client features such as enterprise directory access or enterprise visual voicemail.
- When enterprise collaboration applications such as Cisco Jabber are installed on employee-owned mobile devices, if the enterprise security policy requires the device to be wiped or reset to factory default settings under certain conditions, device owners should be made aware of the policy and encouraged to backup personal data from their device regularly.
- When deploying Cisco Jabber, it is important for the underlying network infrastructure to support, end-to-end , the necessary QoS classes of service, including priority queuing for voice media and dedicated video and signaling bandwidth, to ensure the quality of client application voice and video calls and appropriate behavior of all features.

For further information regarding Cisco Jabber clients, see the product collateral and documentation at: http://www.cisco.com/go/jabber.

For further information regarding Cisco Mobile Unified Communications, see the Cisco Unified Communications System 9.X SRND at: http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/9x/mobilapp.html.

Cisco Virtual Workspace (VXI) Smart Solution

The Cisco Virtual Workspace (VXI) Smart Solution provides an optimized end-to-end infrastructure for desktop virtualization deployments.

Cisco Virtual Workspace (VXI) Architecture

The Cisco Virtual Workspace (VXI) architecture consists of three fundamental building blocks: Cisco Virtualized Data Center, Virtualization-Aware Network, and Virtualized Collaborative Workspace.

Cisco's Virtualized Data Center provides the computing, switching, storage, and virtualization capabilities needed to support a hosted virtual desktop solution from Citrix.

Cisco's Virtualization-Aware Network connects data centers, enterprise campuses, branch offices, and remote workers to help ensure that traffic flowing between end users and their hosted desktops is transported securely, reliably, and efficiently. Virtualization-Aware Networks employ bandwidth optimization, load balancing, quality of service (QoS), security, and other technologies from Cisco's industry-leading portfolio.

Cisco's Virtualized Collaborative Workspace builds on the Cisco Collaboration architecture, extending the reach of the virtual desktop to a wide range of endpoints while supporting critical collaboration capabilities hosted in the data center. Endpoints can be zero clients, thin clients, mobile devices, or thick clients.

Cisco Virtual Workspace (VXI) Smart Solution also supports management tools for both Cisco and ecosystem partner products, as well as a rich services portfolio that helps enterprises make the most of their virtualization investments.

Cisco Virtual Workspace (VXI) Application Virtualization and Citrix

This Cisco Virtual Workspace (VXI) Smart Solution validated design is based on Citrix XenDesktop and XenApp virtualization solutions.

Citrix XenDesktop is a desktop virtualization solution that delivers Windows desktops as an on-demand service to users on any device, anytime, with a high definition user experience.

Citrix XenApp, which is included as part of the XenDesktop license, is an on-demand application delivery solution that enables Windows applications to be virtualized, centralized, and instantly delivered as a service to users anywhere on any device.

For more information about the Cisco Virtual Workspace with Citrix, refer to the Cisco Virtual Workspace (VXI) Smart Solution CVDs at:
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns1100/landing_vxi.html.

License Requirements for BYOD Solution

Cisco ISE comes with several license options, such as Evaluation, Base, Advanced, and Wireless. For this design to be implemented, ISE requires the Advanced license option. To obtain more information on licensing, see:
http://www.cisco.com/en/US/products/ps11640/tsd_products_support_series_home.html.



PART 3

Configuring the Infrastructure



Summary of Configuring the Infrastructure

Revised: August 7, 2013

This part of the CVD section discusses the different infrastructure components that are critical to the deployment of the BYOD design and the configuration steps used for this design guide.

There are numerous ways to enable a BYOD solution based on the unique business requirements of a specific organization. While some organizations may take a more open approach and rely on basic authentication, other organizations will prefer more secure ways to identify, authenticate, and authorize devices. A robust network infrastructure with the capabilities to manage and enforce these policies is critical to a successful BYOD deployment

The following components and the configuration steps are discussed to support different BYOD use cases:

- Wireless Controllers (Unified and Converged Access)
- Access Layer Switches
- Identity Service Engine
- Certification Authority (CA) server
- Integration with Mobile Device Managers

This part of the CVD includes the following chapters:

- BYOD Wireless Infrastructure Design—This section presents different network designs used to support BYOD, including Campus and Branch designs. This section presents both Unified Wireless and Converged Access designs with single or dual SSID configurations.
- Identity Services Engine for BYOD—The Cisco Identity Services Engine plays a critical role in enabling the BYOD model and allows for enforcement of centrally-configured policies across wired and wireless networks. The section focuses on digital certificates, authentication and authorization policies, device profiling, and different ways to on-board devices with either single or dual SSID configurations.
- BYOD Wired Infrastructure Design—This section highlights how to on-board wired devices and how to enforce BYOD policies and network access for wired devices. This section has details for both campus and branch deployments.
- Security Group Access for BYOD—This section presents two different deployment scenarios that rely on Security Group Tags to enforce BYOD policies. These scenarios are not mutually exclusive and may be used together to implement different business use cases.

- Mobile Device Manager Integration for BYOD—This section focuses on how to configure ISE to integrate with third party MDM products through an XML-based API. BYOD Advanced Use Case—Mobile Device Manager Integration expands this configuration to receive device posture information from the MDM.



BYOD Wireless Infrastructure Design

Revised: August 7, 2013

The Cisco Wireless LAN Controller (WLC) is used to automate wireless configuration and management functions and to provide visibility and control of the wireless networks. The WLC is able to interact with the Identity Service Engine to enforce authentication and authorization policies across endpoints.

While designing WLAN networks, the following should be considered:

- The role of the WLAN
- The authentication mechanism for the WLAN
- The number of WLANs present in a network

This design guide logically separates the WLAN into distinct logical functions: device provisioning and secure network access. These two functions can be provided by two different WLANs or combined into a single WLAN. This design guide covers both single and dual SSID deployment models for both the branch and the campus locations. Note that in this design guide wireless guest access is implemented on a different WLAN.

Some considerations when selecting a single versus dual SSID configuration:

- Some organizations prefer having a dedicated SSID for on-boarding devices.
- Others see dual SSID as an extra management burden.
- A second SSID adds channel overhead.
- Enabling too many SSIDs may degrade wireless performance.

The organization's unique requirements and preferences will dictate which model to deploy. The configurations of both the ISE and WLC may be easily modified to support either single or dual SSID deployments.

Campus—Unified Wireless LAN Design

As mentioned in Centralized (Local Mode) Wireless Design in Chapter 5, "Campus and Branch Network Design for BYOD," the two wireless LAN designs for the campus which are discussed within this design guide are Centralized (Local Mode) and Converged Access designs. Clients connecting from the campus wireless infrastructure are served by a dedicated cluster of CT5508 Unified Controllers configured in local mode (central switching) or served by a combination of Catalyst 3850 series switches which provide the Mobility Agent (MA) function, while CT5760 wireless controllers provide the Mobility

Controller function. This section discusses the Unified Wireless LAN Design, while discussion on Converged Access follows. The wireless controllers are configured with the proper SSIDs to provide device on-boarding and secure access. This functionality may be provided via single or dual SSIDs.

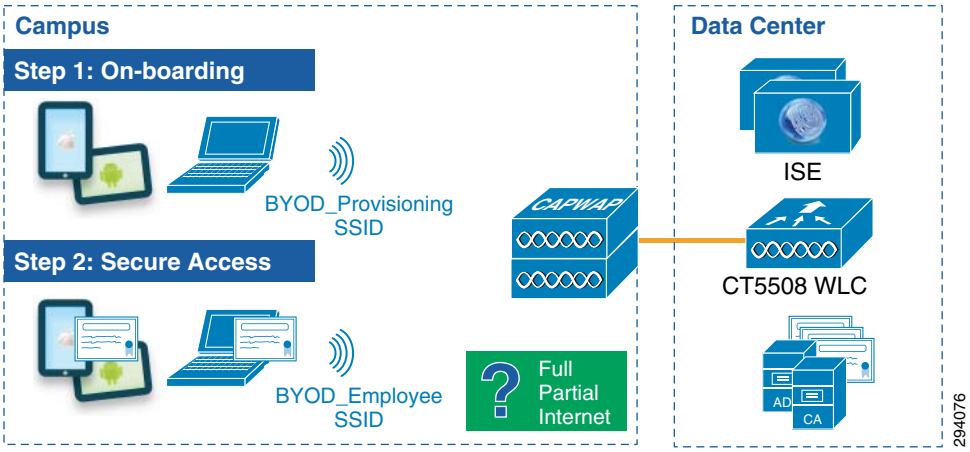
Note The CT5760 wireless controller can also be configured to function as a centralized (Local Mode) wireless controller. As discussed in Campus Migration Path of Chapter 5, "Campus and Branch Network Design for BYOD," this may be a necessary step in migrating from an existing wireless overlay design to a converged access design.

Centralized Campus—Dual SSID Design

In this design there are two SSIDs: one provides enrollment/provisioning and the other provides secure network access. After connecting to the BYOD_Provisioning SSID and completing the enrollment and provisioning steps, the user connects to the BYOD_Employee SSID, which provides network access over a secure EAP-TLS connection.

Figure 9-1 shows the dual SSID design for the campus APs.

Figure 9-1 Campus-Dual SSIDs



In a dual SSID design, there are some additional considerations:

- The provisioning SSID can be either open or password protected. When the provisioning SSID is open, any user can connect to the SSID, whereas if it is password protected, then only users that have credentials, such as AD group membership, are allowed to connect to the SSID. In this design guide, the provisioning SSID is configured to be open and its only purpose is to provide on-boarding services.
- After the device is provisioned, it is assumed that the user will switch to the second SSID for regular network access. To prevent the user from staying connected to the provisioning SSID, an access list that provides only access to ISE, DHCP, and DNS must be enforced on the provisioning SSID. The details of the ACL_Provisioning_Redirect ACL are shown below.
- This design guide makes use of the following SSIDs: BYOD_Provisioning and BYOD_Employee.

The properties of these two SSIDs are highlighted in Table 9-1.

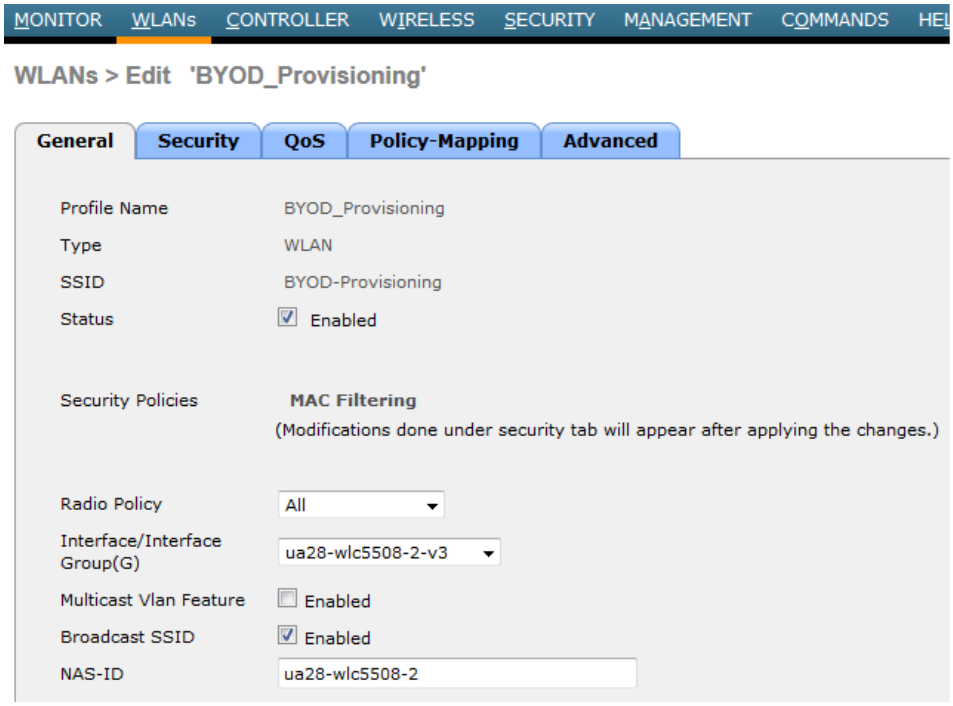
Table 9-1 WLAN Parameters

Attribute	BYOD_Provisioning	BYOD_Employee
Description	Used only for device provisioning	For employees that have completed the on-boarding process
Layer 2 Security	None (for Open SSID)	WPA+WPA2
MAC Filtering	Enabled (for Open SSID)	Disabled
WPA+WPA2 Parameters	None	WPA2 Policy, AES, 802.1X
Layer 3 Security	None	None
AAA Server	Select ISE	Select ISE
Advanced	AAA Override Enabled	AAA Override Enabled
Advanced	NAC State-RADIUS NAC	NAC State-RADIUS NAC
Quality of Service	Best Effort	Platinum
AVC	None	Enabled

To create a WLAN, click **WLANs > Create New > Go** and provide the SSID and profile details. Starting with Figure 9-2 the general configuration steps of the BYOD_Provisioning SSID are highlighted. The steps to configure the BYOD_Employee WLAN are similar, following the settings in Table 9-1.

Note When implementing BYOD solutions using more than one Wireless LAN Controller, WLAN IDs must be kept consistent. WLAN ID is used by ISE in determining which WLAN (SSID) clients are using to connect to the network. Ensuring each WLAN has the same WLAN ID on each WLC is essential for proper operation and security.

Figure 9-2 Creating the BYOD_Provisioning SSID



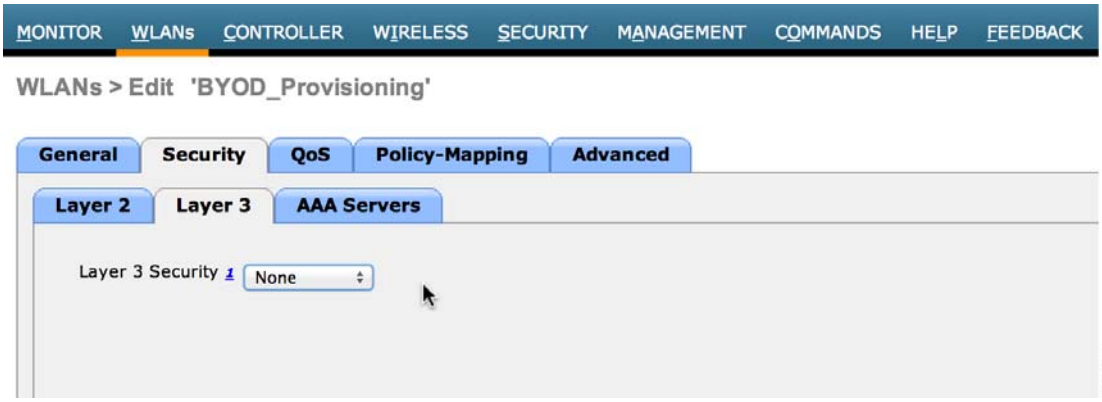
The Layer 2 security settings are configured as **None** since BYOD_PROVISIONING is an open SSID. If the provisioning SSID has to be password-protected, then the Layer 2 security settings must be configured as WPA+WPA2 Enterprise.

Figure 9-3 Layer 2 Security Settings



The Layer 3 Security is configured as **None**, as shown in Figure 9-4.

Figure 9-4 Layer 3 Security Settings



The main configuration in the security settings is to specify the RADIUS server configuration details. Figure 9-5 shows how the ISE's IP address is configured for Authentication and Authorization.

Figure 9-5 AAA Security Settings

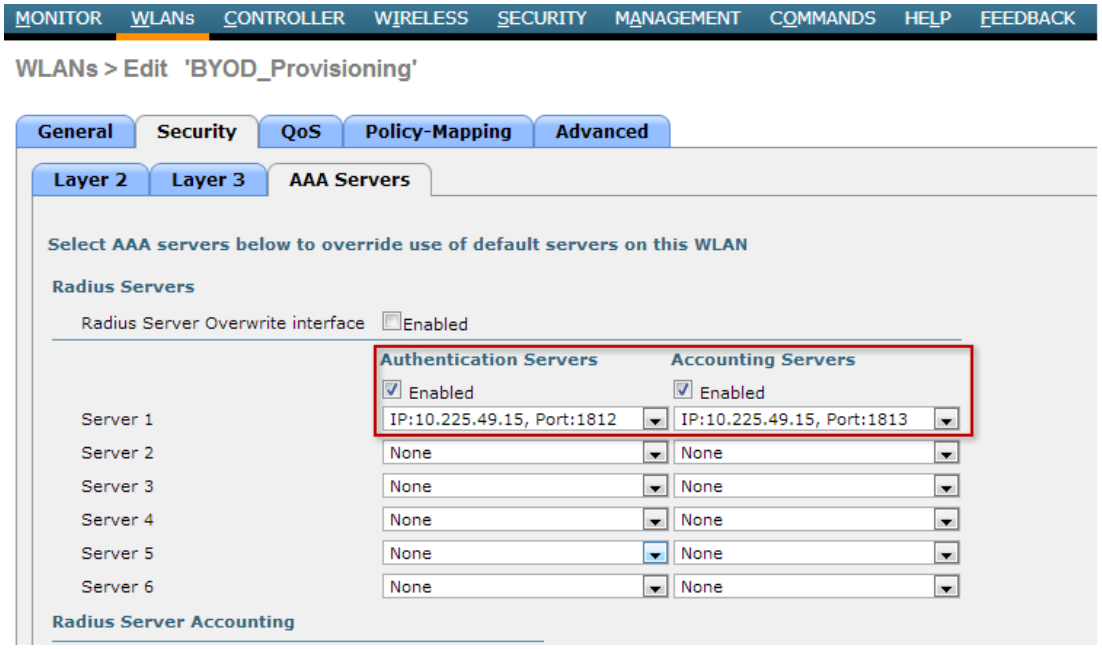
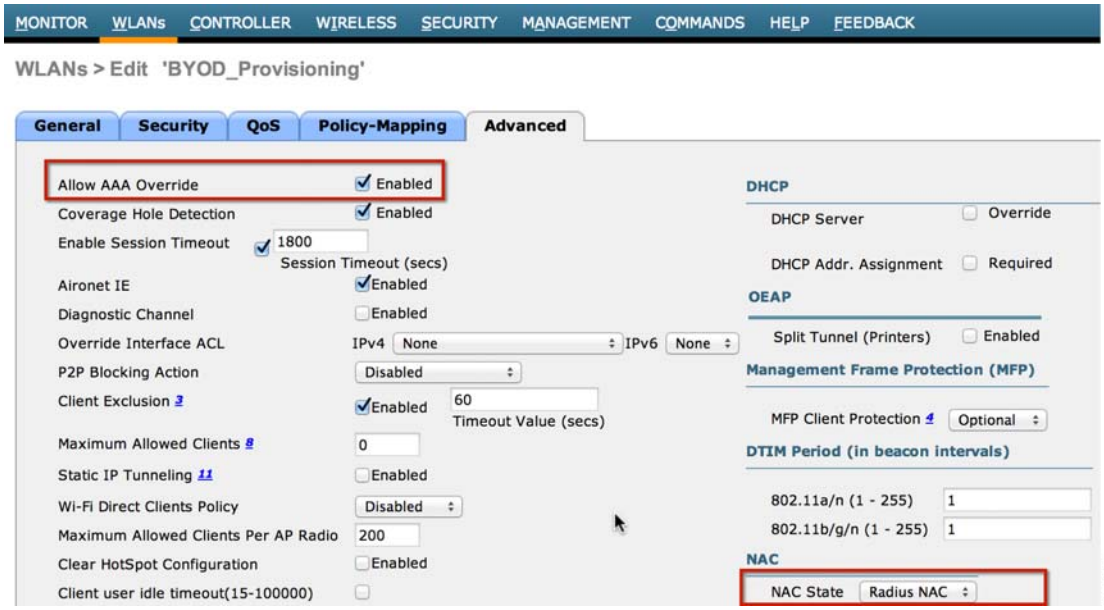


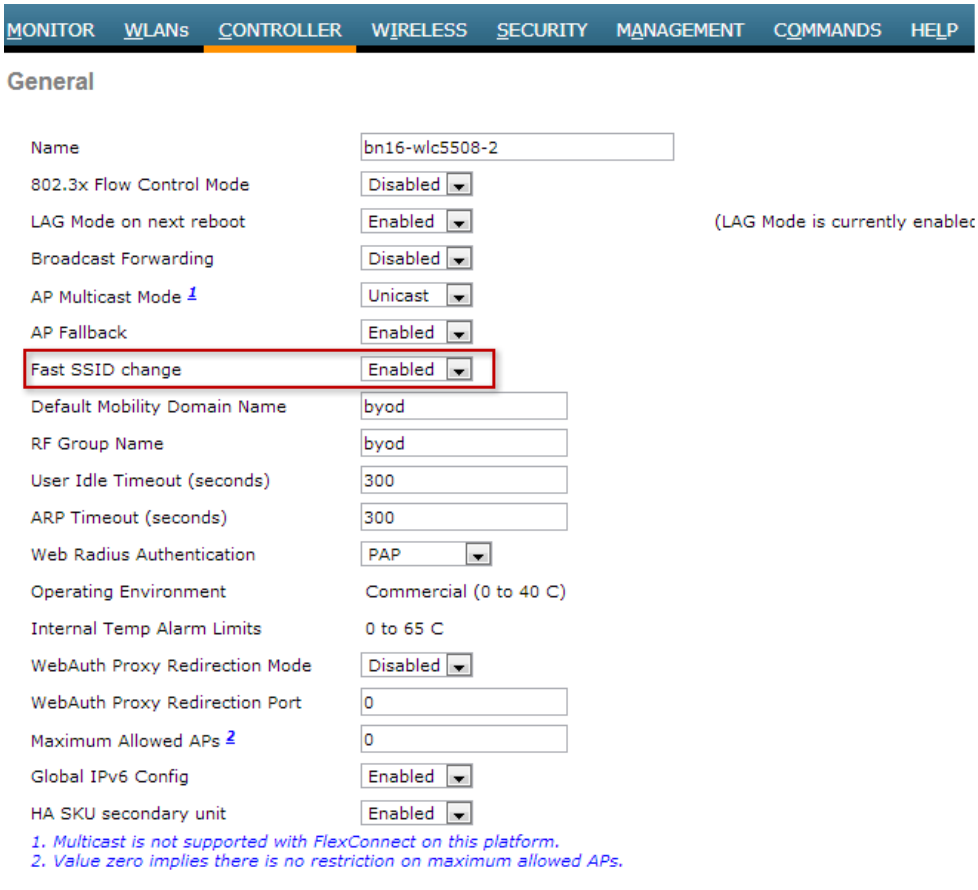
Figure 9-6 shows the advanced settings, including AAA Override and NAC State.

Figure 9-6 Advanced Settings



The Fast SSID Change feature is useful when a device needs to switch from one SSID to another. This applies to the dual SSID BYOD design. After the user completes registration with BYOD_Provisioning, the user is switched to BYOD_Employee SSID. By enabling the FAST SSID Change feature, the user switches immediately to the new SSID without experiencing delays. To enable Fast SSID Change, click **Controller > General > Fast SSID change**, as shown in Figure 9-7.

Figure 9-7 Fast SSID Change

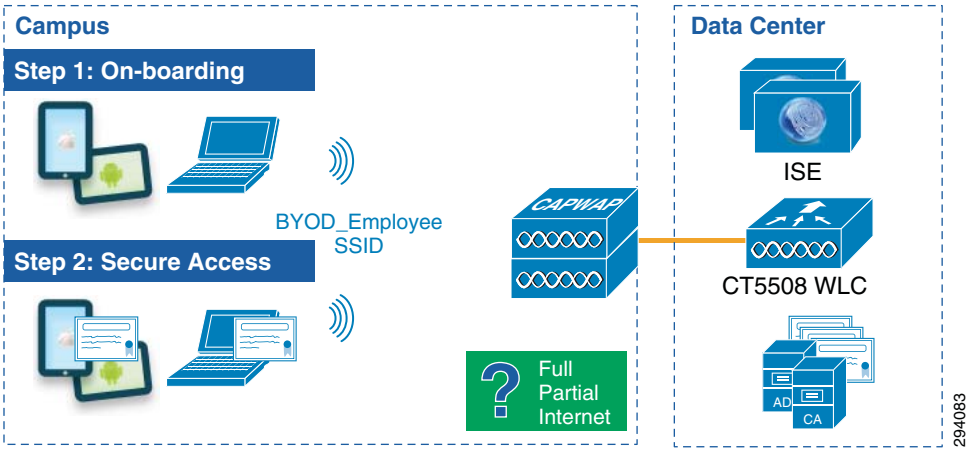


Note Authorization Policies and Profiles in Chapter 10, "Identity Services Engine for BYOD" shows the ACLs and authorization profiles used for dual and single SSID provisioning.

Centralized Campus—Single SSID Design

In a single SSID design the same WLAN (BYOD_Employee) is used for on-boarding and secure network access. Figure 9-8 shows how this design may be implemented using the 5508 Wireless LAN Controller. In this case, the controllers are dedicated to manage the APs in the campus.

Figure 9-8 Campus—Single SSID



Note Authorization Policies and Profiles in Chapter 10, "Identity Services Engine for BYOD" shows the ACLs and authorization profiles used for dual and single SSID provisioning.

Centralized Campus—Policy Enforcement using TrustSec

As discussed in ACL Complexity and Considerations in Chapter 5, "Campus and Branch Network Design for BYOD," past versions of the CVD utilized Named ACLs pre-configured on the wireless controllers to enforce role-based policies for access to network and Data Center resources. This CVD introduces a complimentary technology known as TrustSec and, more specifically, Security Group Access (SGA) to enforce role-based policies through the use of Security Group Tags (SGT) to control access to data center resources. This CVD discusses an approach to slowly migrate to the use of SGT as opposed to, or even in addition to, the use of ACLs through Network Device definitions created in ISE.

Branch—Unified Wireless LAN Design

FlexConnect Wireless LAN Design

In this design guide, endpoints connecting from branch locations are managed by a cluster of Flex 7500 Wireless LAN Controllers or Virtual Wireless LAN Controllers (vWLCs). The vWLC is software which can run on industry standard virtualization infrastructure and is more suitable for small- and medium-sized businesses.

The configuration parameters described in this section apply to both the vWLC and Flex 7500 controllers.

The following link provides more information on how to set up vWLCs using VMware: http://www.cisco.com/en/US/customer/products/ps12723/products_tech_note09186a0080bd2d04.shtml.

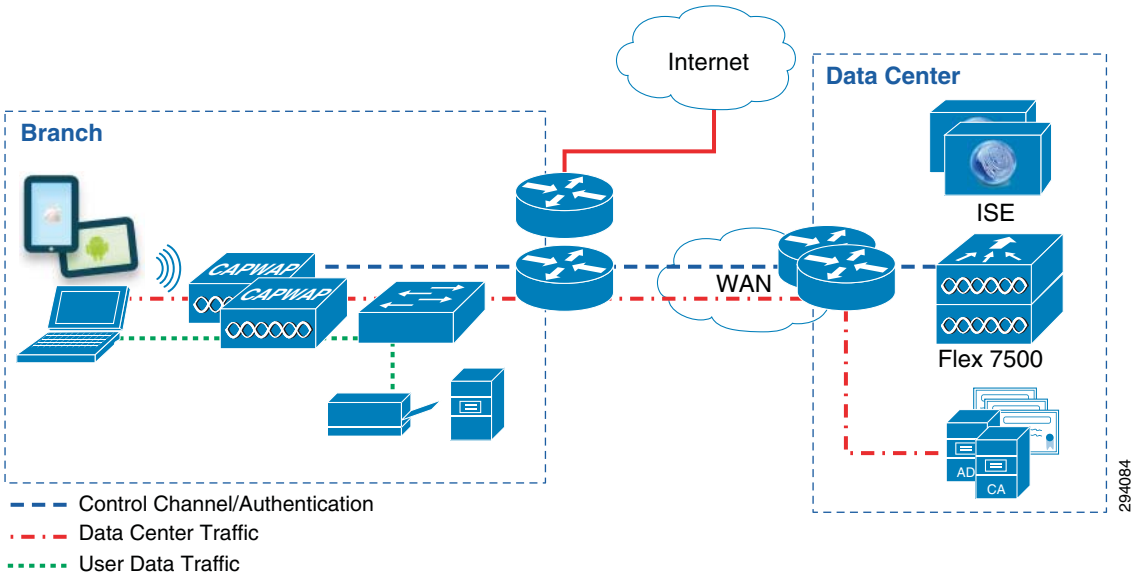
FlexConnect (previously known as Hybrid Remote Edge Access Point or H-REAP) is a wireless solution for branch office and remote office deployments. It enables customers to configure and control access points in a branch or remote office from the corporate office through a wide area network (WAN) link without deploying a controller in each office. The FlexConnect access points can switch client data traffic locally and perform client authentication locally when their connection to the controller is lost.

Distributing client data traffic using the FlexConnect architecture offers some advantages:

- A controller is not required at each branch location.
- Mobility resiliency within branch during WAN link failures.
- Central management and troubleshooting.

The FlexConnect architecture in Figure 9-9 shows different traffic flows originating at the branch.

Figure 9-9 FlexConnect Architecture



When an endpoint associates to a FlexConnect access point, the access point sends all authentication messages to the controller and either switches the data packets locally (locally switched) or sends them to the controller (centrally switched), depending on the WLAN configuration.

With respect to data packet flows, the WLAN can be in any one of the following modes:

- Central switching—Central switched WLANs tunnel both the wireless user traffic and all control traffic to the centralized WLC, where the user traffic is mapped to a dynamic interface or VLAN.
- Local switching—In this mode the FlexConnect access point switches data packets locally by dropping all traffic locally at the wired interface. Wireless user traffic is mapped to discrete VLANs via 802.1Q trunking.

The Flex 7500 Wireless Branch Controller Deployment Guide offers more details: http://www.cisco.com/en/US/products/ps11635/products_tech_note09186a0080b7f141.shtml.

The key strategy for providing differentiated access to users is done by assigning users to different VLANs dynamically. The AAA Override feature for FlexConnect assigns individual clients to specific VLANs, based on the returned RADIUS attributes from the ISE.

The access point must be preconfigured with all of the possible VLANs that can be returned by the ISE server. The VLAN assignment returned by the ISE as part of authorization is applied. If the VLAN that was returned from the ISE is not present on the AP, the client falls back to the default VLAN configured for the WLAN.

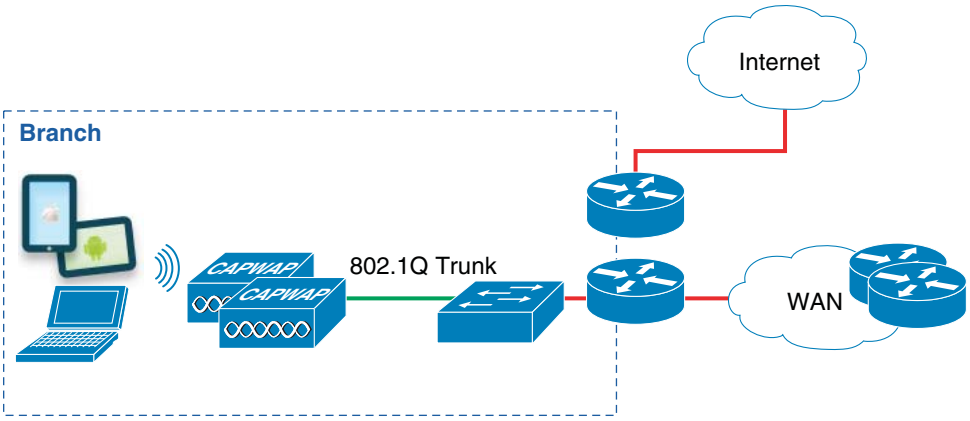
In this design three VLANs have been configured for wireless connectivity on the BYOD_Employee SSID. Table 9-2 illustrates those VLANs and their purpose.

Table 9-2 VLANs and Purpose

VLAN Number	VLAN Name	Description
10	Wireless_Full	Users assigned to this VLAN get full access to campus and branch servers.
11	Wireless_Partial	In addition to Internet access, users assigned to this VLAN access to additional campus and branch resources.
12	Wireless_Internet	Users assigned to this VLAN get only Internet access.
18	AP_Mgmt_Flex	This is the native VLAN that the user will initially be placed into, until the authorization policy determines the appropriate VLAN.

Since more than one VLAN is configured for local switching, FlexConnect APs at the branch must be connected to an 802.1Q trunk link. Both the AP and the upstream switchport need to be configured for 802.1Q trunking. Figure 9-10 shows an example configuration of the access layer switch that connects to the FlexConnect AP.

Figure 9-10 Trunk Configuration



Branch Wireless IP Address Design

Once the device has been dynamically assigned to a VLAN, the endpoint must obtain an IP address from a DHCP server. In the following example the branch router's Layer 3 subinterfaces are configured with the **ip-helper address** command, pointing to a DHCP server:

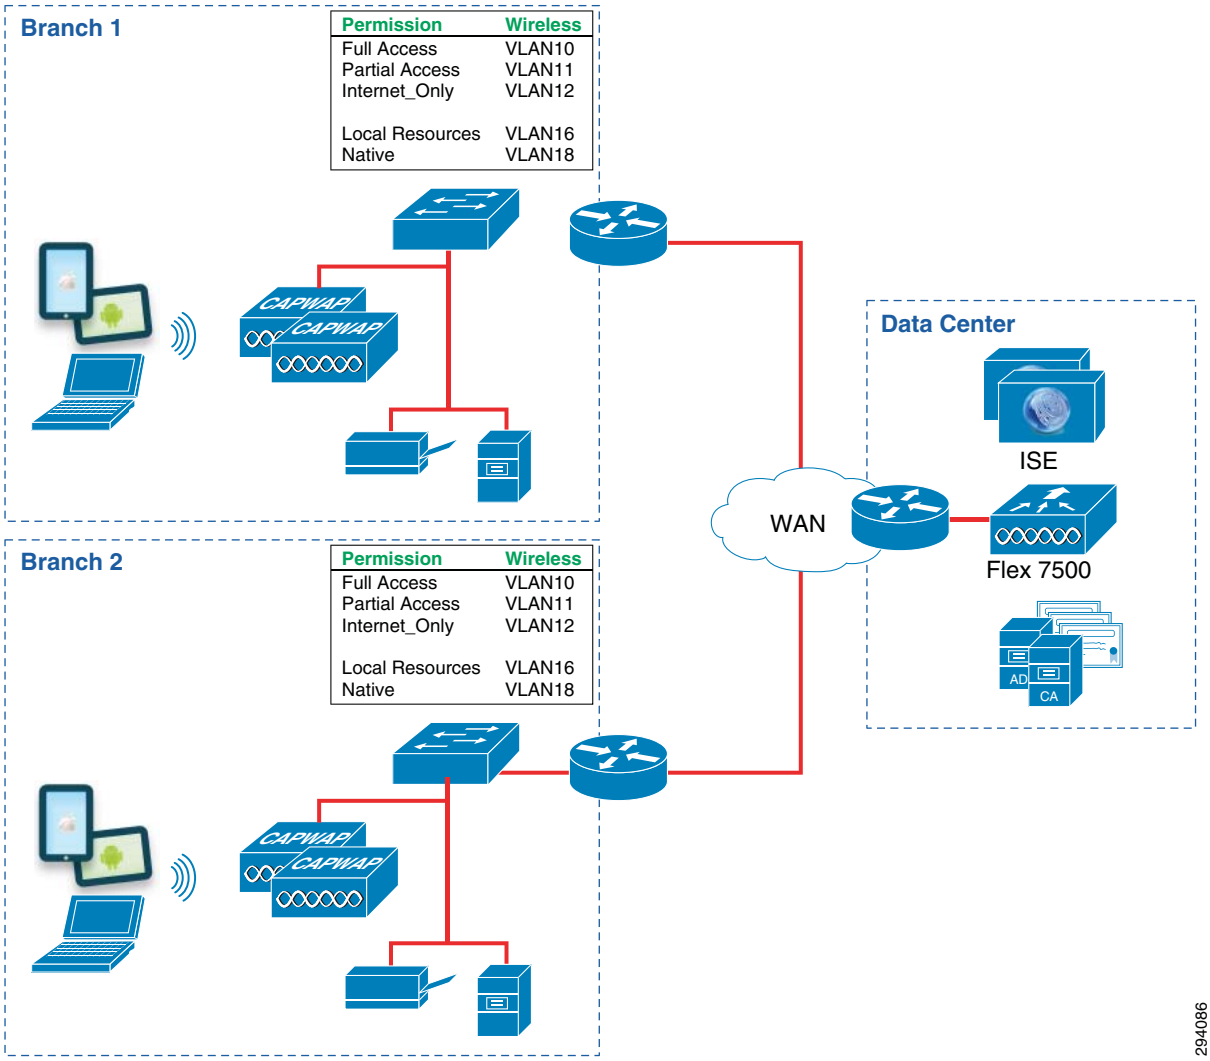
```
interface GigabitEthernet0/1
 description Trunk to branch bn22-3750x-1
 no ip address
 media-type sfp
!
interface GigabitEthernet0/1.10
 encapsulation dot1Q 10
 ip address 10.200.10.2 255.255.255.0
 ip helper-address 10.230.1.61
 standby 10 ip 10.200.10.1
 standby 10 priority 110
 standby 10 preempt
!
interface GigabitEthernet0/1.11
 encapsulation dot1Q 11
 ip address 10.200.11.2 255.255.255.0
 ip helper-address 10.230.1.61
 standby 11 ip 10.200.11.1
 standby 11 priority 110
 standby 11 preempt
!
interface GigabitEthernet0/1.12
 encapsulation dot1Q 12
 ip address 10.200.12.2 255.255.255.0
 ip helper-address 10.230.1.61
 standby 12 ip 10.200.12.1
 standby 12 priority 110
 standby 12 preempt
```

The diagram in Figure 9-11 shows two branch locations utilizing resources from the data center and illustrates the following key points:

- At the branch, endpoints are placed in different VLANs based on the level of access to which they are entitled.
- The wireless infrastructure from the branches is managed by a single cluster of Flex 7500 controllers.
- Endpoints that get assigned to VLAN 10 are granted full access to network resources, VLAN 11 for partial access and VLAN 12 for Internet access.

Based on the matching authorization profile, a user is assigned to a specific VLAN where predefined permissions have been defined.

Figure 9-11 VLANs Used at the Branches

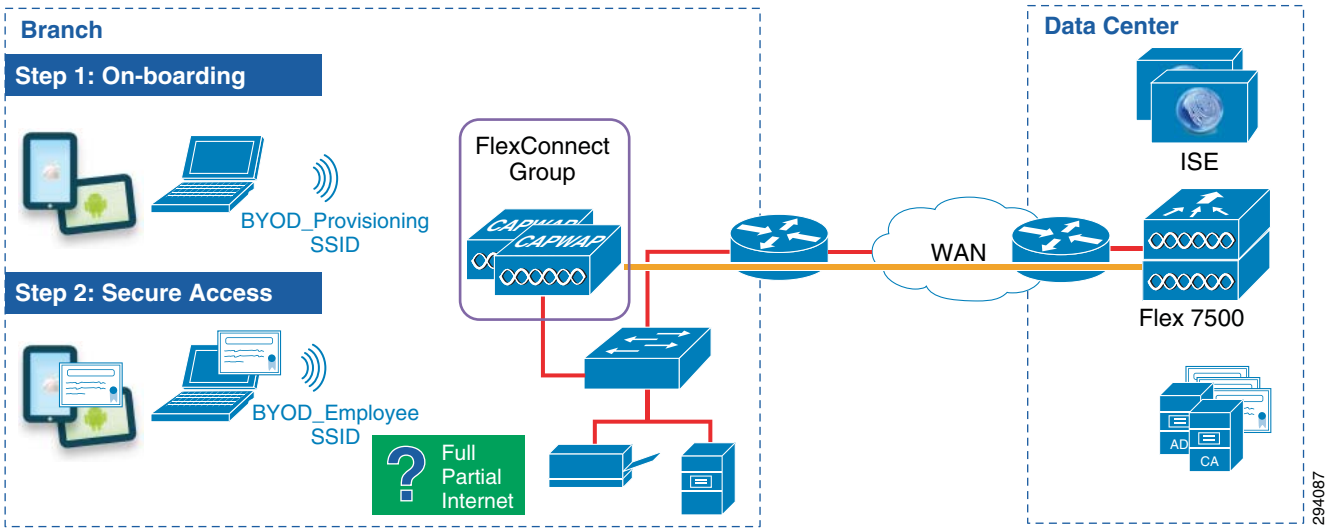


FlexConnect Branch—Dual SSID Design

In the Dual SSID design two SSIDs are configured: one SSID provides enrollment/provisioning while the other provides secure EAP-TLS access. After connecting to the BYOD_Provisioning SSID and completing the enrollment and provisioning steps, the user connects to the BYOD_Employee SSID, which provides secure network access.

Figure 9-12 shows the dual SSID design for the branch APs.

Figure 9-12 Branch-Dual SSIDs



In a dual SSID design, there are some additional considerations:

- The provisioning SSID can be either open or password-protected. When the provisioning SSID is open, any user can connect to the SSID, whereas if it is password protected, then only users that have credentials, such as AD group membership, are allowed to connect to the SSID.
- After the device is provisioned, the user connects via EAP-TLS to the BYOD_Employee SSID for network access. To prevent the user from remaining connected to the provisioning SSID, an access list that provides access only to ISE, DHCP, and DNS must be enforced on the provisioning SSID. The details of this SSID are discussed in the Client Provisioning section.

Table 9-3 shows the WLAN parameters for the SSIDs used in this design guide.

Table 9-3 WLAN Parameters

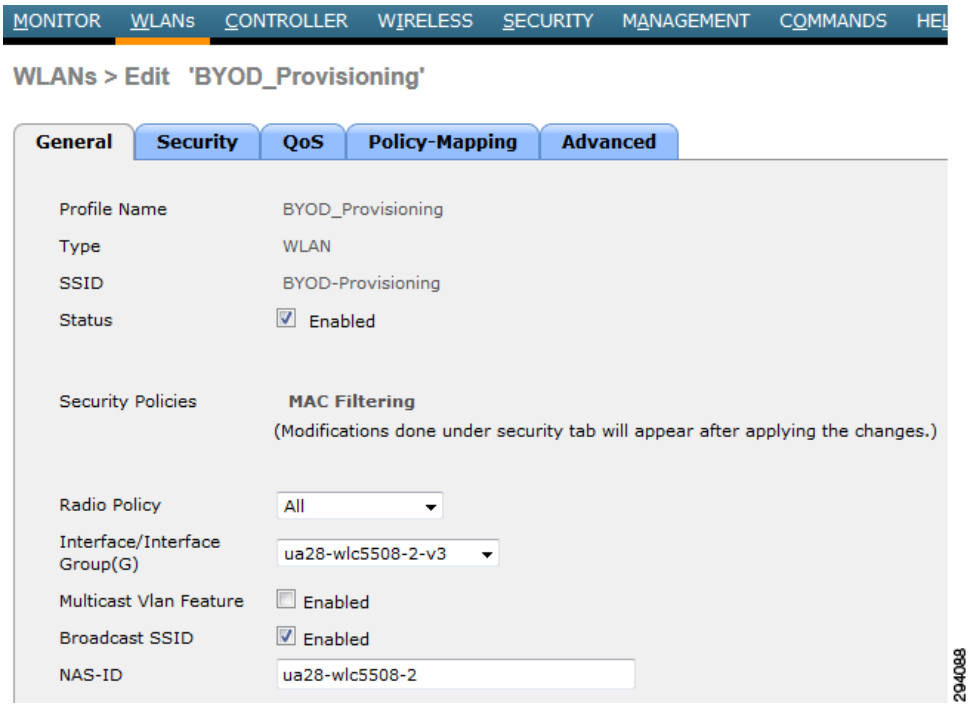
Attribute	BYOD_Provisioning	BYOD_Employee
Description	Used for device provisioning	For employees that have completed the on-boarding process
Layer 2 Security	None (for Open SSID)	WPA+WPA2
MAC Filtering	Enabled (for Open SSID)	Disabled
WPA+WPA2 Parameters	None (for Open SSID)	WPA2 Policy, AES, 802.1X
Layer 3 Security	None	None
AAA Server	Select ISE	Select ISE
Advanced	AAA Override Enabled	AAA Override Enabled
Advanced	NAC State-RADIUS NAC	NAC State-RADIUS NAC
Advanced-FlexConnect Local Switching	Disabled for Central Switching Provisioning Enabled for Local Switching Provisioning	Enabled

Table 9-3 WLAN Parameters

Attribute	BYOD_Provisioning	BYOD_Employee
Quality of Service	Best Effort	Platinum
AVC	Does Not Apply	Does Not Apply

To create a WLAN, click **WLANs > Create New > Go** and provide the SSID and profile details. Figure 9-13 shows the general configuration details of the BYOD_Provisioning SSID.

Figure 9-13 Creating the Branch BYOD_Provisioning SSID



Since BYOD_Provisioning is an open SSID, the Layer 2 security settings in are configured as **None**. If the provisioning SSID had to be password-protected, the Layer 2 security settings would be configured as WPA+WPA2 Enterprise.

Figure 9-14 Layer 2 Security Settings



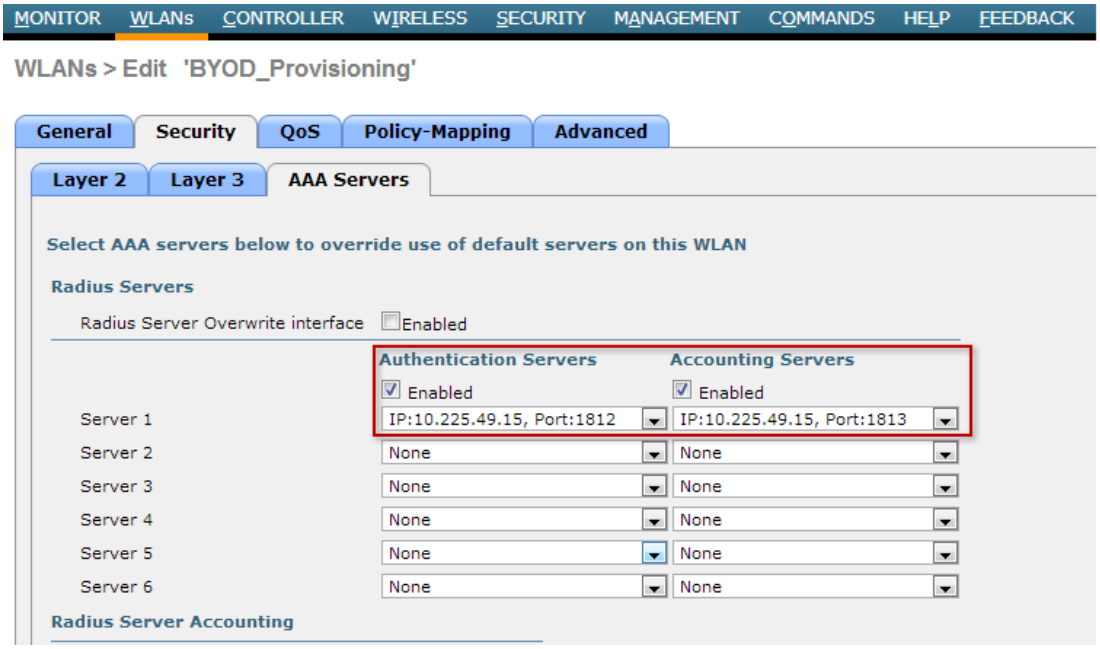
The Layer 3 Security is configured as **None**, as shown in Figure 9-15.

Figure 9-15 Layer 3 Security Settings



Under **Security > AAA servers**, configure the RADIUS server details. Figure 9-16 shows the ISE's IP address configured for Authentication and Authorization.

Figure 9-16 AAA Security Settings



Within the dual SSID deployment there are two possible ways to direct provisioning traffic:

- From the campus or data center—The endpoint receives an IP address from a DHCP scope at the data center and the provisioning traffic is directed through the CAPWAP tunnel between the branch and the Flex 7500 controller.
- At the branch—The endpoint receives an IP address from a DHCP scope at the branch and the provisioning traffic uses the switching and WAN infrastructure for connectivity to data center resources.

Dual SSID—Central Switching Provisioning

Figure 9-17 shows how with central switching provisioning, the endpoint communicates with ISE and data center resources using the CAPWAP tunnel and all traffic is tunneled back to the controller in the data center.

Figure 9-17 Central Switching Provisioning

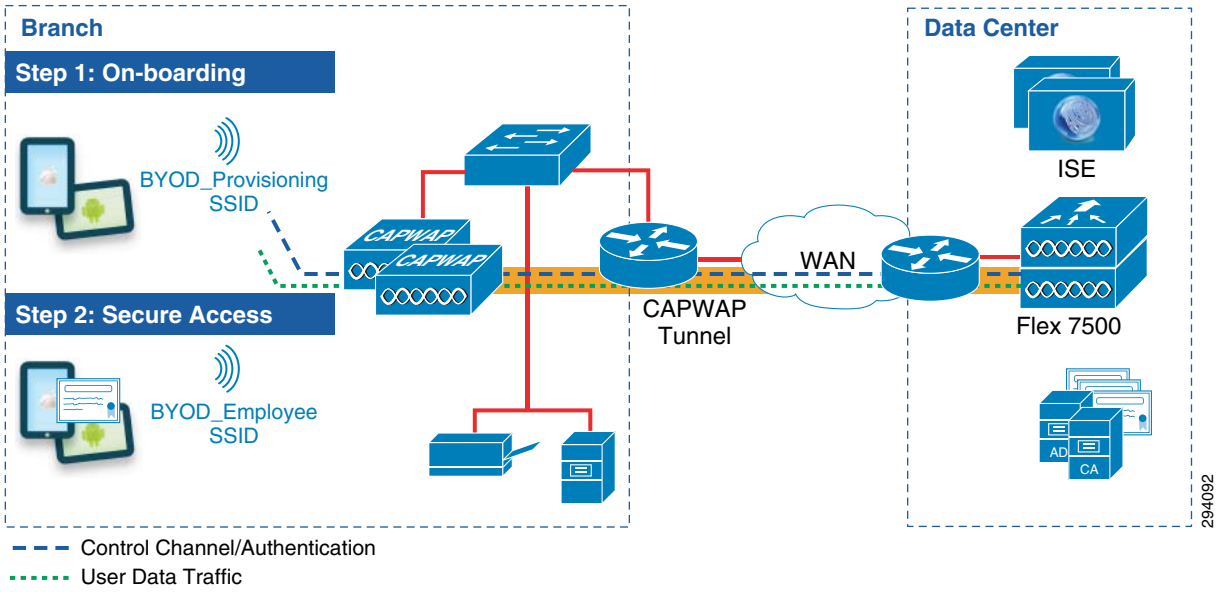
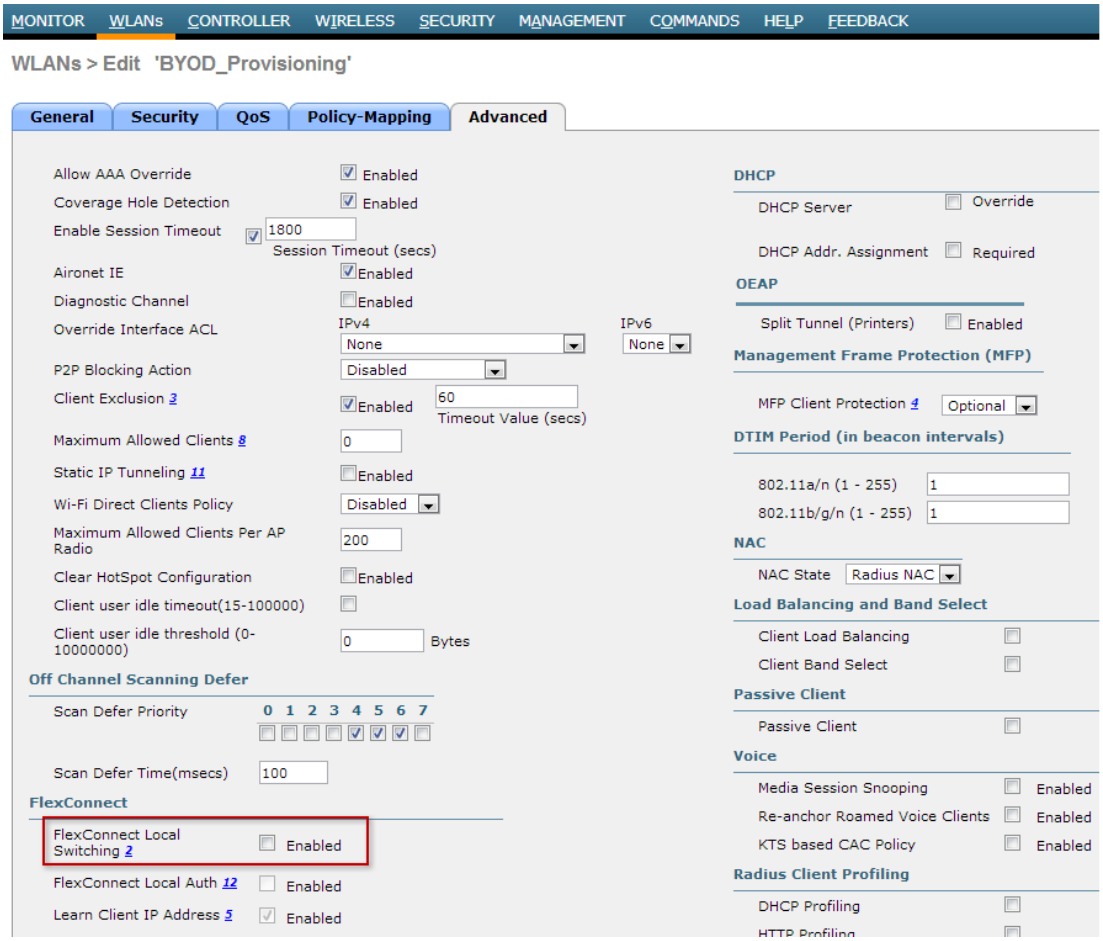


Figure 9-18 shows the advanced settings for BYOD_Provisioning, including the AAA Override and NAC State. The FlexConnect Local Switching setting is disabled for central switching provisioning.

Figure 9-18 Advanced Settings for Central Switching Provisioning



Note Authorization Policies and Profiles in Chapter 10, "Identity Services Engine for BYOD" shows the ACLs and authorization profiles used for dual and single SSID provisioning.

Dual SSID—Local Switching Provisioning

Figure 9-19 shows provisioning with local switching mode. The user data traffic is sent to the switch interface and the endpoint relies on the normal router/WAN infrastructure to reach the ISE and other network resources.

Figure 9-19 Local Switching Provisioning

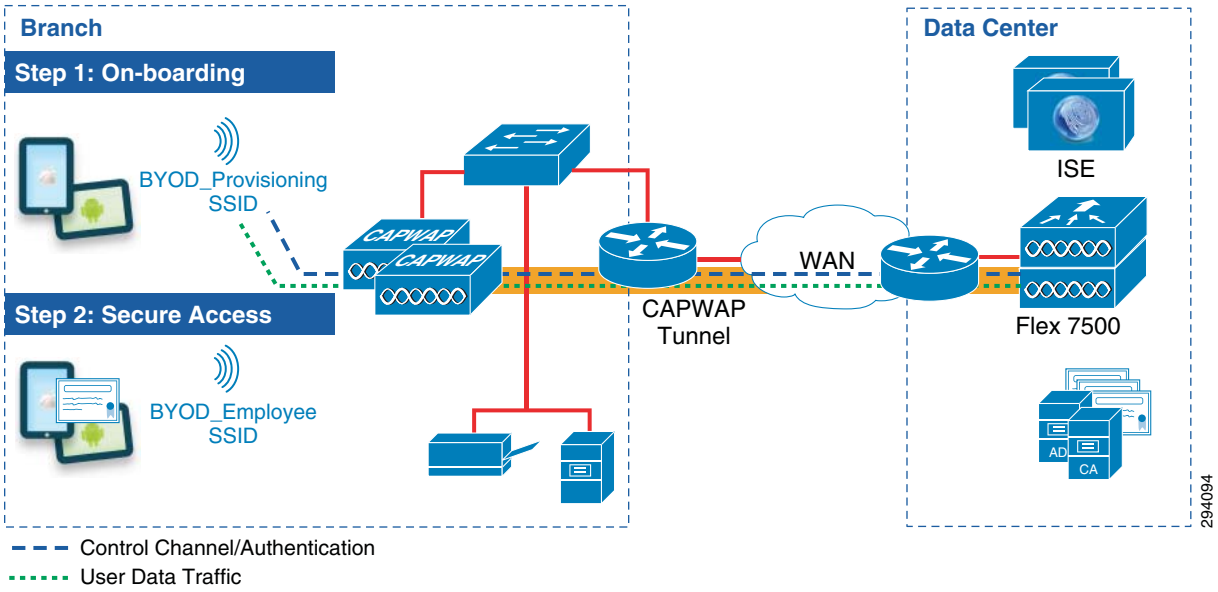
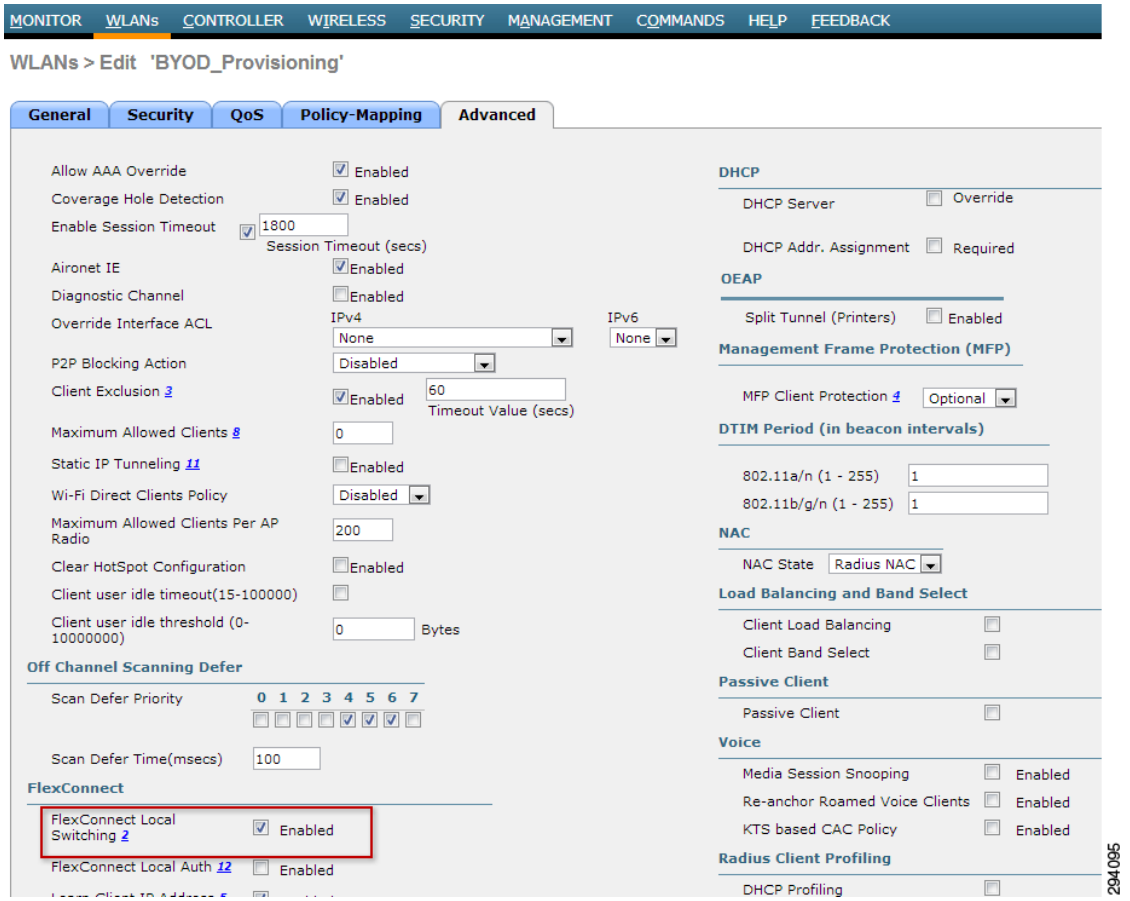


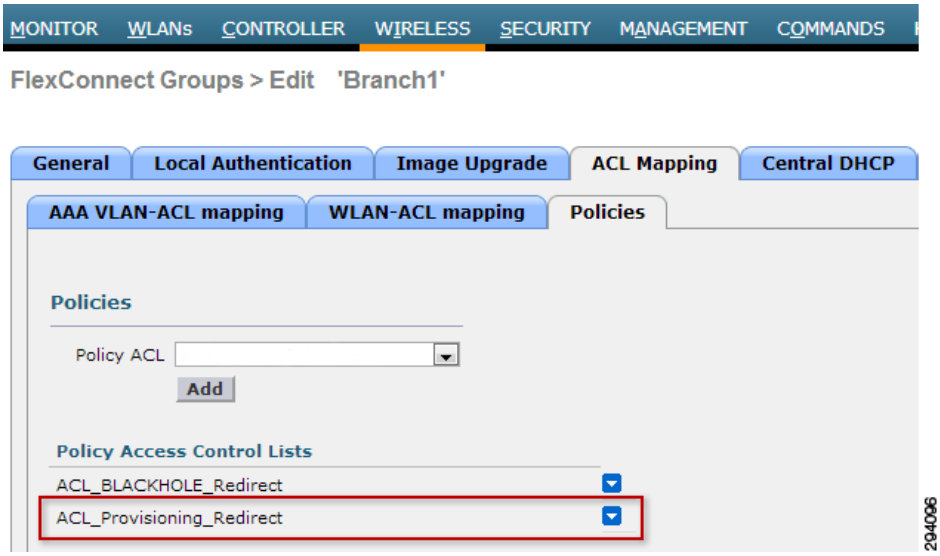
Figure 9-20 shows the advanced settings for BYOD_Provisioning, including the AAA Override and NAC State. The FlexConnect Local Switching is enabled for local switching provisioning.

Figure 9-20 Advanced Settings for Local Switching Provisioning



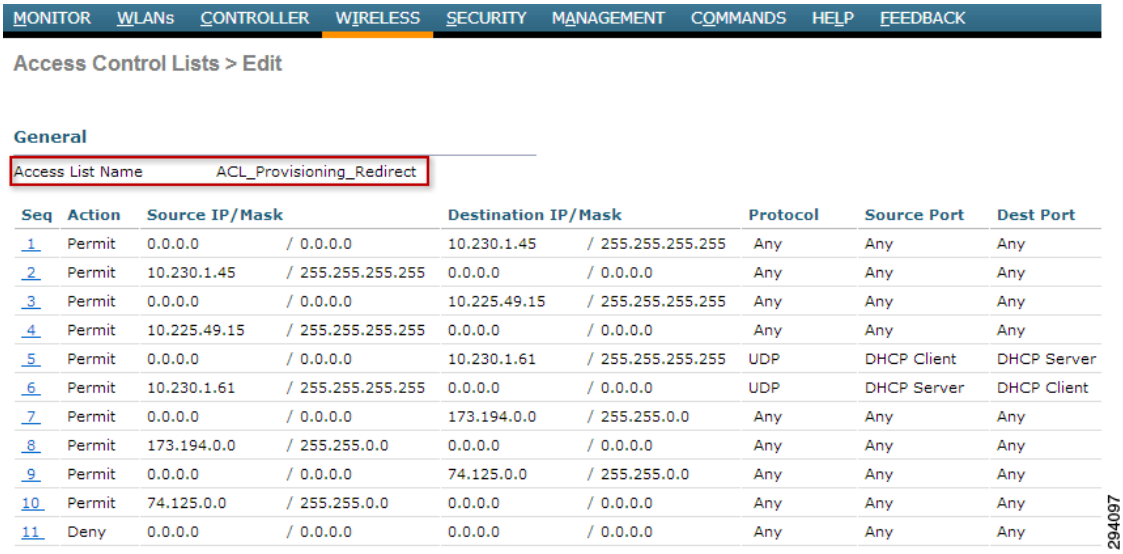
To enforce the redirection to the self-registration portal, a FlexConnect ACL is defined under the Policies tab for the specific FlexConnect group, as shown in Figure 9-21.

Figure 9-21 Policies for FlexConnect Group



The ACL_Provisioning_Redirect FlexConnect ACL shown in Figure 9-22 allows access to ISE, DNS, the Google Play Store, and denies all other traffic. Android devices require access to the Google Play Store to download the SPW package.

Figure 9-22 ACL_Provisioning_Redirect FlexConnect ACL



The ACL_Provisioning_Redirect ACL specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Access to Google Play.

Note The purpose of the ACL shown above is to provide an example that network administrators can use to deploy in the network. The Google and Apple app stores may change their addresses, so it is advisable to validate those addresses before deploying the ACL.

Note ACL_Provisioning_Redirect must redirect all traffic sent to enroll.cisco.com. The Cisco Configuration Assistant for Android devices requires this redirect to discover the IP address of the ISE server.

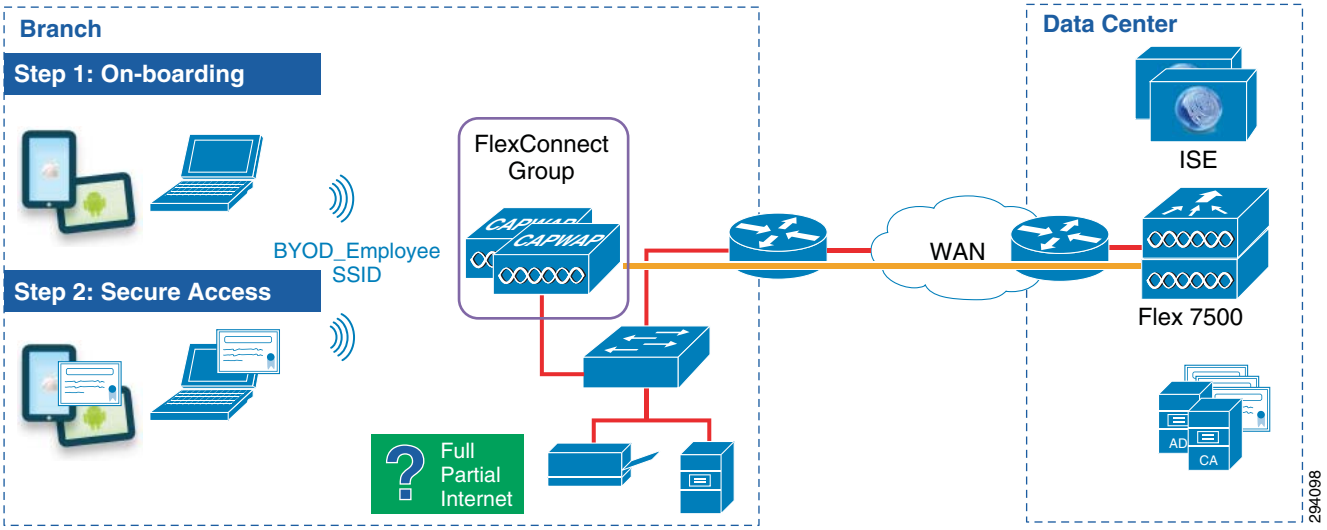
FlexConnect Branch—Single SSID Design

In a single SSID design, the same WLAN is used for certificate enrollment, provisioning (on-boarding process), and secure network access. There are some considerations that should be taken into consideration while deploying a Single SSID solution:

1. Since the authentication method is PEAP, the user is expected to enter the AD credentials before the registration process can begin. In the PEAP protocol, the server presents its identity certificate to the end user. In this design, ISE presents its identity certificate to the endpoint. Some endpoints may reject the certificate if the root certificate is not present in their list of trusted providers. During the registration process, the root CA certificate is installed on the endpoint, but this can't be done if the initial dialog itself fails. Hence, this presents a chicken-and-egg problem. To prevent this from happening the ISE identity certificate must be signed by a third-party trusted provider such as VeriSign.
2. If the above cannot be done, then it is better to deploy dual SSID design.

Figure 9-23 shows how this design uses the BYOD_Employee SSID and is implemented using the Flex 7500 Controller cluster, which is dedicated to manage the APs in the branch locations.

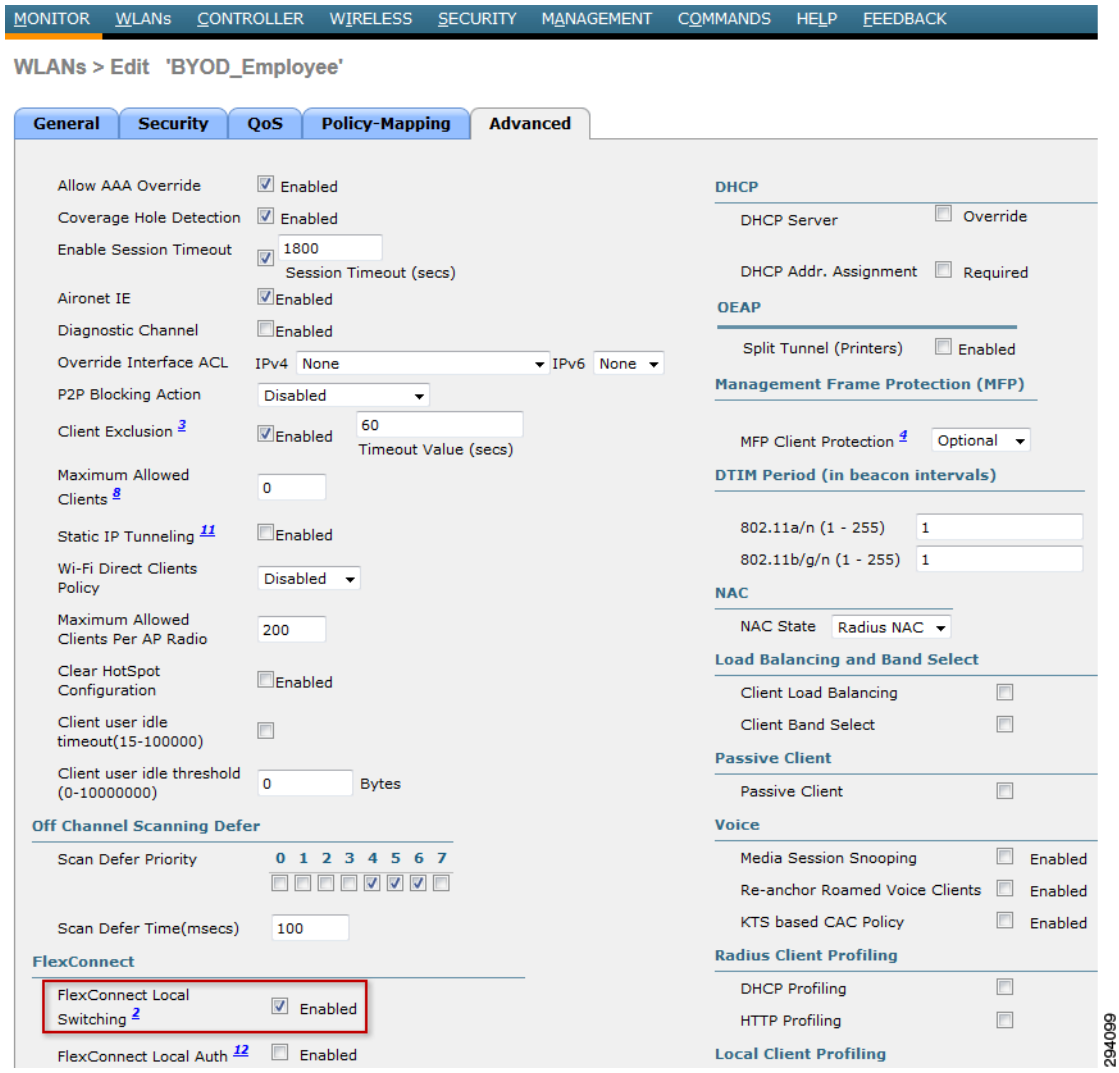
Figure 9-23 Branch-Single SSID



In this scenario the APs associate with the Flex 7500 controller and the FlexConnect capabilities allow the on-boarding and secure access capabilities to be handled by the single BYOD_Employee SSID.

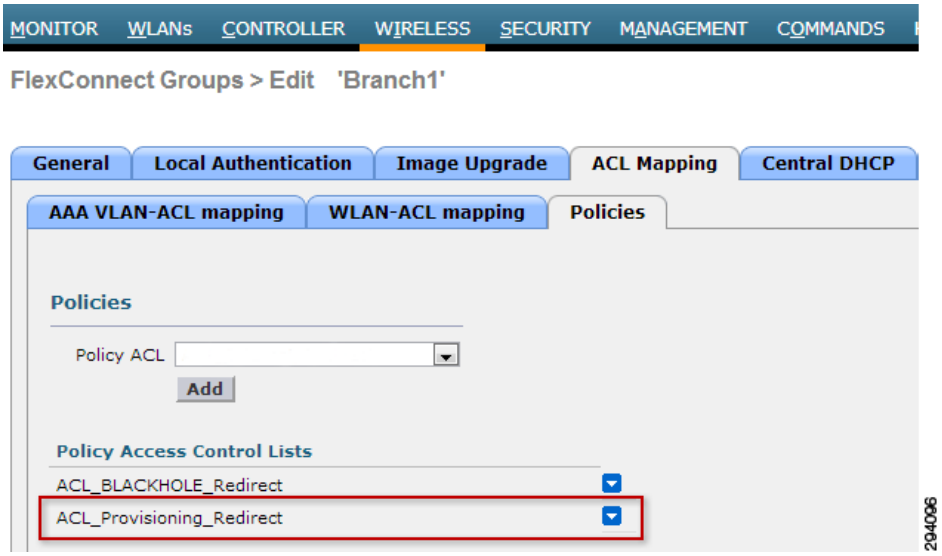
The steps to configure the BYOD_Employee WLAN are similar, following the parameters outlined in Table 9-3. It is important to note that FlexConnect Local Switching is enabled on the BYOD_Employee WLAN, as highlighted in Figure 9-24.

Figure 9-24 FlexConnect Local Switching



To enforce the redirection to the self-registration portal, a FlexConnect ACL is defined under the Policies tab, as shown in Figure 9-25.

Figure 9-25 Policies for FlexConnect Group

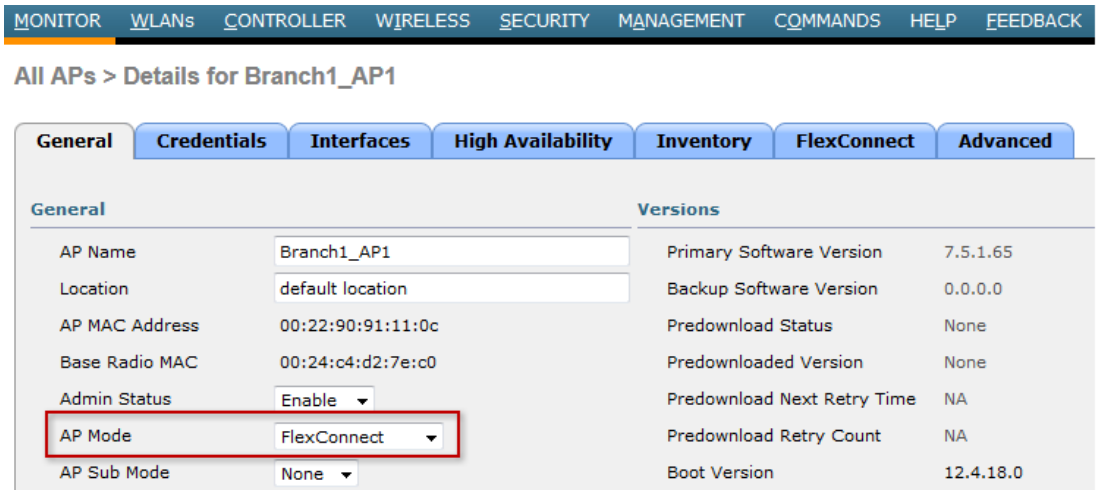


The ACL_Provisioning_Redirect ACL is shown in Figure 9-22 above.

FlexConnect Access Point Configuration

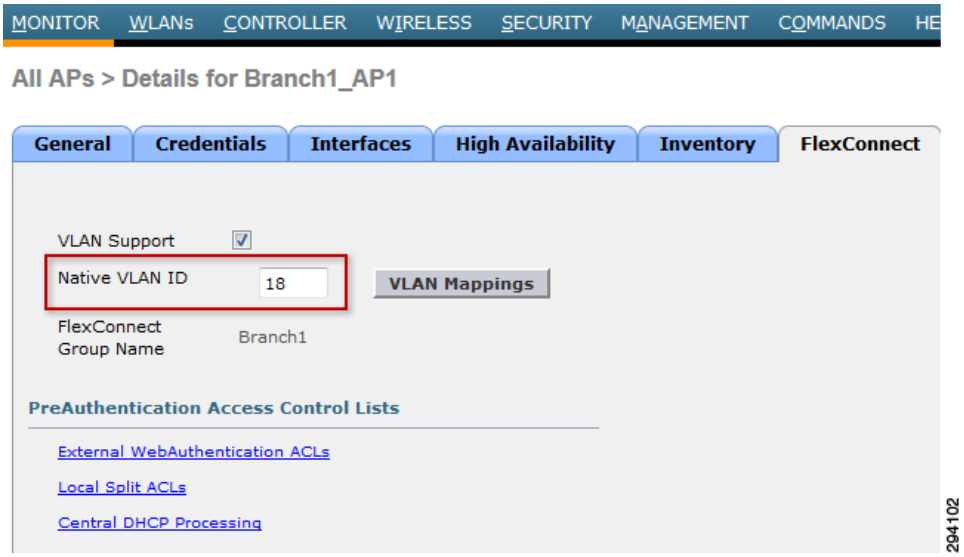
Configure the access point in FlexConnect mode by changing the AP Mode to FlexConnect. Click **Wireless > Access Points** and select the proper branch AP. Figure 9-26 shows the setting for an access point in Branch1.

Figure 9-26 FlexConnect AP Mode



Click the **FlexConnect** tab and specify the Native VLAN for the branch, as shown in Figure 9-27. The access point relies on the native VLAN for IP connectivity.

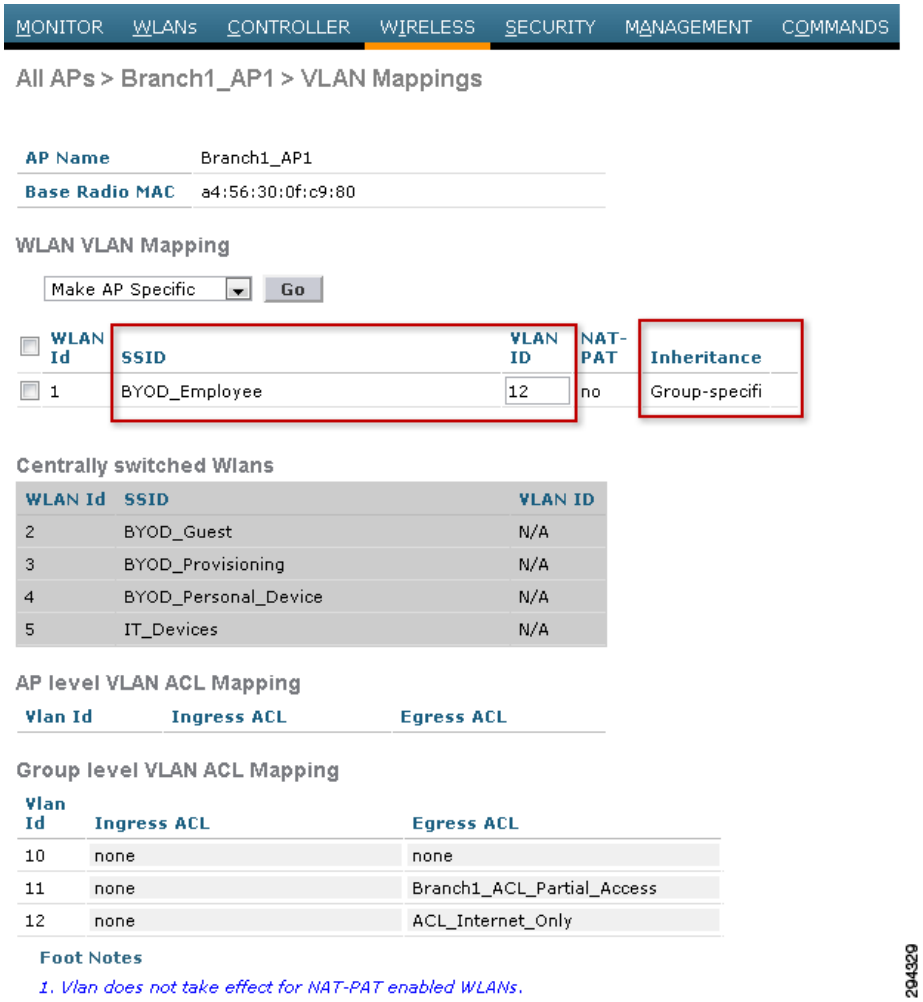
Figure 9-27 Native VLAN ID



Define the VLAN ID to be used for local switching. In Figure 9-28, clients obtain an IP address from VLAN 12 (Internet access) when doing local switching. When using the AAA Overrides for FlexConnect feature, the client is moved to a different VLAN dynamically, based on the matched authorization profile and will obtain an IP address from the defined VLAN.

This setting can be configured at the AP level or the AP can inherit the settings from the FlexConnect Group. FlexConnect Groups are explained in the next section.

Figure 9-28 BYOD_Employee VLAN ID



FlexConnect Groups

FlexConnect groups provide a convenient way to group access points that share the same configuration settings. This is particularly helpful when grouping several FlexConnect access points in remote or branch locations. Instead of configuring each access point separately, FlexConnect groups allow the configuration parameters to be applied to all access points at once. For example, a FlexConnect ACL can be applied to a particular VLAN across all access points within a branch simply by adding the access points to the same FlexConnect group.

For the purpose of this guide, a unique FlexConnect group was defined for each branch, as shown in Figure 9-29.

Figure 9-29 FlexConnect Groups

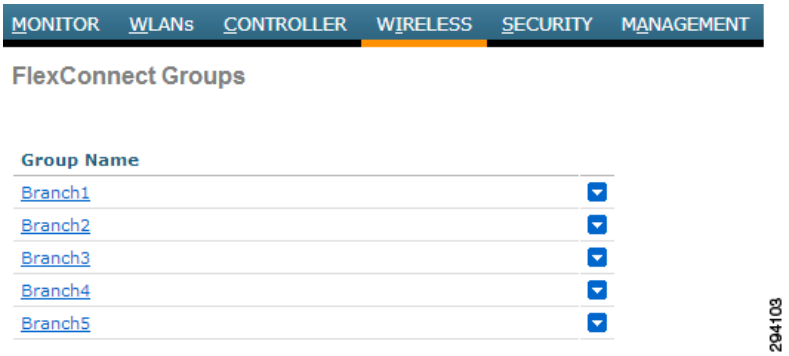
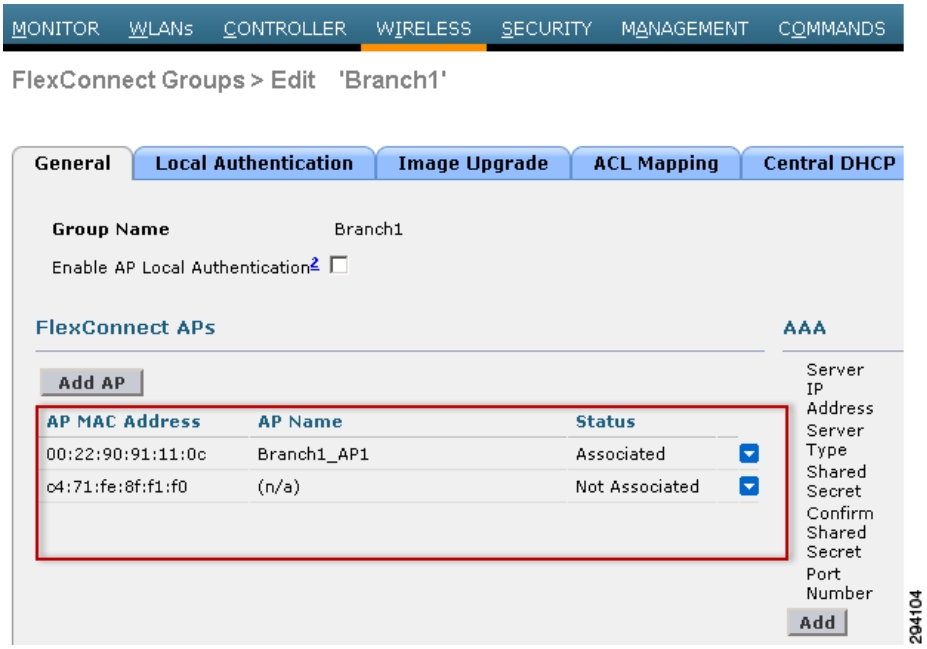


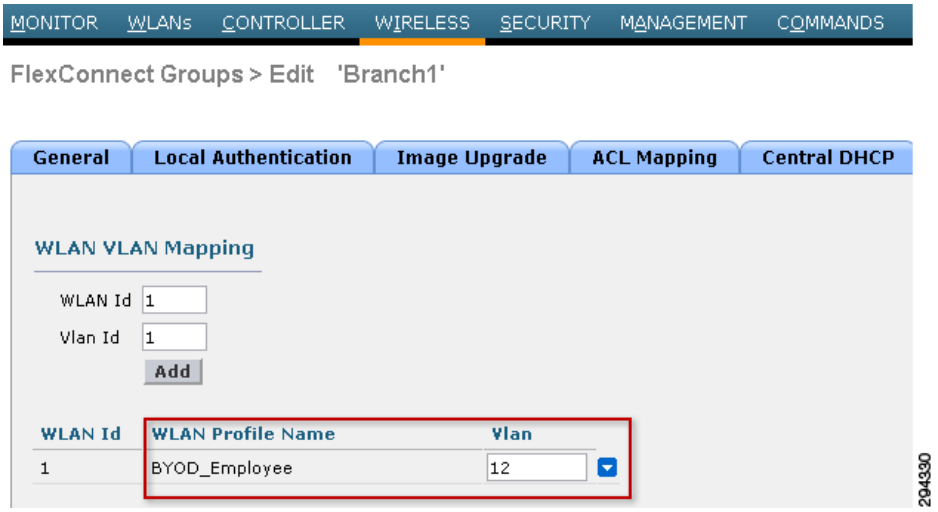
Figure 9-30 shows the access points that have been added to the Branch1 FlexConnect group.

Figure 9-30 Branch1 FlexConnect Group



The VLAN ID used for local switching can be defined at the AP level,as shown in Figure 9-28, or at the FlexConnect Group level, as shown in Figure 9-31. In this example, clients will obtain an IP address from VLAN 12 (Internet access) when doing local switching. When using the AAA Overrides for FlexConnect feature, the client is moved to a different VLAN dynamically based on the matched authorization profile and will obtain an IP address from the defined VLAN.

Figure 9-31 Local Switching VLAN—FlexConnect Group Level

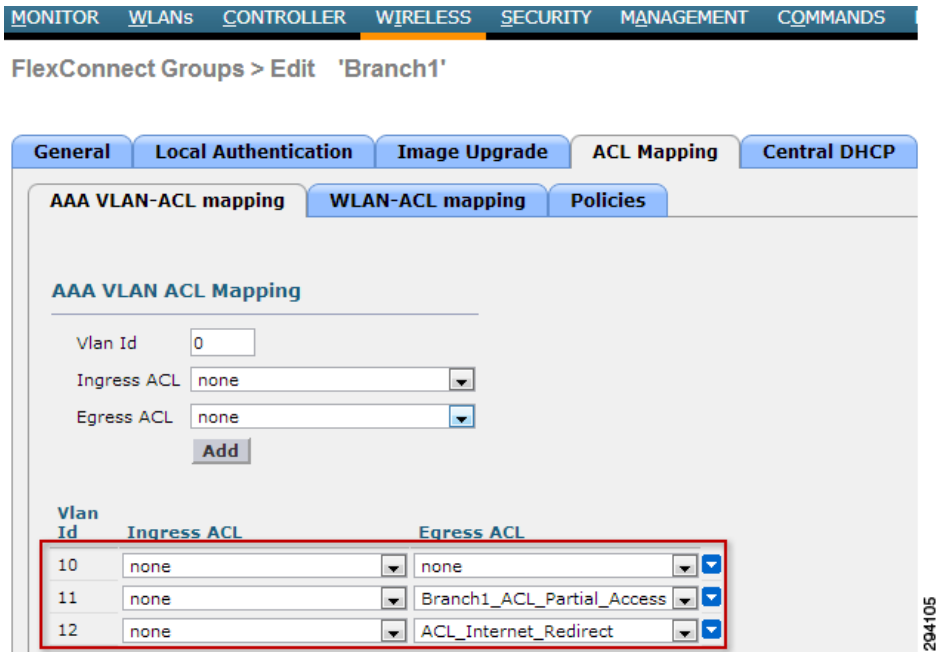


Before ISE can enforce an authorization policy, FlexConnect ACLs must be defined and assigned to each VLAN. By clicking the AAA VLAN-ACL mapping tab, the FlexConnect ACL may be enforced for each VLAN ID. This assumes that every branch location shares the same VLAN ID numbers:

- VLAN 10 for full access
- VLAN 11 for partial access
- VLAN 12 for Internet only access

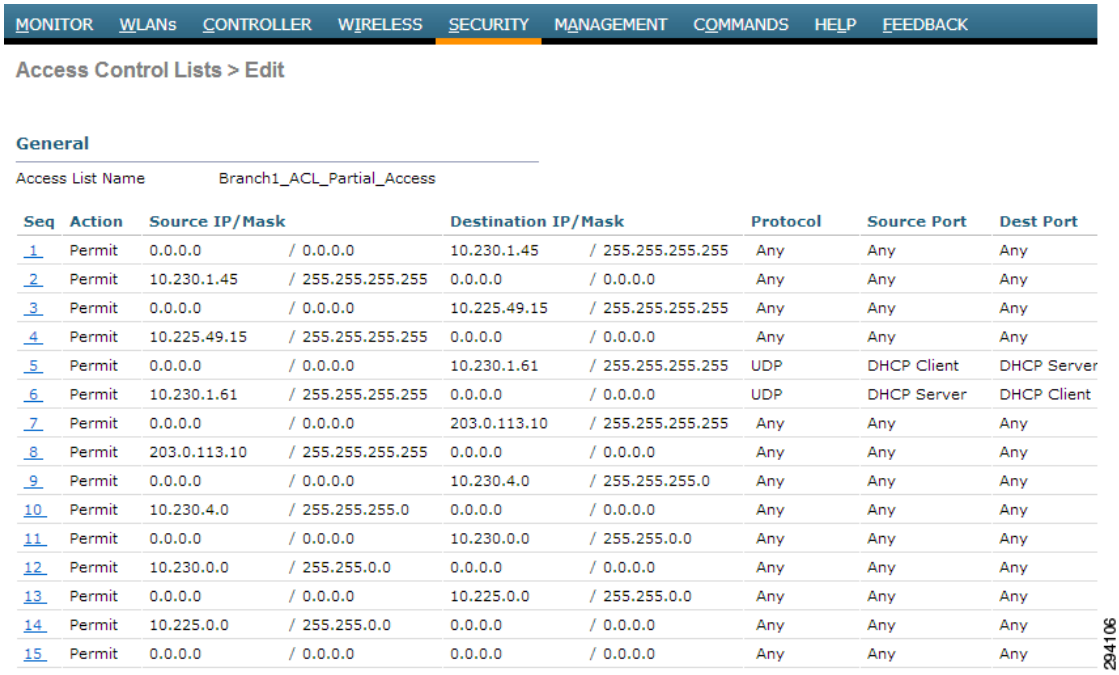
Figure 9-32 shows how the different FlexConnect ACLs have been mapped to each VLAN.

Figure 9-32 VLAN-ACL Mapping



The FlexConnect ACLs shown in Figure 9-33 and Figure 9-34 are explained in more detail in Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices."

Figure 9-33 Branch1_ACL_Partial_Access FlexConnect ACL



MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Access Control Lists > Edit

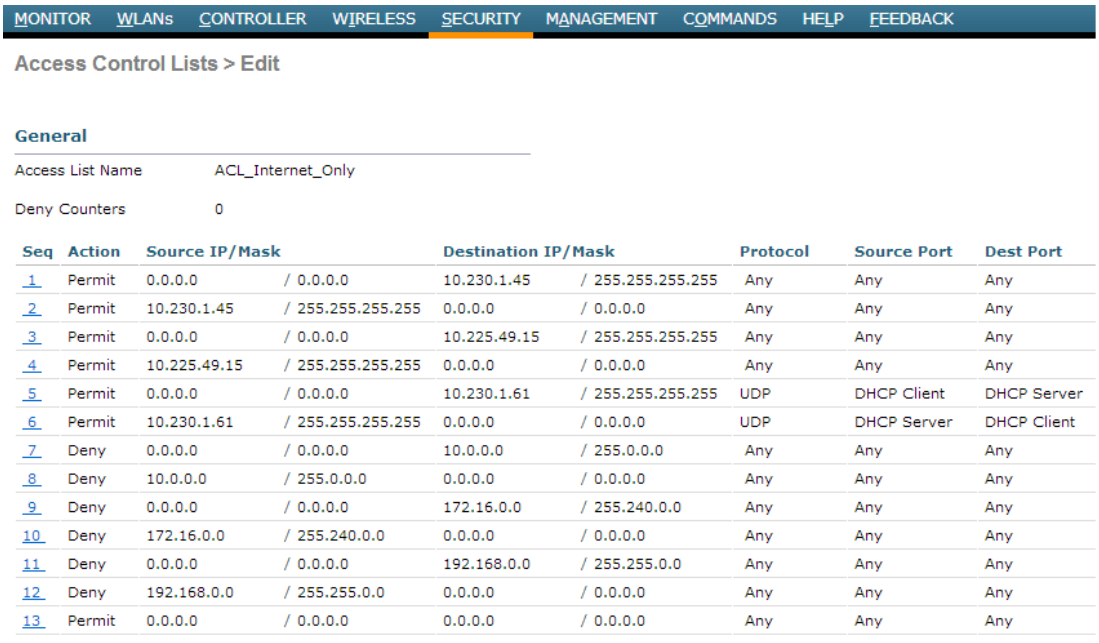
General

Access List Name Branch1_ACL_Partial_Access

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client
7	Permit	0.0.0.0	/ 0.0.0.0	203.0.113.10	/ 255.255.255.255	Any	Any	Any
8	Permit	203.0.113.10	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
9	Permit	0.0.0.0	/ 0.0.0.0	10.230.4.0	/ 255.255.255.0	Any	Any	Any
10	Permit	10.230.4.0	/ 255.255.255.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
11	Permit	0.0.0.0	/ 0.0.0.0	10.230.0.0	/ 255.255.0.0	Any	Any	Any
12	Permit	10.230.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
13	Permit	0.0.0.0	/ 0.0.0.0	10.225.0.0	/ 255.255.0.0	Any	Any	Any
14	Permit	10.225.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
15	Permit	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any

294106

Figure 9-34 ACL_Internet_Only



MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Access Control Lists > Edit

General

Access List Name ACL_Internet_Only

Deny Counters 0

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client
7	Deny	0.0.0.0	/ 0.0.0.0	10.0.0.0	/ 255.0.0.0	Any	Any	Any
8	Deny	10.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
9	Deny	0.0.0.0	/ 0.0.0.0	172.16.0.0	/ 255.240.0.0	Any	Any	Any
10	Deny	172.16.0.0	/ 255.240.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
11	Deny	0.0.0.0	/ 0.0.0.0	192.168.0.0	/ 255.255.0.0	Any	Any	Any
12	Deny	192.168.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
13	Permit	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any

294107

FlexConnect VLAN Override

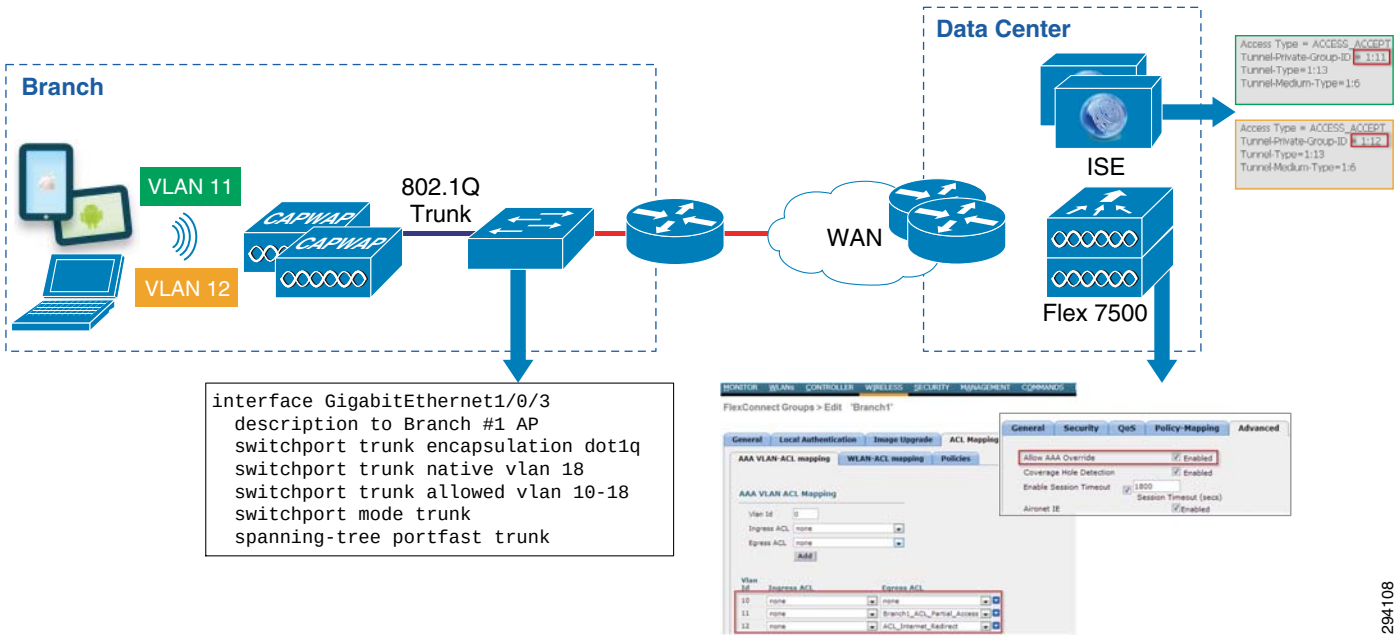
In the current FlexConnect architecture, there is a strict mapping of WLAN to VLAN, so the client getting associated on a particular WLAN on a FlexConnect AP has to abide by the VLAN which is mapped to it. This method has limitations because it requires clients to associate with different SSIDs in order to inherit different VLAN-based policies.

Starting on WLC release 7.2, AAA Override (Dynamic VLAN assignment) of VLANs on individual WLANs configured for local switching is supported. To assign endpoints dynamically to a VLAN, the VLAN IDs are pre-created and the corresponding WLAN-VLAN Mapping on a FlexConnect group is configured, as shown in Figure 9-32.

Figure 9-35 shows the different configuration settings required to dynamically assign endpoints to a branch VLAN, which include:

- The WLAN at the branch configured for local switching mode.
- 802.1Q trunk between the Catalyst switch and the access point.
- A native VLAN and allowed VLANs for the trunk.
- The ISE authorization profile defines what VLAN is assigned to the endpoint.
- The WLAN is configured at the controller to allow AAA Override.
- The VLANs are pre-defined and the VLAN-ACL mapping is defined for the FlexConnect group.

Figure 9-35 FlexConnect VLAN Override



Campus—Converged Access Design

The converged large campus design looks at the hybrid large campus design model, as discussed in Campus Migration Path of Chapter 5, "Campus and Branch Network Design for BYOD." A hybrid large campus design consists of multiple Catalyst 3850s switches or switch stacks deployed at the access layer

of the network, operating in Mobility Agent (MA) Mode. A centralized Cisco CT5760 controller within the campus contains the Mobility Controller (MC) function. A Unified Controller CT5508 exists within the campus controller and forms a mobility group with CT5760s. APs may be connected to the CT5760 or CT5508 controllers via Catalyst 3850 or CT3750 switches. In addition a CT5760 or CT5508 may be used as guest access anchor at the Internet edge of the campus. In this design guide the CT5508 is configured as a guest controller.

This design guide will make the following assumptions for the large campus converged access design:

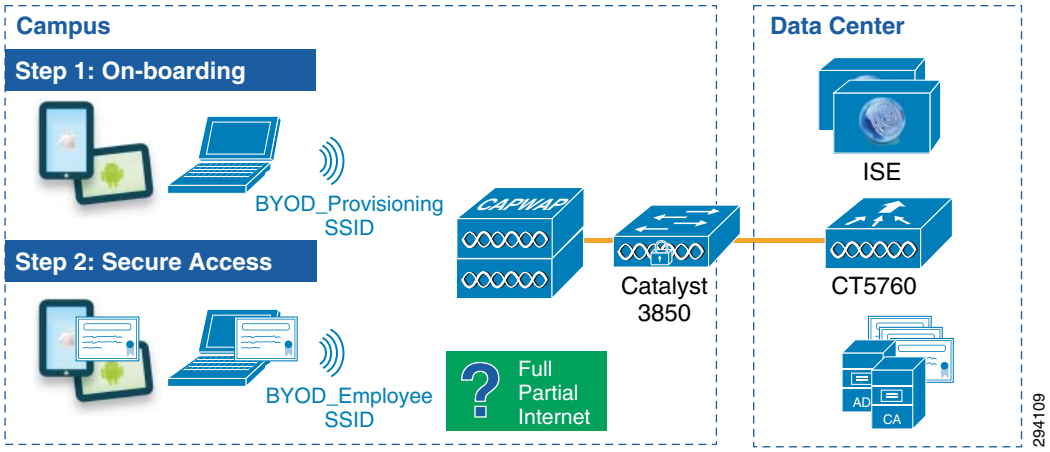
- On-boarded wired and wireless devices will share the same VLAN, and hence the same IP subnet addressing space. It is recognized that customers may implement separate subnets for wired and wireless devices due to issues such as additional security compliance requirements for wireless devices. This is not addressed within this version of the design guidance.
- Catalyst 3850 Series switches are deployed as Layer 2 access switches within the campus. Layer 3 connectivity will be provided by the Catalyst 6500 building distribution switches. Also, in keeping with campus best practices, VLANs will be limited to a single wiring closet. In other words, VLANs will not extend between access-layer switches. Future design guidance may address Catalyst 3850 Series switches deployed as Layer 3 switches or address spanning VLANs across access-layer switches.

Campus Converged Access—Dual SSID Design

In this design there are again two SSIDs: one provides enrollment/provisioning and the other provides secure network access. After connecting to the BYOD_Provisioning SSID and completing the enrollment and provisioning steps, the user connects to the BYOD_Employee SSID, which provides network access over a secure EAP-TLS connection.

Figure 9-36 shows the dual SSID design for the campus APs.

Figure 9-36 Campus—Dual SSID



In the converged access dual SSID design, there are some additional considerations:

- The provisioning SSID can be either open or password protected. When the provisioning SSID is open, any user can connect to the SSID, whereas if it is password protected, then only users that have credentials, such as AD group membership, are allowed to connect to the SSID. In this design guide, the provisioning SSID is configured to be open and its only purpose is to provide on-boarding services.

- After the device is provisioned, it is assumed that the user will switch to the second SSID for regular network access. To prevent the user from staying connected to the provisioning SSID, an access list that provides only access to ISE, DHCP, and DNS must be enforced on the provisioning SSID. The details of the ACL_Provisioning_Redirect ACL are shown below.
- This design guide makes use of the following SSIDs: BYOD_Provisioning and BYOD_Employee.

The properties of these two SSIDs are highlighted in Table 9-4.

Table 9-4 WLAN Parameters

Attribute	BYOD_Provisioning	BYOD_Employee
Description	Used only for device provisioning	For employees that have completed the on-boarding process
Layer 2 Security	None (for Open SSID)	WPA+WPA2
MAC Filtering	Enabled (for Open SSID)	Disabled
WPA+WPA2 Parameters	None	WPA2 Policy, AES, 802.1X
Layer 3 Security	None	None
AAA Server	Select ISE	Select ISE
Advanced	AAA Override Enabled	AAA Override Enabled
Advanced	NAC State- NAC	NAC State- NAC

To configure WLAN BYOD_Provisioning SSID on a CT5760 and Catalyst 3850 follow the steps below. The security on the BYOD_Provisioning SSID is NONE as this is a provisioning SSID through which devices are provisioned on the network. The FAST-SSID feature provides a way for a client to directly switch from BYOD_Provisioning to BYOD_Employee SSID after it has been properly provisioned by ISE.

```
aaa new-model
!
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
!
!
!
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 032A4802120A701E1D5D4C
!
aaa session-id common
!
ip device tracking
!
!
qos wireless-default-untrust
captive-portal-bypass
!
mac access-list extended MAC_ALLOW
 permit any any
!
!
interface Vlan2
 description ### BYOD-Employee Vlan ###
```

```
 ip address 10.231.2.7 255.255.255.0
 load-interval 30
!
interface Vlan3
 description ### BYOD-Provisioning Vlan ###
 ip address 10.231.3.7 255.255.255.0
 load-interval 30
!
ip http server
ip http authentication local
ip http secure-server
!
wireless management interface Vlan47
wireless client fast-ssid-change
wireless rf-network byod
wireless security dot1x radius call-station-id macaddress
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan BYOD-Employee
 nac
 security web-auth parameter-map global
 session-timeout 1800
 no shutdown
wlan BYOD_Provisioning 3 BYOD_Provisioning
 aaa-override
 client vlan BYOD-Provisioning
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 session-timeout 1800
 no shutdown
```

The example configuration shown above must be configured on both the Catalyst 3850 which functions as the Mobility Agent (MA), and the CT5760 which functions as the Mobility Controller (MC) in the campus design. Note that the IP addressing for the VLAN interfaces will be different for the MA and MC, however, since they are deployed in different parts of the network infrastructure. Mobility is handled as a separate topic within this chapter, following the WLAN configuration discussion. Additional configuration lines must be added to the MA and MC, respectively for mobility. These are discussed shortly.

The BYOD_Provisioning SSID has no Layer 2 security, as this is an SSID through which devices are provisioned on the network. Instead the wireless client uses MAC-filtering (basically a wireless version of MAB) to authenticate to the network. A URL re-direction and Centralized Web Authentication (CWA) policy is pushed down to the client from ISE, upon connecting to the network. Hence, the configuration for MAC-filtering, NAC, and AAA override are required on the BYOD_Provisioning SSID.

The security on the BYOD_Employee SSID is WPA2 with AES encryption. Note that this is the default setting for a WLAN on the Converged Access platforms (CT5760 or Catalyst 3850) and therefore does not appear within the configuration. The configuration for NAC and AAA override are required on this SSID in order to support a dynamic ACL assignment to the wireless client. In the case of this design guide, the dynamic ACL is a named ACL configured locally on the Catalyst 3850 switch.



Note The administrative level command **show wlan name *<name_of_wlan>*** can be used to show the details regarding the configuration of any WLAN on either the Catalyst 3850 Series switch or the CT5760 wireless controller. This includes any default settings which do not appear within the configuration.

Even though a CWA policy is pushed to the wireless client from ISE during on-boarding via the BYOD_Provisioning SSID, the HTTP and HTTPS server functionality must be globally enabled on the Catalyst 3850 Series switch. This is in order to support the URL re-direction of web sessions from wireless clients to the ISE provisioning portal. The RADIUS server group configuration points back to ISE as the RADIUS server for authentication and authorization of wireless (and wired) clients. The captive portal bypass functionality must be globally enabled on the Catalyst 3850 Series switch in order to allow Apple devices to on-board successfully. The fast-ssid-change global configuration provides a way for client to switch from BYOD_Provisioning to BYOD_Employee SSID after it has been properly provisioned by ISE.

The wireless mobility configuration commands for the CT5760 which functions as the MC will be different from the Catalyst 3850 which functions as the MA. An example of the global mobility configuration lines for the CT5760 wireless controller is shown below.

```
!
interface Vlan47
 description MGMT VLAN
 ip address 10.225.47.2 255.255.255.0
 load-interval 30
!
wireless mobility controller peer-group 100
wireless mobility controller peer-group 100 member ip 10.203.61.5 public-ip 10.203.61.5
wireless mobility controller peer-group 200
wireless mobility controller peer-group 200 member ip 10.207.61.5 public-ip 10.207.61.5
wireless mobility controller peer-group 200 member ip 10.207.71.5 public-ip 10.207.71.5
wireless mobility group member ip 10.225.50.36 public-ip 10.225.50.36/Points to CT5508
wireless mobility group name byod
wireless management interface Vlan47
wireless rf-network byod
!
```

As can be seen, the CT5760 is configured as the mobility controller (MC) for two switch peer-groups (SPGs)—100 and 200—in the example above. Switch peer-group 100 contains a single Catalyst 3850 switch functioning as a MA. Switch peer-group 200 contains two Catalyst 3850 switches functioning as MAs. An example of the global mobility configuration lines for the Catalyst 3850 is shown below.

```
interface Vlan47
 description MGMT VLAN
 ip address 10.225.61.5 255.255.255.0
 load-interval 30
!

wireless mobility controller ip 10.225.47.2 public-ip 10.225.47.2 / IP Address of 5760 MC
```

The IP address corresponding to the wireless management interface of the Catalyst 3850 series switch shown in the configuration above appears as a member of SPG 200. SPGs are designed to scale mobility within a Converged Access design. Roaming between Catalyst 3850 Series switch mobility agents (MAs) within a single SPG is handled directly by the switches without the involvement of the CT5760 mobility controller (MC). This is done via a full mesh of CAPWAP tunnels between the Catalyst 3850 Series switch mobility agents (MAs) within a single SPG. Roaming between Catalyst 3850 Series switch mobility agents (MAs) across two SPGs is handled by the CT5760 mobility controller (MC). This is done via CAPWAP tunnels between each Catalyst 3850 Series switch mobility agent (MA) and the CT5760 mobility controller (MC).

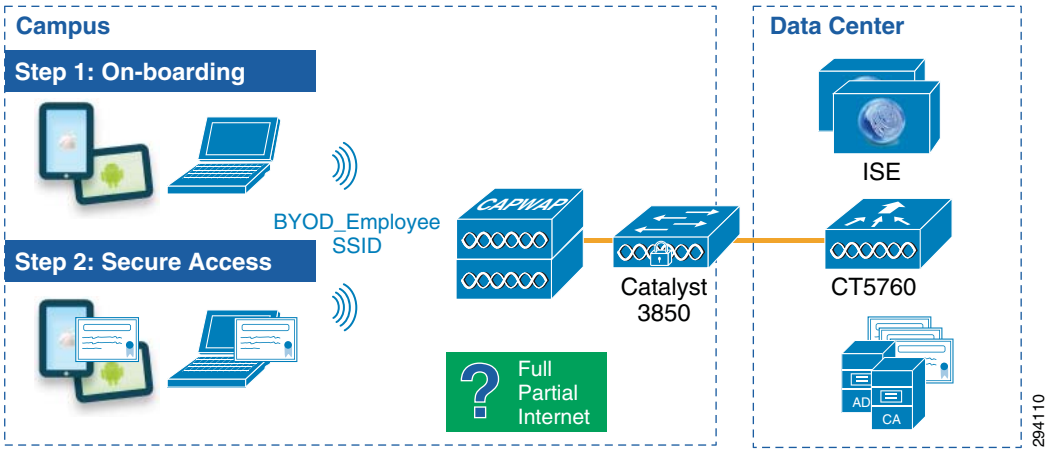
As discussed previously, a hybrid campus design may consist of CT5508 wireless controllers operating in Local Mode, alongside the Converged Access infrastructure. This may be necessary during the migration from a centralized wireless overlay model to a Converged Access deployment model. In order

to support mobility between the CT5508 wireless controller and the CT5760 wireless controller, the IP address of the CT5508 wireless controller has been added as a wireless mobility group member to the configuration of the CT5760 shown above.

Campus Converged Access—Single SSID Design

In a single SSID design the same WLAN (BYOD_Employee) is used for on-boarding and secure network access. Figure 9-37 shows how this design may be implemented using the CT5760 as an MC and Catalyst 3850 as MA.

Figure 9-37 Campus-Single SSID



The configuration for a single SSID converged campus design is almost the same as a dual SSID design but without the use of the BYOD_Provisioning SSID. A snippet of configuration on the CT5760 and the Catalyst 3850 is shown below. Mobility is handled as separate topic following the WLAN configuration discussion.

```
aaa new-model
!
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
!
!
!
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 032A4802120A701E1D5D4C
!
aaa session-id common
!
ip device tracking
!
!
qos wireless-default-untrust
captive-portal-bypass
!
mac access-list extended MAC_ALLOW
 permit any any
!
```

```
!
interface Vlan2
 description ### BYOD-Employee Vlan ###
 ip address 10.231.2.7 255.255.255.0
 load-interval 30
!
interface Vlan3
 description ### BYOD-Provisioning Vlan ###
 ip address 10.231.3.7 255.255.255.0
 load-interval 30
!
ip http server
ip http authentication local
ip http secure-server
!
wireless management interface Vlan47
wireless client fast-ssid-change
wireless rf-network byod
wireless security dot1x radius call-station-id macaddress
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan BYOD-Employee
 nac
 security web-auth parameter-map global
 session-timeout 1800
 no shutdown
```

The mobility configuration for both the MC and MA will remain the same as discussed above for the dual SSID Converged Access design.

Campus Converged Access—Mobility

For the large campus design it is important to understand mobility and roaming considerations.

This design highlights multiple Catalyst 3850 Series switches or switch stacks deployed at the access layer of a large sized campus. Switch stacks form Switch Peer Groups (SPGs) in which all switches contain the Mobility Agent (MA) function. Roaming within a SPG is handled through a full mesh of mobility tunnels between MAs within the SPG. Multiple SPGs exist within the large sized campus. APs must be directly connected to MA and not via an intermediate switch (example: a Catalyst 3750 switch).

A Cisco CT5760 wireless controller deployed within a centralized service module within the campus contains the Mobility Controller (MC) function. Multiple SPGs connecting to a single MC form a Mobility Sub-Domain. Multiple Mobility Sub-Domains exist within the large sized campus. Roaming between SPGs within a Mobility Sub-Domain is done through the Cisco CT5760 and/or CT5508 wireless controller. APs connected to a Catalyst 3850 switch register with the CT5760 MC. APs can also be connected to CT5760 via Catalyst 3750 switches.

Multiple Cisco CT5760 and/or CT5508 wireless controllers form a Mobility Group. Hence, a Mobility Group also consists of multiple Mobility Sub-Domains. Roaming between Mobility Sub-domains is done through the Cisco CT5760 and/or CT5508 wireless controllers within the Mobility Group. A single Mobility Group and hence a single Mobility Domain extends across and are entirely contained within the large campus within this design.

For hybrid models consisting of both a CUWN local-mode and converged access products, either a Cisco CT5760 or a CT5508 also serves as a wireless controllers for access points connected to Catalyst 3750-X Series switches using traditional local mode (centralized switching) wireless connectivity.

Keeping above the considerations in mind, few things should be kept in mind.

By default Catalyst 3850 operates as a Mobility Agent and there is no need of any configuration. A Catalyst 3850 may also operate as a Mobility Controller. This mode is covered as part of Branch Design.

See Appendix C, "Software Versions" for details about the Catalyst 3850 software licensing.

CT5760 wireless controller operates only as a Mobility Controller. Mobility tunnels should be setup between CT5760s and Catalyst 3850s for APs connected on Catalyst 3850s to be registered with the MC (CT5760). A snippet of configuration for MC is as below:

```
wireless mobility controller peer-group 100
wireless mobility controller peer-group 100 member ip 10.203.61.5 public-ip 10.203.61.5
wireless mobility controller peer-group 200
wireless mobility controller peer-group 200 member ip 10.207.61.5 public-ip 10.207.61.5
wireless mobility controller peer-group 200 member ip 10.207.71.5 public-ip 10.207.71.5
```

On each Catalyst 3850 acting as an MA, the configuration below is needed to establish a mobility tunnel with the CT5760 MC or a 5508 MC.

```
wireless mobility controller ip 10.225.47.2 public-ip 10.225.47.2 / IP Address of MC
```

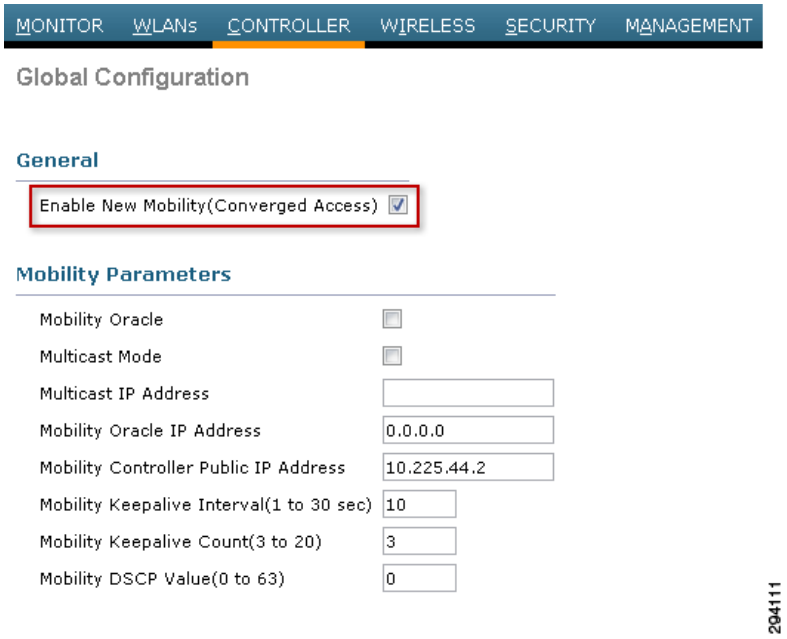
The CT5508 and CT5760 can also form a mobility group. The CT5508 should be upgraded to either 7.3.112 or a version above 7.5 of the WLC to support mobility between converged access and unified access products. The configuration on the CT5508 to enable mobility between the CT5760 and the CT5508 is shown below. The design guide provides guidance for version 7.5 for CT5508.

```
wireless mobility controller/ Enables the MC function, by default turned on CT5760
wireless mobility group name byod/ Create mobility group byod
wireless mobility group member ip 10.225.50.36 public-ip 10.225.50.36/ IP of member CT5508
```

Note Only WLC versions 7.3.112 or 7.5 and above support mobility between converged access products and unified access products. Ensure that you have code version running compatible code. This design guide uses 7.5 release.

To enable mobility between converged access and unified access products, first the New Mobility should be enabled on the WLC as shown in Figure 9-38.

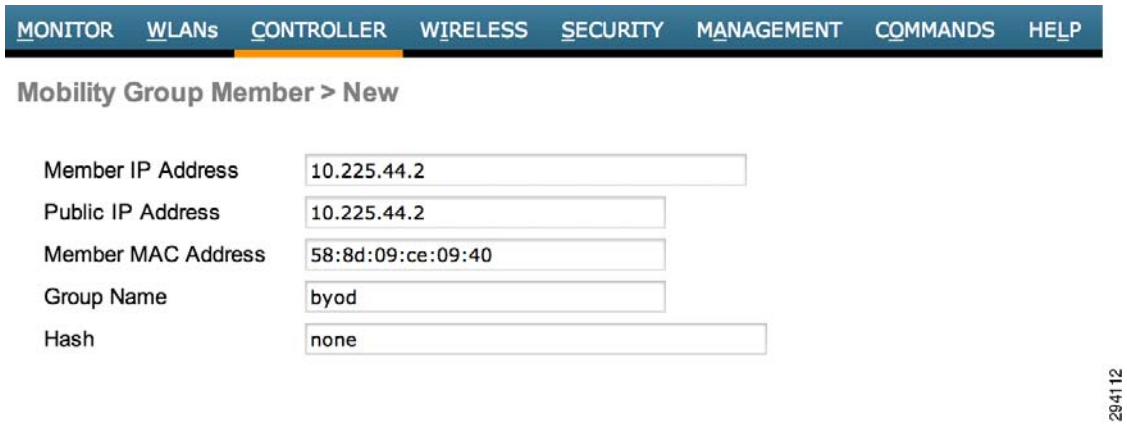
Figure 9-38 Enable New Mobility



After enabling New Mobility and restarting the Wireless LAN Controller, additional options for configuring switch peer groups as well as mobility groups are enabled. For CT5760 and CT5508 to form a group and talk to each other, additional configuration as below is required.

Click **Mobility Management > Mobility Groups** and click **New**, as shown in Figure 9-39.

Figure 9-39 Create New Mobility Group



The Member IP address above should be the CT5760 IP address that enables mobility messaging and CAPWAP tunnels to be set up between CT5760 and CT5508.

Other design considerations while deploying a large campus WLAN infrastructure include the following:

- 802.1X, WLAN, and VLAN configurations should be replicated on all Catalyst 3850s and CT5760s.
- Mobility group name should be the same between CT5760s and CT5508s.

Branch—Converged Access Design

With a converged access design, a centralized FlexConnect wireless controller can be replaced by a Catalyst 3850 switch that operates both as a Mobility Agent (MA) and Mobility Controller (MC). Guest wireless access still utilizes the same model wherein the guest traffic is auto-anchored to a dedicated guest anchor controller located within the Internet Edge of the campus. The guest controller can be a CT5508 controller with a 7.5 version of code, or a CT5760 converged wireless LAN controller.

The integrated controller branch BYOD design guide makes the following assumptions:

- On-boarded wired and wireless devices will share the same VLAN and hence the same IP subnet addressing space. It is possible to use different VLAN and addressing space for wireless and wired clients, however it is not addressed in this design guide.
- The Catalyst 3850 switches are deployed as a Layer 2 switches within the branch location. Layer 3 connectivity within the branch will be provided by ISR routers which also serve as the WAN connectivity point for the branch. (Future design guides may address Catalyst 3850 deployed as Layer 3 switch within the branch location).

Branch Converged Access—Dual SSID Design

In the dual-SSID design, a dedicated open SSID (BYOD_Provisioning) with MAC-filtering (i.e., MAC Authentication Bypass) will be configured for on-boarding devices. The SSID will be statically mapped to a separate Provisioning VLAN on the Catalyst 3850 switch. Figure 9-40 shows the branch converged access for a dual SSID design.

Figure 9-40 Branch Converged Access—Dual SSID

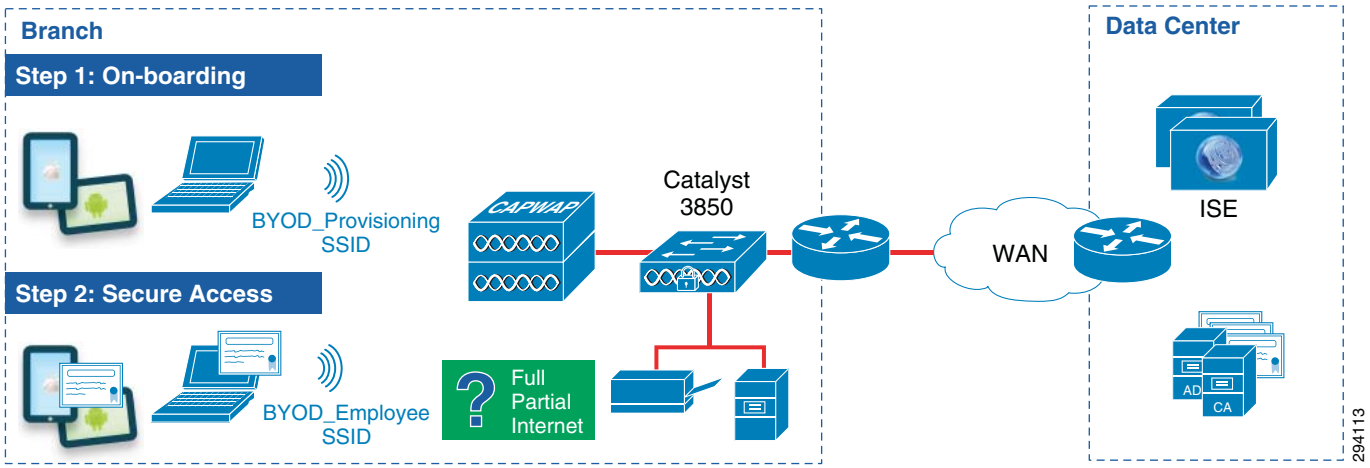


Table 9-5 summarizes the VLANs within the branch when utilizing the dual-SSID BYOD on-boarding design.

Table 9-5 VLANs in Branch with Dual-SSID BYOD On-boarding Design

Description	VLAN	VLAN Name
Wired and wireless corporate access. IT managed devices. Employee managed devices with full, partial, or Internet access.	12	BYOD_Employee
Provisioning VLAN for Dual-SSID wireless on-boarding.	13	BYOD_Provisioning
Separate VLAN for branch servers.	16	Server
Dedicated VLAN for management of network infrastructure.	18	Management
Isolated VLAN for pass through of wireless auto-anchor tunnels. Not trunked to Layer 3 router.	777	BYOD_Guest

The following configuration snippet provides an example of the possible configuration additions to the Catalyst 3850 in order to support on-boarding of wireless devices in a dual-SSID BYOD implementation using MAC-filtering.

```
aaa new-model
!
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting network default start-stop group radius
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 032A4802120A701E1D5D4C
 auth-type any
!
aaa session-id common
!
ip device tracking
!

qos wireless-default-untrust
vtp domain bn
!
mac access-list extended MAC_ALLOW
 permit any any
!
wireless mobility controller
wireless mobility group member ip 10.225.50.36 public-ip 10.225.50.36
wireless mobility group name byod
wireless management interface Vlan18
wireless client fast-ssid-change
wireless rf-network byod
wireless security dot1x radius call-station-id macaddress
wireless broadcast
wireless multicast
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan BYOD_Employee
 nac
 security dot1x authentication-list default
 session-timeout 1800
 no shutdown
wlan BYOD_Guest 2 BYOD_Guest
 aaa-override
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
```

```
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth
 session-timeout 1800
 no shutdown
wlan BYOD_Provisioning 3 BYOD_Provisioning
 aaa-override
 client vlan BYOD_Provisioning
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 session-timeout 1800
 no shutdown
!
```

The following configuration snippet provides a partial example of the possible configuration additions to the branch router configuration in order to support on-boarding of wireless devices in a dual-SSID BYOD implementation using MAC-filtering—when the Catalyst 3850 Series switch is functioning as a Layer 2 switch.

```
!
interface GigabitEthernet0/0
 description CONNECTION TO CATALYST 3850 SWITCH
 no ip address
 load-interval 30
 duplex auto
 speed auto
!
interface GigabitEthernet0/1.13/ Provisioning VLAN
 description CATALYST 3850 PROVISIONING VLAN
 encapsulation dot1Q 13
 ip address 10.200.13.2 255.255.255.0
 ip helper-address 10.230.1.61/ Relay DHCP to the DHCP server
 ip helper-address 10.225.42.15/ Relay DHCP to ISE for profiling
 standby 13 ip 10.200.13.1
 standby 13 priority 110
 standby 13 preempt
!
```

Branch Converged Access—Single SSID Design

In the single SSID design, the corporate SSID (BYOD_Employee) supports authentication via PEAP for non on-boarded devices. Once on-boarding is complete, the corporate SSID supports authentication via EAP-TLS for on-boarded devices. The corporate SSID is statically mapped to a separate Corporate VLAN on the Catalyst 3850 switch. Figure 9-41 shows the branch converged access for a single SSID Design.

Figure 9-41 Branch Converged Access—Single SSID

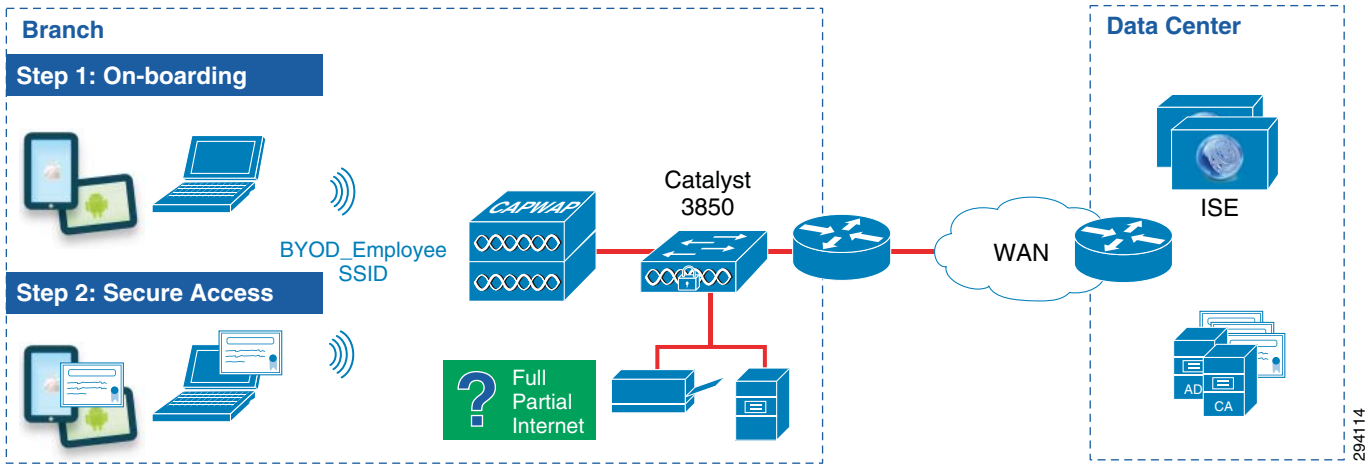


Table 9-6 summarizes the VLANs within the branch when utilizing the single-SSID BYOD on-boarding design.

Table 9-6 VLANs in Branch when Utilizing Single-SSID BYOD On-boarding Design

Description	VLAN	VLAN Name
Wired and wireless corporate access. IT managed devices. Employee managed devices with full, partial, or Internet access.	12	BYOD_Employee
Separate VLAN for branch servers.	16	Server
Dedicated VLAN for management of network infrastructure.	18	Management
Isolated VLAN for past through of wireless auto-anchor tunnels. Not trunked to Layer 3 router.	777	BYOD_Guest

The following configuration shows relevant parts of configuration for the Catalyst 3850 when utilizing a single SSID on-boarding model.

```
aaa new-model
!
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting network default start-stop group radius
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 032A4802120A701E1D5D4C
 auth-type any
!
aaa session-id common
!
ip device tracking
!
!
qos wireless-default-untrust
!
mac access-list extended MAC_ALLOW
 permit any any
!
```

```
wireless mobility controller
wireless mobility group member ip 10.225.50.36 public-ip 10.225.50.36
wireless mobility group name byod
wireless management interface Vlan18
wireless client fast-ssid-change
wireless rf-network byod
wireless security dot1x radius call-station-id macaddress
wireless broadcast
wireless multicast
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan BYOD_Employee
 nac
 security dot1x authentication-list default
 session-timeout 1800
 no shutdown
wlan BYOD_Guest 2 BYOD_Guest
 aaa-override
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth
 session-timeout 1800
 no shutdown
!
?
```



Identity Services Engine for BYOD

Revised: September 27, 2013

The Cisco Identity Services Engine (ISE) allows for enforcement of centrally configured policies across wired and wireless networks to help organizations provide secure unified access. The Cisco ISE plays a critical role in enabling the BYOD model, where employees are allowed to connect their personal devices securely to the network. By integrating with third-party Mobile Device Managers (MDM), additional device posture may be used to enforce permissions into the network.

Cisco ISE provides a highly scalable architecture that supports both standalone and distributed deployments. The configuration guidelines shown in this document reflect a distributed architecture with multiple nodes.

For small BYOD deployments, one or two ISE nodes may be configured in standalone mode. Depending on how the AAA connections are configured across the access layer switches and Wireless LAN Controllers, either an active/backup or load balancing of AAA workflows can be enabled across the redundant standalone ISE nodes.

For larger BYOD deployments, the ISE functionality can be distributed across multiple nodes. Distributed deployments support the following different ISE personas:

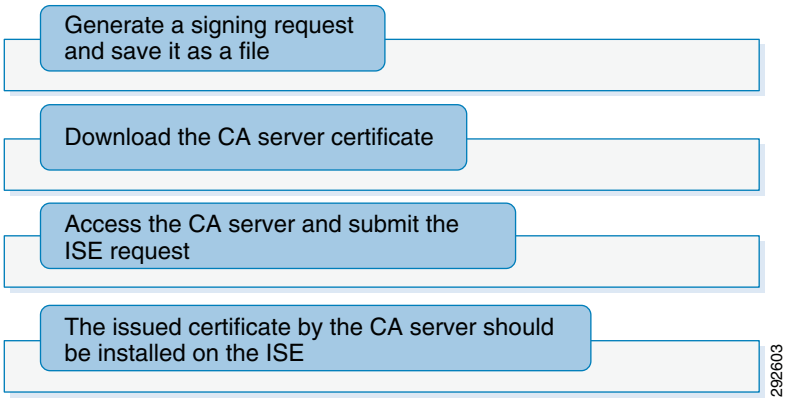
- Administration—The administration node handles all system level configuration. There can be one primary and one secondary administration node in a distributed deployment.
- Monitoring—The monitoring node handles log collection and provides monitoring and troubleshooting tools. There can be one primary and one secondary monitoring node in a distributed deployment.
- Policy Service—The policy service node provides authentication, authorization, guest access, client provisioning, and profiling services. There can be multiple policy services nodes in a distributed deployment.

To support a medium-sized BYOD deployment, both administration and monitoring personas can be deployed on a single node while dedicated policy services nodes can handle AAA functions. For a large BYOD deployment, the monitoring persona can be implemented on a dedicated node providing centralized logging functions.

Identity Certificate for ISE

ISE needs an identity certificate that is signed by a CA server so that it can be trusted by endpoints, gateways, and servers. Figure 10-1 illustrates the steps at a high level.

Figure 10-1 High-Level Steps for Deploying Identity Certificates on ISE



For more details on installing a digital certificate on the Cisco ISE, refer to the TrustSec How-To Guide: http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns744/docs/howto_60_byod_certificates.pdf.

Network Device Definition within ISE

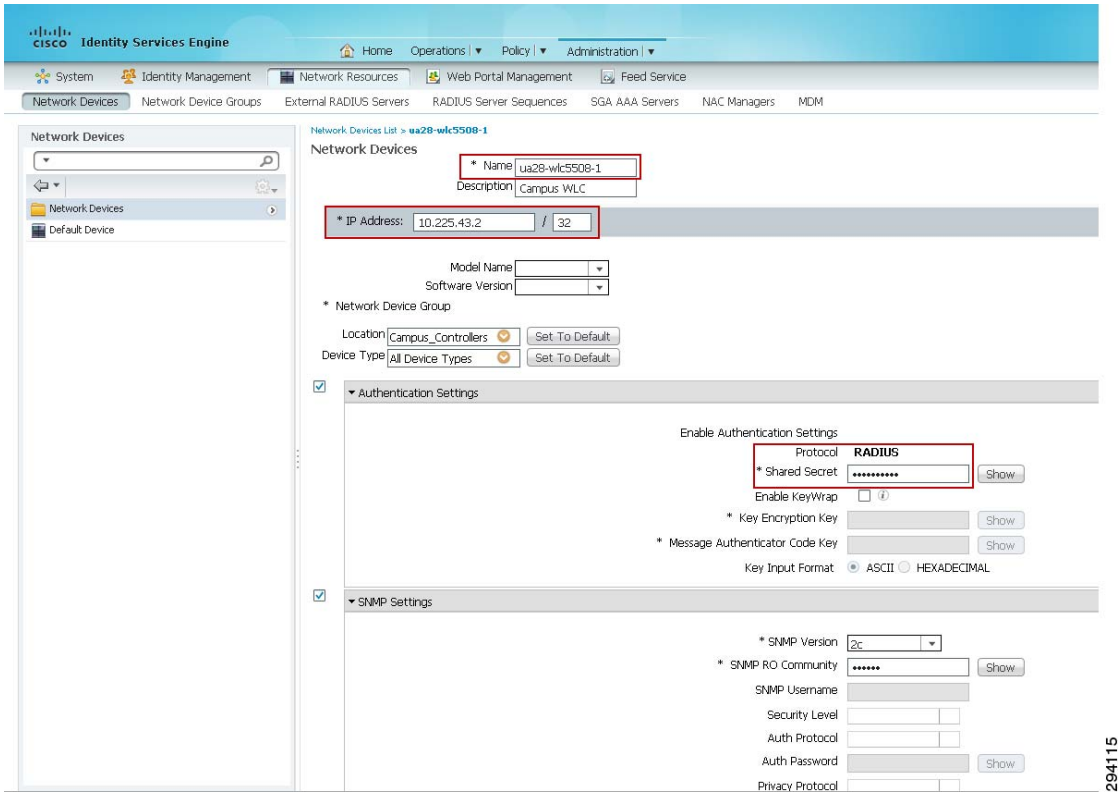
A network device is an authentication, authorization, and accounting (AAA) client through which AAA service requests are attempted, for example, switches, routers, and so on. The network device definition enables the Cisco Identity Services Engine (Cisco ISE) to interact with the network devices that are configured. A network device that is not defined cannot receive AAA services from Cisco ISE.

As users/devices connect to network infrastructure such as wireless controllers and switches enabled for 802.1X authentication, the network device serves as an 802.1X Authenticator to the client's Supplicant. In order for the network device to determine if access is to be granted and what services the device is authorized for, the network device must be able to communicate with the ISE serving as the Authentication Server. To enable this communication, the ISE must be configured with information about that network device as well as credentials to be used to authenticate it.

To configure ISE with this information, refer to Figure 10-2 and the following:

1. At ISE go to **Administration > Network Resources > Network Devices** and click **Add**.
2. Enter the hostname of the device.
3. Enter the IP Address of the network device. This must be the address used to source all RADIUS communications from the device.
4. Change the Network Device Location or Device Type if a custom location/type has been previously defined.
5. Configure the RADIUS Shared Secret. This must match that configured on the network device for the ISE server.
6. Click the down arrow next to SNMP Settings and complete as appropriate.

Figure 10-2 Network Device Configuration in ISE

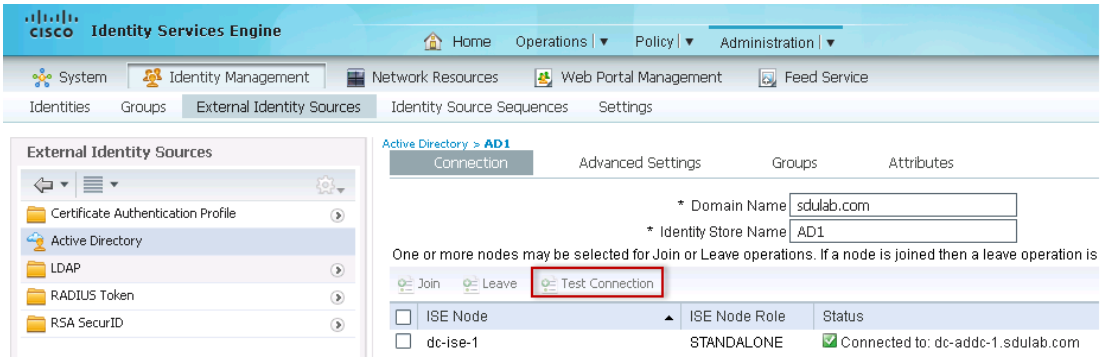


ISE Integration with Active Directory

While the ISE can maintain an internal list of users for authentication purposes, most organizations rely on an external directory as the main identity source. By integrating with Microsoft's Active Directory, objects such as users and groups become critical in the authorization process and can be accessed from a single source.

To integrate with Active Directory, on the ISE click **Administration > External Identity Sources > Active Directory** and specify the domain name, as shown in Figure 10-3. To verify that the ISE node can connect to the Active Directory domain, click **Test Connection** and authenticate with an AD username and password, as shown in Figure 10-3. Click **Join** to join the ISE node to Active Directory.

Figure 10-3 Active Directory Integration



Note The Cisco Identity Services Engine User Guide has detailed configuration steps: http://www.cisco.com/en/US/customer/docs/security/ise/1.2/user_guide/ise_user_guide.html.

Guest and Self-Registration Portals

The Cisco ISE server has the capability to host multiple portals. The BYOD system design relies on the Guest Portal to provide wireless guest access and, for provisioning purposes, the redirection of employees to the Self-Registration portal to on-board their devices. Chapter 21, "BYOD Guest Wireless Access" discusses the use of the Guest Portal for guest wireless access. The default ISE portals have standard Cisco branding that may be customized to identify unique portals for different purposes and with individual policies.

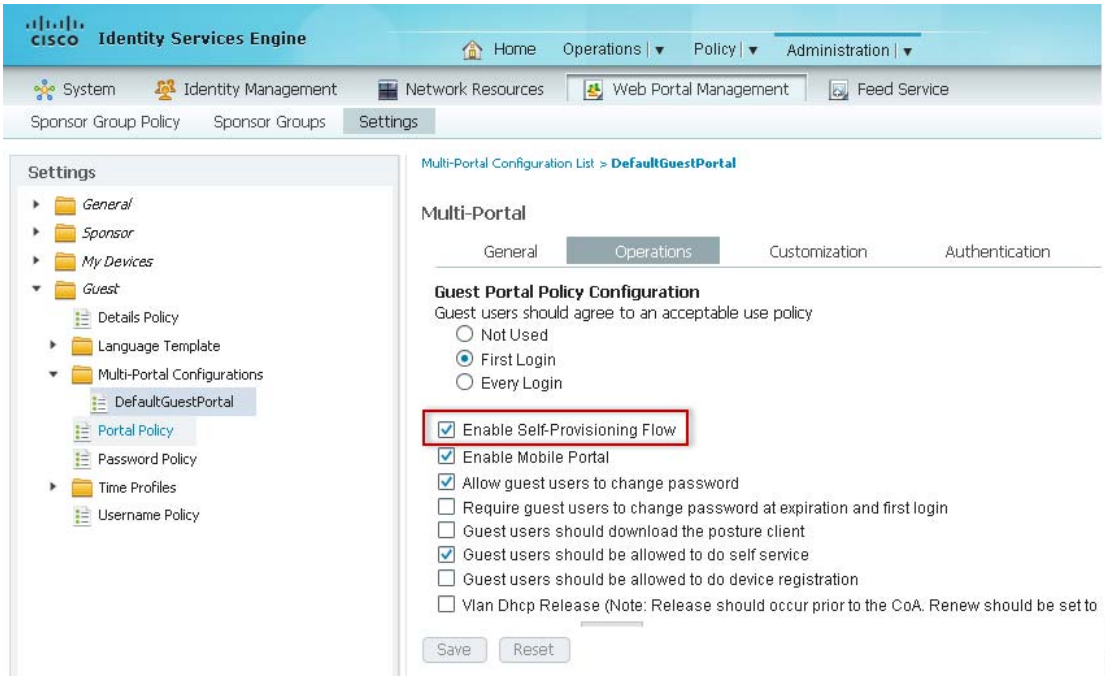
ISE enables self-provisioning, which allows employees to register their personal devices. The ISE provisions the device with its native supplicant during device registration.

The BYOD system leads the employee through the following provisioning steps the first time they bring their personal device to work and register:

1. The employee connects the device to the open SSID (BYOD_Provisioning SSID for dual SSIDs).
2. The device is redirected to the Guest Registration portal.
3. The employee enters credentials and ISE authenticates against Active Directory.
4. If the device is not yet registered on the network, the session is redirected to the self-registration portal.
5. The employee is asked to enter a unique device description and complete the device registration.

To enable Self-Provisioning, configure these portals as follows: click **Administration > Web Portal Management > Settings > Guest > Multi-Portal Configurations**, as shown in Figure 10-4.

Figure 10-4 Portal Settings—Operations

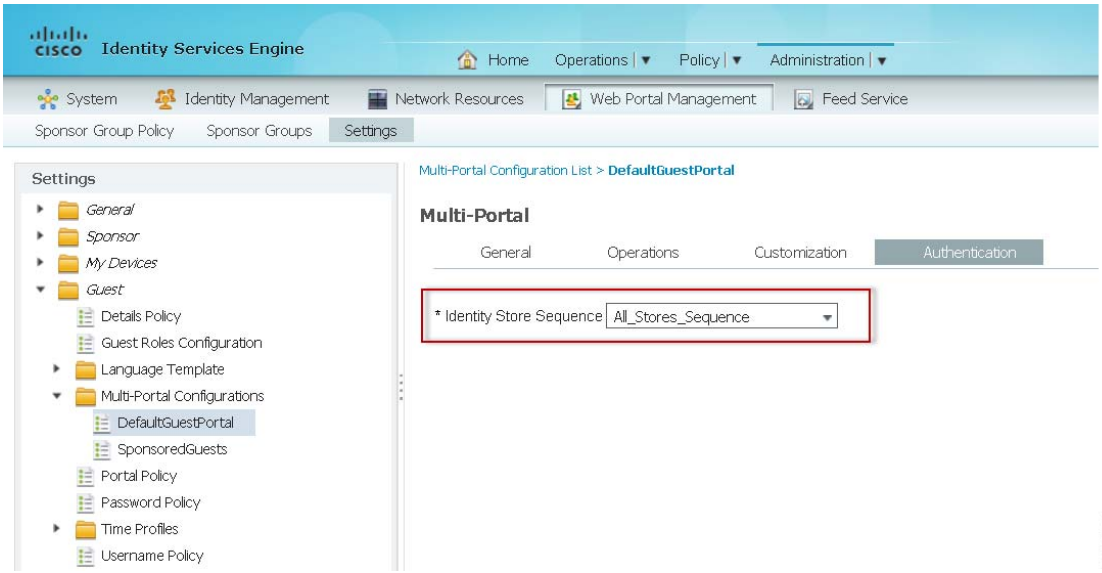


The DefaultGuestPortal refers to the portal used for self-registration—otherwise known as the Self-Registration portal in this document.

To specify how the portal authenticates users, select the Authentication tab within the particular portal, as shown in Figure 10-5, and select the appropriate option:

- Guest—The portal authenticates guest user accounts stored in the local database.
- Central WebAuth—The user is authenticated against the databases specified in the Identity Store Sequence.
- Both—The user is authenticated against a local guest database first. If the user is not found, authentication is attempted using additional databases defined in the Identity Store Sequence.

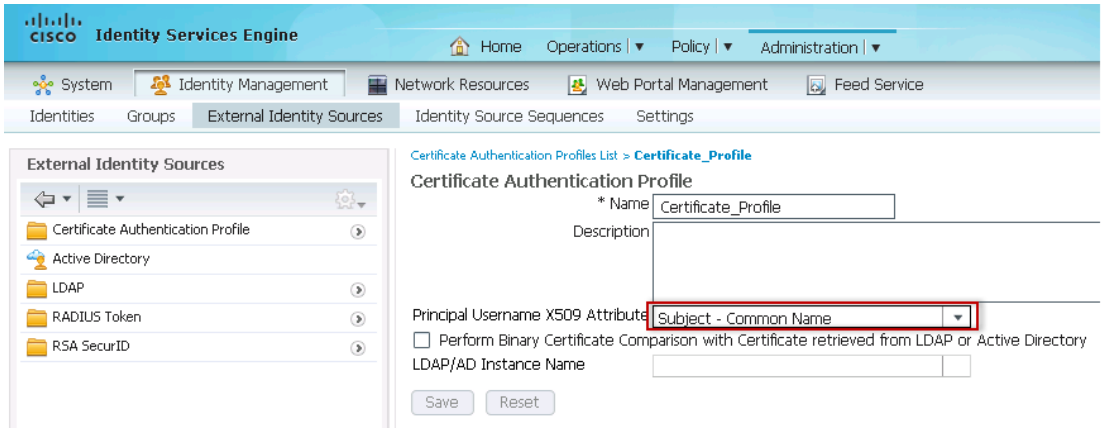
Figure 10-5 Authentication Portal Settings



ISE Using Certificates as an Identity Store

To configure ISE to use certificates as an identity store, choose **Administration > External Identity Sources > Certificate Authentication Profile > Add** and define the Certificate Authentication Profile, as shown in Figure 10-6.

Figure 10-6 Certificate Authentication Profile

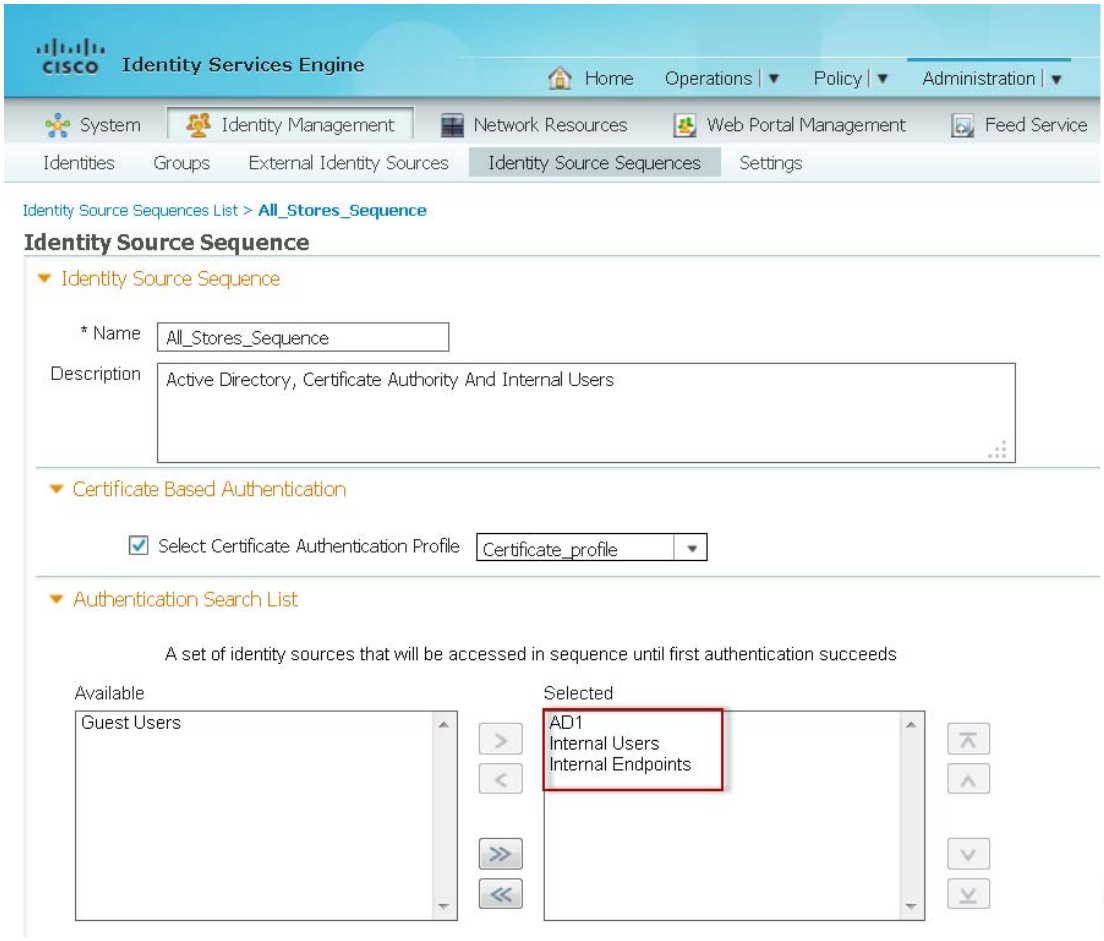


Identity Source Sequences

Identity Source Sequences define the order in which ISE will look for user credentials in the different databases. These databases include Internal Users, Active Directory, LDAP, RSA, etc.

To add a new Identity Source Sequence, click **Administration > Identity Source Sequences > Add**. The configuration shown in Figure 10-7 creates a new Identity Source Sequence named All_Stores_Sequence. It relies on Active Directory (AD1), a certificate profile named "Certificate_profile" and Internal Users.

Figure 10-7 Identity Source Sequence

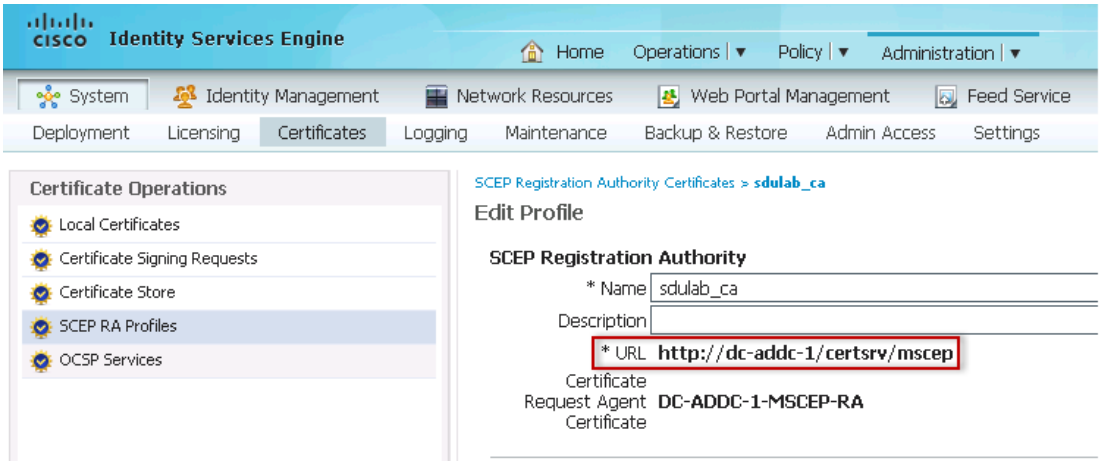


SCEP Profile Configuration on ISE

Within this design, ISE is acting as a Simple Certificate Enrollment Protocol (SCEP) proxy server, thereby allowing mobile clients to obtain their digital certificates from the CA server. This important feature of ISE allows all endpoints, such as iOS, Android, Windows, and MAC, to obtain digital certificates through the ISE. This feature combined with the initial registration process greatly simplifies the provisioning of digital certificates on endpoints.

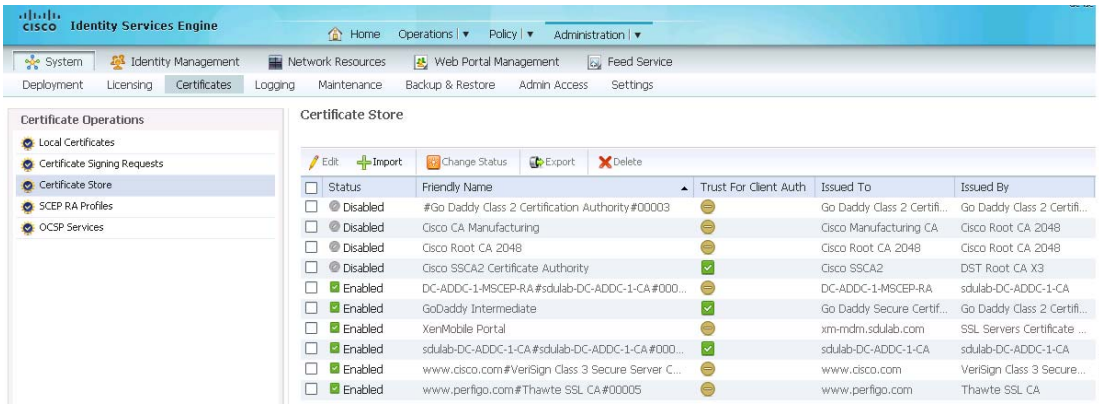
To configure SCEP profile on the ISE, click **Administration > Certificates > SCEP RA Profiles > Add**. Define the SCEP profile, as shown in Figure 10-8.

Figure 10-8 SCEP Profile Configuration



After the configuration is successful, ISE downloads the RA certificate and the root CA certificate of the CA server, as shown in Figure 10-9.

Figure 10-9 Certificate Store



Authentication Policies

Authentication policies are used to define the protocols used by the ISE to communicate with the endpoints and the identity sources to be used for authentication. ISE evaluates the conditions and based on whether the result is true or false, it applies the configured result. An authentication policy includes:

- An allowed protocol service, such as PEAP, EAP-TLS, etc.
- An identity source used for authentication

Similar to the way access lists are processed, authentication rules are processed from the top down. When the first condition is met, processing stops and the assigned identity rule is used.

The rules are evaluated using “If, then, else” logic:

```
IF Wired_802.1X   Then
     Allow default protocols
Elseif   next condition
     Take action
```

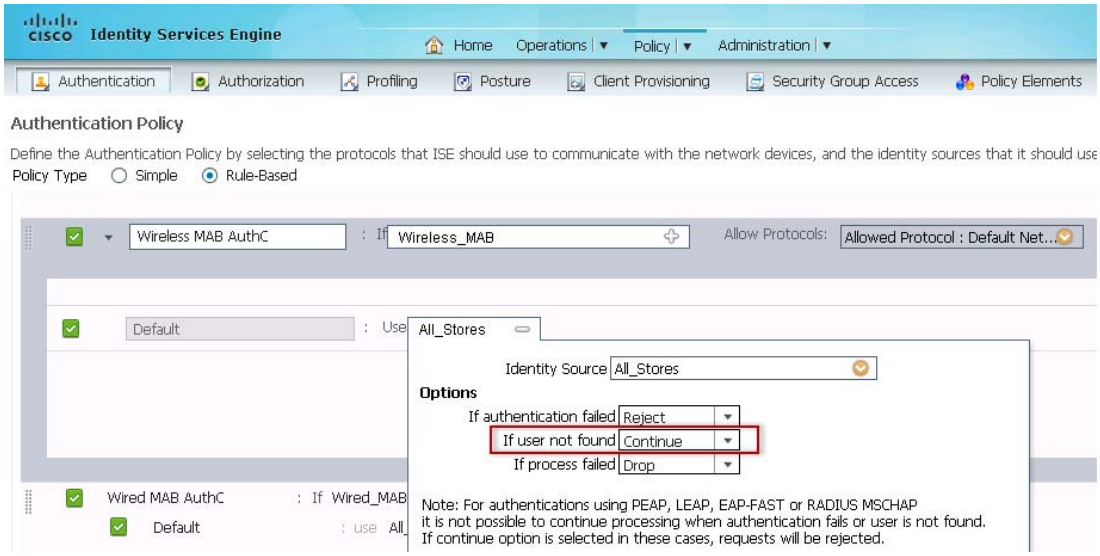
```
Else
     Use Default Rule
```

In BYOD designs discussed throughout this document, ISE authenticates several protocols such as MAB and dot1x against all the Identity Stores. The Identity Stores could be AD, Certificate_Profile, RSA, Internal Users, and Internal Endpoints. The network access medium could be wired, wireless, or remote connection. The network device uses any of the mediums mentioned before, using different protocols to connect to ISE.

MAC Authentication Bypass (MAB) protocol is used to authenticate devices not configured with dot1x. When a brand new device accesses the network it communicates via the MAB protocol and uses its own MAC address as its identity. In a normal scenario, ISE would validate if the MAC address is present in any of its identity stores; if not, it would reject the connection. However in this BYOD design the MAB protocol is used by new devices for on-boarding purposes and it may not be feasible to know the MAC address of the device in advance.

To circumvent this problem, ISE continues the authentication process and redirects the device to the next stage, even if the device's MAC address is not present in any of its identity stores. Figure 10-10 highlights this configuration.

Figure 10-10 Authentication Rule for MAB



In a normal deployment scenario, the endpoints would primarily use the dot1x protocol to communicate with ISE. ISE authenticates these endpoints against an Active Directory or authenticates them via digital certificates. Figure 10-11 depicts the different protocols and how these protocols use different identity stores for authentication.

Figure 10-11 Authentication Policy

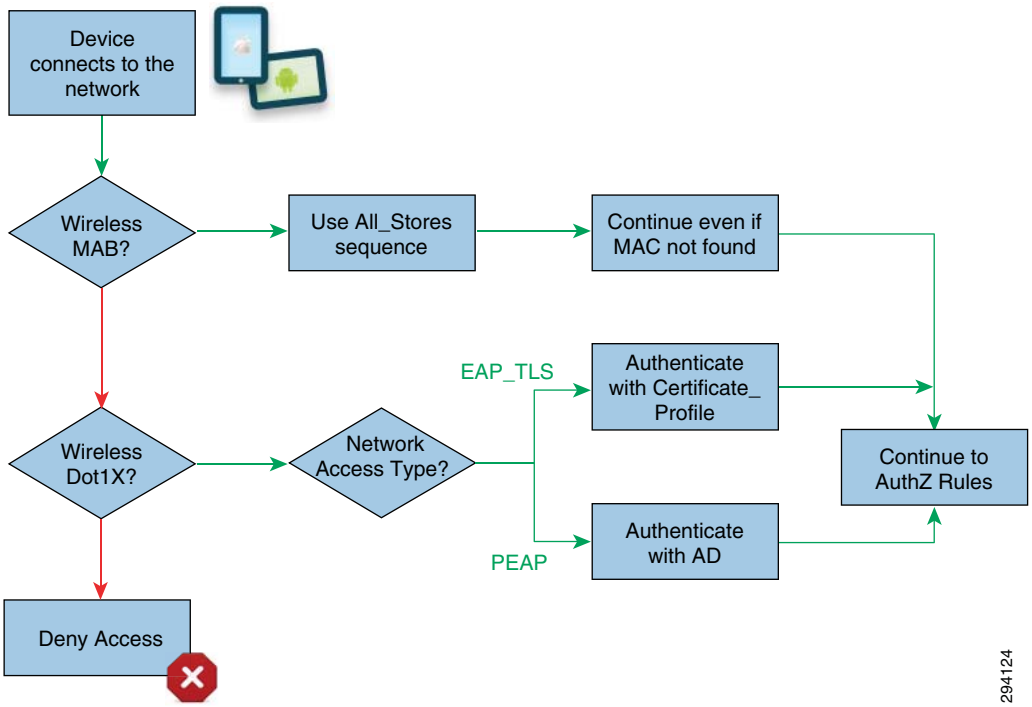


Table 10-1 explains how these rules are implemented in this design guide.

Table 10-1 Authentication Rules

Rule Name	Network Access Medium	Allowed Protocols	Conditions		Identity Store
Wireless MAB AuthC	Wireless MAB	All	Default		All_Stores
Wired MAB AuthC	Wired MAB	All	Default		All_Stores
Wireless Dot1X AuthC	Wireless_8021X	All	Wireless Certificate	EAP_TLS	Certificate_Profile
			Wireless Password	PEAP	All_Stores
Wired Dot1X AuthC	Wired_802.1X	All	Wired Certificate	EAP_TLS	Certificate_Profile
			Wired Password	PEAP	All_Stores
Default					Deny Access

Authentication Policy for Wireless

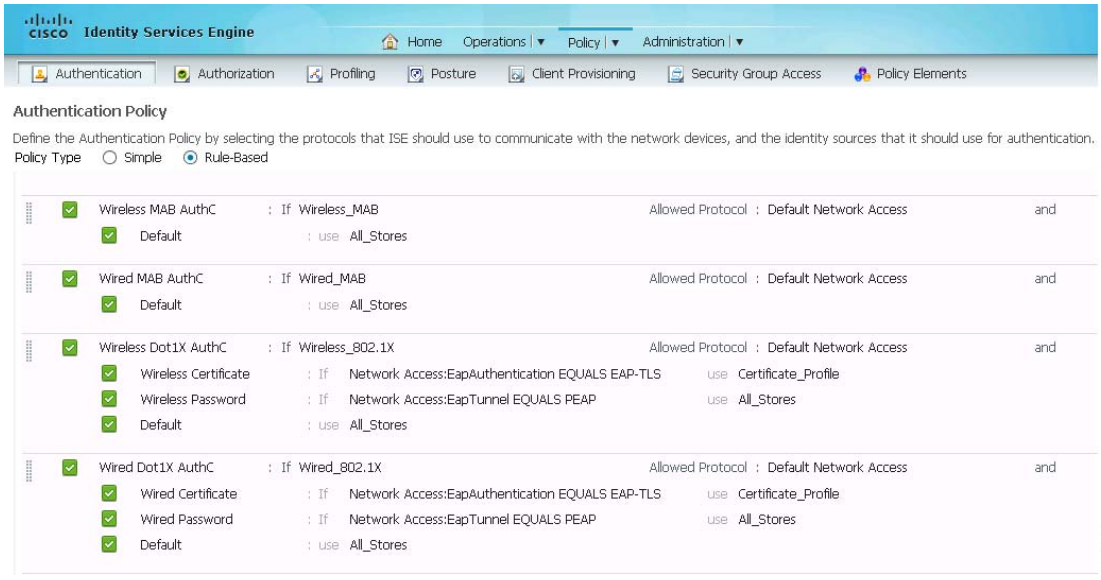
The endpoint devices could use either MAB or dot1x protocol when connecting to the wireless network. The authentication policy for wireless networks using MAB is explained in the previous section. This section explains the authentication policy for wireless medium using dot1X protocol, as shown in Table 10-1.

Wireless Dot1X AuthC is the rule name for wireless_dot1x protocol. This rule matches wireless_dot1x protocol and has two inner rules:

- Wireless Certificate—Matches when the authentication protocol is EAP_TLS and it verifies the digital certificate using the identity store Certificate_Profile.
- Wireless Password—Matches on the PEAP authentication protocol and uses the All_Stores identity store, which includes Active Directory.

Figure 10-12 shows how these rules were configured on the ISE for this design guide.

Figure 10-12 Authentication Rules



Client Provisioning

The Cisco ISE looks at various elements when classifying the end user's device type, including operating system version, browser type, etc. Once the ISE classifies the client machine, it uses client provisioning resource policies to ensure that the client is configured with an appropriate agent version, up-to-date compliance modules and correct agent customization packages and profiles, if necessary. The ISE Profiling service is discussed in Enabling the DHCP and RADIUS Probes. It is important to understand the difference between Client Provisioning Policy and Client Provisioning Resources. Client Provisioning Resources are basically the resources that are pushed to the end device and assist the end device in completing the on-boarding process. Client Provisioning Resources are of two types:

- Native profiles that can be configured on ISE; for example, iOS profile.
- Software Provisioning Wizards that must be downloaded from Cisco site.

Client Provisioning Policy on the other hand links an endpoint device to an appropriate Client Provisioning Resource. Therefore the Client Provisioning Resources must be added to the ISE before configuring the Client Provisioning Policy. This section discusses Client Provisioning Resources and Client Provisioning Policies for iOS, Android, Windows and Mac OS X devices.

The following are considerations for client provisioning on the endpoints:

- Based on the endpoint, push an appropriate Software Provisioning Wizard (SPW) to the device. This Wizard configures the dot1x settings on the endpoint and configures the endpoint to obtain a digital certificate.

- In certain endpoints such as iOS devices, there is no need for SPW package because for iOS devices the native operating system is used to configure the dot1x settings.
- For Android devices, the SPW package needs to be downloaded from Google Play Store.

Client Provisioning Resources—Apple iOS and Android

To configure a client provisioning resource for mobile devices, click **Policy > Policy Elements > Results > Client Provisioning > Resources > Add Native Supplicant Profile**. Figure 10-13 shows the configuration details for the Wireless iOS TLS profile used by Apple iOS devices. This profile is used to configure the parameters required to access to the BYOD_Employee SSID after on-boarding.

Figure 10-13 Wireless iOS TLS Profile

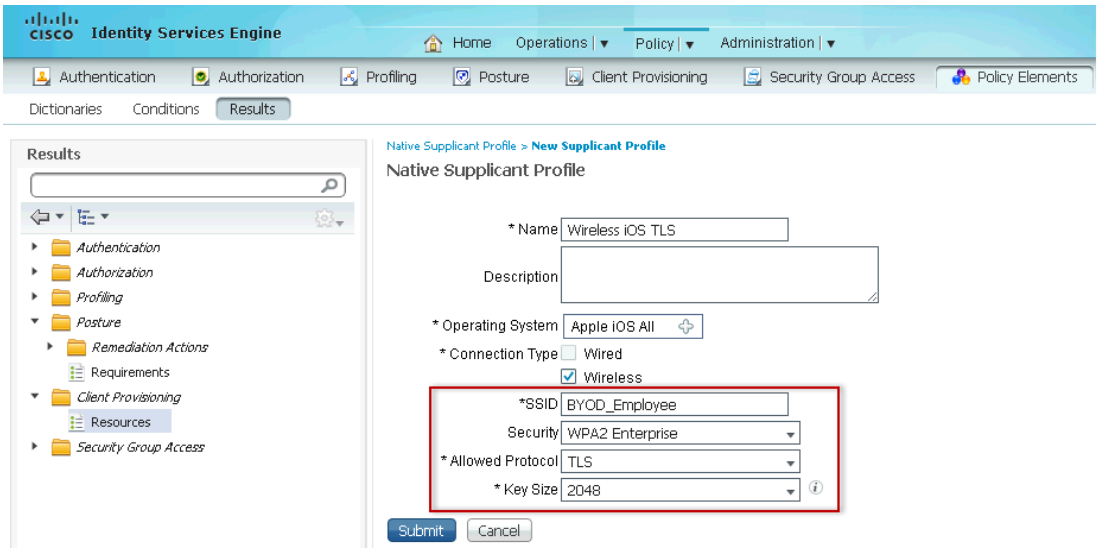
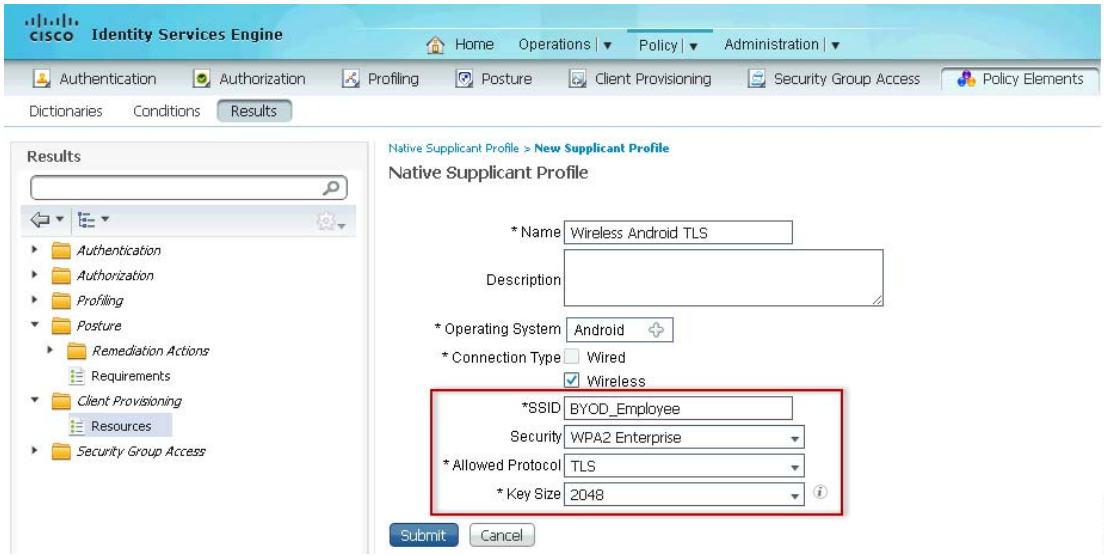


Figure 10-14 shows the configuration details for the Wireless Android TLS profile used by Android devices.

Figure 10-14 Wireless Android TLS

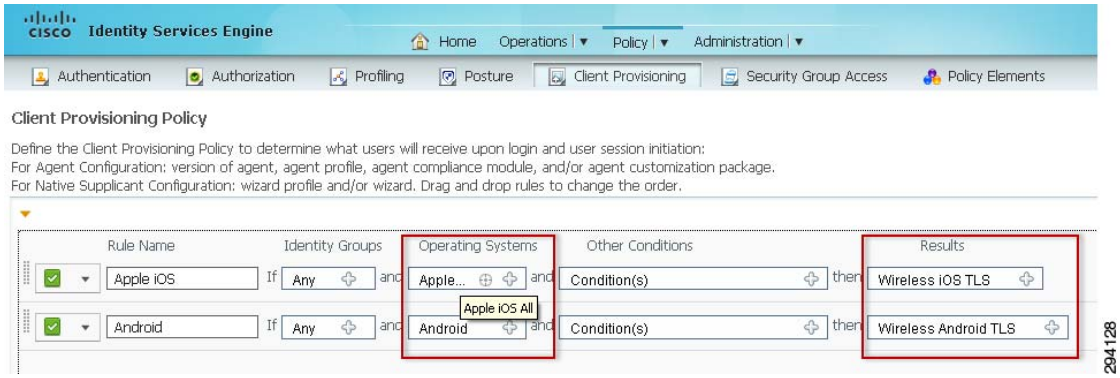


Client Provisioning Policy—Apple iOS and Android Devices

Client provisioning policies determine which users receive which version of resources. After defining the Native Supplicant Profile, the next step is to use the appropriate profile when devices connect to the network by clicking **Policy > Client Provisioning**.

The configuration in Figure 10-15 determines the operating system running on the device and defines which resources to distribute. In this case the previously defined profiles are distributed based on the appropriate operating system.

Figure 10-15 Client Provisioning Policies



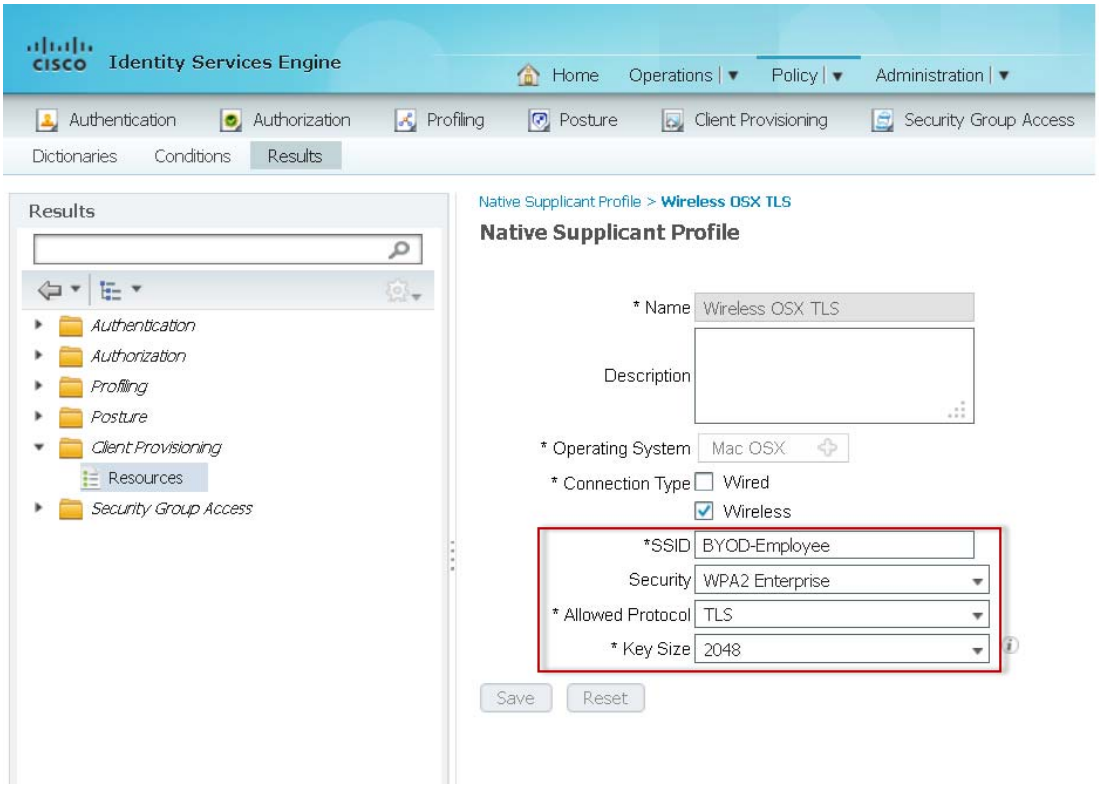
It is important to note that for Android devices the user is also required to download the software from Google's Play Store, since it cannot be distributed by ISE.

Client Provisioning Resources—Mac OS

For MAC OS workstations, the following is required:

- A Native Supplicant profile that determines what kind of configuration should be provisioned on the device, for example the Wireless SSID name. Figure 10-16 shows the native supplicant profile for Mac OSX devices.

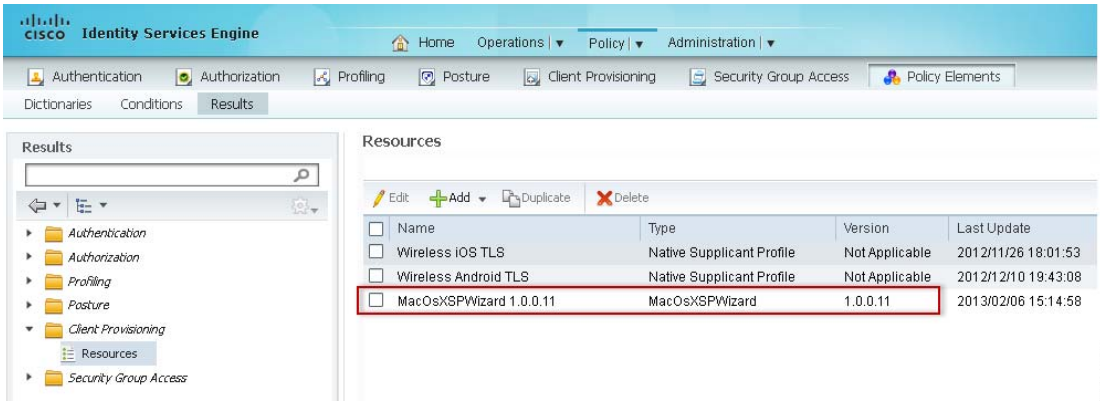
Figure 10-16 Native Supplicant Profile for Mac OSX Devices



- A Wizard Profile—The Supplicant Provisioning Wizard profile is a software agent that may be downloaded from Cisco.

To define the client provisioning resources, click **Policy > Policy Elements > Results > Client Provisioning > Resources > Add > Agent Resources** from the Cisco site and select the **MacOsXSPWizard**. Figure 10-17 shows the MacOsXSPWizard profile.

Figure 10-17 Mac OsXSPWizard Profile



Client Provisioning Policy for Mac OS Devices—Wireless

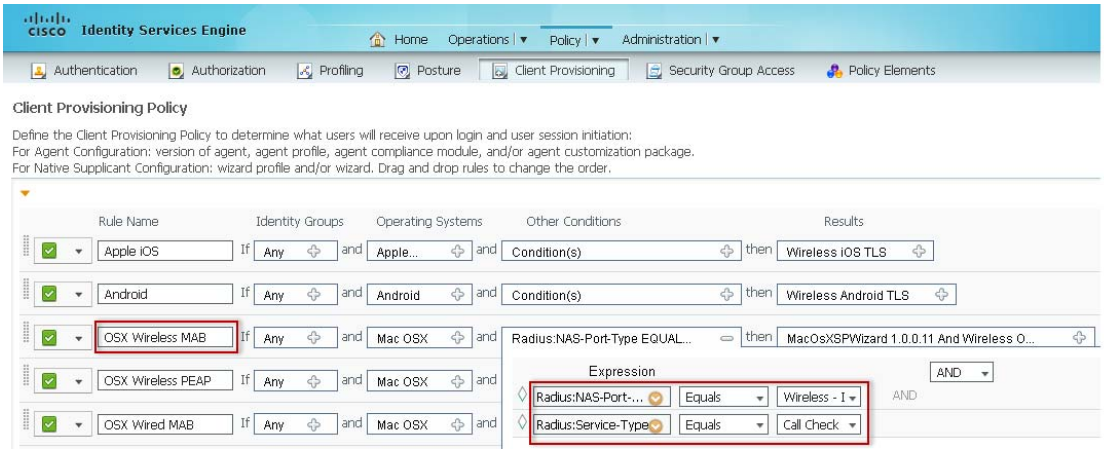
The previous section discussed the resources needed for provisioning Mac OS devices. Once the resources have been configured, the next step is to define under what conditions these resources will be used. The Mac OS X devices can use either MAB or PEAP protocol during the provisioning process. Therefore different conditions have to be configured to match either one of them.

The MAB protocol is matched by the following two conditions:

- RADIUS:NAS-Port-Type EQUALS Wireless—IEEE 802.11
- RADIUS:Service-Type EQUALS Call Check

Figure 10-18 shows the Client Provisioning Policy to match on the MAB protocol.

Figure 10-18 Client Provisioning Policy for MAB

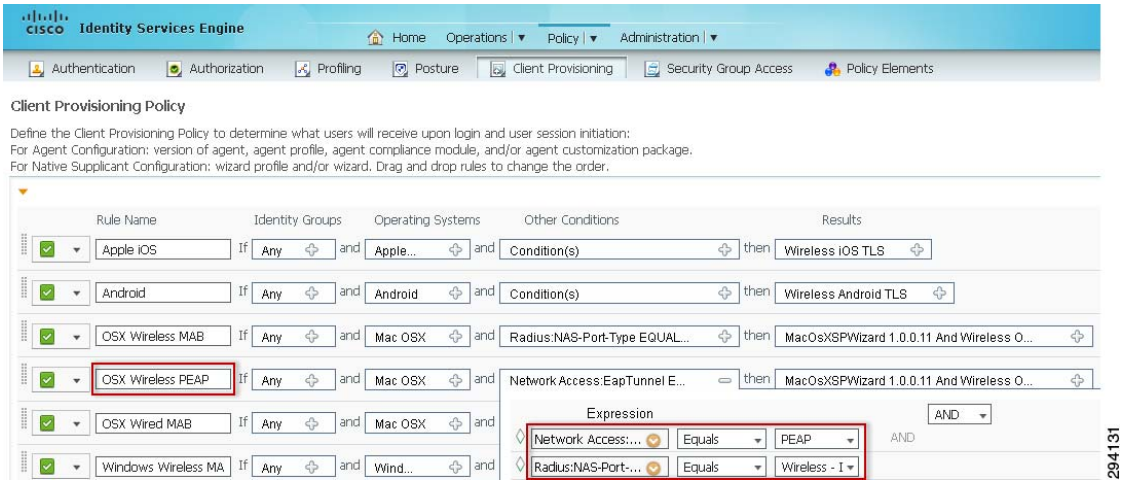


To match a Mac device using the PEAP protocol, the following conditions are needed:

- RADIUS:NAS-Port-Type EQUALS Wireless—IEEE 802.11
- Network Access:EapTunnel EQUALS PEAP

Figure 10-19 shows the condition to match on MAC devices using the PEAP protocol.

Figure 10-19 Client Provisioning Policy for PEAP

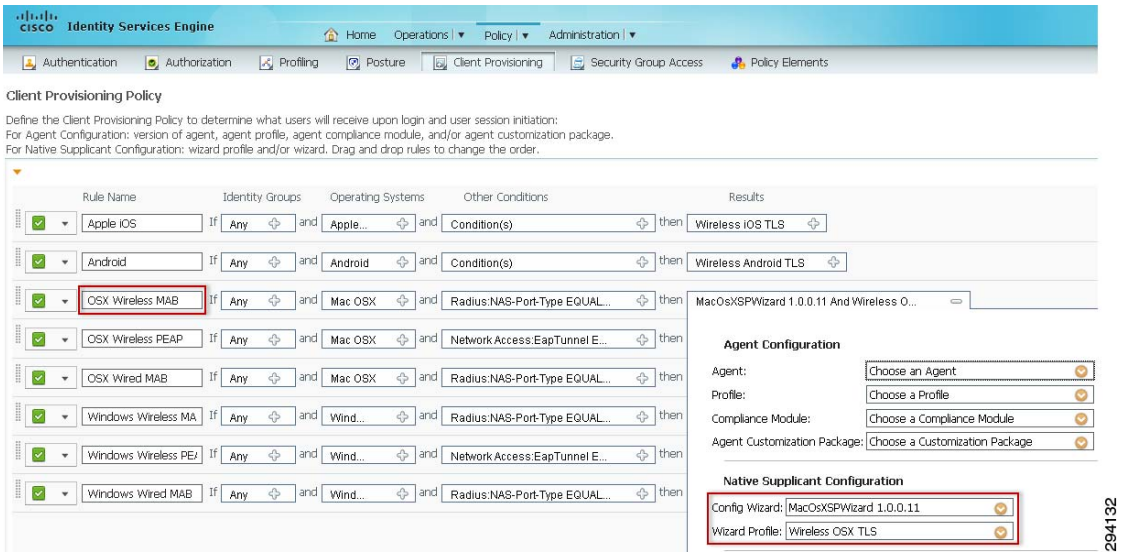


To complete a Client Provisioning policy for MAC_OSX_Wireless devices, the following must be defined:

- The Operating System must be selected as Mac OSX.
- The Conditions should be used to match either MAB or PEAP protocol.
- The result section must contain the Native Supplicant profile and the SPW for Mac OS X devices.

The complete policy is shown in Figure 10-20.

Figure 10-20 Client Provisioning Policy for Mac OS X



Client Provisioning Policy for Windows Devices—Wireless/Wired

The configuration steps for defining the provisioning policy for Windows devices is very similar to Mac OS X or iOS devices, so the same configuration steps are not repeated here. The only difference to point out is that for Windows devices a different SPW package is needed. Figure 10-21 depicts the Client Provisioning Policy for Windows (wireless or wired) devices using either MAB or PEAP.

Figure 10-21 Client Provisioning Policy for Windows

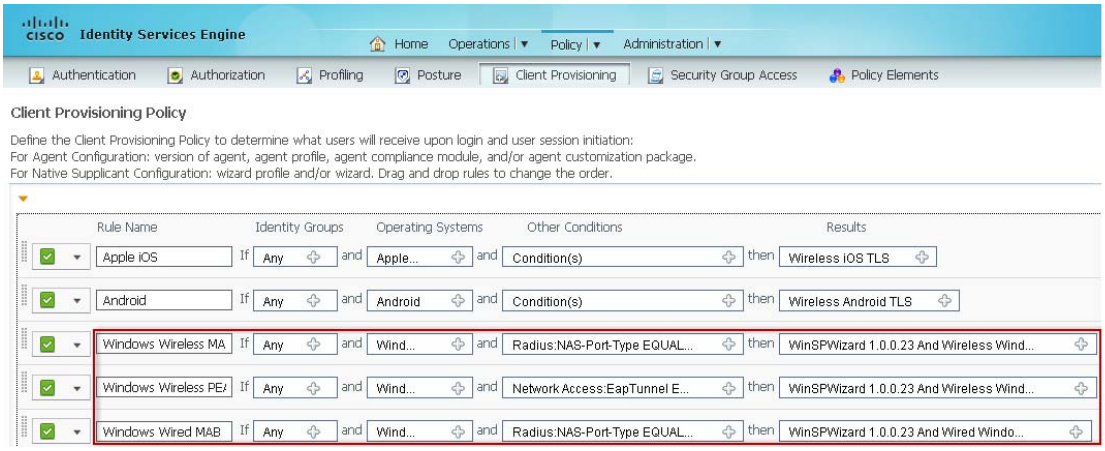
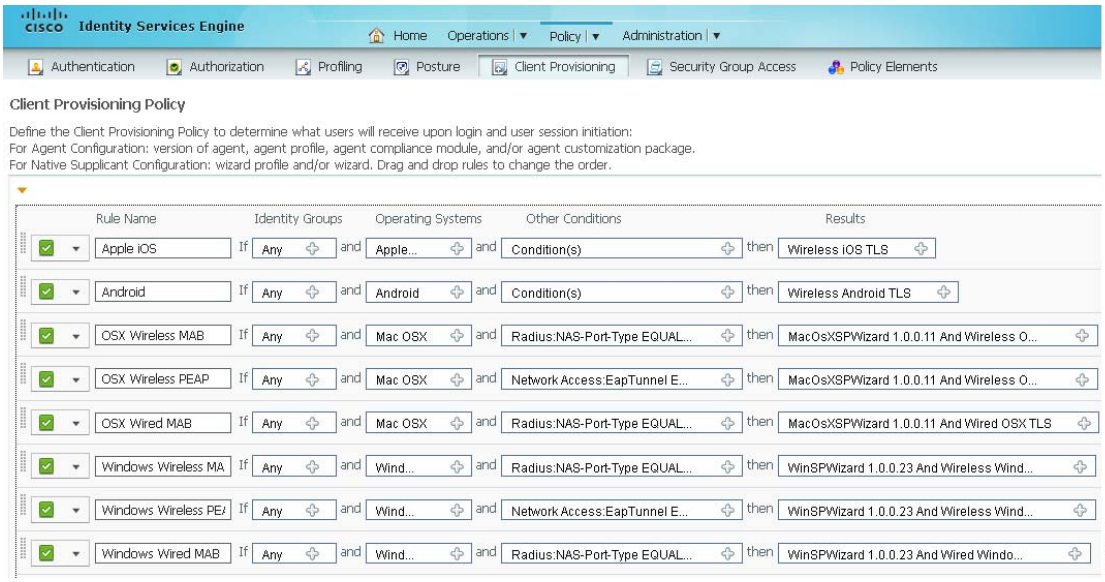


Figure 10-22 shows the complete client provisioning policy used during testing.

Figure 10-22 Complete Client Provisioning Policy



Profiling

Profiling is a key service responsible for identifying, locating, and determining the capabilities of endpoints that attach to the network to deny or enforce specific authorization rules. Two of the main profiling capabilities include:

- Collector—Used to collect network packets from network devices and forward attribute values to the analyzer.
- Analyzer—Used to determine the device type by using configured policies that match attributes.

There are two main methods to collect endpoint information:

- The ISE acting as the collector and analyzer.
- Starting in version 7.3, the WLC can act as the collector and send the required attributes to the ISE, which acts as the analyzer.

Client profiling from a controller running 7.3 or later is supported on access points that are in Local mode and FlexConnect mode. Table 10-2 shows the main differences between the WLC and ISE profiling.

Table 10-2 ISE versus WLC Profiling Support

ISE	WLC
Profiling using a large number of probes, including RADIUS, DHCP, DHCP SPAN, HTTP, DNS, etc.	DHCP and HTTP based profiling only
ISE supports as policy action multiple different attributes	WLC supports VLAN, ACL, session timeout, QoS
Profiling rules may be customized with user-defined attributes	Only default profiling rules may be used

Note This design guide uses the profiling capabilities of the ISE and did not test the controller client profiling capabilities.

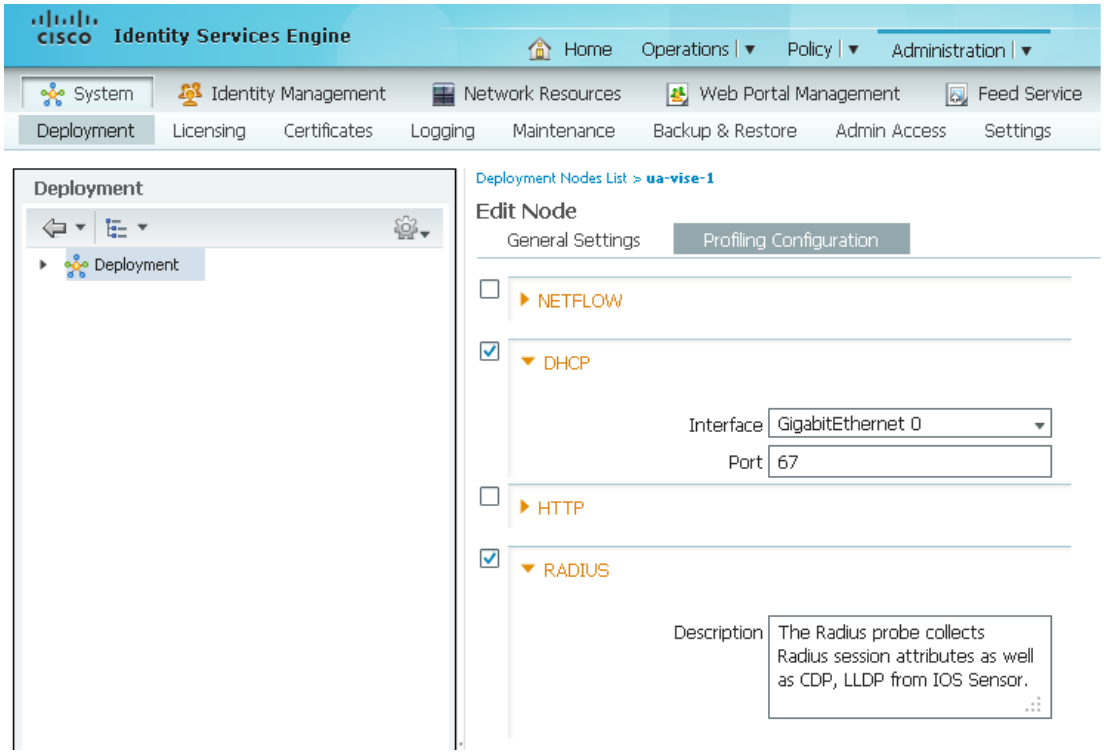
The ISE supports a number of sensors to capture endpoint attributes and classify them according to their profiles. The sensors rely on a number of probes that capture network packets by querying network access devices. Once the endpoints are profiled, different authentication and authorization policies may be enforced. Some examples of using different policies based on the device's profile include:

- Allow employee-owned iPads to access the network, but only for HTTP traffic.
- If the iOS device connecting to the network is a company-owned device, grant full access to the network.
- If an employee-owned iPad has been provisioned with a digital certificate, grant full access to the network.
- Force some devices to register with their Mobile Device Manager.
- Deny access to all iPads or Android devices.

Enabling the DHCP and RADIUS Probes

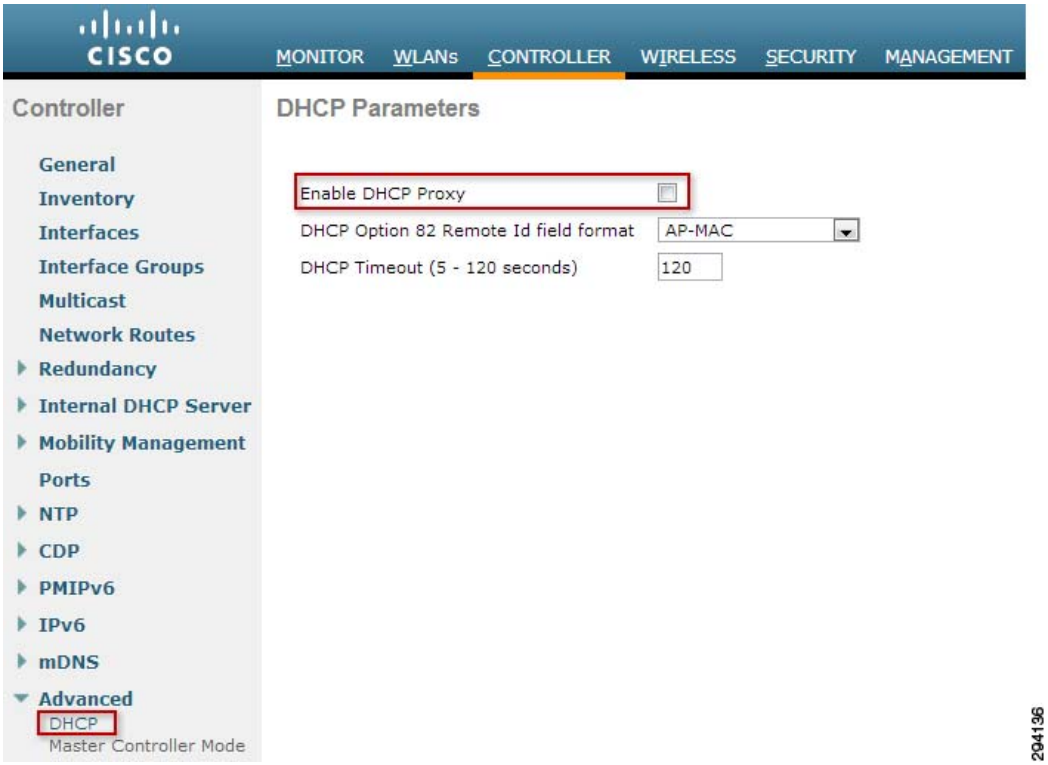
To enable profiling on the ISE, click **Administration > System > Deployment**. Click the ISE hostname and click **Profiling Configuration**. Enable the appropriated probes to listen to packets forwarded from the LAN switch or Wireless LAN Controller, as shown in Figure 10-23.

Figure 10-23 Profiling Probes



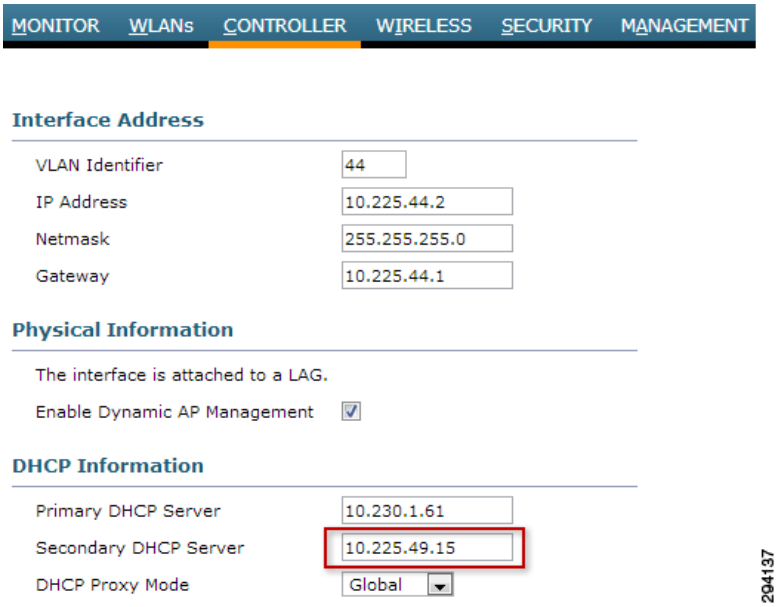
The Wireless LAN Controller should be configured in DHCP bridging mode to forward DHCP packets from the wireless endpoints to the ISE. Click **Controller > Advanced > DHCP** and clear the Enable DHCP Proxy check box, as shown in Figure 10-24.

Figure 10-24 Disable DHCP Proxy



Specify the ISE's IP address as the secondary DHCP server in the WLC by clicking **Controller > Interfaces > Secondary DHCP**, as shown in Figure 10-25.

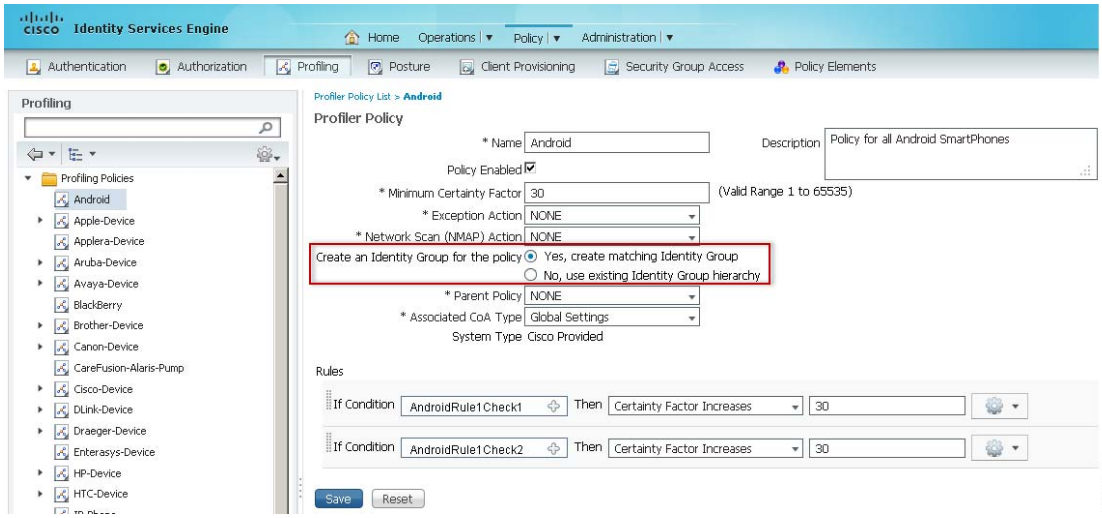
Figure 10-25 Secondary DHCP Server



Profiling Android Devices

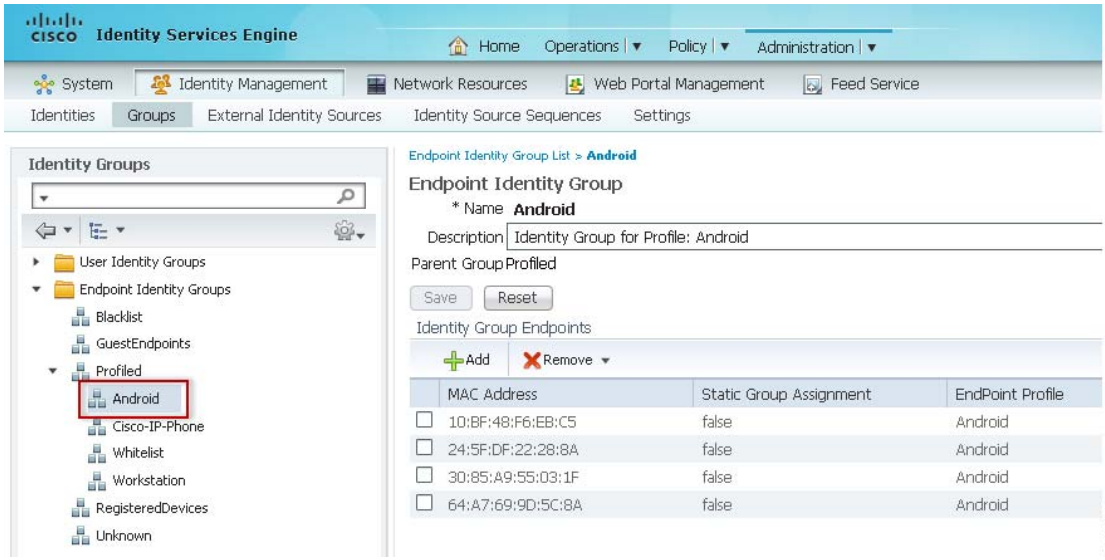
To create an identity group based on the Android policy, click **Policy > Profiling > Profiling Policies > Android** and enable the Create Matching Identity Group, as shown in Figure 10-26.

Figure 10-26 Android Profiling Policy



The Android profiling policy should be listed under Endpoint Identity Groups > Profiled. Click **Administration > Identity Management > Groups** to see a list of Android devices that have been profiled by the ISE, as shown in Figure 10-27.

Figure 10-27 Android Identity Group



Logical Profiles

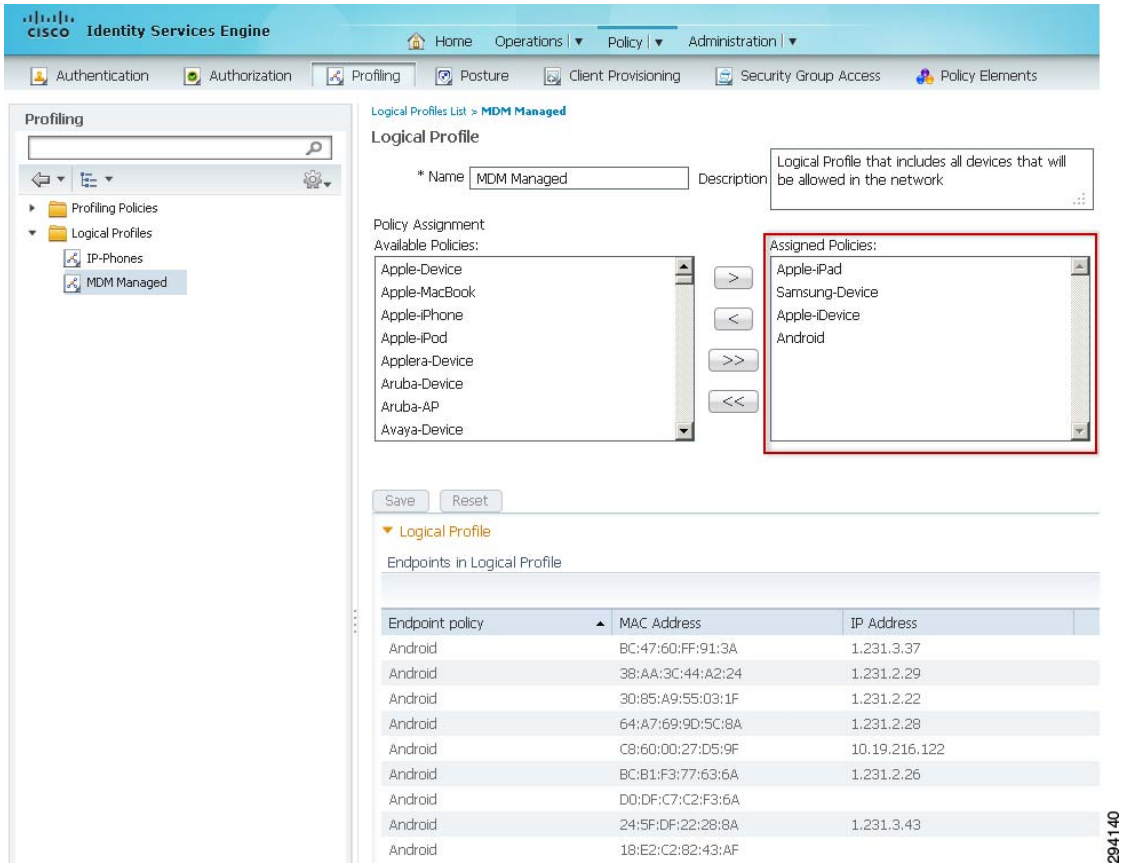
Logical profiles are containers that group different profiles to create an overall category of profiles. Logical profiles provide additional flexibility to the authorization policies, enhancing the overall network access policy.

With logical profiles, a single entry in the authorization rule is able to include several profiles. Before logical profiles were available, a matching identity groups had to be created for each device type.

In this design guide, a logical profile was created to group the mobile devices that are managed by the MDM. This profile combines some mobile devices into a single logical profile that may be invoked from the authorization rules.

To create a logical profile, click **Policy > Profiling > Profiling > Logical Profiles**, as shown in Figure 10-28.

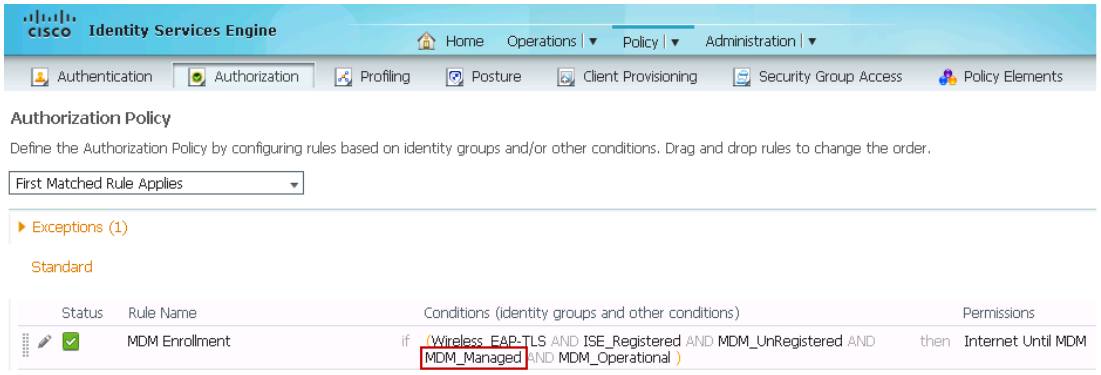
Figure 10-28 MDM Managed Logical Profile



This logical profile provides the flexibility to add new devices at any time without modifying the authorization rules. Figure 10-29 shows how the MDM Managed Logical Profile is used to identify devices supported by the MDM.

This and other authorization rules are explained in more detail later in this design guide.

Figure 10-29 MDM Enrollment Authorization Rule



Authorization Policies and Profiles

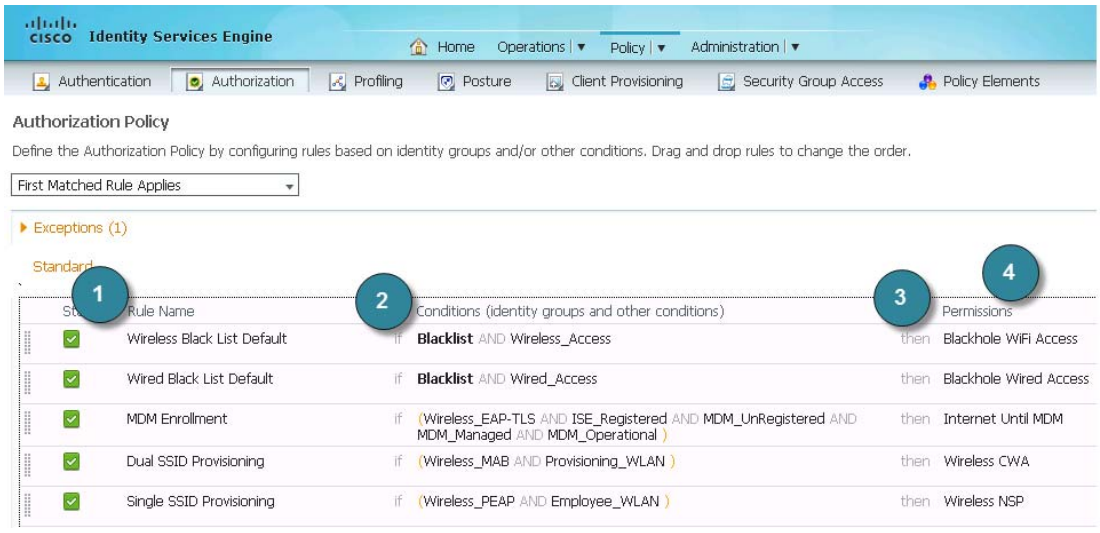
Authorization policies define the overall security policy to access the network. Network authorization controls user access to the network and its resources and what each device can do on the system with those resources. An Authorization Policy is composed of multiple rules.

Authorization rules are defined by three main elements, as shown in Figure 10-30:

- Names (1)
- Conditions (2)
- Permissions (3)
- Authorization Profiles (4)

Permissions are enforced by authorization profiles (4). Similar to the authentication rules, authorization rules are processed from the top down. When the first condition is met, processing stops and the assigned permission dictates what authorization profile to use.

Figure 10-30 Authorization Policy



Authorization Profiles

An authorization profile acts as a container where a number of specific permissions allow access to a set of network services. The authorization profile is where a set of permissions to be granted is defined and can include:

- An associated VLAN.
- An associated downloadable ACL (DACL).
- Wireless LAN Controller attributes such as the use of a Named ACL or Security Group Tag for policy enforcement.
- Advanced settings using attributes contained in dictionaries.

In addition to the standard PermitAccess and DenyAccess authorization profiles, the following are some of the profiles that are defined within this design guide:

- Wireless CWA—This profile is used for redirection of wireless devices to the registration portal for devices using MAB and dual SSIDs.
- Wireless NSP—This profile is used to redirect wireless users to the registration portal when they access the network using dot1x or a single SSID.
- Blackhole WiFi Access—Used to block access to devices reported lost (for more information, see Chapter 22, "Managing a Lost or Stolen Device").

Several other authorization profiles are explained in other chapters of this design guide.

Note Cisco has been made aware of potential incompatibilities introduced by Apple iOS 7. We are working to understand the limitations and design updates will be made to this publication.

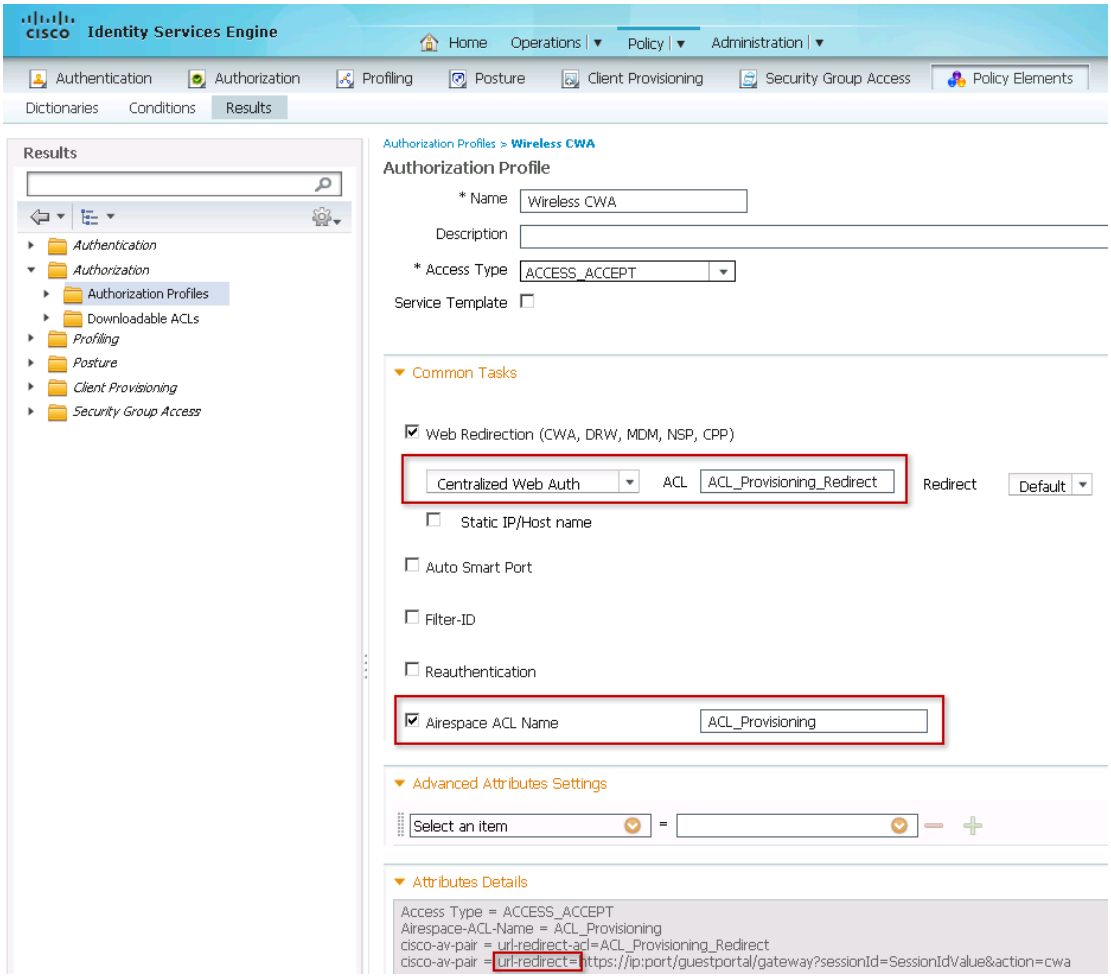
Wireless CWA Authorization Profile for Dual SSID Provisioning

This policy is used in dual SSID configurations to redirect wireless devices to the Self-Registration portal upon connecting to the network. This authorization profile restricts access by triggering the ACL_Provisioning_Redirect access list, which is defined in advance in the Wireless LAN Controller.

When implementing dual SSIDs, the provisioning SSID can be either open or password-protected with Active Directory credentials. In this design guide, the provisioning SSID is open and relies on MAC Authentication Bypass (MAB) to grant access to the network.

To configure this authorization policy, click **Policy > Policy Elements > Results > Authorization Profiles**, as shown in Figure 10-31.

Figure 10-31 Wireless CWA Authorization Profile



To force devices to the self-registration portal, a redirect URL is created with a unique Session ID and pushed to the device:

https://ip:port/guestportal/gateway?sessionId=SessionIdValue&action=cwa

When the user launches a web browser, the device is redirected to the Self-Registration portal. To prevent the user from staying connected to the provisioning SSID, the ACL_Provisioning_Redirect ACL only permits access to the Cisco ISE, DHCP, and Domain Name System (DNS) services.

The Wireless CWA authorization profile relies on two named ACLs previously defined in the Wireless LAN Controller:

- ACL_Provisioning_Redirect—Applied to the Centralized Web Auth setting.
- ACL_Provisioning—Sent to the wireless controller via the Radius:Airespace-ACL-Name attribute value (AV).

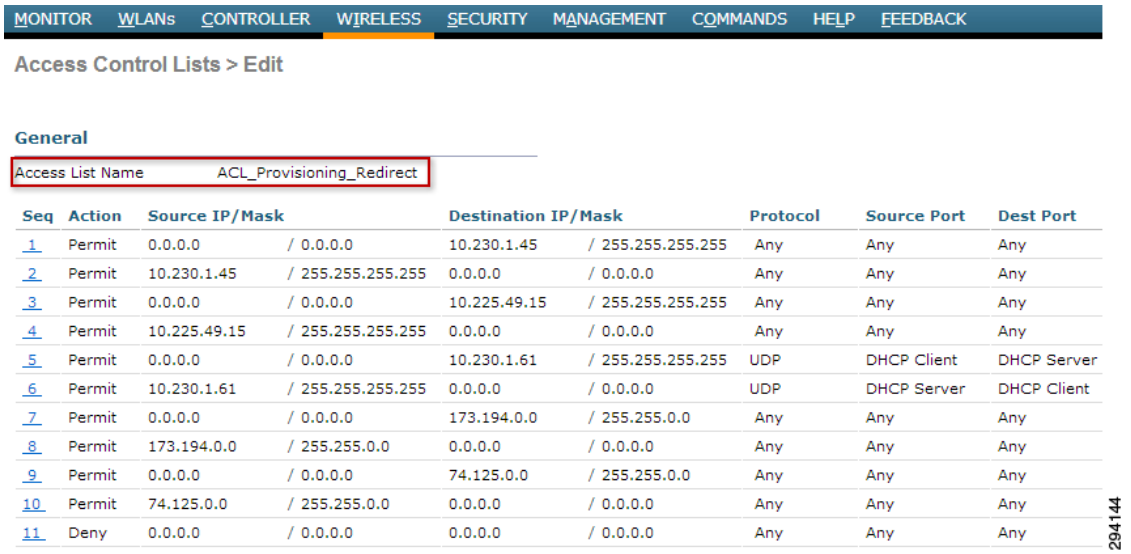
The behavior of the two ACLs is slightly different between wireless controllers:

- For CUWN wireless controllers (e.g., CT5508 and Flex 7500), ACL_Provisioning_Redirect functions as both the ACL which controls web redirection and as the ACL which controls access on the network. ACL_Provisioning serves simply as an extra security configuration and is not used when URL redirection is specified. For CUWN wireless controllers the ACL_Provisioning _Redirect ACL shown in Figure 10-32 can be the same as the ACL_Provisioning.

- For Cisco IOS XE based wireless controllers (e.g., CT5760 and Catalyst 3850), ACL_Provisioning_Redirect functions strictly as the ACL which controls web redirection. ACL_Provisioning functions as the ACL, which controls what the wireless client is allowed to access on the network. Hence IOS XE based wireless controllers make use of both ACLs when URL redirection is specified.

Figure 10-32 displays the configuration for ACL_Provisioning_Redirect on the WLC. This is just an example, since each organization will have unique business policies and security requirements.

Figure 10-32 WLC Access List for Provisioning



The ACL_Provisioning_Redirect ACL specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Access to Google Play.

Note Android devices require access to the Google Play Store to download the SPW package. Modify the ACL to allow endpoints to download the SPW. Analyzing the DNS transactions between the DNS server and the device is one approach to develop and troubleshoot ACL_Provisioning_Redirect.

On the Catalyst 3850 or the CT5760 Controller, the ACL_Provisioning_Redirect is defined as follows:

```
ip access-list extended ACL_Provisioning_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
 permit tcp any any eq www
 permit tcp any any eq 443
```

The ACL_Provisioning_Redirect ACL specifies the following access:

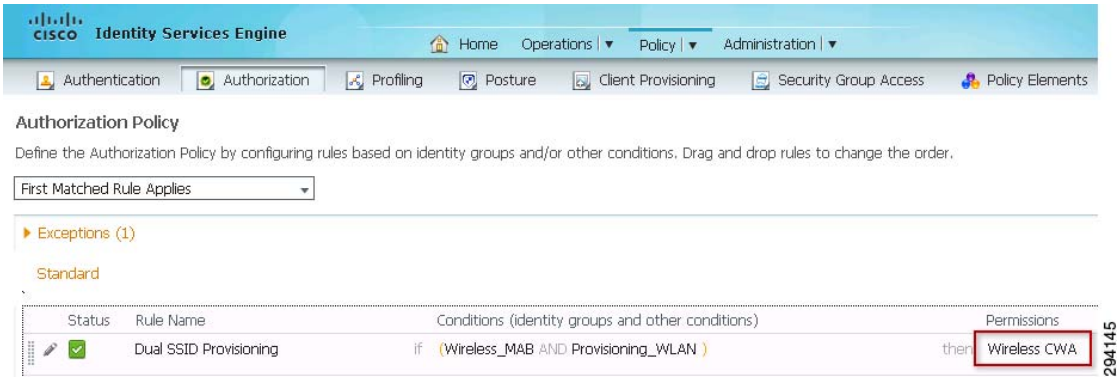
- Deny (do not redirect) IP access to and from the DNS server (10.230.1.45).

- Deny (do not redirect) IP access to and from the ISE Server (10.225.49.15).
- Deny (do not redirect) DHCP Access (bootpc and bootps).
- Permit (redirect) TCP access to any web host.
- Permit (redirect) TCP access to any secure web host.
- Deny (do not redirect) all other access to the Internet.

Dual SSID Provisioning Authorization Rule

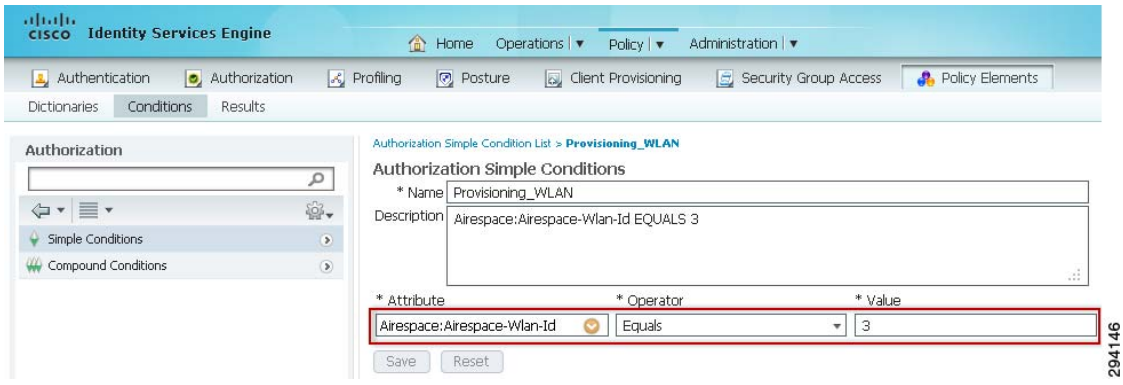
The Dual SSID Provisioning rule links the Wireless CWA authorization profile to the conditions that authorize MAB devices into the Provisioning SSID, as shown in Figure 10-33. It includes two conditions: Wireless_MAB and Provisioning_WLAN.

Figure 10-33 Dual SSID Authorization Rule



The Wireless_MAB condition is a predefined condition in ISE, while the Provisioning_WLAN condition was defined from the menu **Policy > Conditions > Simple Conditions**, as shown in Figure 10-34.

Figure 10-34 Provisioning_WLAN Condition



For the purposes of this CVD, the BYOD_Provisioning SSID number was defined as 3 during testing. The simple condition Provisioning_WLAN matches when the SSID number is 3. The condition is created to improve readability of the rules.

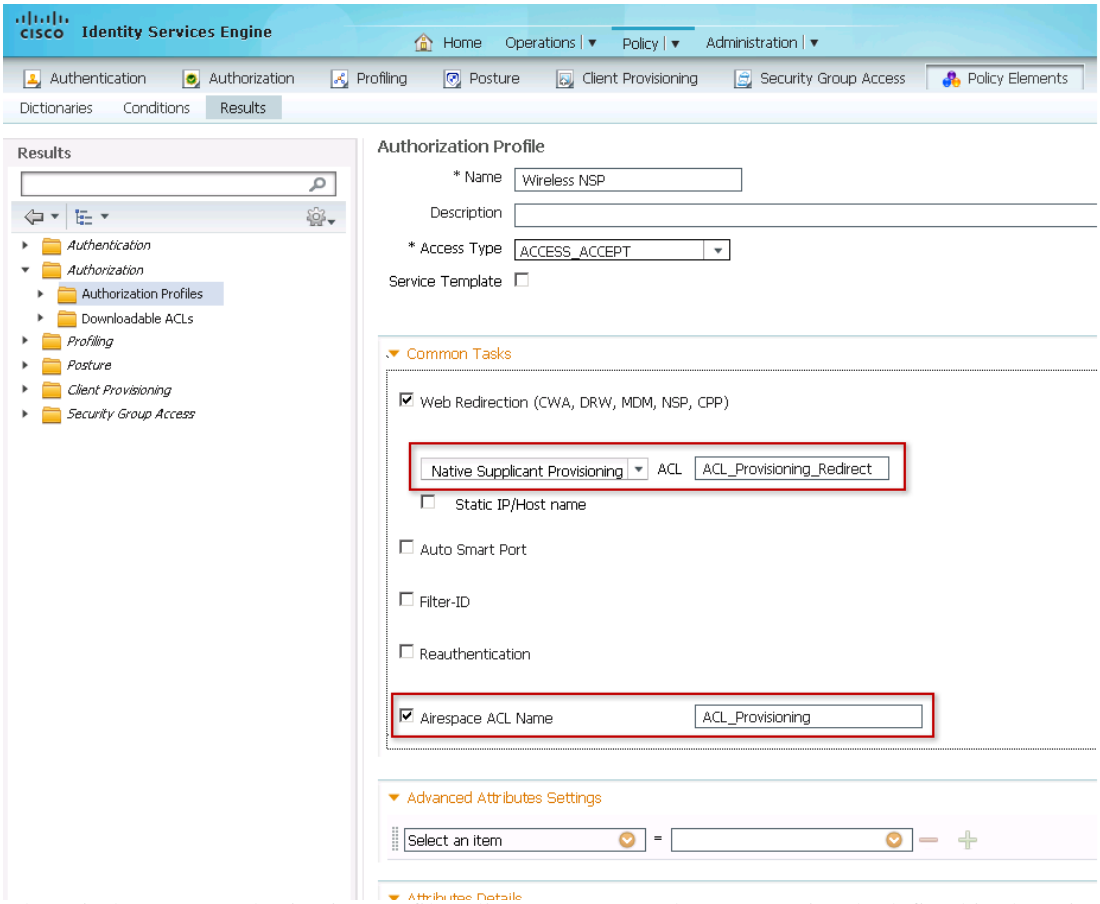
Wireless NSP Authorization Profile for Single SSID Provisioning

The native supplicant flow starts similarly regardless of device type by redirecting employees using a supported personal device to the Guest portal where they are required to enter their user credentials. From there, they are redirected to the Self-Provisioning portal to confirm their device information.

The Wireless NSP authorization profile is used in single SSID configurations to redirect devices to the Guest portal using the PEAP authentication protocol.

To configure this authorization policy, click **Policy > Policy Elements > Results > Authorization Profiles**, as shown in Figure 10-35.

Figure 10-35 Wireless NSP Authorization Profile



The Wireless NSP authorization profile relies on two named ACLs previously defined in the Wireless LAN Controller:

- ACL_Provisioning_Redirect—Applied to the Centralized Web Auth setting.
- ACL_Provisioning—Sent to the wireless controller via the Radius:Airespace-ACL-Name attribute value (AV).

The behavior of the two ACLs is slightly different between wireless controllers:

- For CUWN wireless controllers (e.g., CT5508 and Flex 7500), ACL_Provisioning_Redirect functions as both the ACL which controls web redirection and as the ACL which controls access on the network. ACL_Provisioning serves simply as an extra security configuration and is not used when URL redirection is specified. For CUWN wireless controllers the ACL_Provisioning _Redirect ACL shown in Figure 10-32 can be the same as the ACL_Provisioning.
- For Cisco IOS XE based wireless controllers (e.g., CT5760 and Catalyst 3850), ACL_Provisioning_Redirect functions strictly as the ACL which controls web redirection. ACL_Provisioning functions as the ACL which controls what the wireless client is allowed to access on the network. Hence IOS XE based wireless controllers make use of both ACLs when URL redirection is specified.

Single SSID Provisioning Authorization Rule

The Single SSID Provisioning rule links the Wireless NSP authorization profile to the conditions that authorize wireless devices authenticating via PEAP.

To force devices to the self-registration portal, a redirect URL is created with a unique Session ID and pushed to the device:

https://ip:port/guestportal/gateway?sessionId=SessionIdValue&action=nsp

When the user launches a web browser, the device is redirected to the Self-Registration portal.

Figure 10-36 shows the authorization rule defined under the authorization policies. This rule includes two conditions: Wireless_PEAP and Employee_WLAN.

Figure 10-36 Single SSID Provisioning Authorization Rule

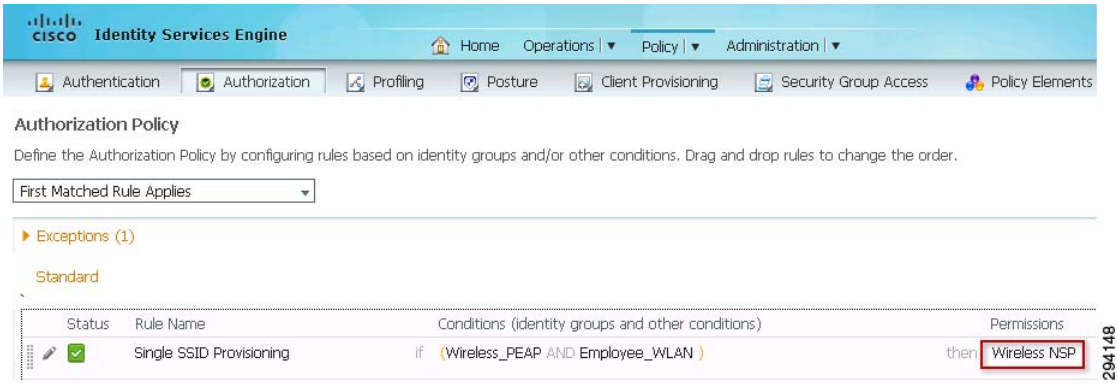
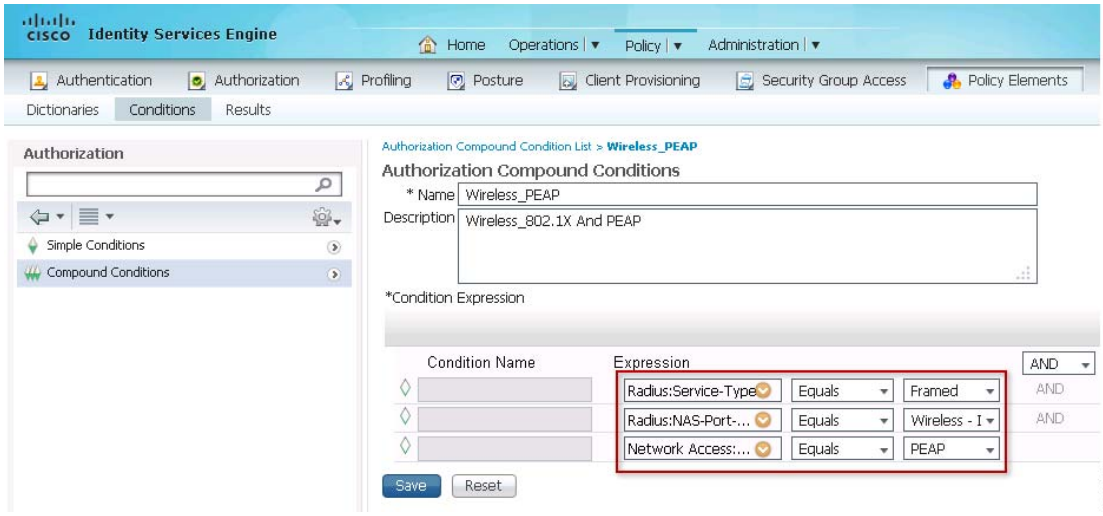


Figure 10-37 shows the Wireless_PEAP compound condition in ISE, which includes these expressions:

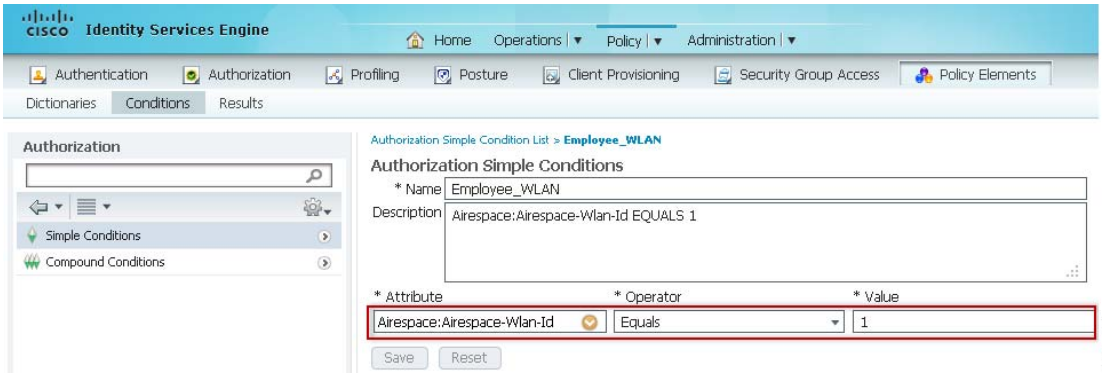
- Radius:Service-Type Equals Framed
- Radius:NAS-Port-Type Equals Wireless—IEEE 802.11
- Network Access: EapTunnel Equals PEAP

Figure 10-37 Wireless_PEAP Compound Condition



For the purposes of this CVD, the BYOD_Employee SSID number was defined as 1 during testing. The simple condition Employee_WLAN matches when the SSID number is 1. The condition is created to improve readability of the rules.

Figure 10-38 Employee_WLAN Condition



Certificate Authority Server

The Certificate Authority server is the central authority for distributing digital certificates. A Windows 2008 CA server was used as the CA server for this solution. This section focuses on:

- Network Device Enrollment Service, which is Microsoft's implementation of SCEP.
- Certificate Templates and how to design them.

NDES Server Configuration for SCEP

The Network Device Enrollment Service (NDES) is the Microsoft implementation of the SCEP, a communication protocol that makes it possible for network devices to enroll for X.509 certificates from a CA. To distribute and deploy digital x.509 client certificates to users, the Microsoft Network Device Enrollment Service (NDES) was utilized in conjunction with a Microsoft CA Server. For more details on how to implement NDES, see: http://technet.microsoft.com/en-us/library/cc753784%28WS.10%29.aspx.

By default, the NDES service is configured to present one-time enrollment passwords for certificate enrollment. The use of one-time passwords by the NDES service is typically used to allow network and IT administrators to enroll certificates for network devices within the IT organization. However, in this solution this feature is disabled because remote endpoints are authenticated by using their RSA SecurID tokens.

Disabling the “one-time password” on the NDES server is configured in the following registry key: Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Cryptography\MSCEP\EnforcePassword.

EnforcePassword value data is set to “0”, which ensures no password is requested by NDES.

Note Windows Server 2003, Microsoft SCEP (MSCEP) required a Resource Kit add-on to be installed on the same computer as the CA. In Windows Server 2008, MSCEP support has been renamed NDES and is part of the operating system. NDES may be installed on a different computer than the CA (http://technet.microsoft.com/en-us/library/cc753784%28WS.10%29.aspx).

The NDES extension to IIS uses the registry to store configuration settings. All settings are stored under one registry key:

HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\MSCEP

Note It is possible for the ISE to generate URLs which are too long for the IIS. To avoid this problem, the default IIS configuration may be modified to allow longer URLs.

The following command should be run on a command line with administrative privileges:

```
%systemroot%\system32\inetsrv\appcmd.exe set config
  /section:system.webServer/security/requestFiltering
  /requestLimits.maxQueryString:"6044"
  /commit:apphost
```

Certificate Template

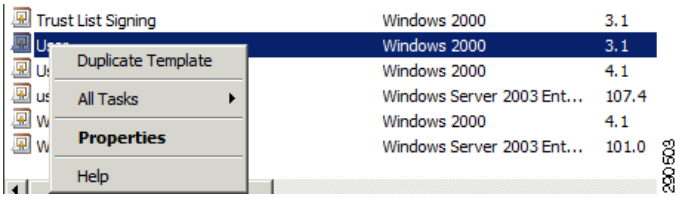
Digital certificates can be used for different purposes like server authentication, secure email, encrypting the file system, and client authentication. Hence it is important that a client is issued a certificate which serves its purpose. For example, a web server may need a certificate for server authentication. Similarly, a normal client needs a certificate mainly for client authentication. Therefore, certificate templates are needed to properly distribute certificates to users based on their specific needs. In this solution, a security template has been created on the Microsoft Windows 2008 CA server so that users can obtain the proper certificate. This section describes important steps to set up the certificate template on the Windows CA server and specific actions needed by the user.

For more information on certificate templates, see:
http://technet.microsoft.com/en-us/library/cc730826%28WS.10%29.aspx.

SCEP is used as a protocol by the endpoints to obtain their digital certificates from the CA server. The endpoints send the certificate requests to ISE, which forwards the requests to the CA server. ISE is configured as SCEP Proxy to handle these requests and once the CA server issues the certificates, ISE sends the certificates back to the clients. The properties of the "User" template are being used. That is a default template in the Microsoft Server 2008 R2 CA Server deployment. Default templates in Microsoft Server 2008 R2 cannot be edited. Therefore, a customized template can be built that gives an administrator more flexibility in defining the certificate options. This section describes how to create a customized template named "user2" in this example.

The first step is to create a duplicate template from the pre-defined list of templates. Figure 10-39 shows how to create a duplicate template.

Figure 10-39 Creating a Duplicate Template

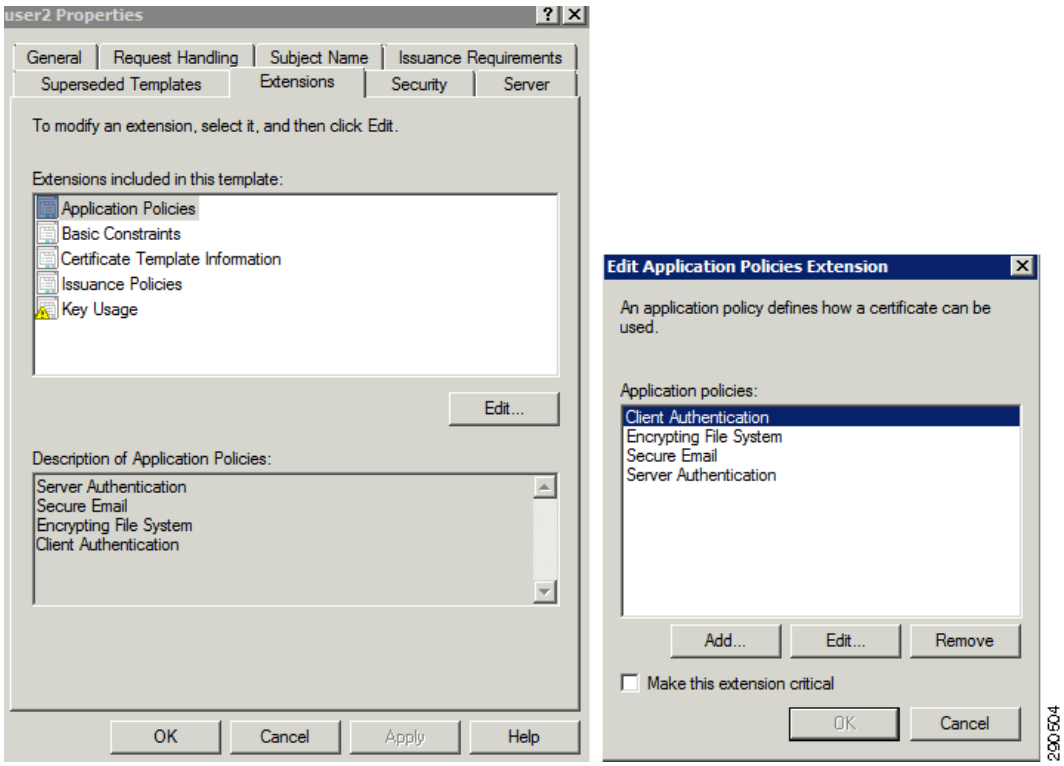


The default "User" template was copied and renamed "user2". Then the "user2" template was used to auto-enroll AnyConnect VPN clients with client certificates using this newly created template.

The next step is to configure the extensions of the certificates that are derived from the "user2" template. The EKU extension and extended property specify and limit the valid uses of a certificate. The extensions are part of the certificate itself. They are set by the issuer of the certificate and are read-only. Certificate-extended properties are values associated with a certificate that can be set in an application. To obtain more information about extended properties, see:
http://msdn.microsoft.com/en-us/library/aa380252%28v=vs.85%29.aspx.

Figure 10-40 describes how to configure the extended properties for the certificates.

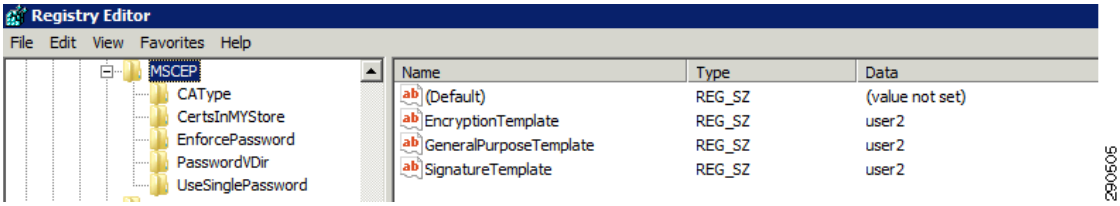
Figure 10-40 Configuring Extended Properties for Certificates



Notice the template named “user2”. This value must be set in the registry as it correlates to the “user2” template, which was copied from the “User” template in the Certificate Templates Console on the CA Server.

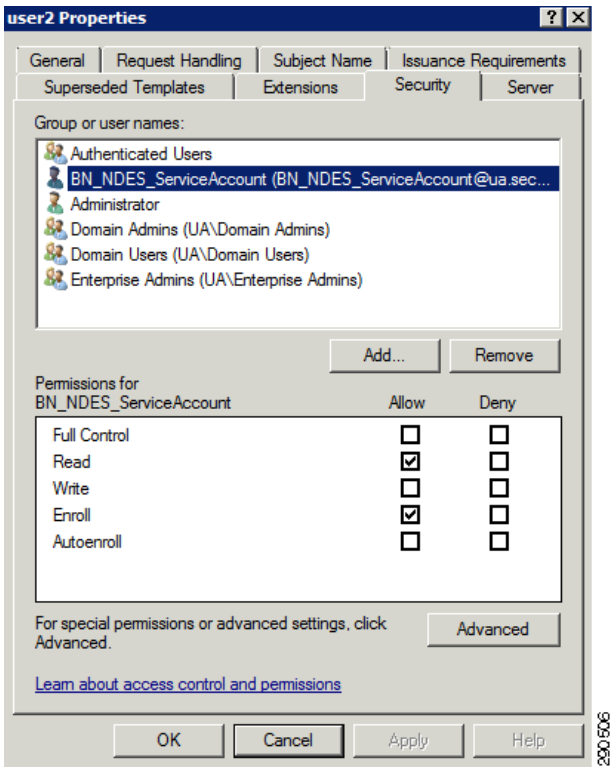
Figure 10-41 describes how the registry setting must be modified to reflect the newly-created template “user2”.

Figure 10-41 Modifying the Registry



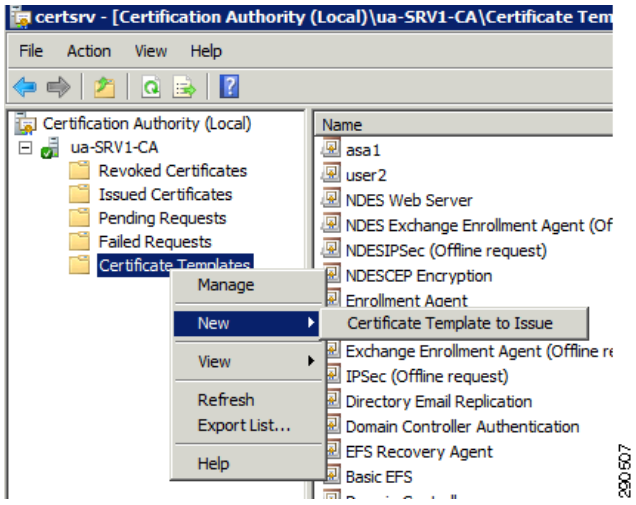
Once the template has been duplicated, the permissions are set for the NDES_ServiceAccount on the “user2” template to Read and Enroll. Figure 10-42 displays the Read and Enroll permissions that have been set for the NDES_ServiceAccount on the “user2” template.

Figure 10-42 Read and Enroll Permissions



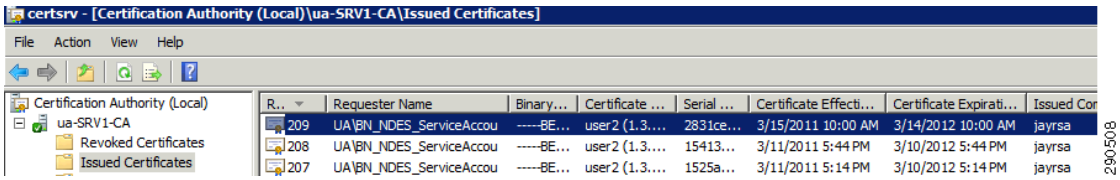
Ensure that the newly created "user2" template is available to be issued via the CA. Right click "user2" and choose the newly-created "User2 Certificate", as shown in Figure 10-43.

Figure 10-43 Ensuring Template is Available From CA



Now the certificate template is fully configured and can be used by users to submit enrollment requests. Figure 10-44 shows a successful enrollment request to the "user2" template that was submitted by a user, "jayrsa".

Figure 10-44 Successful Enrollment Request



A successful auto-enrollment request has occurred on the CA Server. Notice that the requester name is the NDES Service Account that is configured for Read and Enroll permissions and also notice that the "user2" certificate template was chosen.



BYOD Wired Infrastructure Design

Revised: August 7, 2013

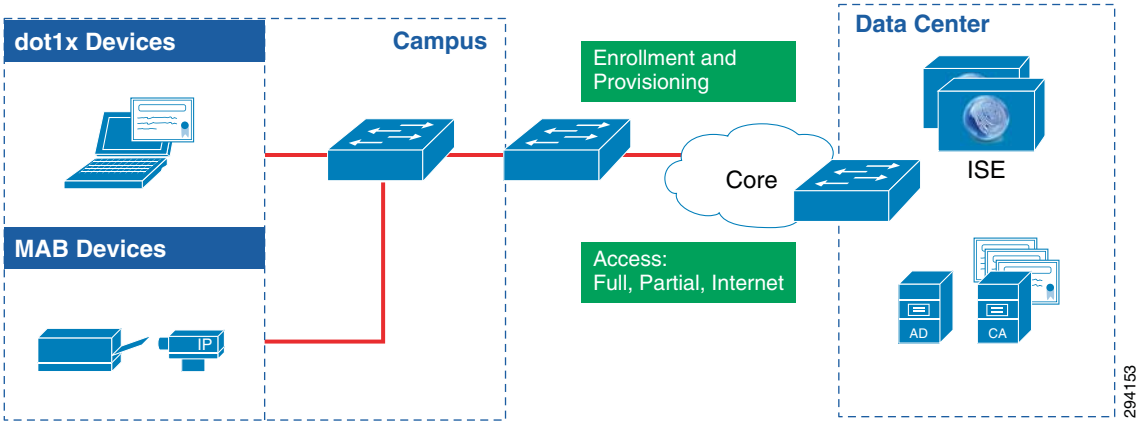
The previous sections discussed how BYOD devices can be on-boarded to the network and also how different policies can be enforced for mobile devices using wireless medium. This section discusses how to design and configure on-boarding and enforcing network access policies for wired devices. These devices can be located at Campus or at Branch location. Moreover, wired devices can connect using either converged access layer switches or by using non-converged access layer switches. This section discusses the design and configuration details for following network architectures:

- Campus (Both Converged Access and non-Converged Access)
- Branch (Both Converged Access and non-Converged Access)

Campus Wired Design

At the campus location there are 802.1X-capable clients that go through the provisioning/enrollment process and there are other types of devices like printers, cameras, etc. which do not have 802.1X capabilities and can only provide their MAC address as their source of authentication. These devices also will need to access the network and this design allows them to authenticate/authorize and obtain their authorization policy from ISE. Figure 11-1 shows an end-to-end network architecture diagram that includes wired device access from campus:

Figure 11-1 Network Diagram for Wired Devices at Campus Location



VLAN Design for Wired Switches at Campus

In the campus BYOD wired designs presented in this document, the VLAN assignment is same for all types of access—Full, Partial, or Internet. This means that the VLAN assignment to the port does not change when the device accessing the port changes. For example, the corporate-owned asset and the personal device would still use the same VLAN number. The policy enforcement is done by a DACL which is pushed from the ISE for non-Converged Access switches. On the other hand, policy enforcement in Converged Access switches is done using a named ACL instead of a DACL. To obtain more information about the different types of DACLs or named ACLs, refer to Chapter 16, "BYOD Limited Use Case—Corporate Devices" or Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices." The following is an example configuration of Layer 2 interface of the access layer switch and is the same either on centralized campus or a converged access campus switch:

```
interface GigabitEthernet1/0/2
 switchport access vlan 42   ! VLAN used in this design is 42
 switchport mode access
 ip access-group ACL-DEFAULT in
 authentication event fail action next-method
 authentication host-mode multi-auth
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 3
 spanning-tree portfast
```

IP Address Design for Campus Wired Infrastructure

In the campus wired network designs discussed in this design guide, the access layer switch performs Layer 2 functions only. The aggregation switch performs Layer 3 routing. The following is an example of part of the configuration of the Layer 3 aggregation switch:

```
ua31-6500-1#show running-config interface vlan 42
Building configuration...

Current configuration : 91 bytes
!
interface Vlan42
 ip address 10.207.42.1 255.255.255.0
 ip helper-address 10.230.1.61
end

ua31-6500-1#
```

As seen above, the Layer 3 interfaces are configured with the ip-helper address command, which helps clients obtain an IP address. For the purposes of this design guide, the DHCP server is located in the data center.

Policy Enforcement in the Campus for Wired Devices

ACLs are the primary method through which policy enforcement is done at access layer switches for wired devices within the campus. There are two distinct sets of ACLs used:

- ACLs for managing the device—These ACLs are used for provisioning the device or managing the device like Blacklisting.
- ACLs that are mainly for enforcing the policies.

The policy enforcement method at campus non-Converged Access switches is done by defining DACLs in the ISE based on the authZ policy and pushing that DACL to the port on the access layer switch. In Converged Access switches, policy enforcement is done through a named ACL sent by the ISE based on the authZ policy. The named ACL must be previously configured on the Converged Access Catalyst 3850 switch. To obtain more information on the authZ profiles used in this design guide, refer to either Chapter 16, "BYOD Limited Use Case—Corporate Devices" or Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices."

ACL Design for Campus Access Layer Switches

This section discusses the set of ACLs that are used for provisioning devices onto the network, protecting against unauthorized access, and blacklisting the device. The ACLs discussed in this section apply to both Converged access layer switches and non-converged access layer switches. Table 11-1 summarizes these ACLs and their purpose.

Table 11-1 Campus ACLs and Purpose

ACL Name	Where it Applies	Purpose
ACL-DEFAULT	Access Layer Switch	To protect against unauthorized access through the switch port
ACL_Provisioning	ISE	To allow an endpoint access to complete on-boarding process
ACL_Provisioning_Redirect	Access Layer Switch	Redirect web traffic initiated by new devices accessing the network. This ACL is a named ACL that is present on the access layer switch.
ACL_BLACKHOLE_Redirect	Access Layer Switch	Redirect web traffic initiated by black listed devices

ACL-DEFAULT—This ACL is configured on the access layer switch and used as a default ACL on the port. Its purpose is to prevent un-authorized access. In an 802.1X authentication/authorization scenario, after the device is authenticated and authorized, if there is no DACL applied to the port or if there is a mistake in the syntax of the downloadable ACL and the switch rejects the DACL sent by ISE, ACL-DEFAULT protects the port in the above mentioned scenarios. In the converged access design, the ACL is a named ACL and is configured on the Catalyst 3850 switch. The ISE sends the name of the ACL to be applied at the port. Again, if the switch rejects the named ACL sent by ISE, ACL-DEFAULT protects the port.

An example of a default ACL on a campus access layer switch is shown below:

```
Extended IP access list ACL-DEFAULT
    10 permit udp any eq bootpc any eq bootps log (2604 matches)
    20 permit udp any host 10.230.1.45 eq domain
    30 permit icmp any any
    40 permit udp any any eq tftp
    50 deny ip any any log (40 matches)
```

As seen from the output above, ACL-DEFAULT allows DHCP, DNS, ICMP, and TFTP traffic and denies everything else.

ACL_Provisioning_Redirect—This ACL is used during on-boarding of wired devices. The ACL triggers a redirection upon HTTP or HTTPS traffic from the client to anywhere, which means that when the user opens a web browser and attempts to access any website, that traffic is re-directed. The example shown below redirects any web traffic initiated by the user. However, this ACL can be modified to allow only certain traffic to be redirected to ISE portal. The underlying assumption in this design is that all the devices must be registered with ISE, therefore when an un-registered device accesses the network, it is redirected to ISE.

An example of ACL_Provisioning_Redirect ACL on a campus switch is shown below:

```
Extended IP access list ACL_Provisioning_Redirect
    10 deny udp any eq bootpc any eq bootps log
    20 deny udp any host 10.230.1.45 eq domain (43 matches)
    30 deny ip any host 10.225.42.15 (27 matches)
    40 permit tcp any any eq www (30 matches)
    50 permit tcp any any eq 443 (240 matches)
```

ACL_BLACKHOLE_Redirect—This ACL is used to redirect devices that have been blacklisted to the ISE portal to let the user know that the device in use has been blacklisted. This ACL is similar to the ACL_Provisioning_Redirect ACL.

An example of ACL_BLACKHOLE_Redirect on a campus switch is shown below:

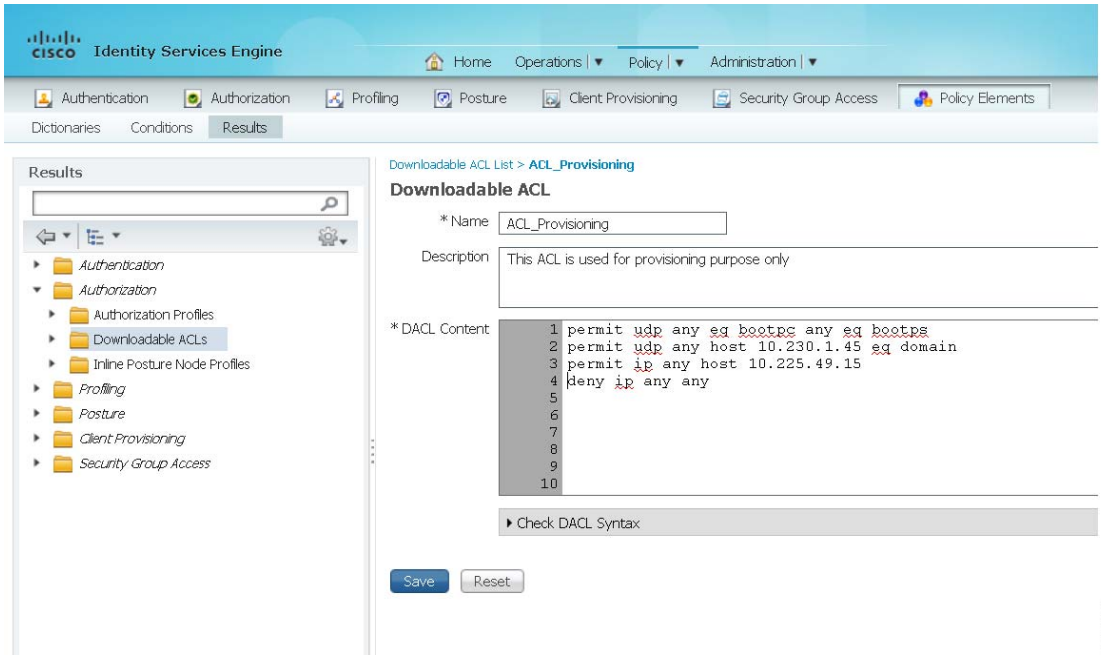
```
Extended IP access list ACL_BLACKHOLE_Redirect
    10 deny udp any eq bootpc any eq bootps
    20 deny udp any host 10.230.1.45 eq domain
    35 deny ip any host 10.225.49.15
    40 permit ip any any
```

Note The converged access layer switches use the same ACL_BLACKHOLE_Redirect for redirecting black listed wired devices.

Provisioning ACL

This ACL is also used during the on-boarding of wired devices. This DACL is downloaded from the ISE and restricts access to only the ISE, DNS, and DHCP server. This ACL is defined on the ISE, as shown in Figure 11-2.

Figure 11-2 ACL_Provisioning



Note The ACL_Provisioning ACL is also used for Converged access layer switches.

802.1X and AAA Configuration for Campus Switches

A Cisco Catalyst Switch is used to provide end user Ethernet connectivity into the network in this design guide. The access layer switch enables 802.1X authentication for the client devices and interacts with the Identity Services Engine using the RADIUS protocol. Based on the results from the authentication process, a user may be allowed restricted or full access into the network using a VLAN assignment and a downloadable Access Control List (DACL). The flex-authentication configuration described below allows for using both 802.1X and MAC Authentication Bypass (MAB) as a fallback mechanism. Flex-auth is useful for devices that do not have 802.1X support such as printers.

This section discusses on the configuration details of enabling AAA on the campus access layer switches, and these switches can be either converged access layer switches non-converged access layer switches.

The following steps are required to configure the access switch for AAA on the Campus Switch:

Step 1 Enable Authentication, Authorization, and Accounting (AAA):

```
ACL(config)# aaa new-model
```

Step 2 Create an authentication method for 802.1X (default use all RADIUS servers for authentication):

```
ACL(config)# aaa authentication dot1x default group radius
```

Step 3 Create an authorization method for 802.1X (enables RADIUS for policy enforcement):

```
ACL(config)# aaa authorization network default group radius
```

Step 4 Create an accounting method for 802.1X (provides additional information about sessions to ISE):

```
ACL(config)# aaa accounting dot1x default start-stop group radius
```

The following steps are required to configure the access switch for RADIUS:

Step 1 Add ISE server to the RADIUS group:

```
ACL(config)# radius-server host 10.225.49.15 auth-port 1812 acct-port 1813 key
shared-secret
```

Step 2 Configure ISE server dead time (15 seconds total-3 retries of 5 second timeout):

```
ACL(config)# radius-server dead-criteria time 5 tries 3
```

Step 3 Configure the switch to send Cisco Vendor-Specific attributes:

```
ACL(config)# radius-server vsa send accounting
ACL(config)# radius-server vsa send authentication
```

Step 4 Configure the Cisco Vendor-Specific attributes:

```
ACL(config)# radius-server attribute 6 on-for-login-auth
ACL(config)# radius-server attribute 8 include-in-access-req
ACL(config)# radius-server attribute 25 access-request include
```

Step 5 Configure IP address to be used to source RADIUS messages:

```
ACL(config)# ip radius source-interface interface-name Vlan4093
```

The following steps are required to configure the access switch for 802.1X:

Step 1 Enable 802.1X globally (command by itself does not enable authentication on the switchports):

```
ACL(config)# dot1x system-auth-control
```

Step 2 Enable IP device tracking:

```
ACL(config)# ip device tracking
```

The following interface level commands enable 802.1X for Flex-Auth:

Step 1 Configure the authentication method priority (dot1x has higher priority over MAB):

```
ACL(config-if)# authentication priority dot1x mab
```

Step 2 Configure the authentication method order (dot1x first):

```
ACL(config-if)# authentication order dot1x mab
```

Step 3 Enable Flex-Auth:

```
ACL(config-if)# authentication event fail action next-method
```

Step 4 Enable support for more than one MAC address on the physical port:

```
ACL(config-if)# authentication host-mode multi-auth
```

Step 5 Configure the violation action (restrict access for additional devices that may fail authentication):

```
ACL(config-if)# authentication violation restrict
```

Step 6 Enable port for 802.1X:

```
ACL(config-if)# dot1x pae authenticator
```

Step 7 Enable port for MAB:

```
ACL(config-if)# mab
```

Step 8 Configure timers (30 seconds (10x3) until falling back to MAB):

```
ACL(config-if)# dot1x timeout tx-period 3
```

Step 9 Turn authentication on:

```
ACL(config-if)# authentication port-control auto
```

Step 10 Enable the ACL-DEFAULT to the port

```
ACL(config-if)# ip access-group ACL-DEFAULT in
```

Step 11 Enable http and https server:

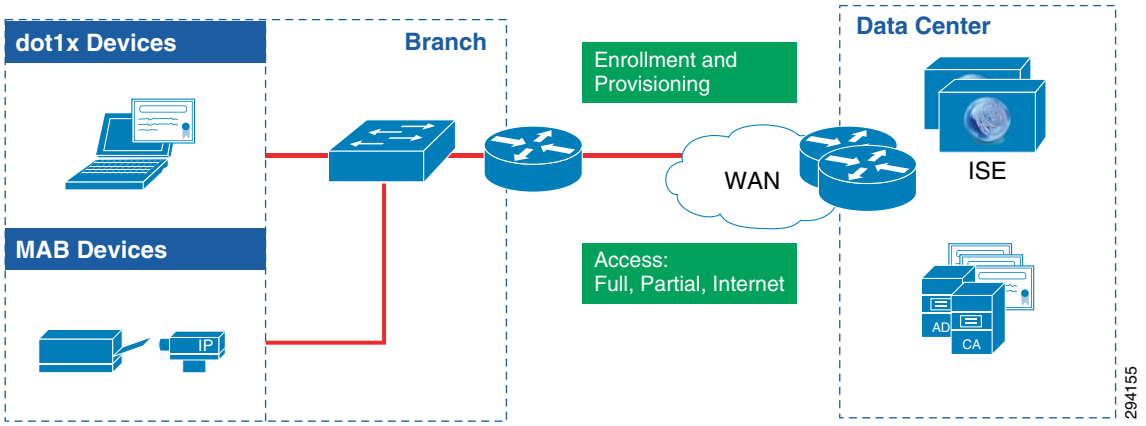
```
ACL (config)# ip http-server
ACL (config)# ip http secure-server
```

Branch Wired Design—Non-Converged Access

At a branch location, there are 802.1X capable clients that go through the provisioning/enrollment process and there are also other types of devices such as printers, cameras, etc. which do not have 802.1X capabilities and can only provide their MAC address as their source of authentication. These devices also need to access the network and this design allows them to authenticate/authorize and obtain their authorization policy from ISE. This section discusses wired designs for branches which do not deploy Converged Access (Catalyst 3850) switches. The branch wired design discussed in this section is meant to accompany FlexConnect-based wireless branch designs. The Converged Access branch wired design is discussed in Branch Wired Design—Converged Access.

Figure 11-3 shows an end-to-end network architecture diagram that includes wired device access from the branch.

Figure 11-3 Network Diagram for Wired Access at Branch Location



VLAN Design at Branch Locations

Four VLANs are implemented for wired devices at the non-Converged Access branch location. Table 11-2 illustrates the names of these VLANs and their purpose.

Table 11-2 VLANs and their Purpose

VLAN Name	VLAN Number	Description
Wired_Full	13	Devices placed in this VLAN get full access to corporate resources and branch local servers.
Wired_Partial	14	Devices placed in this VLAN get restricted access to resources.
Wired_Internet	15	Devices placed in this VLAN get only Internet access only.
Branch_Server	16	Local Servers at branch location are placed in this VLAN.

IP Address Allocation at Branch Location

In the non-converged access branch network design discussed in this design guide, the switch performs Layer 2 functions only and the branch router performs Layer 3 routing. Hence, all the Layer 3 interfaces for the VLANs mentioned above are implemented at the branch router. The following is an example configuration of the branch router:

```
interface GigabitEthernet0/1.13
 encapsulation dot1Q 13
 ip address 10.200.13.2 255.255.255.0
 ip helper-address 10.230.1.61
 standby 13 ip 10.200.13.1
 standby 13 priority 110
 standby 13 preempt
!
interface GigabitEthernet0/1.14
 encapsulation dot1Q 14
 ip address 10.200.14.2 255.255.255.0
 ip access-group Branch1_ACL_Partial_Access in
 ip helper-address 10.230.1.61
 standby 14 ip 10.200.14.1
 standby 14 priority 110
```

```
 standby 14 preempt
!
interface GigabitEthernet0/1.15
 encapsulation dot1Q 15
 ip address 10.200.15.2 255.255.255.0
 ip access-group ACL_Internet_Only in
 ip helper-address 10.230.1.61
 standby 15 ip 10.200.15.1
 standby 15 priority 110
 standby 15 preempt
!
interface GigabitEthernet0/1.16
 encapsulation dot1Q 16
 ip address 10.200.16.2 255.255.255.0
 ip helper-address 10.230.1.61
 standby 16 ip 10.200.16.1
 standby 16 priority 110
 standby 16 preempt
!
```

As seen above, the Layer 3 interfaces are configured with the **ip-helper address** command, which helps branch clients obtain an IP address. For the purposes of this design guide, the DHCP server is in a data center location.

Policy Enforcement in the Branch for Wired Devices

ACLs are the primary method through which policy enforcement is done at access layer switches for wired devices within the branch. There are two distinct sets of ACLs used:

- ACLs for managing the device—These ACLs are used for provisioning the device or managing the device like Blacklisting.
- ACLs that are mainly for enforcing the policies.

When designing the ACLs for branch the following should be considered:

- Configuring static ACLs at every branch router in the network.
- Configuring the ISE to push downloadable ACLs to access layer switches at every branch location.

Table 11-3 gives the advantages and disadvantages of each approach.

Table 11-3 ACL Policy Enforcement

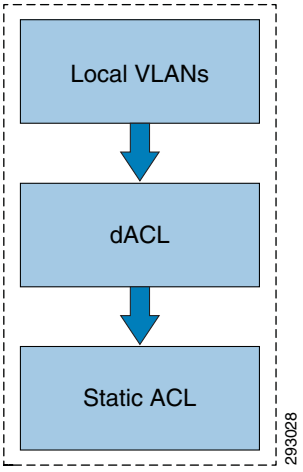
Method	Advantages	Disadvantages
Static ACLs	Modify the ACL based on the branch's needs	Hard to manage each branch policy individually
Downloadable ACLs	Centralized access control	Creating an individual policy at ISE for every branch location would make the policy very large from an administrative perspective. For example, to manage 500 unique branches, 500 ACLs, 500 authZ profiles, and 500 authZ policy rules would need to be defined on the ISE.

Each of the above methods has advantages and disadvantages. This design guide focuses on a combination that includes both methods. The static ACLs are the primary method by which access is restricted. However, the static ACLs are applied at the router only; to override the ACL called

"DEFAULT-ACL" which is present on every port of the access layer switch, a DACL (permit all traffic) is downloaded from ISE. This DACL from the ISE allows the traffic to flow upstream from the access layer switch to the Branch router. The Branch router is pre-configured with the different ACLs that restrict access. These ACLs either provide full, partial, or Internet access to the users.

Figure 11-4 shows how the authorization policy pushes the VLAN information and the DACL (permit all traffic) to the port on the access layer switch, thereby allowing the traffic to reach from the access layer switch up to the router where the traffic will be filtered.

Figure 11-4 Enforcing Permissions



ACL Design at Branch Location

ACLs are very important at the branch location, since they are the main method used to enforce policies. Some ACLs are defined on the Layer 2 switch for provisioning purposes, while others are defined on the branch router. In addition, some ACLs may be downloaded from the ISE.

Table 11-4 summarizes the various ACLs at branch locations and their purpose.

Table 11-4 Branch ACLs and Purpose

ACL Name	Where it Applies	Purpose
ACL_DEFAULT	Switch	To protect against unauthorized access through the switch port
ACL_Provisioning_Redirect	Switch	Redirect web traffic initiated by new devices accessing the network.
ACL_Blackhole	Switch	Redirect web traffic initiated by black listed devices
ACL_Internet_Only	Branch Router	Allow only Internet traffic
ACL_Provisioning	ISE	Used during provisioning process
ACL_Partial_Access	Branch Router	Allow partial access to certain resources

ACL_DEFAULT—This ACL is used as a default ACL on the port and its purpose is to prevent un-authorized access. In an 802.1X authentication/authorization scenario, after the device is authenticated and authorized, if there is no DACL applied to the port or if there is a mistake in the syntax of the downloadable ACL and the switch rejects the DACL sent by ISE, ACL_DEFAULT protects the port in the above mentioned scenarios. An example of a default ACL is shown below:

```
bn22-3750x-1#show ip access-lists
Load for five secs: 13%/0%; one minute: 16%; five minutes: 16%
Time source is NTP, 16:24:50.872 EDT Wed Sep 19 2012

Extended IP access list ACL-DEFAULT
  10 permit udp any eq bootpc any eq bootps
  20 permit udp any any eq domain
  30 permit icmp any any
  40 permit udp any any eq tftp
  50 deny ip any any log
```

As seen from the output above, ACL_DEFAULT allows DHCP, DNS, ICMP, and TFTP traffic and denies everything else.

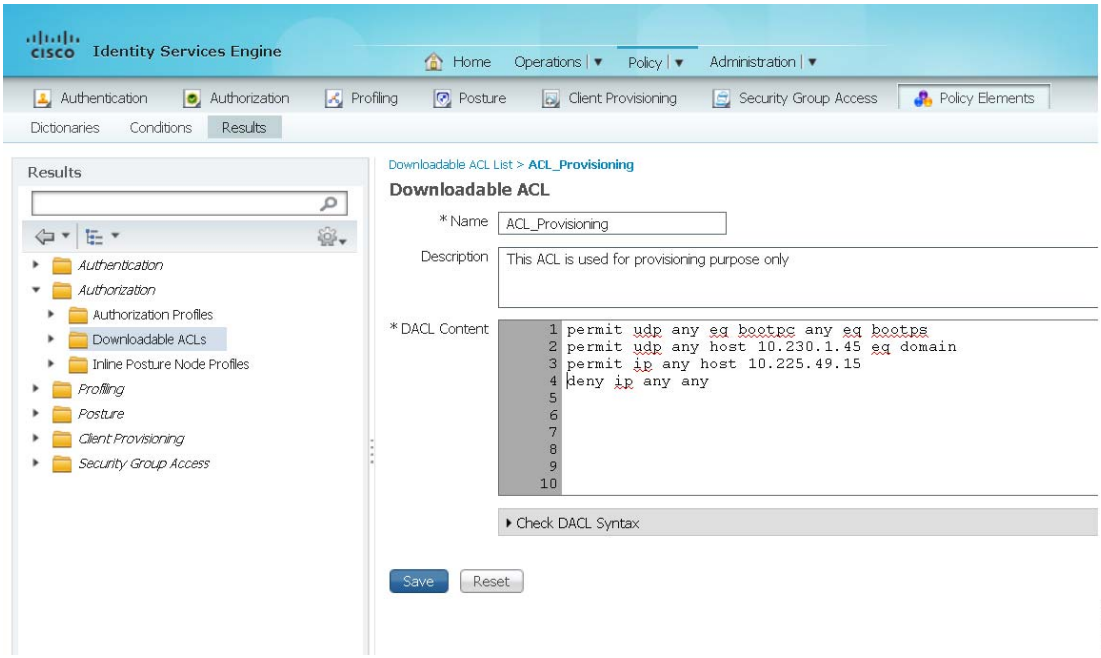
ACL_Provisioning_Redirect—This ACL is used during the on-boarding of wired devices. This ACL triggers a redirection upon HTTP or HTTPS traffic from the client to anywhere, which means that when the user opens a web browser and attempts to access any website, that traffic is re-directed. The example shown below redirects any web traffic initiated by the user. However this ACL can be modified to allow only certain traffic to be redirected to ISE portal. The underlying assumption in this design is that all the devices must be registered with ISE, therefore when an un-registered device accesses the network, it is redirected to ISE.

```
uasl-3750x-1#show ip access-lists | begin ACL_Provisioning_Redirect
Extended IP access list ACL_Provisioning_Redirect
    10 deny udp any eq bootpc any eq bootps log
    20 deny udp any host 10.230.1.45 eq domain (1865 matches)
    30 deny ip any host 10.225.42.15 (839 matches)
    40 deny ip any host 10.225.49.15 (1853 matches)
    50 permit tcp any any eq www (3728 matches)
    60 permit tcp any any eq 443 (4140 matches)
uasl-3750x-1#
```

Provisioning ACL

This ACL is also used during the on-boarding of wired devices. This DACL is downloaded from the ISE and restricts access to only the ISE, DNS, and DHCP server. This ACL is defined on the ISE, as shown in Figure 11-5.

Figure 11-5 ACL_Provisioning



802.1X and AAA Configuration for Branch Switches

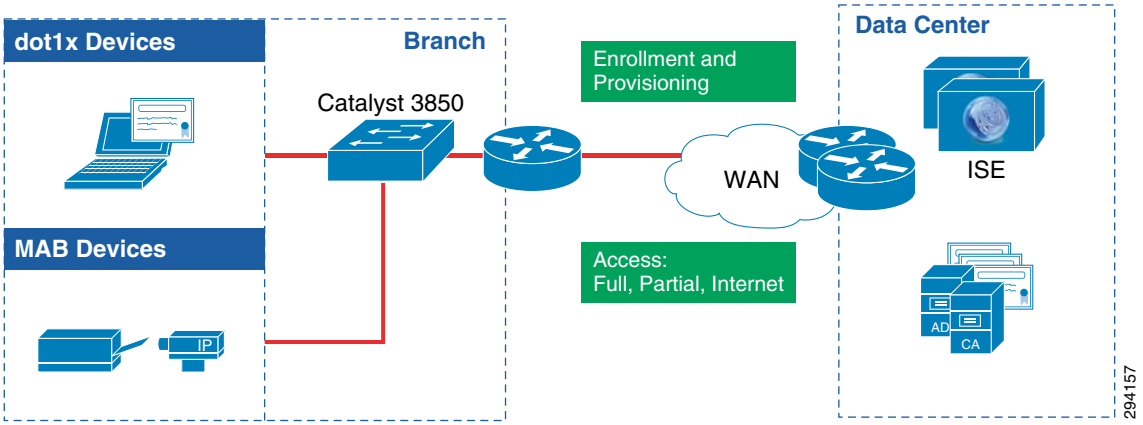
The configuration of 802.1X and AAA for branch non-Converged Access switches is exactly identical to the campus switches. Refer to 802.1X and AAA Configuration for Campus Switches for details.

Branch Wired Design—Converged Access

At a branch location, there are 802.1X capable clients that go through the provisioning/enrollment process and there are other types of devices like printers, cameras, etc. which do not have 802.1X capabilities and can only provide their MAC address as their source of authentication. These devices also need to access the network and this design allows them to authenticate/authorize and obtain their authorization policy from ISE. This section discusses wired designs for branches that deploy Converged Access (Catalyst 3850) switches.

Figure 11-6 shows an end-to-end network architecture diagram that includes wired device access from the branch.

Figure 11-6 Network Diagram for Wired Devices at Branch Location Using Converged Access Switches



VLAN Design at Branch Locations

In the Branch BYOD wired designs presented in this document, the VLAN assignment is same for all types of access—Full, Partial, or Internet. This means that the VLAN assignment to the port does not change when the device accessing the port changes. For example, the corporate-owned asset and the personal device would still use the same VLAN number. The policy enforcement in Converged Access switches is done using a named ACL instead of a DACL. Different named ACLs are applied to each device granting different access to the network. Since the named ACL is configured on the Catalyst 3850 switch specific to the particular branch, a single Cisco ISE policy can be implemented across multiple branches. However the Access Control Entries (ACEs) within the ACL for each branch can be unique to the IP addressing of the branch. This reduces the administrative complexity of the Cisco ISE policy, albeit at the expense of increased complexity of having to configure and maintain ACLs at each branch Catalyst 3850 Series switch.

Three VLANs are implemented for wired devices at the Converged Access branch location. Table 11-5 illustrates the names of these VLANs and their purpose.

Table 11-5 VLANs and their Purpose—Converged Access

VLAN Name	VLAN Number	Description
BYOD_Employee	10	Devices in this VLAN get access to either full, partial or limited access based on named ACL.
BYOD_Provisioning	11	Provisioning VLAN
Branch_Server	16	Local Servers at branch location are placed in this VLAN.

IP Address Allocation at Branch Location

In the converged access branch network design discussed in this design guide, the Catalyst 3850 switch performs Layer 2 functions only. There is no branch router, unlike Branch Wired design for non-Converged Access. The following is an example configuration of a Layer 2 interface of the access layer switch and is the same on either a centralized campus or a converged access campus switch:

```
interface GigabitEthernet1/0/2
 switchport access vlan 42   ! VLAN used in this design is 42
```

```
 switchport mode access
 ip access-group ACL-DEFAULT in
 authentication event fail action next-method
 authentication host-mode multi-auth
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 3
 spanning-tree portfast
```

Layer 3 connectivity within the branch is provided by the ISR routers that also serve as the WAN connectivity point for the branch. The following is an example of part of the configuration of the Layer 3 router:

```
ua31-6500-1#show running-config interface vlan 42
Building configuration...

Current configuration : 91 bytes
!
interface Vlan42
 ip address 10.207.42.1 255.255.255.0
 ip helper-address 10.230.1.61
end
```

As seen above, the Layer 3 interfaces are configured with the **ip-helper address** command, which helps branch clients obtain an IP address.

Policy Enforcement at the Branch Using Converged Access Switches

ACLs are the primary method through which policy enforcement is done at access layer switches for wired devices within the branch. There are two distinct sets of ACLs used:

- ACLs for managing the device—These ACLs are used for provisioning the device or managing the device like Blacklisting.
- ACLs that are mainly for enforcing the policies.

In Converged Access switches, policy enforcement is done through a named ACL sent by the ISE based on the authZ policy. The named ACL must be previously configured on the Converged Access Catalyst 3850 switch. To obtain more information on the authZ profiles used in this design guide, refer to either Chapter 16, “BYOD Limited Use Case—Corporate Devices” or Chapter 15, “BYOD Enhanced Use Case—Personal and Corporate Devices.”

ACL Design at Branch Location for Converged Access Switches

This section discusses both sets of ACLs that are important for a converged access layer switch at a branch location:

- ACL that are used for provisioning.
- ACLs that are used for policy enforcement.

Table 11-6 summarizes the various ACLs in the Converged Access branch wired branch design and their purpose.

Table 11-6 Branch ACLs and Purpose

ACL Name	Where it Applies	Purpose
ACL_DEFAULT	Switch	To protect the switch port
ACL_Blackhole	Switch	Redirect web traffic initiated by black listed devices
ACL_Internet_Only	Switch	Allow only Internet traffic
ACL_Provisioning	ISE	Used during provisioning process
ACL_Partial_Access	Switch	Allow partial access to certain resources
ACL_Full_Access	Switch	Allow full access to all resources

ACL_Default—This ACL is used as a default ACL on the port and its purpose is to prevent un-authorized access. In the Converged Access Design this is done through a named ACL approach. The ACL_DEFAULT resides on the Catalyst 3850 switch.

```
Extended IP access list ACL_DEFAULT
  10 permit udp any eq bootpc any eq bootps
  20 permit udp any any eq domain
  30 permit icmp any any
  40 permit udp any any eq tftp
  50 deny ip any any
```

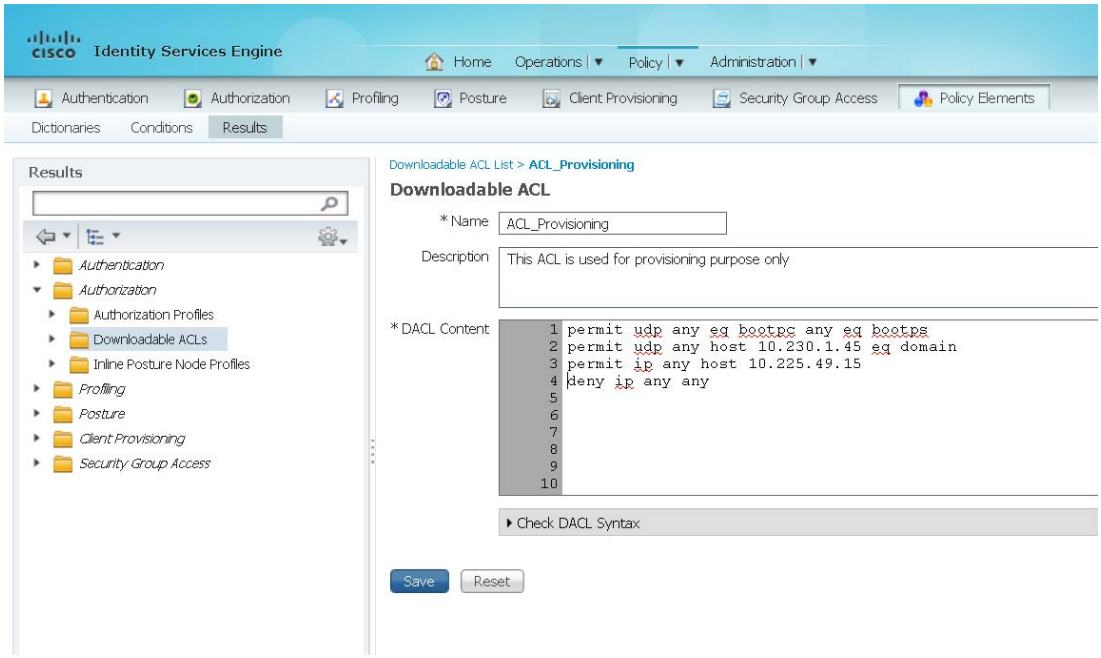
As seen from the output above, ACL_DEFAULT allows DHCP, DNS, ICMP, and TFTP traffic and denies everything else.

ACL_Provisioning_Redirect—This ACL is used during the on-boarding of wired devices. This ACL triggers a redirection upon HTTP or HTTPS traffic from the client to anywhere, which means that when the user opens a web browser and attempts to access any website, that traffic is re-directed. The example shown below redirects any web traffic initiated by the user. However, this ACL can be modified to allow only certain traffic to be redirected to ISE portal. The underlying assumption in this design is that all the devices must be registered with ISE, therefore when an un-registered device accesses the network, it is redirected to ISE.

```
Extended IP access list ACL_Provisioning_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 permit tcp any any eq www
 permit tcp any any eq 443
```

ACL_Provisioning—This ACL is also used during the on-boarding of wired devices. This DACL is downloaded from the ISE and restricts access to only the ISE, DNS, and DHCP server. This ACL is defined on the ISE, as shown in Figure 11-7.

Figure 11-7 ACL_Provisioning



802.1X and AAA Configuration for Branch Switches

The configuration of 802.1X and AAA for branch switches is exactly identical to the campus switches. Refer to 802.1X and AAA Configuration for Campus Switches.

MAB Devices at Branch or at Campus Location

This section discusses how to design access for MAB devices using either converged access switches or traditional access layer switches. MAB devices can be present at either branch or campus locations.

MAB devices are generally those devices that cannot run 802.1X and can only present their mac-address for authentication. It is important to note that BYOD devices also use the MAB protocol during the provisioning process. During the provisioning process, BYOD devices are re-directed to the ISE guest portal to complete the registration process. MAB devices do not need to be registered and therefore do not need to be re-directed. The requirement for MAB devices is to authenticate the device and apply an authorization policy. Here are the high level steps that need to be performed for MAB devices:

1. Configure the access layer switch port or WLC to support the MAB protocol.
2. Import a MAC-address list of all MAB devices as an Identity group in ISE.
3. Configure an authentication policy for Wired and Wireless MAB devices. This same policy will be used to authenticate BYOD devices during provisioning.
4. Configure authorization policy rules in ISE for wired and wireless devices.

When a MAB device connects, the access layer switch sends the authentication request to the ISE using the MAC-address of the device as the source of authentication. An example is shown below.

```
Sep 25 11:09:50.741: %DOT1X-5-FAIL: Authentication failed for client (0050.568f.
1bb2) on Interface Gi1/0/10 AuditSessionID 0AC81304000002212 92C2D59
```

```
Sep 25 11:09:50.741: %AUTHMGR-7-RESULT: Authentication result 'no-response' from
 'dot1x' for client (0050.568f.032b) on Interface Gi1/0/10 AuditSessionID 0AC813
0400000221292C2D59
Sep 25 11:09:50.749: %AUTHMGR-7-FAILOVER: Failing over from 'dot1x' for client
0050.568f.032b) on Interface Gi1/0/10 AuditSessionID 0AC81304000000221292C2D59
Sep 25 11:09:50.749: %AUTHMGR-5-START: Starting 'mab' for client (0050.568f.032b
) on Interface Gi1/0/10 AuditSessionID 0AC81304000000221292C2D59
```

In this design, all the MAC addresses of MAB devices are placed in an internal identity group called MAB_DEVICES so ISE will know this device in advance. To add new MAC addresses to the MAB_DEVICES identity group, click **Administration > Groups > Endpoint Identity Groups**, as shown in Figure 11-8.

Figure 11-8 MAB_DEVICES Identity Group

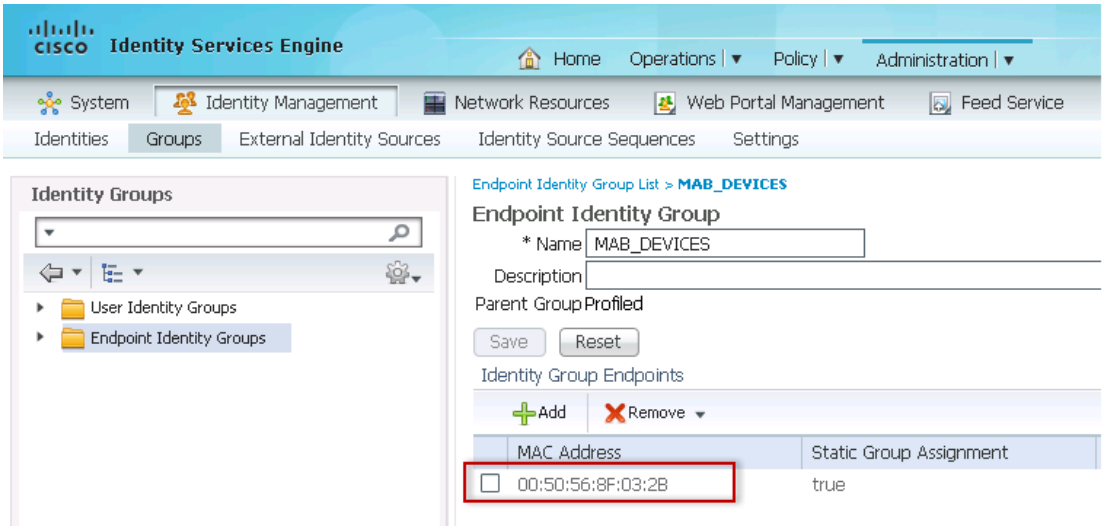


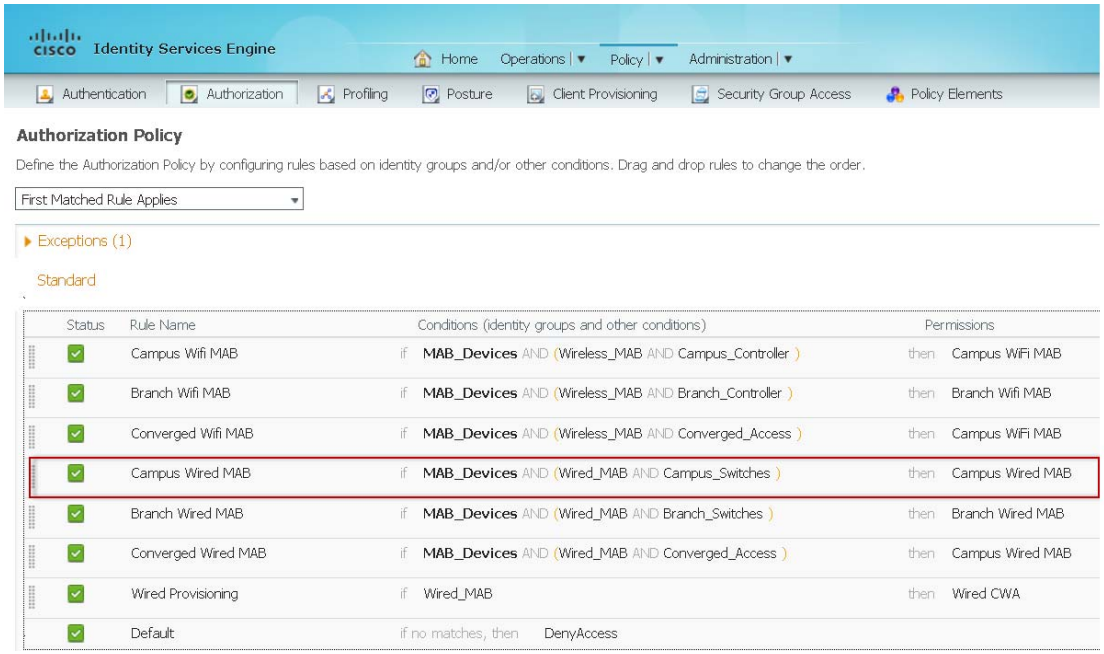
Figure 11-9 shows the authentication policy defined on the ISE for wired MAB devices.

Figure 11-9 WIRED MAB AuthC



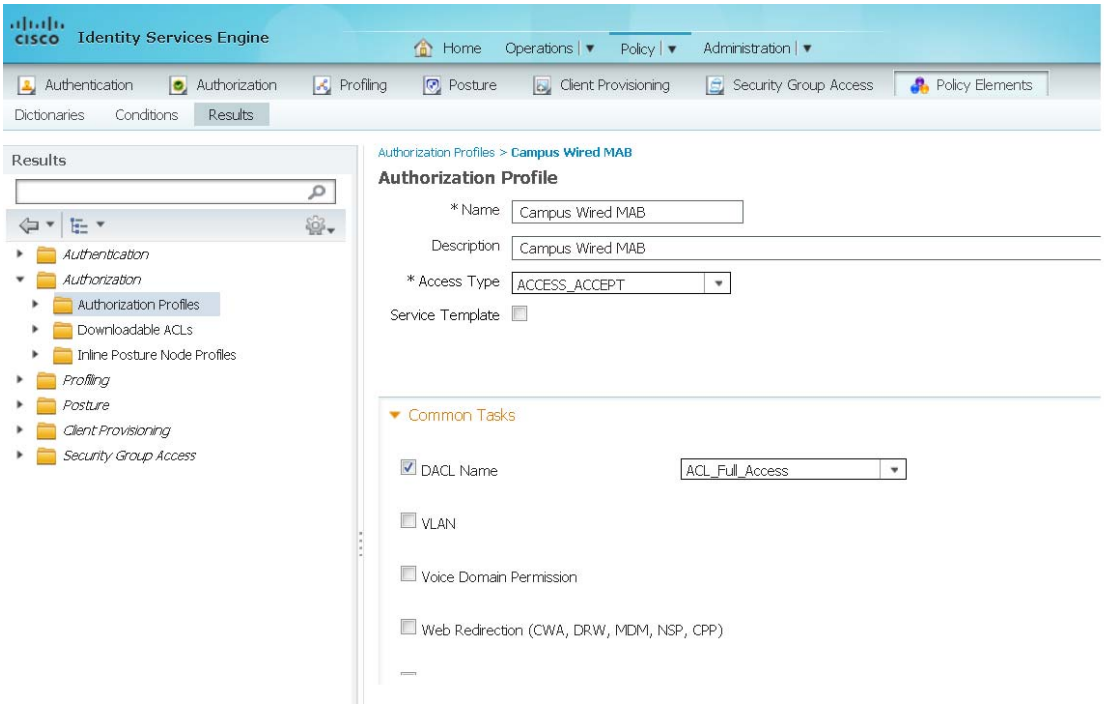
The authorization policy is different for MAB devices originating in the branch design with FlexConnect versus in the campus location. In the branch design in a FlexConnect model, every device is placed in different VLANs, but this is not done in the campus or a branch design with converged access. Hence there are different rules that are defined in the authZ policy to take care of location of the device-branch versus campus. Figure 11-10 shows how the policy rules are defined for campus devices.

Figure 11-10 Authorization Policy for MAB devices



Campus Wired MAB is an authorization profile that pushes the appropriate settings to the access layer switch. Figure 11-11 shows the authorization profile details.

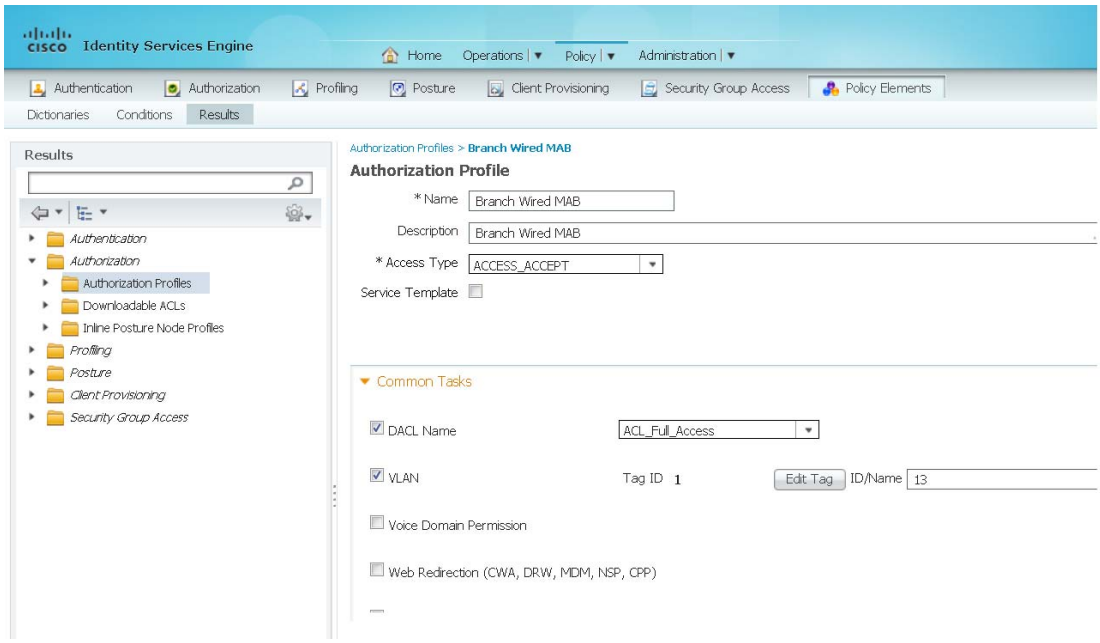
Figure 11-11 Campus Wired MAB Authorization Profile



The Campus Wired MAB authorization profile does not push VLAN information, but rather applies a DACL to the port. The Converged Access design uses the same authorization profile as shown in the Figure 11-11. Also note that in the Converged Access design, for the authorization profile, a DACLs is used for both Campus and Branch designs.

Conversely, Branch_Wired_MAB authorization profile pushes the VLAN information to the access port on the wired switch for designs with branch with FlexConnect. Figure 11-12 shows the Branch_Wired_MAB profile configuration.

Figure 11-12 Branch_Wired_MAB Authorization Profile



The Converged Access Branch Design also uses same authorization profile for MAB devices as shown in the Figure 11-12.



Security Group Access for BYOD

Revised: August 7, 2013

The following section describes the infrastructure used in this CVD and provides an outline of the two deployment scenarios used to enforce policies based on Security Group Tags. These deployment scenarios are not mutually exclusive and may be used together to satisfy an organization's requirements. Configuration details for the infrastructure are also provided.

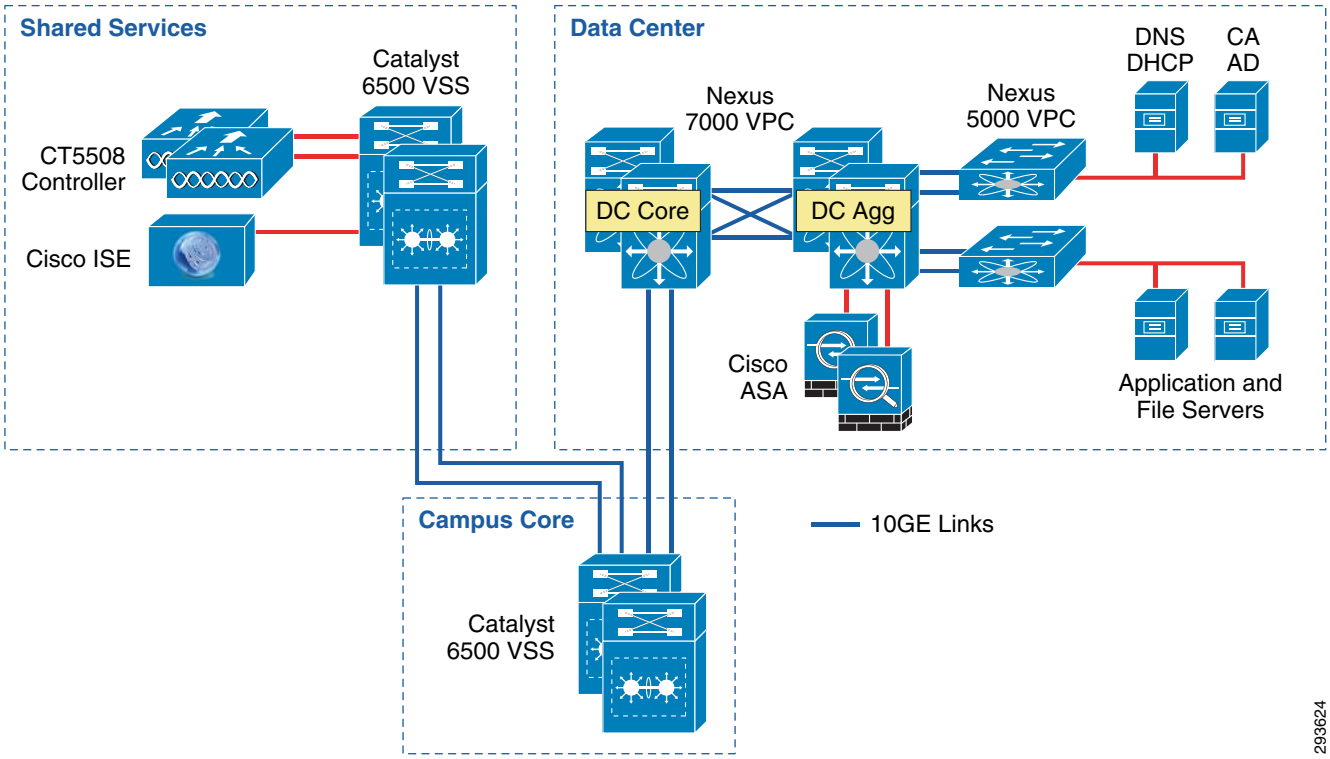
Unified Infrastructure Design to Support SGA

As described in SGT Deployment Scenarios in this CVD in Chapter 5, "Campus and Branch Network Design for BYOD," two specific infrastructure deployment scenarios are examined in this CVD. The first use case uses the SGA Policy defined at the Identity Services Engine and the resulting SGACLs are dynamically exchanged with the Catalyst 6500 and Nexus 7000 infrastructure. The second use case also uses the SGA Policy defined at the Identity Services Engine, but enforces this policy through the configuration of Security Group Firewall (SG-FW) policies defined on an ASA providing secure access to data center resources.

In both scenarios, campus wireless users/devices connecting through centralized CUWN CT5508 controllers configured for local mode have access to data center resources based on their authorized roles and enforced through the use of SGT-based policies as implemented in the two deployment scenarios.

Figure 12-1 depicts the infrastructure that is used for purposes of SGA validation within the CVD.

Figure 12-1 TrustSec Infrastructure for BYOD v2.5 CVD

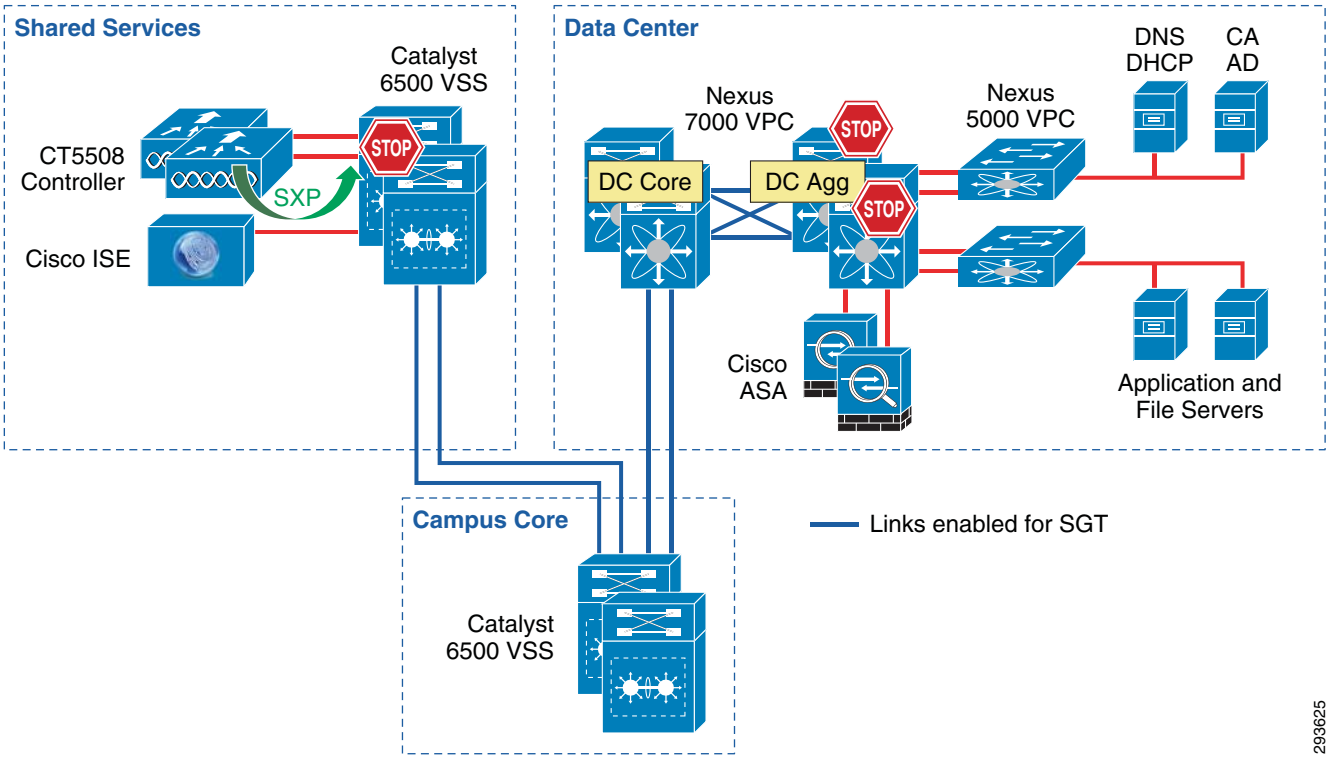


In Figure 12-1, the links extend between the Catalyst 6500 VSS in Shared Services to the Catalyst 6500 VSS in core and extends to the Nexus 7000 are 10GE links. On the Catalyst 6500s, WS-X6904 linecards with the FourX Adapters provide the 10GE interfaces while the N7K-M108X2-12L linecards provide the Nexus 7000 interfaces. The links between the Nexus 7000 and the Nexus 5548 are likewise connected to N7K-M108X2-12L linecards at the N7K and 10GE ports on the Nexus 5548. All other network connectivity for wireless controllers, ASA Firewalls, ISE, and the miscellaneous servers depicted are 1GE links.

Policy Configuration for SGACLs in Scenario 1

For Deployment Scenario 1, refer to Figure 12-2.

Figure 12-2 Infrastructure Deployment Scenario 1 SGT Enforcement



Deployment Scenario 1 requires that Security Group Tags are forwarded from the Shared Services Catalyst 6500 VSS, where the wireless controller is attached, through the core of the BYOD infrastructure enroute to servers located in the data center proper. In Figure 12-2, the links depicted in blue will be configured for SGT forwarding as well as manually configured for 802.1ae MACsec encryption. As previously discussed, the CT5508 wireless controller does not support native tagging on its 1GE interfaces, so a Security Group Tag Exchange Protocol (SXP) connection will be defined between the controller(s) and the Shared Services C6500 VSS switch as depicted above.

In this first scenario, wireless users, upon successful authentication and authorization, will be associated with a specific role and an IP to SGT mapping will be created on the wireless controller with the device's IP Address and the appropriate SGT. SXP will be used to communicate this mapping to the Shared Services Catalyst 6500s to which the wireless controllers are attached. As wireless user traffic egresses the Shared Services Catalyst 6500s, it will be tagged with the appropriate SGT learned via SXP from the wireless controller. As this traffic traverses the SGT-capable Core, this tag will be propagated hop-by-hop enroute to the Nexus 7000s comprising the data center switching infrastructure within which the various servers are located.

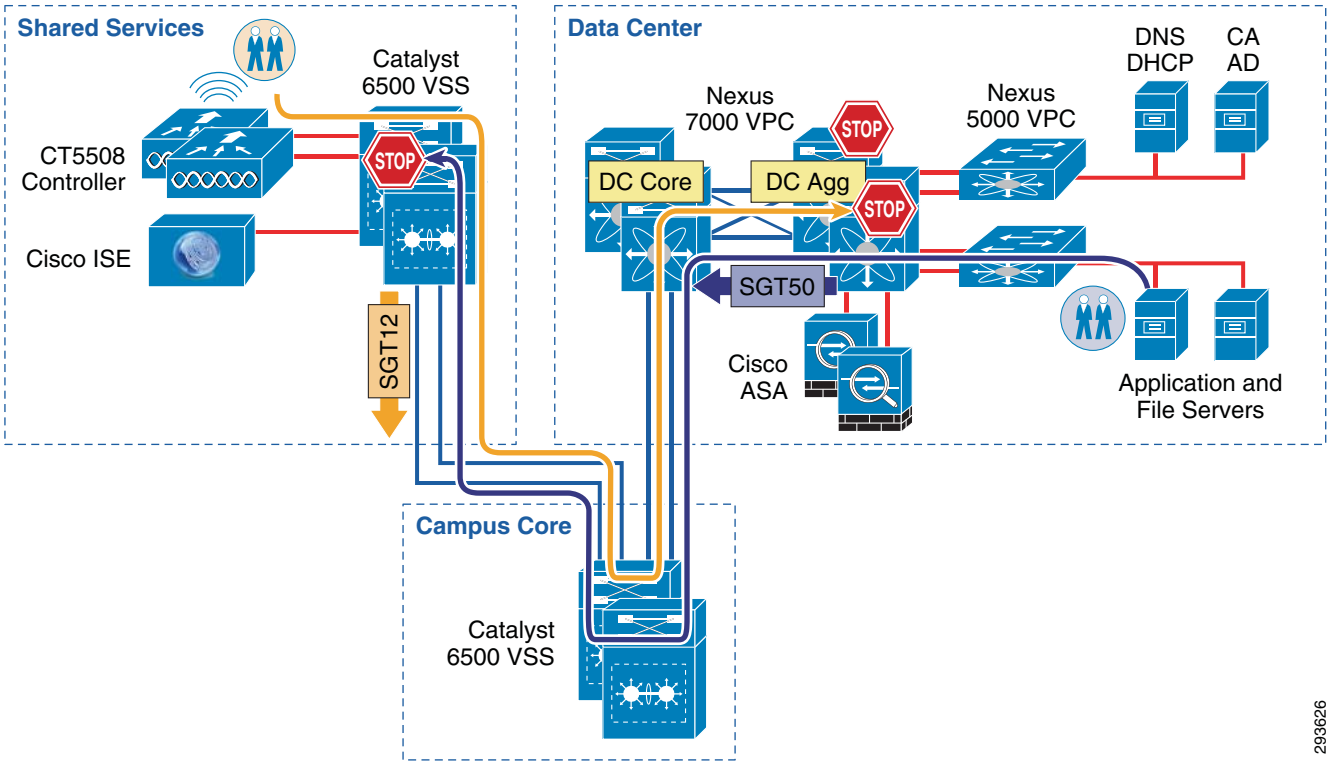
As 802.1X is not used to authenticate the servers residing in the Nexus data center infrastructure, the Server IP Address to SGT mapping can either be manually defined on the Nexus 7000 Data Center Aggregation switch or at the ISE server which would subsequently populate that mapping to the Nexus 7000. For purposes of the CVD, these mappings have been manually defined on the Nexus 7000 DC Aggregation Switch. As discussed in SGT Deployment Scenarios in this CVD in Chapter 5, "Campus and Branch Network Design for BYOD," there are other methods of associating traffic with a specific SGT on the Nexus 7000 platform.

As tagged user traffic arrives at the Nexus 7000 data center switch where the manual SGT mappings for the servers have been created, the traffic will be matched against TrustSec Policy (SGACL) defined either centrally at ISE or locally, as in the case of destination "Unknown" (SGT0), and will be either forwarded or dropped as applicable.

As discussed earlier, all server IP to SGT mappings have been manually created on the Nexus 7000 aggregation switches. As the servers are connected to the Nexus 5548 switches depicted in Figure 12-3, traffic from the Nexus 5548s egresses untagged as no mappings have been created there. Once this traffic passes through the Nexus 7000 Aggregation switch, the resident SGT mappings will be examined and the appropriate SGT imposed upon egress from the aggregation switch. In the event that traffic is initiated by a server associated with an SGT in the data center, the tagged traffic egresses the Nexus 7000 data center switches and traverses the Core and Shared Service infrastructure with the SGT propagated at each hop enroute to the destination, which is the wireless controller attached to the Shared Services 6500. Once the traffic arrives at the Shared Services 6500, the traffic will be matched against TrustSec Policy (SGACL) and will be either forwarded or dropped as defined.

Figure 12-3 depicts where SGACLs will be enforced in the Unified Access infrastructure.

Figure 12-3 Policy Enforcement in Deployment Scenario 1



The following major tasks are required for this deployment scenario and are outlined in the following sub-sections in Chapter 23, "BYOD Policy Enforcement Using Security Group Access":

1. Configuring ISE to Support TrustSec
2. Configuring ISE for Network Access Device Authentication
3. Configuring Network Access Devices for Authentication at ISE
   a. RADIUS Server Configuration on the Wireless Controller
   b. RADIUS Server Configuration on the Catalyst 6500
   c. RADIUS Server Configuration on the Nexus 7000
4. Catalyst 6500 Platform Specific Considerations
5. Configuring Switching Infrastructure to Support TrustSec with 802.1ae MACsec Encryption

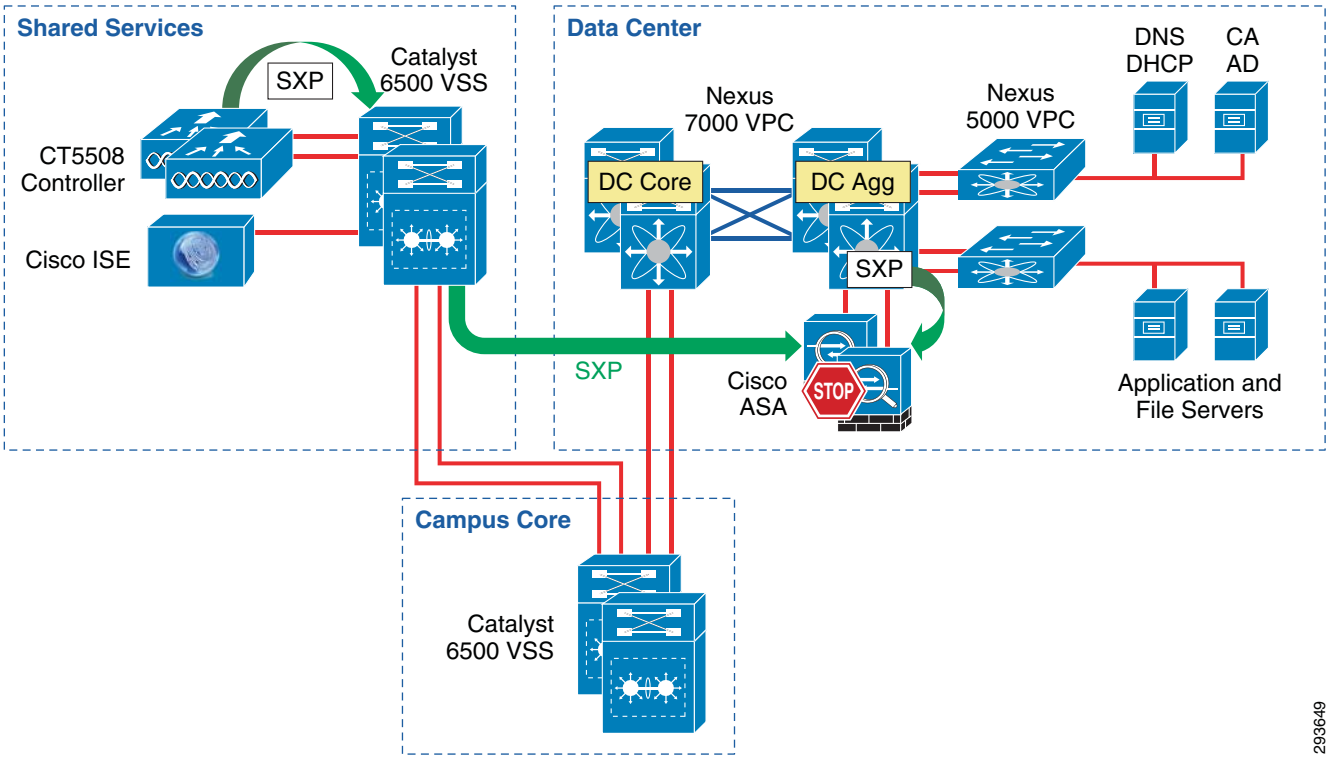
a. Catalyst 6500 Commands
b. Nexus 7000 Commands

6. Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers
   a. Wireless Controller Configuration
   b. Catalyst 6500 SXP Configuration
7. Configuring Static IP/SGT Bindings on Nexus Switches

Policy Configuration in Scenario 2

For the topology used in Deployment Scenario 2, refer to Figure 12-4.

Figure 12-4 Deployment Scenario 2 Configuration



With Deployment Scenario 2 an alternate means other than SGACLs is used to enforce SGA policy. In Scenario 2, the ASA running version 9.0 will be used as a Security Group Firewall (SG-FW) securing data center resources from outside access. As the ASA does not presently support Native SGT Tagging on its Ethernet interfaces, SXP must be used for it to learn IP/SGT mappings from other areas of the network where they have been dynamically learned or statically configured.

As in the case of the first deployment scenario, wireless users, upon successful authentication and authorization, will be associated with a specific role and an IP to SGT mapping will be created on the wireless controller with the device's IP Address and the appropriate SGT. Security Group Tag Exchange Protocol (SXP) will be used to communicate this mapping to the Shared Services Catalyst 6500s to which the wireless controllers are attached.

Unlike Scenario 1 however, the 10GE infrastructure between the Shared Services Catalyst 6500 VSS and the data center does not need to be enabled to support Security Group Tags or SGACLs. SXP will be used instead to re-advertise the mappings the Shared Services Catalyst 6500 VSS learned from the wireless controllers to the ASA Firewall.

The reason that a multi-hop SXP configuration is used from the wireless controller is primarily due to the fact that concentrating all of the controllers SXP advertisements at the 6500 VSS switch in Shared Services for single advertisement elsewhere in the network just provides a cleaner approach; it would be entirely possible to create the SXP peering directly between the wireless controller and the ASA firewall. The only other consideration lies in the fact that the WLC-5508 controllers used in this guide, as well as the WiSM2 only support four SXP connections, whereas the 6500 scales far beyond that.

In addition to the SXP Peering between the Shared Services 6500s and the ASA, the Nexus 7000 aggregation switches will require an SXP Peering to advertise SGT mappings that have been configured on them. It is by virtue of these SXP advertisements that the ASA is capable of inspecting the traffic from various devices and associating the appropriate tag for subsequent SG-FW policy enforcement as the ASA's interfaces are not TrustSec aware and incapable of manipulating the SGT.

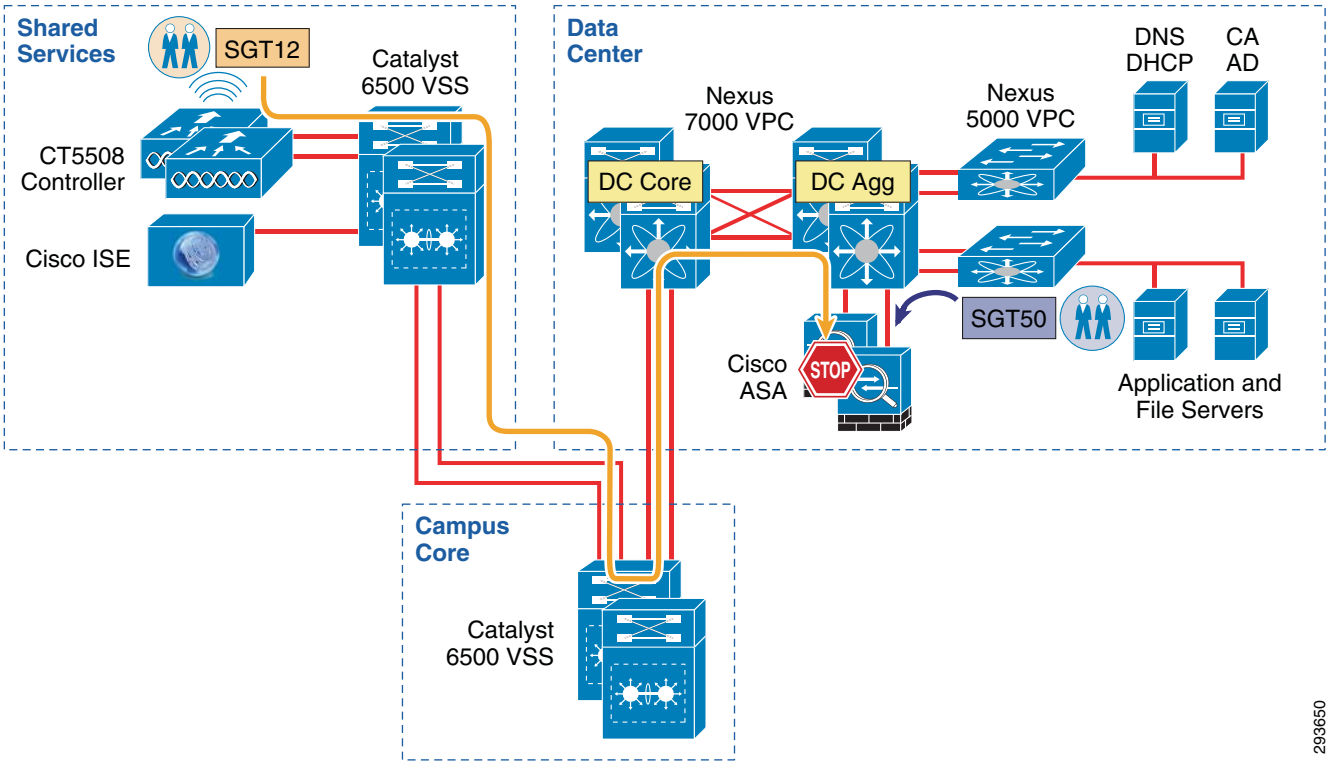
As previously discussed, the ASA Firewall that will be used to enforce SG-FW policies must be manually configured with SGT policies as Network Device Admission Control (NDAC) is presently not supported in the ASA and is therefore unable to acquire these policies dynamically from ISE. For more information regarding NDAC, refer to Chapter 23, "BYOD Policy Enforcement Using Security Group Access." These policies as defined in Chapter 23, "BYOD Policy Enforcement Using Security Group Access" for Deployment Scenario 2 can be configured via CLI, ASDM, or a management platform such as Cisco Security Manager.

As wireless traffic egresses the Shared Services Catalyst 6500s en route to the data center, the traffic will be untagged and will simply pass through the Core, enter the data center switching infrastructure, and ultimately arrive at the ASA Firewall where the appropriate SG-FW policy will be enforced.

Should any traffic be sourced from a server in the data center, it will likewise egress the Nexus 7000 aggregation switch untagged and be forwarded to the ASA firewall where any applicable SG-FW policy will be enforced.

Figure 12-5 depicts the infrastructure used in Deployment Scenario 2 and the means by which security group policies will be enforced.

Figure 12-5 SGA Policy Enforcement Using SXP and SG-FW



The following major tasks are required for this deployment scenario and are outlined in the following sub-sections in Chapter 23, "BYOD Policy Enforcement Using Security Group Access":

1. Configuring ISE to Support TrustSec
2. Configuring ISE for Network Access Device Authentication
3. ISE Configuring the Network Devices for Integration with ISE
   a. RADIUS Server Configuration on the Wireless Controller
   b. RADIUS Server Configuration on the ASA Firewall
   c. RADIUS Server Configuration on the Nexus 7000
4. Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers
   a. Wireless Controller Configuration
   b. Catalyst 6500 SXP Configuration
   c. Nexus 7000 SXP Configuration
   d. ASA SXP Configuration
5. Configuring Static IP/SGT Bindings on Nexus Switches
6. Configuring SG-FW Role-Based Policies at ASA

TrustSec Summary

For information regarding the detailed, platform-specific configuration steps, refer to the TrustSec section in Chapter 23, "BYOD Policy Enforcement Using Security Group Access."



Note Patch 1 for ISE 1.2 **MUST** be installed in order for NDAC (Network Device Admission Control) to function properly between the network device and ISE. Without Patch 1, the network device will be unable to authenticate with ISE in order to derive TrustSec environment data, PAC file, and security group policies when CTS Manual Mode is configured and, additionally, the credentials required to authenticate peers/TrustSec links when CTS Dot1x Mode is configured. Refer to the ISE 1.2 Release Notes for additional information regarding this **very important** information.



Mobile Device Manager Integration for BYOD

Revised: August 7, 2013

The Cisco ISE can be configured to integrate with third-party Mobile Device Manager (MDM) products through an XML-based API. This allows network policy decisions based on mobile device posture that can include PIN lock, storage encryption, or registration status. In this release, both Apple and Android devices are supported. Configuring the infrastructure to support this functionality involves setting up ISE to send API requests to the MDM and configuring the MDM to accept these requests. Chapter 6, "Mobile Device Managers for BYOD" includes a discussion of the communications between the various components. Some MDM configurations, including device compliance policy, are discussed in general terms in this section. Detailed partner specific information can be found in supporting documentation at: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/own_device.html.

An overview of the topology common with the MDM architecture is presented below. The two basic models that are detailed in this section are an On-Premise model and a Cloud-based SaaS model. The components are similar except that the cloud model can also include an on-premise component to facilitate the integration with the enterprise.

Establishing IP Connectivity for an On-Premise MDM

Typically the on-premise MDM resides in the DMZ or some location where mobile devices can establish inbound connections. This allows the MDM to monitor the device's posture while the device is on the outside of the firewall. Without this access, the device would need to be placed on the network and interrogated prior to establishing the posture compliance of the device. The device does not automatically update the server whenever it joins a new network, therefore this interrogation would need to be manually initiated by the enterprise. If the MDM is located in the data center, some provision is required to allow inbound TCP sockets from the Internet. The specific ports vary based on the MDM partner and are detailed in the supporting documentation on Design Zone.

In addition to inbound sessions from the devices, the MDM needs to establish outbound connections to the push servers. The MDM uses the push service to locate and notify the device of changes to the MDM policy. Apple refers their service as the Apple Push Notification Service (APNS) and requires an Apple signed certificate to authenticate the MDM. Google refers to their service as Google Cloud Messaging for Android (GCM). This service replaces the older Cloud to Device Messaging Framework (C2DM). Both Apple and Android incorporate the push service into the device's operating system (OS) to allow the MDM server to communicate with the MDM client application. Apple devices also allow the MDM to communicate with the OS MDM API with the appropriate credentials. Both require the end user to establish an account with either Google or Apple respectively. This account effectively binds a device list to a user.

The MDM will also host a user-centric My Devices Portal to allow users to log into the MDM and manage some aspects of their device. This is similar but distinct from the My Devices Portal offered by ISE and serves a different purpose. Users may attach from either the mobile device or their standard desktop. The MDM web server can be configured with ACLs to restrict access to the My Devices Portal page from specific source address. For example, it is possible to block Internet access to the portal. The same is true for the administrator website.

The MDM will also receive inbound HTTPS session on port 443 by default to support the API used by ISE. In contrast to the MDM placement, ISE should be located in the data center. Firewall policy should be set to allow TCP 443 sessions that are initiated from ISE towards the MDM server. The MDM will have a default route pointing towards an outbound firewall and a more specific route to ISE pointing towards an inbound firewall. The majority of MDM partners support on-premise deployments on VM servers that may support multiple interfaces. It is possible that the route to ISE may be over a dedicated link. The topology of the DMZ should match the established corporate policy for servers. Typically the MDM will also allow the administrator to protect the API with an ACL. In this case, the ACL could be configured to permit ISE but deny any other connections.

ISE supports the use of a proxy for external connections. Currently the proxy configuration is globally configured. If ISE is required to use a proxy for the feed service, then it will also direct MDM requests to the proxy. This could cause connectivity issues between ISE and an on-premise MDM. In this scenario, the proxy configuration will require careful review to ensure that the ISE can connect to the MDM via the proxy.

Establishing IP Connectivity for a Cloud-Based MDM

Subscribing to an online MDM service simplifies many of the connectivity issues, especially between the mobile devices and the device manager. Because personal mobile devices spend the majority of time connected to the public Internet, choosing this model offers some advantages over a traditional on-premise model. The Apple APNS or Google GCM are also simplified when a cloud model is in use. The enterprise will still need to generate a certificate-signing request and present that to Apple prior to using the APNS service. This is explained in the partner-specific supporting documentation. However with the advantages realized with a cloud deployment, there are also challenges with respect to enterprise integration, specifically the corporate directory structure. Without any integration, a separate and dedicated user database would need to be established and maintained on the MDM servers. Typically in the cloud model, the enterprise will establish a small integration server that resides in the DMZ and serves as a proxy to a secure LDAP binding. This is explained in the partner-specific supporting documentation. With the exception of this additional server, all of the other components found in the on-premise model are present in a cloud model.

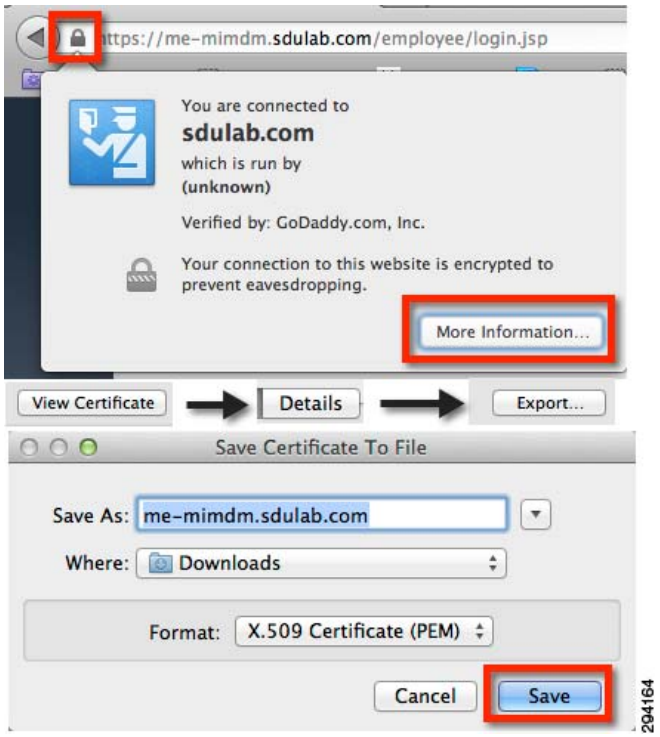
The primary concern is the HTTPS connection between ISE and the cloud-based MDM server, which is outbound from ISE. Corporate firewalls need to allow the ISE server siting in the data center to establish outbound HTTPS servers to the MDM server. The MDM partner may be able to provide a range of destination subnets if outbound sessions are restricted from data center servers such as ISE. Before ISE will trust the MDM server, the MDM server's certificate should be imported into the local ISE certificate store (this is explained below). The MDM service will provide the URL of the API. It is this certificate on this site that should be imported. In addition, users will need to be able to establish outbound HTTPS connections to the My Devices Portal page on the cloud-based MDM server. This would only be a concern in environments where users are not allowed access to Internet websites. If WCS or ScanSafe is in use, then the enterprise should confirm that the MDM site has a reputation score that exceeds the threshold needed for access or the site should be manually added to the permitted whitelist. Routing is straightforward. ISE and the user devices will follow the default route towards the Internet. The session may flow over NAT boundaries without requiring a NAT fix-up.

Connectivity must also be provided to mobile clients that have been quarantined to the MDM and to either the Apple Push Service or Google Cloud Messaging Servers. This allows the MDM to communicate with the device as needed to update the device's posture information on the server. In some situations, the mobile client may also need access to Google Play and the Apple AppStore to download required applications such as the MDM mobile client.

Configure ISE to Authenticate the MDM API

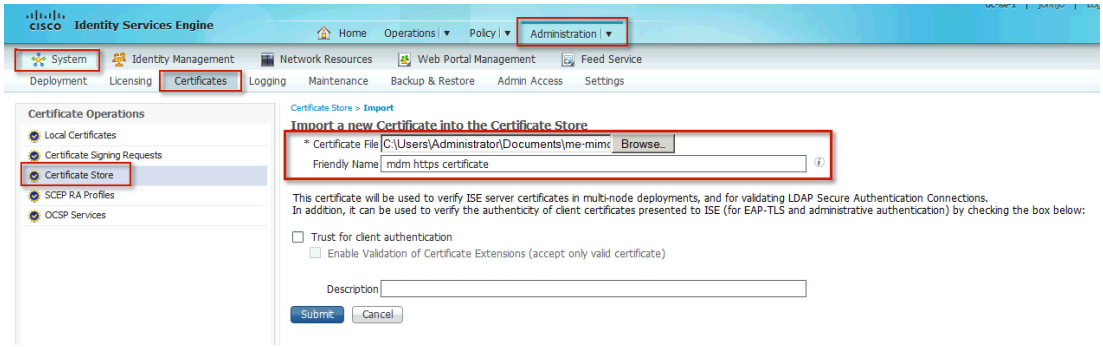
Prior to configuring the MDM, ISE must trust the HTTPS certificate presented by the MDM website. In either the cloud or on-premise deployment model, this can be accomplished by installing the MDM's HTTPS certificate in the ISE certificate store. The easiest method is to browse to the MDM server, export the HTTPS certificate, and then import it into ISE. Figure 13-1 shows this in Firefox, however the procedure may be different for other browsers.

Figure 13-1 Exporting MDM Certificate



Once the certificate has been saved to the local disk, the user will import it into the local certificate store on ISE. By default, the browser will save the certificate file with a name based on identity contained in the certificate, which is typically the FQDN of the site. The file extension could be .com, which is a well-known MS-DOS extension, making the cert more difficult to locate. While this does not affect importing the certificate, it could make browsing for the file on the hard drive less obvious. Importing the certificate into ISE is shown in Figure 13-2.

Figure 13-2 Importing Certificate into ISE



If ISE and the MDM are using the same CA, then importing the MDM SSL certificate may not be required. ISE does not maintain a system list of well-known public root certificates, therefore all trust relationships must be established by the administrator. Installing the MDM SSL certificate is the simplest approach and is shown here to ensure success.

Creating the MDM API User Account

In addition to the certificate, ISE will need a user account on the MDM that will allow access to the API. The previously installed certificate allows ISE to attach to the MDM via HTTPS, which will encrypt all data exchanges between ISE and the MDM, including the API credentials. All of the MDM partners support a local user account that can be granted API privileges. Some vendors may allow the account to be defined in an external data store such as Active Directory. This could be useful if ISE is using the same account to access AD or other resources and centralized machine account management is in use. In all cases, the API user account should be protected by strong passwords. For specific guidance on setting up this account, refer to the partner-specific supporting documentation or the partner MDM Administrator guide.

There are two account issues that may prevent the API from functioning properly:

- Incorrect username or password combination
- Defined user has not been granted API access

Setting Up the MDM Connection

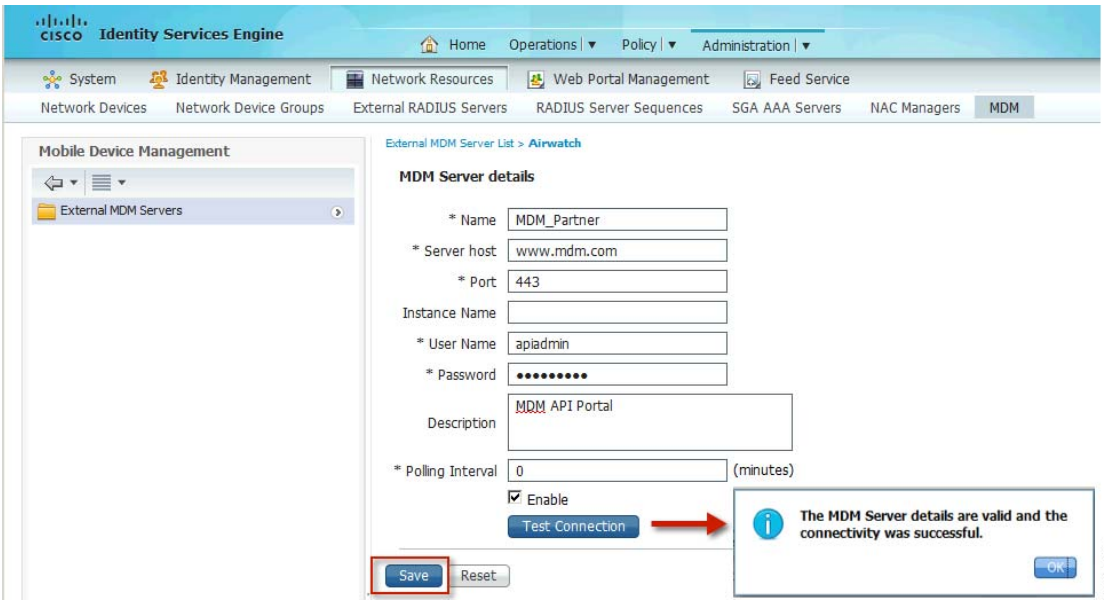
ISE will contact the MDM to gather posture information about devices or to issue device commands such as corporate wipe or lock. The session is initiated from ISE towards the MDM server. The URL for the MDM server is typically the same as the admin page and will be the same website used to export the certificate. The directory path is handled automatically by the system and is not specified as part of the configuration. The instance is used in multi-tenant deployments more commonly found when subscribing to a cloud service. The field should be left blank unless the cloud provider has instructed otherwise. The port will typically be TCP 443 for HTTPS. Typically the MDM cannot be configured to listen on a specific port for API users. Any change will also impact both the admin and user portal pages.

The polling interval specifies how often ISE will query the MDM for changes to device posture. By default, this is set to 0 minutes, effectively disabling polling. Polling can be enabled to periodically check the MDM compliance posture of an endpoint. If the device is found to be out of compliance and the device is associated to the network, then ISE will issue a CoA forcing the device to re-authenticate. Likely the device will need to remediate with the MDM, although this will depend on how the policy is

configured. Note that MDM compliance requirements are configured on the MDM and are independent of the policy configured on ISE. It is possible, although not practical, to set the polling interval even if the ISE policy does not consider this dictionary attribute. The advantage of polling is that if a user takes the device out of MDM compliance, they will be forced to reauthorize that device. The shorter the window, the quicker ISE will discover the condition. There are some considerations to be aware of before setting this value to an aggressively low value. The MDM compliance posture could include a wide range of conditions not specific to network access. For example, the device administrator may want to know when an employee on a corporate device had exceeded 80% of the data plan to avoid overage charges. In this case, blocking network access based solely on this attribute would aggravate the MDM compliance condition and run counter the device administrator intentions. In addition, the CoA will interrupt the user WiFi session, possibly terminating real-time applications such as VoIP calls. The recommendation is to leave the polling interval at 0 until a full understanding of the MDM's configuration is complete. If the polling interval is set, then it should match the device check-in period defined on the MDM. For example, if the MDM is configured such that devices will report their status every four hours, then ISE should be set to the same value and not less than half of this value. Over sampling the device posture will create unnecessary loads on the MDM server and reduced battery life on the mobile devices.

Finally, the enable check box will be set to active on one MDM server. It is possible to save multiple configurations, but only one can be active at a time. Figure 13-3 shows a typical configuration.

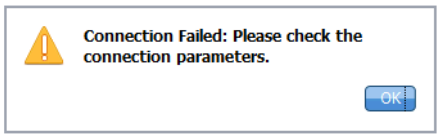
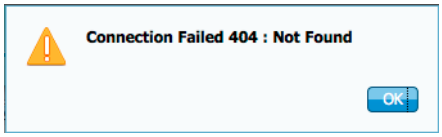
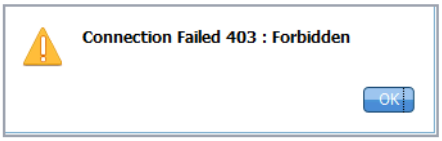
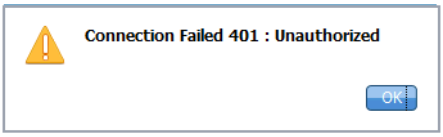
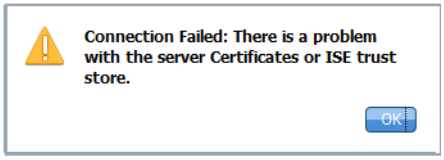
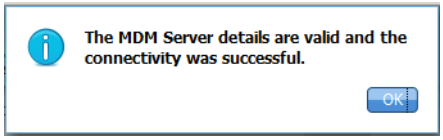
Figure 13-3 MDM Server Details



Verifying the MDM Connection

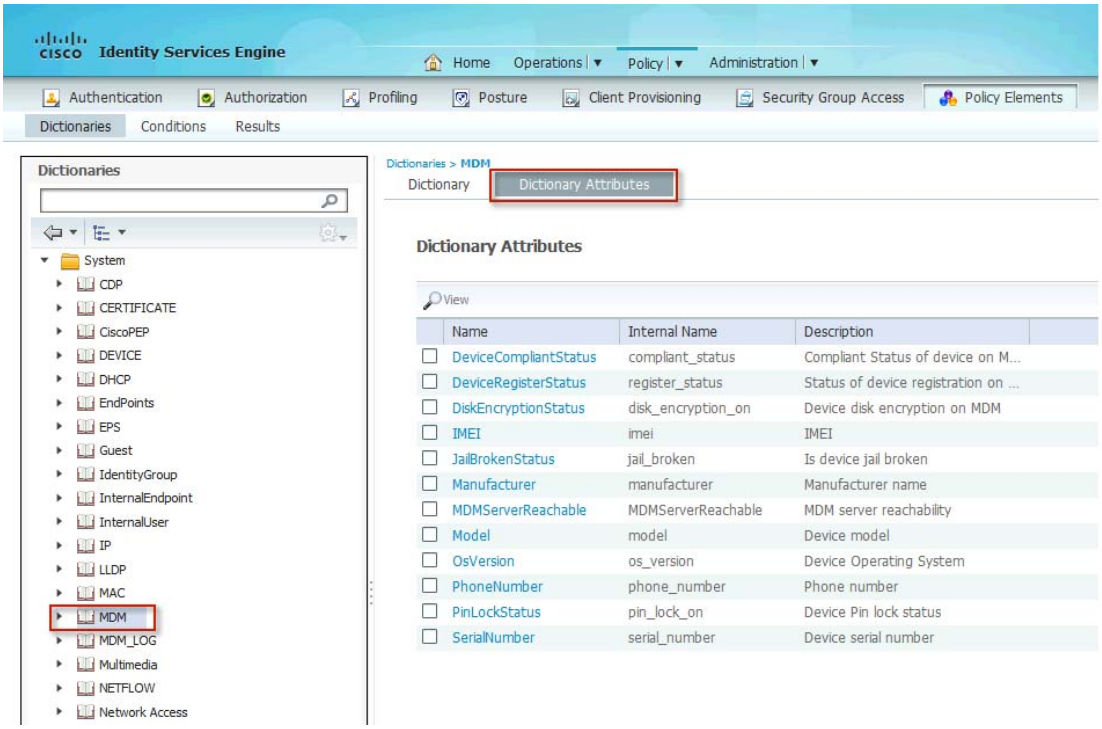
The test button will establish a connection to the MDM and attempt to authenticate using the configured credentials. This should be complete prior to saving the settings. If not, then the save button will validate the settings. If any errors are encountered, the MDM Enable button will be deselected prior to saving. If any error messages are presented, the administrator can refer to Table 13-1 for guidance in correcting the setup. In order to re-run the test on a previously validated server, the user should deselect the Enable checkbox, save, and then re-enable the checkbox.

Table 13-1 Common MDM Connection Error Codes

	A routing or firewall problem exists between ISE located in the data center and the MDM located in either the DMZ or Cloud. The firewall's configuration should be checked to confirm HTTPS is allowed in this direction.
	The most likely cause of an HTML 404 error code is that an instance was configured when it was not required, or that the wrong instance has been configured.
	The user account setup on the MDM server does not have the proper roles associated to it. Validate that the account being used by ISE is assigned the REST API MDM roles as shown above.
	The user name or password is not correct for the account being used by ISE. Another less likely scenario is that the URL entered is a valid MDM site, but not the same site used to configure the MDM account above. Either of these could result in the MDM server returning an HTML code 401 to ISE.
	ISE does not trust the certificate presented by the MDM website. This indicates the certificate was not imported to the ISE certificate store as described above or the certificate has expired since it was imported.
	The connection has successfully been tested. The administrator should also verify the MDM dictionary has been populated with attributes.

After successfully configuring the MDM, the ISE policy dictionary will contain the attributes needed to create policy. The user can verify the dictionary by clicking **Policy > Dictionaries > System > MDM**, as shown in Figure 13-4.

Figure 13-4 Dictionary Attributes



Configuring the MDM

In addition to the API user account needed for ISE, there are several other administrative tasks that need to be accomplished on the MDM, such as signing and installing the APNS certificates before ISE can issue device actions through the API. The partner-specific supporting documentation has additional details on the minimum requirements. The MDM can also be configured to integrate with the corporate directory structure through LDAP. The administrator should review the MDM installation and administration guides to bring the MDM system into a fully functional state.



PART 4

BYOD Use Cases



Summary of BYOD Use Cases

Revised: August 7, 2013

This part of the CVD focuses on implementing the rules previously defined in business policies. These rules translate into four different use cases that focus in providing differentiated access to corporate, personal-owned, and guest devices and highlight the network infrastructure as the enforcement point for BYOD policies.

There are numerous ways to enable a BYOD solution based on the unique business requirements of a specific organization. While some organizations may take a more open approach and rely on basic authentication, other organizations will prefer more secure ways to identify, authenticate, and authorize devices. A robust network infrastructure with the capabilities to manage and enforce these policies is critical to a successful BYOD deployment.

The following components and configuration steps are discussed to support different BYOD use cases:

- Digital Certificates
- Microsoft Active Director authentication
- Wireless Controllers (Unified and Converged Access)
- Identity Services Engine
- Access Layer Switches
- API Integration with Mobile Device Managers

This part of the CVD includes the following chapters:

- BYOD Enhanced Use Case—Personal and Corporate Devices—This use case provides network access for personal devices and corporate-issued devices. It provides unique access (Full, Partial, and Internet Only) based on different conditions analyzed by the ISE. ISE relies on the network infrastructure to enforce unique permissions.
- BYOD Limited Use Case—Corporate Devices—This use case focuses on identifying corporate-issued devices and providing Full Access to Network Resources.
- BYOD Advanced Use Case—Mobile Device Manager Integration—The API integration with third party Mobile Device Managers allows the ISE to query for additional posture information on endpoints. This information translates into more granular authorization rules and allows for more visibility into the endpoint.
- BYOD Basic Access Use Case—An extension all the traditional wireless guest access, this use case presents an alternative where the business policy is not to on-board personal wireless devices but still provides access to network resources.

- User Experience—How To On-board a BYOD Device—Providing a positive user experience is important for any BYOD deployment. Employees should be provided with a simple way to on-board their devices and enable the necessary security features with minimum manual intervention. This chapter captures a typical user interaction with ISE during the on-boarding process.



BYOD Enhanced Use Case—Personal and Corporate Devices

Revised: August 7, 2013

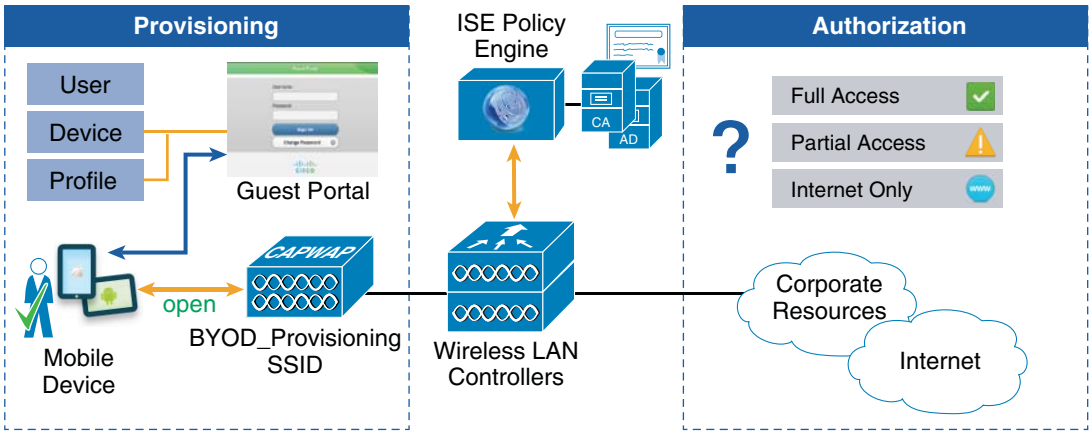
The BYOD Enhanced use case is a super-set of the BYOD Limited use case, covering both personal and corporate devices. This chapter defines the design scenarios for deploying BYOD for personal device access and the design considerations for each scenario. It also highlights how to deny access to personal devices based on their device type. Since Chapter 16, "BYOD Limited Use Case—Corporate Devices" provides design guidance for corporate devices, that design information for corporate devices is not duplicated in this chapter.

One of the main objectives of a BYOD solution is to provide a simple way for employees to on-board their personal devices without requiring assistance from IT. Since the BYOD needs of the majority of employees will be met with either simple Internet access or Partial access, the only time an employee requires assistance from IT is when they require full access to corporate resources.

Cisco ISE provides different ways to define security policies and determine what network resources each employee is allowed to access. The security policies are then enforced throughout the network infrastructure. The ISE feature set is extremely flexible, enabling different business policies to be enforced. This chapter explains the steps to on-board personal devices and how to apply different policies.

Figure 15-1 shows how a personal device is profiled and registered by ISE and the different network components (WLC, AD, CA) that play a role in the process. Different conditions are evaluated to provide the proper authorization and access to network resources, including digital certificates, Active Directory groups, etc.

Figure 15-1 Provisioning Personal Devices



Note Unless otherwise specified, in the figures throughout this chapter "Wireless LAN Controllers" refers to either standalone devices such as the Cisco Flex 7500, CT5508, and CT5760 wireless controller or to wireless LAN controller functionality integrated within Catalyst 3850 Series converged access switches.

Figure 15-2 shows how a personal device is restricted from accessing the network. Once the device connects and regardless of user authentication, ISE profiles the device and enforces the DenyAccess authorization rule.

Figure 15-2 Deny Access

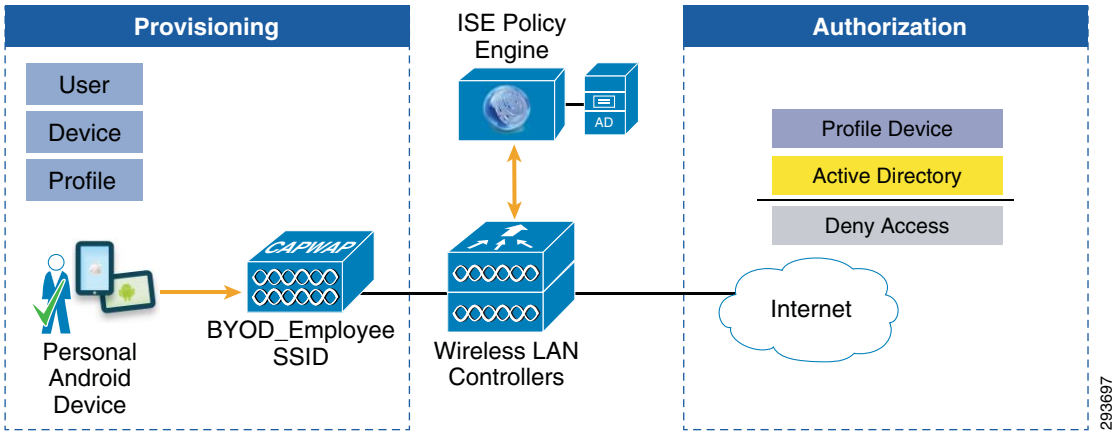


Figure 15-3 shows the different permission levels configured in this chapter. These access levels are enforced using various mechanisms, including Access Control Lists (ACLs) in the Wireless LAN Controllers (WLCs) or Catalyst switches, Security Group Tag (SGT) assignment in the WLCs, and VLAN assignment in the Catalyst switches along with FlexConnect ACLs in access points.

Figure 15-3 Permission Levels for Personal Devices

	Permission	Access
	Full Access	Internet plus all corporate resources
	Partial Access	Internet plus some corporate applications
	Internet Only	Internet Only
	Deny Access	Explicitly deny network access

292642

Active Directory Groups

Active Directory groups can be used as an additional way to grant differentiated access to users. This chapter relies on the following three AD groups:

- BYOD_Full_Access—Members of this group are granted full access to network resources.
- BYOD_Partial_Access—Members of this group are granted partial access to network resources. With this permission, users are able to access the Internet and a subset of corporate applications, such as email, corporate directory, travel tools, etc.
- Domain Users—All users are members of this system-generated group by default. Employees that are not members of the above mentioned groups are granted Internet access.

This model could easily be expanded to include other user groups with similar access requirements. A good example would be to create a new group and access list to grant access to contractors or partners.

Figure 15-4 highlights the different access policies validated for this chapter, along with the different requirements and permissions granted by each policy. These policies, along with detailed configurations, are explained later in this chapter.

Figure 15-4 Access Policies and Permissions

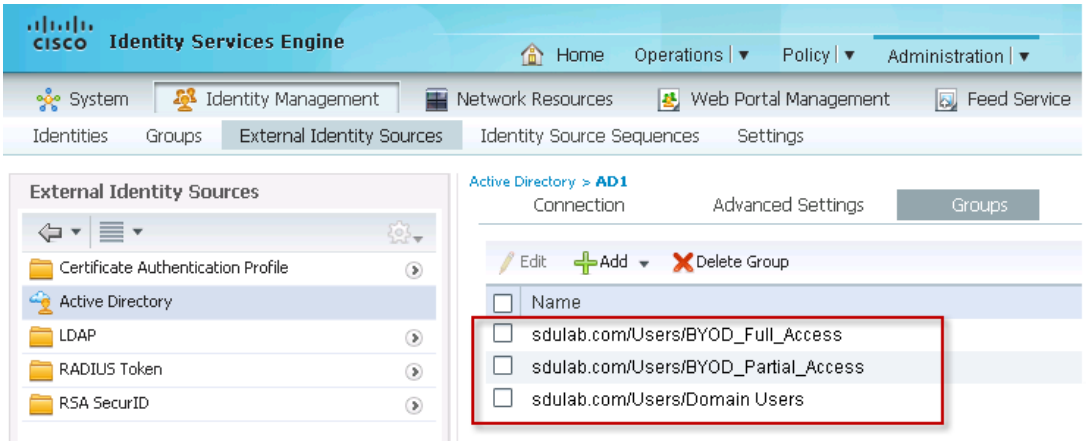
Policy	Location	AD Group	Profile	Permission	
Full Access	Campus/Branch/SGT	BYOD_Full_Access		Full	
Partial Access	Campus/Branch/SGT	BYOD_Partial_Access		Partial	
Internet Only	Campus/Branch/SGT	Domain Users		Internet Only	
Deny Android Devices			Android	Deny	

293698

Note The location SGT refers to a Wireless LAN Controller which is dedicated for the purpose doing SGT only.

To configure the Active Directory groups that can be available for use in the authorization policy conditions, click **Administration > Identity Management > External Identity Sources > Active Directory > Groups** and check the boxes next to the groups that will be used in the policy conditions and rules. Figure 15-5 includes the groups used in this design guide.

Figure 15-5 Active Directory Groups

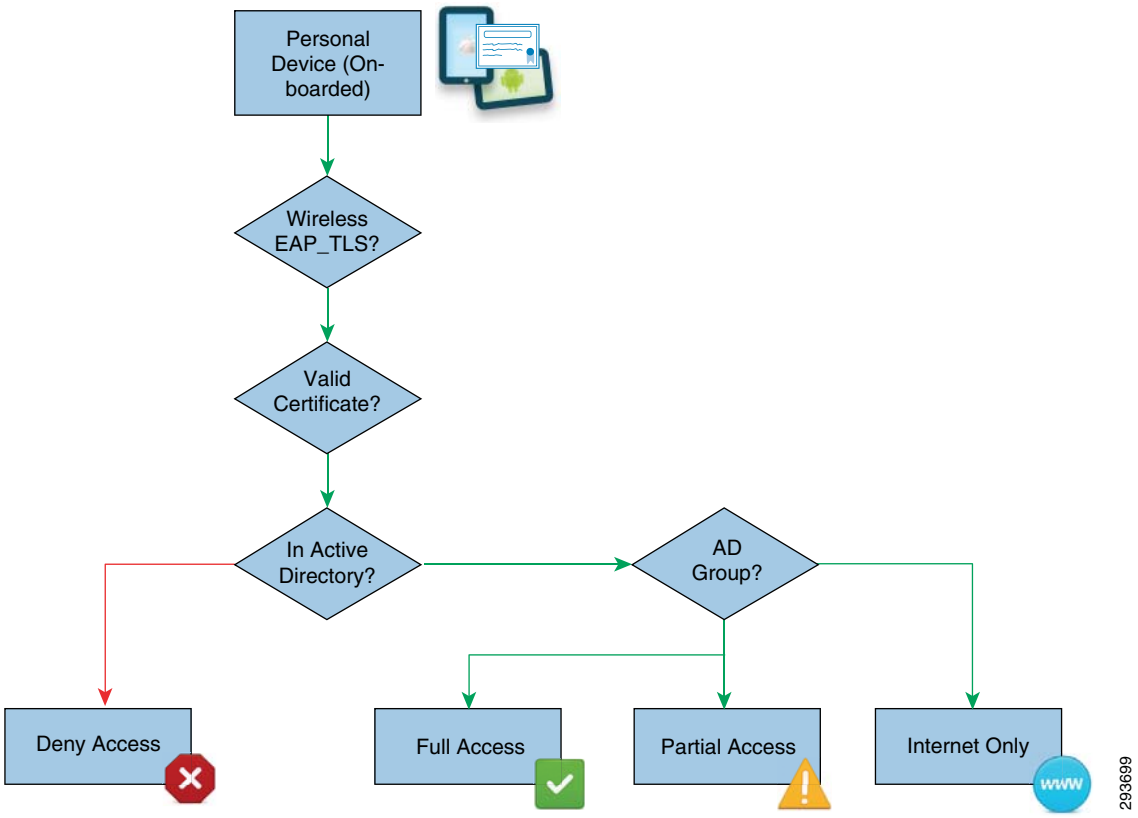


This chapter assumes that after employees have on-boarded their devices, they will connect to the BYOD_Employee SSID. Figure 15-6 highlights the connectivity flow for personal devices.

If the endpoint connected from a wireless 802.1X EAP-TLS SSID and has a valid certificate, based on their AD group membership, employees will get:

- Full Access if they belong to the BYOD_Full_Access group.
- Partial Access if they belong to the BYOD_Partial_Access group.
- Internet Access if they are a valid Active Directory member (Domain Users group)

Figure 15-6 Personal Device BYOD Access



Distributing Digital Certificates

Digital signatures, enabled by public key cryptography, provide a means to authenticate devices and users. In public key cryptography, such as the RSA encryption system, each user has a key pair containing both a public and a private key. The keys act as complements and anything encrypted with one of the keys can be decrypted with the other.

A digital signature is encrypted with the sender's private key. The signature must be verified to confirm the sender's identity. This is done by the receiver, who decrypts the signature with the sender's public key. If the signature sent with the data matches the result of applying the public key to the data, the validity of the message is established.

This process relies on the receiver having a copy of the public key of the sender and a high degree of certainty that this key belongs to the sender, not to someone pretending to be the sender.

Deploying digital certificates on mobile devices requires a unique process, as many of these devices do not natively support all the features and functionality to create, download, and install digital client certificates in the same manner as traditional PC-based devices. At the same time, some endpoints do not support Simple Certificate Enrollment Protocol (SCEP) natively.

For example, for users to install digital client certificates using SCEP on Apple iOS devices, the IT administrator needs to manually create the configuration profile using the iPhone Configuration Utility and distribute the profile to user devices via email, USB, or web pages.

Traditional full-featured PC-based devices are more apt to take advantage of the many services, such as Microsoft's NDES, to provide certificate enrollment. However, with the onset of many Android and Apple iOS devices on the market, it cannot be assumed that these devices can natively interoperate with many of the enterprise services currently deployed.

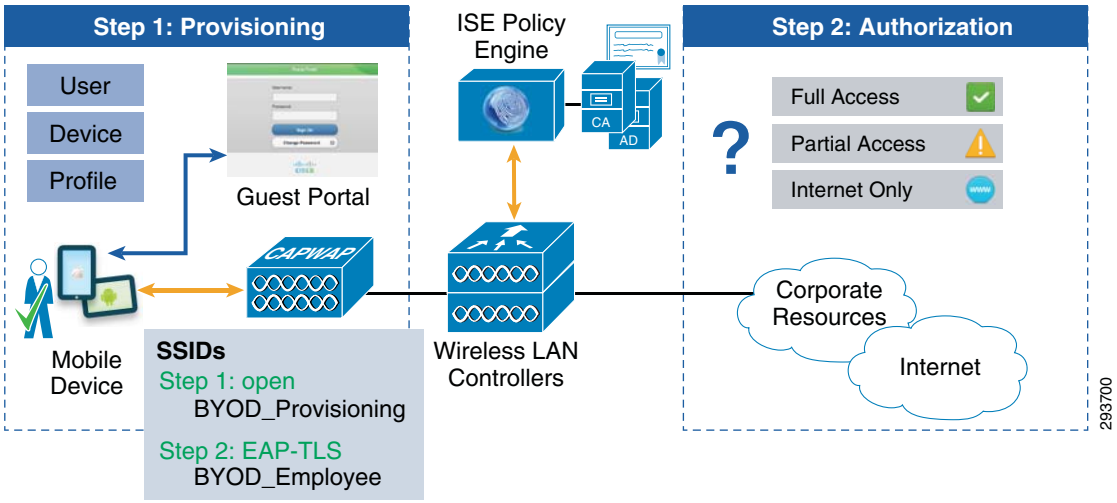
ISE solves this problem by distributing digital certificates to endpoints using the SCEP Proxy feature, which allows endpoints to obtain digital certificates through ISE. Moreover, this feature is combined during the initial registration process, thereby preventing different registration steps. The next section on mobile devices discusses how the endpoints obtain their digital certificates during the registration process.

Mobile Device On-boarding

Deploying digital certificates to endpoint devices requires a network infrastructure that provides the security and flexibility to enforce different security policies. Figure 15-7 highlights the general steps that are followed when a mobile device connects to the network using dual SSIDs.

1. A new device connects to a provisioning SSID. This SSID (open or secured with PEAP) is configured to redirect the user to the Guest Registration portal.
2. The certificate enrollment and profile provisioning begins once the user is properly authenticated.
3. The provisioning service requests information from the mobile device user and provisions the configuration profile, which includes a WiFi profile with the required parameters to connect to the secured Employee SSID.
4. For subsequent connections, the device uses the Employee SSID and is granted access to network resources based on different ISE authorization rules.

Figure 15-7 Enrollment and Provisioning—Dual SSID



The on-boarding steps may also be configured with a single SSID used for provisioning and secure access. The general steps followed when the mobile device connects are similar, redirecting the user to the Guest registration portal and provisioning the device with a digital certificate and configuration profiles.

Provisioning Flows

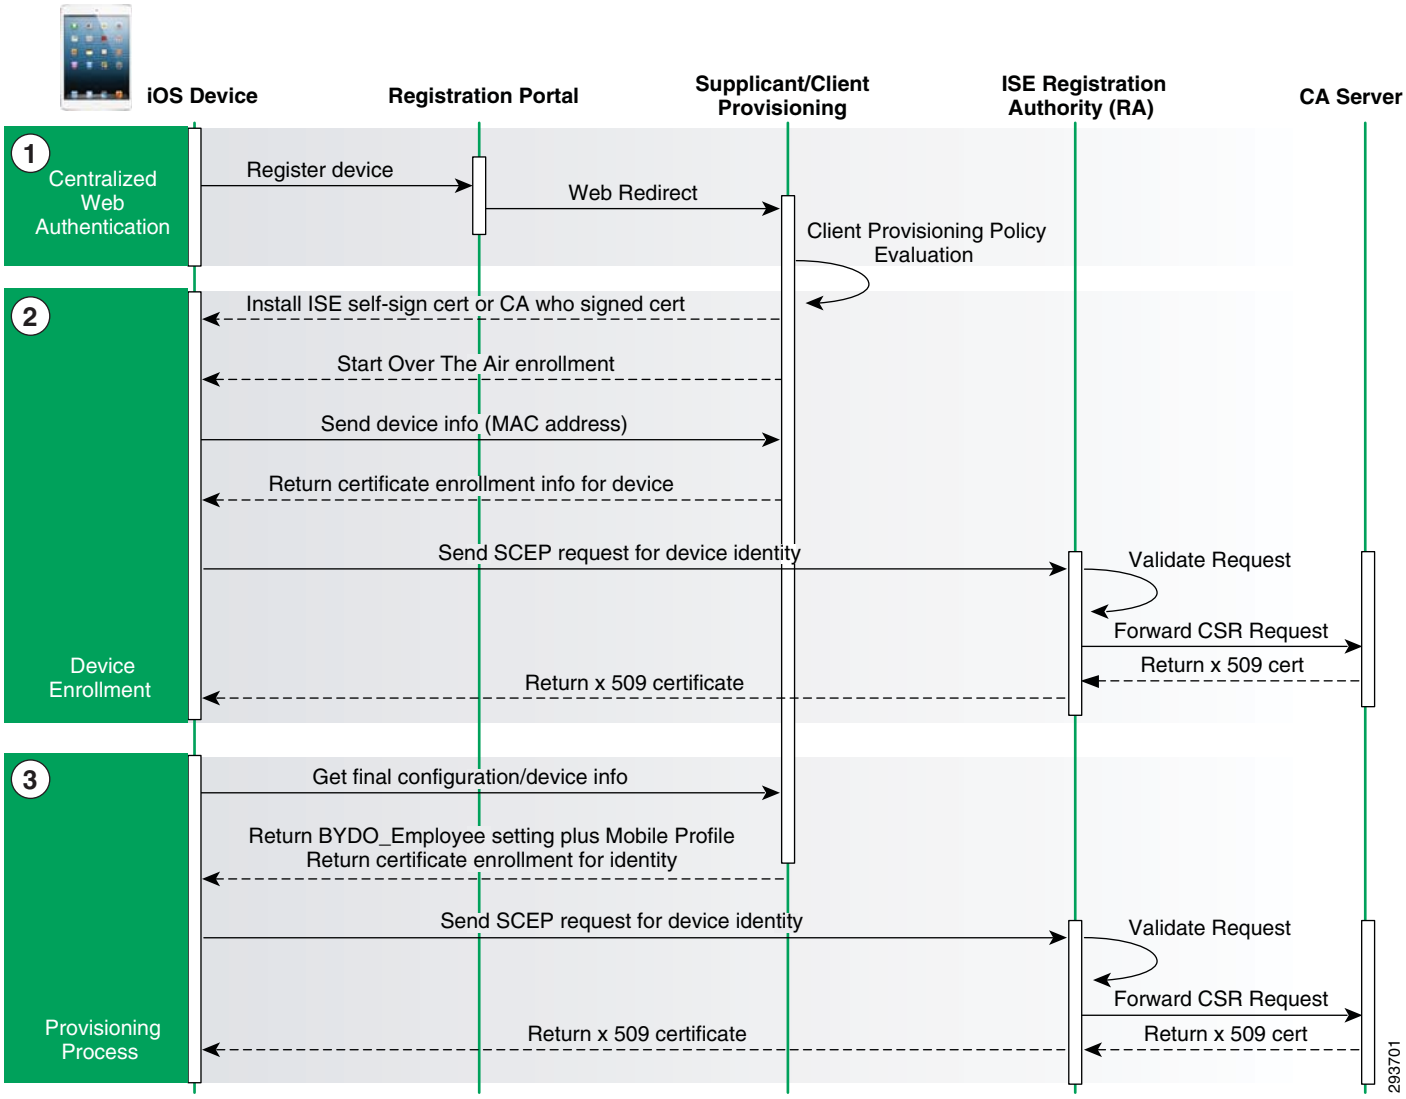
This section explains the interaction between the endpoints and the Guest Registration portal and the steps required to enroll the digital certificate and configuration profile. The way Windows and Mac devices are provisioned is similar. Chapter 8, "Summary of Configuring the Infrastructure" points to the chapters with the steps required to configure the different network components.

Provisioning Apple iOS Devices

The following steps take place while provisioning Apple iOS devices:

1. The device is redirected to the Guest Registration Portal.
2. After successful authentication, the Over-The-Air (OTA) enrollment begins.
3. Device sends unique identifier (MAC address) and other information.
4. Certificate enrollment information is sent to the device.
5. A SCEP request is made to ISE, which returns a certificate.
6. Wireless profile for BYOD_Employee is sent to the device.
7. Once the enrollment is complete, the user manually connects to BYOD_Employee SSID.

Figure 15-8 Provisioning Flow for Apple iOS Devices



For more details on Over-the-Air Enrollment and configuration for iOS devices, review the iPhone OS Enterprise Deployment Guide: http://manuals.info.apple.com/en_US/Enterprise_Deployment_Guide.pdf.

Provisioning Android Devices

The following steps take place while provisioning Android devices:

1. The device is redirected to the Guest Registration portal.
2. After successful authentication, the self-registration portal page redirects the user to the Google Play.
3. The user installs the Supplicant Provisioning Wizard (SPW).
4. The SPW is launched to perform provisioning of the supplicant. The SPW performs the following functions:

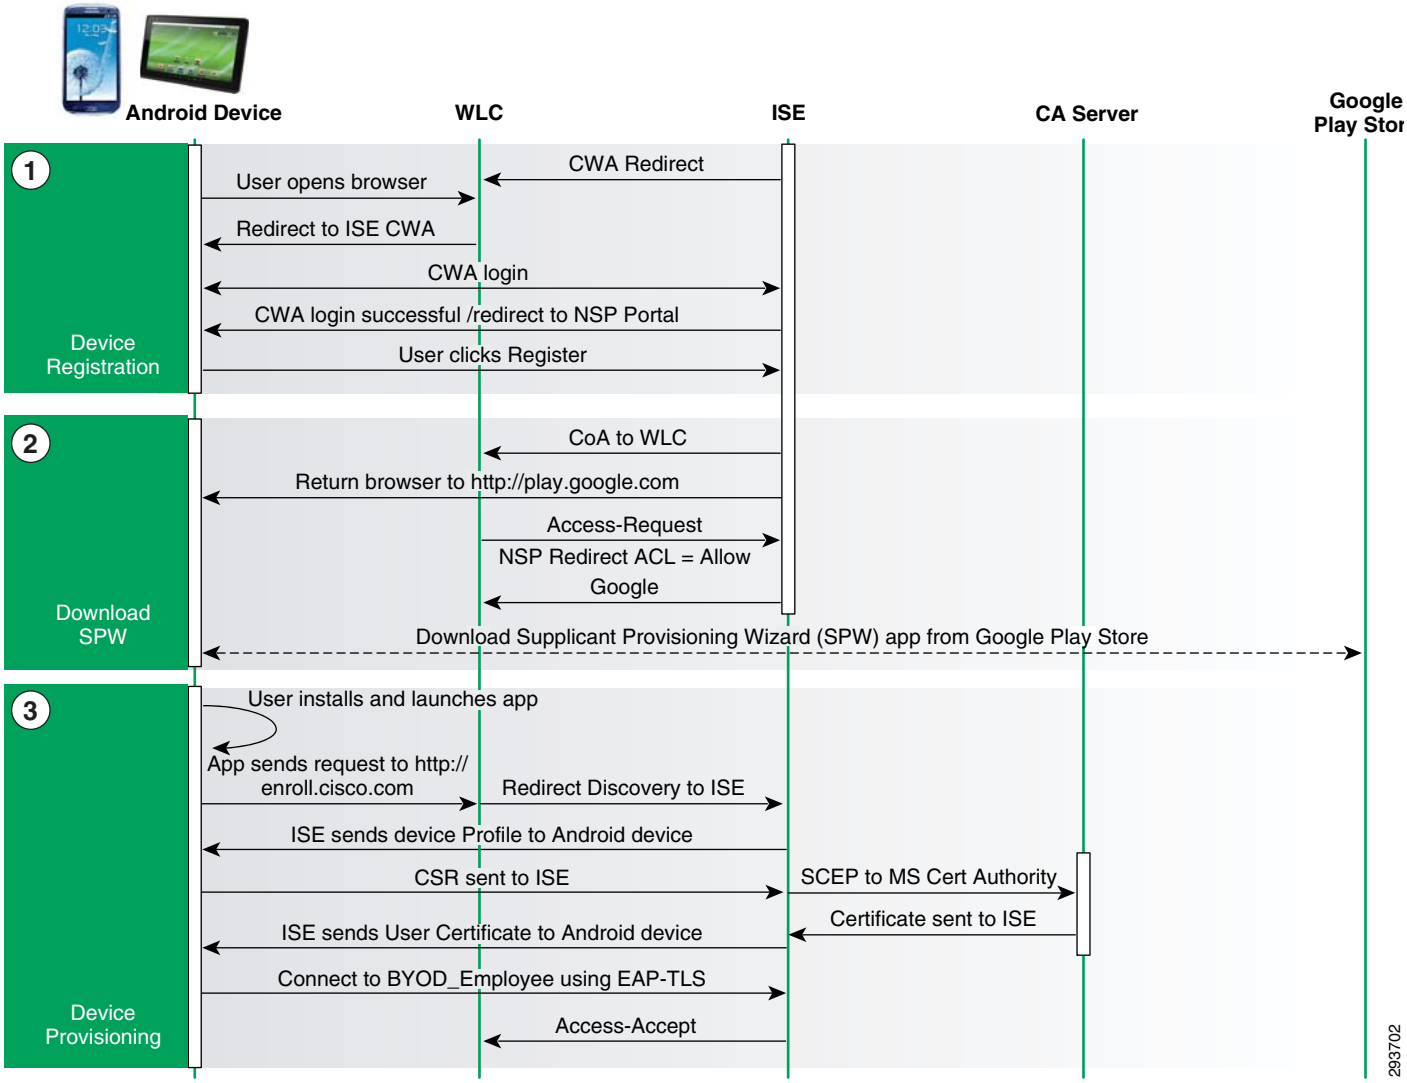
a. Discovers the ISE and downloads the profile from the ISE.
b. Creates a certificate/key pair for EAP TLS.
c. Makes a SCEP proxy request to ISE and gets the certificate.
d. Applies the wireless profile to allow connectivity to BYOD_Employee SSID.

5. The SPW triggers re-authentication and connects to BYOD_Employee SSID automatically.

Note The Android agent must be downloaded from Google Play and is not provisioned by the ISE. Endpoints must be able to reach Google Play on the Internet.

Note ACL_Provisioning_Redirect must redirect all traffic sent to enroll.cisco.com. The Cisco Configuration Assistant for Android devices requires this redirect to discover the ISE server. This is shown as part of step 3 in Figure 15-9. This is not a concern if the guidance presented in the CVD is followed since all Internet traffic except Google Play is redirected back to ISE.

Figure 15-9 Provisioning Flow for Android Devices



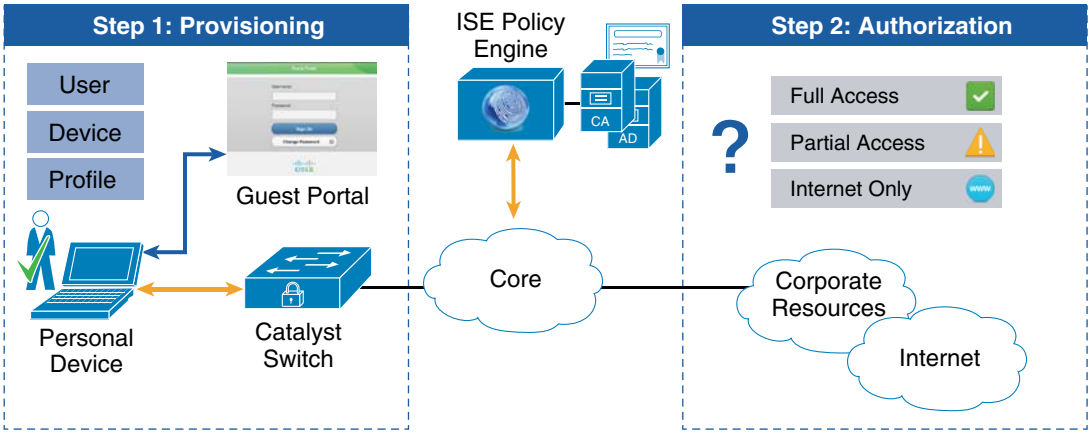
Provisioning Wired Devices

BYOD applies to both wired and wireless devices. Wired devices can be provisioned, registered, authenticated, and authorized in much the same way as wireless devices. The following are some of the advantages of provisioning wired devices:

- Certificate provisioning can be done during the provisioning process, which alleviates the burden on IT to support another model to provision certificates on the devices.
- The native supplicants on the device can be configured with the right protocols during the provisioning process. If this is left to the user, it may often lead to incorrect configurations and additional management overhead for IT.
- Provides easier methods for IT to obtain visibility into who is accessing the network and also methods to remove network access for devices that are lost or stolen.

Figure 15-10 shows a high-level overview of the components used to deploy wired devices. ISE uses several building blocks, such as AD group membership, the EndPoints:BYODRegistration attribute, and Digital Certificates to authenticate and authorize devices. Examples on how to construct these polices are explained in this design guide.

Figure 15-10 Wired Device Deployment



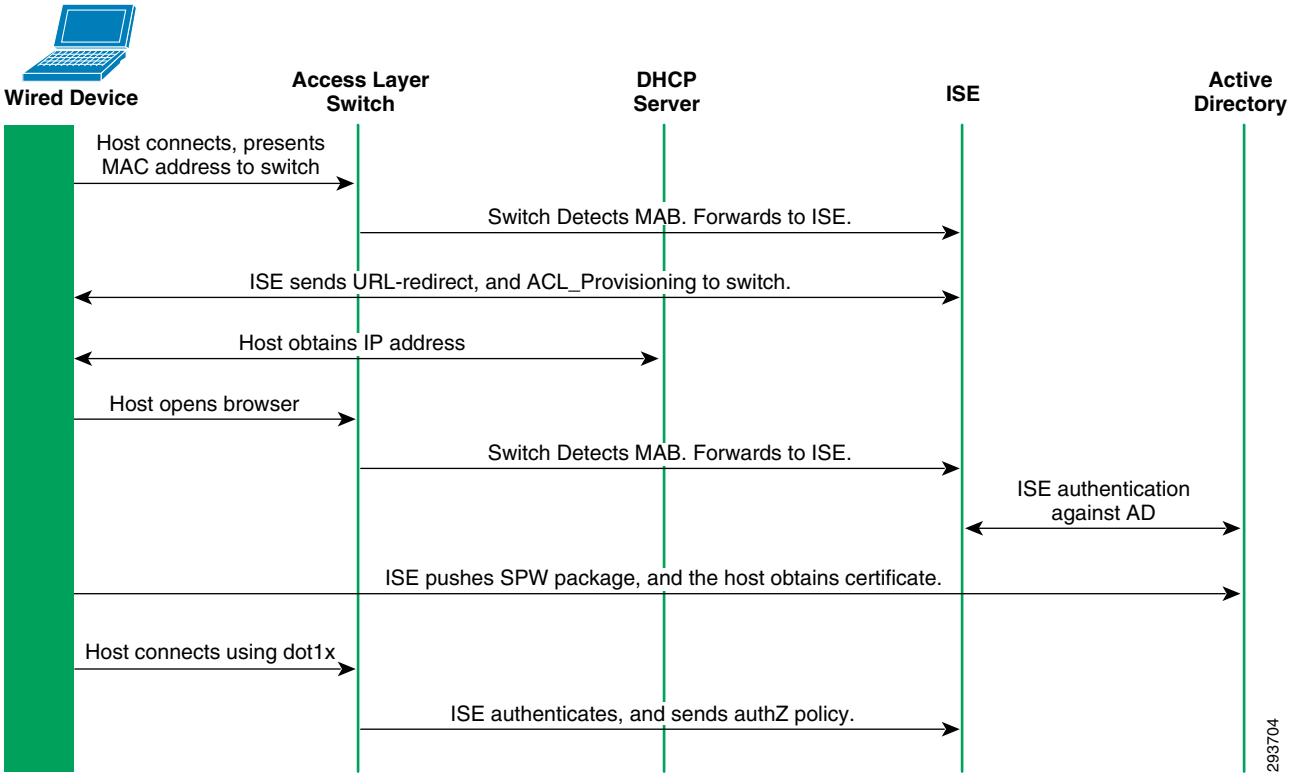
Note Unless otherwise specified, in the figures throughout this chapter "Access Layer Switches" refers to either switches such as the Catalyst 3750X Series and Catalyst 4500 Series or to Catalyst 3850 Series converged access switches.

The following are high-level steps that occur when a wired device connects to the access layer switch:

1. The switch must detect that the wired endpoint is not configured for dot1x and should authenticate using MAB.
2. The ACL_Provisioning_Redirect ACL is used to match web traffic.
3. The URL redirect must point to the ISE Guest Registration portal.
4. The ACL_Provisioning ACL must be downloaded on the port that restricts access in this state.
5. The user opens a browser and attempts to access any resource.
6. The switch redirects the user to an ISE self-registration portal.
7. ISE authenticates the user against AD and pushes the SPW package.
8. The SPW package helps the user register and obtain a digital certificate from ISE.
9. CoA occurs and the user reconnects to the network using the obtained digital certificate.

Figure 15-11 illustrates the flow for wired device provisioning.

Figure 15-11 Wired Device Provisioning Flow



Key and Certificate Storage

Being able to store digital certificates and their associated keys safely is critical for every device. Storage is implemented differently, depending on the operating system or media used. Table 15-1 shows the different platforms tested in this design guide and their certificate stores.

Table 15-1 Platform and Certificate Storage

Device	Certificate Store	How to Access
Microsoft Windows	Machine Certificate Store	Use the Certificates snap-in from the mmc.exe utility
Mac OS	Device Certificate Store	Use the Keychain Access application
Apple iPad	Device Certificate Store	Settings > General > Profile
Android	Credential Storage	Settings > Location & Security

Once provisioned, the certificates have the following attributes, which can be used by ISE to enforce different permissions:

```
Common Name (CN) of the Subject:
    User identity used for authentication

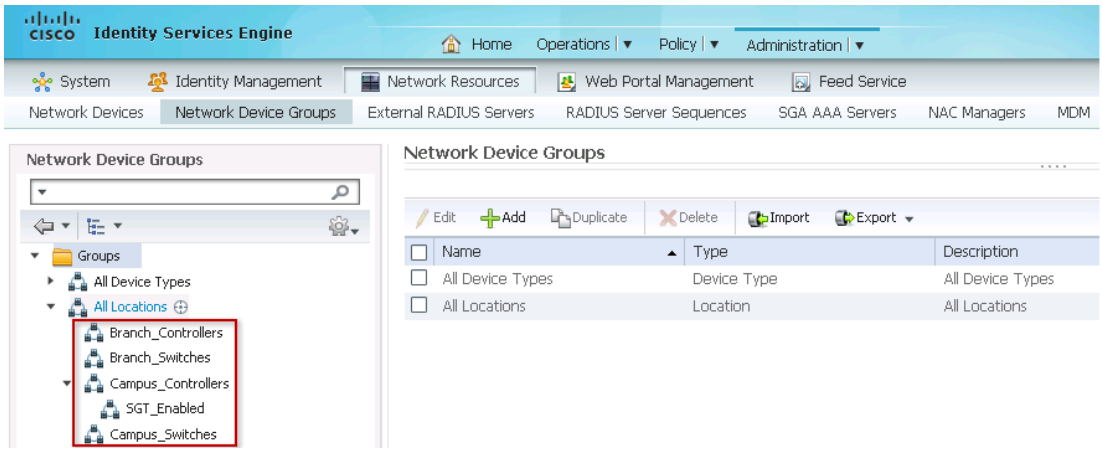
Subject Alternative Name:
    MAC address(es) of the endpoint.
```

Network Device Groups

To differentiate these connections, the ISE relies on Network Device Groups to group WLCs based on their location or device type. This allows a single ISE to enforce policies across different groups of devices. Click **Administration > Network Resources > Network Device Groups** to define locations for branch, campus, and SGT enabled controllers.

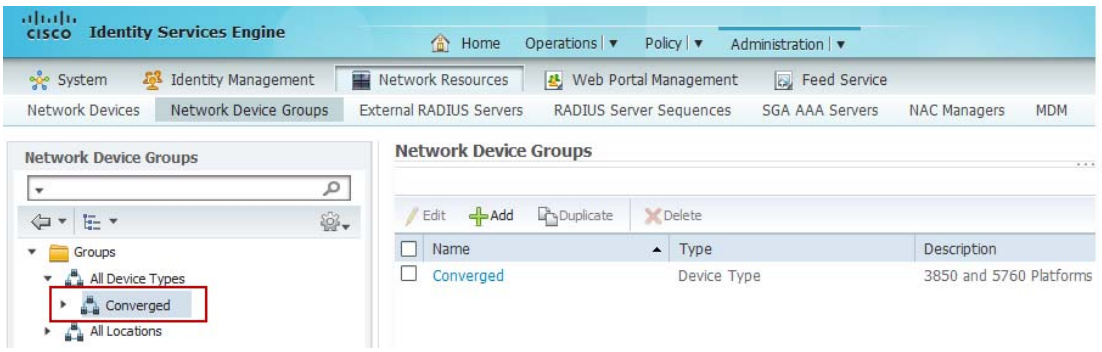
Figure 15-12 shows the different locations used in the authorization policy.

Figure 15-12 ISE Device Groups—Locations



Similarly, Figure 15-13 shows a device type called **Converged** which can be created for Catalyst 3850 Series switches and CT5760 wireless controllers and used in the authorization policy.

Figure 15-13 ISE Device Groups—Device Types

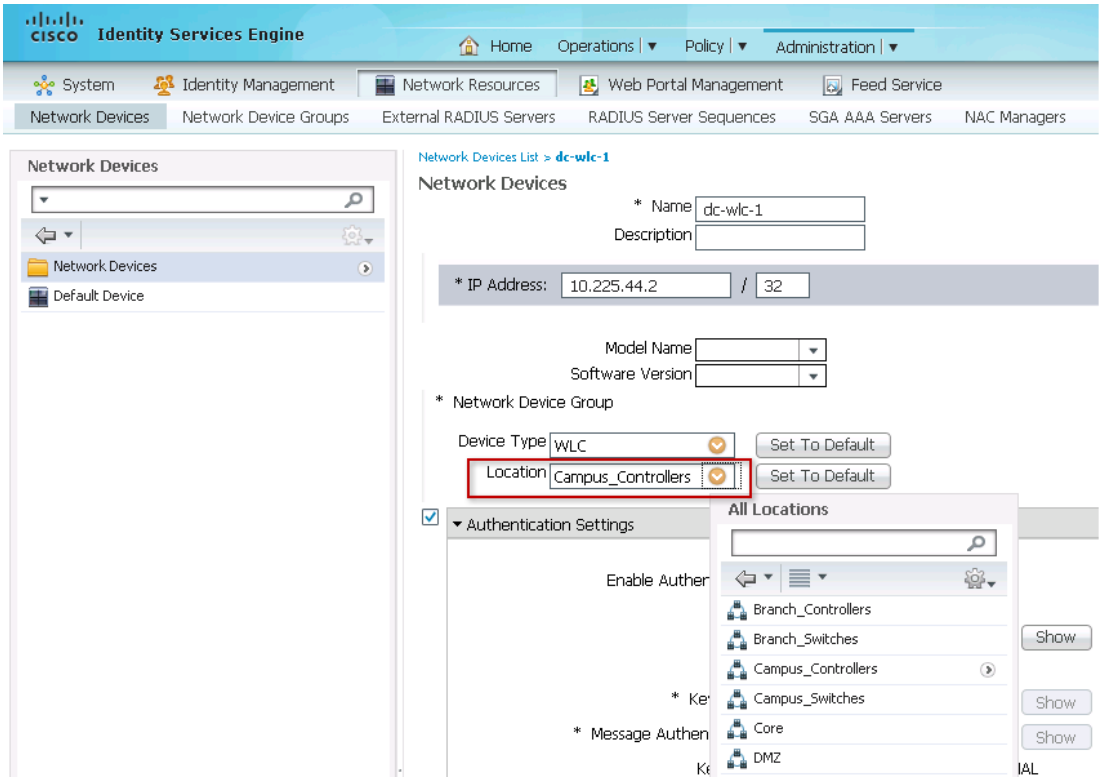


Note One of the reasons a device type is used instead of a location for Converged Access designs is that the same authorization policy rules are used for Converged Access branch and campus designs within this design guide. Hence location—campus versus branch—is not particularly relevant from a Cisco ISE perspective to Converged Access designs presented in this design guide. Note that if location is relevant, the customer can always modify the design presented here and choose to deploy separate authorization policy rules for branch and campus Converged Access designs as well.

Each Wireless LAN Controller, previously defined as a Network Device, needs to be added to the proper device group by clicking **Administration > Network Resources > Network Devices** and specifying the proper location or device type from the pull-down menu.

Figure 15-14 shows how the dc-wlc-1 Wireless LAN Controller belongs to the Campus_Controllers Network Device Group.

Figure 15-14 ISE Network Devices—Campus Controller



Policy Enforcement for Security Group Access

Enforcing policies for SGA in BYOD architecture consist of two main components:

- Defining tags for the endpoints and the destination servers.
- Defining and implementing the Security Group ACL or Security Group Firewall (SG-FW) policy.

In this design, the Security Group ACL is implemented either on the Catalyst and Nexus Data Center switching infrastructure or on the ASA Firewall as an SG-FW policy. These two choices are discussed in Chapter 23, "BYOD Policy Enforcement Using Security Group Access" as two different deployment scenarios.

Security Group Access Tags

The basic idea of Security Group Access is to associate a device or server's IP address with a Security Group Tag and then create role-based policies that either permit or deny traffic flows based on these source and destination SGTs. For example: are devices tagged with an SGT 10 allowed to communicate with a server tagged with an SGT 40? This flow is either allowed or blocked at the enforcement point.

Using this tagging concept, unique tags have been defined for each device accessing the network as a client. As explained in Chapter 4, "BYOD Use Cases," unique permissions are granted to personal and corporate devices and hence unique tags have been defined for each use case.

Table 15-2 illustrates how tags are assigned to different devices.

Table 15-2 Source Tags

Device Type	Tag
Corporate device with Full Access	SGT 10
Personal device with Full Access	SGT 11
Personal device with Partial Access	SGT 12

Similarly, the destination servers also need to be associated with a particular tag. Table 15-3 illustrates how based on their role, servers are assigned with a different tag.

Table 15-3 Destination Tags

Destination Server	Tag
Open Access	SGT 40
Corporate Server	SGT 50

Security Group ACL

After defining the source tags and destination tags, the next logical step is to define the egress policy matrix that defines the enforcement policy between a source and destination tag. Table 15-4 illustrates the egress policy matrix.

Table 15-4 Egress Policy Matrix

	Device Type	SGT 40	SGT 50
SGT 10	Corporate device with Full Access	Yes	Yes
SGT 11	Personal device with Full Access	Yes	Yes
SGT 12	Personal device with Partial Access	Yes	No

As shown in Table 15-4, corporate or personal devices granted full access will be allowed to reach servers that have been tagged with SGT 40 or SGT 50. Similarly, a personal device with partial access is only allowed to connect with a server tagged with SGT 40.

The implementation of the SGACL depends upon the type of the deployment scenario. For the deployment scenario where Nexus is the enforcement point the SGACL is defined in ISE and pushed to the Nexus switch whereas when the deployment scenario is using ASA as the enforcement point then an SGT-based access rule is configured manually on the ASA firewall.

Authorization Polices for SGT

Based on policy matches and authorization profiles, the ISE assigns an SGT tag to endpoints. Centralized Campus with SGT, found later in this chapter, provides information as to the criteria used to determine when an ACL versus an SGT should be returned upon successful authorization based on the network device type of the wireless controller defined in ISE.

For the destination tags assigned to servers, the configuration is done manually at the data center switch, and the actual enforcement happens based on the deployment scenario. To obtain more information on how to configure tags for servers and on the ASA, refer to Chapter 23, "BYOD Policy Enforcement Using Security Group Access."

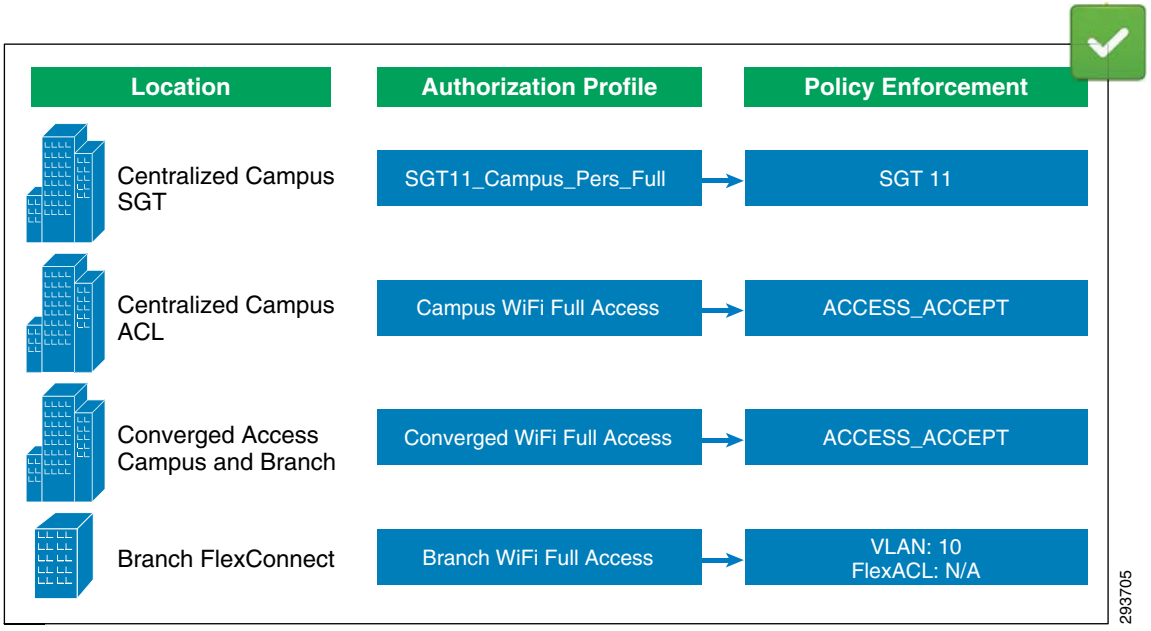
Personal Wireless Devices—Full Access

To provide full access to personal devices, the Cisco ISE verifies the following:

- The employee has completed the on-boarding process through the Guest Registration portal.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate.
- The connection originated using EAP-TLS authentication.
- The user is a member of the BYOD_Full_Access Active Directory group.

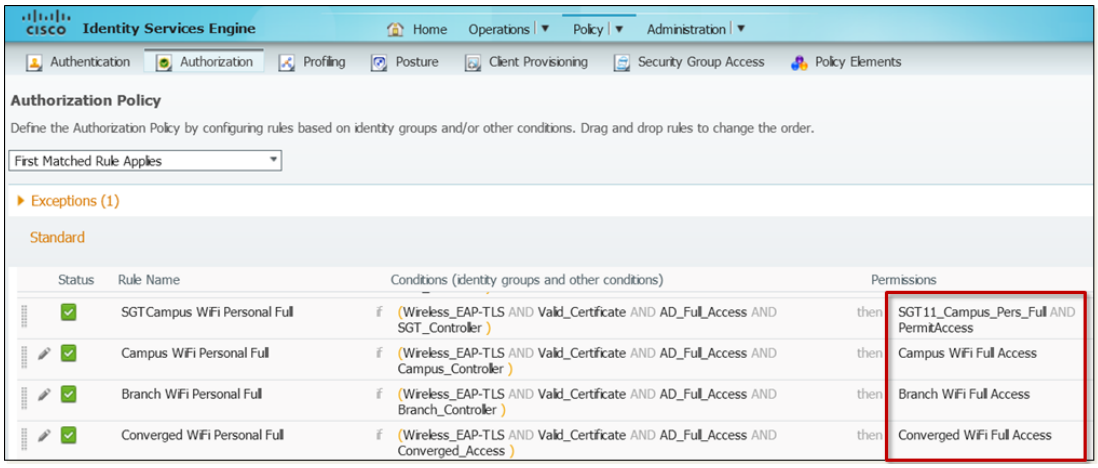
Since the wireless designs presented in this design guide rely on different WLCs for FlexConnect branches, centralized controller campuses, and Converged Access campuses and branches; unique authorization rules are created for connections originating from each design. At a high level, Figure 15-15 shows how different authorization profiles are selected for connections originating from different locations with different wireless designs. Each authorization profile in turn enforces a unique permission.

Figure 15-15 Full Access Enforcement



To configure the authorization rules in ISE, click **Policy > Authorization**. Figure 15-16 highlights the authorization policy to grant full access to personal devices.

Figure 15-16 Authorization Policies for Full Access



Looking at the rules in more detail, ISE evaluates the following conditions:

- Wireless_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name. (defined as a simple condition).
- The user belongs to a specific Active Directory group (defined as a simple condition).
- The RADIUS authentication originated from a wireless controller which was a member of one of the following device groups—Campus_Controller, SGT_Controller, Branch_Controller, or Converged_Access (defined as a simple condition).

Note A wireless controller which is a member of the Converged_Access device group could be either a standalone device, such as a Cisco CT5760 wireless controller, or a switch with integrated wireless controller functionality, such as the Catalyst 3850.

Simple and Compound Conditions

To improve the readability of the authorization policy, simple and compound authorization conditions were defined to group different conditions. These conditions may be reused and modified without changing every authorization rule.

Table 15-5 shows the conditions used in the authorization rules:

Table 15-5 Simple and Compound Conditions

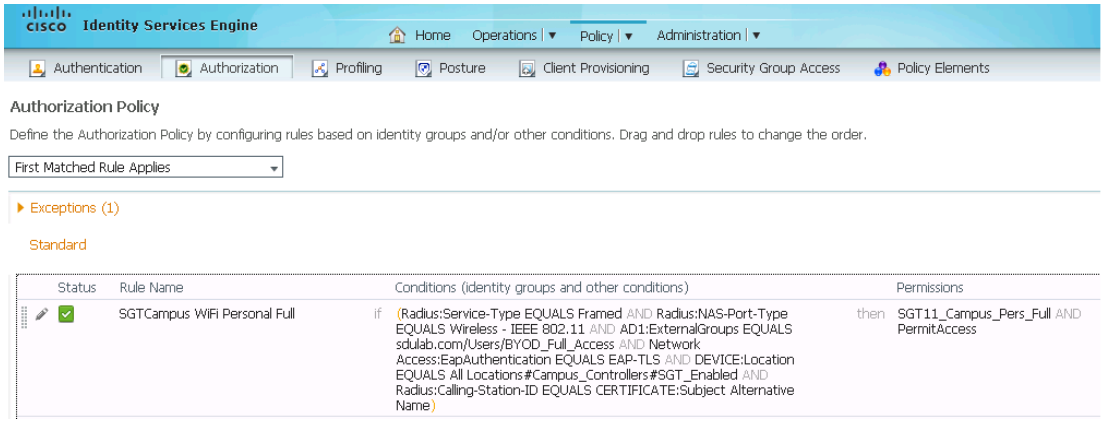
Wireless EAP-TLS (Compound)	
Wireless_EAP-TLS (See Figure 15-18)	Radius:Service-Type Equals Framed AND Radius:NAS-Port-Type Equals Wireless - IEEE 802.11 AND Network Access:EapAuthentication Equals EAP-TLS

Table 15-5 Simple and Compound Conditions

Check for Valid Certificate (Simple)	
Valid_Certificate (See Figure 15-19)	Radius:Calling-Station-ID EQUALS CERTIFICATE:Subject Alternative Name
Active Directory Group (Simple)	
AD_ Full_Access	AD1:ExternalGroups EQUALS sdulab.com/Users/BYOD_Full_Access
AD_ Partial_Access	AD1:ExternalGroups EQUALS sdulab.com/Users/BYOD_Partial_Access
AD_Domain_users	AD1:ExternalGroups EQUALS sdulab.com/Users/Domain User
WLC Location or Device Type (Simple)	
SGT_Controller	DEVICE:Location EQUALS All Locations#Campus_Controllers#SGT_Enabled
Campus_Controller	DEVICE:Location EQUALS All Locations#Campus_Controllers
Branch_Controller	DEVICE:Location EQUALS All Locations#Branch_Controllers
Converged_Access	DEVICE:Device Type EQUALS All Device Types#Converged

To illustrate the value of using simple/compound conditions, Figure 15-17 shows how much longer and harder to read a rule can be when simple/compound conditions are not implemented.

Figure 15-17 Authorization Rule without Conditions



To define a new compound condition, click **Policy > Conditions > Authorization > Compound Conditions**. Figure 15-18 shows how the Wireless_EAP-TLS condition combines several conditions into one.

Figure 15-18 Wireless_EAP-TLS Condition

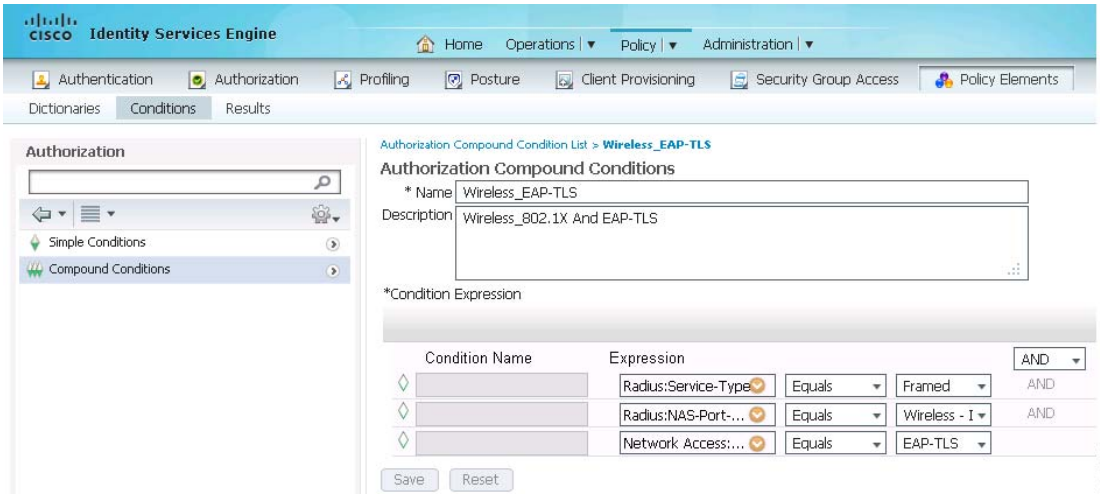


Figure 15-19 shows the Valid_Certificate simple condition.

Figure 15-19 Valid_Certificate Condition

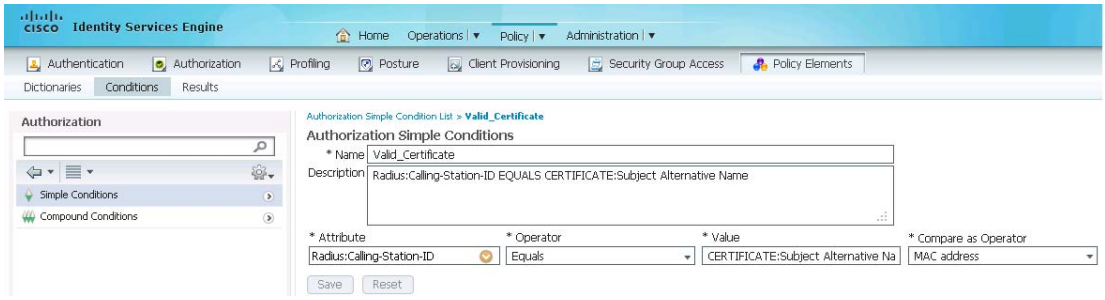
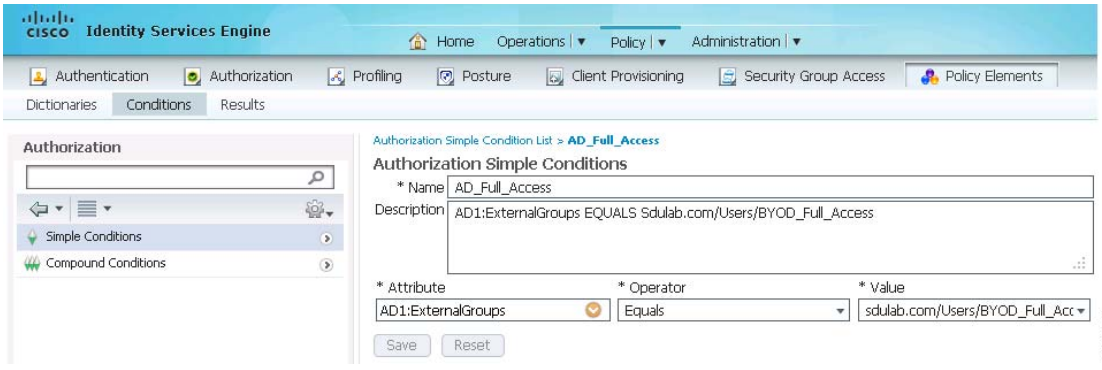


Figure 15-20 shows the AD_Full_Access simple condition.

Figure 15-20 AD_Full_Access Condition



The remaining conditions mentioned in Table 5 are defined in a similar way.

Permissions

Permission is a result of an authorization policy match, and the permission can be of different types such as an authorization profile, or a standard result. Table 15-6 explains the permissions used for full access.

Table 15-6 Permissions for Full Access

Permission name	Permission type	Purpose
SGT11_Campus_Pers_Full	Standard result	To Assign a SGT for 802.1X wireless devices connecting from an SGT-enabled controller allowing full access.
Campus WiFi Full Access	Authorization profile	Provides Full Access for 802.1X wireless devices connecting from a centralized campus controller.
Branch WiFi Full Access	Authorization profile	To push a VLAN for 802.1X wireless devices connecting from a FlexConnect branch controller
Converged WiFi Full Access	Authorization profile	Provides Full Access for 802.1X wireless devices connecting from a Converged Access controller.

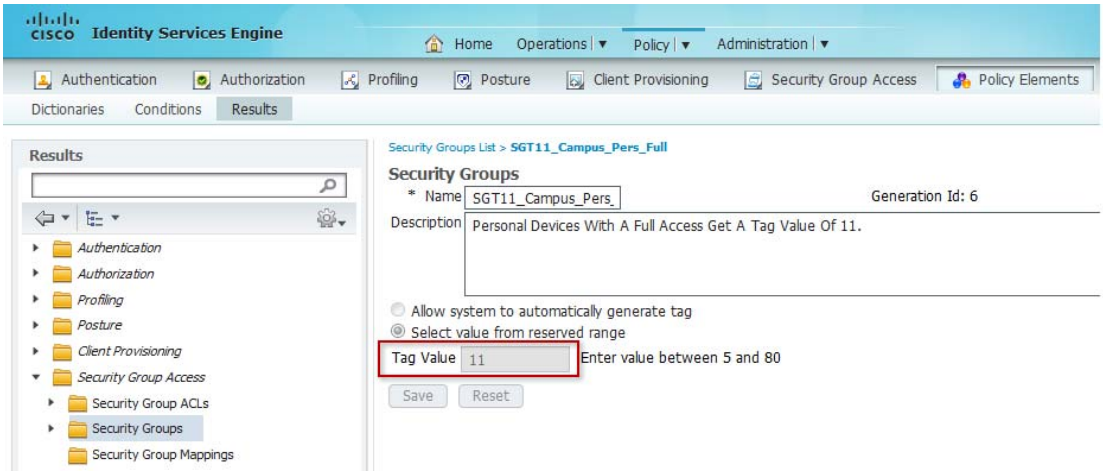
Note A Converged Access infrastructure refers to a branch or campus deployment with Catalyst 3850 Series switches and/or CT5760 wireless controllers within this design guide.

Centralized Campus with SGT

As explained in Policy Enforcement for Security Group Access, the personal device with permissions for full access will be assigned an SGT value of 11. Once the personal device obtains the tag value of 11, it can then connect with all the servers in the data center.

Figure 15-21 shows how the SGT11_Campus_Pers_Full authorization profile uses SGT 11 for personal devices granted full access.

Figure 15-21 SGT11_Campus_Pers_Full

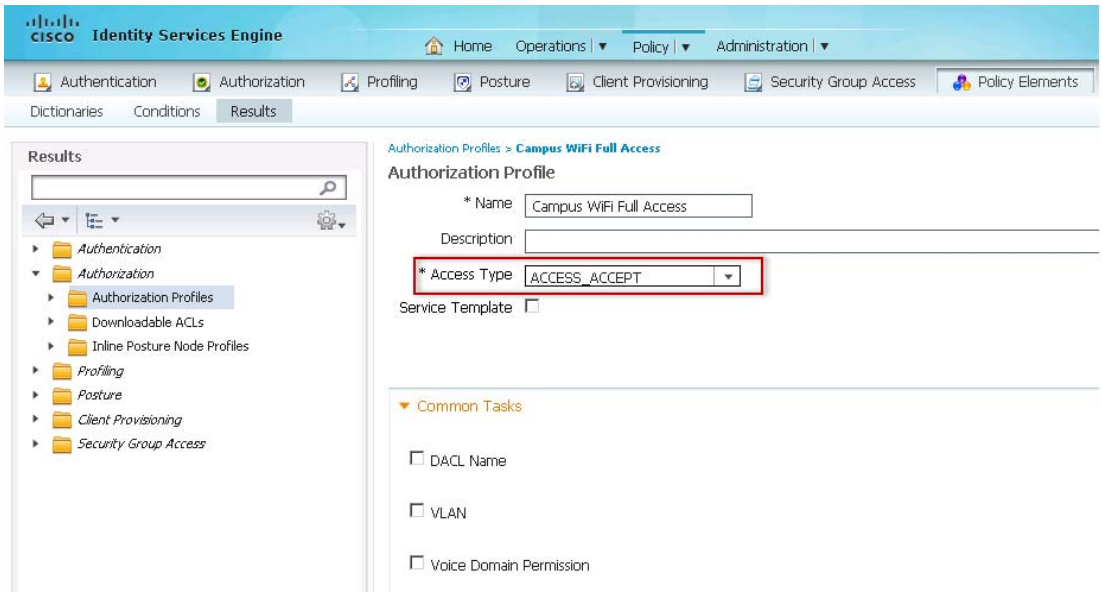


See Policy Enforcement for Security Group Access for details regarding the SGT egress policy matrix which defines the permissions for SGT11 allowing full access for a campus which implements a centralized controller (Local Mode) design with SGTs.

Centralized Campus with ACLs

Figure 15-22 shows how the Campus WiFi Full Access authorization profile is using the ACCESS_ACCEPT Access Type to allow full access.

Figure 15-22 Campus WiFi Full Access

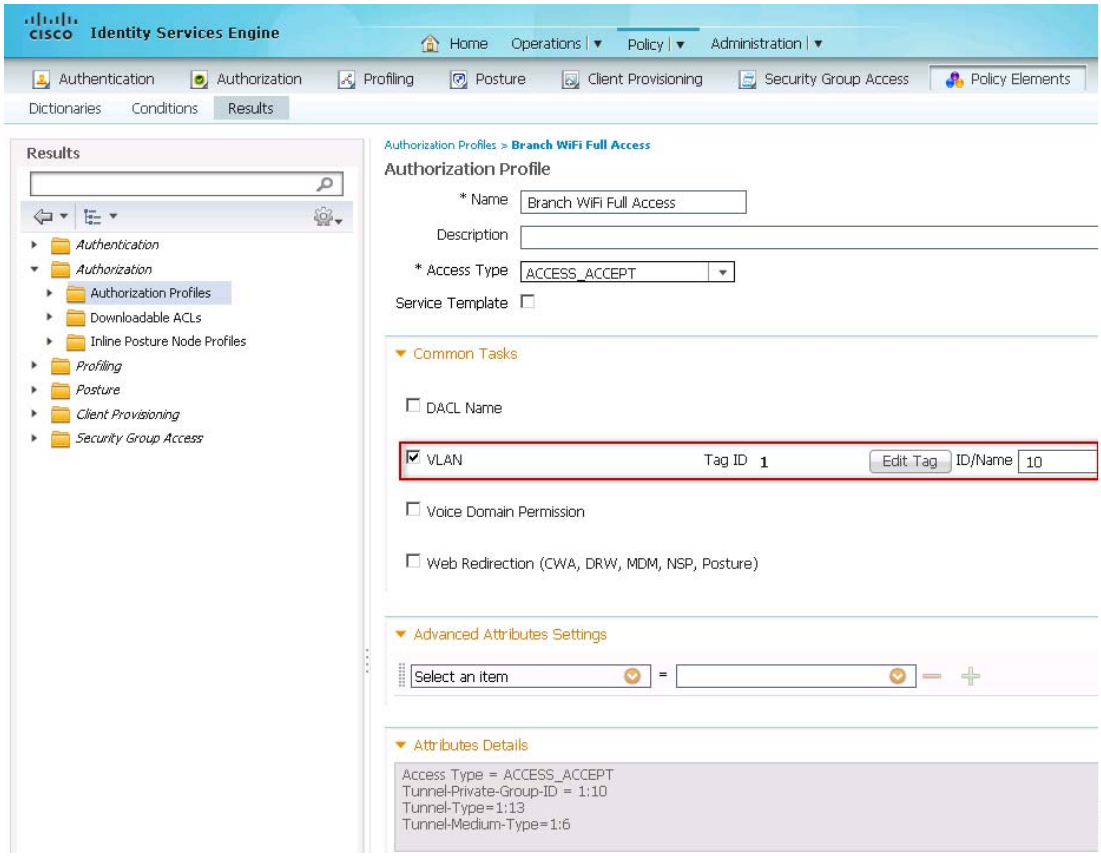


Since full access is allowed with this authorization profile, no Named ACL for access control needs to be specified by ISE.

Branch with FlexConnect

Endpoints connecting from a branch location which implements FlexConnect dynamically get assigned to VLAN 10, which has been configured to provide full access. Figure 15-23 shows the Branch WiFi Full Access authorization profile.

Figure 15-23 Branch WiFi Full Access

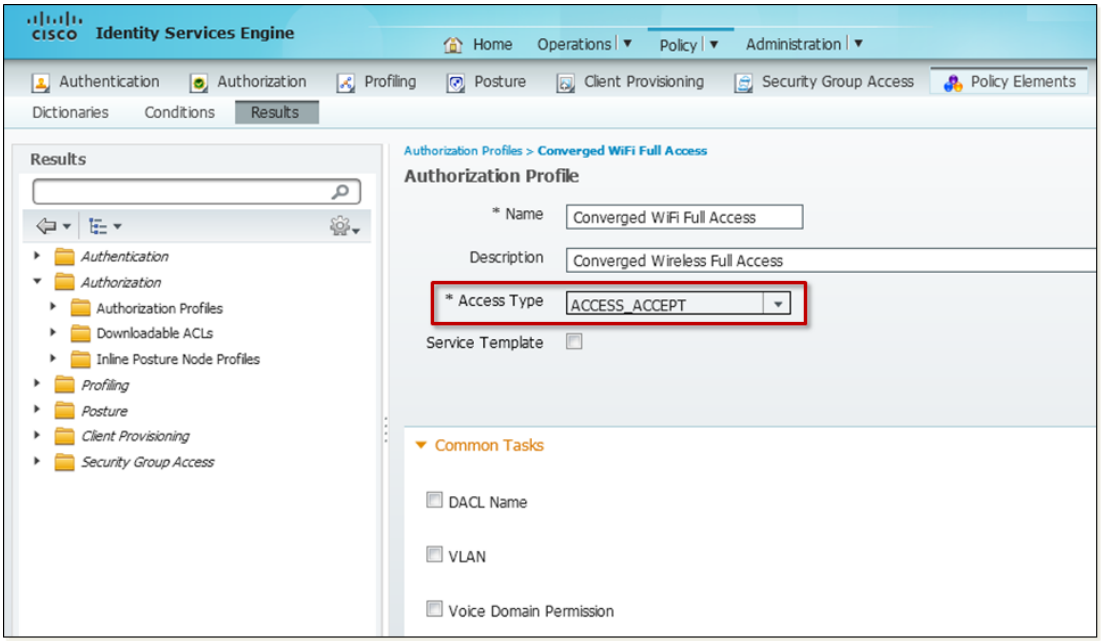


Since full access is allowed with this authorization profile, no FlexConnect ACL for access control needs to be associated with VLAN 10 on the access point.

Converged Access Branch or Campus

Figure 15-24 shows how the Converged WiFi Full Access authorization profile is using the ACCESS_ACCEPT Access Type to allow full access.

Figure 15-24 Converged WiFi Full Access



Again, since full access is allowed with this authorization profile, no Named ACL for access control needs to be specified by ISE.

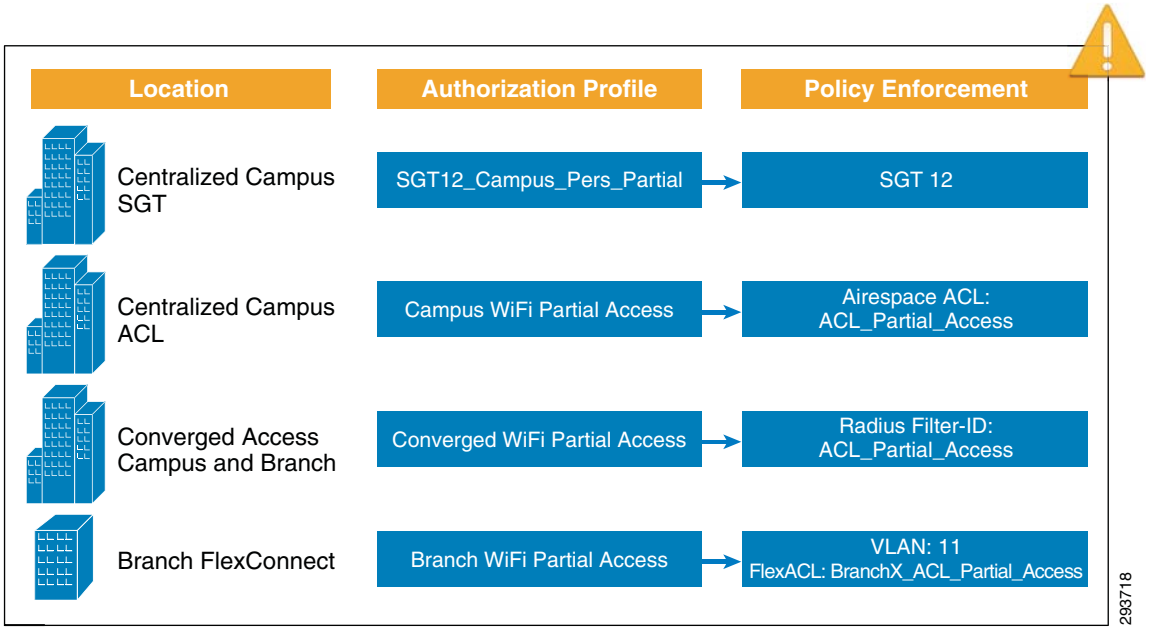
Personal Wireless Devices—Partial Access

To provide partial access to personal devices, the Cisco ISE verifies the following:

- The employee has completed the on-boarding process through the Guest Registration portal.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate, in this case, the MAC address of the endpoint.
- The connection originated using EAP-TLS authentication.
- The user is member of the BYOD_Partial_Access Active Directory group.

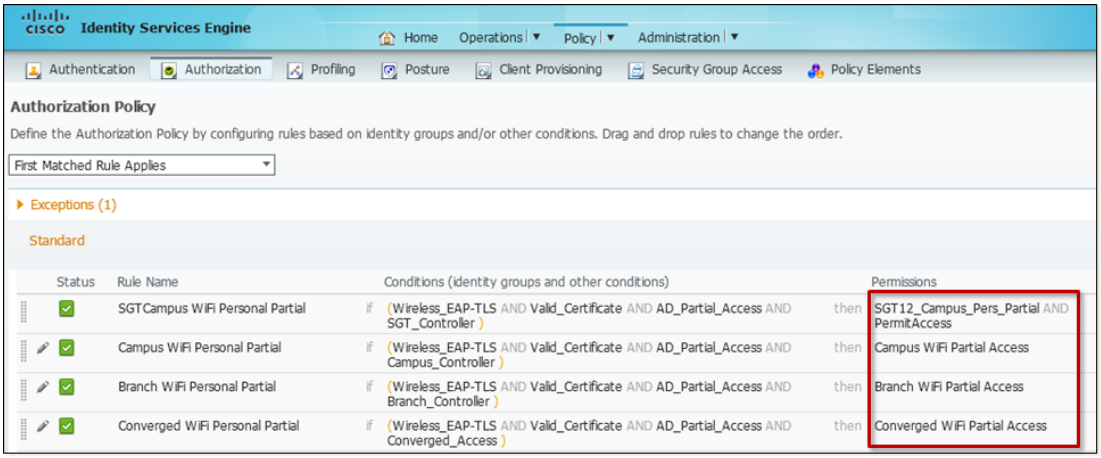
At a high level, Figure 15-25 shows how different authorization profiles are selected for devices accessing the network from different locations with different wireless designs. Each authorization profile in turn enforces a unique permission using either VLANs, SGTs, dynamic ACLs, FlexConnect ACLs, etc.

Figure 15-25 Partial Access Enforcement



To configure the authorization rules in ISE, click **Policy > Authorization**. Figure 15-26 highlights the authorization policy to grant partial access to personal devices.

Figure 15-26 Authorization Policies for Partial Access



Looking at the rules in more detail, ISE evaluates the following conditions:

- Wireless_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name. (defined as a simple condition).
- The user belongs to a specific Active Directory group (defined as a simple condition).
- The RADIUS authentication originated from a wireless controller which was a member of one of the following device groups—Campus_Controller, SGT_Controller, Branch_Controller, or Converged_Access (defined as a simple condition).

Simple and Compound Conditions explains the different conditions used in the rules.

Permissions

Permission is a result of an authorization policy match and the permission can be of different types such as an authorization profile or a standard result. Table 15-7 explains the permissions used for partial access.

Table 15-7 Permissions Used for Wireless Partial Access

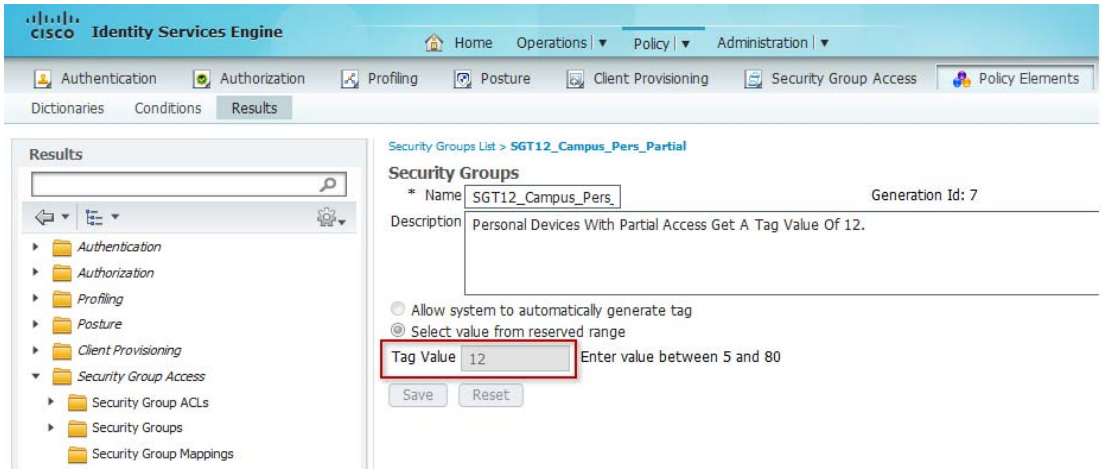
Permission name	Permission type	Purpose
SGT12_Campus_Pers_partial	Standard result	To Assign a SGT for 802.1X wireless devices connecting from an SGT-enabled controller granting partial access to some servers in the data center.
Campus WiFi Partial Access	Authorization profile	To push a named ACL for 802.1X wireless devices connecting from a centralized campus controller.
Branch WiFi Partial Access	Authorization profile	To push a VLAN for 802.1X wireless devices connecting from a FlexConnect branch controller
Converged WiFi Partial Access	Authorization profile	To push a named ACL for 802.1X wireless devices connecting from either a campus or branch location with Converged Access

Centralized Campus with SGT

As explained in Policy Enforcement for Security Group Access, the personal device with permissions for partial access will be assigned an SGT value of 12. Once the personal device obtains the tag value of 12 it can then connect with only those servers with an SGT of 50 in the data center.

Figure 15-27 shows how the SGT12_Campus_Pers_Partial authorization profile is configured to assign SGT 12 to personal devices thus allowing partial access.

Figure 15-27 SGT12_Campus_Pers_Partial

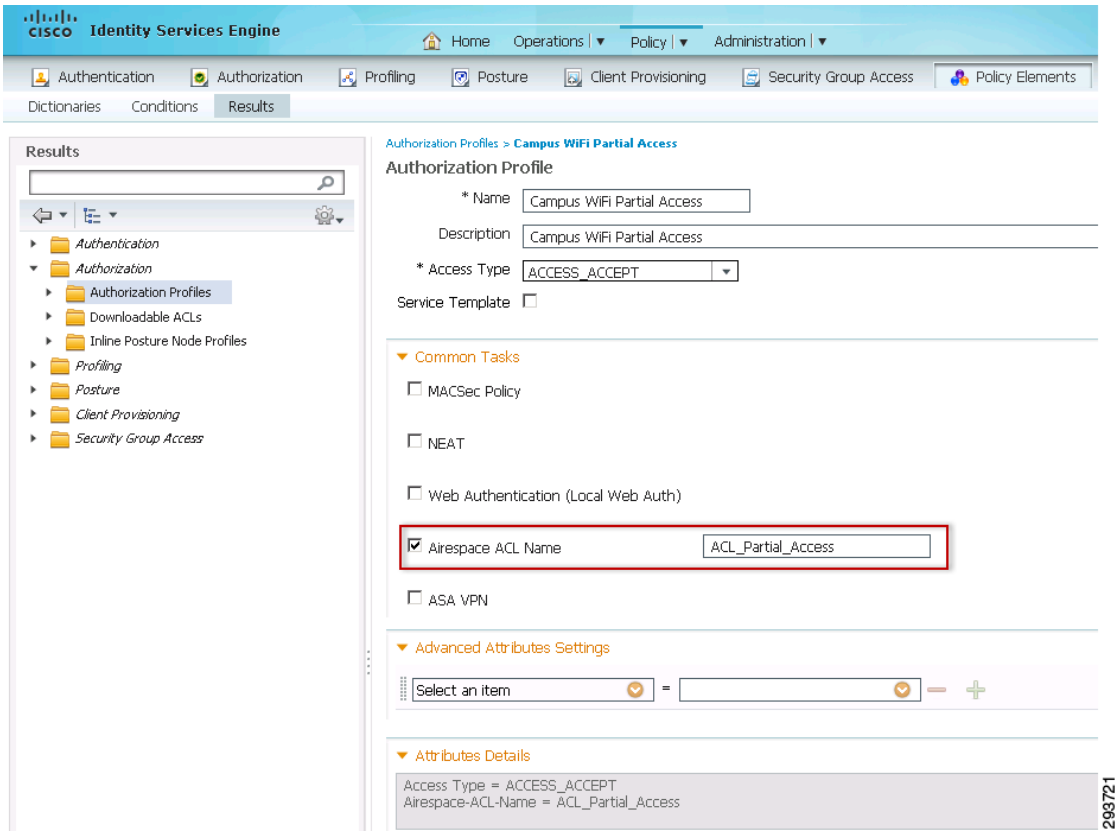


See Policy Enforcement for Security Group Access for details regarding the SGT egress policy matrix which defines the permissions for SGT12 allowing full access for a campus which implements a centralized controller (Local Mode) design with SGTs.

Centralized Campus with ACLs

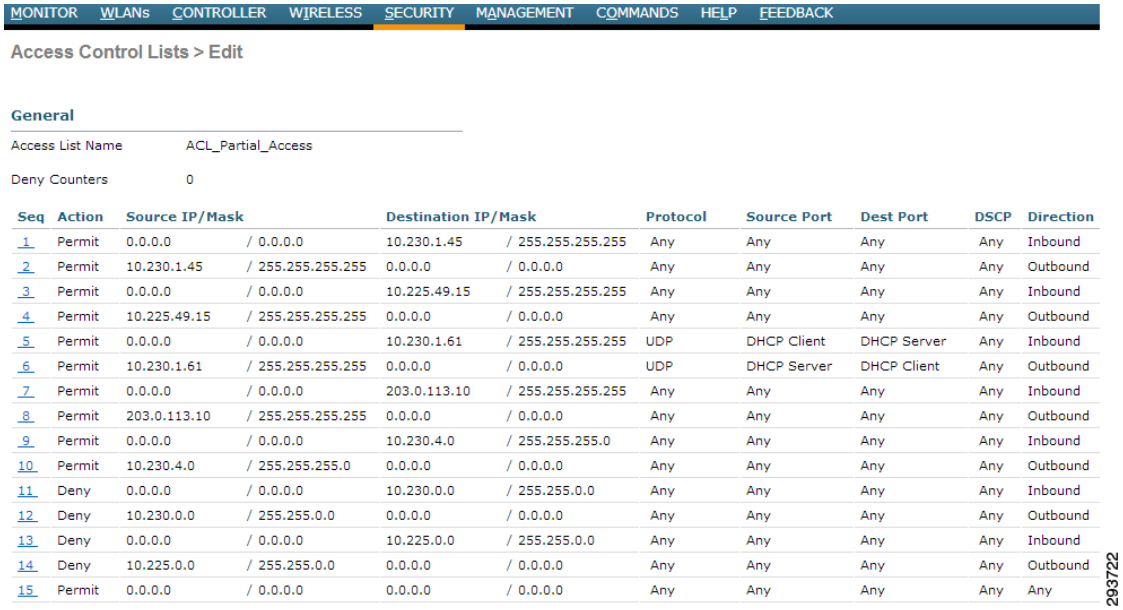
For devices connecting from a campus location which implements a centralized (Local Mode) wireless design, the Campus WiFi Partial Access authorization profile relies on the ACL_Partial_Access access list, enforced by the WLC. Figure 15-28 shows the authorization profile.

Figure 15-28 Campus WiFi Partial Access



Cisco Wireless LAN Controllers support named ACLs, meaning that the ACL must be previously configured on the controller rather than being downloaded from ISE. Using the RADIUS Airespace-ACL Name attribute-value pair, ISE instructs the WLC to apply the ACL_Partial_Access ACL. Figure 15-29 shows the contents of this ACL, as defined in the CT5508 WLC campus controller.

Figure 15-29 ACL_Partial_Access on Wireless Controller



MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Access Control Lists > Edit

General

Access List Name ACL_Partial_Access

Deny Counters 0

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port	DSCP	Direction
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any	Any	Inbound
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any	Any	Inbound
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server	Any	Inbound
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client	Any	Outbound
7	Permit	0.0.0.0	/ 0.0.0.0	203.0.113.10	/ 255.255.255.255	Any	Any	Any	Any	Inbound
8	Permit	203.0.113.10	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
9	Permit	0.0.0.0	/ 0.0.0.0	10.230.4.0	/ 255.255.255.0	Any	Any	Any	Any	Inbound
10	Permit	10.230.4.0	/ 255.255.255.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
11	Deny	0.0.0.0	/ 0.0.0.0	10.230.0.0	/ 255.255.0.0	Any	Any	Any	Any	Inbound
12	Deny	10.230.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
13	Deny	0.0.0.0	/ 0.0.0.0	10.225.0.0	/ 255.255.0.0	Any	Any	Any	Any	Inbound
14	Deny	10.225.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
15	Permit	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Any

293722

The access-list shown in Figure 15-29 specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Allow IP access to and from the MDM Server (203.0.113.10).
- Allow IP access to and from specific subnet (10.230.4.0 /24).
- Deny IP access to and from data center subnets (10.230.0.0 /16).
- Deny IP access to and from campus subnets (10.225.0.0 /16).
- Allow access to and from all other subnets (Internet access).

Note Note that Access Control Entries (ACEs) for the MDM server are included within the ACLs in this chapter. These are not used for the Enhanced Access use case discussed within this chapter. The Advanced Access use case, discussed in Chapter 17, "BYOD Advanced Use Case—Mobile Device Manager Integration," makes use of the same authorization policy rules and authorization profiles.

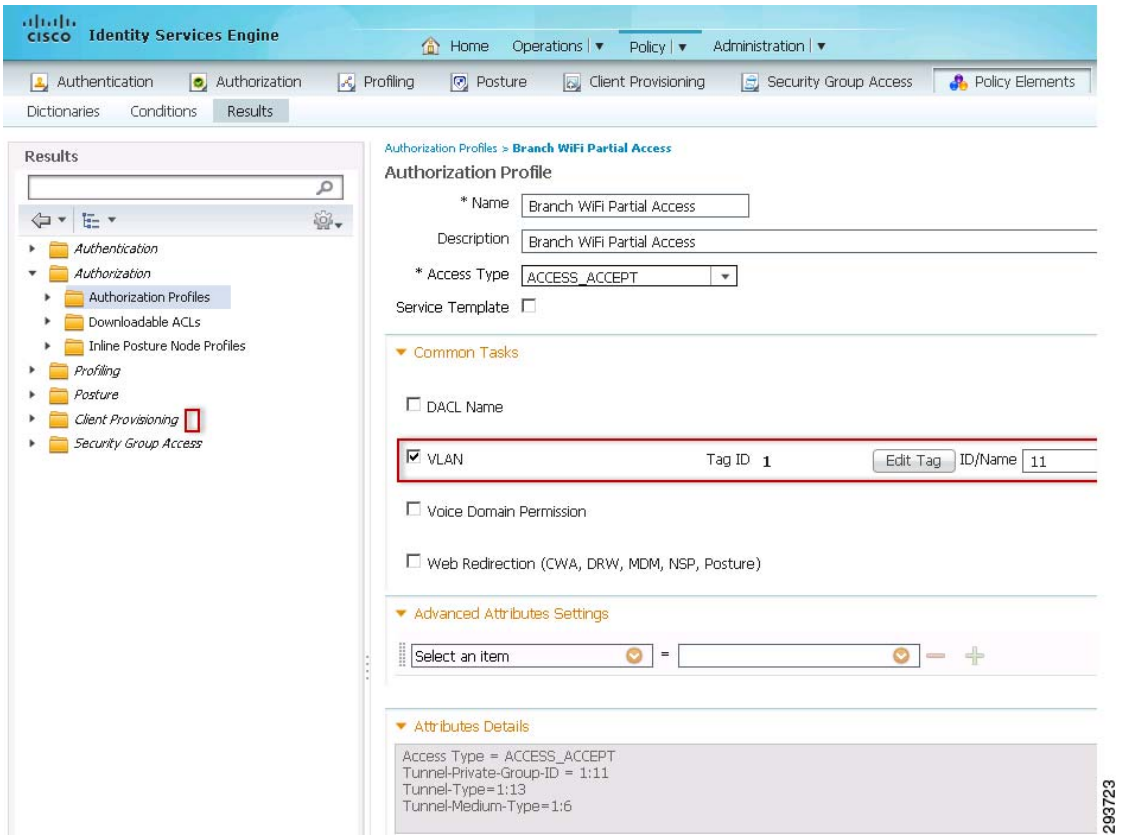
The access list shown in Figure 15-29 is a generic example used to enforce a hypothetical use case and is not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

Note The CT5508 WLC supports a maximum of 64 ACLs with a maximum of 64 lines per ACL.

Branches with FlexConnect

For devices connecting from a branch location which implements a FlexConnect wireless design, the Branch WiFi Partial Access authorization profile dynamically assigns the device to VLAN11, which is dedicated for devices obtaining Partial Access. Figure 15-30 shows this authorization profile.

Figure 15-30 Branch WiFi Partial Access



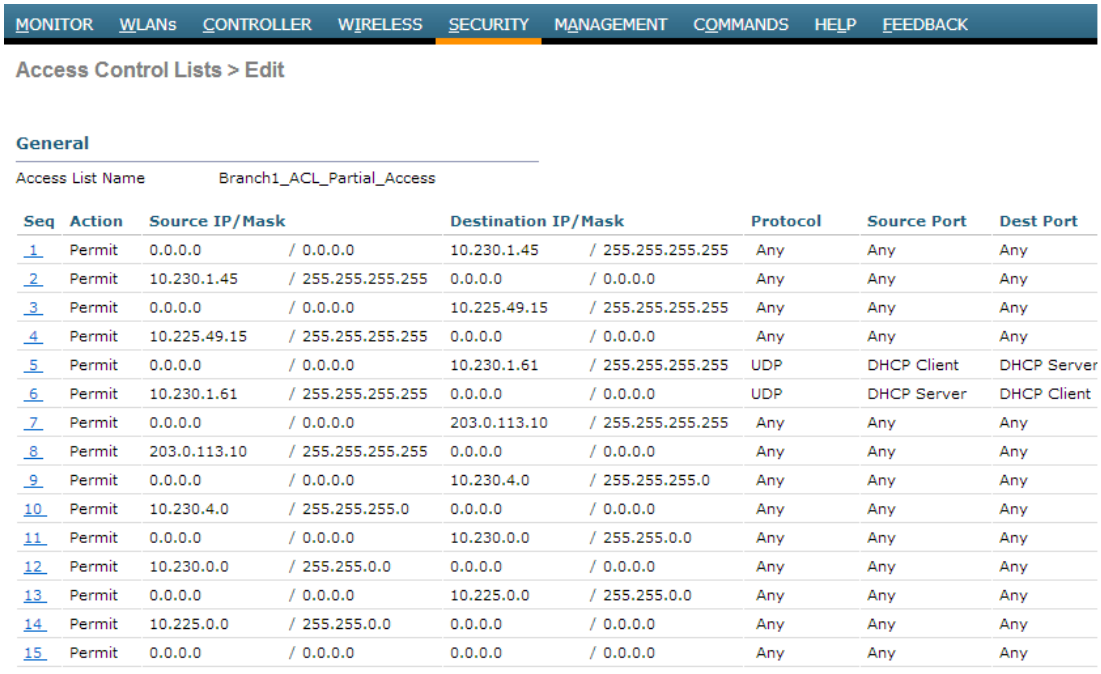
Deploying ACLs on the branch is slightly different than using ACLs with a centralized WLC. For branch locations, the Cisco 7500 Flex Wireless Controller relies on FlexConnect ACLs to enforce policy permissions within this design guide. FlexConnect ACLs are created on the WLC and configured with the VLAN defined on the AP or the FlexConnect Group using the VLAN-ACL mapping for dynamic or AAA override VLANs. These FlexConnect ACLs are pushed to the APs when the authorization policy matches. This design guide relies on FlexConnect groups to enforce Flex ACLs for each VLAN. The steps are as follows:

1. Create a FlexConnect ACL for each branch.
2. Apply the FlexConnect ACL on the FlexConnect group for each branch.
3. Define the VLAN-ACL mapping for each VLAN.

On the FlexConnect 7500 Controller, click **Security > Access Control Lists > FlexConnect ACLs** and define the ACL rules for Partial Access. Figure 15-31 shows the Branch1_ACL_Partial_Access ACL, which allows access to the Internet and some internal resources.

Note A unique ACL may be needed for each branch since each branch location may have its own local resources and unique IP address space.

Figure 15-31 Branch1_ACL_Partial_Access



MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Access Control Lists > Edit

General

Access List Name Branch1_ACL_Partial_Access

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client
7	Permit	0.0.0.0	/ 0.0.0.0	203.0.113.10	/ 255.255.255.255	Any	Any	Any
8	Permit	203.0.113.10	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
9	Permit	0.0.0.0	/ 0.0.0.0	10.230.4.0	/ 255.255.255.0	Any	Any	Any
10	Permit	10.230.4.0	/ 255.255.255.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
11	Permit	0.0.0.0	/ 0.0.0.0	10.230.0.0	/ 255.255.0.0	Any	Any	Any
12	Permit	10.230.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
13	Permit	0.0.0.0	/ 0.0.0.0	10.225.0.0	/ 255.255.0.0	Any	Any	Any
14	Permit	10.225.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
15	Permit	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any

293724

The above ACL specifies the following access:

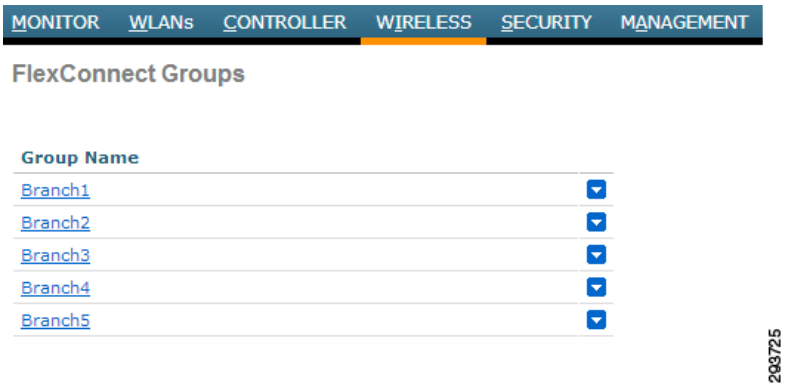
- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Allow IP access to and from the MDM Server (203.0.113.10).
- Allow IP access to and from specific subnet (10.230.4.0 /24). Similar ACL entries could be added to allow access to Branch1 subnets/servers.
- Deny IP access to and from data center subnets (10.230.0.0 /16).
- Deny IP access to and from campus subnets (10.225.0.0 /16).
- Allow access to and from all other subnets (Internet access).

Note The access list shown is a generic example used to implement an arbitrary use case and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

For the purposes of this design guide, a FlexConnect group is defined for each branch, which allows for multiple FlexConnect access points in the branch to share configuration parameters.

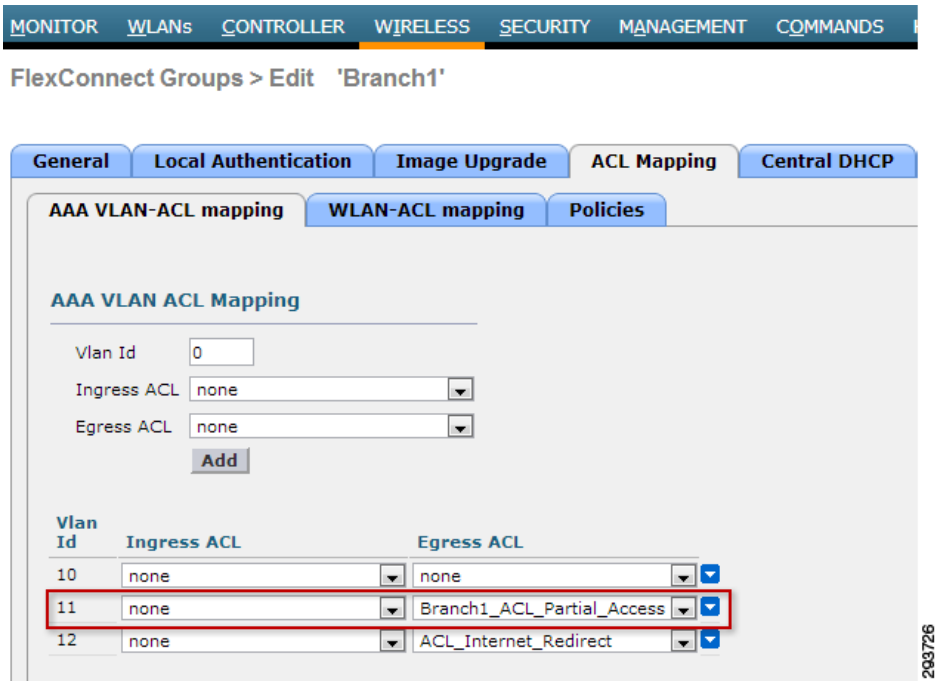
On the FlexConnect 7500 controller, click **Wireless > FlexConnect Groups** and select the FlexConnect group for a particular branch location, as shown in Figure 15-32.

Figure 15-32 FlexConnect Groups



In Figure 15-33, the FlexConnect group for Branch1 is applying the Branch1_ACL_Partial_Access ACL to endpoints connecting to VLAN 11.

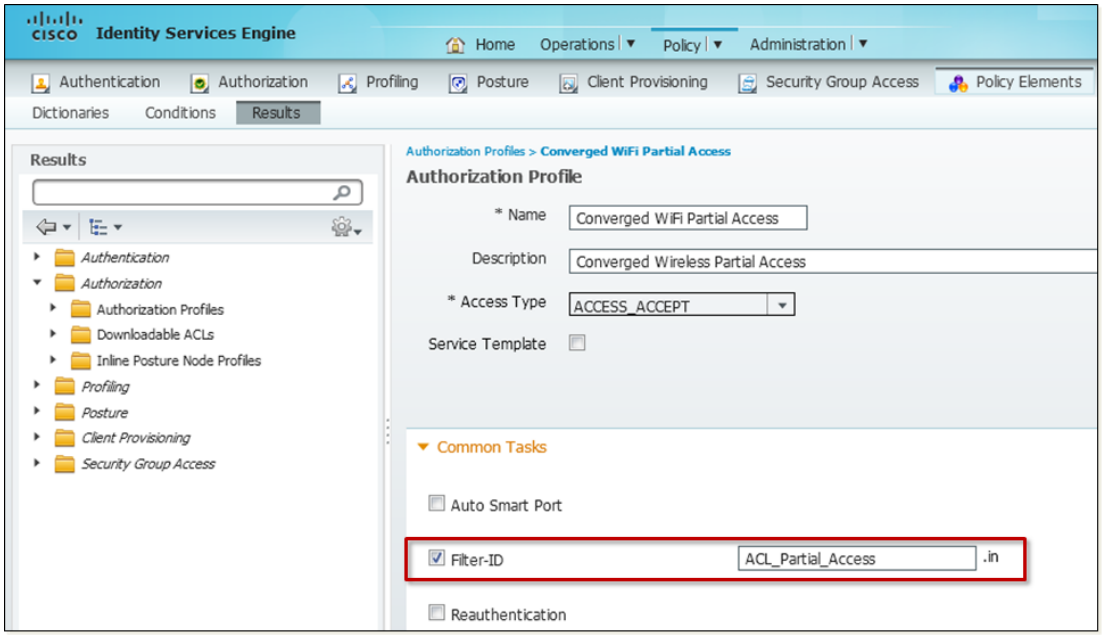
Figure 15-33 FlexConnect Group for Branch1



Converged Access Branch or Campus

For devices connecting from campus or branch locations which implement a Converged Access design, the Converged WiFi Partial Access authorization profile relies on the ACL_Partial_Access access list, enforced by the Catalyst 3850 Series switch. Figure 15-34 shows the authorization profile.

Figure 15-34 Converged WiFi Partial Access



Cisco Catalyst 3850 Series switches (and CT5760 wireless controllers) support both named ACLs and downloadable ACLs. For the Converged Access design a named ACL is implemented, meaning that the ACL must be configured on the Catalyst 3850 switch rather than being downloaded from ISE. Using the RADIUS Filter-ID attribute-value pair, ISE instructs the WLC to apply the ACL_Partial_Access ACL. The following configuration example shows the contents of this ACL, as defined in the Catalyst 3850 switch.

```
!
ip access-list extended ACL_Partial_Access
 permit udp any eq bootpc any eq bootps
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 203.0.113.10
 permit ip any 10.230.4.0 0.0.0.255
 permit ip any host 10.230.6.2
 permit ip any host 10.225.100.10
 deny   ip any 10.230.0.0 0.0.255.255
 deny   ip any 10.225.0.0 0.0.255.255
 deny   ip any 10.200.0.0 0.0.255.255
 permit ip any any
!
```

The access list shown in above similar to the access list shown in Figure 29, but configured on a Catalyst switch instead of a wireless controller. Hence the structure of the access list is slightly different. The access list specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the MDM Server (203.0.113.10).
- Allow IP access to and from specific servers (10.230.6.2 and 10.225.100.10).
- Deny IP access to and from data center subnets (10.230.0.0 /16).
- Deny IP access to and from campus subnets (10.225.0.0 /16).

- Deny IP access to and from branch subnets (10.200.0.0 /16).
- Allow access to and from all other subnets (Internet access).

Again, the access list shown in this example is generic and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

Note This design guide shows the use of the Radius:Airespace-ACL-Name AV pair for specifying a named ACL on CUWN wireless controller platforms, such as the CT5508 and Flex 7500 wireless controllers. However, it shows the use of the Radius:Filter-Id AV pair for specifying a named ACL on Cisco IOS-based wireless controller platforms, such as the CT5760 wireless controller and Catalyst 3850 Series switch. For wireless devices, this is simply to highlight the fact that the network administrator can utilize the legacy Airespace-ACL-Name AV pair or the RADIUS standard method of the Filter-Id AV pair within the BYOD implementation. It is recommended that the network administrator standardize on one of the two methods to specify a named ACL where possible in actual deployments.

Since a single ISE policy rule is defined for partial access of personal wireless devices from either a branch or campus with a converged access infrastructure, the access list can be customized to the specific branch or campus location as needed. The specific example ACL shown above is consistent with the ACL discussed in the campus centralized (Local Mode) controller design previously discussed.

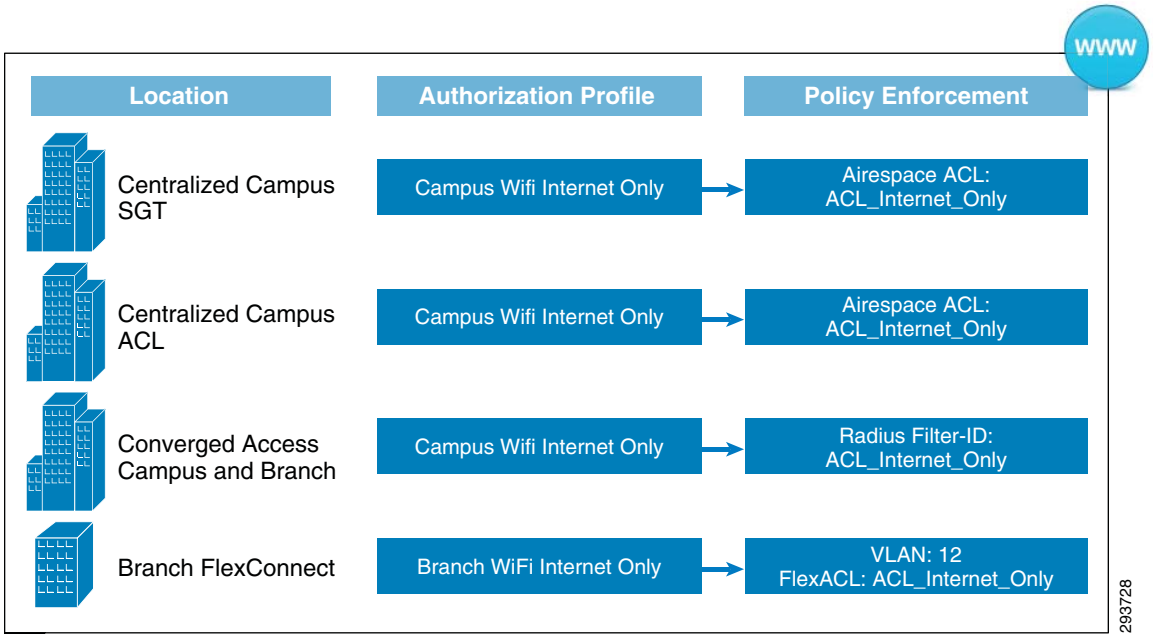
Personal Wireless Devices—Internet Only Access

To provide Internet Only access to personal devices, the Cisco ISE verifies the following:

- The employee has completed the on-boarding process through the Guest Registration portal.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate, in this case, the MAC address of the endpoint.
- The connection originated using EAP-TLS authentication.
- The user is a member of the Domain Users Active Directory group.

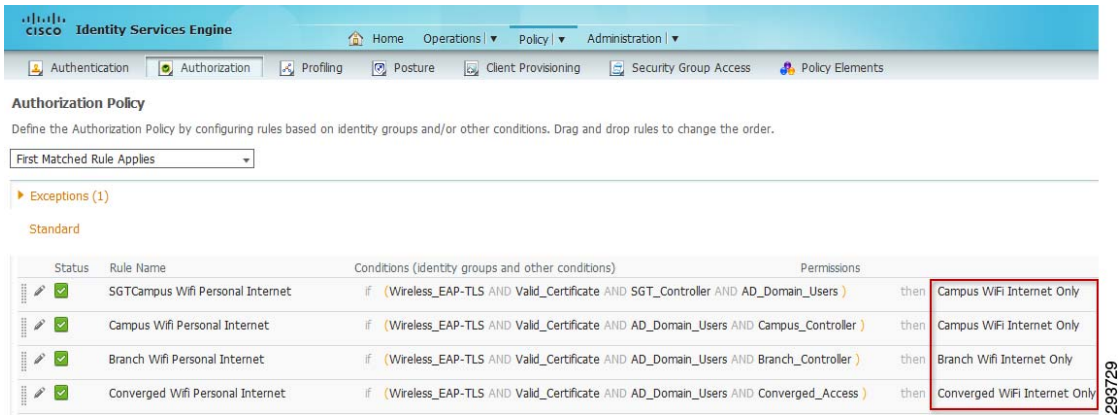
At a high level, Figure 15-35 shows how different authorization profiles are selected for devices coming from different locations with different wireless designs. Each authorization profile in turn enforces a unique permission using VLANs, SGTs, named ACLs, FlexConnect ACLs, etc.

Figure 15-35 Internet Only Enforcement



To configure the authorization rules in ISE, click **Policy > Authorization**. Figure 15-36 highlights the authorization policy to grant Internet Only access to personal devices.

Figure 15-36 Authorization Policies for Internet Only Access



Looking at the rules in more detail, ISE evaluates the following conditions:

- Wireless_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name (defined as a simple condition).
- The user belongs to the Domain Users Active Directory group (defined as a simple condition).
- The RADIUS authentication originated from a wireless controller which was a member of one of the following device groups—Campus_Controller, SGT_Controller, Branch_Controller, or Converged_Access (defined as a simple condition).

Simple and Compound Conditions explains the different conditions used in the rules.

Permissions

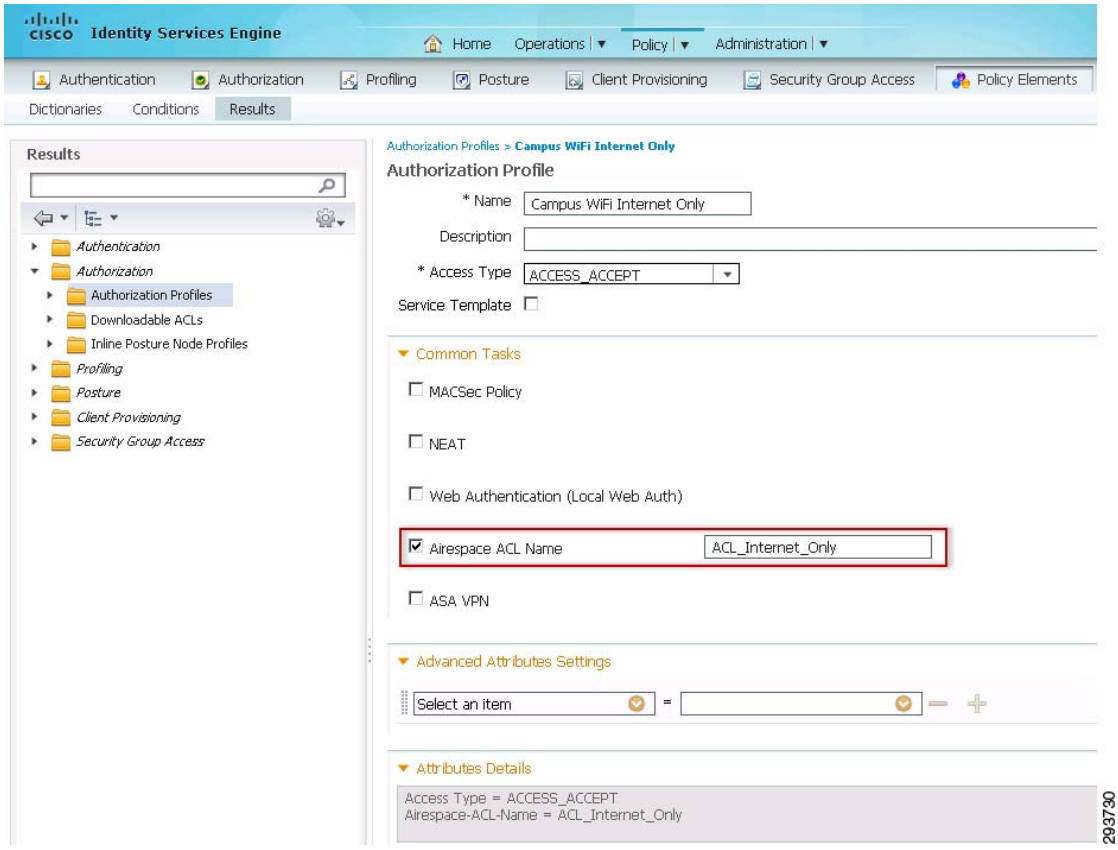
When all conditions in the authorization policy rules match, the rule invokes the proper permission. In the previous sections the permission result was either an authorization profile (non-SGT based access), or a standard result for SGT based access. However, in this section only the authorization profile will be used for the reasons explained below.

- Campus WiFi Internet Only for 802.1X wireless devices connecting from a SGT_Enabled controller or from a campus location with a centralized (Local Mode) wireless design. Both the campus controller and the SGT_Enabled controller use the same authorization profiles. SGTs are not used to tag Internet only traffic since SGT relies on source and destination tags and the destination tag for the Internet is unknown.
- Branch Wireless Internet Only for 802.1X wireless devices connecting from a branch location which implements a Flexconnect wireless design.
- Converged WiFi Internet Only for 802.1X wireless devices connecting from either a campus or branch location which implements a Converged Access infrastructure design.

Centralized Campus with SGT or ACLs

For devices connecting from a campus location which implements a centralized (Local Mode) wireless design or from a SGT_Enabled controller, the Campus WiFi Internet Only authorization profile relies on the ACL_Internet_Only access list enforced by the WLC. Figure 15-37 shows how the authorization profile instructs the WLC to apply the ACL.

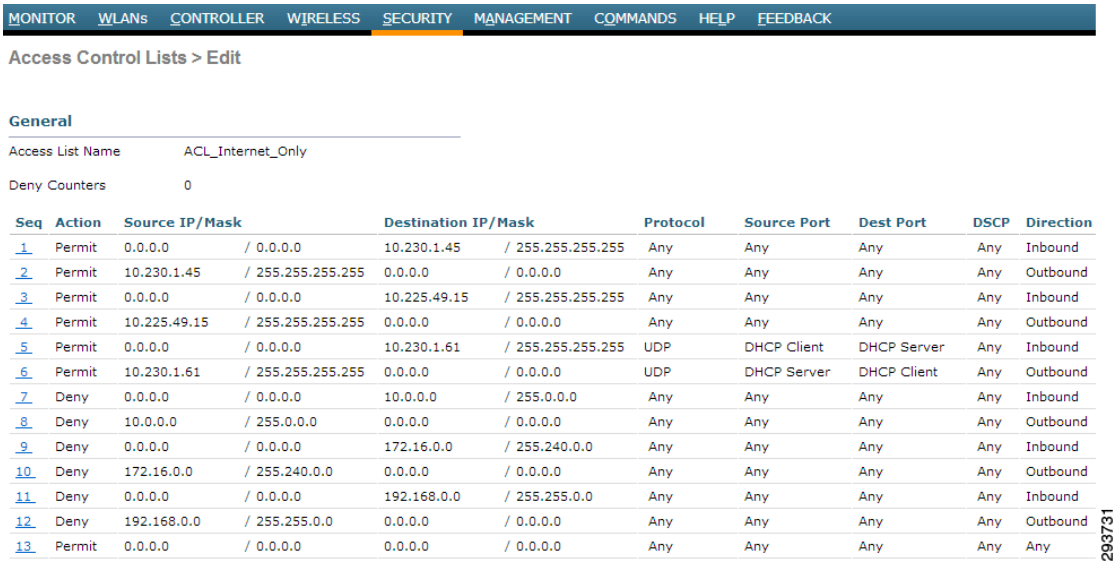
Figure 15-37 Campus WiFi Internet Only



Cisco Wireless LAN Controllers support named ACLs, meaning the ACL must be previously configured on the controller, rather than being downloaded from ISE. Using the RADIUS Airespace-ACL Name attribute-value pair, ISE instructs the WLC to apply the ACL_Internet_Only ACL.

Figure 15-38 shows the contents of this ACL, as defined in the CT5508 WLC campus controller.

Figure 15-38 ACL_Internet_Only



MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Access Control Lists > Edit

General

Access List Name ACL_Internet_Only

Deny Counters 0

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port	DSCP	Direction
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any	Any	Inbound
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any	Any	Inbound
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server	Any	Inbound
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client	Any	Outbound
7	Deny	0.0.0.0	/ 0.0.0.0	10.0.0.0	/ 255.0.0.0	Any	Any	Any	Any	Inbound
8	Deny	10.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
9	Deny	0.0.0.0	/ 0.0.0.0	172.16.0.0	/ 255.240.0.0	Any	Any	Any	Any	Inbound
10	Deny	172.16.0.0	/ 255.240.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
11	Deny	0.0.0.0	/ 0.0.0.0	192.168.0.0	/ 255.255.0.0	Any	Any	Any	Any	Inbound
12	Deny	192.168.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Outbound
13	Permit	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any	Any	Any

293731

The access list specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Deny IP access to and from internal network address space (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).
- Allow access to and from all other subnets (Internet access).

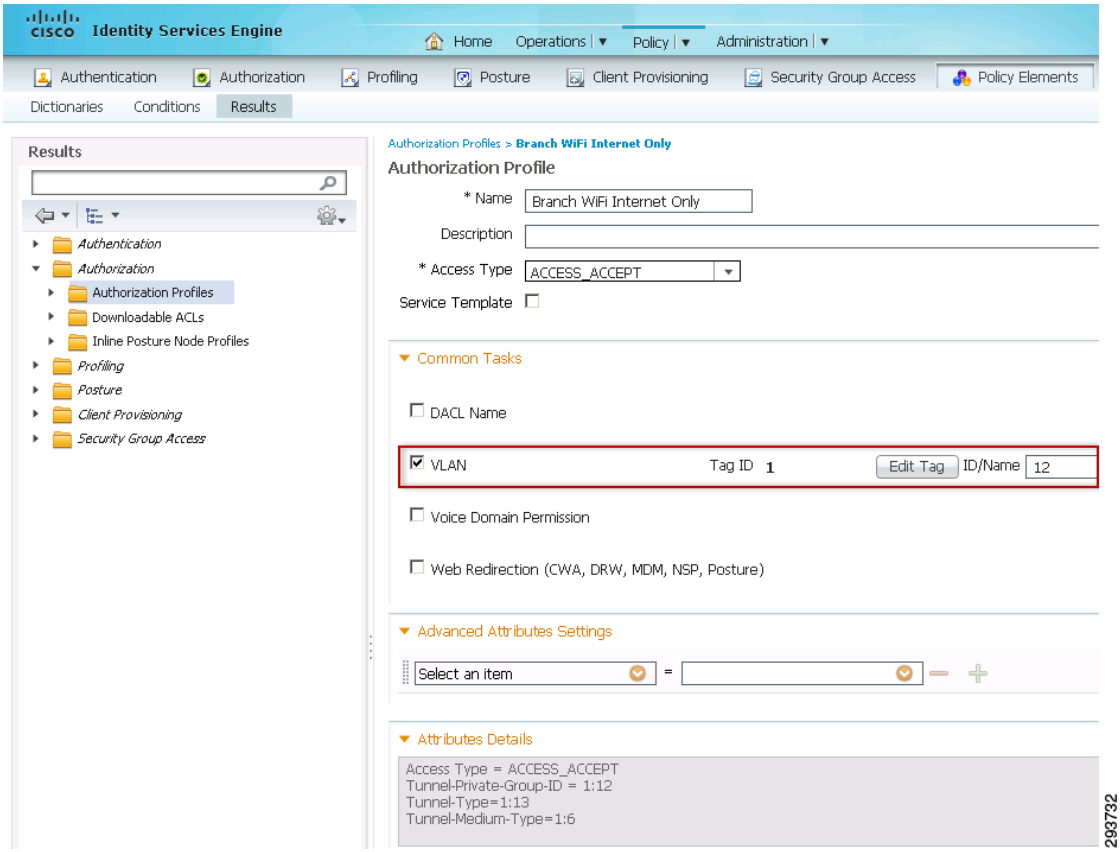
Once again, the access list is generic and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

Branches with FlexConnect

For devices connecting from a branch location which implements a FlexConnect wireless design, the Branch WiFi Internet Only authorization profile dynamically assigns the device to VLAN12, which is dedicated for devices obtaining Internet_Only access.

Figure 15-39 shows this authorization profile.

Figure 15-39 Branch WiFi Internet Only



For branch locations which implement FlexConnect wireless designs, the Cisco 7500 Flex Wireless Controller relies on FlexConnect ACLs to enforce policy permissions. FlexConnect ACLs are created on the WLC and configured with the VLAN defined on the AP or the FlexConnect Group using the VLAN-ACL mapping for dynamic or AAA override VLANs. These FlexConnect ACLs are pushed to the APs when the authorization policy matches.

1. Create a FlexConnect ACL for each branch.
2. Apply the FlexConnect ACL on the FlexConnect group for each branch.
3. Define the VLAN-ACL mapping for each VLAN.

On the FlexConnect 7500 Controller, click **Security > Access Control Lists > FlexConnect ACLs** and define the ACL rules for Internet_Only access. Figure 15-40 shows an example of the Branch_ACL_Internet_Only ACL, which only allows access to the Internet. This ACL is the same for all branches.

Figure 15-40 Branch_ACL_Internet_Only

MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Access Control Lists > Edit

General

Access List Name Branch_ACL_Internet_Only

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client
7	Deny	0.0.0.0	/ 0.0.0.0	10.0.0.0	/ 255.0.0.0	Any	Any	Any
8	Deny	10.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
9	Deny	0.0.0.0	/ 0.0.0.0	172.16.0.0	/ 255.240.0.0	Any	Any	Any
10	Deny	172.16.0.0	/ 255.240.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
11	Deny	0.0.0.0	/ 0.0.0.0	192.168.0.0	/ 255.255.0.0	Any	Any	Any
12	Deny	192.168.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
13	Deny	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any

293733

The access list specifies the following access:

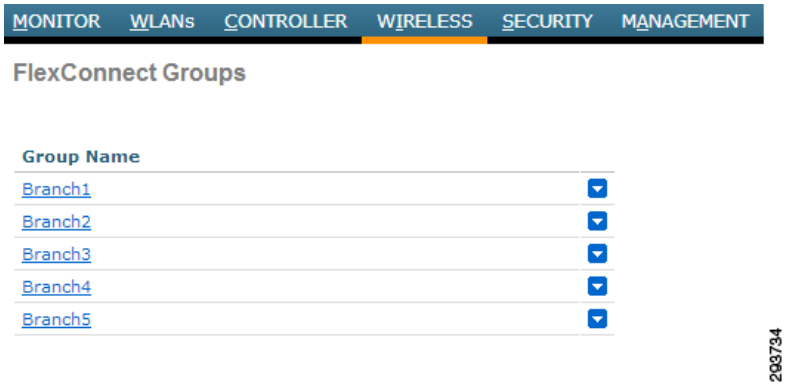
- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Deny IP access to and from internal network address space (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).
- Allow access to and from all other subnets (Internet access).

The access list is generic and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

For the purposes of this design guide, a FlexConnect group is defined for each branch, which allows for multiple FlexConnect access points in the branch to share configuration parameters.

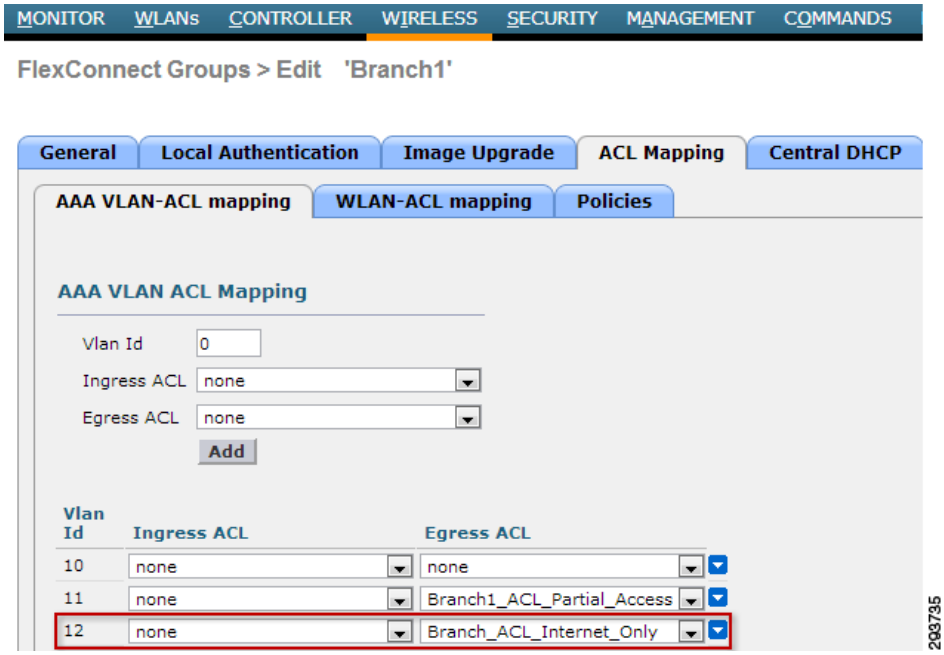
On the FlexConnect 7500 controller, click **Wireless > FlexConnect Groups** and select the FlexConnect group for a particular branch location, as shown in Figure 15-41.

Figure 15-41 FlexConnect Groups



In Figure 15-42, the FlexConnect group for Branch1 is applying the Branch_ACL_Internet_Only ACL to endpoints connecting to VLAN 12.

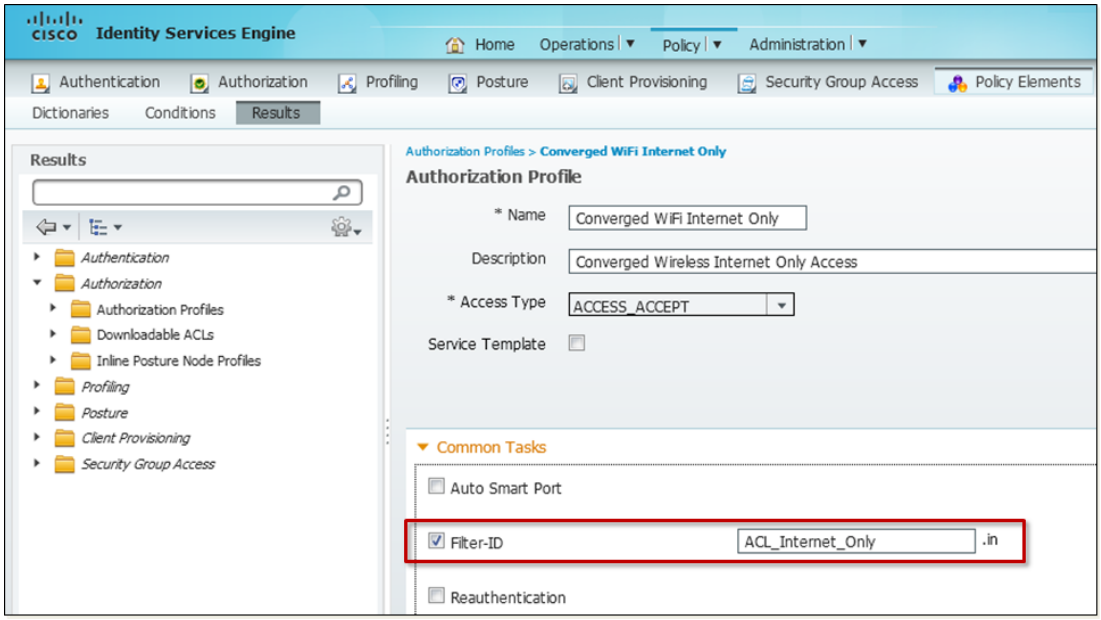
Figure 15-42 FlexConnect Group for Branch1 Internet Only



Converged Access Branch or Campus

For devices connecting from campus or branch locations which implement a Converged Access design, the Converged WiFi Internet Only authorization profile relies on the ACL_Internet_Only access list, enforced by the Catalyst 3850 Series switch. Figure 15-43 shows the authorization profile.

Figure 15-43 Converged WiFi Internet Only



Cisco Catalyst 3850 Series switches (and CT5760 wireless controllers) support both named ACLs and downloadable ACLs. For the Converged Access design a named ACL is implemented, meaning that the ACL must be configured on the Catalyst 3850 switch rather than being downloaded from ISE. Using the RADIUS Filter-ID attribute-value pair, ISE instructs the WLC to apply the ACL_Internet_Only ACL. The following configuration example shows the contents of this ACL, as defined in the Catalyst 3850 switch.

```
ip access-list extended ACL_Internet_Only
 permit udp any eq bootpc any eq bootps
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.225.100.10
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 172.16.0.0 0.15.255.255
 deny   ip any 192.168.0.0 0.0.255.255
 permit ip any any
!
```

The access-list shown in above is similar to the access list shown in Figure 15-38, but configured on a Catalyst switch instead of a wireless controller. Hence the structure of the access-list is slightly different. However the access-list specifies the following access:

- Allow DHCP access (bootpc and bootps).
- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE server (10.225.49.15).
- Deny IP access to and from the rest of the internal network IP address space (10.0.0.0 /8, 172.16.0.0 /12, 192.168.0.0 /16).
- Allow access to all other addresses (Internet addresses).

Note The MDM above is shown with private RFC1918 address. In practice, the MDM must be reachable from the public Internet and may not require a specific line entry in the ACL.

Again, the access list shown in this example is generic and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

Personal Wired Devices—Full Access

To provide full access to personal wired devices, the Cisco ISE verifies the following:

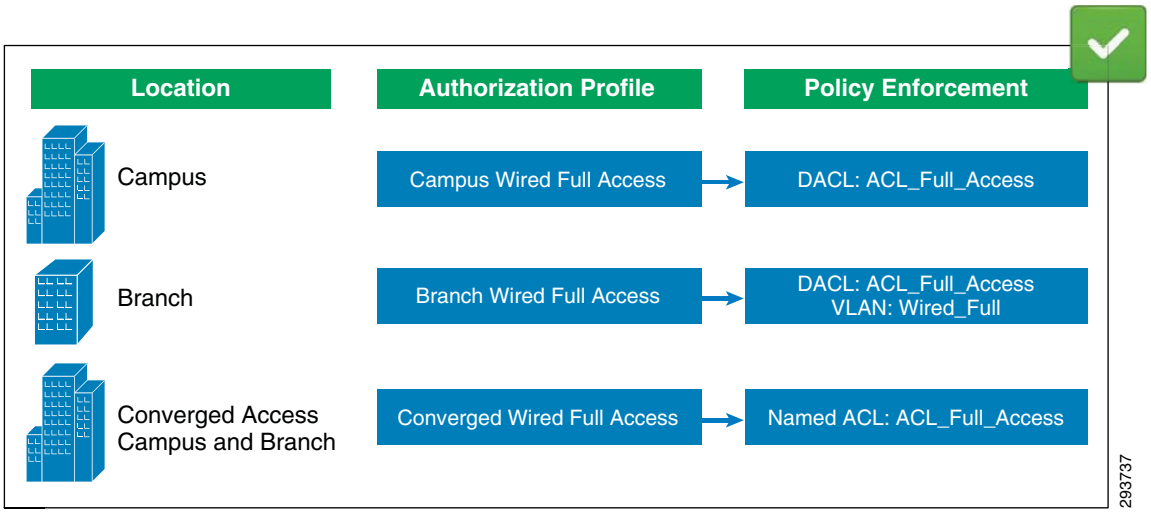
- The employee has completed the on-boarding process through the Guest Registration portal.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate, in this case, the MAC address of the endpoint.
- The connection originated using EAP-TLS authentication.
- The user is a member of the BYOD_Full_Access Active Directory group.

Since the wired designs presented in this design guide rely on slightly different access control mechanisms for Converged Access campuses and branches, for campuses and branches which do not implement Converged Access infrastructures, unique authorization rules are created for connections originating from each design.

Note For the purpose of clarity within this design guide, Converged Access branches refer to designs in which Catalyst 3850 Series switches are deployed at the access-layer of the branch network. From a wired perspective, branches which do not implement Converged Access infrastructures are branches which deploy other Catalyst access-layer switches, such as the Catalyst 3750X Series. Unless otherwise specified, these will be referred to simply as “branches” within the design guide. Similarly, Converged Access campuses refer to designs in which Catalyst 3850 Series switches are deployed at the access-layer of building distribution modules within the campus. From a wired perspective, campuses which do not implement Converged Access infrastructures are campuses which deploy other Catalyst access-layer switches, such as the Catalyst 3750X Series. Unless otherwise specified, these will be referred to simply as “campuses” within this design guide. This is in order to minimize the use of verbose phrases such as “branches which do not implement Converged Access infrastructure” and “campuses which do not implement Converged Access infrastructure”.

At a high level, Figure 15-44 shows how different authorization profiles are selected for connections originating from different locations with different infrastructure designs. Each authorization profile in turn enforces a unique permission using VLANs, dynamic ACLs (either named or downloadable [DACL]), etc.

Figure 15-44 Full Access Wired Enforcement

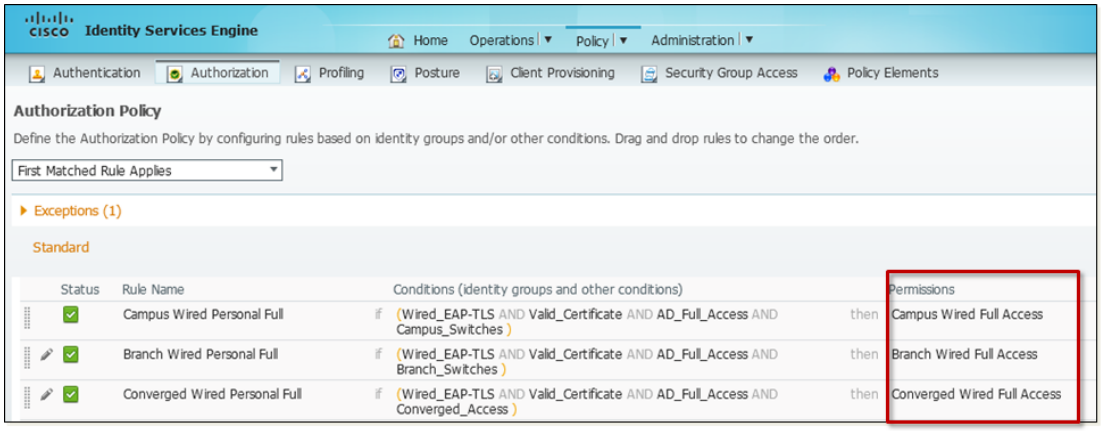


Note Wired assignment of Security Group Tags (SGTs) is not discussed within this version of the design guide. Hence, there is no wired policy enforcement via SGTs in Figure 15-44. Future versions of this design guide may address wired SGT assignment.

To differentiate these connections, the ISE relies on Network Device Groups to group Catalyst switches based on their location or device type. This allows a single ISE to enforce policies across different groups of devices. Each Catalyst switch needs to be added to the proper device group by clicking **Administration > Network Resources > Network Devices** and specifying the proper location or device type from the pull-down menu.

Figure 15-45 shows the details of authorization profile configured in ISE for wired devices.

Figure 15-45 Authorization Policies for Wired Full Access



Looking at the rules in more detail, ISE evaluates the following conditions:

- Wired_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name. (defined as a simple condition).

- The user belongs to a specific Active Directory group (defined as a simple condition).
- The RADIUS authentication originated from a Catalyst switch which was a member of one of the following device groups—Campus_Switches, Branch_Switches, or Converged_Access (defined as a simple condition).

Wired Simple and Compound Conditions

To improve the readability of the authorization policy, simple and compound authorization conditions were defined to group different conditions. These conditions may be reused and modified without changing every authorization rule.

Table 15-8 shows the conditions used in the authorization rules.

Table 15-8 Simple and Compound Conditions

Wired EAP-TLS (Compound)	
Wired_EAP-TLS	Radius:Service-Type Equals Framed Radius:NAS-Port-Type Equals Ethernet Network Access:EapAuthentication Equals EAP-TLS
Check for Valid Certificate (Simple)	
Valid_Certificate	Radius:Calling-Station-ID EQUALS CERTIFICATE:Subject Alternative Name
Active Directory Group (Simple)	
AD_ Full_Access	AD1:ExternalGroups EQUALS sdulab.com/Users/BYOD_Full_Access
AD_ Partial_Access	AD1:ExternalGroups EQUALS sdulab.com/Users/BYOD_Partial_Access
AD_Domain_users	AD1:ExternalGroups EQUALS sdulab.com/Users/Domain User
WLC Location or Device Type (Simple)	
Campus_Switches	DEVICE:Location EQUALS All Locations#Campus_Switches
Branch_Switches	DEVICE:Location EQUALS All Locations#Branch_Switches
Converged_Access	DEVICE:Device Type EQUALS All Device Types#Converged

Permission

When all conditions in the authorization policy rule match, the rule invokes the proper permission. The permissions can be of different forms such as an authorization profile or a standard result. In this design guide, for wired access the authorization policy is used as permission to the policy rules match. Table 15-9 explains the permissions used for the full access for wired users.

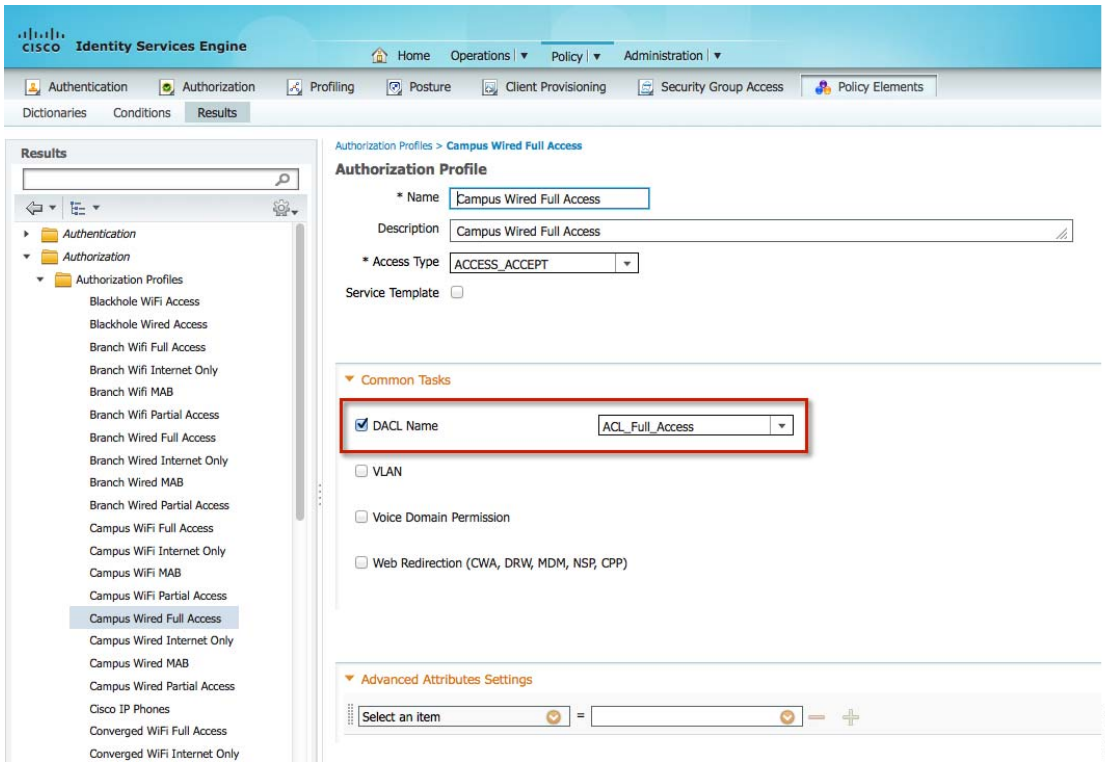
Table 15-9 Permissions Used for Wired Full Access

Permission name	Permission type	Purpose
Campus Wired Full Access	Authorization profile	Provides Full Access for 802.1X wired devices connecting from campus location.
Branch Wired Full Access	Authorization profile	To push a VLAN for 802.1X wired devices connecting from a branch location.
Converged Wired Full Access	Authorization profile	To push a named ACL for 802.1X wired devices connecting from either a campus or branch location with Converged Access.

Campus Wired

Figure 15-46 shows how the Campus Wired Full Access authorization profile is defined in ISE.

Figure 15-46 Campus Wired Full Access Authorization Profile

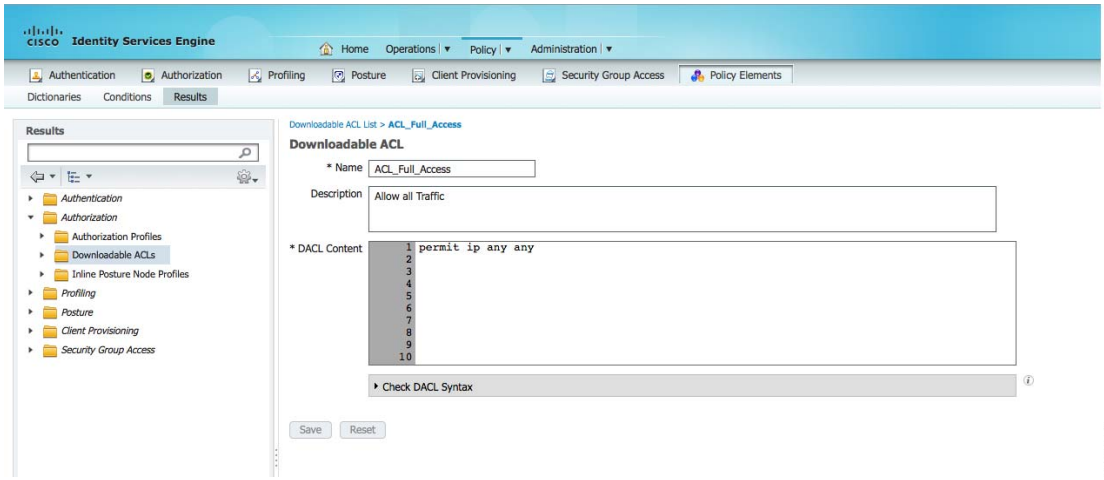


Note Cisco Catalyst switches support both downloadable ACLs (DACLs) and named ACLs. This design guide shows the use of downloadable ACLs for access control of wired devices when implementing a non-Converged Access infrastructure and the use of named ACLs for access control of wired devices when implementing a Converged Access infrastructure. This is done to show the range of capabilities for access control available on Cisco IOS-based platforms. The same wired policy enforcement for Converged Access and non-Converged Access infrastructure can be achieved if the customer desires,

simply by using either downloadable ACLs (DACLs) or named ACLs for both designs. Both downloadable and named ACLs have advantages and disadvantages, depending upon where they are deployed within the network.

The downloadable ACL (DACL), which allows all IP traffic, overrides the default-ACL configured on the switch port. The default-ACL is used as an additional preventative measure in case the downloadable ACL is not applied to the switch port for some reason. Figure 15-47 shows the DACL definition which must be configured in ISE:

Figure 15-47 DACL Definition in ISE that Grants Full Access to the Users

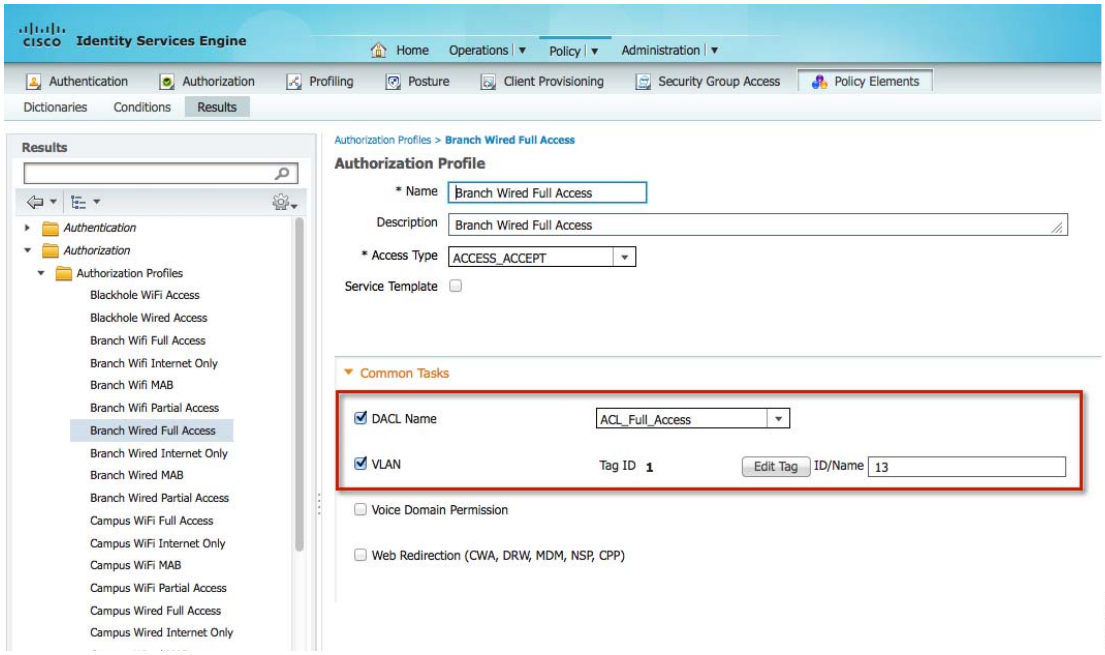


Note ISE 1.2 has the ability to check the DACL syntax. This feature should be utilized in order to minimize the possibility that the DACL is not applied due to a syntax error.

Branch Wired

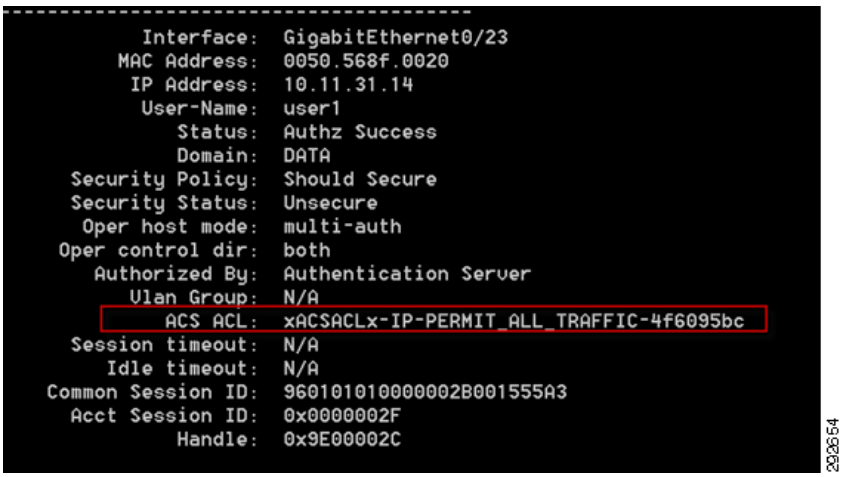
Figure 15-48 shows how a similar authorization profile is constructed for devices connected at a branch which does not have a Converged Access infrastructure. The authorization profile pushes VLAN information—in this case the VLAN name—along with the ACL information.

Figure 15-48 Branch Wired Full Access Authorization Profile



Once this profile is downloaded to the Catalyst switch, the endpoint gets full access to the network. Figure 15-49 shows an example of the state of the port when this profile is downloaded by using the switch command **show authentication session interface Gi0/23**.

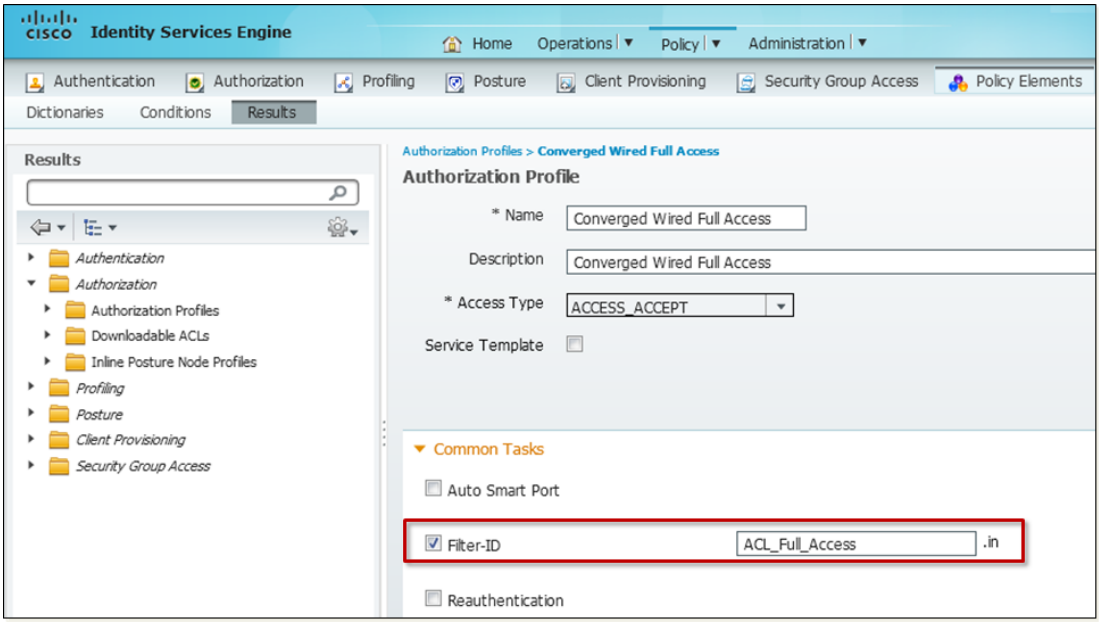
Figure 15-49 Catalyst Switch Port



Converged Access Branch or Campus

Figure 15-50 shows how the Converged Wired Full Access authorization profile is defined in ISE.

Figure 15-50 Converged Wired Full Access Authorization Profile



For the Converged Access design a named ACL is implemented, meaning that the ACL must be configured on the Catalyst 3850 switch rather than being downloaded from ISE. Using the RADIUS Filter-ID attribute-value pair, ISE instructs the converged access switch to apply the ACL_Full_Access ACL. The following configuration example shows the contents of this ACL, as defined in the Catalyst 3850 switch.

```
!
ip access-list extended ACL_Full_Access
 permit ip any any
!
```

This ACL is used to override the default-ACL configured on the switch port. The default-ACL is used as an additional preventative measure in case ACL_Full_Access is not configured on the switch.

Personal Wired Devices—Partial Access

Partial Access grants access to corporate resources, in addition to Internet access. As mentioned in Personal Wireless Devices—Partial Access, once a device authenticates to ISE, an authorization profile is applied. For wired devices, the authorization profile is applied to the access layer switch.

To provide partial access to personal wired devices, the Cisco ISE verifies the following:

- The employee has completed the on-boarding process through the Guest Registration portal.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate, in this case, the MAC address of the endpoint.
- The connection originated using EAP-TLS authentication.
- The user is a member of the AD_Partial_Access Active Directory group.

At a high level, Figure 15-51 shows how different authorization profiles are selected for devices coming from different locations with different wired infrastructure designs. Each authorization profile in turn enforces a unique permission using VLANs, dynamic ACLs (either named or downloadable [DACL]), etc.

Figure 15-51 Wired Partial Access Enforcement

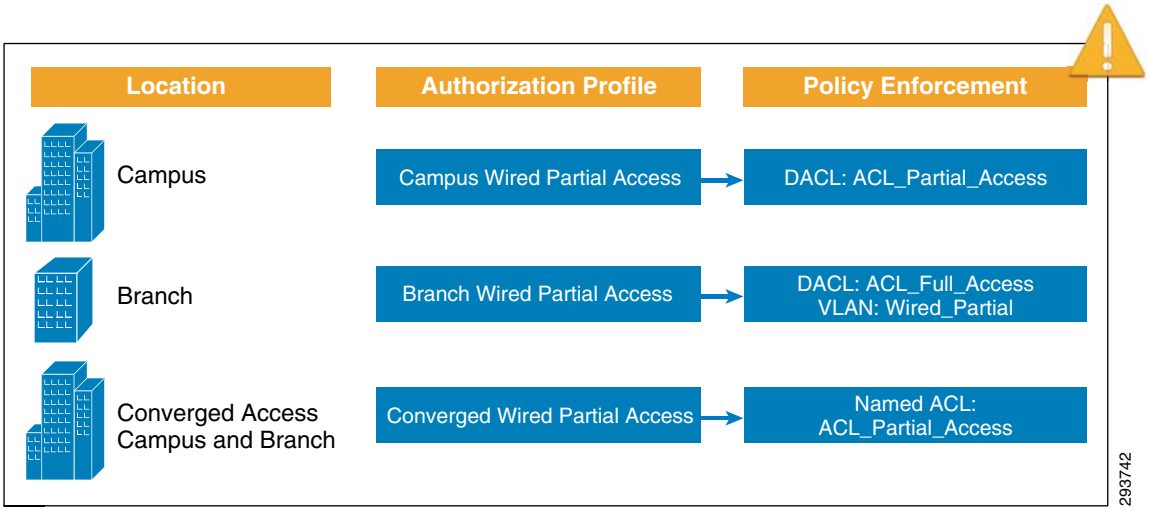
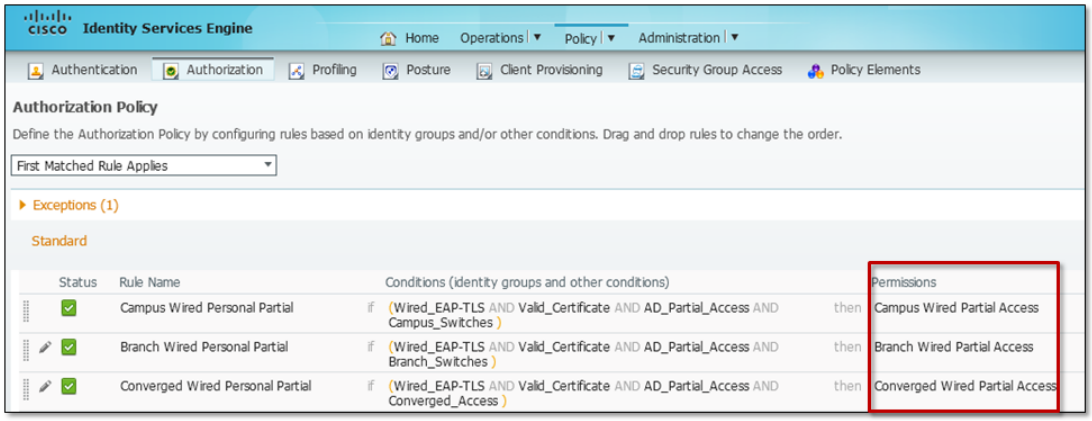


Figure 15-52 shows the details of authorization profile configured in ISE for wired devices.

Figure 15-52 Authorization Policies for Wired Partial Access



Looking at the rules in more detail, ISE evaluates the following conditions:

- Wired_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name. (defined as a simple condition).
- The user belongs to a specific Active Directory group (defined as a simple condition).
- The RADIUS authentication originated from a Catalyst switch which was a member of one of the following device groups—Campus_Switches, Branch_Switches, or Converged_Access (defined as a simple condition).

Wired Simple and Compound Conditions explains the different conditions used in the rules.

Permissions

When all conditions in the authorization policy rule match, the rule invokes the proper permission. The permissions can be of different forms such as an authorization profile or a standard result. In this design guide, for wired access the authorization policy is used as permission to the policy rules match. Table 15-10 explains the permissions used for the partial access for wired users.

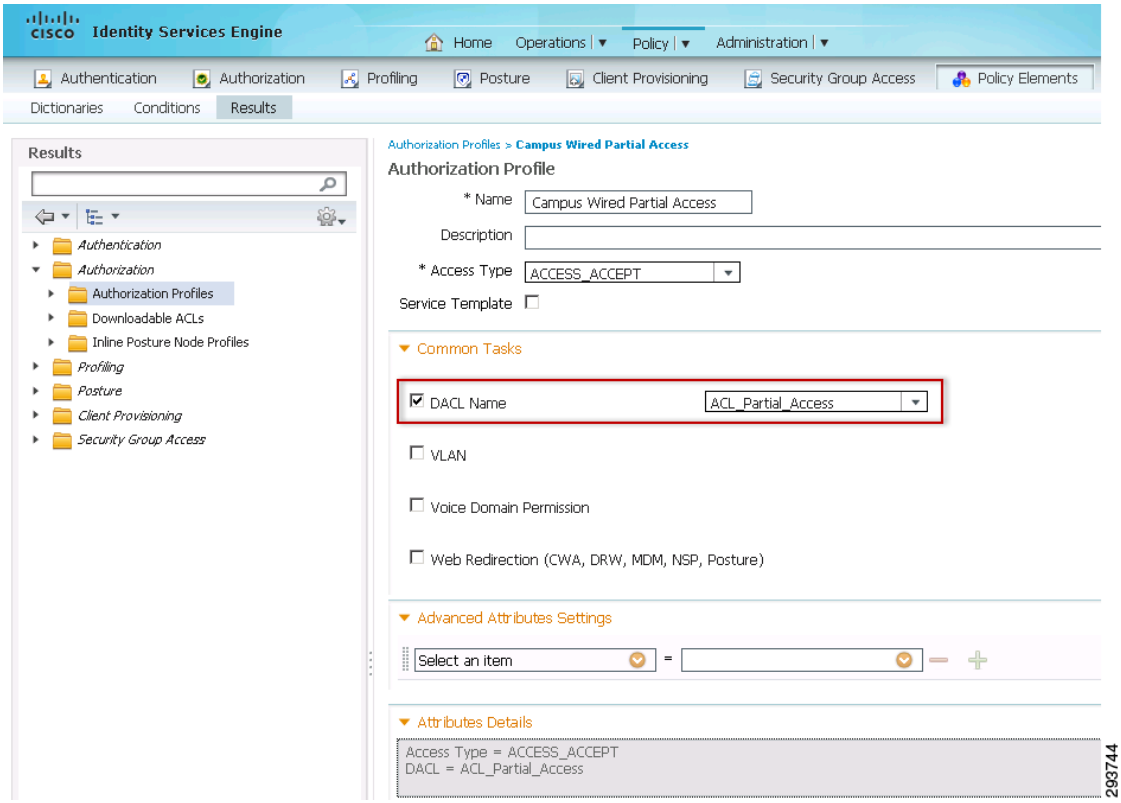
Table 15-10 Permissions Used for Wired Partial Access

Permission name	Permission type	Purpose
Campus Wired Partial Access	Authorization profile	To push a DACL for 802.1X wired devices connecting from campus location.
Branch Wired Partial Access	Authorization profile	To push a VLAN for 802.1X wired devices connecting from a branch location.
Converged Wired Partial Access	Authorization profile	To push a named ACL for 802.1X wired devices connecting from either a campus or branch location with Converged Access.

Campus Wired

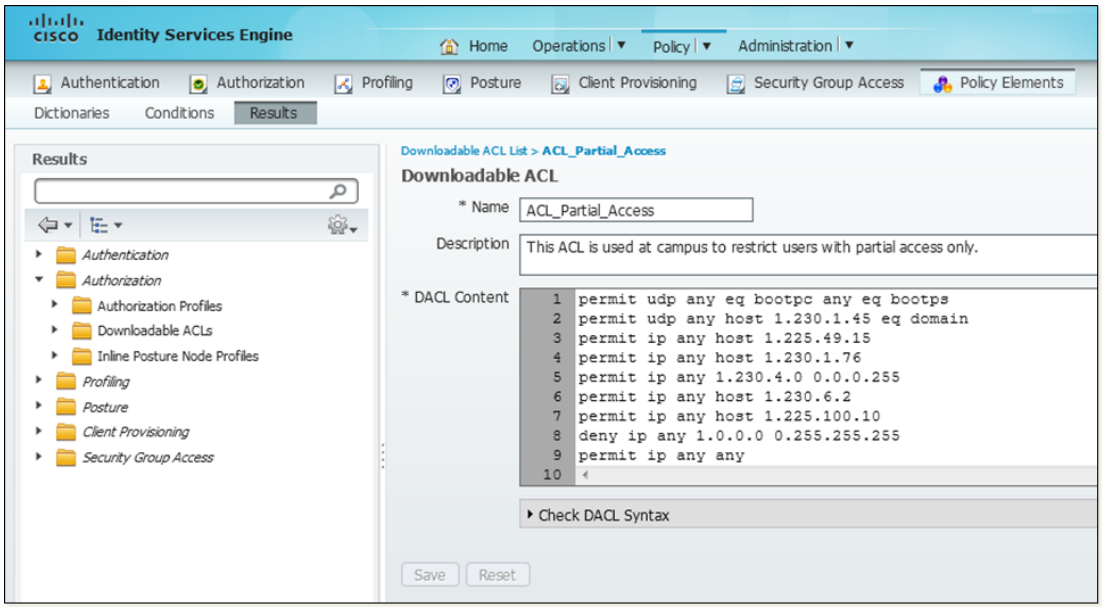
For devices connecting from a campus location, the Campus Wired Partial Access authorization uses a DACL named ACL_Partial_Access enforced by the access layer switch, as shown in Figure 15-53.

Figure 15-53 Campus Wired Partial Access



The DACL overrides the default-ACL configured on the switch. Figure 15-54 shows an example of this ACL, which is configured within ISE.

Figure 15-54 ACL_Partial_Access within ISE



The above ACL specifies the following access:

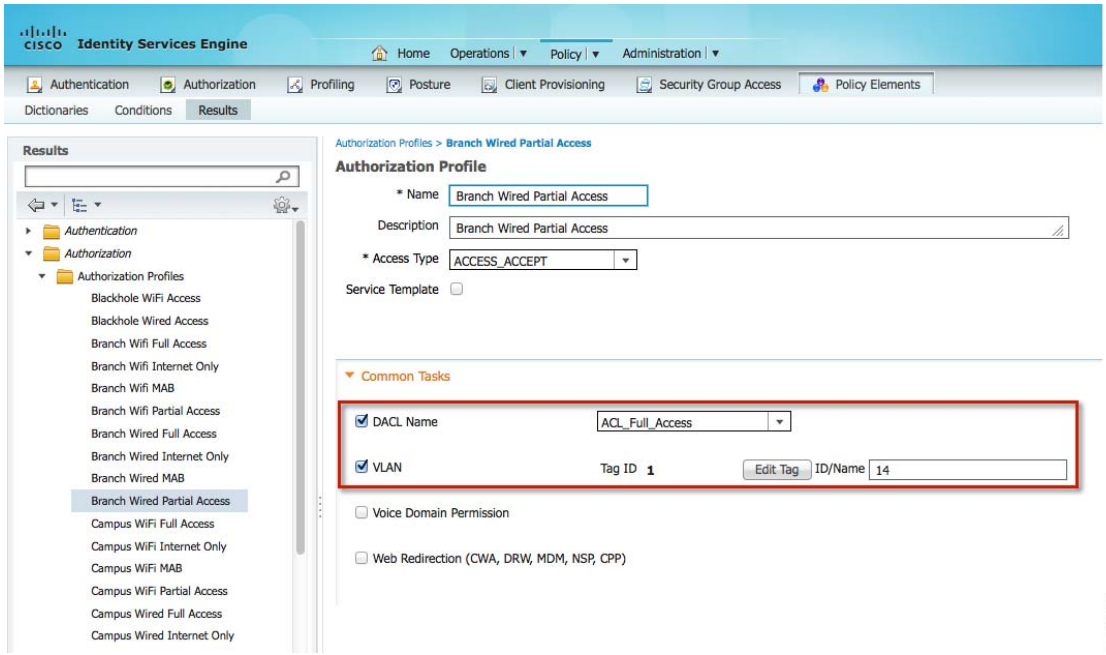
- Allow DHCP access (bootpc and bootps).
- Allow DNS access to the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from specific subnet (10.230.4.0 /24).
- Allow IP access to and from specific servers (10.230.6.2 and 10.225.100.10).
- Deny IP access to and from internal network address space (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).
- Allow access to and from all other subnets (Internet access).

Branch Wired

For devices connecting from a branch location, the Branch Wired Partial Access authorization pushes a VLAN assignment along with a downloadable ACL.

Figure 15-55 shows the details of the authorization profile configured in ISE.

Figure 15-55 Branch Wired Partial Access Authorization Profile



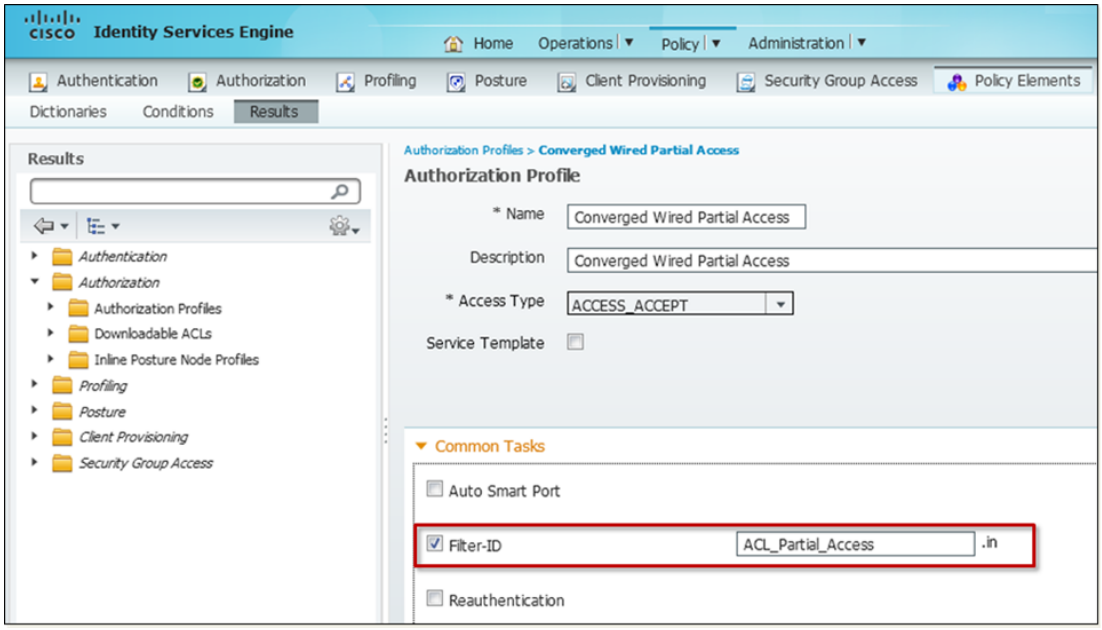
The DACL allows all IP traffic, so that the traffic initiated by the host reaches the branch router where the router-ACL is applied.

See VLAN Design at Branch Locations in Chapter 11, "BYOD Wired Infrastructure Design" for details regarding how this assignment of VLAN 14 is given full access for a branch which does not implement a Converged Access infrastructure.

Converged Access Branch and Campus

Figure 15-56 shows how the Converge Wired Partial Access authorization profile is defined in ISE.

Figure 15-56 Converged Wired Partial Access Authorization Profile



For the Converged Access design a named ACL is implemented. Using the RADIUS Filter-ID attribute-value pair, ISE instructs the converged access switch to apply the ACL_Partial_Access ACL. This is the same ACL discussed for converged access infrastructure in Personal Wireless Devices—Partial Access. For wired devices, the ACL is used to override the default-ACL configured on the switch port. The default-ACL is used as an additional preventative measure in case ACL_Partial_Access is not configured on the switch.

Personal Wired Devices—Internet Only Access

To provide Internet Only access to personal devices, the Cisco ISE verifies the following:

- The employee has completed the on-boarding process through the Guest Registration portal.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate, in this case, the MAC address of the endpoint.
- The connection originated using EAP-TLS authentication.
- The user is a member of the Domain Users Active Directory group.

At a high level, Figure 15-57 shows how different authorization profiles are selected for devices coming from different locations with different wired infrastructure designs. Each authorization profile in turn enforces a unique permission using VLANs, dynamic ACLs (either named or downloadable [DACL]), etc.

Figure 15-57 Wired Internet Only Access Enforcement

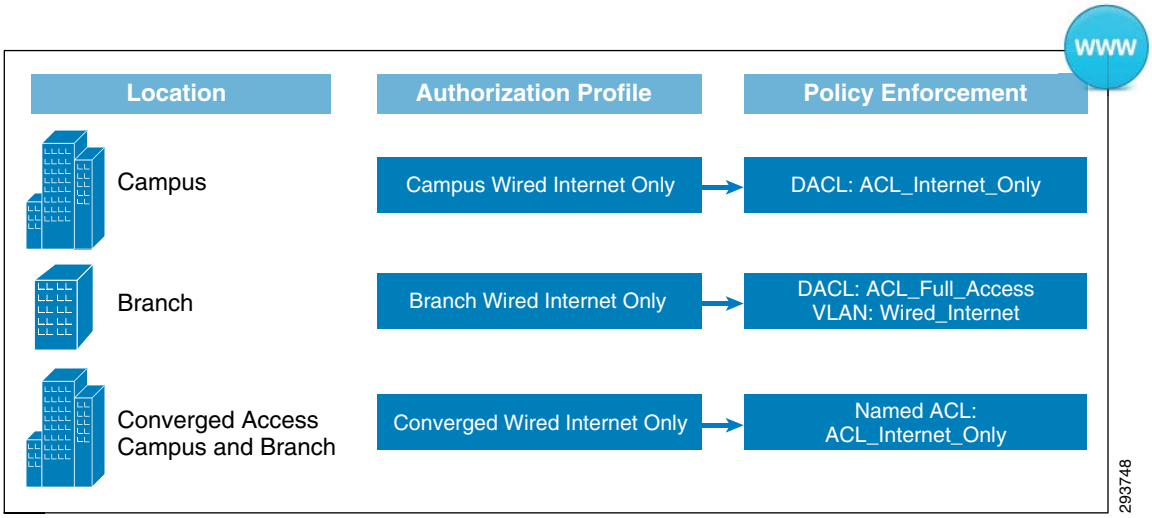
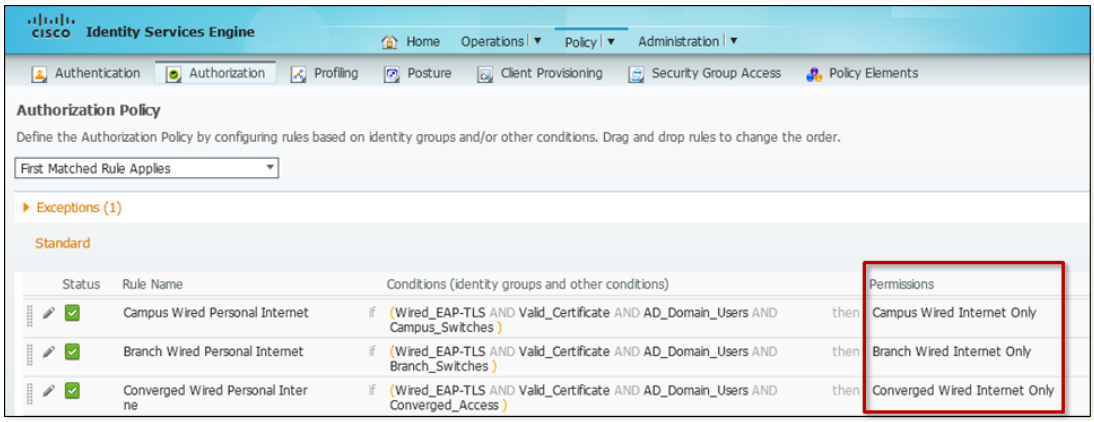


Figure 15-58 highlights the authorization policy to grant Internet Only access to personal wired devices.

Figure 15-58 Authorization Policies for Wired Internet Only Access



Looking at the rules in more detail, ISE evaluates the following conditions:

- Wired_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name. (defined as a simple condition).
- The user belongs to a specific Active Directory group (defined as a simple condition).
- The RADIUS authentication originated from a Catalyst switch which was a member of one of the following device groups—Campus_Switches, Branch_Switches, or Converged_Access (defined as a simple condition).

Wired Simple and Compound Conditions explains the different conditions used in the rules.

Permissions

When all conditions in the authorization policy rule match, the rule invokes the proper permission. The permissions can be of different forms such as an authorization profile or a standard result. In this design guide, for wired access the authorization policy is used as permission to the policy rules match. Table 15-11 explains the permissions used for Internet access.

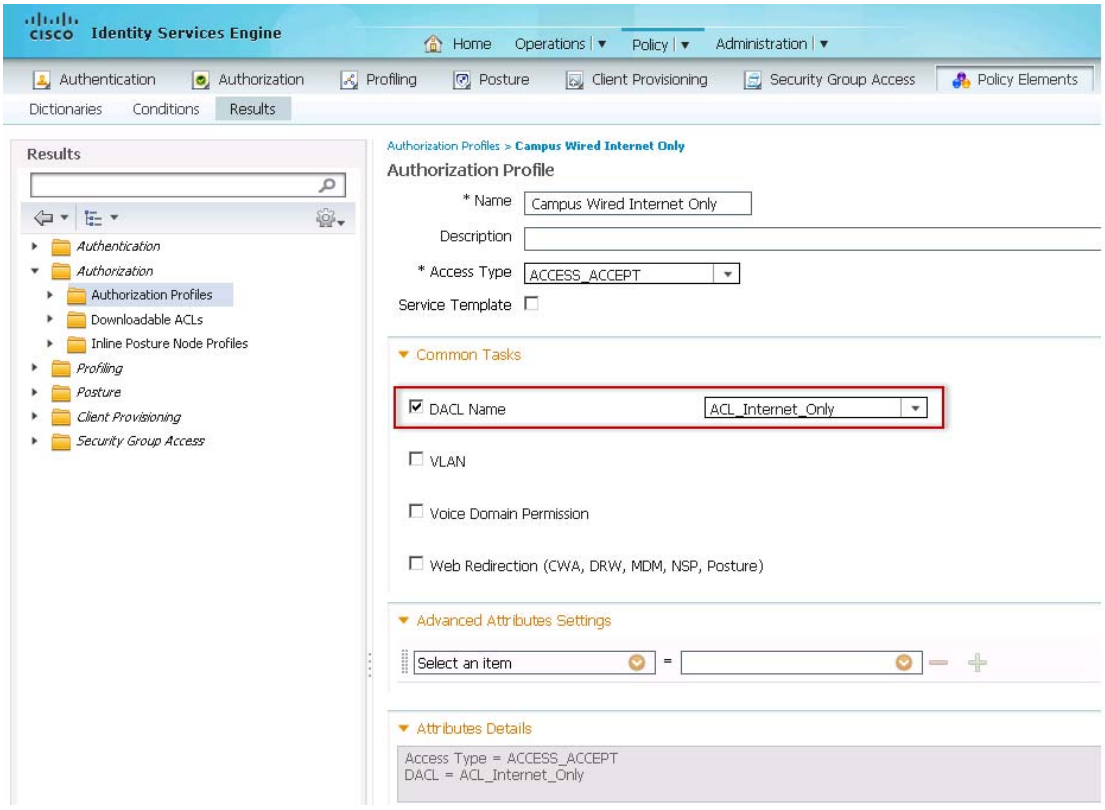
Table 15-11 Permissions Used for Wired Internet Access

Permission name	Permission type	Purpose
Campus Wired Internet Only	Authorization profile	To push a DACL for 802.1X wired devices connecting from a campus location.
Branch Wired Internet Only	Authorization profile	To push a VLAN for 802.1X wired devices connecting from a branch location.
Converged Wired Internet Only	Authorization profile	To push a named ACL for 802.1X wired devices connecting from either a campus or branch location with Converged Access

Wired Campus

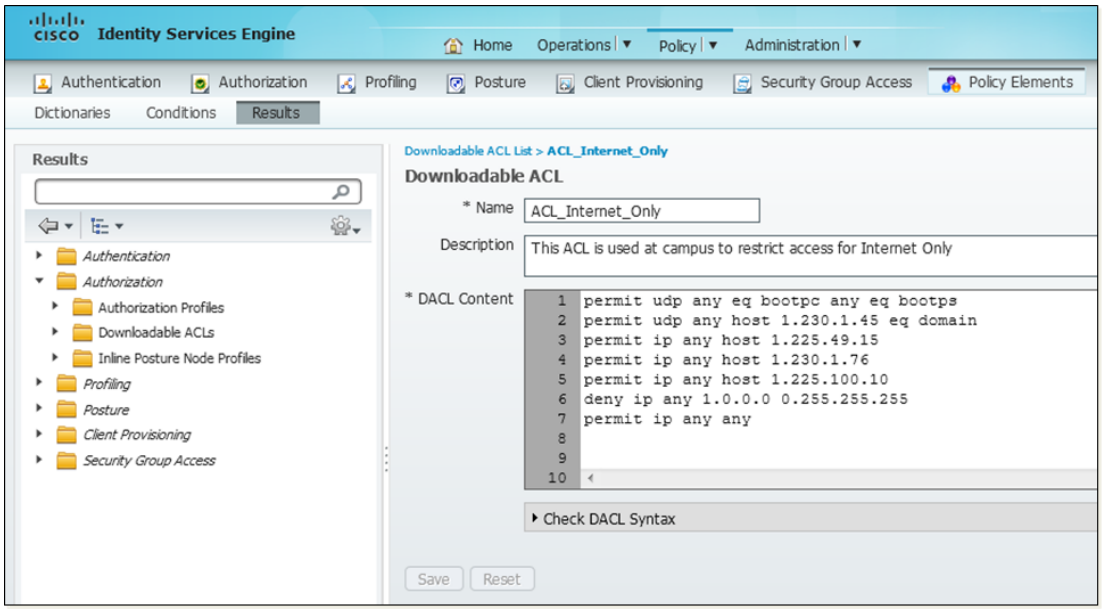
For devices connecting from a campus location, the Campus Wired Internet Only authorization profile uses the DACL named ACL_Internet_Only, which is pushed to the access layer switch port. Figure 15-59 shows the Authorization profile.

Figure 15-59 Campus Wired Internet Only Authorization Profile



The DACL overrides the default-ACL configured on the switch. Figure 15-60 shows an example of this ACL, which is configured within ISE.

Figure 15-60 ACL_Internet_Only DACL



The access list specifies the following access:

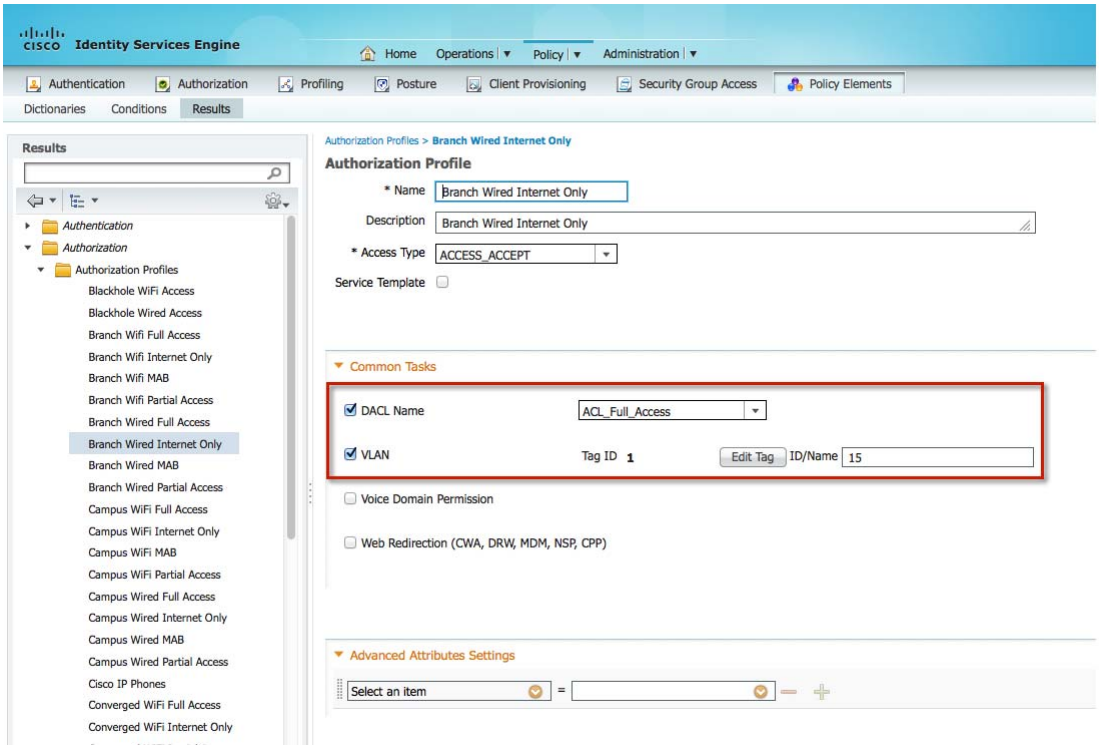
- Allow DHCP access (bootpc and bootps).
- Allow DNS access to the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Deny IP access to and from internal network address space (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).
- Allow access to and from all other subnets (Internet access).

This access list is generic and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

Branch Wired

For devices connecting from a branch location, the Branch Wired Internet Only authorization pushes a VLAN assignment along with a downloadable ACL. Figure 15-61 illustrates the Branch Wired Internet_Only authorization profile, used to grant Internet_Only access to personal devices connecting from the branch.

Figure 15-61 Branch Wired Internet Only Authorization Profile



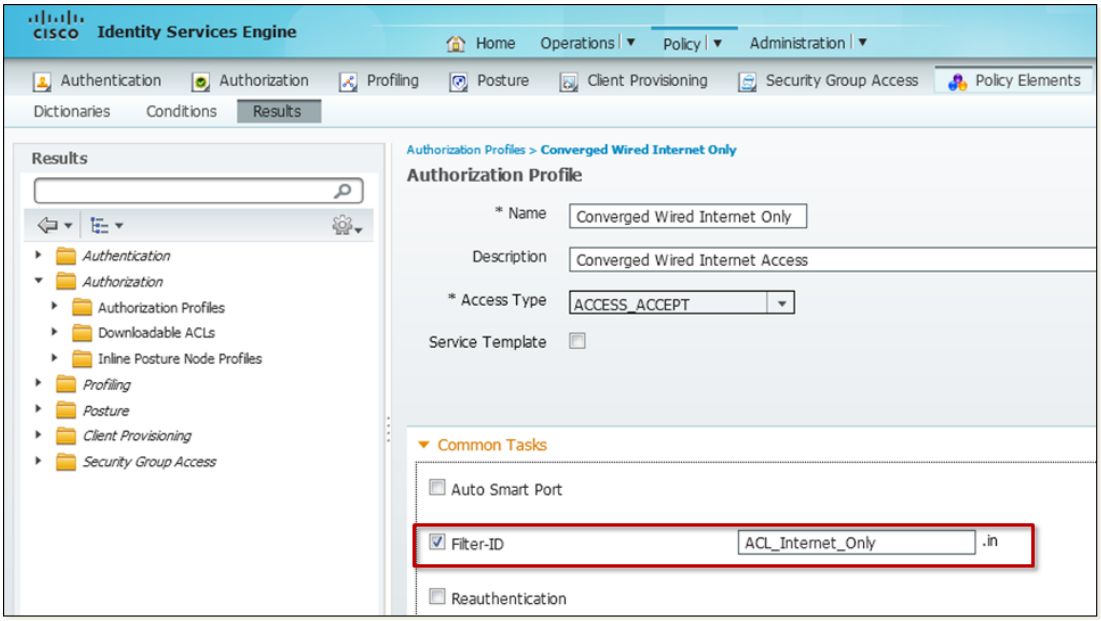
The ACL_Full_Access DACL is pushed to the access-layer switch to override the default-ACL on the port and allow all IP traffic to flow up to the branch router.

See VLAN Design at Branch Locations in Chapter 11, "BYOD Wired Infrastructure Design" for details regarding how this assignment of VLAN 15 is given full access for a branch which does not implement a Converged Access infrastructure.

Converged Access Branch and Campus

Figure 15-62 shows how the Converge Wired Internet Only authorization profile is defined in ISE.

Figure 15-62 Converged Wired Internet Only Authorization Profile



For the Converged Access design a named ACL is implemented. Using the RADIUS Filter-ID attribute-value pair, ISE instructs the converged access switch to apply the ACL_Internet_Only ACL. This is the same ACL discussed for converged access infrastructure in Personal Wireless Devices—Internet Only Access. For wired devices, the ACL is used to override the default-ACL configured on the switch port. The default-ACL is used as an additional preventative measure in case ACL_Internet_Only is not configured on the switch.

Android Devices—Deny Access

Rather than allowing differentiated access to the network which was depicted in the previous use cases, this use case discusses on how to deny access permission for some BYOD devices from connecting to the network. For example, some organizations may decide to have a more restrictive BYOD environment and grant access only to a specific type of device (e.g., Android, Apple iOS, etc.).

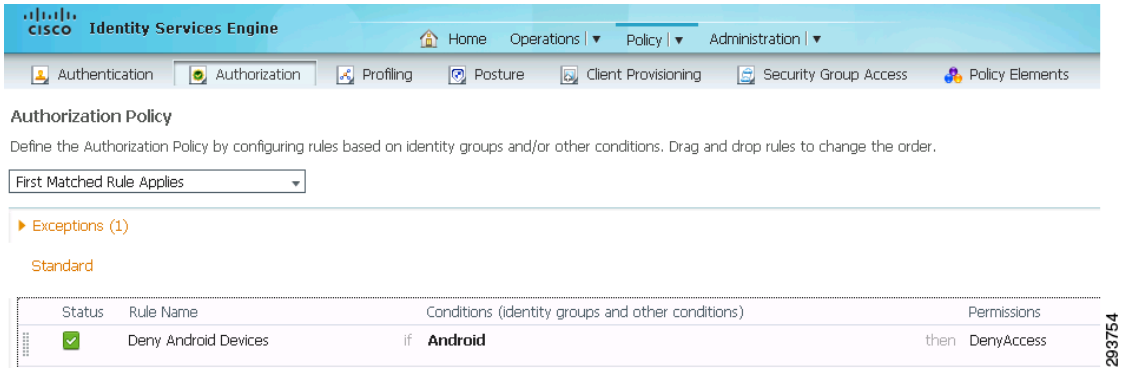
This example focuses on denying access to Android devices, relying on the profiling capabilities of ISE.

To deny access to Android devices, the Cisco ISE verifies the following:

- The employee attempts to connect to the network.
- The ISE profiler identifies the device type.
- If the device type is Android, deny access.

To configure the authorization rules in ISE, click **Policy > Authorization**. Figure 15-63 highlights the authorization policy to deny access to Android devices.

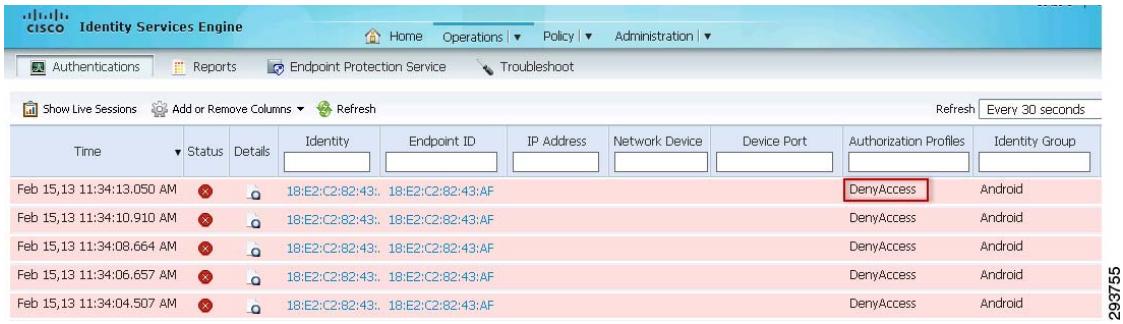
Figure 15-63 Deny Android Devices



The DenyAccess authorization profile is used to enforce the permissions and deny access to Android devices. The DenyAccess profile is a standard ISE profile and cannot be edited. This reserved profile cannot be edited but may be found under **Policy > Results> Authorization Profiles**.

Figure 15-64 shows an entry from ISE's log, highlighting the fact that the device has been profiled as an Android device and the DenyAccess authorization rule has been enforced.

Figure 15-64 DenyAccess

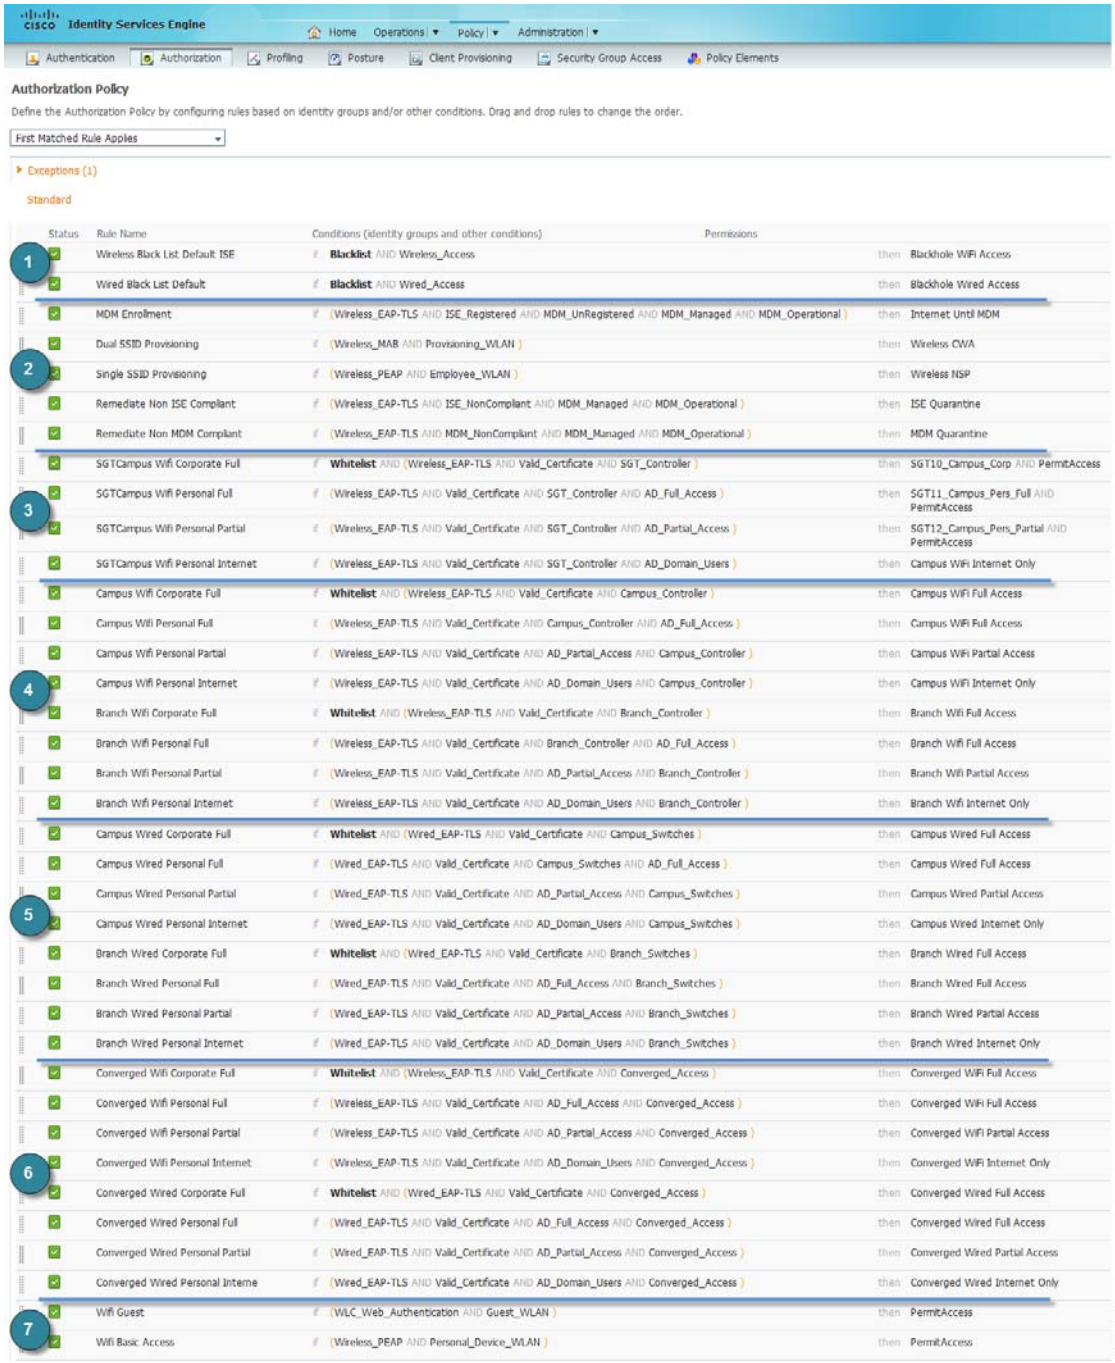


ISE Authorization Policy

For reference purposes, the complete authorization policy used during validation is shown in Figure 15-65. The figure highlights the following sections:

1. Used for blacklisting lost or stolen devices.
2. On-boarding and MDM registration/remediation (required for the advanced use case).
3. Wireless devices connecting from an access point in a SGT_Enabled location.
4. Wireless devices connecting from an access point in the campus or branch locations.
5. Wired devices connecting from a campus or branch locations.
6. Wired and Wireless devices connecting from a converged location.
7. Guest and Basic Access.

Figure 15-65 Complete Authorization Policy





BYOD Limited Use Case—Corporate Devices

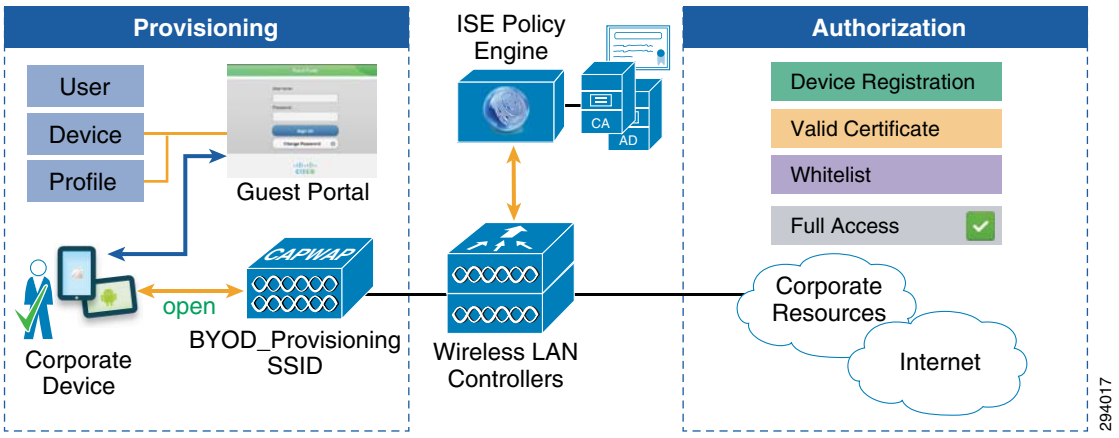
Revised: August 7, 2013

This chapter discusses design considerations and the construction of policy rules for providing network access to corporate devices, as well as the security policies enforced by the Cisco ISE. A corporate device is a corporate asset that is provided by the organization and has the approved configuration profiles and digital certificate to access the network. ISE allows the employee to on-board their corporate devices in a similar way personal devices are on-boarded, as discussed in Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices."

Cisco ISE provides many ways to enforce security policies and determines what network resources each device is allowed to access. This chapter focuses on allowing full access to corporate devices. The ISE feature set is extremely flexible to meet diverse business requirements. The goal of this chapter is to highlight this flexibility of ISE and explain the steps to restrict access for corporate devices.

Figure 16-1 shows the different authorization rules used by ISE to grant full access to corporate devices. Different network components play a role in this process, including the wireless infrastructure, Active Directory, a CA server, etc.

Figure 16-1 Provisioning Corporate Devices



Whitelist Identity Group

An identity group is a logical list used to group endpoints according to their profiles and is an efficient way for ISE to enforce different permissions to different types of devices.

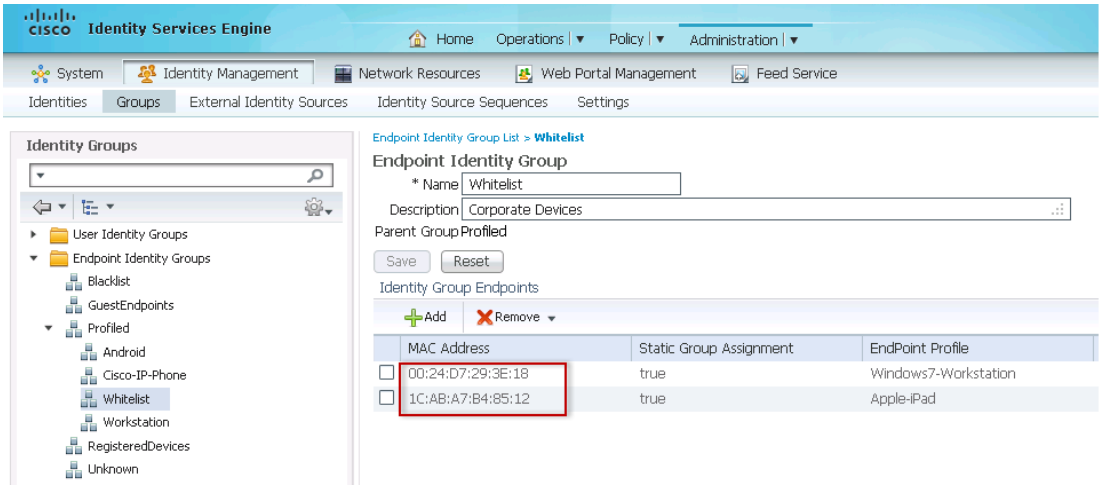
For this design guide, the Whitelist identity group was created for the purpose of uniquely identifying corporate devices. This identity group maintains a list of devices owned by the corporation. The Whitelist is manually updated by the IT administrator and contains the MAC addresses of devices that are granted full access. The assumption is that IT has added the devices to the Whitelist in advance.

The identity group provides an additional check when enforcing authorization rules. Endpoints may be moved to other identity groups, such as the Whitelist identity group or the Blacklist identity group, used when a device is lost or stolen. Chapter 22, "Managing a Lost or Stolen Device" has more details on the Blacklist identity group.

Note Endpoints can only be members of one identity group at a time.

To update an endpoint's identity group, click **Administration > Groups > Endpoint Identity Groups**. Figure 16-2 shows endpoints as members of the Whitelist identity group.

Figure 16-2 Whitelist Identity Group



An Active Directory group could be used as an additional check, but for the purpose of this design guide the Whitelist identity group identifies a device as a corporate device. This model could easily be expanded to include AD groups to enforce additional policy requirements.

Figure 16-3 highlights the policy tested in this chapter, along with the different requirements and permissions. This policies, along with detailed configurations, is explained in this chapter.

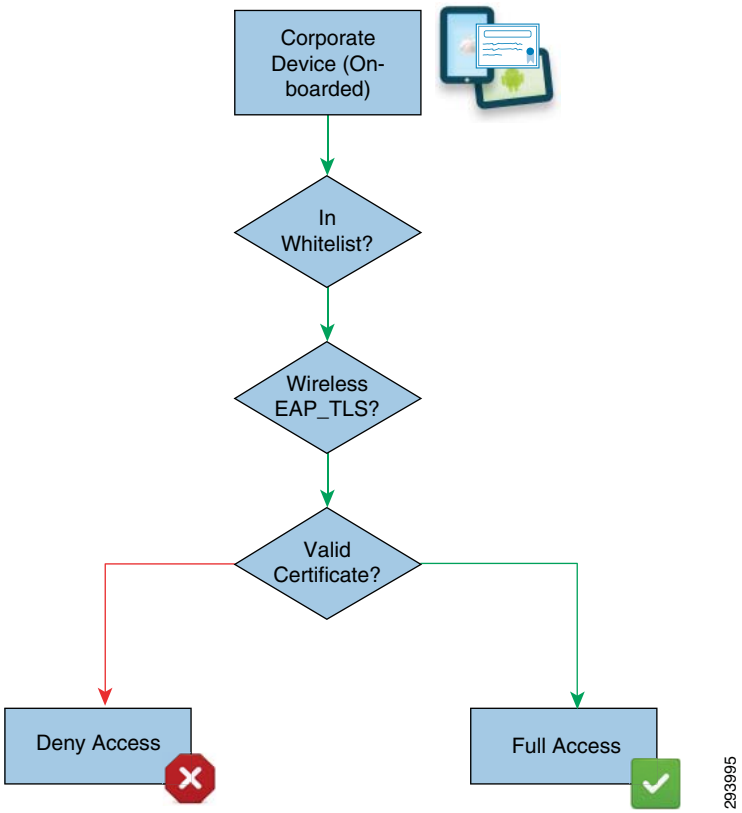
Figure 16-3 Access Policies and Permissions

Policy	Identity Group	Location	Permission
Corporate Owned	Whitelist	Campus/Branch/SGT	Full

294019

This chapter assumes that after employees have on-boarded their devices, they'll connect to the BYOD_Employee. Figure 16-4 highlights the connectivity flow granting full access to corporate devices. If the device has been on-boarded, it belongs to the Whitelist identity group and has a digital certificate, the device is granted full access.

Figure 16-4 Corporate Device BYOD Access



Policy Enforcement for Security Group Access

A new addition to this CVD is the use of Security Group Tags in creating role-based policies in addition to ACLs to enforce policies for campus wireless users. Enforcing policies for SGA in BYOD architecture consist of three main components:

- Defining tags for the endpoints and the destination servers.
- Defining and implementing the Security Group ACL in the switching infrastructure.
- Defining firewall policies based on Security Group Tags if Security Group Firewall is implemented.

Depending on the SGA deployment scenario followed, the Security Group Egress Policy can be defined at the ISE and dynamically pushed to Catalyst 6500 infrastructure switches and the Nexus Data Center switches as Security Group ACLs (SGACL) or configured at the ASA Firewall as a SGT-based policy. These two choices are discussed in Chapter 23, "BYOD Policy Enforcement Using Security Group Access" as two different deployment scenarios.

Security Group Access Tags

The basic idea of Security Group Access is to define the tags for source and destination traffic flows and specify what source tags are able to reach other destination tags. For example, are devices tagged with an SGT 10 allowed to communicate with a server tagged with an SGT 40? This flow is either allowed or blocked at the enforcement point.

Using this tagging concept, unique tags have been defined for different types of source traffic generated by the endpoints. As explained in Chapter 12, "Security Group Access for BYOD," unique permissions are established for personal and corporate devices and unique tags have been defined for each use case.

Table 16-1 illustrates how tags are assigned to different devices.

Table 16-1 Source Tags

Device Type	Tag
Corporate device with Full Access	SGT 10
Personal device with Full Access	SGT 11
Personal device with Partial Access	SGT 12

Similarly, the destination servers also need to be associated with a particular tag. Table 16-2 illustrates how, based on their role, servers are assigned with a different tag.

Table 16-2 Destination Tags

Destination Server	Tag
Open Access	SGT 40
Corporate Server	SGT 50

Security Group ACL

After defining the source and destination tags, the next logical step is to define the egress policy that establishes the permissions for traffic between those tags. Table 16-3 illustrates the egress policy as an enforcement matrix.

Table 16-3 Enforcement Matrix

	Device Type	SGT 40	SGT 50
SGT 10	Corporate device with Full Access	Yes	Yes
SGT 11	Personal device with Full Access	Yes	Yes
SGT 12	Personal device with Partial Access	Yes	No

As shown in Table 16-3, corporate or personal devices granted full access will be allowed to reach servers that have been tagged with SGT 40 or SGT 50. Similarly, a personal device with a partial access is only allowed to connect with a server tagged with SGT 40.

The implementation of the SGACL depends upon the type of the deployment scenario. For the deployment scenario where Nexus is the enforcement point, the SGACL is designed in ISE and pushed to the Nexus switch. When the deployment scenario is using ASA as the enforcement point, then an SGT-based access rule is configured manually on the ASA firewall.

Authorization Polices for SGT

Based on policy matches and authorization profiles, the ISE assigns an SGT to campus wireless devices. Centralized Campus—Policy Enforcement using TrustSec in Chapter 9, "BYOD Wireless Infrastructure Design" provides information about the criteria used to determine when, based on the network device type of the wireless controller defined in ISE, an ACL versus an SGT should be returned upon successful authorization.

For the destination server tags, the configuration is done manually at the data center switch and the actual enforcement happens based on the deployment scenario. For more information about configuring tags for servers and the ASA, see Chapter 23, "BYOD Policy Enforcement Using Security Group Access.".

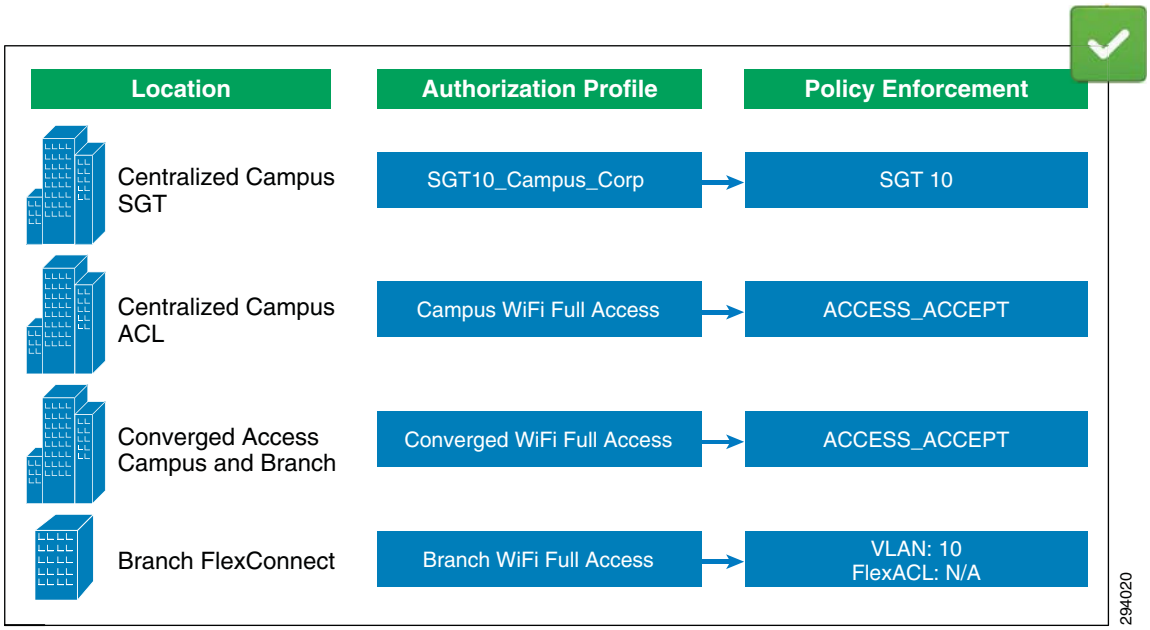
Corporate Devices—Full Access

To provide full access to corporate devices, the Cisco ISE verifies the following:

- The device has been on-boarded and has the right certificate and profile configurations.
- The device has been added to the Whitelist identity group.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate.
- The connection originated using EAP-TLS authentication.

Since the wireless designs presented in this design guide rely on different WLCs for FlexConnect branches, Centralized Controller campuses, and Converged Access campuses and branches, unique authorization rules are created for connections originating from each design. At a high level, Figure 16-5 shows how different authorization profiles are selected for connections originating from different locations with different wireless designs. Each authorization profile in turn enforces a unique permission using VLANs, SGTs, dynamic ACLs (either named or downloadable [DACL]), FlexConnect ACLs, etc.

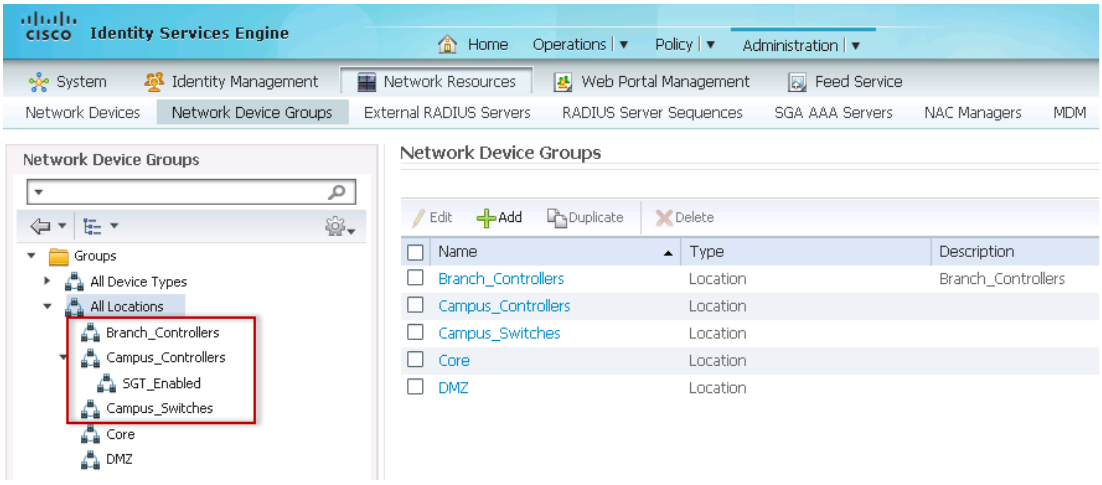
Figure 16-5 Full Access Enforcement



To differentiate these connections, the ISE relies on Network Device Groups to group WLCs based on their location or device type. This allows a single ISE to enforce policies across different groups of devices.

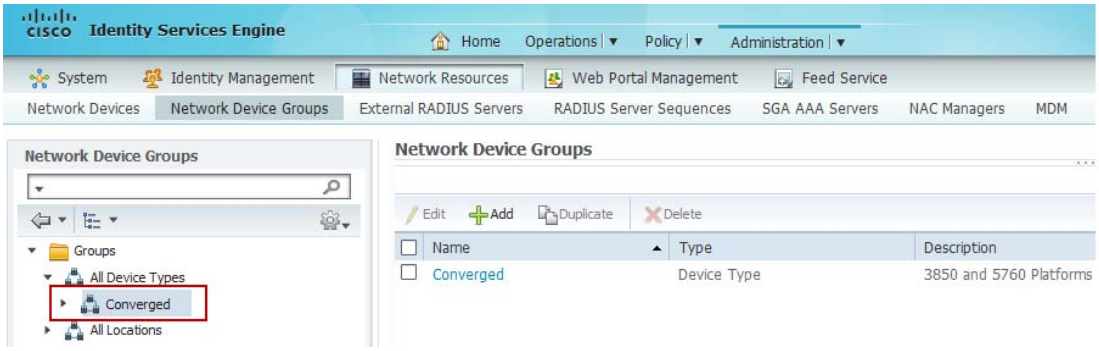
Figure 16-6 shows the different locations created for branch and campus devices.

Figure 16-6 ISE Device Groups—Locations



Similarly, Figure 16-7 shows a device type called Converged which can be created for Catalyst 3850 Series switches and CT5760 wireless controllers and used in the authorization policy.

Figure 16-7 ISE Device Groups—Device Types

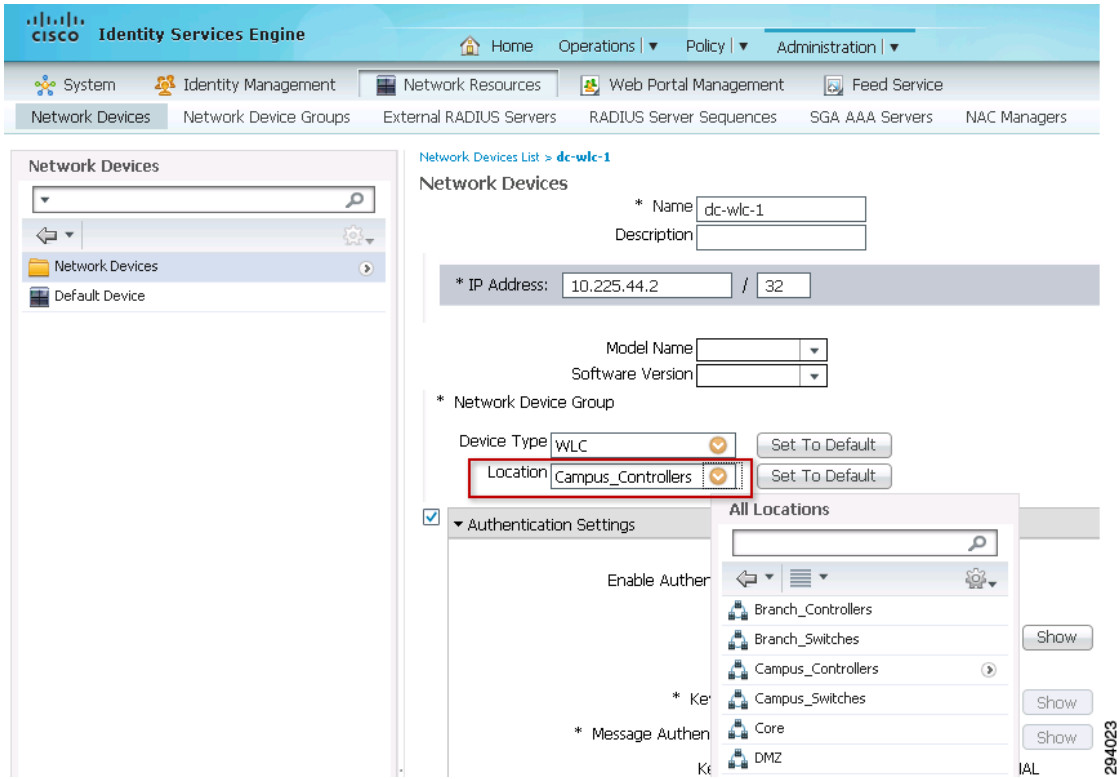


Note One of the reasons a device type is used instead of a location for Converged Access designs is that the same authorization policy rules are used for Converged Access branch and campus designs within this design guide. Hence location—campus versus branch—is not particularly relevant from a Cisco ISE perspective to converged access designs presented in this design guide.

Each Wireless LAN Controller needs to be added to the proper device group by clicking **Administration > Network Resources > Network Devices** and specifying the proper location or device type from the pull-down menu.

Figure 16-8 shows how the dc-wlc-1 belongs to the Campus_Controllers location.

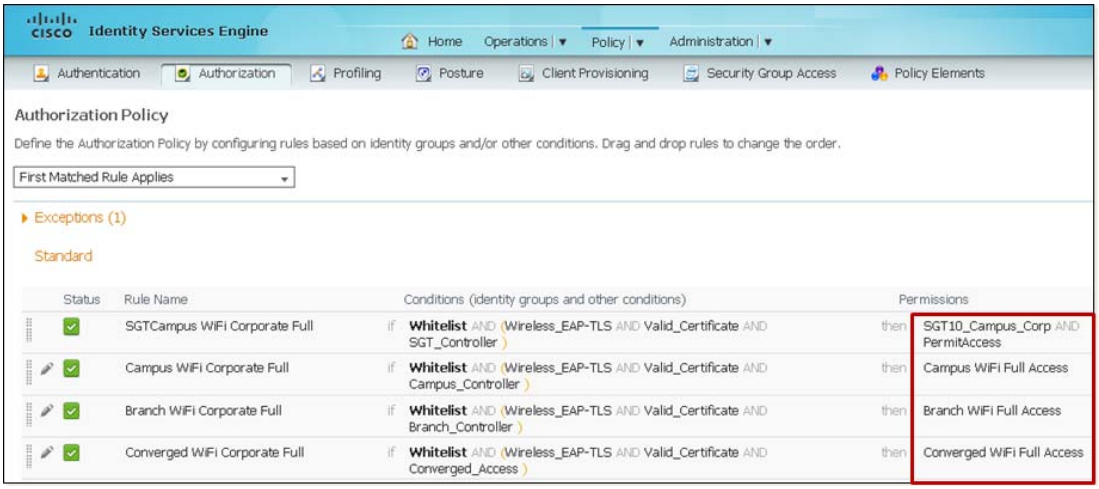
Figure 16-8 Campus Controller



Wireless LAN Controllers implementing a centralized (Local Mode) campus wireless design are assigned to the Campus_Controllers group.

To configure the authorization rules in ISE, click **Policy > Authorization**. Figure 16-9 highlights the authorization policy to grant full access to corporate devices.

Figure 16-9 Authorization Policies for Full Access



Looking at the first rule in more detail, ISE evaluates these conditions:

- Whitelist—The endpoint has been added by an IT Administrator to the Whitelist identity group.
- Wireless_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name (defined as a simple condition).
- The RADIUS authentication originated from a wireless controller which was a member of one of the following device groups: Campus_Controller, SGT_Controller, Branch_Controller, or Converged_Access (defined as a simple condition).

Note A wireless controller which is a member of the Converged_Access device group could be either a standalone device, such as a Cisco CT5760 wireless controller, or a switch with integrated wireless controller functionality, such as the Catalyst 3850.

Simple and Compound Conditions

To improve the readability of the authorization policy, simple and compound authorization conditions were defined to group different conditions. These conditions may be reused and modified without changing every authorization rule.

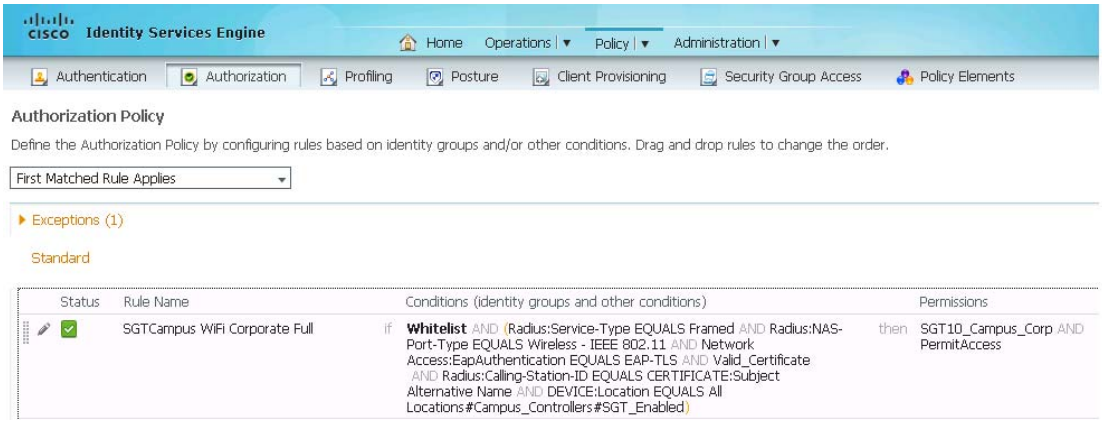
Table 16-4 shows the conditions used in the authorization rules.

Table 16-4 Wireless Simple and Compound Conditions

Wireless EAP-TLS (Compound)	
Wireless_EAP-TLS (See Figure 16-11)	Radius:Service-Type Equals Framed Radius:NAS-Port-Type Equals Wireless - IEEE 802.11 Network Access:EapAuthentication Equals EAP-TLS
Check for Valid Certificate (Simple)	
Valid_Certificate (See Figure 16-12)	Radius:Calling-Station-ID EQUALS CERTIFICATE:Subject Alternative Name
WLC Location or Device Type (Simple)	
SGT_Controller	DEVICE:Location EQUALS All Locations#Campus_Controllers#SGT_Enabled
Campus_Controller	DEVICE:Location EQUALS All Locations#Campus_Controllers
Branch_Controller	DEVICE:Location EQUALS All Locations#Branch_Controllers
Converged_Access	DEVICE:Device Type EQUALS All Device Types#Converged

To illustrate the value of using simple/compound conditions, Figure 16-10 shows how much longer and harder to read a rule can be when simple/compound conditions are not implemented.

Figure 16-10 Authorization Rule without Conditions



To define a new compound condition, click **Policy > Conditions > Authorization > Compound Conditions**. Figure 16-11 shows how the Wireless_EAP-TLS condition combines several conditions into one.

Figure 16-11 Wireless_EAP-TLS Condition

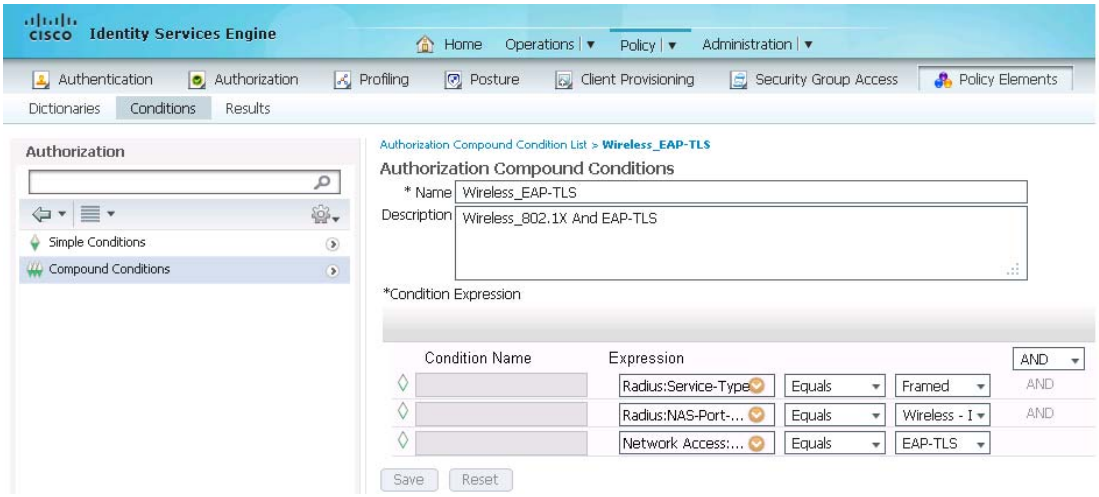
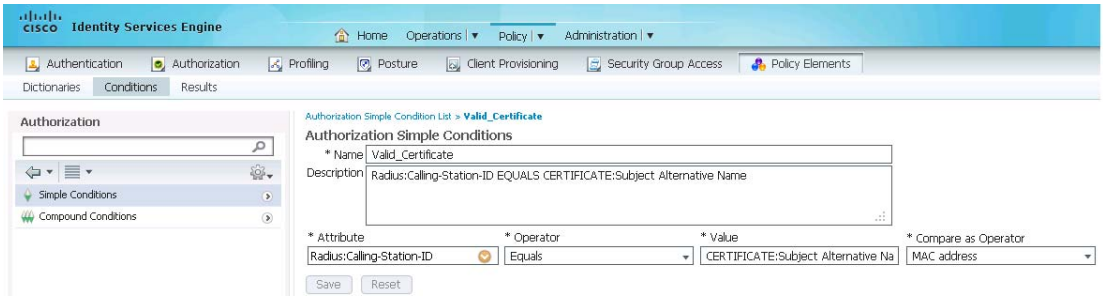


Figure 16-12 shows the Valid_Certificate simple condition.

Figure 16-12 Valid_Certificate Condition



The remaining conditions shown in Table 16-4 are defined in a similar way.

Permissions for Wireless Users

When all conditions in the authorization policy rule match, the rule invokes the proper permissions to provide Full Access, as shown in Table 16-5.

Table 16-5 Permissions for Full Access

Permission name	Permission type	Purpose
SGT10_Campus_Corp	Standard result	To assign a SGT for 802.1X wireless devices connecting from an SGT-enabled controller.
Campus WiFi Full Access	Authorization profile	Provides Full Access for 802.1X wireless devices connecting from a centralized campus controller.

Table 16-5 Permissions for Full Access

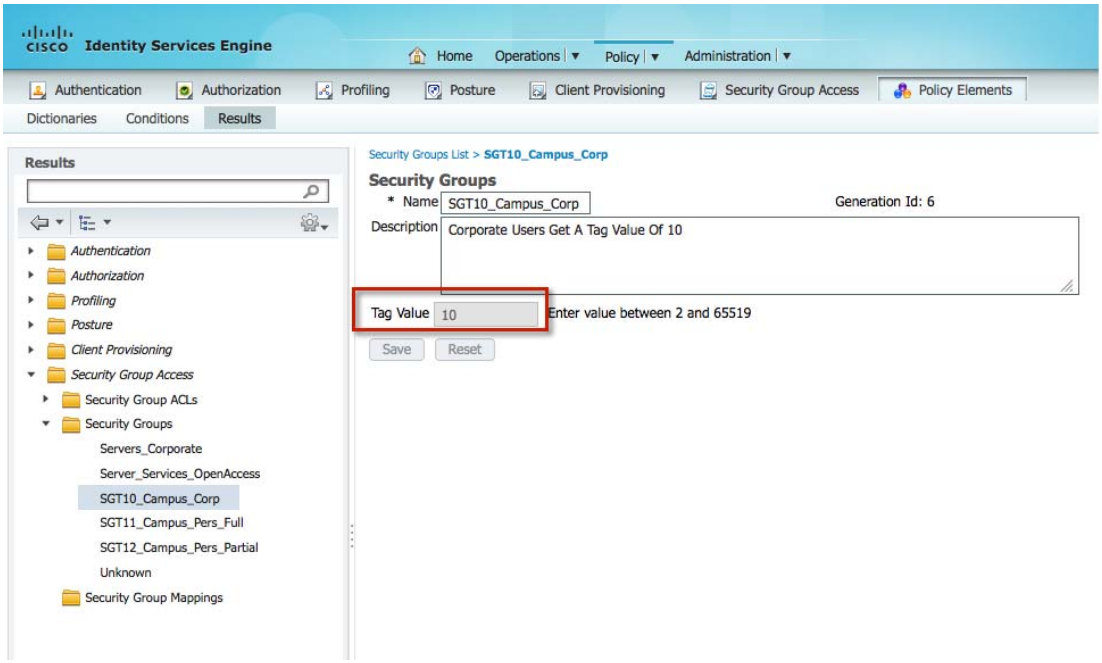
Permission name	Permission type	Purpose
Branch WiFi Full Access	Authorization profile	To push a VLAN for 802.1X wireless devices connecting from a FlexConnect branch controller.
Converged WiFi Full Access	Authorization profile	Provides Full Access for 802.1X wireless devices connecting from a Converged Access controller.

Note A Converged Access infrastructure refers to a branch or campus deployment with Catalyst 3850 Series switches and/or CT5760 wireless controllers within this design guide.

Centralized Campus with SGTs

Figure 16-13 shows how the SGT10_Campus_Corp authorization profile is pushing the Security Group Tag whose value is 10 for a user of a corporate device, who is allowed to obtain full access.

Figure 16-13 SGT10_Campus_Corp

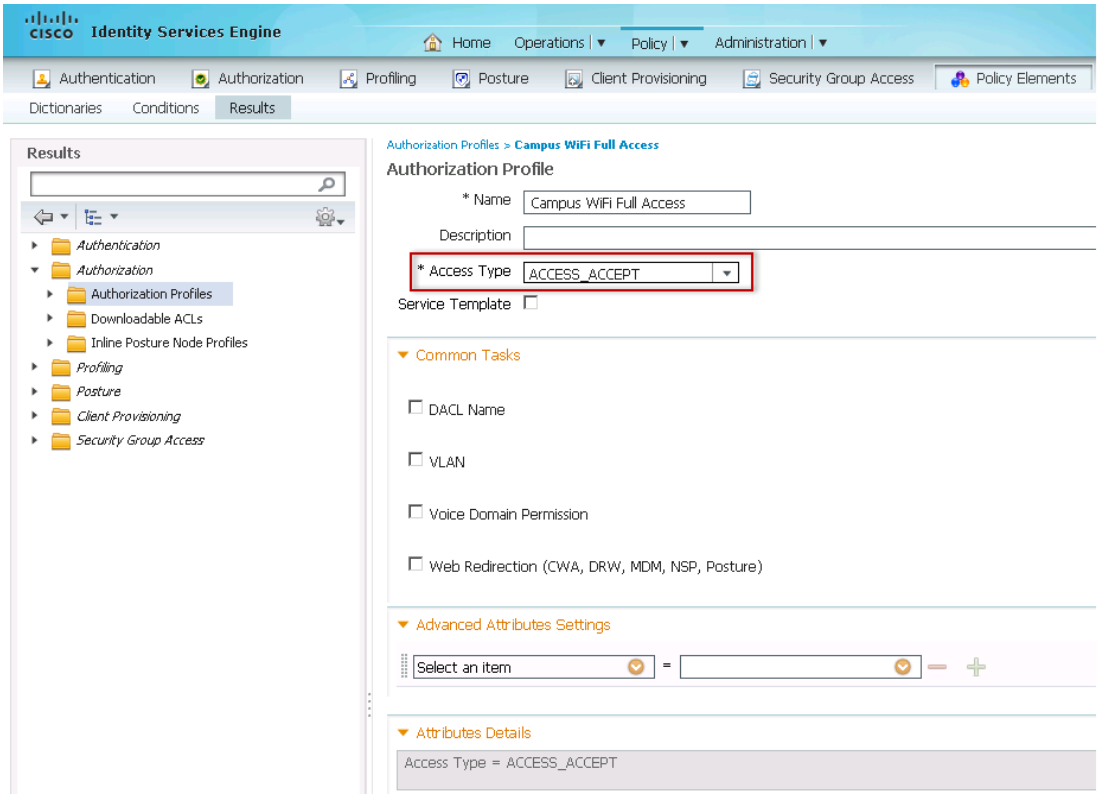


See Chapter 23, "BYOD Policy Enforcement Using Security Group Access" for details regarding how SGT 10 is given full access for a campus which implements a centralized controller (Local Mode) design with SGTs.

Centralized Campus with ACLs

Figure 16-14 shows how the Campus WiFi Full Access authorization profile is using the ACCESS_ACCEPT Access Type to allow full access.

Figure 16-14 Campus WiFi Full Access

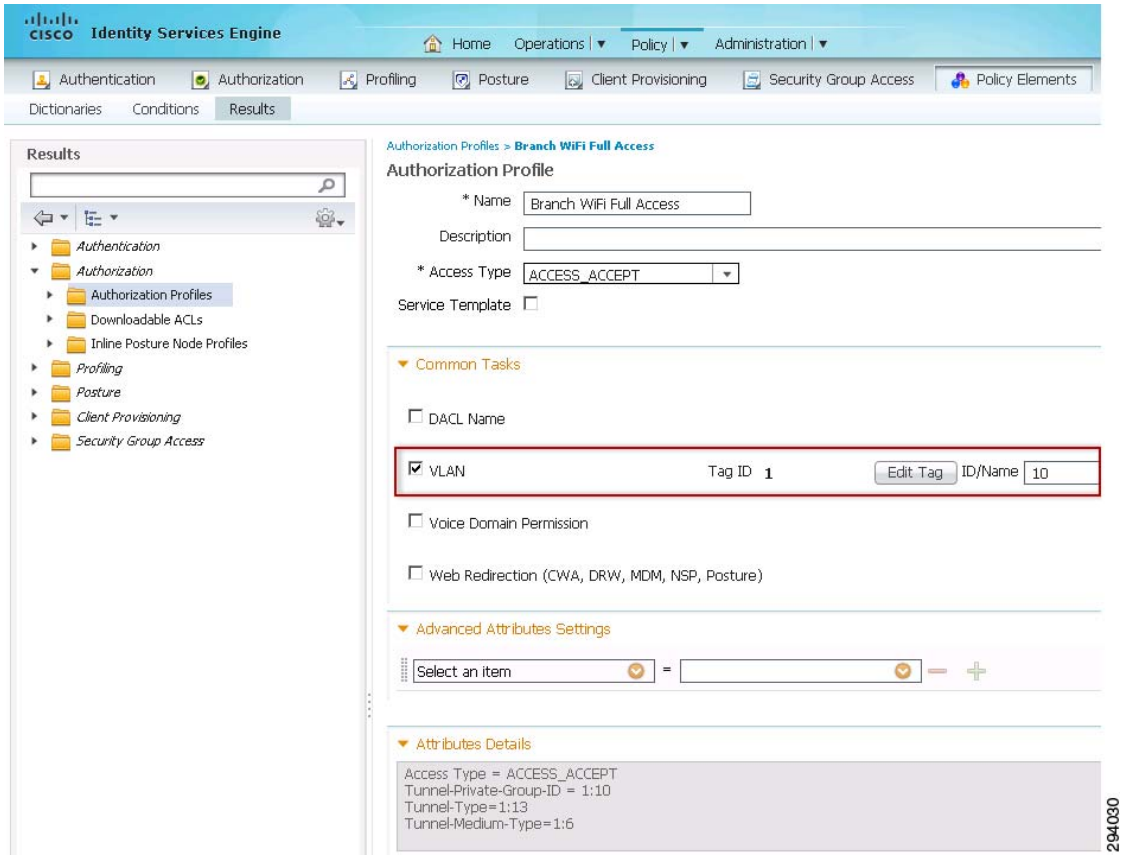


Since full access is allowed with this authorization profile, no Named ACL for access control needs to be specified by ISE.

Branch with FlexConnect

Endpoints connecting from a branch location dynamically get assigned to VLAN 10, which has been configured without an ACLs thus providing full access.

Figure 16-15 Branch WiFi Full Access

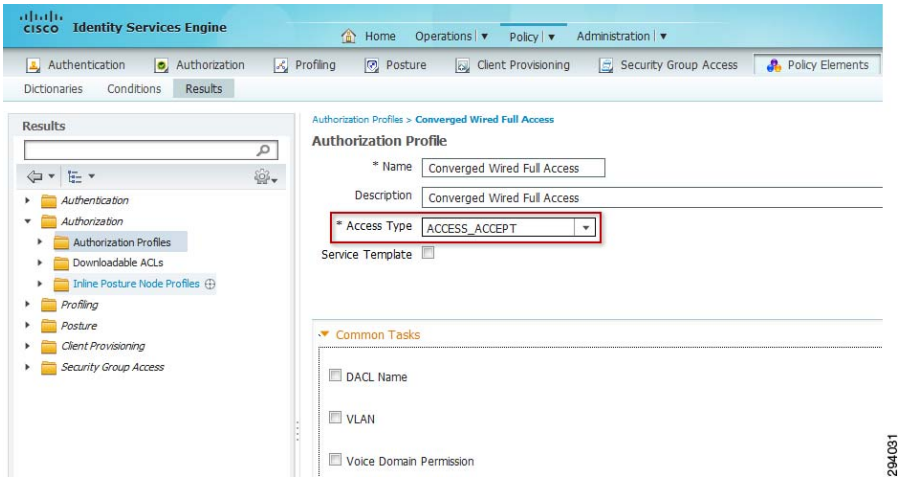


Converged Access Branch or Campus

Within this design guide, a Converged Access infrastructure refers to a branch or campus deployment with Catalyst 3850 Series switches and/or CT5760 wireless controllers.

Figure 16-16 shows how the Converged WiFi Full Access authorization profile is using the ACCESS_ACCEPT Access Type to allow full access.

Figure 16-16 Converged WiFi Full Access



Again, since full access is allowed with this authorization profile, no Named ACL for access control needs to be specified by ISE.

Corporate Wired Devices—Full Access

The authorization policy rules for wired devices follow the same logic as wireless devices. Corporate approved wired devices are assumed to be pre-configured with the correct configuration profiles after being on-boarded.

To provide full access to corporate wired devices, the Cisco ISE verifies the following:

- The device has been on-boarded and has the right certificate and profile configurations.
- The device has been added to the Whitelist identity group.
- To uniquely identify the device and prevent spoofing, the Calling-Station-ID matches the Subject Alternative Name of the certificate, in this case, the MAC address of the endpoint.
- The connection originated using EAP-TLS authentication.

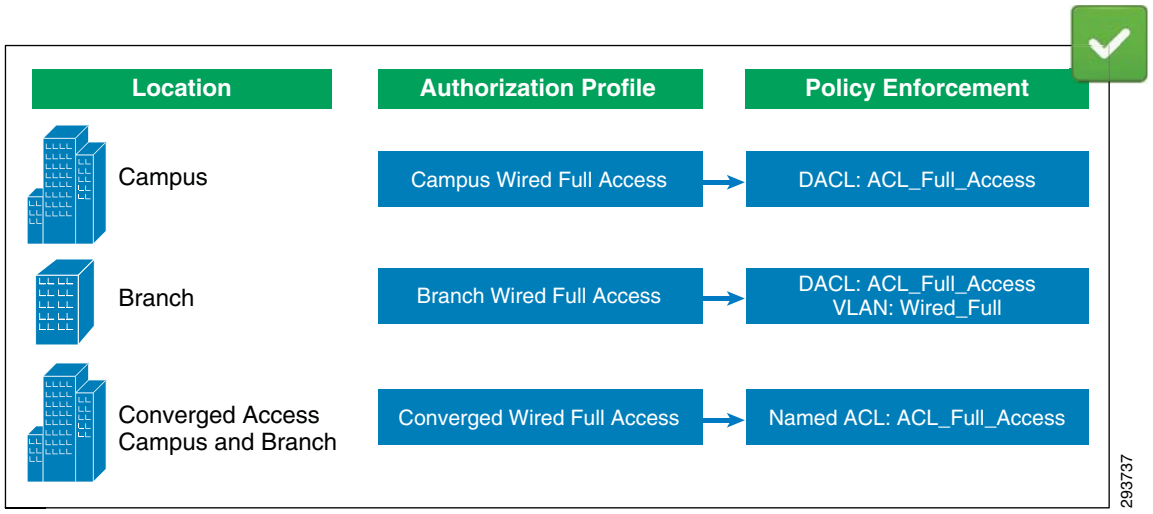
Since the wired designs presented in this design guide rely on slightly different access control mechanisms for converged access campuses and branches, for campuses and branches that do not implement converged access infrastructures unique authorization rules are created for connections originating from each design.

Note For the purpose of clarity within this design guide, converged access branches refer to designs in which Catalyst 3850 Series switches are deployed at the access-layer of the branch network. From a wired perspective, branches which do not implement converged access infrastructures are branches which deploy other Catalyst access-layer switches, such as the Catalyst 3750X Series. Unless otherwise specified, these are referred to simply as “branches” within the design guide. Similarly, converged access campuses refer to designs in which Catalyst 3850 Series switches are deployed at the access-layer of building distribution modules within the campus. From a wired perspective, campuses which do not implement converged access infrastructures are campuses which deploy other Catalyst access-layer switches, such as the Catalyst 3750X Series. Unless otherwise specified, these are referred to simply as

"campuses" within this design guide. This is in order to reduce the use of verbose phrases such as "branches which do not implement converged access infrastructure" and "campuses which do not implement converged access infrastructure".

At a high level, Figure 16-17 shows how different authorization profiles are selected for connections originating from different locations with different infrastructure (converged access versus non-converged access) designs. Each authorization profile in turn enforces a unique permission using VLANs, dynamic ACLs (either named or downloadable [DACL]), etc.

Figure 16-17 Full Access Wired Enforcement

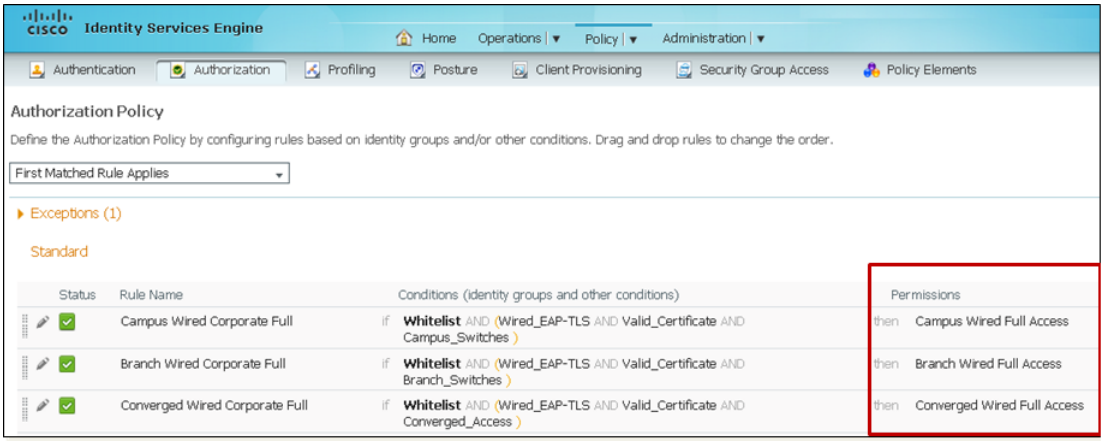


Note Wired assignment of Security Group Tags (SGTs) is not discussed within the design guide. Hence there is no wired policy enforcement via SGTs in Figure 16-17. Future versions of this design guide may address wired SGT assignment.

To differentiate these connections, the ISE relies on Network Device Groups to group Catalyst switches based on their location or device type. This allows a single ISE to enforce policies across different groups of devices. Each Catalyst switch needs to be added to the proper device group by clicking **Administration > Network Resources > Network Devices** and specifying the proper location or device type from the pull-down menu.

Figure 16-18 shows the details of authorization profile configured in ISE for wired devices.

Figure 16-18 Authorization Policies for Wired Full Access



Looking at the rules in more detail, ISE evaluates the following conditions:

- WhiteList—The endpoint has been added by an IT Administrator to the WhiteList identity group.
- Wired_EAP-TLS—The endpoint connected using EAP-TLS (defined as a compound condition).
- The endpoint has a valid certificate. The Calling-Station-ID matches the MAC address included in the certificate's Subject Alternative Name (defined as a simple condition).
- The RADIUS authentication originated from a Catalyst switch which was a member of one of the following device groups: Campus_Switches, Branch_Switches, or Converged_Access (defined as a simple condition).

Wired Simple and Compound Conditions

To improve the readability of the authorization policy, simple and compound authorization conditions were defined to group different conditions. These conditions may be reused and modified without changing every authorization rule.

Table 16-6 shows the conditions used in the authorization rules.

Table 16-6 Wired Simple and Compound Conditions

Wired EAP-TLS (Compound)	
Wired_EAP-TLS	Radius:Service-Type Equals Framed Radius:NAS-Port-Type Equals Ethernet Network Access:EapAuthentication Equals EAP-TLS
Check for Valid Certificate (Simple)	
Valid_Certificate	Radius:Calling-Station-ID EQUALS CERTIFICATE:Subject Alternative Name
WLC Location or Device Type (Simple)	
Campus_Switches	DEVICE:Location EQUALS All Locations#Campus_Switches
Branch_Switches	DEVICE:Location EQUALS All Locations#Branch_Switches
Converged_Access	DEVICE:Device Type EQUALS All Device Types#Converged

Permissions for Wired Users

When all conditions in the authorization policy rule match, the rule invokes the proper permission. The permissions can be of different forms, such as an authorization profile or a standard result. In this design guide, for wired access the authorization policy is used as permission to the policy rules match. Table 16-7 explains the permissions used for the full access for wired users.

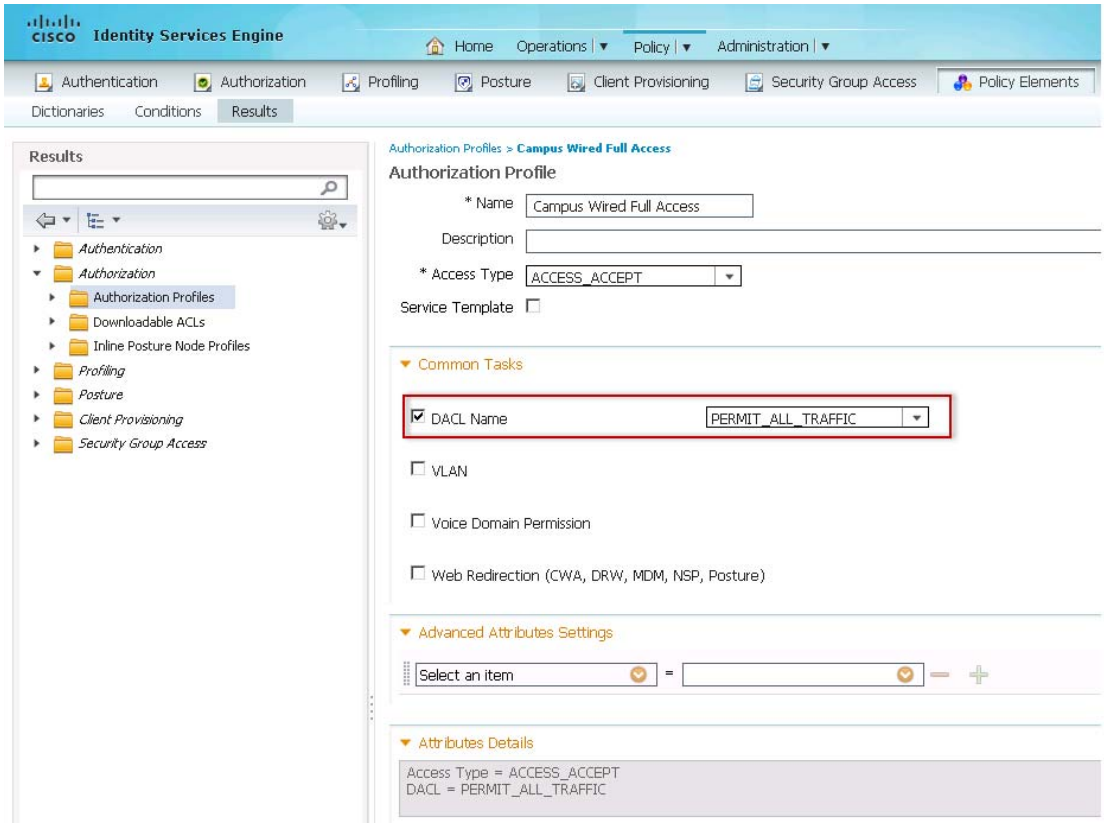
Table 16-7 Permissions Used for Wired Full Access

Permission Name	Permission Type	Purpose
Campus Wired Full Access	Authorization profile	Provides Full Access for 802.1X wired devices connecting from campus location.
Branch Wired Full Access	Authorization profile	To push a VLAN for 802.1X wired devices connecting from a branch location.
Converged Wired Full Access	Authorization profile	To push a named ACL for 802.1X wired devices connecting from either a campus or branch location with Converged Access.

Campus Wired

Figure 16-19 shows how the Campus Wired Full Access authorization profile is defined in ISE.

Figure 16-19 Campus Wired Full Access Authorization Profile



Note Cisco Catalyst switches support both downloadable ACLs (DACLs) and named ACLs. This design guide shows and validates the use of downloadable ACLs for access control of wired devices when implementing a non-converged access infrastructure and the use of named ACLs for access control of wired devices when implementing a converged access infrastructure. This is done to show the depth and range of capabilities for access control available on Cisco IOS-based platforms. Note that the same wired policy enforcement for converged access and non-converged access designs can be achieved if the customer desires by simply using either downloadable ACLs (DACLs) or named ACLs for both designs. Both downloadable and named ACLs have advantages and disadvantages, depending upon where they are deployed within the network. ACL Complexity and Considerations in Chapter 5, "Campus and Branch Network Design for BYOD" discusses some of these advantages and disadvantages.

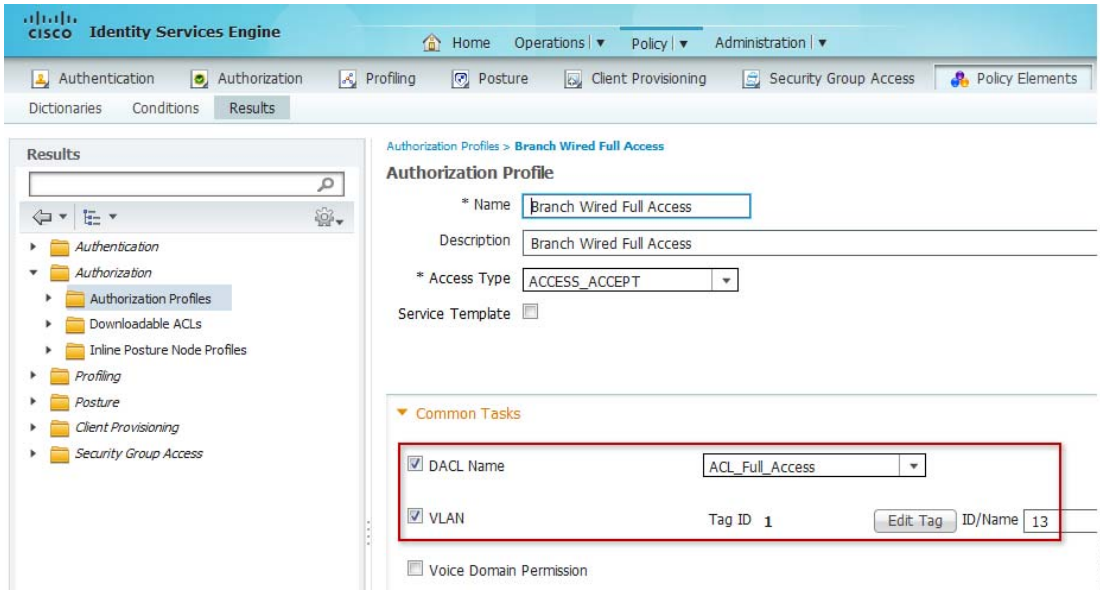
The downloadable ACL (DACL), which allows all IP traffic, overrides the default-ACL configured on the switch port. The default-ACL is used as an additional preventative measure in case the downloadable ACL is not applied to the switch port for some reason.

Note ISE 1.2 has the ability to check the DACL syntax. This feature should be utilized in order to minimize the possibility that the DACL is not applied due to a syntax error.

Branch Wired

Figure 16-20 shows how a similar authorization profile is constructed for devices connected at a branch which does not have a Converged Access infrastructure. The authorization profile pushes VLAN information along with the ACL information.

Figure 16-20 Branch Wired Full Access Authorization Profile

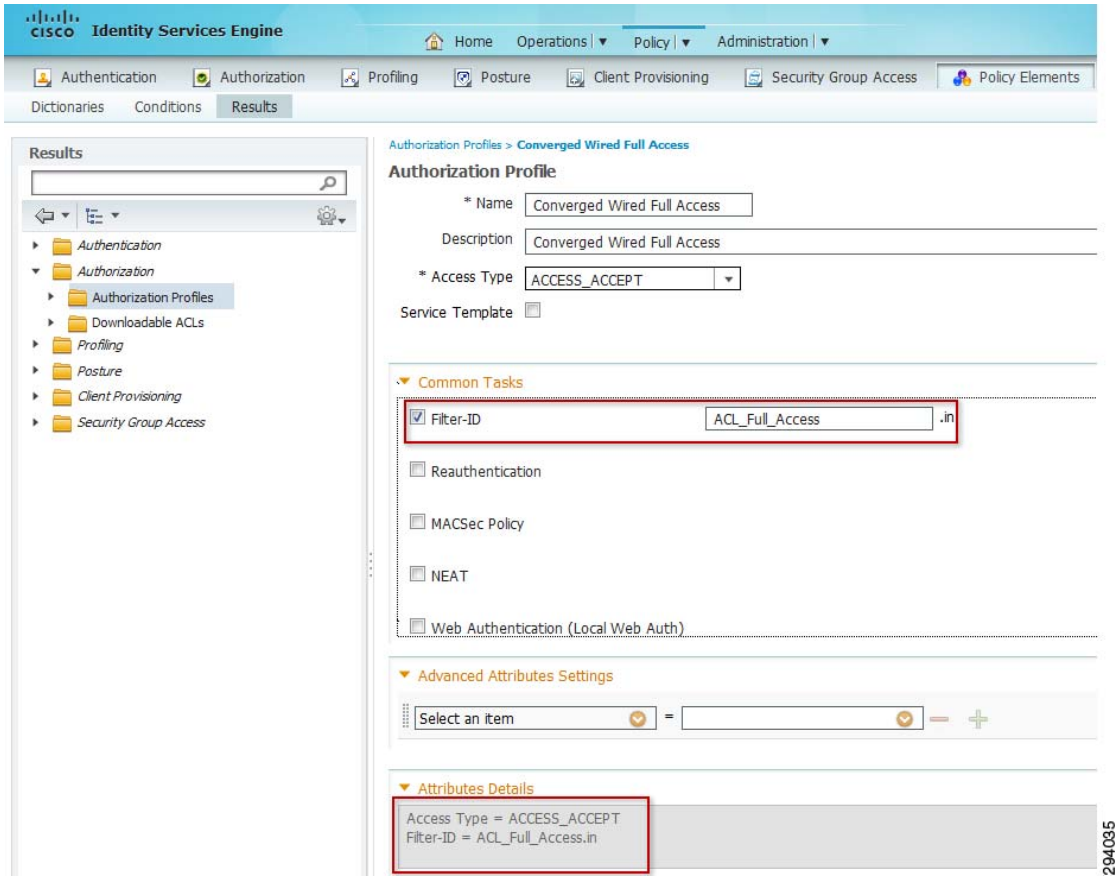


Once this profile is downloaded to the Catalyst switch, the endpoint gets full access to the network.

Converged Access Branch or Campus

Figure 16-21 shows how the Converged Access Wired Full Access authorization profile is defined in ISE.

Figure 16-21 Converged Wired Full Access Authorization Profile



For the converged access design, a named ACL is implemented, meaning that the ACL must be configured on the Catalyst 3850 switch rather than being downloaded from ISE. Using the RADIUS Filter-ID attribute-value pair, ISE instructs the Catalyst 3850 to apply the ACL_Full_Access ACL. The following configuration example shows the contents of this ACL as defined in the Catalyst 3850 switch.

```
!
ip access-list extended ACL_Full_Access
 permit ip any any
!
```

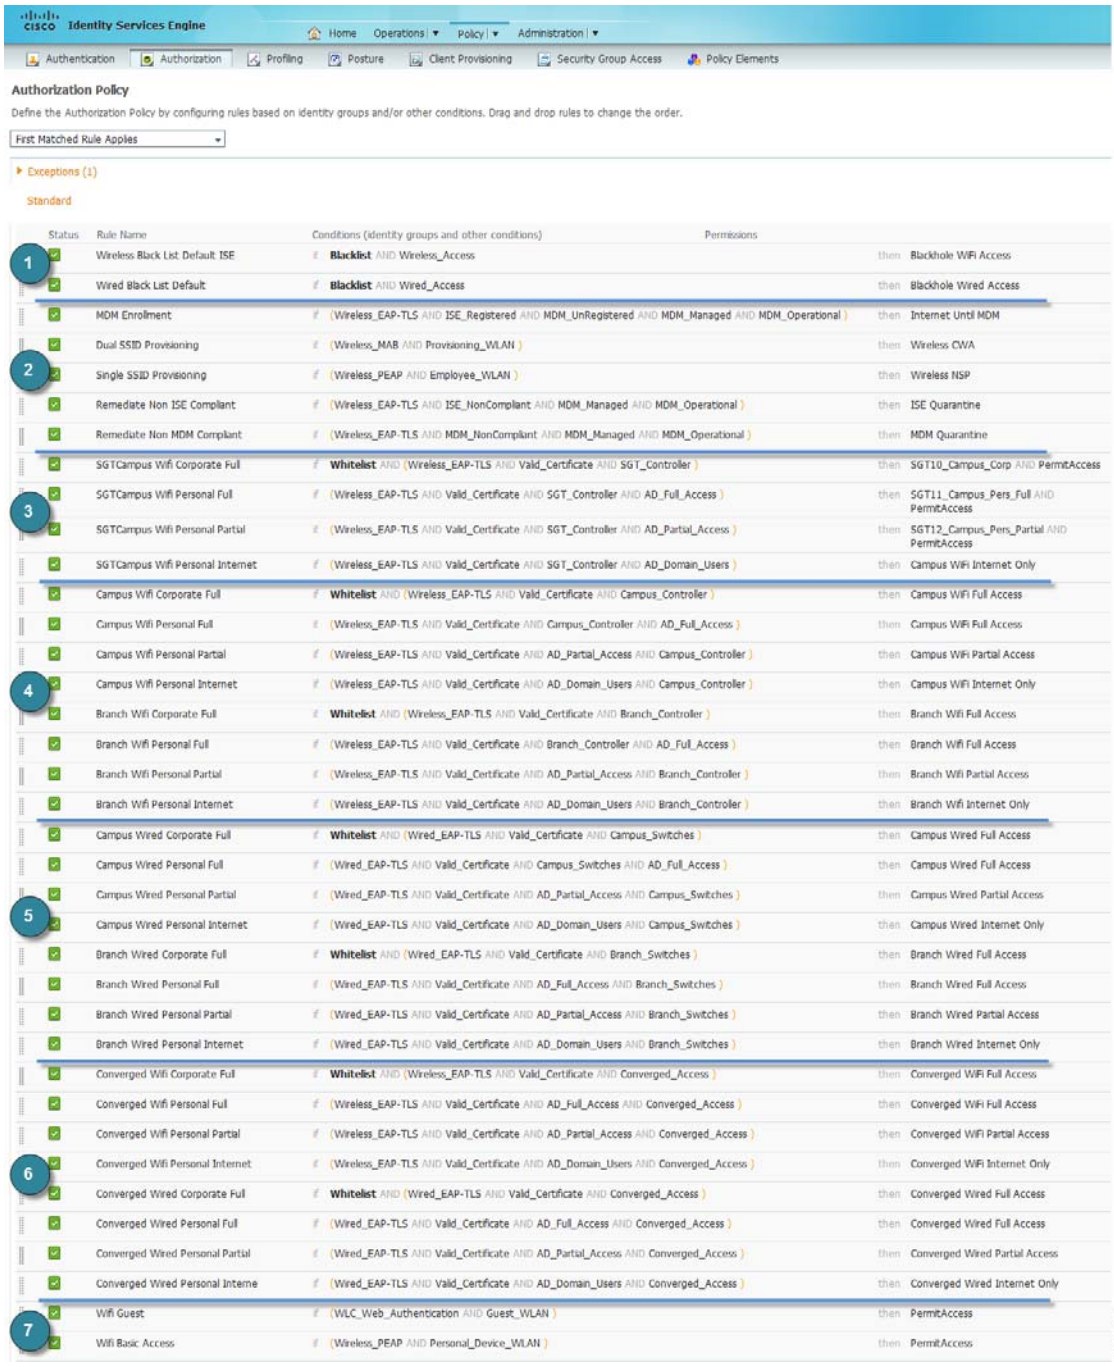
ISE Authorization Policy

For reference purposes, the complete authorization policy used during testing is shown in Figure 16-22. The figure highlights the following sections:

1. Used for blacklisting lost or stolen devices.
2. On-boarding and MDM registration/remediation.
3. Wireless devices connecting from an access point in a SGT_Enabled location.
4. Wireless devices connecting from an access point in the campus or branch locations.
5. Wired devices connecting from a campus or branch locations.

6. Wired and Wireless devices connecting from a converged location.
7. Guest and Basic Access.

Figure 16-22 Complete Authorization Policy





BYOD Advanced Use Case—Mobile Device Manager Integration

Revised: August 7, 2013

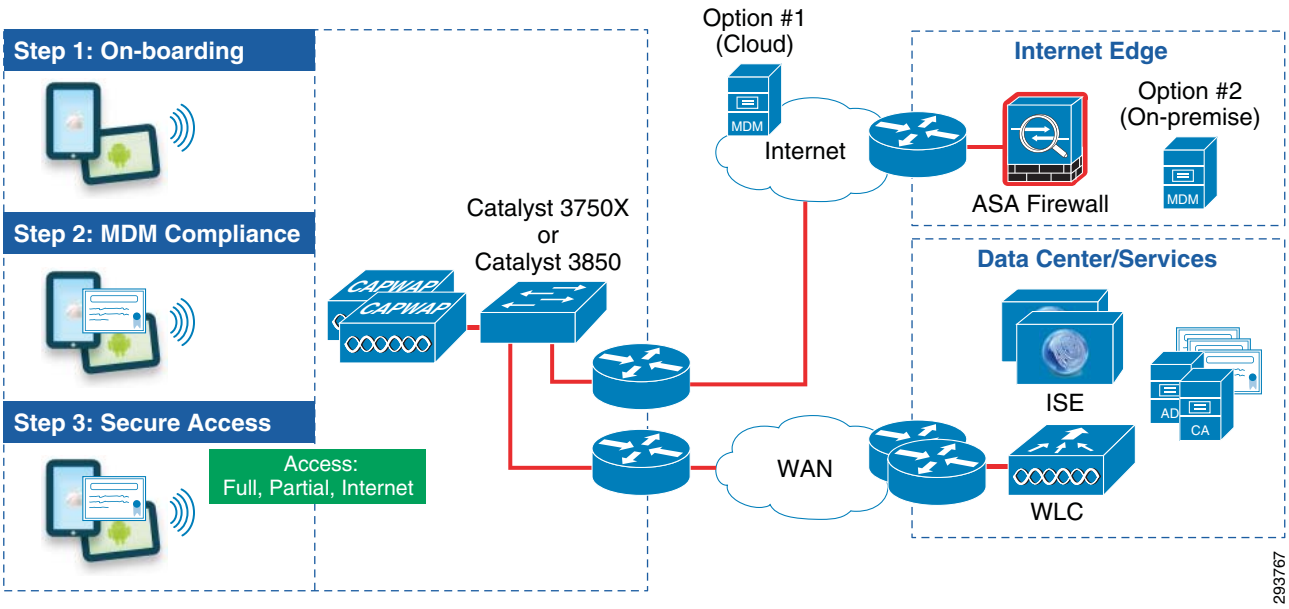
This chapter focuses on getting additional posture information from the integration with third-party Mobile Device Managers (MDMs). While previous chapters focused on on-boarding corporate and personal devices and providing differentiated access, this chapter makes use of more detailed endpoint information to enforce authorization policies.

MDM servers secure, monitor, manage, and support mobile devices to secure and control the use of mobile applications. The network is the only entity that can provide granular access to endpoints based on VLAN assignment, ACLs (named or downloadable [DACL]), SGTs, FlexConnect ACLs, etc. By integrating with third-party MDM servers, the Cisco ISE receives the necessary device attributes to enforce a more granular network access to those endpoints.

Figure 17-1 shows the interoperability between MDMs and the Cisco ISE. Once a device has been on-boarded with ISE, ISE queries the MDM for additional endpoint information. If the endpoint is registered and compliant with MDM policies, the device is granted access, based on other attributes, as described in the previous sections.

Notice that the MDM server can be deployed on-premise or in the cloud.

Figure 17-1 MDM Interoperability



The following steps take place when checking for device compliance:

1. The user connects to an on-boarding SSID and is guided through the registration and on-boarding process with ISE.
2. Once the user has been on-boarded with the proper certificate/profiles, the user connects to the secure Employee SSID.
3. ISE makes an API call to the MDM server. If the device is not registered with the MDM, the user is presented with the appropriate page to proceed to their MDM enrollment page.
4. Once the user completes the enrollment with the MDM server, they return to an enrollment redirect page that includes a continue button. When the user selects the continue option from the page, ISE will issue a Change of Authorization (CoA), forcing the user to re-authenticate. The API should now indicate the user has enrolled with the MDM. The MDM API results are cached by ISE for the duration of the authorization flow.
5. ISE uses the cached MDM information for the specific MAC address to verify the device's posture including its MDM compliance status. If the device is not in compliance with the MDM policies, the user is once again informed and is asked to become compliant.
6. Once the device becomes compliant, the user is authorized to access the network based on the assigned permissions (Full, Partial Access, or Internet Access).
7. ISE can poll the MDM server periodically to get compliance information.

Chapter 13, "Mobile Device Manager Integration for BYOD" provides more details on how to configure the integration between ISE and the MDM.

Supported MDM Functions

Cisco ISE relies on REST API calls to query the external MDM server for additional endpoint information. The communication between ISE and the MDM is mostly unidirectional, where ISE sends different commands to the MDM. Some commands query for device information (model, compliance status, serial number, etc.) while others can invoke an action on the device (Corporate Wipe, Full Wipe, PIN lock, etc.).

The following are some of the functions that ISE performs in conjunction with an MDM server:

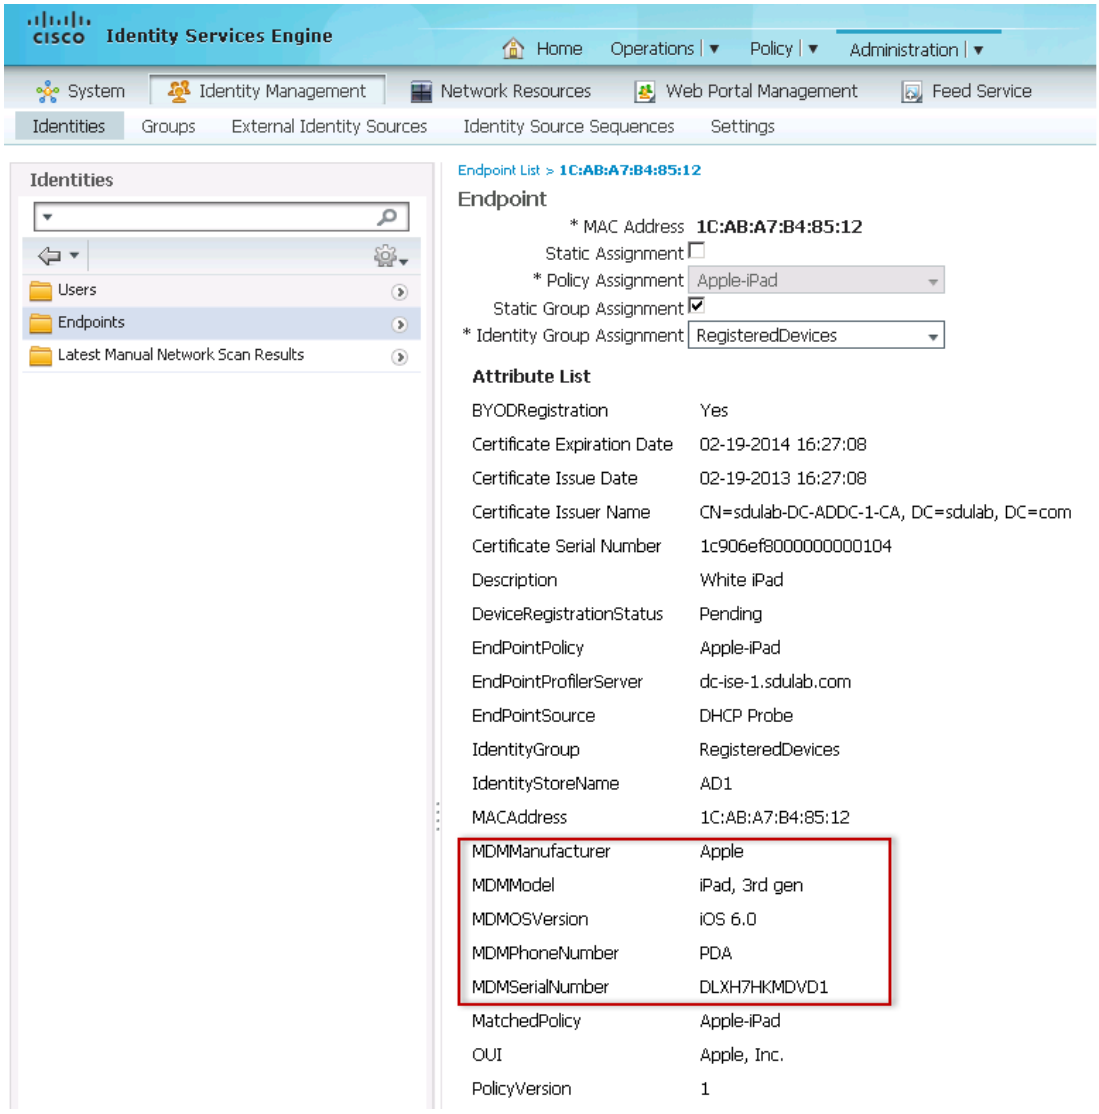
- Device registration—Unregistered endpoints connecting to the network are redirected to a web page hosted on the MDM server to initiate the MDM enrollment process.
- Device remediation—Cisco ISE imposes a captive portal on noncompliant endpoints. The user is redirected to a web page hosted by ISE but populated with device posture information obtained via the MDM API. The page informs the user and what actions to take to become compliant.
- Periodic compliance checks—Cisco ISE polls the MDM server periodically for a list of non-MDM compliant devices. ISE will determine if any device on the list is currently associated with the network and will issue a CoA against those devices.
- Device instructions through the MDM server—Remote actions can be issued for users' devices through the MDM server.

The endpoint database is updated with additional information from the MDM server that cannot be collected using the Cisco ISE Profiler. The following device attributes can be obtained from the MDM:

- MDMManufacturer
- MDMModel
- MDMOSVersion
- MDMPhoneNumber
- MDMSerialNumber
- MDMIMEI

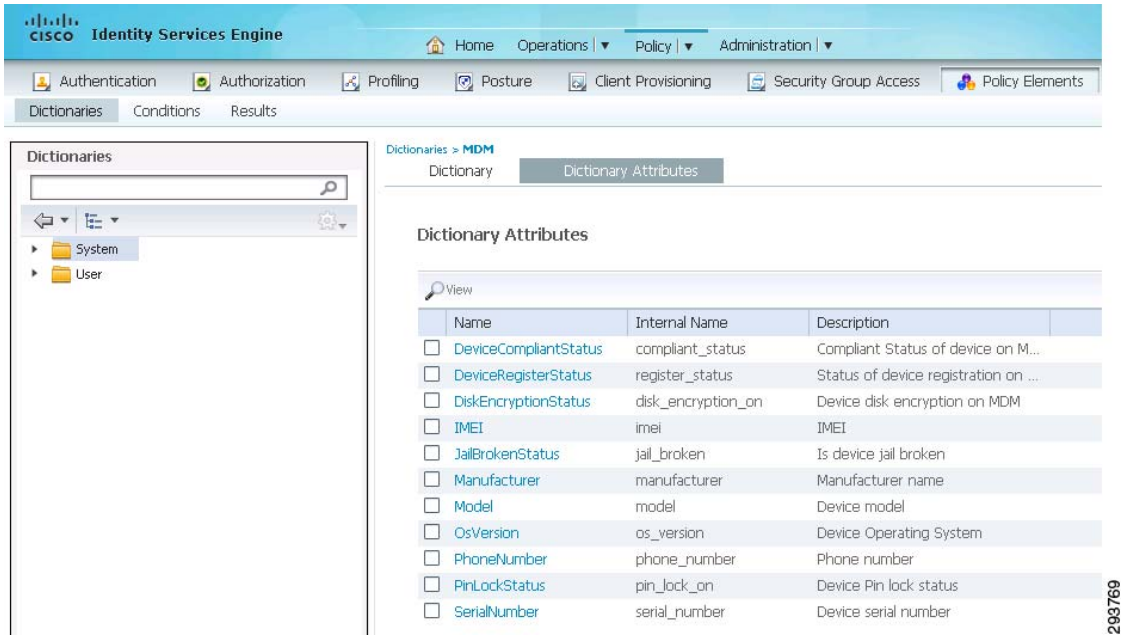
These attributes can be viewed by clicking **Administration > Identity Management > Identities > Endpoints**, as shown in Figure 17-2.

Figure 17-2 MDM Attributes



The integration with an MDM server allows the Cisco ISE to configure policies based on additional MDM attributes. The dictionary attributes can be found under **Policy > Dictionaries > MDM > Dictionary Attributes**, as shown in Figure 17-3.

Figure 17-3 Dictionary Attributes



Some of these attributes are used in several authorization rules in this design guide.

Integration Process Flow

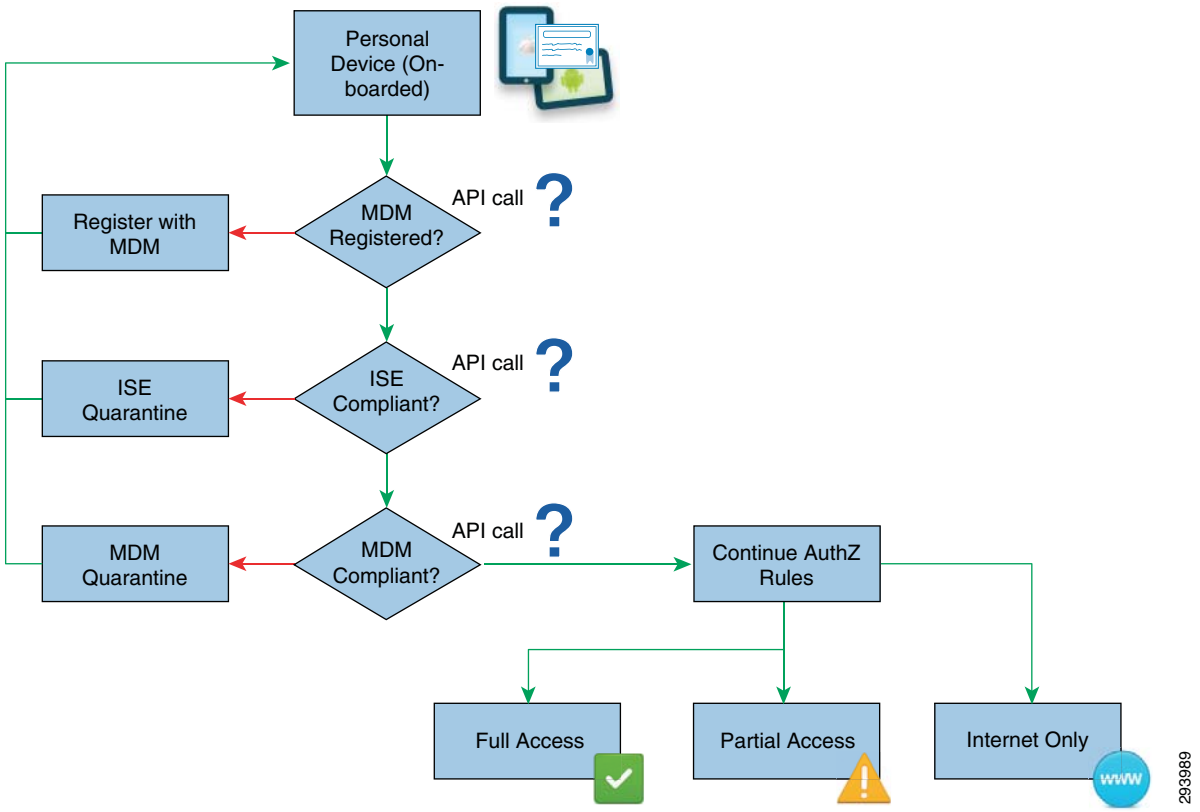
Figure 17-4 shows the integration process between ISE and the MDM:

1. The device has been on-boarded and the user connects to the BYOD_Employee SSID.
2. Cisco ISE makes an API call to the MDM server to verify that the device has been registered with the MDM.
3. If the device is not registered and is required to be registered, the user is asked to enroll with the MDM. This may include installing the appropriate application from the Apple App Store or Google Play.
4. ISE enforces some device attributes before allowing network access. If the device is not ISE Compliant, quarantine the device (explained below).
5. If the device is not compliant with the MDM policies, quarantine the device.

If the user has been on-boarded and is compliant with ISE and MDM policies, allow the proper permissions based on different rules.

Note The previous chapters explained how to grant Full Access, Partial Access and Internet Only access to personal and corporate devices.

Figure 17-4 MDM Compliance Checks



ISE Compliance Check

Once the device has been registered with the MDM, the ISE checks for certain device attributes before allowing access to the network. This can serve as an additional check and the first opportunity to verify some device attributes.

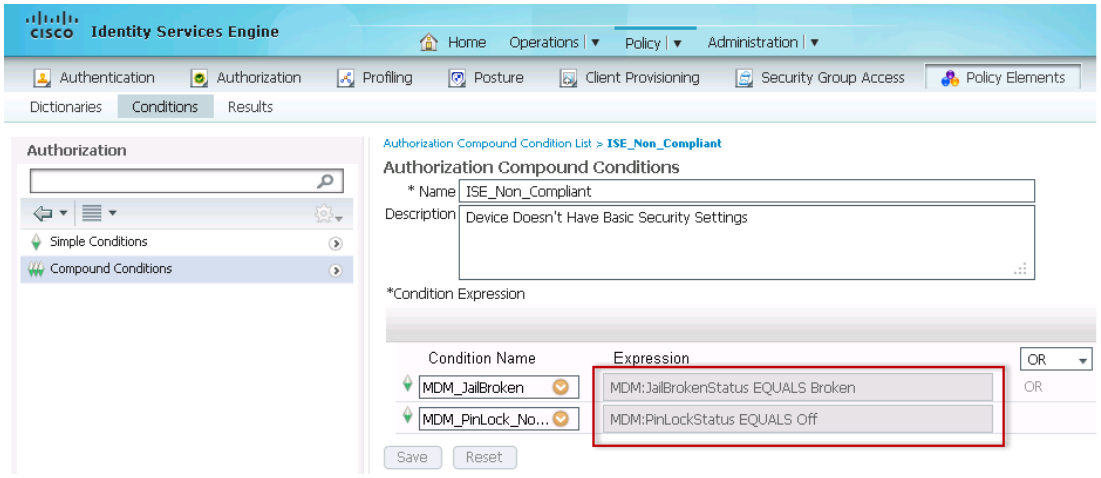
In this design guide, two device attributes were considered as the minimum for devices to obtain access to the network:

- Jailbroken or rooted devices—Not allowed into the network.
- PIN lock enforcement—Devices without a device PIN lock are denied access.

ISE makes the *JailBrokenStatus* and *PinLockStatus* API calls to the MDM to verify these attributes.

A compound condition was defined in the ISE to check for these two attributes. To define this compound condition click **Policy > Policy Elements > Conditions > Compound Conditions**, as shown in Figure 17-5.

Figure 17-5 ISE_Non_Compliant



Additional dictionary attributes could be added to accommodate the security policies of each organization.

MDM Compliance Check

For devices that have been on-boarded and have met the ISE compliance check, an additional API call is made to the MDM to make sure the device has met all the compliance requirements established by the MDM.

If the device is not fully compliant with the MDM, ISE grants access to the Internet and denies access to all internal resources. When the user tries to access an internal resource, ISE redirects the session to a portal, highlighting the steps required to be completed to meet the MDM compliance rules.

ISE makes the DeviceCompliantStatus API call to the MDM to verify compliance. The MDM establishes what conditions result in a device being non-compliant.

ISE Configuration

Several ISE features play a role in verifying for ISE and MDM compliance before permissions can be applied. Some of them include Logical Profiles, authorization rules, ACLs, and API calls to the MDM to receive endpoint attributes.

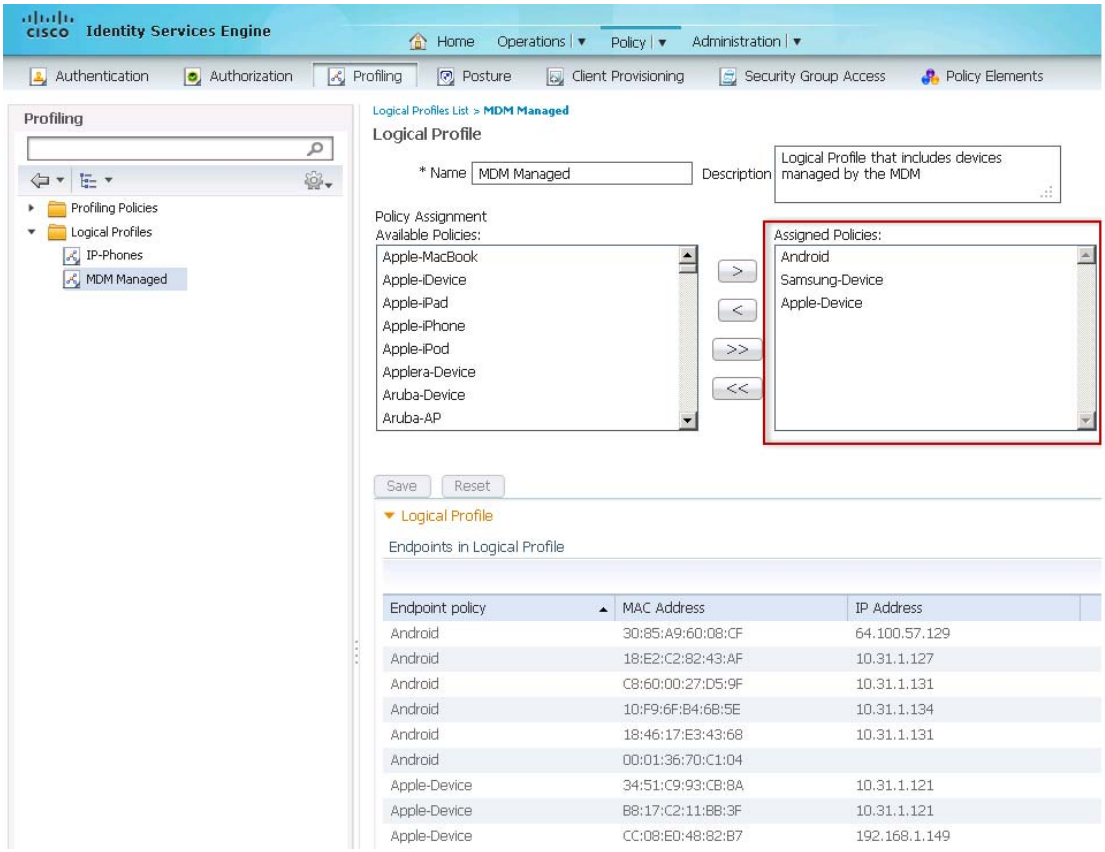
Logical Profile

The profiling service in the Cisco ISE is able to identify devices connecting to the network to grant the appropriate access to endpoints based on their device type. By collecting attributes of endpoints and grouping them according to their profiles, unique policies may be enforced for specific type of devices.

A logical profile is a virtual container of objects that share a common attribute. One example of a logical profile is a set of devices that includes all tablets or all smartphones. A second example is a set of all Android or Apple devices. For the purposes of this design guide, a logical profile was created to include the devices that are managed by the MDM. This allows the administrator to dynamically add or remove devices managed by the MDM.

Figure 17-6 shows the MDM Managed logical profile used to group devices allowed and managed or licensed by the MDM. This logical profile is used when defining the ISE authorization policy and includes devices profiled as Android, Samsung-Device, and Apple-Device. To configure this logical profile on the ISE, click **Policy > Profiling > Logical Profiles**.

Figure 17-6 MDM Managed Logical Profile

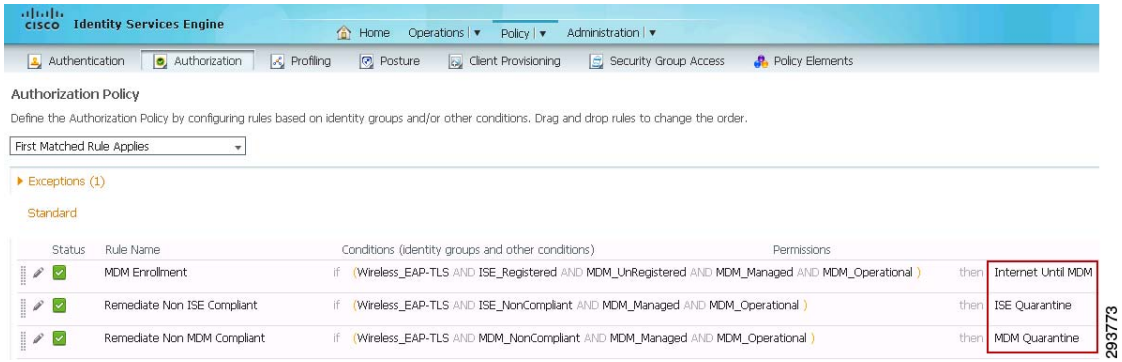


The logical profile could be easily expanded to include other devices supported and managed by the MDM without modifying the ISE authorization policy.

Authorization Policy

Figure 17-7 shows the ISE authorization rules used to enforce MDM and ISE compliance. These authorization rules are executed after the user has on-boarded the mobile device and before additional access (e.g., Full, Partial, Internet) to the network is granted.

Figure 17-7 MDM Compliance Authorization Rules



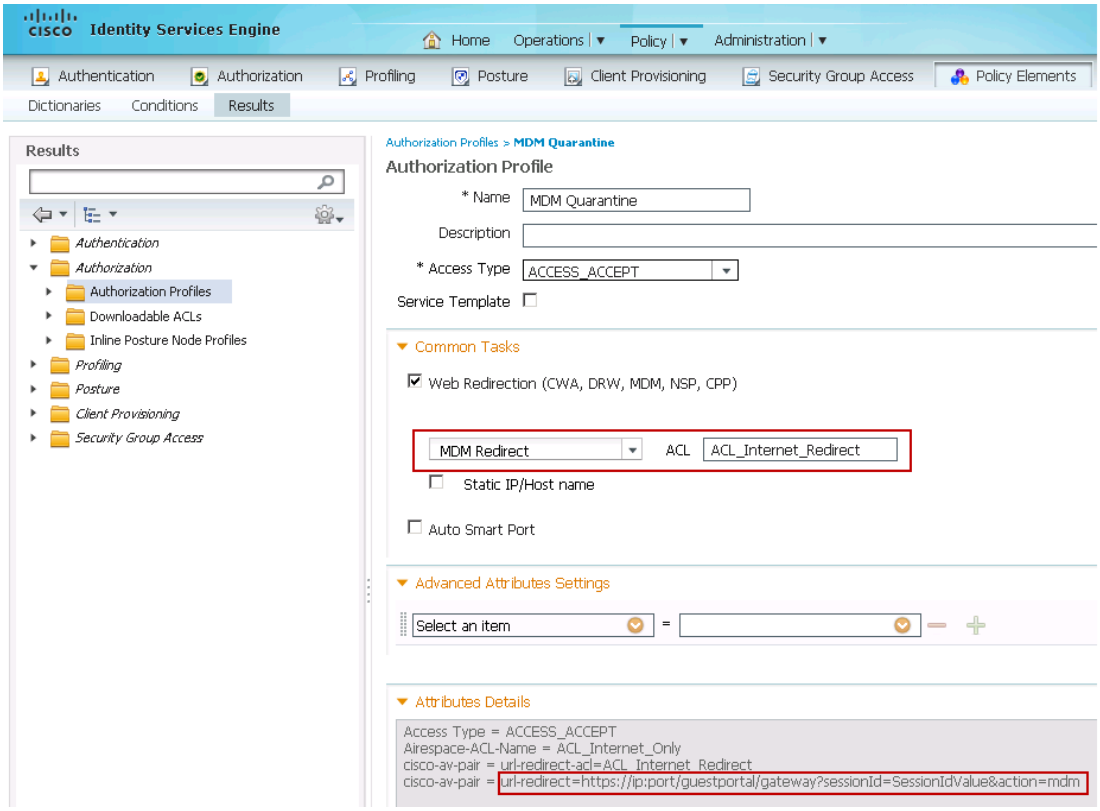
MDM Enrollment Rule

This rule matches when the following conditions are met:

- The endpoint connected via a wireless 802.1X SSID.
- Logical Profile equals MDM Managed—The device is managed and supported by the MDM.
- The device has gone through the on-boarding process and registered with ISE.
- The device has **not** registered with the MDM.

If these conditions match, the *Internet Until MDM* authorization profile is used. This profile is configured to redirect nonregistered devices to the redirect URL highlighted in Figure 17-8.

Figure 17-8 Internet Until MDM

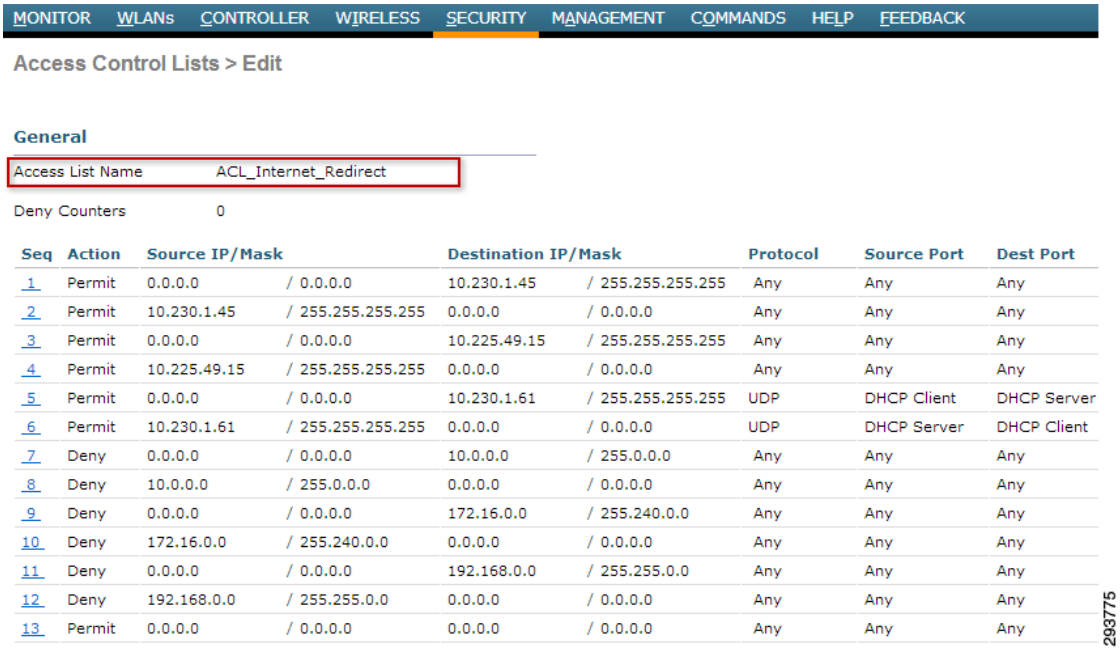


The authorization profile relies on two named ACLs previously defined in the Wireless LAN Controller: ACL_Internet_Redirect and ACL_Internet_Only. ACL_Internet_Redirect is shown as being applied to the MDM Redirect setting in the figure above. ACL_Internet_Only is shown as being sent to the wireless controller via the Radius:Airespace-ACL-Name attribute value (AV) pair in Figure 17-8. The behavior of the two ACLs is slightly different between CUWN wireless controllers, such as the CT5508 and Flex 7500, and IOS XE based controllers, such as the CT5760 and Catalyst 3850.

Note Within this document, wireless LAN Controller refers to either a standalone appliance such as the Cisco CT5508, Flex 7500, or CT5760 wireless controllers or to wireless controller functionality integrated within the Catalyst 3850 Series switch.

For CUWN wireless controllers, ACL_Internet_Redirect functions as both the ACL which controls web redirection, as well as the ACL which controls what the wireless client is allowed to access on the network. ACL_Internet_Only serves simply as a extra security configuration. CUWN wireless controllers do not make use of this ACL when URL redirection is specified. For CUWN wireless controllers the ACL_Internet_Redirect ACL shown in Figure 17-9 can be the same as the ACL_Internet_Only ACL discussed in the previous chapters.

Figure 17-9 ACL_Internet_Redirect



Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client
7	Deny	0.0.0.0	/ 0.0.0.0	10.0.0.0	/ 255.0.0.0	Any	Any	Any
8	Deny	10.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
9	Deny	0.0.0.0	/ 0.0.0.0	172.16.0.0	/ 255.240.0.0	Any	Any	Any
10	Deny	172.16.0.0	/ 255.240.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
11	Deny	0.0.0.0	/ 0.0.0.0	192.168.0.0	/ 255.255.0.0	Any	Any	Any
12	Deny	192.168.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
13	Permit	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any

The ACL specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Deny IP access to and from internal network address space (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).
- Allow access to and from all other subnets (Internet access).

.For Cisco IOS XE based wireless controllers, ACL_Internet_Redirect functions strictly as the ACL which controls web redirection. ACL_Internet_Only functions as the ACL which controls what the wireless client is allowed to access on the network. Hence IOS XE based wireless controllers make use of both ACLs when URL redirection is specified. An example of the ACL_Internet_Redirect ACL for Cisco IOS XE based wireless controllers is shown in Example 17-1.

Example 17-1 Internet Redirect ACL for IOS XE Based Controllers

```
!
ip access-list extended ACL_Internet_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   ip any host 10.230.1.45
 deny   ip any host 10.225.49.15
 permit ip any 10.0.0.0 0.255.255.255
 permit ip any 172.16.0.0 0.15.255.255
 permit ip any 192.168.0.0 0.0.255.255
 deny   ip any any
!
```

The above ACL specifies the following access:

- Deny (do not redirect) DHCP access (bootpc and bootps).
- Deny (do not redirect) IP access to and from the DNS server (10.230.1.45).

- Deny (do not redirect) IP access to and from the ISE server (10.225.49.15).
- Allow (redirect) IP access to and from the rest of the internal network IP address space (10.0.0.0 /8, 172.16.0.0 /12, 192.168.0.0 /16).
- Deny (do not redirect) all other access to the Internet.

An example of the ACL_Internet_Only ACL for Cisco IOS XE based wireless controllers is shown in Example 17-2.

Example 17-2 Access ACL for IOS XE Based Controllers

```
!
ip access-list extended ACL_Internet_Only
 permit udp any eq bootpc any eq bootps
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 172.16.0.0 0.15.255.255
 deny   ip any 192.168.0.0 0.0.255.255
 permit ip any any
!
```

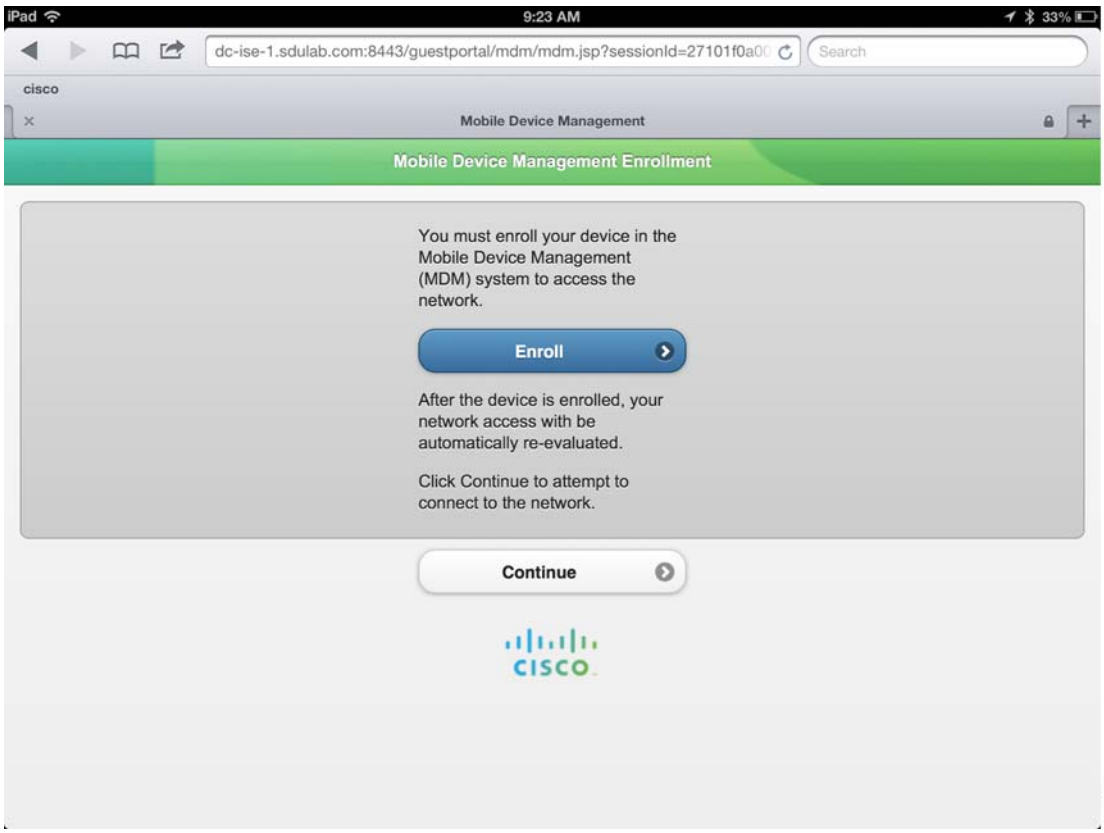
The above access-list specifies the following access:

- Allow DHCP access (bootpc and bootps).
- Allow IP access to and from the DNS server (1.230.1.45).
- Allow IP access to and from the ISE server (1.225.49.15).
- Deny IP access to and from the rest of the internal network IP address space (10.0.0.0 /8, 172.16.0.0 /12, 192.168.0.0 /16).
- Allow access to all other addresses (Internet addresses).

Note The access list shown in Example 17-2 is generic and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

The first time the user tries to browse an internal resource, the session is redirected to a page similar to the one in Figure 17-10 providing a link to register with the appropriate MDM.

Figure 17-10 Register with MDM

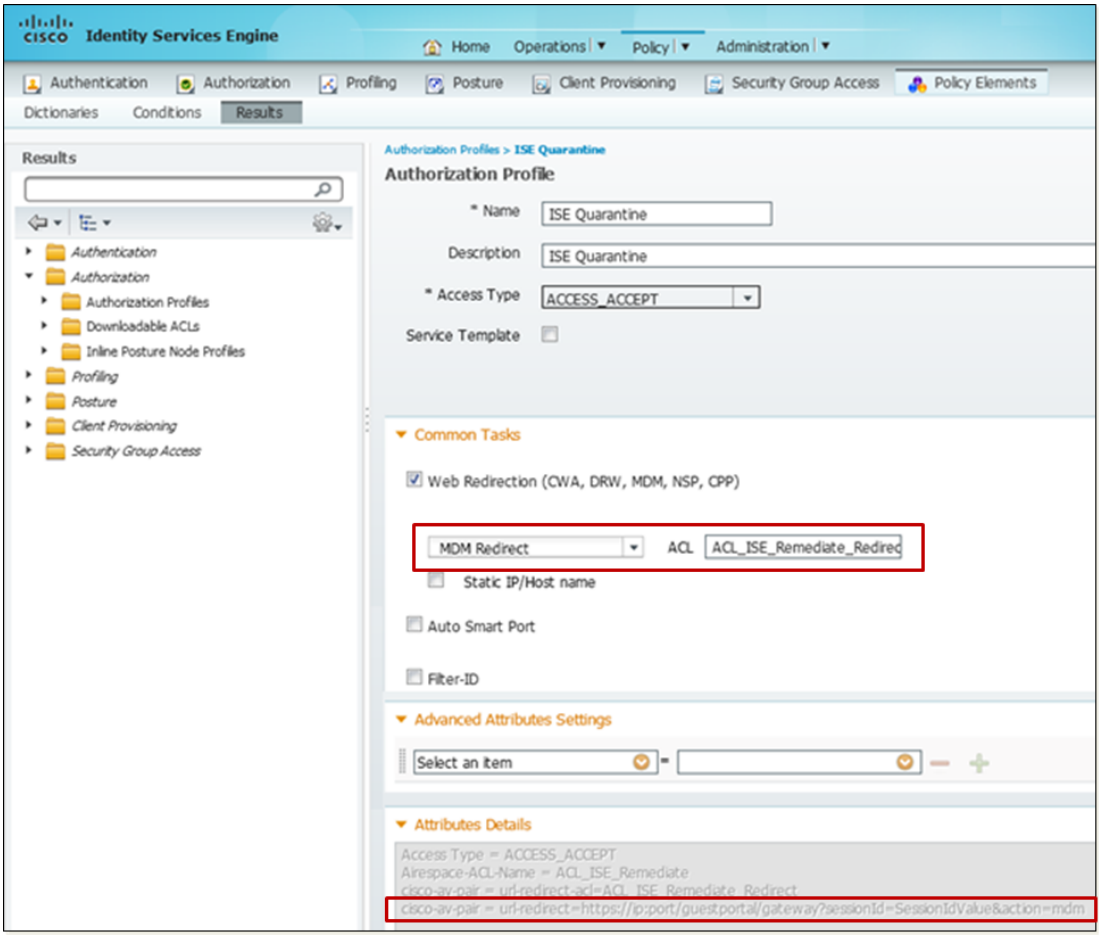


Remediate Non-ISE Compliant Rule

This rule matches when the following conditions are met:

- The connection originated using EAP-TLS authentication (defined as a compound condition, see below).
- The device does not comply with ISE requirements. The ISE_Non_Compliant compound condition is highlighted in Figure 17-5.
- Logical Profile equals MDM Managed—The device is managed and supported by the MDM and is shown in Figure 17-6.
- If these conditions match, the ISE Quarantine authorization profile is used. This profile is configured to redirect nonregistered devices to the redirect URL highlighted in Figure 17-11.

Figure 17-11 ISE Quarantine Authorization Profile

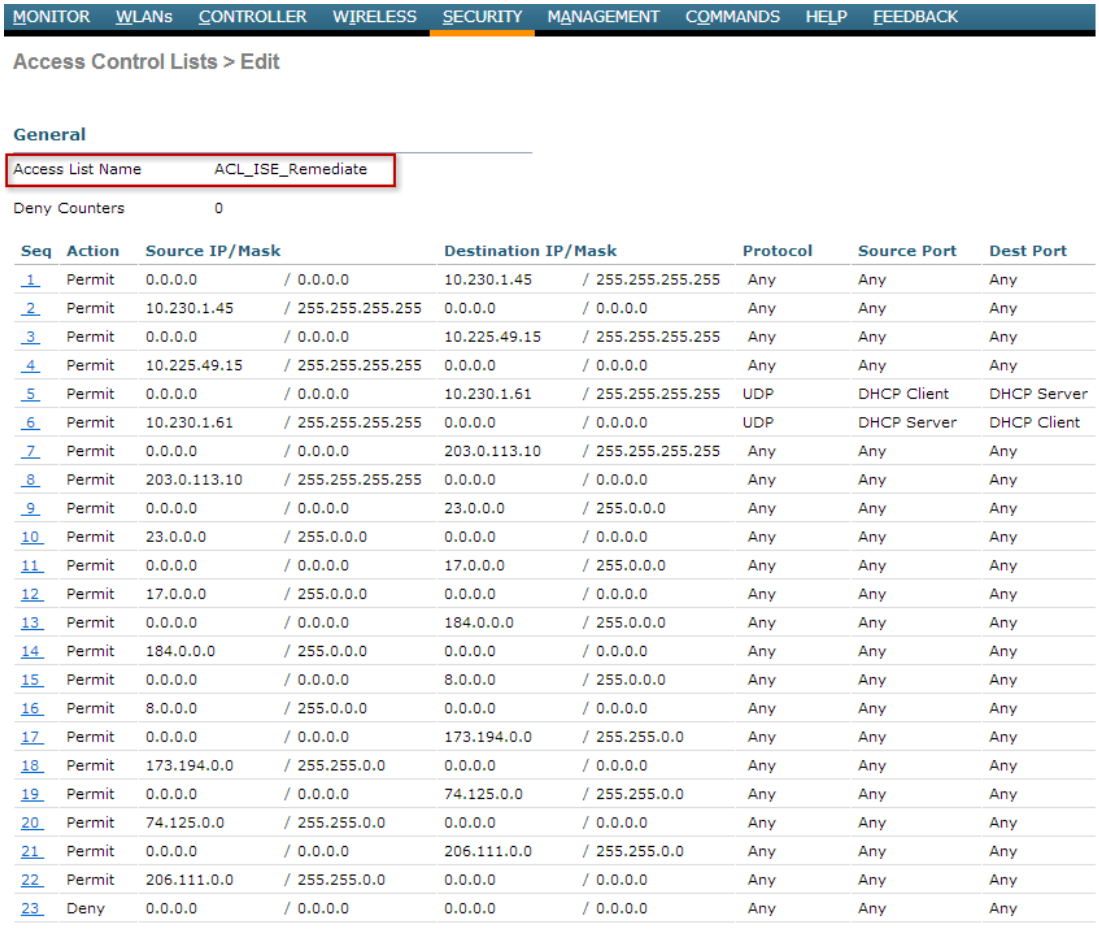


The authorization profile relies on two named ACLs, previously defined in the Wireless LAN Controller: ACL_ISE_Remediate_Redirect and ACL_ISE_Remediate. ACL_ISE_Remediate_Redirect is shown as being applied to the MDM Redirect setting in Figure 17-11. ACL_ISE_Remediate is shown as being sent to the wireless controller via the Radius:Airespace-ACL-Name attribute value (AV) pair in Figure 17-11. The behavior of the two ACLs is slightly different between CUWN wireless controllers, such as the CT5508 and Flex 7500, and IOS XE based controllers such as the CT5760 and Catalyst 3850.

For CUWN wireless controllers, ACL_ISE_Remediate_Redirect functions as both the ACL which controls web redirection, as well as the ACL which controls what the wireless client is allowed to access on the network. ACL_ISE_Remediate serves simply as a extra security configuration. CUWN wireless controllers do not make use of this ACL when URL redirection is specified.

For CUWN wireless controllers the ACL_ISE_Remediate and ACL_ISE_Remediate_Redirect ACLs can be the same. An example of the ACL_ISE_Remediate ACL is shown in Figure 17-12.

Figure 17-12 ACL_ISE_Remediate



MONITOR WLANs CONTROLLER WIRELESS SECURITY MANAGEMENT COMMANDS HELP FEEDBACK

Access Control Lists > Edit

General

Access List Name ACL_ISE_Remediate

Deny Counters 0

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port
1	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.45	/ 255.255.255.255	Any	Any	Any
2	Permit	10.230.1.45	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
3	Permit	0.0.0.0	/ 0.0.0.0	10.225.49.15	/ 255.255.255.255	Any	Any	Any
4	Permit	10.225.49.15	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
5	Permit	0.0.0.0	/ 0.0.0.0	10.230.1.61	/ 255.255.255.255	UDP	DHCP Client	DHCP Server
6	Permit	10.230.1.61	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	UDP	DHCP Server	DHCP Client
7	Permit	0.0.0.0	/ 0.0.0.0	203.0.113.10	/ 255.255.255.255	Any	Any	Any
8	Permit	203.0.113.10	/ 255.255.255.255	0.0.0.0	/ 0.0.0.0	Any	Any	Any
9	Permit	0.0.0.0	/ 0.0.0.0	23.0.0.0	/ 255.0.0.0	Any	Any	Any
10	Permit	23.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
11	Permit	0.0.0.0	/ 0.0.0.0	17.0.0.0	/ 255.0.0.0	Any	Any	Any
12	Permit	17.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
13	Permit	0.0.0.0	/ 0.0.0.0	184.0.0.0	/ 255.0.0.0	Any	Any	Any
14	Permit	184.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
15	Permit	0.0.0.0	/ 0.0.0.0	8.0.0.0	/ 255.0.0.0	Any	Any	Any
16	Permit	8.0.0.0	/ 255.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
17	Permit	0.0.0.0	/ 0.0.0.0	173.194.0.0	/ 255.255.0.0	Any	Any	Any
18	Permit	173.194.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
19	Permit	0.0.0.0	/ 0.0.0.0	74.125.0.0	/ 255.255.0.0	Any	Any	Any
20	Permit	74.125.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
21	Permit	0.0.0.0	/ 0.0.0.0	206.111.0.0	/ 255.255.0.0	Any	Any	Any
22	Permit	206.111.0.0	/ 255.255.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any
23	Deny	0.0.0.0	/ 0.0.0.0	0.0.0.0	/ 0.0.0.0	Any	Any	Any

293778

The ACL specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Allow IP access to and from the DHCP server (10.230.1.61).
- Allow IP access to and from the MDM server (203.0.113.10).
- Allow IP access to Apple Push Notification Server (23.0.0.0 /8 and 17.0.0.0 /8).
- Allow IP access to Google Cloud Messaging (184.0.0.0 /8, 8.0.0.0 /8, 173.194.0.0 /16, 74.125.0.0 /16, 206.111.0.0 /16).
- Deny IP access (redirect) to all other IP addresses.

For Cisco IOS XE based wireless controllers, ACL_Internet_Redirect functions strictly as the ACL which controls web redirection. ACL_Internet_Only functions as the ACL which controls what the wireless client is allowed to access on the network. Hence IOS XE based wireless controllers make use of both ACLs when URL redirection is specified. An example of the ACL_Internet_Redirect ACL for Cisco IOS XE based wireless controllers is shown in Example 17-3.

Example 17-3 Remediate Redirect ACL for IOS XE Controllers

```
!
```

```
ip access-list extended ACL_ISE_Remediate_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   ip any host 1.230.1.45
 deny   ip any host 1.225.49.15
 deny   ip any host 1.230.1.76
 deny   ip any 63.128.76.0 0.0.0.255
 deny   ip any 23.0.0.0 0.255.255.255
 deny   ip any 17.0.0.0 0.255.255.255
 deny   ip any 184.0.0.0 0.255.255.255
 deny   ip any 8.0.0.0 0.255.255.255
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
 deny   ip any host 1.225.100.10
 permit ip any any
!
```

The above ACL specifies the following access:

- Deny (do not redirect) DHCP traffic.
- Deny IP access (do not redirect) to and from the DNS server (1.230.1.45).
- Deny IP access (do not redirect) to and from the ISE server (1.225.42.15).
- Deny IP access (do not redirect) to and from MDM servers (host 1.230.1.76 and subnet 63.128.76.0 /24).
- Deny IP access (do not redirect) to Apple Push Notification Server (23.0.0.0 /8 and 17.0.0.0 /8).
- Deny IP access (do not redirect) to Google Cloud Messaging (184.0.0.0 /8, 8.0.0.0 /8, 74.125.0.0 /16, 173.194.0.0 /16, 206.111.0.0 /16).
- Permit IP access (redirect) traffic to all other IP addresses.

An example of the ACL_ISE_Remediate ACL for Cisco IOS XE based wireless controllers is shown in Example 17-4.

Example 17-4 Remediate Access ACL for Cisco IOS XE Controller

```
!
ip access-list extended ACL_ISE_Remediate
 permit udp any eq bootpc any eq bootps
 permit ip any host 1.230.1.45
 permit ip any host 1.225.49.15
 permit ip any host 1.230.1.76
 permit ip any 63.128.76.0 0.0.0.255
 permit ip any 23.0.0.0 0.255.255.255
 permit ip any 17.0.0.0 0.255.255.255
 permit ip any 184.0.0.0 0.255.255.255
 permit ip any 8.0.0.0 0.255.255.255
 permit ip any 74.125.0.0 0.0.255.255
 permit ip any 173.194.0.0 0.0.255.255
 permit ip any 206.111.0.0 0.0.255.255
 deny   ip any any
!
```

The above access-list specifies the following access:

- Allow DHCP traffic.
- Allow IP access to and from the DNS server (1.230.1.45).
- Allow IP access to and from the ISE server (1.225.42.15).
- Allow IP access to and from MDM servers (host 1.230.1.76 and subnet 63.128.76.0 /8).
- Allow IP access to Apple Push Notification Server (23.0.0.0 /8 and 17.0.0.0 /8).

- Allow IP access to Google Cloud Messaging (184.0.0.0 /8, 8.0.0.0 /8, 74.125.0.0 /16, 173.194.0.0 /16, 206.111.0.0 /16).
- Permit IP access to all other IP addresses.

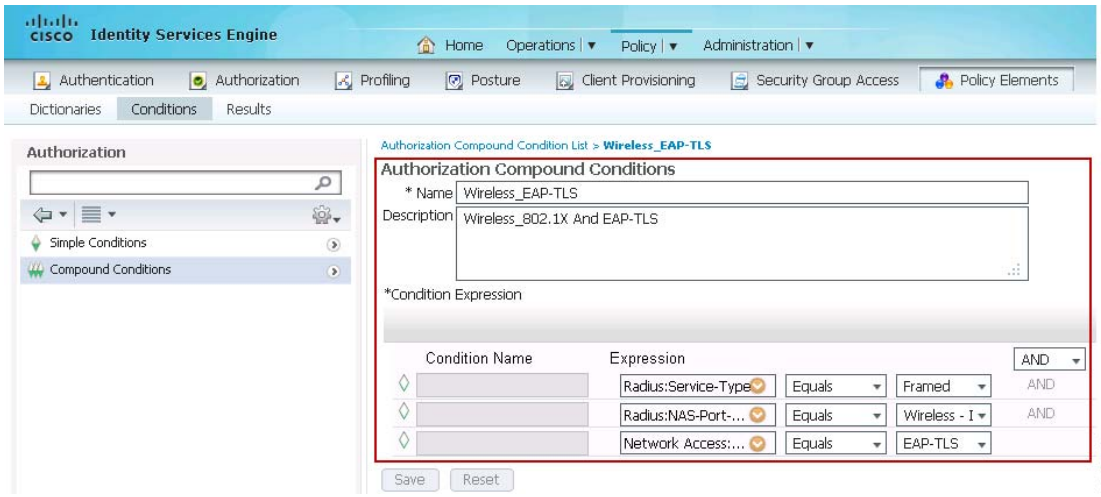
Note The access list shown in Example 17-4 is generic and not intended to work for every organization. An ACL should be more specific and only allow access to specific IP addresses and protocols in the required direction. A common practice is to make the ACLs as detailed as possible and to define every entry down to the port level.

The Wireless_EAP-TLS compound condition checks for the following conditions:

- Radius:Service-Type Equals Framed
- Radius:NAS-Port-Type Equals Wireless - IEEE 802.11
- Network Access:EapAuthentication Equals EAP-TLS

To define this compound condition, click **Policy > Conditions > Authorization > Compound Conditions**. Figure 17-13 shows how the Wireless_EAP-TLS condition combines several conditions into one.

Figure 17-13 Wireless_EAP-TLS Condition



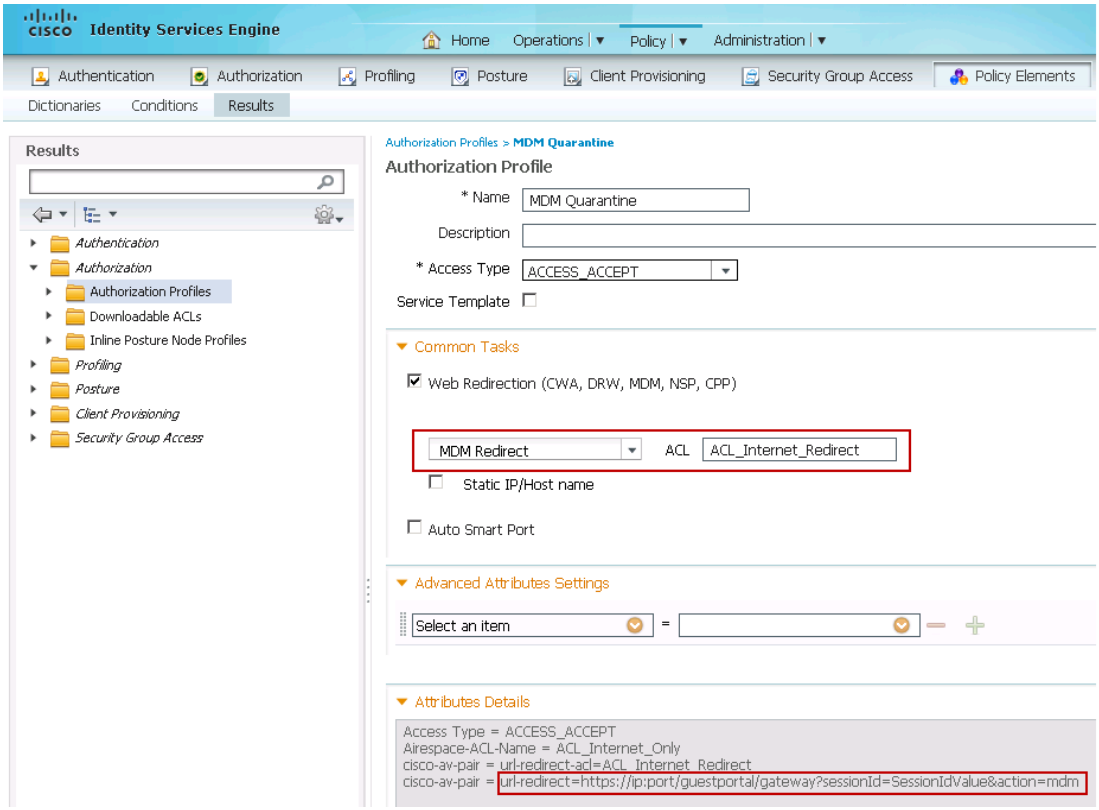
Remediate Non-MDM Compliant Rule

This rule matches when the following conditions are met:

- The connection originated using EAP-TLS authentication, as defined by the Wireless_EAP-TLS compound condition highlighted in Figure 17-13.
- The device does not comply with MDM policies. The ISE makes the DeviceCompliantStatus API call to the MDM to obtain this information.
- Logical Profile equals MDM Managed—The device is managed and supported by the MDM.

If these conditions match, the MDM Quarantine authorization profile is used. This profile is configured to redirect nonregistered devices to the redirect URL highlighted in Figure 17-14.

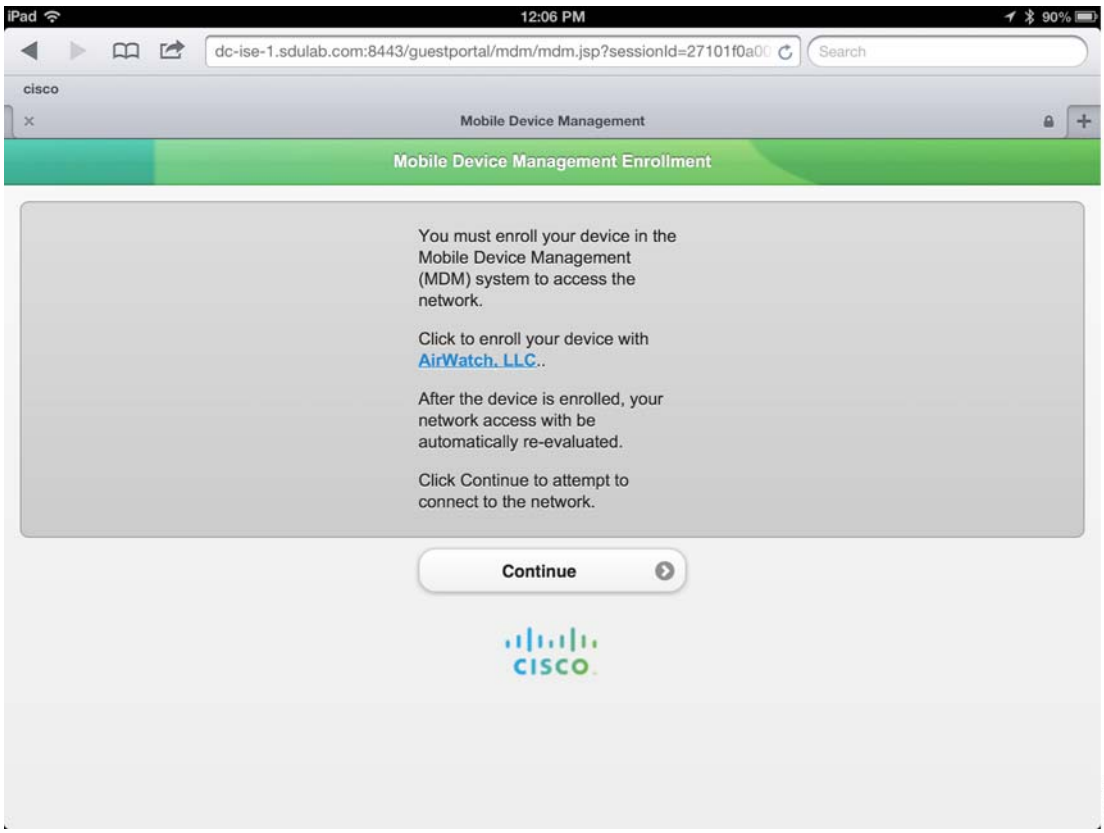
Figure 17-14 MDM Quarantine Authorization Profile



The authorization profile relies on the same two named ACLs: ACL_Internet_Redirect and ACL_Internet_Only previously discussed in Remediate Non-ISE Compliant Rule.

When the user tries to browse an internal resource, the session is redirected to page similar to the one in Figure 17-15, indicating why the device is not compliant with the MDM policies.

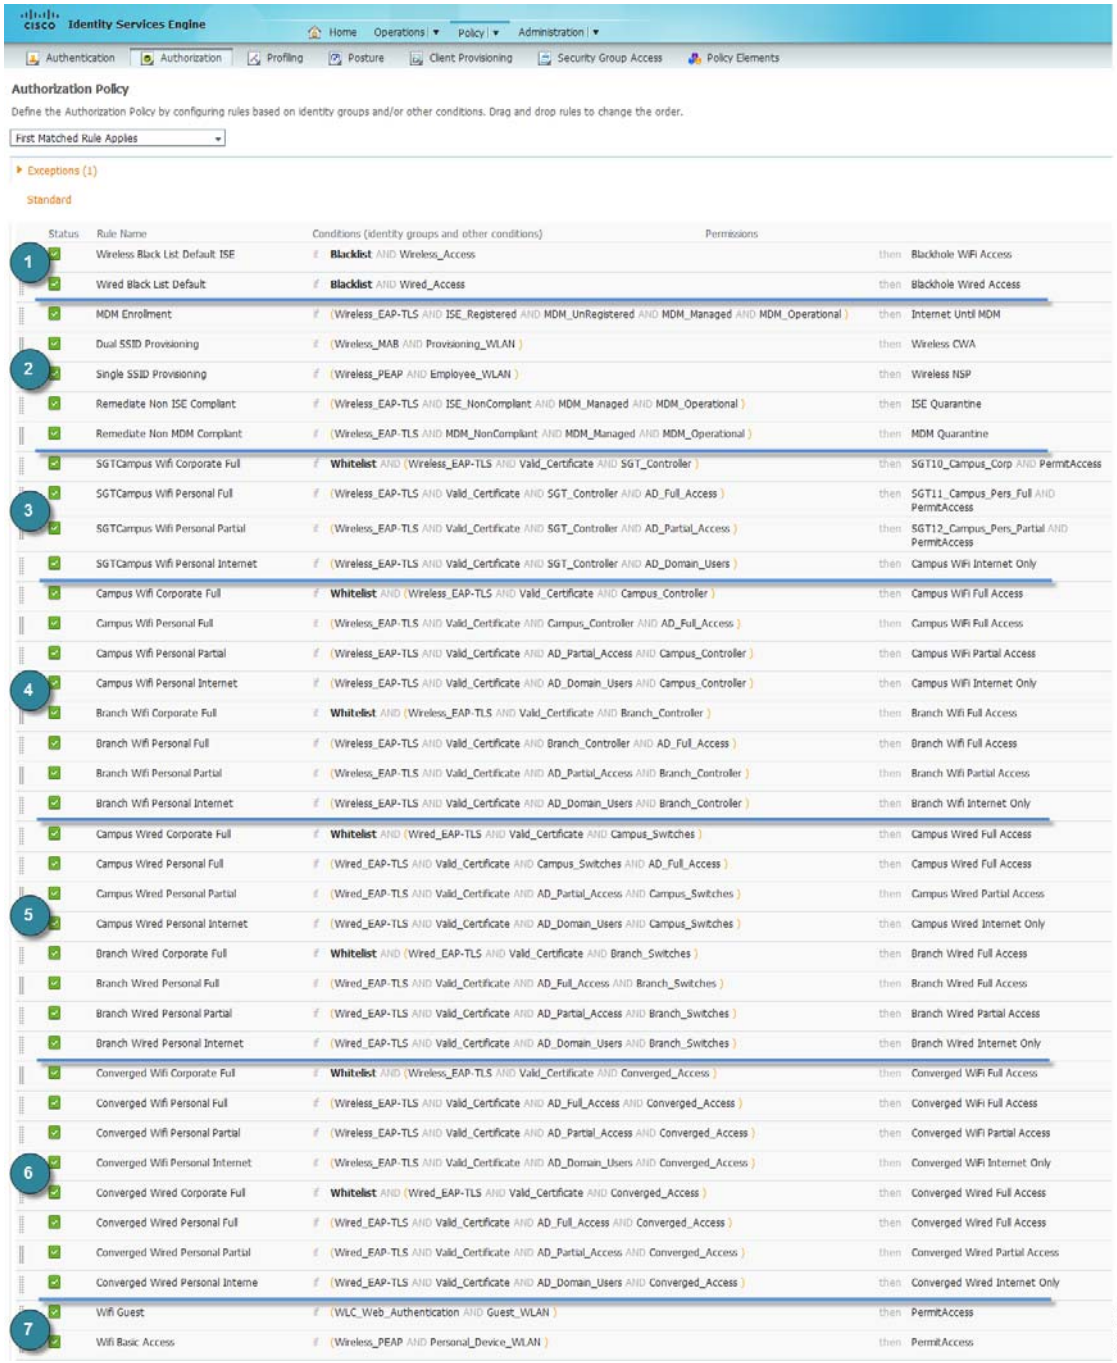
Figure 17-15 MDM Quarantine



For reference purposes, the complete authorization policy used during testing is shown in Figure 17-16. The figure highlights the following sections:

1. Used for blacklisting lost or stolen devices.
2. On-boarding and MDM registration/remediation.
3. Wireless devices connecting from an access point in a SGT_Enabled location.
4. Wireless devices connecting from an access point in the campus or branch locations.
5. Wired devices connecting from a campus or branch locations.
6. Wired and Wireless devices connecting from a converged location.
7. Guest and Basic Access.

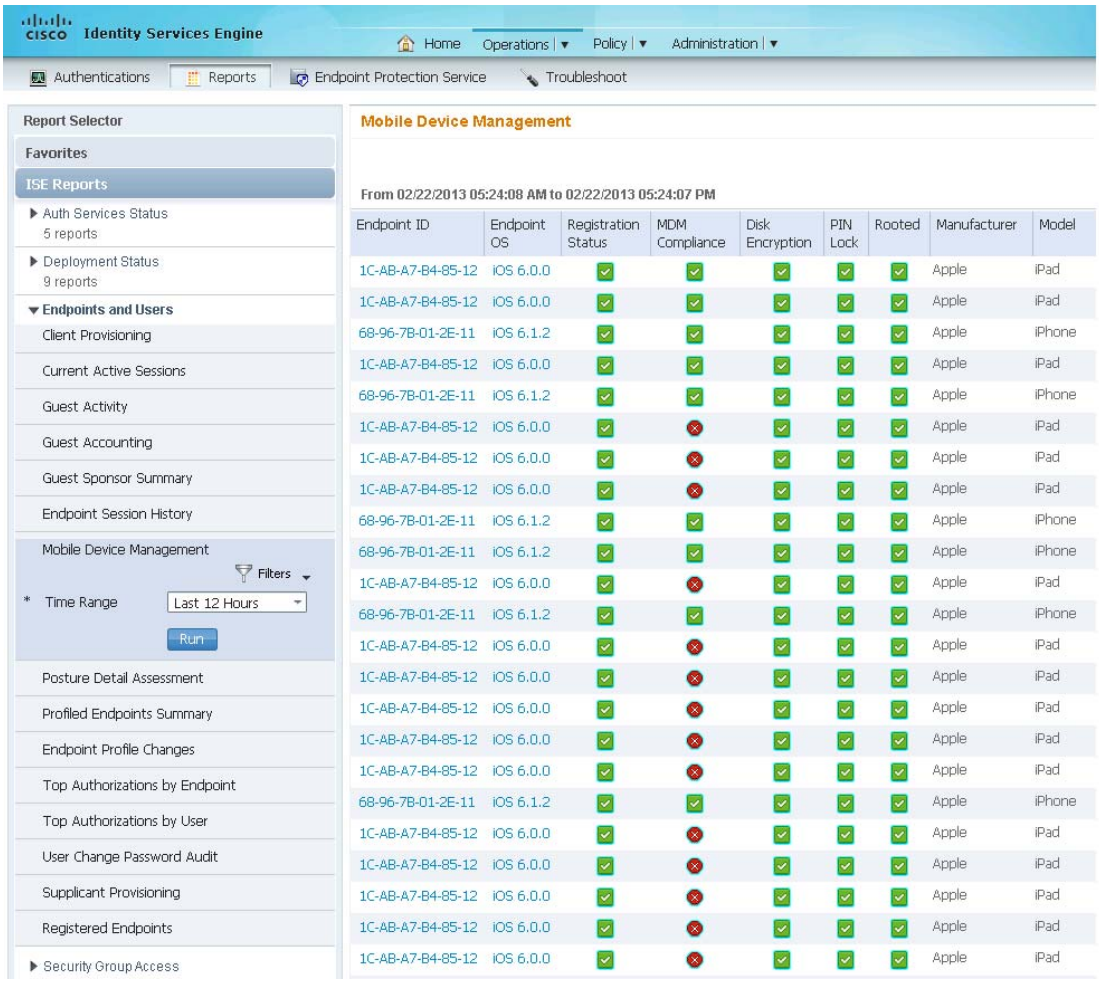
Figure 17-16 Complete Authorization Policy



MDM Reports

Cisco ISE provides a logging mechanism that is used for auditing, fault management, and troubleshooting. Several reports provide information related to endpoints. In Figure 17-17, the Mobile Device Management report shows the endpoints connected to the ISE, and several attributes collected from the MDM. Other reports provide additional information that may be used for reporting and troubleshooting.

Figure 17-17 MDM Report





BYOD Basic Access Use Case

Revised: August 7, 2013

Previous chapters of this design guide have examined on-boarding employee personal devices to provide full, partial, or Internet only access. The use of digital certificates provides an additional level of authentication security by preventing the spoofing of device MAC addresses. Additionally, the use of the guest portal for self-registration and the My-Devices portal streamline the on-boarding and maintenance of employee personal devices, resulting in lower IT operating costs associated with providing BYOD services.

Despite these benefits, a subset of organizations may still decide on a business policy which does not on-board wireless employee personal devices, yet provides some access to corporate services and the Internet for such devices. This may be because of one or more of the following reasons:

- The organization does not have the desire or the ability to deploy digital certificates on employee personal devices.
- The employee may decide to opt-out of having the organization manage their personal device.
- The organization does not wish to administratively manage and maintain separate lists of registered corporate devices and BYOD devices which have full network access.
- The organization may wish to simply restrict employee personal devices to "outside" of the corporate firewall due to an unknown or un-trusted security posture of such devices.

Because of this, the following sections discuss design options for wireless employee personal devices which do not involve on-boarding such devices. The designs are based around extending traditional guest wireless access discussed in Chapter 21, "BYOD Guest Wireless Access" and providing similar guest-like wireless access for employee personal devices.

Note Throughout this chapter, it is assumed that corporate-owned devices will still be on-boarded as discussed in Chapter 16, "BYOD Limited Use Case—Corporate Devices." The use of a whitelist is still necessary to prevent employee personal devices from getting full access to the corporate network.

Extending Guest Wireless Access to Employee Personal Devices

The following sections discuss two methods for extending guest wireless access, discussed in Chapter 21, "BYOD Guest Wireless Access," to also allow employee personal devices access to the guest network:

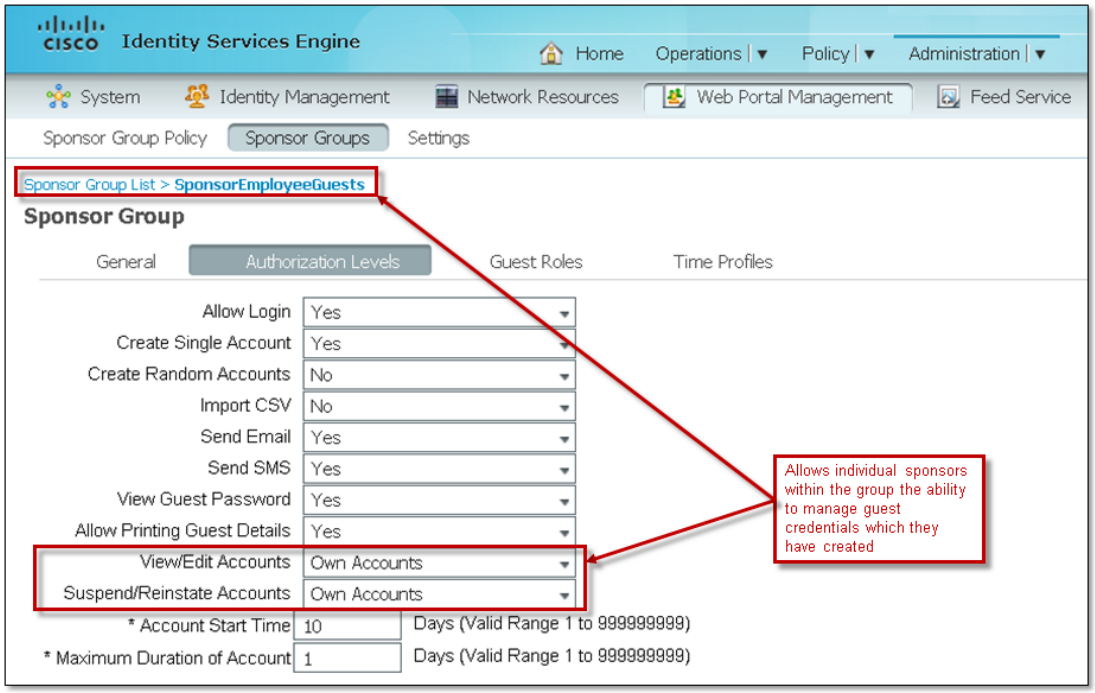
- Allowing employees to provision guest credentials for themselves.
- Extending guest web authentication (Web Auth) to also utilize the Microsoft Active Directory (AD) database when authenticating guests and employees using personal devices.

Allowing Employees to Provision Guest Credentials for Themselves

Chapter 21, "BYOD Guest Wireless Access" discusses provisioning guest credentials with the Cisco ISE sponsor portal. The most basic form of extending guest wireless access to employee personal devices is simply to allow employees to sponsor themselves as guests. Employees then manually connect to the open guest SSID to utilize personal devices on the wireless guest network.

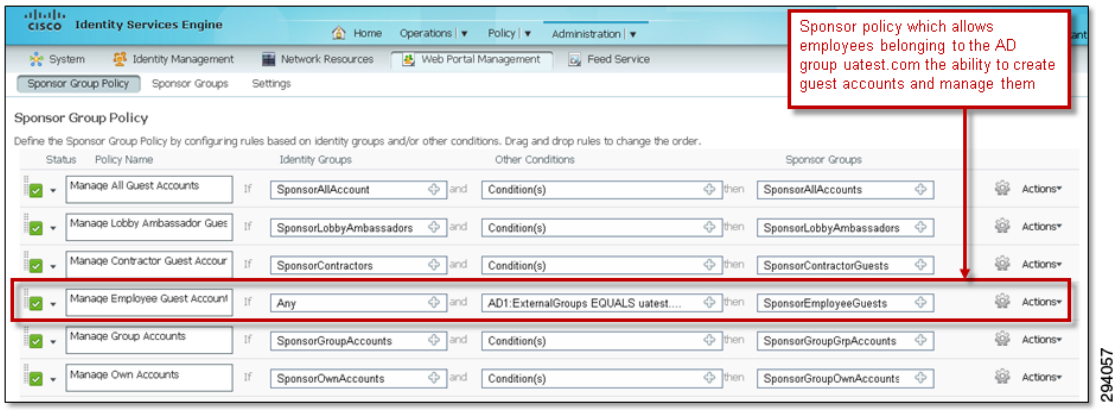
With the Cisco ISE sponsor portal, the sponsor authentication policy can be based on membership within a particular Microsoft AD group. This was discussed in Configuring the Cisco ISE Sponsor Portal, which discussed the use of Microsoft AD groups within the ISE sponsor group policy as a means to limit sponsor access to the Cisco ISE server. This can equally be used to allow broader access to the ISE sponsor portal simply by adding additional employees to those Microsoft AD groups. Alternatively a new sponsor group can be created that allows individual employees to configure guest credentials yet restricts them to being able to only modify credentials that they provisioned. An example is shown in the figures below.

Figure 18-1 Example ISE Sponsor Group for Employees to Create Self Guest Credentials



Membership to this ISE sponsor group can then be restricted to a Microsoft AD group via the ISE sponsor group policy, as shown in Figure 18-2.

Figure 18-2 Example ISE Sponsor Group Policy for Employees to Create Self Guest Credentials



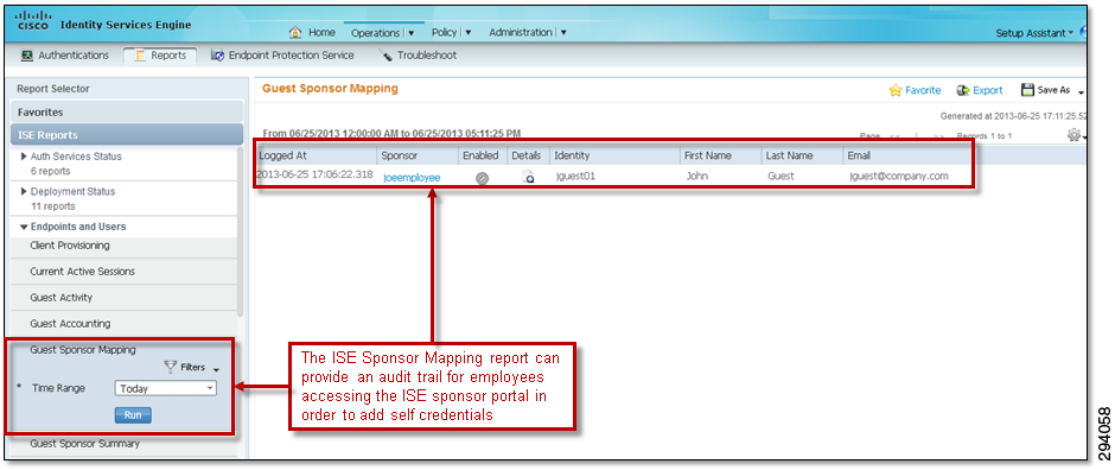
In this example, access to the sponsor group is limited to those members of the Microsoft Active Directory domain who are members of the group "Users/Domain Users". The exact condition for the example is of the form:

```
AD1:ExternalGroups EQUALS uatest.com/Users/Domain Users
```

The Microsoft Active Directory domain is "uatest.com" in this example. Note that the Microsoft Active Directory server must be configured as an external identity source to select this option. In this example it is known by the name "AD1".

One advantage of this design is that an audit trail exists through the ISE Reports for employees authenticating to the ISE sponsor portal. The ISE Sponsor Mapping report shows when the employee created a guest credential as well as the guest credential created. The ISE Sponsor Mapping report can be run over various time ranges from 30 minutes up to 30 days or a custom time range can be requested. This report can be used to gain a rough idea of which employees are accessing the ISE sponsor portal to create guest credentials and how often. An example is shown in Figure 18-3.

Figure 18-3 Example of the Audit Trail for an Employee Accessing the ISE Sponsor Portal



Note that this method of allowing employees the ability to create guest credentials for themselves does not prevent them from creating credentials for true guests who are visiting the organization. Corporate business policy should dictate that true guest credentials only be added by authorized members of

sponsor groups, as discussed in Chapter 21, "BYOD Guest Wireless Access." The next design option eliminates this issue by removing the ability for employees to create guest credentials for themselves altogether.

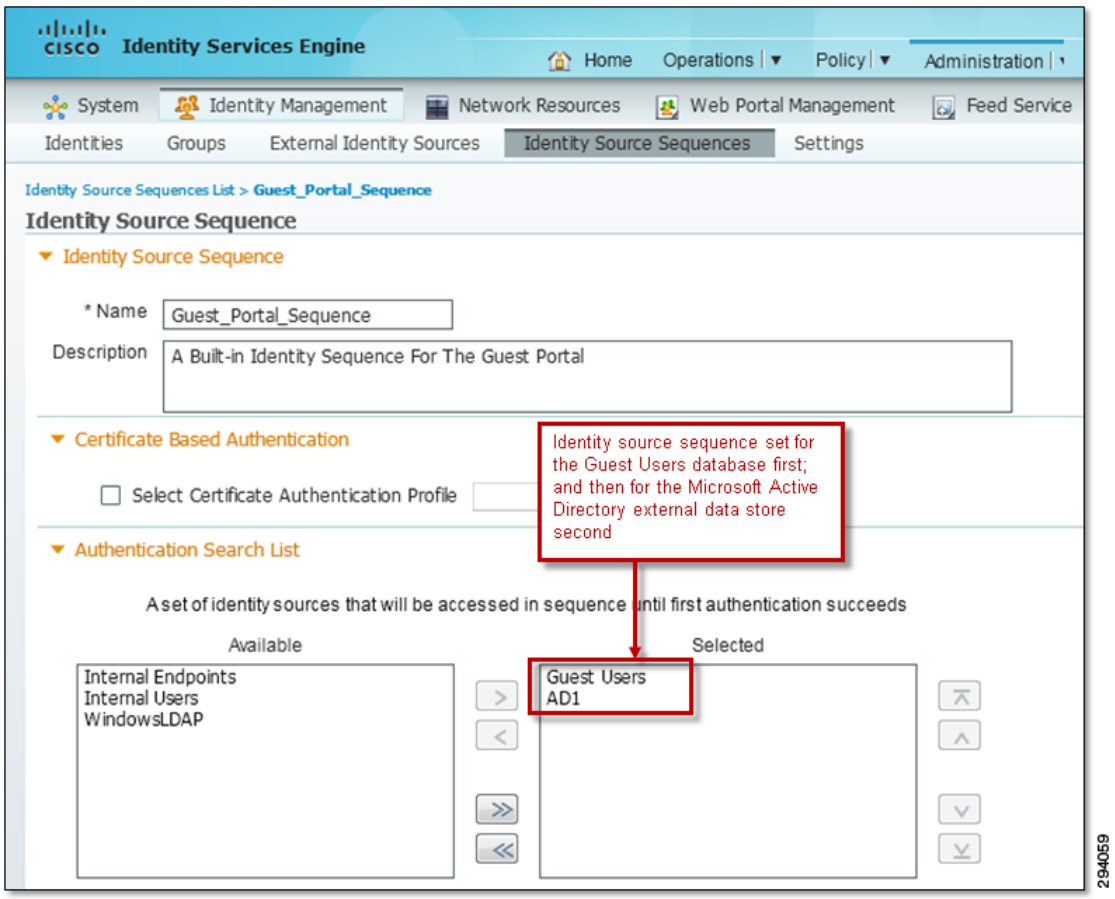
Extending Web Auth to Use Microsoft AD when Authenticating Employees with Personal Devices

The previous section discussed the most basic way of extending guest wireless access to allow employees with personal devices to access the guest network. That method simply allows employees to configure guest credentials for themselves via the Cisco ISE sponsor portal. Although this provides several advantages over methods such as utilizing a shared sponsor account on the guest wireless controller for adding credentials, it still has several shortcomings. Employees must still provision temporary guest credentials for themselves. Finally, there is nothing preventing the employee from sponsoring a true guest, other than corporate business policy.

An alternative means of providing access to the guest wireless network for employee personal devices is to simply allow the ISE server to check multiple identity sources for credentials when performing web authentication (Web Auth). For example, the ISE server could first check its internal identity groups (local database) for guest credentials. If the credentials are not found there, then check the Microsoft AD external identity store to see if the person accessing the guest network is an employee instead of a guest.

Cisco ISE Policy Configuration in Chapter 21, "BYOD Guest Wireless Access" discusses the use of a user-defined identity source sequence called Guest_Portal_Sequence for authenticating guest access via Web Auth. The Guest_Portal_Sequence uses the Internal Users identity source only. This can easily be extended by adding a Microsoft Active Directory (AD1) external identity source to the sequence, as shown in Figure 18-4.

Figure 18-4 Example of Guest_Portal_Access Identity Source Sequence Extended to Include AD



This configuration now allows employees who are in the Microsoft Active Directory database to access the guest wireless network from their personal devices once they have performed web authentication and accepted the Acceptable Use Policy (AUP) or End User Agreement (EUA).

Note This allows employees to access the guest wireless network from corporate-owned devices as well, since the authentication and authorization decision is based on Microsoft Active Directory userid and password only. The device itself is not considered within the authentication and authorization decision.

Deploying Guest-Like Wireless Access for Employee Personal Devices

The previous sections discussed options for extending access to the wireless guest network for employees with personal devices, either by allowing employees to configure guest credentials for themselves or by extending Web Auth to also check the Microsoft AD database where employee credentials are kept. With these options, employee personal devices share the same wireless SSID as guest devices. Employee personal devices also share the same IP subnet address space as guest devices, since they terminate on the same DMZ segment of the ASA firewall. Essentially the employee's personal device is treated as a guest on the network, which can bring up potential concerns.

Since IP subnet space is shared between guest devices and employee personal devices, there can be some concern that employee personal devices could deplete the IP addressing space, limiting the ability of guest devices to gain access to the guest network or vice-versa.

The ASA firewall guest DMZ interface ingress policy can be modified to allow certain application flows inbound from the guest network to a mirror of the company website server, dedicated for employee personal devices, sitting on other DMZ segments. This is discussed further in Accessing Corporate Resources. However, the ASA firewall policy would not be able to distinguish between guest devices and employee personal devices, since they share the same IP subnet address space. Therefore, application level access control at the mirror web server itself is necessary to prevent guests from accessing services on it. Likewise, the ASA firewall guest DMZ interface ingress policy could be modified to allow traffic from a virtual client—such as a Citrix or VMware client—inbound to internal Citrix or VMware servers. Again, the ASA firewall policy would not be able to distinguish between true guest devices and employee personal devices, since they share the same IP subnet address space. Application level access control at the Citrix or VMware server would be necessary to prevent guests from accessing these servers.

Since the guest SSID is typically open with no encryption, traffic from employee personal devices will be in the clear, unless the devices use either secure applications or some form of VPN tunneling, which encrypts the traffic. Although web traffic can be made secure simply by requiring the use of HTTPS, not all applications used by employee personal devices may be encrypted. This leaves some vulnerability which must be considered by security operations personnel, especially any site that authenticates using the employee's corporate login and password. If internal company data is being accessed from employee personal devices via the guest wireless network, then that information should be encrypted to prevent eavesdropping. Allowing employee devices to launch a VPN client to establish an IPsec VPN session-either directly to the ASA firewall or out to the Internet and back to another corporate VPN concentrator-may be one alternative. Another option is the establishment of an SSL VPN tunnel to the ASA firewall for employee devices. Both of these options may also require per-user authentication.



Note Cisco's implementation of Web Auth uses HTTPS when redirecting the web session and requesting credentials.

Because of these concerns, security operations personnel may be hesitant to allow anything but Internet access for any device accessing the guest network, whether it is a true guest device or an employee personal device.

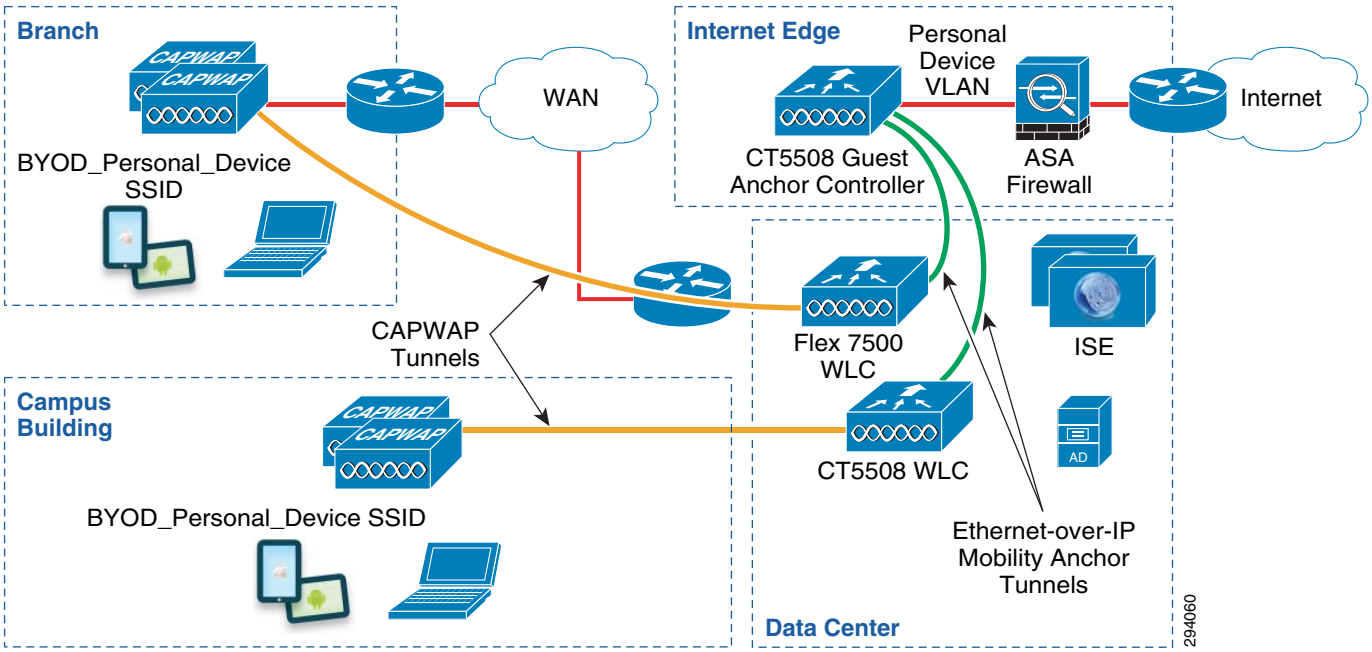
As in Chapter 21, "BYOD Guest Wireless Access," two distinct terminologies are used in this chapter. The first pair of terminologies is guest controller and campus controller. The guest controller is a dedicated controller that is used for dealing with guest and employee personal device traffic. The campus controller is dedicated for handling internal traffic. Note that the term campus controller is used somewhat generically here. The campus controller discussed within this chapter may refer to a standalone wireless controller platform deployed within a campus location or wireless controller functionality integrated within Catalyst 3850 Series switches deployed within a branch location.

The second set of terminologies is foreign controller and anchor controller. These two terminologies are used when a user roams from one controller to another controller. The new controller to which the user associates is the foreign controller. This controller anchors all the traffic to the old controller, which becomes the anchor controller. These terminologies are used in this document.

Dedicated SSID for Employee Personal Devices

This section discusses another option in which a second guest-like wireless SSID is provisioned for employee personal devices. This SSID is auto-anchored to another DMZ segment off of the ASA firewall, in a similar fashion as the wireless guest SSID. An example of this design using CUWN infrastructure is shown in Figure 18-5.

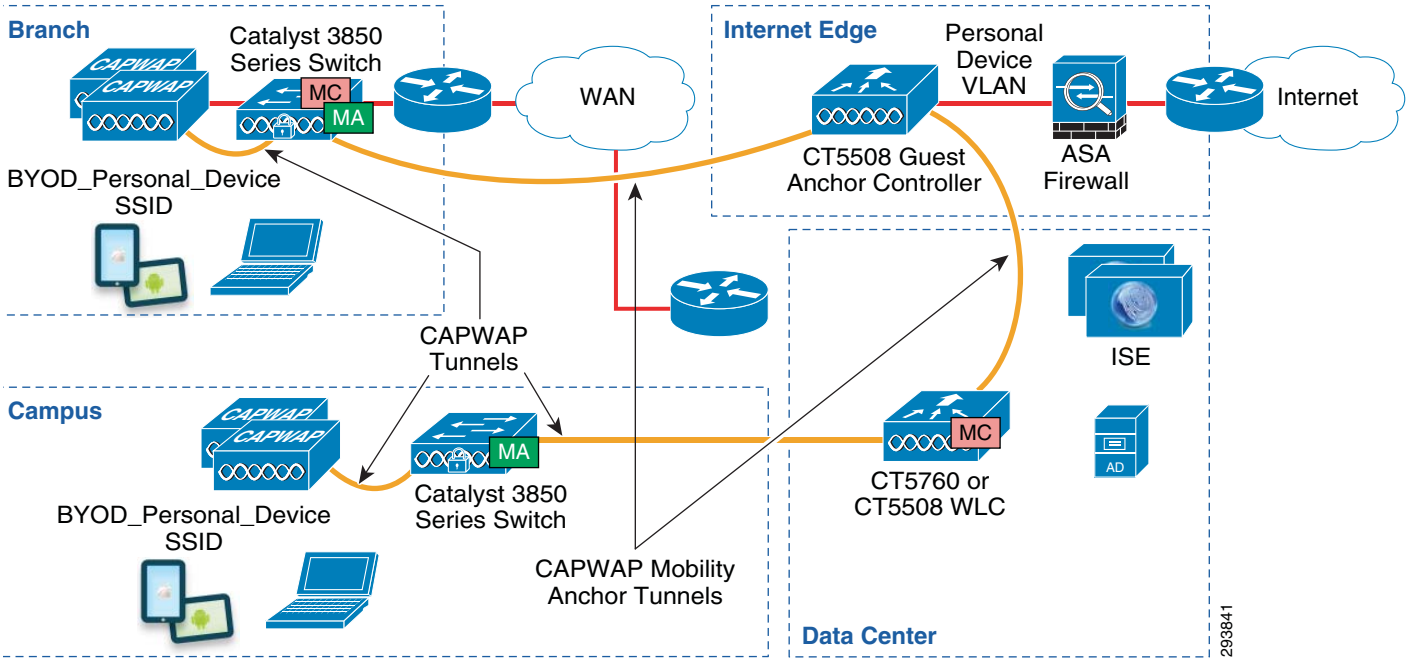
Figure 18-5 Guest-Like Wireless Access for Employee Personal Devices Using CUWN Infrastructure



With the auto-anchor mobility feature of Cisco wireless controllers, packets from the wireless client are encapsulated through a mobility tunnel between the internal wireless controller (known as the foreign controller) to the guest wireless controller (known as the anchor controller), where they are de-capsulated and delivered to the wired network.

A similar example of this design using a Converged Access wireless infrastructure is shown in Figure 18-6.

Figure 18-6 Guest-Like Wireless Access for Employee Personal Devices Using Converged Access Infrastructure



The auto-anchor mobility feature also applies to the Converged Access Infrastructure. Note that with the Converged Access branch design shown in Figure 18-6, the Catalyst 3850 switch within the branch implements both the Mobility Agent (MA) and Mobility Controller (MC) function. With the Converged Access campus design, the Catalyst 3850 switch only implements the MA function. The MC function is implemented by a dedicated CT5760 wireless controller. Since the auto-anchor tunnel is initiated from the MC within a Converged Access design, traffic is first tunneled from the campus Catalyst 3850 switch to the CT5760 wireless controller, where it is then auto-anchored to the guest anchor controller within the DMZ. For additional discussion regarding the MA and MC functionality within Converged Access infrastructure, see Configuring the Infrastructure.

Note In this version of the design guide, it is assumed that employees with personal devices would need to manually associate with this SSID. Future versions may investigate other alternatives that make use of RADIUS change-of-authentication (CoA) functionality or device profiles.

An advantage of implementing a dedicated SSID for employee personal devices is that the SSID does not have to be configured with open access and can be encrypted, unlike the guest SSID discussed throughout this design guide. Instead, the employee personal device SSID can be secured via mechanisms such as 802.1X authentication and WPA-2/AES encryption to prevent eavesdropping of traffic from employee personal devices. Employees can be authenticated via the Cisco ISE server using the external Microsoft AD identity source as they associate with the SSID.

Another advantage of implementing a dedicated SSID for employee personal devices is that the devices can be isolated from guest devices by provisioning separate VLANs for each SSID. Separate DMZ segments-implemented as separate physical interfaces off of the ASA firewall or as separate VLAN sub-interfaces off a single physical interface of the ASA firewall-can be deployed. Each DMZ interface now has a separate IP subnet address space and a separate access-control policy. This expands the IP addressing space deployed for guest devices as well as employee personal devices and removes the issue of employee personal devices causing IP address starvation issues for guest devices and vice-versa. The guest DMZ can be configured to allow only Internet access for guest devices. The employee personal

device DMZ can be configured to allow Internet access, as well as access to a mirror web server sitting on another DMZ segment. Additional access can be allowed by modifying the ASA firewall personal device DMZ interface ingress policy to allow traffic from a virtual client-such as a Citrix or VMware client-inbound to internal Citrix or VMware servers. Access can also be extended by allowing employee personal devices to launch a VPN client to establish an IPsec VPN session, either directly to the ASA firewall or out to the Internet and back to another corporate VPN concentrator. Another option is the establishment of an SSL VPN tunnel to the ASA firewall for employee personal devices.

Wireless Controller Configuration

To deploy this option in which a second guest-like wireless SSID is provisioned for employee personal devices, both the campus (foreign) wireless controllers and the guest (anchor) controller need to be configured with a new WLAN for employee personal devices. This WLAN has a unique SSID different from the corporate WLAN and the guest WLAN.

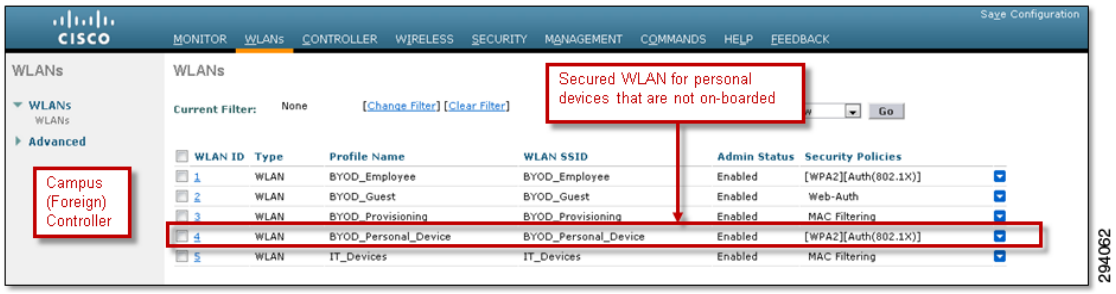
Campus Controller Configuration

This section discusses configurations when using either CUWN wireless controllers or Converged Access (IOS XE based) wireless controllers for the campus controller.

CUWN Wireless Controller Configuration

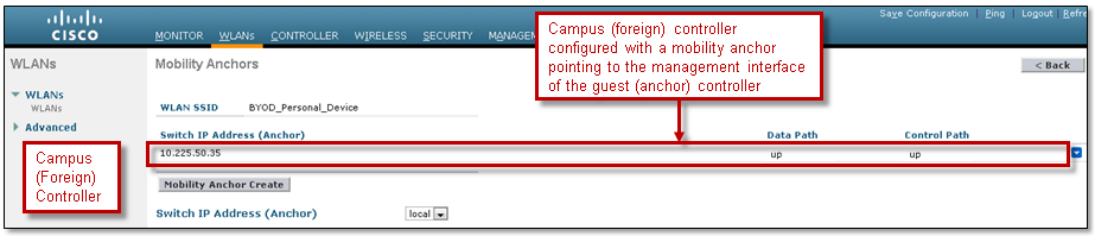
An example of a CUWN campus controller configuration is shown in Figure 18-7, in which the new WLAN is called the BYOD_Personal_Device WLAN.

Figure 18-7 Example Configuration of a CUWN Campus Wireless Controller for the Employee Personal Devices WLAN



The WLAN is configured for WPA2 security with AES as the cipher, along with 802.1X authenticated key management. Next, the WLAN on the campus (foreign) controller needs to be configured with a mobility anchor pointing at the management interface of the guest (anchor) controller. An example is shown in Figure 18-8. Note that in branch scenarios which utilize a FlexConnect wireless design, the controller servicing branch APs will be the foreign controller for the branch personal device WLAN.

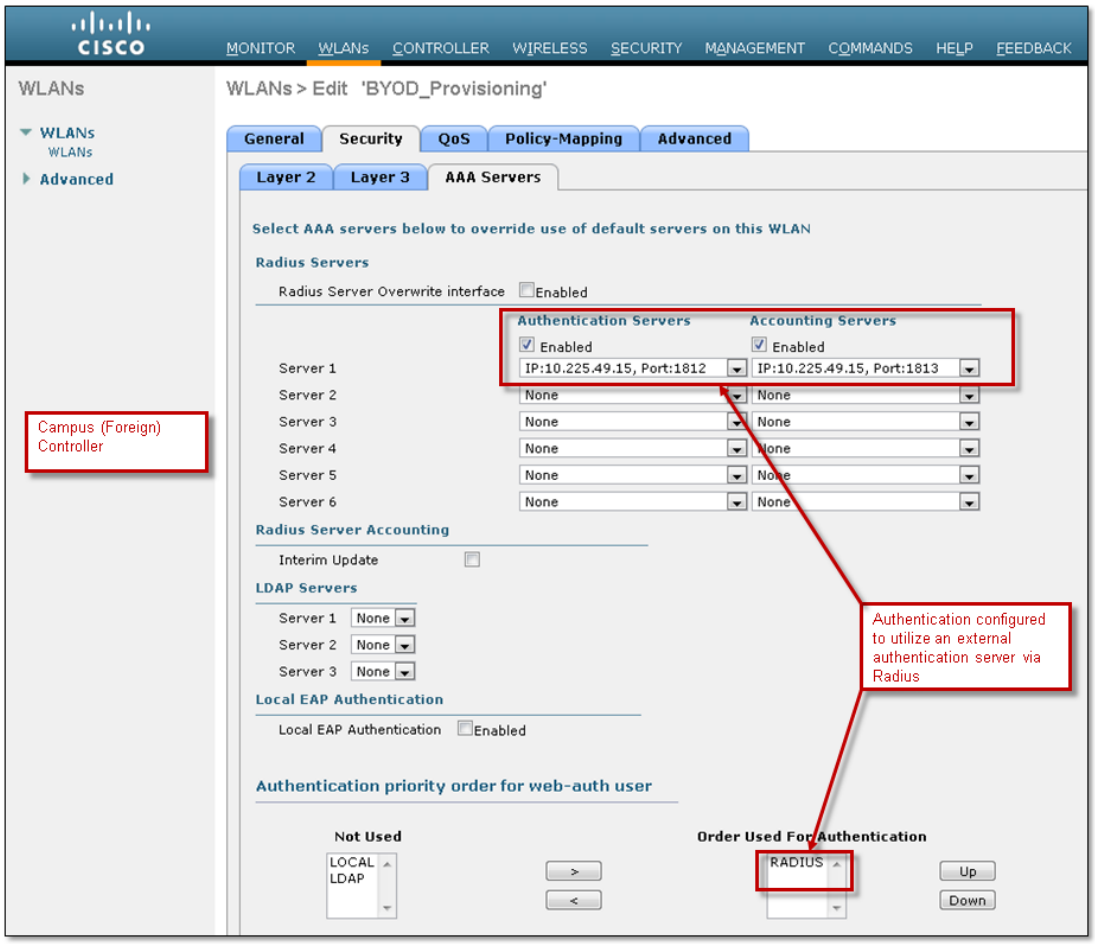
Figure 18-8 Example Configuration of the Mobility Anchor on a Campus CUWN Wireless Controllers



Note that the campus controller and the guest controller must be part of the same mobility group before the mobility anchor can be configured. The mobility anchors establish the mobility tunnel through which packets from the wireless client are automatically encapsulated and sent from the campus controller to the guest controller where they are de-capsulated and delivered to the wired network.

The network administrator must also configure the employee personal devices WLAN to use RADIUS within the campus controller for authentication. This is shown in Figure 18-9.

Figure 18-9 Authentication via RADIUS for the Employee Personal Devices WLAN on the Campus Controller



Since Web Auth is not involved with this option, there is no redirection of web sessions to a guest portal and hence no configuration required within the guest controller for Web Auth or any requirement for a pre-authentication ACL. When an employee personal device connects to the SSID, the RADIUS session is initiated by the campus controller management interface to the Cisco ISE server for authentication and authorization. Upon successful authentication, the employee's personal device is then anchored to the guest controller.

Converged Access (IOS XE Based) Wireless Controller Configuration

The following partial configuration shows an example of the employee personal devices WLAN along with part of the AAA server configuration on an IOS XE based wireless controller.

```
!
aaa new-model
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 032A4802120A701E1D5D4C
!
aaa session-id common
!
dot1x system-auth-control
!
ip http server
!
radius-server attribute 6 on-for-login-auth
radius-server attribute 8 include-in-access-req
radius-server attribute 25 access-request include
radius-server attribute 31 send nas-port-detail
radius-server dead-criteria time 5 tries 3
radius-server host 10.225.49.15 auth-port 1812 acct-port 1813 key 7 1237161E060E5D56797F71
!                   /Radius server in the line above points to ISE
!
wlan BYOD_Personal_Device 4 BYOD_Personal_Device/Defines the personal devices WLAN and
SSID
 client vlan Guest/Static assignment to non-routed (isolated) VLAN
 mobility anchor 10.225.50.35/Creates an anchor tunnel to the guest wireless controller
 session-timeout 1800
 no shutdown    /Enables the employee devices WLAN
!
```

By default WPA2 with AES as the cipher is enabled, along with 802.1X authenticated key management. Hence they do not show up in the configuration. The administrative-level **show wlan id <*wlan ID number*>** command can be used to display additional details regarding the configuration of the WLAN, including default values which do not appear within the configuration.

The employee devices WLAN must be configured on the device which functions as the Mobility Agent (MA) and on the device which functions as the Mobility Controller (MC). Therefore the configuration above is basically the same, regardless of whether a Catalyst 3850 Series switch is deployed as both the MA and MC within a branch deployment, a Catalyst 3850 Series switch is deployed as only an MA within a campus deployment, or a CT5760 wireless controller is deployed as an MA and MC within a campus deployment.

The Guest client VLAN in the configuration above is a VLAN which is isolated on the CT5760 wireless controller or Catalyst 3850 Series switch. It is not trunked to the adjacent Layer 3 device. This isolates any guest devices should the CAPWAP tunnel between the foreign and anchor controllers go down.

The following partial configuration example shows the configuration of the mobility group and mobility group members on an IOS XE based wireless controller.

```
!
wireless mobility group member ip 10.225.50.35 public-ip 10.225.50.35/Guest Controller
wireless mobility group name byod/Mobility group name
!
```

The mobility group name and mobility group peers must appear on the device which functions as the Mobility Controller (MC). Therefore, when a Catalyst 3850 Series switch is deployed as both the MA and MC within a branch deployment, the configuration must include the above two lines. A Catalyst 3850 Series switch deployed as only an MA within a campus deployment would not include the mobility group configuration. Instead the CT5760 wireless controller deployed as a MA and MC within a campus would contain the mobility group configuration. Note that since IOS XE based wireless controllers only support the new hierarchical mobility architecture; no configuration is required to enable it.

Note Cisco wireless controllers currently support two different mobility architectures. The old mobility architecture relies on Ethernet-over-IP tunnels between wireless controllers. The new mobility architecture, also called the hierarchical mobility architecture, relies on CAPWAP tunnels between wireless controllers. The two mobility architectures are not compatible with each other. If mobility (including the auto-anchoring function) is required between wireless controllers, all wireless controllers must be running either the new mobility or the old mobility architecture. The new mobility architecture is supported on Cisco 5508 and WiSM2 wireless controllers with software release 7.3.112 and on the Cisco 5508, WiSM2, and 2504 wireless controllers with software release 7.5. The new mobility architecture is supported on the Cisco 5760 wireless controller and the Catalyst 3850 Series switch with IOS XE software releases 3.2.0SE and 3.2.2SE. CUWN wireless controller release 7.4 and releases below 7.3.112 support only the old mobility architecture. The Cisco Flex 7500, 8500, and vWLC do not support the new mobility architecture. IOS XE based wireless controllers do not support the old mobility architecture. Hence if a network contains both Flex 7500 wireless controllers and Converged Access controllers, then separate sets of guest wireless controllers must be deployed with the DMZ to support both mobility architectures with the guest wireless design discussed in this design guide.

Guest Controller Configuration

The guest controller is the point where all the employee device traffic is terminated. For this version of the design guide, the discussion only includes a CT5508 CUWN wireless controller as the guest controller. An example of a guest controller configuration for the employee personal devices WLAN is shown in Figure 18-10.

Figure 18-10 Example Configuration of a Guest Wireless Controller for the Employee Personal Devices WLAN



As can be seen, the configuration of the WLAN on the guest controller must match the configuration of the WLAN on the campus controller.

The guest controller needs to be configured with a mobility anchor pointing at itself. An example is shown in Figure 18-11.

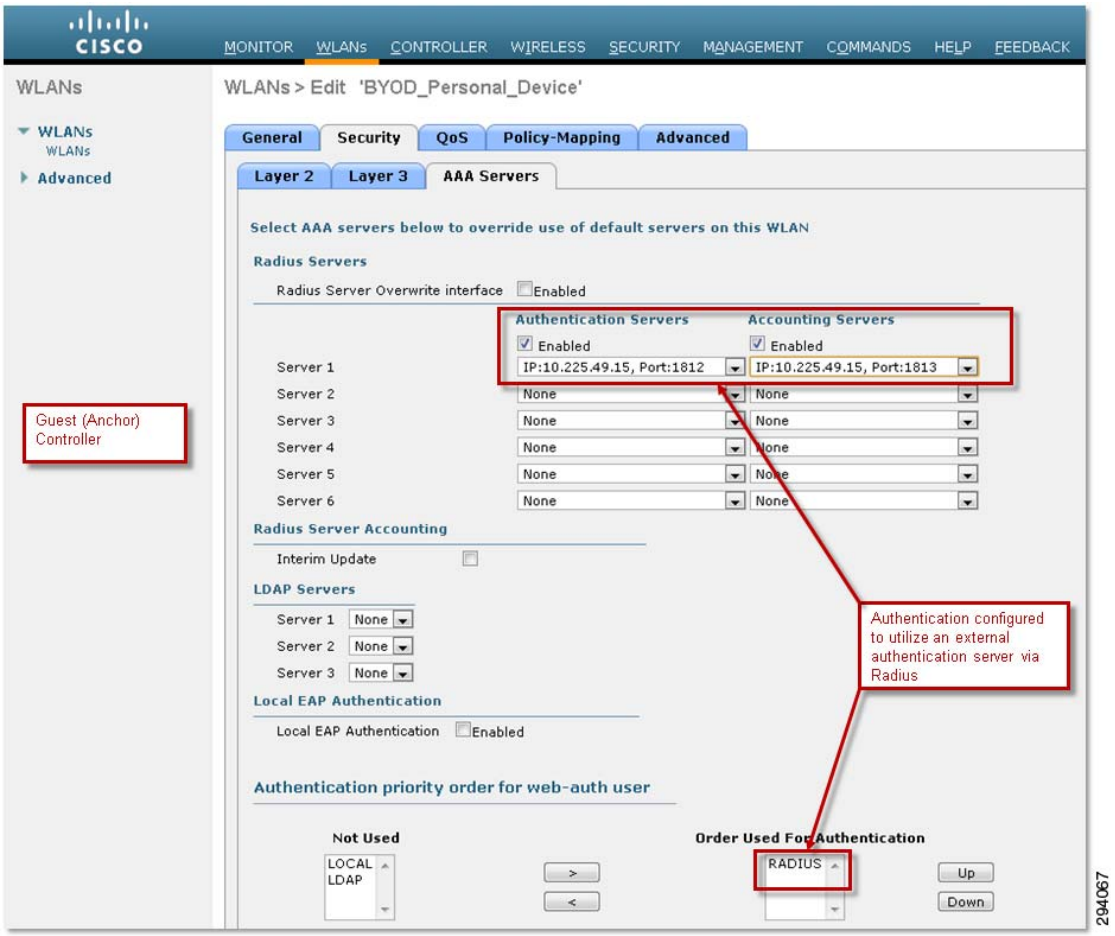
Figure 18-11 Example Configuration of the Mobility Anchor on a Guest CUWN Wireless Controller



Again, this assumes the campus controller and the guest controller are configured to be part of the same mobility group before the mobility anchor is configured.

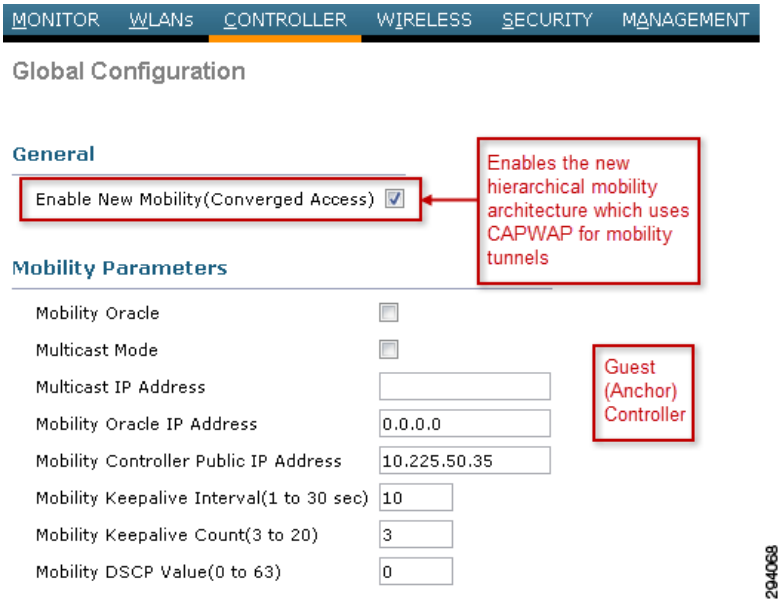
The network administrator must also configure the employee personal devices WLAN to use RADIUS within the guest controller for authentication. This is shown in Figure 18-12.

Figure 18-12 Authentication via RADIUS for the Employee Personal Devices WLAN on the Guest Controller



Finally, in order to support the new mobility architecture (also referred to as the hierarchical mobility architecture), the network administrator must check the Enable Hierarchical Architecture option within the global mobility configuration of both the campus and guest controller. An example is shown in Figure 18-13.

Figure 18-13 Enabling the Hierarchical Mobility Architecture

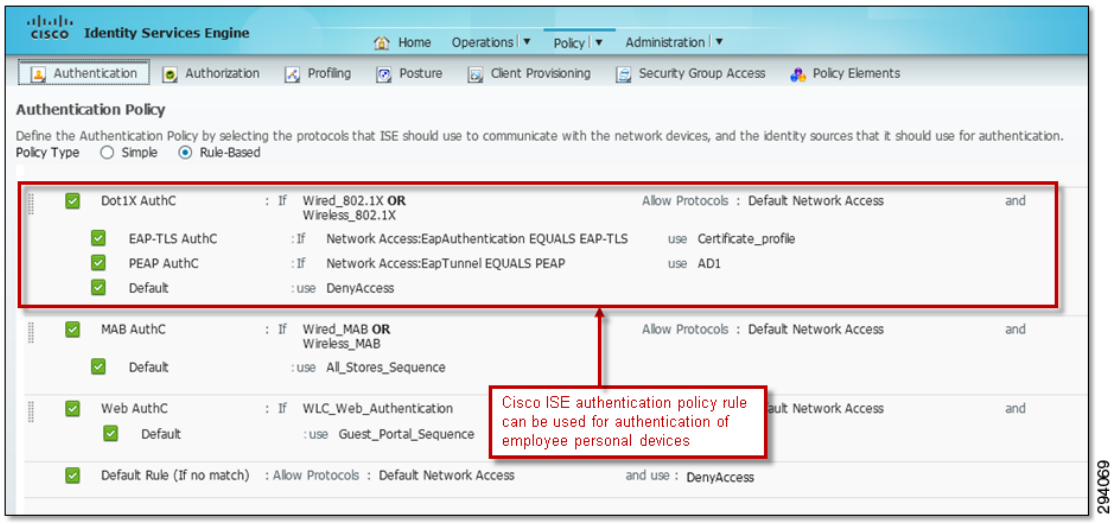


Note Since the Flex 7500 wireless controller does not support the new mobility architecture, this step can be skipped in deployments involving a Flex 7500 as the branch wireless controller.

Cisco ISE Policy Configuration

From a Cisco ISE policy perspective, the existing authentication rule—which is used to support both the on-boarding of corporate-owned (IT-managed) devices and the authentication of on-boarded corporate-owned devices within the Limited Access use case—can also be used to support wireless employee personal devices for the Basic Access use case. An example of such a policy rule is shown in Figure 18-14.

Figure 18-14 Example of Cisco ISE Authentication Policy Allowing Wireless Employee Personal Device Access



The logical format of the authentication policy rule for this example is:

```
IF (Wired_802.1X OR Wireless_802.1X)
    THEN (Allow Default Network Access AND
        IF Network Access:EapAuthentication EQUALS EAP-TLS USE Certificate_Profile
        IF Network Access:EapTunnel EQUALS PEAP USE AD1
        ELSE Default EQUALS DenyAccess)
```

Wired_802.1X is a system-generated compound condition that is used here to match 802.1X-based authentication requests from switches. It matches the following two standard RADIUS dictionary attribute-value (AV) pairs:

```
Service-Type - [6] EQUALS Framed
NAS-Port-Type - [61] EQUALS Ethernet
```

Wireless_802.1X is a system-generated compound condition that is used here to match 802.1X-based authentication requests from Cisco wireless controllers. It matches the following two standard RADIUS dictionary attribute-value (AV) pairs:

```
Service-Type - [6] EQUALS Framed
NAS-Port-Type - [61] EQUALS Wireless - IEEE 802.11
```

Default Network Access is a system-generated authentication result, which allows various EAP protocols to be used for authentication.

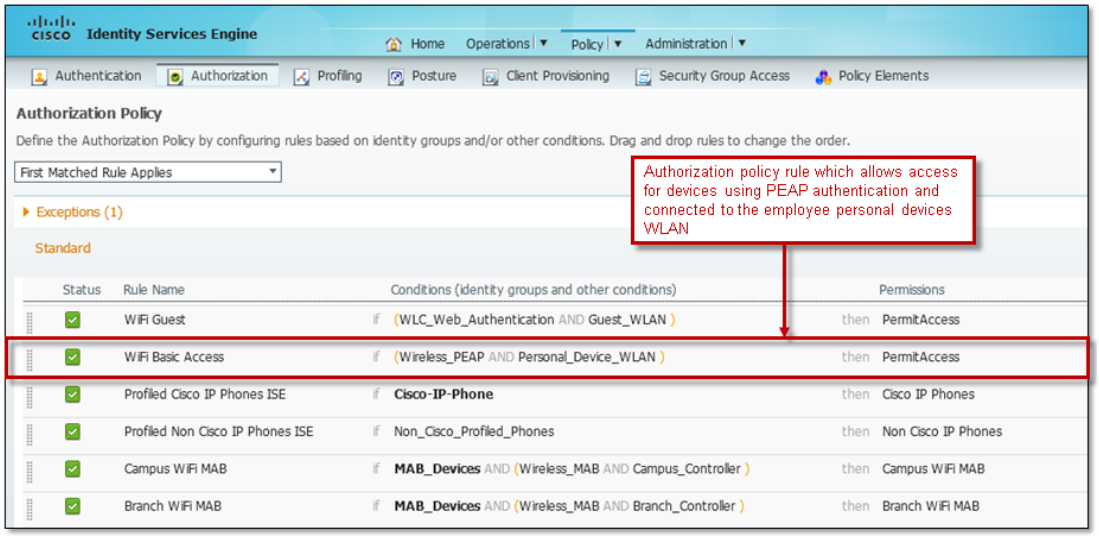
Certificate_Profile is a user-defined Certificate Authentication Profile configured under the External Identity Sources section of the Cisco ISE server.

AD1 corresponds to the Microsoft Active Directory identity store, where employee credentials are typically held within an organization.

For the Basic Access use case, wireless employee personal devices which use PEAP for authentication will match on the second condition—**IF Network Access:EapTunnel EQUALS PEAP USE AD1**—causing these devices to proceed to the authorization phase. Note that the example above works for employee personal devices which utilize PEAP authentication only. If the customer requires other authentication or EAP methods, additional conditions would need to be added to the above authentication policy rule. These conditions could also use the AD1 Microsoft Active Directory identity store to verify employee user credentials. Alternatively, the default condition could be changed from denying access to also using the AD1 identity store.

From a Cisco ISE policy perspective, an additional authorization rule needs to be added in order to support the Basic Access use case. This rule permits access for devices using PEAP authentication, which originate from the SSID corresponding to the employee personal devices WLAN. An example of the policy rule is shown in Figure 18-15.

Figure 18-15 Example of Cisco ISE Authorization Policy Allowing Wireless Employee Personal Device Access



The logical format of the authorization policy rule for this example is:

```
IF (Wireless_PEAP AND Personal_Device_WLAN)
    THEN PermitAccess
```

Wireless_PEAP is a user-defined compound authorization condition that is used here to match authentication requests from wireless PEAP devices. It matches the following two standard RADIUS dictionary attribute-value (AV) pairs, along with a Network Access condition which specifies the use of PEAP.

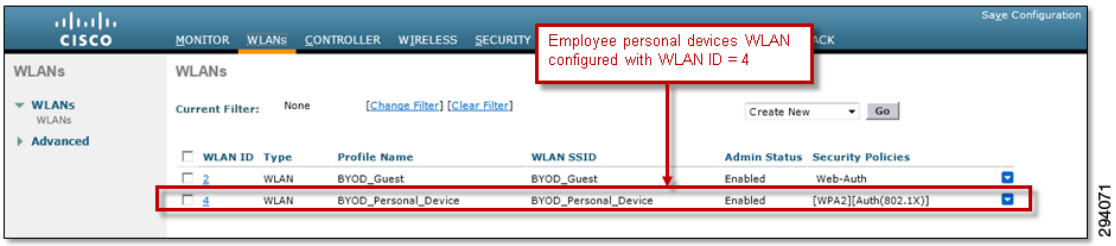
```
Service-Type - [6] EQUALS Framed
NAS-Port-Type - [61] EQUALS Wireless - IEEE 802.11
Network Access:EapTunnel EQUALS PEAP
```

Personal_Device_WLAN is a user-defined simple authorization condition for employee personal devices that access the network via the WLAN corresponding to the employee personal devices SSID. It matches the following RADIUS AV pair from the Airespace dictionary:

```
Airespace-Wlan-Id - [1] EQUALS 4
```

The Airespace-Wlan-Id is the identification number (WLAN ID) of the WLAN corresponding to the employee personal devices SSID. This is shown in Figure 18-16.

Figure 18-16 Example Wireless Controller WLAN IDs Showing Employee Personal Devices WLAN and SSID

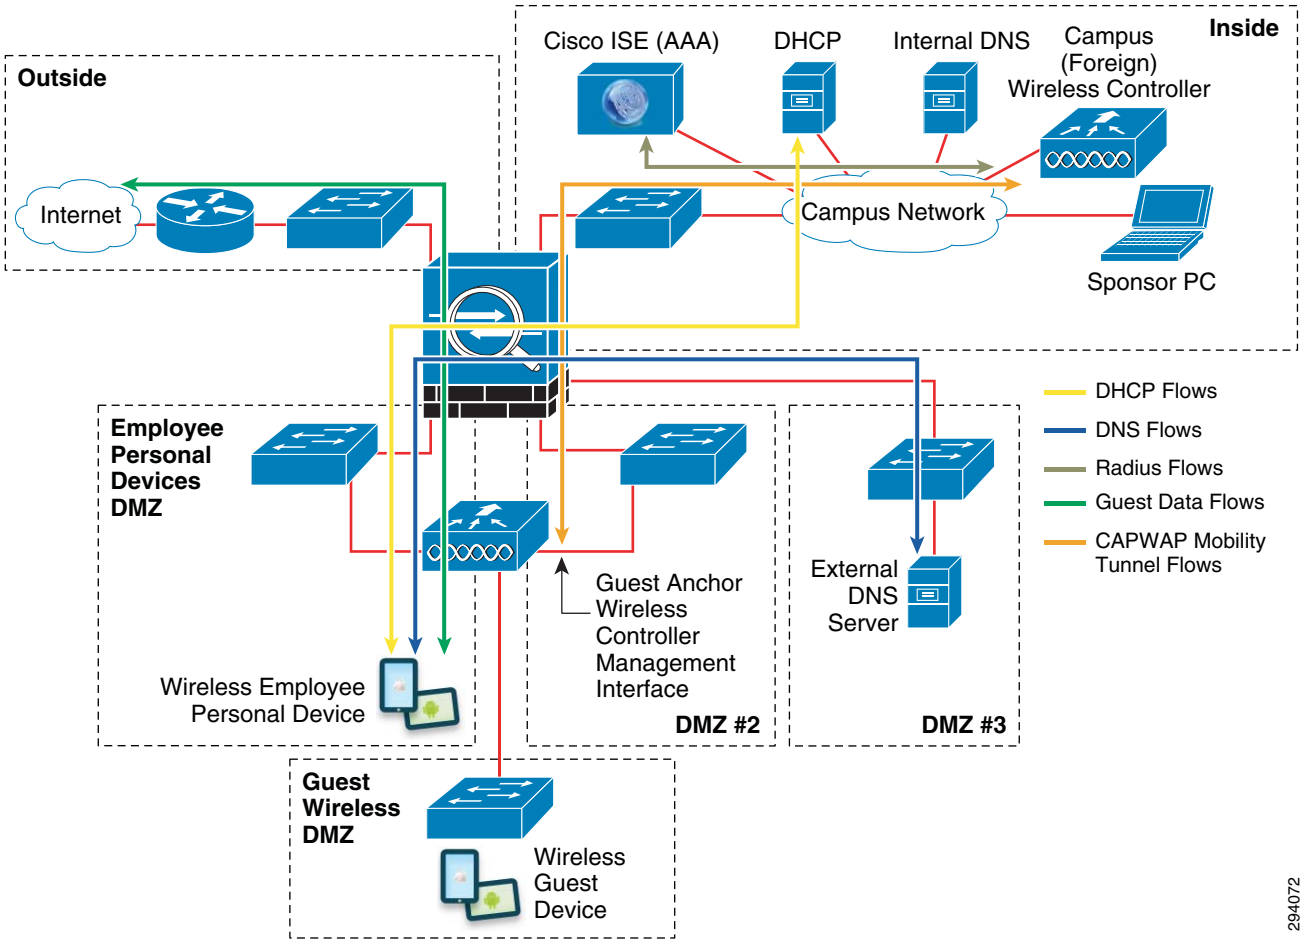


Note that this WLAN ID must be consistent across the entire BYOD deployment. This rule allows the ISE authorization policy to differentiate 802.1X authentication requests coming from the employee personal devices SSID and simply permit access. A WLAN ID cannot be changed on an existing WLAN. To change a WLAN ID, the WLAN must be removed and created again.

ASA Firewall Configuration

Figure 18-17 shows an example of the flows that need to pass through the Cisco ASA firewall to support this option.

Figure 18-17 Example of Flows that Need to Pass Through the Cisco ASA Firewall for Employee Personal Devices



The RADIUS session is initiated by the campus (foreign) wireless controller management interface to the Cisco ISE server for authentication and authorization. Therefore it does not need to be allowed through the ASA firewall for employees who are authenticating with personal devices. If using the newer hierarchical mobility architecture (as indicated in Figure 18-17), a CAPWAP auto-anchor mobility tunnel (UDP ports 5246 and 5247) between the management interfaces of the two wireless controllers must be allowed through the ASA firewall. If using the older mobility tunnel architecture, an Ethernet-over-IP (IP protocol 97) auto-anchor mobility tunnel, as well as the WLAN control port (UDP port 16666) between the management interfaces of the two wireless controllers, must be allowed through the ASA firewall.

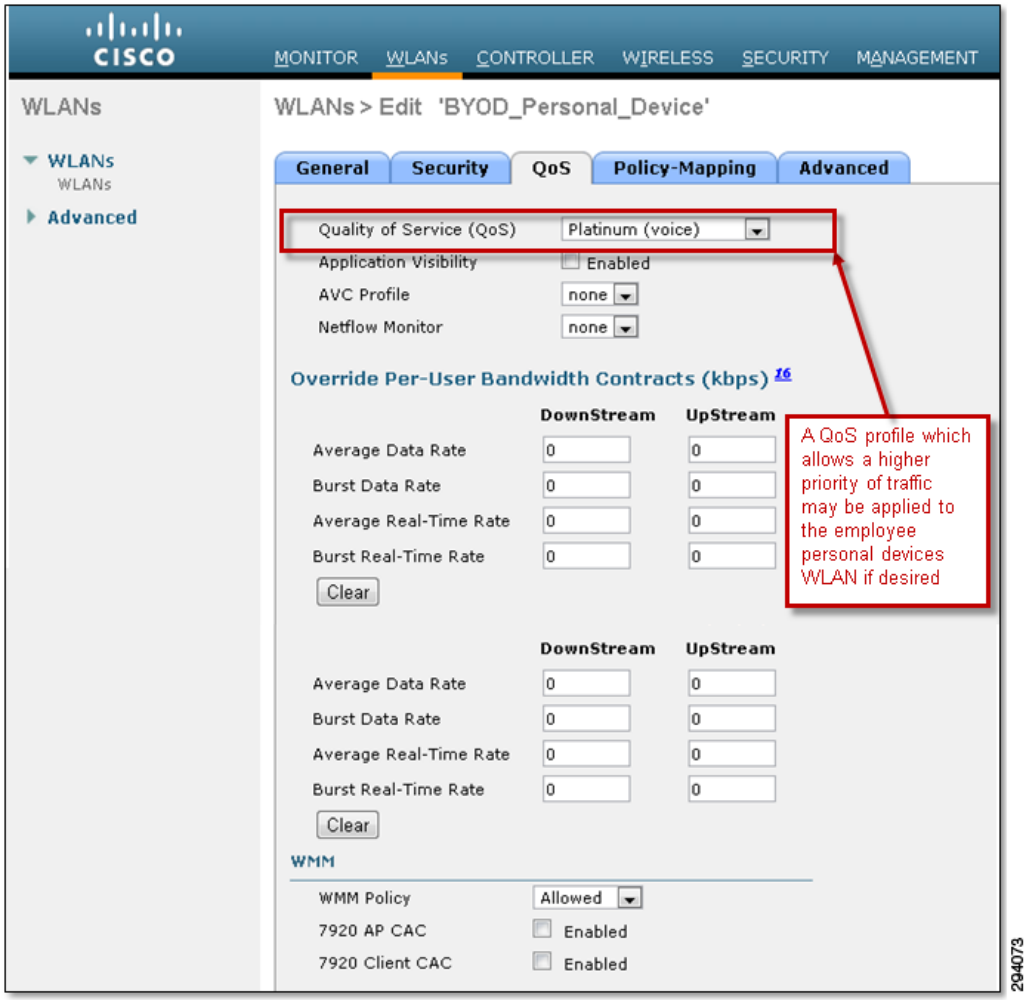
Besides allowing DNS and DHCP (assuming the deployment of an internal DHCP server), the ASA firewall should be configured to block all other traffic generated from employee personal devices onto the internal network. Additional ports can be opened to accommodate the access of corporate resources as discussed in the following section and summarized in Table 21-2.

Note Additional ASA firewall ports may need to be opened to accommodate guest wireless access, depending upon the deployment model discussed in Chapter 21, "BYOD Guest Wireless Access."

Differentiated Quality of Service Treatment

With this deployment model, a separate QoS policy can be applied to employee personal devices, different from guest wireless devices. This is because wireless employee personal devices are terminated on a separate WLAN from wireless guest devices. For CUWN wireless controllers, as of software version 7.2, QoS is applied per WLAN by way of a profile, as shown in Figure 18-18.

Figure 18-18 Example of QoS Profile Applied to a WLAN



Note IOS XE based wireless controllers support much more extensive QoS capabilities. However they also have the ability to support similar QoS per WLAN based on the older profiles model of CUWN wireless controllers. This version of the design guide does not address QoS on IOS XE based wireless controllers. Future versions will discuss this topic more thoroughly.

This may be desirable if employee personal devices are going to be allowed to run virtual desktop client applications such as a Citrix client or VMware client or if employee personal devices run collaboration clients such as Cisco Jabber.

Note Chapter 21, "BYOD Guest Wireless Access" discusses rate limiting per-SSID and per-User. These features can be used for employee personal devices as they were for guest devices. Rate limiting is configured on the campus (foreign) controller and not the guest (anchor) controller.

Accessing Corporate Resources

Because employee devices are associated to the network from a DMZ interface, which is effectively outside of the corporate firewall, they do not have access to company resources located inside the firewall. This may be perfectly acceptable and desirable. Employee devices still have access to the public Internet. This enables them to connect to cloud-based resources such as Cisco WebEx or partner websites. The employee device is afforded some level of usability, making the device useful as a productivity tool.

Companies may want to offer access to additional resources but still maintain the security offered by restricting employee devices to the guest side of the firewall. There are various options available to accomplish this objective, including:

- Setting up a mirror of the company website
- Allowing VPN access
- Allowing virtual desktop client access

Securing Mirror Sites for Personal Devices

One approach to bringing services to employee personal devices accessing the guest network is to set up a mirror of the company website in another DMZ segment, referred to as an employee device security zone (EDSZ) within this section. If employee personal devices connect to the guest wireless network, the website will only be accessible after the user has completed the Web Auth process and accepted an Acceptable Use Policy (AUP) or End User Agreement (EUA). This website does not need to be an exact match of an internal website, but could contain relevant content that employees can use on their personal devices to make them more effective. In addition, the website could include content optimized for smaller mobile displays. Examples of applications that could be offered in the EDSZ include access to email, a team wiki page, or the company news site.

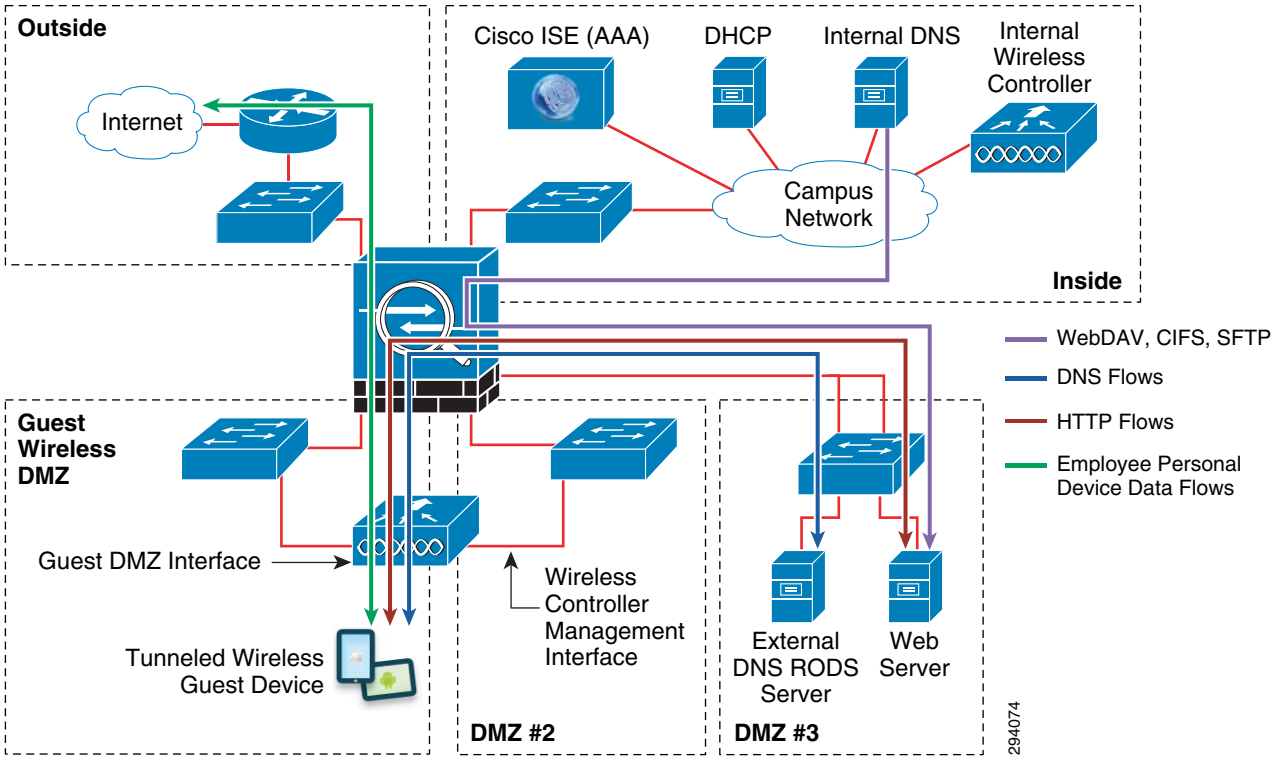
There are several methods available to setup a secure website. In general, the deployment is very similar to a typical DMZ web service, except that rather than residing in an Internet facing DMZ, the server is located on a subnet accessible by employee personal devices. The intent of this section is not to provide detailed guidance on the deployment of a presentation server, application server, and database server. Site administrators should be familiar with the approach that best fits their security environment. Some high level considerations when setting up the server include:

- Dual attachment—Usually this is considered to be more secure. The client-side NIC should implement firewall services that allow inbound connections on TCP port 80 or 443. Only session requests initiated by the client subnets should be allowed towards the server. Session requests initiated by the server should not be allowed towards the client subnets. If the employee device wireless network is encrypted, then some organizations may be comfortable allowing users to attach via HTTP.

- The back-end NIC should be used to move content between the server and the data store or application server. It may also be allowed for remote administration of the site. Single-attached servers are possible although usually a dedicated and secured gateway is setup for the server-to-server communications that is separate from the server-to-web client gateway.
- Content Delivery—Generally content is either static or dynamic. Static content can be pushed overnight or as needed to keep the website current. Dynamic content could allow the employee device to post information to the site. The website could use a local data store that is synchronized with an external data store or have a secured channel to an application server.
- User Authentication—Some method should be used to ensure only authorized and authenticated users are able to view site content. Utilizing Microsoft Active Directory or a local user database are two possible methods. Active Server pages (ASP.NET) can also be used to leverage the login controls used with Microsoft Internet Information Server (IIS) servers and provide a more sophisticated authentication model such as single sign-on (SSO). Local databases are easier to setup but require a high level of administrative overhead and are usually only appropriate in very small organizations.
- Secure Sockets—Websites in the employee device security zone (EDSZ) should implement secure socket layer (SSL) or transport layer security (TLS) if employees are sending sensitive information, such as their login credentials. This can be relaxed somewhat if the EDSZ has been implemented with wireless encryption. On the other hand, if employee devices reside in the traditional guest network where wireless packets are not typically protected with encryption and are co-mingled with actual guest traffic, then SSL/TLS websites are needed. This is particularly important if employees are passing their corporate credentials to a mirror website.
- Web Server software—There is a wide range of web server software available. Deciding which type of server to deploy impacts what security features are available. Common choices include Microsoft IIS, Apple, and Tomcat. Wikipedia has a comparison of web server software that can illustrate the choices available (http://en.wikipedia.org/wiki/Comparison_of_web_server_software). It is also possible to host sites on a secure cloud service. This service could be restricted to IP addresses or require client side certificates. With proper security precautions, a cloud-based site could also allow mobile employees access to some traditional corporate resources such as payroll or benefits that are increasingly finding their way to the cloud.

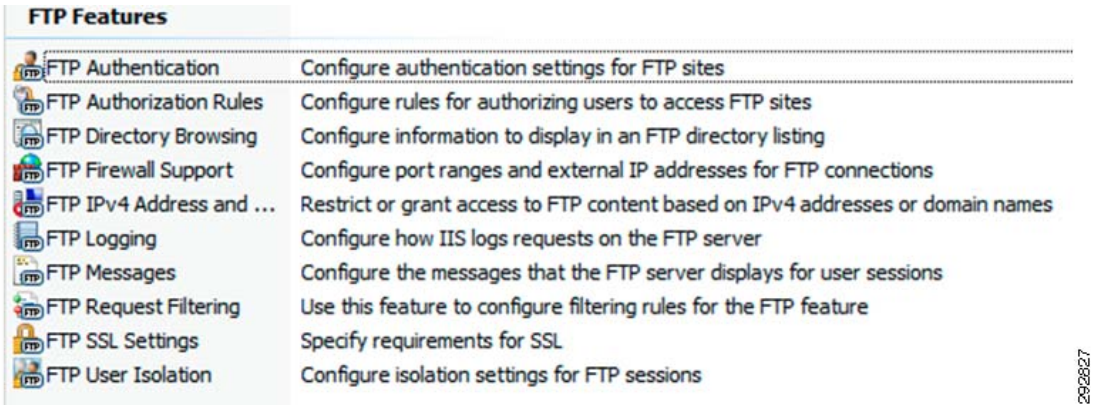
Figure 18-19 illustrates a simple scenario where static content is deployed in a DMZ dedicated for employees using personal devices. A Windows 2008 server is deployed with Microsoft IIS 7.0 as well as some other network services specific to the DMZ, such as DNS and Read-Only Directory Services (RODS). Users are authenticated against the corporate Microsoft Active Directory server. The RODS service requires DNS to be installed on the same server. The web server is typically a dedicated box. However, in some situations it may be desirable to run RODS and IIS on the same server to simplify basic authentication using Microsoft AD. It is more appropriate when supporting a small-to-modest number of employee devices.

Figure 18-19 Example of a Mirror Website for Employee Personal Devices



Moving content to the secured server can be done by various means. One approach is to use FTP, however because FTP is not secured, a better approach is to use either SFTP or FTP over SSL. By default, Microsoft IIS does not ship with a secure FTP server. Microsoft supports FTP over SSL rather than SFTP. Administrators must copy the installation package from Microsoft's website (http://learn.iis.net/page.aspx/310/what-is-new-for-microsoft-and-ftp-in-iis-7/) and install the feature on their server. If FTP is already running, the administrator needs to deselect the FTP feature from the services manager prior to installing the FTP over SSL server. The improved FTP over SSL server offers additional tools not available in the standard FTP package that are used to manage access to the FTP site, as shown in Figure 18-20.

Figure 18-20 Example of Tools Available with the FTP over SSL Server



Anonymous authentication should be disabled and at least basic authentication should be enabled. There are other options that may be appropriate from some organizations.

Instead of FTP over SSL, administrators can choose to use Web Distributed Authoring and Versioning (WebDAV). This method offers more flexibility than FTP because many operating systems allow the connection to be mounted to the file system. By providing a directory handle, web authoring applications as well as other applications can directly use the secured pipe. WebDAV is based on HTTP or HTTPS and provides the ability to authenticate and encrypt data. If employees are placed directly in the guest SSID, then WebDAV HTTPS should be used since passwords are sent. If employee devices are placed in an encrypted EDSZ, then WebDAV HTTPS could be used to provide an additional layer of encryption. The security considerations are detailed in section 20 of RFC4918.

Another option similar to WebDAV is CIFS. This protocol also allows local directory mounts of the remote site. It is commonly found in Microsoft environments, although Samba is available for non-Windows servers. Microsoft also supports SMB2 with Vista, Windows 7, and Windows 8, which is an update to CIFS. There are several other approaches that provide a secure path between either the web authoring site or the application server. A particular enterprise will likely leverage the same methods as the web servers located in the traditional DMZ.

DNS Support

The employee devices need access to a DNS server. If the EDSZ web server is using RODS, then DNS is already available on the Directory Server. It is installed by default when the RODS is setup unless the administrator explicitly chooses not to. Dynamic updates are secured and zone transfers are not enabled. If the web server is not using AD for employee authentication, then DNS will be a standalone service.

Outlook Web Access for Employee Devices

Email is a foundational service that can be offered to employee devices. This can be accomplished by deploying a web interface to the mail server such as Outlook Web Access (OWA) or ActiveSync for exchange environments. Some enterprises may already be offering this service for employees that need email while traveling. In this case, the employee devices can continue to use the current Internet facing OWA server.

Microsoft does not support the OWA server in a DMZ zone. Instead, the OWA server should be behind the firewall. Holes can be opened for port 443. Another option is to setup Apache in the EDSZ as a reverse proxy. The Microsoft recommended approach is to run OWA on the client access server (CAS) and publish the CAS with Microsoft's Internet security and acceleration (ISA) server into the EDSZ. This is a full blown deployment and may not be appropriate as a method to grant employee personal device access due to the complexity involved versus other methods, such as simply opening a hole in the EDSZ DMZ firewall for HTTPS to the enterprise CAS.

Another option is subscribing to Office365, which is Microsoft's cloud-based Exchange and Office environment. In this case, employees would use the public Internet service to gain access to their email or other cloud-based enterprise resources. At this time, there is not a native Office365 application for either Android or iOS devices and users would be restricted to HTTPS access unless they were using a Windows 8-based mobile device. This method would also allow Direct-To-Cloud over 3G/4G or off-premise accesses to the same resources. Microsoft is only one of many cloud based enterprise environments offering email services.

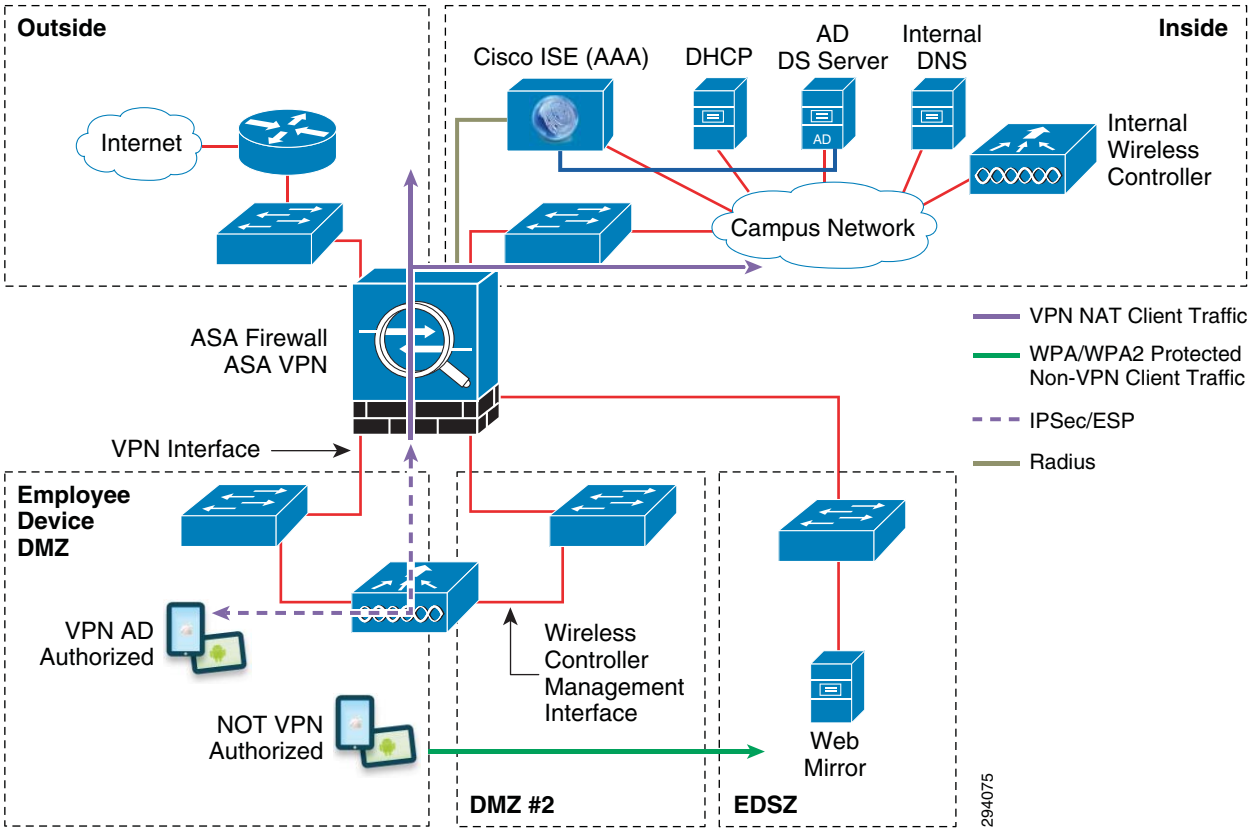
ActiveSync Support

Future versions of this document will discuss mobile device managers (MDM) that are used to manage the configuration profiles of mobile devices and provide additional security features needed for lost or stolen devices. Some of the MDM functionality is licensed from Microsoft, including remote-wipe, PIN lock enforcement, and others. ActiveSync is used to synchronize email, calendar events, and contacts between the Microsoft Exchange server and the mobile device's email application. Administrators may be interested in providing ActiveSync to devices in the employee device security zone (EDSZ). When properly configured and certified, ActiveSync can also provide the MDM security features mentioned previously. Providing this support is similar to OWA. The firewall policy in the EDSZ can be set to allow connections to ActiveSync on the CAS that may be published on an ISA server. ActiveSync is supported by WebDAV and should be used over HTTPS (TCP port 443).

VPN Client

Enterprises that restrict employee devices to a guest or dedicated EDSZ may want to allow a subset of these users the ability to launch the built-in VPN client to connect to the secured part of the network. There are two methods that can be employed. First, the device may already have access to the current Internet-facing VPN concentrator. In this case, the employee device would connect from the guest network out through the public Internet and then back in to the Internet DMZ where the VPN concentrator is located. If employee devices are co-mingled with actual guest traffic, then this may be the best approach. However, if a dedicated and secured SSID is deployed behind an ASA firewall specifically for employee personal devices, then this firewall could also provide VPN access for some privileged users. Alternatively, a dedicated VPN concentrator could be located in the EDSZ. After the device authenticates and joins the secured wireless domain, these users would connect to the VPN concentrator to gain additional secured access. Only employee devices in the dedicated security zone can reach the concentrator. The general layout of the network components are shown in Figure 18-21.

Figure 18-21 Employee Zone VPN Network Components



Apple iOS devices include a built-in Cisco IPSec client allowing ESP tunnel mode and XAUTH. Both Apple and Android devices offer L2TP with IPsec and pre-shared key (PSK). The ASA can be setup to accept both types of VPN clients. For this discussion, the focus is the Cisco VPN client found on Apple iOS devices.

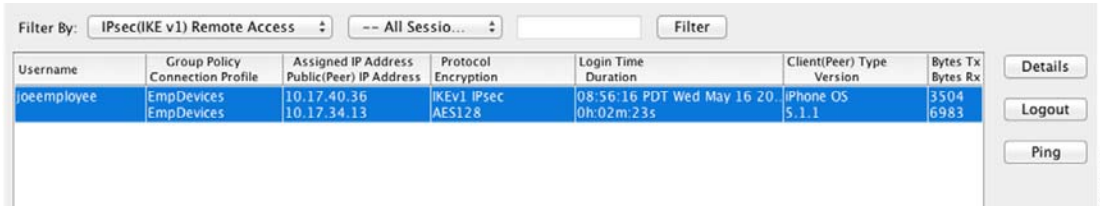
The employee needs to know the name of the VPN concentrator, group, and group secret to configure the VPN client. Future version of this document will illustrate how the VPN configuration can be pushed to the employee device without user intervention. Certificates can also be pushed to the employee device to further secure access to the VPN concentrator. The use of certificates negates the need for a group and group secret.

Figure 18-22 iOS VPN Client Configuration



When the user connects to the VPN concentrator, they are asked for their Microsoft AD credentials. This information is passed by the Cisco ASA to the Cisco ISE server, where a policy decision can be made. This decision can include attributes from Microsoft Active Directory or any of the other parameters Cisco ISE can use to determine policy. If the user is authenticated and authorized by Cisco ISE, the ASA completes the VPN connection. Once connected, the ASA can apply additional security and access restrictions to the tunnel, further controlling what resources the employee device can reach. The ASA can also be used to monitor who is using the VPN portal, as shown in Figure 18-23.

Figure 18-23 ASA Management for VPN Connections



ASA provides additional information for managing VPN connections.

Virtual Desktop Client

Another available deployment model is to allow a virtual desktop to run on the employee device. The actual applications and associated data remain on the secured hosting server. Once the device disconnects from the network, the data is typically no longer available to the user. The enterprise can control which users are able to launch a virtual desktop and what applications are available on that desktop. The firewall between the EDSZ and the hosting server can be configured to allow specific connections.

There are a wide range of possibilities. In its simplest from, the employee could use VNC to connect back to their desktop or a dedicated server. A VNC client is available for both iOS and Android devices. This may be adequate for some small environments where availability and manageability are not a top priority. By default, the connection is not encrypted, which is a concern. Administrators may not have the necessary control over which applications are available on the hosting server. Employees may be tempted to attach to their desktop and send sensitive data to an external account via email or cloud file sharing services such as Dropbox and Google Drive. This temptation only arises as a means to bypass IT policy and should be considered before allowing remote desktops to attach to employee deployed VNC servers. The use case for employee devices with virtual VNC desktops needs careful review. The employee is likely sitting in front of the actual desktop, with a full keyboard and mouse and likely does not need a remote desktop. The best approach may be to block TCP port 5900 from traversing the firewall to unknown destinations.

Cisco offers the Cisco Virtual Workspace (VXI) Smart Solution and partners with several companies that offer a virtual desktop on mobile iOS and Android devices including VMware View, Citrix, and WYSE. Virtual desktops on employee devices are best suited in VXI environments where a centralized UCS server is securing and managing sessions. With VXI in place for corporate devices, extending access to employee devices may allow productivity gains. This can be done by opening the firewall to allow a connection to specific and well known servers. Beyond BYOD, virtual desktops are compelling because of the reduction in IT costs. Adding tablet support maximizes the benefit because the requirement for employee laptops is reduced. A small and lightweight VDI hardware appliance replaces the traditional desktop, and mobile device support un-tethers the employee from the cube. Tablets with virtual desktops can offer much of the same functionality as employee laptops, but at a reduced cost and with stronger tools to address lost and stolen devices plus centralized data security inherent in VXI.

Finally AnyConnect provides a centralized virtual desktop that integrates with the ASA firewall. This is a good approach because security is the foundation of the system. AnyConnect desktops attach to the ASA firewall via SSL. Virtual desktops are evolving in capabilities and will be covered in more detail the next release of this document.

Summary

These types of alternate solutions offer a wide range of options to provide employees compute resources without compromising corporate data. Leveraging the guest environment can serve as part of a migration path to a fully certificate-based BYOD solution. Guest type deployments can be set up fairly quickly without the need to touch a large number of third-party devices and yet still meet the basic requirement of allowing employees to use their personal devices to increase the organization’s productivity.



User Experience—How To On-board a BYOD Device

Revised: September 27, 2013

The Cisco ISE allows employees to be in charge of on-boarding their own devices through a self-registration workflow and simplifies the automatic provisioning of supplicants as well as certificate enrollment for the most common BYOD devices. The workflow supports iOS, Android, Windows, and Mac OS devices and assists in transitioning these devices from an open environment to a secure network with the proper access based on device and user credentials.

The simple workflow provides a positive experience to employees provisioning their own device and allows IT to enforce the appropriate access policies.

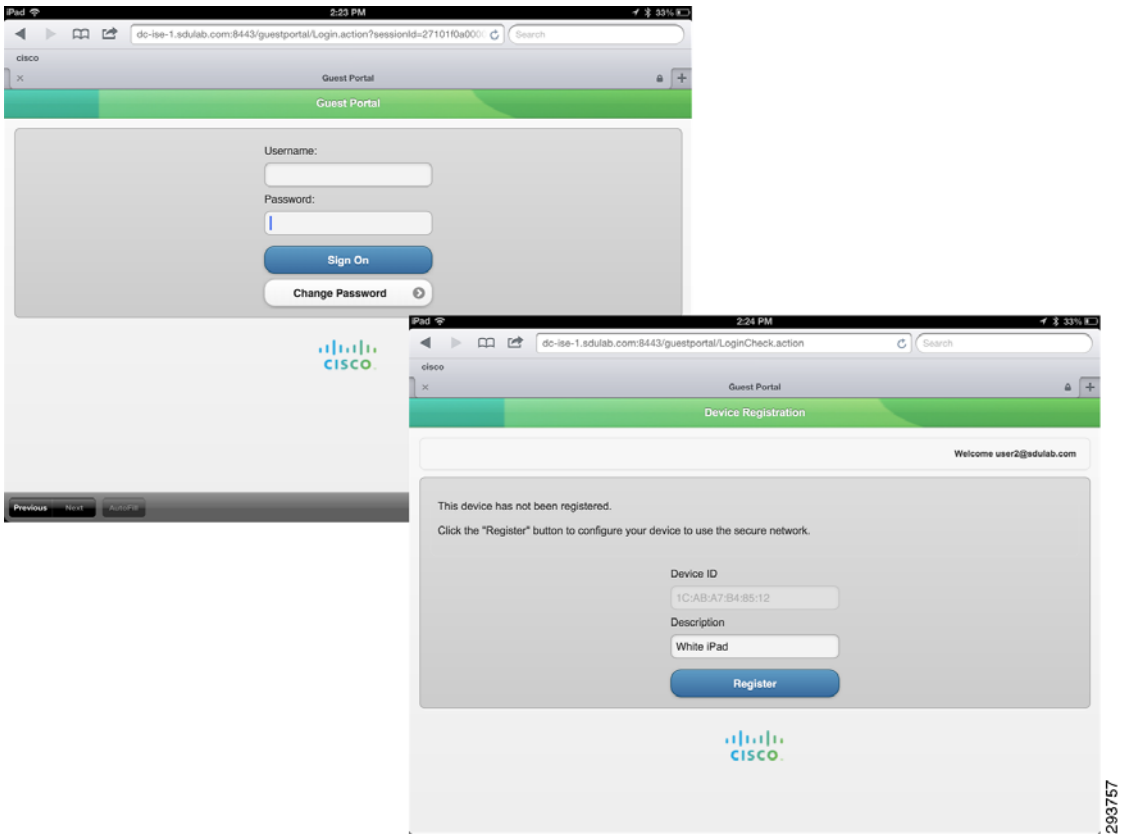
Apple iOS Devices

The employee connects to the provisioning SSID and is redirected to the Guest Registration portal for registration after opening a browser. The employee logs in using their Active Directory credentials.

If the device is not yet registered, the session is redirected to the self-registration portal, where the user is asked to enter a description for the new device. The employee is not allowed to change the Device ID (MAC address), which is automatically discovered by ISE.

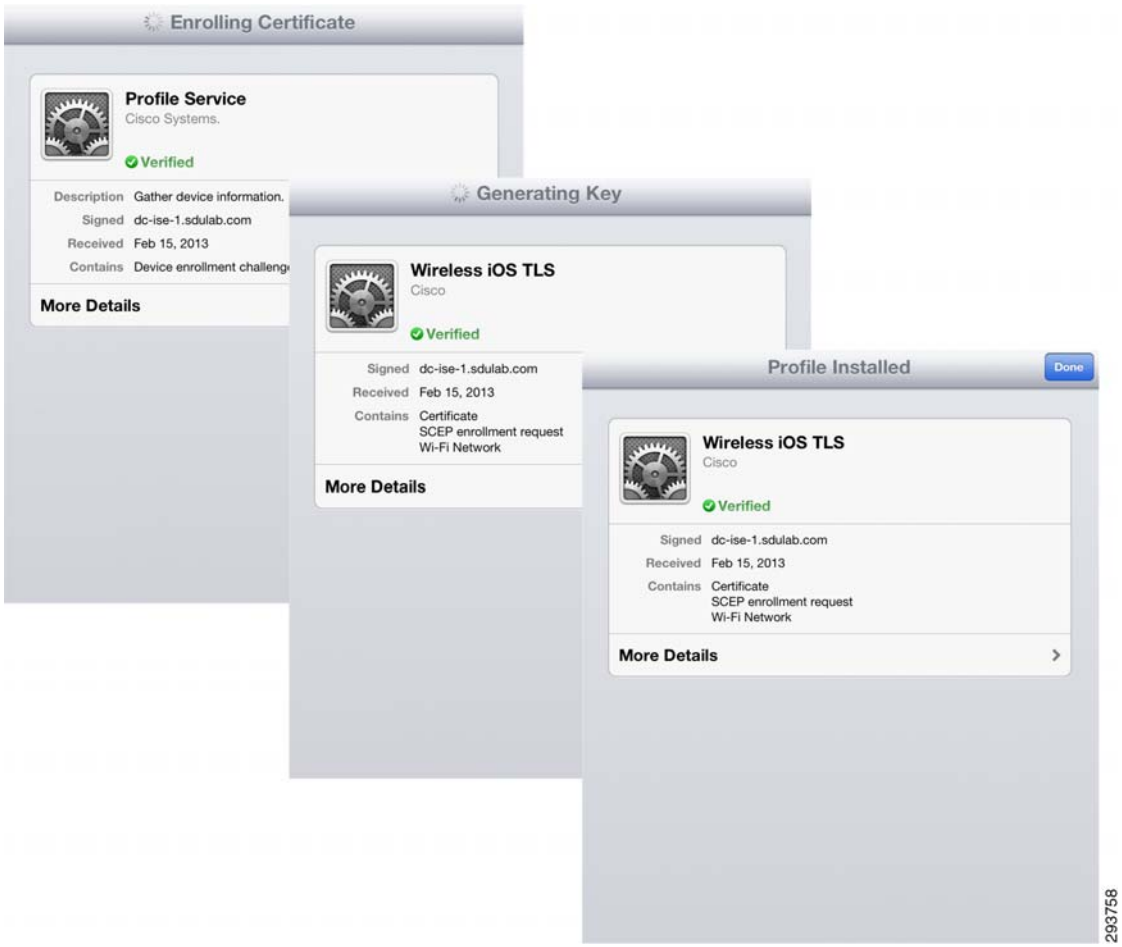
Note Cisco has been made aware of potential incompatibilities introduced by Apple iOS 7. We are working to understand the limitations and design updates will be made to this publication.

Figure 19-1 Guest Portal and Self-Registration Portal



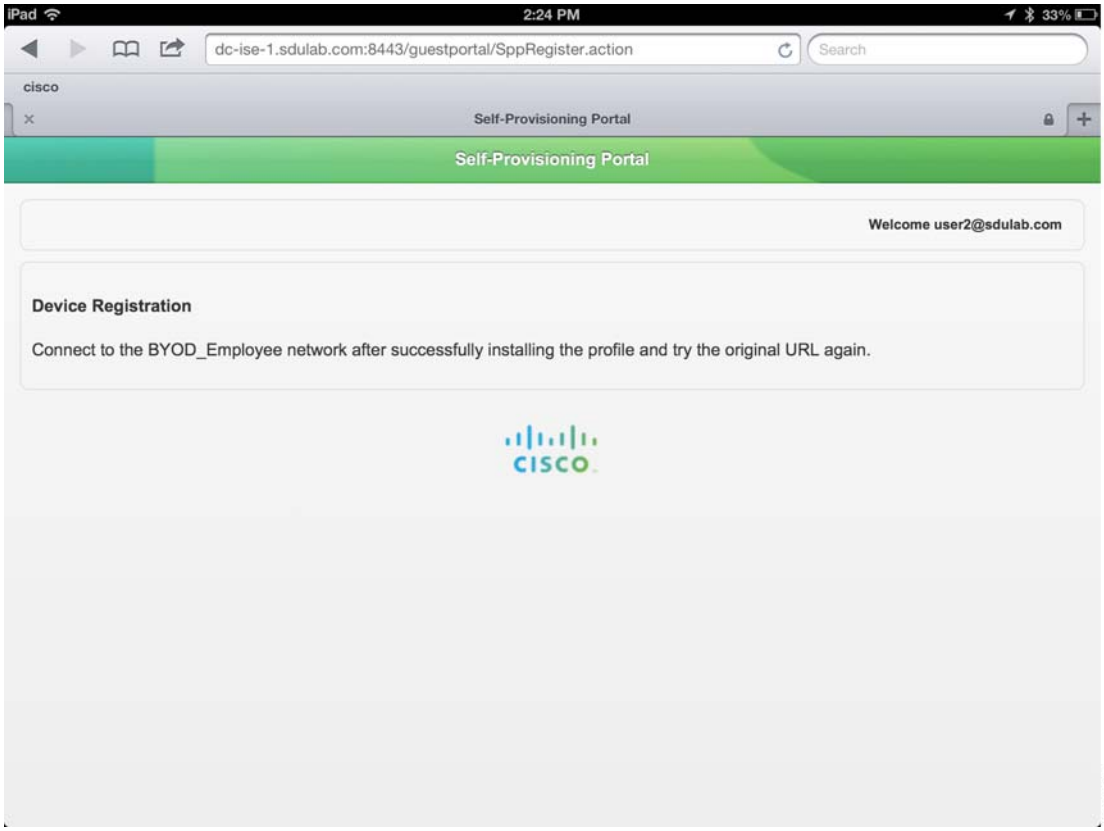
- The supplicant profile is downloaded and installed on the endpoint.
- Keys are generated and the certificate is enrolled.
- The Wi-Fi profile required to connect to the BYOD_Employee is installed.

Figure 19-2 Enrollment and Profile Installation



The employee is notified that registration is complete and is reminded to manually connect to the BYOD_Employee SSID.

Figure 19-3 Device Registration Complete



Note For iOS devices, the employee is required to connect manually to the BYOD_Employee SSID.

The certificates and profile can be viewed by clicking **Settings > General > Profiles** and selecting **Mobile Profile**. Figure 19-4 highlights the Wi-Fi profile to connect to the BYOD_Employee SSID.

Figure 19-4 Mobile Profile Details



As shown in Figure 19-5, the ISE maintains a detailed log of the authentications as they take place:

- The first log shows how the first time the device connects, the MAC address is used for authentication, and the Wireless CWA profile is used for authorization, enabling the redirection to the Guest Registration portal.
- Once the enrollment and provisioning take place, the user connects to the secure BYOD_Employee SSID. ISE grants Partial Access to the device.

Figure 19-5 ISE Authentications Log

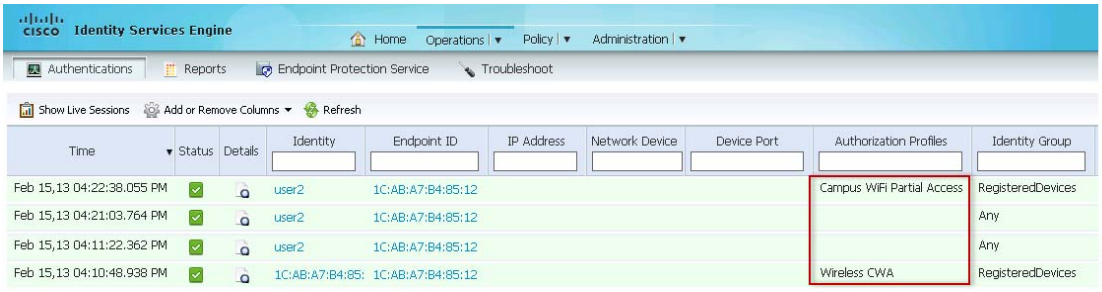
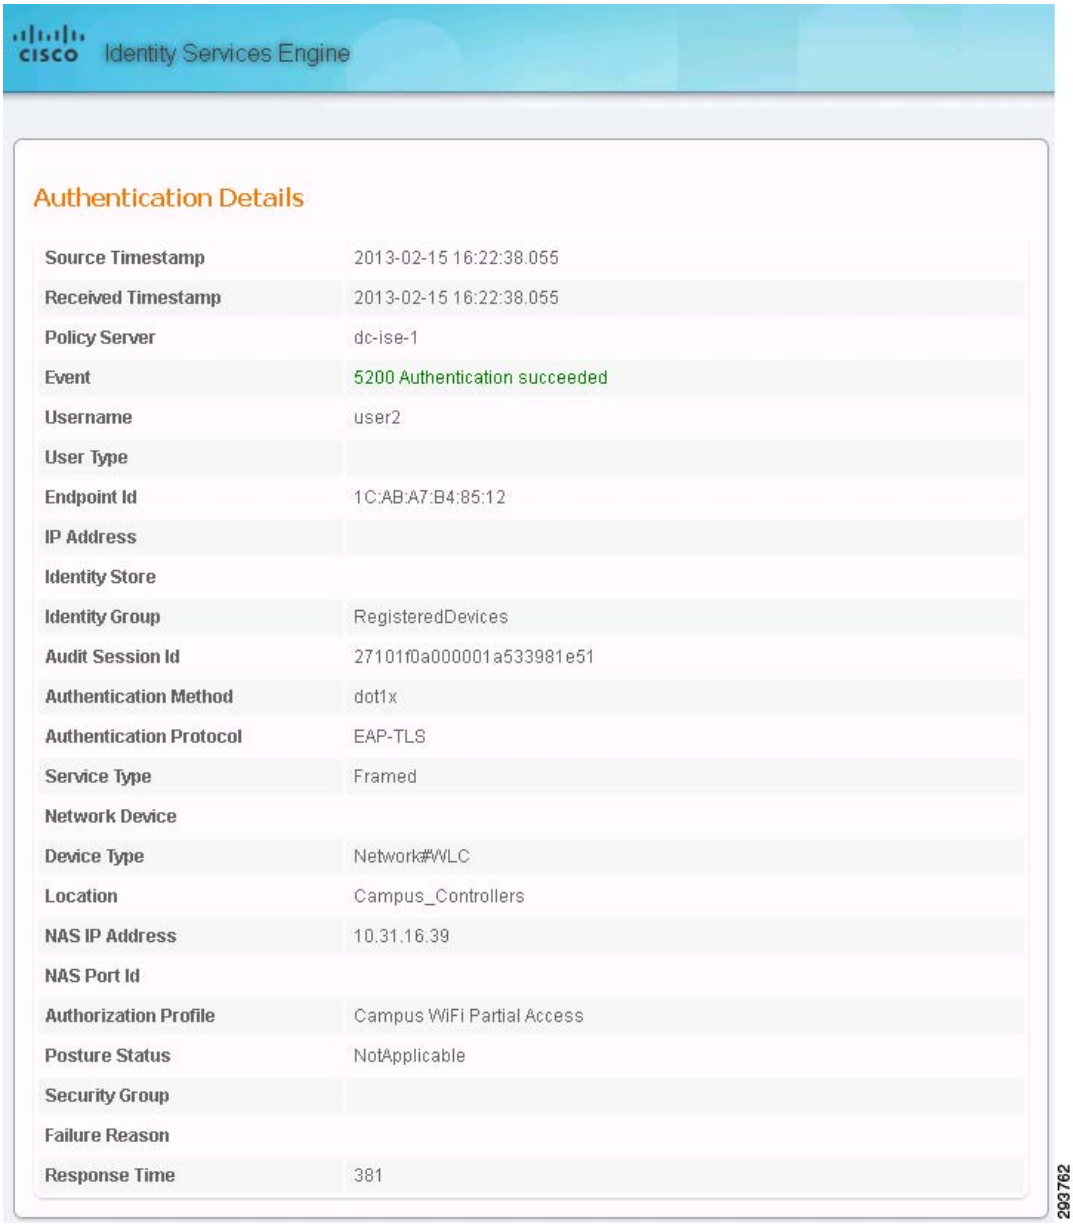


Figure 19-6 shows in more detail the steps that took place and how the rule was evaluated to grant Partial Access.

- Authentication is dot1x and EAP-TLS.
- Username is user2. The Active Directory (AD1) identity store was used.

- MAC Address is discovered.
- The Wireless_Dot1X_AuthC authentication rule was used.
- The Campus WiFi Partial Access authorization rule matched. This rule enforces the access list ACL_Partial_Access in the Wireless LAN Controller.

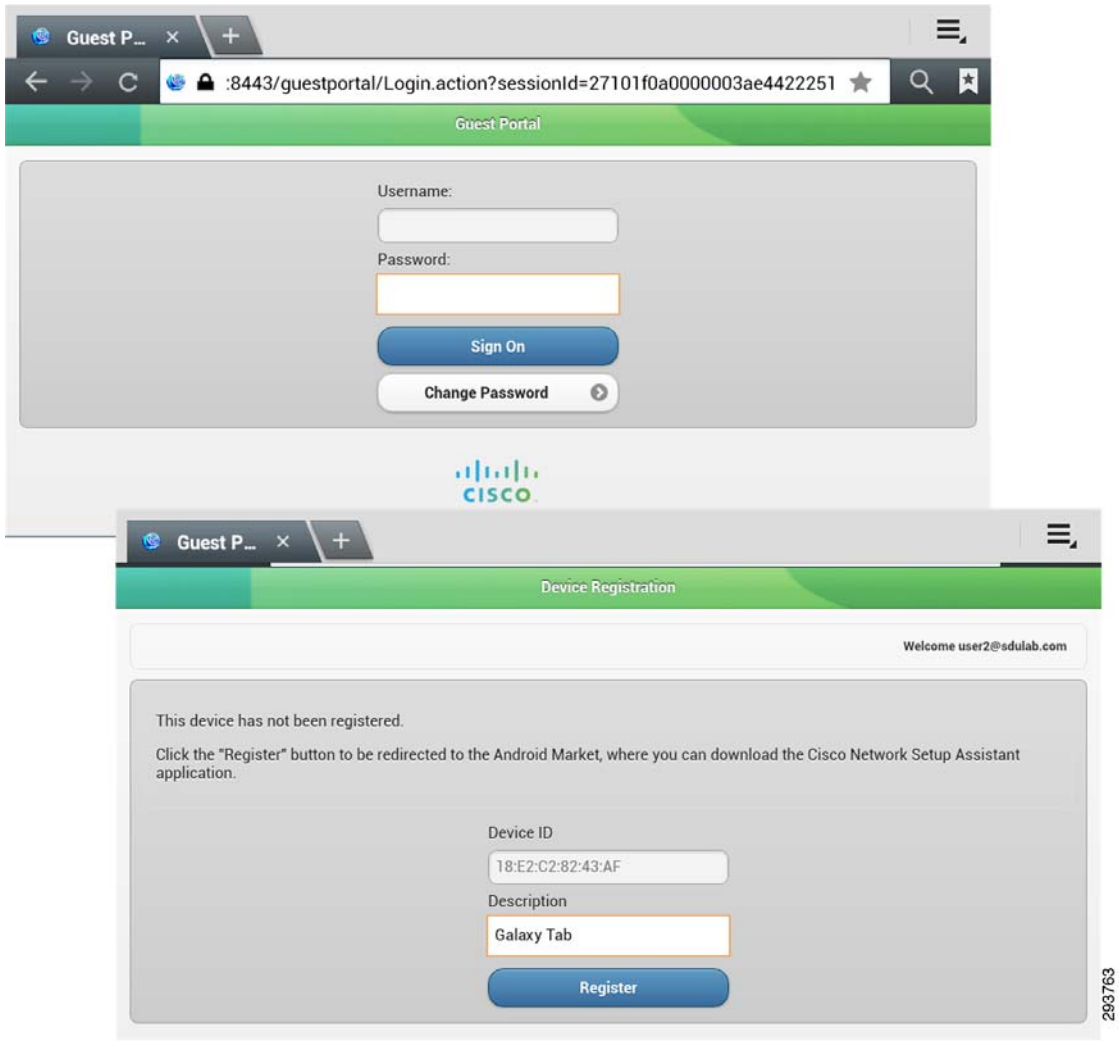
Figure 19-6 ISE Authentication Details



Android Devices

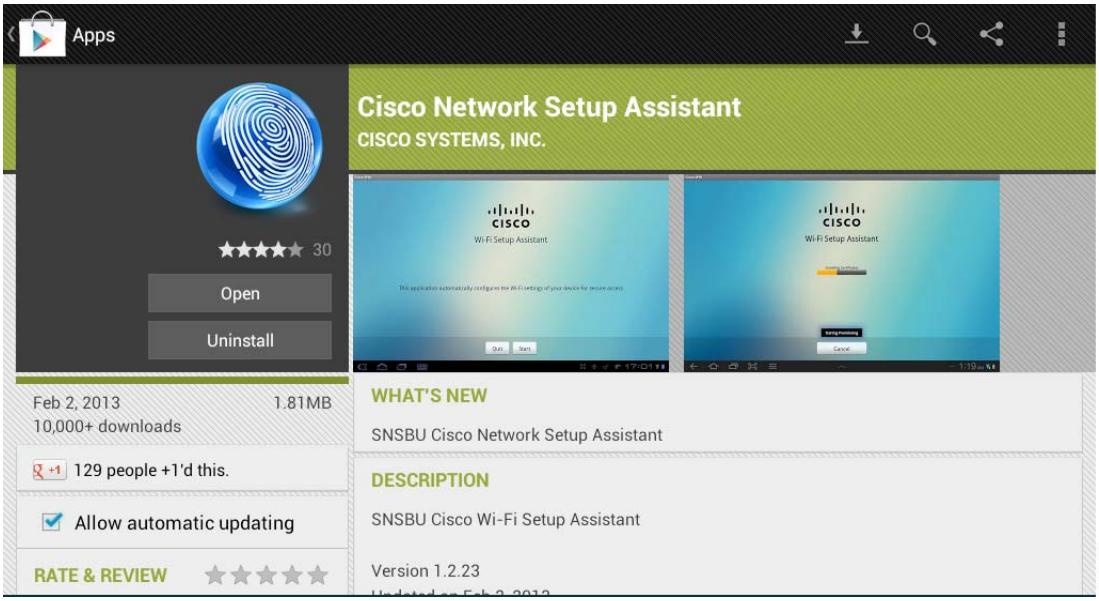
The user experience is very similar when provisioning Android devices. The employee is redirected to the Guest Registration portal and is allowed to enter a description for the new device.

Figure 19-7 Guest Portal and Self-Registration Portal



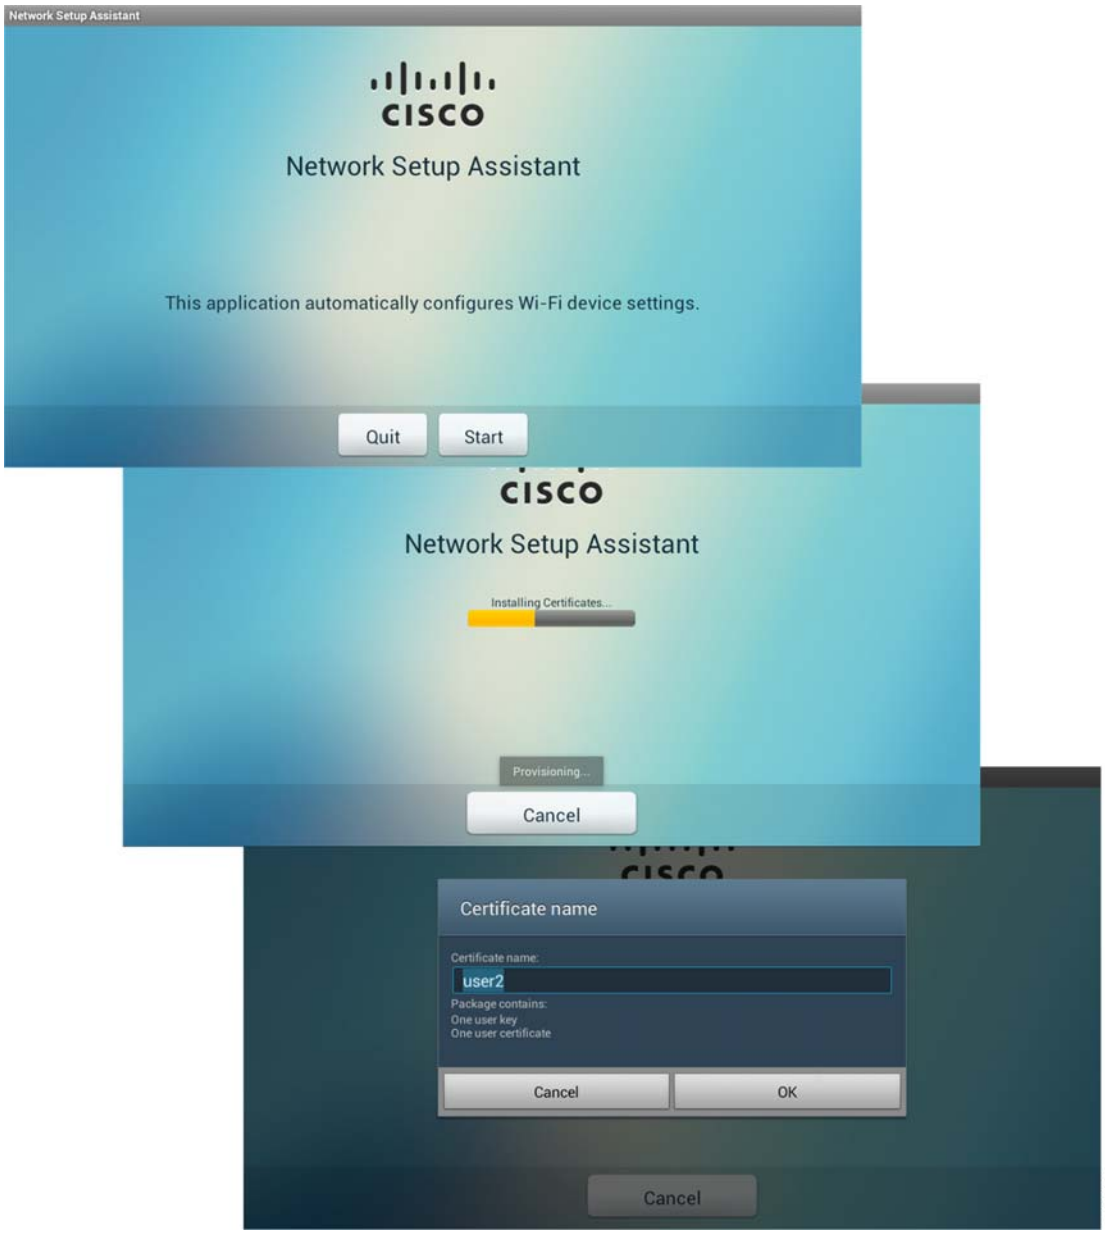
The employee is then redirected to Google Play where the Cisco SPW for Android may be downloaded.

Figure 19-8 Supplicant Provisioning Wizard from Google Play



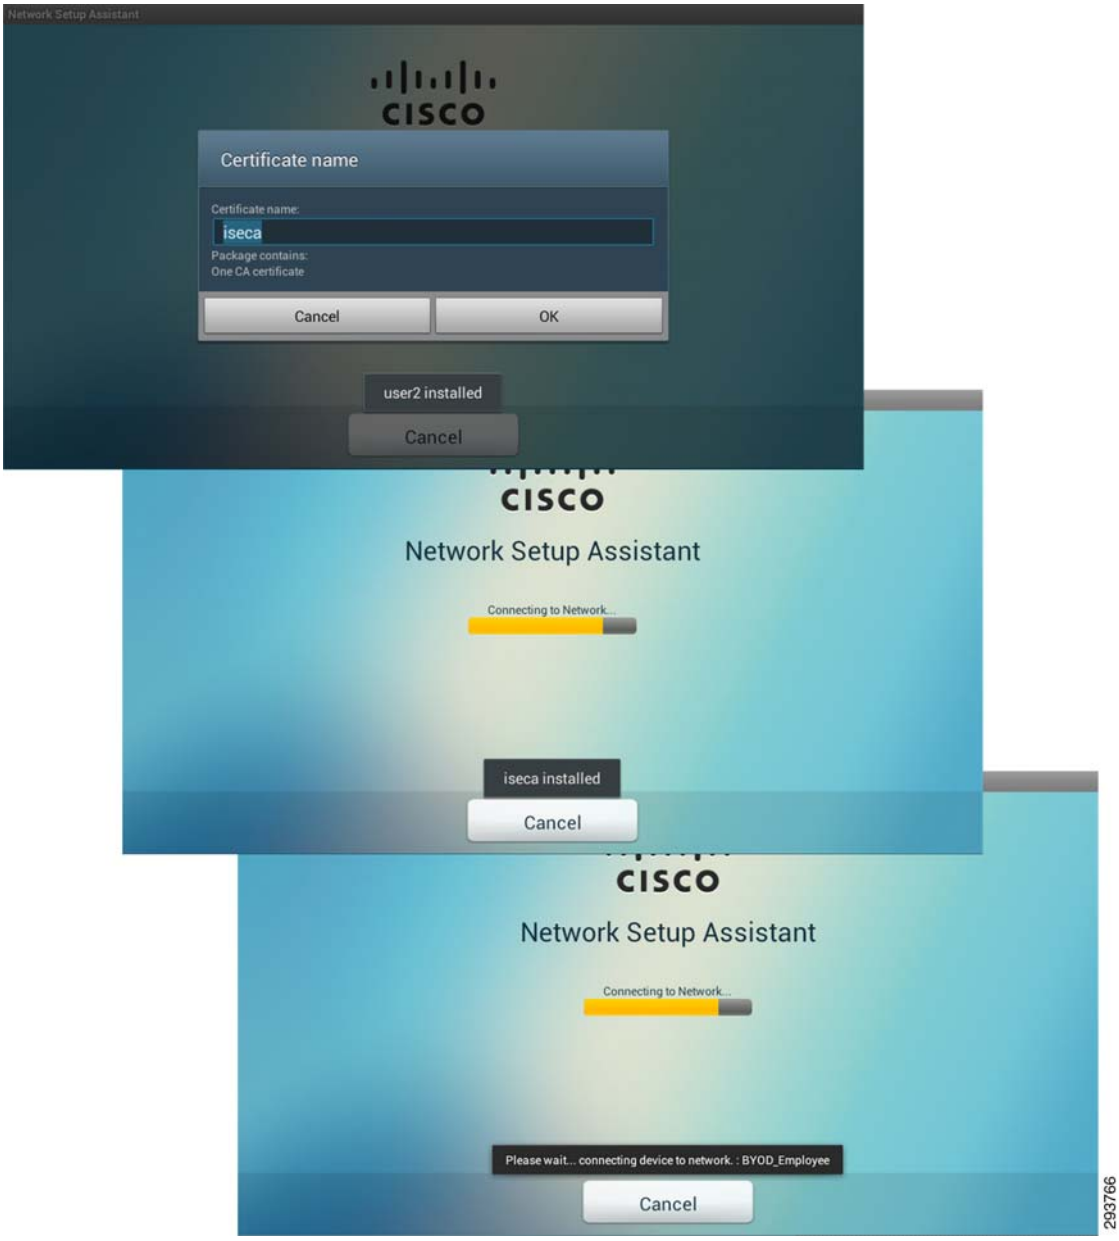
The SPW is launched and the provisioning process begins. The SPW discovers the ISE and begins downloading the profile and installing the certificates.

Figure 19-9 Provisioning Process



The employee is allowed to name the certificate and provides a password for the certificate storage for their device. The Wi-Fi profile to connect to BYOD_Employee is applied.

Figure 19-10 Certificate and Profile



Without employee intervention, the provisioning process automatically connects the Android device to the BYOD_Employee SSID, as shown in Figure 19-11.

Figure 19-11 Automatic Connection to BYOD_Employee

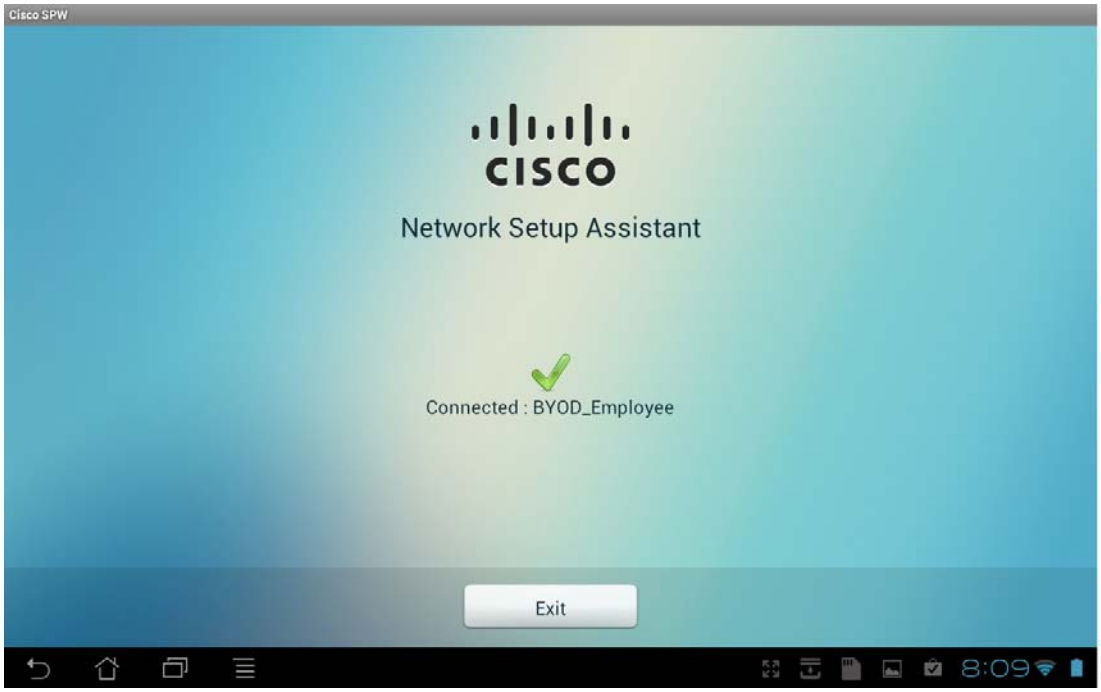
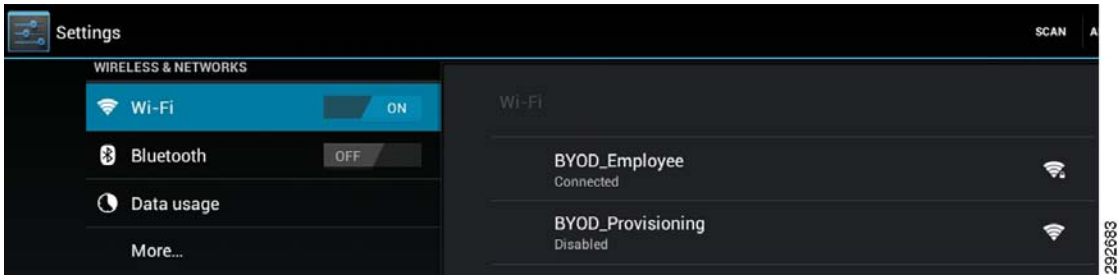


Figure 19-12 shows how the device is automatically connected to the secure BYOD_Employee SSID.

Figure 19-12 BYOD_Employee Secure SSID



Windows Devices

The user experience while provisioning a Windows device is very similar, redirecting the session to the Guest Registration portal and asking the employee for authentication.

Some Windows devices have multiple network adapters, for example, a laptop with both wired and wireless adapter. The network security policy checks that the device mac-address (sent using calling-station-id attribute) matches the SAN field of the device digital certificate before allowing access. This is done to prevent spoofing. Since each adapter has a unique mac-address, the anti-spoofing policy check can cause difficulties for devices with multiple adapters. If a device registers with a wired adapter, it will obtain a digital certificate with the mac-address of the wired adapter. If the same device attempts to later authenticate to the secure wireless network, most operating systems will attempt to the

use the wired adapter certificate for authentication and will fail because the mac-address of the wireless adapter will not match the SAN field of the digital certificate. To avoid this problem, a device with multiple adapters must register both the wired and wireless adapter.

There are two methods for provisioning wireless devices:

- Dual SSID model, which supports MAB on the provisioning SSID and dot1X on the employee SSID.
- Single SSID model, which supports only the dot1X protocol.

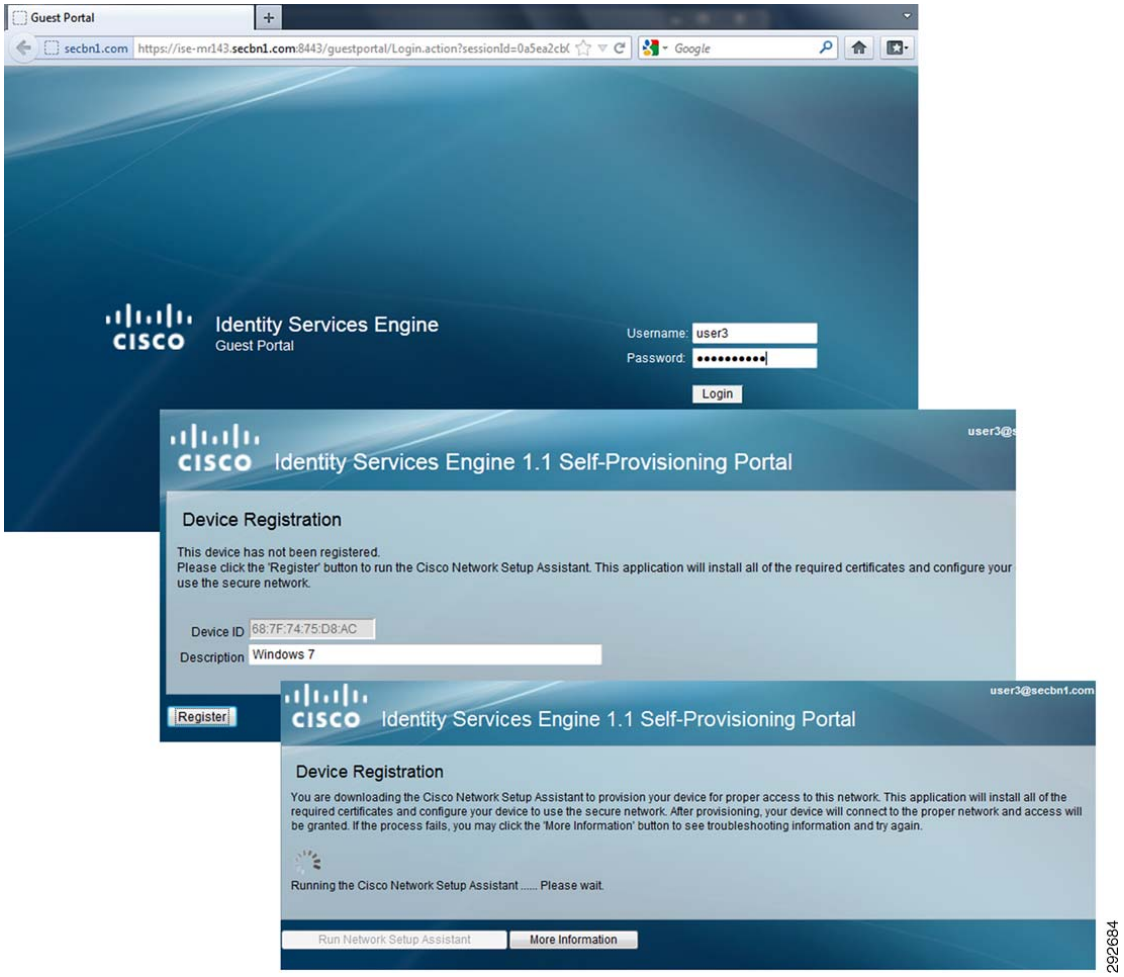
For the Dual SSID method, the wired and wireless adapters may be provisioned in any order. For example, if the device associates with the provisioning SSID (which supports MAB) and is successfully provisioned, then subsequently connects with the wired adapter, 802.1X will fail because of the anti-spoofing check in the policy and the user will be re-directed to complete the provisioning process. Thereafter, the device can access the network with either adapter.

For the Single SSID method, the order in which the wired and wireless adapters are provisioned matters. For example, if the device connects using the wired adapter first and is successfully provisioned, then subsequently connects using the wireless adapter, some operating systems attempt to establish a EAP-TLS connection using the wired digital certificate instead of undergoing the provisioning process. This connection attempt will fail because of the anti-spoofing check in the policy. To prevent this from happening, the user must provision the wireless adapter before connecting with the wired adapter for the single SSID method.

Windows Wireless Devices

After authenticating at the portal and entering a description for the new Windows device, the SPW is downloaded.

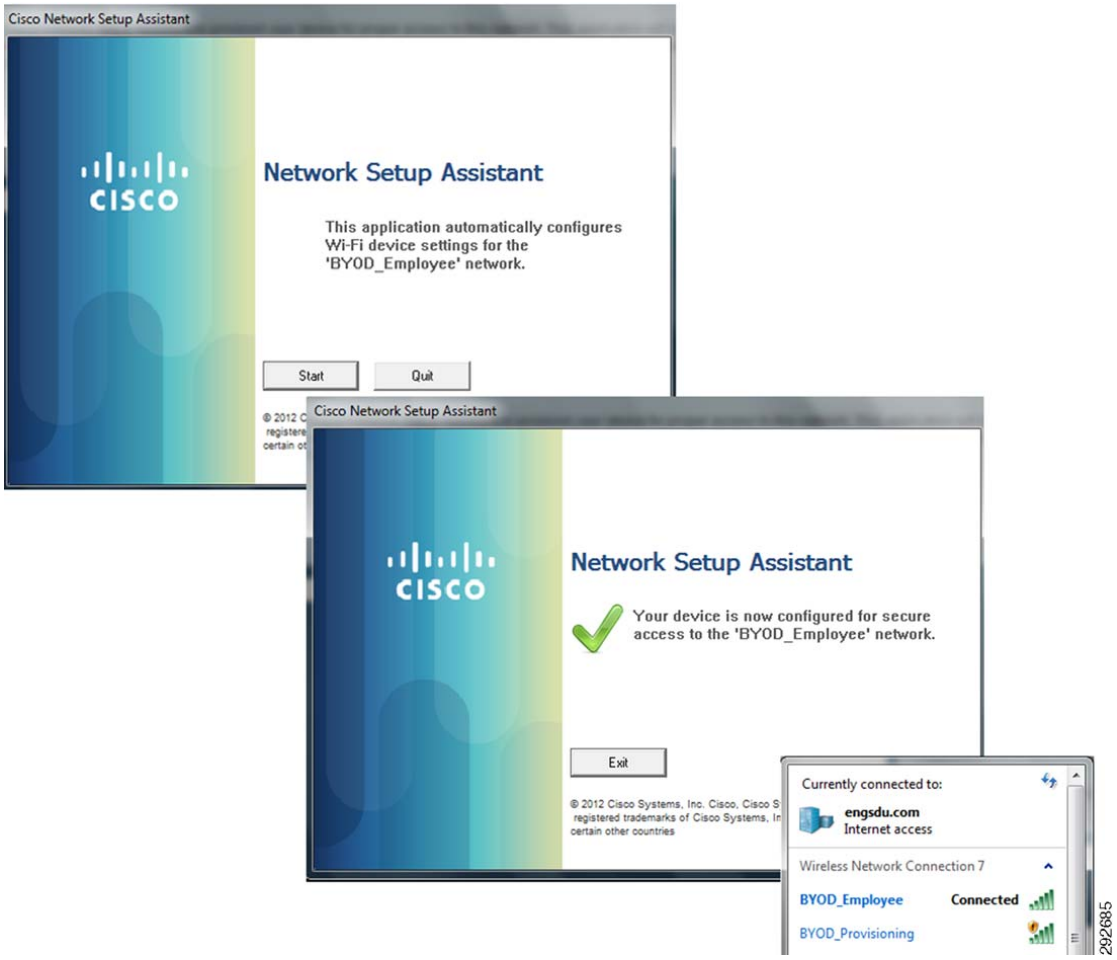
Figure 19-13 Guest Registration Portal



The SPW is launched to install the profile, the keys are generated, and the certificate enrollment takes place.

The SPW installs the BYOD_Employee configuration to connect to the secure SSID. The connection is switched automatically to the BYOD_Employee SSID.

Figure 19-14 SPW and Connection to Secure SSID



Windows Wired Devices

The user experience is very similar, but instead of configuring access to a secure SSID, the SPW configures the devices to connect via a wired connection.

Figure 19-15 Guest Registration Portal



The SPW is downloaded and the proper configurations are applied to the device.

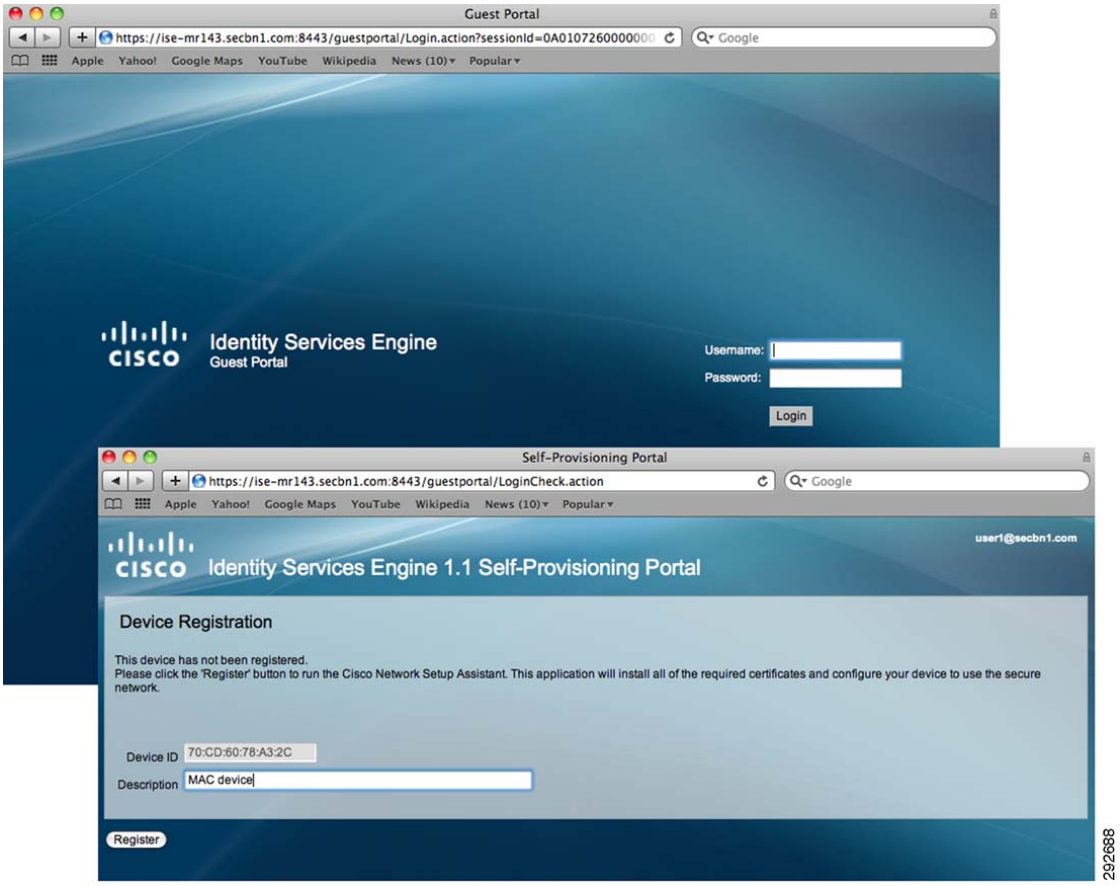
Figure 19-16 SPW and Secure Access



Mac OS/X Devices

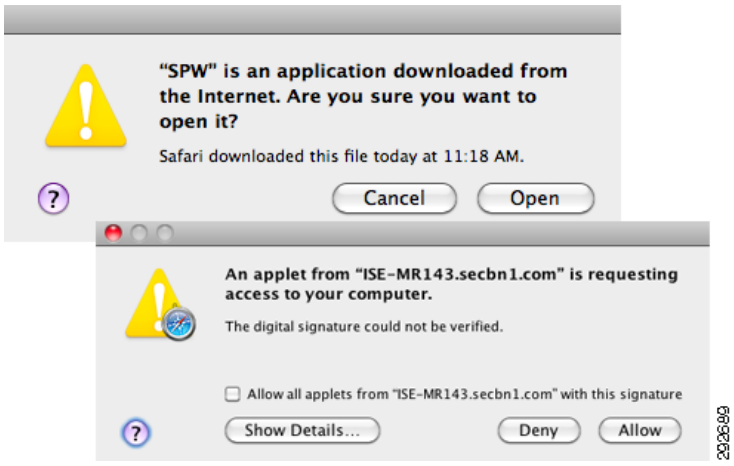
The user experience while provisioning a Mac OS X wired device is also very similar, redirecting the session to the Guest Registration portal and asking the employee for authentication.

Figure 19-17 Guest Registration and Self-Registration Portals



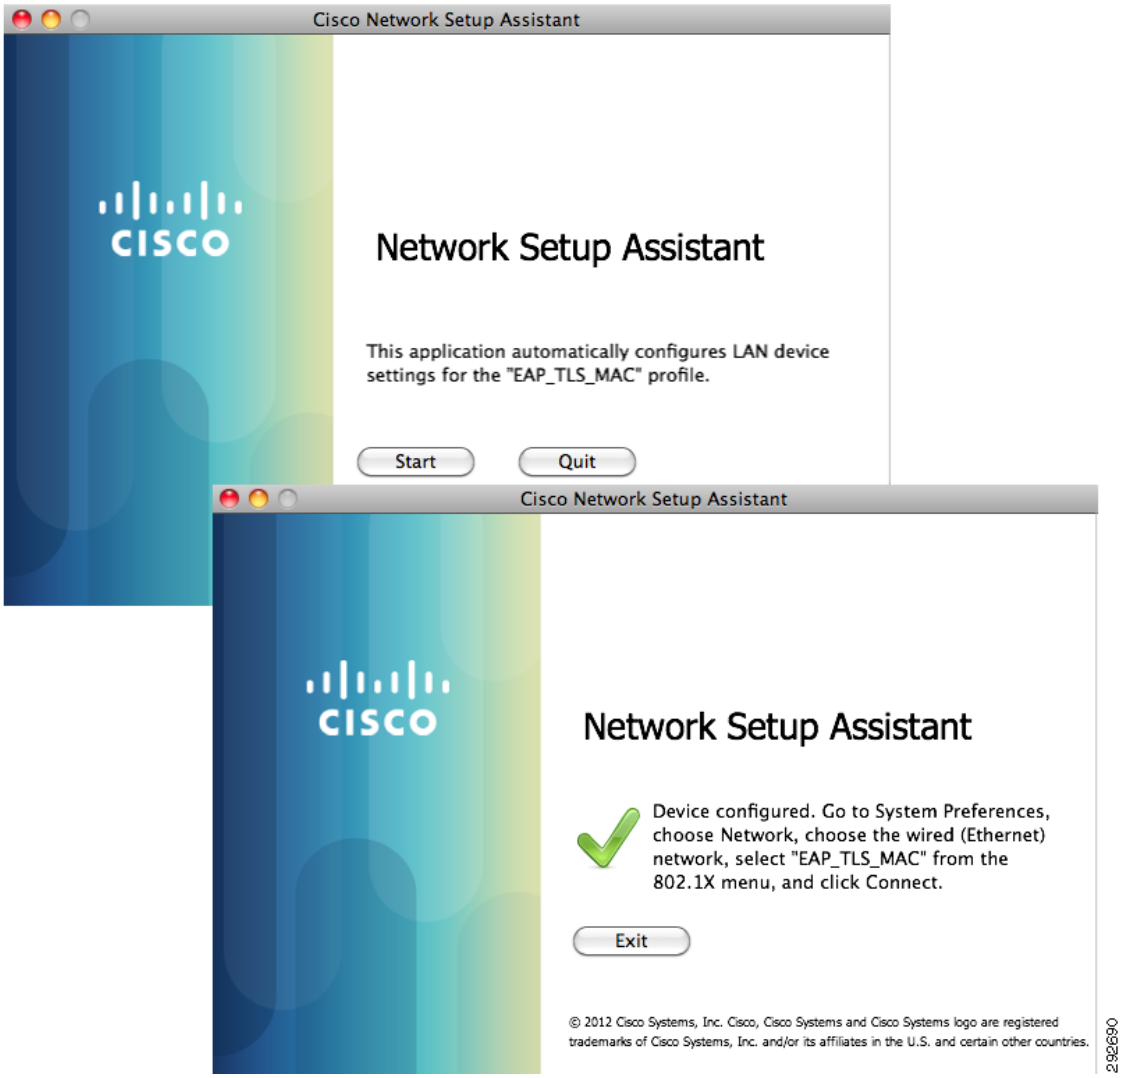
The SPW is downloaded and installed.

Figure 19-18 SPW and Secure Access



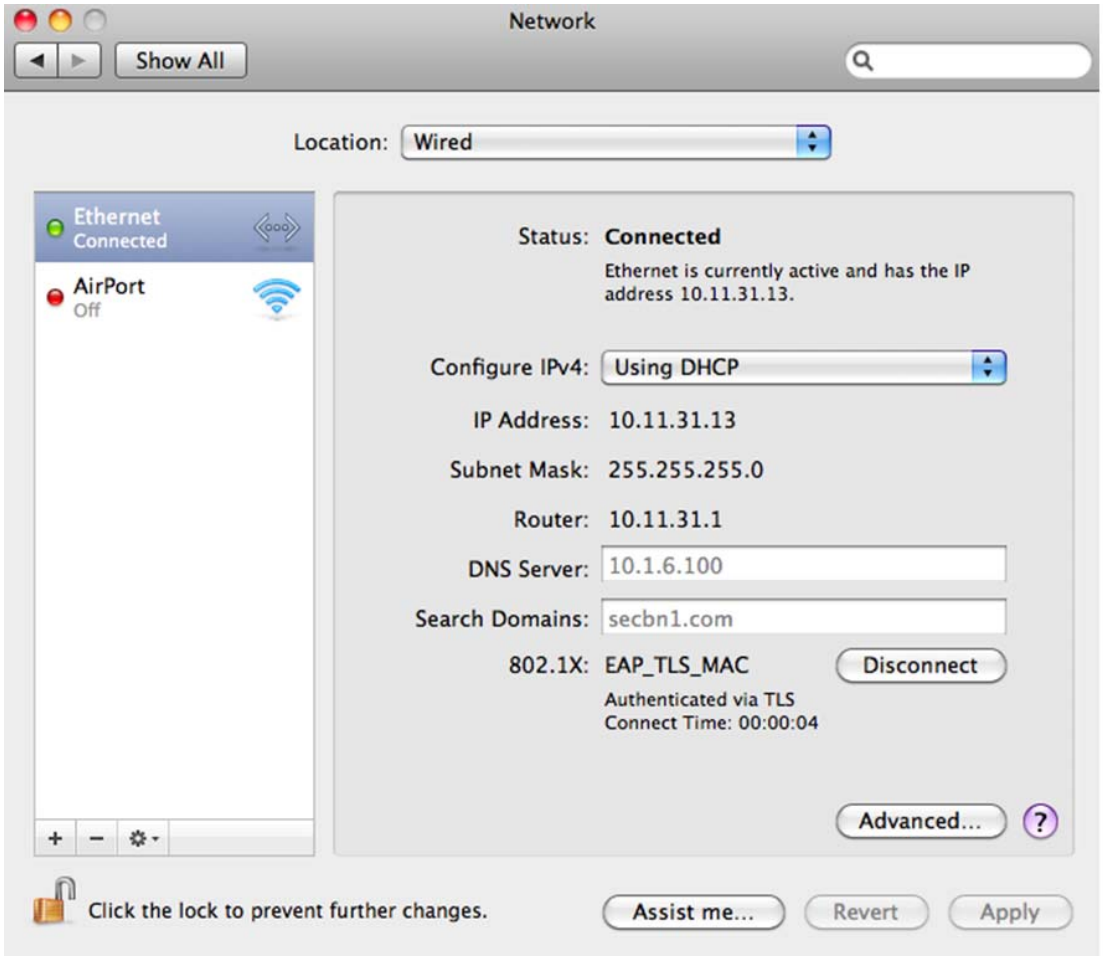
The Network Setup Assistant configures the EAP-TLS_MAC profile for secure access.

Figure 19-19 Network Setup Assistant



The network settings in Figure 19-20 show the new EAP_TLS_MAC configuration.

Figure 19-20 Network Settings





PART 5

BYOD Operations and Services



Summary of Operations and Services

Revised: August 7, 2013

This part of the CVD describes four services in addition to the use cases described in the earlier parts of this CVD. This part highlights how to extend access to guest and remote users and how to manage the BYOD environment and lost or stolen devices.

There are numerous ways to enable a BYOD solution based on the unique business requirements of a specific organization. While some organizations may take a more open approach and rely on basic authentication, other organizations will prefer more secure ways to identify, authenticate, and authorize devices. A robust network infrastructure with the capabilities to manage and enforce these policies is critical to a successful BYOD deployment.

The following components and configuration steps are discussed to support different BYOD use cases:

- Digital Certificates
- Microsoft Active Director authentication
- Wireless Controllers (Unified and Converged Access)
- Identity Services Engine
- Access Layer Switches
- API Integration with Mobile Device Managers

This part of the CVD includes the following chapters:

- BYOD Guest Wireless Access—This chapter describes a traditional wireless guest access solution where users do not have to on-board or register their device with ISE. Internet-only access is granted to guest devices.
- Managing a Lost or Stolen Device—This chapter describes how to deny access to a device that is reported lost or stolen to prevent unauthorized access to the network. By connecting to the My Devices Portal in ISE, users are allowed to manage their devices to prevent unauthorized access or initiate device wipes through the MDM API integration.
- BYOD Policy Enforcement Using Security Group Access—This chapter highlights Security Group Tags as an alternative approach to enforcing policy and traffic restrictions addressing the same use cases addressed in the CVD.
- Mobile Traffic Engineering with Application Visibility and Control (AVC)—This chapter describes different designs that benefit from features such as Quality of Service and the Application Visibility and Control (AVC) on the Cisco WLC. The configuration for different policies is also discussed in detail.

- Managing Bonjour Services for BYOD—This chapter shows how to use the Bonjour Gateway feature of the Cisco WLC to manage Apple's Bonjour protocol in a BYOD enterprise context.
- BYOD Remote Device Access—This chapter describes how to accommodate devices that attempt to connect remotely to access internal resources.
- BYOD Network Management—This chapter describes how to configure and deploy Cisco Prime Infrastructure management suite to manage the BYOD solution.



BYOD Guest Wireless Access

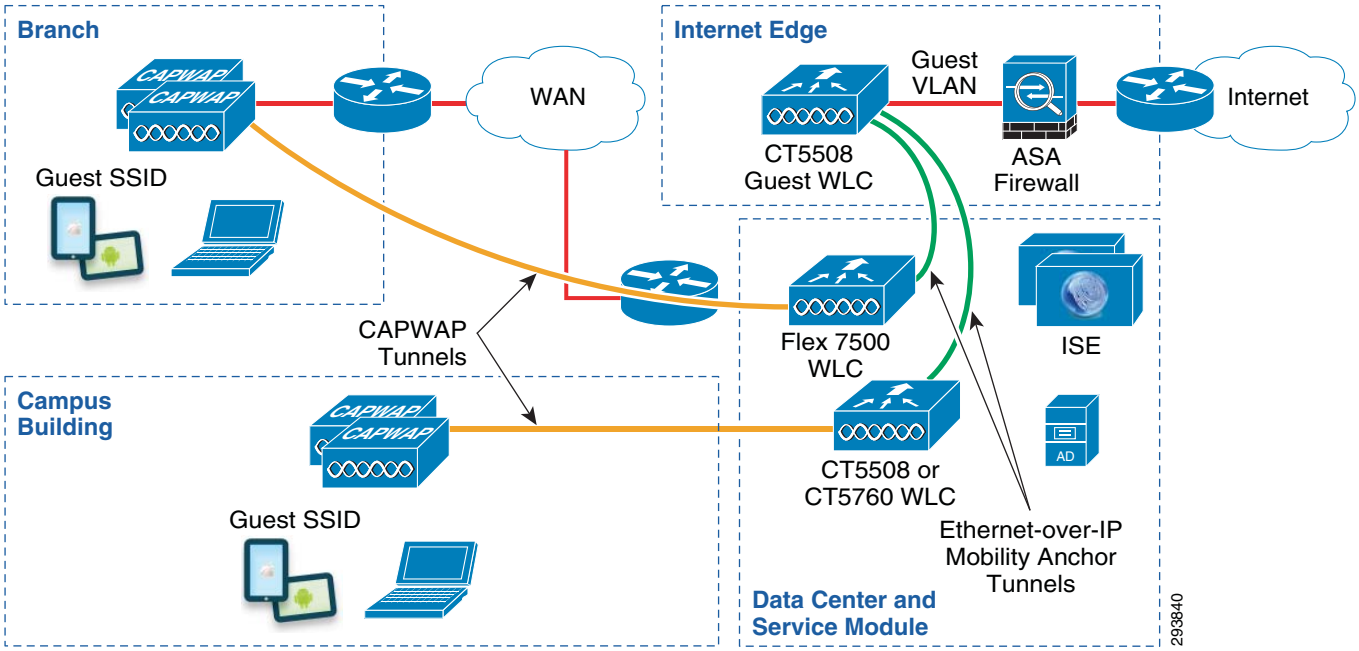
Revised: September 27, 2013

This chapter discusses traditional network access for wireless guest devices and presents various ways to accommodate guest wireless devices within a BYOD implementation. It also provides background information for Chapter 18, "BYOD Basic Access Use Case," which discusses how guest wireless access can be extended to support wireless employee personal devices. Note that within this design guide, guest access refers to temporary Internet access provided for visitors who are sponsored by a representative of the organization being visited.

Overview

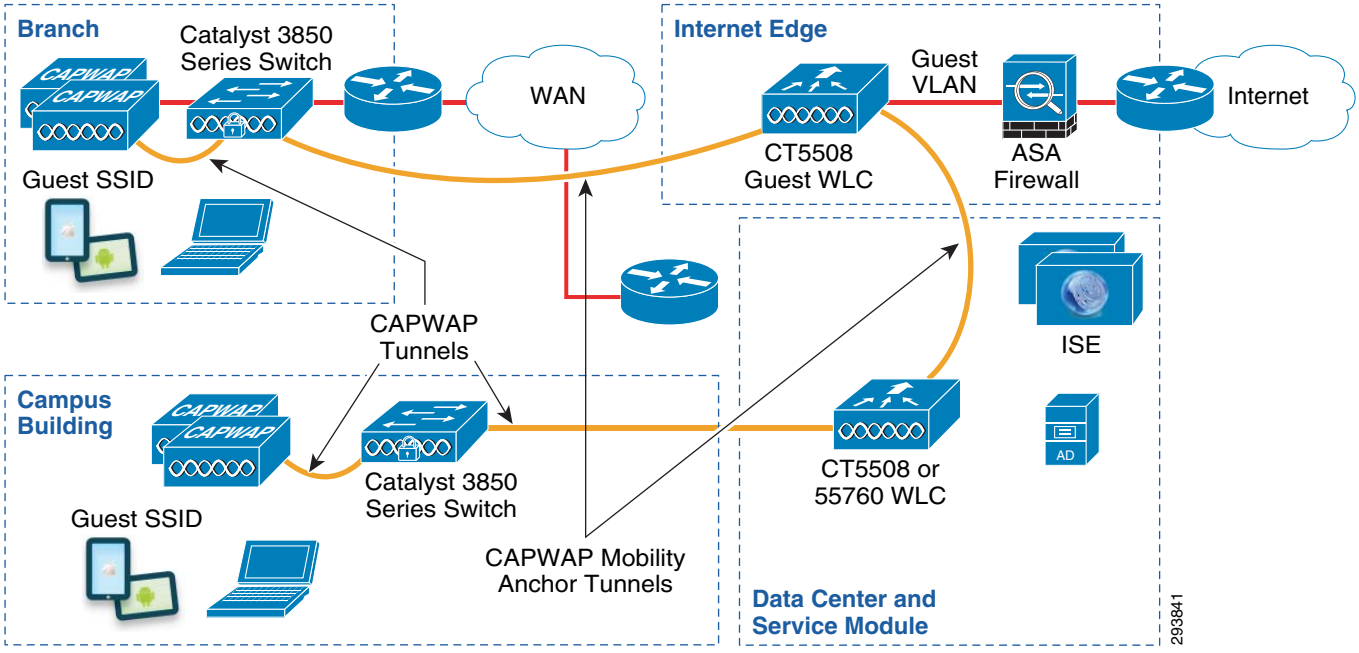
For guest wireless access, a Cisco recommendation has been to deploy a separate, dedicated wireless controller off of a DMZ segment of a Cisco ASA firewall located within the Internet edge module. An example of this design using Cisco Unified Wireless Networking (CUWN) infrastructure is shown in Figure 21-1.

Figure 21-1 Typical Enterprise Guest Wireless Deployment Using CUWN Infrastructure



A similar example of this design using a converged access infrastructure is shown in Figure 21-2.

Figure 21-2 Typical Enterprise Guest Wireless Deployment Using Converged Access Infrastructure



Multiple alternatives for deploying guest access may be deployed. However, this design guide discusses only guest wireless designs based around a dedicated guest SSID configured for open access with no encryption. This is often done because the organization's IT department usually has no knowledge of, or control over, the hardware or software capabilities of the guest wireless device. Hence, open access is the least common denominator applicable to all wireless devices.

Guest wireless traffic from the campus or a branch location is configured to be auto-anchored (tunneled via Ethernet-over-IP or CAPWAP) from the internal wireless controllers to the guest wireless controller. This may provide a somewhat higher level of security, in that guest wireless devices are not terminated on the "inside" of the corporate network. This is often desirable from a customer perspective because the security posture of guest devices cannot be determined.



Note Cisco wireless controllers currently support two different mobility architectures. The old mobility architecture relies on Ethernet-over-IP tunnels between wireless controllers. The new mobility architecture, also called the hierarchical mobility architecture, relies on CAPWAP tunnels between wireless controllers. The two mobility architectures are not compatible with each other. If mobility (including the auto-anchoring function) is required between wireless controllers, all wireless controllers must be running either the new mobility or the old mobility architecture. The new mobility architecture is supported on Cisco 5508 and WiSM2 wireless controllers with software release 7.3.112 and on the Cisco 5508, WiSM2, and 2504 wireless controllers with software release 7.5. The new mobility architecture is supported on the Cisco 5760 wireless controller and the Catalyst 3850 Series switch with IOS XE software releases 3.2.0SE and 3.2.2SE. CUWN wireless controller release 7.4 and releases below 7.3.112 support only the old mobility architecture. The Cisco Flex 7500, 8500, and vWLC do not support the new mobility architecture. IOS XE based wireless controllers do not support the old mobility architecture. Hence if a network contains both Flex 7500 wireless controllers and Converged Access controllers, then separate sets of guest wireless controllers must be deployed with the DMZ to support both mobility architectures with the guest wireless design discussed in this design guide.

There are two distinct sets of terminologies used in this chapter. The first pair of terminologies is guest controller and campus controller. The guest controller is a dedicated controller that is mainly used for dealing with guest wireless traffic and the campus controller is dedicated for handling internal traffic. Note that the term campus controller is used somewhat generically here. The campus controller discussed within this chapter may refer to one or more standalone wireless controller platforms deployed within a campus location or to wireless controller functionality integrated within one or more Catalyst 3850 Series switches deployed within branch locations.

The second set of terminologies is foreign controller and anchor controller. These terminologies are used when a user roams from one controller to another controller. The new controller to which the user associates is the foreign controller and this controller anchors all the traffic to the old controller which becomes the anchor controller.

This design guide chapter discusses wireless guest access primarily from the perspective of how it integrates with the network infrastructure and with the Cisco ISE server for AAA services within an overall BYOD deployment. For details regarding the configuration of wireless controllers for supporting guest access, see the Cisco Unified Wireless Guest Access Services chapter of the Cisco Enterprise Mobility 4.1 Design Guide at: http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html.

IP Addressing and DNS

As with other devices, guest wireless devices require IP addresses and name resolution (DNS) services. A local DHCP server can be deployed on the subnet which supports guest wireless access. This option works well if the ASA firewall performs NAT between the inside and guest wireless DMZ interfaces. Although this may be the most secure option in terms of isolating guest IP addressing from the rest of the corporate network, it is also somewhat cost prohibitive and more difficult to administratively maintain. This cost can be offset by implementing an ASA firewall configured with a DHCP pool to hand back IP addresses directly to wireless clients. The advantage of this option is again the isolation of guest

IP addressing from the rest of the corporate network and the fact that DHCP from guest devices do not have to be allowed through the ASA firewall. The downside is the management of a separate IP addressing pool for guest wireless devices within the ASA firewall.

IP addressing for guest wireless devices can also be provided through a DHCP server on the inside of the corporate network. This option works well if NAT is not implemented between the inside and the guest wireless DMZ interfaces. The remainder of this chapter assumes no NAT functionality for the guest wireless DMZ interface. The advantage of implementing a centralized DHCP server is the centralized control of IP addressing for guest devices. The downside is that DHCP has to be allowed through the ASA firewall to the internal DHCP server.

Cisco wireless controllers can be configured to proxy for wireless clients to an internal DHCP server. This is a common deployment model for wireless controllers. With this configuration the DMZ interface of the ASA firewall needs to allow inbound DHCP packets from the IP address of the wireless controller associated with the guest WLAN interface through the ASA firewall. Alternatively, instead of the guest wireless controller acting as a proxy for wireless devices, the ASA firewall can be configured to relay DHCP to an internal DHCP server. With this configuration, guest wireless clients directly send DHCP through the wireless controller, which are then relayed to an internal DHCP server by the DMZ interface of the ASA firewall. Note that DHCP profiling of end devices via a Cisco ISE server can be accomplished by relaying the DHCP discover to both the internal DHCP server as well as the ISE profiling server. However, there may be no need or desire to profile guest devices, since they require only temporary access.



Note The network administrator should always weigh the benefits achieved from enabling DHCP server or DHCP relay functionality against the incremental risks of enabling such additional features on the ASA firewall to determine the appropriate security policy for the organization.

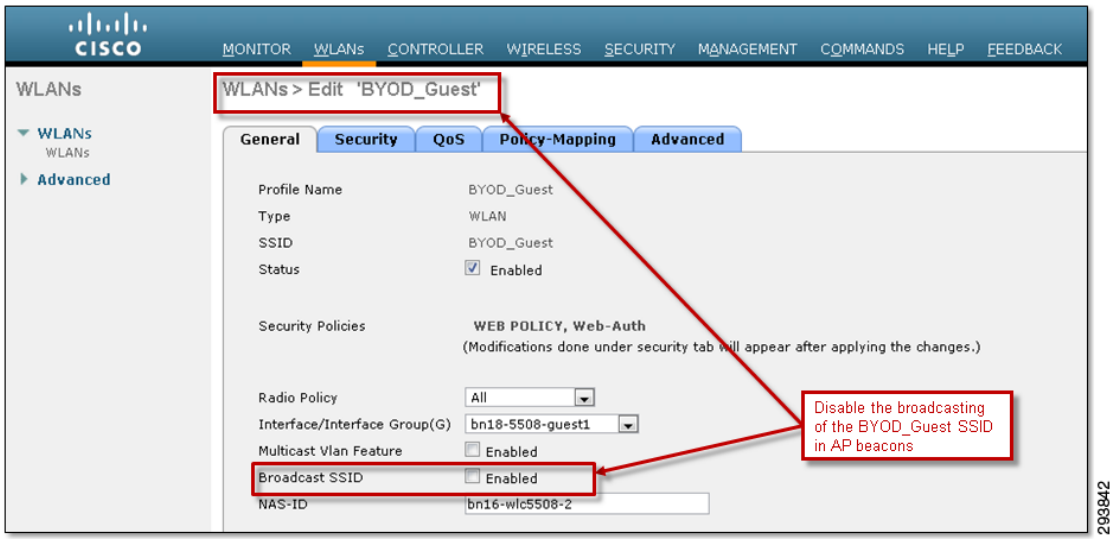
An increasing issue with guest wireless networks is IP address depletion. This may be the result of opening up the traditional guest network to employee personal devices. It may also be the unintentional result of having an office in a densely populated area where the general public is connecting to the open SSID corresponding to the organization's guest WLAN, thinking that it provides “hot spot” wireless services. As the proliferation of consumer wireless devices continues and as organizations continue to adopt BYOD strategies, this problem may become more widespread. If branch locations are offering guest services, the required address pool can become quite large.

There are a number of methods which can be implemented to help alleviate the issue of IP address depletion. From a security perspective, the optimal solution is to try to tune the Access Point (AP) radios such that the SSID corresponding to the guest WLAN-along any of the organization’s other wireless SSIDs for that matter-are not visible outside the physical boundaries of the organization. However, this is not always possible while still maintaining adequate wireless coverage across the entire floor space.

A second method is to decrease the lease time on the DHCP server for the IP subnet corresponding to the guest WLAN. This does not prevent the general public from connecting to the open SSID corresponding to the organization’s guest WLAN. However, when end users realize they do not have the web authentication (Web Auth) credentials needed to access anything, they may reconnect to another SSID. The IP addresses handed-out to these devices are made available to hand-out again more quickly if the DHCP lease time is decreased. The downside is the additional overhead on the DHCP server and slightly additional overhead on the wireless device itself from having to renew leases faster.

A third method is to hide the SSID corresponding to the guest WLAN by not broadcasting it in AP beacons. Cisco Unified Wireless Network (CUWN) controllers provide an easy means of achieving this by simply un-checking the Broadcast SSID checkbox for the WLAN corresponding to the guest SSID, as shown in Figure 21-3.

Figure 21-3 Disabling the Broadcast of the SSID Corresponding to the Guest WLAN in AP Beacons



Similarly, the following configuration example shows only the part of the configuration of the guest SSID on a Converged Access (IOS XE based) wireless controller in which broadcasting of the SSID has been disabled.

```
!
wlan BYOD_Guest 2 BYOD_Guest/Configuration for the guest SSID
 no broadcast-ssid/Disables broadcasting the SSID in AP beacons
!
```

This is by no means a foolproof method of keeping unwanted devices from connecting to the open SSID corresponding to the guest WLAN, since it can still be discovered by other means. However, it does make it harder to find and connect to it, potentially reducing the number of unwanted devices and therefore the number of IP addresses being issued by the DHCP server. The downside is that guests have to manually type in the name of the SSID when trying to connect to the organization's guest wireless network. However, the name of the SSID can also be included with the credentials provided to the guest either prior to or at the time of arrival to organization's site.

Another option is to provision a larger contiguous IP subnet address space for the guest wireless network, simply by changing the IP subnet mask of the existing guest IP address space. This works well if the adjacent IP address space is available and unused. If this cannot be done, a second guest DMZ interface can be provisioned on the wireless controller to increase the IP address space available to hand out to devices on the guest WLAN.

Increasing the pool of available IP addresses is the most direct method to ensure guests are not prevented access due to address depletion. It is worth noting that this approach does not discourage adjacent wireless clients from associating to the wrong network. Web Auth or some other method is needed to control access to guest resources. It is also considered good practice to audit the actual number of guest users with the anticipated number. Comparing the number of guest that have passed through the guest portal with the number of addresses that are leased out from the DHCP server is a good means to determine how many unintentional wireless clients are associating with the network. The lease time can be adjusted down if the number of leased addresses far exceeds the anticipated number of guests.

Wireless guest devices also need name translation services (DNS) to reach locations on the Internet. Also, when Web Auth is implemented, the URL within the guest's web browser must resolve to an IP address. This is necessary for Web Auth to redirect the session to the guest portal to request guest credentials. Name translation services can be provided by allowing guest devices to reach either an external DNS server deployed on another DMZ segment off the ASA firewall or an internal DNS server

deployed on the inside of the corporate network. Allowing guest devices access to an external DNS server provides the advantage that internal sites and services can be hidden from the guest devices. However, if the wireless guest network is extended to include employee personal devices, as discussed in Chapter 18, "BYOD Basic Access Use Case," the network administrator needs to determine if an external DNS server can still provide the necessary name translation services.

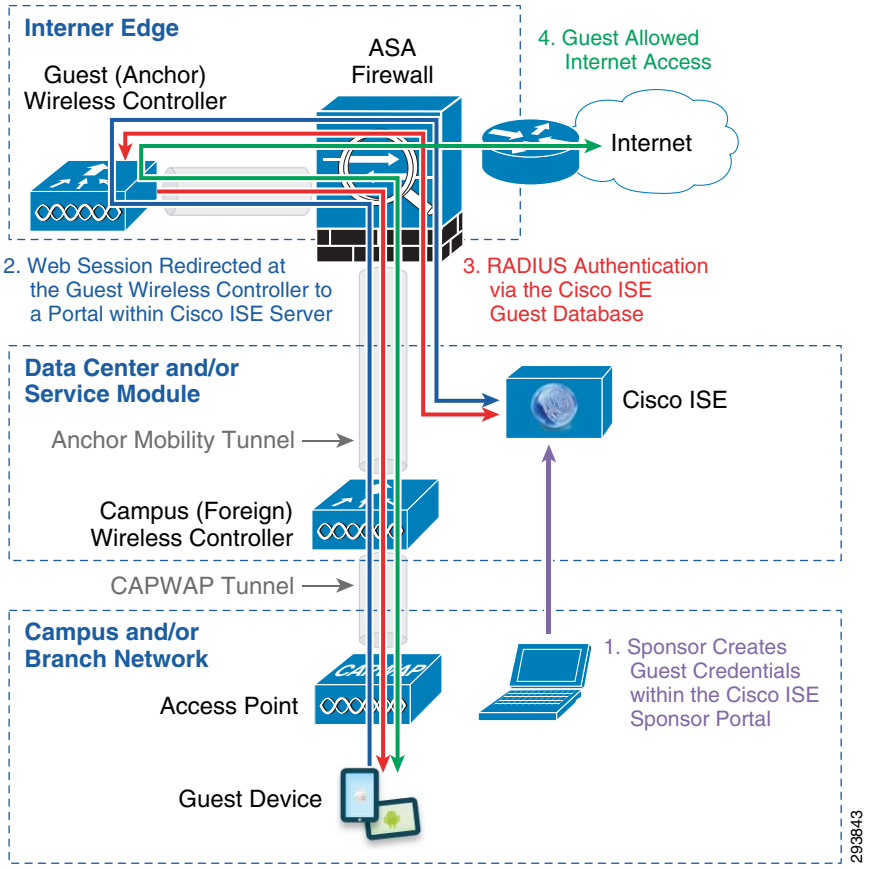
Note DHCP packets from client to server utilize UDP source port 68 and destination port 67. DHCP packets from the server to the client utilize UDP source port 67 and destination port 68. DNS uses UDP port 53. These ports must be allowed through the firewall when internal DNS and DHCP servers are implemented.

Authentication and Authorization

Most organization's IT departments choose to have guest wireless users authenticate first, before allowing access to the Internet. This step is sometimes accompanied with the guest user reading and agreeing to an Acceptable Use Policy (AUP) or End User Agreement (EUA) before accessing the Internet. Since the organization's IT department typically has no control over the hardware or software capabilities of guest wireless devices, the authentication and authorization decision is often based on a guest userid and password only. In other words, from a BYOD perspective, the device with which the guest is accessing the network may not be considered for the policy decision. A typical way of implementing guest user authentication, which is shown in Figure 21-4, is through the guest user's web browser, known as web authentication or Web Auth. With this method of authentication, the wireless guest must first open their web browser to a URL located somewhere within the Internet. The browser session is re-directed to a web portal which contains a login page which requests login credentials. Upon successful authentication, the guest user is either allowed access to the Internet or redirected to another website.

A major requirement of guest access is the ability of a sponsor, such as a lobby administrator, to access a portal to create temporary guest credentials which are valid for a limited time. Hence, this functionality is also included within the discussion below.

Figure 21-4 Guest Wireless Access with Web Authentication



Designing Guest Access for Campus and Branch locations

Implementing the guest access design discussed within this document is very similar for campus and branch networks. Typically the same configuration steps can be used for both. Deploying the guest access solution involves configuring several components such as the wireless controller (WLC), ASA firewall, and Cisco ISE.

WLC Configuration

For the design shown in this document, redirection of the guest web session and the point of authentication from the wireless controller is directed to the Cisco ISE server. Other methods of performing the web authentication are available, but are not discussed in this guide. The guest client's web session is redirected by the guest wireless controller to a portal containing the login screen located within the Cisco ISE server.

Note This option is sometimes referred to as Centralized Web Auth (CWA).

By positioning the Web Auth login page (and optionally the AUP or EUA) in a central location, the network administrator can provide one unified login page for all wireless guest access without having to download the login page to each guest wireless controller.

As discussed in the initial overview, the recommendation from this design is to deploy two different controllers:

- A campus controller which handles all internal wireless traffic.
- A dedicated guest controller that only handles the guest traffic.

These two controllers have a mobility anchor tunnel established between them. This section discusses the configuration details for both of the controllers.

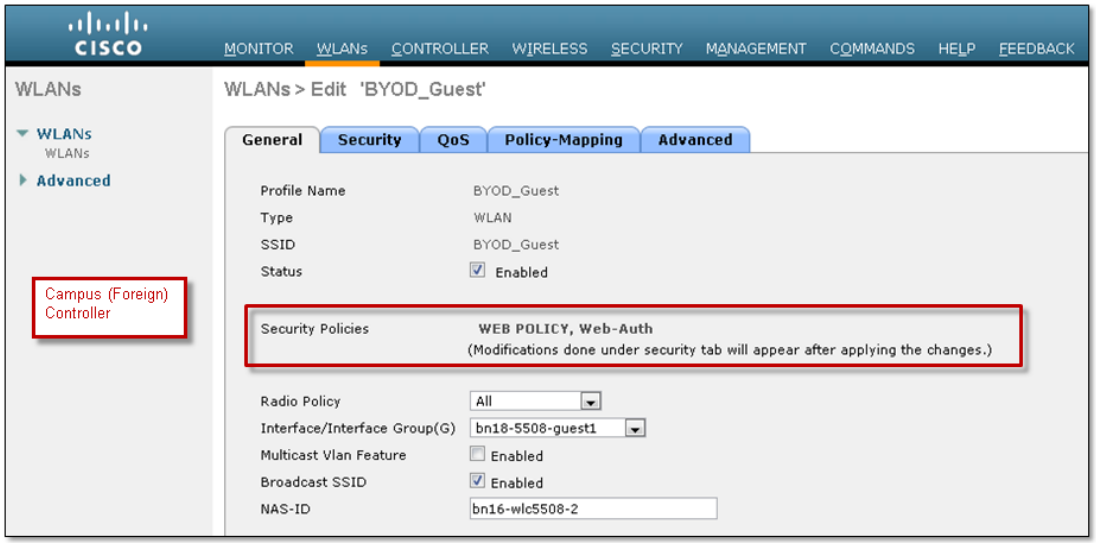
Campus Controller

This section discusses the campus controller configuration when using either CUWN wireless controllers or Converged Access (IOS XE based) wireless controllers.

CUWN Wireless Controllers

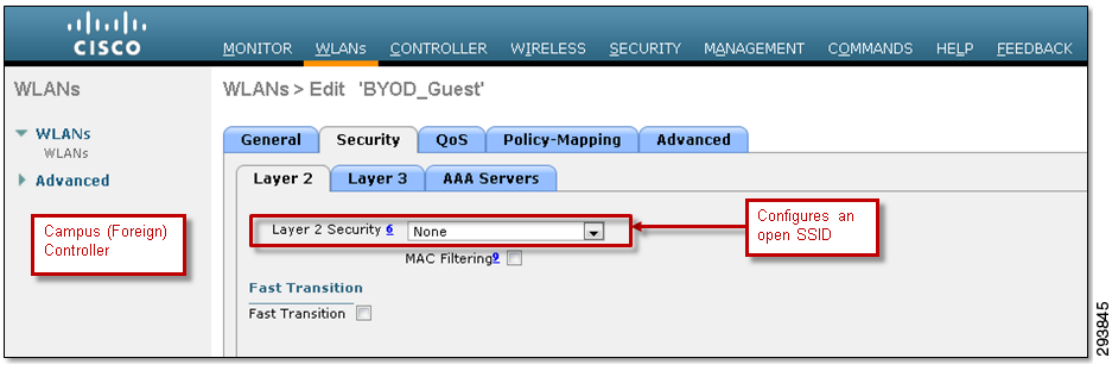
The first step is to configure a guest SSID. Figure 21-5 shows the configuration for the BYOD_Guest SSID. Note that the authentication must be configured for Web Auth.

Figure 21-5 BYOD_Guest SSID Details on the Campus Controller



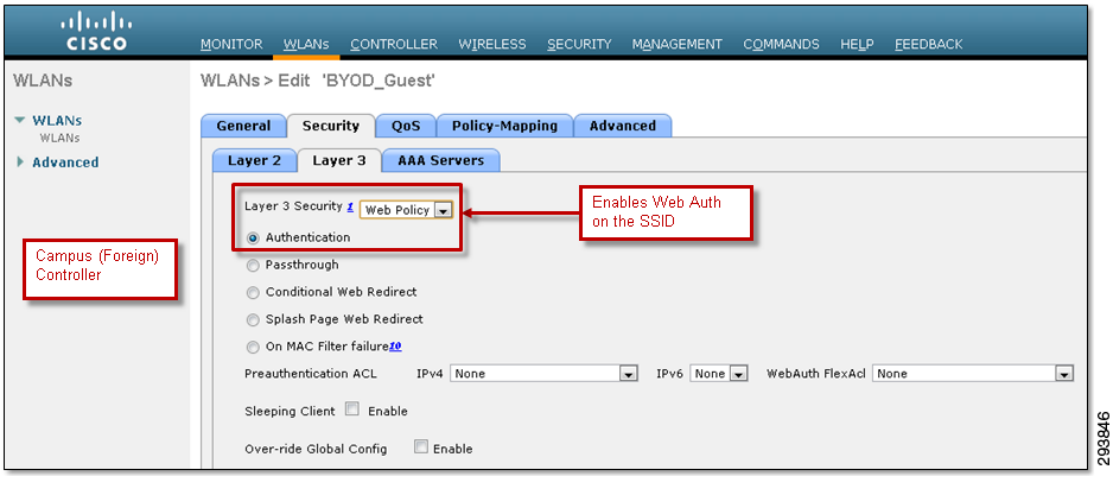
The next step is to configure the Layer 2 and Layer 3 security parameters for this SSID. Figure 21-6 shows the Layer 2 security parameters.

Figure 21-6 Layer 2 Security Details of BYOD_Guest



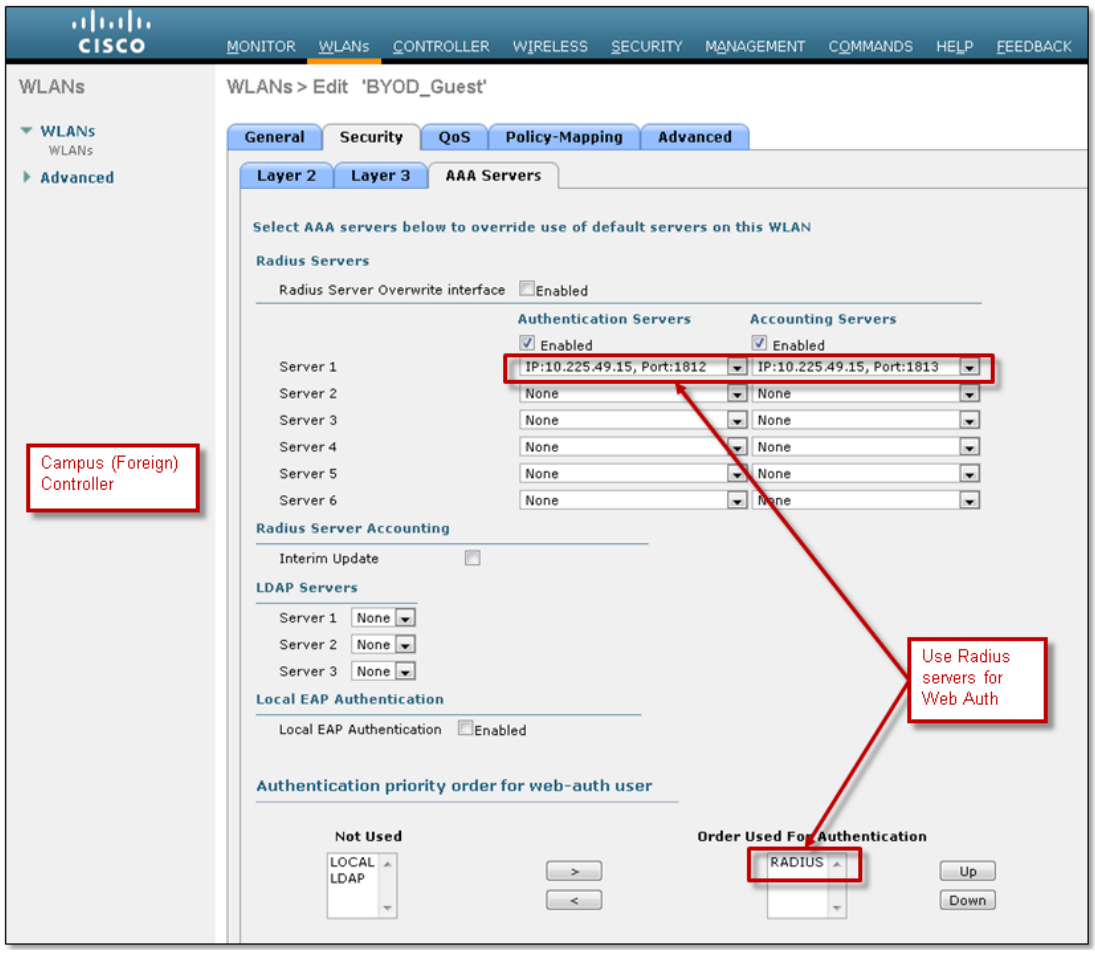
As mentioned before, the Layer 2 security parameters are set for None, indicating an open SSID. The Layer 3 security parameters are shown in Figure 21-7.

Figure 21-7 Layer 3 Security Details of the BYOD_Guest WLAN on the Campus Controller



The Layer 3 security parameters enable web authentication (Web Auth). The next step is to configure the AAA server parameters, which are shown in Figure 21-8.

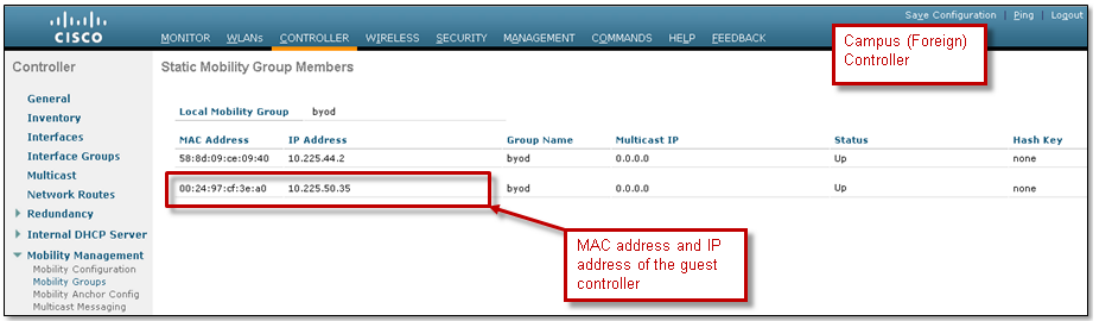
Figure 21-8 AAA Server Configuration for the BYOD_Guest WLAN on the Campus Controller



The AAA server parameters are configured such that Web Auth utilizes the Cisco ISE server for authenticating guests using RADIUS.

The next step is to configure the mobility tunnel between the campus and the guest controller. The guest controller must first be added to the campus controller as a mobility group member. Figure 21-9 shows an example of this.

Figure 21-9 Adding the Guest Controller to the Mobility Group



Note Both the MAC address and the IP address of the management interface of the guest controller are needed in order to add it as a mobility group member.

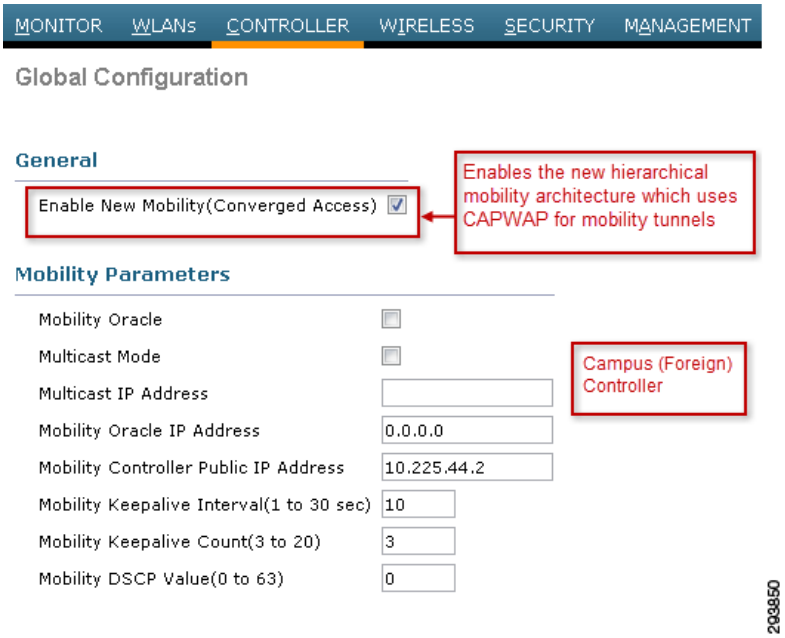
Finally, a mobility anchor is created on the BYOD_Guest SSID which points to the IP address of the management interface of the guest controller. An example is shown in Figure 21-10.

Figure 21-10 Configuring the Mobility Anchor on the Campus Controller



In order to support the new mobility architecture (also referred to as the Hierarchical mobility architecture) the network administrator must check the **Enable Hierarchical Architecture** option within the global mobility configuration of the campus wireless controller. This is shown in Figure 21-11.

Figure 21-11 Enabling the Hierarchical Mobility Architecture in the Campus Controller



Note Since the Flex 7500 wireless controller does not support the new mobility architecture, this step can be skipped when implementing a Flex 7500 as the branch wireless controller.

Converged Access (IOS XE Based) Wireless Controllers

The following partial configuration example shows the configuration of the guest WLAN on a Converged Access (IOS XE based) wireless controller.

```
!
vlan 777        /Isolated VLAN for guest devices if anchor tunnel is down
 name Guest
!
~
!
wlan BYOD_Guest 2 BYOD_Guest/Guest WLAN, WLAN ID, and SSID on the campus controller
 aaa-override
 client vlan Guest/Static assignment to non-routed (isolated) VLAN
 mobility anchor 10.225.50.35/Creates CAPWAP anchor tunnel to guest wireless controller
 no security wpa/Layer 2 security set to none (Open SSID)
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth/Layer 3 security set for web authentication
 session-timeout 1800
 no shutdown    /Enables the Guest WLAN
!
```

Note that the Guest client VLAN in the configuration above is a VLAN which is isolated on the CT5760 wireless controller or Catalyst 3850 Series switch. It is not trunked to the adjacent Layer 3 device. This isolates any guest devices should the CAPWAP tunnel between the foreign and anchor controllers go down.

The guest WLAN must be configured on the device which functions as the Mobility Agent (MA) and on the device which functions as the Mobility Controller (MC). For more information regarding the MA and MC functions, see Chapter 9, "BYOD Wireless Infrastructure Design." Therefore, when the converged access infrastructure consists of a Catalyst 3850 switch configured as the MA with a CT5760 wireless controller configured as the MC within a large campus, the guest WLAN must be configured on both devices. When the converged access infrastructure consists of just a Catalyst 3850 switch configured as both the MA and MC within a branch, the guest WLAN is also configured as shown above.

Note that the wireless mobility configuration will differ, depending upon whether a Catalyst 3850 switch is configured as a MA within a large campus or as both a MA and MC within a branch. This is discussed in Chapter 9, "BYOD Wireless Infrastructure Design."

In order to support guest access, the guest wireless controller must be added as a member of the mobility group within the MC. The following partial configuration shows an example of the configuration of a mobility group and a mobility group member which points to a guest wireless controller.

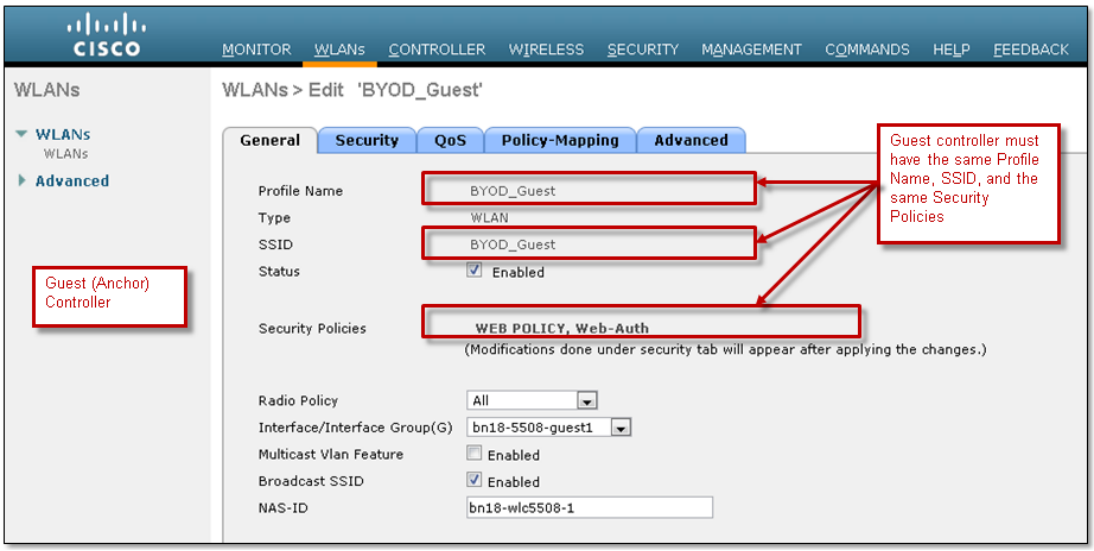
```
!
Wireless mobility controller/Enables the MC function
wireless mobility group member ip 10.225.50.35 public-ip 10.225.50.35/Guest Controller
wireless mobility group name byod/Mobility group name
!
```

The mobility group name and mobility group peer configuration must appear on the device which functions as the Mobility Agent (MA). Therefore, when a Catalyst 3850 Series switch is deployed as both the MA and MC within a branch deployment, the configuration must include similar lines. A Catalyst 3850 Series switch deployed as only an MA within a campus deployment would not include the mobility group configuration. Instead the CT5760 wireless controller deployed as an MA and MC within a campus would contain the mobility group configuration. Note that since IOS XE based wireless controllers only support the new hierarchical mobility architecture, no configuration is required to enable it.

Guest Controller

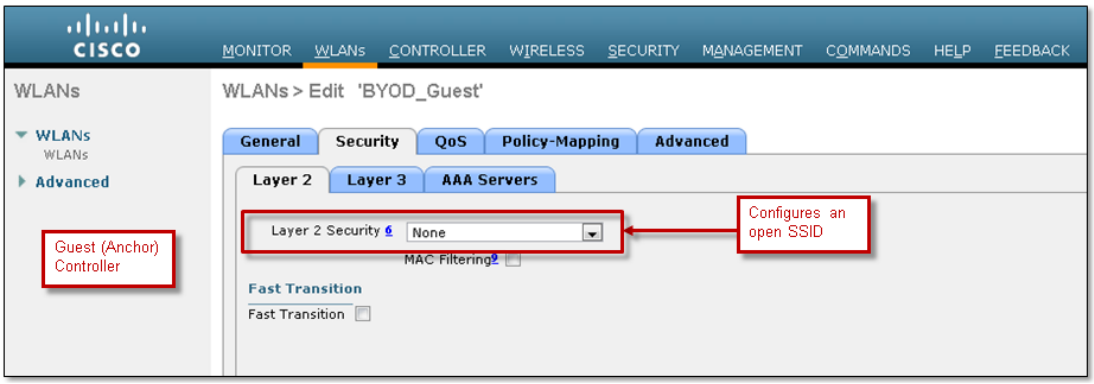
The guest controller is the point where all the guest wireless traffic is terminated. For this version of the design guide, the discussion only includes a CT5508 CUWN wireless controller as the guest controller. As explained in Overview, a mobility anchor tunnel is established between the guest controller and the campus controller. The guest controller authenticates against ISE all guest traffic which is originated from campus or branch controllers. The first step is to define the guest SSID, named BYOD_Guest. The name of this SSID must be identical to the BYOD_Guest defined in the campus controller. Figure 21-12 depicts the details.

Figure 21-12 BYOD_Guest Details on the Guest Controller



The next important tab is the Layer 2 Security details of the BYOD_Guest WLAN, which is shown in Figure 21-13.

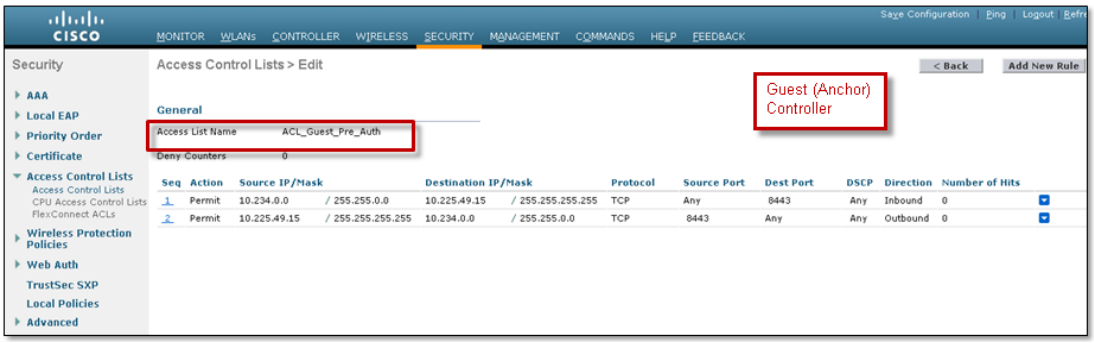
Figure 21-13 Layer 2 Security Details of the BYOD-Guest WLAN on the Guest Controller



Layer 2 security is set for None, indicating an open SSID. The authentication must be configured for Web Auth. Both of these match the configuration of the campus controller.

A Web Auth pre-authentication ACL is necessary when utilizing a remote Cisco ISE guest portal for login and optionally the AUP or EUA. The Web Auth pre-authentication ACL must be configured to allow all possible IP addresses associated with the guest wireless subnet (which can be handed out to guest wireless devices) to be redirected to TCP port 8443 of the Cisco ISE guest portal. An example of a Web Auth pre-authentication ACL is shown in Figure 21-14.

Figure 21-14 Example of a Pre-Authentication ACL for Guest Wireless Access via Web Auth



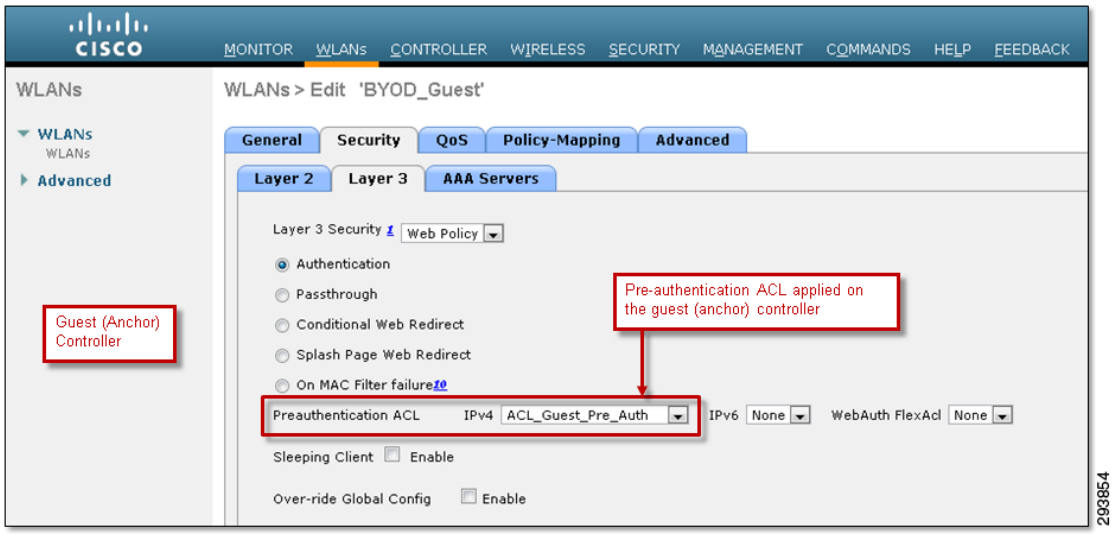
The ACL specifies the following access:

- Allow (do not redirect) traffic from devices on the 10.234.0.0 / 16 network address space to TCP port 8443 of the ISE server (10.225.49.15).

The ACL implicitly denies (redirects) all other traffic to the ISE guest portal. When specifying an ACL down to the port level within the guest wireless controller, both inbound (from the wireless guest devices to the Cisco ISE server) and outbound (from the Cisco ISE server to the wireless guest devices) rules must be configured. Specifying an inbound rule only does not automatically allow return traffic through the wireless controller, as is done with a stateful firewall. Also, specifying a single rule of the form above with a direction of “Any” does not work. The wireless controller does not reverse the source and destination IP addresses for the return traffic.

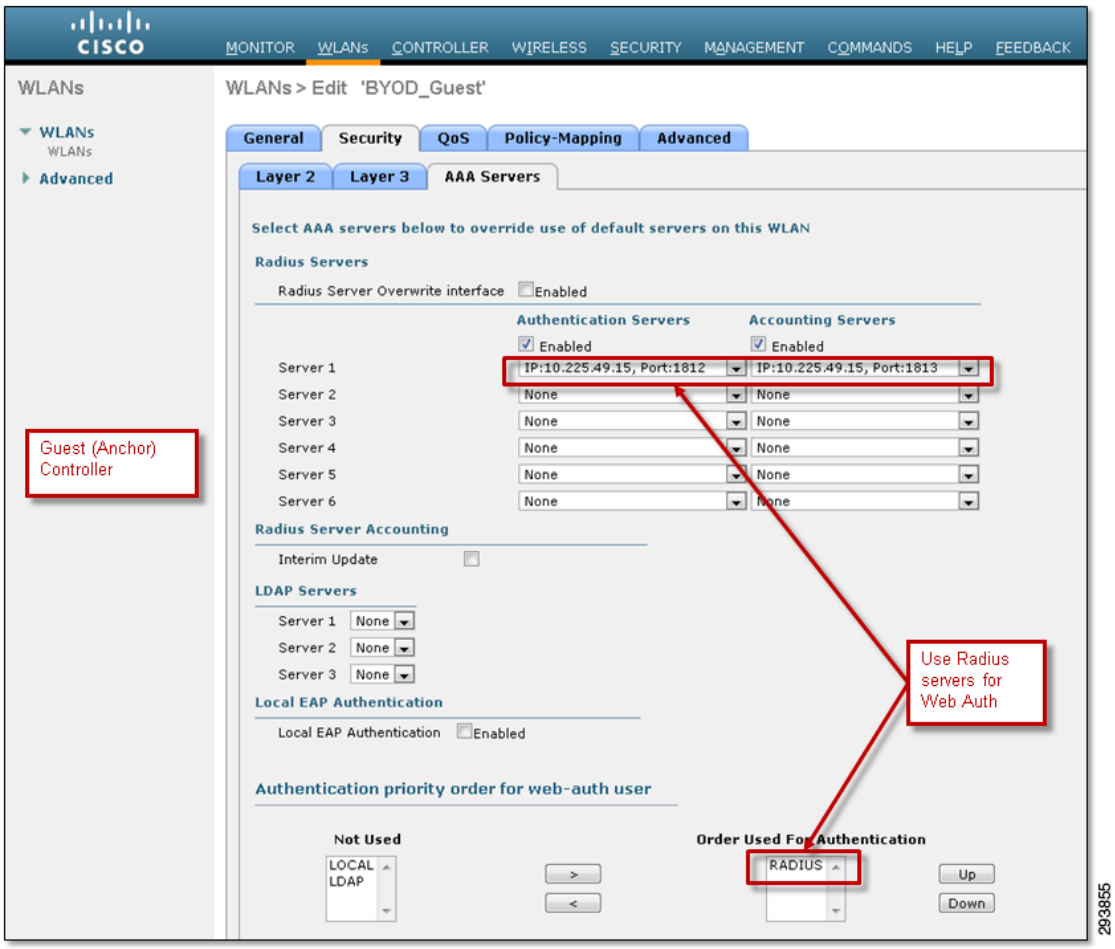
Once the ACL is configured, it must be applied as a Web Auth pre-authentication ACL. This is done in the Guest WLAN Layer 3 Security policy, as shown in Figure 21-15.

Figure 21-15 Applying an ACL as a Web Auth Pre-Authentication ACL



The AAA server configuration details are shown in Figure 21-16.

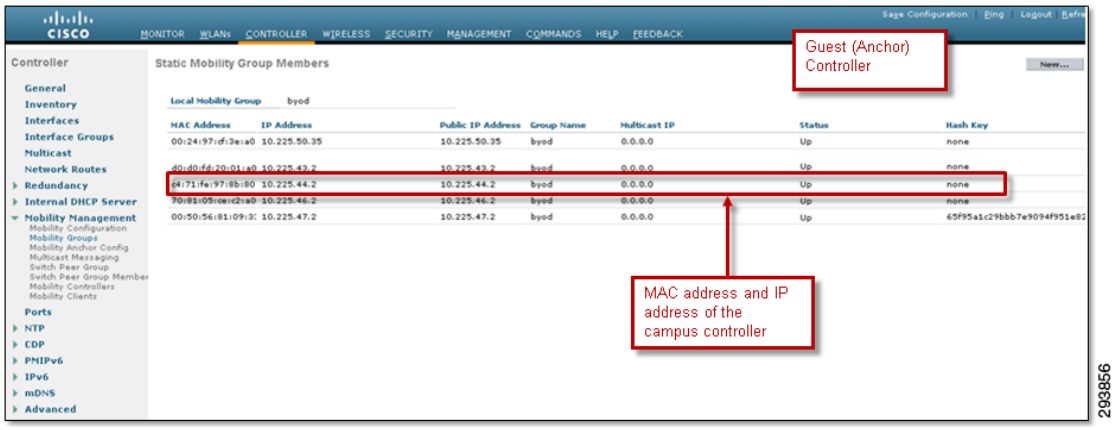
Figure 21-16 AAA Server Configuration for the BYOD-Guest WLAN on the Guest Controller



The AAA server parameters are configured such that Web Auth utilizes the Cisco ISE server for authenticating guests using RADIUS.

The next step is to configure the anchor mobility tunnel between the guest and the campus controller. The campus controller must first be added to the guest controller as a mobility group member. An example is shown in Figure 21-17.

Figure 21-17 Adding the Campus Controller to the Mobility Group



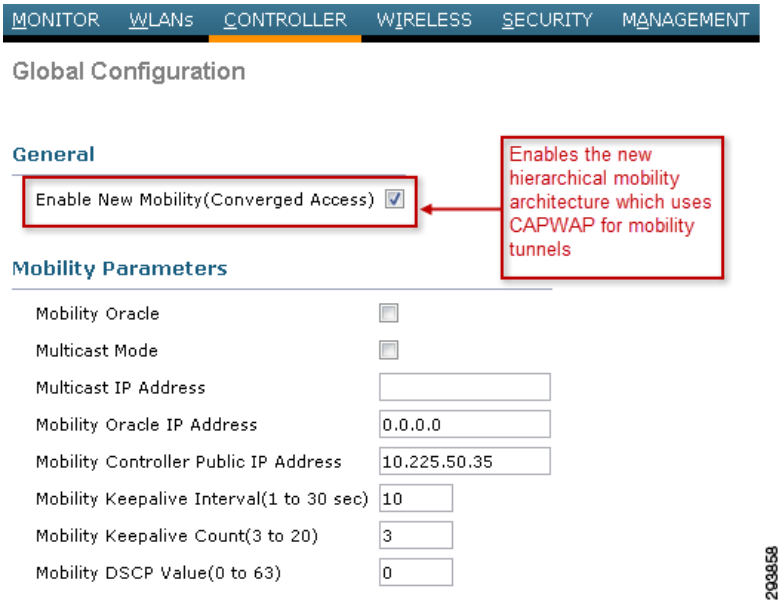
Finally, a mobility anchor is created on the BYOD_Guest SSID. For the guest controller, the mobility anchor points to the local IP address of the management interface of itself. This is different from the campus controller configuration which points to the guest controller. An example is shown in Figure 21-18.

Figure 21-18 Configuring the Mobility Anchor on the Guest Controller



In order to support the new mobility architecture (also referred to as the hierarchical mobility architecture) the network administrator must check the **Enable Hierarchical Architecture** option within the global mobility configuration of the wireless controller. This is shown in Figure 21-19.

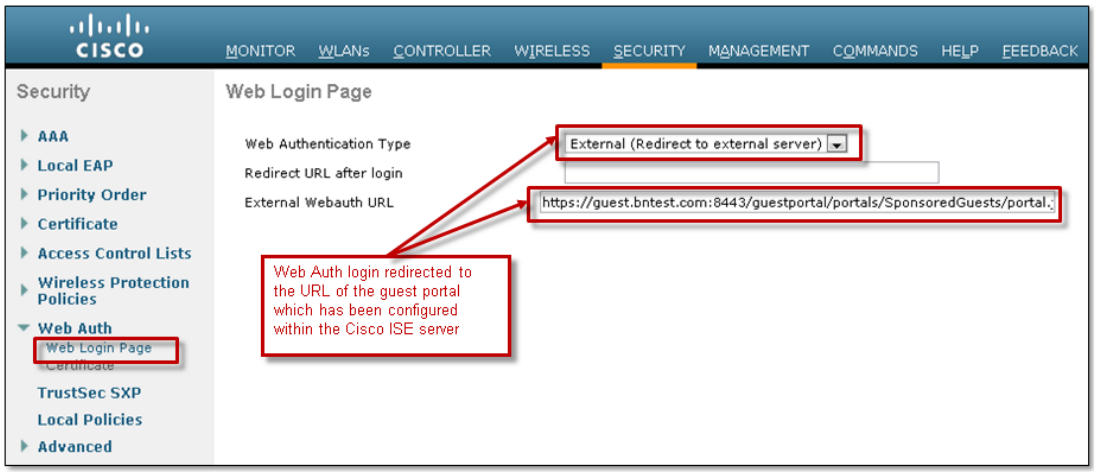
Figure 21-19 Enabling the Hierarchical Mobility Architecture in the Guest Controller



Note Since the Flex 7500 wireless controller does not support the new mobility architecture, this step can be skipped when implementing a guest controller which is auto-anchoring wireless devices from a Flex 7500 foreign controller.

The guest controller authenticates the users against an external server, which is ISE in this design. Hence, the guest controller must be configured to redirect the guest users to the ISE, which is shown in Figure 21-20.

Figure 21-20 Configuration for Redirection to an External Server



In Figure 21-20, the **External Webauth URL** is set to:

https://guest.bntest.com:8443/guestportal/portals/SponsoredGuests/portal.jsp

The name of the server, which in the example above is **guest.bntest.com**, must resolve via DNS to the IP address of ISE, which is 10.225.49.15 in the examples shown in this chapter.

Table 21-1 shows the IP address information of the guest and the campus controllers used in the screen captures and configuration examples shown in this section.

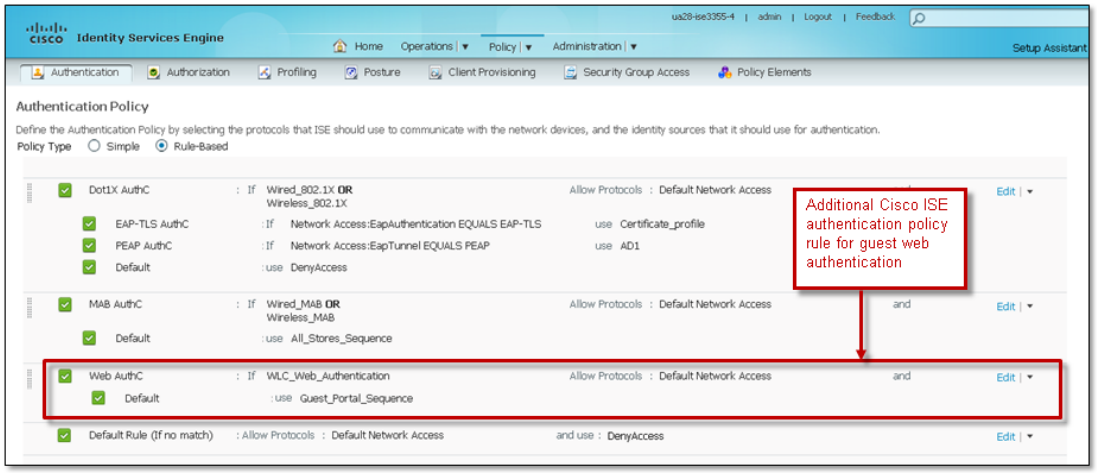
Table 21-1 IP Addresses of Campus (Foreign) and Guest (Anchor) Controllers

Device	Local IP Address	Remote IP Address
Campus CUWN Controller	10.225.44.2	10.225.50.35
Campus Converged Access (IOS XE Based) Controller	10.225.47.2	10.225.50.35
Guest Controller	10.225.50.35	1.225.44.2 and 10.225.47.2

Cisco ISE Policy Configuration

From a Cisco ISE policy perspective, an additional authentication rule needs to be added for guest authentication. This rule allows wireless controller web authentications, originated from the SSID corresponding to the guest WLAN, to utilize a separate Cisco ISE user identity sequence for wireless guest access. An example of such a policy rule is shown in Figure 21-21.

Figure 21-21 Example of Cisco ISE Authentication Policy Allowing Guest Wireless Access



The logical format of the example authentication policy rule is as follows:

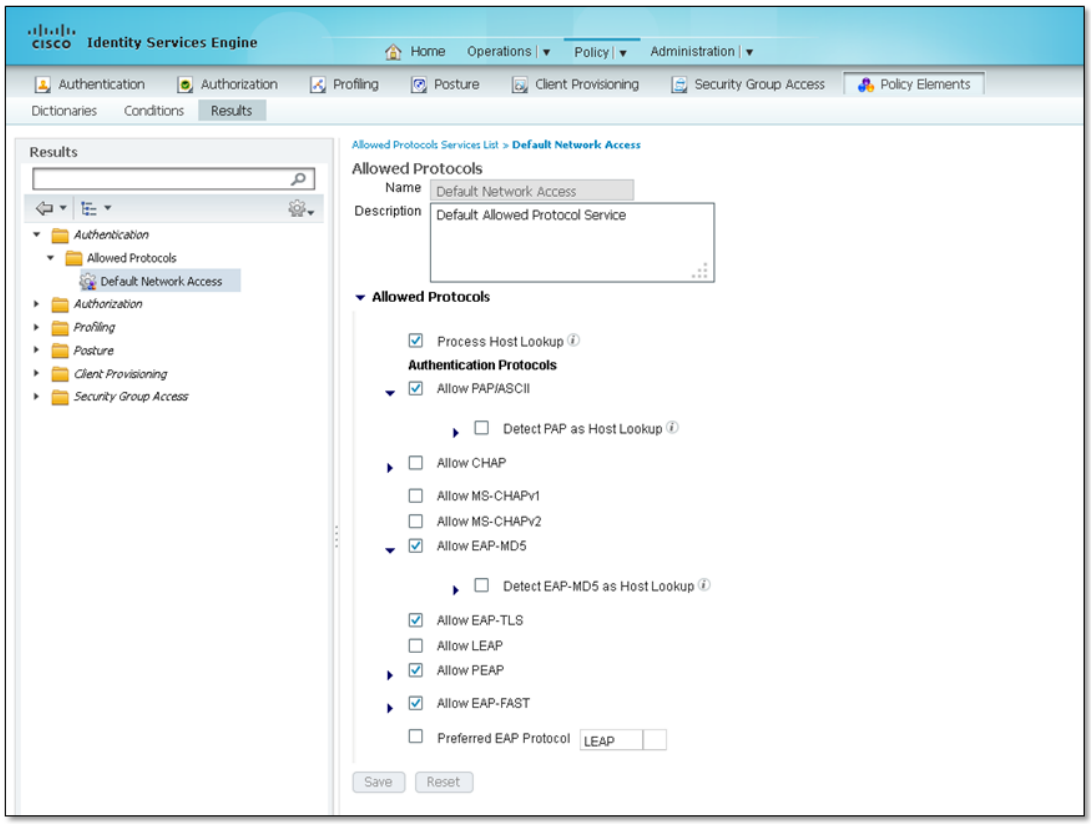
```
IF (WLC_Web_Authentication)
    THEN (Allow Default Network Access AND USE Guest_Portal_Sequence)
```

WLC_Web_Authentication is a system-generated compound condition which is used here to match Web Auth requests from Cisco Wireless LAN Controllers. It matches the following two standard RADIUS dictionary attribute-value (AV) pairs:

```
Service-Type - [6] EQUALS Login
NAS-Port-Type - [61] EQUALS Wireless - IEEE 802.11
```

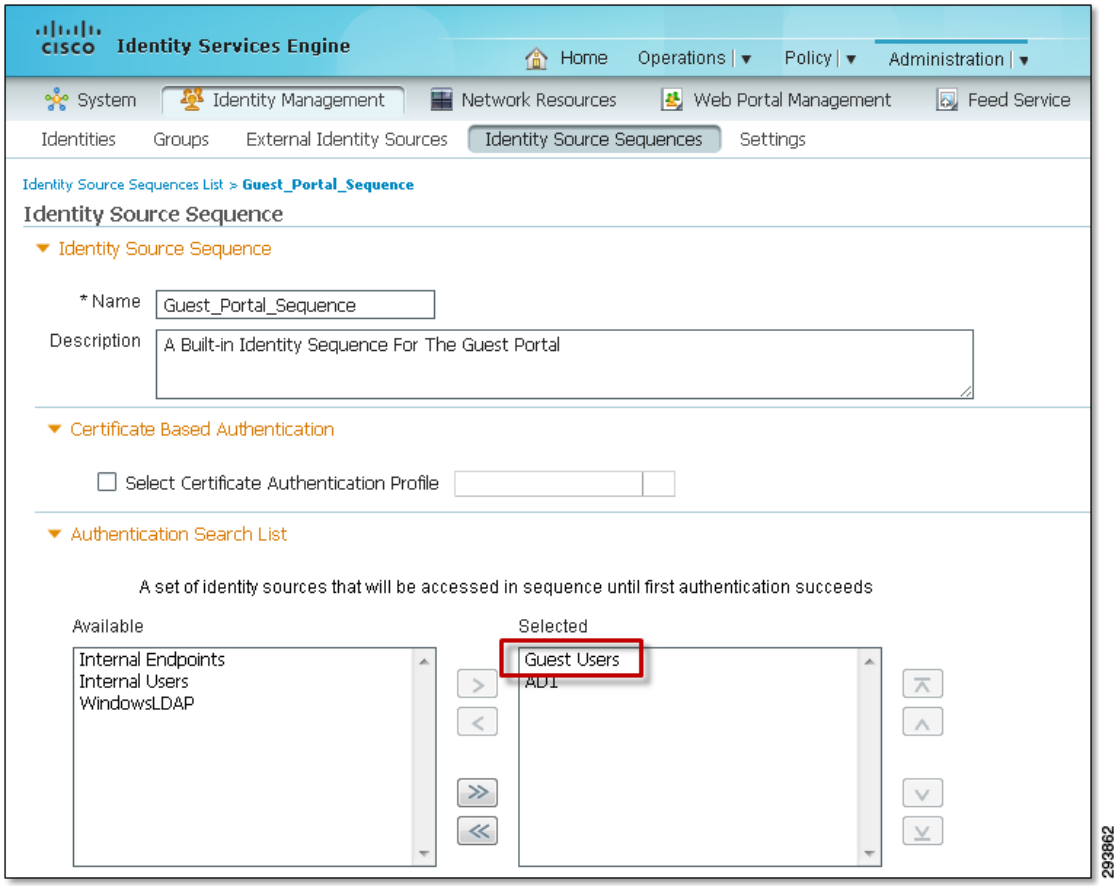
Default Network Access is a system-generated authentication result, which allows various protocols to be used for the Web Auth. An example is shown in Figure 21-22.

Figure 21-22 Example of Allowed Protocols Under Default Network Access



Guest_Portal_Sequence is a user-defined identity source sequence. An example is shown in Figure 21-23.

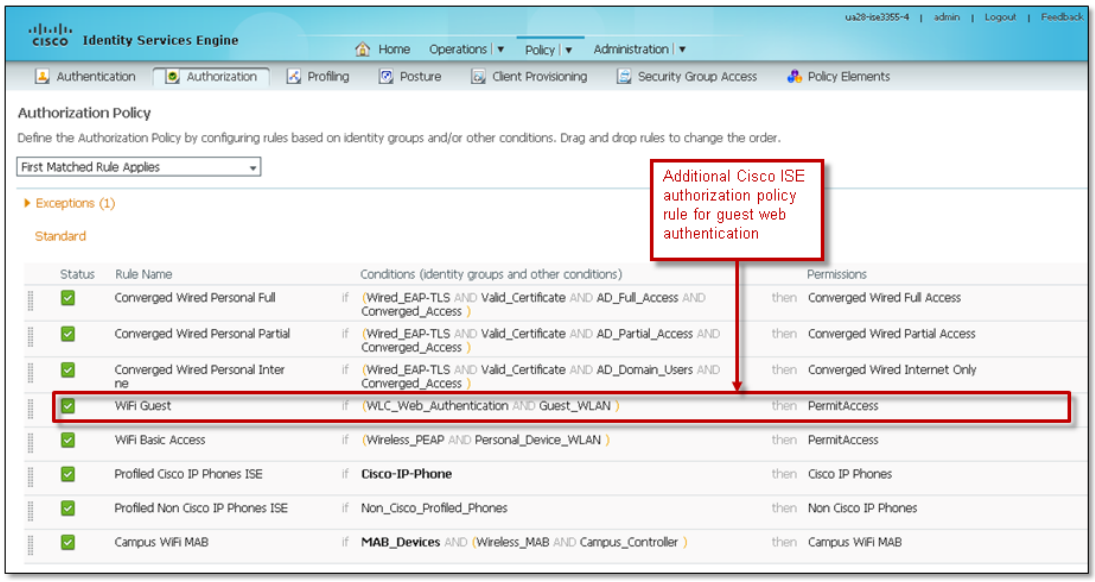
Figure 21-23 Example of Guest_Portal_Access Identity Source Sequence



Guest_Portal_Sequence in the example above uses the Guest Users identity source as the primary source and uses the AD1 group as the next source. Guest Users is a system generated identity source, which is a new feature beginning with ISE 1.2. This identity source is a place where guest credentials are held when they are configured through the Cisco ISE sponsor portal, which is discussed later in this chapter. Although an identity source sequence is not strictly needed when only a single identity source is specified, configuring a sequence allows guest wireless access to be easily extended to include employee personal devices by adding an additional identity source.

From a Cisco ISE policy perspective, an additional authorization rule also needs to be added for guest users. This rule permits access for wireless controller web authentications originated from the SSID corresponding to the guest WLAN. An example of the policy rule is shown in Figure 21-24.

Figure 21-24 Example of Cisco ISE Authorization Policy Allowing Guest Wireless Access



The logical format of the example authorization policy rule is:

```
IF (WLC_Web_Authentication AND Guest_WLAN
    THEN Permit Access
```

WLC_Web_Authentication was discussed with regard to the authentication policy above.

Guest_WLAN is a user-defined simple authorization condition for guests accessing the Internet via web authentication through the WLAN corresponding to the open guest SSID. It matches the following RADIUS AV pair from the Airespace dictionary:

```
    Airespace-Wlan-Id - [1] EQUALS 2
```

The Airespace-Wlan-Id is the identification number (WLAN ID) of the WLAN corresponding to the Guest SSID, as shown in Figure 21-25.

Figure 21-25 Example Guest Wireless Controller WLAN IDS



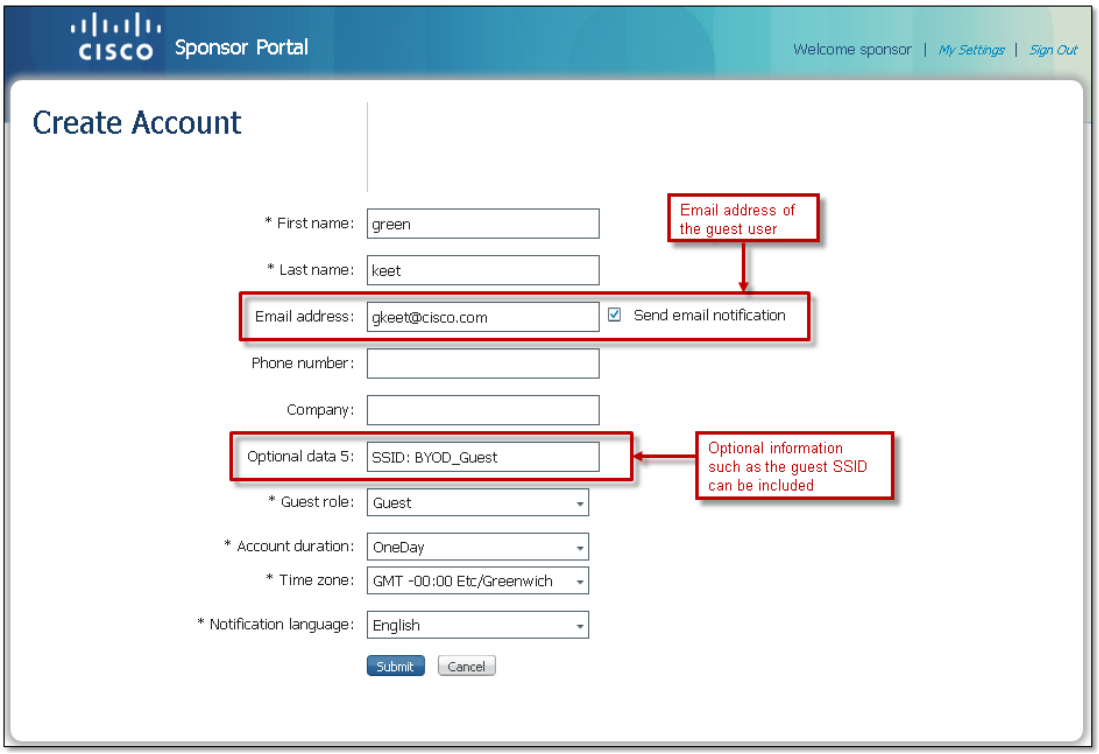
This allows the ISE authorization policy to differentiate Web Auth requests coming from the guest WLAN and permit them.

Note Simple Conditions such as Guest_WLAN are optionally used to give attribute and value pairs a descriptive name. This allows the policy to be more readable and easier to support.

Cisco ISE Sponsor Portal

The Cisco ISE sponsor portal can be accessed at: https://ISE_server:8443/sponsorportal/, where ISE_server is either the IP address or the name of the Cisco ISE server. An example of the web page for creating guest credentials within the Cisco ISE sponsor portal is shown in Figure 21-26.

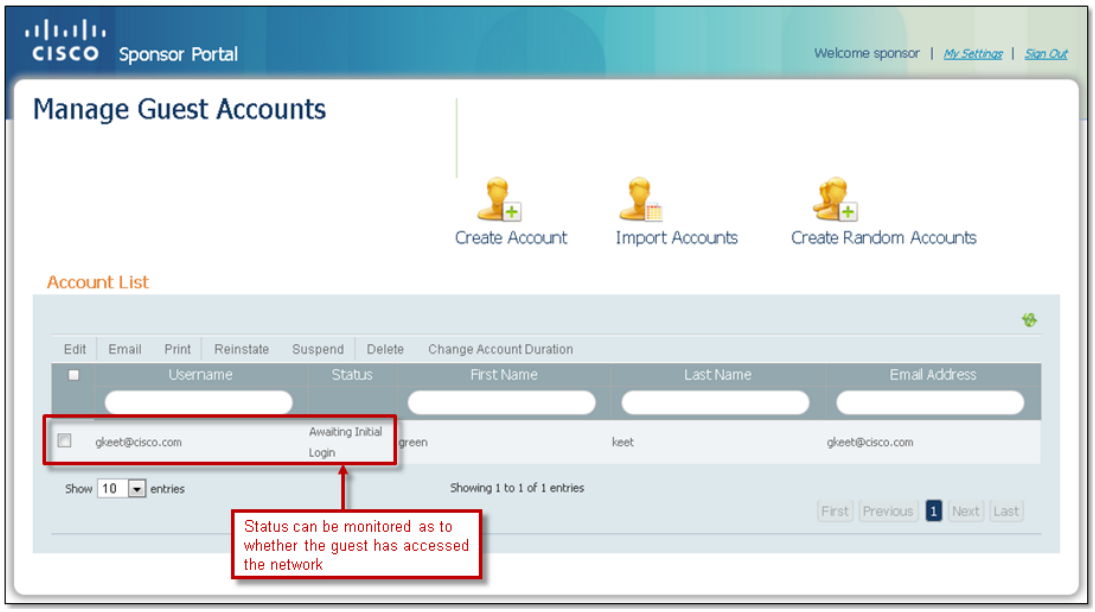
Figure 21-26 Creating Guest Credentials on the Cisco ISE Sponsor Portal



Information such as guest's company name, the guest's email address and phone number, as well as optional user-defined data can be included. Optional data could include the WLAN SSID the guest needs to connect to (if the SSID is hidden), as well as the name, phone number, and department of the sponsor. Depending upon the allowed time profiles, the credentials can be configured to become active at a future date and time and remain active for a period of time. The Cisco ISE sponsor portal also has the capability to deliver the guest credentials to the guest prior to arrival via email or SMS. Sending credentials via email helps ensure the guest has provided a valid email address.

Once guest credentials are created, they can be monitored and managed by the sponsor via the Cisco ISE sponsor portal, as shown in Figure 21-27.

Figure 21-27 Monitoring Guest Credentials from the Cisco ISE Sponsor Portal



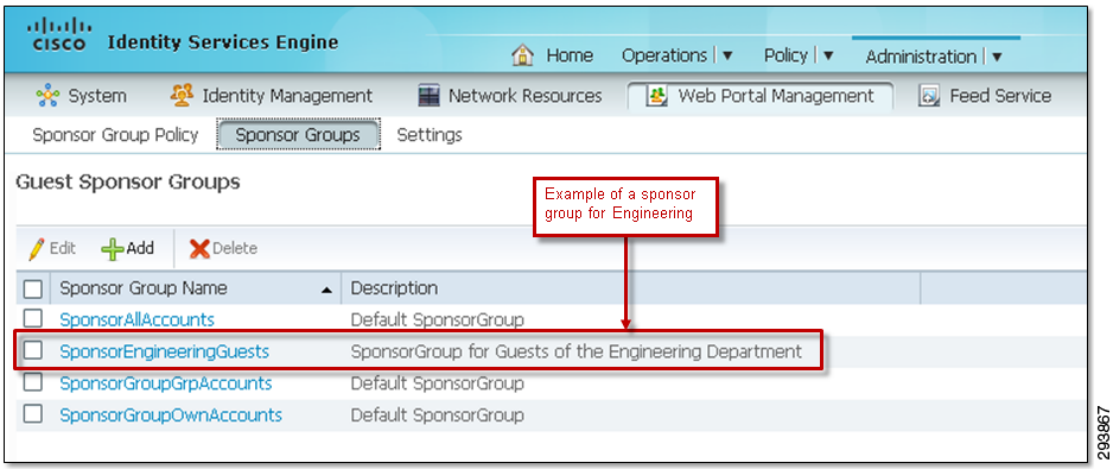
Note that in Figure 21-27 the guest username was based upon an email address versus just the first and last name of the guest. Chapter 18, "BYOD Basic Access Use Case" discusses extending guest wireless access to allow employee personal devices as well. Use of the email address within the guest username is one possible way to differentiate between guests and employees who may have the same first and last names.

Configuring the Cisco ISE Sponsor Portal

Configuration of the Cisco ISE sponsor portal is done through the Web Portal Management section of the Cisco ISE server. Different levels of sponsor responsibility can be created, ranging from individual sponsors who can only view and edit guest accounts they have created, to group sponsors who can view and edit guest accounts for a particular group, to sponsors who can view and edit all guest accounts.

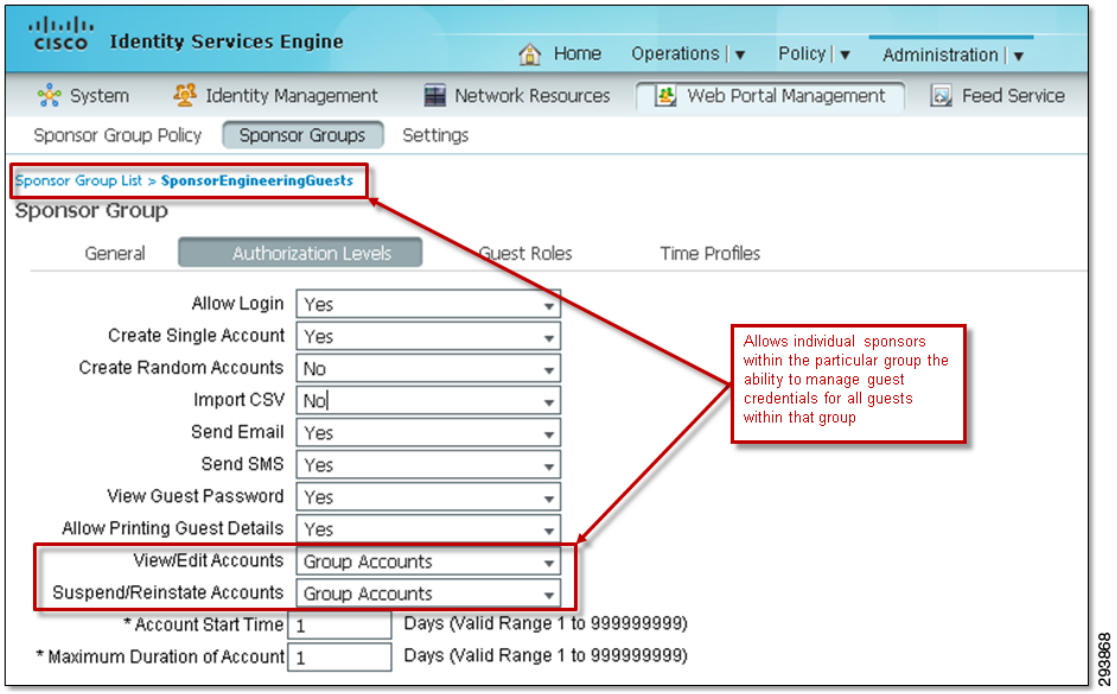
Multiple sponsor groups, each with their own members, can be created through the Sponsor Groups tab under the Web Portal Management section of the Cisco ISE server. Figure 21-28 shows an example where a separate group has been added for guests sponsored by the Engineering department.

Figure 21-28 Example of Multiple ISE Sponsor Groups



Different authorization parameters can then be configured for each sponsor group by selecting the particular sponsor group and selecting the **Authorization Levels** tab, as shown in Figure 21-29.

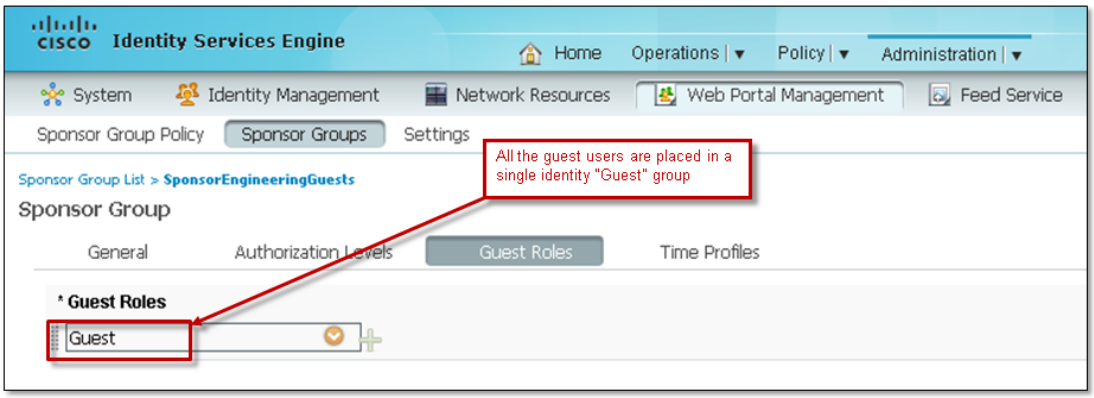
Figure 21-29 Example of Authorization Levels for an Individual Sponsor Group



This example shows a configuration where any member of the sponsor group is allowed to view, edit, suspend, and reinstate a guest credential created by any other member of the sponsor group. However, members of different sponsor groups cannot modify guest credentials created for this group.

The Guest Roles tab is used to select the user identity group (i.e., guest credential database) into which the guest credentials created by a member of this sponsor group are placed. An example is shown in Figure 21-30.

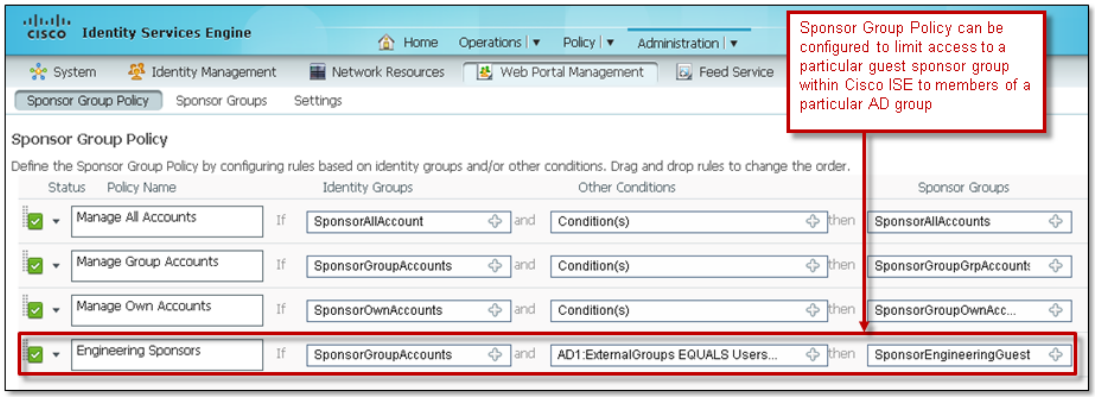
Figure 21-30 Example of Guest Roles for an Individual Sponsor Group



The Time Profiles tab allows the network administrator to determine which time profiles (either default or pre-configured within ISE) are applied to the particular sponsor group.

Once the sponsor groups are created, the Sponsor Group Policy tab can be used to create policies controlling who has access to which sponsor groups. More commonly, the organization may wish to leverage existing Microsoft Active Directory groups to differentiate among different sponsors. Figure 21-31 shows an example of this.

Figure 21-31 Example of Microsoft AD for Sponsor Group Membership



In this example, access to the sponsor group is limited to those members of the Microsoft Active Directory domain who are members of the group called “Users/uatest.com”. Note that the Microsoft Active Directory server must be configured as an external identity source to select this option. In this example it is known by the name “AD1”.

By tightly controlling members of the Microsoft AD groups which have sponsor access to ISE, the network administrator can limit the use of the guest wireless network to its original intended purpose—guest wireless access—instead of employee personal devices, if desired.

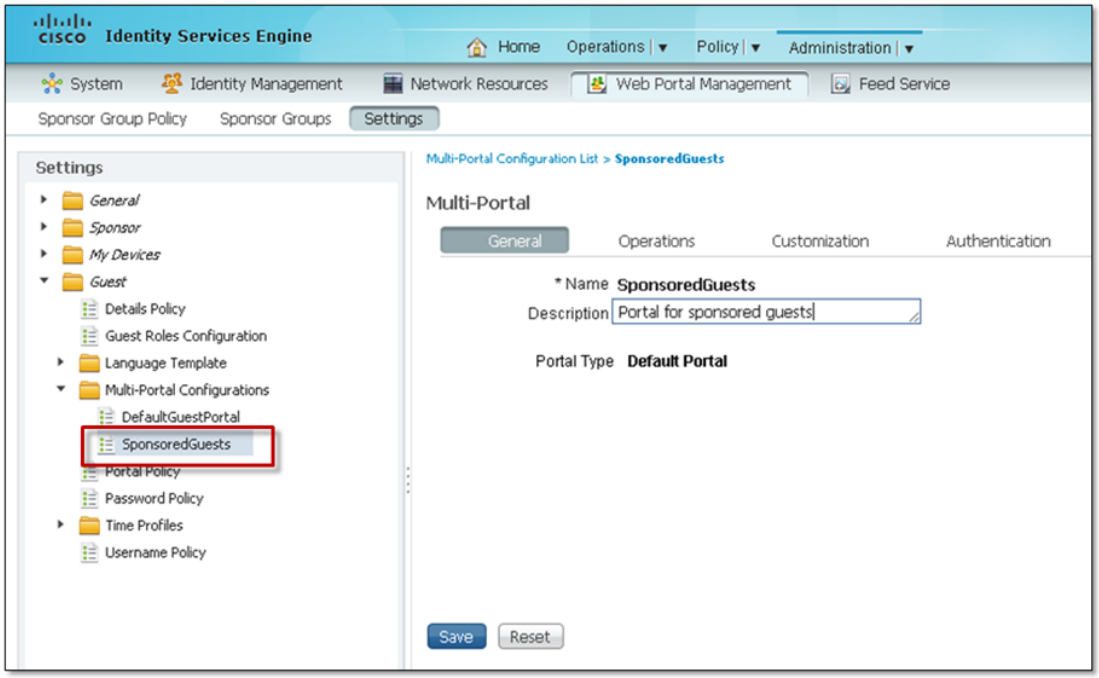
Cisco ISE Guest Portal

As mentioned previously, Cisco ISE has the capability to support multiple guest portals. The Cisco ISE server has a system-generated DefaultGuestPortal configuration. This allows the network administrator to provision a guest portal in order for employees or IT staff to on-board corporate-owned or employee personal devices, as discussed in Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices" and Chapter 16, "BYOD Limited Use Case—Corporate Devices."

Configuring the Cisco ISE Guest Portal

An additional guest portal for wireless guest access can be defined through the Guest > Multi-Portal Configurations. An example is shown in Figure 21-32.

Figure 21-32 Example Multi-Portal BYOD Deployment



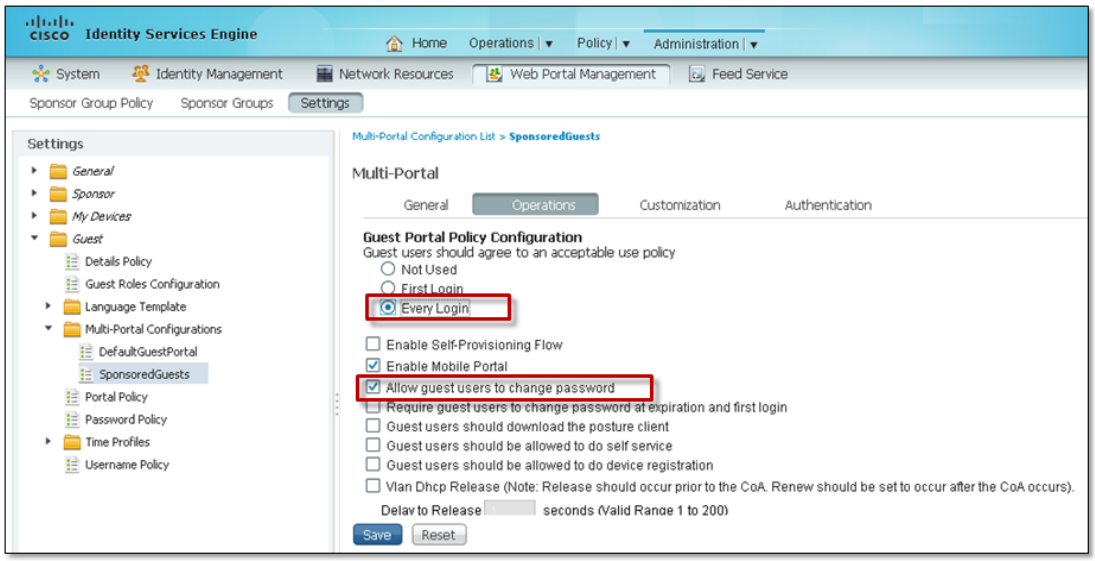
When a user-defined guest portal is implemented, the URL which needs to be configured within the guest wireless controller Web Auth Web Login Page, as shown in Figure 21-20:

http://*ISE_server*:8443/guestportal/portals/*name_of_user-defined_portal*/portal.jsp

ISE_server is either the IP address or the name of the Cisco ISE server. *Name_of_user-defined_portal* is the name of new user-defined guest portal, which is SponsoredGuests in the example above.

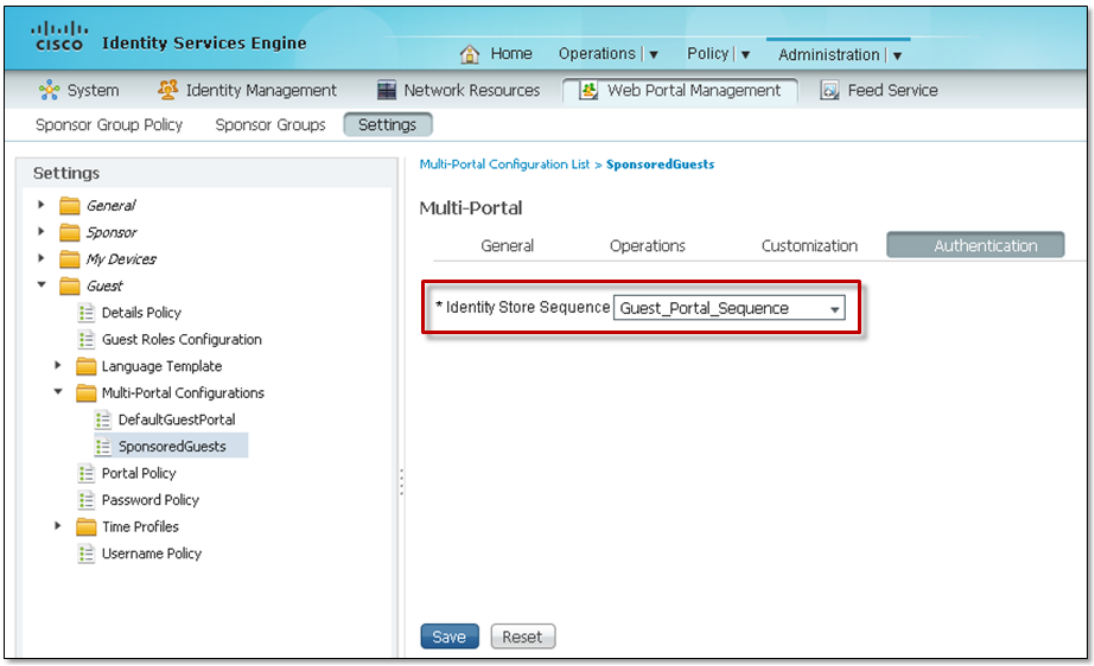
Once the new guest portal is defined, the Operations tab can be used to display an Acceptable Use Policy (also known as an End User Agreement or EUA), as well as control whether the guest can or must change the sponsor provisioned password. Note that the Operations tab can also be used to force the guest to register their devices with the Cisco ISE server before accessing the Internet from the guest wireless network. This design guide assumes that the guest device itself is not considered in the decision to allow access to the guest wireless network. Hence, this use case is not discussed. Figure 21-33 shows an example of the Operations tab.

Figure 21-33 Example of Operations Tab



The Authentication tab determines which identity source sequence is used for the guest credentials. An example is shown in Figure 21-34.

Figure 21-34 Example of Authentication Settings for a User-Defined Guest Portal

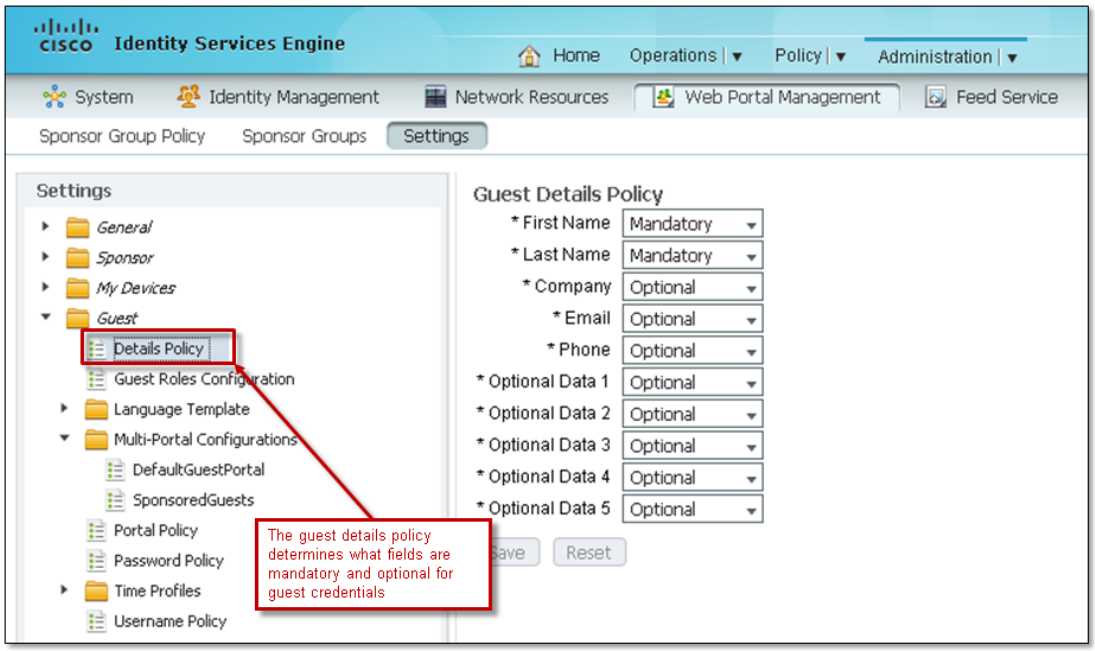


For this example the identity source sequence called **Guest_Portal_Sequence** is chosen. This identity source sequence utilizes **Guest Users** when only wireless guest access is deployed, as shown in Figure 21-23. This allows guests credentials to pass both the guest portal access and the ISE authentication policy. This configuration also allows guest access to be easily extended to include employee personal devices by simply adding Microsoft Active Directory identity store, as discussed in Chapter 18, "BYOD Basic Access Use Case."

Note Cisco ISE authentication logs may show the guest user authentication appearing twice with this configuration, although the guest is only authenticated once via Web Auth.

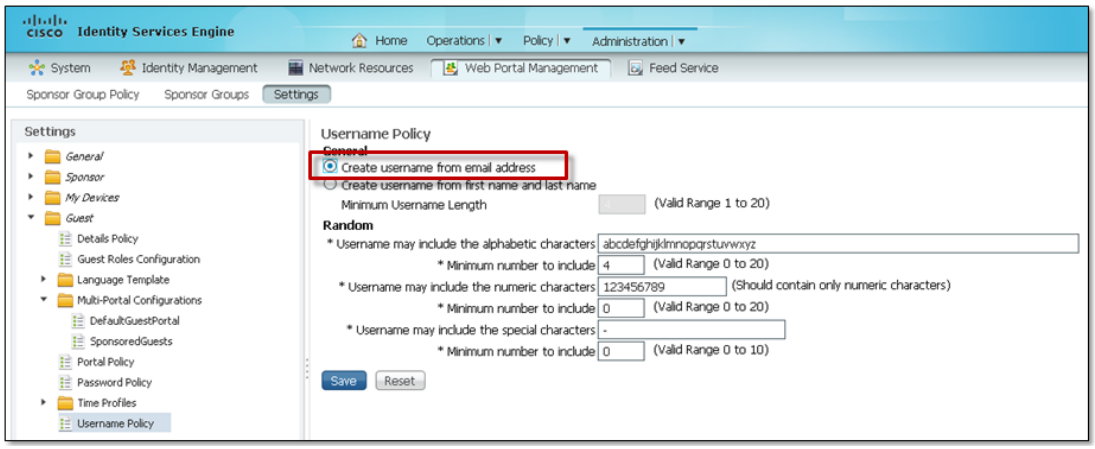
The Guest Details Policy is used to configure additional global guest parameters, including mandatory and optional parameters. An example is shown in Figure 21-35.

Figure 21-35 Example of Guest Details Policy



Additional web pages under the Guest folder control other global guest configuration parameters, such as Username Policy and Password Policy. The Username Policy is where the guest username can be selected to be based upon their Email address, as shown in Figure 21-36.

Figure 21-36 Example of Guest Username Policy

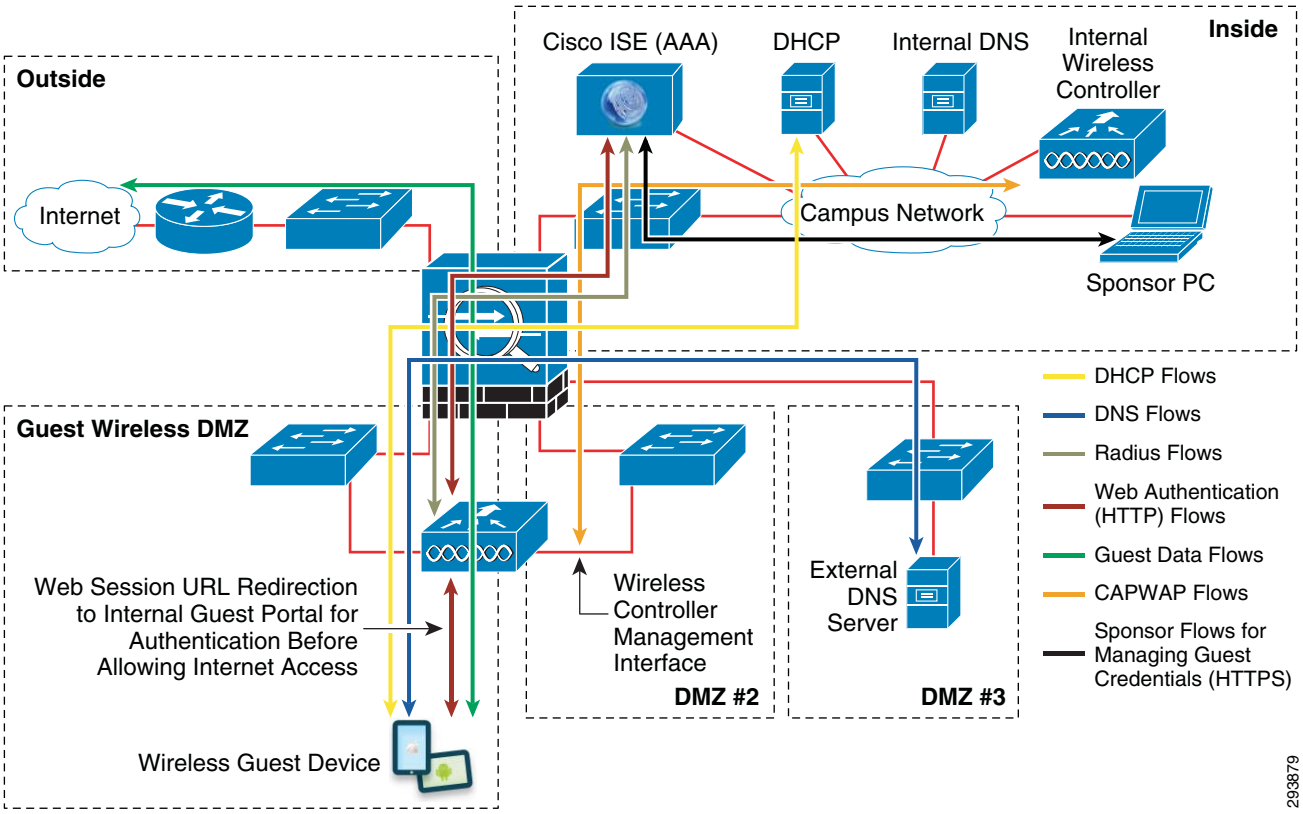


Finally, the Time Profiles folder can be used to select one of the existing time profiles for guest user access or to create a custom time profile. Time profiles are selected by the sponsor when configuring guest credentials to control when the guest user has access to the network and for how long.

ASA Firewall Configuration

Figure 21-37 shows an example of the flows that need to pass through the Cisco ASA firewall to support the design discussed in this chapter.

Figure 21-37 Example of Flows that Need to Pass Through the Cisco ASA Firewall



This design requires a RADIUS session to be allowed through the ASA firewall between the guest wireless controller and the Cisco ISE server. In addition, this design requires the guest web session to be re-directed and allowed through the ASA firewall to the inside of the network, where the Cisco ISE server sits. By default Cisco ISE uses TCP port 8443 for the guest portal. When using the older mobility architecture, an Ethernet-over-IP (IP protocol 97) auto-anchor mobility tunnel, as well as the WLAN control port (UDP port 16666) between the management interfaces of the two wireless controllers, must still be allowed through the ASA firewall. When using the new hierarchical mobility architecture, a CAPWAP (UDP port 5246 for control and UDP port 5247 for data) auto-anchor mobility tunnel between the management interfaces of the two wireless controllers must be allowed through the ASA firewall. Besides allowing DNS, DHCP (assuming the deployment of an internal DHCP server), and TCP port 8443 for the HTTPS redirection, the ASA firewall should be configured to block all other traffic generated from guest wireless devices onto the internal network.

Table 21-2 summarizes the relevant ports that need to be allowed through the ASA firewall.

Table 21-2 Ports to be Allowed through the ASA Firewall

Application	Transport	Port
WLAN Control	IP Protocol 97	-
	UDP	16666
WLAN Control	UDP	16666
	UDP	16667
ISE Guest Portal	TCP	8443
DNS	UDP	53
BOOTPS (DHCP)	UDP	67
BOOTPC (DHCP)	UDP	68
CAPWAP Control Channel (new mobility architecture)	UDP	5246
CAPWAP Data Channel (new mobility architecture)	UDP	5247

Additional Considerations

When implementing guest wireless access for devices such as Apple iOS or Mac OS X Lion, the network administrator should be aware that these devices have implemented a feature which automatically detects the presence of a captive portal deployment. It does this by generating an HTTP request to an Apple website and looking for a response. If a redirect is received, then a captive portal deployment is assumed. This feature only applies to SSIDs which have open access, as is typical with most guest wireless networks. When a captive portal deployment is detected, the iOS or Mac OS X Lion device automatically displays a dialog window for authentication without the end user having to launch the web browser. This feature is intended to make it easier for non-browser based applications to access the Internet, without the end user having to launch a web browser, by performing Web Auth via the pop-up window. Many HTML-based mobile applications do not use the browser as the user interface. This is known as Captive Portal Network Assistance (CPNA) and is effectively a light weight HMTL-based user interface. Unfortunately the interface is not properly interacting with the iOS profiler manager. The symptoms are different based on the version of iOS. In iOS5, the user was not allowed to install the WiFi profile without canceling the CPNA, forcing the device off the provisioning SSID. In iOS6, the user is automatically brought to the profile manager, but after installing the profile, the user is not returned to the CPNA to receive the certificate. In both cases, the CPNA is not able to successfully on-board the device

Cisco wireless controllers have implemented a workaround that bypasses this feature, allowing Apple iOS or Mac OS X Lion devices to operate within a captive portal deployment with HTTPS connectivity to a guest portal with a self-signed certificate. For CUWN wireless controllers, the network administrator needs to establish an SSH session to the guest wireless controller and issue the following command:

```
configure network web-auth captive-bypass enable
```

Note Cisco has been made aware of potential incompatibilities introduced by Apple iOS 7. We are working to understand the limitations and design updates will be made to this publication.

For IOS XE wireless controllers, the network administrator needs to add the following command to the global configuration of the CT5760 wireless controller or Catalyst 3850 Series switch:

```
captive-portal-bypass
```

This command causes the wireless controller to answer back the HTTP request, spoofing the iOS or Mac OS X Lion device into thinking that there is no captive portal deployment. Once the end user opens a browser and attempts to navigate to any site, they are redirected to the portal and prompted for credentials using the normal Web Auth process. Note that non-browser based applications are not able to access the network until the end user opens a web browser and proceeds through the normal Web Auth process. This includes HTML-based applications such as WebEx.

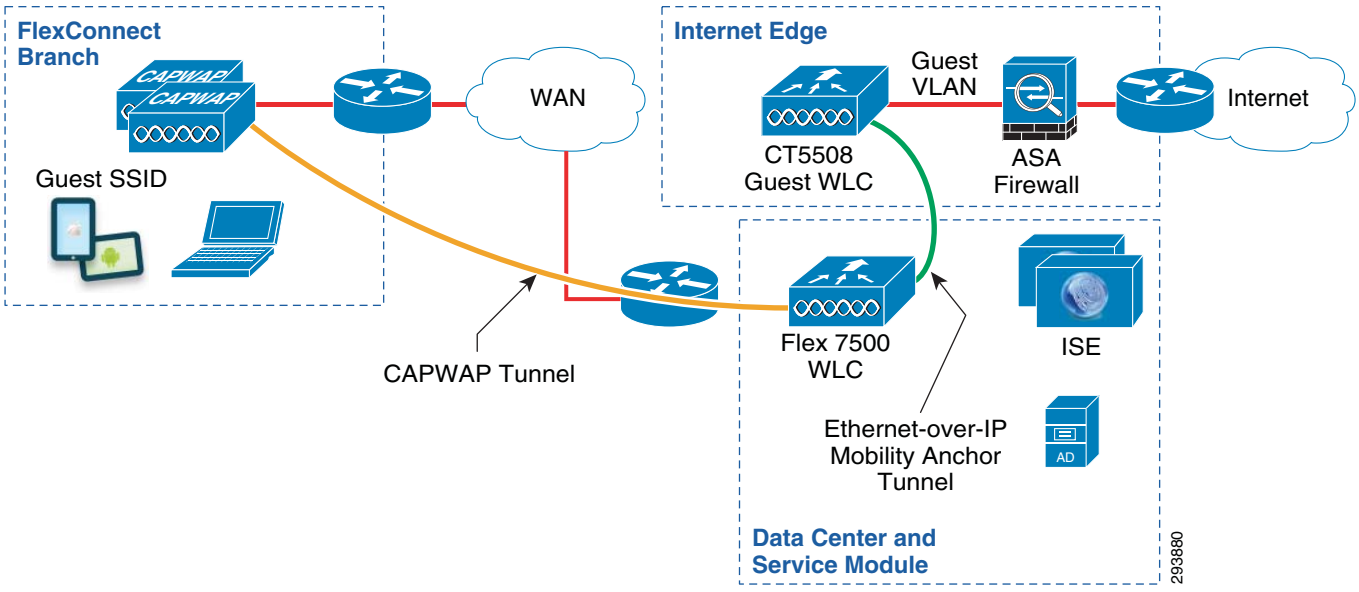
Wireless Guest Access at the Branch

Branch networks frequently offer wireless guest services. There are two basic architectures that can be deployed. The first is a centralized model in which all branch wireless guest traffic is tunneled via CAPWAP to a central controller located within the campus, known as the foreign controller. Wireless guest traffic is then further tunneled via a mobility anchor tunnel to an anchor controller located in the DMZ. This is the method presented in this design guide.

An alternate method is to use either FlexConnect or a converged access infrastructure to locally terminate guest traffic in a secure segment located within the branch. The advantage of the second approach is that guest traffic does not consume expensive corporate WAN bandwidth. Instead guest traffic is isolated within the branch and uses a local branch Internet path. Future versions of this design guide may explore this option. In addition, there are many other possible WAN deployment models that may be leveraged to provide guest users with access to the Internet. A collection of white papers that explain various WAN architectures is available at: http://www.cisco.com/en/US/netsol/ns816/networking_solutions_white_papers_list.html.

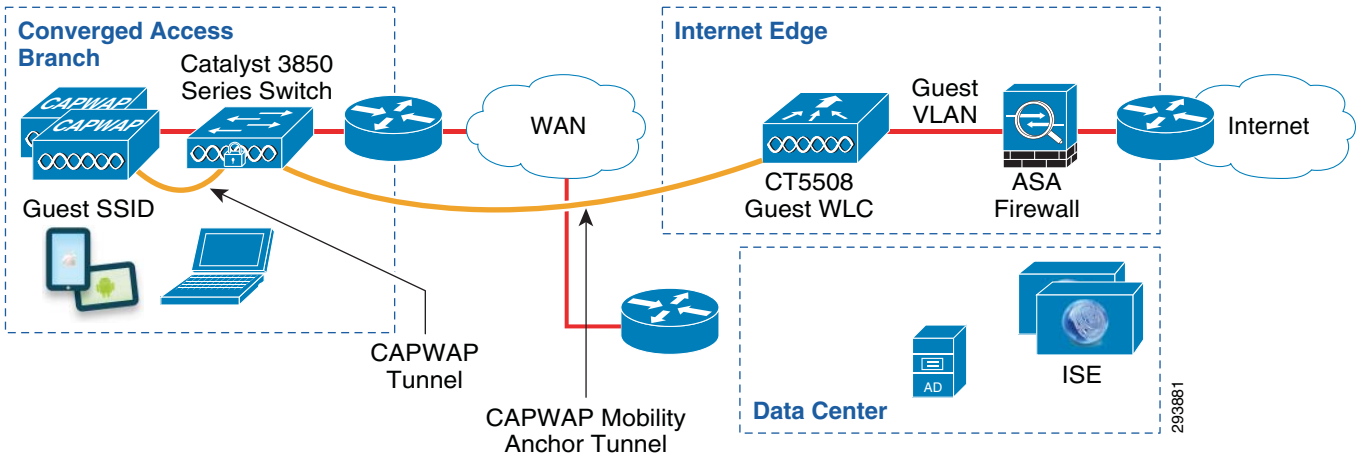
The guidance presented here follows a centralized model. For FlexConnect wireless designs, the FlexConnect wireless controller which services branch locations and which provides on-boarding to branch BYOD devices is also used as a foreign controller to tunnel wireless guest traffic to a guest wireless controller within the campus Internet edge. Figure 21-38 shows the various components required for this model.

Figure 21-38 Guest Wireless Access at the FlexConnect Branch



For Converged Access designs, the Catalyst 3850 Series switch which serves as the wireless controller for the branch locations, and which provides on-boarding to branch BYOD devices, is also used as a foreign controller to tunnel wireless guest traffic to a guest wireless controller within the campus Internet edge. Figure 21-39 shows the various components required for this model.

Figure 21-39 Guest Wireless Access at the Converged Access Branch



Note When deploying converged access wireless designs in which the Catalyst 3850 Series switch functions as the Mobility Controller (MC) and Mobility Agent (MA), it should be noted that the mobility tunnel for wireless guest access initiates from the Catalyst 3850 to the Guest anchor controller located within the DMZ. Hence each branch will initiate a mobility tunnel for wireless guest access with this design. The maximum number of mobility controllers within a mobility domain is 72 for the CT5508 wireless controller. Therefore the maximum number of mobility anchor tunnels is limited to 71 for the CT5508 wireless controller. Therefore the network administrator may need to deploy additional CT5508 guest anchor controllers. Alternatively, the network administrator may look at providing direct Internet access from the branch for guest access. Future versions of this design guide may address such designs.

Because separate wireless controllers are deployed for campus and branch wireless access, the same guest SSID can be configured on both wireless controllers, but with different characteristics, such as rate limiting. This is one advantage of deploying separate wireless controllers for branch and campus locations.

Due to the limited amount of WAN bandwidth available at branches, network administrators often have the requirement to limit the amount of bandwidth that guest users can utilize below that which guest users can utilize within a campus. The next section discusses rate limiting of wireless guest traffic at the branch. Most other aspects of branch wireless guest access are essentially carried over from the campus wireless guest design. For example, branch wireless guests can continue to use Cisco ISE for the guest portal. Logically the wireless topology for branch guest traffic is the same as wireless access for campus guest traffic. The main difference is that the capacity of the transport will vary over the guest SSID to a larger extent than what would be expected in the campus where the physical path is typically supported by gigabit Ethernet.

Rate-Limiting Guest Wireless Access

Note This section applies only to CUWN wireless controller platforms. Future versions of this design guide may extend the discussion to Converged Access (IOS XE based) wireless controllers.

The prevalence of mobile devices and the expectation of universal network access have resulted in a steady increase in the loads on the guest network. This solution offers rate-limiting tools that can be used to manage these loads. Rate limiting can be configured in various ways-per-user or per-SSID as well as upstream and downstream.

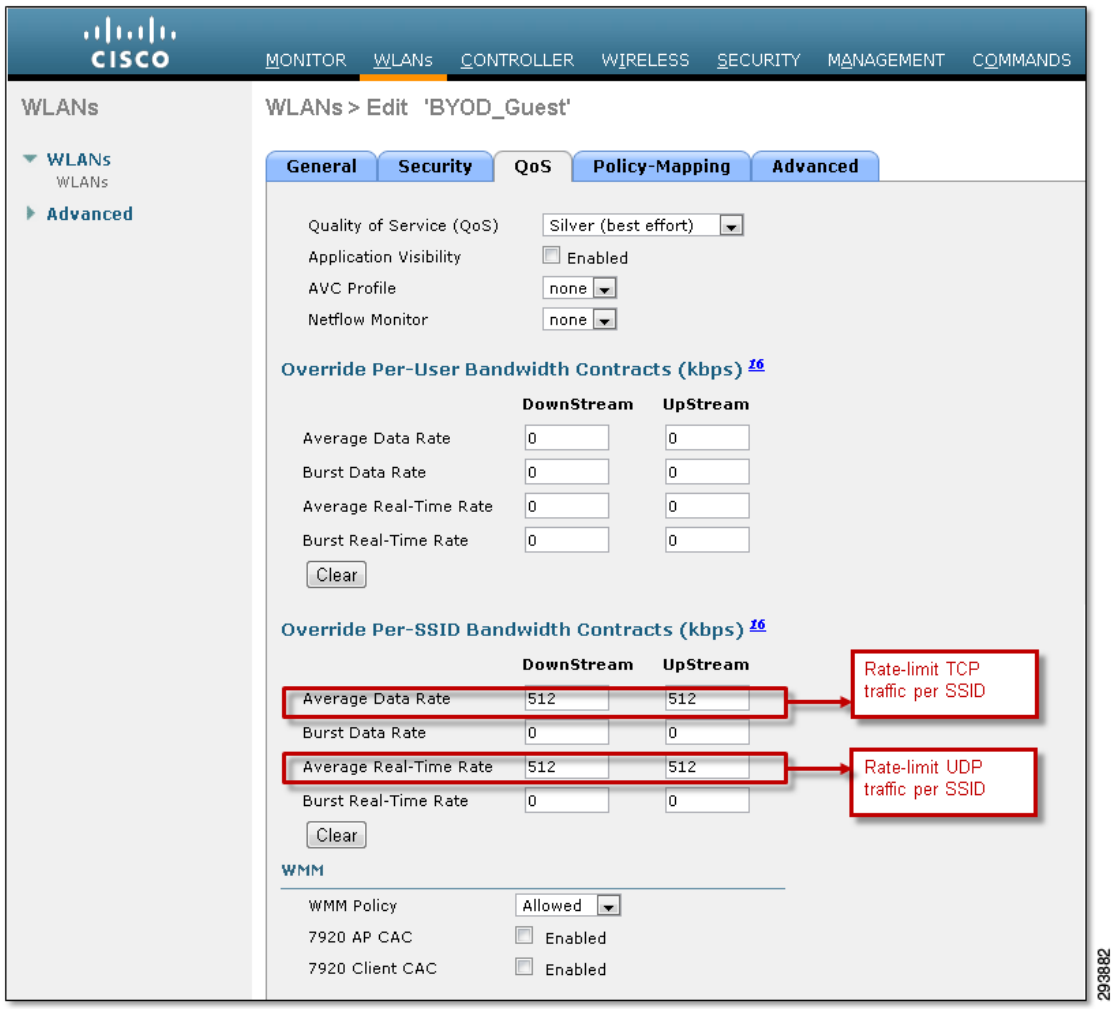
Note Per-SSID rate limiting is actually per-BSSID, since the rate limiting is per SSID per access point per radio. However this design guide refers to this as per-SSID rate limiting.

Per-user rate limiting applies to each specific wireless device. Per-SSID is an aggregate rate shared by all devices within a given SSID. In both cases, upstream rate limiting occurs on the radio. Downstream per-SSID rate limiting also occurs on the radio while downstream per-user rate limiting occurs on the wireless controller.

Rate-limiting in this context is analogous to policing. Packets determined to be in excess of the configured rate are dropped and not metered or buffered. Policers implement a token bucket. The bucket is credited with tokens at a rate that equals CIR. When the bucket is full, no additional tokens are added. Tokens are removed from the bucket when a packet is transmitted, provided a token is available. If no tokens are available, the packet is discarded. The size of the token bucket determines the burst rate. As long as tokens are available, packets are transmitted at line rate. In an effort to keep the configuration intuitive, users configure the burst rate directly and the algorithm determines the appropriate bucket size. If the burst rate is set to 0, a default bucket size is used. An example of how to configure rate limiting of the Guest SSID, by overriding the rate limiting settings of the QoS profile assigned to the SSID, is shown in Figure 21-40.

One unique characteristic of wireless is that not all transmissions occur at a single rate. Signal strength and signal-to-noise ratio (SNR) will determine the actual speed of the physical medium for any single station. Unlike wired networks where the speed is fixed at the port rate, wireless rates can vary for each host on the subnet and may even change as the station moves closer or further from the access point. With wireless rate limiting, the time required to drain a full token bucket depends on the access speed of the wireless client and is not fixed. Stations that associated at 54 Mbps will be able to drain a token bucket faster than those at 1 Mbps. If per-SSID rate-limiting is in place, all clients on a particular AP share a single bucket. If per-User rate-limiting is in effect, then each station is assigned a unique bucket. It is possible to do both per-client and per-SSID rate limiting. In this case a token must be available and is removed from both the shared SSID bucket and per client bucket before the packet is transmitted. While this may provide more fairness to a slower user trying to access a shared token, it increases the amount of state information that must be maintained, increasing processing requirements on the controller. Because many deployments of guest wireless access simply provide best-effort service levels, extra processing requirements are not typically merited. As such, only per-SSID shaping is shown here. There may be other situations where a business case does justify doing both per-user and per-SSID rate limiting simultaneously.

Figure 21-40 Example Configuration for Rate Limiting the Guest SSID



One of the branch designs presented within this design guide uses FlexConnect with local branch termination for corporate wireless clients and central termination for guest traffic. Corporate-approved devices may send data to servers located within the central datacenter. Alternatively, they may send data to a local server. Where access to a local server is required, FlexConnect with local termination can save WAN bandwidth by eliminating the need to transfer data through a CAPWAP tunnel over the WAN to a central controller. Locally terminated traffic may still travel over the WAN when access to servers located within the data center is required, but these packets will not be tunneled within CAPWAP. In this case, normal QoS techniques can be applied. Hence, wireless packets are classified along with wired traffic. This common classification for corporate wired and wireless devices applies in both the upstream and downstream direction. With the design presented within this document, CAPWAP tunnels are used for all guest traffic-traffic from personal devices which have not on-boarded, as well as wireless control traffic (traffic from the wireless controller and the branch access points). Therefore, of all the CAPWAP traffic leaving the branch, the majority of packets will likely belong to guest users. This can help distinguish guest traffic from corporate traffic.

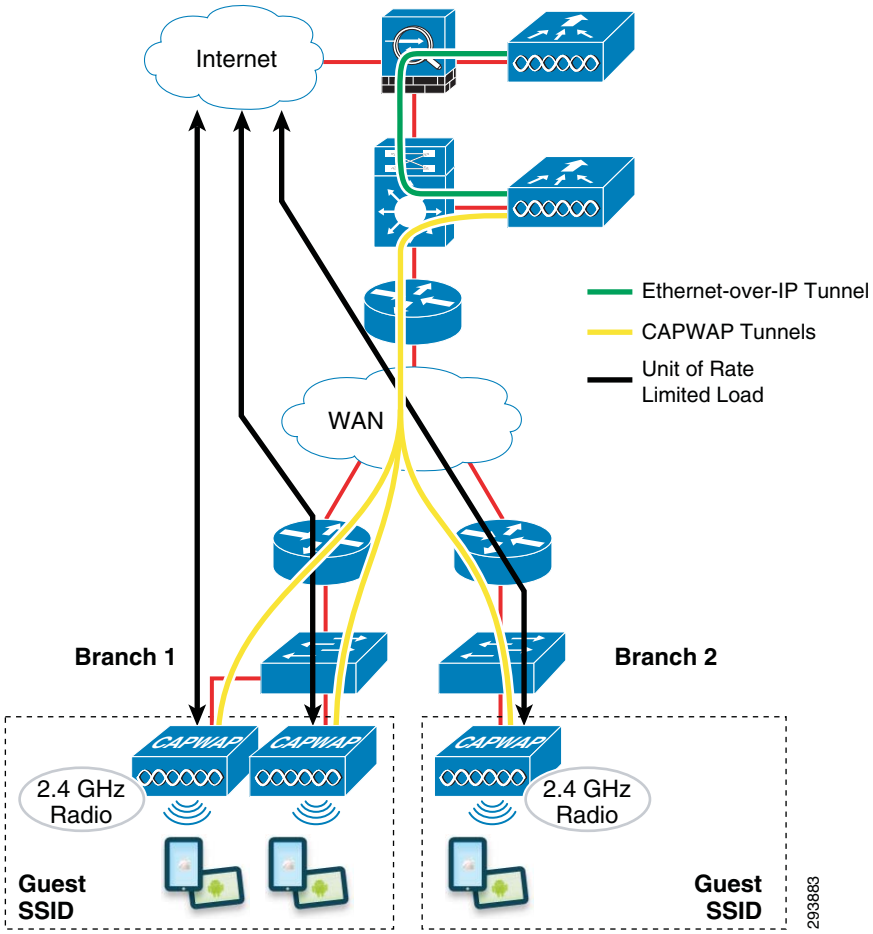
The example configuration shown in Figure 21-40 allows for two classification rates, the data rate and the real-time rate. Within the context of this configuration, Data is meant to be all TCP flows and Real-Time is meant to be all UDP flows. As a best practice for QoS, it is recommended to prevent UDP

and TCP from competing directly with each other for bandwidth due to differences in how dropped packets influence flow. Providing distinct token buckets for each protocol prevents any undesirable interaction between UDP and TCP.

Rate limiting is configured on the foreign controller. When guest access is offered at both the campus and branch locations, there will be two foreign controllers tunneling to an anchor controller. The rate limit configured on each foreign controller can be different and unique for that class of users. Typically the foreign controller servicing campus guests will have a higher bandwidth contract than the foreign controller servicing branch guest users because of the higher campus bandwidth available when compared to the WAN.

There are other caveats to be aware off when rate limiting. Because SSID rate limiting occurs at the radio itself, each radio will limit the SSID to the configured rate. This means that if Branch A and Branch B are each members of the BYOD_Guest SSID, each branch will limit guest traffic without regard to the current load in the Guest SSID at the neighboring branch. However, this means if the Guest SSID is present on two radios at the same branch and the rate is configured for 1 Mbps, the combined rate could be as high as 2 Mbps over the WAN at that single branch. Even within a single AP, if the Guest SSID is using the 2.4 GHz radio and the 5 GHz radio, the total bandwidth could be double the configured guest rate limit. As stated earlier, the rate limiting feature's primary purpose is to protect the radios. Because of this, rate limiting may necessitate the over subscription of WAN bandwidth intended for guest use. One possible method to minimize the extent of oversubscription, is for the Guest SSID not to be enabled on the 5 GHz radio. In addition, the number of APs participating in this SSID should be the minimum required to provide adequate coverage. AP groups can be used to manage which APs are participating. Rate limiting a single BYOD_Guest SSID across all branch locations may result in different WAN rates at different branches, as illustrated in Figure 21-41.

Figure 21-41 Rate-Limiting the Guest SSID



In Figure 21-41 assume the rate limiting of the BYOD_Guest SSID is configured for 1 Mbps. At Branch 1, the local WAN circuit could experience as much as 2 Mbps of guest traffic (due to the two APs), while the WAN aggregation circuit at the head-end could experience up to 3 Mbps of guest load (due to a total of three APs). If a single SSID is in use for guest traffic, then the configured rate should be appropriate for the slowest branch location that will be hosting guest traffic. There are some options available to better manage guest loads at the branch that are discussed below.

Multiple Guest SSIDs and AP Groups

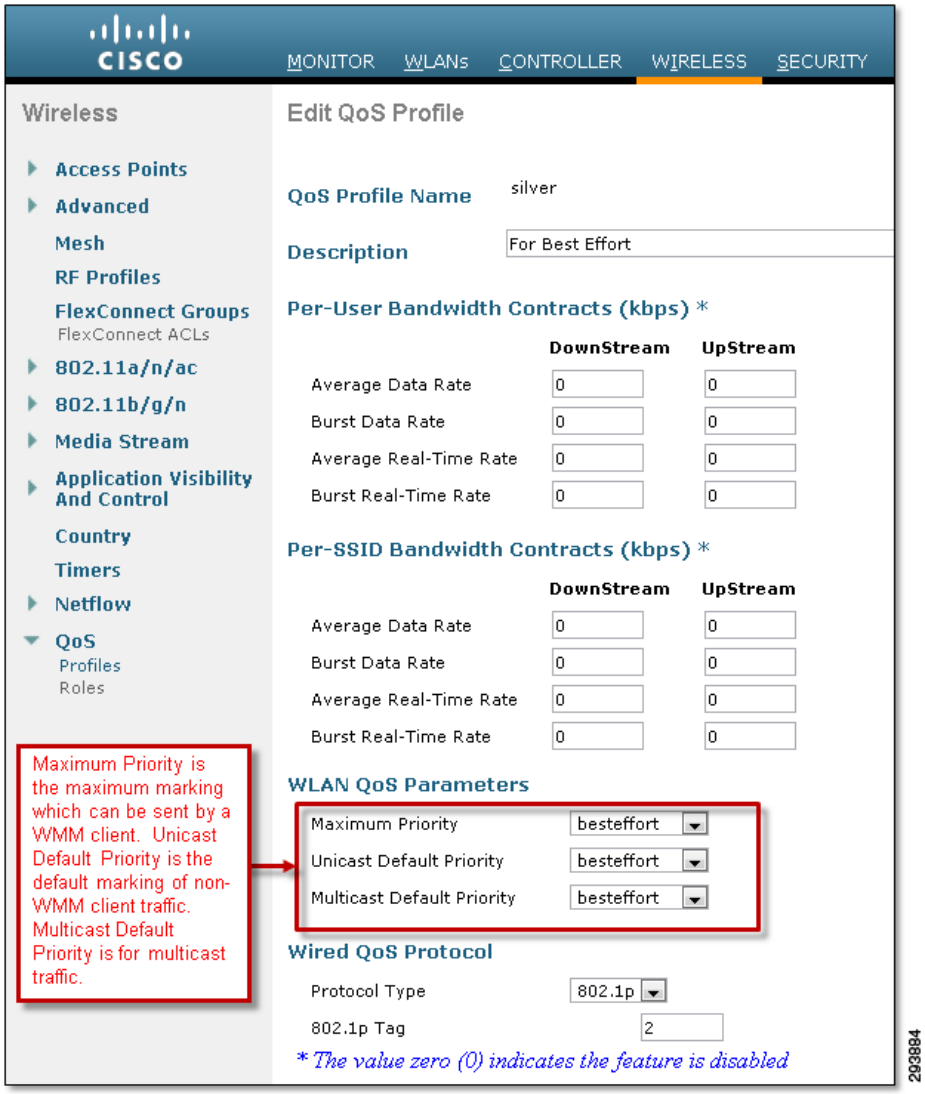
Because traffic limits are established per SSID and because not all branches have the same bandwidth available for guest use, the administrator may want to establish multiple Guest SSIDs based on the configured rate-limit. For example, the GUEST_128 SSID may be rate-limited to 128 Kb/s while the GUEST_256 SSID may be twice as fast. AP groups must be used to ensure both WLANs are not available at all branch locations. If the majority of branch locations have more than one AP that will host guest traffic, then the configured rate limit will be less than the actual desired rate to minimize oversubscription. AP groups can be used to manage how many radios are contributing to the total guest load for that location. Multiple Guest SSIDs in conjunction with AP groups can be used to ensure adequate guest coverage without excessive WAN loads. Creating informative names for the branch APs will simplify creating AP groups.

AP Groups are explained in greater detail in the Flex 7500 Wireless Branch Controller Deployment Guide at:
http://www.cisco.com/en/US/products/ps11635/products_tech_note09186a0080b7f141.shtml#ap-gr.

Managing the Downstream Load

With the FlexConnect design presented within this document, CAPWAP tunnels are used for all guest traffic, traffic from personal devices which have not on-boarded, as well as wireless control traffic (traffic from the wireless controller and the branch access points). Figure 21-40 shows an example where the Silver (best effort) QoS profile applied to the BYOD_Guest SSID. QoS profiles are used to set QoS markings for wireless data traffic encapsulated within the CAPWAP tunnel. Note that CAPWAP control traffic is prioritized separately from the settings within the QoS profile. Figure 21-42 shows an example of the default settings for the Silver (best effort) QoS profile.

Figure 21-42 Default Settings for the Silver (best effort) QoS Profile



The QoS profile can be used to set the following parameters:

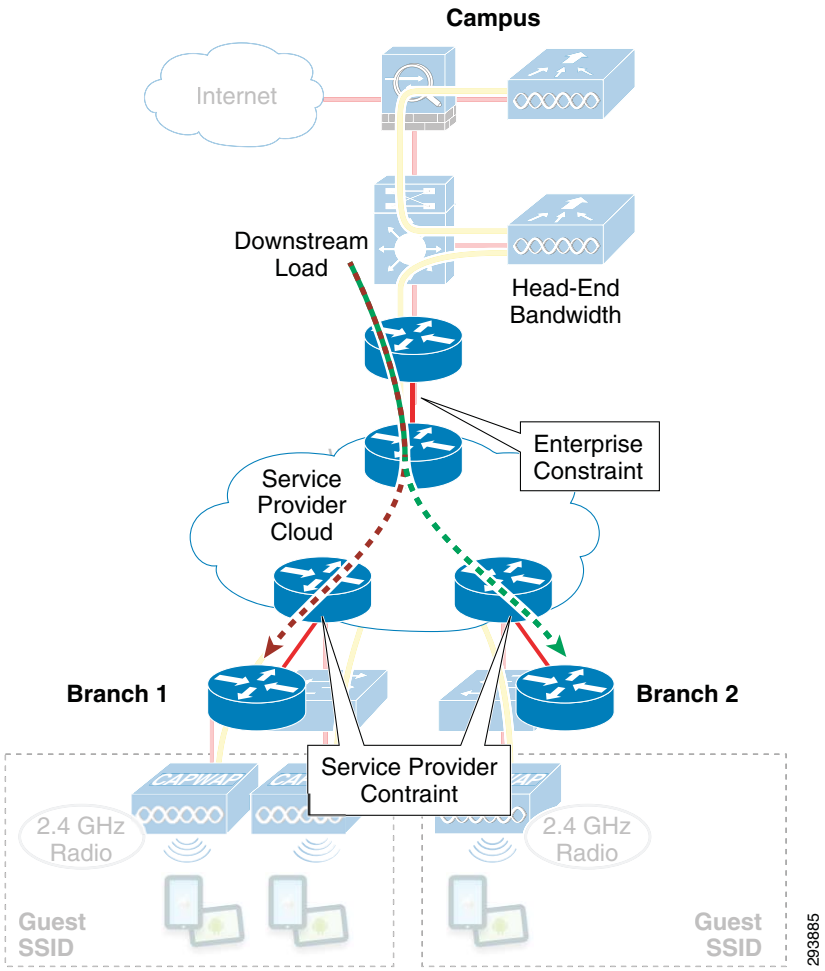
- Maximum Priority—This limits the maximum 802.11 User Priority marking which can be sent by a wireless client which supports WiFi Multimedia (WMM). The use of this parameter implies the SSID is configured to support WMM.
- Unicast Default Priority—This sets the default 802.11 User Priority marking for traffic sent from wireless client devices which do not support WMM.
- Multicast Default Priority—This sets the default 802.11 User Priority marking for multicast traffic.

The 802.11 User Priority value is then used to set the outer DSCP value of traffic encapsulated within the CAPWAP tunnel between the Access Point and the CUWN wireless controller. As can be seen above, the default User Priority is set for best effort, which maps to DSCP 0. Therefore in this example, of all the CAPWAP traffic traveling towards the branch, the majority of packets marked with DSCP 0 will likely belong to guest users. This can help distinguish guest traffic from corporate traffic.

Note The network administrator should note that the default User Priority settings of the Bronze QoS profile are set for Background. Therefore, if the network administrator wishes to set guest traffic to a lower User Priority of Background, which maps to DSCP 8 (corresponding to CS1 which is the Scavenger class), they can do so by assigning the guest SSID to the Bronze QoS profile. The network administrator should take into account the business requirements of the organization in order to determine whether guest traffic should be considered Best Effort or Background. Alternatively, the network administrator could change the default settings of the Silver QoS profile to background. Changing the default settings may not necessarily be the optimal solution, however, since there are only four QoS profiles which can be applied to all SSIDs configured within the CUWN wireless controller.

There are two points in the downstream path where loads imposed by the guest users could impact corporate traffic. These are the outbound interface on the WAN aggregation router and the outbound interface on the PE router adjacent to the branch. Figure 21-43 highlights the areas of concern in the downstream direction.

Figure 21-43 Downstream Congestion Points



The Per-SSID rate limiting discussed in the previous section does not provide direct control of the load on the WAN aggregation head ends imposed by guest users at the branch. The guest load will be proportional to the total number of branch APs hosting the Guest SSID times the per-SSID rate limit of the WLAN. Guest wireless traffic may be distinguishable from other WAN traffic because it will be in a CAPWAP tunnel and marked with the default DSCP setting. Some traffic from employees on-boarding personal devices will also be marked the same way if the same QoS Profile is applied to the dedicated provisioning SSID. However the percentage will be very small. It is possible to construct a policy that will mark CAPWAP packets with default DSCP values into the scavenger class. This will have the effect of setting guest traffic below the priority of default corporate traffic. When the bandwidth of the WAN aggregation circuit begins to saturate, this policy will allow wireless guest traffic to be discarded prior to corporate traffic. If on-boarding traffic is also dropped along with guest traffic, then employees will need to wait until the WAN loads are lowered prior to bringing a new device onto the network. This is implemented with traditional QoS policy maps on the outbound circuits of the WAN Aggregation router. Incidentally the same approach could be used on the branch uplink to manage situations where the number of APs at the branch could unreasonably oversubscribe the uplinks.

The service provider local links to the branch may also come under load as a result of the guest traffic. The per-SSID rate-limiting does benefit the branch WAN links in this direction by limiting the effective guest bandwidth as a result of application-based flow control. An example is TCP-based applications, which will manage their flow to minimize drops. Even though per-SSID rate-limiting in the downstream direction is applied at the radio towards the end station, the client application will throttle down to meet

the rate available over the entire path. The last hop interface on the SP PE router also contributes to application throttling if aggressive policers are used to enforce contracted rates. Assuming wireless guests are remarked scavenger and appropriate DSCP to EXP mappings are being used, then SP policers should disproportionately impact wireless guest TCP applications. Although guest Internet traffic rarely uses UDP, it also generally exhibits the same type of flow control behavior as TCP even though the protocol itself does not implement feedback as part of the transport layer. This is because UDP is often transactional based. When UDP is used for bulk transfer, blocks of data are numbered and acknowledged by the application, for example TFTP. A transmitter will not send a block of data until the receiver has acknowledged the previous block. If a block of data is dropped, the transmitter will wait for a timeout period before retransmitting the previous block. Two exceptions to UDP application based flow control are UDP-based IP video surveillance which may not use RTSP to monitor received data and UDP multicast. Neither of these are typical applications guest will use on the Internet. In any case, per-SSID rate limiting is an effective means to manage guest traffic on the SP's PE routers.



CHAPTER 22

Managing a Lost or Stolen Device

Revised: August 7, 2013

When a previously provisioned device is reported lost or stolen, the device must be denied access to prevent unauthorized access to the network.

A first level of defense to protect against a lost or stolen device is to enforce the use of a PIN lock, a passcode required to access the device that may lock the device automatically after a short period of inactivity (typically five to ten minutes). This can also be enhanced by erasing all data on the mobile device after a number of failed passcode attempts or performing a selective wipe. This and other rules may be enforced by the integration of Cisco ISE with a Mobile Device Manager.

The Cisco ISE offers different ways to prevent a lost or stolen device from connecting to the network. The My Devices Portal allows the employee to mark a device as lost and prevent others from gaining unauthorized access with that device. In addition, if the device is connected to the network when the device is marked as lost, the ISE may issue a Change of Authorization (CoA) to force the endpoint off the network.

The administrator is also able to blacklist a device and force the endpoint off the network. In addition, the administrator is able to use Endpoint Protection Services (EPS) to quarantine an endpoint from the network.

Employees and administrators have different capabilities to block lost or stolen devices:

Employees:

From the My Devices Portal:

- Report devices as lost.
- Enforce a PIN lock through the MDM.
- Initiate a remote device wipe through the MDM.
- Reinstate a device to regain access without having to register the device again.

Note Devices that have been fully wiped cannot be restored by ISE and will need to re-register to recover the certificate and WiFi profile.

Administrators:

- Add the endpoint to the Blacklist Identity Group.
- If the endpoint is connected, force it off the network using the Show Live Sessions screen.
- Enforce a PIN lock through the Endpoints screen in ISE.

- Initiate a remote device wipe through the Endpoints screen in ISE.
- Quarantine the endpoint using the ISE's Endpoint Protection Services feature (employees are not able to reinstate endpoints quarantined by the administrator).
- Revoke the device's digital certificate.
- Disable the RSA SecurID token..

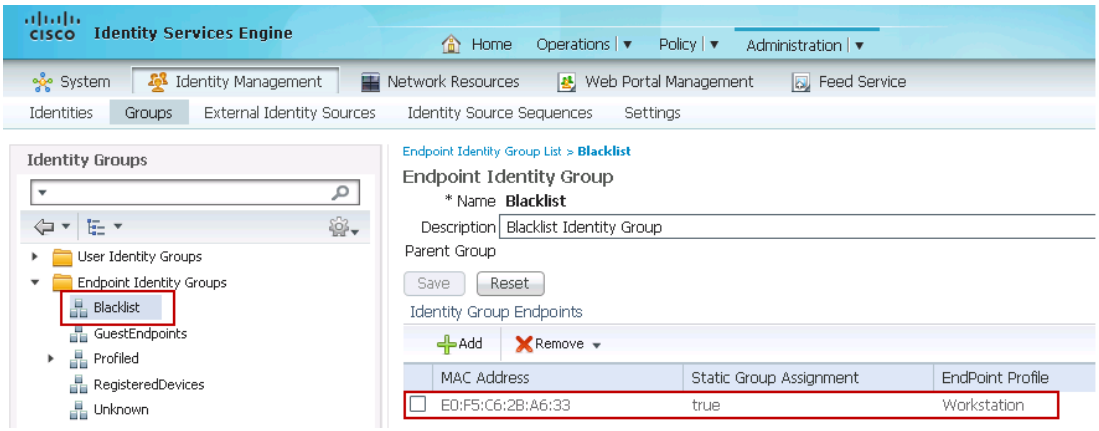
Blacklist Identity Group

The Blacklist identity group is system generated and maintained by ISE to prevent access to lost or stolen devices. In this design guide, two authorization profiles are used to enforce the permissions for wireless and wired devices within the Blacklist:

- Blackhole WiFi Access
- Blackhole Wired Access

The Blacklist Identity Group is displayed by clicking **Administration > Identity Management > Groups > Endpoint Identity Groups**. Figure 22-1 shows an empty Blacklist identity group containing one endpoint.

Figure 22-1 Blacklist Identity Group

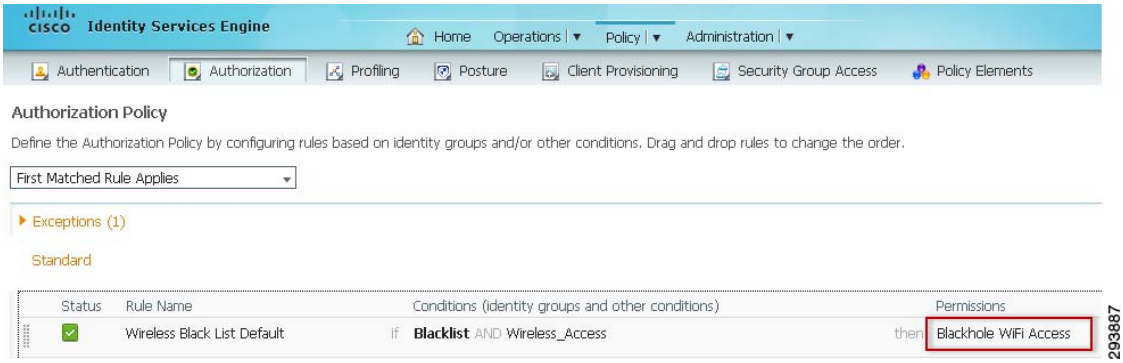


Devices that have been blacklisted are assigned to the Blacklist identity group. Both wired and wireless devices can be placed into the Blacklist identity group. An authorization profile is used to define the access granted to the blacklisted devices. Blacklisted device connection requests are redirected to a web page that informs the user that the device is blacklisted.

Blacklisting Wireless Devices

To enforce the blacklist permissions, an authorization rule is defined under **Policy > Authorization**. Figure 22-2 shows the Wireless Black List Default rule enforcing the Blackhole WiFi Access permissions.

Figure 22-2 Wireless Black List Default Authorization Rule

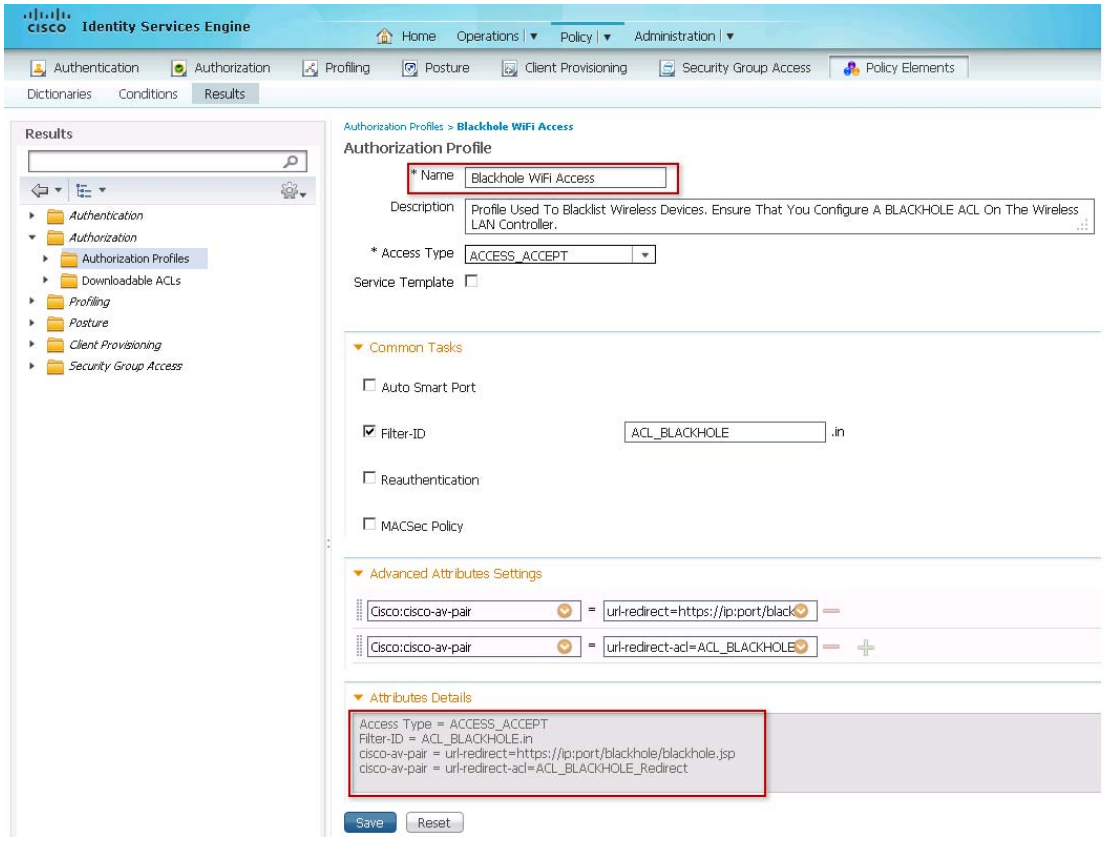


The Blackhole WiFi Access authorization profile is configured under **Policy > Policy > Elements > Results > Authorization Profiles**, as shown in Figure 22-3. The Access Type is defined as ACCESS_ACCEPT and the following cisco-av-pairs are defined:

- cisco-av-pair: url-redirect=https://ip:port/blackhole/blackhole.jsp. The user gets redirected to this page when a device is in the Blacklist identity group.
- cisco-av-pair: url-redirect-acl=ACL_BLACKHOLE_Redirect. The Wireless LAN Controller must have an ACL named ACL_BLACKHOLE_Redirect configured for the redirection to work at the campus and a FlexConnect ACL at the branch with the same name.

This authorization profile only allows access to the ISE “Unauthorized Network Access” page to inform the user that access to the network has been denied for that device.

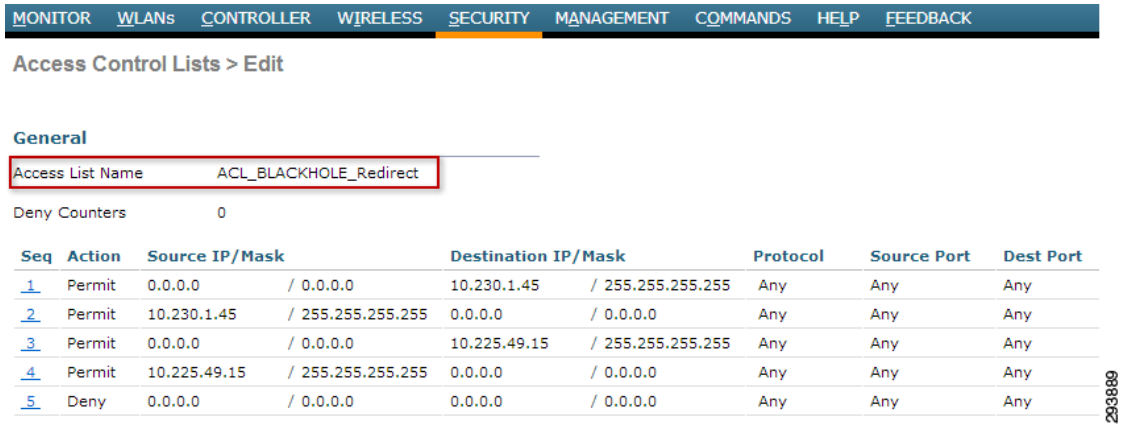
Figure 22-3 Blackhole WiFi Access Authorization Profile



The behavior of the two ACLs in the authorization profile is slightly different between CUWN wireless controllers, such as the CT5508 and Flex 7500, and IOS XE based controllers such as the CT5760 and Catalyst 3850. For CUWN wireless controllers, ACL_BLACKHOLE_Redirect functions as both the ACL which controls web redirection, as well as the ACL which controls what the wireless client is allowed to access on the network.

Figure 22-4 shows how the ACL_BLACKHOLE_Redirect access list is defined in a CUWN WLC to only allow access to the ISE and DNS server. By granting access to DNS and ISE, the endpoint is able to reach the blackhole.jsp web page hosted at the ISE.

Figure 22-4 ACL_BLACKHOLE_Redirect ACL



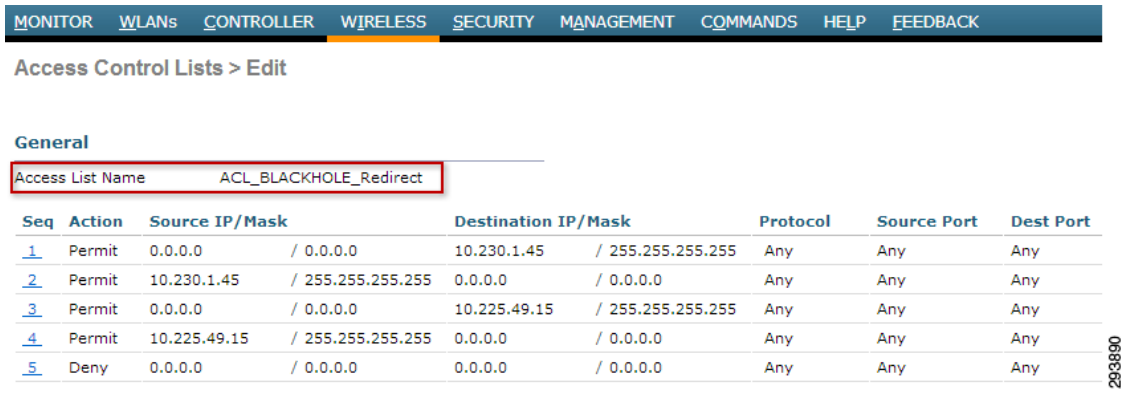
The ACL specifies the following access:

- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE Server (10.225.49.15).
- Deny access to and from all other addresses.

Note ACL_BLACKHOLE serves simply as an extra security configuration. CUWN wireless controllers do not make use of this ACL when URL redirection is specified. For CUWN wireless controllers the ACL_BLACKHOLE ACL can be the same as the ACL_BLACKHOLE_Redirect ACL.

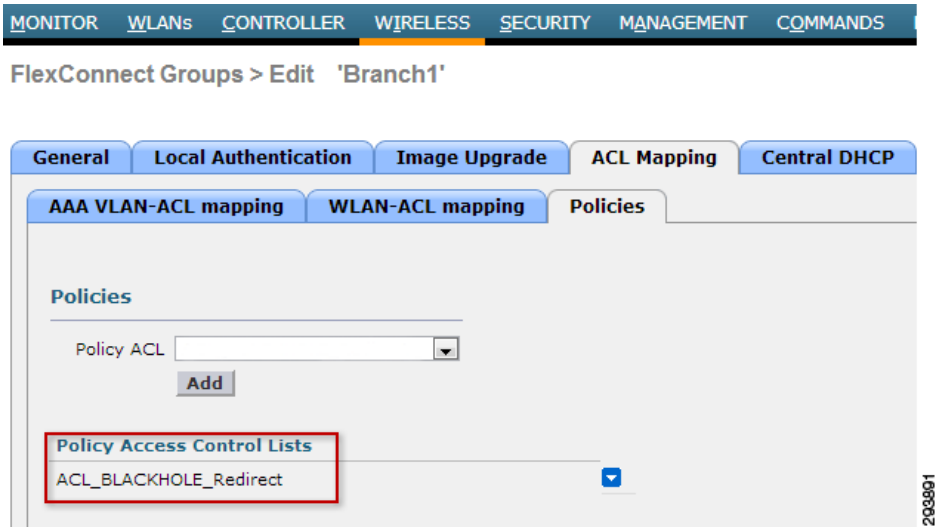
For endpoints connecting to the branch, a similar FlexConnect ACL is defined and applied to the FlexConnect Group. Figure 22-5 shows the ACL_BLACKHOLE_Redirect FlexConnect ACL. This ACL is similar to the one used for campus devices, shown above, but defined under **Security > Access Control Lists > FlexConnect ACLs**.

Figure 22-5 ACL_BLACKHOLE_Redirect FlexConnect ACL



To apply this FlexConnect from the branch, select the appropriate FlexConnect Group and click the **Policies** tab. Add the ACL_BLACKHOLE_Redirect ACL, as shown in Figure 22-6.

Figure 22-6 Policies for Branch1



On converged access products, namely the CT5760 wireless controller or Catalyst 3850 Series switches, both the BLACKHOLE_ACL_Redirect and BLACKHOLE_ACL ACLs must be configured. An example of the BLACKHOLE_ACL_Redirect ACL is shown below.

```
!
ip access-list extended ACL_BLACKHOLE_Redirect/ Blacklisting Redirection ACL
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 permit ip any any
!
```

The above ACL specifies the following access:

- Deny DHCP access (bootpc and bootps).
- Deny IP access to and from the DNS server (10.230.1.45).
- Deny IP access to and from the ISE server (10.225.49.15).
- Allow (redirect) all other IP access.

The ACL above causes any web traffic (HTTP or HTTPS) from any source to any destination to be redirected to the blacklisted devices web page within the Cisco ISE.

The authorization profile also applies a second, RADIUS specified local ACL (BLACKHOLE_ACL) across the WLAN for network access. The CT5760 and Catalyst 3850 Design use named ACLs. The name of the ACL is sent from the ISE to the Catalyst 3850 Series switch or the CT5760 wireless controller via the RADIUS Airespace-ACL-Name attribute-value pair within the Airespace dictionary. The specific form for the example is as follows:

```
Airespace-ACL-Name = ACL_BLACKHOLE
```

The WLAN access-control ACL (ACL_BLACKHOLE) determines what traffic is allowed on the WLAN by the Catalyst 3850 series switch or the CT5760 wireless LAN controller. An example of the BLACKHOLE_ACL ACL is shown below.

```
!
ip access-list extended ACL_BLACKHOLE/ Blacklisting Access Control ACL
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
```

!
The above access-list specifies the following access:

- Allow DHCP access (bootpc and bootps).
- Allow IP access to and from the DNS server (10.230.1.45).
- Allow IP access to and from the ISE server (10.225.49.15).
- Implicitly deny all other IP access.

The ACL above allows traffic from any source to the blacklisted devices webpage within the Cisco ISE.

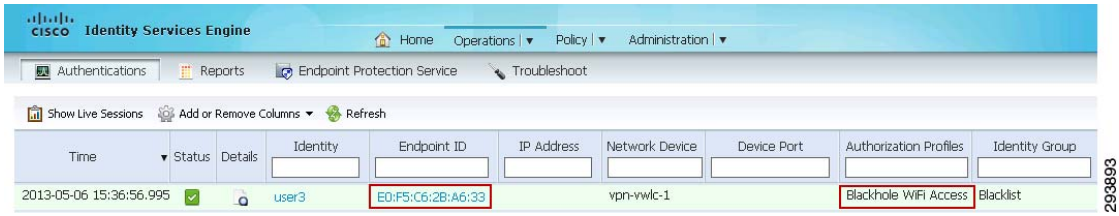
Once a device is in the Blacklist identity group, future attempts to connect to the network are denied. When a user opens a web browser on a blacklisted device, the session is redirected to the page shown in Figure 22-7.

Figure 22-7 Unauthorized Network Access



Figure 22-8 shows how a device in the Blacklist attempts to connect to the network and how the Blackhole WiFi Access authorization profile is applied.

Figure 22-8 Device in Blacklist Identity Group



Blacklisting Wired Devices

The user experience when a wired device is blacklisted is similar to a wireless device that has been blacklisted. The ISE authorization policy rule for blacklisting of on-boarded wired devices will be the same for devices connected via Converged Access products and other Catalyst switches for BYOD.

When a device is blacklisted and the user attempts to access any web page, the device is re-directed to the portal that lets the user know that the device has been identified as lost. The following steps show how to implement this behavior:

Step 1 Create an ACL_BLACKHOLE downloadable DACL that only allows access to the ISE.

Step 2 Create a URL Redirect ACL called ACL_BLACKHOLE_Redirect on the access layer switch that matches any HTTP or HTTPS traffic.

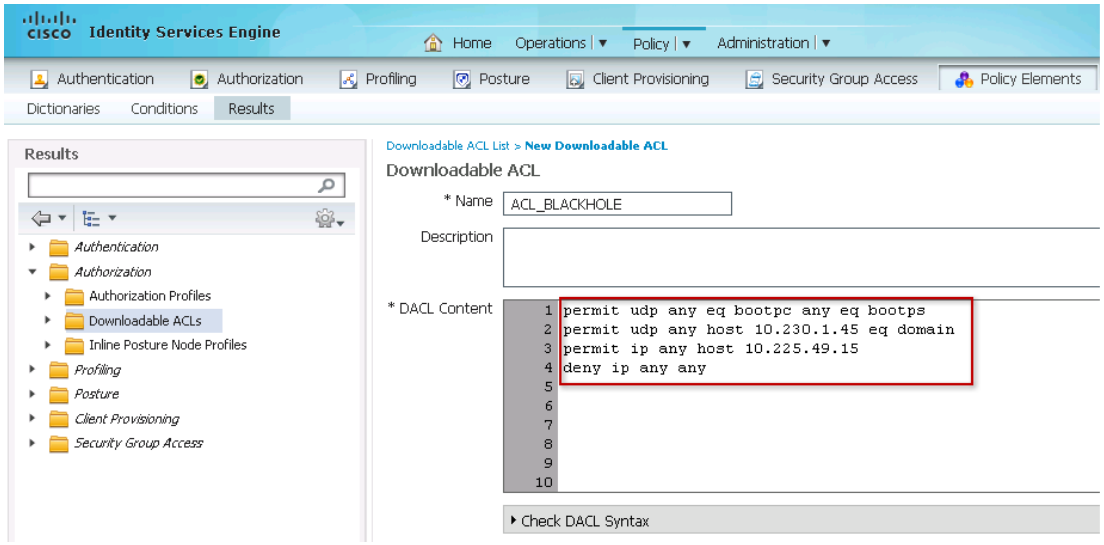
Step 3 Create a Blackhole Wired Access authorization profile that pushes the DACL and redirect-link to the switch.

Step 4 Define a new rule in the Authorization policy that matches on the blacklisted devices and assigns the authorization profile Blackhole Wired Access.

Creating a Downloadable ACL on ISE

The ACL_BLACKHOLE DACL is created under **Policy > Policy Elements > Results > Downloadable ACLs**, as shown in Figure 22-9.

Figure 22-9 ACL_BLACKHOLE DACL



The ACL above allows traffic from any source to the blacklisted devices webpage within the Cisco ISE.

Creating the URL REDIRECT ACL on the Switch

The configuration of the URL redirect ACL is the same, regardless of whether the switch is a Converged Access Catalyst 3850 Series switch, or a Catalyst 3750-X series switch. An example of the configuration for the ACL_BLACKHOLE_Redirect ACL on the wired switch is shown below.

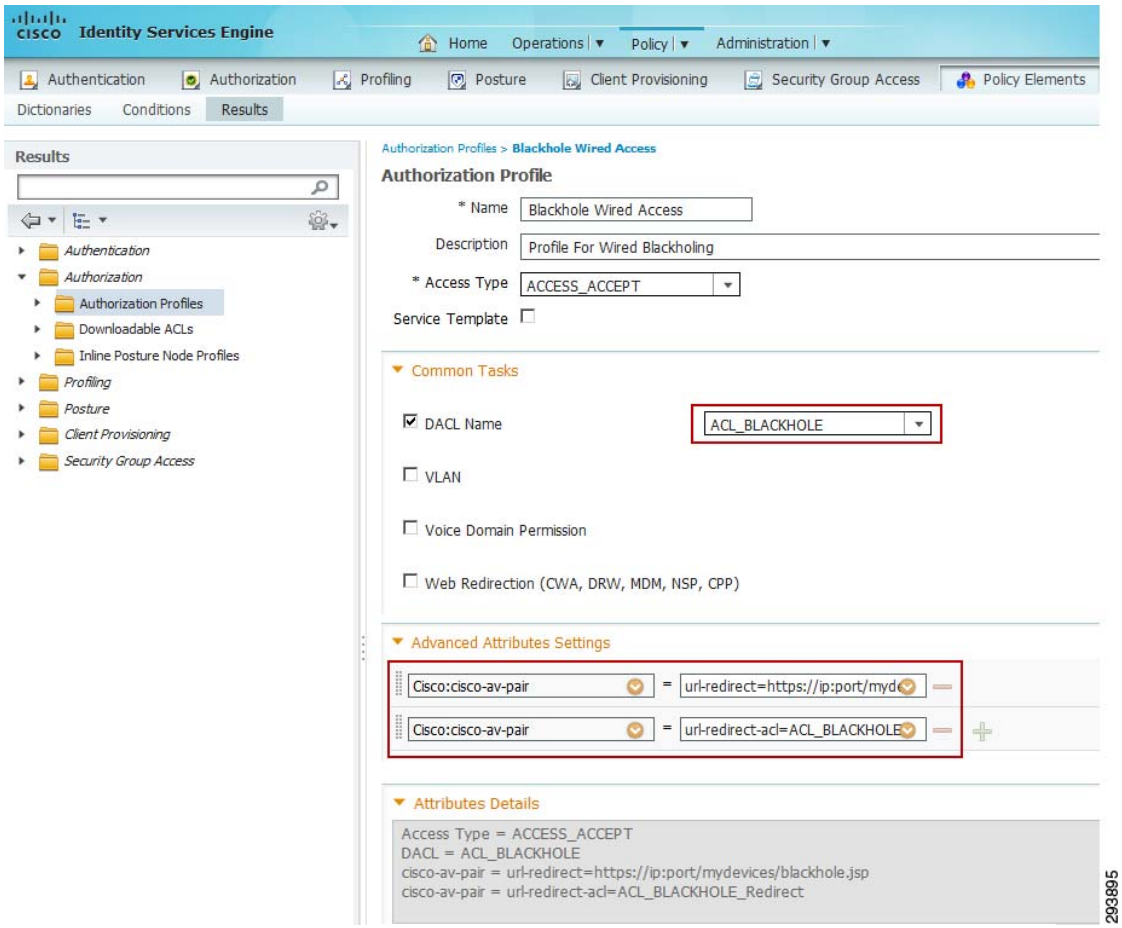
```
!
ip access-list extended ACL_BLACKHOLE_Redirect  / Blacklisting Redirection ACL
 udp any eq bootpc any eq bootps
 udp any host 10.230.1.45 eq domain
 ip any host 10.225.49.15
 permit ip any any
```

Note that this is the same URL redirect ACL used for blacklisted wireless devices for the Catalyst 3850 Series switch, discussed above. Bear in mind that the ACL is only required on the Catalyst 3850 and not on the CT5760 wireless controller because only the Catalyst 3850 supports direct connection of wired clients through its switch ports.

Configuring the Authorization Profile

Define an authorization profile called Blackhole Wired Access under **Policy > Policy Elements > Results > Authorization Profiles**, as shown in Figure 22-10.

Figure 22-10 Blackhole Wired Access Authorization Profile



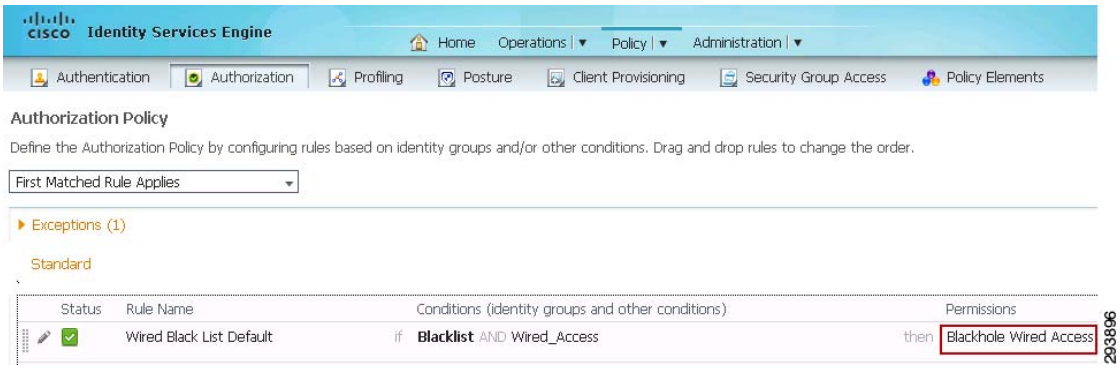
The following cisco-av-pairs are defined:

- cisco-av-pair: url-redirect=https://ip:port/mydevices/blackhole.jsp. The user gets redirected to this page when a device is in the Blacklist identity group.
- cisco-av-pair: url-redirect-acl=ACL_BLACKHOLE_Redirect. The access layer switch must have an ACL named ACL_BLACKHOLE_Redirect configured for the redirection to work.

Creating a Rule in the Authorization Policy

The last configuration step is to create a new rule in the authorization policy that uses the Blackhole Wired Access authorization profile created above when it matches a dot1x wired device which is blacklisted. Figure 22-11 displays the rule.

Figure 22-11 Wired Black List Default



Once the rules are defined, a blacklisted wired device will be denied access to the network.

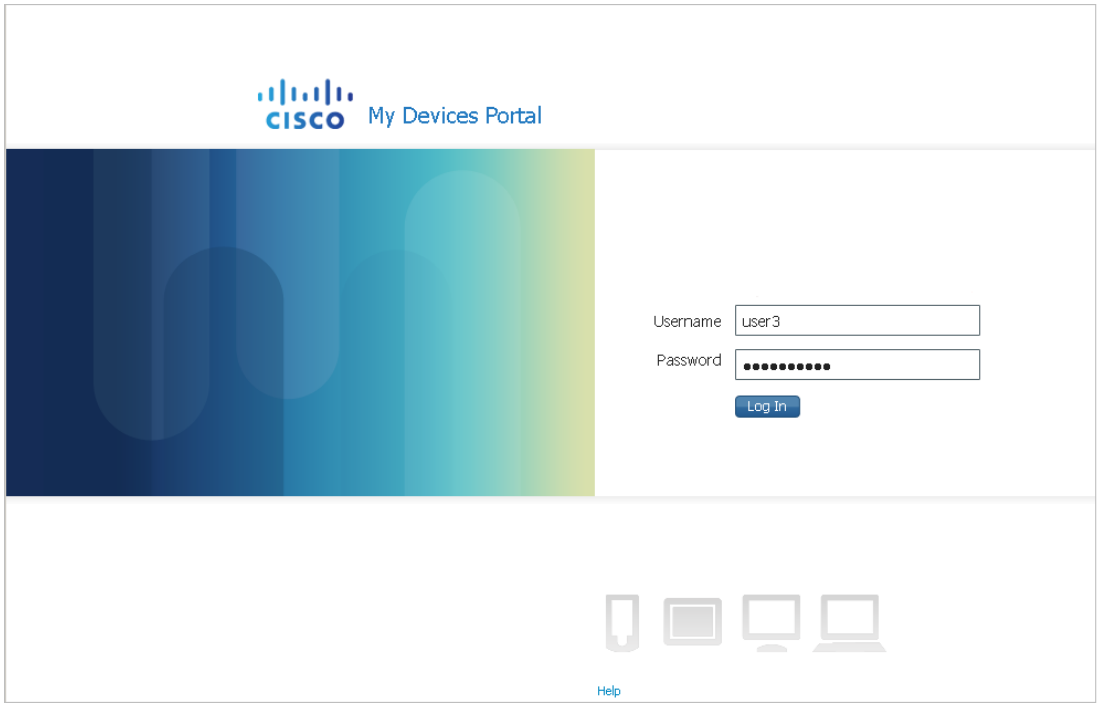
Employees' My Devices Portal

Using the My Devices Portal, employees are able to mark any of their devices as Lost to prevent further network access. The portal requires user authentication and displays the devices that have been added or registered by the employee. The My Devices Portal may be accessed from the following URL:

https://<*ISE_IP_Address*>:8443/mydevices/

In addition to being able to report a device lost or stolen, the portal is used to enforce MDM Actions, such as PIN lock and device wipes through MDM integrations. Figure 22-12 shows the My Devices Portal page.

Figure 22-12 My Devices Portal

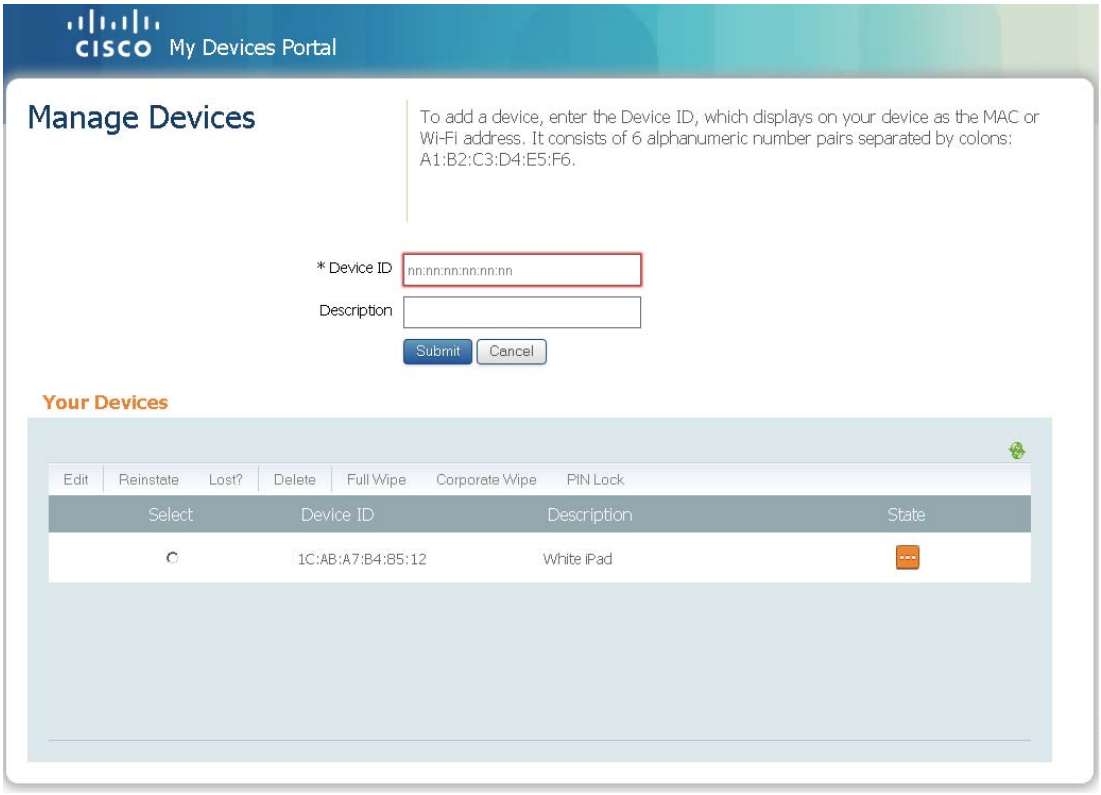


The portal displays devices assigned to the employee or previously registered using the self-registration portal.

Reporting a Device as Lost

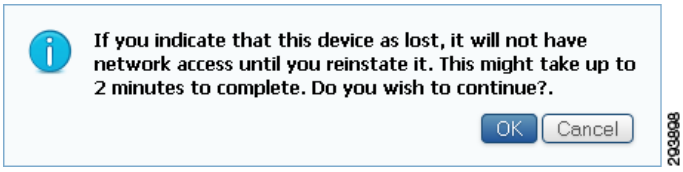
The employee is able to edit the device's description and report a device as lost, as shown in Figure 22-13.

Figure 22-13 Lost Device



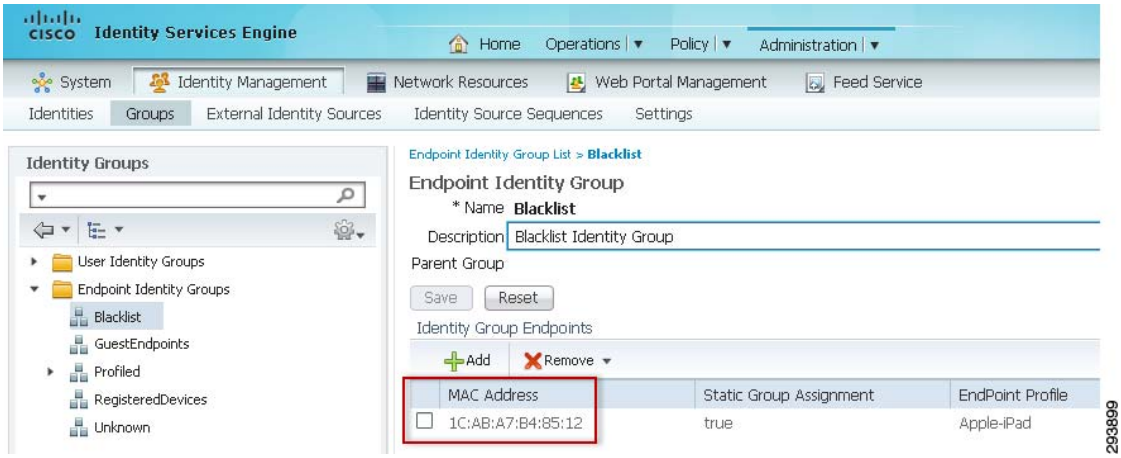
Before blacklisting a device, ISE displays the warning shown in Figure 22-14.

Figure 22-14 Blacklist Warning



Once the device is marked as lost, the device is added to the Blacklist Identity Group. If the device is connected to the network at that time, ISE issues a Change of Authorization (CoA) and forces the device off the network. To verify that the device has been added to the Blacklist Identity group, click **Administration > Identity Management > Groups > Endpoint Identity Groups** and review the Blacklist. Figure 22-15 shows the MAC address added to the Blacklist.

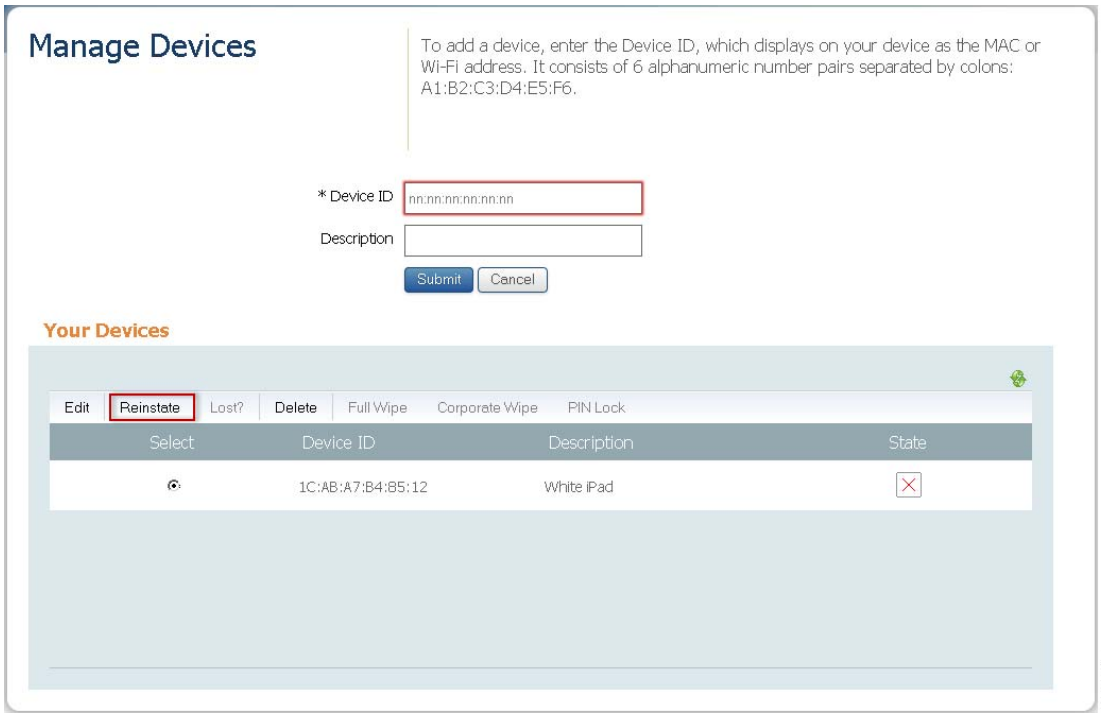
Figure 22-15 Blacklist Identity Group



Reinstating a Device

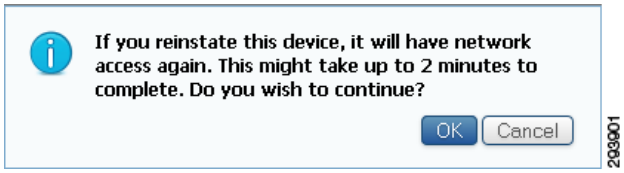
The My Devices Portal also gives the user the option to reinstate a blacklisted device, allowing the device to regain access to the network. Figure 22-16 shows the Reinstate option.

Figure 22-16 Reinstate a Device



Before reinstating a device, ISE displays the warning shown in Figure 22-17.

Figure 22-17 Reinstate Warning



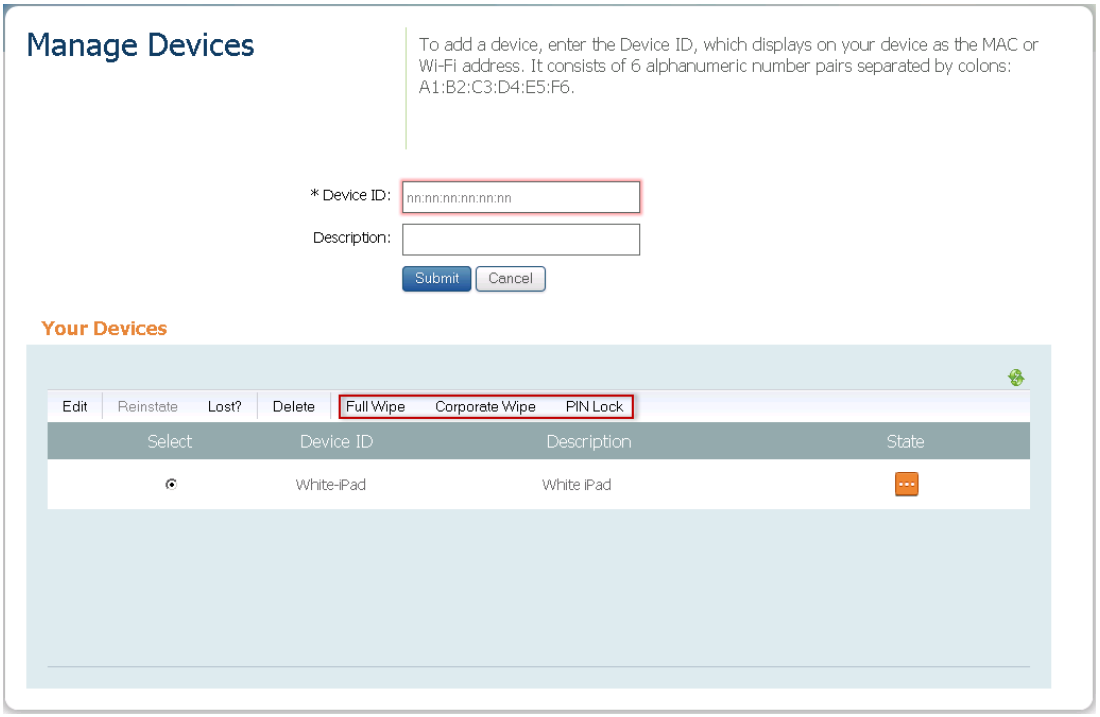
PIN Lock and Device Wipes

The integration with an MDM provides additional device control to the user. With the My Devices Portal, the employee is able to:

- Initiate a Corporate Wipe—Remove settings and applications configured in the MDM policies.
- Initiate a Full Wipe—Remove all information from the device (factory reset).
- Enforce a PIN lock—Lock the device.

Figure 22-18 shows the MDM actions available from the My Devices Portal:

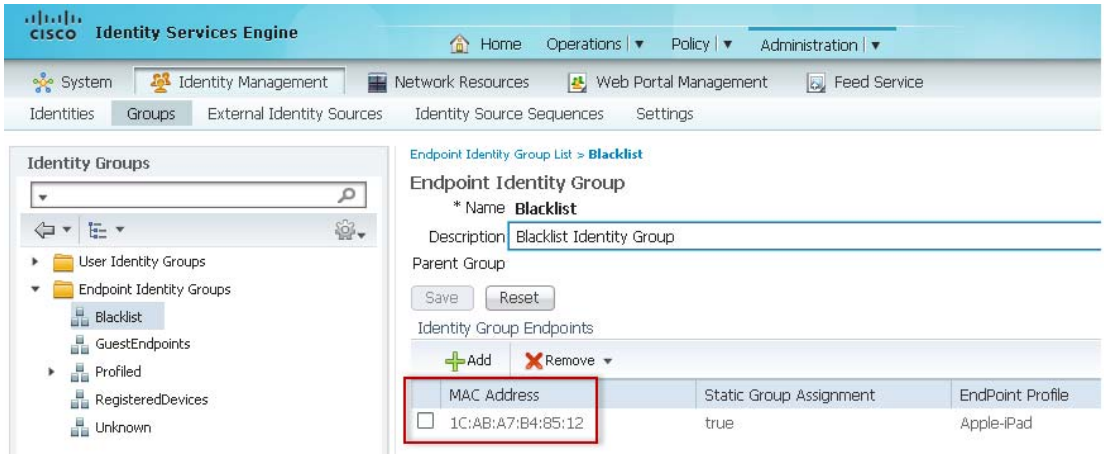
Figure 22-18 MDM Actions from My Devices Portal



Administrators—Blacklisting a Device

Administrators are able to blacklist a device by adding it manually to the BlackList Identity Group. Click **Administration > Groups > Endpoint Identity Groups > Blacklist** and add the MAC address of the device to be blacklisted. Figure 22-19 shows the MAC address added to the identity group.

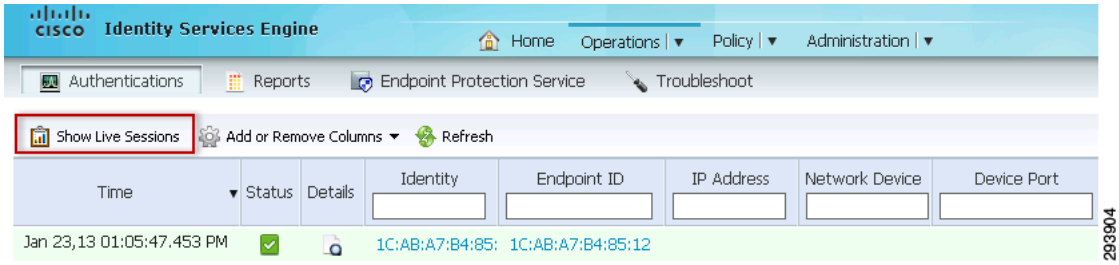
Figure 22-19 Device to be Blacklisted



Note that by adding the device to the Blacklist Identity Group, ISE prevents future attempts to connect to the network, but if the user is currently connected to the network, an additional step is required to force the endpoint off the network.

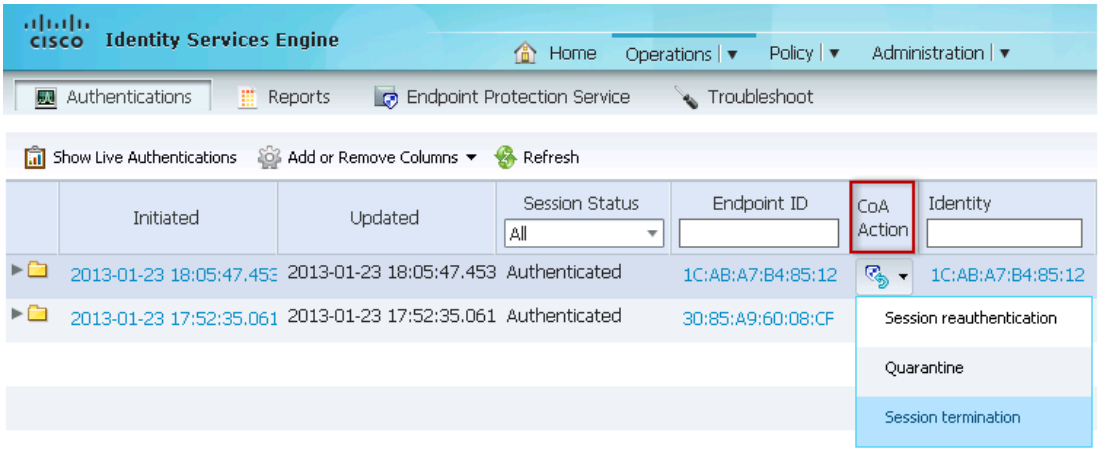
To force a device off the network, click **Operations > Authentications > Show Live Sessions**, as shown in Figure 22-20.

Figure 22-20 Show Live Sessions



To terminate the endpoint's session, select Session termination from the CoA Action menu, as shown in Figure 22-21.

Figure 22-21 Session Termination

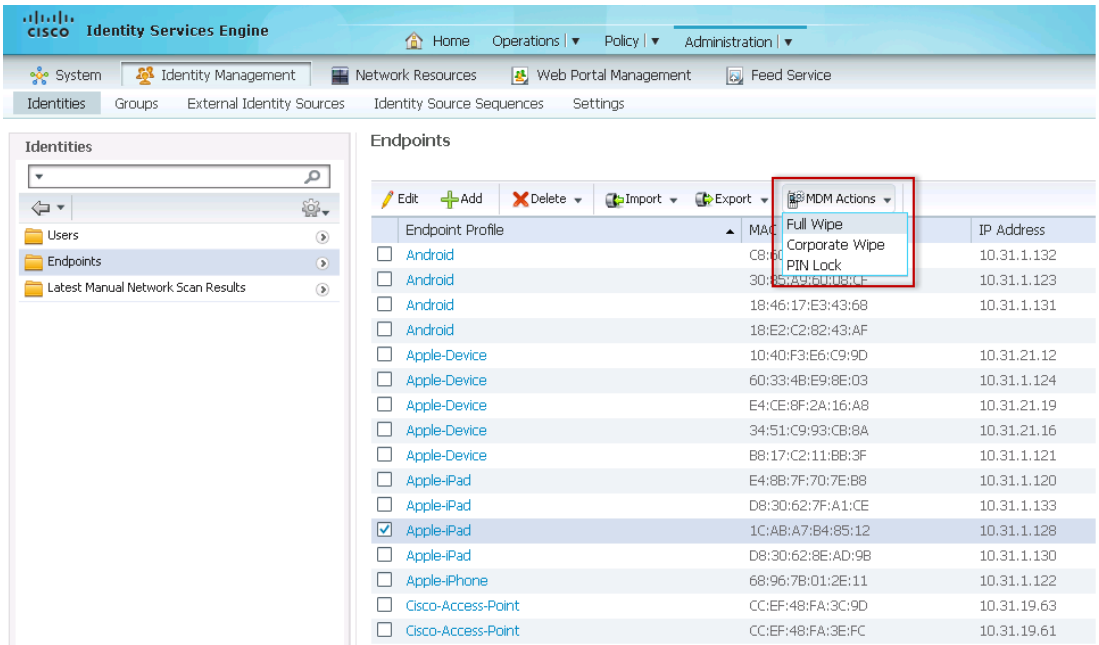


Note The employee has the option to reinstate a device that was blacklisted by the administrator using the My Devices Portal.

MDM Actions

The administrator also has the option to initiate MDM actions on the endpoints. To perform an action, click on **Administration > Identities > Endpoints** and select the appropriate device, as shown in Figure 22-22.

Figure 22-22 MDM Actions



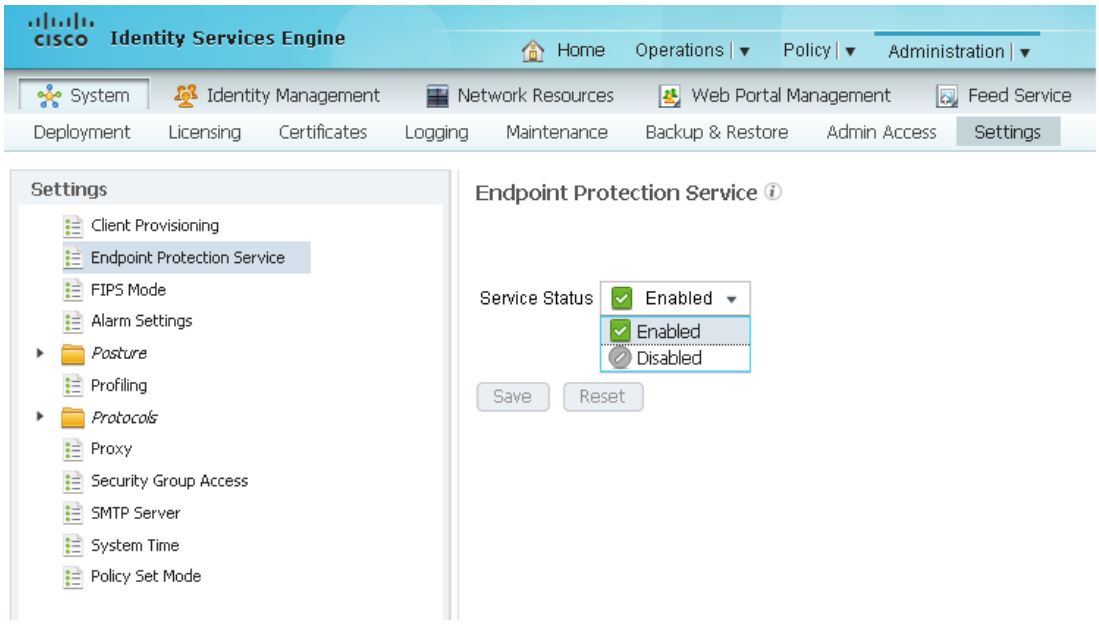
Endpoint Protection Services (EPS)

Endpoint Protection Services is a service provided by the ISE to extend the monitoring and controlling capabilities of endpoints. EPS also monitors and changes the authorization state of endpoints. EPS can be used to change the authorization state of an endpoint without having to modify the overall authorization policy. EPS allows the administrator to quarantine or limit access to a device and unquarantine a device or allow full access to the network to reverse the quarantine status.

Note EPS requires an ISE Advanced license.

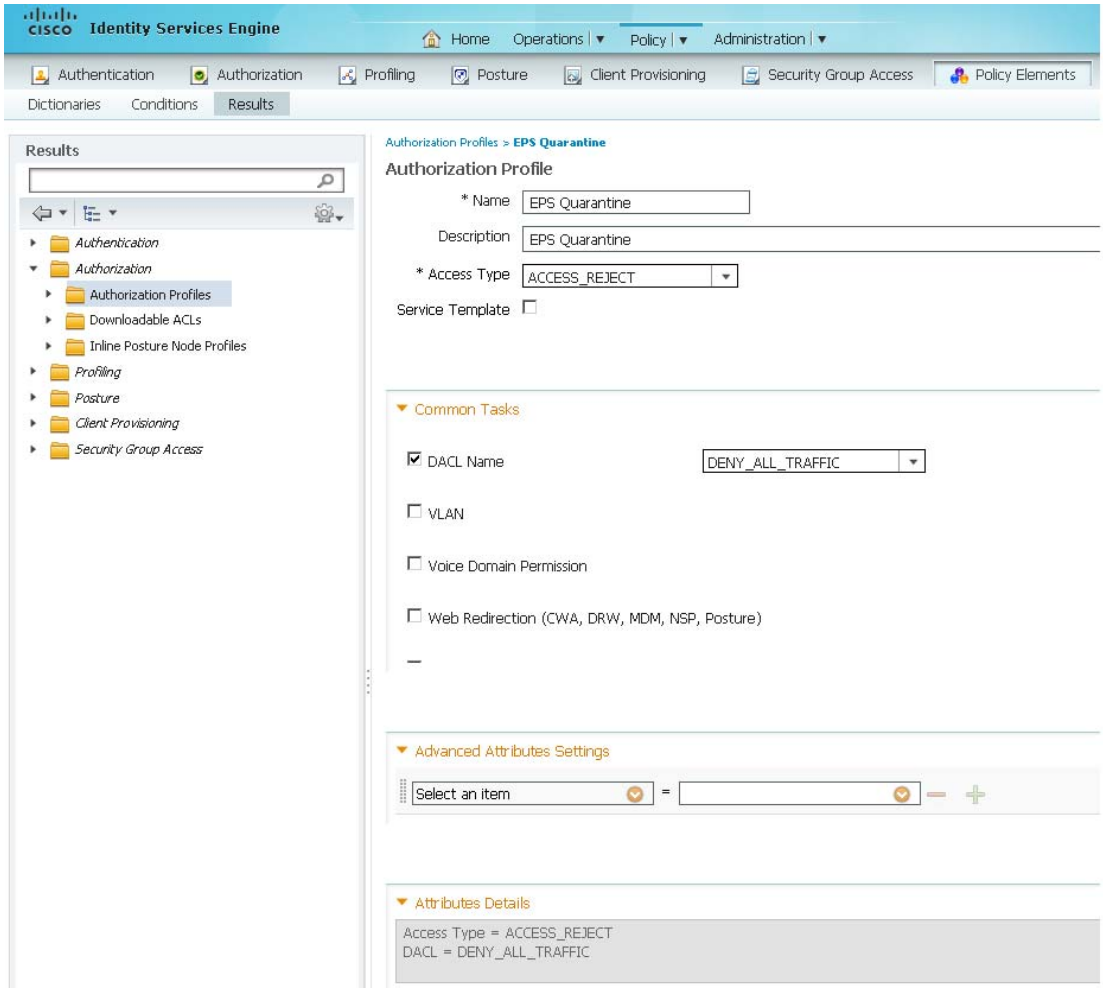
To enable EPS, click **Administration > System > Settings > Endpoint Protection Services** and select **Enabled**, as shown in Figure 22-23.

Figure 22-23 Enable EPS



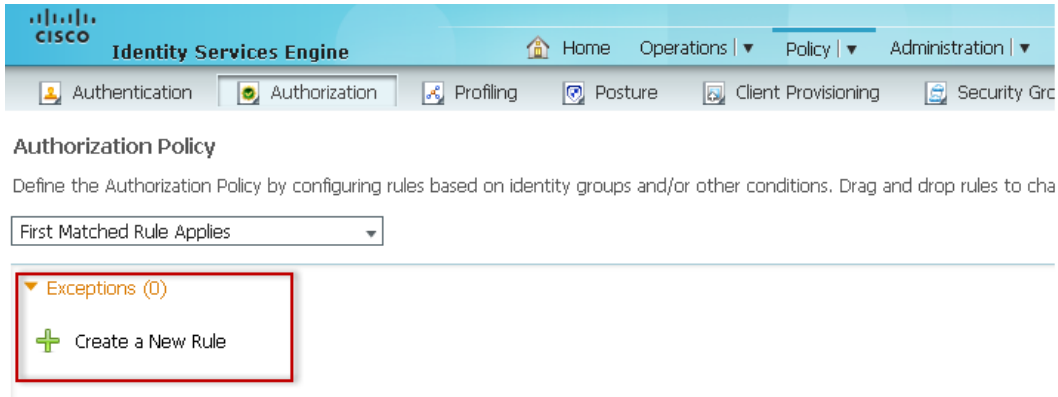
Create an authorization profile to define the permissions to specified network services. Click **Policy > Policy Elements > Results > Authorization > Authorization Profiles** and define a new authorization profile, as shown in Figure 22-24.

Figure 22-24 EPS_Quarantine Authorization Profile



Create an EPS Exception policy and rule to be processed before the standard policies are processed. Click **Policy > Authorization > Exceptions > Create a New Rule**, as shown in Figure 22-25.

Figure 22-25 EPS Exception Policy



Enter a Rule Name and under Conditions create a new condition (**Advanced Option**). Under Expression click **Select Attribute** and select **EPSStatus Equals Quarantine**, as shown in Figure 22-26.

Figure 22-26 EPS Exception Policy



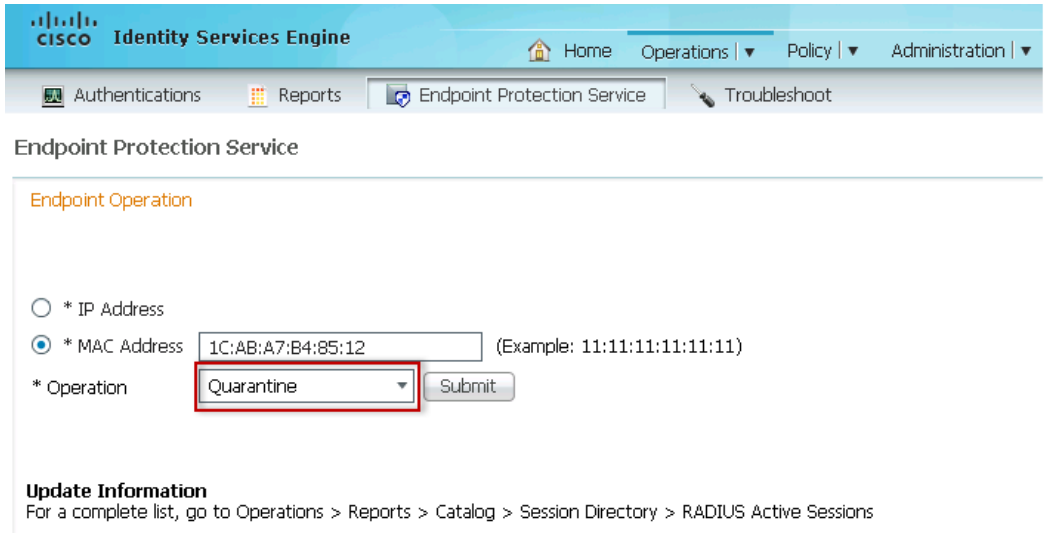
Under Permissions, select the previously defined **EPS_Quarantine** Authorization Profile. Figure 22-27 shows the complete Exception policy.

Figure 22-27 EPS Quarantine Permissions



To quarantine a device, click **Operations > Endpoint Protection Service** and enter the endpoint's MAC Address to be quarantined. Under Operation select **Quarantine**, as shown in Figure 22-28.

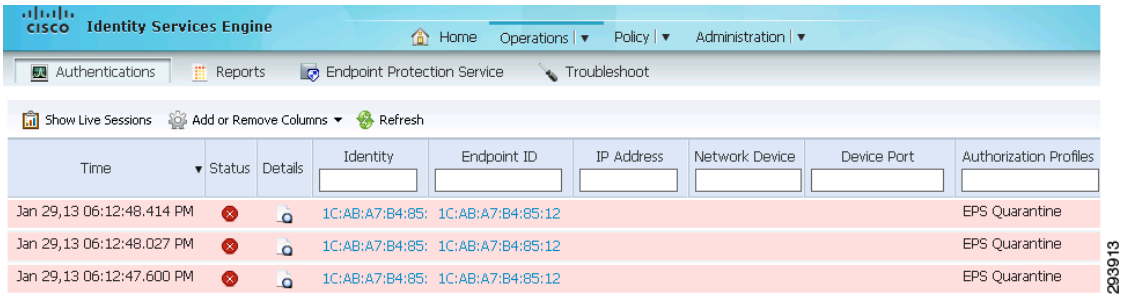
Figure 22-28 EPS



As soon as the administrator clicks **Submit**, the device is forced off the network and future attempts to connect are rejected.

Figure 22-29 shows the EPS_Quarantine authorization profile being applied and how a device is denied access.

Figure 22-29 Quarantined Endpoints

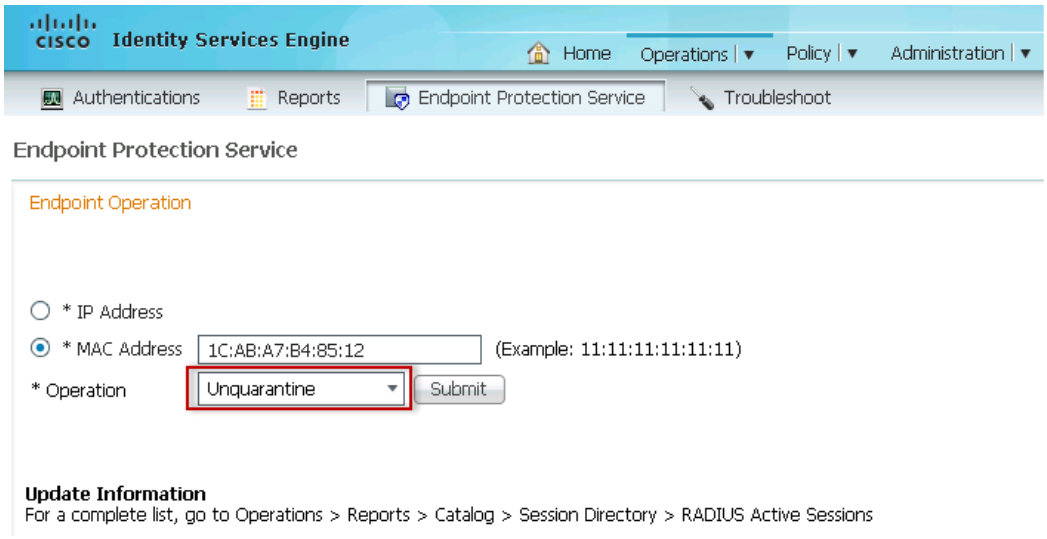


EPS provides an extra layer of control to monitor and change the authorization state of endpoints.

Note Since EPS has higher precedence over blacklisted devices, employees do not have the option to reinstate devices that have been quarantined by the administrator.

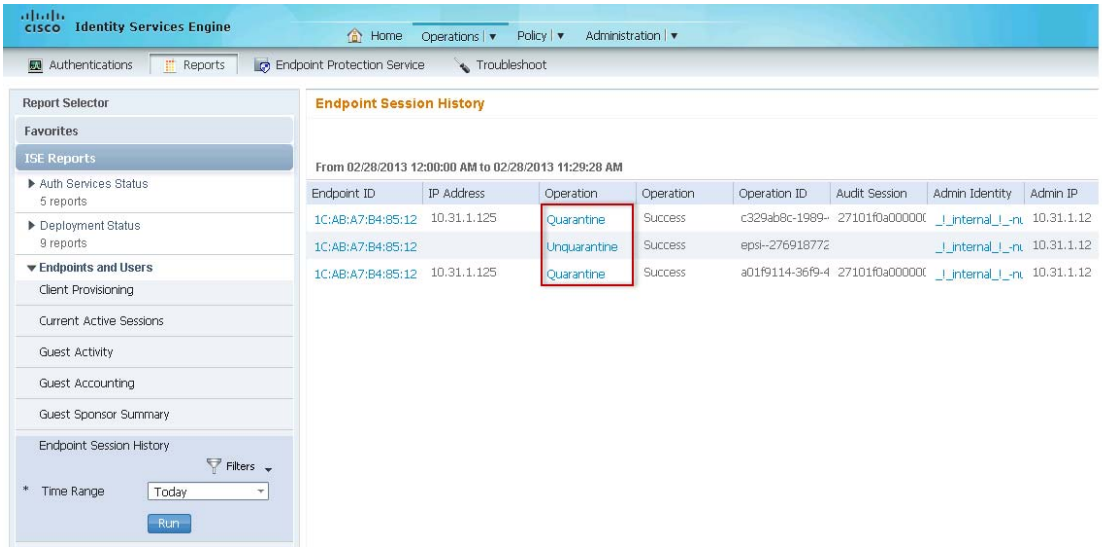
To unquarantine a device, enter the device's MAC address and select **Unquarantine** from the operation pull down menu, as shown in Figure 22-30.

Figure 22-30 Unquarantine a Device



The EPS activities are logged by ISE and can be reviewed by clicking **Operations > Reports > Endpoints and Users > Endpoint Session History**. Figure 22-31 shows one of the reports that includes endpoint information.

Figure 22-31 EPS Logs

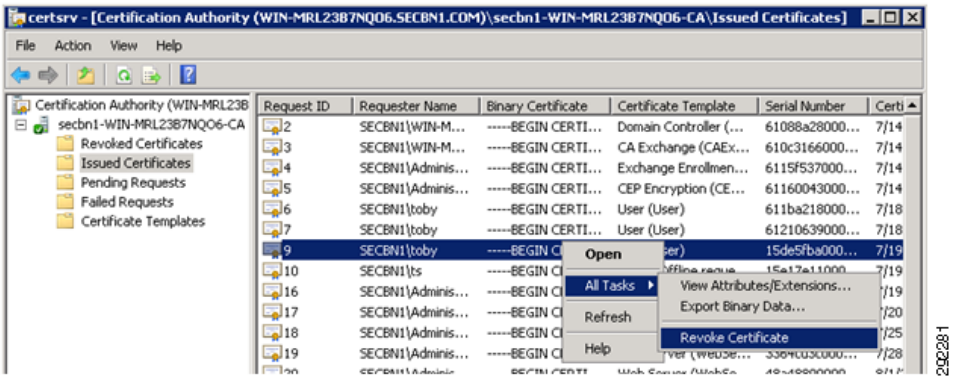


Revocation of Digital Certificate

Administrators also have the option of revoking an employee's digital certificate from the CA server to prevent further use by unauthorized devices. The CA server periodically publishes the Certificate Revocation List (CRL). ISE is configured to validate the certificate presented by the clients against the CRL list. If there is a match, the ISE rejects the digital certificate presented by the client.

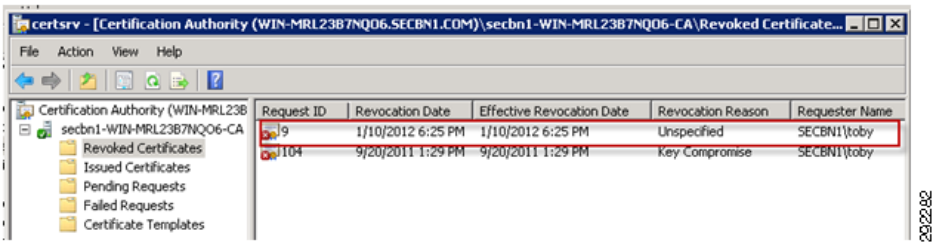
The first step is to revoke the digital certificate from the CA server. Figure 22-32 shows how to revoke the digital certificate for username "toby".

Figure 22-32 Revoking the Digital Certificate on the CA Server



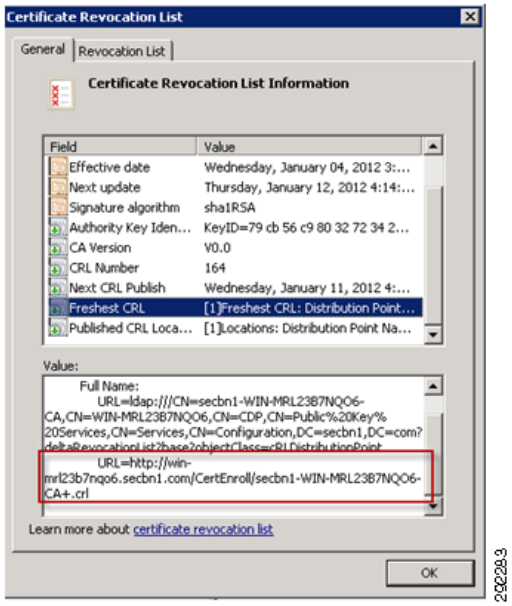
Once the above process is complete, the certificate serial number is added to the Certificate Revocation List (CRL). Figure 22-33 displays the CRL information.

Figure 22-33 Certificate Revocation List



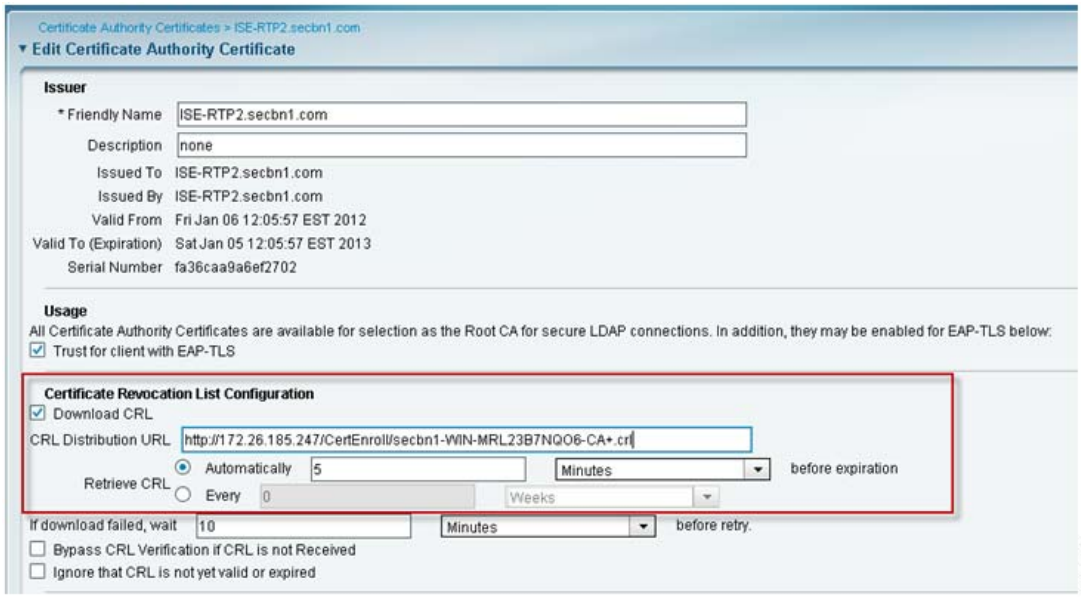
This information is periodically published by the CA server. Figure 22-34 shows the location of the CRL where the ISE can download the list.

Figure 22-34 CRL Distribution Point



The next step is for ISE to be configured with CRL distribution location so that it can periodically download the list and compare it with the certificates presented by the clients. Click **Administration > Certificates > Certificate Authority Certificates** and configure the CRL values, as shown in Figure 22-35.

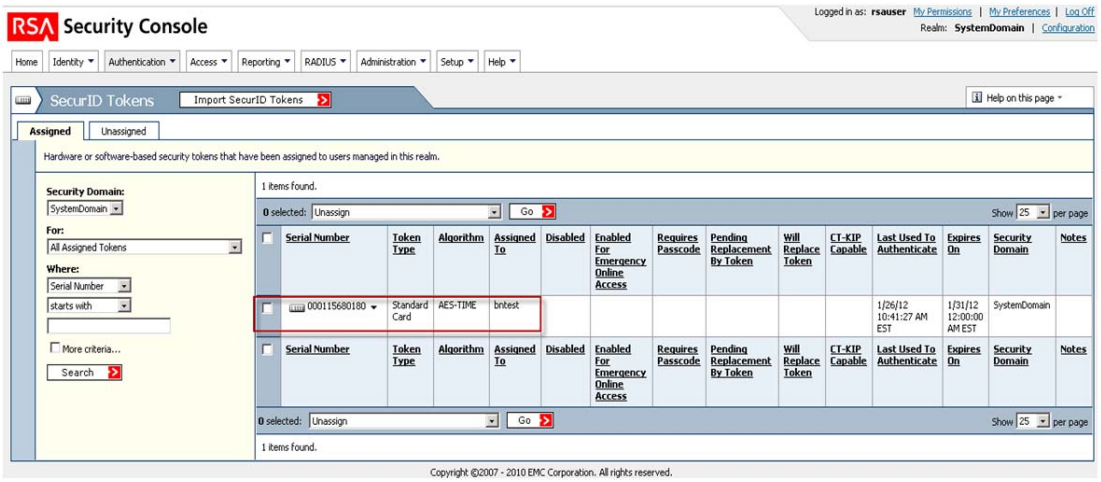
Figure 22-35 CRL Location Information on the ISE



Disable the RSA SecurID Token

When a device that has been previously provisioned is reported lost or stolen, the device must be denied access to prevent unauthorized access to the network. In addition, the remote user's RSA SecurID token must be disabled at the RSA Server so that the remote user cannot use the network. Figure 22-36 shows how to disable the RSA SecurID token at the RSA server.

Figure 22-36 Disabling the RSA SecurID Token





CHAPTER 23

BYOD Policy Enforcement Using Security Group Access

Revised: August 7, 2013

Security Group Tag Overview

The BYOD CVD has primarily relied on Access Controls Lists for policy enforcement to restrict user traffic as appropriate upon successful authentication and authorization. The use of ACLs can become a daunting administrative burden when factoring the number of devices upon which they are applied and the continual maintenance required to securely control network access.

This CVD uses a complimentary technology known as TrustSec and the use of Security Group Tags (SGT). Security Group Tags offer a streamlined and alternative approach to enforcing policy and traffic restrictions with minimal and in some cases, little or no ACLs at all if TCP/UDP port level granularity is not required.

Security Group Tags will be used as an alternative to ACLs for enforcing role-based policies for campus wireless devices where the Cisco Wireless Controllers have been centrally deployed and configured for operation in local mode (wireless traffic locally switched at the controller).

ACL Complexity and Considerations

To date, variations of named ACLs on wireless controllers, static and downloadable ACLs on various routing and switching platforms, as well as FlexACLs for FlexConnect wireless traffic in the branch have been used as a means of enforcing traffic restrictions and policies. In order to configure and deploy these ACLs, a combination of either command line (CLI) access to each device via Telnet/SSH or network management such as Prime Infrastructure have been required and used for statically configured ACLS while the Cisco Identity Services Engine (ISE) has been used to centrally define and push downloadable ACLs (DACL) to switching platforms.

- Unique ACLs may be required for different locations such as branches or regional facilities, where user permissions may need to be enforced for local resources such as printers, servers, etc.
- The operational complexity of ACLs may be impacted by changes in business policies.
- The risk of security breaches increases with potential device misconfigurations.
- ACL definitions become more complex when policy enforcement is based on IP addresses.

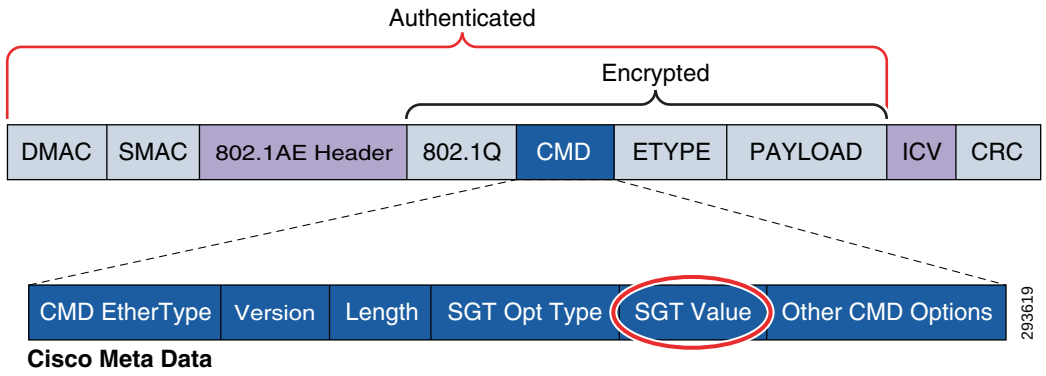
- Platform capabilities, such as processor memory, scalability, or TCAM resources may be impacted by complex ACLs.

Cisco's TrustSec provides a scalable and centralized model for policy enforcement by implementing Cisco's Security Group Access architecture and the use of Security Group Tags.

Security Group Tag

Security Group Tags, or SGT as they are known, allow for the abstraction of a host's IP Address through the arbitrary assignment to a Closed User Group, represented by an arbitrarily defined SGT. These tags are centrally created, managed, and administered by the ISE. The Security Group Tag is a 16-bit value that is transmitted in the Cisco Meta Data field of a Layer 2 Frame as depicted in Figure 23-1.

Figure 23-1 Layer 2 SGT Frame Format



The Security Group Tags are defined by an administrator at Cisco ISE and are represented by a user-defined name and a decimal value between 1 and 65,535 where 0 is reserved for "Unknown". Security Group Tags allow an organization to create policies based on a user's or device's role in the network providing a layer of abstraction in security policies based on a Security Group Tag as opposed to IP Addresses in ACLs.

The SGT is dynamically assigned, or bound, to user/device's IP Address upon successful AAA Authentication and subsequent Authorization to the network via Cisco ISE. This SGT mapping is communicated to and stored at the Network Access Device (NAD) serving as the Authenticator. On Cisco switches, these mappings may be dynamically created through RADIUS Attribute Value (AV) pairs passed down from ISE; or may optionally be defined at the device for a host's IP Address, physical port, VLAN, or subnet, depending on the switching platform's capabilities. In the case of Cisco Unified Wireless Network (CUWN) wireless LAN controllers such as the WiSM2 or CT5508 however, these IP to SGT mappings can only be dynamically created at the controller through the information communicated by ISE. For specific platform capabilities, refer to the appropriate configuration guides found at http://www.cisco.com or in the Cisco TrustSec Switch Configuration Guide at: http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/sgacl_config.html#wp1054201.

For additional information regarding the Security Group Access architecture, refer to the TrustSec Design and Implementation Guide at: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html.

Security Group Access Domain Infrastructure

There are two methods of configuring the infrastructure to support TrustSec enabling the forwarding and policy enforcement of frames with an embedded Security Group Tag.

- TrustSec using 802.1X for Link Encryption
- TrustSec in Manual Mode for Link Encryption

Common to both of these methods is first the ability to employ MACsec (802.1ae) which provides encryption, a message integrity check, and data-path replay protection for links between adjacent network devices thereby protecting the CMD field and the SGT value it contains. Second is configuration for 802.1X authentication between those network devices that will enforce policies based on the SGT and the Cisco Identity Services Engine (ISE) acting as an Authentication Server.

TrustSec using 802.1X for Link Encryption

The first method for configuring an TrustSec infrastructure makes use of 802.1X to establish a domain of authenticated and trusted network devices. Every networking device in the TrustSec Domain must be authenticated either directly with ISE, or through its neighbor acting as an authenticator on behalf of the Supplicant network device.

The first networking device to join a TrustSec Domain is considered to be the "Seed" Device. When first powered on, it acts as an 802.1X supplicant joining the TrustSec Domain through an EAP-FAST exchange with ISE as the Authentication Server. Upon successful authentication with ISE, the network device will, through the EAP provisioning tunnel, receive a Protected Access Credential (PAC) key and secure token generated by ISE. This key is used for all future RADIUS Exchanges with ISE.

Seed devices are configured with the list of ISE servers against which it can authenticate. It is not necessary to provide this list of AAA servers on every device when 802.1X is used and subsequently, as adjacent networking devices configured for TrustSec come up, the seed device will act as an 802.1X Authenticator to its neighbor as a Supplicant. As such these neighbors are considered to be "non-seed" devices. This process is known as Network Device Admission Control or NDAC. Once the networking devices have authenticated against ISE, a common Pairwise Master Key (PMK) is derived for use by both sides during subsequent mutual authentication and MACsec negotiation with optional encryption for each of the interconnecting interfaces.

Finally each device, using the credentials acquired after successful authentication with ISE, will through Secure RADIUS exchange, acquire SGT definitions and policies to be enforced in the network.

For additional information regarding NDAC and MACsec, please refer to the Cisco TrustSec 3.0 document "Introduction to MACSec and NDAC" which can be found in Design Zone on Cisco.com at: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html.

TrustSec in Manual Mode for Link Encryption

The second method, depicted in this CVD, does not rely on 802.1X for device link authentication and encryption. To enable link protection without 802.1X, TrustSec manual mode can be configured such that a common Pairwise Master Key (PMK) is manually configured on the respective interfaces for use by both sides during subsequent MACsec negotiation and optional encryption. This is one of the primary differences with TrustSec with 802.1X wherein this key is dynamically derived from the credentials acquired from ISE after successful authentication.

Another distinction between 802.1X mode and manual mode lies in how the concept of a TrustSec domain of trust is established. When using 802.1X for link authentication, the network device credentials are used during the process of bringing the link up and subsequent authentication with the peer interface; this establishing a trust state with the adjacent device. As this 802.1X-based mechanism is unavailable when configuring manual mode, a static policy must be defined establishing a trusted state on both side of a TrustSec enabled link in addition to the manual PMK configuration.

As in the case of TrustSec with 802.1X link authentication, communication on the links between trusted devices in the TrustSec domain can be secured with a combination of encryption, message integrity checks, and data-path replay protection mechanisms through the use of 802.1ae MACsec. This encryption capability allows the SGT value carried in the Cisco Meta Data (CMD) field of the 802.1q Header to be protected. Today, there are two keying mechanisms available for use with 802.1ae based encryption, the first is a Cisco proprietary protocol known as Security Association Protocol or SAP (similar to 802.1X-2010 MKA) and the second, a standards-based mechanism known as MACsec Key Agreement or MKA. Both use Galois Message Authentication Code (GMAC) as a mechanism to provide authentication and 128-bit AES-GCM (Galois/Counter Mode) symmetric encryption, which is capable of line-rate encryption and decryption for both 1 GB and 10 GB Ethernet interfaces, and provides replay attack protection of every frame. Within this CVD, SAP is used as the keying mechanism for all 10GE MACsec links.

SAP, is a key derivation and exchange protocol based on a draft of IEEE 802.11i which performs the following functions:

- Negotiate cipher suite for data traffic.
- Derive session keys for data traffic.
- Exchange SCIs (Secure Channel Identifier) that will be used by data traffic.
- Ensure that the exchange is being performed with the same devices that participated in authentication.
- Perform the exchange with an acceptable degree of security (i.e. confidentiality, protection against MiM attacks, message integrity, etc.).

SAP supports the following modes:

- gcm-encrypt—GMAC authentication, GCM encryption
- gmac—GMAC, authentication only, no encryption
- no-encap—No encapsulation, no SGT plain Ethernet
- null—Encapsulation/SGT present, no authentication, no encryption

When configuring the TrustSec infrastructure to make use of Manual Mode on the links as opposed to 802.1X mode, there is no requirement to configure every network device to support 802.1X-based link authentication. However, the requirement for 802.1X configuration and network device authentication at ISE still exists for those devices that will require SGT definitions and mappings as well as for SGT-based policy enforcement.

In order to authenticate, the network device(s) will require a configuration identifying the AAA servers (ISE) against which they will authenticate. Once a device has successfully authenticated, secure RADIUS using a PAC key and secure token acquired during authentication is used to communicate with ISE to acquire TrustSec environmental data such as the Security Group Name and the numeric value associated with the tag, an optional SGT used by the device to source packets, and SGT/IP mappings that have been created at ISE. Additionally, policies based on SGTs in the form of SGACLs created at ISE are pushed out to those devices capable of enforcing them such as certain Catalyst and Nexus switching platforms.

Note At the time of this writing, the ASA only receives the SG Name and Tag value and does not support dynamic policy download; ACEs containing SGTs must be locally defined on the ASA. The ASA however must store a PAC key/token which is provisioned at the ASA in order to dynamically acquire Security Group Names and Tag values as created within ISE for later use in defining those policies based on SGT. Wireless controller platforms such as the WiSM2 and 5508, although defined at ISE to support 802.1X wireless client authentication, do not download any TrustSec environment data but merely receive SGT mapping information upon successful client authorization to be discussed later.

In the BYOD v2.5 CVD, this 802.1X configuration for network device authentication will be configured at only two locations in the infrastructure to be discussed later; a Shared Services block where the wireless controllers have been deployed and at the data center aggregation switches. As discussed later, these will be the two locations in the network where policies based on SGTs will be enforced using dynamically acquired SGACLs. Although it is entirely possible to configure every device in the path for 802.1X authentication it is purely optional as SGACLs are enforced upon egress from the device having a corresponding destination IP Host to SGT mapping matching an SGACL. The devices that are in the path between source and destination will not have this mapping typically and hence this TrustSec environment data will not be applicable or enforceable.

Security Group Tag Distribution and Forwarding Mechanisms

In order to impose or forward a frame with an SGT Value, specialized switching ASICs are required for forwarding on Ethernet Ports. A variety of Cisco switching platforms in the Nexus and Catalyst families support the inline tagging of an SGT value on 10G Ethernet Ports and to a lesser extent, 1G Ethernet depending on the platform. This SGT inline tagging capability is sometimes referred to as SGToverEthernet or “native tagging”. In the BYOD v2.5 CVD the Catalyst 6500 with SUP2T and WS-X6904 linecards as well as the Nexus 7000 with SUP1 and M1 linecards will provide the 10GE connectivity supporting SGT imposition and forwarding in the Services, Core, and data center in the campus network. For specific information regarding platform support for SGT inline tagging, refer to the appropriate product documentation at

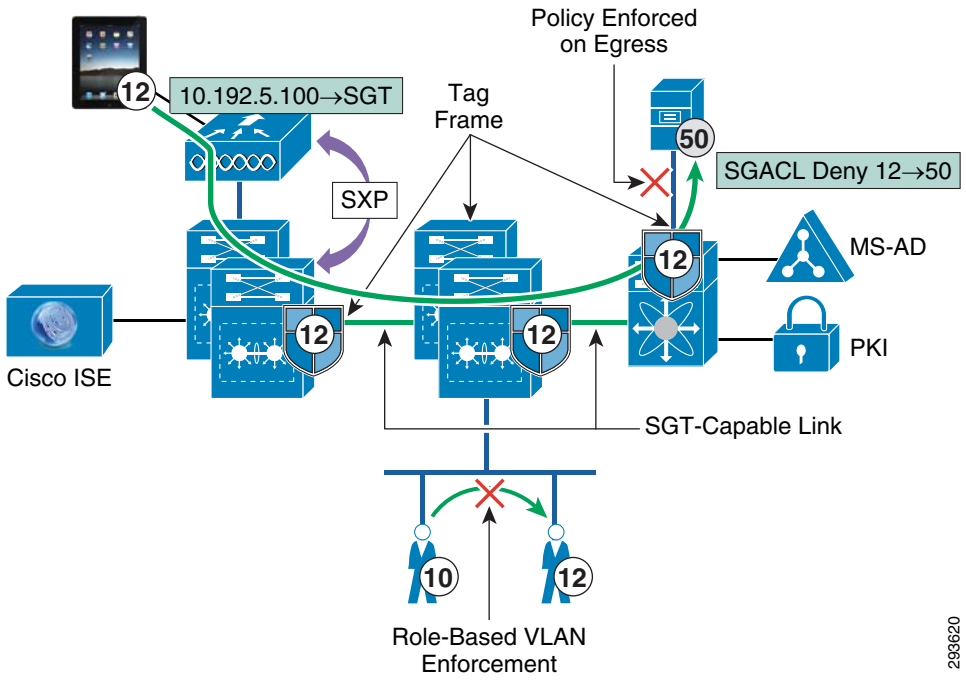
http://www.cisco.com or the “TrustSec Switching Configuration Guide” at: http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/trustsec.html.

For the network access devices (NAD) capable of SGT native tagging, when a user or device is authenticated and a SGT Value has been forwarded as a RADIUS AV to the NAD, this IP to SGT mapping will be created and stored at the access switch. Frames sourced from that user or device will be tagged upon egress from the switch on a physical port or, in the case of an SVI, internally associated with the packet as it egresses the SVI. In either case, if an SGACL is matched prior to egress from the device or SVI, the action defined in the SGACL will be enforced otherwise the packet with the imposed SGT will be forwarded.

Note SGACLs are always enforced upon egress from a device or SVI and only if the network device has a local mapping for the destination SGT enforced by the SGACL.

In certain Catalyst and Nexus switching products, it is also possible to enable role-based enforcement within a VLAN, where an SGACL may be enforced between devices with different SGT values. Refer to Figure 23-2.

Figure 23-2 SXP Advertisement, SGT Imposition, and SGACL Enforcement



Packets that have a Security Group Tag applied can be forwarded throughout an infrastructure as long as those network devices support SGToverEthernet or native tagging and the link has been configured with the appropriate policy defining whether those tags should be trusted or re-written through the use of the <policy static> command on the TrustSec interface. As a packet arrives at a switch that supports SGToverEthernet for example, the switch will remove and inspect the header to perform forwarding lookups, apply any QOS treatment, and act upon any security ACLs configured there. Providing the intermediate device does not have an IP to SGT mapping that is denied in an SGACL at this device, the packet will be forwarded along with the associated SGT towards the destination where an egress SGACL will be enforced (permit or deny).

For those platforms that do not support the native SGT tagging capabilities, the SGT eXchange Protocol or SXP as it is known was created to advertise IP Address to SGT Binding information. On devices that support SXP only, they are considered to be “SGT-Aware”, the 802.1X authentication and authorization of a user is exactly the same as those devices supporting native tagging. On these devices supporting SXP, the IP to SGT mapping is created and maintained and can be advertised to a device where native tagging is supported. This advertisement or SXP “Peering” as it is known can be created to an adjacent device or one that is multiple L2 or L3 hops away as SXP uses TCP as a communications transport between peered devices. Refer to Figure 23-2.

As untagged packets sourced from a device advertising IP to SGT mappings via SXP arrive at a switch that is capable of native tagging, the source IP Address is identified and either the associated SGT can be added to the packet and forwarded or an applicable SGACL enforced.

For a detailed explanation of SXP and SGT, see the TrustSec Design and Implementation Guide at: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html.

User Policies with SGT

When a user/device accesses the network for the first time, whether wired or wireless, as described in the CVD the network access device (NAD), wireless controller or switch, serves as an 802.1X authenticator to start the authentication process. The AAA server against which the user will be authenticated will be the Cisco Identity Services Engine (ISE). As this is the first time the device has been seen on the network, it will first be provisioned with the proper credentials for access to the network. During this provisioning process, access to the network will be restricted though the use of standard ACLs to those services such as DHCP, DNS, ISE, and the Google PlayStore required for on-boarding the device as specifically defined in the BYOD CVD.

Once a device has been successfully on-boarded and provisioned, the network access device (NAD) once again acts as an 802.1X authenticator to start the device/user authentication process with ISE. During this authentication process, the device/user is identified based on credentials offered such as a Digital Certificate or Active Directory Group membership. In past versions of the BYOD CVD, once authenticated, users or devices would have been authorized and based on a matching policy would have either been granted unrestricted access to the network or perhaps partial access, restrictions enforced by a suitable ACL either downloaded in the case of switches or associated with a statically configured (named) ACL on the networking device. As an alternative to this approach, an SGT can be used and upon identifying the appropriate authorization policy in ISE, the NAD will receive and store the appropriate SGT to be associated with the user or device's IP Address, commonly referred to as an IP to SGT Binding.

Based on these Security Group Tags, role-based policies can be enforced on supporting hardware through the use of Security Group ACLs (SGACLs) on Cisco switching infrastructure, policies defined on Security Group Firewalls (SGFW) such as the ASA, or an SGFW implemented on the IOS Zone-Based Firewall (ZBFW) on Cisco Routers. These policies may be as simple as a permit or deny an SGT statement or may include specific IP Port information in addition to source or destination SGT to identify specific applications or traffic.

It should be apparent, that when an abstraction layer such as the TrustSec architecture and SGTs are used, device and virtualized server mobility is greatly enhanced as the IP Address of the device is no longer a consideration in enforcing policies in the network. This is true as long as the SGT value was dynamically assigned through a port profile when using the Nexus 1000V, by ISE based on authorization policy, or if the mapping was the result of a VLAN, L3 interface, or IP Subnet to SGT Binding with the VLAN, L3 interface, or Subnet duplicated on other devices. Now, as an entity moves in the network either through mobile roaming or server vMotion by virtue of the port profile when using the Nexus 1000V, one need not be concerned with having appropriate address-based ACLs defined on the destination device. The policy can follow them based on the SGT they have been assigned. If however the IP Host Address to SGT mapping were statically defined at a networking device, that mapping is only resident on that device and not shared with other devices in the TrustSec Domain.

Through the use of Security Group Tags, it will be possible to eliminate many of the Downloadable and named ACLs required in previous BYOD CVDs with a ubiquitous set of tags applicable to an entire domain and managed centrally at the ISE.

This CVD uses TrustSec as an alternative to named ACLs for campus wireless (centralized) access. As of Cisco Unified Wireless Networking (CUWN) software release 7.4MR1 and 7.5 for the WiSM2 and CT5508, sixty-four ACLs can be configured, each having sixty-four Access Control Entries or ACEs (permit/deny statements) within an ACL. In most organizations this may not be an issue, however in others, this may be too limited when using ACLs to segment the network based on roles or device types or if the devices are a Corporate or Personal asset. When using ACLs, a Named ACL is created on the wireless controller and as a user is authenticated and authorized, ISE pushes down the name of the ACL to be applied to the user in the RADIUS AV pair returned to the controller. If there are more than sixty-four unique roles, or more likely more than sixty-four Access Control Entries in an ACL, the use

of named ACLs will not be possible. For these scenarios, TrustSec offers an extremely scalable alternative where hundreds of roles may be identified for users or devices, thereby eliminating the requirement for the use of specific IP addresses in an ACL.

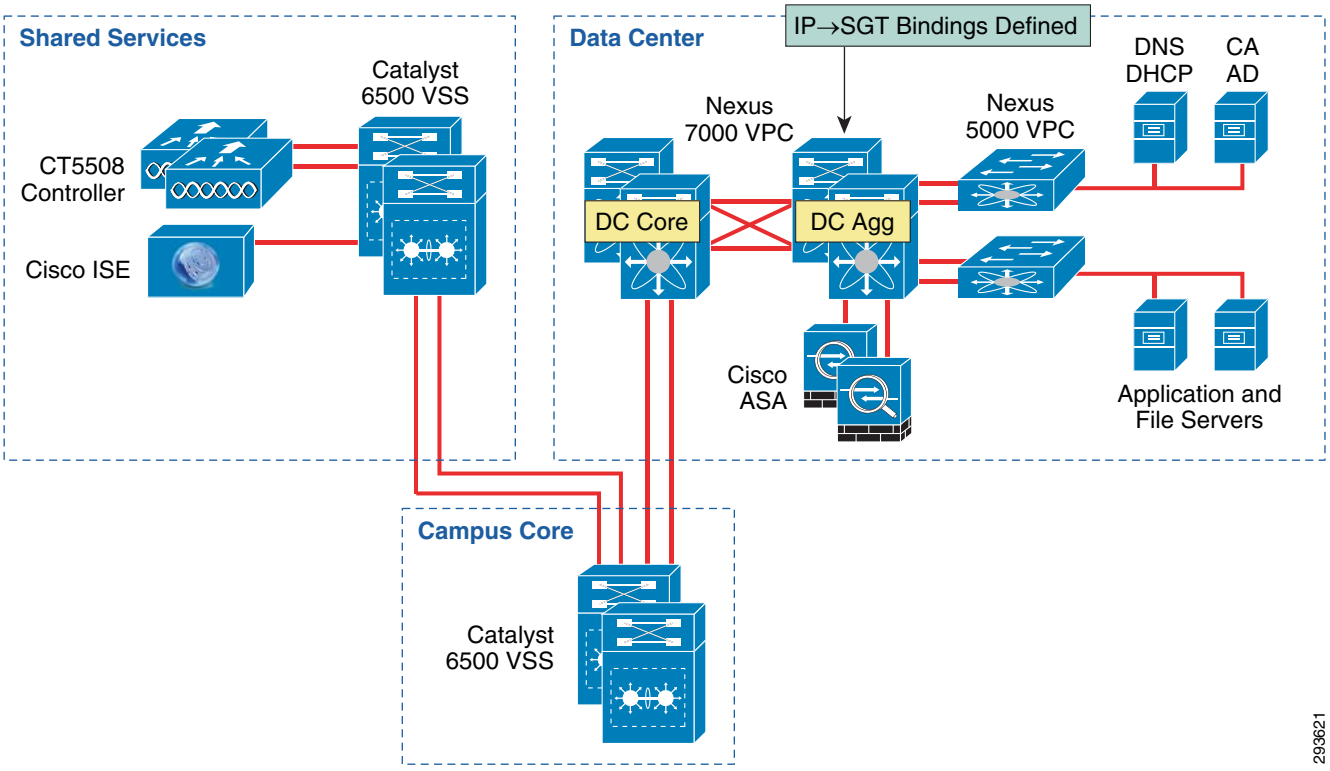
The policies demonstrated in this CVD cover the use case where North/South (wireless access to data center) policies need to be enforced restricting access to resources in the data center. Although entirely possible, East/West (wireless user to user) policy enforcement using SGT will be considered out of scope for this version of the CVD. East/West traffic enforcement will be included in a future version of the BYOD CVD when wired access with SGT is addressed.

SGT Assignment for Data Center Servers

Unlike campus access through Catalyst Switches and Cisco Wireless Controllers where dynamic SGT mappings are communicated and created through 802.1X and RADIUS exchange, the vast majority of organizations do not implement 802.1X for server connectivity. As such, data center switches such as the Cisco Nexus switches provide only limited support for the use of 802.1X and do not specifically support an SGT RADIUS AV as an option. Although 802.1X support would be available if using Catalyst Switches in the data center, this use case is not covered. Therefore, IP Address to SGT mappings for these resources must be either manually defined as in the case of bare metal servers and non-Cisco virtual switches or within port profiles if the Cisco Nexus 1000V virtual switch is deployed.

For purposes of this CVD we have defined our IP to SGT Bindings at a Nexus 7000 data center aggregation layer switch as depicted in Figure 23-3.

Figure 23-3 Server IP to SGT Bindings



In addition to the manual creation of an IP Address to SGT mapping either globally, or within a VLAN for enforcement of intra-vlan traffic between hosts belonging to different Security Groups, the Nexus 7000 also supports the following:

- Assigning an SGT to a port for all data sourced from a host attached to that port.
- Future support (NX-OS v6.2, Summer 2013) for mapping an SGT to a VLAN such that all traffic from hosts within that will be tagged accordingly.

The Nexus 5500 as of this version of the CVD, only supports the manual IP to SGT mapping Globally or within a VLAN as previously described for the Nexus 7000.

For the Nexus 1000V, the SGT can be assigned within the port profile definition and subsequently advertised via SXP to a device such as the Nexus 7000 or 5500 for SGT imposition or SGACL enforcement. The use of the Nexus 1000V is considered out of scope for this release of the BYOD CVD, however additional information can be found in "Segmenting Clients and Servers in the Data Center" at: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html.

Finally, it is also possible to define these SGT host mappings within ISE. These mappings are subsequently pushed to all SGT-capable network devices authenticated within the TrustSec Domain. An SGT-capable device can impose and forward the SGT as well as enforce an SGACL. The exceptions here are the ASA, as its Security Group policies are not dynamically obtained but locally configured, and devices that are only SGT-Aware (only SXP supported). Although an effective means of dynamically deploying these mappings throughout the network, it may be impractical from a scaling perspective as every device would possess every server mapping within the TrustSec Domain. Although possible to access each networking device and remove the mapping manually within the configuration this may prove to be very operationally intensive. Within the BYOD CVD we advocate a locally defined configuration in the Aggregation switches closest to the server and once available on the Nexus 7000, recommend the use of VLAN to SGT mapping; note that VLAN to SGT mapping is currently unavailable for the Nexus 5500.

As the implementation of TrustSec within the data center will most likely be through a migrational approach, we advocate that those resources or servers that all users have access to be some of the first identified for policy enforcement through the use of SGT and assignment to a security group. In doing so, as users or devices are associated with an SGT representing a role that has access to only these devices and nothing else, a policy can be created granting that specific access while denying everything else based on the use of SGT 0 or the "Unknown" tag.

A key concept to be considered when granting access to the data center is that of the SGT value of zero or "Unknown" as it is referred to. If the IP Address of a server has not been mapped to an SGT at the point of enforcement such as the Nexus 7000 in the data center, that server would be considered Unknown and associated with SGT 0. Unlike ACLs with an implicit deny at the end, SGACLs when implemented on a switching platform have an implicit permit to Unknown or all; this is not true on the ASA or IOS ZBFW acting as a SG-FW where an implicit deny is still maintained. Hence on a switch, if there is not a specific tag value assigned to a server, the destination is considered Unknown (SGT 0) and the packet forwarded. This SGT 0 thus allows a migrational approach to tag assignment in the data center.

In a BYOD setting where personal or contractor assets are permitted on the network and only partial access to data center resources is permitted through the use of SGACLs on the switching infrastructure, the task of assigning an SGT to every data center asset will likely prove daunting when first migrating to the use of Security Group Tags as it is not possible to create a switch-based SGACL using both SGT and IP addresses; the destination SGT of the device must be known. Hence it is with this in mind that the concept of the Unknown tag may be used to provide a phased approach to SGT assignment. Throughout the campus infrastructure the default policy permitting any SGT to "Unknown" is left unaltered. However, at the Nexus 7000 Data Center Aggregation switches where policies based on

SGACLs are enforced, an explicit deny of a given source SGT to "Unknown" can be created. A policy that is explicitly (manually) configured at a networking device will take precedence over a policy that is dynamically received from ISE.

In the BYOD CVD the approach followed is to group users and devices with only partial access to the data center and assign them to SGT12 for example. Within the data center, we group those resources that SGT12 will have access to together using an SGT value 40. Servers that these users should not be able to access and that have already been associated with as SGT are assigned SGT 50. In doing so we enforce three policies regarding server access:

- Permit SGT12 to SGT40.
- Deny SGT12 to SGT50.
- Deny SGT12 to Unknown.

For users or devices that have full access to data center resources and are assigned a different SGT, the default, implicit permit to Unknown is left in place and any other SGT-based restrictions can be enforced as required. An example might be in the case of a web server and its corresponding database where a corporate device can access the web server but only the web server and DB Admins can access the database.

When using an ASA firewall to protect data center resources additional flexibility in policy definition is possible as either a source or destination SGT can be used with a destination or source IP Address in the ACE. This now allows one to create policies where a tag assigned to users/devices with partial access can be granted or denied to specific SGTs and or IP Addresses. When using an SG-FW to enforce policy, an ACE can be defined denying a source SGT access to "Unknown" as well.

Prior to implementation of a Security Group Access architecture, many organizations may have already designed their data centers in such a way as to protect those resources by placing them behind a firewall. Others may simply elect to secure them through the use of access control lists where only specific access has been granted while general access from all other sources is denied. Both approaches are demonstrated within this CVD, where campus Wireless BYOD users and devices have role-based access through the use of Security Group Access Control Lists (SGACLs) and the Security Group Firewall (SG-FW) capability found in the ASA. It should be pointed out that these two approaches, SGACLs and SG-FW, are not mutually exclusive and may be used together.

It is beyond the scope of this CVD to provide a detailed approach to developing a migration strategy for the implementation of Security Group Tags in the data center as well as providing architectural guidance for secure, containerized, data centers. For additional information on these topics, refer to the data center document repository in Design Zone on Cisco.com at: http://www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html as well as the TrustSec document repository at: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html.

SGT Deployment Scenarios in this CVD

The BYOD CVD addresses four different Use Cases based on the type of network access that an organization will permit.

- Enhanced Access—This use case provides network access for personal devices, as well as corporate issued devices. It allows a business to build a policy that enables granular role-based application access and extends the security framework on and off-premises.
- Limited Access—This use case enables access exclusively to corporate issued devices.

- Advanced Access—This comprehensive use case also provides network access for personal and corporate issued devices. However, it includes the posture of the device into the network access control decision through integration with third party Mobile Device Managers (MDMs).
- Basic Access—This use case is an extension of traditional wireless guest access. It represents an alternative where the business policy is to not on-board/register employee wireless personal devices, but still provides Internet-only or partial access to the network

Specifically, SGT will be used as a means of policy enforcement for the Enhanced Access Use Case where a campus wireless user/device terminated centrally at a wireless controller in Local Mode is granted either full or partial access to the network. Different classes of servers will be defined to which those users may or may not have access to. The CVD also defines a class that has access to the Internet only. For BYOD v2.5, an ACL on the wireless controller is used to deny access to all but the Internet. The Converged Access products such as the Catalyst 3850 and CT-5760 will not be addressed relative to SGT in this CVD as Security Group Tags and SXP are currently not supported. More about SGT and the Enhanced Use Case will be discussed in the ensuing section discussing the actual authorization policies.

Two deployment scenarios will be depicted within this CVD, the first will make use of SGACLs to enforce policies at the Nexus 7000 data center switches, whereas the second scenario will enforce policies configured at a Cisco ASA configured as a Security Group Firewall (SG-FW). Again, these deployment scenarios are not mutually exclusive and can be used together. This first scenario can be seen in Figure 23-4 and the second scenario follows in Figure 23-5.

Figure 23-4 Policy Enforcement using SGACL

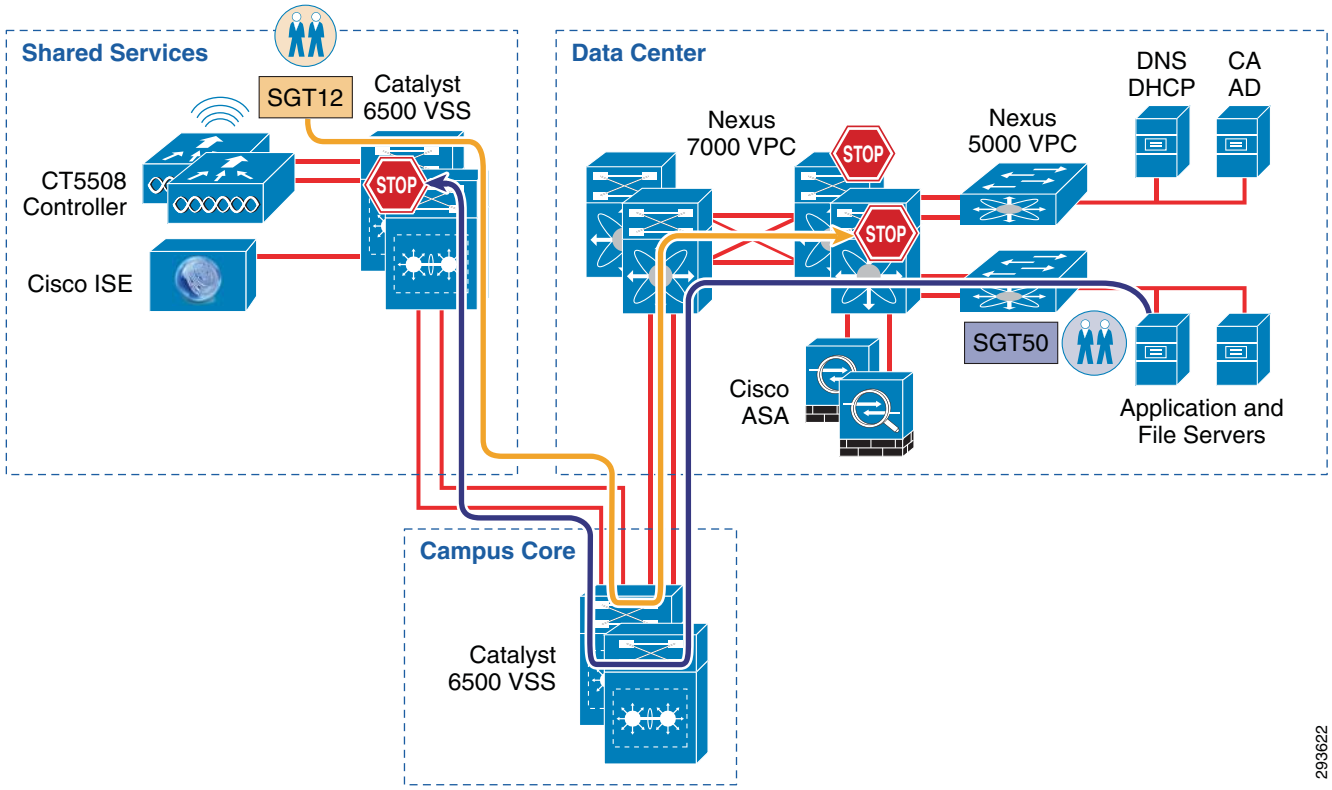
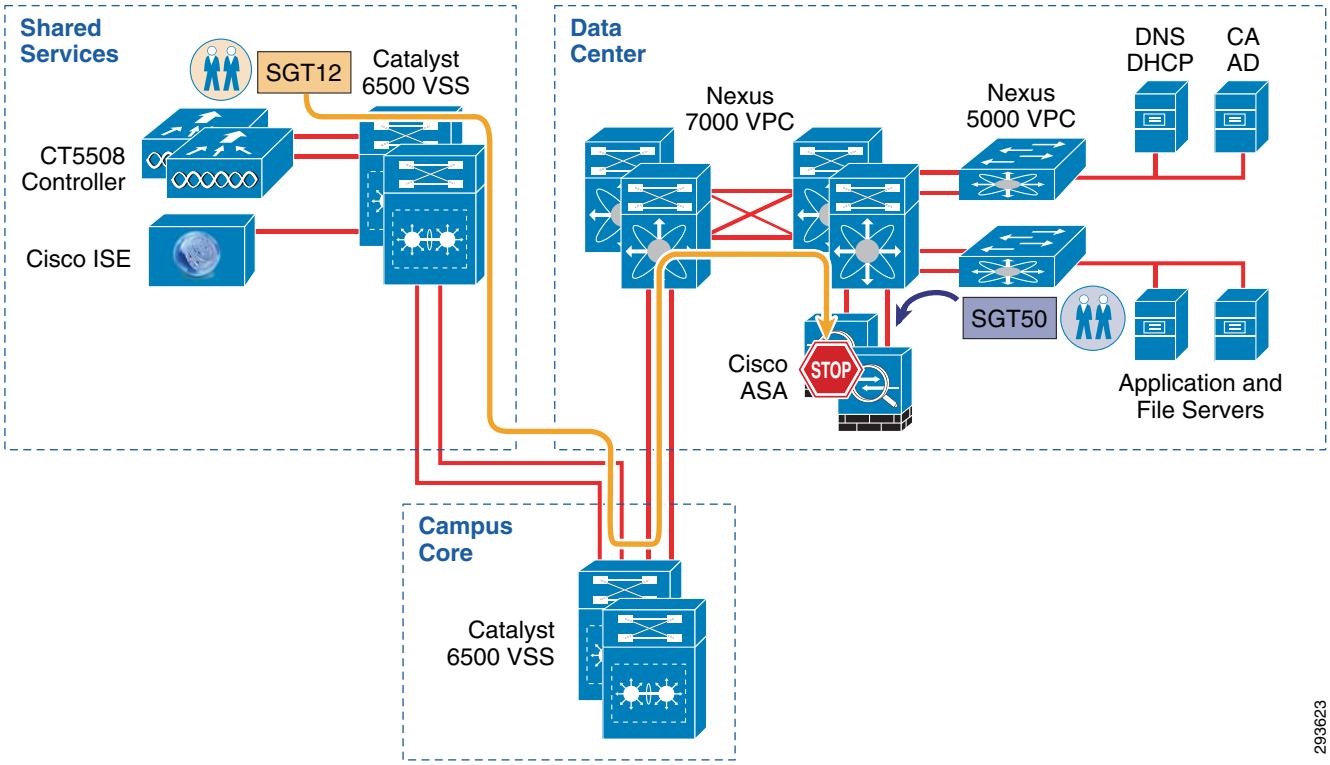


Figure 23-5 Policy Enforcement using SG-FW



Configuring the Infrastructure for TrustSec

The following section describes the infrastructure to be used in this CVD and provides an outline of the two deployment scenarios to be used to enforce policies based on Security Group Tags. These deployment scenarios are not mutually exclusive and may be used together to satisfy an organization's requirements. Configuration details for the infrastructure are provided.

Unified Infrastructure Design to Support TrustSec

As described previously in the Design Overview section, two specific infrastructure deployment scenarios will be examined within this CVD. The first use case will make use of TrustSec Policy defined at the Identity Services Engine and the resulting SGACLs being dynamically exchanged with the Catalyst 6500 and Nexus 7000 infrastructure. The second use case will once again make use of the TrustSec Policy defined at the Identity Services Engine but will enforce this policy through the configuration of Security Group Firewall (SG-FW) policies defined on an ASA providing secure access to data center resources.

In both scenarios, campus wireless users/devices connecting through the centralized CUWN CT5508 controllers configured for local mode, will have access to data center resources based on their authorized roles and enforced through the use of SGT-based policies as implemented in the two deployment scenarios.

The infrastructure components that will be used for the policy enforcement through the use of Security Group Tags in the CVD consist of the following Cisco products.

Wireless Deployment Scenario 1 Components:

- CT5508; CUWN 7.5
- Catalyst 6500
- SUP2-T; 15.1.1-SY1
- WS-X6904 Linecard with 10GE Modules (FourX Adapter)
- Nexus 7000
- SUP1; 5.2(7)
- M1 Linecards; N7K-M108X2-12L
- Nexus 5548
- ISE 1.2 with Patch 1 installed

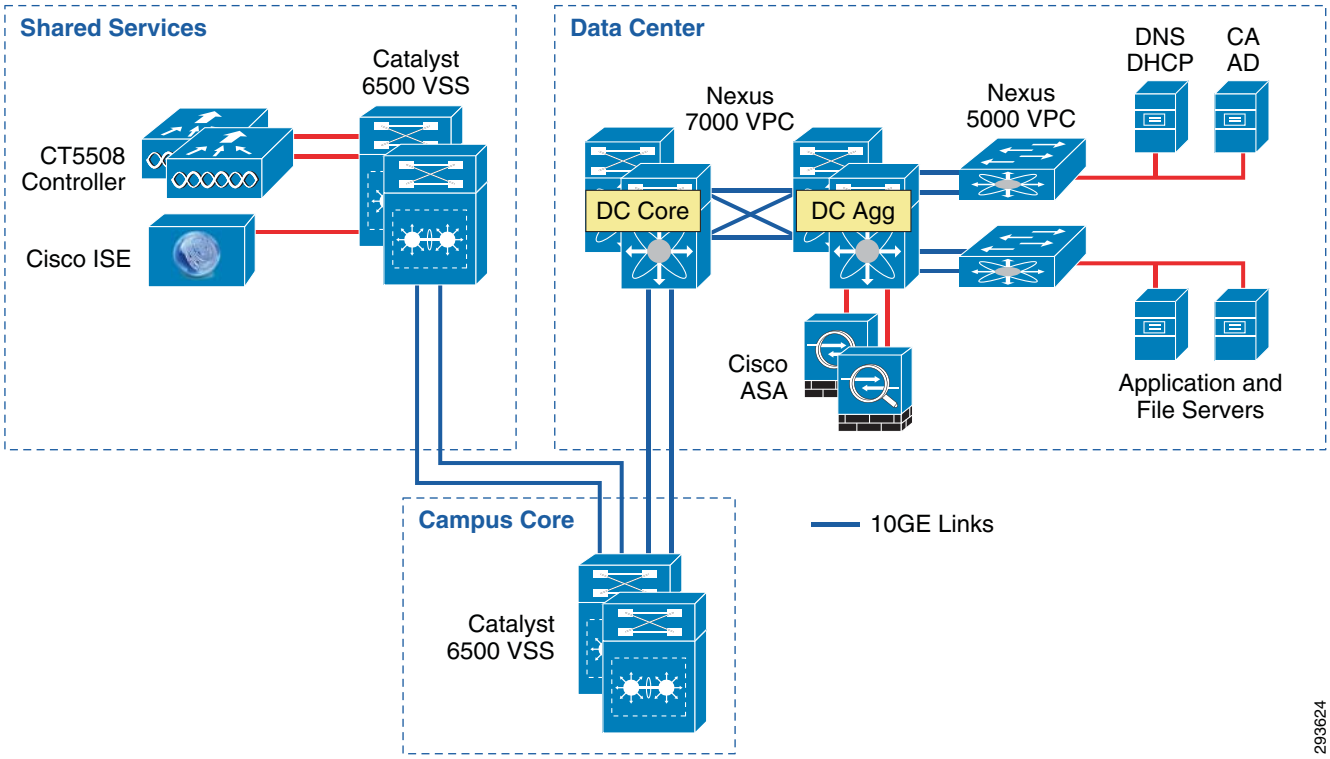
Wireless Deployment Scenario 2 Components:

- CT5508; CUWN 7.5
- Catalyst 6500
- SUP2-T; 15.1.1-SY1
- WS-X6904 Linecard with 10GE Modules (FourX Adapter)
- Nexus 7000
- SUP1; 5.2(7)
- M1 Linecards; N7K-M108X2-12L
- Nexus 5548
- ASA-5520; 9.0(2)
- ISE 1.2 with Patch 1 installed

Note Patch 1 for ISE 1.2 **MUST** be installed in order for NDAC (Network Device Admission Control) to function properly between the network device and ISE. Without Patch 1, the network device will be unable to authenticate with ISE in order to derive TrustSec environment data, PAC file, and security group policies when CTS Manual Mode is configured and, additionally, the credentials required to authenticate peers/TrustSec links when CTS Dot1x Mode is configured. Refer to the ISE 1.2 Release Notes for additional information regarding this **very important** information.

Figure 23-6 depicts the infrastructure that is used for purposes of TrustSec validation within this CVD.

Figure 23-6 TrustSec Infrastructure This BYOD CVD

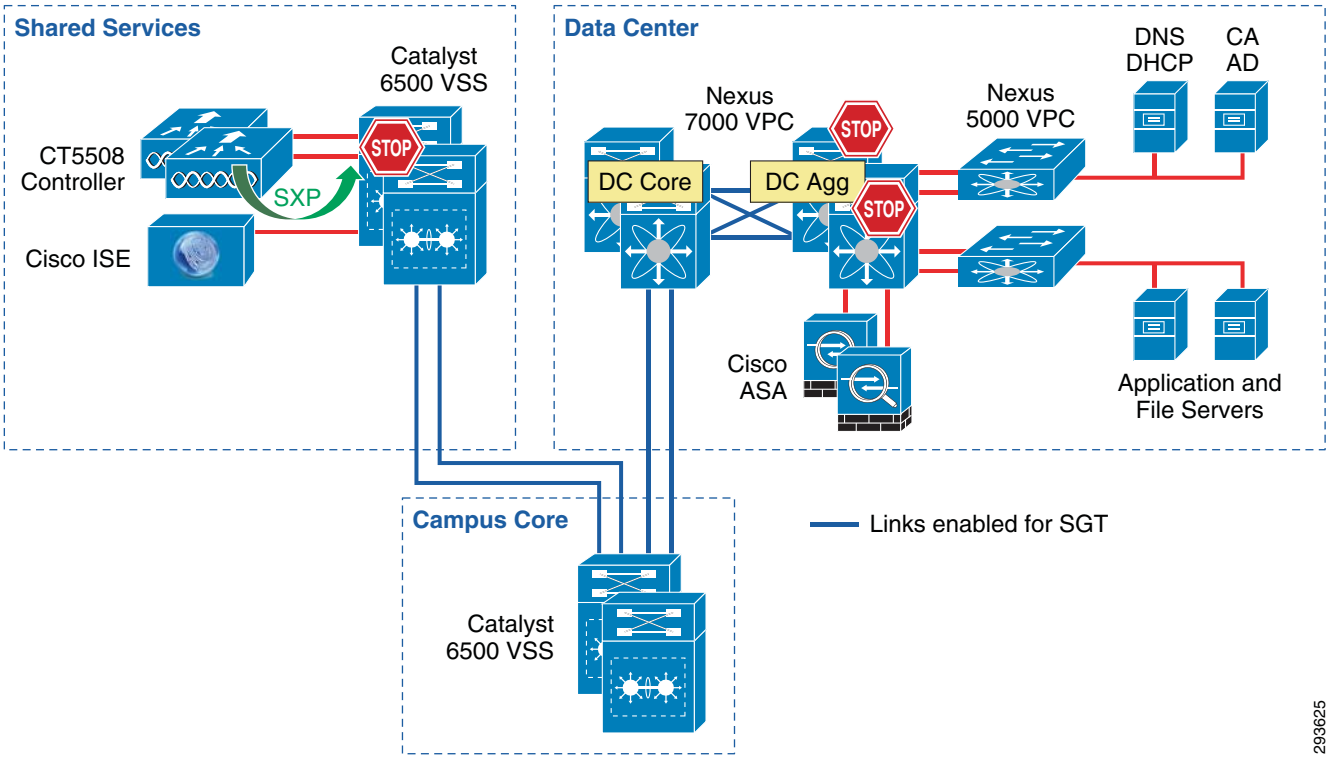


In Figure 23-6 the links extending between the Catalyst 6500 VSS in Shared Services to the Catalyst 6500 VSS in the core and extending to the Nexus 7000 are 10GE links. On the Catalyst 6500s, WS-X6904 linecards with the FourX Adapters provide the 10GE interfaces while the N7K-M108X2-12L linecards provide the Nexus 7000 interfaces. The links between the Nexus 7000 and the Nexus 5548 are likewise connected to N7K-M108X2-12L linecards at the Nexus 7000 and 10GE ports on the Nexus 5548. All other network connectivity for wireless controllers, ASA firewalls, ISE, and the miscellaneous servers depicted are 1GE links.

TrustSec Policy Configuration for SGACLs in Infrastructure Deployment Scenario 1

For Deployment Scenario 1, refer to Figure 23-7.

Figure 23-7 Infrastructure Deployment Scenario 1 SGT Enforcement



Deployment Scenario 1 will require Security Group Tags be forwarded from the Shared Services Catalyst 6500 VSS, where the wireless controller is attached, through the Core of the BYOD infrastructure en route to servers located in the data center proper. In Figure 23-7 above, the links depicted in blue will be configured for SGT forwarding as well as manually configured for 802.1ae MACsec encryption. As previously discussed, the CT5508 wireless controller does not support native tagging on its 1GE interfaces, as such a Security Group Tag Exchange Protocol (SXP) connection will be defined between the controller(s) and the Shared Services Catalyst VSS switch as depicted above.

In this first scenario, wireless users, upon successful authentication and authorization will be associated with a specific role and an IP to SGT Binding will be created on the wireless controller with the device's IP Address and the appropriate SGT. SXP will be used to communicate this mapping to the Shared Services Catalyst 6500s to which the wireless controllers are attached. As wireless user traffic egresses the Shared Services Catalyst 6500s, it will be tagged with the appropriate SGT learned via SXP from the wireless controller. As this traffic traverses the SGT-capable Core, this tag will be propagated hop-by-hop en route to the Nexus 7000s comprising the data center switching infrastructure within which the various servers are located.

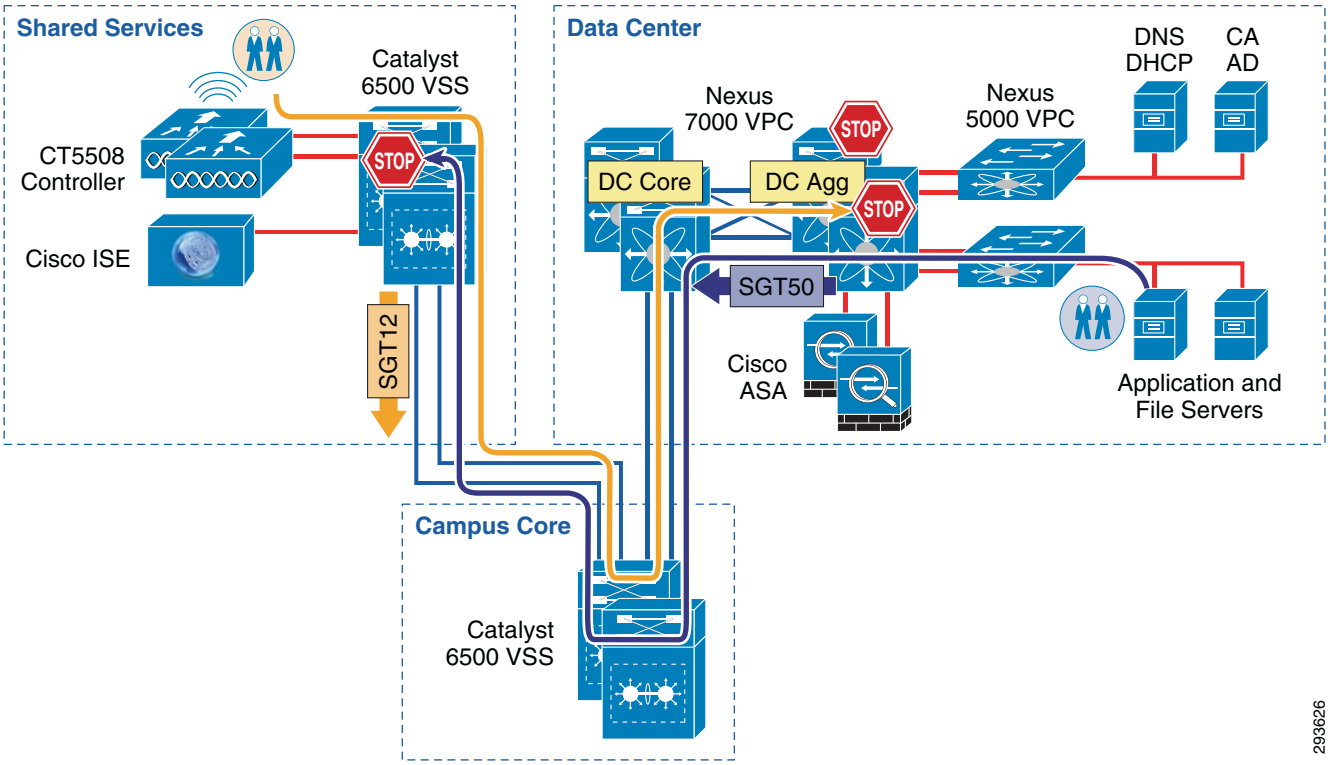
As 802.1X is not used to authenticate the servers residing in the Nexus data center infrastructure, the server IP Address to SGT Binding can either be manually defined on the Nexus 7000 Data Center Aggregation switch or at the ISE server which would subsequently populate that mapping to the Nexus 7000. For purposes of the CVD, these mappings have been manually defined on the Nexus 7000 data center Aggregation Switch. As discussed in the design overview, there are other methods of associating traffic with a specific SGT on the Nexus 7000 platform.

As tagged user traffic arrives at the Nexus 7000 data center switch where the manual SGT mappings for the servers have been created, the traffic will be matched against TrustSec Policy (SGACL) defined either centrally at ISE or locally, as in the case of destination “Unknown” (SGT0), and will be either forwarded or dropped as applicable.

As discussed earlier, all server IP to SGT mappings have been manually created on the Nexus 7000 aggregation switches. As the servers are connected to the Nexus 5548 switches depicted in Figure 23-8, traffic from the Nexus 5548s egresses untagged as no mappings have been created there. Once this traffic passes through the Nexus 7000 Aggregation switch, the resident SGT mappings will be examined and the appropriate SGT imposed upon egress from the aggregation switch. In the event that traffic would be initiated by a server associated with an SGT in the data center, the tagged traffic would then leave the Nexus 7000 data center switches traversing the Catalyst 6500 Core and Shared Service infrastructure with the SGT propagated at each hop towards the destination; the wireless controller attached to the Shared Services 6500. Once the traffic arrives at the Shared Services 6500, the traffic will be matched against TrustSec Policy (SGACL) and will be either forwarded or dropped as defined.

Figure 23-8 depicts where SGACLs will be enforced in the Unified Access infrastructure.

Figure 23-8 Policy Enforcement in Deployment Scenario 1



The following major tasks are required for this deployment scenario and outlined in the following sub-sections:

1. Configuring ISE to support Security Group Access
2. Configuring ISE for Network Access Device Authentication
3. Configuring Network Access Devices for Authentication at ISE
   a. RADIUS server configuration on the CT5508
   b. RADIUS server Configuration on the Catalyst 6500
   c. RADIUS server Configuration on the Nexus 7000
4. Catalyst 6500 Platform Specific Considerations
5. Configuring Switching Infrastructure to Support TrustSec with 802.1ae MACsec Encryption

a. Catalyst 6500 commands

b. Nexus 7000 Commands

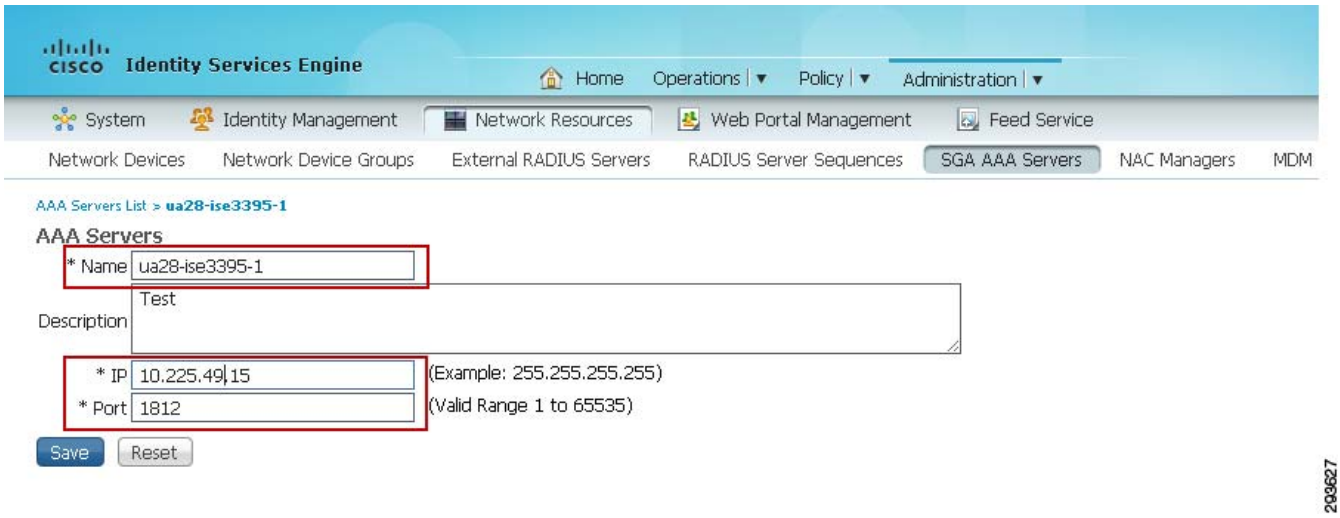
6. Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controller
   a. Wireless Controller Configuration
   b. Shared Services Catalyst 6500 VSS SXP configuration
7. Configuring static IP/SGT Bindings on Nexus switches

Configuring ISE to Support TrustSec

For Cisco ISE to function as a TrustSec server and provide TrustSec services, you must define the following global TrustSec settings. The first step is to define ISE as an TrustSec AAA server as depicted in Figure 23-9.

1. Go to Administration > Network Resources > TrustSec AAA servers and click **Add**.
2. Enter the host name of the Identity Services Engine server or Policy Service Node if ISE Roles have been distributed among dedicated servers.
3. Enter the IP Address of the ISE server.
4. Enter the UDP Port number for RADIUS authentication and click **Save**.

Figure 23-9 ISE TrustSec Server AAA Definition



The next step to be completed is to configure TrustSec Server Protected Access Credential (PAC) Time-to-Live settings and SGT reservations.

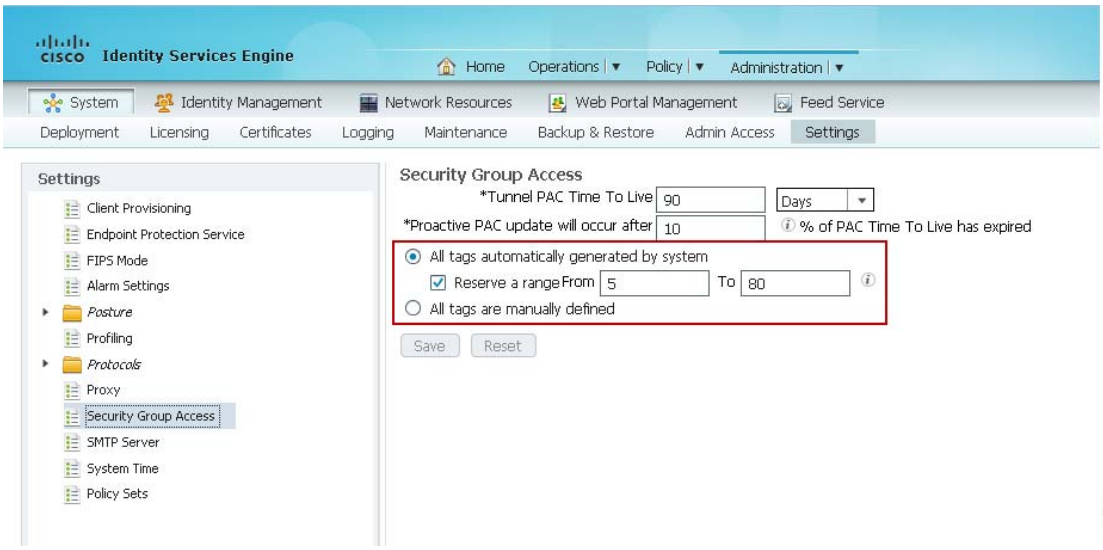
The tunnel PAC generates a tunnel for the EAP-FAST protocol and is used for Secure RADIUS communications with Network Devices for TrustSec environmental data. A new PAC is generated if the network device re-authenticates for any reason or when the TTL expires.

By default Security Group Tags are dynamically assigned a decimal/hex value in ascending order by ISE. It is possible to change this behavior such that all tags must be manually defined or to reserve a range that can be specifically allocated to users, devices, or servers. In the CVD, a range of SGT values is reserved for allocation among users/devices and servers. Figure 23-10 depicts the reservation of Tags 5 through 80 for use in the CVD.

To complete this step:

1. Access the Identity Services Engine and follow the path Administration > System > Settings > Security Group Access.
2. Configure the Tunnel PAC Time to Live.
3. Configure the Proactive PAC update time if desired. In Figure 23-10, the PAC will be renegotiated after 10% of TTL or nine days.
4. By default the system will automatically assign SGT values. If you wish to reserve a range that can be specifically allocated to users, devices, or servers, select the check box next to "Reserve a range From" and specify the Tag values.
5. Save the settings.

Figure 23-10 TrustSec Servers Settings in ISE

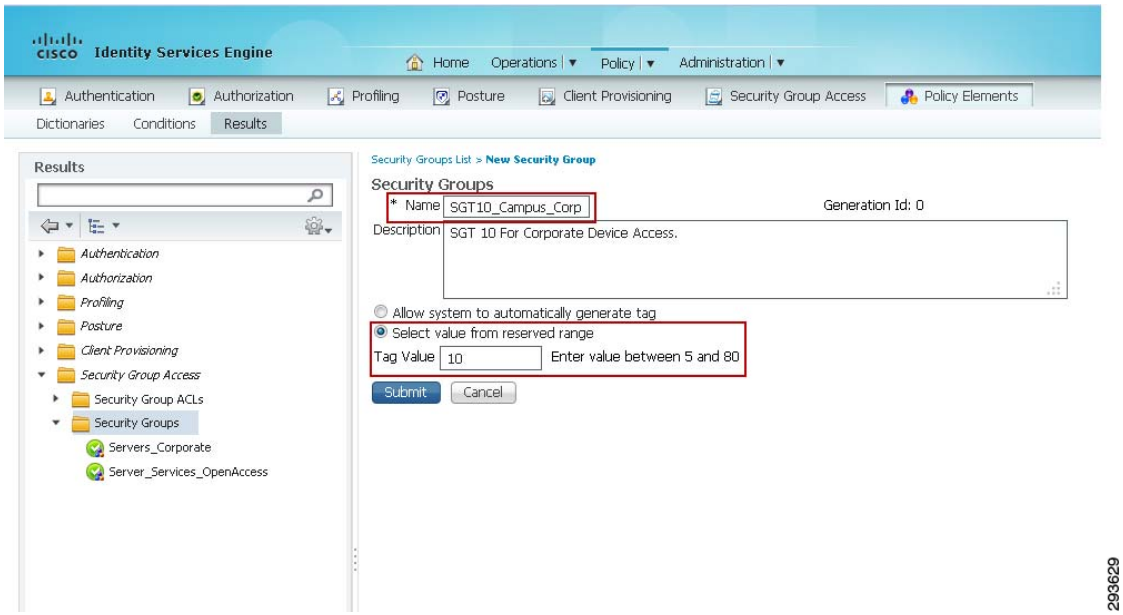


The next step is to define Security Group Tag names and associate them with a numerical value at the Identity Services Engine. The SGT names are periodically pushed to the Network Access Device (NAD) through periodic updates or upon that network device's authentication (NDAC Authentication) with ISE. They may also be manually pushed as well.

To complete this step as depicted in Figure 23-11:

1. Go to Policy > Policy Elements > Results >Security Group Access > Security Groups.
2. Click "**Add**".
3. Define the SGT Name and add an optional description.
4. Click the radio button next to "Select value from reserved range"
5. Enter the desired SGT value from the range defined in the previous step.

Figure 23-11 Security Group Creation



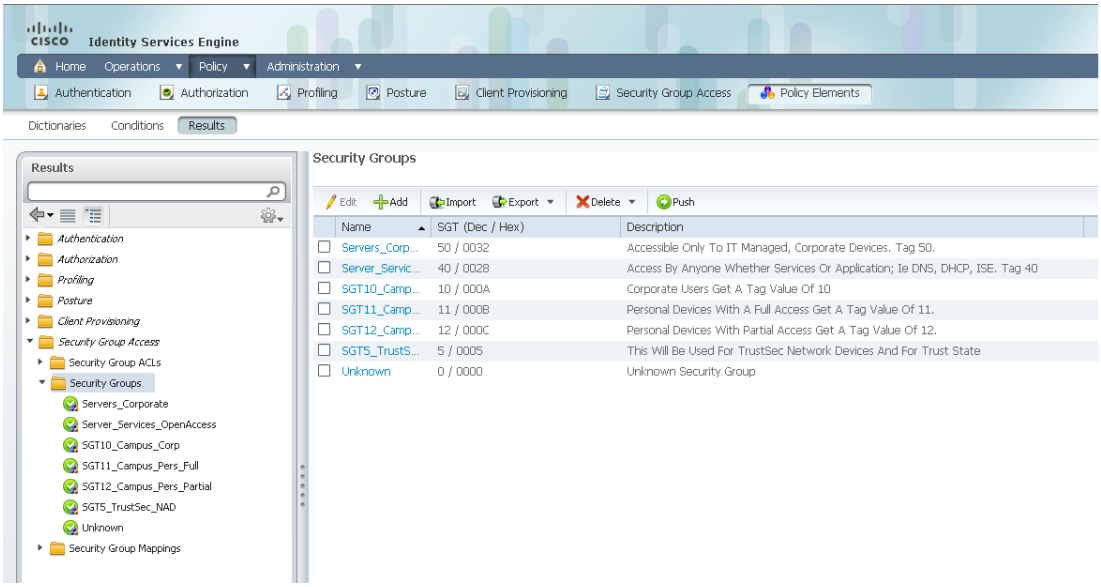
In this CVD, Table 23-1 shows the SGT Names and corresponding Tag Values used.

Table 23-1 SGT Names and Tag Values

SGT Value	SGT Name	Description
5	SGT5_TrustSec_NAD	This is used For TrustSec Network Devices And For Trust State.
10	SGT10_Campus Corp	Corporate device with full access to the network.
11	SGT11_Campus_Pers_Full	Personal device with full access to the network
12	SGT12_Campus_Pers_Partial	Personal device with restricted access to some resources on the network.
40	Servers_Services_OpenAccess	Servers in data center accessible by all devices.
50	Servers_Corporate	Corporate Servers that only Corporate devices or approved personal devices have access to.
0	Unknown	System defined/reserved representing a device (IP Address) not associated with a SGT.

These values will be defined in ISE as can be seen in Figure 23-12.

Figure 23-12 Security Groups Used in CVD



The next step for TrustSec configuration at ISE will be the definition of the actual policy to be enforced. As discussed previously, the TrustSec Policy will only be created for use as an SGACL on Catalyst and Nexus switching products. The ASA family of firewalls as of v9.0 do not support dynamic creation/update of the policies defined in ISE. Only the actual Security Group Names and Tag value are dynamically acquired from ISE. The policy used in this CVD is depicted in Figure 23-13.

Figure 23-13 TrustSec Policy to be Enforced

	SGT5	SGT10	SGT11	SGT12	SGT40	SGT50	Unknown
SGT5	✓	✓	✗	✗	✗	✗	✓
SGT10	✓	✓	✗	✗	✓	✓	✓
SGT11	✗	✗	✓	✗	✓	✓	✓
SGT12	✗	✗	✗	✓	✓	✗	✓ *
SGT40	✗	✓	✓	✓	✓	✓	✓
SGT50	✗	✓	✓	✗	✓	✓	✓
Unknown	✓ **	✓	✓	✓	✓	✓	✓

293631

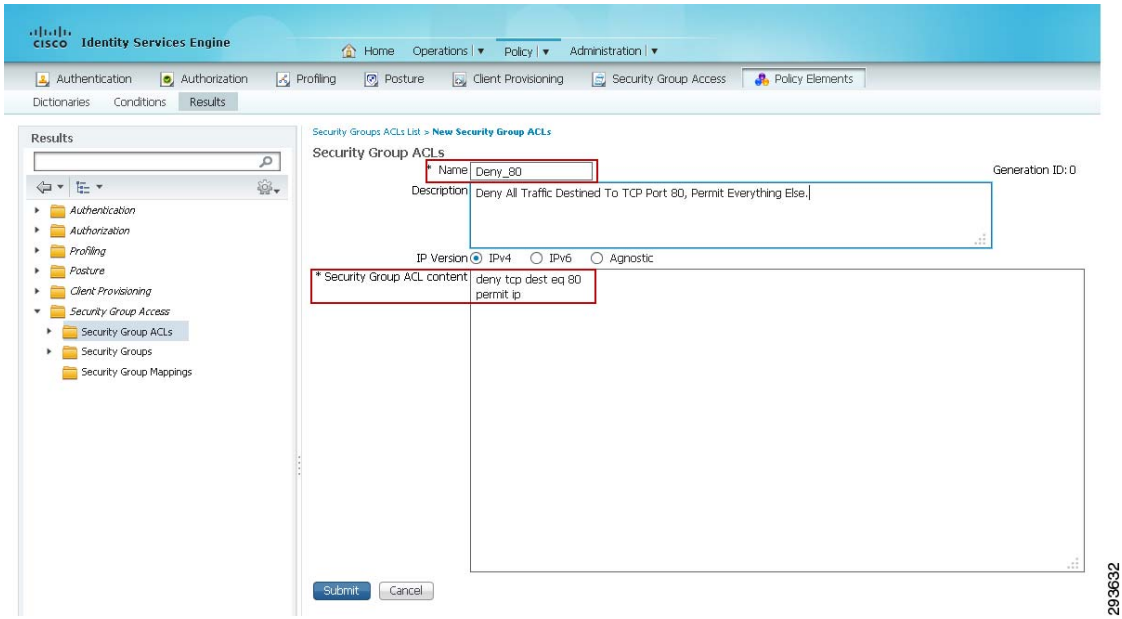
Note A policy allowing SGT12 to Unknown is defined at ISE. However, a policy will be configured locally on the Nexus 7000 aggregation switches denying SGT12 to Unknown. The reasoning for this is essentially that devices granted partial access to the network should only have access to specific data center resources.

Note In this CVD role-based access control through the use of SGT is limited to the access of data center resources by wireless users. In this setting, access via telnet, SSH, or HTTP/HTTPS to network devices is not subject to role-based access control through the use of SGT as it may be assumed that much if not all of the wired infrastructure has not been migrated to the use of Security Group Access. In the CVD, personal devices that are associated with an SGT have been specifically denied access to network devices associated with SGT 5, however users/devices not associated with an SGT, (considered "Unknown") are allowed to communicate with these network devices subject to normal authentication policies at the device.

The basic TrustSec Policy that is used permits or denies a specific source SGT to a specific destination SGT. In addition to this basic policy it is possible to create SGACLs with additional granularity restricting or permitting access to specific TCP or UDP port numbers. Although not utilized in the CVD, the procedure for creating this policy is to first create an SGACL at ISE and then when creating a specific SGT policy, adding the SGACL to the definition. The following steps illustrate the creation of an optional SGACL to then be used in an SGT Policy.

1. Go to Policy > Policy Elements > Results > Security Group ACLs and Add an ACL.
2. Enter the name of the ACL and optionally add a description as depicted in Figure 23-14.
3. Add the ACL.

Figure 23-14 SGACL Creation



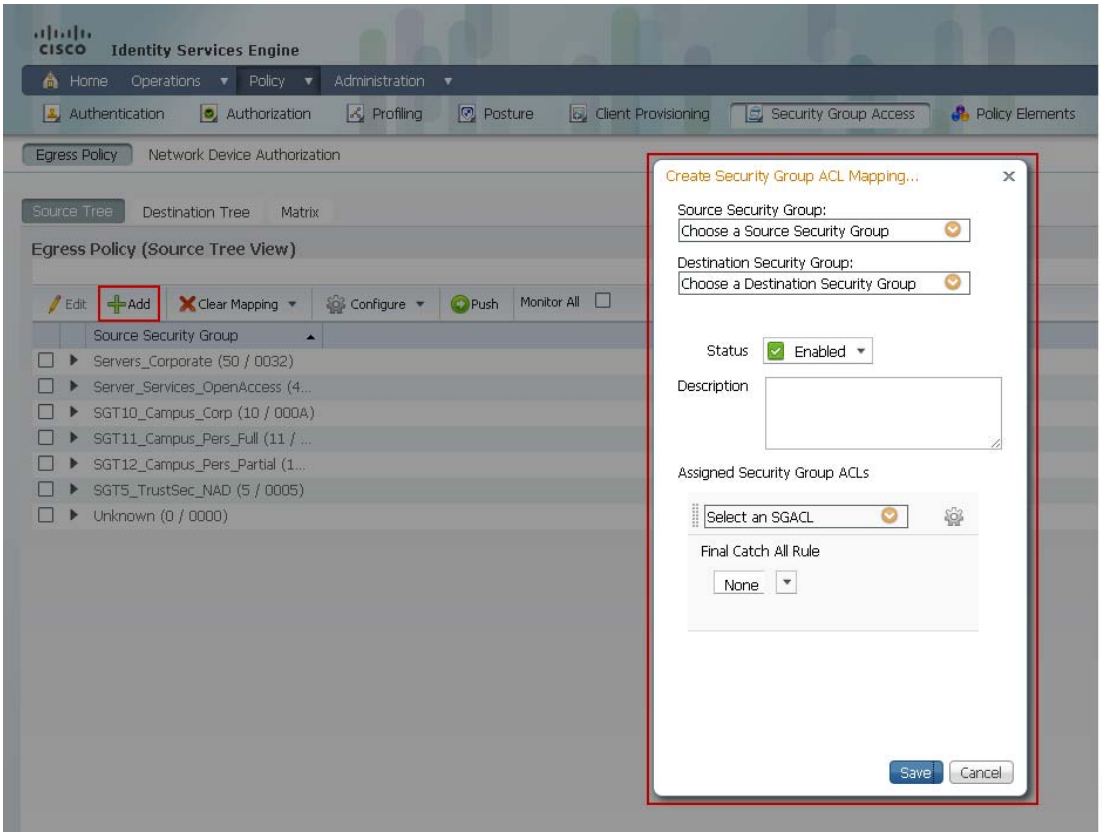
The next step is to define the SGT Policy on the ISE server by creating the appropriate entries to enforce the policy contained in Figure 23-13 earlier in this section. When creating the SGT Egress Policy definitions, it is possible to do this from three unique views:

- Source Tree
- Destination Tree
- Matrix

Note that although the three views display the policy differently the steps for creating the policy are essentially the same. Please refer to the steps outlined below and depicted in Figure 23-15 where the Source Tree view is used.

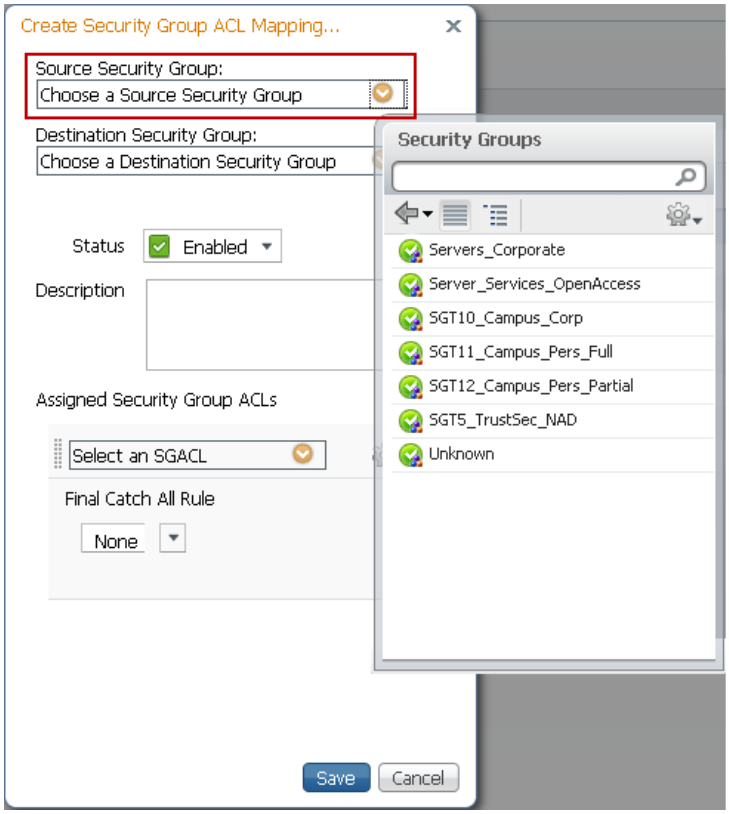
4. Go to Policy > Security Group Access > Egress Policy and select a View. This example uses the Source View.
5. Click **Add**.
6. The following window opens as shown in Figure 23-15.

Figure 23-15 TrustSec Egress Policy Creation



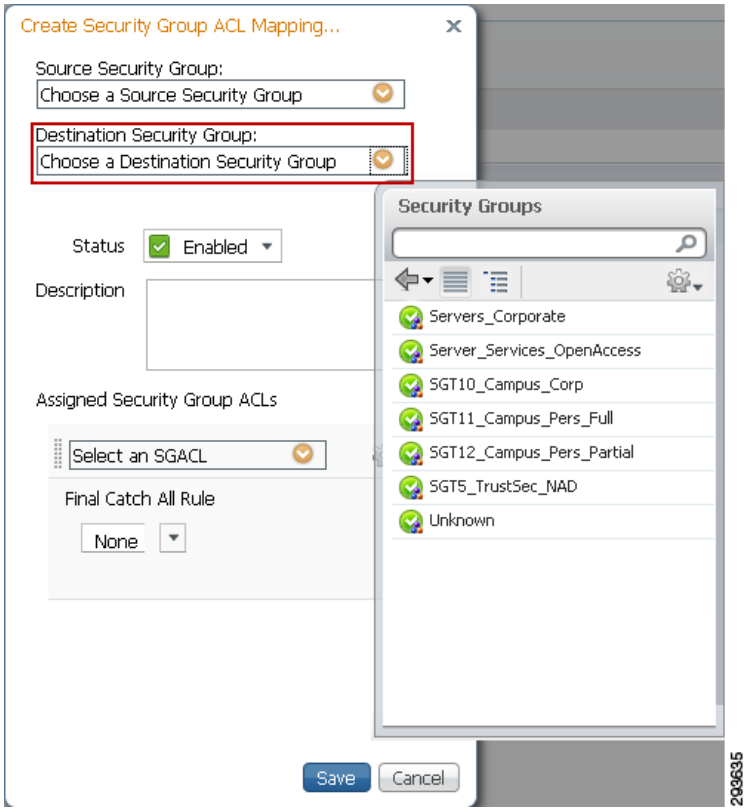
7. Next click the Source Group down arrow and select the appropriate Source Group created earlier as depicted in Figure 23-16.

Figure 23-16 TrustSec Egress Policy Creation Source SGT



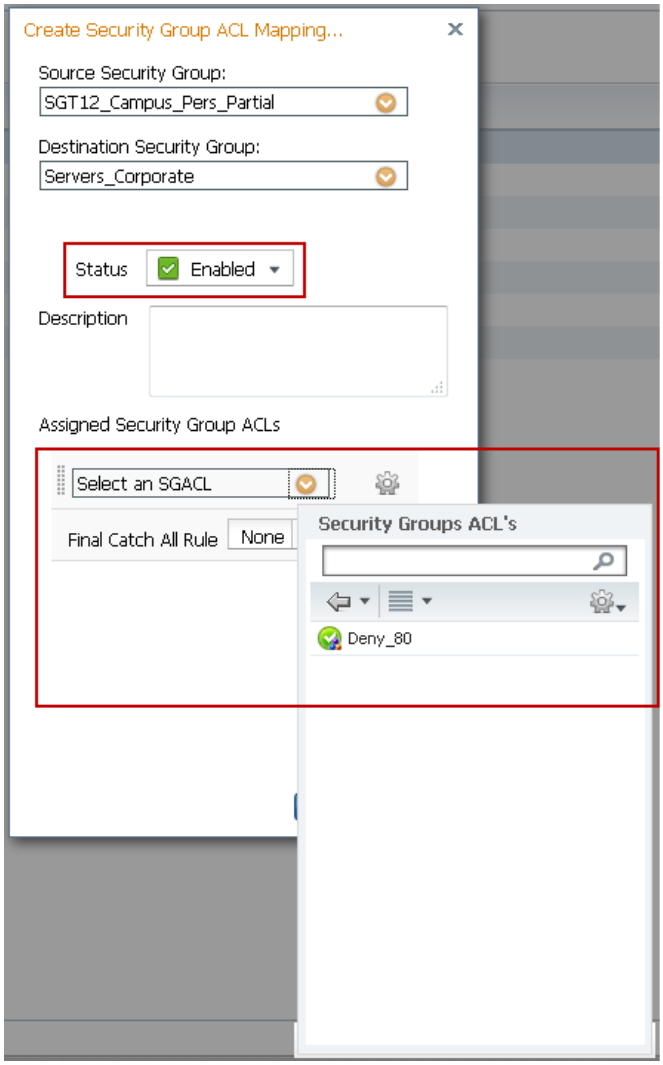
8. Click the Destination Group down arrow and select the appropriate Source Group as depicted in Figure 23-17.

Figure 23-17 TrustSec Egress Policy Creation Destination SGT



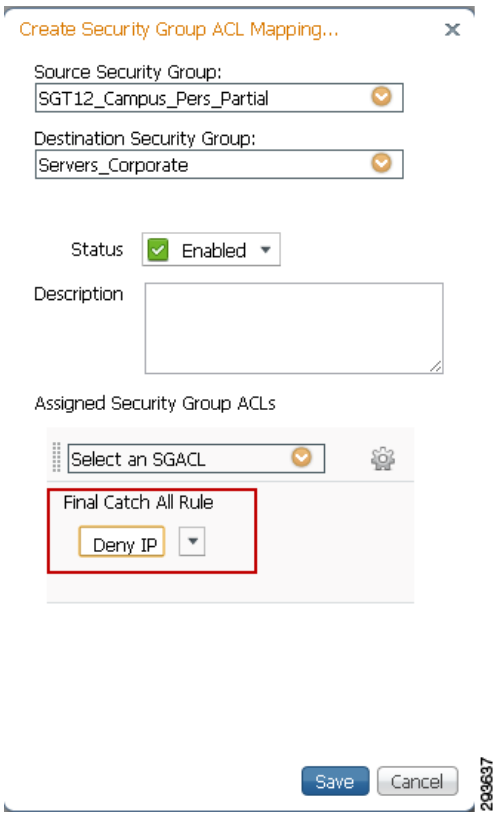
9. Ensure that the policy status is enabled and select an optional SGACL if necessary as in Figure 23-18.

Figure 23-18 Enabling TrustSec Egress Policy with SGACL



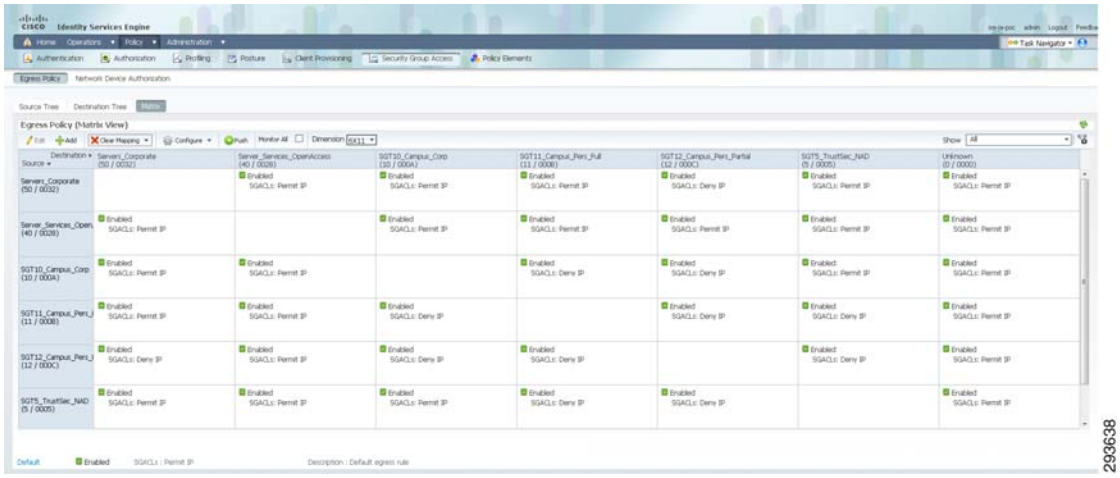
10. Finally define whether the traffic is permitted or denied and click **Save** as shown in Figure 23-19. This policy denies all traffic between a source associated with SGT12 and a destination with SGT50. Note that in Figure 23-19 an SGACL has not been defined. Also note that in creating the policy a like policy for return traffic denying SGT50 to SGT12 is not created automatically.

Figure 23-19 TrustSec Egress Policy Creation Denying Traffic



Once the addition of all egress policies is completed, the definitions can be viewed as a matrix as in Figure 23-20.

Figure 23-20 Egress Policy Matrix View



The following example output from the command **show cts role-based permissions** may be used to verify policies at a Catalyst 6500 once AAA configuration tasks have been completed later in this document.

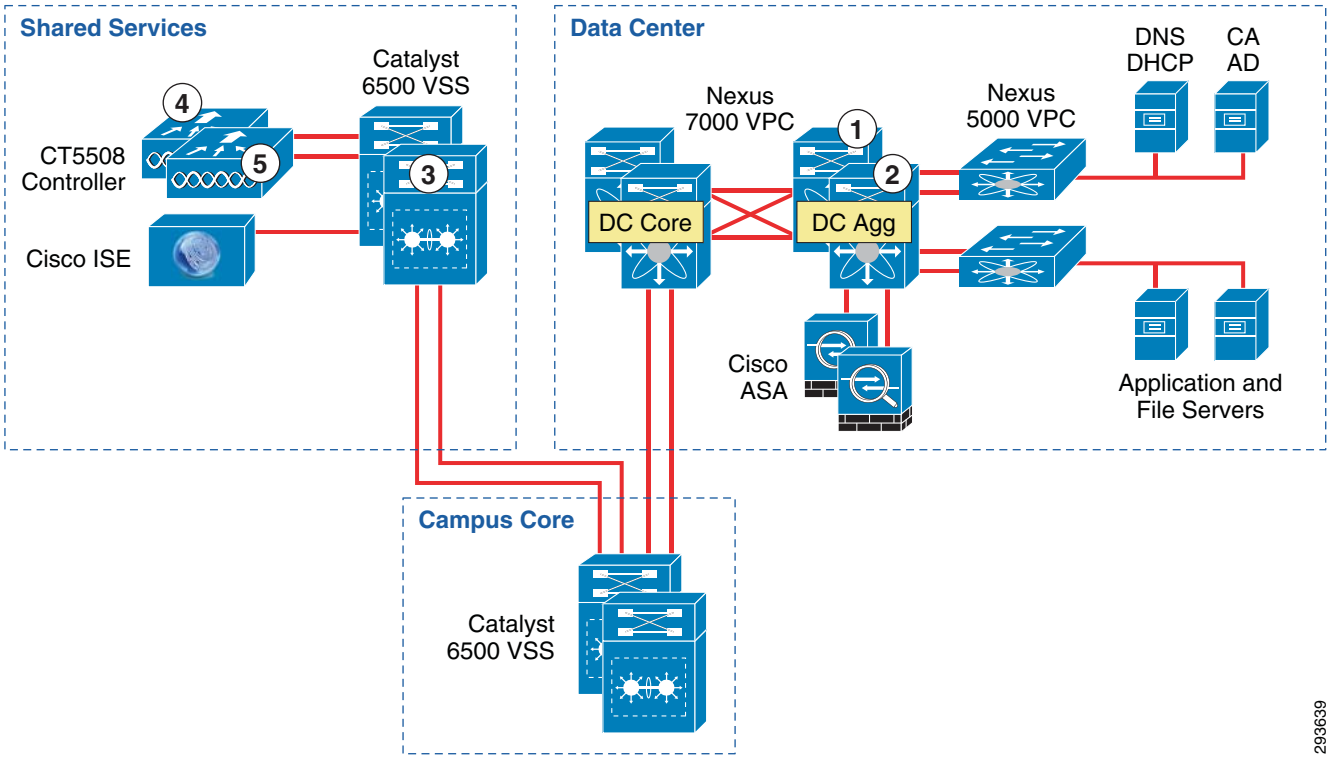
```
ua28-6500-1#sh cts role-based permissions
```

```
IPv4 Role-based permissions default:
        Permit IP-00
IPv4 Role-based permissions from group Unknown to group Unknown:
        Permit IP-00
IPv4 Role-based permissions from group 10:SGT10_Campus_Corp to group Unknown:
        Permit IP-00
IPv4 Role-based permissions from group 11:SGT11_Campus_Pers_Full to group Unknown:
        Permit IP-00
IPv4 Role-based permissions from group 12:SGT12_Campus_Pers_Partial to group Unknown:
        Permit IP-00
IPv4 Role-based permissions from group 40:Server_Services_OpenAccess to group Unknown:
        Permit IP-00
IPv4 Role-based permissions from group 50:Servers_Corporate to group Unknown:
        Permit IP-00
IPv4 Role-based permissions from group Unknown to group 40:Server_Services_OpenAccess:
        Permit IP-00
IPv4 Role-based permissions from group 10:SGT10_Campus_Corp to group
40:Server_Services_OpenAccess:
        Permit IP-00
IPv4 Role-based permissions from group 11:SGT11_Campus_Pers_Full to group
40:Server_Services_OpenAccess:
        Permit IP-00
IPv4 Role-based permissions from group 12:SGT12_Campus_Pers_Partial to group
40:Server_Services_OpenAccess:
        Permit IP-00
IPv4 Role-based permissions from group 40:Server_Services_OpenAccess to group
40:Server_Services_OpenAccess:
        Permit IP-00
IPv4 Role-based permissions from group 50:Servers_Corporate to group
40:Server_Services_OpenAccess:
        Permit IP-00
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE
```

Configuring ISE for Network Access Device Authentication

The next configuration task is to define the Network Devices that will be enforcing TrustSec Egress Policies in ISE and create the necessary AAA configuration on the devices. The purpose for doing this is to create a secure tunnel for EAP-FAST authentication of the device such that TrustSec environment data such as SG Names and associated SGT as well as TrustSec Egress Policies or SGACLs can be exchanged periodically. As discussed in the TrustSec Overview, this process is known as Network Device Admission Control or NDAC. As 802.1X will not be used to authenticate network devices across a link in order to build a trust relationship for SGT forwarding as well as MACsec encryption, it is only necessary to define those devices enforcing SGT policy as well as those wireless controllers serving as an 802.1X Authenticator to Supplicants on wireless devices attempting to Access the network. Therefore in Deployment Scenario 1 it will be necessary to create definitions for minimally five devices based on the infrastructure depicted in Figure 23-21.

Figure 23-21 Network Devices to be Defined in ISE



These devices are:

1. Data Center Nexus 7000 Aggregation Switch #1 (TrustSec policy enforcement)
2. Data Center Nexus 7000 Aggregation Switch #2 (TrustSec policy enforcement)
3. Catalyst 6500 VSS Shared Services Switch (TrustSec policy enforcement)
4. CT5508 Wireless Controller #1 (required for 802.1X wireless device access)
5. CT5508 Wireless Controller #2 (required for 802.1X wireless device access)

As SGACLs will not be enforced at the data center Nexus 7000 Core switches nor the Catalyst 6500 VSS Core, it is not mandatory for these devices to be added to the Network Device List in ISE; the network device does not need to be defined in ISE to simply forward packets with an embedded SGT. It is however recommended to define these devices to accommodate future network changes or enforcement policies as well as define a device SGT to be used by traffic sourced by these switches. The following steps must be taken to define the network devices within ISE as depicted in Figure 23-22:

1. At ISE go to Administration > Network Resources > Network Devices and click **Add**.
2. Enter the hostname of the device. This will be the same name as configured at the network device and documented later with the **cts credential** command on switches and would be the wireless controller name.
3. Enter the IP Address of the network device. This must be the address used to source all RADIUS communications from the device.
4. Change the Network Device Location or Device Type if a custom location/type has been previously defined. Within the CVD the Shared Services, Core, and Data Center switches all make use of the default setting as seen in Figure 23-22. The exceptions to this are the wireless controllers. For the controllers we have specified a custom Location known as “Campus_Controller:SGT_Enabled”. This custom location is configured under “Network Device Groups” as depicted in Figure 23-23.

The significance of the use of a custom device location lies in the ability to use that as an attribute within the Authorization Profile at ISE to determine a result. In this CVD we use the location "Campus_Controller:SGT_Enabled" to determine that an SGT Value should be handed back to the wireless controller after a user or device's successful authorization as opposed to an ACL for policy enforcement. This allows for a migrational approach to deploying infrastructure that will make use of SGT as opposed to ACLs for policy enforcement.

5. Configure the RADIUS Shared Secret. This must match that configured on the network device.
6. Click the down arrow next to SNMP Settings and complete as appropriate.
7. Click the down arrow next to Advanced TrustSec Settings. Note that it is unnecessary to complete the Advanced TrustSec Settings for CT5508 or WiSM2 wireless controllers. These controllers are not capable of supporting anything other than the creation of IP Address to SGT mapping and advertising those mappings to manually configured peers via the SXP protocol. These settings are only used for those devices supporting SGACLs and Native SG Tagging on SGT-capable hardware requiring the download of TrustSec environmental and policy data from ISE.

Figure 23-22 Network Device Generic Definition at ISE

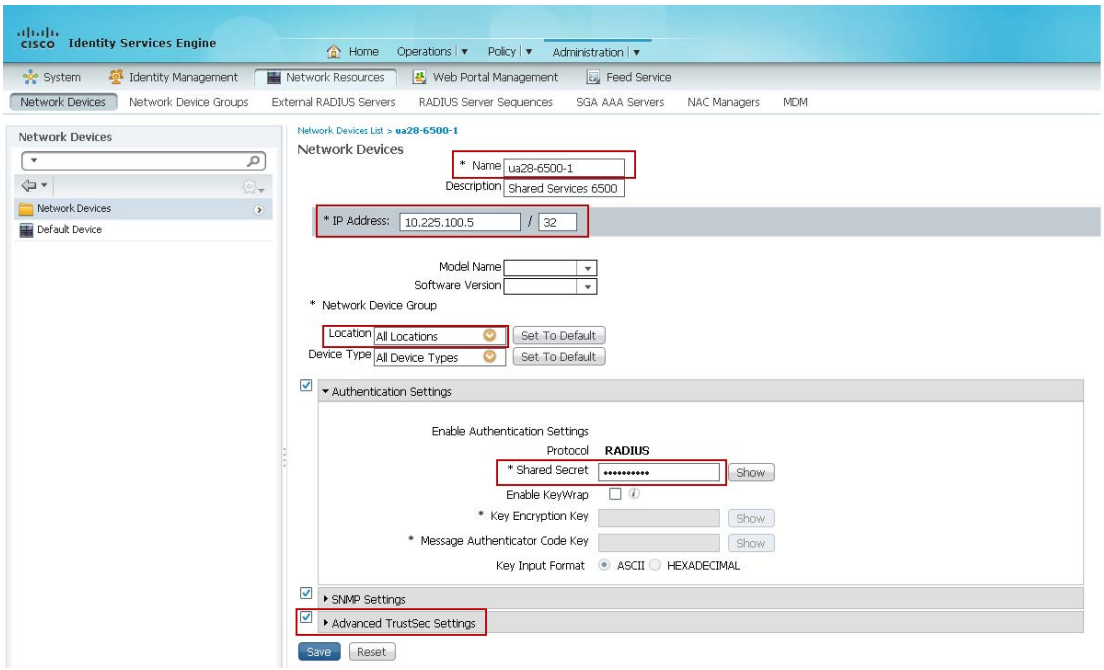
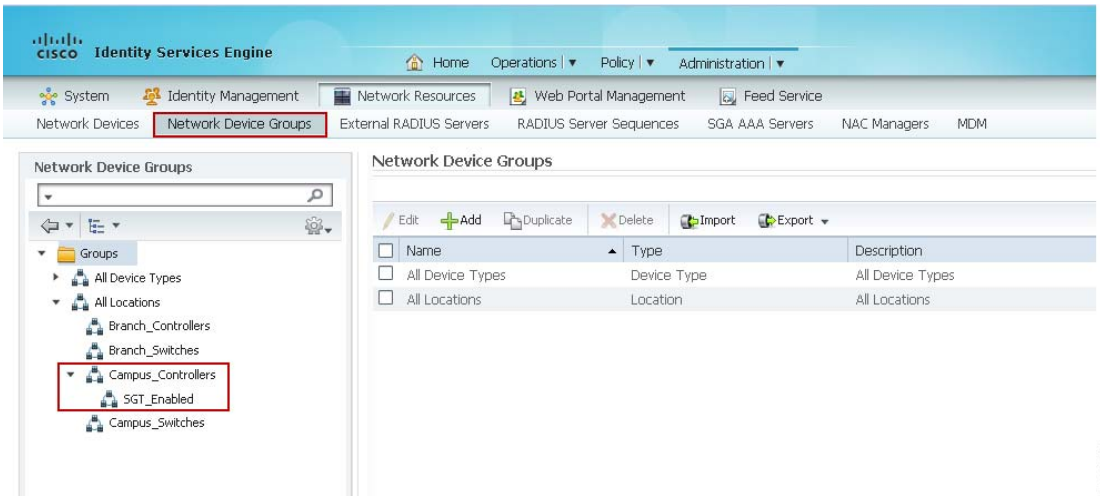
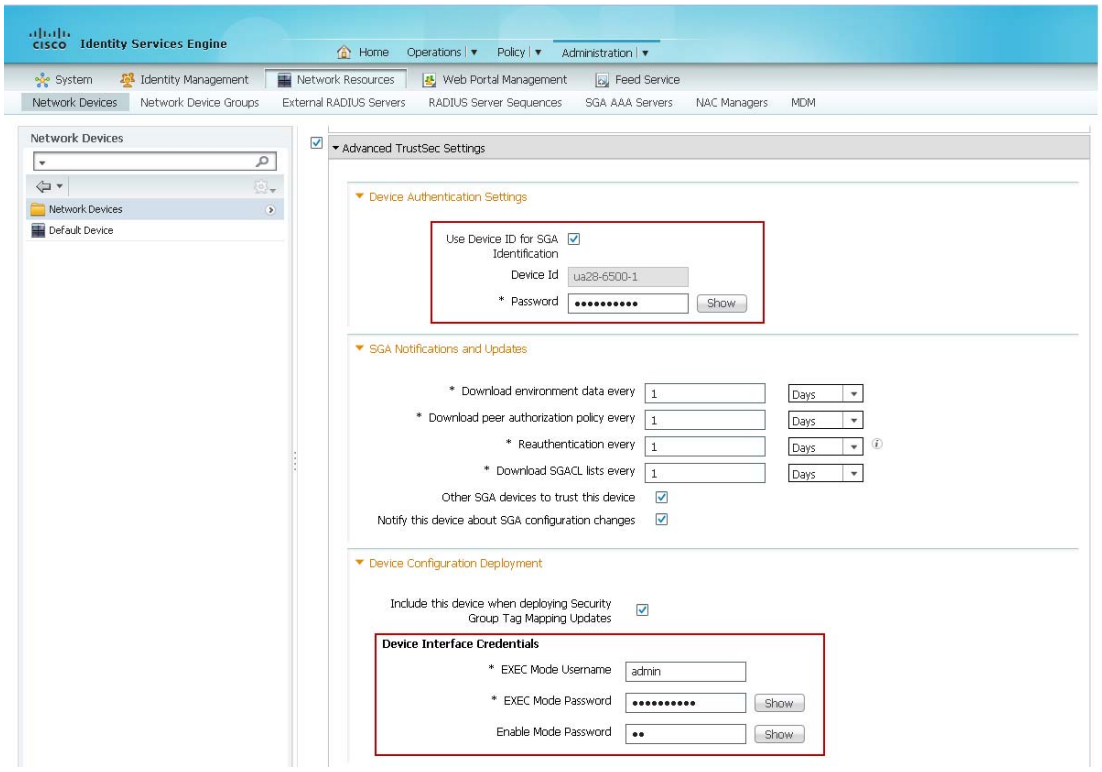


Figure 23-23 Network Device Group definition



8. Once the Advanced TrustSec Settings configuration box has been expanded as seen in Figure 23-24, click the check box next to "Use Device ID for TrustSec Identification"
9. Enter the password that will be configured later on the network device in the **cts credential** command. This can be the same as the RADIUS Shared Secret.
10. Configure the desired settings for "TrustSec Notifications and Updates". Note that these are the settings that determine the frequency of the TrustSec Environment updates to the network device. It is recommended that aggressive timers not be used here and as such these have been left at the default value for one day. Note that this is to configure the automated, periodic update of the pertinent data. In addition to these periodic updates, it is possible to manually push updates for SGT Names/Values, Network Device SGT, SGT Egress Policy (SGACL), and AAA Server List from within ISE to those network devices supporting Change of Authorization (CoA). Note that the Nexus 7000 does not support CoA at the time of this document. To support a manual push of environmental data and policy to the Nexus 7000, it is possible to do so through reissuing the **cts credential** command at the Nexus 7000 discussed later. For further information regarding these parameters and how TrustSec environment data is exchanged, refer to the ISE User Documentation and specifically the "Configuring TrustSec Settings on Switches" and "TrustSec CoA" sections of the chapter "Configuring Cisco Security Group Access Polices" located in the ISE 1.2 User Guide at: http://www.cisco.com/en/US/products/ps11640/products_user_guide_list.html.
11. Enter the credentials to access Exec Mode (if applicable) and the Enable Mode password used by ISE to access the device to manually push updated information.
12. Complete steps one through seven for every wireless controller providing 802.1X-authenticated wireless access to users and steps one through eleven for all network devices that will be enforcing TrustSec Policy requiring SGT Names and SGACL egress policies.

Figure 23-24 Network Device—Advanced TrustSec Settings

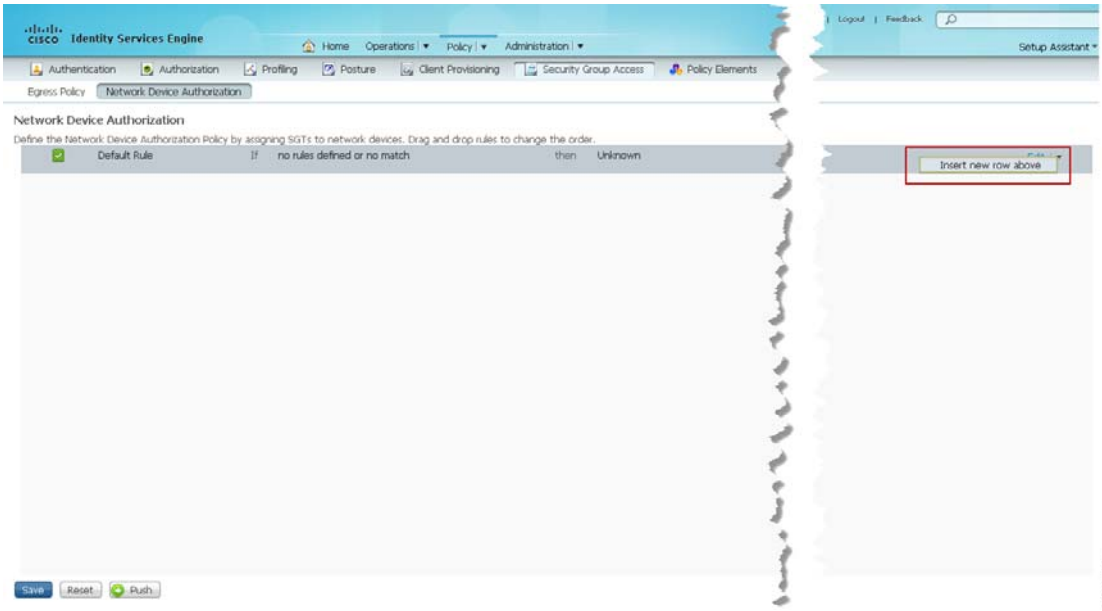


Configuring Network Devices with a Device SGT

The following outlines the configuration steps required to define a device. Once a network device is configured for with a device SGT, any traffic sourced from that device will use the defined SGT. Note that assigning a Security Group Tag to a device is purely optional. Network devices in this CVD are assigned a device SGT of 5. As granular role-based policies using SGT are defined in the network, the assignment of an SGT to network devices will provide an additional level of control over whom or what may access the network infrastructure to poll or modify these devices.

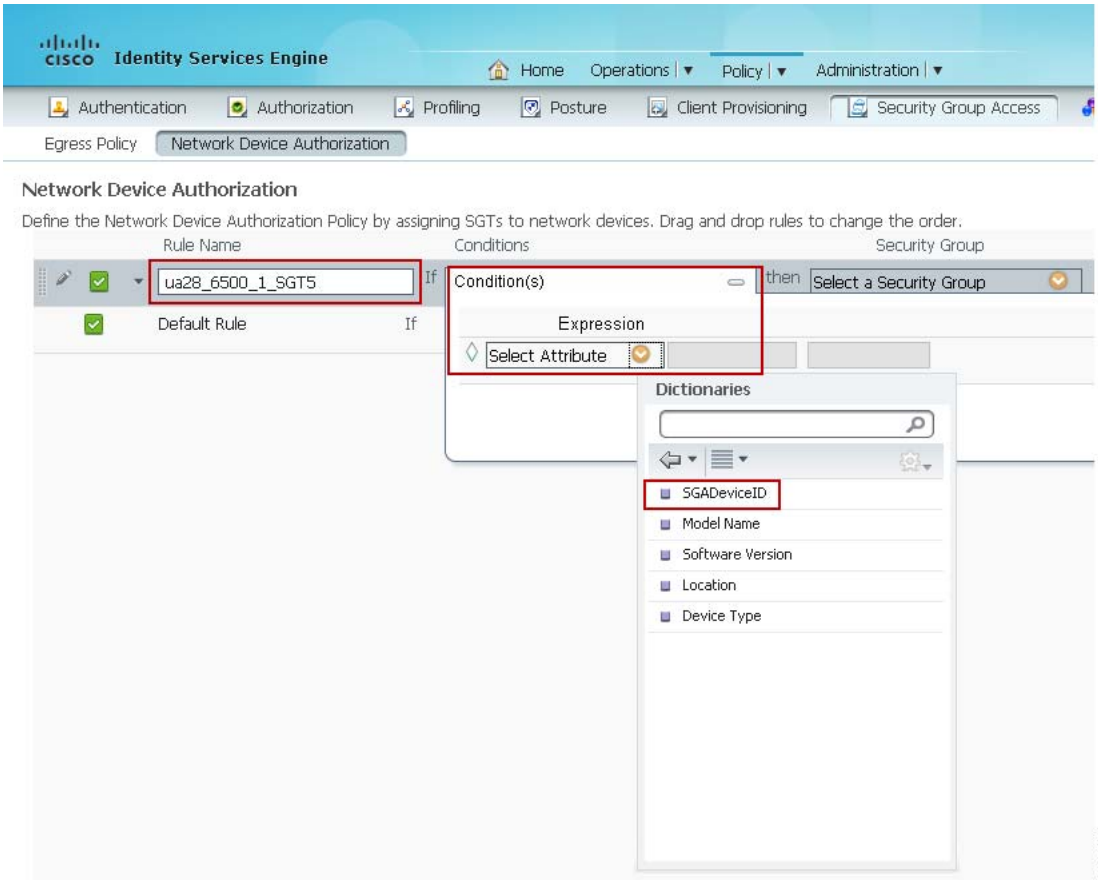
1. At ISE navigate to Policy > Network Device Authorization.
2. Click the drop down box next to edit at the top right of the screen.
3. Select “Insert new row above” as depicted in Figure 23-25.

Figure 23-25 Configuring Network Device Authorization



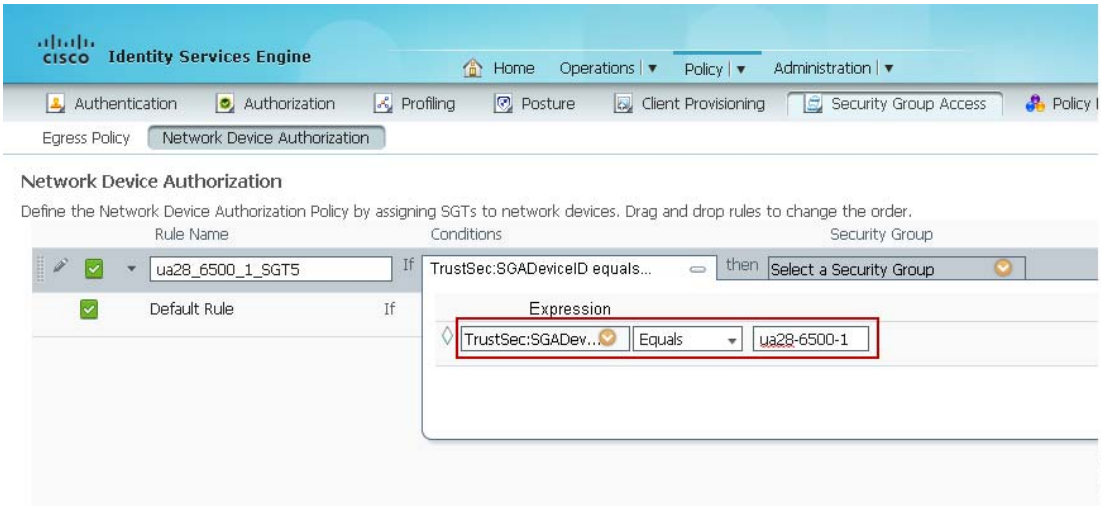
4. A new line will be inserted as seen in Figure 23-26. Enter a rule name.
5. Click the "+" symbol in conditions and a drop down box will appear.
6. Click the arrow next to "Select Attribute" and the Dictionaries drop down box will appear.

Figure 23-26 Defining Authorization Policy for the Network Device



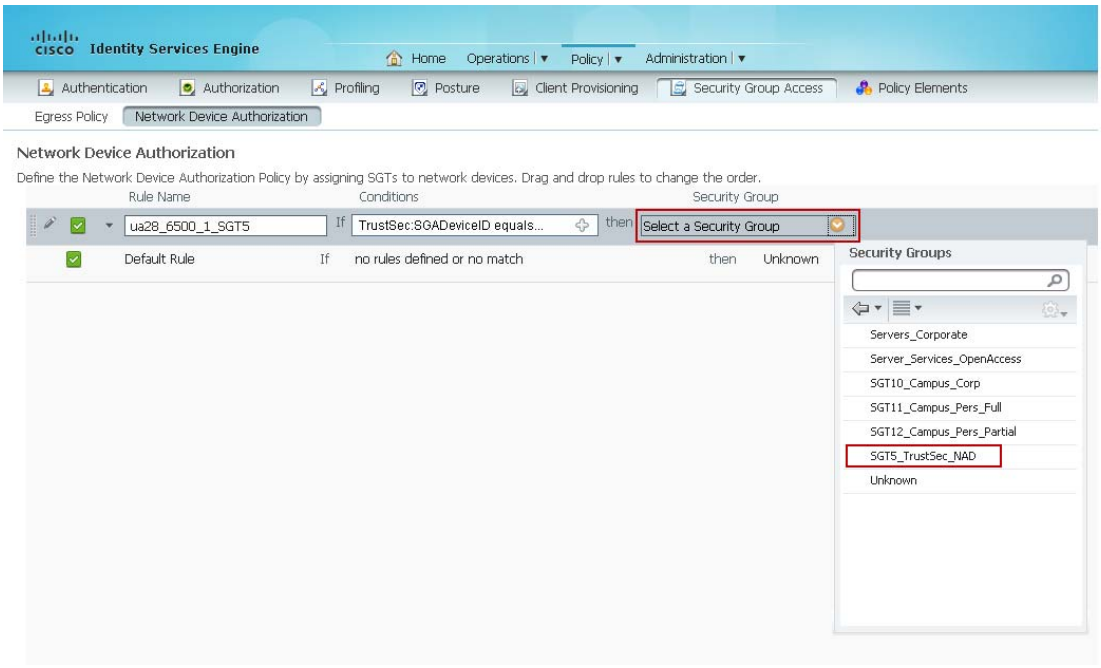
7. Select “SGADeviceID” and as can be seen in Figure 23-27 1 the Expression is populated with SGADeviceID.
8. Ensure “Equals” is displayed in the expression and enter the name of the device previously configured when defining the device as depicted in Figure 23-22.

Figure 23-27 Defining Network Device ID



9. Finally, as in Figure 23-28 define the SGT for the device by clicking the arrow next to "Select a Security Group" and select the appropriate SG Name; the CVD uses "SGT5_TrustSec_NAD".
10. Repeat steps one through nine to define all network devices; as wireless controllers cannot natively tag packets with an SGT on its interfaces this procedure is not required for them.

Figure 23-28 Assigning the SGT to a Network Device



Configuring Network Access Devices for Authentication at ISE

Configuration of the network devices for NDAC support will be outlined in this section. A section will be devoted to each of the Wireless Controllers, Catalyst 6500 VSS and the Nexus 7000 switches.

The following configuration tasks will all be completed at the network devices themselves. This configuration is critical to identify the ISE Primary server as the AAA server from which information regarding TrustSec will be exchanged. Note that for greater resilience, multiple ISE Policy Service Nodes can be listed and will be tried in succession. As previously mentioned, it is only necessary to configure those devices that will provide access for the wireless users and the network devices that will actually enforce TrustSec policies. It is completely optional whether or not this needs to be configured on other devices in the path between the enforcement points.

RADIUS Server Configuration on the Wireless Controller

The configuration of RADIUS server information is required in order for the controller to act as a RADIUS Authenticator for 802.1X-based authentication of wireless clients accessing the network. More than likely these steps have already been completed as this is a basic requirement for securing wireless access regardless of the desire to use ACLs or Security Group Tags.

1. Access the wireless controller either through the local UI or Prime. In Figure 23-29 the controller's GUI is depicted.
2. Go to Security > AAA > Radius > Authentication
3. In the top right of the screen, if the ISE server has not been configured yet, click **New**.
4. A new window will open as seen in Figure 23-30.

Figure 23-29 Wireless Controller RADIUS Configuration

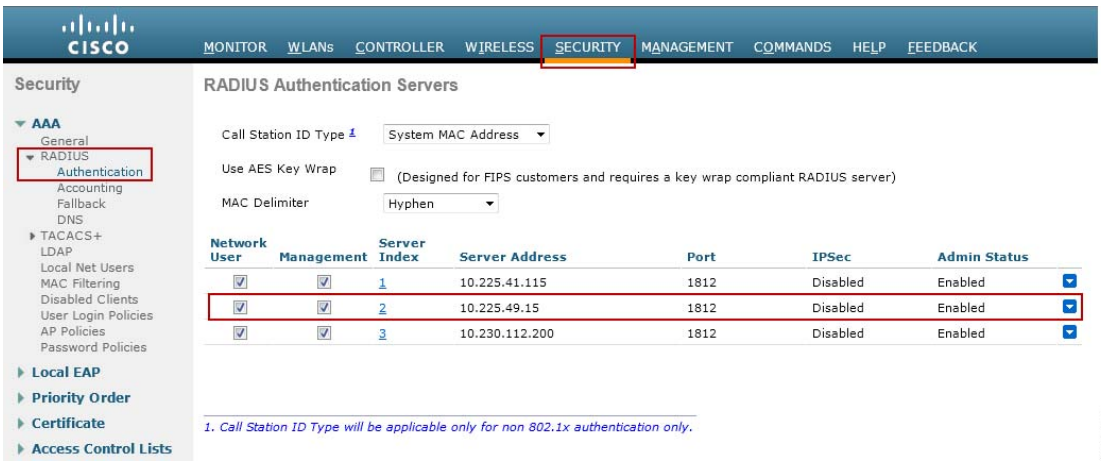
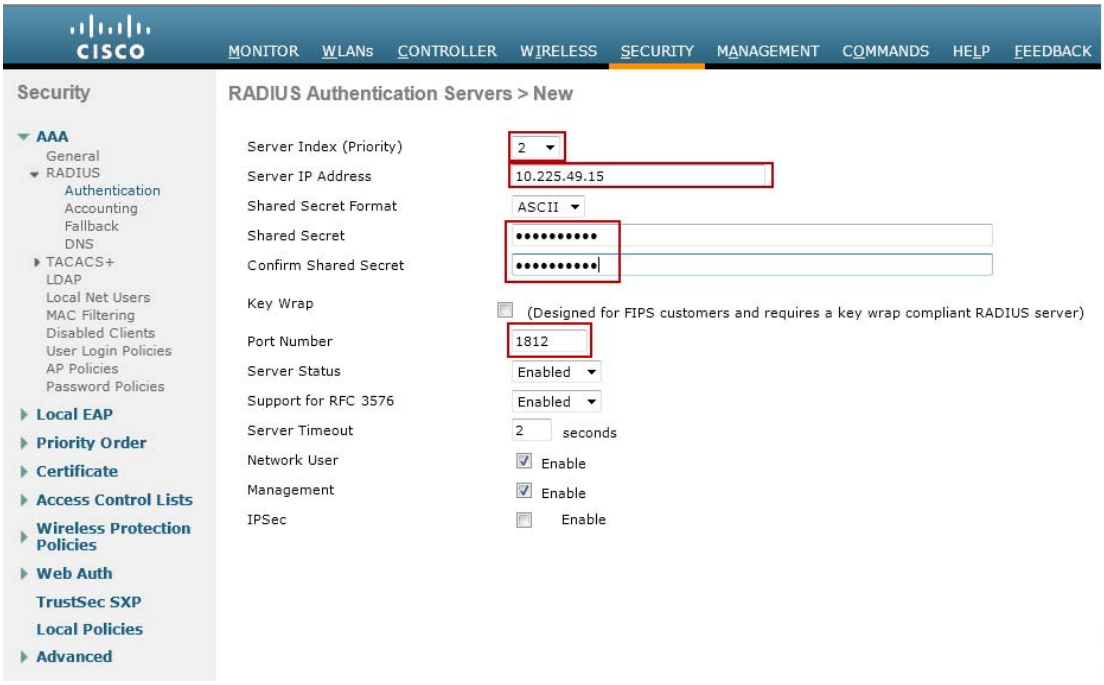


Figure 23-30 Wireless Controller RADIUS Server Configuration



5. Use the drop down to enter the correct priority; lower number is higher priority.
6. Enter the ISE server's IP address.
7. Enter the Shared Secret which must match that configured for the wireless controller as defined in the Network Device List in ISE.
8. Enter the correct RADIUS Authentication UDP port number.
9. Click **Apply**.
10. Configure the controller's RADIUS Accounting information as the Authentication information above. This can be accessed by following Security > AAA > Radius > Accounting.

Configuration of the wireless controller RADIUS server configuration is complete. Repeat these steps if additional controllers need to be configured.

RADIUS Server Configuration on the Catalyst 6500

The following steps outline those tasks necessary to configure ISE as a RADIUS server at the Catalyst 6500. As discussed, this is to establish a secure connection with ISE for the exchange of TrustSec Environment Data and Policies. These steps should be performed on any Catalyst 6500 that will enforce policies based on Security Group Tags. For the infrastructure depicted in this CVD, configuration must be completed on the Catalyst 6500 VSS in Shared Services to which the wireless controllers are attached. As the Catalyst 6500 VSS serving as the Core of the network is merely forwarding tagged packets sourced from either the wireless users or servers themselves and not enforcing any policies, there is no requirement to configure it within ISE. The following provides the configuration commands required on the Catalyst 6500.

```
cts credentials id device-id password <password>/TrustSec credentials for use with ISE;
device ID and password must be the same as at ISE.
aaa new-model/Enables AAA
aaa group server radius <ise>/Creates a AAA Server Group ISE
```

```
 server 10.225.49.15 auth-port 1812 acct-port 1813/Defines 10.225.49.15 as a member of
Group ISE.
aaa authentication dot1x default group radius/Specifies the 802.1X port-based
authentication method as RADIUS.
aaa server radius dynamic-author/Enable CoA on the 6500 to enable updates for SG
Names/Tags, Environment Data, and RBACL
 client 10.225.49.15 server-key/Identifies ISE as AAA server initiating CoA
aaa authorization network <ise> group radius/Configures the switch to use RADIUS
authorization for all network-related service requests using server group <ise>.
aaa accounting dot1x default start-stop group radius/Enables 802.1X accounting using
RADIUS
cts authorization list <ise>/Specifies a Cisco TrustSec AAA server group
ip radius source-interface Loopback0/Matches the IP Address of the device configured in
ISE; uses the IP Address of Lo0 to source all RADIUS.
radius-server host 10.225.49.15 auth-port 1812 acct-port 1813 pac key <secret>/Specifies
the RADIUS authentication server, shared secret must be same as configured for RADIUS
secret at ISE.
radius-server vsa send authentication platform /Configures the switch to recognize and use
vendor-specific attributes (VSAs) in RADIUS Access-Requests generated by the switch during
the authentication phase
dot1x system-auth-control/Globally enables 802.1X port-based authentication
```

RADIUS Server Configuration on the Nexus 7000

The following steps outline those tasks necessary to configure ISE as a RADIUS server at the Nexus 7000. As discussed, this is to establish a secure connection with ISE for the exchange of TrustSec Environment Data and Policies. These steps should be performed on any Nexus 7000 that will enforce policies based on Security Group Tags. For the infrastructure depicted in this CVD, configuration must be completed on the Nexus 7000 Aggregation Switch in the data center. As the Nexus 7000 serving as the core of the data center is merely forwarding tagged packets sourced from either the wireless users or servers themselves and not enforcing any policies, there is no requirement to configure it within ISE.

On the Nexus 7000, enabling TrustSec is slightly different than on the Catalyst 6500 such that the 802.1X and TrustSec Features must be enabled on the Nexus platform through the following commands:

```
feature dot1x/Enable dot1x support.
feature cts/Enable cts (TrustSec) support
```

Once the features have been enabled the AAA servers and TrustSec Device credentials must be enabled through the following commands:

```
cts role-based enforcement/Enables SGACL enforcement on Nexus 7000
cts role-based counters enable/Enable role-based access control list (SGACL) counters
cts device-id device-id password password/TrustSec credentials for use with ISE; device ID
and password much be the same as at ISE.
radius-server host 10.225.49.15 key <secret> pac/Specifies the RADIUS authentication
server, shared secret must be same as configured for RADIUS secret at ISE.
aaa group server radius <ise>/Creates a AAA Server Group ISE
  server 10.225.49.15/Defines 10.225.49.15 as a member of Group ISE
aaa authentication dot1x default group <ise>/Specifies the 802.1X port-based
authentication method as RADIUS.
aaa accounting dot1x default group <ise>/Enables 802.1X accounting using group ISE
aaa authorization cts default group <ise>/To configure the default authentication,
authorization, and accounting (AAA) RADIUS server groups for Cisco TrustSec authorization
ip radius source-interface loopback0/Matches the IP Address of the device configured in
ISE; uses the IP Address of Lo0 to source all RADIUS.
```

Catalyst 6500 Platform Specific Considerations

Prior to discussing aspects of the TrustSec infrastructure configuration, it is necessary to highlight one aspect regarding the Catalyst 6500 when configured for Cisco TrustSec and specifically SG Tagging capability. As previously stated, the SUP2T and WS-X69xx series of linecards are the only SGT-Capable Supervisor and linecard available today for processing and imposition/removal of Security Group Tags in the Catalyst 6500. Specialized ASICs are required in order to forward tagged packets and encrypt the frames using 802.1ae MACsec with 10GE wirespeed performance. This functionality involves changes in the internal forwarding process. In order to support earlier linecards that do not have these newer ASICs (i.e. WS-X68xx, WS-X67xx, and WS-X61xx) in the same chassis when TrustSec has been enabled, two new operating modes have been developed called Ingress and Egress Reflector mode. Ingress Reflector mode is intended for use when the Catalyst 6500 is providing network access; it does not support linecards with a DFC installed. Egress reflector mode provides compatibility with legacy line cards by using the SUP2T forwarding engine's built-in packet replication ASICs to initiate a second packet forwarding decision. This second forwarding decision is used to impose the Cisco TrustSec SGT information on packets egressing the system from an SGT-capable interface. For purposes of the BYOD v2.5 CVD, Egress Reflector Mode will be used on the Catalyst 6500 VSS switches containing legacy linecards.

To enable the Egress Reflector Mode on the Catalyst 6500, it is necessary to perform a reload of the system. It is recommended for obvious reasons that this be performed off hours and that necessary precautions have been put in place to ensure that traffic can be forwarded around the reloading system if required. In the case of Catalyst 6500s configured for VSS, it is recommended that the entire system be reloaded through the **reload** command. The command to enable this mode of operation on the Catalyst 6500 is:

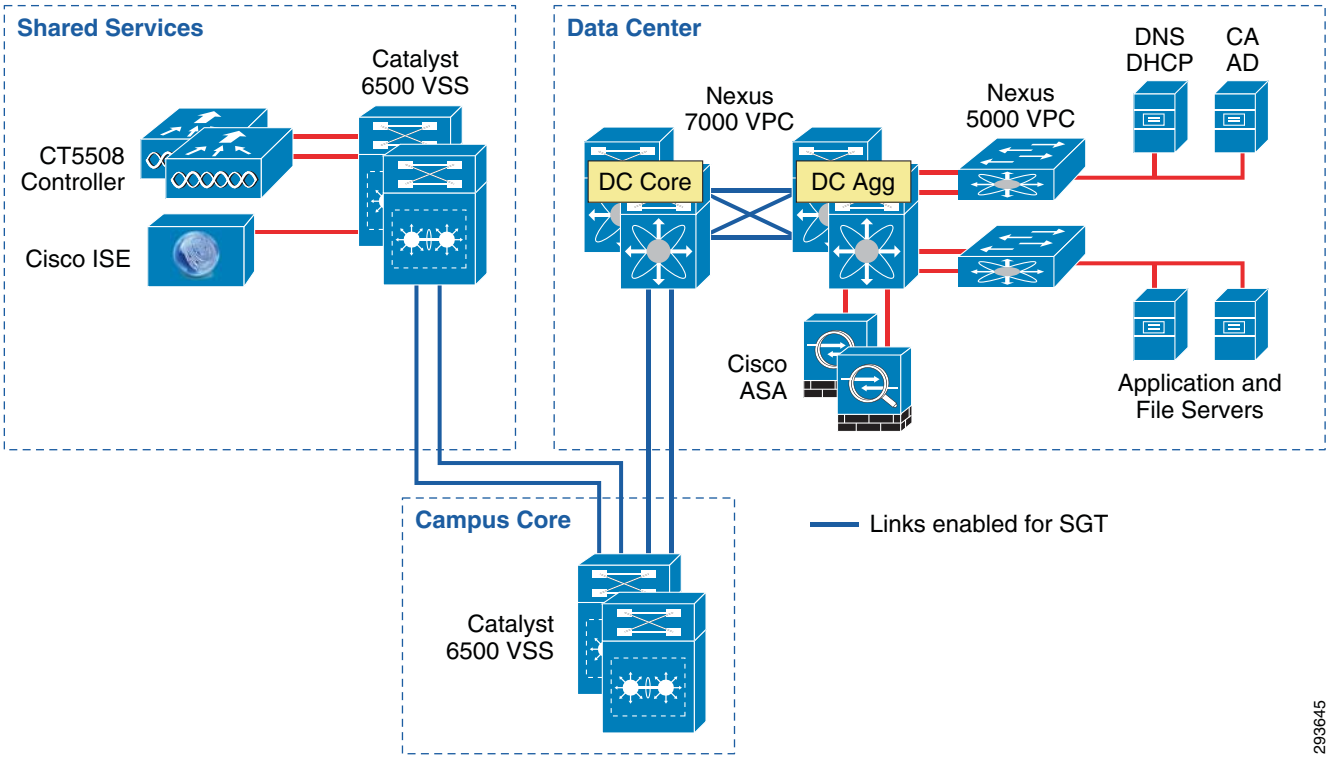
```
platform cts egress
```

For more detailed information, refer to:
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11-658388.html.

Configuring Switching Infrastructure to Support TrustSec with 802.1ae MACsec Encryption

Now that the configuration of ISE and the Network Access Devices' ability to communicate with ISE as the AAA server via RADIUS have been completed, the following steps are required for the configuration of the 10GE links in the switching infrastructure to support TrustSec, specifically SGT insertion, removal, and forwarding as well as the encryption of those links using 802.1ae MACsec. In the BYOD CVD infrastructure depicted in Figure 23-31, it is necessary to configure the blue, 10GE links in the figure for TrustSec using the Catalyst and Nexus switch **cts** command.

Figure 23-31 TrustSec Configuration of 10GE Links



There are two methods, 802.1X Mode and Manual Mode, for configuring 10GE interfaces to support Security Group Access (TrustSec), enabling the forwarding and policy enforcement of frames with an embedded Security Group Tag. CTS 802.1X Mode uses a Pairwise Master Key (PMK) derived through the authentication phase between the network device and ISE for link encryption whereas with CTS Manual Mode, as its name implies, the PMK is manually configured. In this CVD, the Manual Mode of configuration is used.

Common to both of these methods is first the ability to employ MACsec (802.1ae) which provides encryption, a message integrity check, and data-path replay protection for links between adjacent network devices thereby protecting the CMD field and the SGT value it contains. Second, as previously discussed is the configuration for 802.1X authentication of the network devices that will enforce policies based on the SGT with ISE acting as an Authentication Server.

Also common to both CTS Manual and 802.1X Mode is the use of the Security Association Protocol (SAP) which is an encryption key derivation and exchange protocol based on a draft version of the 802.11i IEEE protocol. In a TrustSec configuration, the keys are used for MACsec link-to-link encryption between two interfaces.

In CTS Manual Mode, the Pairwise Master Key (PMK) will be manually configured on each of the two interconnecting 10GE interfaces with the **sap pmk** command. The PMK is a hexadecimal value with an even number of characters and a maximum length of 32 characters. It is not necessary to specify all 32 characters as the value provided will be padded with zeroes. This value MUST be the same on both sides of the link between the two switches. The Catalyst 6500s will pad the PMK provided with leading zeroes by default however the Nexus 7000 will pad the PMK with trailing zeroes by default. It is possible however, at the Nexus 7000 command line, to alter this behavior which is demonstrated below.

Note If configuring 10GE links that have not been defined as a member of a port channel, you may proceed to the commands listed below. If however, these 10GE links are presently active within a port channel, it will be necessary to first remove them from that port channel as otherwise issuing the **cts** command will fail, having not successfully passed a port channel consistency check. Once removed from the Port Channel, TrustSec, through the **cts** command can now be configured.
When migrating port channels to enable TrustSec/MACsec, one possible migration option is to remove the links one at a time, configure TrustSec and MACsec as applicable on both sides of the link and ensure that the link comes back up. Repeat this on each port channel member until the last one is reached. Remove the last remaining member from the port channel. Once the port channel has no remaining links, those configured for TrustSec can then be added sequentially. This type of migrational procedure can be used on both the Catalyst and Nexus switching platforms.

Additionally, for manual mode, trust must be established through the configuration of a policy on the interface. This is only necessary on the Catalyst 6500 as the Nexus 7000 by default trusts the tags forwarded by its peer. If this policy is not created on the Catalyst 6500 interfaces however, the tags embedded in frames sourced from its peer will be untrusted and subsequently removed, and the frame forwarded un-tagged as the peer interface default is “untrusted”. This policy is established through the use of the **policy static sgt id trusted** command.

Note The policy static command is used to specify that the peer should either be trusted or untrusted. For example, when using the command **policy static sgt 5** on an interface, the peer is considered to be untrusted and all traffic that arrives at the interface, tagged or untagged, will be marked with an SGT of 5. On the other hand, the command **policy static sgt 5 trusted** infers that all traffic arriving on an interface from its peer, tagged or un-tagged will be trusted and forwarded with the embedded tag, or in the case of an un-tagged frame will be forwarded without a tag. The use of the SGT value specified in the command has been deprecated and has no effect on how the frame is forwarded relative to re-marking or adding a tag to the frame. This is a cosmetic issue that is under examination but at the time of writing this CVD, an SGT value must be specified in the command. In this CVD it is recommended to use the peer’s device SGT if one has been configured. If network devices have not been assigned an SGT, it would be recommended to use the same SGT as corporate devices or SGT 10 as documented in this CVD. In this CVD all links are configured as trusted.

Catalyst 6500 Commands

At the Ten Gigabit Ethernet interface configuration prompt issue the following.

```
(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
(config-if-cts-manual)# sap pmk ABC123 mode-list gcm-encrypt gmac null/Manually specify
the Pairwise Master Key (PMK) and the Security Association Protocol (SAP) authentication
and encryption modes to negotiate MACsec link encryption between two interfaces.
(config-if-cts-manual)# policy static sgt id trusted/Establishes the trust state of the
peer interface. See note above.
```

When configuring the Security Association Protocol as the keying mechanism for use with MACsec several options are available for authentication and encryption on the link. Table 23-2 provides a summary of each option. When mode-list is specified as above, the devices on either side of the link will negotiate via the SAP protocol, the method supported. As defined above, gcm-encrypt will be tried first and so on.

Table 23-2 Security Association Protocol Options

Mode	Description
gcm-encrypt	Authentication and encryption
gmac	Authentication, no encryption
no-encap[1]	No encapsulation[1]
null	Encapsulation (SGT), no authentication or encryption

1. If the interface is not capable of SGT insertion or data link encryption, **no-encap** is the default and the only available SAP operating mode.

Nexus 7000 Commands

At the Ten Gigabit Ethernet interface configuration prompt issue the following. If this configuration is being applied to a link interconnecting a Catalyst 6500 and a Nexus 7000 note that it will be necessary to alter the <sap pmk> command to change the default method of padding the PMK should less that 32 characters be specified. As noted earlier in the Catalyst 6500 interface configuration, it is unnecessary to specify the trust state of the peer through the use of interface command <policy static> as the default is trusted on the Nexus 7000.

```
(config-if)#cts manual/Manually enable an interface for Cisco TrustSec Security (CTS)
```

Link connecting to another Nexus switch:

```
(config-if-cts-manual)# sap pmk ABC123 mode-list gcm-encrypt gmac null/Manually specify
the Pairwise Master Key (PMK) and the Security Association Protocol (SAP) authentication
and encryption modes to negotiate MACsec link encryption between two interfaces.
```

Link connecting to a Catalyst switch:

```
(config-if-cts-manual)# sap pmk ABC123 left-zero-padded mode-list gcm-encrypt gmac null
/Manually specify the Pairwise Master Key (PMK) and the Security Association Protocol
(SAP) authentication and encryption modes to negotiate MACsec link encryption between two
interfaces.
```

The SAP modes of operation are identical to those of the Catalyst 6500 detailed earlier.

Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers

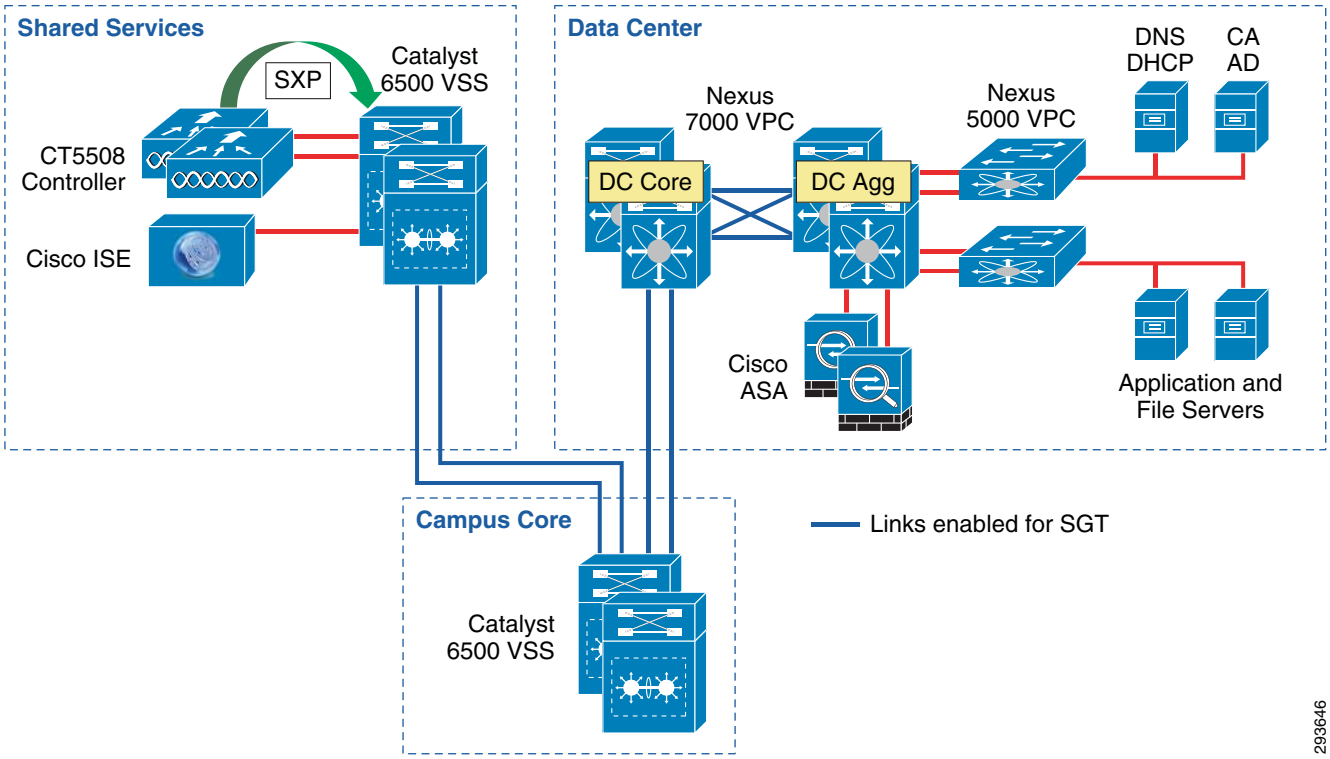
Campus wireless users accessing the network upon successfully matching an authorization profile at the Identity Services Engine will be associated with an SGT. Upon successful authentication and subsequent authorization to the network the Identity Services Engine will pass the appropriate SGT value to the wireless controller through a RADIUS AV. This SGT value is associated with the IP Address of the wireless user obtained through the 802.1X authentication and an IP/SGT mapping/mapping created at the wireless controller.

Wireless controllers such as the CT5508, used in this guide, and the WiSM2 however, do not support the tagging of packets sourced from these wireless users out of the controller through both physical and internal interfaces, in the case of the WiSM2. As such, the Security Group Tag Exchange Protocol will be used to advertise these SGT mappings to the Shared Services Catalyst 6500VSS switch which will impose the appropriate tag upon egress from the switch.

SXP is configured on a device by identifying its peer's IP Address and specifying a password for use in authenticating each side of the connection. SXP supports two modes which can be used either exclusively or combined. The first mode is that of the “Speaker” which as the name suggests advertises IP/SGT mappings. The other mode is “Listener” which also as its name suggests, listens for the Speaker's advertisements. It is possible for a device to be both a “Speaker” and a “Listener”. The CT5508 and WiSM2 obviously only support “Speaker” mode as they do not support SGACLs or SGT Tagging and hence a “Listener” mode is not applicable.

Both sides of the SXP connection must be configured and with Deployment Scenario 1, includes configuration at both of the CT5508s as well as the Shared Services Catalyst 6500 VSS switch. Please refer to Figure 23-32.

Figure 23-32 SXP Peering



Wireless Controller Configuration

Access the wireless controller via a web UI and follow the following steps:

1. Navigate to Security > TrustSec SXP.

   The screen in Figure 23-33 appears.

2. Click the drop down arrow for “SXP State” and select Enabled.
3. Set the “default Password”. This must match that configured on its peer.
4. Click “**New**” (top right).
5. Fill in the IP Address (typically Loopback if possible) of the SXP Peer or the Shared Services Catalyst 6500VSS switch.
6. Click **Apply**.

Once successfully configured, the screen in Figure 23-34 should be presented upon accessing Security > TrustSec SXP. The "Connection Status" will indicate "Off" until the other device is configured.

Figure 23-33 SXP Configuration at Wireless Controller

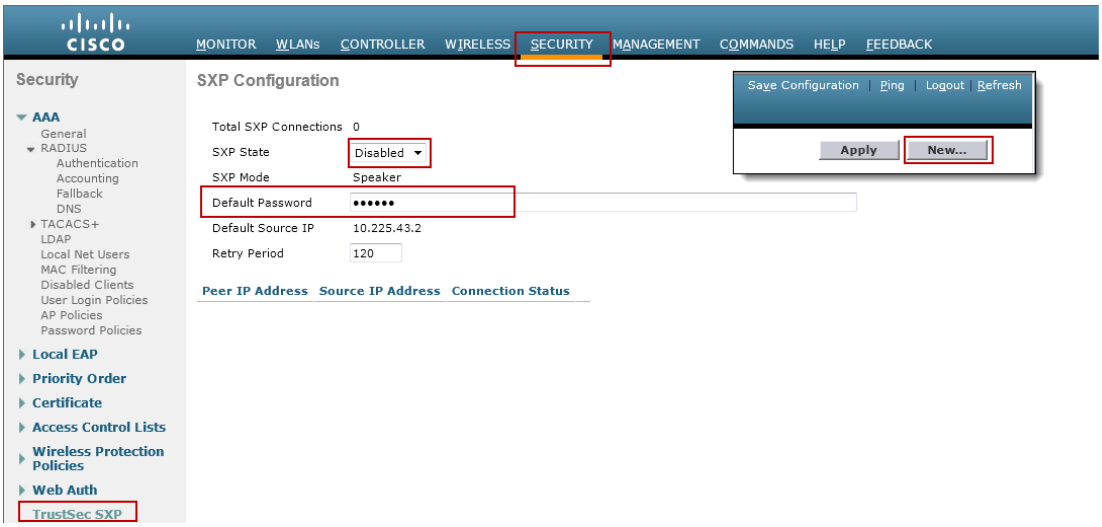
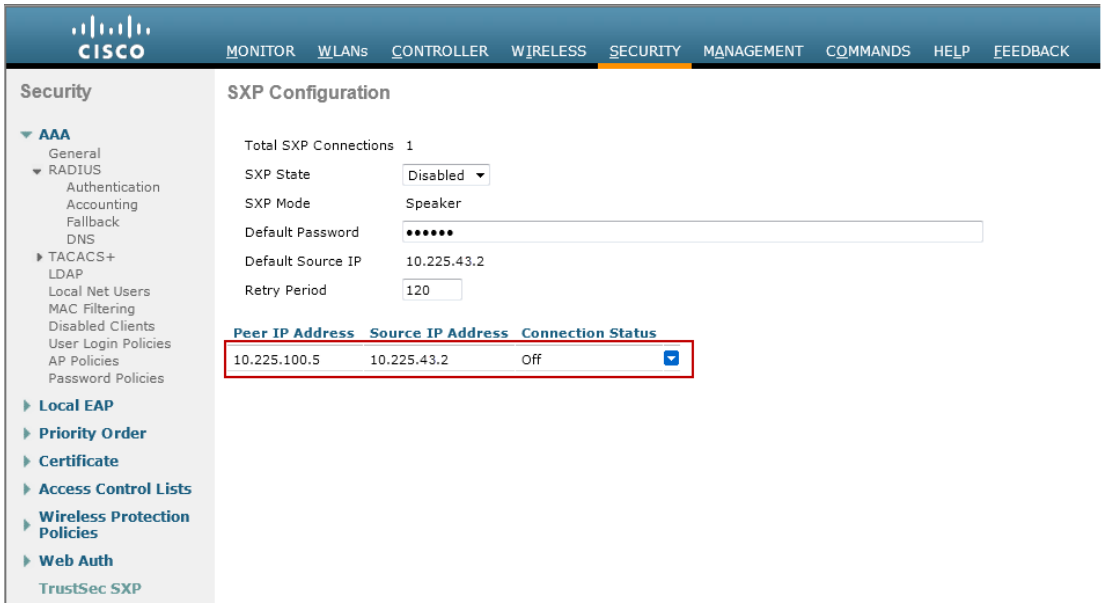


Figure 23-34 SXP Configuration Complete



Catalyst 6500 SXP Configuration

The following commands were used at the Shared Services Catalyst 6500 VSS to enable SXP peering with the CT5508 wireless controllers.

```
cts sxp enable/Enable CTS
cts sxp default source-ip 10.225.100.5/Source SXP connection from 10.225.100.5 (Lo)
```

```
cts sxp default password password/Configured password on the 6500 for incoming SXP
connections
cts sxp connection peer 10.225.43.2 source 10.225.100.5 password default mode local
listener hold-time 0 0/Builds an SXP connection to its peer at 10.225.43.2 using source
address of 10.225.100.5 and the default password defined above. Specifies that this
(local) device is in "Listener" mode. The source IP Address used here is purely optional
as it was specified above.
```

Issuing the command **show cts sxp connection** results in the following output.

```
ua28-6500-1>sh cts sxp connections
 SXP:Enabled
 Highest Version Supported:4
 Default Password :Set
 Default Source IP:10.225.100.5
Connection retry open period:120 secs
Reconcile period:120 secs
Retry open timer is not running
------------------------------------------------
Peer IP:10.225.43.2<--Wireless Controller
Source IP:10.225.100.5
Conn status:On
Conn version:2
Local mode:SXP Listener
Connection inst#:1
TCP conn fd:1
TCP conn password:default SXP password
Duration since last state change:0:21:49:55 (dd:hr:mm:sec)

Total num of SXP Connections = 1
```

Configuring Static IP/SGT Bindings on Nexus Switches

Unlike campus access through Catalyst Switches and Cisco Wireless Controllers where dynamic SGT mappings are communicated and created through 802.1X and RADIUS exchange, the vast majority of organizations do not implement 802.1X for server connectivity. As such, data center switches such as the Cisco Nexus switches provide only limited support for the use of 802.1X and do not specifically support an SGT RADIUS AV as an option. Therefore, IP Address to SGT mappings will be manually defined for bare metal and virtual servers.

For purposes of this CVD we have defined our IPtoSGT Bindings at the Nexus 7000 Data Center Aggregation layer switches as depicted below. Note that as there are two Nexus 7000s comprising the aggregation layer in the data center, both must be configured identically for consistent policy enforcement.

The following commands provide an example of these:

```
cts role-based sgt-map 10.230.4.2 40/Binds 10.230.4.2 to SGT 40
cts role-based sgt-map 10.230.4.22 40/Binds 10.230.4.22 to SGT 40
cts role-based sgt-map 10.230.5.2 50/Binds 10.230.5.2 to SGT 50
cts role-based sgt-map 10.230.6.2 40/Binds 10.230.6.2 to SGT 40
cts role-based sgt-map 10.230.7.2 50/Binds 10.230.7.50 to SGT 50
```

To verify the IP/SGT mappings at the Nexus 7000, issue the command **sh cts role-based sgt-map**.

```
IP ADDRESSSGTVRF/VLANSGT CONFIGURATION
10.225.49.1540vrf:1CLI Configured
10.225.42.1540vrf:1CLI Configured
10.230.1.4540vrf:1CLI Configured
10.230.1.4640vrf:1CLI Configured
10.230.4.240vrf:1CLI Configured
```

```
10.230.4.2240vrf:1CLI Configured
10.230.5.2 50vrf:1CLI Configured
10.230.6.240vrf:1CLI Configured
10.230.7.250 vrf:1CLI Configured
```

A key concept to be considered when granting access to the data center is that of the SGT value of zero or "Unknown" as it is referred to. If the IP Address of a server has not been mapped to an SGT at the point of enforcement such as the Nexus 7000 in the data center, that server would be considered Unknown and associated with SGT0. Unlike ACLs with an implicit deny at the end, SGACLs when implemented on a switching platform have an implicit permit to Unknown or all; this is not true on the ASA or IOS ZBFW acting as a SG-FW where an implicit deny is still maintained. Hence on a switch, if there is not a specific tag value assigned to a server, the destination is considered Unknown (SGT0) and the packet forwarded. This SGT0 thus allows a migrational approach to tag assignment in the data center.

In a BYOD setting where personal or contractor assets are permitted on the network and only partial access to data center resources is permitted through the use of SGACLs on the switching infrastructure, the task of assigning an SGT to every data center asset will likely prove daunting when first migrating to the use of Security Group Tags as it is not possible to create a switch-based SGACL using both SGT and IP addresses; the destination SGT of the device must be known. Hence it is with this in mind that the concept of the Unknown tag may be used to provide a phased approach to SGT assignment. Throughout the campus infrastructure the default policy permitting any SGT to "Unknown" is left unaltered. However, at the Nexus 7000 Data Center Aggregation switches where policies based on SGACLs are enforced, an explicit deny of a given source SGT to "Unknown" can be created. Using this premise one could easily identify and provide a tag for those servers where open access is permitted. These devices with partial access can get to them, but if a tag doesn't exist for any other server they are attempting to access (unknown) the SGACL will deny access. As servers that only some users or other servers should be able to access are associated with a SGT, new policies can be defined restricting that access.

A policy that is explicitly (manually) configured at a networking device will take precedence over a policy that is dynamically received from ISE. Hence By defining this SGACL at these data center aggregation switches only, these partial access devices are permitted access to Unknown anywhere else in the network except in these data center switches.

The following commands configure the SGACL locally on the Nexus 7000 Data Center Aggregation switches:

```
cts role-based access-list block12toUnk/Creates an SGACL "block12Unk"
  deny ip/Action performed by  SGACL "block12Unk"
cts role-based sgt 12 dgt 0 access-list block12toUnk/Manually configure mapping of Cisco
TrustSec Security Group Tags (SGTs) to a security group access control list (SGACL).
Defines source SGT12 to destination SGT0 (Unknown)
```

To verify the role-based access policy at the Nexus 7000, issue the command **sh cts role-based policy**. The following is an excerpt of the entire output showing the previously denied SGACL.

```
ua33-n7k-1-agg# sh cts role-based policy
sgt:10
dgt:unknown     rbacl:Permit IP
        permit ip

sgt:11
dgt:unknown     rbacl:Permit IP
        permit ip

sgt:12
dgt:unknown     rbacl:block12toUnk
        deny ip

sgt:any
```

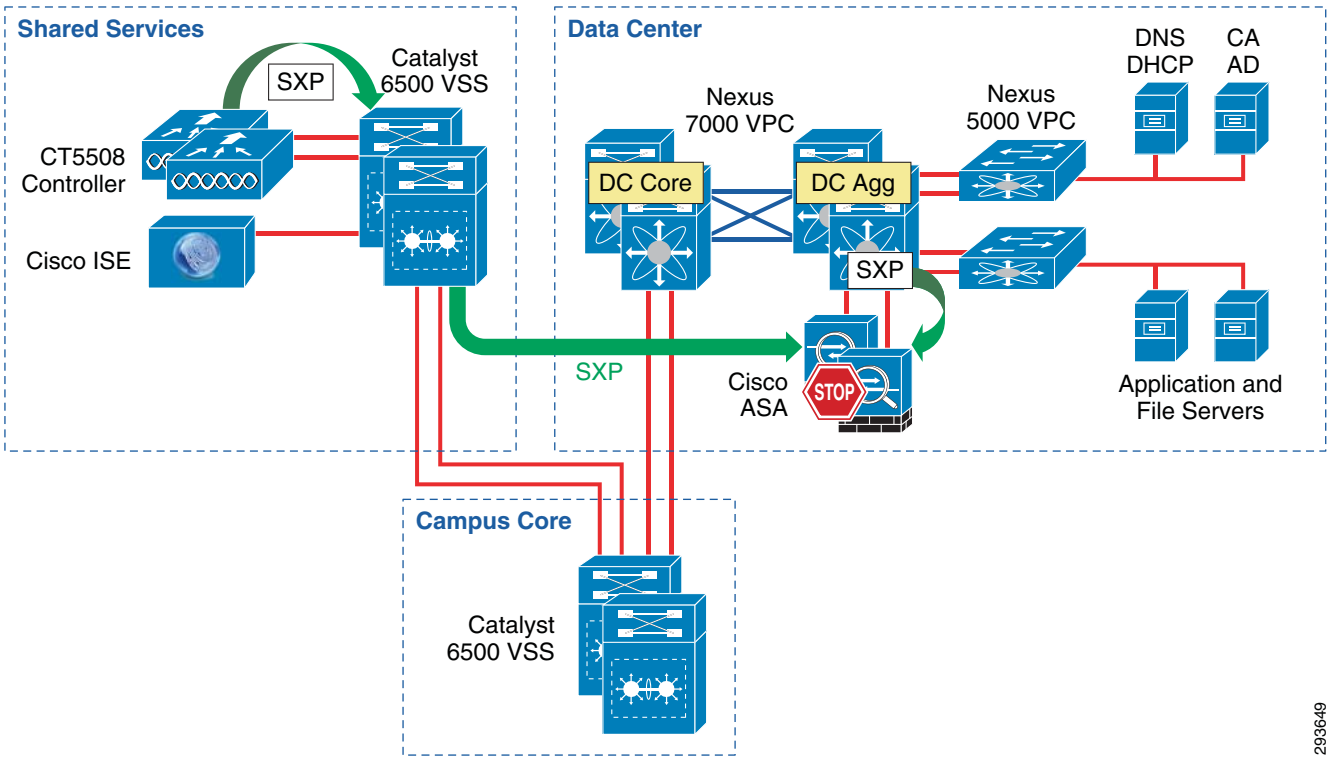
```
dgt:any rbacl:Permit IP
        permit ip
```

This completes the configuration for Deployment Scenario 1. The next steps are to configure the actual user policies as defined in Chapter 16, "BYOD Limited Use Case—Corporate Devices" for corporate devices and Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices" for personal devices.

TrustSec Policy Configuration Using the ASA and Security Group Firewall in the Data Center—Deployment Scenario 2

For the topology used in Deployment Scenario 2, see Figure 23-35.

Figure 23-35 Deployment Scenario 2 Configuration



With Deployment Scenario 2 an alternate means other than SGACLs will be used to enforce TrustSec policy. In Scenario 2 an ASA running version 9.0(2) will be used as a Security Group Firewall (SG-FW) securing data center resources from outside access. Unlike Scenario 1, the 10GE infrastructure between the Shared Services Catalyst 6500 VSS and the data center does not need to be enabled to support Security Group Tag forwarding or SGACLs. As the ASA does not presently support native SGT tagging on its Ethernet interfaces, Security Group Tag Exchange Protocol (SXP) must be used for it to learn IP/SGT mappings from other areas of the network where they have been dynamically learned or statically configured. It is by virtue of these SXP advertisements that the ASA is capable of inspecting the untagged traffic and, through the use of these IP/SGT mappings, that SG-FW policies are enforced

As in the case of the first deployment scenario, wireless users, upon successful authentication and authorization, will be associated with a specific role and an IP to SGT Binding will be created on the wireless controller with the device's IP Address and the appropriate SGT. SXP will be used to communicate this mapping to the Shared Services Catalyst 6500s to which the wireless controllers are attached. SXP will then be used to re-advertise the mappings the Shared Services Catalyst 6500 VSS has learned from the wireless controllers to the ASA firewall. In addition to the SXP Peering between the Shared Services 6500s and the ASA, the Nexus 7000 aggregation switches will also require an SXP Peering with the ASA firewall to advertise SGT mappings that have been statically configured on them.

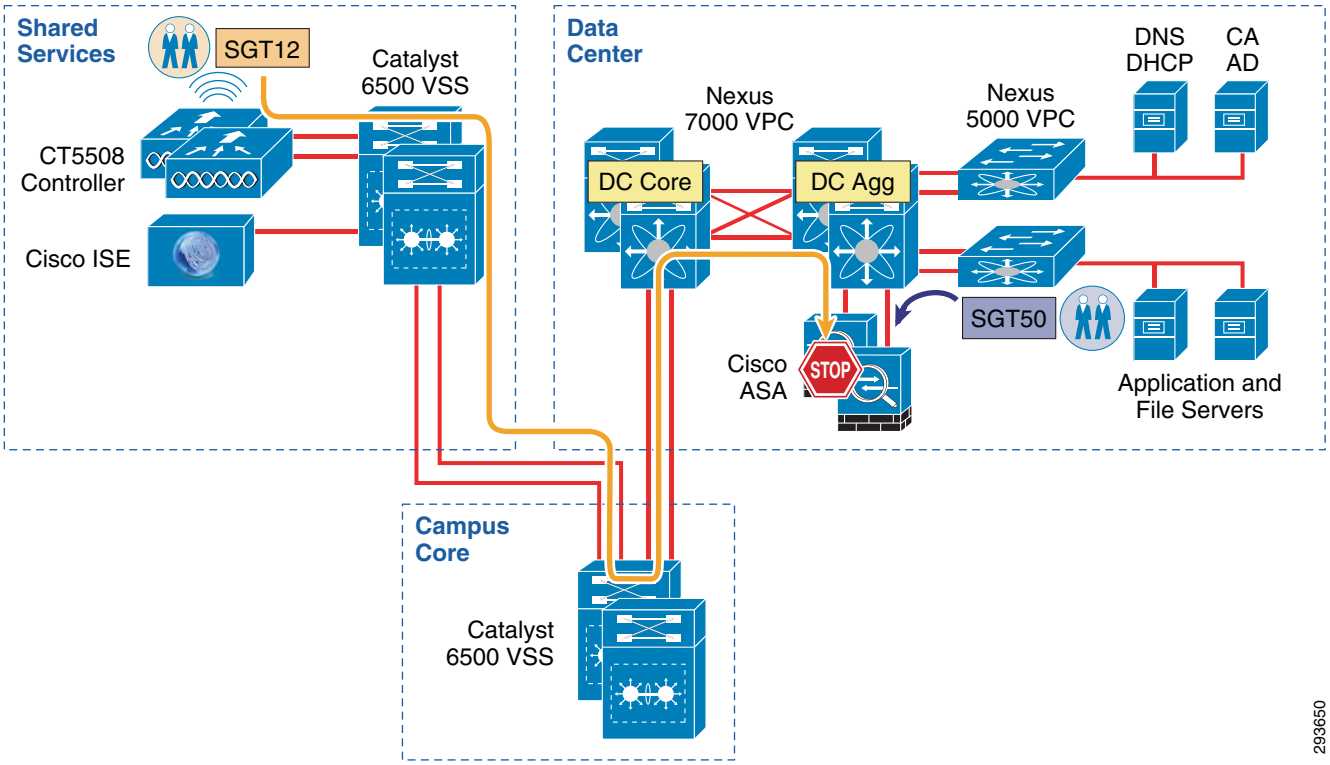
As previously discussed, the ASA firewall that will be used to enforce SG-FW policies must be manually configured with SGT policies as dynamic updates via ISE is presently not supported in the ASA. The details regarding these SG-FW policies will be discussed later.

As wireless traffic egresses the Shared Services Catalyst 6500s en route to the data center, the traffic will be untagged and will simply pass through the Core, enter the data center switching infrastructure, and ultimately arrive at the ASA firewall where the appropriate SG-FW policy will be enforced.

In the unlikely event that any traffic would be sourced from a server in the data center, it would likewise egress the Nexus 7000 aggregation switch untagged and be forwarded to the ASA firewall where any applicable SG-FW policy will be enforced.

Figure 23-36 depicts the infrastructure used in Deployment Scenario 2 and the means by which security group policies will be enforced.

Figure 23-36 TrustSec Policy Enforcement Using SXP and SG-FW



The following major tasks will be required for this deployment scenario and outlined in the following sub-sections:

1. Configuring ISE to support Security Group Access.
2. Configuring ISE for Network Access Device Authentication.

3. Configuring Network Access Devices for Authentication at ISE.
   a. RADIUS server configuration on the CT5508.
   b. RADIUS server configuration at the Nexus 7000.
   c. RADIUS server Configuration on the ASA.
4. Configuring Security Group Tag Exchange Protocol (SXP).
   a. Wireless Controller Configuration.
   b. Shared Services Catalyst 6500 VSS SXP configuration.
   c. Nexus 7000 Data Center Aggregation SXP Configuration.
5. Configuring static IP/SGT Bindings on Nexus switches.
6. Configuring SG-FW policies on the ASA firewall.

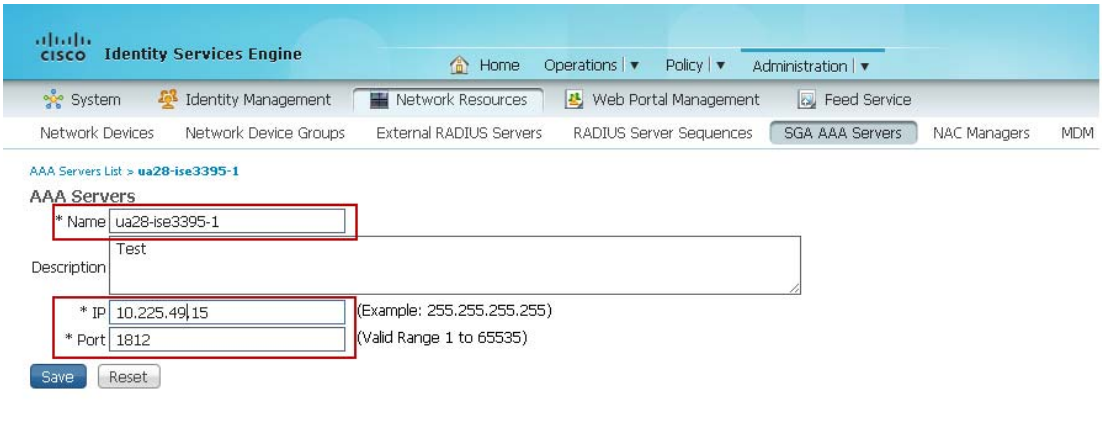
Configuring ISE to Support TrustSec

For Deployment Scenario 2, the ISE configuration required to define the Security Group Tag value/name remains the same as previously described for Scenario 1. However, it will be necessary to configure the ASA SG-FW policy manually at the ASA as there is no NDAC support for acquiring TrustSec egress policies cannot be learned dynamically. As such, unlike Deployment Scenario 1, it is unnecessary to configure the actual policies to be enforced at ISE; only the SG Names or Tables as they are commonly referred to. This will be addressed in an upcoming section.

For Cisco ISE to function as a Security Group Access (TrustSec) server and provide TrustSec services, you must define the following, global TrustSec settings. The first step is to define ISE as an TrustSec AAA server as depicted in Figure 23-37.

1. Go to Administration > Network Resources > TrustSec AAA Servers and click **Add**.
2. Enter the host name of the Identity Services Engine server or Policy Service Node if ISE Roles have been distributed among dedicated servers.
3. Enter the IP Address of the ISE server.
4. Enter the UDP Port number for RADIUS authentication and click **Save**.

Figure 23-37 ISE TrustSec Server AAA Definition



The next step is to configure TrustSec Server Protected Access Credential (PAC) Time-to-Live settings and SGT reservations.

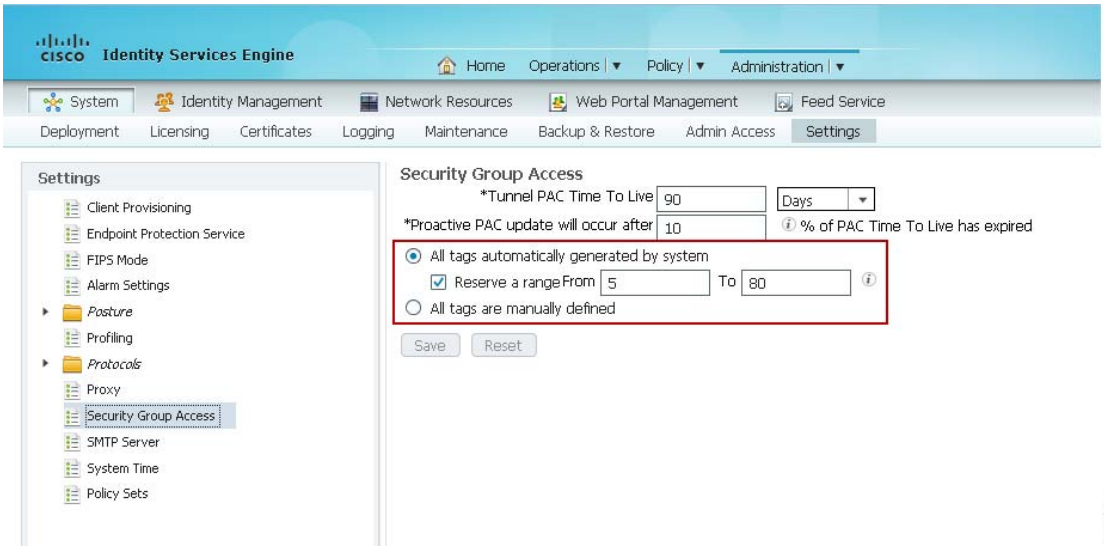
The tunnel PAC generates a tunnel for the EAP-FAST protocol and is used for Secure RADIUS communications with Network Devices for TrustSec environmental data. Unlike other networking devices such as Catalyst and Nexus switches, a new PAC must be manually re-generated and reapplied for the ASA prior to expiration; this will be discussed in the following ASA section as well.

By default Security Group Tags are dynamically assigned a decimal/hex value in ascending order by ISE. It is possible to change this behavior such that all tags must be manually defined or to reserve a range that can be specifically allocated to users, devices, or servers, In the CVD, a range of SGT values will be reserved for allocation among users/Devices and servers. Figure 23-38 depicts the reservation of Tags 5 through 80 for use in the CVD.

To complete this step:

1. Access the Identity Services Engine and follow the path Administration > System > Settings > Security Group Access.
2. Configure the Tunnel PAC Time to Live.
3. Configure the Proactive PAC update time if desired. In Figure 23-38, the PAC will be re-negotiated after 10% of TTL or nine days.
4. By default the system will automatically assign SGT values. If you wish to reserve a range that can be specifically allocated to users, devices, or servers, select the check box next to “Reserve a range From” and specify the Tag values. Here SGT 5 through 80 have been reserved.
5. Save the settings.

Figure 23-38 TrustSec Servers Settings in ISE



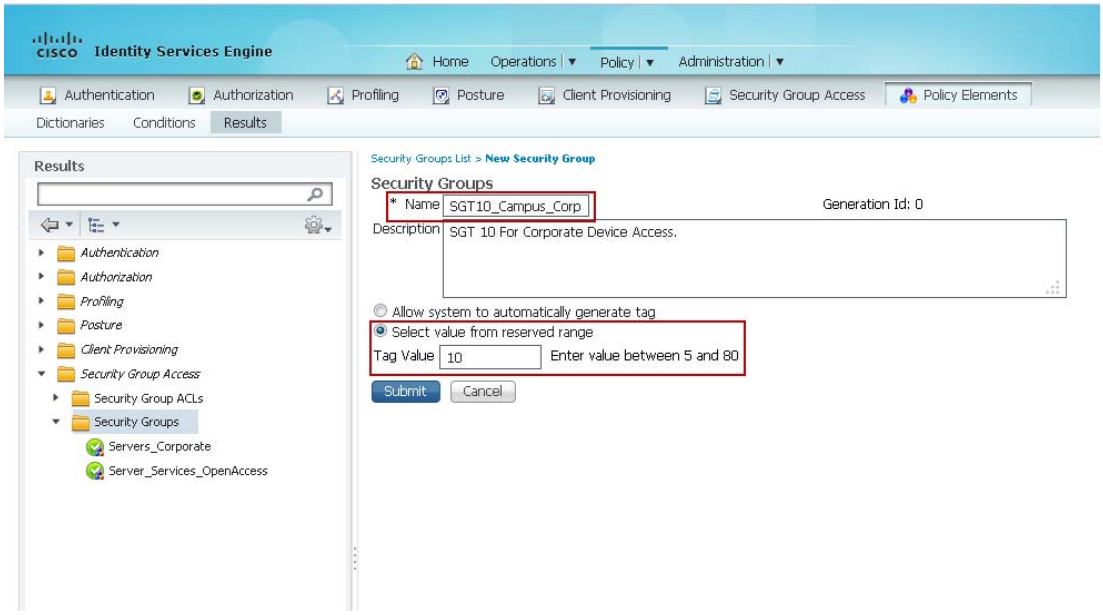
The next step is to define Security Group Tag names and associate them with a numerical value at the Identity Services Engine. The SGT names are periodically pushed to the Network Access Device (NAD) through periodic updates or upon that network device's authentication (NDAC Authentication) with ISE. They may also be manually pushed as well.

To complete this step as depicted in Figure 23-39:

6. Go to Policy > Policy Elements > Results >Security Group Access > Security Groups.
7. Click **Add**.
8. Define the SGT Name and add an optional description.

9. Click the radio button next to “Select value from reserved range”.
10. Enter the desired SGT value from the range defined in the previous step.

Figure 23-39 Security Group Creation



In BYOD v2.5, Table 23-3 shows the SGT Names and corresponding Tag Values used.

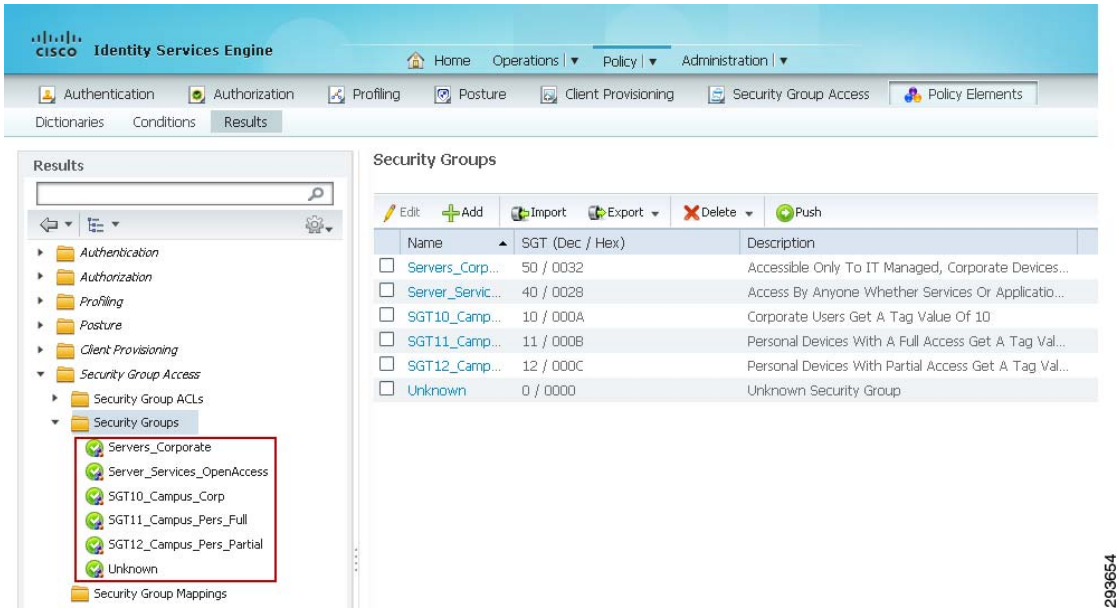
Table 23-3 SGT Names and Tag Values

SGT Value	SGT Name	Description
10	SGT10_Campus Corp	Corporate device with full access to the network.
11	SGT11_Campus_Pers_Full	Personal device with full access to the network
12	SGT12_Campus_Pers_Partial	Personal device with restricted access to some resources on the network.
40	Servers_Services_OpenAccess	Servers in data center accessible by all devices.
50	Servers_Corporate	Corporate Servers that only Corporate devices or approved personal devices have access to.
0	Unknown	System defined/reserved representing a device (IP Address) not associated with a SGT.

Note Deployment Scenario 2 in this CVD does not use device SGTs. If desired, refer to Configuring Network Devices with a Device SGT in Deployment Scenario 1.

Figure 23-40 depicts all of the security group names defined in ISE for this CVD.

Figure 23-40 Security Groups Used in CVD

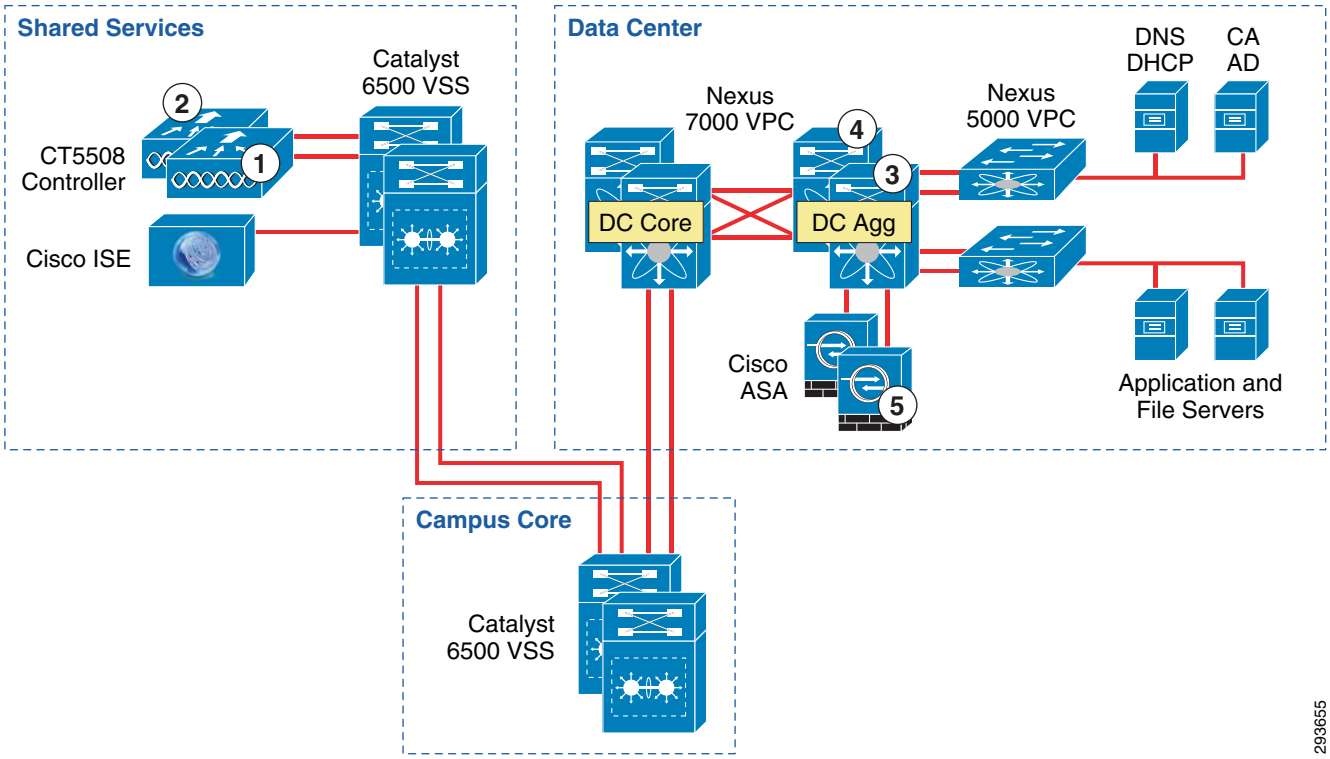


Configuring ISE for Network Access Device Authentication

Network device definitions for the wireless controllers, Nexus switches, and the ASA firewall must be created in ISE for Deployment Scenario 2. Each of the two controllers will need to be defined in ISE as both are active, however as the ASA firewalls are running in an Active/Standby HA mode, only the primary firewall needs to be defined in ISE. These devices are depicted in Figure 23-41 and are:

1. CT5508 Wireless Controller #1 (required for 802.1X wireless device access)
2. CT5508 Wireless Controller #2 (required for 802.1X wireless device access)
3. Data Center Nexus 7000 Aggregation Switch #1 (TrustSec policy enforcement)
4. Data Center Nexus 7000 Aggregation Switch #2 (TrustSec policy enforcement)
5. ASA Firewall

Figure 23-41 Network Devices to be Defined in ISE for Scenario 2



Wireless Controller Configuration

In this deployment scenario, it is necessary to define the wireless controllers as they will serve as 802.1X Authenticators for the wireless devices accessing the network. As such, the controllers upon a device's successful authentication and authorization will within a RADIUS exchange with ISE, receive the appropriate SGT value to be associated with that device's IP Address. More than likely, wireless controllers will likely have already been defined for 802.1X wireless access regardless of the use of SGT or ACLs for enforcing policies. In step 4 below however, we have suggested the additional use of Network Device Location for use within Authorization policies.

The following steps must be taken to define the wireless controllers within ISE as depicted in Figure 23-42; the ASA specific configuration will follow separately:

1. At ISE go to Administration > Network Resources > Network Devices and click **Add**.
2. Enter the hostname of the device. This will be the same name as configured at the wireless controller name.
3. Enter the IP Address of the wireless controller. This must be the address used to source all RADIUS communications from the device.
4. Change the Network Device Location or Device Type if a custom location/type has been previously defined. For the wireless controllers, the CVD uses a custom Location known as "Campus_Controller:SGT_Enabled". This custom location is configured under "Network Device Groups" as depicted in Figure 23-43. The significance of the use of a custom device location lies in the ability to use that as an attribute within the Authorization Profile at ISE to determine a result. The CVD uses the location "Campus_Controller:SGT_Enabled" to determine that an SGT Value

should be handed back to the wireless controller after a user or device's successful authorization as opposed to an ACL for policy enforcement. This allows for a migrational approach to deploying infrastructure that will make use of SGT as opposed to ACLs for policy enforcement.

5. Configure the RADIUS Shared Secret. This must match that configured on the network device.
6. Click the down arrow next to SNMP Settings and complete as appropriate.
7. Complete steps one through six for every wireless controller providing 802.1X-authenticated wireless access to users.

Note that it is unnecessary to complete the Advanced TrustSec Settings for CT5508 or WiSM2 wireless controllers. These controllers are not capable of supporting anything other than the creation of IP Address to SGT mapping and advertising those mappings to manually configured peers via the SXP protocol. These settings are only used for those devices supporting SGACLs and Native SG Tagging on SGT-capable hardware requiring the download of TrustSec environmental and policy data from ISE.

Figure 23-42 Wireless Controller Definition at ISE

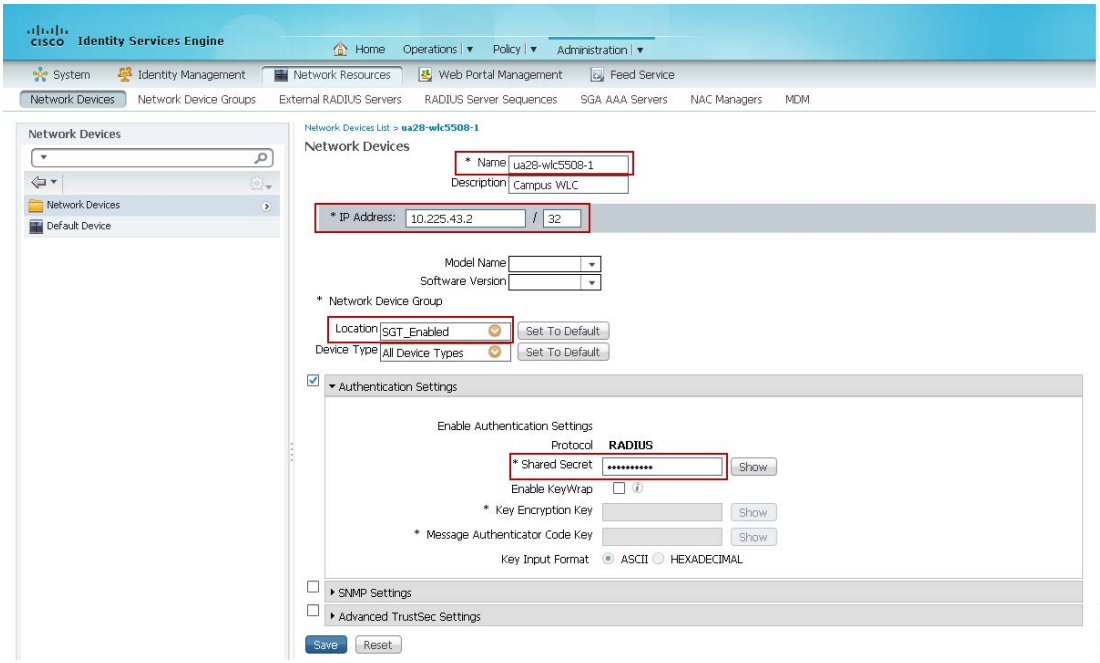
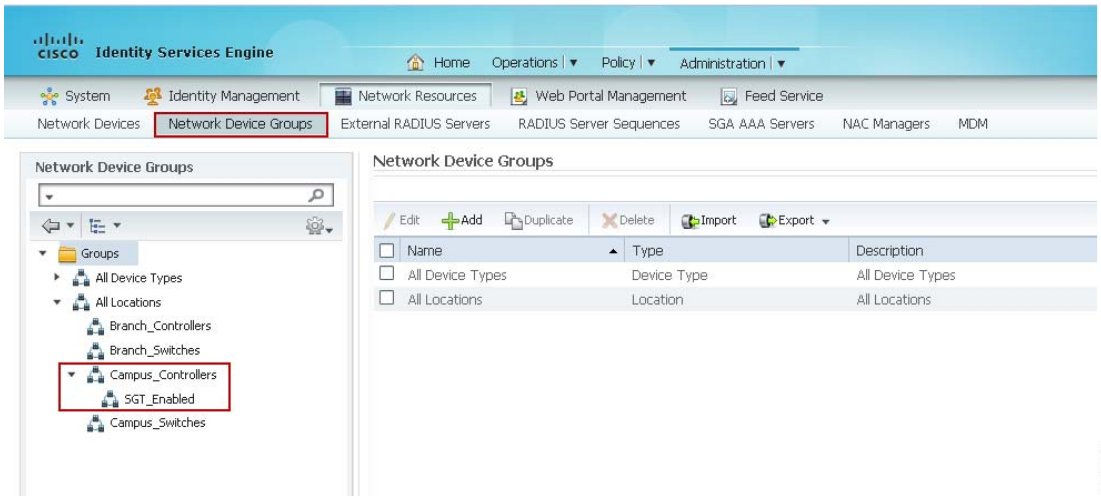


Figure 23-43 Network Device Group Definition



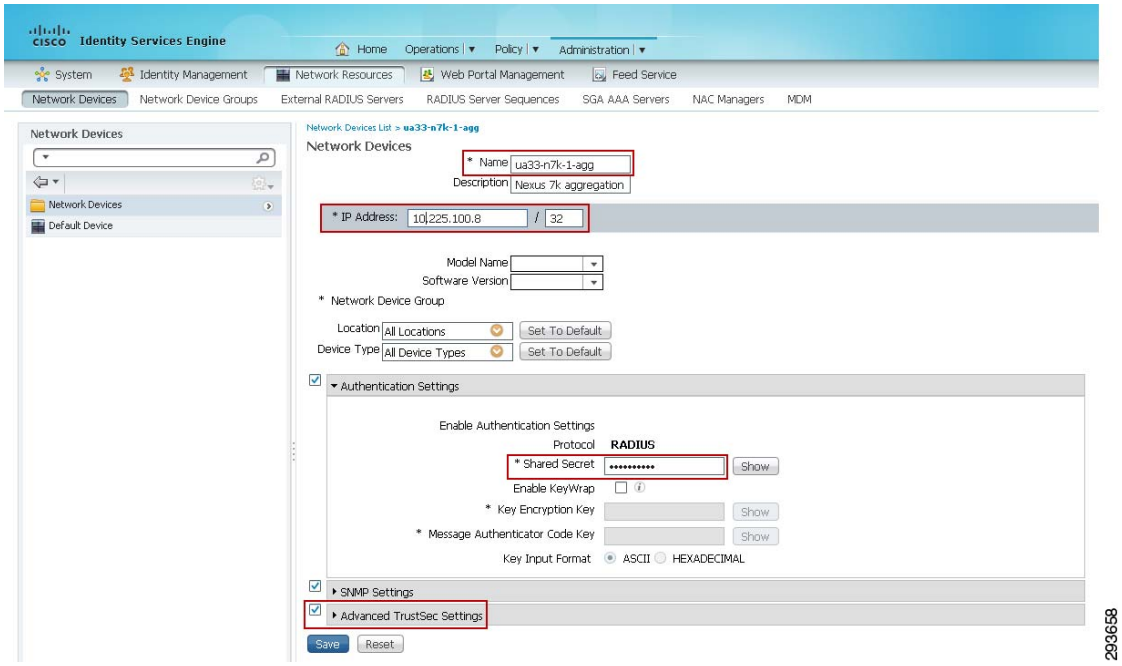
This completes the Network Device Definition for the wireless Controllers in ISE.

Nexus 7000 Configuration

Although SGACLs will not be enforced at the Nexus 7000 Data Center Aggregation switches as in the first deployment scenario, but rather at the ASA configured as a Security Group Firewall, it is still necessary to define the Nexus 7000 aggregation switches in ISE in order to share TrustSec Environment Data in order to learn SG Names and IP/SGT mappings defined centrally at ISE. Note that for the purposes of the CVD, all IP/SGT mappings will be statically configured on the Nexus 7000 aggregation switches as in the first deployment scenario. The following steps must be taken to define the Nexus 7000 Data Center Aggregation switches within ISE as depicted in Figure 23-44:

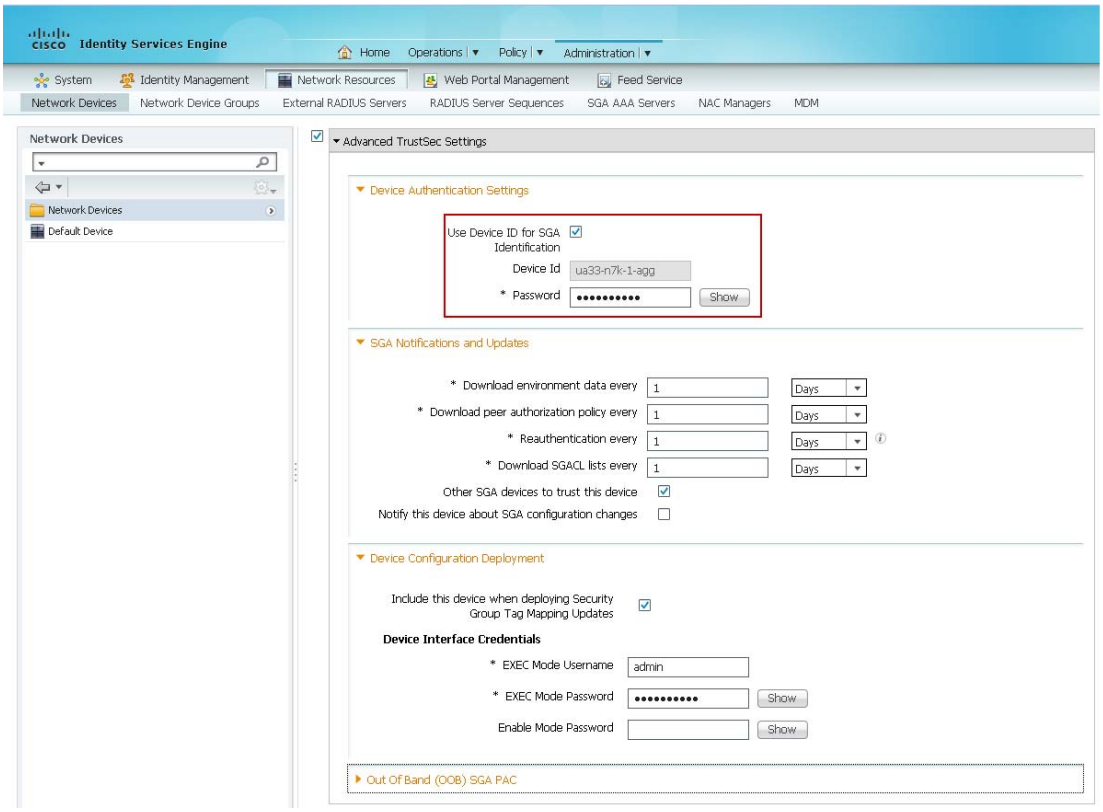
1. At ISE go to Administration > Network Resources > Network Devices and click **Add**.
2. Enter the hostname of the device. This will be the same name as configured at the network device and documented later with the **cts credential** command on switches and would be the wireless controller name.
3. Enter the IP Address of the network device. This must be the address used to source all RADIUS communications from the device.
4. Change the Network Device Location or Device Type if a custom location/type has been previously defined. Within the CVD the Shared Services, Core, and Data Center switches all make use of the default setting as seen in Figure 23-42. The exception to this are the wireless controllers as previously discussed. Configure the RADIUS Shared Secret. This must match that configured on the network device.
5. Click the down arrow next to SNMP Settings and complete as appropriate.

Figure 23-44 Nexus 7000 Generic Definition at ISE



6. Click the down arrow next to Advanced TrustSec Settings. Once the Advanced TrustSec Settings configuration box has been expanded as seen in Figure 23-45, click the check box next to "Use Device ID for TrustSec Identification"
7. Enter the password that will be configured later on the network device in the **cts credential** command. This can be the same as the RADIUS Shared Secret.
8. Configure the desired settings for "TrustSec Notifications and Updates". Note that these are the settings that determine the frequency of the TrustSec Environment updates to the network device. It is recommended that aggressive timers not be used here and as such these have been left at the default value for one day. Note that this is to configure the automated, periodic update of the pertinent data. In addition to these periodic updates, it is possible to manually push updates for SGT Names/Values, Network Device SGT, SGT Egress Policy (SGACL), and AAA Server List from within ISE to those network devices supporting Change of Authorization (CoA). Note that the Nexus 7000 does not support CoA at the time of this document. To support a manual push of environmental data and policy to the Nexus 7000, it is possible to do so through re-issuing the <cts credential> command at the Nexus 7000 discussed later. For further information regarding these parameters and how TrustSec environment data is exchanged, refer to the ISE User Documentation and specifically the "Configuring TrustSec Settings on Switches" and "TrustSec CoA" sections of "Configuring Cisco Security Group Access Polices" located in the ISE 1.2 User Guide at: http://www.cisco.com/en/US/products/ps11640/products_user_guide_list.html.
9. Enter the credentials to access Exec Mode (if applicable) and the Enable Mode password used by ISE to access the device to manually push updated information.

Figure 23-45 Nexus 7000 TrustSec Definition at ISE



ASA Firewall Configuration

Once the wireless controllers and Nexus switches have been defined in ISE, the ASA firewall will be defined next. Whereas Catalyst and Nexus switches can import the PAC file from ISE when authenticating for the first time, The ASA firewall requires that this process be performed manually. Importing the PAC file to the ASA establishes a secure communication channel with ISE. After the channel is established, the ASA initiates a PAC secure RADIUS transaction with ISE and downloads Cisco TrustSec environment data; specifically, the ASA downloads the security group table. As discussed earlier, the ASA firewall does not download SGT Policies only the SG tables; the SGT policies will be manually defined at the ASA. The security group table maps SGTs to security group names. Security group names are created on ISE and provide user-friendly names for security groups.

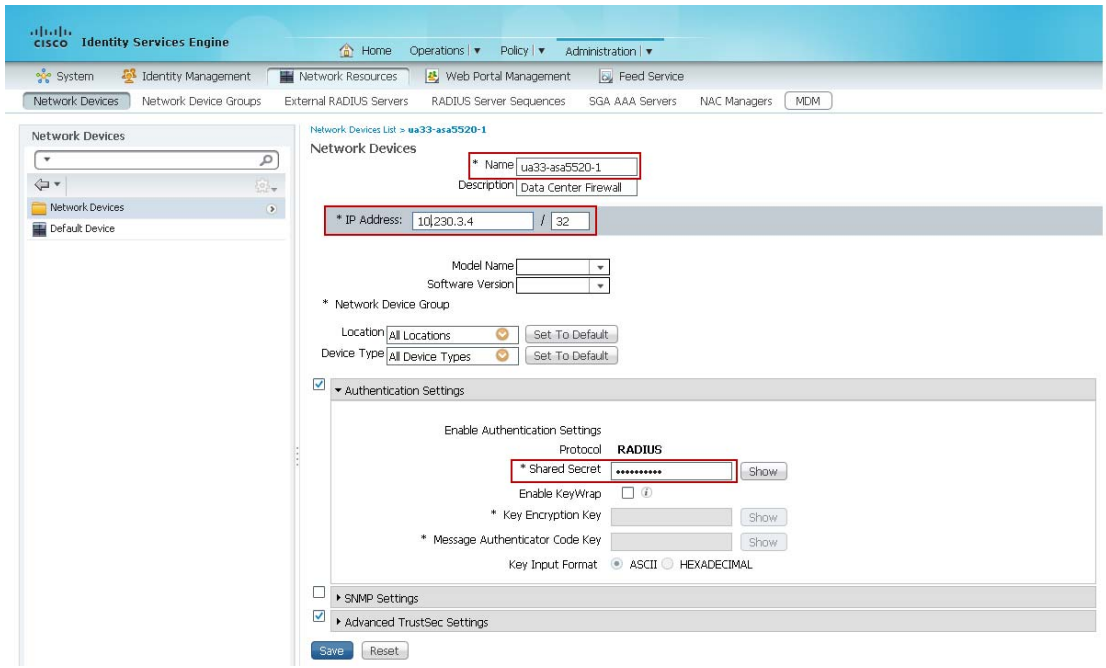
The first time the ASA downloads the security group table, it walks through all entries in the table and resolves all the security group names contained in security policies configured on the ASA; the ASA then activates those security policies locally. If the ASA is unable to resolve a security group name, it generates a system log message for the unknown security group name.

One special consideration needs to be kept in mind regarding the ASA and that is if PAC expiration occurs, a new PAC file will not be automatically downloaded as in the case of Catalyst and Nexus switches and hence updated SGT tables will not be able to be downloaded. At the time of this writing (v9.0.2), a new PAC file must be imported manually prior to the expiration of the existing one. If the ASA cannot download an updated security group table, the ASA continues to enforce security policies based on the last downloaded security group table until a new PAC file is downloaded and the ASA downloads an updated table.

The following steps must be completed to define the ASA firewall within ISE as depicted in Figure 23-46:

1. At ISE go to Administration > Network Resources > Network Devices and click **Add**.
2. Enter the hostname of the ASA firewall.
3. Enter the IP Address of the interface closest to the ISE server. In the case of the CVD, that happened to be the Outside Interface as ISE was located in a Shared Services block logically separated from the data center servers located within protected zones on various inside interfaces of the firewall.
4. Change the Network Device Location if appropriate.
5. Configure the RADIUS Shared Secret. This must match that configured on the ASA firewall.

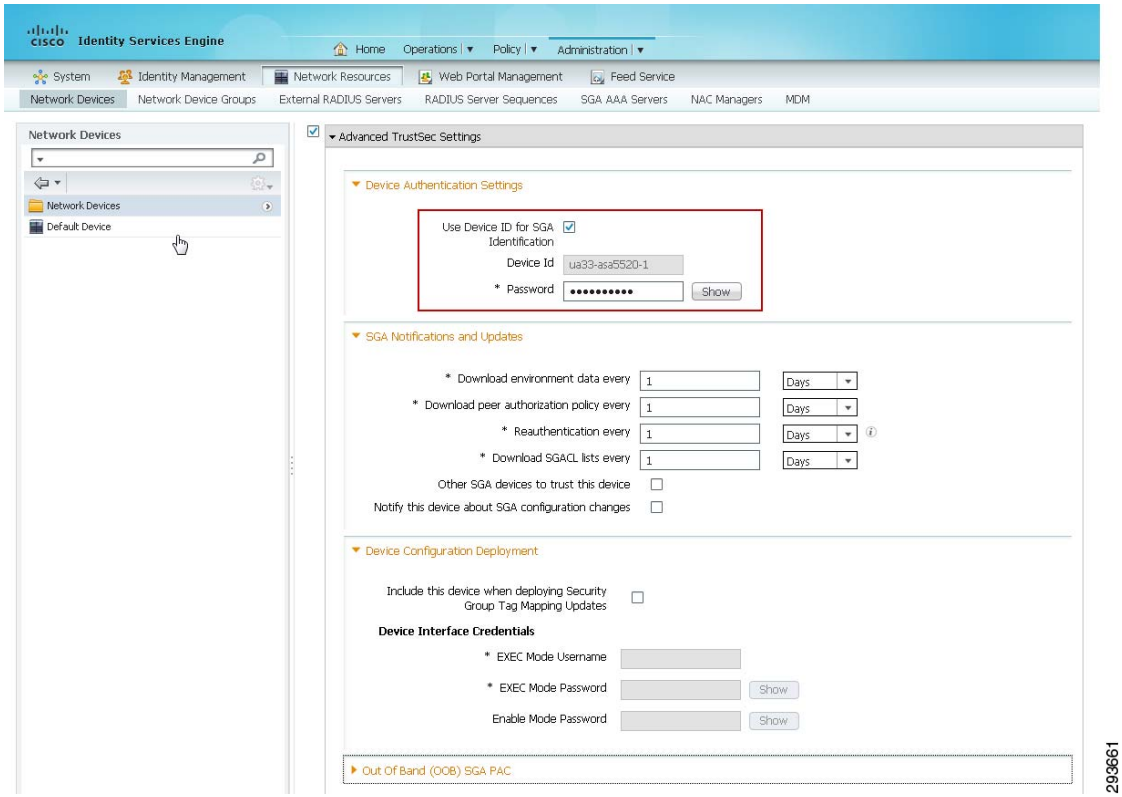
Figure 23-46 ASA Network Device General Settings in ISE



Next complete the following steps for TrustSec configuration as depicted in Figure 23-47:

6. Click the arrow next to “Advanced TrustSec Settings”
7. Enter the Device ID.
8. Enter the password.
9. Note that none of the other settings for “TrustSec Notifications and Updates” and “Device Configuration Deployment” need to be completed. The TrustSec Environment Data and particularly the Security Group Tables/Names are downloaded to the ASA either manually or periodically based on the SXP Reconcile Timer configured at the ASA and covered later. Device Configuration Deployment is used to update Security Group Policies via CoA from ISE to the device. As TrustSec policies cannot be dynamically learned and must be manually defined on the ASA, these setting are not applicable as well.

Figure 23-47 ASA TrustSec Settings in ISE



The final task at ISE is to generate an Out of Band PAC for subsequent importing at ISE as depicted in Figure 23-48:

10. Click the arrow next to “Out of Band (OOB) TrustSec PAC”.
11. Click the “Generate PAC” button.

    A popup will appear as depicted in Figure 23-49.

    The device identity will be pre-populated using the “Device ID for TrustSec” hostname.
12. Define an arbitrary Encryption Key to be used.
13. Set the PAC Time to Live. Notice that this be configured for Days, Weeks, Months, or Years.
14. Click the “Generate PAC” button.

    A PAC File will now be generated and you will be prompted to download and save the PAC file locally for later import and use by the ASA firewall.

Figure 23-48 Generating the TrustSec PAC File for the ASA Firewall at ISE

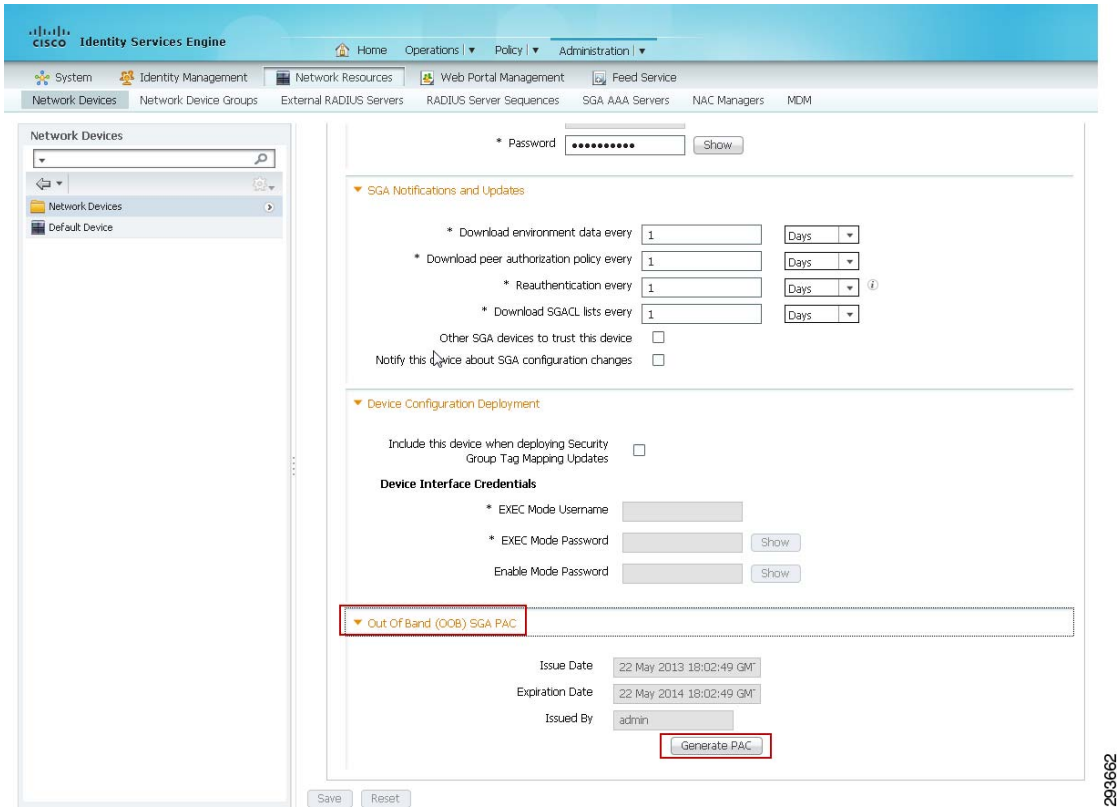
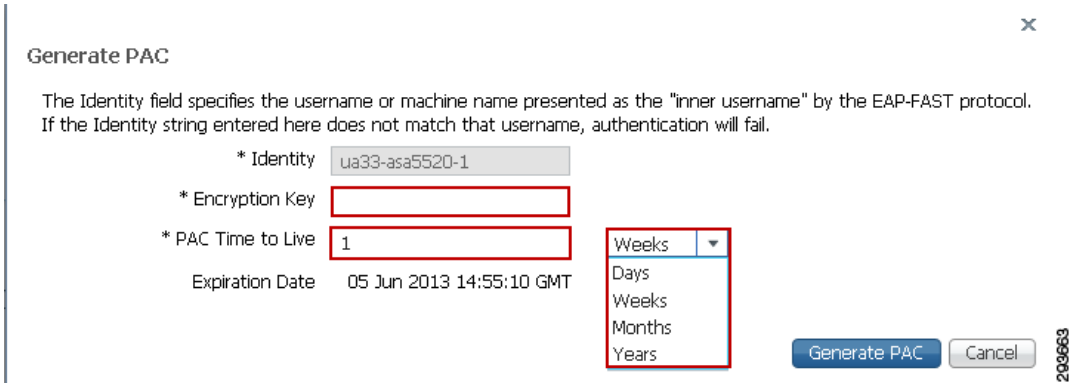


Figure 23-49 TrustSec PAC File Definition for the ASA Firewall at ISE



This completes the Network Device Definition for the ASA firewall in ISE.

Configuring the Network Devices for Integration with ISE

The following configuration tasks will all be completed at the network devices themselves. This configuration is critical to identify the ISE Primary server as the AAA server from which information regarding TrustSec will be exchanged. Note that for greater resilience, multiple ISE Policy Service

Nodes can be listed and will be tried in succession. As previously mentioned, it is only necessary to configure those devices that will provide access for the wireless users and the network devices where static IP/SGT mappings will be defined.

RADIUS Server Configuration on the Wireless Controller

The configuration of RADIUS server information is required in order for the controller to act as a RADIUS Authenticator for 802.1X-based authentication of wireless clients accessing the network. More than likely these steps have already been completed as this is a basic requirement for securing wireless access regardless of the desire to use ACLs or Security Group Tags.

1. Access the wireless controller either through the local UI or Prime. In Figure 23-50 the controller's GUI is depicted.
2. Go to Security > AAA > RADIUS > Authentication.
3. In the top right of the screen, if the ISE server has not been configured yet, click **New**.

   A new window will open as seen in Figure 23-51.

Figure 23-50 Wireless Controller RADIUS Configuration

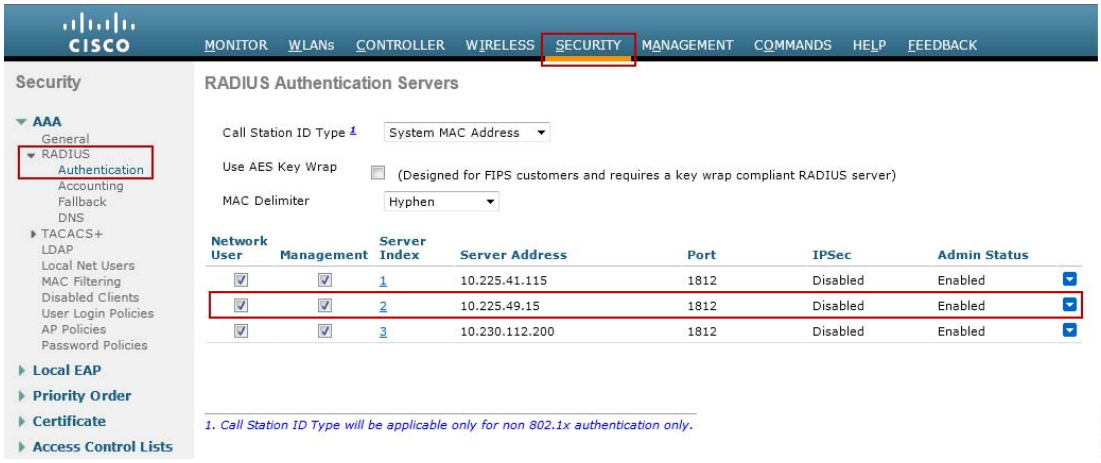
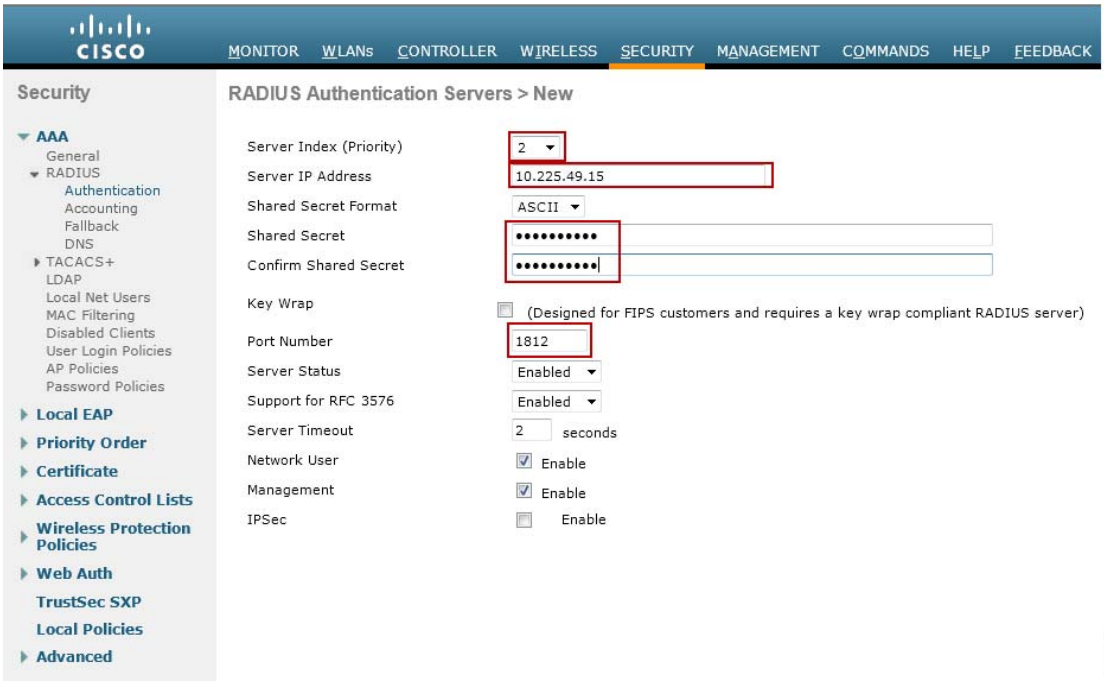


Figure 23-51 Wireless Controller RADIUS Server Configuration



4. Use the drop down to enter the correct priority; lower number is higher priority.
5. Enter the ISE server's IP Address.
6. Enter the Shared Secret which must match that configured for the wireless controller as defined in the Network Device List in ISE.
7. Enter the correct RADIUS Authentication UDP port number.
8. Click **Apply**.
9. Configure the controller's RADIUS Accounting information as the Authentication information above. This can be accessed by following Security > AAA > Radius > Accounting.

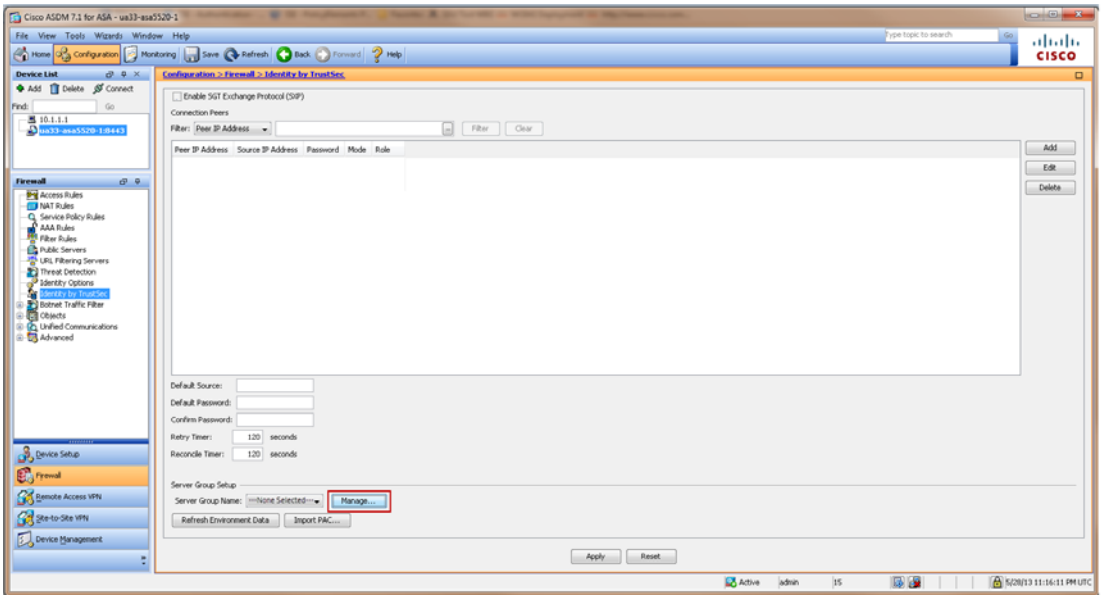
Configuration of the wireless controller RADIUS server configuration is complete. Repeat these steps if additional controllers need to be configured.

RADIUS Server Configuration on the ASA Firewall

The following configuration steps need to be completed in order to establish the ISE server as a AAA server for the ASA and importing the TrustSec PAC File used for secure RADIUS exchange of TrustSec Environment Data and the SGT Tables specifically.

1. Open ASDM.
2. In ASDM, navigate to Configuration > Firewall > "Identity by TrustSec" as depicted in Figure 23-52.
3. Click the **Manage** button in the "Server Group Setup" area.

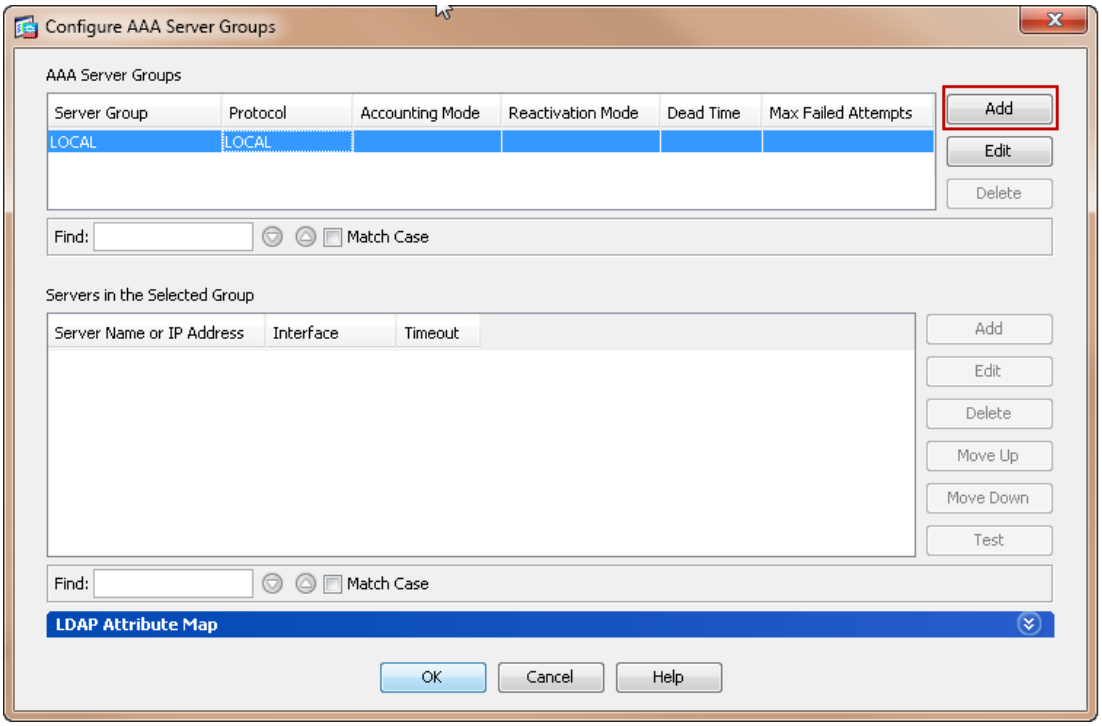
Figure 23-52 AAA Server Configuration in ASA



A popup window will open as can be seen in Figure 23-53.

4. Click the **Add** button.

Figure 23-53 Configuring AAA Server Groups in ASDM

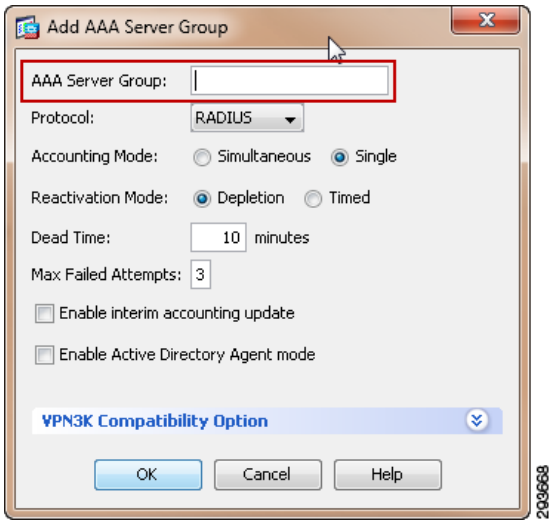


A popup window opens as seen in Figure 23-54.

5. Enter a name for the AAA server Group.

6. Click **OK**.

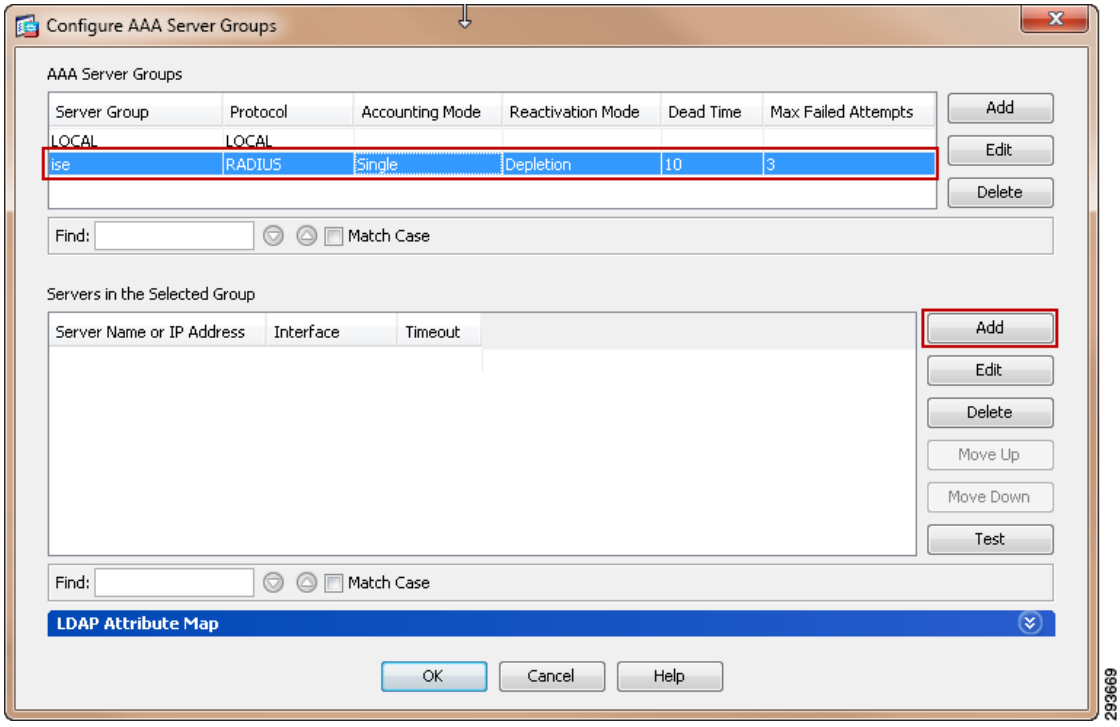
Figure 23-54 Adding AAA Server Group in ASA



Once OK has been clicked the popup window closes and the "Configure AAA Server Groups" window is populated with the new Server Group as seen in Figure 23-55.

7. Click the **Add** button next to the box for "Servers in the Selected Group".

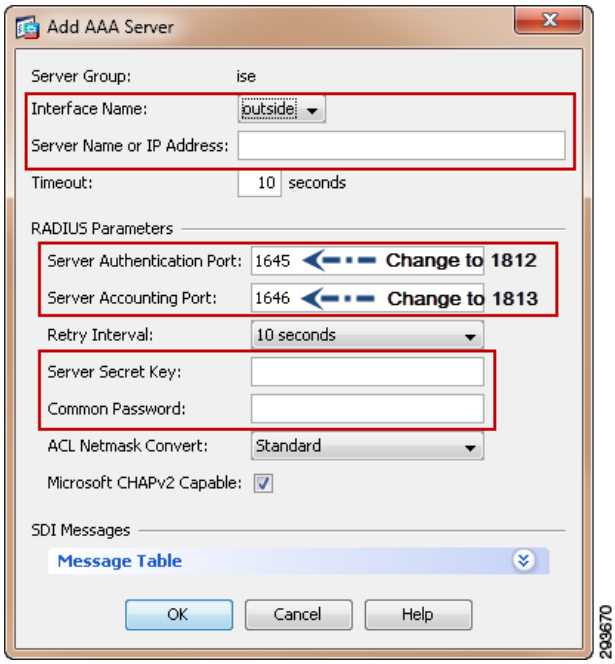
Figure 23-55 Adding AAA Server to Server Group in ASA



A popup window opens as seen in Figure 23-56.

8. From the Interface drop-down select the interface closest to the ISE server as that interface's IP Address will serve as the source address for all RADIUS communications with ISE.
9. Enter the DNS Hostname or IP Address of the ISE server (PSN).
10. Change the Server Authentication Port to 1812 to match that configured at ISE.
11. Change the Server Accounting Port to 1813 to match that configured at ISE.
12. Enter the RADIUS "Server Secret Key" and "Common Password". These will be the same as the shared secret key used earlier to define the ASA in ISE.
13. Click **OK** to add the AAA server to the AAA server Group. If more than one ISE Policy Service Node exist, repeat steps 10 through 15 to add additional AAA servers.

Figure 23-56 Adding AAA Server to Server Group



Once the AAA Server Group has been defined it will now be necessary to import the TrustSec PAC File at the ASA firewall. As depicted in Figure 23-57:

14. Down in the "Server Group Setup" area, select the correct AAA Server Group in the drop-down.
15. Click **Import PAC**.

    A popup window will open as seen in Figure 23-58.
16. Browse to the location where the file was locally stored when generating the PAC at ISE and use the password used during PAC File generation.

Figure 23-57 Importing TrustSec PAC File at ASA

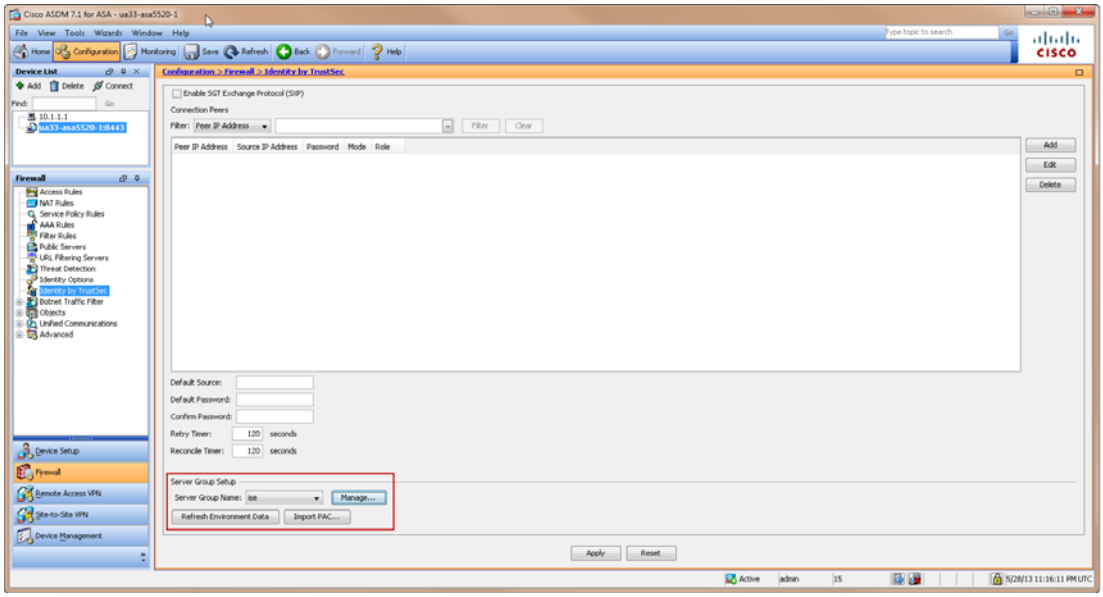
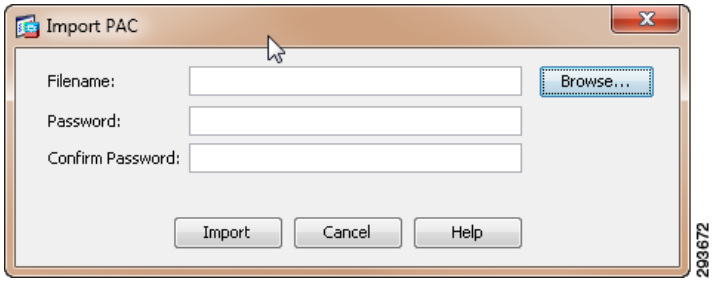


Figure 23-58 Importing TrustSec PAC at ASA



These final steps should be taken to validate that the PAC file was imported correctly at the ASA and that the SGT Tables have been downloaded from ISE. Issue the **show cts pac** command at the ASA firewall. The output should be as shown in Example 23-1.

Example 23-1 ASA show cts pac Output

```
ua33-asa5520-1/pri/act# show cts pac

  PAC-Info:
    Valid until: May 29 2014 21:01:35
    AID: 46094036746cf4dee55688b595c61925
    I-ID: ua33-asa5520-1
    A-ID-Info: Identity Services Engine
    PAC-type: Cisco Trustsec
  PAC-Opaque:
    000200b800030001000400104609403674 6cf4dee55688b595c619250006009c000301
    003104b1f4305546ff989c2e4dae291af500000013519faef000093a804ee94210ca6a
    4b1b5404baaa22eb122daa3f6bf5cdcbb5d19939191f6c10dab85ee437b842716e801c
    f2d1ad4b5d6a89a53605da628efd34ba58357273365d082f391b5642fe9d979df22878
    a5a55fa02b2a28e64967b0060546926bc4c3990b064d833cf2ba3184ad3491ac21438c
    a53cf02dd7b6559af1b4c25582
```

```
ua33-asa5520-1/pri/act#
```

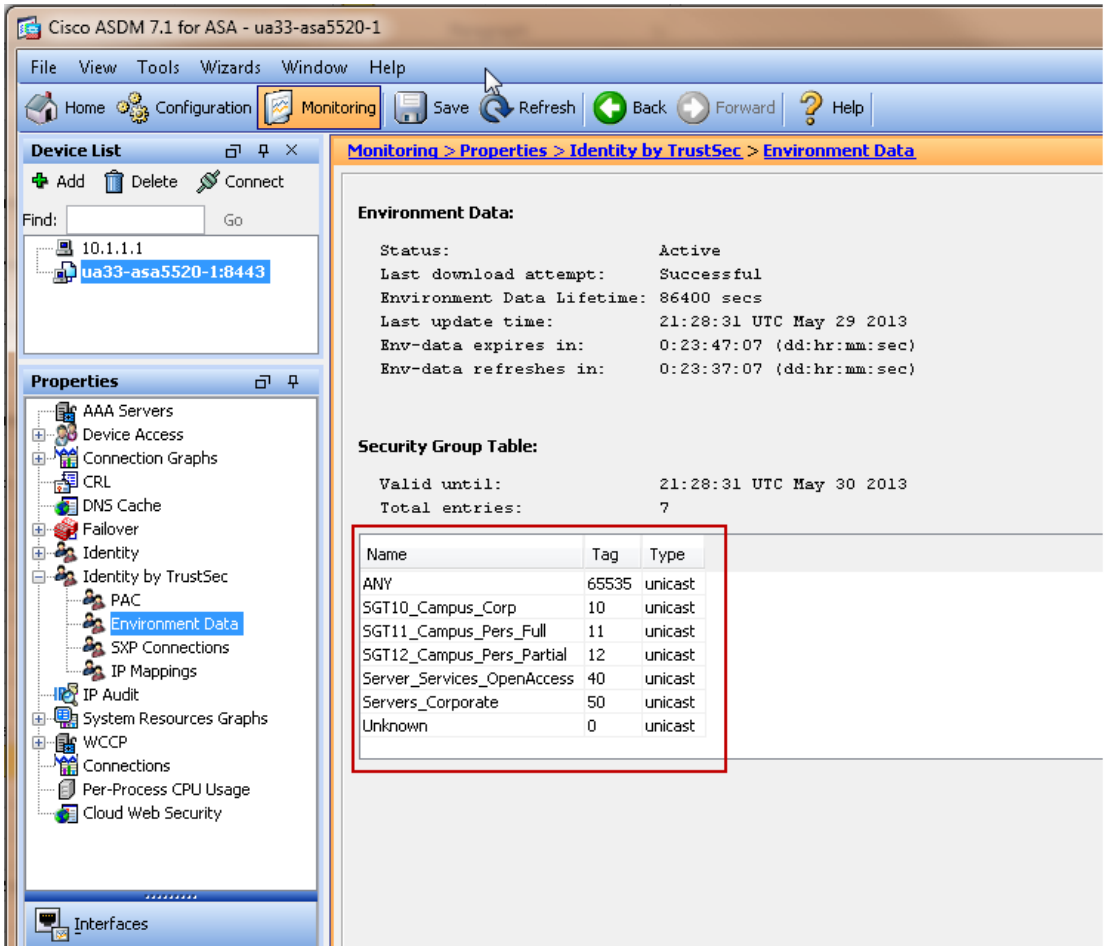
Now check that ASA to ISE communications has been successfully established by issuing the <show cts environment-data> command at the ASA. The following output should be seen.

```
ua33-asa5520-1/pri/act# sh cts environment-data
CTS Environment Data
====================
Status:Active
Last download attempt:Successful
Environment Data Lifetime:86400 secs
Last update time: 21:18:12 UTC May 30 2013
Env-data expires in: 0:08:33:27 (dd:hr:mm:sec)
Env-data refreshes in: 0:08:23:27 (dd:hr:mm:sec)
```

You can also, validate that the TrustSec Environment Data and the SGT Tables have been successfully downloaded using ASDM as can be seen in Figure 23-59:

17. From ASDM go to Monitoring > Properties > Identity by TrustSec > Environment Data.
18. The Security Group names configured earlier in ISE should be present as seen in Figure 23-59.

Figure 23-59 TrustSec Environment Data at ASA



RADIUS Server Configuration on the Nexus 7000

The following steps outline those tasks required to configure ISE as a RADIUS server at the Nexus 7000. As discussed, this is to establish a secure connection with ISE for the exchange of TrustSec Environment Data and specifically the SGT Names for use in creating IP/SGT Bindings at the Nexus 7000 Data Center Aggregation switches; no policies are available for download as they are configured at the ASA as SGACLs are not used in this deployment scenario.

```
feature dot1x/Enable dot1x support.
feature cts/Enable cts (TrustSec) support
```

Once the features have been enabled the AAA servers and TrustSec device credentials must be enabled through the following commands:

```
cts role-based enforcement/Enables SGACL enforcement on Nexus 7000
cts role-based counters enable/Enable role-based access control list (SGACL) counters
cts device-id device-id password password/TrustSec credentials for use with ISE; device ID
and password much be the same as at ISE.
radius-server host 10.225.49.15 key <secret> pac/Specifies the RADIUS authentication
server, shared secret must be same as configured for RADIUS secret at ISE.
aaa group server radius <ise>/Creates a AAA Server Group ISE
  server 10.225.49.15/Defines 10.225.49.15 as a member of Group ISE
aaa authentication dot1x default group <ise>/Specifies the 802.1X port-based
authentication method as RADIUS.
aaa accounting dot1x default group <ise>/Enables 802.1X accounting using group ISE
aaa authorization cts default group <ise>/To configure the default authentication,
authorization, and accounting (AAA) RADIUS server groups for Cisco TrustSec authorization
ip radius source-interface loopback0/Matches the IP Address of the device configured in
ISE; uses the IP Address of Lo0 to source all RADIUS.
```

Configuring Security Group Tag Exchange Protocol (SXP) for Wireless Controllers

Campus wireless users accessing the network upon successfully matching an authorization profile at the Identity Services Engine will be associated with an SGT. Upon successful authentication and subsequent authorization to the network the Identity Services Engine will pass the appropriate SGT value to the wireless controller through a RADIUS AV. This SGT value is associated with the IP Address of the wireless user obtained through the 802.1X authentication and an IP/SGT mapping created at the wireless controller.

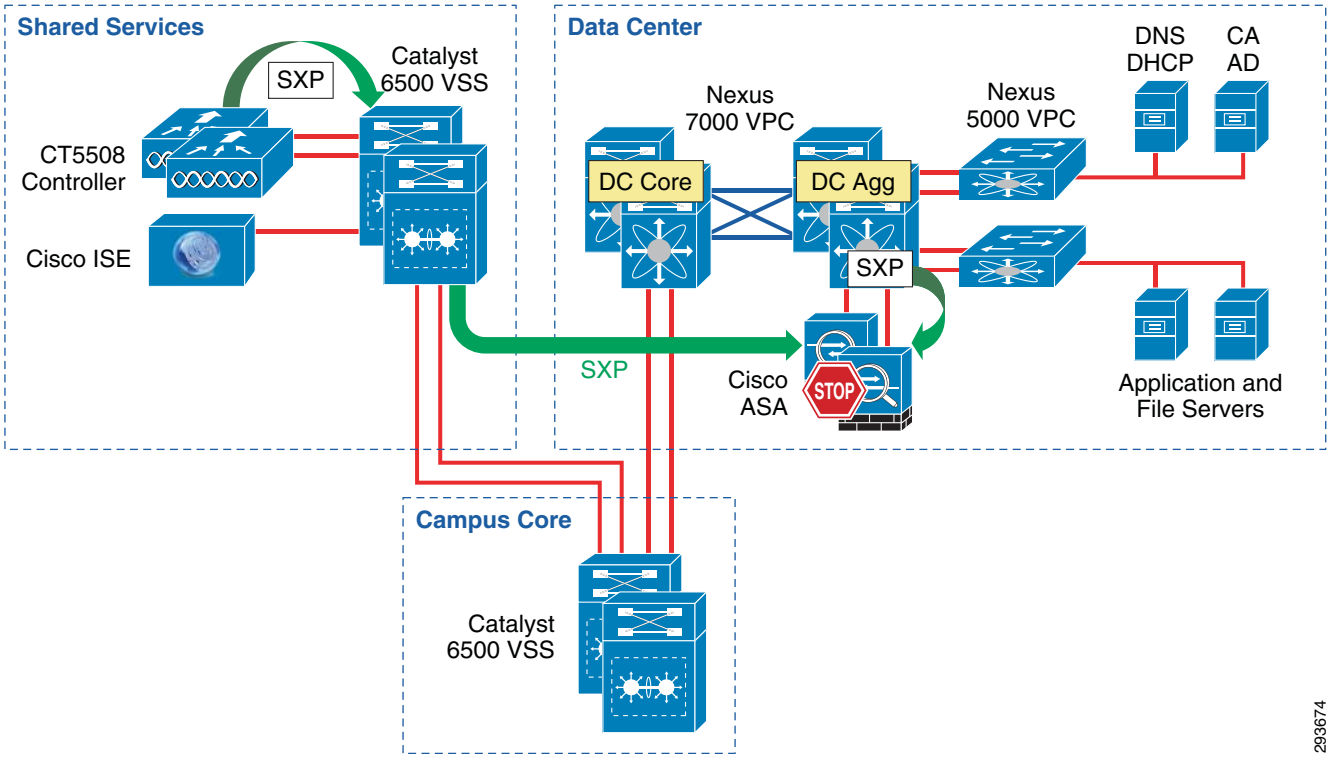
In this deployment scenario there is no requirement to configure any links for forwarding of Security Group Tags nor MACsec. Instead SXP will be used to advertise the IP/SGT mappings created dynamically at the wireless controller to the ASA configured as a Security Group Firewall (SG-FW) as well as those IP/SGT mappings statically created at the Nexus 7000 Data Center Aggregation switch. Refer to Figure 23-60.

SXP is configured on a device by identifying its peer's IP address and specifying a password for use in authenticating each side of the connection. SXP supports two modes which can be used either exclusively or combined. The first mode is that of the "Speaker" which as the name suggests advertises IP/SGT mappings. The other mode is "Listener" which also as its name suggests, listens for the Speaker's advertisements. It is possible for a device to be both a "Speaker" and a "Listener". In this scenario, both the CT5508 and the Nexus 7000 data center Agg switch will be in speaker mode while the Shared Services Catalyst 6500 acts as both a speaker to the ASA and a listener of the controller. The ASA firewall, although capable of both modes, will only be configured for listener mode, receiving the advertisements for both wireless users and servers in the data center.

Per Figure 23-60, the following SXP peering relationships will be established:

1. CT5508 to Shared Services Catalyst 6500 VSS
2. Shared Services Catalyst 6500 VSS to ASA firewall
3. Nexus 7000 Data Center Aggregation switch to ASA firewall

Figure 23-60 SXP Configuration in Deployment Scenario 2



Wireless Controller Configuration

Access the wireless controller via a web UI and follow these steps:

1. Navigate to Security > TrustSec SXP.

   The screen in Figure 23-61 appears.
2. Click the drop down arrow for “SXP State” and select Enabled.
3. Set the “default Password”. This must match that configured on its peer.
4. Click **New** (top right).
5. Fill in the IP Address (typically Loopback if possible) of the SXP Peer or the Shared Services Catalyst 6500VSS switch.
6. Click **Apply**.

Once successfully configured, the screen in Figure 23-62 should be presented upon accessing Security > TrustSec SXP. The “Connection Status” will indicate “Off” until the other device is configured.

Figure 23-61 SXP Configuration at Wireless Controller

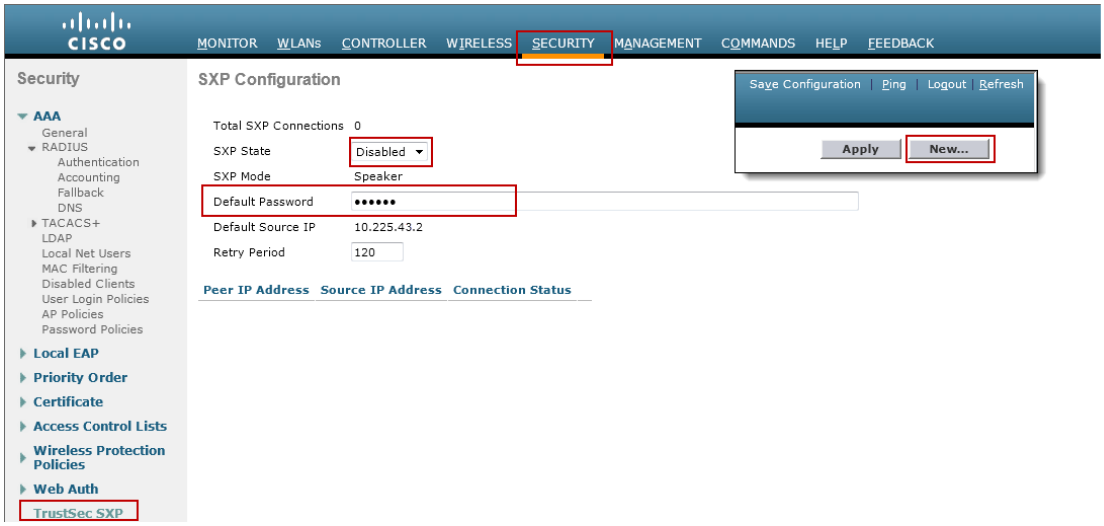
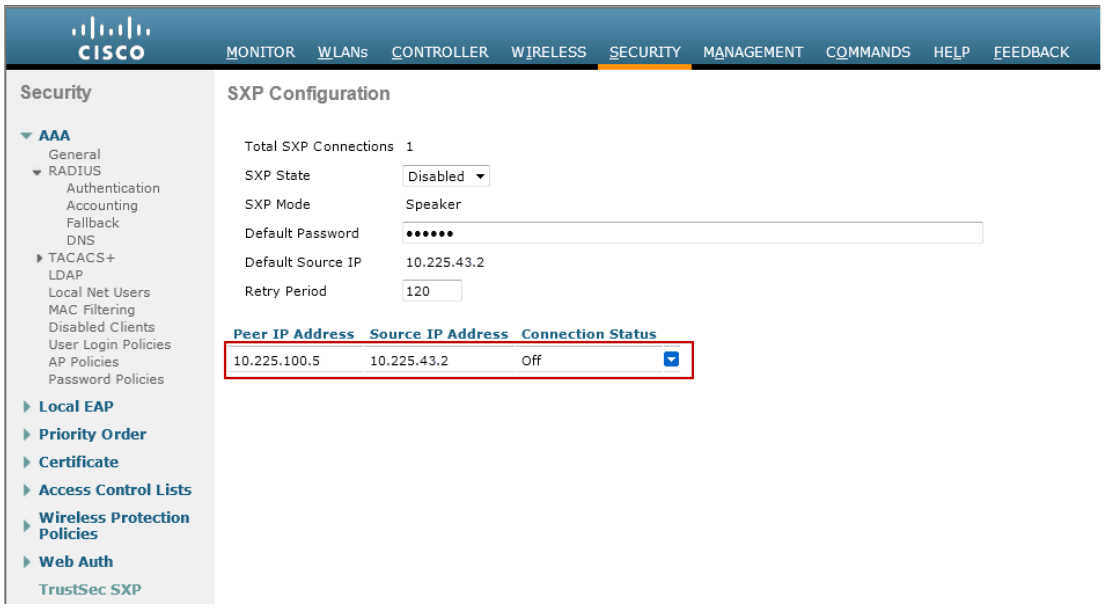


Figure 23-62 SXP Configuration Complete



Catalyst 6500 SXP Configuration

The following commands are used at the Shared Services Catalyst 6500 VSS to enable SXP peering with the CT5508 wireless controllers and the ASA firewall.

```
cts sxp enable/Enable CTS
cts sxp default source-ip 10.225.100.5/Source SXP connection from 10.225.100.5 (Lo)
cts sxp default password password/Configured password on the 6500 for incoming SXP
connections
```

```
cts sxp connection peer 10.225.43.2 source 10.225.100.5 password default mode local
listener hold-time 0 0/Builds an SXP connection to its peer, the wireless controller at
10.225.43.2 using source address of 10.225.100.5 and the default password defined above.
Specifies that this (local) device is in "Listener" mode. The source IP Address used here
is purely optional as it was specified above.
cts sxp connection peer 10.230. 3.4 source 10.225.100.5 password default mode local
speaker hold-time 0 0/Builds an SXP connection to its peer, the ASA at 10.230.3.4 using
source address of 10.225.100.5 and the default password defined above. Specifies that this
(local) device is in "Speaker" mode. The source IP Address used here is purely optional as
it was specified above.
```

Issuing the command **show cts sxp connection** results in the following output.

```
ua28-6500-1>sh cts sxp connections
 SXP:Enabled
 Highest Version Supported:4
 Default Password :Set
 Default Source IP:10.225.100.5
Connection retry open period:120 secs
Reconcile period:120 secs
Retry open timer is not running
------------------------------------------------
Peer IP:10.225.43.2<--Wireless Controller
Source IP:10.225.100.5
Conn status:On
Conn version:2
Local mode:SXP Listener
Connection inst#:1
TCP conn fd:1
TCP conn password:default SXP password
Duration since last state change:0:21:49:55 (dd:hr:mm:sec)
------------------------------------------------
Peer IP:10.230.3.4<--ASA Firewall
Source IP:10.225.100.5
Conn status:On
Conn version:2
Local mode:SXP Speaker
Connection inst#:1
TCP conn fd:2
TCP conn password:default SXP password
Duration since last state change:0:21:49:22 (dd:hr:mm:sec)

Total num of SXP Connections = 2
```

Nexus 7000 SXP Configuration

The following steps are required to configure the SXP connection between the data center aggregation switches and the ASA firewall.

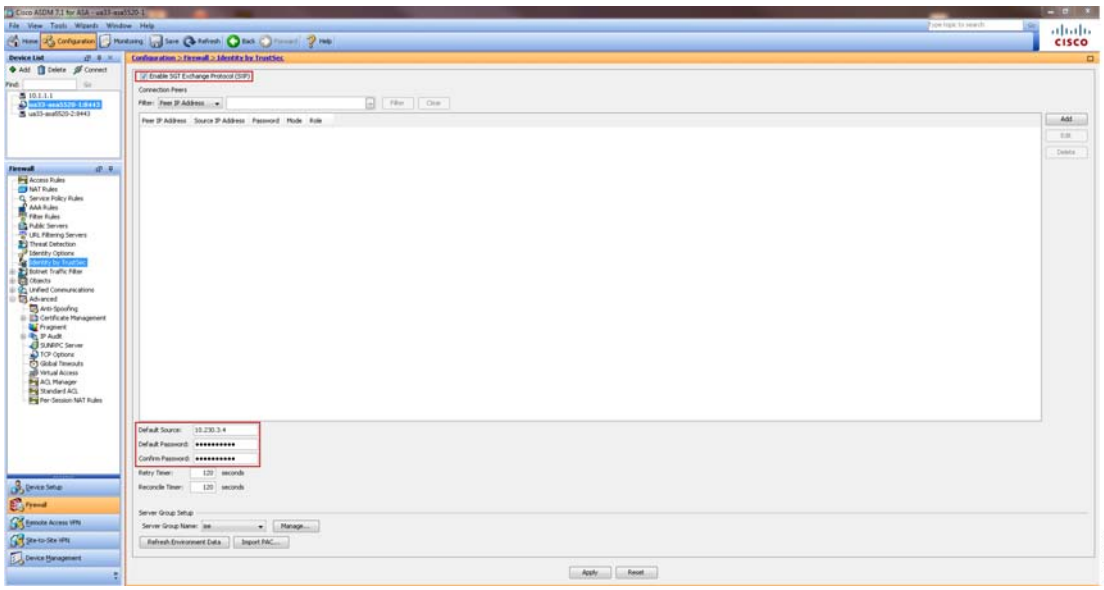
```
cts sxp enable/Enables SXP on the Nexus 7000
cts sxp connection peer 10.230.3.4 source 10.225.100.8 password required password mode
listener/Establishes an SXP connection with the ASA firewall at 10.230.2.4 sourcing this
peering from 10.225.100.8, defines the password, and identifies the peer as in listener
mode.
```

ASA SXP Configuration

The following steps are required to configure the SXP connections between the ASA firewall and Nexus 7000 Data Center Aggregation switch and the ASA firewall and Shared Services Catalyst 6500 VSS switch. Refer to Figure 23-63.

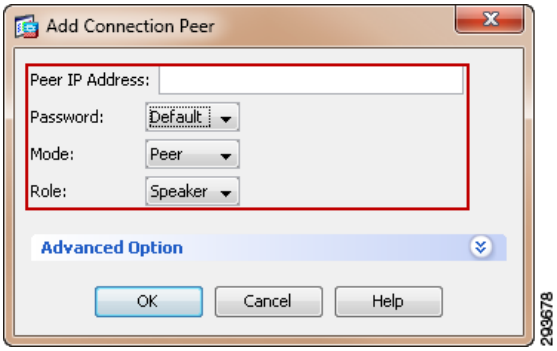
1. Using ASDM to access the ASA firewall, navigate to Configuration > Firewall > Identity by TrustSec.
2. Check the box “Enable SGT Exchange Protocol (SXP)”.
3. Enter the Default Source IP Address the firewall will use. In the CVD, the Outside Interface is the one accessible to the Nexus 7000 and Catalyst 6500 and so it was chosen. Note that this IP Address must match that previously configured on the Nexus and Catalyst switches as their peer address on the firewall.
4. Configure and confirm the password to be used to establish the secure SXP connection. Note that this password must match that used to configure the SXP connection on the Nexus and Catalyst switches previously configured.
5. Click **Add**.

Figure 23-63 Configuring SXP on the ASA



The popup window in Figure 23-64 appears when **Add** is clicked.

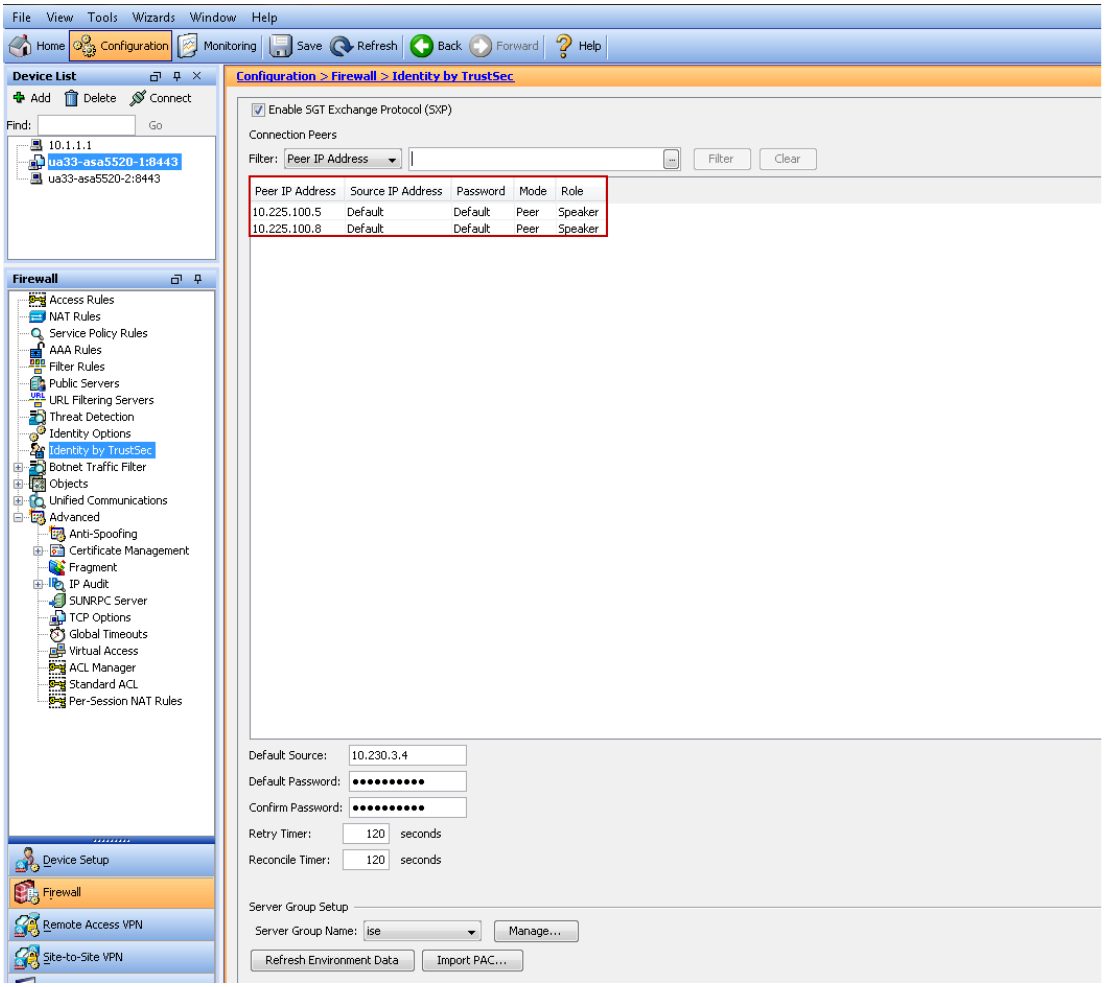
Figure 23-64 Adding SXP Peers at the ASA



6. Enter the Peer's IP Address. This must be the same one specified as the source interface at the Nexus or Catalyst switches.
7. For Password, select "Default" this is the password configured in Step 4 above.
8. For Mode select "Peer" from the drop down box.
9. For Role select "Speaker". This defines the peer as a Speaker and the ASA as a Listener.
10. Add both Shared Services Catalyst 6500 VSS and Nexus 7000 Data Center Aggregation Switches as SXP Peers.

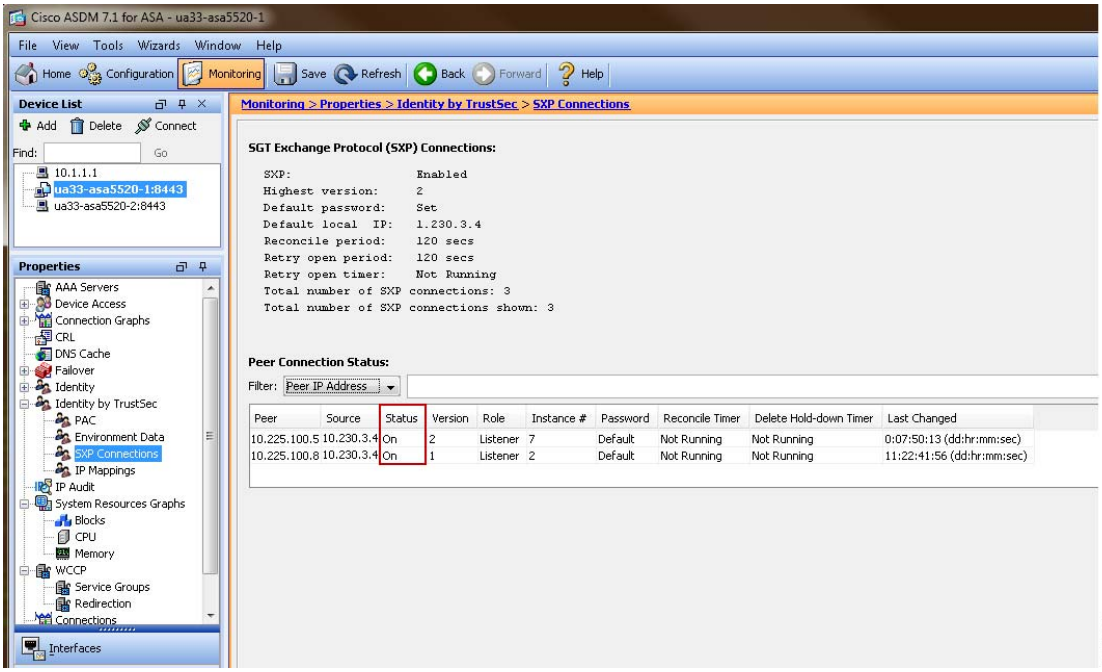
Once completed, the ASA "Identity by TrustSec" window should appear as in Figure 23-65. The two entries that were configured can now be seen in the window.

Figure 23-65 SXP Configured on ASA



In order to check the status of the SXP connections navigate to Monitoring > Properties > Identity by TrustSec > SXP Connections and the status of the connections can be seen as in Figure 23-66.

Figure 23-66 Checking Status of SXP at ASA



Configuring Static IP/SGT Bindings on Nexus Switches

Unlike campus access through Catalyst Switches and Cisco Wireless Controllers where dynamic SGT mappings are communicated and created through 802.1X and RADIUS exchange, the vast majority of organizations do not implement 802.1X for server connectivity. As such, data center switches such as the Cisco Nexus switches provide only limited support for the use of 802.1X and do not specifically support an SGT RADIUS AV as an option. Therefore, IP Address to SGT mappings will be manually defined for bare metal and virtual servers.

For purposes of this CVD we have defined our IPtoSGT Bindings at the Nexus 7000 Data Center Aggregation layer switches as depicted below. Note that as there are two Nexus 7000s composing the aggregation layer in the data center; both must be configured identically for consistent policy enforcement. The following provides an example of these:

```
cts role-based sgt-map 10.230.4.2 40/Binds 10.230.4.2 to SGT 40
cts role-based sgt-map 10.230.4.22 40/Binds 10.230.4.22 to SGT 40
cts role-based sgt-map 10.230.5.2 50/Binds 10.230.5.2 to SGT 50
cts role-based sgt-map 10.230.6.2 40/Binds 10.230.6.2 to SGT 40
cts role-based sgt-map 10.230.7.2 50/Binds 10.230.7.50 to SGT 50
```

Unlike Deployment Scenario 1 where an SGACL was statically defined to restrict SGT12 user traffic to “Unknown” or SGT0, this is unnecessary for this scenario as all policy enforcement is performed at the ASA firewall running as a Security Group Firewall (SG-FW).

To verify the IP/SGT mappings at the Nexus 7000, issue the command **sh cts role-based sgt-map**.

```
IP ADDRESSSGTVRF/VLAN       SGT CONFIGURATION
10.225.49.1540vrf:1CLI Configured
10.225.42.1540vrf:1CLI Configured
10.230.1.4540vrf:1CLI Configured
10.230.1.4640vrf:1CLI Configured
10.230.4.240vrf:1CLI Configured
```

```
10.230.4.2240vrf:1CLI Configured
10.230.5.2 50vrf:1CLI Configured
10.230.6.240vrf:1CLI Configured
10.230.7.250 vrf:1CLI Configured
```

Configuring SG-FW Role-Based Policies at ASA

In Deployment Scenario 1 Egress Policies created at ISE were used almost exclusively for policy enforcement via SGACLs dynamically pushed to the Shared Service Catalyst 6500 VSS and Nexus 7000 Data Center Aggregation Switches in the infrastructure. When using an ASA, running v9.0(2) or higher software as a Security Group Firewall, these policies must be created on the ASA through CLI or ASDM. In order to create these role-based access policies, the SG Names and Tag Values must be first downloaded to the ASA prior to use in a policy. This is accomplished through the previous configuration steps and subsequent importing of the ISE TrustSec PAC file to be used in these exchanges competed earlier.

The table in Figure 23-67 reflects the policy that will be configured at the ASA firewall. One aspect of SG-FW configuration on the ASA is that role-based policies can be created that permit or deny communications between SGT that have been defined or Unknown (SGT 0). Additionally, on the ASA, it is possible to create policies with IP Addresses or network objects as the source or destination. This offers an extremely powerful configuration capability that is different than the Catalyst or Nexus switches where SGACLs are defined using the SGT values for source and destination without support of IP Addresses.

Very much like the Catalyst and Nexus switches as discussed in the Deployment Scenario 1 section, the ASA also supports the “Unknown” SGT value of SGT 0. A key concept to be considered when granting access to the data center is that of the SGT value of zero or “Unknown” as it is referred to. If the IP Address of a server has not been mapped to an SGT at the point of IP/SGT mappings such as the Nexus 7000 in the data center, that server would be considered Unknown and associated with SGT0. Unlike SGACLs when implemented on a switching platform where an implicit permit to Unknown is permitted, the ASA or IOS ZBFW acting as a SG-FW still enforces an implicit deny. Policies can however be created on the ASA SG-FW that permit or deny access to “Unknown” from any give source SGT. and thus can be used to support a migrational approach to tag assignment in the data center.

Between the ability to use “Unknown” and IP Addresses/Network Objects in role-based policies, the ASA offers an excellent platform supporting a migrational approach to implement TrustSec in the data center.

Figure 23-67 SG-FW Policy to be Configured on ASA Firewall

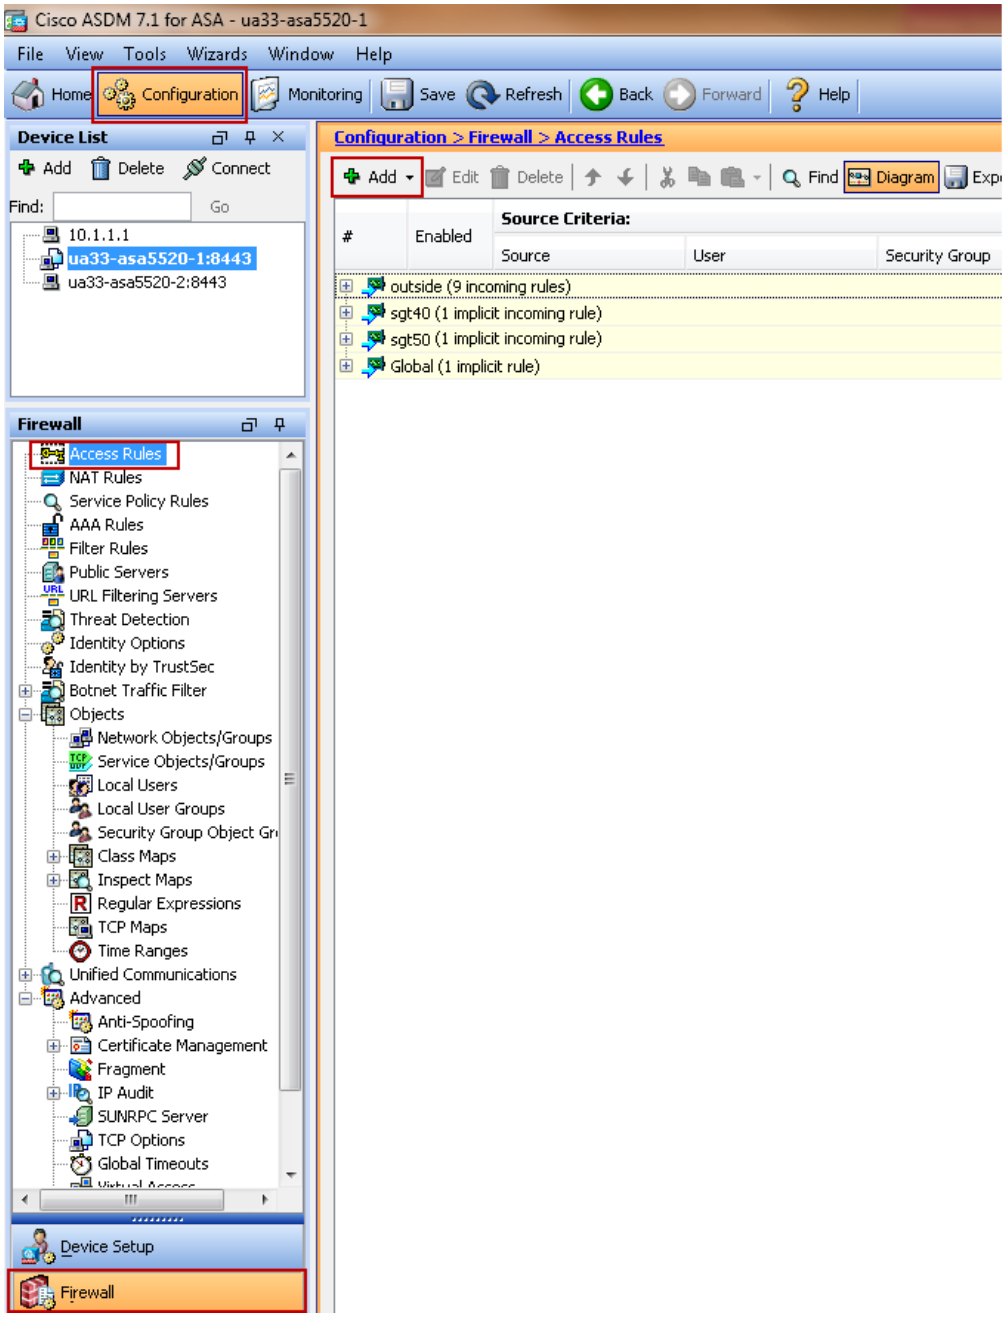
	SGT10	SGT11	SGT12	SGT40	SGT50
SGT10	✓	✗	✗	✓	✓
SGT11	✗	✓	✗	✓	✓
SGT12	✗	✗	✓	✓	✗
SGT40	✓	✓	✓	✓	✓
SGT50	✓	✓	✗	✓	✓

293631

The ASA firewall in scenario two has been configured with three Layer 3 interfaces named Outside, SGT40, and SGT50. This configuration illustrates that a migration from a policy based on IP Addresses can easily be migrated to one based on SGT with only the need to add those new policies while still enforcing existing policies.

To configure role-based policies in the ASA, use ASDM as seen in Figure 23-68.

Figure 23-68 ASDM Role-based Policy Configuration

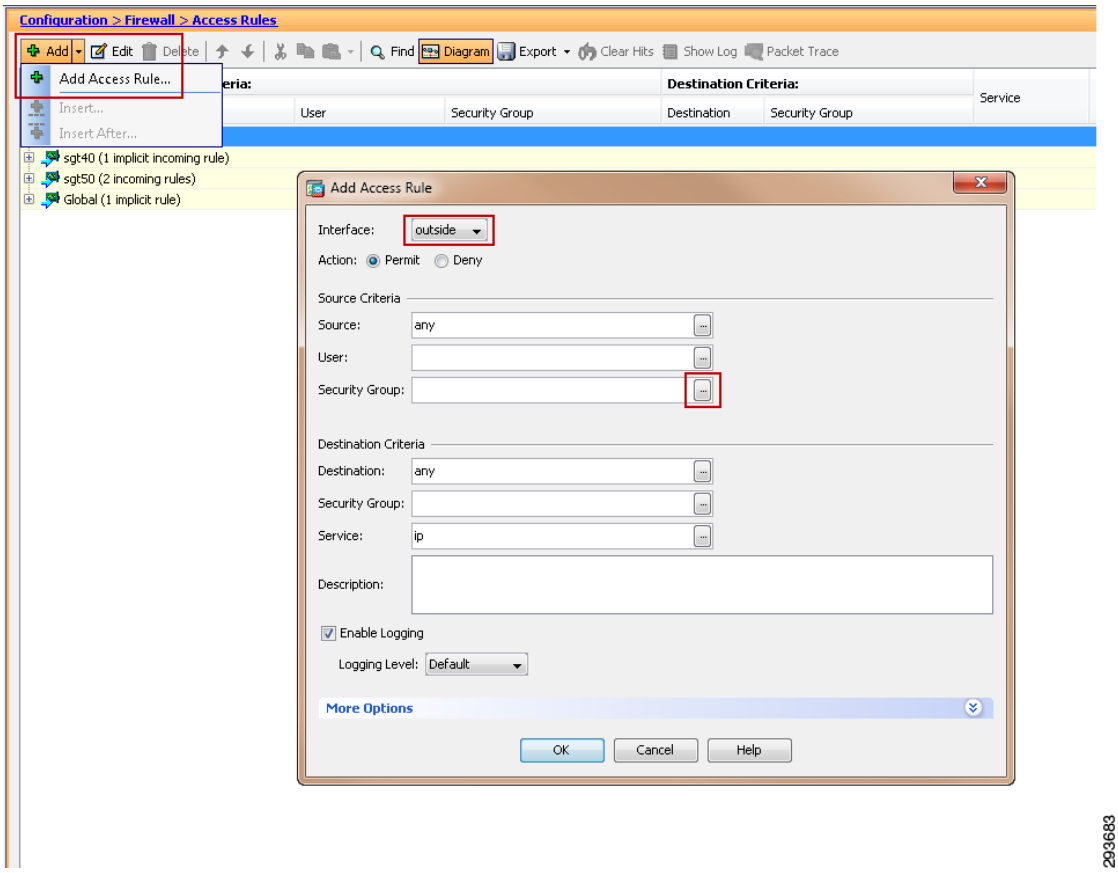


1. Navigate to Configuration > Firewall > Access Rules.

2. Highlight the desired interface to create the access rule on and click **Add**. A drop down box will open as can be seen in Figure 23-69.
3. Click **Add ACL**.

   A popup window will open up.
4. From the Interface drop-down box, select the appropriate interface. In this example it will be the "Outside" interface as this will demonstrate the creation of a policy for user access to the data center.
5. Click the browse button next to the "Source Security Group" box.

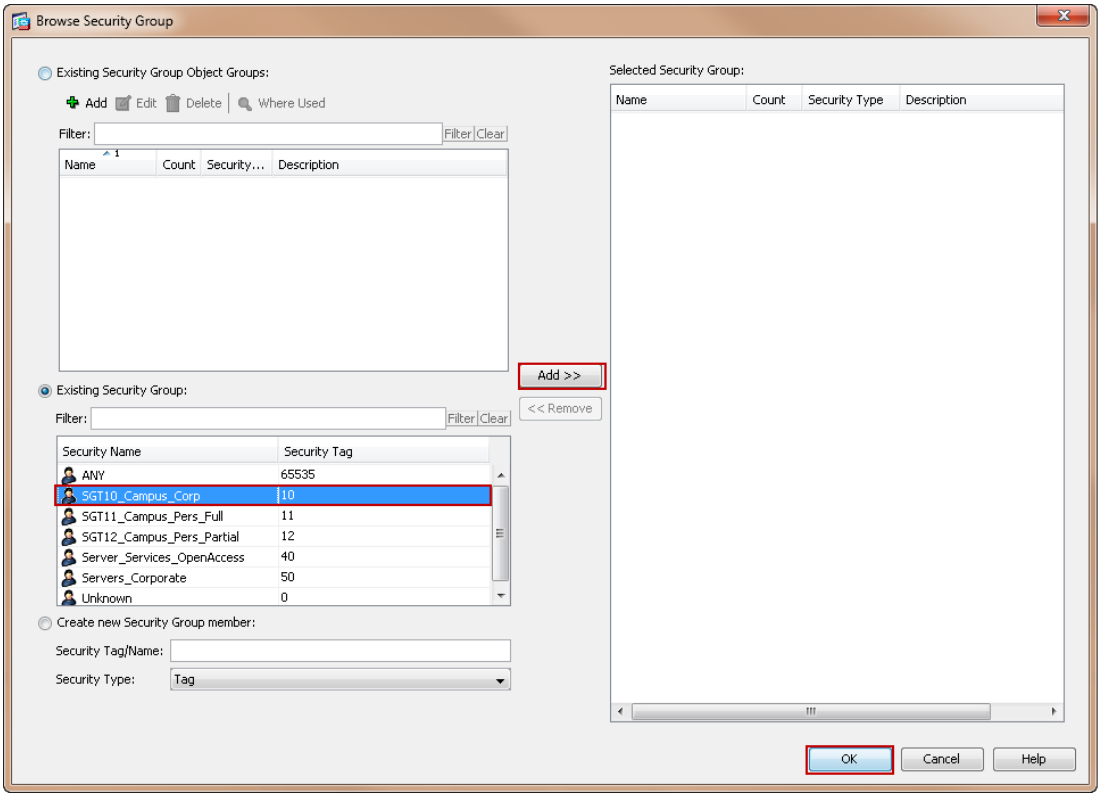
Figure 23-69 Adding an Access Rule at the ASA



A popup window opens as in Figure 23-70.

6. Select the appropriate Security Name from the "Security Group" window.
7. Click **Add** and the selected name will populate the "Selected Security Group" box on the right.
8. Click **OK**.

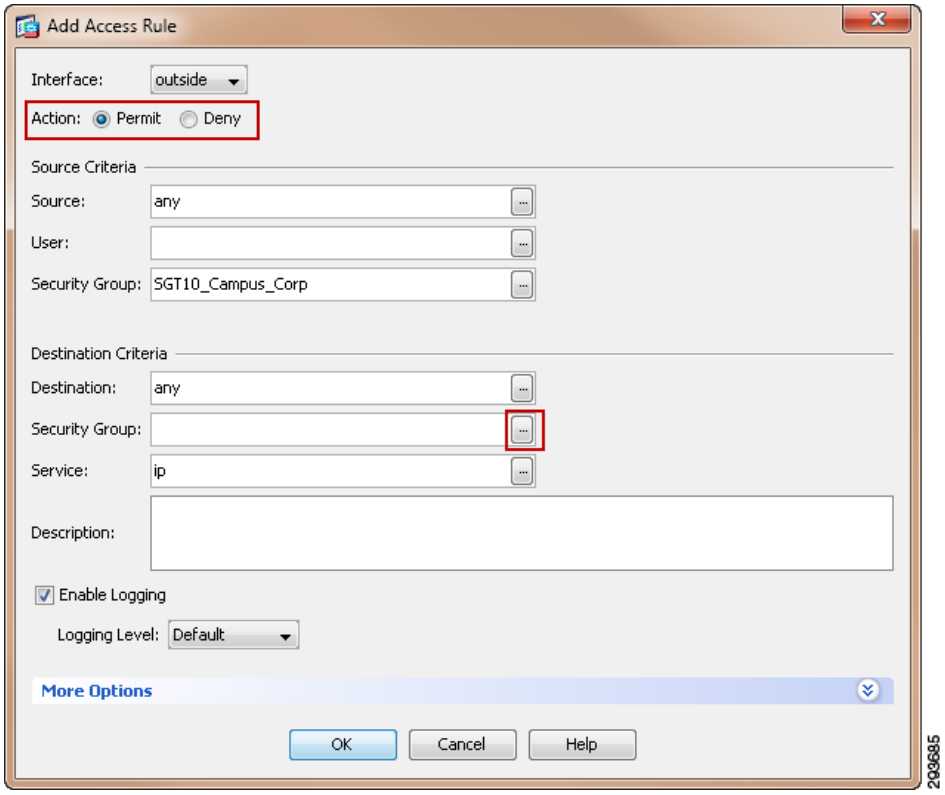
Figure 23-70 Adding a Source Group to an Access Rule at the ASA



A new window will open as depicted in Figure 23-71.

9. Select the appropriate action; Permit or Deny.
10. Click the browse button next to the “Destination Security Group” box.

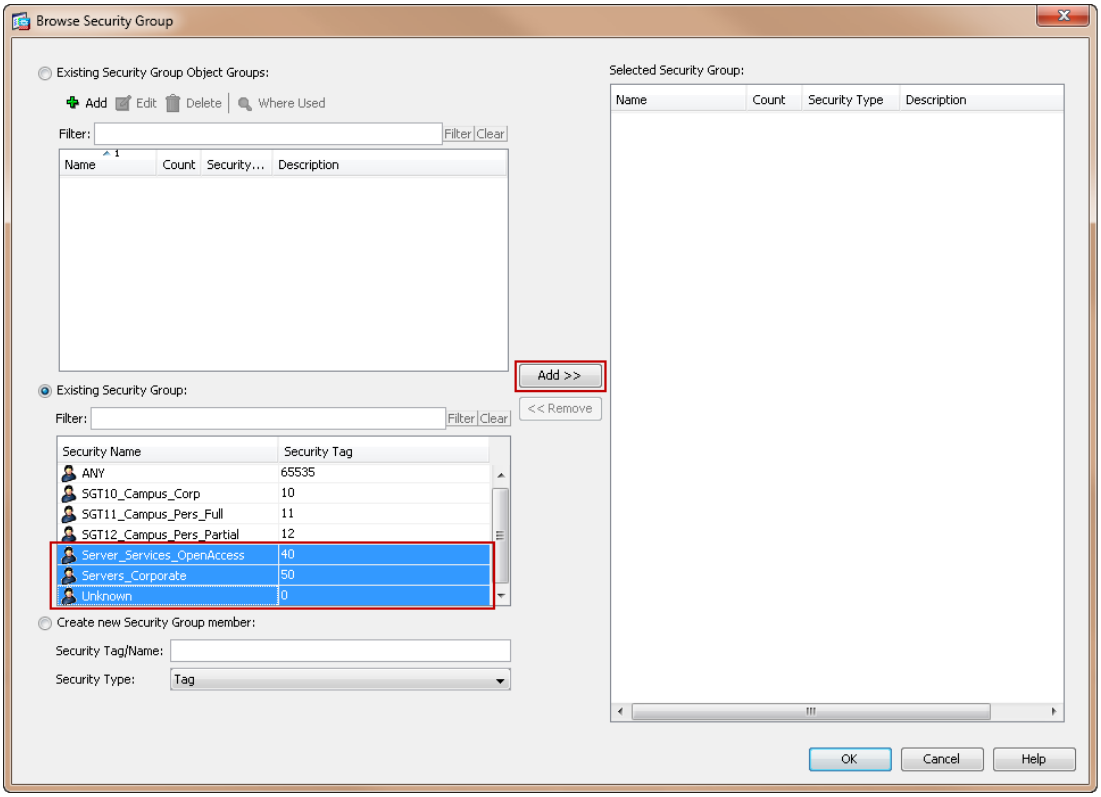
Figure 23-71 Adding Destination Group to Access Rule on ASA



A new popup window opens as in Figure 23-72.

11. Select the destination groups for the policy. In this case three groups have been selected; Server_Services_OpenAccess, Servers_Corporate, and Unknown.
12. Click **Add** and the selected name will populate the "Selected Security Group" box on the right.
13. Click **OK**.

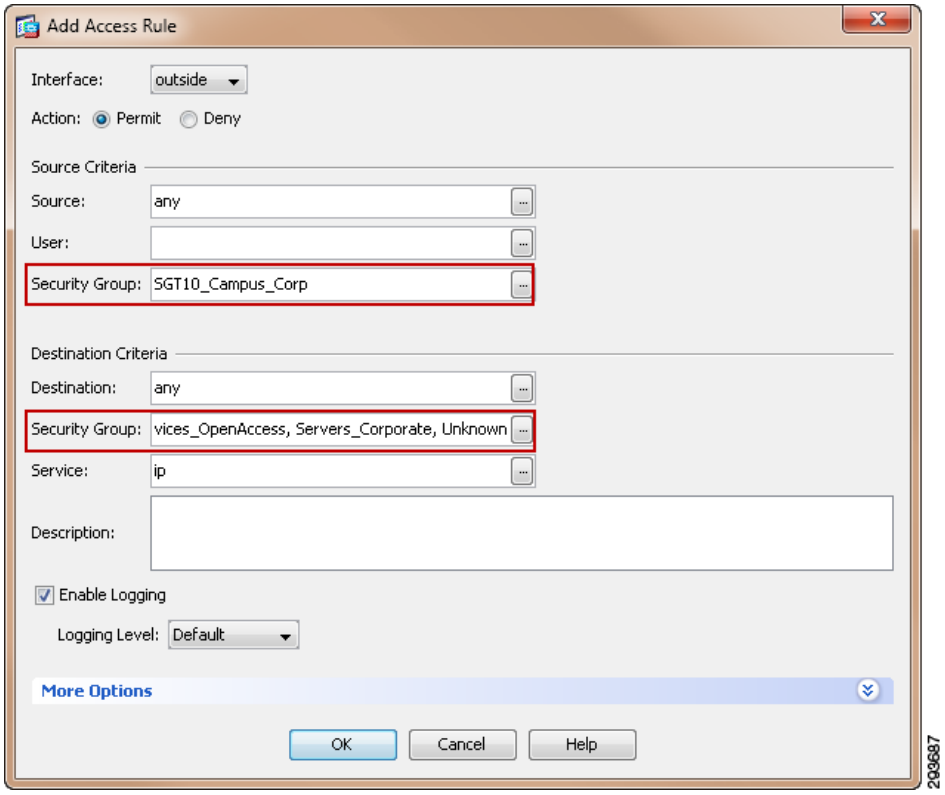
Figure 23-72 Adding a Destination Group to an Access Rule at the ASA



You will be returned to the "Add Access Rule" window and can see that Source and Destination Security Group boxes have been populated as seen in Figure 23-73.

14. Click **OK**. You will be returned to the Access Rule main window.
15. Continue adding additional policies as appropriate.

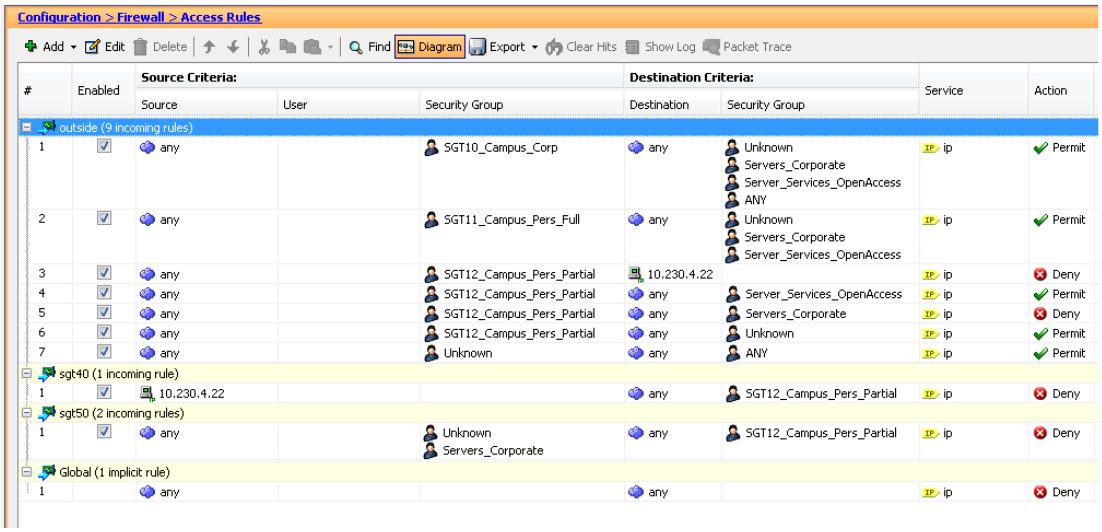
Figure 23-73 Finalizing New Access Rule



Having completed the addition of all of the necessary policies, the Access Rules will appear similar to the example in Figure 23-74 as can be seen from the example below:

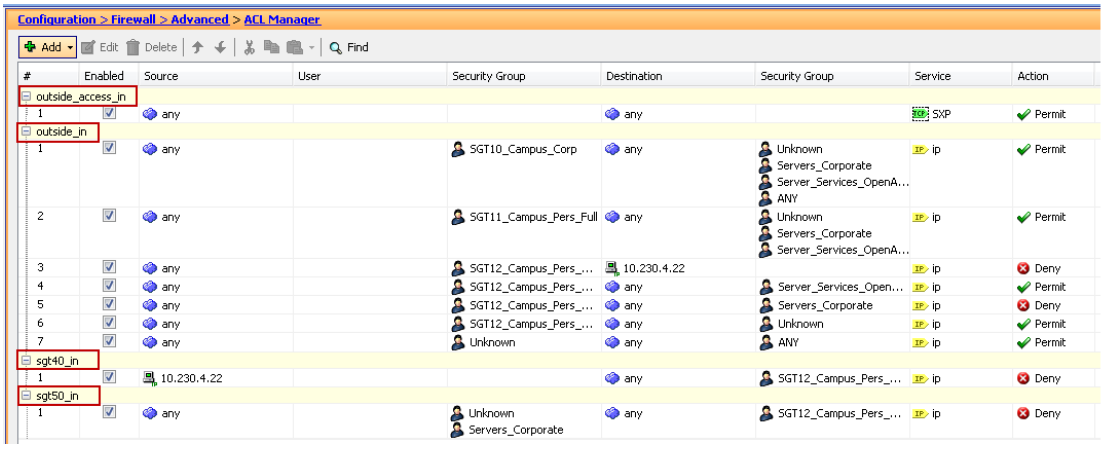
1. Outside Interface—SGT 10 can access Unknown, Servers_Corporate, Server_Services_OpenAccess, and Any. Obviously, rather than specifying the SG Names of the servers, ANY would have sufficed.
2. Outside Interface—SGT 11 can access Unknown, Servers_Corporate, and Server_Services_OpenAccess.
3. Outside Interface—SGT 12 cannot access 10.230.4.22 which is mapped to SGT 40.
4. Outside Interface—SGT 12 can access Server_Services_OpenAccess (SGT 40).
5. Outside Interface—SGT12 cannot access Unknown and Servers_Corporate.
6. Outside Interface—Devices that are not associated with an SGT can get to any server. This may be undesirable and should be examined closely prior to implementing this rule. Essentially, it permits any device to access any server internally.
7. As the SGT 40 and SGT 50 Layer 3 interfaces have a higher priority of fifty as opposed to zero for the Outside interface, an implicit permit exists for traffic sourced from SGT 40 or SGT50 to the Outside by default.
8. SGT 40 Interface—10.230.4.22 cannot access devices mapped to SGT12.
9. SGT 50 Interface—Unknown and Servers_Corporate (SGT 50) cannot access devices mapped to SGT 12.

Figure 23-74 SG-FW Access Rules



In Figure 23-75, by navigating to Configuration > Firewall > Advanced > ACL Manager, the Access List names (access-groups within the CLI configs) can be seen with the associated ACEs assigned. Of particular interest is the first ACL “outside_access_in”. This is automatically created when the SXP peering is defined to allow the session be established to the outside interface.

Figure 23-75 Access Lists with Assigned ACEs



This completes the configuration for Deployment Scenario 2. The next steps will be to configure the actual user policies as defined in the “Limited Use Case” in the CVD for corporate devices and the “Enhanced Use Case” for personal devices.

Device On-boarding, Provisioning, Authentication, and Authorization Policies for TrustSec in ISE

The final steps for configuring the infrastructure to support role-based policy enforcement involve the configuration within ISE of the various attributes and conditions required to support:

- On-boarding a device within the ISE policy server.
- Provisioning the device with the appropriate configuration and credentials to access the network.
- Define authentication and authorization profiles granting the appropriate access to the network.

ISE configuration to support corporate devices can be found in Chapter 16, "BYOD Limited Use Case—Corporate Devices."

ISE configuration to support personal devices can be found in Chapter 15, "BYOD Enhanced Use Case—Personal and Corporate Devices."



CHAPTER 24

Mobile Traffic Engineering with Application Visibility and Control (AVC)

Revised: August 7, 2013

Executive Summary

Mobile wireless traffic is soon about to exceed wired traffic on a global basis and is comprised mainly of mobile video applications. Furthermore, the devices generating wireless traffic are shifting away from traditional laptops towards smartphones and tablets. As such, the volume, composition, and device-shift of mobile application traffic can pose a challenge to network administrators tasked with ensuring their quality.

Business use cases for managing mobile applications include:

- Improving the quality of wireless voice from cellular-quality to toll-quality
- Enhancing the quality of experience for wireless video applications
- Expediting the response times for business critical data applications over wireless devices
- Managing background application traffic by preventing bulky traffic flows from monopolizing bandwidth away from more transaction-oriented flows, thus further improving user productivity
- Controlling non-business applications on wireless networks, including social networking applications, video and media-downloading applications, peer-to-peer sharing applications, gaming applications, etc.

This document presents design considerations relating to wireless (and wired) network quality management by overviewing the underlying technologies involved and showing how these interact to create an end-to-end solution.

However, the main theme of this document is detailed design guidance on best-practice Quality of Service (QoS) configurations for Cisco Wireless LAN Controllers (highlighting the new Cisco Application Visibility and Control feature) as well as for network switches to achieve the business use cases for mobile application management.

Macro Trends and Business Requirements

Mobile application traffic is continuing to explode. According to Cisco's Visual Networking Index Forecast, global mobile data traffic will grow 13-fold from 2012 to 2017.[1]

Additional key trends relating to mobile devices, applications, and traffic include:

- By 2014, wireless IP traffic will exceed wired (and will exceed 60% by 2016).[2]
- By 2016, the number of mobile-connected devices will exceed three-times the world's population.[3]
- By 2016, non-PC devices (such as smartphones and tablets) will generate 30% of all IP traffic.[4]
- By 2017, tablets will account for more than 12% of global mobile data traffic.[5]
- By 2017, mobile video will represent two-thirds of all mobile data traffic.[6]
- By 2017, 45% of global mobile data traffic will be offloaded to fixed networks via WiFi or femtocell.[7]

Therefore (non-PC) mobile devices will generate exponentially more traffic in the coming years, with the bulk of this traffic traversing wireless networks and with the traffic itself being primarily composed of video applications.

Since QoS is critical to the overall Quality of Experience (QoE) of video-based applications, network administrators need to concern themselves with ensuring that the applications traversing their networks—especially their wireless LANs (where the majority of traffic will be sourced from)—is being adequately provisioned.

1. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf
2. Cisco Visual Networking Index: Forecast and Methodology, 2011-2016 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf
3. Cisco Visual Networking Index: Forecast and Methodology, 2011-2016 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf
4. Cisco Visual Networking Index: Forecast and Methodology, 2011-2016 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf
5. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf
6. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf
7. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf

Business use-cases for wireless application quality include:

- Guaranteeing voice quality from wireless applications meets enterprise VoIP requirements. For example, independent third-party testing has shown that wireless VoIP quality over congested wireless networks can be improved from a Mean Opinion Score (MOS) of 3.92 (which is considered cellular quality) to 4.2 (which is considered toll quality) by applying the recommendations detailed later in this document.[1]
- Ensuring video applications—both interactive and streaming—are delivered to/from wireless devices with a high Quality of Experience, so that users can communicate and collaborate more efficiently and effectively—regardless of their location or device. As another example, the same testing has shown video quality to improve from Good (9 fps) to Excellent (14 fps) after respective policies were deployed.[2]
- Provisioning preferred services for business-critical applications running on wireless devices, such as Virtual Desktop applications, sales applications, customer relationship management (CRM) applications, and enterprise resource planning (ERP) applications, etc. Yet another example has shown Citrix traffic latency to decrease by a factor or 7 (from 14 ms to 2 ms) when properly provisioned over the wireless network.[3]
- De-prioritizing "background" application traffic (i.e., applications that send data to/from servers, rather than directly to other users and which do not directly impact user-productivity), such as email, file-transfers, content distribution, backup operations, software updates, etc.
- Identifying, de-prioritizing (or dropping) non-business applications, which can include social networking applications, peer-to-peer file-sharing applications and type of entertainment and/or gaming applications so that network resources are always available for business-oriented applications.

A key facilitating technology for identifying and managing application traffic over wireless networks to meet these business use case requirements is the Cisco Application Visibility and Control (AVC) feature for Cisco Wireless LAN Controllers, which is discussed next.

Cisco Application Visibility and Control (AVC) for Wireless LAN Controllers

Beginning with Cisco WLC software release 7.4, the Application Visibility and Control set of features—already supported on Cisco routing platforms, like ASR 1000s and ISR G2s—became available on WLC platforms, including the Cisco 2500, 5500, 7500, 8500 WLCs, and WiSM2 controllers on Local and Flex Modes (for WLANs configured for central switching only in 7.4 release).

The AVC feature set increases the efficiency, productivity, and manageability of the wireless network. Additionally, the support of AVC embedded within the WLAN infrastructure extends Cisco's application-based QoS solutions end-to-end.

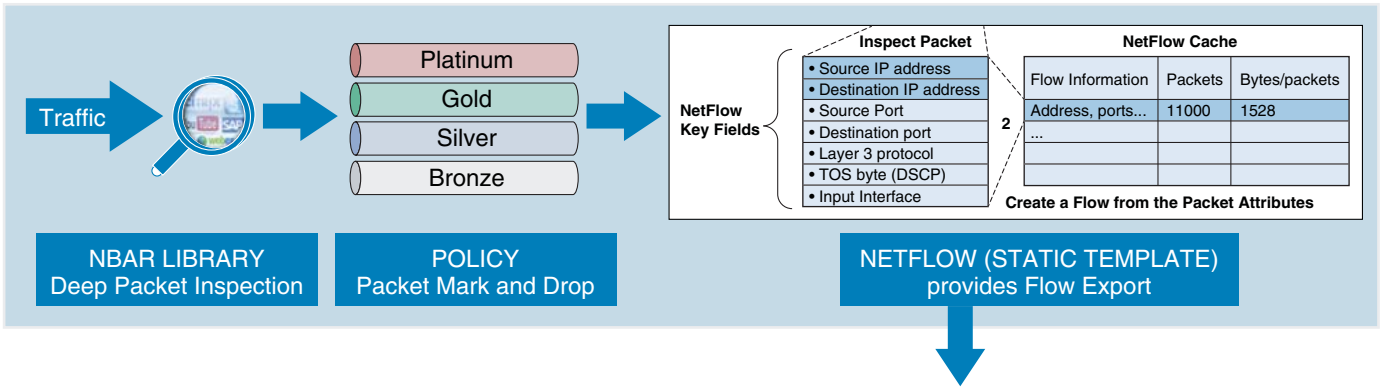
1. Syracuse University: Network Technology Performance Evaluation Cisco Application Visibility and Control (AVC) (February 1, 2013)
http://www.cisco.com/en/US/prod/collateral/wireless/cisco_avc_application_improvement.pdf
2. Syracuse University: Network Technology Performance Evaluation Cisco Application Visibility and Control (AVC) (February 1, 2013)
http://www.cisco.com/en/US/prod/collateral/wireless/cisco_avc_application_improvement.pdf
3. Syracuse University: Network Technology Performance Evaluation Cisco Application Visibility and Control (AVC) (February 1, 2013)
http://www.cisco.com/en/US/prod/collateral/wireless/cisco_avc_application_improvement.pdf

AVC includes these components:

- Next-generation Deep Packet Inspection (DPI) technology called Network Based Application Recognition (NBAR2), which allows for identification and classification of applications. NBAR is a deep-packet inspection technology available on Cisco IOS based platforms, which includes support of stateful L4-L7 classification.
- QoS—Ability to remark applications using DiffServ, which can then be leveraged to prioritize or de-prioritize applications over both the wired and wireless networks.
- A template for Cisco NetFlow v9 to select and export data of interest to Cisco Prime or a third-party NetFlow collector to collect, analyze, and save reports for troubleshooting, capacity planning, and compliance purposes.

These AVC components are shown in Figure 24-1.

Figure 24-1 Cisco AVC Components



AVC on the WLC inherits NBAR2 from Cisco IOS that provides deep packet inspection technology to classify stateful L4-L7 application classification. This is critical technology for application management, as it is no longer a straightforward matter of configuring an access list based on the TCP or UDP port number(s) to positively identify an application. In fact, as applications have matured—particularly over the past decade—an ever increasing number of applications have become opaque to such identification. For example, HTTP protocol (TCP port 80) can carry thousands of potential applications within it and in today's networks seems to function more as a transport protocol, rather than as the OSI application-layer protocol that it was originally designed as. Therefore, to identify applications accurately, Deep Packet Inspection technologies—such as NBAR2—are critical.

Once applications are recognized by the NBAR engine by their discrete protocol signatures, it registers this information in a Common Flow Table so that other WLC features can leverage this classification result. Features include QoS, NetFlow, and Firewall features, all of which can take action based on this detailed classification.

Thus AVC provides:

- Application Visibility on the Cisco WLC by enabling Application Visibility for any WLAN configured. Once Application Visibility is turned on, the NBAR engine classifies Applications on that particular WLAN. Application Visibility on the WLC can be viewed at an overall network level, per WLAN or per client.
- Application Control on the Cisco WLC by creating a AVC profile (or policy) and attaching it to a WLAN. The AVC Profile supports QoS rules per application and provides the following actions to be taken on each classified application: Mark (with DSCP), Permit (and transmit unchanged) or Drop.

Key business use cases for AVC include:

- Classifying and marking wireless mobile device applications—Identifying and differentiating realtime voice, video, or business-critical applications from less important, but potentially bandwidth-hungry-applications so as to prioritize, de-prioritize, or drop specific application traffic.
- Capacity planning and trending—Baselining the network to gain a clearer understanding of what applications are consuming bandwidth and trending application usage to help network administrators plan for infrastructure upgrades.

To better understand how AVC works in WLAN scenarios, an overview of the challenges and tools for managing quality of service over wireless media may be helpful (and is discussed next, along with a summary of the overall strategic recommendations for deploying quality of service policies across an enterprise). Network administrators already familiar with these concepts may find it more efficient to skip the following sections and to proceed directly to Configuring Downstream QoS Policies for Mobile Applications and Configuring Upstream QoS Policies for Mobile Applications.

Challenges and Solutions for Managing Application Quality over Wireless Media

To better understand design recommendations available for managing application quality over wireless media, it is beneficial to lay some context as to the challenges and solutions available for managing traffic over this media. To begin with, it should be noted that the very nature of wireless as a transmission media makes it less predictable and controllable from a quality perspective, as compared to wired networks.

For example, wired campus networks operate at full-duplex mode, with endpoints being able to transmit data at any time at maximum capacity. For example, a server connected to a switch by a 1 Gbps full-duplex link can theoretically both send and receive at 1 Gbps of data simultaneously, without having to contend with other stations for access to the medium.

Wireless networks, on the other hand, operate in half-duplex mode, with endpoints contending among themselves as well (as with the wireless access point) for the opportunity to transmit data. This is because in WLANs, every station associated to a particular access point (AP) must share the radio frequency (RF) with all the other stations; however, only one station—including the AP itself—may transmit at a given time. The result of this is that each station must contend with all the other stations for airtime. WLANs are-by definition-a multiple-access, broadcast medium, meaning that if more than one station transmits at any one time, no other station is able to understand what has been transmitted. Put another way, if two (or more) stations began transmitting data simultaneously over a WLAN, this would result in a collision. This limitation means that to avoid RF interference, full-duplex is simply not possible in WLANs (if both transmitting and receiving are performed on the same channel, as is the case in most WLAN deployments). Before quality can even be addressed over WLANs, the first requirement is finding a solution to avoiding collisions over wireless media.

IEEE 802.11 Distributed Coordination Function (DCF)

A baseline understanding of the Distributed Coordination Function (DCF)—operating at the 802.11 MAC layer (which is responsible for scheduling and transmitting Ethernet frames onto the wireless medium)—is essential to understanding the subsequent enhancements that allow for wireless quality of service.

DCF has the following key components, which are briefly described below:

- Collision Sense Multiple Access/Collision Avoidance (CSMA/CA)
- Short Interframe Space (SIFS)
- DCF Interframe Space (DIFS)
- Contention Window (CW)

Collision Sense Multiple Access/Collision Avoidance (CSMA/CA)

Wi-Fi wireless networks are completely egalitarian, meaning that all wireless stations have equal access to the medium. In fact, even the AP has no more priority to access the medium than the client stations do. For example, a wireless IP phone has to abide by exactly the same principles as a wireless laptop, regardless of the fact that one of them might be transmitting real-time VoIP traffic and the other might be transmitting peer-to-peer (P2P) traffic. Since each client, and thus each application, has an equal opportunity to transmit frames at any given time, there must be an orderly system to coordinate the transmission of packets onto the medium. If no control were implemented, there would be a high probability of a collision. Additionally, the more clients associated to the AP, the higher the likelihood of collisions occurring. Furthermore, each time a collision occurs, stations would reattempt their transmissions, likely causing additional collisions in the process.

A similar problem existed in the early days of wired Ethernet, when half-duplex links and hubs were common. In half-duplex wired Ethernet environments collisions were a common outcome of multiple end-stations trying to transmit frames onto the wire at the same time. To address this situation, a system called Carrier Sense Multiple Access/Collision Detection (CSMA/CD) was developed. CSMA/CD is a set of rules that all end stations are required to follow when trying to transmit a frame onto the medium. For example, if a collision occurred, the stations involved in the collision would follow a strict set of backoff rules, involving random timers that help to reduce the probability of a future collision next time round. CSMA/CD thus proved a relatively effective mechanism to reduce collisions on half-duplex Ethernet networks.



Note While CSMA/CD worked well enough, the problem of collisions was to be obviated altogether in wired networks by introducing switching technology, which provided dedicated collision domains to each network segment. In this manner, no endpoint contended for media access with any other endpoint, as each had a dedicated collision domain between itself and the switch and as such could then operate at full-duplex capacity.

However while it may seem that CSMA/CD might likewise be applicable to wireless networks, there is a key difference: wireless stations have no way to detect a collision.

In a wired network, transmissions are sent as bursts of energy on the wire that can be reflected back to the end stations, thus allowing accurate detection of collisions. In a wireless medium, the RF energy is shared over the air, meaning reflections of the energy wave do not come back to the sending station, thus making collision detection impossible; hence CSMA/CD is an impractical approach for WLANs.

Notwithstanding this, the IEEE modified the CSMA/CD mechanism to accommodate wireless networks, as Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA). The techniques differ in that CSMA/CD deals with what to do after a collision occurs, whereas CSMA/CA works to prevent a collision in the first place.

To understand this better, consider a person participating in an audio conference call. If many people are on the call together, there is a good chance that one person may begin talking at the same time as another, making both of them unintelligible to everyone else (effectively, a “collision”). If the parties used a CSMA/CD approach, they would pause for a few seconds and then try talking again in the hopes that they would not again talk at the same time, so that at least one of them could be understood. On the other

hand, if the parties were following a more polite CSMA/CA approach, then instead of simply starting to talk on the conference call (while hoping that no one else would at the same time), each party would wait patiently for a quiet period. Then, when they were certain that no one else was talking, they would begin. Additionally, the other parties would recognize and respect that one party was talking and would remain silent until they had completed speaking (without interruption).

Thus CSMA/CA opts to listen to the channel first to see if any transmissions are in progress and only when the channel is free does it attempt to send its frame. If a collision does occur even after listening and waiting, the wireless station deals with it in a similar way to CSMA/CD, by waiting for a random backoff period before it tries to resend the frame.

However it is important to note that CSMA/CA can never fully guarantee that a collision won't occur; rather, it reduces the probability that a collision will occur by trying to "avoid" a future collision. CSMA/CA is a bit like arriving in your car at a four-way stop at the same time as three other drivers. Although you might try very hard to avoid a collision by looking both ways very carefully before driving into the intersection, you can never fully guarantee what the other driver will do. If you make a decision to proceed, there is always a slight possibility that the other driver might do the same thing at the same time and thus there is always the possibility of a collision. The same goes for WLAN stations that operate using CSMA/CA.

Short Interframe Space (SIFS)

So how does a sending station know that its transmission succeeded? Since, due to the broadcast nature of the wireless medium, collisions cannot be detected (they can only be avoided with a measure of probability), there is an obvious need to confirm whether a transmission was successful. To solve this problem, DCF ensures that each frame is acknowledged once the transmission is successfully received. Specifically, there is a provision in DCF where all clients keep silent after a transmission finishes so the receiving station has a chance to send the acknowledgement. This is period is called the Short Interframe Space (SIFS). This ensures that the transmitting station knows that it does not need to retransmit and it can move on to its next frame.

DCF Interframe Space (DIFS)

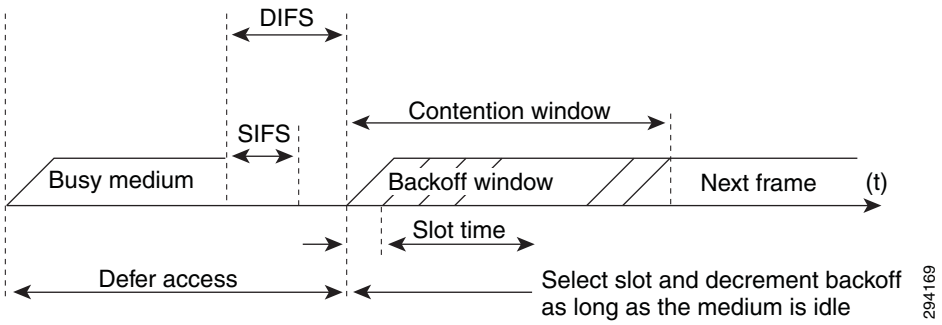
To help control and organize the transmission of frames on the wireless medium, DCF uses some clever rules whereby the contending stations wait for different periods of time before they can transmit their frames onto the channel. A central and key concept to how DCF operates is the DCF Interframe Space (DIFS). DIFS is a pre-established, fixed wait timer observed by all stations before they attempt transmission of a frame onto the channel.

Note There are actually several Interframe Space (IFS) types used in 802.11 networks; however, for the purposes of this discussion attention is focused only on SIFS and DIFS.

As mentioned, CSMA/CA provides a framework of "listen before you talk" for wireless stations. When a wireless station wants to transmit a frame, the first thing it does is to wait the appropriate DIFS time. Once this DIFS countdown has finished, if the medium is still clear, it transmits. DIFS is like a level set for all stations that want to transmit. If they all just started transmitting as soon as they had a frame in the queue, collisions would be plentiful. By waiting the DIFS period it gives a chance for the station to confirm that the channel is indeed clear for transmission.

Figure 24-2 shows the operation of Interframe Spaces (SIFS and DIFS) within DCF.

Figure 24-2 Interframe Spaces Operation



Contention Window (CW)

However, it may be the case that after waiting for the DIFS period to expire, the DCF process detects the medium is not idle. If this is the case, the station waits (technically speaking the station will "defer") for a random period of time, called the contention window (CW). The first time a station needs to defer, the CW random backoff period is set from 0 to a maximum value known as CWmin. There is an obvious advantage to waiting a random period of time, for if multiple stations are all trying to transmit at the exact same time because a collision occurred and they all backed off for the same length of time, collisions would continually occur. After the CW timer expires the station again looks to see if the medium is free and if it is, it begins transmission.

However, if after the CW timer expires the sending station detects that the medium is still not clear, then it will:

- Defer again until the wireless medium is finally clear.
- Wait again for the DIFS period.

  Then:

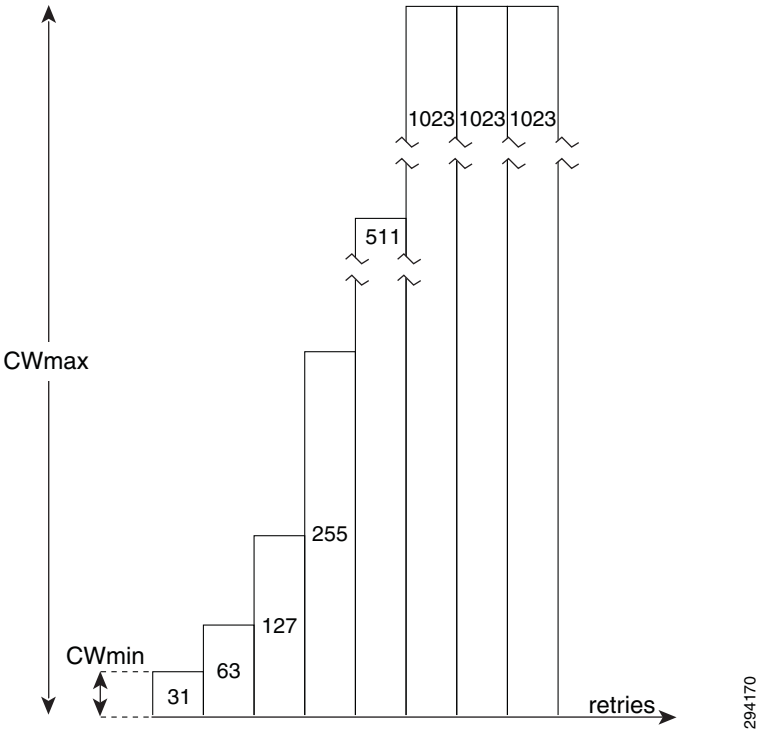
- Wait for another (longer) backoff period.

At first, the station doubles the CW value it used previously. However, if the station continually finds that the medium is not clear, then the station will continue to double the backoff window each time it tries to send the frame. It keeps on doing this up to a maximum amount of time, known as the CWmax value. This continues until it either it transmits the frame or the Time to Live (TTL) expires.

The amount of time that the station counts down is not actually measured in seconds, but rather in slot times. The slot time is a time value derived from the RF characteristics of the radio network and so it is unique for each network (but the actual length of these times is in microseconds, with 802.11 specifying about 20 microseconds for a slot time). As an example, in the case of IEEE 802.11n, the CWmin default is 15 slot times (~300 µs or 0.3 ms) and CWmax is 1023 slot times (~20,460 µs or 20.46 ms).

Due to the variable nature of contention windows, it is easy to see how significant amounts of delay variation (jitter) can be introduced into the packet flows. Figure 24-3 illustrates Contention Window Operation.

Figure 24-3 Contention Window Operation

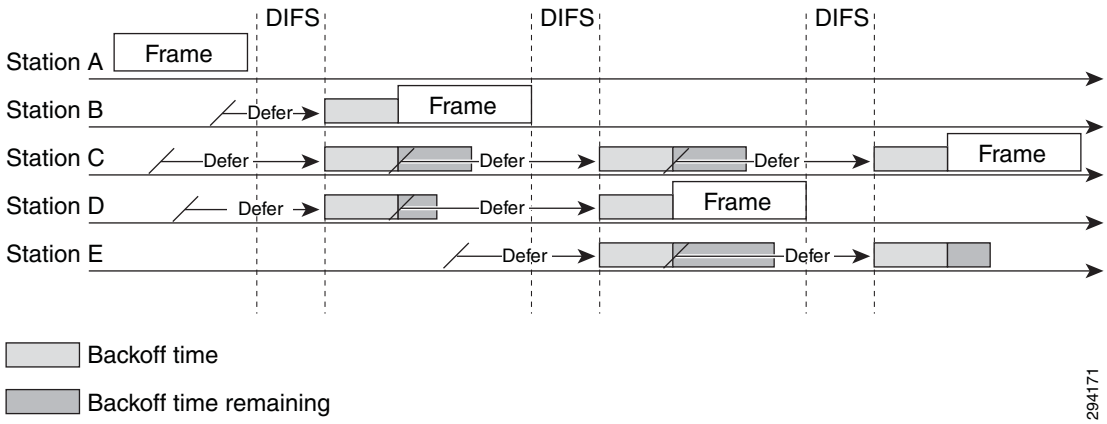


To illustrate how this works, if the random backoff for a station was initially set to 10 slot times, the second backoff would be 20. If the channel is still busy, the CW backoff is increased to 40, then 80, then 160, then 320, and finally 640 (as 640 doubled would exceed the CWmax value of 1023 slot times). At this point, if the frame has not been sent, the process begins again. These retries continue until the packet TTL is reached. This process of doubling the backoff window is referred to as binary exponential backoff.

DCF Operation

Consider an example of the DCF process might apply to a real world contention scenario. As illustrated in Figure 24-4, five stations associated to the same AP and all are trying to send data at approximately the same time.

Figure 24-4 DCF Operation



The DCF operation steps illustrated in Figure 24-4 are as follows:

Step 1 Station A successfully sends a frame; three other stations also want to send frames, but must defer to Station A traffic.

Step 2 After Station A completes the transmission, all the stations must still defer to the DIFS. When the DIFS is complete, stations waiting to send a frame can begin to decrement the backoff counter, once every slot time, and can send their frame.

Step 3 The backoff counter of Station B reaches zero before Stations C and D, and therefore Station B begins transmitting its frame.

Step 4 When Station C and D detect that Station B is transmitting, they must stop decrementing the backoff counters and defer until the frame is transmitted and a DIFS has passed.

Step 5 During the time that Station B is transmitting a frame, Station E receives a frame to transmit, but because Station B is sending a frame, it must defer in the same manner as Stations C and D.

Step 6 When Station B completes transmission and the DIFS has passed, stations with frames to send begin to decrement the backoff counters. In this case, the Station D backoff counter reaches zero first and it begins transmission of its frame.

Step 7 The process continues as traffic arrives on different stations.

IEEE 802.11e/WMM

As can be seen from the previous section's analysis of the DCF model, there is no provision to differentiate service levels by traffic types, thus quality assurance is not possible using legacy DCF. To address this (and other limitations) the IEEE 802.11e task group provided enhancements to the original 802.11 specification that recommended several modifications to the way DCF operates that would facilitate a differentiated services model. These 802.11e modifications to the DCF model have been rolled into the wider 802.11-2007 standard (which is essentially a retrofit of the original 802.11 specification). The IEEE 802.11e model was also certified by the Wi-Fi Alliance as the Wireless Multimedia (WMM) model; as such these terms generally refer to the same mechanisms and are used interchangeably in this paper.

The 802.11e task group has provided many enhancements to the overall 802.11 specification, but the key goal was to introduce an intelligent system of application traffic differentiation on wireless radio interfaces. Two of the most significant changes proposed by this task group were:

- To support marking within wireless frames by supporting a 3 bit marking value known as 802.11e User Priority (UP); 802.11e UP is essentially the same as 802.1p CoS marking, but for wireless frames (as opposed to wired Ethernet tagged frames).
- To replace DCF with a new MAC layer protocol known as Enhanced Distributed Channel Access (EDCA), which introduces the concept of relative prioritization by giving different application traffic varied access levels to the wireless media.

However, a critical point to understand about wireless QoS tools is that—unlike wired QoS tools—these can only offer a greater probability of one traffic type being differentiated from another, not an absolute guarantee of it.

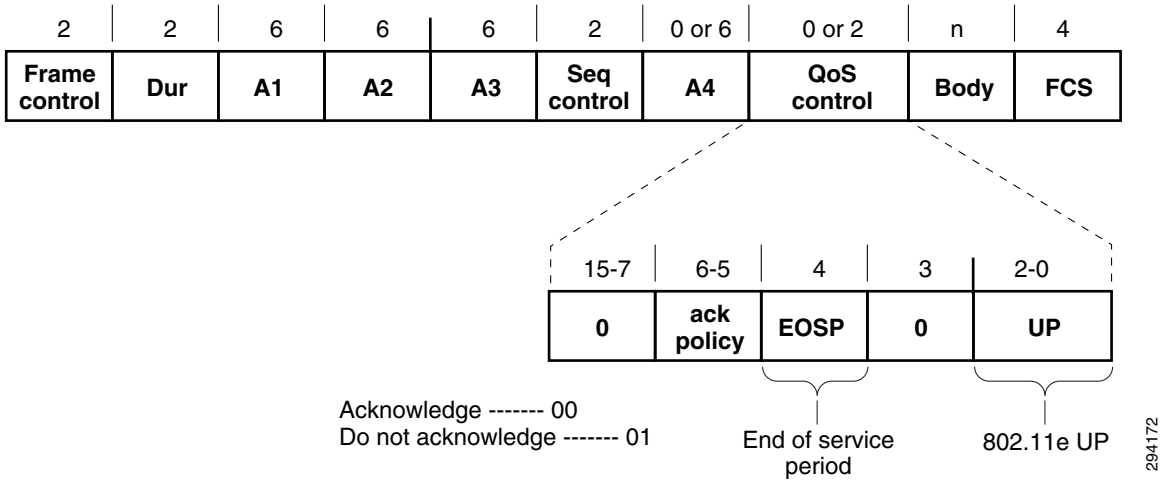
IEEE 802.11e/WMM introduced major enhancements over DCF, which in turn enables a QoS toolset over Wi-Fi networks. Each of these enhancements is briefly described:

- 802.11e User Priorities
- Access Categories (AC)
- Enhanced Distributed Coordination Function (EDCF)
- Arbitration Interframe Spacing (AIFS)
- Contention Window Enhancements
- Transmission Opportunity (TXOP)
- Call Admission Control (TSpec)

802.11e User Priorities

IEEE 802.11e/WMM introduced a new frame format, shown in Figure 24-5, which includes support for a 3 bit marking field—compatible with 802.1D priority and 802.1p CoS—that is referred to as 802.11e User Priority (UP).

Figure 24-5 WMM Frame Format and 802.11e UP



Access Categories (AC)

IEEE 802.11e/WMM specifies four different access categories, which are:

- Voice (AC_VO)
- Video (AC_VI)
- Best-Effort (AC_BE)
- Background (AC_BK)

IEEE 802.11e also supports a default mapping of User Priority markings to these access categories; these mappings are shown in Table 24-1.

Table 24-1 IEEE 802.11e/WMM Access Categories, Mappings and Designations

Relative Priority	802.11e UP	802.11e Access Category (AC)	WMM Designation	Cisco WLC Designation
Highest	7	AC_VO	Voice	Platinum
	6			
	5	AC_VI	Video	Gold
	4			
	3	AC_BE	Best Effort	Silver
Default	0			
	2	AC_BK	Background	Bronze
Lowest	1			

Note An important item to note regarding the 802.11e/WMM AC model shown in Table 24-1 is that several CoS values in this WMM model do not line up with their IETF DSCP counterparts. For example, voice is mapped to a 802.11e UP value of 6. However, in wired networks, voice is typically marked 802.1p CoS value of 5, as this corresponds to the three Most Significant Bits (MSB) of the IETF recommended (six-bit) DSCP marking value for voice traffic of EF/46 (based on RFCs 3246[1] and 4594[2]). The root cause of this marking incompatibility is that the IETF defines Layer 3 marking standards (i.e., DSCP), while the IEEE defines Layer 2 standards (like 802.1p CoS and 802.11e UP). Unfortunately, this sometimes leads to confusion and inconsistencies with the marking schemes that must be dealt with by the network engineer. These mapping considerations are discussed in greater detail later.

Note There is no relative priority between UP or CoS markings assigned to a single access category over the WLAN; for example, flows marked with either UP or CoS values of 4 and 5 are assigned to the video/gold WMM access category, but there is no difference in treatment between these flows.

Note While the WMM uses specific application designations for these access categories (Voice, Video, Best Effort, and Background), Cisco WLC software uses a more generic designation based on precious metals: platinum, gold, silver and bronze. Nonetheless each set of names refers to the same underlying

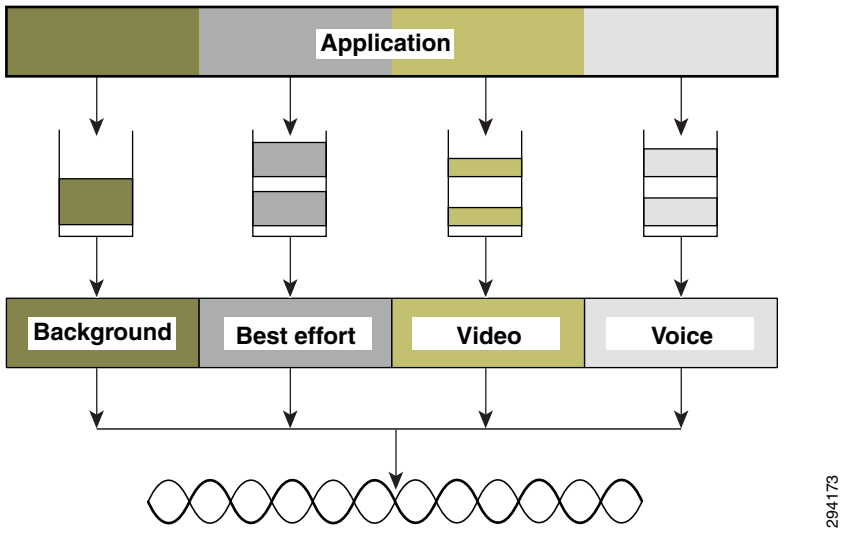
1. An Expedited Forwarding PHB (Per-Hop Behavior) http://www.ietf.org/rfc/rfc3246
2. Configuration Guidelines for DiffServ Service Classes http://www.ietf.org/rfc/rfc4594

access categories. In this paper, the latter (precious-metal-based) designations are used to simplify mapping examples and reduce confusion when describing RFC 4594-based application classes versus WMM access categories.

Enhanced Distributed Coordination Function (EDCF)

Figure 24-6 shows the queuing performed on a WMM client or AP. There are four separate queues, one for each of the access categories. Each of these queues contends for the wireless channel in a similar manner to the DCF mechanism described previously, but with each of the queues using different interframe spaces and with different CWmin and CWmax values. If frames from different access categories collide internally, the frame with the higher priority is sent and the lower priority frame adjusts its backoff parameters as though it had collided with a frame external to the queuing mechanism. This system is called the Enhanced Distributed Coordination Function (EDCF).

Figure 24-6 WMM Queues



Arbitration Interframe Spacing (AIFS)

One of the key limitations of DCF is that the DIFS value is the same for all traffic types. The rule to remember here is that once the channel is declared available, all stations wanting to transmit must wait the DIFS time period following the end of the current station's transmission. The problem is that if there are multiple stations waiting to transmit, they all have to wait the exact same DIFS gap, regardless of how latency-sensitive their data is, thus giving no preferential treatment to either high or low priority traffic.

To address this, EDCA introduces a variable interframe spacing (IFS) period for data and management frames, called the Arbitration Interframe Spacing (AIFS) number. The intention of assigning different IFS values to each AC is that the higher-priority ACs are assigned shorter wait times as compared to the lower-priority ACs. This approach thus gives the high-priority traffic a much better probability of being transmitted first.

While AIFS numbers are configurable, the default values defined in EDCF (as measured in slot times) are shown in Table 24-2.

Table 24-2 EDCA Default AIFS Numbers

Access Category	AIFS (Slot Times)
Voice	2
Video	2
Best Effort	3
Background	7

Contention Window Enhancements

DCF gives no preferential treatment to high-priority traffic during the CW, meaning that all traffic types have the same statistical probability of being the next one to transmit. Therefore an additional enhancement EDCA introduced is to give preferential CW random backoff ranges for higher priority traffic, thus making it much more likely for a voice or video frame to be transmitted before a best-effort or background frame. This is particularly important for latency sensitive traffic, such as voice and video, which suffers greatly if they have to wait up to the CWmax interval. Similar to the different AIFSN values assigned to the different priority queues, different contention window values serve to give higher priority traffic a better probability of having to wait a shorter period of time before having a chance to transmit and also limit the impact of long CW wait times. The default EDCA CW values for 801.11a/g/n are shown in Table 24-3.

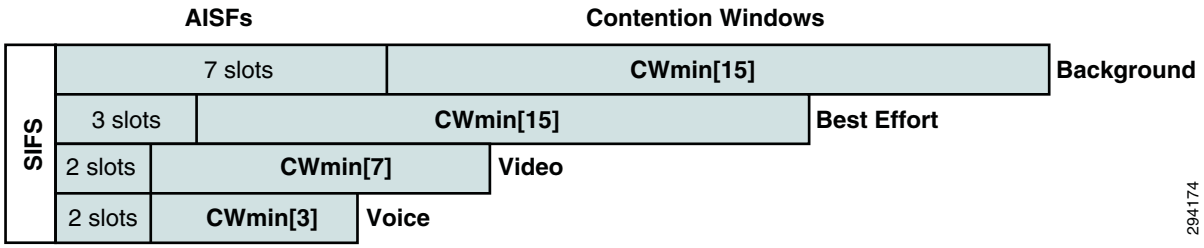
Table 24-3 EDCA/WMM Default Contention Window Values

Access Category	CWmin (Slot Times)	CWmax (Slot Times)
Legacy DCF (for comparison)	15	1023
Voice	3	7
Video	7	15
Best-Effort	15	1023
Background	15	1023

Table 24-3 shows that voice only backs off between 3-7 slot times compared to background traffic (which is still the same as the legacy DCF CW backoff period). Of course, since the CW is randomly generated, there is still a small probability that the lower priority queues might backoff for a shorter period than a higher priority queue; however, over time the higher-priority queues are statistically serviced much more often.

Figure 24-7 shows how AIFS and per-AC CW backoff timers work together to improve the overall handling of the four WMM access categories. In this example, the voice queue waits for five slot times before attempting to send its data onto the channel (2 AIFS slots + a randomly generated CW of 3 slots), thus resulting in a significantly improved probability that the voice traffic is sent over the air before anything else.

Figure 24-7 AIFS and CW Operation for Access Category Priority



Transmission Opportunity (TXOP)

EDCF provides contention-free period access to the wireless medium, called the Transmission Opportunity (TOXP). The TXOP is a set period of time when a wireless station may send as many frames as possible without having to contend with other stations. In the legacy DCF model once a station has access to the medium, it is able to keep sending frames as long as it wants. When a low data-rate station gains access to the medium, it forces all other stations to wait until it is finished its transmission.

With EDCA's TXOP enhancement, each station has a set TXOP time limit where it can transmit. Once the TXOP limit expires, it must give up access to the medium.

Call Admission Control (TSpec)

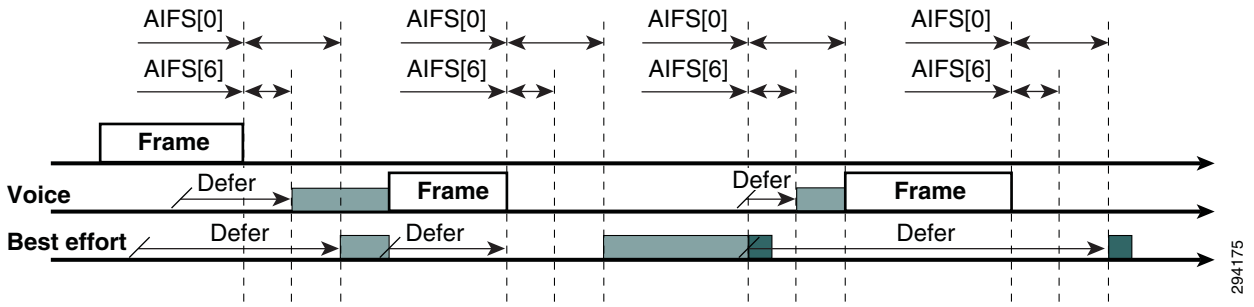
One last major enhancement introduced by 802.11e is a mechanism for Call Admission Control (CAC) called Transmission Specification (TSpec). TSpec allows real-time applications, such as voice calls or video calls in-progress, to be prioritized over requests for new calls. To use this feature of EDCF, TSpec must be configured on the AP and optionally on the client stations.

When running TSPec, a client station signals its traffic requirements (mean data rate, power save mode, frame size, etc.) to the AP. In this way, before a client sends traffic of a certain priority type (AC), it must first request permission via the TSpec mechanism. For example, a WLAN client device wanting to use the voice AC must first make a request for use of that AC to see if there is sufficient space on the network to do so. If the AP decides there is insufficient availability on the network, it denies access for that client station, thus protecting the currently transmitting stations.

EDCF Operation

With all the elements combined, EDCF operation highlighting application-level QoS over wireless media, is presented in Figure 24-8. In this example voice traffic is contending with best effort.

Figure 24-8 EDCF Operation



The EDCF process follows this sequence:

Step 1 While Station X is transmitting its frame, other stations determine that they must also send a frame. Each station defers because a frame was already being transmitted, and each station generates a random backoff.

Step 2 Because the Voice station has a traffic classification of voice, it has an Arbitrated Interframe Space (AIFS) of 2, and uses an initial CWmin of 3, and therefore must defer 5 slot times total before attempting to transmit.

Step 3 Best-effort has an AIFS of 3 and a longer random backoff time, because its CWmin value is 15.

Step 4 Voice has the shortest random backoff time, and therefore starts transmitting first. When Voice starts transmitting, all other stations defer.

Step 5 After the Voice station finishes transmitting, all stations wait their AIFS, then begin to decrement the random backoff counters again.

Step 6 Best-effort then completes decrementing its random backoff counter and begins transmission. All other stations defer. This can happen even though there might be a voice station waiting to transmit. This shows that best-effort traffic is not starved by voice traffic because the random backoff decrementing process eventually brings the best-effort backoff down to similar sizes as high priority traffic, and that the random process might, on occasion, generate a small random backoff number for best-effort traffic.

Step 7 The process continues as other traffic enters the system.

Cisco's Strategic Approach to Application Quality Management

Having reviewed the QoS tools and mechanisms available for managing application quality over WLANs, the administrator is faced with having to make the decisions as to which applications to assign to each of the four WMM access categories available, as well as how to interconnect the QoS policies for the WLAN with the rest of the network so as to create an end-to-end solution.

When making these decisions, it is best to take a step back from the tools and technical elements of the equation and examine the business or organization needs. Focusing on the tools exclusively can be compared to going to a hardware store, coming across a handy tool, and then going home and trying to decide what to build with it. Conversely, a better approach is to understand what needs to be built and then finding the right tool(s) to achieve that objective.

Therefore a network administrator needs to first define the business and organizational objectives to be addressed with QoS policies across their end-to-end network, which includes defining:

- Which applications are viewed as critical for achieving business/organizational objectives?
- What are the respective service-level requirements of these critical applications?
- What applications are present over the network that consume resources away from critical applications and which could be deprioritized?
- How many unique classes of service need to be provisioned in order to meet these per-application service level requirements and thus the business objectives?

Cisco recommends a strategic approach to application quality management in an end-to-end manner that encompasses all Places-in-the-Network (PINs), including the WLAN. This approach is based on IETF RFC 4594.

Cisco's Strategic Application-Class QoS Recommendations

Cisco's strategic approach to application QoS (which is based on IETF RFC 4594) is summarized in Figure 24-9.[1]

Figure 24-9 Cisco's (RFC4594-based) Strategic Application-Class QoS Recommendations

Application Class	Per-Hop Behavior	Admission Control	Queuing and Dropping	Application Examples
Voice	EF	Required	Priority Queue (PQ)	Cisco IP Phones (G.711, G.729
Broadcast Video	CS5	Required	(Optional) PQ	Cisco IP Video Surveillance/Cisco Enterprise TV
Realtime Interactive	CS4	Required	(Optional) PQ	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Jabber, Cisco WebEx
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoDs)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Signaling	CS3		BW Queue	SCCP, SIP, H.323
Ops/Admin/Mgmt (OAM)	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	VDI Apps, ERP Apps, CRM Apps, Database Apps
Bulk Data	AF1		BW Queue + DSCP WRED	E-mail, FTP, Backup Apps, Content Distribution
Best Effort	DF		Default Queue + RED	Default Class
Scavenger	CS1		Min BW Queue (Deferential)	YouTube, iTunes, BitTorrent, Xbox Love

294176

1. Enterprise Medianet Quality of Service Design 4.0-Overview http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSIntro_40.html#wp61104

Note Cisco has adopted RFC 4594 as its general DiffServ QoS strategy with the following exception: Cisco has swapped the marking recommendations of the RFC 4594 Broadcast Video class (of CS3) with the Signaling class (of CS5). Therefore Cisco has decided to mark Broadcast Video traffic as CS5 and Signaling traffic as CS3. This is primarily because Cisco has been marking Signaling to CS3 for over a decade (well before RFC 4594 was even drafted) and lacking a compelling business case to change the defaults on all its voice and video-telephony products, has decided to continue doing so. Furthermore, such a marking change would correspondingly force their customer base to change all their network QoS policies relating to the Signaling class as well. Therefore, Cisco has swapped these marking recommendations. It is important to remember that RFC 4594 is an informational RFC and not a standard and, as such, compliance-in full or in part-is not mandatory.

As shown in Figure 24-9, an enterprise may be required to support up to twelve application classes of that have unique service level requirements:

- Voice—This service class is intended for VoIP telephony (audio media only traffic-VoIP signaling traffic is assigned to the "Signaling" class). Traffic assigned to this class should be marked EF. This class is provisioned with an 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). Example traffic includes 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 ("inelastic" refers to flows that 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. Example traffic includes 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. An example 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 (AF41) and should be provisioned with a guaranteed bandwidth queue with DSCP-based Weighted-Random Early Detect (DSCP-WRED) enabled. Admission to this class should be controlled; additionally, traffic in this class may be subject to policing and re-marking. Example applications include Cisco Jabber and Cisco WebEx.
- 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. Example applications include Cisco Digital Media System Video-on-Demand (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, as network control traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes EIGRP, OSPF, BGP, HSRP, IKE, etc.

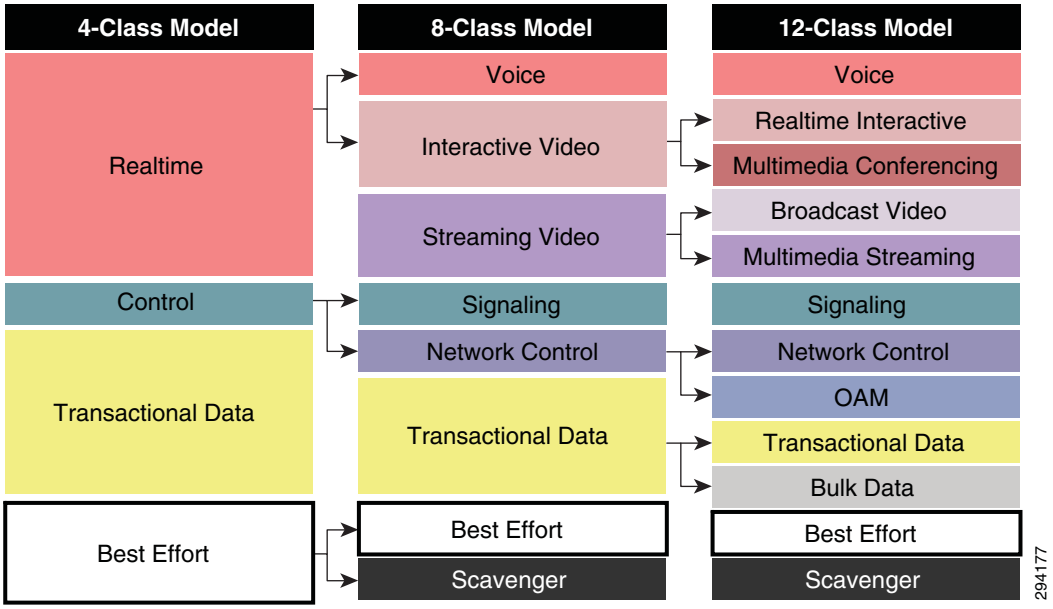
- Signaling—This service class is intended for signaling traffic that supports IP voice and video telephony. Traffic in this class should be marked CS3 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as signaling traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SCCP, SIP, H.323, etc.
- Operations/Administration/Management (OAM)—As the name implies, 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, as OAM traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SSH, SNMP, Syslog, etc.
- 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—in order 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. Example applications include data components of multimedia collaboration applications, Virtual Desktop Infrastructure (VDI) applications, Enterprise Resource Planning (ERP) applications, Customer Relationship Management (CRM) applications, database applications, etc.
- 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—in order 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. Example applications include: email, backup operations, FTP/SFTP transfers, video and content distribution, etc.
- Best Effort (the 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. Example traffic includes YouTube, Facebook, Xbox Live/360 Movies, iTunes, BitTorrent, etc.

Application-Class Expansion

While there are merits to adopting a 12-class model, Cisco recognizes that not all enterprises are ready to do so, whether this be due to business reasons, technical constraints, or other reasons. In fact, most businesses deploying QoS are typically somewhere between a 4 and 8 class model today.

Therefore, rather than considering these QoS recommendations as an all-or-nothing approach, Cisco recommends considering a phased approach to application class expansion, as illustrated in Figure 24-10.

Figure 24-10 Application-Class Expansion Models



By considering such a phased approach to application class expansion network administrators can incrementally implement QoS policies across their infrastructures in a progressive manner, in line with their evolving business needs. Nonetheless, at each phase of QoS deployment, the enterprise needs to clearly define their business objectives to determine how many traffic classes will be required at each phase.

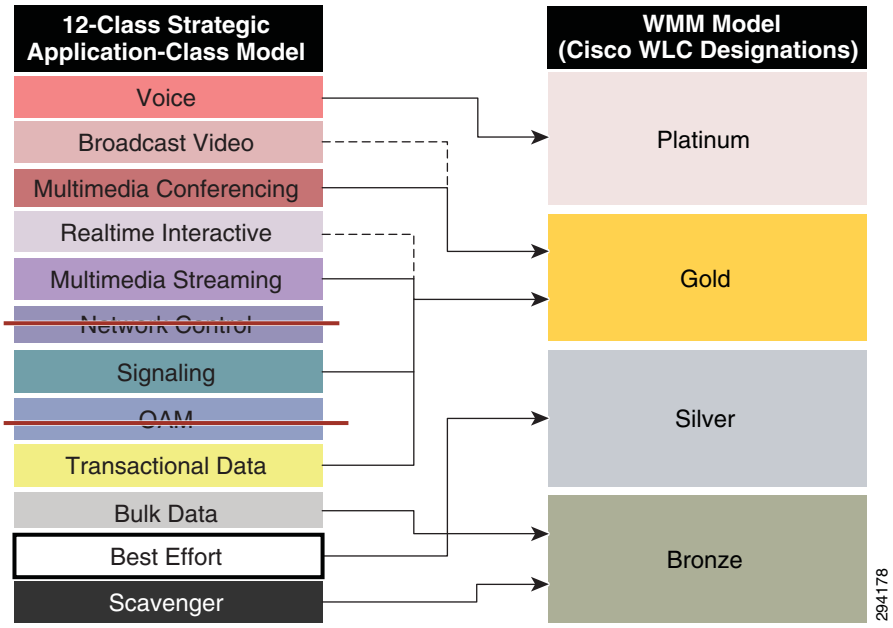
Application-Class Mapping to WMM

As previously mentioned, it is important to base the overall network QoS strategy on business or organizational requirements—first and foremost—and not by the number of classes of service supported on a platform, place-in-the-network, or service provider. This is because platforms, technologies, and service provider models all change over time, yet the QoS strategy should only change as the business evolves. Furthermore, any strategic class-of-service model can be mapped to a reduced number of service classes, as required.

For example, it has already been shown that in the IEEE802.11e/WMM model there are four access categories supported on the WLAN. This should not be taken to mean that an enterprise should never deploy more than four application classes. Rather, they should define their overall end-to-end strategy based on their organizational requirements and then map into these four access categories at the WLAN.

Even highly complex models—such as an RFC 4594-based 12-class model—can be mapped into the four WMM access categories, as illustrated in Figure 24-11.

Figure 24-11 Example Twelve-Class Application-Class Model Mapped into WMM



Note Figure 24-11 serves only to illustrate the concept of application class mapping into a reduced set of traffic classes; additional details are provided later in this document to show best-practice recommendations in mapping a 4-Class model, an 8-Class model and this same 12-Class enterprise model into WMM.

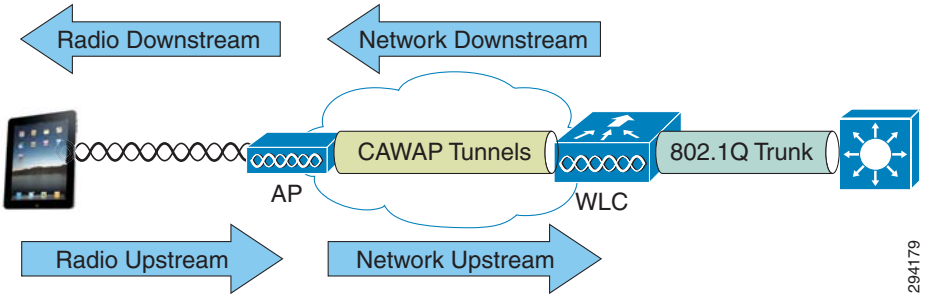
Cisco 802.11e/802.1p/DSCP Mappings

At this point, it may be helpful to clarify some terms that will be used extensively throughout the rest of this paper:

- Downstream—Used to refer to the flow of packets from the wired network infrastructure to the wireless devices, including:
  - Network Downstream—Refers to traffic leaving the WLC traveling to the AP; this traffic is encapsulated within LWAPP. Wired campus QoS policies provision downstream QoS.
  - Radio Downstream—Refers to traffic leaving the AP and traveling to the WLAN clients. WMM provides downstream QoS for WLAN clients.
- Upstream—Used to indicate the flow of packets from the mobile wireless device to the wired network infrastructure, including:
  - Radio Upstream—Refers to traffic transmitted by the WLAN clients and traveling to the AP. WMM provides upstream QoS for WLAN clients.
  - Network upstream—Refers to traffic leaving the AP, traveling to the WLC; this traffic is encapsulated within LWAPP. Wired campus QoS policies provision upstream QoS.

These terms are illustrated in Figure 24-12.

Figure 24-12 Downstream and Upstream QoS



In order for campus wired networks to interoperate with WLAN networks, their respective markings need to be translated or mapped in either direction of flow (upstream and downstream).

Default DSCP-to-CoS/UP Mappings

By default, 6-bit DSCP values are mapped to 3-bit 802.1p CoS and 802.11e UP values by taking the three Most-Significant Bits (MSB) of the DSCP and copying these as the CoS and/or UP values. For example, DSCP EF/46 (binary 101110) is mapped to CoS or UP 5 (binary 101), by default. For example, by default, the network switch that connects to the Cisco WLC will generate 802.1p CoS values (for the 802.1Q trunked traffic) by setting these to match the three MSB of the DSCP values.

Conversely, in the reverse direction, the CoS or UP values are simply multiplied by 8 (in order to shift these three binary bits to the left) to generate a DSCP value. Continuing the example, CoS or UP 5 (binary 101) would be mapped (i.e., multiplied by 8) to DSCP 40 (binary 101000), also known as CS5.

As can be seen in the above pair of examples, because information is being truncated from 6-bits to 3-bits, marking details can get lost in translation. In this example, the original voice packet was sent with DSCP EF, but was received as DSCP CS5 (based solely on default Layer 3/Layer 2 mapping). This needs to be taken into account when mapping from wired-to-wireless and vice-versa.

Cisco WLC/AP QoS Translation Table

As has already been pointed out in the consideration of Table 24-1, IEEE 802.11e and 802.1p application marking values do not always align with IETF-based DSCP-to-CoS mappings. For example, DSCP EF/46 is recommended by the IETF for use for voice, which would map by default to CoS/UP 5; but the IEEE designates CoS/UP 6 for voice. Similarly, the IETF recommends DSCP CS4 or AF4 for realtime or interactive video conferencing, both of which would map by default to CoS 4; but the IEEE designates CoS/UP 5 for video.

In an effort to reconcile the markings recommendations between these independent and disagreeing standards bodies, Cisco has implemented an automatic mapping function within WLC software to automatically convert special marking values to the respective IETF or IEEE marking recommendations, as shown in Table 24-4.

Table 24-4 Cisco WLC/AP DSCP-to-UP Translation Table[1]

Application Class	IETF DSCP	IEEE 802.11e UP	WLC QoS Profile
Network control	56 (CS7)	7	Platinum
Internetwork control	48 (CS6)	7	Platinum
Voice	46 (EF)	6	Platinum
Multimedia Conferencing	34 (AF41)	5	Gold
Multimedia Streaming	26 (AF31)	4	Gold
Transactional Data	18 (AF21)	3	Silver
Bulk Data	10 (AF11)	2	Bronze
Best Effort	0 (BE)	0	Silver

1. Cisco Wireless LAN Controller WLAN Configuration Guide, Release 7.4 - Working with WLANs - Assigning QoS Profiles http://www.cisco.com/en/US/partner/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01010111.html

Residual DSCP-to-WMM Mappings

The IEEE 802.11e UP value for DSCP values that are not mentioned in the Table 5 are calculated by considering 3 MSB bits of DSCP. Thus with the exceptions noted in Table 5, DSCP values will map to the WMM/WLC Access Categories shown in Table 24-5.

Table 24-5 Default DSCP, UP and WLC Access Category Mappings (Excluding the Exceptions Listed in Table 24-4)

DSCP Range	IEEE 802.11e UP	WLC QoS Profile
DSCPs 56-63	7	Platinum
DSCPs 48-55	6	
DSCPs 40-47	5	Gold
DSCPs 32-39	4	
DSCPs 24-31	3	Silver
DSCPs 0-7	0	
DSCPs 16-23	2	Bronze
DSCPs 8-15	1	

WLC AVC/QoS Profile-to-DSCP Mappings

If QoS or AVC Profiles are created on a Cisco WLC and applied to WLANs, then the packets assigned to the access-categories within these profiles will be marked to the DSCP values shown in Table 24-6. For example, if an application is assigned to the Gold profile, then all application traffic will be marked to DSCP 34.

Table 24-6 Default WLC Profile to DSCP Mappings

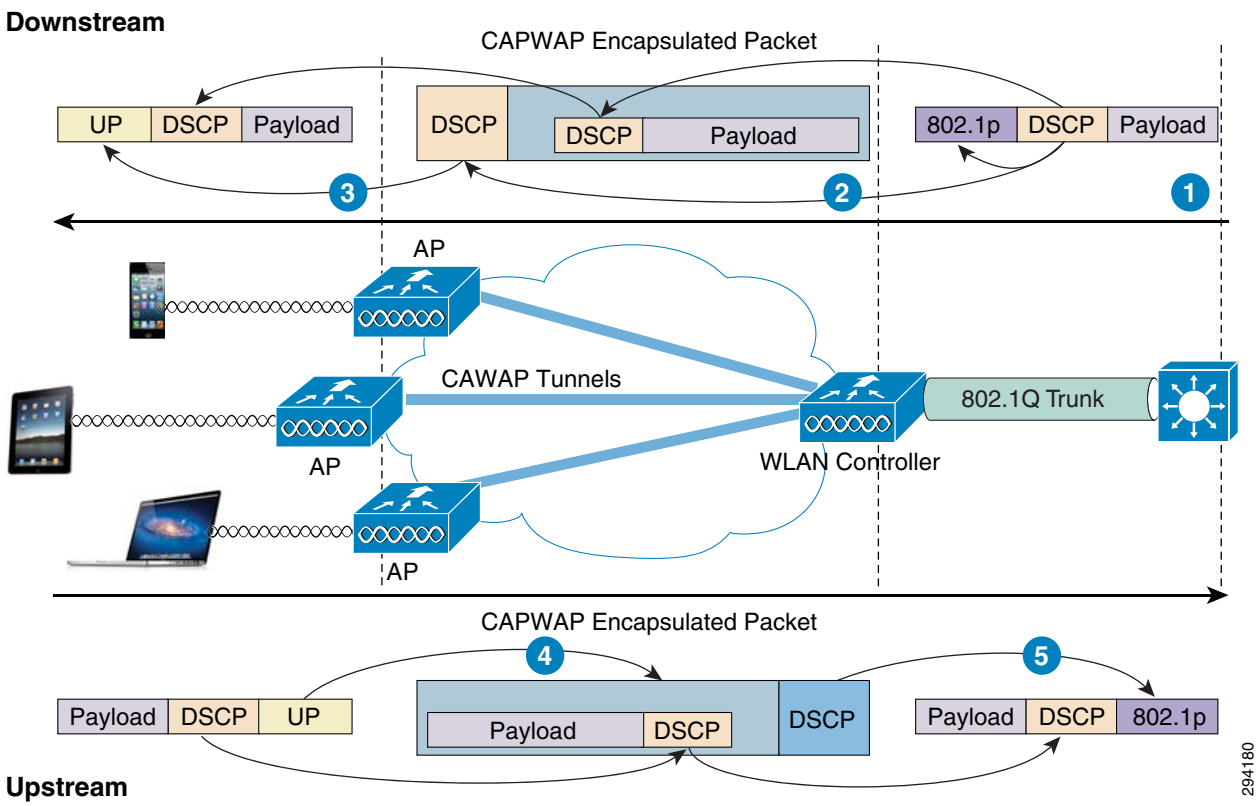
WLC QoS Profile	DSCP
Platinum	EF (DSCP 46)
Gold	AF41 (DSCP 34)
Silver	DF (DSCP 0)
Bronze	AF11 (DSCP 10)

Cisco Wired-Wireless Mapping Points

With the translation-mapping table and default mapping methods in place, the questions that remain are: where is mapping done between wired and wireless networks? And how?

Figure 24-13 is a very important diagram that identifies the various places where wired-and-wireless mapping takes place-in both the downstream direction (top of Figure 24-13) and the upstream direction (bottom of Figure 24-13).

Figure 24-13 Downstream and Upstream Layer 2/Layer 3 Mapping



In Figure 24-13, the following mappings occur in the downstream direction:

Step 1 The network switch maps the DSCP value of an incoming packet (destined to a WLAN via the WLC) to an 802.1p CoS value as it transits the 802.1Q trunks connecting to the WLC (by default this mapping is done by taking the three MSBs of the DSCP and copying these to the 802.1p CoS value).

Step 2 An 802.1Q frame/packet with an 802.1p marking and a DSCP marking arrive at the WLC. The DSCP of the packet is copied to the inner and outer DSCP fields of the CAPWAP packet on egress as it transits toward the destination APs and WLANs.

Note An exception will occur if the DSCP exceeds the Maximum Priority (i.e., the maximum DSCP marking value) defined in the QoS Profile associated with the destination WLAN, in which case both the inner and outer DSCP values will be marked down to this Maximum Priority value.

Step 3 The outer DSCP of the CAPWAP packet arriving at the AP will be mapped to an 802.11e UP marking based on the QoS Translation Table (Table 24-4) or a default mapping (if no explicit mapping is found for the specific DSCP value in the QoS Translation Table). The inner DSCP value is copied to the DSCP-field of the 802.11e frame/packet.

Conversely, in Figure 24-13, the following mappings occur in the upstream direction:

Step 1 The 802.11e UP values and DSCP values of a mobile application are marked in the software of the device as it is transmitted. When the frame/packet arrives at the AP, the UP value will be mapped to the outer DSCP value of the CAPWAP packet; this mapping is based on the QoS Translation Table (Table 24-4) or a default mapping (if no explicit mapping is found for the specific DSCP value in the QoS Translation Table). Additionally, the DSCP value set on the mobile device will be copied to the inner DSCP value of the CAPWAP packet.

Note As previously, the DSCP value assigned to the CAPWAP packet is subject to any Maximum Priority value capping that may be configured within the QoS Profile associated with the WLAN.

Step 2 The outer DSCP of the CAPWAP packet will be mapped to an 802.1p CoS value as the frame/packet leaves the WLC towards the wired network (by default this mapping is done by taking the three MSBs of the DSCP and copying these to the 802.1p CoS value). Additionally, the inner DSCP of the CAPWAP packet will be copied to the DSCP value of the 802.1Q trunked IP packet.

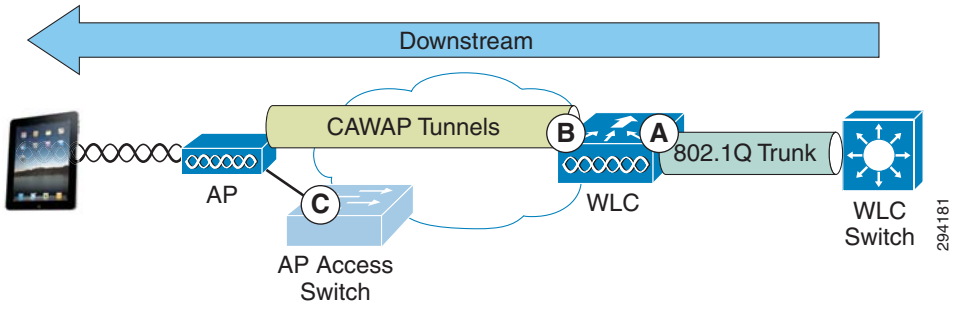
Configuring Downstream QoS Policies for Mobile Applications

Provisioning end-to-end mobile application QoS requires policy configuration at the following points in the wired and wireless networks for downstream flows:

- WLC AVC Profiles—Are configured and assigned to WLANs to identify applications and mark (with DSCP) or drop these packets. Additionally AVC Profiles assign applications to WMM access categories in the radio-downstream direction. AVC Profiles are applied on WLC ingress and are shown as point A in Figure 24-14.
- WLC QoS Profiles—Are assigned to each WLAN and define the (unicast and multicast) Default Priority (i.e., default/Best Effort DSCP value and access category) and the Maximum Priority value for the WLAN. QoS Profiles are applied on WLC egress and are shown as point B in Figure 24-14.

- AP Access Switch QoS Policies—The entire underlying wired-network infrastructure should be configured with QoS policies in line with the strategic application-class model in use for the given enterprise. Additionally, the access-switch to which a given wireless access-point connects to may be used to work-around non-configurable upstream/downstream mapping operations (as is discussed in detail later); these policies are typically applied on access switch egress and are shown as point C in Figure 24-14.

Figure 24-14 Downstream QoS Policy Configuration Points in Wired/Wireless Networks



Wireless Controller QoS Profile GUI Configuration

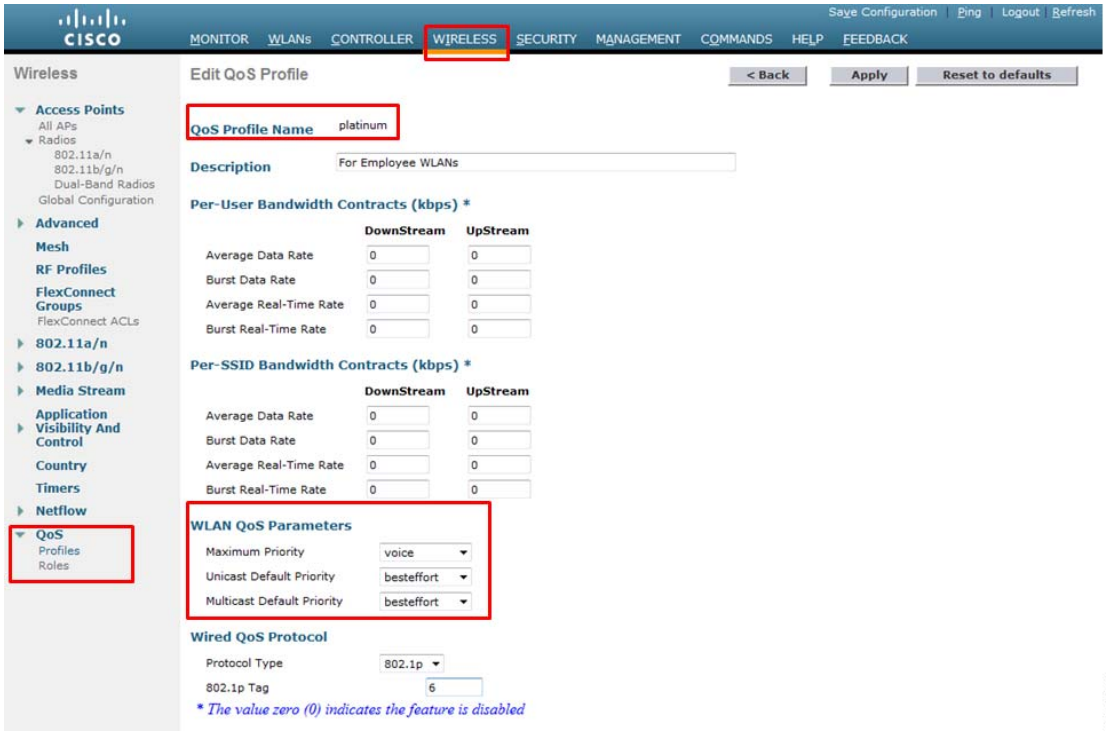
QoS Profiles—like AVC Profiles—are applied to both upstream and downstream flows on WLC egress. It is recommended to complete these steps to define an appropriate QoS Profile for a given WLAN. Among many other parameters, the WLAN QoS Profile defines:

- Per-User Bandwidth Contracts—(Optional) per-user limits for average and peak data and realtime traffic rates.
- Per-SSID Bandwidth Contracts—(Optional) per-SSID limits for average and peak data and realtime traffic rates.
- WLAN Maximum Priority—The highest DSCP marking value that may be used on the WLAN; this value can override AVC policies as well DSCP-values received from the wired network. As such, in multiservice WLANs, it is generally recommended to ensure that the Maximum Priority value be set to voice (i.e., platinum).
- Unicast and Multicast Default Priority—The default DSCP marking value to be used on the WLAN for all traffic not explicitly classified by an overriding AVC Profile. Typically these values are set as best effort (i.e., silver), however there may be cases where this default value may be set to background (i.e., bronze), which is discussed later in the Four-Class Model Mapping Configuration.
- Wired QoS Protocol—Can be set to 802.1p and the maximum CoS value can be defined per WLAN.

Details of a given QoS Profile can be viewed and modified by performing the following steps:

1. Open a web browser to the WLC IP address via HTTPS and login.
2. Click the **WIRELESS** heading bar and expand the **QoS** link on the lower left and click **Profiles**.
3. Click the profile to be viewed/modified (these are listed in alphabetical order: bronze/gold/platinum/silver). Details of the Platinum QoS profile are shown in Figure 24-15.

Figure 24-15 Viewing/Editing the Platinum WLC QoS Profile

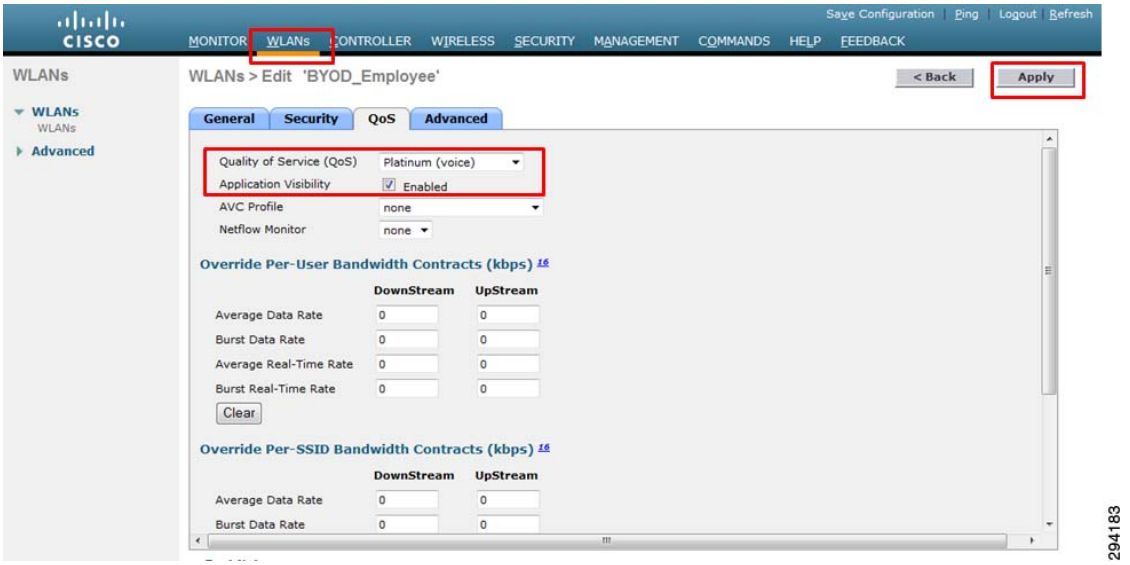


To apply a QoS profile to a WLAN, complete the following steps.

1. Click the **WLANs** heading and select an existing WLAN (or click the **CREATE NEW** button to create a new one).
2. Click the **QoS** tab of the selected WLAN and ensure that:
   a. **Quality of Service (QoS)** is set to **PLATINUM (VOICE)**; this allows for the highest DSCP markings values to be used for AVC policies that are to be attached to the WLAN (in the following section).
   b. **Application Visibility** checkbox is **ENABLED**, as shown in Figure 24-16.
3. Click **Apply**.

Note Applying a QoS policy to a WLAN will momentarily interrupt service.

Figure 24-16 Assigning a QoS Profile to a WLAN and Enabling AVC



WLC QoS Profile CLI Configuration

The following CLI commands will apply the Platinum QoS Profile configuration to WLAN as well as enable AVC Visibility for the WLAN:

Example 24-1 Configuring QoS Profiles for WLANs and Enabling Application Visibility

```
(Cisco WLC) > config wlan disable 1
 ! Disables the WLAN so that the QoS profile may be changed

(Cisco WLC) > config wlan qos 1 platinum
 ! Applies the Platinum QoS profile to the WLAN

(Cisco WLC) > config wlan avc 1 visibility enable
 ! Enables AVC Visibility on WLAn 1

(Cisco WLC) > config wlan enable 1
 ! Re-enables WLAN 1

(Cisco WLC) >
```

This configuration can be verified with:

- **show wlan** (as shown in Example 24-2)

Example 24-2 QoS Profile Verification: show wlan

```
(Cisco WLC) >show wlan 1

WLAN Identifier.................................. 1
Profile Name..................................... BYOD_Employee
Network Name (SSID).............................. BYOD-Employee
Status........................................... Enabled

<snip>
```

```
Quality of Service.............................. Platinum

<snip>

AVC Visibilty.................................... Enabled

(Cisco WLC) >
```

Wireless Controller AVC QoS Profile Configuration

AVC Profiles are applied to both upstream and downstream flows on WLC ingress. While this may simplify the QoS policy configuration on the WLC, it has design implications in upstream/downstream mapping, as has been previously discussed (and which is expanded on in the following sections).

Additionally, each WLAN can have only one AVC profile attached to it to control applications, however an AVC Profile can be attached to multiple WLANs. Also, an AVC Profile can contain a maximum of 32 application rules and a maximum of 16 AVC profiles can be created on a WLC.

Limitations of AVC on the WLC are as noted below and should be kept in mind while deploying AVC on WLC:

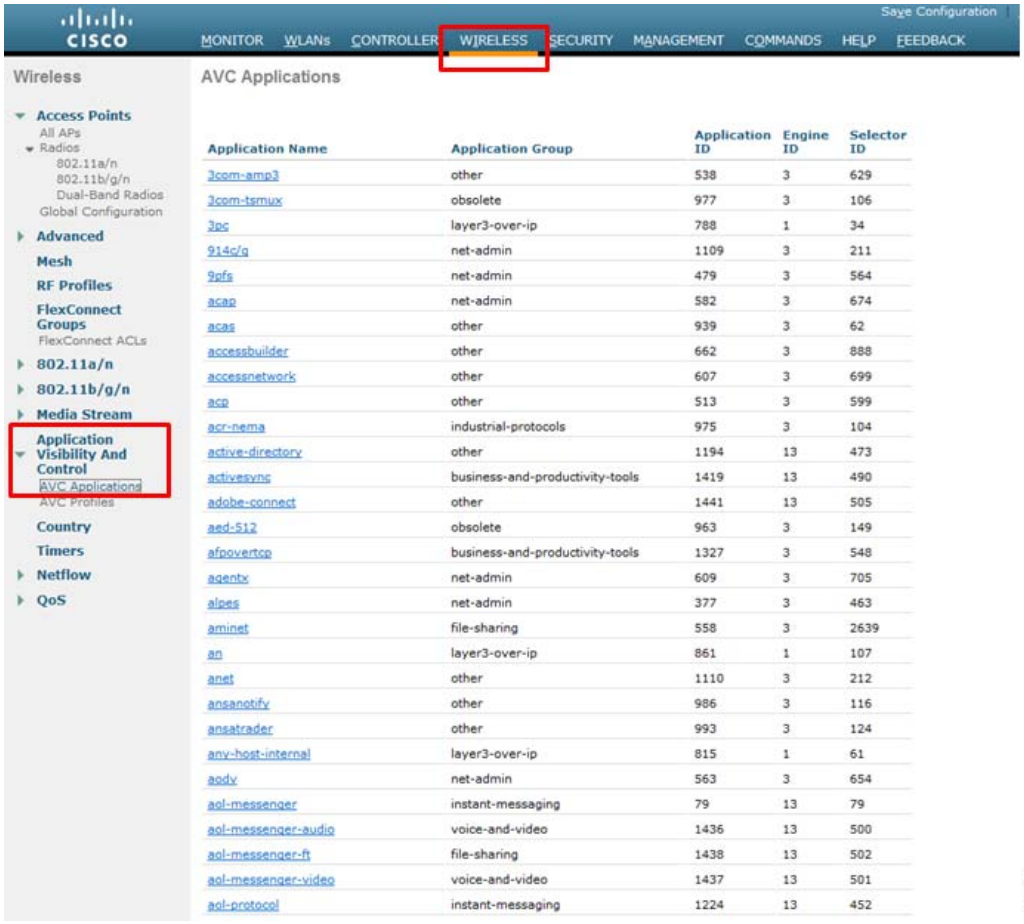
- IPv6 traffic or Multicast traffic cannot be classified.
- AVC is not supported on virtual WLC models.
- AVC is not supported on WLAN configured for Local Switching.
- AVC Profiles cannot be applied on a per controller, VLAN, or client basis (only WLAN).
- The AVC profiles do not support AAA override.
- NBAR based Rate Limiting per Application is not supported, however AVC will work with wireless Bi-Directional Rate Limiting (BDRL) per controller/interface/WLAN/user.
- Any application which is not classified/supported/recognized by the NBAR2 engine is captured as UNCLASSIFIED traffic.

As has been previously discussed, it also is important to note that each WLAN can have both a QoS Profile and an AVC Profile attached to it. The AVC Profile is applied when the packet enters the WLC and the QoS policy is applied when packet exits the WLC. If an AVC profile is mapped to WLAN with an explicit rule for a MARK action, that MARK action will override any default QoS Profile marking configured on the WLAN. However, QoS Profiles may define a Maximum Priority DSCP value for packet marking, which will override any AVC Profile marking policy. Thus care should be taken that QoS and AVC Profiles are correctly configured to complement-and not contradict-one another.

It may be helpful to begin by viewing the list of the over 1000 AVC supported applications by performing the following:

1. Select the **Wireless** heading and then expand the **Application Visibility and Control** link on the lower left.
2. Click the **AVC Applications** link and scroll through the alphabetical application list, as shown in Figure 24-17.

Figure 24-17 AVC Application List



The 1039 AVC applications supported in WLC software version 7.4 are grouped into 16 Application Groups:

- Browsing
- Business-and-productivity-tools
- Email
- File-sharing
- Gaming
- Industrial-protocols
- Instant-messaging
- Internet-privacy
- Layer3-over-ip
- Location-based-services
- Net-admin
- Newsgroup
- Obsolete
- Other

- Trojan
- Voice-and-video

Example 24-3 shows how to display a list of these applications via the WLC CLI.

Example 24-3 WLC CLI-Show AVC Applications

```
(Cisco WLC) > show avc applications

 Application-Name            App-ID  Engine-ID  Selector-ID   Application-Group-Name
 ================            ======  =========  ===========   ======================
 3com-amp3                    538        3       629        other
 3com-tsmux                   977        3       106        obsolete
 3pc                          788        1        34        layer3-over-ip
 914c/g                      1109        3       211        net-admin
 9pfs                         479        3       564        net-admin
 acap                         582        3       674        net-admin
 acas                         939        3        62        other
 accessbuilder                662        3       888        other
 accessnetwork                607        3       699        other
 acp                          513        3       599        other
 acr-nema                      975        3       104         industrial-protocols
…
```

AVC Applications By Business Use Case

Sample AVC applications that relate to the business use cases mentioned at the outset of this document—of provisioning preferential services to protect voice, video, and data applications over wireless networks—are listed below. Additionally applications that might be given a deferential level of service are also listed. However, it is important to note that these are only example applications and do not represent an exhaustive list of applications by class. With over a thousand applications to choose from, these lists are simplified for the sake of brevity and serve only to illustrate the AVC policy concepts.

To ensure voice quality for wireless devices, the cisco-phone application would typically be assigned to the Platinum (Voice) access category via AVC. However, additional VoIP applications may include:

- aol-messenger-audio
- audio-over-http
- fring-voip
- gtalk-voip
- yahoo-voip-messenger
- yahoo-voip-over-sip

Similarly, to protect video and multimedia applications, the following applications might be assigned to the Gold (Video) access-category via AVC:

- cisco-ip-camera
- telepresence-media
- webex-meeting
- ms-lync-media
- aol-messenger-video
- fring-video

- gtalk-video
- livemeeting
- msn-messenger-video
- rhapsody
- skype
- video-over-http

Note It may be that some of these video conferencing applications may be considered non-business in nature (such as Skype and gtalk-video), in which case these may be provisioned into the Bronze (Background) access category.

To deploy AVC policies to protect the signaling protocols relating to these voice and video applications, the following applications might be marked to the Call-Signaling marking of CS3 (DSCP 24) via AVC:

- sip
- sip-tls
- skinny
- telepresence-control
- h323
- rtcp

To deploy policies to protect business-critical applications, the following applications might be marked AF21 (DSCP 18) via AVC:

- citrix
- ms-lync
- ms-dynamics-crm-online
- salesforce
- sap
- oraclenames
- perforce
- phonebook
- semantix
- synergy

On the other hand, some business applications would be best serviced in the background by assigning these to the Bronze (Background) access category via AVC:

- ftp / ftp-data / ftps-data
- cifs
- exchange
- notes
- smtp
- imap/secure imap

- pop3 / secure pop3
- gmail
- hotmail
- yahoo-mail

And finally, many non-business applications can be controlled by either being assigned to the Bronze (Background) access category or dropped via AVC policies:

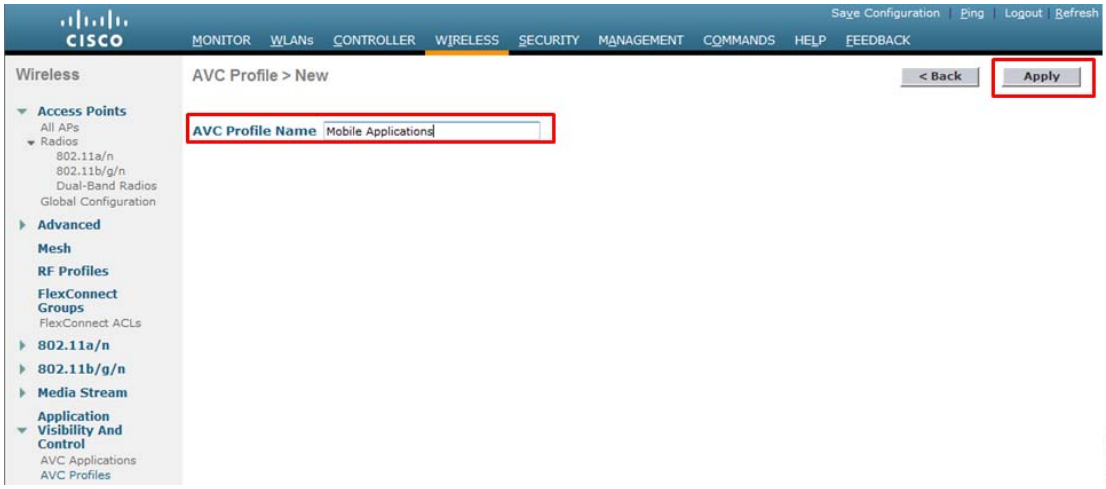
- youtube
- netflix
- facebook
- twitter
- bittorrent
- hulu
- itunes
- picasa
- call-of-duty
- doom
- directplay8

WLC AVC Profile GUI Configuration

To configure an AVC Profile (a set of AVC policies to be applied on a per-WLAN basis), perform the following steps:

1. Click the **AVC Profiles** link on the lower right.
2. Click the **NEW** button on the upper right to create a new AVC profile.
3. Name the new AVC profile and click the **APPLY** button, as shown in Figure 24-18.

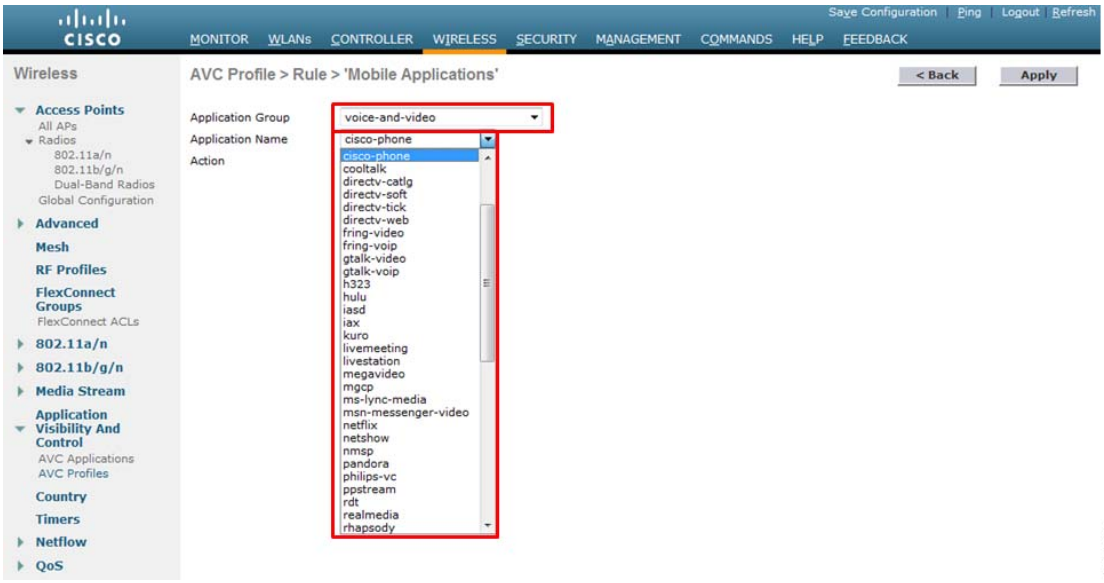
Figure 24-18 Creating an AVC Profile—Part 1



4. Click the name of the new AVC profile (which has become a link).

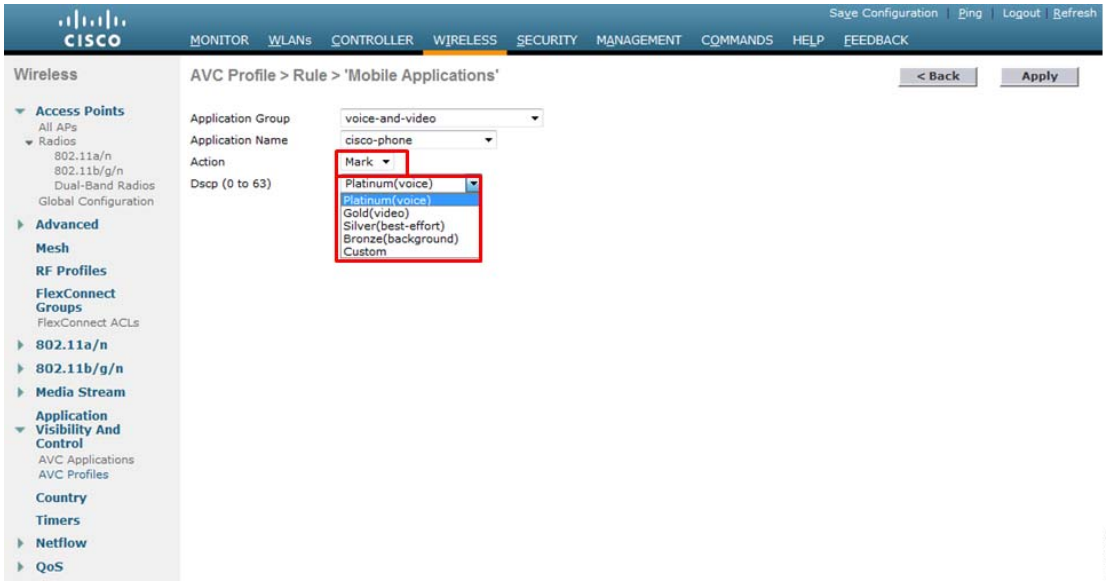
5. Click the **ADD NEW RULE** button on the upper right.
6. Select an **Application Group** to which the application to be controlled belongs.
7. Scroll down the **Application List** to specifically select the application to be controlled, as shown in Figure 24-19.

Figure 24-19 Creating an AVC Profile—Part 2



8. From the **Action** drop-down box select either **DROP** or **MARK**.
9. From the **DSCP** drop-down box select the WMM access category or a **Custom** DSCP marking for the application (as shown in Figure 24-20):
   - Platinum (Voice)—Marks packets to DSCP EF/46.
   - Gold (Video)—Marks packets to DSCP AF41/34.
   - Silver (Best Effort)—Marks packets to the default DSCP DF/0.
   - Bronze (Background)—Marks packets to DSCP AF11/10.
   - Custom [DSCP]—Allows for custom DSCP packet marking between 0 and 63.

Figure 24-20 Creating an AVC Profile—Part 3



10. Click the **Apply** button.

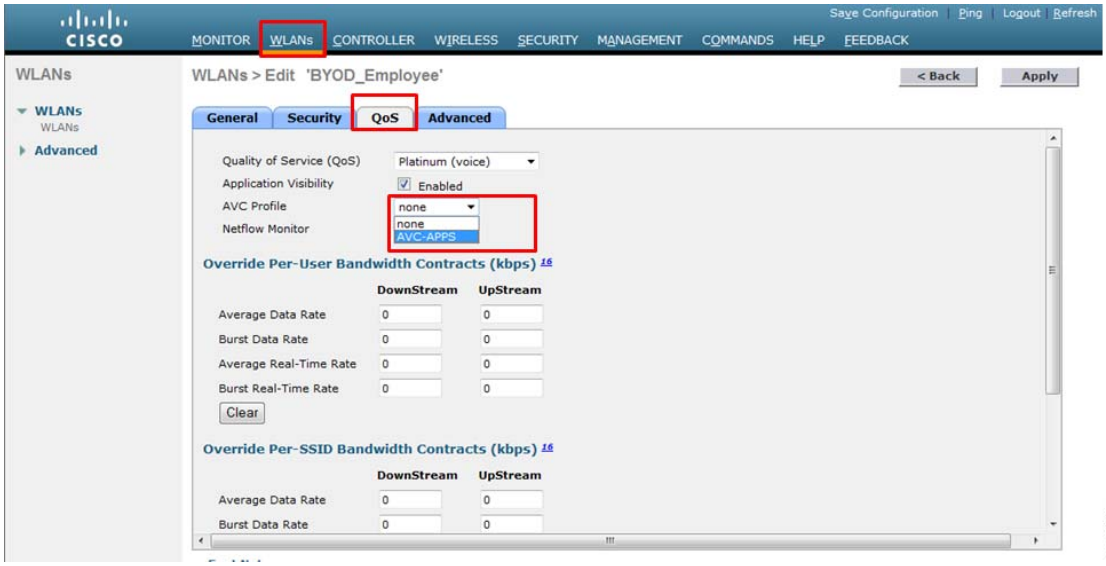
Note Applying a QoS policy to a WLAN will momentarily interrupt service.

Steps 4 through 10 can be repeated to add other applications to the AVC profile; up to 32 rules can be configured within an AVC profile.

Care should be taken to align DSCP marking policies with the enterprise strategic QoS model in use (4/8/12 class models).

Figure 24-21 shows an extended AVC profile that matches a 12-class model's marking scheme and includes temporary markings for the Signaling class and the Transactional Data class.

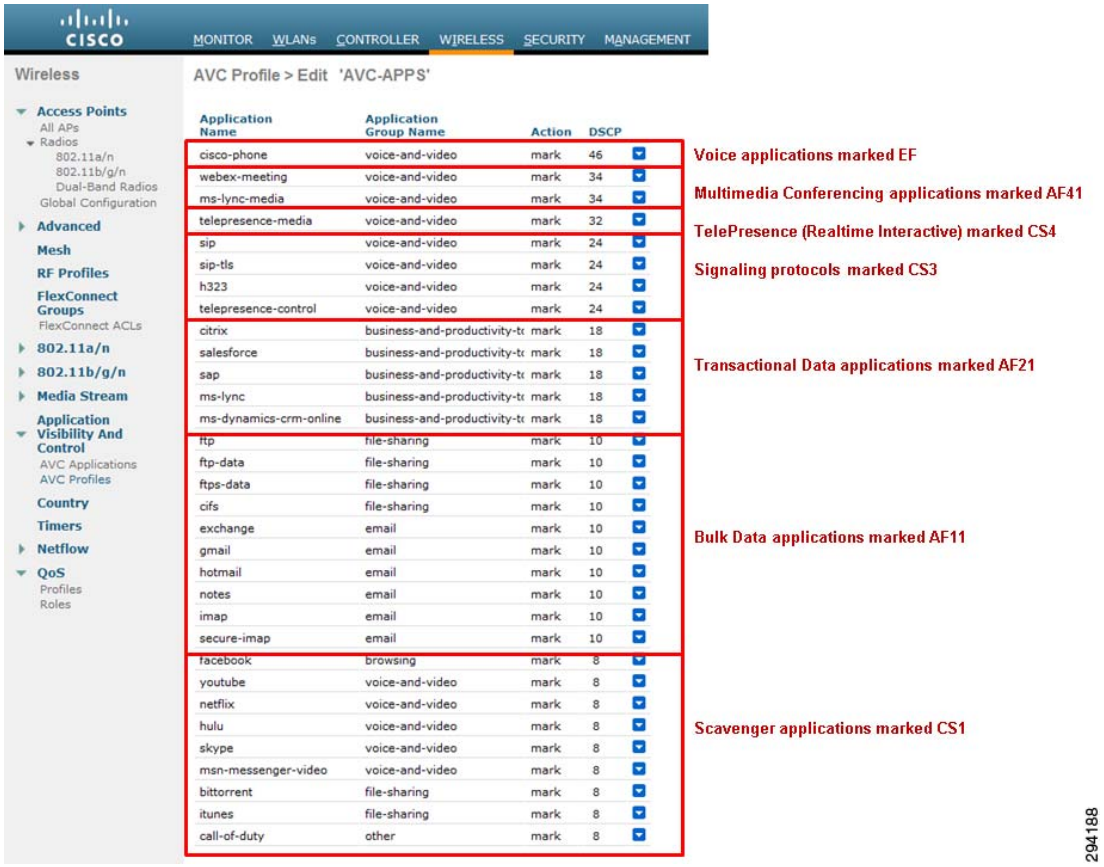
Figure 24-21 Creating an AVC Profile—Part 4



Once the AVC profile has been completed with all the applications to be classified (or the maximum of 32 applications has been reached), then complete the following steps to apply the profile to the WLAN.

1. Click the **WLANS** heading again.
2. Select the WLAN the AVC profile is to be applied to by clicking its WLAN ID link.
3. Click the **QoS** tab.
4. Select the AVC profile to be applied to the WLAN from the AVC Profile drop-down list, as shown in Figure 24-22.

Figure 24-22 Attaching an AVC Profile to a WLAN



A WLAN can only have a single AVC profile applied to it, however AVC profiles may be applied to multiple WLANs.

WLC AVC Profile CLI Configuration

AVC profiles can also be created via the WLC CLI, as shown in Example 24-4.

Example 24-4 WLC CLI—AVC Profile Creation and Definition Example

```
 ! This section creates the "AVC-APPS" AVC profile
(Cisco WLC) > config avc profile AVC-APPS create


 ! This section configures AVC for Voice
 ! Marking Cisco Phone Voice as DSCP EF and assigning it to the Platinum WMM AC
(Cisco WLC) > config avc profile AVC-AVC-APPS rule add application cisco-phone mark 46


 ! This section configures AVC for Multimedia Conferencing applications
 ! Marking these as DSCP AF41 and assigning them to the Gold WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application webex-meeting mark 34
(Cisco WLC) > config avc profile AVC-APPS rule add application ms-lync-media mark 34


 ! This section configures AVC for Realtime Interactive applications
 ! Marking these as DSCP CS4 and assigning them to the Gold WMM AC
```

```
(Cisco WLC) > config avc profile AVC-AVC-APPS rule add application telepresence-media mark
32

 ! This section configures AVC for Signaling protocols
 ! Marking these to DSCP 24 and assigning them to the Gold WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application sip mark 24
(Cisco WLC) > config avc profile AVC-APPS rule add application sip-tls mark 24
(Cisco WLC) > config avc profile AVC-APPS rule add application h323 mark 24
(Cisco WLC) > config avc profile AVC-APPS rule add application telepresence-control mark
24

 ! This section configures AVC for Transactional Data applications
 ! Marking these to DSCP 18 and assigning them to the Gold WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application citrix mark 18
(Cisco WLC) > config avc profile AVC-APPS rule add application salesforce mark 18
(Cisco WLC) > config avc profile AVC-APPS rule add application sap mark 18
(Cisco WLC) > config avc profile AVC-APPS rule add application ms-lync mark 18
(Cisco WLC) > config avc profile AVC-APPS rule add application ms-dynamics-crm-online mark
18

 ! This section configures AVC for Bulk Data applications
 ! Marking these to DSCP AF11 and assigning them to the Bronze WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application ftp mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application ftp-data mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application ftps-data mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application cifs mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application exchange mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application gmail mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application hotmail mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application notes mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application imap mark 10
(Cisco WLC) > config avc profile AVC-APPS rule add application secure-imap mark 10

 ! This section configures AVC for Scavenger applications
 ! Marking these to DSCP CS8 and assigning them to the Bronze WMM AC
(Cisco WLC) > config avc profile AVC-APPS rule add application facebook mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application youtube mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application netflix mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application hulu mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application skype mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application msn-messenger-video mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application bittorrent mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application itunes mark 8
(Cisco WLC) > config avc profile AVC-APPS rule add application call-of-duty mark 8
```

This configuration can be verified with:

- **show avc profile summary** (as shown in Example 24-5)
- **show avc profile detailed** (as shown in Example 24-6)

Example 24-5 AVC Profile Verification: show avc profile summary

```
(Cisco WLC) > show avc profile summary
  Profile-Name                  Number of Rules
  ============                  ===============
  AVC-APPS                             32
(Cisco WLC) >
```

Example 24-6 AVC Profile Verification: show avc profile detailed

```
(Cisco WLC) > show avc profile detailed AVC-APPS

  Application-Name          Application-Group-Name            Action  DSCP
  ================          ======================            ======  ====
  cisco-phone               voice-and-video                    Mark     46
  webex-meeting             voice-and-video                    Mark     34
  ms-lync-media             voice-and-video                    Mark     34
  telepresence-media        voice-and-video                    Mark     32
  sip                       voice-and-video                    Mark     24
  sip-tls                   voice-and-video                    Mark     24
  h323                      voice-and-video                    Mark     24
  telepresence-control      voice-and-video                    Mark     24
  citrix                    business-and-productivity-tools    Mark     18
  salesforce                business-and-productivity-tools    Mark     18
  sap                       business-and-productivity-tools    Mark     18
  ms-lync                   business-and-productivity-tools    Mark     18
  ms-dynamics-crm-online    business-and-productivity-tools    Mark     18
  ftp                       file-sharing                       Mark     10
  ftp-data                  file-sharing                       Mark     10
  ftps-data                 file-sharing                       Mark     10
  cifs                      file-sharing                       Mark     10
  exchange                  email                              Mark     10
  gmail                     email                              Mark     10
  hotmail                   email                              Mark     10
  notes                     email                              Mark     10
  imap                      email                              Mark     10
  secure-imap               email                              Mark     10
  facebook                  browsing                           Mark      8
  youtube                   voice-and-video                    Mark      8
  netflix                   voice-and-video                    Mark      8
  hulu                      voice-and-video                    Mark      8
  skype                     voice-and-video                    Mark      8
  msn-messenger-video       voice-and-video                    Mark      8
  bittorrent                file-sharing                       Mark      8
  itunes                    file-sharing                       Mark      8
  call-of-duty              other                              Mark      8

  Associated WLAN IDs       :
  Associated Remote LAN IDs :
  Associated Guest LAN IDs  :

(Cisco WLC) >
```

AVC profiles can also be attached to the WLAN via the WLC CLI, as shown in Example 24-7.

Example 24-7 WLC CLI-Attaching AVC Profiles to WLANs Example

```
(Cisco WLC) > config wlan avc 1 profile AVC-APPS enable
 ! This command applies the AVC profile "AVC-APPS" to WLAN ID 1
```

This configuration can be verified with:

- **show avc profile detailed** (as shown in Example 24-8)
- **show wlan** (as shown in Example 24-9)

Example 24-8 AVC Profile Verification: show avc profile detailed

```
(Cisco WLC) > show avc profile detailed AVC-APPS

  Application-Name          Application-Group-Name            Action  DSCP
```

```
  ================          ======================          ======  ====
  cisco-phone               voice-and-video                  Mark     46

<snip>

  Associated WLAN IDs        : 1
  Associated Remote LAN IDs  :
  Associated Guest LAN IDs   :

(Cisco WLC) >
```

Example 24-9 AVC Profile Verification: show wlan

```
(Cisco WLC) >show wlan 1

WLAN Identifier.................................. 1
Profile Name..................................... BYOD_Employee
Network Name (SSID).............................. BYOD-Employee
Status........................................... Enabled

<snip>

Quality of Service.............................. Platinum

<snip>

AVC Visibilty.................................... Enabled
AVC Profile Name................................. AVC-APPS

(Cisco WLC) >
```

AP Access Switch Downstream QoS Policies

The purpose of network QoS policies is to ensure that enterprise application classes will map into the correct WMM access categories (and vice versa). Since the admission to the WMM access categories is based on 802.11e UP values, which in turn are based on DSCP-to-UP translations (as covered in Step 3 of Figure 24-13), the packets need to enter the AP with the necessary (outer) DSCP values (of the CAPWAP packet) to achieve the desired mappings. This requirement is underscored by the fact that the WLC/AP QoS Translation Table (Table 24-4) is not modifiable and neither is the default DSCP-to-CoS/UP mapping (Table 24-5). Therefore the final point that an administrator could correctly set (or reset) DSCP values in packets before handing these off to the AP to ensure the traffic is placed in the desired WMM access category is at the egress interface of the network access switch to which the AP is connected (shown as point C in Figure 24-14).

Note Granted, it would be possible to perform DSCP remarking within the Cisco WLC via an AVC policy. Such an AVC policy would match on an application type(s) and then set this to a differing DSCP value as used in the enterprise campus network. However AVC policies apply in both directions simultaneously (and there is no option to apply these in the downstream or upstream direction only). Therefore such an AVC policy—intended for downstream remarking—would include the undesired effect of marking these application types incompatibly with the enterprise models for traffic in the upstream direction towards the wired network.

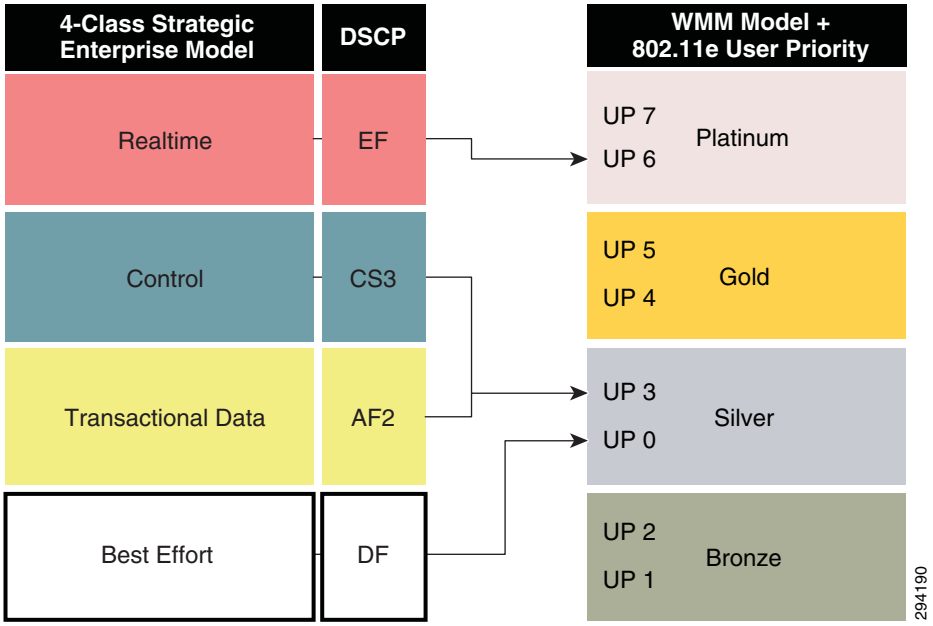
The details of the network downstream DSCP remarking policy will vary according to the enterprise application-class models in use. To illustrate a cross-section of marking-model variations—as well as Catalyst switching platform variations—the following mappings examples will be considered in detail:

- Four-class enterprise model mapped to WMM on a Catalyst 3750-X series switch (Four-Class Enterprise to WMM Mapping Model)
- Eight-class enterprise model mapped to WMM on a Catalyst 4500E series switch with Supervisor 7-E (Eight-Class Enterprise to WMM Mapping Model)
- Twelve-class enterprise model mapped to WMM on a Catalyst 6500 series switch with Supervisor 2T (Twelve-Class Enterprise to WMM Mapping Model)

Four-Class Enterprise to WMM Mapping Model

If the enterprise has deployed a four-class strategic model, then this allows for a 1:1 mapping into the four WMM access categories. The default mapping for a four-class enterprise model to WMM is shown in Figure 24-23.

Figure 24-23 Default Downstream Four-Class Enterprise Model Mapping to WMM



As can be seen in Figure 24-23, if the QoS Translation Table (Table 24-4) and default mappings (Table 24-5) were applied on these four enterprise QoS classes, these would map to only two access categories: Platinum (Voice) and Silver (Best Effort). The remaining two access categories would not even be used. While this default mapping will ensure voice quality, no other application will benefit from wireless QoS.

An alternative approach would be to perform remarking at the network access switch (connecting to the wireless AP) at the egress interface in the outbound direction such that the Control/Signaling application class is remarked to a DSCP value that will map to an UP value (on the AP) that will, in turn, be assigned to the Gold WMM AC. Similarly, Best Effort traffic may be mapped to a DSCP value that will map to an UP value that will, in turn, be assigned to the Bronze WMM AC. This re-mapping approach would

then leave the Silver WMM AC to exclusively service Transactional Data applications and thereby reflect the same overall relative priority of servicing as the original model. The modified four-class mapping model is shown in Figure 24-25.

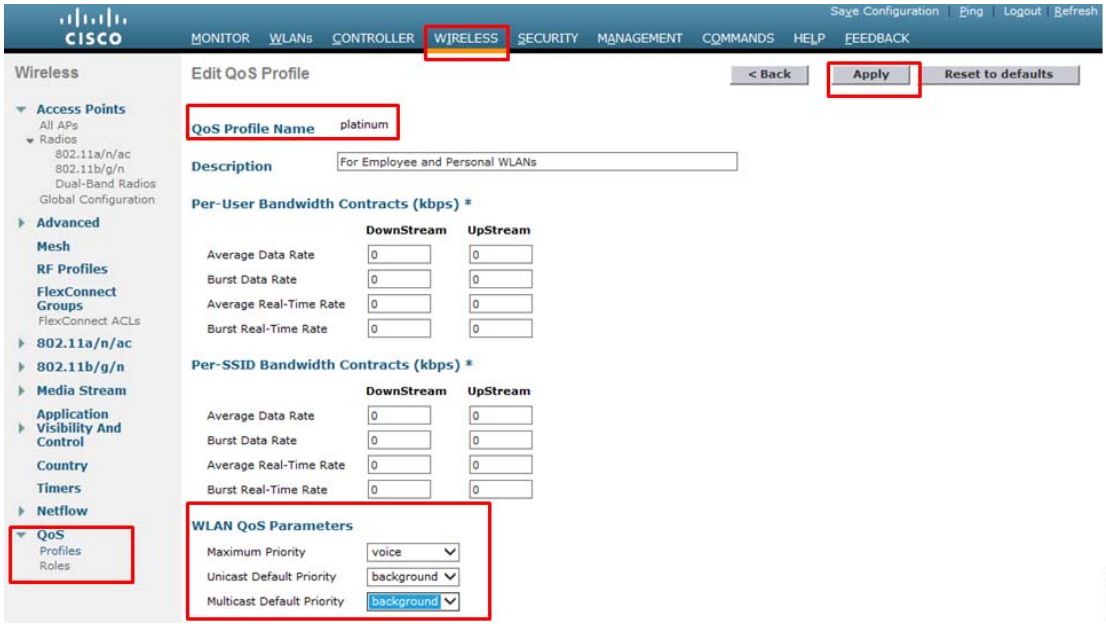
Perhaps the question may arise: why not just configure an AVC policy on the WLC to match signaling traffic and remark it to a DSCP value that will map to the Gold WMM AC—like say CS4/DSCP 32—rather than configuring mapping policies on the access switch? To answer this, it should be kept in mind that such an AVC marking policy on the WLC—which operates both in the upstream and downstream direction—will interfere with QoS policies on the rest of the wired network (which includes both the upstream network as well as the transit network between the WLC and the APs). Specifically, while such a policy may achieve the intended result in the downstream direction, it will include the unintended effect of marking signaling to CS4 (rather than CS3) in the upstream direction also, which would be incompatible with the rest of the enterprise strategic policy and would have to be remapped at the WLC upstream switch. Furthermore, Signaling traffic marked on the WLC to CS4 by AVC and transiting in the downstream direction would have no QoS applied over the wired network between the WLC and the APs unless the administrator modified all the policies in the path to accommodate. Therefore, a much simpler approach to achieve the same end result is to include a simple remarking policy on the final access switch connecting to the AP.

However, rather than remapping signaling traffic to a code-point that may be used for another application, it may be better to remark signaling to a non-standard codepoint for this one-time operation, such as DSCP 33. DSCP 33 would map (by default) to the Gold WMM AC and would uniquely identify this remapped signaling traffic.

Similarly the default Best Effort class could be remapped on the network access switch to a non-standard codepoint that would be assigned to the Bronze WMM AC-such as DSCP 9. However, some platforms, such as the Catalyst 3750, do not support egress marking policies and only support marking/remarking/DSCP-mutation policies on ingress. In such a case, remarking Best Effort traffic on access switch ingress will affect BE marking on all interfaces and not just the interface connecting to the access point.

In such a case, a more elegant method would be to modify the QoS Profile for the WLAN so that the Unicast/Multicast Default Priority for the WLAN is set to Background instead of Best Effort, as shown in Figure 24-24.

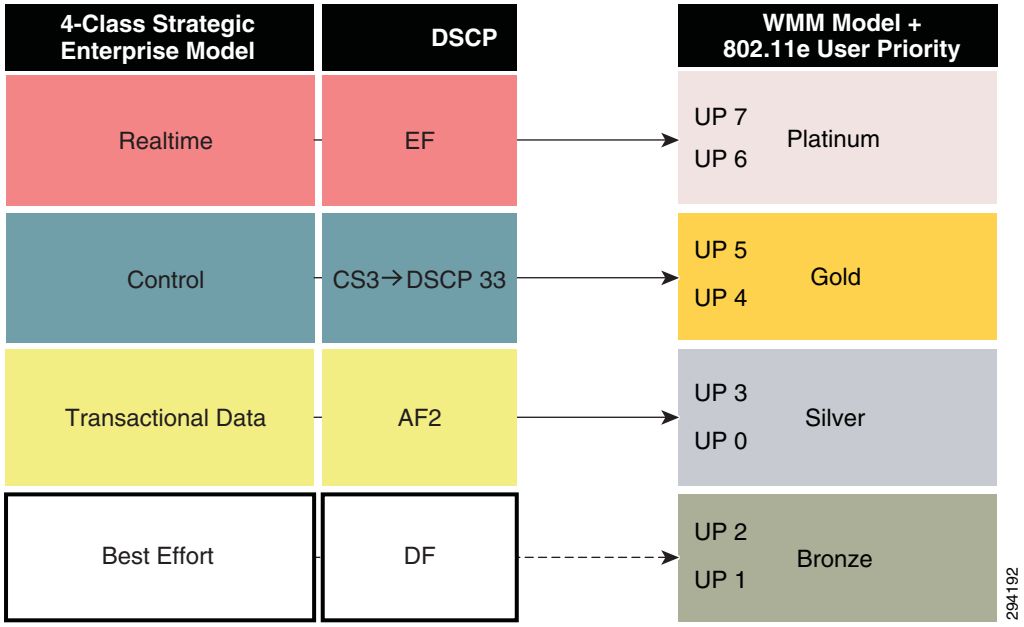
Figure 24-24 Modifying the Platinum WLC QoS Profile-Best Effort 'Background



It bears noting that assigning the default class to the Background WMM category will marginally delay best effort traffic, however this is the nature of QoS as there is always a tradeoff. In this scenario, this is the necessary cost of providing superior levels of service to the signaling and transactional data application classes, thus providing the four levels of service required to meet their strategic business objectives of QoS.

The modified mapping of a four-class enterprise model to WMM is shown in Figure 24-25.

Figure 24-25 Modified Downstream Four-Class Enterprise Model Mapping to WMM

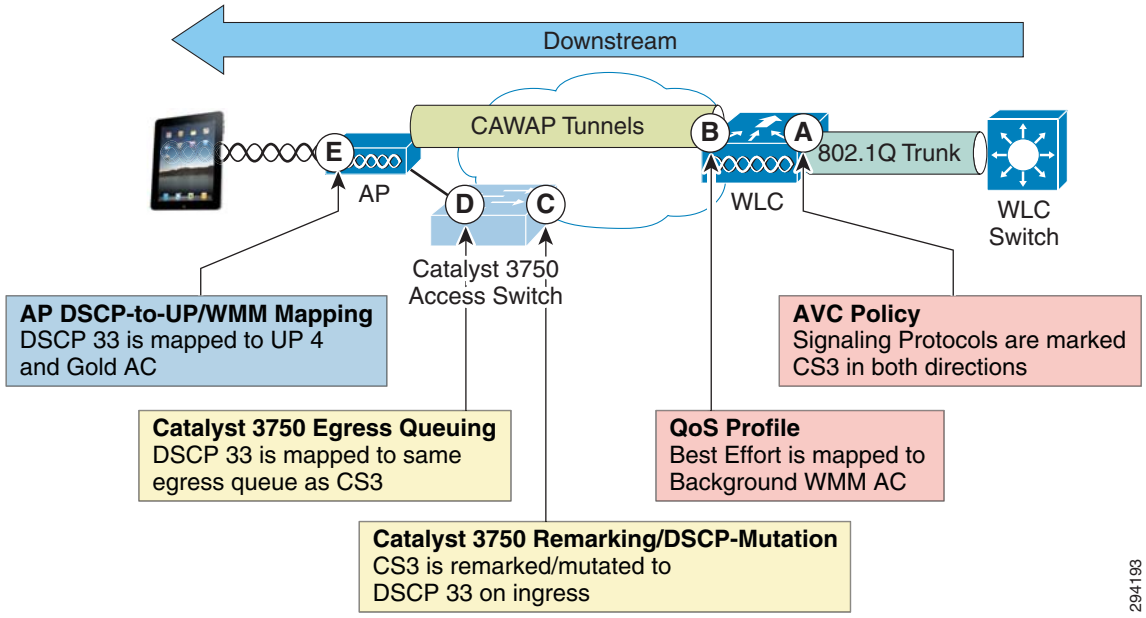


Catalyst 3750 Configuration Example

To highlight platform-specific syntax, features, and limitations, each mapping model is presented on a different Catalyst switching platform. This four-class enterprise mapping to WMM AC is presented on the Catalyst 3750 (which will be identical to configuration for the Catalyst 2960 and 3560 switches also).

As noted, the Catalyst 3750 does not support egress remarking; only ingress marking or ingress DSCP-mutation are supported, as shown in Figure 24-26.

Figure 24-26 Catalyst 3750 Four-Class Enterprise to WMM Marking and Mapping Policies



The Catalyst 3750 employs Multilayer Switch QoS (MLS QoS), and as such DSCP-remarking policies can be configured in one of two ways:

- Class-based marking policy (as shown in Example 24-10)
- DSCP-Mutation (as shown in Example 24-11)

Example 24-10 Catalyst 3750 Downstream Four-Class Enterprise Model Mapping Policy to WMM via Class-Based Marking

```
 ! This section configures the class-map
C3750-X(config-cmap)# class-map match-all SIGNALING
C3750-X(config-cmap)#  match ip dscp cs3
 ! Signaling traffic is matched on DSCP CS3


 ! This section configures the Network Downstream Remarking policy-map
C3750-X(config-cmap)# policy-map DOWNSTREAM-WMM-REMARKING
C3750-X(config-pmap-c)# class SIGNALING
C3750-X(config-pmap-c)#  set dscp 33
 ! Signaling is remarked DSCP 33 to map into WMM Gold downstream


 ! This section attaches the Downstream policy to the campus-side interface
C3750-X(config)# interface TenGigabitEthernet2/1/1
C3750-X(config-if)# mls qos trust dscp
```

```
 ! Configures the port to statically trust DSCP on ingress
C3750-X(config-if)# service-policy input DOWNSTREAM-WMM-REMARKING
 ! Attaches the Downstream DSCP remarking policy to the interface on ingress
```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Example 24-11 Downstream Four-Class Enterprise Model Mapping Policy to WMM via DSCP-Mutation

```
 ! This section configures the Downstream DSCP Mutation map
C3750-X(config)# mls qos map dscp-mutation DOWNSTREAM-WMM-MUTATION 24 to 33


 ! This section attaches the Downstream policy to the campus-side interface
C3750-X(config)# interface TenGigabitEthernet2/1/1
C3750-X(config-if)# mls qos trust dscp
 ! Configures the port to statically trust DSCP on ingress
C3750-X(config-if)# mls qos dscp-mutation DOWNSTREAM-WMM-MUTATION
 ! Attaches the Downstream DSCP mutation map to the interface on ingress
```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show mls qos maps dscp-mutation**

Additionally, this non-standard DSCP representing signaling traffic should be mapped to the same egress queue as regular signaling traffic (marked CS3), as shown in Example 24-12.

Example 24-12 Catalyst 3750 Egress Queuing Configuration

```
! This section configures buffers and thresholds on Q1 through Q4
C3750(config)# mls qos queue-set output 1 buffers 15 30 35 20
 ! Queue buffers are allocated
C3750(config)# mls qos queue-set output 1 threshold 1 100 100 100 100
 ! All Q1 (PQ) Thresholds are set to 100%
C3750(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%
C3750(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%
C3750(config)# mls qos queue-set output 1 threshold 4 60 100 100 400
 ! Q4T1 is set to 60%; Q4T2 is set to 100%
 ! Q4 Reserve Threshold is set to 100%;
 ! Q4 Maximum (Overflow) Threshold is set to 400%
 ! This section configures egress CoS-to-Queue mappings (if required)
C3750(config)# mls qos srr-queue output cos-map queue 1 threshold 3 4 5
 ! CoS 4 and 5 are mapped to egress Q1T3 (the tail of the PQ)
C3750(config)# mls qos srr-queue output cos-map queue 2 threshold 1 2
 ! CoS 2 is mapped to egress Q2T1
C3750(config)#mls qos srr-queue output cos-map queue 2 threshold 2 3
 ! CoS 3 is mapped to egress Q2T2
C3750(config)# mls qos srr-queue output cos-map queue 2 threshold 3 6 7
 ! CoS 6 and 7 are mapped to Q2T3
```

```
C3750(config)# mls qos srr-queue output cos-map queue 3 threshold 3 0
 ! CoS 0 is mapped to Q3T3 (the tail of the default queue)
C3750(config)# mls qos srr-queue output cos-map queue 4 threshold 3 1
 ! CoS 1 is mapped to Q4T3 (tail of the less-than-best-effort queue)


 ! This section configures egress DSCP-to-Queue mappings
C3750(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)
C3750(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
C3750(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
C3750(config)# mls qos srr-queue output dscp-map queue 2 threshold 2 24 33
 ! DSCP CS3 and DSCP 33 is mapped to egress Q2T2
C3750(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
 ! DSCP CS6 and CS7 are mapped to egress Q2T3
C3750(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)
C3750(config)# mls qos srr-queue output dscp-map queue 4 threshold 1 8
 ! DSCP CS1 is mapped to egress Q4T1
C3750(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
 ! DSCP AF1 is mapped to Q4T3 (tail of the less-than-best-effort queue)


 ! This section configures interface egress queuing parameters
C3750(config)# interface range GigabitEthernet1/0/1-48
C3750(config-if-range)# queue-set 1
 ! The interface(s) is assigned to queue-set 1
C3750(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
C3750(config-if-range)# priority-queue out
 ! Q1 is enabled as a strict priority queue
```

This configuration can be verified with the commands:

- **show mls qos queue-set**
- **show mls qos maps cos-output-q**
- **show mls qos maps dscp-output-q**
- **show mls qos interface interface x/y queueing**
- **show mls qos interface interface x/y statistics**



Note Additional design recommendations for the Catalyst 3750 can be found in the Medianet Campus QoS Design 4.0 design chapter at: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html#wp1099462.

Eight-Class Enterprise to WMM Mapping Model

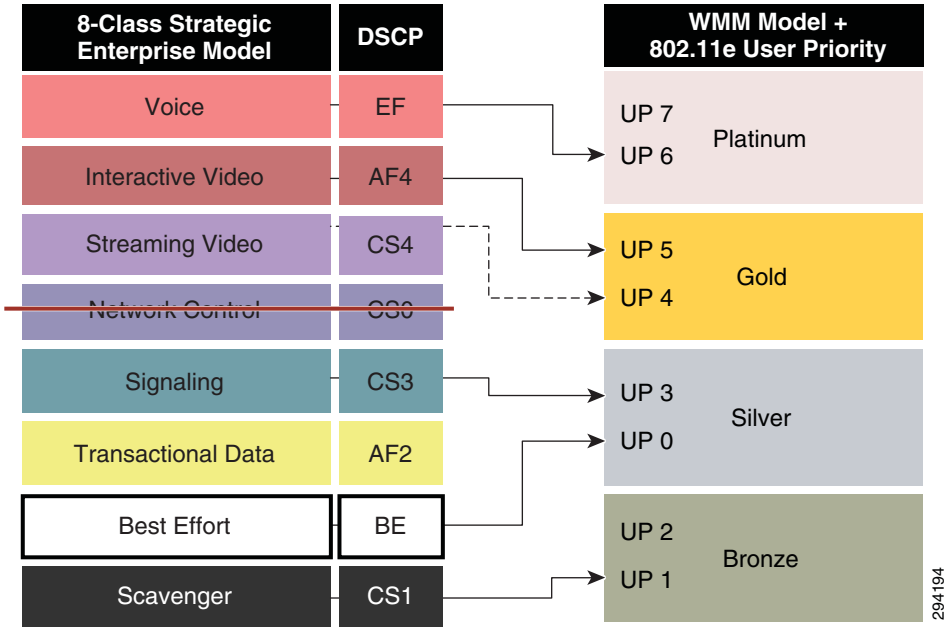
If the enterprise has deployed an eight-class strategic model, then a simple 1:1 mapping is no longer possible and some additional considerations need to be taken into account.

One such consideration is whether all eight traffic classes will have application traffic generated to and/or from wireless mobile devices. For example, no mobile clients or devices should be transmitting or receiving Network Control traffic (as this application class is intended for servicing the control plane requirements of the network infrastructure).

Note It also can be argued that no Streaming Video traffic should be sourced from mobile devices (although these devices might be receivers of such streams). Nonetheless, since traffic for this class may be present—primarily in the downstream direction—and as such will be included in the mapping model (albeit with a dashed line to indicate that this application traffic is typically unidirectional in the downstream direction only).

The default mapping for a four-class enterprise model to WMM is shown in Figure 24-27.

Figure 24-27 Default Downstream Eight-Class Enterprise Model Mapping to WMM



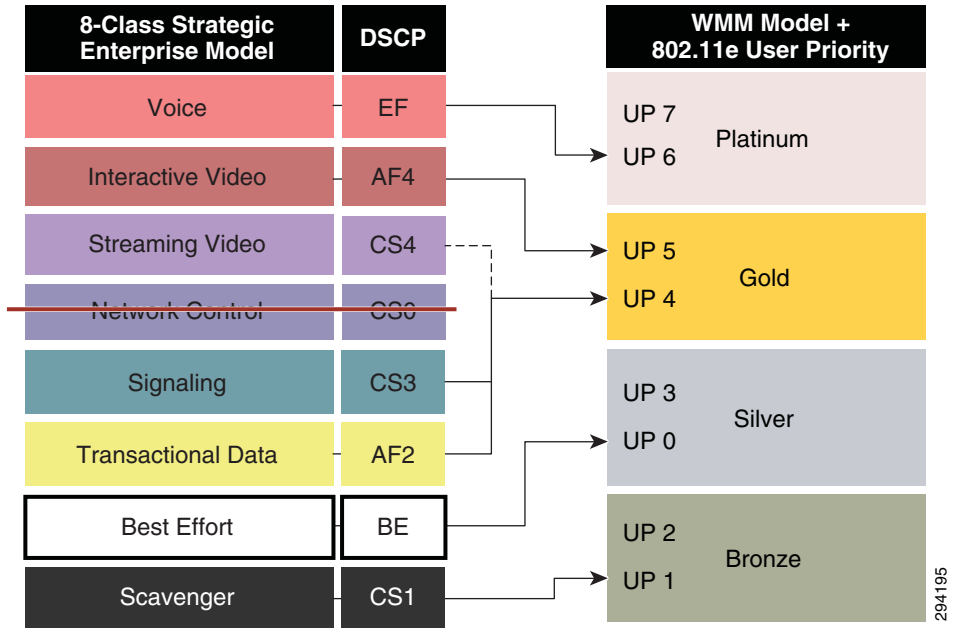
While this default Eight-Class mapping may be adequate for some environments, it should be noted that there is no QoS provisioned for Transactional Data traffic (that may include VDI applications like Citrix XenDesktop or VMware View) nor for Signaling traffic (which is control plane traffic for IP telephone and/or IP video telephony applications)—other than merely protecting these application classes from Scavenger traffic.

Therefore, it may be desirable to remap these applications to the Gold access-category, using non-standard codepoints (as in the previous example). In this case, Signaling traffic can again be mapped to DSCP 33 and Transactional Data can be mapped to DSCP 35 at the AP access switch in the egress direction,

Note As both DSCP 33 and 35 map to the same UP value of 4, there is no advantage/disadvantage to marking these applications to a higher/lower DSCP value, provided these maps to the desired WMM AC. The key is keeping these values unique and distinct for traffic management purposes.

The modified eight-class mapping model is shown in Figure 24-28.

Figure 24-28 Modified Downstream Eight-Class Enterprise Model Mapping to WMM



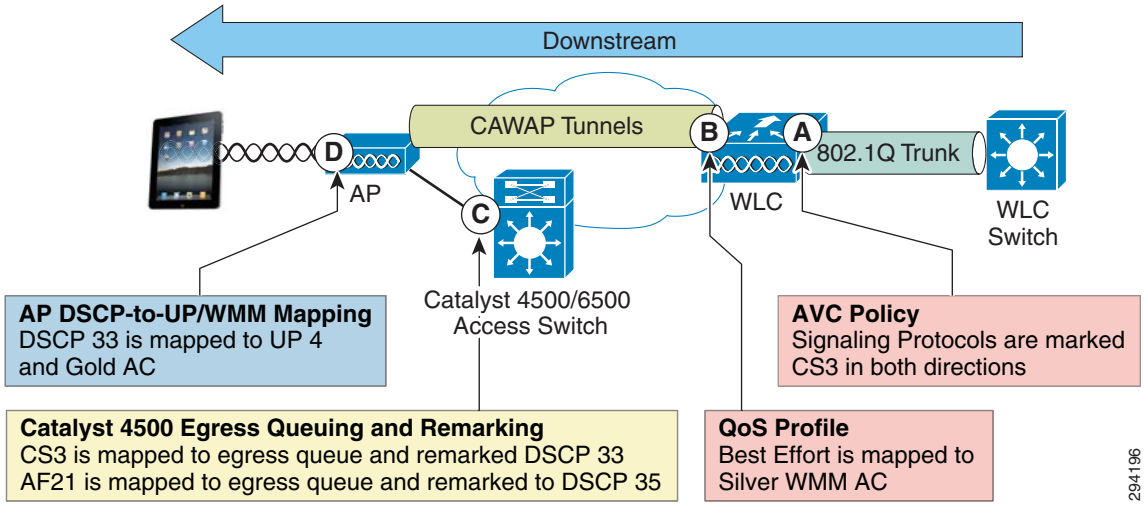
Note Best Effort traffic is assigned to the default Silver (Best Effort) WMM Access Category in the WLC WLAN QoS Profile for this mapping model.

Catalyst 4500E Supervisor 7-E Example

This DSCP remarking tactical design adapted for the Catalyst 4500 is illustrated in Figure 24-29. Since the Catalyst 4500 uses MQC QoS, it supports egress remarking. As such, only a single policy is required on this switch: namely adapting the final egress queuing policy to remark the Signaling and Transactional Data application classes to DSCP 33 and 35, respectively.

Note Incidentally, the remarking policies and enforcement points for an Eight-Class Enterprise to WMM Mapping Model on a Catalyst 4500 are the same as a Twelve-Class Enterprise to WMM Mapping Model on a Catalyst 6500. As such, both are shown in a single diagram (Figure 24-29).

Figure 24-29 Catalyst 4500/6500 Eight-Class/Twelve-Class Enterprise to WMM Marking and Mapping Policies



This modified eight-class mapping configuration example is presented on the Catalyst 4500-E series switch (with Supervisor 7-E), as shown in Example 24-13. Since MQC only supports one service-policy statement attached to a given interface in a single direction, the egress remarking policy must be combined with the (eight-class) egress queuing policy as a single MQC service-policy in the output direction, as shown in Example 24-13.

Example 24-13 Catalyst 4500 Downstream Eight-Class Enterprise Model Queuing and Mapping Policy to WMM

```
 ! This section configures the class-maps for the egress queuing policy
C4500(config)# class-map match-any PRIORITY-QUEUE
C4500(config-cmap)# match dscp ef
 ! VoIP (EF) is mapped to the PQ
C4500(config)# class-map match-all INTERACTIVE-VIDEO-QUEUE
C4500(config-cmap)# match dscp af41 af42 af43
 ! Interactive-Video (AF4) is assigned a dedicated queue
C4500(config)# class-map match-all STREAMING-VIDEO-QUEUE
C4500(config-cmap)# match dscp af31 af32 af33
 ! Streaming-Video (AF3) is assigned a dedicated queue
C4500(config)# class-map match-any CONTROL-QUEUE
C4500(config-cmap)# match dscp cs6
 ! Network Control (CS6) is mapped to a dedicated queue
C4500(config)# class-map match-any SIGNALING-QUEUE
C4500(config-cmap)# match dscp cs3
 ! Signaling (CS3) is mapped to a dedicated queue
C4500(config)# class-map match-all TRANSACTIONAL-DATA-QUEUE
C4500(config-cmap)# match dscp af21 af22 af23
 ! Transactional Data (AF2) is assigned a dedicated queue
C4500(config)# class-map match-all SCAVENGER-QUEUE
C4500(config-cmap)# match dscp cs1
 ! Scavenger (CS1) is assigned a dedicated queue


 ! This section configures the 1P7Q1T+DBL egress queuing policy-map
C4500(config)# policy-map 1P7Q1T+DBL+DOWNSTREAM-MAPPING
C4500(config-pmap-c)# class PRIORITY-QUEUE
C4500(config-pmap-c)#  priority
```

```
 ! Defines a priority queue
C4500(config-pmap-c)# class INTERACTIVE-VIDEO-QUEUE
C4500(config-pmap-c)#  bandwidth remaining percent 23
C4500(config-pmap-c)#  dbl
 ! Defines a interactive-video queue with 23% BW remaining + DBL
C4500(config-pmap-c)# class STREAMING-VIDEO-QUEUE
C4500(config-pmap-c)# bandwidth remaining percent 10
C4500(config-pmap-c)#  dbl
 ! Defines a streaming-video queue with 10% BW remaining + DBL
C4500(config-pmap-c)# class CONTROL-QUEUE
C4500(config-pmap-c)#  bandwidth remaining percent 5
 ! Defines a control/management queue with 5% BW remaining
C4500(config-pmap-c)# class SIGNALING-QUEUE
C4500(config-pmap-c)#  bandwidth remaining percent 2
 ! Defines a signaling queue with 2% BW remaining
C4500(config-pmap-c)#  set dscp 33
 ! Remarks signaling traffic to DSCP 33 for WMM Gold downstream mapping
C4500(config-pmap-c)# class TRANSACTIONAL-DATA-QUEUE
C4500(config-pmap-c)#  bandwidth remaining percent 24
C4500(config-pmap-c)#  dbl
 ! Defines a transactional data queue with 24% BW remaining + DBL
C4500(config-pmap-c)#  set dscp 35
 ! Remarks transactional data to DSCP 35 for WMM Gold downstream mapping
C4500(config-pmap-c)# class SCAVENGER-QUEUE
C4500(config-pmap-c)#  bandwidth remaining percent 1
 ! Defines a (minimal) scavenger queue with 1% BW remaining/limit
C4500(config-pmap-c)# class class-default
C4500(config-pmap-c)#  bandwidth remaining percent 25
C4500(config-pmap-c)#  dbl
 ! Provisions the default/Best Effort queue with 25% BW remaining + DBL


 ! This section attaches the egress queuing & mapping policy to AP interface
C4500(config)# interface GigabitEthernet 3/1
C4500(config-if)# service-policy output 1P7Q1T+DBL+DOWNSTREAM-MAPPING
 ! Attaches the combined egress queuing and egress remarking policy to the int
```

Note Additional detail on Catalyst 4500 queuing policy recommendations can be found in Medianet Campus QoS Design 4.0 design chapter at: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html#wp1100873.

Note Class-maps defined for egress-queuing policies require unique names from any ingress-policy class-maps; otherwise classification errors may occur due to overlapping classification-logic. Therefore the class-maps names in this example have "-QUEUE" appended to them.

Note It is recommended to use **match dscp** and **set dscp**—as opposed to **match ip dscp** and **set ip dscp**—as the former will match on both IPv4 and IPv6 packets, whereas the latter will match only on IPv4 packets. Although AVC does not yet classify IPv6 traffic, these packets can still be properly mapped at this node to the correct downstream WMM access category.

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**

- **show policy-map interface**

Twelve-Class Enterprise to WMM Mapping Model

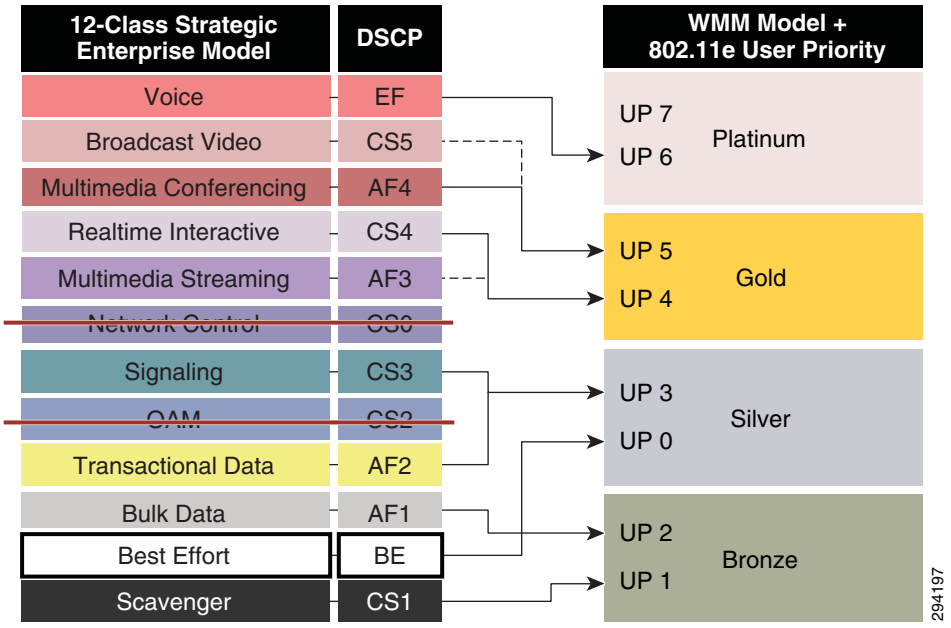
If the enterprise has deployed a twelve-class strategic model, then further class-pruning and application-class collapsing needs to take place.

To this end, both the Network Control and the Operations/Administration/Management application traffic classes can be pruned out of the mapping model, as wireless endpoint devices should not be generating nor receiving traffic from these classes (as these classes are intended as control plane classes for the network infrastructure).

Note It can be argued that neither Broadcast Video nor Streaming Video traffic should be sourced from mobile devices (although these devices might be receivers of such streams). Nonetheless, since traffic for this class may be present, it is included in the mapping model (albeit with a dashed line to indicate that this application traffic is typically unidirectional in the downstream direction).

The default mapping for a twelve-class enterprise model to WMM is shown in Figure 24-30.

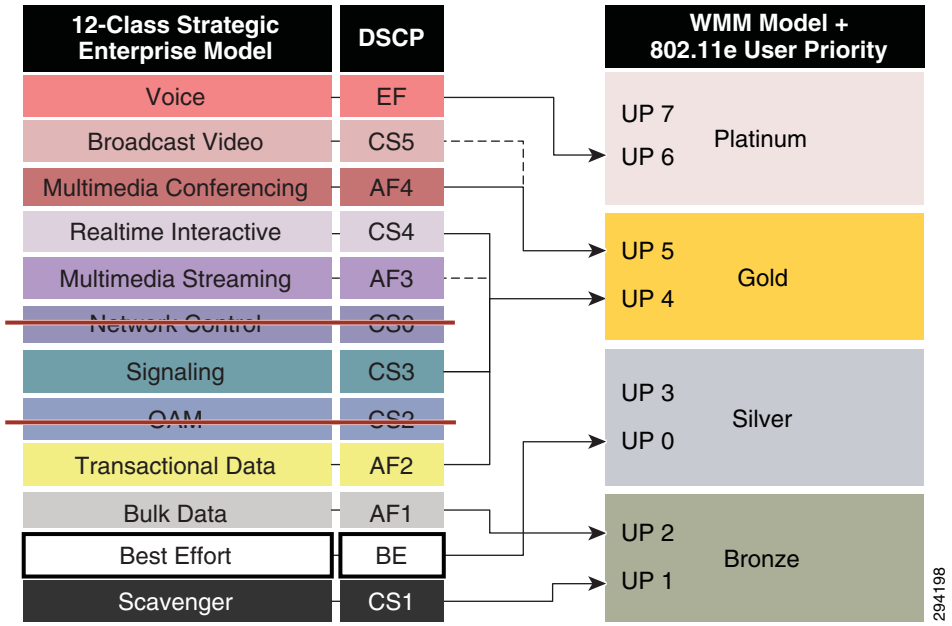
Figure 24-30 Default Downstream Twelve-Class Enterprise Model Mapping to WMM



As with the previous Eight-Class Model, the default Twelve-Class model provisions no QoS for Transactional Data traffic (that may include VDI applications like Citrix XenDesktop or VMware View) nor for Signaling traffic (which is control plane traffic for IP telephone and/or IP video telephony applications)—other than merely protecting these from Scavenger traffic.

Therefore it may be desirable to remap these applications to the Gold WMM access-category, as shown in Figure 24-31.

Figure 24-31 Modified Downstream Twelve-Class Enterprise Model Mapping to WMM



Catalyst 6500 Supervisor 2T Example

This modified twelve-class mapping example is presented on the Catalyst 6500 series switch (with Supervisor 2T). The Catalyst 6500 employs Cisco Common Classification Policy Language (C3PL) QoS and, as such, supports egress class-based marking. Therefore class-based marking policies are only applied on the AP-side network interface, as with the Catalyst 4500, and are shown in Figure 24-30.

Unlike IOS MQC, C3PL does not allow for the combining of class-based marking policies with queuing policies. However both may be simultaneously applied to an interface in a given direction because queuing policies require the additional keywords **type lan-queuing** in their respective class-maps and policy-maps (and class-based marking policies do not).

The Catalyst 6500 supports ingress queuing, but such a policy would technically not be required on an interface connecting to an AP as this interface would not be oversubscribed (due to the AP's capacity being less than 1 Gbps).

Furthermore, the egress queuing policy applied to the interface would remain unchanged, as queuing policies on Catalyst 6500 10/100/1000 interfaces are CoS-based and as such remarked Signaling and Transactional-Data traffic cannot be assigned to the same queues as Signaling. This is because remarked Signaling (DSCP 33) and remarked Transactional Data (DSCP 35) would map to CoS 4, whereas Signaling traffic (DSCP CS3/24) maps to CoS 3. These remarked flows will receive a preferential QoS treatment, but it will be the one ordinarily intended for the video application classes that traditionally align to CoS 4, rather than the Data and Signaling classes that map to CoS 3.

Thus two policies will be present on the AP interface on the Catalyst 6500:

- A class-based marking policy to remap and to restore temporary DSCP values for Signaling and Transactional Data traffic (as shown in Example 24-14).
- An egress C3PL queuing policies to map 12-application classes into a 4 hardware queue (1P3Q8T) model and provide intra-queue QoS via DSCP-based WRED.

Note Since the egress queuing policy remains unchanged, it is not included here for the sake of brevity, but can be referenced from the Medianet Campus QoS Design Guide at http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html.

Therefore, altogether two separate service-policy statements will be applied to a Catalyst 6500 10/100/1000 interface connecting to an AP, as shown in Example 24-14.

- **service-policy output** DOWNSTREAM-WMM-REMARKING
- **service-policy type lan-queuing output** EGRESS-1P3Q8T

Example 24-14 Catalyst 6500 Downstream Class-Based Marking Mapping Policies for WMM

```
 ! This section configures the class-maps
C6500(config-cmap)# class-map match-any SIGNALING
C6500(config-cmap)#  match dscp cs3
 ! Signaling from campus is matched on CS3
C6500(config-cmap)# class-map match-all TRANSACTIONAL-DATA
C6500(config-cmap)#  match dscp af21 af22 af23
 ! Transactional Data from campus is matched on AF2


 ! This section configures the Downstream DSCP-Restoration Policy
C6500(config)# policy-map DOWNSTREAM-WMM-REMARKING
C6500(config-pmap)# class SIGNALING
C6500(config-pmap-c)#  set dscp 33
 ! Signaling is mapped to DSCP 33 for Downstream WMM Gold AC
C6500(config-pmap-c)# class TRANSACTIONAL-DATA
C6500(config-pmap-c)#  set dscp 35
 ! Transactional-Data is mapped to DSCP 35 for Downstream WMM Gold AC


 ! This section applies the downstream remarking policies to AP interface
C6500(config-pmap-c)# interface GigabitEthernet1/10
C6500(config-if)# service-policy output DOWNSTREAM-WMM-REMARKING
 ! Attaches the DOWNSTREAM-WMM-REMARKING policy to the AP interface
C6500(config-if)# service-policy type lan-queuing output EGRESS-1P7Q4T
 ! Attaches the EGRESS-1P3Q8T queuing policy to the interface
```

Note It is recommended to use **match dscp** and **set dscp**—as opposed to **match ip dscp** and **set ip dscp**—as the former will match on both IPv4 and IPv6 packets, whereas the latter will match only on IPv4 packets. Although AVC does not yet classify IPv6 traffic, these packets can still be properly mapped at this node to the correct downstream WMM access category.

This configuration can be verified with the commands:

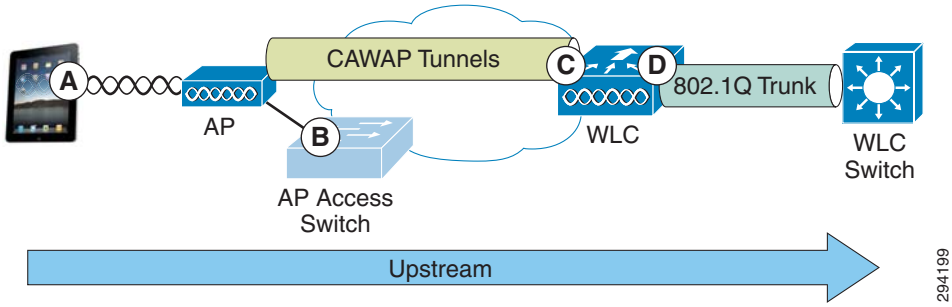
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Configuring Upstream QoS Policies for Mobile Applications

The following policies are required for upstream flows:

- Wireless Device Marking and WMM Policies—QoS policies are embedded within mobile application software and operating systems which can mark UP and DSCP values as well as assign application traffic into the respective WMM access categories in the radio-upstream direction. These policies are shown as point A in Figure 24-32.
- AP Access-Switch QoS Policies—These policies are typically applied on access switch ingress and are shown as point B in Figure 24-32.
- WLC AVC Profiles—Applied on WLC ingress and are shown as point C in Figure 24-15; when configured, these policies apply in both directions-therefore no additional configuration is required to configure AVC QoS in the upstream direction.
- WLC QoS Profiles—Applied on WLC egress and are shown as point D in Figure 24-15; when configured, these policies apply in both directions-therefore no additional configuration is required to configure QoS in the upstream direction.

Figure 24-32 Upstream QoS Policy Configuration Points in Wired/Wireless Networks



Each of these sets of policies will be covered in detail in the respective following sections.

Wireless Device Marking and WMM Policies

The initial upstream over-the-air WMM policies are up to the mobile application vendor to include within their application software and the network administrator has no control over these policies. However, they really only have effect in the upstream "first-hop" to the AP (whereas the majority of traffic to wireless devices flows in the downstream direction, which administrators can control via WLC AVC and QoS profiles).

Nonetheless, it is not a matter of every application on every BYOD device behaving in the same manner. Consider Figure 24-33, which shows various Cisco, Apple, and Samsung phones and their corresponding voice and signaling IP and DSCP marking values.

Figure 24-33 Mobile Wireless IP Phone/Smartphone 802.11 UP and DSCP Marking Values by Hardware Vendor[1]

Device Model	Application	Traffic Type	QoS marking done by Client		DSCP in CAPWAP	DSCP in Wired Packet
			802.11 UP	IP DSCP		
Cisco 7921	N/A	RTP-Vocie	UP-6 (Voice)	EF	EF	EF - VOICE
Cisco 7921	N/A	SCCP-Vocie signalling	UP-4 (Cont Load)	CS3	AF31	N/A
Cisco7925G	N/A	RTP-Vocie	UP-6 (Voice)	EF	EF	EF - VOICE
Cisco7925G	N/A	SCCP-Vocie signalling	UP-4 (Cont Load)	CS3	AF31	N/A
iPhone4	Jabber	RTP-Vocie	UP-5 (Video)	EF	AF41	EF - VOICE
iPhone4	Jabber	SCCP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPhone5	Jabber	RTP-Vocie /Video	UP-5 (Video)	EF	AF41	EF - VOICE
iPhone5	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
Galaxy SII	Jabber	RTP-Vocie	UP-4 (Cont Load)	EF	AF31	EF - VOICE
Galaxy SII	Jabber	SIP-Vocie signalling	UP-3 (Excelent Effort)	CS3	AF21	N/A
HTC Desire C	Jabber	RTP-Vocie				
HTC Desire C	Jabber	SIP-Vocie signalling				

294244

Additionally, 802.11 UP and DSCP markings are virtually non-existent on tablet devices, as shown by Figure 24-34.

Figure 24-34 Mobile Wireless Tablet 802.11 UP and DSCP Marking Values by Hardware Vendor[2]

Device Model	Application	Traffic Type	QoS marking done by Client		DSCP in CAPWAP	DSCP in Wired Packet
			802.11 UP	IP DSCP		
iPad2	Jabber	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad2	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPad2	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad2	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad3	Jabber	RTP-Voice/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad3	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPad3	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad3	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad mini	Jabber	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad mini	Jabber	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A
iPad mini	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
iPad mini	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
Nexus 7						
Nexus 7						
Samsung Galaxy	Policom RealPresence	RTP-Vocie/Video	UP-0 (Best Effort)	00	00	00 - Best Effort
Samsung Galaxy	Policom RealPresence	H.245-Signalling	UP-0 (Best Effort)	00	00	00 - Best Effort

294245

And finally UP and DSCP markings for laptop applications are shown in Figure 24-35, which are similarly absent.

Figure 24-35 Mobile Laptop 802.11 UP and DSCP Marking Values by Hardware Vendor[3]

Device Model	Application	Traffic Type	QoS marking done by Client		DSCP in CAPWAP	DSCP in Wired Packet
			802.11 UP	IP DSCP		
Dell 6420	Jabber (9.1.0)	RTP-Vocie	UP-0 (Best Effort)	00	00	00 - Best Effort
Dell 6420	Jabber (9.1.0)	SCCP-Vocie signalling	UP-0 (Best Effort)	00	00	00 - Best Effort
MacBook Pro	Jabber (8.6.5)	RTP-Vocie/Video	UP-5 (Video)	EF	AF41	EF - VOICE
MacBook Pro	Jabber (8.6.5)	SIP-Vocie signalling	UP-0 (Best Effort)	CS3	00	N/A

294246

1. BYOD with QoS http://mrncciew.com/2013/01/08/byod-with-qos/
2. BYOD with QoS http://mrncciew.com/2013/01/08/byod-with-qos/
3. BYOD with QoS http://mrncciew.com/2013/01/08/byod-with-qos/

Without WMM markings on these wireless client devices, these multimedia applications have reduced quality on their radio upstream connections.

AP Access Switch Upstream QoS Policies

As shown in the above tables, most applications do not yet mark media and signaling flows. However those that do typically mark signaling flows to UP 4 and DSCP CS3. This presents a slight misalignment in QoS policies, as UP 4 will be mapped to (an outer CAPWAP DSCP value of) AF31 (as shown in Figure 24-13) as it is received by the access point (refer to Table 24-4). As such, signaling traffic will not have the proper QoS applied to it in the upstream direction in transit over the wired network between the AP and the WLC. Only when it exits the WLC in the upstream direction is the original inner DSCP value (of CS3) going to be used.

Therefore, to correct this, administrators can configure an ingress marking policy on the access switch interface(s) connecting to wireless access points that receives traffic marked AF31 (which is derived by the AP mapping for UP 4) and remarking these to DSCP CS3. Such remarking policies will now be presented for the same access switches as before, specifically, the:

- Catalyst 3750
- Catalyst 4500
- Catalyst 6500

Catalyst 3750 Configuration Example

The upstream ingress configuration required on a Catalyst 3750 switch to remap Signaling traffic from (the AP UP 4 mapped marking of) AF31 to the enterprise marking of CS3 is shown in Example 24-15.

Example 24-15 Catalyst 3750 Upstream Signaling-Marking Restoration

```
 ! This section configures the class-map
C3750-X(config-cmap)# class-map match-all AP-SIGNALING
C3750-X(config-cmap)#  match ip dscp af31
 ! Signaling traffic from the AP is matched on DSCP AF31


 ! This section configures the Network Upstream Remarking policy-map
C3750-X(config-cmap)# policy-map UPSTREAM-WMM-REMARKING
C3750-X(config-pmap-c)# class AP-SIGNALING
C3750-X(config-pmap-c)#  set dscp cs3
 ! Signaling is remarked to DSCP 24 to align to enterprise QoS model


! This section attaches the upstream policy to the AP-connected interface
C3750-X(config)# interface GigabitEthernet1/0/10
C3750-X(config-if)# mls qos trust dscp
 ! Configures the port to statically trust DSCP on ingress
C3750-X(config-if)# service-policy input UPSTREAM-WMM-REMARKING
 ! Attaches the upstream DSCP remarking policy to the AP interface on ingress
```

This configuration can be verified with the commands:

- **show mls qos interface**
- **show class-map**
- **show policy-map**
- **show policy-map interface**

Catalyst 4500 Configuration Example

The upstream ingress configuration required on a Catalyst 4500 switch to remap Signaling traffic from (the AP UP 4 mapped marking of) AF31 to the enterprise marking of CS3 is shown in Example 24-16.

Example 24-16 Catalyst 4500 Upstream Signaling-Marking Restoration

```
 ! This section configures the class-map
C4500(config-cmap)# class-map match-all AP-SIGNALING
C4500(config-cmap)#  match dscp af31
 ! Signaling traffic from the AP is matched on DSCP AF31


 ! This section configures the Network Downstream Remarking policy-map
C4500(config-cmap)# policy-map UPSTREAM-MAPPING
C4500(config-pmap-c)# class AP-SIGNALING
C4500(config-pmap-c)#  set dscp cs3
 ! Signaling is remarked to DSCP 24 to align to enterprise QoS model


 ! This section attaches the upstream and downstream policies to AP interface
C4500(config)# interface GigabitEthernet 3/1
C4500(config-if)# service-policy input UPSTREAM-MAPPING
 ! Attaches the upstream DSCP remarking policy to the AP interface on ingress
C4500(config-if)# service-policy output 1P7Q1T+DBL+DOWNSTREAM-MAPPING
 ! Attaches the combined egress queuing and remarking policy to the AP interface
```

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map interface**

Catalyst 6500 Configuration Example

The upstream ingress configuration required on a Catalyst 6500 switch to remap Signaling traffic from (the AP UP 4 mapped marking of) AF31 to the enterprise marking of CS3 is shown in Example 24-17.

Example 24-17 Catalyst 6500 Upstream Signaling-Marking Restoration

```
 ! This section configures the class-map
C6500(config-cmap)# class-map match-all AP-SIGNALING
C6500(config-cmap)#  match dscp af31
 ! Signaling traffic from the AP is matched on DSCP AF31


 ! This section configures the Network Downstream Remarking policy-map
C6500(config-cmap)# policy-map UPSTREAM-WMM-REMARKING
C6500(config-pmap-c)# class AP-SIGNALING
C6500(config-pmap-c)#  set dscp cs3
 ! Signaling is remarked to DSCP 24 to align to enterprise QoS model


 ! This section applies the upstream & downstream remarking policies to AP interface
C6500(config-pmap-c)# interface GigabitEthernet1/10
C6500(config-if)# service-policy input UPSTREAM-WMM-REMARKING
 ! Attaches the UPSTREAM-WMM-REMARKING policy to the AP interface on ingress
C6500(config-if)# service-policy output DOWNSTREAM-WMM-REMARKING
 ! Attaches the DOWNSTREAM-WMM-REMARKING policy to the AP interface on egress
C6500(config-if)# service-policy type lan-queuing output EGRESS-1P7Q4T
```

```
! Attaches the EGRESS-1P3Q8T queuing policy to the interface on egress
```

This configuration can be verified with the commands:

- **show class-map**
- **show policy-map**
- **show policy-map interface**

Application Visibility and Management

Cisco AVC for wireless LAN controllers provides application visibility:

- Globally
- On an WLAN basis
- On an individual client basis

Additionally, AVC supports the export of application statistics via Netflow.

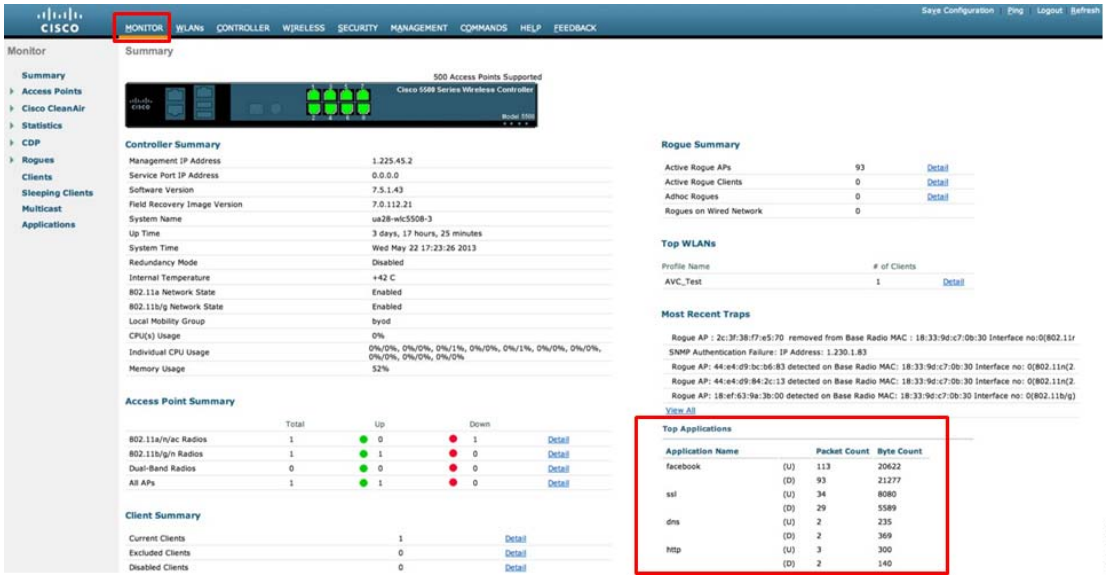
Each of these options is presented in turn.

Note It should be noted that in order to see any application visibility statistics on either a global, WLAN, or client level, the Application Visibility feature must be enabled on at least one WLAN, as shown in Figure 24-17.

Global Application Visibility

To see application traffic at a global level for all traffic traversing the WLC, click the **MONITOR** heading bar and from the main **Summary** screen (in the bottom right corner) the Top 10 applications are summarized by name, as well as by Packet and Byte Count, as shown in Figure 24-36.

Figure 24-36 Global Application Visibility Monitoring



This information can also be collected via CLI by issuing the command:

- **show avc statistics top-apps { upstream | downstream}**

Per-WLAN Application Visibility

To see application traffic at a WLAN-level, perform the following steps.

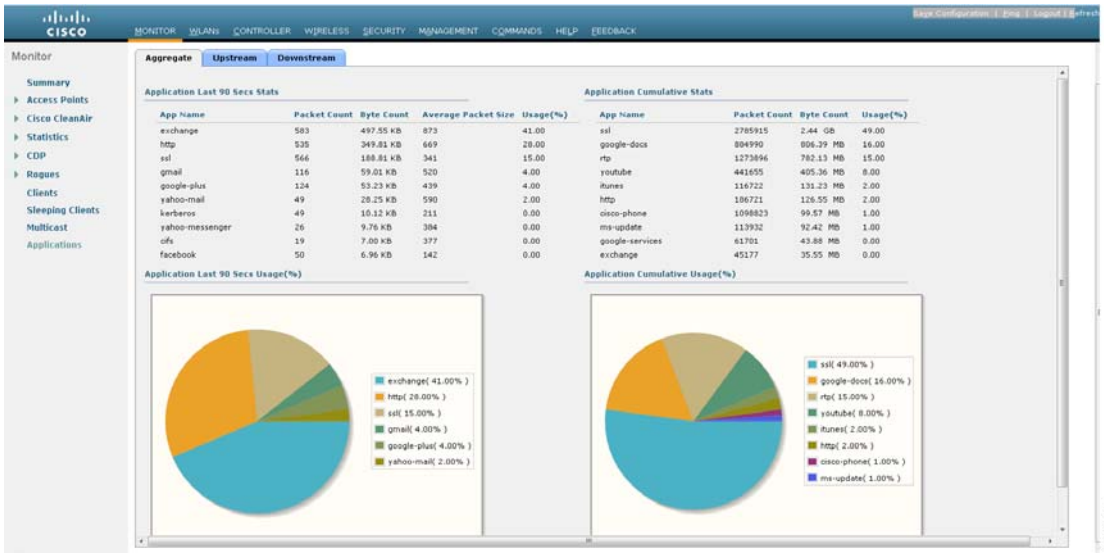
1. Click the **Monitor** heading bar.
2. Select the **Applications** link on the lower-left.
3. Select the **WLAN ID** to be monitored.

Application traffic for the WLAN is summarized in the following tabs:

- **Aggregate**—Includes both upstream and downstream application traffic.
- **Upstream**—Upstream-only application traffic.
- **Downstream**—Downstream-only application traffic.

Per-WLAN application visibility monitoring is shown in Figure 24-37.

Figure 24-37 Per-WLAN Application Visibility Monitoring



This information can also be collected via CLI by issuing the commands:

- **show avc statistics wlan** [*WLAN_Number*] **top-apps { upstream | downstream}**
- **show avc statistics wlan** [*WLAN_Number*] **top-app-groups { upstream | downstream}**
- **show avc statistics wlan** [*WLAN_Number*] **application** [*application_name*]

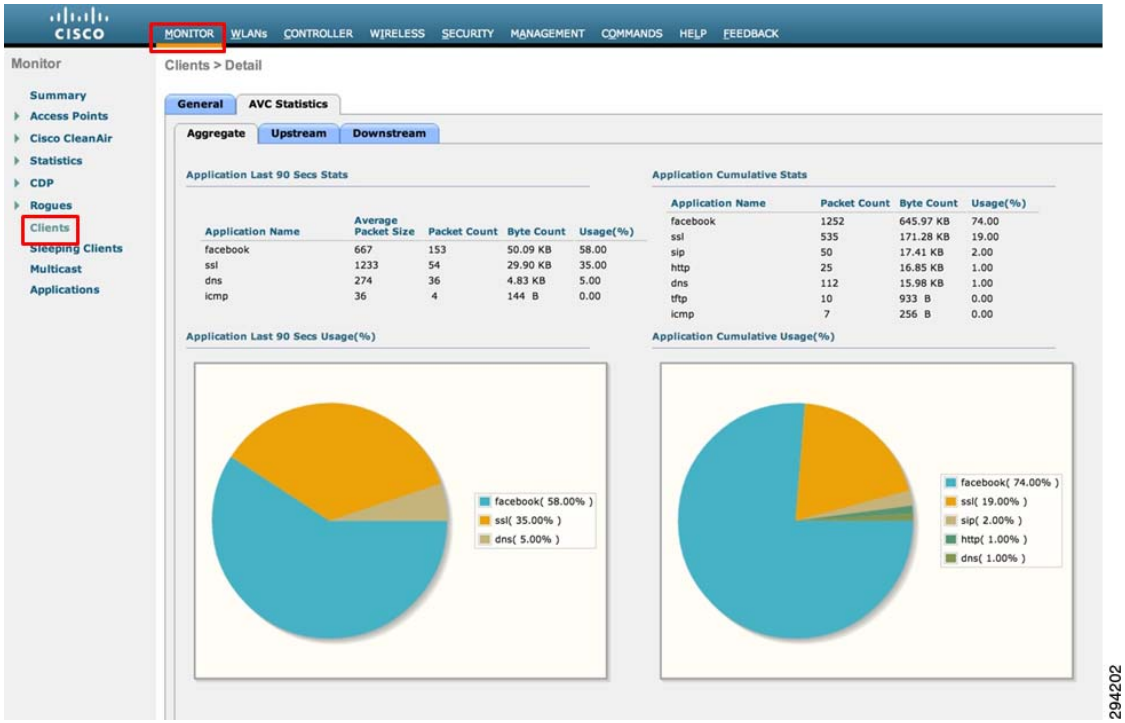
Per-Client Application Visibility

To see application traffic at a client-level, perform the following steps.

1. Click the **Monitor** heading bar.
2. Select the **Clients** link on the left.
3. Select the **Client MAC Addr** of the client to be monitored.

Per-Client application visibility monitoring is shown in Figure 24-38.

Figure 24-38 Per-Client Application Visibility Monitoring



This information can also be collected via CLI by issuing the commands:

- **show avc statistics client** [*client_mac*]
- **show avc statistics client** [*client_mac*] **application** [*application_name*]

NetFlow

NetFlow is a protocol that provides information about network users and applications, peak usage times, and traffic routing. The NetFlow protocol collects IP traffic information from network devices to monitor traffic. The NetFlow architecture consists of the following components:

- Collector—Entity that collects all the IP traffic information from various network elements.
- Exporter—Network entity that exports the template with the IP traffic information. The WLC can be configured to act as an exporter.

NetFlow Configuration-GUI

Netflow Export can be configured on the WLC via the GUI by performing the following steps:

1. Click the **Wireless** heading bar and expand the **Netflow** link and click **Exporter**.
2. Click **New**.
3. Enter the *Exporter name*, *IP address*, and the *port number.*
4. Click **Apply**.
5. Click **Save Configuration**.

NetFlow Monitoring can be configured by following these steps:

1. Click the **Wireless** heading bar and expand the **Netflow** link and click **Monitor**.
2. Click **New** and enter the *Monitor name*.
3. On the Monitor List page, click the Monitor name to open the Netflow Monitor > Edit page.
4. Choose the Exporter name and the Record name from the respective drop-down lists.
5. Click **Apply**.
6. Click **Save Configuration**.

A NetFlow Monitor can be associated with a WLAN by following these steps:

1. Click the **WLANs** heading bar and click the WLAN ID to open the WLANs > Edit page.
2. In the QoS tab, choose the **NetFlow Monitor** from the Netflow Monitor drop-down list.
3. Click **Apply**.
4. Click **Save Configuration**.

NetFlow Configuration-CLI

Create an Exporter by entering this command:

```
config flow create exporter exporter-name ip-addr port-number
```

Create a NetFlow Monitor by entering this command:

```
config flow create monitor monitor-name
```

Associate a NetFlow Monitor with an Exporter by entering this command:

```
config flow add monitor monitor-name exporter exporter-name
```

Associate a NetFlow Monitor with a Record by entering this command:

```
config flow add monitor monitor-name record ipv4_client_app_flow_record
```

Associate a NetFlow Monitor with a WLAN by entering this command:

```
config wlan flow wlan-id monitor monitor-name enable
```

NetFlow configurations can be verified with the following show commands:

- **show flow monitor summary**
- **show flow exporter {summary | statistics}**

Summary

This design document presented the business case for managing applications over wireless networks, highlighting the macro trends in wireless traffic volume growth and application trends. Benefits of applying policies to manage application quality were presented, including increasing voice and video quality, business-critical application responsiveness, as well as controlling background applications and non-business applications over wireless networks.

Next, to set design context, wireless QoS tools were overviewed to show how these evolved and operate, highlighting both their capabilities as well as their limitations. Following this, a discussion of how Layer 2 and Layer 3 mapping works over Cisco wired and wireless networks was presented, including a QoS Translation Table that performs non-default mappings to reconcile IEEE Layer 2 markings with IETF Layer 3 markings in Cisco WLCs and APs.

Subsequently, the various policies required to ensure QoS in both the downstream and upstream directions were summarized, including:

- QoS Profiles
- AVC Profiles
- Network switch DSCP-mapping
- Mobile Device WMM marking

Each of these main policy elements were then discussed in detail to show how these could be configured. WLC policies configuration—namely QoS and AVC Profiles—were presented both in GUI format and in CLI commands.

Wired network switch policies for the access switches connecting to Cisco wireless access points were presented for a Four-Class, Eight-Class, and Twelve-Class enterprise strategic application-class models mapped to WMM. Additionally, these policies configurations were shown for Catalyst 3750, 4500, and 6500 series switches, highlighting the platform-specific idiosyncrasies in function and CLI relating to the policy configurations.

Finally, administrators were shown how to monitor application visibility on a global level, WLAN level, and a client level, as well as how to configure NetFlow Export and Monitoring for network management purposes.

Additional Reading

- Cisco Wireless LAN Controller Configuration Guide, Release 7.4
  http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED.html
  - Configuring QoS
    http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01110.html
  - Configuring Application Visibility and Control
    http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01111.html
  - Working With WLANS-Assigning QoS Profiles
    http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01010111.html
- Medianet QoS 4.0 Design:
  - Strategic QoS Design Overview 4.0
    http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSIntro_40.html
  - Campus QoS Design 4.0
    http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html#wp1099462



CHAPTER 25

Managing Bonjour Services for BYOD

Revised: August 7, 2013

Executive Summary

This chapter focuses on how to use the Cisco Wireless LAN Controller software Bonjour Gateway feature to manage Apple's Bonjour protocol in a BYOD enterprise context.

Bonjour is Apple's zero-configuration protocol for advertising, discovering, and connecting to network services like file sharing, print sharing, media sharing, etc. The Bonjour protocol was originally designed for home network use and utilizes Multicast Domain Name Services (mDNS) via link-local multicasting to share network services. While this approach works well in home networks, a limitation of link-local multicasting is that these network services will only be shared within a single Layer 2 domain (such as a VLAN or WLAN). In a BYOD enterprise scenario, different WLANs and VLANs are used for different classes of devices, including corporate devices, employee devices, personal devices, and guest devices (as well as quarantine WLANs for unapproved devices). As such, basic Bonjour operations—such as printing to a wired printer from a wireless LAN—may not be natively supported.

To address this limitation and to facilitate the user demand of BYOD for Apple devices within the enterprise, Cisco has developed the Bonjour Gateway feature for its Wireless LAN Controllers (WLCs). This feature was introduced in Cisco WLC software version 7.4 and solves the Layer 2 domain limitation for Bonjour by allowing the WLC to snoop, cache, and proxy-respond to Bonjour service requests that may reside on different Layer 2 domains. Additionally, these responses may be selectively controlled by administrative policies, so that only certain Bonjour services will be permitted in specific Layer 2 domains.

This chapter provides an overview of the Bonjour protocol and shows how the Bonjour Gateway feature functions, as well as how it can be practically deployed in an enterprise BYOD context to manage Bonjour services. To this end, step-by-step configuration guidance and verification commands are presented, both for the Cisco WLC GUI as well as the Command Line Interface (CLI).

Why Bonjour?

Bonjour is Apple's implementation of a suite of zero-configuration networking protocols and is supported on both Mac OS X devices (such as laptops and desktops), as well as on Apple iOS devices (such as iPhones and iPads). Bonjour is designed to make network configuration easier for users.

For example, consider enabling IP-based print services. Each printer needs a unique IP address, whether statically assigned or dynamically assigned (by a DHCP server). Since dynamically-assigned addresses can change, most printers are manually configured with a static address so that computers on the network can reach them using the same address every time. In this case, each client device must know the statically configured IP address of the printer(s) in order to use these. To make the process more user friendly, network administrators may configure DNS records so that clients can access printers by name, rather than by specific IP addresses. Even so, the clients must know the specific DNS name of each printer they are trying to access. Thus, the seemingly minor task of enabling IP-based printing can require significant client and server configuration. Additionally, in a home network environment, people who do not fit the traditional role of the network administrator often set up networks (e.g., families connecting their laptops and personal devices to the Internet over a shared router). As such, this level of configuration simply is not practical in such a setting.

Consider the same example in a network running Bonjour. Bonjour lets you connect a printer to your network without assigning it a specific IP address or manually entering that address into each computer. With zero-configuration networking, nearby computers can discover its existence and automatically determine the printer's IP address. If that address is a dynamically assigned address that changes, they can automatically discover the new address in the future.

Bonjour functionality is not limited to printing and includes:

- File Sharing Services
- Remote Desktop Services
- Full screen Mirroring (Apple iOS v5.0+ for iPad2, iPhone4S, or later)
- iTunes Services:
  - iTunes File Sharing
  - iTunes Wireless iDevice Syncing (Apple iOS v5.0+)
  - Music broadcasting (Apple iOS v4.2+)
  - Video broadcasting (Apple iOS v4.3+)

Bonjour's zero-configuration networking services benefit not only users (who will no longer have to assign IP addresses or host names to access network services), but also applications (as applications can leverage Bonjour to automatically detect required services or to interact with other applications to allow for automatic connection, communication, and data exchange, all without any user configuration).

Bonjour Overview

Bonjour offers zero-configuration solutions for three areas of IP networking:

- Addressing (allocating IP addresses to hosts)—Bonjour Addressing
- Naming (using names to refer to hosts instead of IP addresses)—Bonjour Naming
- Service discovery (finding services on the network automatically)—Bonjour Service Discovery

Each of these areas is discussed in turn, as well as how Bonjour optimizes the delivery of these solutions.

Bonjour Addressing

Bonjour solves the addressing problem of allocating IP addresses to hosts by leveraging self-assigned link-local addressing. Link-local addressing uses a range of addresses reserved for the local network and is achieved differently by IPv6 and IPv4:

- IPv6 includes self-assigned link-local addressing as part of the protocol
- IPv4 self-assigned addressing works by picking a random IP address in the link-local range and testing it. If the address is not in use, it becomes the local address. If it is already in use, the computer or other device chooses another address at random and tries again.

Any user or service on a computer or iOS device that supports link-local addressing benefits from this feature automatically. When a host computer joins a local network, it finds an unused local address and adopts it. No user action or configuration is required.

Bonjour Naming

Bonjour leverages Multicast DNS (mDNS) for name-to-address translation, which sends DNS-format queries over the local network using an IP multicast address. Because these DNS queries are sent to a multicast address, no single DNS server with global knowledge is required to answer the queries. Each service or device can provide its own DNS capability—when it sees a query for its own name, it provides a DNS response with its own address.

Actually, Bonjour goes a bit further than basic mDNS functionality by including a responder that handles mDNS queries for any network service on the host computer or iOS device. This relieves an application of the need to interpret and respond to mDNS messages. Once a service is registered with the Bonjour process, Bonjour automatically advertises the availability of the service so that any queries for it are directed to the correct IP address and port number automatically.

Note Registration is performed using one of the Bonjour APIs. This functionality is available only to services running on the host OS X computer or iOS device. Services running on other devices, such as printers, need to implement a simple mDNS responder daemon that handles queries for services provided by that device (which is included on printers supporting the Apple AirPrint feature).

Bonjour also provides built-in support for the NAT port mapping protocol (NAT-PMP). If the upstream router supports this protocol, OS X and iOS applications can create and destroy port mappings to allow hosts on the other side of the firewall to connect to the provided services.

For name-to-address translation to work properly, a unique name on the local network is necessary. Unlike conventional DNS host names, the local name only has significance on the local network or LAN segment. A local name can be assigned much the same way as a self-assigned a local address: a name is chosen and if it is not already in use, it gets used. If it is unavailable, then the name can be modified slightly and re-tested for availability. For example, if a printer with the default name XYZ-LaserPrinter.local attaches to a local network with two other identical printers already installed, it tests for XYZ-LaserPrinter.local, then XYZ-LaserPrinter-**2**.local, then XYZ-LaserPrinter-**3**.local, which is unused and which becomes its name.

Bonjour Naming Rules

This section explains the Bonjour local “domain” and the naming rules for Bonjour service instances and service types. These service names are snooped by and presented within the Cisco WLC and as such are helpful for an administrator to understanding.

Bonjour protocols deal primarily with local link service advertisements. A host’s link-local network includes itself and all other hosts that can exchange packets without IP header data being modified (i.e., hosts sharing a single layer 2 domain/VLAN). In practice, this includes all hosts not separated by a router. On Bonjour systems, “local.” is used to indicate a name that should be looked up using an mDNS query on the local IP network.

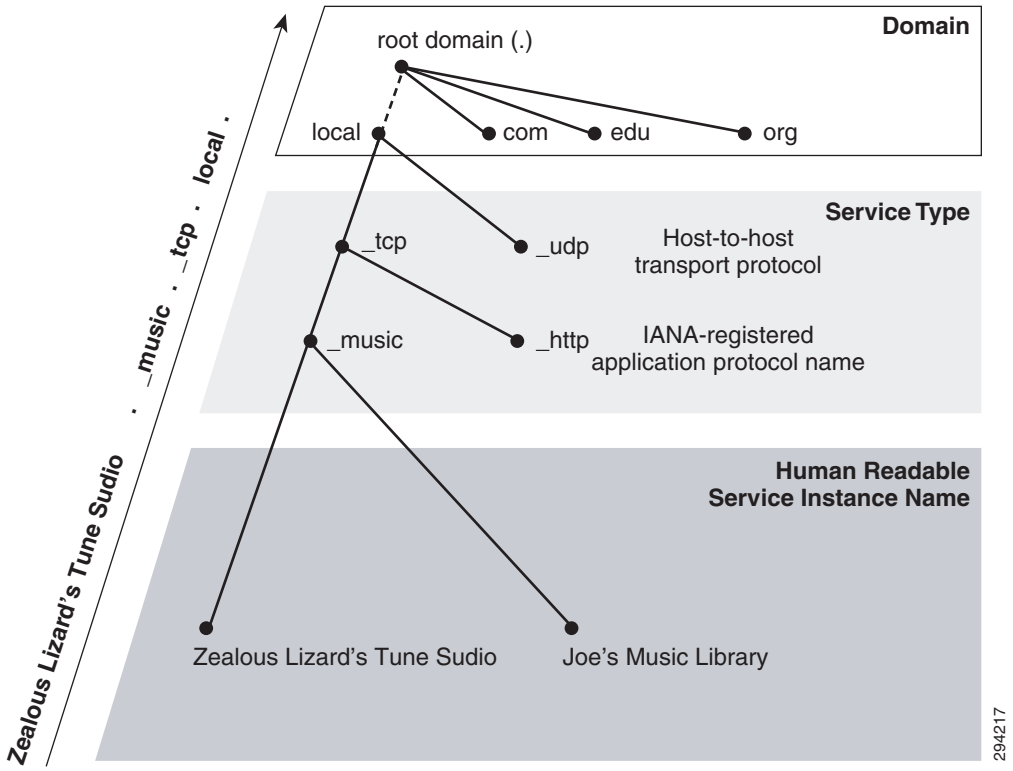
Note that “local.” is not really a domain, but rather a pseudo-domain. It differs from conventional DNS domains in a fundamental way: names within DNS domains are globally unique; link-local domain names are not. As such, local names are useful only on the local network. In many cases this is adequate, as these provide a way to refer to network devices using names instead of IP numbers and of course they require less effort to coordinate and administer as compared to globally unique names.

Locally unique names are particularly useful on networks that have no connection to the global Internet, either by design or because of interruption, and on small, temporary networks, such as a pair of computers linked by a crossover cable or a few people playing network games using laptops on the wireless network of a home or cafe.

Note If a name collision on the local network occurs, a Bonjour host finds a new name automatically (in the case of an iOS device) or by asking the user (in the case of an OS X personal computer).

Bonjour service instance names are intended to be user-readable strings with descriptive names. Figure 25-1 illustrates the organization of the name of a Bonjour service instance. At the top level of the tree is the domain, such as “local.” for the local network. Below the domain is the registration type, which consists of the service type preceded by an underscore (_music) and the transport protocol, also preceded by an underscore (_tcp). At the bottom of the tree is the human-readable service instance name, such as Zealous Lizard's Tune Studio. The complete name is a path along the tree from bottom to top, with each component separated by a dot.

Figure 25-1 Bonjour Service Name Hierarchy and Organization



Other Bonjour service name suffixes include:

- _ipp._tcp.local. for AirPrint Printers
- _printer._tcp.local. for generic IP Printers

- _airplay._tcp.local. for AppleTV

Bonjour Service Discovery

The final element of Bonjour is service discovery. Service discovery allows applications to find all available instances of a particular type of service and to maintain a list of named services and port numbers. The application can then resolve the service hostname to a list of IPv4 and IPv6 addresses, as previously described.

The list of named services provides a layer of indirection between a service and its current DNS name and port number. Indirection allows applications keep a persistent list of available services and resolve an actual network address just prior to using a service. The list allows services to be relocated dynamically without generating a lot of network traffic announcing the change.

Service discovery in Bonjour is accomplished by "browsing." An mDNS query is sent out for a given service type and domain, and any matching services reply with their names. The result is a list of available services to choose from.

This is very different from the traditional device-centric paradigm of network services, which describes services in terms of physical hardware. In a device-centric view, the network consists of a number of devices or hosts, each with a set of services. In a device-centric browsing scheme, a client queries the server for what services it is running, gets back a list (FTP, HTTP, print-services and so on), and decides which service to use. The interface reflects the way the physical system is organized. But this is not necessarily what the user logically wants or needs.

On the other hand, a service-centric paradigm is typically more logical and efficient from a user-perspective. Users typically want to accomplish a certain task, not query a list of devices to find out what services are running. It makes far more sense for a client to ask a single question, "What print services are available?" than to query each available device with the question, "What services are you running?" and sift through the results looking for printers. The device-centric approach is not only time-consuming, but it also generates a significant amount of irrelevant network traffic. In contrast, the service-centric approach sends a single query, generating only relevant replies.

Bonjour takes the service-oriented view. Queries are made according to the type of service needed, not the hosts providing them. Applications store service instance names, not addresses, so if the IP address, port number, or even host name has changed, the application can still connect. By concentrating on services rather than devices, the user's browsing experience becomes more relevant and efficient.

Bonjour Optimization

Server-free addressing, naming, and service discovery have the potential to create a significant amount of excess network traffic, but Bonjour uses several mechanisms to reduce this traffic to a minimum to avoid unnecessary "chattiness", including:

- Caching
- Suppression of Duplicate Responses
- Exponential Back-Off and Service Announcement

Each of these Bonjour optimization mechanisms is briefly described in the following sections.

Caching

Bonjour uses a cache of mDNS records to prevent hosts from requesting information that has already been requested. For example, when one host requests, say, a list of print spoolers, the list of printers comes back via multicast, so all local hosts see it. The next time a host needs a list of print spoolers, it already has the list in its cache and does not need to reissue the query.

Suppression of Duplicate Responses

To prevent repeated answers to the same query, Bonjour service queries include a list of known answers. For example, if a host is browsing for printers, the first query includes no print services and gets, say, twelve replies from available print servers. The next time the host queries for print services, the query includes a list of known servers. Print servers already on the list do not respond.

Bonjour also suppresses duplicate responses in another way. If a host is about to respond, and notices that another host has already responded with the same information, the host suppresses its response.

Exponential Back-off and Service Announcement

When a host is browsing for services, it does not continually send queries to see if new services are available. Instead, the host issues an initial query and sends subsequent queries exponentially less often, for example: after 1 second, 3 seconds, 9 seconds, 27 seconds, and so on, up to a maximum interval of one hour.

This does not mean that it can take over an hour for a browser to see a new service. When a service starts up on the network, it announces its presence a few times using a similar exponential back-off algorithm. This way, network traffic for service announcement and discovery is kept to a minimum, but new services are seen very quickly.

Cisco Bonjour Gateway Solution in WLC 7.4+

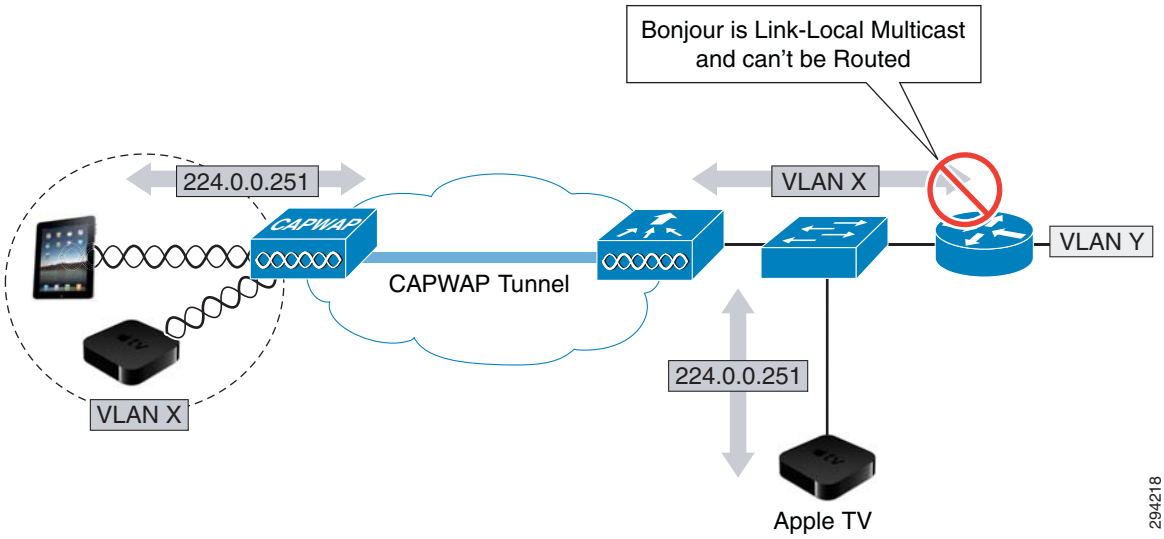
As previously discussed, the Bonjour protocol uses mDNS queries. These queries are sent over UDP port 5353 to the reserved group addresses listed below:

- IPv4 Group Address: 224.0.0.251
- IPv6 Group Address: FF02::FB

However it should be noted that the mDNS addresses used by Bonjour are link-local multicast addresses and are only forwarded within the local Layer 2 domain, as link-local multicast is meant to stay local by design. Furthermore, routers cannot even use multicast routing to redirect the mDNS queries, because the time-to-live (TTL) of these packets is set to 1.

Bonjour was originally developed with home networks in mind. As such, since most home networks consist of a single Layer 2 domain, this link-local limitation of mDNS rarely posed any practical deployment constraints. However in an enterprise context, where large numbers of (wired and wireless) Layer 2 domains exist, this limitation severely handicaps Bonjour functionality, as Bonjour clients would only see locally-hosted services and would not be able to see or connect to services hosted on other subnets. This link-local multicast limitation of Bonjour mDNS is illustrated in Figure 25-2.

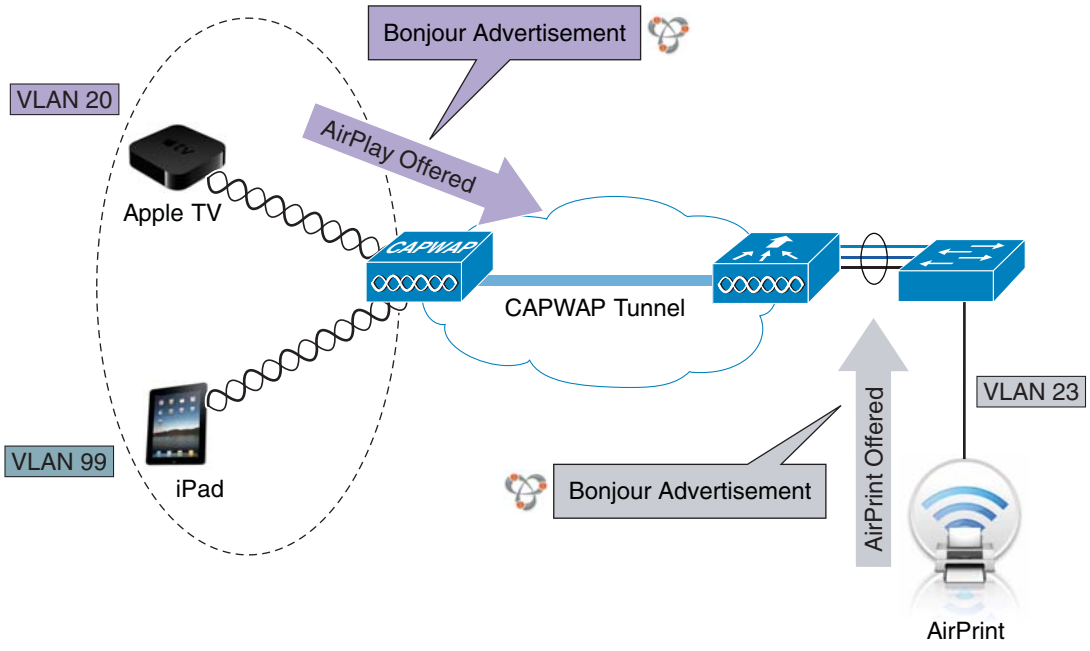
Figure 25-2 Bonjour Deployment Limitation in Enterprise Networks



To address this limitation and to facilitate BYOD functionality on enterprise networks, Cisco released a Bonjour Gateway feature in WLC 7.4+ software. The Bonjour Gateway feature (technically speaking a mDNS gateway feature, but most relevantly applied to Bonjour) snoops and caches all Bonjour service advertisements across multiple VLANs and can be configured to (selectively) reply to Bonjour queries. Figure 25-3 through Figure 25-5 illustrate the operation of the Bonjour Gateway.

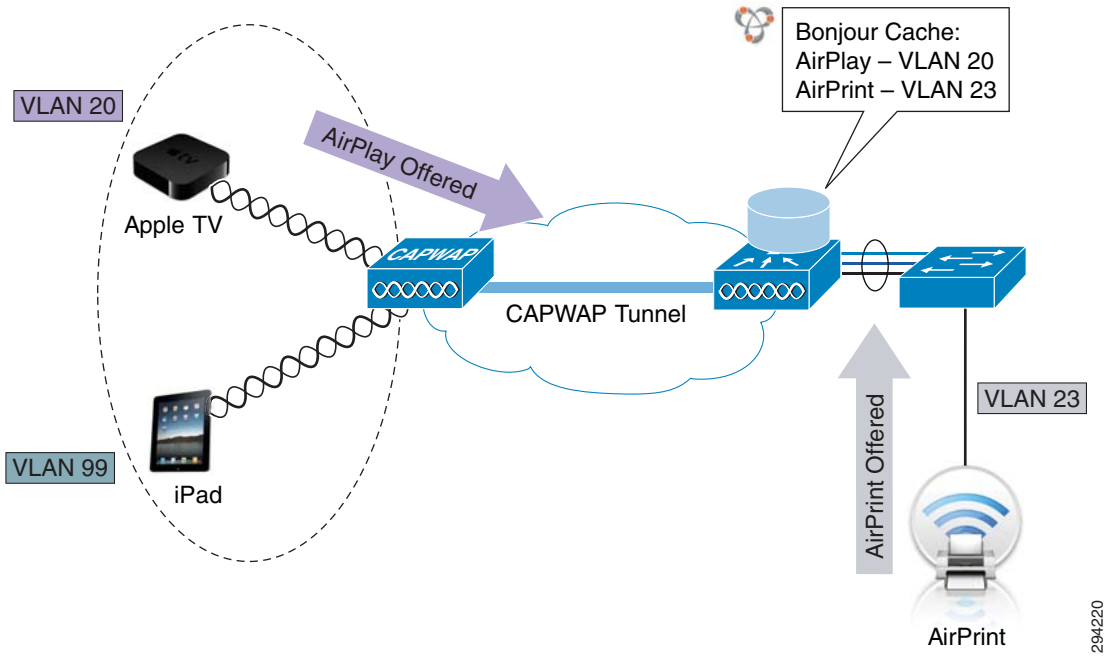
In Figure 25-3, the Bonjour Gateway listens/snoops all Bonjour advertisements.

Figure 25-3 Cisco WLC Bonjour Gateway Operation—Step 1—Bonjour Service Advertisement Snooping



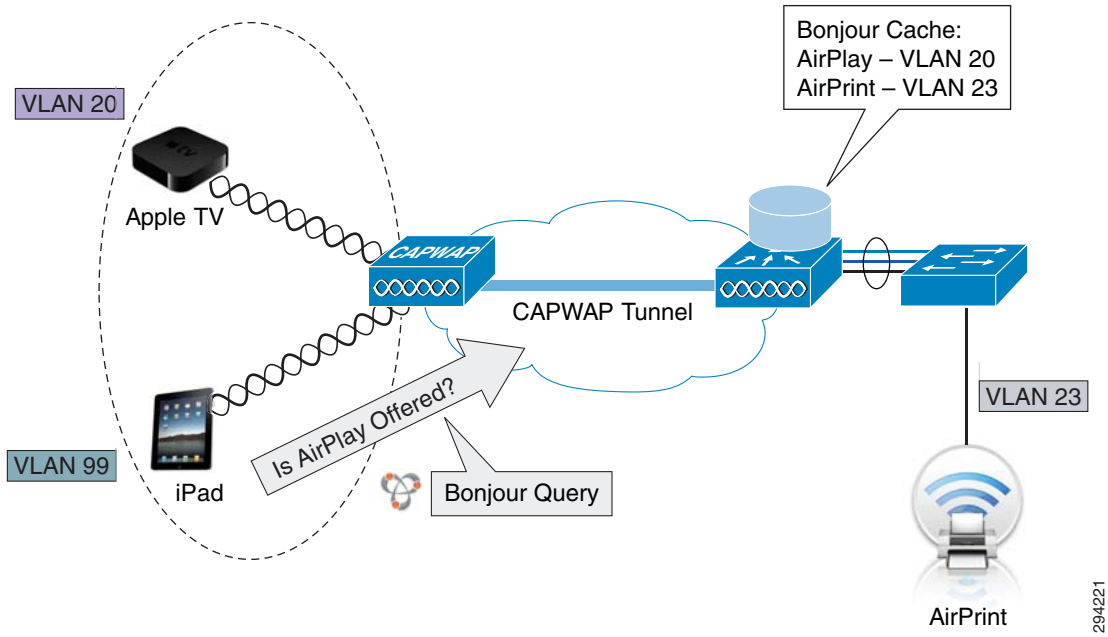
Next, the Bonjour Gateway caches all these service advertisements, as shown in Figure 25-4. Incidentally, the WLC 7.4 release supports up to 64 services and 100 service providers per service type. Each service provider is registered in the WLC as its domain name. Additionally, each Bonjour service has an advertised TTL (which is different from a packet's TTL) and the controller asks the device for an update at 85% of this TTL.

Figure 25-4 Cisco WLC Bonjour Gateway Operation—Step 2—Service Advertisement Caching



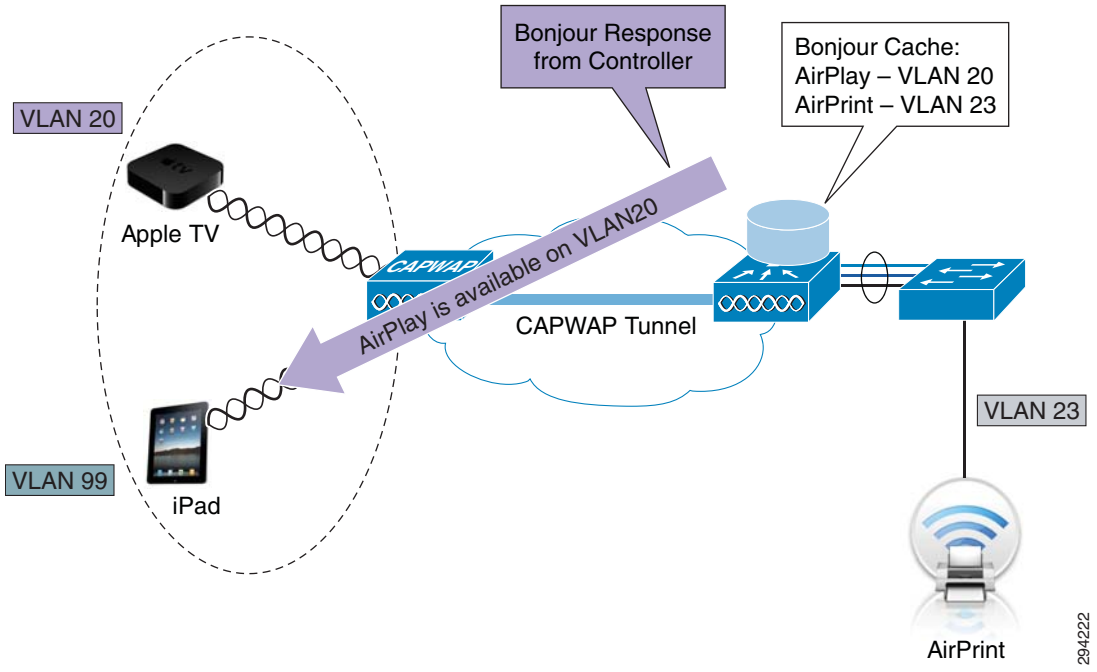
In addition to listening to service advertisements, the WLC is always listening for client queries for services, as illustrated in Figure 25-5.

Figure 25-5 Cisco WLC Bonjour Gateway Operation—Step 3—Bonjour Query Snooping



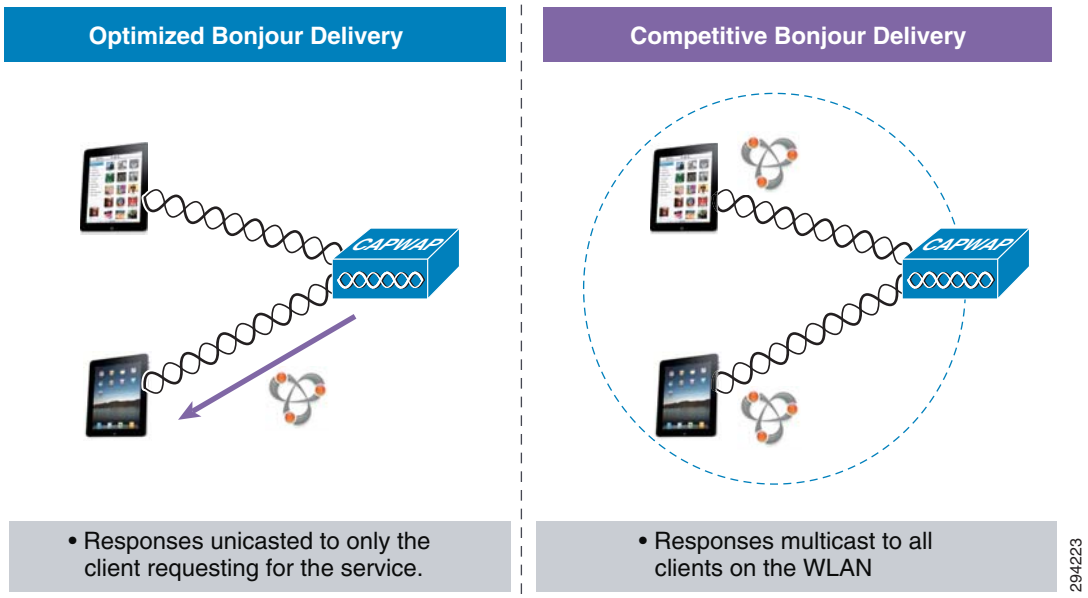
Clients that request locally-hosted services will receive unicast replies from the service provider; however clients that request services that may be hosted on other VLANs will receive unicast responses from the WLC, as shown in Figure 25-6.

Figure 25-6 Cisco WLC Bonjour Gateway Operation—Step 4—Bonjour Query Response (from Cache)



And finally, the Bonjour Gateway service can serve to further optimize Bonjour traffic by unicasting replies directly to clients requesting a given service (as opposed to multicasting replies like some competitive solutions), making more efficient use of network resources, as shown in Figure 25-7.

Figure 25-7 Cisco WLC Bonjour Gateway Operation versus Competitive Offering Operation



Bonjour Gateway Service Policy Deployment Options

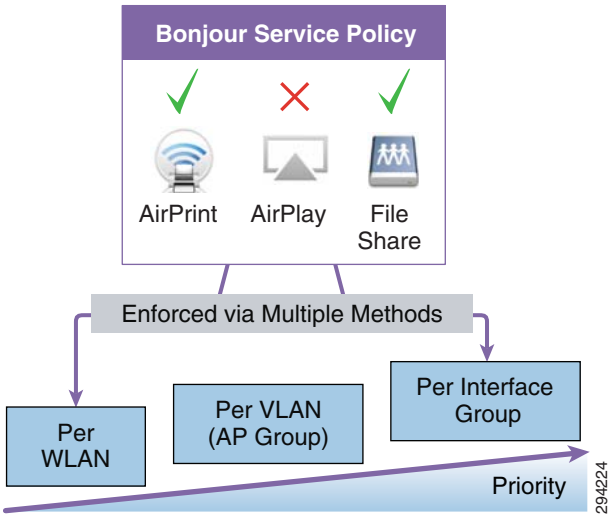
A key functional advantage of the Bonjour Gateway is that it can be configured to selectively reply to Bonjour service requests, thus allowing for administrative control of Bonjour services within the enterprise. Bonjour policies can be applied on the following basis:

- Per WLAN
- Per VLAN
- Per Interface/Interface-Group

Note Per-User Bonjour policy application is planned for a future release via RADIUS AAA-Override.

These Bonjour service policy options are illustrated in Figure 25-8.

Figure 25-8 Cisco WLC Bonjour Gateway Service Policy Deployment Options

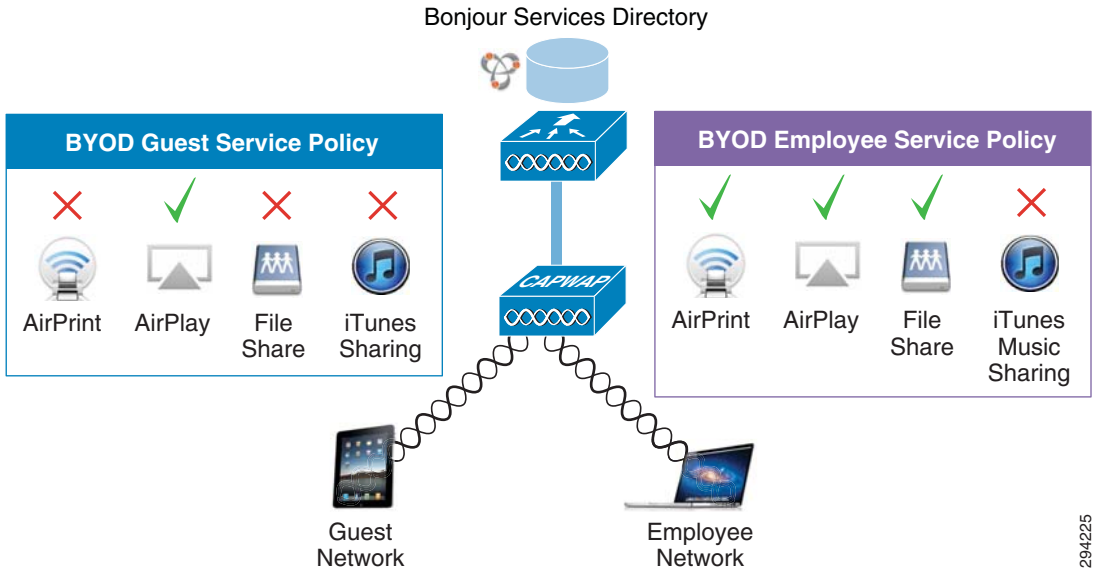


Consider a few examples of how such Bonjour service policies may be deployed. For instance, in an BYOD enterprise context, you can configure Bonjour policies such that employees can take advantage of Bonjour services that enhance productivity (such as AirPrint, AirPlay, and File Sharing), but block entertainment-oriented Bonjour services (such as iTunes Sharing).

Additionally, stricter limitations could be placed on Guest WLANs. For example, inter-domain Bonjour services could be limited to AirPlay only—such that guest devices may be allowed to connect to (wired or wireless) AppleTVs that reside on the production network—so that guests could share presentations, videos, demonstrations, etc.

These example Bonjour service policies for an enterprise BYOD deployment context are illustrated in Figure 25-9.

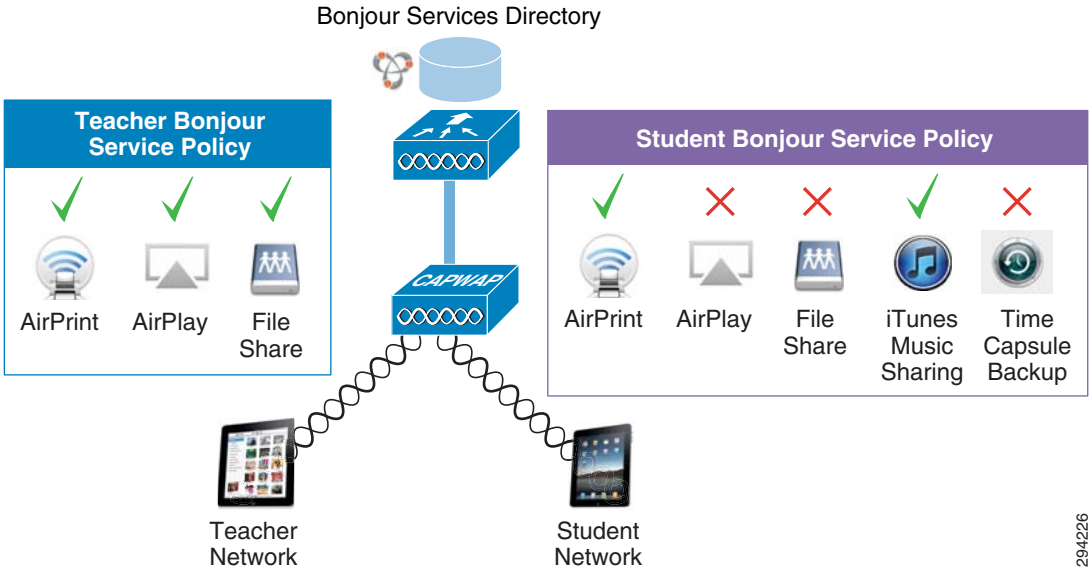
Figure 25-9 Cisco WLC Bonjour Gateway Service Policy Deployment Example 1—A BYOD Enterprise



It is important to note that these Bonjour service policy examples are not a one-size-fits-all solution. The policy-specifics will likely vary according to deployment contexts. As a second example consider a college/university deployment context. In this example, assume separate WLANs for teachers and students. Teachers would likely have all Bonjour productivity-oriented services enabled, such as AirPrint, AirPlay, and File Sharing. However you may wish to limit AirPlay on student networks, as this may prevent significant volumes of traffic traversing different WLANs as students may host full-length HD movies on one network while streaming them to devices on another. Similarly, Time Capsule traffic may be another service to limit from spanning WLANs—again due to the significant traffic loads these typically entail. However, consideration may be extended students by permitting iTunes Music Sharing (as music files are significantly smaller than videos or Time-Capsule backups).

These example Bonjour service policies for a university BYOD deployment context are illustrated in Figure 25-10.

Figure 25-10 Cisco WLC Bonjour Gateway Service Policy Deployment Example 2—A BYOD University



While the specifics of a Bonjour service policy may differ according to deployment context, there are two broad use cases for Bonjour Gateway deployments that are discussed next.

Bonjour Gateway BYOD Use Cases and Configuration Examples

There are effectively two general use cases for Bonjour Gateway service policy deployments:

- Wireless-to-Wired Bonjour Gateway Service Policies—The primary use case is enabling wireless BYOD devices to print to wired AirPrint printers.
- Wireless-to-Wireless Bonjour Gateway Service Policies—Enables Bonjour services to be shared among devices in separate WLANs; an example use case would be to allow guest devices to access wireless AppleTVs to share presentations (even though these devices may reside in different WLANs).

Bonjour service policies on Cisco WLCs can be configured using one of two approaches:

- Editing the default mDNS profile

- Creating new mDNS profiles

Also, mDNS profiles can be applied directly to:

- Interfaces/Interface-Groups
- VLANs
- WLANs

To highlight deployment options, the examples that follow utilize a variety of these options.

Design configuration are presented both via the Cisco WLC GUI and the Cisco WLC CLI. CLI examples show both the general syntax of a command (which is highlighted in **blue**) and the specific variation needed in the design example (which is highlighted in **red**).

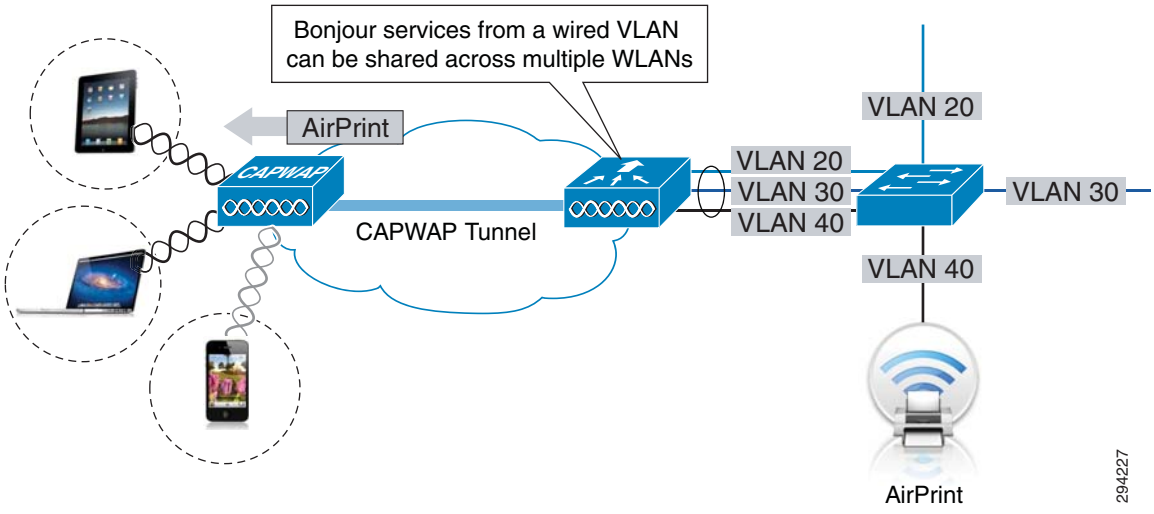
Note In these design examples, it is assumed that the network infrastructure and WLC have been configured in accordance with the best-practice BYOD designs presented in this CVD.

Use Case 1—Wireless-to-Wired Bonjour Gateway Service Policy—BYOD Employee AirPrint Example

In this primary Bonjour Gateway use case, wireless BYOD employee devices are permitted to access AirPrint-enabled printers that are deployed on separate wired networks. Incidentally, this design will also support wireless printing from wireless clients across separate WLANs.

A prerequisite of this design is that the wired VLANs hosting AirPrint printers must be trunked to the Cisco WLC controller, as shown in Figure 25-11.

Figure 25-11 Use-Case 1—Cisco WLC Bonjour Gateway Wireless-to-Wired Design Example



Multiple design and configuration options exist to enable Bonjour service policies. In this example, Bonjour service policies will be configured by:

- Step 1—Globally enabling mDNS snooping
- Step 2—Editing the default mDNS profile
- Step 3—Applying the default profile to an interface

Each of these steps is detailed in turn (with additional design options being presented in the following example).

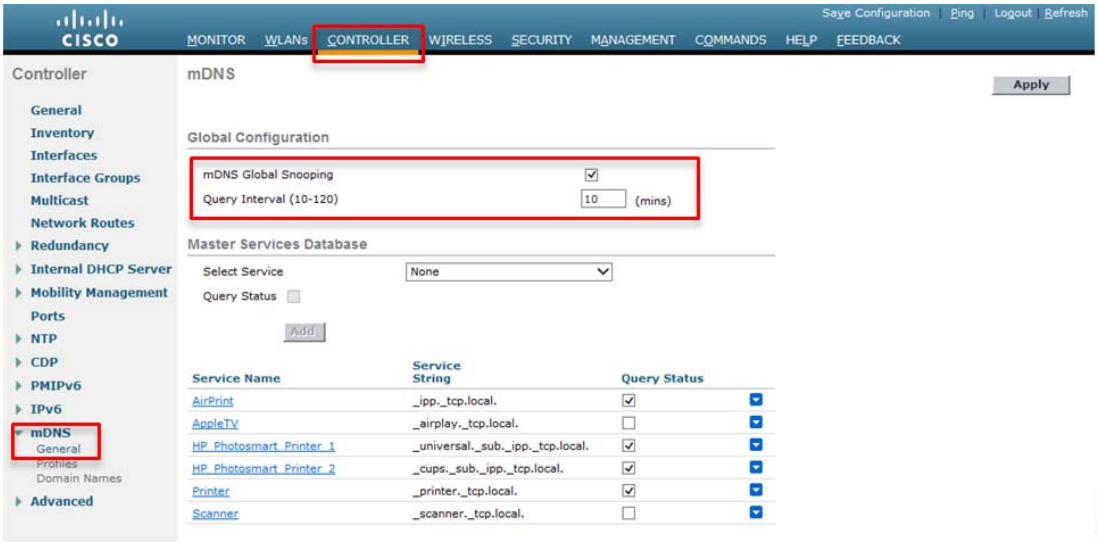
Step 1—Enable mDNS Global Snooping

The first step is to globally enable mDNS snooping by doing the following:

1. Open a web browser to the Cisco WLC IP address via HTTPS and login.
2. Click the **CONTROLLER** heading-bar and expand the **mDNS** link on the lower left and click **General**.
3. Under the **Global Configuration** heading, select the checkbox to enable **mDNS Global Snooping**.
4. Optionally the mDNS Snooping **Query Interval** can be tuned (from 10 min. to 120 min.).

These steps are shown in Figure 25-12.

Figure 25-12 Use-Case 1—Step 1—Enabling mDNS Global Snooping



The corresponding Cisco WLC CLI for globally enabling mDNS snooping is shown in Example 25-1.

Example 25-1 Enabling mDNS Global Snooping

```
General command/specific example:
(Cisco Controller) >config mdns snooping enable
! Globally enables mDNS snooping
```

The mDNS snooping query interval can be tuned with the command shown in Example 25-2 (again the range is 10 to 120 minutes). Example 25-2 shows both the general version of this command and the specific syntax to set the mDNS query interval to 10 minutes.

Example 25-2 Tuning the mDNS Query Interval

```
General command:
(Cisco Controller) >config mdns query interval minutes

Specific example:
(Cisco Controller) >config mdns query interval 10
```

```
! Sets the mDNS query interval to 10 minutes
```

These mDNS configuration commands can be verified by the **show network summary** command output, as illustrated in Example 25-3.

Example 25-3 Verifying mDNS Global Snooping and Query Interval—show network summary

```
(Cisco Controller) >show network summary

RF-Network Name.............................. byod
Web Mode..................................... Disable
Secure Web Mode.............................. Enable
Secure Web Mode Cipher-Option High.......... Disable
Secure Web Mode Cipher-Option SSLv2......... Disable
Secure Web Mode RC4 Cipher Preference....... Disable
OCSP......................................... Disabled
OCSP responder URL...........................
Secure Shell (ssh)........................... Enable
Telnet....................................... Enable
Ethernet Multicast Forwarding................ Disable
Ethernet Broadcast Forwarding................ Disable
IPv4 AP Multicast/Broadcast Mode............. Unicast
IGMP snooping................................ Disabled
IGMP timeout................................. 60 seconds
IGMP Query Interval.......................... 20 seconds
MLD snooping................................. Disabled
MLD timeout.................................. 60 seconds
MLD query interval........................... 20 seconds
User Idle Timeout............................ 300 seconds
ARP Idle Timeout............................. 300 seconds
Cisco AP Default Master...................... Disable
AP Join Priority............................. Disable
Mgmt Via Wireless Interface.................. Enable
Mgmt Via Dynamic Interface................... Disable
Bridge MAC filter Config..................... Enable
Bridge Security Mode......................... EAP
Mesh Full Sector DFS......................... Enable
AP Fallback ................................. Enable
Web Auth CMCC Support ....................... Disabled
Web Auth Redirect Ports ..................... 80
Web Auth Proxy Redirect  .................... Disable
Web Auth Captive-Bypass   ................... Enable
Web Auth Secure Web  ........................ Enable
Fast SSID Change ............................ Enabled
AP Discovery - NAT IP Only .................. Enabled
IP/MAC Addr Binding Check ................... Enabled
CCX-lite status ............................. Disable
oeap-600 dual-rlan-ports .................... Disable
oeap-600 local-network ...................... Enable
oeap-600 Split Tunneling (Printers)......... Disable
WebPortal Online Client ..................... 0
mDNS snooping................................ Enabled
mDNS Query Interval.......................... 10 minutes
<snip>
```

Step 2—Editing the Default mDNS Profile

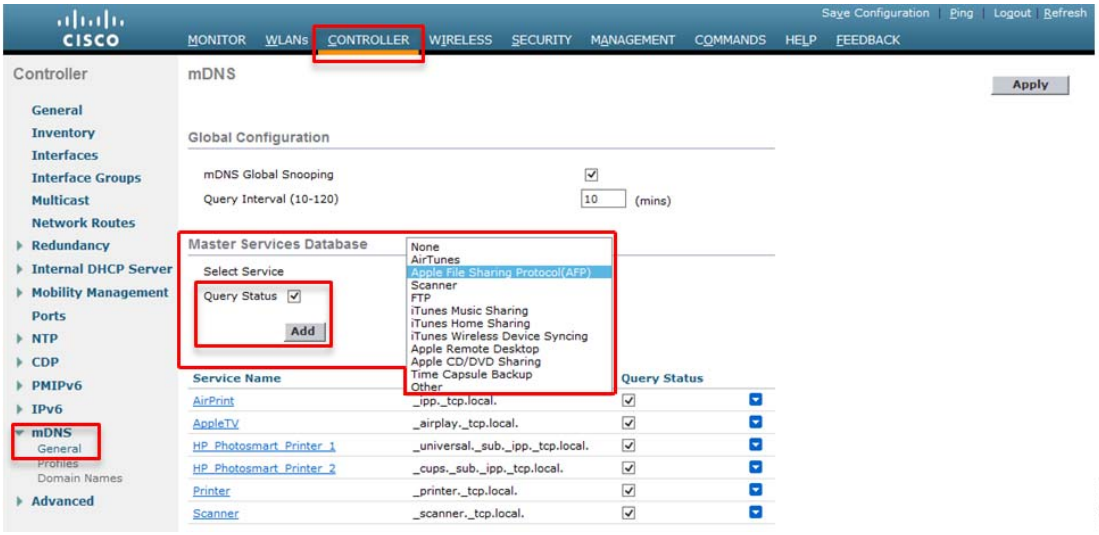
Additional Bonjour services may be added to the default mDNS profile (or even removed from it). To add additional Bonjour services, perform the following:

1. Select the Bonjour Service to be added from the **Master Services Database** drop-down list.

2. Enable the **Query Status Checkbox** for the service.
3. Click the **Add** button.
4. The added service will subsequently appear under the Service Name bar (in alphabetical order).

Figure 25-13 shows the Apple File Sharing Protocol (AFP) service being added to the default mDNS profile.

Figure 25-13 Use-Case 1—Step 2—Adding Bonjour Services to the Default mDNS Profile



Bonjour services can be added to the default (or non-default) profiles with the command shown in Example 25-4. The Profile Name of the default mDNS profile is "**default-mdns-profile**".

Example 25-4 Adding Bonjour Services to a mDNS Profile

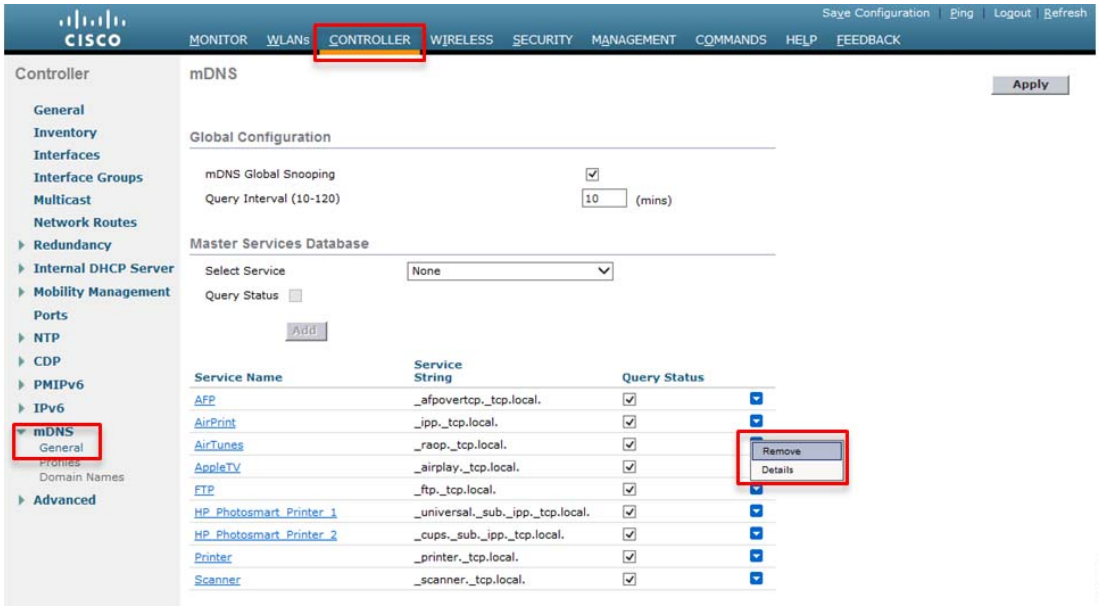
```
General Command:
(Cisco Controller) >config mdns profile service add mdns-profile-name mdns-service-name

Specific example:
(Cisco Controller) >config mdns profile service add default-mdns-profile AirPrint
! Adds the Apple AirPrint service to the default mDNS profile
```

Conversely, services can be removed from the default mDNS profile by clicking the blue-box at the end of the row for the service and then selecting **Remove**.

This is shown in Figure 25-14 where the AirPlay service (the service that allows for iTunes music to be streamed to a remote Apple Airport Express device, which in turn can supply an audio signal of the music to speakers) is removed from the default mDNS profile.

Figure 25-14 Use Case 1—Step 2b—Removing Bonjour Services from the Default mDNS Profile



Bonjour services can also be removed from the default (or non-default) profiles with the command shown in Example 25-5.

Example 25-5 Removing Bonjour Services from a mDNS Profile

```
General Command:
(Cisco Controller) >config mdns profile service delete mdns-profile-name mdns-service-name

Specific example:
(Cisco Controller) >config mdns profile service delete default-mdns-profile AirTunes
! Deletes the Apple AirTunes service from the default mDNS profile
```

The addition/removal of services to a mDNS profile can be verified by the **show mdns profile** command, which can either show a summary of configured profiles or a detailed view of a specific profile, as shown in Example 25-6 and Example 25-7, respectively.

Example 25-6 Verifying mDNS Profiles—show mdns profile summary

```
(Cisco Controller) >show mdns profile summary
Number of Profiles.............................. 1

ProfileName                       No. Of Services
--------------------------------  ---------------
default-mdns-profile                    6

(Cisco Controller) >
```

Example 25-7 Verifying mDNS Profiles—show mdns profile detailed Profile-Name

```
(Cisco Controller) >show mdns profile detailed default-mdns-profile

Profile Name.................................... default-mdns-profile
Profile Id...................................... 2
No of Services.................................. 6
```

```
Services..................................... AirPrint
                                              AppleTV
                                              HP_Photosmart_Printer_1
                                              HP_Photosmart_Printer_2
                                              Printer
                                              Scanner

No. Interfaces Attached...................... 1
Interfaces................................... dynamic

No. Interface Groups Attached................ 0
No. Wlans Attached........................... 4
Wlan Ids..................................... 1
                                              3
                                              4
                                              5

(Cisco Controller) >
```

Step 3—Apply the Default mDNS Profile an Interface (or Interface-Group)

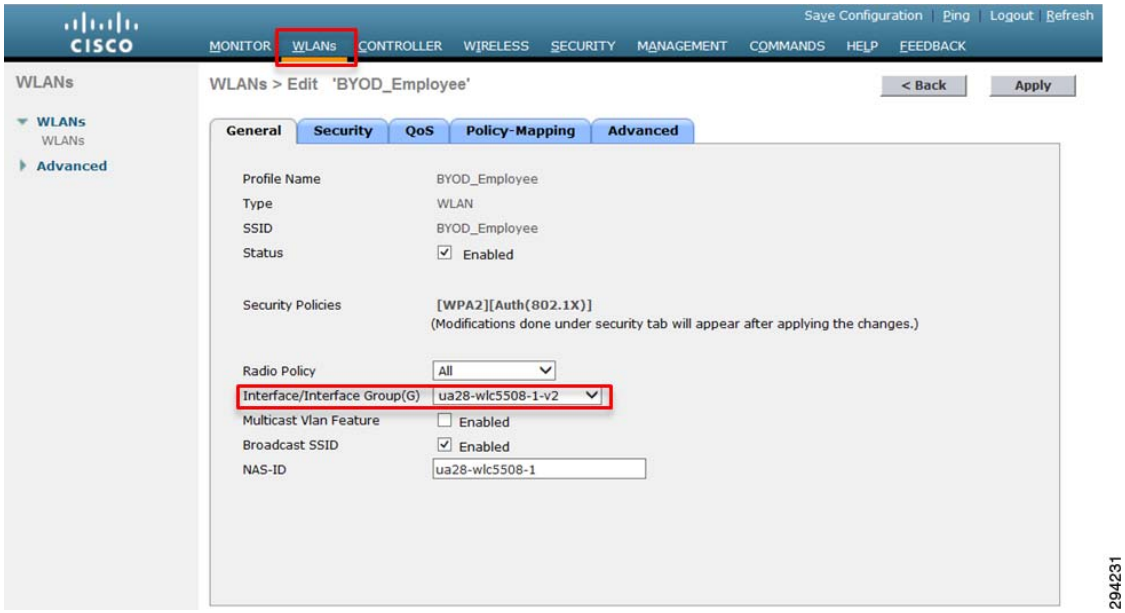
Bonjour service policies may be applied to interfaces, VLANs, or WLANs. In this example the Bonjour policies (as represented in the Default mDNS Profile) are attached to an interface.

There are five types of interfaces are available on the Cisco WLC controller. Four of these are static and are configured at setup time and the fifth type is dynamic and user-defined:

- Management interface (static and configured at setup time; mandatory)
- AP-manager interface (static and configured at setup time; mandatory)
- Virtual interface (static and configured at setup time; mandatory)
- Service-port interface (static and configured at setup time; optional)
- Dynamic interface (user-defined)

In this case, it is assumed that the ua28-wlc5508-1-v2 interface is applied to the BYOD_Employee WLAN (in line with the recommendations in Chapter 9, "BYOD Wireless Infrastructure Design"), as shown in Figure 25-15. If this is not the case, then the policies should be applied to whatever (static or dynamic) interface is associated with the WLAN. This association is verified by selecting the **WLANs** heading bar and then selecting the WLAN number that corresponds to the BYOD_Employee WLAN.

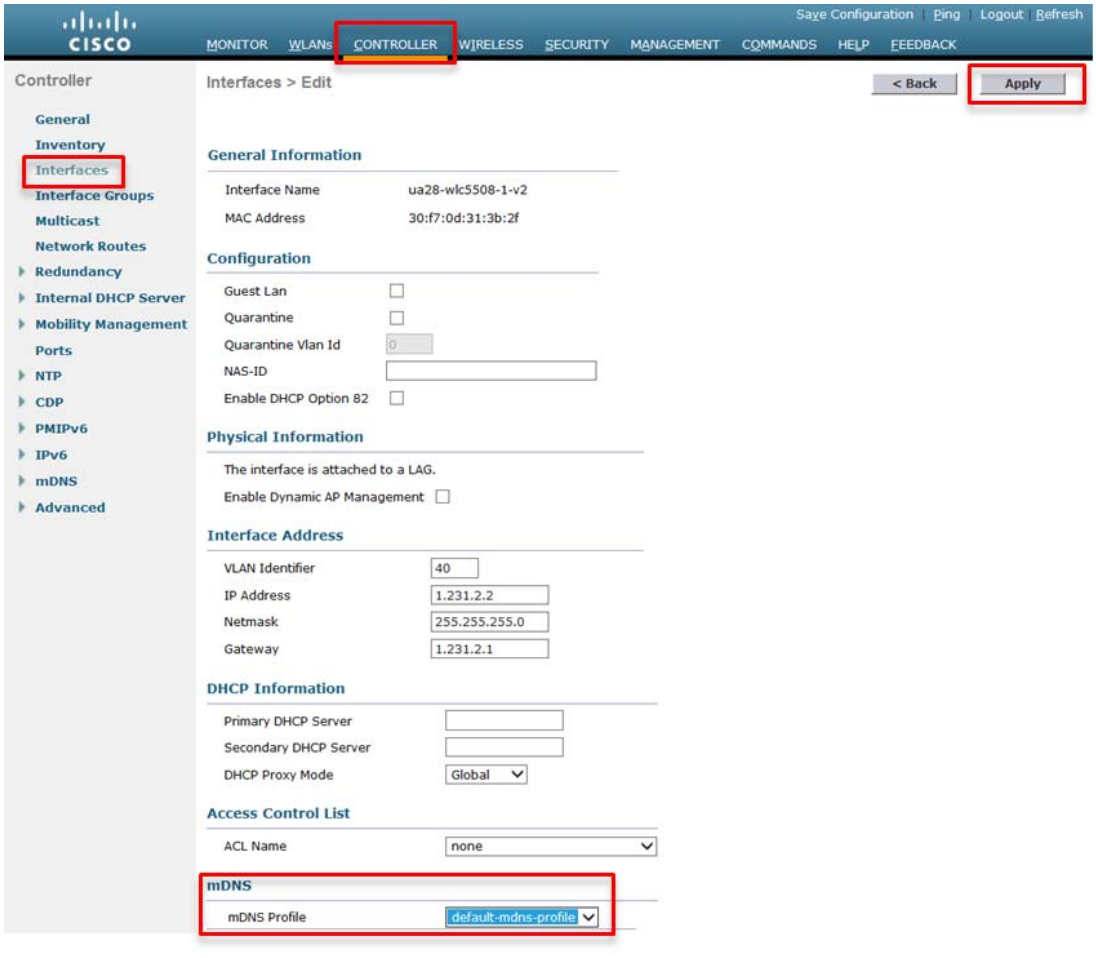
Figure 25-15 Use Case 1—Verifying WLAN/Interface Association



To apply the Default mDNS policies to an interface, perform the following:

1. Click the **CONTROLLER** heading-bar and then the **Interfaces** (or **Interface Group**) link on the left.
2. Select the interface that corresponds to the VLAN/WLAN to which the Bonjour service policies are to be applied.
3. At the bottom of the **Interface > Edit** page, select the **default-mdns-profile** from the **mDNS Profile** drop-down list.
4. Click the **Apply** button at the top-right of the page.

Figure 25-16 Use Case 1—Step 3—Applying the Default mDNS Profile to an Interface



As Figure 25-16 shows (in this case) the ua28-wlc5508-1-v2 interface corresponds to VLAN 40, which is where the wired AirPrint printer(s) reside. Bonjour service advertisements from these printers will now be shared with other WLANs/VLANs.

The Default mDNS profile can be added to the interface associated with the WLAN with the commands shown in Example 25-8.

Example 25-8 Adding a mDNS Profile to an Interface

```
General command:
(Cisco Controller) >config interface mdns-profile {interface-name | all} mdns-profile-name

Specific example:
(Cisco Controller) >config interface mdns-profile ua28-wlc5508-1-v2 default-mdns-profile
! Adds the default mDNS profile to the "ua28-wlc5508-1-v2" interface
```

The mDNS profile attached to an interface can be verified by the command **show interface detailed** *interface-name*, as shown in Example 25-9. Alternatively, if the mDNS profile is attached to an interface-group, then the show command would be **show interface group detailed** *interface-group-name*.

Example 25-9 Verifying Interface mDNS Profiles—show interface detailed interface-name

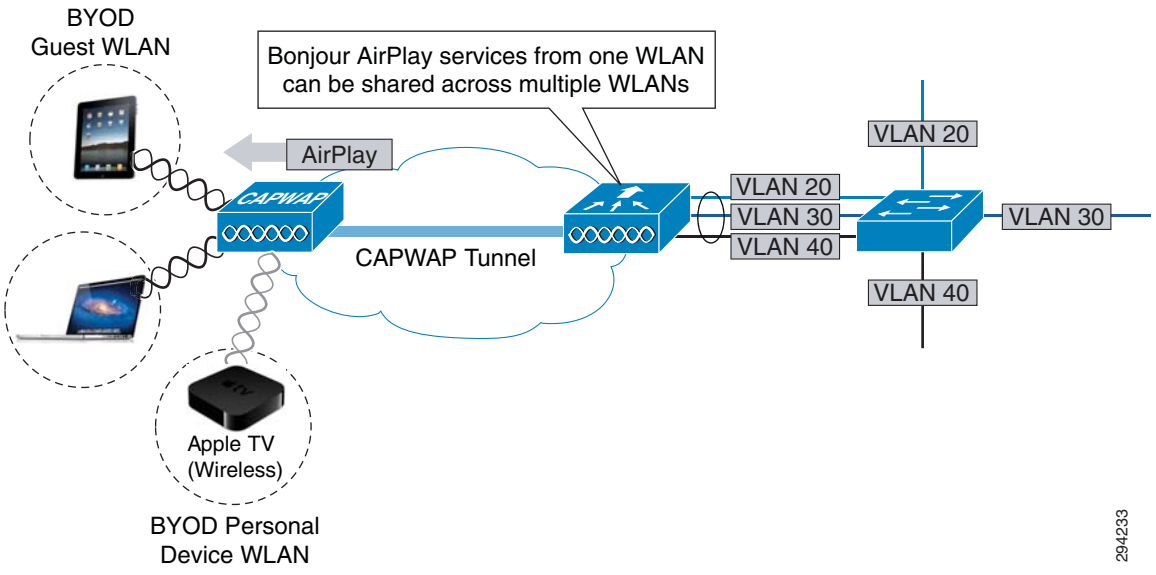
```
(Cisco Controller) >show interface detailed ua28-wlc5508-1-v2

Interface Name................................... ua28-wlc5508-1-v2
MAC Address...................................... 30:f7:0d:31:3b:2f
IP Address....................................... 10.225.43.2
IP Netmask....................................... 255.255.255.0
IP Gateway....................................... 10.225.43.1
External NAT IP State............................ Disabled
External NAT IP Address.......................... 0.0.0.0
VLAN............................................. 40
Quarantine-vlan.................................. 0
Active Physical Port............................. LAG (13)
Primary Physical Port............................ LAG (13)
Backup Physical Port............................. Unconfigured
DHCP Proxy Mode.................................. Global
Primary DHCP Server.............................. 10.230.1.61
Secondary DHCP Server............................ Unconfigured
DHCP Option 82................................... Disabled
IPv4 ACL......................................... Unconfigured
IPv6 ACL......................................... Unconfigured
mDNS Profile Name................................ default-mdns-profile
<snip>
```

Use Case 2—Wireless-to-Wireless Bonjour Gateway Service Policy—BYOD Guest AirPlay Example

In this secondary Bonjour Gateway use case, wireless guest devices are permitted to access Apple TV devices (using AirPlay) so that guests may share presentations, video, or other content with employees. Incidentally, Apple TVs, like some AirPrint printers, may be connected via wired or wireless connections; this design supports both options. However in this case, assume the Apple TV is residing in the BYOD Personal Devices WLAN, as shown in Figure 25-17.

Figure 25-17 Use Case 2—Cisco WLC Bonjour Gateway Wireless-to-Wireless Design Example



To highlight design and deployment options, in this example Bonjour service policies are configured by:

- Step 1—Creating a new mDNS profile.
- Step 2—Adding Bonjour services to the new mDNS profile.
- Step 3—Enabling mDNS snooping and the new mDNS profile directly on the WLAN.

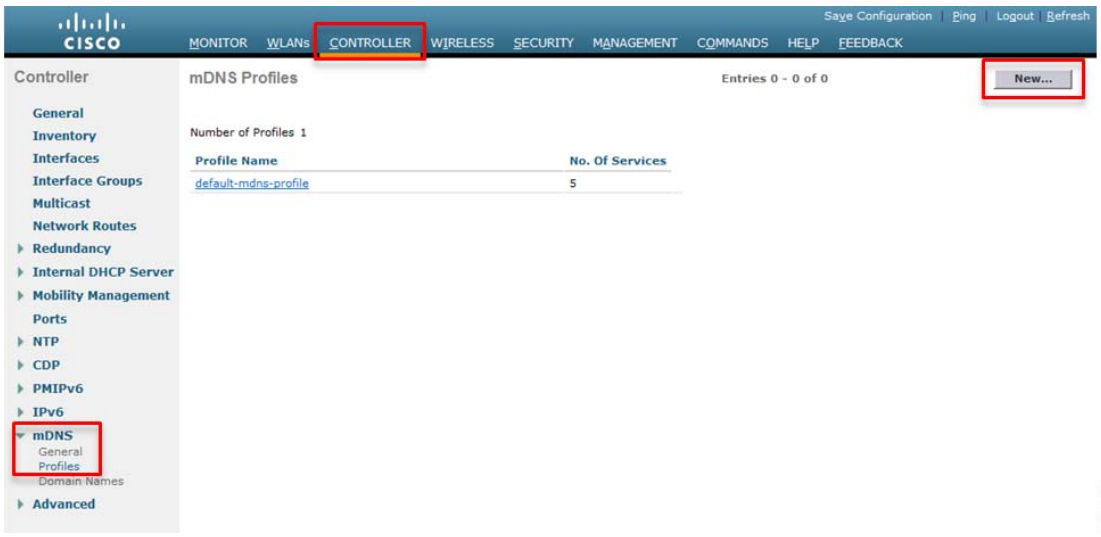
Each of these steps is detailed in turn.

Step 1—Creating a New mDNS Profile

The first step in this example is to create a new mDNS profile, which can be done by performing the following:

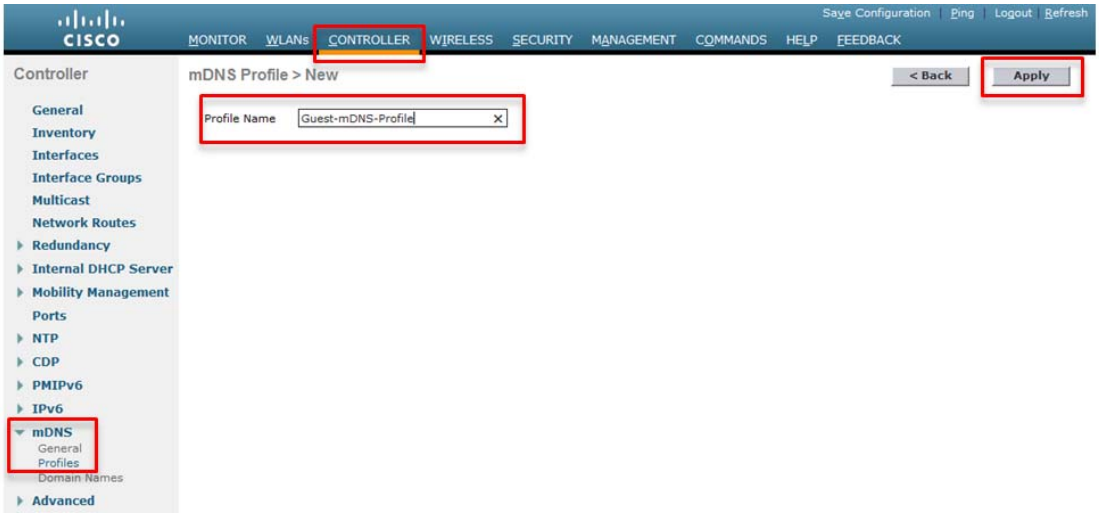
1. Click the **CONTROLLER** heading-bar and expand the mDNS link on the lower left and click **Profiles**.
2. Click the **New** button at the top-right, as shown in Figure 25-18.

Figure 25-18 Use Case 2—Step 1a—Creating a New mDNS Profile



3. Give the new profile a name and click the **Apply** button, as shown in Figure 25-19.

Figure 25-19 Use Case 2—Step 1b—Naming the New mDNS Profile



The corresponding Cisco WLC CLI for creating a new mDNS profile is shown in Example 25-10, which creates a new mDNS profile named "**Guest-mDNS-Profile**".

Example 25-10 Creating a New mDNS Profile

```
General command:
(Cisco Controller) >config mdns profile create mdns-profile-name

Specific example:
(Cisco Controller) >config mdns profile create Guest-mDNS-Profile
! Creates a new mDNS profile named "Guest-mDNS-Profile"
```

Newly created mDNS profiles will be displayed by the **show mdns profile summary verification** command, as shown in Example 25-11.

Example 25-11 Verifying mDNS Profiles—show mdns profile summary

```
(Cisco Controller) >show mdns profile summary
Number of Profiles.............................. 2

ProfileName                       No. Of Services
--------------------------------  ---------------
Guest-mDNS-Profile                      0
default-mdns-profile                    6

(Cisco Controller) >
```

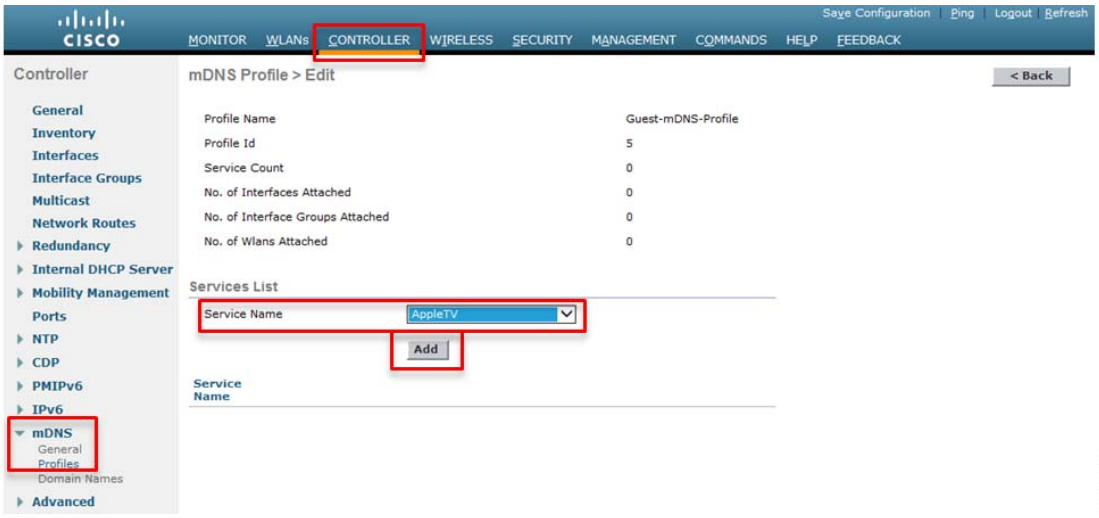
Step 2—Adding Bonjour Services to the New mDNS Profile

In this particular use case, only the AirPlay service will be offered to BYOD guest devices. Therefore the Bonjour AirPlay service needs to be added to the new mDNS Profile, which is done by performing the following:

1. Select and click the new mDNS profile.
2. Select the desired Bonjour service(s) from the **Services List** drop-down list and click the **Add** button, as shown in Figure 25-20.

3. The added service will subsequently appear under the **Service Name** bar.

Figure 25-20 Use Case 2—Step 2—Adding Bonjour Services to the New mDNS Profile



The corresponding Cisco WLC CLI for adding Bonjour services to a profile is shown in Example 25-12.

Example 25-12 Adding Bonjour Services to a mDNS Profile

```
General command:
(Cisco Controller) >config mdns profile service add mdns-profile-name mdns-service-name

Specific example:
(Cisco Controller) >config mdns profile service add Guest-mDNS-Profile AppleTV
! Adds the AppleTV service to the "Guest-mDNS-Profile" profile
```

Services within a mDNS profile can be verified by the **show mdns profile detailed** command, as presented in Example 25-13.

Example 25-13 Verifying mDNS Profiles—show mdns profile detailed Profile-Name

```
(Cisco Controller) >show mdns profile detailed Guest-mDNS-Profile

Profile Name..................................... Guest-mDNS-Profile
Profile Id....................................... 1
No of Services................................... 1
Services......................................... AppleTV

No. Interfaces Attached.......................... 0
No. Interface Groups Attached.................... 0
No. Wlans & Guest-LANs Attached.................. 0

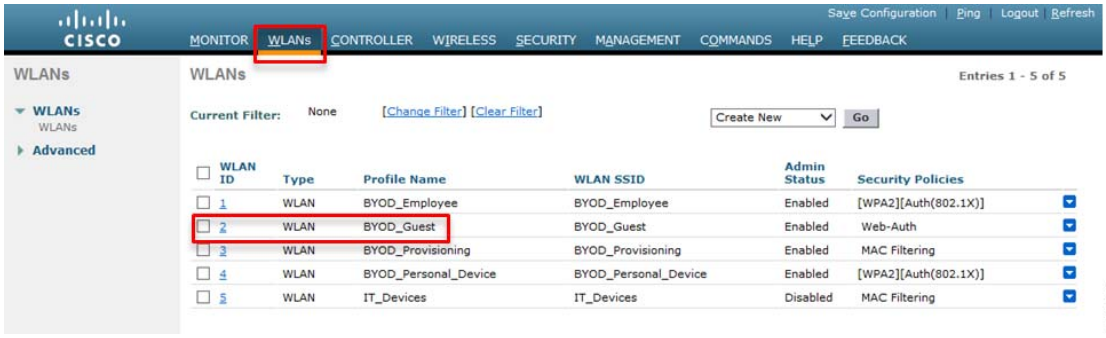
(Cisco Controller) >
```

Step 3—Enable mDNS Snooping and the New mDNS Profile on the WLAN

Once all the Bonjour services have been added to the new profile, it can be added to the desired WLAN (in this case, the BYOD_Guest WLAN) by performing the following:

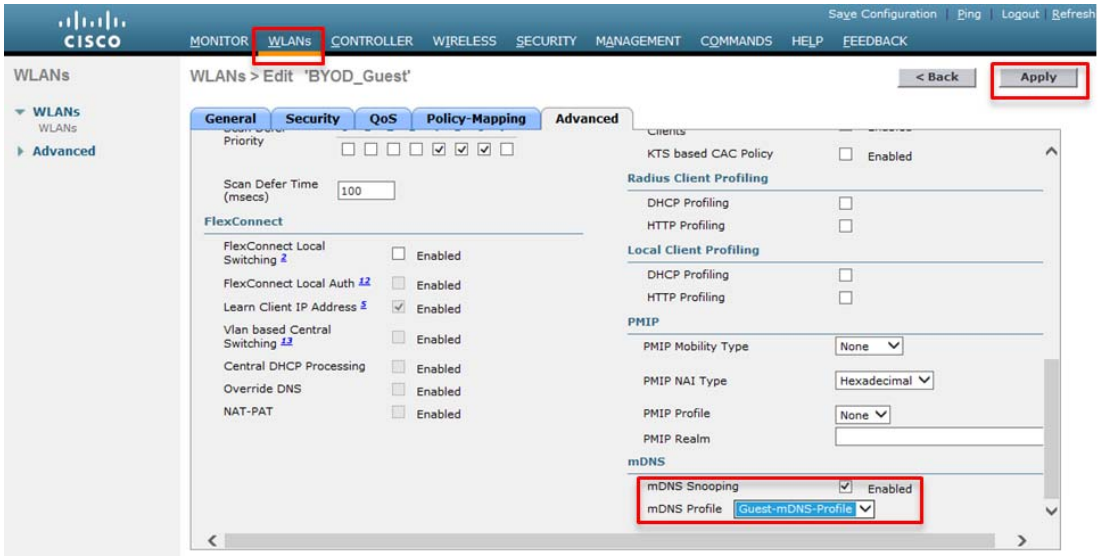
1. Click the **WLANs** heading-bar and select the desired WLAN (in this case, the BYOD_Guest WLAN, as shown in Figure 25-21).

Figure 25-21 Use Case 2—Step 3a—Selecting the WLAN to which the New mDNS Profile Will Be Applied



2. Click the **Advanced** tab and scroll to the bottom.
3. Ensure that the **mDNS Snooping** checkbox is selected.
4. Select the **mDNS Profile** from the drop-down list.
5. Click the **Apply** button at the top-left.

Figure 25-22 Use Case 2—Step 3b—Enabling mDNS Snooping on the WLAN and Applying the New mDNS Profile



The corresponding Cisco WLC CLI for these steps of enabling mDNS snooping and a specific mDNS profile on a WLAN is shown in Example 25-14 and Example 25-15, respectively.

Example 25-14 Enabling mDNS Snooping on a WLAN

```
General command/specific example:
(Cisco Controller) > config wlan mdns enable
```

Example 25-15 Adding a mDNS Profile to a WLAN

```
General command:
(Cisco Controller) >config wlan mdns profile {wlan-id | all } mdns-profile-name

Specific example:
(Cisco Controller) >config wlan mdns profile 2 Guest-mDNS-Profile
! Adds the "Guest-mDNS-Profile" to WLAN 2 (the BYOD_Guest WLAN, as shown in Figure 21)
```

The mDNS settings of a WLAN can be verified by the **show wlan** *wlan-id* verification command, as shown in Example 25-16.

Example 25-16 Verifying WLAN mDNS Settings—show wlan

```
(Cisco Controller) >show wlan 2

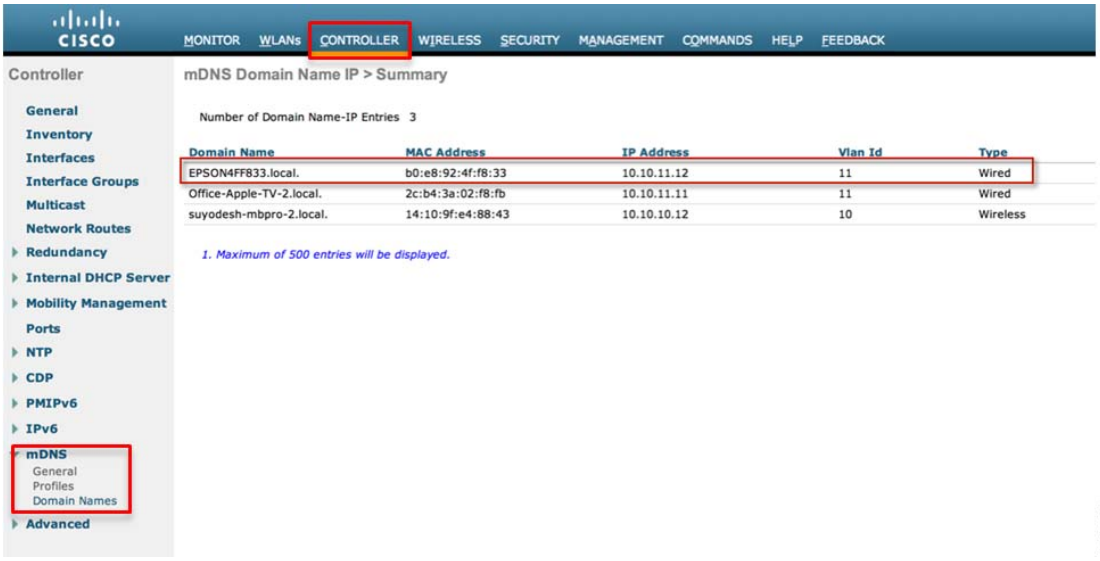

WLAN Identifier.................................. 2
Profile Name..................................... BYOD_Guest
Network Name (SSID).............................. BYOD_Guest
Status........................................... Enabled
MAC Filtering.................................... Disabled
Broadcast SSID................................... Enabled
AAA Policy Override.............................. Enabled
Network Admission Control
Client Profiling Status
    Radius Profiling ............................ Disabled
     DHCP ....................................... Disabled
     HTTP ....................................... Disabled
    Local Profiling ............................. Disabled
     DHCP ....................................... Disabled
     HTTP ....................................... Disabled
  Radius-NAC State............................... Disabled
  SNMP-NAC State................................. Disabled
  Quarantine VLAN................................ 0
Maximum number of Associated Clients.............. 0
Maximum number of Clients per AP Radio........... 200
Number of Active Clients......................... 0
Exclusionlist Timeout............................ 60 seconds
Session Timeout.................................. 1800 seconds
User Idle Timeout................................ Disabled
Sleep Client..................................... disable
Sleep Client Timeout............................. 12 hours
User Idle Threshold.............................. 0 Bytes
NAS-identifier................................... ua28-wlc5508-1
CHD per WLAN..................................... Enabled
Webauth DHCP exclusion........................... Disabled
Interface........................................ ua27-5508-2-guest
Multicast Interface.............................. Not Configured
WLAN IPv4 ACL.................................... unconfigured
WLAN IPv6 ACL.................................... unconfigured
WLAN Layer2 ACL.................................. unconfigured
mDNS Status...................................... Enabled
mDNS Profile Name................................ Guest-mDNS-Profile
<snip>
```

Verifying Bonjour Gateway Operation

In addition to the GUI and CLI configuration-verification screenshots and commands that have been highlighted in the previous sections, Cisco WLC software has some additional options for verifying Bonjour Gateway operation, which we now discuss.

For instance, a summary of all mDNS records can be shown by clicking the **CONTROLLER** heading-bar, expanding the **mDNS** link on the lower left, and then clicking **Domain Names**, as shown in Figure 25-23.

Figure 25-23 Verifying mDNS Domain Names



A summary of mDNS records can also be provided via the CLI with the command **show mdns domain-name-ip summary**, as shown in Example 25-17.

Example 25-17 Verifying mDNS Records—show mdns domain-name-ip summary

```
(Cisco Controller) >show mdns domain-name-ip summary

Number of Domain Name-IP Entries.................. 3

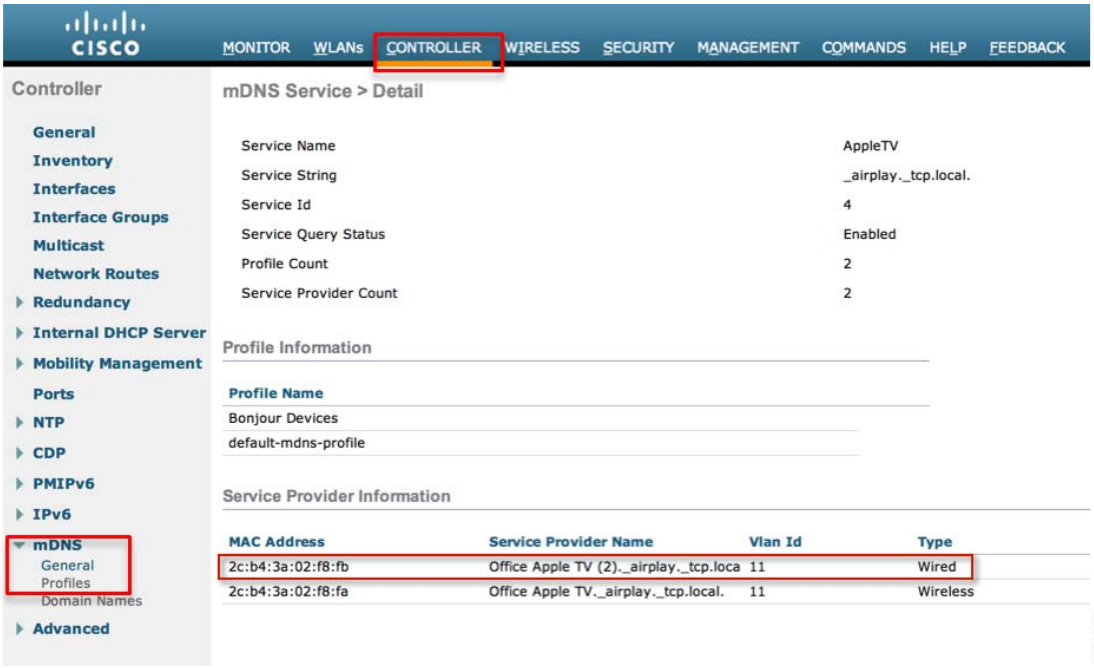
DomainName              MAC Address        IP Address    Vlan Id  Type      TTL  Time left
                                                                             (sec) (sec)
--------------------    ----------------   -----------   -------  ------    -----  -----

EPSON4FF833.local.      b0:e8:92:4f:f8:33  10.10.10.12   11       Wired     4725   4354
Office-Apple-TV.local.  2c:b4:3a:02:f8:fb  10.10.11.11   11       Wired     4725   4712
suyodesh-mbpro-2.local. 14:10:9f:e4:88:43  10.10.10.12   10       Wireless  4725   3753

(Cisco Controller) >
```

Also, clicking on any service listed within an mDNS profile will display a **mDNS Service>Detail** screen that will display device-level details—including MAC address, VLAN, and network-type (wired or wireless) for any and all devices providing that Bonjour service. For example, Figure 25-24 shows that the Apple TV service is available both via the wired and wireless networks.

Figure 25-24 Verifying mDNS Service Details



Device-level mDNS service detail is also available via the CLI using the command **show mdns service detailed** *mdns-service-name*, as demonstrated in Example 25-18.

Example 25-18 Verifying mDNS Service Details—show mdns service detailed

```
(Cisco Controller) >show mdns service detailed AppleTV

Service Name..................................... AppleTV
Service Id....................................... 4
Service query status............................. Enabled
Service LSS status............................... Disabled
Service learn origin............................. Wireless and Wired
Number of Profiles............................... 2
Profile.......................................... Guest-mDNS-Profile
                                                  default-mdns-profile

Number of Service Providers ..................... 1
Number of priority MAC addresses ................ 0
ServiceProvider                                   MAC Address       AP Radio MAC
Vlan Id    Type       TTL     Time left

(sec)      (sec)
--------------------                              ----------------  ----------------
-------    ------     -----   ---------

Office Apple TV._airplay._tcp.local.              2c:b4:3a:02:f8:fa  04:da:d2:b2:47:10
11     Wireless     4500      4460
```

Additionally, the CLI allows for a summary of mDNS services to be displayed via the **show mdns service summary** command, as shown in Example 25-19.

Example 25-19 Verifying mDNS Service Summary—show mdns service summary

```
(Cisco Controller) >show mdns service summary
```

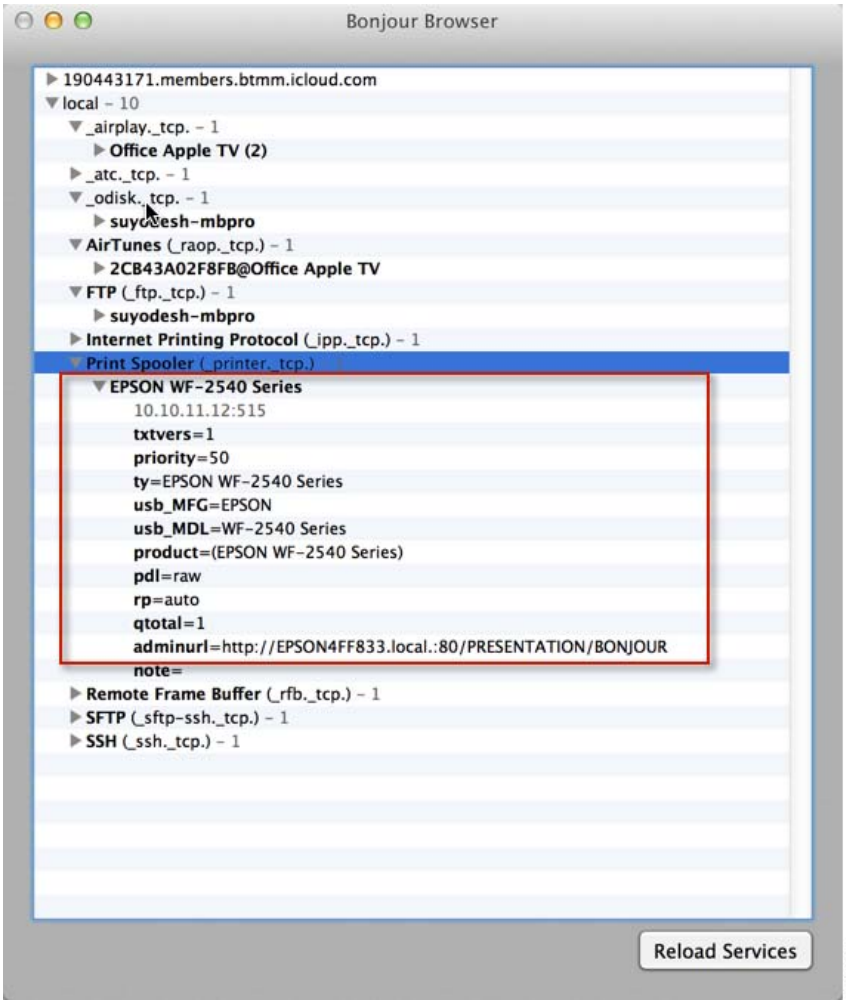
```
Number of Services.............................. 11

Service-Name                      LSS   Origin     No SP    Service-string
--------------------------------- ----  ----------  -----   ---------------
AFP                               No    All           0   _afpovertcp._tcp.local.
AirPrint                          No    All           1   _ipp._tcp.local.
AirTunes                          No    All           1   _raop._tcp.local.
AppleTV                           No    All           1   _airplay._tcp.local.
FTP                               No    All           1   _ftp._tcp.local.
HP_Photosmart_Printer_1           No    All           1   _universal._sub._ipp._tcp.local.
HP_Photosmart_Printer_2           No    All           0   _cups._sub._ipp._tcp.local.
Printer                           No    All           1   _printer._tcp.local.
Scanner                           No    All           1   _scanner._tcp.local.
TimeCapsuleBackup                 No    All           0   _adisk._tcp.local.
iTuneHomeSharing                  No    All           0   _home-sharing._tcp.local.

(Cisco Controller) >
```

Finally, it bears mentioning that third-party tools are also available to verify mDNS operations. For example, Figure 25-25 shows Tildesoft's "Bonjour Browser" displaying mDNS details for the Epson wired AirPrint printer.

Figure 25-25 Verifying Bonjour Gateway Operation via Third-Party Tools—Tildesoft Bonjour Browser Example



Advanced Bonjour Gateway Scenario Operation

This section will briefly overview Bonjour Gateway operation in three additional scenarios:

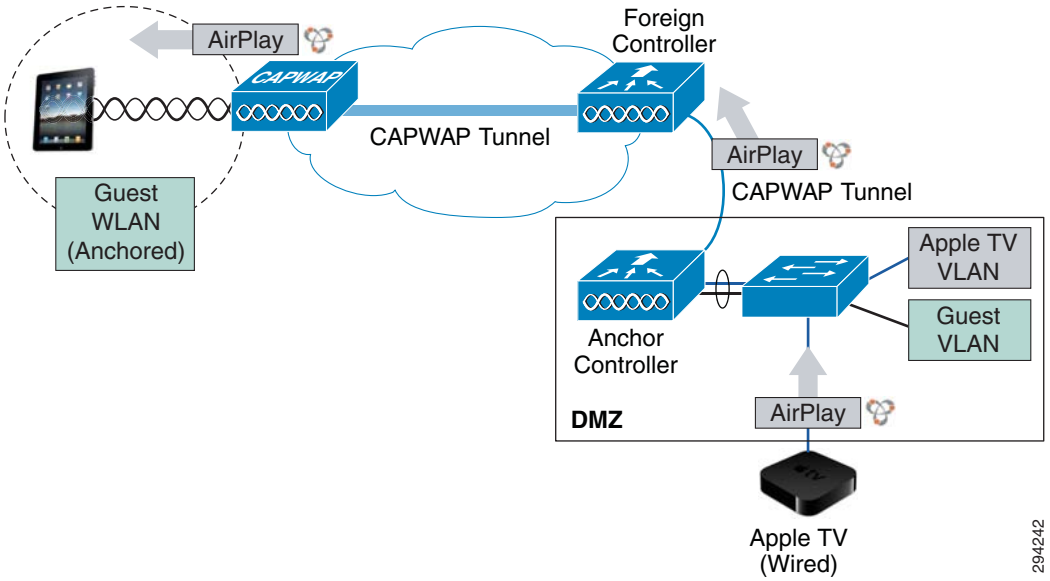
- Guest Anchoring
- Layer 3 Roaming
- FlexConnect

While the configuration and verification of the Bonjour Gateway feature remains the same for these scenarios, it may be helpful for network administrators to understand how this feature operates in these contexts.

Guest Anchoring

In guest anchoring scenarios, the guest WLAN is able to see Bonjour services advertised to the anchor controller. This is because the Bonjour queries and advertisements are sent inside the Control and Provisioning of Wireless Access Points (CAPWAP) tunnel, as shown in Figure 25-26.

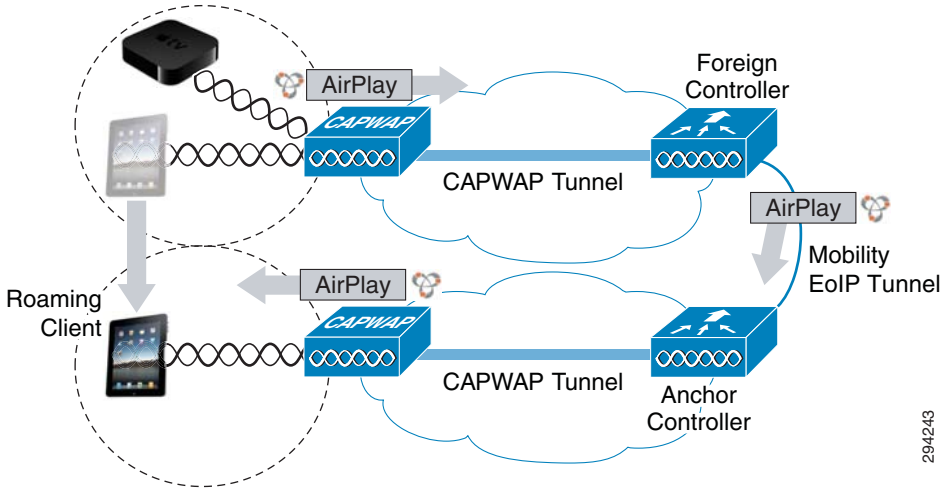
Figure 25-26 Bonjour Gateway Operation in Guest Anchoring Scenarios



Layer 3 Roaming

Bonjour Gateway with Layer 3 roaming works across Ethernet over IP (EoIP) tunnels to ensure that users moving among access points (APs) on different controllers continue to see the devices they saw on the original controller. The Bonjour services on the anchor controller are displayed to the client, including both wired and wireless devices, as shown in Figure 25-27.

Figure 25-27 Bonjour Gateway Operation in Layer 3 Roaming Scenarios



FlexConnect

For centrally-switched WLANs, the behavior for Bonjour is the same as if the AP was in local mode. In this case, Bonjour queries from the client are sent to the controller and Bonjour responses from the controller are sent back to the AP in the unicast CAPWAP tunnel. This means FlexConnect APs will not require "Multicast-Unicast" mode to support Bonjour.

For locally switched WLANs, the behavior for Bonjour will continue to work for a single subnet only.

Summary

This paper overviewed Apple's Bonjour protocol—a zero-configuration protocol for advertising, discovering, and connecting to network services—and how it can be effectively managed within a BYOD enterprise context.

The design limitation of Bonjour's use of link-local multicasting was discussed, showing how it limited the usefulness of the protocol to only a single Layer 2 domain. To enable the use of Bonjour in (multi-WLAN/VLAN) BYOD enterprise networks, the Cisco WLC Bonjour Gateway was introduced. Next, an overview of the operation of the Bonjour Gateway feature was provided, showing how it can be used to snoop, cache, and proxy-respond to Bonjour service requests. Additionally, it was shown how these responses could be selectively enabled and disabled, allowing for administrative policy-based control of Bonjour services.

Following this, deployment details of this feature were presented by considering two main use-case scenarios:

- Printing from wireless devices to wired printers.
- Sharing Bonjour services between wireless devices in different WLANs.

Step-by-step configuration guidance was presented for each scenario, using slightly different approaches to highlight the various configuration options available. Each step was presented for not only the Cisco WLC GUI configuration and verification, but also for the Cisco WLC CLI.

Additional verification options were also highlighted, as well as how the Bonjour Gateway operates in various advanced scenarios, including guest anchoring, Layer 3 roaming, and FlexConnect deployments.

References

- Cisco Wireless LAN Controller Configuration Guide, Release 7.4
  http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED.html
- Cisco Wireless LAN Controller Configuration Guide, Release 7.4—Configuring Multicast Domain Name System
  http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01011.html#d75540e531a1635
- Cisco WLC Bonjour Gateway Deployment Guide
  http://www.cisco.com/en/US/docs/wireless/technology/bonjour/Bonjour_Deployment.html

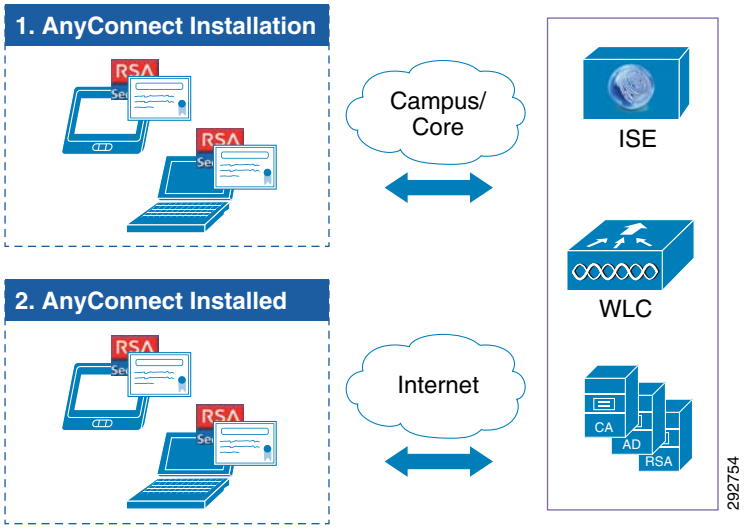


BYOD Remote Device Access

Revised: August 7, 2013

A BYOD design should be able to accommodate devices that attempt to connect remotely to access the internal resources. The device could be a workstation, tablet, smartphone, or any other device which is allowed to connect securely to the network. In this design, the Cisco ASA is used as a VPN gateway for establishing an SSL VPN session to the remote endpoints. The ASA authenticates the user's digital certificate. Cisco ISE then authenticates the user via an RSA SecurID token. The combination of both allows the device onto the network. Figure 26-1 shows the network components involved in remote device access.

Figure 26-1 Remote Device Network Components



This design assumes the following for providing remote access capability:

- Devices that want to connect to the network remotely must be corporate approved devices. The corporate approved devices are the ones that have been provisioned with a digital certificate by the IT organization. To understand more about corporate approved devices, see Chapter 16, "BYOD Limited Use Case—Corporate Devices."
- Devices must be provisioned at the campus. The provisioning process consists of:
  - Installing the AnyConnect Client
  - Configuring the VPN gateway IP address

  - Setting up one-time-password scheme for the user.

  These steps must be completed at the campus location before it can be used remotely. This design does not allow remote provisioning of the devices.

- Devices connecting remotely are subjected to two factor authentication, which means the user should provide two forms of credentials.

Solution Components

The following components play a role in providing connectivity to remote clients:

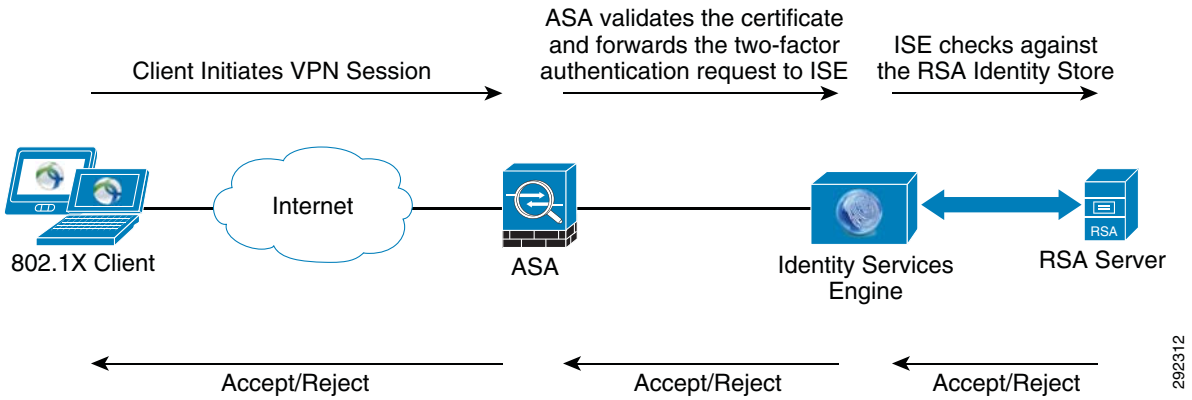
- Cisco Adaptive Security Appliance (ASA)—Functions as an SSL VPN concentrator for terminating VPN sessions.
- Cisco AnyConnect—Acts as a VPN client installed on the remote device.
- Cisco Identity Services Engine (ISE)—Acts as an intermediary to an external identity source for authentication between remote endpoints and the RSA Server. The token gets forwarded from the client to the ASA, from the ASA to the ISE, and then from the ISE to the RSA.
- RSA SecurID—Acts as the authentication server for tokens generated by the client.

RSA SecurID

VPN security is only as strong as the methods used to authenticate users (and device endpoints) at the remote end of the VPN connection. Simple authentication methods based on static passwords are subject to password "cracking" attacks, eavesdropping, or even social engineering attacks. Two-factor authentication, which consists of "something you know" and "something you have", is a minimum requirement for providing secure remote access to the corporate network. For more details, see: http://www.cisco.com/web/about/security/intelligence/05_08_SSL-VPN-Security.html.

This design includes the RSA SecurID Authentication Server 7.1, along with RSA SecurID hardware tokens, to provide two-factor authentication. The passcode that the user presents is a combination of their secret PIN and the one time password (OTP) code that is displayed on their token at that moment in time. This design utilizes both RSA SecurID (two-factor authentication) in conjunction with the deployment and use of x.509 client digital certificates. Figure 26-2 shows how RSA is used for two-factor authentication.

Figure 26-2 RSA Used for Two-Factor Authentication

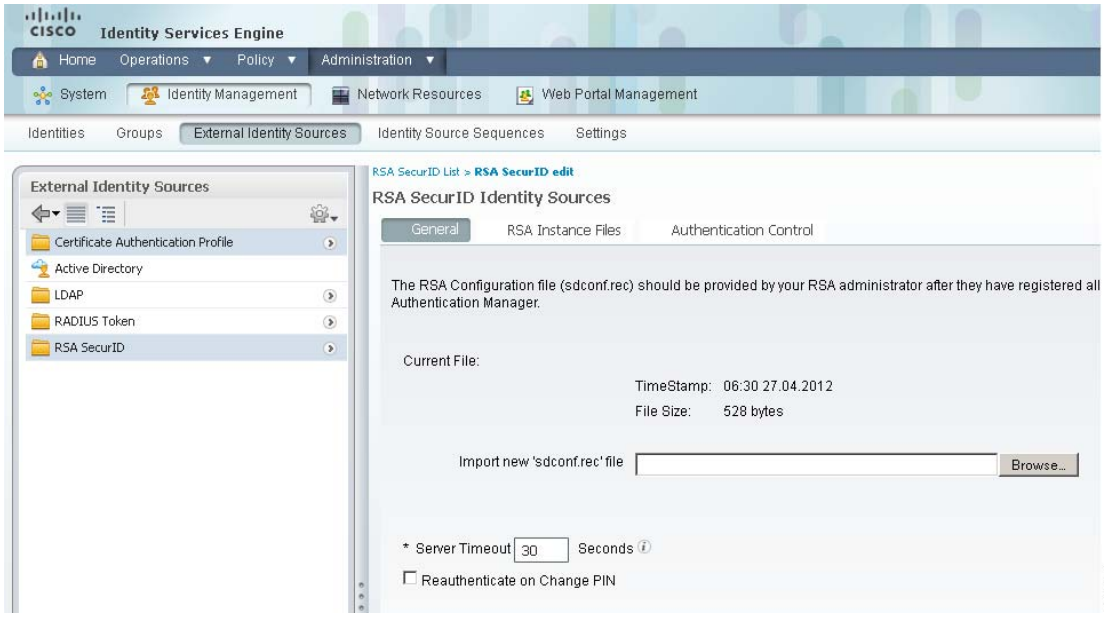


For information on configuring the RSA Secure Authentication Manager, see: http://www.emc.com/security/rsa-securid.htm.

ISE Integration with RSA

The RSA Identity Store is used primarily to authenticate remote users. The remote users are authenticated initially by their digital certificates and then they must provide a one-time password using a RSA SecurID token. To configure the RSA as an identity store, click **Administration > Identify Management > External Identity Sources > RSA SecurID > Add**, as shown in Figure 26-3.

Figure 26-3 RSA Server as an Identity Store for ISE



VPN Design Considerations

This section discusses the primary role of the ASA for this design, which is to terminate SSL VPN connections. The following are some of the many design considerations when implementing the SSL VPN:

- How do remote users trust the VPN gateway?
- How does the VPN gateway identify remote users?
- How to organize different types of users in groups so that different kinds of services can be provided?
- What kind of mobility client solution is needed for a particular client?
- Once the right kind of VPN solution is identified, how will the mobility client be installed on the remote device?
- How to centralize the policy settings for VPN users? It is not always easy or convenient for remote users to configure a mobile device for VPN functionality.

The Cisco ASA coupled with the Cisco AnyConnect client addresses the considerations mentioned above. The Cisco AnyConnect client 3.0 is used to meet the needs of wired, wireless, and remote users. The Cisco AnyConnect Secure Mobility client is the next-generation VPN client, providing remote users with secure IPsec (IKEv2) or SSL VPN connections to the Cisco 5500 Series Adaptive Security Appliance (ASA). AnyConnect provides end users with a connectivity experience that is intelligent, seamless, and always-on, with secure mobility across today's proliferating managed and unmanaged mobile devices.

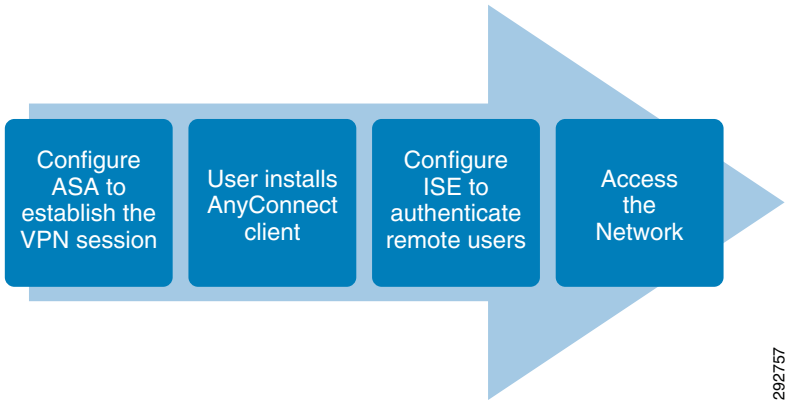
The Cisco AnyConnect Secure Mobility client integrates new modules into the AnyConnect client package:

- Network Access Manager (NAM)—Formerly called the Cisco Secure Services Client, this module provides Layer 2 device management and authentication for access to both wired and wireless networks.
- Posture Assessment—The AnyConnect Posture Module provides the AnyConnect Secure Mobility Client with the ability to identify the operating system, antivirus, antispyware, and firewall software installed on the host prior to creating a remote access connection to the ASA. Based on this pre-login evaluation, you can control which hosts are allowed to create a remote access connection to the security appliance. The Host Scan application is delivered with the posture module and is the application that gathers this information.
- Telemetry—Sends information about the origin of malicious content detected by the antivirus software to the web filtering infrastructure of the Cisco IronPort Web Security Appliance (WSA), which uses this data to provide better URL filtering rules.
- Web Security—Routes HTTP and HTTPS traffic to the ScanSafe Web Security scanning proxy server for content analysis, detection of malware, and administration of acceptable use policies.
- Diagnostic and Reporting Tool (DART)—Captures a snapshot of system logs and other diagnostic information and creates a .zip file on your desktop so you can conveniently send troubleshooting information to Cisco TAC.
- Start Before Logon (SBL)—Forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting AnyConnect before the Windows login dialog box appears.

For more information on the Cisco AnyConnect 3.0 client, see: http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect30/administration/guide/ac01intro.html.

Figure 26-4 shows the steps to provide VPN connectivity.

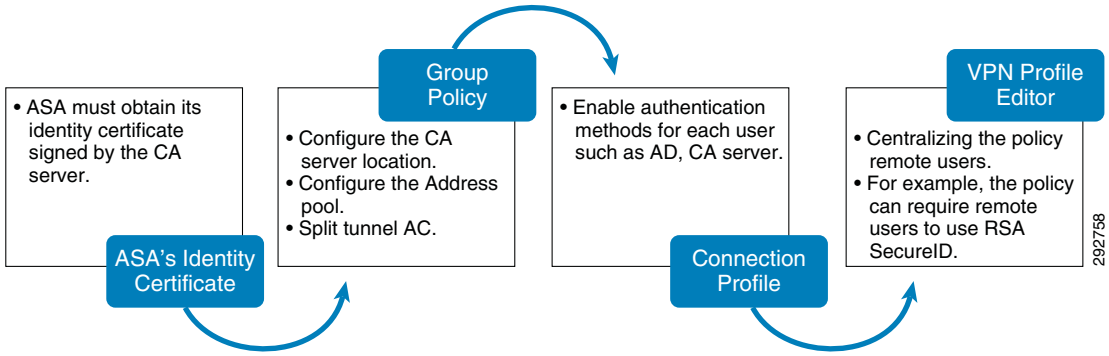
Figure 26-4 High Level Steps for Providing VPN Connectivity



ASA Configuration

The configuration of the ASA involves many steps. Figure 26-5 shows, at a high level, the steps required to configure the ASA.

Figure 26-5 Configuration of ASA

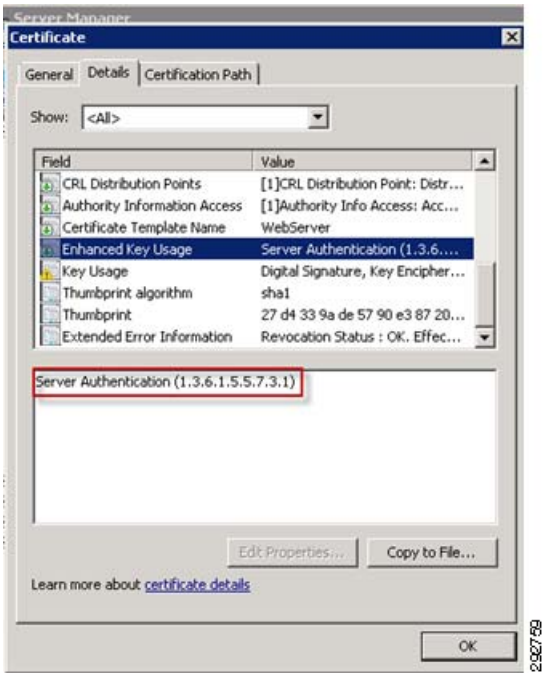


ASA's Identity Certificate

The ASA needs to present a digital certificate as means of authenticating itself to the clients. The remote clients validate the digital certificate and if the validation is successful, then they proceed to the next steps of establishing a VPN connection.

The digital certificate provided by the ASA must be issued by a trusted third-party like VeriSign or it could be also issued by an internal CA, which is signed by a trusted third-party. Instead, if the ASA presents a self-signed certificate, then the clients cannot validate the certificate because the signing authority (ASA for self-signed) is not in the list of trusted CAs in the client browser. Hence for greater security, it is recommended that the ASA's digital certificate is either issued by a trusted third-party or by an internal CA which is signed by a trusted third-party. When using a Microsoft CA as internal CA, it is important to verify that the certificate properties support Server Authentication. Figure 26-6 shows the certificate that can be used for server authentication. The certificate should contain EKU of Server Authentication, as indicated in Figure 26-6.

Figure 26-6 Certificate for Server Authentication



The ASA can obtain the certificate from the CA server by using SCEP or by a manual cut-and-paste method. To obtain more information on deploying certificates on the ASA, see: http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert_cfg.html.

The following example shows the ASA configuration for certificate enrollment:

```
crypto ca trustpoint WIN2K-CA
 enrollment terminal
 subject-name CN=ASA-remote1
 serial-number
 ip-address 172.26.185.195
 keypair ssl
 no client-types
 crl configure
```

The above trust is used by ASA to obtain its own Identity Certificate from the CA server. In this design the enrollment method is terminal.

The following command shows the digital certificate issued by the CA server to the ASA:

```
ASA-remote1(config)# show crypto ca certificates
Certificate
  Status: Available
  Certificate Serial Number: 1594b5d9000000000213
  Certificate Usage: General Purpose
  Public Key Type: RSA (2048 bits)
  Signature Algorithm: SHA1 with RSA Encryption
  Issuer Name:
    cn=secbn1-WIN-MRL23B7NQO6-CA
    dc=secbn1
    dc=com
  Subject Name:
    cn=ASA-remote1
    hostname=ASA-remote1.secbn1.com
    ipaddress=172.26.185.195
    serialNumber=JMX1215L1KF
```

```
  CRL Distribution Points:
    [1]  ldap:///CN=secbn1-WIN-MRL23B7NQO6-CA,CN=WIN-MRL23B7NQO6,CN=CDP,CN=Publi
c%20Key%20Services,CN=Services,CN=Configuration,DC=secbn1,DC=com?certificateRevo
cationList?base?objectClass=cRLDistributionPoint
    [2]  http://win-mrl23b7nqo6.secbn1.com/CertEnroll/secbn1-WIN-MRL23B7NQO6-CA.
crl
  Validity Date:
    start date: 09:29:35 EST May 30 2012
    end   date: 09:29:35 EST May 30 2014
  Associated Trustpoints: WIN2K-CA
```



Note The client must have network connectivity to the CRL distribution point as provided in the certificate.

ASA Trust Point to Authenticate Remote Users

ASA also needs a trust point to authenticate remote users' identity certificates. The following is the configuration of the trust point:

```
crypto ca trustpoint Validate
 enrollment terminal
 crl configure
```

The above trust point "Validate" is used to copy the root CA certificate. To understand more about how to cut-and-paste certificates using terminal method, see:
http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert_cfg.html.

Creating Groups for Different Types of Users

Group policy is an important building block for designing an effective access mechanism for users. The needs of specific users can differ. For example, one user might like to have a domain value of xyz.com and have 1.1.1.1 and 2.2.2.2 as their DNS servers. Another user might have similar requirements, but in addition might need a proxy server configured for their user name. If you have to attach all these attributes to each individual user, the configuration might become very large and complex. To solve this problem, multiple groups can be created, each with its set of individual attributes. In this case you can simply associate a user with a group name, rather than the large number of attributes, thus minimizing the configuration complexity when you have multiple users.

By default, the Cisco ASA creates DftGrpPolicy and the other group polices that inherit most of the common attributes. Only very specific attributes need to be configured explicitly for each group.

For more information about configuring tunnel groups, group policies, and users, see:
http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/vpngrp.html.

In the group policy definition for this design guide, the main attributes needed are vpn-tunnel-protocol, split-tunnel-network-list, and address pool location. The following example shows how this group policy was defined:

```
group-policy SSLClientPolicy internal     !This group policy is defined internally not
downloaded from radius.
group-policy SSLClientPolicy attributes
 wins-server value 10.1.6.100          ! WINS server IP address
 dns-server value 10.1.6.100          ! DNS server IP address
 vpn-tunnel-protocol ssl-client ssl-clientless
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value split_ACL  ! split_ACL prevents some local network
traffic from getting into VPN traffic.
 default-domain value secbn1.com
```

```
 address-pools value testpool        ! The IP address pool value.
```

Connection Profile Configuration

While group policies define the attributes for a group, the connection profile specifies the attributes specific to a connection. For example, a connection profile for AnyConnect specifies if the users belonging to this connection are authenticated by a RADIUS server or locally. The connection profile also points to the group profile to which it belongs. If no connection profile is defined on the system, the ASA points to a default connection profile, but to make administration simple it is better to define a specific group and connection profiles. The following example shows the ASA configuration for the AnyConnect connection profile:

```
tunnel-group SSLClientProfile type remote-access
tunnel-group SSLClientProfile general-attributes
 authentication-server-group ISE    ! The remote sessions are authenticated with ISE.
 default-group-policy SSLClientPolicy ! The parent group policy used by this connection
profile.

tunnel-group SSLClientProfile webvpn-attributes
 authentication aaa certificate        ! The remote users are authenticated by AAA and
Digital Certificate.
 group-alias SSLVPNClient enable  ! The remote users are presented with this alias name
during the session.
 group-url https://172.26.185.195/SSLVPNClient enable
 group-url https://192.168.167.225/SSLVPNClient disable
!
```

The above configuration steps illustrate how to configure SSLVPN sessions with AnyConnect. The same configuration can also be done using ASDM editor or by another management tool. To obtain more information about the configuration using other tools, refer to ASA configuration editor at: http://www.cisco.com/en/US/docs/security/asa/asa84/configuration/guide/vpn_anyconnect.html#wp1090443.

Enabling AnyConnect VPN on the ASA

After defining the group-policy and connection profile on the ASA, the last step is to enable the AnyConnect VPN feature on the ASA. After enabling AnyConnect, the administrator can also configure additional features, such as pointing to the AnyConnect image software, NAM profile, and VPN Profile. The following example shows the configuration commands to enable AnyConnect modules:

```
webvpn
 enable outside
 anyconnect keep-installer installed ! This forces the anyconnect to remain installed on
the endpoint device, after the session is terminated.
 anyconnect modules none
```

Provisioning a Windows Device to Rmotely Connect to the Network

For a corporate approved device to obtain remote access capability, the user must perform the following steps at the campus location:

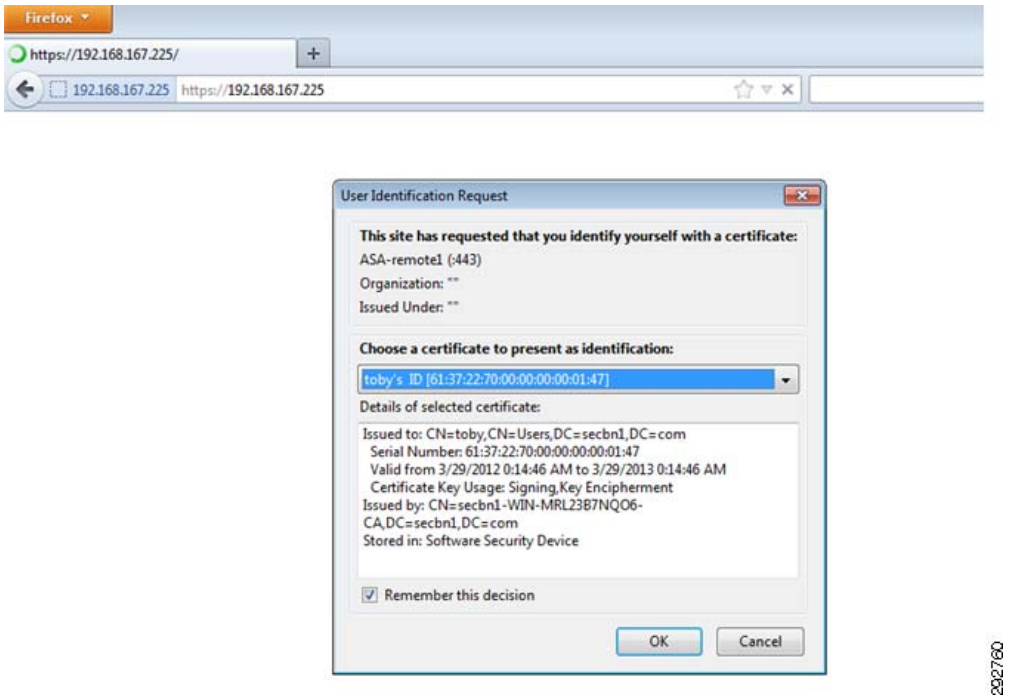
Step 1 Install the RSA SecurID application on the remote device and, with IT support, provision the software on the device.

Step 2 It is assumed that before the AnyConnect installation begins, the workstation has successfully completed the enrollment and provisioning process, which implies that the workstation has a valid digital certificate issued by the CA server.

The steps shown below are for one time installation. After the installation is completed, the user is never prompted for these steps.

Step 3 Initiate an SSL VPN session using a web browser to the ASA VPN gateway IP address, which is shown in Figure 26-7.

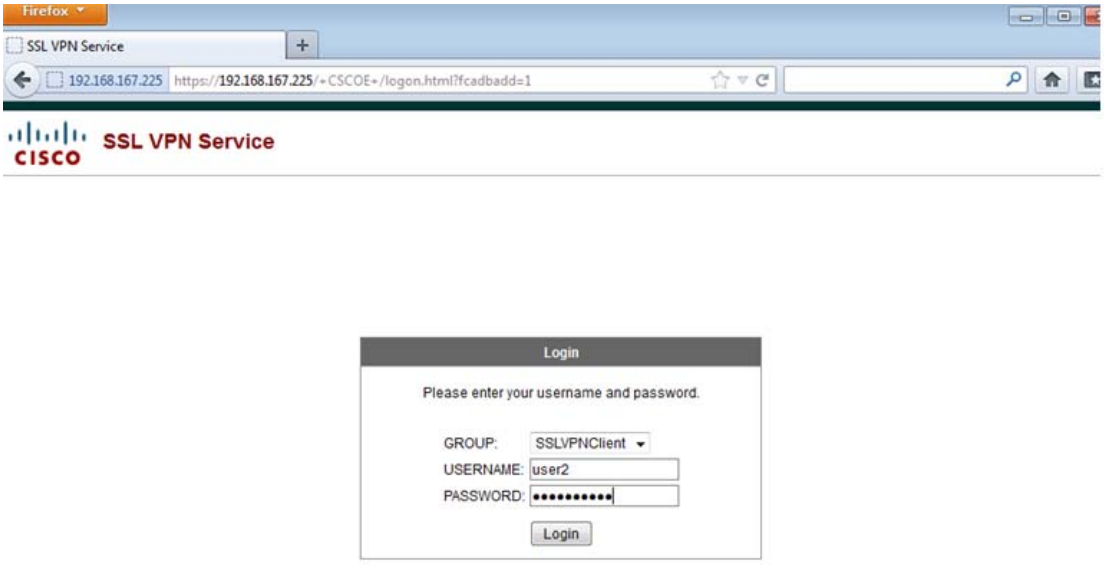
Figure 26-7 SSL VPN Session to ASA VPN Gateway IP Address



Note The certificates presented by the ASA remote endpoints and the identity certificate of the ASA must be signed by the same root CA server.

Step 4 The user is presented with login screen and the user needs to select the Group to which they belong. The user is expected to select the group-policy name and the valid credentials, as shown in Figure 26-8.

Figure 26-8 Selecting Group



Step 5 After the user credentials are validated, Cisco AnyConnect installation begins, which is depicted in Figure 26-9.

Figure 26-9 Cisco AnyConnect Installation

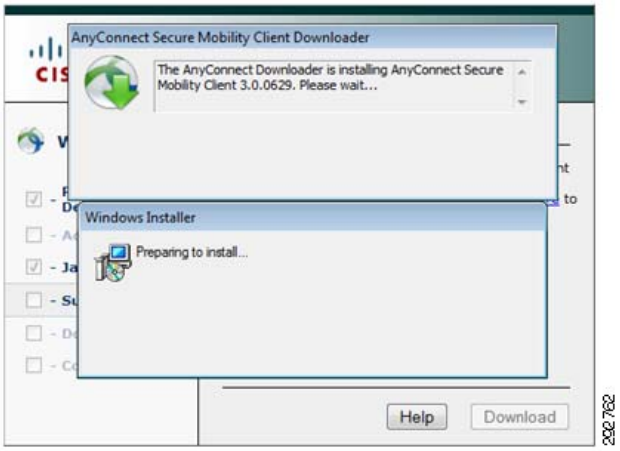


Figure 26-10 depicts successful installation of Cisco AnyConnect on the workstation:

Figure 26-10 Successful Installation of Cisco AnyConnect



Figure 26-11 shows the Windows workstation establishing a session.

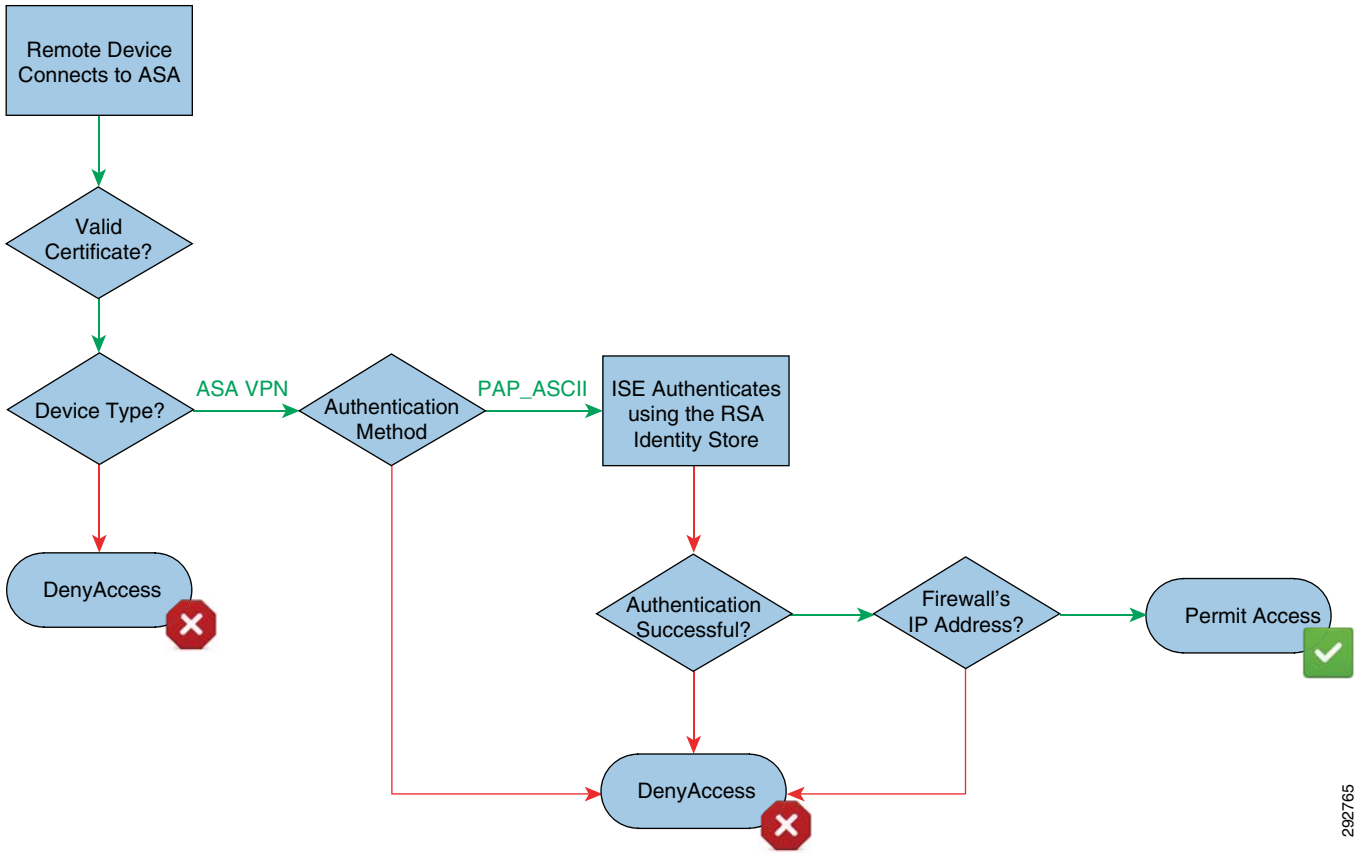
Figure 26-11 AnyConnect Initiates VPN Connection



As explained in the Connection Profile Configuration, when a remote worker connects to the network both ISE and the ASA authenticate the device. ASA initially validates the digital certificate of the remote user. If the certificate is valid, then the next step of authentication happens, which is through the RSA SecurID token. The remote worker is allowed to access the network if both authentications are valid.

The logic flow in Figure 26-12 shows what happens when a remote device accesses the network.

Figure 26-12 Logic Flow for a Remote Worker Accessing the Network



Verifying What ISE Policy Rules Are Applied

As shown in Figure 26-12, ISE validates the remote user in the following sequence:

1. Validate if the Device Type Equals ASA VPN. This is to ensure that only devices that are configured as VPN Type can initiate the communication with ISE.
2. Validate if the authentication protocol is PAP_ASCII. This is the protocol used by ASA to send the RSA Secure ID token passwords to the ISE, which the ISE sends to RSA Secure ID server for authentication.
3. In the authorization Rule, ISE validates the Source IP address of the ASA as a means to authorize the connection. In this design remote VPN users are only authenticated by ASA and ISE, and there is no authorization taking place. Figure 26-13 and Figure 26-14 detail the authentication and authorization rules in ISE.

The ISE rule shown in Figure 26-13 uses the RSA SecurID identity store for authentication.

Figure 26-13 Authentication Rule



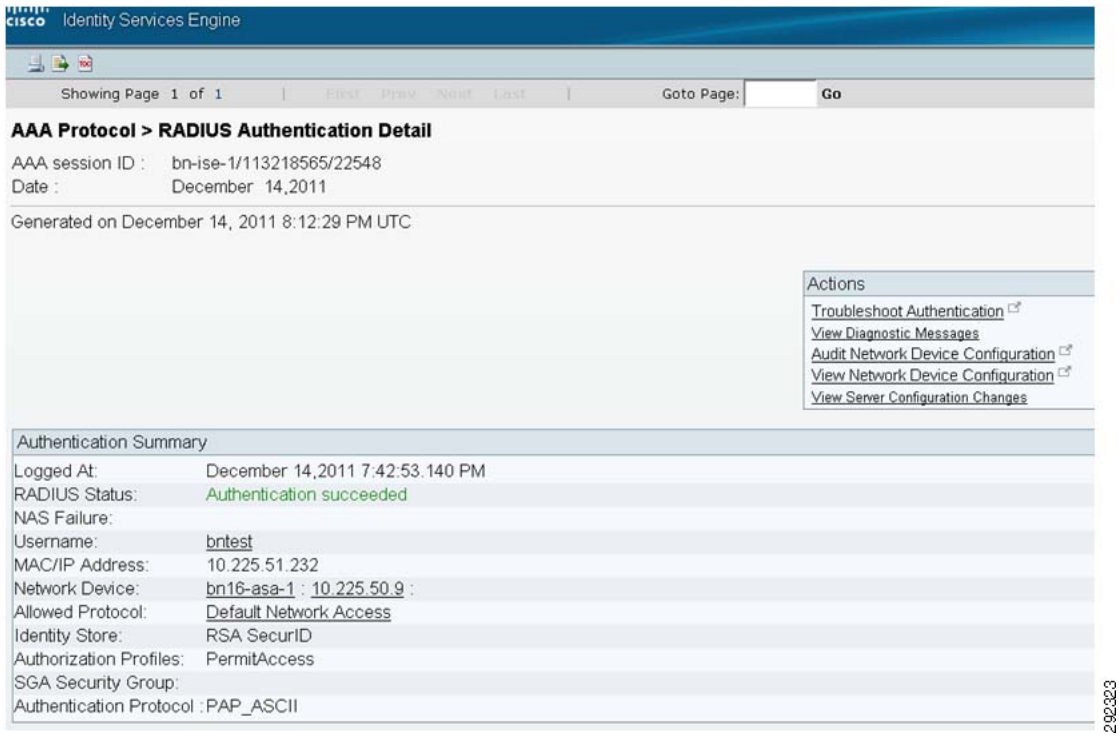
The ISE authorization rule shown in Figure 26-14 matched since a remote device connected and the Network Access: Device IP Address matches the ASA firewall's address.

Figure 26-14 Authorization Rule



To verify what rules were applied on the ISE, click **Monitor > Authentication**, as shown in Figure 26-15.

Figure 26-15 Log Information on ISE for Successful Remote Worker Authentication



Provisioning an Apple iOS Device to Remotely Connect to the Network

Similar to workstations, an Apple iOS device that is provisioned at the campus can establish an SSL VPN connection to the campus network remotely using Cisco AnyConnect. The following steps must be completed by the user before establishing SSL VPN connectivity:

Step 1 The iOS device should already have a digital certificate installed. Figure 26-16 shows an example of a provisioned device.

Figure 26-16 Provisioned Device



Step 2 The user should install the Cisco AnyConnect from Apple's App Store.

Step 3 Configure the profile on AnyConnect and select the certificate which is already installed in the device (Certificate Provisioning must happen before initiation remote VPN communication), as shown in Figure 26-17.

Figure 26-17 Configure Profile and Select Certificate



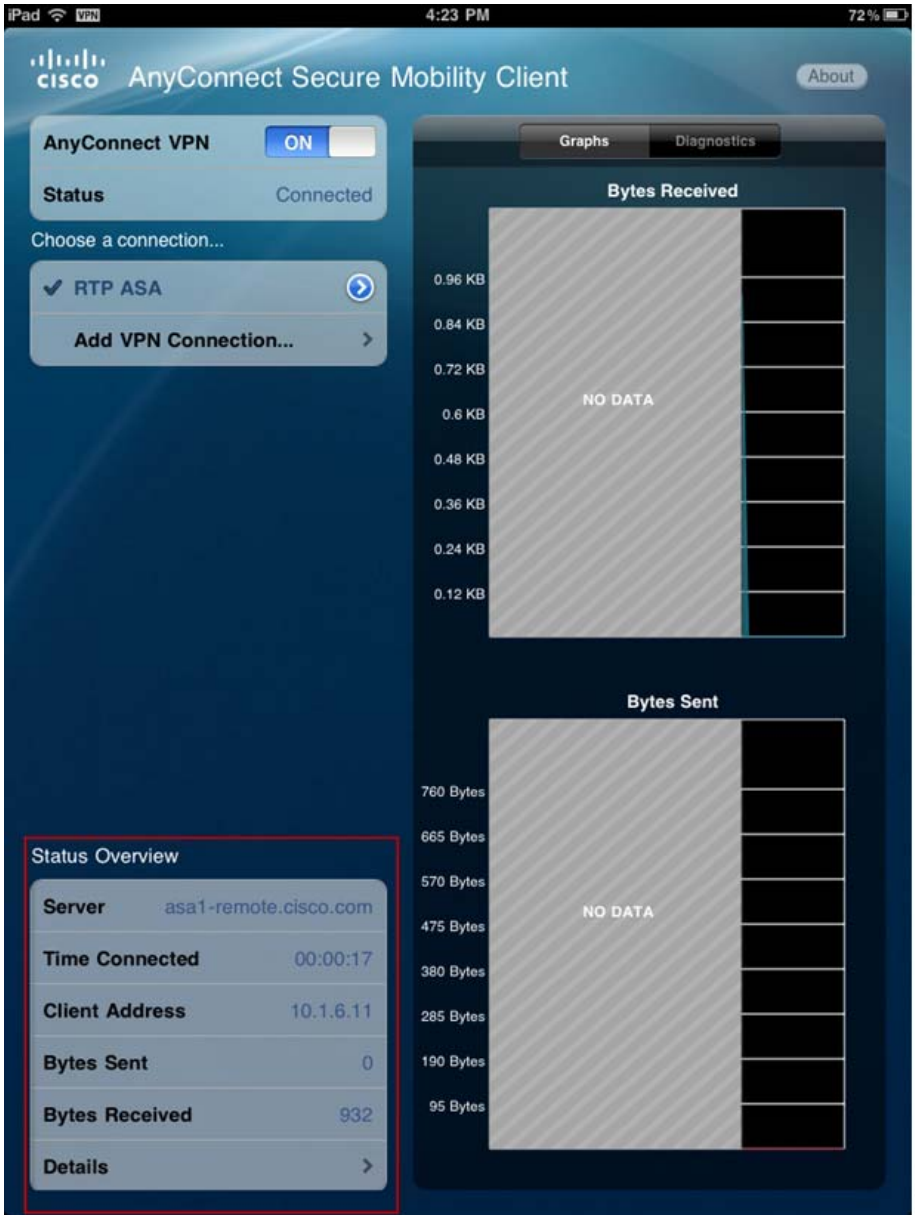
Step 4 Select the VPN Group, which the Network Administrator must inform the user about the right group to select, and enter the user credentials. ASA VPN authentication requires the user certificate and the user credentials. Hence the user credentials have to be entered, as shown in Figure 26-18.

Figure 26-18 User Credentials



Figure 26-19 shows a successfully connected SSL VPN session

Figure 26-19 Connected SSL VPN Session





BYOD Network Management

Revised: August 7, 2013

Network Management for BYOD is separated into three sections:

- Cisco Prime Infrastructure Overview—A brief overview that covers the basic capabilities of Cisco Prime Infrastructure. The subsequent sections are focused on specific abilities of Prime Infrastructure directly related to BYOD.
- BYOD User and Device Tracking—Uses information from multiple components, consolidated by Cisco Prime Infrastructure, to identify and track end users and end devices on the network.
- BYOD Template-Based Configuration—Covers using Cisco Prime Infrastructure as a management tool for configuring and maintaining the BYOD wireless configurations across Cisco Wireless LAN Controllers (WLC).

This document does not cover the basic implementation of Prime Infrastructure and assumes the WLCs are already managed by Prime Infrastructure. For further information on Prime Infrastructure implementation, refer to the Prime Infrastructure Configuration Guide: http://www.cisco.com/en/US/products/ps12239/products_installation_and_configuration_guides_list.html.

Key Acronyms and Terminology

Table 27-1 Acronyms and Terminology

Key Term	Explanation
Prime	Prime refers to Cisco Prime Infrastructure in this document. The Cisco Prime Family includes other products not covered in this document.
End Device	Also referred to as Endpoint. Both wired and wireless devices such as Android and Apple tablets and smartphones, wired IP phones, and laptops.
End User	Also referred to as User. Identified by “username” of one or more end devices.
WLC	Also referred to as Controller. Wireless LAN Controller
WLAN/SSID	WLAN (Wireless LAN) and SSID have a one-to-one relation and can be thought of as the same thing in this section.

Cisco Prime Infrastructure Overview

Cisco Prime Infrastructure is an exciting new offering from Cisco aimed at managing wireless and wired infrastructure while consolidating information from multiple components in one place. While allowing management of the infrastructure, Prime Infrastructure gives a single point to discover who is on the network, what devices they are using, where they are, and when. The capabilities of Prime Infrastructure and the other components featured go far beyond the focus of this document. A brief overview of Cisco Prime Infrastructure and supporting components follows.

Cisco Prime Infrastructure 1.2 is the evolution of Cisco Prime Network Control System 1.1 (NCS), adding additional infrastructure and wired device management and configuration capabilities while improving on existing capabilities in NCS 1.1.

Table 27-2 Cisco Prime Infrastructure 1.2

Prime NCS 1.1	Prime Infrastructure 1.2
• Wireless Network Management • Wireless and Wired Monitoring and Reporting • Template-based Configuration of WLCs • User and Device Tracking for Wired and Wireless	• All NCS 1.1 Capabilities • Improved UI • Template-based Configuration of Switches and Routers • Application and Media Performance Monitoring Capabilities

Prime Infrastructure and Supporting Components

Cisco Prime Infrastructure interacts with many other components to be a central management and monitoring portal. Prime Infrastructure has integration directly with two other appliance-based Cisco products, the Cisco Mobility Services Engine and Identity Services Engine for information consolidation. Prime Infrastructure controls, configures, and monitors all Cisco Wireless LAN Controllers (WLCs), and by extension, all Cisco Access Points on the network. Prime Infrastructure also configures and monitors Cisco Catalyst switches and Cisco routers.

Figure 27-1 Prime Infrastructure Component Interaction Summary

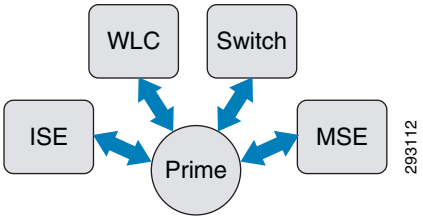


Table 27-3 Prime Infrastructure Components

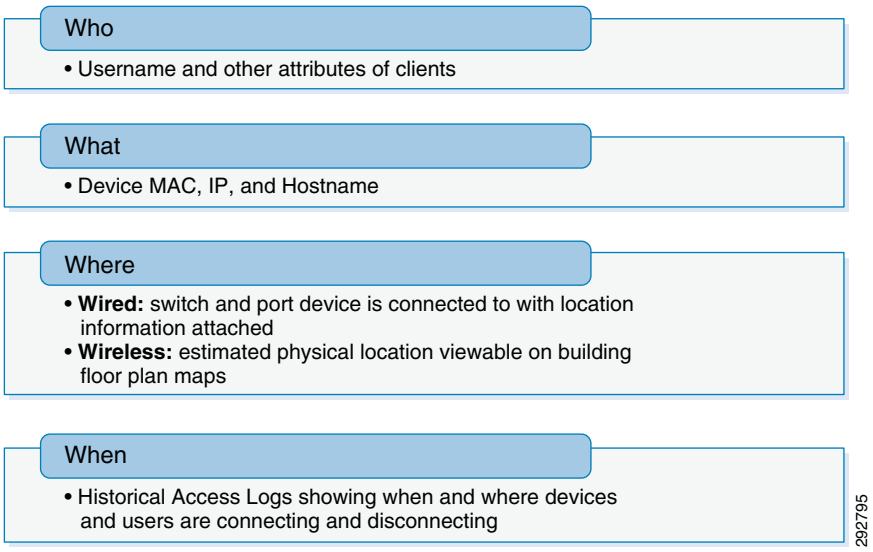
Prime Infrastructure	Cisco Prime Infrastructure is the core component that sends information to and consolidates information from the other four component types.
ISE	Cisco Identity Services Engine is the core component of BYOD for user and device authorization and access to the network. ISE provides user information to Prime Infrastructure.
WLC	Cisco Wireless LAN Controller is configured, controlled, and monitored by Prime Infrastructure. WLCs provide Prime Infrastructure with a wealth of real-time wireless environment and client device information.
Switch/Router	Cisco switches and routers are configured, controlled, and monitored by Prime Infrastructure. Wired device information is provided to Prime Infrastructure to be consolidated with wireless device information.
MSE	Cisco Mobility Services Engine complements Prime Infrastructure with current and historical location, usage, and other information for all devices Prime Infrastructure sees.

The following link has more information about Cisco Prime Infrastructure and the rest of the Cisco Prime family of products: http://www.cisco.com/go/prime.

BYOD User and Device Tracking

The ability to track users and devices on the wired and wireless networks is critical to knowing who is accessing the network, with what they are accessing it, where are they accessing it, and when they accessed it.

Figure 27-2 Who, What, Where, and When Summary



Understanding who is accessing the corporate network, what they are using, and where they are connected allows customers to better understand:

- Location and movement of employees and devices on the network
- Suspicious or unauthorized access of the network
- Location of missing or stolen assets, such as in a college campus environment
- Location of unknown devices on the network
- Current utilization of the network

Adding historical logging of when users and devices access the network allows:

- Persistent records of when users and devices accessed the network and their specific locations
- Searchable historical data of user and device access for tracking and troubleshooting issues
- Historical port utilization data

Components

Cisco Prime Infrastructure is the central portal for user and device tracking. Prime Infrastructure uses information from multiple places to give a single, consolidated view of current and historical user and device access to the network. Figure 27-3 adds to Figure 27-1 showing how the components cover the Who, What, Where, and When aspects of user and device tracking.

Figure 27-3 Prime Infrastructure Component Interaction Summary for User and Device Tracking

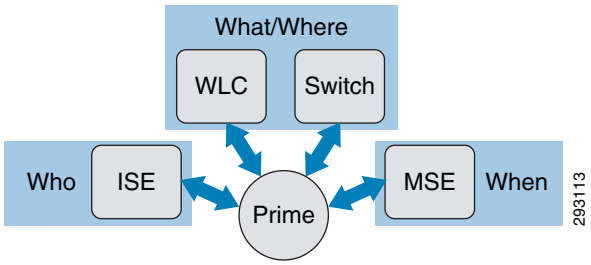


Figure 27-4 shows a more detailed view of how Prime Infrastructure interacts with the rest of the architecture. The five main components needed for User and Device tracking of both wired and wireless users are listed below with brief summaries of each.

Figure 27-4 Prime Infrastructure Interaction with Infrastructure Components

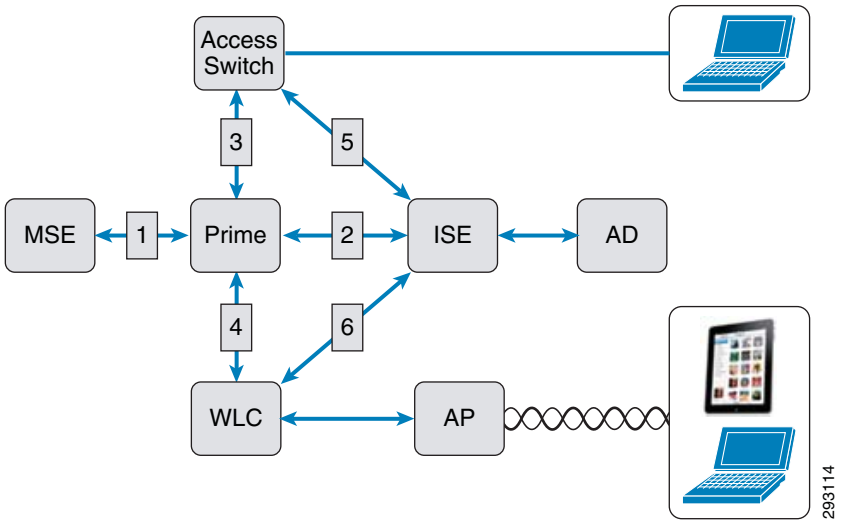


Table 27-4 Prime Infrastructure Interaction with Other Infrastructure Components

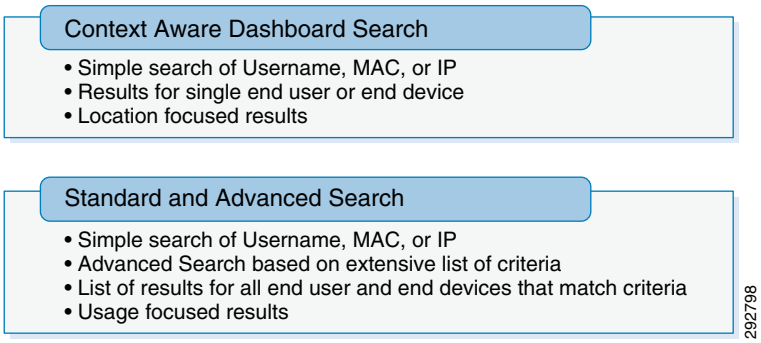
	Components	Communication
1	Prime—MSE	Prime receives current and historical location information for mobile devices
2	Prime—ISE	Prime receives user information including username and device MAC and authentication history
3	Prime—Switch/Router	Prime receives wired device information including port and MAC. Prime sends/receives component configuration.
4	Prime—WLC	Prime receives wireless user and extensive device information. Prime sends/receives component configuration.
5	Switch—ISE	RADIUS authentications
6	WLC—ISE	RADIUS authentications

To locate and track users and devices, Prime Infrastructure pulls information from all of these sources, consolidating it based mainly on common MAC. Prime Infrastructure is device focused and displays detailed reports based on a particular device. Prime Infrastructure also has the ability to show all devices with which a particular user accesses the network, giving the ability to track a particular user across multiple wireless and wired devices.

Locating Users and Devices

There are two basic ways to display information on users and devices. Both options are considered "Search" options, although the abilities to filter and display based on an extensive list of criteria goes far beyond what most would consider a simple search option.

Figure 27-5 Types of Search

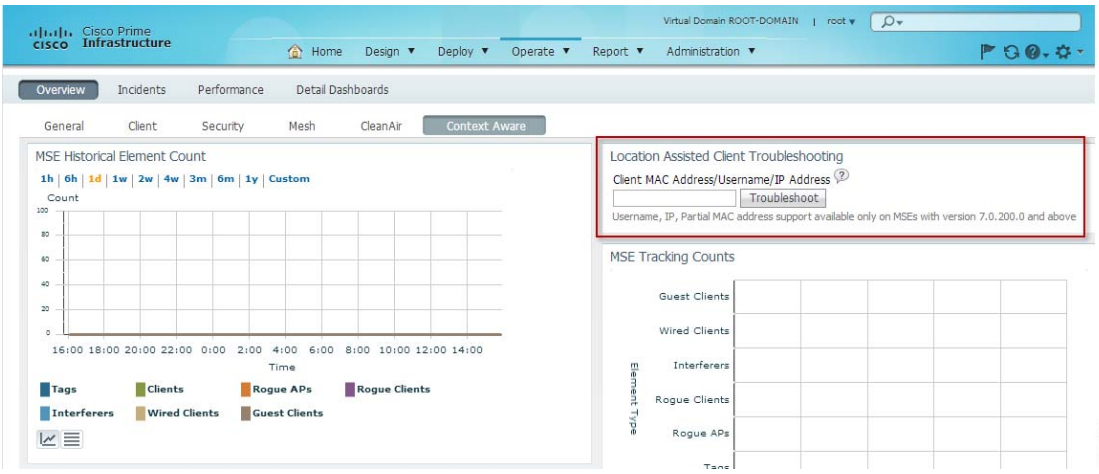


Context Aware Dashboard Search

Context Aware search is used to display information on a single device based on current MAC, IP, or Username of the end user of the device. While limited in how you can search and what is displayed, this option does give you a slightly different view of location information compared to the standard search.

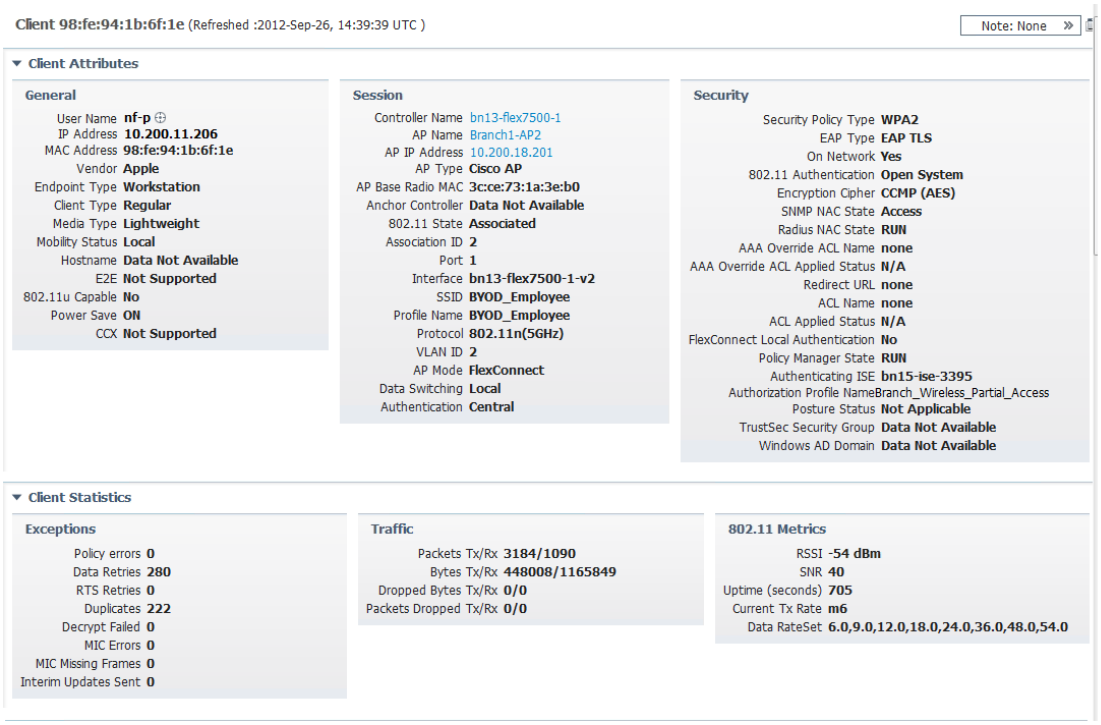
The Context Aware Dashboard in Figure 27-6 has a search box titled "Location Assisted Client Troubleshooting", which is where the search is executed. The search instantly resolves the MAC, IP, or Username to the device and display that device only.

Figure 27-6 Context Aware Dashboard



The results shown in Figure 27-7 are common to both types of search and give quite a bit of current information about the device and end user of it, if there is one.

Figure 27-7 Context Aware Search Standard Results

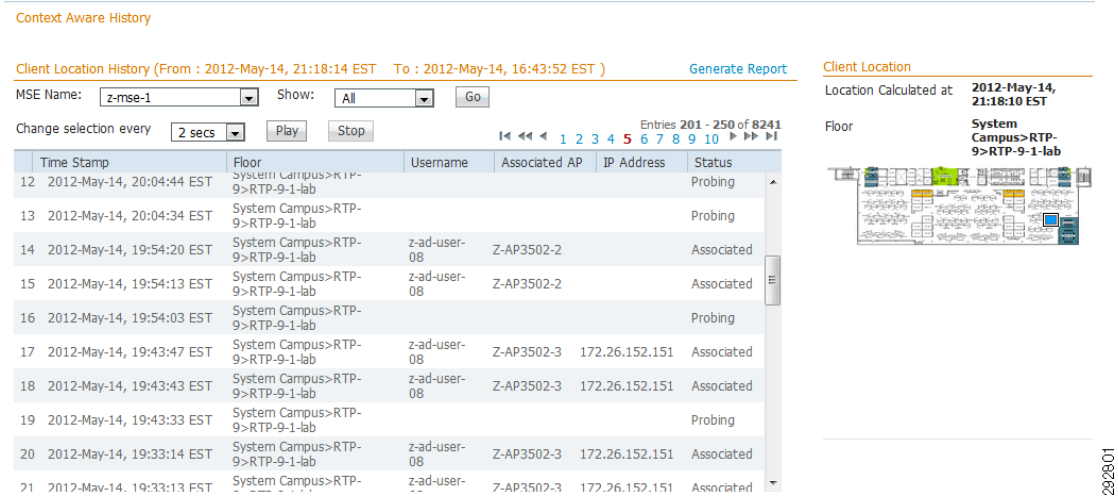


The information in the search results unique to the Context Aware Search is location based. Standard and Advanced Search show “Association History”, which includes location, but not in the same format.

Context Aware Search results give you the ability to easily see exactly where a device was at a given time as well as show historical motion of the device. Using the “Play” feature, the device location is shown being updated on a map for a visual representation of movement, which can be accurate down to several feet in a properly implemented wireless network.

Figure 27-8 shows the location results with a blue square showing current location on the floor plan map adjacent. Pressing “Play” would show the blue square moving as location references were cycled through.

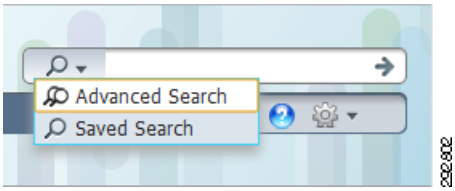
Figure 27-8 Context Aware Search Location Results



Standard and Advanced Search

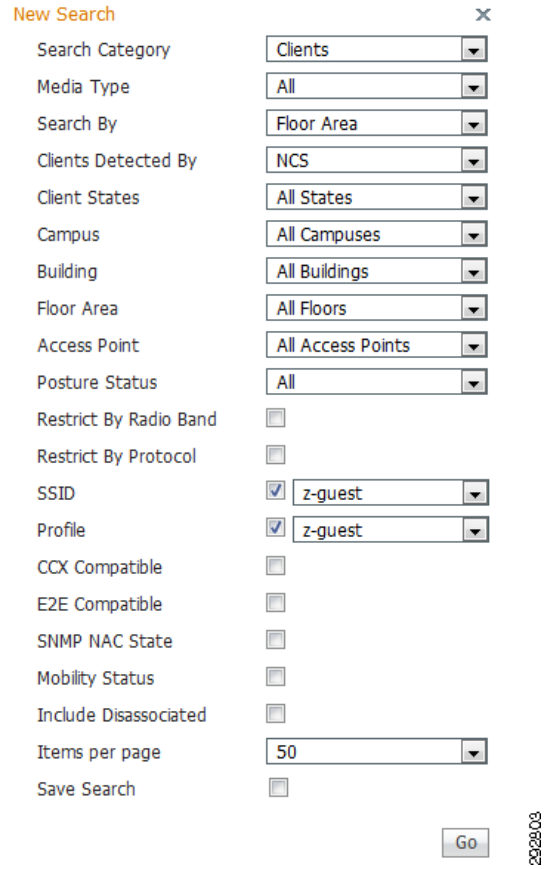
In addition to searching for clients or devices using the Context Aware Dashboard, Standard and Advanced Search may be performed in the top right corner of the Prime Infrastructure interface, visible at any time. Much more granular searches may be performed using the Advanced Search option shown in Figure 27-9. Search results are more device usage focused with the Standard and Advanced Search, but still contain location information.

Figure 27-9 Standard and Advanced Search Box



With Advanced Search, results for many end users or end devices that meet a particular set of criteria may be displayed instead of looking for one particular user or device. Parameters such as physical location, type of user, SSID, and even posture/authentication status may be used. Figure 27-10 shows a subset of criteria available.

Figure 27-10 Advanced Search Criteria

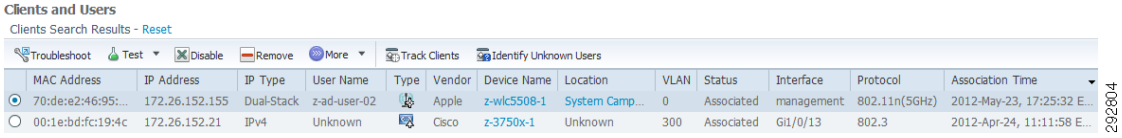


The form above is dynamic, changing as selections are made, which means this image shows only a subset of the criteria available for the "Clients" category, which in this case refers to both end devices and end users.

Additional search criteria and information may be found in the Cisco Prime Infrastructure Configuration Guide: http://www.cisco.com/en/US/products/ps12239/products_installation_and_configuration_guides_list.html.

Figure 27-11 shows the search results list, which contains both wired and wireless users unless one type is filtered out. In this example two devices are shown, the first a wireless device and the second wired.

Figure 27-11 Standard and Advanced Search Results List



An extensive amount of information is available from just the search results screen. The result columns may be customized and results list sorted by any of those columns. Figure 27-12 shows a list of available columns.

Figure 27-12 Search Results Columns

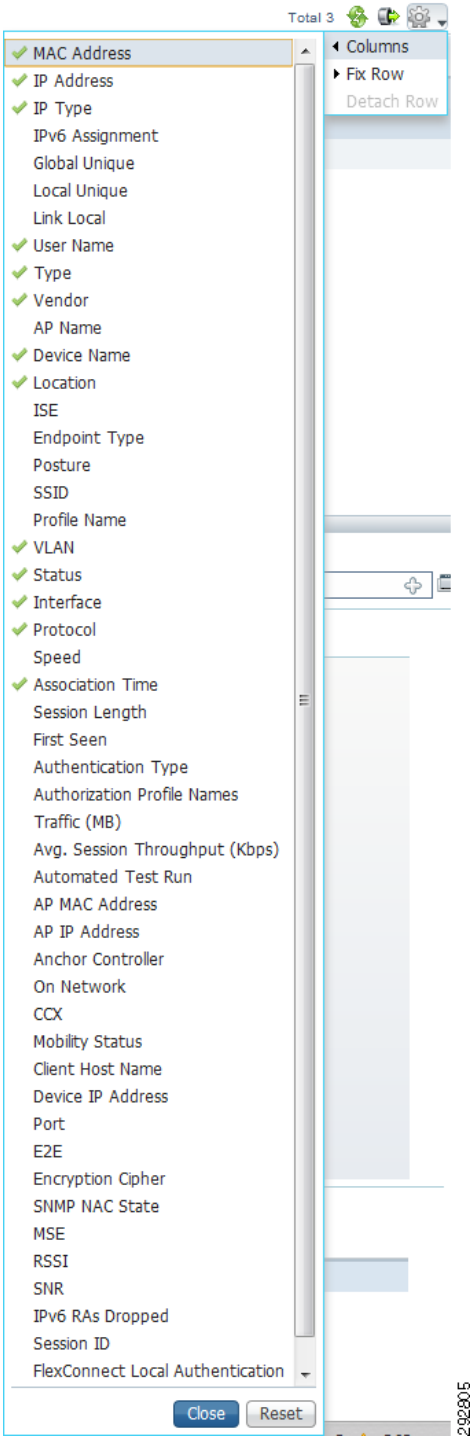


Figure 27-13 through Figure 27-18 show examples of basic and extended information shown for individual devices selected from the search list.

Figure 27-13 shows the same basic end user and end device information as the Context Aware Search discussed earlier.

Figure 27-13 General User and Device Details

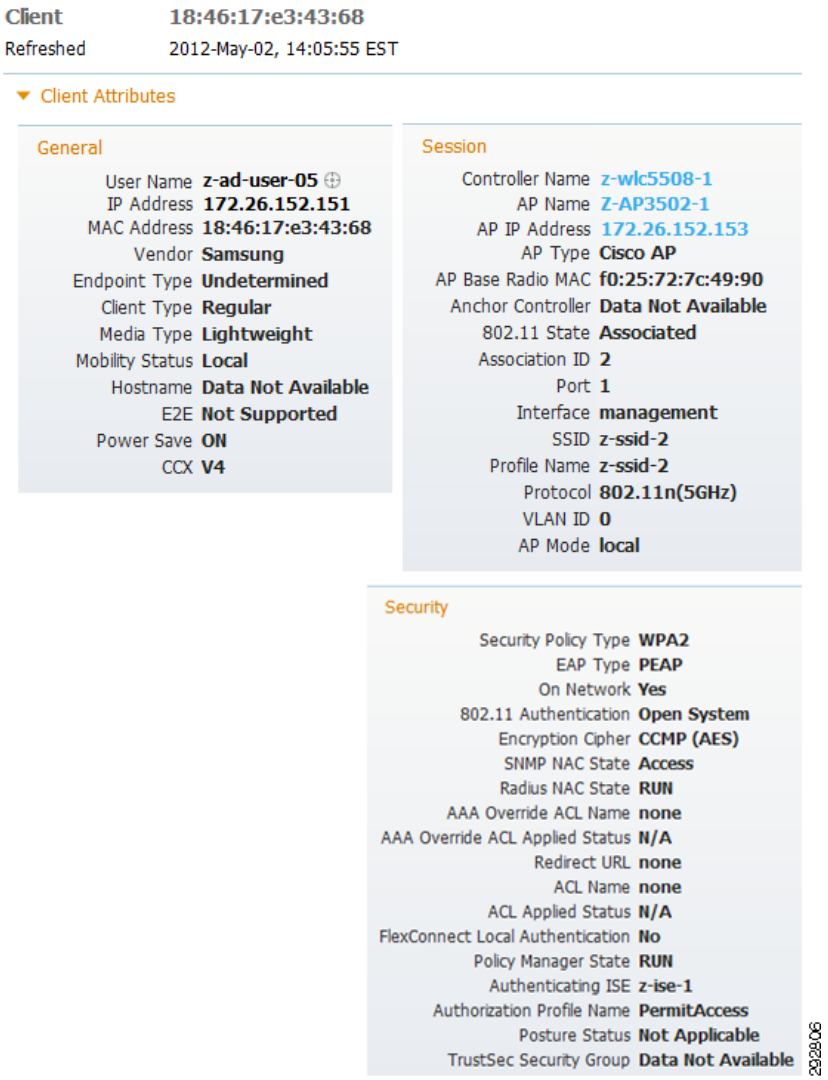


Figure 27-14 shows association times, durations, and locations, which is similar but not the same as the location history with the Context Aware Search.

Figure 27-14 Device Association History

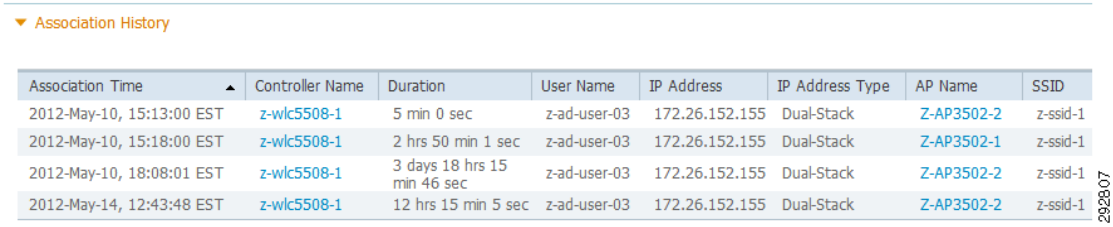


Figure 27-15 is pulled directly from ISE and shows recent authentication successes and failures.

Figure 27-15 Device Authentication History

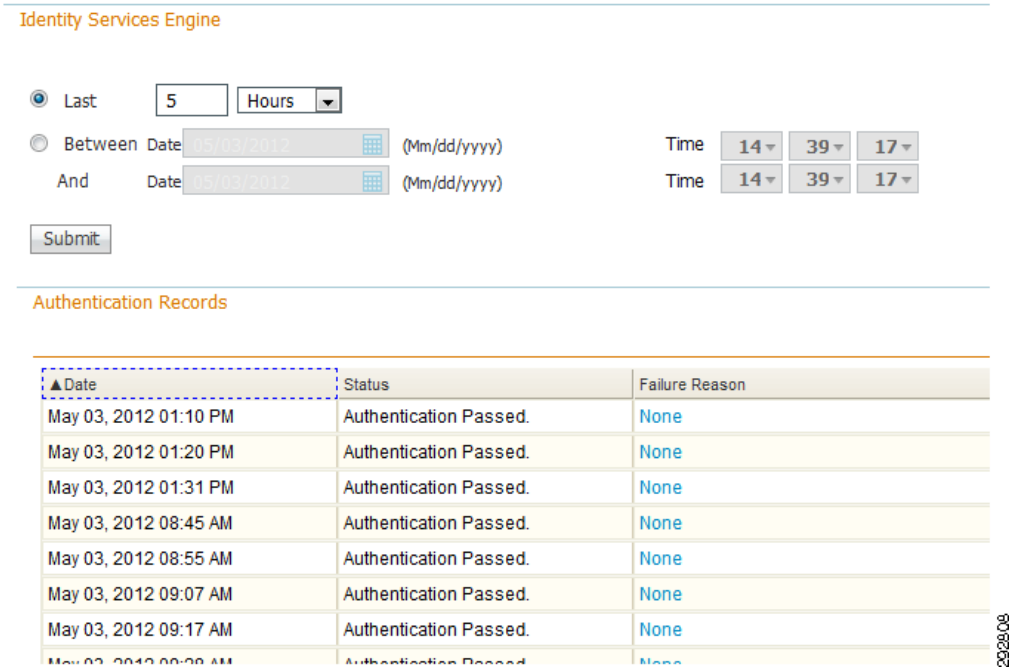


Figure 27-16 shows signal quality for various, changeable time frames in graph format.

Figure 27-16 Device Signal Quality and Usage History



Figure 27-17 shows the device in its current location, along with any additional information chosen. In this example, only heat map and AP location is selected, but many other items are available for display, such as interfering devices and other clients.

Figure 27-17 Floor Plan Heat Map with APs and Client Device

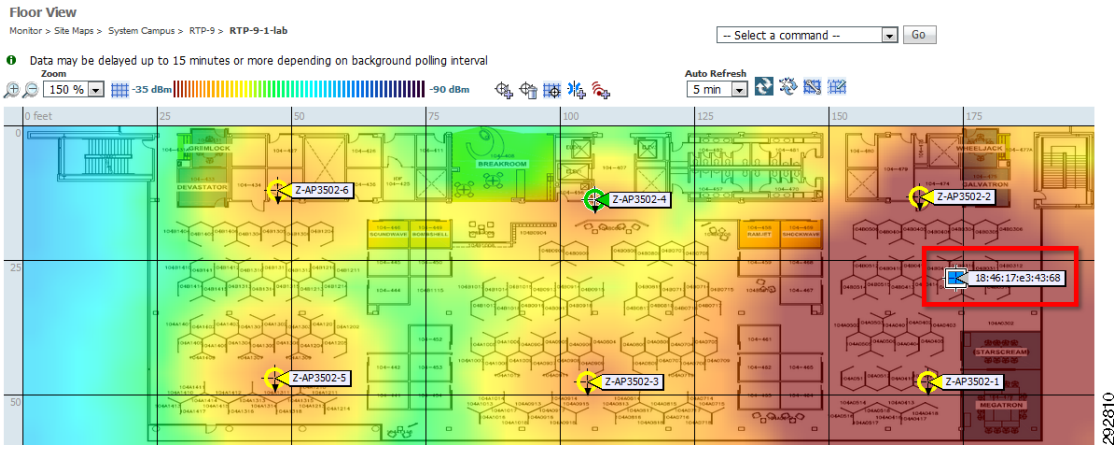
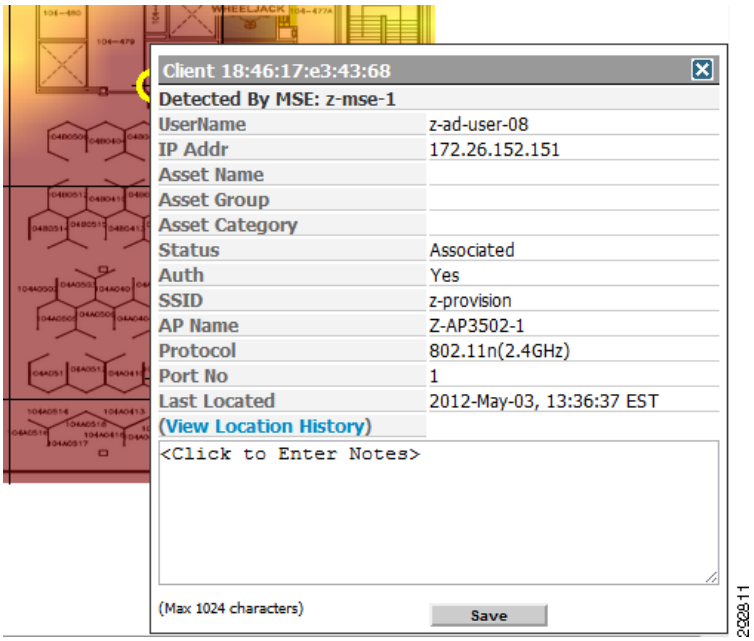


Figure 27-18 shows details of any device shown on the heat map.

Figure 27-18 Device Detail Pop-up from Heat Map



BYOD Template-Based Configuration

This section covers the use of Cisco Prime Infrastructure to deploy and maintain configurations to Cisco Wireless LAN Controllers (WLCs) matching the BYOD configurations referenced in this document.

Cisco Prime Infrastructure has the ability to control configuration of Cisco Wireless LAN Controllers (WLCs) directly or through the use of templates. One template will not configure the entire controller. Templates are separated out to granularly cover each feature of the controller. Templates exist for just about every small feature that can be implemented on the controllers and many portions of the templates can be modified during deployment to accommodate unique settings in WLCs. Templates can be configured for a common configuration across all WLCs as well as be implemented across a sub-set of WLCs or individual WLCs.

Note Each WLAN has exactly one SSID and the two terms may be thought of as being the same thing for simplicity in understanding this content: **WLAN = SSID**.

Template-based configuration has a number of advantages compared to individual configuration of WLCs:

- Consistent Configuration of WLCs
- Multiple Templates for Variations of Deployment
- Rapid Deployment of New or Replacement Components
- Staged Rollout of Configuration Changes with Rapid Rollback

Consistent Configuration of WLCs

Inconsistencies in configuration can easily occur when configuring multiple WLCs through their local web-based administrative interfaces. Inconsistencies can have far reaching negative impact on WLAN functionality, security, and performance.

Inconsistencies in even the order of configuration can sometimes have serious impacts. For example, configuring multiple WLANs in different orders on different controllers will cause the WLAN IDs (an integer that uniquely identifies each WLAN) to be inconsistent. WLAN IDs are used by ISE to determine how the client should be treated. Inconsistent WLAN IDs may result in a client attaching to a particular SSID and being assigned access as if they were attached to a different SSID.

One important note here, however, is that Prime Infrastructure will have the controller auto-assign the WLAN ID. If the base configuration of the controllers starts in an inconsistent state, such as a WLAN existing on one controller that does not exist on another, the WLAN IDs will be set inconsistently when applied from Prime Infrastructure. Checks should be done to ensure the WLAN IDs are consistent across all controllers.

Multiple Templates for Variations of Deployment

Variations in deployment may be required for some features of WLCs based on model or location in the network. If a WLC is being used for dedicated guest access, its configuration for certain features would differ from other WLCs on the network, requiring some variation of templates.

Prime Infrastructure supports multiple templates for the same feature, allowing templates to be created with variation for WLCs. Templates may be applied to all WLCs or individually selected WLCs at the time of application.

Rapid Deployment of New or Replacement Components

By creating templates for WLC configuration, new and replacement WLCs may be rapidly configured from the latest templates, reducing time to deploy and eliminating errors from misconfigurations.

Staged Rollout of Configuration Changes with Rapid Rollback

Multiple templates for a specific feature can be created allowing an altered configuration to co-exist with the current configuration in template form. The new configuration template can then be tested on one or more WLCs with ease of rollback to the previous configuration template should issues arise.

Note The acronym WLC (Wireless LAN Controller) is frequently used in this document while some of the interfaces shown use the common term “Controller”. In this document “WLC” and “Controller” refer to the same thing: **WLC = Controller**.

Template Creation and Implementation

Template creation and implementation is a fairly straightforward process in Prime Infrastructure with a few caveats. The templates and configurations that follow are specific to the BYOD solution in this document and are but a tiny subset of the many settings and features that are needed for implementing an enterprise wireless network.

This section assumes the WLCs are already managed by Prime Infrastructure. For further information on Prime Infrastructure implementation, refer to the Prime Infrastructure Configuration Guide: http://www.cisco.com/en/US/products/ps12239/products_installation_and_configuration_guides_list.html.

Template Creation

Template creation for WLC configuration may be handled in one of three ways:

1. Create new templates directly in Prime Infrastructure.
2. Configure a WLC through Prime Infrastructure and then create templates from that configuration.
3. Configure a WLC through the local WLC web interface and then create templates from that configuration.

The following material focuses on option 2, configuring a WLC through Prime Infrastructure followed by creation of templates to configure additional WLCs and make changes to the original WLC. This approach would seem the most logical since it incorporates option 3 as well if the WLC were already configured.

This approach would also be appealing to someone wanting to learn the interface and operation of Prime Infrastructure and the WLC at the same time, using a separate WLC to create the configurations and templates to be deployed in production at a future date. For base template creation, functional trials, and understanding of the solution, most of the features covered in this document may be deployed using a relatively inexpensive Cisco 2504 with the base license and a single AP. Two key features the 2504 lacks as part of the BYOD solution are the ability to act as a DMZ Guest WLC and the ability to rate limit traffic. Both of those features are covered in Chapter 21, “BYOD Guest Wireless Access.” All other features and abilities in the BYOD solution are supported on this platform.

Due to the extensiveness of configuration for a FlexConnect environment, not every step is shown for the initial configuration. Using the Prime Infrastructure interface instead of the WLC interface directly should be fairly straightforward. Minor differences in location of options and features in the Prime Infrastructure interface are shown.

Using option 2 from above (Configure a WLC through Prime Infrastructure and then create templates from that configuration) the following steps are used. If beginning with a WLC that was directly configured, just skip steps 1 and 3.

- Step 1—Configure Base Network Connectivity on the New WLC
- Step 2—Add the WLC as a Managed Device to Prime Infrastructure
- Step 3—Using Prime Infrastructure, Directly Configure the WLC
- Step 4—Create Templates from the Configured WLC
- Step 5—Deploy Templates on One or More WLCs

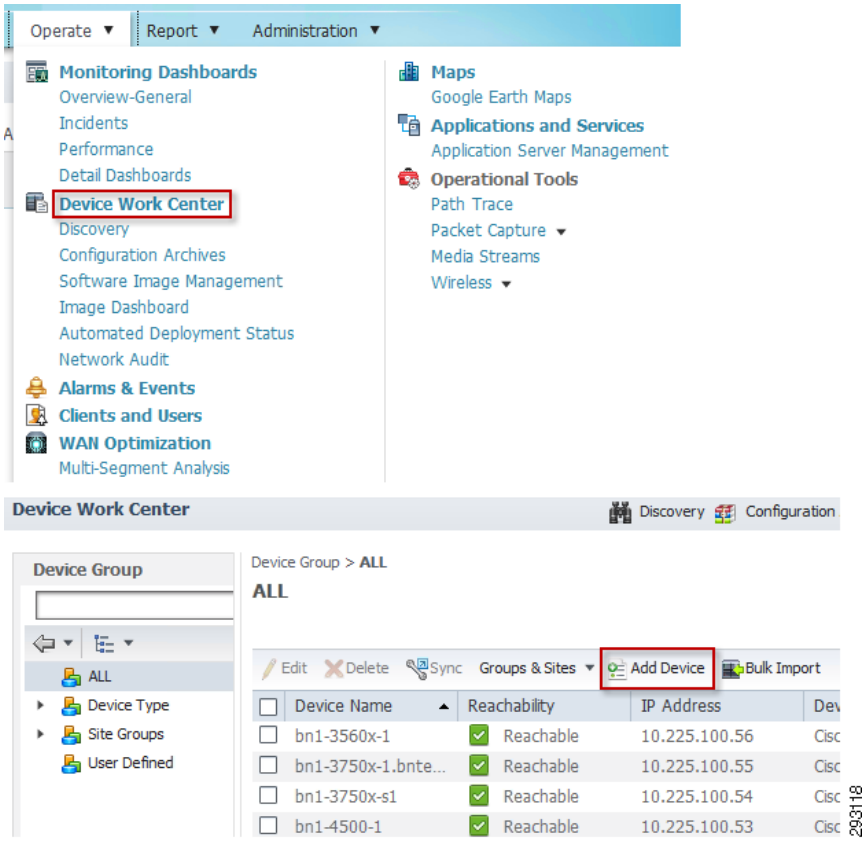
Step 1—Configure Base Network Connectivity on the New WLC

This step should be completed with documentation for the WLC you are implementing. Documentation for all Cisco WLCs can be found at: http://www.cisco.com/web/tsweb/redirects/mm/support/wireless.html.

Step 2—Add the WLC as a Managed Device to Prime Infrastructure

In Prime Infrastructure, use the Device Work Center and manually add the device, as shown in Figure 27-19.

Figure 27-19 Device Work Center—Add a Device



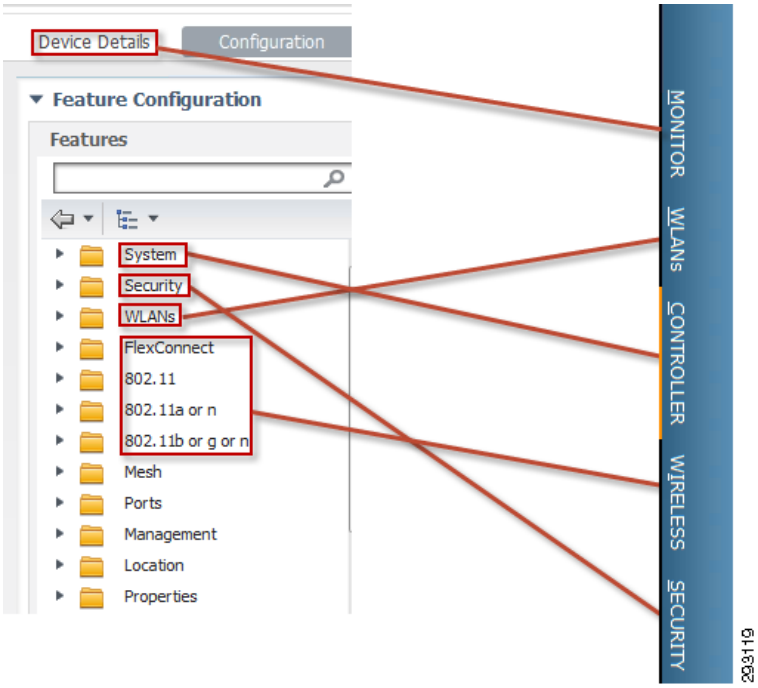
You do not need to specify the WLC type as Prime Infrastructure will determine it during the synchronization process. Alternatively, the WLC may be added through the discovery process, which is not shown.

After the WLC is added it will synchronize any existing configuration with Prime Infrastructure. This process should take only a couple of minutes and show a Device Status of "Managed" in the Device Work Center. It will also be placed in the appropriate Device Type folder, which can be expanded on the left side of the Device Work Center screen shown in Figure 27-19.

Step 3—Using Prime Infrastructure, Directly Configure the WLC

The WLC may now be directly configured in the Device Work Center, similar to being on the web-based interface of the WLC itself, by selecting the WLC and then the **Configuration** tab in the section below. The configuration interface is very similar, but not exactly the same. Figure 27-20 shows how the WLC interface main categories map to the Prime Infrastructure categories.

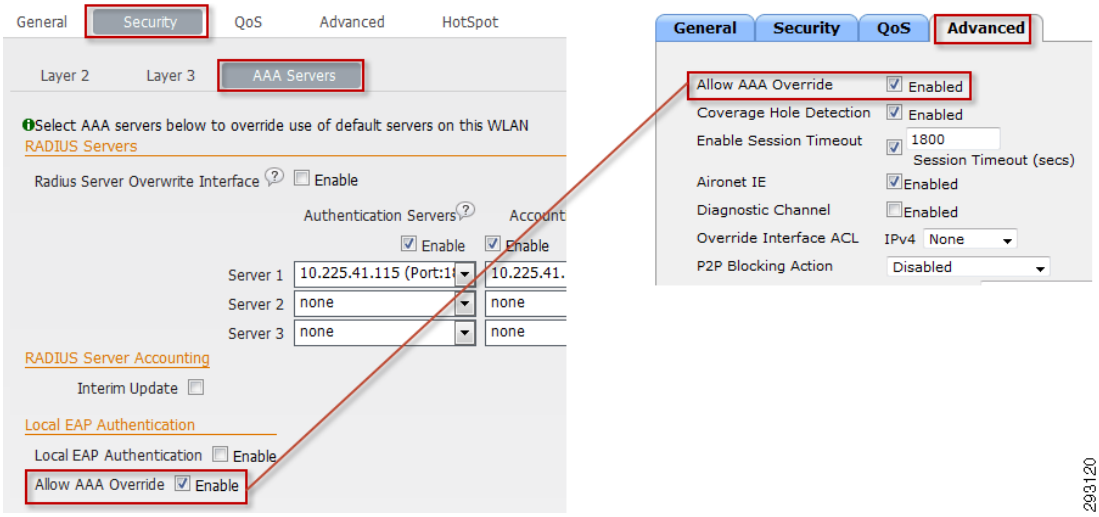
Figure 27-20 Mapping of WLC Interface Categories to Prime Infrastructure Categories



Be aware that one significant feature, AAA Override, is in a different location.

The feature AAA Override is shown in the **Advanced** tab of the WLAN settings when configuring through the WLC interface. This same feature is in the **Security** tab of the WLAN settings when configuring through Prime Infrastructure, as shown in Figure 27-21.

Figure 27-21 AAA Override on Advanced Tab of WLAN Settings

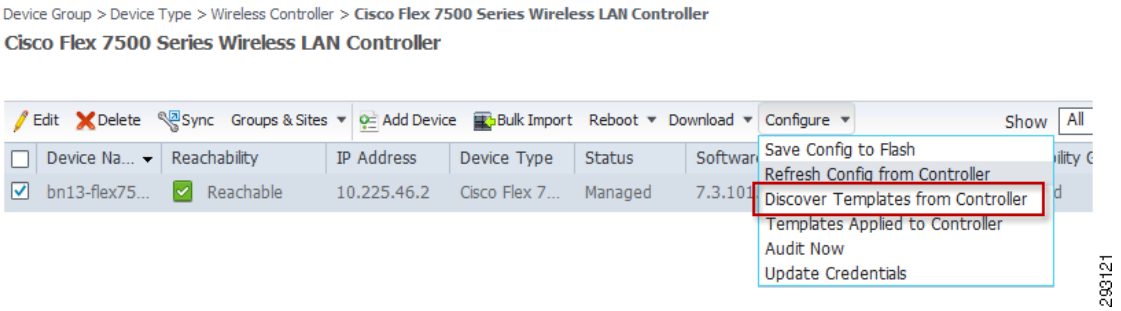


Take note of the WLAN ID caveat at the end of this section as it is important to both the creation of WLANs initially as well as template-based deployment of WLANs.

Step 4—Create Templates from the Configured WLC

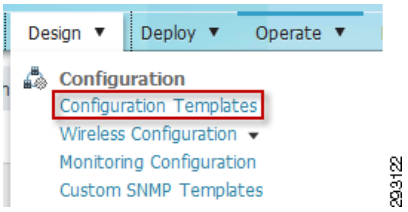
Creating templates from a configured WLC is a fairly simple process. An automated process creates templates of everything in the WLC that can have a template created. To accomplish this, go to the device in **Device Work Center**, the same place as the last step. Select the configured WLC and choose **Configure** and then **Discover Templates from Controller**, as shown in Figure 27-22.

Figure 27-22 Discover Templates from Controller



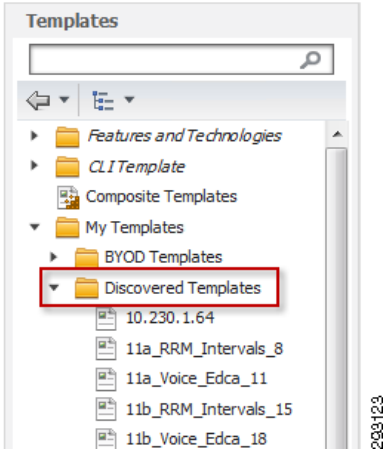
After the template discovery process completes, the templates are found in the **Configuration Templates** section in the **Design** section from the top menu, as shown in Figure 27-23.

Figure 27-23 Configuration Templates



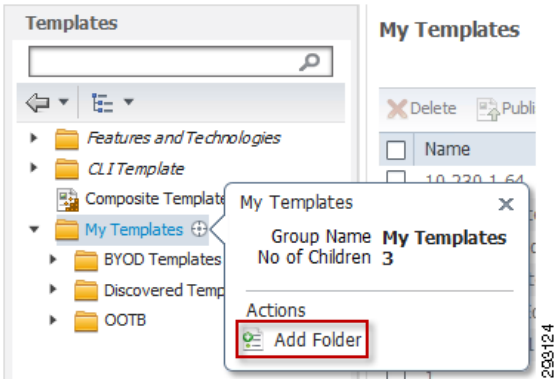
The newly discovered templates will be shown under **My Templates**, then **Discovered Templates** as shown in Figure 27-24.

Figure 27-24 Discovered Templates



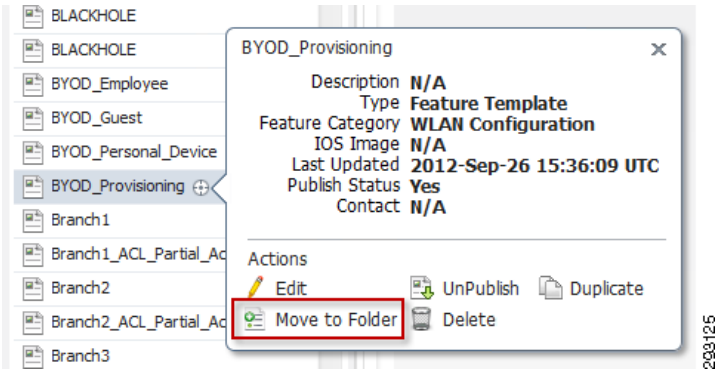
There will be many templates shown in this section, but only a small number of them are really needed. Before any customization or deployment of templates occurs, it is highly recommended to organize the needed templates into a custom folder. First, create a new folder by placing the mouse pointer next to **My Templates**, which pops up a box. Click **Add Folder**, as shown in Figure 27-25.

Figure 27-25 Add Folder



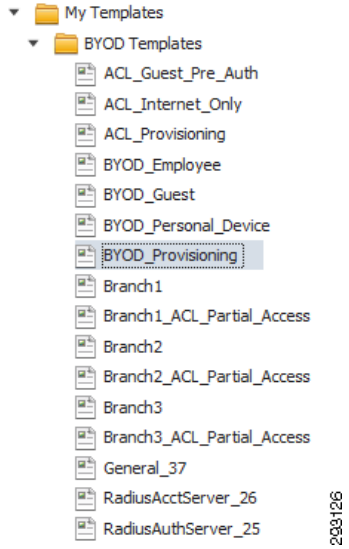
After creating the folder, in this case named BYOD Templates, place the pointer next to each of the desired templates, one at a time, and click **Move to Folder**, moving them to the newly created folder, as shown in .

Figure 27-26 Moving Templates



Once complete, all the desired templates will show up in the new folder, ready for editing and deployment, as shown in Figure 27-27.

Figure 27-27 BYOD Templates



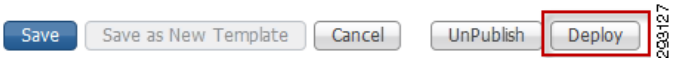
Step 5—Deploy Templates on One or More WLCs

It is fairly straightforward to deploy standard templates with no unique settings. Deploying templates that require unique configurations, such as FlexConnect Groups, is more involved.

A FlexConnect Group has specific APs associated with it, which will be different from WLC to WLC. A simple static template would not be particularly useful and the deployment must accommodate customization. The following template deployment is of a FlexConnect Group, showing the most complex type of deployment as an example.

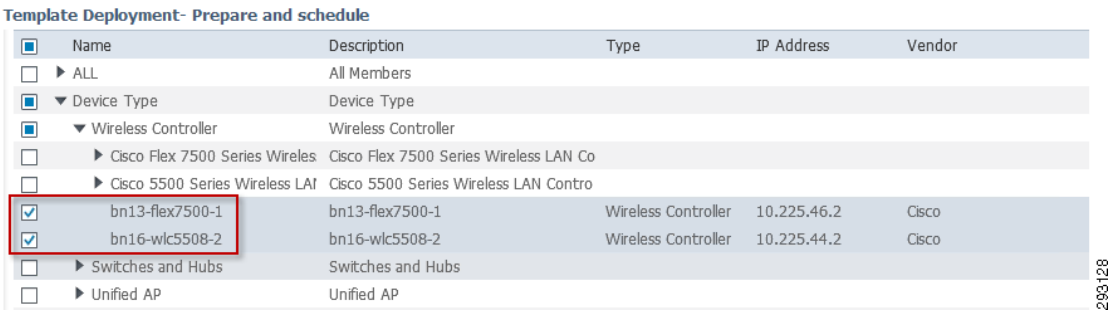
At the bottom of every template is a button to deploy it, shown in Figure 27-28, which when clicked shows the deployment screen.

Figure 27-28 Deploy Button



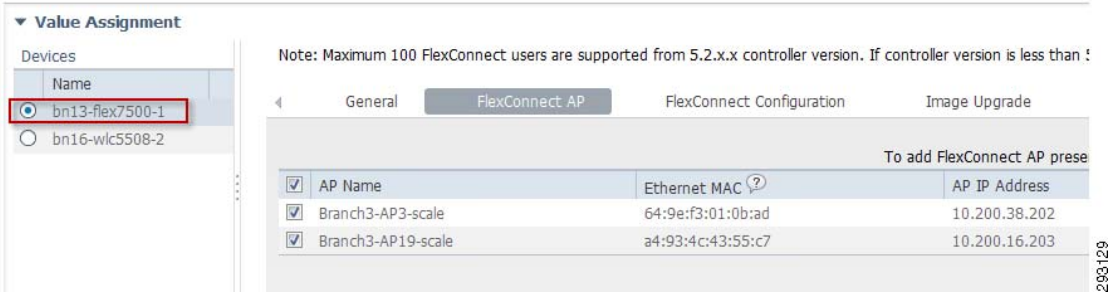
When a FlexConnect Group template is launched, the target WLCs must be chosen. In this example two WLCs of different types are selected, as shown in Figure 27-29.

Figure 27-29 Deployment Screen



When selected, you see the **Value Assignment** screen, shown in Figure 27-30. This section allows you to assign values and resources to each WLC independently. In this example APs may be added to the FlexConnect Group separately for each WLC. The APs are added by clicking **Add AP** on the far right of the screen (not shown) which will bring up a list of all APs that are visible to Prime Infrastructure.

Figure 27-30 Value Assignment



After completion of customization, the template can be deployed immediately or scheduled.

Caveat 1—WLAN ID

WLAN ID is used by ISE in determining what SSID (WLAN) clients are using to connect to the network. This ID is unique to each WLAN on each controller, so ensuring each WLAN has the same WLAN ID on each controller is essential for proper operation and security.

Ensuring this can become complex for large enterprise customers with multiple WLCs. Take note of the following:

- Prime Infrastructure cannot set the WLAN ID and lets the WLC assign the WLAN ID.
- WLCs with existing WLANs increment to the next available integer.
- Creating a WLAN using the WLC web interface directly allows the WLAN ID to be chosen.
- WLAN IDs cannot be changed once the WLAN is created.

The following simple example shows the issue:

- WLC A has no WLANs defined.
- WLC B has WLAN “Special-SSID” with WLAN ID 1.

Using Prime Infrastructure to create a new WLAN, “Employee-SSID”, across all WLCs results in it being assigned WLAN ID 1 on WLC A and WLAN ID 2 on WLC B.

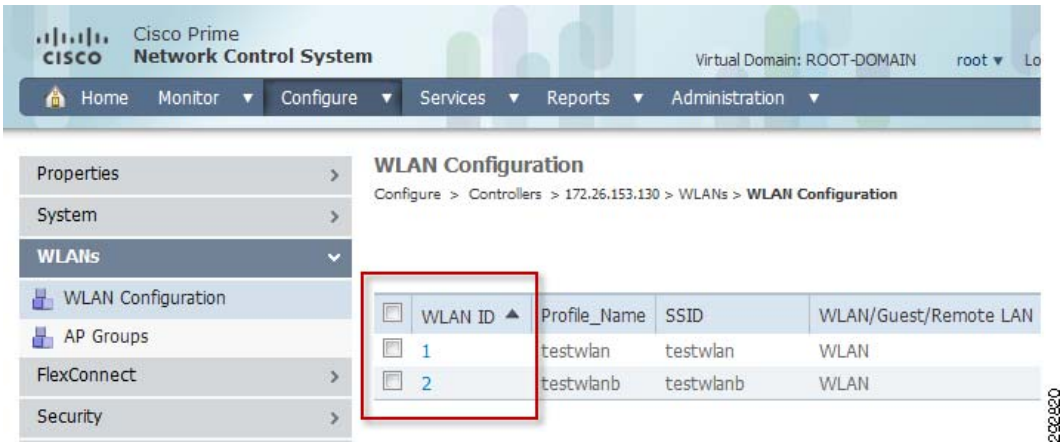
- WLC A
  - WLAN “Employee-SSID” WLAN ID 1
- WLC B
  - WLAN “Special-SSID” WLAN ID 1
  - WLAN “Employee-SSID” WLAN ID 2

To avoid this potentially serious mismatch, it is essential to audit existing WLCs for WLANs and prepare the WLCs for template-based WLAN configuration. Using only Prime Infrastructure and not the WLC interface, the following summarized steps (followed by detailed steps) prevent WLAN ID inconsistencies.

Detailed Steps for Ensuring WLAN ID Consistency

1. Add all WLCs to Prime Infrastructure and synchronize their configurations.
2. Using Prime Infrastructure, look at the WLANs on each WLC to determine the highest WLAN ID in existence, as shown in the example in Figure 27-31. In this example, two WLANs exist on a particular WLC.

Figure 27-31 WLAN ID List



3. Create disabled dummy WLAN templates and apply to WLCs to bring them all up to the same highest WLAN ID.

   The dummy WLAN settings are irrelevant as long as they are created in the "disabled" state. In this example, two dummy WLAN templates must be created and applied to all WLCs with no WLANs.

   As an alternative to creating the dummy WLANs, the existing WLANs may be deleted and re-created with higher WLAN IDs manually set directly on the WLC. Deletion and re-creation is the only method currently available for changing the WLAN ID. Changing the WLAN ID on an existing WLAN is not possible.
4. Create the new WLAN templates for BYOD configurations and apply to all WLCs.
5. Check WLCs to ensure WLAN IDs are consistently assigned across all WLCs.
6. After WLAN templates are applied, dummy WLANs may be deleted if desired.

Note When adding a new WLC to the network, dummy WLAN templates must be applied to them before applying BYOD WLAN templates. Since the WLAN ID is assigned sequentially, BYOD WLAN templates must always be applied in the same order.



PART 6

Appendices



BYOD Converged Access Configurations

Revised: August 7, 2013

Converged Access—Campus

The Converged Access Campus consists of CT5760 as the Mobility Controller (MC) and the Catalyst 3850 as the Mobility Agent (MA).

An example configuration of a CT5760 in a campus design acting as a MC is shown below:

```
aaa new-model
!
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
!
!
!
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 032A4802120A701E1D5D4C
!
aaa session-id common
!
ip device tracking
!
!
qos wireless-default-untrust
captive-portal-bypass
!
dot1x system-auth-control
!
!
!
!
mac access-list extended MAC_ALLOW
 permit any any
spanning-tree mode pvst
spanning-tree extend system-id
!
!

interface Vlan2
 description ### BYOD-Employee Vlan ###
```

```
 ip address 10.231.2.7 255.255.255.0
 load-interval 30
!
interface Vlan3
 description ### BYOD-Provisioning Vlan ###
 ip address 10.231.3.7 255.255.255.0
 load-interval 30
!
interface Vlan47
 description ### Mgmt Vlan ###
 ip address 10.225.47.2 255.255.255.0
 load-interval 30
!
ip http server
ip http authentication local
ip http secure-server
!
ip access-list extended ACL_BLACKHOLE
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15

ip access-list extended ACL_BLACKHOLE_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 permit ip any any

ip access-list extended ACL_Full_Access
 permit ip any any

ip access-list extended ACL_ISE_Remediate
 permit udp any eq bootpc any eq bootps
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 63.128.76.0 0.0.0.255
 permit ip any 23.0.0.0 0.255.255.255
 permit ip any 17.0.0.0 0.255.255.255
 permit ip any 184.0.0.0 0.255.255.255
 permit ip any 8.0.0.0 0.255.255.255
 permit ip any 74.125.0.0 0.0.255.255
 permit ip any 173.194.0.0 0.0.255.255
 permit ip any 206.111.0.0 0.0.255.255
 permit ip any host 10.225.100.10
 permit ip any 173.223.0.0 0.0.255.255
 deny   ip any any
ip access-list extended ACL_ISE_Remediate_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   ip any host 10.230.1.45
 deny   ip any host 10.225.49.15
 deny   ip any host 10.230.1.76
 deny   ip any 63.128.76.0 0.0.0.255
 deny   ip any 23.0.0.0 0.255.255.255
 deny   ip any 17.0.0.0 0.255.255.255
 deny   ip any 184.0.0.0 0.255.255.255
 deny   ip any 8.0.0.0 0.255.255.255
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
 deny   ip any host 10.225.100.10
 deny   ip any 173.223.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Internet_Only
```

```
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 63.128.76.0 0.0.0.255
 permit ip any host 10.225.100.10
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 172.16.0.0 0.15.255.255
 deny   ip any 192.168.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Internet_Redirect
 deny   ip any host 10.230.1.45
 deny   ip any host 10.225.49.15
 deny   ip any host 10.230.1.76
 deny   ip any 63.128.76.0 0.0.0.255
 deny   ip any host 10.225.100.10
 permit ip any 10.0.0.0 0.255.255.255
 permit ip any 10.0.0.0 0.255.255.255
 permit ip any 172.16.0.0 0.15.255.255
 permit ip any 192.168.0.0 0.0.255.255
 deny   ip any any
ip access-list extended ACL_Partial_Access
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 10.230.4.0 0.0.0.255
 permit ip any host 10.230.6.2
 permit ip any host 10.225.100.10
 deny   ip any 10.230.0.0 0.0.255.255
 deny   ip any 10.225.0.0 0.0.255.255
 deny   ip any 10.200.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Provisioning
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
 permit ip any 74.125.0.0 0.0.255.255
 permit ip any 173.194.0.0 0.0.255.255
 permit ip any 206.111.0.0 0.0.255.255
ip access-list extended ACL_Provisioning_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
 permit tcp any any eq www
 permit tcp any any eq 443
ip access-list extended BLACKHOLE_ACL
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
!
radius-server attribute 6 on-for-login-auth
radius-server attribute 8 include-in-access-req
radius-server attribute 25 access-request include
radius-server attribute 31 send nas-port-detail
radius-server dead-criteria time 5 tries 3
radius-server host 10.225.49.15 auth-port 1812 acct-port 1813 key 7 1237161E060E5D56797F71
!
wireless mobility controller peer-group 100
wireless mobility controller peer-group 100 bridge-domain-id 1
wireless mobility controller peer-group 100 member ip 10.203.61.5 public-ip 10.203.61.5
wireless mobility controller peer-group 100 member ip 10.203.71.5 public-ip 10.203.71.5
```

```
wireless mobility controller peer-group 200
wireless mobility controller peer-group 200 bridge-domain-id 1
wireless mobility controller peer-group 200 member ip 10.207.61.5 public-ip 10.207.61.5
wireless mobility controller peer-group 200 member ip 10.207.71.5 public-ip 10.207.71.5
wireless mobility controller peer-group 200 member ip 10.207.81.5 public-ip 10.207.81.5
wireless mobility controller peer-group 300
wireless mobility controller peer-group 300 bridge-domain-id 1
wireless mobility controller peer-group 300 member ip 10.211.61.5 public-ip 10.211.61.5
wireless mobility controller peer-group 300 member ip 10.211.71.5 public-ip 10.211.71.5
wireless mobility group member ip 10.225.50.36 public-ip 10.225.50.36
wireless mobility group member ip 10.225.45.2 public-ip 10.225.45.2
wireless mobility group name byod
wireless management interface Vlan47
wireless client fast-ssid-change
wireless rf-network byod
wireless security dot1x radius call-station-id macaddress
wireless exclusionlist 1CB0.9414.9077 description gregg
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan BYOD-Employee
 nac
 security web-auth parameter-map global
 session-timeout 1800
 no shutdown
wlan BYOD_Guest 2 BYOD_Guest
 aaa-override
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth
 session-timeout 1800
 no shutdown
wlan BYOD_Provisioning 3 BYOD_Provisioning
 aaa-override
 client vlan BYOD-Provisioning
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 session-timeout 1800
 no shutdown
wlan BYOD_Personal_Device 4 BYOD_Personal_Device
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
 security web-auth parameter-map global
 session-timeout 1800
 no shutdown
wlan IT_Devices 5 IT_Devices
 aaa-override
 client vlan BYOD-Employee
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth parameter-map global
 session-timeout 1800
 no shutdown
```

An example configuration of Converged Access Catalyst 3850 in a campus design acting as a Mobility Agent is shown below:

```
aaa new-model
!
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
!
!
!
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 0525150635491F5B4A5142
 auth-type any
!
aaa session-id common
!
ip device tracking
!
!
qos wireless-default-untrust

captive-portal-bypass
!
!
dot1x system-auth-control
!
mac access-list extended MAC_ALLOW
 permit any any
spanning-tree mode pvst
spanning-tree extend system-id
!

vlan 57
 name Employee
!
vlan 58
 name Provisioning
!
vlan 59-60
!
vlan 61
 name Access_Point
!
vlan 777
 name Guest
!
!
class-map match-any non-client-nrt-class
  match non-client-nrt
!
policy-map port_child_policy
 class non-client-nrt-class
    bandwidth remaining ratio 10
!
!

interface GigabitEthernet1/0/1
 switchport access vlan 57
 switchport mode access
```

```
 ip access-group ACL-DEFAULT in
 authentication event fail action next-method
 authentication host-mode multi-auth
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 3
 spanning-tree portfastinterface Vlan61
 ip address 10.207.61.5 255.255.255.0
!
ip http server
ip http authentication local
ip http secure-server
ip http active-session-modules none
!
ip access-list extended ACL-DEFAULT
 permit udp any eq bootpc any eq bootps
 permit udp any any eq domain
 permit icmp any any
 permit udp any any eq tftp
 deny   ip any any log
ip access-list extended ACL_BLACKHOLE
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
ip access-list extended ACL_BLACKHOLE_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 permit ip any any
ip access-list extended ACL_Full_Access
 permit ip any any
ip access-list extended ACL_ISE_Remediate
 permit udp any eq bootpc any eq bootps
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 63.128.76.0 0.0.0.255
 permit ip any 23.0.0.0 0.255.255.255
 permit ip any 17.0.0.0 0.255.255.255
 permit ip any 184.0.0.0 0.255.255.255
 permit ip any 8.0.0.0 0.255.255.255
 permit ip any 74.125.0.0 0.0.255.255
 permit ip any 173.194.0.0 0.0.255.255
 permit ip any 206.111.0.0 0.0.255.255
 permit ip any host 10.225.100.10
 deny   ip any any
ip access-list extended ACL_ISE_Remediate_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   ip any host 10.230.1.45
 deny   ip any host 10.225.49.15
 deny   ip any host 10.230.1.76
 deny   ip any 63.128.76.0 0.0.0.255
 deny   ip any 23.0.0.0 0.255.255.255
 deny   ip any 17.0.0.0 0.255.255.255
 deny   ip any 184.0.0.0 0.255.255.255
 deny   ip any 8.0.0.0 0.255.255.255
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
```

```
 deny   ip any host 10.225.100.10
 permit ip any any
ip access-list extended ACL_Internet_Only
 permit udp any eq bootpc any eq bootps
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 63.128.76.0 0.0.0.255
 permit ip any host 10.225.100.10
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 172.16.0.0 0.15.255.255
 deny   ip any 192.168.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Internet_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   ip any host 10.230.1.45
 deny   ip any host 10.225.49.15
 deny   ip any host 10.230.1.76
 deny   ip any 63.128.76.0 0.0.0.255
 deny   ip any host 10.225.100.10
 permit ip any 10.0.0.0 0.255.255.255
 permit ip any 10.0.0.0 0.255.255.255
 permit ip any 172.16.0.0 0.15.255.255
 permit ip any 192.168.0.0 0.0.255.255
 deny   ip any any
ip access-list extended ACL_Partial_Access
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 10.230.4.0 0.0.0.255
 permit ip any host 10.230.6.2
 permit ip any host 10.225.100.10
 deny   ip any 10.230.0.0 0.0.255.255
 deny   ip any 10.225.0.0 0.0.255.255
 deny   ip any 10.200.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Provisioning
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
 permit ip any 74.125.0.0 0.0.255.255
 permit ip any 173.194.0.0 0.0.255.255
 permit ip any 206.111.0.0 0.0.255.255
ip access-list extended ACL_Provisioning_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
 permit tcp any any eq www
 permit tcp any any eq 443
!
ip radius source-interface Vlan61
logging 10.230.1.83
!
!
radius-server attribute 6 on-for-login-auth
radius-server attribute 8 include-in-access-req
radius-server attribute 25 access-request include
radius-server attribute 31 send nas-port-detail
radius-server dead-criteria time 5 tries 3
radius-server host 10.225.49.15 auth-port 1812 acct-port 1813 key 7 153C1805102F7A767B6760
```

```
!
!
wireless mobility controller ip 10.225.47.2 public-ip 10.225.47.2
wireless management interface Vlan61
wireless client fast-ssid-change
wireless rf-network byod
wireless security dot1x radius call-station-id macaddress
wireless broadcast
wireless multicast
wireless mgmt-via-wireless
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan Employee
 nac
 session-timeout 300
 no shutdown
wlan BYOD_Guest 2 BYOD_Guest
 aaa-override
 client vlan Guest
 mobility anchor 10.225.50.36
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth
 session-timeout 1800
 no shutdown
wlan BYOD_Provisioning 3 BYOD_Provisioning
 aaa-override
 client vlan Provisioning
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 session-timeout 1800
 no shutdown
wlan BYOD_Personal_Device 4 BYOD_Personal_Device
 client vlan Guest
 mobility anchor 10.225.50.36
 session-timeout 1800
 no shutdown
wlan IT_Devices 5 IT_Devices
 aaa-override
 client vlan Employee
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 session-timeout 1800
 no shutdown
```

Converged Access—Branch

An example configuration of a Converged Access Catalyst 3850 in a branch design is shown below. Note that in a branch design, the Catalyst 3850 acts both as a Mobility Controller (MC) and a Mobility Agent (MA) in a single switch.

```
aaa new-model
!
!
aaa authentication login default enable
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting network default start-stop group radius
!
!
!
!
aaa server radius dynamic-author
 client 10.225.49.15 server-key 7 032A4802120A701E1D5D4C
 auth-type any
!
aaa session-id common
switch 1 provision ws-c3850-24p
!
ip device tracking
!
!
qos wireless-default-untrust
captive-portal-bypass
!
!
mac access-list extended MAC_ALLOW
 permit any any
!
vlan 10
 name BYOD-Employee
!
vlan 11
 name BYOD-Provisioning
!
vlan 17
 name AP_Management
!
vlan 18
!
vlan 777
 name BYOD_Guest
!

!
interface GigabitEthernet1/0/1
 switchport access vlan 10
 switchport mode access
 ip access-group ACL-DEFAULT in
 authentication event fail action next-method
 authentication host-mode multi-auth
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 3
 spanning-tree portfast
!
…
….

interface GigabitEthernet1/0/6
!
```

```
!
interface Vlan17
 ip address 10.200.17.5 255.255.255.0
!
ip http server
ip http authentication local
ip http secure-server
!
ip access-list extended ACL-DEFAULT
 permit udp any eq bootpc any eq bootps
 permit udp any any eq domain
 permit icmp any any
 permit udp any any eq tftp
 deny   ip any any
ip access-list extended ACL_BLACKHOLE
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
ip access-list extended ACL_BLACKHOLE_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 permit ip any any
ip access-list extended ACL_Full_Access
 permit ip any any
ip access-list extended ACL_ISE_Remediate
 permit udp any eq bootpc any eq bootps
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 63.128.76.0 0.0.0.255
 permit ip any 23.0.0.0 0.255.255.255
 permit ip any 17.0.0.0 0.255.255.255
 permit ip any 184.0.0.0 0.255.255.255
 permit ip any 8.0.0.0 0.255.255.255
 permit ip any 74.125.0.0 0.0.255.255
 permit ip any 173.194.0.0 0.0.255.255
 permit ip any 206.111.0.0 0.0.255.255
 permit ip any host 10.225.100.10
 permit ip any 173.223.0.0 0.0.255.255
 deny   ip any any
ip access-list extended ACL_ISE_Remediate_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   ip any host 10.230.1.45
 deny   ip any host 10.225.49.15
 deny   ip any host 10.230.1.76
 deny   ip any 63.128.76.0 0.0.0.255
 deny   ip any 23.0.0.0 0.255.255.255
 deny   ip any 17.0.0.0 0.255.255.255
 deny   ip any 184.0.0.0 0.255.255.255
 deny   ip any 8.0.0.0 0.255.255.255
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
 deny   ip any host 10.225.100.10
 deny   ip any 173.223.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Internet_Only
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 63.128.76.0 0.0.0.255
```

```
 permit ip any host 10.225.100.10
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 172.16.0.0 0.15.255.255
 deny   ip any 192.168.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Internet_Redirect
 deny   ip any host 10.230.1.45
 deny   ip any host 10.225.49.15
 deny   ip any host 10.230.1.76
 deny   ip any 63.128.76.0 0.0.0.255
 deny   ip any host 10.225.100.10
 permit ip any 10.0.0.0 0.255.255.255
 permit ip any 10.0.0.0 0.255.255.255
 permit ip any 172.16.0.0 0.15.255.255
 permit ip any 192.168.0.0 0.0.255.255
 deny   ip any any
ip access-list extended ACL_Partial_Access
 permit ip any host 10.230.1.45
 permit ip any host 10.225.49.15
 permit ip any host 10.230.1.76
 permit ip any 10.230.4.0 0.0.0.255
 permit ip any host 10.230.6.2
 permit ip any host 10.225.100.10
 deny   ip any 10.230.0.0 0.0.255.255
 deny   ip any 10.225.0.0 0.0.255.255
 deny   ip any 10.200.0.0 0.0.255.255
 permit ip any any
ip access-list extended ACL_Provisioning
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
 permit ip any 74.125.0.0 0.0.255.255
 permit ip any 173.194.0.0 0.0.255.255
 permit ip any 206.111.0.0 0.0.255.255
ip access-list extended ACL_Provisioning_Redirect
 deny   udp any eq bootpc any eq bootps
 deny   udp any host 10.230.1.45 eq domain
 deny   ip any host 10.225.49.15
 deny   ip any 74.125.0.0 0.0.255.255
 deny   ip any 173.194.0.0 0.0.255.255
 deny   ip any 206.111.0.0 0.0.255.255
 permit tcp any any eq www
 permit tcp any any eq 443
ip access-list extended BLACKHOLE_ACL
 permit udp any eq bootpc any eq bootps
 permit udp any host 10.230.1.45 eq domain
 permit ip any host 10.225.49.15
!
ip radius source-interface Vlan17
!
!
radius-server attribute 6 on-for-login-auth
radius-server attribute 8 include-in-access-req
radius-server attribute 25 access-request include
radius-server attribute 31 send nas-port-detail
radius-server dead-criteria time 5 tries 3
radius-server host 10.225.49.15 auth-port 1812 acct-port 1813 key 7 153C1805102F7A767B6760
!
!

wireless mobility controller
wireless mobility group member ip 10.225.50.36 public-ip 10.225.50.36
wireless mobility group name byod
```

```
wireless management interface Vlan17
wireless client fast-ssid-change
wireless rf-network byod
wireless security dot1x radius call-station-id macaddress
wireless exclusionlist 1CB0.9414.9077 description gregg
wireless broadcast
wireless multicast
wlan BYOD_Employee 1 BYOD_Employee
 aaa-override
 client vlan BYOD-Employee
 nac
 security dot1x authentication-list default
 session-timeout 1800
 no shutdown
wlan BYOD_Guest 2 BYOD_Guest
 aaa-override
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth
 session-timeout 1800
 no shutdown
wlan BYOD_Provisioning 3 BYOD_Provisioning
 aaa-override
 client vlan BYOD-Provisioning
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 session-timeout 1800
 no shutdown
wlan BYOD_Personal_Device 4 BYOD_Personal_Device
 client vlan BYOD_Guest
 mobility anchor 10.225.50.36
 security web-auth parameter-map global
 session-timeout 1800
 no shutdown
wlan IT_Devices 5 IT_Devices
 aaa-override
 client vlan BYOD-Employee
 mac-filtering MAC_ALLOW
 nac
 no security wpa
 no security wpa akm dot1x
 no security wpa wpa2
 no security wpa wpa2 ciphers aes
 security web-auth parameter-map global
 session-timeout 1800
 no shutdown
end
```



References

Revised: August 7, 2013

Cisco BYOD Smart Solution:
http://www.cisco.com/go/byod/

TrustSec Guides:
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html

Flex 7500 Wireless Branch Controller Deployment Guide:
http://www.cisco.com/en/US/products/ps11635/products_tech_note09186a0080b7f141.shtml#override

Wireless BYOD for FlexConnect Deployment Guide:
http://www.cisco.com/en/US/products/ps10315/products_tech_note09186a0080bcb905.shtml

FlexConnect Feature Matrix:
http://www.cisco.com/en/US/products/ps10315/products_tech_note09186a0080b3690b.shtml

Cisco Identity Services Engine Documentation:
http://www.cisco.com/en/US/customer/products/ps11640/products_user_guide_list.html

Cisco 5700 Series Wireless LAN Controllers—Configuration Guides:
www.cisco.com/en/US/customer/products/ps12598/products_installation_and_configuration_guides_list.html

Catalyst 3850 Series Switches Configuration Guides:
http://www.cisco.com/en/US/customer/products/ps12686/products_installation_and_configuration_guides_list.html

Catalyst 4500 Series Switch Software Configuration Guide:
http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/15.1/XE_330SG/configuration/guide/dot1x.html#wp1324657

Catalyst 3750 Switch Software Configuration Guide:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/15.0_2_se/configuration/guide/sw8021x.html#wp1434612

Cisco TrustSec Configuration Guide:
http://www.cisco.com/en/US/customer/docs/ios-xml/ios/sec_usr_cts/configuration/15-sy/sec-cts-15-sy-book.html

Cisco Medianet Campus QoS Design 4.0:
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html

Cisco Borderless Campus 1.0 Cisco Validated Design Guide:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/Borderless_Campus_Network_1.0/Borderless_Campus_1.0_Design_Guide.html

Configuring Certificates:
http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert_cfg.html

Configuring Tunnel Groups, Group Policies, and Users:
http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/vpngrp.html

Apple iPhone OS Enterprise Deployment Guide:
http://manuals.info.apple.com/en_US/Enterprise_Deployment_Guide.pdf

Cisco Enterprise Mobility 4.1 Design Guide:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html

Cisco Jabber clients—product collateral and documentation:
http://www.cisco.com/go/jabber

Cisco Unified Communications System 9.X SRND:
http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/9x/mobilapp.html

Cisco Visual Networking Index:
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html

Cisco Wireless LAN Controller Configuration Guide, Release 7.4:
http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED.html

Cisco Wireless LAN Controller Configuration Guide, Release 7.4—Configuring Multicast Domain Name System:
http://www.cisco.com/en/US/docs/wireless/controller/7.4/configuration/guides/consolidated/b_cg74_CONSOLIDATED_chapter_01011.html#d75540e531a1635

Cisco WLC Bonjour Gateway Deployment Guide:
http://www.cisco.com/en/US/docs/wireless/technology/bonjour/Bonjour_Deployment.html



Software Versions

Revised: August 7, 2013

The following tables highlight the hardware and software components used for validation testing in this design guide. While other software versions provide the required BYOD functionality, the Validated Software Version column indicates the software version used during the validation.

For specific feature support, consult the Cisco Feature Navigator: http://www.cisco.com/go/fn.

Table C-1 Security

Component	Role	Validated Software Version
Microsoft Active Directory	Server	Windows 2008 Server R2
Microsoft Certificate Authority	Server	Windows 2008 Server R2
RSA Authentication Manager	Server	7.1
Cisco Identity Services Engine	Authentication/Policy Server	ISE SW 1.2
Cisco Prime Infrastructure	Management	1.2

Table C-2 Wireless Access

Component	Role	Validated Software Version
Cisco AP3502	Access Point	N/A
Cisco AP3602	Access Point	N/A
Cisco AP2602	Access Point	N/A
Cisco 5508 Wireless Controller	Wireless LAN Controller	AireOS 7.5
Cisco Flex 7500 Wireless Controller	Wireless LAN Controller	AireOS 7.5
Cisco 5760	Wireless LAN Controller	IOS XE 3.2.2
Cisco Catalyst 3850	Campus and Branch Wireless Access	IOS XE 3.2.2
Apple iPad	Wireless Device	iOS 6.1.3
Apple iPhone	Wireless Device	iOS 6.1.3
Asus Transformer Prime 201	Wireless Device	Android 4.0.3

Table C-2 Wireless Access

Component	Role	Validated Software Version
Samsung GT-P7510 Galaxy Tab	Wireless Device	Android 3.2
Samsung Galaxy Tab 8.9 LTE	Wireless Device	Android 3.2
Samsung Galaxy SIII	Wireless Device	Android 4.2.2
Samsung Galaxy Tab2	Wireless Device	Android 3.2

The document Cisco Identity Services Engine Network Component Compatibility, Release 1.2 provides a detailed description of which end user devices are supported (http://www.cisco.com/en/US/docs/security/ise/1.2/compatibility/ise_sdt.html).

Table C-3 Wired Access

Component	Role	Validated Software Version
Catalyst 3850	Campus/Branch	IOS XE 3.2.2
Catalyst 3750-X	Wired Access	IOS 15.0(2)SE
Catalyst 4500E Sup7-E	Wired Access/Distribution	IOS XE 3.3.0SG
Catalyst 6500 VSS 1400 Catalyst 6500-E VSS4T	Campus Distribution/Core	IOS 15.1(1)SY1
Cisco Nexus® 7000 Sup1	Campus Core/Data Center	NX-OS 5.2.7
Apple MacBook Pro	Wired/Wireless Device	OSX 10.7
Windows Laptop	Wired/Wireless Device	Microsoft Windows 7 (64 bit)

Table C-4 Remote Access

Component	Role	Software Version
ASA 5520	VPN Concentrator	8.4(2)
AnyConnect	VPN Client Software	3.0

The three feature sets available with Cisco Catalyst 3850 Series Switches are:

- LAN Base—The LAN Base feature set offers enhanced intelligent services that include comprehensive Layer 2 features with up to 255 VLANs. It does not support either the wireless Mobility Agent (MA) or the Mobility Controller (MC) function.
- IP Base—Enterprise access Layer 3 switching features. The IP Base feature set provides entry-level enterprise services in addition to all LAN Base features with 1000 VLANs. IP Base also includes the support for wireless controller functionality (Mobility Agent (MA) and Mobility Controller (MC) role; additional access point license required for Mobility Controller role), routed access, smart operations, FNF, etc.
- IP Services—The IP Services feature set provides full enterprise services that include advanced Layer 3 features such as EIGRP, OSPF, BGP, PIM, and IPv6 routing such as OSPFv3 and EIGRPv6. Also supports the wireless controller functionality.

All software feature sets support advanced security and MQC-based QoS.

Customers can transparently upgrade the software feature set in the Cisco Catalyst 3850 Series Switches through Cisco IOS Software CLI using the right to use (RTU)-based software upgrade process. Software activation enables the Cisco IOS Software feature sets. Based on the license's type, Cisco IOS Software activates the appropriate feature set. License types can be changed, or upgraded, to activate a different feature set.

An access point license is required for Cisco Catalyst 3850 operating in Mobility Controller (MC) mode. No access point license is required for 3850 operating in Mobility Agent (MA) mode. This functionality is included in the IP Base feature set. Access point licenses can be transferred only between two Catalyst 3850 switches or between Catalyst 3850 switches and CT5760 wireless controllers and vice versa.

The Catalyst 3850 and CT 5760 running IOS XE 3.2.2 software is compatible with CUWN software version 7.3.112, or 7.5 and higher—in order to support mobility interoperability.



BYOD System Release Notes

Revised: August 7, 2013

This appendix contains information regarding known caveats within Cisco product software revisions which could impact the Cisco BYOD solution documented in this design guide. The caveats are organized around product and software releases within this appendix for ease of reference.

IOS XE 3.2.2 SE Software Release—Catalyst 3850 Series Switches and Cisco 5760 Wireless Controllers

- Wireless clients associated with Access Points connected to Catalyst 3850 Series switches may unexpectedly go into an IDLE state. This rare occurrence can happen after a client has been on-boarded and is connected to the employee SSID with full, partial, or Internet only access, during on-boarding (using a single or dual-SSID design), while a client is blacklisted, or during MDM or ISE remediation.

  The network administrator can view a wireless client in this state via the **show wireless client summary** CLI command. An example is shown below in which the two highlighted wireless clients are in an IDLE state.

```
uasl-3850-2# show wireless client summary
Number of Local Clients : 5

MAC Address    AP Name                          WLAN State              Protocol
-------------------------------------------------------------------------------
38aa.3c44.a224 test_ap                          1    UP                 11n(2.4)
789e.d0cd.f8ac test_ap                          1    Idle               11n(2.4)
8832.9b0d.e554 test_ap                          1    WEBAUTH_PEND       11n(2.4)
c860.001a.a77e test_ap                          3    WEBAUTH_PEND       11n(2.4)
f0b4.7953.26b8 test_ap                          1    Idle               11n(2.4)
```

  Once in this state, the wireless client cannot disconnect and re-connect to the wireless network without assistance from a network administrator. The network administrator can clear an IDLE client by issuing the **wireless client mac-address** <*mac-address of the wireless client*> **deauthenticate forced** CLI command. Note that this command does not appear within the existing Cisco IOS 3.2.xSE command reference. For example, in order to clear the wireless client with mac address f0b4.7953.26b8 shown in the example above, the network administrator would issue the following command:

```
uasl-3850-2# wireless client mac-address f0b4.7953.26b8 deauthenticate forced
```

- Blacklisted wireless clients associated with Access Points connected to Catalyst 3850 Series switches may occasionally require manual intervention in order to regain access to network resources when attempting to reinstate the client. This rare occurrence happens when the Catalyst 3850 Series switch is unable to process the CoA request to remove the ACL_BLACKHOLE and ACL_BLACKHOLE_Redirect ACLs applied to a Blacklisted device as part of the Cisco BYOD solution. The end user can toggle WiFi access on the mobile device in order to regain network access. Alternatively, the network administrator can de-authenticate the wireless client via the **wireless client mac-address** <*mac-address of the wireless client*> **deauthenticate** CLI command and then allow the device to normally reconnect to the network. For example, in order to clear a wireless client with mac address b4f0.abce.fa66, the network administrator would issue the following command:

  ```
  uasl-3850-2# wireless client mac-address f0b4.7953.26b8 deauthenticate
  ```

  This caveat is documented as bug ID CSCuh66924 within the Cisco Bug Toolkit.

Additional Observations

The Release Notes for the Catalyst 3850 Series Switch, Cisco IOS XE Release 3.2.xSE contains additional caveats specific to the Catalyst 3850 Series Switch and can be found at: http://www.cisco.com/en/US/docs/switches/lan/catalyst3850/software/release/3.2_0_se/release_notes/OL28114.html#wp224029