

Official Cert Guide

Advance your IT career with hands-on learning

Video Training

Review Exercises

CCNP Security Identity Management

SISE 300-715

AARON WOLAND, CCIE® No. 20113
KATHERINE McNAMARA, CCIE® No. 50931

ciscopress.com

CCNP Security Identity Management

SISE 300-715

Official Cert Guide

Enhance Your Exam Preparation

Save 80% on Premium Edition eBook and Practice Test

The *CCNP Security Identity Management SISE 300-715 Official Cert Guide Premium Edition and Practice Test* provides three eBook files (PDF, EPUB, and MOBI/Kindle) to read on your preferred device and an enhanced edition of the Pearson Test Prep practice test software. You also receive two additional practice exams with links for every question mapped to the PDF eBook.

See the card insert in the back of the book for your Pearson

CCNP Security Identity Management

SISE 300-715

Official Cert Guide

AARON WOLAND, CCIE No. 20113

KATHERINE MCNAMARA, CCIE No. 50931

Cisco Press

CCNP Security Identity Management SISE 300-715 Official Cert Guide

Aaron Woland

Katherine McNamara

Copyright© 2021 Cisco Systems, Inc.

Published by:
Cisco Press

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions.

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ScoutAutomatedPrintCode

Library of Congress Control Number: 2020913602

ISBN-10: 0-13-664294-2

ISBN-13: 978-0-13-664294-7

Warning and Disclaimer

This book is designed to provide information about the Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) exam. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Editor-in-Chief: Mark Taub

Alliances Manager, Cisco Press: Arezou Gol

Director, ITP Product Management: Brett Bartow

Executive Editor: James Manly

Managing Editor: Sandra Schroeder

Development Editor: Christopher A. Cleveland

Project Editor: Mandie Frank

Copy Editor: Kitty Wilson

Technical Editor: Akhil Behl

Editorial Assistant: Cindy Teeters

Designer: Chuti Prasertsith

Composition: codeMantra

Indexer: Erika Millen

Proofreader: Donna Mulder



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: 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. (1110R)

Credits

Figure Number	Attribution/Credit Line
Figure 2-7	Screenshot of ISE LDAP Identity Source Configuration © Microsoft 2020
Figure 2-11	Screenshot of Authy One-Time Password Soft Token App © 2020 Twilio, Inc.
Figure 3-6	Screenshot of Windows Services Control Applet © Microsoft 2020
Figure 3-7	Screenshot of Wired AutoConfig Properties © Microsoft 2020
Figure 3-8	Screenshot of Local-Area Connection Properties © Microsoft 2020
Figure 3-9	Screenshot of Authentication Tab © Microsoft 2020
Figure 3-10	Screenshot of Protected EAP Properties Dialog © Microsoft 2020
Figure 3-11	Screenshot of EAP MSCHAPv2 Properties Dialog © Microsoft 2020
Figure 3-12	Screenshot of EAP MSCHAPv2 Properties Dialog© Microsoft 2020
Figure 3-13	Screenshot of Back to the Protected EAP Properties © Microsoft 2020
Figure 3-14	Screenshot of Back to the Authentication Tab © Microsoft 2020
Figure 3-15	Screenshot of Advanced Settings Tab © Microsoft 2020
Figure 3-16	Screenshot of Authentication Mode Choices © Microsoft 2020
Figure 15-16	Screenshot of Connecting to a CA's Website © Microsoft 2020
Figure 15-17	Screenshot of Downloading a Public Certificate © Microsoft 2020
Figure 21-66	Screenshot of Launching DART from the Start Menu © Microsoft 2020
Figure 23-10	Screenshot of Users Must Be Part of the Security Manager Role in Tenable.SC © 2020 Tenable, Inc.
Figure 23-12	Screenshot of Tenable.SC Repositories © 2020 Tenable, Inc.
Figure 23-13	Screenshot of Tenable.SC Scan Results © 2020 Tenable, Inc.
Figure 23-29	Screenshot of Tenable.SC Scan Results © 2020 Tenable, Inc.

Contents at a Glance

Introduction xxxvi

Part I Authentication, Authorization, and Accounting

Chapter 1 Fundamentals of AAA 2

Chapter 2 Identity Management 18

Chapter 3 Extensible Authentication Protocol (EAP) over LAN: 802.1X 38

Chapter 4 Non-802.1X Authentication 76

Chapter 5 Introduction to Advanced Concepts 92

Part II Cisco Identity Services Engine

Chapter 6 Cisco Identity Services Engine Architecture 104

Chapter 7 A Guided Tour of the Cisco ISE Graphical User Interface (GUI) 122

Chapter 8 Initial Configuration of Cisco ISE 174

Chapter 9 Authentication Policies 206

Chapter 10 Authorization Policies 232

Part III Implementing Secure Network Access

Chapter 11 Implement Wired and Wireless Authentication 258

Chapter 12 Web Authentication 306

Chapter 13 Guest Services 334

Chapter 14 Profiling 402

Part IV Advanced Secure Network Access

Chapter 15 Certificate-Based Authentication 460

Chapter 16 Bring Your Own Device 482

Chapter 17 TrustSec and MACSec 548

Chapter 18 Posture Assessment 626

Part V Safely Deploying in the Enterprise

Chapter 19 Deploying Safely 714

Chapter 20 ISE Scale and High Availability 734

Chapter 21 Troubleshooting Tools 764

Part VI Extending Secure Access Control

Chapter 22 ISE Context Sharing and Remediation 818

Chapter 23 Threat Centric NAC 868

Part VII Device Administration AAA

Chapter 24 Device Administration AAA with ISE 906

Chapter 25 Configuring Device Administration AAA with Cisco IOS 930

Chapter 26 Configuring Device Admin AAA with the Cisco WLC 968

Part VIII Final Preparation

Chapter 27 Final Preparation 988

Part IX Appendixes

Glossary of Key Terms 991

Appendix A Answers to the "Do I Know This Already?" Quizzes and Q&A Sections 1002

Appendix B CCNP Security Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) Exam Updates 1032

Appendix C Sample Switch Configurations 1034

Index 1062

Online Element

Appendix D Study Planner

Reader Services

Register your copy at www.ciscopress.com/title/9780136642947 for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account.* Enter the product ISBN 9780136642947 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Contents

Introduction xxxvi

Part I Authentication, Authorization, and Accounting

Chapter 1 Fundamentals of AAA 2

"Do I Know This Already?" Quiz 3

Foundation Topics 5

Comparing and Selecting AAA Options 5

Device Administration AAA 6

Network Access AAA 7

TACACS+ 7

TACACS+ Authentication Messages 9

TACACS+ Authorization and Accounting Messages 10

RADIUS 12

AV Pairs 15

Change of Authorization (CoA) 16

Comparing RADIUS and TACACS+ 16

Exam Preparation Tasks 16

Review All Key Topics 16

Define Key Terms 17

Q&A 17

Chapter 2 Identity Management 18

"Do I Know This Already?" Quiz 18

Foundation Topics 20

What Is an Identity? 20

Identity Stores 20

Internal Identity Stores 21

Internal User Identities 21

User Identity Groups 22

Endpoint Groups 22

External Identity Stores 23

Active Directory 24

LDAP 25

Multifactor Authentication 26

One-Time Password (OTP) Services 29

Smart Cards 29

Certificate Authorities 30

Has the Digital Certificate Been Signed by a Trusted CA? 31
Has the Certificate Expired? 32
Has the Certificate Been Revoked? 33
Identity Source Sequences 34
Special Identity Sources 35
SAML IdPs 35
Social Login 35
Exam Preparation Tasks 36
Review All Key Topics 36
Define Key Terms 36
Q&A 36

Chapter 3 Extensible Authentication Protocol (EAP) over LAN: 802.1X 38
"Do I Know This Already?" Quiz 38
Foundation Topics 41
Extensible Authentication Protocol 41
EAP over LAN (802.1X) 41
EAP Types 42
Native EAP Types (Non-Tunneled EAP) 43
Tunneled EAP Types 44
All Tunneled EAP Types 47
EAP Authentication Type Identity Store Comparison 49
Network Access Devices 49
Supplicant Options 50
Windows Native Supplicant 50
User Authentication 58
Machine Authentication (Computer Authentication) 58
Cisco AnyConnect NAM Supplicant 59
Client Policy 61
Authentication Policy 62
Networks 62
Network Groups 71
Implementing AnyConnect NAM Profiles 71
EAP Chaining 73
Exam Preparation Topics 73
Review All Key Topics 73
Define Key Terms 74
Q&A 74

Chapter 4 Non-802.1X Authentication 76
"Do I Know This Already?" Quiz 76
Foundation Topics 79
Devices Without a Supplicant 79
MAC Authentication Bypass 80
Web Authentication 83
Local Web Authentication 84
Local Web Authentication with a Centralized Portal 84
Centralized Web Authentication 85
Centralized Web Authentication with Third-Party Network Device Support 87
Remote-Access Connections 88
EasyConnect 89
Exam Preparation Tasks 90
Review All Key Topics 90
Define Key Terms 91
Q&A 91

Chapter 5 Introduction to Advanced Concepts 92
"Do I Know This Already?" Quiz 92
Foundation Topics 95
Change of Authorization 95
Automating MAC Authentication Bypass (MAB) 96
Posture Assessment 99
Mobile Device Management (MDM) 101
Exam Preparation Tasks 102
Review All Key Topics 102
Define Key Terms 102
Q&A 103

Part II Cisco Identity Services Engine

Chapter 6 Cisco Identity Services Engine Architecture 104
"Do I Know This Already?" Quiz 104
Foundation Topics 106
What Is Cisco ISE? 106
Personas 108
Administration Persona 109
Monitoring Persona 109

Policy Services Persona 110
pxGrid Persona 111
Physical or Virtual Appliances 111
ISE Deployment Scenarios 113
Single-Node Deployment 113
Two-Node Deployment 114
Distributed Deployments 116
Exam Preparation Tasks 120
Review All Key Topics 120
Define Key Terms 120
Q&A 120

Chapter 7 A Guided Tour of the Cisco ISE Graphical User Interface (GUI) 122
"Do I Know This Already?" Quiz 123
Foundation Topics 125
Logging in to ISE 125
Initial Login 125
ISE Home Dashboards 132
Administration Portal 137
Global Search for Endpoints 139
Help Menu 140
ISE Setup Wizards 141
Settings Menu 142
Organization of the ISE GUI 142
Context Visibility 143
Operations 143
RADIUS 144
Threat-Centric NAC Live Logs 146
TACACS Live Log 147
Troubleshoot 147
Adaptive Network Control 148
Reports 150
Policy 150
Policy Sets 150
Profiling 152
Posture 152
Client Provisioning 153
Policy Elements 154

Administration 155
System 155
Identity Management 161
Network Resources 163
Device Portal Management 166
pxGrid Services 169
Feed Service 169
Threat Centric NAC 170
Work Centers 170
Types of Policies in ISE 171
Authentication 171
Authorization 171
Profiling 172
Posture 172
Client Provisioning 172
TrustSec 172
Exam Preparation Tasks 173
Review All Key Topics 173
Define Key Term 173
Q&A 173

Chapter 8 Initial Configuration of Cisco ISE 174
"Do I Know This Already?" Quiz 174
Foundation Topics 177
Cisco Identity Services Engine Form Factors 177
Bootstrapping Cisco ISE 177
Where Are Certificates Used with Cisco Identity Services Engine? 181
Self-Signed Certificates 181
CA-Signed Certificates 182
Network Devices 192
Network Device Groups 192
Network Access Devices 192
ISE Identity Stores 194
Local User Identity Groups 194
Local Endpoint Groups 195
Local Users 195
External Identity Stores 196
Active Directory 196

Prerequisites for Joining an Active Directory Domain 196
Joining an Active Directory Domain 197
Certificate Authentication Profile (CAP) 202
Identity Source Sequences 202
Exam Preparation Topics 204
Review All Key Topics 204
Define Key Terms 204
Q&A 205

Chapter 9 Authentication Policies 206
"Do I Know This Already?" Quiz 207
Foundation Topics 209
The Relationship Between Authentication and Authorization 209
Authentication Policy 210
Goal 1: Accept Only Allowed Protocols 210
Goal 2: Select the Correct Identity Store 210
Goal 3: Validate the Identity 211
Goal 4: Pass the Request to the Authorization Policy 211
Understanding Policy Sets 211
Allowed Protocols 213
Understanding Authentication Policies 216
Conditions 217
Identity Store 219
Options 220
Common Authentication Policy Examples 220
Using the Wireless SSID 220
Remote Access VPN 223
Alternative ID Stores Based on EAP Type 224
More on MAB 227
Restore the Authentication Policy 229
Exam Preparation Tasks 230
Review All Key Topics 230
Q&A 230

Chapter 10 Authorization Policies 232
"Do I Know This Already?" Quiz 232
Foundation Topics 235
Authentication Versus Authorization 235

Authorization Policies 235
Goals of Authorization Policies 235
Understanding Authorization Policies 236
Role-Specific Authorization Rules 241
Authorization Policy Example 241
Employee Full Access Rule 241
Internet Only for Smart Devices Rule 243
Employee Limited Access Rule 246
Saving Conditions for Reuse 249
Combining AND with OR Operators 252
Exam Preparation Tasks 256
Review All Key Topics 256
Define Key Terms 256
Q&A 256

Part III Implementing Secure Network Access

Chapter 11 Implement Wired and Wireless Authentication 258
"Do I Know This Already?" Quiz 259
Foundation Topics 261
Authentication Configuration on Wired Switches 261
Global Configuration AAA Commands 261
Global Configuration RADIUS Commands 262
IOS 12.2.x 262
IOS 15.x and IOS XE 263
IOS 12.2.x, 15.x, and IOS XE 264
Global 802.1X Commands 266
Device Tracking in IOS XE 16.x and Later 267
Creating Local Access Control Lists 268
Interface Configuration Settings for All Cisco Switches 269
Configure Interfaces as Switch Ports 269
Configure Flexible Authentication and High Availability 269
Host Mode of the Switch Port 272
Configure Authentication Settings 274
Configure Authentication Timers 275
Apply the Initial ACL to the Port and Enable Authentication 275
Authentication Configuration on WLCs 276
Configure the AAA Servers 276
Add the RADIUS Authentication Servers 277

Add the RADIUS Accounting Servers 278
Configure RADIUS Fallback (High Availability) 279
Configure the Airespace ACLs 280
Create the Web Authentication Redirection ACL 280
Add Google URLs for ACL Bypass 282
Create the Posture Agent Redirection ACL 283
Create the Dynamic Interfaces for the Client VLANs 284
Create the Employee Dynamic Interface 284
Create the Guest Dynamic Interface 285
Create the Wireless LANs 286
Create the Guest WLAN 287
Create the Corporate WLAN 291
Verifying Dot1x and MAB 295
Endpoint Supplicant Verification 295
Network Access Device Verification 296
Verifying Authentications with Cisco Switches 296
Sending Syslog to ISE 299
Verifying Authentications with Cisco WLCs 300
Cisco ISE Verification 302
RADIUS Live Log 302
Live Sessions 303
Looking Forward 303
Exam Preparation Tasks 303
Review All Key Topics 303
Define Key Terms 304
Q&A 304

Chapter 12 Web Authentication 306
"Do I Know This Already?" Quiz 306
Foundation Topics 309
Web Authentication Scenarios 309
Local Web Authentication (LWA) 310
Centralized Web Authentication (CWA) 311
Configuring Centralized Web Authentication 313
Cisco Switch Configuration 313
Configure Certificates on the Switch 313
Enable the Switch HTTP/HTTPS Server 314
Verify the URL-Redirect ACL 314

Cisco WLC Configuration 315
Validate That MAC Filtering Is Enabled on the WLAN 315
Validate That ISE NAC Is Enabled on the WLAN 315
Validate That the URL-Redirection ACL Is Configured 316
Configure ISE for Centralized Web Authentication 317
Configure MAB Continue for the Authentication 318
Verify the Web Authentication Identity Source Sequence 319
Configure a dACL for Pre-WebAuth Authorization 319
Configure an Authorization Profile 320
Building CWA Authorization Policies 322
Create the Rule to Redirect Users to the CWA Portal 323
Create the Rules to Authorize Users Who Authenticate via CWA 323
Verifying Centralized Web Authentication 324
Check the Experience from the Client 324
Verify CWA Through the ISE UI 327
Check Live Log 327
Check the NAD 327
***show** Commands on the Wired Switch 328*
Viewing the Client Details on the WLC 329
Exam Preparation Tasks 331
Review All Key Topics 331
Define Key Terms 331
Q&A 332

Chapter 13 Guest Services 334
"Do I Know This Already?" Quiz 334
Foundation Topics 337
Guest Services Overview 337
Portals, Portals, and More Portals! 341
Guest Portal Types 341
Hotspot Guest Portal 342
Self-Registered Guest Portal 342
Sponsored Guest Portal 342
Guest Types 343
Contractor 344
Daily 346
Weekly 347
Social 348
Guest Portals and Authorization Policy Rules 348

Configuring Guest Portals and Authorization Rules 351
Configuring a Hotspot Guest Portal 351
Portal Behavior and Flow Settings 351
Portal Page Customization 358
Authorization Rule Configuration 362
Configuring a Self-Registered Guest Portal 365
Portal Settings 366
Login Page Settings 367
Registration Form Settings 368
Self-Registration Success 371
Guest Change Password Settings and Guest Device Registration Settings 371
BYOD Settings 372
Guest Device Compliance Settings 373
Authorization Rule Configuration 373
Configuring a Sponsored Guest Portal 380
Sponsors 381
Sponsor Groups 381
Sponsor Portals 384
Portal Settings 385
Login Settings and AUP Page Settings 386
The Remaining Settings 387
Notification Services 388
SMTP Servers 388
SMS Gateway Providers 388
Provisioning Guest Accounts from a Sponsor Portal 389
SAML Authentication 394
Call to Action 400
Exam Preparation Tasks 400
Review All Key Topics 400
Define Key Terms 401
Q&A 401

Chapter 14 Profiling 402
"Do I Know This Already?" Quiz 402
Foundation Topics 404
ISE Profiler 404
Anomalous Behaviour 406
Cisco ISE Probes 409

Probe Configuration 409
DHCP and DHCPSPAN 411
RADIUS 414
Network Scan (Nmap) 415
DNS 417
SNMPQUERY and SNMPTRAP 417
NETFLOW 419
HTTP Probe 420
Active Directory Probe 422
pxGrid Probe 423
Infrastructure Configuration 424
DHCP Helper 424
SPAN Configuration 424
VLAN Access Control Lists (VACLs) 425
Device Sensor 426
VMware Configurations to Allow Promiscuous Mode 427
Profiling Policies 429
Profiling Feed Service 429
Configuring the Profiler Feed Service 429
Verifying the Profiler Feed Service 429
Endpoint Profile Policies 431
Logical Profiles 441
ISE Profiler and CoA 442
Global CoA 442
Per-Profile CoA 443
Global Profiler Settings 444
Configure SNMP Settings for Probes 444
Endpoint Attribute Filtering 444
Custom Attributes for Profiling 445
Publishing Endpoint Probe Data on pxGrid 450
Profiles in Authorization Policies 450
Endpoint Identity Groups 450
EndPointPolicy 453
Verify Profiling 454
The Dashboard 454
Global Search 454
Endpoint Identities 455
Device Sensor show Commands 457

Exam Preparation Topics 458
Review All Key Topics 458
Define Key Terms 458
Q&A 458

Part IV Advanced Secure Network Access

Chapter15 Certificate-Based Authentication 460
"Do I Know This Already?" Quiz 460
Foundation Topics 463
Certificate Authentication Primer 463
Determine If a Trusted Authority Has Signed the Digital Certificate 463
Examine Both the Start and End Dates to Determine If the Certificate Has Expired 465
Verify If the Certificate Has Been Revoked 466
Validate That the Client Has Provided Proof of Possession 468
A Common Misconception About Active Directory 469
EAP-TLS 470
Configuring ISE for Certificate-Based Authentications 470
Validate Allowed Protocols 470
Certificate Authentication Profile 471
Verify the Authentication Policy Is Using the CAP 472
Authorization Policies 474
Ensure the Client Certificates Are Trusted 475
Import the Certificate Authority's Public Certificate 476
Configure Certificate Status Verification (Optional) 478
Exam Preparation Tasks 479
Review All Key Topics 479
Define Key Terms 480
Q&A 480

Chapter 16 Bring Your Own Device 482
"Do I Know This Already?" Quiz 483
Foundation Topics 485
BYOD Challenges 485
Onboarding Process 487
BYOD Onboarding 487
Dual SSID 487
Single SSID 488

Configuring NADs for Onboarding 489
Configuring a WLC for Dual SSID Onboarding 489
Review of the WLAN Configuration 490
Verify the Required ACLs 492
ISE Configuration for Onboarding 495
The End-User Experience 496
Single SSID with Apple iOS Example 496
Dual SSID with Android Example 503
Unsupported Mobile Device: BlackBerry Example 508
Configuring ISE for Onboarding 510
Creating the Native Supplicant Profile 510
Configure the Client Provisioning Policy 512
Configure the WebAuth 514
Verify Default Unavailable Client Provisioning Policy Action 515
Create the Authorization Profiles 516
Create the Authorization Rules for Onboarding 517
Create the Authorization Rules for the EAP-TLS Authentications 518
ISE as a Certificate Authority 519
Configuring SCEP 520
Configuring ISE as an Intermediate CA 521
BYOD Onboarding Process Detailed 523
iOS Onboarding Flow 523
Phase 1: Device Registration 523
Phase 2: Device Enrollment 525
Phase 3: Device Provisioning 526
Android Flow 526
Phase 1: Device Registration 526
Phase 2: NSP App Download App 528
Phase 3: Device Provisioning 529
Windows and macOS Flow 531
Phase 1: Device Registration 531
Phase 2: Device Provisioning 532
Verifying BYOD Flows 534
RADIUS Live Logs 534
Reports 534
Identity Group 535

MDM Onboarding 535
Integration Points 536
Configuring MDM Integration 537
Configuring MDM Onboarding Rules 539
Create the Authorization Profile 539
Create the Authorization Rules 540
Managing Endpoints 542
Self-Management 543
Administrative Management 545
The Opposite of BYOD: Identify Corporate Systems 545
Exam Preparation Topics 546
Review All Key Topics 547
Define Key Terms 547
Q&A 547

Chapter 17 TrustSec and MACsec 548
"Do I Know This Already?" Quiz 548
Foundation Topics 551
Ingress Access Control Challenges 551
VLAN Assignment 551
Ingress Access Control Lists 553
East–West Segmentation 554
What Is TrustSec? 555
What Is a Security Group Tag? 556
What Is the TrustSec Architecture? 557
TrustSec-Enabled Network Access Devices 558
Defining the TrustSec Settings for a Network Access Device 559
Configuring an IOS XE Switch for TrustSec 560
Configuring an ASA for TrustSec 564
Network Device Admission Control (NDAC) 566
Configuring the Seed Device 566
Configuring the Non-Seed Device 567
Defining the SGTs 572
Classification 575
Dynamically Assigning SGT via 802.1X 577
Manually Assigning SGTs to a Port 577
Manually Binding IP Addresses to SGTs in ISE 578
Access-Layer Devices That Do Not Support SGTs 580

Mapping a Subnet to an SGT 580
Mapping a VLAN to an SGT 580
Transport: SGT Exchange Protocol (SXP) 581
SXP Design 582
Configuring SXP on ISE 584
Configuring SXP on IOS Devices 587
Configuring SXP on Wireless LAN Controllers 590
Configuring SXP on Cisco ASA 591
Verifying SXP Connections in ASDM 592
Transport: Native Tagging 593
Configuring Native SGT Propagation (Tagging) 594
Configuring Manual SGT Propagation on Cisco IOS XE Switches 595
Enforcement 597
SGACL 597
Configuring Security Group ACLs 601
TrustSec Policy Matrix 604
Configuring the TrustSec Policy Matrix 605
Security Group Firewalls 611
Security Group Firewall on the ASA 612
Security Group Firewall on the Firepower 612
Security Group Firewall on the ISR and ASR 613
Software-Defined Access (SD-Access) 613
MACsec 614
Downlink MACsec 616
Switch Configuration Modes 618
ISE Configuration 619
Uplink MACsec 619
Manually Configuring Uplink MACsec 620
Verifying the Manual Configuration 622
Exam Preparation Tasks 623
Review All Key Topics 623
Define Key Terms 623
Q&A 624

Chapter 18 Posture Assessment 626
"Do I Know This Already?" Quiz 626
Foundation Topics 629

Posture Assessment with ISE 629
A Bit of a History Lesson 629
ISE Posture Flows 633
Configuring Posture 636
Update the Compliance Modules 637
Configure Client Provisioning 638
Protect Your Sanity 638
Download AnyConnect 640
Upload AnyConnect Headend Deployment Packages to ISE 642
Configure the Client Provisioning Portal 650
Configure the Client Provisioning Policy 652
Configuring Posture Policy Elements 653
Conditions 654
Remediations 679
Requirements 687
Configure Posture Policies 688
Other Important Posture Settings 690
Posture Lease 691
Cache Last Known Posture Compliant Status 691
Reassessment Configurations 691
Authorization Rules 693
Create an Authorization Profile for Redirection 693
Create the Authorization Rules 694
The Endpoint Experience 695
Scenario 1: AnyConnect Not Installed on Endpoint Yet 696
Scenario 2: AnyConnect Already Installed, Endpoint Not Compliant 700
Scenario 3: Stealth Mode 703
Scenario 4: Temporal Agent and Posture Compliant 705
Mobile Posture 707
Create Mobile Posture Authorization Conditions 709
Create Mobile Posture Authorization Rules 710
Exam Preparation Tasks 713
Review All Key Topics 713
Define Key Terms 713
Q&A 713

Part V Safely Deploying in the Enterprise

Chapter 19 Deploying Safely 714

"Do I Know This Already?" Quiz 714
Foundation Topics 717
Why Use a Phased Approach? 717
Comparing authentication open to Standard 802.1X 719
Prepare ISE for a Staged Deployment 720
Monitor Mode 722
Low-Impact Mode 725
Closed Mode 728
Transitioning from Monitor Mode to Your End State 730
Wireless Networks 731
Exam Preparation Tasks 731
Review All Key Topics 731
Q&A 732

Chapter 20 ISE Scale and High Availability 734

"Do I Know This Already?" Quiz 734
Foundation Topics 737
Configuring ISE Nodes in a Distributed Environment 737
Make the First Node a Primary Device 738
Registering an ISE Node to the Deployment 739
Ensure That the Persona of Each Node Is Accurate 742
Understanding the High Availability Options Available 743
Primary and Secondary Nodes 743
Monitoring and Troubleshooting Nodes 743
Policy Administration Nodes 745
Promoting the Secondary PAN to Primary 745
Auto PAN Switchover 745
Configuring Automatic Failover for the Primary PAN 746
Licensing in a Multi-Node ISE Cube 747
Node Groups 748
Add the Policy Services Nodes to the Node Group 750
Using Load Balancers 751
General Guidelines 752
Failure Scenarios 753
Anycast High Availability for ISE PSNs 753
IOS Load Balancing 756

Maintaining ISE Deployments 757
Patching ISE 757
Backup and Restore 759
Exam Preparation Tasks 761
Review All Key Topics 761
Define Key Term 761
Q&A 762

Chapter 21 Troubleshooting Tools 764

"Do I Know This Already?" Quiz 764
Foundation Topics 766
Logging 766
Live Log 766
Advanced Filtering 771
Authentication Details Report 771
The Blank Lines 774
Live Sessions 776
Logging and Remote Logging 777
Logging Targets 777
Logging Categories 778
Debug Logs 779
Downloading Debug Logs from the GUI 780
Viewing Log Files from the CLI 781
Support Bundles 782
Diagnostic Tools 785
RADIUS Authentication Troubleshooting Tool 785
Execute Network Device Command 787
Evaluate Configuration Validator 788
Posture Troubleshooting 794
Endpoint Debug 796
TCP Dump 798
Session Trace Tests 801
Troubleshooting Methodology 804
Log De-duplication 805
The USERNAME User 807
Troubleshooting Outside of ISE 808
Endpoint Diagnostics 809
Cisco AnyConnect Diagnostics and Reporting Tool (DART) 809

Supplicant Provisioning Logs 812
Network Device Troubleshooting 812
Show Authentication Session Interface 812
Viewing Client Details on the WLC 813
Debug Commands 815
Exam Preparation Tasks 815
Review All Key Topics 815
Q&A 816

Part VI Extending Secure Access Control

Chapter 22 ISE Context Sharing and Remediation 818
"Do I Know This Already?" Quiz 818
Foundation Topics 820
Integration Types in the ISE Ecosystem 820
MDM Integration 820
Rapid Threat Containment 821
Platform Exchange Grid 824
pxGrid 825
pxGrid in Action 826
Context-In 827
Configuring ISE for pxGrid 828
Configuring pxGrid Participants 831
Configuring Firepower Management Center for Identity with pxGrid 831
Configuring the Web Security Appliance 850
Integrating Stealthwatch and ISE 857
Exam Preparation Tasks 867
Review All Key Topics 867
Define Key Terms 867
Q&A 867

Chapter 23 Threat Centric NAC 868
"Do I Know This Already?" Quiz 868
Foundation Topics 871
Vulnerabilities and Threats, Oh My! 871
Integrating Vulnerability Assessment Sources 872
TC-NAC Flows 873
Enable TC-NAC 874

Configure the Integration with a Vulnerability Assessment Vendor 878
Authorization Profile and Authorization Rules 884
Seeing TC-NAC with Vulnerability Scanners in Action 887
Verifying What Happened 888
Integrating with Threat Sources 890
Cognitive Threat Analytics (CTA) 890
Create a CTA STIX/TAXII API Account 892
Create a CTA Integration for TC-NAC 894
Using CTA with Authorization 896
AMP for Endpoints 897
Normalized Events 899
Configuring the AMP Adapter 900
Exam Preparation Tasks 904
Review All Key Topics 904
Define Key Terms 905
Q&A 905

Part VII Device Administration AAA

Chapter 24 Device Administration AAA with ISE 906
"Do I Know This Already?" Quiz 906
Foundation Topics 909
Device Administration AAA Refresher 909
Device Administration in ISE 910
Device Administration Design 911
Large Deployments 912
Medium Deployments 913
Small Deployments 913
Enabling TACACS+ in ISE 914
Network Devices 916
Device Administration Global Settings 917
Connection Settings 918
Password Change Control 918
Session Key Assignment 918
Device Administration Work Center 919
Identities 920
Network Resources 921
Policy Elements 922

TACACS Command Sets 922
TACACS Profiles 923
Policy Sets 925
Reports 927
Exam Preparation Tasks 928
Review All Key Topics 928
Q&A 928

Chapter 25 Configuring Device Administration AAA with Cisco IOS 930
"Do I Know This Already?" Quiz 930
Foundation Topics 932
Overview of IOS Device Administration AAA 932
TACACS Profile 932
TACACS+ Command Sets 934
Configure ISE and an IOS Device for Device Administration AAA 936
Prepare ISE for IOS Device Administration AAA 937
Ensure That the Device Administration Service Is Enabled 937
Prepare the Network Device 937
Prepare the Policy 939
Configure the TACACS Profiles 939
Configure the TACACS Command Sets 941
Configure the Policy Set 943
IOS Configuration for TACACS+ 946
Configure TACACS+ Authentication and Fallback 946
Configure TACACS+ Command Authorization 948
Configure TACACS+ Command Accountings 951
Testing and Troubleshooting 951
Testing and Troubleshooting in ISE 952
Troubleshooting at the IOS Command Line 954
Exam Preparation Tasks 966
Review All Key Topics 966
Define Key Terms 967
Q&A 967

Chapter 26 Configuring Device Admin AAA with the Cisco WLC 968
"Do I Know This Already?" Quiz 968
Foundation Topics 971
Overview of WLC Device Administration AAA 971

Configure ISE and the WLC for Device Administration AAA 972
Prepare ISE for WLC Device Administration AAA 972
Prepare the Network Device 972
Prepare the Policy Results 974
Configure the Policy Set 977
Adding ISE to the WLC TACACS+ Servers 979
Testing and Troubleshooting 981
Exam Preparation Tasks 986
Review All Key Topics 986
Q&A 987

Part VIII Final Preparation

Chapter 27 Final Preparation 988
Hands-on Activities 988
Suggested Plan for Final Review and Study 988
Summary 989

Part IX Appendixes

Glossary of Key Terms 991

Appendix A Answers to the "Do I Know This Already?" Quizzes and Q&A Sections 1002

Appendix B CCNP Security Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) Exam Updates 1032

Appendix C Sample Switch Configurations 1034

Index 1062

Online Element

Appendix D Study Planner

About the Authors

Aaron Woland, CCIE No. 20113, is a Principal Engineer in Cisco's Advanced Threat Security & Integrations group and works with Cisco's Largest Customers all over the world. His primary job responsibilities include security design, solution enhancements, standards development, advanced threat solution design, endpoint security, and futures.

Aaron joined Cisco in 2005 and is currently a member of numerous security advisory boards and standards body working groups. Prior to joining Cisco, Aaron spent 12 years as a Consultant and Technical Trainer.

Aaron's other publications include *Integrated Security Technologies and Solutions*, Volume I; both Volumes I and II of the *Cisco ISE for BYOD and Secure Unified Access* book; the *All-in-one Cisco ASA Firepower Services, NGIPS and AMP* book; the *CCNP Security SISAS 300-208 Official Cert Guide*; the *CCNA Security 210-260 Complete Video Course*; and many published white papers and design guides.

Aaron is one of only five inaugural members of the Hall of Fame Elite for Distinguished Speakers at Cisco Live and is a security columnist for Network World where he blogs on all things related to security. His other certifications include GHIC, GCFE, GSEC, Certified Ethical Hacker, MCSE, VCP, CCSP, CCNP, CCDP, and many other industry certifications.

You can follow Aaron on Twitter: @aaronwoland.

Katherine McNamara, CCIE No. 50931, is a Cybersecurity Technical Solutions Architect at Cisco Systems and has worked with large enterprise and public sector customers.

Katherine joined Cisco in 2014 and has worked in IT since 2007 in multiple networking and security roles. She graduated with a Bachelor of Science in IT Security and a Master of Science in Information Security and Assurance. Her many certifications include CCIE Data Center, CCIE Security, MCSE, VCP, CISSP, CCNP, CCDP, and more.

Outside of her day job, she runs a blog called network-node.com, which provides training articles and videos about Cisco Security products. She also helps co-organize the largest Cisco study Meetup group in the world named Routergods.

You can follow Katherine on Twitter: @kmcnam1

About the Technical Reviewer

Akhil Behl, CCIE No. 19564, is a passionate IT executive with a key focus on cloud and security. He has more than 16 years of experience in the IT industry working in several leadership, advisory, consultancy, and business development profiles with various organizations. His technology and business specialization includes cloud, security, infrastructure, data center, and business communication technologies.

Akhil is a published author. Over the span of the past few years, Akhil authored multiple titles on security and business communication technologies. He has contributed as technical editor for over a dozen books on security, networking, and information technology. He has published several research papers in national and international journals, including *IEEE Xplore*, and presented at various IEEE conferences, as well as other prominent ICT, security, and telecom events. Writing and mentoring are his passion and a part of his life.

He holds CCIE (Collaboration and Security), CCSK, CHFI, PMP, ITIL, VCP, TOGAF, CEH, ISM, CCDP, and many other industry certifications. He has Bachelor in Technology and Masters in Business Administration degrees.

Dedications

From Aaron and Katherine: This book was written largely during the rise of the COVID-19 worldwide pandemic. We spent some of this time not knowing if there was going to be a world to live in, much less if there was going to be any readers to learn from this material. So, this book is dedicated to all the people of the world who did not survive the ordeal, to our friends, family, and colleagues who did contract the virus and pulled through, and the many of us who survived the trials and tribulations of the quarantines.

From Aaron: This book is dedicated to my amazing best friend, fellow adventurer, and wife, Suzanne. Thank you for your continued support, encouragement, and patience and for putting up with all the long nights I had to be writing while you let the twins have a "sleepover" in our room, and let me sleep in to make up for it and for always believing in me and supporting me. You are beyond amazing.

To Mom and Pop. You have always believed in me and supported me in absolutely everything I've ever pursued and showed pride in my accomplishments (no matter how small). I hope I can continue to fill your lives with pride, happiness, "nachas"; and if I succeed in my endeavors to make you proud, it will still only be a fraction of what you deserve.

To my four absolutely incredible daughters, Eden, Nyah, Netanya, and Cassandra: You girls are my inspiration, pride, and joy! I can only hope that one day you will look back at the ridiculous man that raised you and feel a level of pride.

—Aaron

From Katherine: This book is dedicated to my wonderful, amazing, and supportive wife, Dianne. This book is as much yours as it is mine. I could not have done it without the support and love you gave me through it all. Thank you for being so amazing every day. I know life with me has been a crazy rollercoaster of adventures but there is truly no one else I'd rather have by my side though all of life's surprises. The luckiest day of my life was the day I saw you walking toward me in Colorado all those years ago. You're my foundation, rock, muse, and soulmate. I love you forever. Also, Brock is your cat.

I also want to dedicate this book to my parents—Robert and Evelyne McNamara—and my siblings—Ryan, Jack, Cai, Mimi, Grace, Molly, and Little Bob. I love you all, and I can't wait until the time we can see each other again in person.

To Rakel—It's been a crazy few years. You have had the monumental job of putting up with me for which I am thankful for all you do. Thank you for grounding me so many times when I needed it and always being a listening ear (or text message). As far as I'm concerned, you are a part of my tribe. Thanks for being my Jiminy Cricket. I wish you all the lavender, crackpie, stressballs, monster salads, hot yoga, travel, love, and happiness in the world.

To Gordon—Who will never see this. You wanted me to be a better person and gave me another chance at life. I hope if you were still here, you would be proud of me.

To my Secret Ninja Pirates—Dustin Schuemann, Tim McConnaughy, Matthew McGee, Joel Sprague, Hieu Phan, Renee Kostreva, Steven McNutt, Anthony Sytnik, Brad Johnson,

Joshua Burget, JP Cedeno, and even David Gaytan—thank you for the many laughs and vent sessions.

I want to thank Bill Boyles Jr., Amanda Boyles, Jessica Rojas, Lily Speerbrecker, Marshall DuVal, Simone Hirsekorn, Jenna Bulis, Chelsea Filer, James Ziegenbalg, and Amanda La Ford. You guys are my real-life superheroes and inspiration to push myself in so many ways.

I want to give thanks to my very awesome SWS team who has made me feel so welcomed in this amazing team/family. Thank you, Isabella Yani, Oli Laurent, Greg Evans, Blake Fletcher, Cristin Beckendorf, Dani Hemmings, Glen Oltmanns, James Nolan, Jeff Hubbell, Jeremy Stephens, Jerod Atkins, Tammy McKeever, Patrick Taylor, Matt Rhebeck, Shane Hanner, Dan Turner, and Thomas Archuleta.

To Miguel de Zubeldia: You're an amazing work partner. Thanks to you I can never look at spicy tacos the same again. Thank you for the all the amazing work, laughs, and accepting me as the other half of this pant-less dynamic duo.

Thank you to everyone in Routergods and in Cisco who supported me along the way—especially Humphrey Cheung, Dmitry Figol, James Schallau, Nicole Wajer, Lindsay Simancek, Nick Russo, Francois Caen, Alex Shkolnik, Brad Edgeworth, Carolina Terrazas, John Behen, David Peñaloza, Joe Astorino, Kyle Winters, Joey Muniz, Moses Frost, Russell Pope, Jeff Denton, and so many more.

To my amazing professional mentors—my amazing co-author Aaron Woland who honored me so much by choosing me as a co-author. Jeff Fanelli who always gives me amazing advice and encouragement. You make me feel like I can do anything I set my mind to. Thank you for taking a chance with me at Cisco Live and I'll always strive to never let you down. Denise Fishburne who has been an awesome friend in the last few years. You always offer great and wise advice and I am always thankful for your insight. Willow Young—your bravery and wit always inspire me. Jeff Moncrief who should never shave his beard off and who vouched for me to make the transition to a security specialist. You're my favorite furry, and I'm honored to call you friend. Neno Spasov, who got me interested in ISE in the first place, has always been patient with all my questions and literally was in the trenches with me. I love you, big guy.

Acknowledgments

From Aaron: to my co-author, Katherine. Since you joined me at Cisco, you have reinvigorated my passion for ISE and security. I truly treasure our professional and personal relationships and I especially love seeing Dianne's and your fun interactions online. Thank you for agreeing to do the lion's share of the work for this book; it's been a true honor collaborating with you on this book and it is also an honor to call you a colleague and even more so to call you a friend.

To Chris Cleveland, you and I have worked on so many projects together at this point, it is like we know each other on a personal level, even though we've never met in person. You are an amazing editor and Pearson is truly blessed to have you! As soon as we can meet in person, my friend, the drinks are on me!

From Aaron and Katherine:

To editors Chris Cleveland, Mandie Frank, and Kitty Wilson: You are amazing. I hope the readers appreciate how much you all have done to make this book what it is today: correcting all the mistakes we made and keeping us aligned with the Cisco Press requirements for style and content. It must have been like herding cats!

To the technical editor, Akhil Behl, you have amazing insights into the security industry, and we look forward to reading your doctoral thesis someday.

To our employer, Cisco: You really are the greatest place to work in the world. Not only are we both passionate for your technology, but also we are patinate about the place we call home. Chuck's leadership in the community and in the company is second to none.

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([ ]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ([{ }]) indicate a required choice within an optional element.

Introduction

The Cisco Certified Network Professional (CCNP) certification program has several technology tracks including Enterprise, Security, Data Center, Service Provider, and, last but not least, Collaboration. This book will focus on one of the optional concentration exams to achieve your CCNP Security certification – Implementing and Configuring Cisco Identity Services Engine (SISE 300-715).

You may already have other Cisco certifications in other networking technologies or this may be your first foray into the Cisco certification process. You may instead be reading this book to enrich your skillset for your job and not even take the exam. Whichever the case, you have chosen a great resource to further your learning and we wish you the best of luck in your studies.

CCNP Security Certification Overview

Security is an ever evolving and growing networking technology—a technology that will likely be needed for generations to come. As the protocols, applications, and user bas that communicate over a network change and evolve, so must the security approach that is implemented. Network security requires a holistic approach whereby a single chink in the security armor can equal a significant compromise of intellectual property and may result in costly network downtime.

The CCNP Security certification track provides a solid basis in core Cisco security technologies and optional concentration exams that focus on operating a variety of security technologies and concepts – Email Security Appliance, Next-Generation Firewall/IPS, Web Security, Virtual Private Networks (VPN), Identity Services Engine, and automation for Cisco Security Solutions. As highlighted above, the focus of this book will be on the implementation and configuration of Identity Services Engine (Cisco Certification 300-715 SISE). Table I-1 lists the optional concentration exams one may take in addition to the 300-701 SCOR exam in order to receive the CCNP Security Certification.

Table I-1 CCNP Security Concentration Exams

Concentration Exam(s)	Recommended Training
300-715 SNCF	Securing Networks with Cisco Firepower Next-Generation Firewall (SSNGFW) Securing Networks with Cisco Firepower Next-Generation IPS (SSFIPS)
300-720 SESA	Securing Email with Cisco Email Security Appliance (SESA)
300-715 SISE	Implementing and Configuring Cisco Identity Services Engine (SISE)
300-725 SWSA	Securing the Web with Cisco Web Security Appliance (SWSA)
300-730 SVPN	Implementing Secure Solutions with Virtual Private Networks (SVPN)
300-735 SAUTO	Implementing Automation for Cisco Security Solutions (SAUI)

By educating yourself in these areas of the Cisco security solutions portfolio, you will be well equipped to implement a well-rounded security infrastructure onto your network.

Contents of the CCNP Security SISE Exam

In order to study effectively for an exam, it is important to know what is actually going to be on the exam. Cisco fully understands this need and provides a "blueprint" for each of its certification exams. These blueprints give a high-level overview as to what is going to be covered on the exam. By diving deeper into each of these blueprint topics, you will become better prepared for your certification exam.

To view the blueprints for the complete CCNP exam certification tracks, you can browse to http://www.cisco.com/go/ccnp. This webpage contains links to each of the CCNP certification tracks – including the CCNP Security track. The link to go directly to the CCNP Security certification track is https://www.cisco.com/c/en/us/training-events/training-certifications/certifications/professional/ccnp-security-v2.html.

To drill down specifically to the SISE exam blueprint, click the link under the "Concentration exams (choose one)" corresponding to the SISE exam. On this page, you will find several links that provide a high-level description of the SISE exam, Exam Topics, Exam Policies, and sample exam questions. As you review the blueprint (under Exam Topics) and other content pertaining to the SISE exam, you may find that some topics overlap with other Cisco certifications – namely, the CCNA. You may choose to enhance your studies by reviewing some of the topics covered in these other exams to refresh your core knowledge.

The topics contained on the CCNP Security SISE exam are provided in Table I-2.

Table I-2 CCNP Security SISE Exam (300-715) Topics

Certification Guide Chapter	Exam Domain/Topic
	Architecture and Deployment (10% of exam)
8, 18, 20–21, 22–24	**Configure Personas**
6, 20, 24	**Describe Deployment Options**
	Policy Enforcement (25% of exam)
8	**Configure native AD and LDAP**
2, 13	**Describe identity store options**
11	**Configure wired/wireless 802.1x network access**
19	**Configure 802.1x phasing deployment**
11	**Configure network access devices**
11	**Implement MAB**
17	**Configure Cisco TrustSec**

Certification Guide Chapter	Exam Domain/Topic
9–18, 23	**Configure policies including authentication and authorization profiles**
	Web Auth and Guest Services (15% of exam)
12	**Configure web authentication**
13	**Configure guest access services**
13	**Configure sponsor and guest portals**
	Profiler (15% of exam)
14	**Implement profiler services**
14	**Implement probes**
14	**Implement CoA**
14–16, 18	**Configure endpoint identity management**
	BYOD (15% of exam)
16	**Describe Cisco BYOD functionality**
16	**Configure BYOD device onboarding using internal CA with Cisco switches and Cisco wireless LAN controllers**
15–16	**Configure certificates for BYOD**
16	**Configure blacklist/whitelist**
	Endpoint Compliance (10% of exam)
18	**Describe endpoint compliance, posture services, and client provisioning**
18	**Configure posture conditions and policy and client provisioning**
18	**Configure the compliance module**
18	**Configure Cisco ISE posture agents and operational modes**
1	**Describe supplicant, supplicant options, authenticator, and server**
	Network Access Device Management (10% of exam)
1, 24	**Compare AAA protocols**
24–26	**Configure TACACS+ device administration and command authorization**

Besides the training resources provided on the SISE exam page, you might also find additional study resources at the links provided in Table I-3. Other unofficial texts, video, and online training resources can also be found via your favorite online search engine.

Table 1-3 Additional Training Resources

Resource	URL
The Cisco Learning Network SISE page	https://www.cisco.com/c/en/us/training-events/training-certifications/exams/current-list/sise-300-715.html
Cisco Support Forums	https://supportforums.cisco.com
Cisco Press	http://www.ciscopress.com
Cisco ISE for BYOD and Secure Unified Access	http://www.ciscopress.com/store/cisco-ise-for-byod-and-secure-unified-access-9781587144738
BYOD Networking LiveLessons	http://www.ciscopress.com/store/cisco-bring-your-own-device-byod-networking-livelessons-9781587144219

How to Take the SISE Exam

To take the CCNP Security SISE Exam, browse to https://www.cisco.com/go/ccnp. Click on the link for the CCNP Security certification and then the link for SISE. You will find information about the exam including the languages in which the exam will be offered, duration of the exam, as well as a link to register for the exam. At the time of publication of this book, the only approved testing vendor for the SISE exam is Pearson VUE (www.vue.com). To register, click on the Pearson VUE link, create an account, and register for the 300-715 SISE exam. You will then be allowed to select a time and testing center that is most convenient to you.

Who Should Take This Exam and Read This Book?

The SISE 300-715 Exam is just one piece of the CCNP Security certification track. For this reason, the primary audience for this book is those people who are working toward the CCNP Security certification. Furthermore, this book can either be used as the totality of the study material or supplement other study resources (other texts, videos, instructor-led training, online training). Whether you are participating in formalized training for the SISE exam or studying on your own, this text is for you.

Those who take the CCNP Security certification or other CCNP exams are often those individuals who require this level of expertise in their job or their intended career path. Sometimes, the CCNP-level exams are the pinnacle of an individual's intended training—once their CCNP certification is achieved, the recipient chooses to not pursue additional certifications. Other times, the CCNP exams are used as a stepping-stone to higher certifications. In this latter case, the next step in the certification progression is to take the CCIE in the relevant discipline. If the CCNA is the bachelor's degree equivalent of the certification hierarchy and the specialist certifications are a minor in a particular discipline, the CCNP of that discipline is a master's degree. If we were to continue this analogy, the CCIE would be the PhD of the specific technology. See Table I-4 for a comparison chart.

Table I-4 Security Certification Comparison Chart

Certification Name	Years of Experience	Job Role	Product/Technology	Number of Exams	Prerequisites
CCNA	1–3	Network Specialist, Network Support Engineer	Network fundamentals, Network access, IP connectivity, IP services, Security fundamentals, Automation and programmability	1	None
CCNP Security	3–5	Network Security Engineer	Network security, Cloud security, Endpoint protection and detection, Security network access, Visibility and enforcement, Next-Generation Firewall/IPS (Optional), Identity Services Engine (Optional), Email Security Appliance (Optional), Web Security Appliance (Optional), VPN (Optional), C	2	None
CCIE Security	7+	Network Security Engineer	Cisco Adaptive Security Appliance (ASA) Firewall, Firepower Threat Defense (FTD), Firepower Management Center (FMC), Cisco IOS Security, Virtual Private Networks (VPN), LAN Security, Identity Services Engine (ISE), Web Security Appliance (WSA), Email Security Appliance (ESA), AnyConnect, Advanced Malware Protection (AMP), Umbrella, Cognitive Threat Analytics (CTA), Cisco Threat Response (CTR), Security automation and programmability	2	None

Certification Name	Years of Experience	Job Role	Product/Technology	Number of Exams	Prerequisites
Security Specialists					
Cisco Certified Specialist – Network Security Firepower	1–3	Security Administrators, Security Consultants, Network Administrators	Cisco Firepower Threat Defense and Firepower 7000 and 8000 Series virtual appliances	1	None
Cisco Certified Specialist – Security Identity Management Implementation	1–3	Network Security Engineers, ISE Administrators, Wireless Network Security Engineers	Identity Services Engine (ISE)	1	None
Cisco Certified Specialist – Email Content Security	1–3	Enterprise Messaging Managers, Email System Designers, System Administrators	Email Security Appliance (ESA)	1	None
Cisco Certified Specialist – Web Content Security	1–3	Security Architects, System Designers, Network Administrators, Operations Engineers, Security Technicians	Web Security Appliance (WSA)	1	None
Cisco Certified Specialist – Network Security VPN Implementation	1–3	Network Security Engineers, Network Engineers, Network Administrators	Virtual Private Networks (VPN) including GETVPN, DMVPN, FlexVPN, Site-to-Site VPN, Clientless VPN, Remote Access VPN	1	None
Cisco Certified DevNet Specialist – Security Automation and Programmability	1–3	Security Engineers, DevOps Engineers, Network Engineers	Programming concepts, RESTful APIs, Data Models	1	None

Format of the CCNP Security SISE Exam

If you have taken other Cisco Certification Exams, this exam format will not be much different. After registering for the SISE exam, you will have a date and location where you will take your exam. It is recommended that you arrive at the testing center 15–20 minutes ahead of your testing schedule. You will then be asked to present two forms of personal identification—a government-issued picture ID and a second that has at least your signature. You will then be asked to put all of your personal effects into a locker or other secure area as you walk into the testing room. As all Cisco Certification Exams are "closed book," you will not be allowed to take any study materials into the exam room.

The testing room contains a number of testing PCs—often isolated in their own cubicle to encourage privacy and to minimize any interruptions between those who are taking exams. Your testing proctor will escort you into the testing room. You will be provided earplugs and two sheets of writing material (front and back of each sheet is usually available). Oftentimes, these are laminated sheets with a white-erase marker and eraser—allowing you to reuse the sheets as often as you require during your exam. Further details about your testing experience will be provided at the base of the Confirmation Letter as you schedule your exam.

When you start your exam, you will be given the option of taking a sample quiz. This sample quiz will allow you to become familiar with the exam's format. If you are familiar with using a computer, the sample quiz test engine, and that of the actual exam, will likely be easy to navigate.

The CCNP-level exams follow the same format and construction as the CCNA and include the following question types:

- Multiple-Choice
  - Single-Answer
  - Multiple-Answer
- Drag and Drop
- Fill-in-the-Blank
- Testlet
- Simlet
- Simulated Lab

With the multiple-choice questions, these can take on one of two formats—single-answer and multiple-answer. With the single-answer, multiple-choice questions, you will be given a question with several options for the correct answer. You will be asked to select only one of these options using a round radio button to the left of the chosen answer—pointing your mouse icon at the radio button and left-clicking the mouse. For the multiple-answer, multiple-choice questions, you will still be given a question with several options for the correct answer. However, you will usually be asked to select a prescribed number of correct answers—for instance, "Choose 3." These will be selected

using a square radio button to the left of the chosen answers. If you attempt to choose too many answers, you will be prompted to choose only the prescribed amount.

Drag and Drop questions will test your ability to match or put into order a number of words/concepts. You will select one option by left-clicking the option and then, while still maintaining the left-click, move the option to another part of the screen. Often, you will be matching an option from one side of the screen to a related option on the other side of the screen. At times, there may be more "answers" on the left than there are slots to fill on the right. In this case, you have to narrow down your choices to those answers that best match the slots on the right.

Although very uncommon, the Cisco certification testing environment does allow for the Fill-in-the-Blank question format. In this type of question, a question is asked and the tester is to expected to input the correct answer into the Fill-in-the-Blank box.

A Testlet is a question whereby a scenario is given. The examinee is given multiple choices to choose from to address the given scenario.

The Simlet questions will provide a simulated scenario. With this scenario, you will be asked a number of questions—usually multiple-choice questions. After answering all of the multiple-choice questions, you can submit your collective answer from the Simlet. Be sure that you have answered all of the multiple-choice answers before submitting the Simlet.

The final question format is a Simulated Lab. The exam software has the ability to emulate a number of different Cisco devices interconnected in a simulated network. As part of this Simulated Lab question type, you will be asked to configure the relevant network devices. You will interact with the simulated device in a manner similar to how you would interact with the device in a real-live network. If a graphical user interface (GUI) is the normal method of configuring the test device, you will need to use the GUI to affect the configuration and behavior of the affected device. If you normally use the command-line interface (CLI) to configure a device, the CLI may be the best way to configure the device during your exam. In this Simulated Lab environment, not all commands are going to be available and the standard '?' context-sensitive help available on Cisco Routers and Switches or Tab-completion for commands may not be available. However, all commands that are needed to complete the question adequately should be available.

Again, the format of the CCNP-level tests is very similar to the format of the CCNA. There are examples of the question formats available on Cisco's Learning Network. The direct link to this Exam Tutorial can be found at http://www.cisco.com/web/learning/wwtraining/certprog/training/cert_exam_tutorial.html.

CCNP Security SISE 300-715 Official Certification Guide

As you review the contents of this book, take every opportunity you can to apply the information to your daily job, your studies, and any supplemental training that you may do. By applying the information within this book whenever possible, it will help to reinforce the material—making it more relevant to your particular application and, hopefully, making it easier to remember when you take the actual certification exam.

In Part I of the book, the focus will be on Identity Management and Secure Access. In this part, we will be discussing how to manage the users as well as how to allow them secure access to the network. The chapters in Part I help present the basis of Authentication, Authorization, and Accounting—AAA. We'll cover the management of users—leveraging the internal user database of Cisco's Identity Services Engine (ISE)—as well as third-party enterprise databases. The verification of the user via one of these databases—internal or external—is called Authentication.

There are a number of methods that can be used to authenticate users when they are joining the network. We'll cover a number of these authentication methods and the underlying protocols during this first part of the book. We'll cover how to authenticate a wired and wireless user using 802.1X, MAC Authentication Bypass, as well as non-standard flows including Local and Centralized Web Authentication.

Once we've authenticated the user, we'll need to dictate the level of access that the user will be given on the network. This process is called authorization. Authorization oftentimes leverages the authentication step—providing differentiated access to each endpoint based as much on the user who owns the device as the device itself.

We'll round out this part of the book by discussing some advanced concepts—diving more deeply into some of the details of how ISE and the supporting network infrastructure accomplish what needs to be accomplished. By the end of this first section, you should have a pretty good overview of the end-to-end AAA process.

Part II of the book will focus on Cisco's Identity Services Engine (ISE) and its configuration. We'll discuss the specific roles that each persona plays in the ISE architecture and several common deployment scenarios. After this overview of ISE architecture, we'll walk you through the ISE GUI and do some initial configuration of ISE including certificate generation and assignment as well as identity stores—those internal and external databases that provide us the authentication function.

After we have firmly established a complete understanding of AAA concepts and constructs, we'll consider the policy on ISE for both authentication and authorization. We'll walk you step-by-step through how ISE is configured for authentication policies and authorization policies—highlighting all of the building blocks that are required for a typical enterprise deployment.

Depending on the method of access (for example, wired versus wireless), the manner in which we enforce the level of access may change. For instance, the enforcement mechanisms (VLANs, Access Control Lists, Security Group Access, etc.) may be different depending on the method of access. By combining the authentication method (802.1X, MAB, and so on), the method of access (wired versus wireless), endpoint posturing, and profiling, we'll be able to leverage ISE to granularly apply differentiated access to each endpoint individually.

Part III of the book will move most of its focus away from ISE and onto the individual network devices that form the network infrastructure—the switches and wireless LAN controllers. We'll review how to configure the various Switching and Wireless platforms to put our AAA policy into action—leveraging 802.1X, MAB, as well as Local and Centralized Web Authentication.

We'll finish off Part III by reviewing some special use cases—how to configure guest services within ISE as well as how to profile devices as they try to join the network. Configuring guest services can be essential to an enterprise deployment—either by providing basic Internet access to employees or access to vendors and visitors. Profiling is a process whereby ISE can make an intelligent guess as to what type of device is joining the network—making granular authorization decisions based on device type. By the end of Part III, you should have a pretty solid understanding of how to secure your network leveraging ISE as the AAA server and the infrastructure devices to enforce the ISE's policy.

As we get into the Part IV of the book, Advanced Secure Network Access, we'll start to apply more of our knowledge in an advanced manner. Up to this point, we were doing basic configuration and basic policy enforcement. In the chapters in Part IV, we'll incorporate certificate-based user authentication—authenticating a user based on an X.509 certificate, either issued by ISE or by a third-party device. The ability to use certificates to validate a user can greatly enhance the level of security in the authentication process.

Bring your own device (BYOD) is also an advanced topic that we'll cover in this part of the book. BYOD is a process and security infrastructure that allows a user to bring her personal smart device onto the corporate network. The BYOD onboarding process allows a user to self-manage his device and registers the device to the corporate network. There are a number of special portals and configurations that are required to allow for an effective BYOD deployment. To ensure that this personal device doesn't adversely affect the network or gain access to unauthorized resources, ISE can provide differentiated access to the endpoint based on a number of key factors.

The next advanced topic that we'll review in Part IV is TrustSec and MACSec. We'll do a quick overview of these two topics and highlight some of benefits as well as the constructs and configurations that affect the Security Group Access configuration and enforcement both on the device and within ISE.

The final topic that we'll address in Part IV is Posture Assessment. Posturing and profiling are sometimes used interchangeably, but that is not accurate. Profiling often leverages information that is readily available via protocols that run over the network—including protocols such as RADIUS, DHCP, HTTP, as well as MAC addresses that are provided within the RADIUS exchange protocol. By replicating or otherwise sending this data to ISE as a client joins the network, profiling is able to make an intelligent decision as to what device is trying to join the network—without ever actively probing the device. Posturing is a little more entrenched at the client/endpoint level. Posturing will leverage information that is contained deep in the configuration of the endpoint—requiring a posturing agent to be run on the endpoint. Once key information is read from the endpoint via this agent, the ISE will make a decision as to whether the device/user is compliant to be allowed access to the network and, if so, what level of access the user should be given.

Part V of this book is geared toward the operational aspects of having ISE. As part of this chapter, we'll discuss how to slowly roll out your ISE deployment to minimize network outages. By leveraging deployment phasing, a network administrator can be in "monitor mode" whereby a device will not be denied access to the network but simply a log is thrown if the user doesn't match an available policy. This allows network administrators

to fully discover and understand the endpoints on their network—without having an adverse effect on the users. Once the network administrators are confident that they have reasonably triaged any unknown endpoints, they can gradually increase the level of policy enforcement.

A second important topic covered in Part V is ISE scale and high availability. This part will highlight how to configure and deploy a distributed ISE architecture in order to accommodate additional load, demand, and possible additional features. Each instance of ISE has an upper limit based on the platform and particular software that it is running on. By providing a distributed deployment architecture, the ISE deployment can grow as a company grows—incorporating a new ISE appliance whenever needed.

As we round out Part V of the CCNP-Security SISE 300-715 Official Certification Guide, we'll provide you with some tips and tricks to troubleshoot ISE. Some of these tools include a configuration validator, Live Logs, as well as a TCP dump. In the right hands, these tools can provide all of the necessary information to isolate any quality or network issues.

In Part VI of the book, we'll dive into turning ISE into the center of a full security ecosystem and extending the access control with other security products using the platform exchange grid (pxGrid) and adding some much needed security operations value to posture by extending network access control with threat and vulnerability data using threat-centric NAC.

Part VII rounds out the exam learning topics of the book with the other half of authentication, authorization, and accounting (AAA) device administration. This is the ability to control access to the network devices like Cisco routers, switches, and wireless controllers.

In the final section, Part VIII, we'll describe the steps that you'll need to take in order to prepare for the CCNP Security SISE.

Objectives and Methods

This book uses several key methodologies to help you discover the exam topics on which you need more review, to help you fully understand and remember those details, and to help you prove to yourself that you have retained your knowledge of those topics. This book does not try to help you pass the exam only by memorization; it seeks to help you to truly learn and understand the topics. This book is designed to help you pass the Security Identity Management SISE (300-715) exam by using the following methods:

- Helping you discover which exam topics you have not mastered
- Providing explanations and information to fill in your knowledge gaps
- Supplying exercises that enhance your ability to recall and deduce the answers to test questions
- Providing practice exercises on the topics and the testing process via test questions on the companion website

Book Features

To help you customize your study time using this book, the core chapters have several features that help you make the best use of your time:

- **Foundation Topics:** These are the core sections of each chapter. They explain the concepts for the topics in that chapter.
- **Exam Preparation Tasks:** After the "Foundation Topics" section of each chapter, the "Exam Preparation Tasks" section lists a series of study activities that you should do at the end of the chapter:
  - **Review All Key Topics:** The Key Topic icon appears next to the most important items in the "Foundation Topics" section of the chapter. The Review All Key Topics activity lists the key topics from the chapter, along with their page numbers. Although the contents of the entire chapter could be on the exam, you should definitely know the information listed in each key topic, so you should review these.
  - **Define Key Terms:** Although the PenTest+ exam may be unlikely to ask a question such as "Define this term," the exam does require that you learn and know a lot of pentest-related terminology. This section lists the most important terms from the chapter, asking you to write a short definition and compare your answer to the glossary at the end of the book.
  - **Review Questions:** Confirm that you understand the content that you just covered by answering these questions and reading the answer explanations.
- **Web-based practice exam:** The companion website includes the Pearson Cert Practice Test engine that allows you to take practice exam questions. Use it to prepare with a sample exam and to pinpoint topics where you need more study.

The Companion Website for Online Content Review

All the electronic review elements, as well as other electronic components of the book, exist on this book's companion website.

To access the companion website, which gives you access to the electronic content with this book, start by establishing a login at www.ciscopress.com and register your book.

To do so, simply go to www.ciscopress.com/register and enter the ISBN of the print book: 9780136642947. After you have registered your book, go to your account page and click the **Registered Products** tab. From there, click the **Access Bonus Content** link to get access to the book's companion website.

Note that if you buy the *Premium Edition eBook and Practice Test* version of this book from Cisco Press, your book will automatically be registered on your account page. Simply go to your account page, click the **Registered Products** tab, and select **Access Bonus Content** to access the book's companion website.

Please note that many of our companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following the steps at left, please visit www.pearsonITcertification.com/contact and select the **Site Problems/ Comments** option. Our customer service representatives will assist you.

How to Access the Pearson Test Prep (PTP) App

You have two options for installing and using the Pearson Test Prep application: a web app and a desktop app. To use the Pearson Test Prep application, start by finding the registration code that comes with the book. You can find the code in these ways:

- **Print book:** Look in the cardboard sleeve in the back of the book for a piece of paper with your book's unique PTP code.
- **Premium Edition:** If you purchase the Premium Edition eBook and Practice Test directly from the Cisco Press website, the code will be populated on your account page after purchase. Just log in at www.ciscopress.com, click **account** to see details of your account, and click the **digital purchases** tab.
- **Amazon Kindle:** For those who purchase a Kindle edition from Amazon, the access code will be supplied directly from Amazon.
- **Other Bookseller E-books:** Note that if you purchase an e-book version from any other source, the practice test is not included because other vendors to date have not chosen to vend the required unique access code.

NOTE Do not lose the activation code because it is the only means with which you can access the QA content with the book.

Once you have the access code, to find instructions about both the PTP web app and the desktop app, follow these steps:

Step 1. Open this book's companion website, as was shown earlier in this Introduction under the heading "How to Access the Companion Website."

Step 2. Click the **Practice Exams** button.

Step 3. Follow the instructions listed there both for installing the desktop app and for using the web app.

Note that if you want to use the web app only at this point, just navigate to www.pearsontestprep.com, establish a free login if you do not already have one, and register this book's practice tests using the registration code you just found. The process should take only a couple of minutes.

NOTE Amazon eBook (Kindle) customers: It is easy to miss Amazon's email that lists your PTP access code. Soon after you purchase the Kindle eBook, Amazon should send an email. However, the email uses very generic text and makes no specific mention of PTP or practice exams. To find your code, read every email from Amazon after you purchase the book. Also do the usual checks for ensuring your email arrives like checking your spam folder.

NOTE Other eBook customers: As of the time of publication, only the publisher and Amazon supply PTP access codes when you purchase their eBook editions of this book.

Customizing Your Exams

Once you are in the exam settings screen, you can choose to take exams in one of three modes:

- **Study mode:** Allows you to fully customize your exams and review answers as you are taking the exam. This is typically the mode you would use first to assess your knowledge and identify information gaps.
- **Practice Exam mode:** Locks certain customization options, as it is presenting a realistic exam experience. Use this mode when you are preparing to test your exam readiness.
- **Flash Card mode:** Strips out the answers and presents you with only the question stem. This mode is great for late-stage preparation when you really want to challenge yourself to provide answers without the benefit of seeing multiple-choice options. This mode does not provide the detailed score reports that the other two modes do, so you should not use it if you are trying to identify knowledge gaps.

In addition to these three modes, you will be able to select the source of your questions. You can choose to take exams that cover all of the chapters or you can narrow your selection to just a single chapter or the chapters that make up specific parts in the book. All chapters are selected by default. If you want to narrow your focus to individual chapters, simply deselect all the chapters and then select only those on which you wish to focus in the Objectives area.

You can also select the exam banks on which to focus. Each exam bank comes complete with a full exam of questions that cover topics in every chapter. The two exams printed in the book are available to you as well as two additional exams of unique questions. You can have the test engine serve up exams from all four banks or just from one individual bank by selecting the desired banks in the exam bank area.

There are several other customizations you can make to your exam from the exam settings screen, such as the time of the exam, the number of questions served up, whether to randomize questions and answers, whether to show the number of correct answers for multiple-answer questions, and whether to serve up only specific types of questions. You can also create custom test banks by selecting only questions that you have marked or questions on which you have added notes.

Updating Your Exams

If you are using the online version of the Pearson Test Prep software, you should always have access to the latest version of the software as well as the exam data. If you are using the Windows desktop version, every time you launch the software while connected to the Internet, it checks if there are any updates to your exam data and automatically downloads any changes that were made since the last time you used the software.

Sometimes, due to many factors, the exam data may not fully download when you activate your exam. If you find that figures or exhibits are missing, you may need to manually update your exams. To update a particular exam you have already activated and downloaded, simply click the **Tools** tab and click the **Update Products** button. Again, this is only an issue with the desktop Windows application.

If you wish to check for updates to the Pearson Test Prep exam engine software, Windows desktop version, simply click the **Tools** tab and click the **Update Application** button. This ensures that you are running the latest version of the software engine.

This page intentionally left blank



CHAPTER 1

Fundamentals of AAA

This chapter covers the following topics:

AAA options

TACACS+

RADIUS

Features and functionality of authentication and authorization

RADIUS flows

AV pairs

Device administration

In the world of security, we can only be as secure as the security controls we deploy permit us to be. Laws in the United States define what an airplane passenger is permitted to bring onboard. If the Transportation Security Administration (TSA) agents weren't operating the metal detectors and x-ray machines (and all the other things that slow us down when trying to reach our planes), how would the airport ever really enforce those policies?

With technology, we face similar challenges. We need to have controls in place to ensure that only the correct entities are using our technological gadgets. The same concepts can be applied to many use cases, including human interaction with a computer, a computer's interaction with a network, a human's interaction with an application, and an application's interaction with data.

Key Topic

An important security principle is known as authentication, authorization, and accounting (AAA), and it is often referred to as "triple A." Before allowing an entity to perform certain actions, you must ensure that you know who that entity actually is (authentication) and whether the entity is authorized to perform that action (authorization). In addition, you need to ensure that accurate records are maintained, showing that the action has occurred, so you keep a security log of the events (accounting).

The concepts of AAA may be applied to many different aspects of a technology lifecycle. However, the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam and, therefore, this book focus on the two main aspects of AAA related to networking:

- **Device administration:** ***Device administration AAA*** involves controlling access to who can log in to a network device console, Telnet session, or Secure Shell (SSH) session. AAA for device administration requires unique policy constructs.

- **Network access:** *Network access AAA* involves securing network access to ensuring the identity of the device or user before permitting the entity to communicate with the network. The SISE 300-715 exam and this book puts the most focus on this type of AAA.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 1-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 1-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Comparing and Selecting AAA Options	1–8
TACACS+	3, 6–8
RADIUS	4, 5, 7
Comparing RADIUS and TACACS+	1, 2, 9

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following best describes the difference between authentication and authorization?
   - **a.** There is no difference between authentication and authorization.
   - **b.** Authorization determines what a user may do, while authentication determines what devices the user may interact with.
   - **c.** Authentication is used with both network access and device administration, while authorization only applies to device administration.
   - **d.** Authentication involves validating a user's identity, while authorization involves determining what a user is permitted to do.
2. Which of the following are types of AAA covered on the SISE 300-715 exam? (Choose two.)
   - **a.** Device administration
   - **b.** Device access
   - **c.** A division of minor league baseball
   - **d.** Network access
   - **e.** Network administration

3. Which of the following protocols is best suited for granular command-level control with device administration AAA?
   a. DIAMETER
   b. TACACS+
   c. RADIUS
   d. RADIUS+
4. Which of the following protocols is best suited for authenticating and authorizing a user for network access AAA?
   a. TACACS+
   b. CHAP
   c. RADIUS
   d. MS-CHAPv2
5. Which of the following is most suited for device administration AAA?
   a. RADIUS
   b. SAML
   c. TACACS+
   d. OAuth 2
6. Which of the following Cisco products should be used for device administration with TACACS+?
   a. Cisco Secure Identity Control Server (ICS)
   b. Cisco Identity Services Engine
   c. Cisco TACACS+ Control Server (TCS)
   d. Cisco Centri
7. Why is RADIUS or TACACS+ needed? Why can't an end user authenticate directly to the authentication server?
   a. The added level of complexity enables Cisco and other vendors to sell more products.
   b. They are not needed.
   c. RADIUS and TACACS+ are used between the end user and the authentication server.
   d. Both RADIUS and TACACS+ extend the Layer 2 authentication protocols, allowing the end user to communicate with an authentication server that is not Layer 2 adjacent.
8. Which of the following is a TACACS+ message that is sent from the AAA client to the AAA server?
   a. START
   b. REPLY
   c. CHALLENGE
   d. REQUEST

9. When using RADIUS, what tells the AAA server what type of action is being authenticated?
   a. The TACACS+ service
   b. The Service-Type field
   c. RADIUS does not distinguish different services.
   d. The action AV pair
10. Which of the following best describes an AV pair?
   a. When communicating with an AAA protocol, the AV pair stipulates a common attribute or object and its assigned value.
   b. Cisco invented the term "AV pair" to confuse people who are studying for this exam.
   c. The AV pair is used to choose either TACACS+ or RADIUS.
   d. The AV pair is used to specify the quality of service (QoS) levels for audio and video traffic.

Foundation Topics

Authentication involves validating an identity or a credential. It is a very important step in the process of performing any sort of secure access control, regardless of what you are controlling. Forget about technology for a second and consider paying for groceries with a credit card. As a credit card owner, you have the choice to sign the back of the card or to write "check ID" on the back. The more secure method is to force the validation of the ID of the person using that card and ensure that the credential matches the name on the front of the credit card.

Having a cashier check the identity of the card user is authentication. Ensuring that the identity matches the name printed on the card is authorization. Say that Katherine McNamara goes into a retail store and hands the cashier a credit card to pay for the $10,000 of electronics she is purchasing. She passes her driver's license to the cashier, and the picture matches her face. It is certainly her identity. However, the name printed on the credit card is Aaron Woland. Should the transaction go through? Of course not: It should not be authorized.

After Katherine attempts to use Aaron's credit card, records of the attempt appear in the log files of the point-of-sale system, the video monitoring system of the store, and other systems. This is the accounting portion of AAA. It's a critical piece that is required for reporting, audits, and more.

It will become paramount for you as a security professional and as someone who is taking the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam to understand the differences between and purposes of all three As in AAA.

Comparing and Selecting AAA Options

As mentioned earlier in this chapter, the SISE 300-715 exam and, therefore, this book focus on two types of AAA: device administration AAA and network access AAA.

Device Administration AAA

Key Topic

Device administration is a method of AAA for controlling access to a network device console, Telnet session, Secure Shell (SSH) session, or device operating system for purposes of configuration. For example, say that your company has an Active Directory group named Cisco Administrators, and it should have full access (privilege level 15) to the Cisco switches in the company's network. The members of this group should therefore be able to make changes to VLANs, see the entire running configuration of the device, and more. There could be another group named Cisco Operators who should only be allowed to view the output of **show** commands but should not be allowed to configure anything in the device.

Device administration AAA can get granular. The Cisco Identity Services Engine (ISE) has a capability to provide command sets, which are lists of commands that an authenticated user is either permitted or not permitted to run. In other words, device administration means that a user may authenticate to the Cisco IOS shell, and ISE can permit or deny the user's execution of individual commands.

Figure 1-1 illustrates device administration graphically.

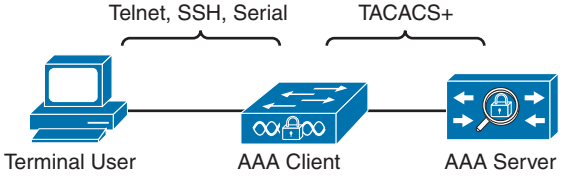


Figure 1-1 *Device Administration*

Device administration can be very interactive in nature, with the need to authenticate once but authorize many times during a single administrative session at the command line of a device. It therefore lends itself well to use of the Terminal Access Controller Access Control System (TACACS) client/server protocol and, to a lesser extent, Remote Authentication Dial-In User Service (RADIUS).

TACACS was designed for device administration AAA, to authenticate and authorize users into mainframe and UNIX terminals and other terminals or consoles.

Both the TACACS and RADIUS protocols are discussed in more depth later in this chapter, but for now you just need to know that TACACS separates out the authorization portion of AAA, allowing for a single authentication and multiple authorizations within the same session, which is why it lends itself to device administration. RADIUS, in contrast, does not provide the ability to control which commands can be executed.

NOTE To learn more about device administration AAA, beyond what you need to know to pass the SISE 300-715 exam, see the Cisco Press book *AAA Identity Management Security*.

Network Access AAA

Key Topic

Secure network access is essentially about learning the identity of the user or endpoint before permitting that entity to communicate within the network. This type of AAA is the main focus of the SISE 300-715 exam and this book. Network access AAA really took a strong hold back in the day of modems and dial-up networking with basic telephone service.

In the days of dial-up, people would gain Internet access by using their modems and dial-up. Basically, all that was needed was a modem and a phone line. Companies soon realized that it was not safe to allow anyone to dial in to their networks just by dialing their modem's phone numbers. They needed to authenticate a user before allowing a connection. RADIUS was created to deal with this need, and its name is a reminder of those days: Remote Access Dial-In User Service. RADIUS was used between a network access device (NAD) and an authentication server. The authentication was normally Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), or Microsoft CHAP (MS-CHAP).

Figure 1-2 illustrates dial-up remote access graphically.

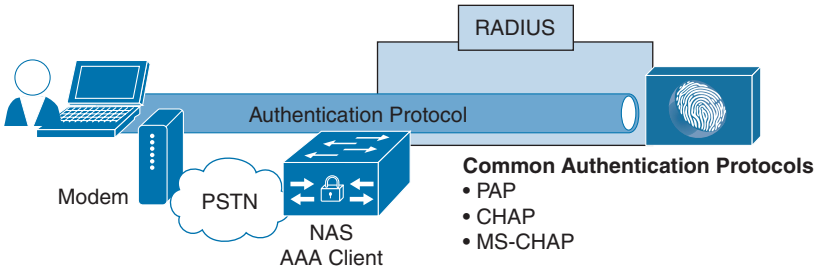


Figure 1-2 *Dial-Up Remote Access*

As technology continued to evolve, direct dial-in to a company was replaced by virtual private networks (VPNs), Wi-Fi became prevalent, and the IEEE standardized on a method to use Extensible Authentication Protocol (EAP) over local-area networks (IEEE 802.1X) using RADIUS to carry the authentication traffic. In fact, IEEE 802.1X cannot use TACACS; it must use RADIUS.

NOTE Another protocol similar to RADIUS, known as DIAMETER, may also be used with 802.1X; however, it is found mostly in the service provider space and is not covered on the SISE 300-715 exam.

RADIUS is the protocol used almost exclusively with network access AAA, although many security appliances use RADIUS to control access to their management consoles. This is mostly due to the prevalence of open-source RADIUS modules available for reuse and a significant misunderstanding of AAA in general.

TACACS+

Key Topic

Terminal Access Controller Access Control System (TACACS) is a protocol set created and intended for controlling access to UNIX terminals. Cisco created a new protocol called TACACS+, which was released as an open standard in the early 1990s. TACACS+ may be

derived from TACACS, but it is a completely separate and non-backward-compatible protocol designed for AAA. While TACACS+ is mainly used for device administration AAA, it is possible to use it for some types of network access AAA.

It is very important to understand that TACACS+ was not supported by Cisco ISE until ISE Version 2.0. Until that point, the Cisco Secure Access Control Server (ACS) was Cisco's primary authentication server product for enterprises that need to use TACACS+ for device administration AAA. Cisco ACS is now end of sale and fast approaching end of life.

NOTE Other Cisco products support TACACS+, such as the Cisco Access Registrar solution. However, those products are geared toward service providers and are not germane to the SISE 300-715 exam or this book.

TACACS+ uses Transmission Control Protocol (TCP) port 49 to communicate between the TACACS+ client and the TACACS+ server. For example, say that a Cisco switch is authenticating and authorizing administrative access to the switch's IOS CLI. The switch is the TACACS+ client, and ISE is the server, as illustrated in Figure 1-3.

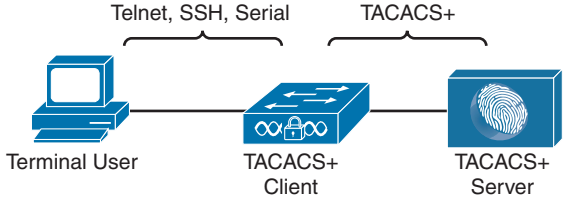


Figure 1-3 *TACACS+ Client/Server Communication*

One of the key differentiators of TACACS+ is its ability to separate authentication, authorization, and accounting as separate and independent functions. This is why TACACS+ is so commonly used for device administration, even though RADIUS is still certainly capable of providing device administration AAA.

Device administration can be very interactive in nature, possibly involving the need to authenticate once but authorize many times during a single administrative session in the command line of a device. A router or switch may need to authorize a user's activity on a per-command basis. TACACS+ is designed to accommodate that type of authorization need.

As its name indicates, TACACS+ was designed for device administration AAA, to authenticate and authorize users into mainframe and UNIX terminals and other terminals or consoles.

TACACS+ communication between the client and server uses different message types, depending on the function. In other words, different messages may be used for authentication than are used for authorization and accounting. Another very interesting point to know, especially when you take the SISE 300-715 exam, is that TACACS+ communication encrypts an entire packet.

TACACS+ Authentication Messages

When using TACACS+ for authentication, there are only three types of packets exchanged between the client (the network device) and the server:

- **START:** This packet is used to begin the authentication request between the AAA client and the AAA server.
- **REPLY:** These messages are sent from the AAA server to the AAA client.
- **CONTINUE:** These messages from the AAA client are responses to the AAA server requests for usernames and passwords.

This section looks at these messages in the context of the authentication process.

When an authentication request is sent from the client to the server, it begins with a START message from the network device to the server. The START message tells the server that an authentication request is coming. All messages from the server to the network device (client) during authentication are REPLY messages. The server sends a REPLY message asking for the client to retrieve the username. The client sends the username to the server in a CONTINUE message.

Once the server receives the username, it sends a REPLY message back to the client, requesting the password, which is sent back to the server in a CONTINUE message. The server sends a final REPLY message with the pass or fail status of the authentication request.

The final REPLY message from the AAA server to the AAA client can include the following possible values:

- **ACCEPT:** The user authentication succeeded, and the authorization process may begin, if the AAA client is configured for authorization.
- **REJECT:** The user authentication failed. The login is denied or the end user is prompted to try again, depending on the configuration of the AAA client.
- **ERROR:** An error occurred at some point during the authentication. AAA clients typically attempt to authenticate the user again or attempt a different method of authenticating the user.
- **CONTINUE:** The user is prompted for additional information. This value sent from the AAA server within a REPLY message, indicating that more information is required, should not be confused with the CONTINUE message sent from the AAA client to the AAA server.

Figure 1-4 illustrates the authentication messages between the client and server.

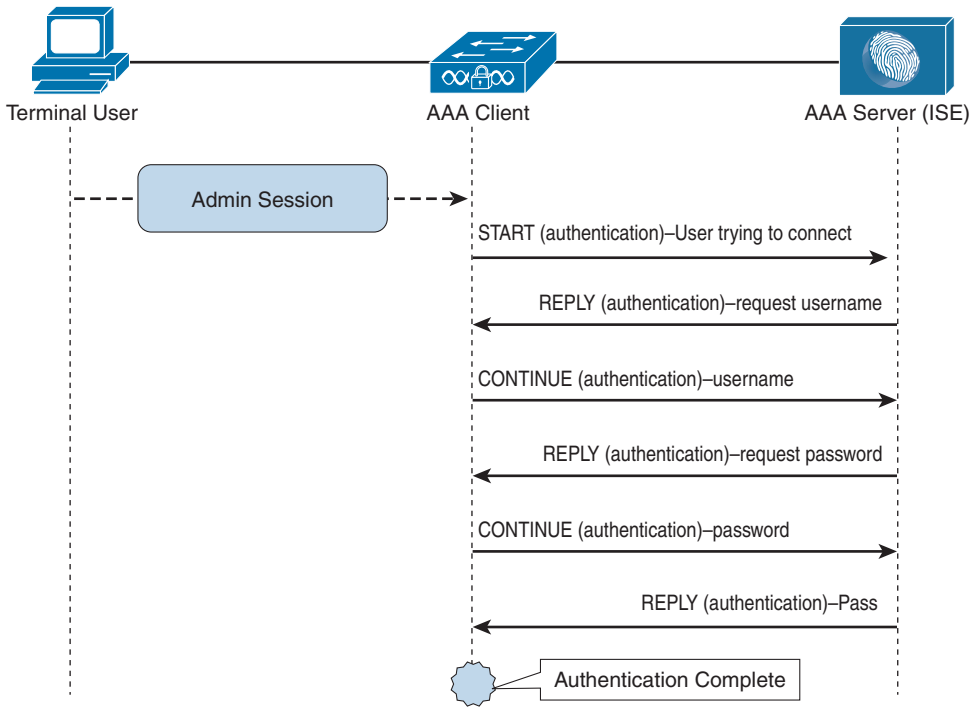


Figure 1-4 *TACACS+ Authentication Communication Flows*

TACACS+ Authorization and Accounting Messages

When using TACACS+ for authorization, only two messages are used between the AAA client and the AAA server:

- **REQUEST:** This message is sent from the AAA client to the AAA server to request authorization. The authorization may be related to access to a CLI shell or possibly to authorize a specific command. The protocol communication does not discriminate. The function requested is known as a "service." For example, the service would be "shell" for CLI access to a device running Cisco IOS. Each service may be communicated with attribute/value (AV) pairs, which are discussed later in this chapter. You can find more about specific TACACS+ AV pairs by searching online for "TACACS AV Pairs."
- **RESPONSE:** This message is sent from the AAA server back to the AAA client with the result of the authorization request, including specific details, such as the privilege level assigned to the end user. A RESPONSE message may contain one of the following replies:
  - **FAIL:** This response indicates that the user should be denied access to the requested service.
  - **PASS_ADD:** This response indicates successful authorization, and the information contained within the RESPONSE message should be used in addition to the

requested information. If no additional arguments are returned by the AAA server within the RESPONSE message, then the request is simply authorized as is.

- **PASS_REPL:** This response indicates successful authorization, but the server has chosen to ignore the REQUEST message and is replacing it with the information sent back in the RESPONSE message.
- **FOLLOW:** This reply indicates that the AAA server wants the AAA client to send the authorization request to a different server. The new server information will be listed in the RESPONSE packet. The AAA client may use that new server or treat the response as a FAIL.
- **ERROR:** This response indicates a problem occurring on the AAA server. Further troubleshooting needs to occur.

A key function of AAA that cannot be overlooked is accounting. It is crucial to security to have a record of what has transpired. In addition to the authorization request being sent to the AAA server, there should be accounting records of the activities of the user.

Much as with authorization messages, there are only two message types used in accounting:

- **REQUEST:** This message is sent from the AAA client to the AAA server to indicate a notification of activity. Three values that may be included in a REQUEST message:
  - **START:** This value indicates that a service has begun.
  - **STOP:** This value indicates that the service has ended.
  - **CONTINUE:** The CONTINUE record, sometimes referred to as a Watchdog or UPDATE record, is sent when a service has already started and is in progress, but there is updated information to provide in relationship to the service.
- **RESPONSE:** This message is sent from the AAA server back to the AAA client with the result of the accounting REQUEST message. It may contain one of three replies:
  - **SUCCESS:** This value indicates that the server received the record from the client.
  - **ERROR:** This value indicates an error on the server and that the record was not stored.
  - **FOLLOW:** This value indicates that the server wants the client to send the record to a different AAA server and includes that server's information in the RESPONSE message.

Figure 1-5 illustrates an end user being authorized to access the IOS exec CLI. This figure continues the process begun in Figure 1-4, where the authentication occurred. In Figure 1-5, the end user is authorized to enter the IOS exec and run the **show run** command.

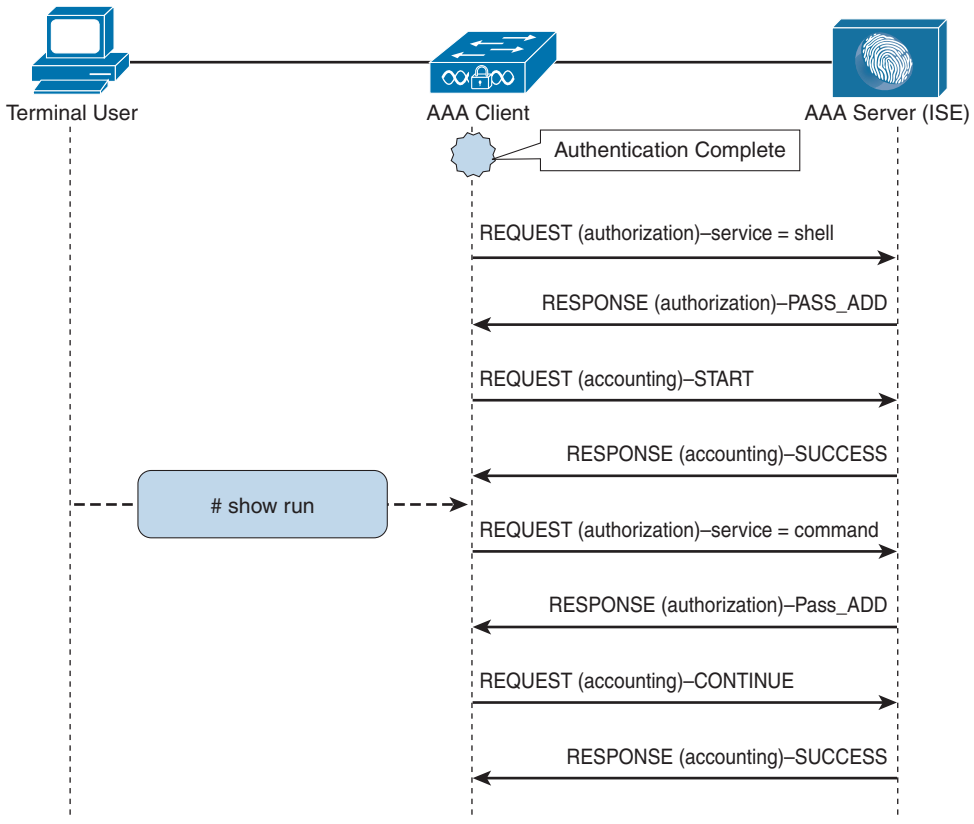


Figure 1-5 *TACACS+ Authorization and Accounting Communication Flows*

NOTE For more information on TACACS+, see the Cisco Press book *AAA Identity Management Security*.

RADIUS

Key Topic

Remote Access Dial-In User Service (RADIUS) is an IETF standard for AAA. Like TACACS+, it is based on a client/server model, where the client initiates requests to the server. RADIUS is the protocol of choice for network access AAA. If you connect to a secure wireless network regularly, RADIUS is most likely being used between the wireless device and the AAA server because RADIUS is the transport protocol for Extensible Authentication Protocol (EAP), as well as for many other authentication protocols.

Originally, RADIUS was used to extend authentication from the Layer 2 Point-to-Point Protocol (PPP) used between the end user and the network access server and carry that authentication traffic from the network access server to the AAA server performing the authentication. This allowed a Layer 2 authentication protocol to be extended across Layer 3 boundaries to a centralized authentication server.

As mentioned earlier in this chapter, RADIUS has evolved far beyond just the dial-up networking use cases it was originally created for. Today it is still used in the same way, carrying the authentication traffic from the network device to the authentication server. With IEEE 802.1X, RADIUS is used to extend the Layer 2 EAP from the end user to the authentication server, as illustrated in Figure 1-6.

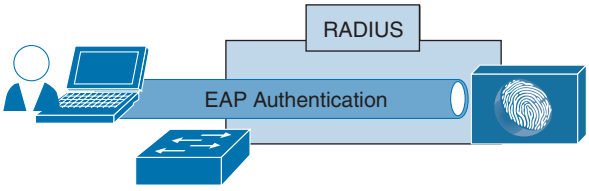


Figure 1-6 *RADIUS Carrying the Layer 2 EAP Communication*

There are many differences between RADIUS and TACACS+. One of them is that authentication and authorization are not separated in a RADIUS transaction. When an authentication request is sent to a AAA server, the AAA client expects to have the authorization result sent back in reply.

There are only a few message types with RADIUS authentication and authorization:

- **Access-Request:** This message is sent from the AAA client to the AAA server to request authentication and authorization. The request could be for network access or for device shell access; RADIUS does not discriminate. The function requested is included in the Service-Type field. For example, the Service-Type field value might be Framed for an IEEE 802.1X authentication. Table 1-2 lists some common RADIUS Service-Type values. For a more complete list of Service-Type values, search online for "RADIUS Service-Types."

Table 1-2 Common RADIUS Service-Type Values

Service-Type Field Value	Name	Commonly Used For
1	Login	Login requests; often used with web authentication with network equipment from vendors other than Cisco
2	Framed	IEEE 802.1X
5	Outbound	Local web authentication
10	Call-Check	MAC Authentication Bypass (MAB)

- **Access-Accept:** This message is sent from the AAA server to the AAA client to signal a passed authentication. The authorization result is included as AV pairs. The AV pairs may include items such as the assigned VLAN, a downloadable access control list (dACL), or a security group tag (SGT).
- **Access-Reject:** This message is sent from the AAA server to the AAA client to signal authentication failure and indicate that no authorization has been granted.
- **Access-Challenge:** This optional message may be sent from the AAA server to the AAA client when additional information is needed, such as a second password for two-factor authentications.

Figure 1-7 illustrates a sample RADIUS flow. In Figure 1-7, note that authentication and authorization are combined with RADIUS. The Access-Accept message includes the AV pairs defining what the user is authorized to do.

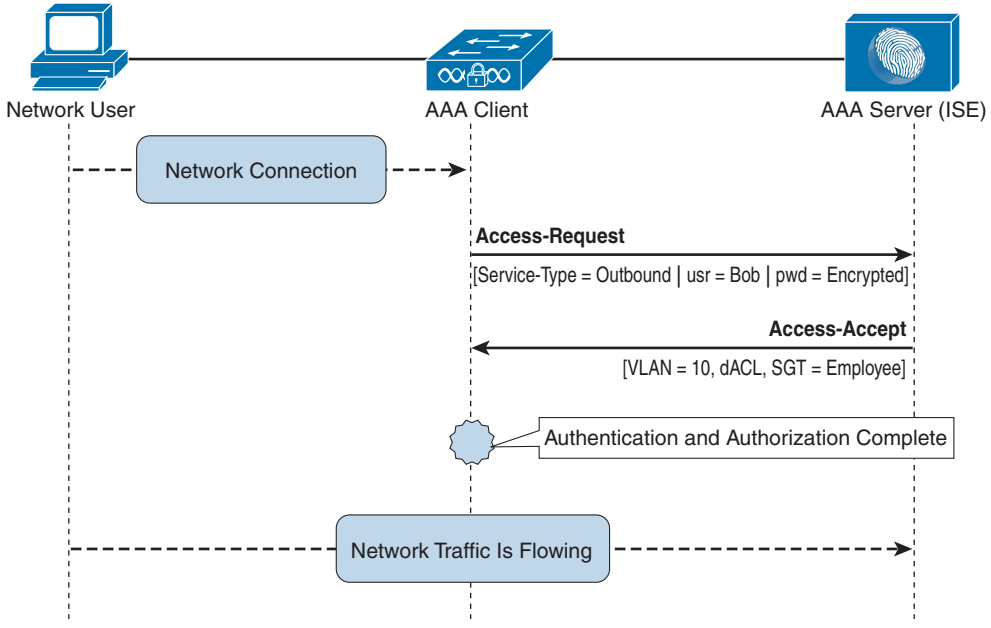


Figure 1-7 *RADIUS Authentication and Authorization Communication Flows*

The accounting function of AAA is crucial to security as it provides a record of what has transpired. In addition to the authorization request being sent to the AAA server, there should be accounting records of the activities of the user.

Only two message types are used in accounting:

- **Accounting-Request:** This message is sent by the AAA client to the AAA server. It may include time, packets, DHCP information, CDP information, and so on. The message may be a START message indicating that service has begun or a STOP message indicating the service has ended.
- **Accounting-Response:** This message acts as an acknowledgment of receipt, so the AAA client knows the accounting message was received by the AAA server.

Figure 1-8 illustrates a sample RADIUS accounting flow. The figure continues the process shown in Figure 1-7, where the authentication and authorization occurred.

Unlike TACACS+, RADIUS uses UDP as the transmission protocol. The standard ports used by RADIUS are UDP port 1812 for authentication and UDP port 1813 for accounting. Cisco supported RADIUS before the standard was ratified, and the ports used were UDP port 1645 for authentication and UDP port 1646 for accounting. Most Cisco devices support using either set of ports to ensure backward compatibility.

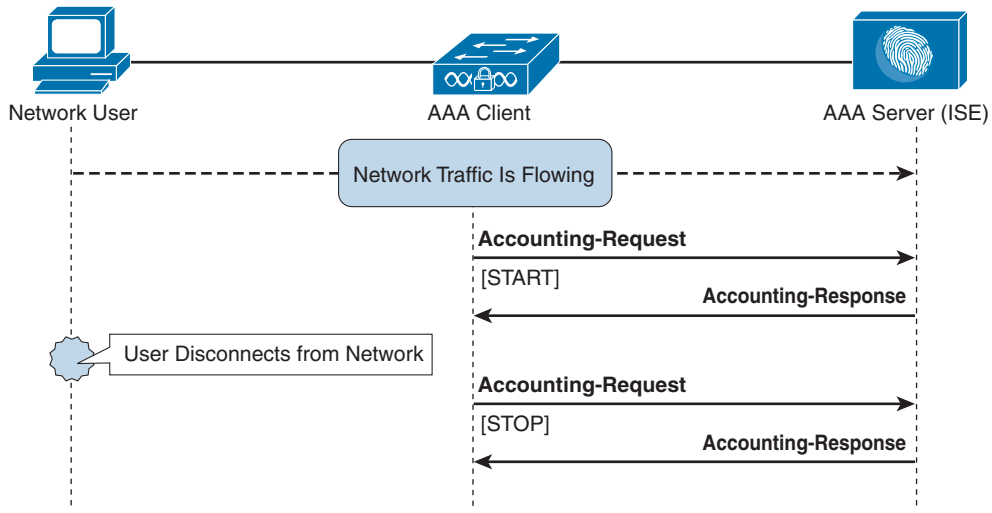


Figure 1-8 *RADIUS Authentication and Authorization Communication Flows*

AV Pairs

Key Topic

When communicating with an AAA protocol, many attributes can be referenced to clearly indicate answers or results. A RADIUS server may assign an attribute to the authentication session, such as for a VLAN. In this case, the VLAN placeholder is the attribute, and the assigned VLAN number is the value for that placeholder. The attribute in the AAA communication and its assigned value are paired together and referred to as an attribute/value pair (AV pair). Figure 1-9 illustrates the concept of AV pairs by using a table with attribute and value columns to represent the pairings.

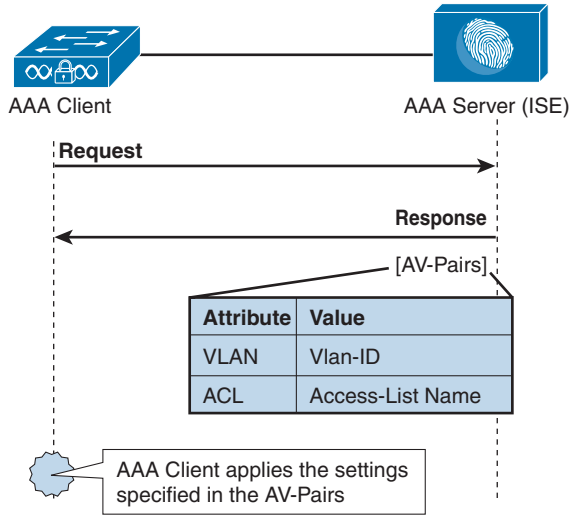


Figure 1-9 *AV Pairs*

Change of Authorization (CoA)

RADIUS was always defined for client/server architectures, with the client always initiating the conversation, and it became challenging for the AAA server to take action. As RADIUS was defined, the AAA server could only assign an authorization in reply to an authentication request.

As technology advanced, many new demands appeared, including the ability for the network to deal with misbehaving clients by rejecting them, quarantining them, or changing their access. It was difficult for such demands to be met when the network access used a RADIUS control plane, and the AAA client always needed to initiate the RADIUS conversations. RFC 3576 and its successor RFC 5176 defined a new enhancement to RADIUS known as Dynamic Authorization Extensions to RADIUS, more commonly called *Change of Authorization* (*CoA*).

Key Topic

CoA allows a RADIUS server to initiate a conversation with a network device and disconnect a user's session, bounce the port (perform a shut/no-shut), or even tell the device to reauthenticate the user. As you learn more about Cisco ISE and the advanced functionality it brings to network access AAA, you will see how critically important CoA is.

Comparing RADIUS and TACACS+

Table 1-3, which provides a reference for your studies, summarizes the two main AAA protocols: RADIUS and TACACS+.

Table 1-3 Comparing RADIUS and TACACS+

Feature	RADIUS	TACACS+
Protocol and port(s) used	UDP: 1812 and 1813 or UDP: 1645 and 1646	TCP: 49
Security	Hashes the password field only	Encrypts the entire payload
Authentication and authorization	Combines authentication and authorization	Separates authentication and authorization
Primary use	Network access	Device administration

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 1-4 lists these key topics and the page number on which each is found.

Key Topic

Table 1-4 Key Topics

Key Topic Element	Description	Page
Paragraph	The principles of AAA	2
Section	Device administration	6
Section	Network access	7
Section	TACACS+	7
Section	RADIUS	12
Section	AV pairs	15
Paragraph	CoA	16

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

device administration AAA, network access AAA

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What AAA protocol is best suited for device administration?
2. Which AAA protocol supported by ISE combines authentication and authorization in a single transaction?
3. Which AAA protocol supported by ISE uses TCP for transport?
4. In AAA communication, what is the name of the combination of an attribute and its assigned value?
5. What is the common name for the technology extension to RADIUS that allows communication to be initiated from the authentication server to the authenticator (the NAD)?



CHAPTER 2

Identity Management

This chapter covers the following topics:

- Identity store options, including LDAP, AD, PKI, OTP, smart card, and local options
- Special identity sources

What is an identity? How do you use identities in the design of your Identity Services Engine (ISE) infrastructure? This chapter answers these questions and introduces several key concepts regarding the use of identity management in ISE, including identity stores, identity store sequences, and internal and external identity sources. This chapter also looks at the various ISE authentication configuration options, including Active Directory, LDAP, one-time token passwords, smart cards, and the internal ISE user database. This chapter also includes an in-depth discussion of the use of public key infrastructure (PKI) with X.509 certificates for authentication with ISE deployments.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 2-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 2-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
What Is an Identity?	1
Identity Stores	2–4, 6–10
Identity Source Sequences	5

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What types of identities are used in Cisco ISE? (Choose two.)
   a. SSID
   b. MAC address
   c. Username
   d. IP address
2. Which of the following are types of identity stores used by Cisco ISE? (Choose two.)
   a. Temporary
   b. External
   c. Internal
   d. Permanent
3. Cisco ISE identity stores are used to authentication which of the following? (Choose two.)
   a. Endpoints
   b. AD security groups
   c. RADIUS
   d. Users
4. What identity store attributes can be used in an ISE authorization policy? (Choose two.)
   a. User
   b. Time
   c. Accounting
   d. Machine
5. Which of the following is a list of identity sources that are checked in order as part of authentication?
   a. Authentication source
   b. Identity database
   c. Identity source sequence
   d. Authentication database
6. How is an identity store sequence processed?
   a. Bottom-to-top
   b. Left-to-right
   c. Top-to-bottom
   d. In any order
7. Which of the following does ISE support for authentication? (Choose three.)
   a. LDAP
   b. DIAMETER
   c. Microsoft Active Directory
   d. RADIUS proxy servers

8. Which of the following can be used with an internal identity store? (Choose two.)
   a. SSID
   b. Guest
   c. Time
   d. MAB
9. What types of internal identity stores are used in ISE? (Choose two.)
   a. User database
   b. Endpoint database
   c. System database
   d. Admin database
10. What are the primary reasons for using external identity stores? (Choose two.)
   a. Performance
   b. Monitoring
   c. Scalability
   d. Management

Foundation Topics

What Is an Identity?

Who are you? We're not asking one of those existential questions that requires deep introspection while sweating it out in a tent with 100 other people during a heat wave. We're merely asking for your identity. Your ***identity*** is a representation of who you are. It's a simple question, and your response might be as simple as your name. But if you meet someone new and tell them your name, how do they know you are who you really say you are? What about the reverse situation? Someone could walk up to you and say that their name is "Suzanne," but how would you know that's the truth? Perhaps you could look at some form of "trusted" evidence, such as a driver's license or passport, to verify their identity. Such evidence of an identity is known as a ***credential***.

One of the foremost uses of credentials in the world of security and access control is to represent the identity of a person or a device. There are typically many factors that can make up an identity, including an employee's username and password that would be used to access corporate network resources. Another form of identity could be the use of a Media Access Control (MAC) address to uniquely identify a specific endpoint, such as a network-based printer.

Identity Stores

An ***identity store*** is basically a database that houses credentials of users or endpoints. Because it contains those values, it can be used to authenticate the identity of a user or an endpoint. The identity store could be an internal database that resides on the AAA server or an external database that houses the identities.

Identity stores can also be used for the retrieval of user or machine attributes used in authorization policies, as discussed in Chapter 10, "Authorization Policies." Each individual identity store is referred to as an ***identity source***.

Key Topic

Cisco ISE can leverage both internal and external identity sources. External identity sources include Microsoft Active Directory (AD), Lightweight Directory Access Protocol (LDAP), one-time password (OTP) servers, smart cards, certificate authorities, and even Security Assertion Markup Language (SAML) identity providers (IdPs) for special cases.

Internal Identity Stores

Cisco ISE is primarily used as an identity authentication and authorization platform. In fact, it came very close to being named Policy Services Engine. The majority of the identity stores that companies use are external to ISE, but there are some internal stores to enable some critical use cases.

Internal User Identities

Cisco ISE has a built-in ***internal user database*** that can be used for a limited number of user accounts, referred to as *internal users*, also (unfortunately) referred to as *network access users*.

One example for use of this internal identity store would be to authenticate and authorize a preconfigured account for a contractor or other temporary worker to access your corporation's wireless network, without needing to add those contractors to your company's main identity store. Another example would be to create internal user accounts that could be used for logging in to and administering Cisco routers through device administration AAA.

Figure 2-1 shows an example of creating an internal user account.

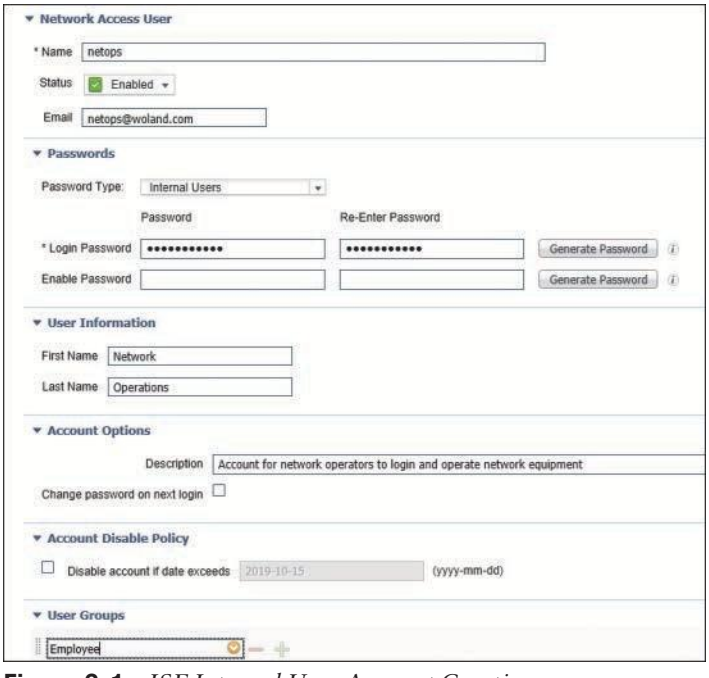


Figure 2-1 *ISE Internal User Account Creation*

Any internal user accounts configured can then be associated with a defined user ***identity group*** that sets the type of internal account being created. There are several preconfigured user identity groups that can be used when creating internal user accounts, including the Employees group and a few groups that are used for guest access.

User Identity Groups

You are not limited to using only the prebuilt user groups but can add new groups as needed. Keep in mind that most organizations use external identity sources (such as Microsoft's Active Directory) as their main identity stores. You should plan not to use a lot of internal identities or groups unless there is a special need to do so.

Figure 2-2 shows an example of creating a new group called NetOps, and Figure 2-3 shows a list of user groups.

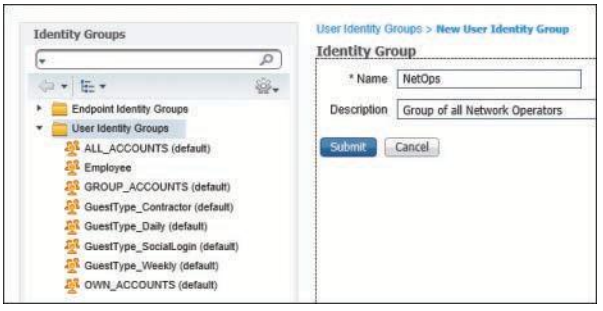


Figure 2-2 *Creating a New User Group*

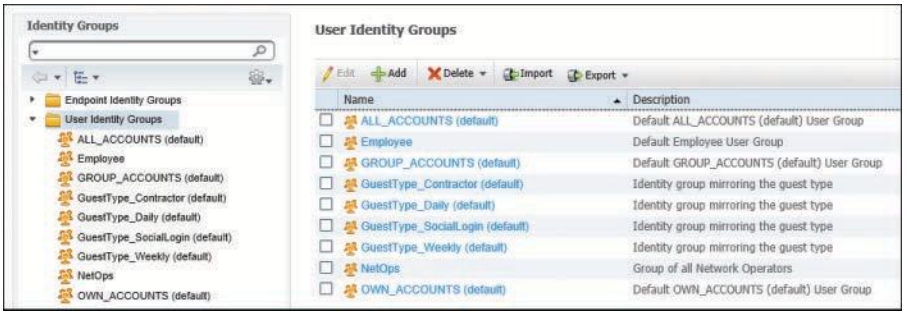


Figure 2-3 *ISE Internal User Identity Groups*

Endpoint Groups

Users are not the only entity on a network that have identities. Endpoints also have identities, and Cisco ISE maintains an ***internal endpoint database*** that stores information about all the endpoint devices that connect to the network infrastructure.

See Figure 2-4 for an example of endpoint objects in ISE's Internal Endpoint Database.

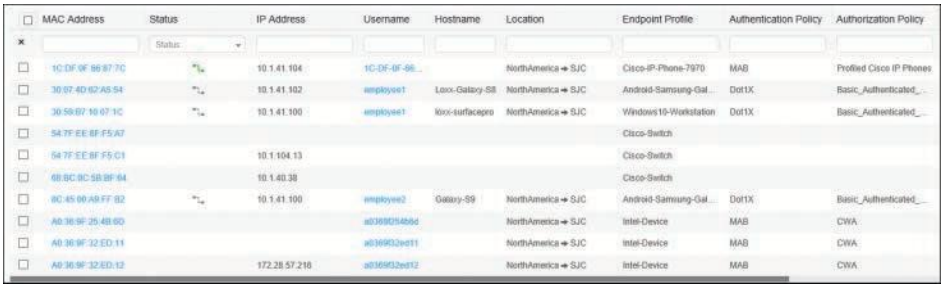


Figure 2-4 *ISE Internal Endpoint Database*

External Identity Stores

An organization is likely to have an existing identity store that is the master database for user identities. Many organizations refer to this store as the single source of truth for identity. ISE can integrate with that existing user database, and ISE refers to it as an ***external identity store.***

As user and machine accounts are created, modified, or deleted in the external identity store, you do not have to worry about these processes being duplicated for the Cisco ISE deployment. External identity stores allow for better scalability and management for ISE authentication and authorization processes, as well as for integrating other authentication services (such as for website or application authentication).

A Cisco ISE deployment connects to the external identity source to verify username and password credentials for authentication policies. Certificate information can also be retrieved from an external identity store.

Key Topic

Examples of external identity stores include Certificate Authentication Profiles, Active Directory (AD), Lightweight Directory Access Protocol (LDAP), RADIUS token servers, and RSA SecurID. Security Assertion Markup Language (SAML) and social login (Facebook only) are supported, but only for guest portals; they cannot be used as standard external identity sources for network access.

Figure 2-5 shows an example of external identity stores available for configuration in ISE.



Figure 2-5 *External Identity Stores*

Active Directory

Microsoft's ***Active Directory (AD)*** is one of the most popular external identity stores supported with Cisco ISE. In fact, the best estimate from Cisco's security business unit is that AD accounts for more than 90% of ISE deployments.

AD is arguably one of the most successful, highly available distributed database deployments in existence today. It follows the X.500 structure, just as DNS does. In fact, AD is very tightly coupled with DNS and requires DNS to be robust for AD to function normally.

Key Topic

For a security professional looking to implement network access control with ISE leveraging AD, it is quite important to understand Active Directory design. AD leverages ***AD domains***, which are groupings of users and machine accounts, permissions, and policies.

AD domains belong to an ***AD forest***, which is a collection of domains with trust relationships between them. In other words, a single forest may exist that contains multiple domains, and each of those domains may have implicit and explicit trusts between them, allowing them to work together.

Figure 2-6 provides an example of an AD forest that contains three different domains. The root domain, which can also be called the parent domain, is cisco.com. In the cisco.com forest, europe.cisco.com and america.cisco.com are both child domains.

Example AD Forest

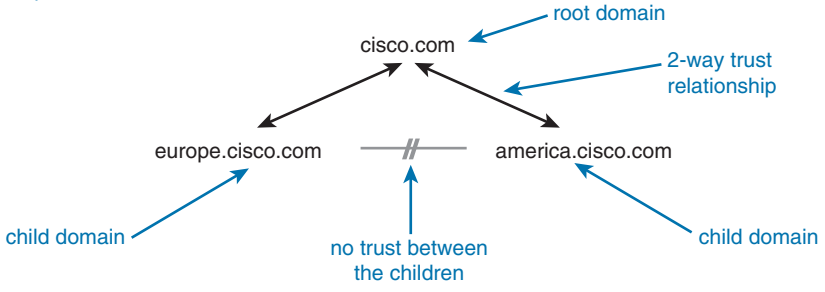


Figure 2-6 *AD Forest Example*

In the example illustrated in Figure 2-6, there is two-way trust between the parent domain and its children; the child domains do not trust one another at all. These trust relationships can be configured any way the AD administrator would like; this illustration is merely an example.

As a security professional implementing ISE, why should you care about these AD domain trust relationships? It's important because user and machine accounts (sometimes referred to as *objects*) can exist with any of the domains. For example, users bob@cisco.com, mary@americas.cisco.com, and lee@europe.cisco.com might all try to log in to a network where access is controlled by ISE. In this case, it would be very important to understand how ISE joins the forest and how it queries AD when performing authentications and authorizations.

ISE joins an AD domain. If that domain does not have a valid trust relationship with the other domains, authentication of those domain members will not be possible. So, ISE may have to join more than one domain.

Refer to Figure 2-6 for an example of ISE Active Directory identity source configuration. The configuration of ISE and Active Directory are covered even more in Chapter 8, "Initial Configuration of Cisco ISE."

NOTE As of ISE Version 2.7, **Microsoft's** Azure AD was not a natively supported identity store.

LDAP

Lightweight Directory Access Protocol (LDAP) is a standard protocol used to connect to X.500 directory services databases and retrieve information from them; these databases include Microsoft's Active Directory, NetIQ eDirectory (formerly Novell Directory Services [NDS]), Sun Directory Server, Oracle Identity Manager, IBM Tivoli Identity Manager (TIM), and many other identity stores.

NOTE LDAP is not used very often with ISE to connect to Active Directory because ISE has a more powerful connector built in.

The use of LDAP as an external identity store is supported for user authentication, as are retrieval of group membership information and account attributes that can be used in the Cisco ISE authentication and authorization policies.

LDAP uses TCP port 389, and all communications (including username and passwords) are sent in plaintext by default. To prevent plaintext credentials from going across the wire, it is highly recommended to secure LDAP communications by using Secure Sockets Layer (SSL), which uses TCP port 636.

X.500 directories are very organized and powerful. They can distribute data across multiple servers, and namespaces make it possible to quickly determine the locations of objects within a directory.

Figure 2-7 shows an example of an LDAP client displaying the namespace structure of a Windows domain, as well as some specific attributes of one of the user accounts (doctor1). The globally unique identifier of the object—the distinguishedName attribute (CN=doctor1, CN=Users,DC=securitydemo,DC=net)—helps illustrate the directory namespace.

A directory object is referred to by its common name and is represented with the CN designation. The DC designation represents the domain component of the namespace. In the example shown in Figure 2-7, the doctor1 user account exists in the Users organizational unit (OU), which is part of the securitydemo.net domain.

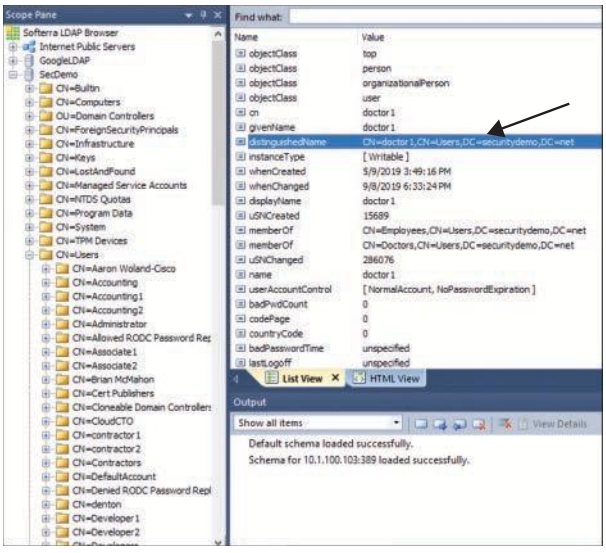


Figure 2-7 *ISE LDAP Identity Source Configuration*

NOTE The LDAP browser used in Figure 2-7 is the Softerra LDAP browser, which is a free tool you can use to test LDAP connections and view attributes. It's highly recommended to add it, or something similar, to your arsenal of tools when deploying ISE.

Multifactor Authentication

In traditional secure enterprises, the typical user identity is based on a username and some type of static password that is likely required by security policy to be changed periodically.

Passwords are often weak and, even worse, commonly reused across multiple websites or applications. When any of those applications are compromised, such as LinkedIn or Facebook, the bad guys have a password that may work in other places as well. The threat actors are likely to add the password to a giant database of usernames and passwords and use some form of machine learning or artificial intelligence to help parse that vast amount of data and come up with likely passwords for more targeted attacks. (For more on stolen password databases, check out the "have I been pwned" service at http://www.haveibeenpwned.com.)

Two types of authentication that add a lot more security are ***multifactor authentication (MFA) and two-factor authentication (2FA)***. MFA is a security enhancement that allows you to present two or more pieces of evidence when logging in to an account. MFA requires that your credentials (evidence) fall into any of these three categories:

- Something you know, such as a password or PIN
- Something you have, like a smart card or a smartphone that receives push notifications or text messages
- Something you are, like your fingerprint or faceprint (biometrics)

With MFA, your credentials must come from two different categories, so entering two different passwords would not be considered multifactor authentication. Just remember that

MFA and 2FA require something *you have* and a separate required something *you know* or *you are*.

Biometrics are becoming much more commonplace. Technology such as Windows Hello, which uses facial recognition or your fingerprint to log you in to Windows 10 devices, Apple's "faceprint" in the iPhone X (and newer), and the use of fingerprints on most Android devices have put biometric authentication in the hand of the average user.

In October 2018, Cisco acquired Duo Security for $2.35 billion. Duo got its start in the security space by creating a very simple-to-use and powerful MFA solution. It then expanded its product portfolio further into the zero trust space. Let's look at an example of MFA using Cisco's Duo Security solution. Say that a user logs in to a web application, such as the Cisco AMP for Endpoints portal shown in Figure 2-8. The user's browser is redirected to the Duo Login prompt, where the username and password are verified, and then a push notification is sent to the Duo Security app on the user's phone. Figure 2-9 shows the Duo Login prompt after the user's credentials are accepted, and Figure 2-10 shows the push notification on the user's mobile phone.

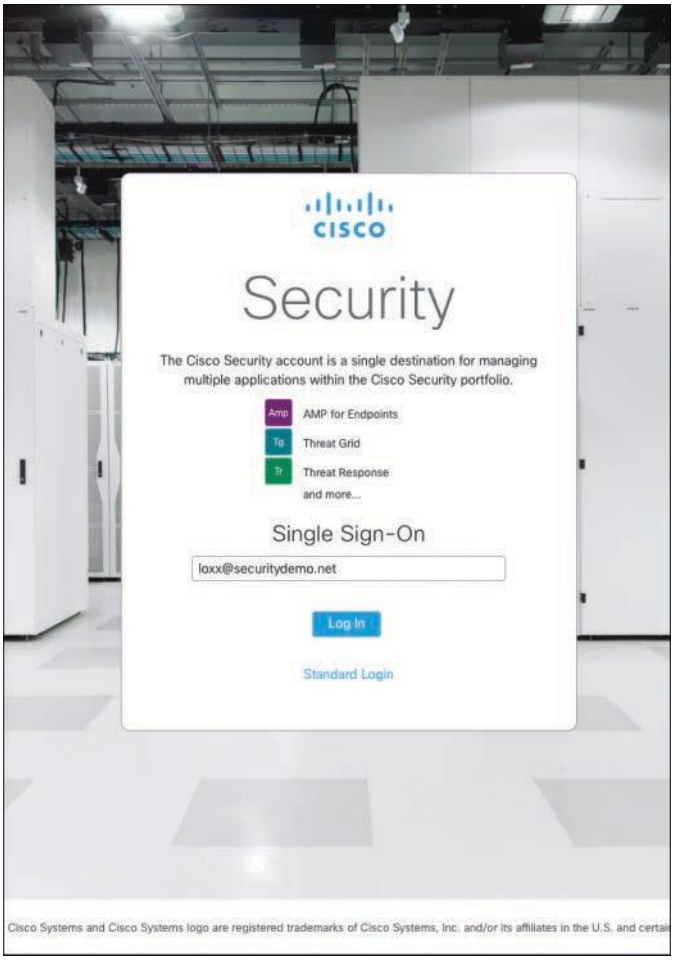


Figure 2-8 *Login Page for a Web Application*

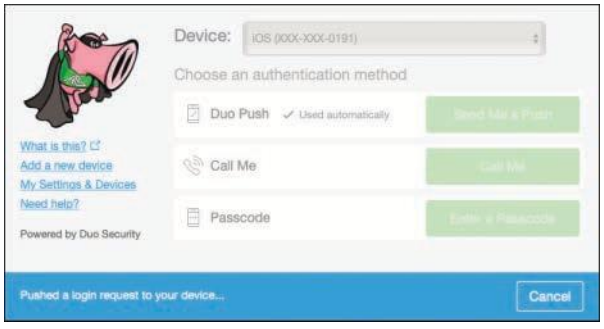


Figure 2-9 *Duo Security Login Page*

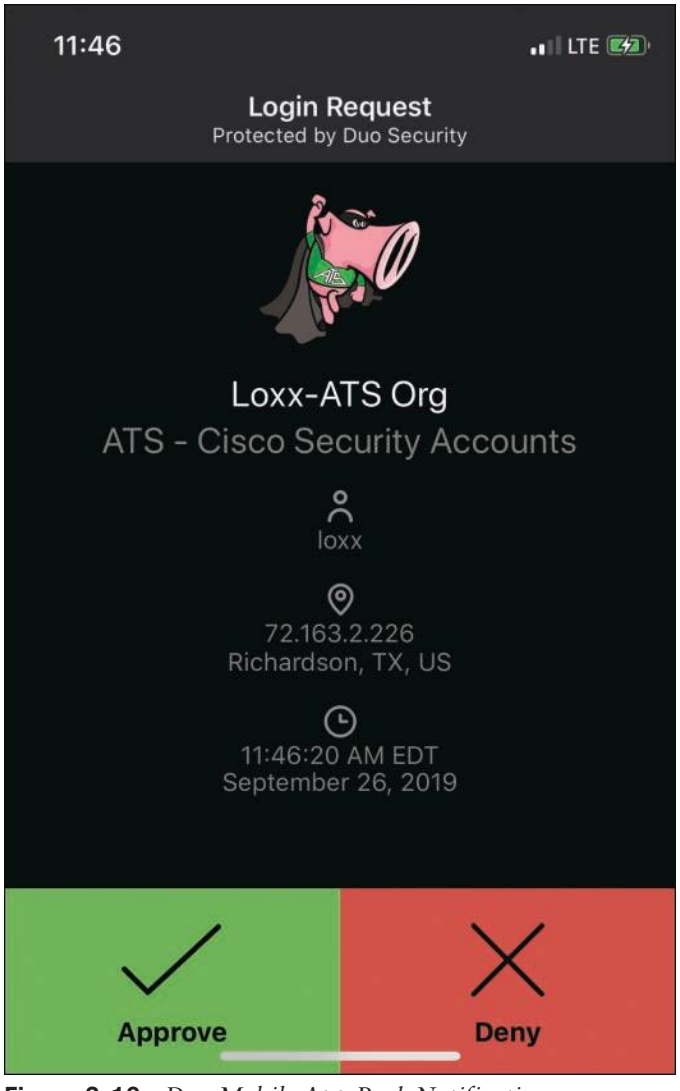


Figure 2-10 *Duo Mobile App Push Notification*

Assuming that the user is authorized to use the application, after the user verifies the second factor by clicking the green checkmark button on the Duo Security app, the user is granted access to the web application.

One-Time Password (OTP) Services

There is a specific version of MFA called ***one-time password (OTP)***. In the case of OTP, the "something you have" is a hardware or software token generator. Hardware token generators come in various forms, including credit cards, USB token keys, or key fobs. Soft tokens, also known as virtual tokens, serve the same role as hard tokens but run as software on your computer or mobile device, so you don't have to carry around a separate physical token card.

Both hardware and software token generators leverage an algorithm to create a code that can be used only one time—and usually for a limited time. The IETF standardized on OTP as a way to create these single-use codes in RFC 6238, and it is commonly used in products like Google Authenticator and Authy.

The key thing to understand about OTP services is that each code may only be used one time. The time-based OTPs rotate at a specific interval. Other solutions use the algorithm to generate a list of codes that must be used in order and only one time. Thanks to the adoption of RFC 6238 for most major online services, including Facebook and Amazon, time-based OTPs have become much more common.

Figure 2-11 shows a screenshot of a mobile application called Authy, which is an RFC 6238–compliant soft token application (see https://authy.com/). It is safe to show you this screenshot of the application because the code displayed on the screen was valid for only 30 seconds, and as you can see in the figure, only 5 seconds remained when the screenshot was taken.

NOTE The terms MFA and OTP have blended in today's world and are often used interchangeably, but they are, in fact, distinct. OTP is a form of MFA, but not all MFA uses OTP.

Smart Cards

Another type of identity store that Cisco ISE supports is smart cards with embedded integrated circuits. A smart card can be used as a form of identification to authenticate for secure network access control systems. Smart cards come in variety of formats, including bank credit cards (commonly called *chip cards*), driver's licenses, and hospital patient ID cards.

A ***Common Access Card (CAC)*** is a particular type of identification badge smart card. You can insert a CAC into a card reader and then enter your PIN. Again, you have to "have something" (the CAC) and "know something" (the PIN) for the authentication to be valid.

Smart cards store a set of X.509 certificates and use a ***public key infrastructure (PKI)*** to store the encrypted digital certificates for the authentication process.

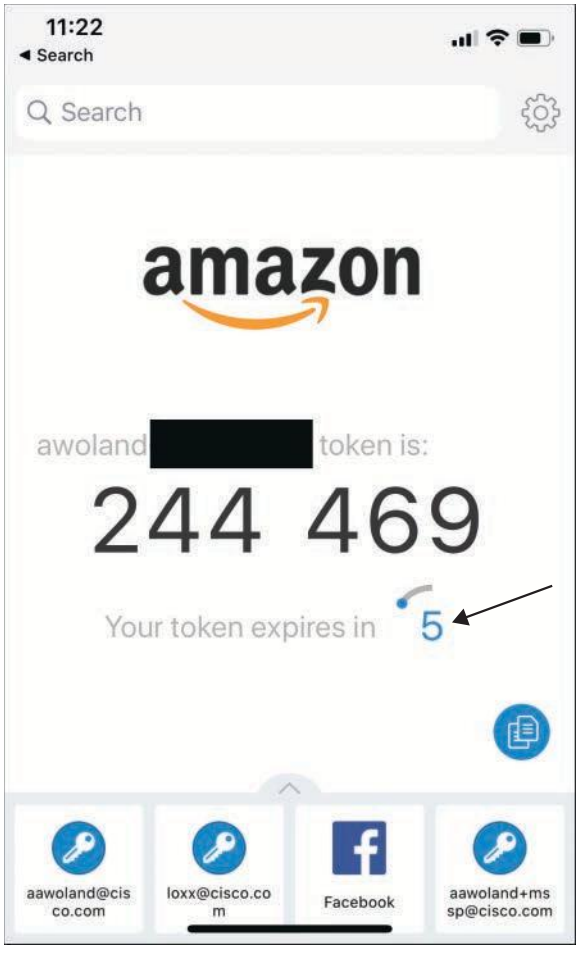


Figure 2-11 *Authy One-Time Password Soft Token App*

Certificate Authorities

In general terms, a *certificate* is an electronic document designed to identify an entity (such as a device, a user, or a website) that uses a digital signature to validate a document. These identity documents are used within X.509, which is a standard for PKI that specifies the format for public key certificates, certificate authority hierarchy, revocation checking, and more.

To understand the topic of certificates, it can be helpful to consider a driver's license example. When you take a driving exam and pass the test, you are given a driver's license that is "signed" by a recognized and trusted authority, such as a state agency known as the Department of Motor Vehicles (DMV). The authority that issues your driver's license can use various methods to sign your license to indicate that it came from a trusted source and is valid and trusted. For example, a hologram may be embedded on the driver's license to make it difficult to duplicate or forge the license.

An X.509 certificate is very much like a driver's license. It is used to identify an entity such as a user, a device, or an application. The recognized authority signs the certificate; with X.509, this authority is a known trusted ***certificate authority (CA)***.

The comparison between X.509 certificates and driver's licenses does not end here. What happens when a police officer stops a speeding driver? The officer asks the driver to provide a valid driver's license, which the police officer uses to perform the following tasks:

Step 1. Validate that the driver's license appears to be a real document and not a forgery.

Step 2. Ensure that the picture does in fact look like the person in the driver's seat.

Step 3. Ensure that the driver's license is valid and not expired.

Step 4. Verify via computer query that the driver's license has not been suspended or revoked based on past offenses.

Key Topic

X.509 certificates work in a very similar fashion. Think of an authentication server, such as Cisco ISE, as the police officer. When the authentication server is presented with an identity certificate, it performs the following tasks (at a minimum):

Step 1. Ensure that the digital certificate was issued and properly signed by a known and trusted certificate authority.

Step 2. Ensure that the certificate is not expired.

Step 3. Ensure that the certificate has not been revoked by the certificate authority.

Step 4. Ensure that the client has provided proof of possession.

The following sections look at the details of these four steps.

Has the Digital Certificate Been Signed by a Trusted CA?

When an identity certificate is presented, you must be able to trust the signing authority. Unless trust is established, the authentication server cannot validate the identity certificate.

The signing of a certificate involves two requirements. The first requirement is that the certificate must be presented in the proper format, according to the X.509 standard. If the certificate is not formatted properly, the authentication server discards the identity certificate and terminates the authentication process. The second requirement is that the trusted certificate authority's ***public certificate*** must be stored in the trusted certificates database of the authentication server, and that certificate must be trusted for purposes of authentication.

With Cisco ISE, the Trust for Client Authentication checkbox needs to be selected for the trusted certificate, as shown in Figure 2-12.

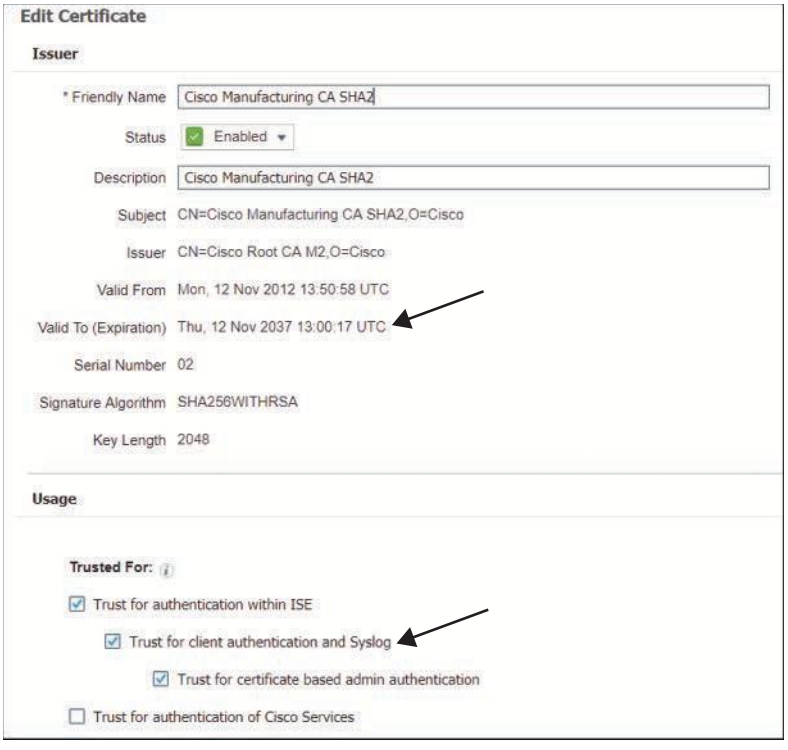


Figure 2-12 *Trust for Client Authentication Checkbox*

Not only does Cisco ISE trust the certificates that have been signed by the certificate authority, it must trust those certificates for the specific required use case, such as client authentication, as in this example. If a client presents a certificate that has not been signed by a certificate authority that is trusted for client authentication, the authentication fails immediately. This authentication process is comparable to someone entering the wrong password for a user account.

Has the Certificate Expired?

Just as a driver's license or a passport has an issued date and an expiration date, a certificate has validity dates. A certificate lists the "valid from" date, which is the starting date on which the certificate is first valid. The "valid to" date is the date after which the certificate will no longer be valid (see Figure 2-12).

After a certificate expires, it has to be renewed with the trusted certificate authority. An authentication server then examines the dates to ensure that the identity certificate is valid for the date and time that the authentication request is processed.

When working with certificates, it is critical to synchronize your network infrastructure, including the certificate authority used in creating new certificates and the authentication servers used in processing certificates for authentication. The best practice is to use ***Network Time Protocol (NTP)*** (UDP port 123) with a trusted and authenticated network time server.

Many of the issues involved with creating, implementing, and processing certificates involve time being out of sync, such as when the certificate "valid from" date and timestamp are

not yet valid (for example, when a certificate is presented for authentication on September 9, 2020, at 10:13 a.m., but its "valid from" date is listed as September 10, 2020, at 5:13 p.m.). Such issues typically arise because the certificate authority's clock was not in sync when the certificate was originally generated.

Has the Certificate Been Revoked?

Let's return to the earlier driver's license example. Remember that an officer can look at your driver's license and determine whether the expiration date is still valid, but the license itself doesn't tell the officer whether the driver's license has been revoked.

In order to verify that the driver's license has not been revoked, the police officer has to run your driver's license number against the DMV database. When the police officer runs your driver's license number, say that he finds out that you have received many speeding tickets, and your driver's license has been revoked. In this case, he immediately arrests you for driving on a revoked license.

Key Topic

When presented with a certificate for authentication, the authentication server has the same capability as the police officer to verify that the trusted certificate authority has not revoked the identity certificate. Every certificate authority should also have a service to publish a list of certificates that have been revoked. As discussed in the following sections, there are two primary ways to check for certificates that have been revoked: using a certificate revocation list and using Online Certificate Status Protocol (OSCP).

Certificate Revocation List

A certificate authority publishes a ***certificate revocation list (CRL)***, which is basically a signed list of the revoked certificate serial numbers. Any device or application can check the CRL to validate whether a certificate has been invalidated. This file is stored on a website and can be periodically downloaded and stored locally on the authentication server. When a certificate is presented for authentication, the authentication server examines the CRL to see if the client's certificate serial number is on the list; if it is, it means the certificate has been revoked.

Using a CRL involves some security concerns that you need to be aware of. Different CAs update their CRLs at different frequencies. If a CRL is updated only once a week, there is a possibility that the CRL will not include a recently revoked certificate.

Another issue is the frequency at which the authentication server downloads the CRL from the certificate authority. Again, if the CRL is not updated locally on the authentication server, then the list of certificate serial numbers that are revoked could be stale.

Online Certificate Status Protocol (OCSP) can be used to avoid these security concerns.

Online Certificate Status Protocol

Using ***Online Certificate Status Protocol (OCSP)*** is the preferred method for revocation checks in most environments today because it provides near-real-time updates. OCSP allows the authentication server to send a real-time request that is similar to an HTTP web request to the service running on the CA or to another device and check the status of the certificate immediately. Using OCSP for revocation checks could be compared to a police officer using a computer in her squad car to a live lookup on the DMV's database. If the certificate has been revoked, access is denied.

Figure 2-13 shows the configuration screen for a certificate authority in ISE. The administrator can configure a location to check for OCSP and/or the CRL when a certificate is signed by a particular CA or its subordinates.



Figure 2-13 *Certificate CRL and OCSP Configuration*

Identity Source Sequences

You can combine multiple identity sources into an ordered list called an ***identity source sequence***, which can then be processed for user and endpoint authentication and authorization. The identity source sequence can then process the identity stores in the order in which they are defined, in top-to-bottom succession. Figure 2-14 shows an example of an ISE identity source sequence configuration.

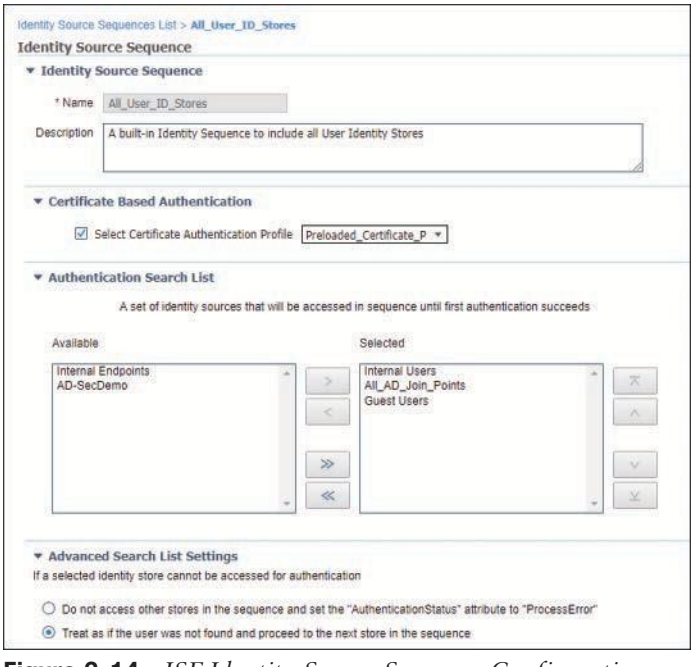


Figure 2-14 *ISE Identity Source Sequence Configuration*

Special Identity Sources

Two identity sources can only be used for ISE's guest services: SAML identity providers (IdPs) and social login. These identity stores are only used for authentication to web authentication and sponsor portals hosted on ISE.

SAML IdPs

Beginning in Version 2.6, ISE provides the ability to authenticate a user by leveraging Security Assertion Markup Language (SAML) IdPs. SAML is an XML-based open standard for authentication (AuthN) and authorization (AuthZ) for web-enabled single sign-on (SSO).

Some noteworthy definitions go along with SAML:

- **SAML service provider (SP):** The SAML SP is the application or service that is being accessed (for example, the ISE web portal). AuthN and AuthZ are required to access the SAML SP.
- **SAML assertion:** Much like a web cookie, a SAML assertion contains the identity and attributes of the user and the user's authorization for the service, which is passed from the IdP to the SP.
- **SAML IdP:** The IdP is a policy engine that determines authorization and issues SAML assertions. The IdP is the identity provider against which ISE is authenticating.
- **User agent:** The user agent is the web browser or app that is requesting access to the service on behalf of the human user.

Using a SAML authentication source on a web portal allows users to authenticate with single sign-on to sponsor and guest portals. If a portal is configured to use SAML for authentication, then this is the only type of authentication available because SAML providers cannot be used within an identity source sequence.

NOTE As this book went to press, SAML was not an option for the administrative login to ISE.

Social Login

Social login is a special identity source for use with guest portals only. With social login, a guest user can click a link such as "Log in with Facebook" and leverage his Facebook credentials to gain guest access.

As of Version 2.6, Facebook was the only social media source ISE supported for guest logins.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 2-2 lists these key topics and the page number on which each is found.

Key Topic

Table 2-2 Key Topics

Key Topic Element	Description	Page Number
Paragraph	ISE can leverage internal and external ID stores	21
Paragraph	Examples of external ID stores	23
Paragraph	Description of AD domains	24
Paragraph	CA services must publish list of revoked certs	31
Paragraph	Checking the validity of a certificate	33

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

identity, credential, identity store, identity source, identity source sequence, internal user database, user identity group, internal endpoint database, external identity store, Active Directory (AD), Lightweight Directory Access Protocol (LDAP), Remote Authentication Dial-In User Service (RADIUS), Certificate Authentication Profile (CAP), multifactor authentication (MFA), two-factor authentication (2FA), one-time password (OTP), Common Access Card (CAC), public key infrastructure (PKI), certificate authority (CA), public certificate, registration authority (RA), Network Time Protocol (NTP), certificate revocation list (CRL), Online Certificate Status Protocol (OCSP)

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What are the three possible factors for multifactor authentication?
2. Name two identity stores that are used only for guest authentications.
3. What combines multiple identity sources into an ordered list for authentication?
4. What authentication source is used for more than 90% of production ISE deployments?
5. What is the term for the representation of a user or device?

This page intentionally left blank



CHAPTER 3

Extensible Authentication Protocol (EAP) over LAN: 802.1X

This chapter covers the following topics:

- EAP types
- Supplicant, authenticator, and authentication server
- Supplicant options
- Network access devices

In the early 2000s, the IEEE standardized 802.1X as a solution for port-based network access control. It was predicted to revolutionize networking, and no device would be able to plug in and communicate on a network without the users identifying themselves and being authorized again. It took about a decade for 802.1X to really start catching on. This book went to press in 2020, nearly two decades after the IEEE standardized 802.1X, and network admins are still rather fearful of it—but they do not need to be!

There are three fundamental components of 802.1X: the supplicant, the authenticator, and the authentication server. This chapter explains those components, along with critical elements of an 802.1X solution, such as the different EAP types that can be used.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 3-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 3-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
EAP over LAN (802.1X)	1, 3, 4, 7, 9, 10
Supplicant Options	2, 5, 6, 8

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is true?
   a. The authenticator decides if the supplicant is allowed on the network.
   b. EAP communication occurs between the supplicant and the authentication server.
   c. The supplicant uses RADIUS to communicate the user's identity to the authentication server.
   d. The authenticator uses EAP to send the user's credentials to the authentication server.
2. Which supplicant is capable of EAP chaining?
   a. Windows native supplicant
   b. Cisco AnyConnect NAM
   c. Cisco Secure Services Client (CSSC)
   d. Odyssey client
3. What is the purpose of the outer identity?
   a. The outer identity is used for dual-factor authentication, such as use of a username and password combined with a one-time-password (OTP).
   b. The outer identity provides a mechanism to modify the actual identity of the end user or device to allow for identity spoofing.
   c. The outer identity provides a mechanism to authenticate the identity of the endpoint during the tunnel establishment phase.
   d. The outer identity represents the machine, and the inner identity represents the user during EAP chaining.
4. Which AAA protocol can IEEE 802.1X use to communicate the EAP identity to the authentication server?
   a. RADIUS
   b. SAML
   c. TACACS+
   d. OAuth2

5. What is required for a supplicant to form a TLS tunnel with the authentication server within which the EAP transaction will occur?
   a. The name of the authentication server must be resolvable by both forward and reverse DNS.
   b. The supplicant must be using a certificate issued to it by the authentication server.
   c. The name of the authentication server must be resolvable by forward DNS, but reverse DNS is not required.
   d. The supplicant must trust the EAP certificate of the authentication server.
6. What is the name of the secure cookie used with EAP-FAST that can be used in lieu of or in addition to a certificate?
   a. Protected Password File (PPF)
   b. Shadow Credential File (SCF)
   c. Private Authorization Credential (PAC)
   d. Protected Access Credential (PAC)
7. With what types of authentication types can MS-CHAPv2 be used when the identity store has an LDAP connection to Active Directory?
   a. MS-CHAPv2 cannot be used with ISE when the identity store is LDAP.
   b. Machine authentication can be used.
   c. User authentication can be used.
   d. Both user and machine authentication can be used.
8. Which of the following is an untrue statement regarding machine authentication?
   a. Windows machine authentication can occur with any tunneled EAP authentication.
   b. Windows machines can authenticate only if the tunneled EAP type is PEAP.
   c. Windows machines can be authenticated using certificates and EAP-TLS.
   d. Windows machines can be authenticated with MS-CHAPv2.
9. What are the three main components of IEEE 802.1X?
   a. Agent, broker, and authentication server
   b. Supplicant, authorizer, and authorization server
   c. Authentication server, supplicant, and authenticator
   d. EAP, RADIUS, and TLS
10. Which of the following are EAP types? (Choose two.)
    a. EAP chains
    b. Tunneled EAP
    c. Native Extensible Authentication Protocol
    d. Machine EAP

Foundation Topics

Extensible Authentication Protocol

Extensible Authentication Protocol (EAP) is an authentication framework that defines the transport and usage of identity credentials. EAP encapsulates the usernames, passwords, certificates, tokens, one-time passwords, and so on that a client is sending for purposes of authentication.

EAP has become the de facto standard of authentication protocols. It is used for many different authentication methods, including over VPNs, but most importantly, IEEE 802.1X is a ratified standard for using ***EAP over LAN (EAPoL)***.

EAP over LAN (802.1X)

IEEE 802.1X (commonly referred to as Dot1x) is defined as a standard for port-based network access control for local-area and metropolitan-area networks. The standardization of a network-based authentication framework was the catalyst for all identity-based networking that we see today. There are three main components to 802.1X: the supplicant, the authenticator, and the authentication server, as illustrated in Figure 3-1 and described in Table 3-2.

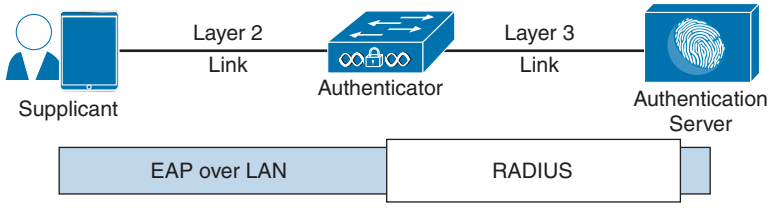


Figure 3-1 *Components of 802.1X*

Key Topic

Table 3-2 802.1X Components

Component Name	Description
Supplicant	Software on the endpoint (which the IETF also calls a *peer*) that communicates with EAP at Layer 2. This software responds to the authenticator and provides the identity credentials with the EAP communication.
Authenticator	The network device that controls physical access to the network, based on the authentication status of the endpoint. The authenticator acts as the middleman, encapsulating Layer 2 EAP communication from the supplicant in RADIUS, directed at the active authentication server. The most common authenticators with a Cisco ISE deployment are LAN switches and wireless LAN controllers (WLCs). Cisco ISE refers to these authenticators generically as *network access devices* (*NADs*).
Authentication server	The server that is performing the authentication of the client. The authentication server validates the identity of the endpoint and provides the authenticator with a result, such as accept or deny. Cisco ISE is an authentication server.

While reviewing Figure 3-1 and Table 3-2, keep in mind that the authenticator (that is, the switch or WLC) acts as the middleman or proxy. The actual EAP identity exchange and authentication occurs between the supplicant and the authentication server. The switch or WLC has no idea of what EAP type is in use or whether the user's credentials are valid; it simply encapsulates the unmodified EAP frame within the RADIUS packet sent to the authentication server and authorizes the port if the authentication server tells it to.

Therefore, EAP authentication is completely transparent to the authenticator. Figure 3-2 displays the communication through a successful authentication.

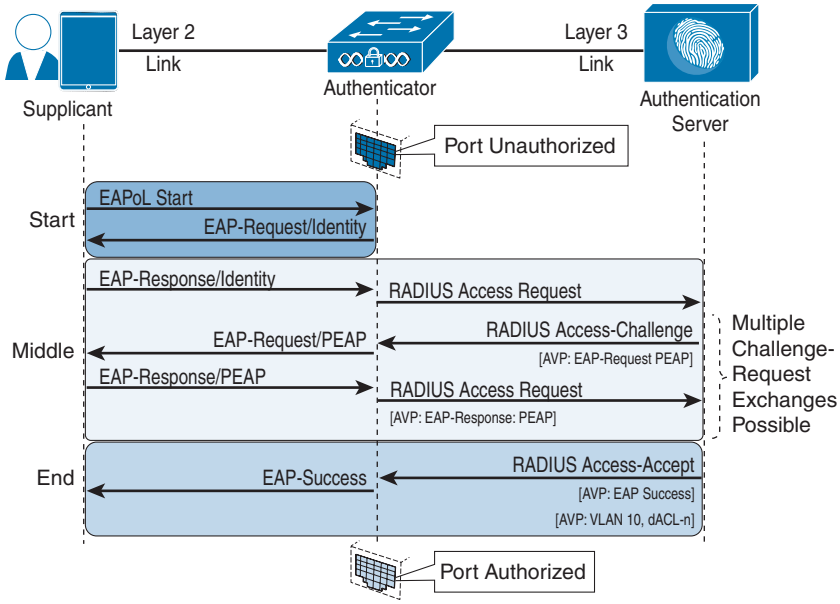


Figure 3-2 *Authentication Communication*

The authentication can be initiated by either the authenticator or the supplicant. The authenticator initiates authentication when the link state changes from down to up or periodically, as long as the port remains up and unauthenticated. The switch sends an EAP-Request/Identity frame to the endpoint to request its identity. Upon receipt of the frame, the client's supplicant responds with an EAP response/identity frame. However, enhancements were made to allow the supplicant to trigger the authenticator to request an identity by sending an EAPoL-Start message at any time. This enhancement provided for a much better end-user experience with 802.1X.

EAP Types

There are many different EAP types, each with benefits and downsides. The EAP type defines the authentication mechanism to be used with EAP, and the type's name typically indicates the mechanism. Most of the many EAP types are not discussed in this book due to lack of adoption or lack of inclusion in the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam blueprint (for example, EAP-Kerberos).

The EAP types can be broken down into two categories: *native EAP types* and *tunneled EAP types*. A tunneled EAP type simply uses a ***native EAP*** inside a TLS tunnel between the supplicant and the authenticator.

Native EAP Types (Non-Tunneled EAP)

Figure 3-3 shows a graphical representation of native EAP types, which are described in more detail in the list that follows:

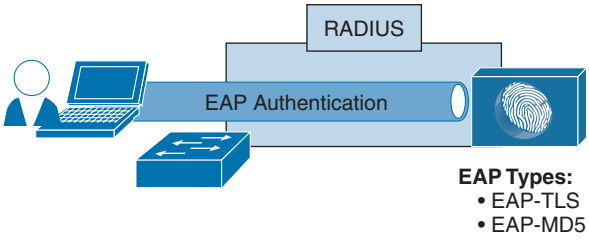


Figure 3-3 *Native EAP Types*

- **EAP-MD5:** This EAP type uses a message digest algorithm to hide the credentials in a hash. The hash is sent to the server, where it is compared to a local hash to see if the credentials were accurate. However, EAP-MD5 does not have a mechanism for mutual authentication. This means the server is validating the client, but the client does not authenticate the server (that is, it does not check to see if it should trust the server).

  EAP-MD5 is common on IP phones, and some switches may send MAC Authentication Bypass (MAB) requests using EAP-MD5.

- **EAP-TLS:** This EAP type uses Transport Layer Security (TLS) to provide the secure identity transaction. TLS is very similar to SSL, and it handles the encryption between your web browser and a secure website. EAP-TLS has the benefit of being an open IETF standard and is considered universally supported.

  EAP-TLS uses X.509 certificates and provides the ability to support mutual authentication, where the client must trust the server's certificate and vice versa. It is considered among the most secure EAP types, since password capture is not an option; the endpoint must still have the private key.

NOTE EAP-TLS is quickly becoming the EAP type of choice when supporting BYOD in the enterprise.

- **EAP-MS-CHAPv2:** With this EAP type, the client's credentials are sent to the server encrypted within an MS-CHAPv2 session. This allows for simple transmission of username and password (or even computer name and computer password) to the RADIUS server, which in turn authenticates them to Active Directory.

NOTE Cisco ISE currently supports EAP-MS-CHAPv2 only within tunneled EAP and not as native EAP.

- **EAP-GTC:** Cisco created EAP-Generic Token Card (GTC) as an alternative to MS-CHAPv2 that allows generic authentication to virtually any identity store, including one-time password (OTP) token servers, LDAP, and Novell eDirectory.

NOTE Cisco ISE currently supports EAP-GTC only when it is inside tunneled EAP and not as native EAP.

Tunneled EAP Types

The native EAP types send their credentials immediately, while ***tunneled EAP*** types form an encrypted tunnel first and then transmit the credentials within that tunnel. Figure 3-4 shows a graphical representation of tunneled EAP, which is described in more detail in the sections that follow.

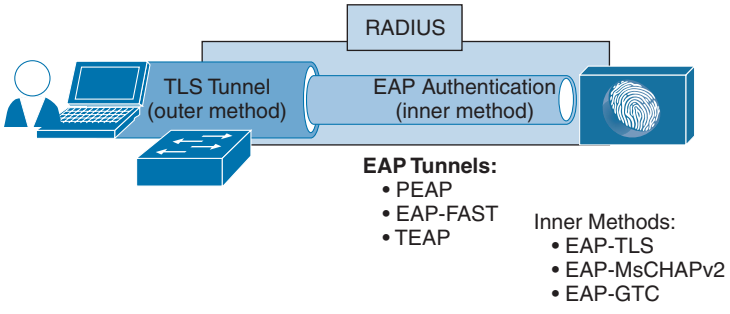


Figure 3-4 *Tunneled EAP Methods*

Key Topic

NOTE It is important to understand that tunneled EAP methods work very similarly to the way a secure tunnel is formed between a web browser and a secure website (such as a banking site). The tunnel is formed first, normally using only the certificate of the server (one-way trust), and then the user enters banking login credentials within that secure tunnel.

PEAP

Protected EAP (PEAP) was originally proposed by Microsoft. This EAP tunnel type has quickly become the most popular and widely deployed EAP method in the world. PEAP forms a potentially encrypted TLS tunnel between the client and server, using the X.509 certificate on the server in much the same way the SSL tunnel is established between a web browser and a secure website.

After the tunnel has been formed, PEAP uses another EAP type as an "inner method"—authenticating the client using EAP within the outer tunnel.

Three inner methods are supported for PEAP:

- **EAP-MS-CHAPv2:** Using this inner method, the client's credentials are sent to the server encrypted within an MS-CHAPv2 session. This is the most common inner method, as it allows for simple transmission of username and password (or even computer name and computer-password) to the RADIUS server, which in turn authenticates them to Active Directory.
- **EAP-GTC:** Cisco created this inner method as an alternative to MS-CHAPv2 to allow generic authentication to virtually any identity store, including OTP token servers, LDAP, and Novell eDirectory.
- **EAP-TLS:** While rarely used and not widely known, PEAP is capable of using EAP-TLS as an inner method.

EAP-FAST

Flexible Authentication via Secure Tunneling (FAST) is very similar to PEAP. Cisco created FAST as an alternative to PEAP to allow for faster re-authentication and support faster wireless roaming.

Just like PEAP, FAST forms a TLS outer tunnel and then transmits the client credentials within that TLS tunnel. Where FAST differs from PEAP is in the ability to use Protected Access Credentials (PACs). A PAC can be thought of as a secure cookie that is stored locally on the host as proof of successful authentication. The PAC is used for rapid session resumption when a client roams or is disconnected from the network temporarily and to enable EAP chaining (which is discussed later in this chapter).

FAST uses a native EAP type as an inner method to authenticate the client using EAP within the outer tunnel. Three inner methods are supported for FAST:

- **EAP-MS-CHAPv2:** Using this inner method, the client's credentials are sent to the server encrypted within an MS-CHAPv2 session. This is the most common inner method, as it allows for simple transmission of username and password (or even computer name and computer-password) to the RADIUS server, which in turn authenticates them to Active Directory.
- **EAP-GTC:** Cisco created this inner method as an alternative to MS-CHAPv2 to allow generic authentication to virtually any identity store, including OTP token servers, LDAP, and Novell eDirectory.
- **EAP-TLS:** EAP-FAST is capable of using EAP-TLS as an inner method. This has become quite popular with EAP chaining.

EAP-TTLS

EAP Tunneled Transport Layer Security (EAP-TTLS) is another tunneled EAP type developed mainly by Funk Software (which has since been acquired by Juniper). It is very similar to PEAP and EAP-FAST, as it creates a TLS outer tunnel first and then performs the authentication within that TLS tunnel.

There are some benefits to using TTLS over PEAP: It provides some enhancements to the simplicity of the setup, and it has enhancements to the TLS tunnel's capability to resume a session.

TTLS uses a native EAP type as an inner method—authenticating the client using EAP within the outer tunnel.

EAP-TTLS support was added in ISE Version 2.0. Six inner methods are supported for EAP-TTLS:

- **PAP/ASCII**: This method is not an EAP type. It simply uses native Password Authentication Protocol (PAP) to negotiate the authentication and transmits the username and password in plaintext within the encrypted EAP-TTLS outer tunnel.
- **CHAP**: This method is not an EAP type but uses native Challenge Handshake Authentication Protocol (CHAP) to negotiate the authentication with added replay attack prevention. The username is sent in plaintext within the EAP-TTLS tunnel, and a calculated hash of the password is sent as well (and the actual password is never sent).
- **MS-CHAPv1**: Again, this is not an EAP type but a native authentication protocol. Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) is an enhancement to CHAP that enhances the protocol with more functions, such as password change mechanisms, defined failure codes, and retry mechanisms.
- **MS-CHAPv2**: This is a newer version of MS-CHAP with additional security enhancements. Again, this is not an EAP type but a native authentication protocol.
- **EAP-MD5**: This EAP type uses a message digest algorithm to hide the credentials in a hash. The hash is sent to the server, where it is compared to a local hash to see if the credentials were accurate. However, EAP-MD5 does not have a mechanism for mutual authentication. This means the server is validating the client, but the client does not authenticate the server (that is, it does not check to see if it should trust the server). This may be acceptable to some implementations—especially if the TTLS outer method has already performed the mutual authentication.
- **EAP-MS-CHAPv2**: Using this inner method, the client's credentials are sent to the server encrypted within an MS-CHAPv2 session. This is the most common inner method, as it allows for simple transmission of username and password (or even computer name and computer-password) to the RADIUS server, which in turn authenticates them to Active Directory.

NOTE Many believe that the main benefit of supporting EAP-TTLS is simply that the eduroam initiative (www.eduroam.org) mandates the use of EAP-TTLS. The eduroam initiative is designed to allow federating college student identities to roam between college campuses using their same credential.

"One EAP Type to Rule Them All": TEAP

Tunnel EAP (TEAP), defined in RFC 7170, was designed as the EAP type to end the EAP wars. Every vendor seemed to have its own EAP type and tried to claim superiority based on that EAP. Microsoft had PEAP, Cisco had FAST, Funk/Juniper had EAP-TTLS, and so on.

The IETF working group that designed TEAP had active members from many vendors in this space, with a heavy effort being shared between Cisco and Juniper. All the benefits of FAST, TTLS, and PEAP were put into this single tunneled EAP method, as well as more enhancements based on what the industry had learned. In addition to the standard bits we expect to find in a protocol like PEAP or TTLS, TEAP adds other unique capabilities, such as the following:

- EAP chaining (covered later in this chapter)
- Certificate provisioning within the outer tunnel
- Certificate renewal within the outer tunnel
- Distribution of the list of trusted EAP servers' certificates to the clients, which can fix a lot of today's networking challenges
- Channel binding, which enhances security by allowing the client and server to verify that they are communicating over the same switch, preventing man-in-the-middle attacks

TEAP support was added to ISE in Version 2.7. This was part of a joint Microsoft/Cisco initiative to improve wired/wireless security and usability. TEAP support was added to Windows 10 at the same time.

Just like the other tunneled EAP types, TEAP forms a TLS outer tunnel and then transmits the client credentials within that TLS tunnel. While TEAP is capable of using PACs, using PACs is not required to achieve fast roaming and resume sessions.

While RFC 7170 for TEAP defines a lot of advancements for network authentication, not all of those functions made it into Windows 10 or ISE in the first release. As this book went to press, neither ISE nor Windows 10 supported any of the certificate-provisioning functions of TEAP; they also did not support the distribution of the list of EAP server certificates for the client to trust. Product development is an iterative process, however, and you should expect both Microsoft and Cisco to continue to enhance the TEAP support.

As this book went to press, Windows 10 and ISE supported two inner methods: EAP-MS-CHAPv2 and EAP-TLS.

All Tunneled EAP Types

Tunneled EAP involves the concept of inner and outer identities. The inner identity is easier to explain. It is the user's or device's actual credentials, sent with the native EAP or authentication protocol. The outer identity, which is typically set to anonymous, is the identity that is used between the supplicant and the authentication server for the initial TLS tunnel setup.

Cisco ISE is able to read that outer identity and use it to help make identity store selection decisions. Put simply, that outer identity may contain information (such as the domain name) that tells Cisco ISE to submit the credentials to Active Directory or LDAP or some other identity store.

Most supplicants hide this option from the end user, and only administrators see the outer identity. However, one supplicant that does expose it to the end user is the native Android

supplicant (see Figure 3-5). Note that the Android supplicant refers to the outer identity as the "anonymous identity"—an amusing oxymoron.

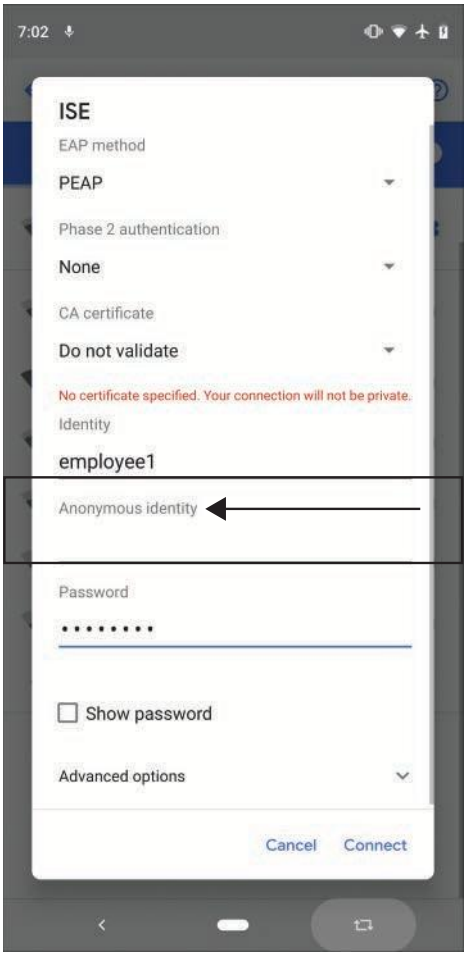


Figure 3-5 *Android Supplicant Exposing the Outer Identity*

Table 3-3 provides a comparison of tunneled EAP types.

Key Topic

Table 3-3 Tunneled EAP Types Comparison Table

Feature	EAP-FASTv2 (proprietary)	EAP-PEAP (RFC draft)	EAP-TTLS (RFC 5281)	EAP-TEAP (RFC 7170)
User and machine EAP chaining	✓	—	—	✓
FAST reconnect with PAC file	✓	—	—	✓*
Fast reconnect with server	—	—	✓	✓

Feature	EAP-FASTv2 (proprietary)	EAP-PEAP (RFC draft)	EAP-TTLS (RFC 5281)	EAP-TEAP (RFC 7170)
Channel binding to prevent MITM	✓	—	—	✓*
Protection of credential transmission with TLS	✓	✓	✓	✓
Multiple inner method choices	✓	✓	✓	✓
Mandated use by eduroam initiative	—	—	✓	—
Certificate provisioning	—	—	—	✓*
Distribution of EAP server trust list	—	—	—	✓*
Certificate renewals	—	—	—	✓*

* Not supported in ISE as of Version 2.7 but included in the RFC standard.

EAP Authentication Type Identity Store Comparison

The appropriate EAP type depends on the operating system, 802.1X supplicant, and supported backend credential database or identity store. Table 3-4 compares the EAP authentication type identity store.

Table 3-4 EAP Authentication Type Identity Store Comparison

Authentication Type (native EAP or inner method of TTLS, PEAP, FAST, or TEAP)	Identity Store			
	Active Directory	LDAP	Internal DB	RADIUS Token, OTP, RSA
EAP-GTCPAP / ASCII	✓	✓	✓	✓
MS-CHAPv1/2 (TTLS only) EAP-MS-CHAPv2LEAP	✓	—	✓	—
EAP-TLS* (native)EAP-TLS* (inner method)	✓	✓	—	—
EAP-MD5CHAP	—	—	✓	—

* TLS authentication only validates the certificate, but the authorization function may use these identity stores.

Network Access Devices

Cisco ISE refers to the authenticator role as a network access device (NAD). The NAD serves multiple roles. It is an authenticator for 802.1X, and it proxies EAP communications from a supplicant to the authentication server. A NAD is also commonly referred to as a policy enforcement point. The NAD is responsible for enforcing whatever authorization result it receives from the authentication server (for example, Cisco ISE).

In simple terms, a NAD is an access-layer device but can be any device that is going to send RADIUS authentication requests to Cisco ISE. Common NAD types include the following:

- Wired Ethernet switch
- Wireless LAN controller (WLC)
- Cisco Adaptive Security Appliance (ASA)
- Cisco Firepower device
- Load balancer
- Software application that uses ISE for AAA

Enforcement types are covered in more detail in other chapters, but for now you just need to know that the following are some examples of common enforcement types:

- Dynamic VLAN (dVLAN) assignment
- Downloadable access control list (dACL)
- Security group tag (SGT)
- Airespace ACL name (for use with WLCs)
- URL redirection

NADs are an incredibly important part of any secure network access design. NADs do much more than simply pass frames at Layer 2. These devices are the security enforcement devices that apply the policy to the end user, and they perform functions that used to be available exclusively in overlay appliances. Such functions include URL redirection and agent discovery of the active server.

These topics are covered in much more depth in many of the other chapters. At this point, we want to stress that these devices are key to the success of an 802.1X deployment. Intelligence at the edge is an absolute requirement to ensure success.

Supplicant Options

The bulk of this chapter is about 802.1X. This topic would not be complete without a discussion of the client side of the authentication flow: the supplicant. As described earlier in this chapter, a supplicant is the software on an endpoint that understands how to communicate with EAP on the LAN. A supplicant must be configured to use credentials—either stored credentials (such as a certificate) or a username and password.

Windows Native Supplicant

The most common supplicant in wired networks is the Windows native supplicant. One reason this supplicant is so popular is that it is built in to the most common business desktop operating system in the world. In addition, it offers one truly beneficial feature that no other supplicant has been able to compete with: It is fully controllable from a central Active

Directory group policy. That fact alone has made this supplicant appear very attractive to desktop managers of corporations all over the world.

In order to use the Windows native supplicant, the service must first be started. It is unintuitively named Wired AutoConfig, and the default state of the service is Manual. You need to change the state to Automatic in order for the supplicant to be enabled at each boot. There is also a WLAN AutoConfig service for the wireless supplicant, and it is set to Automatic by default.

The following pages provide a guided tour of the Windows native supplicant to drive home some of the points made previously surrounding EAP types and inner versus outer methods. If you have a Windows 10 device at your disposal, feel free to follow along:

Step 1. In the Windows services control applet, locate **Wired AutoConfig** in the list of services, as shown in Figure 3-6, and double-click it to enter the service properties.

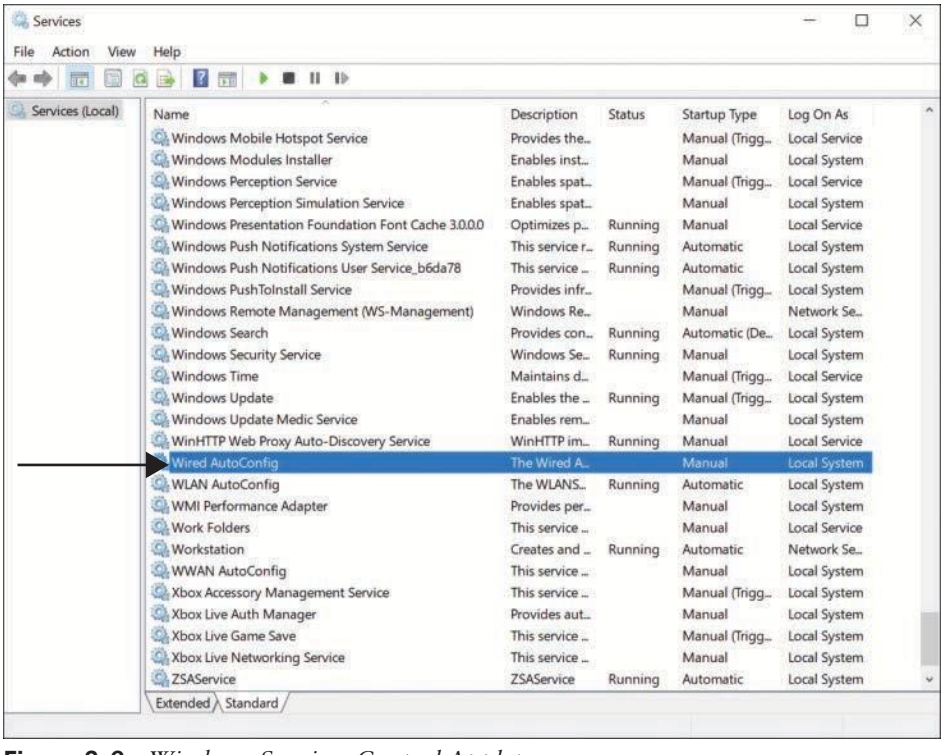


Figure 3-6 *Windows Services Control Applet*

Step 2. Change the Startup Type value to **Automatic**, as shown in Figure 3-7.

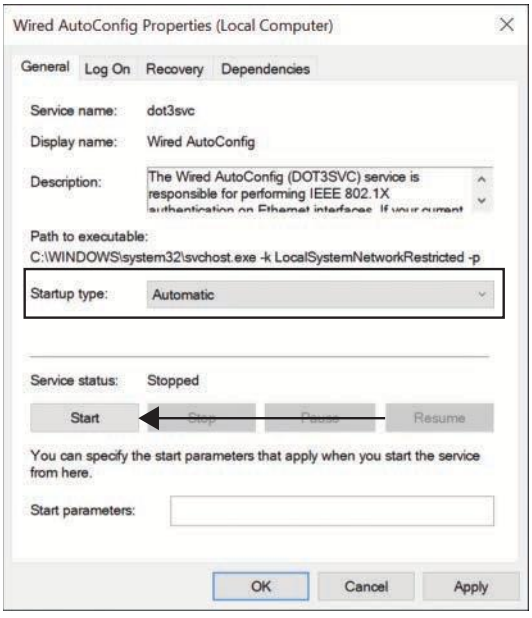


Figure 3-7 *Wired AutoConfig Properties*

Step 3. Click **Start** to enable the supplicant and populate the wired NIC properties page with an authentication tab (see Figure 3-8).

Then navigate to **Control Panel > Network and Internet > Network Connections.**

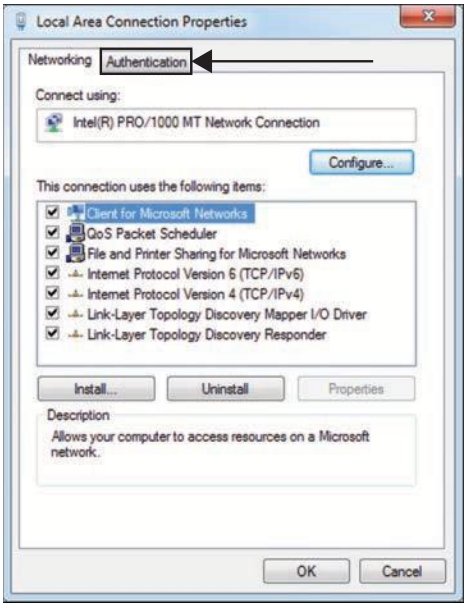


Figure 3-8 *Local-Area Connection Properties*

Step 4. Select the Authentication tab, as shown in Figure 3-9.

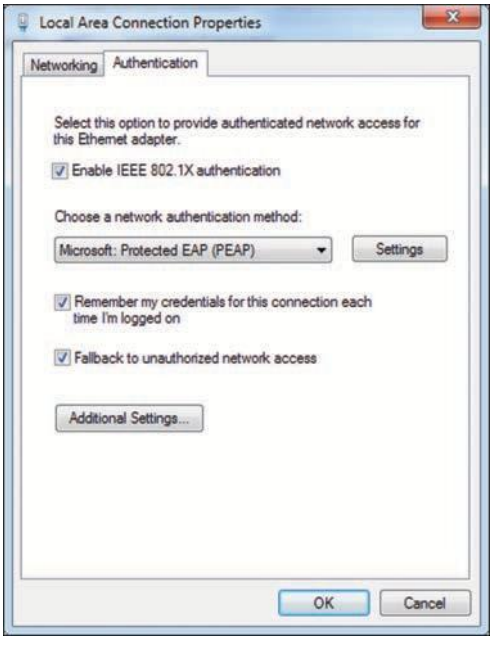


Figure 3-9 *Authentication Tab*

As shown in Figure 3-9, there is a checkbox on the Authentication tab to enable IEEE 802.1X authentication. There is also a checkbox to allow the supplicant to store (remember) the credentials so that the end user will not be prompted for credentials at each network login.

The last checkbox is for falling back to unauthorized network access. This setting basically states that if the network device is not an authenticator (that is, if it does not send EAP identity requests), the network connection should still work. If this setting is unchecked and the user plugs in to a non-802.1X-enabled switch, the connection is treated as if the authentication has failed, and the user cannot access the network. It's a good idea to leave this setting enabled to ensure a positive user experience.

To see additional settings, follow these steps:

Step 1. Click the **Settings** button. The Protected EAP Properties page shown in Figure 3-10 opens.

Select the **Validate Server Certificate** checkbox to require the supplicant to trust the certificate from the RADIUS server (802.1X authentication server) before it forms the secure TLS tunnel. Disabling this setting turns off that level of authentication, and the supplicant forms a tunnel and sends EAP credentials to any RADIUS server.

This dialog also allows you to specify the specific servers that are allowed—not only the certificates to trust. This gives you even more strict security control over which RADIUS servers are allowed to receive the supplicant's credentials, although it is rarely used because of the ever-changing environments at customer sites.

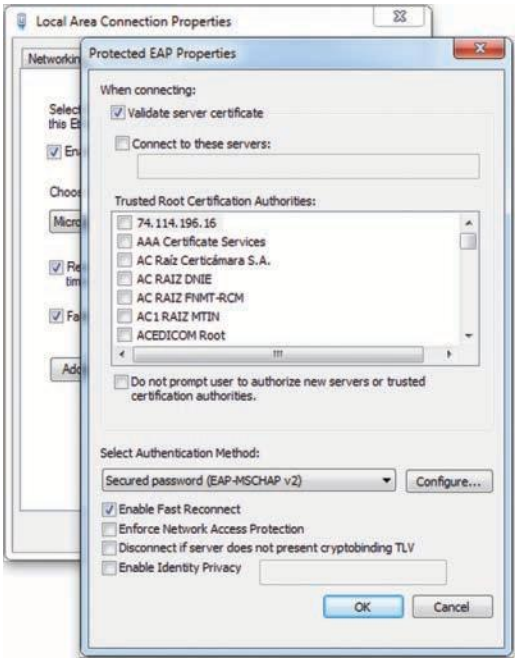


Figure 3-10 *Protected EAP Properties Dialog*

In the Trusted Root Certificate Authorities area of the dialog, you can select which specific root CAs are trusted for authentication. You can also select the checkbox below this area if you do not want to prompt users to authorize new servers or trusted certificate authorities. If you leave this checkbox disabled, users are not prompted to trust certificates that are not explicitly trusted in the list above the checkbox.

In the Select Authentication Method area of the dialog, you can select the inner method for PEAP: either Secured Password (EAP-MSCHAPv2) or Smart Card or Other Certificate (to use EAP-TLS as the inner method for certificate authentication).

Step 2. In the Select Authentication Method area, select **Secured Password (EAP-MSCHAPv2)** and click **Configure**, as the dialog shown in Figure 3-11 appears.



Figure 3-11 *EAP MSCHAPv2 Properties Dialog*

Step 3. Select the checkbox to enable single sign-on and click **OK** to close the EAP MSCHAPv2 Properties dialog.

Step 4. Back in the Protected EAP Properties dialog, set Select Authentication Method to **Smart Card or Other Certificate** and click the **Configure** button to bring up the properties window shown in Figure 3-12.

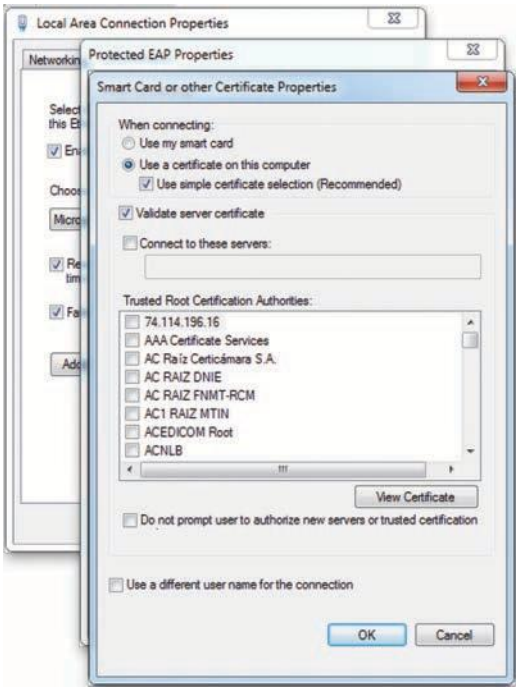


Figure 3-12 *EAP MS-CHAPv2 Properties*

In the Smart Card or Other Certificate Properties dialog, you can select the use of a smart card or a certificate that is stored on the local computer. Along with this choice is a selection for simple certificate selection. A Windows machine may store many different certificates that have multiple purposes. If you enable simple certificate selection, the list of certificates is filtered down to only identity certificates, thereby simplifying the experience of choosing the correct certificate.

As with the outer method PEAP, with the inner method EAP-TLS, you have options to validate server certificates, specify specific servers to trust, select trusted root certificate authorities, and control user prompts to trust new servers. You can also specify a different username for this connection (inner identity).

Step 5. Click **OK** or **Cancel** to return to the Protected EAP Properties dialog.

Step 6. Under Select Authentication Method, once again select **Secured Password (EAP-MSCHAPv2)**, as shown in Figure 3-13. The only other relevant setting on this page is the Enable Identity Privacy checkbox. This option allows you to set the outer identity. When you leave this box unchecked, the outer identity is set to Anonymous.

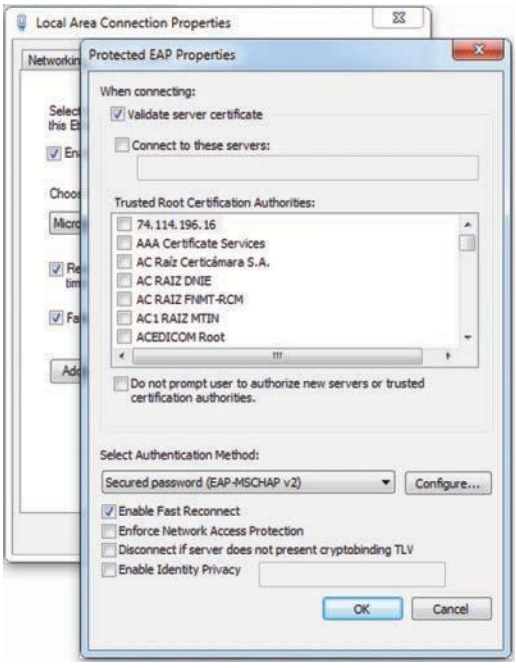


Figure 3-13 *Back to the Protected EAP Properties*

Step 7. Click **Cancel** to return to the Authentication tab, shown in Figure 3-14.

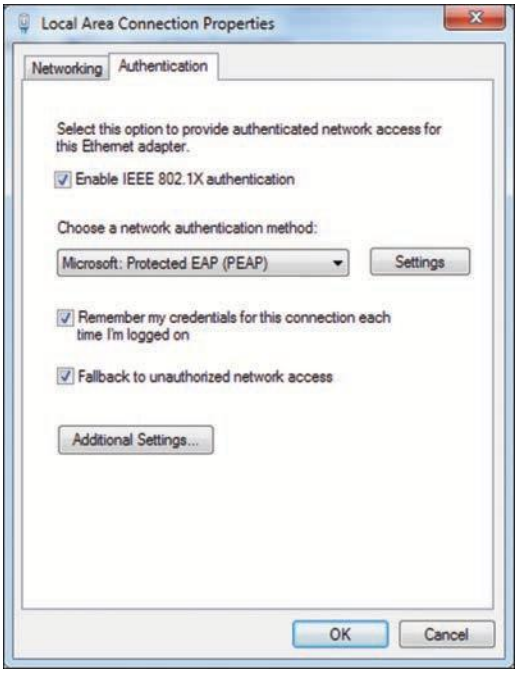


Figure 3-14 *Back to the Authentication Tab*

Step 8. Click the **Additional Settings** button. You now see some additional authentication options for the supplicant, related to mode and sign-on timing (see Figure 3-15).

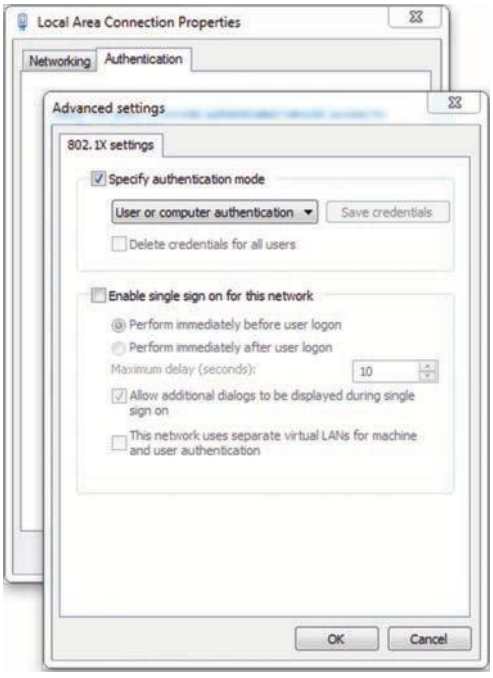


Figure 3-15 *Advanced Settings Tab*

Let us cover these settings in reverse order, starting with the checkbox Enable Single Sign on for This Network. This setting allows you to allow (or disallow) users to log in to the workstation before the 802.1X authentication occurs. You can choose the option Perform Immediately Before User Logon to send the EAPoL-Start message and allow the supplicant to perform the EAP exchange before allowing the user to interact with the workstation interface (that is, before the start button and desktop are displayed to the end user). The Maximum Delay setting puts a timer on how long the supplicant should wait for success before allowing the user to interact with the desktop or log the user off again.

You can select the option Perform Immediately After User Logon to allow the user to interact with the desktop immediately, which could include the ability to respond to prompts where the supplicant is asking the user to enter the username and password. At the bottom of the tab is the option This Network Uses Separate Virtual LANs for Machine and User Authentication. This option will make more sense after you have read the next section. Basically, it forces the supplicant to do an IP release/renew and attempt to get a new IP address when the user logs in.

Now, let's return to the first item of note in the Advanced Settings tab: the Specify Authentication Mode checkbox. The options under this checkbox are User or Computer Authentication, Computer Authentication, User Authentication, and Guest Authentication, as shown in Figure 3-16.

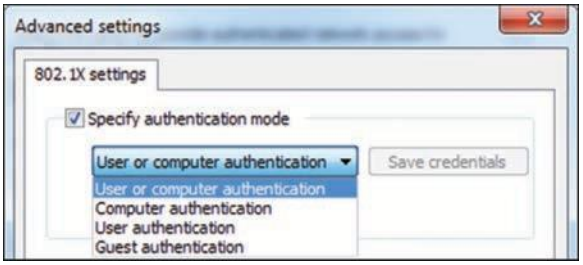


Figure 3-16 *Authentication Mode Choices*

User Authentication

User authentication is the authentication that people commonly think of when discussing 802.1X and network access control. This type of authentication provided the identity credentials of the user to the authentication server and allows for role-based access control to the network. When the 802.1X standard was created, it was important to know who was accessing networks before allowing them network access.

User authentication on a Microsoft Windows device can occur with a username or password, a user-issued certificate, or even a smart card. With Windows devices, there is a separate certificate store for user-issued certificates. Each user who logs in to a Windows workstation has a unique certificate store for authenticating to the network.

Machine Authentication (Computer Authentication)

Key Topic

Microsoft Windows workstations have a very powerful management system, known as Active Directory (AD). AD needs to remain in contact with the computers that it manages for policy updates and other management tasks. When the IEEE 802.1X standard came out, if no user was logged in to the computer, no network access was granted. This concept broke the communication between the AD managed computers, and it also prevented a user's Kerberos authentication from reaching Active Directory, so the credentials could not be passed down to the supplicant itself, especially in instances where the user's password was expired.

This was a lot like a denial of service (DoS). Microsoft quite brilliantly designed a way to solve this problem and prevent this chicken-and-egg scenario. It created multiple states for its supplicant: a machine state and a user state.

Whenever there was no interactive user logged in to the workstation (that is, when no one had pressed Ctrl+Alt+Delete and logged in to the computer), the machine was able to log in to the network with its own credentials. As soon as a user interactively logged in to the system, the supplicant would send a new EAPoL-Start message on the network, triggering a whole new authentication of the user's credentials.

Active Directory negotiates a password with a Windows workstation when the workstation joins AD, and the machine state supplicant can use this password as the credential. In addition, the computer maintains its own system store for certificates that is separate from the users' certificate stores. Therefore, machine authentication (also known as computer authentication) is capable of using the computer name and password (PEAP-MS-CHAPv2) or a machine certificate (EAP-TLS or PEAP-EAP-TLS).

Figure 3-17 illustrates the computer starting with machine authentication and then a user logging in to the system, thereby causing a new machine authentication with user authentication.

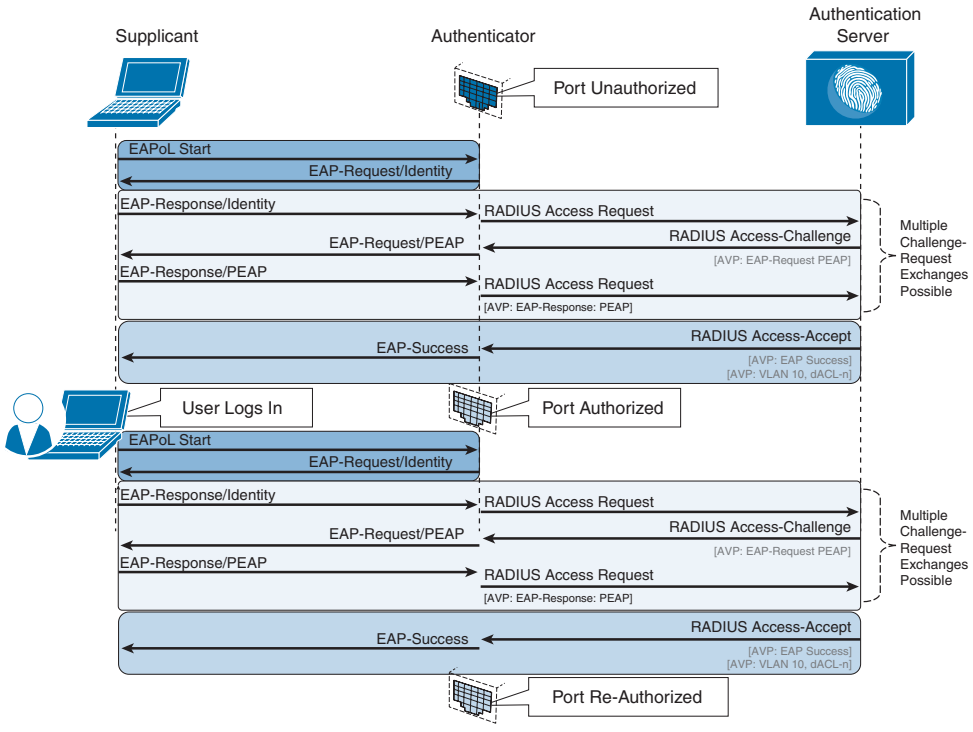


Figure 3-17 *Multiple Supplicant States*

Microsoft and Cisco worked together to update the Windows 10 supplicant and Cisco ISE Version 2.7 with native support for TEAP specifically to support EAP chaining use cases.

Cisco AnyConnect NAM Supplicant

Key Topic

The Cisco AnyConnect Secure Mobility Client is really multiple products in one. Most people are familiar with it as Cisco's premier VPN client, but it is actually more than that. The software is modularized, and several can be installed on the system directly or even pushed down in an update from a Cisco ASA to the software loaded on the Windows workstation. Figure 3-18 shows the module choices when installing AnyConnect on a Windows system.

The key modules that are of interest with Cisco ISE and the SISE 300-715 exam in particular are the ***AnyConnect Network Access Manager (NAM)*** and the ***AnyConnect Diagnostics and Reporting Tool (DART)***. It may help to understand a little bit of history: Many years ago, in a galaxy far far away, there was a product known as Cisco Trust Agent. This product has since been announced as end of life, so don't waste any time remembering the name. The point of interest of this product is that it contained an OEM supplicant from a company called Meeting House. Since that time, Cisco has acquired Meeting House and packaged that supplicant as the Cisco Secure Services Client (CSSC). We have bothered to provide this little tidbit of history because the CSSC client is one of the most widely deployed non-native supplicants for Windows in the world, and Cisco used to charge around $50 per seat for it. Starting with AnyConnect 3.1, that same enterprise-class supplicant was integrated into the AnyConnect client as AnyConnect Network Access Manager (NAM).



Figure 3-18 *Cisco AnyConnect Secure Mobility Client Modules*

There are two main ways to configure the NAM supplicant. One is to use a standalone AnyConnect profile editor for NAM. The other is to use the profiler editor on a Cisco ASA.

Let's walk through the configuration of AnyConnect NAM. The standalone profiler editor allows you to build configurations, and the end user will never see those screens. The editor contains the following views:

- **Client Policy**: This view is for configuration options for connection, wired/wireless/3G mobile broadband, end-user, and administrative settings. Many of the items in this section should look familiar as many of the options are also in the Windows native supplicant.
- **Authentication Policy**: This view is for configuration options related to authentication requirements for user-created networks. In other words, it is where an administrator can restrict what types of networks the end user is permitted to connect to when not at the corporate location.
- **Networks**: This view is where the administrator defines networks available to all groups or specific groups. In other words, this is where the administrator can define the corporate network and the security to use on the corporate wired and wireless networks.
- **Network Groups**: This view is where defined network connections can be assigned to a particular group, which enables easier management of network connections.

The following sections provide a guided tour of the standalone profiler editor. If you have a Windows 10 device at your disposal, feel free to download the standalone editor from Cisco.com and follow along.

Client Policy

The Client Policy view, shown in Figure 3-19, is for configuration options for connection, wired/wireless/3G mobile broadband, end-user, and administrative settings. Many of the items in this section should look familiar as a number of them are also in the Windows native supplicant.

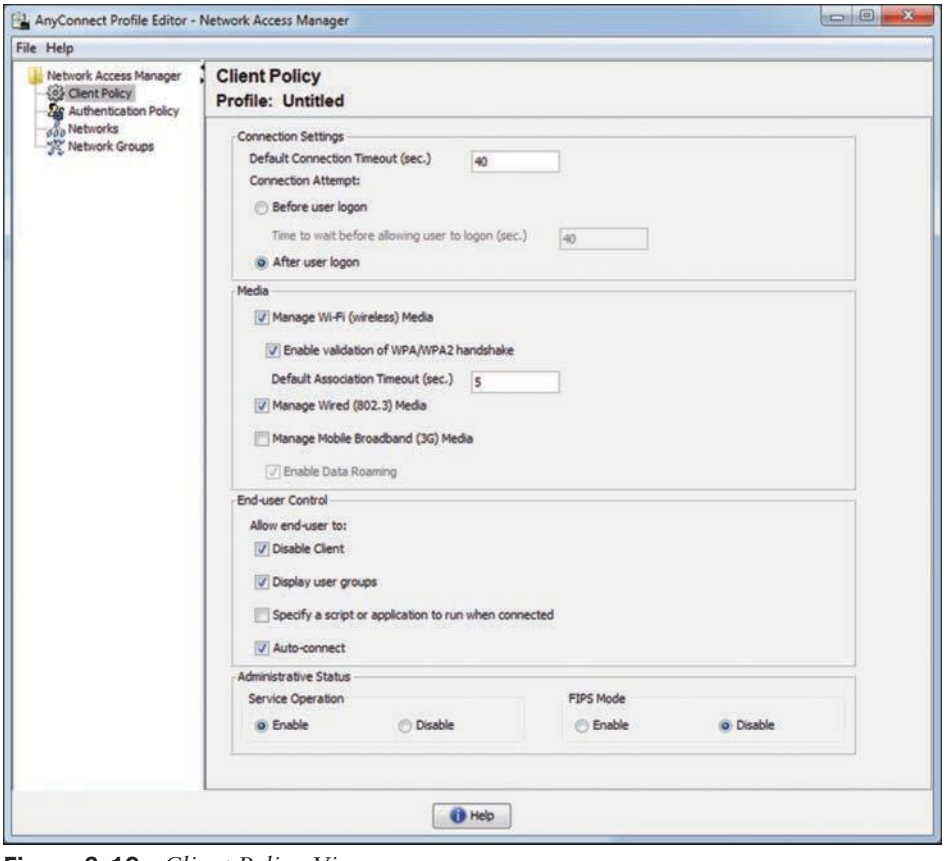


Figure 3-19 *Client Policy View*

The first section, Connection Settings, is very similar to single sign-on settings section for the native supplicant, where you can define whether the supplicant performs authentication before or after allowing the user to interact with the desktop (refer to Figure 3-15).

The next section, Media, allows you to define whether end users should be able to use AnyConnect NAM and join wired and wireless networks.

In the End-User Control section, you can define whether the end user can perform certain functions. Disable Client allows end users to disable NAM and use the Windows native supplicant. Display User Group makes user-created groups created from CSSC 5.x visible and capable of connection; even though they do not correspond to administrator-defined groups, they appear under the local group. In addition, user-created networks defined in AnyConnect NAM appear here.

If you check the box Specify a Script or Application to Run When Connected, users can specify a script or an application to run after the network connection is successfully established. The scripting settings are specific to one user-configured network and allow the user to specify a local file (.exe, .bat, or .cmd) to run when that network gets to a connected state. Many enterprises use this to trigger group policy object (GPO) refreshes from Active Directory or to map network drives.

Auto-Connect allows NAM to automatically connect to administratively defined network connections without user interaction.

The Administrative Status section has two settings. If you set Service Operation to Disable, the system can use the native supplicant and connection manager when the device is not on an administrative network instead of having AnyConnect NAM take over as the connection manager and supplicant for the entire system all the time. The FIPS Mode setting is for Federal Information Processing Standard (FIPS 140-2 Level 1), which is a U.S. government standard that specifies security requirements for cryptography modules. If you set FIPS Mode to Enable, NAM performs cryptographic operations in a way that meets the government requirements.

NOTE There are many more complexities to FIPS mode, and simply enabling this setting is not enough. Please see the official Cisco AnyConnect administrative documentation on www.cisco.com for more information on FIPS mode.

Authentication Policy

The Authentication Policy view allows you to specify what methods are permissible for transmitting credentials to the network. This policy is very similar to the Allowed Protocols authentication result in ISE.

Figure 3-20 shows how you can use the Authentication Policy view to control all sorts of network authentication, including the type of wireless network that an employee can connect to. For example, the company could prevent the managed laptops from connecting to open wireless networks (such as public hotspots) to prevent unencrypted communications.

A policy could be created to deny weaker encryption types, such as WPA/TKIP, and force WPA2/AES only to ensure stronger security. Such a policy could control wired connectivity as well, forbidding wired network connections that do not use 802.1X (although it would be very rare to see this policy) or even deny any connections that don't use 802.1X and Layer 2 encryption (802.11AE), known as MACsec.

Networks

In the Networks view, you can define the corporate network and the security to use on the corporate wired and wireless networks. Figure 3-21 shows a Networks view with a default network named wired. This default network has all security disabled and is there to ensure that the supplicant will work in a plain-Jane non-authenticating wired network.

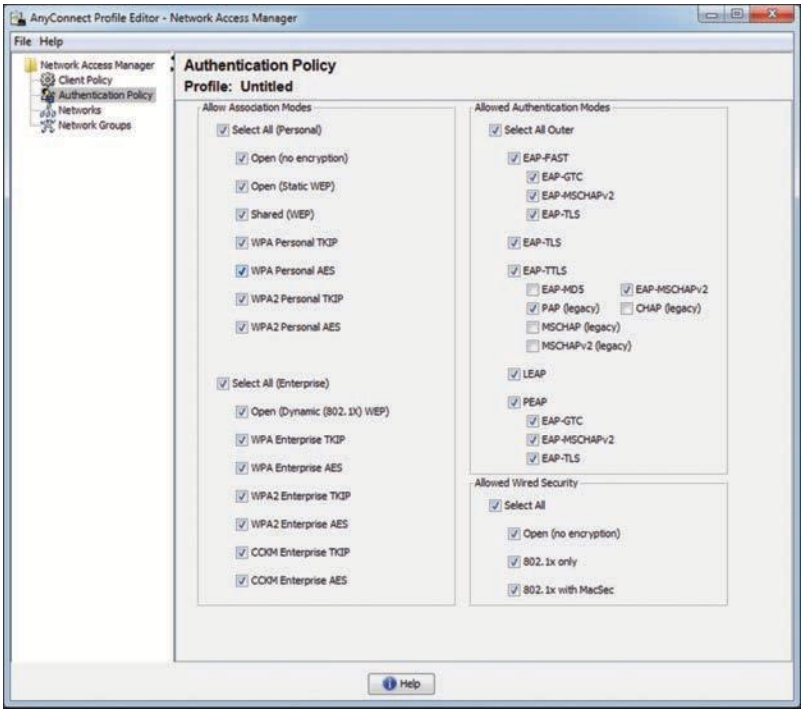


Figure 3-20 *AnyConnect NAM Authentication Policy*

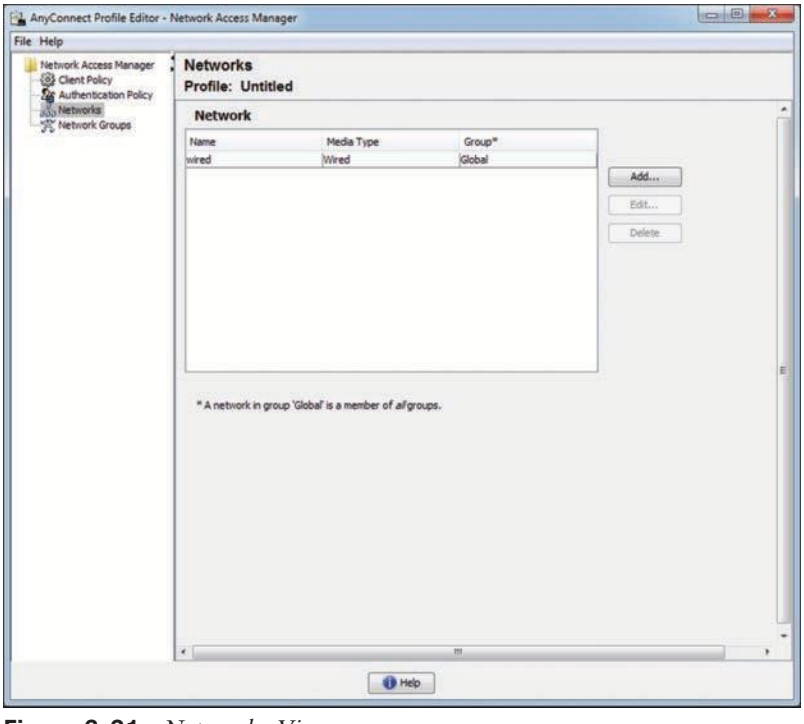


Figure 3-21 *Networks View*

In the Networks view, things get interesting when you click Add. A wizard then walks you through the process of creating administratively defined networks, as shown in Figure 3-22.

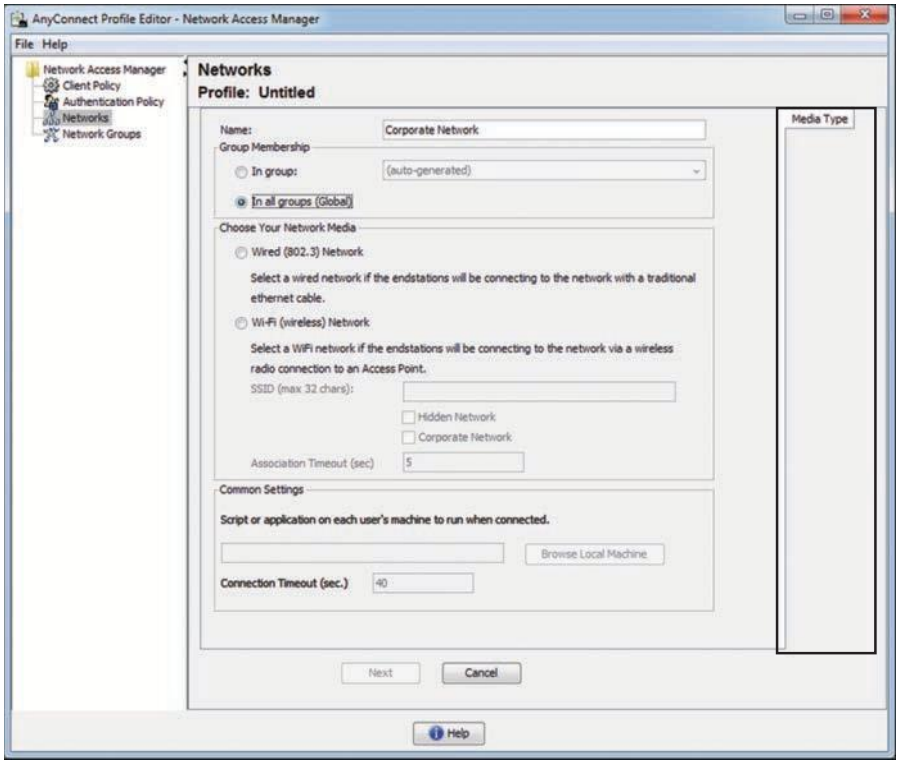


Figure 3-22 *Adding a New Network*

As you build a network profile, the wizard dynamically adds new tabs to the right side, as highlighted in Figure 3-22. If you select a wired (802.3) network, you get a new Security Level tab, as shown in Figure 3-23.

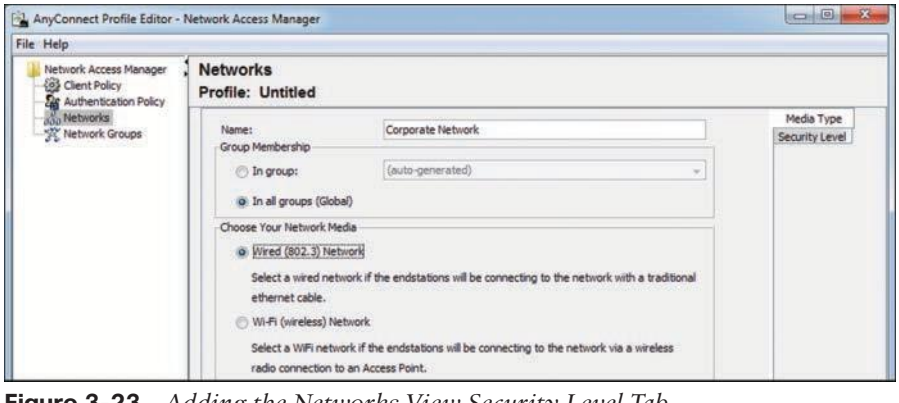


Figure 3-23 *Adding the Networks View Security Level Tab*

Clicking Next at the bottom of the Networks view takes you to the next tab, or you can click a tab to go there directly. Figure 3-24 shows the result.

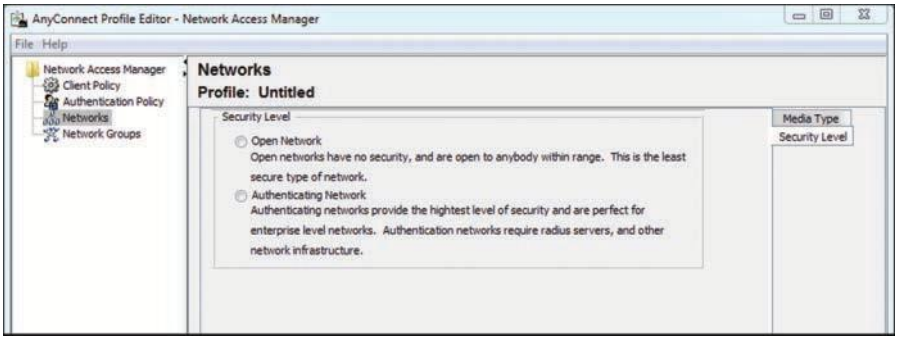


Figure 3-24 *Networks View Security Level Tab*

Because we are talking about corporate networks, and the SISE 300-715 exam covers secure network access, you are likely to choose the Authenticating Network option. As you select this option, a Connection Type tab appears on the right side (see Figure 3-25).

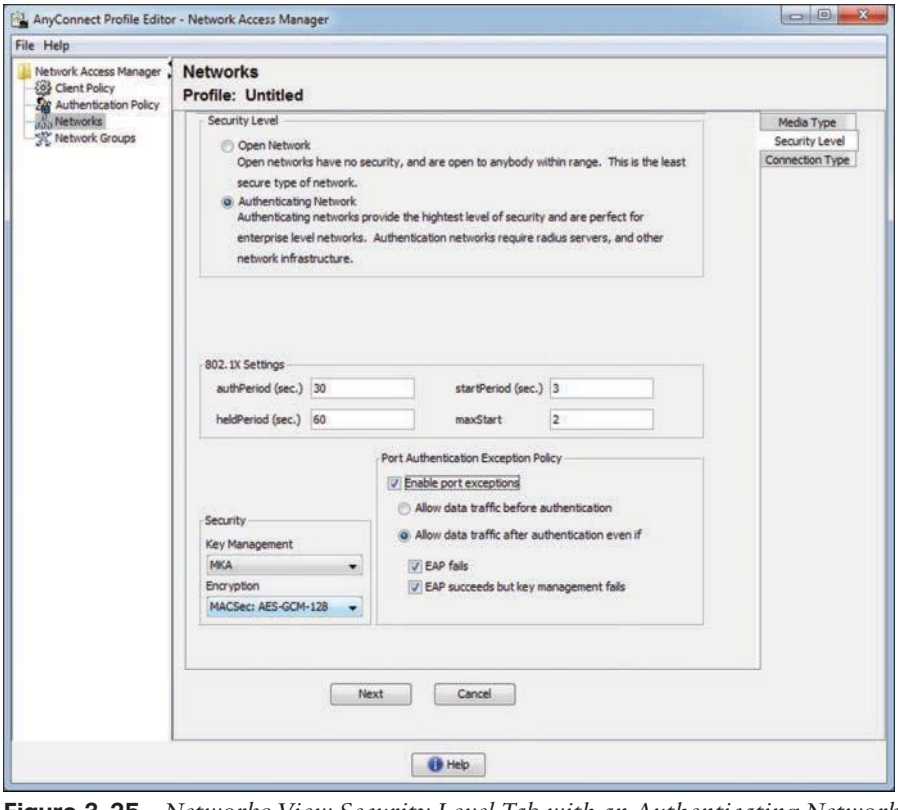


Figure 3-25 *Networks View Security Level Tab with an Authenticating Network*

Key Topic

In Figure 3-25, notice the 802.1X timers you can tune. Also pay attention to the Security section, where key management may be turned on for Layer 2 MACsec encryption between the supplicant and the switch. This provides AES-GCM 128-bit encryption over the wire. Finally, the Port Authentication Exception Policy, when enabled, allows the supplicant to control whether the client can send data immediately upon link (prior to an authentication) or after the authentication, with options about sending data even if EAP fails and even if MACsec fails to negotiate. Environments such as certain government institutions require such strict controls that they must be able to deny traffic if it cannot be encrypted.

Clicking Next takes you to the Connection Type tab. The settings here should seem familiar as they are related to machine authentication, user authentication, or both. If you select Machine and User Connection, you get four new tabs to the right side. Clicking Next takes you to the Machine Auth tab (see Figure 3-26).

When you select the EAP method, the lower portion of the tab becomes populated. For example, if you choose EAP-FAST as the tunneled EAP method, you can then customize the behavior of the supplicant. The inner method is fully selectable for EAP-MS-CHAPv2 or EAP-GTC or EAP-TLS (authenticating using a certificate).

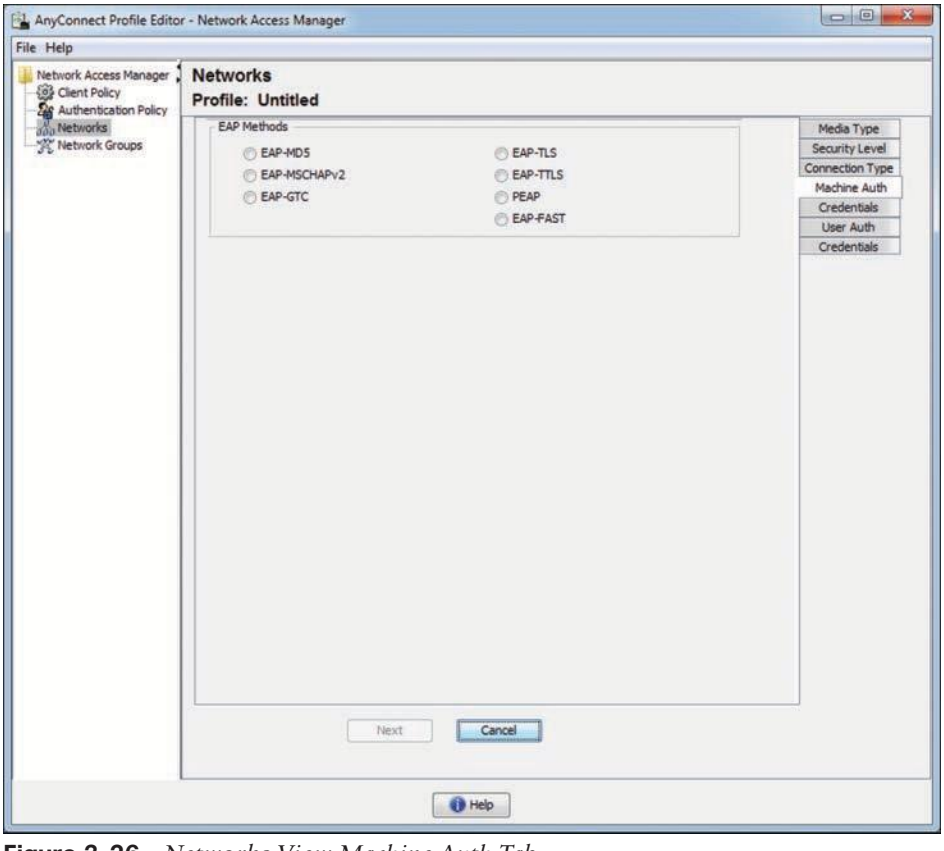


Figure 3-26 *Networks View Machine Auth Tab*

Clicking Next takes you to the Certificates tab, shown in Figure 3-27. You go to the Certificates tab at this point if you chose a password-based authentication method. This tab allows you to define the certificates to trust for the outer method when the TLS tunnel is formed. AnyConnect NAM provides very powerful and flexible rules for filtering what certificates to trust, although these rules are beyond the scope of the SISE 300-715 exam and this book. The Certificates tab also gives you the option to trust any root CA that is already trusted by the operating system or to specify the trusted root CA in the profile.

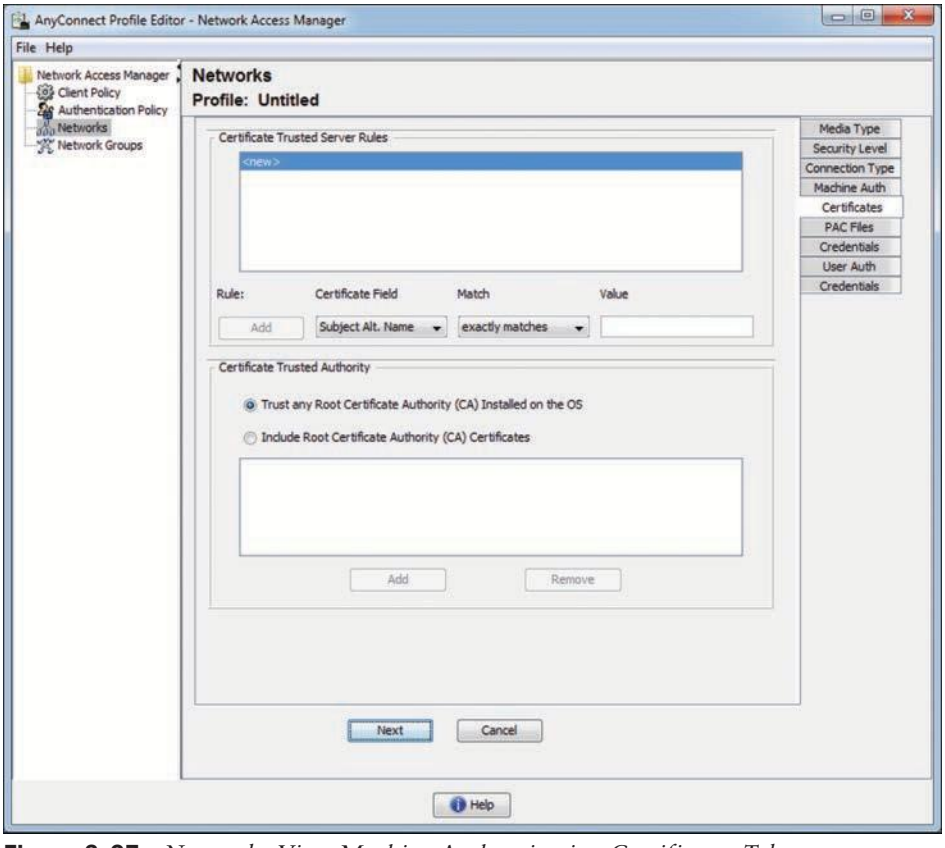


Figure 3-27 *Networks View Machine Authentication Certificates Tab*

Clicking Next takes you to the PAC Files tab, shown in Figure 3-28. This tab, which is specific to EAP-FAST, offers the option of using PAC files. As mentioned earlier in this chapter, a PAC can be thought of as a secure cookie; it is an encrypted trusted file that can be used in lieu of or in addition to a certificate for setting up a secure communication channel with the RADIUS server. This configuration option gives you a way to distribute the server's PAC file in the NAM profile for secure tunnel establishment without the need for certificates.

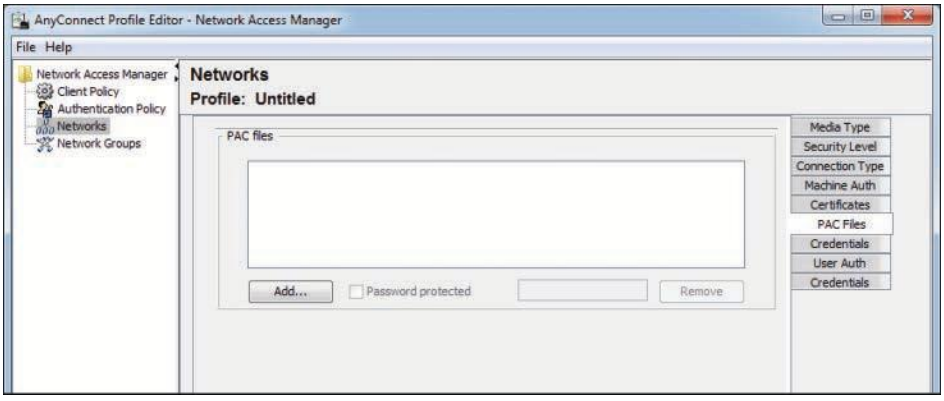


Figure 3-28 *Networks View Machine Authentication PAC Files Tab*

Clicking Next takes you to the Credentials tab, shown in Figure 3-29. This tab provides options for the inner and outer identities with a tunneled EAP. In the case of EAP-FAST, the unprotected identity is the outer identity. Because this is a machine authentication, the outer identity begins with host/, and you can see that the outer identity is set to anonymous by default. You can choose to add a value to this outer identity to aid ISE in making the identity store selection for the authentication. For example, by appending @[domain] to the end of the unprotected identity, you could configure a rule on ISE to route the authentication to the correct Active Directory or LDAP instance. The protected identity pattern is the inner identity; in other words, it is the Active Directory machine account name. With the machine credentials, the default is to use the machine's password that was automatically generated when the computer joined AD. However, if something special is required, you can enter a static password.

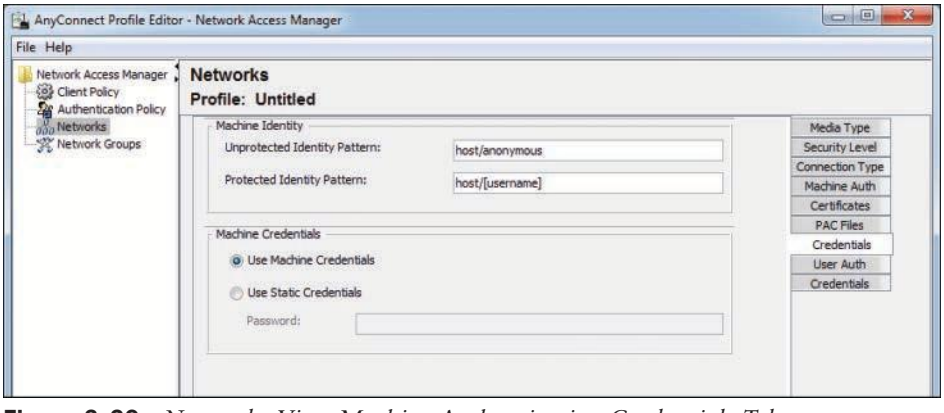


Figure 3-29 *Networks View Machine Authentication Credentials Tab*

Clicking Next takes you to User Auth, the first tab of the user authentication, shown in Figure 3-30. Here you can see that a different EAP method can be selected for the user authentication. It does not have to be the same as the machine authentication. If you select

EAP-FAST, you see all the same choices available as with the machine auth—but with one exception. There is an option to extend user connection beyond logoff. This option permits the supplicant to remain authenticated as that user, even when the user is no longer logged in.

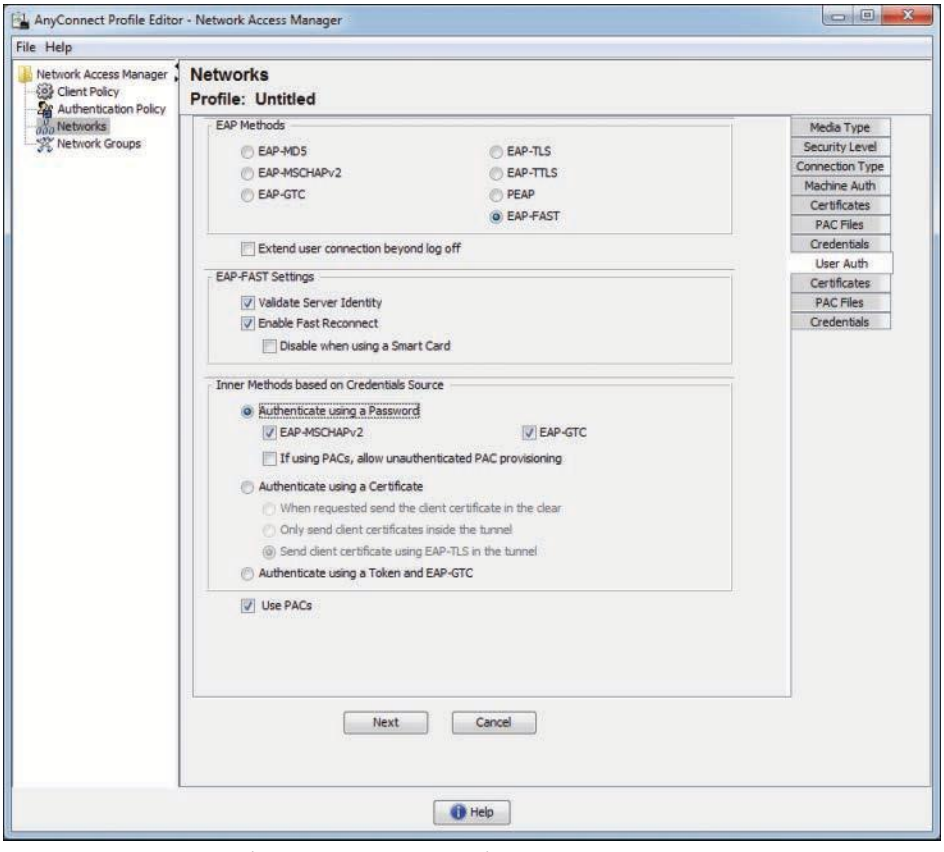


Figure 3-30 *Networks View User Auth Tab*

Clicking Next takes you to the Certificates tab, for user certificates. Just as the machine certificates, you can specify certificates, include them with the profile, and more. Click Next again, and you see the PAC Files tab for users, which is exactly like the machine auth PAC Files tab.

Clicking Next takes you to the final tab for the Networks view: the Credentials tab, shown in Figure 3-31. As with the machine auth Credentials tab, there are fields for the unprotected (outer) and protected (inner) identities. There are also options for using single sign-on credentials, prompting for the user credentials, and using static credentials.

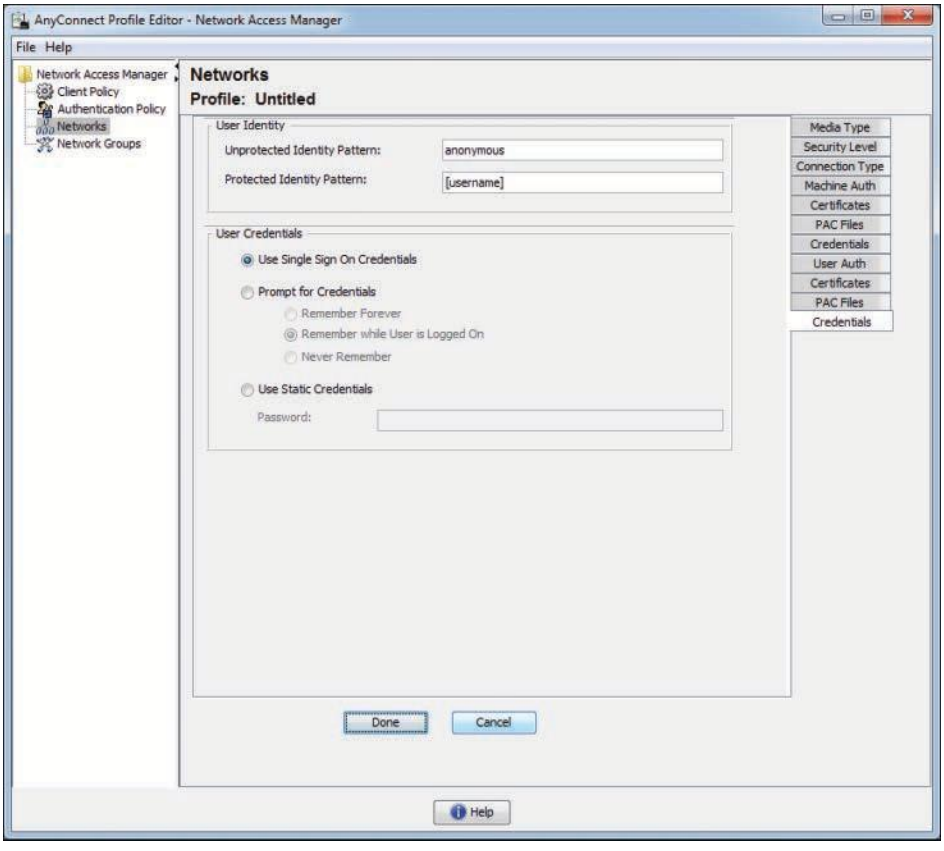


Figure 3-31 *Networks View User Authentication Credentials Tab*

Click Done and notice that the new network configuration is added to the list, as shown in Figure 3-32.

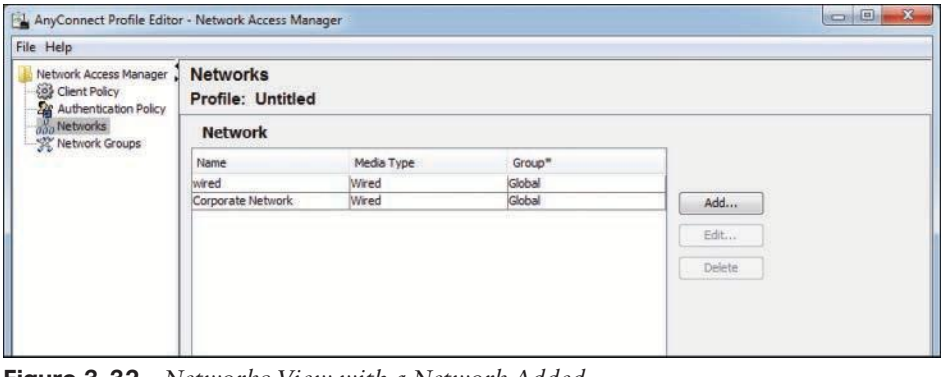


Figure 3-32 *Networks View with a Network Added*

Network Groups

As shown in Figure 3-33, the Network Groups view allows you to organize network profiles into logical groupings for easier use.

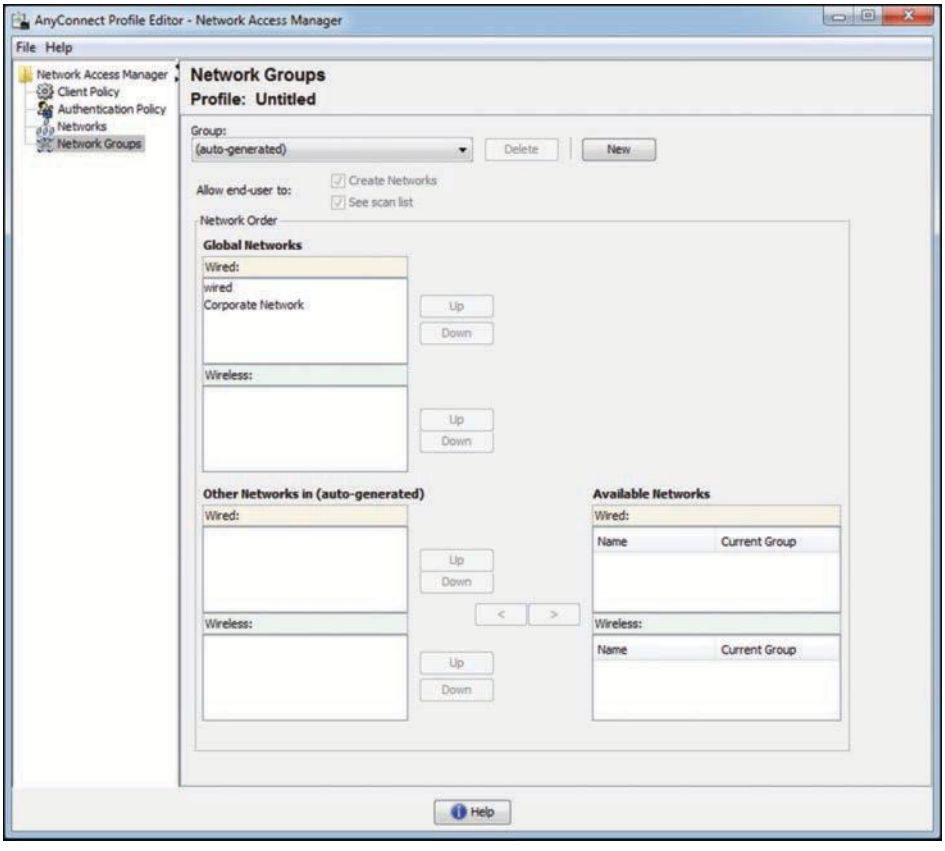


Figure 3-33 *Network Groups View*

Implementing AnyConnect NAM Profiles

After a custom AnyConnect NAM profile has been created, it must be pushed out to all the NAM clients in the network. There is no magical configuration management solution from Cisco that ensures that all endpoints get the latest profile. This is left up to the company's existing software deployment mechanisms, such as Active Directory Group Policy, SCCM, or IBM Big Fix. Another option is to publish the configuration on Cisco ASA, where the profiles get updated the next time a user enters the network through a VPN.

An AnyConnect NAM profile file must be saved with the filename of configuration.xml and needs to be stored in the %SystemDrive%\ProgramData\Application Data\Cisco\Cisco AnyConnect Secure Mobility Client\Network Access Manager\newConfigFiles directory, as shown in Figure 3-34. When NAM starts, it reads in this configuration data. The location of the file depends on the Windows version as well as the AnyConnect version.

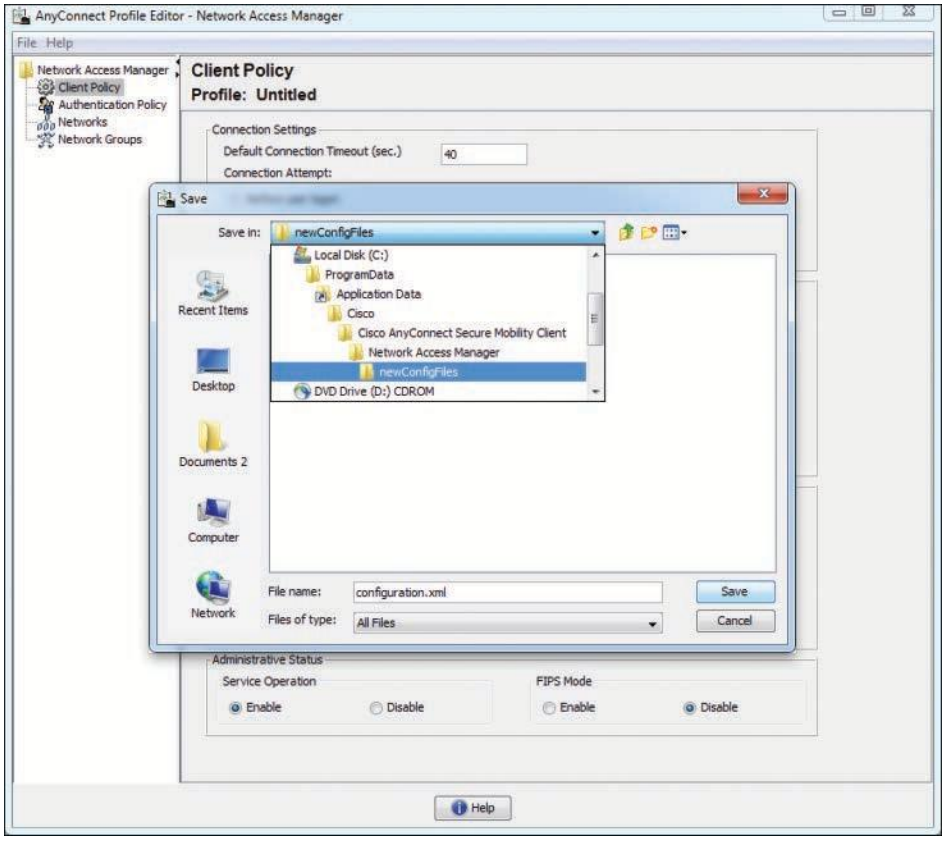


Figure 3-34 *Saving configuration.xml*

The end-user experience is clean and simple. Figure 3-35 shows the GUI that the end user interacts with.

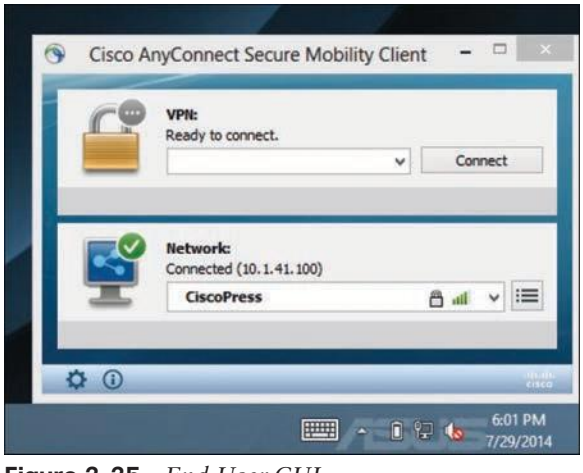


Figure 3-35 *End-User GUI*

EAP Chaining

Key Topic

Having machine authentication and user authentication sounds great. If a machine is able to successfully authenticate to AD, then it must be a company asset and, therefore, managed by the company (using official company software packages and so on). In addition, knowing the user is authorized to be on the company network is, of course, desirable. Providing such assurance is the point of 802.1X.

However, the machine and user states are completely separate. Each EAP session is completely unaware of the other EAP session. Enhancements to RADIUS servers over the years have attempted to track the machine and user authentication states and provide a way to join them together. Many of Cisco's customers, including some of the very largest corporations in the world, have needed to be able to authorize not only users but machines and to ensure that the two are authorized together.

Cisco went beyond what is available with industry standards and created an enhancement to EAP-FAST. With EAP-FASTv2, a differentiation is made between a user PAC and a machine PAC. After a successful machine authentication, ISE issues a machine PAC to the client. Then when processing a user authentication, ISE requests the machine PAC to prove that the machine was successfully authenticated, too. This is the first time in 802.1X history that multiple credentials have been able to be authenticated within a single EAP transaction, and it is known as ***EAP chaining***.

As described earlier in this chapter, the IETF created a new open standard based on EAP-FASTv2 called TEAP. Windows 10 added support for TEAP with EAP chaining, and as this book went to press, Windows was the only supplicant in the world to support TEAP and EAP chaining.

Cisco AnyConnect NAM supports EAP chaining—but not with TEAP. NAM supports EAP-FASTv2, and as this book went to press, NAM had not been updated to support TEAP since Windows offered native support.

NOTE EAP chaining is revisited in Chapter 16, "Bring Your Own Device."

Exam Preparation Topics

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 3-5 lists these key topics and the page number on which each is found.

Key Topic

Table 3-5 Key Topics

Key Topic Element	Description	Page
Table 3-2	802.1X components	41
Section	Tunneled and native EAP types	44
Table 3-3	Tunneled EAP authentication types	48
Paragraph	Machine authentication (computer authentication)	58
Paragraph	AnyConnect NAM supplicant	59
Paragraph	802.1AE (MACsec) encryption and timers	66
Paragraph	EAP chaining concept	73

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Extensible Authentication Protocol (EAP), EAP over LAN (EAPoL), native EAP, tunneled EAP, supplicant, authenticator, authentication server, network access device (NAD), AnyConnect Network Access Manager (NAM), AnyConnect Diagnostics and Reporting Tool (DART)

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. Which two EAP types support the ability to chain together machine authentication and user authentication?
2. What software on an endpoint is used to authenticate to a network using EAP?
3. Which module of AnyConnect can be used to collect detailed logs related to AnyConnect and the host system?
4. What operating system can perform an authentication to the network before a user credential is known or in use?
5. What supplicant is capable of preventing a managed computer from joining open networks (that is, wireless networks without security)?



CHAPTER 4

Non-802.1X Authentication

This chapter covers the following topics:

Describe supplicants and devices without a supplicant

MAC Authentication Bypass

Describe the MAB process within an 802.1X framework

Centralized Web Authentication

EasyConnect

As described in Chapter 3, "Extensible Authentication Protocol (EAP) over LAN: 802.1X," the IEEE standardized on 802.1X as a solution for port-based network access control in the early 2000s. Many predicted that 802.1X would revolutionize networking as we knew it, and no device would be able to plug in and communicate on a network without users identifying themselves and being authorized again. Well, here we are over a decade later, and 802.1X has caught on.

Some thought that all ports would be authenticating with 802.1X, but this really is not practical, and it has not come to be. Ultimately, there can be complications with managing a large number of devices that are not managed PCs, and with the explosion of IoT devices, there is no way that every device that connects to a network can have a configured supplicant. Think about it: Who manages all the printers in a company, and what management is available? Does the team that manages the printers know what a supplicant is, how to configure it, or how to get a certificate onto that device? Most of the time, the answer is no. In addition, many IoT devices are unmanaged and do not support 802.1X, so there must be other ways to identify those devices. The solution could be to list the switch ports that have these devices on them and disable 802.1X on those switch ports—and only those switch ports. Then what? Should you repeat the disabling of security on any port that might have a headless device, for things like IP cameras, badge readers, and digital signage?

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 4-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 4-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Devices Without a Supplicant	1
MAC Authentication Bypass	2, 3, 7, 9
Web Authentication	6, 8, 10
EasyConnect	4, 5, 6, 8

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. In order to allow headless endpoints to connect to a network where IEEE 802.1X has been enabled, what must be configured on the port?
   a. LWA
   b. MAB
   c. CWA
   d. EasyConnect
2. Which of the following is true?
   a. With non-authenticating endpoints, the authenticator takes over the EAP communication from the endpoint.
   b. With non-authenticating endpoints, the authenticator may be configured to send the MAC address of the endpoint as the identity to the authentication server in a RADIUS Access-Request message.
   c. The endpoint's supplicant uses RADIUS to communicate the endpoint's MAC address to the authentication server.
   d. The authenticator can use TACACS+ to send the endpoint's MAC address to the authentication server.
3. Which of following is an accurate statement related to using MAC Authentication Bypass (MAB)?
   a. An administrator is limited in the types of authorization results that can be sent and is restricted to a simple permit-all or deny-all result.
   b. An administrator may assign all authorization results except for VLAN assignment.
   c. An administrator may assign all authorization results except for Security Group Tags (SGTs).
   d. An administrator is not limited in the types of authorization results that can be sent, which may include dACL, VLAN assignment, SGTs, and others.

4. Which of following accurately describes EasyConnect?
   a. EasyConnect provides a quick wizard to allow, configure, and deploy 802.1X certificate authentication.
   b. With EasyConnect, ISE uses Microsoft Active Directory logins to passively map user information onto existing network sessions that have been initiated with MAB.
   c. With EasyConnect, an 802.1X supplicant is required for user authentication.
   d. EasyConnect provides a simple way to deploy 802.1X authentication without the need for a public key infrastructure (PKI) for trusted credential transport.
5. Which of the following does EasyConnect require to function with Active Directory? (Choose two.)
   a. An AD account with permissions to access WMI remotely
   b. A domain administrator account
   c. Access to read the security event log of the Active Directory domain controller
   d. An agent installed on every Active Directory domain controller
6. Which of the following are non-802.1X authentications?
   a. WebAuth, MAB, EasyConnect, remote-access VPNs
   b. Remote access, WebAuth, EAP-MSChapV2
   c. PAP, LWA, RA VPN
   d. WebAuth, EAP-GTC, HTTP POST
7. Which of the following can be used in conjunction with MAC Authentication Bypass (MAB) to help identify an endpoint type?
   a. 802.1X
   b. Profiling
   c. WebAuth
   d. MAC whitelisting
8. Which two non-802.1X authentication methods use specialized authorization results to connect a user's credentials to a MAB session?
   a. Remote access
   b. Local Web Authentication (LWA) with a centralized portal
   c. Centralized Web Authentication (CWA)
   d. Local Web Authentication (LWA)
   e. EasyConnect
9. What is one of the main reasons that MAB is used in modern networks?
   a. Most endpoints, such as printers and IP phones, do not have supplicants and therefore cannot use 802.1X.
   b. An endpoint may have a supplicant, but the enablement and configuration of that supplicant could be very complicated or operationally difficult for the company. Therefore, the company might opt to use MAB instead.
   c. The endpoints mostly do have supplicants, but they are not compatible with Cisco networks.
   d. MAB is as secure as 802.1X, and it is therefore chosen often to save the company the operational difficulties of configuring the supplicants on such disparate endpoints.

10. In ISE Version 2.1, what option was introduced for web redirection with third-party network access devices?
    a. CWA with DHCP and DNS services
    b. CWA with CoA
    c. LWA
    d. LWA with centralized portal

Foundation Topics

Devices Without a Supplicant

As described in Chapter 3 and the introduction to this chapter, the IEEE standardized 802.1X as a port-based network access control solution. More than a decade later, we have seen 802.1X truly take hold and become the de facto solution deployed everywhere, on both wired and wireless networks.

With 802.1X, a supplicant and an authenticator exchange EAP messages. However, if the endpoint that connected to the authenticator (switch) did not have a supplicant, the EAP identity request messages would go unanswered. This would result in an authentication timeout, and with the original concept of identity networking and 802.1X, the endpoint would be denied access to the company network. In other words, only devices that can authenticate and have authenticated to the network will be allowed on the network.

Although it was thought that all ports would need to authenticate with 802.1X in order to gain access, this did not happen for a myriad of reasons. When designing a secure network access solution, there is a tendency to consider only the managed desktops, laptops, and now tablet-type devices when thinking about network authentication. However, the reality is that organizations tend to possess a plethora of devices beyond those, such as printers, IP cameras, IP phones, thin client terminals, and building automation and other "headless" or IoT devices that exist in networks today.

Take Cisco IP Phones as an example. Such a device does have a supplicant, which can be configured individually at the keypad, but this solution does not scale very well; imagine having to configure a supplicant on every phone in an organization with hundreds of thousands of phones. Cisco IP Phone supplicants may also be configured centrally from the Call Control Server (formerly named the Cisco Call Manager). What about the other devices? Do they also have a central management platform that can configure each supplicant across large numbers of devices deployed at scale? In most cases, probably not.

One of the original ways to deal with these non-authenticating devices was to not configure 802.1X on the individual switch ports where the non-authenticating endpoint would be plugged in to the network. This "solution" would allow anyone to simply unplug the non-authenticating endpoint and plug their laptop into the port to gain full network access without any challenge whatsoever.

What about when a device (such as a printer) needs to be moved? The network team might need to be involved for that move to enable 802.1X on the old switch port and disable 802.1X on the new switch port; this would impose a significant management burden on the IT department, it would just not be scalable, and it would defeat the point of network access control.

What happens when an employee who should have network access has a misconfigured supplicant or an expired credential or is using a temporary device? What about guest users who only need access to the Internet? A series of special fixes would need to be created to deal with all these variations, and it might not be possible to deal with them in a sustainable way.

MAC Authentication Bypass

Key Topic

The first special fix to help with non-authenticating endpoints is for the authenticator to act on behalf of the endpoint that does not have a supplicant. In this scenario, the authenticator crafts a RADIUS Access-Request message and sends it to the authentication server. The authenticator uses the endpoint's MAC address as the identity.

The authentication server (the RADIUS server) performs an authentication lookup, using that MAC address as the credential. If that MAC address is in a list to be given access to the network, a RADIUS Access-Accept message is sent back from the authentication server to the authenticator. This process is known as *MAC Authentication Bypass (MAB)*.

Figure 4-1 illustrates the MAB process.

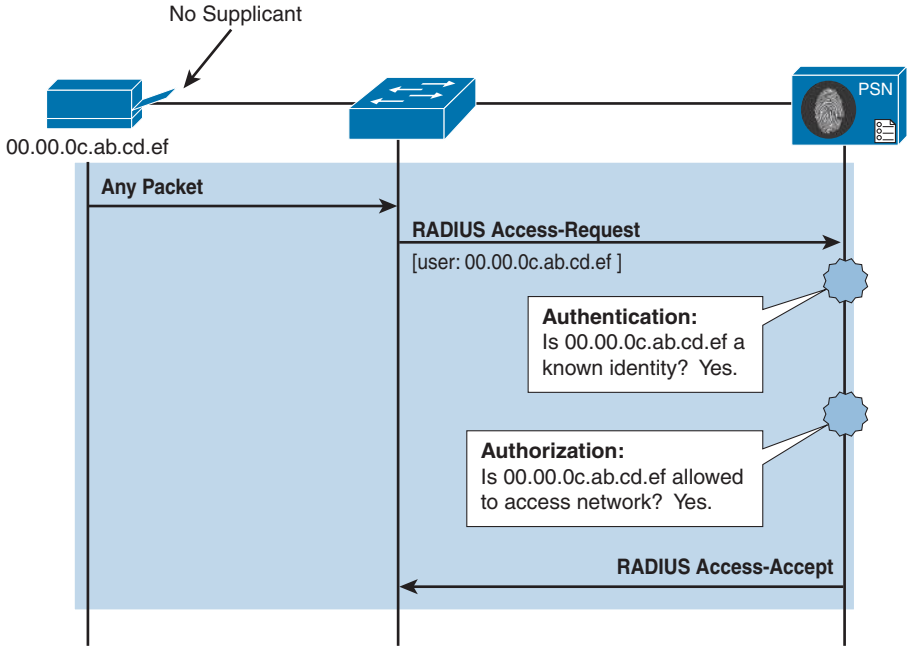


Figure 4-1 *MAB Sending MAC Address as the Identity*

In Figure 4-1, you can see that there is a non-authenticating endpoint (a printer) with the MAC address 00.00.0c.ab.cd.ef. There is also a switch, which is the authenticator, and an ISE server is acting as the authentication server. The following steps occur:

Step 1. Because the printer does not have a supplicant, the authenticator (switch) crafts a RADIUS Access-Request message, using the printer's MAC address as the identity.

Step 2. The authentication server (ISE) receives the RADIUS Access-Request and performs an identity lookup to determine whether the address is a known MAC address.

Step 3. The authentication server (ISE) determines whether the device should be granted access to the network and, if so, what level of access to provide.

Step 4. The authentication server (ISE) sends the RADIUS response (Access-Accept) to the authenticator, allowing the printer to access the network.

It is important to note that while 802.1X is a standard, MAB is not. Each vendor can implement MAB differently; vendors only need to ensure that the RADIUS communication complies with the standard for RADIUS.

How does a switch (authenticator) know when the endpoint that is plugged into it does not have a supplicant? Following the 802.1X standard, the method is simply a timeout. The authenticator is meant to send EAP over LAN identity request frames every 30 seconds, by default. After three timeouts—a period of 90 seconds, by default—it is accepted that the endpoint must not have a supplicant. Like most other Cisco switch features, timers are adjustable. Recommended practice is to change the default timer on the port if a switch port is configured for a mix of MAB and 802.1X.

Figure 4-2 shows timeouts occurring three times before MAB begins.

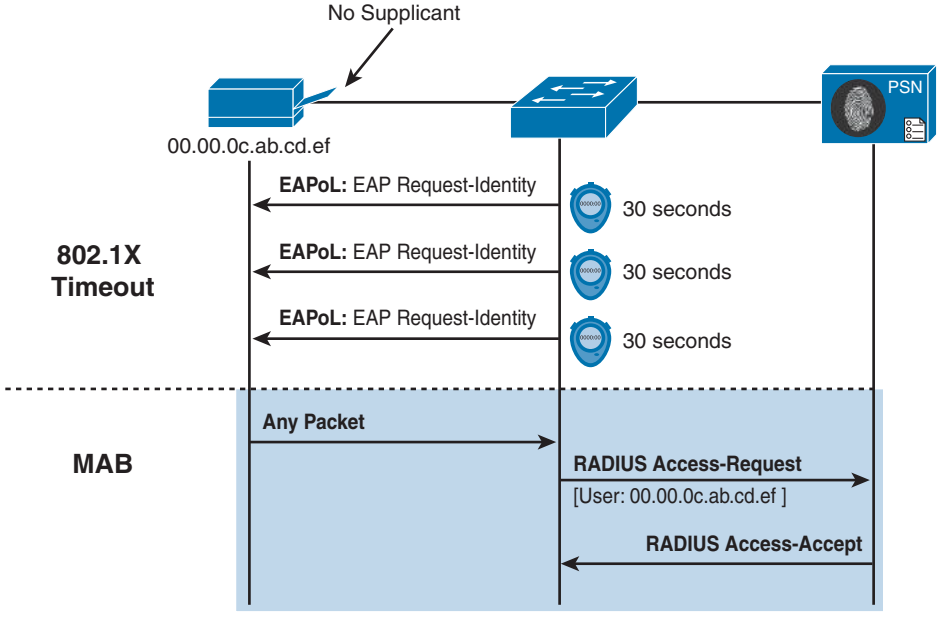


Figure 4-2 *EAP Process Timing Out Causing Failover to MAB*

Key Topic

Keep in mind that MAB is inherently not a secure authentication method. When implementing MAB, you are bypassing the stronger security of 802.1X by allowing specific MAC addresses to gain access without authentication. MAC addresses alone are easily spoofed,

which means it is easy to configure an endpoint to use a MAC address other than the one burned into the hardware. When using MAB, always follow a defense-in-depth approach. This means when authorizing a MAB device for network access, the endpoint should be granted access only to the networks and services that a device is required to speak to. In other words, don't provide full access to devices that have been bypassed using MAB; instead, provide them with an authorization that is more limited. Because MAB is a standard RADIUS authentication, and the authorization decision is being sent from the authentication server (ISE), there really are no limitations on the type of authorization results that can be sent to the authenticator. Some examples include but are not limited to the following:

- Downloadable ACLs (dACLs)
- Dynamic VLAN assignment (dVLAN)
- URL redirection
- Security Group Tags (SGTs)
- Smartport macros

Keep in mind that if an endpoint does not have a supplicant, it is not recommended to ever change its VLAN. When you change a VLAN assigned to an endpoint, that endpoint must know (somehow) to renew the DHCP request. The best solution is to not use VLAN changes on open networks because there is nothing on the client to detect the VLAN change and trigger the DHCP renewal. When the network uses 802.1X, there is a supplicant on the client to do the VLAN change detection (such as determining whether the gateway is reachable) and trigger the DHCP renewal.

If you still choose to change the VLAN on open networks, then you have only a few choices (none of which are considered best practices). You can set the DHCP lease time to something *very* low, to renew the address frequently. There is also an option to use an ActiveX or Java applet on the portal to do the VLAN change detection in lieu of a supplicant. You can also configure a port bounce as part of the Change of Authorization (CoA) but this can also create a new set of problems when the endpoint is a PoE device and has to reboot when the port is bounced, so it is generally not recommended. It should also be noted that even if CoA performs a port bounce, a reauthorization will be sent instead if there are multiple 802.1X sessions on the same switch port.

Another way to add security to MAB is to combine it with profiling to authorize an endpoint. With profiling, ISE collects attributes about the connecting device beyond just the MAC address to form a fingerprint of what type of device it might be. Based on that information and the type of device that is connecting, authorization on the network can be granted.

NOTE We dig more deeply into profiling in Chapter 14, "Profiling." Defense in depth, authorization results, and applying the correct level of access to the correct device types are covered in much more detail in Chapter 10, "Authorization Policies."

Web Authentication

Just because there is no configured supplicant on an endpoint does not mean the user of that endpoint does not need to authenticate. Consider the use cases of guests or visitors, or maybe just a misconfiguration or an expired credential for the end user. The user might still require network access and might be granted access to the network.

This is where ***Web Authentication***, commonly referred to as just ***WebAuth***, comes in. An authenticator would be able to send a user to a locally or centrally hosted web page. In other words, this web page could be hosted on the local device itself—the switch, wireless controller, or even the firewall or VPN concentrator—or it could be hosted on the authentication server (ISE). It could be just a very basic web page that allows a user to submit a username and password.

The username and password submitted to the web portal are then sent to the authentication server (ISE) for authentication. If the portal is locally stored in the switch, in a very similar fashion to what occurs with MAB, the switch sends the request for the endpoint because the endpoint is not authenticating to the switch. Figure 4-3 illustrates the WebAuth concept.

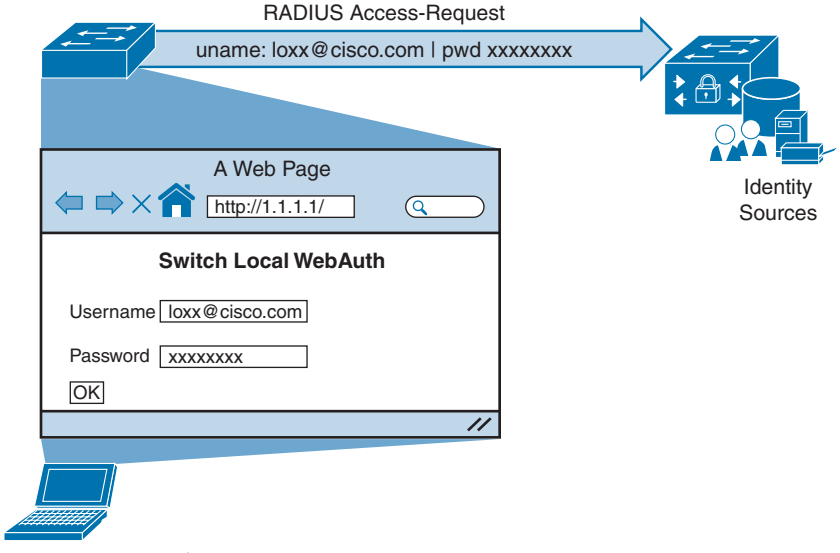


Figure 4-3 *Web Authentication*

The credentials that get submitted through the WebAuth page could be the Active Directory credentials of an employee, or they could be guest credentials for someone who is only temporarily allowed to have Internet access (and no other access). The use of WebAuth is really not limited to one type or another.

Keep in mind that WebAuth is only an effective authentication method for a device that is interactive. In other words, it would not make sense to try to use WebAuth for a printer as there would be no user to interactively enter credentials and click Submit.

Neither MAB nor WebAuth is a standard.

There are multiple ways to perform WebAuth, and each has benefits and drawbacks.

Local Web Authentication

Local Web Authentication (LWA) is the original form of WebAuth. As described in the preceding paragraphs, the authenticator redirects web traffic (HTTP and/or HTTPS) to a locally hosted web portal, where a user can enter a username and password.



The credentials are submitted through the web portal, and the authenticator (for example, a switch, a wireless controller) sends the RADIUS Access-Request to the authentication server, using the username and password from the form. It is key to remember that any time the switch is sending the credentials for the user, the process is considered LWA.

Local Web Authentication isn't used today as it once was for several reasons. Storing the web authentication portal on the local authenticator can often limit the customization capabilities for the web pages and is not very dynamic. On a Cisco switch, the pages are almost not customizable at all. Many organizations not only prefer but require that the web portals be customized to match the corporate branding. For those companies, traditional LWA is not usually an acceptable solution—at least not for wired WebAuth.

When using LWA with Cisco switches, there is no native support for advanced services, including acceptable use policy acceptance pages, client provisioning, password changing, self-registration, or device registration. For such advanced capabilities, a company needs to consider using Centralized Web Authentication, discussed later in this chapter.

Cisco switches as well as a variety of other 802.1X-compliant switches have a configuration option (called the Guest VLAN) that involves assigning a special VLAN to endpoints when the authentication timer expires; this means they don't have a supplicant. This option was available before powerful policy servers like Cisco ISE existed. There are still many production deployments of 802.1X today that use the Guest VLAN option to provide wired guest access to the Internet. However, it is important to note that once a switch makes a local decision, such as assigning the Guest VLAN, LWA is no longer an option.

In addition to LWA and the Guest VLAN being mutually exclusive, there are some other interesting bits of information about LWA. For example, LWA does not support VLAN assignment, so you are basically limited to ACL assignment. LWA is also restricted from Change of Authorization (COA) support; therefore, access policy cannot be changed based on posture or profiling state—or even an administrative change as a result of malware or any other need to quarantine the endpoint.

Local Web Authentication with a Centralized Portal

There is an option with many modern authenticators to redirect LWA to a centralized web portal. Utilizing a centralized portal allows an organization to customize the portal with corporate branding and provide the right look and feel for the organization.

With Cisco switches, the locally stored web pages may contain the redirection string that sends user web traffic to the centrally hosted portal. That hosted portal is configured to send the credentials entered into it to the source NAD through an HTTP POST method or returned to the NAD through a hidden iframe.

Figure 4-4 illustrates the process of the user entering credentials into the centrally located portal (which in this case is hosted on ISE) and the HTTP POST occurring from the browser to the network device (authenticator).

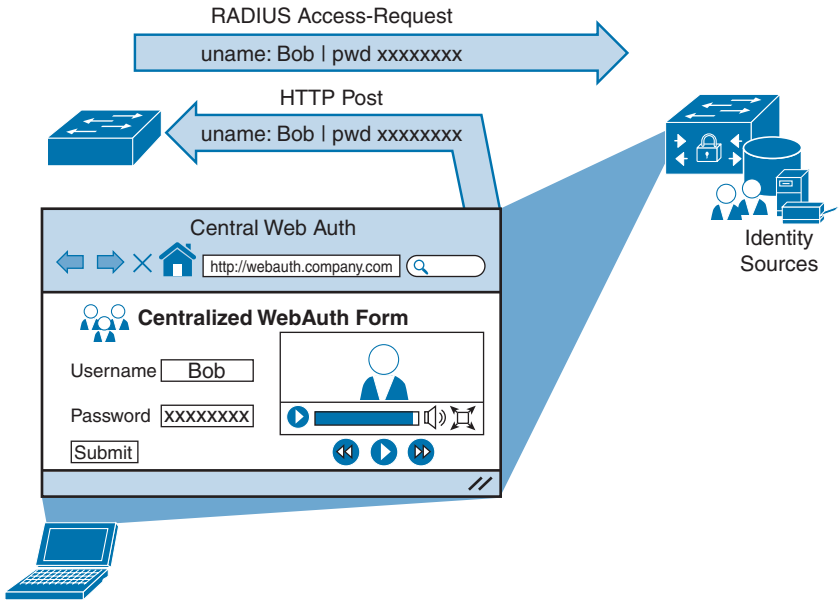


Figure 4-4 *Centrally Hosted Portal for LWA: HTTP POST Occurs from Browser to Authenticator*

There are pros and cons to the POST method and the iframe method. The iframe method does not work with newer browsers. The POST method does not work with some mobile browsers, and it is very insecure because the user's credentials are passed from the central portal through the network, destined to the authenticator. The web portal needs to have the intelligence to determine which methods to use with each browser type by examining the user agent string from the browser.

Cisco Wireless LAN Controller (WLC) Version 7.0 required the use of either POST or iframe method in order to support LWA with a centralized portal on ISE. Traditional LWA was also still available. Even though the portal is centralized, the same restrictions associated with traditional LWA are still in effect. CoA does not function, and client provisioning and other advanced functionality are also not available. This method isn't widely adopted or used for these reasons.

Centralized Web Authentication

Cisco ISE uses ***Centralized Web Authentication (CWA)*** throughout the Secure Access solution. While Cisco ISE is still capable of supporting LWA methods, those methods were historically used for non-Cisco network access devices before DHCP and DNS services were added in ISE Version 2.1. (We explore this option later in this chapter.)

Just like LWA, CWA is only for interactive users that have web browsers. In addition, a user manually enters a username and password, and the WebAuth and Guest VLAN functions remain mutually exclusive.

CoA works fully with CWA, which leads to support for all the authorization results, such as ACL and VLAN authorization. Keep in mind that any time you change VLANs on an endpoint, the endpoint must be able to detect the VLAN change and trigger an IP address

renewal. With 802.1X, the supplicant takes care of the VLAN change detection and address renewal. However, when using CWA, there is normally no supplicant on the endpoint. Therefore, the portal must use an ActiveX or Java applet to handle the renewal of the IP address after the VLAN assignment. One thing to keep in mind with DHCP changes and the use of applets is that the guest machine or browser might be locked down by the employer. For this reason, VLAN changes are not always preferable when other options such as dACLs, ACLs, and SGTs exist.

CWA supports advanced services such as client provisioning, posture assessments, acceptable use policies, password changing, self-registration, and device registration.

Key Topic

Now that you've read about all the things CWA can do, you must be wondering how it works and what makes it different from the other WebAuth options. The authenticator only sees a MAB, and the rest is handled on the authentication server (ISE). Figure 4-5 shows MAB occurring with a redirection to the centralized portal, and the switch still only sees a MAB request, while ISE is maintaining the user authentication.

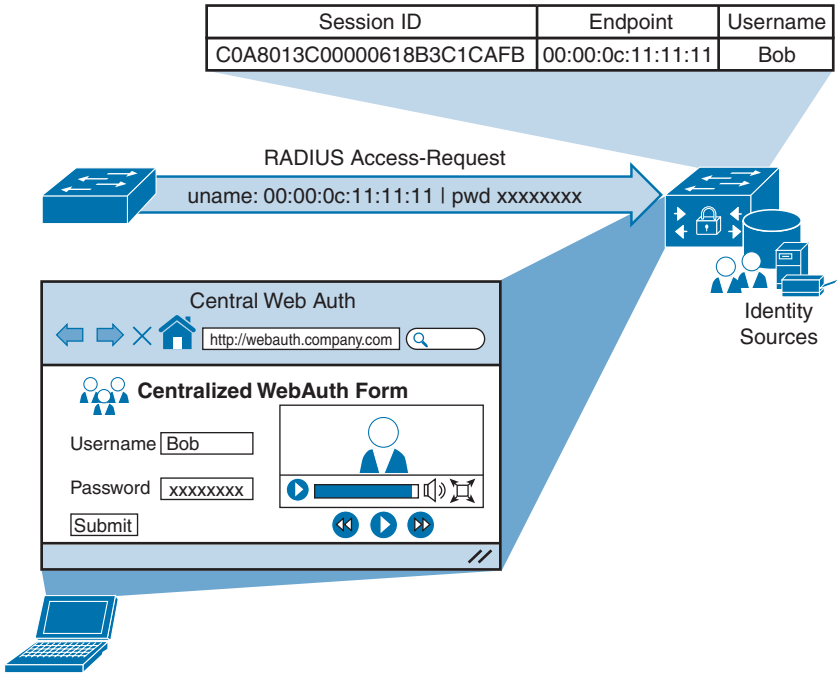


Figure 4-5 *URL-Redirected MAB with Credentials Never Sent to the Authenticator*

The process shown in Figure 4-5 includes the following steps:

Step 1. The endpoint entering the network does not have a supplicant configured.

Step 2. The authenticator performs MAB, sending the RADIUS Access-Request to Cisco ISE (the authentication server).

Step 3. The authentication server (ISE) sends the RADIUS result, including a URL redirection, to the centralized portal on the ISE server.

Step 4. The end user enters credentials into the centralized portal. Unlike with LWA, the credentials are never sent to the switch but are stored in the ISE session directory and tied together with the MAB coming from the switch.

Step 5. ISE sends a reauthentication CoA (CoA-reauth) to the switch. This causes the switch to send a new MAB request with the same session ID to ISE, and that request is processed.

Step 6. ISE sends the final authorization result to the switch for the end user.

Chapter 12, "Web Authentication," examines the configuration of CWA and the policies involved.

Centralized Web Authentication with Third-Party Network Device Support

ISE Version 2.1 added support for third-party network access devices, which included DHCP and DNS services. Since many third-party network devices do not support static or dynamic URL redirection natively, this important enhancement gave ISE the capability to provide an authentication VLAN (or auth VLAN, as it is commonly called) to simulate that URL redirection flow. Figure 4-6 illustrates the URL redirection mechanism using the auth VLAN. The following list provides details of the steps shown in the figure:

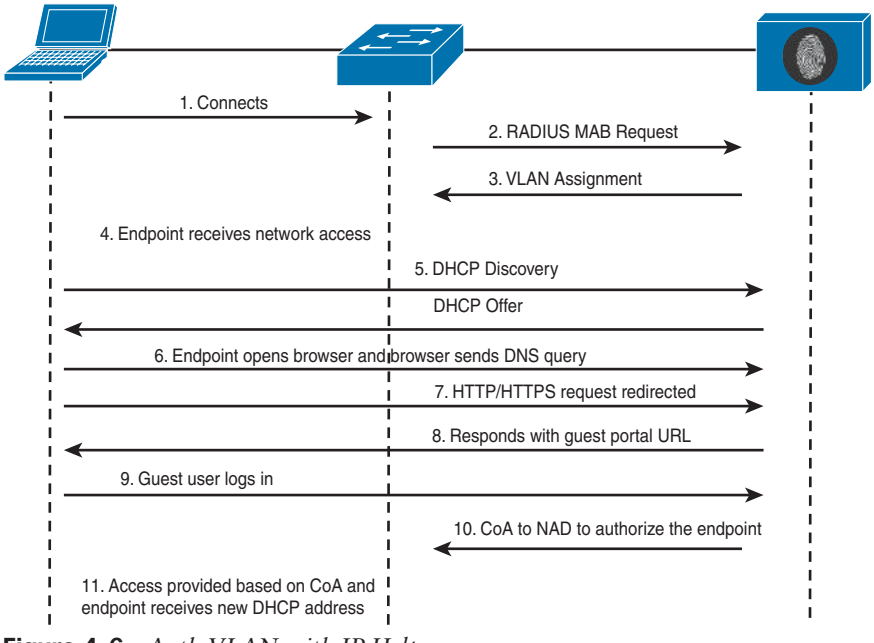


Figure 4-6 *Auth VLAN with IP Helper*

Step 1. The guest endpoint connects to the authenticator (third-party network access device).

Step 2. The authenticator sends the RADIUS/MAB request to ISE.

Step 3. The authentication server (ISE) sends the RADIUS access/accept message containing the auth VLAN ID.

Step 4. The guest endpoint receives network access.

Step 5. The guest endpoint broadcasts a DHCP request. The endpoint obtains a client IP address as well as ISE's sinkhole DNS IP address from ISE's DHCP service.

Step 6. The guest endpoint opens a browser. The browser sends a DNS query and receives ISE's IP address.

Step 7. The endpoint HTTP or HTTPS request is redirected to the authentication server (ISE).

Step 8. The authentication server responds with an HTTP 301 (Moved) message and provides the guest portal URL. The endpoint's browser then redirects to the centralized guest portal page.

Step 9. The guest user logs in for authentication.

Step 10. The authentication server (ISE) responds to the network access device, sends the CoA, authorizes the endpoint, and bypasses the sinkhole.

Step 11. Access is provided to the user based on the CoA, and the endpoint receives an IP address from the enterprise DHCP server. The user can now use the network.

Remote-Access Connections

So far, this chapter has focused on wired and wireless connections where the endpoint is not using 802.1X. However, there is a third network access use case: remote access. Today this trinity is commonly referred to as "wired, wireless, and VPN."

Connecting an endpoint to a corporate network from another location is known as remote access (RA). RA used to be handled mainly by dial-up networking, and this is actually where RADIUS got its name. (RADIUS stands for Remote Authentication Dial-In User Service.) Today it is much more common to use a virtual private network (VPN) to join an endpoint to a company's network through the public Internet. With this type of connection, the VPN concentrator (usually a Cisco ASA) is the authenticator, and ISE is the authentication server. The remote client forms a secure tunnel between the VPN software (usually AnyConnect) and the VPN headend (usually a Cisco ASA). The tunnel may be formed using IPsec or SSL, and the user's credentials are passed inside that secure tunnel.

The VPN headend sends the RADIUS Access-Request to ISE, which directs the identity lookup to the appropriate identity store—usually a one-time password (OTP) server. With a remote-access VPN, the user's credentials are either sent through PAP or MS-CHAPv2 from the ASA to ISE.

NOTE 802.1X is not involved with remote-access VPNs.

EasyConnect

EasyConnect is a feature that was added in ISE Version 2.1 as an alternative for user authentication without 802.1X or WebAuth. Like 802.1X, it provides port-based authentication, but it is much easier to implement. EasyConnect allows users to connect a wired endpoint to the network and grants just enough access so users can authenticate through Active Directory. ISE collects the user authentication information from Active Directory and issues a CoA to the network device for further network access. EasyConnect can be implemented on the same port as 802.1X and can provide initial authentication to a network as an enterprise is migrating to 802.1X. Table 4-2 outlines the advantages and disadvantages of EasyConnect.

Table 4-2 EasyConnect Advantages and Disadvantages

Advantages	Disadvantages
No endpoint supplicant configuration is required.	Only Cisco NADs are supported.
No PKI configuration is required.	IPv6 is not supported.
ISE issues the CoA after Active Directory authenticates the user.	Wireless connections are not supported.
	EasyConnect enables only user authentication and does not support machine authentication.
	Only Windows-installed endpoints are supported.
	Only logon events are captured. User logoff detection doesn't exist for EasyConnect. The session ends when the session ends (port disconnect) or a new user logs into the computer.

In order to intercept the user authentication information from the Active Directory domain controller, ISE must be able to connect to the domain controller using the Microsoft WMI interface and query logs from Windows event messaging. It can be tedious to configure this manually on the domain controller but, thankfully, ISE provides a one-click way to push the required configuration to the domain controllers. The session information gathered from the domain controller is added to ISE's session directory and can also be published to third-party systems with pxGrid. (We discuss pxGrid further in Chapter 22, "ISE Context Sharing and Remediation.")

Figure 4-7 illustrates the EasyConnect flow, and the numerals in the figure correspond with the following steps:

Step 1. The endpoint connects to the wired authenticator (switch).

Step 2. The authenticator sends an access request to the authentication server (ISE), and the authentication server responds with access based on user configuration, allowing the endpoint to access Active Directory, DNS, and DHCP.

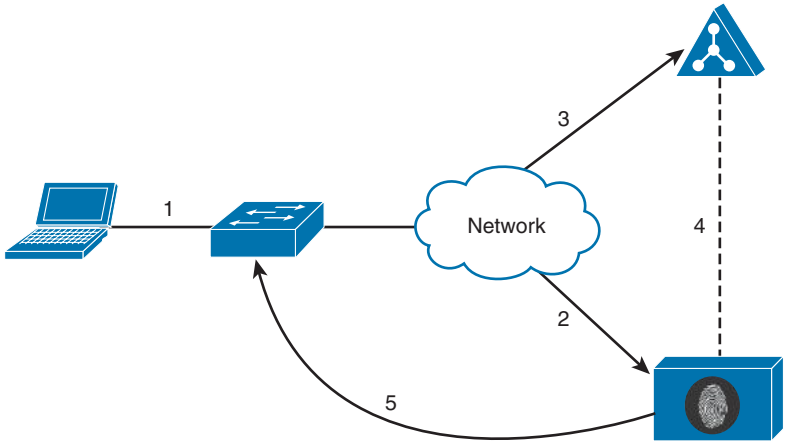


Figure 4-7 *EasyConnect Basic Flow*

Step 3. The user logs in to the domain, and the domain controller sends a security audit event to the authentication server (ISE).

Step 4. The authentication server (ISE) collects the MAC address from RADIUS, and it collects the IP address and domain name, as well as login information about the user, from the security audit event.

Step 5. Once all the data is collected and merged in the ISE session directory, the authentication server (ISE) issues a CoA to the authenticator, and the endpoint is provided access to the network.

Exam Preparation Tasks

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 4-3 lists these key topics and the page number on which each is found.

Key Topic

Table 4-3 Key Topics

Key Topic Element	Description	Page Number
Section	MAC Authentication Bypass authenticating endpoints	80
Paragraph	The MAB process in the 802.1X framework	81
Paragraph	Local Web Authentication (LWA)	84
Paragraph	Centralized Web Authentication (CWA)	86

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

MAC Authentication Bypass (MAB), Web Authentication (WebAuth), Local Web Authentication (LWA), Centralized Web Authentication (CWA), EasyConnect

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. Why is there a need to authenticate endpoints without using 802.1X?
2. What are some of the advantages of using EasyConnect in an enterprise?
3. Why is MAB alone considered insecure?
4. What are some of the advantages of CWA over LWA?
5. For remote access, how are credentials typically sent from the VPN headend to ISE?



CHAPTER 5

Introduction to Advanced Concepts

This chapter covers the following topics:

Change of Authorization

Automating MAB

Posture assessment

Mobile device management (MDM)

The features and capabilities available in an authentication server are designed to accommodate a wide variety of deployments—from simple to complex. These features and capabilities allow network administrators to appropriately stage their deployments, implementing only the basic authentication and authorization features that are required in a network and eventually adding more advanced features. Advanced features provide for more comprehensive and often more granular deployments as they provide greater depth and breadth for network security. This chapter focuses on several of these advanced concepts.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 5-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 5-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Change of Authorization (CoA)	1, 2
Automating MAC Authentication Bypass (MAB)	3–6
Posture Assessments	7, 8, 9
Mobile Device Management (MDM)	10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following does a RADIUS Change of Authorization (CoA) allow an authentication server to do?
   a. Escalate an administrative user's access level within the server's administration portal
   b. Grant context-appropriate network access after initial access has previously been granted
   c. Gain root-level access to all network devices
   d. Grant context-appropriate operating system access
2. Which of the following are three possible options for Change of Authorization actions?
   a. IKEv1, IKEv2, SSL
   b. HTTP, FTP, Telnet
   c. No COA, port bounce, reauthenticate
   d. User mode, privileged mode, configuration mode
3. MAC Authentication Bypass (MAB) is a process by which a device does which of the following?
   a. Bypasses all authentication and authorization processes by using a supplicant
   b. Authenticates with an X.509 certificate to establish a secure tunnel with the network
   c. Authenticates without a 802.1X supplicant on the endpoint by using its MAC address as the RADIUS identity
   d. Hides its MAC address from being discovered on the network
4. A MAC address is six octets in length. Which of the following is true of the first three octets?
   a. They are a duplicate of the IP address subnet in hexadecimal format.
   b. They are always the same across all network devices.
   c. They are assigned dynamically upon connection to the network.
   d. They are an organizationally unique identifier (OUI) that indicates the device's vendor.
   e. They are all F's (that is, FF:FF:FF).
5. What devices often lack an 802.1X supplicant?
   a. Printers
   b. Laptops
   c. Cell phones
   d. All of these answers are correct.

6. Prior to the introduction of MAB, a switch port with a non-802.1X client would be configured without 802.1X. Why did this present issues?
   a. A broadcast storm would be created as the endpoint device was plugged in to the interface.
   b. A non-802.1X client would still not be able to gain network access.
   c. A rogue user could unplug the non-802.1X endpoint and gain unauthorized access to the network.
   d. Rebooting the device would cause the switch port to go into an error disable state.
7. Posture checking can check for which of the following?
   a. File conditions, including existence, date, and version
   b. Registry condition (that is, whether a registry entry is present) on Windows-based endpoints
   c. Service condition (that is, whether a service is running) on Windows-based endpoints
   d. All of these answers are correct.
8. When configuring authorization policy based on posture assessment outcome, which of the following values are available for the PostureStatus attribute?
   a. Permit, Deny, Drop
   b. Compliant, NonCompliant, Unchecked
   c. Internet Only, Partial Access, Full Access
   d. Compliant, NonCompliant, Unknown
   e. AntiVirusNotPresent, AntiVirusNeedsUpdate, AntiVirusCurrent
9. To remediate noncompliant endpoints, a redirect ACL must be defined _____, and the web redirection must be destined to a(n) _____ portal on the authentication server.
   a. as a dACL; remediation
   b. on the switch; remediation
   c. as a dACL; profiling mitigation
   d. on the switch; profiling mitigation
   e. as a dACL; authentication DMZ
   f. on the switch; authentication DMZ
10. Which of the following best characterizes a mobile device manager?
   a. A network administrator responsible for onboarding all mobile devices into the authentication server
   b. An application that runs on a mobile device, allowing the user or endpoint to manage the authentication server and other network devices
   c. A wireless access point that detects rogue mobile endpoints
   d. A third-party appliance or service that provides advanced posture assessment for mobile endpoints

Foundation Topics

Change of Authorization

In order for a network to be secure, it is paramount that no individuals or endpoints be allowed access to the network or network resources until they have been sufficiently authenticated. As discussed in Chapter 1, "Fundamentals of AAA," through Chapter 4, "Non-802.1X Authentication," a number of different identity stores and authentication methods can be used to authenticate a user or device. This authentication establishes who the user or endpoint is. For instance, through authentication, it is possible to determine whether a user is part of the network administrators group of users or whether the user is a more basic user.

Following—and often related to or based on—authentication, the user or endpoint is authorized to access the network at some level. Again, as mentioned in earlier chapters, authorization is the process of determining what resources an authenticated user can have access to. This level of authorization could allow the user unfettered access to all network resources, access to none of the network resources, or a level of access that is somewhere between these two extremes. Being able to provide different levels of access to network administrators and to basic users is critical for network security; it is a best practice to give users the least level of access that they require in order to accomplish their duties.

However, the determination of a user's access needs is not always clear-cut. Sometimes a user, an endpoint, or the network changes after the initial authentication/authorization process. For example, an endpoint could be stolen, become infected, or otherwise fall out of the good graces of the network administrators. The network must be able to react to such changes and provide updated levels of access, or authorization, to these endpoints.

Key Topic

Change of Authorization (CoA) is a standard defined in RFC 5176 that provides a mechanism to change the AAA attributes of a session after it has already been authenticated. It functions as a push model in which the request originates from the authentication server (ISE) and enables dynamic reconfiguration of the session. As an endpoint's level of authorization changes, the authentication server can elect to do nothing (that is, not issue a CoA), perform a port bounce (that is, determine whether to shut the port), or instruct the network access device to reauthenticate the user.

The following per-session requests can be initiated as part of CoA:

- Session reauthentication
- Session reauthentication with rerun
- Session termination
- Session termination with port shutdown
- Session termination with port bounce
- Security and password
- Accounting

This CoA process gives you an opportunity, via the authentication server, to rethink or reexecute the authorization of an endpoint as the network or device changes—and implement the advanced concepts discussed in this chapter.

Automating MAC Authentication Bypass (MAB)

As discussed as part of Chapter 4, MAC Authentication Bypass (MAB) is the process by which an endpoint bypasses authentication by using its MAC address alone. To accomplish this, the authenticator network device crafts a RADIUS Access-Request message on behalf of the endpoint. This Access-Request uses the endpoint's MAC address in both the Identity and Password fields with the RADIUS Service-Type attribute set to Call Check. This indicates to the RADIUS server that the endpoint does not have or is currently not using an 802.1X supplicant. These devices are sometimes referred to as "dumb," or headless, devices.

An authentication server may be configured to set the RADIUS Service-Type attribute to Call Check and authenticate the device based on its MAC address. If the device's MAC address is present in the authorized dumb devices database on the authentication server, the device is authenticated and subsequently authorized according to the configured authorization policy. If desired, devices can be subdivided into different endpoint identity groups (for example, printers, IP phones) and provided different levels of network access upon successful authentication, depending on the chosen authorization policy.

Key Topic

NOTE As discussed at length in Chapter 4, MAB is not a secure technology, and it is best to leverage MAB as a secondary "authentication" method—not the primary method for allowing network access. Please keep this in mind as you implement MAB on your network.

Managing MAB devices can quickly become a network administrator's worst nightmare. The dumb devices that require MAB include IP phones, IP cameras, and printers—among many other devices. Managing these devices manually in an organization can quickly get out of hand even for a small to medium-sized business. Any time a new dumb device is placed on the network, its MAC address needs to be added to an internal endpoints database on the authentication server, sometimes through the MAB Devices list. If a user gets a new IP phone at his desk because he poured coffee on his last one or if the printer management company replaces a printer, a network administrator must manually update the MAB database with the new device.

Some authentication servers can help automate this process for you via device profiling. ***Profiling*** is the process by which the authentication server can make educated observations about what the device is, based on various attributes of the endpoint. We discuss some of the profiling process now and leave a more in-depth discussion of profiling for Chapter 14, "Profiling."

As a device connects to the network, the network access device (NAD) is configured for 802.1X, oftentimes with MAB as a backup authentication method but rarely with just MAB. (Remember that MAB is not secure!) In either case, the NAD begins the RADIUS exchange with the authentication server, forwarding this server critical information about the authentication session, including the MAC address of the endpoint, which is associated with the authentication session ID. The Calling-Station-Id field in the RADIUS packets provides the MAC address of the endpoint requesting access to the network, and the session ID is a unique identifier associated with that RADIUS session (and, therefore, uniquely associated to a MAC address).

As the device is initially connecting to a network, it goes through its normal network processes, including requesting and receiving an IP address, discovering and advertising its neighbors, and so on. At each step along the way, these various network exchanges can produce attributes that can be sent to the authentication server and associated with the MAC address of the device. These attributes form a "fingerprint" of the device that can give clues about what type of device is connecting to the network.

To be most effective at profiling an endpoint, a number of profiling options must be enabled on the authentication server in order to facilitate and gather the information that is being provided by the endpoint. The two profiling options that are most useful when automating the MAB enlistment process are RADIUS and DHCP.

With RADIUS profiling, the authentication server can gather, record, and correlate the information contained within a RADIUS packet into an endpoint database, making intelligent observations about the device along the way. From the very first RADIUS packet that is received on the RADIUS server for the endpoint, observations are made and recorded. With each RADIUS exchange, the RADIUS Calling-Station-Id and session ID are provided; this makes it easy for the RADIUS server to know which endpoint the newly found information is about.

The MAC address alone can reveal quite a bit of information about an endpoint. As shown in Figure 5-1, a MAC address is divided into two main portions: The first three octets (for example, 00:1D:71 in Figure 5-1) are the organizationally unique identifier (OUI), and the last three octets (9B:06:40 in Figure 5-1) are the network interface card (NIC)–specific portion. Where a governing body such as the IEEE issues the OUI, the NIC-specific portion is unique per device and is typically the serial number that is assigned by the device manufacturer.

With the OUI portion of the MAC address, some observations can immediately be made. Certain manufacturers are associated with producing a certain type of network device, such as a network printer or IP phone. Therefore, whenever that particular OUI is seen on the network, it is a safe assumption that the owner of that MAC address is a network printer. The authentication server maintains a database of the OUIs associated with certain device vendors as well as the specific device types that are associated with these vendors, allowing the RADIUS server to dynamically associate a device with a particular network function or device group.



Figure 5-1 *MAC Address*

For vendors that produce a large number of device types under a single OUI, such as Cisco (which makes routers, switches, IP phones, IP cameras, and so on), additional profiling attributes are often collected and leveraged. Some network access devices, such as switches and wireless LAN controllers (WLCs), can provide the authentication server with additional attributes by using a feature called Device Sensor. The Device Sensor feature was introduced in Cisco switches starting with IOS Version 15.2(1)E, IOS XE 3.3.0SG, and later releases. With the Device Sensor feature, the network access device can gather additional information from the endpoint and provide it to the authentication server for classification. "Neighbor" protocols such as Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) provide some very useful information to the authentication server. These protocols are often confined to the Layer 2 domain; if the authentication server is often not on the same Layer 2 domain, then another device (such as the ones using Device Sensor) must tunnel, port mirror, or otherwise relay this information to the authentication server. Within these two protocols, additional attributes such as the device capabilities, a description of the device, the hostname, platform type, and other useful information can be learned. Device Sensor may also optionally sniff DHCP packets using DHCP snooping and send the information to the authentication server. All this information can then be encapsulated into a RADIUS accounting packet or similar mechanism and sent to the authentication server. The authentication server can then parse these packets and marry the additional information that was learned with the original MAC address and RADIUS session ID, further evolving the dossier for the device.

Any other additional field, or combination of fields, contained within the RADIUS authentication, authorization, or accounting exchange can also be leveraged to make decisions about devices' final profiling. From the totality of information gathered from the RADIUS processes for the endpoint, many MAB devices can be automatically added to the appropriate endpoint list(s) and assigned the appropriate level of authorization.

DHCP profiling can also be a useful mechanism for developing an endpoint's dossier. As most devices join a network, they request an IP address using DHCP. This DHCP process can easily be sent to the authentication server—simply by adding an additional "IP helper" address to the interface on the switch and pointing to the profiling server or turning on Device Sensor on the NAD. As the authentication server receives the DHCP packets from the endpoint, the server leverages fields from the DHCP packets (which also, subsequently, contain the MAC address of the endpoint) to determine what type of device the endpoint could be. Some of the information within DHCP that the authentication server can leverage include the device DHCP client identifier (or MAC address—which contains the device OUI), the DHCP requested address (which allows the authentication server to further correlate MAC-to-IP mapping), and the DHCP parameter request list (which can also act as a device-type fingerprint).

By analyzing the complete DHCP process and exchange, the authentication server can further establish an endpoint's identity with a reasonable level of certainty in its database.

As a device moves through this RADIUS and DHCP profiling process, the level of access to the network is likely to need to change. By leveraging RADIUS and DHCP profiling, a network administrator can quickly and automatically identify many dumb network devices as printers, IP phones, IP cameras, and so on and permanently associate each device with a particular endpoint list(s). This process can also be used to profile "smart" devices, allowing the network administrator to make authorization decisions specific to a particular device type (for example, a printer that is only allowed access to the network on certain ports, an IP phone that is allowed access to a certain Cisco Unified Call Manager). Devices that are not easily identified with certainty could even be assigned to an Unknown device list, given a basic level of network access, and further triaged manually, as needed.

In order to trigger a new level of authorization, you can leverage a Change of Authorization (CoA). Say that a device was given basic access to the network in order for the authentication server to gather the required RADIUS and DHCP information from the device. Now that the authentication server has a more updated view of the device, the authentication server can, through the NAD, provide a more open level or a more strict level of access to the endpoint, as necessary. To invoke this reauthorization process, the authentication server sends a CoA to the NAD with the relevant session ID for the endpoint. The NAD then forces the device to reauthenticate, and it subsequently reauthorizes the device with a different level of network access.

Posture Assessment



Authentication is often focused on the user of an endpoint. As the user brings a device onto the network, network administrators often look at who is using device without putting any focus on what device is being used. Historically, the endpoint devices on a network were heavily managed by IT, and the operating system was sufficiently locked down so that the end user could effectively change nothing.

This change control—allowing only IT to modify settings and applications on the PC—slowly started to dwindle as end devices became more mobile. Now an end user often requires some level of flexibility to modify his or her laptop PC without IT intervention. This is important for allowing the end user to add home Wi-Fi parameters, install productivity applications, and install device drivers when working remotely from a home office or using a hotspot.

How do network administrators ensure that the changes to the laptop made by the end users in the sanctity of their homes are not going to compromise the security of the corporate network? How do they ensure that the device continues to comply with the company security policy? How do they react to a device that is out of compliance? This is where device *posture* assessment comes in.

A posture assessment, or simply posturing, is a process by which an application (such as the Cisco AnyConnect ISE Posture Module) running on an endpoint or a temporary executable file (such as the Cisco Temporal Agent) running on the client provides critical information about the software or operating system that is actively running on the device. Several conditions can be checked as part of posture assessment. The availability of these conditions or

their exact function may vary depending on the authentication or posturing server that is being used:

- **File condition:** Checks the existence, date, or version of a file on the endpoint.
- **Firewall condition:** Checks whether a specific firewall product is running on the endpoint.
- **Registry condition:** Checks the existence or value of a registry key on the endpoint.
- **Application condition:** Checks whether an application is running or installed on the endpoint.
- **Service condition:** Checks whether or not a service is running on the endpoint.
- **Dictionary simple condition:** Checks an attribute that is associated to a user and value.
- **Windows Update verification:** Confirms if Windows Update is enabled or not.
- **Antivirus application verification:** Confirms that the endpoint has an antivirus solution installed and confirms whether specific software is required.
- **Antivirus definition verification:** Confirms that the endpoint has the virus definition files present and that they are newer than a particular date or the latest definition file.
- **Antispyware application verification:** Confirms that the endpoint has an antispyware solution installed.
- **Antispyware** definition verification – Confirms that the endpoint has an antispyware definition file present and that it is newer than a particular date.
- **Antimalware application verification:** Confirms that the endpoint has an antimalware solution installed and confirms whether specific software is required.
- **Antimalware definition verification:** Confirms that the endpoint has the antimalware definition files present and that they are newer than a particular date.
- **Patch management condition:** Checks that the endpoint has selected patch management products installed, enabled, and up to date.
- **USB condition:** Checks whether a USB storage device is connected to the endpoint.
- **Disk encryption condition:** Checks that disk encryption software is installed and checks the encryption state of the endpoint.
- **Hardware attributes condition:** Checks the hardware attributes of the endpoint.
- **External data source condition:** Checks Active Directory for specified attributes for the endpoint.

This information about the endpoint is sent to the authentication server. The server compares the data to the security policy as configured and assesses the device as Compliant, Non-Compliant, or Unknown. This assessment outcome can then be used as part of the authorization decision for the endpoint to determine what this endpoint will have access to. If the endpoint is compliant with configured security policy, the device can be given access to the

network based on the selected authorization policy. If the endpoint is assessed NonCompliant or Unknown, the resulting authorization policy could redirect the endpoint to a remediation portal by using a redirect ACL that is defined on the switch. If the posture status is Unknown, the posture agent is installed or run on the endpoint so a proper posture assessment can be completed.

After a device has been appropriately remediated via the portal page, a CoA can be issued to the NAD to reauthenticate the user. A more in-depth review of posturing configuration is provided in Chapter 18, "Posture Assessment."

Mobile Device Management (MDM)

Key Topic

A ***mobile device manager*** is a security software system that allows an administrator to configure and secure a mobile device, regardless of where the device is located. As an endpoint initially onboards with a mobile device manager, the owner of the endpoint installs MDM software and gives it permission to access the device and gather any required information. This is accomplished by leveraging a mobile device manager agent on the endpoint—either a separate application or an agent that is built into the OS—and a server. The MDM software performs a number of assessments on the device, identifying the current state of security, software, and networking, among other factors. This information is then passed to the MDM server for processing. This MDM server can reside on-premises at the corporation or in the cloud. In addition, other supporting servers that reside out in the cloud supplement the operation of the mobile device manager, including Apple's Push Notification servers and Google's Client Messaging servers.

There are a number of MDM vendors available. Each of these vendors may provide comprehensive device provisioning, detailed user and device context, and increased device and application security. Depending on the user base, cost, preferred deployment model (on-premises or in the cloud), and security policy, you might choose to go with one MDM vendor over the other.

The following are some of the features that are supported almost ubiquitously across all MDM vendors:

- **Support for most mobile and laptop platforms:** For example, Android, Apple iOS, Mac OSX, or Windows
- **Application provisioning:** For example, pushing required applications to a device
- **Security posturing:** For example, whether a device has PIN lock enabled, has encryption enabled, or is jailbroken
- **Modify device capabilities:** For example, whether Wi-Fi, GPS, Bluetooth, the camera, and so on are enabled
- **Remote wipe:** A partial or full wipe on a device that is lost or stolen
- **Active Directory support:** Ability to authenticate MDM users based on AD credentials
- **VPN provisioning:** Deploying VPN configurations so the endpoint can securely connect to the corporate network
- **Device backup:** Ensuring that the device has been backed up recently

Many of the features listed here are not directly available to the authentication server's posture assessment process and provide additional control over the endpoint even when it is off the corporate network. Depending on the authentication server, some features may be explicitly communicated to the authentication server. Otherwise, in absence of these explicit features, the MDM server can provide a final summary assessment of Compliant or Non-Compliant to the authentication server.

Mobile device managers, either via an application on an endpoint or via software features embedded natively in the operating system, can see a great amount of detail. Mobile device managers also provide a greater level of security remotely, for user and administrators alike—allowing an endpoint to be locked remotely and wiped securely from a portal available within the mobile device manager or by proxy via the authentication server.

To take advantage of MDM features, an authentication server's authorization policy can query the status of an endpoint as compliant or noncompliant, according to the mobile device manager, as well as check the status of several other explicit endpoint statuses from the mobile device manager (such as jailbroken, PIN lock, and so on). If a device needs to be redirected to register with the mobile device manager or for noncompliance, the authentication server's authorization policy might perform a redirect, sending the endpoint to an MDM portal page for information on further mitigation and remediation.

As corporations continue to accept mobile devices as viable and often critical business tools, the network availability from these devices, the applications available from these tools, and the security of these devices remain paramount. A mobile device manager is one more tool to manage and enforce security policy on these devices.

Exam Preparation Tasks

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 5-2 lists these key topics and the page number on which each is found.

Key Topic

Table 5-2 Key Topics

Key Topic Element	Description	Page
Paragraph	Change of Authorization (CoA)	95
Note	MAC Authentication Bypass (MAB)	96
Section	Posture assessment	99
Section	Mobile device management	101

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Change of Authorization (CoA), profiling, posture, mobile device manager

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is a Change of Authorization?
2. What can be used to help automate MAB?
3. What does profiling do?
4. What does posturing help accomplish?
5. What are some of the benefits that a mobile device manager can provide that cannot be done with posturing?



CHAPTER 6

Cisco Identity Services Engine Architecture

This chapter covers the following topics:

Cisco ISE basics

Personas

Physical or virtual appliances

Cisco ISE deployment scenarios

Up to this point in the book, we've looked generically at several key components of a network security architecture, highlighting fundamentals such as AAA, identity management, and EAP/802.1X. This chapter looks specifically at Cisco's Identity Services Engine. It provides a general description of Cisco ISE, the various personas that ISE can assume, physical and virtual instantiations of ISE, and several common Cisco ISE deployment scenarios.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 6-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 6-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
What Is Cisco ISE?	1
Personas	2–6
Physical or Virtual Appliances	7
ISE Deployment Scenarios	8–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Cisco Identity Services Engine (ISE) is:
   a. A sports medicine approach to heal swollen and aching muscles
   b. A network management platform
   c. A network security and policy platform
   d. A unified computing system that incorporates virtualization of endpoints
2. Which of the following are Cisco ISE personas? (Select all that apply.)
   a. Administration
   b. Authentication Server
   c. File Download
   d. Monitoring
   e. Policy Services
   f. Identity Management
   g. pxGrid
3. Which of the following is true of the Cisco ISE Administration persona?
   a. It is where policy configuration changes are made.
   b. It is the network management platform for the network.
   c. It is the engine where policy decisions are made.
   d. It is responsible for logging and reporting data.
4. Which of the following is true of the Cisco ISE Monitoring persona?
   a. It is the node where policy configuration changes are made.
   b. It is the network management platform for the network.
   c. It is the engine where policy decisions are made.
   d. It is responsible for logging and reporting data.
5. Which of the following is true of the Cisco ISE Policy Services persona?
   a. It is the node where policy configuration changes are made.
   b. It is the network management platform for the network.
   c. It is the engine where policy decisions are made.
   d. It is responsible for logging and reporting data.
6. Which of the following is true of how Cisco ISE responds to the Network Access Device (NAD)?
   a. ISE responds to the NAD's RADIUS request with RADIUS attribute/value pairs (AV pairs) that contain the security policy details.
   b. ISE uses a REST API to respond to the NAD and configure the security policy.
   c. ISE communicates to the Cisco Security Agent on the endpoint to control network access.
   d. ISE and the NAD do not communicate.

7. Which of the following characteristics should a virtual ISE appliance have? (Select two.)
   a. It should be as small as possible for speed and agility.
   b. It should be sized appropriately to match the equivalent physical appliance.
   c. Appropriate resources should be reserved to ensure that other virtualized network functions do not cannibalize the ISE resources.
   d. It should leverage shared storage, such as a NAS or SAN.
8. Which of the following is true for a single-node, or standalone, Cisco ISE deployment?
   a. Each Cisco ISE appliance services a single network access device.
   b. Each Cisco ISE appliance services only a single ISE persona.
   c. All endpoints bypass authentication.
   d. All core Cisco ISE personas reside on a single ISE appliance.
9. In a hybrid Cisco ISE deployment, the ____ and ____ personas are combined on two of the appliances, and the ____ persona is by itself on each of the other two appliances.
   a. Administration, Policy Services, Monitoring
   b. Administration, pxGrid, Monitoring
   c. Policy Services, Monitoring, pxGrid
   d. Policy Services, Administration, Monitoring
   e. Administration, Monitoring, pxGrid
   f. Administration, Monitoring, Policy Services
10. The maximum number of Policy Services nodes supported with ISE Version 2.7 in a fully distributed deployment model is ____, resulting in a maximum of _____ supported endpoints.
   a. 5; 5000
   b. 5; 10,000
   c. 5; 50,000
   d. 50; 500,000
   e. 40; 20,000
   f. 50; 2,000,000

Foundation Topics

What Is Cisco ISE?

The Cisco Identity Services Engine (ISE) is a security policy management platform and a key component of the Cisco's zero trust architecture. Its role in the network is to allow the network to be the enforcement point for network security.

Often, network security is delegated to individual devices in the network. For instance, you might allow unfettered access for all endpoints in the core of your corporate network and enforce the access policy at the edge firewall. For wireless users, you might choose to

enforce a single policy for all users—such as allowing every wireless user access to HTTP, HTTPS, SSH, and/or Telnet—and implement this policy at the access point (in autonomous mode) or at the wireless LAN controller (in lightweight mode). This one-size-fits-all approach is not the ideal way to implement network security.

Cisco ISE helps facilitate the paradigm shift to making the network (rather than a single device) the enforcement point for network security. This is possible thanks to Remote Authentication Dial-In User Service (RADIUS), which provides the authentication, authorization, and accounting (AAA) functions for the network (as discussed in Chapter 2, "Identity Management"). Cisco ISE acts as a centralized network security policy platform and RADIUS server, extending the AAA function to all network devices. As a user or an endpoint attempts to access the network, the network access device (NAD) forwards all relevant authentication parameters to Cisco ISE. Cisco ISE responds back to the NAD with the resulting security policy to be applied to the user/endpoint by using RADIUS attribute/value pairs (AV pairs).

The authorization policy that is sent to the NAD is implemented on a per-user basis. Tightly coupling the authorization of the user with the authentication allows a network administrator to provide differentiated network access to all users. If you are a basic user, your level of network access is likely going to be different than if you are a network administrator. Your level of access may also be different if you are accessing the network using a wireless connection than if you are on a wired connection, and it may be different if you are accessing the network from a remote branch location than if you are accessing it from the campus network. The granular policy building blocks that are available in Cisco ISE ensure that you can implement the best network security policy—on either a per-user or per-group basis.

Cisco ISE also provides a number of advanced functions. Profiling, which is natively integrated in Cisco ISE, can dynamically identify endpoints. As an endpoint attempts to access the network, ISE can profile the device—identifying, managing, and taking inventory of the devices accessing the network based on a predefined security policy. This profiling operation can determine the manufacturer of the endpoint, the function of the endpoint (IP phone, IP camera, or network printer) and conduct other network-level assessments of the endpoint. The result of this profiling operation can also be used as one more criterion in the authorization policy that is applied to the endpoint.

Another advanced function that is available with Cisco ISE is posturing. Posturing can go slightly deeper than profiling, and it uses a slightly different approach. Whereas profiling uses network-level communications to determine information about the endpoint, posturing leverages information that resides on the endpoint. Using a network access control (NAC) agent that resides on the endpoint directly, posturing ensures that the endpoint is adhering to security policies that have been deployed to the endpoint—policies such as antivirus software, antispyware software, and other security (and even non-security) software and services. If the posture adheres to the implemented security policy on Cisco ISE, additional network access can be allowed via the authorization policy. Conversely, if the endpoint posture does not adhere to the current security policy, a reduced level of network access can be deployed to the NAD on behalf of the endpoint—possibly forcing the endpoint to remediate its noncompliance before it can gain access to the network.

The depth of posturing that can be leveraged via Cisco ISE can be further enhanced with third-party software security systems called mobile device management (MDM) systems, as described in Chapter 18, "Posture Assessment." Depending on the MDM system, Cisco ISE can leverage additional industry compliance requirements (including HIPAA and PCI-DSS requirements).

Much of the network security policy that we have discussed thus far pertains to corporate-owned assets: assets that are owned and managed by the enterprise IT department. However, Cisco ISE has a number of features that allow a granular network security policy to be extended to employee-owned or guest endpoints. The employee-owned endpoints, depending on the corporate security policy, may be allowed trusted access equivalent to that of a corporate-owned asset or basic Internet-level guest access.

Trusted access can be implemented in a number of different ways—either by dynamically onboarding the endpoint into the Cisco ISE database (possibly by using an identity management system to identify the user of the device) or by manually adding the device to a trusted-device database. Both of these approaches are discussed in other chapters, such as Chapter 16, "Bring Your Own Device," and Chapter 18, "Posture Assessment."

Guest-level access, leveraging built-in guest portals, is also available natively within ISE. The guest portals available in ISE allow a trusted employee to create temporary guest credentials for visitors, contractors, or other temporary users. These guest credentials are then given to the guest user, either manually or via email, to log in to the network. As the guest user attempts to access the network, a second network portal (or network-level authentication such as PEAP) is be used to authenticate the user—pushing the relevant guest-user access policy to the network access device (wired or wireless) from which the user is accessing the network.

Because Cisco ISE provides a high level of differentiated access to corporate endpoints, employee endpoints, and guest endpoints—enforcing profiling, posturing, and MDM influenced authorization policy—it is paramount that network administrators be able to monitor the level of access granted to users/endpoints. This, too, is a feature that is native to Cisco ISE. As each authentication, authorization, profiling, or posturing action occurs, Cisco ISE creates a logging event. These logging events can be filtered, sent to syslog servers, and otherwise searched by network administrators to see details on a user's network authentication and authorization.

Cisco ISE is a one-stop shop for network security policy. By combining the industry-trusted RADIUS protocol with the granular, context-based policy that is available with ISE, a network administrator can distribute a unified security policy to the far reaches of the network and ensure that the security policy deployed to each user/endpoint is as secure as possible.

Personas

Given the many roles that Cisco ISE may play in an enterprise environment, it requires a distributed architecture to handle those responsibilities at scale; these responsibilities are known as ***personas***.

Cisco ISE is designed with scalability in mind. To achieve scalability, it divides the roles it provides into separate personas. Each persona is a different function within a Cisco ISE deployment that is required for proper operation of the platform, and personas run within ISE nodes.

An ISE ***node*** is an individual appliance (virtual or physical) configured to run one or more of the four personas that exist within an ISE deployment. (An ISE deployment is affectionately referred to an ***ISE cube***.) A single ISE node may be configured with just one persona, with multiple personas, and even with all personas in the case of a standalone or two-node deployment. The four personas are as follows:

- Administration
- Monitoring
- Policy Services
- pxGrid

The sections that follow describe these four personas in detail.

Administration Persona

The first, and possibly the most important, persona is the Administration persona. This function of Cisco ISE is important, as it enables a network administrator to configure and administer the network security policy.

To make a change to the security policy—whether it is the authentication, authorization, profiling, posturing, or other policy—on Cisco ISE, a network administrator needs to access the Administrative persona's graphical user interface (GUI).

NOTE The administrative portal used to be a Flash-based web interface, but all Flash was finally removed in ISE Version 2.4.

An ISE node that runs the Administration persona is referred to as simply the *Admin node* or as the ***Policy Administration node (PAN)***; it is the central control center for Cisco ISE, and all policies in ISE are configured in the PAN and then synchronized to all other ISE nodes/personas in the ISE cube. The PAN also includes the licensing function for the platform, allowing the network administrator to enable advanced features or additional functionalities in Cisco ISE by installing purchased licenses.

There can be a maximum of two nodes with the Administration persona: primary PAN and a secondary PAN for redundancy. A PAN also acts as the root for the built-in certificate authority.

Monitoring Persona

The Monitoring persona acts as the centralized logging server for the ISE cube. An ISE node with this persona enabled is referred to as a ***Monitoring and Troubleshooting (MnT) node.*** There can be a maximum of two MnT nodes in an ISE cube: a primary MnT node and secondary MnT node for redundancy.

The function of this node is exactly as it sounds: to provide monitoring and troubleshooting functions for the Cisco ISE cube. As an endpoint authenticates to the Policy Services node (PSN), events are created to keep track of the authentication and authorization process. These events are forwarded to the MnT node, which consolidates and processes these events into a legible format. When a network administrator requires reports to be created

for whatever purpose—managerial slides and presentations, access reports, and so on—this function is also provided by the MnT node.

The second function of the MnT node is troubleshooting. As each event is forwarded to the MnT node, there can be a success or failure of the authentication or authorization process. Due to the detailed event tracking that happens on the PSNs, the content of the logs present on the MnT node are also quite detailed. If a user/endpoint is having any issue accessing the network, a network administrator can leverage the MnT node to see the cause of the failure as the node tracks the user's/endpoint's authentication in a play-by-play manner.

There can be a large number of PSNs deployed in a single ISE cube, and the MnT node acts as a centralized recipient of the logs. Without the central monitoring and troubleshooting function that the MnT node provides, each PSN would be responsible for its own reporting—while actively participating in ongoing authentication and authorization processes. This can be very resource prohibitive. Furthermore, network administrators need to check the logs on each PSN for event correlation; the MnT node provides this function in a single, one-stop-shop approach.

Policy Services Persona

The Policy Services persona serves many functions and is often referred to as "the workhorse." A node running the Policy Services persona acts as the RADIUS server for the ISE cube, handing the authentication requests, performing the identity lookups and policy evaluation, and issuing the resulting authorization result. A node running the Policy Services persona is referred to as a ***Policy Services node (PSN)***.

After an administrator makes policy changes on the PAN, all policies are synchronized to the PSNs. Each NAD is configured to leverage one or more of the PSNs as RADIUS servers. Additional scalability can be achieved by leveraging a network load balancer, taking into account key implementation requirements, including No-NAT, Certificate Subject Alternative Names, and PSN "stickiness," among other factors.

When a user authenticates to the network, the NAD sends the RADIUS Access-Request and subsequent packets to the configured PSN(s). The PSN(s), having the complete security policy as pushed by the PAN, can authenticate the user and provide the requisite authorization policy in an autonomous fashion (that is, it does not need to communicate back to the PAN in order to perform the function).

If the user endpoint is a candidate for any profiling and/or posturing that is done by ISE, the PSN is also responsible for deploying a ***Change of Authorization (CoA)*** message to the NAD. The NAD needs to also have the dynamic authorization configuration present, which enables the PSN to force a reauthentication. The CoA is created by the PSN and sent to the NAD without any active participation by the PAN.

The PSN is also responsible for hosting the different web portals for guest access and sponsorship, as well as being the issuing CA for the built-in ISE CA. A PSN is also the ISE persona handling the profiling of endpoints, and it may also provide additional services, such as Threat Centric Network Access Control (TC-NAC), SGT Exchange Protocol, device administration (TACACS+), and the passive identity service.

There can be up to 50 PSNs per ISE cube for scalability and distribution of PSN functions and load.

pxGrid Persona

The final persona in an ISE cube is the ***pxGrid*** persona. pxGrid is Cisco's premier publish and subscribe (pub/sub) communication bus, and it was designed from the ground up to be a scalable secure data-sharing system. In other words, it is designed to share a lot of security data in an efficient and scalable manor.

Physical or Virtual Appliances

Now that you are aware of the various ISE personas, you are better equipped to start the deployment of Cisco ISE on a network. However, there are still a few more decisions you need to make as you plan for your deployment. The first of these decisions is whether to deploy physical ISE appliances or whether to virtualize ISE on your network. This section discusses both options.

Your decision of whether to go with a physical and/or virtual appliance ISE deployment depends on a number of factors. Several of the key factors impacting this decision include deployment size, budget, scalability, current network topology, and, most importantly, the politics of your organizational structure (that is, who has ownership of the VMware environment in your organization).

As we start this discussion of physical versus virtual ISE deployments, we should first highlight the options for each approach. The physical appliances are called Secure Network Server (SNS) appliances. This line of physical appliance has gone through a number of hardware refreshes over the years. It started with the SNS-33xx models, and the latest models as this book went to press were the SNS-36xx appliances.

The SNS-36xx appliance comes in three form factors: small (SNS-3615), medium (SNS-3655), and large (SNS-3695). Table 6-2 outlines the physical appliance specifications.

NOTE See the Secure Network Server data sheet at www.cisco.com/c/en/us/products/collateral/security/identity-services-engine/datasheet-c78-726524.html for more information.

Table 6-2 Cisco ISE Physical Appliance Specifications

Feature	SNS-3615	SNS-3655	SNS-3695
Processor	1 Intel Xeon 2.10 GHz 4110	1 Intel Xeon 2.10 GHz 4116	1 Intel Xeon 2.10 GHz 4116
Cores per processor	8	12	12
Memory	32 GB (2 16 GB)	96 GB (6 16 GB)	256 GB (8 32 GB)
Hard disk	1 2.5-inch 600 GB 6 Gbps SAS 10K RPM	4 2.5-inch 600 GB 6 Gbps SAS 10K RPM	8 2.5-inch 600 GB 6 Gbps SAS 10K RPM
Hardware RAID	No	Level 10 Cisco 12 Gbps SAS Modular RAID controller	Level 10 Cisco 12 Gbps SAS Modular RAID controller

Feature	SNS-3615	SNS-3655	SNS-3695
Network interfaces	2 10GBASE-T 4 1GBASE-T	2 10GBASE-T 4 1GBASE-T	2 10GBASE-T 4 1GBASE-T
Power supplies	1 770W	2 770W	2 770W

If you choose to implement a virtual appliance, you need to take great care to ensure that virtual machines are created with hardware resource reservations that are equal to the specs of the physical appliances. This is very critical. We repeat: This is very critical. There have been a tremendous number of TAC cases around ISE deployments and corrupt databases (or similar). The vast majority of those cases have been related to the virtual machines not having the correct hardware reservations. This problem is so common that Cisco started posting open virtual appliances (OVAs) with hard-coded preconfigured resource reservations for customer download and use instead of manual installation. Using the OVAs ensures that the reservations are statically configured; they should remain configured. Table 6-3 lists the files available for creating an ISE Version 2.7 virtual appliance.

Table 6-3 Files Available for Creating an ISE 2.7 Virtual Appliance

Filename	Description
ise-2.7.0.356.SPA.x86_64.iso	ISO file to download for manual installations.
ISE-2.7.0.356-virtual-SNS3615-SNS3655-300.ova	Small form factor with 300 GB disk. Used for evaluations or dedicated PSNs.
ISE-2.7.0.356-virtual-SNS3615-SNS3655-600.ova	Small form factor with 600 GB disk. Used for small deployments with PANs or MnT nodes.
ISE-2.7.0.356-virtual-SNS3655-SNS3695-1200.ova	Larger form factor with 1.2 TB disk. Use for large deployments with PANs or MnT nodes.
ISE-2.7.0.356-virtual-SNS3695-2400.ova	Largest form factor with 2.4 TB disk. Use this for a large MnT node.

Table 6-4 lists the different SNS appliance models and their respective scalability numbers.

Table 6-4 SNS Appliance Scalability

Endpoints Supported	SNS-3615	SNS-3655	SNS-3695
Standalone configuration	10,000	25,000	50,000
Per PSN	10,000	50,000	100,000

If your company has a highly virtualized server infrastructure already, you might consider the virtual appliance approach, leveraging existing Cisco UCS infrastructure. With the virtualized appliance approach for your ISE deployment, it is best to ensure resource reservations on the virtual machine for the entire number of virtual processors, memory, and hard disk space (refer to Table 6-2). This will guarantee that other virtual machines are not impacting the proper operation of ISE. Failure to reserve the recommended resources may greatly impact the overall performance and scalability of the Cisco ISE cube.

Whether you choose to use a virtual or physical ISE appliance, be sure to properly scale your deployment for the anticipated user base and any redundancy that will be required. Deployment scenarios and redundancy are addressed in greater detail in the next section.

It is also important to note that ISE supports VMware, Microsoft Hyper-V, and Linux KVM virtual environments. There are distinct restrictions and guidelines for using the advanced services of these virtual environments (such as VMotion), and it is recommended that you read the "Virtual Machine Requirements" section of the ISE administration guide for the version you are installing.

ISE Deployment Scenarios

The basic architecture for network access AAA does not really change, regardless of scale. You have an endpoint that is attempting to connect to a network through an access-layer network device called a ***network access device (NAD)***, which sends the authentication request to an ISE PSN over RADIUS.

Some things do change, however, including how many ISE nodes are deployed, which personas are running on each node, and where those nodes exist in the network design. ISE can be designed as a single-node (or standalone), two-node, or distributed deployment. As a candidate for the SISE 300-715 exam, you need to be intimately familiar with ISE design options, what each persona does, and what all the ISE services are.

Single-Node Deployment

A single-node, or standalone, deployment is the most basic Cisco ISE deployment type. In this type of deployment, the entire ISE infrastructure resides on a single appliance—either physical or virtual. As discussed earlier in this chapter, each persona has a specific role to serve in the Cisco ISE ecosystem; a single-node ISE deployment is responsible for performing all these personas on a single device (see Figure 6-1).

In the standalone deployment illustrated in Figure 6-1, a single ISE node is running the Administration, Monitoring, and Policy Services personas, so the NAD sends all access requests to the single ISE node. That single node is responsible for all the administrative functions, logging, profiling, pxGrid, certificate authority, hosting of all web portals, and authentication. The other optional services, such as pxGrid, device administration, and TC-NAC, may also be run on the single node, but if they are, scalability will be affected.

With a single-node deployment, there is no redundancy. Therefore, if the appliance fails or loses network connectivity for any reason, the ability to provide network authentication and authorization may not exist. The maximum number of supported endpoints in a single-node deployment is 10,000 for the SNS-3615 (or equivalent virtual appliance) or 50,000 for the SNS-3695 (or equivalent virtual appliance).

As there is no redundancy in this solution and the availability of network resources may be impacted in the event of an outage, this deployment is suggested only for testing in a lab environment.

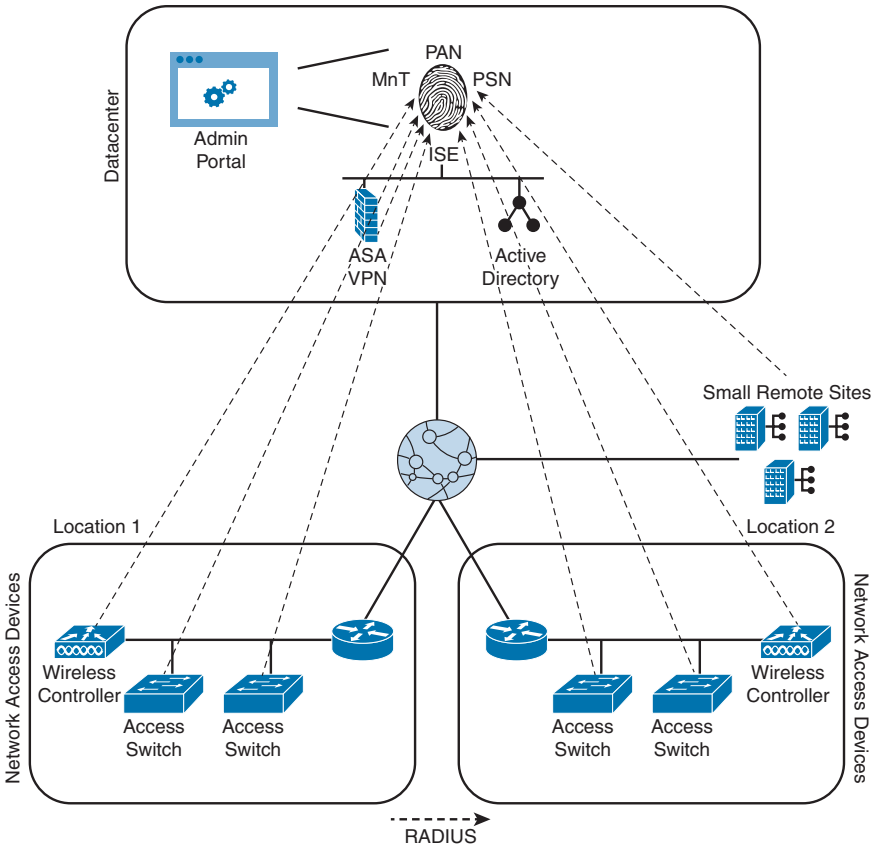


Figure 6-1 *Single-Node ISE Deployment with Persona Assignment*

Two-Node Deployment

A two-node Cisco ISE deployment incorporates redundancy, which is highly recommended in a production network. In this deployment type, the various ISE personas are distributed across two appliances. As with a single-node deployment, these appliances may be either physical or virtual appliances. A mixture of physical and virtual appliances may also be considered, depending on the network topology and application. Each of the personas must be assigned to one or more of the two available appliances. When redundancy is available in this manner, it is sometimes referred to as an *ISE distributed deployment*.

The Administration and Monitoring personas both have the concept of primary and secondary nodes. However, the Policy Services persona does not have primary and secondary designations; a PSN is always "active" as soon as the service is enabled on the appliance. It is up to the NAD to determine which PSN the RADIUS requests are sent to.

Figure 6-2 illustrates how both nodes in a two-node deployment run all services and are exact mirrors of one another. In this type of deployment, if one node goes down, the remaining node can still perform the authentications, and the end users attempting to access the network should not be affected.

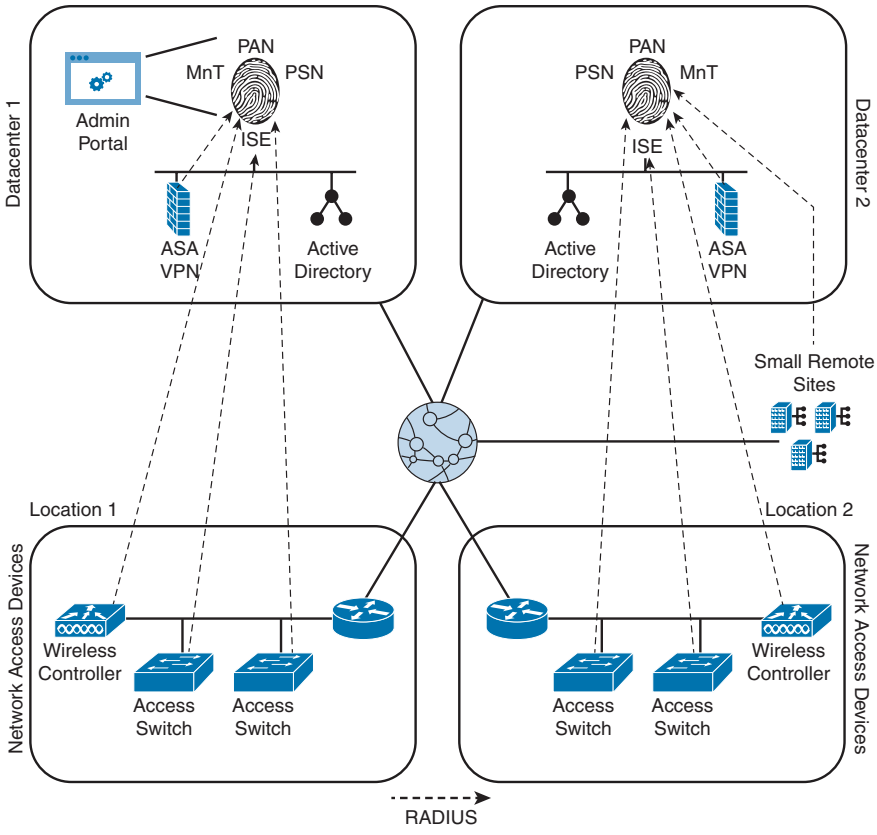


Figure 6-2 *Basic Two-Node ISE Deployment*

By balancing the primary and secondary roles across the two appliances, each persona is logically and physically distributed across the two appliances, providing the most resilient deployment using only two appliances. If all personas and nodes are working without issue, the loads of the ISE servers are, for the most part, equally distributed, with each of the nodes actively performing a role in the ISE ecosystem. In the event of the failure of a single persona, the secondary instance of the persona on the other node can take over that single role. If there is a failure of a single appliance, every persona's role can be provided by the other node.

Key Topic

If a failure of the primary PAN occurs, no synchronization between the PAN and PSNs will occur until the secondary PAN is promoted to primary. Once it becomes primary, synchronization can resume. This is sometimes referred to as a "warm spare."

Because a PSN is always active, each NAD on the network should point to each PSN. Each NAD contains a complete RADIUS configuration pointing to both ISE 1 and ISE 2. Depending on the order of the ISE PSN configuration on the NAD, the NAD will either choose to use ISE 1 first and ISE 2 as a secondary/backup or vice versa. To provide the best load balancing of the PSN persona across ISE 1 and ISE 2, you may consider listing ISE 1 as the primary PSN on 50% of the NADs on your network and ISE 2 as the primary PSN on the other 50% of your NADs—using the other ISE as the secondary PSN.

If there is a failure of the Policy Services persona (and, therefore, RADIUS processing) or a failure of a single appliance, the NAD detects the failure of the RADIUS service on one of

the PSNs and directs 100% of future RADIUS queries from that NAD to the remaining PSN. The availability, failover, and recovery detection criteria of the RADIUS service are configured on the NAD. If you want to ensure 100% redundancy of ISE services on a network in a two-node deployment in the event of a single logical or physical failure—causing a single ISE node to take on 100% of ISE functionality—the maximum number of supported endpoints would be 10,000 for the SNS-3615 or 50,000 for the SNS-3695. If there are more than the specified number of endpoints, the remaining, operational, ISE appliance might not be able to single-handedly manage the load.

Distributed Deployments

Key Topic

As a deployment grows beyond two nodes running the Policy Services persona, you might choose to have dedicated PSNs. You can have no more than five PSNs in an ISE cube where the Administration and Monitoring personas are running concurrently on a single node. This deployment model maintains a maximum concurrent session count of 50,000 endpoints, but it allows you to distribute the Policy Services functions for redundancy, reduce round-trip time (getting the PSN closer to the NADs), and dedicate PSNs to functions such as TC-NAC, pxGrid, and SXP.

Figure 6-3 shows an ISE deployment in which the Administration and Monitoring personas are running together on the same node, with two of those nodes for redundancy. In addition, it has five dedicated PSNs, one of which has been designated to run the pxGrid and TC-NAC functions. This is commonly referred to as a *hybrid deployment*.

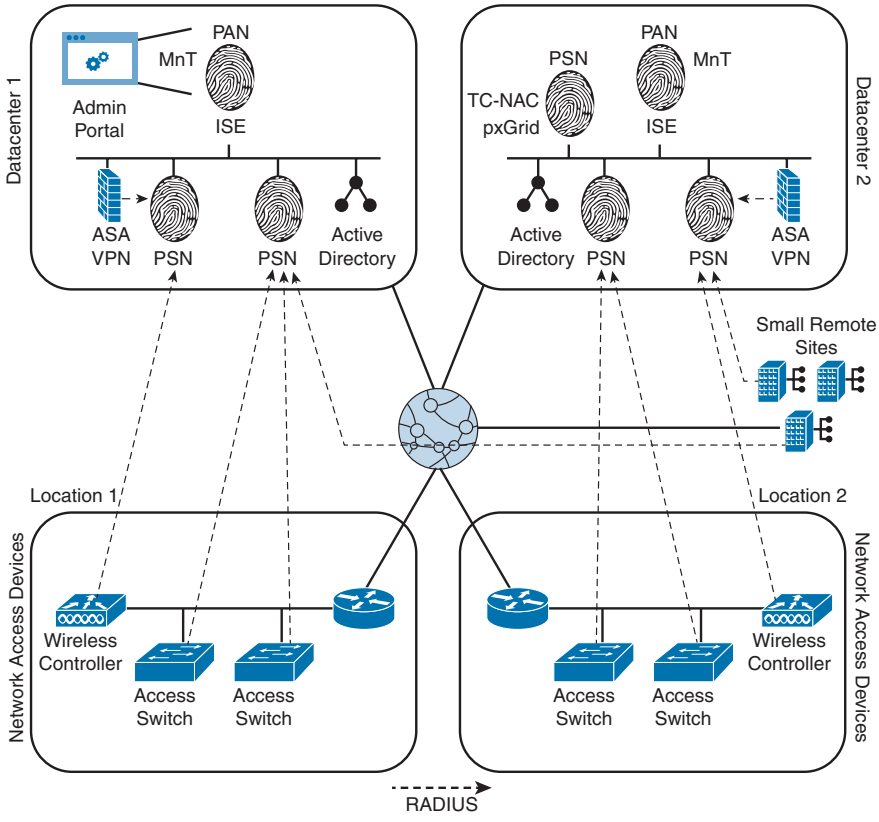


Figure 6-3 *Seven-Node Hybrid Deployment*

By separating the Policy Services persona from the Administration and Monitoring personas, RADIUS calls to a PSN will continue to function even if the Administration and Monitoring personas fail. Despite having all Policy Services functions dedicated to separate appliances, the number of supported endpoints remains the same as in a two-node deployment: 10,000 for the SNS-3415 and 50,000 for the SNS-3695.

To grow the scale beyond 50,000 concurrent active sessions or beyond 5 PSNs, the Administration and Monitoring personas must be running on dedicated nodes. Once all personas are divided up to dedicated nodes, the scale can reach 2,000,000 active concurrent sessions in ISE Version 2.7 and the SNS-3695 appliance, with a maximum of 50 PSNs plus up to 4 dedicated pxGrid PSNs.

Figure 6-4 illustrates a fully distributed deployment.

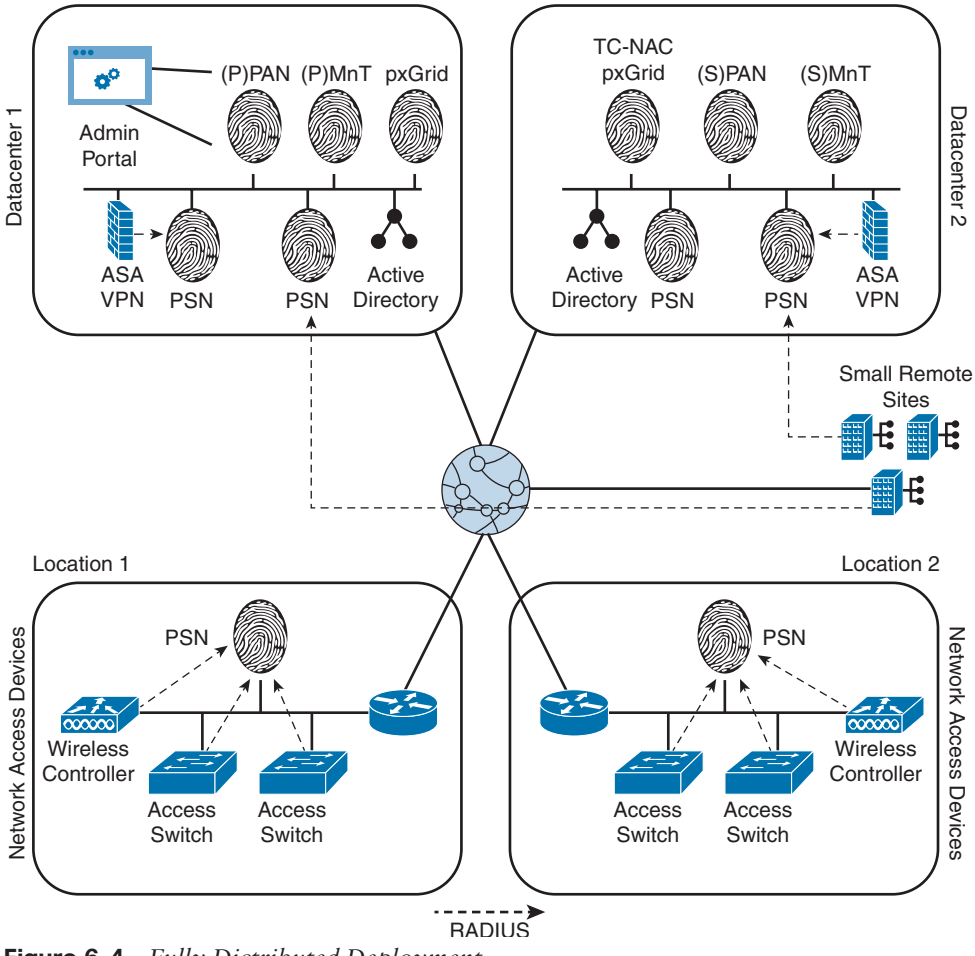


Figure 6-4 *Fully Distributed Deployment*

NOTE As of ISE Version 2.7, only one node is able to run the TC-NAC function, regardless of the deployment size.

Table 6-5 summarizes the ISE scale limits for Versions 2.2 through 2.7.

Key Topic

Table 6-5 ISE Deployment Scale and Limits

Characteristic	ISE 2.2 Maximums	ISE 2.4 Maximums	ISE 2.6 and 2.7 Maximums
Concurrent sessions in a fully distributed deployment (separate PAN, MnT, and PSN)	250,000: 3495 PAN 3495 MnT	34xx not supported	2,000,000: 3695 PAN 3695 MnT
	500,000: 3595 PAN 3595 MnT	500,000: 3595 PAN 3595 MnT	500,000: 3595 PAN 3595 MnT
Concurrent sessions in a hybrid deployment (PAN and MnT on a single node and dedicated PSNs)	5000: 3415 PAN+MnT	34xx not supported	10,000: 3615 PAN+MnT
	10,000: 3495 PAN+MnT	7500: 3515 PAN+MnT	25,000: 3655 PAN+MnT
	7500: 3515 PAN+MnT	20,000: 3595 PAN+MnT	50,000: 3695 PAN+MnT
			7500: 3515 PAN+MnT
	20,000: 3595 PAN+MnT	—	20,000: 3595 PAN+MnT
Concurrent sessions in a standalone deployment (PAN, MnT, and PSN all on a single node)	5000: 3415	34xx not supported	10,000: 3615
	10,000: 3495	7500: 3515	25,000: 3655
	7500: 3515	20,000: 3595	50,000: 3695
	20,000: 3595		7500: 3515
			20,000: 3595
PSNs in a fully distributed deployment (separate PAN, MnT, and PSN nodes)	40: 3495 PAN	50: 3595 PAN	50: 3595 PAN 50: 3695 PAN
PSNs in a hybrid deployment (PAN and MnT on a single node and dedicated PSNs)	5	5	5
pxGrid nodes in a fully distributed deployment (separate PAN, MnT, PXG, and PSN nodes)	2	4 (all active)	4 (all active)

Characteristic	ISE 2.2 Maximums	ISE 2.4 Maximums	ISE 2.6 and 2.7 Maximums
pxGrid nodes in a hybrid deployment (PAN and MnT on a single node and dedicated PSNs)	2 (either collocate PXG on PAN+MNT node or else dedicate PXG nodes and reduce PSN count by up to 2 nodes)	2 (either collocate PXG on PAN+MNT node or else dedicate PXG nodes and reduce PSN count by up to 2 nodes)	2 (either collocate PXG on PAN+MNT node or else dedicate PXG nodes and reduce PSN count by up to 2 nodes)
NADs	100,000	100,000	100,000
Network device groups	10,000	10,000	10,000
Active Directory forests (join points)	50	50	50
Active Directory controllers (WMI query)	100	100	100
Internal users	300,000	300,000	300,000
Internal guests	1,000,000 Expect latency for admin GUI + user auth >500,000 guests	1,000,000 Expect latency for admin GUI + user auth >500,000 guests	1,000,000 Expect latency for admin GUI + user auth >500,000 guests
User certificates	1,000,000	1,000,000	1,000,000
Server certificates	1000	1000	1000
Trusted certificates	1000	1000	1000
User portals (guest, BYOD, MDM, cert, posture)	600	600	600
Endpoints	1,500,000	1,500,000	2,000,000
Policy sets	100	200	200
Authentication rules	100 (simple policy mode) 200 (policy set mode)	N/A (simple policy mode) 1000 (policy set mode)	N/A (simple policy mode) 1000 (policy set mode)
Authorization rules	600 (simple policy mode) 700 (policy set mode)	N/A (simple policy mode) 3000 (policy set mode) with 3200 Authz profiles	N/A (simple policy mode) 3000 (policy set mode) with 3200 Authz profiles
User identity groups	1000	1000	1000
Endpoint identity groups	1000	1000	1000

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 6-6 lists these key topics and the page number on which each is found.

Key Topic

Table 6-6 Key Topics

Key Topic Element	Description	Page Number
Paragraph	Failover scenario with PANs	115
Paragraph	Distributed deployment types	116
Table 6-5	ISE deployment scale and limits	118

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

persona, Policy Administration node (PAN), Monitoring and Troubleshooting (MnT) node, Policy Services node (PSN), ISE cube, nodes, network access device (NAD), Change of Authorization (CoA), pxGrid

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What are the four ISE personas in ISE Version 2.7?
2. What is the maximum number of endpoints supported in a single ISE cube, with all personas running on SNS-3695 physical appliances?
3. What is commonly referred to as a hybrid deployment of ISE?
4. Which persona is responsible for receiving RADIUS Access-Requests, authenticating the entity, and performing the policy decisions for authorization?
5. Which persona hosts the web portals used for guest user authentication?

This page intentionally left blank



CHAPTER 7

A Guided Tour of the Cisco ISE Graphical User Interface (GUI)

This chapter covers the following topics:

Logging in to ISE

Organization of the ISE GUI

Types of policies in ISE

In Chapter 6, "Cisco Identity Services Engine Architecture," we discussed the roles of the ISE personas and how ISE is deployed within the network. After you have determined what deployment model to use in a network, you need to begin the configuration of the ISE security policy. To do this, you need to be familiar with the ISE administration portal.

As you saw in Chapter 6, ISE's administrative portal is a web interface that you use for ISE administration. To use the portal, you need to use a supported web browser. Early versions of ISE required browser support for Adobe Flash, but as of ISE Version 2.4, Flash has been completely removed, and HTML 5 is used instead. ISE Version 2.6 supports the following web browsers:

- **Mozilla Firefox**: Version 69 and earlier versions, ESR 60.9 and earlier versions
- **Google Chrome**: Version 77 and earlier versions
- **Microsoft Edge**: Beta Version 77 and earlier versions
- **Microsoft Internet Explorer**: Versions 10.x and 11.x

The minimum required screen resolution for all browsers is 1280×800 pixels.

NOTE If you are using Internet Explorer 10.x, you need to enable TLS 1.1 and TLS 1.2 and disable SSL 3.0 and TLS 1.0 under Internet Options > Advanced.

In this chapter, you will see how to log in to ISE, take a look at the organization of the ISE *GUI*, and learn about the various policies that are available in Cisco ISE.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 7-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 7-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Logging into ISE	1–2
Organization of the ISE GUI	3–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. To ensure the highest level of security, which of the following does the ISE administrative GUI use?
   - **a.** SSH
   - **b.** SCP
   - **c.** HTTP
   - **d.** HTTPS
2. Which of the following best characterizes the initial certificate presented by the ISE administrative GUI?
   - **a.** It is signed by a trusted public certificate authority.
   - **b.** It is a self-signed certificate that is automatically generated by ISE.
   - **c.** It is delivered in a separate envelope from the ISE appliance.
   - **d.** It is put in a frame and hung over your desk at work.
3. What do ISE operations allow an administrator to do?
   - **a.** Actively monitor, report, and troubleshoot active authentication and authorization sessions
   - **b.** Configure how ISE operates on the network
   - **c.** Create the web portals for client provisioning
   - **d.** Modify the security policy of ISE

4. Which of the following features cannot be configurated by an administrator in the Policy tab of the Cisco ISE GUI?
   a. Authorization
   b. Client provisioning
   c. Web portals
   d. Security group access
5. Which of the following items can you configure in the Administration component of Cisco ISE? (Choose two.)
   a. Policy elements
   b. Certificates
   c. Dictionaries
   d. Network devices
6. When adding a network access device (NAD) to Cisco ISE, which of the following details can be configured for the NAD? (Choose three.)
   a. MAC address
   b. IP address
   c. Device name
   d. RADIUS server IP address
   e. RADIUS shared secret key
   f. Mobile device manager
   g. SGA AAA servers
7. Profiling policy in ISE can leverage all except which of the following protocols to determine the type of endpoint that is accessing the network? (Choose two.)
   a. DHCP
   b. RADIUS (by proxy)
   c. SSH
   d. HTTP(S)
   e. FTP
8. Client provisioning is a process by which all necessary _____ are deployed to the endpoint, allowing the endpoint to more easily—and maybe even automatically—join the network in the future.
   a. credentials and configurations
   b. regulations and policies
   c. IP addresses and ACLs
   d. protocols and processes

9. Which of the following does pxGrid provide?
   a. Integration between ISE and MDM systems
   b. Location-based authorization by integrating Cisco's MSE
   c. Integration with third-party ISE ecosystem partners
   d. Integration with vulnerability and threat adapters
10. When an endpoint is quarantined with Adaptive Network Control (ANC), what happens to the endpoint?
   a. The endpoint is blocked from accessing the network, based on MAC address.
   b. The endpoint is given Internet-only access until it is unquarantined.
   c. Nothing happens without a corresponding security policy.
   d. The endpoint is blocked from accessing the network, based on IP address.

Foundation Topics

Logging in to ISE

The ISE Admin portal provides easy access to configuration, reporting, and troubleshooting elements. As a security policy management platform, it acts as a Swiss Army knife for network access control, with many menus, configuration options, and features. Due to its robust capabilities, the interface can be daunting to navigate for someone new to ISE. This chapter reviews how to navigate and find the most important elements.

Initial Login

To log in to Cisco ISE, you need to open a supported web browser. If you are logging in to a version of ISE that predates ISE Version 2.4, you need to ensure that the browser is running a supported version of Adobe Flash, as described in the ISE Release Notes for that version. If your web browser does not meet these requirements, your ability to access or configure the various components of Cisco ISE may be drastically affected.

NOTE To find the minimum browser recommendations for ISE, check the documentation for the ISE version at https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html.

Once you have opened a supported web browser, you need to access Cisco ISE. In order to maintain the highest level of security, the Cisco ISE GUI leverages HTTPS for administrative access. Therefore, the URL to access Cisco ISE is https://<*ISE_FQDN*>. If you have not added ISE to your DNS server, either by hostname or FQDN, you may need to type the ISE's IP address explicitly rather than use *ISE_FQDN* in order to access the administrative GUI.

Key Topic

When you first access Cisco ISE, you might receive a security alert from your browser. The default generated X.509 certificate that is used to identify the HTTPS ISE administrative portal is a self-signed certificate. After this initial login, you can choose to permanently accept the certificate in your web browser or load a new X.509 certificate that is pre-trusted by your web browser. (The process by which the certificate is permanently accepted by the web browser is browser specific and, therefore, beyond the scope of this chapter.)

NOTE The use of X.509 certificates with ISE is discussed further in Chapter 8, "Initial Configuration of Cisco ISE."

As you first browse to the provided URL, the login screen shown in Figure 7-1 appears. This login screen helps authenticate you as a valid administrative user. The username and password used here define what level of access you receive as you access the administrative GUI. The initial administrative user (configured during the initial bootstrapping of Cisco ISE) receives complete access to all components of the administrative portal. (This is the "super admin" role; see Table 7-2.) The administrative username is admin.

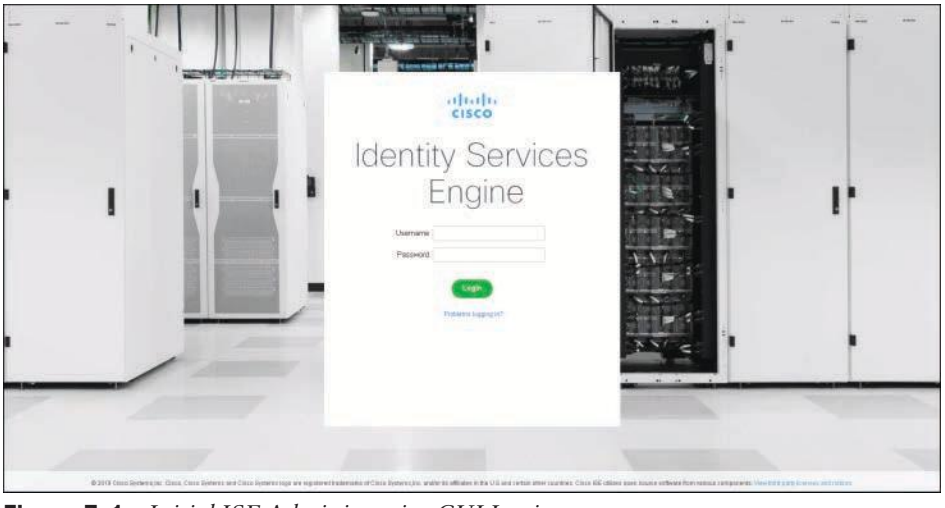


Figure 7-1 *Initial ISE Administrative GUI Login*

You can create additional administrative users after the initial login. Each additional administrative user can be assigned a specific administrative role. As mentioned previously, the username and password—and, therefore, the resulting administrative role—provided at login determine your access to various functions in the administrative portal. The built-in administrative roles are defined in Table 7-2.

Table 7-2 Cisco ISE Admin Groups, Access Levels, Permissions, and Restrictions

Admin Group Role	Access Level	Permissions	Restrictions
Customization admin	Manage sponsor, guest, and personal devices portals	Configure guest and sponsor access Manage guest access settings Customize end-user web portals	Cannot perform any policy management or identity management or system-level configuration tasks in Cisco ISE Cannot view any reports
Helpdesk admin	Query monitoring and troubleshooting operations	Run all reports Run all troubleshooting flows View the Cisco ISE dashboard and Live Log View alarms	Cannot create, update, or delete reports, troubleshooting flows, live authentications, or alarms
Identity admin	Manage user accounts and endpoints Manage identity sources	Add, edit, and delete user accounts and endpoints Add, edit, and delete identity sources Add, edit, and delete identity source sequences Configure general settings for user accounts (attributes and password policy) View the Cisco ISE dashboard, Live Log, alarms, and reports Run all troubleshooting flows	Cannot perform any policy management or system-level configuration tasks in Cisco ISE
MnT admin	Perform all monitoring and troubleshooting operations	Manage (run, create, and delete) all reports Run all troubleshooting flows View the Cisco ISE dashboard and Live Log Manage (create, update, view, and delete) alarms	Cannot perform any policy management or identity management or system-level configuration tasks in Cisco ISE

Admin Group Role	Access Level	Permissions	Restrictions
Network device admin	Manage Cisco ISE network devices and the network device repository	Read and write permissions on network devices Read and write permissions on NDGs and all network resources object types View the Cisco ISE dashboard, Live Log, alarms, and reports Run all troubleshooting flows	Cannot perform any policy management or identity management or system-level configuration tasks in Cisco ISE
Policy admin	Create and manage policies for all Cisco ISE services across the network that are related to authentication, authorization, posture, profiler, and client provisioning	Read and write permissions on all the elements used in policies, such as authorization profiles, NDGs, and conditions Read and write permissions on identities, endpoints, and identity groups (user identity groups and endpoint identity groups) Read and write permissions on services, policies, and settings View the Cisco ISE dashboard, Live Log, alarms, and reports Run all troubleshooting flows Access to Device Administration work center Permission for TACACS policy conditions and results Network device permissions for TACACS proxy and proxy sequences	Cannot perform any identity management or system-level configuration tasks in Cisco ISE No guarantee of access to the subordinate links in the Access to the Device Administration work center
RBAC admin	All tasks under the Operations menu except for Endpoint Protection Services Adaptive Network Control and partial access to some menu items under Administration	View the authentication details Enable or disable endpoint protection services Adaptive Network Control Create, edit, and delete alarms; generate and view reports; and use Cisco ISE to troubleshoot problems in the network Read permissions on administrator account settings and admin group settings	Cannot perform any identity management or system-level configuration tasks in Cisco ISE

Admin Group Role	Access Level	Permissions	Restrictions
		View permissions on admin access and data access permissions, along with the RBAC policy page View the Cisco ISE dashboard, Live Log, alarms, and reports Run all troubleshooting flows	
Read-only admin	Read-only access to the ISE GUI	View and use the functions of the dashboard, reports, and Live Log sessions, such as filtering data, querying, saving options, printing, and exporting data Change passwords of their own accounts Query ISE using global search, reports, and Live Log sessions Filter and save data based on the attributes Export data pertaining to authentication policies, profile policies, users, endpoints, network devices, network device groups, identities (including groups), and other configurations Customize, save, print, and export reports Generate custom report queries, save, print, or export the results Save UI settings for future reference Download logs, such as ise-psc-log, from the Operations > Troubleshoot > Download Logs page	Perform any configuration changes, such as create, update, delete, import, quarantine, and MDM actions of any objects, such as authorization policies, authentication policies, posture policies, profiler policies, endpoints, and users Perform system operations, such as backup/restore, registration/deregistration of nodes, synchronization of nodes, creation/editing/deletion of node groups, or upgrade/installation of patches Import data pertaining to policies, network devices, network device groups, identities (including groups), and other configurations Perform operations such as CoA, endpoint debugging, modification of collection filters, bypassing of suppression on live sessions data, modification of the PAN-HA failover settings, and editing of the personas or services of any Cisco ISE node Run commands and tools that might negatively impact performance (for example, TCP Dump) Generate support bundles

Admin Group Role	Access Level	Permissions	Restrictions
Super admin	All Cisco ISE administrative functions (The default administrator account belongs to this group.)	Create, read, update, delete, and execute (CRUDX) permissions on all Cisco ISE resources Note that the super admin user cannot modify the default system-generated RBAC policies and permissions. To do this, you must create new RBAC policies with the necessary permissions based on your needs and map these policies to any admin group. Access to Device Administration Work Centers Permission for TACACS policy conditions and results Network device permissions for TACACS proxy and proxy sequences and permission to enable TACACS global protocol settings	No guarantee of access to subordinate links in the Device Administration work center. Only the admin user from the default super admin group can modify or delete other admin users. (Even an externally mapped user who is part of an Admin group cloned with the menu and data access privileges of the super admin group cannot modify or delete an admin user.)
System admin	All Cisco ISE configuration and maintenance tasks	Full access (read and write permissions) to perform all activities under the Operations tab and partial access to some menu items under the Administration tab Read permissions on administrator account settings and administrator group settings Read permissions on admin access and data access permissions along with the RBAC policy page Read and write permissions for all options under the Administration > System menu View the authentication details Enable or disable endpoint protection services Adaptive Network Control Create, edit, and delete alarms; generate and view reports; and use Cisco ISE to troubleshoot problems in the network Permission to enable TACACS global protocol settings	Cannot perform any policy management or system-level configuration tasks in Cisco ISE

Admin Group Role	Access Level	Permissions	Restrictions
Elevated system admin (available in Cisco ISE Version 2.6, patch 2 and above)	All Cisco ISE configuration and maintenance tasks	In addition to all of the privileges of the system admin, an elevated system admin can create admin users	Cannot create or delete super admin user Cannot manage the super admin groups
External RESTful Services (ERS) admin	Full access to all ERS API requests, such as GET, POST, DELETE, and PUT	Create, read, update, and delete ERS API requests	Meant only for ERS authorization supporting internal users, identity groups, endpoints, endpoint groups, and SGTs
External RESTful Services (ERS) operator	Read-only access to ERS API (only GET)	Can only read ERS API requests	Meant only for ERS authorization supporting internal users, identity groups, endpoints, endpoint groups, and SGTs
TACACS+ admin	Full access	Access to the Device Administration work center Enable TACACS+ services Access to the external identity stores Access to Operations > TACACS Live Logs page	

To access these administrative roles in Cisco ISE, you must browse to Administration > System > Admin Access > Administrators > Admin Groups (see Figure 7-2). To add an administrative user with the given role, go to Administration > System > Admin Access > Administrators > Admin Users (see Figure 7-3).

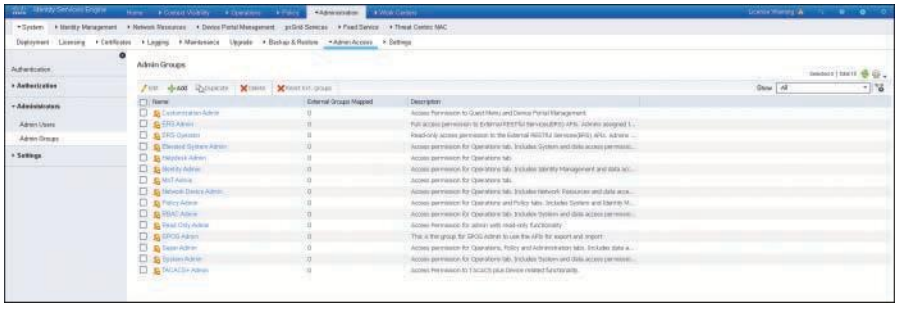


Figure 7-2 *Additional Administrative Roles*

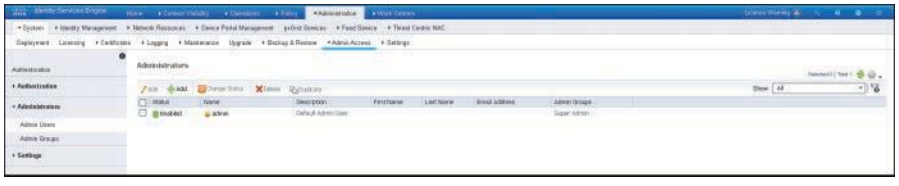


Figure 7-3 *Adding Administrative Users*

ISE Home Dashboards

Upon logging in to ISE, you are greeted with the ISE Home dashboard, as shown in Figure 7-4. This dashboard is analogous to the dashboard of your car in that it provides a condensed at-a-glance overview of the ISE deployment with real-time data. This dashboard can be an administrator's first stop when checking the status of an entire ISE deployment. The ISE dashboard can only be viewed from the primary Policy Administration node (PAN).

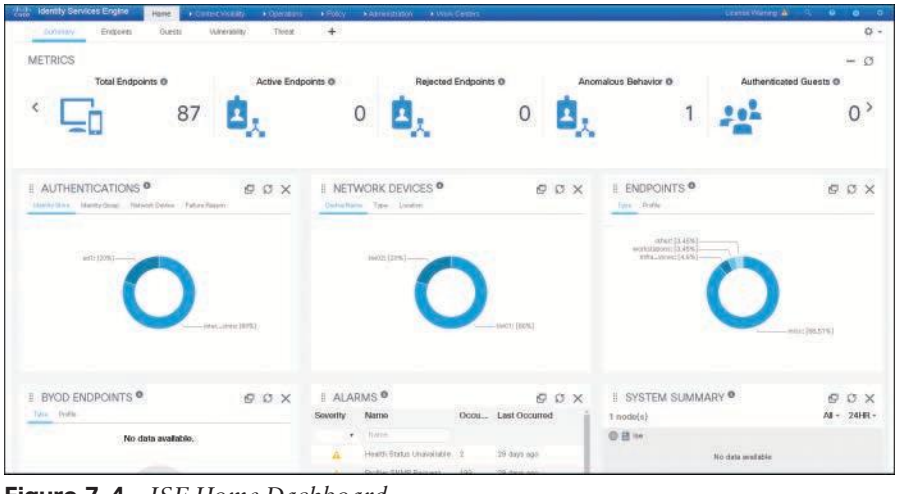


Figure 7-4 *ISE Home Dashboard*

The ISE Home page has five default tabs:

- **Summary:** This tab provides high-level information about the endpoints and overall deployment health.
- **Endpoints:** This tab provides the high-level status of endpoints, endpoint categories, and network devices.
- **Guests:** This tab provides information about the guest user types, logon failures, and locations.
- **Vulnerability:** If ISE is integrated with a vulnerability adapter via Threat Centric NAC, this tab provides information reported to ISE by that adapter.
- **Threat:** If ISE is integrated with a threat adapter via Threat Centric NAC, this tab provides information reported to ISE by that adapter.

You can create additional customized dashboards by clicking on the plus (+) sign in the dashboard bar to open the window shown in Figure 7-5 or by clicking on the gear icon in the dashboard to open the Dashboard Settings menu, as shown in Figure 7-6. After you give a new dashboard a name, a popup window provides a list of optional dashlets you can add to the new dashboard.

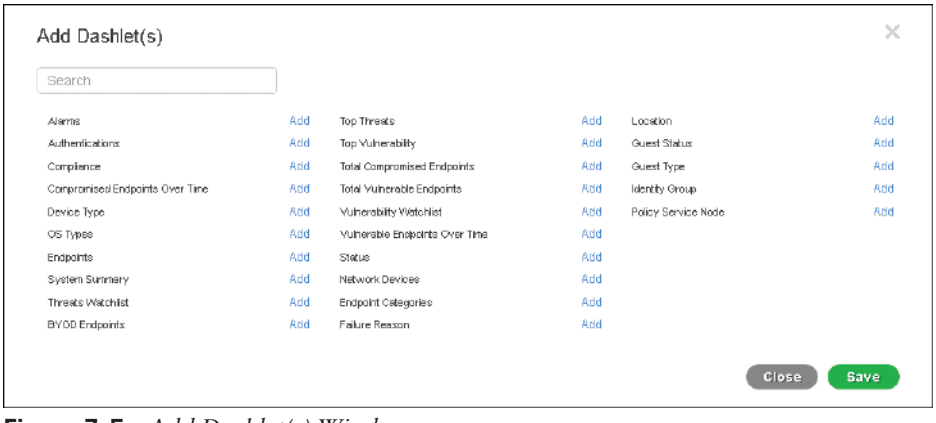


Figure 7-5 *Add Dashlet(s) Window*

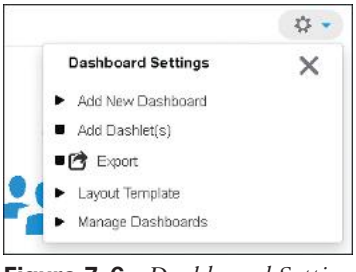


Figure 7-6 *Dashboard Settings Window*

In the Dashboard Settings menu, the following options are available:

- **Add New Dashboard:** Use this option to create a new dashboard. You can add a maximum of 15 dashboards, in addition to the 5 default dashboards.
- **Add Dashlet(s):** Use this option to add a dashlet to the existing dashboard.
- **Export:** Use this option to export the dashlet data either in PDF or CSV format.
- **Layout Template:** Use this option to change the layout of the dashlets as they are displayed in the dashboard.
- **Manage Dashboards:** Use this option to mark a dashboard as a default dashboard (home page) or reset all the dashboards.

The Summary tab (refer to Figure 7-4) displays the following dashlets:

- **Metrics:** This dashlet shows a high-level summary of the total number of endpoints that have been seen by ISE, how many endpoints are active, how many endpoints have been rejected, how many endpoints have been detected exhibiting anomalous behavior, the number of authenticated guests, the number of BYOD endpoints, and how many endpoints have been postured as compliant.
- **Authentications:** This dashlet provides a summary of passed and failed authentications, filtered by identity store, identity group, network device, and failure reason (if applicable).
- **Network Devices:** This dashlet provides information on endpoints authenticated by network devices and can be filtered by either device name, device type, or device location.
- **Endpoints:** If profiling is enabled, this dashlet provides an overview of endpoints profiled by network function and can be filtered by endpoint type or endpoint profile.
- **BYOD Endpoints:** This dashlet provides profiling information for BYOD devices and can be filtered by endpoint type or endpoint profile.
- **Alarms:** This dashlet provides any alarms or anomalous behaviors that have been seen by ISE and is refreshed automatically every 30 seconds. A few examples of alarms are authentication inactivity, NTP sync issues, and insufficient virtual machine failures.
- **System Summary:** This dashlet provides a quick status overview of all ISE appliances in the deployment and is refreshed every 60 seconds. This summary includes the health status, the CPU level, the memory usage, and the authentication latency for each ISE appliance.

Figure 7-7 shows the Endpoints tab, which includes the following dashlets:

- **Status:** This dashlet provides information about the network connection status of endpoints, based on session information. A pie chart shows connected and disconnected endpoints.

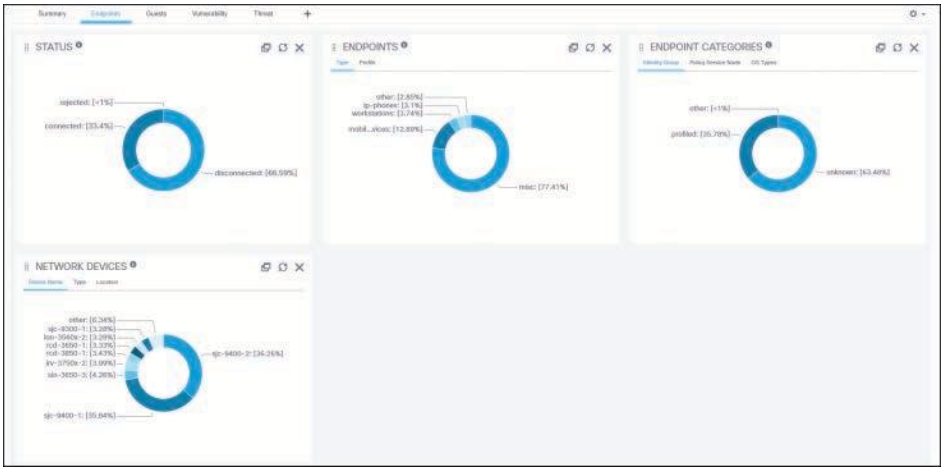


Figure 7-7 *Endpoints Tab*

- **Endpoints:** If profiling is enabled, this dashlet provides an overview of endpoints profiled by network function and can be filtered by endpoint type or endpoint profile.
- **Endpoint Categories:** If profiling is enabled, this dashlet provides information about profiled endpoints and can be filtered by identity group, Policy Services node (PSN), and operating system type.
- **Network Devices:** This dashlet provides information on endpoints authenticated by network devices and can be filtered by either device name, device type, or device location.

Figure 7-8 shows the Guests tab, which includes the following dashlets:

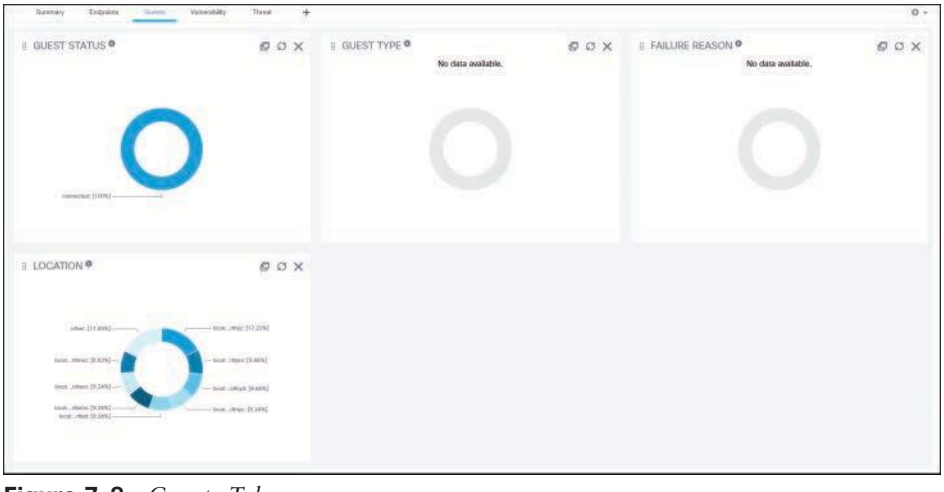


Figure 7-8 *Guests Tab*

- **Guest Status:** This dashlet provides information about the network connection status of guest endpoints, based on session information. A pie chart shows connected and disconnected endpoints.
- **Guest Type:** This dashlet provides a summary of guest authentication, based on guest type.
- **Failure Reason:** This dashlet provides failure reason information for guest endpoints (if applicable).
- **Location:** This dashlet provides information about guest connections, by network device location.

Figure 7-9 shows the Vulnerability tab, which includes the following dashlets:

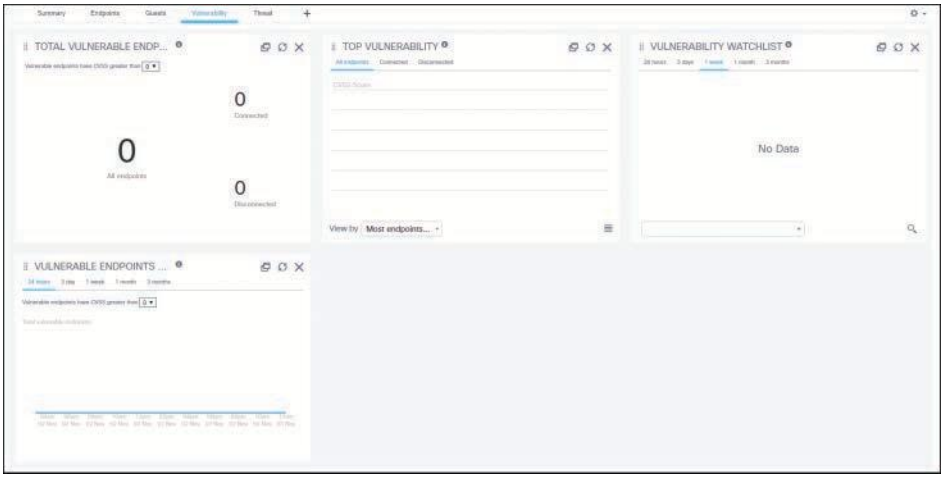


Figure 7-9 *Vulnerability Tab*

- **Total Vulnerable Endpoints:** This dashlet displays the total number of endpoints that have a Common Vulnerability Scoring System (CVSS) score greater than the specified value. You can choose the value from a dropdown in the dashlet to filter this information. The dashlet also displays the number of disconnected and connected endpoints, based on that CVSS score.
- **Top Vulnerability:** This dashlet provides the top vulnerabilities, based on the number of endpoints impacted or the severity of the vulnerability.
- **Vulnerability Watchlist:** This dashlet helps you analyze the trend of selected vulnerabilities over time. It can be viewed by 24 hours, 3 days, 1 week, 1 month, and 3 months.
- **Vulnerable Endpoints Over Time:** This dashlet displays a historical view of the impact on endpoints over time. It can be viewed by 24 hours, 3 days, 1 week, 1 month, and 3 months.

Figure 7-10 shows the Threat tab, which includes the following dashlets:

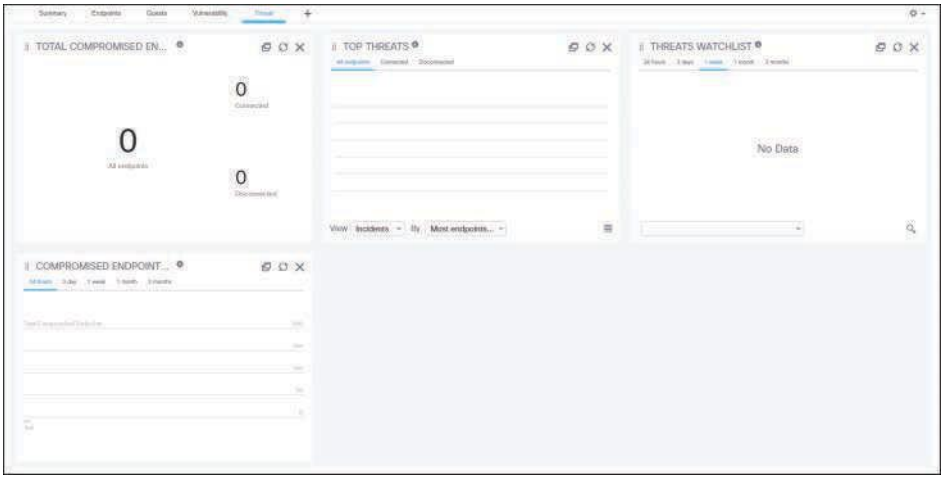


Figure 7-10 *Threat Tab*

- **Total Compromised Endpoints:** This dashlet displays the total number of endpoints, both connected and disconnected, that are currently impacted.
- **Top Threats:** This dashlet displays the top threats, based on the number of endpoints that have been impacted and the severity of the detected threat.
- **Threat Watchlist:** This dashlet helps you analyze the trend of selected events over time. It can be viewed by 24 hours, 3 days, 1 week, 1 month, and 3 months.
- **Compromised Endpoints Over Time:** This dashlet displays a historical view of the impact on endpoints over a specific period. It can be viewed by 24 hours, 3 days, 1 week, 1 month, and 3 months.

Administration Portal

In addition to the ISE Home dashboard, the ISE Admin portal provides access to a number of configuration and reporting tools, as shown in Figure 7-11.

Along the top of the screen, you can see several interactive elements. As indicated by the number 1 in Figure 7-11, a number of tabs are available:

- **Context Visibility:** This tab displays information about endpoints, users, and NADs. The information can be segmented by features, applications, BYOD, and other categories, depending on your license. The Context Visibility menus use a central database and gather information from database tables, caches, and buffers, which makes updates to context dashlets and list content very fast. The Context Visibility menus consist of dashlets at the top and a list of information at the bottom. As you filter data by modifying the column attributes in the list, the dashlets are refreshed to show the changed content.

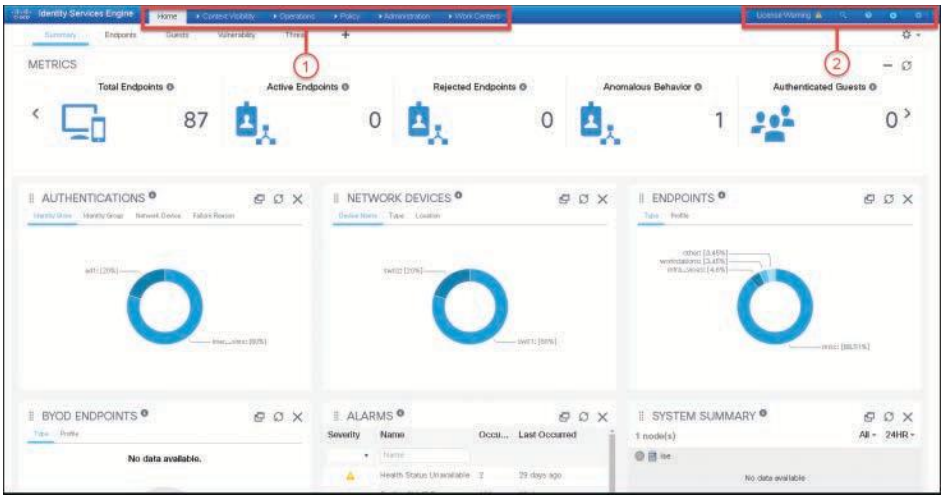


Figure 7-11 *Administration Portal*

- **Policy:** This tab provides access to tools for managing network security in the areas of authentication, authorization, profiling, posture, and client provisioning.
- **Administration:** This tab provides access to tools for managing Cisco ISE nodes, licenses, certificates, network devices, users, endpoints, and guest services.

As indicated by the number 2 in Figure 7-11, a number of options are available in the top-right menu in the ISE Admin portal:

- **Search:** The Search icon allows you to search for endpoints and displays their distribution by profiles, failures, identity stores, location, device type, and so on.
- **Help:** You can click the question mark icon to access online help for the currently displayed page and find links to the ISE community, ISE Portal Builder, and more.
- **PassiveID Setup:** The PassiveID Setup option launches the PassiveID Setup wizard, which helps you set up passive identities using Active Directory. You can configure the server to gather user identities and IP addresses from external authentication servers and deliver the authenticated IP addresses to the corresponding subscriber.
- **Visibility Setup:** The Visibility Setup option is a proof of value (PoV) service that collects endpoint data—such as applications, hardware inventory, USB status, firewall status, and the overall compliance state of Windows endpoints—and sends it to Cisco ISE. When you launch the ISE Visibility Setup Wizard, you can specify an IP address range to run endpoint discovery for a preferred segment of the network or a group of endpoints.

  The PoV service uses the Cisco Stealth temporal agent to collect endpoint posture data. Cisco ISE pushes the Cisco Stealth temporal agent to computers running Windows with an administrator account type and automatically runs a temporary executable file to collect context; then the agent removes itself. To experience the optional debugging capabilities of the Cisco Stealth temporal agent, go to Visibility

Setup > Posture and check the Endpoint Logging checkbox to save the debug logs in an endpoint or in multiple endpoints. You can view the logs in either of the following locations:

C:\WINDOWS\syswow64\config\systemprofile\ (for 64-bit operating system)

C:\WINDOWS\system32\config\systemprofile\ (for 32-bit operating system)

- **Wireless Setup (BETA):** The Wireless Setup (BETA) option provides an easy way to set up wireless flows for 802.1X, guest, and BYOD. This option also provides workflows to configure and customize each portal for guests and BYOD.
- **System activities:** You can bring up the online help and configure account settings.

Global Search for Endpoints

The Search icon at the top of the ISE Admin portal allows you to perform a global search for endpoints, based on a variety of criteria, as shown in Figure 7-12.

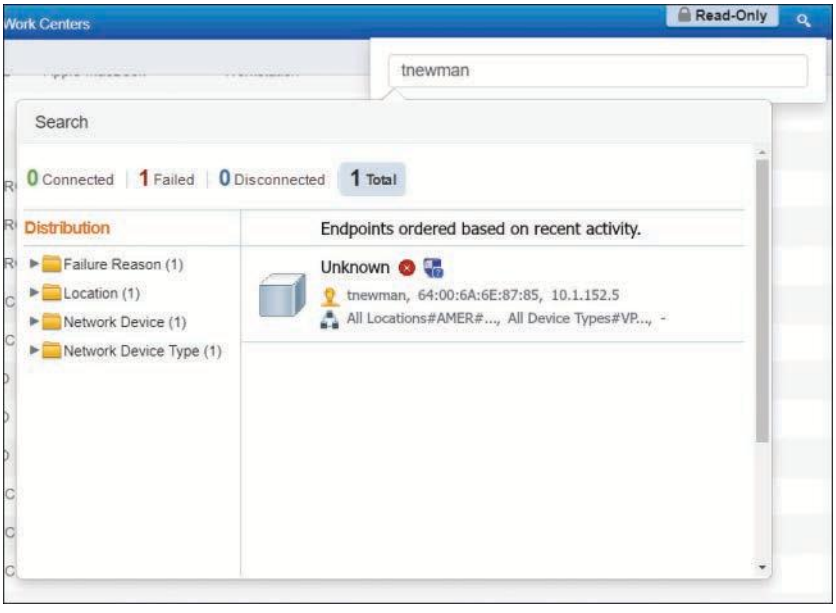


Figure 7-12 *Search Results for Endpoints*

In order to conduct a search, you need to enter at least three characters in the Search field. The results provide at-a-glance information that you can drill into for troubleshooting. The following criteria can be used for a global endpoint search:

- Username
- User type
- MAC address
- IP address

- Authorization profile
- Endpoint profile
- Failure reason
- Identity group
- Identity store
- Network device name
- Network device type
- Operating system
- Posture status
- Location
- Security group

Help Menu

In the top-right corner of the Admin portal, the Help icon provides a number of useful resources to help you manage Cisco ISE, as illustrated in Figure 7-13:

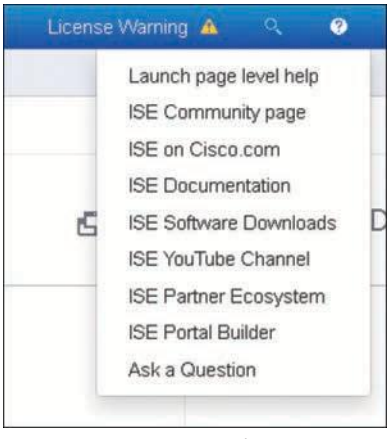


Figure 7-13 *ISE Help Menu*

- **Launch Page Level Help:** This option provides page-level help. When you choose this option, ISE immediately launches the associated help documentation for the current page in ISE. For example, if you are currently navigated to Context Visibility in ISE and want more information on this page, you can choose the Launch Page Level Help option to get more information about this page.

- **ISE Community Page:** This option takes you to the Cisco ISE Community, which is an online repository of ISE resources, including product documentation, data sheets, how-to guides, tech notes, and much more. It is also possible to post questions to the ISE Community to receive answers from other community members or Cisco ISE business unit resources.
- **ISE on Cisco.com:** This option takes you to an overview product page that provides a high-level marketing view of ISE.
- **ISE Documentation:** This option takes you to the Cisco Support page for ISE, where all the official documentation is posted.
- **ISE Software Downloads:** This option takes you to the Cisco page to download ISE patches, upgrade bundles, and images.
- **ISE YouTube Channel:** This option takes you to the ISE YouTube channel, which provides educational videos about features and walkthroughs of configuring ISE.
- **ISE Partner Ecosystem:** This option takes you to Cisco's Security Technical Alliance Partners page, which provides information about integrations between ISE and other Cisco and third-party products.
- **ISE Portal Builder:** This option takes you to Portal Builder, which is a Cisco-provided tool you can use to easily create customized user-facing guest, BYOD, MDM, and client provisioning portals. This is an excellent resource for branding your ISE portals without having web development experience. Creating a portal is as easy as uploading logos and arranging them in ISE Portal Builder. After completing your ISE portal design, you download the file and upload it to your ISE deployment.
- **Ask a Question:** This option opens a page in the ISE Community to start a new discussion where you can ask questions and receive replies from ISE subject matter experts.

ISE Setup Wizards

To streamline the configuration of certain functions, ISE provides several built-in wizards you can access at the top of the Admin portal to quickly deploy profiling, basic posturing of endpoints, and setup of guest and corporate wireless access control. Figure 7-14 shows the dropdown for the ISE Setup Wizard menu.



Figure 7-14 *ISE Setup Wizard Menu*

Settings Menu

The Settings menu at the top of Admin portal provides information on high-level system settings and help, as illustrated in Figure 7-15.

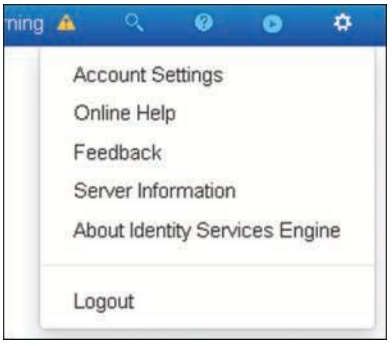


Figure 7-15 *Settings Menu*

Organization of the ISE GUI

The Cisco ISE GUI is divided into five functional components in addition to the Home screen:

- **Context Visibility:** This page, like the Home screen, provides several prebuilt dashboards you can use to get granular information about endpoints, users, network devices, and applications. You can filter the information in these dashboards to search for specific information.
- **Operations:** Operations are ISE components that allow you to actively monitor, report, and troubleshoot ongoing authentication and authorization sessions. This screen is a place where you can monitor, report, and troubleshoot on network devices and policies that are already configured on ISE.
- **Policy:** The Policy page enables you to configure the security policy with functions such as authentication, authorization, profiling, posture, client provisioning, and security group access policy, as well as the supplementary building blocks that are used in these policies. As a network device authenticates and authorizes to ISE, ISE processes the credentials provided by the NAD through this policy and provides the resulting authorization security policy back to the NAD.
- **Administration:** The Administration screen focuses on the configuration of ISE: what, who, and how users and devices can access ISE. This screen allows you to define how the ISE deployment behaves, what external identity resources are to be used, what devices can use the ISE security policy, what services ISE provides to the user base, and how often ISE updates its device databases.

- **Work Centers:** Work Centers is the last major screen of the Cisco ISE GUI. Work centers enable you to configure common use cases in ISE. This screen provides everything needed to configure, monitor, and troubleshoot use cases in ISE.

The following sections dive a little deeper into each of these major functional components, describing how each smaller part contributes to the complete ISE ecosystem.

Context Visibility

The Context Visibility screen provides prebuilt contextual dashboards with data on endpoints in ISE. From these dashboards, you can create custom filters to customize views further and even export the data. As illustrated in Figure 7-16, this screen is divided into four main tabs:

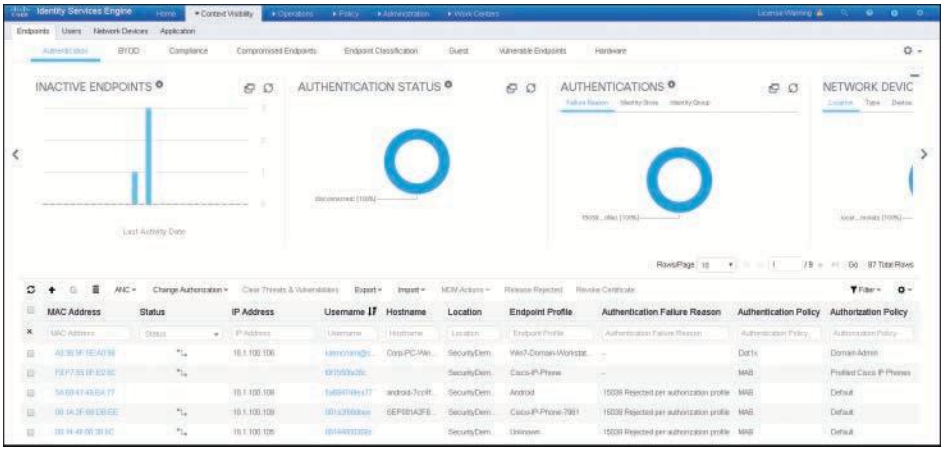


Figure 7-16 *Context Visibility Dashboard*

- **Endpoints:** This tab allows you to filter endpoints based on the types of devices, compliance status, authentication type, hardware inventory, and much more. From here, you can view detailed information as well as statically assign the endpoint to an endpoint policy (profile) or identity group.
- **Users:** This tab displays information by user identity source and type.
- **Network Devices:** This tab displays the network devices and the endpoints connected to them. You can use this tab to filter the connected endpoints by network device.
- **Application:** This tab lists applications discovered via posturing and the number of devices that have those applications installed.

Operations

The Operations screen has six main tabs: RADIUS, Threat-Centric NAC Live Logs, TACACS, Troubleshoot, Adaptive Network Control, and Reports. The following sections describe these tabs.

RADIUS

Key Topic

The Live Logs tab under RADIUS displays real-time information about recent RADIUS authentication. This screen can be configured to update in a periodic fashion to provide a live view of the authentications that are occurring and the security policies that are actively being issued to the endpoints. When testing new policies or troubleshooting authentication issues, this tab is an administrator's first stop in order to ascertain whether an endpoint has successfully authenticated. From this screen, an administrator can view further troubleshooting details about successful or rejected authentications.

In the Live Logs tab (see Figure 7-17), you can see users (Identity column), devices (Endpoints column), and NADs (Network Device column) as they attempt to authenticate to ISE. If a policy is successful, the resulting security policy (Authorization Profile column) is sent to the NAD. Each authentication event is logged on this tab.

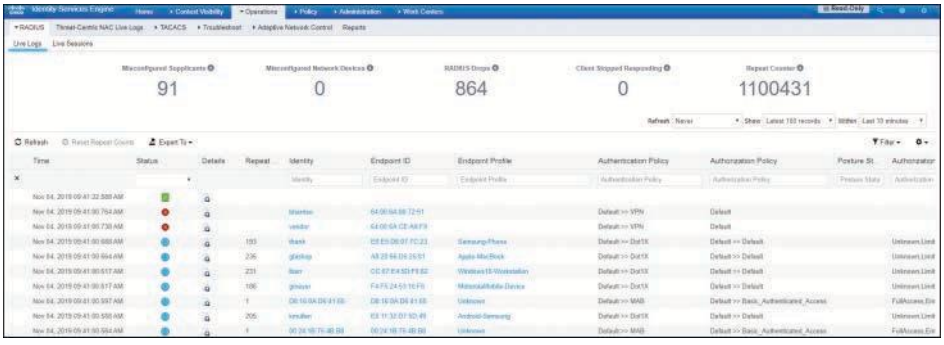


Figure 7-17 *Live Logs*

The Live Logs tab provides only a summary of the authentication—not all of the details. However, the details of the authentication session are sometimes as important as (or more important than) the final result. For instance, if an authentication fails directly or fails indirectly (because it hit the wrong policy), those authentication details can help an administrator determine the nature of the problem. To see the details of an authentication session, click the magnifying glass icon in the Details column for the relevant authentication event.

If you click an icon in the Details column, you see a large amount of information about the authentication session, including details of the RADIUS exchange that resulted in the authentication event, details about the endpoint, details about the NAD, and details about the identity store (see Figure 7-18). By reviewing portions of the Live Log details and knowing the expected policy outcome, you should be able to determine the reason that a particular authentication, authorization, posturing, or profiling policy was chosen (correctly or incorrectly) as the result of the authentication session.



Figure 7-18 *Live Log Details*

The Live Sessions tab under RADIUS shows the status of active sessions with ISE, as illustrated in Figure 7-19. These sessions might be sessions that have just started or sessions that are ongoing.

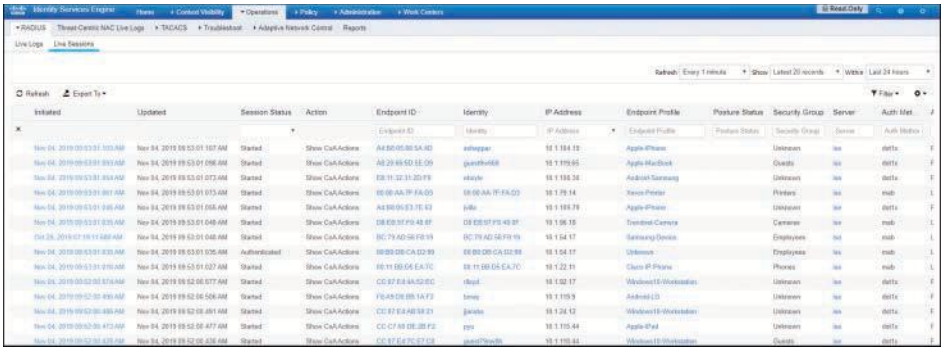


Figure 7-19 *Live Sessions*

By selecting a session or filtering for the relevant session that you are interested in, you can see a play-by-play of the authentication session, as shown in Figure 7-20. If a session is having issues, this view should provide a sufficient timeline for when each step of the authentication took place as well as which steps succeeded or failed.

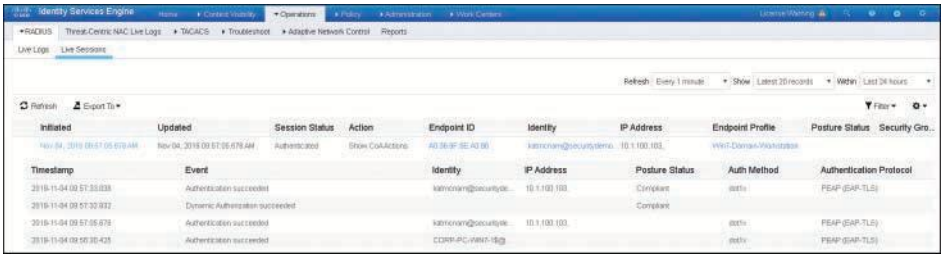


Figure 7-20 *Live Sessions Play-by-Play*

The Live Sessions tab enables you to force the termination of a session or force a session to reauthenticate. To take advantage of this ability, you select the CoA Action icon for the relevant session, as shown in Figure 7-21.

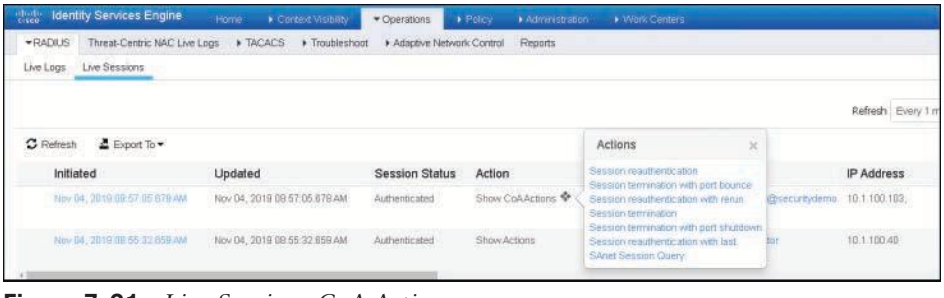


Figure 7-21 *Live Sessions CoA Actions*

Threat-Centric NAC Live Logs

The Threat-Centric NAC Live Logs tab lists threat and vulnerability events provided by the threat and vulnerability adapter (see Figure 7-22).

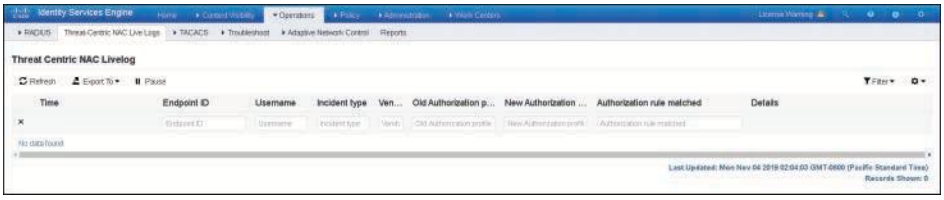


Figure 7-22 *Threat-Centric NAC Live Logs Tab*

TACACS Live Log

The TACACS Live Log provides real-time information on TACACS authentications and is similar to the RADIUS Live Logs screen. You can click on the details of a TACACS authentication to see more information about a single TACACS event, as shown in Figure 7-23.

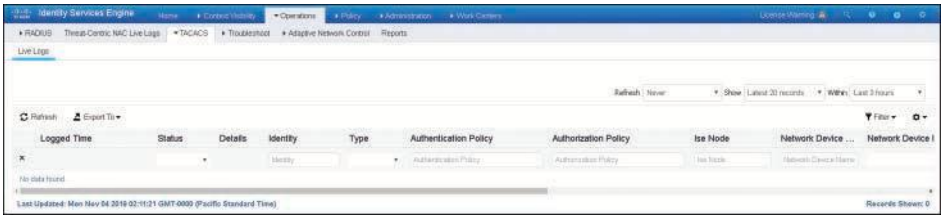


Figure 7-23 *TACACS Live Logs*

Troubleshoot

The Troubleshoot tab under Operations provides a number of tools a network administrator can use (see Figure 7-24). These tools provide several useful functions, including a tool to search for RADIUS authentications (RADIUS Authentication Troubleshooting), a diagnostic tool that runs **show** commands on any network device (Execute Network Device Command), a configuration validation tool (Evaluate Configuration Validator), a tool to find the cause of posture-check failures (Posture Troubleshooting), an endpoint debugging tool (Endpoint Debug), a packet capture utility (TCP Dump), a tool to test policy flow to verify policy configuration (Session Trace Tests), and various other TrustSec tools.

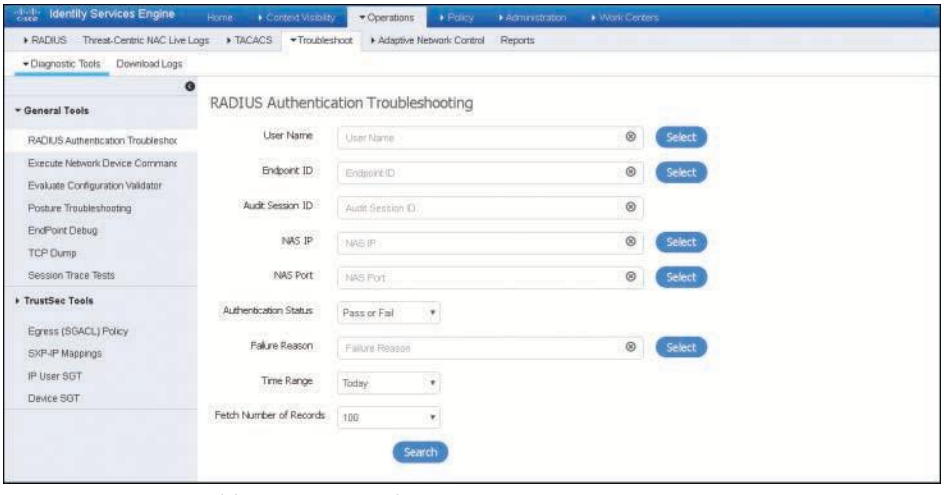


Figure 7-24 *Troubleshooting Tools*

If you cannot resolve an issue by monitoring Live Log or troubleshooting on your own, you might need to open a support case with Cisco TAC. In order to provide the customer support engineer the necessary information, you might be asked to gather the logs from ISE at the time of the issue. These logs are available as a support bundle (see Figure 7-25) or as individual debug logs (see Figure 7-26)—both under the Download Logs tab.

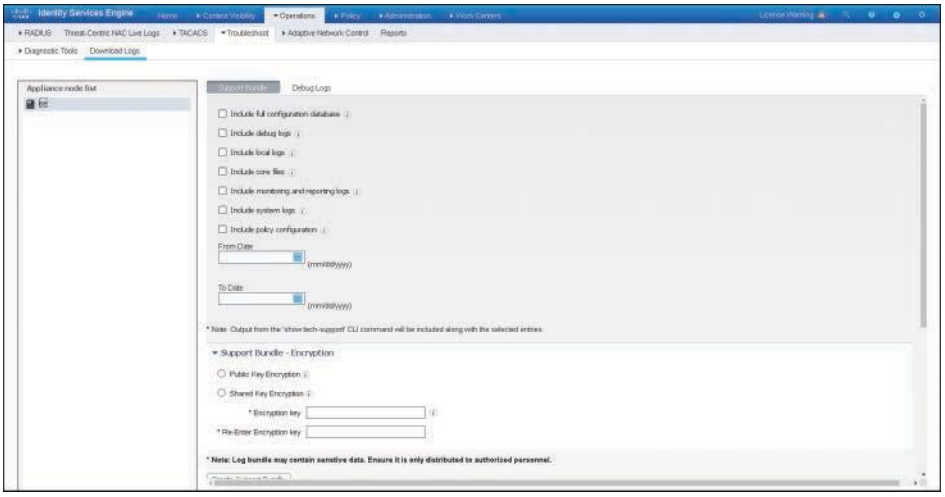


Figure 7-25 *Support Bundle*

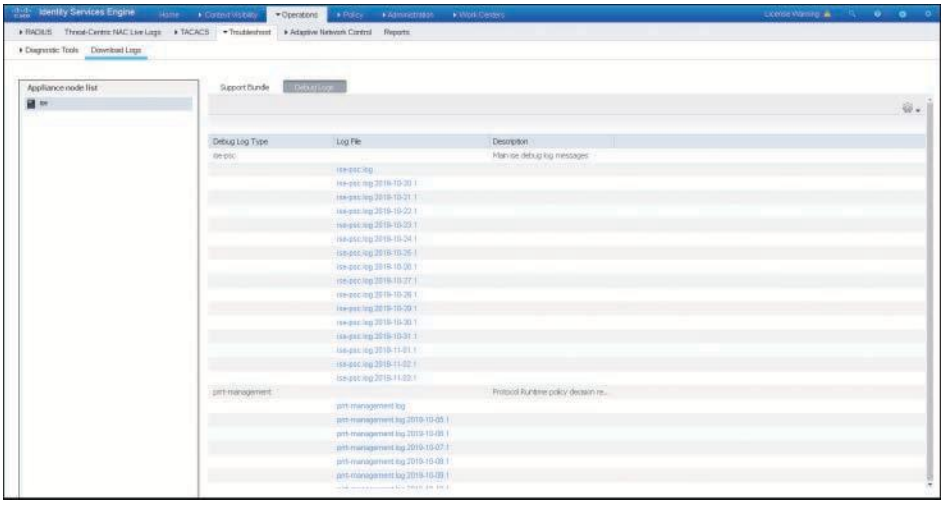


Figure 7-26 *Debug Logs*

Adaptive Network Control

Adaptive Network Control (ANC) is a very useful tool that allows an administrator to change the network access of an endpoint to quarantine, unquarantined, or shutdown. A network administrator can quarantine an endpoint by using its MAC address (which is most effective) or IP address.

The term "quarantine" in ISE is a little misleading. Quarantine is just an internal label that can be assigned to endpoints. What ISE does with an endpoint that has been labeled with "quarantine" varies depending on how a corresponding security policy is configured. For some enterprises, quarantine might mean granting Internet access only and specifying that in the security policy. For others, quarantine might mean completely rejecting a user from the network and stating that in the security policy. If there is no security policy that tells ISE what to do with quarantined endpoints, nothing happens when an endpoint is given that label. That endpoint simply continues to have the same level of access that it had prior to being given the quarantine label.

If an endpoint is quarantined—with the appropriate security policy configured—all requests for authentication are refused, resulting in the endpoint being restricted to the Layer 2 network between the endpoint and the network access device. In the case of a wired endpoint, this is usually the physical connectivity to the switch; with wireless, it is usually the connection between the endpoint and the wireless LAN controller (WLC) from which the RADIUS calls are being made. This is a very strong-handed approach to keeping a device off the network and is often used to address the most egregious security offenses on the network. In this case, the user or the endpoint is simply refused network access—without being given any concrete feedback about why (such as a redirection to a web page that provides detailed feedback).

The Adaptive Network Control component has two tabs:

- **Policy List**: This is where ANC policies, or "labels," can be created (see Figure 7-27). An ANC policy must be given a name and an action. There are three available actions: Quarantine, Shut_Down, and Port_Bounce. It's important to remember that a policy cannot do anything without a corresponding authorization policy in ISE.
- **Endpoint Assignment**: You can use this tab to view, add, and remove endpoints from an ANC policy (see Figure 7-28).

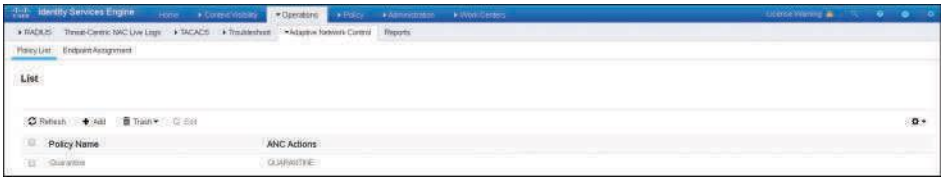


Figure 7-27 *Adaptive Network Control Policy List*

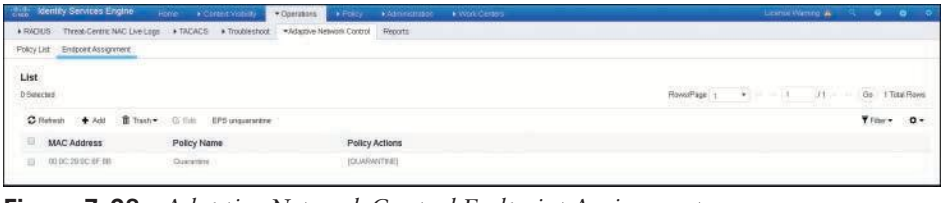


Figure 7-28 *Adaptive Network Control Endpoint Assignment*

Reports

You use the Reports tab under Operations to generate reports for ISE functions and sessions. Figure 7-29 shows the following report categories:

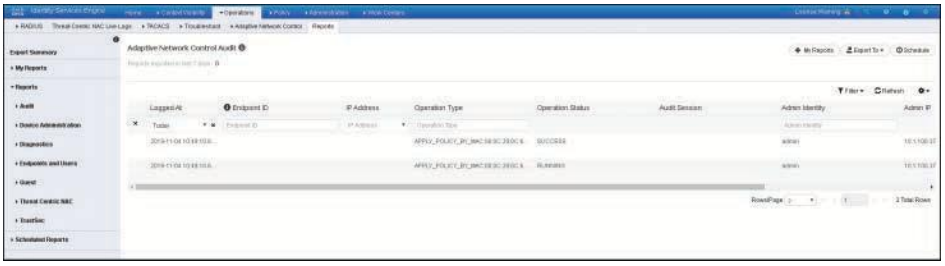


Figure 7-29 *Available Reports*

- **Audit:** Reports in this category provide details about administrator operations, change management, and the configuration of ISE.
- **Device Administration:** Reports in this category provide details about TACACS device administration authentications, authorizations, accounting, and session information.
- **Diagnostics:** Reports in this category provide information about the ISE deployment health.
- **Endpoints and Users:** Reports in this category provide details about user provisioning, endpoint profiling, and user management. These reports focus on Cisco ISE end users and endpoints.
- **Guest:** Reports in this category provide details about guest authentications, AUP reporting, and device audits.
- **Threat Centric NAC:** Reports in this category provide details specific to Threat Centric NAC functionality and security events.
- **TrustSec:** Reports in this category provide details about TrustSec operations.

Policy

You use the Policy tab in the ISE GUI to configure policy sets, profiling, posturing, client provisioning, and policy elements.

Policy Sets

Key Topic

The Policy Sets tab under Policy provides an interface to configure authentication and authorization policy for network access control in ISE (see Figure 7-30). *Authentication*, as described in Chapter 1, "Fundamentals of AAA," is the process by which the identity of a user or device is determined. The authentication policy defines the rules by which ISE identifies the user—leveraging the various RADIUS attribute/value pairs sent by the NAD to ISE, specific information about the NAD itself, and other unique criteria from the authentication session. *Authorization* is the process of determining what an endpoint device will have access to on the network once it has been authenticated. The authorization policy, much like the authentication policy, leverages the RADIUS attribute/value pairs that are sent from the NAD, specific information about the NAD itself, and other unique criteria.

Policy sets, introduced in ISE Version 1.3, provide an administrator the ability to logically group authentication and authorizing policies together in the same set. You can manage several different network access use cases—such as wireless, wired, and guest—in separate security policies. The policy sets can be configured with conditions that must be met—such as how or where the endpoint is connecting to the network—for the endpoint to be evaluated against that policy. Policy sets are hierarchical, which means they are evaluated in a top-down first-match method when a RADIUS request is received.

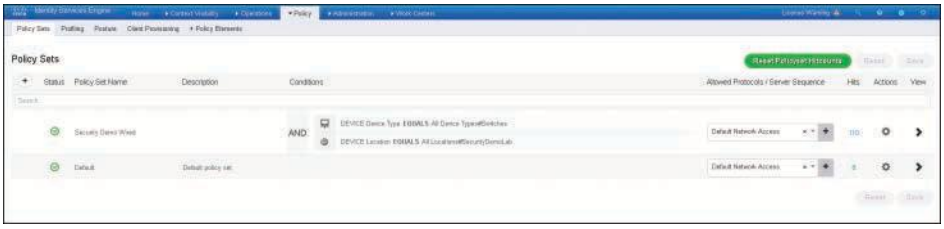


Figure 7-30 *Policy Sets*

Inside the Policy Sets tab are four main subsections that may contain rules. These rules are also evaluated in a top-down first-match manner just as the policy sets are (see Figure 7-31):

- **Authentication Policy**: You use this section to configure rules and conditions for identifying the endpoint and user.
- **Authorization Policy—Local Exceptions**: You use this section to configure local authorization exception rules that are specific to a policy set.
- **Authorization Policy—Global Exceptions**: You use this section to configure global authorization exception rules that are shared across all policy sets.
- **Authorization Policy**: You use this section to configure requirements and conditions to meet in order to grant some result or access.

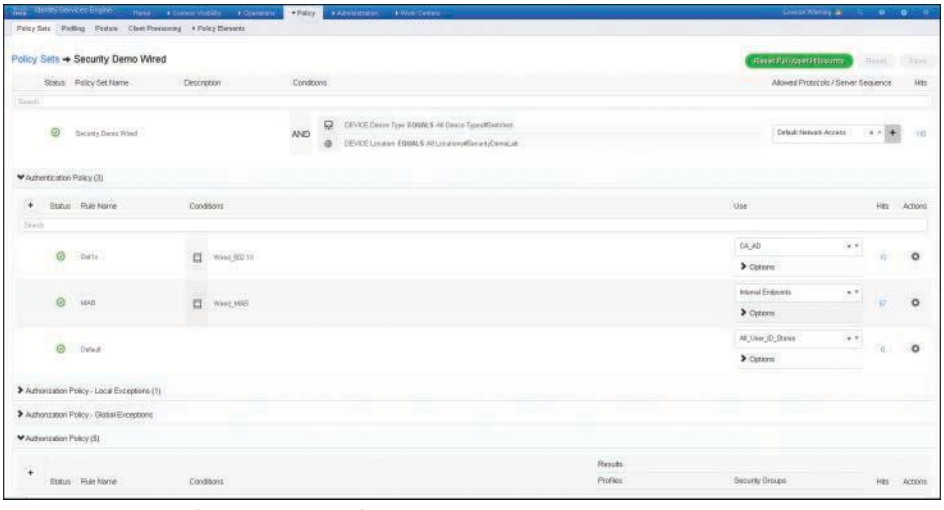


Figure 7-31 *Policy Set Example*

Profiling

ISE has the ability to determine (or at least make a very educated guess about) what type of endpoint is attempting to access the network. This process, called *profiling*, is accomplished through ISE's monitoring of protocols sent by the endpoint onto the network (DHCP, HTTP, and, by proxy, RADIUS, among a few others). The profiling done by ISE can sometimes be general in nature (such as identifying just the vendor or just the device type) or very specific (such as identifying the endpoint's vendor, device model, and function on the network). You can use the Profiling tab to create custom profiles and to view the existing profiles in ISE (see Figure 7-32).

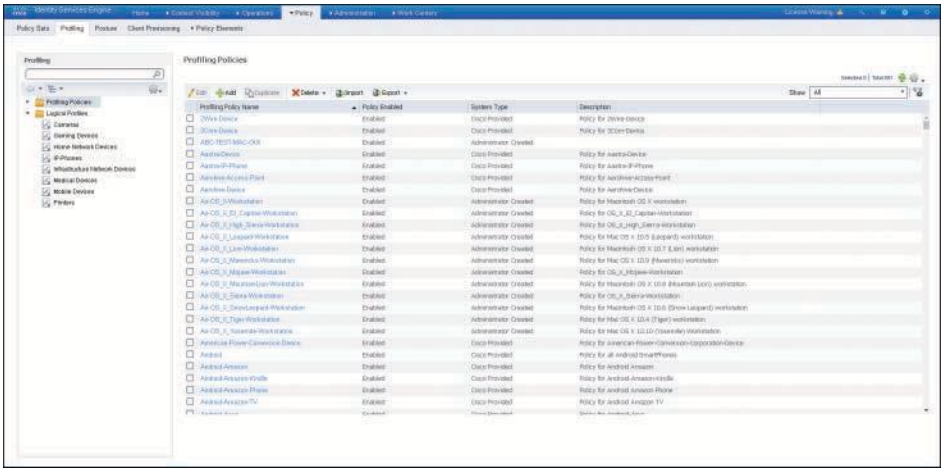


Figure 7-32 *Profiling Policy*

The Profiling Policies section, shown in Figure 7-32, allows an administrator to leverage the Cisco preconfigured profiling policy to identify endpoints and/or create new policies by defining the rules or criteria by which the type of endpoint that is accessing the network is determined.

Posture

To configure a posture policy in ISE, you need to access the Posture tab under Policy. Posturing is a mechanism by which a software supplicant or an agent on an endpoint provides ISE detailed information about the endpoint's software and hardware configuration. The information provided by the endpoint may include information about the operating system, antivirus/antispyware/antimalware software, and current OS patch levels (see Figure 7-33). Depending on the operating system, other important information can be gathered from the endpoint as part of the posturing process.

Based on the final assessment of the endpoint's posture compared to the configured posture policy in ISE, ISE may quarantine an endpoint that has a noncompliant or unknown posture—possibly providing the endpoint reduced network access in order to remediate its posture. After taking all remediation steps needed to become compliant, the endpoint may be allowed appropriate network access.

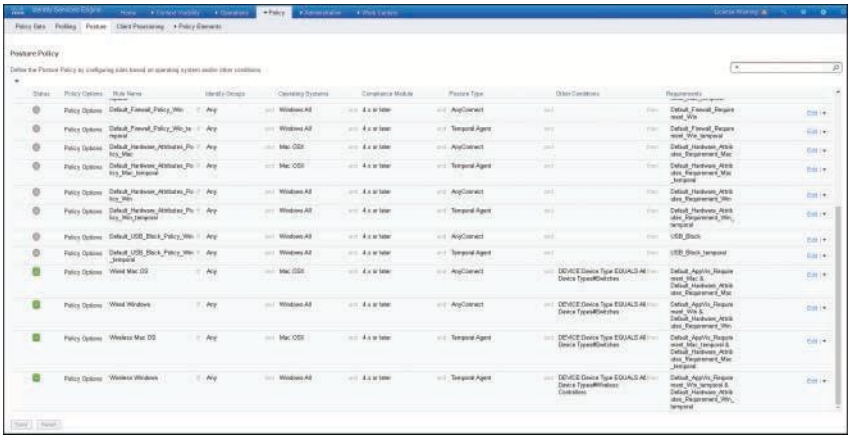


Figure 7-33 *Posture Policy*

Client Provisioning

In the Client Provisioning tab under Policy, you can configure how each endpoint type is provisioned to join the network. The client provisioning policy configures the endpoint supplicant for BYOD and posturing. This is sometimes be referred to as *onboarding* the endpoint. As an endpoint attempts to onboard onto the network, the appropriate security credentials (including X.509 certificates and wireless profiles) will be downloaded onto a BYOD endpoint. If the endpoint must be postured, the posture client must be downloaded and installed on the client. After the endpoint is onboarded, the user of the endpoint can join the network in the future with minimal manual interaction.

In order to accomplish this client provisioning, ISE must provide the appropriate information in a manner understood by the endpoint. Via the client provisioning policy (see Figure 7-34), based on the OS and other criteria provided by the endpoint and the NAD, ISE presents the required credentials (X.509 certificates and wireless profiles) to the endpoint to store for future use. This might sometimes require a supplicant built in to the OS or one that is automatically downloaded from ISE (often a Java or ActiveX-based supplicant).

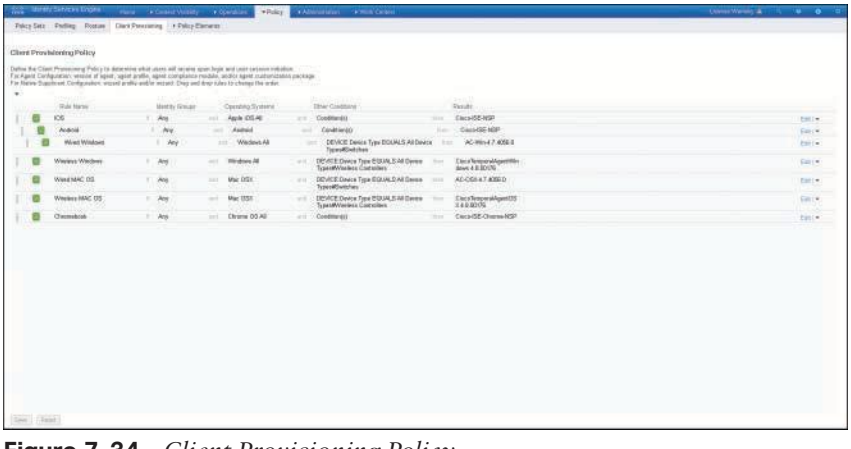


Figure 7-34 *Client Provisioning Policy*

Policy Elements

Key Topic

Within the Policy screen are a number of different policies that can be configured. These policies are often composed of numerous building blocks, or elements, that help simplify the overall configuration of ISE policy—either by facilitating reuse across multiple policies or by providing elements that are more human readable. The Policy Elements tab under Policy is the place in ISE where you can add these reusable building blocks.

Figure 7-35 shows that there are three tabs under Policy Elements:

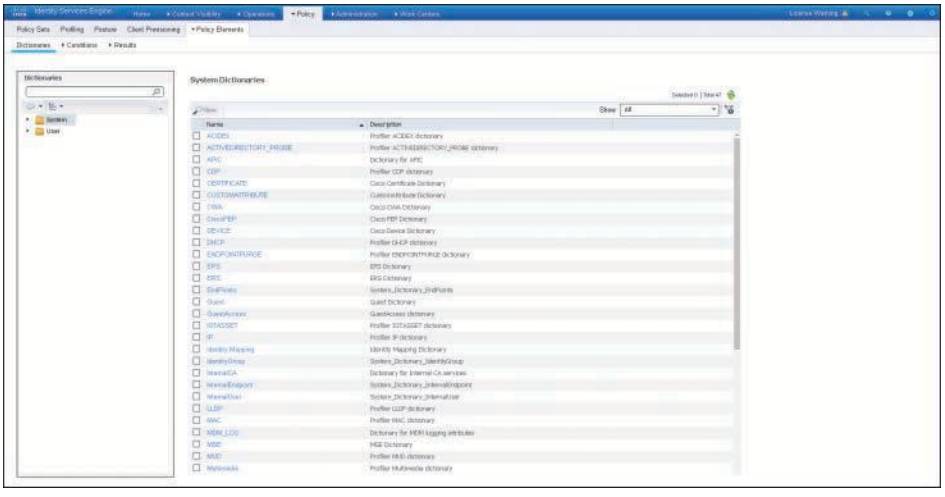


Figure 7-35 *Policy Elements*

- **Dictionaries:** This tab provides a list of system- and user-defined elements and their available attributes. The individual dictionary elements leverage an attribute/value pairs approach, similar to RADIUS, whereby each chosen attribute is assigned a value either by the endpoint, the NAD, or ISE. These attributes and their assigned values can be used as the smallest building blocks in the ISE policy.
- **Conditions:** Many of the policy constructs in ISE leverage an IF-THEN programmatic approach. With IF-THEN clauses, the IF portion of the clause is a conditional statement—a "test" statement that can take on either a TRUE or FALSE value. If the totality of this conditional statement is TRUE, everything following the THEN portion of the clause is executed. The Conditions tab under Policy Elements is where you define the IF portions of ISE policy. Conditions may be nested and combined in a policy to create complex rules that match almost any use case.
- **Results:** The Results tab under Policy Elements is where you define the THEN portions of ISE policy. Therefore, if an IF statement (the conditions) results in a TRUE outcome, everything following the THEN (the results) will be executed.

Administration

The Administration tab in the ISE GUI includes a number of setup-type functions. These functions are typically set up one time and rarely modified thereafter. The Administration section of the ISE GUI includes the tabs System, Identity Management, Network Resources, Device Portal Management, pxGrid Services, Feed Service, and Threat Centric NAC.

System

The System tab under Administration allows an administrator to configure the following elements, which affect ISE's behavior directly:

- **Deployment:** You use the Deployment tab (see Figure 7-36) to configure all ISE nodes. Each node must be defined and registered in the deployment configuration on the PAN. As discussed in Chapter 6, each node is assigned at least one persona. In the Deployment tab, you can assign at least one persona, role (primary, secondary, or standalone), and IP address or DNS-resolvable hostname to each node.

  The Deployment tab is also where you enable the profiling function on a per-node basis—to indicate which PSN nodes actively participate in endpoint profiling.

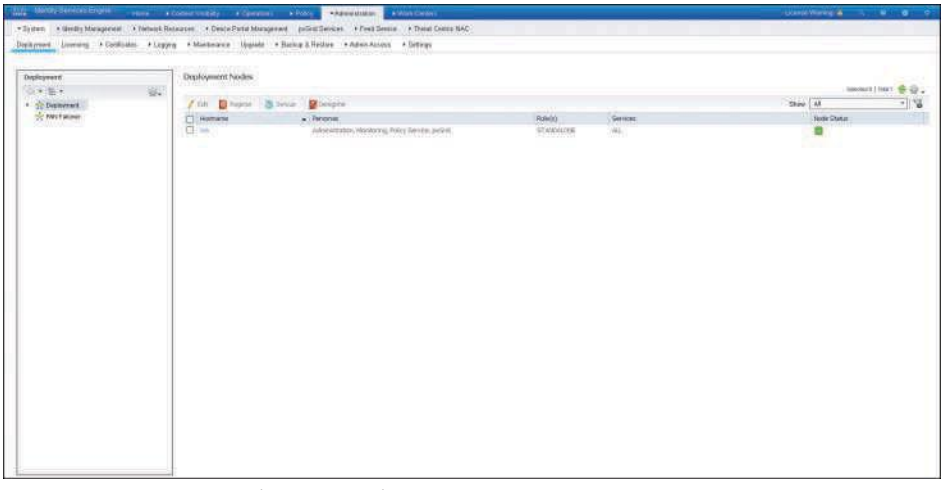


Figure 7-36 *ISE Deployment Tab*

- **Licensing:** Every feature in ISE requires a license—whether that feature is a basic feature or a more advanced feature. Many of the common features that do not require external resources are grouped under the Base license. Features that require periodic updates from Cisco require either a Plus license or an Apex license. A Device Administration license is required for each PSN that will be responding to TACACS requests. By default, ISE comes with a 90-day Evaluation license. Table 7-3 breaks down the features covered with each license.

Table 7-3 Cisco ISE License Packages

License Package	Perpetual or Subscription	ISE Functionality Covered	Notes
Base	Subscription (1, 3, or 5 years)	Basic network access: AAA, IEEE-802.1X	Passive identity services are available as part of the upgrade from ISE-PIC to a Base license, including limited pxGrid features available to Cisco subscribers only.
		Guest management	
		Link encryption (MACsec)	
		TrustSec	
		ISE application programming interfaces	
Plus	Subscription (1, 3, or 5 years)	Bring your own device (BYOD) when consuming either a built-in or external certificate authority	Does not include Base services. A Base license is required to install the Plus license. When onboarding an endpoint with the BYOD flow, the Plus services are consumed on the active session even when related BYOD attributes are not in use. Plus licenses are supposed to be consumed when profiling-related authorization policies contain *IdentityGroup*:*Name*.
		MSE integration for location services	
		Profiling and feed services	
		Adaptive Network Control (ANC)	
		Cisco pxGrid	
Apex	Subscription (1, 3, or 5 years)	Third-party mobile device management (MDM) integration	Does not include Base services. A Base license is required to install the Apex license. Note that when you use Cisco AnyConnect as the unified posture agent across wired, wireless, and VPN deployments, you need Cisco AnyConnect Apex user licenses in addition to Cisco ISE Apex licenses.
		Posture compliance	
		TC-NAC	
Mobility	Subscription (1, 3, or 5 years)	A Mobility license turns on the functionality of Base and Apex licenses for wireless LAN deployments.	Cannot coexist on a Cisco PAN with Base, Plus, or Apex licenses.

License Package	Perpetual or Subscription	ISE Functionality Covered	Notes
Mobility Upgrade	Subscription (1, 3, or 5 years)	A Mobility Upgrade license turns on the functionality of Base and Apex licenses for all wireless and non-wireless client-access methods, including wired and VPN concentrator access.	You can only install a Mobility Upgrade license on top of an existing Mobility license.
Device Administration	Perpetual	TACACS+	A Base or Mobility license is required to install the Device Administration license. The number of Device Administration licenses must be equal to the number of PSNs with the TACACS+ persona enabled on them.
ISE-PIC	Perpetual	Passive identity services	One license per node. Each license supports up to 3000 parallel sessions.
ISE-PIC Upgrade	Perpetual	This license can have up to 300,000 additional parallel sessions enabled. It can also be upgraded to a full ISE-PIC license.	One license per node. Each license supports up to 300,000 parallel sessions. After installing this license, the upgraded node can join an existing ISE deployment or, alternatively, Base licenses can be installed on the node to make it function as the PAN. Passive identity services available as part of the upgrade to a Base license include limited pxGrid features available to Cisco subscribers only.
Evaluation	Temporary (90 days)	Full Cisco ISE functionality is provided for 100 endpoints.	Limited use of Cisco ISE product for presale customer evaluation. All Cisco ISE appliances are supplied with an Evaluation license.

Figure 7-37 shows the ISE Licensing tab, where licenses can be added or ISE clusters can be connected to a Cisco Smart Licensing account.

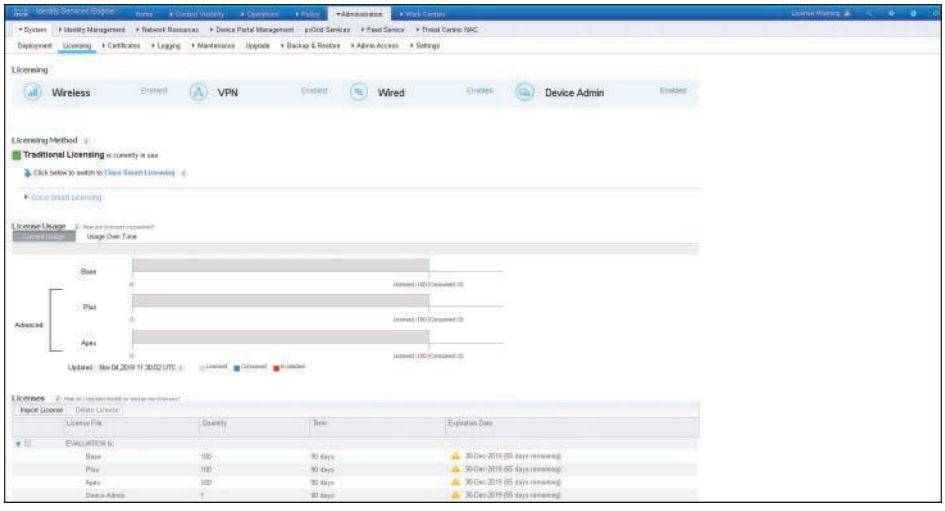


Figure 7-37 *ISE Licensing Tab*

- **Certificates:** X.509 Certificates (often referred to simply as *certificates*) are the key components to a public-key infrastructure (PKI), providing a scalable deployment option that marries each endpoint to the user and also providing nonrepudiation. In order to ensure that ISE can properly authenticate endpoints and the endpoints can properly authenticate ISE, you can leverage the X.509 certificates for both server and client authentication.

  The Certificates tab (see Figure 7-38) allows an administrator to configure the certificates that will be used for the Admin portal as well as any portal that is served to the client as part of the onboarding or remediation process. The certificate authorities needed to authenticate the endpoints or infrastructure devices (such as mobile device managers) must also be loaded here. Certificates are discussed in greater detail in Chapter 8.

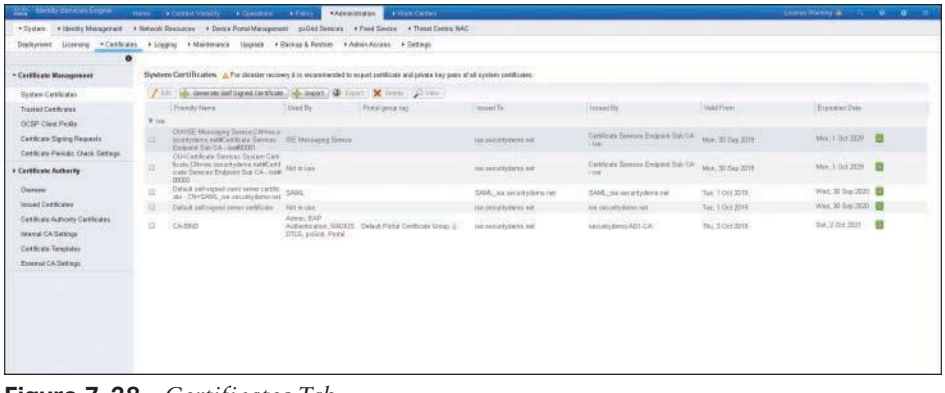


Figure 7-38 *Certificates Tab*

- **Logging:** The Logging tab (see Figure 7-39) is where an administrator configures the specifics of how ISE creates syslogs. In this tab, the administrator can define how long logs will be kept, how detailed the logs will be, and where to send the logs. These logs may be required as part of a compliance program (such as PCI-DSS or HIPAA), a security policy, or to simply monitor how ISE is performing. If the logging level chosen is too low, the amount of detail provided in these logs may not be enough to address any issues as they occur. If the logging level chosen is too high, the amount of log information generated may adversely affect storage space on the ISE appliance and/or affect the overall performance of the appliance.

Key Topic

The level of logging that is enabled should be carefully chosen and well balanced to facilitate ongoing operations while providing enough detail to identify any issues as they arise. In case of any issues, Cisco TAC may request detailed logs at the time of the issue. If debug-level logs are requested, it is best to enable these logs temporarily to gather the requested level of detail. As soon as the requested information is gathered with the debug-level logs, it is advised to return the logging level to a more neutral level.

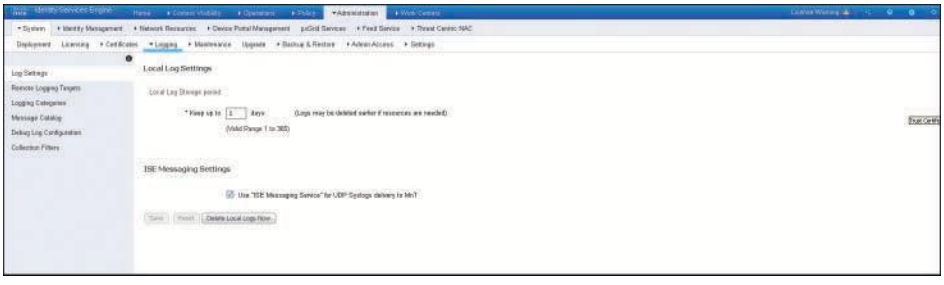


Figure 7-39 *Logging Tab*

- **Maintenance:** As with software systems from other vendors, Cisco may periodically release software updates to add additional features to ISE or to resolve any defects. Cisco periodically compiles patches to address a number of defects that have been found in previous ISE releases. To load these patches, to manage download locations for software maintenance, or to manage how often operations data is purged to maintain optimal system performance, you access the Maintenance tab under System, as illustrated in Figure 7-40. This tab is also where a repository may be configured. A repository is an off-system storage location that ISE can use to upload or download files to. ISE uses repositories to upload ISE backups and download upgrade packages.

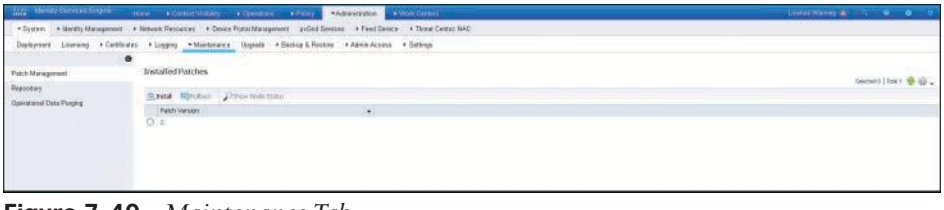


Figure 7-40 *Maintenance Tab*

- **Upgrade:** ISE provides a GUI-based centralized upgrade page that lists all the nodes in the deployment, the personas enabled on them, and the version of ISE installed and the status of each node. You can begin the upgrade process from this screen, which is shown in Figure 7-41.



Figure 7-41 *Upgrade Tab*

- **Backup & Restore:** As part of any disaster recovery process, it is important for an administrator to periodically perform a backup of all network configurations. The Backup & Restore subsection in the ISE GUI (see Figure 7-42) allows an administrator to back up both the ISE configuration and ISE operational data, either automatically on a preconfigured schedule or manually, sending the data to a chosen repository. Then, if disaster strikes, the configuration can be restored.

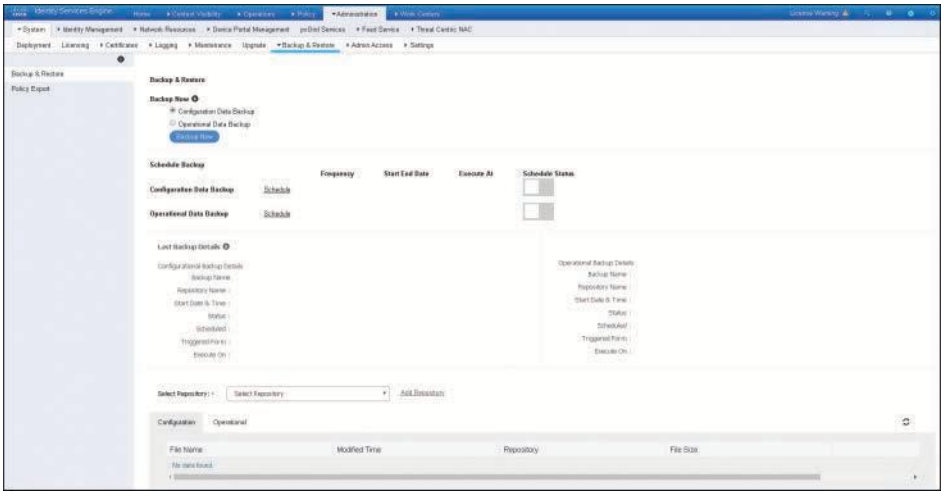


Figure 7-42 *Backup & Restore Tab*

- **Admin Access:** The Admin Access tab (see Figure 7-43) does exactly as the name implies: It defines how an administrator can access ISE. Some of the options available in this tab include usernames and passwords for administrative users, what administrative groups are defined, the requisite password policy for administrative users, and from where and for how long administrators are allowed to access the ISE Admin portal. Additional options are also available in the Admin Access tab.

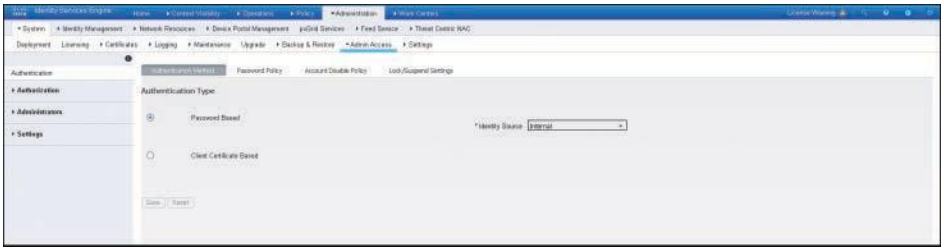


Figure 7-43 *Admin Access Tab*

- **Settings:** The Settings tab under System (see Figure 7-44) allows an administrator to define how the system will behave. Some features that are configured via the Settings tab include client provisioning support, endpoint protection services, FIPS mode, posture timers, and profiling CoA behavior. In addition, proxy, SMTP, and NTP servers can be defined here.

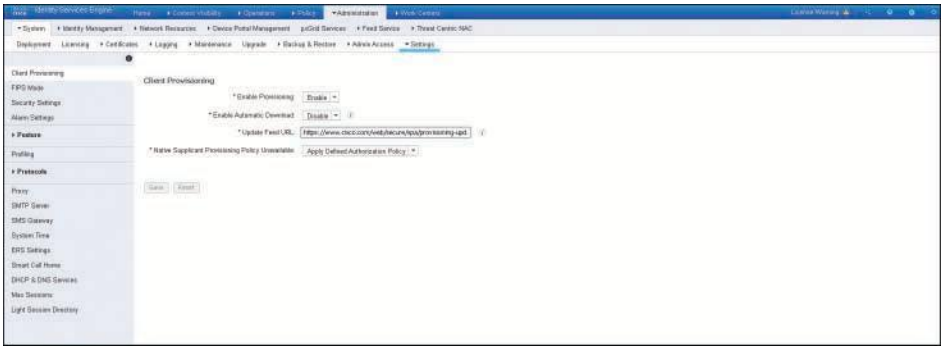


Figure 7-44 *Settings Tab*

Identity Management

Identity management is a key construct in ISE. Nearly every step in the ISE authentication and authorization processes requires at least one identity as the subject of the authentication or authorization operations. The Identity Management tab under Administration (see Figure 7-45) allows an administrator to manage all ISE configurations that pertain to identities. It includes the following tabs:

- **Identities:** The Identities tab is used to manage both local ISE non-guest network access users and endpoints discovered via manual NMAP scans.

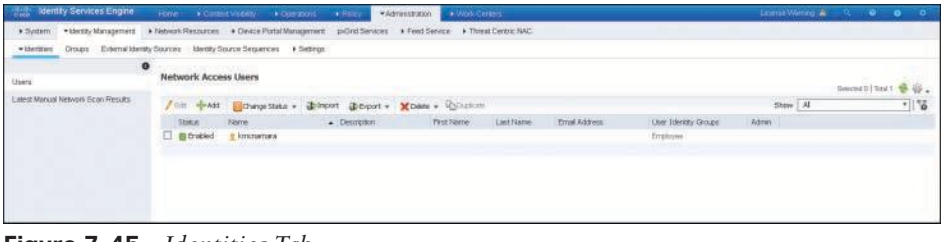


Figure 7-45 *Identities Tab*

- **Groups:** Identity groups allow administrators to subdivide and manage users or endpoints based on common criteria (see Figure 7-46). By grouping users or endpoints, the administrator can more easily enforce any associated policy based on which group a particular user or endpoint is part of. For instance, all smartphones may have one security policy, while all laptop computers may have a slightly different security policy.

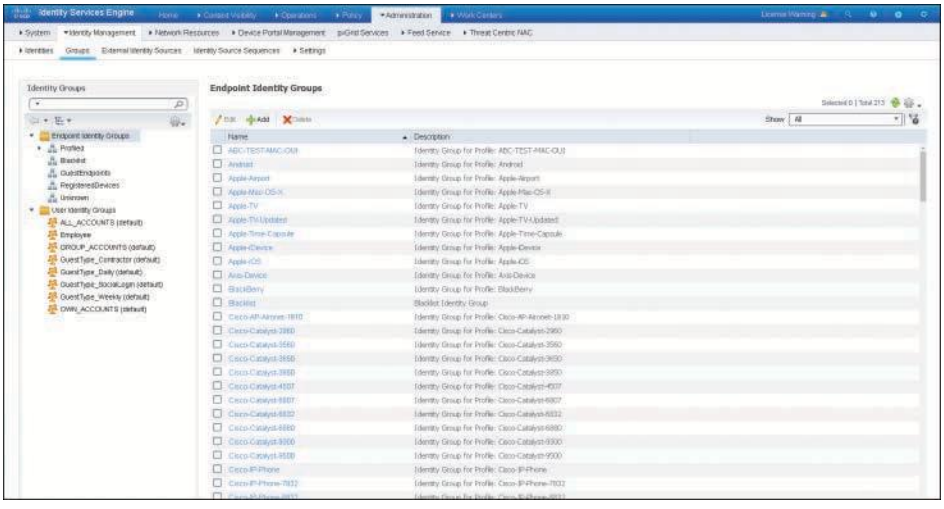


Figure 7-46 *Groups Tab*

- **External Identity Sources:** In many ISE deployment scenarios, ISE is deployed on a network that has been running for a while. Often, the user identities and user management are stored outside ISE instead of each identity being imported into ISE, and the users must be managed in two different places. Therefore, ISE allows for external identity sources (see Figure 7-47). These external identity sources may include Active Directory, LDAP, RSA SecurID (token-based two-factor authentication), ODBC, RADIUS Token, SAML ID, and social media login. Depending on the authentication policy configured in ISE, a user's or endpoint's credentials may be authenticated to any one of these external identity sources. We discuss authentication in more detail in Chapter 8, "Initial Configuration of Cisco ISE," and Chapter 9, "Authentication Policies."

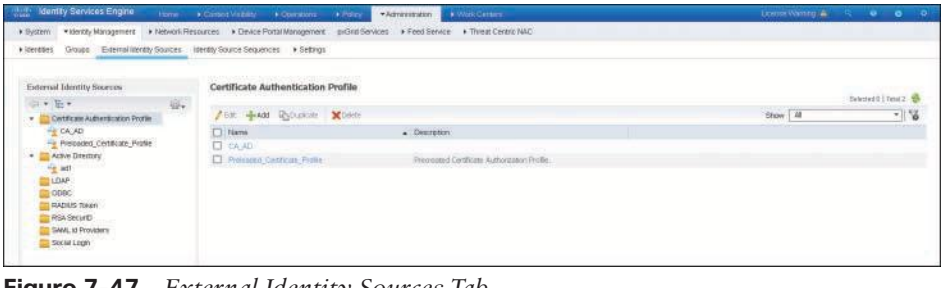


Figure 7-47 *External Identity Sources Tab*

- **Identity Source Sequences:** Used in conjunction with the local and external identity sources just mentioned, identity source sequences (see Figure 7-48) allow an administrator to sequentially check a number of identity sources for a successful authentication. This type of sequential search can be useful if users or endpoints can exist in multiple identity databases, as the administrator can use a single rule to check all identity stores sequentially rather than search each identity store separately. This approach can greatly simplify authentication policy. The authentication process ends upon the first successful authentication, based on the chosen authentication policy. We discuss authentication in more detail in Chapters 9 and 10.

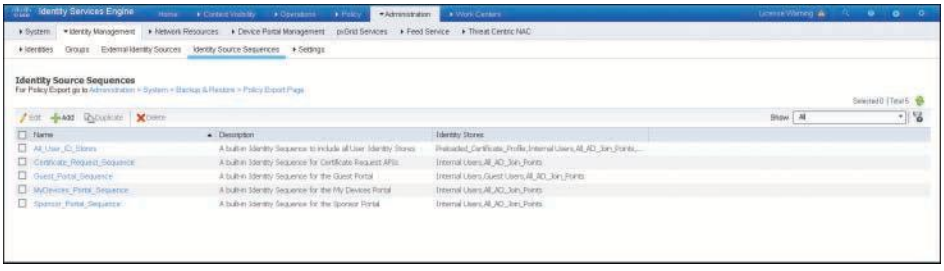


Figure 7-48 *Identity Source Sequences Tab*

- **Settings:** The Settings tab under Identity Management (see Figure 7-49) is where you define custom attributes that can be associated to each user, endpoint purge setting, and user password policy.

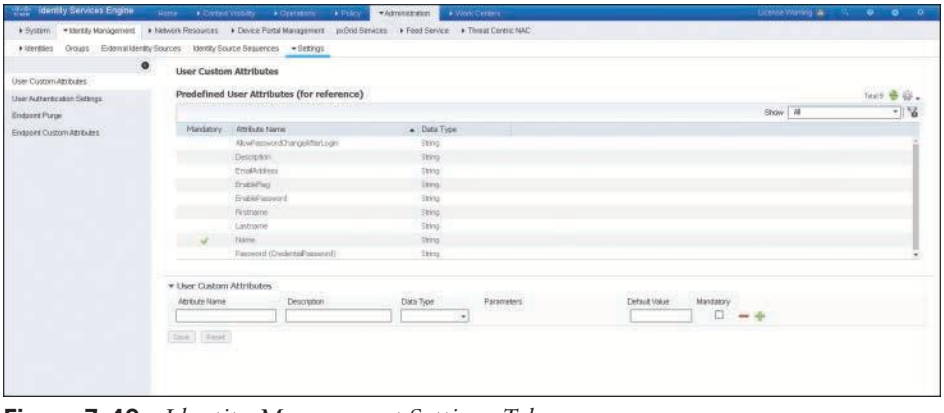


Figure 7-49 *Identity Management Settings Tab*

Network Resources

The Network Resources tab under Administration is where all external supplementary resources are configured. These resources include all the network access devices (NADs), external RADIUS servers, NAC managers, and MDMs. You can use the following tabs under Network Resources:

- **Network Devices:** Whenever a new NAD—such as a switch, router, firewall, or wireless LAN controller (WLC)—is added to a network, it may be necessary to add the

device's credentials to ISE in order to allow the NAD to authenticate endpoints to ISE. For each NAD that is added to the network devices database, additional information can be provided about the NAD, including the IP address(es), model name, software version, device location, and device type. In addition, the appropriate network services can be configured, including RADIUS, SNMP, and TrustSec. Without the requisite protocol being selected and appropriately configured, the NAD in question will not be able to access that service on ISE.

Unless it is added to the Network Devices tab (see Figure 7-50), the only other way for a NAD to use ISE for AAA is if a default device is configured. Default device is disabled by default. If it is enabled and configured, ISE accepts RADIUS requests from all NADs that share the configured default device RADIUS shared secret.

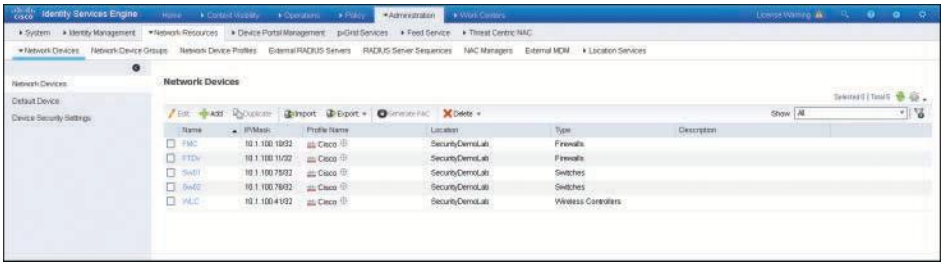


Figure 7-50 *Network Devices Tab*

- **Network Device Groups**: As the number of NADs on a network increases, the ability to manage these network devices can become more and more difficult. Furthermore, as the ISE security policy is configured, grouping the NADs according to certain criteria may greatly simplify the policy rules.

  By default, ISE allows for each device to have a device type and device location component—and additional criteria can be added. The device type may refer to the functional role of the device on the network, such as a router, switch, firewall, or WLC. The device location may refer to the geographic location of the device—such as in Ohio or in North Carolina—or the location within the network topology—such as at a branch location or on the campus. In either case, the device type and device location can be very general or as granular as you wish.

  As shown in Figure 7-51, network device groups can be used within the security policy to provide a differentiated level of access to the endpoints—for example, allowing for a different level of access (more secure or less secure) if the endpoint is authenticating to a NAD at a branch location versus at the campus location. Using device groups also gives a network administrator flexibility in terms of the policy that is pushed to the NAD. If certain NADs support more features (such as TrustSec-capable devices) or if certain device types require specific RADIUS attribute/value pairs (such as wireless access lists), you might choose to separate those devices into their own device group (either by leveraging the default device type group or adding your own custom device group).

You'll learn more about differentiating levels of access in Chapters 9 and 10.

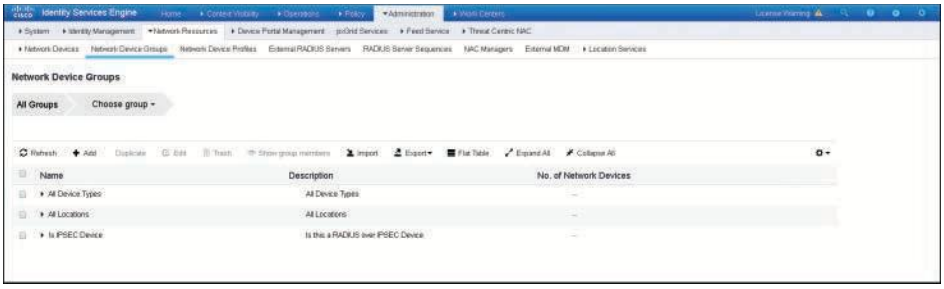


Figure 7-51 *Network Device Groups Tab*

- **Network Device Profiles:** ISE supports third-party NADs through the use of network device profiles. The NAD profiles are configured to define the capabilities of these third-party devices.
- **External RADIUS Servers:** If your organization already has a RADIUS server in place, you can continue to leverage and supplement this RADIUS server as you implement ISE. If an external RADIUS server is in use, ISE can receive the RADIUS requests and proxy a RADIUS request to the external server. The response from the external RADIUS server is also proxied by ISE back to the original NAD.
- **RADIUS Server Sequences:** RADIUS server sequences leverages external RADIUS servers. After configuring external RADIUS servers, you can select one or more of these servers to create a RADIUS server sequence. Furthermore, you can elect whether to modify the RADIUS request before it is forwarded to the external RADIUS server and whether to modify the RADIUS Accept before it is forwarded to the NAD. Again, ISE is acting as a RADIUS proxy in both the forward and reverse directions.
- **NAC Managers:** NAC (Network Admission Control) managers are the Cisco clean access managers on the network that are used to authenticate, authorize, evaluate, and remediate wired, wireless, and remote users and their machines before users are allowed on the network. The solution identifies whether machines are compliant with security policies and repairs vulnerabilities before permitting access to the network.
- **External MDM:** Mobile device managers are third-party security software systems that allow an administrator to configure and secure an endpoint, regardless of where the device is located (see Figure 7-52). This is accomplished by leveraging an MDM agent on the endpoint—either a separate application or an agent that is built in to the OS—and an MDM server. This MDM server can reside on-premises to the corporation or in the cloud.

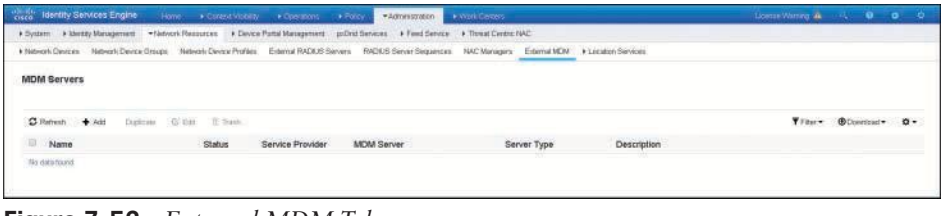


Figure 7-52 *External MDM Tab*

- **Location Services:** ISE can be integrated with Cisco Mobility Services Engine (MSE) to provide location-based authorization. Using the information provided by MSE, ISE can adjust an endpoint's network access based on the physical location of the user.

Device Portal Management

The Device Portal Management tab under Administration is where all the personal device portals are configured. These portals can be used for blacklisted endpoints, BYOD, certificate provisioning, client provisioning, mobile device management, My Device management, and the upload of custom portal files. Under Device Portal Management, you can access the following tabs:

- **Blacklist:** The Blacklist portal provides information to the endpoints that have been blacklisted and cannot gain further access to the network beyond the redirection to this portal. There is only one default Blacklist portal that may be customized but no additional Blacklist portals may be created (see Figure 7-53).

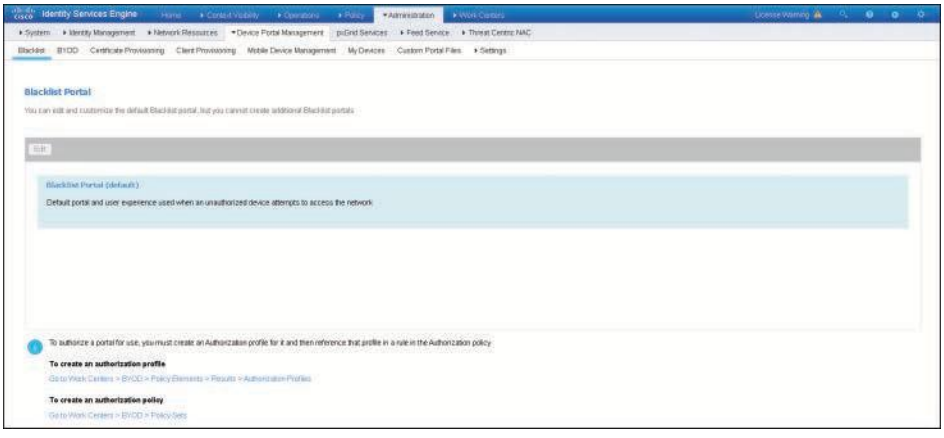


Figure 7-53 *Blacklist Tab*

- **BYOD:** BYOD portals enable employees to register their personal devices by using the supplicant provisioning functionality (see Figure 7-54).
- **Certificate Provisioning:** This tab provides employees and administrators the ability to request user or device certificates for devices to be onboarded (see Figure 7-55).
- **Client Provisioning:** The client provisioning portal forces users to download into their devices a posture agent that checks for compliance (see Figure 7-56).

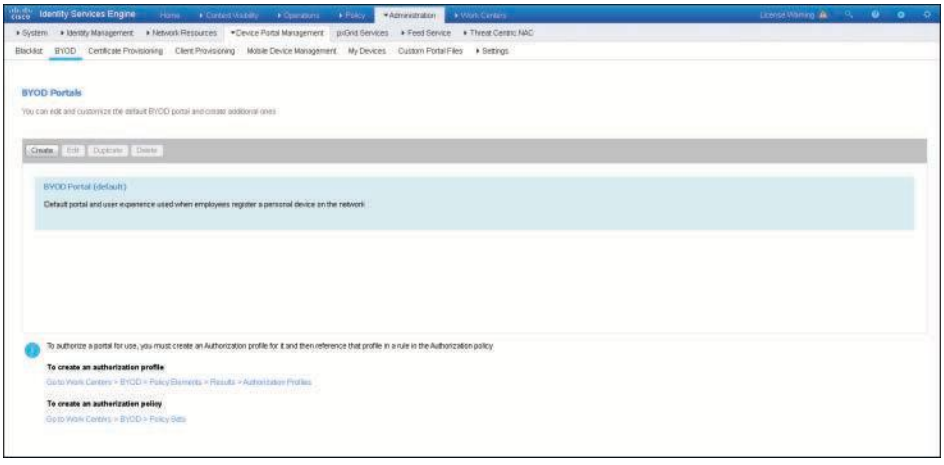


Figure 7-54 *BYOD Tab*

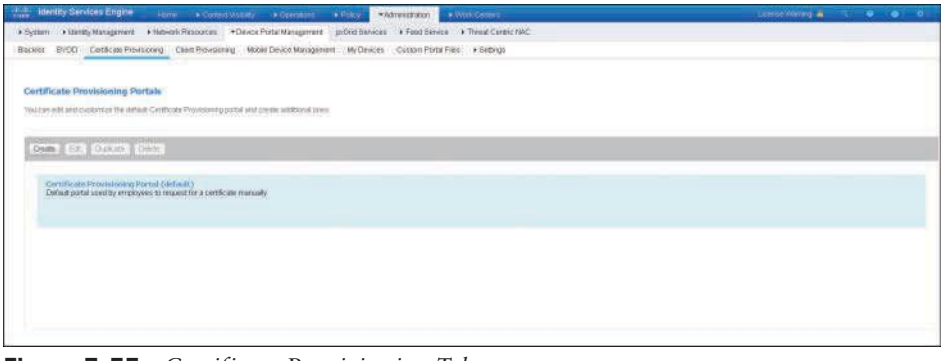


Figure 7-55 *Certificate Provisioning Tab*

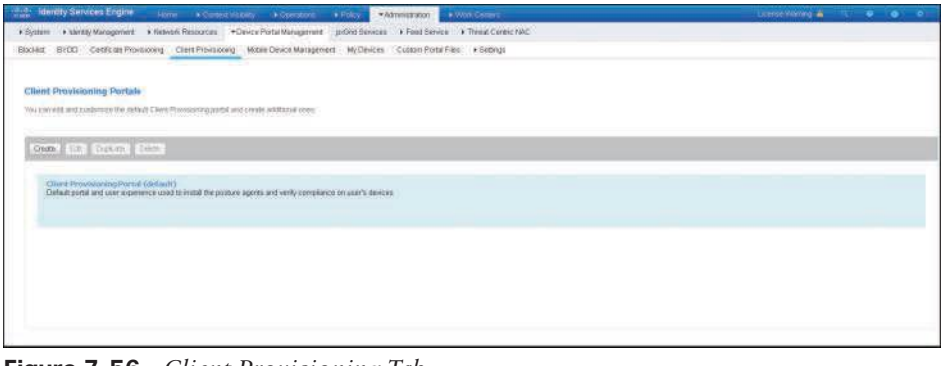


Figure 7-56 *Client Provisioning Tab*

- **Mobile Device Management:** The MDM portal provides employees with the ability to enroll their devices with an external MDM system (see Figure 7-57).

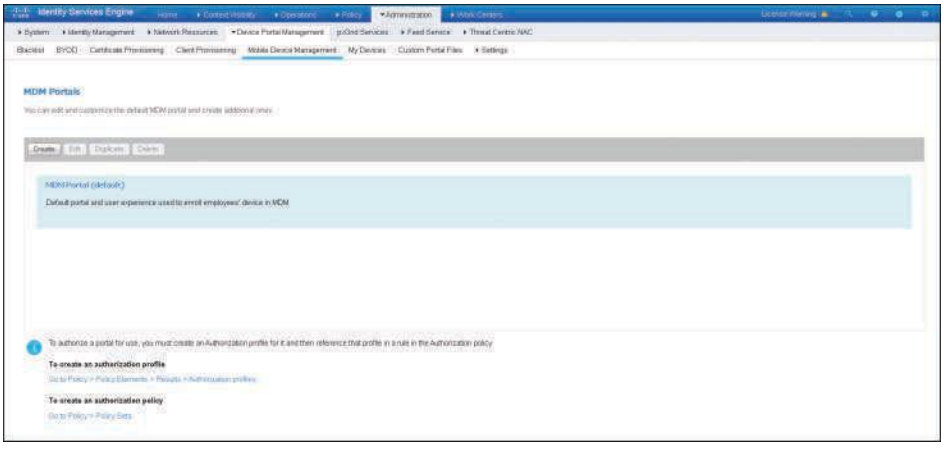


Figure 7-57 *Mobile Device Management Tab*

- **My Devices:** The My Devices portal empowers employees to add, register, and manage their personal devices, including those that do not support a native supplicant (see Figure 7-58).

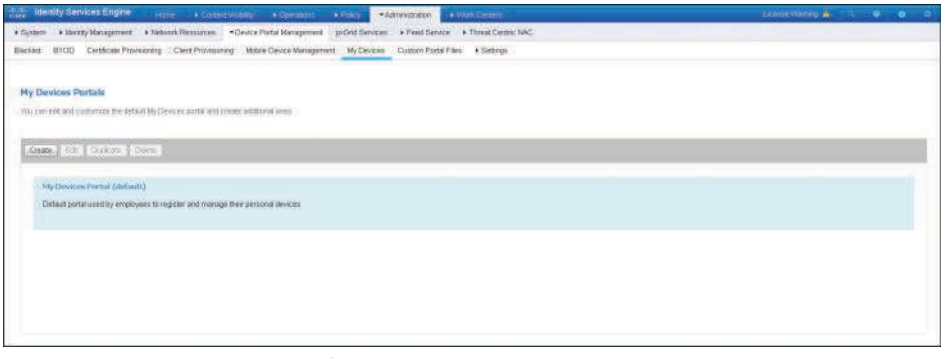


Figure 7-58 *My Devices Tab*

- **Custom Portal Files:** This tab, shown in Figure 7-59, lets administrators upload their own files to the ISE server, and these files can be used to customize all user-facing portals with the exception of the Admin portal.

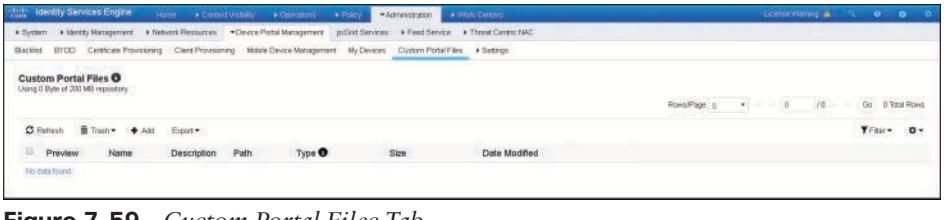


Figure 7-59 *Custom Portal Files Tab*

- **Settings:** The Settings tab under Device Portal Management allows a user to customize the operational portals of ISE (see Figure 7-60).

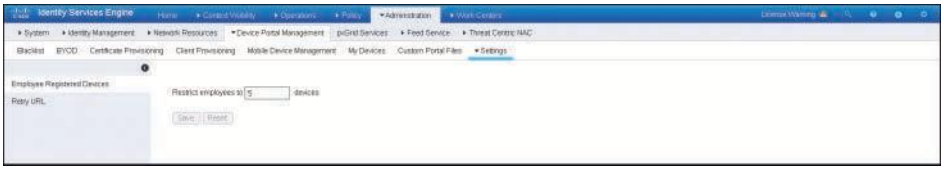


Figure 7-60 *Device Portal Management Tab*

pxGrid Services

Cisco pxGrid is a function of ISE that can be used to share context-sensitive information from ISE's session direction with other systems in the ISE partner ecosystem. Furthermore, pxGrid integration can be used to trigger adaptive network control actions to quarantine endpoints in response to network or security events. To integrate ISE with a third-party system via pxGrid, ISE and the other system must form a certificate trust. On the pxGrid Services tab (see Figure 7-61), you can view the subscribed third-party pxGrid clients, issue certificates, view Live Log for a pxGrid event, configure permissions for a pxGrid client, and more.

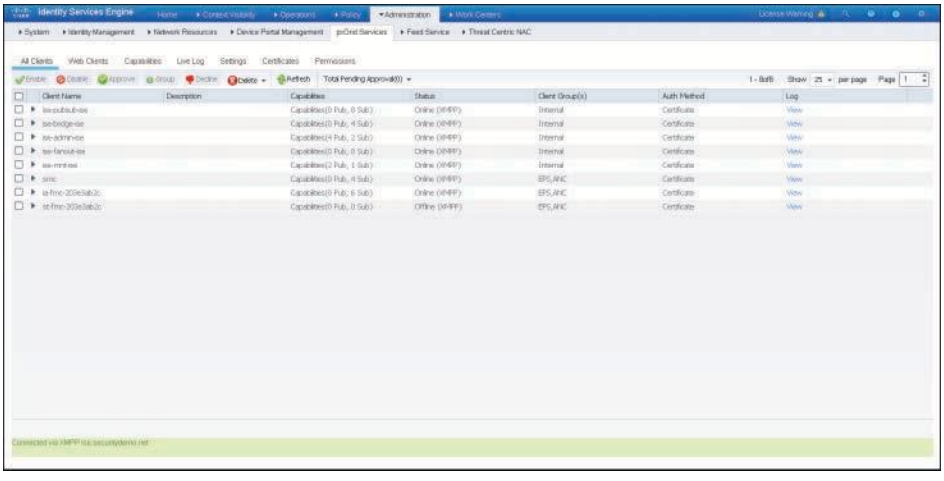


Figure 7-61 *pxGrid Services Tab*

Feed Service

The Feed Service tab enables periodic updates to ISE's Profiler service (see Figure 7-62). The Profiler service makes it possible for ISE to identify the specifics of endpoints that are accessing the network. As discussed in Chapter 5, "Introduction to Advanced Concepts," ISE profiles endpoints based on RADIUS attribute/value pairs (such as the MAC address of the endpoint), DHCP traffic, HTTP traffic, and other network communications. As endpoint vendors develop new endpoints, ISE can download the updated database of the criteria (for example, MAC address, DHCP parameters) needed to identify the new endpoint.

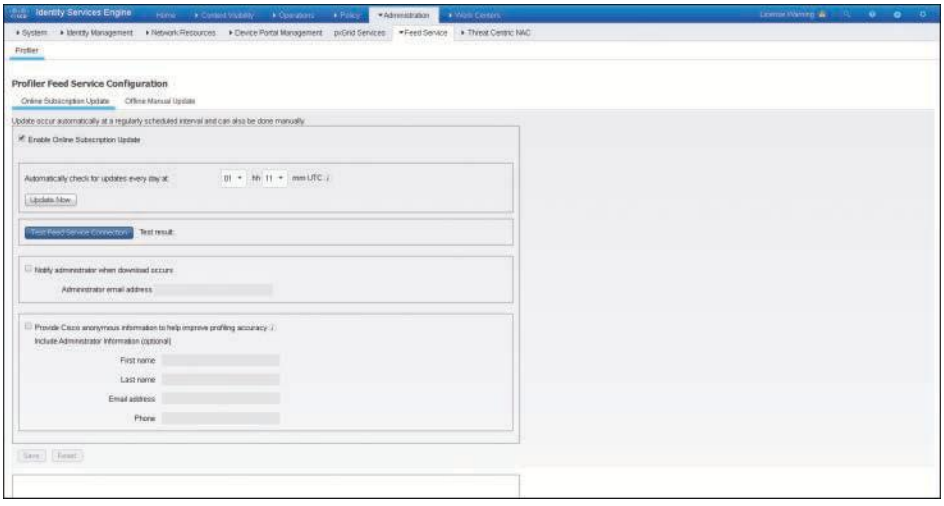


Figure 7-62 *Feed Service Tab*

Threat Centric NAC

The Threat Centric Network Access Control (TC-NAC) service provides the ability to integrate ISE with vulnerability and threat scanners and create authorization policies based on the attributes provided by those third-party systems. The Threat Centric NAC tab (see Figure 7-63) is where you configure the external vulnerability and threat scanners to connect to ISE.

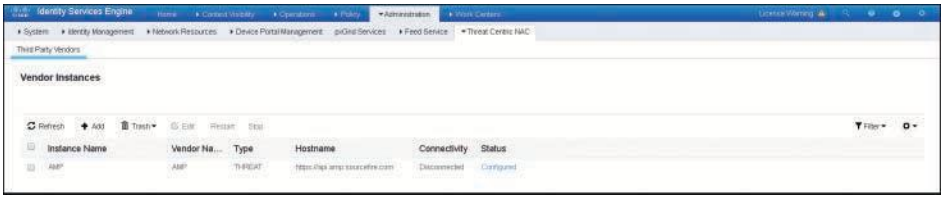


Figure 7-63 *Threat Centric NAC Tab*

Work Centers

Work centers logically group together all the pages related to configuring, monitoring, and reporting for common ISE functions. There are eight work centers to choose from (see Figure 7-64):

- Network Access
- Guest Access
- TrustSec
- BYOD
- Profiler
- Posture

- Device Administration
- PassiveID



Figure 7-64 *Work Centers*

Types of Policies in ISE

ISE is a policy management platform for wired, wireless, and remote access endpoints. As each of these endpoints attempts to access the network, the endpoint, via the NAD, is exposed to a number of different policies. Each of these policies serves a purpose in terms of limiting or appropriately managing the level of access given to the endpoints. Many of these policies take in the programmatic construct referred to as an IF-THEN statement.

Authentication

Authentication refers to the ability to identify an endpoint or the user of an endpoint as it connects to the network. By determining the identity of the endpoint or user, a network administrator can make additional decisions about what level of access should be given to the endpoint or user. We discuss authentication policy in more detail in Chapter 9.

Authorization

Once ISE has determined the identity of a user or an endpoint, ISE can use this identity to determine what the endpoint should be able to access. The granting of—and even the denial of—access is referred to as *authorization*. An authorization policy allows a network administrator to make a final determination based on any number of criteria, such as information provided via authentication, by the endpoint, or by the NAD from which the endpoint is accessing the network. We discuss authorization policy in more detail in Chapter 10.

Profiling

Profiling is the ability to determine the type of device that is accessing the network. It is often accomplished based on endpoint information provided via the RADIUS exchange between ISE and the NAD or via other network protocols that the endpoint leverages as it attempts to join the network (for example, DHCP, HTTP). We discuss profiling in more detail in Chapter 14, "Profiling."

Posture

A posture assessment of an endpoint is an in-depth review of the security policy, including the operating system status and supplementary security software (such as antivirus software, antimalware software, antispyware software, and OS patches), present on the endpoint. If the endpoint's posture does not adhere to the company's security policy, the device can be labeled as noncompliant and remediated as necessary. Posturing is covered in Chapter 18, "Posture Assessment."

Client Provisioning

Client provisioning is a process whereby all necessary credentials and configurations are deployed to an endpoint—allowing the endpoint to more easily (and maybe even automatically) join the network in the future. Provisioning is sometimes referred to as *onboarding*. The onboarding process depends on the endpoint's operating system, the NAD where the endpoint is joining the network, the method of access (wired or wireless), and a number of other criteria. The client provisioning policy provides a decision tree showing what credentials and configurations are deployed to the endpoint in these scenarios. Client provisioning is discussed in greater detail in Chapter 16, "Bring Your Own Device," and Chapter 18, "Posture Assessment."

TrustSec

TrustSec is a security solution in which a Security Group Tag (SGT) is assigned to an endpoint following authentication. SGTs are often designed to subdivide and easily distinguish different user scenarios—for instance, full access versus partial access, employee versus guest, engineering versus marketing. Every packet that has the same SGT is given the same level of access on the network. The NAD inserts the SGT into each packet sent by the endpoint, allowing each TrustSec-enabled upstream device to know what group of users or endpoints was responsible for sending the packet and to enforce the appropriate security policy.

The TrustSec policy in ISE determines what SGT to assign to an endpoint based on the user of the endpoint, the type of endpoint, the method of access (wired versus wireless), and other criteria that can be determined from the authentication process. TrustSec is discussed in detail in Chapter 17, "TrustSec and MACsec."

Exam Preparation Tasks

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 7-4 lists these key topics and the page number on which each is found.

Key Topic

Table 7-4 Key Topics

Key Topic Element	Description	Page
Paragraph	Self-signed certificates for the Admin portal	126
Section	RADIUS Live Log	144
Section	Policy Sets	150
Section	Policy Elements	154
Paragraph	Logging	159

Define Key Term

Define the following key term from this chapter and check your answers in the glossary:

graphical user interface (GUI)

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. Name some of the criteria that can be used for a global endpoint search in ISE.
2. What is important about the Context Visibility dashboards?
3. How does RADIUS Live Log assist in troubleshooting?
4. What are the three primary elements of a policy?
5. What are the different types of policies in ISE?



CHAPTER 8

Initial Configuration of Cisco ISE

This chapter covers the following topics:

- Describe native AD and LDAP
- Describe identity store options
- Network access devices
- ISE endpoint identity configuration
- Troubleshoot using Cisco ISE diagnostic tools
- Identity management

When you buy a Cisco Identity Services Engine appliance, you can take it out of the box and power it on, and the appliance is ready to control access to the network, right? Of course not. Cisco ISE is a very powerful policy engine, and you need to do a fair amount of preparation to teach it how to communicate in your environment. We are not simply talking about the standard requirement of ensuring that a device has a valid IP address so it can communicate on a network. While that is absolutely a requirement with Cisco ISE, there is much more.

Cisco ISE needs to be configured for working with name resolution (DNS) and time synchronization (NTP), and there must be a valid and resolvable fully qualified domain name (FQDN) that any network-attached device can use to resolve the address. Cisco ISE likely needs to have an installed certificate that has been signed by the organization's root certificate authority or a public certificate authority. It also needs to be configured with the IP addresses and RADIUS shared secret keys of all the network devices that will use it. In addition, it probably needs to be joined to an Active Directory infrastructure to use as one of the identity sources.

The blueprint for the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam focuses mainly on the state of Cisco ISE when the administrative graphical user interface (GUI) is running and available. However, some critical steps during the initial installation can make a deployment successful or make it more difficult. This chapter focuses on all these areas.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 8-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 8-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Bootstrapping Cisco ISE	2, 3, 5, 7
Network Devices	6, 8
ISE Identity Stores	1, 9, 10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you incorrectly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which rights and permissions are required for the account used to join Cisco ISE to the Active Directory domain?
   a. Search Active Directory, Remove Workstation from Domain, Change Passwords
   b. Write to Active Directory, Add Workstation to Organizational Unit, Read Properties of Computer Objects
   c. Search Active Directory, Add Workstation to Domain, Set Attributes on the New Machine Account
   d. Write to Active Directory, Add Workstation to Domain, Read Properties of Computer Objects
2. What CLI command lists all the ISE processes and the status of each one?
   a. **show status ise**
   b. **show application status ise**
   c. **show application status**
   d. **show version**
3. What two functions does a certificate fulfill when used with HTTPS and EAP over LAN?
   a. It authenticates the server to the client, and the encryption method is embedded in the transform-set field in the certificate.
   b. It identifies the client to the NAD and is used as the basis for the encrypted transport between the client and the NAD.
   c. It authenticates the server to the client and is used as the basis for the encrypted transport between the client and server.
   d. It authenticates the client to the NAD, and the encryption method is embedded in the transform-set field in the certificate.

4. Which of the following are form factors on which ISE appliances can be installed? (Choose two.)
   a. Container
   b. Virtual machine
   c. Physical appliance
   d. Native cloud
5. In the ISE command-line interface, what command can be entered to show the running application processes?
   a. **configure application ise**
   b. **show application status ise**
   c. **show cpu**
   d. **show processes**
6. What is a valid use of a network device group?
   a. As the condition by which to build different policy sets for the staged deployment of ISE
   b. As the condition to process the incoming authentication protocol type
   c. As the condition by which to determine which ISE policy node to route the authentication request to
   d. As the condition to allow the user to log in and control devices in the assigned network device group
7. Which certificate usage would you select if you wished to use the certificate for the Admin portal?
   a. RADIUS DTLS
   b. Admin
   c. Portal
   d. pxGrid
8. Which of the following top-level network device groups are created by default in ISE? (Choose three.)
   a. All Device Types
   b. Deployment Stage
   c. All Locations
   d. Is IPSEC Device
   e. DTLS Enabled
9. What is the purpose of a certificate authentication profile?
   a. Defines which certificate authority to use for revocation checking via either certificate revocation lists (CRLs) or Online Certificate Status Protocol (OCSP)
   b. Used with MS-CHAPv2 for a client to validate the authentication server
   c. Serves as the identity source for certificate authentications and defines the field of a certificate whose data will be extracted and used as the principal identity for the authorization process
   d. Used with EAP-FAST to allow for faster reauthentication and secure transport without the use of X.509 certificates

10. In order to join ISE to an Active Directory (AD) domain, what permissions must the AD account that will be used to join have? (Choose three.)
    a. Search Active Directory
    b. Domain Admins AD Group
    c. WMI Read-Only
    d. Add Workstation to Domain
    e. Set Attributes on New Machine Account (such as OS type and version)

Foundation Topics

Cisco Identity Services Engine Form Factors

Before Cisco Identity Services Engine can be configured, it must first be installed. The installation of Cisco ISE is beyond the scope of the SISE 300-715 exam and therefore is beyond the scope of this book. However, this chapter touches on some of the installation requirements. There are two form factors for Cisco ISE:

NOTE For instructions on installation and configuration, see the Cisco Identity Services Engine installation guides at https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-guides-list.html.

- **Physical appliance:** The physical appliance is called the Cisco Secure Network Server (SNS). As this book went to press, there were two supported series of ISE physical appliances: SNS-3500 and SNS-3600.
- **Virtual appliance:** It is important to recognize that a virtual appliance is different from a traditional virtual machine. Whereas a virtual machine is meant to leverage the shared physical compute resources of a host server, a virtual appliance must be configured to be an exact replica of the physical appliance it is representing. This means the hardware cannot be oversubscribed, and reservations for the memory and CPU must be configured with values equal to those of the physical SNS appliances.

NOTE The distinction between virtual appliances and virtual machines is important because in a number of escalation cases, the root of the problem has been misconfiguration of the VMware virtual machines. Two misconfigurations are most common:

- There are no resource reservations set for the memory and CPU.
- The VM storage is much slower than the physical drives of the SNS appliance, and delayed writes cause database corruption.

Ensure that the virtual appliances are always configured with the correct reservations and storage in place.

Bootstrapping Cisco ISE

As mentioned earlier in this chapter, Cisco ISE is a very powerful policy engine, and you need to do a fair amount of preparation to teach it how to communicate in your environment.

When you use the virtual appliance or install from a DVD (or an .iso file), you see the installation choices menu shown in Figure 8-1. The SNS appliance comes preinstalled with Cisco ISE and waiting at the setup screen (as shown in Figure 8-2).



Figure 8-1 *Installation Choices*

The underlying operating system that is installed is a very customized version of Linux called ADE-OS. After the operating system is completely installed, the appliance restarts and boots into the setup prompt, as shown in Figure 8-2.

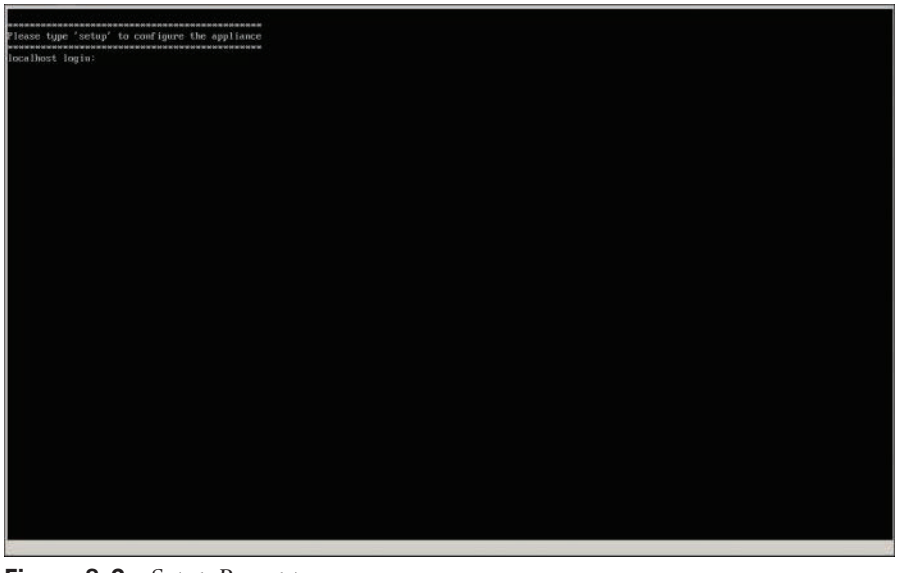


Figure 8-2 *Setup Prompt*

Next, you can enter the **setup** command to launch the command-line setup wizard. Then you can simply follow the prompts to enter all the bootstrapping information that the setup wizard asks for, as shown in Figure 8-3.

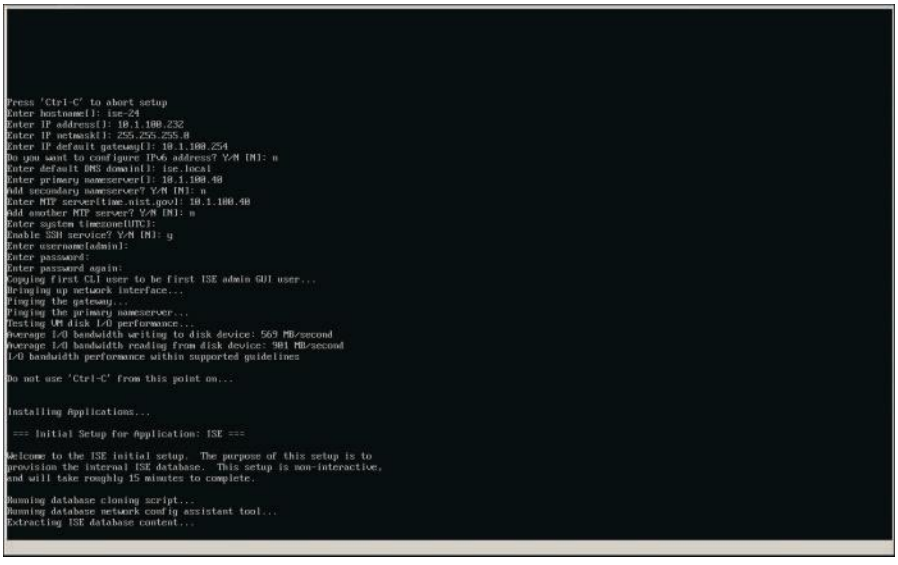


Figure 8-3 *CLI Setup Wizard*

After the settings have been validated, the ISE application is installed. This part of the process takes quite a bit of time, and you can go drink a few cups of coffee and maybe catch up on a streaming television episode while you wait. When the ISE application is fully installed, the device reboots. You should log in to the CLI one more time and ensure that the ISE application has started successfully before attempting to connect to the web interface. The command **show application status ise** lists all the ISE processes and the status of each of them, as shown in Figure 8-4. When the application service is running, ISE is ready for you to connect to the administrative GUI.

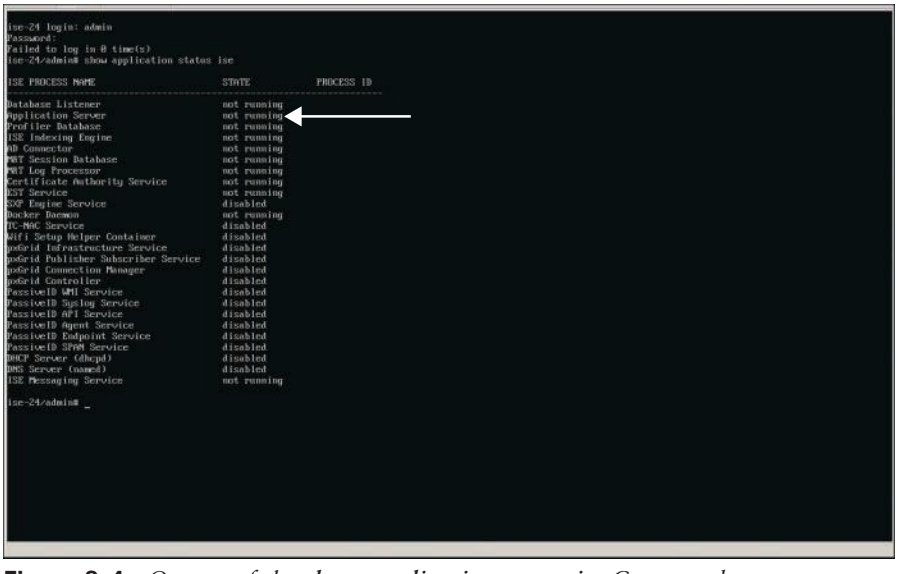


Figure 8-4 *Output of the* ***show application status ise*** *Command*

You connect to the Cisco ISE GUI by using Secure HTTP (HTTPS), with the URL https://<*ise-fqdn*>. When you connect to the ISE GUI for the first time, your web browser is likely to present an error about the certificate being unknown or untrusted, as shown in Figure 8-5. This error message is due to the self-signed certificate that is created during installation.

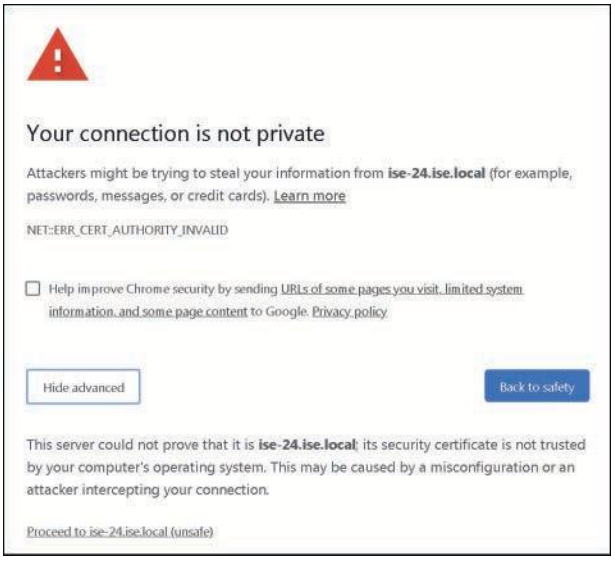


Figure 8-5 *Untrusted Certificate Error*

You can safely accept and trust this certificate to gain access to the administrative GUI login. You then see the screen shown in Figure 8-6, where you can enter the administrator username and password created in the installation setup wizard.

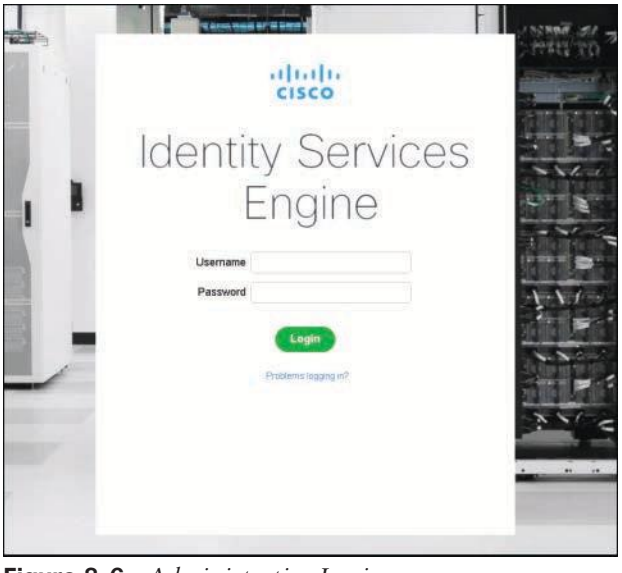


Figure 8-6 *Administrative Login*

Where Are Certificates Used with Cisco Identity Services Engine?

When you connect to a website secured with Secure Sockets Layer (SSL, which indicates HTTPS), the website uses an X.509 certificate to establish the secure communication. The certificate is used for the following functions:

- To prove the "identity" of the website
- To serve as the basis of encryption between the client (web browser) and the website

Figure 8-7 shows the process of securing a website. Numerous web portals are hosted on ISE, including the Admin portal shown in Figure 8-6, and a number of end-user-facing portals for web authentication, endpoint management, and other uses. All the ISE portals require encryption by default. In other words, all portals in ISE use HTTPS, and a certificate on ISE is used to secure the communication.

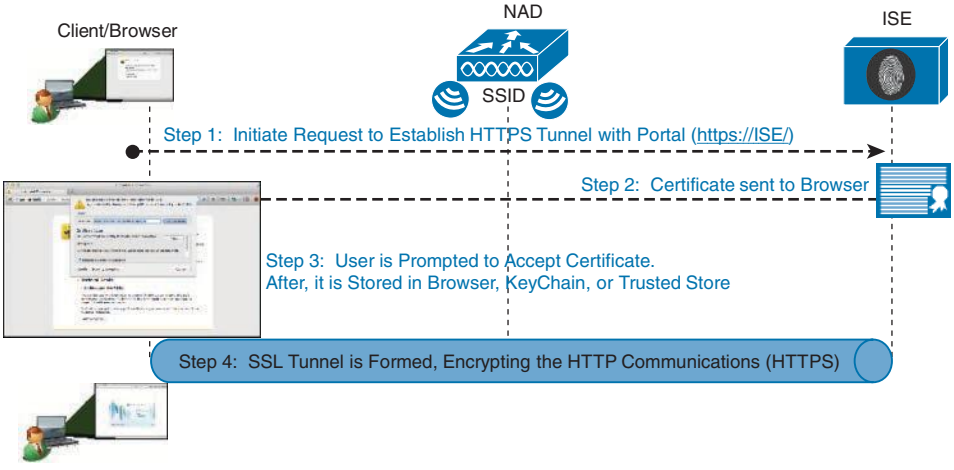


Figure 8-7 *Securing Web Communications with SSL*

Certificates are also used to secure Extensible Authentication Protocol communications, as described in great detail in Chapter 3, "Extensible Authentication Protocol (EAP) over LAN: 802.1X." Certificates are used to form a Transport Layer Security (TLS) tunnel between an endpoint and a RADIUS server, within which EAP authentication will occur, as shown in Figure 8-8. The X.509 certificate has multiple functions with EAP: It is used to identify the RADIUS server (ISE) to the client (supplicant) as well as to form the basis of the encrypted communication.

Self-Signed Certificates

Key Topic

When you initially install ISE, the application generates a self-signed certificate, which means no certificate authority (CA) signed the certificate and vouched that it should indeed be trusted. By default, many browsers and operating systems come with a number of pre-installed CA certificates from the popular CAs. Because a well-known and well-trusted CA was not used to sign a self-signed certificate, clients (browsers and supplicants alike) do not trust the certificate, and they display warning messages similar to the message shown in Figure 8-5.

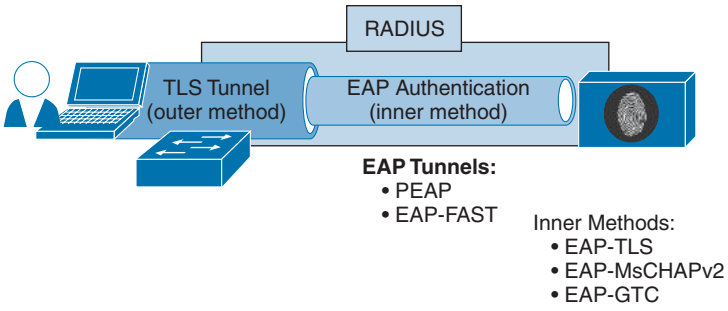


Figure 8-8 *Securing EAP Communications with TLS*

NOTE Before you can join ISE nodes in what is commonly referred to an ISE cube, each ISE node must trust the certificate(s) used by the other ISE nodes.

There are numerous complexities involved with running a production deployment with self-signed certificates. For one thing, each ISE node needs to trust each other ISE node (and this problem will increase exponentially as the deployment grows). For this and many other reasons, it is recommended to use a certificate signed by a trusted certificate authority. This does not mean the certificate must be signed by a public trusted root; a certificate signed by the organization's own internal CA would suffice.

CA-Signed Certificates

Key Topic

There are a few ways to implement a signed certificate on a device. If you have the public and private keys (known as the *key pair*) backed up and stored somewhere secure, those keys can simply be imported into whatever device you want to use those keys. This is not normally what happens when you first install ISE.

When initially installing ISE, a more common way to receive a signed certificate from a CA is to generate a ***certificate signing request (CSR)*** and then submit that CSR to the CA. The CA then signs the public key based on the CSR and returns that public key to you, as shown in Figure 8-9.

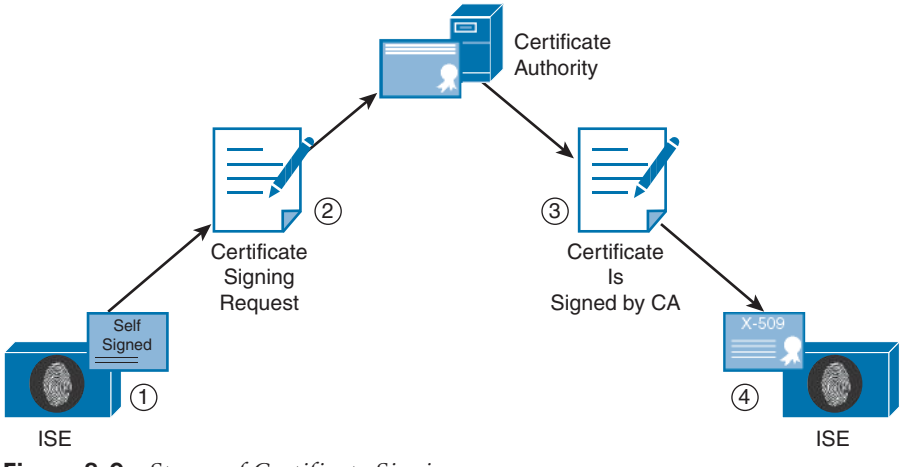


Figure 8-9 *Stages of Certificate Signing*

For ISE, all certificate management is handled in the ISE GUI under Administration > System > Certificates. The initial screen, shown in Figure 8-10, displays the local certificates, including the default self-signed certificate.

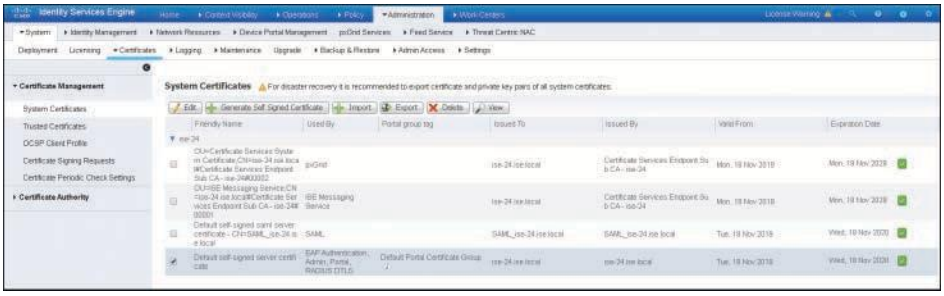


Figure 8-10 *Administration > System > Certificates > System Certificates*

To initiate the generation of a CSR, navigate to Administration > System > Certificates > Certificate Signing Requests and click Generate Certificate Signing Requests (CSR), as shown in Figure 8-11.

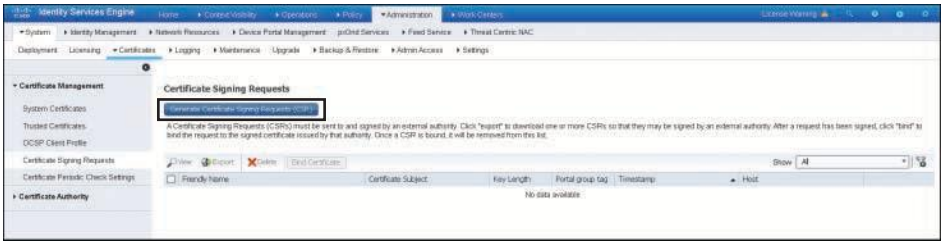


Figure 8-11 *Administration > System > Certificates > Certificate Signing Requests*

In the Generate Certificate Signing Request page (shown in Figure 8-12), fill in the certificate subject information for the CSR. You can customize the following fields for the certificate subject:

- **Common Name (CN):** This is the only mandatory field for the certificate subject. The CN should be a fully qualified domain name (FQDN) that will resolve via DNS to the IP address of the ISE node. The standard is to use the hostname and domain name configured during the initial setup of ISE (for example, ise01.ise.local).

NOTE Pay close attention to capitalization because it matters when you bind the signed certificate.

- **Organization (O):** This should be the organization the certificate is meant to represent—that is, the name of the organization that is implementing ISE.
- **Organizational Unit (OU):** This is the division of the organization that is controlling the device for which the certificate is being used.
- **Locality (L):** This is often used to denote the city where the entity resides.

- **State (ST):** This is the state or province the entity resides in.
- **Country (C):** This is the country that the entity resides in.

The only one of these fields that is required is Common Name (CN).

The ISE GUI also allows you to customize the Subject Alternative Name (SAN) field, which is meant to carry a number of different attributes, each designed to aid in the identification of the certificate. You can add the following customizations for the SAN field:

- **DNS Name:** This is the FQDN of the ISE node. If the certificate will be a wildcard certificate, be sure to specify the wildcard notation (*).
- **IP Address:** This is the IP address of the ISE node associated with the certificate.
- **Uniform Resource Name:** This is the URI associated with the certificate.
- **Directory Name:** This is a string representation of the distinguished name (DN). DNs are defined in RFC 2253.



A client that is presented with a certificate compares the name of the host with the CN field of the certificate subject. If the hostname or FQDN used to reach the host does not match what is in the CN field, a certificate name mismatch occurs. For example, if there is a secure website hosted at https://www.ciscopress.com, and the subject of that certificate has CN set to www.ciscopress.com, the name and the CN match. However, the FQDN www.ciscopress.com resolves to the IP address 159.182.165.54, and that address may also resolve to the FQDN of www.pearsonpublishing.com. If you were to put the URL https://159.182.165.54 or https://www.pearsonpublishing.com into the web browser, there would be a certificate mismatch error.

This brings the conversation to wildcard certificates. A wildcard certificate uses a wildcard notation in the CN (an asterisk and period before the domain name), which allows the certificate to be shared across multiple hosts in an organization. An example of a CN value for a wildcard certificate's subject name would be *woland.com. If you configure a wildcard certificate to use *woland.com, that certificate can be used to secure any host whose DNS name ends in woland.com, such as:

aaa.woland.com

psn.woland.com

mydevices.woland.com

sponsor.woland.com

While a wildcard certificate would solve the issue of a different hostname in the CN, the domain name must still be the same. So, why doesn't ISE just use wildcard certificates for everything? Wildcard certificates are not used for all functions because Microsoft supplicants do not trust wildcard certificates for EAP communications.

As an alternative to wildcard certificates, you can use the Subject Alternative Name (SAN) field, which is designed to allow for these name mismatches. The certificate is "authorized" to represent the name in the CN field, as well as the name(s) in the SAN field. The options

could be DNSName or IPAddress to match a different FQDN or even the IP address of the host. The RFC defining the X.509 certificate fields (RFC 6125) describes the requirement for the FQDN from the CN field to also be listed as a DNSName in the SAN field.

NOTE Bear in mind that a certificate is not limited to a single SAN. If you were to add additional SANs to a certificate, that same certificate could be used across multiple ISE nodes, as long as the certificate had a SAN field with that ISE node's DNS name.

As shown in Figure 8-12, a wildcard value can also be added as a DNSName entry in the SAN field. This is a "best of both worlds" approach for matching multiple values: It allows the use of exactly the same public/private key pair to be installed across all the ISE nodes, which enables a much smoother user experience with mobile devices.

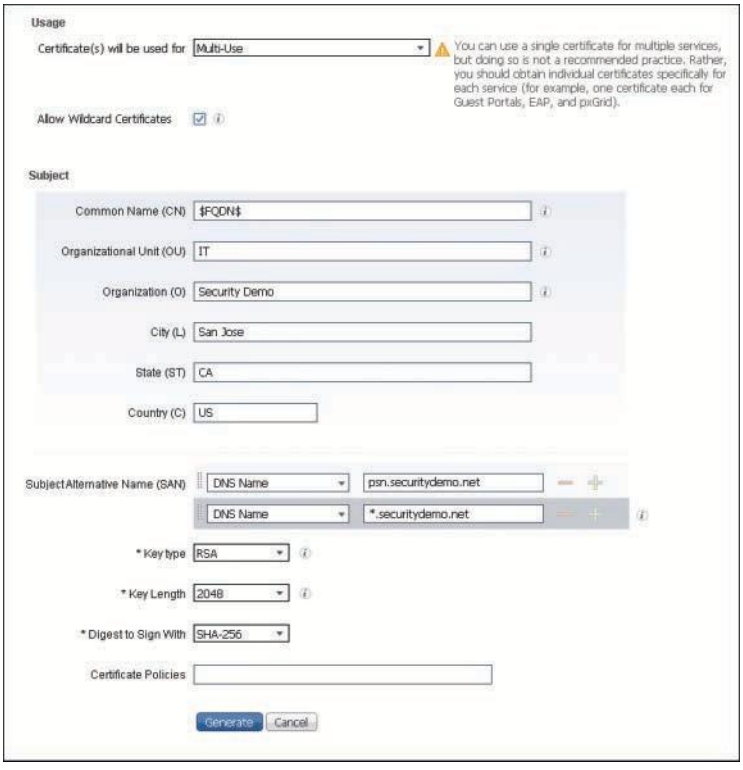


Figure 8-12 *Certificate Signing Request Form*

When you create a public/private key pair, private key is kept secret and never distributed to any third party. The public key and the X.509 information are combined to form the CSR, and this CSR is stored under Administration > System > Certificates > Certificate Signing Requests (see Figure 8-13).

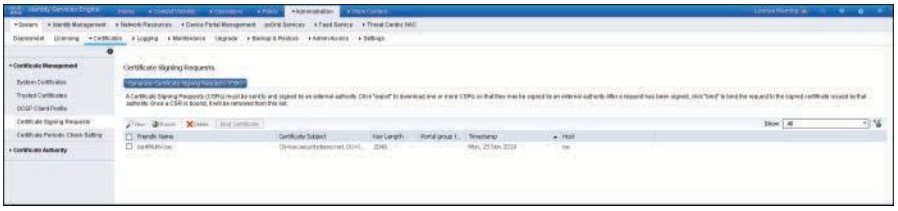


Figure 8-13 *Administration > System > Certificates > Certificate Signing Requests*

If you select the certificate signing request and click Export, ISE prompts you to save the CSR to your local drive, as shown in Figure 8-14.

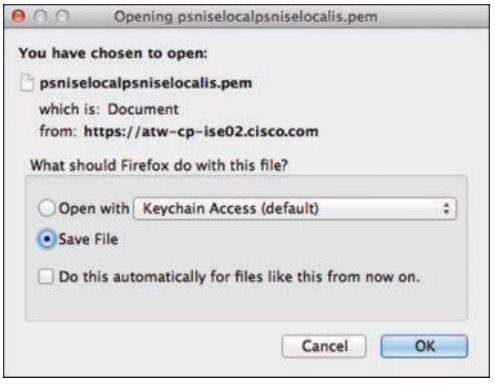


Figure 8-14 *Downloading the CSR*

The CSR is a Base64-encoded plaintext file known as a *PEM file*. It can be opened in a standard text editor, as shown in Figure 8-15. You will need to use all the text in this file, so select all the text, including -----BEGIN CERTIFICATE REQUEST----- and -----END CERTIFICATE REQUEST-----, but do not include any extra spaces (if there are any).

Keep the PEM file open and open a web browser and navigate to the certificate authority (for example, http://atw-cp-ad.ise.local/certsrv, which is the CA that is built in to my Active Directory in my lab). Figure 8-16 shows the page that appears. (If you are not using a CA server, a different process might be required.)



Figure 8-15 *PEM File*

Click Request a Certificate, as shown in Figure 8-16.

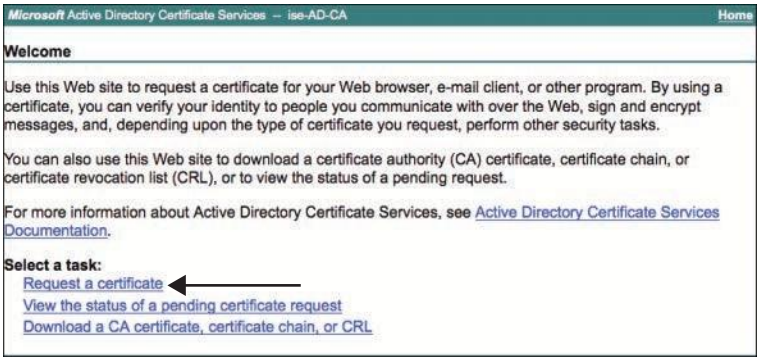


Figure 8-16 *Requesting a Certificate*

Click Advanced Certificate Request, as shown in Figure 8-17.

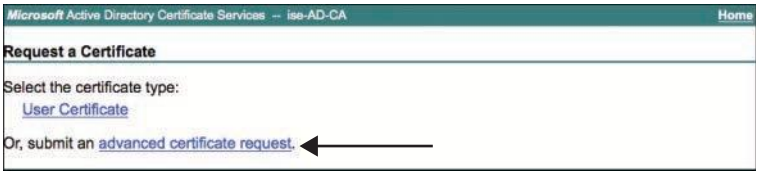


Figure 8-17 *Advanced Certificate Request*

Paste the contents of the PEM file into the form, again ensuring that there are no extra spaces. Choose the Web Server certificate template from the dropdown, as shown in Figure 8-18, and click Submit.

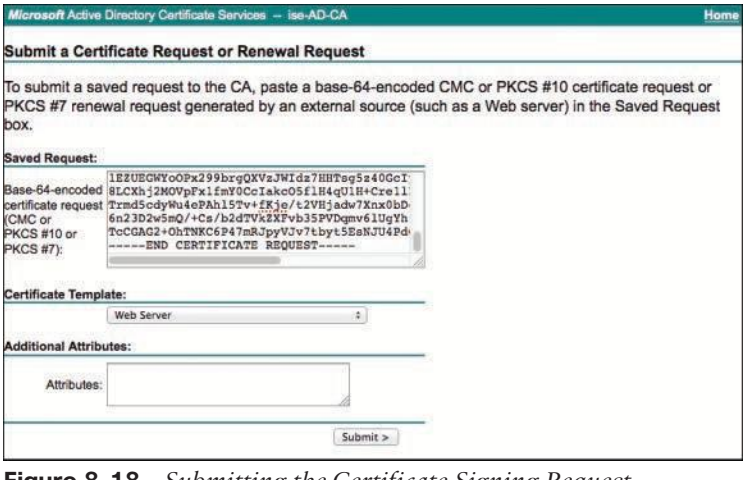


Figure 8-18 *Submitting the Certificate Signing Request*

As shown in Figure 8-19, choose the Base 64 encoded option, which is the PEM format, and click Download Certificate.

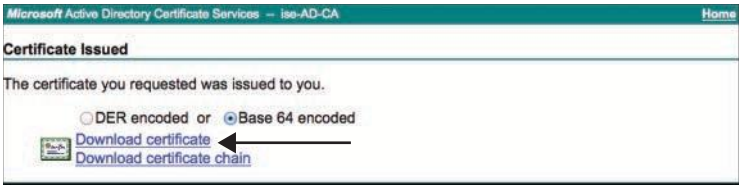


Figure 8-19 *Downloading the Signed Certificate*

> **NOTE** Do not download the certificate chain. Cisco recommends importing all certificates separately instead of using certificate chains (that is, PKCS files). This is due to the difference in behavior of the numerous client supplicants and how those supplicants authenticate the certificate chains.

Navigate back to the main page of the certificate authority (refer to Figure 8-16) and click Download a CA Certificate, Certificate Chain, or CRL. This makes it possible to download the public certificate of the CA. In the window that appears, shown in Figure 8-20, select Base 64 and click Download CA Certificate.



Figure 8-20 *Downloading the CA's Public Certificate*

Now you can import the CA's certificate into the ISE trusted store by navigating to Administration > System > Certificates > Trusted Certificates in the ISE GUI, as shown in Figure 8-21.

Key Topic

The Trusted Certificates page shows the locations of all the certificates that ISE will trust, the functions for which certificates from that CA are trusted, and where the revocation checking is defined.

Click Choose File and select a CA's certificate. Provide a friendly name to help identify the certificate in the list and select Trust for Client Authentication or Secure Syslog, as shown in Figure 8-22. Then click Submit.

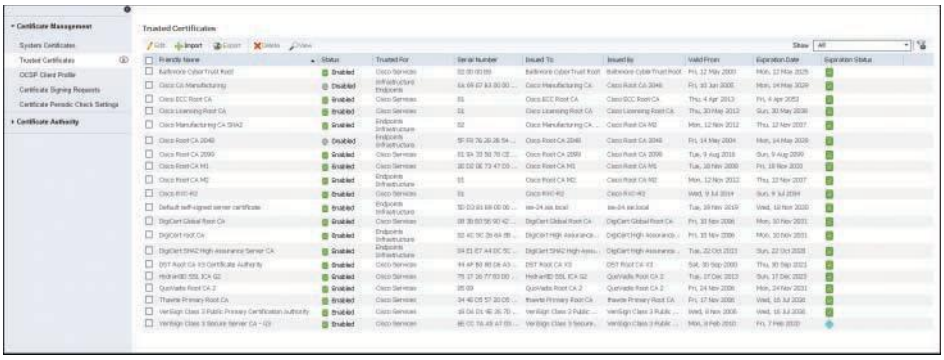


Figure 8-21 *Administration > System > Certificates > Trusted Certificates*

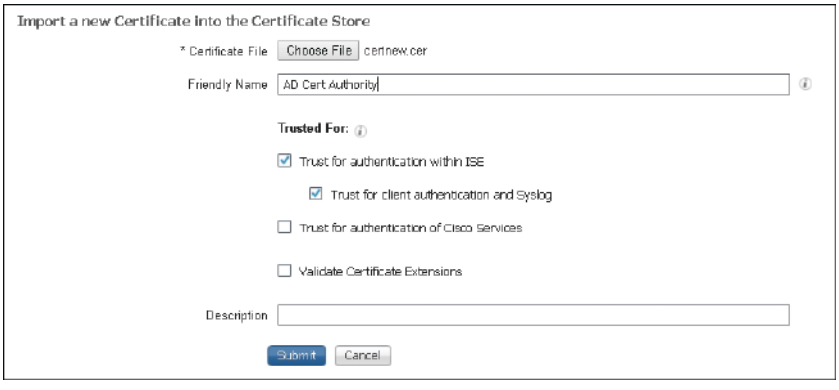


Figure 8-22 *Importing the CA's Certificate to Be Trusted*

Now that ISE trusts certificates signed by the Active Directory CA, it is time to bind the CA Signed Identity Certificate to the CSR so that ISE can use the certificate. Navigate to Administration > System > Certificates > Certificate Signing Requests. Check the box next to the previously created CSR and then click the Bind Certificate button.

Click Browse and select the signed identity certificate. Provide a friendly name and ensure that the Allow Wildcard Certificates checkbox is checked if you are using a wildcard certificate and you provided a CN that is not exactly the same as the ISE node's FQDN. Check the following usage options, shown in Figure 8-23, to establish for which services ISE will use this certificate:

- **Admin:** If this option is checked, the certificate will be used for the ISE Admin portal.
- **EAP Authentication:** If this option is checked, ISE will use the certificate for EAP protocols that use SSL/TLS tunneling.

NOTE While an administrator can only assign one certificate for EAP authentication on ISE, this does not mean that ISE will only accept EAP authentications using only this certificate. ISE will accept certificates presented to it for EAP authentication if ISE trusts the certificate authorities that issued the certificates. This is how ISE can accept EAP authentications from certificates issued both by trusted external CAs and self-signed certificates issued by ISE's own internal CA for BYOD devices.

- **RADIUS DTLS:** The original RADIUS protocol only encrypts the passwords in the RADIUS packets. RADIUS over Datagram Transport Layer Security (DTLS) was developed to further secure RADIUS and encrypt the full packet. The certificate can be used to establish and secure the communication between the ISE server and the network access device (NAD).
- **pxGrid:** If this option is checked, this certificate will be used to establish trust and build a secure connection between ISE and third-party pxGrid clients to share contextual information. If using a certificate issued by your Microsoft CA, one thing to note is that the web server certificate template, by default, only has the server authentication EKU. In order to use the certificate for pxGrid services, the template must be modified to include both server and client authentication in the EKU.
- **Portal:** If this option is checked, this certificate can be used for ISE's user portals.

After selecting the services, click Submit.

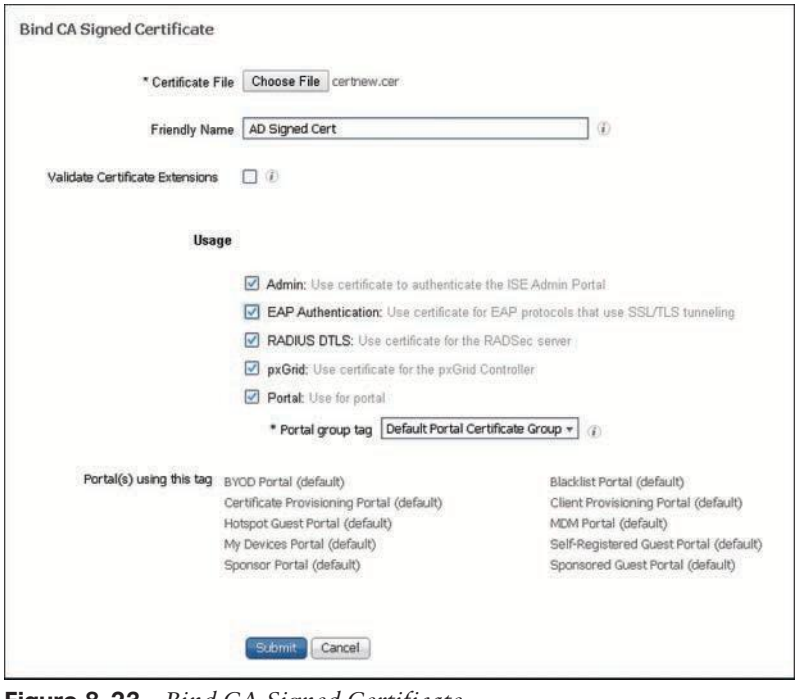


Figure 8-23 *Bind CA-Signed Certificate*

If the Admin usage was chosen, the ISE node needs to restart all services in order to use the newly applied certificate. Otherwise, the certificate should be applied right after you click Submit. When you connect after the services are running, you should also notice that the certificate used to protect and identify the ISE Admin portal is now the signed certificate. Figure 8-24 shows the identity certificate in use.



Figure 8-24 *The Signed Certificate in Use*

One more step should be completed (for various reasons): The public/private key pair should be exported and stored in a very secure and safe location. This step ensures that the same certificates can be used on any additional ISE nodes and also that there is a safe backup in case of an emergency in which the ISE node must be rebuilt.

From the ISE GUI, navigate to Administration > System > Certificates > System Certificates. Select the signed certificate and click Export. As shown in Figure 8-25, choose Export Certificate and Private Key and provide a secure password that can be used to unencrypt the keys for future use.

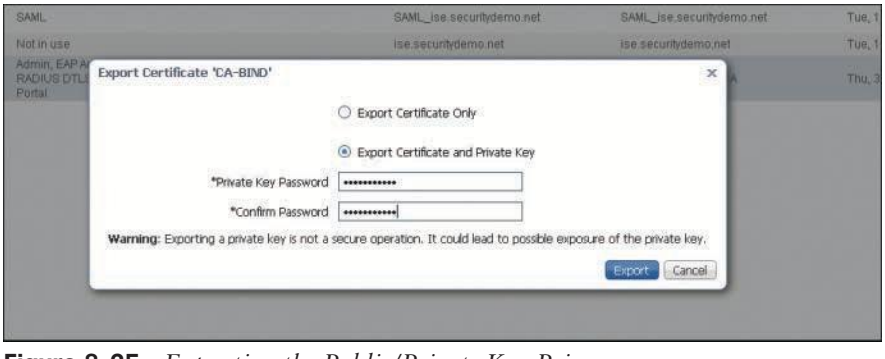


Figure 8-25 *Exporting the Public/Private Key Pair*

Network Devices

Cisco ISE is a policy server. In industry standard terms, ISE would be considered a policy administration point (the admin node) and a policy decision point (the PSN). The policy enforcement point is the ***network access device (NAD)***. The NAD is the device that applies the access control to the endpoint.

The NAD—specifically the NAD's capabilities—is a critical part of any secure network access strategy. The more capable the network access devices (switches, wireless controllers, VPN concentrators), the more flexibility and power the ISE administrator has to accomplish the business goals.

Network Device Groups

Before adding network access devices to ISE, it is highly recommended that you create some logical ***network device groups (NDGs)*** under Administration > Network Resources > Network Device Groups. NDGs can be very powerful tools when used correctly. You can, for example, use NDGs to create policy based on the type of network device, its location, or any other logical grouping you might want.

ISE is prebuilt with three top-level groups named All Device Types, All Locations, and Is IPSEC Device. An ISE administrator can create other top-level groups, as required, such as Line of Business or Deployment (basically anything that would be an organizational unit for the business or deployment structure).

Figure 8-26 shows an example of the NDG hierarchy.

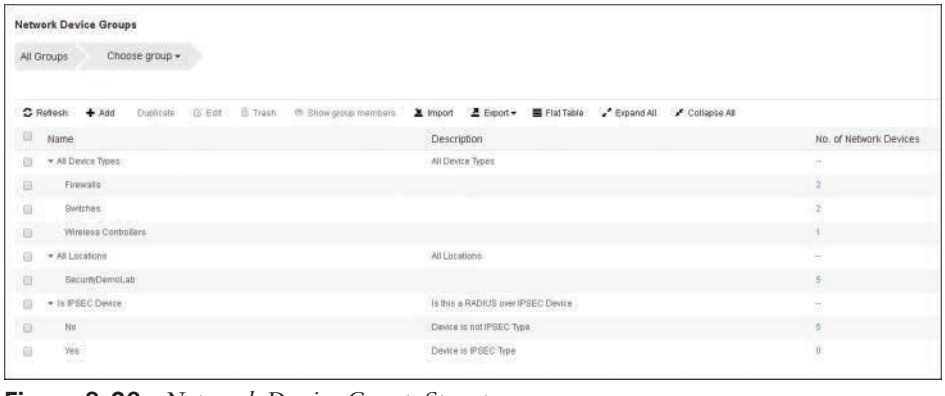


Figure 8-26 *Network Device Group Structure*

Network Access Devices

When the NDG structure is in place, it is time to add network access devices. When you add a NAD (such as a switch, a wireless LAN controller, or a VPN) to ISE, you need to teach ISE about its IP address, configure a shared secret key for the RADIUS communication, and maybe even configure some SNMP strings to pull data from that device (for profiling data). This process teaches ISE how to communicate to the device that will be doing all the enforcement of ISE policy.

To add a NAD to ISE, navigate to Administration > Network Resources > Network Devices and click Add. The network device is added to ISE as an object. That object has numerous

important attributes that configure ISE to uniquely identify each NAD, as well as the shared secret key, TACACS authentication settings, SNMP community strings, and TrustSec configuration. The RADIUS shared secret key must match the configured key on the NAD exactly.

As shown in Figure 8-27, each NAD object must be configured with one or more IP addresses that ISE uses to communicate with the NAD and to identify which NAD is the source of an incoming request. An entire subnet can be associated with a single entry; however, this configuration could make troubleshooting problematic in the future as it would be difficult to isolate which NAD on the configured subnet authenticated the endpoint. The screen shown in Figure 8-27 is also where the NDGs are selected.

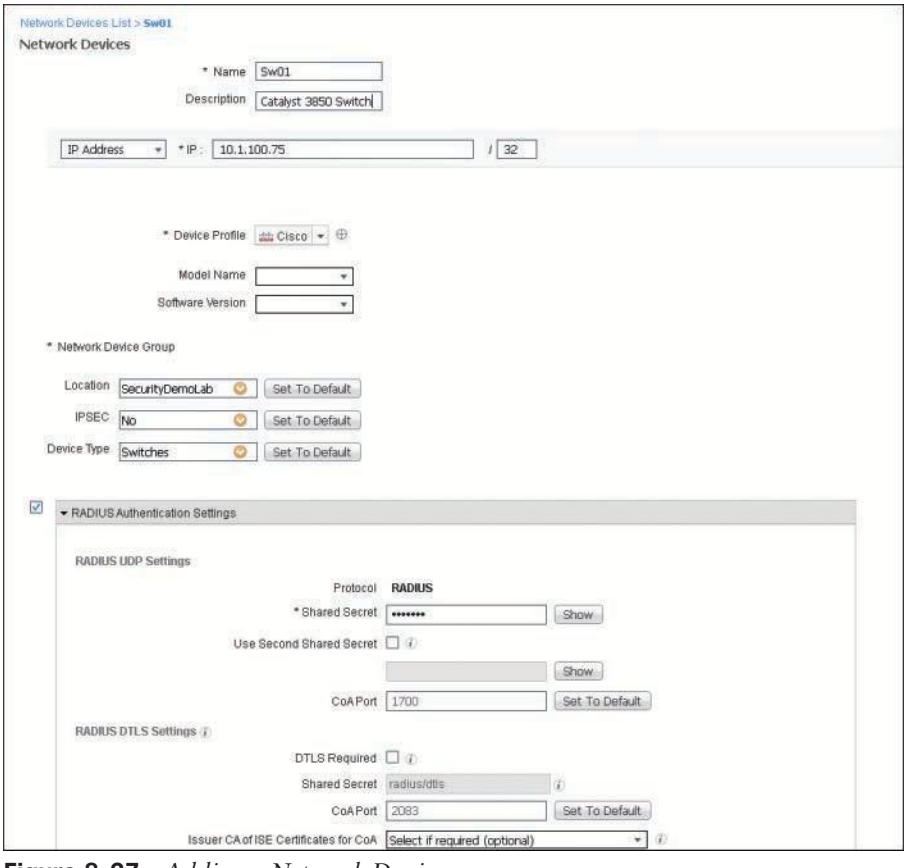


Figure 8-27 *Adding a Network Device*

Each NAD can be assigned to one NDG of each type—in other words, one Location group, one Device Type group, and one of each manually created top-level NDGs, such as Stage.

The section RADIUS Authentication Settings is very important. This is where the RADIUS shared secret is configured. The SNMP Settings section is where you configure the SNMP community for the network device, which ISE will use for querying the NAD for profiling attributes such as CDP and LLDP cache information. You should configure SNMP settings only for devices that do not use the device sensor to send ISE profiling data within RADIUS

accounting packets. You use the Advanced TrustSec Settings section to configure the TrustSec and Network Device Admission Control (NDAC) settings.

NOTE It is possible and common to use a comma-separated values (CSV) file to do bulk imports of NDGs and NADs. However, these functions are beyond the scope of this book.

ISE Identity Stores

Local User Identity Groups

Local users and local user groups are forms of identity stores. Identity stores are covered in much more detail in Chapter 2, "Identity Management," and this section discusses instances in which local users and groups are needed during initial setup for a small proof of concept or some other small deployment.

Local user groups are for users, including guest users, that are created and stored locally in ISE instead of in an external identity store. The groups may be used to define specific levels of access and other attributes that should be applied to a group of local users instead of individuals. Group uses could be used, for example, for local sponsors of guest users, guest user types, and more. These special groups are covered in more detail in Chapter 13, "Guest Services."

As shown in Figure 8-28, you can see a list of prebuilt groups by navigating to Administration > Identity Management > Groups > User Identity Groups. These identity groups can be used directly in an authorization policy, and there is even a special location in the policy where a user or an endpoint identity group can be used.

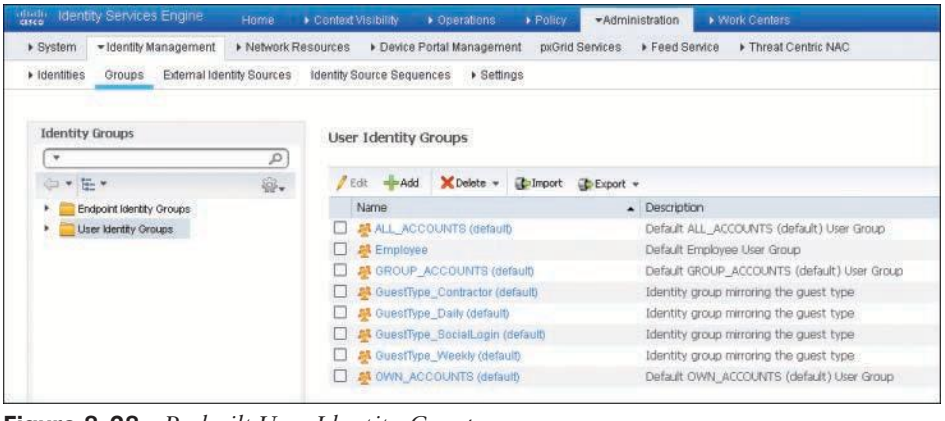


Figure 8-28 *Prebuilt User Identity Groups*

NOTE Unlike endpoint identity groups, user identity groups cannot be nested. This means you cannot create a group that is a member of another group.

Local Endpoint Groups

Key Topic

Endpoints and endpoint identity groups are forms of identity stores. Endpoint identities, covered in detail in Chapter 14, "Profiling," are best used for what many organizations call MAC address management. For example, an endpoint identity group named Corporate Asset could be created to hold the MAC addresses of corporate-owned tablets, laptops, desktops, and other devices.

NOTE An endpoint can exist in one and only one identity group. However, identity groups can be nested. For example, you can create an identity group named iPads that is a member of the Corporate Asset group.

Before ISE Version 1.2, an administrator had to create a matching endpoint identity group for every profile that the administrator wanted to use in a policy. In ISE Version 1.2 and later, the profile attribute itself is used in the authorization policy, so there is no need to waste endpoint identity groups on a profile; now they can be used for pure MAC address management.

Figure 8-29 shows the preconfigured endpoint identity groups located under Administration > Identity Management > Groups > Endpoint Identity Groups.

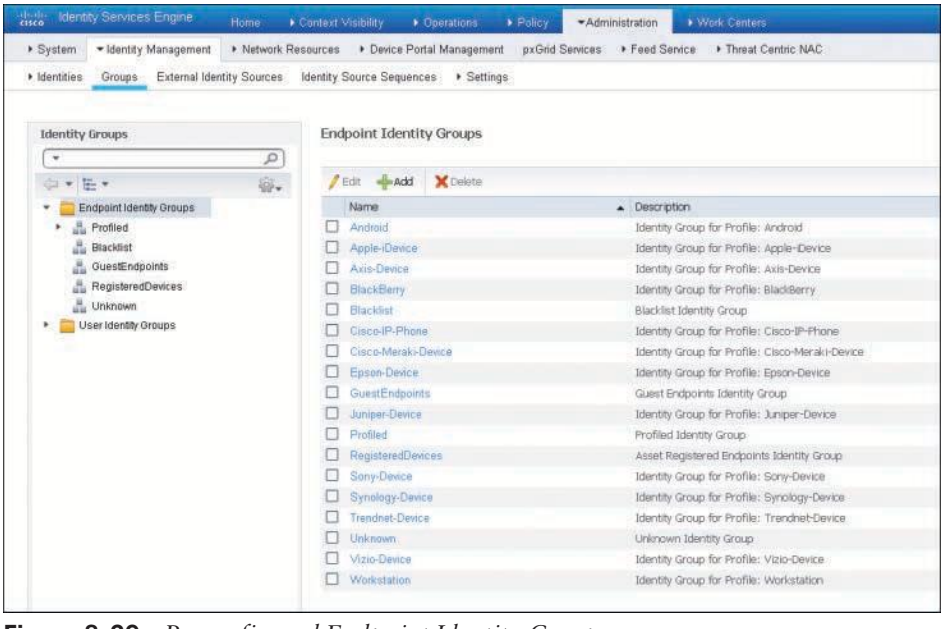


Figure 8-29 *Preconfigured Endpoint Identity Groups*

Local Users

Under Administration > Identity Management > Identities, you can access local users, which are identity stores where user accounts can be created for use with network logins, as sponsor accounts for guest users, or to create test accounts for service probes. While it is possible to use local accounts for an entire deployment of ISE, it is not very common. The majority

of deployments use an external identity store such as Active Directory as the "single source of truth" and leverage locally configured users for administrative purposes.

For more on identity stores, see Chapter 2.

External Identity Stores

While it would be easy to continue to think of a policy server as maintaining all the user accounts locally, in reality, this would not be practical. With large enterprises, there is usually a need for what is commonly called a "single source of truth" for identity as there are often many different identity stores used within an organization.

You are likely to need to configure one or more ***external identity stores*** as you configure ISE. The following sections describe the most common external identity stores and how to configure them. For more on identity source types, see Chapter 2.

Active Directory

Microsoft Active Directory is the single most common external identity store used with ISE deployments. Cisco ISE joins Active Directory (AD), creates a machine account, and uses native AD communication to query the identity store for user accounts and their attributes.

Key Topic

Cisco ISE Version 2.4 supports up to 50 Active Directory domain joins and can query other domains within the AD forest, as long as trust relationships exist. For more on Active Directory as an identity source, see Chapter 2.

Prerequisites for Joining an Active Directory Domain

To join Cisco ISE to an AD domain, you don't have to use the administrator account. Any account can be used, as long as it has the following rights and permissions:

- Search Active Directory (to see if the ISE machine account already exists)
- Add Workstation to Domain (if it does not already exist)
- Set Attributes on the New Machine Account (OS type and version; this is optional)

Certain communication ports must be open between ISE and the domain controllers. Table 8-2 lists the ports that are required for Active Directory communication.

Table 8-2 Active Directory Communication Ports

Protocol	Port (Remote-Local)	Authenticated	Notes
DNS (TCP/UDP)	53	No	May use DNSSEC
MSRPC	445	Yes	
Kerberos (TCP/UDP)	88	Yes (Kerberos)	MS AD/KDC
Kpasswd	464	No	
LDAP (TCP/UDP)	389	Yes	
NTP	123	No	

NOTE It is critical to ensure that good time synchronization is in place before joining an AD domain. Time synchronization issues are the number-one reason ISE fails to join an AD domain.

Joining an Active Directory Domain

The majority of AD join and leave operations occur within the ISE GUI, under Administration > Identity Management > External Identity Sources > Active Directory.

On the Connection tab, shown in Figure 8-30, you need to name the AD domain and click Submit to save it. The domain name should be in a DNS-style format, such as ciscopress.com or ise.local. After you save the domain name, the Connection tab shows that the domain has not been joined yet, as shown in Figure 8-31.

NOTE For the sake of the examples in this book, we'll retain the default identity store name ad1 and reference this value in future chapters as we define the ISE policy.

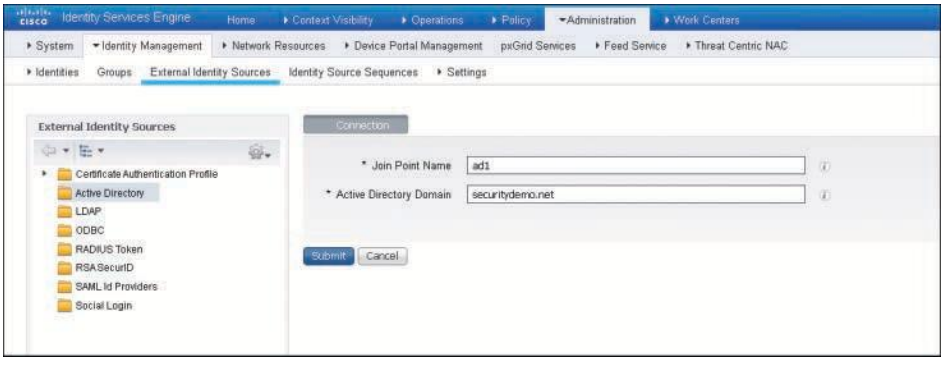


Figure 8-30 *Active Directory > Connection Tab*

Although the Connection tab shown in Figure 8-31 displays only a single ISE node, the screen displays the status of all ISE nodes in the ISE cube related to the AD connection(s). In other words, if the ISE deployment (or ISE cube) had 34 nodes in it, this tab displays all 34 nodes and the current status of each node's connection to the configured Active Directory domain. This is necessary because every ISE node joins the Active Directory domain independently of the other nodes, although the entire configuration is handled from this centralized GUI.

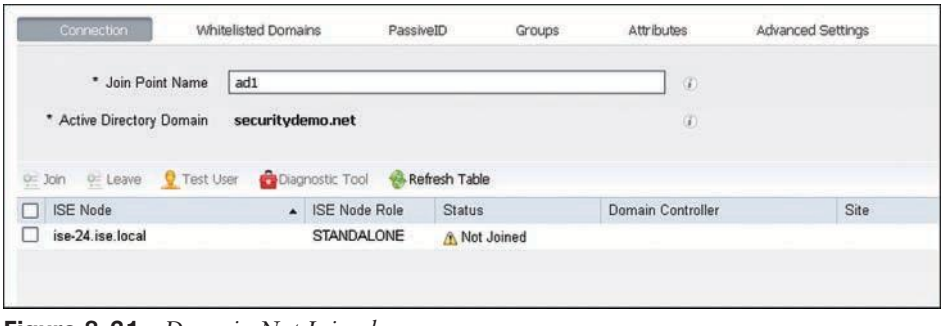


Figure 8-31 *Domain Not Joined*

Choose the ISE node(s) and click the Join button on the screen shown in Figure 8-31; a dialog box pops up for the username and password of an account with the correct rights and permissions to join the machine to AD (see Figure 8-32). Enter the correct username and password and click OK.

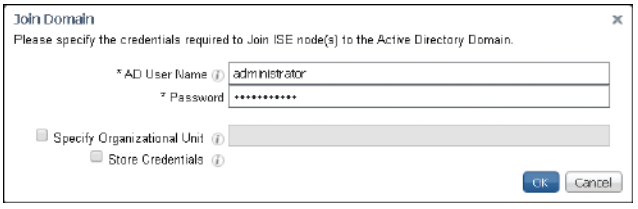


Figure 8-32 *Joining a Domain*

Figure 8-33 shows the Connection tab after the ISE node has been joined to the domain.

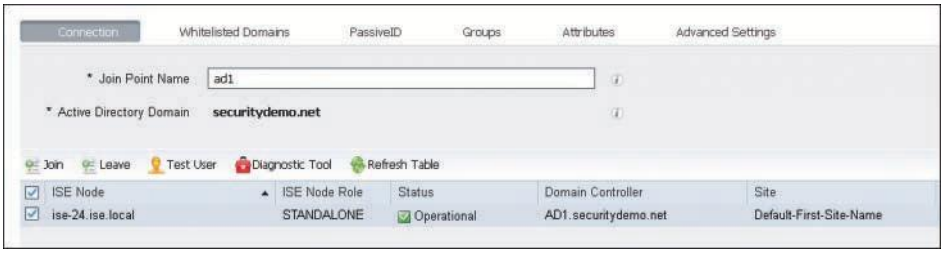


Figure 8-33 *Active Directory > Connection Tab (After the Domain Join)*

Now that ISE has been joined to Active Directory, other items can make the overall deployment easier, such as preselecting the AD groups of interest or preselecting specific account attributes to use in the access policy.

To preselect the groups of interest, click on the Groups tab under Administration > Identity Management > External Identity Sources > Active Directory. Click Add > Select Groups from Directory and the Select Directory Groups popup appears, as shown in Figure 8-34. From this menu, we may retrieve all the Active Directory groups or do a search for only select ones. After checking each of the groups of interest, click OK to accept the choices.

Now, when you are building a policy that will use AD group membership as a condition, you are presented with this list of groups instead of having to sort through the list of all AD groups. Figure 8-35 shows such a selection of groups.

In addition to group membership, other attributes in Active Directory may be of interest for the creation of access policies. You can select these attributes in the Attributes tab by clicking Add > Select Attributes from Directory. To make the administration easier, ISE requires the name of a sample user account. Then, when you click the Retrieve Attributes button, all attributes of that sample user account are displayed. Select the attributes of interest and click OK. Don't forget to save the configuration when you are finished.

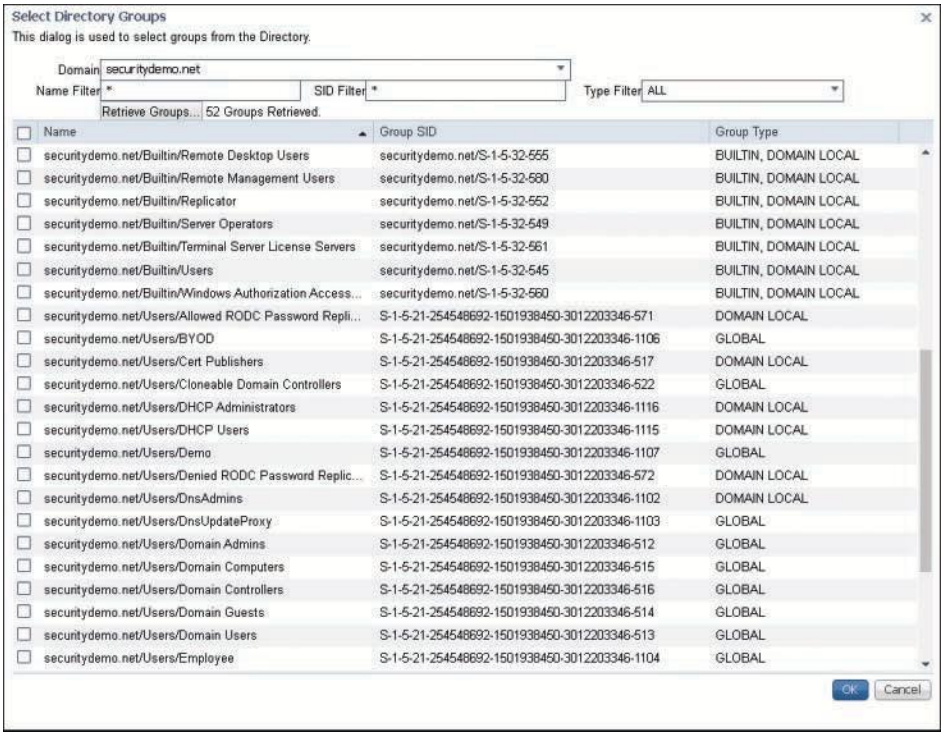


Figure 8-34 *Selecting Directory Groups*

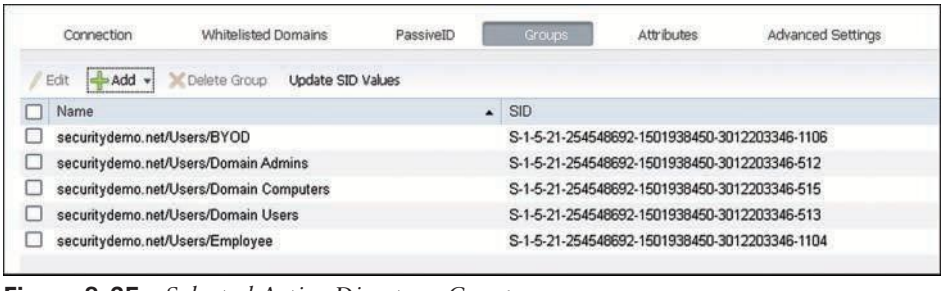


Figure 8-35 *Selected Active Directory Groups*

Figure 8-36 shows attributes being selected, using katmcnam as the sample account, and Figure 8-37 displays the final selection of attributes.

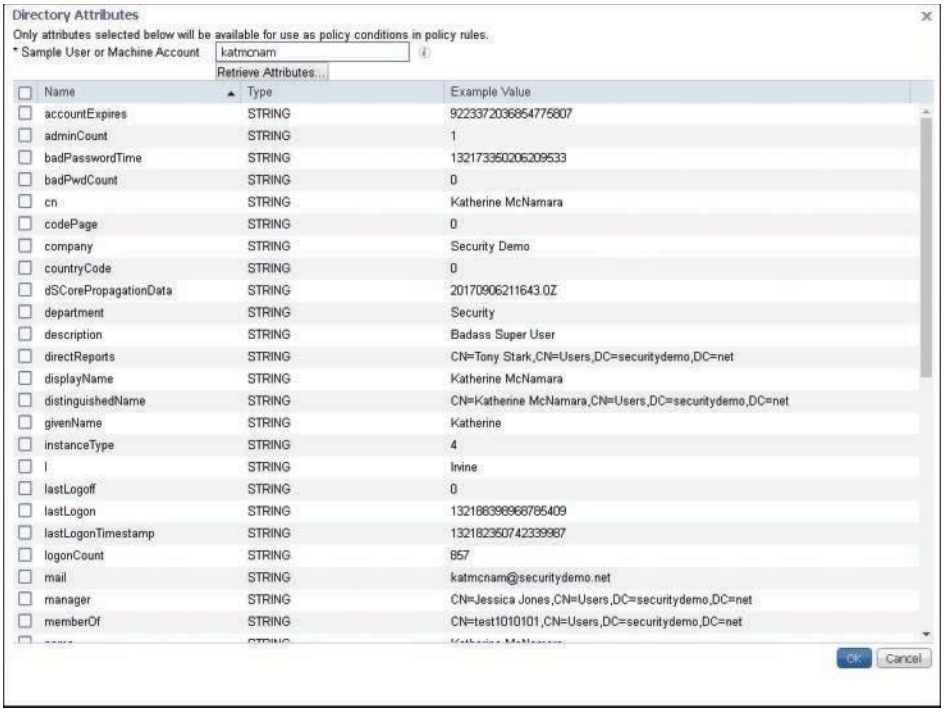


Figure 8-36 *Selected Active Directory Groups*

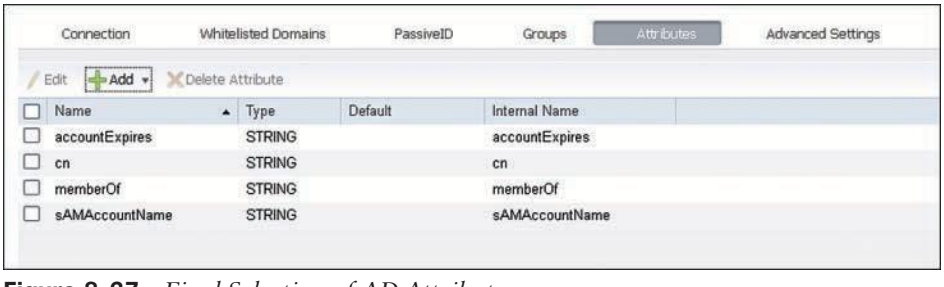


Figure 8-37 *Final Selection of AD Attributes*

The Connection tab provides two tools you can use to test an AD connection: Test User and Diagnostic Tool (see Figure 8-38). ISE administrators as well as Cisco TAC commonly use these tools in determining the root cause of an issue with the Active Directory communication. Figure 8-39 shows the output of the test results using the katmcnam user. Figure 8-40 shows the options for the Active Directory Diagnostic Tool.



Figure 8-38 *User Authentication Tools*

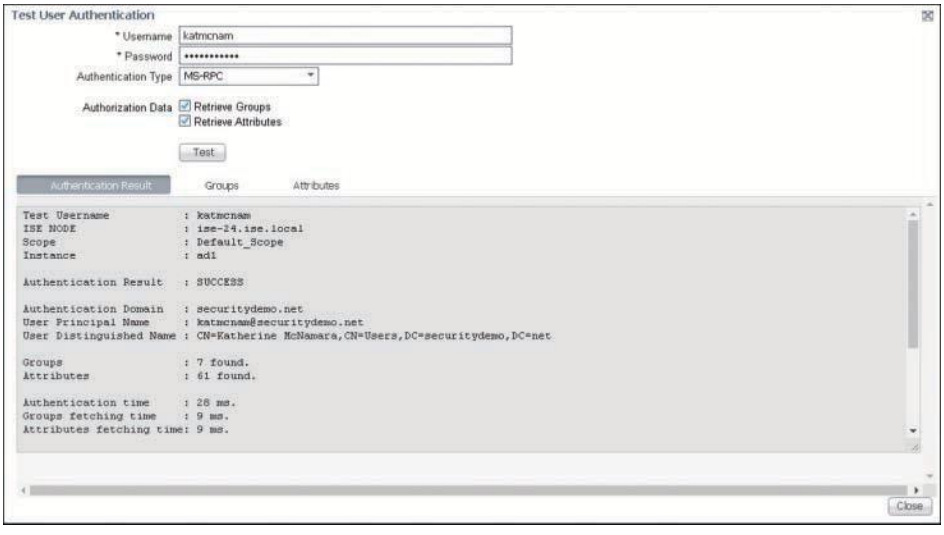


Figure 8-39 *Test User Authentication Tool Output*

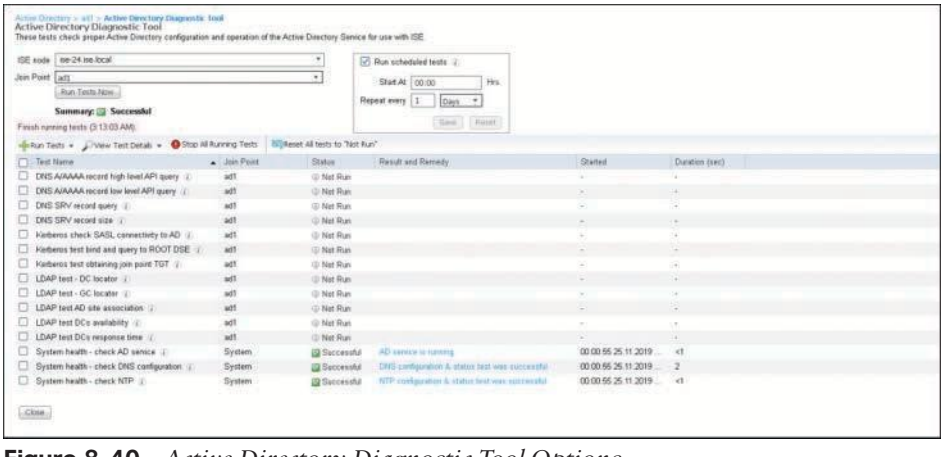


Figure 8-40 *Active Directory Diagnostic Tool Options*

8

NOTE It's important to note that Active Directory may be used for authentication as well as authorization. However, it is very common to use Active Directory for the authorization operation only and to use another source for the authentication operation.

Certificate Authentication Profile (CAP)

As described in the previous section, many different identity stores can be used for authentication in addition to Active Directory. However, the authorization of a user or device may require some information that is available in Active Directory or another ID store. This is the case with digital certificates, for example.

When you break it down, the authentication of a digital certificate is basically nothing more than ISE validating that a certificate was signed by a trusted authority and that the certificate is still valid. However, to perform authorization for a device or user that is being represented by that certificate, ISE can examine values from the certificate and compare those values to attributes that exist in another identity store, such as AD. Therefore, the identity source for certificate-based authentication is a certificate authentication profile (CAP). The CAP tells ISE what value from the certificate should be used as the principal identity for comparison during the authorization phase.

Figure 8-41 shows a very common CAP in which the certificate subject's Common Name (CN) field is being assigned as the principal identity.

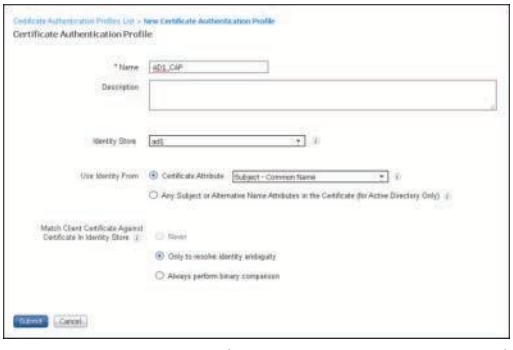


Figure 8-41 *Certificate Authentication Profile*

The second major function of a CAP is to determine if a binary comparison of the certificate should be performed and, if so, what LDAP server to use for that comparison. A binary comparison takes the public certificate used by the user or device attempting access and performs a bit-for-bit comparison to a copy stored elsewhere (usually on the issuing CA). This setting is configured in the CAP by choosing the Perform Binary Certificate Comparison with Certificate Retrieved from LDAP or Active Directory option and selecting which LDAP or AD store will contain the copies of the public certificates.

Identity Source Sequences

To keep the policies of ISE more flexible and easier to work with, you can create an ***identity source sequence (ISS)***. An ISS checks a series of identity sources from top to bottom (as configured), enabling a single authentication rule to check a number of identity sources instead of just one.

Identity source sequences are configured in the ISE GUI under Administration > Identity Management > Identity Source Sequences. With ISE Version 2.4, there are five preexisting ISSs, as shown in Figure 8-42, and more can be created as needed.

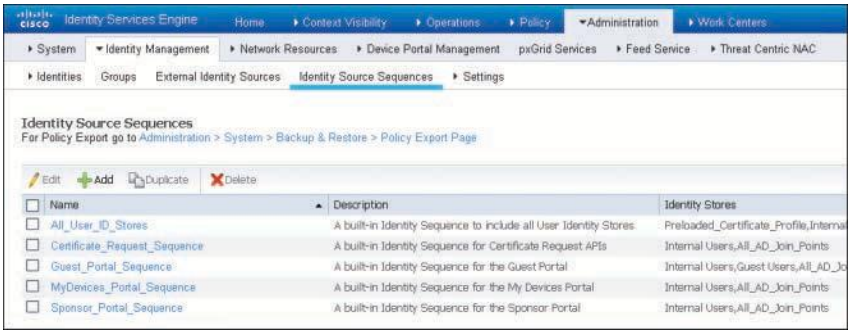


Figure 8-42 *Identity Source Sequences*

To add a new ISS, click Add. Provide a name for the sequence, such as ALL_ID_STORES. As with all other objects created in ISE, you should provide a good description to help document the configuration for any other administrators of the system as well as your future self. Be sure to use a solid description that will help the viewer understand why an ISS was created.

You can select a CAP to include certificate-based authentications, along with the other selected identity stores. Figure 8-43 shows the first part of an ISS that authenticates against a CAP, followed by the internal users, guest users, and finally Active Directory.

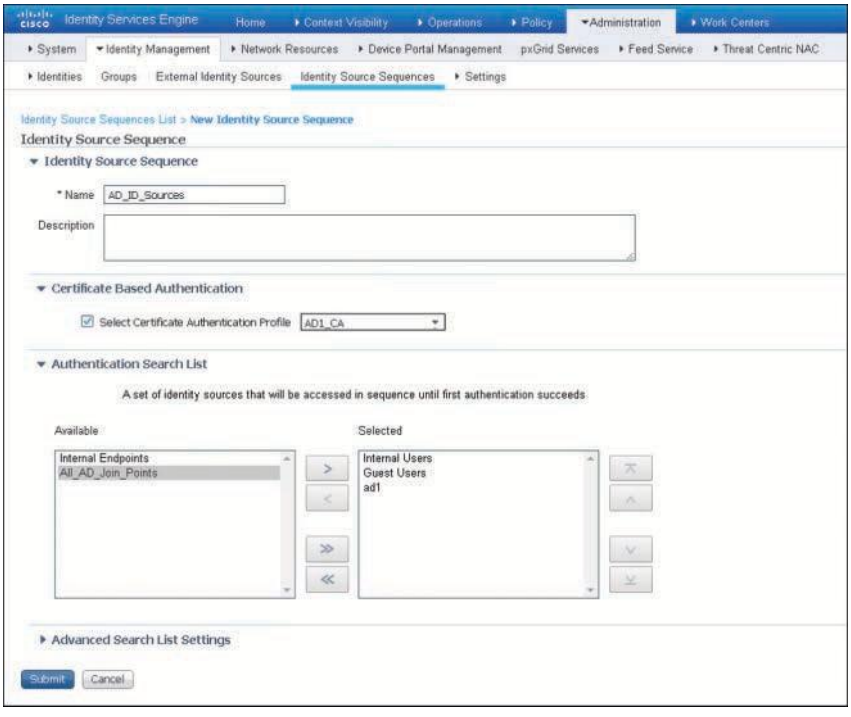


Figure 8-43 *All Identity Stores ISS, Part 1*

Figure 8-44 shows the second part of the ISS, where you configure the action to take if a selected identity store cannot be accessed. The ISS can terminate the lookups and trigger a process error, or it can act as if the user was not found and continue to the next identity store in the sequence.

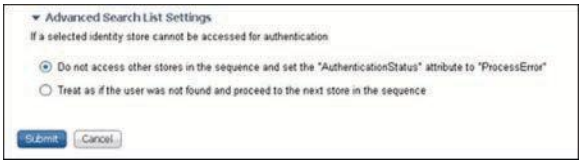


Figure 8-44 *All Identity Stores ISS, Part 2*

This chapter discusses the initial configuration of ISE and the creation of identity sources for use in authentication policies. Chapter 9, "Authentication Policies," details the use of those identity sources in authentication policies.

Exam Preparation Topics

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 8-3 lists these key topics and the page number on which each is found.

Key Topic

Table 8-3 Key Topics

Key Topic Element	Description	Page
Section	Self-signed certificates	181
Section	CA-signed certificates	182
Paragraph	Certificates and the CN field	184
Paragraph	Trusted Certificates page	188
Section	Local endpoint groups	195
Section	Joining Cisco ISE to domains, including inter-forest domains	196

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

network device group (NDG), certificate signing request (CSR), network access device (NAD), external identity store, identity source sequence (ISS)

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. When initially setting up ISE at the command line, what key fields are required to finish installing ISE?
2. What are some of the attributes that a SAN field can include?
3. What are the advantages and disadvantages of using a wildcard certificate?
4. What is the primary use case for network device groups?
5. If an attempt to join an Active Directory domain fails initially, what are the three things to check first when troubleshooting?



CHAPTER 9

Authentication Policies

This chapter covers the following topics:

Describe identity store options (LDAP, AD, PKI, OTP, smart card, local)

Implement wired/wireless 802.1X

EAP types

Implement MAB

Describe the MAB process within an 802.1X framework

ISE authentication/authorization policies

ISE endpoint identity configuration

An *authentication* is simply the validation of a credential. It is a very important part of the process of performing network access control. When thinking about authentication, it can help to relate the topic to something that occurs in real life, such as being pulled over for speeding.

After a highway patrol officer has a driver pull her car over to the side of the road, the officer walks up to the driver's window and asks for her driver's license and proof of insurance. (At least that is what happens in the United States.) The driver hands over these documents for the officer to inspect. The officer then examines the driver's license and determines whether it appears to be real. He sees the correct hologram and watermarks on the license, and the picture on the license looks like the driver. The license hasn't expired. The officer goes back to the squad car and performs a lookup into the Department of Motor Vehicles (DMV) database and finds that the license has not been suspended. All checks have passed. This is a valid ID. The authentication was successful.

Authentication policies have a few goals:

- They drop traffic that isn't allowed and prevent it from taking up any more processing power. (Similarly, the officer would immediately reject a library card, as that is not an allowed form of ID for a driver.)
- They route authentication requests to the correct identity store. (The officer does a lookup at the appropriate DMV, such as the North Carolina DMV or the New York DMV.)
- They validate the identity. (The officer determines whether the license is valid for the driver in question.)
- They pass successful authentications over to the authorization policy. (The officer records the laws the driver has broken, such as exceeding the speed limit.)

When thinking about authentication for network access, it often helps to relate the topic to an example like this one. Doing so can help you understand the difference between authentication and authorization, which we discuss in more detail in this chapter.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 9-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 9-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
The Relationship Between Authentication and Authorization	3
Authentication Policy	3
Understanding Policy Sets	2, 4, 5, 7
Common Authentication Policy Examples	6, 9, 10
More on MAB	1, 8

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is required to perform MAB from a Cisco network device?
   - **a.** The RADIUS packet must have Service-Type set to Login and Called-Station-Id populated with the MAC address of the endpoint.
   - **b.** The RADIUS packet must have Service-Type set to Call-Check and Calling-Station-Id populated with the MAC address of the endpoint.
   - **c.** The RADIUS packet must have Service-Type set to Call-Check and Called-Station-Id with the MAC address of the endpoint.
   - **d.** The RADIUS packet must have Service-Type set to Login and Calling-Station-Id populated with the MAC address of the endpoint.
2. Which EAP type is capable of performing EAP chaining?
   - **a.** PEAP
   - **b.** EAP-FAST
   - **c.** EAP-TLS
   - **d.** EAP-MD5

3. Which of the following are purposes of an authentication policy?
   a. To permit or deny access to the network, based on the incoming authentication request
   b. To apply access control filters such as dACLs or Security Group Tags (SGTs) to the network device to limit traffic
   c. To drop requests using an incorrect authentication method, route authentication requests to the correct identity store, validate the identity, and pass successful authentications over to the authorization policy
   d. To terminate encrypted tunnels for purposes of remote access into the network
4. What option needs to be enabled in the authentication allowed protocols list in order to accept MAB?
   a. PAP/ASCII
   b. Process Host Lookup
   c. LEAP
   d. CHAP
5. Where are reusable conditions saved?
   a. Library
   b. Saved Conditions
   c. Conditions Studio
   d. Smart Conditions
6. What method will work effectively to allow a different identity store to be selected for each EAP type used?
   a. It is not possible to allow a different identity store to be selected for each EAP type used. The first rule to match 802.1X is used, and no further rules can be used.
   b. Create one authentication rule that matches the Service-Type Framed for each of the EAP protocols. Each authentication rule should have one subrule that matches the EapAuthentication setting (such as EAP-TLS or EAP-FAST).
   c. This is possible only for the main EAP types. If there is an inner method of EAP-MS-CHAPv2 with PEAP, it must be sent to the same identity store as the EAP-MS-CHAPv2 inner method of EAP-FAST.
   d. Create a rule for each EAP type under the policy set's authentication policy that points to the appropriate identity store for each rule.
7. What RADIUS attribute is used to match the SSID?
   a. Calling-Station-ID
   b. Source-Wireless-SSID
   c. Framed-Station-ID
   d. Called-Station-ID
8. Which RADIUS attribute contains the MAC address of the endpoint?
   a. Calling-Station-ID
   b. Source-Wireless-SSID
   c. Framed-Station-ID
   d. Called-Station-ID

9. What is the purpose of the continue option of an authentication rule?
   a. The continue option is used to send an authentication down the list of rules in an authentication policy until there is a match.
   b. The continue option sends an authentication to the next subrule within the same authentication rule.
   c. The continue option is used to send an authentication to the authorization policy, even if the authentication was not successful.
   d. The continue option sends an authentication to the selected identity store.
10. In the Conditions Studio, which of the following options can you not do? (Choose three.)
   a. Create and edit the authentication rule result
   b. Create and edit simple and compound conditions
   c. Create and edit library conditions
   d. Select the attribute and attribute value to use

Foundation Topics

The Relationship Between Authentication and Authorization

It is important to understand the differences between authentication and authorization. Many IT professionals, especially those who do not have a security background, confuse these terms. It's important to understand the differences between them so you can understand how the ISE policies are structured. An ISE policy set has two parts: an authentication policy and an authorization policy.

Wireless was the first place in the network where 802.1X really took hold, and today it is still the most prevalent use case of 802.1X authentication. The vast majority of wireless environments provide employees full network access as long as their usernames and passwords are correct (so that authentication is successful).

Key Topic

An authentication, simply put, is the verification of an identity. If you go to a bank and request a withdrawal from an account, the bank teller asks for your identification. You pass your ID to the bank teller, and the teller looks at the ID and goes through a checklist of sorts:

- Is the ID from a recognized authority (for example, a driver's license or a military ID)?
- Does the picture on the ID look like the person in front of the teller's window?

If you are able to hand the teller a valid ID (so that authentication is successful), does that mean you are *entitled* to withdraw money from the bank account in question? No. Successful authentication is only the first step and does not say that the user is entitled to access. The bank teller now needs to determine whether you are authorized to touch the money in that account, so she checks the account and finds that you are allowed to withdraw up to $1,000. This is the process of *authorization*.

Most of the time spent working in a product like Cisco ISE is spent setting up and tuning the authorization policy. The authorization policy is where the bulk of the decisions are made because it is where you specify what access should be granted after a successful authentication occurs.

Authentication Policy

Authentication policies are the first opportunity for Cisco ISE to interact with the RADIUS Access-Request coming from a network access device (NAD). The main goal of an authentication policy is to process the authentication request quickly so it can be dropped if invalid, denied immediately if the credentials are incorrect, or forwarded to be run through the authorization policies if successful. In addition, authentication policies have a few other goals:

- Goal 1: Drop traffic that isn't allowed and prevent it from taking up any more processing power.
- Goal 2: Route authentication requests to the correct identity store (sometimes called a policy information point [PIP]).
- Goal 3: Validate the identity.
- Goal 4: Pass successful authentications over to the authorization policy.

Goal 1: Accept Only Allowed Protocols

By default, ISE allows nearly all supported authentication protocols. However, it is recommended to allow only the protocols that are expected and supported. This serves a few purposes: It reduces the load on the Policy Service nodes (PSNs), uses the authentication protocol to help choose the right identity store, and prevents authentications from vulnerable or obsolete protocols. For example, say that a corporation only wants to support EAP-TLS for its corporate SSID. When an authentication comes in for a device attempting to join the corporate SSID, the allowed protocols could be set to only allow EAP-TLS, so the PSNs would not waste time processing PEAP requests from a device that was not configured with a certificate.

Keep in mind that an organization is best served by having its security policy dictate which authentication protocols meet the security requirements of the organization. Disabling the less secure protocols ensures that any protocol that is more easily compromised cannot be used. Allowing only certain protocols means not only defining which set of protocols should be permitted but also defining the specific tuning of those protocols as well. For example, EAP-FAST may be in the allowed protocol list, but you need to configure options for EAP-FAST, such as whether to allow in-band PAC provisioning or use EAP chaining.

Goal 2: Select the Correct Identity Store

Once an authentication has been accepted, ISE must make an identity store selection decision, which you could think of as an identity routing decision. Based on the attributes of an incoming authentication, ISE must determine which identity store should be used to check the credentials. Obviously, if a certificate is being presented, ISE should not try to validate that certificate against the internal users database that is expecting usernames and passwords.

If your organization has multiple lines of business or acquisitions, it may have more than one Active Directory domain or more than one LDAP store. By using attributes from the authentication request, ISE can determine which domain or LDAP store to check.

Goal 3: Validate the Identity

Once the correct identity store has been identified, ISE must make sure the credentials are valid. In the case of password-based authentications, it considers the following:

- Is the username valid?
- Does the user's password match what is in the directory store?

For certificate-based authentications, it considers the following:

- Has the digital certificate been issued and signed by a trusted certificate authority (CA)?
- Is the certificate expired? (ISE checks both the start and end dates.)
- Has the certificate been revoked?
- Has the client provided proof of possession?
- Does the certificate presented have the correct key usage, critical extensions, and extended key usage (EKU) values present?

Goal 4: Pass the Request to the Authorization Policy

If an authentication fails, the authentication policy can reject the request without wasting CPU cycles comparing the request to the authorization policy. If the request does not match any of the configured authentication rules, the policy might be to send a reject message. However, if a request passes authentication, it is time for the authorization policy to be evaluated.

Understanding Policy Sets

Now that you understand the four main goals of an authentication policy, it will be easier to understand policy sets, which are described in this section.

In the ISE GUI, navigate to Policy > Policy Sets, where you see the default policy set (see Figure 9-1).



Figure 9-1 *Policy Sets*

To expand a policy set in order to view the policies within it, click on the right-facing arrow at the far right of the policy set name. Figure 9-2 shows the expanded default policy set.



Figure 9-2 *Default Policy Set*

In older versions of ISE, there was a single overarching authentication policy and a single overarching authorization policy for an entire deployment. That meant that each policy had a single running list of rules for all the different network access use cases being used. These rules are hierarchical, which means they are evaluated from the top down. If a less restrictive or more restrictive rule was reorganized in that policy, it could affect the entire deployment in adverse ways. This made troubleshooting or making policy changes risky. Thankfully, the creators of ISE helped mitigate this by creating policy sets.

A policy set has its own authentication and authorization policies within it and acts as a logical container for the policies. Each policy set has a top-level rule with conditions that must be matched. If the conditions are matched by a RADIUS request, that request is passed to be evaluated against the authentication and authorization policy within that policy set. For example, if you create a policy set with a top-level rule that matches against any RADIUS request with Called-Station-ID set to CiscoPress-SSID, this rule ensures that this policy set will only be evaluated against RADIUS requests for that SSID. It also eases the administrative burden by localizing the policies to segmented logical containers. Any future changes to these authentication and authorization policies will only affect RADIUS requests from that SSID, reducing the risk of a misconfiguration affecting other RADIUS requests.

The top-level rule for a policy set is located above the authentication policy:

```
IF conditions THEN ALLOW PROTOCOLS IN LIST AllowedProtocolList

AND EVALUATE POLICY SET
```

Earlier in this chapter, you saw that one of the goals of an authentication policy is to specify which allowed protocols can be used. It can be confusing to see the allowed protocol list above the authentication policy. In versions of ISE before policy sets were introduced, every rule in an overarching authentication policy needed to have an allowed protocol list configured. With policy sets, the allowed protocol list is declared in the top-level rule for a specific policy set. After the allowed protocol list is defined, you can specify an authentication protocol in the conditions of the authentication rules and also specify how to handle it. (You will see how to configure this later in this chapter.)

In the policy set pictured in Figure 9-2, there are no specific conditions defined in the top-level rule except for the allowed protocol list. By default, the Default Network Access list is populated, but you can change this to a more restrictive allowed protocol list.

Just like the rules in a policy, policy sets are hierarchical. If two policy sets have top-level rules with conditions that a RADIUS request could match, the policy set that appears higher in the list is chosen.

Allowed Protocols

After conditions are matched in a top-level rule, the rule then dictates what authentication protocols are permitted for that policy set. The Default Network Access list of allowed protocols has almost every supported authentication protocol selected. You can also create a more restrictive allowed protocols list and use a different one in each policy set if you wish to do so.

To examine the default allowed protocols, take the following steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Elements > Results > Authentication > Allowed Protocols.**

Step 2. Select **Default Network Access.**

Key Topic

In Figure 9-3, the long list of supported protocols and their options is so extensive that it extends past the viewable area in the screenshot. This default list is very inclusive, with the intention of making deployments work easily for customers, but security best practice is to limit this list to only the protocols needed for the particular policy set. Be sure to only allow the protocols that are consistent with your corporate security policy and ensure that the most secure protocol is chosen for each particular application.

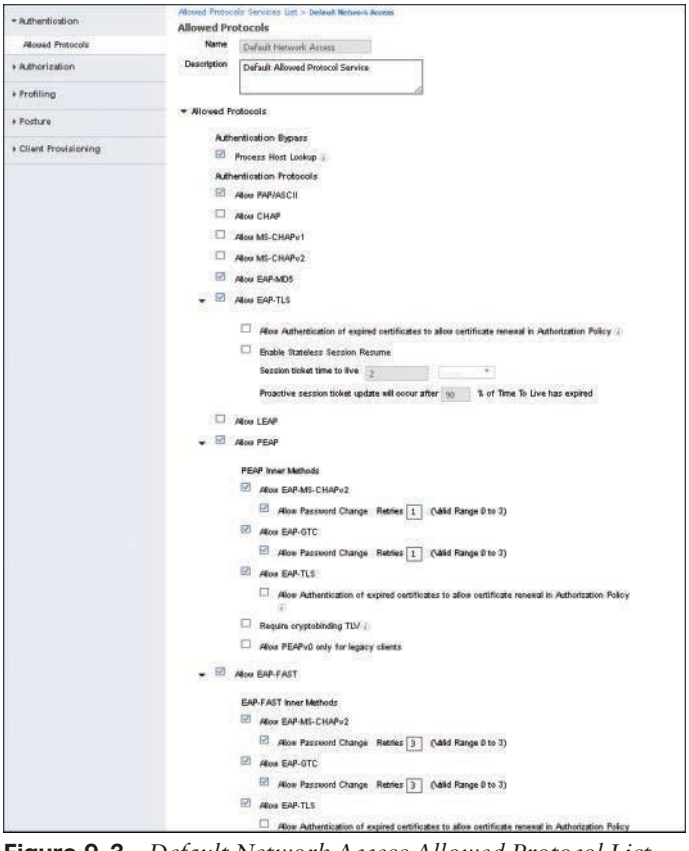


Figure 9-3 *Default Network Access Allowed Protocol List*

If you understand the most common authentication protocols and their uses, you will be able to create a more specific list of allowed protocols for your deployment. The following list looks at the default allowed protocol list from top down (refer to Figure 9-3):

- **Process Host Lookup:** This option is used for MAC Authentication Bypass (MAB). When this option is selected, ISE takes the username and password as the MAC address and checks in the internal database.
- **PAP:** With Password Authentication Protocol (PAP), the username is sent in plaintext, and the password is optionally encrypted.
- **CHAP:** With Challenge Handshake Authentication Protocol (CHAP), the username and password are encrypted using a challenge sent from the server. CHAP is not often used with network access.
- **Extensible Authentication Protocol (EAP) types:** EAP is an authentication framework that provides for the transport and use of identity credentials. EAP encapsulates the usernames, passwords, and certificates that a client is sending for purposes of authentication. There are many different EAP types, each with benefits and downsides. (As an interesting side-note, 802.1X defines EAP over LAN.)
  - **EAP-MD5:** EAP-MD5 uses a Message Digest algorithm to hide the credentials in a hash. The hash is sent to the server, where it is compared to a local hash to see if the credentials are accurate. Unfortunately, EAP-MD5 does not have a mechanism for mutual authentication. This means the server is validating the client, but the client does not authenticate the server (that is, does not check to see if it should trust the server). EAP-MD5 is common on some IP phones, and it is also possible for some switches to send MAB requests within EAP-MD5.
  - **EAP-TLS:** This EAP type uses TLS (Transport Layer Security) to provide secure identity transactions. TLS is very similar to SSL and the way encryption is formed between your web browser and a secure website. EAP-TLS is considered universally supported across many different operating systems and mobile devices. EAP-TLS uses X.509 certificates and provides the ability to support mutual authentication, where the client must trust the server's certificate and vice versa. It is considered among the most secure EAP types because password capture is not an option; the endpoint must have the private key. EAP-TLS is considered the gold standard of EAP types for network access.
- **Tunneled EAP types:** The previous EAP types transmit credentials immediately. These following EAP types form encrypted secure tunnels first and then transmit the credentials within the tunnel (see Figure 9-4):

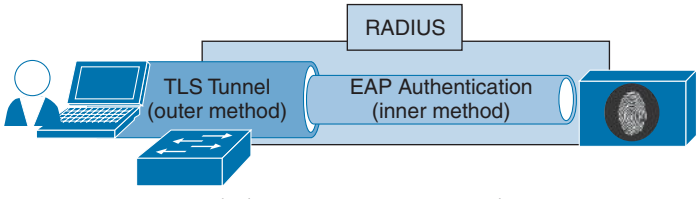


Figure 9-4 *Tunneled EAP Types (PEAP and FAST)*

- **PEAP:** Protected EAP (PEAP), originally proposed by Microsoft, is an EAP tunnel type that has quickly become the most popular and widely deployed EAP method in the world. PEAP forms a potentially encrypted TLS tunnel between the client and server using the X.509 certificate on the server in much the same way an SSL tunnel is established between a web browser and a secure website. After the tunnel has been formed, PEAP uses another EAP type as an "inner method" to authenticate the client using EAP within the outer tunnel. PEAP supports the following inner methods:
  - **EAP-MS-CHAPv2:** Using this inner method, the client's credentials are sent to the server encrypted within an MS-CHAPv2 session. This is the most common inner method as it allows for simple transmission of username and password or even computer name and computer password to the RADIUS server, which authenticates them to Active Directory.
  - **EAP-GTC:** EAP-Generic Token Card (GTC) is an inner method created by Cisco as an alternative to MS-CHAPv2 that allows generic authentications to virtually any identity store, including one-time-password (OTP) token servers, LDAP, Novell eDirectory, and more.
  - **EAP-TLS:** PEAP is capable of using EAP-TLS as an inner method, although this method is rarely used and not widely known. EAP-TLS adds an additional layer of security by tunneling the EAP authentication.
- **EAP-FAST:** Flexible Authentication via Secure Tunneling (FAST) is very similar to PEAP. Cisco created FAST as an alternative to PEAP that allows for faster reauthentication and supports faster wireless roaming. Just like PEAP, FAST forms a TLS outer tunnel and then transmits the client credentials within that TLS tunnel. Where FAST differs from the PEAP is in the ability to use Protected Access Credentials (PACs). A PAC can be thought of as a secure cookie that is stored locally on the host as proof of a successful authentication. EAP-FAST is supported natively on Apple OS X 10.4.8 and newer. The Windows native supplicant does not support EAP-FAST. Instead, the AnyConnect software client with the Network Access Manager (NAM) can be installed on endpoints running Windows Vista and later. NAM acts as the supplicant for the endpoint, which allows Windows to support a range of new EAP options. EAP-FAST supports the following inner methods:
  - **EAP-MS-CHAPv2:** With this inner method, the client's credentials are sent to the server encrypted within an MS-CHAPv2 session. This is the most common inner method, as it allows for simple transmission of username and password or even computer name and computer password to the RADIUS server, which authenticates them to Active Directory.
  - **EAP-GTC:** EAP-Generic Token Card (GTC) is an inner method created by Cisco as an alternative to MS-CHAPv2 that allows generic authentications to virtually any identity store, including one-time-password (OTP) token servers, LDAP, Novell eDirectory, and more.
  - **EAP-TLS:** EAP-FAST is capable of using EAP-TLS as an inner method. This became quite popular with EAP chaining.

- **EAP chaining with EAP-FASTv2:** This EAP method can only be used with the Windows AnyConnect NAM software. As an enhancement to EAP-FAST, a differentiation was made to have both a user credential and a machine credential authenticated. These credentials can be certificates or username and password. This provides ISE administrators the ability to validate that a machine is an authenticated domain computer and also to validate the credentials of the user. EAP chaining makes it possible, for the first time in 802.1X history, for multiple credentials to be authenticated within a single EAP transaction.
- **EAP-TTLS:** Tunneled TLS EAP (EAP-TTLS) is similar to PEAP in the way it works and the features it provides. It encapsulates a TLS session by using a handshake phase and a data phase. During the handshake, the server authenticates itself to the client or they both mutually authenticate using standard TLS. The server can then use that established tunnel to authenticate the client. Versions of Windows prior to Windows 8 do not support EAP-TTLS natively. However, Microsoft has included native supplicant support in Windows 8 and later. The EAP inner methods that are accepted with EAP-TTLS include PAP, CHAP, MS-CHAPv1, MS-CHAPv2, and EAP-MD5.
- **TEAP:** This EAP method is a new addition to Microsoft Windows 10. It enables secure communication between a peer and server using TLS to establish a mutually authenticated tunnel. While TEAP is not widely adopted yet, it does introduce the possibility of EAP-Chaining natively in the Windows supplicant as of Windows 10 build 2004 (May 2020) and ISE 2.7 with patch 1. While configuring EAP-Chaining with TEAP is outside of the scope of the SISE exam given its relatively new addition to Windows 10, it is an option to keep an eye on in the future.

Understanding Authentication Policies

To better understand authentication policies, this section examines the authentication policy in the default policy set. Figure 9-5 shows the expanded authentication policy.

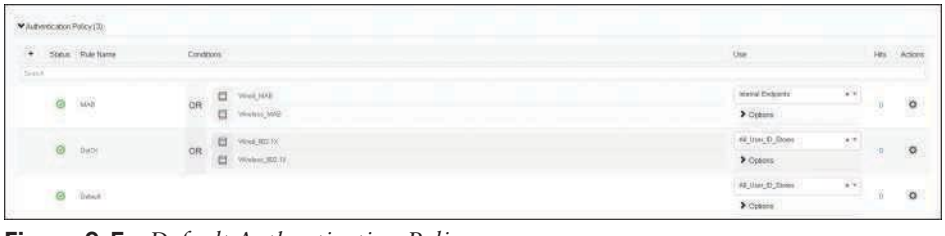


Figure 9-5 *Default Authentication Policy*

Basic authentication policy rules are logically organized in this manner:

```
IF conditions THEN

CHECK THE IDENTITY STORE IN LIST IdentityStore
```

Rules are processed in a top-down, first-match order—just as in a firewall policy. If the conditions do not match, the authentication is compared to the next rule in the policy.

As shown in Figure 9-5, the default authentication policy is preconfigured with a default rule for MAC Authentication Bypass (MAB). MAB is used for a number of things, such as allowing non-802.1X endpoints onto the network, guest access, BYOD, and more, as discussed in later chapters. For now, we are going to use this rule to dig into authentication rules and how they work. If you have a live ISE system, it might help to follow along on it with the text.

Figure 9-6 demonstrates a MAB rule in flowchart format.

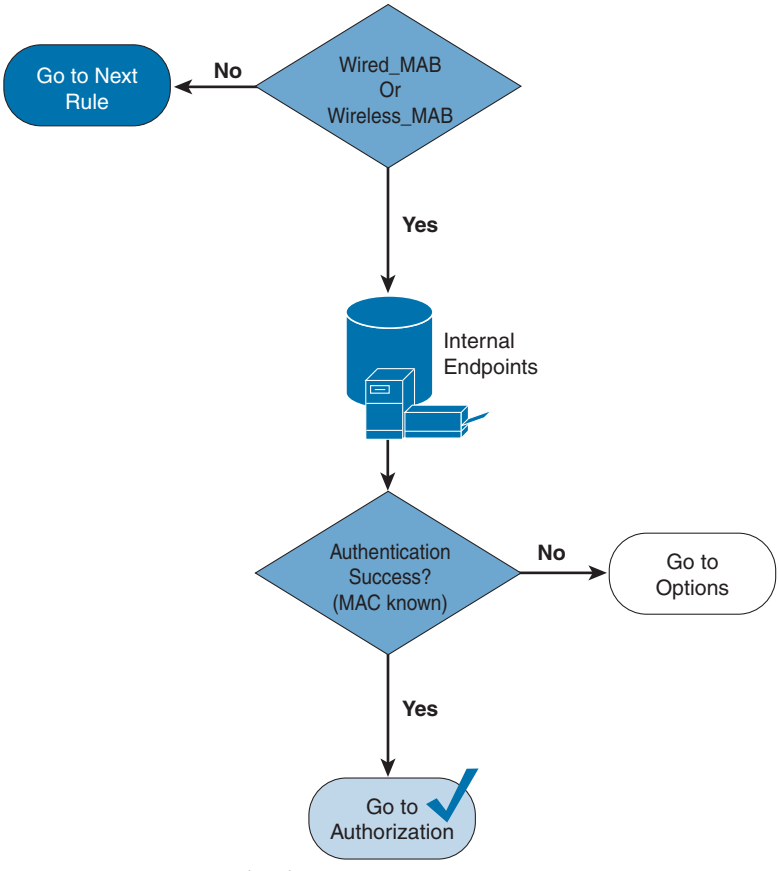


Figure 9-6 *MAB Rule Flowchart*

Conditions

The conditions of the MAB rule shown in Figure 9-6 state that if the authentication request is Wired_MAB or Wireless_MAB, it will match this rule. Wired_MAB and Wireless_MAB are predefined smart conditions that contain one or more specific conditions and can be reused across different policies and rules.

To see more details, you can click on a condition displayed in the rule. When you do, the Conditions Studio opens (see Figure 9-7). The Conditions Studio, which was introduced in ISE Version 2.3, gives an ISE administrator a workspace to configure and edit simple and complex policies and save them for reuse in the library.

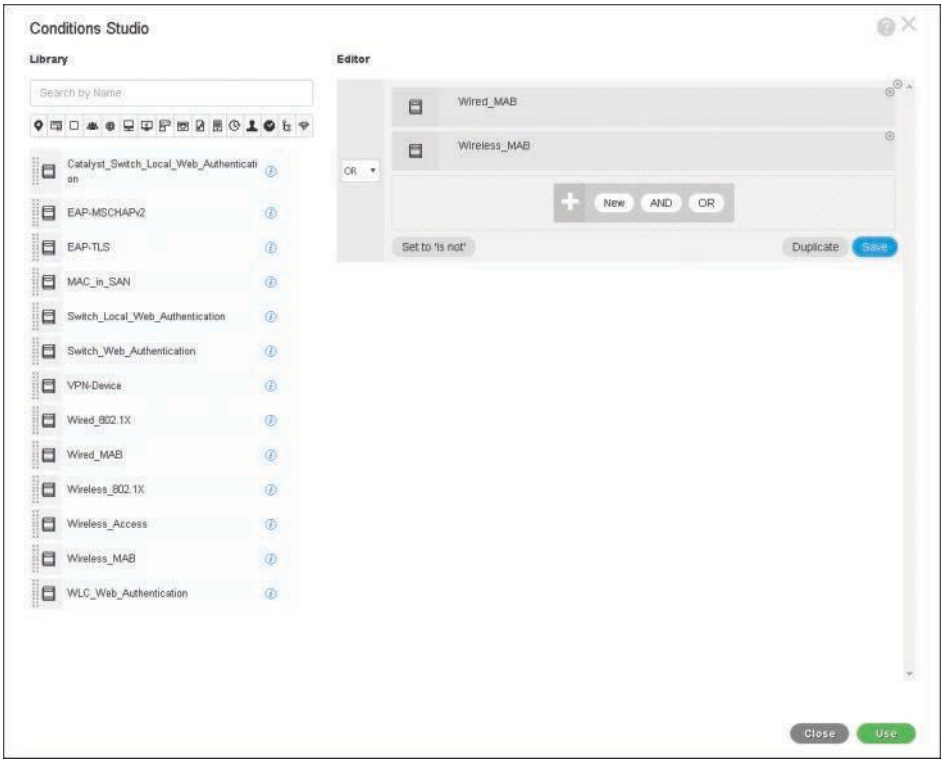


Figure 9-7 *Conditions Studio*

The Conditions Studio has two main parts:

- **Editor**: The editor, in the right pane of the Conditions Studio, is where you can create new and edit existing conditions. The editor is divided into different columns and rows, where the rows represent individual rule conditions and the columns represent additional hierarchical levels that can be added to create complex conditions with the AND and OR operation options. For example, the MAB rule in Figure 9-7 is set to match authentications that are Wired_MAB OR Wireless_MAB.

  You can also set any condition to "Is Not." For example, if you want to have a rule match for authentications from every NAD except NADs located in London, the Is Not option would provide an easy way to accomplish that goal. You would add the condition DEVICE > Location EQUALS "London" and then set that condition to Is Not. This rule would then match on all incoming authentications from all NADs except those in London.

  If you create a condition block that you would like to reuse in other rules or policies, you can click the Save button in the editor to add it to the library.

- **Library**: The library, in the left pane of the Conditions Studio, displays all the saved condition blocks. To use a condition block, you drag it from the library to the editor in the Conditions Studio.

To see more information about a predefined smart condition, you can mouse over the smart condition in the Conditions Studio. A blue I icon becomes visible, and you can see the granular details of what encompasses the smart condition by moving the mouse over that icon.

Alternatively, if you want to see the library of smart conditions without having to navigate all the way to a policy set, you can navigate to Policy > Policy Elements > Conditions > Library Conditions.

As you can see in Figure 9-8, Wired_MAB is looking for the normalized RADIUS flow type to be Wired_MAB.



Figure 9-8 *Wired_MAB Condition*

In Figure 9-9, you can see that Wireless_MAB is very similar to Wired_MAB. It uses the normalized RADIUS flow type Wireless_MAB.



Figure 9-9 *Wireless_MAB Condition*

Identity Store

After matching the conditions in the rule, an authentication request is then authenticated against the chosen identity store or, in the case of MAB, compared to the internal endpoints database (the list of MAC addresses stored locally on ISE).

If the MAC address is known (that is, present in the provided endpoints database), it is considered to be a successful MAB. (Notice that we did not say that it is considered a successful *authentication*.) MAC Authentication Bypass does what its name says: It bypasses authentication and it is not considered a secure authentication on its own.

The selected identity source may be an identity source sequence (ISS), which tries a series of identity stores, in order. This is discussed in more detail in Chapter 20, "ISE Scale and High Availability."

Options

Every authentication rule has a set of options that is stored with the identity store selection. These options tell ISE what to do if an authentication fails, if the user/device is unknown, or if the process fails. Figure 9-10 shows the different options available.



Figure 9-10 *Identity Store Options*

These options are as follows:

- **Reject:** Send Access-Reject back to the NAD.
- **Continue:** Continue to the authorization policy, regardless of whether authentication passes or fails. (This option is used with Web Authentication.)
- **Drop:** Do not respond to the NAD; the NAD then acts as though the RADIUS server is dead.

See Chapters 19, "Deploying Safely," 20, "ISE Scale and High Availability," and 21, "Troubleshooting Tools," for more on when to use these options.

Common Authentication Policy Examples

This section provides a few quick examples of authentication policies based on common use cases.

Using the Wireless SSID

One of the most common authentication policy requests is to treat authentications differently based on the SSID of the wireless network. Creating such a policy is not difficult, but it can be challenging to identify the attribute to use because Source-SSID is not a field in a RADIUS packet. The attribute you need to use is Called-Station-Id as this is the field that describes the wireless SSID name.

Key Topic

This section shows how to build a rule for the SSID CiscoPress that will be configured to do the following:

- Match only authentications coming from the SSID CiscoPress
- Allow only PEAP-MS-CHAPv2 authentications
- Authenticate against Active Directory

To create this rule, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Insert a new policy set above the default policy set by clicking the + button located under the Policy Sets header.

Step 3. Provide a name for the policy set—in this case, **CiscoPress SSID**.

Step 4. For the top-level rule, choose **RADIUS > Called-Station-ID** in the Conditions Studio editor.

Step 5. Select **Ends With**.

Step 6. Type in the SSID name **CiscoPress**, as shown in Figure 9-11.



Figure 9-11 *Called-Station-ID Ends With CiscoPress*

Step 7. To create a new allowed protocols list that only allows PEAP-MS-CHAPv2, click the + button next to the Allowed Protocols dropdown and select **Create a New Allowed Protocol**, as shown in Figure 9-12.



Figure 9-12 *Create a New Allowed Protocol*

Step 8. Enter the name **PEAP-MSCHAPv2 ONLY**.

Step 9. Optionally, provide a description.

Step 10. From the top down, ensure that all the checkboxes are unchecked except for Allow PEAP.

Step 11. Ensure that **Allow EAP-MS-CHAPv2** is enabled under PEAP Inner Methods, as shown in Figure 9-13.

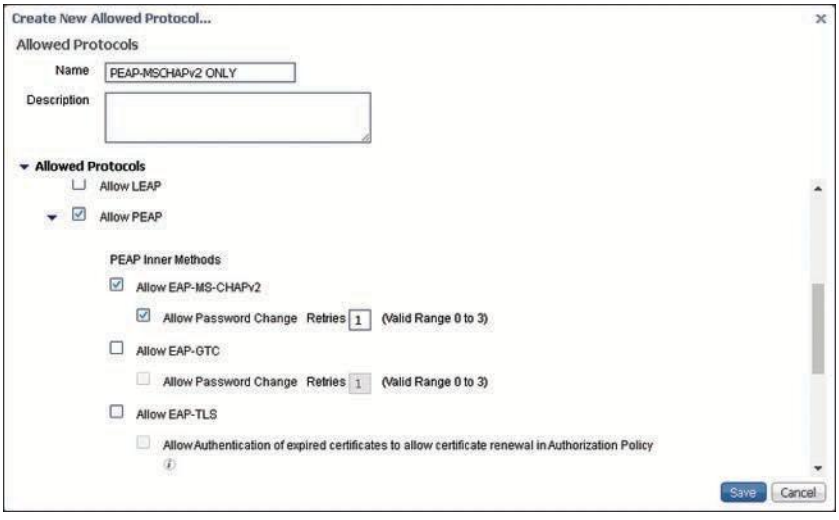


Figure 9-13 *Allowed Protocols*

Step 12. Click **Save** to save the allowed protocols list.

Step 13. Because the new allowed protocols list does not prepopulate in the rule, select it from dropdown in the rule.

Step 14. Click **Save** and then expand the policy set (see Figure 9-14).

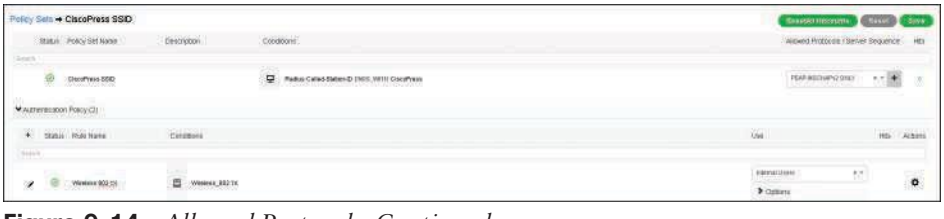


Figure 9-14 *Allowed Protocols, Continued*

Step 15. In the authentication policy, click on the + button above the default rule to create a rule.

Step 16. Name the new rule **Wireless 802.1X**.

Step 17. Click on the + button on the conditions field to bring up the Conditions Studio.

Step 18. From the library, drag and drop the Wireless_802.1X condition into the editor and click **Use**.

Step 19. Under USE, select **ad1** from the dropdown for our identity server (see Figure 9-15). This is the identity store that will be used for this rule, which will be authenticating against Active Directory.

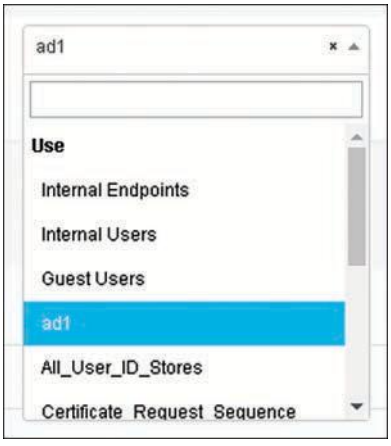


Figure 9-15 *Selecting the AD Identity Source*

Step 20. Leave the default options and click **Save**.

Figure 9-16 shows the completed authentication rule.



Figure 9-16 *Completed Authentication Rule*

The authentication rule is now complete, and the passed authentication can be handed to the authorization policy to determine what actions are taken.

Remote Access VPN

Very often authentications for a remote-access VPN connection are routed to an OTP server, such as RSA SecurID. This section shows how to build a rule for remote-access VPN authentications that will be configured to do the following:

- Match only authentications coming from the VPN device
- Route that authentication to an OTP server

To create this rule, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Insert a new policy set above the previously configured CiscoPress SSID policy set.

Step 3. Name the new policy set **RA VPN**.

Step 4. For the top-level rule, configure the condition as **DEVICE > Device Type**.

Step 5. Set the operator to **Equals**.

Step 6. Select **VPN** as the attribute value.

Step 7. Save the condition to the library by clicking on the **Save** button in the editor window.

Step 8. Save the condition as a new library condition and name it **VPN-Device**.

Step 9. Click **Use**.

Figure 9-17 shows the completed condition.



Figure 9-17 *Device Type Equals VPN*

Step 10. Select the **Default Network Access** allowed protocol list.

Step 11. Click **Save**.

Step 12. Expand the new policy set to view the authentication policy.

Step 13. For the default authentication rule, change the identity store to the OTP server named **ATW_OTP** that was previously configured in Administration > Identity Management > External Identity Sources > RADIUS Token (ATW_OTP).

Step 14. Click **Save**.

Figure 9-18 shows the completed authentication policy.



Figure 9-18 *Completed Authentication Policy*

Alternative ID Stores Based on EAP Type

Thanks to BYOD and mobility, it is very common to have multiple user and device types connecting to the same wireless SSID. In such scenarios, corporate users with corporate laptops often authenticate using EAP-FAST with EAP chaining, and BYOD type devices need to use certificates and EAP-TLS. Anyone authenticating with PEAP would be recognized as a non-corporate and nonregistered asset and would be sent to a device registration portal instead of being permitted network access.

This section shows how to edit the existing CiscoPress SSID policy set. The authentication policy will be configured to do the following:

- Route EAP-TLS authentications to a certificate authentication profile (CAP)
- Route PEAP authentications to an LDAP server
- Route EAP-FAST authentications to Active Directory
- Route EAP-MD5 authentications to internal endpoints for host lookup as a MAB request

To create this rule, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Change the allowed protocol list to the **Default Network Access** list for the CiscoPress SSID Policy Set.

Step 3. Click **Save**.

Step 4. Expand the **CiscoPress SSID** policy set to view the authentication policy.

Step 5. Rename the preconfigured Dot1X rule **Wireless EAP-TLS**.

Step 6. Click on the condition to open the Conditions Studio.

Step 7. Click **New** to add a new attribute.

Step 8. Select **Network Access > EapAuthentication** as the attribute.

TIP To navigate to this attribute more quickly, you can type **eap** into the attribute filter.

Step 9. Set the operator to **Equals**.

Step 10. Select **EAP-TLS** as the attribute value.

Step 11. Click **Use**.

Step 12. For this rule's identity store, choose the previously configured certificate authentication profile **AD1_CA** (see Figure 9-19).



Figure 9-19 *Network Access:EapAuthentication Equals EAP-TLS*

Step 13. Click on the cog on the right side of this authentication rule and select **Duplicate Below** from the menu.

Step 14. Rename the new rule (which is a duplicate of the Wireless EAP-TLS rule) **Wireless PEAP**.

Step 15. Click on the condition to open the Conditions Studio.

Step 16. Click on the **Network Access > EAP Authentication** attribute and change it to **Network Access > EAP Tunnel**.

Step 17. Keep the operator set to **Equals**.

Step 18. Select **PEAP** as the attribute value.

Step 19. Click **Use**.

Step 20. For this rule's identity store, choose the previously configured LDAP server **LDAP1** (see Figure 9-20).



Figure 9-20 *Network Access:EapTunnel Equals PEAP*

Step 21. Click on the cog next to this authentication rule and choose **Duplicate Below**.

Step 22. Change the name of the duplicated rule to **Wireless EAP-FAST**.

Step 23. Click on the condition to open the Conditions Studio.

Step 24. Change the attribute value to **EAP-FAST**.

Step 25. Click **Use**.

Step 26. For this rule's identity store, choose the previously configured Active Directory domain **AD1** (see Figure 9-21).



Figure 9-21 *Network Access:EapTunnel Equals EAP-FAST*

Step 27. Click on the cog next to this authentication rule and choose **Duplicate Below**.

Step 28. Change the name of this rule to **Wireless EAP-MD5**.

Step 29. Click on the condition to open the Conditions Studio.

Step 30. Click on the **Network Access > EAP Tunnel** attribute and change it to **Network Access > EAP Authentication**.

Step 31. Keep the operator set to **Equals**.

Step 32. Select **EAP-MD5** as the attribute value.

Step 33. Click **Use**.

Step 34. Select **Internal Endpoints** for the identity source.

Step 35. Change the default rule's identity store (bottom row) to **Deny Access**.

Step 36. Click **Save**.

Figure 9-22 shows the completed rule and subrules.



Figure 9-22 *Completed Authentication Rule and Subrules*

More on MAB

Something that is often not understood, especially when looking to mix access device vendors, is MAC Authentication Bypass (MAB). There is no standard for MAB. Different vendors implement MAB in different ways. Ultimately, the goal of MAB is to allow the supplicant in the switch to run an authentication request for the endpoint because the endpoint obviously must not have a supplicant.

Some vendors send a RADIUS message with Service-Type set to Login, and some send a RADIUS message with Service-Type set to Framed. Cisco uses the Service-Type setting Call-Check for MAB. Why would Cisco use Call-Check if no other vendor does? Why does Cisco do MAB differently than everyone else? The quick answer is for security.

Key Topic

Many years ago, before Cisco released Cisco ISE or the Cisco ACS 5.x server, there was a possible security vulnerability with MAB. That security vulnerability is still possible with other solutions and other network devices. The issue was/is the lack of differentiation between a MAB request and a Local Web Authentication request. Both requests come from the network device with the same Service-Type setting and the same format. There was no database separation of user IDs from endpoint IDs (that is, MAC addresses). As shown in

Figure 9-23, a malicious user could enter a MAC address into the username and password fields for Web Authentication or maybe even into the endpoint supplicant and gain access to the network.

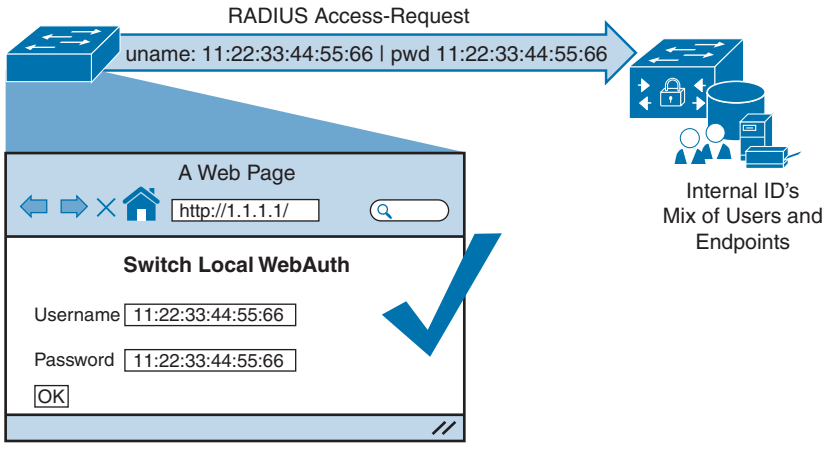


Figure 9-23 *Web Authentication with MAC Address Instead of Username*

In an effort to close this security hole and make MAB a bit more secure, Cisco changed the way it does MAB. It implemented the following unique requirements:

- For authentication requests to be processed as MAB (by default), Service-Type must be Call-Check.
- RADIUS servers (ISE) maintain a separate endpoint database.
- The Calling-Station-Id value is compared to the endpoint database, and the username and password fields of the MAB request are ignored.

All supported Cisco network access devices use Call-Check for Service-Type for MAB requests. They also ensure that the Calling-Station-Id field is populated with the MAC address of the endpoint. In addition, Cisco ISE has a simple checkbox in the allowed protocols configuration that makes it possible to permit or deny access to the endpoint database for a MAB request, as shown in Figure 9-24.

Keep in mind that MAB is inherently not a secure technology. When you implement MAB, you are bypassing the stronger security of 802.1X by allowing specific MAC addresses to gain access without authentication. When using MAB, always follow a defense-in-depth approach. This means a device that has been authorized to use the network from a MAB request should be granted access only to the networks and services that device is required to speak to. In other words, don't provide full access to devices that have gone through successful MAB but instead provide them with an authorization that is more limited. This topic is discussed in more detail in Chapter 10, "Authorization Policies."

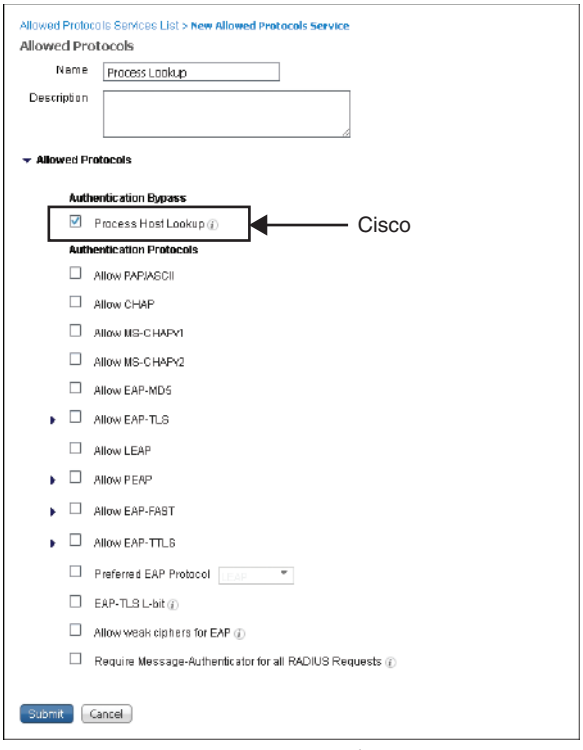


Figure 9-24 *Process Host Lookup*

Restore the Authentication Policy

In this chapter, you have seen how to create a complex and specific authentication policy. If you followed along and created this policy yourself, the process helped you learn how authentication policies work. However, having this policy on your system might make things a bit too complicated for you as you work through the future chapters. To keep things simple, follow these steps in the ISE GUI to restore your authentication policy to a more simple policy that will work for all use cases remaining in this book:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Delete the policy set named **CiscoPress SSID**.

Step 3. Delete the policy set named **RA-VPN**.

Step 4. Click **Save**.

Chapter 10 takes an in-depth look at authorization policies and common authorization rules.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 9-2 lists these key topics and the page number on which each is found.

Key Topic

Table 9-2 Key Topics

Key Topic Element	Description	Page
Paragraph	Authentication versus authorization	209
Figure 9-3	Default network access allowed protocol list	213
List	Using the wireless SSID	221
Paragraph	MAC Authentication Bypass (MAB)	227

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is authentication?
2. What is authorization?
3. What are the goals of an authentication policy?
4. What are policy sets?
5. How do Cisco NADs do MAB differently than other vendors' NADs?

This page intentionally left blank



CHAPTER 10

Authorization Policies

This chapter covers the following topics:

- Implement wired/wireless 802.1X
- AV pairs
- Implement MAB
- ISE authorization policies
- Implement network authorization enforcement
- Network authorization enforcement

As discussed in Chapter 9, "Authentication Policies," an authentication is the validation of a credential or an identity. The other major portion of secure network access control is the control portion. *Authorization* validates the level of entitlement or access that should be granted. That is where authorization policies come in.

Consider a bank example: Walk into a bank and provide a government-issued passport as identification. It checks out completely, the seal is valid, and the picture matches, so there is a successful authentication. However, does this successful authentication mean that you are entitled to withdraw money from the bank? Does it allow you to withdraw money from any account in the bank? As discussed in Chapter 9, the entitlement to withdraw money from the bank would be known as an authorization. In secure network access, the authorization policy determines what level of network access an incoming user or device should be granted, based on the who, what, where, when, and how of the authentication, which is what Cisco calls the *security context* of the user or device.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 10-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 10-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Authentication Versus Authorization	1
Authorization Policies	2–4, 6, 9, 10
Saving Conditions for Reuse	5, 7, 8

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What is an authorization profile?
   a. An authorization profile is a rule in the policy table in the format "IF condition THEN result."
   b. An authorization profile is created to determine which identity store to use for validating the credentials.
   c. An authorization profile is a sequential list of identity stores for validating the credentials.
   d. An authorization profile is the mandatory result of an authorization rule.
2. What is the purpose of an authorization profile?
   a. It contains the TACACS+ response Access-Accept or Access-Reject along with the additional authorization attributes to be sent to the network device for enforcement.
   b. It contains the RADIUS response Access-Accept or Access-Reject along with the additional authorization attributes to be sent to the network device for enforcement.
   c. It contains the RADIUS response Continue or Terminate along with additional authorization attributes to be sent to the network device for enforcement.
   d. It contains the TACACS+ response Continue or Terminate along with additional authorization attributes to be sent to the network device for enforcement.
3. Which of the following are common tasks in an authorization profile?
   a. Access type (Continue or Terminate), dACL name, web redirection, auto smart port
   b. Access type (Accept or Reject), dACL name, web redirection, auto smart port
   c. dACL name, role assignment, Local WebAuth, auto smart port
   d. dACL name, web redirection, Local WebAuth, auto smart port
4. Which of the following is correct?
   a. An authorization policy contains authorization rules, and each rule has at least one authorization profile.
   b. An authorization rule contains authorization policies, and each policy has at least one authorization profile.
   c. An authentication policy contains authorization rules, and each rule has an authentication result.
   d. An authentication rule contains the authorization profiles, and each profile contains one authentication result.

5. What is the primary purpose of saving conditions to the library?
   a. For future use and improved readability
   b. For complex conditions
   c. To combine AND and OR operators
   d. To create multilevel compound conditions
6. What is special about the authorization profile required for an IP phone?
   a. It contains the DNS name or IP address of the Cisco Call Manager server.
   b. It contains the voice domain permission AV pair, which authorizes the endpoint to access the voice VLAN assigned to the interface.
   c. It contains the value for DHCP option 43, which provides the IP address of the Cisco Call Manager server.
   d. It contains the voice domain permission macro, which reconfigures the switch port to be a voice interface.
7. What is the difference between simple conditions and compound conditions?
   a. Simple conditions are easier to use than compound conditions.
   b. Simple conditions are created on the fly in the expression builder, whereas compound conditions must be created separately.
   c. A simple condition contains only one attribute. A compound condition contains multiple attributes along with an operator such as AND or OR.
   d. Both simple conditions and compound conditions can contain multiple attributes, but compound conditions can mix operators such as AND or OR.
8. If you need to combine OR and AND operators with conditions to grant access, what do you do?
   a. Create hierarchical compound conditions
   b. Create multiple compound conditions
   c. Create multiple simple conditions
   d. Create multiple authorization rules
9. What should be the end goal of a Cisco Secure Access deployment?
   a. To provide full access to the network so security devices such as an ASA firewall can provide defense-in-depth
   b. To provide full access to the network, as long as the authentication is successful, to provide limited access to any failed authentications
   c. To secure the network by purchasing Cisco ISE, thereby increasing the stock value of the organization
   d. To provide very specific permissions to any authorization to provide defense-in-depth

10. What is unique about Cisco's downloadable access control lists (dACLs)?
    a. Cisco dACLs allow the RADIUS server to apply ACLs that exist on the switch simply by sending the name of the ACL in the RADIUS AV pairs, and non-Cisco network devices cannot apply ACLs.
    b. Cisco dACLs are created by experts at Cisco and published to Cisco.com, where Cisco ISE can download the ACLs.
    c. Cisco dACLs are created entirely on the RADIUS server, and a full ACL is sent down to the network device within RADIUS AV pairs. Non-Cisco network devices must create ACLs on individual local network devices if they do not support the NAS-Filter-Rule RADIUS attribute.
    d. Cisco dACLs are unique because they are downloaded from ISE and applied to the Cisco ASA that is in the network path, relieving the network device of the burden of traffic control.

Foundation Topics

Authentication Versus Authorization

At the beginning of Chapter 9, we examined the relationship between authentication and authorization. Remember that an authentication is the validation of credentials or an identity. Consider the BYOD use case, where a tablet connects to the network and authenticates with a certificate via EAP-TLS. The authentication policy terminates the EAP session and validates that the certificate was signed by a trusted source and has not expired or been revoked. The authentication succeeds; however, just having a successful authentication does not mean that access should be allowed. Another process needs to be performed to determine what level of access should be granted to the tablet. That access can be based on virtually any attribute, such as device type, location, business unit, user identity extracted from the certificate, Active Directory group membership of the user, the CA that signed the certificate, time of day, or even hair and eye color. The conglomeration of attributes is often referred to as the ***security context*** of the device. A different level of access may be granted as long as there is a rule in the authorization policy that matches the context.

These rich policies are part of the authorization policy. The flexibility of being able to assign a different level of access based on context is why most of the time spent working in a product like Cisco ISE is spent setting up and tuning the authorization policy. The authorization policy is where the bulk of the decisions are made because it is where you specify what access should be granted after a successful authentication occurs.

Authorization Policies

After an endpoint has been successfully authenticated, the authenticated process is passed to the authorization policy, which determines the final access results.

Goals of Authorization Policies

Key Topic

Authorization policies have one main goal: to examine conditions in authorization rules in order to ultimately send an authorization result to a network access device (NAD). What conditions? Well, what did you have in mind?

Common conditions could include internal and external attributes such as Active Directory group membership or internal group membership in ISE. Policies can be built using a myriad of attributes, such as location, time, whether a device was registered, whether a mobile device has been jail broken, or nearly any other attribute imaginable. Even the authentication can be used as an attribute: Was authentication successful? Which authentication protocol was used? What is the content of specific fields of the certificate that was used?

The authorization policy compares the conditions with the explicit goal of providing an authorization result. The result may be as simple as returning a standard RADIUS Access-Accept or Access-Reject message. It can also include more advanced factors, such as VLAN assignment, downloadable access control lists (dACLs), Security Group Tag (SGTs), URL redirection, and more.

An authorization policy not only allows or denies access to the network but can also include any and all restrictions for limiting network access for the user or endpoint.

Understanding Authorization Policies

Now that you understand the main responsibilities of an authorization policy, it will be easier to understand the rest of the material in this chapter. We will now go through a few authorization policies and examine them to understand them better.

Basic authorization policy rules are logically organized in this manner:

```
IF [conditions] THEN [AssignThesePermissions]
```

Just as in an authentication policy, the authorization policy rules are processed in a top-down, first-match order. So if the conditions do not match for a rule, the authentication is compared to the next rule in the policy.

ISE comes preconfigured with several authorization rules in the default policy set. These rules address some of the common use cases that enterprises might need, such as rules for basic endpoint posturing, blacklisted endpoints, profiled IP phones, 802.1X, BYOD, and guest access.

This section examines the Cisco IP Phone and blacklist rules in order to show how they work. If you have a live ISE system, it might help to follow along on it with the text.

Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the default policy set.

Step 3. Expand the authorization policy, as shown in Figure 10-1.

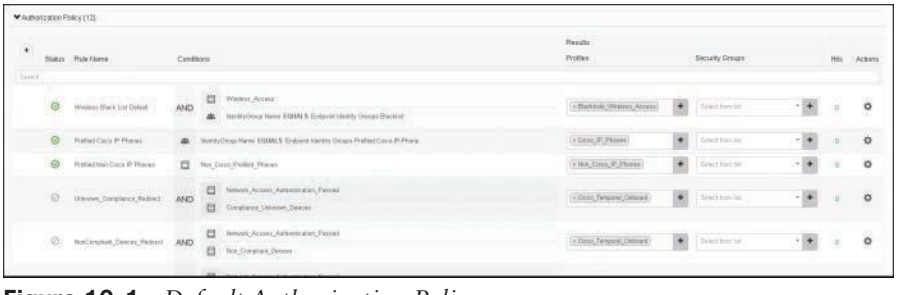


Figure 10-1 *Default Authorization Policy*

For the Profiled Cisco IP Phones rule, the only condition that needs to be matched is that the endpoint needs to be part of an identity group named Cisco-IP-Phone. The authorization profile permits access to any device profiled as a Cisco IP Phone by sending a RADIUS Access-Accept message and an attribute/value pair (AV pair) that permits the phone into the voice VLAN.

Step 4. To examine the permissions (profile) that is used, navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**, as shown in Figure 10-2. An authorization profile is a set of authorization results that should be sent together. Every authorization rule must have a result.

NOTE Notice that there is one more category under Authorization: Downloadable ACLs. We talk more about downloadable ACLs (dACLs) later in this chapter.



Figure 10-2 *Default Authorization Profiles*

Step 5. Select the **Cisco_IP_Phones** authorization profile. The authorization result needs to be RADIUS attributes. In order to make that easier for users of ISE, Cisco has included a Common Tasks section that presents the options in a plain English format. The following are some of the common tasks that are available:

- dACL (IPv4 and IPv6 options)
- ACL (IPv4 and IPv6 options)
- Security Group Tag
- VLAN
- Voice Domain Permissions
- Web Redirection (Local and Centralized options)
- Auto Smart Port
- Assess Vulnerabilities
- Reauthentication
- MACsec Policy
- NEAT
- Interface Template
- Airespace ACL Name (IPv4 and IPv6 options)

- ASA VPN
- AVC Profile

NOTE We look at only some of these common tasks in this book as some of them are beyond the scope of the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam.

The Advanced Attributes Settings section in the authorization profile is for use when an AV pair needs to be returned and is not covered by a common task. For most use cases and deployments, an ISE administrator does not ever have to use this part of the authorization profile.

The Attributes Details section at the bottom of the screen shows the raw RADIUS result that will be sent to the NAD (see Figure 10-3).

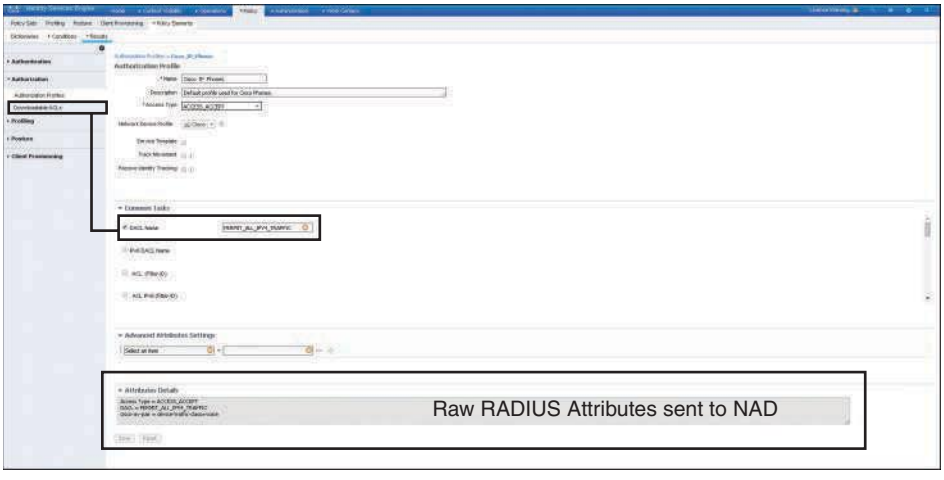


Figure 10-3 *Cisco_IP_Phones Authorization Profile*

Note in Figure 10-3 that you can use the DACL Name dropdown to select a dACL that was created and stored in ISE. The Voice Domain Permission checkbox is required for the switch to permit the phone into the voice VLAN on the switch.

Step 6. Notice in the Attributes Details section that this authorization result will send a RADIUS Access-Accept message, a dACL that permits all traffic, and the voice domain vendor-specific attribute (VSA) to permit the phone to join the voice VLAN.

Next, let's examine the wireless blacklist default rule:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the default policy set.

Step 3. Expand the authorization policy.

Step 4. Select the rule named **Wireless Black List Default**. Notice that the Wireless_Access conditions and the identity group are separated by the AND operator. If both the conditions match, ISE returns the results contained in the authorization profile. Blackhole_Wireless_Access is the authorization profile used for this rule. This profile says to return a RADIUS Access-Accept message with an AV pair that will apply an ACL and a URL redirect. This particular rule is meant to prevent devices that have been marked lost or stolen from accessing the network.

Step 5. To examine the authorization conditions, click the condition in the rule to open the Conditions Studio (see Figure 10-4).

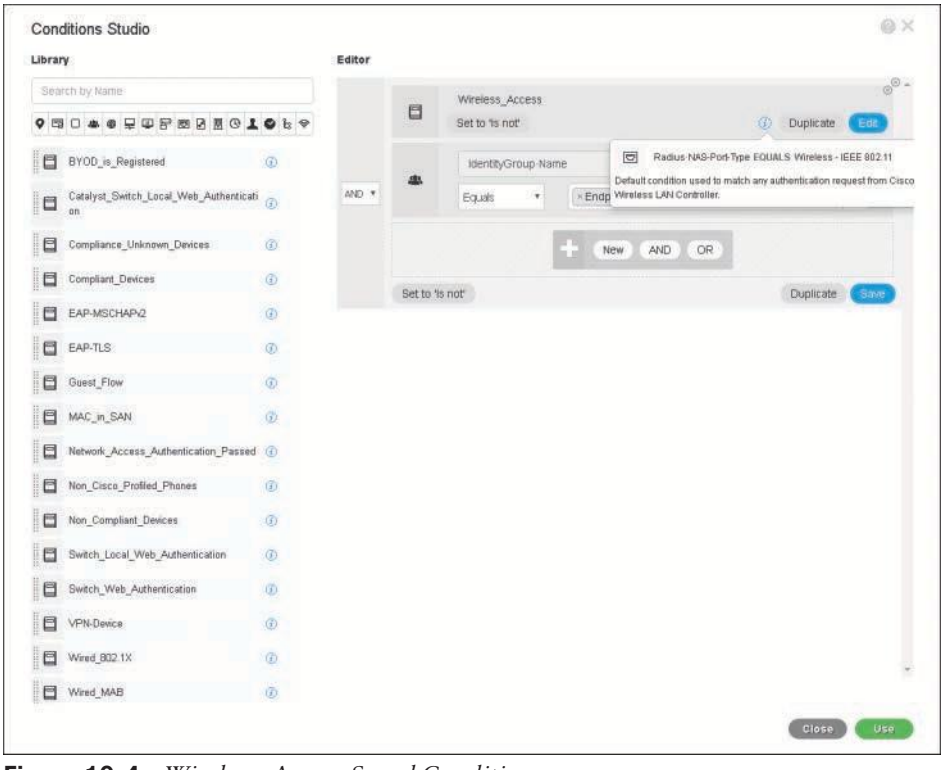


Figure 10-4 *Wireless_Access Saved Condition*

A ***simple condition*** is just a single attribute that has been saved to the library for reuse. A ***compound condition*** is more than one condition joined with an AND or OR operator. You can save a compound condition to the library for reuse by clicking the **Save** button at the bottom of the editor. The Wireless_Access condition shown in Figure 10-4 is a simple condition. However, if you click the **Save** button, the Conditions Studio saves the Wireless_Access and the Blacklist endpoint identity group together as a compound condition that you could later reuse.

Figure 10-5 shows the Save Condition window, where you can add a new saved condition or replace an existing saved condition.

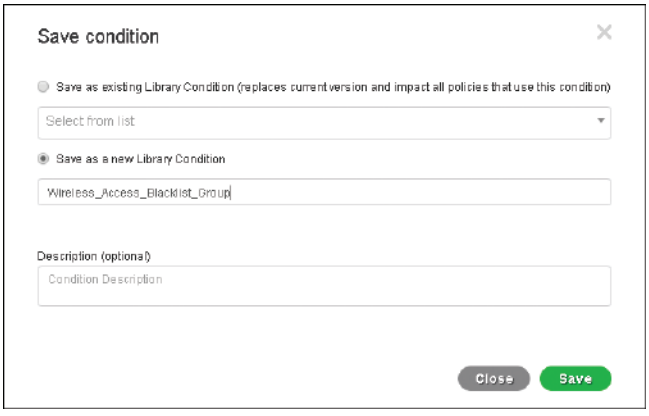


Figure 10-5 *Saving a New Compound Condition*

Step 6. To examine the authorization result that is being sent for this authorization rule, navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles.**

Step 7. Select **Blackhole_Wireless_Access.** As shown in Figure 10-6, the Blackhole_Wireless_Access authorization profile does not use any of the common tasks. Instead, it uses the advanced attribute settings to send a URL-Redirect and URL-Redirect-ACL result to the WLC, along with an Access-Accept. So, this result is allowing the devices onto the network but is forcing all traffic to redirect to a special blacklist web page.



Figure 10-6 *Blackhole_Wireless_Access Authorization Profile*

The two authorization rules reviewed in this section demonstrate a variety of rules. Later in this chapter, we will examine a few common authorization policies.

Role-Specific Authorization Rules

Key Topic

Defense-in-depth is an important concept. Just because you have implemented 802.1X at the access layer with Cisco ISE as the policy controller does not mean the network is immediately secure or that segmentation has been achieved. Everything depends on the security policies built by you, the administrator.

For example, organizations often use a combination of ISE profiling capabilities and MAB to provide network access to printers. ISE profiling is used to assign a printer profile to a MAC address belonging to a device that appears to be a printer. MAB is then used for the network access device to send an authentication request containing the MAC address of the printer to ISE. From there, ISE should send down the Access-Accept result.

If there is a successful MAB, and the printer needs to have access to the network, does that mean the printer deserves full access to the network? Absolutely not! The printer should be granted only enough access to perform its necessary functions. That way, a malicious user who spoofs the MAC address of the printer does not have completely unfettered access to the network.

The end goal of a secure access deployment should be to provide very specific permissions to any authorization. However, that should always be handled in a staged approach in order to limit the impact to the end users. Chapter 19, "Deploying Safely," is dedicated to this phased approach.

Authorization Policy Example

This section provides an example of an authorization policy made up of numerous rules based on common use cases. These use cases were selected to show multiple aspects of the authorization policy and help solidify your working knowledge of the parts of an authorization policy and the workflows associated with creating these policies.

This example shows how to configure three authorization rules. One of the rules assigns full access to an employee who authenticated successfully with EAP chaining, another rule assigns more limited access to the same employee authenticating with a noncorporate machine, and the final rule assigns Internet-only access to the same employee authenticating on a mobile device.

Employee Full Access Rule

This rule assigns full access permissions to an employee who is authenticating from a valid corporate asset. To configure it, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Insert a new policy set above the default policy set.

Step 3. Name the new policy set **Wired Access**.

Step 4. For the top-level condition of the policy set, drag and drop the **Wired_MAB** and **Wired_802.1X** conditions from the library into the editor in the Conditions Studio.

Step 5. Select the **Default Network Access** allowed protocol list from the dropdown.

Step 6. Click **Save**.

Step 7. Expand the new policy set.

Step 8. Expand the authorization policy.

Step 9. Click the + button to add a new authorization rule above the default rule.

Step 10. Name the new rule **Employee and CorpMachine**.

Step 11. Click on the + button to open the Conditions Studio.

Step 12. Select **Network Access > EapChainingResult** for the attribute.

Step 13. Choose **Equals**.

Step 14. Select **User and Machine both Succeeded** as the attribute value.

Step 15. Click the **New** button to add a new condition.

Step 16. Select **AD1 > External Groups Equals Employees** (or another AD group of your choosing) as the condition.

Step 17. For the authorization profile, click the + button.

Step 18. Select **Create a New Authorization Profile**.

Step 19. Name the new authorization profile **Employee Full Access**.

Step 20. Optionally add a description.

Step 21. Set Access Type to **Access_Accept**.

Step 22. Check the box next to **DACL Name** and select **PERMIT_ALL_IPV4_TRAFFIC** from the dropdown, as shown in Figure 10-7.



Figure 10-7 *Employee Full Access Authorization Profile*

Step 23. Click **Save**.

Step 24. Select the new authorization profile from the dropdown menu.

Step 25. Click **Save** to save the authorization policy.

Figure 10-8 shows the completed authorization rule.

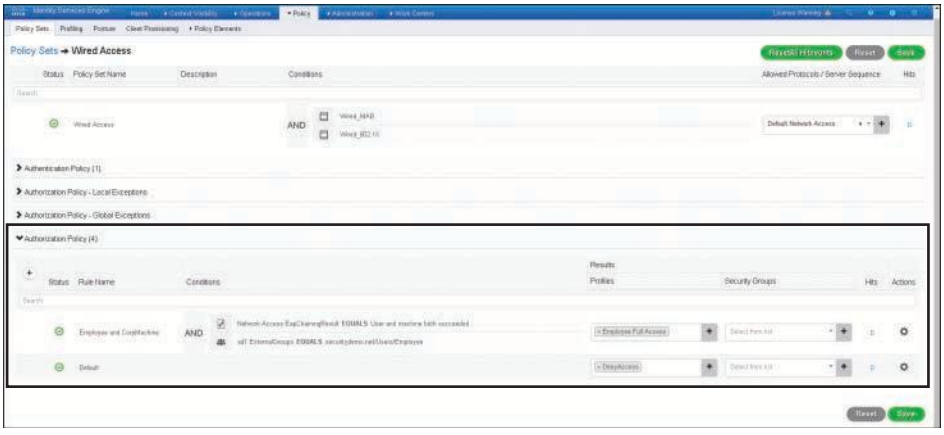


Figure 10-8 *Completed Employee and CorpMachine Rule*

Internet Only for Smart Devices Rule

Now that the rule for employees with corporate devices has been created, you need to create a rule below it to provide Internet access only to employee authentications on mobile devices.

To begin this rule, you first create a new dACL that will be applied to switches. Then you create the authorization result. Finally, you go back into the authorization policy and build the rule. To configure this rule, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Downloadable ACLs**.

Step 2. Click **Add**.

Step 3. Name the ACL **Internet_Only**.

Step 4. Optionally provide a description.

Step 5. In the DACL Content box, provide an ACL that permits required traffic and denies traffic destined to the corporate network. Figure 10-9 provides an example of this rule. You might want to customize this rule for your own organization.

NOTE If these are devices that will be on wireless instead of wired, you would configure an ACL on the wireless controller and specify the name of the Airespace ACL in the Authorization Profile instead.

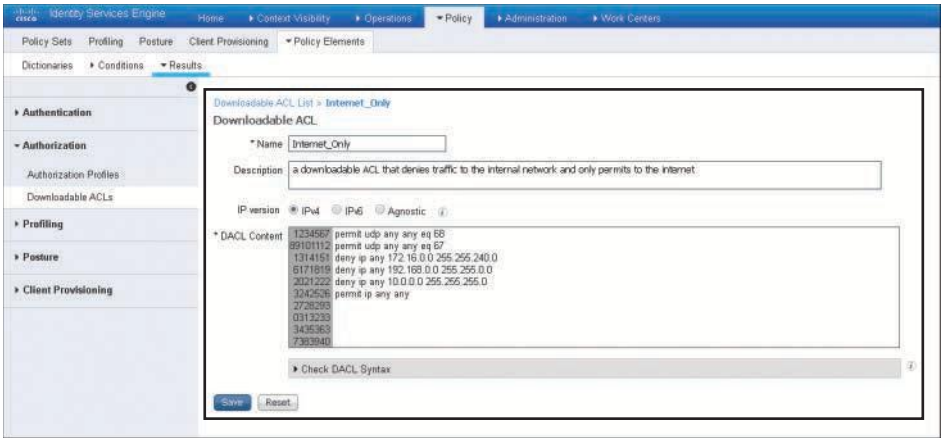


Figure 10-9 *Internet_Only dACL*

Step 6. Click **Save**.

Now that the dACL has been created, it's time to create the authorization profile. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

Step 2. Click **Add**.

Step 3. Name the authorization profile **Internet Only**.

Step 4. Optionally provide a description.

Step 5. Set Access Type to **ACCESS_ACCEPT**.

Step 6. Check the box next to **DACL Name** and select **Internet Only** from the drop-down. If you were going to reuse this authorization profile for wireless access, you could select **Airespace ACL Name** and fill in the name of the ACL on the controller that will provide only Internet access. Since you are referencing this ACL by name only, this ACL must already exist on the wireless LAN controller; this is referred to as a named ACL.

Step 7. Optionally provide a different VLAN assignment for the end user. Keep in mind that this VLAN name or ID will be used for both wired and wireless devices. An alternative is to create separate rules for wired and wireless, so the user can be assigned to a VLAN on wireless but not wired, for example. In either case, the VLAN must already exist on the NAD.

Step 8. Click **Submit**. Figure 10-10 shows the completed authorization profile.



Figure 10-10 *Internet Only Authorization Profile*

Before you build the authorization policy, you are going to create a logical profiling policy that encompasses all mobile device platforms. This will make the policy building much easier and provide a reusable policy object. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Profiling > Logical Profiles.**

Step 2. Click **Add.**

Step 3. Name the Logical Policy **SmartDevice.**

Step 4. Optionally provide a description.

Step 5. Select all the mobile platforms from the Available Devices side and click **>** to move them to the Assigned Policies side, as shown in Figure 10-11.

Step 6. Click **Submit.**

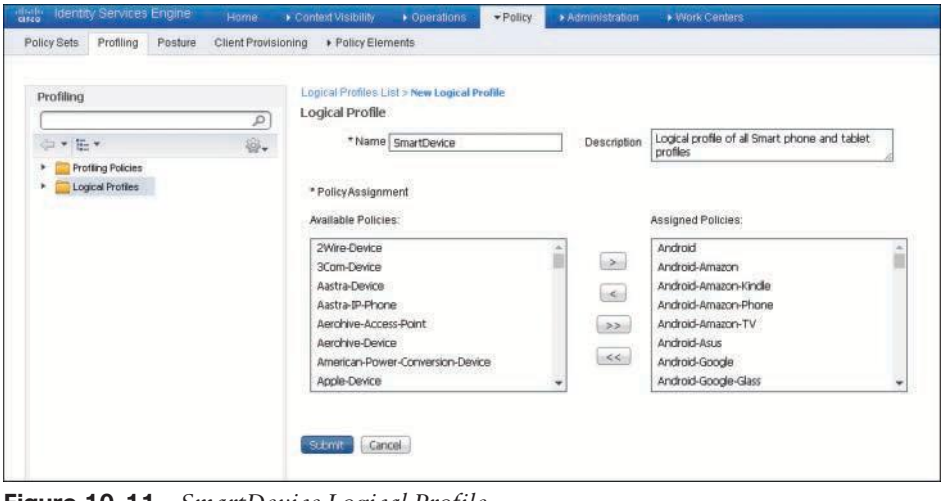


Figure 10-11 *SmartDevice Logical Profile*

Finally, it is now time to create the authorization rule. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the **Wired Access** policy set.

Step 3. Expand the authorization policy.

Step 4. Add a new authorization rule below the Employee and CorpMachine rule by clicking on the cog to the right of this rule and choosing **Insert New Row Below**.

Step 5. Name the rule **Employee SmartDevices**.

Step 6. Click on the + button to bring up the Conditions Studio and select **Endpoints > LogicalProfile** as the attribute.

Step 7. Select **Equals**.

Step 8. Choose **SmartDevice** as the attribute value.

Step 9. Click the **Add** button to add a new condition.

Step 10. Select **AD1 > External Groups Equals "Employees"** or another AD Group of your choosing.

Step 11. Click **Use**.

Step 12. For the authorization profile, choose **Internet Only** from the dropdown.

Step 13. Click **Save**.

Figure 10-12 shows the completed authorization rule.

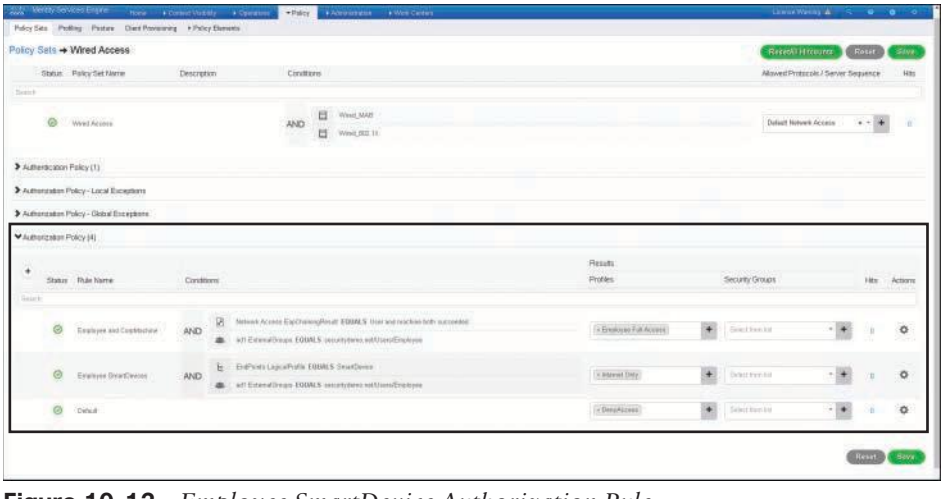


Figure 10-12 *Employee SmartDevice Authorization Rule*

Employee Limited Access Rule

Now that the rule for employees connecting with mobile devices has been created, you need to create a rule below it to provide limited access to employee authentications on any other device.

To begin this rule, you will first create a new dACL that will be applied to switches. Then you will create the authorization result. Finally, you will go back into the authorization policy and build the rule. To configure this rule, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Downloadable ACLs**.

Step 2. Click **Add**.

Step 3. Name the ACL **Employee_Limited**.

Step 4. Optionally provide a description.

Step 5. In the DACL Content box, provide an ACL that permits required traffic and denies traffic destined to anything else. For this example, you can allow traffic to reach only Virtual Desktop Infrastructure and essential services such DNS, as shown in Figure 10-13.

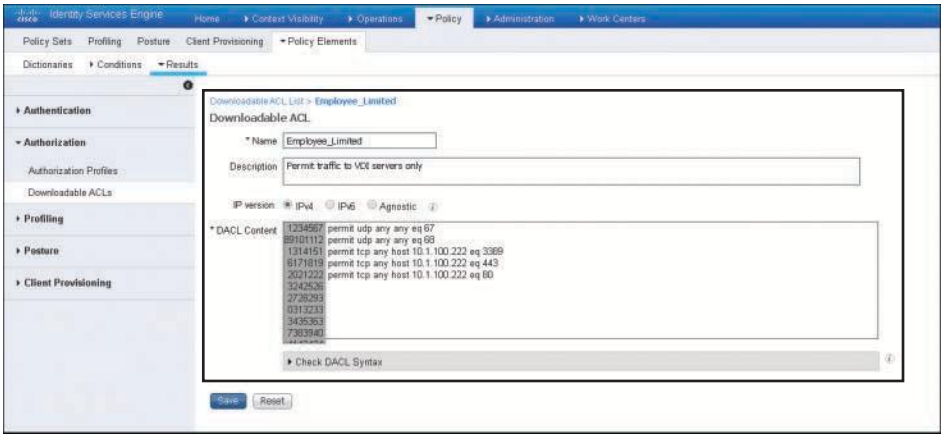


Figure 10-13 *Employee Limited dACL*

Step 6. Click **Save**.

Now that the dACL is created, you can build the authorization policy to permit network access and apply that dACL. To do so, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

Step 2. Click **Add**.

Step 3. Name the authorization profile **Employee Limited**.

Step 4. Optionally provide a description.

Step 5. Set Access Type to **ACCESS_ACCEPT**.

Step 6. Check the box next to **DACL Name** and select **Employee_Limited** from the dropdown.

NOTE You do not need to assign a different VLAN for this authorization.

Step 7. In order to reuse this authorization profile in a wireless policy in the future, select **Airespace ACL Name** and fill in the name of the ACL on the controller that will provide Internet only access.

Step 8. Click **Submit**.

Figure 10-14 shows the completed authorization profile.

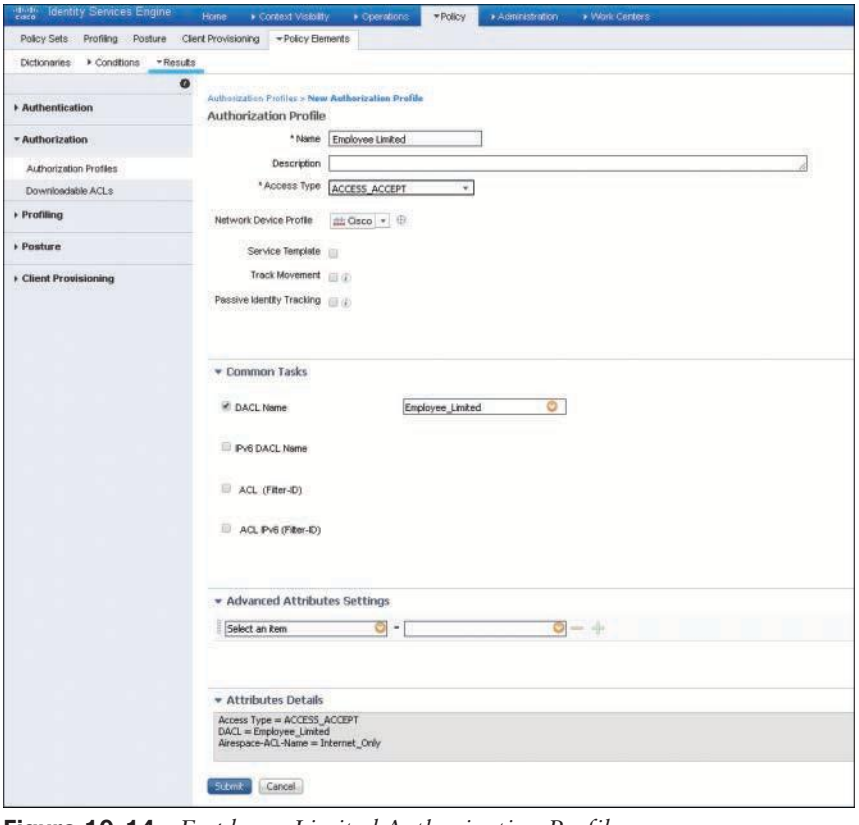


Figure 10-14 *Employee Limited Authorization Profile*

Now, you can create the authorization rule that will use the authorization profile you just created. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the **Wired Access** policy set.

Step 3. Expand the authorization policy.

Step 4. Click on the cog next to the Employee SmartDevices rule and choose **Insert New Row Below**.

Step 5. Name the rule **Employee Limited**.

Step 6. Click the **+** sign to open the Conditions Studio.

Step 7. Select **AD1 > External Groups Equals "Employees"** (or another AD group of your choosing) as the condition.

Step 8. Select **Employee Limited** from the dropdown for the authorization profile.

Step 9. Click **Save**.

Figure 10-15 shows the completed authorization rule.



Figure 10-15 *Employee Limited Authorization Rule*

Saving Conditions for Reuse

Key Topic

ISE offers the ability to save conditions to the library to make it much easier to reuse them in other policies. To see how this works, this section returns to the authorization policy example and shows how to save a few of the conditions. Not only does saving conditions make it easier to find the conditions in the future, it helps clean up the policies and makes them easier to read.

Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the **Wired Access** policy set and the authorization policy.

Step 3. Edit the previously created **Employee and CorpMachine** authorization rule.

Step 4. Click on the condition to open the Conditions Studio.

Step 5. In the editor, click **Save** and name this condition **EmployeeFullEAPChain**. Figure 10-16 shows where you save the full set of conditions, including the AND operator, to the library.

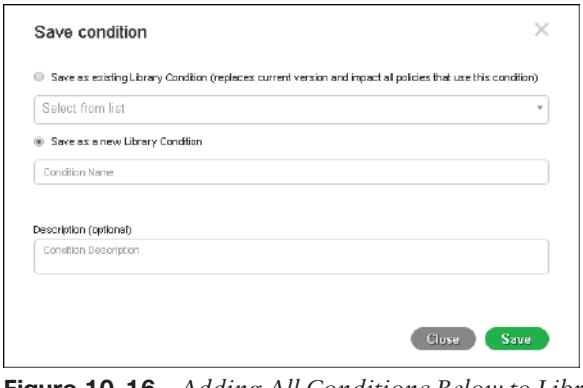


Figure 10-16 *Adding All Conditions Below to Library*

Step 6. Name this new saved condition **EmployeeFullEAPChain**.

Step 7. Click **Save**.

As shown in Figure 10-17, the authorization policy text is simplified now, with the name of the saved conditions instead of the raw attributes, and the saved condition is now displayed in the library.

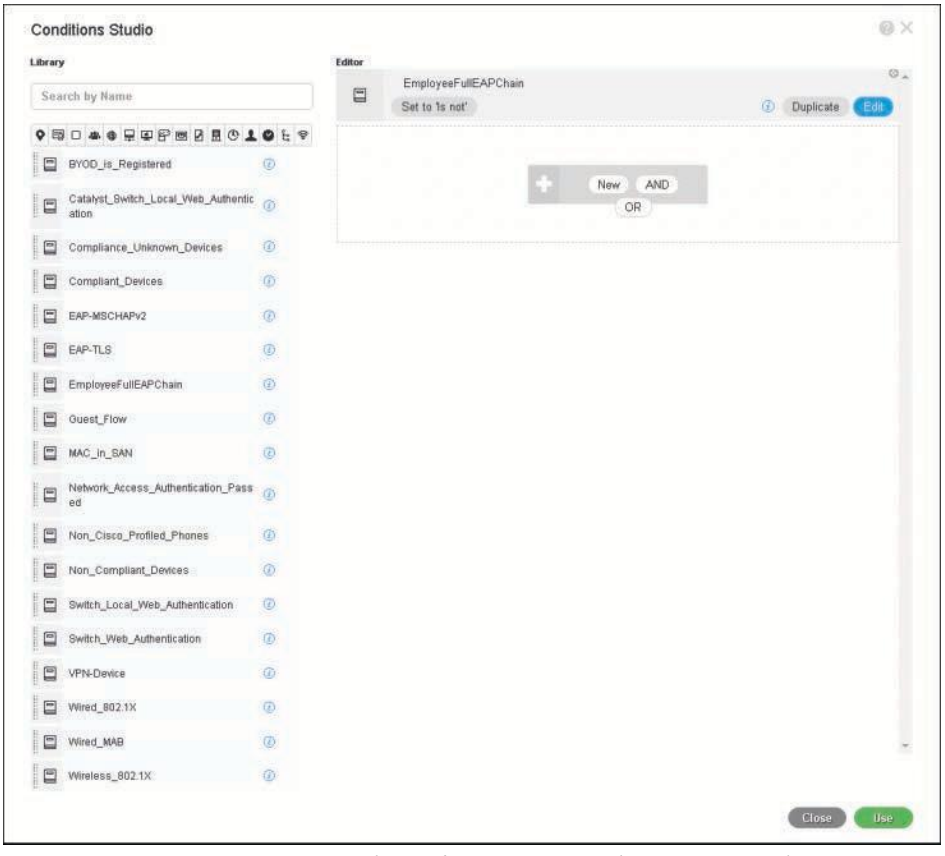


Figure 10-17 *Authorization Policy After Saving Conditions to the Library*

Although Figure 10-17 does not show it, conditions can be saved individually. It is often desirable to save individual conditions to the library, as they can be mixed and matched with other library conditions or raw attributes, which provides more flexibility. Next, you will save the employees group for Active Directory as a condition because it is a common condition used in many rules. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the **Wired Access** policy set and the authorization policy within it.

Step 3. Edit the conditions for the **Employee SmartDevice** authorization rule.

Step 4. Click on the individual condition for the Active Directory group.

Step 5. Click the **Save** button that appears directly underneath this condition to save the single condition to the library.

Step 6. Name the condition **Employees**, as shown in Figure 10-18.

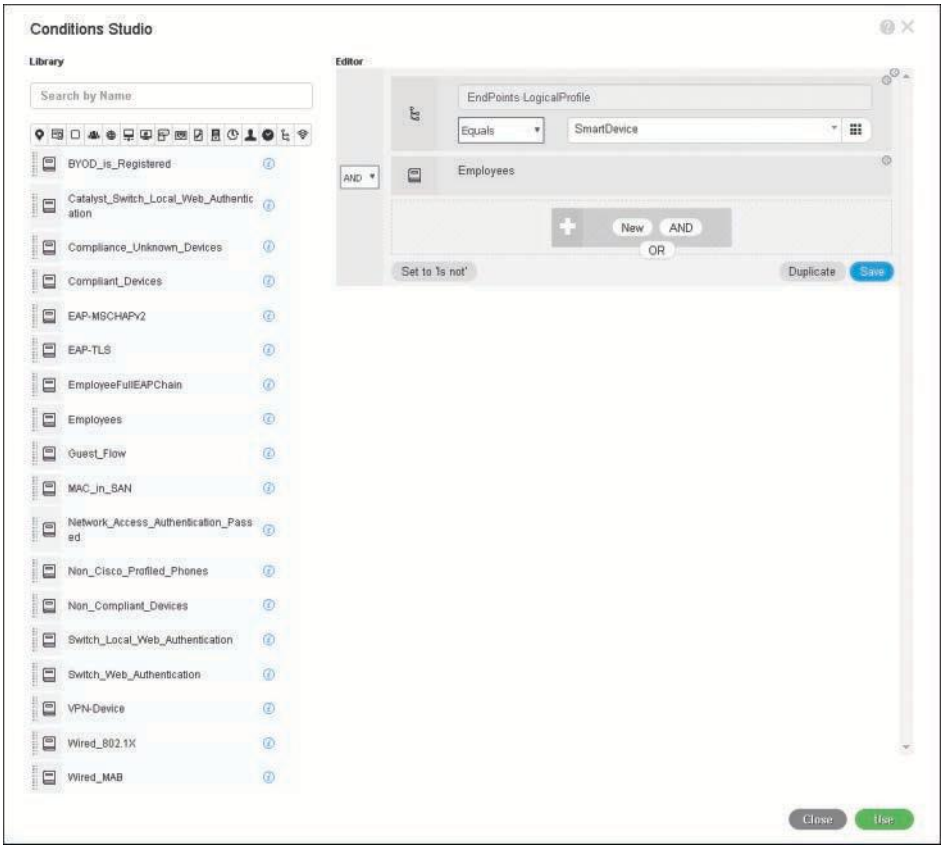


Figure 10-18 *Saving the Employees Condition to the Library*

Step 7. Click **Use** to finish editing the condition.

Step 8. Click **Save**.

Figure 10-19 shows the final authorization policy.

In Figure 10-19, it is easy to see how saving the raw attributes as simple or compound conditions can enhance the readability of the policies. Note that these reusable conditions exist for many policies in ISE, not only the authorization policy.

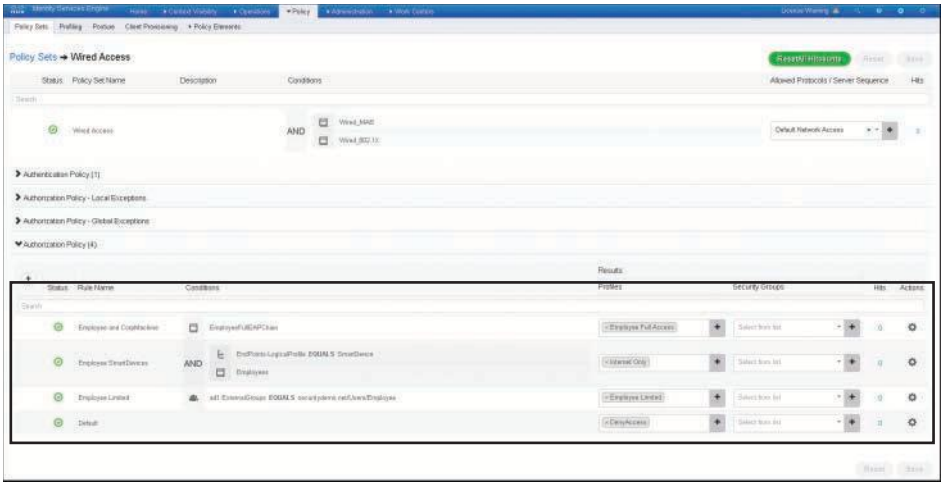


Figure 10-19 *Final Authorization Policy*

Combining AND with OR Operators

When creating policy rules in an enterprise, often the conditions that need to be met are not as binary as a single AND or OR operator. In such situations, you can create conditions blocks in the Conditions Studio to form hierarchical compound conditions to meet these needs. You can create a policy that is as flexible or complex as needed. For example, if you need to create a rule that says that the user must be in at least one of two Active Directory groups *and* needs to be connecting via wired 802.1X, you can accomplish this with a single authorization rule and hierarchical compound conditions.

For this example, you will build an authorization rule in the default policy set. This rule will be configured to allow full access to a user who is a member of either the Domain Admins or IT-Users Active Directory groups and who is also connecting to the network via wired 802.1X.

To create this authorization rule, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the default policy set.

Step 3. Expand the authorization policy.

Step 4. Click on the cog to the right side of the Wireless Black List Default authorization rule and choose **Insert New Row Above**.

Step 5. Name this new authorization rule **IT Users Access**.

Step 6. Click on the + button to open the Conditions Studio.

Step 7. Notice that a single blank condition is ready to be configured in the editor. In order to accomplish the goals of the policy, you need an AND operator on the top of the condition hierarchy, and you need to add a condition block within that hierarchy to use the OR operator. Click the **AND** button in the editor to create the first condition block.

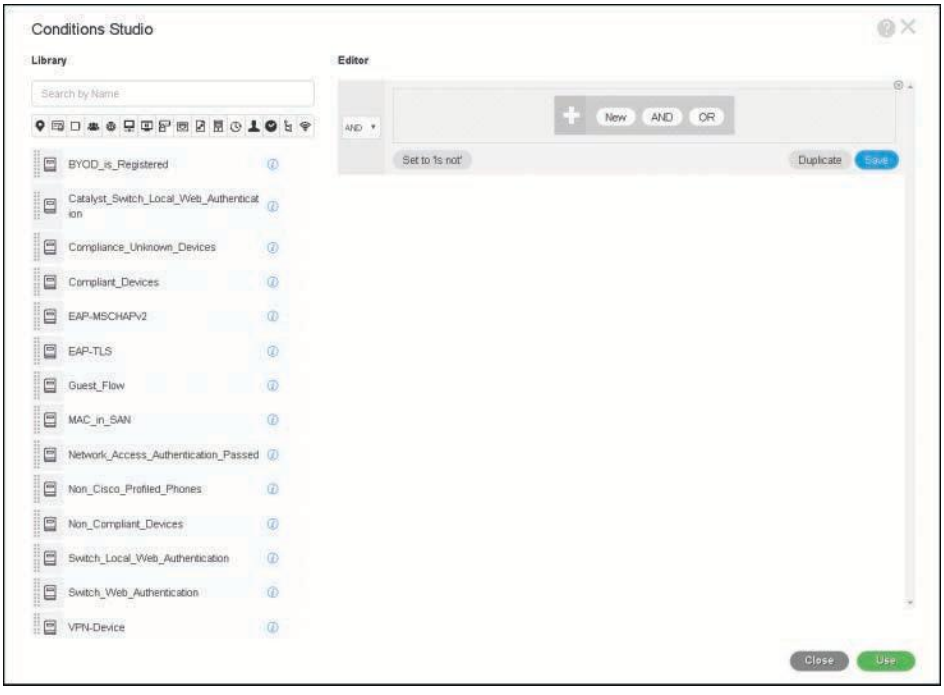


Figure 10-20 *Creating a New Condition Block*

Figure 10-20 shows the newly created conditions level. The AND operator is established at the top of the condition hierarchy. To the right of that operator is the first condition block. You can now either add a condition to this block or add another sublevel to the hierarchy.

Step 8. Click the **OR** button within the condition block to create a new level under the initial condition block. Figure 10-21 shows the new level in the hierarchy. The rules under this level will be evaluated against the OR operator.

Step 9. Click the **New** button twice in the newly created level to add two new conditions.

Step 10. For the first condition, click on attribute to display the attribute menu, choose the Active Directory domain name from the dictionary dropdown, and select the **ExternalGroups** as the attribute. (Another way to pull it up quickly is to use the filter at the top of the attribute selection window. You can type part of the attribute name, such as **external**, to find **ExternalGroups** in the filtered results.)

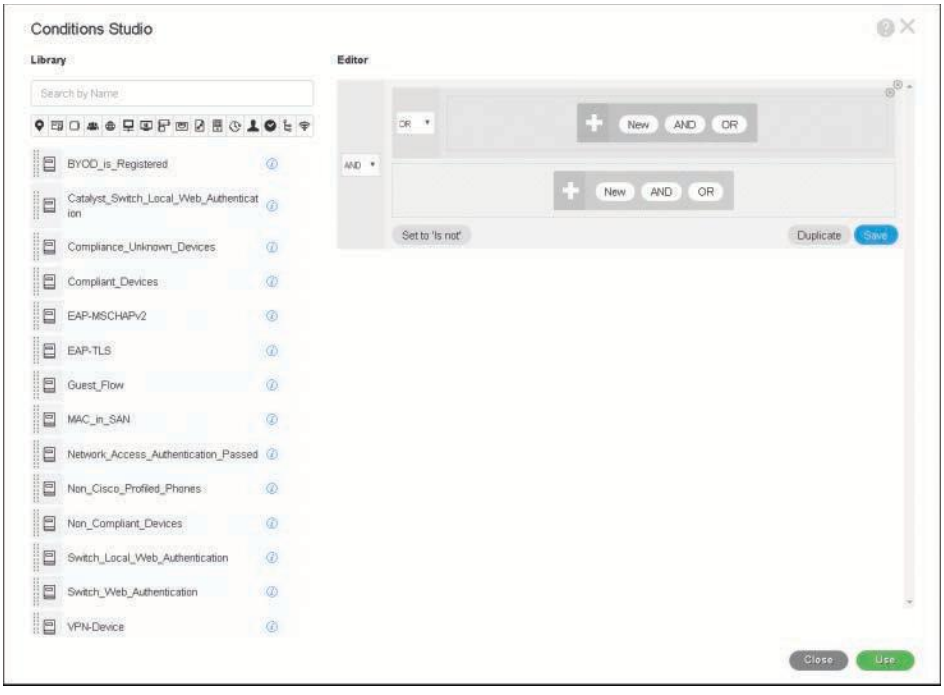


Figure 10-21 *Adding a Level to the Condition Hierarchy*

Step 11. Select **Equals** as the operator.

Step 12. Select **IT-Users** as the attribute value and **ExternalGroups** as the attribute.

Step 13. Select **Equals** as the operator.

Step 14. Select **Domain Admins** as the attribute value. Figure 10-22 shows the newly created compound condition block populated with the attributes. The logic of this block is simple: In order for this rule to match, the user must be part of either the IT-Users AD group or the Domain Admins AD group.

Step 15. From the library, drag the **Wired_802.1X** saved condition into the separate empty condition block. Figure 10-23 shows the completed hierarchical compound condition blocks.

Step 16. Click **Use**.

Step 17. For the authorization result, select the prebuilt **PermitAccess** authorization profile from the dropdown.

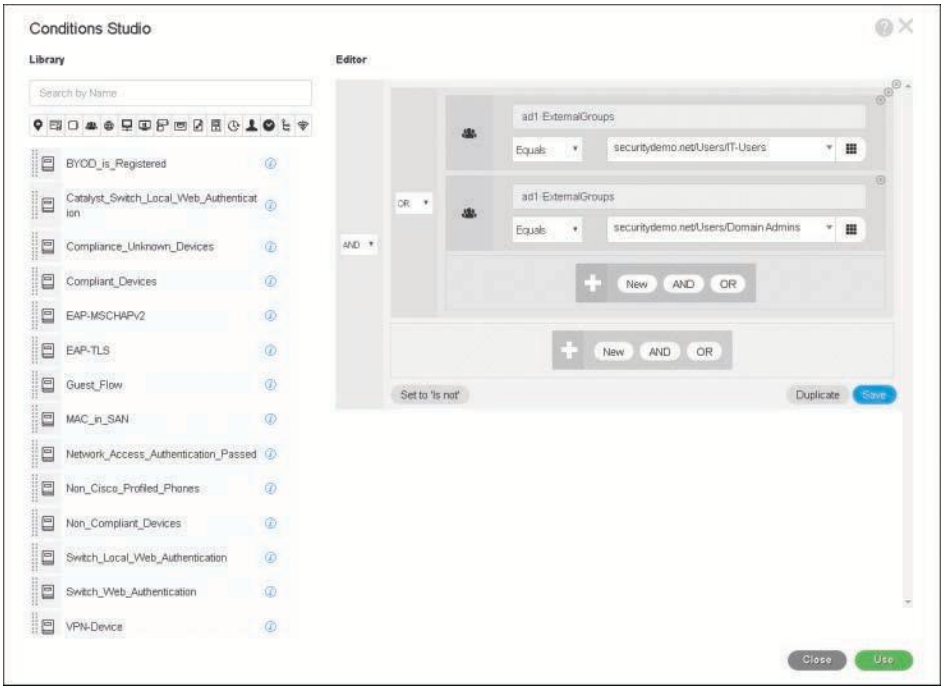


Figure 10-22 *Compound Condition Block*

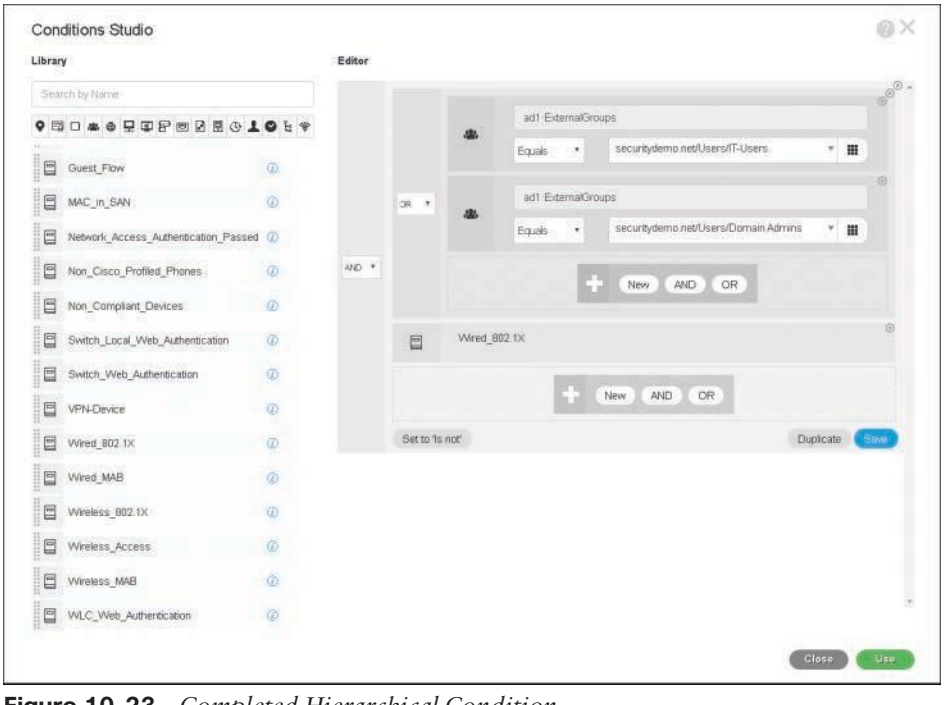


Figure 10-23 *Completed Hierarchical Condition*

Figure 10-24 shows the completed authorization rule. The logic is simple here: For this authorization rule to be selected, the user must be a member of either the IT-Users AD group or the Domain Admins AD group, and the user must be connecting to the network via 802.1X wired access.

In this chapter you have seen just one example of how flexible and customizable ISE policy condition can be. With compound and multilevel conditions, you can create almost endless condition variations to meet an enterprise's security requirements.



Figure 10-24 *Completed Authorization Rule*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 10-2 lists these key topics and the page number on which each is found.

Key Topic

Table 10-2 Key Topics

Key Topic Element	Description	Page
Section	Goals of authorization policies	235
Section	Role-specific authorization rules	241
Section	Saving conditions for reuse	249

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

security context, simple condition, compound condition

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What are the goals of an authorization policy?
2. What kind of conditions can be added to an authorization policy?
3. What is the basic logic of an authorization policy?
4. Does authentication make a network segmented and secure?
5. What is the goal in using hierarchical compound conditions?

This page intentionally left blank



CHAPTER 11

Implement Wired and Wireless Authentication

This chapter covers the following topics:

- Configuring network access devices for wired/wireless 802.1X
- AV pairs
- Implement MAB
- Implement network authorization enforcement
- Troubleshooting, monitoring, and reporting tools

In Chapter 8, "Initial Configuration of Cisco ISE," during the initial configuration of Cisco ISE, you added a ***network access device (NAD***) definition to ISE and were introduced to the concept of importing multiple NADs defined in a comma-separated values (CSV) file. In Chapter 9, "Authentication Policies," and Chapter 10, "Authorization Policies," you learned how to define some basic authentication and authorization policies in ISE.

This chapter discusses one of the most important aspects of secure network access: the configuration of the NADs themselves (switches, wireless controllers, and so on), also commonly referred to as the *access-layer devices*.

A NAD is the point of access for an endpoint. It's the first network device through which an endpoint communicates. It is therefore a critical part of any secure network architecture, and it is the main policy enforcement point.

When an organization begins project planning for a secure network access project (often called a network access control [NAC] project), the basic use cases are always considered—items such as how corporate assets will authenticate to the network and the plan for guest users. A lab can be built, and technology can be tested in a proof-of-concept environment (usually at small scale). If everything ends up working, the proof of concept is considered successful, and everyone celebrates the win.

However, a project team may neglect many important considerations until it truly starts drilling into the deployment at scale—for instance, how to handle WAN outages when the policy server is centrally located, how to handle the desktop reimaging process, how to use thin clients that don't have a local operating system but boot to the network and download an OS or even connect to a centrally located terminal server. (How will devices get access to the network if they cannot authenticate, and how does the team ensure that these devices can

reach the appropriate locations to download their OS or the reimaging server while restricting access to anything unnecessary?) The policy server is obviously a lynchpin in a successful deployment. However, intelligence in the access-layer device is crucial and can make or break an entire project.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 11-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 11-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Authentication Configuration on Wired Switches	1, 4–8
Authentication Configuration on WLCs	2
Verifying Dot1x and MAB	3, 9, 10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. When configuring a Cisco switch for 802.1X, at what level of the configuration do the 802.1X-related commands exist?
   - **a.** Global configuration only
   - **b.** Interface configuration only
   - **c.** Both global configuration and interface configuration
   - **d.** A dot1x subconfiguration mode where all related commands are entered
2. When configuring a Cisco WLC for communication with ISE, what must be configured for the WLAN? (Choose two.)
   - **a.** The authentication and authorization RADIUS servers may be pointed to different ISE PSNs, as long as those PSNs are part of a node group.
   - **b.** The authentication and accounting RADIUS servers must be pointed to the same ISE PSN.
   - **c.** The WLAN must be configured for SNMP NAC.
   - **d.** The WLAN must be configured for RADIUS NAC.

3. If a switch is configured to send syslog messages, which ISE persona should the switch send these messages to?
   a. Administration
   b. Monitoring
   c. Policy Services
   d. pxGrid
4. What is the purpose of adding a user with the **username** *radius-test* **password** *password* command?
   a. The switch can send periodic RADIUS Access-Request messages to the AAA servers in order to verify whether they are still alive. The username and password are used for that test.
   b. The username and password are used for the local RADIUS server available in the switch, which is used in WAN down scenarios.
   c. The username and password are used for the supplicant's outer identity to authenticate against the switch's local users database.
   d. Without the local username and password in the configuration, an administrator may be locked out of the switch when the RADIUS server is unavailable.
5. What feature must be enabled in IOS in order for dACLs to work with a switch?
   a. EPM logging
   b. Device tracking
   c. dot1x system-auth-control
   d. RADIUS VSAs
6. Which of the following technologies enables an administrator to maintain the same configuration on all access ports on all switches, regardless of the type of device connecting to the network?
   a. AnyConnect
   b. Multi-auth
   c. FlexAuth
   d. FlexConnect
7. Which host mode permits a virtually unlimited number of endpoints per port, allowing all subsequent MAC addresses to share the authorization result of the first endpoint authorized?
   a. Single-Host
   b. MDA
   c. Multi-Auth
   d. Multi-Hosts

8. Which interface-level command basically turns on authentication?
   a. **authentication port-control auto**
   b. **dot1x system-auth-control**
   c. **ip device-tracking**
   d. **aaa server radius dynamic-author**
9. Which command on a Cisco switch displays the current status of the AAA server(s)?
   a. **show authentication servers**
   b. **show radius servers**
   c. **show aaa servers**
   d. **show ise servers**
10. Which command validates that authentications are being attempted, shows which authentications are successful, and shows what authorization results have been assigned?
   a. **show authentication method dot1x**
   b. **show aaa servers**
   c. **show authentication statistics**
   d. **show authentication session interface** *interface*

Foundation Topics

Authentication Configuration on Wired Switches

With Cisco switches, numerous configuration sections must be completed. Some aspects must be configured from global configuration mode, and others are configured at the individual interfaces. An authentication subsystem must be enabled across an entire switch (global configuration), and there are interface configurations that must be entered for authentication settings (for timers, enabling 802.1X communication, and more).

Global Configuration AAA Commands

Key Topic

By default, the AAA subsystem of the Cisco switch is disabled. Until the AAA subsystem is enabled, none of the required commands are available as configuration options. To configure the global AAA configuration on a Cisco switch, follow these steps:

Step 1. Enable authentication, authorization, and accounting on the access switch(es) by entering the following command:

```
C3560X(config)# aaa new-model
```

An authentication method is required to tell the switch which group of RADIUS servers to use for 802.1X authentication requests.

Step 2. Create an authentication method for 802.1X by entering the following command:

```
C3560X(config)# aaa authentication dot1x default group
radius
```

The method created in this step enables the user/device identity (username/password or certificate) to be validated by the RADIUS server. However, simply having valid credentials is not enough. There must be an authorization as well. The authorization indicates that the user or device is actually allowed to access the network as well as the level of access permitted.

Step 3. Create an authorization method for 802.1X by entering the following command:

```
C3560X(config)# aaa authorization network default group
radius
```

RADIUS accounting packets are extremely useful and, in many cases, are required. These types of packets help ensure that the RADIUS server (ISE) knows the exact state of the switch port and endpoint. Without the accounting packets, ISE would have knowledge only of the authentication and authorization communication. Accounting packets provide information on when to terminate a live session as well as local decisions made by the switch (such as AuthFail, VLAN assignment, and so on).

RADIUS accounting packets can and should include the Framed-IP-Address field, which is populated with the IP address of the endpoint and updates the RADIUS server (ISE) so it has the ability to correlate the MAC address (which it has from the authentication request) to the IP address that is most often dynamically assigned. If the switch supports IOS Device Sensor, the sensor data is sent to ISE using the RADIUS accounting configuration.

Step 4. Create an accounting method for 802.1X by entering the following command:

```
C3560X(config)# aaa accounting dot1x default start-stop
group radius
```

Global Configuration RADIUS Commands

There was a change in the way RADIUS servers are configured in IOS between Version 12.2.x and Version 15.x and IOS XE code. Where the configuration differs in the following steps, it is broken out into separate sections.

IOS 12.2.x

Cisco switches provide a proactive method to check the availability of the RADIUS server. If this option is configured, the switch sends periodic test authentication messages to the RADIUS server (ISE). The switch is simply looking for a RADIUS response from the server. A success message is not necessary; a failed authentication suffices because that failed authentication response confirms that the server is alive.

If you want to receive an Access-Accept response, the username that is configured on the switch needs to already exist in either ISE or another identity store. This is an optional step if log filtering is not enabled. You might consider this step if you want to reduce the number of failed authentications in ISE logs. Because the RADIUS test continuously runs, it may show a large number of logged failed authentications. This account will be used in a later step, where you define the RADIUS server. You may also configure log filtering as well to reduce the amount of failed authentications from the configured test authentication on the switch if you do not wish to create an account in ISE or another identity store for this purpose.

To configure the RADIUS server for a Cisco switch using IOS Version 12.2, follow these steps:

Step 1. In global configuration mode, add a username and password for the RADIUS keepalive by entering the following command:

```
C3560X(config)# username radius-test password password
```

Key Topic

Step 2. Add the Cisco ISE servers to the RADIUS group by entering the following command:

```
C3560X(config)# radius-server host ise_ip_address
auth-port 1812 acct-port 1813 test username radius-test
key shared_secret
```

In step 2, you are adding each ISE Policy Services node (PSN) to the switch configuration using the test account you created in step 1. The server will proactively be checked for responses one time per hour, in addition to any authentications or authorizations occurring through normal processes.

Step 3. Repeat step 2 for each PSN.

The switch has been configured to proactively check the Cisco ISE server for RADIUS responses. Now you can configure the counters on the switch to determine whether the server is alive or dead. The configuration steps for this do not differ between IOS versions and are provided in the section "IOS 12.2.x, 15.x, and IOS XE," later in this chapter.

IOS 15.x and IOS XE

Cisco switches provide a proactive method to check the availability of the RADIUS server. If this option is configured, the switch sends periodic test authentication messages to the RADIUS server (ISE). The switch is simply looking for a RADIUS response from the server. A success message is not necessary; a failed authentication suffices because that failed authentication response confirms that the server is alive.

If you want to receive an Access-Accept response, the username that is configured on the switch needs to already exist in either ISE or another identity store. This is an optional step. You might consider this step if you want to reduce the number of failed authentications in ISE logs. Because the RADIUS test continuously runs, it may show a large number of logged failed authentications. This account will be used in a later step, where you define the RADIUS server.

To configure the RADIUS server for a Cisco switch using IOS Version 15 and above or IOS-XE, follow these steps:

Key Topic

Step 1. In global configuration mode, add a username and password for the RADIUS keepalive by entering the following command:

```
C3850(config)# username radius-test secret password
```

Step 2. Add the Cisco ISE server by entering the following command:

```
C3560X(config)# radius server [name]
```

In step 2, you are adding each ISE PSN to the switch configuration using the test account you created in step 1. This is where the configuration really differs between IOS 12.2.x and IOS 15.x. In step 2 for IOS 15.x and IOS XE, you

create a new RADIUS server object. The configuration related to that server will be contained within this object. Next, you configure the IP address, ports, and shared secret.

Step 3. Configure the IP address, authentication port, and accounting port by entering the following command:

```
C3560X(config-radius-server)# address ipv4 [ip-address]
auth-port 1812 acct-port 1813
```

Step 4. Configure the shared secret by entering the following command:

```
C3560X(config-radius-server)# key [shared-secret]
```

Using the test account you created in step 1, the server is proactively checked for responses one time per hour in addition to any authentications or authorizations occurring through normal processes. Repeat step 2 if there are multiple PSNs you would like to configure on the switch to fall back upon if the first is not available.

Step 5. Configure the switch to test the RADIUS server object by entering the following command:

```
C3560X(config-radius-server)# automate-tester username
[username-from-step-1]
```

NOTE If you are using a Cisco 3850, 3650, 9300, 9500, or 5760 Series switch, the network access device sends its Calling-Station-Id in the format of *xxxx.xxxx.xxxx* by default, instead of *xx-xx-xx-xx-xx-xx*. You can optionally change this with the **radius-server attribute 31 mac format** global configuration command.

Step 6. (Optional) Configure the format for the switch to send the Calling-Station-Id by entering the following command:

```
C3560X(config)# radius-server attribute 31 mac format
ietf lower-case
```

IOS 12.2.x, 15.x, and IOS XE

You have already configured the switch to proactively check the ISE server for RADIUS responses. Now you can configure the counters on the switch to determine whether the server is alive or dead. In this section, you will see how to configure the switch to wait 5 seconds for a response from the RADIUS server and attempt the test three times before marking the server dead. If that RADIUS server doesn't send a valid response within 15 seconds, it is marked as dead. (High availability is covered in more detail in Chapter 20, "ISE Scale and High Availability.")

To configure the RADIUS settings on the Cisco switch, follow these steps:

Step 1. Set the dead criteria by entering the following command:

```
C3560X(config)# radius-server dead-criteria time 5 tries 3
```

Key Topic

Step 2. Enable *Change of Authorization (CoA)* by entering the following command:

```
C3560X(config)# aaa server radius dynamic-author
C3560X(config-locsvr-da-radius)# client ise_ip_address
server-key shared_secret
```

You previously defined the IP address of a RADIUS server to which the switch will send RADIUS messages. However, you must define the servers that are allowed to perform CoA (see RFC 3576 and RFC 5176) operations in a different configuration that is also in global configuration mode.

Here you configure the switch to send any defined ***vendor-specific attributes (VSAs)*** to ISE PSNs during authentication requests and accounting updates.

Step 3. Configure the switch to use the Cisco VSAs by entering the following command:

```
C3560X(config)# radius-server vsa send authentication
C3560X(config)# radius-server vsa send accounting
```

NOTE In later versions of IOS 15.x and IOS XE, the VSAs are configured to send to the RADIUS server by default. In this case, these commands may not appear in the normal running configuration after step 3. Certain default configurations are hidden in the normal **show running-config** output. In order to verify that VSAs are being sent to the RADIUS server, issue the **show running-config all | inc radius-server** command.

Step 4. Enable the vendor-specific attributes (VSAs) by entering the following command:

```
C3560X(config)# radius-server attribute 6 on-for-
login-auth
C3560X(config)# radius-server attribute 8 include-in-
access-req
C3560X(config)# radius-server attribute 25 access-request
include
```

These are the three VSAs you are enabling here:

- **Service-Type (RADIUS attribute 6):** This VSA ensures that login requests will include the Service-Type attribute, which indicates to ISE which authentication method was chosen for the switch port and, ultimately, the endpoint:
  - Service-Type:Framed = 802.1X
  - Service-Type:Call-Check = MAB
  - Service-Type:Login = Local WebAuth
- **Framed-IP-Address (RADIUS attribute 8):** This VSA ensures that the IP address is included (if one exists). ISE can use this IP address to correlate the user with other ISE profiling techniques (discussed further in Chapter 14, "Profiling").

- **Class (RADIUS attribute 25):** This VSA should be sent in access request messages. This RADIUS attribute can be used for a number of features, including NAC and VPNs.

Switches often have multiple IP addresses associated with them. Therefore, it is a best practice to always force any management communications to occur via a specific interface. The interface IP address must match the IP address defined in the ISE network device object that was created to represent the switch.

The applicable traffic that can be sourced may be RADIUS and SNMP. SNMP is optional and used mostly for sending SNMP traps or information for endpoint profiling.

Step 5. Ensure that the switch always sends traffic from the correct interface by entering the following command:

```
C3560X(config)# ip radius source-interface interface_name
C3560X(config)# snmp-server trap-source interface_name
C3560X(config)# snmp-server source-interface informs
interface_name
```

Global 802.1X Commands

Key Topic

Enabling 802.1X globally on a switch does not actually enable authentication on any of the switch ports. Authentication is configured, but it is not enabled until the interface configuration occurs.

To configure the 802.1x globally and IP device tracking, follow these steps:

Step 1. Enable 802.1X globally on the switch.

```
C3560X(config)# dot1x system-auth-control
```

Downloadable access control lists (dACLs) are a very common enforcement mechanism in ISE deployments. In order for dACLs to function properly on a switch, a function called IP device tracking must be enabled globally. The **dot1x system-auth-control** command allows for the **any** source in the provided dACL to be replaced with the IP address of the single device connected to the switch port.

Step 2. Enable dACLs to function for IOS Version 12.2 and Version 15.x by entering the following command (although it is enabled by default in 15.x version):

```
C3560X(config)# ip device tracking
```

NOTE In IOS Version 15.x, this command is enabled by default and may be seen in the output of **show running-config all | inc ip device tracking**.

Step 3. IP device tracking discovers the IP address for an endpoint using an ARP probe. If the switch sends out the ARP probe for the client while the client is in its duplicate address detection phase, the client detects a duplicate IP address on the network. This causes a popup to display on Windows machines, stating that

a duplicate IP address has been detected on the network. This would create a lot of confusion for the end users. To prevent this issue from happening, there was an additional recommended command to go along with enabling IP device tracking.

In IOS 15.2(2) or IOS XE 03.06.00 and later versions, the **auto-source** command was added for finding another suitable source to use:

```
C3560X(config)# ip device tracking probe auto-source
```

For earlier IOS/IOS XE releases, the **probe delay** command is used to create a delay of 10 seconds when the switch detects a link flap, which minimizes the possibility of the probe being sent while the host is checking for a duplicate IP address:

```
C3560X(config)# ip device tracking probe delay 10
```

Step 4. (Optional) Verify that IP device tracking is working by entering the following command to display the discovered IP addresses:

```
C3560X# show ip device tracking all
```

Device Tracking in IOS XE 16.x and Later

For switches using IOS XE Version 16.x and later, the legacy IP device tracking feature was removed and replaced with Switch Integrated Security Features (SISF) tracking, which is part of a suite of first-hop security features. The result of this change is that device tracking is no longer enabled by default, and the commands to enable it changed considerably. IOS XE now requires the creation of a SISF policy to enable tracking, and that policy must be attached to the applicable interfaces.

To configure device tracking in IOS-XE 16.x and later, follow these steps:

Step 1. Create the SISF policy globally by entering the following command:

```
CAT9300(config)# device-tracking policy policy-name
```

Step 2. Enable tracking in the policy by entering the following command:

```
CAT9300(config-device-tracking)# tracking enable
```

Step 3. Attach the newly created policy to the interface at the interface-level configuration by entering the following command:

```
CAT9300(config)# interface g1/0/1
CAT9300(config-if)# device-tracking attach-policy
policy-name
```

Step 4. In the global configuration, configure device tracking to auto-source the probe by entering the following command:

```
CAT9300(config)# device-tracking tracking auto-source
```

Step 5. (Optional) Enter the following command to verify that the device tracking policy is working by displaying the discovered IP addresses:

```
CAT9300# show device-tracking database details
```

Creating Local Access Control Lists

Certain functions, such as URL redirection, require the use of locally configured access control lists (ACLs) on a switch. Some of these ACLs are used immediately, and some may not be used until a much later phase of deployment. This section shows how to prepare the switches for all possible deployment models at one time and limit the operational expense of repeated switch configuration.

NOTE The details of these ACLs and why they are created this way are discussed in future chapters for configuring web authentication and deployment phasing.

To configure the common local access control lists that ISE may invoke for various policies, follow these steps:

Step 1. Add the following ACL to be used on switch ports in Monitor mode:

```
C3560X(config)# ip access-list extended ACL-ALLOW
C3560X(config-ext-nacl)# permit ip any any
```

Step 2. Add the following ACL to be used on switch ports in Low-Impact mode:

```
C3560X(config)# ip access-list ext ACL-DEFAULT
C3560X(config-ext-nacl)# remark DHCP
C3560X(config-ext-nacl)# permit udp any eq bootpc any eq
bootps
C3560X(config-ext-nacl)# remark DNS
C3560X(config-ext-nacl)# permit udp any any eq domain
C3560X(config-ext-nacl)# remark Ping
C3560X(config-ext-nacl)# permit icmp any any
C3560X(config-ext-nacl)# remark PXE / TFTP
C3560X(config-ext-nacl)# permit udp any any eq tftp
C3560X(config-ext-nacl)# remark Drop all the rest
C3560X(config-ext-nacl)# deny ip any any log
```

Step 3. Add the following ACL to be used for URL redirection with Web Authentication:

```
C3560X(config)# ip access-list ext ACL-WEBAUTH-REDIRECT
C3560X(config-ext-nacl)# remark explicitly deny DNS from
being redirected to address a bug
C3560X(config-ext-nacl)# deny udp any any eq 53
C3560X(config-ext-nacl)# remark redirect all applicable
traffic to the ISE Server
C3560X(config-ext-nacl)# permit tcp any any eq 80
C3560X(config-ext-nacl)# permit tcp any any eq 443
C3560X(config-ext-nacl)# remark all other traffic will be
implicitly denied from the redirection
```

Step 4. Add the following ACL to be used for URL redirection with the posture agent:

```
C3560X(config)# ip access-list ext ACL-AGENT-REDIRECT
C3560X(config-ext-nacl)# remark explicitly deny DNS and
DHCP from being redirected
C3560X(config-ext-nacl)# deny udp any any eq 53 bootps
C3560X(config-ext-nacl)# remark redirect HTTP traffic only
C3560X(config-ext-nacl)# permit tcp any any eq 80
C3560X(config-ext-nacl)# remark all other traffic will be
implicitly denied from the redirection
```

Interface Configuration Settings for All Cisco Switches

So far in this chapter you have completed the global configuration RADIUS and AAA settings on the access-layer switches. This section focuses on building a single-port configuration that can be used across an entire secure unified access deployment, regardless of the switch type, deployment stage, or deployment model you choose.

Configure Interfaces as Switch Ports

One of the first things to do before configuring any of the authentication settings on a switch port is to ensure that the switch port is configured as a Layer 2 port instead of a Layer 3 port. The following steps detail how to configure a switch port, including ensuring that it is a Layer 2 port:

Step 1. Enter interface configuration mode for the switch port range:

```
C3560X(config)# interface range first_interface -
last_interface
```

Step 2. Ensure that the ports are Layer 2 switch ports by entering the following command:

```
C3560X(config-if-range)# switchport
```

Step 3. Configure the port as a Layer 2 edge port by entering the following command:

```
C3560X(config-if-range)# switchport host
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
```

The switch has some predefined macros that run a set of commands in a particular order. The **host** macro automatically runs three commands for you: It configures the port to be an access port (non-trunk), disables channel groups, and configures spanning tree to be in portfast mode.

Configure Flexible Authentication and High Availability

The default behavior of 802.1X is to deny access to the network when an authentication fails. This behavior was discovered to be undesirable in many customer deployments because it does not allow for guest access, and it also does not allow employees to remediate their computer systems and gain full network access. The next phase in handling 802.1X authentication failures was to provide an Auth-Fail VLAN to allow a device or user that failed

authentication to be granted access to a VLAN that provided limited resources. This was a step in the right direction, but it was still missing some practicality, especially in environments that must use ***MAC Authentication Bypass (MAB)*** for all the printers and other non-authenticating devices. With the default behavior of 802.1X, administrators would have to use different configurations for ports for printers and other devices that do not have supplicants and for ports where they planned to do authentication.

Key Topic

Therefore, Cisco created ***Flexible Authentication (FlexAuth)***, which allows a network administrator to set an authentication order and priority on a switch port. It allows a port to attempt 802.1X, MAB, and then WebAuth in a configured order. All these functions are provided with a single configuration on all access ports, which provides a much simpler operational model for customers than traditional 802.1X deployments.

As mentioned previously, there are multiple methods of authentication available on a switch port: 802.1X (also called *dot1x*), MAB, and Web Authentication (WebAuth). For the purpose of the wired deployment we are looking at in this section, we focus on the first two methods: 802.1X and MAB. With 802.1X authentication, the switch sends an identity request (EAP-Identity-Request) periodically after the link state has changed to up. In addition, the endpoint supplicant periodically sends EAP over LAN Start (EAPoL-Start) messages to the switch port to speed up authentication. If a device is not able to authenticate, it must wait until the dot1x timeout occurs, and MAB occurs. Assuming that the device MAC address is in the correct database, the device is then authorized to access the network. Figure 11-1 illustrates this concept.

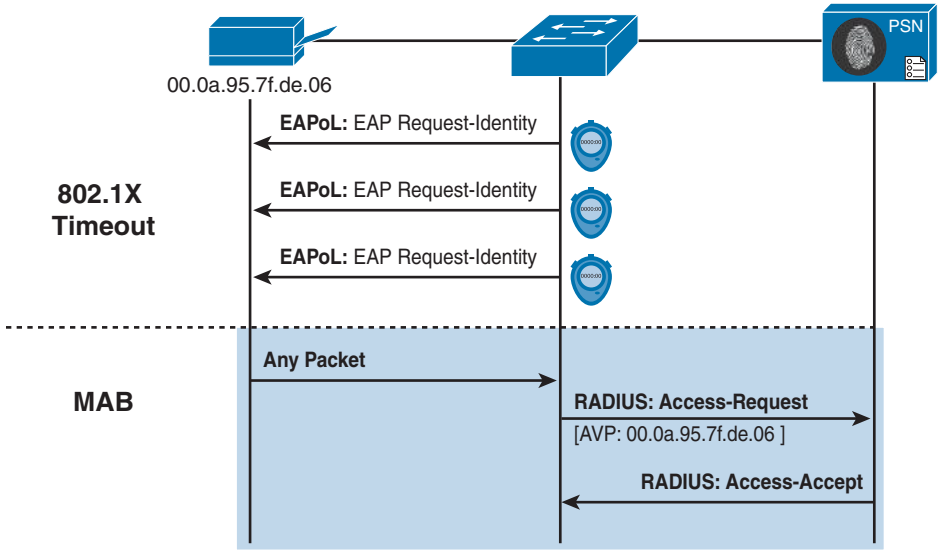


Figure 11-1 *Flexible Authentication*

NOTE With FlexAuth, you can toggle whether dot1x or MAB has the higher priority. The best practice is to always prefer the stronger authentication method, which is dot1x. The dot1x method is also the default on all Catalyst Cisco switches.

To configure FlexAuth and the configurable actions for authentication high availability, follow these steps:

Step 1. Configure the authentication method priority on the switch ports by entering the following command:

```
C3560X(config-if-range)# authentication priority dot1x
mab
```

NOTE With certain deployment methods, where MAB should occur before 802.1X authentication (for example, in an environment where the majority of devices are non-authenticating or where EasyConnect is being used to authenticate the user). For these environments, FlexAuth allows for a network administrator to set a user-definable authentication order. However, the best practice is to maintain the order dot1x and then MAB. Keep in mind that if the order were to be MAB followed by dot1x with the priority being dot1x followed by MAB, any 802.1X communication would stop the MAB and immediately switch to an 802.1X transaction. The configured order on a switch port simply tells the switch port which authentication method to attempt first. However, the priority tells the network access device which authentication method to prefer over others if it takes place on the switch port. This is an important difference to understand.

Step 2. Configure the authentication method order on the switch ports by entering the following command:

```
C3560X(config-if-range)# authentication order dot1x mab
```

Just setting the order and priority alone is not what actually provides the FlexAuth technology. There is still a key command that must be enabled—the command that dictates how the switch should behave when the RADIUS server (ISE) sends an Access-Reject response. The **authentication event fail action** [*action*] command determines what should occur in the event of an Access-Reject message being received.

One option could be to follow the original design of dot1x and deny access to the network. Another option is to assign a VLAN. This would be known as using an Auth-Fail VLAN. In the early days of dot1x, this was how guest access was handled. If the authentication failed, the switch would then assign the Guest VLAN. It's very important to note that in an ISE secure access deployment, **fail action** should be set to **next-method**. This is the third and final piece of the puzzle to make up the FlexAuth technology.

NOTE If a switch makes a local decision, such as assigning a VLAN instead of denying access, the switch no longer accepts CoA commands from the RADIUS server. Think about it this way: ISE now thinks that the switch port is down, but the switch has actually permitted the endpoint to have access to the guest VLAN, so the state is now "out of sync."

Step 3. Configure the port to use FlexAuth by entering the following command:

```
C3560X(config-if-range)# authentication event fail action
next-method
```

Earlier in this chapter, you configured the RADIUS server entry to use a test account that proactively alerts the switch when Cisco ISE has stopped responding to RADIUS requests. Now you can configure the switch port to locally authorize the port when that server is found to be dead and re-initialize authentication when the server becomes alive again. (The topic of high availability is covered in more detail in Chapter 20.)

Step 4. Configure the port to use a local VLAN for voice and data when the RADIUS server is dead (that is, when it stops responding) by entering the following commands:

```
C3560X(config-if-range)# authentication event server dead
action authorize vlan vlan-id
```

```
C3560X(config-if-range)# authentication event server dead
action authorize voice
```

```
C3560X(config-if-range)# authentication event server
alive action reinitialize
```

Historically, all authenticated hosts (when the RADIUS server was **up**) got full access to the network, and the others (the new hosts) did not get access to the network. However, there were concerns regarding multiple authenticating hosts on a single port when a portion of them had already authenticated while the RADIUS server was operational and others (new hosts) were trying to authenticate when the RADIUS server was down.

Cisco added the IOS command **authentication event server dead action reinitialize vlan** *vlan-id* to ensure that the port is re-initialized immediately when new hosts try to access the network while the RADIUS server is down and all the hosts on this port are placed into the same VLAN.

Step 5. (Optional) Configure the port to use a local VLAN when the RADIUS server is dead and allow existing and new hosts by entering the following command:

```
C3560X(config-if-range)# authentication event server dead
action reinitialize vlan vlan-id
```

Host Mode of the Switch Port

Key Topic

The default behavior of an 802.1X-enabled port is to authorize only a single MAC address per port. By design, the standard originally relied heavily on the link state of the connection to determine the beginning and end of the authentication session. Since the inception of 802.1X, there have been enhancements to allow one or more devices per port. The following host modes are available:

- **Single-Host**: ***Single-Host*** mode is the default mode of 802.1X. Only a single MAC address is permitted per interface. In the early days of 802.1X, the protocol expected to rely on the link state for authentication control. When the link comes up, the authenticator (switch) starts sending the EAPoL-ID-Request messages to an endpoint. When the link drops, the session is killed, and a RADIUS accounting stop message is sent to the RADIUS server.

- **Multidomain Authentication (MDA):** *MDA* extends 802.1X to support IP telephony by modifying 802.1X to understand two separate domains: a data domain and a voice domain. A switch port in MDA mode allows a single MAC address in each domain (that is, one phone and one endpoint behind the phone). It's important to note that link state cannot be relied on completely because the switch port has more than one device.

  Cisco has enhanced Cisco IP Phones to support ***CDP Second Port Disconnect*** messages from the phone to the switch. These messages inform the switch of the endpoint unplugging and indicate to terminate any authentications for that endpoint. This works for all authentication types (dot1x, MAB, and WebAuth).

- **The 802.1X standard supports a different method, known as EAPoL-Proxy-Logoff:** *EAPoL-Proxy-Logoff* allows an IP Phone to issue a logoff method for the endpoint supplicant that was authenticated behind the phone. However, this works only for dot1x authentications and not for MAB or WebAuth.

- **Multiauthentication (Multi-Auth):** *Multi-Auth* is an extension to MDA. It supports the dual-domain solution while allowing virtually unlimited endpoints to be in the data domain. Each endpoint is required to authenticate, and each authentication session on the port is managed separately. The IP Phone solutions described for MDA are also applicable in Multi-Auth mode.

- **Multiple Hosts (Multi-Hosts):** *Multi-Hosts* mode is not commonly used but is a valid option. Much like Multi-Auth mode, Multi-Hosts mode is an extension to MDA. There is one authentication on the voice domain and one authentication on the data domain. All other hosts on the data domain are allowed onto the network using the first successful authentication. It's an "authenticate one, allow the rest" model. This model is not often adopted in deployments due to the fact that the whole idea around network access control is to ensure that all the endpoints connecting to the network are authenticated. The IP Phone solutions described for MDA are also applicable to Multi-Auth mode.

NOTE Using *Port Security* to limit the number of MAC addresses allowed on a port is not compatible with 802.1X because 802.1X handles this function natively.

During the initial phases of any secure access deployment, it is often desirable to use Multi-Auth mode to ensure that there is no denial of service while deploying 802.1X. When the deployment moves into an enforcement phase, it is then recommended to use MDA mode. This configuration example assumes that you are in the initial phases of your deployment.

To configure the host mode and violation action on the port, follow these steps:

Step 1. Set the host mode of the port by entering the following command:

```
C3560X(config-if-range)# authentication host-mode
multi-auth
```

Step 2. Configure the violation action by entering the following command:

```
C3560X(config-if-range)# authentication violation restrict
```

When an authentication violation occurs, such as more MAC addresses than are allowed on the port, the default action is to put the port into an err-disabled state. Although this behavior may seem to be a secure behavior, it can create an accidental denial of service, especially during the initial phases of deployment. Therefore, you set the violation action to **restrict**. This mode of operation allows the first authenticated device to continue with its authorization and deny any additional devices.

Configure Authentication Settings

802.1X is designed to be binary by default. Successful authentication means the user is authorized to access the network. Unsuccessful authentication means the user has no access to the network. This paradigm does not lend itself very well to a modern organization. Most organizations need to do workstation imaging with Pre-Execution Environment (PXE) or may have some thin clients that need to boot with DHCP and don't have an 802.1X supplicant.

In addition, when early adopters of 802.1X would deploy authentication throughout the organization, there were repercussions. Many supplicants were misconfigured. In addition, many unknown devices could not authenticate because those endpoints lacked supplicants.

Cisco created open authentication to aid with deployments. Open authentication allows all traffic to flow through the switch port even without the port being authorized. This feature allows authentication to be configured across the entire organization without denying access to any device. (You might wonder how you can ensure the security of your network when you are not denying access to connecting endpoints. Fear not: We cover this shortly.)

Figure 11-2 shows the difference between a switch port configured with the default behavior of 802.1X and a switch port with open authentication configured. Open authentication is a key feature that enables the phased approach to deploying authentication that is discussed in great length in Chapter 19.

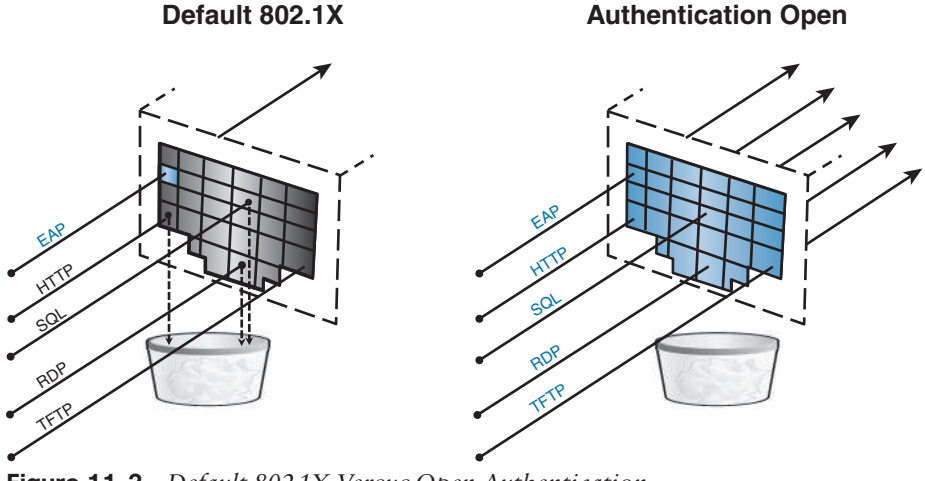


Figure 11-2 *Default 802.1X Versus Open Authentication*

To configure the authentication mode and enable MAB, follow these steps:

Step 1. Set the port for open authentication by entering the following command:

```
C3560X(config-if-range)# authentication open
```

Step 2. Enable MAB on the port by entering the following command:

```
C3560X(config-if-range)# mab
```

Step 3. Enable the port to do IEEE 802.1X authentication by entering the following command:

```
C3560X(config-if-range)# dot1x pae authenticator
```

Earlier you enabled dot1x globally on the switch, but no authentication would actually occur yet. The reason for this is that the individual switch ports have not been configured to do the 802.1X communication. To enable the 802.1X communication to and from a switch port, you use the command **dot1x pae authenticator**.

Even though you have configured a lot of commands on the switch already, you are not ready to process any authentication requests yet. There is a key command that you will learn how to configure shortly.

Configure Authentication Timers

Key Topic

Many timers can be modified as needed in a deployment. Unless you are experiencing a specific problem in which the adjustment of a timer might correct some unwanted behavior, it is recommended to leave all timers at their default values except for the 802.1X transmit timer (tx-period).

The tx-period timer defaults to a value of 30 seconds. Leaving this value at 30 seconds provides a default wait of 90 seconds (three times the tx-period) before a switch port begins the next method of authentication and the MAB process begins for non-authenticating devices.

Based on numerous deployment experiences, the recommendation is to set the tx-period value to 10 seconds to provide the most optimal time for MAB devices. Setting the value below 10 seconds may result in unwanted behavior, and setting the value greater than 10 seconds may result in DHCP timeouts. To configure the tx-period timer, enter the following command:

```
C3560X(config-if-range)# dot1x timeout tx-period 10
```

The recommendation of 10 seconds may need to be reduced or increased depending on the needs of your specific deployment. However, 10 seconds has been a very successful timer for a vast majority of deployments.

Apply the Initial ACL to the Port and Enable Authentication

This section shows you how to prepare the port for Monitor mode by applying a default ACL on the port. This is how you can limit traffic on a port with open authentication. The pre-authentication ACL might vary depending on the organization, but the goal should be to allow the traffic needed prior to the authentication and authorization result. This might include DHCP traffic for the device to receive an IP address, TFTP for PXE boot, and so

on. After an authorization is handed down from ISE, the authorization result bypasses the applied ACL on the port.

To apply the initial ACL to the port and enable authentication, follow these steps:

Step 1. Apply the initial ACL (ACL-ALLOW) by entering the following command:

```
C3560X(config-if-range)#ip access-group ACL-ALLOW in
```

Key Topic

If you wish to enable authentication, you can do so now. There is a key interface command needed to truly enable all these authentication commands: **authentication port-control auto**. Without this command, everything may appear to work, but no authentications will be sent to the RADIUS server.

Step 2. Turn on authentication by entering the following command:

```
C3560X(config-if-range)#authentication port-control auto
```

There are three options for the **port-control** command:

- **auto**: Enables authentication on this port and allows the authorization result to be sent from the RADIUS server. Basically, it turns on authentication. This is the most widely used option for this command.
- **force-authorized**: Hard-codes the port to be in an authorized state. This basically disables the process of authentication but allows all endpoints to connect to the network. Basically, it allows all endpoints connected to this switch port.
- **force-unauthorized**: Hard-codes the port to be in an unauthorized state. This disables the process of authentication and denies network access to all endpoints. Basically, it denies service to the endpoints connected to this switch port.

Authentication Configuration on WLCs

This section reviews the configuration for Cisco Wireless LAN Controller. The focus is on Version 8.0 and above. All the Cisco Wireless LAN Controller screenshots in this chapter are from Version 8.5. However, there should not be any large discernable difference in appearance between various 8.x releases.

NOTE This chapter assumes that you have established basic connectivity with the NAD and are now to the point of bootstrapping the WLC for use with ISE. Some configuration in this section is globally applicable, meaning it applies to the entire system, and other configuration is per wireless LAN (per SSID).

Configure the AAA Servers

Much as with wired NADs, the first step in configuring a wireless LAN controller (WLC) is to add the ISE Policy Services nodes (PSNs) to the WLC as the RADIUS authentication and accounting servers. This configuration affects the entire WLC as a system, much like global configuration with a switch. The AAA server configuration for a WLC takes place in two sections: authentication servers and accounting servers.

Add the RADIUS Authentication Servers

To add the RADIUS authentication servers, start by adding the ISE PSN to the authentication servers. Follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Security > RADIUS > Authentication**.

Step 2. Ensure that MAC Delimiter is set to **Hyphen** (which should already be the default). Figure 11-3 shows the final configuration for the Call Station ID and MAC Delimiter settings.

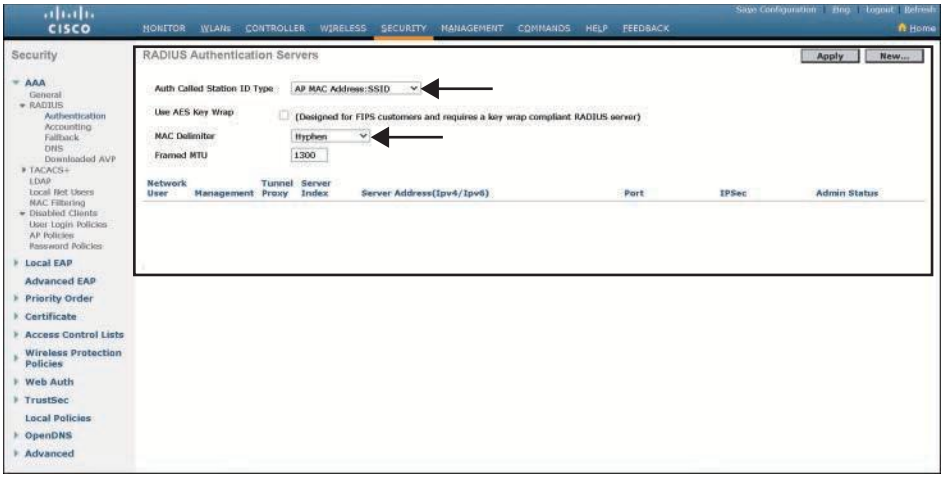


Figure 11-3 *Security > RADIUS > Authentication*

Step 3. Click **New** to add the ISE PSN.

Step 4. Fill in the IP address of the PSN (or the VIP address if using a load balancer).

Step 5. Fill in the shared secret, which *must* match what is configured the ISE Network Device object.

Step 6. Ensure that the port is 1812 for authentication.

Step 7. Ensure that Server Status is set to **Enabled**.

Step 8. Ensure that Support for CoA is set to **Enable**. (For some Wireless LAN Controller versions, there is a checkbox in the new RADIUS authentication server menu to apply the Cisco ISE default settings. If this box is checked, it enables support for CoA as well.)

Step 9. Accept the default Server Timeout setting of 5 seconds.

Step 10. Ensure that Network User is set to **Enable**. This simply means that the RADIUS server may be used for network authentications.

Step 11. Click **Apply** in the upper-right corner.

Step 12. Save the configuration.

Figure 11-4 shows a completed server configuration.

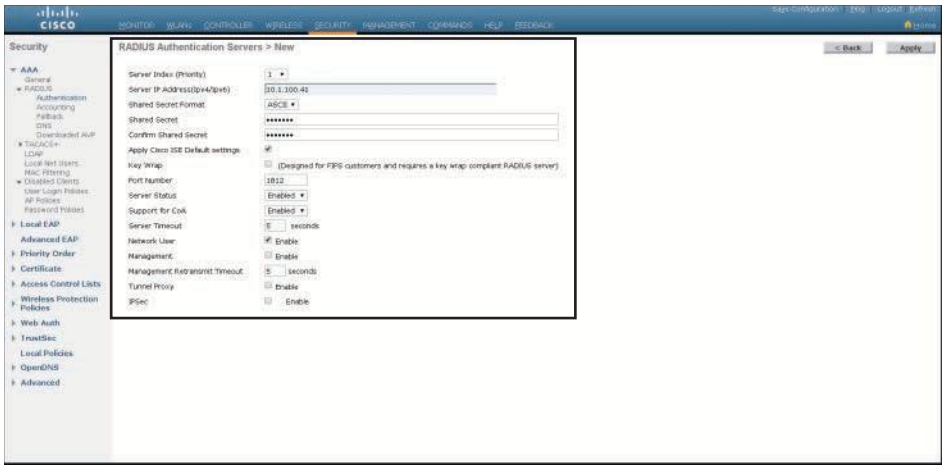


Figure 11-4 *Completed Server Configuration*

Repeat these steps for each PSN you need to add.

Add the RADIUS Accounting Servers

To define the ISE PSNs for the accounting server section, follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Security > RADIUS > Accounting.**

Step 2. Ensure that Acct Called Station ID Type is set to **AP MAC Address:SSID** to match the RADIUS Authentication server settings.

Step 3. Ensure that MAC Delimiter is set to **Hyphen**, as shown in Figure 11-5.

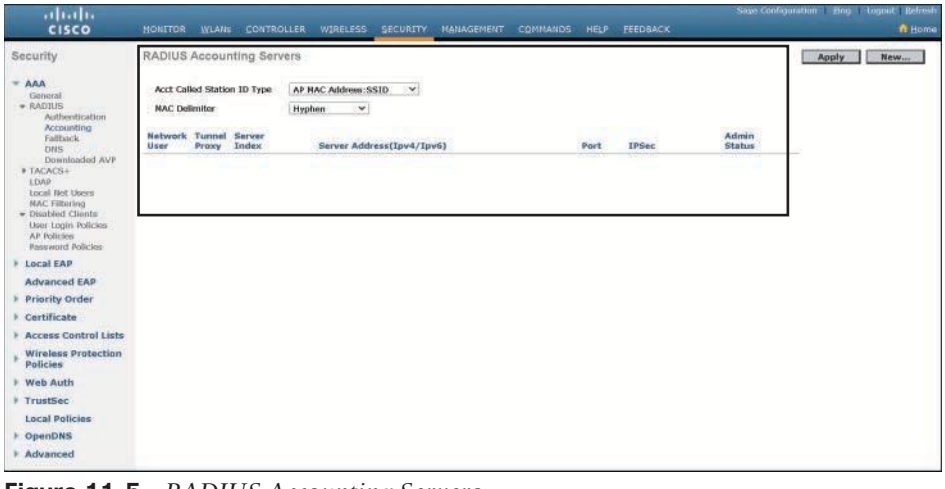


Figure 11-5 *RADIUS Accounting Servers*

Step 4. Click **New** to add the ISE PSN.

Step 5. Enter the IP address of the PSN.

Step 6. Enter the shared secret to match what is configured on ISE.

Step 7. Ensure that Port Number is set to **1813**.

Step 8. Ensure that Server Status is set to **Enabled**.

Step 9. Verify that the **Network User** checkbox is selected. Figure 11-6 shows a completed accounting server configuration

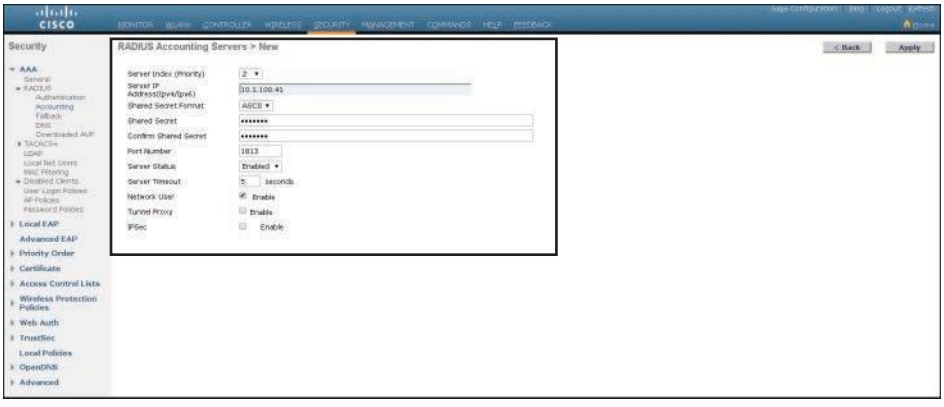


Figure 11-6 *Completed Accounting Server Configuration*

Step 10. Click **Apply** in the upper-right corner.

Step 11. Save the configuration.

Repeat these steps for each PSN that you need to add.

Configure RADIUS Fallback (High Availability)

The primary RADIUS server (the server with the lowest server index) is assumed to be the most preferable server for a Cisco WLC. If the primary server becomes unresponsive, the controller switches to the next active server (the server with the next lowest server index). The controller continues to use this backup server unless you configure the controller to fall back to the primary RADIUS server when it recovers and becomes responsive or to a more preferable server from the available backup servers.

To configure RADIUS fallback, follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Security > AAA > RADIUS > Fallback**.

Step 2. Set Fallback Mode to **Active**. By doing this, you set the WLC to revert to a server with a lower priority from the available backup servers by using RADIUS probe messages to proactively determine whether a server that has been marked inactive is back online.

Step 3. For Username, enter **radius-test** as the name to be sent in the inactive server probes. It's important to note that you do not need to enter a password for this test user account because the system simply looks for a response from the RADIUS server; pass or fail does not matter.

Step 4. Enter **300** as the value for the Interval in Sec field. The interval is the inactive time in passive mode and the probe interval in active mode. The valid range is 180 to 3600 seconds, and the default value is 300 seconds. Figure 11-7 shows the completed RADIUS fallback settings.

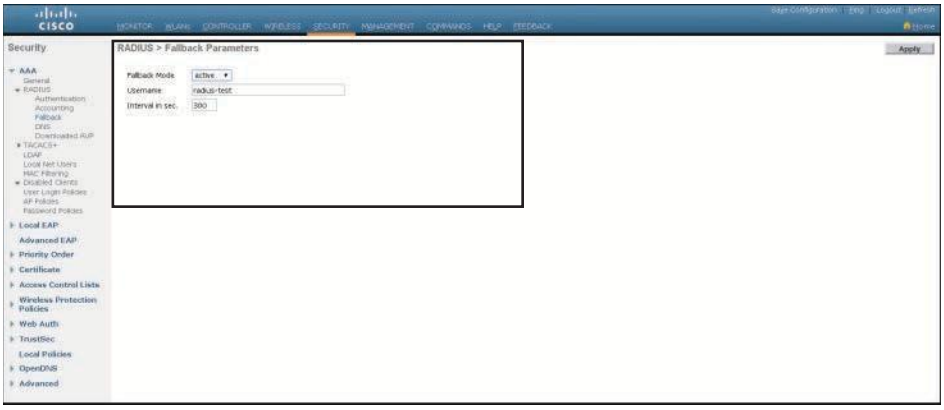


Figure 11-7 *RADIUS Fallback Settings*

Configure the Airespace ACLs

Just as you did with the Cisco Catalyst switches, you can pre-stage a WLC with some access lists. In this case, you will use the following ACLs:

- Web Authentication Redirection (ACL-WEBAUTH-REDIRECT)
- Posture Agent Redirection (ACL-AGENT-REDIRECT)

Notice that these ACL names are the same as the names of the Cisco switches. This is not required, but this naming convention ensures consistency where possible.

Create the Web Authentication Redirection ACL

As with a Catalyst switch, you need a local ACL on a WLC that will redirect web traffic to the Centralized Web Authentication (CWA) portal. With a Catalyst switch, a **permit** statement means that the traffic should be redirected, and a **deny** statement describes traffic that should not be redirected. With the switch, you need two ACLs: one to define what gets redirected and a second one to filter traffic (that is, permit or deny traffic flow).

With Cisco Wireless LAN Controller, there is a single ACL that does double duty: It permits and denies traffic flow, and at the same time, it redirects the traffic that is denied to the CWA portal.

To create the Web Authentication Redirection ACL, follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Security > Access Control Lists > Access Control Lists.**

Step 2. Click **New** to add a new ACL, as shown in Figure 11-8.

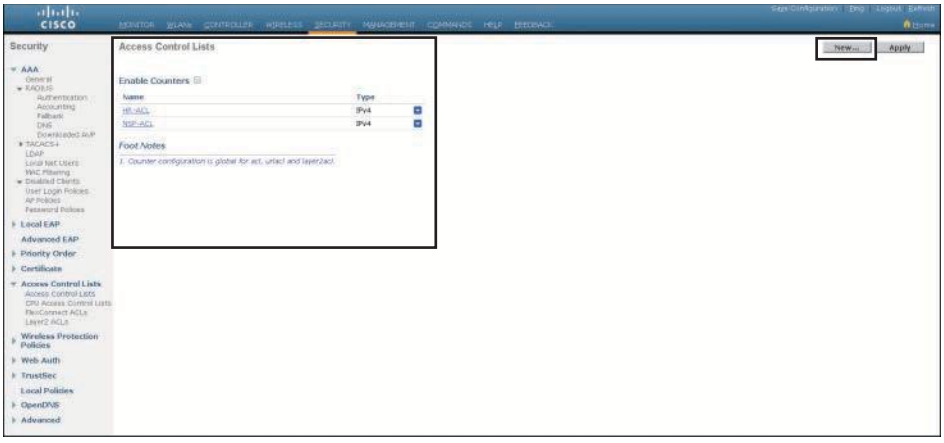


Figure 11-8 *Access Control Lists*

Step 3. Fill in the name **ACL-WEBAUTH-REDIRECT.**

Step 4. Click **Apply**, as shown in Figure 11-9. You are returned to the main ACL screen.

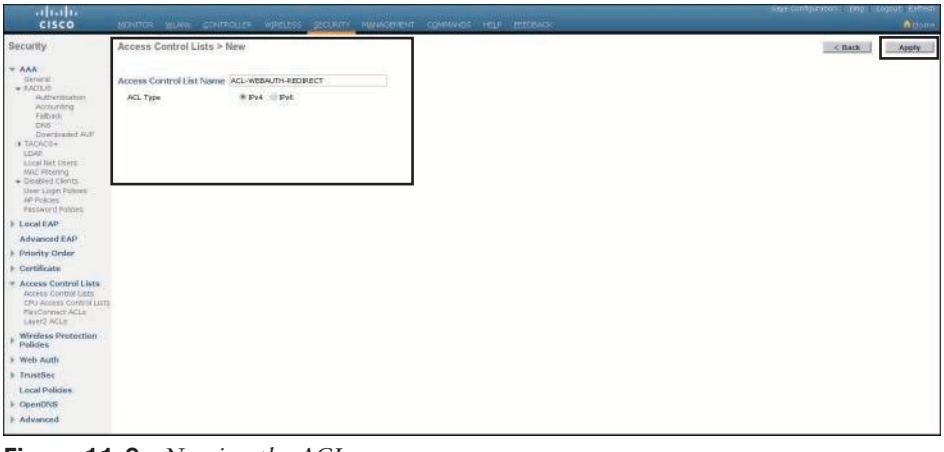


Figure 11-9 *Naming the ACL*

Step 5. Click on the new entry **ACL-WEBAUTH-REDIRECT** (see Figure 11-10).

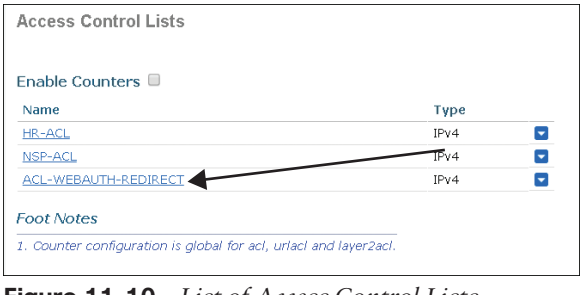


Figure 11-10 *List of Access Control Lists*

Step 6. Click **Add New Rule** in the upper-right corner. (A rule in a WLC is the equivalent of an access control entry in a switch. It is a line in an ACL.)

Step 7. Create a set of rules for this ACL that does the following:

- Permits all traffic outbound (toward the client)
- Permits DHCP and DNS inbound
- Permits TCP port 8443 to the ISE servers
- Denies all other traffic (and redirects it)

Figure 11-11 shows an example of a sample completed ACL.

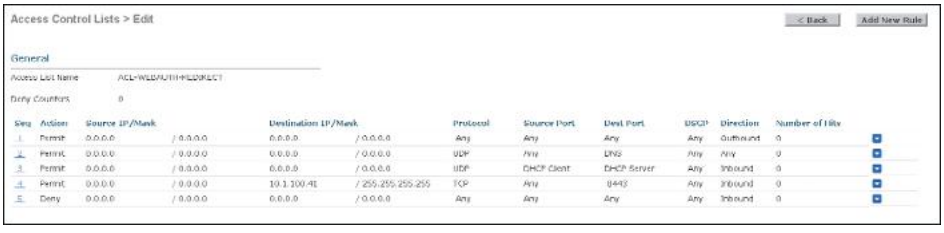


Figure 11-11 *ACL-WEBAUTH-REDIRECT Example*

Step 8. After finishing the configuration of the access control list, click the **Apply** button.

Add Google URLs for ACL Bypass

When Android endpoints go through the BYOD onboarding process, they have to access the Google Play Store to download the Network Setup Assistant app. Allowing this access is not as simple as allowing an IP address in the URL. There might be hundreds of IP addresses that resolve for the DNS names used by the Google Play Store. Instead of using IP addresses, a WLC allows you to use DNS-based ACLs in the form of URL lists that can be added to an existing ACL.

To add Google URLs for ACL bypass, follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Security > Access Control Lists > Access Control Lists**.

Step 2. As shown in Figure 11-12, hover your mouse pointer over the blue-and-white dropdown arrow to the right of the ACL-WEBAUTH-DIRECT ACL created in the previous section.

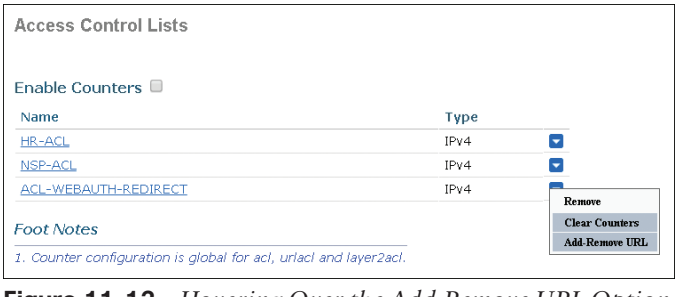


Figure 11-12 *Hovering Over the Add-Remove URL Option*

Step 3. Click **Add-Remove URL**. You are now brought to the URL list. The URLs that you enter here are configured with an implicit wildcard in the first portion. Matches to these URL entries will be permitted through the ACL.

Step 4. Enter the URLs that are to be permitted through the ACL, as shown in Figure 11-13. Entering google.com and clients.android.google.com usually is enough in the United States. If this does not succeed, you can use wildcards in the form *.* for domain extensions.



Figure 11-13 *URL List*

Create the Posture Agent Redirection ACL

Just as with the Web Authentication traffic, you want to have an ACL for posture assessments. When a client successfully authenticates to the network but that client's posture is not known, the resulting authorization should be to put the client in a Posture_Req state on the controller and use a URL redirection with an ACL to redirect traffic to ISE.

For posture assessment, you must allow some additional ports in the ACL that are not required with Web Authentication. Technically, there is nothing to prevent you from consolidating Web Authentication and posture agent redirection to a single ACL if you are willing to allow the required ports for both use cases. In order to maintain consistency throughout our deployment examples, however, this section shows how to create a separate ACL.

To create the Posture Agent Redirection ACL, follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Security > Access Control Lists > Access Control Lists**.

Step 2. Click **New** to add a new ACL.

Step 3. Fill in the name **ACL-AGENT-REDIRECT**.

Step 4. Click **Apply**.

Step 5. Click on the new entry **ACL-AGENT-REDIRECT**.

Step 6. Click **Add New Rule** in the upper-right corner.

Step 7. Build a rule set for this ACL that accomplishes the following:

- Permits all traffic outbound (toward the client)
- Permits DHCP and DNS inbound
- Permits TCP port 8443 to the ISE server
- Permits TCP and UDP ports 8905 and 8909 to the ISE server

- Permits applicable traffic to the remediation servers, which fix the client if it fails the posture assessment. (The remediation servers could be patch-management systems or even external websites. In the ACL shown in Figure 11-14, the remediation server IP address is 10.1.100.100.)
- Denies all other traffic (and redirects it)

Figure 11-14 show an example of a sample completed ACL.

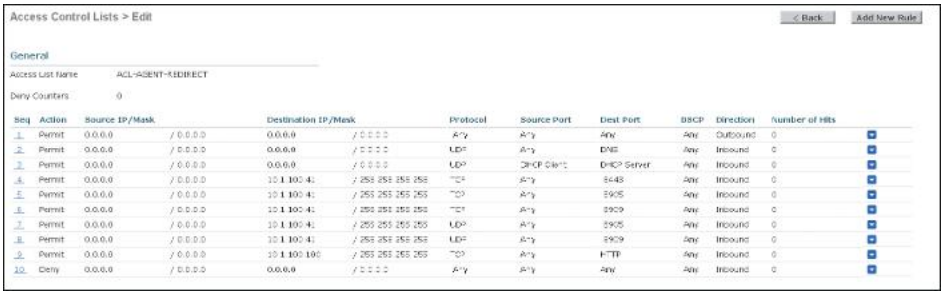


Figure 11-14 *ACL-AGENT-REDIRECT Example*

Create the Dynamic Interfaces for the Client VLANs

When you want to assign a user or device to a VLAN on a Catalyst switch, you just assign the VLAN to the port, and the entire switch port is assigned to that particular VLAN.

The wireless controller has only a few physical connections to the wired network, and it must bridge all wireless users from their RF network (Wi-Fi) to the physical wired network. It must also have the ability to assign a different VLAN for each authenticated session (if necessary). You may be thinking that it just needs to be connected with a trunk, right? Well, yes, that's true.

The WLC can be configured to use 802.1Q to tag traffic for a specific VLAN as that traffic exits the controller. However, the controller calls this a dynamic interface because of the way it can assign a physical interface to traffic or assign a tag.

The following sections show how to create two dynamic interfaces: one for employee traffic and one for guest traffic.

Create the Employee Dynamic Interface

This employee dynamic interface will be used for all successful authentications to the corporate wireless LAN, providing full access to the entire network. To create it, follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Controller > Interfaces**.

Step 2. Click **New**.

Step 3. Name the **employee**.

Step 4. Provide the VLAN ID to be used in the 802.1Q tag.

Step 5. Click **Apply**. You are taken to the Edit menu for that interface.

Step 6. Provide an IP address, netmask, and gateway for the VLAN.

Step 7. Provide the DHCP server address.

Step 8. Click **Apply**.

Figure 11-15 shows this configuration.

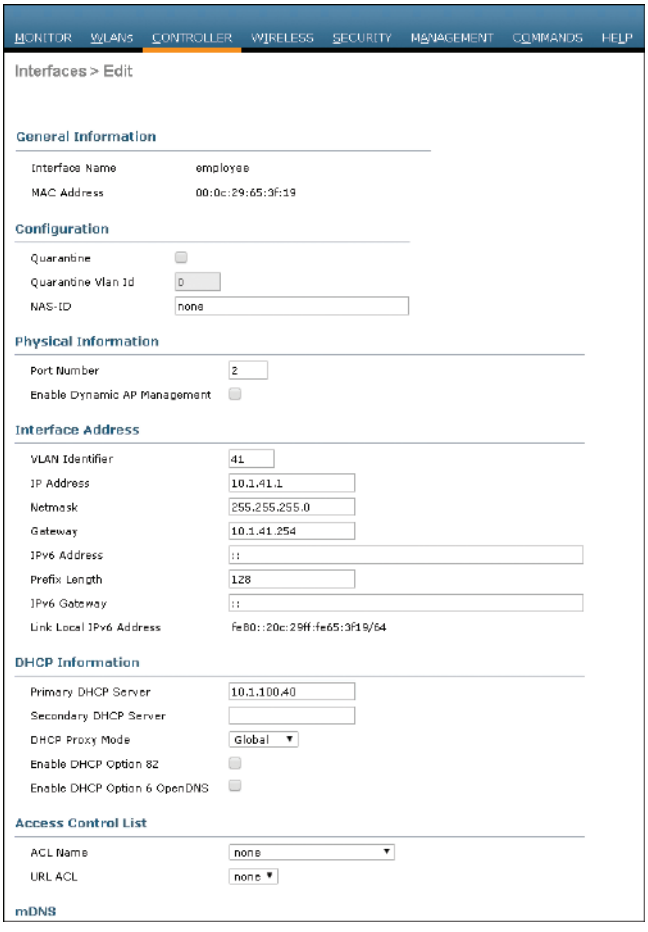


Figure 11-15 *Employee Dynamic Interface*

Create the Guest Dynamic Interface

The guest dynamic interface will be used for all devices connecting to the guest WLAN, as well as unsuccessful or unauthorized authentications to the corporate wireless LAN. This interface will provide Internet access only. To create it, follow these steps in the Wireless LAN Controller GUI:

Step 1. Navigate to **Controller > Interfaces**.

Step 2. Click **New**.

Step 3. Name the interface **guest**.

Step 4. Provide the VLAN ID to be used in the 802.1Q tag.

Step 5. Click **Apply**. Once again, you are taken to the Edit menu for this interface.

Step 6. Provide an IP address, netmask, and gateway for the VLAN.

Step 7. Provide the DHCP server address.

Step 8. Click **Apply**.

Figure 11-16 shows this configuration.

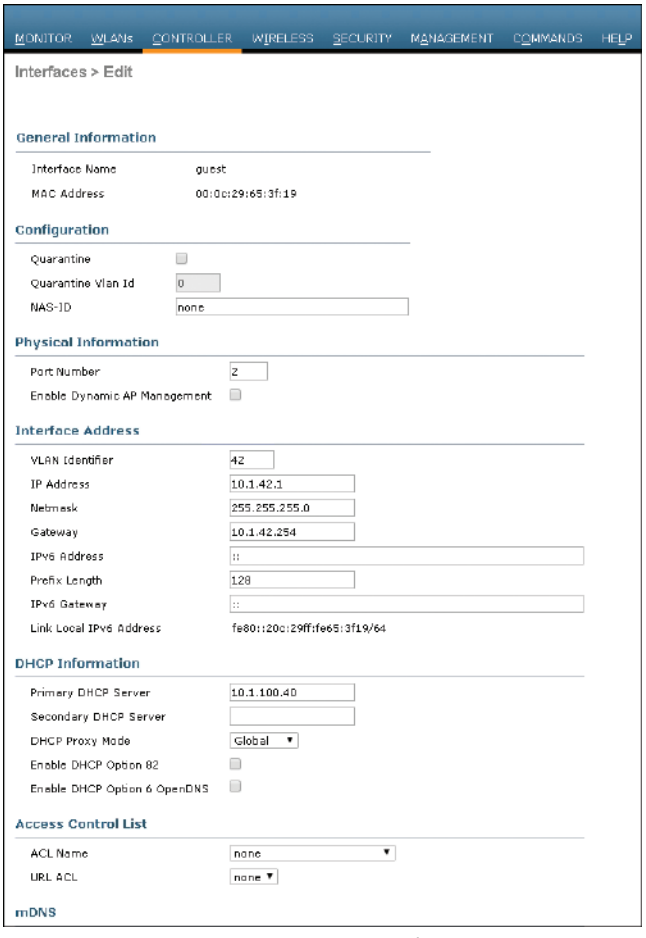


Figure 11-16 *Guest Dynamic Interface*

Create the Wireless LANs

Now that the RADIUS servers, ACLs, and dynamic interfaces are all created and configured, you can create two WLANs: a guest WLAN and a corporate WLAN. The guest WLAN will be an open WLAN, and the corporate WLAN will be configured to use 802.1X to authenticate devices.

Create the Guest WLAN

You can create the guest WLAN as an open SSID WLAN but send the endpoint MAC addresses to ISE for MAB—just as you did for the wired networks.

To create the guest WLAN, follow these steps in the Wireless LAN Controller GUI:

Step 1. Using the headers across the top, navigate to **WLANs.**

Step 2. Select **Create New** from the dropdown menu.

Step 3. Click **Go**.

Step 4. Select **WLAN**.

Step 5. Give the WLAN profile the name **CP-Guest** (see Figure 11-17).

Step 6. Provide the SSID name **CiscoPress-Guest**.

NOTE The WLAN ID is a unique identifier for each WLAN/SSID on a controller. The default setting should be fine in most cases. However, the WLAN ID may be used as a condition within an authorization policy. If you choose to use the WLAN ID in your policy, ensure that the WLAN ID is the same across all WLCs, or you may get unwanted results.

Step 7. Click **Apply**.



Figure 11-17 *Guest WLAN Creation*

Initially, the SSID is created with the default of WPA-2 authentication but little else. In the following sections, you will assign the VLAN or VLAN group to the SSID, define the RADIUS servers, and allow the AAA server to make changes when a successful authentication occurs.

General Tab

You use the General tab to enable or disable a WLAN and select the default interface. You can assign the correct VLAN or VLAN group for the SSID, the name of the SSID, and the SSID number.

Select the General tab and follow these steps:

Step 1. If you are ready to put the current SSID in an operative state, ensure that the **Enabled** checkbox is selected.

Step 2. From the Interface/Interface Group (G) dropdown, choose the previously created **guest** interface.

Figure 11-18 shows the completed General tab.



Figure 11-18 *Guest WLAN General Tab*

Layer 2 Security Tab

You use the Layer 2 Security tab to specify the manner of protection by which devices associate to the wireless network and what encryption should be used (if any). You also use it to define the use of Wi-Fi Protected Access (WPA), Wired Equivalent Privacy (WEP), and MAC filtering. For an open guest network, you can select no Layer 2 security and instead enable MAC filtering. This is how you enable wireless MAB.

Select the Layer 2 Security tab and follow these steps:

Step 1. Change the Layer 2 Security dropdown from the default (WPA) to **None**.

Step 2. Check the MAC Filtering box (to use MAB).

Figure 11-19 shows the completed Layer 2 Security tab.

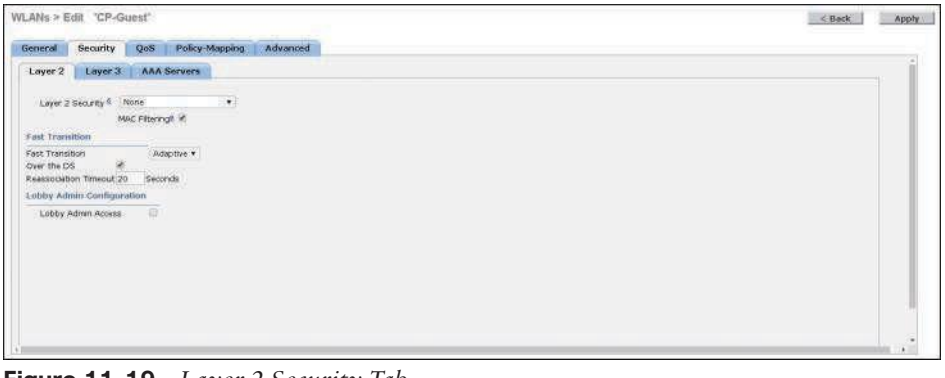


Figure 11-19 *Layer 2 Security Tab*

Layer 3 Security Tab

You use the Layer 3 Security tab options for Local Web Authentication. You should set Layer 3 Security to None for an ISE deployment with Centralized Web Authentication, as shown in Figure 11-20.



Figure 11-20 *Layer 3 Security Tab*

AAA Servers Tab

In order for Cisco Wireless LAN Controller to work correctly for all ISE services, the authentication and accounting servers must be set identically. While the Wireless LAN Controller GUI permits different settings for differing servers, setting the servers differently will break certain functionality related to ISE services.

Select the AAA Servers tab and follow these steps:

Step 1. Select your ISE PSN(s) for both authentication and accounting.

NOTE Starting in WLC Version 8, Cisco aimed to simplify the ISE configuration on a WLAN. Starting in this version, there is an Apply Cisco ISE Default Settings option that can be enabled on the AAA Servers tab of the WLAN settings. If this option is enabled, the following settings are applied:

- The Layer 2 security of the WLAN is set to WPA+WPA2.
- 802.1X is the default AKM.
- MAC filtering is enabled if Layer 2 Security is set to None.

Step 2. Select **Interim Update** box.

Step 3. Set Interim Interval to **0** seconds.

NOTE For WLC Versions 7.6 and earlier, the recommendation was to disable interim updates. However, for Version 8.0 and later, interim updates should be enabled with an interval of 0 seconds. This way, an accounting update is sent only when the client IP address changes and does not impact Device Sensor updates.

Step 4. Click Apply.

Figure 11-21 shows the completed AAA Servers tab.

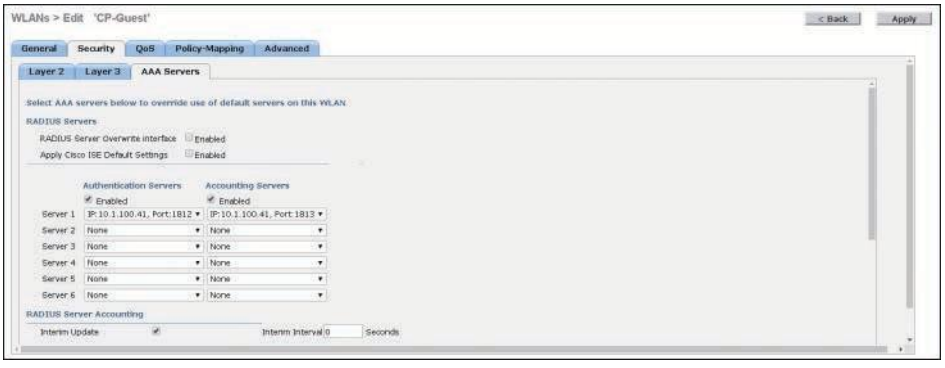


Figure 11-21 *AAA Servers Tab*

Advanced Tab

In order to allow ISE to override the assigned VLAN or interface or even to assign an ACL, the Allow AAA Override setting must be enabled.

Select the Advanced tab and follow these steps:

Step 1. Set Allow AAA Override to **Enable**.

Step 2. Set DHCP Addr. Assignment to **Required**. The WLC has built-in capabilities for the profiling endpoints and sending that profiling data to ISE. The DHCP Addr. Assignment checkbox needs to be enabled for the DHCP Device Sensor built in to the WLC to function correctly.

Step 3. Change the NAC State to **ISE NAC**. This setting is critical to allow for URL redirection, Centralized Web Authentication, posture assessment, native supplicant provisioning, and more. This is a critical setting for the interaction of ISE with the WLC.

NOTE Older versions of WLC code may display **"RADIUS NAC"** instead.

Step 4. Scroll down and select **DHCP Profiling** and **HTTP Profiling**. (We will revisit this in Chapter 14, which covers endpoint profiling in detail.)

Step 5. Click **Apply**.

Step 6. Save the configuration.

Figures 11-22 and 11-23 show the completed Advanced tab.

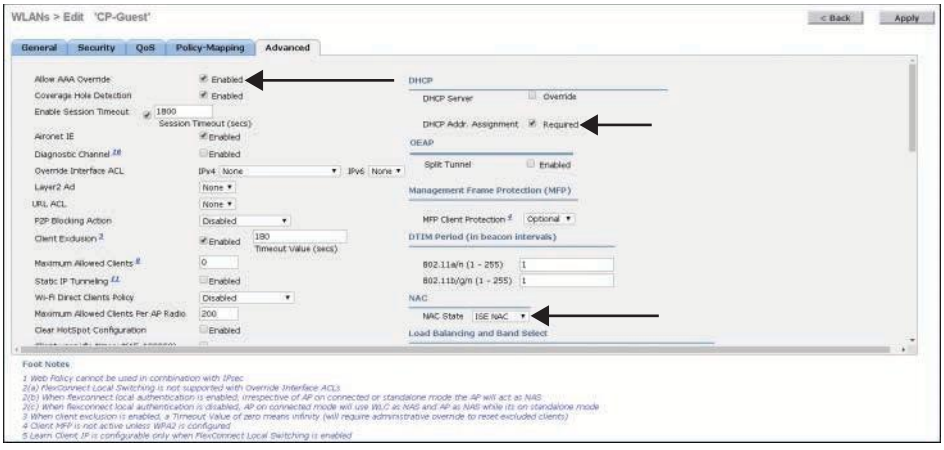


Figure 11-22 *Advanced Tab, Part 1*



Figure 11-23 *Advanced Tab, Part 2*

Create the Corporate WLAN

You can create the corporate WLAN as a closed SSID and require 802.1X authentication in order for an endpoint to associate to the SSID. Unlike wired networks, wireless networks have the added benefit of truly rejecting all access without successful authentication. Users are attuned to the requirement of configuring software in order to connect to a wireless network. This is very much not the case when it comes to wired networks.

To create the corporate WLAN, follow these steps in the Wireless LAN Controller GUI:

Step 1. Using the headers across the top, navigate to **WLANs**.

Step 2. Select **Create New**.

Step 3. Click **Go**.

Step 4. Select **WLAN**.

Step 5. Give the WLAN profile the name **CiscoPress**.

Step 6. Provide the SSID name **CiscoPress**.

Step 7. Click **Apply**.

Figure 11-24 shows the new WLAN being added.

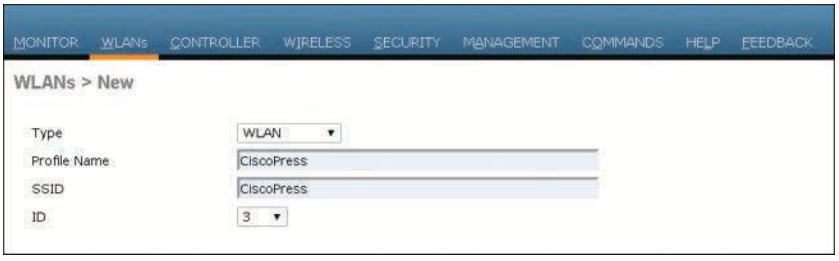


Figure 11-24 *CiscoPress WLAN*

As with the previous SSID configuration, the corporate LAN SSID is created with a default of WPA-2 authentication but little else. The following sections show how to assign the VLAN or VLAN group to the SSID, define the RADIUS servers, and allow the AAA server to make changes upon a successful authentication.

General Tab

You use the General tab to enable or disable a WLAN and select the default interface. You can assign the correct VLAN or VLAN group for the SSID, the name of the SSID, and the SSID number.

Select the General tab and follow these steps:

Step 1. If you are ready to put the current SSID into an operative state, ensure that the **Enable** checkbox is selected.

Step 2. From the Interface/Interface Group (G) dropdown, choose the **employee** interface that you created earlier.

Figure 11-25 shows the completed General tab.

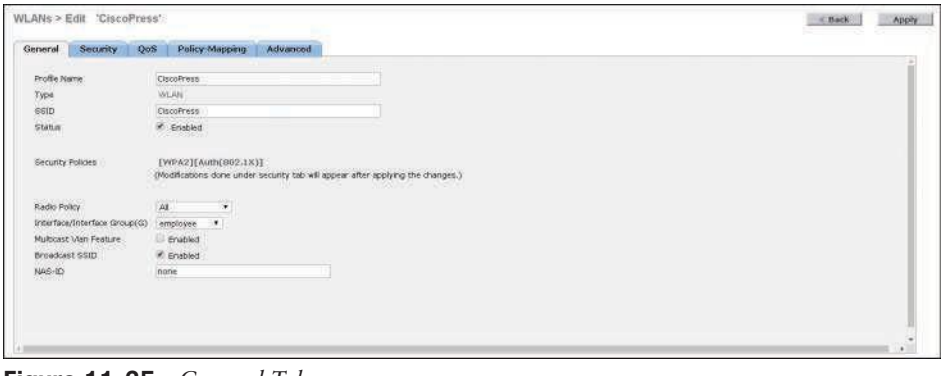


Figure 11-25 *General Tab*

Layer 2 Security Tab

You use the Layer 2 Security tab to specify the manner of protection by which devices associate to the wireless network and what encryption should be used (if any). You also use

it to define the use of Wi-Fi Protected Access (WPA), Wired Equivalent Privacy (WEP), and MAC filtering. For an open guest network, you can select no Layer 2 security and instead enable MAC filtering.

Select the Layer 2 Security tab and follow these steps:

Step 1. Keep the default WPA+WPA2 setting.

Step 2. Ensure that the **MAC Filtering** box is not selected.

Step 3. Ensure that 802.1X is set to **Enable**.

Figure 11-26 shows the completed Layer 2 Security tab.

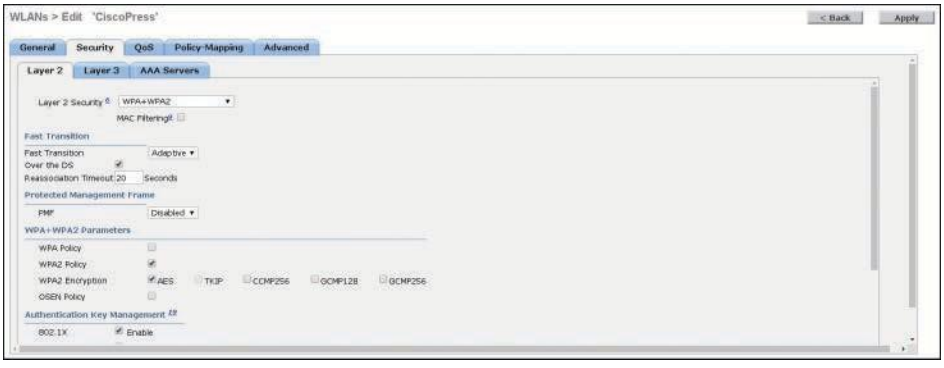


Figure 11-26 *Layer 2 Security Tab*

Layer 3 Security Tab

You use the Layer 3 Security tab options for Local Web Authentication. You should set Layer 3 Security to None for an ISE deployment with Centralized Web Authentication, as shown in Figure 11-27.

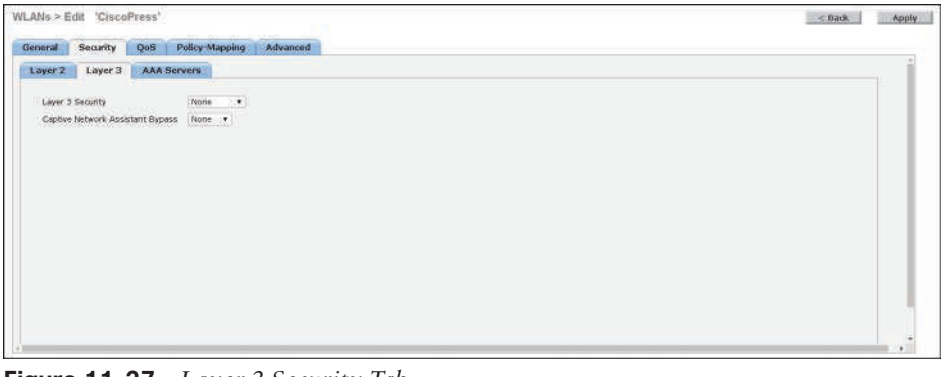


Figure 11-27 *Layer 3 Security Tab*

AAA Servers Tab

In order for Cisco Wireless LAN Controller to work correctly for all ISE services, the authentication and accounting servers must be set identically. While the Wireless LAN

Controller GUI permits different settings for differing servers, setting the servers differently will break certain functionality related to ISE services.

Select the AAA Servers tab and follow these steps:

Step 1. Select your ISE PSN(s) for both authentication and accounting.

Step 2. Select **Interim Update**.

Step 3. Set Interim Interval to **0** seconds.

Step 4. Click **Apply**.

Figure 11-28 shows the completed AAA Servers tab.



Figure 11-28 *AAA Servers Tab*

Advanced Tab

In order to allow ISE to override the assigned VLAN or interface or even to assign an ACL, the Allow AAA Override setting must be enabled.

Select the Advanced tab and follow these steps:

Step 1. Set Allow AAA Override to **Enable**.

Step 2. Set DHCP Addr. Assignment to **Required**. The WLC has built-in capabilities for the profiling endpoints and sending that profiling data to ISE. The DHCP Addr. Assignment checkbox needs to be enabled for the DHCP Device Sensor built in to the WLC to function correctly.

Step 3. Change the NAC State to **ISE NAC**. This setting is critical to allow for URL redirection, Centralized Web Authentication, posture assessment, native supplicant provisioning, and more. This is a critical setting for the interaction of ISE with the WLC.

Step 4. Scroll down and select **DHCP Profiling** and **HTTP Profiling**. (We will revisit this in Chapter 14, which covers endpoint profiling in detail.)

Step 5. Click **Apply**.

Step 6. Save the configuration.

Figures 11-29 and 11-30 show the completed Advanced tab.

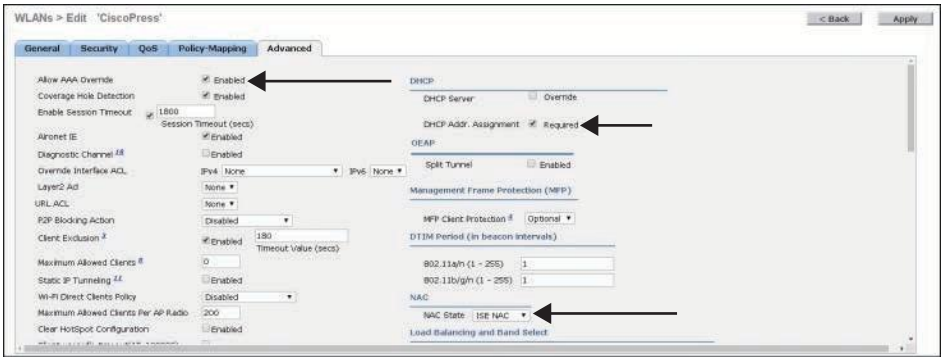


Figure 11-29 *Advanced Tab, Part 1*

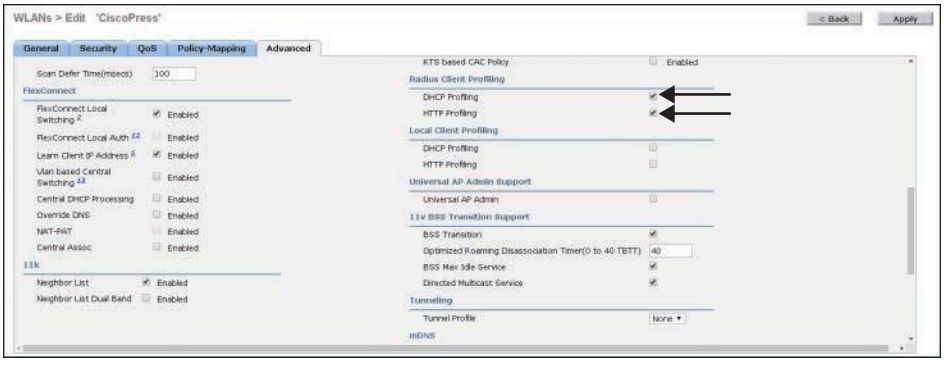


Figure 11-30 *Advanced Tab, Part 2*

Verifying Dot1x and MAB

There are numerous ways to verify the authentication operations of switches and wireless controllers. There are always three locations that must be examined to validate a complete end-to-end transaction. Two of the three locations are much more common and easy to use than the third. These are the three locations:

- **Endpoint supplicant:** For 802.1X authentications
- **NAD:** For all authentications
- **Cisco ISE:** For all authentications

Endpoint Supplicant Verification

Verifying authentications from the supplicant is a bit outside the SISE 300-715 exam, so this book does not focus on it much. With Cisco AnyConnect NAM as your supplicant, you can use DART to get detailed communication and even perform packet captures at the endpoint. If the supplicant is an Apple supplicant (macOS or iOS), you must use the Apple iPhone Configuration Utility to extract and examine the supplicant logs.

With Windows, no supplicant logging is on by default. You must use the command **netsh ras set tracing * enable** to enable the supplicant's logging capabilities. Once logging is enabled, the logs are added to the *%systemroot%\tracing* folder.

Network Access Device Verification

In this section we focus on two NADs: Cisco switches and Cisco wireless LAN controllers. They are quite different in how they verify authentications.

Verifying Authentications with Cisco Switches

There are many items to test with a Cisco switch, and many tools are provided in Cisco IOS. The ones that are used most often are described in this section.

show aaa servers Command

One of the first things to check with a Cisco switch is the status of the RADIUS server (ISE). Using the **show aaa servers** command is a quick and simple way to see the current status of the ISE server from the switch's perspective. Example 11-1 shows the use of this command and its output. The main item of interest with this command's output is the State field. In Example 11-1, the current state is UP. Use this command to validate that the server is up. If it is down, you know that communication to the RADIUS server will not occur.

Example 11-1 *show aaa servers Command*

```
CAT9300# show aaa servers

RADIUS: id 1, priority 1, host 10.1.100.41, auth-port 1812, acct-port 1813
     State: current UP, duration 93974s, previous duration 0s
     Dead: total time 0s, count 0
     Quarantined: No
     Authen: request 29, timeouts 0, failover 0, retransmission 0
             Response: accept 28, reject 0, challenge 0
             Response: unexpected 0, server error 0, incorrect 0, time 247795ms
             Transaction: success 29, failure 0
             Throttled: transaction 0, timeout 0, failure 0
     Author: request 0, timeouts 0, failover 0, retransmission 0
             Response: accept 0, reject 0, challenge 0
             Response: unexpected 0, server error 0, incorrect 0, time 0ms
             Transaction: success 0, failure 0
             Throttled: transaction 0, timeout 0, failure 0
     Account: request 35, timeouts 0, failover 0, retransmission 0
             Request: start 4, interim 4, stop 1
             Response: start 4, interim 4, stop 1
             Response: unexpected 0, server error 0, incorrect 0, time 16ms
             Transaction: success 35, failure 0
             Throttled: transaction 0, timeout 0, failure 0
     Elapsed time since counters last cleared: 1d2h6m
     Estimated Outstanding Access Transactions: 0
     Estimated Outstanding Accounting Transactions: 0
     Estimated Throttled Access Transactions: 0
     Estimated Throttled Accounting Transactions: 0
```

```
     Maximum Throttled Transactions: access 0, accounting 0
     Requests per minute past 24 hours:
             high - 2 hours, 15 minutes ago: 3
             low  - 2 hours, 6 minutes ago: 0
             average: 0
CAT9300#
```

test AAA Command

Cisco switches have a built-in mechanism to send test authentications to the AAA servers they are configured to use. By using the **test aaa** command, you can verify that an authentication is successfully sent to and received by the RADIUS server. Example 11-2 shows the use of the **test aaa** command and the successful response. The **test** command sends an authentication request using PAP_ASCII and returns a RADIUS Access-Accept if successful or an Access-Reject if the password was incorrect. If no response is received, the communication between the switch and the RADIUS server is not occurring. It is also possible that the authentication Allowed Protocols may not permit PAP_ASCII. So in order to conduct this test, you need to ensure that the authentication is not being rejected for that reason.

Example 11-2 *test aaa Command*

```
CAT9300# test aaa group radius employee1 Cisco123 legacy
Attempting authentication test to server-group radius using radius
User was successfully authenticated.

CAT9300#
```

show authentication session interface Command

One of the go-to commands in every implementer's bag of tools is the **show authentication session interface** <*interface*> **detail** command. Yes, the *interface* and **detail** options are added to the base command **show authentication session** to provide more detail. Example 11-3 shows the use of this command and output that indicates a successful MAB authentication. Use this command to validate that the authentications are being attempted, which of them are successful, what authorization results have been assigned, and much more.

Example 11-3 *show authentication session interface details Command*

```
CAT9300# show authentication session int g1/0/2 details
            Interface: GigabitEthernet1/0/2
          MAC Address: 0050.5687.0004
           IP Address: 10.1.10.50
            User-Name: 00-50-56-87-00-04
               Status: Authz Success
               Domain: DATA
      Security Policy: Should Secure
      Security Status: Unsecure
       Oper host mode: multi-auth
     Oper control dir: both
```

```
     Authorized By: Authentication Server
          Vlan Policy: N/A
     Session timeout: N/A
         Idle timeout: N/A
    Common Session ID: 0A0130020000000110073D1F6
      Acct Session ID: 0x00000002
               Handle: 0xA9000012

Runnable methods list:
       Method   State
       mab      Authc Success
       dot1x    Not run
```



Many facets of an authentication session are displayed in this commands output. As this is one of the most important commands, the most important portions of the output are described in this section:

- **Interface**: This is the switch interface where the authentication is taking place.
- **MAC Address**: This is the MAC address of the endpoint being authenticated.
- **IP Address**: The **ip device tracking** command enables the switch to keep track of what IP addresses are associated to the endpoints connected to the switch interface. This applies to both static IP addresses and DHCP assigned addresses. Once the switch has learned the endpoint's IP address, it is listed here.
- **User-Name**: The RADIUS username is displayed here when using 802.1X. When the authentication method is MAB, the username is the same as the MAC address.
- **Status**: This is a list of the status of the authentication session. It may be any of the following: Idle, Running, No Methods, Authc Success, Authc Failed, Authz Success, or Authz Failed.
- **Domain**: This field refers to the domains related to the host mode of the switch interface. With Multi-Auth, MDA, and Multi-Hosts modes, there are two domains: DATA and VOICE. Each authentication session may be assigned to one and only one of the domains.
- **Security Policy**: A better name for this would be MACsec Policy as that is exactly what this field refers to. MACsec is the friendly name for IEEE 802.1AE, a Layer 2 encryption standard. The three options for this field are Should Secure, Must Secure, and Must Not Secure.
- **Security Status**: This field displays the MACsec encryption that is currently applied. Secure means encryption, and Unsecure means no encryption.
- **Oper host mode**: This field lists the host mode of the switch interface. Single-Host, MDA, Multi-Auth, and Multi-Hosts are the available modes of operation.
- **Common Session ID**: The session ID is used to correlate authentication session information between the NAD and the Cisco RADIUS server. When troubleshooting, it is often helpful to compare this value with the one in ISE.

- **Runnable methods**: The available methods are mab, dot1x, and webauth. Possible states of the methods are Not Run, Running, Failed Over, Authc Succeeded, and Authc Failed.

Sending Syslog to ISE

Syslog messages may be generated on Cisco IOS Software in many events. Some of the syslog messages can be sent to the ISE Monitoring and Troubleshooting node (MnT) to be used for troubleshooting purposes. The ISE MnT node correlates the syslog data with the RADIUS data and displays them together in a report. This can be very useful when checking to see if downloadable ACLs (dACLs) have been applied correctly, along with other important validations.

Key Topic

It is not recommended to enable the sending of syslog messages to ISE from all NADs at all times as doing so could add unnecessary load to the MnT node. Instead, you should enable it only when troubleshooting an issue.

To ensure that Cisco ISE is able to compile appropriate syslog messages from the switch, use the following commands:

Step 1. Enable syslog on the switch by using the following commands:

```
CAT9300(config)# logging monitor informational
CAT9300(config)# logging origin-id ip
CAT9300(config)# logging source-interface <interface_id>
CAT9300(config)# logging host <ISE_MNT_PERSONA_IP_
Address_x> transport udp port 20514
```

EPM is a part of the Cisco IOS Software module responsible for features such as Web Authentication and dACLs. Enabling EPM logging generates a syslog related to dACL authorization, and part of the log can be correlated inside ISE when such logs are sent to it.

Step 2. Set up standard logging functions on the switch to support possible troubleshooting/recording for Cisco ISE functions by using the following command:

```
CAT9300(config)# epm logging
```

ISE collects and uses only the following NAD syslog messages:

- AP-6-AUTH_PROXY_AUDIT_START
- AP-6-AUTH_PROXY_AUDIT_STOP
- AP-1-AUTH_PROXY_DOS_ATTACK
- AP-1-AUTH_PROXY_RETRIES_EXCEEDED
- AP-1-AUTH_PROXY_FALLBACK_REQ
- AP-1-AUTH_PROXY_AAA_DOWN
- AUTHMGR-5-MACMOVE
- AUTHMGR-5-MACREPLACE
- MKA-5-SESSION_START

- MKA-5-SESSION_STOP
- MKA-5-SESSION_REAUTH
- MKA-5-SESSION_UNSECURED
- MKA-5-SESSION_SECURED
- MKA-5-KEEPALIVE_TIMEOUT
- DOT1X-5-SUCCESS / FAIL
- MAB-5-SUCCESS / FAIL
- AUTHMGR-5-START / SUCCESS / FAIL
- AUTHMGR-SP-5-VLANASSIGN / VLANASSIGNERR
- EPM-6-POLICY_REQ
- EPM-6-POLICY_APP_SUCCESS / FAILURE
- EPM-6-IPEVENT:
- DOT1X_SWITCH-5-ERR_VLAN_NOT_FOUND
- RADIUS-4-RADIUS_DEAD

Verifying Authentications with Cisco WLCs

Cisco wireless controllers have a number of built-in mechanisms that can be used to verify authentications.

Current Clients

From the Cisco Wireless LAN Controller GUI, navigate to Monitor > Clients, as shown in Figure 11-31.



Figure 11-31 *Monitor Menu*

As shown in Figure 11-32, the Clients screen shows all current clients associated with the WLC, along with very valuable information such as the following:

- **IP Address:** This field shows the IP address of the endpoint, when known.
- **AP Name:** This field shows the name of the access point to which the endpoint is associated.

- **WLAN Profile:** This field shows the name of the WLAN profile created in the WLC.
- **WLAN SSID:** This field shows the name of the SSID for the WLAN profile with which the endpoint has associated.
- **User Name:** With 802.1X, the username is displayed. When the endpoint is authenticated via wireless MAB, the MAC address is displayed.
- **Auth:** If the endpoint/supplicant has authenticated successfully, it is listed here.
- **Device Type:** When the endpoint profile of the device is known, it is displayed in this field.



Figure 11-32 *Monitor > Clients*

To get more detail, you can click on the endpoint's MAC address. This brings up the details related to the individual endpoint's wireless session. As shown in Figure 11-33, a key value to verify for authentication is that Policy Manager is in the RUN state. This means that "all systems are go," and the endpoint's traffic will flow through the wireless controller normally.

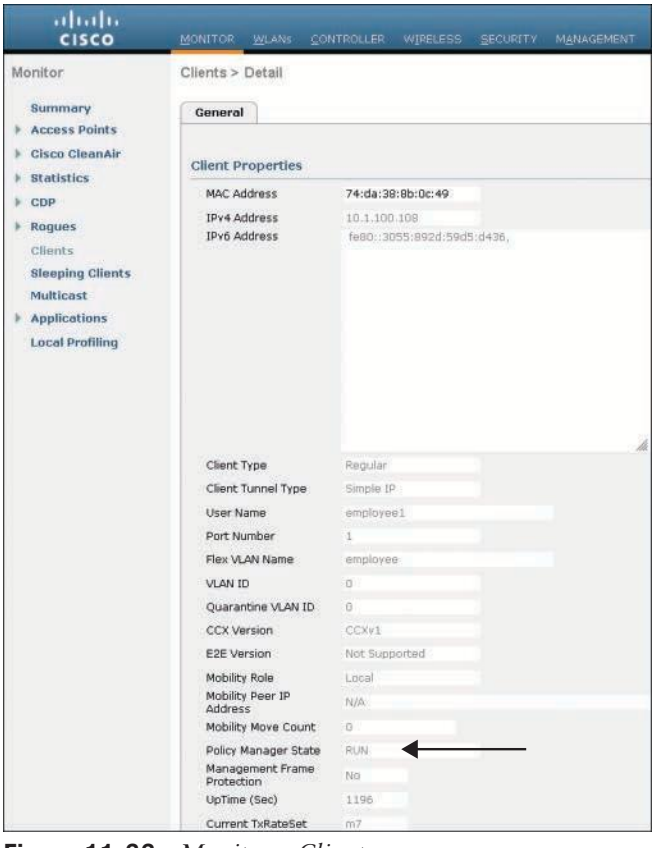


Figure 11-33 *Monitor > Clients*

An endpoint may be in a number of different states. We examine those states in Chapters 12 and 16.

Debug Client

Cisco Wireless LAN Controller provides a few very useful debug commands, including **debug dot1x** and **debug client** <*mac-address*>. **debug dot1x** can be a bit overwhelming on a live network, but the **debug client** command shows only events related to a specific endpoint. Example 11-4 shows an example of the **debug client** command.

Example 11-4 *debug client mac-address Output*

```
(Cisco Controller) >debug client 10:bf:48:d0:05:67


*Dot1x_NW_MsgTask_7: Jul 12 18:40:28.059: 10:bf:48:d0:05:67 Received EAPOL EAPPKT
from mobile 10:bf:48:d0:05:67
*Dot1x_NW_MsgTask_7: Jul 12 18:40:28.059: 10:bf:48:d0:05:67 Received Identity
Response (count=1) from mobile 10:bf:48:d0:05:67
*Dot1x_NW_MsgTask_7: Jul 12 18:40:28.059: 10:bf:48:d0:05:67 Resetting reauth count 1
to 0 for mobile 10:bf:48:d0:05:67
*Dot1x_NW_MsgTask_7: Jul 12 18:40:28.059: 10:bf:48:d0:05:67 EAP State update from
Connecting to Authenticating for mobile 10:bf:48:d0:05:67
*Dot1x_NW_MsgTask_7: Jul 12 18:40:28.059: 10:bf:48:d0:05:67 dot1x - moving mobile
10:bf:48:d0:05:67 into Authenticating state
*Dot1x_NW_MsgTask_7: Jul 12 18:40:28.059: 10:bf:48:d0:05:67 Entering Backend Auth
Response state for mobile 10:bf:48:d0:05:67
*Dot1x_NW_MsgTask_7: Jul 12 18:40:28.067: 10:bf:48:d0:05:67 Processing
Access-Challenge for mobile 10:bf:48:d0:05:67
```

Cisco ISE Verification

Validating an authentication from a NAD has a lot of value, but many times it is preferable to see the authentications from a central console. Cisco ISE is that central console. ISE has RADIUS Live Logs (commonly known as Live Logs) and Live Sessions Log (commonly known as Live Sessions).

RADIUS Live Log

To view RADIUS Live Log, navigate to **Operations > RADIUS > Live Logs**, as shown in Figure 11-34.

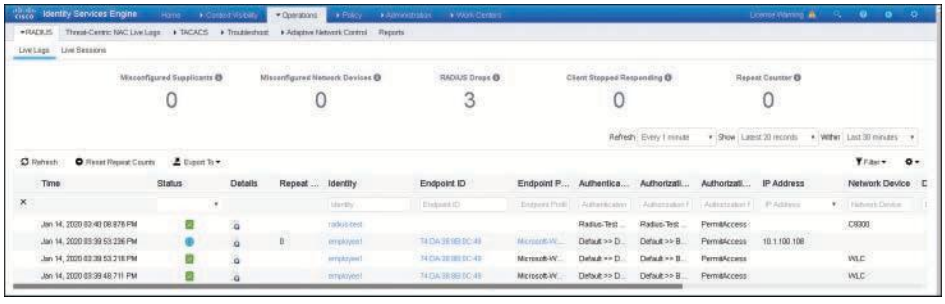


Figure 11-34 *Live Authentications Log*

As shown in Figure 11-34, Live Logs displays a near-real-time display of authentication activity for the ISE cube (that is, ISE deployment). In this dashboard, you can see successful authentications from the radius-test account, which represents the automated-testing performed by the wired switches. In addition, there are successful wireless authentications from employee1.

The Live Logs dashboard is covered in greater detail in other chapters.

Live Sessions

To view Live Sessions, navigate to Operations > RADIUS > Live Sessions. Live Sessions displays a correlated view of all activity related to an authentication session. To see more detail about a specific live session, click on the initiated time of the session, as shown in Figure 11-35. If an endpoint changes state, that change is listed in this view.

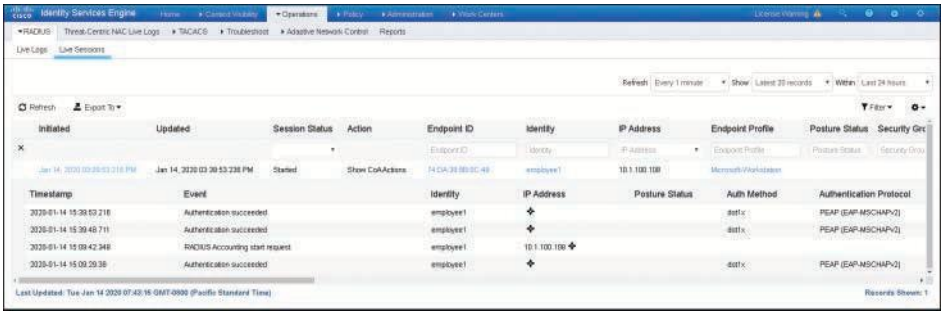


Figure 11-35 *Live Sessions Log*

The Live Sessions log is covered in greater detail throughout this book.

Looking Forward

While this chapter has covered the required configurations for switches and wireless controllers to perform 802.1X and MAB authentications successfully with ISE, these NADs are responsible for much more than just basic authentication and authorization. Local ACLs need to be deployed to define what traffic should be redirected for all use cases, possibly configuring SGT Exchange Protocol (SXP) and other specific items. Those specific NAD configurations are covered in Chapters 12, "Web Authentication," 13, "Guest Services," 16, "Bring Your Own Device," and 17, "TrustSec and MACsec."

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 11-2 lists these key topics and the page number on which each is found.

Key Topic

Table 11-2 Key Topics

Key Topic Element	Description	Page
List	Creating global AAA methods	261
List	Adding ISE PSNs to the switch configuration with IOS 12.2.x	263
List	Adding ISE PSNs to the switch configuration with IOS 15.x and IOS XE	263
List	CoA configuration for switches	265
List	Global 802.1x configuration	266
Section	Understanding FlexAuth	270
List	Host modes	272
Paragraph	Understanding dot1x timers	275
List	Enabling authentication on a switch port	276
List	**show authentication session interface detail** command	298
Paragraph	Sending syslog messages from NADs to MnT nodes for troubleshooting	299

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Flexible Authentication (FlexAuth), network access device (NAD), Change of Authorization (CoA), vendor-specific attribute (VSA), MAC Authentication Bypass (MAB), Single-Host, MDA, Multi-Auth, Multi-Hosts, CDP Second Port Disconnect, EAPoL-Proxy-Logoff, Port Security

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What function do the global AAA configurations serve?
2. What do RADIUS server attributes 6, 8, and 25 provide to ISE?
3. What problem did FlexAuth solve in ISE deployments?
4. What different host modes are available?
5. What is the essential command needed to turn on authentication on a switch port?

This page intentionally left blank



CHAPTER 12

Web Authentication

This chapter covers the following topics:

- Web Authentication scenarios
- Configuring Centralized Web Authentication
- Building CWA authorization rules
- Verifying Centralized Web Authentication

As discussed in Chapter 4, "Non-802.1X Authentications," just because there is no configured supplicant on an endpoint does not mean the user of that endpoint does not need to authenticate. Consider the use cases of guests or visitors, or maybe just a misconfiguration or an expired credential for an end user. The user may still require access to the network.

Enter ***Web Authentication***, commonly referred to as just ***WebAuth***. With WebAuth, an authenticator can send a user to a locally hosted web page—that is, a web page hosted on the local device itself (the switch, wireless controller, or even the firewall or VPN concentrator) where a user can submit a username and password.

As mentioned in Chapter 4, there are multiple types of WebAuth, and Centralized WebAuth (CWA) is the type used with Cisco Secure Access and ISE. CWA is the focus of the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam and, therefore, the main focus of this book.

NOTE This chapter was written based on the assumption that the switches and WLCs have been configured as described in Chapter 11, "Implement Wired and Wireless Authentication." If you have not already configured your network devices for authentication, none of the configuration in this chapter will work, and you should revisit Chapter 11.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 12-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 12-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Web Authentication Scenarios	5
Configuring Centralized Web Authentication	1, 3–6
Building CWA Authorization Policies	3
Verifying Centralized Web Authentication	2, 7–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Before a Cisco switch can generate a self-signed certificate, what configuration is required?
   a. The internal CA must be enabled.
   b. An IPv6 address must be configured.
   c. A Cisco switch cannot generate a self-signed certificate.
   d. A domain name must be configured.
2. Which statement about URL-Redirect ACLs is true?
   a. A URL redirection ACL can be downloaded from ISE to a NAD.
   b. A URL redirection must be preconfigured locally on the NAD, and ISE applies it through the use of RADIUS attribute/value pairs (AV pairs).
   c. There is no ACL needed for URL redirection.
   d. A URL redirection ACL and its ACEs must be configured both in ISE and on the NAD.
3. Which of the following settings is required for a WLAN to support CWA on the Cisco WLC?
   a. SNMP NAC
   b. Layer 3 authentication
   c. ISE NAC
   d. Fast transition
4. For wired and wireless MAB, which option must be configured for unknown identities?
   a. Drop
   b. Continue
   c. Reject
   d. Pass

5. Which of the following rule types need to be created for CWA? (Choose two.)
   a. A WebAuth authentication rule must be created for the authentication through the web portal.
   b. An authorization rule must be created to redirect the user to the CWA portal.
   c. An authentication rule must be created to permit access to users who have successfully authorized through the CWA portal.
   d. An authorization rule must be created to permit access to users who have successfully authenticated through the CWA portal.
   e. A WebAuth authentication rule must be created to redirect the end user to the CWA portal.
6. Which statement is true regarding network segmentation and Web Authentication?
   a. Network segmentation should never be used with Web Authentication; they are mutually exclusive technologies.
   b. VLAN changes may be used, and TrustSec SGTs may be used, but VLAN changes and SGTs can never be used together.
   c. Only TrustSec SGTs can be used with Web Authentication to provide segmentation.
   d. VLAN changes should only be used with devices that can recognize a change and request a new DHCP address.
7. Which of the following statements about CWA is true?
   a. CWA is configured exactly the same for both wired and wireless NADs.
   b. CWA must leverage different policy sets when configured for wired and wireless.
   c. With CWA, the switch isn't aware of the Web Authentication and only identifies the session as using MAB.
   d. CWA stands for Cisco Wide-area Authorization.
8. Which command on a NAD displays information about a URL-redirected session, including the MAC address, IP address, dACL, URL-Redirect ACL, and the URL the end user is being redirected to?
   a. **show epm redirection**
   b. **show authentication sessions**
   c. **show epm authentication | include redirection**
   d. **show authentication session interface** [*interface-name*]
9. Which of the following locations in the ISE GUI is the best one to examine to validate that CWA is working?
   a. Policy > Policy Elements > Results > Authorization
   b. Operations > RADIUS > Live Log
   c. Policy > Policy Elements > Results > Authentication
   d. Operations > Results

10. Which of the following statements most accurately describes the use of Change of Authorization (CoA) in relation to CWA?
    a. The CoA-Reauth causes the NAD to reauthenticate the endpoint within the same session, and ISE is then able to tie together the MAB and CWA authentications.
    b. The CoA sends a Packet of Disconnect (PoD) to the NAD, which starts a new session based on the web credentials.
    c. The CoA-Reauth causes the NAD to reauthenticate the endpoint, which starts a new session based on the web credentials.
    d. The CoA sends a PoD to the NAD, and ISE is able to tie the original MAB session to the new Web Authentication session by correlating the MAC addresses from both authentication sessions.

Foundation Topics

Web Authentication Scenarios

There are a number of reasons that a company may choose to implement a WebAuth strategy. One of the most common reasons is to provide Internet access to visitors (also known as guests), as detailed in Chapter 13, "Guest Services." In addition, as newer versions of ISE come out, many companies are looking to add interactive logins to capture usernames and passwords as additional credentials to certificate-based authentication (think two-factor authentication).

The end user is presented with a web portal to input a username and password. The credentials are then sent from the authenticator to ISE in a standard RADIUS Access-Request packet. So, in a very similar fashion to what occurs with MAC Authentication Bypass (MAB), the switch sends the request for the endpoint, and the endpoint itself does not participate in authentication. Figure 12-1 illustrates the WebAuth concept.

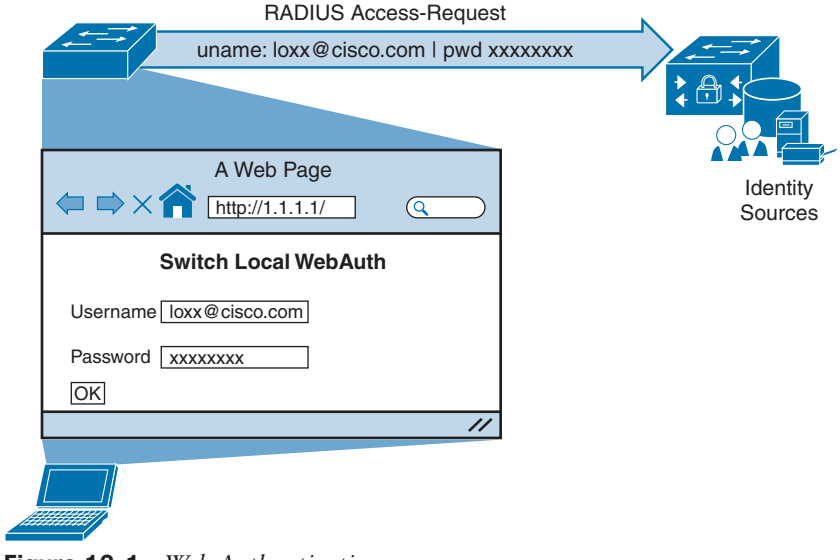


Figure 12-1 *Web Authentication*

The credential that gets submitted through the WebAuth page could be the Active Directory credentials of an employee. The credentials could be guest credentials for someone who is only temporarily allowed to have Internet access (and no other access). The use of WebAuth is really not limited to any specific type of user account.

Keep in mind that WebAuth is only an effective authentication method for a device that has an interactive user. In other words, it would not make sense to try to use WebAuth for a printer as there would be no user to interact with the web portal, enter credentials, and click Submit.

Like MAB, WebAuth is not a standard. There are multiple ways to perform WebAuth, with benefits and downsides to each one.

NOTE It is important not to confuse the term *WebAuth* with the term *WebAuthN*. WebAuthN refers to a new Internet standard for Web Authentication and the use of Web Authentication pages in combination with authentication protocols such as FIDO2 with tokens like YubiKeys, Windows Hello, and Apple's Touch ID. These topics are beyond the scope of the SISE 300-715 exam and, therefore, this book.

Local Web Authentication (LWA)

Local Web Authentication (LWA) is the original WebAuth. With LWA, the authenticator redirects web browser traffic to a locally hosted web portal where a user can enter a username and password.



The credentials are submitted through the switch or wireless controller, which sends the RADIUS Access-Request to the authentication server, using the username and password from the web portal's form. It is key to remember that any time the switch is sending the credentials for the user, it is considered Local Web Authentication.

On a Cisco Catalyst switch, the locally hosted web pages are not very customizable. Many companies require that web portals be customized to match the corporate branding. For those companies, traditional LWA is not usually an acceptable solution—at least not for WebAuth with wired connections.

In addition, when using LWA with Cisco switches, there is no native support for advanced services such as the following:

- Acceptable use policy acceptance pages
- Client provisioning
- Password-changing capabilities
- Self-registration

- Device registration
- BYOD onboarding

For advanced capabilities like these, a company truly needs to consider using Centralized Web Authentication.

NOTE For more details on LWA, see Chapter 4.

Centralized Web Authentication (CWA)

Cisco ISE uses ***Centralized Web Authentication (CWA)*** almost exclusively. While Cisco ISE is capable of supporting LWA methods, those methods are typically reserved for non-Cisco network devices.

Like other forms of Web Auth, CWA is only for interactive users with web browsers, who need to manually enter usernames and passwords.

Change of Authorization (CoA) works fully with CWA, which contributes to support for all the authorization results, such as ACL and VLAN authorization. Keep in mind that any time you change VLANs on an endpoint, the endpoint must be able to detect the VLAN change and trigger an IP address renewal. With 802.1X, the supplicant takes care of the VLAN change detection and address renewal. However, when using WebAuth, a supplicant does not typically exist on the endpoint. Therefore, the DHCP scope length must be set to renew the address quickly, or the portal must use an ActiveX or Java applet to handle the renewal of the IP address after the VLAN assignment, which is not a popular option due to the security concerns related to using Java or ActiveX applets.

CWA also supports advanced services such as the following:

- Client provisioning
- Posture assessments
- Acceptable use policies (AUPs)
- Password changing
- Self-registration
- Device registration
- BYOD onboarding



As described in Chapter 4, a switch or wireless controller only sees MAB, and the rest is handled on the authentication server (ISE). Figure 12-2 shows the MAB occurring with a redirection to the centralized portal, and Figure 12-3 shows how the switch still sees only a MAB request, with ISE maintaining the user authentication.

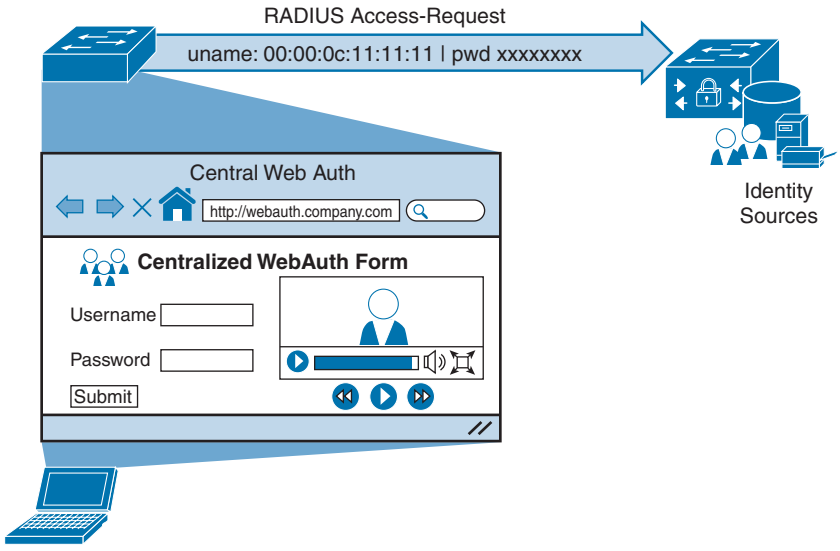


Figure 12-2 *URL-Redirected MAC Authentication Bypass*

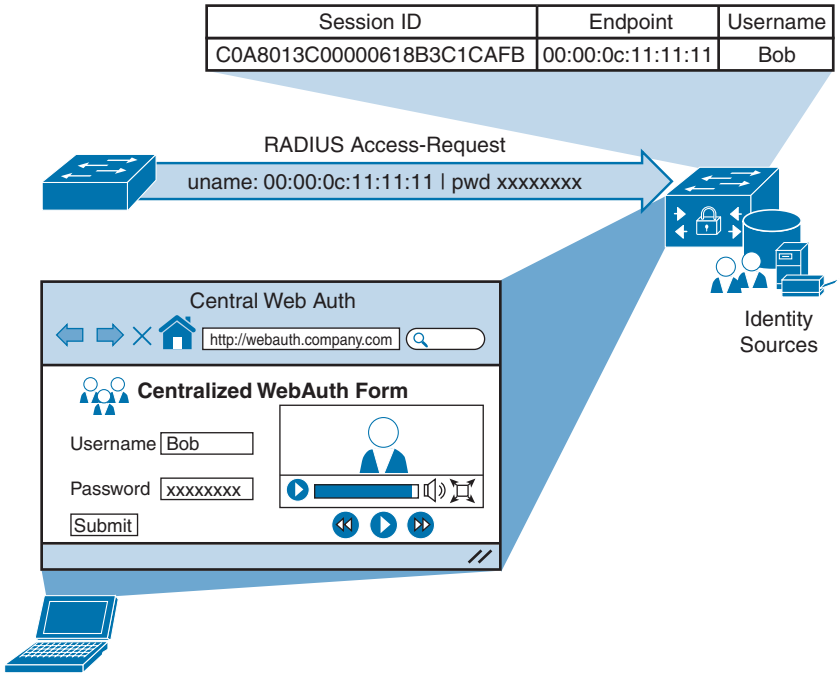


Figure 12-3 *Credentials Never Sent to the Authenticator*

The following steps detail what occurs in Figures 12-2 and 12-3:

Step 1. The endpoint entering the network does not have a supplicant.

Step 2. The authenticator performs MAB, sending the RADIUS Access-Request to Cisco ISE (the authentication server).

Step 3. The authentication server (ISE) sends the RADIUS result, including a URL redirection, to the centralized portal on the ISE server.

Step 4. The end user enters credentials into the centralized portal. Unlike the LWA options, the credentials are never sent to the switch; instead, they are stored within the ISE session directory and tied together with the MAB coming from the switch.

Step 5. ISE sends a reauthentication Change of Authorization (CoA-reauth) to the switch. This causes the switch to send a new MAB request with the same SessionID to ISE, and it is processed.

Step 6. ISE sends the final authorization result to the switch for the end user.

CWA and the URL-redirection capability in the switches and wireless devices are the basis for many of the other solutions in ISE, including Device Registration WebAuth, BYOD onboarding, MDM onboarding, and posture assessment.

Configuring Centralized Web Authentication

Multiple devices need to be configured to enable CWA. The network access device (NAD) requires some special configuration, such as a redirection ACL; in addition, ISE needs authentication and authorization rules set up for CWA. The following sections look at these configurations.

Cisco Switch Configuration

With secure network access using ISE, the switch performs the URL redirection for Web Authentication and also redirects the discovery traffic from the posture agent to the ISE policy service node.

Performing URL redirection at the Layer 2 access (edge) device is a vast improvement over previous NAC solutions, which requires an appliance (such as the inline device) to capture web traffic and perform redirection to a Web Authentication page. URL redirection at the Layer 2 access device simplifies Web Authentication deployment, device onboarding, and the posture agent discovery process.

Configure Certificates on the Switch

In order to redirect HTTPS traffic, there is a prerequisite for the switch to have its own certificate. To configure a certificate, perform the following tasks in global configuration mode on a switch:

NOTE Cisco IOS does not allow for certificates or even self-generated keys to be created and installed until a DNS domain name is defined on the device.

Step 1. To set the DNS domain name on the switch, type **ip domain-name** *domain-name* at the global configuration prompt. Now that the domain name is configured, and the keys can be generated.

Step 2. To generate keys to be used for HTTPS, type **crypto key generate rsa general-keys mod 2048** at the global configuration prompt.

Enable the Switch HTTP/HTTPS Server

The embedded HTTP/HTTPS server in IOS is used to grab HTTP traffic from the user and redirect that user's browser to the Centralized Web Authentication portal or to a device registration portal or even to the mobile device management onboarding portal. This same function is used for redirecting the posture agent's traffic to the Policy Services node. Follow these steps to enable the switch HTTP/HTTPS server:

Step 1. Enable the HTTP server by entering the following command in global configuration mode:

```
C3850(config)# ip http server
```

Step 2. Enable the HTTP secure server by entering the following command:

```
C3850(config)# ip http secure-server
```

Key Topic

Many organizations need to ensure that this redirection process, which is using the switch's internal HTTP server, is decoupled from the management of the switch itself. To disconnect the HTTP management process from the URL-redirection process, run the following two commands in global configuration mode:

```
C3850(config)# ip http active-session-modules none
C3850(config)# ip http secure-active-session-modules none
```

Verify the URL-Redirect ACL

In Chapter 11, you created an access list named ACL-WEBAUTH-REDIRECT, which is used to determine what traffic is redirected to the CWA portal with the **permit** statement. Any traffic that is denied is not redirected.

Key Topic

Contrary to the way a wireless LAN controller works, the URL-Redirect ACL on a switch is used only to determine what traffic is redirected and what traffic is not redirected. If network traffic is denied from redirection, it is not necessarily denied the ability to traverse the network. The traffic-filtering capability comes from the downloadable ACL (dACL) that is sent to the switch from ISE as part of the authorization result.

The use of dual ACLs is limited to IOS-based wired and wireless devices. (The AirespaceOS wireless controllers behave differently and are covered later in this chapter.) Follow these steps to verify the URL-Redirect ACL:

Step 1. Validate whether the ACL-WEBAUTH-REDIRECT ACL is configured on the NAD by entering the following command:

```
C3850# show ip access-list ACL-WEBAUTH-REDIRECT
Extended IP access list ACL-WEBAUTH-REDIRECT
 10 deny udp any any eq domain
 20 permit tcp any any eq www
 30 permit tcp any any eq 443
```

If the ACL is not there or needs to be modified, continue to step 2.

Step 2. Add the following ACL to be used for URL redirection with Web Authentication:

```
C3850(config)# ip access-list ext ACL-WEBAUTH-REDIRECT
C3850(config-ext-nacl)# remark explicitly deny DNS from
being redirected to address a bug
C3850(config-ext-nacl)# deny udp any any eq 53
C3850(config-ext-nacl)# remark redirect all applicable
traffic to the ISE Server
C3850(config-ext-nacl)# permit tcp any any eq 80
C3850(config-ext-nacl)# permit tcp any any eq 443
C3850(config-ext-nacl)# remark all other traffic will be
implicitly denied from the redirection
```

Cisco WLC Configuration

Cisco switches are responsible for redirecting web browser traffic to the centralized portal(s), and Cisco WLCs must do the same thing.

NOTE As stated in the introduction to this chapter, you are expected to have already configured the WLC according to the directions in Chapter 11. In Chapter 11, you should have created an "open" WLAN with MAC filtering enabled and the NAC state configured for ISE NAC. In addition, you created an access list named ACL-WEBAUTH-REDIRECT.

Validate That MAC Filtering Is Enabled on the WLAN

The MAC Filtering option for an open wireless network configures a WLAN for wireless MAB. This is necessary to ensure that an authentication is sent from the WLC to ISE, so ISE can return the URL redirection in the authorization result.

From the WLC's GUI, navigate to the WLANs tab, examine the list of WLANs, and ensure that MAC Filtering is listed in the Security Policies column, as shown for the CiscoPress-Guest SSID in Figure 12-4.

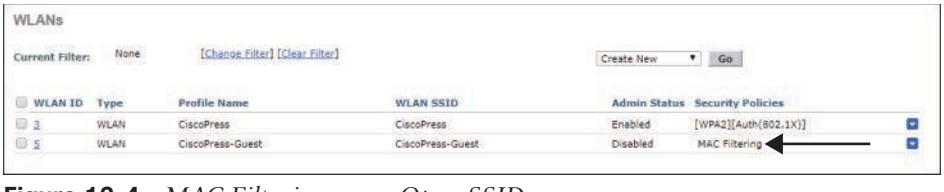


Figure 12-4 *MAC Filtering on an Open SSID*

Validate That ISE NAC Is Enabled on the WLAN

The ISE NAC feature is a very important setting. It is critical to allow for URL redirection, Centralized Web Authentication, posture assessment, native supplicant provisioning, and more.

From the WLC GUI, follow these steps:

Step 1. Navigate to **WLANs >** and select your open SSID.

Step 2. Click on the **Advanced** tab.

Step 3. Ensure that NAC State is forest to **ISE NAC**, as shown in Figure 12-5. In the same screen, ensure that Allow AAA Override is set to **Enabled**.

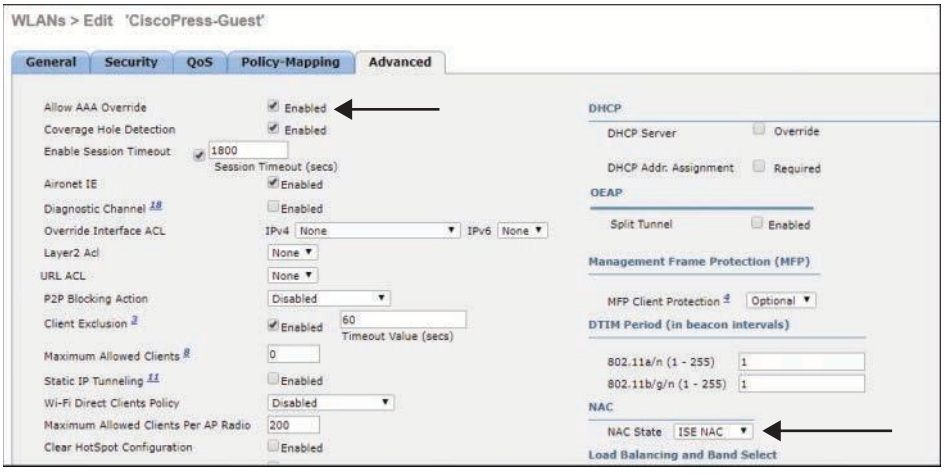


Figure 12-5 *ISE NAC Setting*

Validate That the URL-Redirection ACL Is Configured

The last critical item you need to ensure exists in the WLC configuration is an ACL to use for URL redirections. In Chapter 11, you created an ACL named ACL-WEBAUTH-REDIRECT, which is used to determine what traffic is redirected to the CWA portal with the **deny** statement. Any traffic that is permitted is not redirected.

Key Topic

Unlike IOS-based NADs, AirespaceOS-based wireless controllers use a single ACL to determine which traffic to redirect and which traffic to permit through. In other words, both redirection and traffic filtering are handled by a single ACL. Therefore, the logistics of which traffic is redirected are not the same as with IOS-based devices. With Cisco WLCs, a **deny** statement means that traffic should be redirected. A **permit** statement allows the traffic through the WLC and bypasses the redirection.

In the WLC GUI, follow these steps:

Step 1. Navigate to **Security > Access-Control-Lists > Access-Control Lists**. Ensure that the ACL-WEBAUTH-REDIRECT ACL is in the list, as shown in Figure 12-6.

Step 2. Click this access list and ensure that the entries for your environment are there, as shown in Figure 12-7.

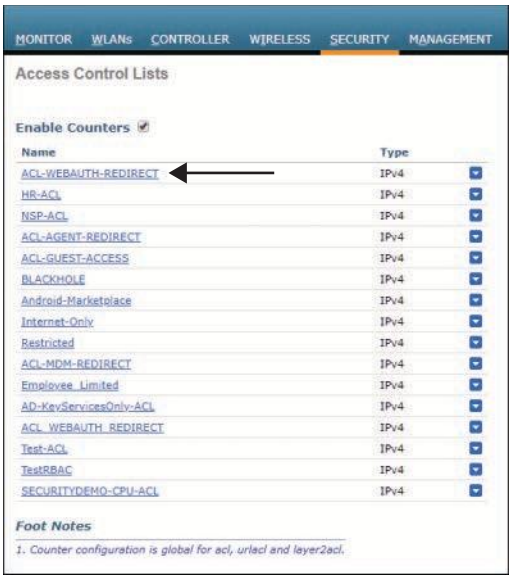
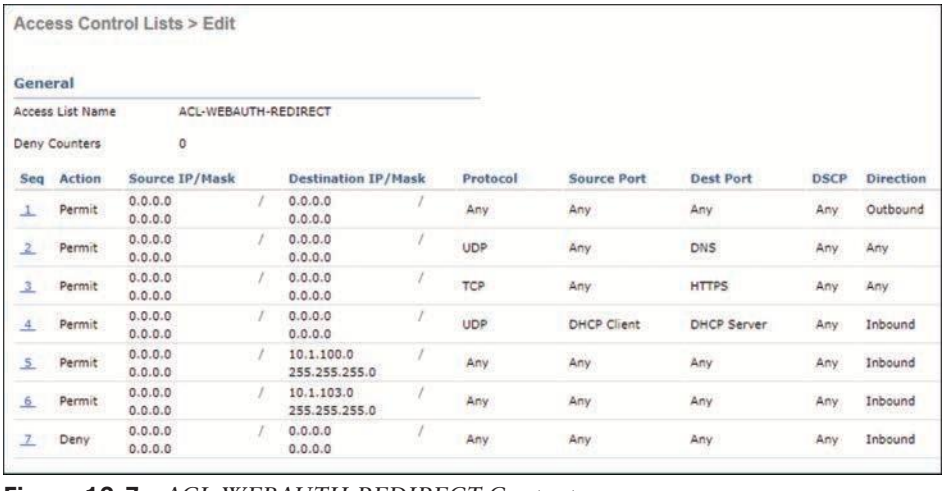


Figure 12-6 *Access Control Lists*



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	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Outbound
2	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	UDP	Any	DNS	Any	Any
3	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	TCP	Any	HTTPS	Any	Any
4	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	UDP	DHCP Client	DHCP Server	Any	Inbound
5	Permit	0.0.0.0 / 0.0.0.0	10.1.100.0 / 255.255.255.0	Any	Any	Any	Any	Inbound
6	Permit	0.0.0.0 / 0.0.0.0	10.1.103.0 / 255.255.255.0	Any	Any	Any	Any	Inbound
7	Deny	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Inbound

Figure 12-7 *ACL-WEBAUTH-REDIRECT Contents*

Configure ISE for Centralized Web Authentication

When you're sure the key elements of the network devices are configured correctly, it's time to ensure that ISE is configured correctly, too. A key change must be made in the authentication policy: An identity source sequence that uses all the appropriate identity stores and the appropriate traffic-filtering dACLs need to be configured. In addition, you need to create the appropriate authorization rules for both before and after Web Authentication.

NOTE Beginning with ISE Version 2.0, ISE contains "smart default" policies. These are preconfigured policies that help customers deploy things like CWA, BYOD, and posture. Due to a communication error between the software developers and Aaron Woland (the man who drove the idea behind smart defaults), those smart defaults include using an ACL named ACL_WEBAUTH_REDIRECT. Notice the underscore instead of the dash. This section does not use the prebuilt rules but shows how to create new rules.

The sections that follow describe the key steps in configuring ISE for Centralized Web Authentication:

Step 1. Configure MAB Continue for the Authentication.

Step 2. Verify the Web Authentication identity source sequence.

Step 3. Configure a dACL for pre-WebAuth authorization.

Step 4. Configure an authorization profile.

Configure MAB Continue for the Authentication

WebAuth is often used for guest access, which means an endpoint is likely to be unknown to ISE when a guest attaches to the network. It is therefore critical to set the identity options to continue when the MAC address is unknown. This has been the default for MAB since ISE Version 2.0, but we examine it anyway to better understand the situation.

In the ISE GUI, follow these steps (see Figure 12-8):

Step 1. Navigate to **Work Centers > Network Access > Policy Sets**.

Step 2. Select the **Default** policy set.

Step 3. Expand the **Authentication Policy** section.

Step 4. In the MAB rule, click **Options** underneath Internal Endpoints.

Step 5. Ensure that If User Not Found is set to **CONTINUE**.

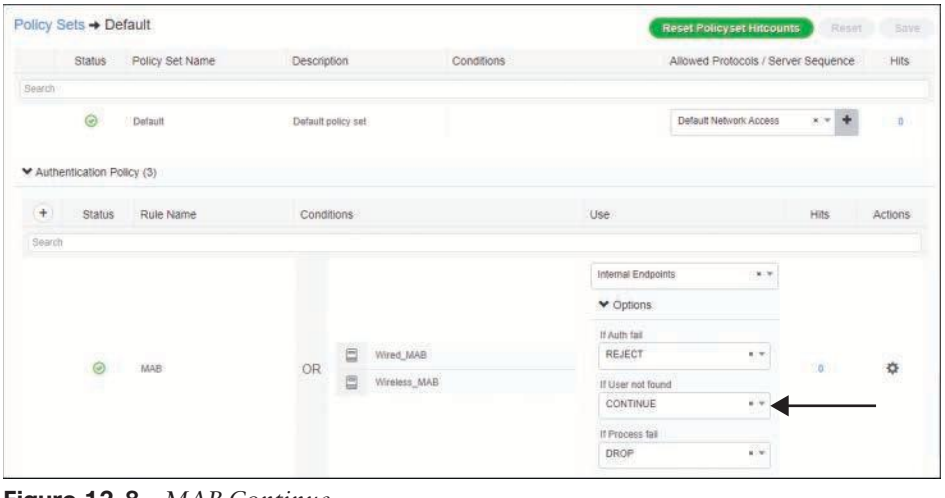


Figure 12-8 *MAB Continue*

Verify the Web Authentication Identity Source Sequence

There is a preconfigured identity source sequence (ISS) named Guest_Portal_Sequence. The default Web Authentication portal is configured to use this ISS. It is configured by default to check Internal Users, Guest Users, and All_AD_Join_Points—in that order.

There is no configuration change required. This sequence, and all the preconfigured sequences since ISE Version 2.0, are ready to use as is. However, just to be sure it does what you need, in the ISE GUI, navigate to Work Centers > Network Access > Identities > Identity Source Sequence and select the Guest_Portal_Sequence, as shown in Figure 12-9.

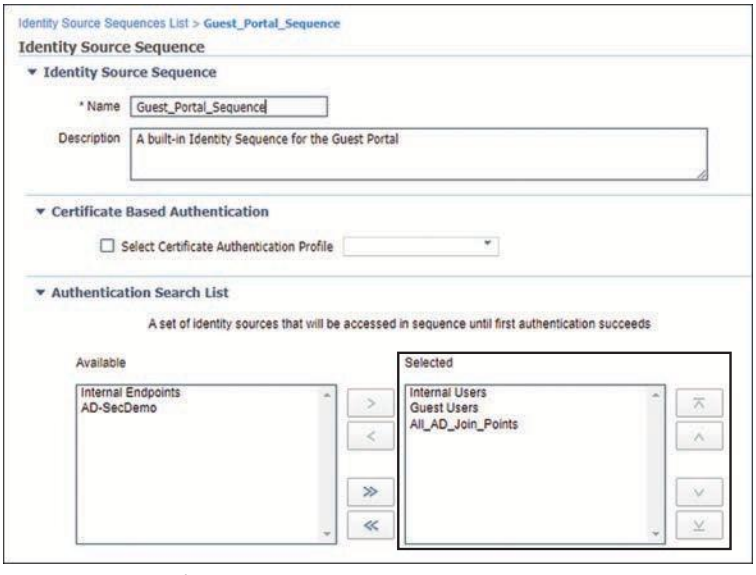


Figure 12-9 *Identity Source Sequences*

Configure a dACL for Pre-WebAuth Authorization

Before a device can reach the CWA portal, it first has to be permitted onto the network. Full network access is not desirable in most cases. For IOS-based devices, a dACL can and should be used to limit the network access.

In the ISE GUI, follow these steps:

Step 1. Navigate to **Work Centers > Network Access > Policy Elements > Results > Downloadable ACLs**, as shown in Figure 12-10.

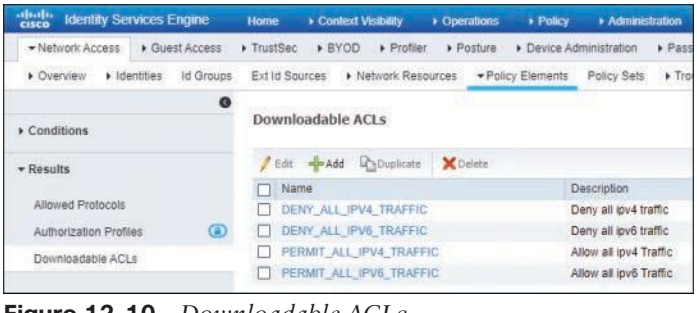


Figure 12-10 *Downloadable ACLs*

Step 2. Click **Add**.

Step 3. Name the new dACL **WebAuth**.

Step 4. Add a description.

Step 5. Configure the ACL to permit traffic to the ISE policy service nodes but deny access to the remainder of the internal network. Figure 12-11 shows what this might look like.



Figure 12-11 *Sample WebAuth dACL*

Step 6. Click **Submit**.

Configure an Authorization Profile

At this point, you are ready to build the authorization profile to allow the end user onto the network, apply the URL redirection to the default CWA portal with the correct URL-Redirect ACL, and apply the dACL to limit network traffic.

In the ISE GUI, follow these steps:

Step 1. Navigate to **Work Centers > Network Access > Policy Elements > Results > Authorization Profiles**, as shown in Figure 12-12.

Step 2. Click **Add**.

Step 3. Name the new Authorization Profile **CWA**.

Step 4. Select the **WebAuth** dACL.

Step 5. Select the **Web Redirection** checkbox and choose **Centralized Web Auth** from the first dropdown. In the ACL text box, type **ACL-WEBAUTH-REDIRECT**. You are using a default WebAuth portal, so ensure that **Sponsored Guest Portal (default)** is selected from the Value dropdown.

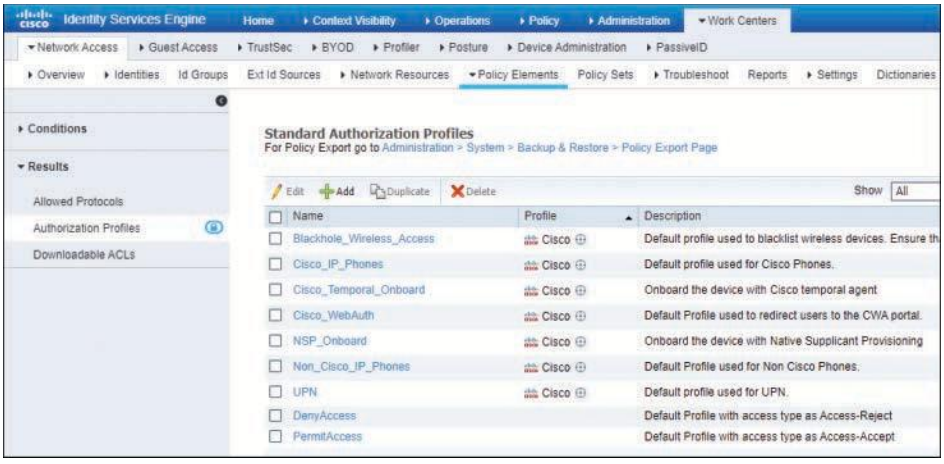


Figure 12-12 *Downloadable ACLs*

Figure 12-13 is a composite image that shows all the key parts in one graphic, including the complete authorization profile.

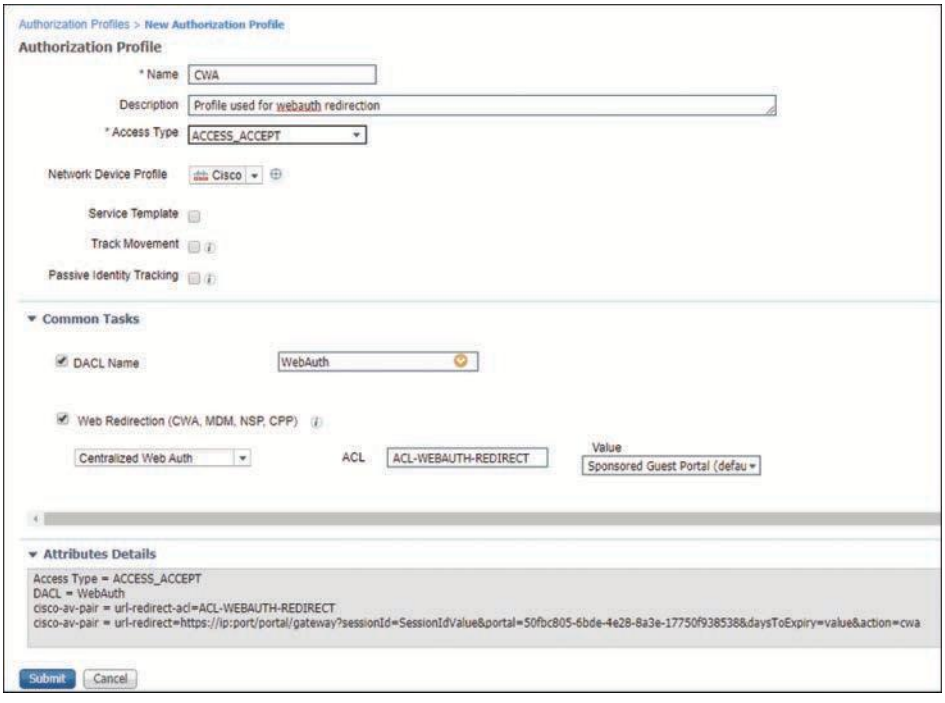


Figure 12-13 *Complete CWA Authorization Profile*

Key Topic

Many organizations implement segmentation with ISE, as it is a security best practice. After all, this is a major use case for ISE and probably a big reason you are reading this book now. With segmentation, users and devices that have an 802.1X supplicant and perform strong authentication should be admitted to the network and kept separate (segmented) from users and devices that have not gone through strong authentication.

If you are going to use VLANs for segmentation, go for it. However, you don't want to change VLANs on a device that is not using a supplicant unless you absolutely have no other choice.

A supplicant detects VLAN changes and renews its IP address in the new VLAN. Most devices that are not using a supplicant don't detect that change of VLAN, and they therefore hold on to the wrong IP address, effectively ensuring that the device will be unable to communicate on the network at all.

Some tricks of the trade can be employed here. (Some in the industry like to call it network gymnastics.) Java and ActiveX applets can be used with WebAuth, and you can play tricks with the DHCP scope in the initial VLAN. The DHCP option is one of the better ones, as a lease can be issued for a very short time, such as 5 minutes. The rules of DHCP client behavior dictate that a client will try to renew its IP address at 50% of the lease time (in this case, 2.5 minutes). ISE has a built-in DHCP server for just such use cases.

However, the recommendation from most implementors is to change VLANs only on clients who use supplicants. Therefore, the default VLAN of your switch ports should remain in a common network segment for guest users and the like, while an authorization profile for employees who gain access via 802.1X should include the VLAN change, and an employee who gains access via CWA should land in the default VLAN of the switch.

Building CWA Authorization Policies

Configuring the authorization policy for centralized Web Authentication is ultimately a two-rule process. This section shows how to create two different authorization rules that will exist toward the end of your authorization policy. They appear at the end of the policy because of the top-down nature of ISE policies and to ensure that CWA is leveraged only when a more specific authorization rule does not apply. If an explicit authorization does not occur, ISE uses the CWA rule to redirect the user to the CWA portal.

The second rule must exist above the redirection rule because this rule is used to assign the right level of access to a user who successfully authenticates to the CWA portal. The second rule must exist above the first rule, or the user will end up in a CWA loop.

Cisco has included preconfigured authorization rules with ISE for wireless guest access Web Authentication. These rules, which are shown in Figure 12-14, are disabled by default.



Figure 12-14 *Preconfigured Web Authentication Rules*

You can leverage these prebuilt rules, but how would that help you learn and prepare for the SISE 300-715 exam? Instead of leveraging those prebuilt rules, which could shorten your configuration time dramatically, in the following sections you will see how to build your own rules.

Create the Rule to Redirect Users to the CWA Portal

The first rule to create is one that redirects unauthenticated users to the CWA portal, where they are required to authenticate interactively.

In the ISE GUI, follow these steps:

Step 1. Navigate to **Work Centers > Network Access > Policy Sets**.

Step 2. Drill down into your default policy set (or the policy set that is in use for your deployment at this time).

Step 3. Insert a new rule above the Basic_Authenticated_Access rule and name the new rule **WebAuth**, as shown in Figure 12-15.

Step 4. For the conditions, select two existing compound conditions from the library: **Wired_MAB** and **Wireless_MAB**. Ensure that the **OR** operator is used with the conditions, as shown in Figure 12-15.

Step 5. Use the **CWA** authorization profile you created previously for the result, as shown in Figure 12-15.

Step 6. Click **Save**.

Figure 12-15 shows the completed WebAuth authorization rule.

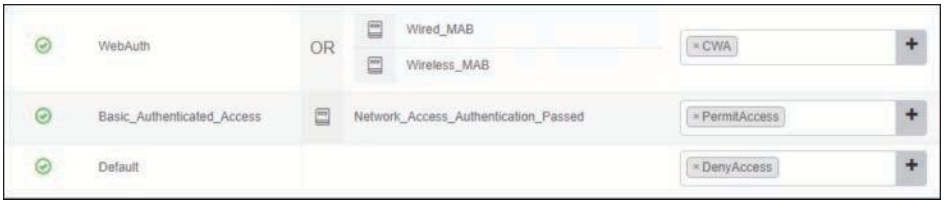


Figure 12-15 *Completed WebAuth Authorization Rule*

Create the Rules to Authorize Users Who Authenticate via CWA

The second rule needs to allow a user who authenticates via WebAuth to have specific access to the network. The number of rules created depends on the needs of your organization. For the purposes of this chapter, you will create only one rule, for employees. (Guest users are covered in Chapter 13.)

In this case, you need to construct a new authorization rule that will allow employees (users who are members of the Employees group in Active Directory) who have successfully authenticated through the web portal to have network access.

To accomplish this task, you can use a dictionary item named ***Guest Flow*** in your rule. ISE uses this dictionary item to identify when an authentication has occurred via an ISE web portal.

Technically, you are not required to use the Guest Flow attribute in your conditions, and an employee logging in through CWA will still land on any rule that matches your employee condition. However, for good security practice, you should be specific and construct an authorization rule that allows employees (users who belong to the Active Directory group named Employees) who have successfully authenticated through the web portal to have Internet-only network access.

In the ISE GUI, follow these steps:

Step 1. Navigate to **Work Centers > Network Access > Policy Sets**.

Step 2. Drill down into your default policy set (or the policy set that is in use for your deployment at this time).

Step 3. Insert a new rule above the WebAuth rule and name it **Employee CWA**.

Step 4. Use **GuestFlow** as the first condition for the rule.

Step 5. Add another condition with the **AND** operator.

Step 6. Select the Active Directory group named **Employees** as the second condition.

Step 7. Use the previously created authorization profile named **Internet-Only**.

Step 8. Select the **Employees** security group tag.

Step 9. Click **Save**.

Figure 12-16 shows the completed Employee_CWA rule.



Figure 12-16 *Completed Employee_CWA Rule*

NOTE It is impossible to stress enough times that you should not leverage VLAN assignment with CWA. The authors of this book have met with countless customers who have deployed ISE and have helped many of them deal with exactly that problem.

Verifying Centralized Web Authentication

You've already gone through a fair bit of configuration in this chapter. Now that you've completed the CWA configurations, you're ready to see it all in action.

There are a number of locations to verify the actions. You can examine the effect the user experiences on the client, check Live Log in ISE, run **show** commands on the wired switch, or even examine the client details on the WLC.

Check the Experience from the Client

A quick way to see if your configuration is working is to try opening a web browser on the client machine and see if the browser is redirected to a portal.

Figure 12-17 shows the client experience on a wired Windows client being redirected to the CWA portal and the user entering credentials in the login form.

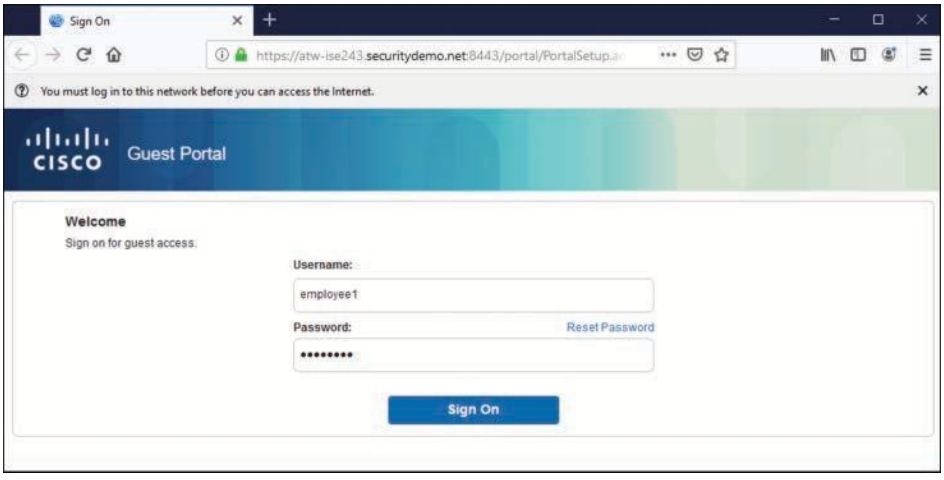


Figure 12-17 *Browser Redirected to the CWA Portal*

Figure 12-18 displays the acceptable use policy that is shown to a user who submits valid authentication credentials. Figures 12-19 and 12-20 show the screens that follow, which indicate that it is now possible for the user to access the Internet. Finally, Figure 12-21 shows the successful connection to the Internet, as the user browses to http://www.cisco.com.

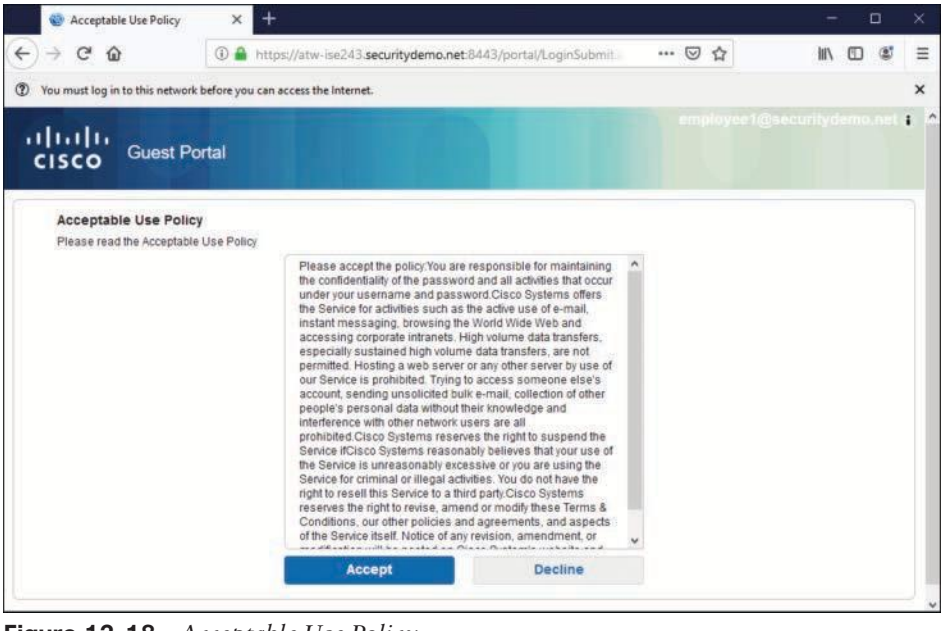


Figure 12-18 *Acceptable Use Policy*

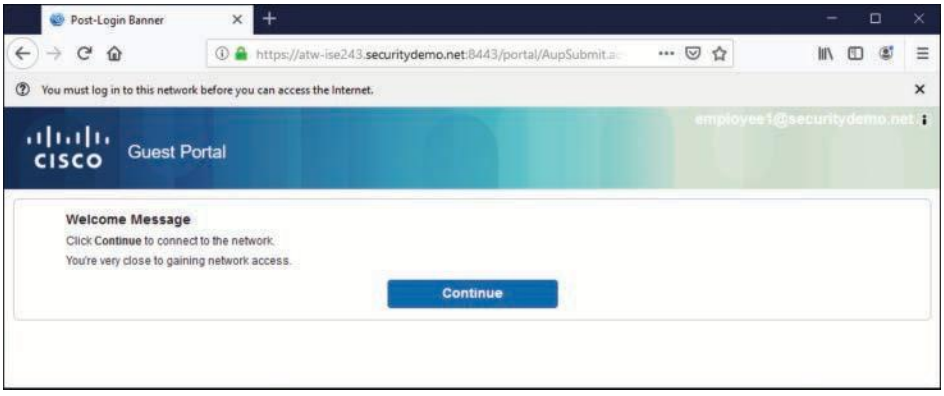


Figure 12-19 *Browser Redirected to the CWA Portal*



Figure 12-20 *Success*

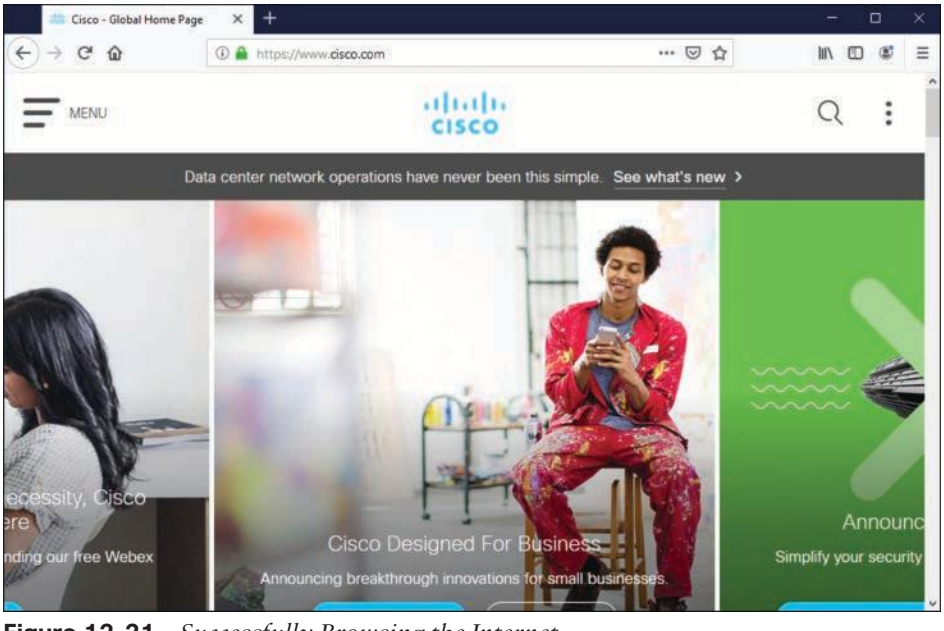


Figure 12-21 *Successfully Browsing the Internet*

Verify CWA Through the ISE UI

A logical way to verify WebAuth configuration would be to look at the centralized policy server. ISE has a number of tools that can be used to verify WebAuth. The most common is Live Log, but many other tools can be used, including TCPDump (although this chapter covers only Live Log). For more on those other tools, see Chapter 21, "Troubleshooting Tools."

Check Live Log

Cisco ISE has a phenomenally useful built-in tool called Live Log. Live Log provides a near-real-time view of all incoming authentications, Change of Authorization (CoA), and more. In this section, you will follow the client experience from the ISE management console.

Key Topic

Figure 12-22 highlights the process.

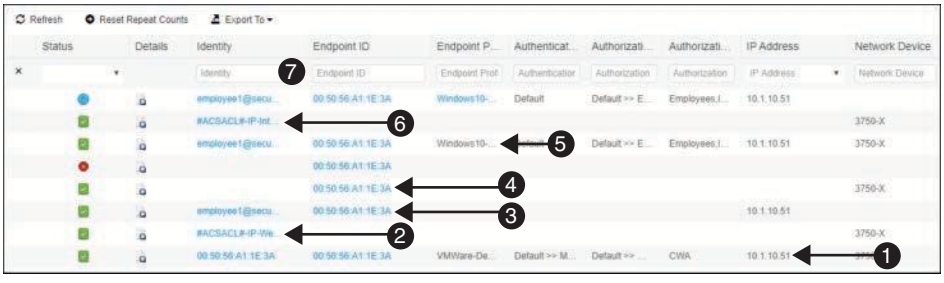


Figure 12-22 *Live Log*

The following points correspond to the numbers in the figure:

1. The initial MAB has been assigned the CWA authorization result.
2. Immediately following the successful authorization, you see the successful download of the dACL.
3. After the end user enters credentials and clicks Submit, those credentials are recorded.
4. Immediately after the credentials are authenticated, a CoA-ReAuth is sent to the switch. The CoA is a key function that causes the switch to reauthenticate the endpoint without starting a new session.
5. That reauthentication means the switch sends another MAB request to ISE, where the Web Authentication from the centralized portal is bound to the MAB request from the switch.
6. Due to the correlation of the centralized Web Authentication to the MAB authentication request, the employee is assigned the Internet_Only authorization profile, which is followed immediately by the successful download of the Internet_Only dACL.
7. Finally, a RADIUS accounting packet is sent from the switch to ISE, confirming the full session establishment.

Check the NAD

Checking the device that is performing the enforcement should be a good way to confirm that CWA is working. In this section, you will see how to examine the authorization result on a Cisco switch and a Cisco WLC.

show Commands on the Wired Switch

The go-to command on a Cisco switch is **show authentication session interface** [*interface-name*]. This provides a lot of valuable information. Example 12-1 shows the command and its output for the endpoint as it was being redirected to the CWA portal.

Highlighted in this example are the MAC address, IP address, dACL (listed as an ACS ACL), URL-Redirect ACL, and the URL the end user is being redirected to. Notice that the username is the MAC address of the endpoint, which is a clear sign that the authentication performed was really a MAB. At the end of the screen output, you also see that MAB has the state Authc Success.

Example 12-1 *show authentication session interface g1/0/3 Command Output*

```
3750-X#show authen sessions int g1/0/1
            Interface:  GigabitEthernet1/0/1
          MAC Address:  0050.56a1.1e3a
           IP Address:  10.1.10.51
            User-Name:  00-50-56-A1-1E-3A
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
        Authorized By:  Authentication Server
          Vlan Policy:  N/A
              ACS ACL:  xACSACLx-IP-WebAuth-5e2a155e
     URL Redirect ACL:  ACL-WEBAUTH-REDIRECT
         URL Redirect:  https://atw-ise243.securitydemo.net:8443/portal/
gateway?sessionId=C0A8FE0100000271FEF432A8&portal=50fbc805-6bde-4e28-8a3e-17750f9385
38&action=cwa&token=b296cdf51b7985efc8adace571ce4c29
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  C0A8FE0100000271FEF432A8
      Acct Session ID:  0x000003AE
               Handle:  0xBF000272

Runnable methods list:

       Method   State
       mab      Authc Success
       dot1x    Not run
```

Example 12-2 shows the command and its output for the endpoint after the user has successfully completed the Web Authentication.

Example 12-2 *show authentication session interface g1/0/3 Command Output*

```
3750-X#show authen sessions int g1/0/1
            Interface:  GigabitEthernet1/0/1
          MAC Address:  0050.56a1.1e3a
           IP Address:  10.1.10.51
            User-Name:  employee1@securitydemo.net
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
        Authorized By:  Authentication Server
          Vlan Policy:  N/A
              ACS ACL:  xACSACLx-IP-Internet_Only-5e606c90
                  SGT:  0004-00
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  C0A8FE0100000271FEF432A8
      Acct Session ID:  0x000003AE
               Handle:  0xBF000272

Runnable methods list:
       Method   State
       mab      Authc Success
       dot1x    Not run
```

Key Topic

Notice the differences between Examples 12-1 and 12-2. Specifically, notice that in Example 12-2, the username is filled in (and no longer listed as the device's MAC address), and there is no longer any redirection happening. However, the authentication method for mab is still listed as Authc Success. This is because a switch still considers CWA to be MAB rather than a separate authentication. ISE is responsible for binding the username to the MAB session.

Viewing the Client Details on the WLC

With the WLC, you can navigate to Monitor > Clients to see a list of all clients currently associated to that controller. Clicking the MAC address brings up the details about the client and its association, including authentication information such as the redirection and run state.

Key Topic

When you implement ISE, in addition to needing to know how ISE works, it is especially important to have a clear understanding of how the network devices work with ISE.

Figure 12-23 is a screenshot from the WLC GUI that shows the details for a client that is currently being redirected to a CWA portal, with a cropped focus on the Security Information section.

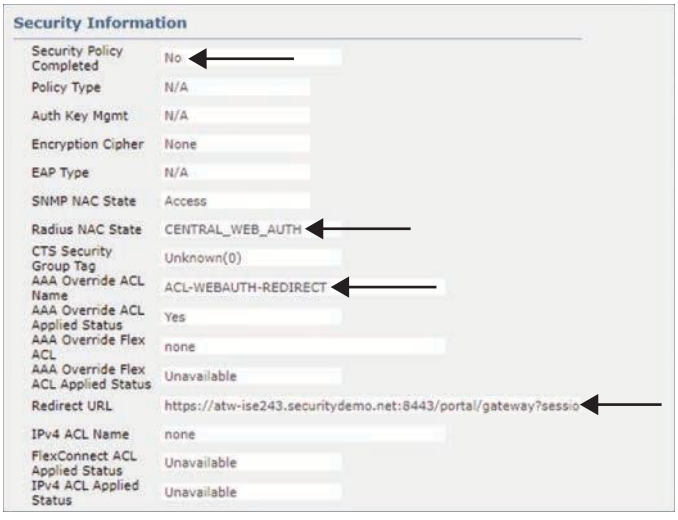


Figure 12-23 *Security Information Section of Client Details – CWA*

Figure 12-23 highlights the important sections:

- The RADIUS NAC state is set to **CENTRAL_WEB_AUTH**.
- The Security Policy Completed state is currently **No**.
- There is an AAA override ACL named **ACL-WEBAUTH-REDIRECT**.
- The redirect URL contains the dynamic URL of the active ISE PSN for this client's session.

Figure 12-24 is a screenshot from the WLC GUI that shows the details for the same client after it has gone through a successful Web Authentication. The screenshot has a cropped focus on the Security Information section.

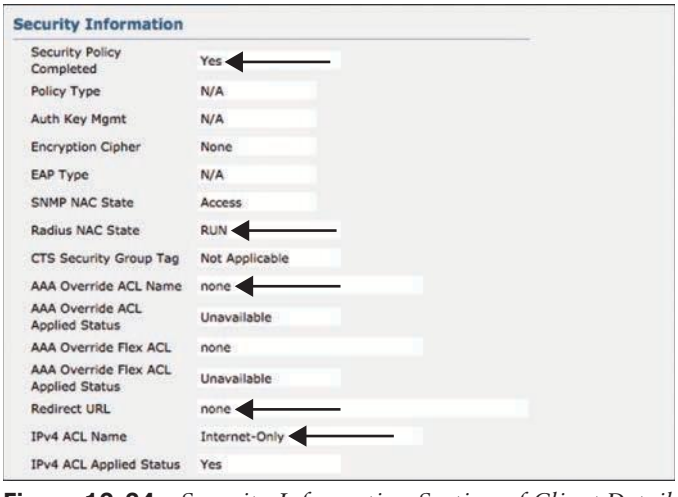


Figure 12-24 *Security Information Section of Client Details – Run State*

Figure 12-24 highlights the important sections:

- The RADIUS NAC state is now **RUN**. This is a key setting: The redirection will never work if the state is RUN since RUN is the final state.
- The Security Policy Completed state is currently **Yes**.
- There is no AAA override ACL in the RUN state.
- There is no redirect URL in the RUN state.
- There is an Internet-only ACL applied.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 12-2 lists these key topics and the page number on which each is found.

Key Topic

Table 12-2 Key Topics

Key Topic Element	Description	Page
Paragraph	Local Web Authentication	310
Paragraph	Disconnecting management traffic from the web server in order to isolate and protect a switch	311
Paragraph	Using MAB with CWA	314
Paragraph	Traffic filtering and traffic matching	314
Paragraph	Traffic filtering and traffic matching combined with the WLC	316
Paragraph	Segmentation	321
Figure 12-22	The steps involved with CWA	327
Paragraph	A switch recognizing CWA as a MAB flow	329
Paragraph	The steps involved with CWA on the WLC	329

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

Web Authentication (WebAuth), Local Web Authentication (LWA), Centralized Web Authentication (CWA), Guest Flow

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is the final state of a client connected to a Cisco wireless LAN controller?
2. What capability in a Cisco NAD enables the client to be sent to a Web Authentication portal?
3. What authentication method is displayed on a switch for a user who has successfully authenticated via CWA?
4. Where is the URL-Redirect ACL created?
5. What is different about URL redirection when comparing how a switch uses ACLs to how a WLC uses ACLs?

This page intentionally left blank



CHAPTER 13

Guest Services

This chapter covers the following topics:

- Web Authentication scenarios
- Configuring Centralized Web Authentication
- Building CWA authorization rules
- Verifying Centralized Web Authentication

In corporate settings, the concept of providing access to guests has been around for a very long time. Consider it from a physical access perspective: A network administrator may need to invite a Cisco account team to meet in a conference room to learn about the latest offerings and features from Cisco. When the Cisco account team members arrive, they may have to check in with a human being who runs the "front desk" in the lobby, provide some form of identification, and tell the lobby administrator what company employee they are there to visit. The lobby administrator then calls the employee to say that the guests have arrived.

In many cases, each visitor is given a temporary badge or sticker that allows entry to the building and access to an elevator and identifies the person as a guest to any employees that pass in the hallways.

Once upon a time, providing such visitor access to the Internet was considered a luxury, but today it is now considered a business-critical function. In fact, guest Internet access was a major driving force in the adoption of Wi-Fi, and solutions like ISE, all over the world. Financial institutions used to install anti-Wi-Fi solutions that would prevent the radio frequencies from operating. Guest access was a driving need that changed the approach dramatically, and now institutions have come to embrace Wi-Fi.

In this chapter, you will learn about the different types of guest access supported by ISE and how to configure them.

NOTE This chapter was written based on the assumption that the switches and WLCs have been configured as described in Chapter 11, "Implement Wired and Wireless Authentication." If you have not already configured your network devices for authentication, none of the configurations in this chapter will work, and you should revisit Chapter 11.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge

of the topics, read the entire chapter. Table 13-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 13-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Guest Services Overview	1
Guest Portal Types	2, 3, 4, 9
Configuring Guest Portals and Authorization Rules	5
Sponsors	6–8
SAML Authentication	10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. In the Guest Access work center, why are endpoint identities listed?
   - **a.** They are listed for visibility, so you can run the report on the endpoints guests like to use.
   - **b.** An administrator must add each endpoint to ISE so that a guest can gain network access.
   - **c.** Endpoints are not listed in the Guest Access work center.
   - **d.** An endpoint is assigned network access when it joins a network.
2. Which statement best describes a guest from ISE's perspective?
   - **a.** Someone who comes to your house for dinner
   - **b.** A user or device that is not a corporate employee or corporate-owned endpoint
   - **c.** An employee who brings a home laptop into the office
   - **d.** A network for employees to use when they want to bypass the corporate network security controls
3. What are the default guest portal types? (Choose three.)
   - **a.** Self-registered guest portal
   - **b.** SAML portal
   - **c.** Sponsored guest portal
   - **d.** Hotspot guest portal
   - **e.** Social login guest portal

4. What are default guest types? (Choose four.)
   a. Social
   b. SAML
   c. Daily
   d. Monthly
   e. Contractor
   f. Weekly
5. Which of the following statements is most accurate regarding a guest and a guest portal?
   a. A guest who connects to the network must type guest.cisco.com into the browser to enter the guest portal.
   b. A guest who connects to the network with an unknown device is redirected to the guest portal.
   c. A sponsor tells the guest what address to type into a browser to reach the guest portal.
   d. Guest portals are used only with wireless networks.
6. Which of the following statements best describes the GROUP_ACCOUNTS sponsor group?
   a. Sponsors of this type can edit only groups, not users.
   b. Sponsors in this group can edit all accounts but can only delete accounts from their own group.
   c. Sponsors of this type may only approve account requests that are routed to their group.
   d. Sponsors in this group can edit the accounts they create, as well as any other accounts created by sponsors in the same group.
7. What is the result when a sponsor who is a member of multiple sponsor groups logs in to the sponsor portal?
   a. Sponsor group permissions are assigned as the least privilege.
   b. Sponsor group permissions are additive, and the sponsor receives the permissions of all the groups.
   c. A sponsor may be a member of only one group.
   d. There are sponsors, but there are no sponsor groups.
8. Which of the following statements is most accurate regarding a sponsor and the sponsor portal?
   a. When a guest connects to the network, the sponsor is notified and must come to the lobby to escort the guest.
   b. There is no sponsor portal; there are only guest portals.
   c. The sponsor portal provides a field for the FQDN of the sponsor portal. The sponsor enters that FQDN in a browser and is directed to the correct portal.
   d. The sponsor must either click a link or type in the FQDN of a PSN with the correct port (:8445 by default).

9. Which of the following statements about the guest and sponsor portals is most accurate?
   a. The portals are fully customizable and have prebuilt templates for basic customizations.
   b. The portals are fully customizable, but you must import all the files after externally editing them.
   c. The portals are customizable only if they are hosted on an external server, not the PSN.
   d. The guest portal is fully customizable, but the sponsor portal is not.
10. Which of the following is a true statement about the use of SAML authentication?
   a. SAML is appended to the identity source sequence to attempt SAML last, after trying the other identity sources.
   b. SAML authentication can be used by the sponsor portal but not the guest portals.
   c. SAML authentication may be used by the sponsor portal and the guest portals.
   d. SAML is prepended to the identity source sequence to attempt SAML first, before trying the other identity sources.

Foundation Topics

Guest Services Overview

ISE Version 1.0 was revolutionary: It was a single solution that combined a powerful RADIUS server with a profiling solution, a posture service, and a set of guest services. However, those guest services left a lot to be desired. To address the issues, Chip Current and Paul Russell led a large team of very talented people who changed ISE forever. This team completely changed the guest services in ISE Version 1.3.

It was an amazingly exciting time, and a full user-experience (UX) team came together with product managers and technical marketing engineers to interview countless customers on how they used guest services, what types of visitors they had coming to their offices, and what they needed within their portals. The usability of the product in development was tested over and again with blind user testing to determine how intuitive the workflow was.

This work resulted in changes to all of ISE's UX, even though that wasn't the initial intent. The need for work centers came out of this effort, and Guest Access was the first work center in ISE. In ISE Version 2.7, the work centers are well understood and commonplace, but back when work centers were created, the idea of a single location within the UI where all tasks related to guest services could be accomplished was brand new.

In ISE, all work centers are designed to include every aspect required to complete the related tasks. The tasks are also in a specific order, ordered from left to right and top to bottom.

To begin using the Guest Access work center, navigate to Work Centers > Guest Access in the ISE UI. As shown in Figure 13-1, you first see an overview page that illustrates the steps needed to get guest access up and running. The page has three top-level tasks: Prepare, Define, and Go Live & Monitor.

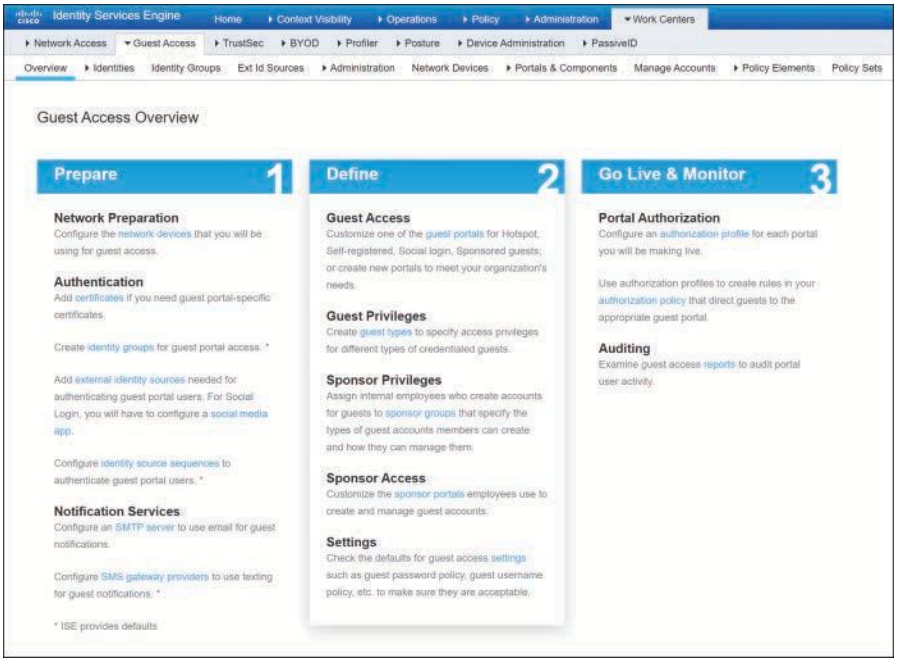


Figure 13-1 *Guest Access Work Center Overview Page*

Following the left-to-right, top-down approach to navigating a work center, the next tab is the Identities tab (see Figure 13-2). The left-hand navigation pane lists identity types, starting with endpoint identities. The main part of the screen shows a filtered view of the context visibility inventory, focusing on endpoints being used by guests. These endpoints are often added to groups after a guest registers in order to assign the guest access to the endpoint when it joins a network.

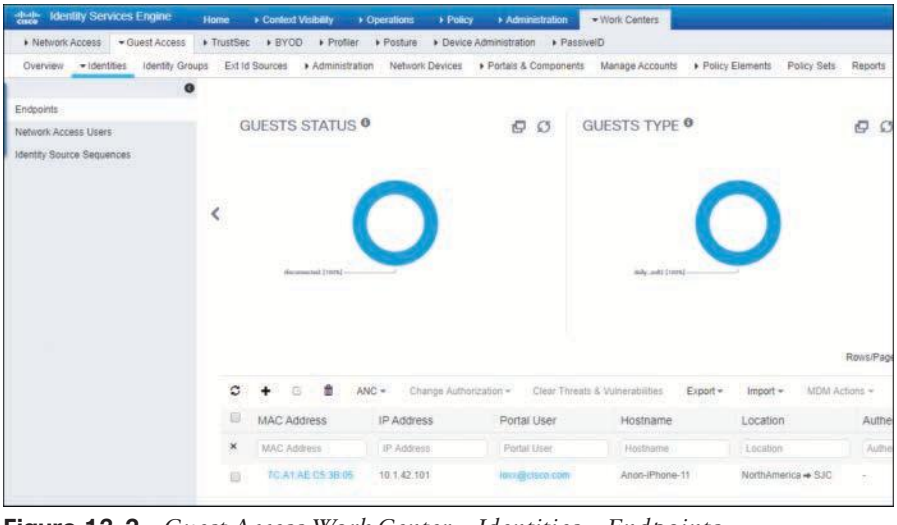


Figure 13-2 *Guest Access Work Center > Identities > Endpoints*

The next identity type as we move from the top down in the left-hand navigation pane is the network access users. These users are not leveraged very often with ISE deployments, except with service accounts or for device administration. You might think that this is where you would find guest users when those accounts get provisioned, but that is not the case. There is an identity store for guest users, but it is separate from this internal identity store.

The next identity type in the left-hand navigation pane is identity source sequence (ISS) identities. By examining the ISSs, you can determine whether there is an identity store for those guest user accounts. Figure 13-3 shows a list of preconfigured ISSs, and Figure 13-4 shows the details of the ISS that you used in Chapter 12, "Web Authentication," when you configured Centralized Web Authentication.

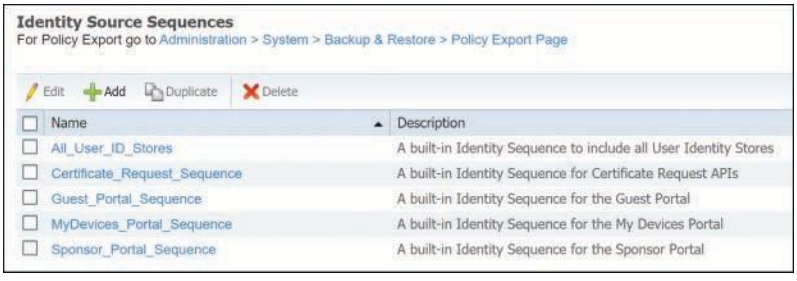


Figure 13-3 *Guest Access Work Center > Identities > Identity Source Sequences*



Figure 13-4 *All_User_ID_Stores ISS*

Notice in Figure 13-4 that Guest Users is an identity source in the selected list on the right-hand side.

There is a prebuilt ISS named Guest_Portal_Sequence that is nearly identical to the All_User_ID_Stores ISS. The biggest difference is that Guest_Portal_Sequence does not include any certificate authentication profiles.

The next tab in the Guest Access work center is Identity Groups. This is not specific to guest access but is exactly the same as the Identity Groups tab in the Network Access work center. Click on User Identity Groups on the left side, as shown in Figure 13-5, and you see groups for each of the guest account types. We will go over these different guest types later in this chapter.

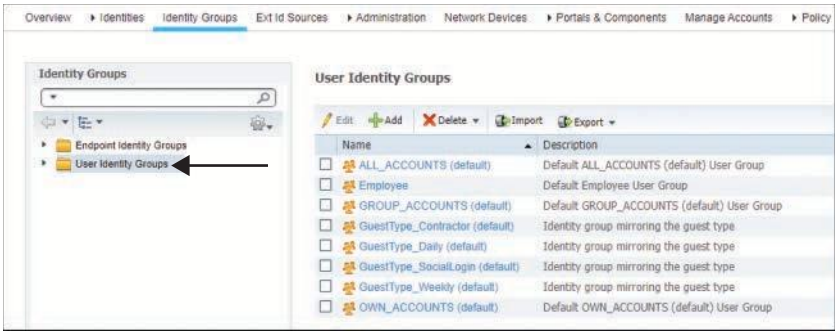


Figure 13-5 *User Identity Groups*

Continuing the left-to-right navigation in the work center, Ext Id Sources is the next tab. We will come back to this tab later, when we look at configuring SAML.

Key Topic

The next tab, Administration, has a section that focuses on certificates, which are critical to guest services. As you saw in Chapter 12, users get redirected to web pages (also known as portals), and those web pages are protected and identified using certificates. It's important to understand that ISE hosts a lot of web pages, and not a single one of them is accessible without HTTPS.

You must ensure that your guest users do not experience any errors when being redirected to these secure guest portals, and the only way to ensure a smooth transition is to leverage a signed certificate from a trusted public certificate authority.

Figure 13-6 shows the System Certificates section of the Administration tab and highlights the use of a certificate for guest portals that has been signed by a trusted Internet root CA: SSL.com.

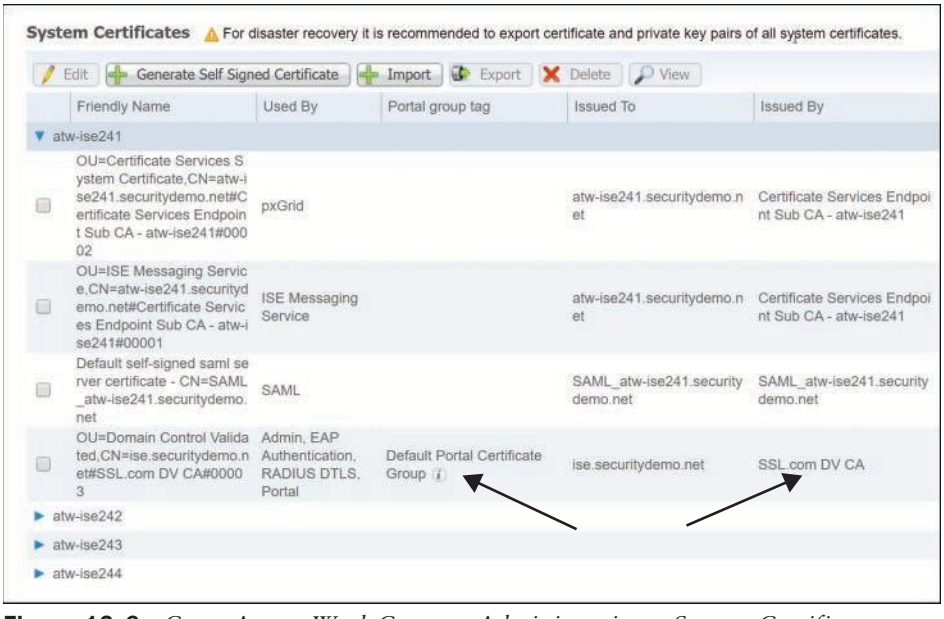


Figure 13-6 *Guest Access Work Center > Administration > System Certificates*



The Administration tab also includes a section for configuration of SMTP servers and SMS gateway providers. These are both quite important as they help guests know what their credentials are, but we save discussion of them for later in the chapter.

You have learned about authentication, including WebAuth, in earlier chapters, so there is no need to dive into the Network Devices tab. The following section discusses the Portals & Components tab.

Portals, Portals, and More Portals!

In the Guest Access work center, Portals & Components is one of the most important tabs. The majority of the "guest" work is really done here, whereas network setup and other supporting configuration processes are accomplished in the other tabs.

Reconsider the scenario presented earlier in this chapter, in which members of the Cisco account team go onsite to visit a Cisco customer. In this scenario, the Cisco account team members are the guests, and the network administrator, who is an employee of the Cisco customer, is sponsoring the team members' access to the building or campus.

Just as these guests need physical guest access, many guests need guest Internet access. Vendors, contractors, and visitors may need access to a company's physical property as well as the Internet (and perhaps the company network); they need credentials (such as a guest badge) and, in many cases, also need a sponsoring employee who is responsible for the actions of the visitors while they are using company property and Internet access.



The Portals & Components tab of the Guest Access work center has two different configuration categories:

- **Guest:** A *guest* is a user or device that is not a corporate employee or corporate-owned endpoint. This user type could be a visitor, contractor, consultant, or customer. Guest accounts are created to allow these users to authenticate and gain Internet access.
- **Sponsor:** A *sponsor* is an employee who introduces and supports a guest. This is also an individual who can be identified as the point person during an investigation. For example, say that an employee of Bob's Company creates a guest account for a fake guest, and then that employee proceeds to leverage the guest account to exfiltrate data or break the code of business conduct; the HR and/or incident response teams can then contact the sponsor to see who he or she supposedly created that account for.

Guest Portal Types

A *portal* is basically a website. It is a series of hosted web pages that an end user can interact with to accomplish a goal. It may also be a series of web pages that employees can use to manage their own devices or guests that they have sponsored. There are three main types of guest portals—hotspot guest portals, self-registered guest portals, and sponsored guest portals—and each of these types provides different user experiences. Figure 13-7 shows the Portals & Components page, from which you can access the three portal types.

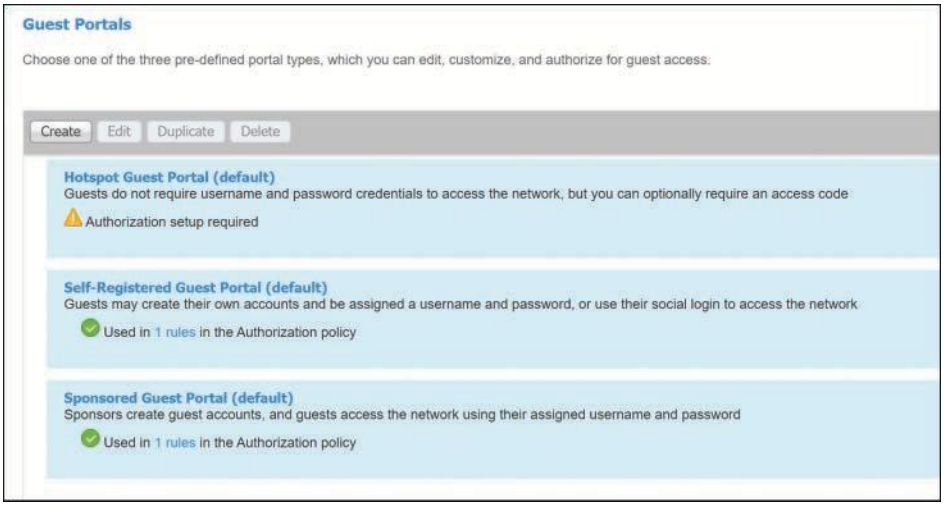


Figure 13-7 *Guest Portal Types*

The following sections discuss these portal types, and later in the chapter you will see how to configure each type.

Hotspot Guest Portal

What used to be known as a device registration portal (DRP) is now known as a hotspot guest portal. A hotspot portal is designed to allow a guest user to access the Internet without needing credentials such as a username and password. With a hotspot portal, a user can just click a button and perhaps agree to an acceptable use policy (AUP) to get uninterrupted access to the Internet; behind the scenes, the user's MAC address is recorded. Hotspot portals are very common and provide what is often referred to as a "coffee shop experience" for accessing the Internet. Chances are that you have experienced this type of portal yourself.

Self-Registered Guest Portal

With a self-registered portal, guests are permitted to come onsite and request access in the portal. Guests may create their own accounts and be assigned a username and password or use a social login to access the network. No sponsorship is required.

Unlike a hotspot portal, a self-registered portal provides some sort of audit trail, but it does not provide as much accountability as a sponsored guest portal. For example, how would you prevent or even audit that an employee used the self-registered guest access function to break company conduct regulations?

Sponsored Guest Portal

The sponsored portal type, which requires credentials (username and password), is among the most commonly used portals and is also the portal type used with Centralized Web Authentication (CWA). In fact, you used a sponsored guest portal with CWA in Chapter 12. With this type of portal, sponsors create or approve accounts, and guests sign in with those credentials.

Guest Types

There are different types of guest accounts for different needs. From a security perspective, businesses typically want guest accounts to expire so they no longer need to be maintained and so there is no concern about usernames and passwords being shared around.

Key Topic

The primary difference between guest types is in how long the accounts are permitted to be active. Default guest types help ISE customers by providing templates, based on the most popular needs from Cisco's customer research, but every setting is configurable. Different guest types have different settings for how many simultaneous logins may be used, the number of endpoints each guest may register with ISE, and the messaging to use with the guest-related account creation and renewal.

The authors of this book have worked with countless organizations around the world, and while each organization is unique, organizations across the board have many similar needs. Let's look at an example of a common practice that requires more than one guest type. Say that any employee is authorized to invite a guest and provide Internet access for a short amount of time, but a manager or director must provide access for any time period longer than a few days or a week.

Figure 13-8 shows the Portals & Components > Guest Types page, which includes four built-in types of guest accounts.

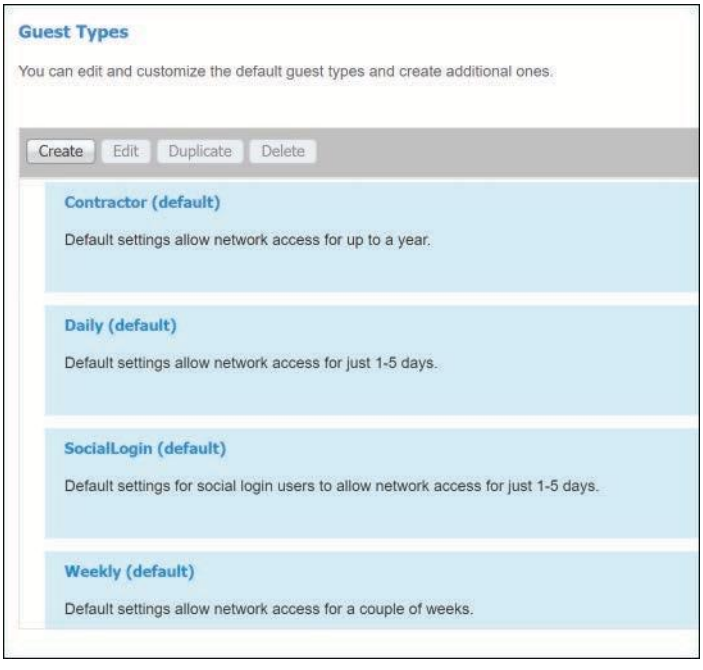


Figure 13-8 *Guest Types*

Key Topic

The following sections dive into the default guest types that ISE provides out of the box.

Contractor

The contractor type of guest account is designed to be active for long periods of time—up to a year. This type of account is for a non-employee who is acting like an employee, such as an onsite software developer. Such a person is not an employee and so wouldn't get an account in the HR system or the Active Directory identity store.

To edit the contractor guest type, navigate to Work Centers > Guest Access > Portals & Components > Guest Types, as shown in Figure 13-9.

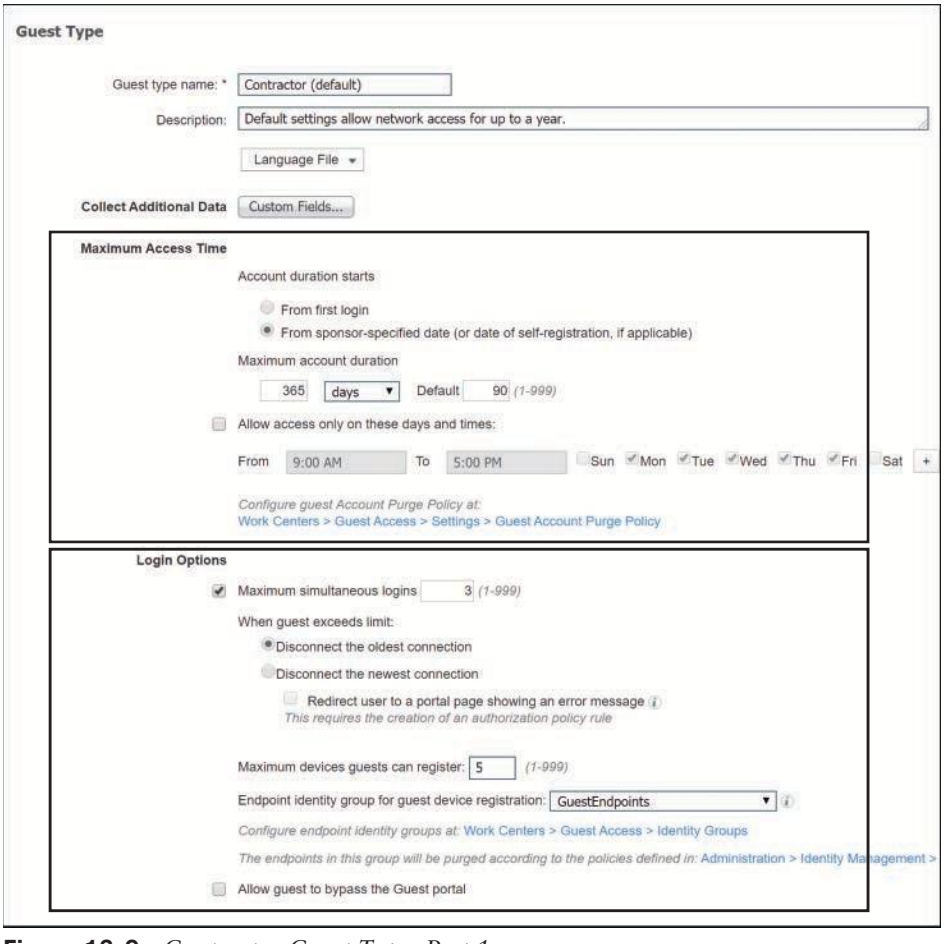


Figure 13-9 *Contractor Guest Type, Part 1*

Following along on your own ISE UI, or examining the Maximum Access Time section in the middle of Figure 13-9, notice the max access time. The default duration for a contractor account is 90 days, and the maximum is set to 365 days. Both duration values could be set to any number between 1 and 999.

In addition to configuring account duration, you can restrict the account type to being used only during work hours and workdays.

The bottom section in Figure 13-9 shows the login options for the accounts. The Maximum Simultaneous Logins setting is important because a guest user may have multiple devices, such as a laptop, a tablet, and a phone. Organizations might want to allow multiple devices but limit them to a reasonable number to prevent misuse of the guest account.

An often overlooked but very important setting is the When Guest Exceeds Limit setting. Disconnecting the oldest session is the default and most common use case. It assumes that the newest session is the most important one. The other option, Disconnect the Newest Session, should be used when you wish to redirect the user to a portal and display an error message. Note that this redirection and error message will work only if you create the appropriate corresponding authorization rule.

The number of devices a guest user can register and which endpoint identity group stores the device's MAC address is also quite important. By registering the MAC address within an endpoint identity group, you can use the Allow Guest to Bypass the Guest Portal setting to create simple authorization to permit the devices to access the network without forcing the guest user to go through Web Authentication every time the guest connects to the Wi-Fi or roams to a different wireless controller. In other words, it ensures that the guest's experience on the Wi-Fi is smooth.

The endpoint identity groups are occasionally purged, and you configure the policies that dictate how often each endpoint group is purged under Administration > Identity Management > Settings > Endpoint Purge, as shown in Figure 13-10.

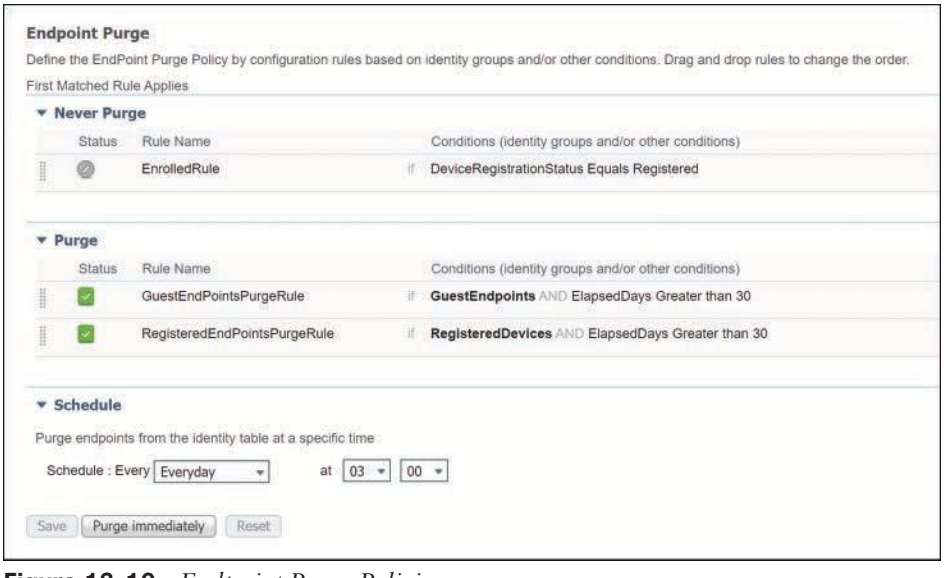


Figure 13-10 *Endpoint Purge Policies*

The purge policies are configurable per endpoint identity group, and specific endpoints (such as endpoints with attribute DeviceRegistrationStatus equal to Registered) can also be configured to never purge, as shown in Figure 13-10. You can also see in Figure 13-10 the default purge status of the GuestEndpoints identity group, which is set to purge any endpoints older than 30 days each day at 3:00 a.m.

You can continue to follow along in your own ISE UI or refer to Figure 13-11 to examine the settings for the contractor guest account type.

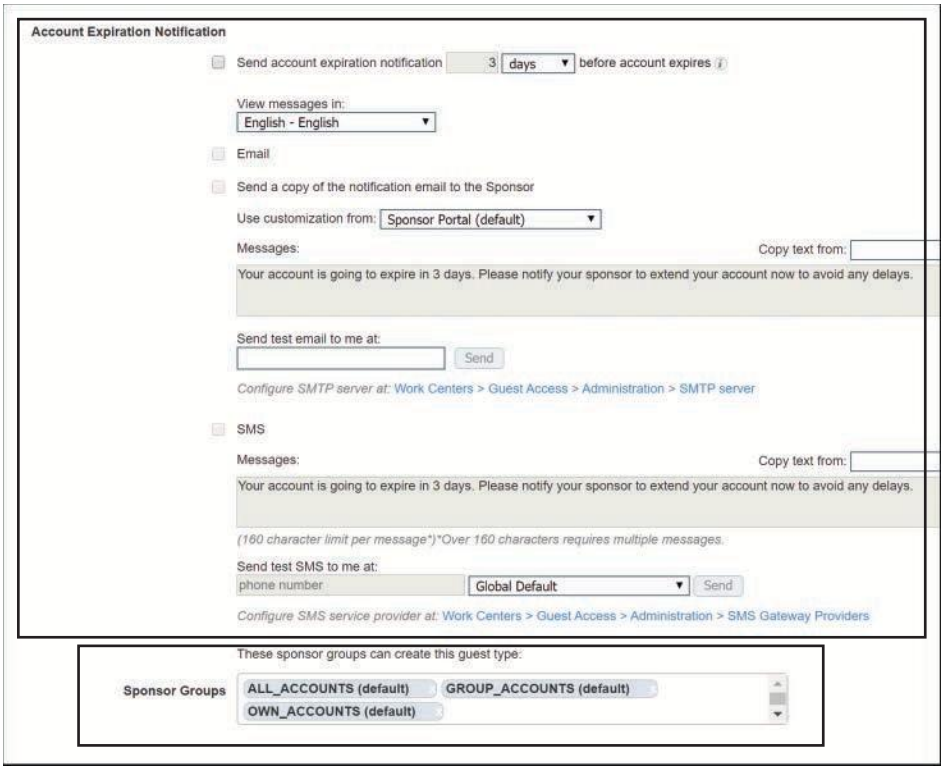


Figure 13-11 *Contractor Guest Type, Part 2*

The Account Expiration Notification section focuses on the settings related to notifying guest users when their accounts are expiring or have expired. You can set the number of days' notice to give guests about expiration, choose whether the notification should be emailed or sent via SMS or both, and set whether to also send a notification to the sponsor. The messages are fully customizable, and you can copy the text from any of the other guest types to avoid duplicating work effort.

Key Topic

The Sponsor Groups section shows which sponsor types can be authorized to create accounts for this guest type. Guest types are assigned to sponsor types. Organizations commonly require a manager or higher-level employee to provision or approve a guest account that is allowed access for a long time.

Daily

A daily guest account expires after only a few days. This type of guest account can typically be created by any employee of the company. You can examine this guest type by navigating to Work Centers > Guest Access > Portals & Components > Guest Types and selecting Daily (default).

Keep in mind that these guest types are the preconfigured default types for the most common use cases, based on customer research. The only differences in the account types are the preconfigured default settings. If you wanted to and had the proper permissions, you could set the daily account type to last 365 days.

The key difference between the preconfigured contractor guest type and the preconfigured daily guest type is the maximum access time. Daily accounts default to a one-day lifetime, with five days as the maximum setting. Again, you can change this if you choose to.

Figure 13-12 highlights the Maximum Access Time and Login Options sections for the daily guest type.

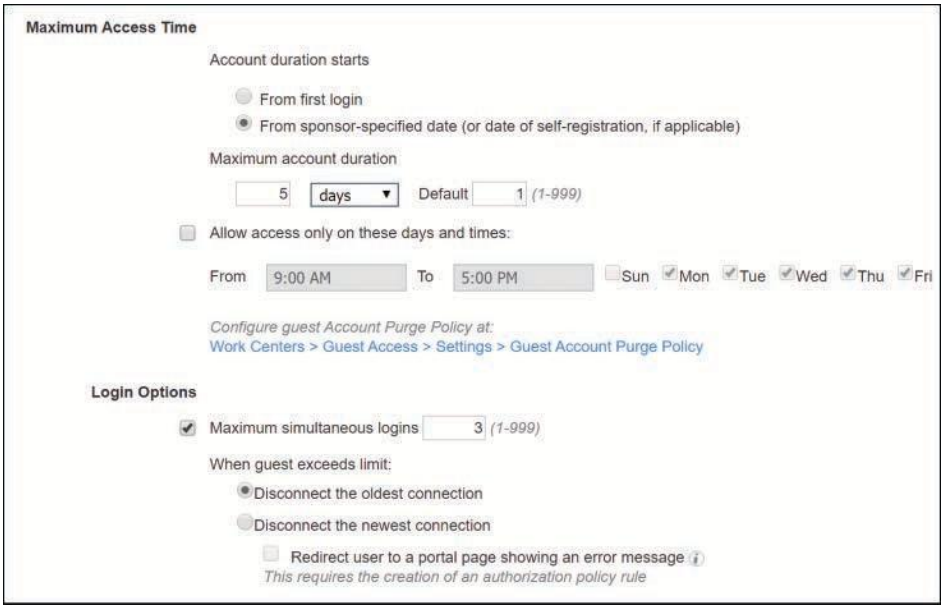


Figure 13-12 *Maximum Access Time and Login Options for Daily Guest Types*

Weekly

The preconfigured weekly guest type expires within a maximum of a few weeks. You can examine this guest type by navigating to Work Centers > Guest Access > Portals & Components > Guest Types and then choosing Weekly (default).

The only difference between the preconfigured daily guest type and the preconfigured weekly guest type is the maximum access time. Daily accounts default to a 1-day lifetime, with 5-days as the maximum setting, while weekly accounts default to a 5-day lifetime (1 business week), with 14-days as the maximum setting for a sponsor to grant access.

Figure 13-13 shows the Maximum Access Time section for the weekly guest type.

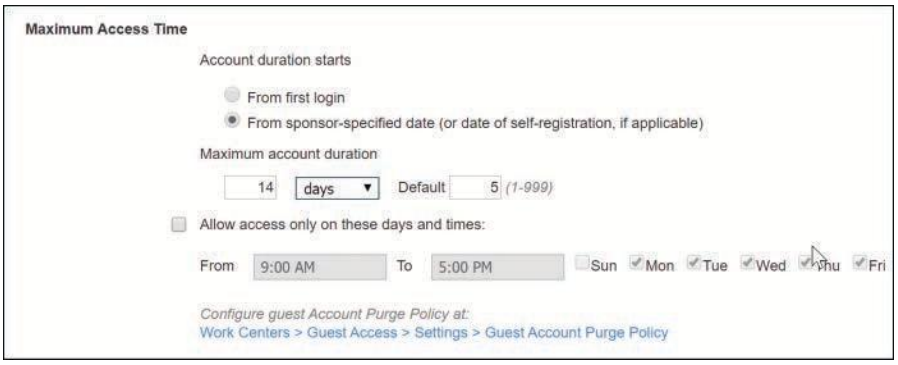


Figure 13-13 *Maximum Access Time for Weekly Guest Types*

Social

Social login guest types are special. These logins do not require a sponsor to preconfigure or approve account creation. These accounts are ultimately mere copies of a social network identity, such as a Facebook identity.

You might want to enable use of social login to ensure that guests are real persons and to have some trackability and accountability for the actions the guests perform while attached to the network. Another use case for social login is for marketing and lead tracking—to be able to run reports on the types of users who have been connecting to the guest network and maybe even to track them as sales leads.

Within ISE, the social guest type is a copy of the daily guest type. It has a five-day maximum lifetime, with a default of one day. The difference between a daily guest and a social guest is the lack of assigned sponsor groups for social login; a social login does not require sponsorship from an employee.

Figure 13-14 shows the Sponsor Groups section for the social guest type. Notice that it is empty.

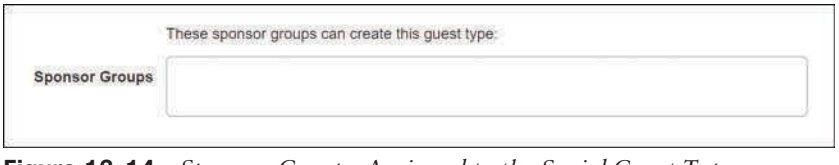


Figure 13-14 *Sponsor Groups Assigned to the Social Guest Type*

The other difference between social login guests and all other types of guest accounts is the requirement for a special identity store to be configured. A social login is a special identity source for use with guest portals only: A guest user can click "Log in with Facebook" and leverage his or her Facebook credentials to gain guest access.

As of ISE Version 2.7, Facebook is the only social media source ISE supports for guest logins.

Guest Portals and Authorization Policy Rules

There is a major commonality between all the different types of guests. Guest users typically connect to a network without strong authentication (802.1X) and are redirected to the guest portal. There are currently only two guest portal types in use.

Key Topic

In order to use a guest portal, there must be an authorization rule that redirects users to that portal. Navigate to Work Centers > Guest Access > Portals & Components > Guest Portals, as shown in Figure 13-15. Notice that the list of portals informs the administrator when a portal is not in use (see the arrow with the hotspot portal). When a portal is leveraged by an authorization policy, such as the self-registered guest portal, the number of authorization rules that use that portal is be displayed as a link (see the bottom two arrows in Figure 13-16). If you click that link, a pop-up message appears, with a list of where the portal is being used, as shown in Figure 13-16.

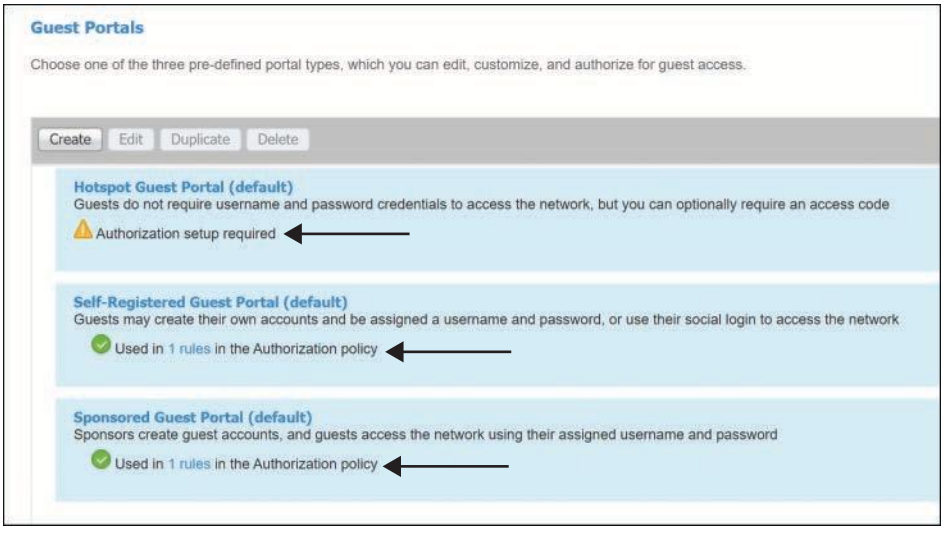


Figure 13-15 *Guest Portal List Informing Administrators of Portal Usage*

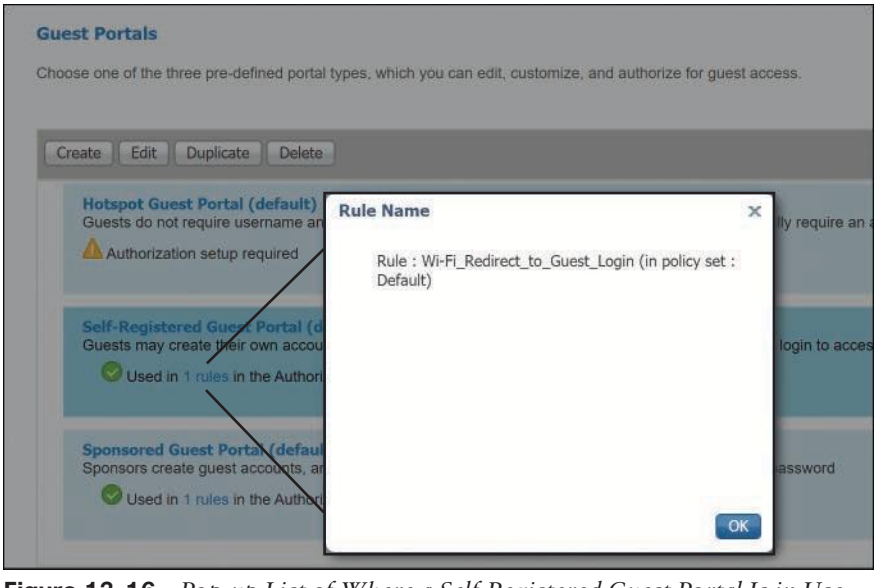


Figure 13-16 *Pop-up List of Where a Self-Registered Guest Portal Is in Use*

If you want to use all three portal types, how can you differentiate the access requests so you know which portal to direct each user to if you don't know the user's identity yet? One method that is commonly leveraged is to use a different SSID for each of the different guest types that your organization needs to deploy. Keep in mind, though, that not every portal type is appropriate for every installation. As a designer of an ISE deployment, you should determine the needs of the organization in terms of Internet access for non-employees and deploy the appropriate solution.

Another consideration—and a recommendation from the authors of this book—is to use the guest endpoint identity group(s) for authorization rules to help prevent guest users from being redirected to login portals every time they roam to a different WLC or come back in to the office again. To understand this recommendation, let's look at the common Wi-Fi experience of being a guest in a hotel:

Day 1: Check in to the front desk.

- The hotel employee checks your ID to ensure that you are the individual who signed up for the room and issues physical keys that permit you to enter your assigned room.
- Along with the keys, the employee also provides you with instructions for accessing the Wi-Fi.
  - You might be able to log in with your customer loyalty account.
  - You might log in with a combination of your last name and room number.
  - You might need to use an access code that is valid for a specific amount of time.
- You retire to your room for the night.
  - You connect your laptop, phone, or tablet (or maybe all three) to the Wi-Fi.
  - You work, stream a movie, and Facetime with loved ones at home.

Day 2: You have to rejoin the Wi-Fi and provide the same credentials all over again.

- You probably have to log in with each device.
- The network may have forgotten your device in the middle of the night, and whatever movies you were downloading to the iPad to watch on your next flight stall out and do not finish downloading because your credentials expired.

Day 3: You have the same problems.

Why can't the hotel remember your device for longer than a day? This is why the authors of this book, who feel that pain with you, recommend the use of guest endpoint identity group(s) in authorization rules and ensuring that those identity groups do not get purged too often. Following that guidance will help prevent your guest users from being redirected to login portals too frequently, and the world will be a better place.

NOTE While it is true that different hotel chains—and even individual hotels—have different experiences, this example use case is still widely prevalent. If you work in or own a hotel with a better experience than described in this section, our example is not meant to offend.

Configuring Guest Portals and Authorization Rules

When configuring a guest portal, there are a few workstreams to go through. You can configure the portal itself, defining the flow that guest users should experience. You can also customize the look and feel of that portal. However, you must use the portal in an authorization profile and leverage that authorization profile in an authorization group.

The following sections show how to configure separate SSIDs for the different portal types. Refer to Chapter 11, "Implement Wired and Wireless Authentication," if you need a refresher on how to configure WLANs on WLCs.

Configuring a Hotspot Guest Portal

To configure the hotspot guest portal, navigate to Work Centers > Guest Access > Portals & Components > Guest Portals and select Hotspot Guest Portal (default). You should see two main sections for configuration, Portal Behavior and Flow Settings and Portal Page Customization, as shown in Figure 13-17, which is a composite image.

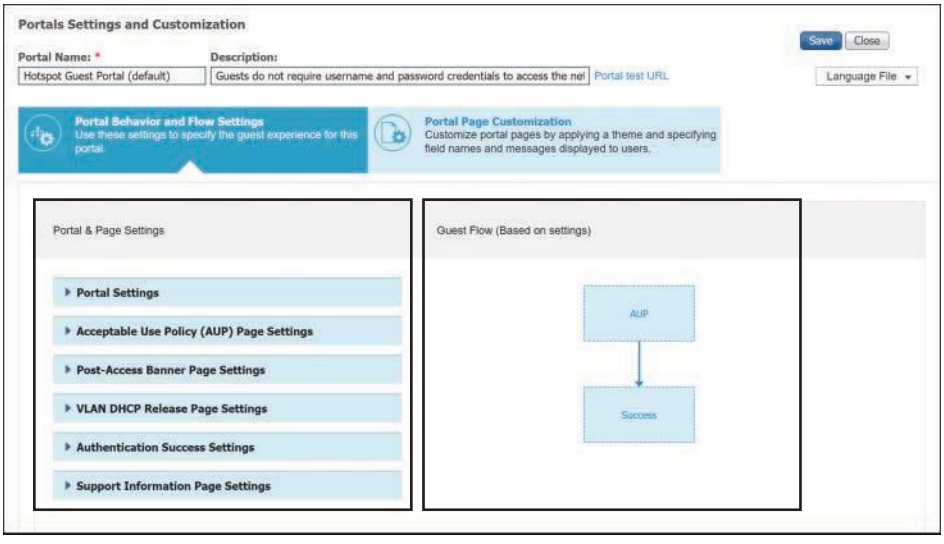


Figure 13-17 *Hotspot Portal Configuration*

In Figure 13-17, the Portals & Page Settings section includes configuration settings for portal behavior. The Guest Flow section in Figure 13-17 shows a dynamic flowchart that is based on the settings in the Portals & Page Settings section; this flowchart helps you understand the implications of the portal settings chosen. (Let's all thank Chip Current for this beautiful user experience.)

Portal Behavior and Flow Settings

Remember that portals are web pages, and an ISE PSN hosts these web pages as a web server. Each PSN hosts an identical set of portals, so the configuration of a portal must be applied for all of the PSNs simultaneously. It would do no good to specify an IP address, since each PSN has its own IP address. Each PSN may also have a different configuration of interfaces. Some PSNs might use interface binding to ensure high availability of the PSN services, and others might not have the binding configured.

Portal Settings

In the Portals & Page Settings section, expand the Portal Settings section, as shown in Figure 13-18. In this section you can see where the portal is assigned a port and where interfaces are selected.

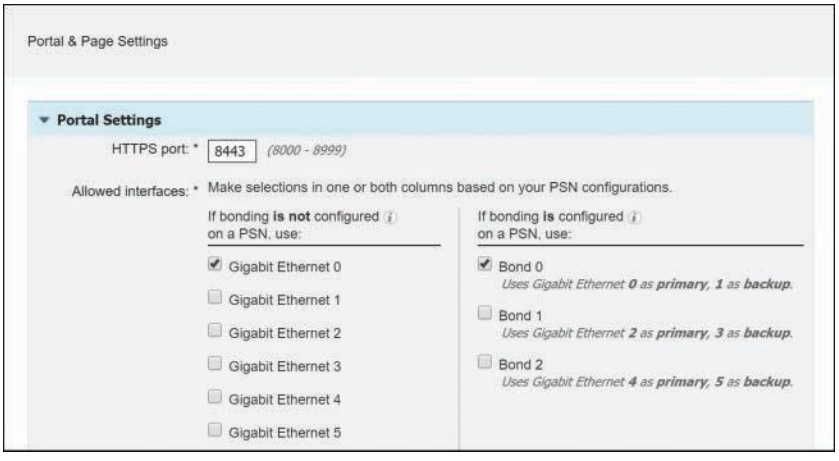


Figure 13-18 *Hotspot Portal Configuration > Portal Settings, Part 1*

As you examine the screen shown in Figure 13-18, take note of the interface selection. By default, all portals are only hosted on the Gig 0 interface, but you can select any combination of the six interfaces. Does this mean all PSNs must have six interfaces for this portal to work? Absolutely not. If a PSN does not have an interface that matches this configuration, it simply ignores the setting and does not try to host the portal on a nonexistent interface.

Interface bonding was added to ISE Version 2.1 to aid with high availability. Customers may have multiple switches in the data center and might want to plug an interface into each switch for redundancy. Keep in mind that this is not an active/active scenario where both interfaces are answering for the same IP address; it is an active/passive situation only. If a switch is lost, the backup interface in the bonded pair takes over the IP address and MAC address of the primary interface.

Interface bonding is configured only at the CLI of each node, one interface at a time, as shown in Example 13-1. The GUI shown in Figure 13-18 is not configuring the bond but instead describing which bond interface(s) the portal should be hosted on if a bond exists on the PSN. If a bond does not exist on a PSN, the PSN ignores the configuration to host the port on the nonexistent bond interface.

Example 13-1 shows the CLI commands to configure a bonded interface pair.

Example 13-1 *Configuring Interface Bonding*

```
atw-ise243/admin(config)# int GigabitEthernet 2
atw-ise243/admin(config-GigabitEthernet)# backup interface GigabitEthernet 3
```

When you configure the NIC bonding feature, Cisco ISE pairs fixed physical NICs to form bonded NICs. Table 13-2 shows which NICs can be bonded together to form a bonded interface. Figure 13-18 also shows the same information in a prettier form.

Table 13-2 NIC Bonding

Cisco ISE Physical NIC Name	Role in Bonded NIC	Bonded NIC Name
Gigabit Ethernet 0	Primary	Bond 0
Gigabit Ethernet 1	Backup	
Gigabit Ethernet 2	Primary	Bond 1
Gigabit Ethernet 3	Backup	
Gigabit Ethernet 4	Primary	Bond 2
Gigabit Ethernet 5	Backup	

Much the way the guest user interface greatly simplifies the myriad possible interface combinations, the user interface greatly simplifies the myriad certificate possibilities.

The ISE PSN hosts all web pages securely, which means a certificate is required to provide the secure HTTPS connection. Each PSN can have its own certificate. Actually, each PSN can have many different certificates.

In order to simplify this greatly, the concept of a portal group tag was added in ISE Version 1.3, along with a new guest portal structure. Every guest portal must be assigned to a portal certificate group tag, as shown in Figure 13-19.

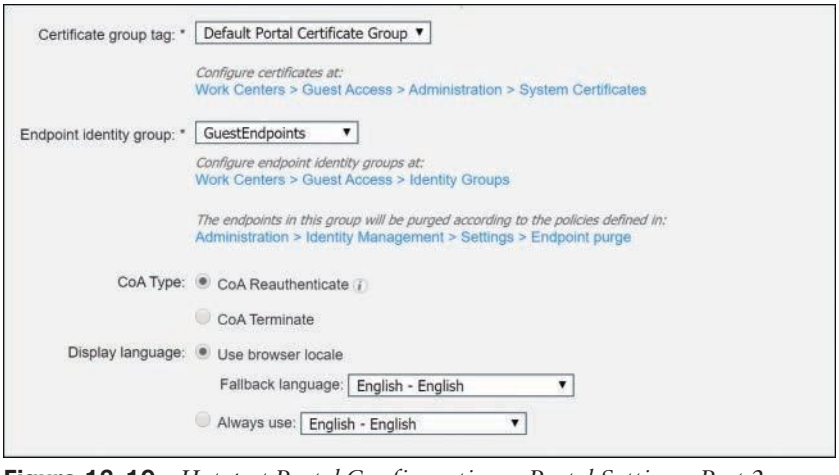


Figure 13-19 *Hotspot Portal Configuration > Portal Settings, Part 2*

The certificate portal group tags are not created within the portal. They are created when you add a new certificate or configure certificate usage, in the System Certificates section of the ISE UI. Any certificate that has the usage type portal assigned must be part of a portal group tag, as shown in Figure 13-20.



Figure 13-20 *Certificate Usage Settings*

Key Topic

The purpose of a portal group tag is to ensure that the correct certificate for each portal is presented when a guest connects to the portal, without requiring someone to manually select each certificate on each PSN individually. Since each portal can be identified and protected by only a single certificate per listener, only one certificate per PSN may be assigned to that portal group tag. A listener could be defined as a single interface and port combination, and a single certificate might be used per listener. If you think the process through, you realize that it would be impossible to know which portal the traffic is destined for until after decrypting the HTTPS headers. This means you cannot house more than one portal on a single listener with different certificates.

Look again at Figure 13-19, and you see these portal settings:

- **Endpoint Identity Group:** This setting is very important, especially for a hotspot portal. A hotspot portal differs greatly from a sponsored guest portal in the way that a user is redirected to the page. With a hotspot portal, a user can just click a button and perhaps agree to an acceptable use policy (AUP) to get uninterrupted access to the Internet; behind the scenes, the user's MAC address is added to this endpoint identity group, and then a Change of Authorization (CoA) is sent (where the endpoint will be granted the access, based on that endpoint identity group).
- **CoA Type:** It is important to choose the correct CoA type. The default, CoA Reauthenticate, is normally the best choice. The only reason to change it to CoA Terminate would be if you plan to change the VLAN on the endpoint. However, in Chapter 12, we cautioned against performing VLAN changes to clients that are not using supplicants.
- **Display Language:** This is a very important setting for many customers who have locations in multiple countries, and you might be surprised how challenging it can be to support multiple languages. The default setting is to use the locale from the user's browser with a fallback to English when the locale does not meet one that ISE is configured for or to force all portals to always use English.

Acceptable Use Policy (AUP) Page Settings

An AUP, which typically comes from an organization's legal department, is a statement that informs guest users about their obligations related to using the Internet connection.

Figure 13-21 shows the guest flow of the portal. If you uncheck Include an AUP page, the flow diagram dynamically updates, as shown in Figure 13-22.

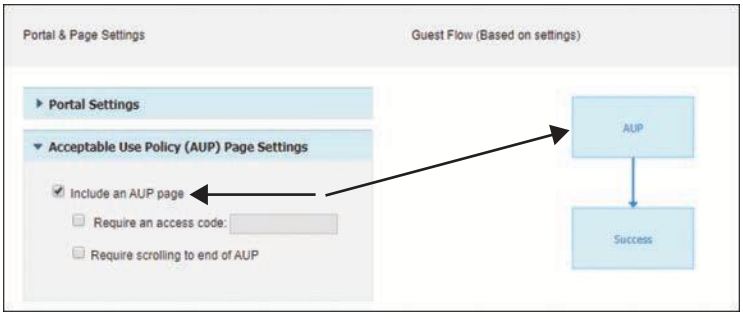


Figure 13-21 *AUP Included in Flow Diagram*

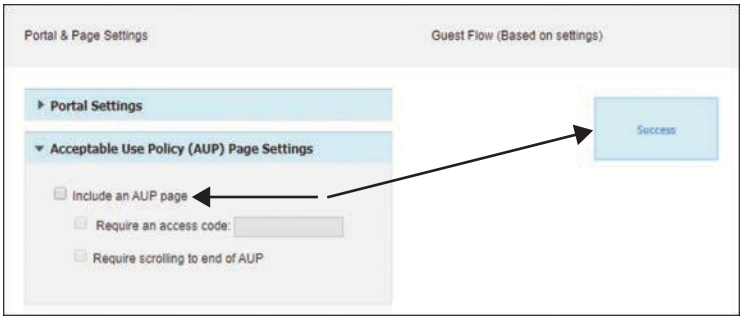


Figure 13-22 *Flow Diagram Dynamically Updates*

In this section, you should be sure that Include an AUP page is selected, as this page is far more than just about AUPs. When the guest accepts the AUP, the acceptance causes the endpoint's MAC address to be added to the guest endpoint's identity group.

In addition to including an AUP on the page, you can add an access code, which is a common passphrase for guest users to use. This option is very popular because it adds another level of control, ensuring, for example, that the guest is on the premises and has seen the posted signs with Wi-Fi instructions—and isn't a hacker in a black hoodie sitting in the parking lot.

As shown in Figure 13-23, the final setting allows you to force the guest user to scroll all the way to the end of the AUP and not just click Accept.

Post-Access Banner Page Settings

There is only a single setting with a single checkbox in the Post-Access Banner Page Settings section: Include a Post-Access Banner page. This banner is basically a "Welcome to our guest network" type of message that you can configure in the Portal Page Customization section. Figure 13-24 shows this setting and the change to the guest flow when it is selected.

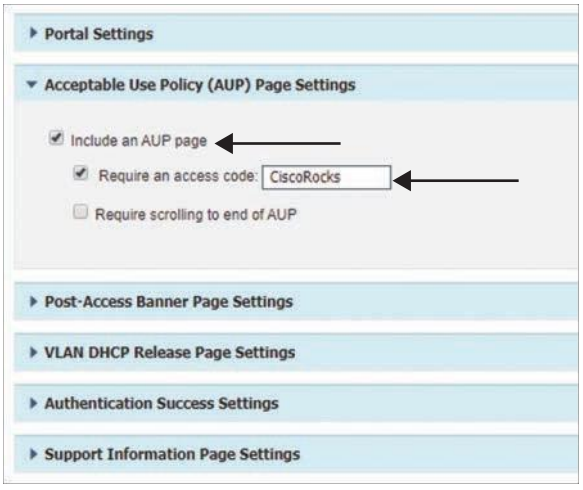


Figure 13-23 *AUP Page Settings*

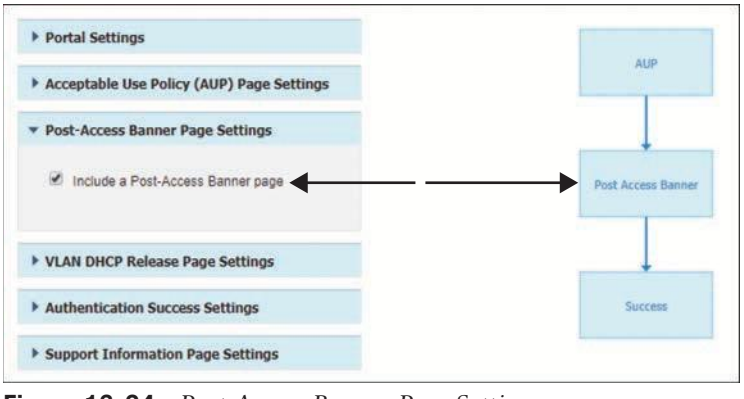


Figure 13-24 *Post-Access Banner Page Settings*

VLAN DHCP Release Page Settings

At first glance, this section might look like it flies in the face of the advice you received in Chapter 12 and earlier in this chapter to change VLANs only when an endpoint has a supplicant and never when the endpoint is using Web Authentication.

In Figure 13-25, it appears as though this is the answer to that very basic technical limitation, where an endpoint does not recognize that the VLAN has changed and therefore does not get a new IP address. However, this setting leverages an ActiveX or Java applet to accomplish the releasing of the current DHCP address and retrieval of the new address. It has been years since ActiveX applets have been permitted to run in Internet Explorer, and it has been almost as long since Java applets were disabled by default on most systems. This setting is, therefore, completely unusable in the vast majority of organizations.

If you absolutely must change VLANs on guest users, the best approach is to use a short-lived DHCP lease time, such as 5 minutes, so the endpoint will try to renew its IP address every 2.5 minutes.

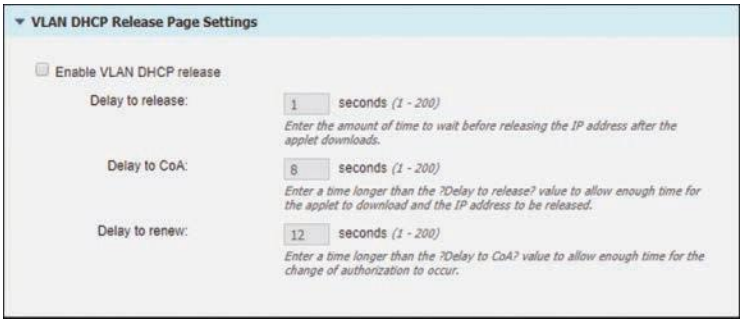


Figure 13-25 *The Obsolete Setting*

Authentication Success Settings

What happens after a guest user has authenticated successfully? Do you want to return the user to the original URL that was redirected to the guest portal? Should you display a simple success page? Maybe you want to redirect the user to your company website. You are likely to want to return the user to the original URL, but not all NADs support that ability.

Figure 13-26 shows the available settings in the Authentication Success Settings section and the change to the guest flow that occurs when you set a URL.

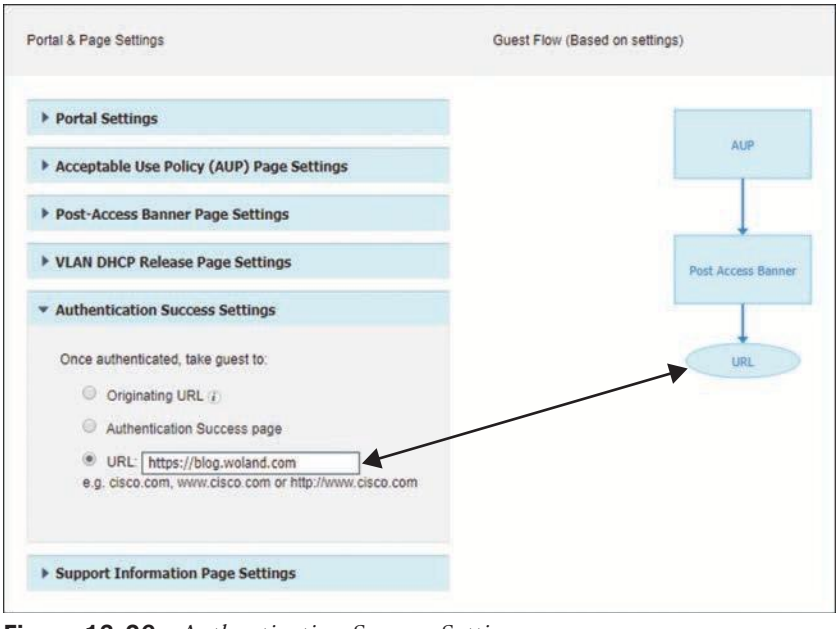


Figure 13-26 *Authentication Success Settings*

Support Information Page Settings

The final configuration section under Portal & Page Settings is Support Information Page Settings. A support information page enables a guest user to click if something is going wrong and provides information for the help desk.

Figure 13-27 shows the settings available for the support page. As you can see, it can include IP and MAC addresses, browser type, and user agent information, as well as which ISE PSN is in use and the failure code.

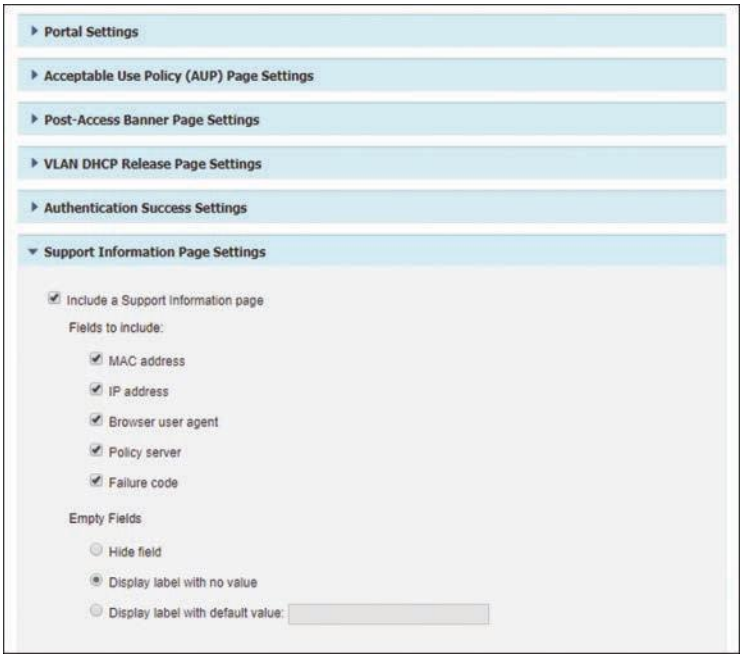


Figure 13-27 *Support Information Page Settings*

Portal Page Customization

At this point, you have been through all the configuration sections for a hotspot portal, but you have yet to see what such a portal actually looks like, so navigate to Work Centers > Guest Access > Portals & Components > Guest Portals > Hotspot Guest Portal (default) and marvel at this amazing tool. This tool offers tremendous power and simply cannot be rendered in a book in a single screenshot, so Figure 13-28 shows it cut off about halfway down.

Let's take a closer look at what is available in terms of portal customization. To do so, you can examine Figure 13-28 or, even better, examine your own ISE screen as you read along. In the lower-right part of the screen, notice the portal preview. As you make changes to the fields in this view, the preview rendering updates. You can also click the Desktop Preview link below it to launch a pop-up window with a live rendering of the portal page as it would look on a larger screen.

Figure 13-29 shows the mobile preview, and Figure 13-30 shows the desktop preview. Keep in mind that as you make changes, the preview updates.

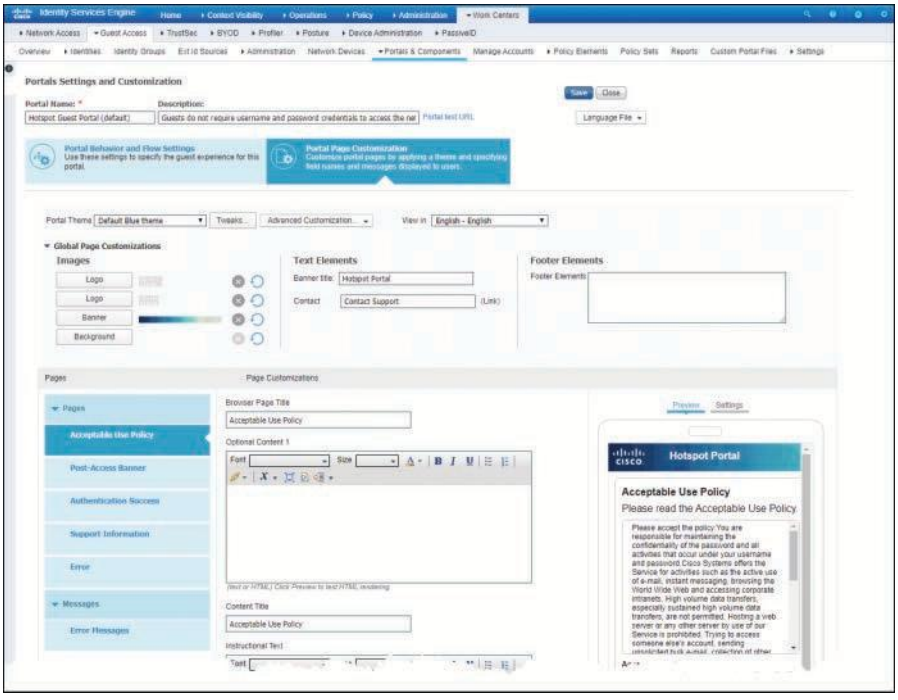


Figure 13-28 *Cropped Portal Page Customization Screen*

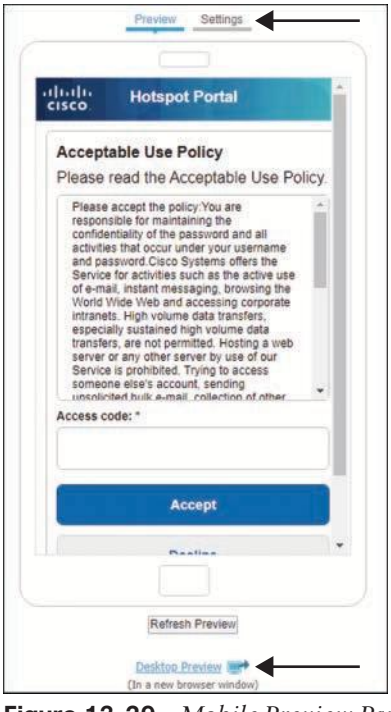


Figure 13-29 *Mobile Preview Pane*

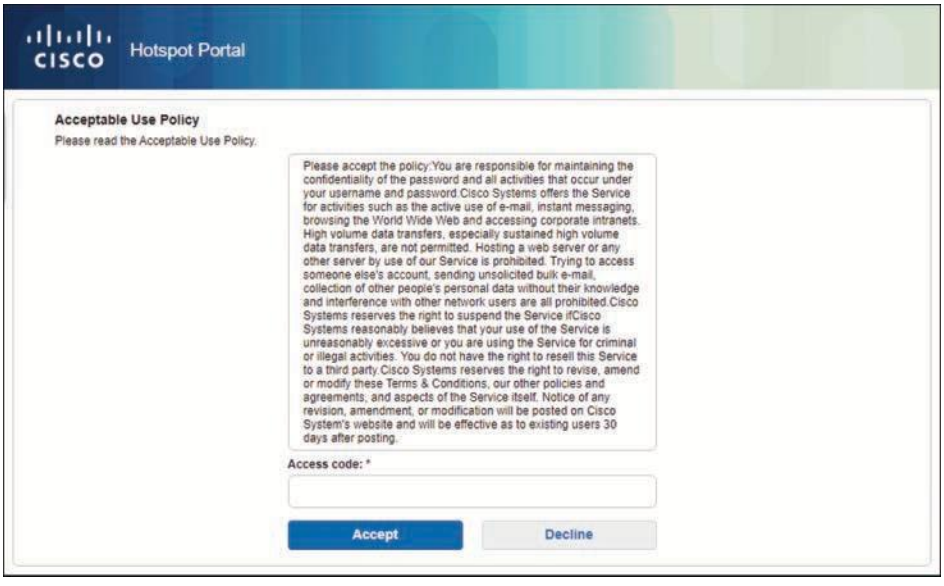


Figure 13-30 *Desktop Preview Page*

If you click the Settings tab at the top of the mobile preview that is called out in Figure 13-29, you can view what is available, as shown in Figure 13-31.

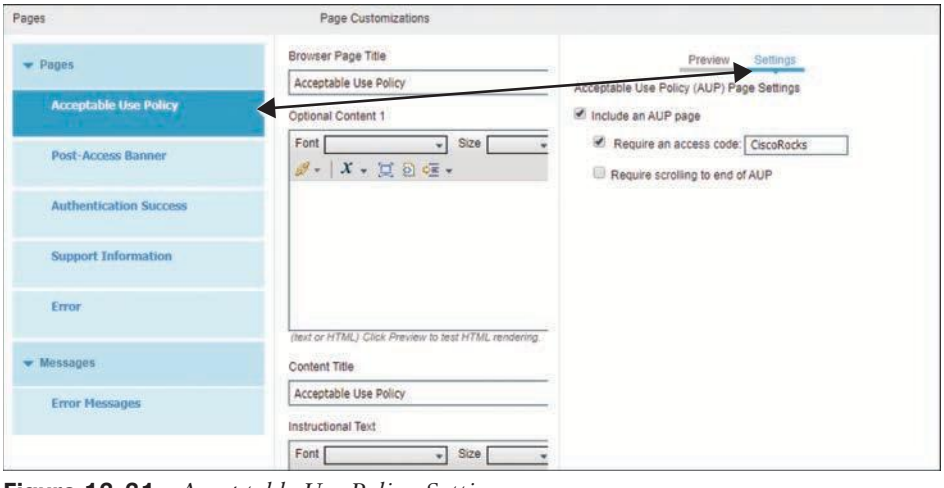


Figure 13-31 *Acceptable Use Policy Settings*

As shown in Figure 13-31, the Settings tab displays the appropriate settings for the page that is currently being used. If you were editing the Post-Access Banner page, for example, the Settings tab would reflect those settings accordingly. This provides for a clean user experience that prevents the administrator from having to click back to the Portal Behavior and Flow Settings tab.

Now let's look at the rest of the settings, starting at the top of the screen and working our way down.

Some default portal themes are prebuilt in ISE. These themes, which affect the entire portal, are available in a dropdown at the top of the customization screen. You can export and import these themes for advanced customizations (see Figure 13-32), and you can make basic tweaks for simple changes that can be implemented by clicking the Tweaks button, as shown in Figure 13-33.

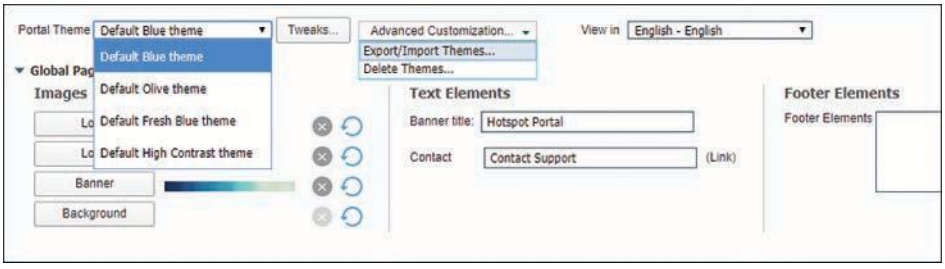


Figure 13-32 *Portal Themes*

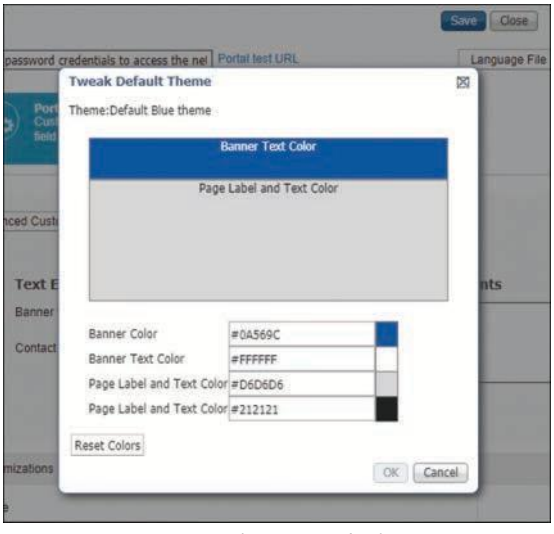


Figure 13-33 *Tweaking a Default Theme*

The Global Page Customizations section allows you to change the images, such as logos for mobile and desktop pages, banner images, and page background. When you change any of these images, the changes affect all pages in the portal. There are text elements for the banner title, a contact link, and the footer to display on the pages. Figure 13-34 shows some customizations made in this section, and Figure 13-35 illustrates those changes in the mobile preview pane.

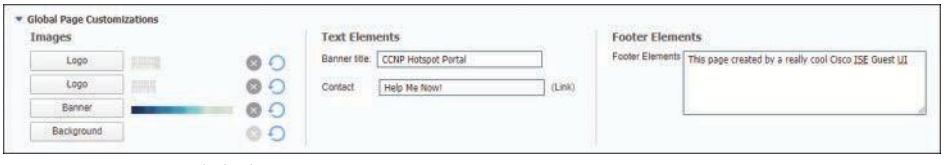


Figure 13-34 *Global Page Customizations*

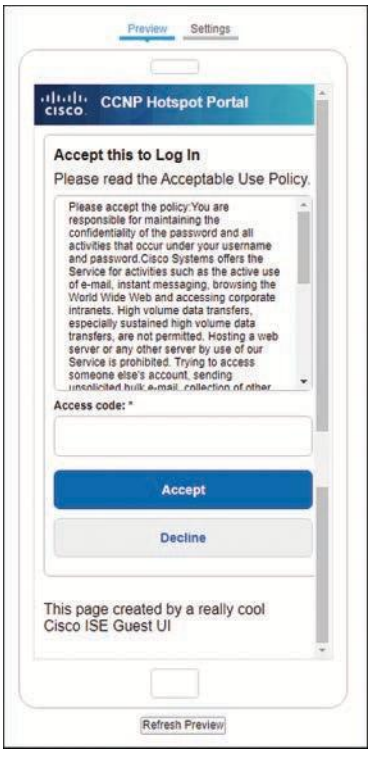


Figure 13-35 *Showing the Page Customizations in the Mobile Preview Pane*

The remainder of the portal page customization options are fairly self-explanatory, and the best way to get to know them is to try them out. All the text and images for a portal are configurable.

Authorization Rule Configuration

This chapter shows how to configure three different types of guest portals. A small or mid-sized business would not typically need to configure all three, but a larger enterprise might.

To help you in your educational journey toward passing the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam, this section shows how to configure a separate SSID for each of the portal types. This will enable you to use the called-station-id RADIUS attribute as a condition in policies to steer the traffic to the correct portal.

For the hotspot portal, you can use a WLAN that was pre-created on the WLC, with SSID ISE-Hotspot and assigned to the Guest VLAN. (See Chapter 11 for a refresher on creating a WLAN, if needed.)

You need to configure an authorization profile to redirect the user to the hotspot portal you configured in this chapter. Follow these steps in the ISE GUI (see Figure 13-36):

Step 1. Navigate to **Work Centers > Guest Access > Policy Elements > Results > Authorization Profiles**.

Step 2. Duplicate the CWA profile that you created in Chapter 12.

Step 3. Name the new profile **Hotspot CWA**.

Step 4. Select **Hot Spot** from the Web Redirection dropdown.

Step 5. Leave the ACL field set to **ACL-WEBAUTH-REDIRECT**.

Step 6. Select **Hotspot Guest Portal (default)** from the Value dropdown.

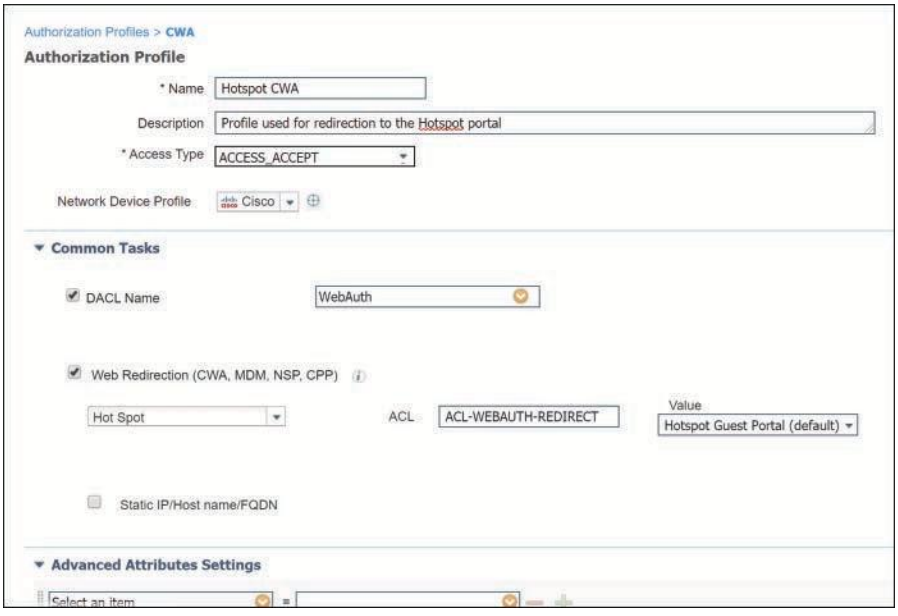


Figure 13-36 *Hotspot CWA Authorization Profile*

You could create a new policy set just for the Hotspot SSID if you wanted. However, for simplicity, here you can create an authorization rule above the CWA rules created in Chapter 12 so that traffic from the ISE-Hotspot SSID receives the Hotspot CWA authorization profile. Follow these steps in the ISE GUI:

Step 1. Navigate to **Work Centers > Guest Access > Policy Elements > Policy Sets**.

Step 2. Select the policy set that you will be working with.

Step 3. Insert a new authorization rule above the Employee_CWA rule and name it **Hotspot CWA**.

Step 4. Ensure that the conditions of the new rule include **Wireless_MAB**.

Step 5. Add the condition **Radius Called-station-id CONTAINS ISE-Hotspot**.

Step 6. In the Authorization Profile column, select the **Hotspot CWA** profile.

Step 7. Duplicate the new rule above the Hotspot CWA rule.

Step 8. Name the new rule **Guest Endpoints - Hotspot**.

Step 9. Include an additional condition for **Identity Group Name EQUALS Endpoint Identity Groups:Guest Endpoints**.

Step 10. Select **Internet-Only** for the authorization profile.

Step 11. Save the policy set.

Figure 13-37 shows the new authorization rules in the policy set.

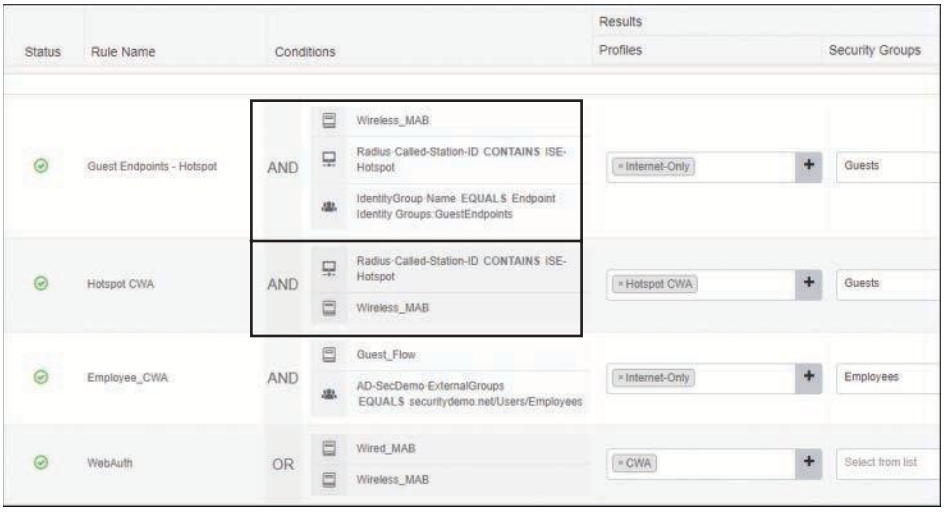


Figure 13-37 *Authorization Rules for Hotspot and Guest Endpoints*

Key Topic

Much as for the rules you created in Chapter 12, the order of these rules is very important. Once a user clicks to accept the AUP, the MAC address of the guest's endpoint is added to the Guest Endpoints identity group. The top rule permits access to the Internet for members of that group coming from the ISE-Hotspot wireless network. If a user is coming from the ISE-Hotspot wireless network and the endpoint is not already in the Guest Endpoints group, the user gets redirected to the hotspot portal.

Figure 13-38 is a flowchart that explains this flow graphically, and Figure 13-39 shows the live log, where the number 1 shows the endpoint being sent through the hotspot first, and the number 2 indicates the result of the user typing in the CiscoRocks access code and accepting the AUP.

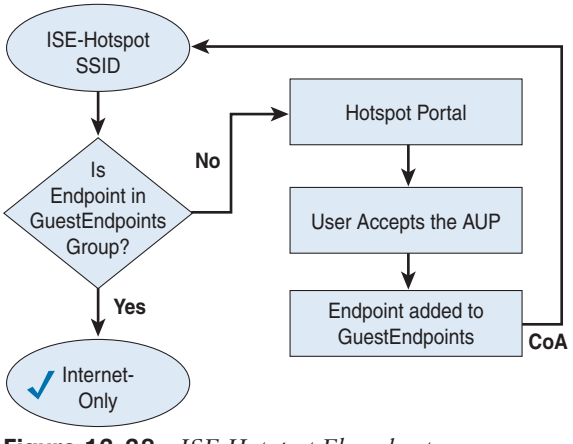


Figure 13-38 *ISE-Hotspot Flowchart*

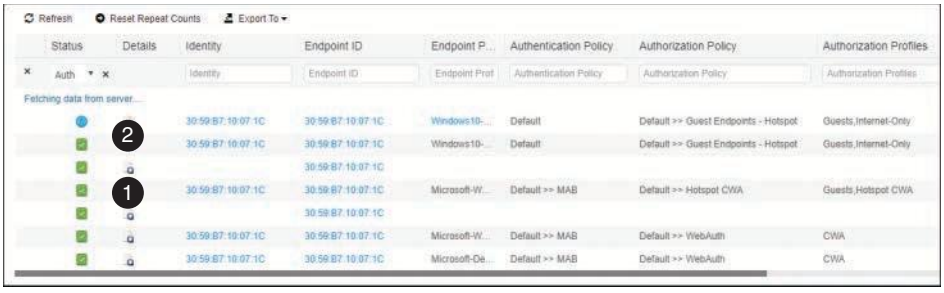


Figure 13-39 *Live Log for Hotspot Flows*

Configuring a Self-Registered Guest Portal

Unlike the preconfigured guest types, which are all created equally, the portal types are not just preconfigured options. They really are different and provide different options.

Figure 13-40 is a composite image that shows the self-registered guest portal (the section labeled 1) and the hotspot portal (the section labeled 2) side-by-side for comparison. Figure 13-41 shows the dynamic flow diagram that is built in the portal.

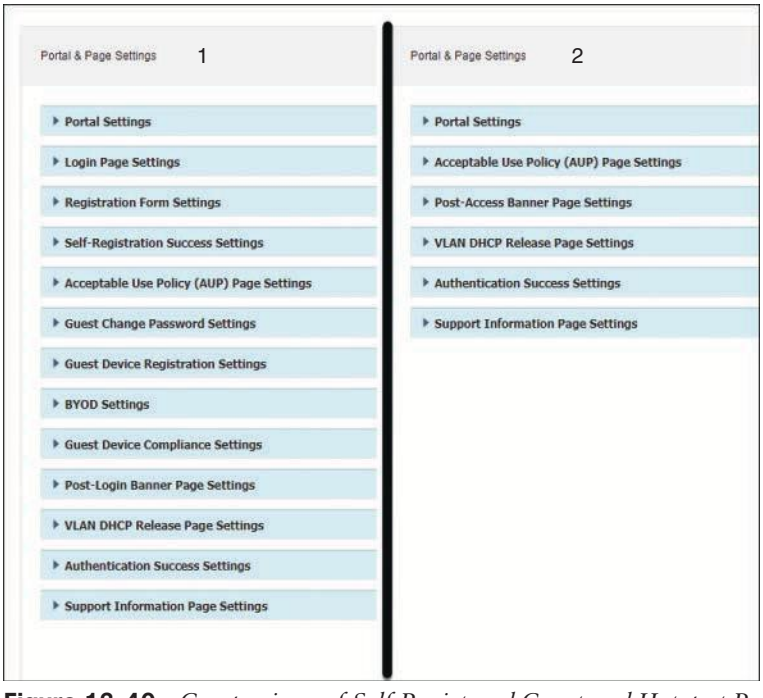


Figure 13-40 *Comparison of Self-Registered Guest and Hotspot Portals*

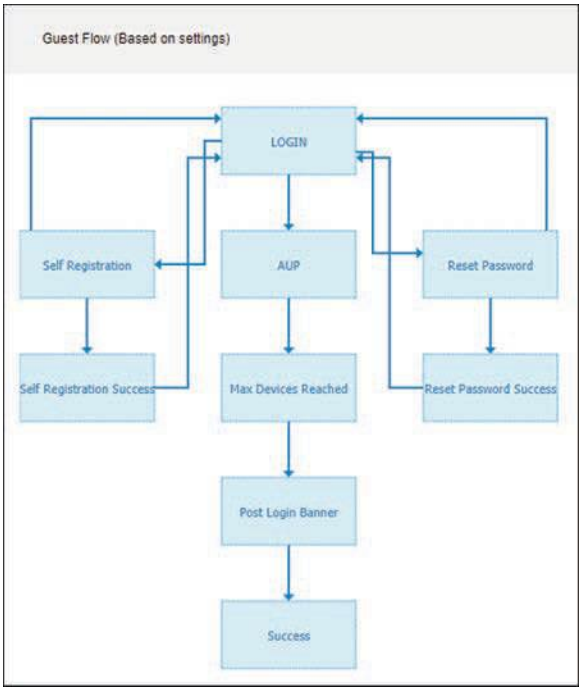


Figure 13-41 *Default Flow for the Self-Registered Guest Portal*

Now that you have all the background you need, it's time to move on to configuration.

Portal Settings

Let's look at the different settings in the Portal Settings section. Follow these steps in the ISE UI:

Step 1. Navigate to **Work Centers > Guest Access > Portals & Components > Guest Portals.**

Step 2. Select **Edit Self-Registered Guest Portal.**

Step 3. Expand **Portal Settings.**

Notice that the self-registered guest portal is hosted on the same interfaces and same ports as the hotspot portal. This means both portals are on the same listener. How is that possible?

Both portals are using the default portal certificate group, which means the PSN can decrypt the session and read the headers to see which portal was being used in the redirection; it can therefore provide the correct portal.

Figure 13-42 shows the portal settings for the self-registered guest portal.

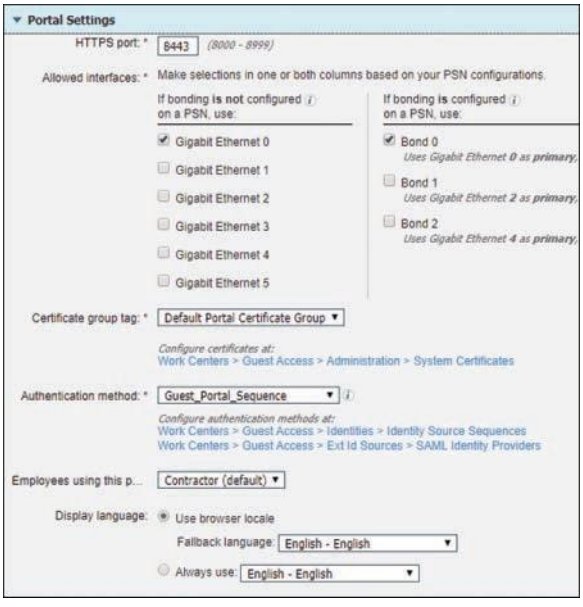


Figure 13-42 *Self-Registered Guest Portal Settings*

There are two notable differences in the Portal Settings section for the self-registered and hotspot portals:

- **Authentication method:** In the hotspot portal, there is no authentication. That is ultimately the goal of a hotspot portal: Allow the user to accept the AUP and get connected. With self-registered guests, the visitor is issued credentials to be properly identified.
- **Employees Using This Portal as Guests Inherit Options From:** Keep in mind that the portal you are editing is a centralized web authentication portal. There are username and password fields, and the authorization of accounts does not need to be limited to visitors. If a user who is not a member of a guest user identity group authenticates through this portal, that user is treated as having the same settings as the contractor guest type. Naturally, this is configurable and can be changed to any guest type you want.

Login Page Settings

The Login Page Settings section does not exist for a hotspot portal because there is no login page with a hotspot. Figure 13-43 shows the Login Page Settings section for a self-registered guest portal.

The login page is one of the main pages for the self-registered guest type. It contains an AUP, and you can also add an access code to add one more level of control. You have a configurable option to slow down or prevent brute-force hacking so that after five failed login attempts, ISE limits logins to have a two-minute waiting period between attempts.



Figure 13-43 *Self-Registered Guest Portal Login Page Settings*

Allowing guests to create their own accounts is the main purpose of a self-registered guest portal, so naturally that setting is checked by default. Guests may also request a password reset in case they lose their credentials.

Just because guests can create their own accounts does not mean that your organization will want them to set their own passwords. Those passwords are automatically generated by ISE, based on the password policy you configure. Allowing a guest to change the password to something more memorable can provide for a better overall experience, but that might not work with the security practices in your organization.

Social login can be enabled on a self-registered guest portal even though it is not a social login guest portal. This is an option of convenience to offer the visitor a choice: *Either create your own account by providing your email address or provide us with your social media identity. Either one is acceptable, but you must pick one to gain access*. You must have a social login identity store configured to use the social login option.

Security Assertion Markup Language (SAML) is another option for this portal type, if you have the SAML identity provider (IdP) configured already. This option allows this login portal to participate in a single sign-on (SSO) ecosystem.

Social and SAML identity sources are configured later in this chapter.

Registration Form Settings

The registration form is arguably the most important page in a self-registered guest portal, and the Registration Form Settings section certainly has a lot of settings. Some settings may be duplicated in other pages, but where that is the case, the same setting is reflected across the pages.

As already mentioned, there are many form settings, and we need to show a number of figures to present them all. Figure 13-44 shows the guest type selection and the fields to collect from the user.

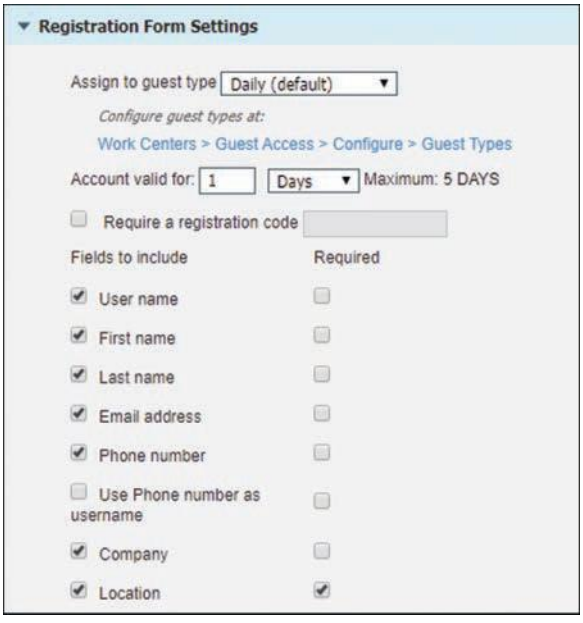


Figure 13-44 *Self-Registered Guest Portal Registration Form Settings, Part 1*

In this portion of the settings, you assign the guest type for users who register through this portal. The default type is daily, and these accounts are valid for one day with a maximum of five days. You can also specify an access code, and this setting can be accessed in multiple areas.

There are multiple fields to include in the registration form, and these fields are pulled from the guest type. If you have custom fields for a guest, they are included in the list. As an administrator, you can determine which of these fields are required and which are optional.

Figure 13-45 shows the ***guest location*** setting for the portal. You can use this setting to indicate which time zones are available for a guest to select when creating a guest account.

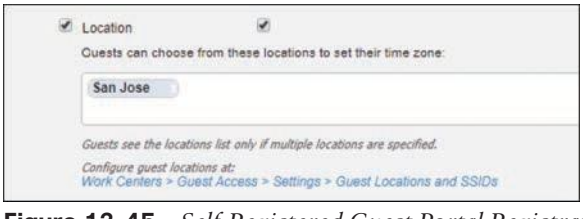


Figure 13-45 *Self-Registered Guest Portal Registration Form Settings, Part 2*

Key Topic

The time zone is a critical setting. Guest accounts have a configurable lifetime, and if the time zone is set incorrectly, a guest might not be able to access the network on the first try. The list of possible time zones is extensive, and to make it easier for guest administrators, ISE allows you to set friendly names that correspond to the company locations, under Work Centers > Guest Access > Settings > Guest Locations and SSIDs, as shown in Figure 13-46.

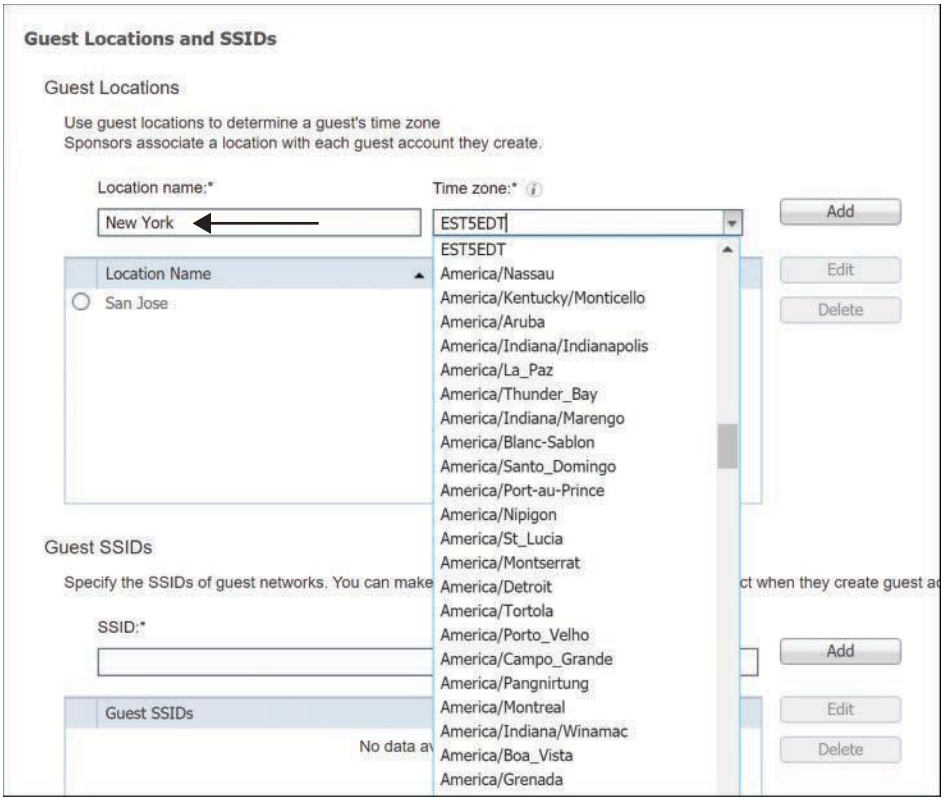


Figure 13-46 *Guest Locations (Time Zones)*

Figure 13-47 shows the selection of SMS service providers to use when sending guests credentials via text message. There is also a setting for the guest to provide the name of the person being visited and a reason for the visitation. Finally, you can require a guest to be approved by a sponsor.

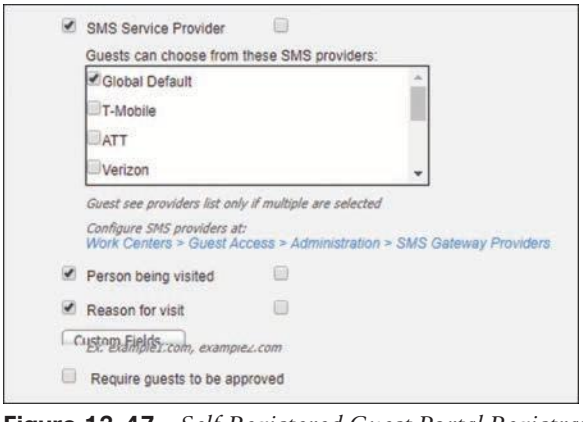


Figure 13-47 *Self-Registered Guest Portal Registration Form Settings, Part 3*

Figure 13-48 shows the rest of the settings for the registration form. The options here involve where to direct guests after they submit the registration form and how to send the credentials to the guest.

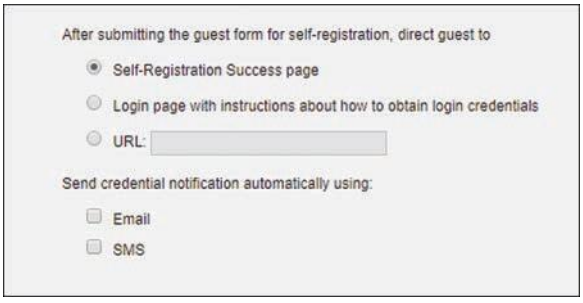


Figure 13-48 *Self-Registered Guest Portal Registration Form Settings, Part 4*

Self-Registration Success

The Self-Registration Success page is displayed to a guest who successfully registers. As shown in Figure 13-49, this page can display all the fields that the guest entered into the registration page. There is also a setting for how guests are permitted to send their credentials to themselves: print, SMS, or email. Finally, there is an option to allow guests to log in directly from the success message page and immediately begin a session.

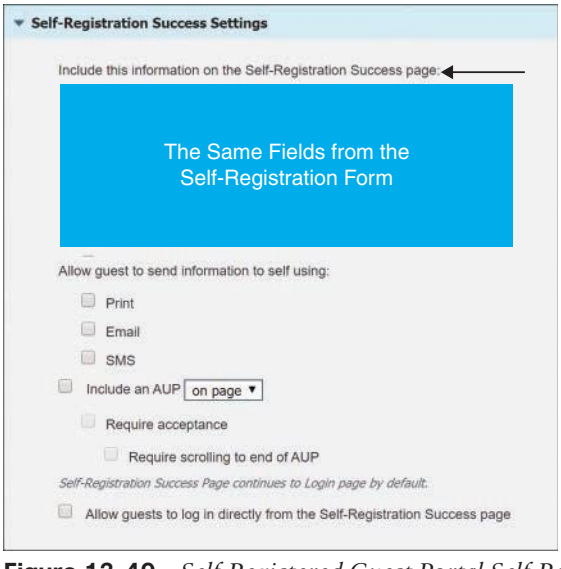


Figure 13-49 *Self-Registered Guest Portal Self-Registration Success Settings*

Guest Change Password Settings and Guest Device Registration Settings

As shown in Figure 13-50, you can require guests to change their passwords after first login, and you can automatically register their devices in order to keep their user experience smooth.

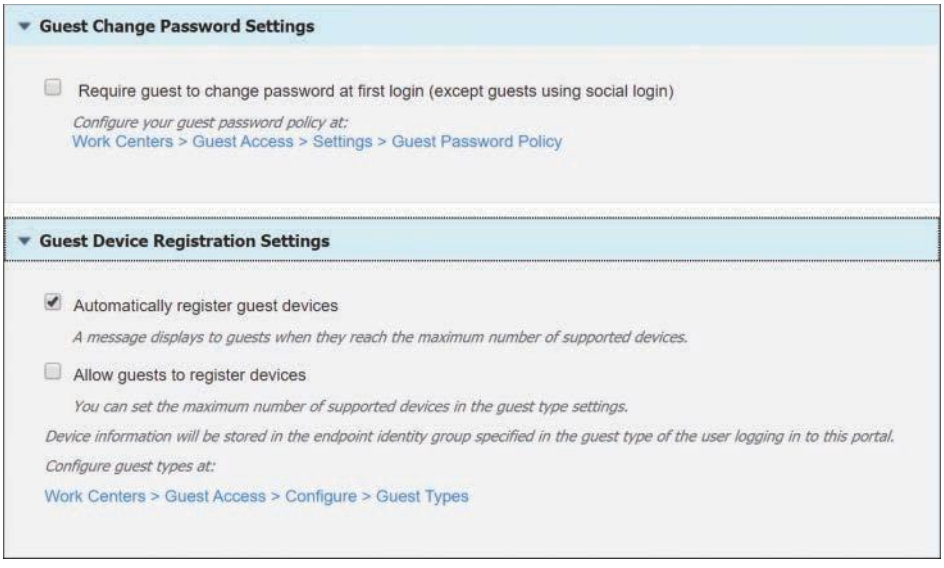


Figure 13-50 *Self-Registered Guest Portal: Change Password and Device Registration Settings*

BYOD Settings

BYOD settings are for employees, not guests. As shown in Figure 13-51, enabling these settings allows employees to register their own devices and obtain Internet access through the network. You will learn more about BYOD in Chapter 16, "Bring Your Own Device."

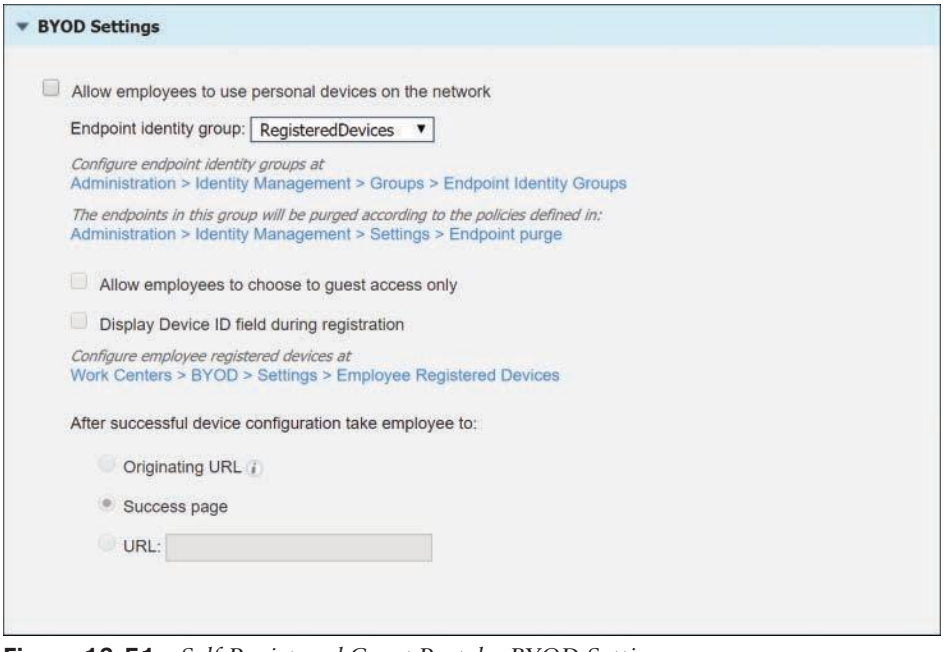


Figure 13-51 *Self-Registered Guest Portal – BYOD Settings*

Guest Device Compliance Settings

Guest Device Compliance Settings is the final page of settings not already covered for the hotspot portal. As you can see in Figure 13-52, the setting in this section can be selected to indicate that guest endpoints need to be posture assessed. You will learn about posture in Chapter 18, "Posture Assessment."

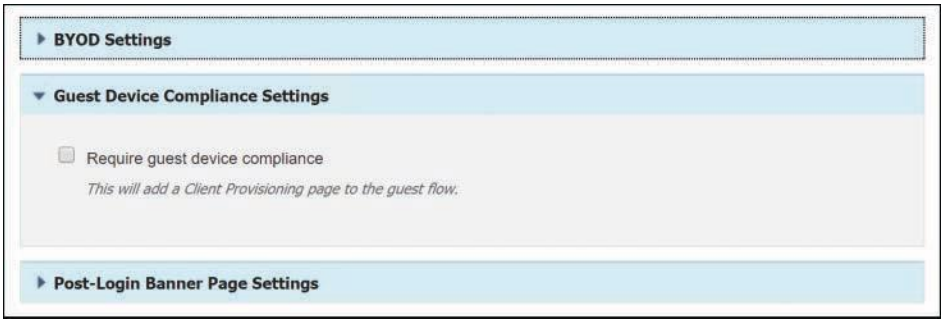


Figure 13-52 *Self-Registered Guest Portal Guest Device Compliance Settings*

Authorization Rule Configuration

So far you've been configuring separate SSIDs for the different guest portal types, mainly so it is easy to see and configure things differently and gain the familiarity you need for the SISE 300-715 exam.

In this section, for a self-registered guest portal, you use a WLAN that was pre-created on the WLC, with SSID ISE-SelfGuest, and assigned to the Guest VLAN. (See Chapter 11 for a refresher on creating a WLAN, if needed.)

You need to configure an authorization profile to redirect the user to the self-registered guest portal you configured in this chapter. Follow these steps in the ISE GUI (see Figure 13-53):

Step 1. Navigate to **Work Centers > Guest Access > Policy Elements > Results > Authorization Profiles**.

Step 2. Duplicate the **CWA** profile that you created in Chapter 12.

Step 3. Name the new profile **Self-Registered-CWA**.

Step 4. Leave the ACL field set to **ACL-WEBAUTH-REDIRECT**.

Step 5. Select **Self-Registered Guest Portal (default)** from the dropdown.

At this point, you should be feeling more comfortable with the technology, so in Figure 13-53 we show the Attributes Details section, where you can see that the URL the guest is redirected to includes a portal identifier—in this case, **deaaa863-1df0-4198-baf1-8d5b690d4361**. This is a unique identifier, and it's leveraged to ensure that when multiple portals are hosted on the same listener, the correct portal is presented to the guest.

With the authorization profile created, you can move forward with the authorization rule. To change things up a bit from what you did for the hotspot portal, here you create a new policy set just for the self-registered guest flows, where traffic from the ISE-Self-Registered SSID is redirected to the Self-Registered-CWA authorization profile.

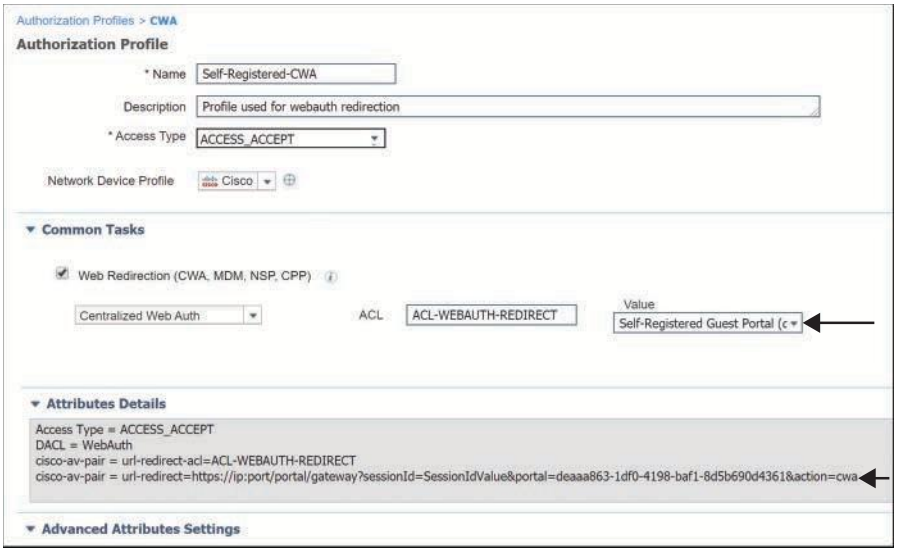


Figure 13-53 *Self-Registered-CWA Authorization Profile*

Follow these steps in the ISE GUI (see Figure 13-54):

Step 1. Navigate to **Work Centers > Guest Access > Policy Elements > Policy Sets**.

Step 2. Create a new policy set above the default set and name it **Self-Registered-Guest**.

Step 3. Use the **Radius Called-Station-ID** attribute, select **Contains**, and enter **ISE-Self-Registered**.



Figure 13-54 *Self-Registered Guest Policy Set*

Step 4. Save the condition to the library for use later.

Step 5. Insert a rule in the authentication policy and name it **MAB**.

Step 6. Use the **Wireless_MAB** library condition.

Step 7. Set the resulting identity store to **Internal Endpoints**.

Step 8. Select **CONTINUE** from the If User not found dropdown, as shown in Figure 13-55.

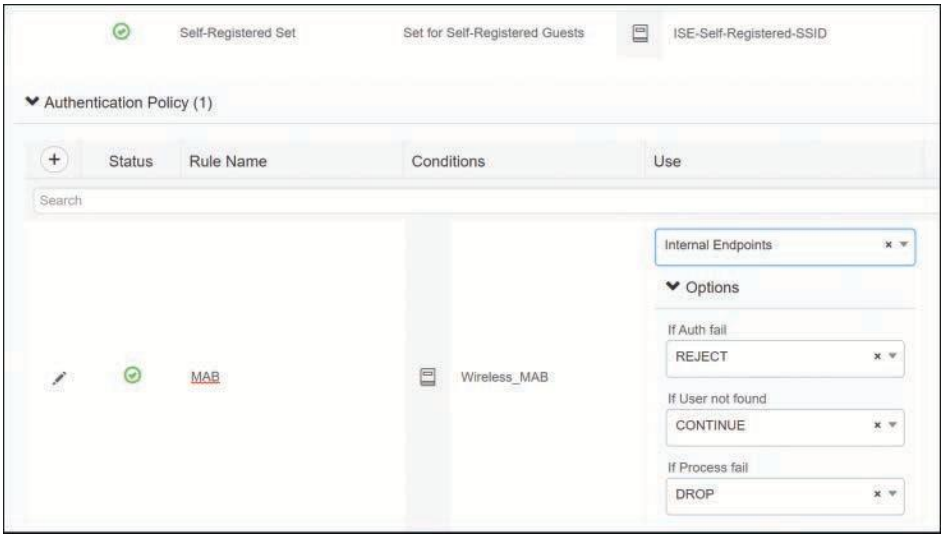


Figure 13-55 *Authentication Policy*

Step 9. Insert a new authorization rule above the default rule and name the rule **Redirect to CWA.**

Step 10. Ensure that the conditions of the new rule include the **ISE-Self-Registered-SSID** condition.

Step 11. In the Authorization Profile column, select the **Self-Registered-CWA** profile.

Step 12. Duplicate the new rule above the Redirect to CWA rule and name the new rule **Internet Only.**

Step 13. Include an additional condition for **Identity Group Name EQUALS Endpoint Identity Groups:Guest Endpoints.**

Step 14. Select **Internet-Only** for the authorization profile.

Step 15. Save the policy set.

Figure 13-56 shows the new authorization rules in the policy set.

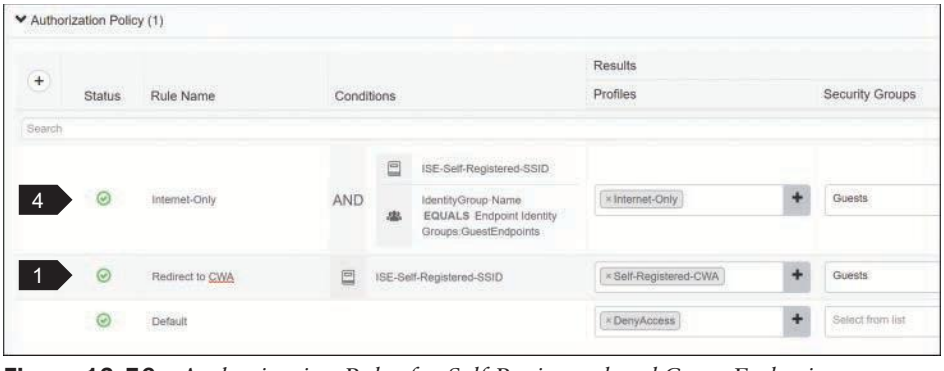


Figure 13-56 *Authorization Rules for Self-Registered and Guest Endpoints*

Figure 13-57 shows the RADIUS Live Log for the self-registered guest flows. The Live Log entry labeled 1 corresponds to the authorization rule labeled 1, where the unknown endpoint is redirected to the self-registered guest portal. The entry labeled 2 is where the end user goes to enter credentials for login to the network, and 3 is where the CoA occurs. Finally, the live log entry labeled 4 corresponds to the authorization rule that is also labeled 4, which is where the endpoint is given access to the Internet.



Figure 13-57 *Live Log for Self-Registered Flows*

Figures 13-58 through 13-64 show the end-user experience from a mobile device. Figure 13-58 shows the portal when the user is first redirected there. A user who does not yet have an account needs to click the Or Register for Guest Access link to begin the registration process.

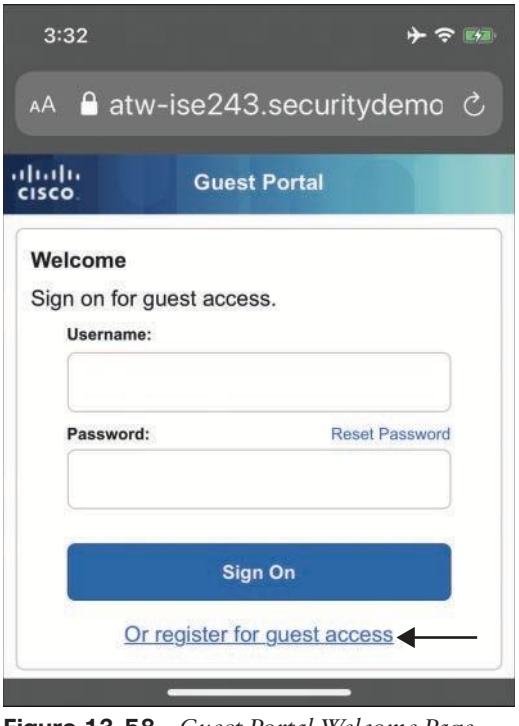


Figure 13-58 *Guest Portal Welcome Page*

Figure 13-59 shows the registration screen, with the fields filled out.

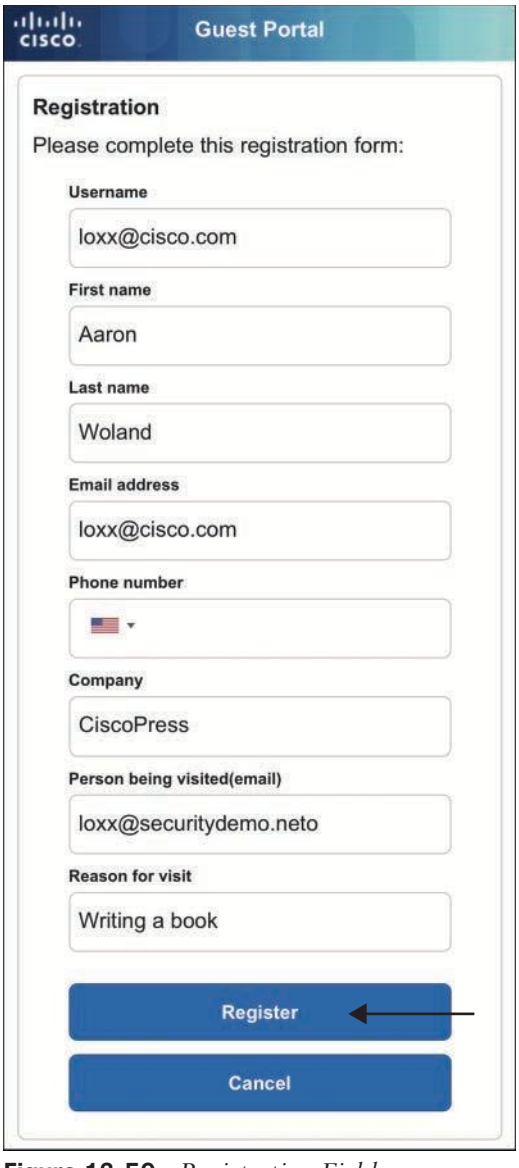


Figure 13-59 *Registration Fields*

Figure 13-60 shows a successfully created account, with the auto-generated password, and provides the guest the ability to click Sign On to enter those credentials without dropping and rejoining the guest network.

Figure 13-61 shows the portal again, but this time the user is able to log in with the newly created credentials.



Figure 13-60 *Successful Creation of Guest Account*

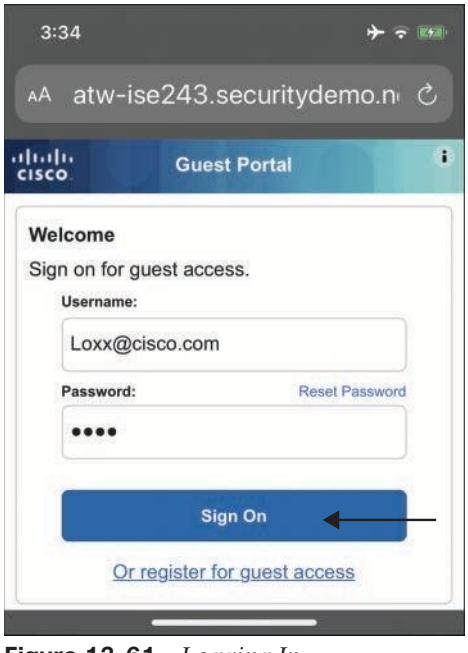


Figure 13-61 *Logging In*

Figure 13-62 shows the acceptable use policy. When the user clicks Accept, the endpoint's MAC address is added to the Guest Endpoints identity group, and the CoA is sent to the NAD.

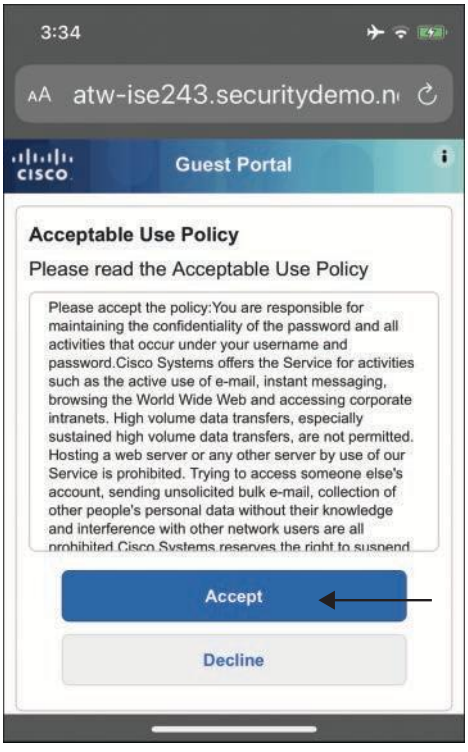


Figure 13-62 *Registration Fields*

Figure 13-63 shows the continue banner.

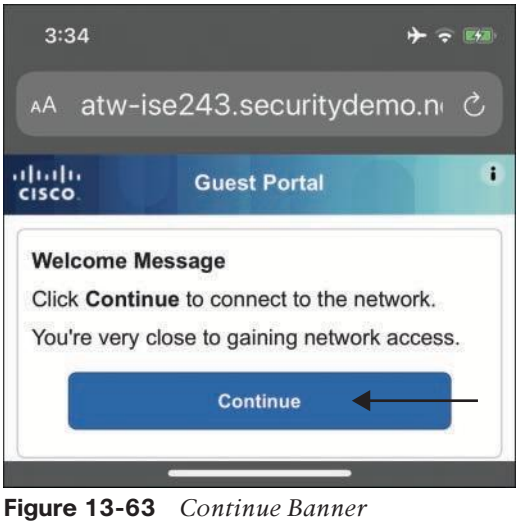


Figure 13-63 *Continue Banner*

Figure 13-64 shows the success message. If you configured a different result, such as directing the user to the company home page, that would be displayed instead.

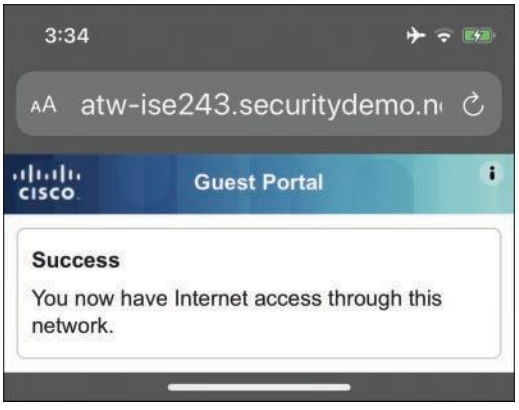


Figure 13-64 *Success Message*

Configuring a Sponsored Guest Portal

The sponsored guest portal seems basic in comparison to the self-registered guest portal, doesn't it? In fact, you already used the sponsored guest portal with the authorization rules created in Chapter 12.

Figure 13-65 shows a comparison of the default flow for a sponsored guest portal and the default flow for the self-registered guest portal. It is fairly obvious that the self-registered portal flows are more complex.

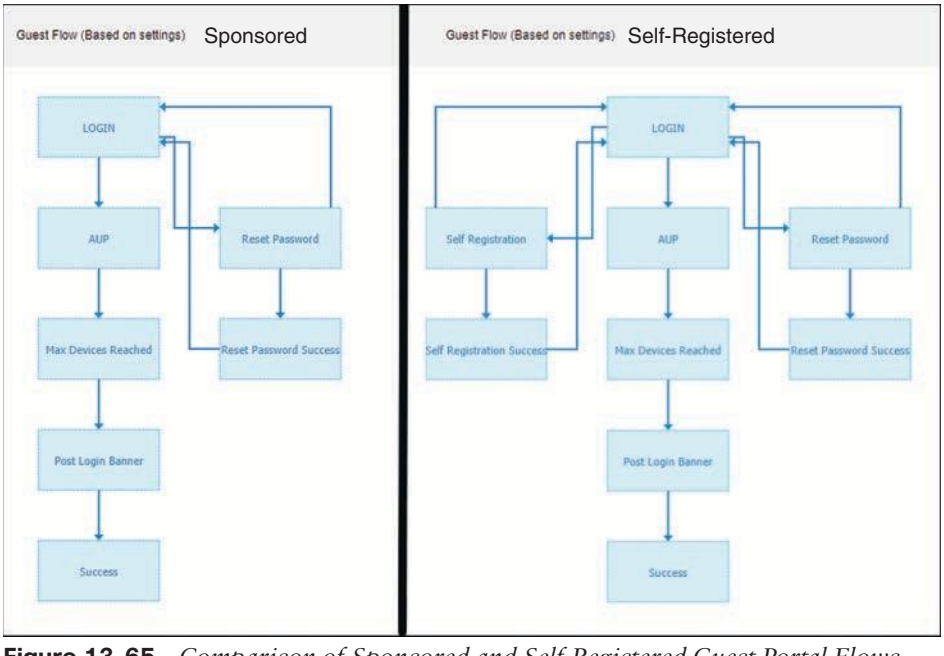


Figure 13-65 *Comparison of Sponsored and Self-Registered Guest Portal Flows*

The good news is that there are no more portal settings for the sponsored guest portal that you have not already examined in this chapter, so we do not have to dive into any more portal settings. Figure 13-66 illustrates the portal settings for sponsored and self-registered guest portals.



Figure 13-66 *Comparison of Sponsored and Self-Registered Guest Portal Settings*

The next section covers sponsors and their portals. We will come back to the sponsored guest portal when we test and validate the use of sponsored guest accounts, later in this chapter.

Sponsors

A sponsor is an employee who introduces and, ultimately, takes responsibility for a guest. The term *responsibility* in this instance means the employee acts as the point person during an incident response or an HR-related investigation involving the guest user.

Sponsor Groups

It is possible to have different sponsor groups, with each one designed to manage a different set of guest users. Figure 13-67 shows the three default sponsor groups:

- **ALL_ACCOUNTS**: Members of this group can edit, delete, reinstate, and approve all guest accounts.
- **GROUP_ACCOUNTS**: Members of this group can edit, delete, reinstate, and approve all guests created by sponsors who are members of the same group.

- **OWN_ACCOUNTS**: Members of this group can only manage the guest accounts they have sponsored themselves.

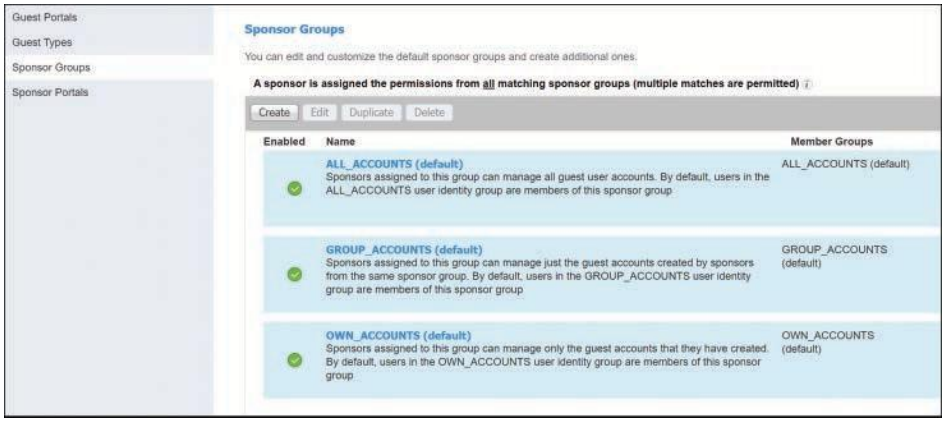


Figure 13-67 *Sponsor Groups*

Sponsor accounts could be local user identities created in ISE and assigned to the default groups, or you can leverage a more detailed policy to define group membership. The most common method is to use an Active Directory group; this option is so common that it appears right on the main configuration page, as shown in Figure 13-68.

Figures 13-68 through 13-70 show the details of the ALL_ACCOUNTS default sponsor group. Clicking the highlighted Members button brings up a list of Active Directory groups to add, and the figure shows that the AD group named managers is added to this sponsor group, as shown in Figure 13-69. If you want to use any other condition (such as LDAP attribute), you must click Create New Condition.

In addition to defining which users belong to the sponsor group, the settings on this page define which guest types the sponsor can manage and in which company locations (time zones).

More sponsor configurations exist, as shown in Figure 13-70. These settings determine whether a sponsor can create bulk accounts and the limits that exist for those bulk creations. These bulk accounts can be imported from a comma-separated values (CSV) document or created at random.

Each sponsor group is assigned different permissions for the guest accounts. This screen indicates whether a sponsor can edit, delete, update, or even view the details of a guest account, including the password. In addition, you can specify whether a sponsor account is able to use the Guest API.

If a sponsor matches multiple sponsor rules, the permissions are additive. So, if one group has permissions to approve contractors, and another group has permissions to approve daily, the resulting permissions for the sponsor would be both Contractors and Daily.

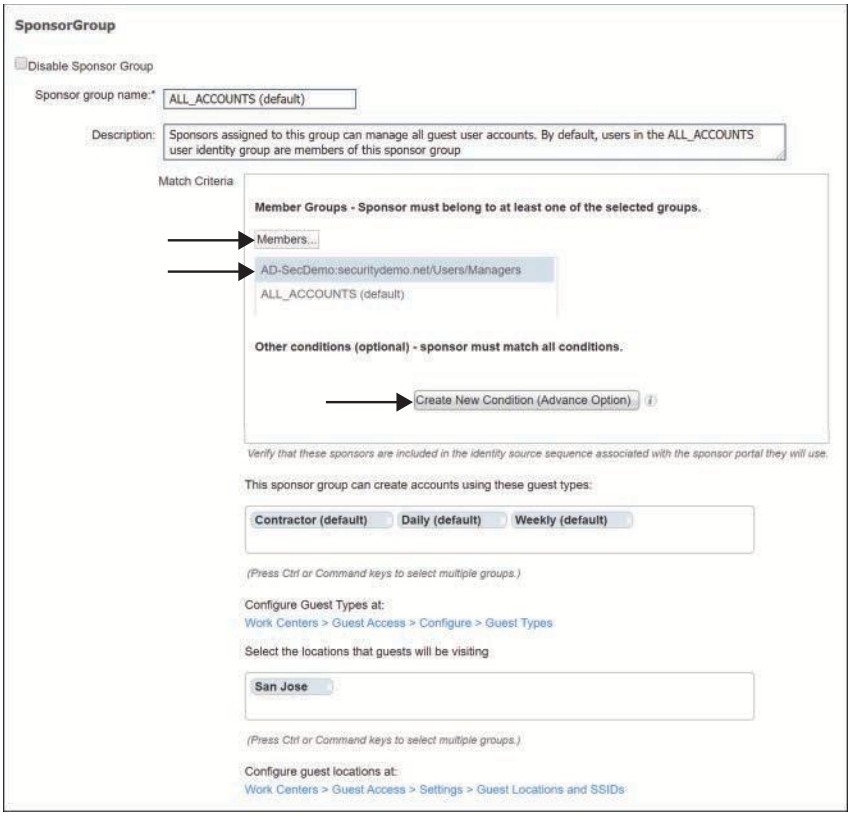


Figure 13-68 *ALL_ACCOUNTS Sponsor Group, Part 1*

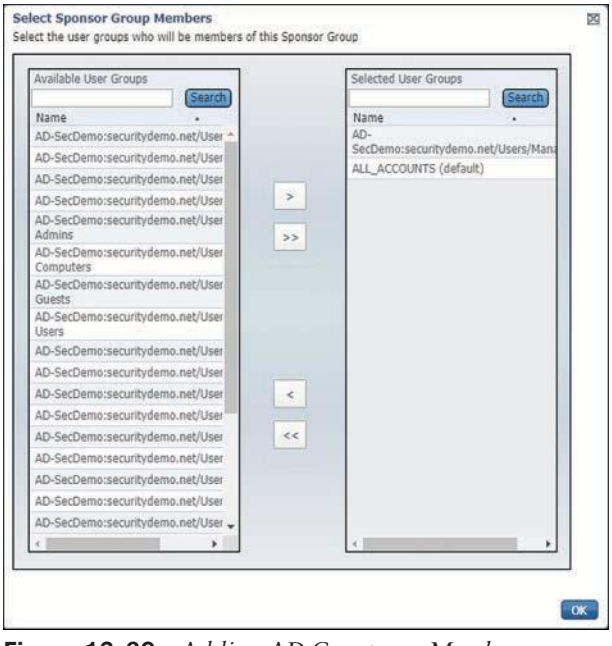


Figure 13-69 *Adding AD Groups as Members*

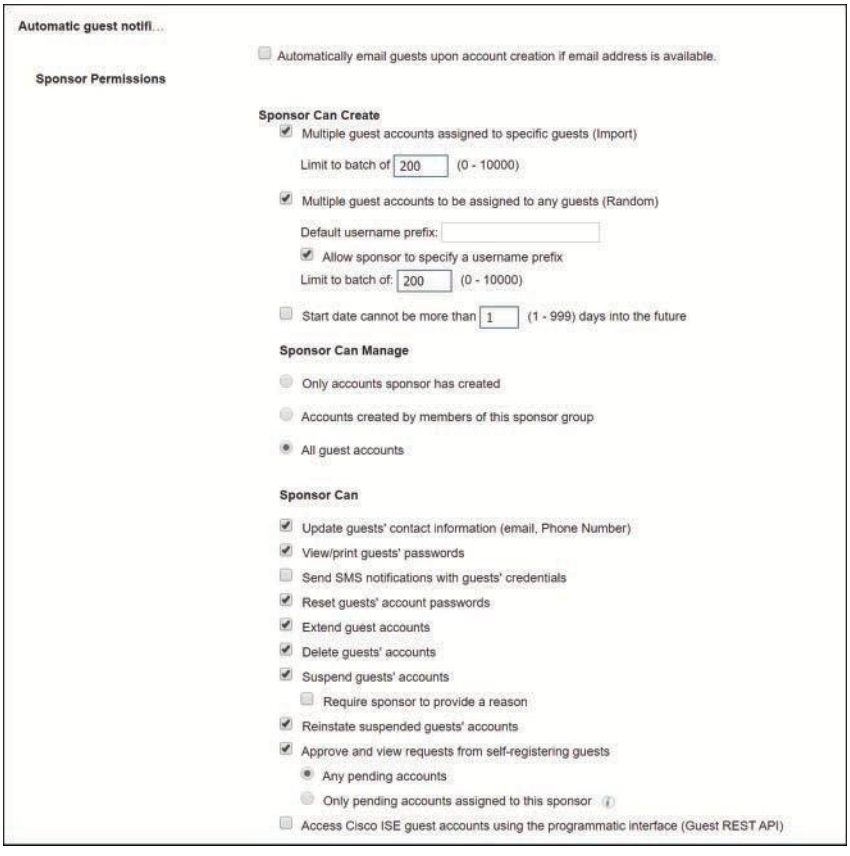


Figure 13-70 *ALL_ACCOUNTS Sponsor Group, Part 2*

Sponsor Portals

Just as ISE offers portals for guests, it offers sponsor portals. These portals are the series of web pages that a sponsor can access to generate, approve, and manage guest accounts. Unlike the guest portals, there is only one sponsor portal created by default. With guest portals, the portal type itself can dictate much of the allowed content and the flows within the portal. With sponsor portals, the sponsor group dictates all the settings.

You can take a look at the default sponsor portal by navigating to Work Centers > Guest Access > Portals & Components > Sponsor Portals > Sponsor Portal (default), as shown in Figure 13-71.

As Figure 13-72 shows, like a guest portal, a sponsor portal has page settings and a dynamically updated flow diagram, and it can be customized.

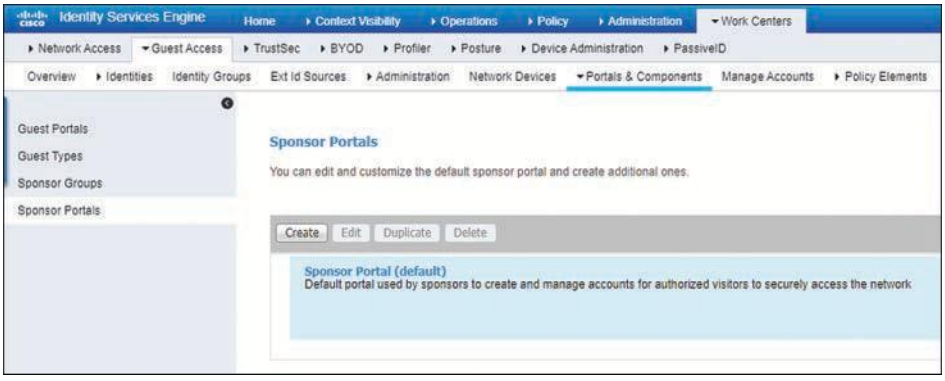


Figure 13-71 *Sponsor Portals*

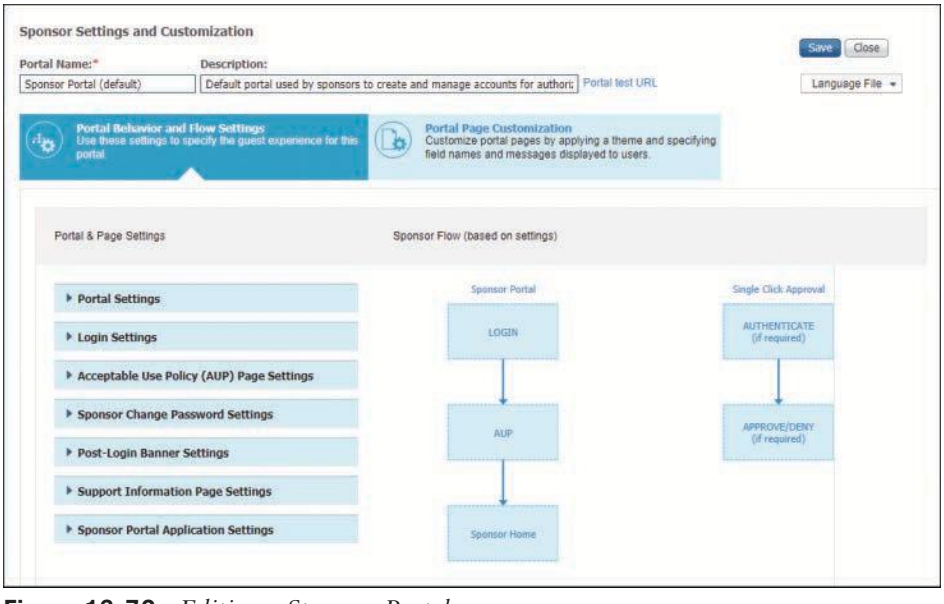


Figure 13-72 *Editing a Sponsor Portal*

Portal Settings

It is not the ISE administrator or IT expert who normally logs in to the sponsor portals. Rather, employees of the organization are the end users.

The sponsor portal settings are very similar to the guest portal settings, as you can see in Figure 13-73. A port must be specified, and TCP port 8445 is the default port. Interface selection and bonded interface choices are exactly the same as for the guest portals. Because this portal will be accessed by end users, the same guidance around using a publicly signed certificate holds true.

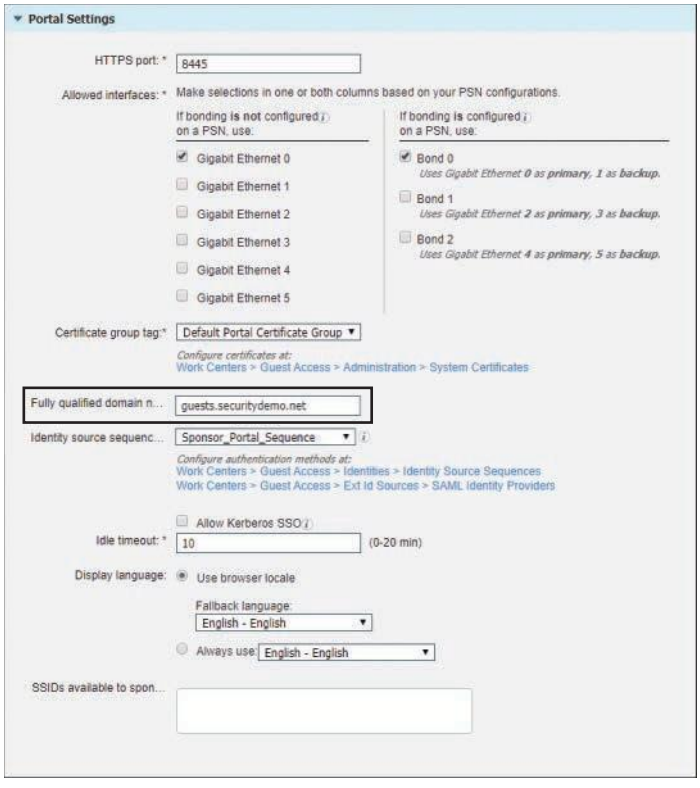


Figure 13-73 *Editing a Sponsor Portal*

A guest who joins the network is redirected to a specific portal based on the authorization result. Sponsor access to a portal is a unique paradigm: A sponsor needs to be able to simply enter a friendly URL into a browser and access the right portal. In order to make that possible, you must fill in the Fully Qualified Domain Name field in the Portal Settings section, as highlighted in Figure 13-73.

The selected identity source sequence can be an ISS, or SAML, but the two choices are mutually exclusive. In other words, SAML cannot be combined with an ISS. SAML is used for SSO and is a very popular option for sponsor portals in many organizations because it prevents the end user from using yet another login.

Login Settings and AUP Page Settings

Much as with guest accounts, you can specify a maximum number of failed logins, after which there is a time limit between login attempts. Notice in Figure 13-74 that the AUP settings can be configured within the Login Settings and the Acceptable Use Policy (AUP) Page Settings areas. This is just another way the user experience is designed to keep things easy.

You might require an acceptable use policy for a sponsor for much the same reason you would include one for a guest, but the wording and requests for guests and sponsors are very different. The AUP for a sponsor is not about what is or isn't acceptable behavior for their Internet access but instead is focused on what is or isn't acceptable behavior for creating a guest account.

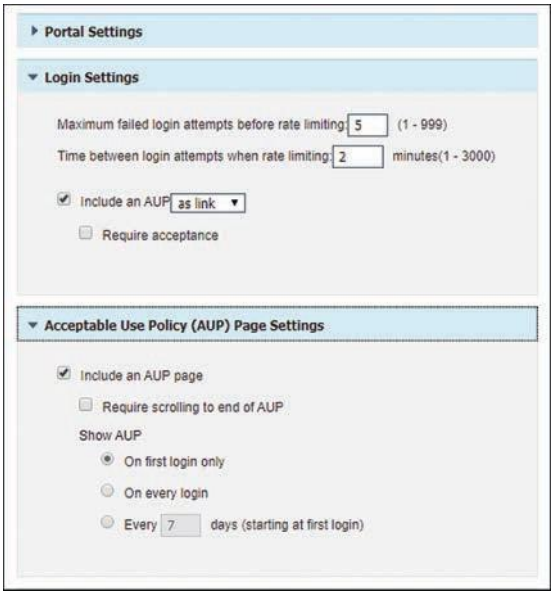


Figure 13-74 *Login Settings and Acceptable Use Policy (AUP) Settings*

The Remaining Settings

Figure 13-75 shows the remaining sponsor portal settings. You can permit or deny sponsors to change their passwords through the portal, display a post-login banner, and display support information (exactly as with the guest portal). Finally, there is a Sponsor Portal Application Settings field, but it is configured only when you customize the pages.

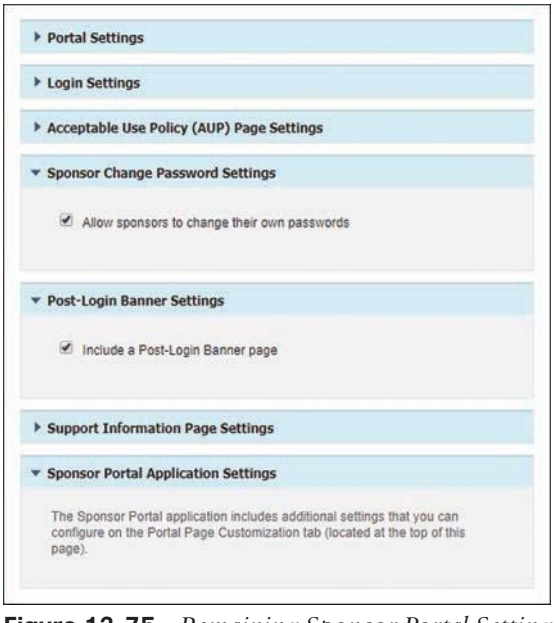


Figure 13-75 *Remaining Sponsor Portal Settings*

Notification Services

When you enable guests to create accounts, the guests need to be able to learn what their credentials are. That is where the SMTP and SMS gateway settings under the Administration menu come in to play.

SMTP Servers

To send a guest credentials via email, you must configure an SMTP server under Work Centers > Guest Access > Administration > SMTP Server, as shown in Figure 13-76.

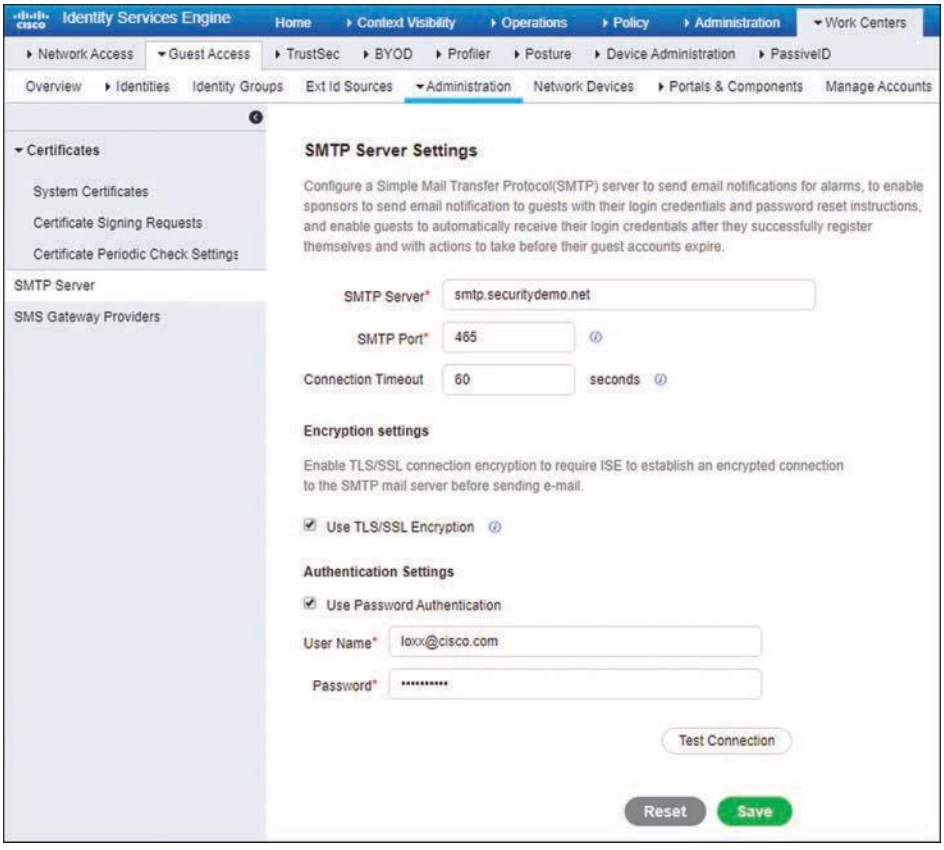


Figure 13-76 *SMTP Server Settings*

As shown in Figure 13-76, ISE can be configured to use standard SMTP, and you can also provide support for TLS/SSL encryption authentication. However, do not go looking for authenticated SMTP in ISE before Version 2.7 as Version 2.7 is the first version that supports secure SMTP.

SMS Gateway Providers

Email is not the only way to provide guests with their credentials. It is possible, of course, to print them out, but a more commonly used method is to send credentials in a text message.

This requires an SMS gateway. Numerous services provide SMS gateways, and many of them are listed in the ISE UI, as shown in Figure 13-77.

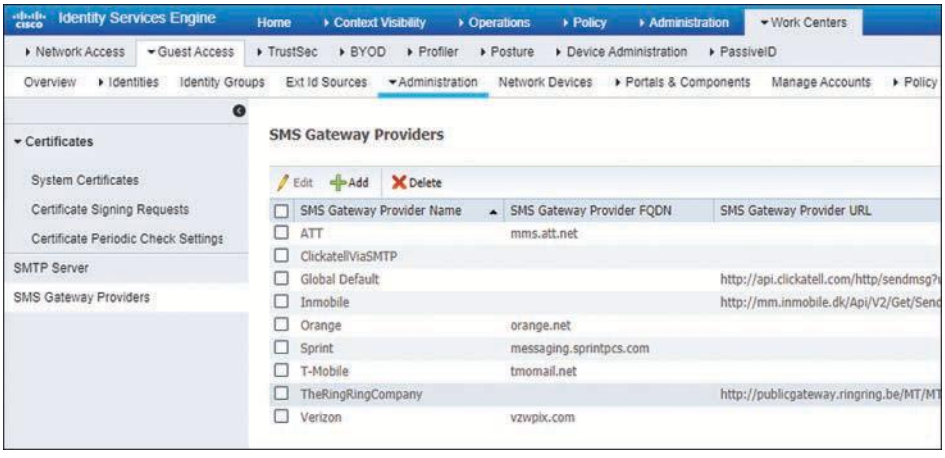


Figure 13-77 *SMS Gateway Providers*

Provisioning Guest Accounts from a Sponsor Portal

Now that you've looked at dozens of configuration screens, you're ready to see an example of adding a guest. Figure 13-77 shows the fully qualified domain name entered as a friendly name. That friendly name must correspond to a DNS entry that points to the PSNs.

You could always use multiple A records for DNS round-robin load balancing or leverage a load balancer, or maybe you could even use Anycast. The bottom line is that when an end user types that friendly name into a browser, it needs to resolve to the ISE PSNs.

When the browser connects to the PSN with the friendly name in the host header, the browser gets redirected to port 8445, where the sponsor portal is being hosted. Figure 13-78 illustrates this redirection, and Figure 13-79 shows the end user being presented with the sign-on page for the sponsor portal.

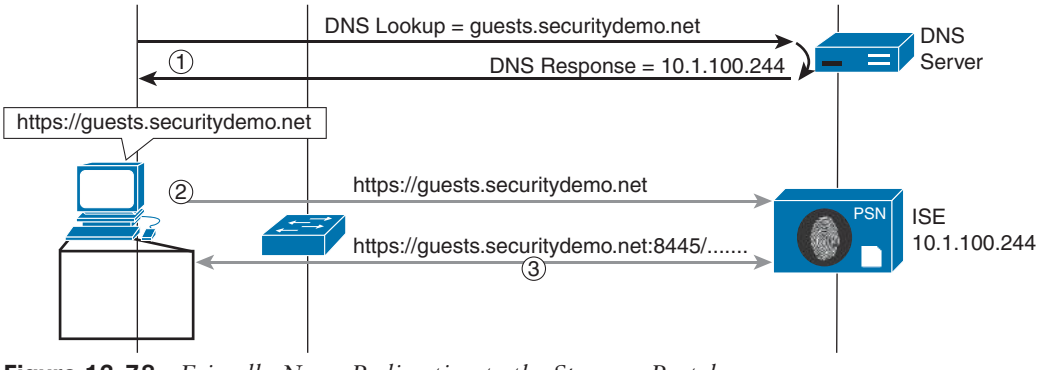


Figure 13-78 *Friendly Name Redirection to the Sponsor Portal*

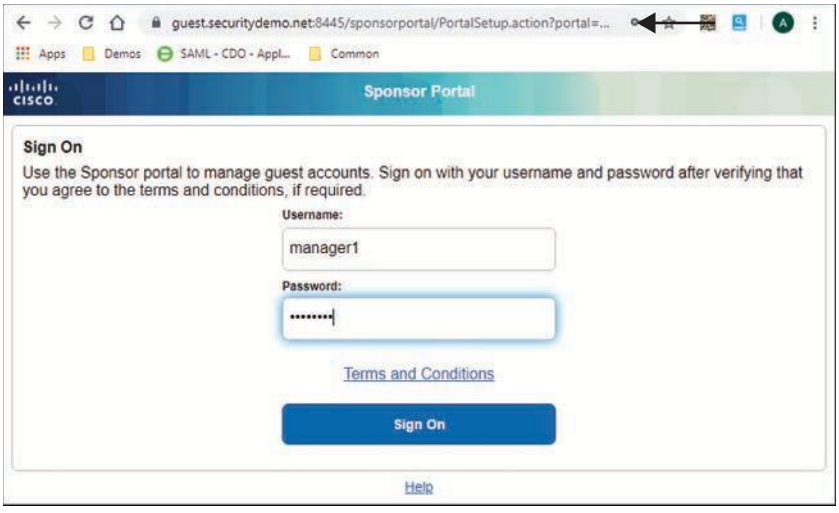


Figure 13-79 *Sponsor Portal Sign-On Page*

After logging in, the sponsor is presented with the AUP, as shown in Figure 13-80.

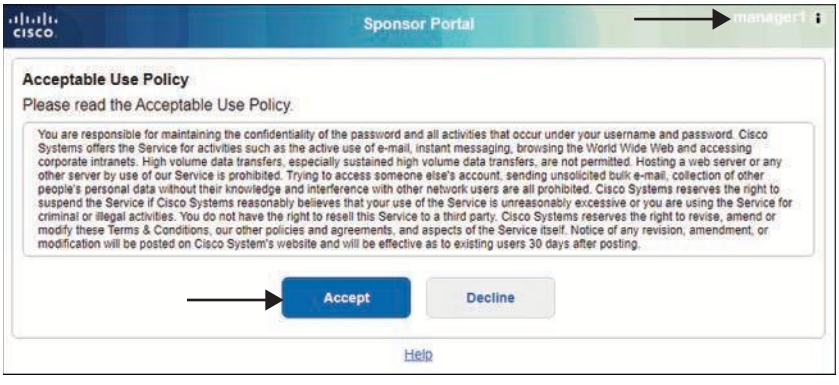


Figure 13-80 *Sponsor AUP*

After accepting the AUP, the sponsor is shown the post-access banner that you configured and then directed to the main page, as shown in Figures 13-81 and 13-82.

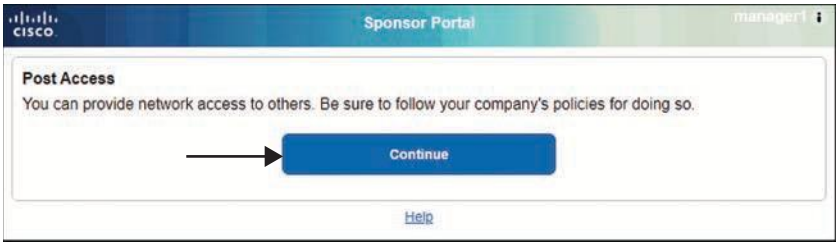


Figure 13-81 *Sponsor Post-Access Banner*

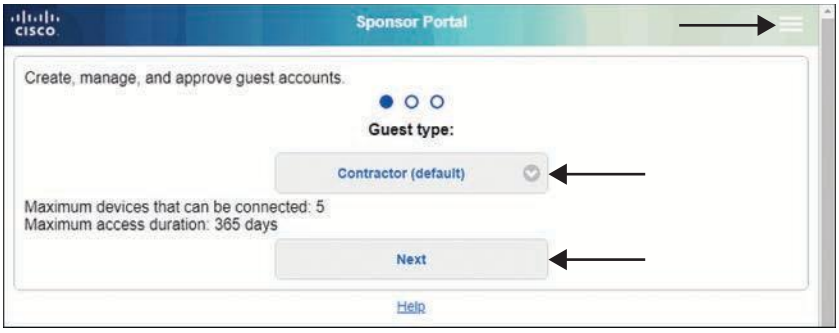


Figure 13-82 *Sponsor Portal Main Page*

From the main page, you can select the guest type to create, or you can click the menu icon in the upper-right corner. Figure 13-83 shows the menu, and Figure 13-84 shows the screen for creating a new contractor guest account.

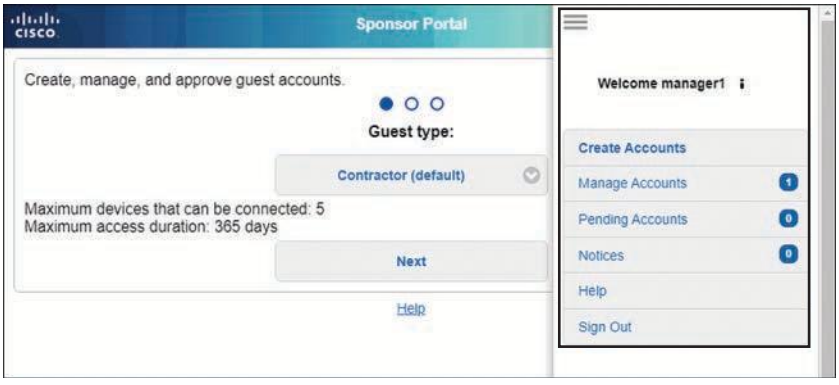


Figure 13-83 *Sponsor Menu*

In Figure 13-84, you can see that this manager is able to type the specific information for a known guest, create random guests, or import a CSV list of guests for an event or a similar bulk-creation need.

Figure 13-85 shows the assignment of the guest account's lifetime. Because this is a contractor account, the maximum is 365 days, but the default is only 90 days. Figure 13-86 shows the account details after creation.

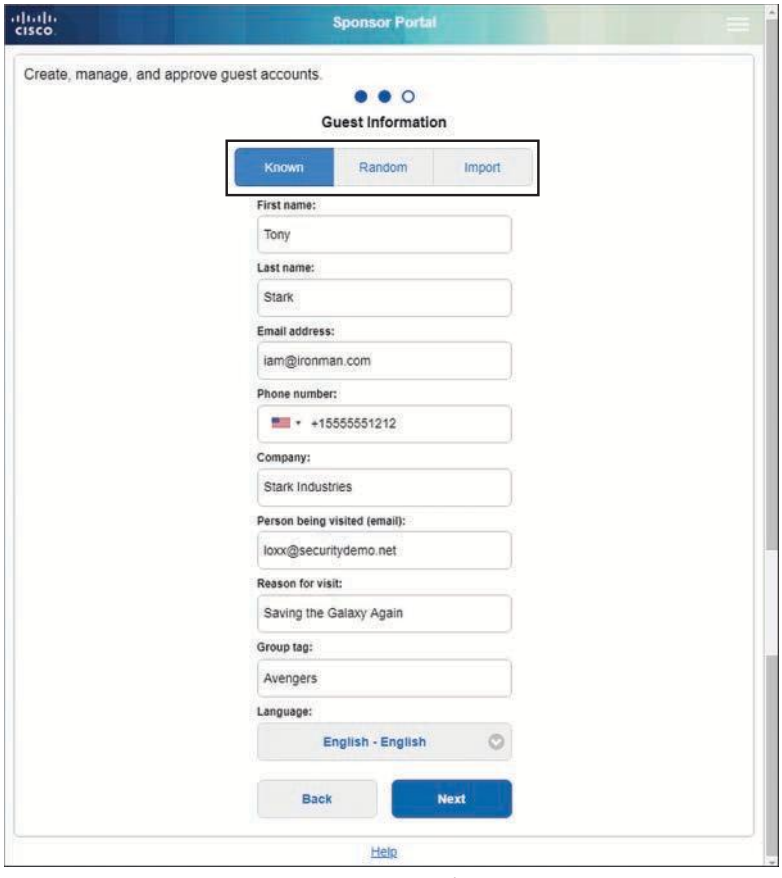


Figure 13-84 *Using the Sponsor Portal to Create a Contractor Guest Account*

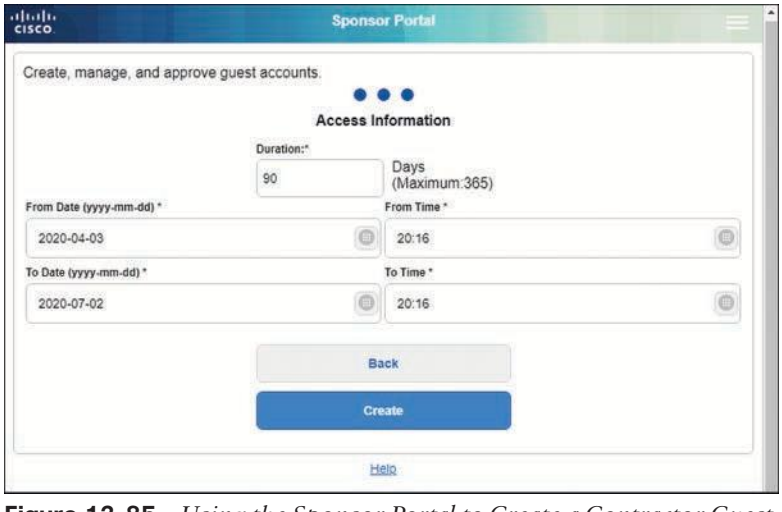


Figure 13-85 *Using the Sponsor Portal to Create a Contractor Guest Account's Access Duration*



Figure 13-86 *Sponsor Portal Guest Credentials*

Clicking Notify brings up a pop-up for the sponsor to select the notification method. As shown in Figure 13-87, any configured notification methods are displayed, such as Print, Email, and SMS.

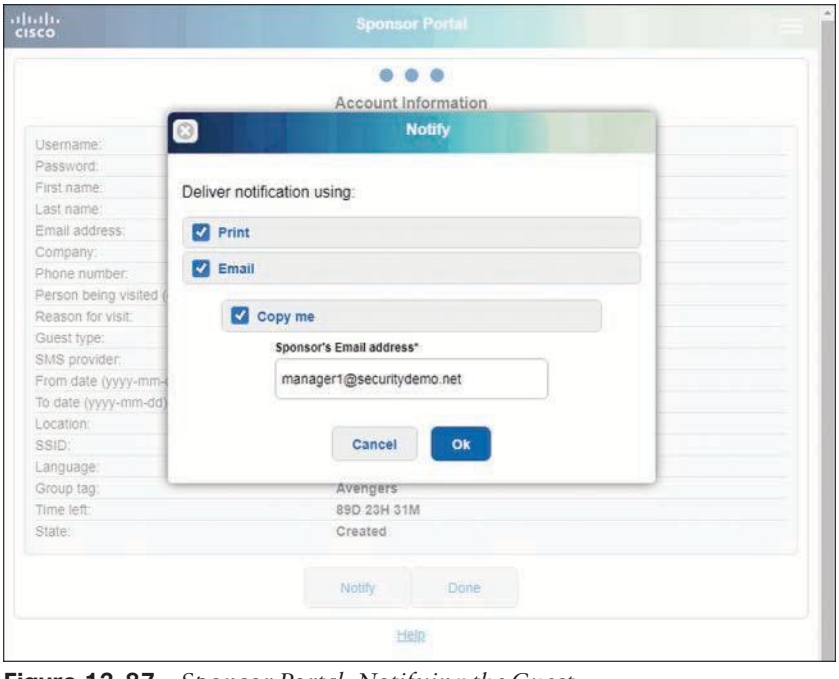


Figure 13-87 *Sponsor Portal: Notifying the Guest*

A sponsor is not only capable of creating guest accounts but can also manage those accounts. Using the menu in the upper-right corner of the sponsor portal, a sponsor can select to manage existing accounts, approve pending accounts, or respond to notices. Figure 13-88 shows the Manage Accounts screen, showing the two guest accounts created in this chapter.

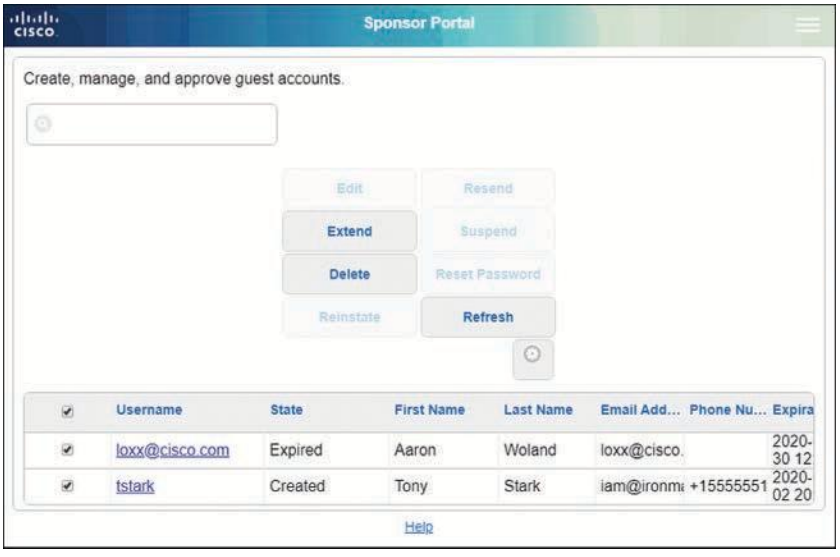


Figure 13-88 *Sponsor Portal: Managing Accounts*

The first account was created during the self-registration process, and the second account was created just now. When you select a guest account to manage, any available actions become active, as shown for Extend, Delete, and Refresh in Figure 13-88.

SAML Authentication

Now that you understand guests and sponsors, we are going to backtrack to the Ext Id Sources tab of the Guest Access work center and configure a SAML ID provider.

Security Assertion Markup Language (SAML) is an XML-based open standard for authentication and authorization that can be leveraged for web-enabled SSO.

Some noteworthy definitions go along with SAML:

- **SAML service provider (SP):** The *SAML SP* is the application or service that is being accessed (for example, the ISE web portal). AuthN and AuthZ are required to access the SAML SP.
- **SAML identity provider (IdP):** The *SAML IdP* is a policy engine that determines authorization and issues SAML assertions. The IdP is the identity provider against which ISE is authenticating.

- **SAML assertion:** Much like a web cookie, a ***SAML assertion*** contains the identity and attributes of the user and the user's authorization for the service, which is passed from the IdP to the SP.

Using SAML authentication on a web portal makes it possible for users to authenticate to the sponsor portal with SSO. Cisco ISE is SAMLv2 compliant and supports all SAMLv2-compliant IdPs. It's important to note that those IdPs must use Base64-encoded certificates.

Figure 13-89 illustrates the SAML flow for login to the guest portal. The sponsor never even sees the sponsor portal login screen. The first screen the user sees is the login for the IdP (if this is first login), and then the user sees the acceptable use policy. If the user agent (browser) already has a valid cookie from the IdP, the user does not even see the login screen but immediately sees the AUP for the sponsor portal.

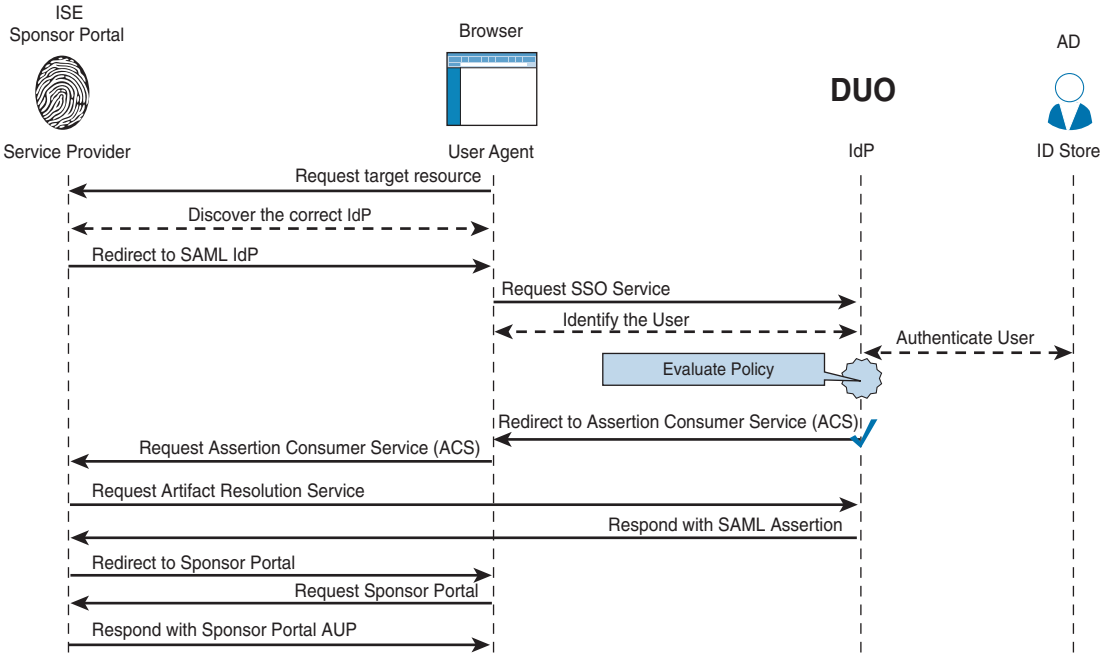


Figure 13-89 *SAML Flow for the Sponsor Portal*

NOTE Before you can add a SAML IdP, you must first trust that IdP's certificate. From the ISE UI, navigate to Administration > System > Certificates > Trusted Certificates > Import.

This section shows how to add Duo Access Gateway (DAG) as the IdP. Figure 13-90 shows the DAG screen you use to download the metadata and the certificate for the DAG. Whatever IdP you use should have a similar download location.



Figure 13-90 *Downloading the Certificate and Metadata from the IdP*

After you have imported the IdP's certificate into ISE's trusted certificate store, you can add a SAML IdP by following these steps:

Step 1. Navigate to **Work Centers > Guest Access > Ext Id Sources > SAML Id Providers** and click **Add.**

Step 2. Enter a friendly name for the IdP, as shown in Figure 13-91.



Figure 13-91 *Adding an IDP: General Tab*

Step 3. Click the **Identity Provider Config** tab.

Step 4. Choose the XML metadata file downloaded from your IdP. This parsed information is displayed on the screen, as shown in Figure 13-92.

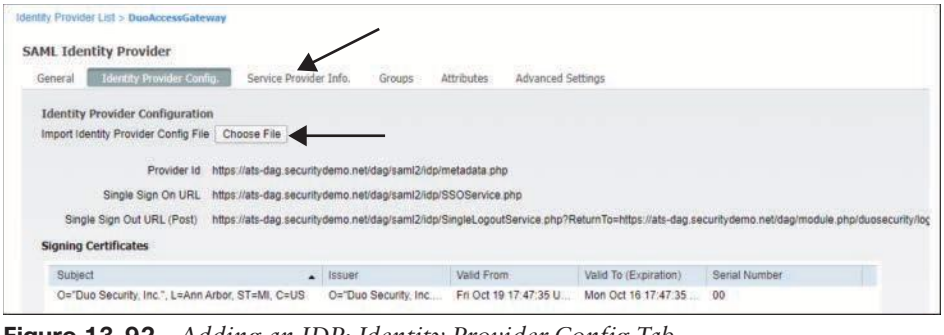


Figure 13-92 *Adding an IDP: Identity Provider Config Tab*

Step 5. Click **Save**. You must leave this area of the configuration and configure the sponsor portal to use SAML instead of the ISS, as shown in Figure 13-93, and then you can come back to step 6. By selecting the IdP in the portal settings, you tell this SAML configuration to protect the sponsor portal, and the service provider information is available to download.

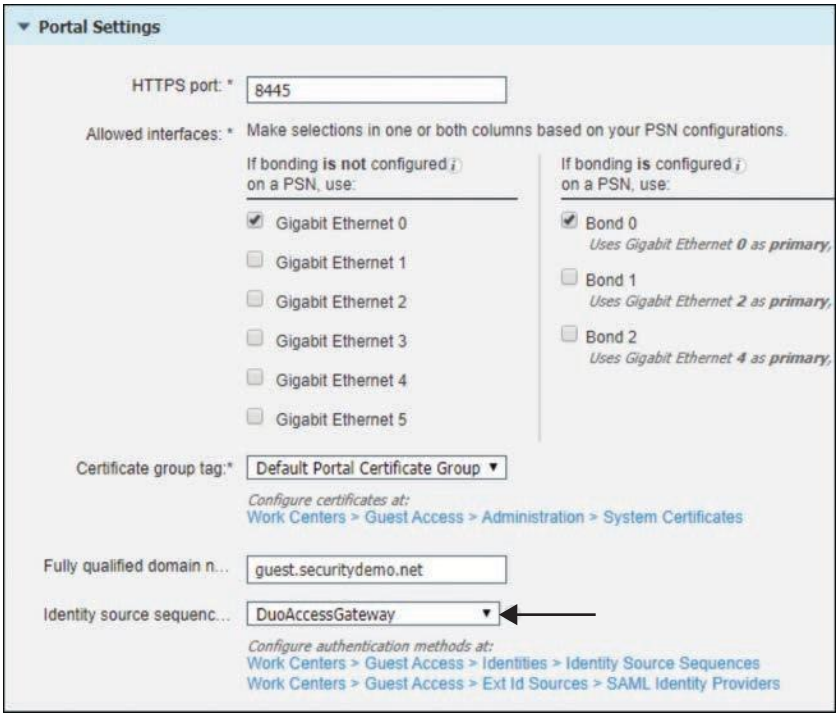


Figure 13-93 *Selecting the SAML IdP in the Portal Settings Section*

Step 6. Click the **Service Provider Info** tab.

Step 7. Click **Export** to download the SAML-SP XML file, as shown in Figure 13-94.



Figure 13-94 *Adding an IDP: Exporting Service Provider Information*

Step 8. Extract the resulting zip file, which contains the XML configuration and a README file for instructions for common IdPs.

Step 9. Import the XML file into your IdP or follow the instructions in the README file. The DAG does not support the importing of the XML, so when you use it, you must manually extract the correct fields from the file and paste them into the DAG, as shown in Figure 13-95.

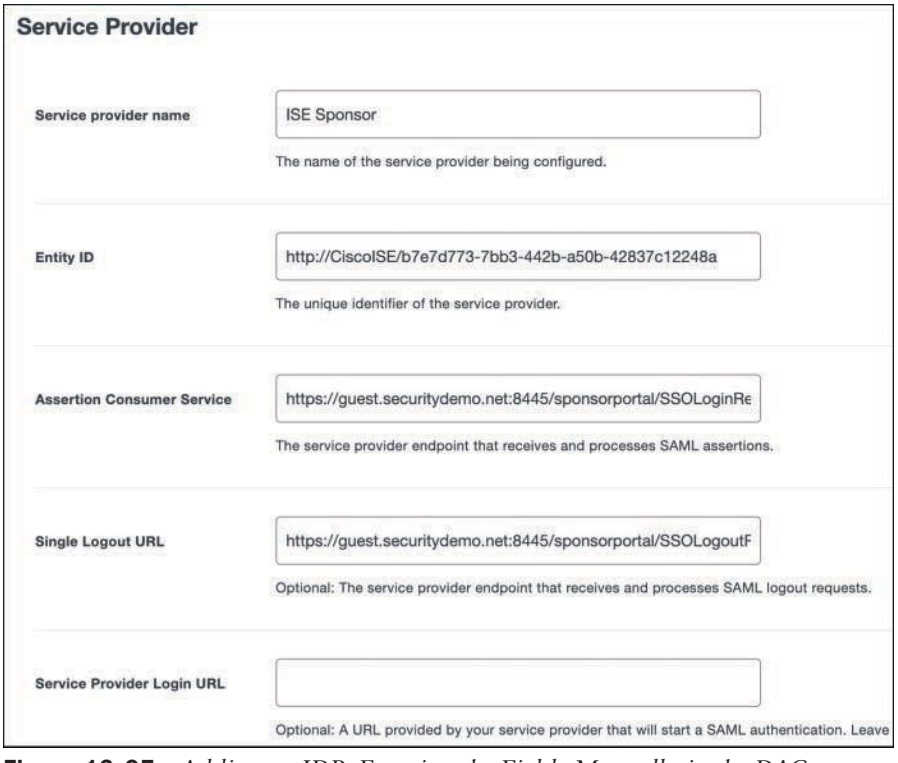


Figure 13-95 *Adding an IDP: Entering the Fields Manually in the DAG*

Step 10. Click **Save** and ensure you are saving the configuration in your IdP as well.

Figure 13-96 shows the manager1 user navigating to the friendly name the guest.securitydemo.net, and Figure 13-97 shows the user being automatically redirected to the login for the IdP. After successful login, manager1 must perform with Duo-Push (see Figure 13-98) because that is how the IdP was configured. When the MFA is successful, manager1 is redirected to the AUP, as shown in Figure 13-99.

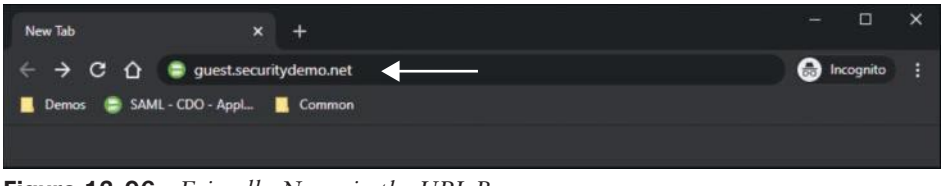


Figure 13-96 *Friendly Name in the URL Bar*

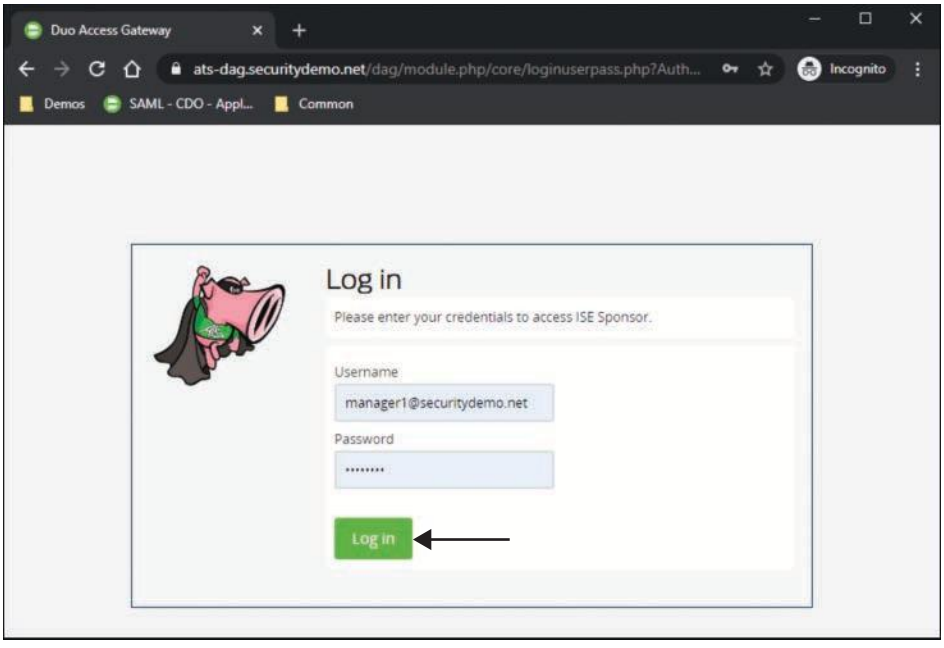


Figure 13-97 *IdP Login*

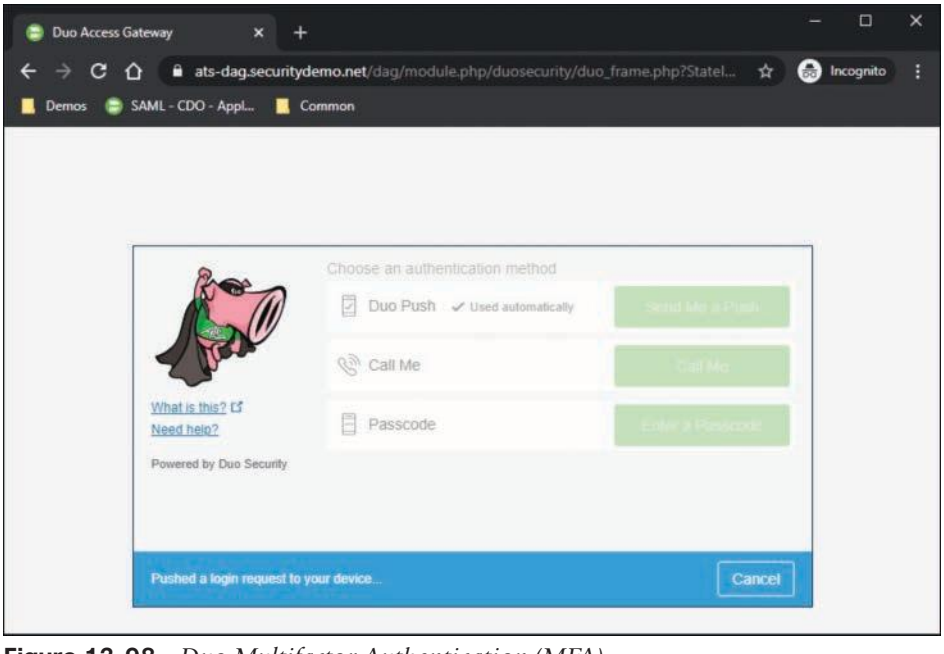


Figure 13-98 *Duo Multifactor Authentication (MFA)*

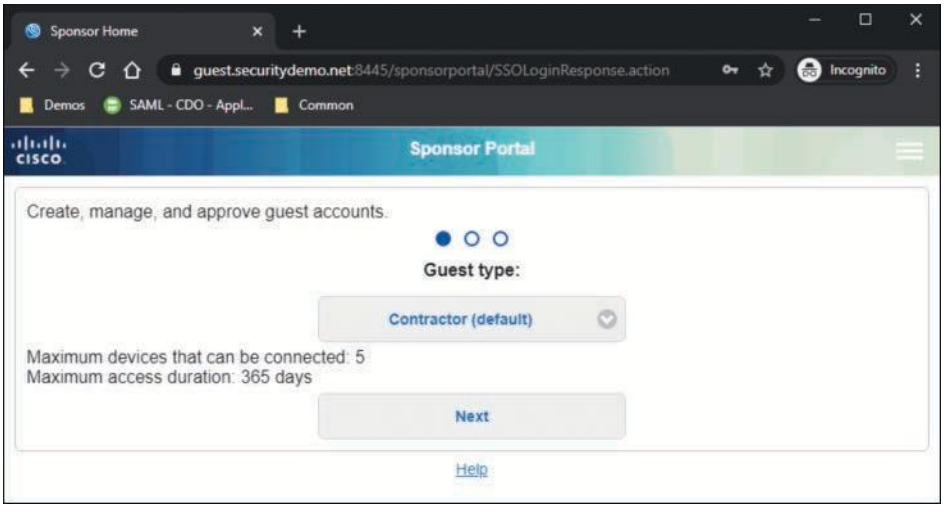


Figure 13-99 *Sponsor Portal AUP*

Call to Action

Guest access is a very common reason to deploy ISE. It is one of the premier tasks asked of an implementor of ISE. The authors of this book have seen many terrific deployments that are nearly perfect; however, they have also seen a few deployments that did not provide good end-user experiences because the implementors were clearly not very well versed on guest access capabilities, flows, and customizations.

As you are studying not only to pass the SISE 300-715 exam but also to gain a complete understanding of how ISE works in order to deploy it in production, play around! Don't stop with what is covered in this chapter only but continue to try different things.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 13-3 lists these key topics and the page number on which each is found.

Key Topic

Table 13-3 Key Topics

Key Topic Element	Description	Page
Paragraph	Certificates	340
Paragraph	Guest access configuration categories	341
Paragraph	Guest and sponsor configurations	341
Paragraph	Guest account types	343

Key Topic Element	Description	Page
Paragraph	Default guest types and where to use them	343
Paragraph	Authorizing sponsor types to create accounts for guest types	346
Paragraph	Authorization rule for redirecting users to a guest portal	349
Paragraph	Certificates for portal group tags and to protect a web portal	354
Paragraph	The importance of the rule order	364
Paragraph	The importance of time zones for guest accounts	369

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

guest, sponsor, portal, guest location, SAML service provider (SP), SAML identity provider (IdP), SAML assertion

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. Why would a company choose to use SAML authentication for a guest portal?
2. What is the purpose of having different guest types?
3. Why does a hotspot portal not include a username or password field?
4. How do visitors get redirected back to the hotspot portal after their time has expired?
5. What is the difference between a self-registered guest portal and a sponsored guest portal?



CHAPTER 14

Profiling

This chapter covers the following topics:

- Enable the profiling services
- Network probes
- IOS Device Sensor
- Feed Service
- Profiling policy rules
- Utilize profile assignment in authorization policies
- Verify profiling operation

There is often confusion about the difference between profiling and posture, and the terms tend to get mixed up. In Chapter 5, "Introduction to Advanced Concepts," you were exposed to both technologies. *Profiling* is the technology used to determine what a device appears to be from an external perspective, whereas *posture* is the technology used to perform an internal examination of a system and report the facts. Posture is covered in Chapter 18, "Posture Assessment."

Profiling is used to build a database of endpoints that are on the network and assign them to categories, based on their external attributes. While profiling started out as a technology used to automate MAC Authentication Bypass (MAB), it has become an integral part of access control.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 14-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 14-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
ISE Profiler	1, 3, 4, 8, 9
Profiling Policies	5, 6, 10

Foundation Topics Section	Questions
ISE Profiler and CoA	5
Profiles in Authorization Policies	2
Verify Profiling	7

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Where in the ISE GUI are profiling probes enabled?
   a. Administration > System > Deployment
   b. Administration > Feed Service > Profiler
   c. Work Centers > Profiler > Overview
   d. Work Centers > Profiler > Settings
2. Which of the following enable an endpoint profile to be used in an authorization policy rule? (Choose three.)
   a. Logical profiles
   b. Identity groups
   c. Nmap OS scan result
   d. EndPointPolicy attribute
   e. EndPointProfile attribute
3. Which probes are used to trigger the SNMPQUERY probe to query a NAD? (Choose two.)
   a. RADIUS
   b. SNMPQUERY
   c. HTTP
   d. SNMPTRAP
4. Which of the following are mechanisms that a Cisco wireless LAN controller can use to send DHCP information to ISE? (Choose two.)
   a. DHCP proxy
   b. Device Sensor
   c. Bridging DHCP to the LAN
   d. SPAN
5. How are updated profiles distributed to customer ISE deployments?
   a. Cisco's Profiler Feed Service distributes the new profiles.
   b. Each new version of ISE and/or ISE patch includes new profile policies.
   c. The profiles are distributed together with the posture checks and compliance modules.
   d. Update packs are downloaded from Cisco.com and then imported.

6. What determines when an endpoint is assigned to a profile?
   a. The profile that matches the most conditions is assigned.
   b. All profiles are manually assigned by the administrator.
   c. The certainty value must equal or exceed the minimum certainty value of the profile.
   d. The ISE posture agent identifies the profile of an endpoint to ISE.
7. Where in ISE can an administrator drill down into the profiles that have been assigned to locate a specific endpoint that has a particular profile?
   a. Endpoint Classification
   b. Cisco Endpoint Profiling Examination Tool (CEPET)
   c. Profiled Endpoints Counter
   d. Profiler Activity Window
8. Which of the following are ways to collect HTTP User-Agent strings? (Choose two.)
   a. Through the AnyConnect HTTP User Agent Reporting Tool
   b. Using SPAN port mirroring
   c. Using Cisco WSA Device Sensor
   d. Directly from ISE web portals
   e. Using Device Sensor in the switch
9. With IOS Device Sensor, which of the following probes become redundant? (Choose two.)
   a. NETFLOW
   b. DHCP
   c. SNMPQUERY
   d. DNS
10. What happens when an ISE administrator has modified a profile and then a Feed Service update is downloaded that contains an updated version of that profile?
   a. The profile is overwritten with the version in the Feed Service update.
   b. The administrator is prompted to choose to overwrite or ignore the profile update.
   c. All non-conflicting profiles are downloaded and installed. The conflicting profiles are ignored.
   d. The update fails, and an alarm is triggered on the dashboard and sent by email.

Foundation Topics

ISE Profiler

Cisco ***ISE Profiler*** is the component of the Cisco ISE platform that is responsible for endpoint detection and classification. It uses a probe or series of probes to collect attributes about an endpoint and then compares the collected attributes to predefined device signatures to locate a match.

In the early days of identity-based networks and 802.1X, humans spent countless hours identifying all the devices that did not have supplicants—that is, devices that could not authenticate to the network using 802.1X, such as printers and fax machines. After identifying the network port where a printer was plugged in, it was possible to configure the port to either not use 802.1X or use MAC Authentication Bypass (MAB). MAB is an extension to 802.1X that allows a switch to send the device's MAC address to the authentication server. If that MAC address is in the "approved list" of devices, the authentication server sends back an "accept" result, thereby allowing specific MAC addresses to bypass authentication.

Countless hours were spent collecting and maintaining this list of MAC addresses. A manual onboarding procedure would eventually have to be created so that when a new printer (or another valid device) was added to the network, its MAC address was added to the list. Obviously, some enhancements to this manual process were required. There had to be some way to build this list more dynamically to avoid spending so much time on prep and maintenance. As described in Chapter 5, this is where profiling technology entered the picture. Profiling allows you to collect attributes about devices from sources such as DHCP, NetFlow, HTTP User-Agent strings, DNS, Active Directory, and more. Those collected attributes are then compared to a set of signatures; this is similar to the way an intrusion prevention system (IPS) signature works. In ISE, these signatures are commonly referred to as *profiles*.

For example consider an endpoint connecting to wireless. While connected, the wireless controller is able to collect and send the MAC address and HTTP User-Agent string to ISE. ISE then determines that the MAC address's OUI belongs to Apple and the HTTP User-Agent string contains the word ipad. Therefore, ISE then assigns that device to the profile Apple-ipad.

Key Topic

Profiling technology has evolved, as technology tends to do. Today an authentication server (ISE) has the ability to use profiling data for much more than just building the MAB list. Cisco ISE now uses the collection and classification data from ISE Profiler as conditions in the authorization policy. Now you can build an authorization policy that looks at much more than identity credentials. You can combine a user's identity with a profiling result as a condition for authorization.

Figure 14-1 and Figure 14-2 show an example of a differentiated authorization policy based on profiling.

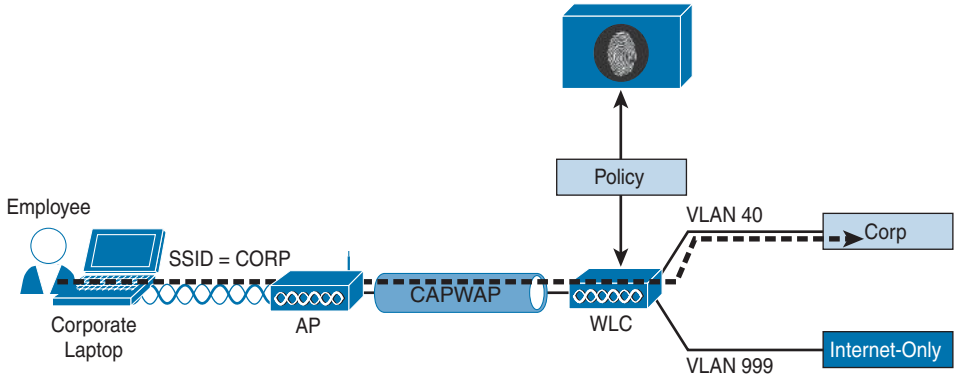


Figure 14-1 *Employee Using a Corporate Laptop to Gain Full Access*

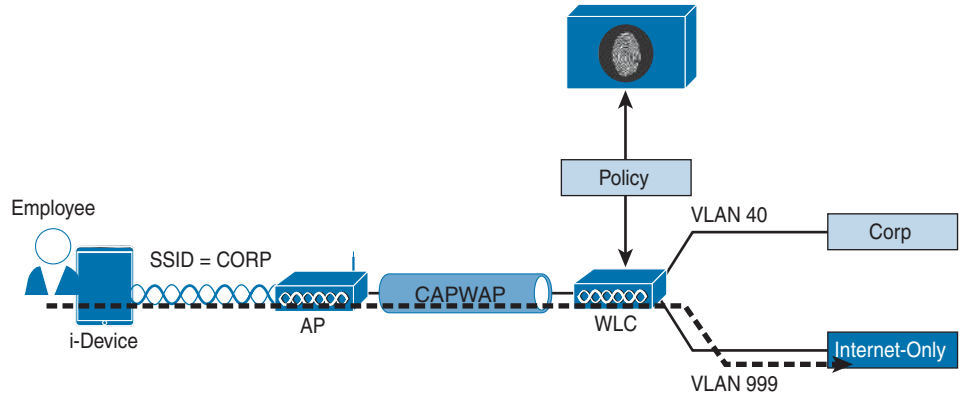


Figure 14-2 *The Same Employee Credentials on an iOS Device Get Limited Access*

Users using the same wireless SSID and the same credentials can be associated to different wired VLAN interfaces, based on the device profile:

- Employees using corporate laptops with their Active Directory usernames are assigned to the corporate VLAN and given full access to the network.
- Employees using iOS devices with their same Active Directory usernames are assigned to a guest VLAN and provided Internet access only.

While it might be quite intuitive to visualize the types of access policies you will create based on the device's profile, the design of where and how ISE Profiler collects the data about the endpoints requires some thought and planning.

Anomalous Behaviour

One of the first questions a security team may ask when beginning to use profiling with any network access control solutions is "Can we use this as an anti-spoofing solution?" Remember that MAC Authentication Bypass is really a very limited replacement for a strong authentication. It would be fairly easy for a malicious user to unplug a printer from the wall, configure a laptop to use the same MAC address as the printer (spoofing), and gain access to the network. To help mitigate this potential issue, ISE added two new features starting in ISE Version 2.2: ***Anomalous Behaviour Detection*** and Anomalous Behaviour Enforcement.

It is important to keep in mind that profiling (without the anomalous behaviour features) is a technology that compares collected attributes about an endpoint to a set of signatures called *profiling policies* to make a best guess about what a device is. Can profiling be used to prevent spoofing? Maybe. However, it is often very difficult to accomplish anti-spoofing with profiling alone. It would require a lot of tuning, trial and error, and constant adjustment, which makes it too operationally expensive and untenable.

The purpose of Anomalous Behaviour Detection is to reduce the operational expense in detecting potential MAC spoofing. With Anomalous Behaviour Detection enabled, ISE checks profiling data and looks for contradictions or certain changes. In order to accomplish this, Anomalous Behaviour Detection checks the following:

- **Port type:** The port type determines an endpoint's method of access. For example, if a MAC address was originally used on a wired connection, that same MAC address should not be connecting to the wireless network. Even on the same computer, the

wireless NIC should have a separate MAC address. This would be a clear indication that a MAC address is being spoofed.

- **DHCP Class-Identifier**: This profiling attribute can help determine the type of client or vendor for an endpoint. If this changes for a detected endpoint, it could be an indication of a spoofed MAC address. For example, if an endpoint is configured with a static IP address, there should be no DHCP Class-Identifier attribute populated in ISE. If that MAC address is later spoofed and new endpoints used DHCP, ISE would detect that change as a potential anomalous behavior. Another example that would take it a step further is if a previously profiled endpoint used DHCP and the Class-Identifier was originally Cisco-Telepresence. If ISE later received the DHCP Class-Identifier MSFT 5.0 for that endpoint, this could also be an indication that the MAC address has been spoofed.
- **Endpoint policy**: The profile policy for an endpoint is determined when ISE compares the collected profile attributes to the defined profiles and finds a match. The endpoint is assigned the endpoint policy. An example of anomalous behavior would be an endpoint being assigned the endpoint policy Cisco-IP-Phone-8831 that later changes to Windows7-Workstation for that MAC address. This kind of change might indicate that someone is spoofing the MAC address of the endpoint.

The Anomalous Behaviour Detection feature does not automatically change the network access of an endpoint when detected. If an endpoint sees certain profile data change, it sets the AnomalousBehaviour attribute for that endpoint to true. In order to change the network access based on Anomalous Behaviour Detection, Anomalous Behaviour Enforcement must be enabled. This allows ISE to issue a Change of Authorization (CoA) upon the detection of the anomalous behavior. The suspicious endpoints can then be reauthorized based on a configured authorization policy.

Key Topic

When it comes to profiling, it is still considered best practice to use a least-privilege strategy, regardless of whether the anomalous behavior features are enabled. If a malicious actor is successful in spoofing the MAC address of a printer and gains network access, what level of network access do you want that printer to have? The authorization policy for printers should not provide full network access but instead should provide a very limited subset of access. For example, a printer should only be permitted to communicate using network ports critical to printer operations (such as TCP port 9100 or port 9600) and should only communicate directly with a print server.

Minimal configuration is required on the ISE side to enable Anomalous Behaviour Detection and Anomalous Behaviour Enforcement:

Step 1. In the ISE GUI, navigate to **Administration > System > Settings > Profiling**.

Step 2. From the CoA Type dropdown, choose **Reauth**.

Step 3. Check the **Enable** checkbox next to Enable Anomalous Behaviour Detection. The AnomalousBehaviour attribute is then set on endpoints, but no Change of Authorization (CoA) is sent when ISE detects anomalous behavior.

Step 4. Check the **Enable** checkbox next to Enable Anomalous Behaviour Enforcement. ISE now allows you to create an authorization policy based on the Anomalous-Behaviour flag and sends a CoA when anomalous behavior is detected. Figure 14-3 shows the completed Profiler Configuration page.

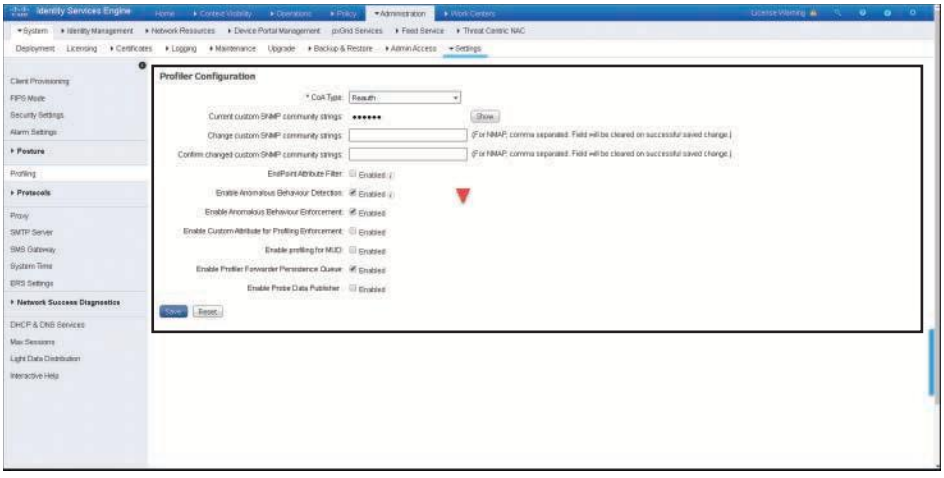


Figure 14-3 *Completed Profiler Configuration Page*

Step 5. Navigate to **Policy > Policy Sets** and expand any active policy set.

Step 6. In the ISE GUI, navigate to **Administration > System > Settings > Profiling**.

Step 7. Expand the **Authorization Policy - Local Exceptions** policy.

NOTE There are two special exception authorization policies in every policy set. The local exceptions authorization policy is for authentications that are being evaluated on the specific policy set. The global exceptions authorization policy is evaluated globally. This means that if an endpoint that is evaluated by any policy set on this ISE deployment, it is also evaluated against any rules in the global exception authorization policy, regardless of which policy set the rules were created in.

Step 8. Click the + button to create a new rule.

Step 9. Give the rule a name.

Step 10. Click the + button to open the Conditions Studio.

Step 11. Choose **Endpoints:AnomalousBehaviour** as the attribute.

Step 12. Select **Equals** as the operator.

Step 13. Type in **true** as the attribute value.

Step 14. Select an appropriate authorization result.

Step 15. Click **Save** to save the policy.

Now when anomalous behavior is detected, ISE reauthenticates the endpoint and denies it access. Figure 14-4 shows the completed authorization rule.

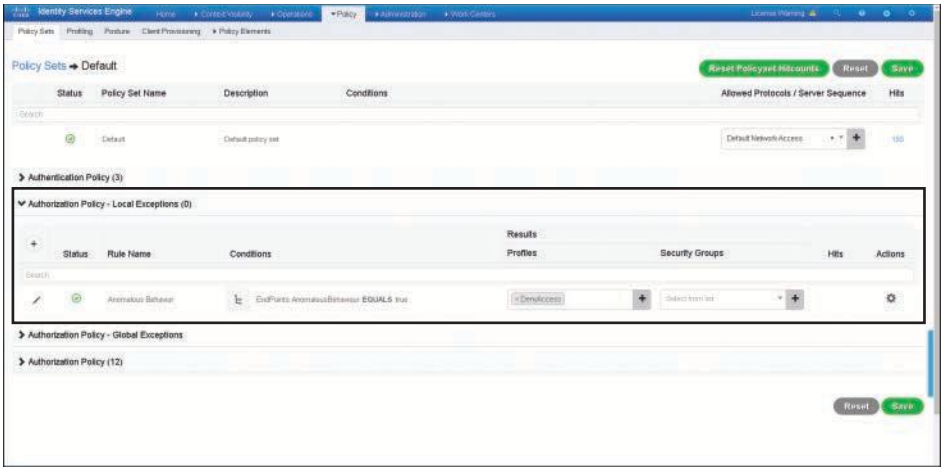


Figure 14-4 *Local Exceptions Authorization Rule for Anomalous Behaviour Enforcement*

Cisco ISE Probes

Cisco ISE is capable of providing access policies where the decisions may be made based on who, what, where, when, how, and more. Profiling is focused on the *what* elements of the policy. In order for the policy engine to know what a device is, you must first collect data.

Cisco ISE uses a number of collection mechanisms known as *probes*. A probe is software designed to collect data to be used in a profiling decision. For example, the HTTP probe captures HTTP traffic and makes it possible for Profiler to examine attributes from the traffic, such as HTTP User-Agent strings.

Probe Configuration

You enable probes on each Policy Services node (PSN), where appropriate. In the Policy Administration node (PAN) GUI, navigate to Administration > System > Deployment. From here, you can select the PSN for which you are configuring probes. Follow these steps for each PSN in your deployment:

Step 1. Select one of the PSNs, as shown in Figure 14-5. In this case, the node is still standalone, which means it is a single node running all personas (Administration, Monitoring, Policy Services, and pxGrid).



Figure 14-5 *ISE Deployment Screen*

Step 2. Ensure that the **Enable Profiling Service** checkbox is selected. This service is enabled by default on all PSNs and is not configurable when in standalone mode, as shown in Figure 14-6.

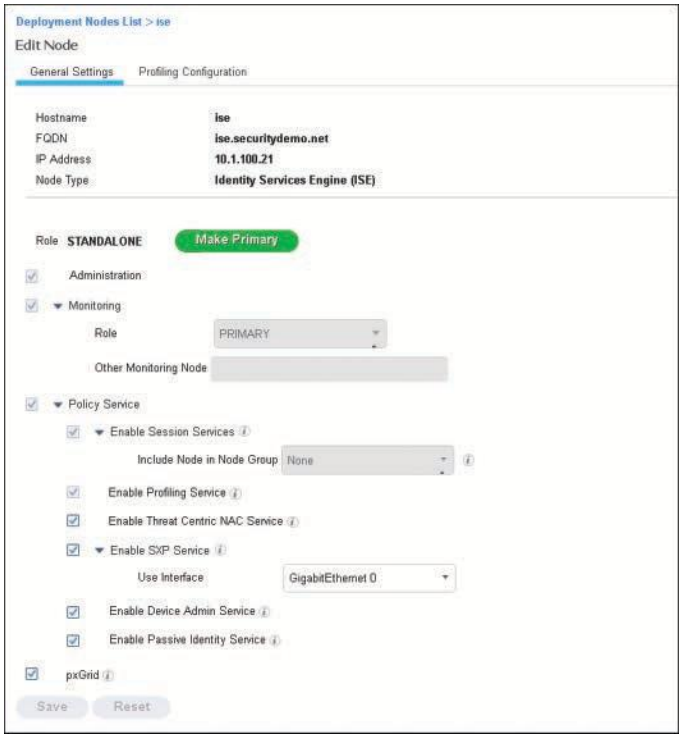


Figure 14-6 *General Settings*

Step 3. Select the **Profiling Configuration** tab, as shown in Figure 14-7.



Figure 14-7 *Profiling Configuration*

As you can see in Figure 14-7, 11 probes can be enabled on each PSN:

- NETFLOW
- DHCP
- DHCPSPAN
- HTTP
- RADIUS
- Network Scan (NMAP)
- DNS
- SNMPQUERY
- SNMPTRAP
- Active Directory
- pxGrid

The following sections examine these probes.

DHCP and DHCPSPAN

DHCP can be one of the most useful data sources for an endpoint device. A primary use of DHCP in profiling has been to capture the device MAC address, but it can also be used to capture other probes as well. The true value in the DHCP probe is that many other attributes can be gleaned to help identify the type of endpoint. As HTTP carries User-Agent field, DHCP requests carry a Class-Identifier field that helps identify the operating system of a device. Some organizations have been known to use a custom DHCP Class-Identifier string to help identify a device as a corporate asset.

The following are some of the most useful DHCP attributes for classifying devices:

- **Class-Identifier**: This optional string may include information about the operating system, client type, client configuration, or other identifying information about an endpoint.
- **Client-FQDN**: This optional field allows the client to convey its FQDN to the DHCP server so it can update the IP address–to–FQDN mapping.
- **Host-Name**: This field may be populated with the hostname of the client. A client may not offer an FQDN, but this field might provide more information about the endpoint.
- **Domain-Name**: This option specifies the domain name the client should use when resolving hostnames via DNS.
- **Parameter-Request-List**: A client can use this option to request values for specified configuration parameters. The client may list the options in an order of preference. Typically, certain operating systems, such as macOS and Windows, have parameter request lists that are unique to the version of the operating system. This attribute can be very useful for identifying the difference between a Windows XP client and a Windows 10 client or an OS X client and an iOS device.

There are two DHCP probes and they work in slightly different ways: DHCP and DHCPSPAN.

DHCP

The DHCP probe requires that DHCP requests be sent directly to the ISE PSN(s). This is often done by using the **ip helper-address** interface configuration command (see Figure 14-8).

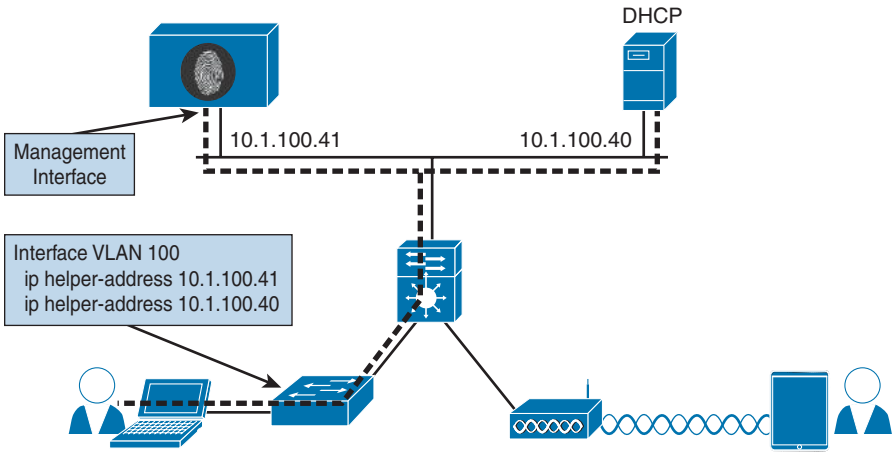


Figure 14-8 *DHCP with* ***ip helper-address*** *Logical Design*

The **ip helper-address** command on a Layer 3 interface converts a DHCP broadcast (which is a Layer 2 broadcast) to a unicast or directed broadcast (which involves sending the broadcast to all hosts on a specific subnet). You can simply add the IP addresses of your PSNs to the list of helper addresses, and the list will be copied on all DHCP requests. Adding another helper address will not delay the return of DHCP requests because the network access device simply sends a copy of the DHCP request to all the helper addresses that are configured at the same time and waits for the first response. By default, ISE does not respond to that DHCP request and only uses the incoming data to profile endpoints. The only exception where ISE might respond to a DHCP request is when DHCP and DNS services are configured on ISE so that you can provide URL redirection on third-party NADs but only for specific subnets that you configure. DHCP and DNS services are covered in more detail in Chapter 18.

DHCPSPAN

Another way for ISE to glean DHCP requests and even DHCP responses is through the use of a Switched Port Analyzer (SPAN) session in true promiscuous mode. A SPAN session copies all traffic between a source interface and the destination interface on a switch, which would be one of the ISE interfaces assigned to the DHCPSPAN probe. A SPAN logical design is illustrated in Figure 14-9.

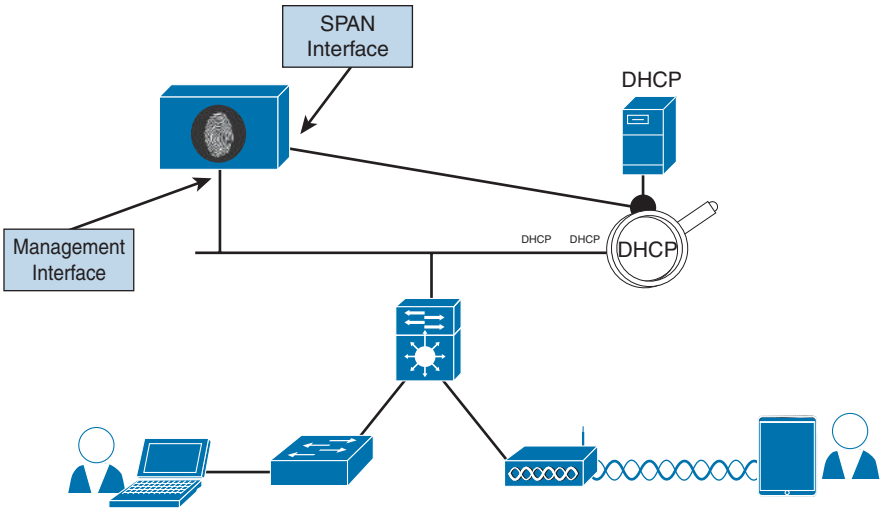


Figure 14-9 *DHCP SPAN Logical Design*

When using the SPAN method, you need to determine the best location for creating the SPAN session and gather the data. One recommended location is the DHCP server, where the DHCP probe can see both ends of the conversation (request and response). However, there are caveats to this placement. For example, if an organization uses distributed DHCP servers, this needs to be factored into the design. For most ISE deployments, the non-SPAN method is preferred and most commonly deployed.

Considerations with Cisco WLCs

By default, a wireless LAN controller (WLC) is configured to act as a RADIUS proxy, which is the WLC's own native way of having an IP helper address. The WLC is involved in all DHCP transactions, sending them to the DHCP server. Unfortunately, this WLC native method is not compatible with the DHCP probe, and therefore, the DHCP proxy must be disabled on the WLC. Then, the DHCP requests from wireless endpoints appear as broadcast messages on the VLAN, and the **ip helper-address** command configured on the Layer 3 interface on the switch or router sends the DHCP request to ISE. This offloads the IP helper functionality to another network access device instead of the WLC. However, this method is not the only option available for sending DHCP information from the WLC. Cisco WLCs can also send DHCP information to ISE by using Device Sensor, as you will see later in this chapter.

Probe Configuration

Minimal configuration is required on the ISE side to enable DHCP probes. In the Profiling Configuration tab you saw earlier in this chapter, in Figure 14-7, follow these steps:

Step 1. Click the checkbox next to **DHCP** and **DHCPSPAN** to enable these probes.

Step 2. Select either a specific interface or all interfaces. (You cannot select multiple interfaces individually. Your choices are a single interface or all interfaces.)

Step 3. (Option) Customize the UDP port on which the DHCP requests will be received. Unless there is a requirement to change the port, it is recommended to keep the default.

Figure 14-10 shows the DHCP probes enabled. There should never be a need to enable both probes for the same interface, however, as doing so would cause double processing of DHCP packets and would be somewhat wasteful of system resources.

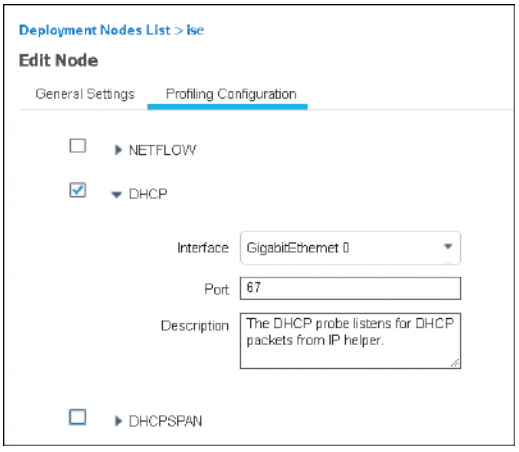


Figure 14-10 *DHCP Probes*

NOTE If you are using only Device Sensor–capable infrastructure, neither DHCP probe needs to be enabled.

RADIUS

RADIUS is the primary communication mechanism from a network access device (NAD) to the authentication server (ISE). RADIUS communication provides very useful data that can help classify a device.

Originally, the RADIUS probe was focused on the MAC address and IP address of a device. By having this data conveyed in the RADIUS packet, ISE is able to build the all-important MAC-to-IP address bindings. Because the endpoint database uses MAC addresses as unique identifiers for all endpoints, these bindings are absolutely critical. Without them, Layer 3 probes such as HTTP and Nmap scanning would not work correctly.

Some of the RADIUS attributes that are most useful in profiling or creating access policies include the following:

- **Calling-Station-Id:** This attribute provides the endpoint's MAC address.
- **Framed-IP-Address:** This attribute provides the IP address of the endpoint.
- **Called-Station-Id:** While it is not very important for profiling, this RADIUS attribute is very important with endpoints that are connecting on a wireless SSID. Typically, this attribute stores the bridge's or access point's MAC address. Appended to that MAC address is the SSID. If you wish to create a rule in an access control policy that is specific to an SSID, you can use the Called-Station-Id attribute to identify authentications on that SSID.

- **NAS-Port-Id:** This RADIUS attribute identifies the physical port of the NAD that the endpoint is connecting to.
- **NAS-Port-Type:** This attribute indicates the type of port that is being used to authenticate the endpoint.
- **Service-Type:** This attribute indicates the type of service requested or provided. This can tell you if the authentication is MAB or 802.1X. For MAB, Cisco switches send the Call-Check value for the Service-Type attribute. For 802.1X, the Service-Type is set to Framed.

In addition, the RADIUS probe triggers the SNMPQUERY probe to poll the NAD (as discussed later in this chapter, in the section "SNMPQUERY and SNMPTRAP").

Key Topic

As Device Sensor–capable switches and wireless controllers have proliferated, the RADIUS probe has become especially critical. Device Sensor is a feature in a switch or controller that collects endpoint attributes locally and then sends those attributes to ISE within the RADIUS accounting packets.

By allowing a network device to proactively send profiling data to ISE, the architecture has placed the collection agents as close to the endpoint as possible, at the point of access to the network. This has eliminated the need to send the IP helper address information to ISE as well as the need to reactively query the switches for CDP/LLDP information using the SNMPQUERY probe.

Considerations with the RADIUS Probe

All NADs in a secure unified access deployment should be configured to send RADIUS accounting packets. It is important to note that either a Cisco switch must learn an endpoint's IP address via DHCP snooping or device tracking must be enabled in order to fill in the Framed-IP-Address field.

Probe Configuration

Minimal configuration is required on the ISE side to enable RADIUS probes. In the Profiling Configuration tab, just click the checkbox next to RADIUS to enable the RADIUS probe, as shown in Figure 14-11.

Notice that there is not really any configuration possible with this probe in ISE. Even so, it is one of the most useful probes, especially when combined with Device Sensor.

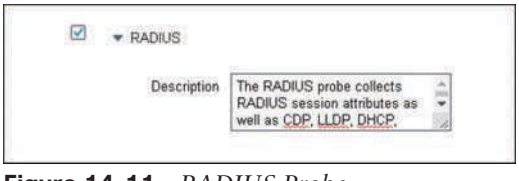


Figure 14-11 *RADIUS Probe*

Network Scan (Nmap)

Nmap is a tool that uses port scans, SNMP, SMB, and other mechanisms to identify a device's operating system, computer name, domain name, or other device attributes. The NMAP probe can be run against a single IP address or a whole subnet.

Importantly, the Profiler engine can be configured to react to a profiling event with a reactive NMAP probe. For example, say that an endpoint is discovered to be an Apple device, and ISE automatically launches an Nmap OS scan against that endpoint to determine if it is running macOS or iOS. Based on the results of that scan, ISE can further classify the device as a Mac or an iOS device.

There are several different options within ISE's Nmap utility:

- **OS:** This scan option enables operating system detection using TCP/IP fingerprinting.
- **SNMP Ports:** This option checks to see if the SNMP port is open on the endpoint. Useful attributes such as the following can be gleaned from a successful SNMP scan:
  - **hrDeviceDescr:** This description of a device may optionally include the device's manufacturer, revision, and serial number.
  - **sysDescr:** This description of a device may optionally include the full name and version identification of the system's hardware type, software operating system, and networking software.
  - **sysContact:** This optional attribute identifies the contact person for this managed device. It might also include contact information for the person.
  - **sysLocation:** This is an optional description of the physical location of a device.
  - **sysName:** This is the administratively assigned name for this managed device; it is often the device's fully qualified domain name (FQDN).
- **Common Ports:** This option scans a selection of well-known common ports.
- **Custom Ports:** This option enables an administrator to specify custom UDP and TCP ports.
- **Include Service Version Information:** A vendor may provide a detailed description of the device in the version information that is found when different services (port/protocols) are queried in a scan. Enabling this option improves the efficiency of scans of custom or common ports.
- **Run SMB Discovery Script:** When this option is selected, Nmap uses the SMB protocol on ports 445 and 139 to attempt to determine the operating system, computer name, domain name, forest name, FQDN, NetBIOS computer name, NetBIOS domain name, workgroup, system time, and time zone of a scanned device.
- **Skip NMAP Host Discovery:** Nmap host discovery causes Nmap to probe for active endpoints before performing heavier scanning. This option skips the initial host discovery state of a manual Nmap scan.

Considerations with the NMAP Probe

The NMAP probe is executed against an IP address or range of IP addresses. However, it is absolutely crucial to keep in mind that the endpoint database uses a MAC address as the unique identifier of any endpoint. The PSN relies on the MAC-to-IP-address binding to update an endpoint's attributes with the results of an Nmap scan. Therefore, it is critical for the PSN to receive valid information from the other probes in order to run a scan.

The NMAP probe can be manually run against a single IP address or subnet. The most common option is for an Nmap scan to be triggered as an action of a profile event. It should be noted that for a manual NMAP scan to report the results, the endpoint must already exist in the endpoint database in ISE. If the endpoint does not exist in the database or detected on the network by ISE, no result will be provided.

Probe Configuration

As with all the other probes, only minimal configuration is needed for the NMAP probe in the ISE GUI. From the Profiling Configuration tab, just click the checkbox next to Network Scan (NMAP) to enable this probe. You may enable this probe as shown in Figure 14-12.

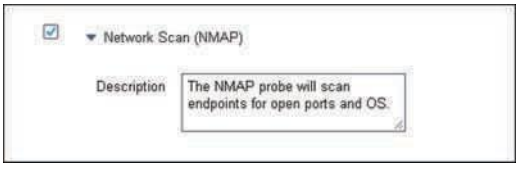


Figure 14-12 *Network Scan (NMAP) Probe*

DNS

The DNS probe is used to collect the FQDN of an endpoint using a reverse lookup for the static or dynamic DNS registration of that endpoint. This probe is quite useful when looking for a specific DNS name format among corporate assets (Active Directory members).

A reverse DNS lookup is completed only when an endpoint is detected by at least one of the following probes: DHCP, RADIUS, HTTP, or SNMP.

To enable the DNS probe, simply click the checkbox next to DNS in the Profiling Configuration tab, as shown in Figure 14-13. This probe uses the name server that is configured on the ISE node.

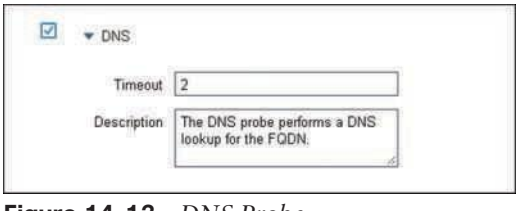


Figure 14-13 *DNS Probe*

SNMPQUERY and SNMPTRAP

SNMP is used to query NADs that do not support Cisco's Device Sensor. After you enable the SNMPQUERY probe, ISE polls all the SNMP-enabled NADs at the configured polling interval.

NOTE It is recommended to remove SNMP settings from IOS Device Sensor–enabled NADs to avoid double work and wasted processing.

There are two SNMP probes: SNMPQUERY and SNMPTRAP.

SNMPTRAP

The SNMPTRAP probe receives from the configured NAD(s) information that supports MAC notification, linkup, linkdown, and informs. The purpose of this probe is twofold: It is used to trigger the SNMPQUERY probe, and it is used as a toggle switch to allow the SNMPQUERY probe to reactively query a NAD instead of waiting for the periodic polling interval. Therefore, you must also enable the SNMPQUERY probe in order for SNMPTRAP to be functional.

The SNMPTRAP probe receives information from a specific NAD(s) when the MAC address table changes or when the link state changes on a switch port. In order to make this feature functional, you must configure the NAD to send SNMP traps or inform.

Example 14-1 shows a sample SNMPTRAPS configuration on a Cisco switch.

Example 14-1 *Configuring SNMPTRAPS Probe*

```
CAT9300(config)#interface g1/0/1
CAT9300(config-if)#snmp trap mac-notification change added
CAT9300(config-if)#snmp trap mac-notification change removed
CAT9300(config-if)#exit
CAT9300(config)#mac address-table notification change
CAT9300(config)#mac address-table notification mac-move
CAT9300(config)#snmp-server trap-source vlan 100
CAT9300(config)#snmp-server enable traps snmp linkdown linkup
CAT9300(config)#snmp-server enable traps mac-notification change move
CAT9300(config)#snmp-server host <ISE-PSN>version 2c <community-string>
```

SNMPQUERY

SNMPQUERY does the bulk of the profiling work by providing the most information about the endpoint. There are actually three different kinds of SNMPQUERY probes:

- **System:** This probe polls all the NADs that are configured with SNMP at the configured intervals.
- **Interface:** This probe fires in response to an SNMPTRAP or RADIUS accounting START packet (only if the SNMPTRAP probe is enabled).
- **Network Scan (NMAP):** This probe triggers the SNMP walk of the endpoint.

When querying a NAD, ISE looks for interface data (such as which interface or which VLAN), session data (if the interface is Ethernet), and Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) data (if these protocols are enabled). CDP and LLDP data can be very useful in identifying a device type by its registered capabilities and similar attributes. For Cisco switches, CDP should be enabled by default. However, LLDP needs to be enabled globally and on the interfaces of the switch.

NOTE For distributed deployments, NAD polling is distributed among all PSNs enabled for SNMPQUERY probes.

The following is an example of an SNMP query configuration on a Cisco switch:

```
CAT9300(config)#snmp-server community <string> RO
```

SNMPQUERY and SNMPTRAP Probe Configuration

Where there is a bit of configuration possible for the SNMPQUERY and SNMPTRAP probes, such as the trap types to examine and the SNMP port, it is recommended to leave these settings at their defaults unless directed otherwise by Cisco TAC. Follow these steps to configure the SNMPQUERY and SNMPTRAP probes:

Step 1. Click the checkboxes next to the **SNMPQUERY** and **SNMPTRAP** to enable these probes.

Step 2. For the SNMPTRAP probe, select either the Gigabit 0 interface or all interfaces. (You cannot select multiple interfaces individually. Your choices are a single interface or all interfaces.)

Figure 14-14 shows the SNMP probes enabled and the default settings for these probes.

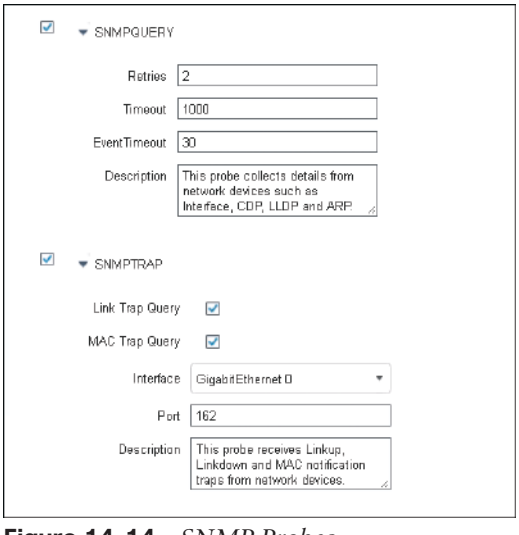


Figure 14-14 *SNMP Probes*

NETFLOW

NetFlow is an incredibly useful and undervalued security tool. It is similar to a phone bill as it provides a summary of all calls sent and received (rather than a recording of conversations in their entirety).

Cisco routers and switches support NetFlow, which sends a record of each flow that is seen by the device. This record includes the ports, source IP addresses, destination IP addresses, and other very usable information.

However, just enabling NetFlow in your infrastructure and forwarding all its information to ISE can quickly overwhelm the PSN. If you are planning to use the NETFLOW probe, it is highly recommended that you have a NetFlow collector, such as Stealthwatch, to filter out

any unnecessary data and send only what is truly needed to ISE. The NETFLOW probe is widely used for profiling in many ISE deployments. The Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam does not focus heavily on the NETFLOW probe, so neither does this book. It is recommended to perform extensive planning prior to using this probe.

Configuring the NETFLOW probe is limited to enabling the checkbox next to NETFLOW in the Profiling Configuration tab and selecting either the Gigabit 0 interface or all interfaces. Figure 14-15 shows the enabled NETFLOW probe.

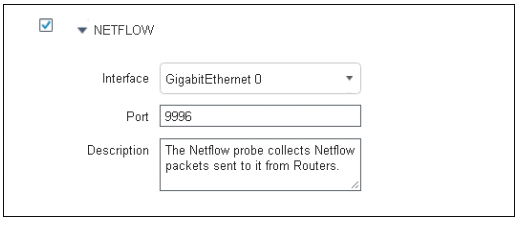


Figure 14-15 *NETFLOW Probe*

HTTP Probe

When an application such as a web browser uses HTTP, the application typically identifies itself with its application type, operating system, software vendor, and software revision by submitting an identification string to its operating peer. This information is transmitted in an HTTP request-header field called the User-Agent field.

ISE uses the information in HTTP packets for profiling; it uses the User-Agent field to help match signatures to determine what profile a device belongs in.

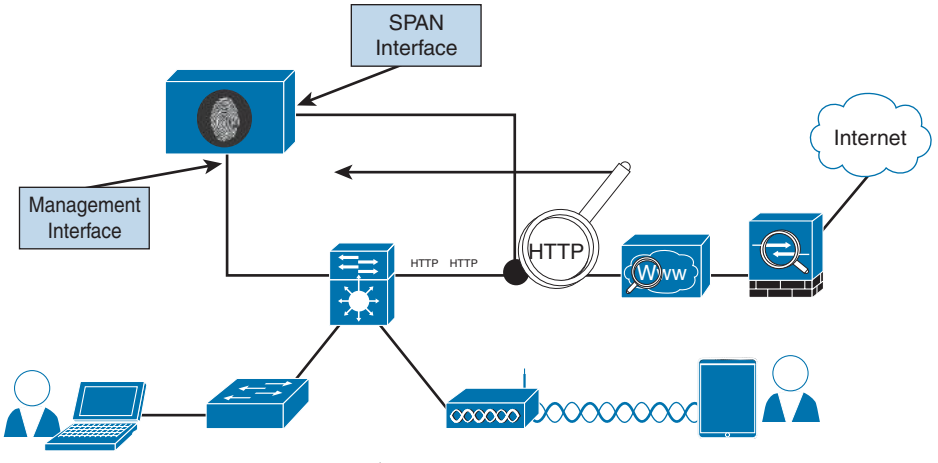


Figure 14-16 *HTTP SPAN Logical Design*

There are two primary mechanisms by which the HTTP probe can collect the HTTP traffic:

- **Using a SPAN session in true promiscuous mode or in conjunction with a filter to limit traffic:** When using the SPAN method, you need to determine the best place

to create the SPAN session and gather the data. One recommended location is the Internet edge, where a network organization would typically deploy a web security appliance such as the Cisco IronPort WSA. Figure 14-16 displays the HTTP SPAN design with ISE.

- **Using a SPAN session in conjunction with a filter to limit the traffic that is visible to ISE:** Another option with a SPAN design is to use VLAN ACLs (VACLs) on a Catalyst 6500 or ACL-based SPAN sessions on a Nexus 7000. These options allow you to build an ACL that defines exactly what traffic you want to capture and send along to ISE; this is different from a pure promiscuous SPAN, in which the ISE interface sees all traffic. This is a better way to manage the resource utilization on your ISE server than sending all unfiltered traffic to the interface.

As you can see, there are multiple ways to use the HTTP probe, and you should consider what works best for your environment and then deploy using that approach. In many environments, it is best to not use SPAN at all but to leverage ISE's own portals to capture the User-Agent strings.

To configure the HTTP probe, click the checkbox next to HTTP in the Profiling Configuration tab, as shown in Figure 14-17. Select either the Gigabit 0 interface or all interfaces.

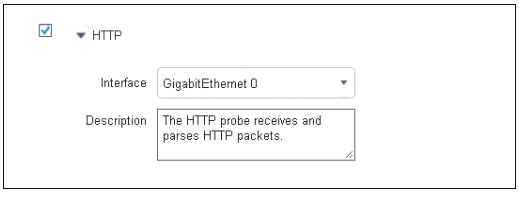


Figure 14-17 *HTTP Probe*

HTTP Profiling Without Probes

The web portal system in ISE has been outfitted to collect the user agent details from the web browser that is communicating with an ISE portal. This occurs regardless of whether HTTP profiling is enabled. The User-Agent field is used to determine which operating system is connecting and, therefore, which agent or client to send to the endpoint with client and native supplicant provisioning.

When any portal collects that user agent, it is automatically passed over to the profiling engine in ISE, without requiring the HTTP probe to be enabled. This is a simple and efficient way to get the extremely valuable User-Agent string without having to rely on the computationally expensive SPAN methods.

HTTP Profiling with the RADIUS Probe

For certain types of Device Sensor–enabled devices, HTTP information is collected and sent via the RADIUS probe. For example, the Cisco Wireless LAN Controller (WLC) Device Sensor is capable of passively collecting HTTP information without the use of a portal. Even without the HTTP probe enabled, the HTTP User-Agent string is collected as part of the RADIUS probe and used for profiling.

Active Directory Probe

The Active Directory probe was added in ISE Version 2.1 to increase the fidelity of operating system information for domain-joined endpoints. This probe can help distinguish between corporate and non-corporate assets and can gather a wealth of information from Active Directory. Microsoft Active Directory tracks detailed operating system information about domain-joined computers such as the version information and the service pack levels; this information can be used as profiling attributes within ISE. As a prerequisite, ISE must be joined to the domains that it will be querying.

The Active Directory probe provides the following specific attributes to ISE:

- **AD-Host-Exists**: This attribute simply states whether the computer exists in Active Directory. If the value of this attribute is true, then this is a domain-joined endpoint.
- **AD-Join-Point**: This attribute defines the Active Directory domain where the endpoint is located.
- **AD-Operating-System**: This attribute indicates the operating system version of the endpoint.
- **AD-OS-Version**: This attribute indicates the version of the endpoint's operating system.
- **AD-Service-Pack**: This attribute indicates the service pack version of the endpoint.

For ISE to query Active Directory for information about an endpoint, ISE must learn the hostname of the endpoint. It can learn this information through the DHCP or DNS probes. After ISE learns the hostname, Active Directory is queried. After the initial query, ISE does not re-query Active Directory for that specific endpoint unless additional profiling activity is detected for the endpoint or the rescan timer expires.

Minimal configuration is required on the ISE side to enable the Active Directory probe in ISE. Follow these steps in the Profiling Configuration tab:

Step 1. Click the checkbox next to the **Active Directory** to enable this probe.

Step 2. (Optional) Change the rescan interval. The default is 1 day before rescan, and it is recommended to keep this default setting.

Figure 14-18 shows the Active Directory probe enabled.



Figure 14-18 *Active Directory Probe*

pxGrid Probe

The Cisco Platform Exchange Grid (pxGrid) is a multivendor, cross-platform network system that allows the sharing of device and user context across a wide range of network and security solutions, such as security monitoring and detection systems, network policy platforms, asset and configuration management, and identity and access management platforms. pxGrid makes it possible to share and receive context from other Cisco and third-party systems. This probe uses that function for receiving endpoint context from external sources and using it for profiling.

The pxGrid probe is based on the pxGrid Version 2 specification, using the Endpoint Asset topic. The available topic attributes are as follows:

- **assetId:** Asset ID
- **assetName:** Asset name
- **assetIpAddress:** IP address
- **assetMacAddress:** MAC address
- **assetVendor:** Manufacturer
- **assetProductId:** Product code
- **assetSerialNumber:** Serial number
- **assetDeviceType:** Device type
- **assetSwRevision:** Software revision number
- **assetHwRevision:** Hardware revision number
- **assetProtocol:** Protocol
- **assetConnectedLinks:** Array of network link objects
- **assetCustomAttributes:** Array of custom names-value pairs

pxGrid is covered more in detail in Chapter 22, "ISE Context Sharing and Remediation."

Minimal configuration is required on the ISE side to enable the pxGrid probe in ISE. In the Profiling Configuration tab, check the box next to **pxGrid** to enable this probe. There are no additional settings to configure on this page.

Figure 14-19 shows the pxGrid probe enabled.

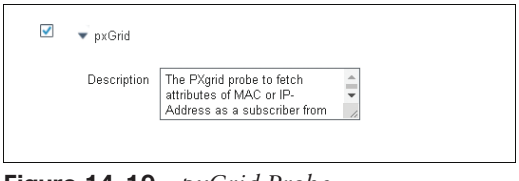


Figure 14-19 *pxGrid Probe*

Infrastructure Configuration

As an overall best practice, it is recommended to do a cost/benefit analysis of the use of processor-intensive probes or probe designs. For example, it is often recommended to use DHCP helper addresses instead of configuring a SPAN session and examining a multitude of traffic that may or may not be relevant.

For example, HTTP traffic is extremely useful for identifying the operating system on a client endpoint. However, an HTTP SPAN probe can consume a large amount of system resources on the PSN. In addition, it may not be critical to have full visibility into the User-Agent strings of all devices, such as corporate-managed Windows devices.

Some deployments make use of VLAN ACL (VACL) captures that can limit what traffic is sent to the SPAN session. Other deployments use the authorization policy in ISE to send unknown devices to an ISE portal, which allows the portal to update ISE's profiling data.

DHCP Helper

As shown in Figure 14-8, the **ip helper-address** commands are configured on the default gateway for each of the access-layer VLANs. To configure destination addresses for a NAD to send a copy of DHCP requests to, you enter interface configuration mode of the Layer 3 interface for the VLAN or subnet and enter the **ip helper-address** *[ip-address]* command. Then you can repeat this process until you have added the DHCP server(s) and all applicable ISE PSNs to the list of helper addresses on that interface. Generally, it's best practice to add no more than two PSN IP addresses as help addresses on a Layer 3 interface.

Example 14-2 shows an example of the output from a Layer 3 interface that is configured to send DHCP requests to the DHCP server in addition to two different ISE PSNs.

Example 14-2 *ip helper-address Settings*

```
C6K-DIST#sho run int VLAN41
Building configuration…

Current configuration : 152 bytes
!
interface Vlan41
  ip address 10.1.41.1 255.255.255.0
! DHCP Server:
  ip helper-address 10.1.100.100
! Policy Services Nodes:  ip helper-address 10.1.100.3
  ip helper-address 10.1.100.4
end

C6K-DIST#
```

SPAN Configuration

A monitor session is configured in the global configuration mode of a switch or router. It can be configured as a local (SPAN) or remote (RSPAN) session. Example 14-3 shows a SPAN configuration in which an Internet-facing VLAN is the source of the session and an interface on the PSN is the destination.

NOTE For more on SPAN configuration, see http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a008015c612.shtml.

Example 14-3 *monitor session Command Input*

```
DC-4948(config)#monitor session [1-4] source [interface | vlan] [rx | tx]
DC-4948(config)#monitor session [1-4] destination interface [interface_name]
```


Example 14-4 shows the output of the **show monitor** command. In this output, you can see the source and destination of the session.

Example 14-4 *Example monitor session (SPAN) Configuration*

```
DC-4948#show monitor session 1
Session 1
Type                                                : Local Session
Source VLANs                                  :
! Source is VLAN 100:
  Both                                              : 100
! Destination is Interface Gig 1/47:
Destination Ports                               : Gi1/47
  Encapsulation                                  : Native
             Ingress                                    : Disabled
             Learning                                 : Disabled
Filter Pkt Type                                  :
   RX Only                                          : Good


DC-4948#
```

VLAN Access Control Lists (VACLs)

VACL configuration is a multistep process. It requires an access control list to specify which traffic should be captured and another access control list to specify which traffic should be ignored. An access map must be configured to apply those ACLs to the traffic. The access map gets assigned to the VLAN(s), and then you define the destination. To configure a VACL to filter HTTP traffic to the SPAN session, follow these steps:

Step 1. Build an ACL to classify the traffic you want to capture, as shown in the following example:

```
C9K-DIST(config)#ip access-list extended HTTP_TRAFFIC
C9K-DIST(config-ext-nacl)#permit tcp any any eq www
```

Step 2. Build an ACL for all the rest of the traffic, as shown in the following example:

```
C9K-DIST(config)#ip access-list extended ALL_TRAFFIC
C9K-DIST(config-ext-nacl)#permit ip any any
```

Step 3. Create a VLAN access map sequence to "capture" HTTP traffic, as shown in the following example:

```
C9K-DIST(config)#vlan access-map HTTP_MAP 10
C9K-DIST(config-access-map)#match ip address HTTP_TRAFFIC
C9K-DIST(config-access-map)#action forward capture
```

Step 4. Add a new sequence to the access map to forward all other traffic, as shown in the following example:

```
C9K-DIST(config)#vlan access-map HTTP_MAP 20
C9K-DIST(config-access-map)#match ip address ALL_TRAFFIC
C9K-DIST(config-access-map)#action forward
```

Step 5. Apply the VLAN access map to the VLAN list as follows:

```
C9K-DIST(config)#vlan filter HTTP_MAP vlan-list 41,42
```

Step 6. Configure the destination port for the PSN's SPAN interface as follows:

```
C9K-DIST(config)#switchport capture allowed vlan 41
C9K-DIST(config)#switchport capture allowed vlan add 42
C9K-DIST(config)#switchport capture
```

Device Sensor

IOS Device Sensor requires a multipart configuration. The first part is to configure the Device Sensor filter lists, which inform Device Sensor of which items to consider important for each of the different protocols.

IOS Device Sensor support three protocols: DHCP, CDP, and LLDP. You must create one list for each protocol, as described in the following steps:

Step 1. Create a list for DHCP. The options you should configure for the DHCP list include the hostname, class identifier, parameter request list, and client identifier, as demonstrated here:

```
CAT9300(config)#device-sensor filter-list dhcp list
<dhcp_list_name>
CAT9300(config-sensor-dhcplist)#option name host-name
CAT9300(config-sensor-dhcplist)#option name class-
identifier
CAT9300(config-sensor-dhcplist)#option name client-
identifier
CAT9300(config-sensor-dhcplist)#option name parameter-
request-list
```

Step 2. Create a list for CDP and configure the device name, version type, capabilities type, and platform type, as demonstrated here:

```
CAT9300(config)#device-sensor filter-list cdp list
<cdp_list_name>
```

```
CAT9300(config-sensor-cdplist)#tlv name device-name
CAT9300(config-sensor-cdplist)#tlv name platform-type
CAT9300(config-sensor-cdplist)#tlv name version-type
CAT9300(config-sensor-cdplist)#tlv name capabilities-type
```

Step 3. Create a list for LLDP that includes the port ID, system name, and system description, as demonstrated here:

```
CAT9300(config)#device-sensor filter-list lldp list
<lldp_list_name>
CAT9300(config-sensor-lldplist)#tlv name port-id
CAT9300(config-sensor-lldplist)#tlv name system-name
CAT9300(config-sensor-lldplist)#tlv name system-
description
```

You have now defined the options that Device Sensor should store.

Step 4. Configure Device Sensor to use the lists created in steps 1–3 as demonstrated here:

```
CAT9300(config)#device-sensor filter-spec dhcp include
list <dhcp_list_name>
CAT9300(config)#device-sensor filter-spec lldp include
list <cdp_list_name>
CAT9300(config)#device-sensor filter-spec cdp include
list <lldp_list_name>
```

Step 5. Enable Device Sensor to run on the switch and configure when it should send updates, as demonstrated here:

```
CAT9300(config)#device-sensor accounting
CAT9300(config)#device-sensor notify all-changes
```

VMware Configurations to Allow Promiscuous Mode

VMware vSwitches reject promiscuous mode by default (see Figure 14-20). In order to use SPAN probes with ISE, you need to configure a vSwitch to allow promiscuous connections.

Follow these steps to modify the configuration and enable promiscuous traffic:

Step 1. Click on the **Edit Settings** link in the vSwitch.

Step 2. Expand the **Security** field.

Step 3. Set the Promiscuous Mode radio button to **Accept**, as shown in Figure 14-21.

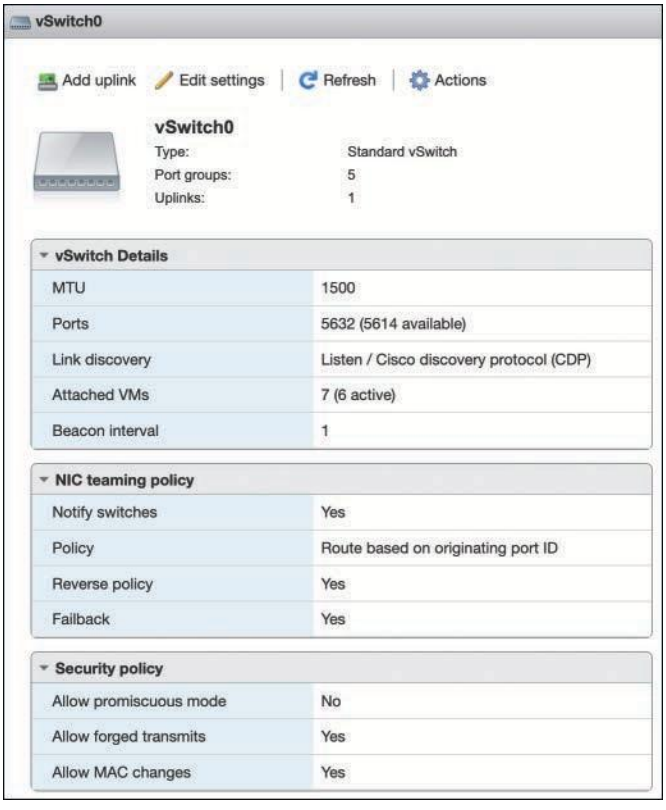


Figure 14-20 *Default vSwitch Configuration*

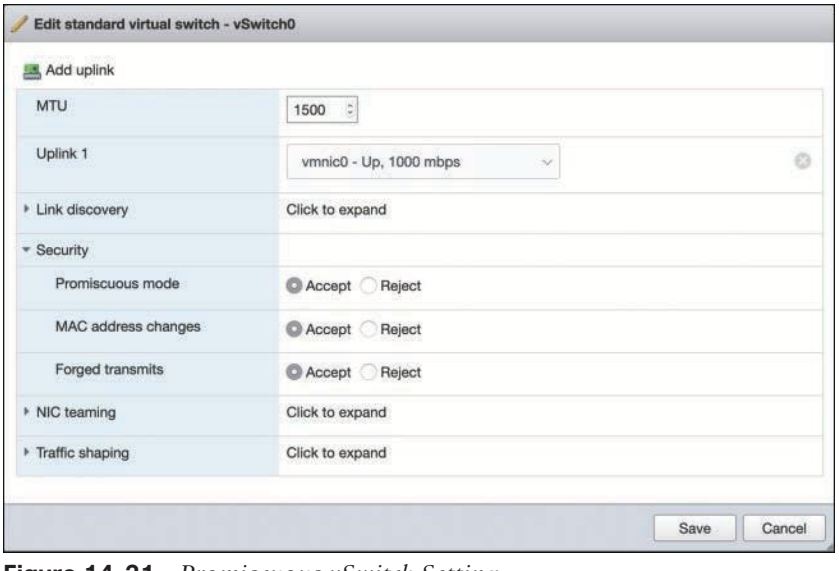


Figure 14-21 *Promiscuous vSwitch Setting*

Profiling Policies

Collecting data for profiling is only part of what needs to happen. In addition, you need endpoint signatures and a policy engine to compare the collected attributes to those signatures, which then leads to the assignment of the endpoint profile.

The profiling engine works a lot like an intrusion detection system (IDS), where traffic is compared to a set of signatures to identify suspicious activity. The profiling engine has hundreds of built-in signatures called *profiles* that are designed to match based on certain attributes. Much like an IDS system, with profiling there is also an update service in ISE to allow the engine to download new signatures.

Profiling Feed Service

Although ISE comes with a very large and comprehensive list of signatures to classify endpoints (profiles), many more devices are produced almost daily. For example, think of how often a new smartphone or version of the phone's operating system is released. New profiles are created and old ones are updated with new information that should be shared to the ISE deployments of the world. This is why Cisco created ISE Profiler Feed Service. When new devices are released to market, Cisco (and Cisco partners such as manufacturers) can create profiles for them. The ISE Profiler Feed Service is used to distribute new profiles after the quality assurance team has approved them.

Configuring the Profiler Feed Service

Configuring ISE Profiler Feed Service is straightforward. Once it is enabled, it reaches out to Cisco.com at the configured time interval and downloads any published profiles. Profiler Feed Service has an option to send an email alert to the administrator when an update occurs, an undo button for reversing the latest update, a link to view a report on the latest updates, and a section to send your information to Cisco to help with understanding how many customers are using Profiler Feed Service.

In case ISE is air-gapped so that it cannot directly connect to the Internet, there is an Offline Manual Update option that allows a user to download the latest profile pack and manually upload it to the ISE system.

Figure 14-22 shows an enabled and configured Profiler Feed Service screen, which you reach by navigating to Administration > Feed Service > Profiler.

In case you don't want to wait for a configured interval for the Feed Service to run, there is an Update Now button. Be cautious about manually updating profiles during a production workday. When the profiles are updated, all endpoints in the endpoint database are compared against the new list of profiles. In other words, a re-profiling of endpoints occurs, and it can be processor intensive.

Verifying the Profiler Feed Service

The easiest way to verify whether Feed Service is operational is to click the Test Feed Service Connection button on the configuration screen located at Administration > Feed Service > Profiler (see Figure 14-23).

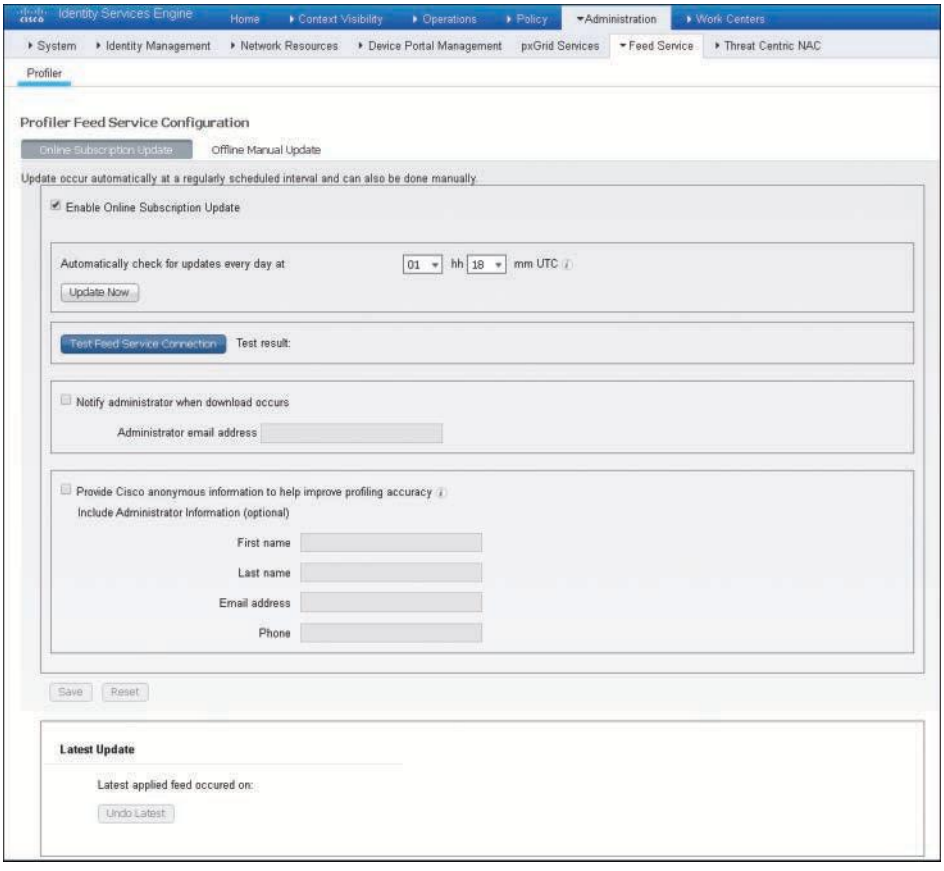


Figure 14-22 *Configured Profiler Feed Service*



Figure 14-23 *Test Feed Service Connection Button*

By clicking this button, you can run a test query to Profiler Feed Service. Figure 14-25 shows the result of a successful test.

Because ISE must be able to reach Cisco.com for the ISE Profiler Feed Service to function, it is quite possible that you may need to configure ISE to use a proxy server to reach the Internet. This optional configuration is located at Administration > System > Settings > Proxy (see Figure 14-24).

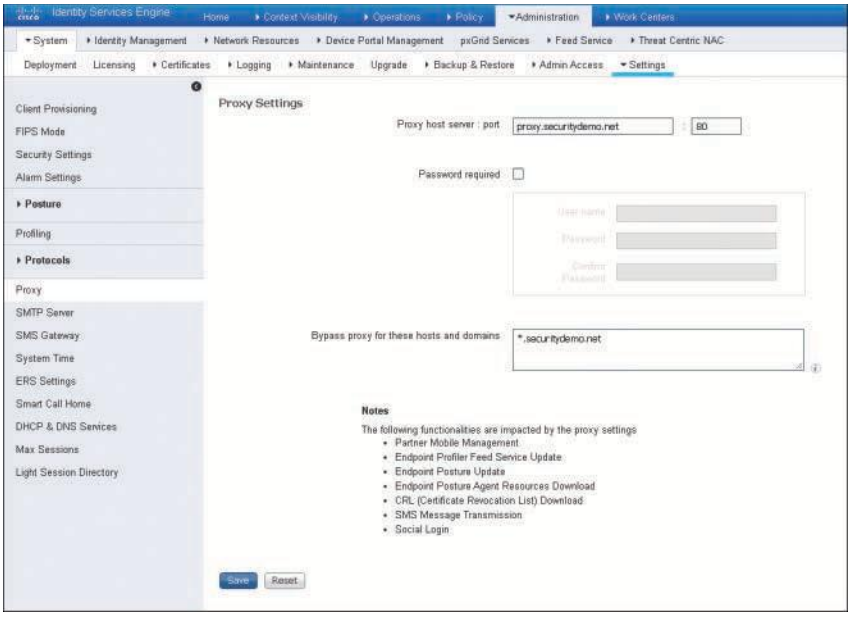


Figure 14-24 *Proxy Setting*

Endpoint Profile Policies

Profiler Feed Service probes collect attributes of endpoints to compare against the profile policies. As mentioned earlier, the profile policies are similar to signatures that define the endpoint profiles. For example, in order to match an Apple-Device profile, the endpoint must have a MAC address that begins with Apple's OUI.

Each endpoint profile policy defines a set of attributes that must be matched for a device to be classified as that endpoint type. ISE has a very large number of predefined profile policies, and you have just read about Feed Service, which is used to update those policies and provide new ones.

You can view the endpoint profile policies by navigating to Policy > Profiling, as shown in Figure 14-25.

Each profile has a system type of either Cisco provided, administrator modified, or administrator created. The administrator modified type is used when an ISE administrator modifies one of the default Cisco-provided profiles. The administrator created profile policies are custom profiles manually created by the ISE administrator. This classification ensures that Feed Service will not override a profile that has been changed or created by the administrator.

Profiles are hierarchical and inclusive in nature. You can pick any level in the hierarchy to use in your authorization policies, which means you can be specific or broad in your rules. In Figure 14-26, you can see a parent policy named Cisco-Device, which has a child policy named Cisco-IP-Phone, which in turn has a child policy named Cisco-IP-Phone-6921. When building an authorization policy, you can choose to use the profile at any point in that chain. If you were to select Cisco-Device, it would apply to all devices classified as Cisco-Device as well as anything classified as a child profile of Cisco-Device.

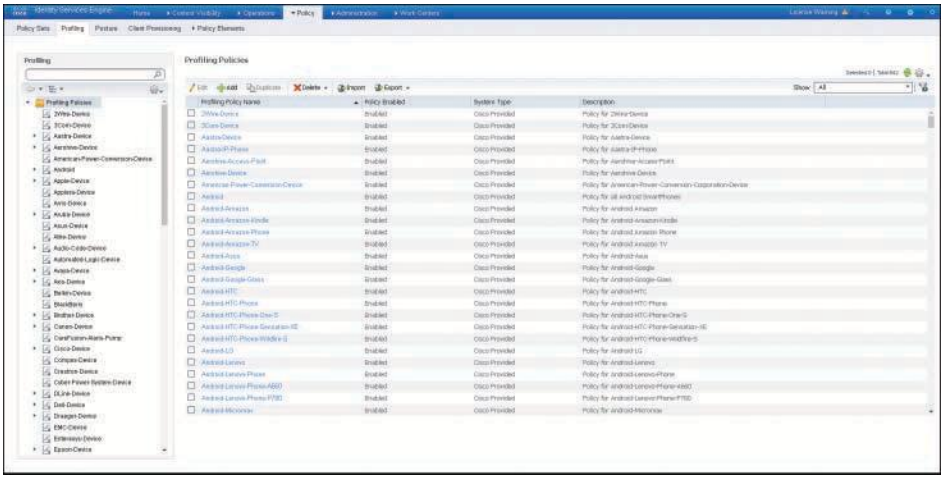


Figure 14-25 *Policy > Profiling*

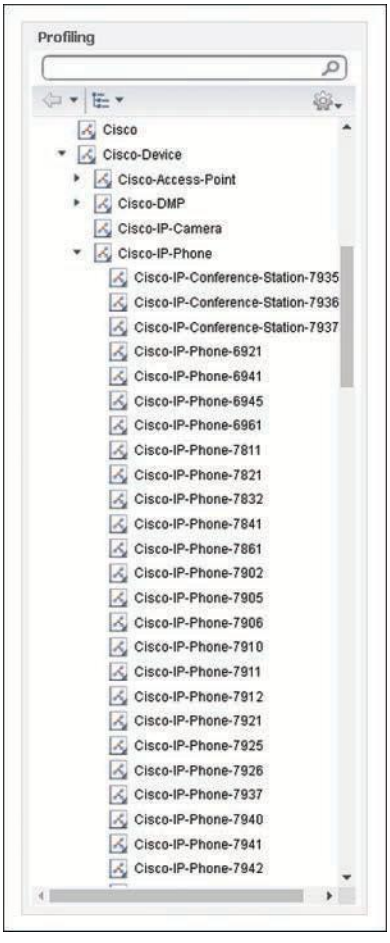


Figure 14-26 *Cisco-Device Parent Profile*

Figure 14-27 graphically displays the way a profile hierarchy is built in ISE.

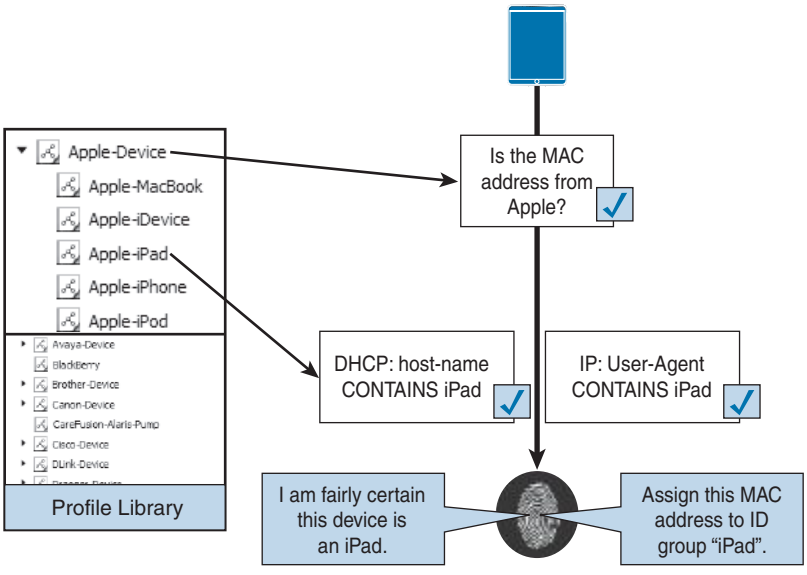


Figure 14-27 *Profiling Hierarchy Illustrated*

Figure 14-28 shows a predefined authorization rule named Profiled Cisco IP Phones in the default policy set. This rule permits full access to the network and assigns permission to join the voice VLAN—for all devices that are profiled as a Cisco-IP-Phone parent profile and any subsequent child profiles.



Figure 14-28 *Default Profiled Cisco IP Phones Rule*

To continue examining policy hierarchy, the following example involves a network where there are a couple Cisco 9971 IP Phones that have been granted access from the Profiled Cisco IP Phones default authorization rule, which permits all devices matching a Cisco-IP-Phone profile identity group. To verify an endpoint's assigned profile and the attributes for that profile, follow the following steps. Start by examining the endpoint attributes and comparing them to the profiling policies:

Step 1. Navigate to **Context Visibility > Endpoints > Authentication** and find the filter for Cisco-IP-Phone-9971 (see Figure 14-29).

Step 2. Click on one of the phone's MAC addresses to see the endpoint details.

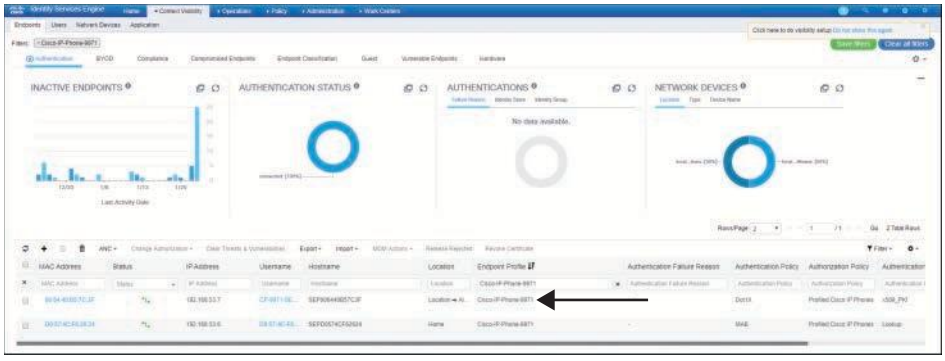


Figure 14-29 *Filtered Endpoint Identities*

Step 3. Click on the **Attributes** tab to view the endpoint attributes. You immediately see that the endpoint policy (that is, the profile of the device) is Cisco-IP-Phone-9971. However, the identity group assignment is Cisco-IP-Phone. In addition, you see that AuthorizationPolicyMatchedRule is Profiled Cisco IP Phones, which means the correct authorization rule is being matched (see Figure 14-30).

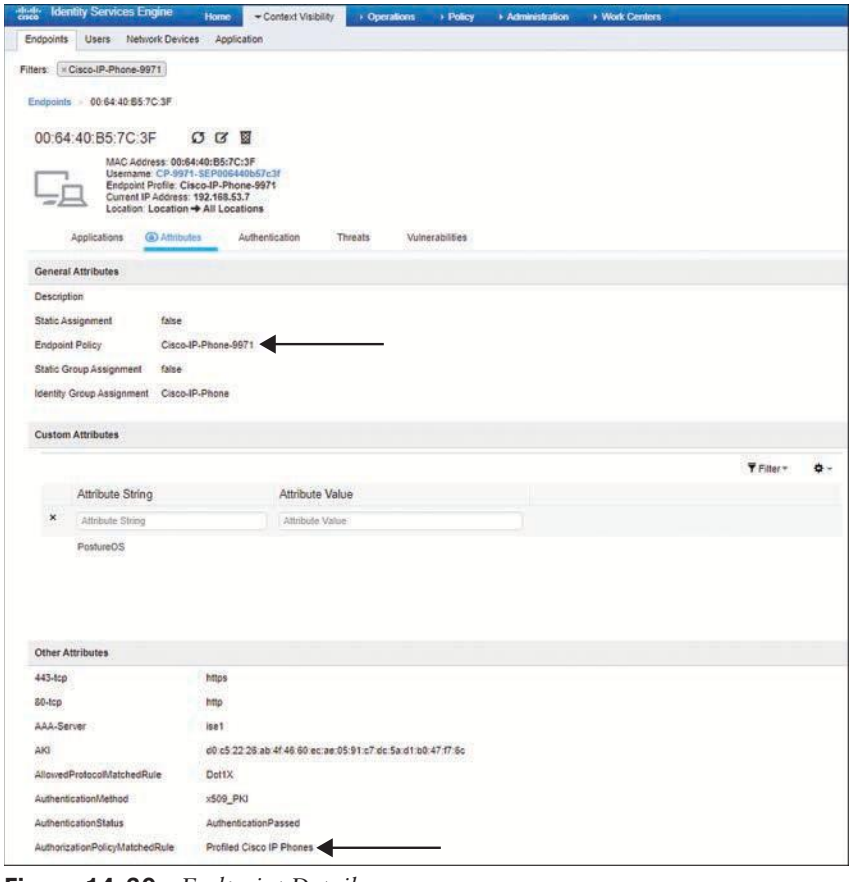


Figure 14-30 *Endpoint Details*

Step 4. Scroll down in the endpoint details and notice that the EndPointSource is SNMPQuery Probe (see Figure 14-31). This means the SNMPQUERY probe provided the most information to classify this device.

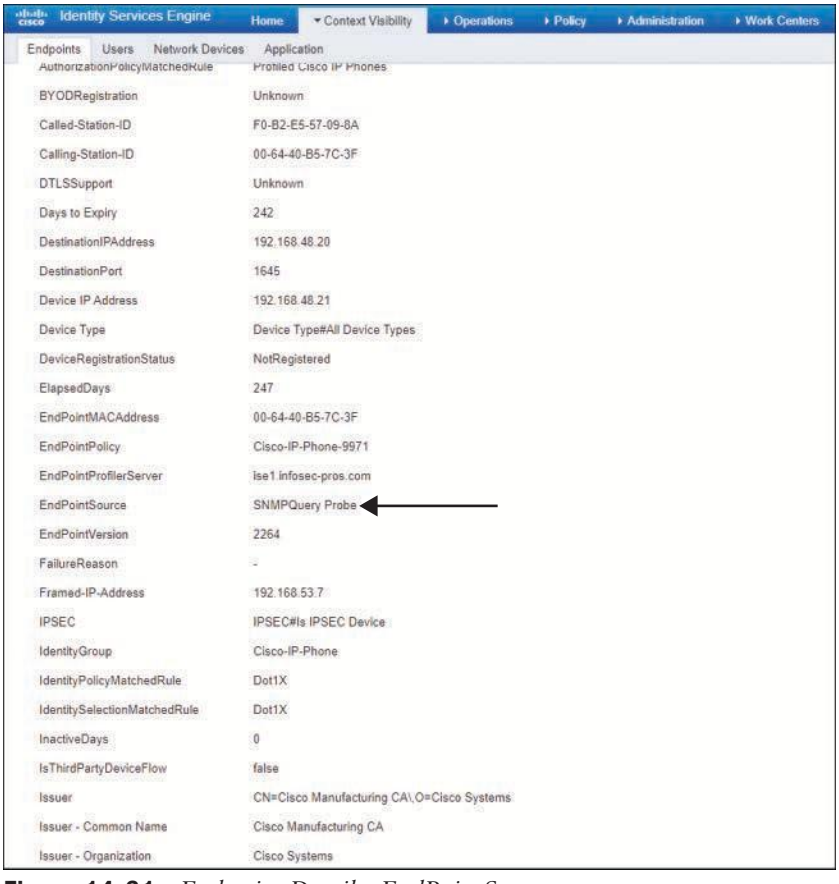


Figure 14-31 *Endpoint Details: EndPointSource*

Step 5. Continue down the list of attributes and notice the CDP cached data from the switch that was sent to ISE via the SNMPQUERY probe (see Figure 14-32).

You can see these attributes of the endpoint, but how is this information used in the profiling policy? To answer this question, you need to examine the profiling policy hierarchy for the endpoint.

Step 6. Navigate to **Policy > Profiling > Cisco-Device**. Cisco-Device is the top level of the profiling hierarchy for this endpoint (see Figure 14-33).

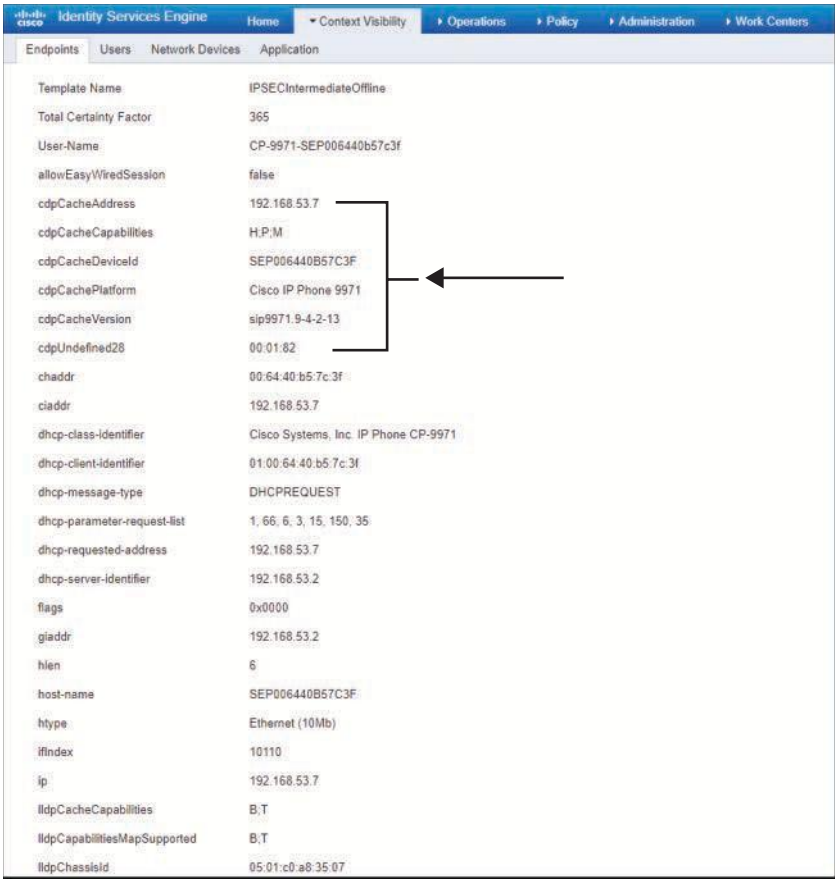


Figure 14-32 *Endpoint Details: CDP Data*

Using Figure 14-33 as a reference point, take note of the following details:

- The Minimum Certainty Factor of this profile is 10. Certainty factor is an aggregate value between 1 and 65535. All of the conditions at the bottom of the policy that are matched add up to equal the endpoint's certainty value. The higher the value, the more certain ISE Profiler is that an endpoint matches the specific profile.
- Note in Figure 14-34 that the OUI is Cisco Systems, Inc; this will match the condition for Cisco-Device shown in Figure 14-35. This is one possible mapping that meets the minimum certainty value and should match the endpoint to this parent policy.

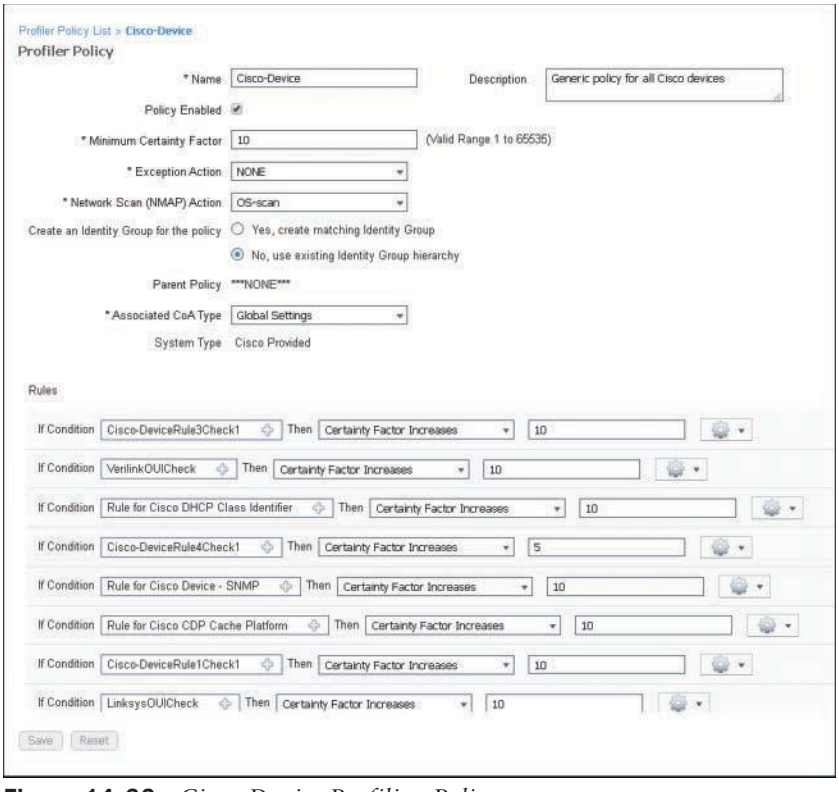


Figure 14-33 *Cisco-Device Profiling Policy*

- Network Scan (NMAP) Action is set to OS-scan. In order for this action to occur, there must be a condition in the profile that has a result to trigger the network scan. Figure 14-36 shows this mapping of the condition to the action. Figure 14-36 shows the MAC OUI contains CISCO, then ISE should take run the Nmap scan action that is defined.
- This profiling policy has a tremendous number of conditions, and most of them a certainty value of 5 or 10. The certainty value only needs to be a minimum of 10 to match the profile, so matching any one of these conditions will most likely create a match.

Step 7. Navigate to **Policy > Profiling > Cisco-Device > Cisco-IP-Phone** (see Figure 14-37) to examine the conditions that are used for this policy.

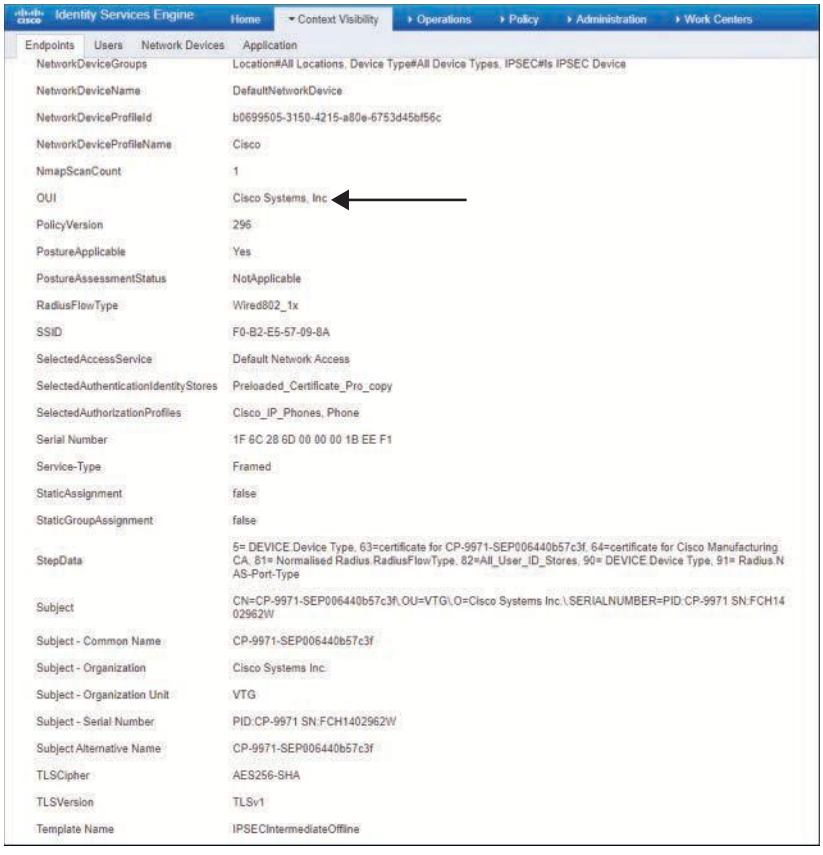


Figure 14-34 *OUI of the Endpoint*

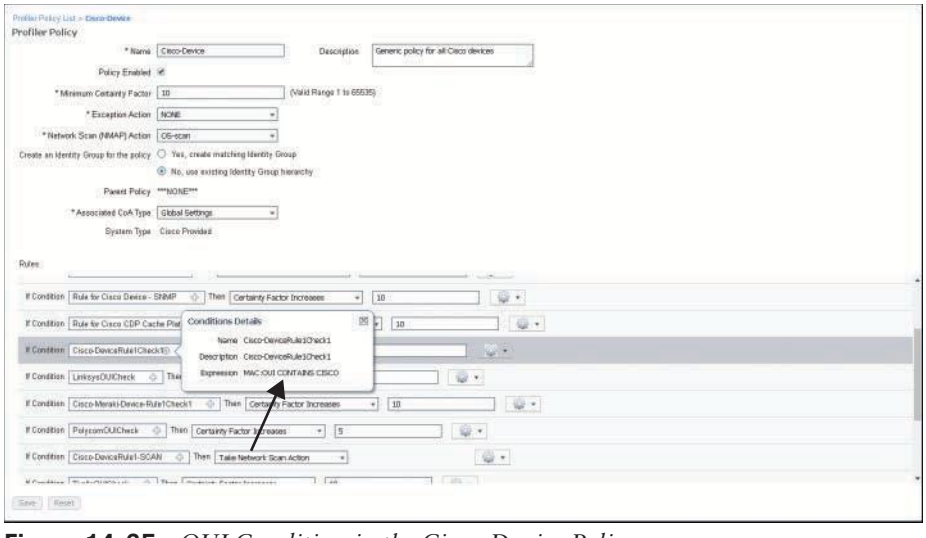


Figure 14-35 *OUI Condition in the Cisco-Device Policy*

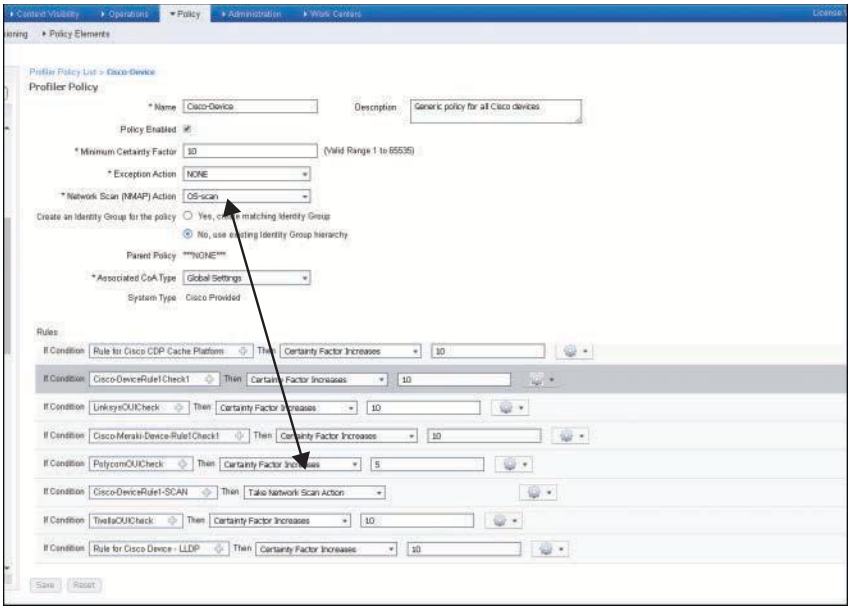


Figure 14-36 *Network Scan Action and Condition with Scan Result*

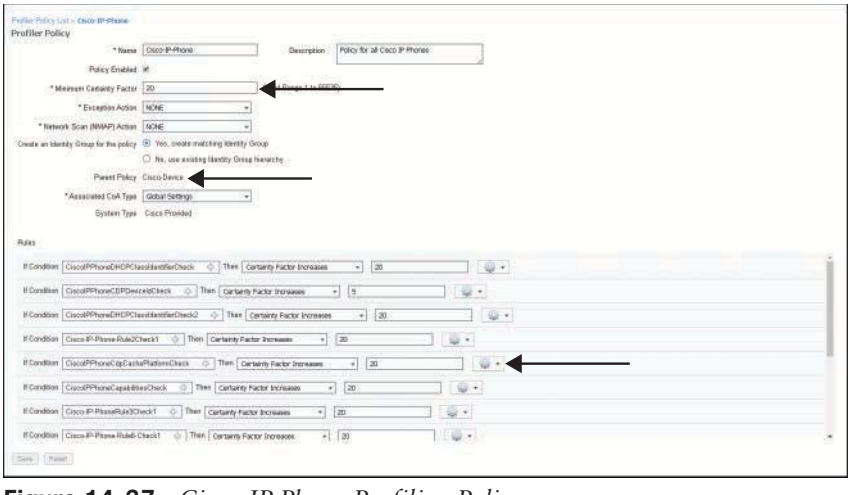


Figure 14-37 *Cisco-IP-Phone Profiling Policy*

Using this profile as a reference point, take note of the following details:

- In order to even be compared to this profiling policy, a device must have first matched its parent policy. In this case, the device had to match the Cisco-Device policy before these conditions could be examined.
- The Minimum Certainty Factor of this profile is 20, as shown in Figure 14-37. The certainty factor is an aggregate value between 1 and 65535. Each of the conditions at the bottom of the policy may add more certainty to this profile, if they are matched.

- The CDP cache in Figure 14-38 shows that the cdpCachePlatform attribute is sent as Cisco IP Phone 9971.
- Figure 14-38 shows that the Cisco IP Phone profile policy uses a condition that is looking for the cdpCachePlatform value to contain Cisco IP Phone and increase the certainty by 20.

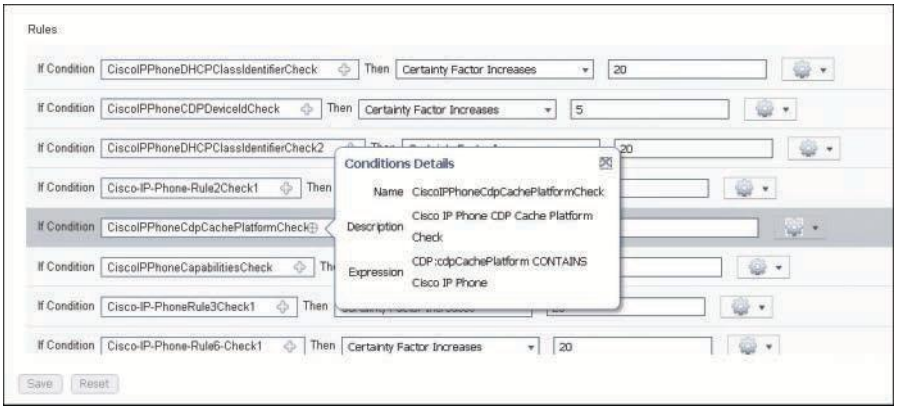


Figure 14-38 *cdpCachePlatform Condition*

Step 6 examined Cisco-IP-Phone, but what does it take to go further and reach the final Cisco-IP-Phone-9971 profile?

Step 8. Navigate to **Policy > Profiling > Cisco-Device > Cisco-IP-Phone > Cisco-IP-Phone-9971** so you can examine the conditions used for this policy (see Figure 14-39).

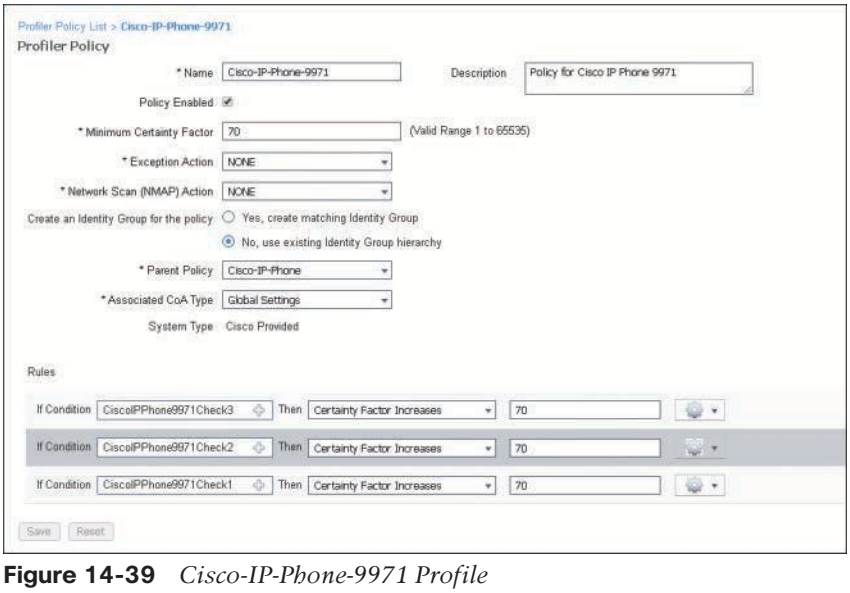


Figure 14-39 *Cisco-IP-Phone-9971 Profile*

Using Figure 14-39 as a reference point, note the following details:

- In order to even be compared to this profiling policy, the device must have first matched its parent policy. In this case, the device had to match the Cisco-Device and Cisco-IP-Phone policies before these conditions can be looked at.
- The profile has three conditions:
  - The DHCP Class-Identifier contains CP-9971.
  - The cdpCachePlatform contains Cisco IP Phone 9971.
  - The lldpSystemDescription contains Cisco IP Phone 9971.
- If any of these conditions are met, the endpoint matches the Cisco-IP-Phone-9971 profile.

NOTE Despite what you see in this example, you would rarely build an authorization policy that is specific to the point of the model number of the Cisco IP Phone. Instead, you would just use the Cisco-IP-Phone parent policy in your authorization policies.

Logical Profiles

When ISE Version 1.0 was first released, many customers quickly asked to have a grouping of profiles that was not hierarchical. For example, they wanted to be able to create a profile group named IP-Phones that contained all the individual profiles of IP phones, both from Cisco and other vendors.

Key Topic

Cisco answered that request in ISE Version 1.2 and later by introducing ***logical profiles***. A logical profile is exactly what customers requested: a grouping of profiles.

To examine the only prebuilt logical profile in ISE, navigate to Policy > Profiling > Logical Profiles and select IP-Phones, as shown in Figure 14-40.

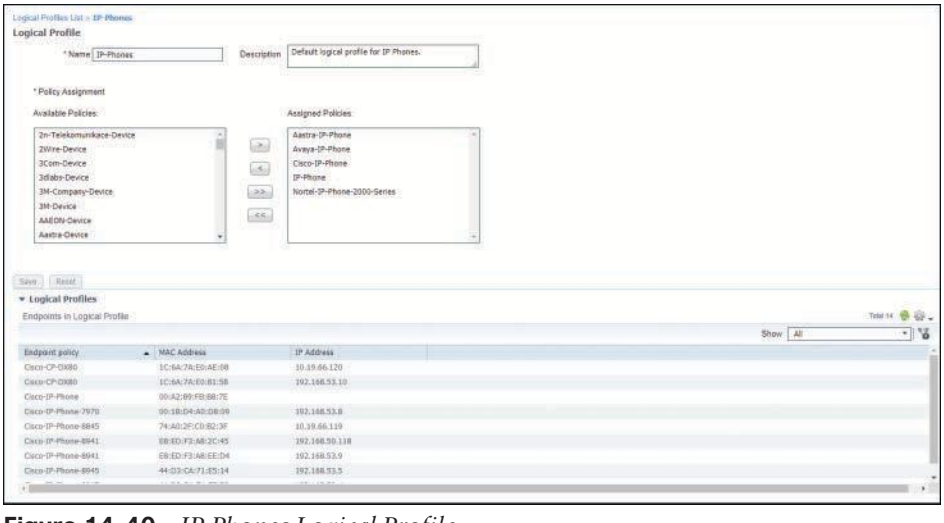


Figure 14-40 *IP-Phones Logical Profile*

Notice in Figure 14-40 that the logical profile contains Avaya, Cisco, Nortel, and generic IP phone profiles.

ISE Profiler and CoA

When you use an endpoint profile as an attribute in an authorization policy, you provide differentiated results for specific profiles. However, there is often a "chicken and the egg" phenomenon happening simultaneously. You cannot provide the right access to a device without knowing what that device is, and you cannot find out what the device is without providing some level of access so that the endpoint will be active on the network and ISE can identify the endpoint profile. This is where Change of Authorization (CoA) becomes important. (See Chapter 5 for more information about CoA.)

Before CoA, the only time a policy server like ISE was permitted to send a command to a NAD was during a response to an authentication request. This created numerous issues because there was not a way to disconnect a bad actor from the network or change the level of access an endpoint was permitted to have based on newer data learned at the policy engine. The current authorization to the network would have to remain in place until the next time the endpoint had to authenticate.

An authorization policy can be configured to send different results for an endpoint before it is profiled and then send another level of authorization after the endpoint profile becomes more solidified. This is due to the fact that once the final result of the endpoint profile is known, you should not wait for the next authentication request before changing the level of access as that would prevent much-needed access. Instead, the profiling engine can use CoA to change the level for each state the endpoint goes through.

There are two main ways to configure CoA with profiling: Use a ***global CoA*** for profiling in the ISE deployment or configure a CoA on a per-profile basis.

Global CoA

To enable CoA for profiling in the ISE cube and to configure the CoA type used by profiling globally, navigate to Administration > System > Settings > Profiler, as shown in Figure 14-41.

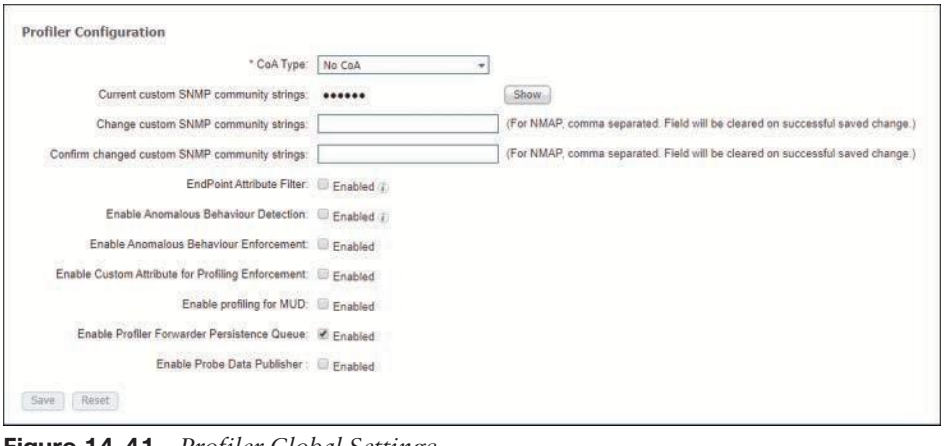


Figure 14-41 *Profiler Global Settings*

As shown in Figure 14-41, the default setting is No CoA. Click the CoA Type dropdown to see the other choices: Port Bounce and Reauth (see Figure 14-42).

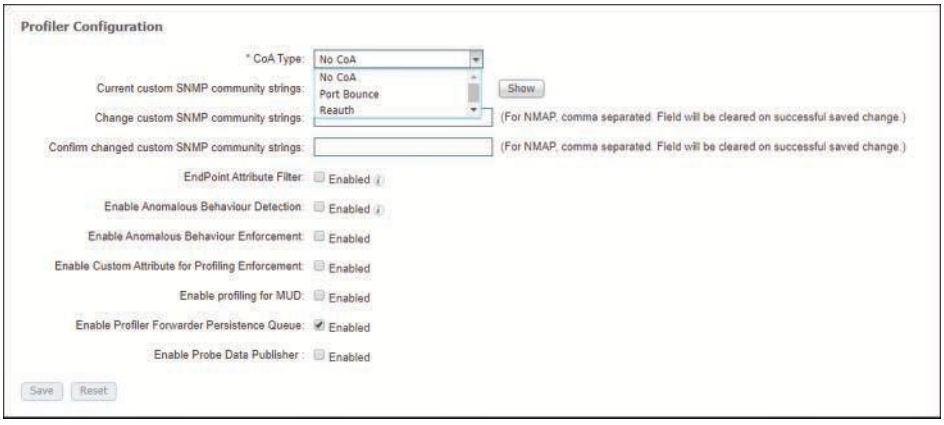


Figure 14-42 *Profiler CoA Types*

The Port Bounce CoA shuts down the switch port and then performs a no shutdown to reenable it. This causes the link state to change, which simulates the unplugging and plugging in of the network cable. The benefit to this type of CoA is that many devices try to renew their DHCP-assigned IP addresses when the link state changes. This can be particularly useful for headless devices and IoT if there is a requirement to place them in a separate VLAN. In addition, there is a built-in failsafe to never send a port bounce when there is more than one MAC address on a switch port. This failsafe ensures that there is no negative impact on IP telephony. When more than one MAC address exists on a switch port, a Reauth CoA is sent instead.

The Reauth CoA instructs the NAD to initiate a new authentication to the endpoint by sending another EAPoL Start message to trigger the supplicant to send the credentials again. In the case of MAB, the NAD resends a RADIUS authentication with the endpoint MAC address as the identity credential. Either way, there will be a new authentication, but that authentication maintains the same authentication session ID. By maintaining the session ID, ISE is able to tie together the multiple states of the endpoint.

Key Topic

With any CoA type, ISE forces a new authentication attempt so that a different authorization result can be sent to the NAD to provide the correct level of network access with the latest profiling information being used. However, setting a global CoA type to Port Bounce is not recommended. The safer bet is to use the Reauth option.

When the profiler CoA is enabled globally, a CoA is automatically sent for any endpoint that transitions from unknown to any known profile.

Per-Profile CoA

The individual profile policies have a setting that allows administrators to control their own destiny with CoA. The per-profile CoA addresses the need to send a Port Bounce CoA for certain devices only while using the global Reauth CoA for the remaining endpoints. A per-profile policy CoA takes priority over a global CoA.

When a CoA type is configured for a profile, it is used when an endpoint is classified as that profile type. Figure 14-43 shows the settings.

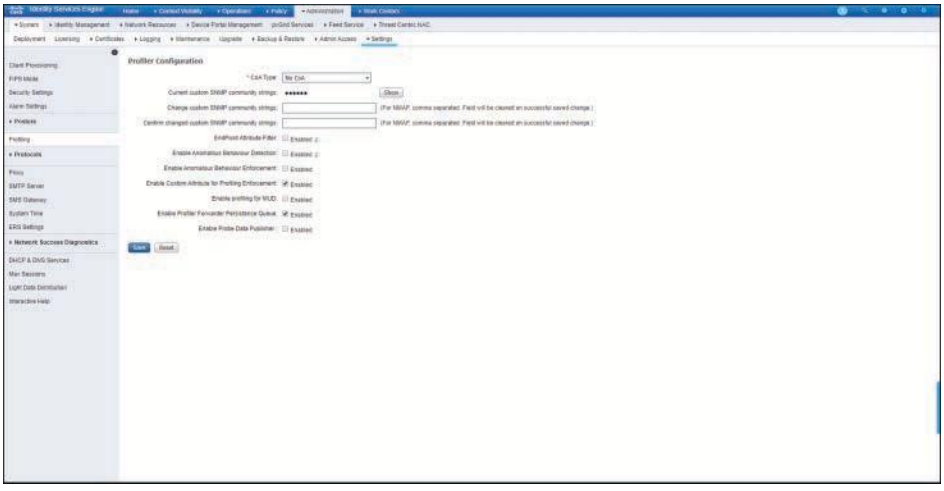


Figure 14-43 *Per-Profile CoA*

As shown in Figure 14-43, a device that is profiled as a Xerox-Device triggers a Port Bounce CoA. This causes the link to go down and back up again, which should trigger the endpoint to request a new IP address from the DHCP server. This is very useful when a device is using MAB and needs to be assigned to a different VLAN.

Global Profiler Settings

Additional settings related to profiling (beyond the global CoA type) are set at the global (systemwide) level: SNMP community strings for Nmap SNMP walks and the enable setting for endpoint attribute filtering.

Configure SNMP Settings for Probes

The SNMPQUERY probe uses the SNMP community strings that are defined as part of the NAD entry under Administration > Network Resources > Network Devices. Each NAD could theoretically have a different community string.

As described earlier in this chapter, Nmap uses SNMP to examine endpoints. In order for this to function, ISE Profiler must know what SNMP community strings to use. The community strings to use are configured under Administration > System > Settings > Profiler, and the community strings are listed one-by-one, with commas separating the values. When they are saved, the two text boxes are erased, and you must click the "Show" button to see the configured strings. The string will be displayed on a popup as shown in Figure 14-44.

Endpoint Attribute Filtering

ISE Profiler can and does collect a lot of data about endpoints. It stores all that data and replicates the data to the other ISE nodes in the deployment. To minimize the replication traffic, ISE offers the EndPoint Attribute Filter setting, which is enabled under Administration > System > Settings > Profiler (see Figure 14-45).

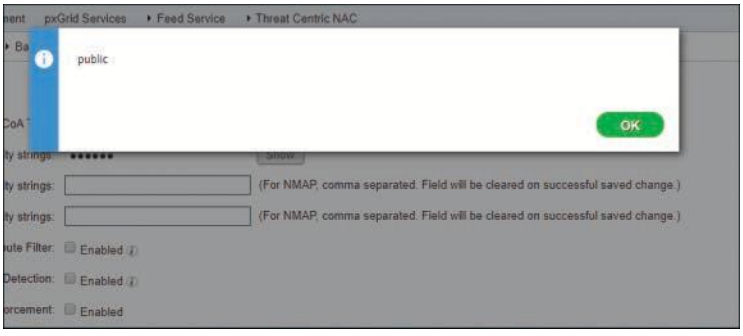


Figure 14-44 *Global SNMP Settings for the NMAP Probe*

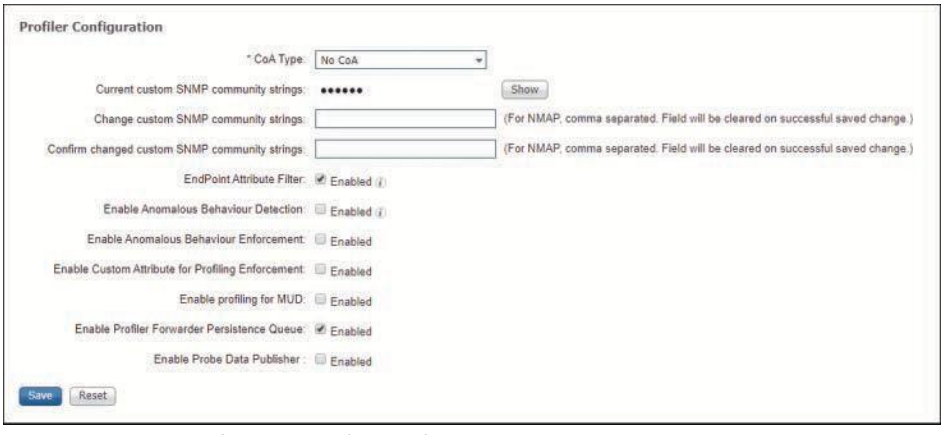


Figure 14-45 *Endpoint Attribute Filter*

When this filtering is enabled, Profiler builds a whitelist of attributes that are used in the existing Profiler policies. In other words, Profiler examines every policy that is enabled and creates a list of attributes that are needed for those policies. Only those attributes are collected and stored in the endpoint database.

Use of the EndPoint Attribute Filter option is highly recommended but only after a deployment has been up and running properly for an extended period of time. You should not enable it immediately because it might take some time to build custom profiles and policies after you set up ISE.

Custom Attributes for Profiling

An ISE administrator can assign custom attributes to an endpoint in addition to the attributes that are gathered from the probes. These custom endpoint attributes can then be used in an authorization policy to profile endpoints. You can create a maximum of 100 endpoint custom attributes.

To create a new custom attribute for use with endpoints, navigate to Administration > Identity Management > Settings > Endpoint Custom Attributes, as shown in Figure 14-46.

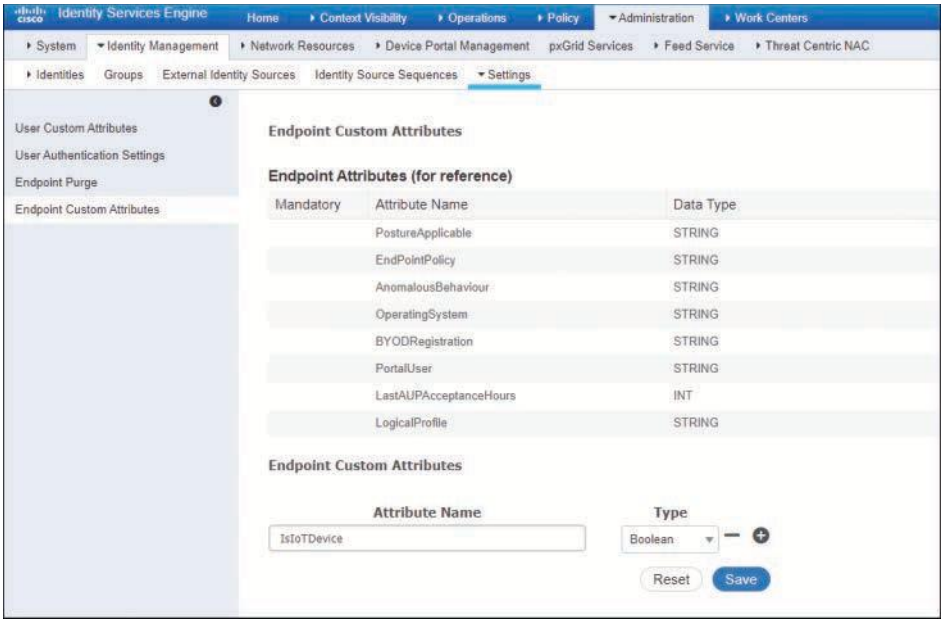


Figure 14-46 *Endpoint Custom Attributes*

Figure 14-46 shows the custom Boolean attribute IsIoTDevice added as an endpoint attribute.

To define the custom attribute under an endpoint, follow these steps:

Step 1. Navigate to **Context Visibility > Endpoints > Authentication**.

Step 2. Click on the MAC address of an endpoint.

Step 3. Navigate to the Attributes tab for that endpoint. As shown in Figure 14-47, the Custom Attributes section now contains any undefined administrator-created attributes.

Step 4. To define the attribute you have just created, click on the notepad and pen icon next to the MAC address heading on the top of the endpoint attributes screen. You can now edit the customizable features for this specific endpoint, as shown in Figure 14-48.

Step 5. Click on the notepad and pen icon next to the IsIoTDevice custom attribute to edit the attribute.

Step 6. Enter an attribute value and click the checkbox next to the custom attribute to set it. Since this custom attribute is a Boolean attribute, the options to enter are true or false. Figure 14-49 shows the added attribute value true for the IsIoT-Device custom attribute.

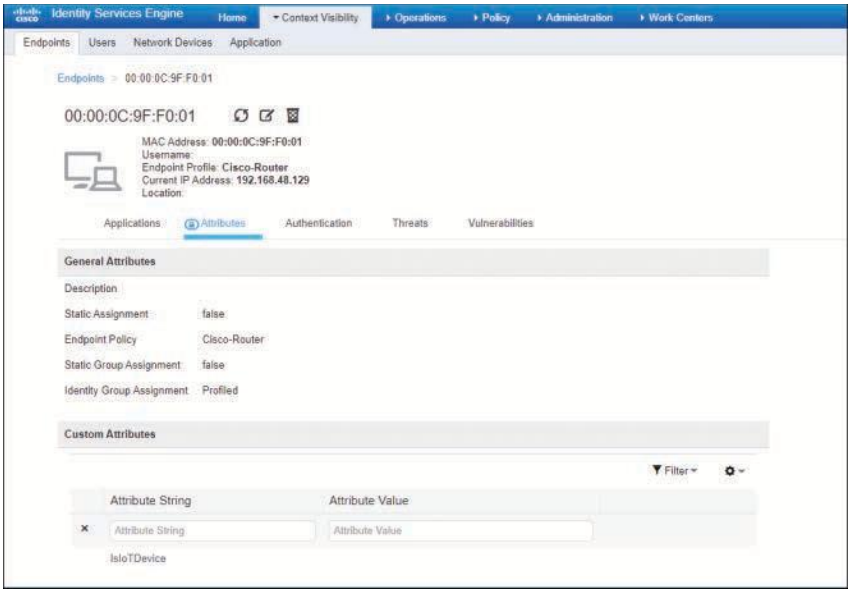


Figure 14-47 *Endpoint Attributes*

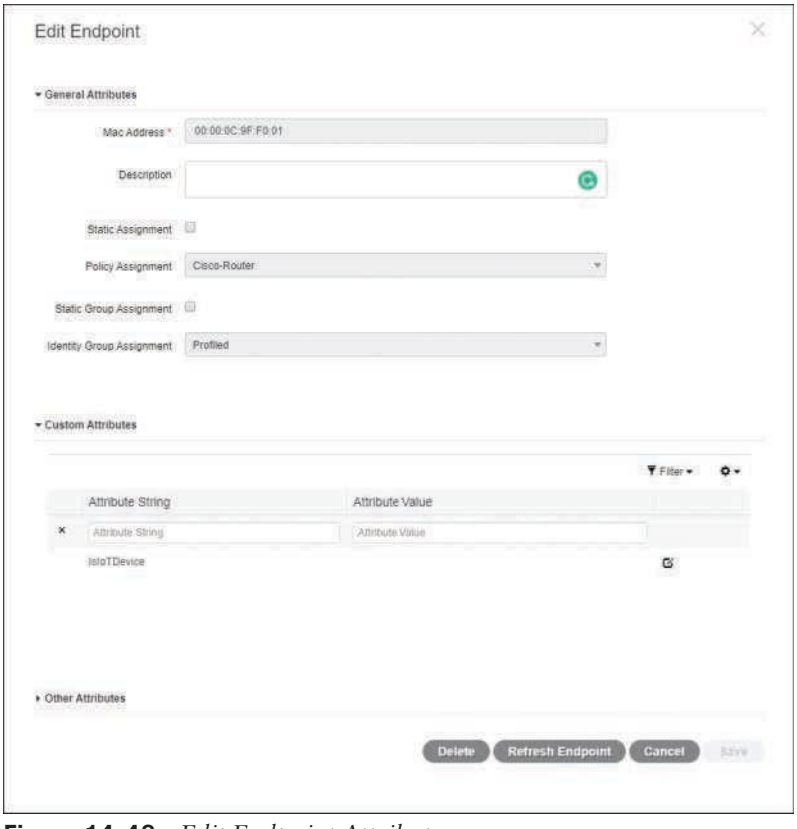


Figure 14-48 *Edit Endpoint Attributes*

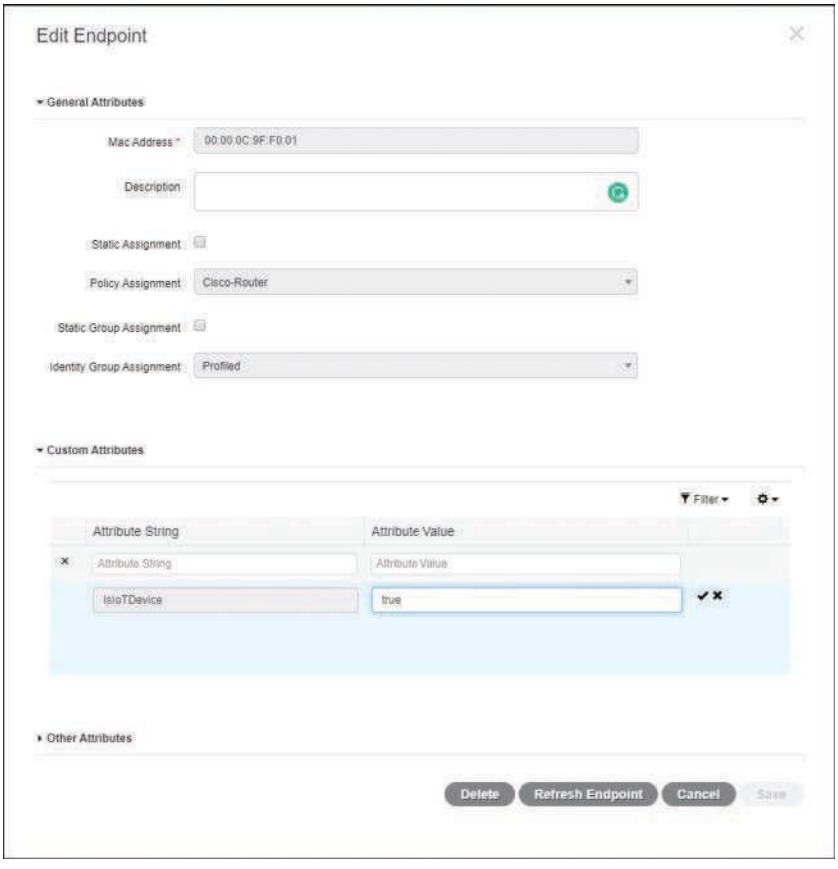


Figure 14-49 *Editing a Custom Attribute*

Step 7. Click **Save** to save your change. Figure 14-50 shows the saved endpoint attributes, including the custom attribute set to true.

Custom endpoint attributes can be manually set by an ISE administrator or set through the RESTful API feature in ISE. The API can be helpful when integrating ISE with third-party applications. Using the API to alter the custom attributes is beyond the scope of this book.

To use custom attributes for profiling purposes, navigate to Administration > System > Settings > Profiling and check the box next to Enable Custom Attribute for Profiling Enforce, as shown in Figure 14-51.

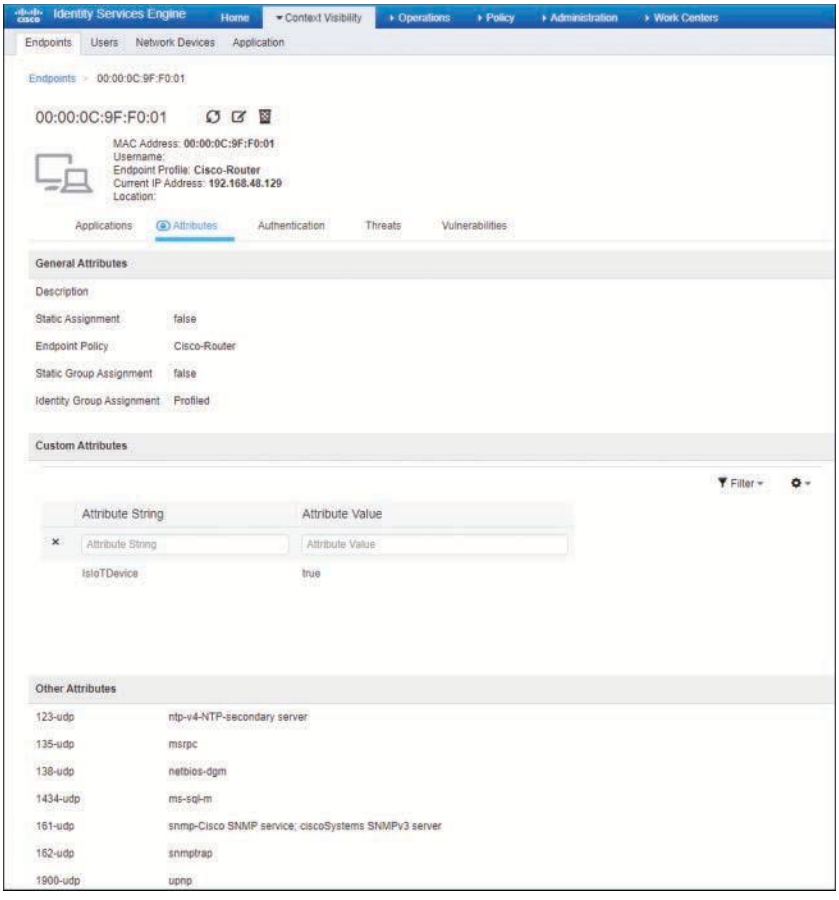


Figure 14-50 *Endpoint Attributes*

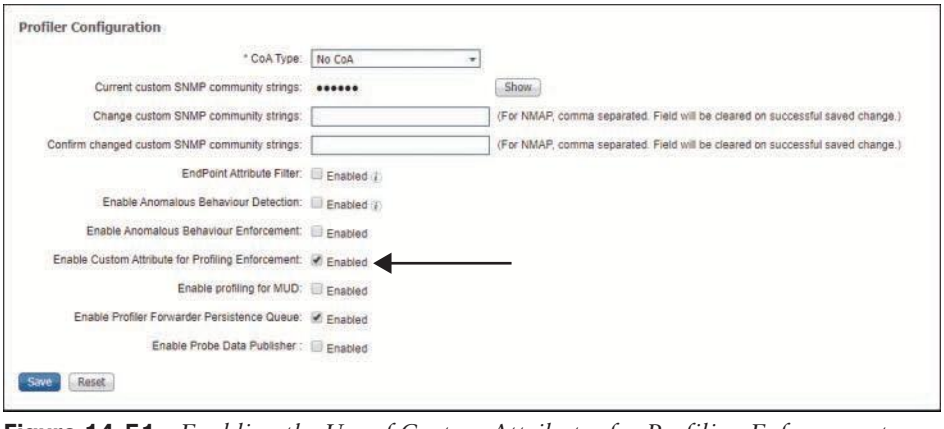


Figure 14-51 *Enabling the Use of Custom Attributes for Profiling Enforcement*

Publishing Endpoint Probe Data on pxGrid

An ISE administrator may optionally publish endpoint probe data to pxGrid subscribers (third-party and other Cisco systems) to add contextual information to classify endpoints. ISE sends the endpoint records to the pxGrid subscribers whenever they are updated in the ISE PAN. In order to enable this feature, navigate to Administration > System > Settings > Profiling and check the box next to Enable Probe Data Publisher, as shown in Figure 14-52.

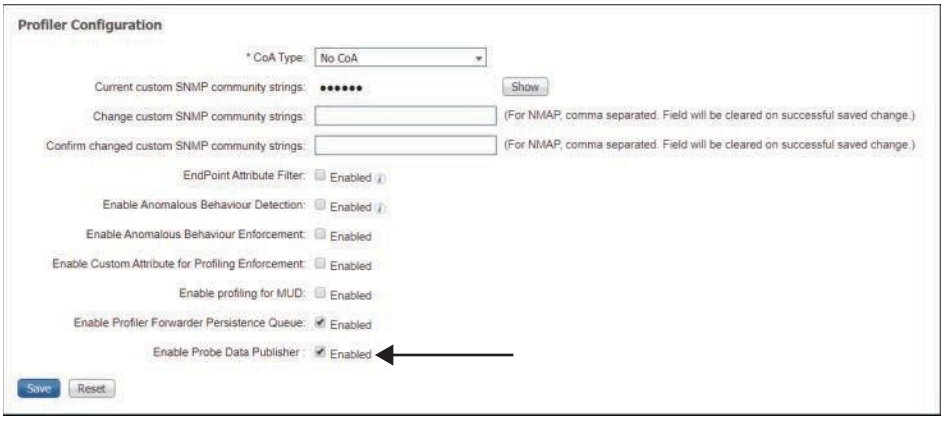


Figure 14-52 *Enabling Probe Data Publisher*

Profiles in Authorization Policies

As you saw with the Profiled Cisco IP Phone authorization rule earlier in this chapter, a profile may be used as a condition of an authorization policy rule in the form of an identity group. Originally, ISE required an identity group in order to use any of the profiling policies in the rule. However, it is now possible to use a profile directly, through the EndPointPolicy attribute.

Endpoint Identity Groups

Local identities in the ISE database can be in the form of user identities or endpoint identities. In addition, identity groups may contain multiple identities, although an identity (a user or an endpoint) may be a member of only one identity group at a time.

To create an identity group based on a profile, you simply select the option Yes, Create Matching Identity Group, as shown in Figure 14-53.

When this option is selected, the matching identity group can be found under Administration > Identity Management > Groups > Endpoint Identity Groups > Profiled (see Figure 14-54).

Even though endpoint identity groups were useful in very early versions of ISE, their use for profiling has been deprecated in favor of using actual endpoint profiles or logical profiles directly in the authorization policy.

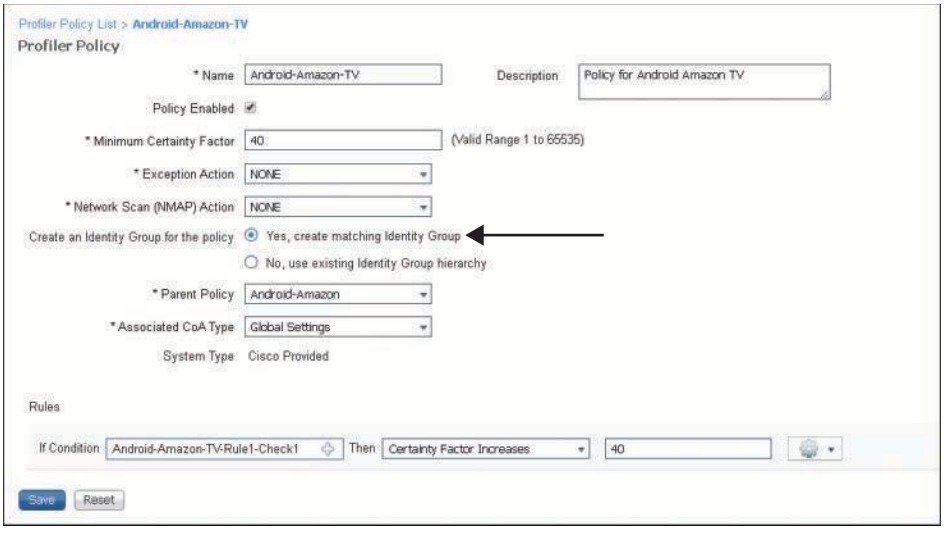


Figure 14-53 *Creating a Matching Identity Group*

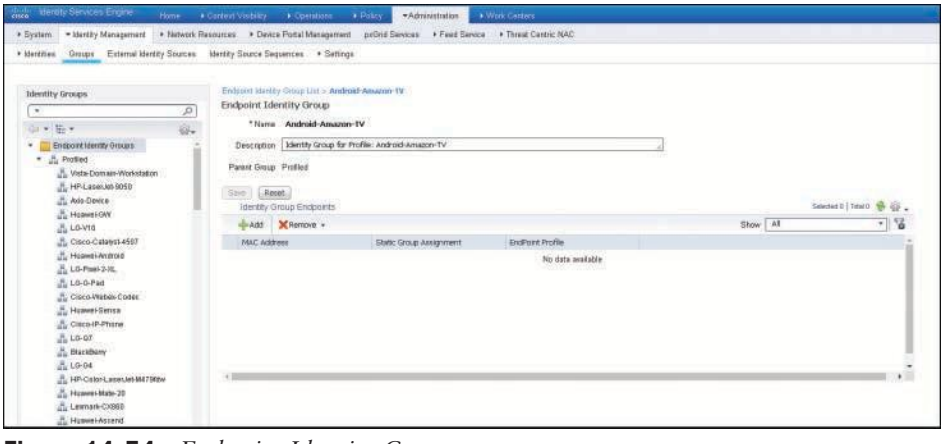


Figure 14-54 *Endpoint Identity Groups*

Today, endpoint identity groups are used for a different purpose. They are used for a MAC Address Management (MAM) model, where you can create a static list of MAC addresses to be authorized specifically (for example, a list of all Apple iPads that are owned by the company so they can be differentiated from personally owned iPads).

For example, consider the Blacklist identity group. If a user loses her personal device, she can log in to the My Devices portal and mark a device as lost (see Figure 14-55). The endpoint is immediately added to the Blacklist group, as shown in in Figure 14-56, and then the endpoint is denied network access by default, as shown in in Figure 14-57.



Figure 14-55 *My Devices: Marking an Endpoint as Lost*



Figure 14-56 *Blacklisted Devices*

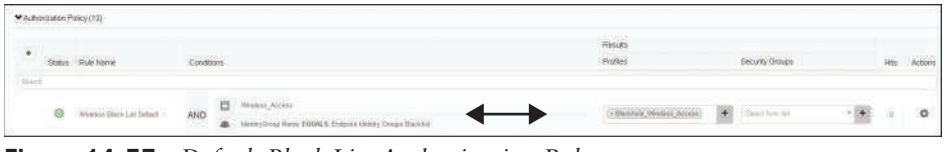


Figure 14-57 *Default Black List Authorization Rule*

In the My Devices portal, it is possible to click Reinstate to move the device from the Blacklist group to the RegisteredDevices group. Figure 14-58 shows the My Devices portal, and Figure 14-59 shows the RegisteredDevices group.



Figure 14-58 *My Devices: Reinstating an Endpoint*

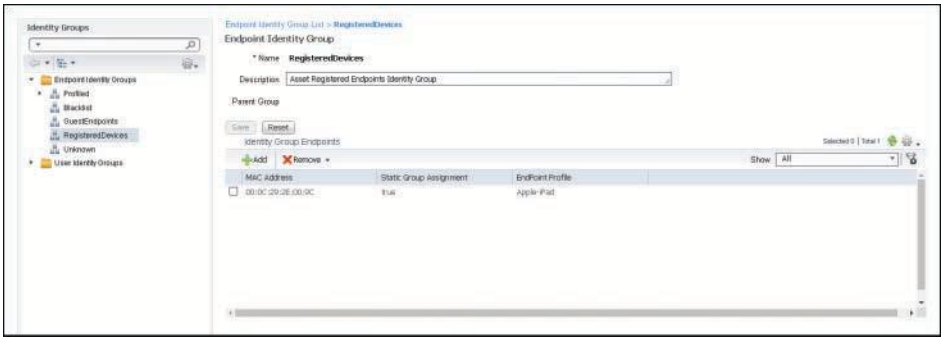


Figure 14-59 *RegisteredDevices Endpoint Identity Group*

EndPointPolicy

As mentioned earlier, using identity groups is no longer the best way to apply policy based on an endpoint's profile. Policy is now built with an actual profile through the use of the Endpoints:EndPointPolicy attribute.

In order to see the use of a profile in a policy, the following example shows how to duplicate the existing NSP authorization rule and then modify that new rule to use the Android profile, which does not have a corresponding identity group. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the default policy set.

Step 3. Duplicate the disabled Employee_Onboarding rule by clicking the cog icon and choosing **Duplicate Above**.

Step 4. Rename the rule **Employee_Onboard_Android**.

Step 5. Add a new condition to the existing EAP-MS-CHAP condition by selecting **Endpoints > EndPointPolicy:Android**.

Step 6. Enable the authorization rule by clicking the grayed-out circle on the left side of the policy and choosing **Enable** from the dropdown.

Step 7. Click **Save** to save the policy set.

Figure 14-60 shows the resulting authorization rule.

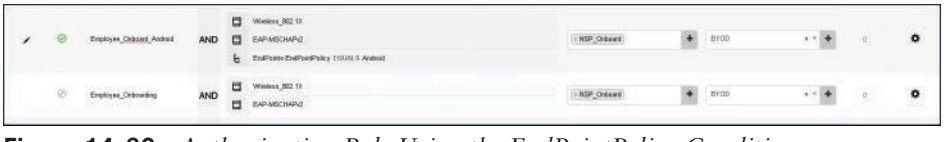


Figure 14-60 *Authorization Rule Using the EndPointPolicy Condition*

Keep in mind that logical profiles can streamline your authorization policies and limit the number of individual EndPointPolicy conditions that you need to add to your authorization rules.

Verify Profiling

There are a few key places to check when verifying profiling operation. Most of them are in the ISE GUI, and of course you can a network device itself.

The Dashboard

By navigating to Context Visibility > Endpoints > Endpoint Classification, you can glean a lot of information (see Figure 14-61).

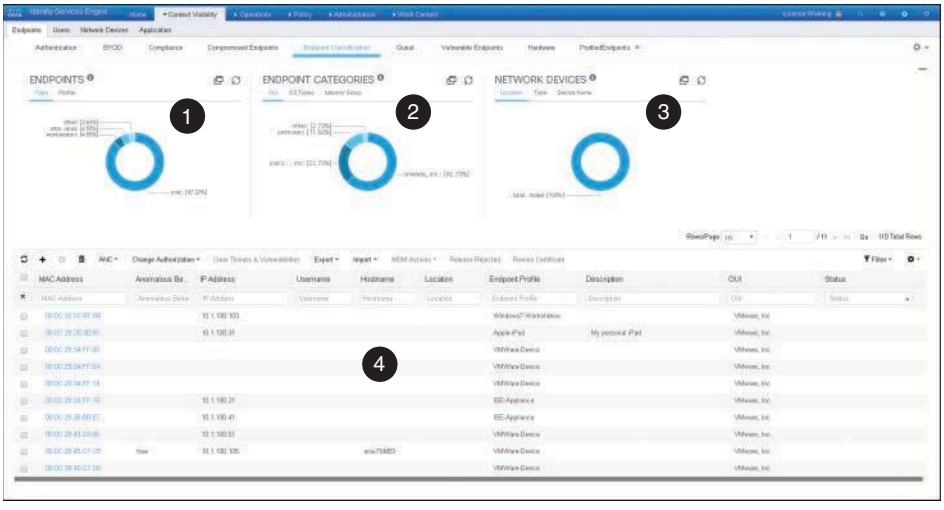


Figure 14-61 *Context Visibility Endpoint Classification*

In Figure 14-61, several areas of this dashboard are marked with numerals that correspond with the following list:

1. **Endpoints widget**: This widget shows the high-level type of endpoints at profile types.
2. **Endpoint Categories widget**: This widget shows the high-level endpoint types by MAC OUI, OS types, and identity groups.
3. **Network Devices widget**: This widget displays the different NADs that endpoints are connecting to, based on NAD location, type, and device name.
4. **Endpoint list**: This list allows an ISE administrator to view the individual endpoints and filter by a variety of types of endpoint data.

Global Search

The Global Search tool is available from any screen in the ISE GUI. When you type in a profile name, you see a list of results that you can drill into further. Figure 14-62 shows the use of the Global Search tool to find profiles containing “apple.”

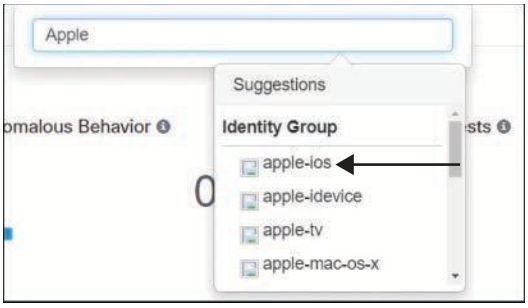


Figure 14-62 *Global Searching for Apple*

If you select apple-device from the list in the Global Search tool, you can see a more focused list of endpoints matching that profile, as shown in Figure 14-63.

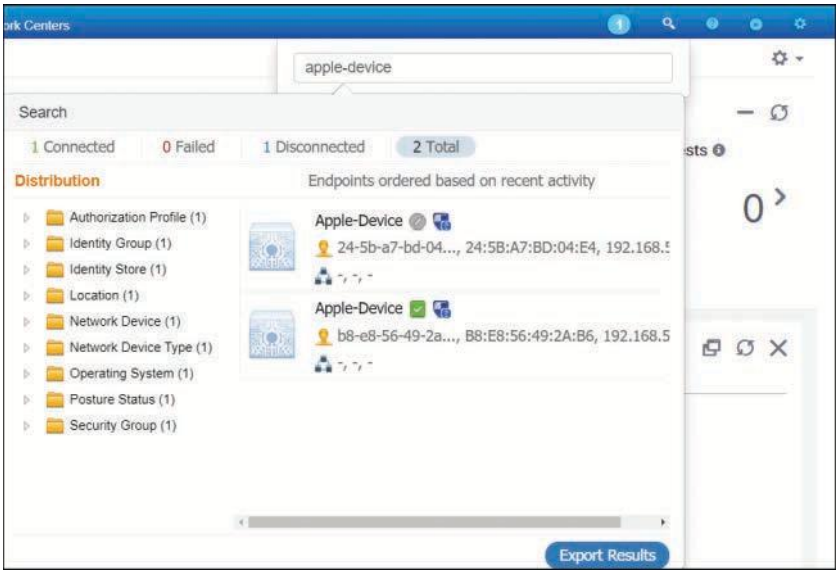


Figure 14-63 *Apple-Device Profile*

Endpoint Identities

The ultimate source of truth about endpoints is the endpoints database, which you can view in Context Visibility. For a graphical view of a list of endpoints, you can navigate to Context Visibility > Endpoints > Authentication. You can then see all the endpoints, with the profile of an endpoint shown in the Endpoint Profile column (see Figure 14-64).

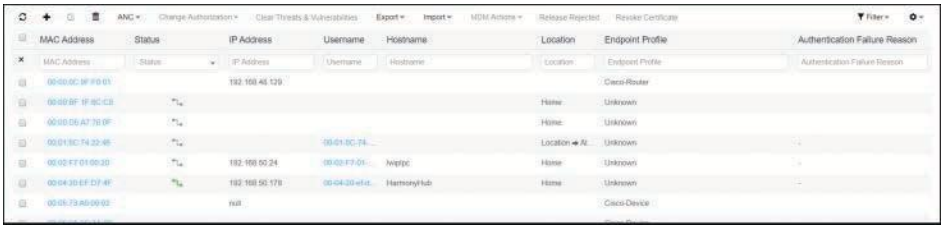


Figure 14-64 *Context Visibility Authentications Dashboard*

Clicking on the MAC address of an endpoint in the table brings up that endpoint's details. By navigating to the Attributes tab, you can view detailed information about the profiled attributes for the device. The list of details maintained for each endpoint can be quite large, especially if the EndPoint Attribute Filter option is not enabled. Figure 14-65 shows the endpoint details screen for an Apple iPhone, and Figure 14-66 shows some of the details further down this list, including some of the attributes profiled.

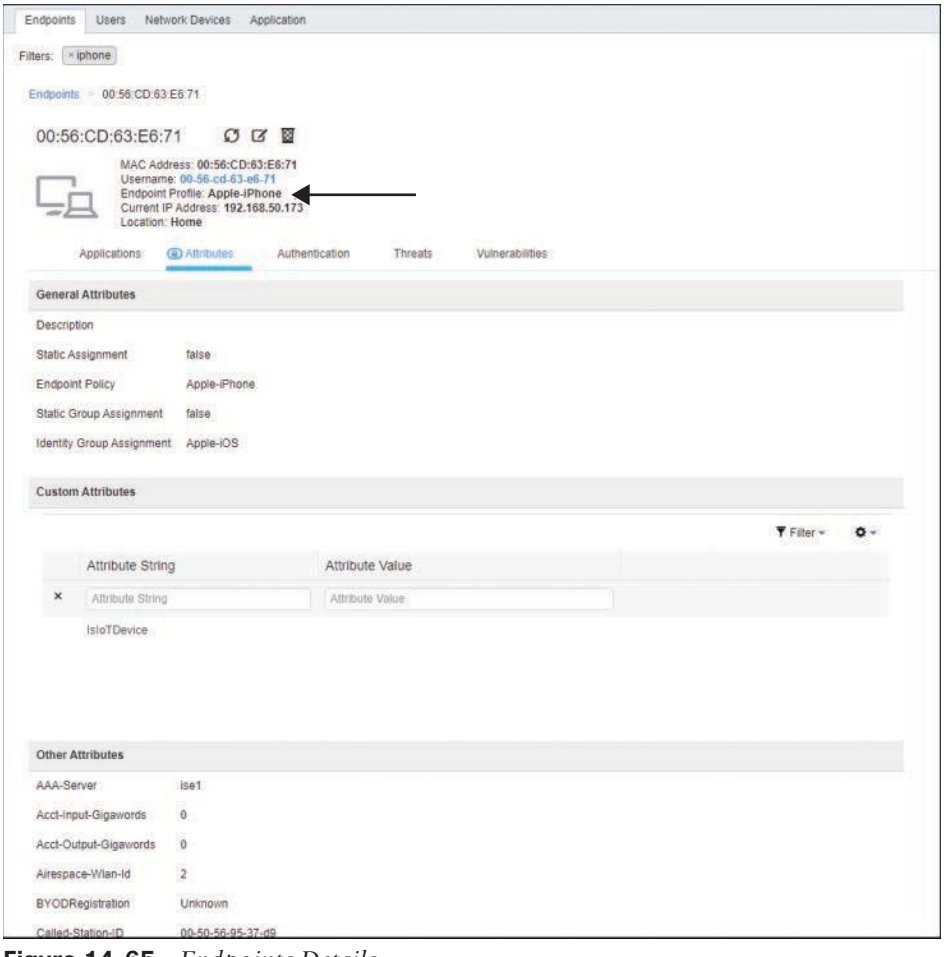


Figure 14-65 *Endpoints Details*

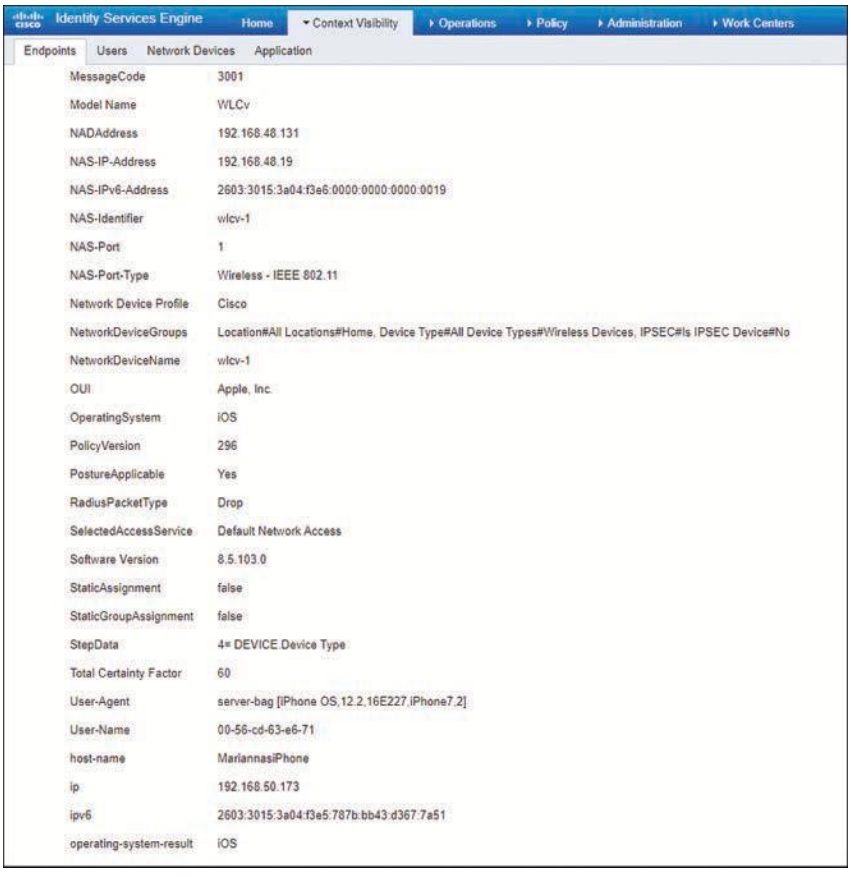


Figure 14-66 *Details in the Endpoints Details List*

Device Sensor show Commands

With Cisco switches that run Device Sensor, there is a **show** command specifically for the **device-sensor** capability in the switch: **show device-sensor cache [all | mac]**. Example 14-5 shows the output of this **show** command. While the values may not make a lot of sense to a human being, you can see in this output that Device Sensor is collecting and caching profiling data.

Example 14-5 *show device-sensor cache all Command Output*

```
3750-X#show device-sensor cache all
Device: 0050.5687.0004 on port GigabitEthernet1/0/2
-------------------------------------------------
Proto Type:Name                          Len Value
dhcp    43:vendor-encapsulated-optio     5 2B 03 DC 01 00
dhcp    55:parameter-request-list       14 37 0C 01 0F 03 06 2C 2E 2F 1F 21 F9 2B FC
dhcp    60:class-identifier             10 3C 08 4D 53 46 54 20 35 2E 30
dhcp    12:host-name                    12 0C 0A 58 59 5A 2D 42 69 6F 4D 65 64
dhcp    61:client-identifier             9 3D 07 01 00 50 56 87 00 04
dhcp    77:user-class-id                13 4D 0B 73 79 6D 75 6E 75 73 2D 62 69 6F
```

Exam Preparation Topics

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 14-2 lists these key topics and the page number on which each is found.

Key Topic

Table 14-2 Key Topics

Key Topic Element	Description	Page Number
Paragraph	Using profiling to identify devices and authorize them beyond the MAC address	405
Paragraph	Using a least-privilege strategy with profiling	407
Paragraph	Device Sensor and profiling	415
Paragraph	Logical profiles	441
Paragraph	Global and per-profile CoA types	443

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

ISE Profiler, Anomalous Behaviour Detection, logical profiles, global CoA

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is ISE Profiler?
2. What does Anomalous Behaviour Detection provide in ISE?
3. What are the 11 different profiling probes that can be enabled in ISE?
4. What are the three global CoA options?
5. What can changing the per-profile CoA options accomplish?

This page intentionally left blank



CHAPTER 15

Certificate-Based Authentication

This chapter covers the following topics:

Describe identity store options

EAP types

Network authorization enforcement

Identify issues using authentication event details in Cisco ISE

Redirect ACL

URL redirect policy

Describe features and functionality of authentication and authorization

Certificates used to be reserved for organizations with large security operations (SecOps) teams or for closed-loop systems that hid the certificate usage from the administrator, such as Active Directory. Many IT professionals hear words like *certificate*, *X.509*, or *PKI* and run away quietly, pretending they didn't hear.

In today's world of mobile devices, bring your own device (BYOD) IT, and borderless networks, certificates are becoming more and more common. Certificates don't need to be complicated; the concepts can be and are very relatable to everyday real-world aspects of your daily life.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 15-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 15-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Certificate Authentication Primer	1–5
A Common Misconception About Active Directory	6
EAP-TLS	7, 8
Configuring ISE for Certificate-Based Authentications	8–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is required for ISE to trust a client certificate?
   a. The client's private key must be imported into ISE's certificate store.
   b. The signing CA's public key must be imported into ISE's certificate store.
   c. The signing CA's private key must be imported into ISE's certificate store.
   d. The signing CA must be part of the Internet's master PKI hierarchy.
2. What determines a digital certificate's validity period?
   a. A certificate is valid any time leading up to the date listed in the Certificate Expiration field of the X.509 certificate.
   b. A certificate is always valid until it is added to the certificate revocation list (CRL).
   c. A certificate is valid any time leading up to the date listed in the Revocation Date field of the X.509 certificate.
   d. A certificate is valid during the time span between the dates listed in the Valid-From and Valid-To fields of the X.509 certificate.
3. Which revocation status mechanisms does ISE support? (Choose two.)
   a. Secure Authentication Mechanism Language (SAML)
   b. Certificate revocation list (CRL)
   c. Open Authentication 2 (Oauth2)
   d. Online Certificate Status Protocol (OCSP)
4. X.509 certificates contain a field that defines a certificate revocation list (CRL) distribution point. How does ISE leverage that certificate attribute?
   a. ISE uses OCSP to download the CRL from the defined distribution point.
   b. ISE uses an HTTP GET request to download the CRL from the defined distribution point.
   c. ISE ignores the CRL distribution point listed in the certificate.
   d. ISE leverages TAXII to download the CRL from the defined distribution point.
5. How does ISE validate proof of possession for a client's certificate?
   a. ISE encrypts data with a combination of ISE's private key and the client's public key.
   b. ISE encrypts data with a combination of ISE's public key and the client's private key.
   c. ISE sends a message to the end user, requesting a screenshot of the private key.
   d. ISE encrypts data with a combination of ISE's private key and the client's private key.

6. Which of the following accurately describes how an Active Directory user is authorized when using certificate-based authentication?
   a. When Active Directory is the certificate authority (CA), ISE sends the full certificate to the CA and cross-references it to the end user the certificate was issued to; it then returns the AD group membership and other attributes to ISE.
   b. It is not possible to perform Active Directory user authorization when performing certificate-based authentication.
   c. Cisco ISE uses CAP to identify the principal identity from the X.509 attributes and then performs the lookup in Active Directory using that identity. Active Directory returns the AD group membership and other attributes to ISE.
   d. This process requires a dual authentication. The first authentication is for the digital certificate, and then the user is prompted for his or her username and password for the Active Directory component.
7. Which of the following is the most common authentication protocol for network access when using certificates?
   a. EAP-TTLS
   b. EAP-TLS
   c. EAP-FAST
   d. EAP-GTC
8. Which of the following lists accurately describes the components required for ISE to process certificate-based authentications?
   a. ISE is capable of processing certificate-based authentications by default, and no additional configuration is required.
   b. EAP-TLS enabled in the allowed protocols, a CAP, the signing CA's public certificate added to the certificate store with the Trust for Client Authentication attribute enabled, and either CRL or OCSP configured.
   c. EAP-TLS enabled in the allowed protocols, a CAP, the signing CA's public certificate added to the certificate store with the Trust for Client Authentication attribute enabled, and an authorization rule for the extracted identity.
   d. EAP-TLS enabled in the allowed protocols, a CAP, and the signing CA's public certificate added to the certificate store with the Trust for Client Authentication attribute enabled.
9. What does the Download CA Certificate Chain link on the Microsoft CA provide for an ISE administrator?
   a. This link provides a form the administrator can use to ask the CA administrator to send the public key, including any intermediary CAs.
   b. This link enables the administrator to configure the Windows client to provide the signer's public key during the authentication process, along with the administrator's own key (hence the certificate chain).
   c. This link enables the administrator to download a PKCS file, which is a certificate chain file that contains the public certificates for the CA and any intermediate CA in the hierarchy.
   d. This link redirects the administrator to a new page where he or she can purchase the public key from the certificate authority.

10. Live Log provides a lot of information at a glance, including a brief failure reason. What should an administrator do to find a more detailed explanation of a failed certificate authentication and a possible resolution?
    a. From Live Log, navigate to Operations > Reports > Failed Authentications.
    b. From Live Log, click on the details icon to launch the authentication details report.
    c. Immediately email Aaron Woland from Cisco and ask him why this isn't working.
    d. Call Cisco TAC because if the detail is not in Live Log, it doesn't exist.

Foundation Topics

Certificate Authentication Primer

The subject of certificate-based authentication does not have to be scary. There are a few universal truths when discussing certificate-based authentication, and standard certificate authentication is the same regardless of whether the certificate authority is built by Microsoft, Cisco, Symantec, Entrust, or another vendor.

A *certificate* is a signed document that represents an identity. You can equate a certificate to a passport, a driver's license, or some other personal identification card. Just as that identification card is meant to represent you and prove you are who you say you are, a certificate provides this function for an identity.

Key Topic

This chapter discusses how certificate-based authentication actually works. When presented with a certificate, an authentication server performs the following checks (at a minimum):

Step 1. It determines whether a ***trusted authority*** has signed the digital certificate.

Step 2. It examines both the start and end dates to determine if the certificate has expired.

Step 3. It verifies whether the certificate has been revoked, based on either Online Certificate Status Protocol (OCSP) or a certificate revocation list (CRL) check.

Step 4. It validates that the client has provided proof of possession.

The following sections examine these four steps one at a time.

Determine If a Trusted Authority Has Signed the Digital Certificate

The signing of a digital certificate has two parts:

- The certificate must have been signed correctly (that is, following the correct format). If it is not signed correctly, it is discarded immediately.
- The signing *certificate authority*'s public key must be in the list of trusted certificates, and that certificate must be trusted for purposes of authentication.

Using Cisco ISE as an example, a copy of the signing certificate authority's public key must be stored at Administration > System > Certificates > Trusted Certificates, and it needs to have the Trust for Client Authentication use case.

Figure 15-1 shows ISE's certificate store and the Trusted for Client Auth field.

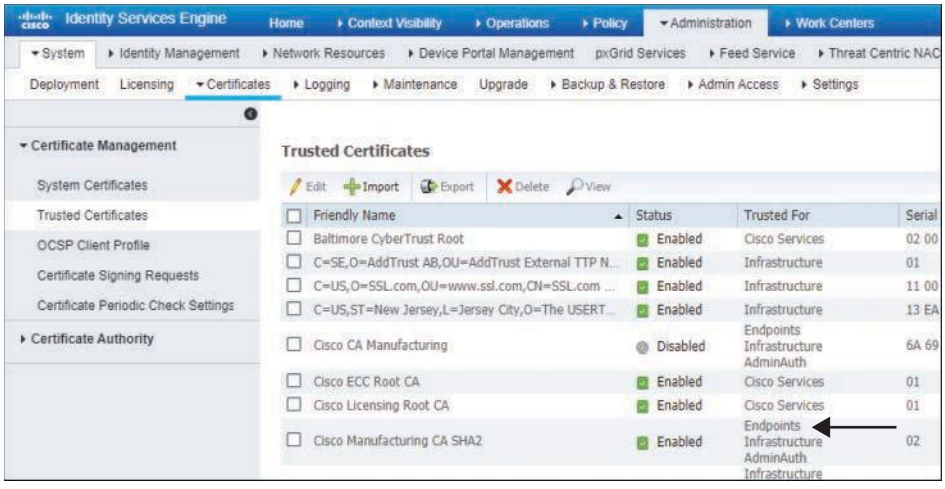


Figure 15-1 *Administration > System > Certificates > Trusted Certificates*

Figure 15-2 shows the details of a public key belonging to Cisco's hardware manufacturing division and has the Trust for Client Authentication checkbox enabled.

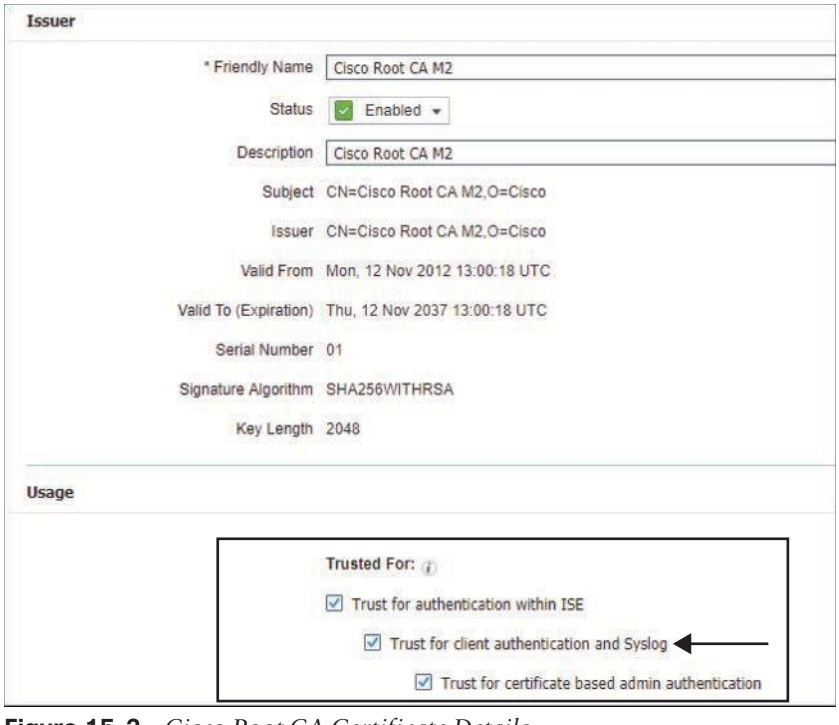


Figure 15-2 *Cisco Root CA Certificate Details*

Not only does ISE trust certificates that have been signed by this CA, it trusts those certificates for a specific use case (client authentication). If a client presents a certificate, and that certificate has not been signed by a CA that is trusted for client authentication, *the authentication fails*. Access is rejected, just as it is when someone enters the wrong password or a bartender spots a fake ID.

Examine Both the Start and End Dates to Determine If the Certificate Has Expired

Continuing to use the driver's license or passport analogy, there are usually two dates listed on this type of identifying document: the date it was issued and the date it expires. If you try to use an expired license or an expired passport, it should be denied. For example, if you are trying to get through security at an airport and present an expired driver's license, the TSA agent is not supposed to allow you through the gate. Even if the picture on the license is still clearly you, and it was obviously a valid identification card once, the signing authority (the Department of Motor Vehicles) is no longer warrantying that the document is valid.

A digital certificate also has two dates listed: the date it was issued and the date until which it is valid (that is, when it expires). The authentication server performs a check similar to the TSA agent's check and should reject any certificate that is not within its validity range. In other words, the authentication server checks to see if the certificate is valid for the date and time that the authentication request comes in.

Network Time Protocol (NTP) is very important when working with certificates. Problems often arise when time is out of sync. For example, say that a certificate is presented on January 10, 2020, at 11:11 a.m., but its valid-from value starts on January 10 at 11:30 a.m. In this example, the certificate was issued for a time 19 minutes in the future. This occurred because of a time sync issue in which the certificate authority thought it was 20 minutes later than the authentication server thought it was, and the brand-new certificate was not valid yet.

Figure 15-3 shows a certificate with the valid-from and valid-to attributes highlighted.

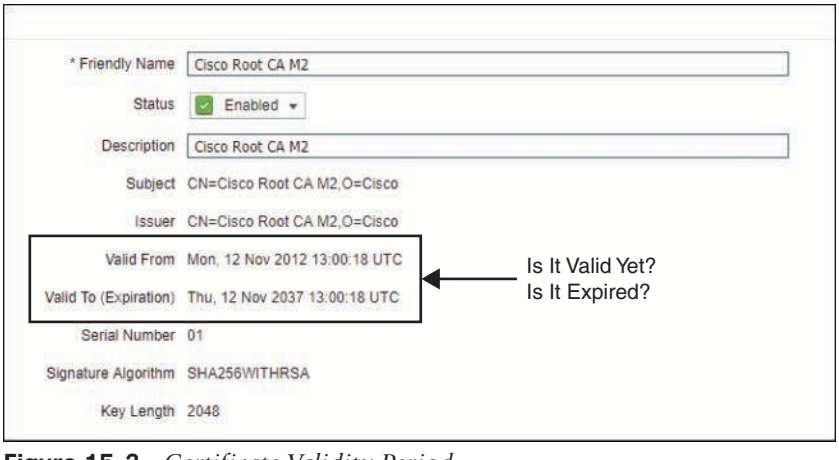


Figure 15-3 *Certificate Validity Period*

Verify If the Certificate Has Been Revoked

Pretend for a moment that you are driving down the road, and you are pulled over by a police officer. The officer asks for your driver's license and proof of insurance. You hand the officer a driver's license, which the officer immediately checks for evidence of authenticity—that is, whether it looks like a valid driver's license or a forgery. If it's not a forgery, the officer next checks the expiration date and finds that it is not expired. Next, the officer asks you to wait while he goes back to his squad car.

While in the squad car, the officer performs some authorization checks (to determine whether you are a registered owner of the car you are driving and so on). Those checks are not important for this hypothetical scenario, though. What is important is that the officer must make sure your valid driver's license has not been revoked by the Department of Motor Vehicles (DMV). A quick computer lookup into the DMV records shows that your driver's license was revoked because you committed too many violations. The cold steel of the handcuffs and the rough shove into the backseat of the squad car as you are hauled off to jail might prompt you to reevaluate your life choices.



Certificate authentication does the same sort of thing as the officer in this scenario (but not the handcuffs). Much as a police officer needs to be able to do lookups into the DMV records to check on revocation status, every certificate authority should also have a service to publish a list of certificates that have been revoked. There are two main ways to do this today:

- **Certificate revocation list (CRL):** A CRL is basically a signed list that a CA publishes on a website that can be read by authentication servers. The file is periodically downloaded and stored locally on the authentication server, and when a certificate is being authenticated, the server examines the CRL to see if the client's certificate has been revoked. A CRL could be compared to the police officer having a list of suspended drivers in his squad car.
- **Online Certificate Status Protocol (OCSP):** Using OCSP is the preferred method for revocation checks in most environments today because it provides near-real-time updates. OCSP allows the authentication server to send a real-time request (similar to an HTTP web request) to the service running on the CA or another device and checking the status of the certificate right then and there. OCSP could be compared to the officer using the computer in the squad car and doing a lookup into the DMV's database. If the certificate has been revoked, access is denied.

Figure 15-4 shows an example of the configuration screen for a trusted certificate in the Cisco ISE GUI. With each trusted certificate, you have the ability to configure where to check for OCSP or a CRL for each trusted CA basis. When ISE is authenticating a certificate that has been signed by a particular root (or its subordinates), ISE uses the configured service to look up revocation status.

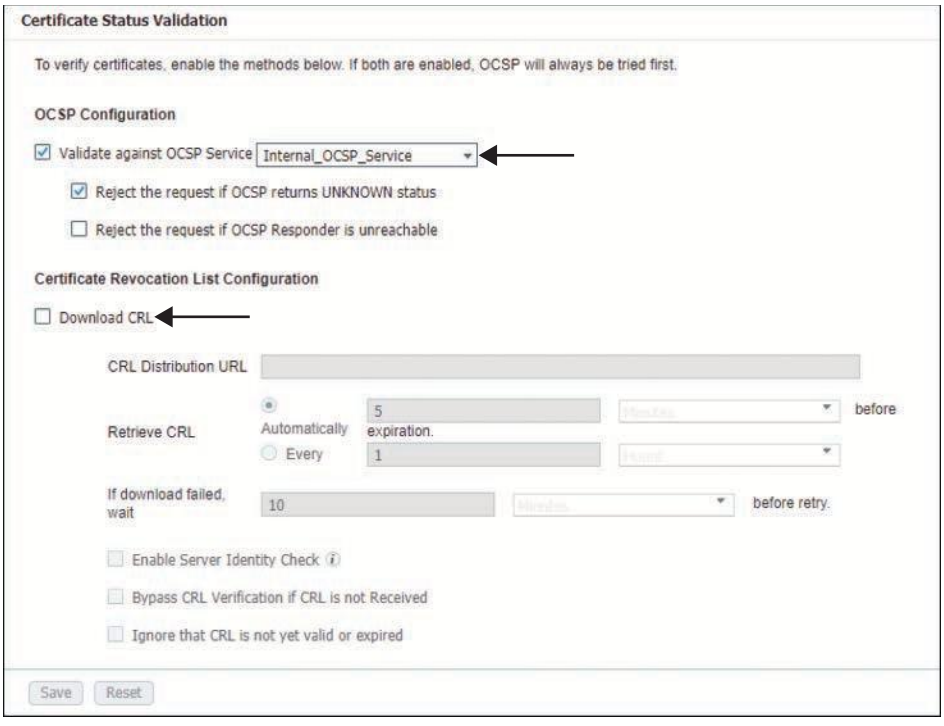


Figure 15-4 *Configuration for Certificate Revocation Checking*

It's very important for you to understand that checking for certificate revocation is an option. Notice in Figure 15-4 that neither the checkbox for CRL nor the checkbox for OCSP is selected by default; an administrator must configure the URL or the service location. It is also critical to understand what behavior will happen if the service is not available or if the status of the certificate is unknown. How does the authentication policy handle exceptions? Is it configured to "fail open" or "fail closed"?

A client's certificate has an attribute called CRL Distribution Points, which can be populated with the URI where the authentication server may locate the CRL. Figure 15-5 illustrates this certificate attribute. However, it is important to note that Cisco ISE ignores this field and uses only the configured CRL point illustrated in Figure 15-4.

Another interesting factoid about managing revocation lists is that in the previous section, where we discussed the certificate expiration, we looked at the fields Valid From and Valid To. These fields form the *validity period*, which defines the period of time that the signing CA warrants it will maintain revocation information regarding the certificate. When the certificate expires, the CA is no longer responsible for maintaining a record of its revocation status. This helps keep CRL and OCSP lists manageable sizes.

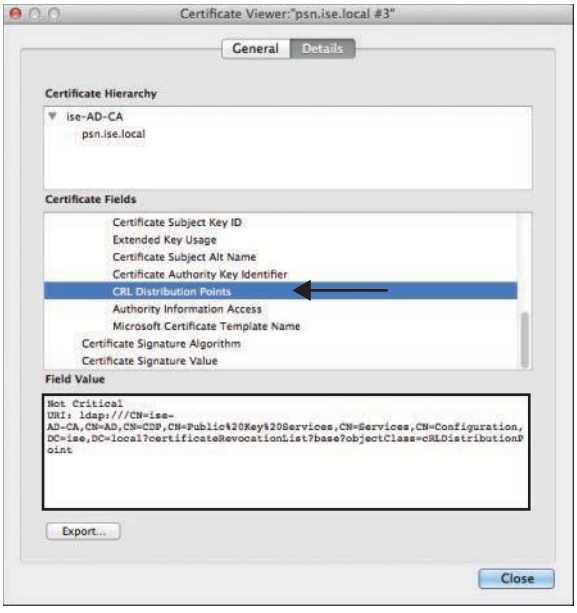


Figure 15-5 *CRL Distribution Points Attribute*

Validate That the Client Has Provided Proof of Possession

Is there a way for an authentication server to be sure the client truly owns a certificate? Think of it this way: If you are pulled over by a police officer for speeding (again), you hand the officer a driver's license, and she evaluates it. The evaluation goes like this:

- It is a valid driver's license and has been issued by a trusted signer (the state DMV), so there is no problem here.
- It has not expired yet, so that is no problem.
- The officer calls the DMV and learns that the driver's license has not been revoked. There's no problem here either.

All three of these evaluation checks pass. However, the picture on the driver's license is a picture of a woman with long flowing brown hair and hazel eyes, and you are an elderly bald man. Oops. This "valid" driver's license was not issued to you, so the proof of possession fails, and you get another ride to jail.

Certificate authentications do something similar. There might be some throwaway piece of data that must be encrypted and decrypted. Successfully encrypting and decrypting that data ensures that the client has both the public and private keys and, therefore, proves possession. This ensures that someone has not just grabbed the client's public key and tried to present it as being their own. If the client cannot provide proof of possession, the authentication fails.



Certificates (specifically, X.509 certificates) are part of a ***public key infrastructure (PKI)*** and can be used for asymmetric encryption. Each certificate is really part of a ***key pair***,

consisting of a ***public key*** and a ***private key***. The public key is a certificate that is seen by anyone and everyone. The private key is one that only the owner should ever see or possess.

What makes PKI interesting is that something that has been encrypted using the private key can only be decrypted using the public key. In addition, if something was encrypted using the public key, it can only be decrypted using the private key. Proof of possession works to ensure that the endpoint truly has the full key pair and not just a copy of the public key. The "throwaway data" is encrypted with the combination of the authentication server's private key and the client's public key (that is, the certificate sent for authentication). Then the endpoint must decrypt the data with the combination of its private key and the server's public key.

A Common Misconception About Active Directory

Microsoft Active Directory has a certificate authority built into it, and this CA can be enabled for full enterprisewide PKI. In fact, it's the single most common CA type used today. When the AD certificate authority is used to issue the certificates, many administrators think this is the same as an Active Directory authentication, but this is not the case.

What often confuses people with regard to AAA is the difference between authentication and authorization—and they do often blend together. A certificate issued by Active Directory Certificate Services is still an X.509 certificate. It goes through all the same authentication validation as any other certificate, regardless of the fact that the CA was integrated into AD. Authenticating the certificate does not involve using Kerberos, as AD authentications do, and it doesn't involve retrieving any Active Directory attributes, such as group membership. In addition, no AD permissions or rights are assigned because of the certificate itself.

However, it is possible to examine a field of a certificate and then to do a separate lookup into AD, based on that field, during the authorization phase. For example, say that a certificate with the subject Aaron is sent via EAP-TLS. The certificate is validated through the four functions of certificate authentication, and it passes all four. Now it's time for the authorization. ISE takes the certificate subject (Aaron) and does a lookup into AD for that username. This is where group membership and other policy conditions are examined and the specific authorization result is issued.

Cisco ISE uses a ***certificate authentication profile (CAP)*** to examine a specific field and map it to a username for authorization. (Figure 15-7, later in this chapter, shows an example of a CAP.)

NOTE Cisco ISE also does a courtesy check to validate whether the machine or account has been disabled in AD. If an account was disabled in AD, access is denied.

Note that the process described here is a very different process from an Active Directory authentication, which uses Kerberos (so AD logs are recorded differently). There are solutions on the market that examine AD log files and use the information obtained to help tie together usernames and IP addresses for single sign-on to web proxy servers and identity-enabled firewalls and other services. Cisco Context Directory Agent (CDA) is an example of this type of solution.

If the authentication used is certificate-based authentication (such as EAP-TLS) but the user was authorized from an AD lookup, the process is unlikely to be able to provide the right types of logging for those identity-enabled firewalls or web proxies.

EAP-TLS

More often than not, when talking about certificate authentication with wired and wireless LANs, EAP-TLS is used as the authentication protocol. Yes, it's completely possible to use EAP-GTC or maybe some other EAP type, but EAP-TLS is the mainstream protocol. It's used natively or as the inner method with tunneled EAP types such as PEAP, EAP-FAST, EAP-TTLS, and TEAP.

With Cisco ISE, you could enable the use of EAP-TLS from a preconfigured supplicant, or you could use the BYOD onboarding functionality with ISE to provision the native supplicant for Windows, MAC, iOS, and Android devices in addition to issuing that device a certificate for use with EAP-TLS.

NOTE For more on EAP methods, see Chapter 3, "Extensible Authentication Protocol (EAP) over LAN: 802.1X." For more on device onboarding, see Chapter 16, "Bring Your Own Device."

Configuring ISE for Certificate-Based Authentications

For ISE to process a certificate-based authentication, some basic configuration is required:

- The allowed protocols for the authentication must allow a protocol that uses certificates.
- A CAP must be created and used in the authentication rule.
  - The CAP should define which certificate field is the ***principal username X.509 attribute***.
  - Optionally, you add that CAP to an identity source sequence (ISS) that includes Active Directory for authorization purposes.
- Cisco ISE must be configured to trust the client certificates.
  - Those certificates must be trusted for client authentication.
  - You need to configure OCSP or CRL checking for the CAs.
- You need to configure the authorization policy.

Validate Allowed Protocols

In order to process a certificate-based authentication, the allowed protocols definition needs to permit EAP-TLS. The definition might be configured to allow native EAP-TLS itself or might be configured within a tunneled EAP type, such as PEAP, TEAP, or EAP-FAST. It is possible that all options might be allowed, regardless of whether you're using native or tunneled transport.

To validate allowed protocols, follow these steps in the ISE GUI (see Figure 15-6):

Step 1. Navigate to **Policy > Policy Elements > Results > Authentication > Allowed Protocols**.

Step 2. Click **Default Network Access** or whichever allowed protocols configuration you are using in your policy.

Step 3. Ensure that EAP-TLS is enabled natively, as well as within the PEAP and EAP-FAST tunnels.

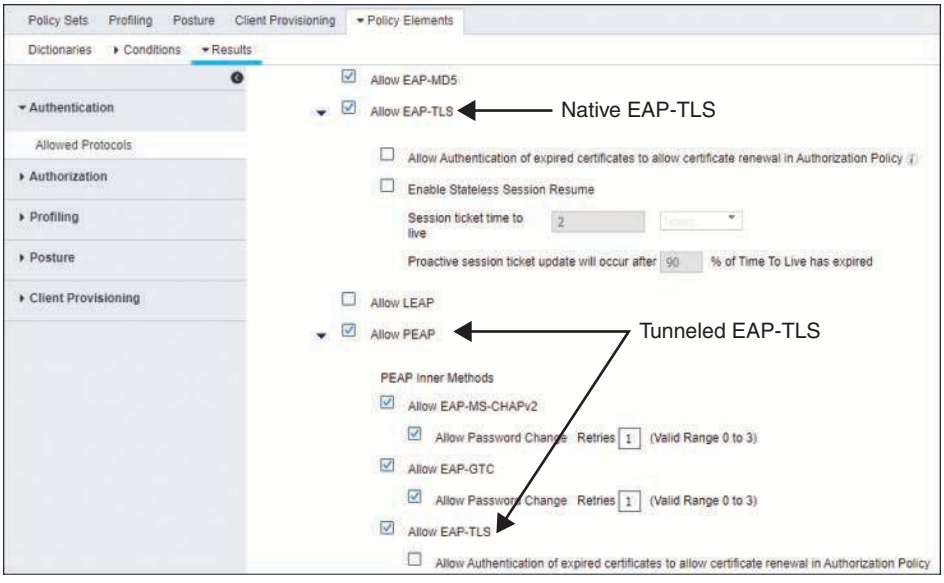


Figure 15-6 *EAP-TLS Enabled Natively and Within the EAP Tunnels*

When you are certain that EAP-TLS is allowed in this particular allowed protocols definition, you can ensure that there is a CAP for the authentication identity store.

Certificate Authentication Profile

As described in Chapter 9, "Authentication Policies," ISE uses a certificate authentication profile (CAP) as the identity source for certificate-based authentication. In reality, the authentication is simply validation of the certificate, as described in detail in this chapter.

Keep in mind that a certificate is really just a signed document that has many fields that can be read. A CAP is used to capture the identity from one of the fields of the certificate. This is known as the principal username X.509 attribute.

To verify that a CAP has been created correctly, you need to take a look at the preconfigured CAP. Follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > Identity Management > External Identity Sources > Certificate Authentication Profile**.

Step 2. Click **Preloaded_Certificate_Profile**.

Figure 15-7 shows the preconfigured CAP that is one of the "smart defaults" with an ISE installation. Notice in the figure that the principal username X.509 attribute is configured to use the Common Name attribute from the certificate's subject field as the principal username.

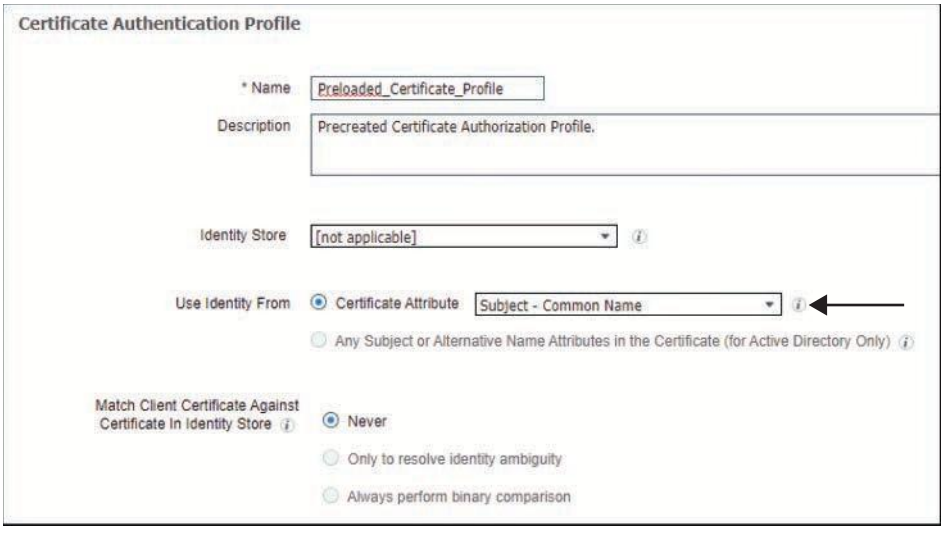


Figure 15-7 *Certificate Authentication Profile*

As described throughout this chapter, certificate-based authentication validates that a trusted authority signed the certificate and that the certificate has not been revoked. There is an additional capability to use AD or LDAP and to perform a bit-for-bit comparison of the client certificate against the copy of the certificate stored in AD. This is not a very common use case, but it does exist for those corner cases where it's required.

NOTE When you use certificate-based authentication for Windows computers that are also joined to AD and you wish to use Machine Access Restriction (MAR), a special configuration is required. Because MAR is a property of Active Directory, if the machine is using certificates (EAP-TLS) for authentication, the certificate authorization profile needs to enable the option named "*Perform Binary Certificate Comparison with Certificate retrieved from LDAP or Active Directory*" using Active Directory as the source for comparison.

Verify the Authentication Policy Is Using the CAP

You have just ensured that there is a CAP to use with the authentication policy. Now you can examine the authentication policy and ensure that the CAP is being used for EAP-TLS traffic.

In this section, you will examine the preconfigured identity source sequence (ISS) to verify that a CAP is being leveraged for the certificate-based authentication. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Click the **>** symbol for the policy set you are using. Figure 15-8 highlights this symbol for the Default policy set.

Step 3. Visually take note of the identity store in the Dot1X rule. Figure 15-9 shows the default, All_User_ID_Stores.

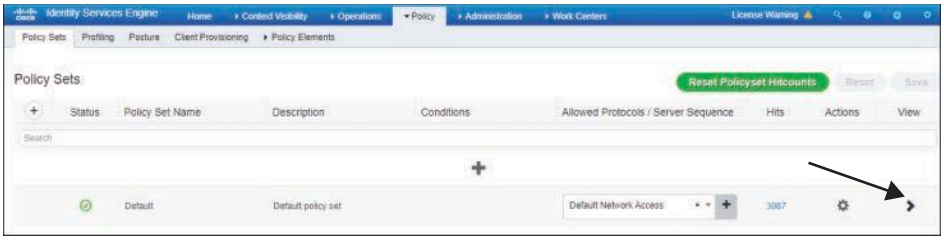


Figure 15-8 *Policy Sets*



Figure 15-9 *All_User_ID_Sources ISS Is the Identity Store for the Dot1X Rule*

Step 4. Examine the details of the All_User_ID_Stores ISS by navigating to **Administration > Identity Management > Identity Source Sequences** and edit **All_User_ID_Stores** as shown in Figure 15-10.

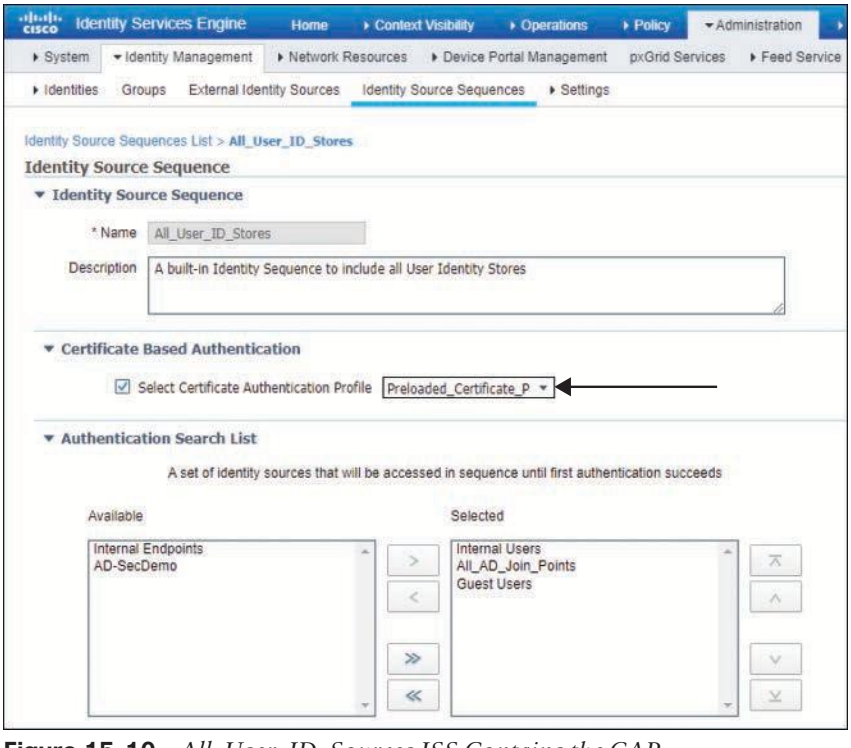


Figure 15-10 *All_User_ID_Sources ISS Contains the CAP*

15

You've now verified that the Dot1X authentication rule will permit EAP-TLS and will authenticate to an identity source sequence that includes the CAP. The authentication rule has been confirmed as ready, and you can now think about the authorization policy.

Authorization Policies

With the CAP, you specified which field to use for the principal identity. That identity may be used in the authorization policy, just as if a user typed in an ID and a password. You can examine AD group membership or other AD attributes for that identity, and you can compare fields in the certificate itself. You may choose to authorize the session based on the signing CA, or you might assign a particular VLAN because the OU of the certificate states that it is for the HR department (for example).

To get a better understanding of certificate-based authorizations, this section shows how to create an authorization rule for employees to gain access to the network. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets > Default > Authorization Policy**. The figures show the default policy set, but you should feel free to use any policy set you want.

Step 2. Duplicate the preconfigured rule named Employee_EAP-TLS by selecting **Duplicate Above** from the cog wheel at the right side of the rule, as shown in Figure 15-11.



Figure 15-11 *Duplicating the Preconfigured Rule*

Step 3. Rename the duplicated rule **Employee_Limited_Certificate** and enable the rule, as highlighted on the left side of Figure 15-12. Take note of the preexisting authorization conditions in the SAN field for wireless 802.1X, BYOD registration, EAP-TLS, and MAC address. This is a prebuilt rule designed for BYOD deployments, and you will change some of these conditions.

Step 4. Edit the conditions by clicking the **Edit** icon, highlighted on the right side of Figure 15-12.

Step 5. In the conditions studio that appears, remove BYOD_is_Registered and MAC_in_SAN from the selected conditions by clicking the **x** in corner of each one.

Step 6. Add another new condition to match an AD group for the employees of your organization; it should look similar to the one shown in Figure 15-13, which will use the CAP to extract the username of the user out of the certificate and then examine the group memberships of that user.

Step 7. Click **Save** before closing out of the policy set.

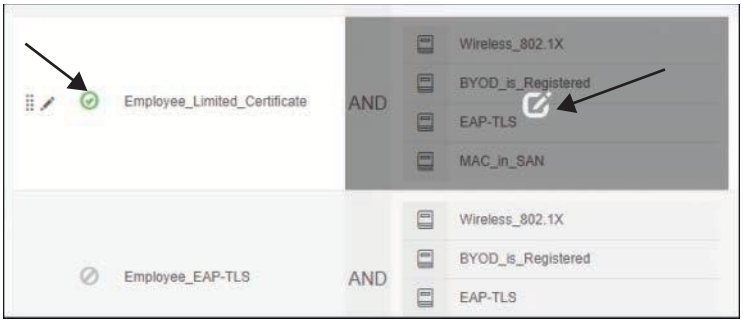


Figure 15-12 *Editing the Authorization Rule Conditions*

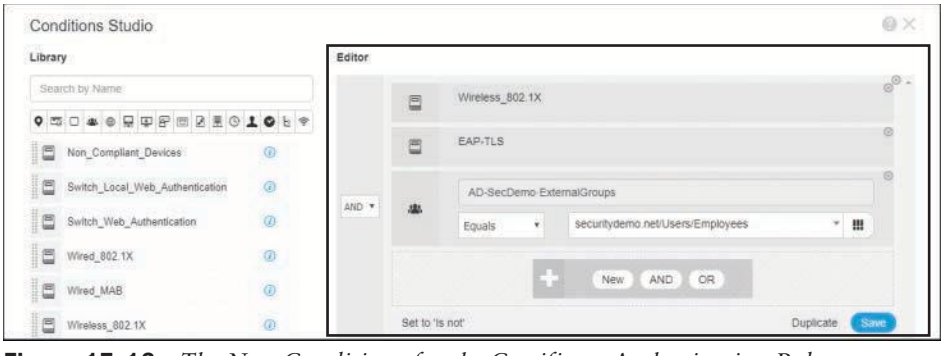


Figure 15-13 *The New Conditions for the Certificate Authorization Rule*

Figure 15-14 shows the Employee_Limited_Certificate rule and conditions. Your end result should look similar.



Figure 15-14 *Employee_Limited_Certificate Conditions*

You've just created a policy that uses a signed document and an X.509 certificate and that performs authorization against the AD membership of a user ID, using the Employees group from Active Directory. The CAP extracts that identity from a field in that certificate and used that extracted identity for the authorization.

You will learn more about these types of policies in Chapter 16.

Ensure the Client Certificates Are Trusted

A step that is often overlooked is ensuring that ISE trusts a client certificate. In order for a certificate to be trusted for client authentication, you must trust the CA that has signed the certificate.

Key Topic

Keep in mind that X.509 provides for a hierarchy that enables scale. A certificate authority may actually belong to a full tree of CAs, all stemming from the root CA. Figure 15-15 illustrates this concept with a root CA that has two subordinate CAs. Below the subordinate CA

is a certificate holder (a tablet). When ISE is presented with the certificate, it should trust all the CAs in the path. In the example shown in Figure 15-15, ISE should trust the subordinate CA and the root CA.

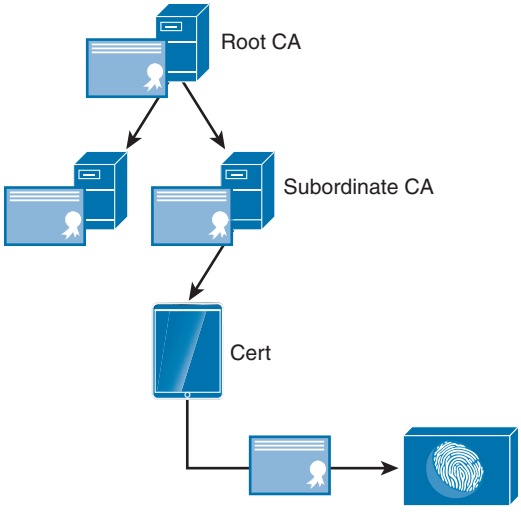


Figure 15-15 *X.509 PKI Hierarchy*

When trusting a certificate authority, keep in mind that each of the CAs in the signing hierarchy must also be trusted.

Import the Certificate Authority's Public Certificate

To trust a certificate authority, you start by downloading the CA's public certificate and adding that certificate to the certificate store under Administration > System > Certificates > Trusted Certificates. (This process is covered in detail in Chapter 8, "Initial Configuration of Cisco ISE.")

Figure 15-16 shows a simple example of connecting to the website of a CA to begin the download process. (The URL for a Windows CA is http://[fqdn-of-CA]/certsrv/.)

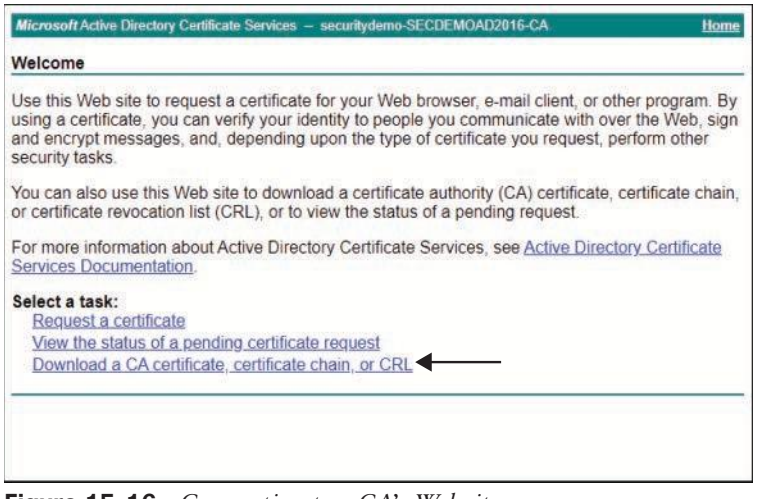


Figure 15-16 *Connecting to a CA's Website*

Figure 15-17 shows a simple example of downloading the public certificate from a Windows CA.



Figure 15-17 *Downloading a Public Certificate*

NOTE When possible, avoid using certificate-chain files (which have a .pkcs extension). Always prefer Base64-encoded files instead of DER-encoded files to prevent the strange behavior that occurs with some Android clients.

Once you have downloaded the public certificates of your root CA and all the intermediate CAs, you can add them to ISE. In the ISE GUI, repeat these steps for the root CA and each intermediate CA:

Step 1. Navigate to **Administration > System > Certificates > Trusted Certificates**.

Step 2. Click **Import**.

Step 3. Click **Choose File**.

Step 4. Select the downloaded public certificate.

Step 5. Provide a friendly name to help identify the certificate in the list.

Step 6. Select **Trust for Client Authentication or Secure Syslog Services**.

Step 7. Provide a description.

Step 8. Click **Submit**.

Figure 15-18 shows a public certificate being imported from the root CA.

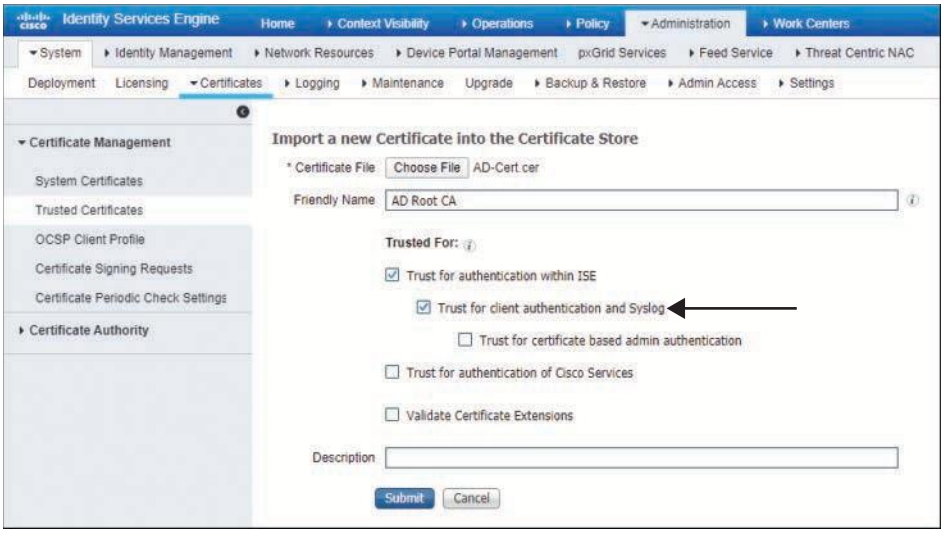


Figure 15-18 *Importing a Public Certificate*

Configure Certificate Status Verification (Optional)

At this point, all certificates signed by the CA are trusted as valid through authentication. In production environments, it is often not good enough to just trust the certificate; you must also know if those certificates have been revoked, which requires the use of either a CRL or OCSP.

Within the certificate trust you just created for client authentication, you can specify the CRL or OCSP configuration. Follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > System > Certificates > Trusted Certificates**.

Step 2. Select the certificate that you just added for client authentication.

Step 3. Scroll down to **Certificate Status Validation**.

Step 4. Select **Validate against OCSP Service**.

Step 5. Choose the correct OCSP provider from the dropdown. (If you have not configured an OCSP provider yet, see Chapter 8 for details.)

Step 6. Click **Save**.

Figure 15-19 shows a public certificate being imported from the root CA.

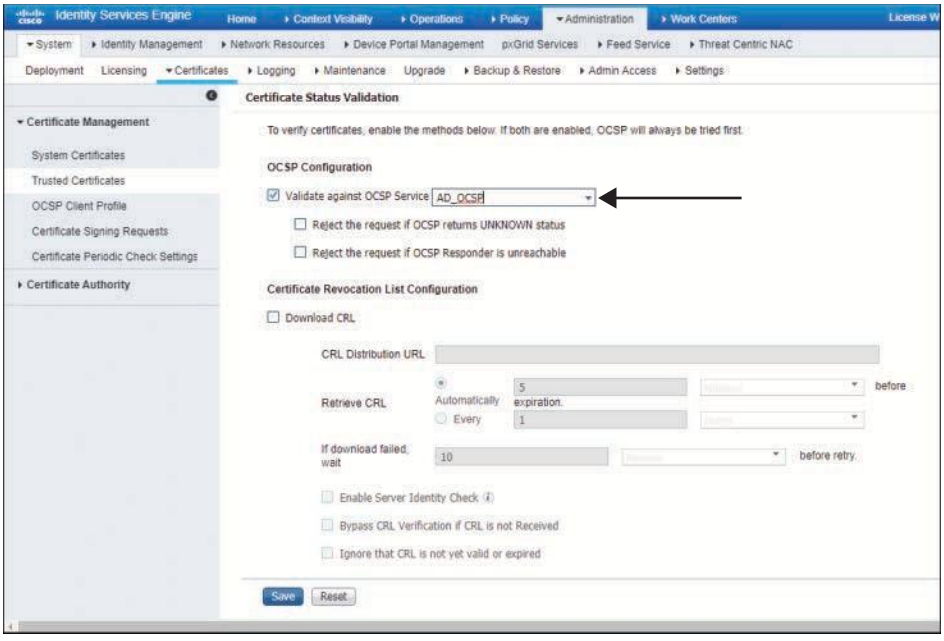


Figure 15-19 *Importing a Public Certificate – Certificate Status Validation*

You are now ready to begin fully authenticating client certificates when they are issued by the Active Directory CA. ISE is preconfigured to authenticate client certificates when they are issued by the internal ISE CA, which is part of the BYOD flow.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 15-2 lists these key topics and the page number on which each is found.

Key Topic

Table 15-2 Key Topics

Key Topic Element	Description	Page
List	Minimal checking to authenticate a certificate	463
List	Steps for revocation checking	466
Paragraph	Public/private key pairs and asymmetric encryption	468
Paragraph	The X.509 PKI hierarchy	475

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

certificate, trusted authority, certificate authority, public key infrastructure (PKI), key pair, public key, private key, certificate authentication profile (CAP), principal username X.509 attribute

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. A certificate file with a .cer extension may be encoded in different ways. Which encoding method should you use when possible?
2. A client is issued a certificate from a signing CA, which is subordinate to a node CA, which is subordinate to the root CA. Which public certificate or certificates must be added to ISE's trusted certificate store?
3. What is the URL for downloading the certificate of a Windows CA?
4. What portion of a certificate key pair should never leave the host?
5. What configuration object does ISE use to compare the fields of a certificate and extract information for use in authorization policy?

This page intentionally left blank



CHAPTER 16

Bring Your Own Device

This chapter covers the following topics:

- Implement BYOD access
- Describe elements of a BYOD policy
- Device registration
- My Devices portal
- Describe supplicant provisioning

In January 2007, Steve Jobs made a world-altering announcement: He introduced a brand-new device called the iPhone and suggested that Apple was shooting for 1% of the mobile device market. The device had a multi-touch screen interface, boasted a real browser instead of the cut-down versions that had been used on mobile devices up to that point, and changed the game for what users expected from their mobile devices. In January 2010, the iPad was released. In June of that same year, Aaron Woland was at Cisco Live! presenting a session on network access control (NAC).

In that NAC session, Aaron asked the audience members a question: Which of their organizations would allow users to bring in iPhones and iPad-type devices and connect to the corporate network for purposes of doing work from those devices? There were about 600 attendees in the room, and only around 10% of the attendees said yes.

In June 2012, just 2 short years later, when Aaron asked the same question about allowing personal devices, only 10% of respondents said their organizations would not allow personal devices. What a difference 2 years can make! Today, bring-your-own-device (BYOD) is an absolute reality.

As shown in Figure 16-1, today we have not just BYOD but choose-your-own-device (CYOD) and even a bring-your-own-app (BYOA) models.

Employees demand to use the devices that make them more productive, with native applications running on those platforms that provide the user experience they are used to and love. BYOD introduces a whole new paradigm for security, requiring the identification of the user, device, the location of the user, and more.

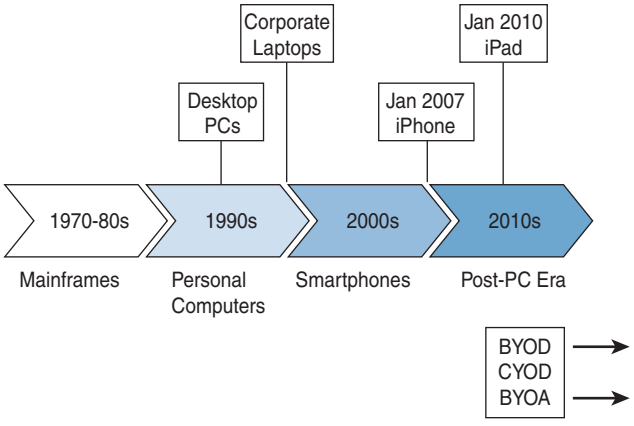


Figure 16-1 *BYOD Timeline*

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 16-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 16-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
BYOD Challenges	1
Configuring NADs for Onboarding	
Onboarding Process	1, 2, 4, 6
ISE Configuration for Onboarding	5, 10
BYOD Onboarding Process Detailed	3, 9
Verifying BYOD Flows	8
MDM Onboarding	
Managing Endpoints	7
The Opposite of BYOD: Identify Corporate Systems	

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What is the process of onboarding as it relates to BYOD?
   a. It is a form of torture used in military interrogations.
   b. It involves preparing an endpoint for network access with supplicant configuration and possibly certificate provisioning.
   c. It is a process in which an IT department pre-stages an endpoint for corporate use before issuing the endpoint to an end user.
   d. It involves preparing an endpoint for network access by preconfiguring an installation package that the end user runs with administrator privileges to configure the endpoint.
2. With a single SSID model for BYOD onboarding, how does the supplicant begin using the new certificate-based credentials?
   a. The endpoint continues to use the initial credentials until the next reauthentication interval.
   b. ISE sends a CoA-DM, causing a new authentication.
   c. ISE sends a CoA-Reauth, causing a new authentication.
   d. The endpoint continues to use the initial credentials until the endpoint is disassociated from the network and reassociates.
3. With dual SSID onboarding, what stops a guest user from receiving a certificate and a supplicant profile?
   a. It is hard-coded in ISE to not permit a guest user to enter the provisioning flow.
   b. This is a configurable option; nothing prevents guests from receiving the certificate and supplicant profile.
   c. This is a configurable option, based on the authorization result given to the user.
   d. The client provisioning policies can permit guests to enter the provisioning flow.
4. The same ACL may be used for all endpoints being onboarded. However, the security of the ACL needs to be relaxed for Android devices. Why?
   a. Google just feels it is so special that Android devices require special access to keep up.
   b. Androids require access to the local app store in ISE.
   c. Android is inherently an insecure operating system and needs a less secure ACL.
   d. Android devices require access to their app store to download and execute Cisco's Network Setup Assistant app.
5. What are an ISE administrator's options for dealing with endpoints that are not supported by the BYOD onboarding process?
   a. Cisco ISE rejects an authentication from any endpoint that cannot go through the onboarding process.
   b. The administrator has configurable choices to deny access to any nonconfigured endpoint that reaches the supplicant provisioning flow or to leave it in the current authorization state.
   c. Cisco ISE automatically permits access to any device that can't be onboarded.
   d. Once the BYOD onboarding flow is enabled, every device must be onboarded. Custom templates make it possible to push profiles to any device that is not natively supported.

6. From where does an iOS-based device download the iOS Setup Assistant?
   a. From the Apple App Store
   b. Nowhere; iOS uses the native OTA functionality
   c. From ISE directly
   d. From the Cisco App Store
7. Where would end users add or remove their own devices from ISE?
   a. Endpoints dashboard
   b. My Devices portal
   c. BYOD portal
   d. MDM portal
8. Which of following lists most accurately describes the portions of BYOD onboarding that can be verified in the RADIUS Live Logs tab?
   a. Entries exist for the initial authentication, the CoA, and the final authentication.
   b. Entries exist for the initial authentication, successful launch of the NSA app, and final authentication.
   c. Entries exist for the initial authentication, successful endpoint registration, download of the NSA app, and final authentication.
   d. Entries exist for the initial authentication, successful endpoint registration, CoA, and final authentication.
9. From where do Windows and macOS endpoints download their setup assistant apps?
   a. Windows downloads the NSA app from the Microsoft App Store. macOS uses the native OTA.
   b. Neither Windows nor macOS uses NSA; they use native capabilities instead.
   c. Windows uses native capabilities, but macOS uses a Java applet downloaded from the Client Provisioning portal.
   d. Windows and macOS use an application that is downloaded from the client provisioning portal hosted on ISE.
10. In which of the following locations does an ISE administrator determine which native supplicant profile to send to a client, based on any number of attributes, including operating system?
    a. **Policy > Onboarding**
    b. **Policy > Client Provisioning Policy**
    c. **Policy > Policy Elements > Results > Client Provisioning**
    d. **Policy > BYOD**

Foundation Topics

BYOD Challenges

Because user identity is typically based on a single identity credential, an IT department does not have the ability to create and enforce a rigorous secure access policy. Although a user might be authorized, the device, location, time, and access medium can prompt company policy violations or even regulatory compliance violations that cannot be adequately detected or enforced.

This chapter focuses on the technical challenges of providing a secure BYOD access model. Say that a user buys a new tablet or device and wants to connect it to the corporate Wi-Fi to be productive on that consumer device. IT has the challenge of identifying the device as a non-corporate device and providing limited access to it; this is the process of *onboarding*, and it is one of the most common challenges of BYOD.

Figure 16-2 illustrates what many companies used ISE to do in the earlier days.

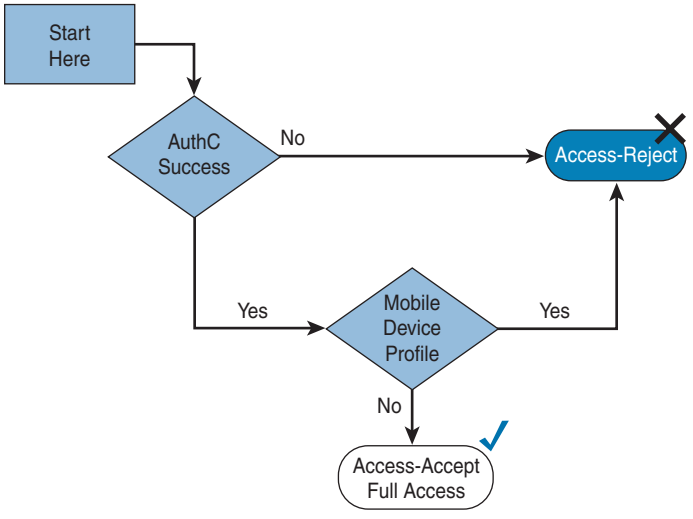


Figure 16-2 *Old-Style Policy*

Then mobile device management (MDM) solutions came into play. These MDM solutions were able to manage mobile platforms to some extent. MDM policies would ensure that devices had a minimum level of security enabled, such as encryption, remote wipe capabilities, and screen lock (PIN lock). The MDM solution would even provision certificates down to the device as well as configurations for the device's supplicant for the device to gain network access. Then ISE would provide the correct level of access for the devices, based on the certificate a device had. Figure 16-3 illustrates this type of policy logic.

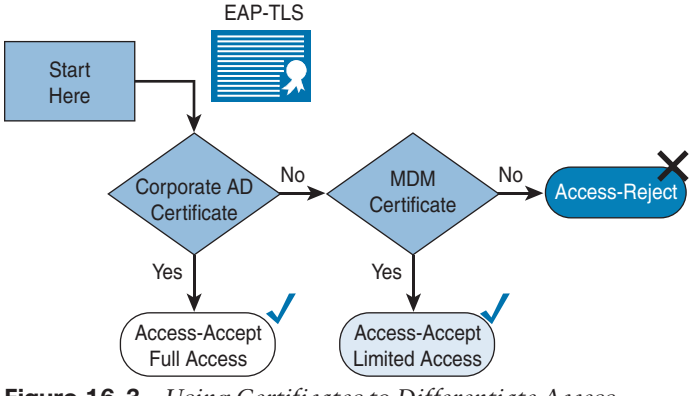


Figure 16-3 *Using Certificates to Differentiate Access*

MDM platforms typically charge per managed device. Many companies were only looking for a good way to provision certificates and configure the device supplicant to use that certificate with 802.1X. Therefore, the cost of provisioning the certificate and getting the device on the network was often prohibitive. Cisco customers were looking for a much easier and cheaper way to accomplish the onboarding aspect of network access. ISE's BYOD functionality makes onboarding much easier and more cost-effective.

Onboarding Process

In this chapter we focus on two types of onboarding:

- **BYOD onboarding**: This includes registering a device with ISE, provisioning the certificate to the device, and configuring the device's supplicant. BYOD onboarding uses the native supplicant within the device's operating system.
- **MDM onboarding**: This is the process of registering a device with a mobile device manager, installing the MDM client software, and enforcing the security policy on that device. The key to successful onboarding in a company is to make it self-service and not require manual involvement of IT.

BYOD Onboarding

As introduced in Chapter 12, "Web Authentication," ISE provides a My Devices portal, which allows users to register devices and manage the devices they have registered. It is possible for a device to simply be registered by the user, which may provide one level of authorization, such as only Internet access. It is also possible for a device to go through the full-blown onboarding and provisioning process, where the supplicant configuration is installed into the device along with the optional certificate for EAP-TLS connectivity.

Regardless of your choice to use device registration only or to use the full onboarding process, the process can be accomplished with a single SSID or dual SSID approach as a wired option.

Figure 16-4 shows the dual SSID approach, and Figure 16-5 shows the single SSID approach. A quick comparison of the approaches follows.

Dual SSID

Dual SSID onboarding is characterized as follows:

- The employee does not need to configure the supplicant on the device.
- The employee authenticates to a web form.
- The employee connects to Open SSID before the provisioning process and then must connect to the corporate SSID after the process completes. This reconnection to the second SSID may be automated or manual.

Figure 16-4 depicts the dual SSID approach.

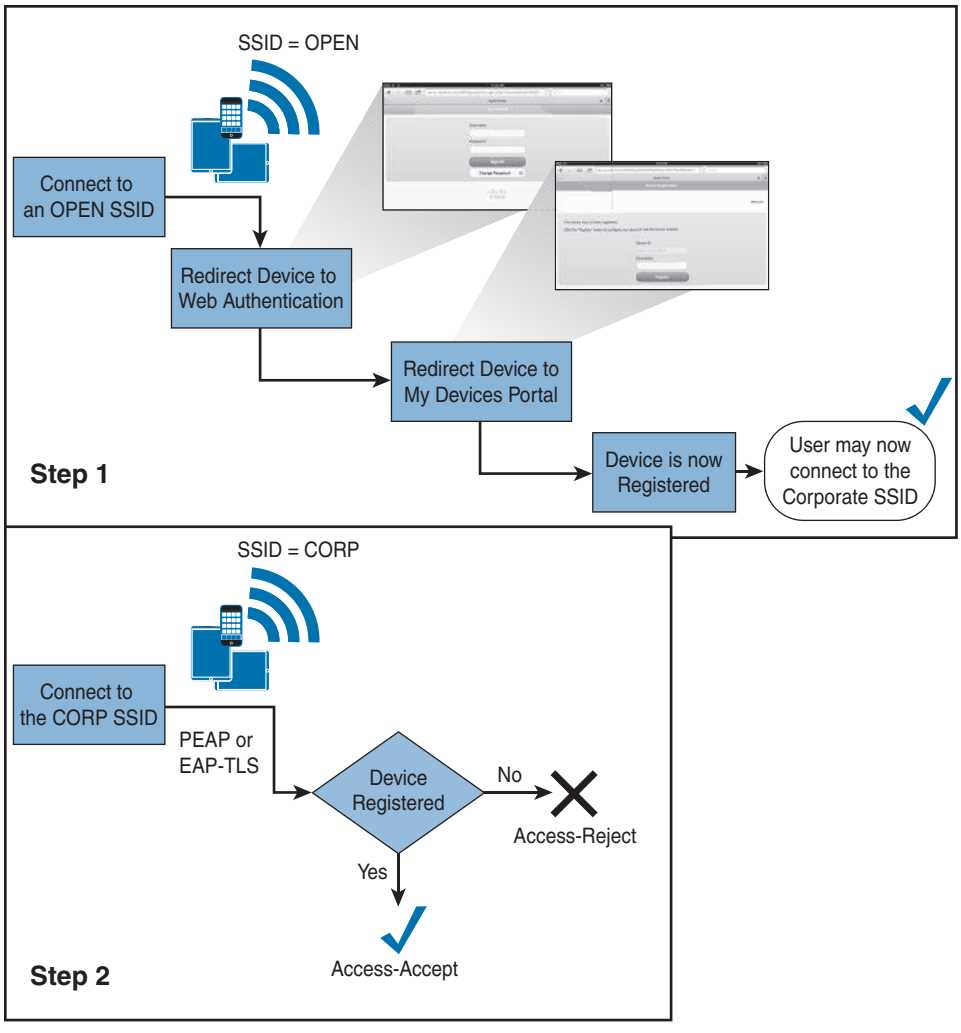


Figure 16-4 *Dual SSID Flow*

Single SSID

Single SSID onboarding is characterized as follows:

- The employee must configure the supplicant on the device to connect to the corporate SSID.
- The authentication used to connect to the corporate SSID is used for single sign-on to the onboarding and provisioning process.
- A Change of Authorization (CoA) is used to provide full access after the provisioning process, without requiring the employee to reconnect to the network.

Figure 16-5 depicts the single SSID approach.

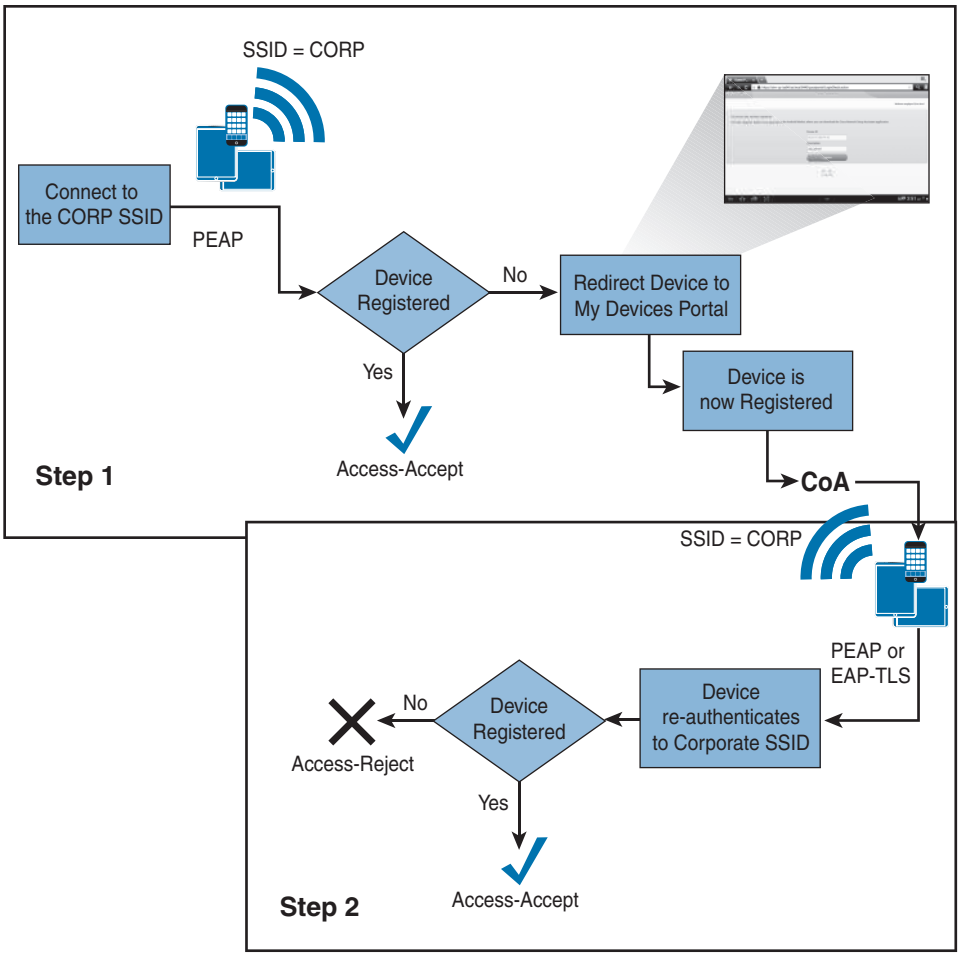


Figure 16-5 *Single SSID Flow*

Configuring NADs for Onboarding

The following sections show how to configure NADs for both dual and single SSID onboarding.

Configuring a WLC for Dual SSID Onboarding

Dual SSID onboarding uses an open WLAN configured for ISE RADIUS, and the security settings only use MAC filtering (also known as *wireless MAB*). You may have already created this network in Chapter 11, "Implement Wired and Wireless Authentication," and reviewed it in Chapter 12, "Web Authentication." The following section walks you through the process of verifying the wireless LAN controller (WLC) settings to ensure that everything is in place.

Review of the WLAN Configuration

In the Cisco Wireless LAN Controller GUI, you can edit the previously created open WLAN (named CP-Guest with SSID of CiscoPress-Guest). This SSID should have a RADIUS server defined that is configured for Change of Authorization (CoA).

The General tab of the WLAN should provide an SSID (CiscoPress-Guest in this example) and profile name, as shown in Figure 16-6.

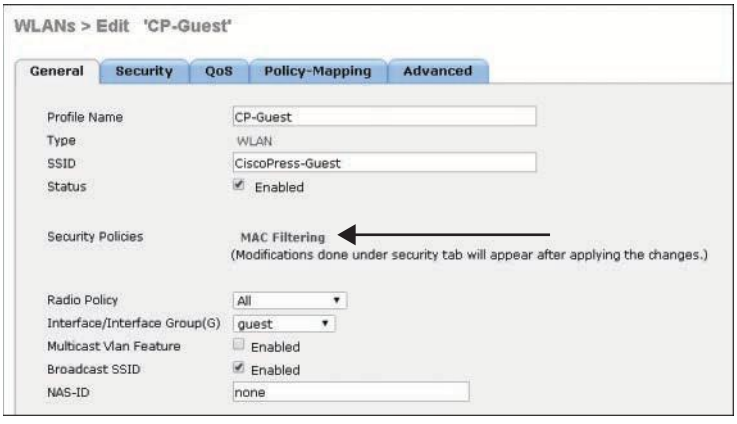


Figure 16-6 *Open WLAN General Tab*

Under Security > Layer 2, Layer 2 Security should be set to None. The MAC Filtering checkbox should be enabled, as shown in Figure 16-7.

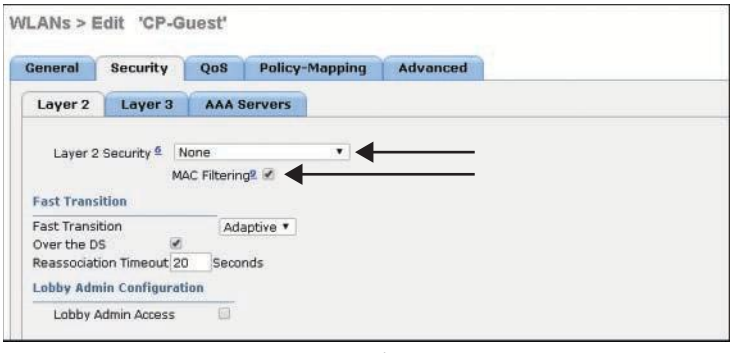


Figure 16-7 *Layer 2 Security Tab*

Under Security > Layer 3, Layer 3 security should be set to None, as shown in Figure 16-8.

For both the open SSID named CiscoPress-Guest and the secured SSID named CiscoPress, you should validate that the next three items are configured correctly. Under Security > AAA Servers, the ISE Policy Services node(s) must be selected for authentication and accounting servers, as shown in Figure 16-9.

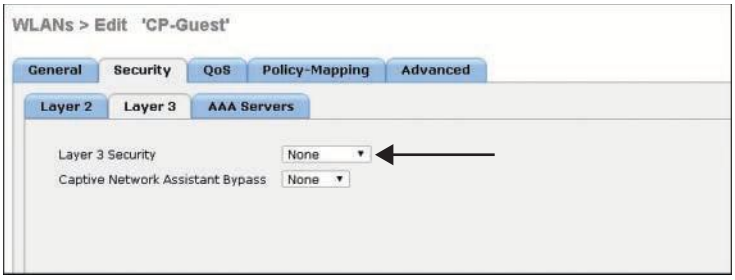


Figure 16-8 *Layer 3 Security Tab*

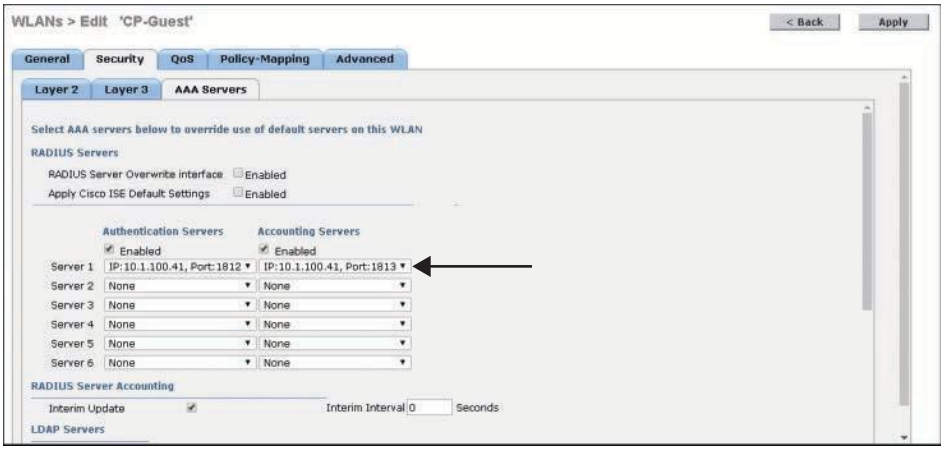


Figure 16-9 *AAA Servers Security Tab*

For the Advanced tab, Allow AAA Override must be enabled, and NAC State should be set to ISE NAC, as shown in Figure 16-10.

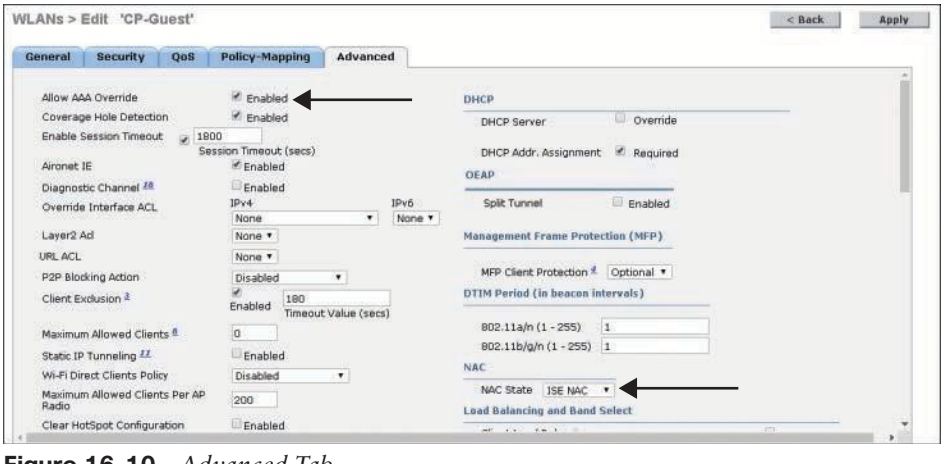


Figure 16-10 *Advanced Tab*

You should also double-check that the RADIUS server definition under Security > AAA > RADIUS > Authentication is configured to allow support for CoA, as shown in Figure 16-11.

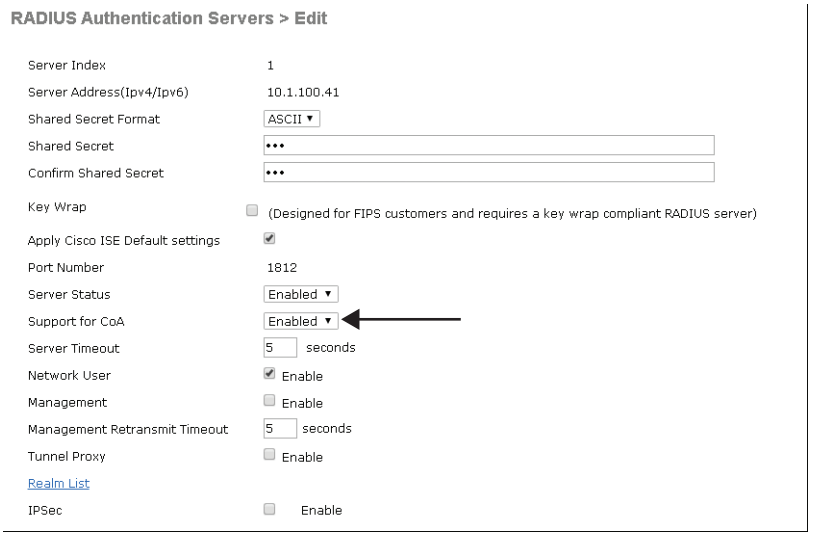


Figure 16-11 *RADIUS Authentication Servers Tab*

Verify the Required ACLs

If you followed along with the configuration in Chapter 13, "Guest Services," you should already have an access control list named ACL-WEBAUTH-REDIRECT on the switches and the wireless controllers that permits DHCP, DNS, and traffic to ISE and denies most other traffic.

When onboarding with iOS, Windows, and macOS, the endpoint only needs to communicate with ISE. iOS uses its native over-the-air (OTA) provisioning process. Windows and macOS both use a Java-based wizard that is downloaded from ISE through the device's browser. Because the communication is limited to ISE only, ACL-WEBAUTH-REDIRECT ACL can be repurposed to act as the onboarding ACL.

Key Topic

Onboarding of Android devices works differently from onboarding of devices running other operating systems. Android devices inherently do not trust apps being installed from any app store other than Google Play Store. Therefore, by default ISE is not allowed to host an app for Android devices. In order to provision an Android endpoint, you must ensure that the Android device has access to the Google Play Store to download the Cisco Network Setup Assistant app.

The Google Play Store (play.google.com) is a cloud service that could use hundreds of IP addresses. It would not be an easy task to permit and maintain hundreds of IP addresses in an ACL. Instead, you can focus on the DNS names they resolve to and use ***DNS-based ACLs*** instead. This feature was introduced in WLC Version 7.6. It involves using DNS snooping, which allows the enforcement of the DNS-based ACLs. You would create a URL list that would be bound to the ACL. If you are located in the United States, the list would likely include the following URLs:

- accounts.youtube.com
- gstatic.com

- .googleapis.com
- .appspot.com
- ggpht.com
- android.pool.ntp.org
- market.android.com
- .google.co
- gvt1.com
- captive.apple.com – Note: This URL is for Apple iOS/MacOS and not Android. However, it is best practice to put this in your URL list as it will prevent the network detection assistant from popping up which is desired since ISE cannot profile the device using the slimmed down browser it uses for detection.

In other countries, URLs might need to be added. In many cases, you can simply add an asterisk (*) to the end of the above URLs the correct top-level domain for the appropriate country.

The ACLs described here can be used for both single and dual SSID onboarding.

Create the NSP ACL on the Wireless LAN Controller

In the Cisco Wireless LAN Controller GUI, under Security > Access Control Lists, create an ACL named NSP-ACL and configure the ACL as shown in Figure 16-12.

The ACL shown in this example is permitting all traffic from the inside network to speak to the client, and it is allowing the client to communicate into the network for DNS and DHCP, as well as TCP traffic destined to the ISE servers on ports 8443, 8905, and 8084.

Access Control Lists > Edit

< Back | Add New Rule

General

Access List Name NSP-ACL

Deny Counters 0

Seq	Action	Source IP/Mask		Destination IP/Mask		Protocol	Source Port	Dest Port	DSCP	Direction	Number of Hits	
1	Permit	0.0.0.0	0.0.0.0	0.0.0.0	0.0.0.0	Any	Any	Any	Any	Outbound	0	
2	Permit	0.0.0.0	0.0.0.0	0.0.0.0	0.0.0.0	UDP	DHCP Client	DHCP Server	Any	Inbound	0	
3	Permit	0.0.0.0	0.0.0.0	0.0.0.0	0.0.0.0	UDP	DNS	Any	Any	Outbound	0	
4	Permit	0.0.0.0	0.0.0.0	0.0.0.0	0.0.0.0	UDP	Any	DNS	Any	Inbound	0	
5	Permit	10.1.100.41	255.255.255.255	0.0.0.0	0.0.0.0	TCP	8443	Any	Any	Any	0	
6	Permit	0.0.0.0	0.0.0.0	10.1.100.41	255.255.255.255	TCP	Any	8443	Any	Inbound	0	
7	Permit	10.1.100.41	255.255.255.255	0.0.0.0	0.0.0.0	TCP	8905	Any	Any	Outbound	0	
8	Permit	0.0.0.0	0.0.0.0	10.1.100.41	255.255.255.255	TCP	Any	8905	Any	Inbound	0	
9	Permit	10.1.100.41	255.255.255.255	0.0.0.0	0.0.0.0	TCP	8084	Any	Any	Outbound	0	
10	Permit	0.0.0.0	0.0.0.0	10.1.100.41	255.255.255.255	TCP	Any	8084	Any	Inbound	0	

Figure 16-12 *NSP-ACL*

Add the DNS Values to the Android ACL on the Wireless LAN Controller

The DNS ACLs on the WLC work in an interesting way. The wireless access point (AP) itself performs DNS snooping to see the response that is sent to the endpoint. The AP does not have to examine all DNS requests and responses; instead, it is configured to only consider certain domain names interesting.



The AP updates the controller to allow specific IP addresses through the ACL. Figure 16-13 illustrates the DNS snooping functionality.

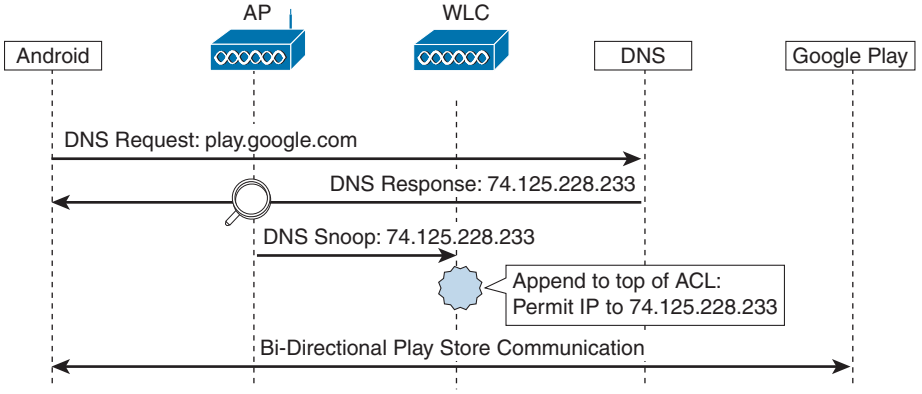


Figure 16-13 *DNS Snooping ACL*

To add a URL list to the existing ACL on the wireless controller, click the blue down arrow next to the ACL and select Add-Remove URL, as shown in Figure 16-14.

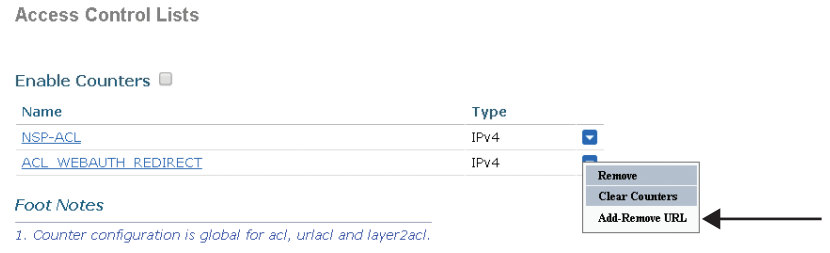


Figure 16-14 *Add-Remove URL*

When you add the URL strings, a wildcard (*) appears in front of the domain. So, for example, google.com will actually match *.google.com. Putting in google.* is a way to ensure that the DNS snooping will match for all country codes (for example, *.google.co.uk). If you are outside the United States, you should use an asterisk (*) instead of .com.

Figure 16-15 shows the URL strings for the ACL.

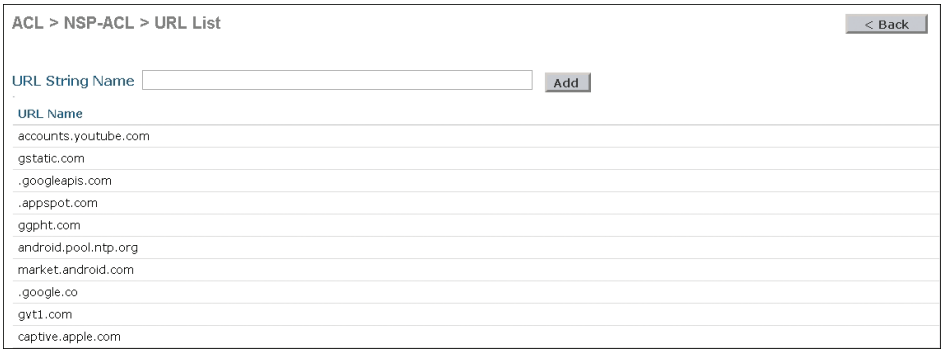


Figure 16-15 *Adding URL Strings to the ACL*

Create the ACL on the IOS-Based NADs

From global configuration mode, add a new extended ACL named NSP-ACL and configure the ACL like the one in Example 16-1. This ACL is allowing DNS to bypass redirection along with five of the known common Google Play Store address ranges. All other web traffic will be redirected to ISE.

Example 16-1 *NSP-ACL for Switches*

```
ip access-list extended Android-Marketplace
 deny   udp any any eq domain
 deny   tcp any 74.125.0.0 0.0.255.255
 deny   tcp any 173.194.0.0 0.0.255.255
 deny   tcp any 172.217.0.0 0.0.255.255
 deny   tcp any 173.227.0.0 0.0.255.255
 deny   tcp any 206.111.0.0 0.0.255.255
 permit tcp any any eq www
 permit tcp any any eq 443
```

ISE Configuration for Onboarding

Now that the NADs are prepared for the onboarding process, it's time to build the logic in the ISE policy set for both the dual and single SSID onboarding models.

Key Topic

The easier one to set up and understand is the single SSID model. This model assumes that in order for a user or an endpoint to be successfully admitted to the network, it must have authenticated with a certificate via EAP-TLS. If an authentication occurs with only a username and password (say, using MS-CHAPv2 as the inner method), then you know the device must not have been onboarded yet.

When you use this setup, when an employee shows up to work with a new mobile device and decides to try to connect to the corporate Wi-Fi, the user is prompted for a username and password. The employee then enters his Active Directory credentials (as would be expected), opens a browser on the mobile device, and is redirected to the My Devices portal, where he can begin the onboarding process. ISE uses the single sign-on credentials from the MS-CHAPv2 authentication and does not need to prompt the end user for his credentials again during the onboarding. The process is simple, quick, and intuitive to most end users today.

Key Topic

Dual SSID onboarding works differently. Instead of joining an SSID that is protected with 802.1X and WPA, the end user joins an open wireless network. That open network redirects the user to the Centralized Web Authentication (CWA) portal, where the end user enters her credentials. If the end user is a guest, she will not be permitted to enter the provisioning process. However, if the end user is not a guest user, she will be redirected to the My Devices portal, where she can begin the onboarding process. ISE uses the credentials entered into the CWA portal for the entire process, and the end user does not have to authenticate again.

The End-User Experience

To fully understand the configuration of BYOD with ISE, you need to experience the end user experience for both single and dual SSID onboarding. The information you gain in this process will aid you in understanding each policy that must be created and each choice you have to make. In order to demonstrate multiple user experiences, the following sections provide examples using Apple iOS and using Android. However, each onboarding method could be used with any of the supported clients (iOS, Android, macOS, and Windows).

Single SSID with Apple iOS Example

In this example, an Apple iOS device (an iPad) is connecting to a secured network with the SSID CiscoPress-Corp. This SSID does not match the one that you have been using so far because you have not completed the entire configuration. For this example, don't worry about following along with your own lab but just try to see and understand the demonstrated end-user experience and why it is happening:

Step 1. The user comes in to the organization with his iOS device, opens Settings, and connects to the corporate Wi-Fi, as shown in Figure 16-16.

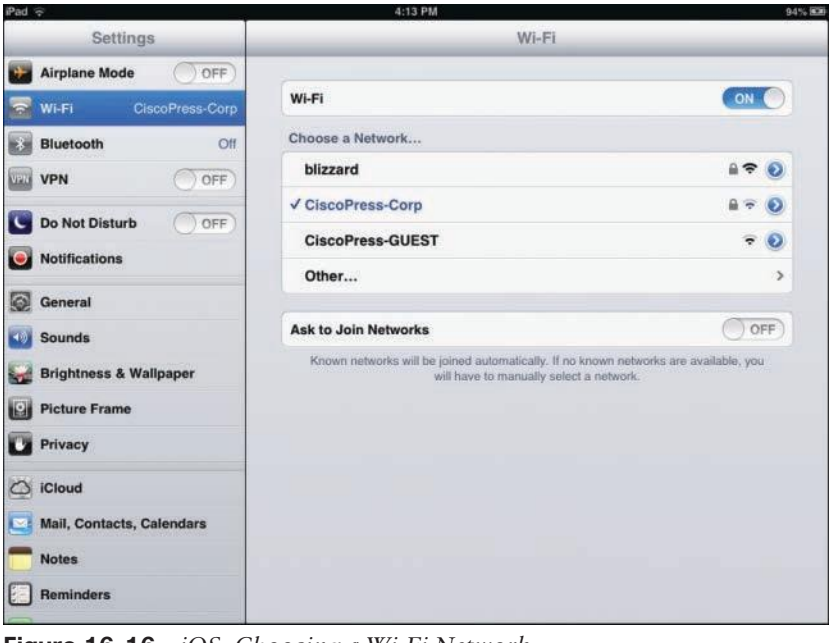


Figure 16-16 *iOS: Choosing a Wi-Fi Network*

Step 2. When the user is prompted to input a username and password, he uses his Active Directory username and password, as shown in Figure 16-17.



Figure 16-17 *iOS: Entering Credentials*

Step 3. If ISE's certificate is not signed by a trusted root, the user is prompted to accept (trust) the certificate used by ISE, as shown in Figure 16-18.

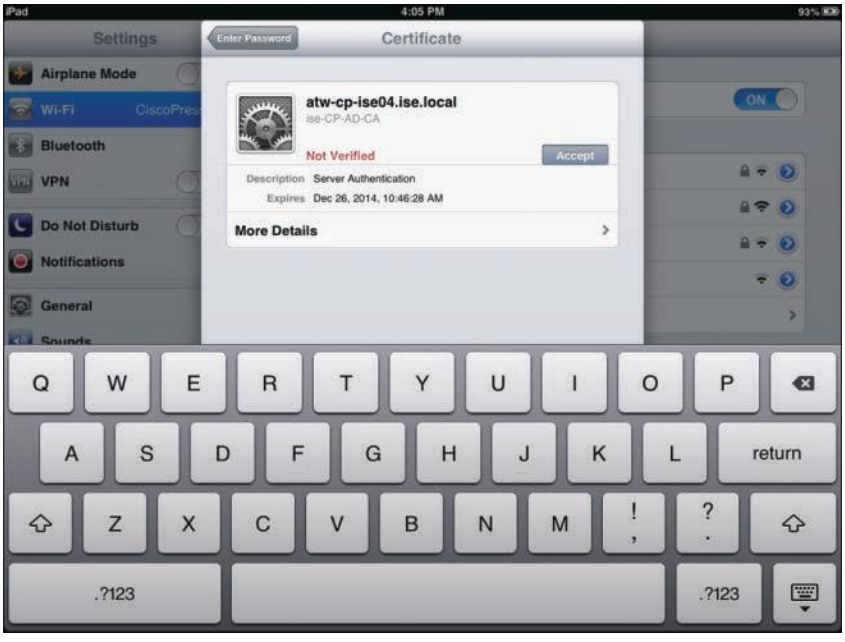


Figure 16-18 *iOS: Trusting the ISE Certificate*

Step 4. The user now sees that he is successfully connected to the corporate network, as shown in Figure 16-19. He does not know that his access is actually limited.

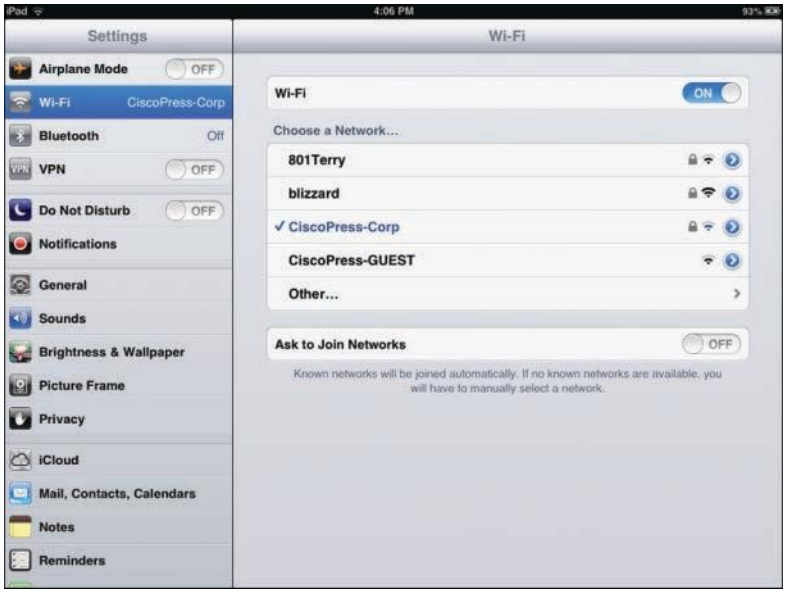


Figure 16-19 *iOS: Connected to the Corporate Wi-Fi*

Step 5. When the end user opens a web browser, he is redirected to the Client Provisioning portal, where the OTA provisioning begins.

Step 6. The first step with OTA is to send the root CA's certificate to the iOS device to be trusted for OTA provisioning, as shown in Figure 16-20.

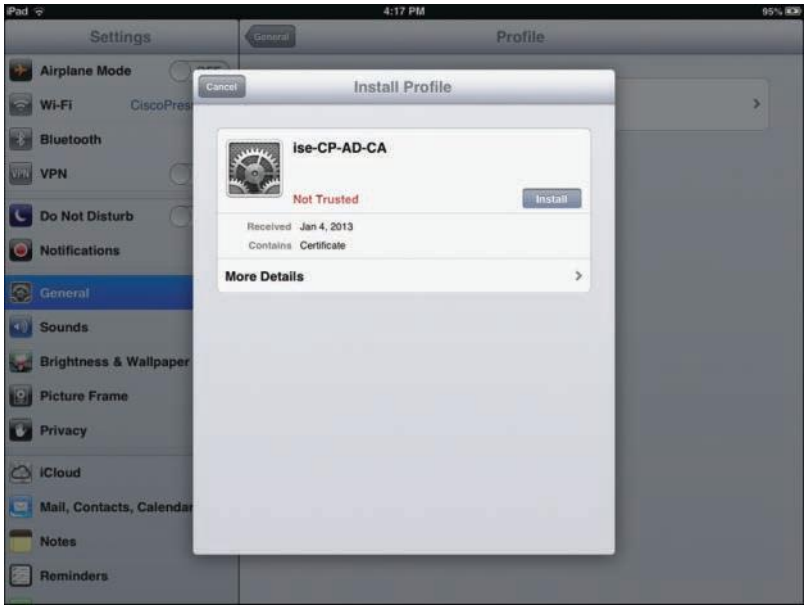


Figure 16-20 *iOS: OTA Provisioning Trusting the Root CA*

Step 7. The user clicks **Install**. A warning message is displayed about the root CA being added to the list of trusted certificates on the device, along with a warning about the profile itself (if the certificate was not in the trusted store already), as shown in Figure 16-21.



Figure 16-21 *iOS: OTA Provisioning Warning Message*

Step 8. The certificate and profile to allow OTA provisioning is successfully installed, and the user clicks **Done**, as shown in Figure 16-22.



Figure 16-22 *iOS: OTA Provisioning Profile Installed*

Step 9. The user is returned to his browser window, which displays the Device Registration page in the My Devices portal. The device's MAC address is prepopulated and non-editable. The user fills out the Description field, as shown in Figure 16-23.

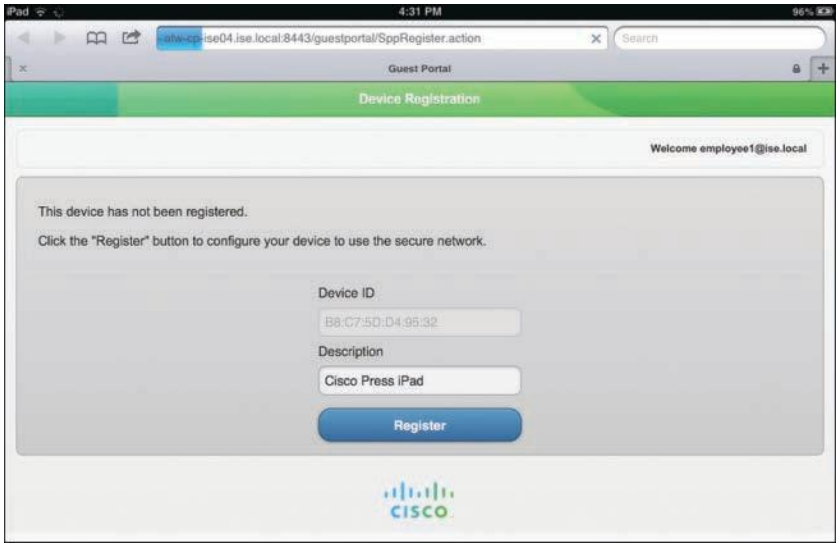


Figure 16-23 *iOS: Device Registration Page*

Step 10. The user clicks **Register**, and the Install Profile window appears, as shown in Figure 16-24.



Figure 16-24 *iOS: Install Profile Window*

Step 11. As shown in Figure 16-25, the user is warned that clicking **Install** will cause a profile to actually be installed. The user clicks **Install Now**.

Step 12. The profile begins to install, generates a certificate using ***Simple Certificate Enrollment Protocol (SCEP)***, and prepares the device to be connected to the corporate SSID, as shown in Figure 16-26 through Figure 16-29.

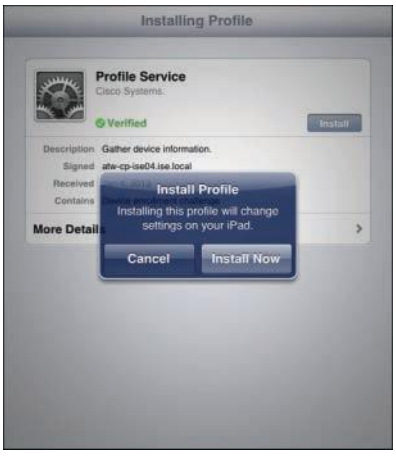


Figure 16-25 *iOS: Redundant Warning Message*



Figure 16-26 *iOS: Gathering Device Information*



Figure 16-27 *iOS: Generating a Certificate Request*



Figure 16-28 *iOS: Installing the Profile*



Figure 16-29 *iOS: Profile Is Installed*

Step 13. When the profile is installed, the user clicks **Done** and is returned to his web browser, where the success message is waiting for him (see Figure 16-30).



Figure 16-30 *iOS: Success Message*

Step 14. The user is now able to browse resources on the network, as shown in Figure 16-31.



Figure 16-31 *iOS: Final Network Access*

Dual SSID with Android Example

In this example, an Android device is connecting to a secured network with the SSID CiscoPress-Corp. This SSID does not match the one that you have been using so far because you have not completed the entire configuration. For this example, don't worry about following along with your own lab but just try to see and understand the demonstrated end-user experience and why it is happening:

Step 1. The user comes in to the organization with her Android device, opens Settings, and connects to the guest Wi-Fi, as shown in Figure 16-32.

Step 2. Because you have protected the guest Wi-Fi network by requiring a login (guest or Active Directory), the user is redirected to the web authentication page shown in Figure 16-33.

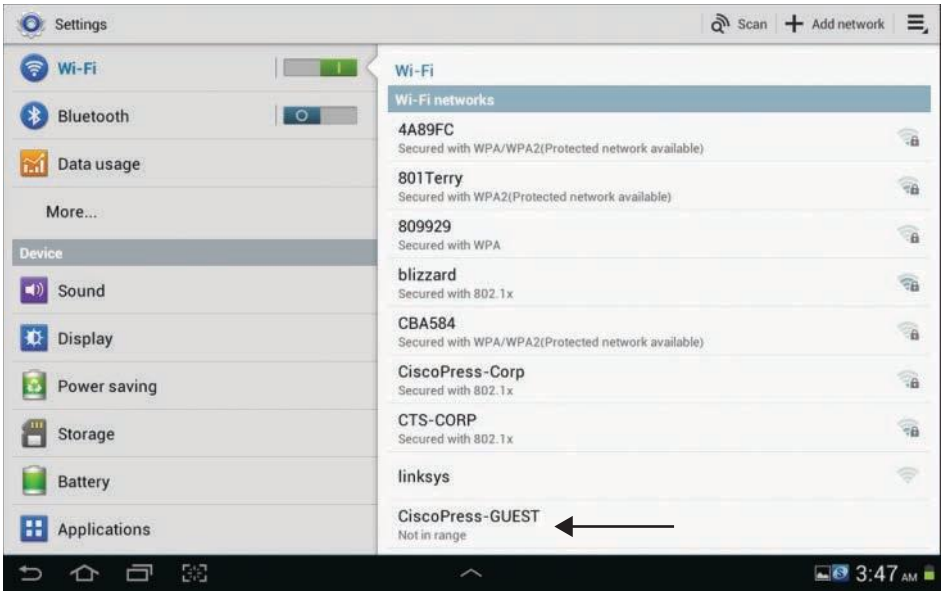


Figure 16-32 *Android: Choosing a Wi-Fi Network*

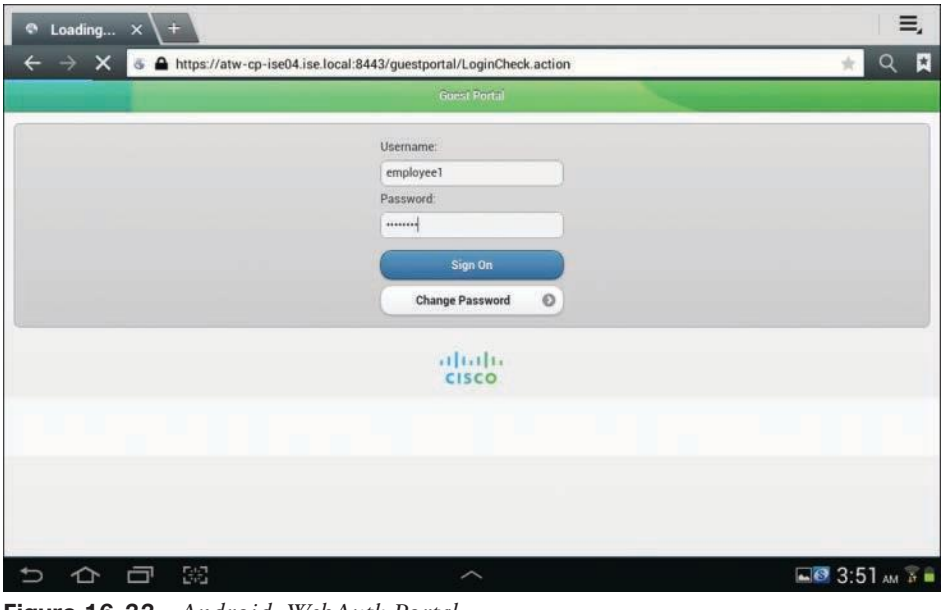


Figure 16-33 *Android: WebAuth Portal*

Step 3. After the user logs in to the WebAuth portal, she is redirected to the device registration page, where the device ID is predefined, as shown in Figure 16-34.

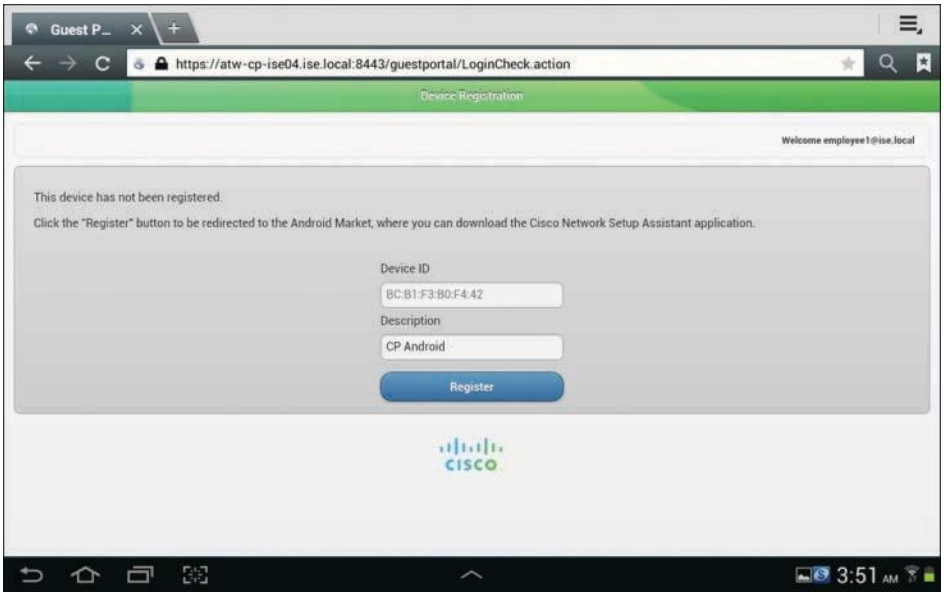


Figure 16-34 *Android: My Device Registration*

Step 4. When the user clicks **Register**, she is prompted to connect to the Google Play Store (play.google.com) and given a choice between using the Internet and using the app, as shown in Figure 16-35.

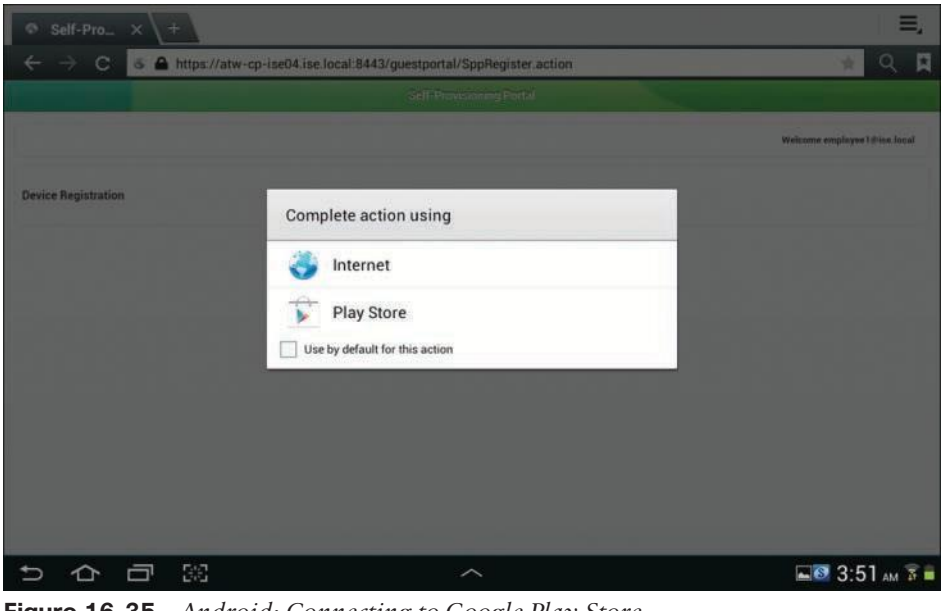


Figure 16-35 *Android: Connecting to Google Play Store*

Step 5. The user downloads the app from the Google Play Store, as shown in Figure 16-36.

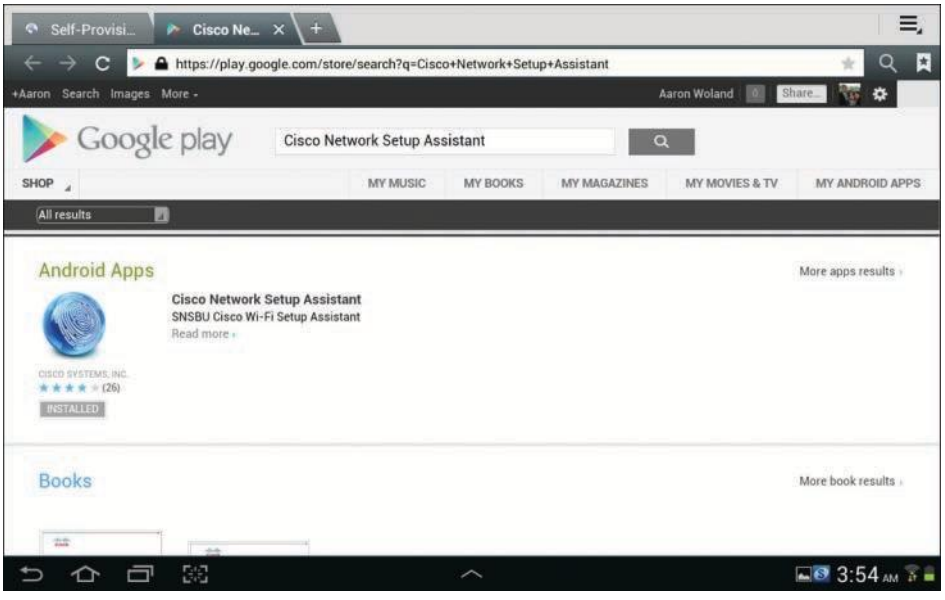


Figure 16-36 *Android: Downloading the Cisco Network Setup Assistant App*

Step 6. The user runs the app and clicks **Start**, as shown in Figure 16-37.



Figure 16-37 *Android: Running the Network Setup Assistant App*

Step 7. The Network Setup Assistant app downloads the profile from ISE, as shown in Figure 16-38.



Figure 16-38 *Android: Network Setup Assistant Downloading Profile*

Step 8. The Network Setup Assistant app automatically changes the network connection to the corporate SSID and authenticates with the new certificate by using EAP-TLS, as shown in Figure 16-39.

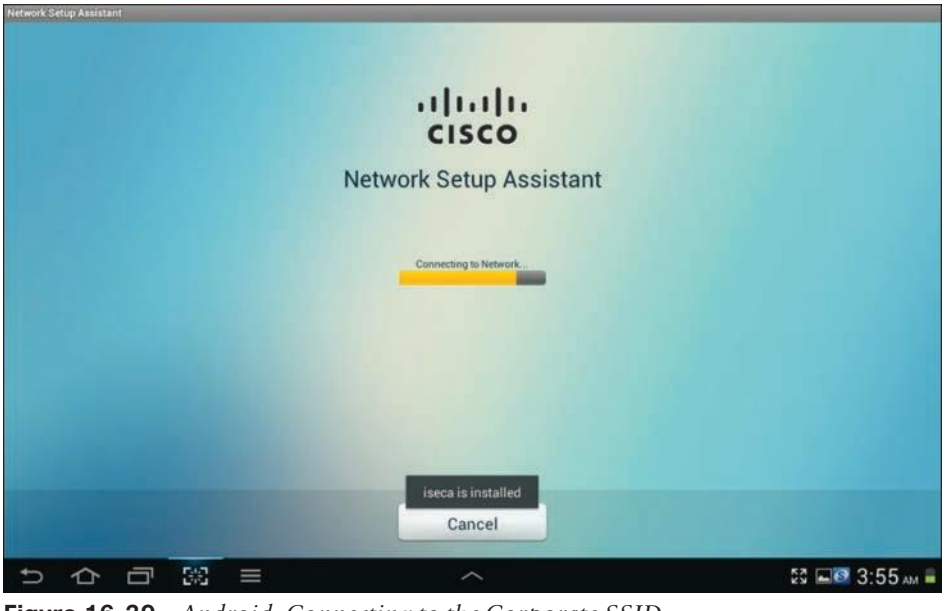


Figure 16-39 *Android: Connecting to the Corporate SSID*

Step 9. The user's Android device is now ready to be used regularly on the corporate network, as shown in Figure 16-40. The onboarding needed to occur only this one time.

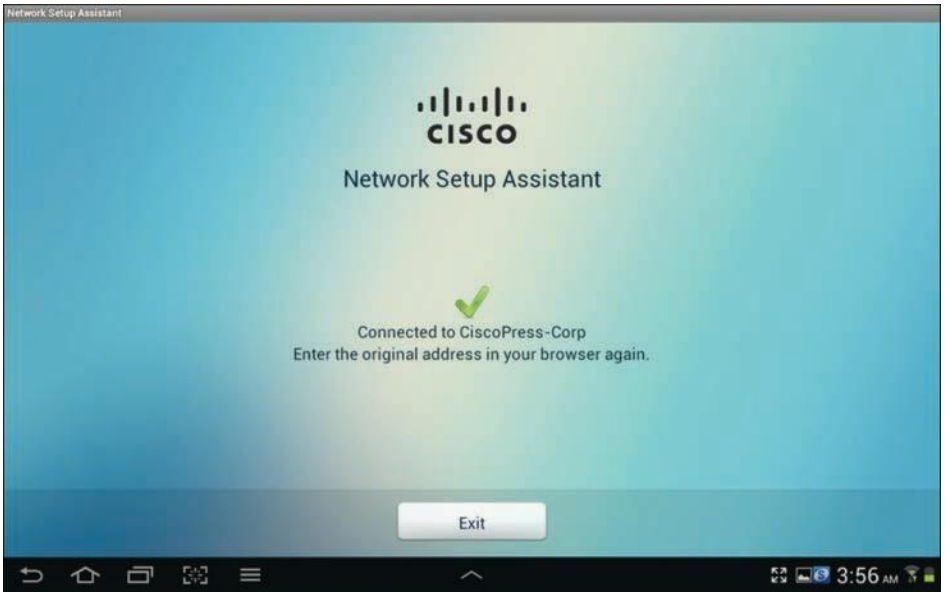


Figure 16-40 *Android: Done*

Unsupported Mobile Device: BlackBerry Example

You have now seen both single and dual SSID onboarding with supported devices. This section presents a brief example of a device that is not supported:

Step 1. A user brings his BlackBerry mobile device to work and selects the guest wireless network, as shown in Figure 16-41.



Figure 16-41 *BlackBerry: Selecting the Guest SSID*

Step 2. The web browser is redirected to the Web Authentication page, where the employee enters his Active Directory credentials, as shown in Figure 16-42.



Figure 16-42 *BlackBerry: Web Authentication*

Step 3. If the global setting is Apply Defined Authorization Policy, the user receives the message shown in Figure 16-43 and cannot gain network access (as discussed in more detail later in the chapter).



Figure 16-43 *BlackBerry: Unable to Register*

Step 4. If the global setting is Allow Network Access, the user receives a notification that he can register the device through the My Devices portal, as shown in Figure 16-44 (as discussed in more detail later in the chapter).

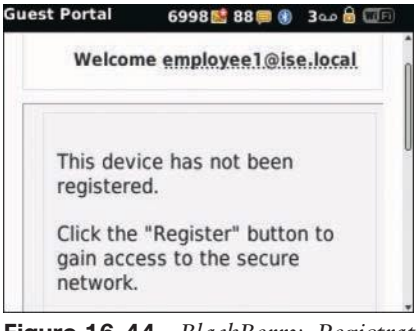


Figure 16-44 *BlackBerry: Registration Permitted*

Step 5. The user is now allowed to register the device and gain network access by manually configuring his supplicant, as shown in Figure 16-45.



Figure 16-45 *BlackBerry: Registering the BlackBerry Device*

Configuring ISE for Onboarding

In most cases, you want the end user's onboarding experience to be very straightforward and easy to follow so that no interaction with the IT department is needed. To make this process easy for the end user, you need to do quite a bit of up-front configuration. The section describes the required configuration.

The first step is to create the native supplicant profile that will be sent to an end user's device. Then you can configure the Client Provisioning portal so the correct profiles are sent for the correct operating systems. Next, you need to configure the default action that should be taken when an unsupported device attempts to connect and cannot be provisioned (as in the BlackBerry example that you just saw).

Creating the Native Supplicant Profile

The native supplicant profile is a required object for BYOD onboarding that defines the following:

- The wireless SSID
- The EAP type to use (PEAP or EAP-TLS)
- The key size for certificates
- The level of wireless security
- Whether the profile applies to wired, wireless, or both

To create the native supplicant profile, follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Elements > Results > Client Provisioning > Resources**.

Step 2. Select **Add > Agent Resources from Cisco Site**.

Step 3. From the screen to download agent resources from Cisco's site, shown in Figure 16-46, select the latest versions of the clients and wizards.

Step 4. Click **Save**.

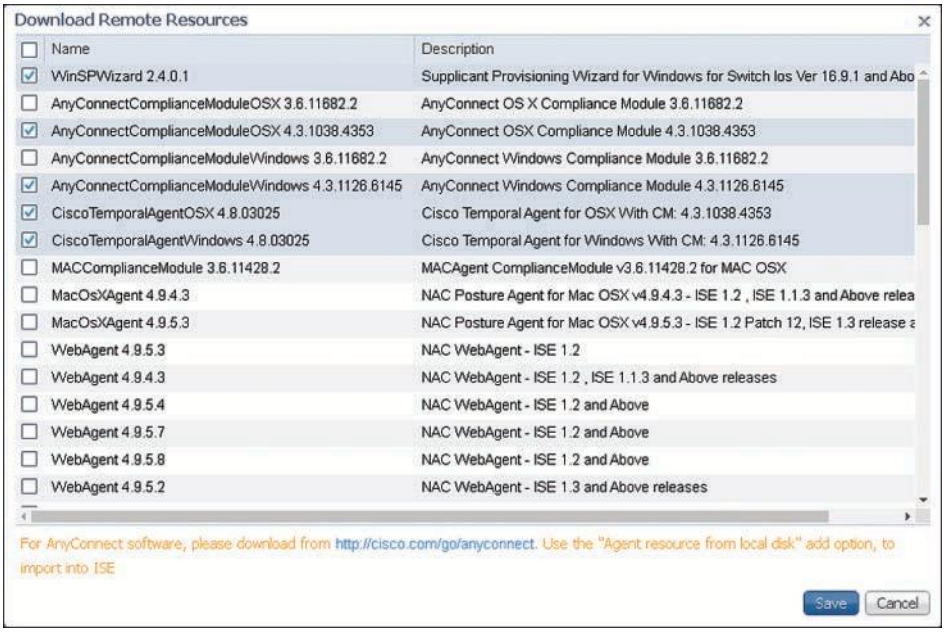


Figure 16-46 *Agent Resources from the Cisco Site*

Step 5. Select **Add > Native Supplicant Profile.**

Step 6. Name the native supplicant profile **CiscoPress-TLS**, as shown in Figure 16-47.



Figure 16-47 *Native Supplicant Profile for EAP-TLS*

Step 7. Leave Operating System set to **ALL.**

Step 8. Click **Add** to add a wireless SSID.

Step 9. Provide the SSID for the corporate wireless network.

16

Step 10. Optionally, add a URL to retrieve a proxy auto-config file and proxy information, if applicable.

Step 11. Set Security Level to **WPA2 Enterprise**.

Step 12. Set Allowed Protocol to **TLS**.

Step 13. Select **2048** as the certificate size.

Step 14. Set Certificate Template to **EAP_Authetication_Certificate_Template**.

Step 15. Under Optional Settings, set Authentication Mode to **User or Computer** and select any additional settings you want.

Step 16. Click **Submit**.

Figure 16-48 shows the completed native supplicant profile using EAP-TLS.

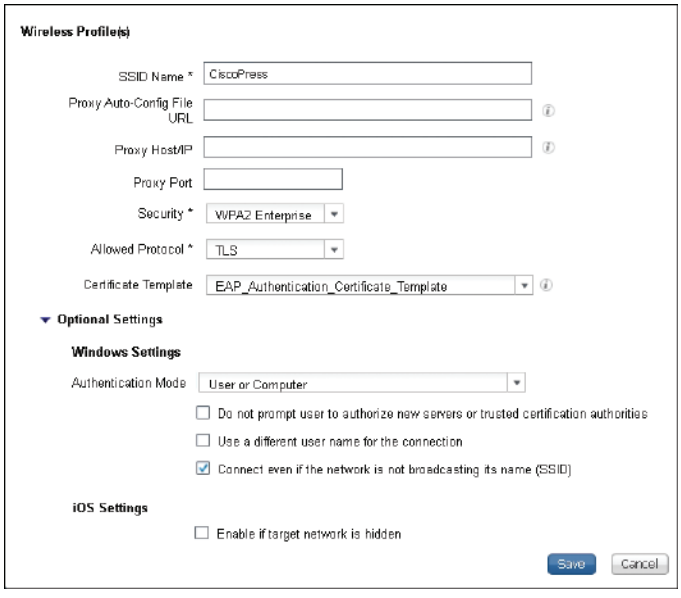


Figure 16-48 *Wireless Profile Configuration*

Configure the Client Provisioning Policy

You can configure a client provisioning policy to dictate what software and profiles should be downloaded and installed based on the operating system of the endpoint as well as a multitude of other possible attributes.

For example, you might configure a policy for Android to be provisioned for the CORP-SSID wireless network when an employee is going through the provisioning process; the same policy might provision CONTRACTOR-SSID for all vendors and contractors who are going through the provisioning process. A client provisioning policy has default rules for iOS, Android, Windows, macOS, and Chromebook endpoints. Most ISE administrators either modify the built-in rules or use them with their default settings. The following configuration example shows how to create new rules.

Follow these steps to create one client provisioning policy per OS:

Step 1. Navigate to **Policy > Client Provisioning**.

Step 2. Edit the built-in rule named **IOS** by clicking the edit hyperlink next to the rule and click on the green checkbox to disable this rule.

Step 3. Click on the downward-pointing arrow next to the IOS rule and choose **Insert New Policy Above** from the menu.

Step 4. Name the new rule **CiscoPress-IOS**.

Step 5. Set Operating System to **Apple iOS All**.

Step 6. Select **CiscoPress-TLS** as the supplicant profile from the Result dropdown.

NOTE The ISE Client Provisioning portal automatically uses the OTA provisioning process that is native to iOS for Apple iOS, so there is no need to specify that here.

Figure 16-49 shows the completed client provisioning policy rule for iOS.



Figure 16-49 *Client Provisioning Policy for CiscoPress-IOS*

Step 7. Disable the default Android rule and insert a new rule below it. Name that rule **CiscoPress-Android**.

Step 8. Set Operating System to **Android**.

Step 9. Select **CiscoPress-TLS** as the supplicant profile from the Result dropdown.

NOTE The ISE Client Provisioning portal automatically redirects Android devices to play.google.com to download the supplicant-provisioning app. There is no ability to specify a different app store at this screen.

Step 10. Disable the default Windows rule and create a new rule below it. Name that rule **CiscoPress-Windows**.

Step 11. Set Operating System to **Windows All**.

Step 12. Configure the results to use **WinSPWizard** and **CiscoPress-TLS**.

NOTE When you choose **WinSPWizard**, you are telling Windows to use the Cisco Supplicant Provisioning Wizard (which is a Java applet) to do the provisioning.

Figure 16-50 shows the completed client provisioning policy (CPP) for Windows.



Figure 16-50 *CPP Results for Windows Operating Systems*

Step 13. Disable the built-in MAC OS rule and insert a new rule below it. Name the new rule **CiscoPress-MAC-OS**.

Step 14. Set Operating System to **Mac OSX**.

Step 15. Configure the results to use **MacSPWizard** and **CiscoPress-TLS**.

Note As with Windows, with macOS you have to specify the supplicant wizard—in this case **MacSPWizard**—to configure the native supplicant.

Step 16. Click **Save**.

Figure 16-51 shows the final client provisioning policy with the four new rules configured and the disabled built-in rules.



Figure 16-51 *Client Provisioning Policy*

Configure the WebAuth

The client provisioning policy you have just created will be used with both dual SSID provisioning and single SSID provisioning. However, you must ensure that the WebAuth portal page is ready for the dual SSID flow.

You have a plethora of options when it comes to WebAuth and supplicant provisioning. For instance, it is possible to configure different web portals based on a number of attributes available from the authentication request (such as source SSID). This way, you can enable the device registration and supplicant provisioning to occur per use case if you so choose.

For simplicity, the following example uses Self-Registered Guest Portal (default). Follow these steps to configure WebAuth in this portal:

Step 1. Navigate to **Work Centers> Guest Access > Portals & Components**.

Step 2. Click on the **Self-Registered Guest Portal (default)** portal.

Step 3. Expand **BYOD Settings** for this portal.

Step 4. Check the box next to **Allow Employees to Use Personal Devices on the Network** (see Figure 16-52). The BYOD provisioning workflow appears on the right side of the settings.

Step 5. Click **Save**.

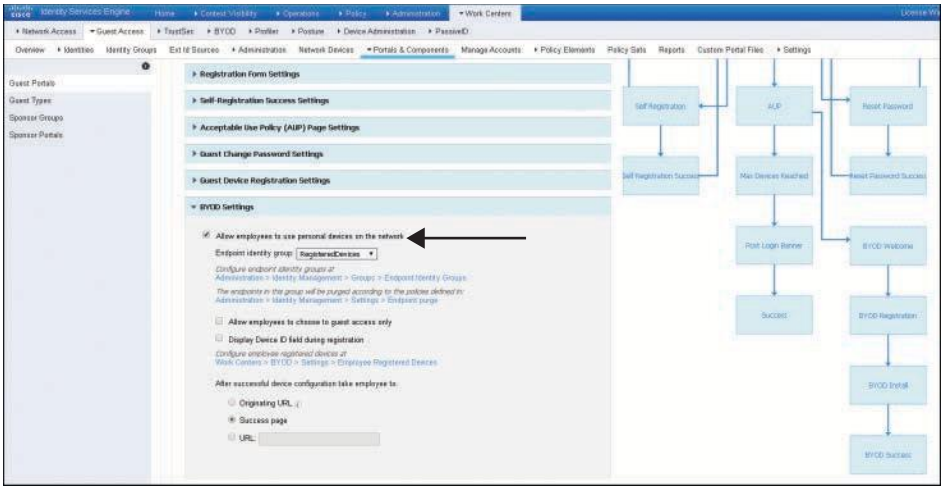


Figure 16-52 *Web Portal Configuration*

Verify Default Unavailable Client Provisioning Policy Action

ISE supports iOS, Android, Windows, Chromebook, and macOS. However, it is quite possible for an end user to attempt access with a client that is not supported by ISE Native Supplicant Provisioning (such as attempting with an older BlackBerry or a Windows Mobile device). ISE offers two options for such a situation:

- **Allow network access**: With this option, users are allowed to register their devices through the My Devices portal and gain network access without having to install and launch a native supplicant wizard. This means users have to interact and configure the supplicant in an out-of-band fashion. In other words, the users need to configure their supplicants themselves. This option may be very attractive if the end users are capable of requesting and installing their own certificates.
- **Apply defined authorization policy**: Basically, this option leaves the client in the current state—which is a state of limited access. This is the default setting.

Figure 16-53 shows the available client provisioning policy (CCP) action that can be configured under Administration > System > Settings > Client Provisioning.

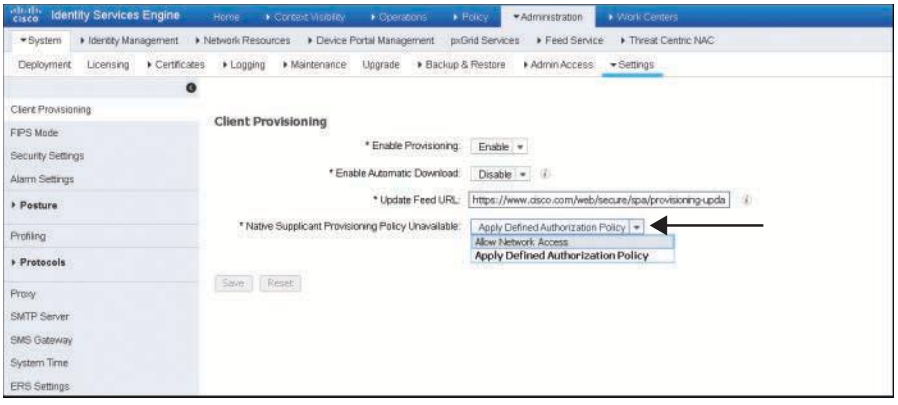


Figure 16-53 *Default Unavailable Client Provisioning Policy Action*

Create the Authorization Profiles

Follow these steps to create an authorization profile that will be used for client provisioning with the ACL that permits Android devices to reach the Google Play Store:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles.**

Step 2. Add a new authorization profile named **NSP-Redirect.**

Step 3. Check the **Web Redirection (CWA, DRW, MDM, NSP, CPP)** checkbox.

Step 4. From the dropdown, choose **Native Supplicant Provisioning.**

Step 5. Set ACL to **NSP-ACL.**

Step 6. For the Value, choose BYOD Portal (default) from the dropdown.

Step 7. Click **Submit.**

Figure 16-54 shows the completed Android native supplicant authorization profile.

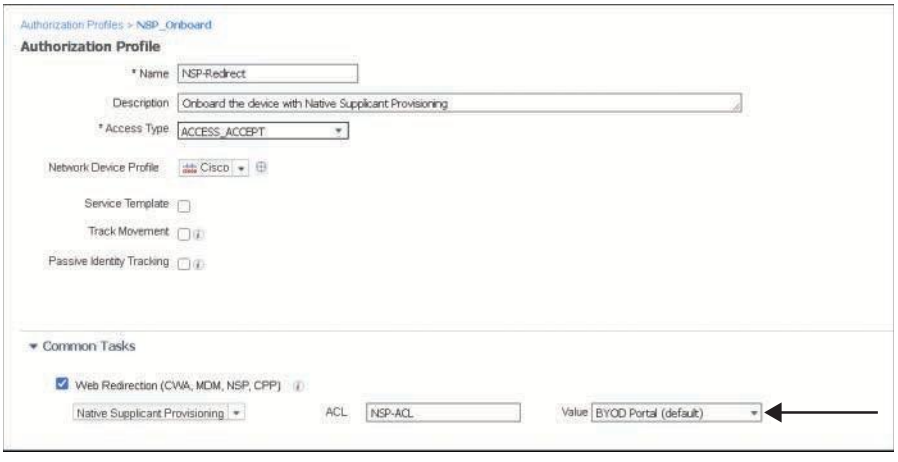


Figure 16-54 *Native Supplicant Authorization Profile*

Create the Authorization Rules for Onboarding

Now that the authorization profiles are ready, you can create two different authorization rules: one for Android and one for the other devices. Follow these steps:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the previously created **CiscoPress SSID** policy set.

Step 3. Expand the authorization policy, click on the cog icon on the topmost rule, and choose **Insert new role above**.

Step 4. Name the rule **NSP**.

Step 5. Set the conditions as follows:

```
Network Access:AuthenticationMethod EQUALS MSCHAPV2
```

Step 6. Set the authorization result to be the **NSP-Redirect** profile created earlier.

Step 7. Click **Save**. Figure 16-55 shows the native supplicant profile rule.



Figure 16-55 *Native Supplicant Profile Authorization Conditions*

Optionally, if you wanted to have a separate ACL for Android and other devices, you could create another rule above this rule that is specifically for Android devices. In this example, we have consolidated the ACL to a single wireless ACL. However, if you would like to tighten the security controls further to create a separate ACL and resulting authorization profile, you can continue with the following steps.

Step 8. Click on the cog icon for the **NSP** rule and choose **Insert new row above**.

Step 9. Name the rule **NSP Android**.

Step 10. Set the conditions as follows:

```
Network Access:AuthenticationMethod EQUALS MSCHAPV2
AND
Session:Device-OS EQUALS Android
```

Step 11. Set the authorization result to be an Android-specific authorization profile if one was created.

Step 12. Click **Done**.

Step 13. Click **Save**.

Figure 16-56 shows an optional separate Android native supplicant profile authorization rule.



Figure 16-56 *Optional Android Native Supplicant Profile Authorization Rule*

Create the Authorization Rules for the EAP-TLS Authentications

At this point you have created an authorization rule that will handle the onboarding of the endpoints. After these endpoints are onboarded, they will be authenticating with EAP-TLS and certificates.

In the CiscoPress SSID policy set, you can now create a new rule that looks for certificate-based authentications and that also ensures that the MAC address of an endpoint matches the MAC address burned into the certificate during the onboarding process. As an additional metric, the rule will also ensure that the endpoint has been through the BYOD registration process. This is a Cisco recommended configuration for BYOD with Cisco ISE.

Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand the **CiscoPress SSID** policy set.

Step 3. Expand the authorization policy.

Step 4. Add a new rule just above the **NSP** rule.

Step 5. Name the new rule **Employee BYOD**.

Step 6. Set the conditions as follows:

```
Network Access:AuthenticationMethod EQUALS x509_PKIANDEn
dpoints:BYODRegistration EQUALS YesANDAD1:ExternalGroups
EQUALS ise.local/Users/Employees
```

Step 7. Set Authorization Result to **Employee Full Access**.

Step 8. Click **Save**.

Figure 16-57 shows the Employee BYOD conditions. As shown in this figure, you may optionally save the conditions individually or together into the library for reuse in other rules.



Figure 16-57 *Employee BYOD Conditions*

ISE as a Certificate Authority

Key Topic

As of Version 1.3, ISE can function as a certificate authority (CA). The advantage of using ISE as a CA is that it eliminates the cost and overhead of public key infrastructure (PKI) maintenance; in addition, it is easy to deploy and enables an ISE administrator to manage endpoint certificates in a single portal. Right out of the box, ISE is able to provide a light-weight PKI that can directly provision certificates to BYOD endpoints and MDM solutions. ISE's CA can also act as an Online Certificate Status Protocol (OSCP) responder, which means it can determine a certificate's revocation status.

When using ISE as a CA, the Policy Administration node (PAN) is the root CA in ISE's PKI. The Policy Services nodes (PSNs) are subordinate CAs and are issued subordinate CA certificates by the PAN. The PSNs may also act as SCEP registration authorities, which perform validation of a SCEP requester and forward the certificate requests to the CA. By default, the CA root certificate in ISE's PKI is a certificate that is self-signed by the primary PAN.

The following are some of the features provided by ISE's certificate services:

- ISE issues certificates to BYOD endpoints
- ISE makes it possible to view the endpoint certificate contents instead of having to reference an external certificate authority.
- It is possible to revoke certificates issued by ISE in the same GUI without having the reference an external PKI.
- ISE acts as an intermediate CA.
- ISE performs OCSP checks for its own PKI.
- ISE can renew OCSP responder certificates for its own PKI.

To view ISE's CA certificates, navigate to Administration > System > Certificates > Certificate Authority, as shown in Figure 16-58.

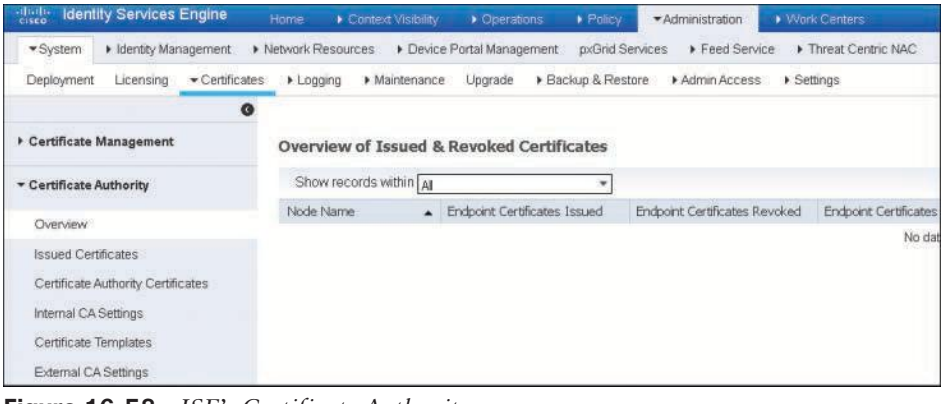


Figure 16-58 *ISE's Certificate Authority*

16

As you can see in Figure 16-58, there are several sections under Certificate Authority on the left side of the screen:

- **Overview:** This menu gives you a view of the CA nodes and the overall counts of certificates that have been issued, requested, revoked, and failed.
- **Issued Certificates:** This menu lists all the endpoint certificates that have been issued by ISE's internal CA.
- **Certificate Authority Certificates:** This menu gives you details about the CA certificates on ISE.
- **Internal CA Settings:** This menu allows you to see the CA settings and disable ISE's internal certificate authority if you wish to do so.
- **Certificate Templates:** This menu provides a quick overview of templates that are used by ISE's CA to generate certificates. There are three prebuilt templates:
  - **CA_SERVICE_Certificate_Template:** This certificate template is used when other network services use ISE as a CA.
  - **EAP_Authentication_Certificate_Template:** This certificate template is used when ISE's CA issues certificates to endpoints to use for EAP authentication.
  - **pxGrid_Certificate_Template:** This certificate template is used when an ISE administrator generates a certificate for pxGrid clients.
- **External CA Settings:** If using an external CA, this is where you can configure the SCEP information for that CA.

Configuring SCEP

Another option besides using the internal CA in ISE is to configure ISE to use an external CA that uses SCEP to issue certificates to endpoints. Multiple different external CA types may be used. The single most common CA deployed is the Microsoft Active Directory certificate authority. This CA, which has support for SCEP, is referred to as Network Device Enrollment Services (NDES). It requires Windows 2008 R2 Enterprise or newer, and that server must be part of a domain.

While not all CAs have been tested, you may theoretically use any CA of your choosing if it meets the following requirements:

- Supports SCEP
- Supports an automated or automatic issuing of the requested certificates
- Does not require a shared secret to be configured between the registration authority and the certificate authority

To configure ISE to use an external CA, follow these simple steps in the ISE GUI:

Step 1. Navigate to **Administration > System > Certificates > Certificate Authority > External CA Settings.**

Step 2. Click **Add**.

Step 3. Name the CA **Active_Directory_CA**, as shown in Figure 16-59.

Step 4. Add the URL, in the format http://*ip-address-of-ca*/certsrv/mscep/.

Step 5. Click **Test Connection**.

Step 6. Click **Submit**.

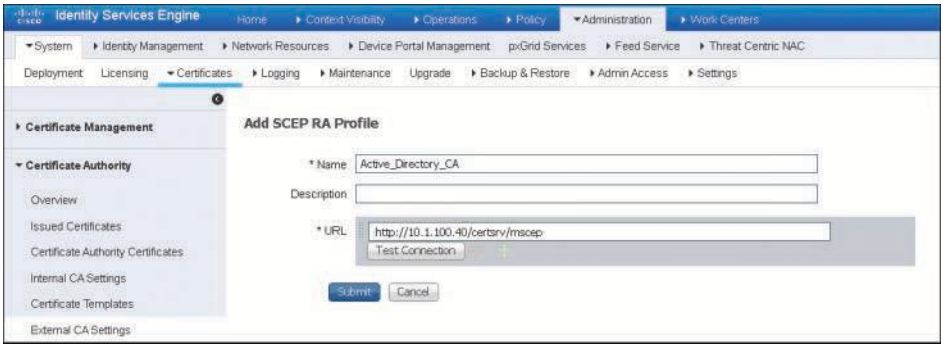


Figure 16-59 *Completed SCEP RA Profile*

When you add a new SCEP profile to ISE, the CA's public key is automatically imported into the certificate store, and the Trust for Client Authentication or Secure Syslog Services setting is enabled. This means that any endpoint for which the CA provisions a certificate will be trusted. If you choose to use CRL or OCSP, you need to edit this certificate and set that configuration. See Chapter 15, "Certificate-Based Authentication," for more details about the certificate store, trust, and revocation settings.

If the configuration on your CA is complete and accurate, your deployment is now ready to do BYOD onboarding. (See Appendix C, "Sample Switch Configurations," for details on configuring CA.) As you join the open network and authenticate or join the closed network and are directed to the onboarding process, you experience firsthand the client experience outlined earlier in this chapter.

Configuring ISE as an Intermediate CA

ISE can work seamlessly with an external CA by becoming an intermediate CA (subordinate CA). This option provides the best of both worlds: It enables an ISE administrator to have ISE issue certificates signed by the PKI and manage them from the same dashboard in ISE.

Follow these steps to configure ISE as an intermediate certificate authority:

Step 1. Navigate to **Administration > System > Certificates > Certificate Management > Certificate Signing Request**, as shown in Figure 16-60.

Step 2. Click **Generate Certificate Signing Requests (CSR)**.

Step 3. Choose **ISE Intermediate CA** from the dropdown and click **Generate**.

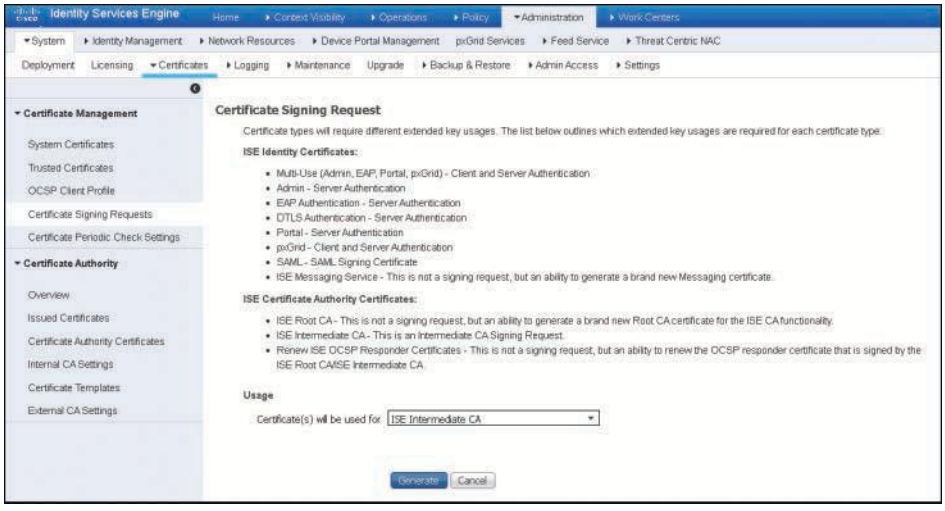


Figure 16-60 *Issuing the Certificate Signing Request (CSR)*

Step 4. When the CSR is downloaded, use it with your certificate authority to create an intermediate CA certificate, as shown in Figure 16-61. (If you are using an external Windows CA, you choose the Subordinate Certificate Authority certificate template when creating the certificate.) Give it a friendly name.

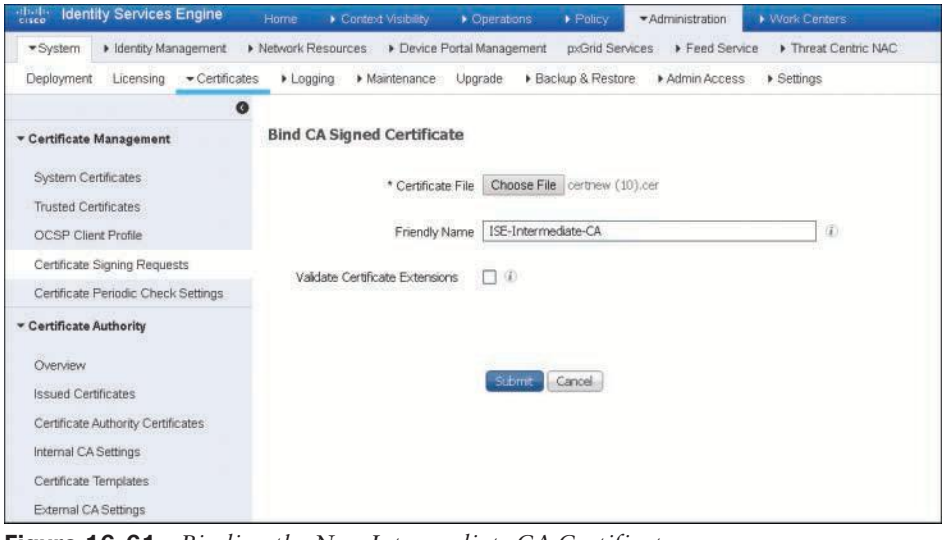


Figure 16-61 *Binding the New Intermediate CA Certificate*

Step 5. In ISE, navigate back to **Administration > System > Certificates > Certificate Management > Certificate Signing Request** and bind the new certificate to ISE.

Step 6. The newly added intermediate CA certificate should now be added to ISE's Certificate Authority Certificates page. Navigate to **Administration > System > Certificates > Certificate Authority > Certificate Authority**

Certificates to verify that the new certificate has been imported into this store, as shown in Figure 16-62. ISE's internal CA can now work seamlessly with your external CA.

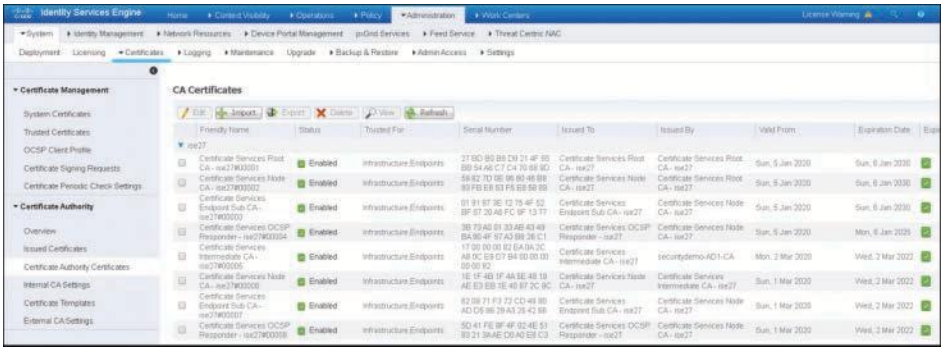


Figure 16-62 *Imported Intermediate Certificate*

BYOD Onboarding Process Detailed

Yes, this chapter is getting very long. However, it is our hope that you will find all this information very useful if you ever find yourself in a spot where you need to perform any troubleshooting of this process or need to describe it for the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam. You have seen that the user experience is quite simple and straightforward; the process behind the scenes is definitely a bit more complex.

iOS Onboarding Flow

In this section, we examine, in detail, the experience with iOS endpoints and onboarding. To do so, we focus on a single SSID onboarding experience in three phases. The end user should only have to complete four actions, as noted in Figure 16-63. However, this section describes everything that occurs behind the scenes.

Phase 1: Device Registration

The first phase of the onboarding process is device registration. Remember that the overall process may have a lot of steps, but the end user should only have to complete four actions, as illustrated in Figure 16-63. The following steps give you a look at all the items that occur behind the scenes:

Step 1. The user joins the corporate SSID, and the iOS device prompts the user for credentials (*Action 1 for the end-user experience*).

Step 2. The user enters her AD username and password (*Action 2 for the end-user experience*).

Step 3. The EAP login request is sent to the wireless controller, which wraps the request in a RADIUS Access-Accept request to ISE.

Step 4. The authorization result from ISE includes a URL redirection to the Native Supplicant Provisioning (NSP) portal.

Step 5. The user opens her web browser, which is redirected to the NSP portal on ISE and displays the device registration page with the Device ID field prepopulated with the MAC address (*Action 3 for the end-user experience*).

Step 6. The user clicks **Register** (*Action 4 for the end-user experience*), which triggers three events:

- Sets the BYODRegistration flag for the endpoint identity to Yes
- Adds the endpoint to the RegisteredDevices identity group
- Sends the CA's certificate to the iOS device for it to trust OTA provisioning

Figure 16-63 illustrates Phase 1.

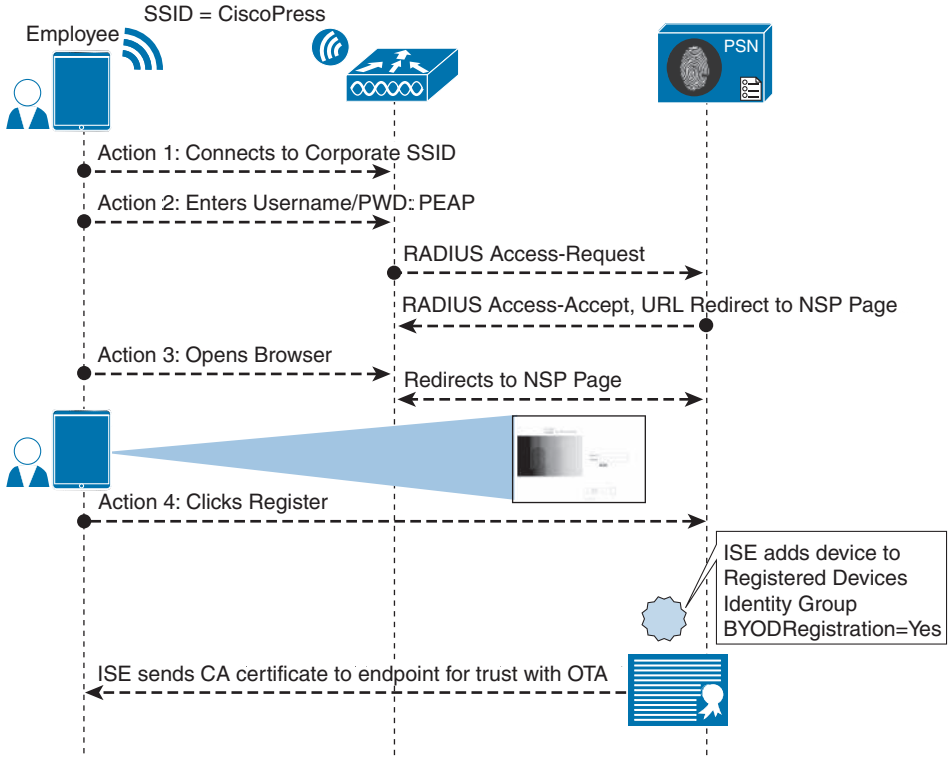


Figure 16-63 *Phase 1: Device Registration*

Phase 2: Device Enrollment

During Phase 2, a device is enrolled in the BYOD system of ISE. The device get its own certificate, but the authentication uses only the user certificate that is part of Phase 3. The following steps occur during Phase 2:

Step 7. ISE sends a profile service to the iOS device via OTA provisioning.

Step 8. The profile instructs iOS to generate a certificate signing request (CSR) using the unique device identifier (UDID) as the certificate's subject and the MAC address as the Subject Alternative Name (SAN) field:

- CN=*device-UDID*
- SAN=*MAC-Address*

Step 9. The device CSR is sent to ISE, which either issues the certificate with its own internal CA or (optionally) uses SCEP to proxy the enrollment request to the CA.

Step 10. Whether using SCEP or the internal CA, the issued certificate is sent back to ISE, which sends it to the endpoint through the OTA service.

Figure 16-64 illustrates Phase 2.

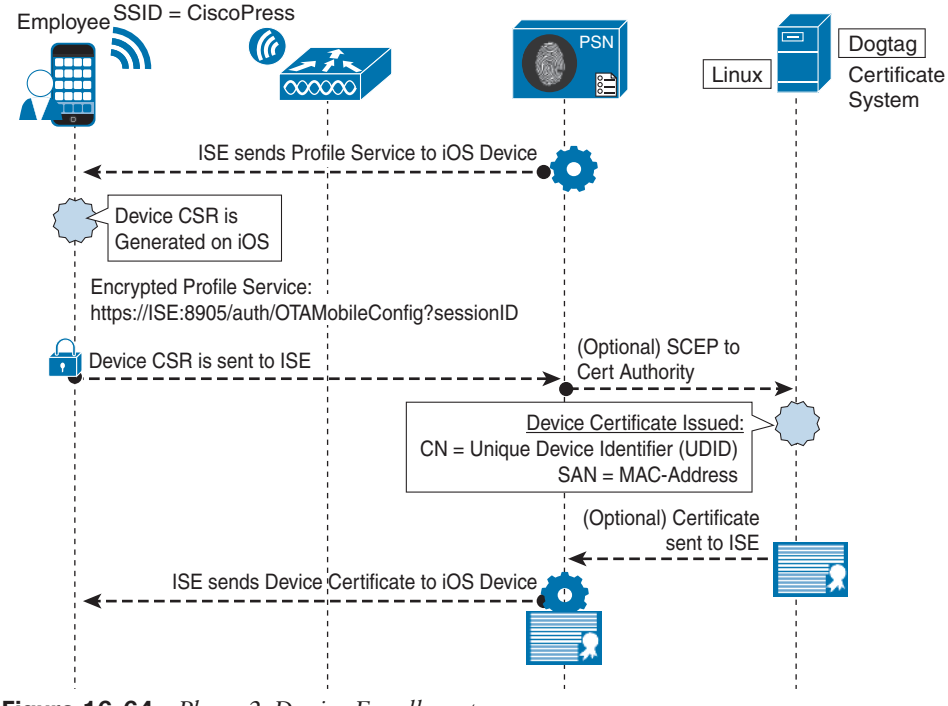


Figure 16-64 *Phase 2: Device Enrollment*

Phase 3: Device Provisioning

Now that the device has a device certificate, it's time to get a user-based certificate issued to the endpoint as well. In addition, this is the phase when the final authentication and authorization based on the user certificate and EAP-TLS occur. The following steps occur during Phase 3:

Step 11. The iOS device generates another CSR using the employee's credentials (given to iOS by ISE via the OTA service) as the certificate's subject and the MAC address as the Subject Alternative Name (SAN) field:

- CN=*Username*
- SAN=*MAC-Address*

Step 12. The user CSR is sent to ISE, which uses its internal CA to issue the certificate or uses SCEP to proxy the certificate enrollment request to an external certificate authority. If the request is proxied to an external CA, the certificate authority automatically issues the certificate and sends it back to ISE.

Step 13. ISE sends the certificate to the device through the OTA service. Included in that profile is the Wi-Fi configuration, which details the SSID and indicates to use EAP-TLS.

Step 14. ISE sends a ReAuth CoA to the NAD, which causes a new authentication.

Step 15. The endpoint authenticates to the corporate SSID using the certificate via EAP-TLS.

Figure 16-65 illustrates Phase 3.

Android Flow

To detail the flow of onboarding with Android, this section looks at the dual SSID approach. Android is certainly capable of handling a single SSID approach as well, but you just saw that experience with iOS. Remember that with the dual SSID approach, the user must log in to a CWA portal. That portal is hard-coded in ISE to launch the onboarding process only if the user is not a guest user.

With the dual SSID approach, the end user should have to complete only six actions, as shown in Figures 16-66 through 16-69. However, this section takes a look at everything that occurs behind the scenes in a series of three phases.

Phase 1: Device Registration

The first phase of the onboarding process for Android devices is device registration. Although the overall process may have a lot of steps, the end user should only have to complete six actions in total. The Android process has more steps than the iOS process because of the requirement to use an app for onboarding rather than using the native OTA process.

The following steps occur in Phase 1:

Step 1. The user joins the open SSID (*Action 1 for the end-user experience*).

Step 2. The WLC sends a MAC Authentication Bypass (MAB) request to ISE.

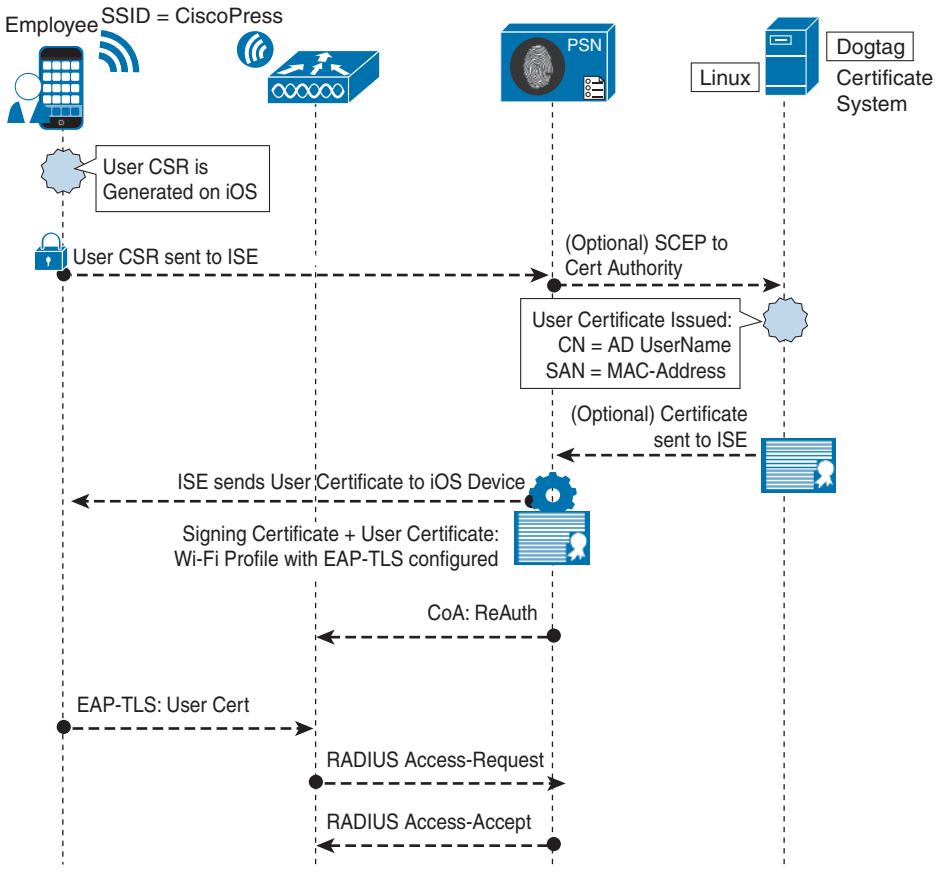


Figure 16-65 *Phase 3: Device Provisioning*

Step 3. The authorization result from ISE includes a URL redirection to the CWA portal.

Step 4. The user opens a browser and is redirected to the CWA portal (*Action 2 for the end-user experience*).

Step 5. The user enters his AD username and password (*Action 3 for the end-user experience*).

Step 6. The successful WebAuth triggers two events:

- The web page changes to the NSP portal.
- A CoA is sent to the WLC and includes a redirection to the NSP portal.

Step 7. The NSP portal displays the device registration page with the Device ID field prepopulated with the MAC address of the endpoint.

Step 8. The user clicks **Register** (*Action 4 for the end-user experience*), which triggers five events:

- Sets the BYODRegistration flag for the endpoint identity to Yes
- Adds the endpoint to the RegisteredDevices identity group

- Sets the Session:Device-OS attribute to Android (which is a temporary attribute and used only for the provisioning process)
- Sends a CoA to the WLC to apply the correct ACL, which allows Google Play Store access for the Android device
- Sends the browser to the Google Play Store

Figure 16-66 illustrates Phase 1 of the dual SSID onboarding process with Android.

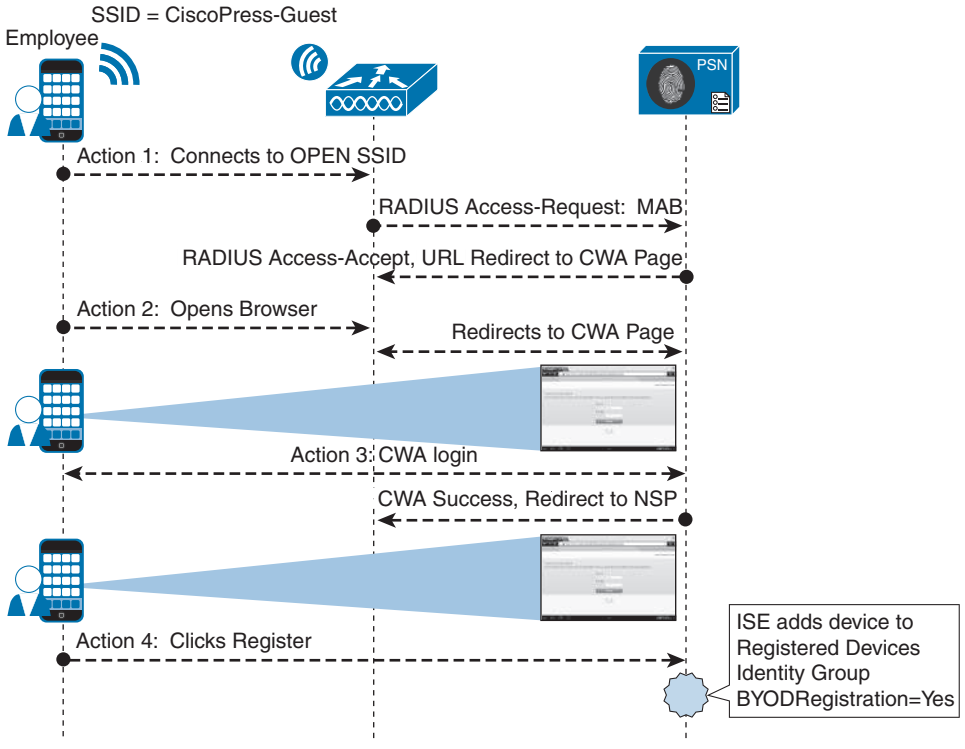


Figure 16-66 *Phase 1: Device Registration*

Phase 2: NSP App Download App

During Phase 2, the end user is sent to the Google Play Store to download the Cisco Network Setup Assistant (NSA) app. The device gets its own certificate, although the authentication process uses only the user certificate that is part of Phase 3, as illustrated in Figure 16-67. Phase 2 consists of the following steps:

Step 9. The CoA from Phase 1 applies an ACL that permits traffic to the Google Play Store (the NSP-ACL ACL).

Step 10. The browser is automatically sent to the Google Play Store, and the Android device prompts the user to choose the Internet browser or Google Play Store App to complete the request.

Step 11. The user is prompted to log in to the Google Play Store.

Step 12. The user clicks to install the Cisco Network Setup Assistant (NSA) app (*Action 5 for the end-user experience*).

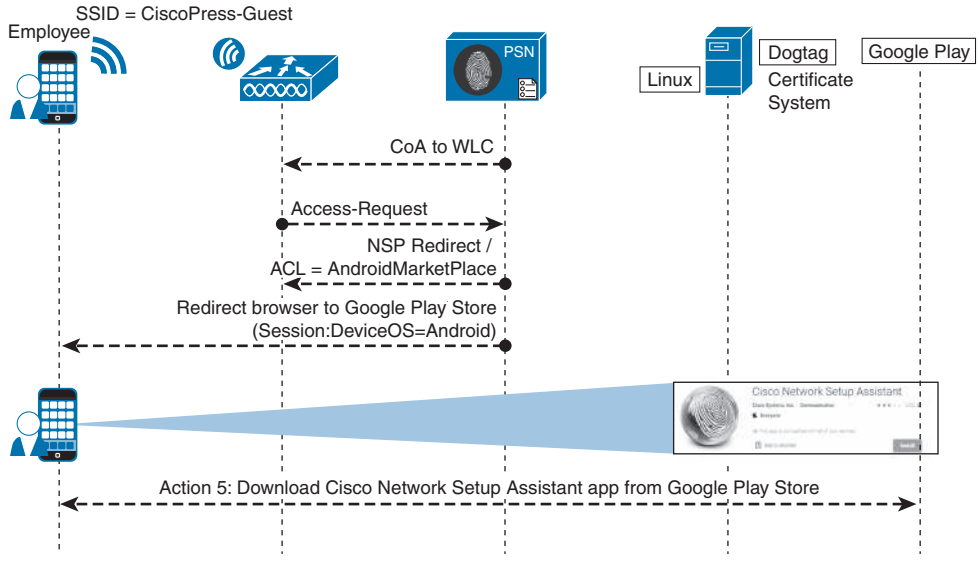


Figure 16-67 *Phase 2: Download NSA*

Phase 3: Device Provisioning

During Phase 3, the end user runs the NSA app, which handles the CSR generation and network profile installation in a similar fashion to what iOS devices do natively. The sixth and final end user action occurs in this phase, as illustrated in Figure 16-68. Phase 3 consists of the following steps:

Step 13. The user runs the Network Setup Assistant (NSA) app (*Action 6 for the end-user experience*).

Step 14. The NSA app sends a discovery message to http://*default-gateway*/auth/discovery.

Step 15. The WLC redirects that HTTP message to the ISE Native Supplicant Provisioning portal, based on the NSP-Redirect result in the authorization from ISE.

Step 16. ISE sends the Android profile based on the CiscoPress-TLS supplicant NSP to the endpoint.

Step 17. The NSA app generates the CSR by using the employee's credentials as the certificate's subject and the MAC address as the Subject Alternative Name (SAN) field:

- CN=*Username*
- SAN=*MAC-Address*

Step 18. The CSR is sent to ISE, which either uses its own internal CA or uses SCEP to proxy the certificate enrollment request to an external CA.

Step 19. The CA automatically issues the certificate.

Step 20. ISE sends the certificate and profile to the NSA app. Included in that profile is the Wi-Fi configuration, which details the SSID and indicates to use EAP-TLS.

Step 21. The NSA app connects the endpoint to the corporate SSID.

Step 22. The endpoint authenticates to the corporate SSID by using the certificate via EAP-TLS.

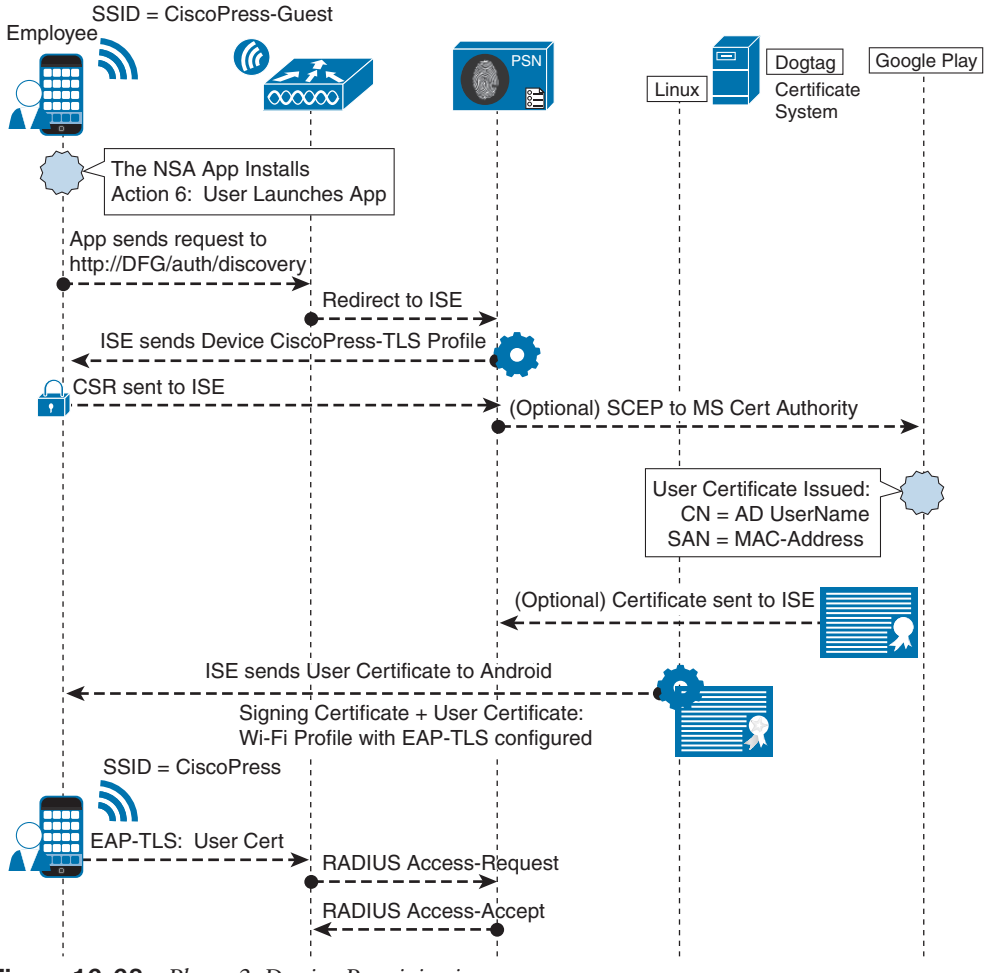


Figure 16-68 *Phase 3: Device Provisioning*

Windows and macOS Flow

macOS and Windows both use a wizard called the Cisco Native Supplicant Provisioning Wizard to accomplish onboarding and provisioning. The wizard takes care of triggering the CSR from the operating system and installing the supplicant profile.

Whereas mobile operating systems go through a three-phase process, the desktop operating systems go through a two-phase process. This section uses a single SSID onboarding example to demonstrate this flow. In a perfect world, the end user needs to take only five actions. (However, we all know that this is not a perfect world.)

The initial supplicant provisioning wizard was a Java-based applet. Due to security vulnerabilities and attacks native to Java, the security levels in the client-side Java have been tightened, making the end-user experience increasingly challenging, with different security warning prompts and workarounds. In response, Cisco has created native applications for both macOS and Windows operating systems that don't use the client-side Java.

Phase 1: Device Registration

Phase 1 for desktop operating systems is much like the Phase 1 for mobile devices. The user joins the network in the same fashion, enters credentials, and is redirected to download a piece of software that does the onboarding for the operating system. This process is similar to the flows for Android, with one major difference: There is no requirement for access to an app store as the software can be hosted directly on ISE. Phase 1 involves the following steps:

Step 1. The user joins the corporate SSID (*Action 1 for the end-user experience*), and the Windows or Mac device prompts the user for her username and password.

Step 2. The user enters her AD username and password (*Action 2 for the end-user experience*).

Step 3. The EAP login request is sent to the wireless controller, which wraps the request in a RADIUS Access-Request to ISE.

Step 4. The authorization result from ISE includes a URL redirection to the NSP portal.

Step 5. The user opens her web browser (*Action 3 for the end-user experience*), which is redirected to the NSP portal on ISE. This portal displays the device registration page with the Device ID field prepopulated with the MAC address.

Step 6. The user clicks **Register** (*Action 4 for the end-user experience*), which triggers two events:

- Sets the BYODRegistration flag for the endpoint identity to Yes
- Adds the endpoint to the RegisteredDevices identity group

Figure 16-69 illustrates Phase 1.

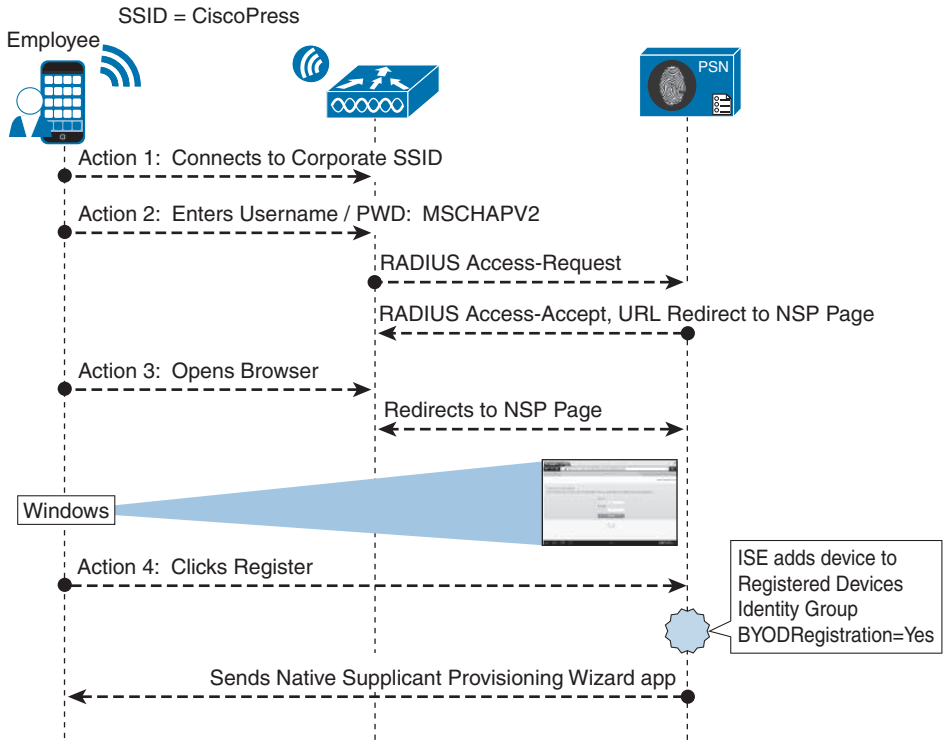


Figure 16-69 *Phase 1: Device Registration*

Phase 2: Device Provisioning

Phase 2 is the final phase of the macOS and Windows onboarding process and is much like the final phase in the Android process. The app is downloaded and executed, and it performs the CSR and provisioning for the endpoint. Phase 2 consists of the following steps:

Step 1. The Native Supplicant Provisioning Wizard is downloaded and runs.

Step 2. The user clicks **Start** (*Action 5 for the end-user experience*).

Step 3. The NSP Wizard sends a discovery message to http://*default-gateway*/auth/discovery.

Step 4. The WLC redirects that HTTP message to the ISE Native Supplicant Provisioning portal, based on the URL redirect result in the authorization result from ISE.

Step 5. ISE sends the NSP profile, based on the CiscoPress-TLS supplicant profile to the endpoint.

Step 6. The NSP Wizard generates the CSR, using the employee's credentials as the certificate's subject and the MAC address as the Subject Alternative Name (SAN) field:

- CN=*Username*
- SAN=*MAC-Address*

Step 7. The CSR is sent to ISE, which uses either its own internal CA or uses SCEP to proxy the certificate enrollment request to an external CA.

Step 8. The CA automatically issues the certificate.

Step 9. The certificate is sent down to the NSP Wizard. Included in that profile is the Wi-Fi configuration, which details the SSID and indicates to use EAP-TLS.

Step 10. ISE sends a ReAuth CoA to the NAD, which causes a new authentication. If this were a dual-SSID situation, the app would automatically connect the endpoint to the corporate SSID at this point.

Step 11. The endpoint authenticates to the corporate SSID, using the certificate via EAP-TLS.

Figure 16-70 shows the steps in Phase 2.

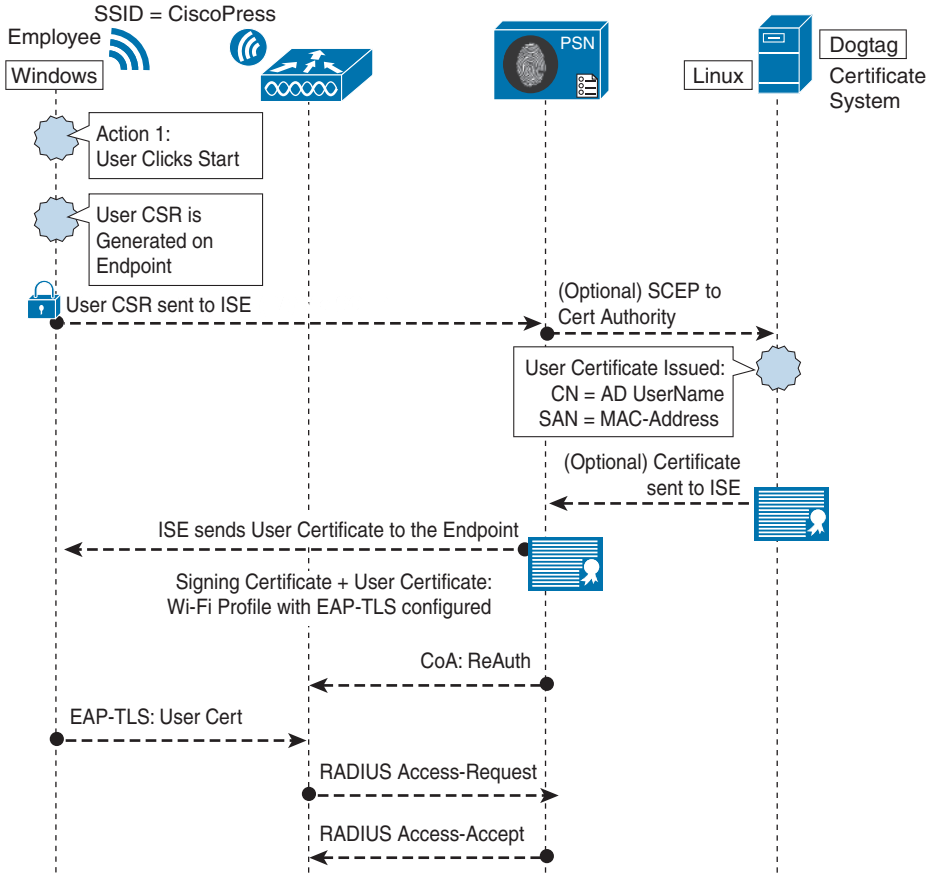


Figure 16-70 *Phase 2: Device Provisioning*

Verifying BYOD Flows

This section shows a few ways to verify that BYOD flows are operating successfully. This topic is covered in detail in Chapter 21, "Troubleshooting Tools," but this section provides enough basic information for you to get started.

An ISE administrator should look in numerous locations to see if something is functioning correctly. The first is the RADIUS Live Logs tab. In addition, reports and the endpoint identities database, as well as the My Devices portal, can provide information. The My Devices portal is covered in more detail in the "Self-Management" section, later in this chapter.

RADIUS Live Logs

As described numerous times throughout this book, the first place an ISE administrator should always look to verify functionality is the RADIUS Live Logs tab. This tab shows, at a glance, the different authentications of an endpoint and the authorization result of each of them.

Figure 16-71 shows a screen capture of the RADIUS Live Logs tab during a single SSID onboarding flow.

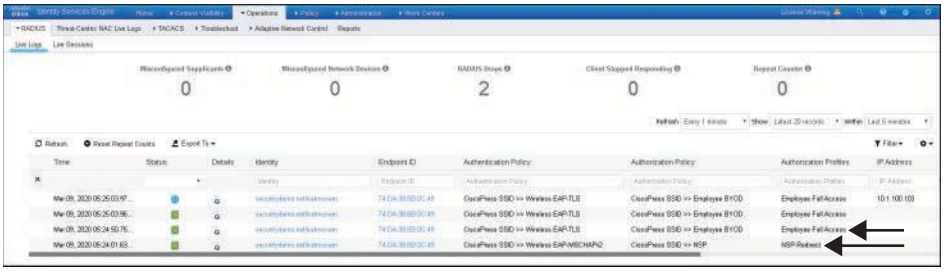


Figure 16-71 *RADIUS Live Logs Tab Showing Single SSID Onboarding*

In Figure 16-71, two logged events are highlighted:

- You see the initial authentication that uses PEAP (EAP-MS-CHAPv2) for authentication and is given the NSP authorization profile. This is the authorization profile that you created to redirect the user to the provisioning portal for device registration and provisioning.
- You see a successful authentication that uses EAP-TLS and is assigned the Employee_Limited authorization profile.

Reports

Cisco ISE provides a number of prebuilt reports that exist under Operations > Reports. They are divided into the logical groupings Audit, Device Administration, Diagnostics, Endpoints and Users, Guest, Threat Centric NAC, and TrustSec. There is a report dedicated to supplicant provisioning. To run this report, follow these steps in the ISE GUI:

Step 1. Navigate to **Operations > Reports > Reports > Endpoints and Users > Supplicant Provisioning.**

Step 2. Choose **Today** as the time range.

Step 3. Click **Run.**

Figure 16-72 shows the output of this report.

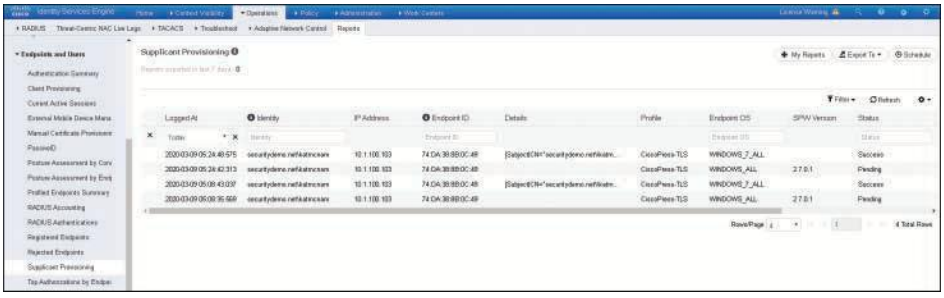


Figure 16-72 *Supplicant Provisioning Report*

Identity Group

Another way to verify the onboarding is to examine an endpoint in the endpoint identity groups database. To do so, follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > Identity Management > Groups > Endpoint Identity Groups > RegisteredDevices**.

Step 2. (Optional) Apply a quick filter to make the list more manageable.

Step 3. Examine the MAC address and static group assignment.

Figure 16-73 shows the endpoints database with a quick filter enabled.



Figure 16-73 *RegisteredDevices Endpoint Identity Group*

MDM Onboarding

Many organizations use mobile device management (MDM) solutions to provide endpoint management for a plethora of devices. MDM solutions help enforce specific security requirements, such as endpoint encryption, PIN lock, jail-break detection, and remote wipe capabilities. Many mobile device managers can even provision supplicants and certificates to devices as part of their management package.

In the past, a user bringing a mobile device into an organization, in order to gain access to the network, would have to call the help desk for instructions on how to onboard the device

with the mobile device manager. There are some significant downsides to this process, including the following:

- Users were required to manually connect to the MDM solution to begin the onboarding process.
- There was no enforcement to help steer the user toward the correct solution.
- A mobile device manager license was required for every device the organization would provision and allow to have network access, but this could be cost prohibitive.

However, Cisco recognized that MDM vendors had mobile device management capabilities that could work with Cisco's onboarding network access policy and enforcement mechanisms.

ISE integrates with many of the industry's top MDM vendors' products, including VMware AirWatch, Mobile Iron, Citrix XenMobile, Meraki Systems Manager, Globo, IBM MaaS360, and Microsoft Intune. These vendors have implemented an application programming interface (API) written by Cisco to enable scalable bidirectional communication between their solution and ISE. By using the same API for all vendors instead of having a custom integration for each vendor, Cisco can ensure consistency and quality.

Integration Points

Cisco's MDM API enables ISE to use MDM attributes in the authorization policies. The authorization may use a macro-level attribute stating that the device is in compliance with the MDM policy, or micro-level attributes such as jail-break status, PIN lock, or endpoint encryption. Table 16-2 describes the various MDM attributes that can be used as part of the authorization policy in ISE.

Table 16-2 MDM Attributes Used in Authorization Policies in ISE

MDM Attribute	Description	Possible Values
DeviceRegisterStatus	Indicates whether the device is registered with the MDM	Unregistered Registered
DeviceComplianceStatus	Macro-level attribute that indicates whether the device meets the security policy of the MDM	NonCompliant Compliant
DiskEncryptionStatus	Indicates whether encryption is enabled on the storage of the device	On Off
PinLockStatus	Indicates whether the device has an automatic lock, requiring a PIN or password to unlock the device	On Off
JailBrokenStatus	Indicates whether the device has been jail broken	Unbroken Broken
Manufacturer	The manufacturer of the device	Text field or can be compared to an attribute from AD or LDAP

MDM Attribute	Description	Possible Values
Model	The model of the device	Text field or can be compared to an attribute from AD or LDAP
IMEI	The unique ID	Text field or can be compared to an attribute from AD or LDAP
SerialNumber	The serial number	Text field or can be compared to an attribute from AD or LDAP
OSVersion	The version of the operating system	Text field or can be compared to an attribute from AD or LDAP
PhoneNumber	The phone number	Text field or can be compared to an attribute from AD or LDAP
DaysSinceLastCheckIn	The number of days since the device last checked in with the server	Text field or can be compared to an attribute from AD or LDAP
MDMFailureReason	The failure reason for MDM	Text field or can be compared to an attribute from AD or LDAP
MDMServerName	The ISE check-in server FQDN	Dropdown of the available MDM servers
MDMServerReachable	Indicates whether the MDM server is reachable	Reachable Unreachable
MEID	The device MEID	Text field or can be compared to an attribute from AD or LDAP
UDID	The device UDID	Text field or can be compared to an attribute from AD or LDAP
UserNotified	Indicates whether the user has been notified	Yes No

Configuring MDM Integration

Before you can configure ISE to communicate with the mobile device manager, ISE needs to trust the certificate of the MDM for the SSL-encrypted communications. To enable this, follow these steps in the ISE GUI (see Figure 16-74):

Step 1. Navigate to **Administration > System > Certificates > Certificate Management > Trusted Certificates**.

Step 2. Import the certificate of the mobile device manager as a trusted certificate.

NOTE If the mobile device manager will be provisioning certificates to endpoints with their own CA instead of using the BYOD onboarding with ISE and your CA, you also need to trust the certificate for client authentication.

Step 3. Navigate to **Administration > Network Resources > External MDM** to add the MDM server to ISE.

Step 4. Click **Add**.

Step 5. Enter a name for the connection to the MDM.

Step 6. Enter the hostname or IP address of the MDM server.

Step 7. Ensure that the port is 443 unless otherwise instructed by the MDM vendor.

Step 8. Ensure that an instance name is used, if needed. (This may be the case when the vendor is multitenant capable, such as when it is a cloud-based MDM solution.)

Step 9. Enter the administrator name that the API should connect with to enroll all the mobile devices.

Step 10. Enter the password for the API account.

Step 11. Add an optional description.

Step 12. (Optional) Change the polling interval and time interval for compliance. (This is the amount of time that ISE polls the MDM server for compliance check information.)

Step 13. Click **Test Connection**.

Step 14. Click **Submit**.

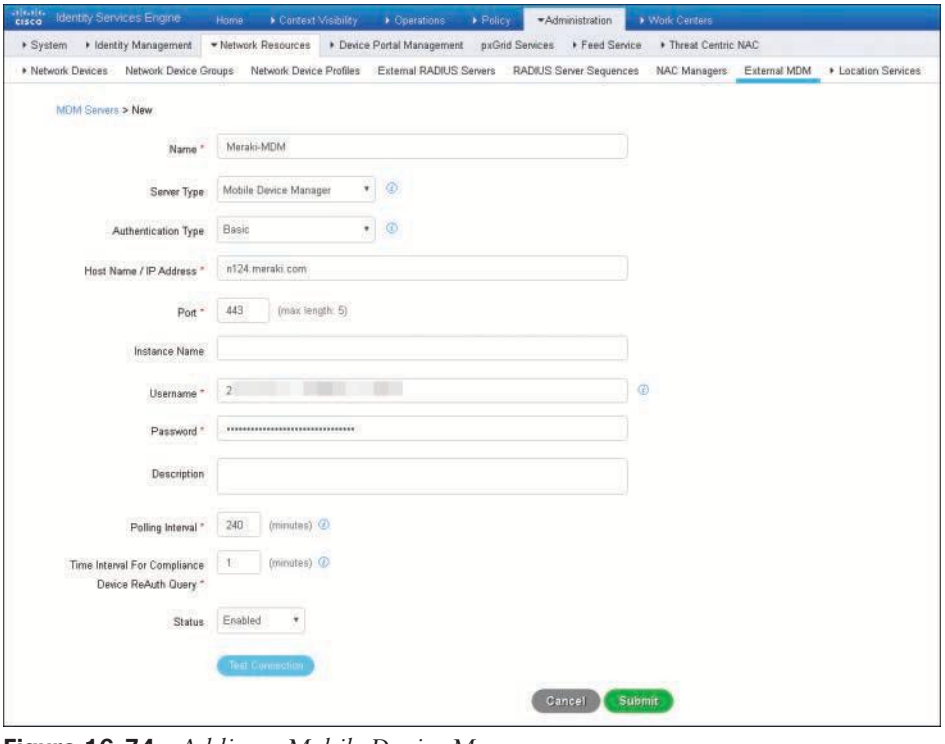


Figure 16-74 *Adding a Mobile Device Manager*

Configuring MDM Onboarding Rules

After a mobile device manager has been added to ISE and the dictionaries for MDM attributes have been added for use in ISE authorization policies, it is time to create the MDM onboarding authorization rules.

MDM onboarding rules are configured much like the ISE BYOD onboarding rules. The authorization rules need to be configured to redirect the appropriate endpoints to an onboarding portal; in this case, however, the destination is the MDM onboarding portal.

One example of where to place an MDM onboarding policy is just below the BYOD onboarding rules but above the rule that permits final access. Some organizations would not want to send all devices to the mobile device manager but would want only specific devices to be included. One way to achieve this is to maintain a separate list of MAC addresses belonging to corporate-owned assets and add that list to an endpoint identity group. This type of methodology is commonly referred to in the industry as MAC address management (MAM). Another way of steering the correct devices toward the mobile device manager is based on Active Directory group membership and device type, such as an iOS endpoint. The example in this section uses endpoint identity groups.

Create the Authorization Profile

The first step in configuring MDM onboarding rules is to create the authorization profile that redirects the endpoint to the mobile device manager's provisioning portal for onboarding. To do so, follow these steps in the ISE GUI (see Figure 16-75):

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization >Authorization Profiles.**

Step 2. Click **Add** to add a new authorization profile and name it **MDM Onboard.**

Step 3. Ensure that Access Type is set to **ACCESS_ACCEPT.**

Step 4. Check the **Web Redirection** checkbox under Common Tasks.

Step 5. Under Web Redirection, select **MDM Redirect** from the dropdown.

Step 6. Enter the ACL name **ACL-MDM-REDIRECT** to specify an Airespace ACL located on the wireless controller. The Web Redirection ACL should reference an ACL that permits access to the mobile device manager's and ISE but denies access to the rest of the Internet. (You might need to create a new one for use in this policy.)

Step 7. From the Value dropdown, choose **MDM Portal (default).** This is the default built-in MDM web portal. Alternatively, you can create a new custom portal by navigating to **Administration > Device Portal Management > Mobile Device Management.**

Step 8. From the MDM Server dropdown, choose **Meraki-MDM** as the MDM server that this authorization profile will be calling on.

Step 9. Click **Submit.**

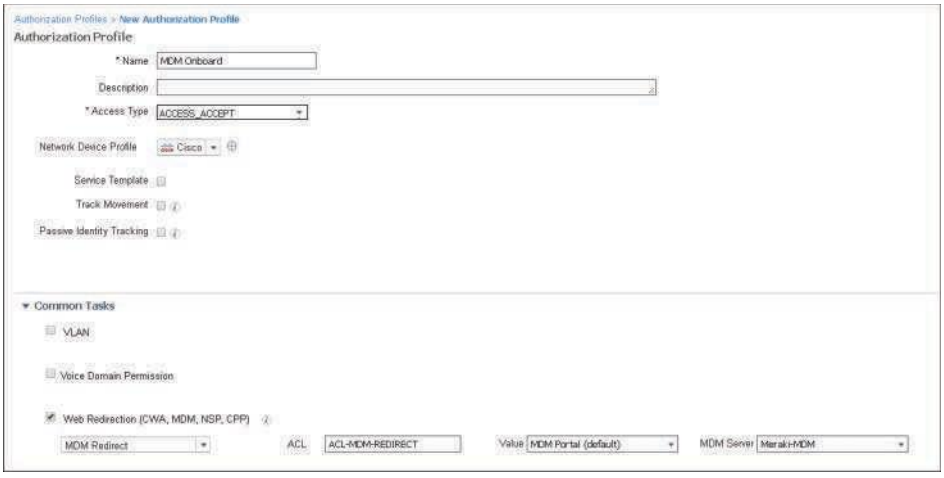


Figure 16-75 *MDM Onboarding Authorization Profile*

Figures 16-76 and 16-77 show an example of an ACL with DNS snooping.

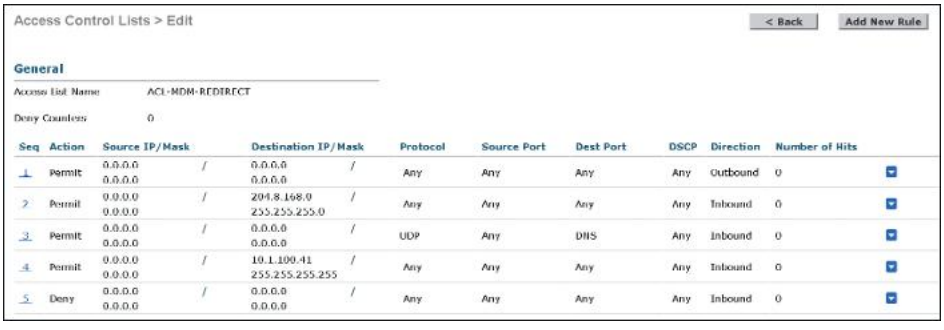


Figure 16-76 *ACL-MDM-REDIRECT*



Figure 16-77 *ACL-MDM-REDIRECT URLs for DNS Snooping*

Create the Authorization Rules

When the authorization results are ready for onboarding, you can create an authorization rule to send devices to the MDM vendor's provisioning portal for onboarding. As mentioned earlier in this chapter, the placement of these rules is important. This example has you place the redirection rule above the Employee BYOD rule. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Open the policy set for the CiscoPress SSID and expand the authorization policy.

Step 3. Insert a new rule above the Employee BYOD rule.

Step 4. Name the new rule **MDM On-boarding**.

Step 5. (Optional) Select the identity group **CorpAssets** (or a group that matches your corporate assets).

Step 6. Set the conditions as follows:

```
Network Access:AuthenticationMethod EQUALS x509_PKI
AND
CERTIFICATE:Subject Alternative Name CONTAINS Radius:
Calling-Station-ID
AND
Endpoints:BYODRegistration EQUALS Yes
AND
Ad1:ExternalGroups EQUALS securitydemo.net/Users/Domain
Users
AND
MDM:DeviceRegisterStatus EQUALS Unregistered
```

Step 7. Set Result to **MDM Onboard**.

Step 8. Click **Done**.

Figure 16-78 show the completed rule.



Figure 16-78 *MDM On-boarding Rule*

Next, you can add another rule below the MDM On-boarding rule that permits access to devices that are registered and meet the MDM compliance guidelines and therefore should be provided with full access to the network. Follow these steps to add the new rule:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Open the policy set for the **CiscoPress** SSID and expand the authorization policy.

Step 3. Click the cog icon for the MDM On-boarding rule and choose **Duplicate Below**.

Step 4. Name the new rule **MDM Permit**.

Step 5. Delete the **BYODRegistration** condition.

Step 6. Set MDM:DeviceRegisterStatus to **EQUALS Registered**.

Step 7. Add a new condition for **MDM:ComlianceStatus EQUALS Compliant**. The final condition set should look like this:

```
Network Access:AuthenticationMethod EQUALS x509_PKI
AND
CERTIFICATE:Subject Alternative Name CONTAINS
Radius:Calling-Station-ID
AND
Ad1:ExternalGroups EQUALS securitydemo.net/Users/Domain
Users
AND
MDM:DeviceRegisterStatus EQUALS Registered
AND
MDM:DeviceComplianceStatus EQUAL Compliant
```

Step 8. Set Result to **Permit Access**.

Step 9. Click **Done**.

Step 10. Click **Save**. Figure 16-79 shows the completed rule.

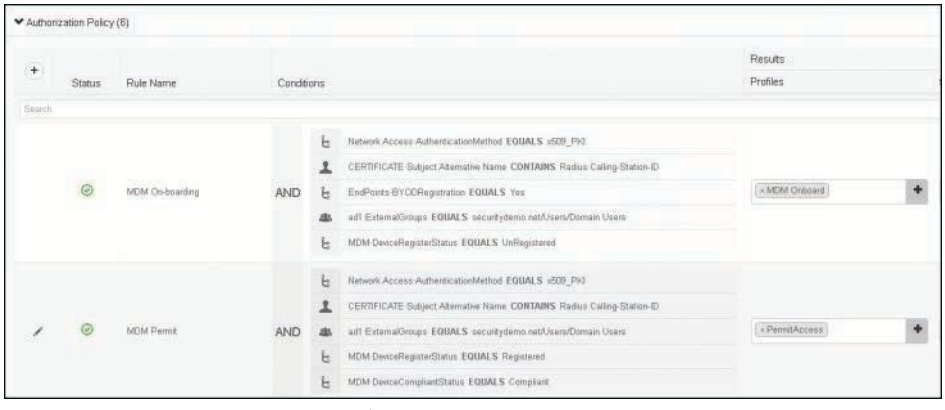


Figure 16-79 *MDM Permit Rule*

When using MDM attributes as part of the authorization policy, ISE checks with the mobile device manager at every authorization. If the mobile device manager is unavailable, the rule will never match. In addition, a bulk download of data from the mobile device manager occurs every four hours, which provides detailed endpoint status. If an endpoint is marked noncompliant during such a download, a CoA is sent, and the device is forced to reauthenticate and provide a different result (such as quarantine).

Managing Endpoints

Each user can use the My Devices portal to manage the devices he or she has personally onboarded, or the administrator can manage user devices from the Endpoints screen. Table 16-3 outlines the options for a device.

Table 16-3 Endpoint Management Options

Registered Device Option	Description
Lost	Adds the device to the Blacklist endpoint identity group and denied further access to the network until it is reinstated
Reinstate	Moves the device from the Blacklist identity group to the RegisteredDevices identity group so that it can be permitted access to the network again
Delete	Removes the endpoint from the endpoint directory
Full Wipe	Wipes the entire endpoint
Corporate Wipe	Removes only the containerized data on the endpoint considered corporate—or part of a work profile
PIN Lock	Initiates a remote lock on the device

Self-Management

End users can come to the My Devices portal to manage any endpoints they have registered and also to register new endpoints.

An employee can log in to the My Devices portal by navigating to either the default URL https://*ISE_fqdn*:8443/mydevices/ or the friendly URL configured at Administration > Device Portal Management > My Devices and clicking the My Devices Portal link.

Figure 16-80 shows the My Devices portal settings page where the friendly URL is configured.

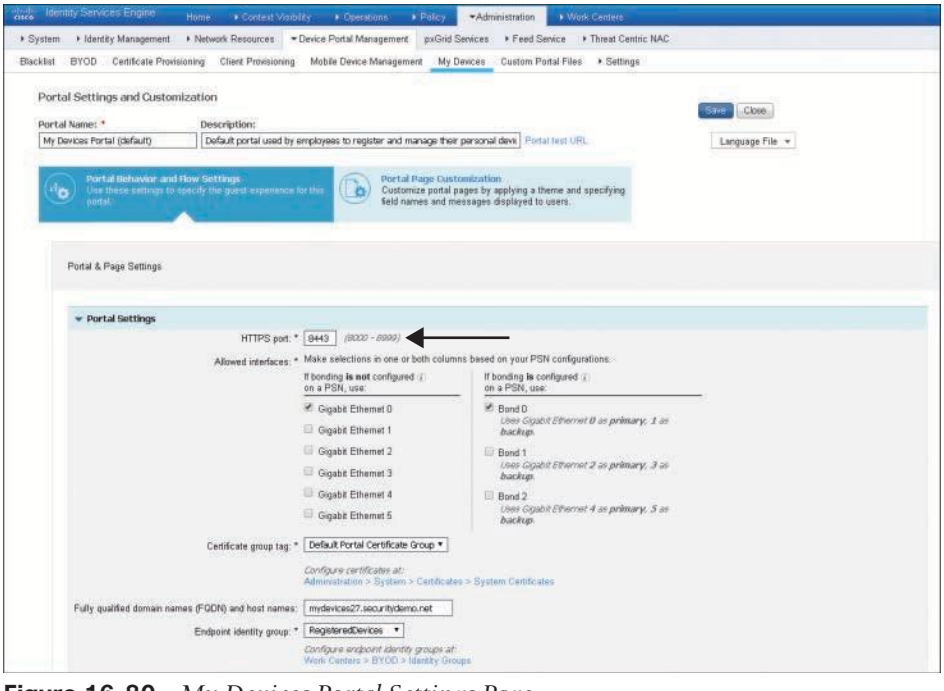


Figure 16-80 *My Devices Portal Settings Page*

By default, the My Devices portal uses an identity source sequence (ISS) called MyDevices_Portal_Sequence, which includes internal users and all Active Directory join points by default.

In the ISE GUI, you can navigate to Administration > Identity Management > Identity Source Sequences and select the MyDevices_Portal_Sequence to view the identity source sequence (see Figure 16-81).

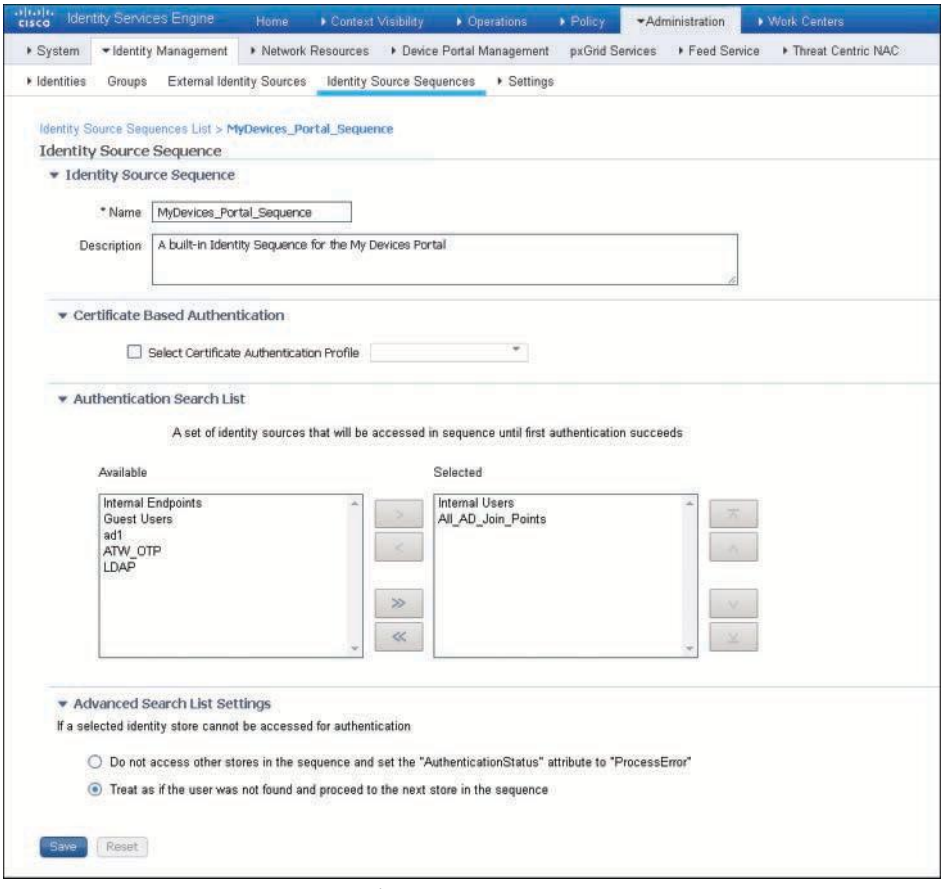


Figure 16-81 *MyDevices_Portal_Sequence*

After a user logs in to the My Devices portal, a list of registered devices is displayed. As shown in Figure 16-82, the user can select a device and initiate one of the options, such as Full Wipe.



Figure 16-82 *My Devices Portal*

Administrative Management

From an administrative perspective, BYOD registered endpoints are administered just like any other endpoints, at Work Centers > BYOD > Identities > Endpoints. From here, an administrator can initiate actions against the registered devices, as shown in Figure 16-83.

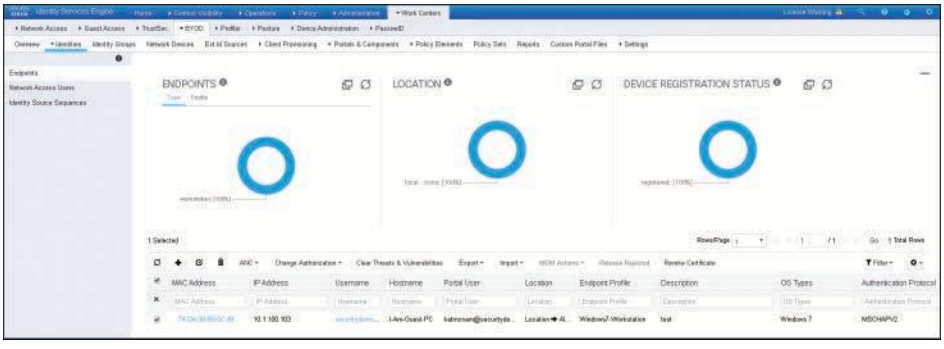


Figure 16-83 *Endpoint Identities*

The Opposite of BYOD: Identify Corporate Systems

For many years, Cisco has had customers explain their business need to identify a machine as an authorized asset in addition to the user being an authorized user. To meet this need, Microsoft Windows has both a user state and a machine state, which allows a device to be authenticated to the network with what is commonly known as Machine Authentication as well as the ability to have the interactive user authenticated to the network.

The issue is that EAP is designed to transport a single credential. Machine Authentication occurs when there is no interactive user or if the supplicant profile is configured to only issue the machine's credentials. When the user logs in to the system, the state changes to a user state, and the system issues the credentials associated to the user. Standard RADIUS and standard EAP provide no way to join those authentications together.

To answer the issue, Cisco enhanced EAP-FAST with the ability to do EAP chaining, which makes it possible to authenticate both a machine and a user within the same authentication session. EAP-FASTv2 has being standardized by the IETF and is known as EAP-TEAP (RFC 7170). As this book went to press, only Windows 10 build 2004 and later had implemented TEAP, but the majority of vendors had plans to implement it further.

With EAP-FASTv2 and EAP chaining, the machine is issued a Protected Access Credential (PAC), which is similar to a secure cookie, and so is the user. ISE may request the machine PAC during the user authentication process, and the authorization policy is capable of using the results of either or both authentications.

The authorization condition is NetworkAccess:EAPChainingResult, and the following options are available:

- No chaining
- User and machine both failed
- User and machine both succeeded
- User failed and machine succeeded
- User succeeded and machine failed

With this level of flexibility, an authorization result may permit limited access to remediate a single failure, no access if neither user nor machine succeeds, and full access if both user and machine succeed (for example).

A real-world deployment used EAP chaining to identify corporate-owned and -managed devices. The authorization rule in that deployment was as follows:

```
IF
 the device and user authentication both succeed
 and the endpoint posture is compliant
 and the user is a member of the PCI group in Active Directory
 and the location of the endpoint is on a corporate campus
THEN
 Permit Full Access
 and assign the PCI Security Group Tag (SGT)
```

This authorization rule allowed only the listed devices to communicate to the servers housing credit card data.

Exam Preparation Topics

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 16-4 lists these key topics and the page number on which each is found.

Key Topic

Table 16-4 Key Topics

Key Topic Element	Description	Page Number
Paragraph	Android onboarding	492
Section	DNS snooping on a wireless controller and ACLs	494
Paragraph	Single SSID onboarding	495
Paragraph	Dual SSID onboarding	496
Paragraph	ISE's internal certificate authority and BYOD	519

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

single SSID onboarding, dual SSID onboarding, DNS-based ACL, Simple Certificate Enrollment Protocol (SCEP)

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is the purpose of BYOD services?
2. What do mobile device management (MDM) solutions provide beyond regular BYOD onboarding?
3. What are the pros and cons of using a single SSID approach to BYOD?
4. What are the pros and cons of using a dual SSID approach to BYOD?
5. What are the options available for BYOD certificates?



CHAPTER 17

TrustSec and MACsec

This chapter covers the following topics:

Describe SGA

Describe TrustSec architecture

SGT classification: Dynamic/static

SGT transport: Inline tagging and SXP

SGT enforcement: SGACL and SGFW

MACsec

Throughout this book, you have been exposed to many different ways of controlling network access based on the context of a user and a device. With VLAN assignment, access is controlled at the Layer 3 edge or by isolating that VLAN into a segmented virtual network (VRF instances). In addition, with ACL assignment, a local ACL can be called into action by a RADIUS attribute or a downloadable ACL (dACL) can be used. Both of these types of ACLs are applied in the ingress direction at the switch port or at a virtual port on a wireless LAN controller (WLC).

These are all very good access control methods, but controlling access only at the point of network ingress is not always the most desirable or scalable solution. In this chapter, you will learn about a powerful Cisco solution called ***TrustSec***, which makes access control more scalable and powerful. Cisco has redefined the term TrustSec numerous times, but this chapter discusses it in the context expected for the CCNP Security Identity Management SISE 300-715 exam. In this context, TrustSec is a policy-defined segmentation solution that provides advanced enforcement; this solution was formerly referred to as Security Group Access (SGA).

TrustSec is a natural progression from the IEEE 802.1AE industry standard, which is commonly referred to as ***MACsec***. This standard defines Layer 2 hop-by-hop encryption for Ethernet environments very similar to Wi-Fi Protected Access (WPA).

This chapter focuses on the fundamentals of TrustSec and MACsec and also shows how to configure the many different devices available for use in a TrustSec environment. This chapter also presents basic use cases for security group ACLs and security group firewalls.

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 17-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections.”

Table 17-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
What Is TrustSec?	1
What Is a Security Group Tag?	
What Is the TrustSec Architecture?	
TrustSec-Enabled Network Access Devices	
Network Device Admission Control (NDAC)	
Defining the SGTs	2
Classification	3
Transport: SGT Exchange Protocol (SXP)	4–6
Transport: Native Tagging	7, 8
Enforcement	9
Software-Defined Access (SD-Access)	
MACsec	10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you incorrectly guess skews your self-assessment results and might provide you with a false sense of security.

1. What is a Security Group Tag (SGT)?
   - **a.** A luggage tag applied by TSA workers at airports to flag bags as they enter security checkpoints
   - **b.** An internal assignment used in ISE to represent a local copy of an Active Directory group
   - **c.** A 16-bit value that represents the context of a user and/or a device
   - **d.** An RFID tag used to identify a wireless asset to ISE
2. Where are security groups defined in the ISE administrative GUI?
   - **a.** Administration > System > Security Group Access > Security Group
   - **b.** Work Centers > TrustSec > Components > Security Groups
   - **c.** Policy > Policy Elements > Dictionaries > System > Security Group Access
   - **d.** Work Centers > TrustSec > TrustSec Policy > Security Groups
3. Which of the following are ways that an SGT can be assigned to network traffic? (Choose three.)
   - **a.** Through manual binding of the IP address to an SGT
   - **b.** Through manual configuration on the switch port
   - **c.** Through dynamic assignment by the network access device
   - **d.** Through dynamic assignment by the 802.1X authorization result
   - **e.** Through manual configuration in the NAC Agent profile
   - **f.** Through dynamic assignment by AnyConnect Network Access Manager

4. What is the minimum SXP version for version negotiation?
   a. Version 2
   b. Version 3
   c. Version 1
   d. Version 4
5. What peering protocol can be used to transmit a mapping of IP addresses to SGTs between SGT-capable devices when traffic is crossing non-SGT-capable network segments?
   a. Enhanced Interior Gateway Routing Protocol (EIGRP)
   b. Intermediate System-to-Intermediate System (IS-IS)
   c. Border Gateway Protocol (BGP)
   d. SGT Exchange Protocol (SXP)
6. Which of the following are modes of SXP peers? (Choose two.)
   a. Speaker
   b. SGT-Reflector
   c. Listener
   d. SGT-Sender
7. How is an SGT transmitted when using native tagging?
   a. The SGT is included in the Cisco MetaData (CMD) portion of the Layer 2 frame.
   b. The SGT is included in 802.1Q trunking.
   c. The SGT is included in Inter-Switch-Link (ISL) trunking.
   d. The SGT is carried in Cisco Discovery Protocol (CDP) messages.
8. When using native tagging of SGTs, how can an administrator ensure confidentiality and integrity of the tag?
   a. By enabling MD5 authentication between SGT peers
   b. By enabling IEEE 802.1AE (MACsec) between the switches
   c. By enabling IEEE 802.1AE (MACsec) between the endpoint and the access switch
   d. By configuring peer-to-peer GRE tunnels between the switches
9. Which of the following are methods of enforcement with SGTs? (Choose two.)
   a. SGACLs on switches
   b. SGACLs on routers
   c. SGFWs
   d. SGAPPs
   e. SGTs are not enforced

10. What is the difference between uplink MACsec and downlink MACsec?
    a. Uplink MACsec defines the encrypted traffic entering the switch from the endpoint, and downlink MACsec is the encrypted traffic leaving the switch, destined to the endpoint.
    b. There is no difference between uplink MACsec and downlink MACsec.
    c. The difference is solely based on the encryption algorithm used.
    d. Uplink MACsec defines the encrypted connection between network infrastructure components, and downlink MACsec defines the encrypted connection between the access-layer device and the endpoint.

Foundation Topics

Ingress Access Control Challenges

VLAN assignment and downloadable ACLs (dACL) are fantastic ways of controlling access to a network. However, when a network grows, so do the challenges of keeping up with the ingress access controls. This section examines several standard use cases and discusses the challenges each of them faces.

VLAN Assignment

VLAN assignment based on the context of a user or device is a very common way to control access to a network. Let's use a hypothetical scenario of controlling access to servers that contain credit card data, which falls under the realm of Payment Card Industry (PCI) regulatory compliance. In this situation, VLAN assignment might work as follows:

1. A user is a member of the Retail-Managers group in Active Directory.
2. The posture of the system is compliant.
3. Therefore, ISE assigns the user to the PCI-Allowed VLAN on the switch or WLC.

In order to use this VLAN assignment to control access to the servers that house that PCI data, an ACL must be applied somewhere. Let's assume that the ACL is applied at a firewall between the campus/branch networks and the data center. The following process would occur:

1. The ACL on the data center firewall must be updated to include all the source IP addresses of PCI-Allowed VLANs throughout the entire network infrastructure. Figure 17-1 illustrates the concept of VLAN-based access control on a single switch.
2. Next, the company has decided to control access to the HR server so that only members of the HR department can talk to HR servers. Another set of rules needs to be built to assign the HR VLAN, and another set of entries needs to be configured in the access list. Figure 17-2 illustrates the concept of controlling access with two VLANs on a single switch.

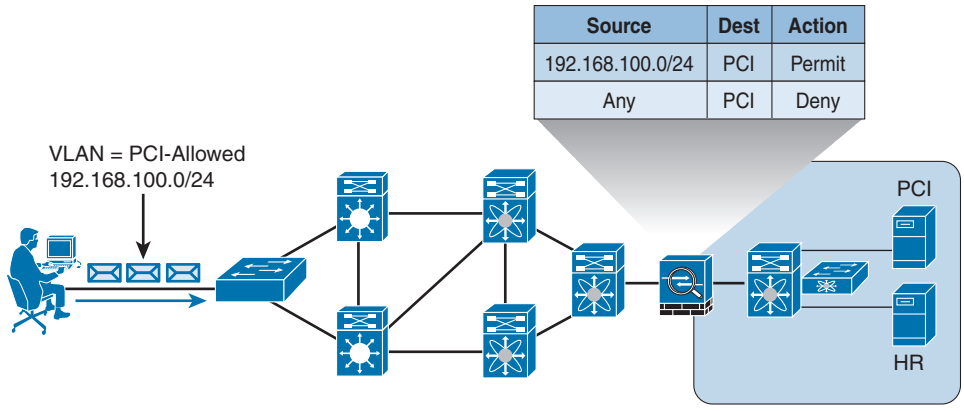


Figure 17-1 *Controlling VLAN Access on a Single Switch*

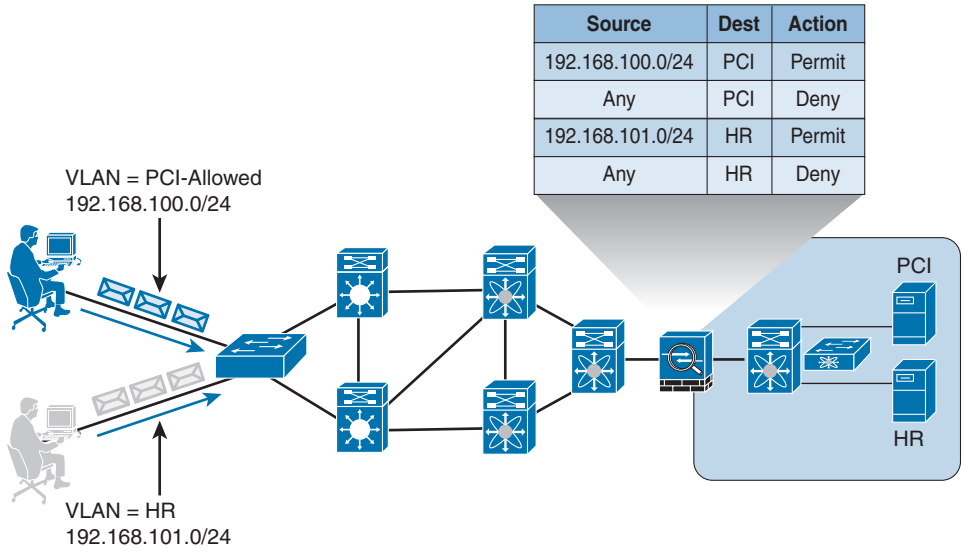


Figure 17-2 *Controlling Access with Two VLANs on a Single Switch*

How would this approach scale if the company continued to add VLANs and continued to add both switches and WLCs to the equation? What if the company had 100 remote sites? That would be 100 new IP subnets and that could easily modify the existing route summarization strategy. The route summarization alone would cause a network redesign, which would add even more operational cost. Figure 17-3 illustrates the growing complexity as the network grows.

NOTE In one real-life scenario, an organization had more than 50,000 switches in the access layer of its environment. That is a tremendous number of VLANs to create and an overwhelming number of addresses to maintain on the firewall's access control list. That same organization had 15 full-time employees managing the firewall rules and needed to find a better mechanism to control access and reduce operating expenses.

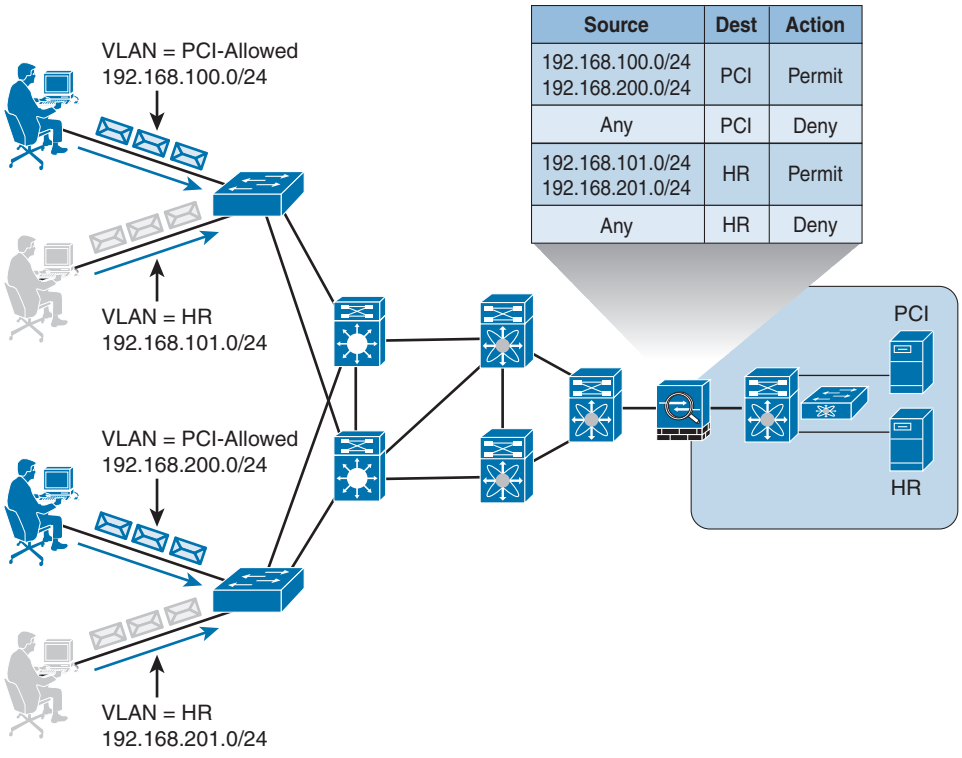


Figure 17-3 *VLAN Control Can Be Operationally Expensive*

There is a formula to determine the number of access control entries (ACE) in an access control list (ACL):

Number of sources × Number of destinations × Permissions = Number of ACEs

So, with the environment depicted in Figure 17-3, 4 sources × 2 destinations × 4 permissions would require 32 ACEs. This is obviously just a small example, and we will examine more in the following sections of this chapter.

Ingress Access Control Lists

Besides using VLAN assignment, another way to control access is to use ACLs applied at the point of ingress (inbound) at the port (or virtual port) that the user or device is using to access the network. The ACLs could be locally defined ACLs that are called by using the Filter-Id RADIUS attribute or downloadable ACLs (dACLs), where the entire ACL is defined on ISE and downloaded to the port.

Obviously, dACLs provide a better operational model than locally defined ACLs, as there is only one place to update an ACL when a change needs to be made. In addition, the number of ACEs required is lower when applying an ACL to a switch port than when applying the ACL to a centralized location since the ACL is being applied at the point of ingress, and there would theoretically be only a single source IP address. Cisco switches perform source substitution on these ACLs to make it even easier. With source substitution, the any

keyword in the Source field of an ACL is replaced with the actual IP address of the host on the switch port.

Say that a dACL solution required you to use 6 destinations and 4 permissions. This would mean 24 ACES:

1 source × 6 destinations × 4 permissions = 24 ACEs

However, there are a few complications with using ACLs on access-layer devices. Two major complications that exist are the regular maintenance of the access lists and the size of the access lists.

If access control lists are used to explicitly defend hosts, they must be updated regularly for all new destinations that get added to the network. Maintaining the lists and ensuring that they get updated correctly can require an exorbitant amount of operating expense. In addition, a switch can apply only a limited number of ACEs.

ACLs get loaded into and executed from Ternary CAM (TCAM). Access-layer switches have a limited amount of TCAM, and it is usually assigned per ASIC. Therefore, the number of ACEs that can be loaded depends on a number of factors, such as the number of hosts per ASIC and the amount of free TCAM space.

Because there is a limited amount of TCAM, ACLs cannot be overly large, especially when the access layer may have a mixture of different switches, each switch having a different level of TCAM space per ASIC. The best practice recommendation is to use fewer than 64 ACEs per dACL. This number may need to be adjusted for your specific environment, but it is a good place to start. Figure 17-4 illustrates the concept of ingress ACLs.

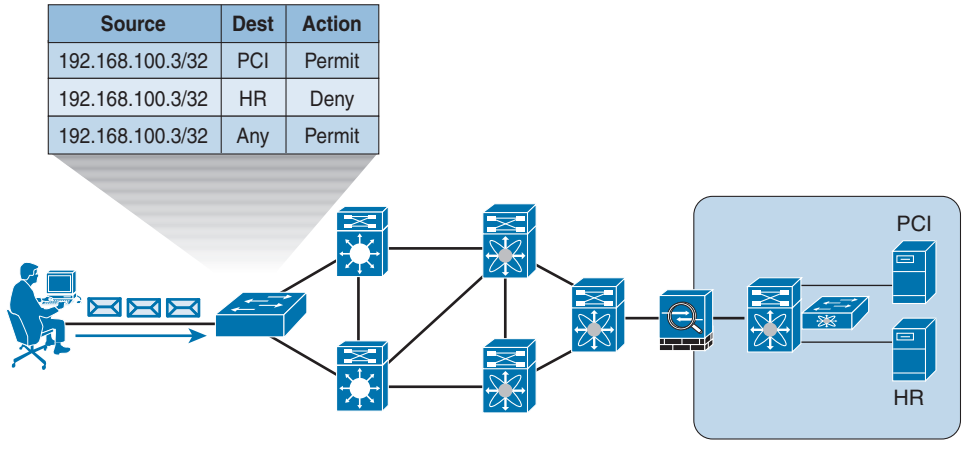


Figure 17-4 *Ingress ACLs*

East–West Segmentation

In May 2017, the WannaCry ransomware attack targeted computers around the globe. The virus affected more than 200,000 computers across 150 countries, with estimates of the damages ranging from hundreds of millions to billions of dollars. WannaCry targeted unpatched computers running Microsoft Windows. After it was executed on a compromised computer, it would exploit an SMB vulnerability to further attempt to spread to other

computers on the same network. WannaCry is just one example of many types of malware that take advantage of gaps in segmentation to further propagate.

To prevent similar attacks from successfully propagating, east–west traffic should be properly filtered based on the host or user role, including hosts that are contained in the same VLAN or subnet. VLAN segmentation is limited due to the focus on filtering north–south communication, but the hosts on the same VLAN may still talk to each other. ACLs or dACLs could potentially limit the unwanted east–west traffic on the same network segment, but that filtering would add to the complexity of those ACLs or dACLs. For example, if communication is allowed between hosts on the same subnet but only on certain ports, there will be more ACEs and complexity in that access list. Due to these limitations of ACLs, Cisco determined that there needed to be a better way to segment endpoints while simplifying the security policy and found an answer in TrustSec.

What Is TrustSec?

Key Topic

TrustSec, formerly known as Security Group Access (SGA), is a next-generation access control enforcement and segmentation technology that was created to address the growing operational expenses related to maintaining firewall rules and access lists. SGA is a complementary enforcement technology that removes the concerns of TCAM space and ACE explosion.

The ultimate goal of TrustSec is to assign a tag known as a ***Security Group Tag (SGT,*** or Scalable Group Tag) to a user's or device's traffic at the point of ingress (inbound into the network) and then enforce the access elsewhere in the infrastructure (for example, in the data center). TrustSec assigns a tag at login and enforces that tag elsewhere in the network; this is called egress enforcement.

An SGT should be representative of some overarching role(s) of the endpoint or user. For instance, an SGT may be assigned to a guest user so that guest traffic can be isolated from non-guest traffic throughout the infrastructure. The following are some very common SGTs:

- **Network Infrastructure:** This SGT gets assigned to all the switches, routers, WLCs, and firewalls in an organization.
- **Network Services:** This SGT is assigned to the servers providing common services that almost everyone should be able to reach (DNS, DHCP, NTP, and so on).
- **Executive:** An organization may use this SGT to classify their executives into their own SGT to ensure that the executives will never be denied access to anything.
- **Sales:** This SGT is for the sales force.
- **Finance:** This SGT is for the bean counters.
- **HR:** This SGT is for the group of employees who have access to your personal and financial data.
- **IoT:** This SGT is for headless devices that might be used for automation, HVAC, and so on that should be completely segmented from the company's user base.
- **Line-of-Business-1 and Line-of-Business-2:** These SGTs are used quite often when an umbrella company has many different lines of business and those lines of business cannot have access to each other's data.

The trick with SGTs is to use them for bulk access control and to do fine-grained access control within application security. In addition, each end user or end device should only be assigned a single SGT. Ideally, you do not want to create too many roles or tags and should focus on simplifying when possible.

What Is a Security Group Tag?

Key Topic

A Security Group Tag (SGT) is a 16-bit value that ISE assigns to a user's or an endpoint's session upon login. The network infrastructure views the SGT as another attribute to assign to the session and inserts the Layer 2 tag into all traffic from that session. The SGT can represent the context of the user and device.

For example, a retail organization that accepted credit cards from customers was subject to Payment Card Industry (PCI) compliance. Access to any server housing credit card data had to be protected as strictly as any technology would allow. This organization defined a rule in ISE that looked for machine and user authentication (EAP chaining) and verified that the user was a member of a PCI group in Active Directory and the machine's posture was compliant. If the user and machine met all these conditions, then an SGT named PCI was assigned. No access was granted to PCI servers without the PCI SGT.

So, as you can see, an SGT can be applied based on the full context of the authentication or simply based on a single condition, such as Guest.

NOTE The endpoint itself is not aware of the tag. It is known in the network infrastructure.

Figures 17-5 and 17-6 show the assignment of an SGT to an authentication session on a Cisco switch and on a Cisco WLC.

```
Sw01#show authentication sessions interface gl/0/47 details
            Interface:  GigabitEthernet1/0/47
                IIF-ID:  0x1073C00000000AE
          MAC Address:  a036.9f5e.a096
         IPv6 Address:  Unknown
         IPv4 Address:  10.1.100.108
            User-Name:  CORP-WIN-10-2$@securitydemo.net
               Status:  Authorized
               Domain:  DATA
      Oper host mode:  multi-auth
    Oper control dir:  both
     Session timeout:  N/A
     Restart timeout:  N/A
   Common Session ID:  0A01644B00000FDA0553B934
     Acct Session ID:  0x00000FAA
              Handle:  0x4E00002E
      Current Policy:  POLICY_Gi1/0/47

Local Policies:
       Service Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
     Security Policy:  Should Secure
     Security Status:  Link Unsecure

Server Policies:
             ACS ACL:  xACSACLx-IP-Computer_ACL-5e9240f9
           SGT Value:  4

Method status list:
      Method           State
      dot1x            Authc Success
```

Figure 17-5 *SGT Applied to a Session on a Switch*

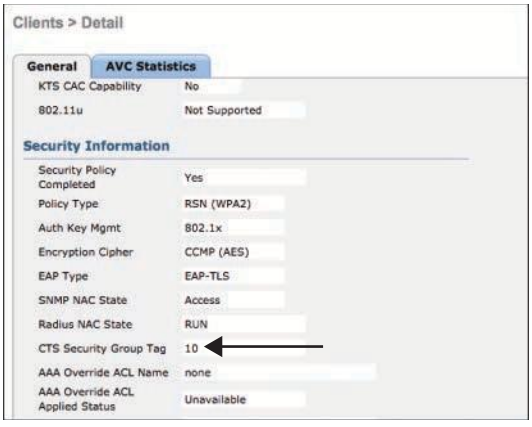


Figure 17-6 *SGT Applied to a Session on a WLC*

What Is the TrustSec Architecture?

Key Topic

TrustSec grants the ability to create secure networks by establishing a chain of trust and providing a consistent and optimized policy across a defined ***TrustSec domain***. With TrustSec, SGT and security group ACL (SGACL) policies are downloaded from a trusted source (ISE). Having these policies centrally managed by a single source of truth reduces the administrative overhead by removing the need to manually configure each switch every time an SGACL is updated or a new SGT is created.

Not only does TrustSec push a consistent policy across all the NADs in a TrustSec domain, it also optimizes how that policy is used. In a non-TrustSec environment, the IP ACL or dACL is instantiated on the NAD as a per-user ACL for that specific RADIUS session and stored in TCAM space. Using device tracking, the NAD can fill in the source IP address of the per-user ACL with the authenticated host's IP address. This per-user ACL will exist for the duration of the RADIUS session. If the port is configured for multidomain authentication, multiple hosts may be connected on the same port, and each host will have its own RADIUS session with its own per-user ACL. Even if all the endpoints have the same authorization result with the same ACL, they will each be instantiated as separate per-user ACLs and consume more TCAM space. In this scenario, dACLs and ACLs may only scale as far as the TCAM space will allow. This limits the ability to provide granular segmentation at scale. TrustSec instead provides an elegant solution to scale past these limitations.

With a TrustSec environment, the security policy is downloaded to a switch and instantiated only once in memory. There is no need to instantiate a per-user ACL for each RADIUS session, as enforcement is not bound to a specific source or to the destination IP address. Instead, the NAD only refers to its downloaded SGACL policy, which defines what source and destination tags may communicate on which ports.

SGTs are tagged or added to the frame on ingress into the TrustSec cloud, and filtering via SGACLs is performed on egress. In order to enforce a TrustSec security policy on a switch, the switch must be aware of four things:

- **The source SGT:** This may be learned inline, statically configured on the NAD, dynamically assigned during the authorization of an endpoint, or through SGT Exchange Protocol (SXP), which we cover later in this chapter.

- **The destination SGT:** This may be learned through SXP or statically configured on the switch, or the endpoint might be directly connected to the switch and dynamically assigned an SGT as an authorization result.
- **The SGACL policies:** These may be either configured manually on a switch or downloaded directly from ISE. SGACLs can be used to define how access is to be granted or how it should be controlled.
- **The environment data:** This is downloaded from ISE. The environment data includes the RADIUS server list, the NAD security group, the configured SGTs, and how often (based on the expiration time) the NAD should refresh its data from ISE.

Initially, the switch downloads the SGACL policies and ***environment data*** from ISE. When the switch receives a new frame, it verifies the source SGT. Next, the switch checks whether it knows the SGT for the destination. If the switch does not know the destination tag, it forwards the frame to the next Layer 2 hop, as it would any other frame, and the enforcement takes place at the destination switch, which has the destination tag of its directly connected endpoints. However, if the NAD has the destination SGT, it refers to its SGACL policy to validate that the source and destination SGTs can communicate. If allowed, the frame is forwarded to the endpoint. Otherwise, it is dropped.

TrustSec-Enabled Network Access Devices

Network access devices (NADs) need to perform secure RADIUS for authentication, authorization, session association, traffic filtering, and encryption in order to participate in TrustSec. The secure RAIUS configuration on the NAD is slightly different from the RADIUS configuration on the NAD. Secure RADIUS needs to have a Protected Access Credential (PAC) provisioned. On a switch, you configure a PAC key as part of the RADIUS server configuration. The switch then uses EAP-FAST to establish a TLS tunnel between it and ISE to automatically provision the PAC.

The PAC is used as a unique shared credential between ISE and the NAD. It is associated to a specific username and server authority identifier (A-ID). The PAC provisioning process is as follows:

Step 1. Server A-ID (ISE) maintains a local key (master key) that is known only by the server.

Step 2. When a client/NAD (also known as the initiator identity [I-ID] for this process) requests a PAC from ISE, ISE generates a randomly unique PAC key and PAC-Opaque field for this client/NAD.

Step 3. The PAC-Opaque field contains a randomly generated PAC key as well as other information, such as an I-ID and a key lifetime.

Step 4. The PAC key, I-ID, and Lifetime in the PAC-Opaque field are encrypted with the master key.

Step 5. A PAC-Info field that contains the A-ID is created.

Step 6. The PAC is distributed automatically or is manually imported to the client. For Cisco IOS XE switches, the PAC is provisioned inline. For an ASA, it must be downloaded from ISE and imported directly into the ASA.

After the PAC is provisioned, the PAC is used to authenticate the NAD via a secure tunnel and to download control data such as the SGACL policy and TrustSec environment data.

Defining the TrustSec Settings for a Network Access Device

Before a NAD can be configured, you need to define it in ISE. To do so, follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > Network Resources > Network Devices**.

Step 2. Click **Add** to define a new NAD.

Step 3. Give the NAD a name, define the IP address from which ISE should expect to see RADIUS requests, and (optionally) add the NAD to any network device groups (NDGs).

Step 4. Check the **RADIUS Authentication Settings** checkbox and enter a RADIUS shared secret.

Step 5. Check the **Advanced TrustSec Settings** checkbox to define TrustSec settings.

Step 6. Check the **Use Device ID for TrustSec Identification** checkbox to use the name that you defined for the device identifier and enter a password that will be used to authenticate this NAD for PAC provisioning.

Step 7. In the TrustSec Notifications and Updates section, ensure that the **Other TrustSec Devices to Trust This Device** checkbox is checked so that peer devices will not change the SGTs on packets arriving from this NAD.

Step 8. Check the box next to **Send Configuration Changes to Device** and choose **CoA** to indicate how ISE should notify the NAD that there are updates to the TrustSec policy and trigger the NAD to refresh its policy.

Step 9. In the Device Configuration Deployment section, check the **Include This Device When Deploying Security Group Tag Mapping Updates** checkbox to enable ISE to send defined IP address-to-SGT mappings using the NAD's credentials.

Step 10. Under Device Interface Credentials, enter a username, a password, and an enable password for an account that will be able to use SSH to remotely access the NAD. Whether you are using a local account on the NAD or a TACACS+ account, use one that has permissions to use SSH to remotely make changes to the configuration. Figure 17-7 shows the completed configuration.

As part of the policy acquisition phase, TrustSec-capable devices receive a Device SGT from ISE. This SGT is for the NAD and is exchanged with neighboring trusted devices. In order to have an SGT assigned to the device, you need to define it in the Network Device Authorization rules section.

Step 11. Navigate to **Work Centers > TrustSec > TrustSec Policy > Network Device Authorization** and create a rule above the default rule. This rule should encompass all the NADs that will be engaging in TrustSec, so a network device group would be a good condition to use to have NADs match this rule. Figure 17-8 shows the Network Device Authorization section with the TrustSec Devices rule.

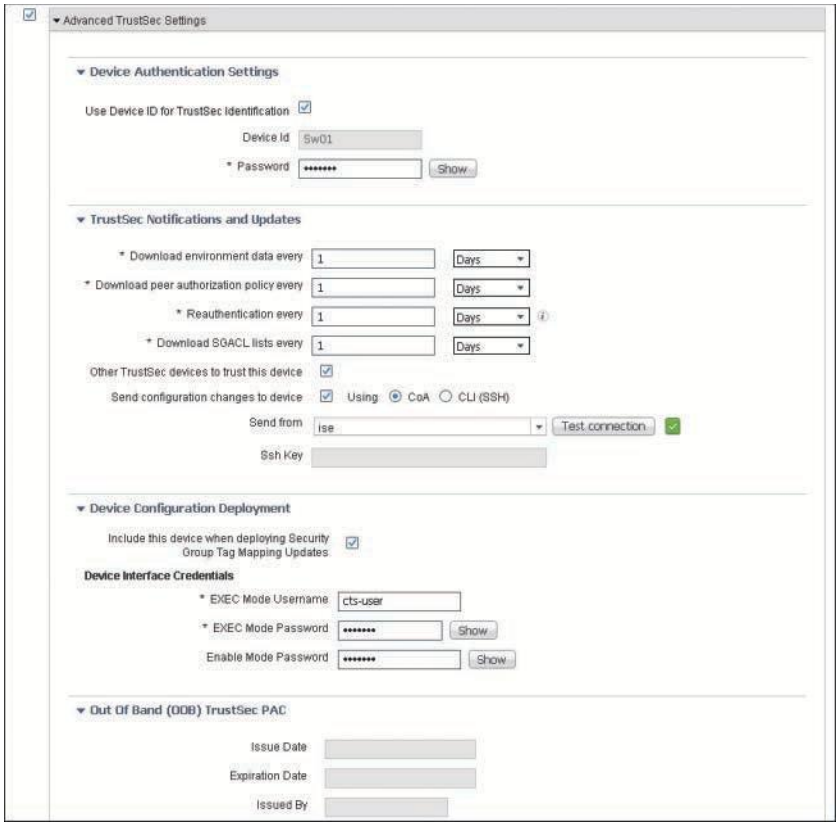


Figure 17-7 *Advanced TrustSec Settings*

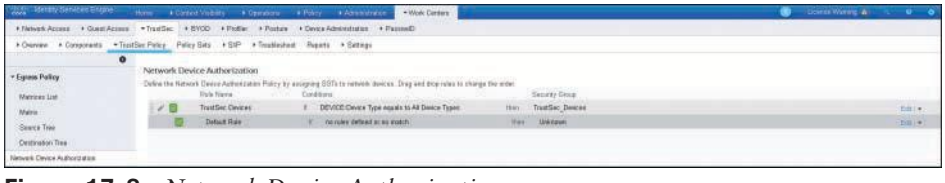


Figure 17-8 *Network Device Authorization*

Configuring an IOS XE Switch for TrustSec

Use SSH to remotely access the switch that you will be configuring and then follow these steps at the switch's CLI:

Step 1. (Optional) If you will be using a local account for ISE to send IPaddress-to-SGT mappings, create the account and set the enable password to match what you defined in the NAD configuration in ISE:

```
C9300(config)# username cts-user privilege 15 secret
Cisco123
C9300(config)# enable secret Cisco123
```

Step 2. In exec mode, define the TrustSec device ID and password:

```
C9300# cts credentials id Sw01 password Cisco123
```

This should match the device ID and password set under the NAD's Advanced TrustSec Settings in ISE. The device uses this ID to authenticate with other Cisco TrustSec devices using EAP-FAST.

Step 3. Enable AAA on the switch:

```
C9300(config)# aaa new-model
```

Step 4. Configure your RADIUS server (ISE) and define the PAC key. Ensure that the new RADIUS ports (UDP 1812/1813) are defined. If they are not defined, the old legacy RADIUS ports (UDP 1645/1646) will be used, which will cause issues with TrustSec:

```
C9300(config)# radius server ise-1
C9300(config)# address ipv4 10.1.100.21 auth-port 1812
acct-port 1813
C9300(config)# pac key Cisco123
```

Step 5. Verify that the PAC has been provisioned from IS:

```
C9300# show cts pac
AID: 95A54E9778870BF50F8184D24E67C998PAC-Info:
  PAC-type = Cisco Trustsec
  AID: 95A54E9778870BF50F8184D24E67C998
  I-ID: Sw03
  A-ID-Info: Identity Services Engine
  Credential Lifetime: 14:51:37 UTC Sat Aug 01 2020
PAC-Opaque:
000200B000030001000400109 5A54E9778870BF50F8184D24E67
C998000600940003010078B5D25F3BFF9F4C5643E0D5872CF7081
000000135EA6FD1A00093A8019B4AB6BD8C878FD05E8A504FAFF
EB6805CF573E865E805A10A05A525BF3D260E7D2BADFACBDFC9E016
53E1D09BA2105D80173E63885AE237E12346D2FC518E5F85299
BB7F9E3CCA6248E5EEB4AD5FE15C3E1F84C429358D8B357297
A3B463EC4F50B31452E1C2A0131A7A2008D27268F4865792
Refresh timer is set for 12w5d
```

Step 6. (Optional) Add the RADIUS servers to a AAA server group:

```
C9300(config)# aaa group server radius ise-group
C9300(config)# server name ise-1
C9300(config)# server name ise-2
```

This makes it possible to configure the ISE servers in a deterministic order. With the configuration shown here, the ise-1 server will be used first. If that server is dead, ise-2 will then be tried.

Step 7. Enable vendor-specific attributes (if they are not enabled by default on the switch already):

```
C9300(config)# radius-server vsa send authentication
C9300(config)# radius-server vsa send accounting
```

Step 8. Configure AAA authentication and accounting method lists:

```
C9300(config)# aaa authentication dot1x default group
ise-group
C9300(config)# aaa accounting system default start-stop
group ise-group
C9300(config)# aaa accounting dot1x default start-stop
group ise-group
```

Step 9. Define a Cisco TrustSec authorization method list:

```
C9300(config)# aaa authorization network cts-list group
ise-group
C9300(config)# cts authorization list cts-list
```

Step 10. Enable Cisco TrustSec enforcement with the following command:

```
C9300(config)# cts role-based enforcement
```

Step 11. Verify that the environment data and SGACL policies have been downloaded. Because you have not configured the SGACL policy, it should at first be blank:

```
C9300# show cts environment-data
CTS Environment Data
====================
Current state = COMPLETE
Last status = Successful
Local Device SGT:
  SGT tag = 2-00:TrustSec_Devices
Server List Info:
Installed list: CTSServerList1-0001, 1 server(s):
 *Server: 10.1.100.21, port 1812, A-ID 95A54E9778870BF50F
8184D24E67C998
           Status = ALIVE
           auto-test = FALSE, keywrap-enable = FALSE,
idle-time = 60 mins, deadtime = 20 secs
Multicast Group SGT Table:
Security Group Name Table:
    0-00:Unknown
    2-00:TrustSec_Devices
    3-00:Network_Services
```

```
        4-00:Employees
        5-00:Contractors
        6-00:Guests
        7-00:Production_Users
        8-00:Developers
        9-00:Auditors
        10-00:Point_of_Sale_Systems
        11-00:Production_Servers
        12-00:Development_Servers
        13-00:Test_Servers
        14-00:PCI_Servers
        15-00:BYOD
        16-00:Cisco_Phones
        17-00:IP_Phones
        255-00:Quarantined_Systems
    Environment Data Lifetime = 86400 secs
    Last update time = 10:08:42 UTC Sun May 3 2020
    Env-data expires in   0:23:54:40 (dd:hr:mm:sec)
    Env-data refreshes in 0:23:54:40 (dd:hr:mm:sec)
    Cache data applied           = NONE
    State Machine is running
    C9300# show cts role-based permissions
    IPv4 Role-based permissions default:
            Permit IP-00
    RBACL Monitor All for Dynamic Policies : FALSE
    RBACL Monitor All for Configured Policies : FALSE
```

Step 12. Globally enable dot1x on the switch:

```
C9300(config)# dot1x system-auth-control
```

Step 13. Configure Change of Authorization (CoA) on the switch:

```
C9300(config)# aaa server radius dynamic-author
C9300(config)# client 10.1.100.21 server-key Cisco123
C9300(config)# auth-type any
```

Step 14. (Optional) To verify that the PAC has been provisioned and the CTS data has been downloaded in ISE, navigate to **Operations > RADIUS > Live Logs**, and you should see authentication entries for the switch to initially authenticate itself as well as to authenticate subsequent TrustSec requests. Figure 17-9 shows the Cisco TrustSec Requests in the RADIUS Live Log.



Figure 17-9 *RADIUS Live Log*

Configuring an ASA for TrustSec

The following steps show how to configure the ASA for TrustSec. Because the ASA does not automatically provision the PAC, you will have to import it manually. Follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > Network Resources > Network Devices** and define the ASA just as you would a normal NAD with a RADIUS shared secret.

Step 2. Check the **Advanced TrustSec Settings** checkbox.

Step 3. Check the **Use Device ID for TrustSec Identification** checkbox and enter a password.

Step 4. Click the **Generate PAC** button, give it a password, and download the PAC.

Step 5. In the ASDM, navigate to **Configuration > Device Management > Users/AAA > AAA Server Groups**.

Step 6. Click **Add** under AAA Server Groups to create a new server group. Check the **Enable Interim Accounting Update** and **Enable Dynamic Authorization** boxes.

Step 7. Click **Add** in the Servers in the Selected Group section and add your ISE server in this newly created group. Enter the same shared secret and use the new RADIUS ports (1812/1813).

Step 8. Click **Apply**. Figure 17-10 shows the completed configuration with ISE added as a RADIUS server on the ASA.

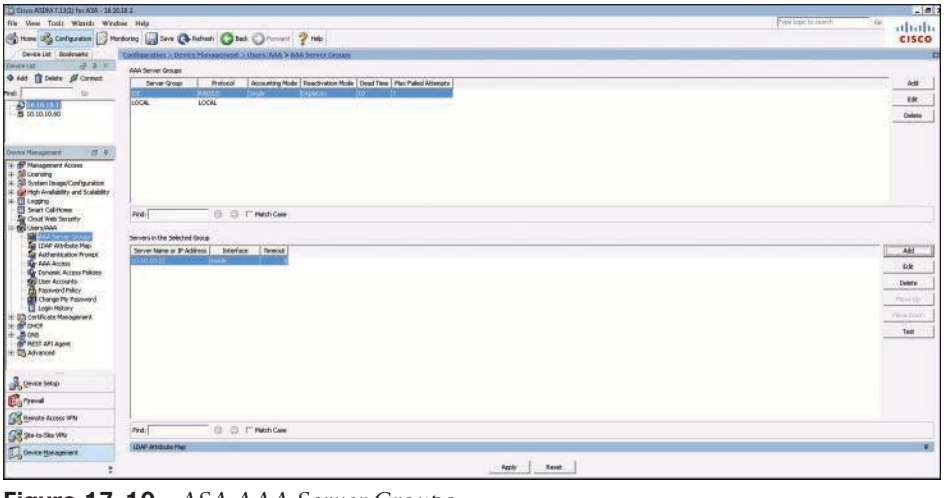


Figure 17-10 *ASA AAA Server Groups*

Step 9. Navigate to **Configuration > Firewall > Identity by TrustSec**.

Step 10. Choose the newly created server group from the dropdown, as shown in Figure 17-11.

Step 11. Click the **Import PAC** button to import the PAC.

Step 12. Click **Apply**.

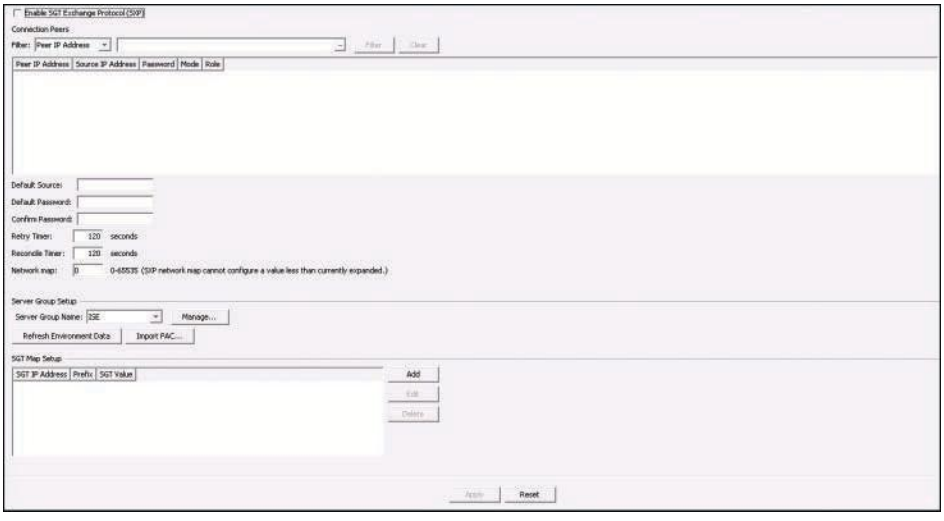


Figure 17-11 *Defining the Server Group for TrustSec*

17

Step 13. Navigate to **Monitoring > Properties > Identity By TrustSec > Environment Data** and verify that the environment data downloaded, as shown in Figure 17-12.

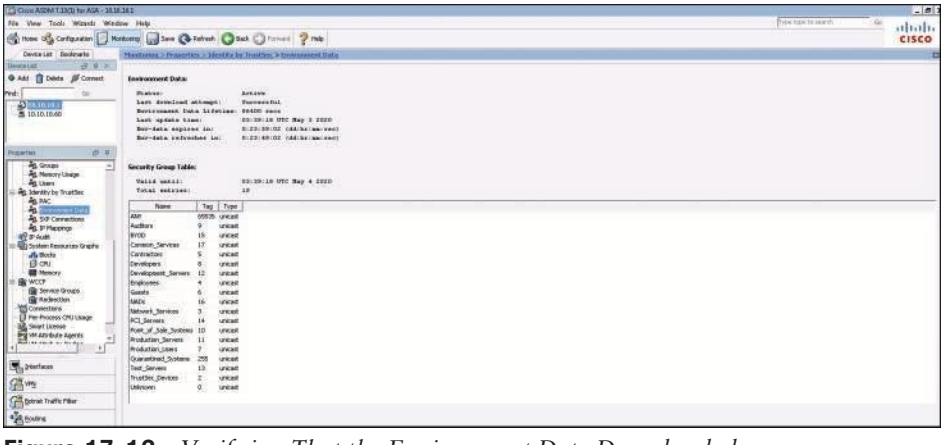


Figure 17-12 *Verifying That the Environment Data Downloaded*

Network Device Admission Control (NDAC)

TrustSec provides the ability to authenticate network access devices joining the TrustSec domain. The purpose of this authentication is to create a chain of trust within a TrustSec domain and allow only trusted network devices. This feature is called ***Network Device Admission Control (NDAC)***. Using NDAC, TrustSec authenticates the network device into the TrustSec domain using 802.1X with EAP-FAST.

The first switch to join the Cisco TrustSec domain is called the *seed switch*. Once it is authorized, this switch can provide AAA configuration information to the non-seed devices. The NDAC non-seed switches authenticate to the seed switch as the 802.1X supplicants. During the initial EAP-FAST exchange, ISE provisions a ***Protected Access Credential (PAC)***, which contains a shared key and an encrypted token that the switch will use for ongoing secure communication with ISE.

NDAC is an optional feature in TrustSec that is not supported across all Cisco switches or IOS XE versions. NDAC links are no longer supported on switching platforms supported in IOS XE Denali and Everest releases.

Configuring the Seed Device

This section walks through the configuration of a TrustSec ***seed device***. Here you start configuring from the point where you left off in the section "Configuring an IOS XE Switch for TrustSec," earlier in this chapter. The seed device will be tasked with authenticating new switches into the TrustSec domain. You need to configure the switch port facing the non-seed switch. Follow these steps at the switch's CLI:

Step 1. Move to the interface configuration mode for the interface facing the non-seed switch:

```
C9300(config)# interface g1/0/1
```

Step 2. Configure the interface for trunk mode:

```
C9300(config-if)# switchport mode trunk
```

Step 3. Enable dot1x on the interface for Cisco TrustSec:

```
C9300(config-if)# cts dot1x
```

Step 4. (Optional) Set the SAP mode for MACsec. If you do not plan on using MACsec, you can use either **no-encap** or **null**. You should mirror the encapsulation mode on the interface on the other switch:

```
C9300(config-if-cts-dot1x)# sap mode-list no-encap
```

Step 5. Enable SGT propagation on the port:

```
C9300(config-if-cts-dot1x)# propagate sgt
```

Step 6. Bring the interface down and then back up:

```
C9300(config-if-cts-dot1x)# shutdown
C9300(config-if)# no shutdown
```

After the port is brought back up, it should not be passing traffic because the connected non-seed switch should need to be configured and authenticated first before traffic would be passed.

Configuring the Non-Seed Device

This section walks through the configuration of a TrustSec ***non-seed device***. The switch must first be defined in ISE, as detailed earlier in this chapter, in the section "Defining the TrustSec Settings for the Network Access Device." No preconfiguration is needed on the non-seed switch. The seed switch should be able to get its RADIUS servers, PAC, and other information through authentication to the network. Follow these steps in the switch's CLI:

Step 1. (Optional) If you will be using a local account for ISE to send IP address-to-SGT mappings, create the account and set the enable password to match what you defined in the NAD configuration in ISE:

```
C9300(config)# username cts-user privilege 15 secret
Cisco123
C9300(config)# enable secret Cisco123
```

Step 2. In exec mode, define the TrustSec device ID and password, which should match the device ID and password set in the NAD's advanced TrustSec settings in ISE:

```
C9300# cts credentials id Sw02 password Cisco123
```

The device uses this ID to authenticate with other Cisco TrustSec devices using EAP-FAST.

Step 3. Enable AAA on the switch:

```
C9300(config)# aaa new-model
```

Step 4. With a non-seed switch, the RADIUS servers do not need to be defined as the switch adds them as part of authentication to the network, so define the AAA method lists and specify RADIUS:

```
C9300(config)# aaa authentication dot1x default group
radius
C9300(config)# aaa accounting dot1x default start-stop
group radius
C9300(config)# aaa accounting system default start-stop
group radius
C9300(config)# aaa authorization network cts-list group
radius
```

Step 5. Enable vendor-specific attributes (VSAs) on the switch (which may be enabled by default on the switch already):

```
C9300(config)# radius-server vsa send authentication
C9300(config)# radius-server vsa send accounting
```

Step 6. Globally enable dot1x on the switch:

```
C9300(config)# dot1x system-auth-control
```

Step 7. Move to the interface configuration mode for the interface facing the non-seed switch:

```
C9300(config)# interface g1/0/1
```

Step 8. Configure the interface for trunk mode:

```
C9300(config-if)# switchport mode trunk
```

Step 9. Enable dot1x on the interface for Cisco TrustSec:

```
C9300(config-if)# cts dot1x
```

Step 10. (Optional) Set the SAP mode for MACsec. If you do not plan on using MACsec, you can use either **no-encap** or **null**. You should mirror the encapsulation mode on the interface on the other switch:

```
C9300(config-if-cts-dot1x)# sap mode-list no-encap
```

Step 11. Enable SGT propagation on the port:

```
C9300(config-if-cts-dot1x)# propagate sgt
```

Step 12. Bring the interface down and then back up:

```
C9300(config-if-cts-dot1x)# shutdown
C9300(config-if)# no shutdown
```

Step 13. Verify that the PAC and CTS environment data was downloaded on the non-seed switch after authenticating to the network:

```
C9300# show cts pac
  AID: 95A54E9778870BF50F8184D24E67C998
  PAC-Info:
    PAC-type = Cisco Trustsec
    AID: 95A54E9778870BF50F8184D24E67C998
    I-ID: Sw02
    A-ID-Info: Identity Services Engine
    Credential Lifetime: 16:58:32 UTC Aug 1 2020
  PAC-Opaque:
000200B000030001000400109 5A54E9778870BF50F8184D24E67C9980
00600940003010 0CE129FE647139CEF14B80F6A73A5C38E000000135
EA6FD1A00093A8019B4AB6BD8C878FD05E8A504FAFFEB68B0CE2D538
EA268EDFE2D1CC000C9635770BC92EA108A9CFC2DF7F5641DBEFBE
34BCA34DF737794A691A59EC0A9426683F1C4A03557BE43E59E8D4C
3740F15BC438394551DBE5A42B86CE062F49FD1D798AF89025CD2
A40A32E179DDCC8C4C9B1F0366751
  Refresh timer is set for 12w5d
C9300# show cts environment-data
CTS Environment Data
====================
Current state = COMPLETE
Last status = Successful
Local Device SGT:
```

```
  SGT tag = 2-00:TrustSec_Devices
Server List Info:
Installed list: CTSServerList1-0001, 1 server(s):
 *Server: 10.1.100.21, port 1812, A-ID 95A54E9778870BF50F
8184D24E67C998
           Status = ALIVE
           auto-test = FALSE, keywrap-enable = FALSE,
idle-time = 60 mins, deadtime = 20 secs
Multicast Group SGT Table:
Security Group Name Table:
    0-00:Unknown
    2-00:TrustSec_Devices
    3-00:Network_Services
    4-00:Employees
    5-00:Contractors
    6-00:Guests
    7-00:Production_Users
    8-00:Developers
    9-00:Auditors
    10-00:Point_of_Sale_Systems
    11-00:Production_Servers
    12-00:Development_Servers
    13-00:Test_Servers
    14-00:PCI_Servers
    15-00:BYOD
    16-00:Cisco_Phones
    17-00:IP_Phones
    255-00:Quarantined_Systems
Environment Data Lifetime = 86400 secs
Last update time = 10:17:00 UTC Sun May 3 2020
Env-data expires in   0:23:59:41 (dd:hr:mm:sec)
Env-data refreshes in 0:23:59:41 (dd:hr:mm:sec)
Cache data applied           = NONE
State Machine is running
```

Step 14. You have not configured the RADIUS server, CoA server, or much else. After authentication to the TrustSec domain, the AAA configuration should not change in the switch as seen in the output of the below show command.

The RADIUS servers are not defined in the running configuration. However, they are listed in the environment data that was downloaded:

```
C9300# show running-config aaa
!
aaa authentication dot1x default group radius
aaa authorization network cts-list group radius
aaa accounting system default start-stop group radius
aaa accounting dot1x default start-stop group radius
username admin privilege 15 secret 5 $1$.
tao$9RBJkWqe3PVqC0nbfDdFn/
username cts-user privilege 15 secret 5
$1$lLcC$JER9rQMhXnsLF4cvzEZ/V.
!
!
!
!
!
!
!
aaa new-model
aaa session-id common
```

Step 15. Verify the policy peer and the interface on the non-seed switch:

```
C9300# show cts policy peer
CTS Peer Policy
===============

Peer name: Sw01
Peer SGT: 2-00:TrustSec_Devices
Trusted Peer: TRUE
Peer Policy Flag: 0x51408003
Peer Policy Lifetime = 86400 secs
Peer Last update time = 10:26:53 UTC Sun May 3 2020
Policy expires in 0:23:57:20 (dd:hr:mm:sec)
Policy refreshes in 0:23:57:20 (dd:hr:mm:sec)
Cache data applied = NONE
```

From this output, you can see that the TrustSec peer device is Sw01, and its device SGT is the TrustSec_Devices tag with number 2.

Next, verify the interface that authenticated to the TrustSec domain by entering the following command:

```
C9300# show cts interface g1/0/1
Global Dot1x feature is Enabled
Interface GigabitEthernet1/0/1:
    CTS is enabled, mode:    DOT1X
    IFC state:               OPEN
    Interface Active for 00:06:25.360
    Authentication Status:   SUCCEEDED
        Peer identity:       "Sw01"
        Peer's advertised capabilities: "sap"
        802.1X role:         Supplicant
        Reauth period applied to link:  Not applicable
to Supplicant role
    Authorization Status:    SUCCEEDED
        Peer SGT:            2:TrustSec_Devices
        Peer SGT assignment: Trusted

[...truncated]

Dot1x Info for GigabitEthernet1/0/1
-----------------------------------
PAE                         = SUPPLICANT
StartPeriod                 = 3
AuthPeriod                  = 30
HeldPeriod                  = 60
MaxStart                    = 3
Credentials profile         = CTS-ID-profile
EAP profile                 = CTS-EAP-profile
```

In this output, you can see that dot1x succeeded on this interface with Sw01, and its device SGT is the TrustSec_Devices tag with number 2. Based on the output, you can also see that Sw02 is the dot1x supplicant in this exchange.

If you ran the same **show** command on Sw01, you'd see the following output, which would indicate that it is the authenticator in the exchange:

```
Dot1x Info for GigabitEthernet1/0/1
-----------------------------------
PAE                          = AUTHENTICATOR
QuietPeriod                  = 60
```

```
ServerTimeout                = 30
SuppTimeout                  = 30
ReAuthMax                    = 2
MaxReq                       = 2
TxPeriod                     = 30
```

Defining the SGTs

ISE serves as the single source of truth for what SGTs exist, and ISE considers an SGT a policy result. Therefore, you can create one SGT result for each SGT you wish to define in the environment. Follow these steps in the ISE GUI:

Step 1. Navigate to **Work Centers > TrustSec > Components >Security Groups**. Figure 17-13 shows the Security Groups list. Notice that there is a default SGT of 0, for "unknown." This is the tag that will be used if untagged traffic arrives. In other words, even the lack of an SGT can be used in the security policy.

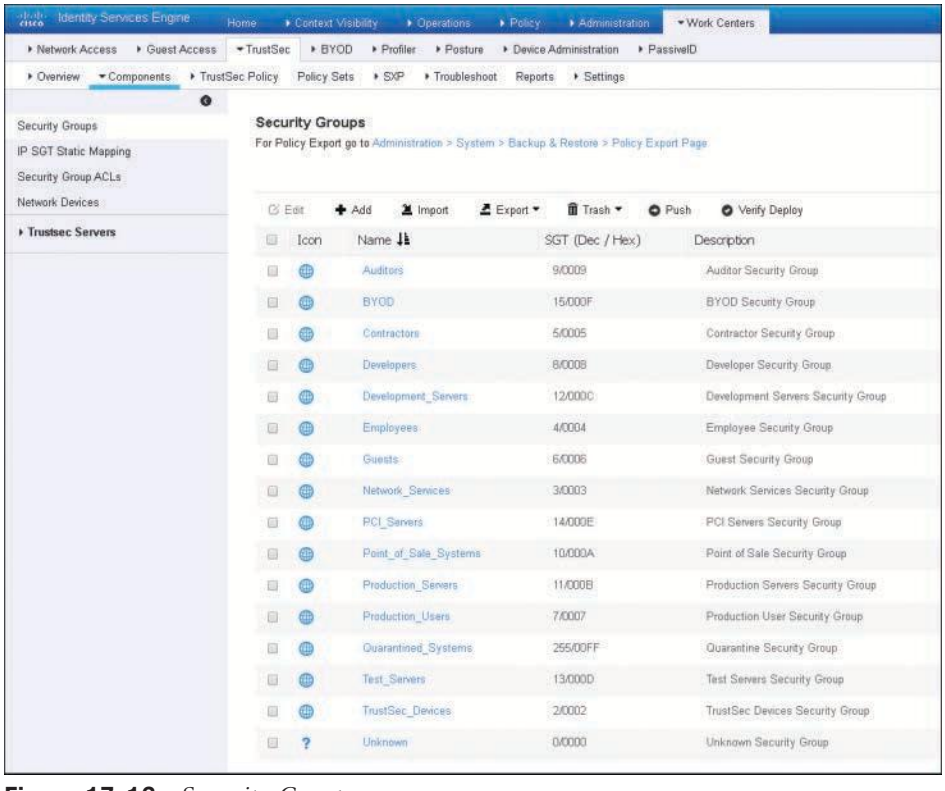


Figure 17-13 *Security Groups*

Next, you can add new security groups to be used in your network infrastructure. You begin by creating a security group for network access devices.

Step 2. Click **Add**.

Step 3. Name the new SGT **NADs**. Figure 17-14 shows the addition of a new SGT for the NADs. Notice in the figure that the Security Group Tag value is predetermined. By default, ISE automatically assigns the values in order, from 1 to 65535. However, this setting can be changed.

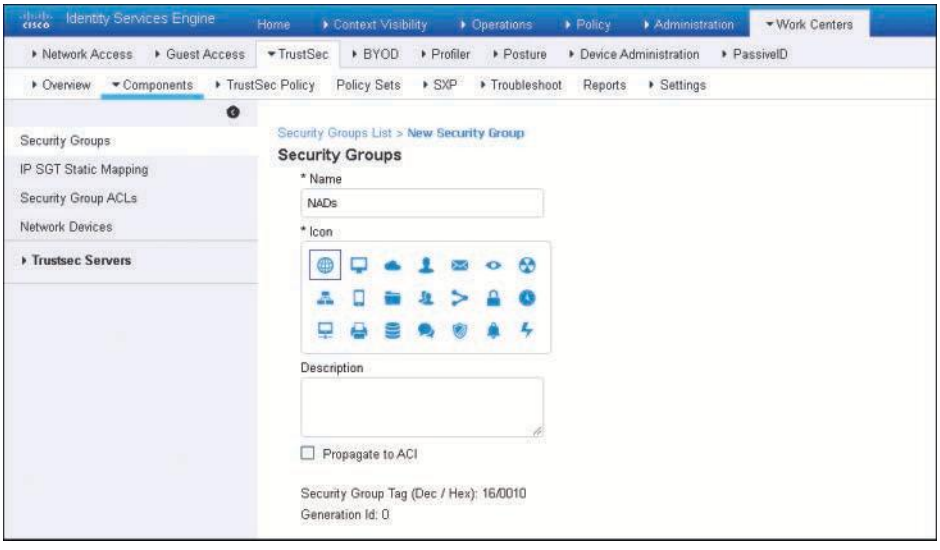


Figure 17-14 *Adding a Security Group for NADs*

17

NOTE If ISE is integrated with ACI, you can also check the box Propagate to SGT to propagate the SGT to ACI. When you do this, you automatically create an endpoint group (EPG) in ACI that maps to this SGT. However, you still need to create contracts within ACI to define how this EPG should communicate with other EPGs.

Step 4. Click **Submit** to save.

NOTE It is considered a best practice to also have a security group for all the common network services that will exist on a network. These are services that should always be accessible by any device, such as DNS and DHCP. The following steps walk through creation of such a group.

Step 5. Click **Add**.

Step 6. Name the new group **Common_Services**.

Step 7. Click **Submit** to save.

Step 8. Repeat steps 6–8 to add a new group named **Admin**, as well as any other security groups you need for your organization's security policy.

Figure 17-15 shows an example of a set of security groups.

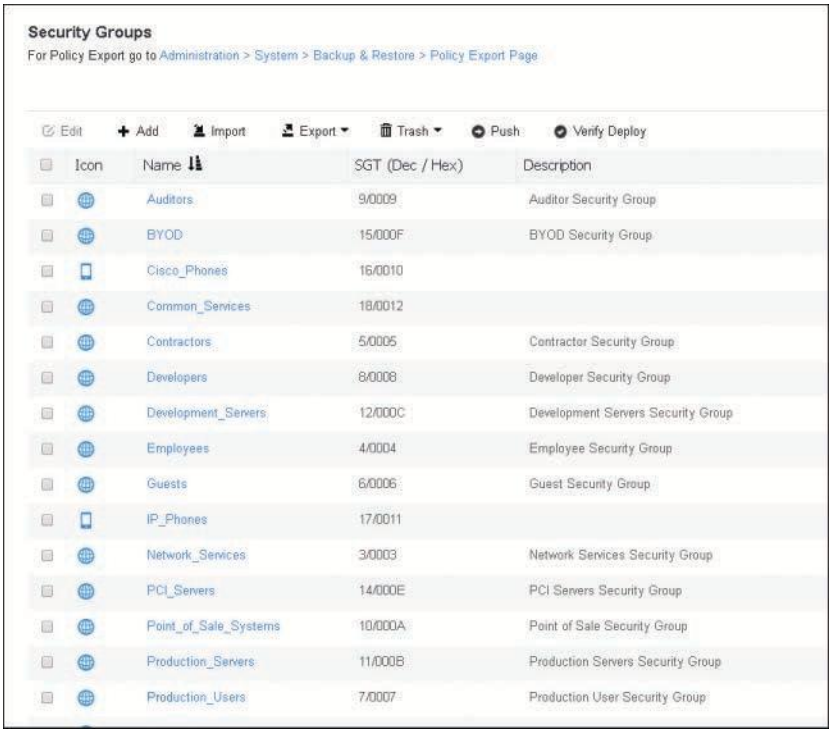


Figure 17-15 *Security Groups Example*

Alternatively, you can change the settings in ISE to automatically create SGTs as you create authorization rules in the policy sets.

To have the security groups automatically created in ISE, follow these steps in the ISE GUI:

Step 1. Navigate to **Work Centers > TrustSec > Settings > General TrustSec Settings**.

Step 2. Check the box next to **Auto Create Security Groups When Creating Authorization Rules** to enable this feature (see Figure 17-16). Optionally, you can define the SGT number range that will be used for all the system-generated SGTs.

Step 3. (Optional) Define how the generated SGTs will be named: based on the authorization rule name, SGT number, or both. You can also add a custom prefix or suffix to the SGT name. Define your SGTs based on the rule name and add the prefix **SGT**.

Step 4. Click **Save**.

Step 5. Navigate to **Policy >Policy Sets** and expand one of the policy sets. In the authorization policy, you should now have the option to automatically create a security group for each authorization rule. Figure 17-17 displays the authorization rules with the option to create security groups.

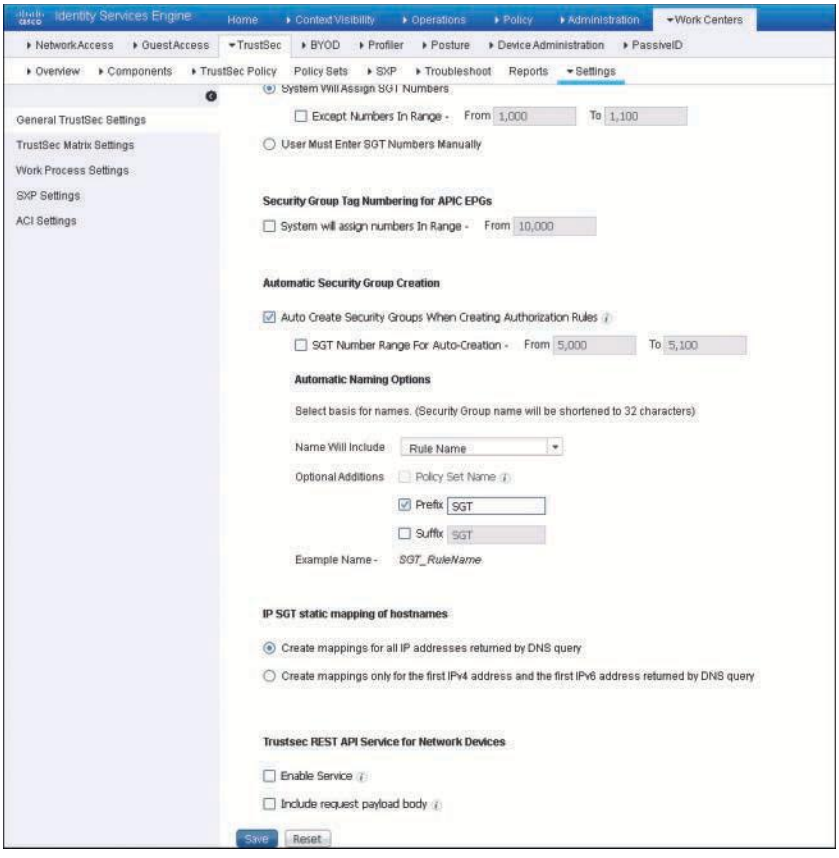


Figure 17-16 *TrustSec Settings*

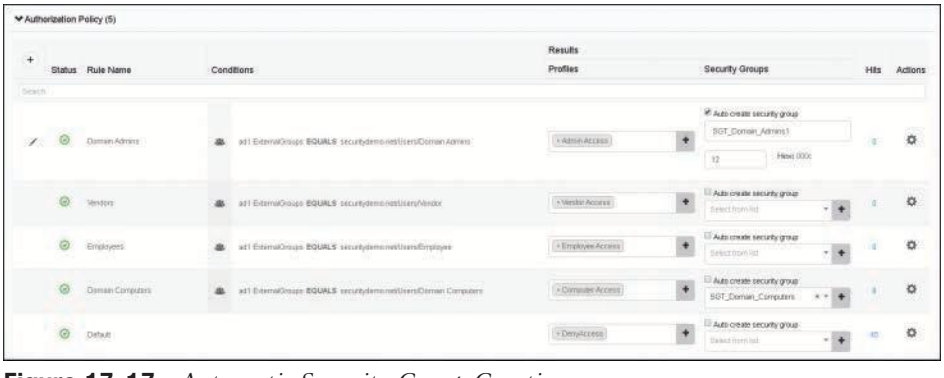


Figure 17-17 *Automatic Security Group Creation*

Classification

Key Topic

In order to use SGTs in your network infrastructure, the network devices must be able to understand and process SGTs. All Cisco Secure Access–supported switches and wireless

controllers support the assignment of SGTs. The assignment of an SGT to an authentication session is called *classification*. The process of communicating the assigned SGT upstream into the network is called *transport*, and it can either occur via native inline tagging or via a peering protocol called *SGT Exchange Protocol* (*SXP*).

Figure 17-18 shows an example of one access switch that has native tagging. The packets get tagged on the uplink port and remain tagged as the packets move through the infrastructure. This figure also shows a non-native-tagging-capable switch that uses SXP to update the upstream switch with the mapping of IP addresses to SGT tags. In both cases, the upstream switch continues to tag the traffic throughout the infrastructure.

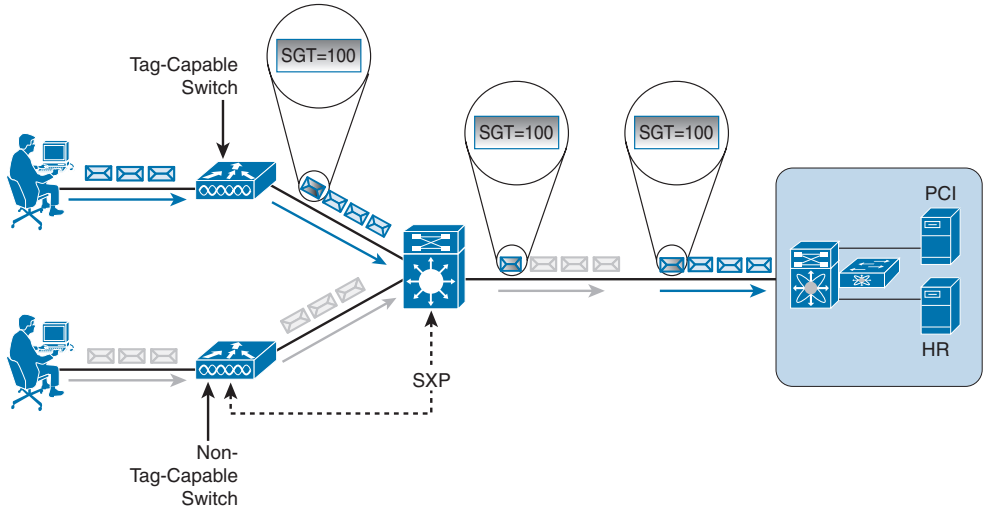


Figure 17-18 *Security Group Tagging*

An SGT needs to be assigned (classified) before it can be used on the network. This may happen dynamically as an authorization result from ISE. The endpoint may be authenticated through 802.1X, WebAuth, or MAB and still deliver an SGT as an authorization result. ISE delivers the tag to the NAD by using the following AVP:

cisco-av-pair=cts:security-group-tag=*tag-number*

SGTs can also be assigned statically. Static assignments are usually reserved for hosts that are not changing their IP addresses regularly, such as data center servers or topology-based policies. There are numerous options available for static assignment:

- **Layer 2 port**: You can configure a port to have a statically assigned SGT. This means that any device connecting to this port will have its incoming packets tagged with the statically configured SGT.
- **Layer 3 port**: You can configure a Layer 3 interface that maps directly to an SGT. The Layer 3 interface may be a routed port, an SVI, a Layer 3 subinterface, or even a tunnel interface.
- **VLAN**: Hosts in a specific VLAN can be mapped to a specific static SGT. This method can be used when there are third-party switches or Cisco switches that do not support

TrustSec. A TrustSec-enabled switch would discover the IP addresses of hosts on that VLAN via ARP and tag the packets coming from hosts on that VLAN.

- **IP address or subnet**: An SGT can be manually bound to an IP address or a subnet on the switch or bound using ISE.

The binding options for static assignment tend to vary depending on what the network access device assigning the tag supports.

Dynamically Assigning SGT via 802.1X

Assigning a tag is as simple as adding it as another permission or as a result of an authorization rule in the authorization policy. Follow these steps in the ISE GUI:

Step 1. Navigate to **Policy > Policy Sets**.

Step 2. Expand an existing policy set.

Step 3. Expand the authorization policy.

Step 4. Click the **+** button in the Security Groups column next to the rule to create a new SGT or select an existing SGT from the dropdown list.

Step 5. Click **Save**.

Figure 17-19 shows some of the authorization rules, using both a standard authorization profile result and a security group tag.

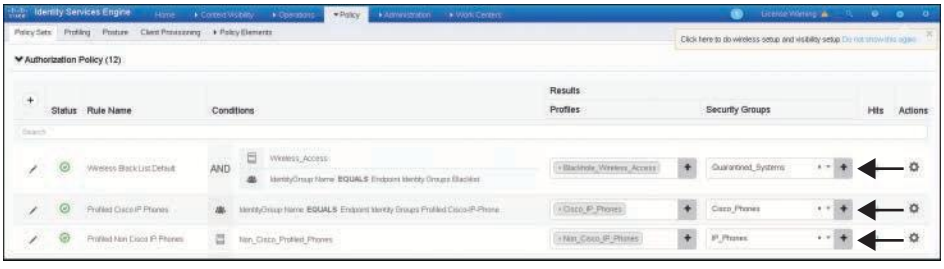


Figure 17-19 *Authorization Rules with SGT Assignment*

Manually Assigning SGTs to a Port

In most cases, 802.1X is not used in the data center. Servers are not usually required to authenticate themselves to the data center switch as the data center is normally considered physically secure, and there is no network access control applied there. However, the servers are the destinations for traffic coming from the campus and from within the data center.

Because 802.1X is not typically used in the data center, you need a way to manually apply an SGT. You can configure an SGT at the interface level of a supported Nexus switch and manually apply it to the port by entering the following commands:

```
N7K-DIST(config)#  int eth1/3

N7K-DIST(config-if)# cts manual

N7K-DIST(config-if-cts-manual)# policy static sgt 0x3
```

These commands manually assign the SGT 3 to the port on the Nexus 7000. However, some Nexus switch models do not support SGTs or are limited in terms of the TrustSec features they support.

Manually Binding IP Addresses to SGTs in ISE

As an alternative to assigning an SGT to a port, ISE enables you to centrally configure a database of IP addresses and their corresponding SGTs, and then SGT-capable devices can download that list from ISE. To manually bind an IP address to an SGT in ISE, follow these steps in the ISE GUI:

Step 1. Navigate to **Work Centers > TrustSec > Components > IP SGT Static Mapping**.

Step 2. Click **Add**.

Step 3. From the dropdown menu, choose **IP Address(s)** if you want to enter an IP address or a subnet to statically assign an SGT to. Optionally, you can choose **Hostname** from the dropdown to have ISE map the SGT to the IP address(es) returned by the DNS query.

Step 4. Select the appropriate SGT from the dropdown.

Step 5. (Optional) If you have SXP configured, you can send the static mapping to an SXP domain by selecting it from the dropdown.

Step 6. To send the IP SGT static mapping to the network access devices, choose the network device group that these bindings should be sent to from the **Deploy to Devices** dropdown.

Step 7. Click **Submit** to save.

Figures 17-20 through 17-22 provide examples of mapping an IP address, subnet, and hostname to an SGT named Production_Servers and setting ISE to deploy this static mapping to all network access devices in the Switches device group.

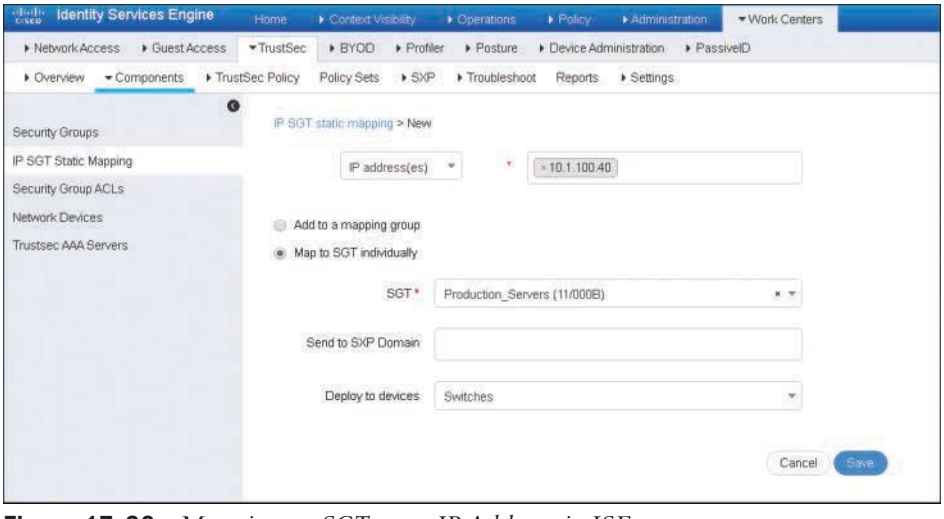


Figure 17-20 *Mapping an SGT to an IP Address in ISE*

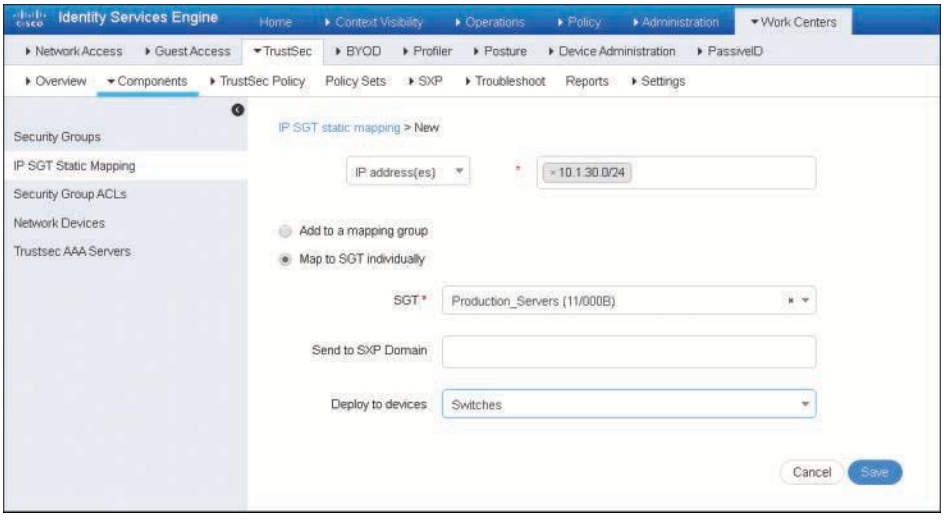


Figure 17-21 *Mapping an SGT to a Subnet in ISE*

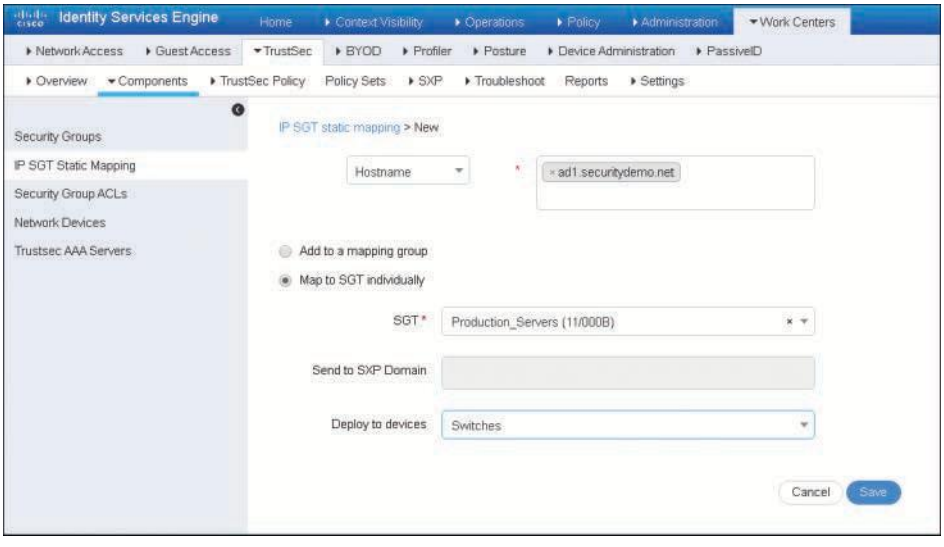


Figure 17-22 *Mapping an SGT to a Hostname in ISE*

Repeat steps 2 through 5 for any other static IP address- or hostname-to-SGT mappings required for your organization. Figure 17-23 shows the IP SGT Static Mapping screen with multiple mappings.

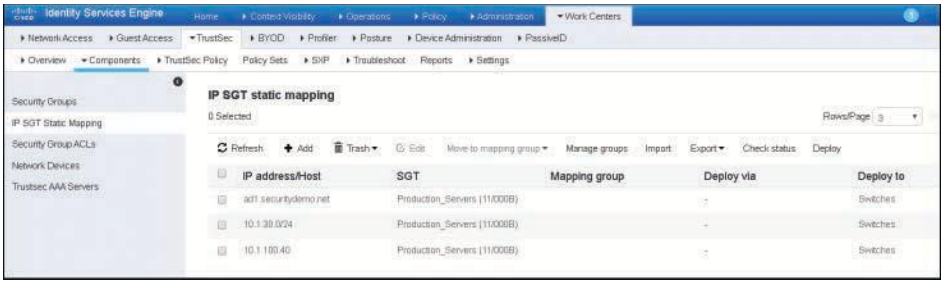


Figure 17-23 *IP SGT Static Mapping Screen with Multiple IP Address-to-SGT Mappings*

Now that this mapping exists, it can be deployed to the network access devices.

Access-Layer Devices That Do Not Support SGTs

Because we do not live in a perfect world and not all the equipment on a network will be the latest and greatest, you might find yourself in situations where you need another way to classify endpoint traffic. For example, a third-party switch might not support TrustSec and therefore might not be able to accept the SGT classification from ISE or send an update via SXP.

In addition, a VPN concentrator or some other type of third-party equipment might have found its way into the deployment. While that gear might not support the classification and transport natively, it may be capable of assigning different VLANs or IP addresses per authorization result.

You can map subnets and VLANs and assign all source IP addresses from the subnet or VLAN to a specific tag at a distribution-layer device. Doing so ensures that tagging takes place past the access layer.

NOTE To view what hardware supports TrustSec and which features it supports, check the Cisco TrustSec Platform Capability Matrix at https://www.cisco.com/c/en/us/solutions/enterprise-networks/trustsec/solution-overview-listing.html.

Mapping a Subnet to an SGT

The **cts role-based sgt-map** [*ipv4-subnet* | *ipv6-subnet*] **sgt** *tag-value* command enables subnet-to-SGT binding:

```
C9K-DIST(config)# cts role-based sgt-map 192.168.26.0/24 sgt 4
```

When you use this command, the device tracking feature in the switch is used to identify matches and assign the SGT.

Mapping a VLAN to an SGT

The **cts role-based sgt-map vlan-list** *vlans* **sgt** *tag-value* command enables VLAN-to-SGT binding:

```
C9K-DIST(config)# cts role-based sgt-map vlan-list 40 sgt 4
```

When you use this command, the device tracking feature in the switch is used to identify matches and assign the SGT.

Transport: SGT Exchange Protocol (SXP)

In a perfect world, all of your network access devices would support tagging the users traffic natively. Unfortunately, we do not live in a perfect world, and not all devices support native tagging. However, all supported Cisco switches do support the assignment of SGTs to authentication sessions (through the process of classification).

In many environments, network equipment supports SGTs, but the WAN connection does not. Or an environment may have a Cisco access layer, and their core or firewall may not support SGTs. In such situations, the SGT needs to be stripped from the packet before being passed to these devices or the WAN, and then it can be retagged by the TrustSec-supported devices after it crosses the unsupported device or network. In order for a TrustSec-supported device to retag a packet, it needs to know the IP address-to-SGT binding. This is where SXP comes in.

Key Topic

Cisco has developed a peering protocol (similar to BGP) to allow devices to communicate their database of IP address-to-SGT mappings to one another. This peering protocol is called Scalable Group Tag (SGT) Exchange Protocol (SXP). Because SXP is a peering protocol, it is possible to be very specific and deterministic about which devices send updates and which ones receive updates. SXP enables TrustSec to communicate IP address-to-SGT bindings across non-TrustSec clouds, a WAN, or brownfield environments that are not composed entirely of Cisco network access devices.

Another use case for SXP is to create strategic TrustSec enforcement points. As mentioned earlier in this chapter, a network access device needs to know the source and destination SGTs in order to enforce its TrustSec security policy. With inline tagging, transit network devices only know the source tag in the header until it reaches the destination network access device that has the destination host connected and the SGT of that host in its table. That network device then verifies and enforces the security policy. However, this might not be optimal. For example, if the destination is across the WAN and the security policy states that the packets should be dropped, it might be more desirable to make that enforcement decision before sending the traffic across the WAN. Strategic enforcement points on a campus or branch core or distribution switch can be created by using SXP. Since this switch would be populated by SXP with the IP address-to-SGT bindings of remote hosts, it may enforce the security policy locally.

An SXP peer may be a speaker, a listener, or both. A *speaker* is a device that sends the IP address-to-SGT bindings. A *listener* is a device that receives the bindings.

SXP communicates on TCP port 64999 for the connection initiation between peers. It uses an MD5 hash of a configured password for an authentication integrity check. In order to form a connection, the two SXP peers must negotiate and agree on the SXP version. As this book went to press, there were four SXP versions:

- **Version 1**: This is the initial SXP version, which supported IPv4 binding propagation.
- **Version 2**: This version adds support for IPv6 binding propagation and version negotiation.
- **Version 3**: This version adds support for subnet SGT binding propagation.
- **Version 4**: This version adds support for bidirectional SXP, loop detection and prevention, capability exchange, and a built-in keepalive mechanism.

TrustSec and SXP are Cisco-proprietary protocols. However, Cisco has submitted a draft to the IETF to standardize it (see https://tools.ietf.org/html/draft-smith-kandula-sxp-09).

SXP Design

Because SXP uses TCP as its transport, the peers may be Layer 2 adjacent or multiple Layer 3 hops away, as long as the TCP packets between the peers are routable and the peers are reachable. A network device can peer with an upstream switch for tag propagation or can even peer directly with the enforcement device (the data center switch or security group firewall). Figure 17-24 illustrates an example of this SXP peering over a non-SGT supported core.

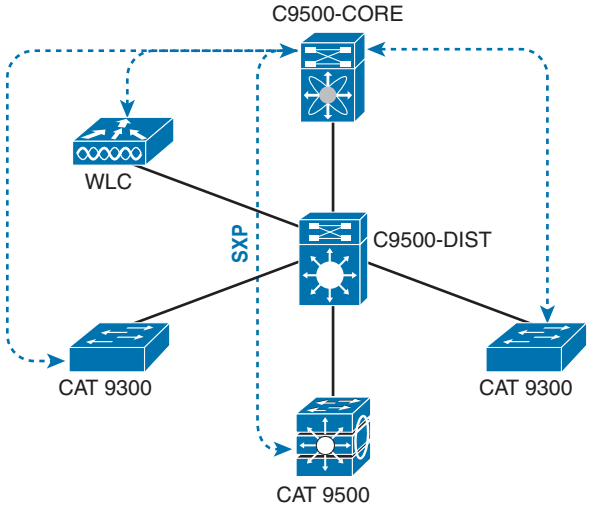


Figure 17-24 *SXP Peering*

Routing protocols have a limitation in terms of the number of neighbors they can scale, and so does SXP. Due to the limitations of scale and the number of peers, SXP design may be architected to be multi-hop, which allows for aggregation points. Distribution or core switches and edge routers can be good choices for this. These aggregation points can be great enforcement points as well. With SGTs, the enforcement device can be any device that knows the source and destination SGTs and can determine whether the traffic should be allowed or denied. Figure 17-25 illustrates SXP multi-hop.

There are numerous benefits to such a design. In addition to not requiring SXP-aware infrastructure along every hop in the network path, it provides a deterministic scalable design.

As of ISE Version 2.0, ISE can participate as an SXP peer and provide a central point in the SXP architecture to deliver the IP address-to-SGT binding. This can help in preventing the need for a fully meshed SXP design, which could become complicated. Figure 17-26 shows an example of a fully meshed SXP design in which every edge branch router needs to be aware of every IP address-to-SGT mapping for the other branches that will communicate with it. In this design, each branch has to be peered with the other branch routers.

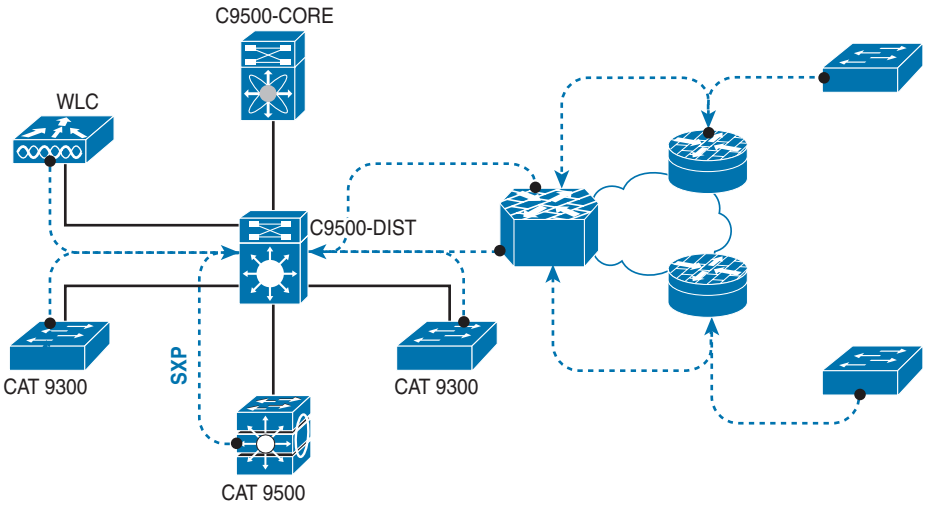


Figure 17-25 *SXP Multi-Hop*

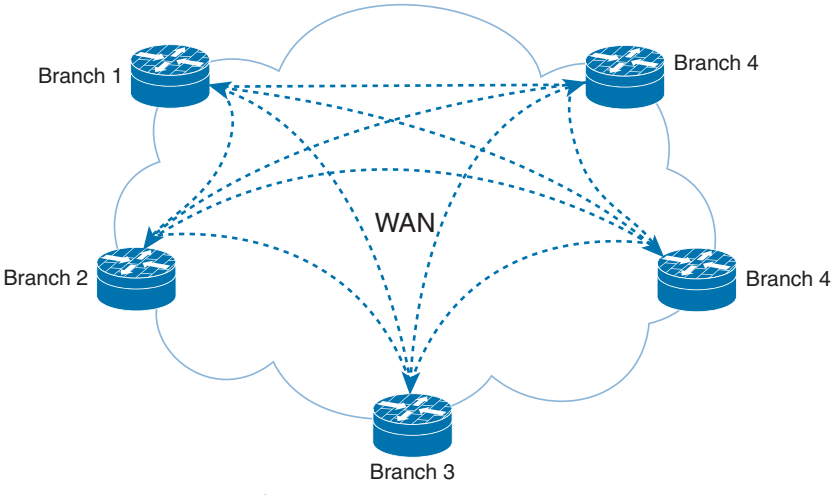


Figure 17-26 *Meshed SXP Design*

However, this design can be simplified and centralized by using ISE as the SXP peer. In this case, each branch router is peered with ISE via SXP. The routers share their IP address-to-SGT binding information with ISE, and ISE, in turn, shares its IP address-to-SGT table with each router. This centralizes the SXP design by having ISE become the single source of truth for those bindings. In ISE, you also have the option to statically configure IP address-to-SGT bindings and add them to all the network devices in the SXP domain. As of Version 2.4, ISE can scale up to 800 SXP peers. Figure 17-27 illustrates this centralized ISE SXP design.

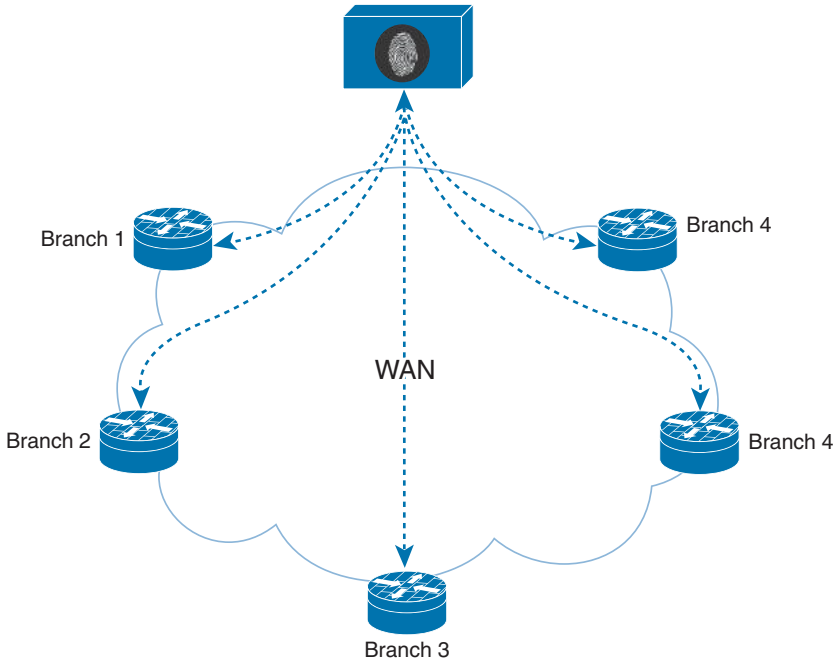


Figure 17-27 *Centralized ISE SXP Design*

Configuring SXP on ISE

In this exercise, we are going to configure ISE as both an SXP listener and speaker and connect it to a switch.

Step 1. Navigate to **Administration > System > Deployment** and click on the name of the PSN that will be running SXP services.

Step 2. In the Deployment menu, check the box next to **Enable SXP Service** (see Figure 17-28) and click **Save.** This enables the SXP service globally for ISE. Without this service enabled, you cannot configure SXP in ISE.

Step 3. In ISE, navigate to **Work Centers > TrustSec > Settings > SXP Settings.** This menu (see Figure 17-29) is where global SXP settings can be configured, such as timers and the global password. The following SXP timers are available:

- **Minimum acceptable hold time:** This is the shortest period of time a speaker is willing to send keepalive messages for keeping the connection alive.
- **Reconciliation timer:** If the SXP connection goes down and it recovers before the delete hold time, the reconciliation timer is started to clean up the old bindings from before it went down.
- **Maximum and minimum hold time:** This the maximum and minimum acceptable hold time period for the listener.

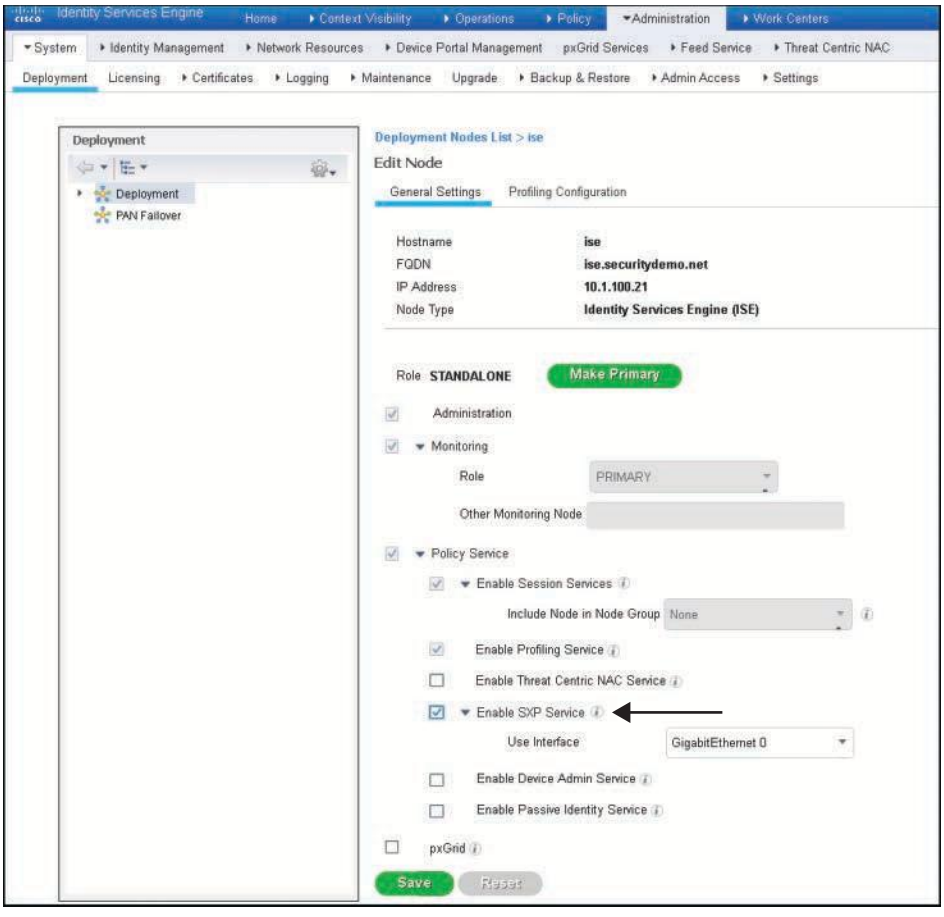


Figure 17-28 *Enabling the SXP Service*

- **Retry Open timer:** This is triggered when there is at least one SXP connection that is not up. It continues until the SXP connection is set up or the retry timer is zero.

The recommendation is to keep these timers set to their defaults unless something different is required.

Step 4. Set the global password to **ISEc0ld**.

Step 5. Check the box next to **Publish SXP Bindings on PxGrid** to ensure that the SXP IP address-to-SGT mappings will be propagated to the pxGrid clients that are subscribed to ISE for session data.

Step 6. Click **Save**.

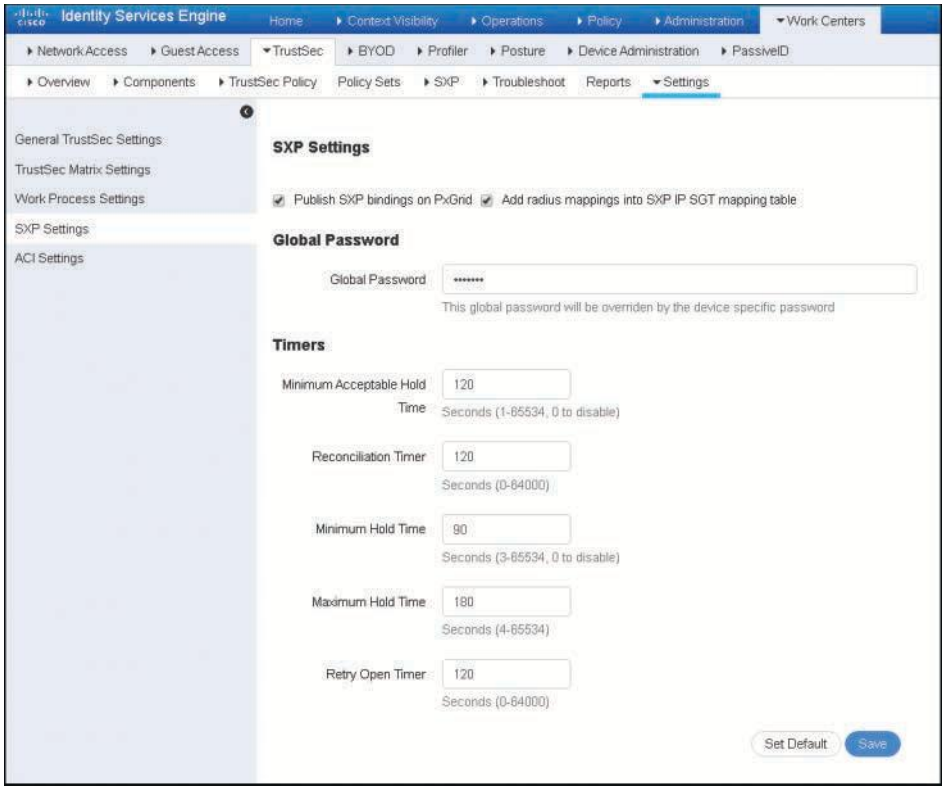


Figure 17-29 *SXP Settings Menu*

Step 7. Navigate to **Work Centers > TrustSec > SXP > SXP Devices**.

Step 8. Click **Add**.

Step 9. Add **Sw01** as the device name for your lab switch and type the IP address of your lab switch in the IP Address field.

Step 10. Select **BOTH** for the peer role.

Step 11. From the connected PSNs dropdown, choose a PSN that has the SXP service enabled.

Step 12. Ensure that SXP Domain is set to **default**.

Step 13. Ensure that Status is set to **Enabled**.

Step 14. Leave Password Type set to **Default** to use the global password that you previously configured to connect with the SXP peer.

Step 15. Leave Version set to **4**. It should negotiate automatically.

Step 16. (Optional) Expand **Advanced Settings** to modify the timers if you would like to change them for just this specific SXP peer.

Step 17. Click the **Save** button to save this SXP configuration. Figure 17-30 shows the completed SXP configuration for a switch.

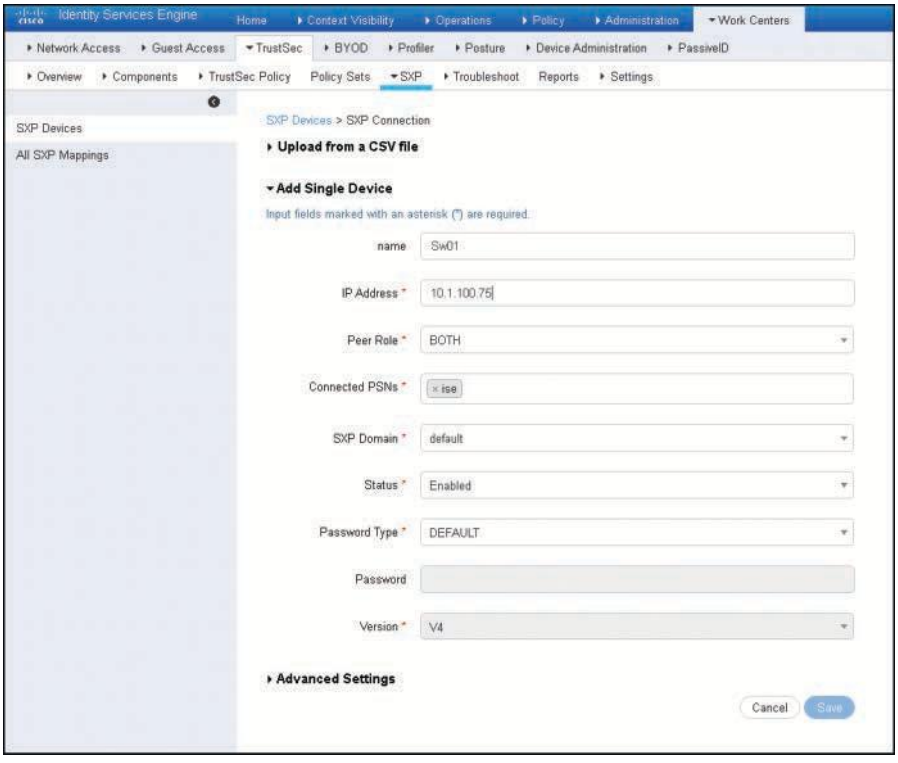


Figure 17-30 *SXP Peer Configuration Screen*

Configuring SXP on IOS Devices

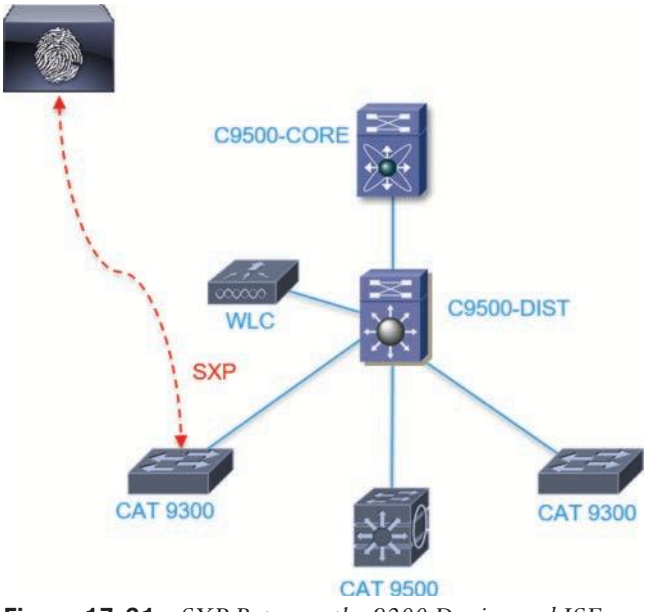


Figure 17-31 *SXP Between the 9300 Device and ISE*

You enable SXP globally from global configuration mode. Each peer needs to be added individually, and you need to set a global default SXP password. Figure 17-31 illustrates the connection. Follow these steps in global configuration mode on an IOS device:

Step 1. Enter **cts sxp enable**.

Step 2. (Optional) To configure the default SXP password, enter **cts sxp default password** *password*.

Step 3. (Optional) To configure the default source IP address that the IOS device should use to connect to the SXP peer, enter **cts sxp default source-ip** *ip-address*.

Step 4. (Optional) To make changes to the timer settings on the local IOS device, use the following commands:

- **Minimum acceptable hold time: cts sxp speaker hold-time** *seconds*
- **Reconciliation timer : cts sxp reconciliation period** *seconds*
- **Maximum and minimum hold time : cts sxp listener hold-time** *min-allowed-seconds max-allowed-seconds*
- **Retry Open timer: cts sxp retry period** *seconds*

Step 5. Enter **cts sxp connection peer** [*peer-ip-address*] **password** [**default | none**] **mode** [**local | peer**] [**listener | speaker | both**] to define the SXP peer. The options are as follows:

- **password default:** This states to use the password is defined globally for all SXP connections. It is currently not possible to have different SXP passwords for different peers.
- **password none:** Do not use a password with this SXP peer.
- **mode local:** States that the following argument defines the local side of the connection.
- **mode peer:** States that the following argument defines the peer's side of the connection.
- **listener:** Indicates that the specified device (local or peer) will receive SXP updates through this connection.
- **speaker:** Indicates that the specified device (local or peer) will send SXP updates through this connection.
- **both:** Indicates that the specific device (local or peer) will form a bidirectional SXP connection. This means that the device will act both as a listener and as a speaker.

For the IP address of the peer, use the IP address of your ISE server so you can form an SXP connection to it.

Step 6. In ISE, navigate to **Work Centers > TrustSec > SXP > SXP Devices** and ensure that the status is set to **ON** (see Figure 17-32). If you see the status **PENDING_ON**, wait approximately 30 seconds.



Figure 17-32 *SXP Between Sw01 and ISE*

Steps 2 through 4 are optional steps for defining connection passwords, source IP addresses, and timers. However, you can set up the SXP connection without changing these options.

Examples 17-1 and 17-2 show the commands for setting up the SXP connection between a 9300 (access-layer device) and ISE and then verifying it in the command line.

Example 17-1 *Enabling SXP on a 9300 Device*

```
Sw01-9300(config)# cts sxp enable
Sw01-9300 (config)#
*Apr 29 14:01:48.103:%CTS-5-SXP_STATE_CHANGE: CTS SXP enabled
Sw01-9300(config)# cts sxp connection peer 10.1.100.21 password default mode peer
speaker
Sw01-9300(config)#
*Apr 29 14:06:11.804: %CTS-6-SXP_TIMER_START: Connection <0.0.0.0, 0.0.0.0> retry
open timer started.
*Apr 29 14:06:11.804: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.100.21,
10.1.100.75>-1 state changed from Off to Pending_On.
Sw01-9300(config)#
*Apr 29 14:06:57.155: %CTS-3-SXP_CONN_STATE_CHG_OFF: Connection <10.1.100.21,
10.1.100.75>-1 state changed from Pending_On to Off.
Sw01-9300(config)# cts sxp default password ISEc0ld
Sw01-9300(config)#
*Apr 29 14:06:59.216: %CTS-5-SXP_DFT_PASSWORD_CHANGE: CTS SXP password changed.
```

Example 17-2 *Verifying the SXP Connection and Downloaded IP Address-to-SGT Bindings*

```
Sw01-9300(config)# show cts sxp connections brief
SXP                : Enabled
 Highest Version Supported: 4
 Default Password : Set
 Default Source IP: 10.1.100.75
Connection retry open period: 120 secs
Reconcile period: 120 secs
Retry open timer is not running
Peer-Sequence traverse limit for export: Not Set
Peer-Sequence traverse limit for import: Not Set
```

```
-------------------------------------------------------------------------------
Peer_IP            Source_IP               Conn Status                 Duration
-------------------------------------------------------------------------------
10.1.100.21      10.1.100.75      On(Speaker)::On(Listener)              0:00:11:44
(dd:hr:mm:sec)::0:00:11:44 (dd:hr:mm:sec)

Total num of SXP Connections = 1
Sw01-9300(config)# show cts sxp sgt-map brief
SXP Node ID(generated):0x0A01644B(10.1.100.75)
IP-SGT Mappings as follows:
IPv4,SGT: <10.1.100.44 , 8>
Total number of IP-SGT Mappings: 1
```

Configuring SXP on Wireless LAN Controllers

Cisco's Wireless LAN Controller added support for SGT classification and SXP beginning with Version 7.2. A wireless LAN controller can only act as a speaker with SXP to inform downstream network devices or ISE of connecting clients' IP address-to-SGT mapping. Follow these steps in the Wireless LAN Controller user interface to configure SXP on a WLC:

Step 1. Navigate to the **Security** tab.

Step 2. On the left-hand side, expand **TrustSec** and select **SXP Config**.

Step 3. Set SXP State to **Enabled**.

Step 4. Set the default password to the password of your choosing. All network devices that connect to this WLC as an SXP peer need to have the same password.

Step 5. Click **Apply** to apply the default password and enable SXP. You have now enabled SXP globally. Each peer needs to be added individually, as described in the following steps.

Step 6. Still on the SXP Configuration screen, enter the IP address of your SXP peer. In this example, you may add ISE as the SXP peer to test it.

Step 7. Click **Add**.

Step 8. Click **Apply** to apply the configuration.

The added peers are displayed on this SXP Configuration page. The SXP connection status is listed next to the IP address of the peer. Once the peer is configured on the other side of the connection, the status should change from Off to On.

Figure 17-33 shows a completed SXP Configuration page waiting for the connection to be established.

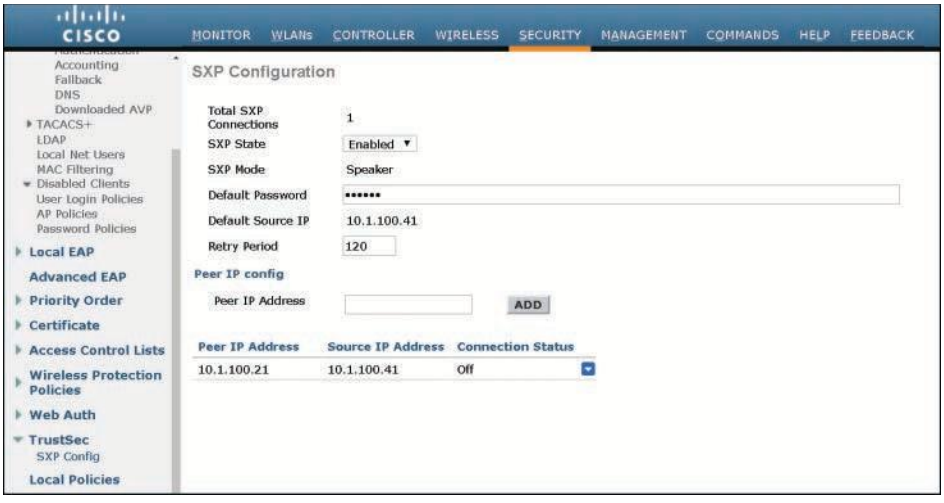


Figure 17-33 *TrustSec SXP Page: Peer Status*

It is also possible to verify an SXP connection from the peer side, as shown in Example 17-3.

Example 17-3 *Verifying the Connection Between a WLC and a Catalyst 9500 Device*

```
C9K-DIST# show cts sxp connections brief
SXP              : Enabled
 Highest Version Supported: 4
 Default Password : Set
 Default Source IP: 10.1.100.75
Connection retry open period: 120 secs
Reconcile period: 120 secs
Retry open timer is not running
Peer-Sequence traverse limit for export: Not Set
Peer-Sequence traverse limit for import: Not Set
---------------------------------------------------------------------------
Peer_IP          Source_IP        Conn Status       Duration
---------------------------------------------------------------------------
10.1.40.2        10.1.40.1        On                4:06:36:24 (dd:hr:mm:sec)
10.1.60.2        10.1.60.1        On                0:00:03:31 (dd:hr:mm:sec)

Total num of SXP Connections = 2
```

Configuring SXP on Cisco ASA

Cisco Adaptive Security Appliance (ASA) supports SGT enforcement through a security group firewall (SGFW). The ASA has supported TrustSec since Version 9.0.1. However, Cisco only added support for native inline tagging with Version 9.3.1 in the ASA 5500X series and Firepower appliances running ASA code. Most commonly, SXP is the method used for the transport of IP address-to-SGT bindings for the ASA, and this is the solution covered on the SISE 300-715 exam.

It is important to note that the ASA has multiple functions, including deep packet inspection (DPI) firewalling and remote-access VPNs. ASA Version 9.2.1 and newer can enforce SGTs, receive SGTs (transport), and assign SGTs to remote-access VPN users (classification). The ASA can be configured as an SXP speaker or as a listener.

Follow these steps in the ASA Device Manager (ASDM) to configure SXP on Cisco ASA:

Step 1. Navigate to **Configuration > Firewall > Identity by TrustSec**.

Step 2. Globally enable SXP by checking the **Enable SGT Exchange Protocol (SXP)** checkbox in the upper left.

Step 3. Click **Add** to add a new SXP peer.

Step 4. In the Add Connection Peer popup, add the IP address of the remote peer.

Step 5. Set Password to **Default** (unless you will not be using passwords).

Step 6. Set Mode to **Peer**.

Step 7. Set Role to **Speaker**.

Step 8. Click **OK**, and you are returned to the main Identity by TrustSec page. At this point, you have SXP enabled and a single peer defined but no default password yet.

Step 9. (Optional) Add the IP address of the source interface that you would like to use to communicate with the SXP peer.

Step 10. Enter the default password for your SXP deployment.

Step 11. Click **Apply** to apply the new configuration. Figure 17-34 shows the completed SXP configuration in the ASA.

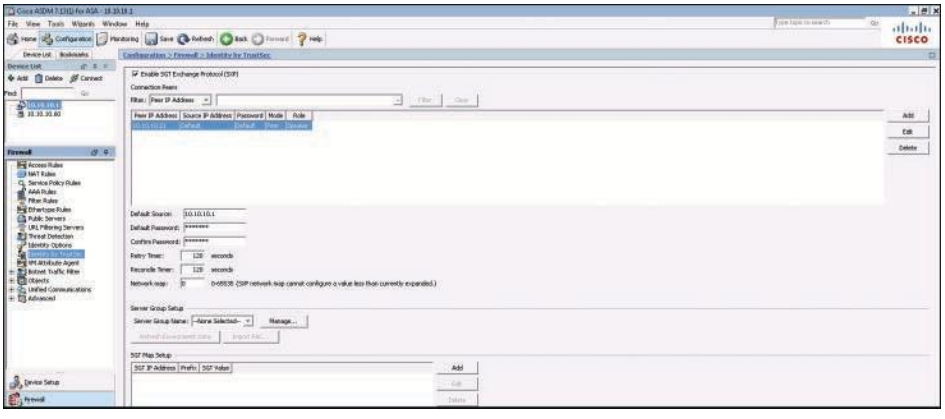


Figure 17-34 *SXP Connection Configuration in ASDM*

Verifying SXP Connections in ASDM

To verify SXP connections in ASDM, follow these steps in the ASDM:

Step 1. Navigate to **Monitoring > Properties > Identity by TrustSec**.

Step 2. Click on **SXP Connections** to see the configured peers and the status of each peer (see Figure 17-35).

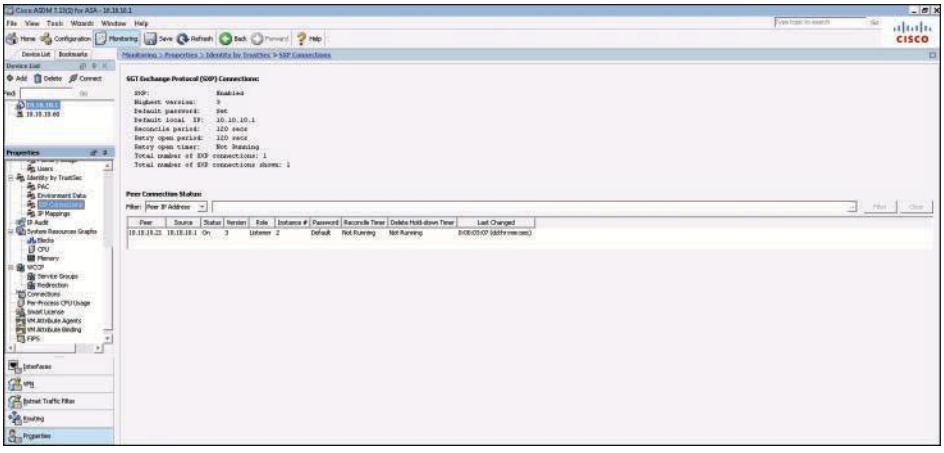


Figure 17-35 *Monitoring SXP Connections in ASDM*

Step 3. Click on **IP Mappings** to see any IP address-to-SGT mappings that the ASA has learned about (see Figure 17-36).

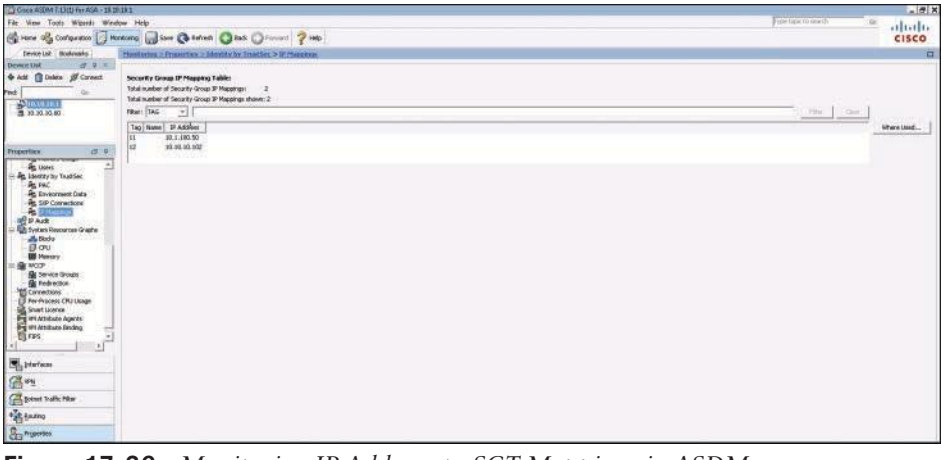


Figure 17-36 *Monitoring IP Address-to-SGT Mappings in ASDM*

Transport: Native Tagging

Native SGT tagging is the ultimate goal with any TrustSec deployment. With this approach, the access layer is capable of applying an SGT to the Layer 2 frame and carrying it across the wire to the upstream host. The upstream host retains the tag in the frame and ensures that the tag is applied as the traffic is sent out the egress interface toward the next hop and so on. Ideally, the tag is present throughout the entire infrastructure.

Figure 17-37 shows the format of the Layer 2 frame with the SGT value that exists in the Cisco Meta Data (CMD) field circled.

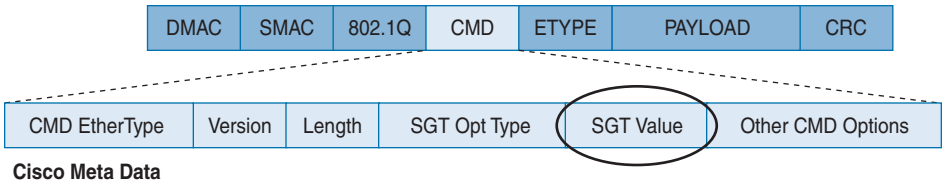


Figure 17-37 *Layer 2 Frame Format with SGT*

Native tagging makes it possible for the technology to scale to virtually endless size and remain completely independent of any Layer 3 protocol. In other words, architecturally speaking, it does not matter if the traffic is IPv4 or IPv6. The tag is completely independent. Figure 17-38 illustrates the pervasive tagging concept.

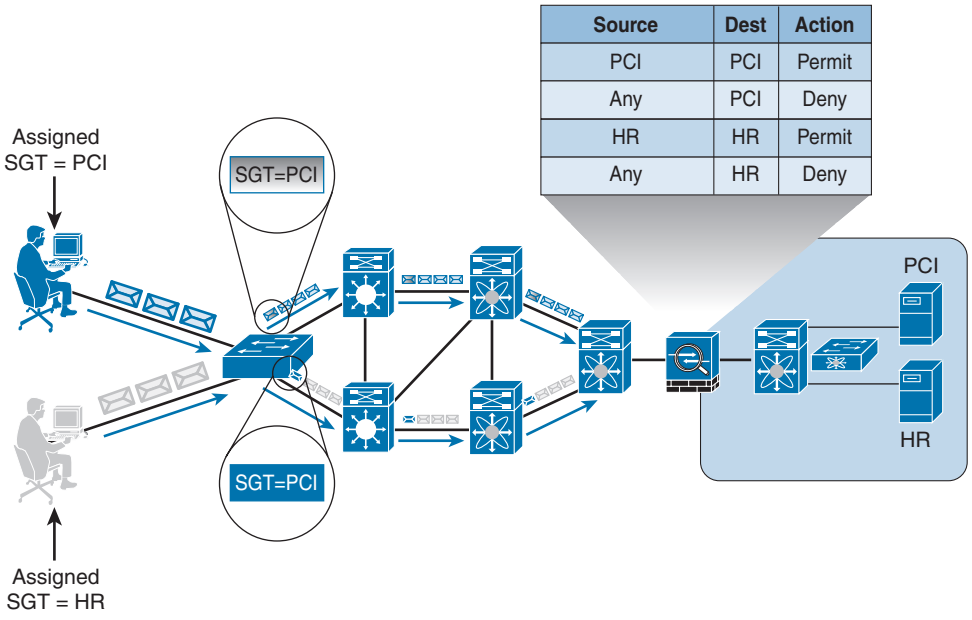


Figure 17-38 *Pervasive Tagging*

As you can see in Figure 17-38, when native tags are supported pervasively in the infrastructure, the SGT is communicated hop-by-hop. This provides for end-to-end segmentation and tremendous scalability. With the tag being applied to the traffic at every Layer 2 link, you can enforce policy at any point in the infrastructure without facing size limitations of the IP address-to-SGT mapping database as it is not being used at all.

For added security, an SGT can be encrypted with MACsec or IPsec, and the network infrastructure can be authenticated before SGTs are sent or received.

Configuring Native SGT Propagation (Tagging)

This section shows how to enable native tagging on an IOS XE switch. Figure 17-39 shows the logical network layout used in the following configuration examples.

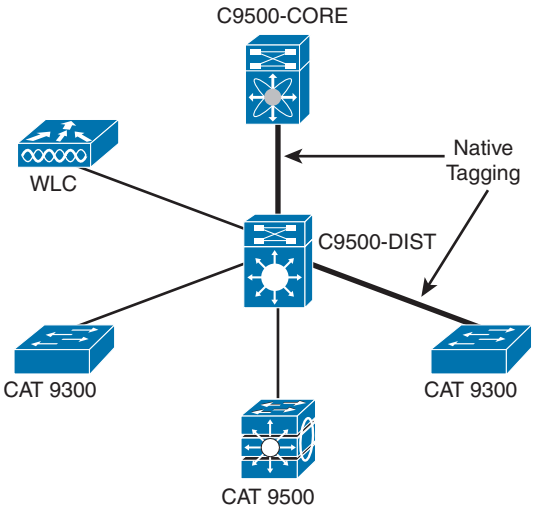


Figure 17-39 *SGTs from Access to Distribution to Core*

Configuring Manual SGT Propagation on Cisco IOS XE Switches

This section discusses the configuration of SGT propagation on access-layer switches such as the Catalyst 9300 and 9500 switches that have the ability to use native tags.

When it comes to inserting a tag into Layer 2 traffic, there is a fundamental choice to make: to use encryption or not to use encryption. For simplicity, this section focuses on the easy way, which is to not use encryption. (You will learn about encrypting the tag later in this chapter, when we discuss MACsec.)

Follow these steps in global configuration on a Catalyst 9300 switch:

Step 1. Enable Cisco TrustSec role-based enforcement on the switch:

```
CAT9300(config)# cts role-based enforcement
```

This globally enables the tagging of frames. It also makes it possible to enforce SGACLs (discussed later in this chapter). Without this command in the global configuration, the switch does not tag the Layer 2 traffic.

Step 2. Enter interface configuration mode on the tagging-capable port:

```
CAT9300(config)# interface g1/0/1
```

Step 3. Enable Cisco TrustSec manually on the interface:

```
CAT9300(config-if)# cts manual
```

Here you use **cts manual** because you are not using Network Device Admission Control (NDAC). The **cts manual** mode of operation allows you to send the applied tags to the Layer 2 frame without needing to negotiate encryption or requiring a fully trusted domain of Cisco switches, as you would need with NDAC.

Step 4. Configure the SGT for the port and trusted tags received on this port:

```
CAT9300(config-if-cts-manual)# policy static sgt
sgt-value trusted
```

Configuring this on the port defines what SGT to apply to incoming untagged or non-trusted traffic. The **trusted** keyword allows you to trust the received SGTs that the peer is sending.

When you created security groups earlier in this chapter, you created a special group for network access devices, called it NADs, and made the value of that group 2 (0x02). That is the value you are applying here with the **policy static sgt 2 trusted** command. The **trusted** keyword in this command ensures that no changes are made to the incoming tags, as they are from a trusted source. Example 7-4 provides an example of the configuration of static SGT assignment to a port. Example 17-5 demonstrates the verification of the SGT assignment of a port.

Example 17-4 *Enabling Tagging on a 9300 Series Access Switch*

```
C9300(config)# cts role-based enforcement
C9300(config)# interface g1/0/1
C9300(config-if)# cts manual
C9300(config-if-cts-manual)# policy static sgt 2 trusted
```

Example 17-5 *Verifying Tagging on a 9300 Series Access Switch*

```
C9300# sho cts interface Gig 1/1/1
Global Dot1x feature is Enabled
Interface GigabitEthernet1/0/3:
    CTS is enabled, mode:    MANUAL
    IFC state:               OPEN
    Interface Active for 00:05:23.250
    Authentication Status:   NOT APPLICABLE
        Peer identity:       "unknown"
        Peer's advertised capabilities: ""
    Authorization Status:    SUCCEEDED
        Peer SGT:            2:TrustSec_Devices
        Peer SGT assignment: Trusted
    SAP Status:              NOT APPLICABLE
    Propagate SGT:           Enabled
    Cache Info:
        Expiration            : N/A
        Cache applied to link : NONE

    Critical-Authentication: Disabled
        Peer SGT: 0
        Peer SGT assignment: Untrusted
        Default PMK: Not Configured
        Default SGACL:
        Fail-Open: Enabled
    Statistics:
        authc success:              0
        authc reject:               0
        authc failure:              0
```

```
        authc no response:          0
        authc logoff:               0
        sap success:                0
        sap fail:                   0
        authz success:              2
        authz fail:                 0
        port auth fail:             0

    L3 IPM:   disabled.
```

Enforcement

Now that you have security groups being assigned (classification), and they are being transmitted across the network (transportation), it is time to focus on the third staple of security group access: enforcement.

There are multiple ways to enforce traffic, based on the tag, but, ultimately, they can be summarized into two major types:

- Enforcement on a switch (SGACL)
- Enforcement on a firewall (SGFW)

SGACL

Key Topic

At first, enforcement with ***security group ACLs (SGACLs)*** was the only option available with TrustSec. Major benefits to the use of SGACLs are the consolidation of ACEs and the operational savings involved with maintenance of traditional ACLs.

An SGACL looks very similar to a spreadsheet. It is always based on a source tag to a destination tag. Figure 17-40 shows an example of an SGACL policy on ISE that represents the SGACLs in a columns-and-rows format.

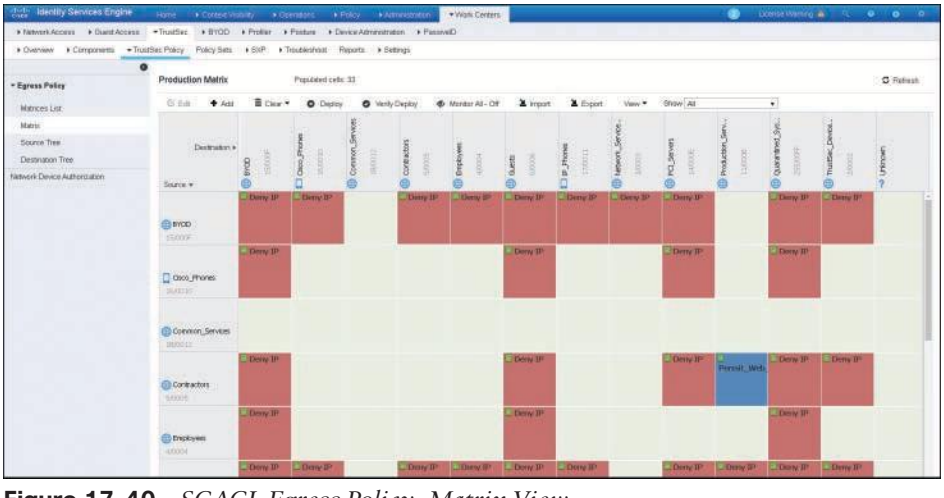


Figure 17-40 *SGACL Egress Policy: Matrix View*

The box that intersects between Contractors and Production_Servers on the matrix shows that if any traffic with the source tag Contractors (5) attempts to reach a destination with the Production_Servers tag (11), an SGACL named Permit_WEB_RDP will be applied, along with the catchall Deny IP. You can see the contents of the Permit_WEB_RDP SGACL by hovering the mouse cursor over the bull's-eye icon in the highlighted cell. In Figure 17-41, you can see that only HTTP, HTTPS, and RDP traffic is permitted.



Figure 17-41 *Permit_WEB_RDP SGACL Contents*

As you can see in Figures 17-40 and 17-41, the resulting SGACL permits HTTP (TCP port 80), HTTP (TCO port 443), and RDP (TCP port 3389) and denies all other traffic. This traffic is applied to the egress of the switch where the SGACL is configured.

This form of traffic enforcement can provide a tremendous savings in terms of the complexity and number of ACEs to maintain. Recall from earlier in this chapter that there is a general formula you can use to see the savings:

Number of sources × Number of destinations × Permissions = Number of ACEs

With a traditional ACL on a firewall you might see the following formula:

4 VLANs (src) × 30 (dst) × 4 permission = 480 ACEs

This is what you would see for the source IP address on the port using a dACL:

1 group (source) × 30 (dst) × 4 permission = 120 ACEs

With SGACLs, the number of ACEs is an order of magnitude smaller:

4 SGT (src) × 3 SGT (dst) × 4 permission = 48 ACEs

There are two main conceptual paths for deploying SGACLs:

- **North–south**: North–south refers to the use case of a user or device being classified at the access layer but enforcement with the SGACL occurring at the data center. For example, a guest entering the access layer is assigned a Guest SGT. Traffic with a Guest SGT is dropped if it tries to reach a server with financial data.
- **East–west**: East–west refers to the use case of an SGACL protecting resources that exist on the same switch or LAN. For example, a development server and a production

server may reside on the same switch in the data center, and an SGACL may be deployed to prevent the development server from ever communicating with the production server. Another east–west example is a guest and an employee both using the same access-layer switch. Traffic could be filtered between these two devices so the guest cannot communicate to the employee who is in the same VLAN on the same switch. This is also referred to as peer-to-peer traffic.

Figure 17-42 illustrates the concepts of north–south and east–west traffic flow enforcement.

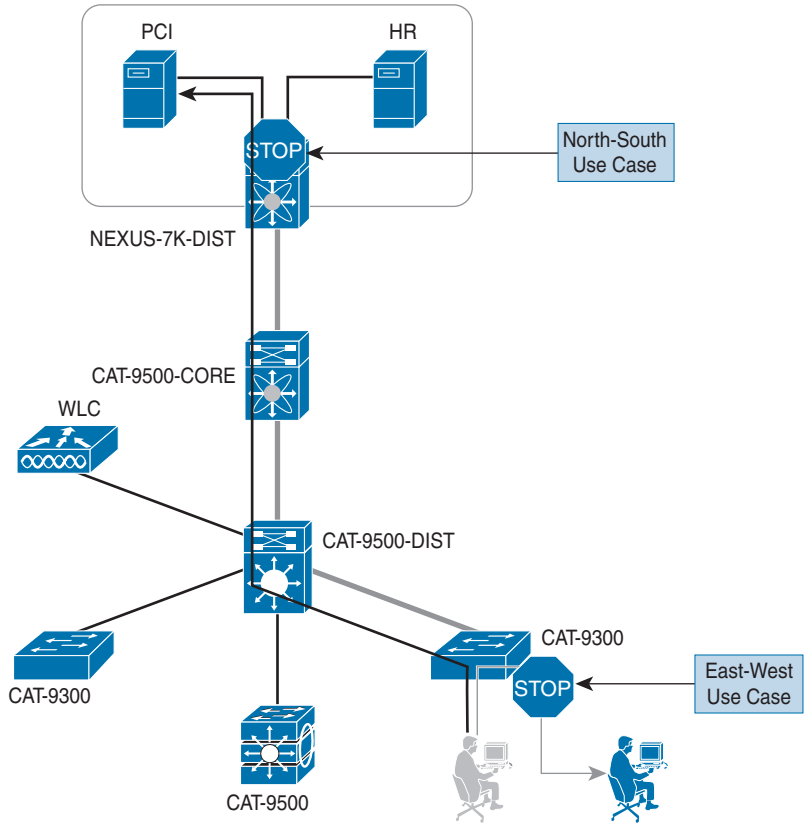


Figure 17-42 *North–South and East–West Traffic Flow Enforcement*

SGACLs are created in ISE. However, they are not part of the security policy until they are applied to a cell in the TrustSec policy matrix, illustrated in Figures 17-40 and 17-41, and that TrustSec policy matrix or SGACL policy is applied to the network devices.

SGACLs created in ISE are configured before they are applied to the TrustSec policy matrix. The syntax of the SGACLs configured in ISE is unlike that of IP ACLs. Since the source and destination are defined by the cell position in the TrustSec Policy Matrix, the SGACL is written without a source or destination. Table 17-2 lists the syntax for SGACLs for IOS, IOS XE, and NX-OS operating systems.

Table 17-2 SGACL Syntax

SGACL CLI and ACEs	Syntax Common Across IOS, IOS XE, and NX-OS
deny permit	ahp, eigrp, gre, icmp, igmp, ip, nos, ospf, pcp, pim, tcp, udp
deny tcp deny tcp src deny tcp dst	dst, log, src
deny tcp dst eq deny tcp src eq	0–65535
deny udp deny udp src deny udp dst	dst, log, src
deny tcp dst eq www deny udp src eq www	0–65535

After an SGACL is created, it is inserted into a cell in the TrustSec policy matrix, which defines the source and destination tag. The TrustSec policy matrix has a default rule: permit or deny. If a cell does not have an SGACL filled in, the SGT source and destination will communicate based on this default rule. Figure 17-43 illustrates this situation. In this figure, there are no SGACLs applied to any of the cells in the matrix, so all communication will be allowed as the default rule is to permit everything. The only way to override this default rule is to explicitly define inbound and outbound egress policies in the TrustSec policy matrix.

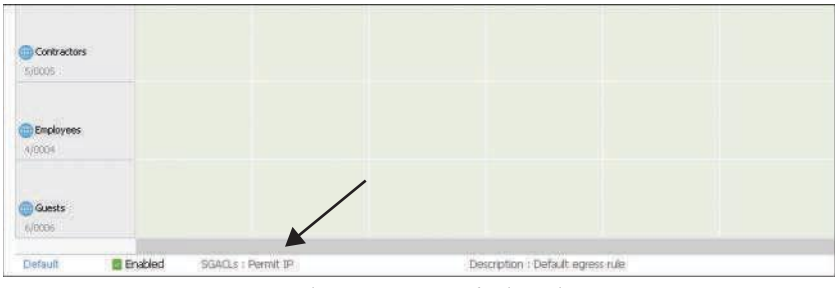


Figure 17-43 *TrustSec Policy Matrix Default Rule*

A cell can be populated with only a single SGACL (which is the default), or ISE can be configured to allow multiple SGACLs in that same cell.

It is possible to have multiple TrustSec matrices. For example, there might be certain locations in high-risk countries that need more restrictive policies, but the same users should receive the normal corporate policies when they are at the corporate office. Instead of creating more SGTs and a more complex TrustSec matrix for these high-risk users, you can create another TrustSec policy matrix and apply it to only the network devices at the location that needs it.

The ability to have multiple matrices means you can have a defcon matrix that can be quickly deployed in an emergency if there is a breach and network access needs to be locked down even further than it is with the regular TrustSec policy while the security team responds.

Configuring Security Group ACLs

The following steps walk through the configuration of a couple SGACLs in ISE for use in the TrustSec policy:

Step 1. In the ISE GUI, navigate to **Work Centers > TrustSec > Components > Security Group ACLs.**

Step 2. Click **Add** to create a new SGACL.

Step 3. Name the new ACL **Permit_Mgmt** and grant this ACL web, SSH, and RDP destination traffic (see Figure 17-44). Click **Save** when completed.

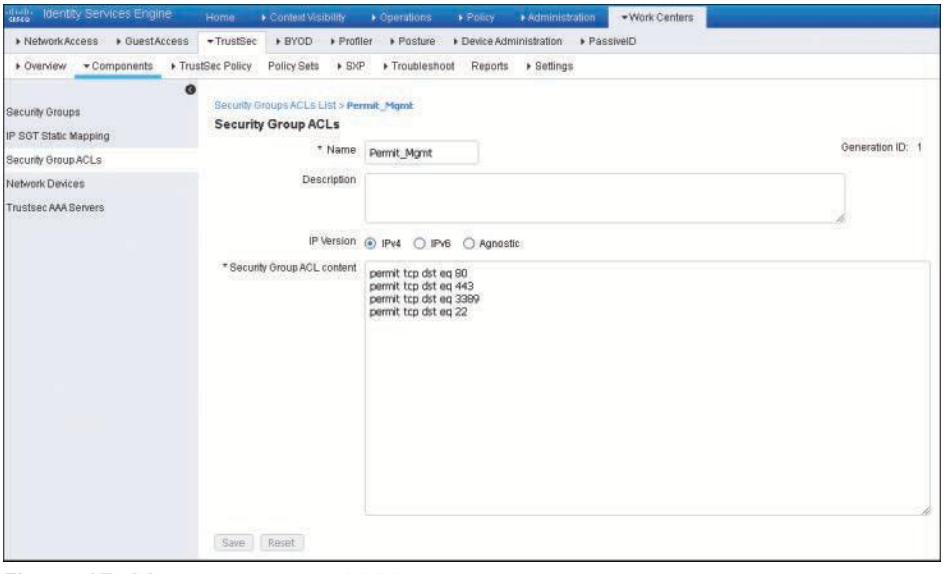


Figure 17-44 *Permit_Mgmt SGACL*

Step 4. Click **Add** again to create another new SGACL.

Step 5. Name the new ACL **Deny_All** and put a deny statement in it to block all traffic (see Figure 17-45).

Step 6. Click **Add** to add another SGACL.

Step 7. Name this new ACL **Permit_HTTP_HTTPS** and permit HTTP and HTTPS for the destination (see Figure 17-46).

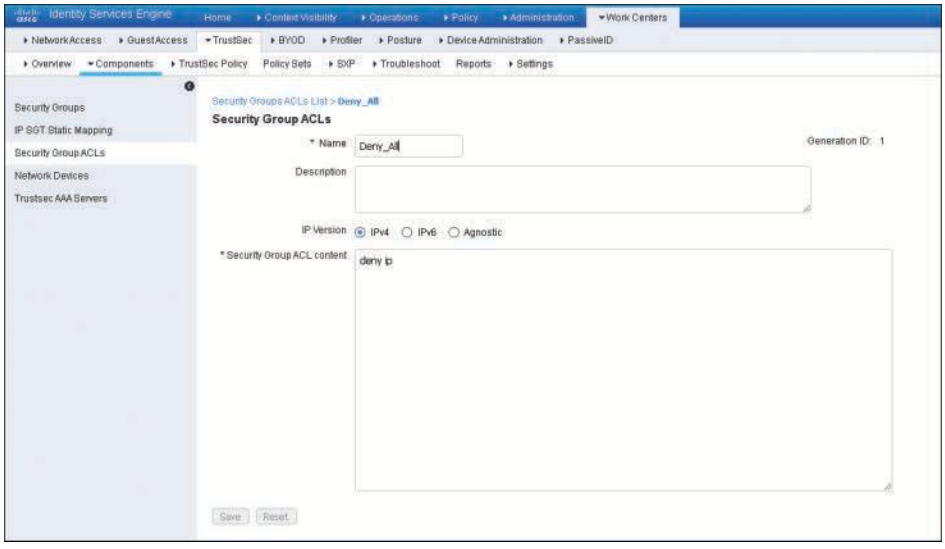


Figure 17-45 *Deny_All SGACL*

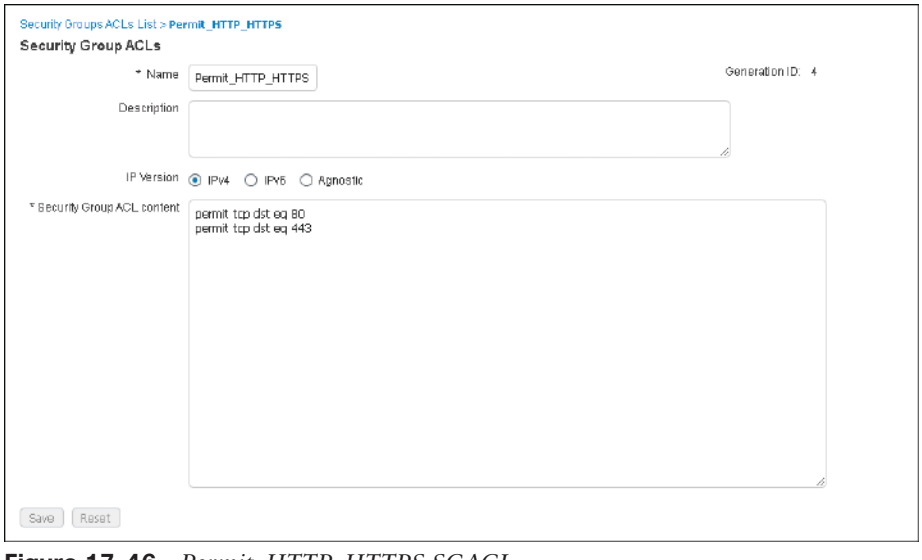


Figure 17-46 *Permit_HTTP_HTTPS SGACL*

Step 8. Click **Add** to add another SGACL.

Step 9. Name this new ACL **Permit_ICMP** and permit ICMP traffic (see Figure 17-47).

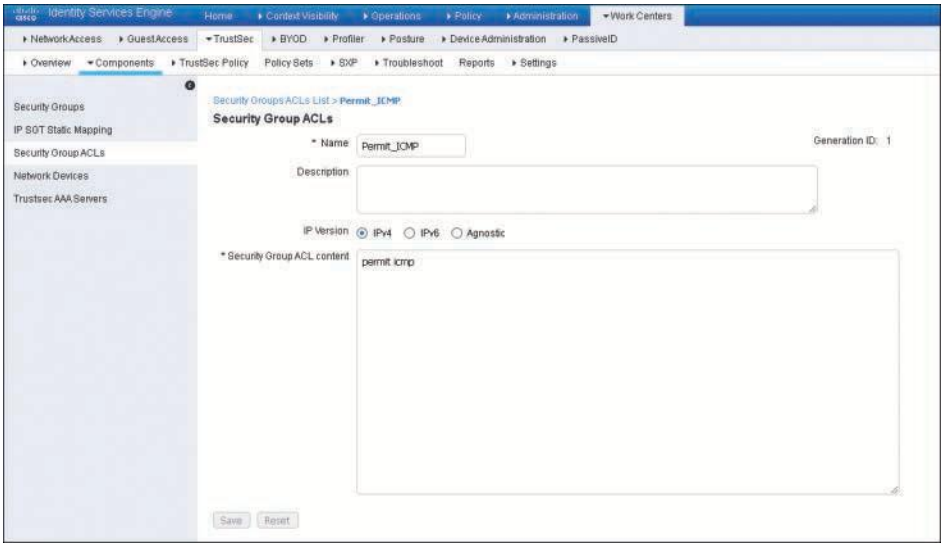


Figure 17-47 *Permit_ICMP SGACL*

Step 10. Click **Add** to add another SGACL.

Step 11. Name this new ACL **Permit_SRC_HTTP_HTTPS** and permit access with a source port of HTTP or HTTPS (see Figure 17-48).

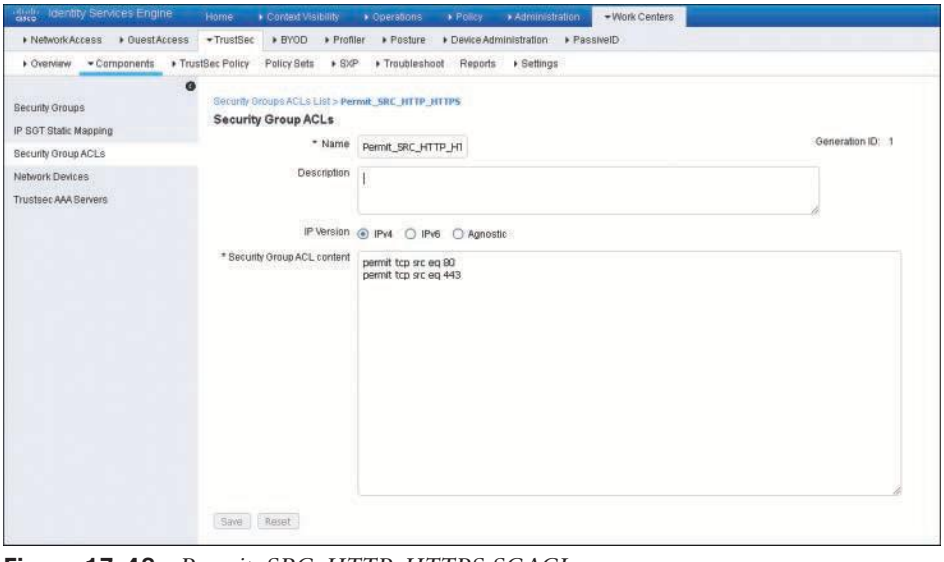


Figure 17-48 *Permit_SRC_HTTP_HTTPS SGACL*

NOTE SGACLs on switches do not function as stateful firewalls. If you permit access in one direction, you need to ensure that access is permitted in the opposite direction. In many cases, filtering in one direction is enough, but if you want to filter the exact posts in either direction, you can create an SGACL based on the source ports.

Figure 17-49 shows the SGACLs created in this example. The completed SGACLs are not enforced in any capacity until they are added to the TrustSec policy matrix, as described in the next section. Then they can be incorporated into an active egress policy.

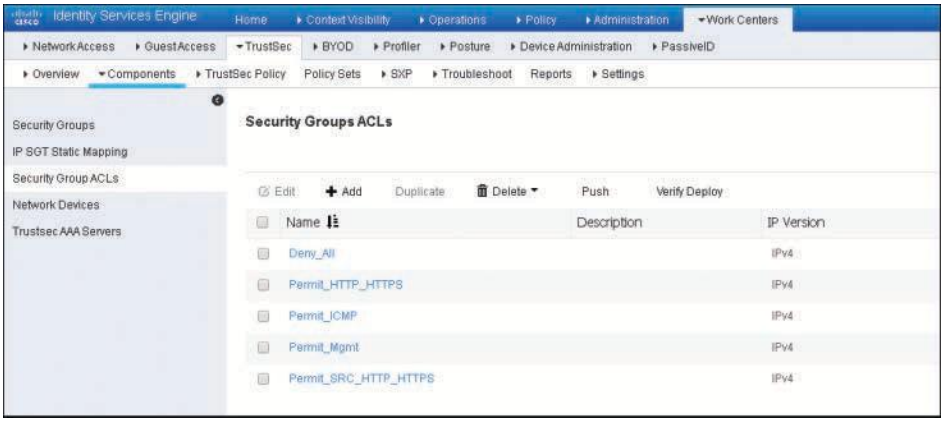


Figure 17-49 *Completed SGACLs*

TrustSec Policy Matrix

With ISE, you can use three views to view and configure the TrustSec egress policy:

- **Source tree:** This is a compact view that lists source SGTs that are mapped to destination SGTs with the corresponding SGACL applied.
- **Destination tree:** This is a compact view that lists the destination SGTs and the source SGTs that they are mapped to, with the corresponding SGACLs applied.
- **Matrix:** This is the most popular view, and it is used for the configuration of the egress policy. The TrustSec policy matrix looks similar to a spreadsheet that maps source SGT to destination SGT in an easy-to-understand view. Because the TrustSec policy matrix is the most commonly used view, we focus on it here.

The main difference between these views is how the security policy is presented.

In earlier versions of ISE, there was only a single TrustSec policy matrix to configure. Starting in Version 2.1, ISE introduced the concept of a staging matrix, which can be used for testing before deployment to the network access devices. ISE Version 2.2 introduced the ability to create multiple matrices so that you can maintain and assign different matrices to different network devices. (For example, one matrix can be deployed to the network access devices in a Hong Kong branch and a different matrix can be deployed to the network access devices in a corporate headquarters in the United States.) ISE Version 2.2 also introduced the

defcon matrix, which makes it possible for a matrix to be sent to all network devices in the event of an emergency or some other specific situation.

There are a few key things to understand with a TrustSec policy matrix. The source axis is the vertical axis, and it lists the source SGT. The destination axis is the horizontal axis, and, of course, it lists the destination SGT. The mapping of a source SGT and a destination SGT is represented by a cell. There are two types of cells in the matrix view: a mapped cell, which is a cell that has one or more SGACLs set in it, and an unmapped cell, which has no set SGACL. If a cell is unmapped, the access allowed is defined in the default TrustSec matrix rule.

Configuring the TrustSec Policy Matrix

Follow these steps in the ISE GUI to configure the TrustSec policy matrix and assign it to a set of network access devices:

Step 1. To support multiple TrustSec matrices and multiple SGACLs, navigate to **Work Centers > TrustSec > Settings > TrustSec Matrix Settings.**

Step 2. Check the **Allow Multiple SGACLs** checkbox to allow the use of multiple SGACLs in a single matrix cell. Optionally, you can also set the color coding of the TrustSec matrix cells here (see Figure 17-50). Click **Save.**

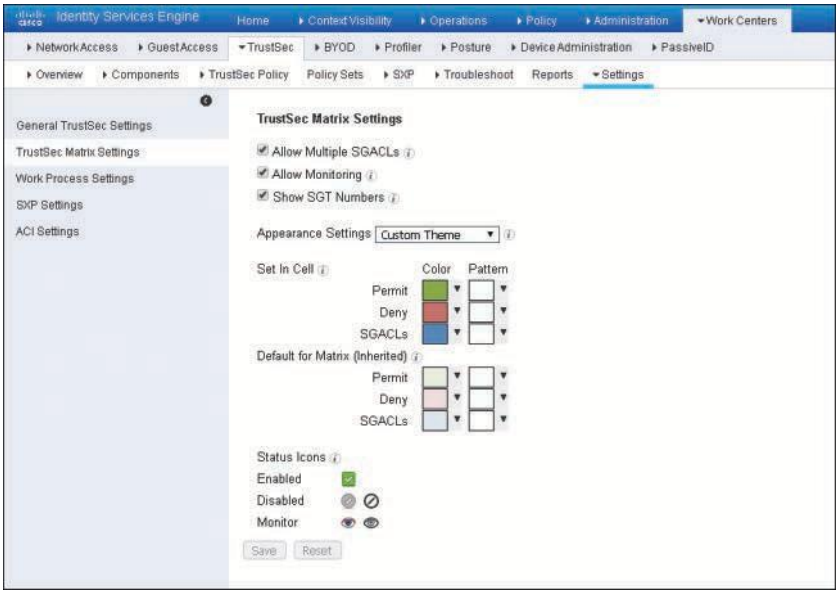


Figure 17-50 *TrustSec Matrix Settings*

Step 3. To configure multiple matrices, navigate to **Work Centers > TrustSec > Settings > Work Process Settings.**

Step 4. Choose the radio button for **Multiple Matrices.** Optionally, you can check the **Use DEFCONS** checkbox to enable defcon matrices (see Figure 17-51). Click **Save.**

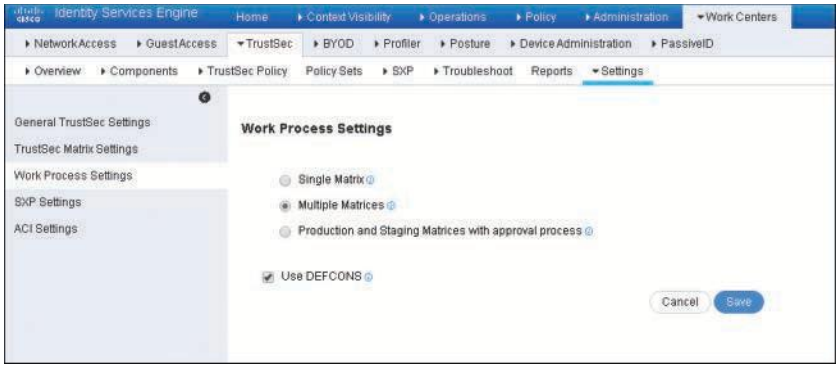


Figure 17-51 *Work Process Settings*

Step 5. To create the TrustSec policy matrix to be used, navigate to **Work Centers > TrustSec > TrustSec Policy > Egress Policy > Matrices List.**

Step 6. Click **Add** to add a new matrix. Name it **CiscoPress_Production** and set Copy Policy From to **Empty** (see Figure 17-52). Click **Submit.**



Figure 17-52 *Matrix Creation*

Step 7. Check the box next to the newly created **CiscoPress_Production** matrix and click the **Assign NADs** button to assign network access devices that this policy will be applied to. Check the box next to the **NADs**, choose the matrix to assign to these devices, and click the **Assign** button (see Figure 17-53).

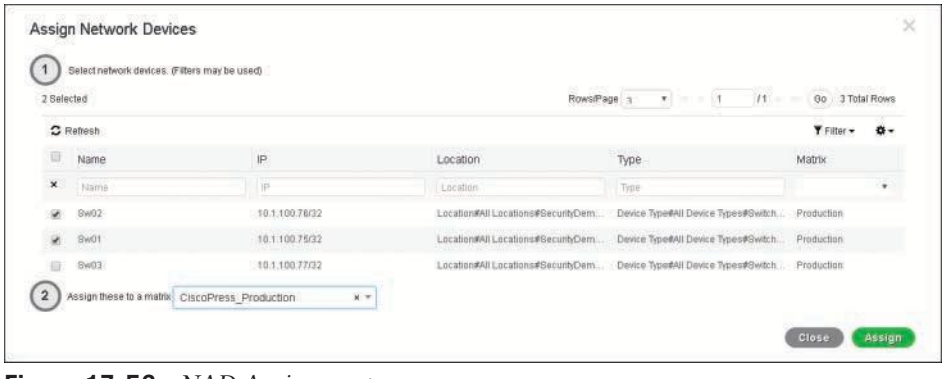


Figure 17-53 *NAD Assignment*

Step 8. (Optional) Click the **Close & Send** button to send the policy to the NADs.

Step 9. Navigate to **Work Centers > TrustSec > TrustSec Policy > Egress Policy > Matrix** to start the configuration of the TrustSec policy matrix. On the top of the TrustSec policy matrix, from the dropdown that allows you to select the matrix you will be editing, choose the **CiscoPress_Production** matrix.

Step 10. On the TrustSec policy matrix, highlight a cell you want to edit and then click the pencil that appears to edit the individual cell. Figure 17-54 displays the edit screen.



Figure 17-54 *Edit Permissions*

The Edit Permissions window pops up when you click the pencil icon. Here you can add the SGACLs that you would like to configure as well as a final catchall rule. If you do not configure a catchall and your SGACLs do not end with a final permit or deny all statement, the default rule of the TrustSec policy determines the outcome. By default, this rule is to permit IP addresses. It is a good idea to leave this setting as is until you tune your policy.

Step 11. Configure the Employee to Unknown permissions to allow both the **Permit_HTTP_HTTPS** and **Permit_Mgmt** SGACLs. Set the final catchall rule to **Deny IP** (see Figure 17-55). Click **Save.**

Step 12. Edit the Contractor to Unknown permissions by adding the **Permit_HTTP_HTTPS** SGACL and setting the final catchall rule to **Deny IP** (see Figure 17-56). Click **Save.**



Figure 17-55 *Employee to Unknown Permissions*



Figure 17-56 *Contractor to Unknown Permissions*

Step 13. Set the following permissions the same way you changed permissions in steps 11 and 12:

- **Admin to Employees:** Permit_ICMP, Deny IP catchall
- **Admin to Contractors:** Permit_ICMP, Deny IP catchall
- **Employees to Employee:** Deny IP catchall

- **Employees to Contractors:** Deny IP catchall
- **Contractors to Employees:** Deny IP catchall
- **Contractors to Contractors:** Deny IP catchall

With this policy in place, employees' devices and contractors' devices cannot communicate with each other even if they are on the same VLAN. However, they can respond to ICMP requests from an administrator. The contractors can communicate with the Unknown tag—which indicates the Internet—on only TCP ports 80 and 443. However, the devices with the Employee tag can also use SSH and RDP, as well as web traffic toward the Internet. Figure 17-57 shows the final TrustSec policy.

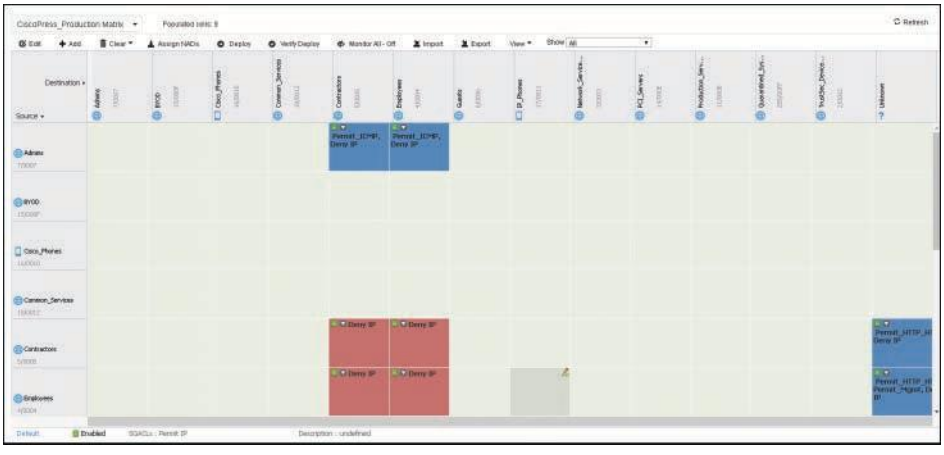


Figure 17-57 *Final TrustSec Policy*

Step 14. At the top of the TrustSec policy matrix, click **Deploy** to push this new policy to the assigned network access devices. If there is an endpoint connected to the VRF instance with the SGT assigned, the applicable policy is downloaded to the network access devices.

Step 15. Verify the policy download from the command line of the switch:

```
C9300# show cts role-based permissions
IPv4 Role-based permissions default:
        Permit IP-00
IPv4 Role-based permissions from group 4:Employees to
group Unknown:
        Permit_HTTP_HTTPS-04
        Permit_Mgmt-01
        Deny IP-00
IPv4 Role-based permissions from group 5:Contractors to
group Unknown:
```

```
        Permit_HTTP_HTTPS-04
        Deny IP-00
IPv4 Role-based permissions from group 4:Employees to
group 4:Employees:
        Deny IP-00
IPv4 Role-based permissions from group 5:Contractors to
group 4:Employees:
        Deny IP-00
IPv4 Role-based permissions from group 7:Admins to group
4:Employees:
        Permit_ICMP-03
        Deny IP-00
IPv4 Role-based permissions from group 4:Employees to
group 5:Contractors:
        Deny IP-00
IPv4 Role-based permissions from group 5:Contractors to
group 5:Contractors:
        Deny IP-00
IPv4 Role-based permissions from group 7:Admins to group
5:Contractors:
        Permit_ICMP-03
        Deny IP-00
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE
C9300# show cts rbacl
CTS RBACL Policy
================
RBACL IP Version Supported: IPv4 & IPv6
  name    = Permit_HTTP_HTTPS-04
  IP protocol version = IPV4
  refcnt = 3
  flag    = 0x41000000
  stale   = FALSE
  RBACL ACEs:
    permit tcp dst eq 80
    permit tcp dst eq 443

  name    = Deny IP-00
  IP protocol version = IPV4, IPV6
```

```
refcnt = 11
flag   = 0xC1000000
stale  = FALSE
RBACL ACEs:
  deny ip

name   = Permit_Mgmt-01
IP protocol version = IPV4
refcnt = 2
flag   = 0x41000000
stale  = FALSE
RBACL ACEs:
  permit tcp dst eq 80
  permit tcp dst eq 443
  permit tcp dst eq 3389
  permit tcp dst eq 22

name   = Permit IP-00
IP protocol version = IPV4, IPV6
refcnt = 2
flag   = 0xC1000000
stale  = FALSE
RBACL ACEs:
  permit ip

name   = Permit_ICMP-03
IP protocol version = IPV4
refcnt = 4
flag   = 0x41000000
stale  = FALSE
RBACL ACEs:
  permit icmp
```


Security Group Firewalls

Some organizations prefer to do traffic enforcement on the switching infrastructure, and some prefer to do it on a firewall, which is a device that is custom-built to do traffic filtering. Cisco has added the ability to enforce traffic on firewalls by implementing the ***security group firewall (SGFW)***. There are two different types of SGFWs: the ASA-based or Firepower-based SGFW and the router-based SGFW. This makes sense because the routers use a zone-based firewall (ZBF), and the ASA does not.

Security Group Firewall on the ASA

Beginning with ASA Version 9.0, the ASA firewall added SGFW functionality. ASDM supports the full configuration, and therefore the ASA has a GUI.

The SGFW in the ASA is a very simple concept. The powerful firewall policy in the firewall has been expanded to include source and destination security groups in the decision. As you can see in Figure 17-58, there is a Security Group column in the Source Criteria and Destination Criteria sections.

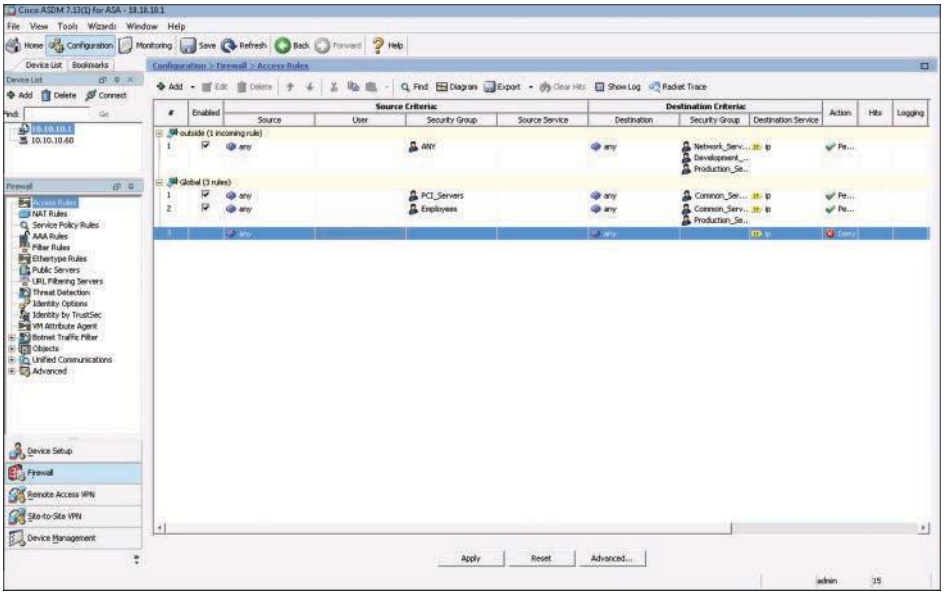


Figure 17-58 *ASDM Firewall Policy*

Security Group Firewall on the Firepower

Starting in Version 6.1, Firepower began to support inline tagging. This provided Firepower Threat Defense with the ability to create access control rules based on the source SGT. These access control rules provided the flexibility to base access control on a tag, much as with the ASA, and also provided the ability to apply next-generation features. SGTs can be used in access control policy rules to specify the following policies should be applied by tag:

- **IPS policy**: Specific IPS policies can be applied or completely bypassed, based on the source tag.
- **File policy**: File inspection or blocking can be applied or completely bypassed, based on the source tag.
- **URL policy**: Web filtering can be applied or completely bypassed, based on the source tag.

pxGrid enables the use of source tag policy for integration between ISE and Firepower. With this integration, the Firepower Management Center is populated with the session data,

including SGTs, the device profile, and user-to-IP address mapping. Access control rules can be created based on these attributes.

Firepower Version 6.5 introduced the ability to learn SXP mapping information from ISE. This new feature made it possible to create access control rules based on source and destination tags. Figure 17-59 shows a Firepower access control policy with two SGT-based rules. The first rule permits guest SGT traffic outbound from the inside interface to the outside interface toward anything with an unknown tag. The second access control rule denies any traffic from an employee SGT toward a device with the PCI_Server SGT.

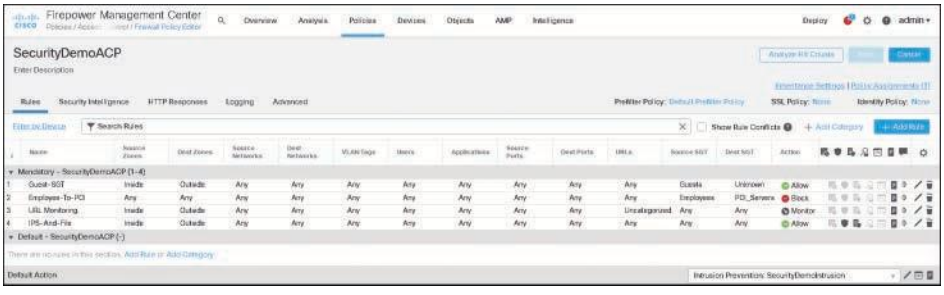


Figure 17-59 *Firepower Access Control Policy*

Security Group Firewall on the ISR and ASR

The ASA is not the only security group firewall on the market. Both the Integrated Services Router (ISR) and the Aggregation Services Router (ASR) have very powerful zone-based firewall capabilities.

The Cisco ISR began support for SGFWs as of Version 15.2(2)T, and the Cisco ASR 1000 added support for SGFWs as of IOS XE Version 3.4. Both routers are capable of running SGT Exchange Protocol (SXP), and the ASR has native tagging capabilities.

Software-Defined Access (SD-Access)

Cisco Software-Defined Access (SD-Access) is an intent-based networking solution for the enterprise that provides consistent management of wired and wireless networks, automates network segmentation, collects contextual insights, and provides open and programmable interfaces. It is based on a campus network fabric that is built on top of a Layer 3 routed underlay. The overlay network is built over that underlay, using VXLAN for encapsulation. LISP is used as the control plane protocol to perform endpoint-to-location mapping throughout the fabric.

The SD-Access fabric provides a simple way to implement hierarchical network segmentation:

- **Macro-segmentation:** With microsegmentation, the fabric logically separates a network topology into smaller virtual networks, using a unique network identifier and separate forwarding tables. This is performed using virtual routing and forwarding (VRF) instances.

- **Micro-segmentation:** Microsegmentation logically separates users and devices further within a VRF by enforcing source-to-destination access control policies. This is done by using SGTs and TrustSec to provide address-agnostic group-based policies.

The complexities of deploying and operationalizing these technologies is abstracted using the DNA Center appliance, which acts as a controller for the fabric. This controller integrates with ISE to push and receive TrustSec policy and environment data. It automates much of the configuration for TrustSec and the network access devices.

The following are some of the benefits of SD-Access:

- Consistent management of wired and wireless network provisioning and policy.
- Automated network segmentation and group-based policy.
- Contextual insights for faster issue resolution and capacity planning with streaming telemetry data to the DNA Center controller.
- Open and programmable interfaces for integration with third-party solutions.
- A Layer 3 routed underlay that eliminates the need for large spanning tree domains and increases network stability.
- A stretched Layer 2 domain, which gives users complete mobility so they can wander around the campus without having their IP addresses change.
- No CAPWAP tunnels, so if a WLC and APs are part of the fabric, they no longer need to tunnel back to the WLC, which removes that bottleneck. Instead, the client traffic is directly handed off to the wired network from the AP.
- SGT tagging configured only on the access-layer network devices of the fabric. Because the overlay uses VXLAN for encapsulation, the SGT is encapsulated, and the packet is treated just like any other IP packet being routed. When the VXLAN is received by the destination network access device, that NAD decapsulates the packet and reads the source SGT tag.

Cisco Software-Defined Access is beyond the scope of the SISE 300-715 exam, so this chapter does not explore it further. It is simply important to understand that TrustSec and ISE are critical components of SD-Access.

MACsec

Long ago, when Wi-Fi was first being introduced into the consumer and corporate space, security concerns arose related to sensitive data being transmitted through the air without any level of confidentiality. The temporary solution was to use Wired Equivalency Protection (WEP). WEP was not very secure (as history has proven), but it provided an extra level of protection designed to bring wireless connections to the same security level available in a wired network.

Wireless networks quickly became even more secure than wired networks. As described many times in this book so far, 802.1X authentication took off like a rocket in wireless networks, and enhancements were made to the encryption and keying mechanisms, such as Wi-Fi Protected Access (WPA/WPA2) using AES encryption. Wireless networks had full encryption mechanisms to provide the confidentiality and integrity of data traversing the Layer 2 hop from the endpoint into the network infrastructure, in addition to the strong identity capabilities of 802.1X.

We needed "wireless equivalency" for wired networks to provide equivalent confidentiality and integrity, but how could we achieve it? Should we use IPsec on every endpoint to every other endpoint, encrypting the entire communication from end to end? If we did that, how would we provide strong levels of quality of service (QoS) since we would not be able to see the content of a packet? We would be simply encrypting both good and bad traffic across the network if we used end-to-end IPsec.

Key Topic

The solution was to layer on the confidentiality and integrity by using IEEE 802.1AE, also known as MACsec. MACsec provides Layer 2 encryption on a LAN between endpoints and the switch as well as between switches. Figure 17-60 illustrates a MACsec configuration from the endpoint through the access layer, distribution layer, and core of a network.

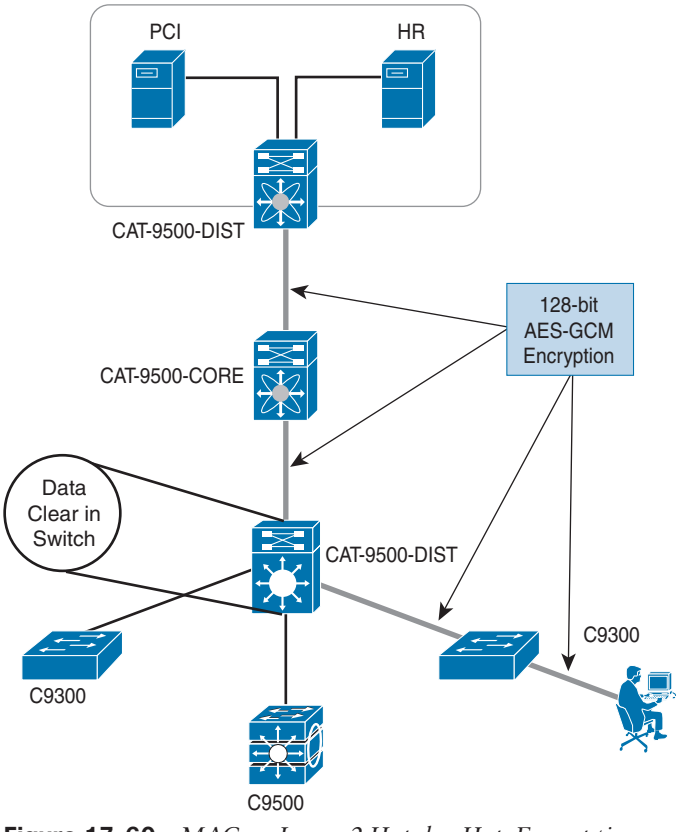


Figure 17-60 *MACsec Layer 2 Hop-by-Hop Encryption*

MACsec provides Layer 2 encryption on the LAN. The encryption also encapsulates and protects the Cisco Meta Data (CMD) field, which carries the SGT, as described earlier in this chapter.

As this book went to press, there were two keying mechanisms available, using 128-bit and 256-bit AES-GCM (Galois/Counter Mode) symmetric encryption that is capable of line-rate encryption and decryption for both Gigabit Ethernet and 10 Gigabit Ethernet interfaces and provides replay attack protection for each and every frame. The keying mechanisms are Security Association Protocol (SAP) and MAC Security Key Agreement (MKA). SAP is a Cisco-proprietary keying protocol used between Cisco switches, and MKA, which is currently used between endpoints and Cisco switches, is becoming the industry standard. SAP is used between switches to negotiate trust. MKA is used between an endpoint and a switch.

Downlink MACsec

Downlink MACsec is the term used to describe the encrypted link between an endpoint and a switch. The encryption between the endpoint and the switch is handled by the MKA keying protocol. Downlink MACsec requires a MACsec-capable switch (such as a Cisco Catalyst 9300) and a MACsec-capable supplicant on the endpoint (such as Cisco AnyConnect Network Access Manager). The encryption on the endpoint can be handled in hardware (if the endpoint possesses the correct hardware) or in software, using the main CPU for the encryption and decryption.

A Cisco switch can force encryption, make it optional, or can force non-encryption. This setting can be configured manually per port (which is not very common) or dynamically as an authorization result from ISE (which is much more common). If ISE returns an encryption policy with the authorization result, the policy issued by ISE overrides anything set using the switch CLI.

Figure 17-61 shows the MACsec policy within an authorization profile on ISE. Notice at the bottom of the figure that the switch sends the attribute cisco-av-pair=subscriber:linksec-policy, followed by the policy itself. As you can see, there are three options in the MACsec Policy dropdown:

- **must-not-secure:** Disables encryption
- **must-secure:** Enables/Forces encryption
- **should-secure:** Prefers encryption, but it is still seen as optional

Example 17-6 shows these options on the switch CLI, and Table 17-3 shows the resulting policy, based on the supplicant policy and switch policy.

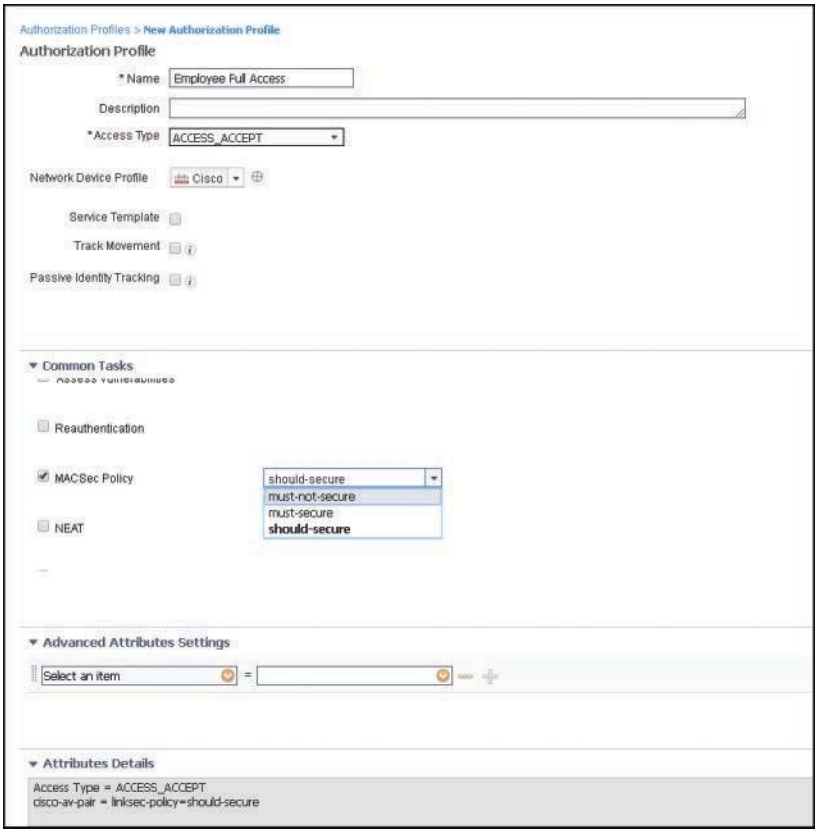


Figure 17-61 *Authorization Profile with MACsec Policy*

Example 17-6 *MACsec Policy at the Switch CLI*

```
CAT9300(config-if)# authentication linksec policy ?
  must-not-secure  Never secure sessions
  must-secure      Always secure sessions
  should-secure    OPTIONALLY secure sessions
```

Table 17-3 Resulting MACsec Policies

Supplicant Policy	Switch Policy	Resulting Policy
Client Supplicant Not Capable of MACsec		
Not MACsec capable	Not MACsec capable	Not secured
Not MACsec capable	must-not-secure	Not secured
Not MACsec capable	should-secure	Not secured
Not MACsec capable	must-secure	Blocked or fallback

Supplicant Policy	Switch Policy	Resulting Policy
Client Supplicant Configured as must-not-secure		
must-not-secure	Not MACsec capable	Not secured
must-not-secure	must-not-secure	Not secured
must-not-secure	should-secure	Not secured
must-not-secure	must-secure	Blocked or fallback
Client Supplicant Configured as should-secure		
should-secure	Not MACsec capable	Not secured
should-secure	must-not-secure	Not secured
should-secure	should-secure	Secured
should-secure	must-secure	Secured
Client Supplicant Configured as must-secure		
must-secure	Not MACsec capable	Blocked
must-secure	must-not-secure	Blocked
must-secure	should-secure	Secured
must-secure	must-secure	Secured

If the authentication server does not return the appropriate attribute/value pair to set the policy, the switch uses the configured policy on the port. If no policy is specified in the switch configuration, the switch reverts to the default policy, which is should-secure.

Switch Configuration Modes

A few configurations on a switch interface have implications for a MACsec deployment, such as the authentication host mode. The host mode plays an important role because it determines the number of endpoints that can be connected to a single switch interface. The following modes are available:

- **Single-Host mode:** MACsec is fully supported in Single-Host mode. In Single-Host mode, only a single MAC address or IP address can be authenticated and secured with MACsec. If a different MAC address is detected on the port after an endpoint has authenticated, a security violation is triggered on the port.
- **Multidomain Authentication (MDA) mode:** With this mode, a single endpoint may be on the data domain, and another endpoint may be on the voice domain. MACsec is fully supported in MDA mode. If both endpoints are MACsec capable, each is secured by its own independent MACsec session. If only one endpoint is MACsec capable, that endpoint can be secured, and the other endpoint sends traffic in plaintext.
- **Multiauthentication mode:** With this mode, a virtually unlimited number of endpoints can be authenticated to a single switch port. MACsec is not supported in this mode.
- **Multi-Hosts mode:** While MACsec use with this mode may technically be possible, it is not recommended. With Multi-Hosts mode, the first endpoint on the port authenticates, and then any additional endpoints are permitted onto the network via the first authorization. So, MACsec works with the first connected host, but no other endpoint's traffic can actually pass since it would not be encrypted traffic.

Example 17-7 shows a switch interface configuration for MACsec-enabled endpoints. The example uses the default MACsec policy, should-secure, and therefore the default setting is not displayed.

Example 17-7 *Switch Interface Configuration for MACsec*

```
interface X
 switchport access vlan 10
 switchport mode access
 switchport voice vlan 99
 ip access-group ACL-ALLOW in
 authentication event fail action next-method
 authentication event server dead action authorize vlan 2274
 authentication event server alive action reinitialize
 authentication event linksec fail action next-method
 authentication host-mode multi-domain
 authentication open
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 macsec
 mka default-policy
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 10
 spanning-tree portfast
end
```


ISE Configuration

Downlink MACsec is configured as an attribute in the authorization profile (as a result of an authorization). To add this result to an authorization profile, follow these steps:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles.**

Step 2. Edit an existing authorization profile.

Step 3. Under Common Tasks, check the box next to **MACsec Policy**.

Step 4. Select **must-secure**, **should-secure**, or **must-not-secure** from the dropdown.

Step 5. Click **Save** to save the change.

Uplink MACsec

Uplink MACsec is the term used to describe encryption of the link between switches with IEEE 802.1AE. At the time this book was written, the switch-to-switch encryption used Cisco's proprietary Security Association Protocol (SAP) instead of MKA, which is used with downlink MACsec toward the endpoint. AES-GCM-128 encryption is used with both uplink and downlink MACsec.

Uplink MACsec can be achieved manually or dynamically. Dynamic MACsec requires 802.1X between the switches, as discussed earlier in this chapter. This section focuses on manual mode.

Manually Configuring Uplink MACsec

Uplink MACsec can be layered on top of the SGTs configured manually earlier in this chapter. It encrypts the inter-switch links. Example 7-8 shows the configuration of an uplink interface from earlier in this chapter.

Example 17-8 *Uplink Configuration*

```
CAT9300# show run int Ten 1/1/1
Building configuration...

Current configuration : 286 bytes
!
interface TenGigabitEthernet1/1/1
 description Cat6K Ten1/5
 no switchport
 ip address 10.1.48.2 255.255.255.252
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 EIGRP
 load-interval 60
 cts manual
  policy static sgt 2 trusted
 no macro auto processing
end
```

With the configuration shown in Example 17-8, the uplink between the Catalyst 9300 and the Catalyst 9500 is set up to use manual keying without any encryption at all, and it applies SGTs to the frames. Now you can layer encryption on top of this to provide confidentiality and integrity of the SGTs and the data. Figure 17-62 shows the relevant infrastructure configuration used for this example.

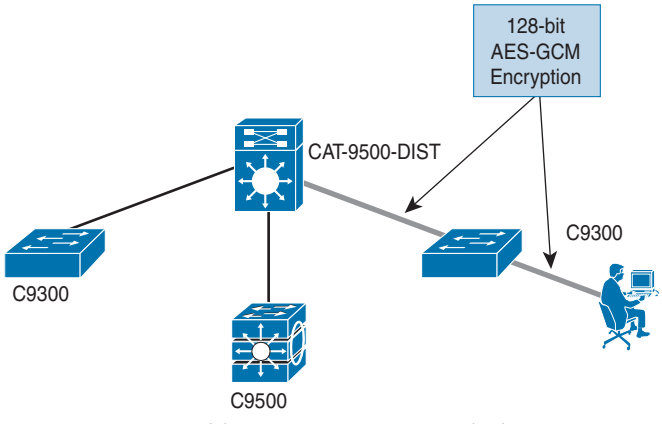


Figure 17-62 *Adding MACsec to the Uplink*

To manually configure uplink MACsec, follow these steps in interface configuration mode at the CLI of the 9300 Series switch:

Step 1. Enter the **cts manual** command.

Step 2. Enable encryption with the **sap pmk** *pairwise-master-key* **mode-list gcm-encrypt** command.

Step 3. Add the **sap pmk** *pairwise-master-key* **mode-list gcm-encrypt** command to the other side of the link. The Pairwise Master Key (PMK) should be a hexadecimal value configured to be the same on both sides of the link. If the PMK is less than 64 hex characters, it will be padded with zeroes (0). Nexus (NX-OS) and Catalyst (IOS/IOS XE) pad differently. NX-OS pads to the right (123000….) while IOS pads to the left (…000123). This can cause a mismatch when configuring MACSec between NX-OS based and IOS based switches. If the padding issue comes up, you can use a PML with 32 characters or issue the **left-zero-padded** keyword when configuring this step in NX-OS. This master key can be compared to a RADIUS shared secret between a NAD and ISE or even the pre-shared key used with IPsec encryption.

Example 17-9 shows these configuration steps, and Example 17-10 shows the final configuration for the uplink port on the Catalyst 9300 switch.

Example 17-9 *Adding Encryption to the Uplink Interface*

```
CAT9300# conf t
Enter configuration commands, one per line. End with CNTL/Z.
CAT9300(config)# int Ten1/1/1
CAT9300(config-if)# cts manual
CAT9300(config-if-cts-manual)# sap pmk 26 mode-list gcm-encrypt
CAT9300(config-if-cts-manual)# end
CAT9300#
```

Example 17-10 *Final Configuration for the Uplink Interface*

```
CAT9300# sho run int ten1/1/1
Building configuration...

Current configuration : 386 bytes
!
interface TenGigabitEthernet1/1/1
 description Cat6K Ten1/5
 no switchport
 ip address 10.1.48.2 255.255.255.252
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 EIGRP
 load-interval 60
```

```
 cts manual
  policy static sgt 2 trusted
  sap pmk 00000000000000000000000000000000000000000000000000000000000000000026 mode-list
gcm-encrypt
 no macro auto processing
end
```

Verifying the Manual Configuration

You can verify that the manual encryption on the uplink was successful by using the **show cts interface** command, as shown in Example 17-11. SAP status is the status of the encryption, and you can see in the example that SAP succeeded, the pairwise cipher is using gcm-encrypt, and replay protection is enabled.

Example 17-11 *Output of the show cts interface Command*

```
CAT9300# show cts interface TenGigabitEthernet 1/1/1
Global Dot1x feature is Enabled
Interface TenGigabitEthernet1/1/1:
    CTS is enabled, mode:    MANUAL
    IFC state:               OPEN
    Authentication Status:   NOT APPLICABLE
        Peer identity:       "unknown"
        Peer's advertised capabilities: "sap"
    Authorization Status:    SUCCEEDED
        Peer SGT:            2
        Peer SGT assignment: Trusted
    SAP Status:              SUCCEEDED
        Version:             2
        Configured pairwise ciphers:
            gcm-encrypt

        Replay protection:      enabled
        Replay protection mode: STRICT

        Selected cipher:        gcm-encrypt

    Propagate SGT:           Enabled
    Cache Info:
        Cache applied to link : NONE

    Statistics:
        authc success:              0
        authc reject:               0
        authc failure:              0
        authc no response:          0
        authc logoff:               0
```

```
        sap success:                2
        sap fail:                   0
        authz success:              5
        authz fail:                 0
        port auth fail:             0

    L3 IPM:    disabled.

CAT9300#
```

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 17-4 lists these key topics and the page number on which each is found.

Key Topic

Table 17-4 Key Topics

Key Topic Element	Description	Page Number
Section	SGA and TrustSec	555
Paragraph	Security Group Tags	556
Section	The TrustSec architecture	557
Section	Classification in a TrustSec architecture	575
Paragraph	SXP in the TrustSec architecture	581
Section	SGT enforcement using SGACLs	597
Paragraph	MACsec as Layer 2 hop-by-hop encryption	615

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

TrustSec, Security Group Tag (SGT), TrustSec domain, Network Device Admission Control (NDAC), seed device, non-seed device, Protected Access Credential (PAC), environment data, policy matrix, SGT Exchange Protocol (SXP), security group access control list (SGACL), security group firewall (SGFW), MACsec

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is the purpose of SXP?
2. What is the goal of TrustSec?
3. What are the three main staples of TrustSec?
4. What are the different ways that an SGT may be assigned?
5. What is required for a switch to perform SGT enforcement?

This page intentionally left blank



CHAPTER 18

Posture Assessment

This chapter covers the following topics:

- Posture assessment
- Posture assessment with ISE
- Configuring posture
- The endpoint experience
- Mobile posture

If you ask three different people in the information security world what the term *posture* means, you will get at least five different answers. The definition of the term really depends on your point of view. Some security professionals think the term refers to the health and security status of the entire organization. However, Cisco introduced this term for endpoints in 2002, with the Network Admission Control (NAC) framework, and it has spread through the network and security world like wildfire. The IETF's Network Endpoint Assessment (NEA) Working Group also used the term *posture*, but other solutions in the industry have called it different things; for example, Microsoft's defunct Network Access Protection (NAP) referred to it as *endpoint health*. You may even see the term *compliance* used, but regardless of the term, what we are discussing here is evaluating the current state of an endpoint and ensuring that it meets the company's security policy requirements before it is admitted to the network. Cisco calls this process *posture assessment*.

Whether you want to prevent an authorized user from using an unauthorized device such as a home laptop or reduce the threats from computer infections and malware, the goal of posture assessment is to validate an endpoint before authorizing its access to the network.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 18-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 18-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Posture Assessment with ISE Configuring Posture	1
Update the Compliance Modules	6
Configure Client Provisioning	10
Configure Posture Policy Elements	2, 4, 8
Configure Posture Policies	9
Other Important Posture Settings	5
Authorization Rules	1
The Endpoint Experience	3
Mobile Posture	7

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you incorrectly guess skews your self-assessment results and might provide you with a false sense of security.

1. What are the three possible posture results that can be used in an authorization policy?
   - **a.** Location, time of day, and operating system
   - **b.** Routes, translations, and permissions
   - **c.** Authentication, authorization, and accounting
   - **d.** Compliant, noncompliant, and unknown
2. What does the file condition for posture do?
   - **a.** Checks for the existence of a file
   - **b.** Checks the date of a file
   - **c.** Checks the version of a file on the client
   - **d.** All of these answers are correct.
3. What is the stealth mode agent?
   - **a.** The stealth mode agent is a special version of AnyConnect that is downloaded from Cisco.com and uploaded to ISE.
   - **b.** The stealth mode agent is used to scrape user-to-IP bindings from Active Directory and send them to ISE for EasyConnect passive identity.
   - **c.** The stealth mode agent is a configuration of AnyConnect that hides the user interface and runs as a service for posture assessment to appear agentless.
   - **d.** The stealth mode agent is loaded into memory from the Client Provisioning portal and is removed when the browser is closed.

4. Which of the following describes the remediation process?
   a. An endpoint that fails a posture check corrects the failure and becomes compliant.
   b. ISE communicates to the ASA firewall to block known attackers.
   c. ISE confirms the identity of the end user, based on the associated endpoint.
   d. An endpoint is moved from compliant state to unknown.
5. Which of the following best describes a posture lease?
   a. Instead of purchasing the posture assessment outright, the client can make lower monthly payments and return the assessment at the end of the term.
   b. As with a DHCP lease, the endpoint requests a posture when it joins the network and refreshes on VLAN change.
   c. ISE uses the last known posture state and avoids a new posture assessment or reassessment during the lease time.
   d. ISE uses the last known posture state and avoids a new posture assessment or reassessment during the lease time for the AnyConnect client only.
6. Antimalware products are updated often. Their versions change, their names change, and processes and data files for their signatures change. Which of the following statements is the most accurate regarding how ISE can keep up with the changes for posture policies?
   a. The compliance modules in the ISE posture runtime must be kept up to date with the latest from Cisco.com.
   b. The compliance modules on both ISE and the AnyConnect System Scan module must be kept up to date with the latest from Cisco.com.
   c. The compliance modules in the AnyConnect System Scan module must be kept up to date with the latest from Cisco.com.
   d. The compliance modules on both ISE and the AnyConnect System Scan module must be kept up to date with the latest from the IETF Network Endpoint Assessment (NEA) group.
7. Which statement related to posture of mobile devices is most accurate?
   a. You must download the iOS and Android versions of AnyConnect and upload them to ISE for CPP. AnyConnect is then used as the posture agent.
   b. iOS and Android endpoints must install AnyConnect from the respective app stores. AnyConnect is then used as the posture agent.
   c. ISE does not perform the posture assessment of the endpoint; it leverages mobile device management (MDM) integrations to accomplish that task.
   d. Posture assessment of mobile devices is not possible.
8. Which statement best describes what the hardware attributes condition is used for?
   a. This unconfigurable condition is used to collect attributes about the endpoint for use in context visibility.
   b. This condition is used to permit network access to endpoints that meet the specific hardware requirements specified.

c. This unconfigurable condition is used to collect attributes about the endpoint for use in authorization policies.

d. There is no hardware attributes condition.

9. Which statement about posture policy elements is most accurate?

a. The posture policy leverages conditions and remediations to determine which posture rules apply to which users and devices.

b. Requirements and remediations are built separately. They are both components that make up conditions, which are used in the posture policy.

c. Conditions and remediations are built separately. They are both components that make up requirements, which are used in the posture policy.

10. Which of these statements best describes the relationship between the AnyConnect client and ISE?

a. AnyConnect is a lightweight client and requires a headend to provide the configurations and modules to the endpoint. ISE is one of the possible headends.

b. AnyConnect is a lightweight client that requires the Cisco cloud to provide the configurations and modules. The configuration for the posture module points to ISE.

c. AnyConnect and ISE do not have a direct relationship.

d. AnyConnect is a lightweight client that requires a Cisco ASA to provide the configurations and modules. The configuration for the posture module points to ISE.

Foundation Topics

Posture Assessment with ISE

This book and the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam are focused on ISE and, therefore, cover posture assessment for ISE. You might run into organizations that use ISE for posture and some that do not. It would be nice if Cisco had just one answer for posture assessment, but that is just not the reality today. It is possible to do posture assessment with application access control using the Cisco Duo Security solution, and it is also possible to perform endpoint posture assessment with a Cisco ASA, using the AnyConnect HostScan module, which ties in to the ASA's Dynamic Access Policy (DAP) engine.

However, the use of the HostScan method was mostly replaced with use of the AnyConnect System Scan module with ISE for the policy server; in addition, when using Firepower Threat Defense (FTD) VPN headends, ISE is the only option for posture assessment because HostScan and DAP were never ported over.

Cisco made the right call: Posture should be centralized from a policy server, such as ISE, and not distributed to local policy engines within each firewall. This central management ensures that an organization has a consistent posture policy that applies to all network access control methods: wired, wireless, and remote access VPNs (RA-VPNs).

A Bit of a History Lesson

To understand why and how ISE works the way it does and what issues would exist if things were done differently, you must first understand the lessons learned so far. Cisco has been

in the posture game longer than anyone else and deployed exponentially more endpoints worldwide. That kind of success does not come without learning some hard lessons.

The concept of authenticating users or devices before allowing network access was brought to the mainstream with IEEE 802.1X, a standard for port access control that provides for strong authentication with extensible authentication protocol (EAP). Cisco was one of the major forces behind 802.1X but did not stop there. 802.1X was released during what many consider the original heyday of computer viruses.

Cisco extended the strong identity controls of 802.1X with a new concept for endpoint posture, ensuring that any endpoints attempting to join the network were compliant with the company's security policy requirements. Validating that an endpoint has antivirus software installed (and running), personal firewalls enabled, and other important checks added a lot more to the authorization aspects of 802.1X in order to keep a company safe from malware. This solution was called ***Cisco Network Admission Control (NAC)***, and it was designed to use 802.1X, so it could leverage Cisco Secure Access Control Server (ACS), the world's most widely deployed RADIUS server, as the centralized brain. Today, we refer to this solution as the "NAC Framework" solution.

In the NAC Framework era, the industry was trying to overload EAP with more than it was ever designed to do. As discussed all throughout this book, 802.1X defines a method for authentication to a network using EAP, which operates at Layer 2. The NAC Framework tried to send the posture data through that EAP communication, similar to the illustration in Figure 18-1, so that the policy server would have the posture data when making the access control decision.

It turns out that EAP does not like this. Overloading the EAP communication with this much data did not work out well. In fact, the IETF eventually created the new protocols PT-EAP and PT-TLS specifically for posture transport. As challenging as it was to use EAP to carry the posture data, the original concept of using a NAD to send the data to the policy server was absolutely spot on.

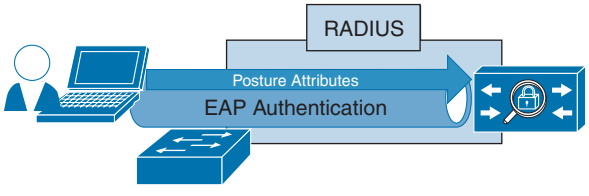


Figure 18-1 *Posture with EAP*

The concept of NAC took off in the industry. Everyone wanted to protect their networks from infected devices, and the industry eventually standardized on the term ***network access control***. Many startup companies were performing network access control using methods other than 802.1X.

The clear leader in the network access control space was a company called Perfigo, which created a Clean Machines inline appliance solution. Cisco acquired Perfigo in 2004 and renamed the solution NAC Appliance. This gave Cisco a solution with strong 802.1X identity control with limited posture capabilities and a separate appliance-based solution for customers looking for stronger posture control with less secure identity control.

Where the NAC Appliance clearly led the industry was in the simplicity of its posture controls. It used SNMP as the control plane to the network, using VLAN switching tricks to steer customer traffic through the inline appliances until the device was permitted to have full network access, and then the endpoint would be switched into the normal data VLAN that did not send all traffic through the inline device.

Figures 18-2 and 18-3 illustrate the NAC Appliance flows, where traffic starts out in the dirty network and is moved to the employees' network after posture assessment is complete. The "dirty" network could be equated to an isolated sandbox where endpoints are kept with limited controlled access until they are determined to be compliant and approved.

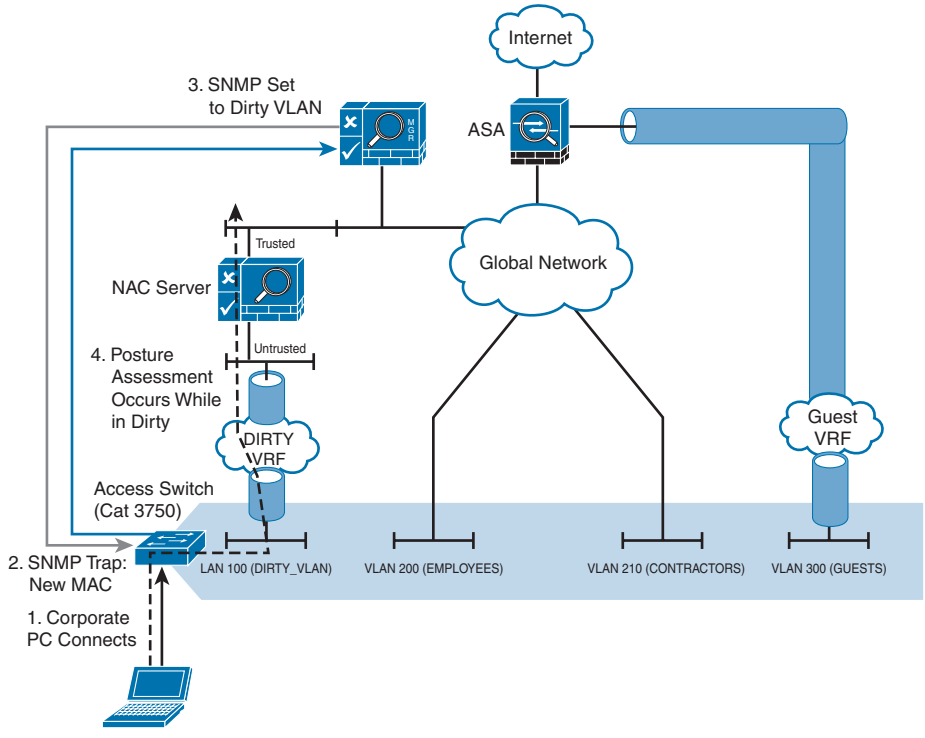


Figure 18-2 *Assessing Posture with Non-802.1X-Based Solutions*

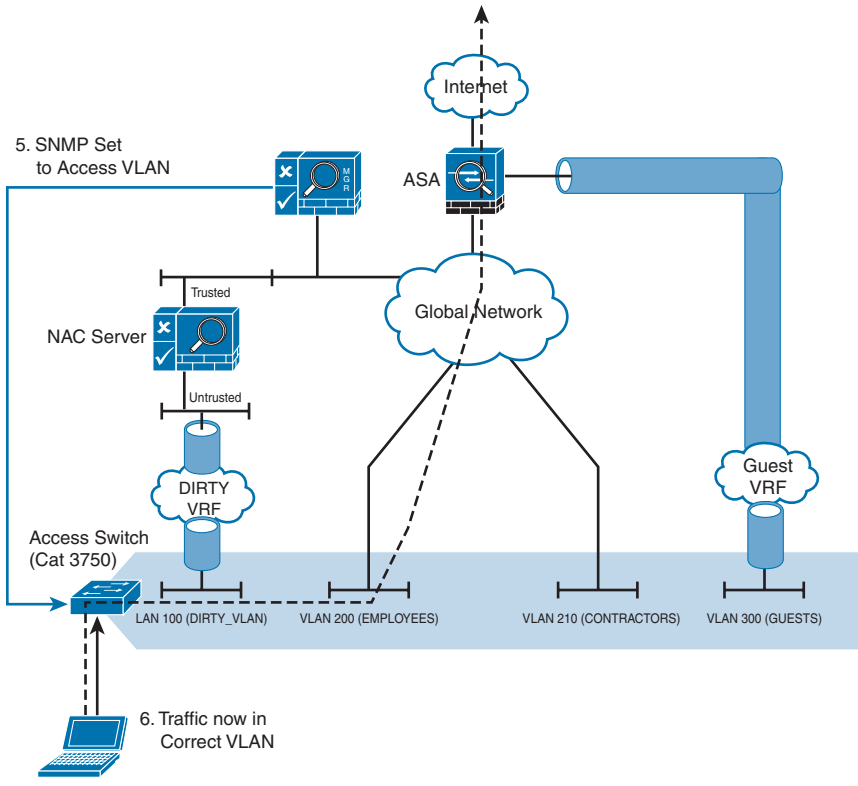


Figure 18-3 *Postposture Admittance with Non-802.1X-Based Solutions*

Figures 18-2 and 18-3 illustrate these six steps:

Step 1. A computer connects to the network.

Step 2. A NAD sends an SNMP trap to the NAC Manager.

Step 3. The NAC Manager sends an SNMP **set** command to put the computer's port into the dirty VLAN.

Step 4. Traffic in the dirty VLAN goes through the inline appliance, and the NAC Agent starts communicating with the NAC server, where posture assessment occurs.

Step 5. When posture assessment is finished, the NAC Manager sends an SNMP **set** command to the NAD, setting the VLAN to the corporate access VLAN.

Step 6. Traffic is now flowing normally.

This form of control had unique complexities, challenges, and shortcomings—especially when deployed at scale. In Chapter 12, "Web Authentication," we cautioned many times about the complexities of changing VLANs on endpoints that do not have agents. Solutions leveraging non-RADIUS control planes had challenges with the following:

- Race conditions related to VLAN changes and DHCP address assignments
- Dealing with failure scenarios (such as fail-open and fail-closed)

- Agents communicating with the enforcement point(s) after being admitted
- Devices without the agent not being directed to the NAC server to download the agent
- Devices without agents not detecting VLAN changes
- Handling multiple devices per switch port

This mixture of out-of-band control plane management with VLAN swapping and temporary-to-permanent inline traffic control worked but was not ideal, especially for large deployments.

ISE Posture Flows

Most of the services in ISE were designed after years of lessons learned with 802.1X, NAC Framework, and NAC Appliance deployments. Engineers like Darrin Miller, Paul Forbes, Craig Hyps, Aaron Woland, and many more brought to the table the experiences of countless customer deployments and the valuable knowledge of what worked well and what didn't work quite as well.

One of the major technological advancements that has enabled solutions (such as posture) that were previously more challenging is the URL-redirection and "sessionization" capability in the network infrastructure. This is a common theme that has been raised numerous times throughout this book with other technologies, such as Centralized Web Authentication (CWA) and profiling.

The enhancements that were made transformed a number of disparate technologies into a real system. The network devices and the policy server share a common session identifier and leverage that session identifier to bridge the gap between a Layer 2 communication (EAP) and a Layer 7 communication of the posture data over a TLS session.

With ISE, the posture data is sent from the endpoint to the active server through a TLS encrypted session, not over EAP. This means that the device must be on the network before posture can be assessed. This is accomplished by providing temporary access for the assessment and any required remediation.

The endpoint needs an agent to perform the compliance interrogation. That is where Cisco's AnyConnect comes in. It is a modular agent, and one of its modules, System Scan, is the posture agent for ISE. Figure 18-4 shows an AnyConnect user interface and highlights the System Scan module.

AnyConnect is on the client side of the posture communication, and ISE's posture subsystem that runs on the Policy Services nodes (PSNs) is on the server side. This subsystem contains a policy engine and a compliance module. The compliance module is the library of supported software, conditions, and operating systems. This module exists both in the ISE posture subsystem and the AnyConnect System Scan module, and it is, in essence, the software that instructs them both how to perform the posture assessment.

Figure 18-5 illustrates the steps involved for access control with ISE, including endpoint posture assessment.

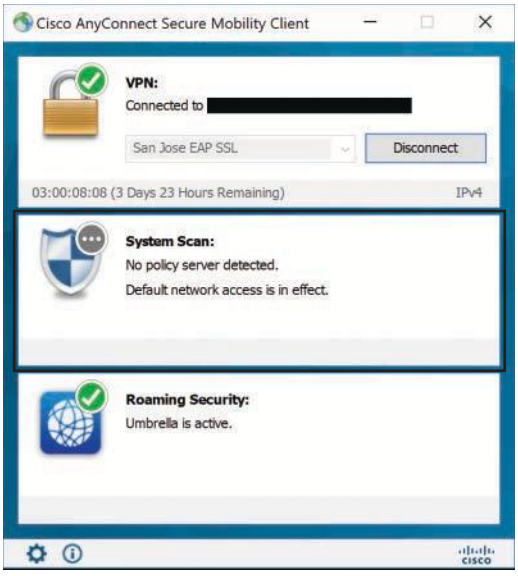


Figure 18-4 *AnyConnect Secure Mobility Client*

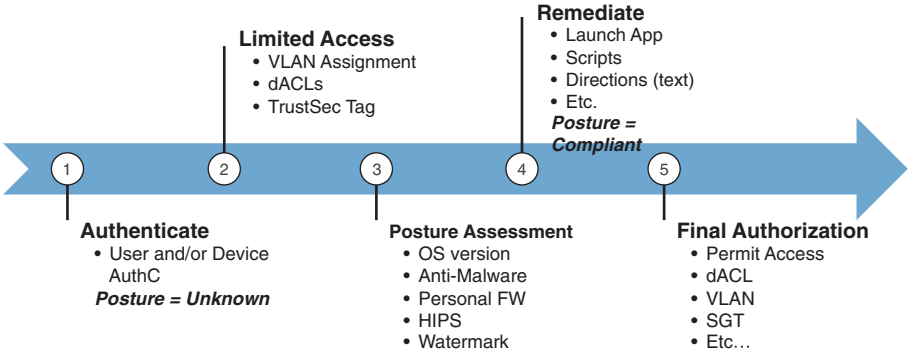


Figure 18-5 *Posture Assessment Steps with ISE*

The following steps for authentication correspond with the numerals in Figure 18-5:

1. A device or user is authenticated.
2. Because the posture is unknown, ISE responds with an authorization that limits the endpoint's access along with URL redirection to the active PSN.
3. Posture discovery occurs, with the System Scan module communicating to the PSN; compliance checking ensues between AnyConnect and the PSN.
4. Any items requiring correction are remediated. Remediation could be automatic, where an application update occurs, or it could be as simple as text being displayed to the end user with instructions.
5. The final authorization result is sent to the NAD, permitting the posture-compliant endpoint into the network.

The URL redirection is important. By leveraging this capability, the NAD can ensure that the authentication traffic and the posture traffic are both sent to the active ISE PSN, along with the identified session ID, to bind the authentication with the data being communicated between the AnyConnect module and the ISE posture subsystem. Figure 18-6 illustrates the posture flows between the agent and ISE.

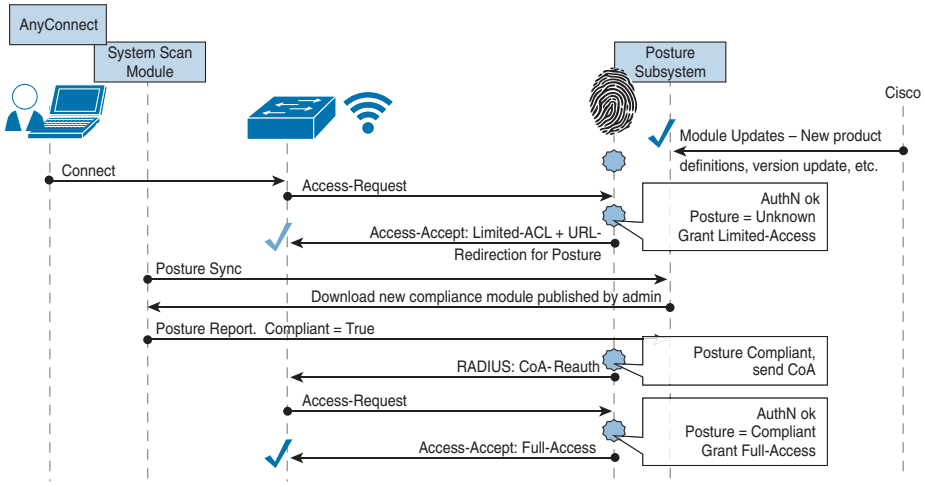


Figure 18-6 *Posture Flows Between ISE and AnyConnect*

Figure 18-7 illustrates what would happen in a system where 802.1X is used, and the posture data does not automatically get sent to the same RADIUS server that is handling the authentication of the endpoint. A race condition ensues, and the posture data must get replicated between all the nodes before the access control decision is made. It is truly inefficient and causes issues at scale.

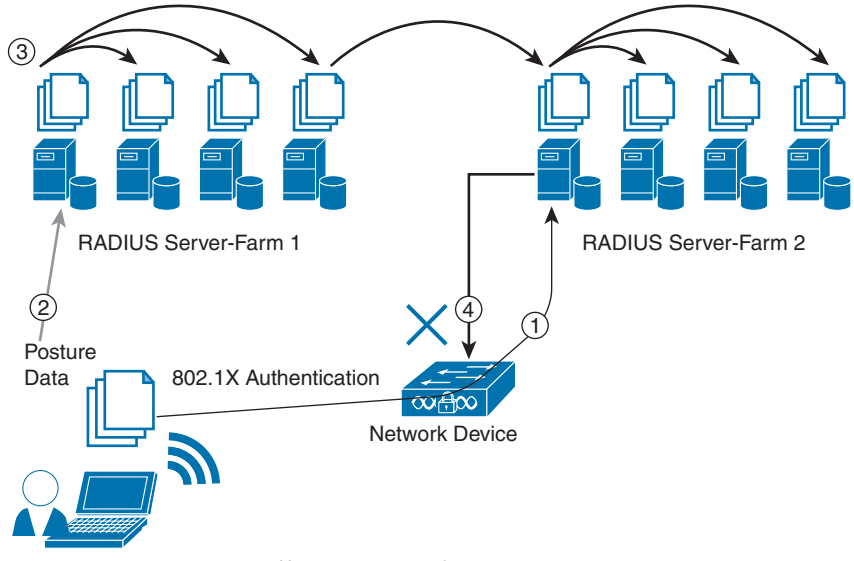


Figure 18-7 *Posture Inefficiency at Scale Without Sessions*

The following steps correspond with the numerals in Figure 18-7 to illustrate that inefficiency:

1. 802.1X authentication occurs with Server 1 in RADIUS Server Farm 2.
2. The posture data is sent to Server 1 in RADIUS Server Farm 1.
3. The posture data must be replicated to all nodes to ensure that the node handling the authorization receives the data before the access control decision is made.
4. Server 1 in RADIUS Server Farm 2 sends the Access-Reject to the network because it did not receive the posture attributes in a timely manner.

The scenario depicted in Figure 18-7 is not a fictional one. It is a reality that was dealt with when designing the ISE system. The team did not want to force ISE to handle the replication timing issues, nor deal with concerns over how busy the PSNs are and affect any replication queues. There are too many unknown variables to leave it to chance in a system designed to support up to the largest of large enterprises.

With the sessionization system complemented by the URL redirection, ISE is able to ensure that posture data is sent directly to the PSN that is handling the 802.1X authentication and the authorization. Figure 18-8 illustrates this gained efficiency.

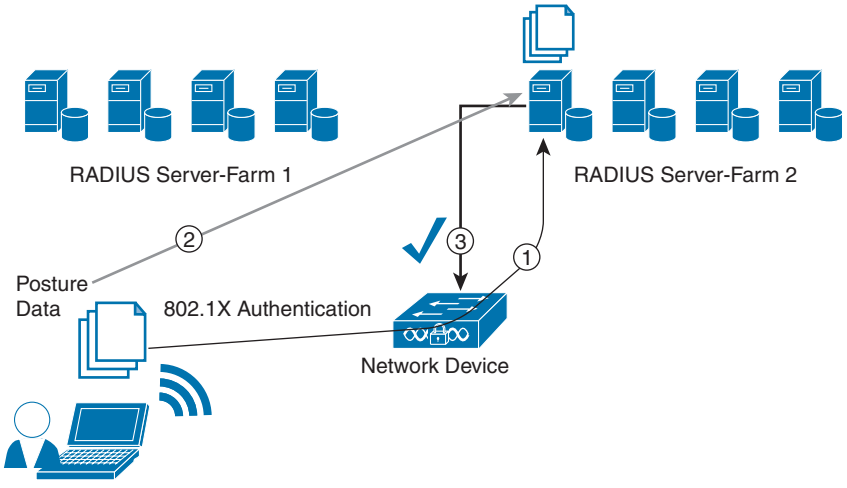


Figure 18-8 *Posture Efficiency at Scale with Sessions*

The following steps correspond with the numerals in Figure 18-8 and illustrate the flows:

1. 802.1X authentication occurs with Server 1 in RADIUS Server Farm 2.
2. The posture data is sent to Server 1 in RADIUS Server Farm 2.
3. The posture data is at the authorizing PSN for the correct authorization to be sent to the NAD.

Configuring Posture

In this section, we move forward with configuring posture assessment and explain it as we progress through the work center.

Like the other work centers in the ISE user interface, the Posture work center has an overview page that lists what you need to accomplish to be successful. Navigate to Work Centers > Posture, and you see the default overview page, shown in Figure 18-9.

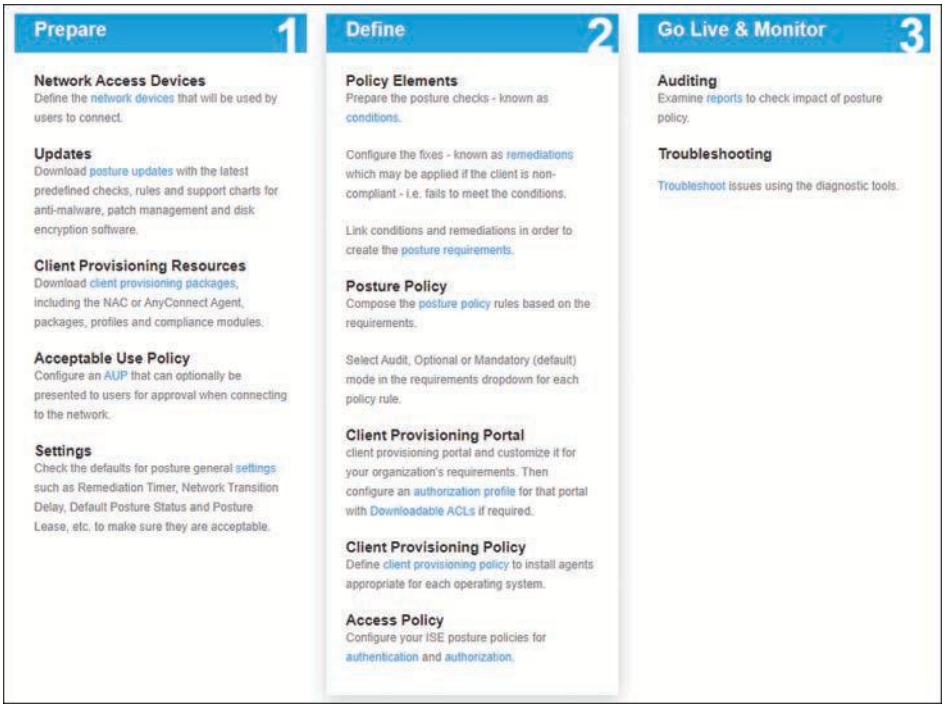


Figure 18-9 *Posture Work Center Overview Page*

The suggested steps on this overview page will help you complete the required tasks to get up and running successfully. In the following sections, you will work with the network access devices you configured and have been running since Chapter 11, "Implement Wired and Wireless Authentication." Therefore, you have already completed the first preparation step shown on the overview page.

Update the Compliance Modules

The next step in the overview page (refer to Figure 18-9) is to download posture updates, which contain the latest definition files for what endpoint protection software suites exist, what their latest versions are, the latest data on the Microsoft "super Tuesday" hotfixes, and a tremendous number of preconfigured objects from Cisco. In other words, you need to download the latest compliance modules, as illustrated in the posture flows in Figure 18-6.

Key Topic

When using posture assessment, it is critical to keep the compliance modules as up to date as possible. Many TAC cases are opened due to an endpoint failing posture checks. The remedy for such cases is simply to update the ISE and AnyConnect compliance modules because the client is running a newer version of some antivirus software, and a new compliance module is required to detect that new version correctly.

The recommended approach to updating ISE with the latest posture packages from Cisco is to download them directly from Cisco.com; this is the default setting. By default, ISE performs the posture update only when you manually click Update Now. You can also configure ISE to periodically check for new updates and download them automatically.

Navigate to Work Centers > Posture > Settings > Software Updates > Posture Updates. Figure 18-10 shows the default Posture Updates screen, and Figure 18-11 shows the screen after a successful update with periodic update checking enabled.

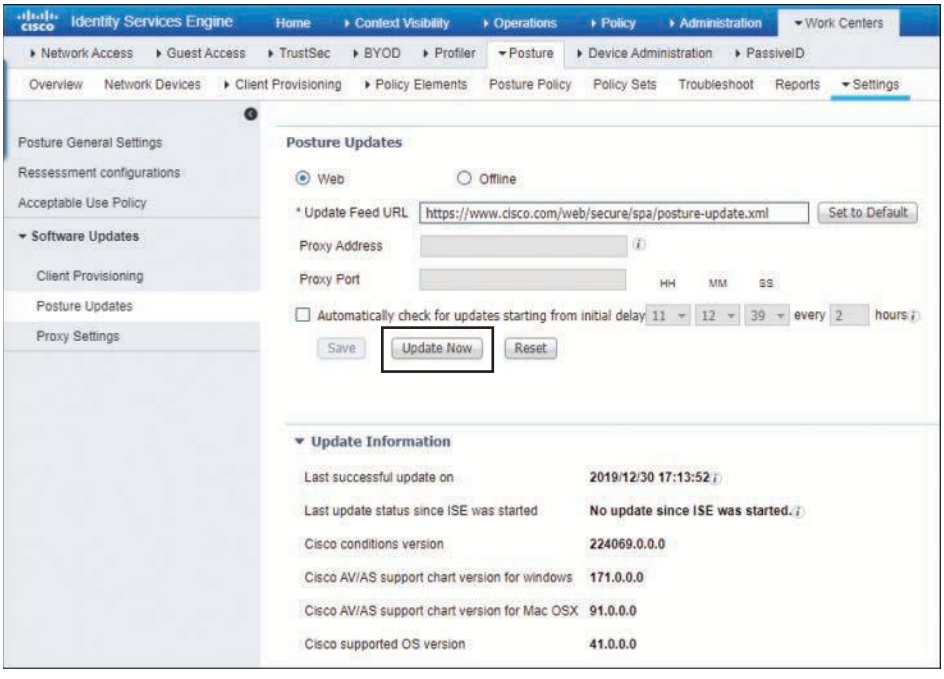


Figure 18-10 *Default Posture Updates Configuration*

If you compare Figure 18-10 with Figure 18-11, you see that Figure 18-11 shows a higher Cisco conditions version, a new version for the Cisco AV/AS support charts for Windows and macOS, and a newer supported OS version. Are you beginning to see the importance of updating the posture module?

Configure Client Provisioning

The next item is to configure the client provisioning policy (CPP), which is not as straightforward as you might hope.

Protect Your Sanity

This section of the configuration can be confusing. The logical order to follow does not necessarily match the user interface. Following the order of operations suggested in this section will go a long way toward preserving your sanity.

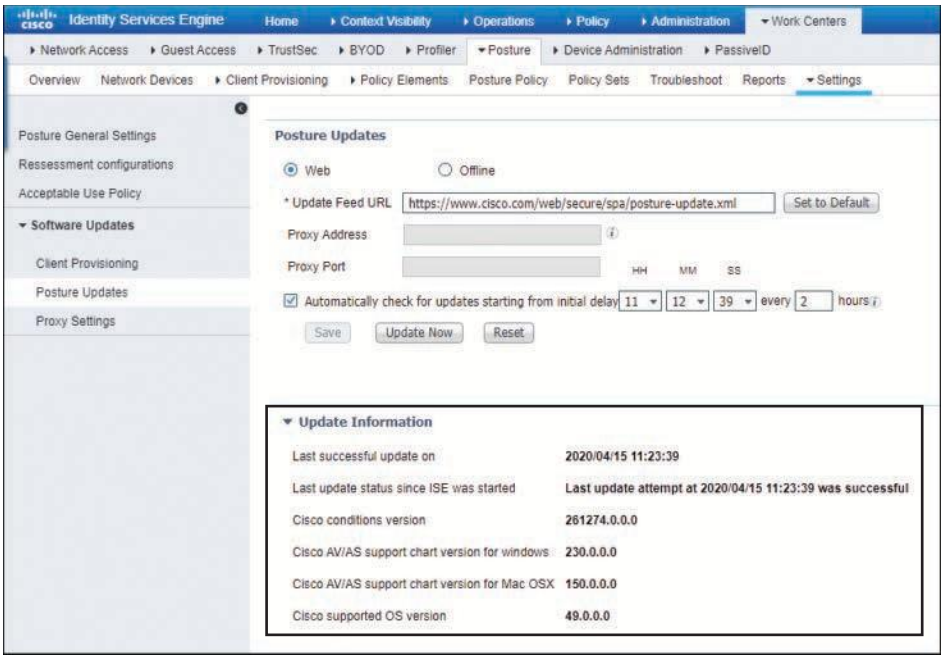


Figure 18-11 *Posture Updated and Periodic Checking Enabled*

Figure 18-12 illustrates the suggested order of operations, and the following list describes the steps:

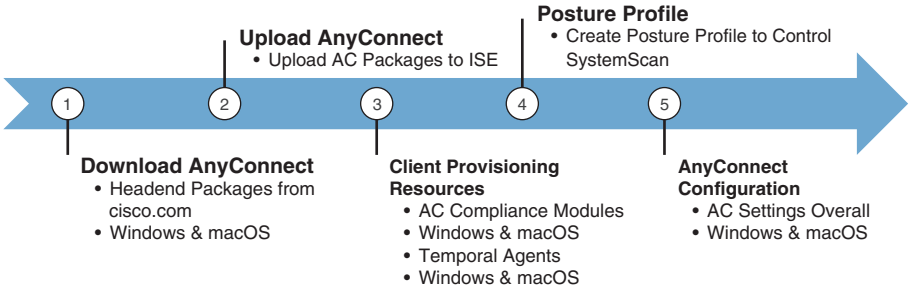


Figure 18-12 *Suggested Order of Operations for Client Provisioning Resources*

Step 1. Download the AnyConnect headend deployment packages from Cisco.com. These are the same .pkg files that you would install on the Cisco ASA for a remote access VPN.

Step 2. Upload the AnyConnect headend deployment packages to ISE.

Step 3. Add the AnyConnect compliance modules in the Client Provisioning Resources screen.

Step 4. Create an AnyConnect posture profile in the Client Provisioning Resources screen, which is used to control all aspects of the AnyConnect System Scan module.

Step 5. Create two AnyConnect configurations in the Client Provisioning Resources screen: one for Windows and the other for macOS.

Download AnyConnect

The AnyConnect Secure Mobility Client is a lightweight modular client that is designed to be controlled by a headend. The headend is responsible for pushing all configuration settings to the client and even for updating the software.

When AnyConnect first hit the market, it was only a VPN client, and the only headend in existence for it was Cisco's ASA. When a user connected to the VPN, if the administrator had published a new policy to the ASA, AnyConnect automatically updated itself with the new policy. The same was true for the AnyConnect software version. If the administrator published a new version, that administrator could also indicate for the client to be upgraded (inline) when the user connected to the ASA.

AnyConnect has since been extended to have modules for Cisco Umbrella, posture assessment with the ASA, posture assessment with ISE, network visibility, a network supplicant, and more.

As the number of modules increased, so did the number of possible headends. With ISE performing posture analysis, it can and does become one of those headends that sends configurations, modules, and software updates to AnyConnect, and it does this through the ISE Client Provisioning portal.

NOTE Beware of the battle of conflicting masters! With so many headends (ASA, FTD, ISE, and Umbrella), it is quite possible for the headends to overwrite one another's policies. It is recommended to keep your VPN headends and ISE in sync at all times.

Before you can configure ISE to be an AnyConnect headend that can deploy and control AnyConnect and its modules, you must first download AnyConnect from Cisco.com and upload it to ISE. Unfortunately, packages cannot be downloaded through the ISE user interface, so you cannot leverage the same process as for other Cisco.com entities.

The files to download from Cisco.com are the ones labeled *headend deployment package*, and these are the same .pkg files that you install on the Cisco ASA for a remote access VPN. Figure 18-13 shows the Cisco.com download page with the headend deployment packages for Windows and macOS highlighted.

Download the files for Windows and macOS and save them to be uploaded into ISE later.

Configure Client Provisioning Resources

The ISE UI page and part of this process should feel rather familiar to you since you configured CPP in Chapter 16, "Bring Your Own Device." From within the Posture work center, navigate to Client Provisioning > Resources, as shown in Figure 18-14.

As you can see in Figure 18-14, some resources come preinstalled with ISE. Most of these are related to native supplicant provisioning, and you worked with them in Chapter 16.

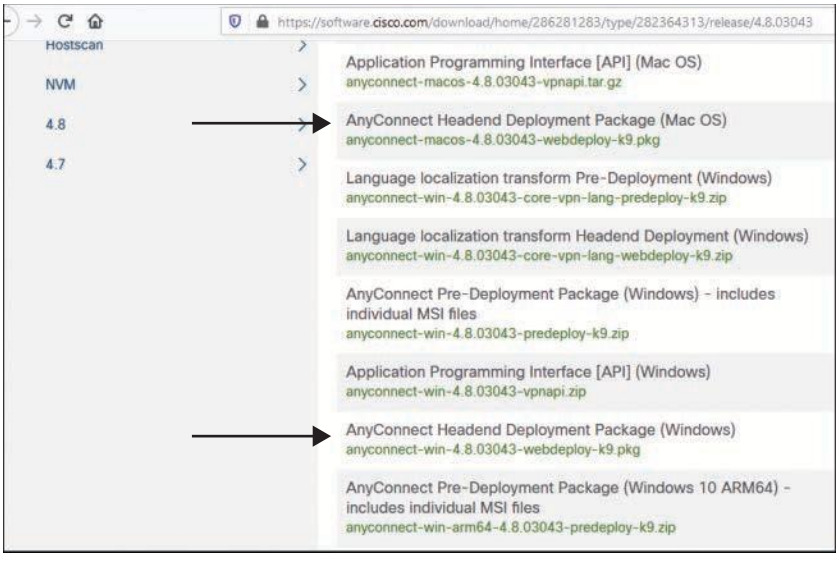


Figure 18-13 *AnyConnect Headend Deployment Packages on Cisco.com*

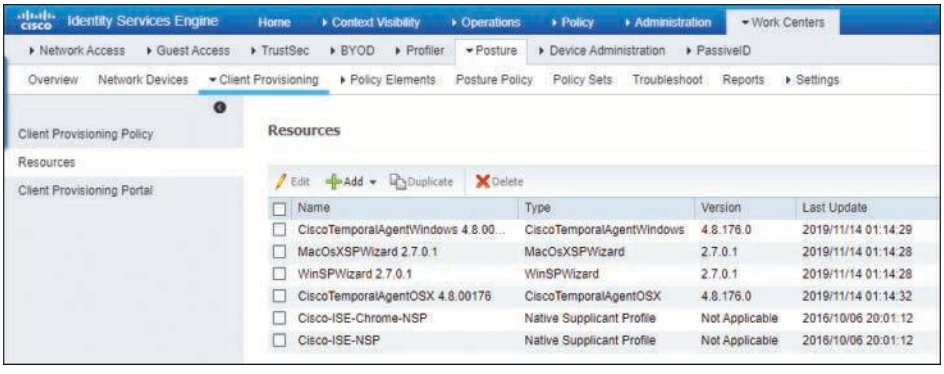


Figure 18-14 *Client Provisioning Resources*

There are two preinstalled resources related to posture: the temporal posture agents for macOS and Windows. (The Windows agent is also known as the "dissolvable agent.") These agents are used for clients that either do not or cannot install the AnyConnect client with its System Scan module, and they provide full posture checking capability, but there is no automatic remediation available with these temporal agents.

Click Add to see the types of resources that can be provisioned from the CPP (see Figure 18-15).

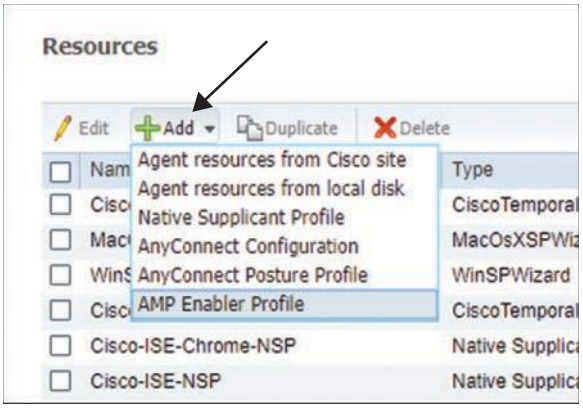


Figure 18-15 *Client Provisioning Resources Available*

As you can see in Figure 18-15, the following types of resources are available:

- **Agent resources from Cisco site:** This option reaches out to the URL specified in the settings and presents the list of available resources. You will use this option in this chapter.
- **Agent resources from local disk:** Some ISE deployments are air gapped—that is, not permitted to communicate to the Internet—and this option is for those customers, and it is the option you will use to upload the AnyConnect packages.
- **Native Supplicant Profile:** This option is used for BYOD onboarding, as discussed in Chapter 16.
- **AnyConnect Configuration:** These configurations are built per operating system (Windows and macOS) and control what AnyConnect modules should be provisioned through the CPP and what profiles should be leveraged per module.
- **AnyConnect Posture Profile:** This is the posture configuration for the System Scan module, where you control all aspects of the posture module, but it is not called a configuration; rather, it is called a *profile* for alignment with the AnyConnect naming convention in Cisco's ASA.
- **AMP Enabler Profile:** This is a configuration for the AMP Enabler module to help deploy AMP for Endpoints through AnyConnect.

The following section shows how to configure the different settings for posture in the Client Provisioning Resources area of the work center.

Upload AnyConnect Headend Deployment Packages to ISE

The second step outlined in Figure 18-12 is to upload the AnyConnect packages to ISE. While it would be terrific if you could just select them here by clicking Add > Agent Resources from Cisco Site, that is not an option. Instead, go to Client Provisioning > Resources and follow these steps:

Step 1. Click **Add > Agent Resources from Local Disk**. As shown in Figure 18-16, the agent resources can be Cisco-provided packages or customer-created packages. A Cisco-provided package is something you download from Cisco.com, such as the AnyConnect software packages. A customer-created package is a profile or a configuration that you created yourself outside the ISE user interface and want to upload to ISE for use with posture assessment.

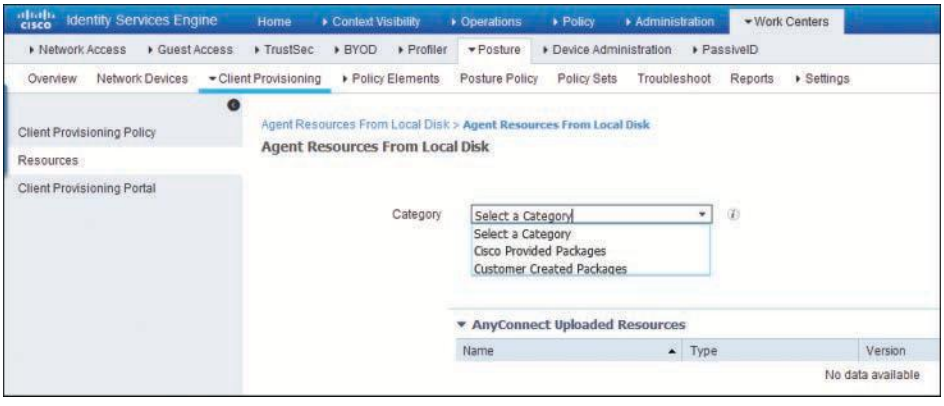


Figure 18-16 *Agent Resources from Local Disk*

Step 2. Select **Cisco Provided Packages** from the Category dropdown.

Step 3. Click **Browse**.

Step 4. Choose one of the AnyConnect packages that you downloaded in the last section. The AnyConnect image is processed, and the information about the package is displayed, as shown in Figure 18-17.

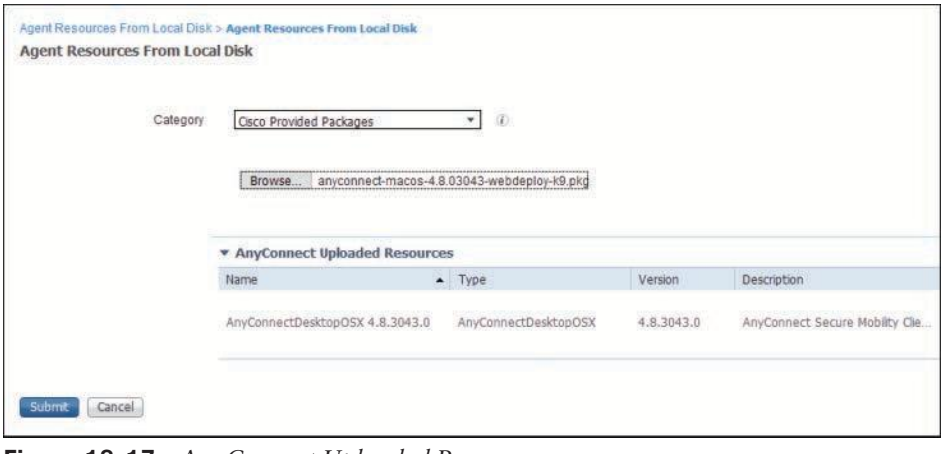


Figure 18-17 *AnyConnect Uploaded Resource*

Step 5. Click **Submit**. Now that AnyConnect is uploaded to ISE, you can have ISE reach out and get the other client resources from Cisco.com.

Step 6. Click **Add > Agent Resources from Cisco Site**. It takes a minute for the window to populate as ISE reaches out to Cisco.com and retrieves a manifest of all the published resources for client provisioning.

Step 7. Select the latest AnyConnect compliance modules for Windows and Mac.

Step 8. Select the latest temporal agents for Windows and Mac.

Step 9. Click **Save**.

Figure 18-18 shows the selected compliance modules and temporal agents.

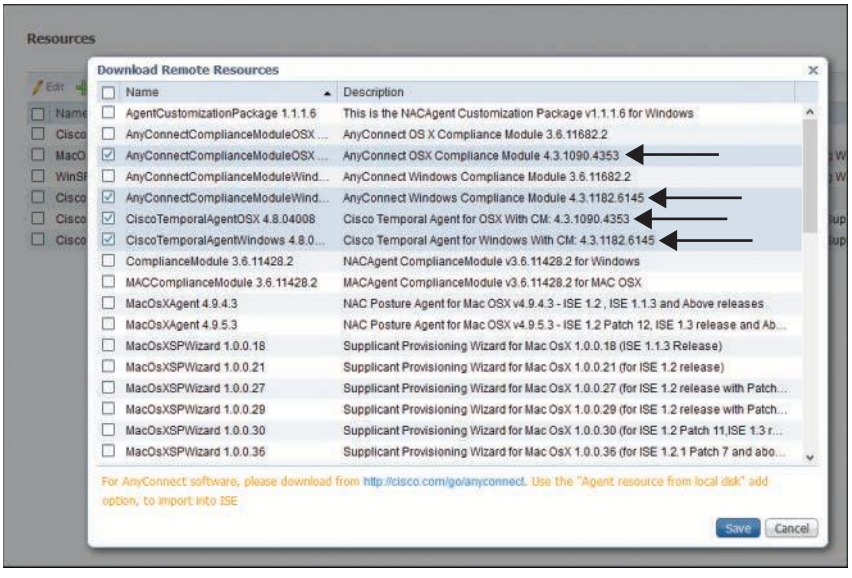


Figure 18-18 *Downloading Remote Resources*

At this point, you have uploaded and downloaded all the required parts. It is now time to build the configuration and profiles required to use those components.

Configure the AnyConnect Posture Profile

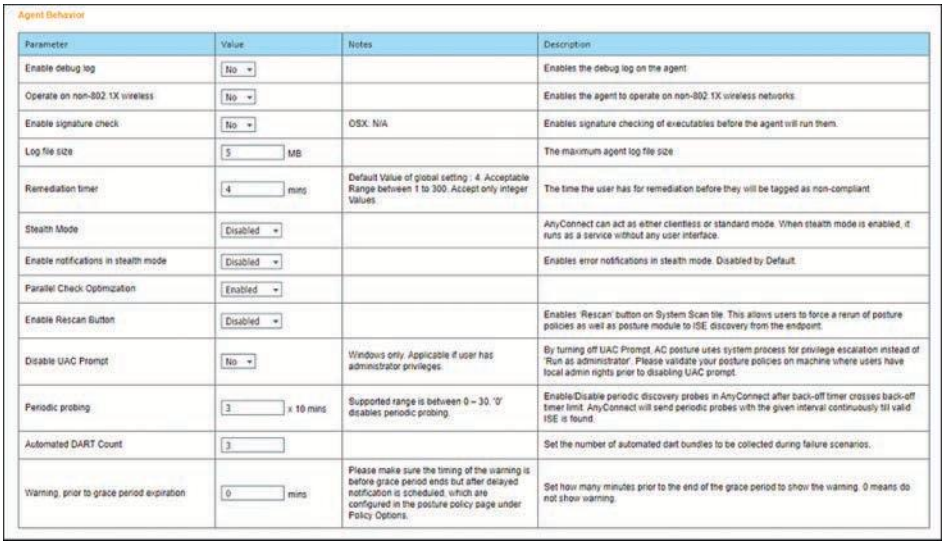
It might seem incorrect to configure an AnyConnect posture profile at this point; however, the posture profile will be used when you create the AnyConnect configuration in the next section.

Click Add > AnyConnect Posture Profile to start building the posture configuration.

The posture profile is (for lack of a better word) gigantic. It configures everything about the System Scan module, which is the evolution of the old NAC Agent from the NAC Appliance days. The configuration is so large that we have broken it up into the following sections.

Agent Behavior

Figure 18-19 shows the Agent Behavior section of the AnyConnect posture profile. Instead of examining every single option, we instead focus on the most important settings:



Agent Behavior

Parameter	Value	Notes	Description
Enable debug log	No		Enables the debug log on the agent
Operate on non-802.1X wireless	No		Enables the agent to operate on non-802.1X wireless networks
Enable signature check	No	OSX: N/A	Enables signature checking of executables before the agent will run them.
Log file size	5 MB		The maximum agent log file size
Remediation timer	4 mins	Default Value of global setting : 4. Acceptable Range between 1 to 300. Accept only integer Values.	The time the user has for remediation before they will be tagged as non-compliant
Stealth Mode	Disabled		AnyConnect can act as either clientless or standard mode. When stealth mode is enabled, it runs as a service without any user interface.
Enable notifications in stealth mode	Disabled		Enables error notifications in stealth mode. Disabled by Default.
Parallel Check Optimization	Enabled		
Enable Rescan Button	Disabled		Enables 'Rescan' button on System Scan tile. This allows users to force a rerun of posture policies as well as posture module to ISE discovery from the endpoint.
Disable UAC Prompt	No	Windows only. Applicable if user has administrator privileges.	By turning off UAC Prompt, AC posture uses system process for privilege escalation instead of 'Run as administrator'. Please validate your posture policies on machine where users have local admin rights prior to disabling UAC prompt.
Periodic probing	3 x 10 mins	Supported range is between 0 – 30. '0' disables periodic probing.	Enable/Disable periodic discovery probes in AnyConnect after back-off timer crosses back-off timer limit. AnyConnect will send periodic probes with the given interval continuously till valid ISE is found.
Automated DART Count	3		Set the number of automated dart bundles to be collected during failure scenarios.
Warning, prior to grace period expiration	0 mins	Please make sure the timing of the warning is before grace period ends but after delayed notification is scheduled, which are configured in the posture policy page under Policy Options.	Set how many minutes prior to the end of the grace period to show the warning. 0 means do not show warning.

Figure 18-19 *AnyConnect Posture Profile: Agent Behavior Section*

- **Operate on non-802.1X wireless networks:** Keep in mind that System Scan is meant to be used with the ISE PSN that is actively authenticating and authorizing the network access session. This usually means 802.1X authentication, but it is possible that ISE posture could be used with open networks or WPA2-Personal networks.
- **Enable signature check:** This setting only applies to Windows. Posture is not only about the checking of conditions; remediation is also a key responsibility. A remediation can include automatic execution of a script binary executable file, and requiring that the file be digitally signed with a valid signature is an important security consideration for some organizations.
- **Remediation timer:** When an endpoint's posture is unknown, the endpoint is put through a posture assessment flow. It takes time to remediate failed posture checks; the default time is 4 minutes before marking the endpoint as noncompliant, but the values can range from 1 to 300 minutes (5 hours).
- **Stealth mode:** This mode was added in direct response to customer feedback requesting it. In many organizations, end users do not have enough rights to remediate their own endpoints; some organizations prefer to keep track of the compliance status and never concern the end user with popups or any notification whatsoever. Stealth mode exists for such organizations.
- **Enable notifications in stealth mode:** This setting permits error messages to be displayed to the end user. In other words, the agent remains in stealth mode except when it comes to errors.
- **Enable Rescan Button:** This setting allows the end user to force a new posture discovery process from the agent, which also triggers the posture policies to be reevaluated.

- **Disable UAC Prompt:** User Access Control (UAC) is a Windows security feature in which most user accounts run with "standard" privileges, and administrator consent is requested for all elevated functions. If you turn off the UAC prompt, System Scan uses the system process instead of Run as Administrator for privilege escalation. (See https://docs.microsoft.com/en-us/windows/security/identity-protection/user-account-control/how-user-account-control-works for more information on UAC.)

IP Address Change

All throughout this book, you have been warned about the complications that result from changing VLANs after an endpoint has an DHCP address. The System Scan module helps deal with those VLAN changes by triggering a DHCP renewal. For 802.1X networks, the supplicant typically handles the renewing of IP addresses, and therefore the agent is not required. However, for non-802.1X networks where VLAN changes will be used, this agent capability is very important.

Figure 18-20 shows the IP Address Change section of the AnyConnect posture profile. These are the most important settings:

IP Address Change

Parameter	Value	Notes	Description
Enable agent IP refresh	Yes	Enables VLAN change detection	Sets the Vlan change detection flag on the server, to transmit the configured dhcp release delay, and the dhcp renew delay values from the server to the client.
VLAN detection interval	0 secs	0 means VLAN detection is disabled	The interval at which the agent will check for a VLAN change
Ping or ARP	Ping	0=Ping, 1=ARP, 2=Ping then ARP	Method for detecting IP address change.
Maximum timeout for ping	1 secs		Ping timeout.
DHCP renew delay	1 secs		
DHCP release delay	4 secs		
Network transition delay	3 secs	The default is empty which means uses the global setting. The default of global setting is 3.	The period for which the agent suspends network monitoring so it can wait for a planned IP change to happen

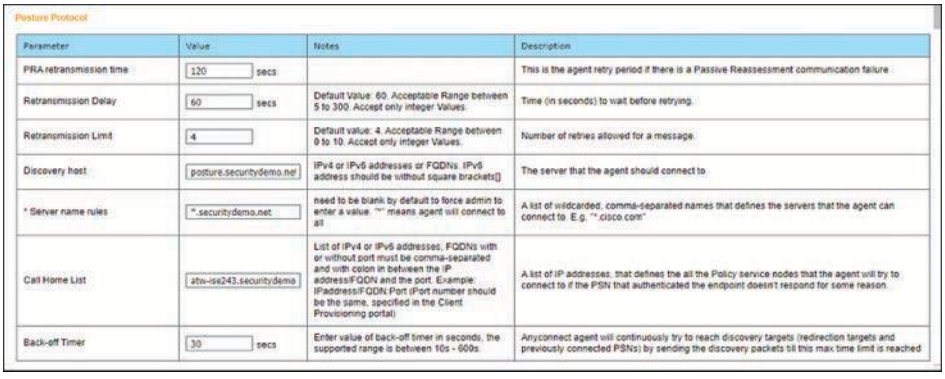
Figure 18-20 *AnyConnect Posture Profile: IP Address Change Section*

- **Enable agent IP refresh:** As it sounds, this setting enables the functionality in the module where VLAN changes are detected and IP addresses are renewed.
- **VLAN detection interval:** This setting enables you to set the number of seconds the module will wait between probing for VLAN changes.
- **Ping or ARP:** This is the actual VLAN change detection method. The agent can ping the default gateway or monitor the ARP cache for the default gateway's entry to timeout or both.

Posture Protocol

The title of this section is a bit misleading as it implies that this section is about the communication protocol itself, but it is not. This section shows how to configure agent behavior related to the agent discovery process and posture reassessment. The URL redirection function usually ensures that posture discovery works smoothly, but posture reassessment occurs when an endpoint has already been granted its full network access and the redirection is no longer in place.

In this section, we look at settings for configuring features to aid in discovery and communication with ISE PSNs. Figure 18-21 highlights the Posture Protocol section of the AnyConnect posture profile. The following settings are the most important ones to understand:



Posture Protocol

Parameter	Value	Notes	Description
PRA retransmission time	120 secs		This is the agent retry period if there is a Passive Reassessment communication failure
Retransmission Delay	60 secs	Default Value: 60. Acceptable Range between 5 to 300. Accept only integer Values	Time (in seconds) to wait before retrying.
Retransmission Limit	4	Default value: 4. Acceptable Range between 0 to 10. Accept only integer Values.	Number of retries allowed for a message.
Discovery host	posture.securitydemo.net	IPv4 or IPv6 addresses or FQDNs. IPv6 address should be without square brackets[]	The server that the agent should connect to
* Server name rules	*.securitydemo.net	need to be blank by default to force admin to enter a value. "*" means agent will connect to all	A list of wildcarded, comma-separated names that defines the servers that the agent can connect to. E.g. "*.cisco.com"
Call Home List	atw-ise243.securitydemo	List of IPv4 or IPv6 addresses, FQDNs with or without port must be comma-separated and with colon in between the IP address/FQDN and the port. Example: IPaddress/FQDN:Port (Port number should be the same, specified in the Client Provisioning portal)	A list of IP addresses, that defines the all the Policy service nodes that the agent will try to connect to if the PSN that authenticated the endpoint doesn't respond for some reason.
Back-off Timer	30 secs	Enter value of back-off timer in seconds, the supported range is between 10s - 600s.	Anyconnect agent will continuously try to reach discovery targets (redirection targets and previously connected PSNs) by sending the discovery packets till this max time limit is reached

Figure 18-21 *AnyConnect Posture Profile: Posture Protocol Section*

- **PRA retransmission time:** This setting determines how quickly the agent should retry when something goes wrong with posture reassessment (PRA) communication.
- **Discovery host:** The discovery host is just an IP address or FQDN that the agent tries to connect to. Originally, this was a way to trigger the NAC Agent to communicate *through* the inline NAC Appliance. The need is very similar with ISE: The traffic being sent to the discovery host can trigger the communication to be redirected to the active PSN.
- **Server name rules:** System Scan could theoretically communicate with any ISE PSN, and this setting enables an administrator to limit the communication to nodes whose names are in this wildcard list (for example, *.securitydemo.net).
- **Call Home List:** This is a list of IP addresses and ports that defines all the PSNs that the agent will try to connect to if the PSN that authenticated the endpoint doesn't respond for some reason. The port must correspond to the port setting in the Client Provisioning portal (for example, atw-ise243.securitydemo.net:8443, atw-ise244.securitydemo.net:8443).

Custom Messages

It is important to provide clear messaging to end users. The final two sections of the configuration allow you to customize the messages that the end user will see when the endpoint is in the noncompliant and grace period states.

Figure 18-22 highlights the two custom message sections of the AnyConnect posture profile.

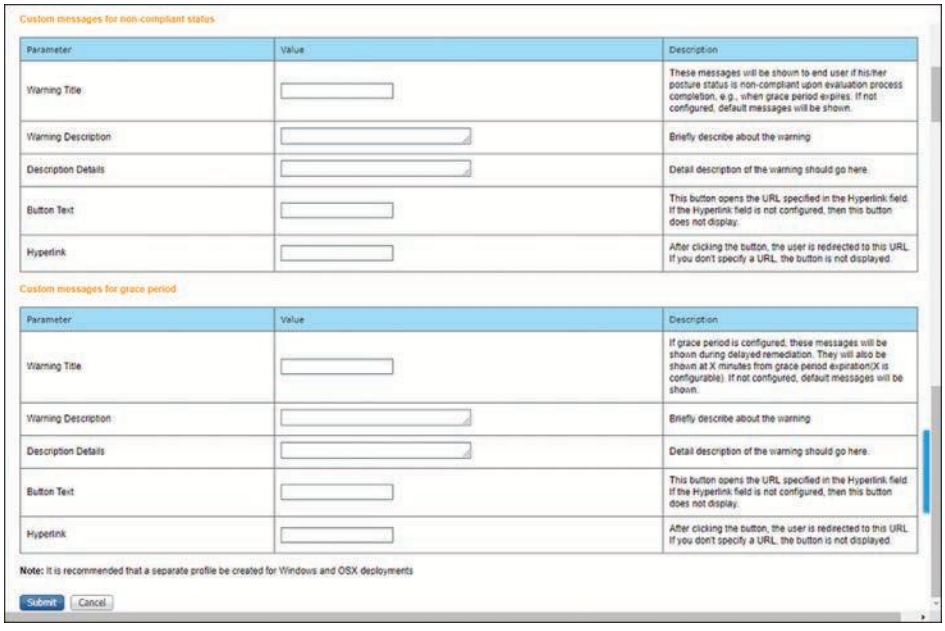


Figure 18-22 *AnyConnect Posture Profile: Custom Messages Sections*

After you have added any customizations to the posture profile, click Submit to save your customized profile. You will use it in the next section as you build the AnyConnect configuration.

Build the AnyConnect Configuration

Key Topic

For a professional who is designing, deploying, or administering ISE and its components, it's a good idea to remember that ISE is a headend for AnyConnect. As such, it has the ability to deploy AnyConnect and its modules to clients. With the ASA, this deployment capability happens inline when a VPN connection is authenticated to the ASA. The deployment functionality with ISE is all handled through the Client Provisioning portal. These are the steps:

Step 1. Click **Add > AnyConnect Configuration**

Step 2. Select the appropriate option from the Windows AnyConnect Package drop-down (see Figure 18-23).

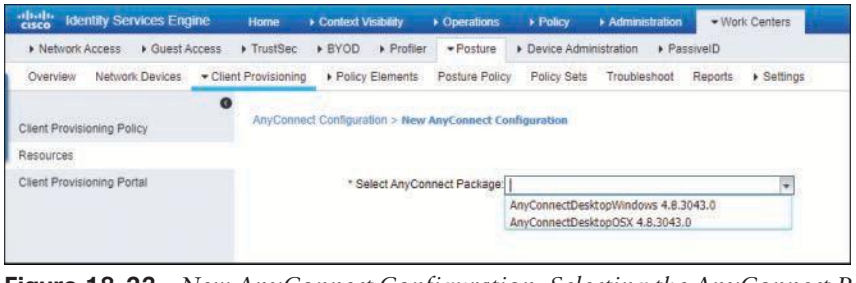


Figure 18-23 *New AnyConnect Configuration: Selecting the AnyConnect Package*

Step 3. Provide a name for this configuration and, optionally, a description.

Step 4. Select the Windows option from the AnyConnect Compliance Module dropdown.

NOTE If you have not downloaded the agent resources already, the AnyConnect Compliance Module dropdown is empty and unusable. This is one of many reasons for the recommended order of operations in Figure 18-12.

Step 5. From the AnyConnect Module Selection section, shown in Figure 18-24, select the modules that should be deployed when clients are provisioned through ISE's CPP.

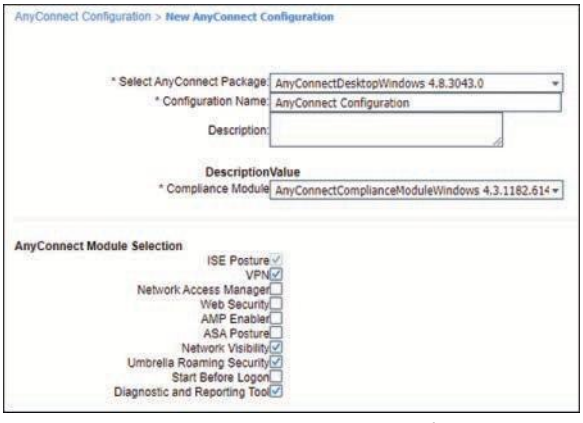


Figure 18-24 *New AnyConnect Configuration: Package and Module Selection Sections*

NOTE Just as the ISE posture module (System Scan) has a profile to configure all aspects of the module, so do all the other modules. You can create the profiles offline and upload them just as you would the AnyConnect packages. The AnyConnect UI can also be customized and localized for different languages. Those bundles can also be created offline and uploaded to ISE.

Step 6. Select the posture profile you created and any additional profiles for each module, if they have been uploaded to ISE already, as shown in Figure 18-25.

Step 7. In the Deferred Updates and Installation Options sections, shown in Figure 18-26, make the appropriate selections. In these sections, you can permit end users to opt out of the AnyConnect software and compliance module upgrades when a newer version has been published to the Client Provisioning portal. Also, as part of the migration from NAC Appliance to ISE deployments, you can configure AnyConnect to uninstall the older NAC Agent. When you are finished making selections, click **Submit**.

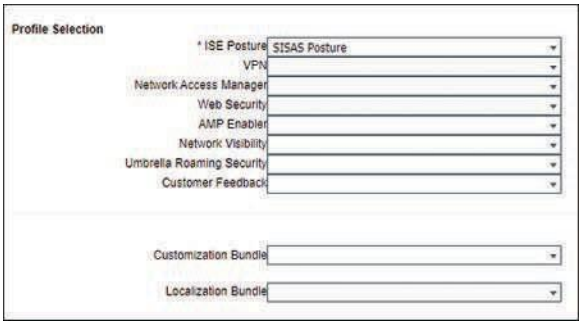


Figure 18-25 *New AnyConnect Configuration: Profile Selection and Customization Sections*

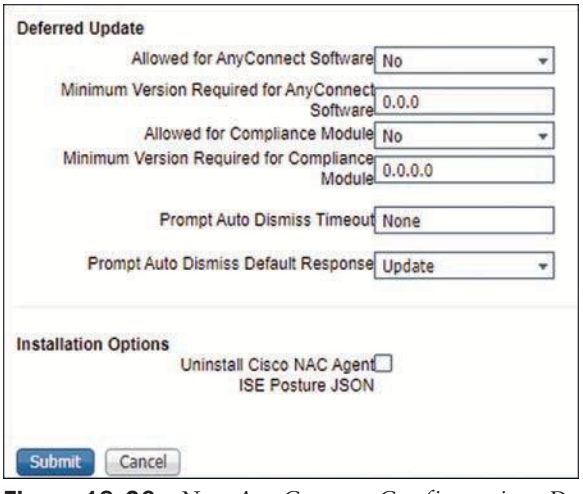


Figure 18-26 *New AnyConnect Configuration: Deferred Update and Installation Options Sections*

Configure the Client Provisioning Portal

The Client Provisioning portal should seem rather familiar to you. All the portal configurations within ISE are all built using exactly the same framework. So, the same customizations and workflows hold true for this portal as for all the portals you worked with in Chapter 13, "Guest Services."

To go to the Client Provisioning portal, select Work Centers > Posture > Client Provisioning > Client Provisioning Portals, as shown in Figure 18-27.

Most of the configuration choices for the Client Provisioning portal are just like the options for the other portals you have seen, although there are some key differences, which are highlighted here. Figure 18-28 shows the Portal Settings section, where you can select the interface and port, as well as the groups that are authorized to the page.

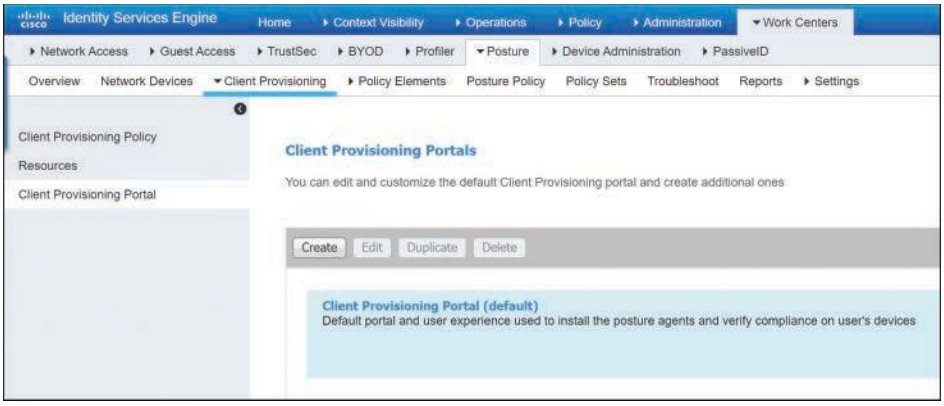


Figure 18-27 *Client Provisioning Portal*

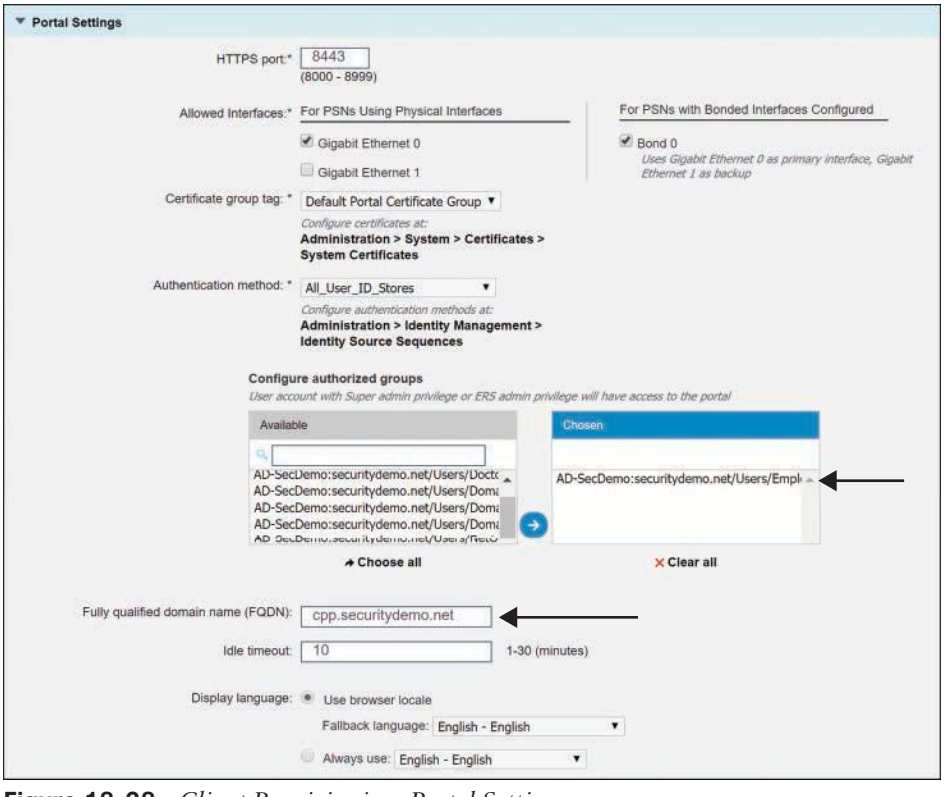


Figure 18-28 *Client Provisioning: Portal Settings*

The Login Page Settings section has an option you have not seen in other configurations to this point: Enable Auto Login (see Figure 18-29). You can select this setting to leverage the ISE session directory to ensure that the user is logged in and determine the user's identity instead of forcing the user to log in to the portal web page separately.

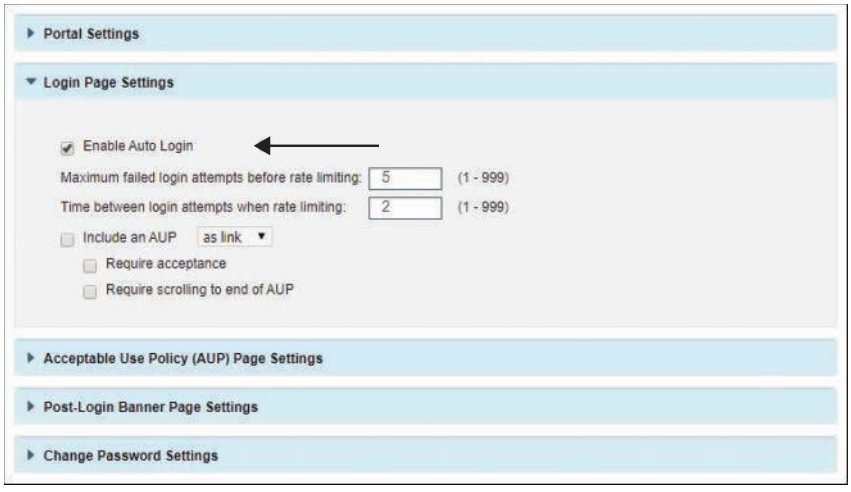


Figure 18-29 *Client Provisioning: Login Page Settings*

Configure the Client Provisioning Policy

The client provisioning policy (CPP) you worked with in Chapter 16 is the same CPP leveraged for posture. This policy has conditions that are matched to determine what software needs to be deployed and leveraged for those endpoints.

Figure 18-30 shows the default client provisioning policy.

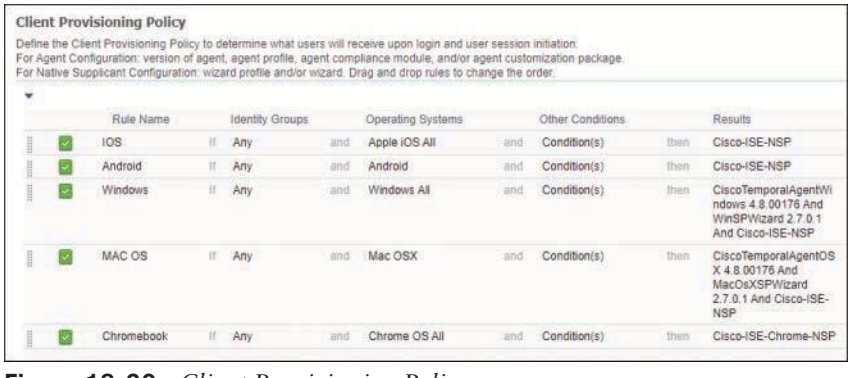


Figure 18-30 *Client Provisioning Policy*

You can use the default temporal agent, or you can follow these steps to prepare a new rule for Windows devices used by employees and provision AnyConnect to them:

Step 1. Click the down arrow next to the Windows rule in the CPP and choose **Duplicate Above**, as shown in Figure 18-31.

Step 2. Name the rule **Windows Employees**.

Step 3. Under Other Conditions, set the Active Directory group to **Employees**.

Step 4. For Results, select the **AnyConnect Configuration** as the agent.

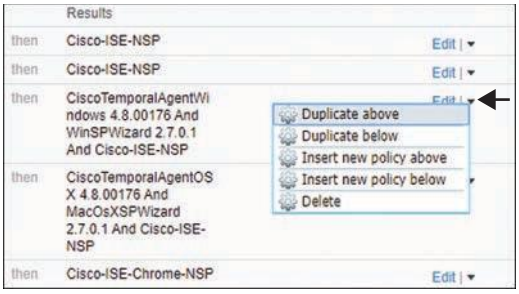


Figure 18-31 *Duplicating a CPP Rule*

NOTE In this case you do not see a compliance module dropdown because it is configured as part of the AnyConnect configuration.

Step 5. Click **Done**.

Step 6. Click **Save**.

Figure 18-32 shows the new Windows Employees rule being added to the CPP.

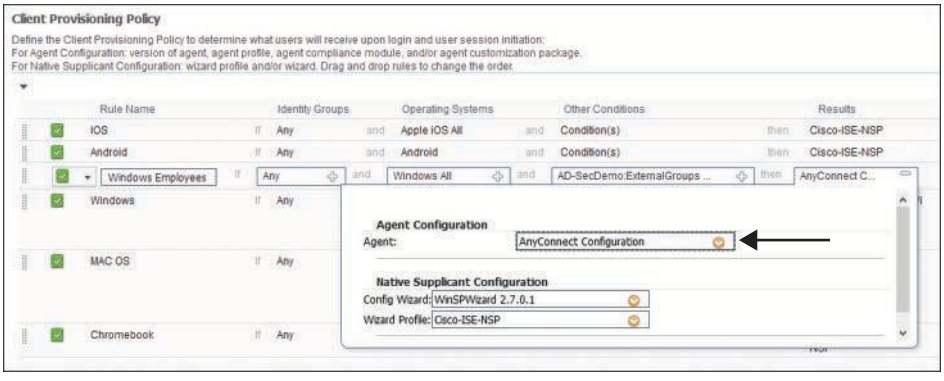


Figure 18-32 *Windows Employees CPP Rule*

Now that the client provisioning is configured, we can move on to configuring the elements used in the posture policy.

Configuring Posture Policy Elements

The posture runtime in ISE was directly ported from the Perfigo acquisition in 2004 that became the NAC Appliance. The relationship between the elements is the same as well, as illustrated in Figure 18-33.

The posture elements contain conditions, remediations, and requirements. Conditions and remediations are linked together into requirements. Requirements are used in the posture policy.

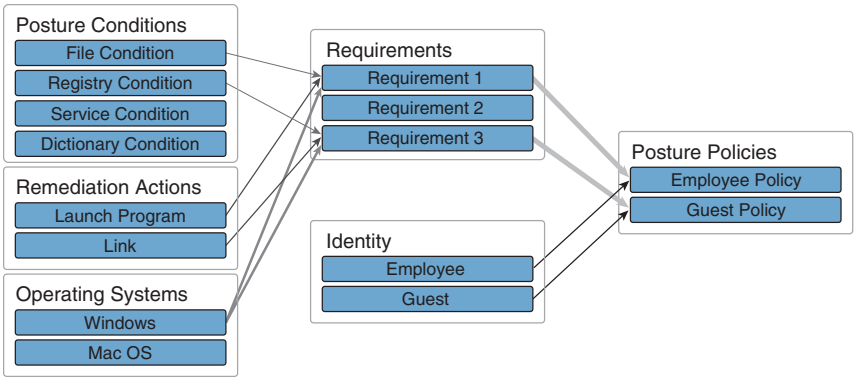


Figure 18-33 *Posture Elements and Their Relationship*

Conditions

Conditions used to be called "checks" because that is the function they perform. A condition checks an endpoint for the existence of certain artifacts or performs a bulk collection of those artifacts.

Figure 18-34 shows the types of conditions that exist.



Figure 18-34 *Condition Types*

The following sections review the condition types and provide some examples.

Hardware Attributes Condition

The hardware attributes condition is preconfigured to collect details about the endpoint hardware to aid in the endpoint context visibility aspects of ISE and can be viewed in the Context Visibility > Endpoints > Hardware dashboard.

Figure 18-35 shows the hardware attributes condition. As you can see, it cannot be edited or modified.

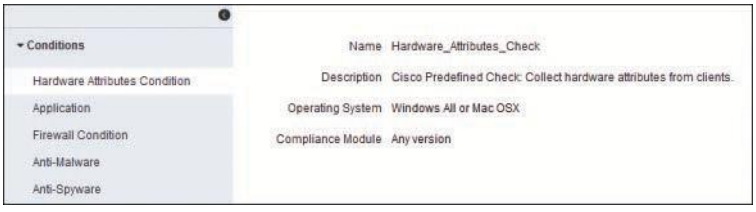


Figure 18-35 *Hardware Attributes Condition*

Figure 18-36 shows the Context Visibility > Endpoints > Hardware dashboard.

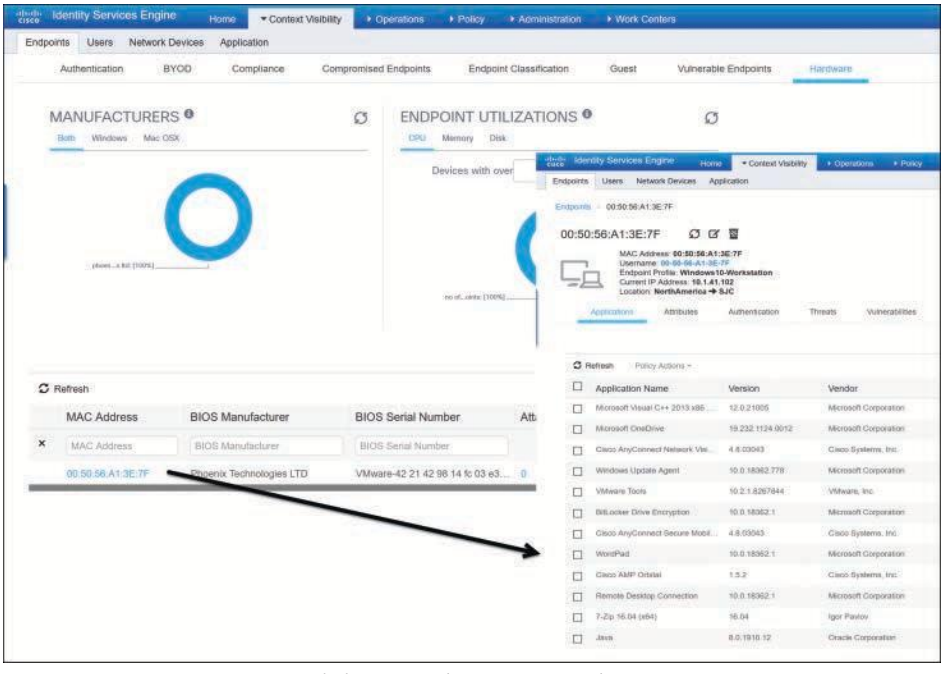


Figure 18-36 *Context Visibility > Endpoints > Hardware*

Application

Application conditions have two purposes:

- To collect application inventory and pass that inventory to ISE for context visibility of the endpoints
- To check for a specific application that is running or installed for posture compliance

To see application conditions, navigate within the Posture work center to Policy Elements > Conditions > Application. Figure 18-37 shows that ISE has predefined checks (conditions) that are part of the posture updates from Cisco.com.

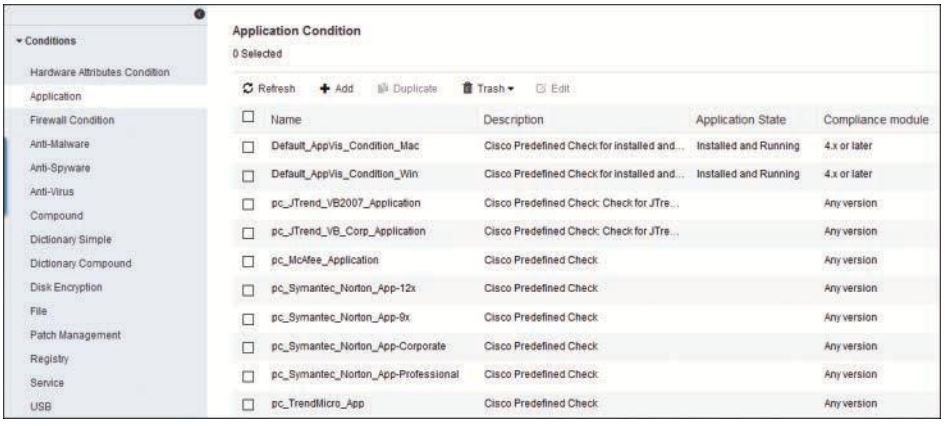


Figure 18-37 *Predefined Application Conditions*

Click Add to create a new condition. Then you can provide a name and, optionally, a description. Clicking the plus sign next to Operating System and select Windows All, as shown in Figure 18-38. (Alternatively, you can choose a very specific version of Windows, as shown in Figure 18-39.)

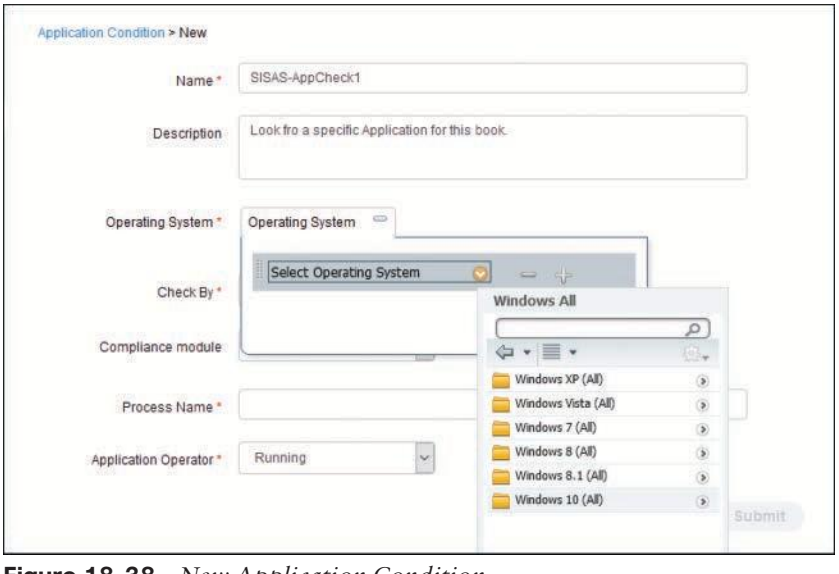


Figure 18-38 *New Application Condition*

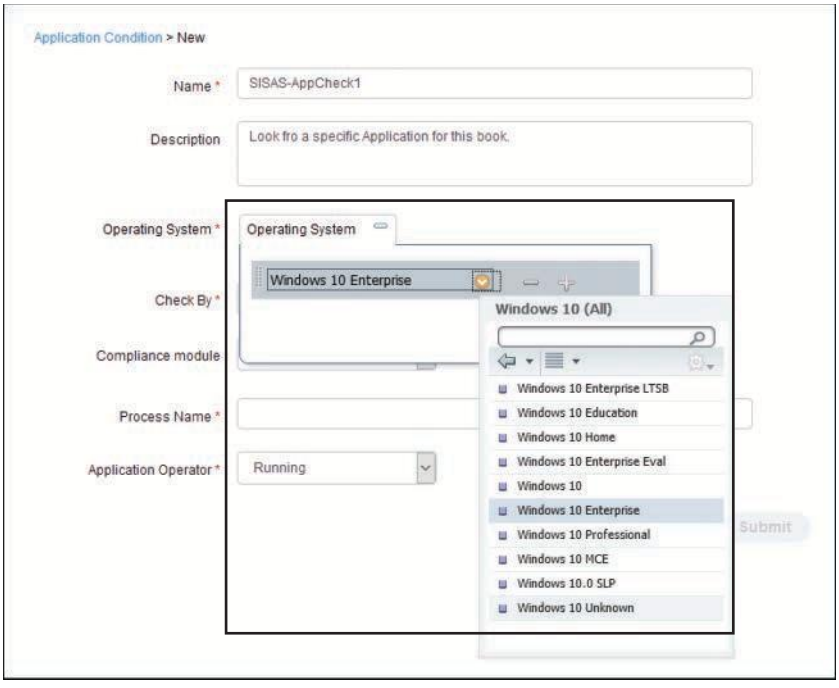


Figure 18-39 *Specific OS Versions*

Next, you can select to check by either application or process. For now, in the Check By dropdown, select Choose Application (see Figure 18-40). This type of check is available only when Compliance Module is set to 4.x or Later, so you are unable to make any other selection from that dropdown. For Application State, you can indicate whether to check the state of the application being installed and identify whether it is running. Finally, if you leave the Provision By dropdown set to Everything, AnyConnect collects information on every installed application and determines whether those applications are running. That information is collected for context visibility, and it is not used when determining the posture compliance of an endpoint.

If you change the Provision By dropdown to Category, the list of categories appears below the dropdown, as shown in Figure 18-41. You can then check the categories for which a posture check collects and reports data to ISE for context visibility.

If you change the Provision By dropdown to Category, the same list of categories appears below the dropdown, as shown in Figure 18-42. This time, when you select a category, you can also select a specific vendor of software within that category, and you can even specify versions of the software to collect and report back to ISE for the Context Visibility inventory.

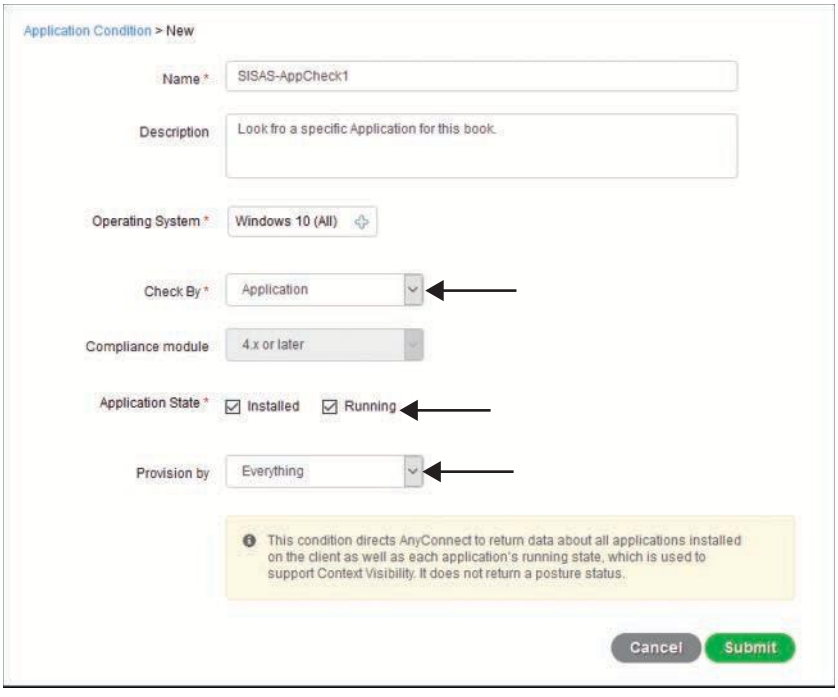


Figure 18-40 *Application Condition: Checking by Application and Provisioning by Everything*

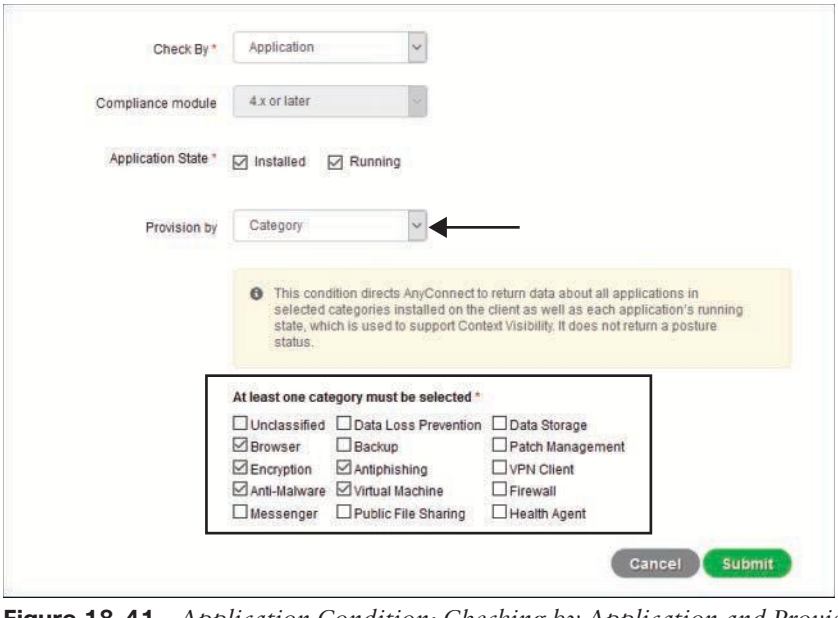


Figure 18-41 *Application Condition: Checking by Application and Provisioning by Category*

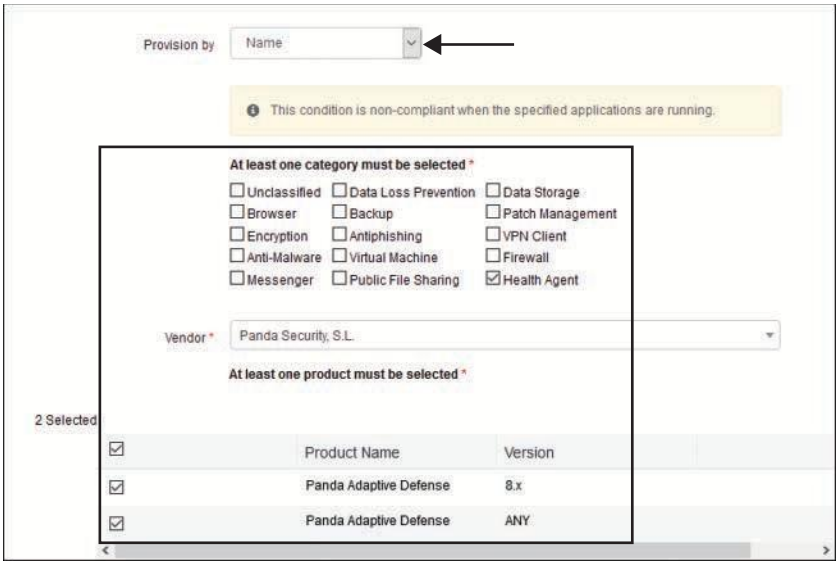


Figure 18-42 *Application Condition: Checking by Application and Provisioning by Name*

If you change the Check By dropdown to Process, you can type in the name of a process and choose to check whether the application is running or not. This allows you to configure ISE to return a noncompliant verdict if the application is not in the state you want it to be in. For example, ISE can deny access if an endpoint is running a virtualization product or if an endpoint is not running a specific corporate application. Figure 18-43 shows an application condition that is looking for notepad.exe to be running.

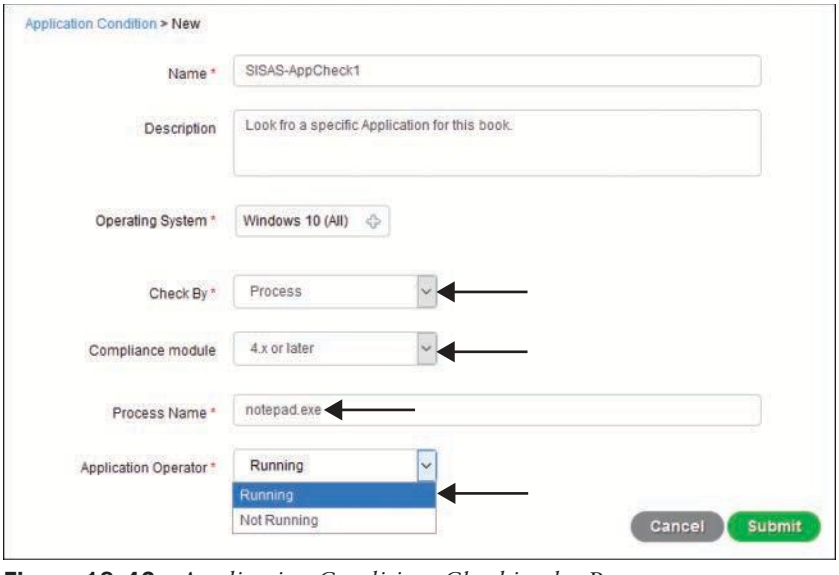


Figure 18-43 *Application Condition: Checking by Process*

When you have finished making selections in this section, click Submit to save the process check for use in the posture policy later.

Firewall Condition

The next condition type in the list is the firewall condition. As shown in Figure 18-44, there are preconfigured checks for this condition to make your life a little easier.

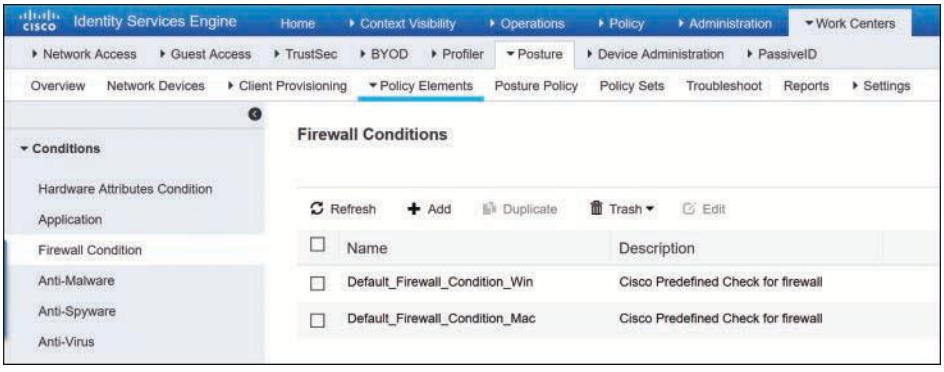


Figure 18-44 *Predefined Firewall Conditions*

If you edit one of the preconfigured rules, you see that the rules are configured to identify firewalls from any vendor (see Figure 18-45). Setting Vendor to Any really makes things easier for you when it comes to configuring checks. (It can be painful to try doing the same thing with ASA and DAP policies.)

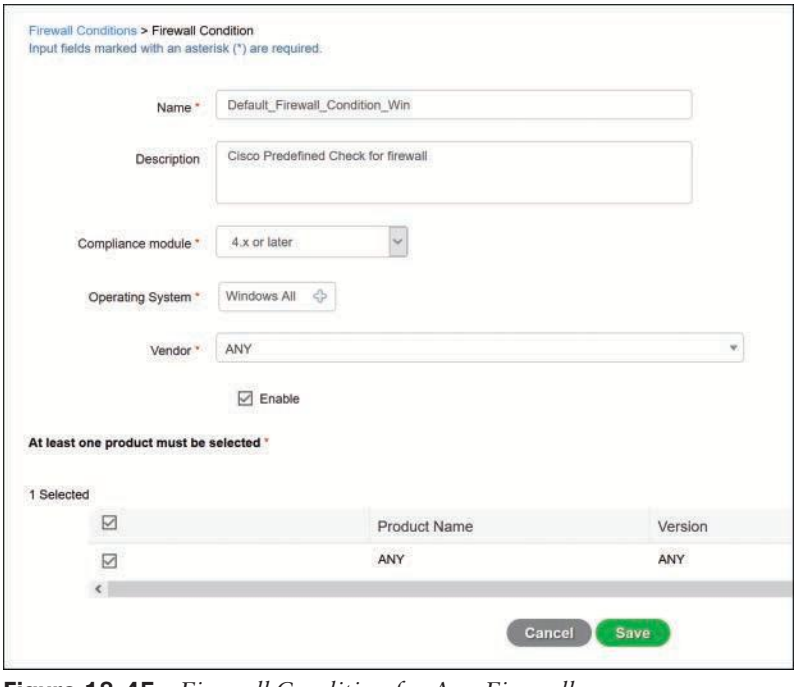


Figure 18-45 *Firewall Condition for Any Firewall*

If you click Add, you can create a custom check for a host-based firewall, as shown in Figure 18-46. Note that the list of vendors in the Vendor dropdown can be a bit overwhelming, as Figure 18-47 illustrates.

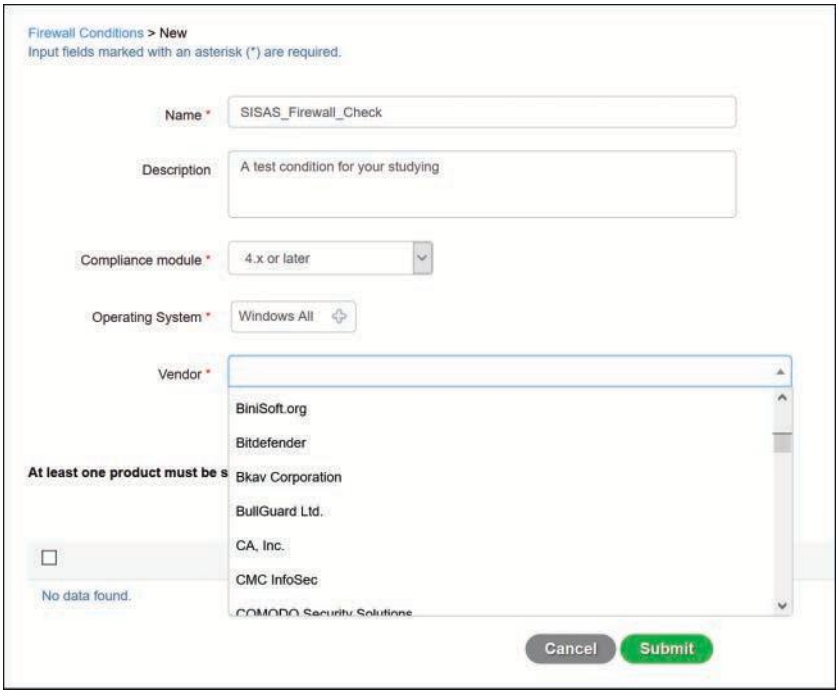


Figure 18-46 *Adding a New Firewall Condition*

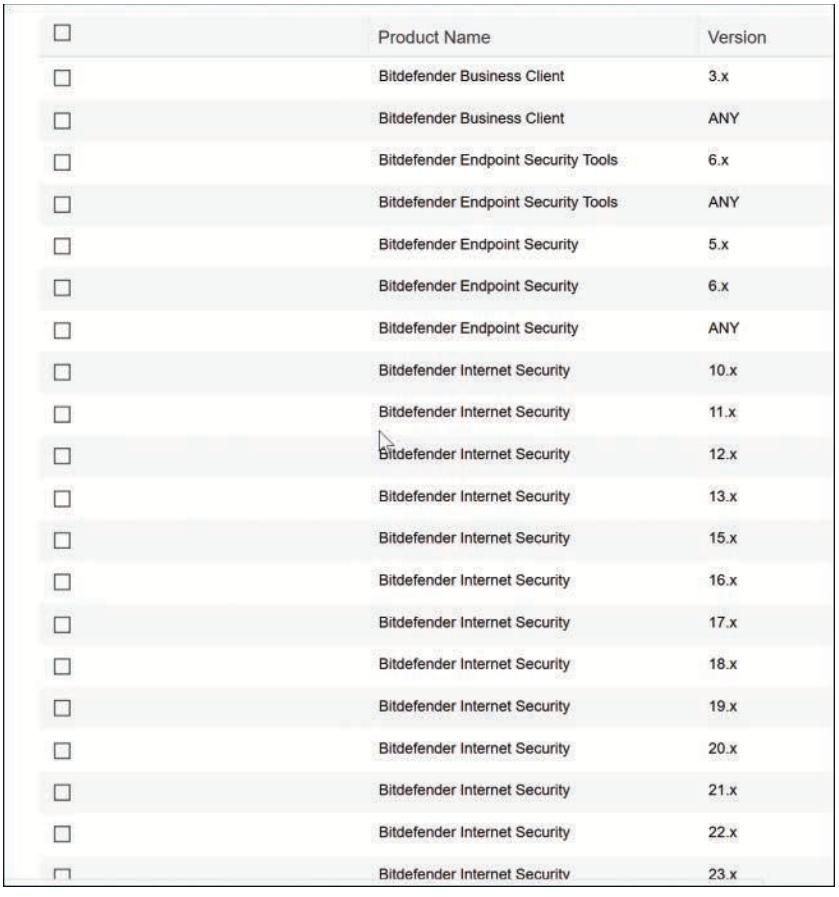
Unless your organization requires the use of a very specific version of a very specific vendor's firewall, you should just use the prebuilt checks for Any.

Anti-Malware

Anti-malware is a single condition that replaces both the anti-spyware and anti-virus conditions, starting with compliance module 4.0. The condition checks for both installation and the definition version of an anti-malware product.

An endpoint detection and response (EDR) platform should not rely very heavily on definitions, as it should be more focused on behaviors or communicating to the cloud. Cisco's AMP for Endpoints is such an EDR client. Because there is no need for definition files, the posture check only needs to ensure that the client is installed and running the right version.

Traditional endpoint protection platforms (EPPs), such as anti-virus and anti-spyware, rely on signatures, much as an intrusion prevention system (IPS) does. The product is only as good as the signatures it has present to detect the viruses, and that is why it is also important to ensure that the definition files are as up to date as possible.



	Product Name	Version
☐	Bitdefender Business Client	3.x
☐	Bitdefender Business Client	ANY
☐	Bitdefender Endpoint Security Tools	6.x
☐	Bitdefender Endpoint Security Tools	ANY
☐	Bitdefender Endpoint Security	5.x
☐	Bitdefender Endpoint Security	6.x
☐	Bitdefender Endpoint Security	ANY
☐	Bitdefender Internet Security	10.x
☐	Bitdefender Internet Security	11.x
☐	Bitdefender Internet Security	12.x
☐	Bitdefender Internet Security	13.x
☐	Bitdefender Internet Security	15.x
☐	Bitdefender Internet Security	16.x
☐	Bitdefender Internet Security	17.x
☐	Bitdefender Internet Security	18.x
☐	Bitdefender Internet Security	19.x
☐	Bitdefender Internet Security	20.x
☐	Bitdefender Internet Security	21.x
☐	Bitdefender Internet Security	22.x
☐	Bitdefender Internet Security	23.x

Figure 18-47 *An Overwhelming Number of Choices per Vendor*

Figure 18-48 shows the preconfigured checks for the anti-malware condition, including checks for installation and definitions for all anti-malware products for Windows and macOS.

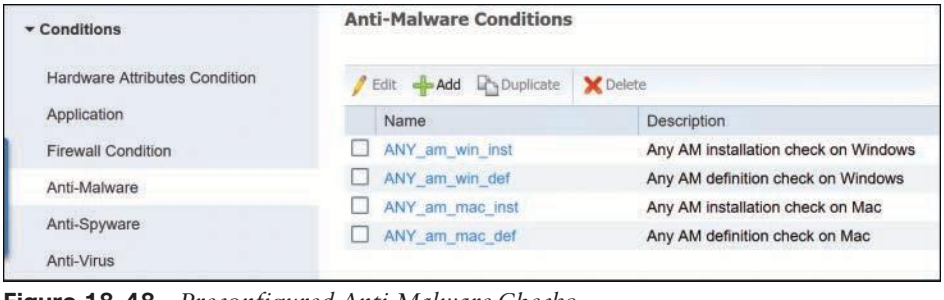


Figure 18-48 *Preconfigured Anti-Malware Checks*

Adding a new condition can help you see the options. For example, the settings shown in Figure 18-49 show some of the options that are available for this type of check.

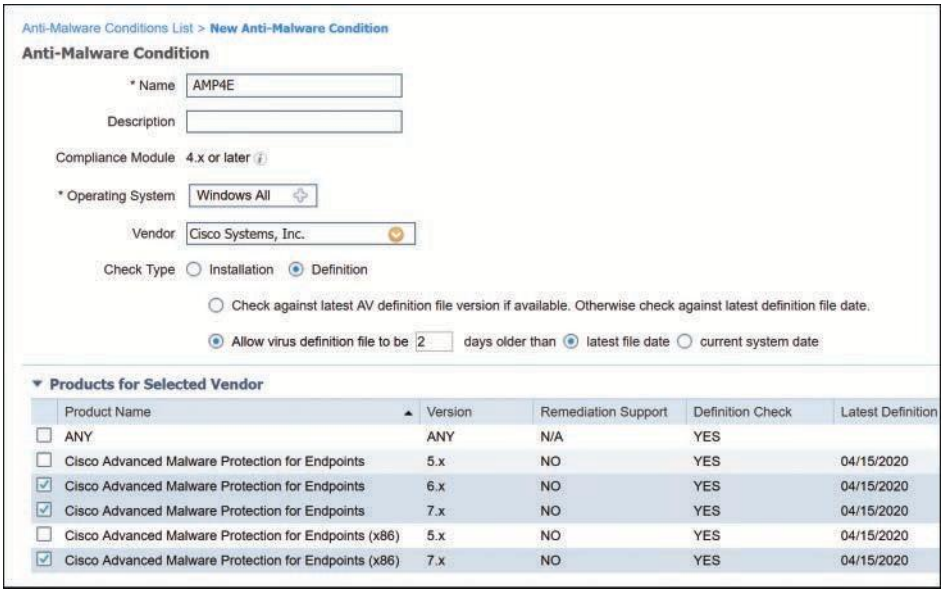


Figure 18-49 *Adding a New Anti-Malware Condition*

Figure 18-49 shows the Cisco AMP for Endpoints product selected. You can choose the version or select Any to select all versions. AMP is both an EDR and an EPP, so it does have a definition file that can be examined as well as the software installation and version. Notice in Figure 18-49 that you can create a buffer of allowed time for definition files to be out of date, and you can also force the definition file to always be up to date.

Anti-Spyware

The anti-spyware condition is valid only when you use compliance module 3.x or earlier. The anti-malware condition replaces it in compliance modules 4.0 and later.

Some preconfigured Any anti-spyware checks come as part of the posture updates, just as with the with anti-malware and firewall conditions, but they only work with compliance module 3.x and earlier.

Anti-Virus

The anti-virus condition is valid only when you use compliance module 3.x or earlier. The anti-malware condition replaces it in compliance modules 4.0 and later.

Some preconfigured Any anti-virus checks come as part of the posture updates, just as with the anti-malware and firewall conditions, but they only work with compliance module 3.x and earlier.

Dictionary Simple

At this point in the chapter, we are breaking from the order in which the conditions are presented in the user interface because the order in the UI does not make logical sense.

Compound conditions require other conditions to be created first, so in this section, we review the simple conditions first. By default, there are no dictionary simple conditions, but you can add them.

To see a simple condition, you can add one by clicking Add. Then you can provide a name and, optionally, a description. Then you choose a condition. In the example shown in Figure 18-50, a condition is being created for Network Access:AuthenticationIdentityStore, which contains the value Demo. Figure 18-51 shows an example of another simple condition being added—in this case, the condition of Endpoints:OperatingSystem, which matches the value ndows.

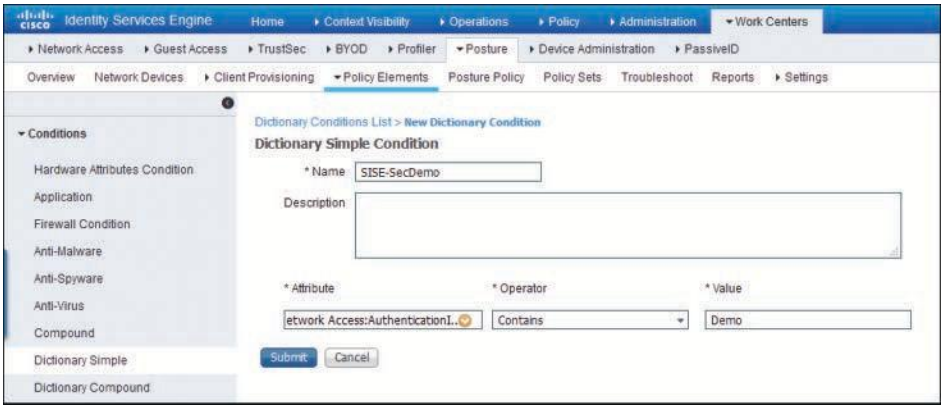


Figure 18-50 *First Dictionary Simple Condition*

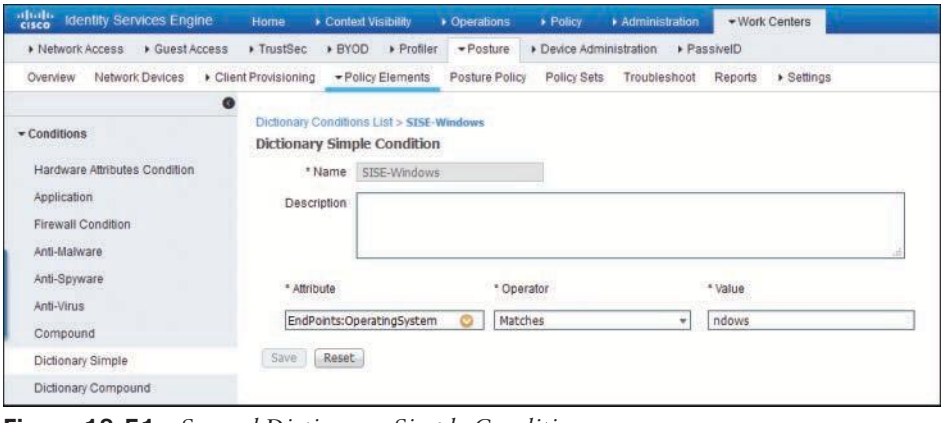


Figure 18-51 *Second Dictionary Simple Condition*

After you create a condition, be sure to click Submit to save it.

Dictionary Compound

A dictionary compound condition is made up of one or more dictionary simple conditions. By default, there are no dictionary compound conditions, but you can add them.

Click Add to create a new compound condition. Then provide a name and, optionally, a description, as shown in Figure 18-52.

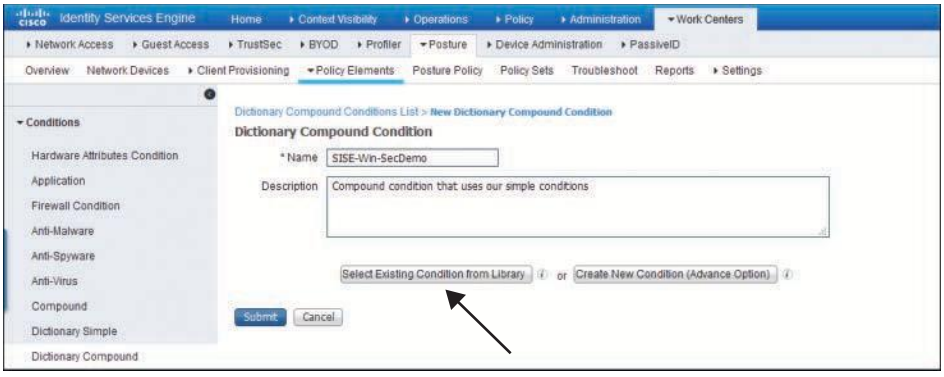


Figure 18-52 *Adding a New Dictionary Compound Condition*

Click Select Existing Condition from Library. Choose each of the conditions that you created and set the operator dropdown to AND, as shown in Figure 18-53.

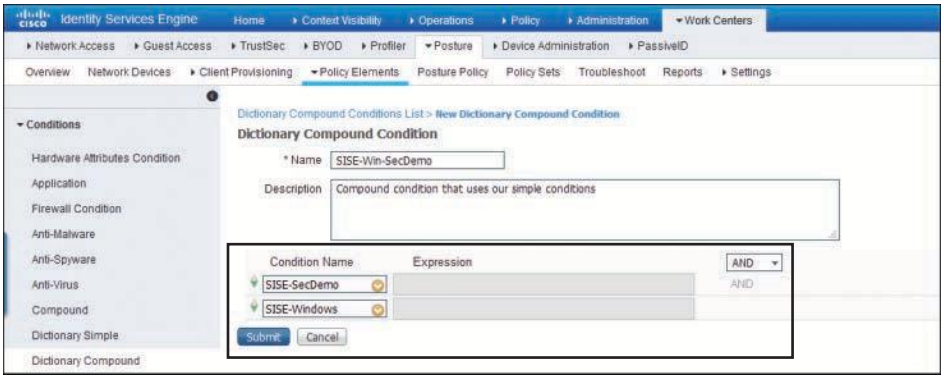


Figure 18-53 *Adding Simple Conditions to a Compound Condition*

Click Submit to save the condition.

Disk Encryption

There are many different vendors for disk encryption for both Windows and macOS, and an endpoint might have different drives, each with a different state encryption condition.

To add a disk encryption condition, click Add. Name the condition, select Mac OSX as the operating system, and set Compliance Module to 4.x or Later. Then click the Vendor Name dropdown to see what vendors the posture module can look for (see Figure 18-54).

Select Apple Inc and select the FireVault product name. Then select Encryption State and ensure that Location is set to All Internal Drives and Is Fully Encrypted is checked, as shown in Figure 18-55.

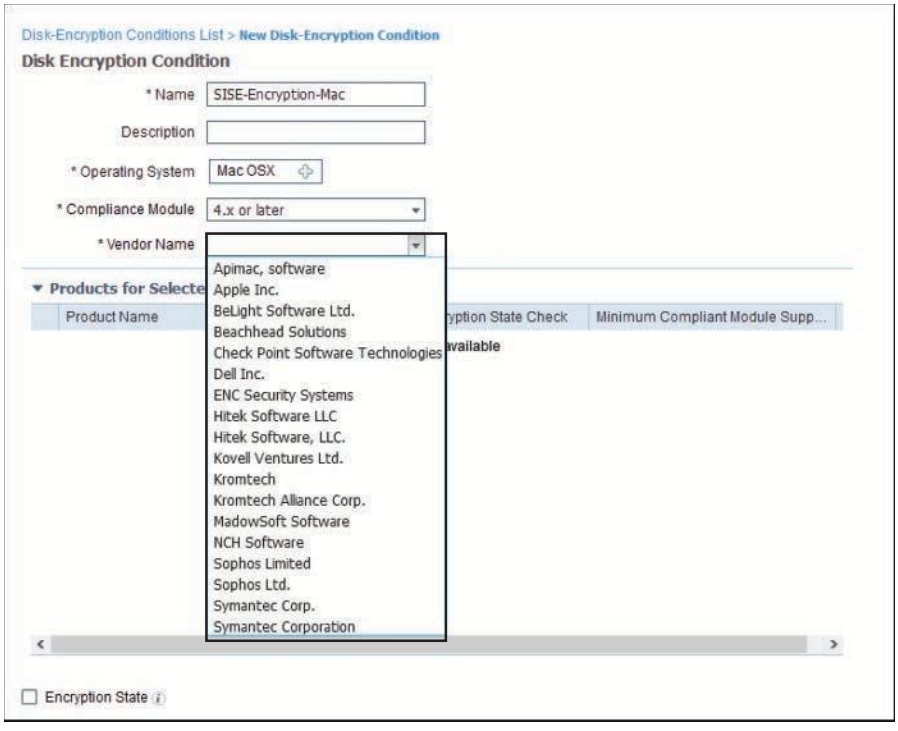


Figure 18-54 *Adding a Disk Encryption Condition*

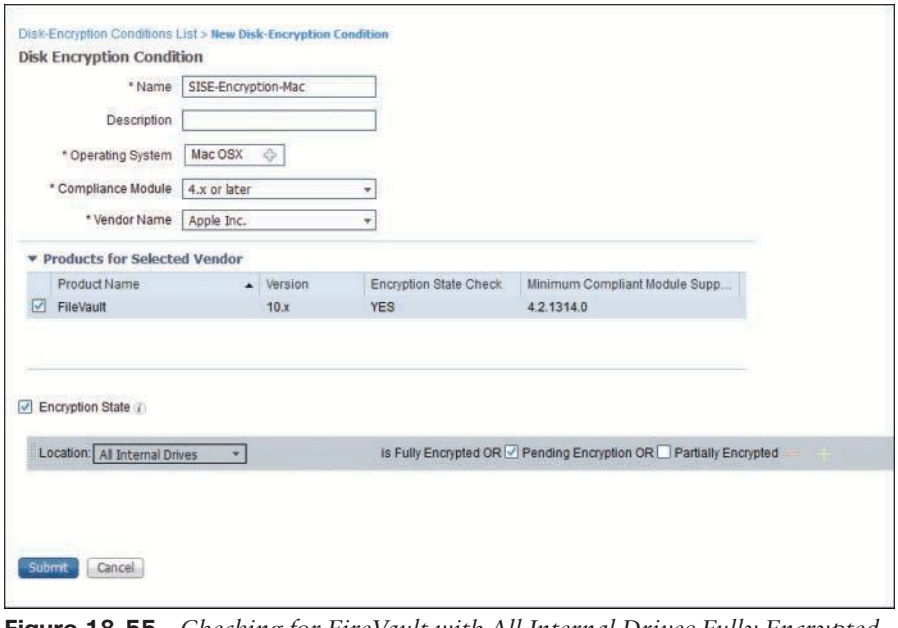


Figure 18-55 *Checking for FireVault with All Internal Drives Fully Encrypted*

Click Submit to save the condition.

File

File checks are among the most popular posture conditions leveraged by ISE customers. The file condition is often used to locate a watermark file, which is a custom file located somewhere very obscure on system endpoints to help identify the endpoints that are truly managed by the company.

In addition, there are a lot of preconfigured file conditions. In fact, the screen shown in Figure 18-56 indicates that there are 370 of them.

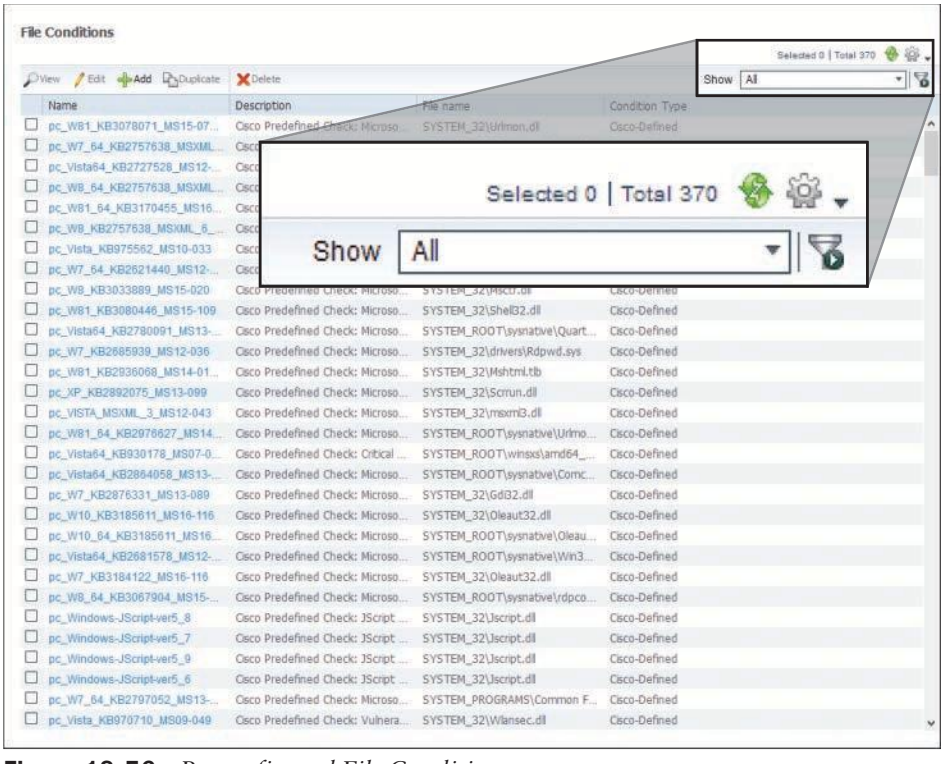


Figure 18-56 *Preconfigured File Conditions*

Click on pc_W7_64_KB2757638_MSXML_6_MS13-002, which is a predefined check for the Windows 7 MS13-002 patch, as shown in Figure 18-57. As you can see, selecting a file condition can help you see what the condition can check for.

In Figure 18-57, you can see that this check is designated for Windows 7 only. It can work with any compliance module, and it is looking for a file modification date later than 10/30/2012 for the C:\Windows\synative\msxml6.dll file. But where did the C:\Windows\ path come from? In Figure 18-57, SYSTEM_ROOT is selected from the File Path dropdown. This means the search path will start in the directory in which Windows is installed.

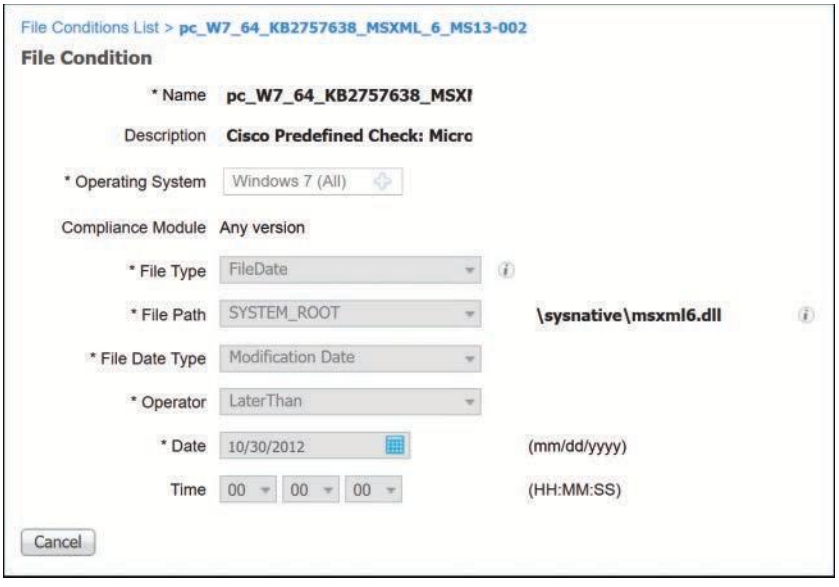


Figure 18-57 *Example Predefined Check for MS13-002 Patch*

Click Cancel to return to the file conditions. Click Add to create a new condition. Provide a name and select Windows All from the Operating System dropdown. When you click the File Type dropdown, you see the different ways to configure the file check, as shown in Figure 18-58.

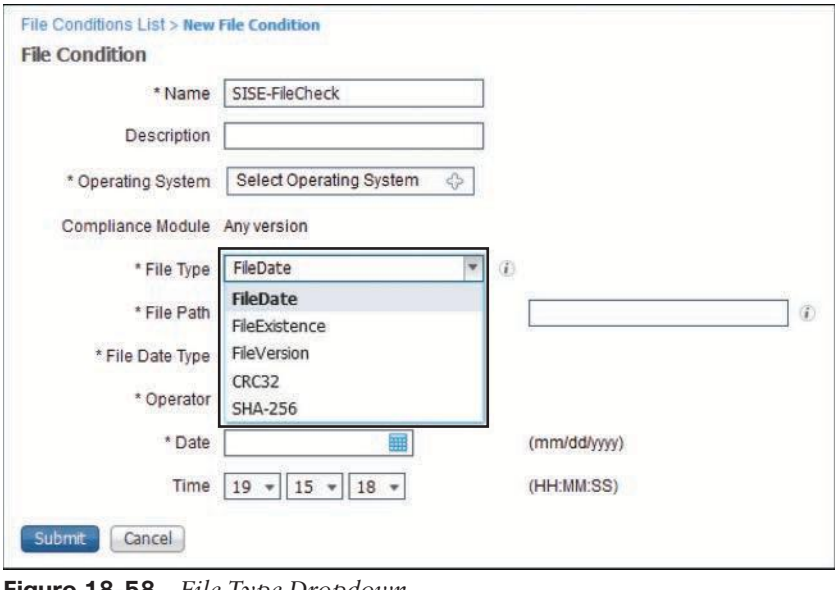


Figure 18-58 *File Type Dropdown*

Figure 18-58 shows the five file type options:

- **FileDate:** When you select this file type, you then need to select the location and name of the file and specify the file path. There are seven predefined File Path drop-down options (see Figure 18-59):

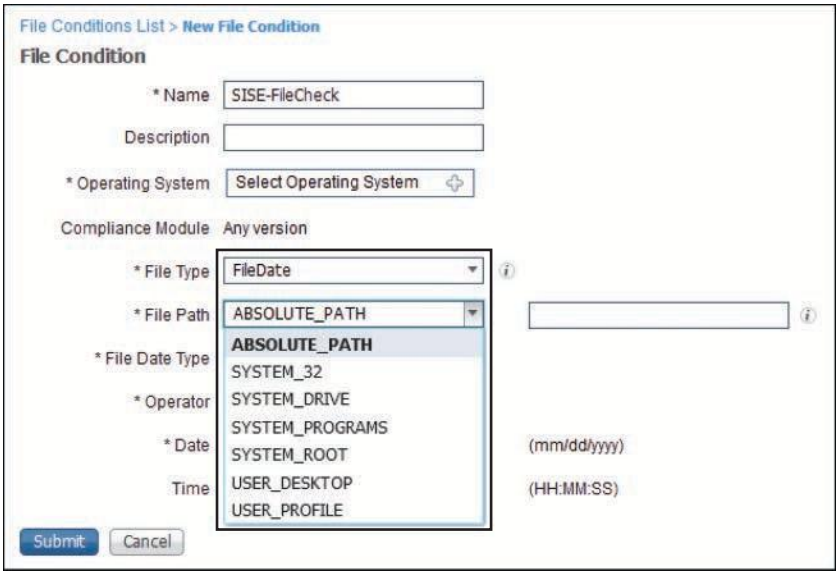


Figure 18-59 *File Path Options*

- **ABSOLUTE_PATH:** As the name implies, with this file type, you enter the full path to the file in the system (for example, *C:\CompanyName\watermark.xml*), as shown in Figure 18-60.
- **SYSTEM_32:** This option checks for files starting in the file path with the System32 subdirectory of the Windows folder—wherever that Windows installation folder exists (for example, D:\Windows\System32). After selecting SYSTEM_32, you type the filename and path in the box, as shown in Figure 18-61.
- **SYSTEM_DRIVE:** This option checks for files starting in the root of the drive on which Windows is installed (for example, C:\ or D:\).
- **SYSTEM_PROGRAMS:** This option checks for files starting in the \Program Files\ directory. It can also check in the system drive or on another drive, depending on how customized the Windows installation is.
- **SYSTEM_ROOT:** This option begins the file check in the \Windows\ directory—wherever it is installed (for example, C:\Windows\ or D:\Windows\).
- **USER_DESKTOP:** This option checks whether the file exists on the user's desktop.
- **USER_PROFILE:** This option checks whether the file exists in the user's local profile directory.

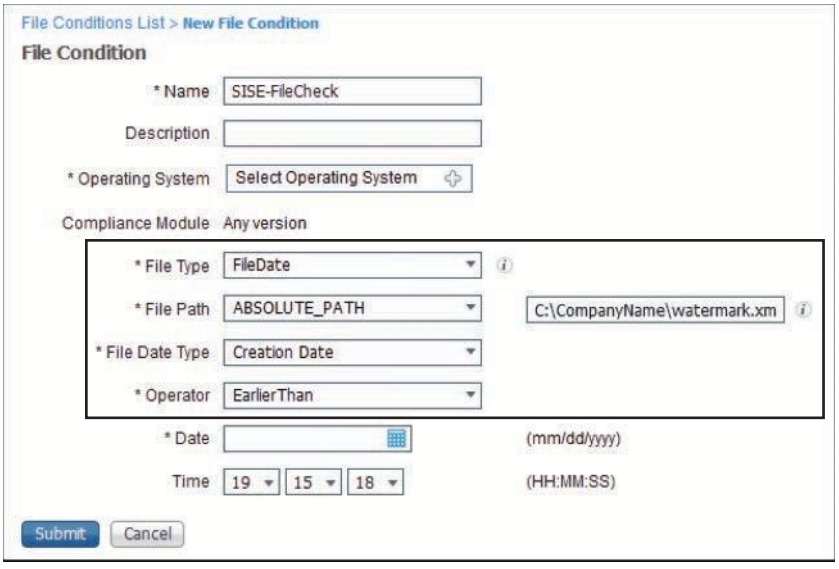


Figure 18-60 *ABSOLUTE_PATH File Path Option*

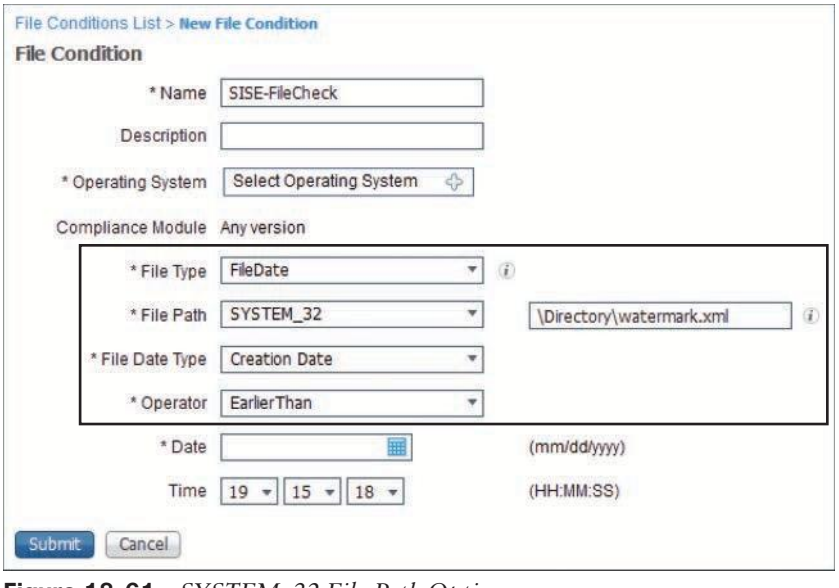


Figure 18-61 *SYSTEM_32 File Path Option*

After choosing the file path, you must choose the file date type, which can be the creation date or modification date that the condition checks for. You also select an operator, which can be EarlierThan, LaterThan, EqualTo, or Within a time range, as shown in Figure 18-62.

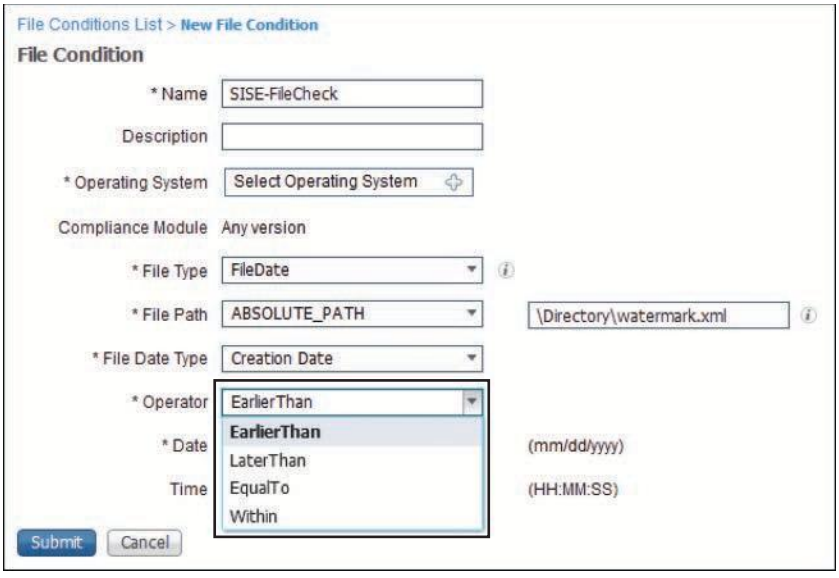


Figure 18-62 *Operator Choices*

Figure 18-63 shows a completed file condition using the FileDate file type.

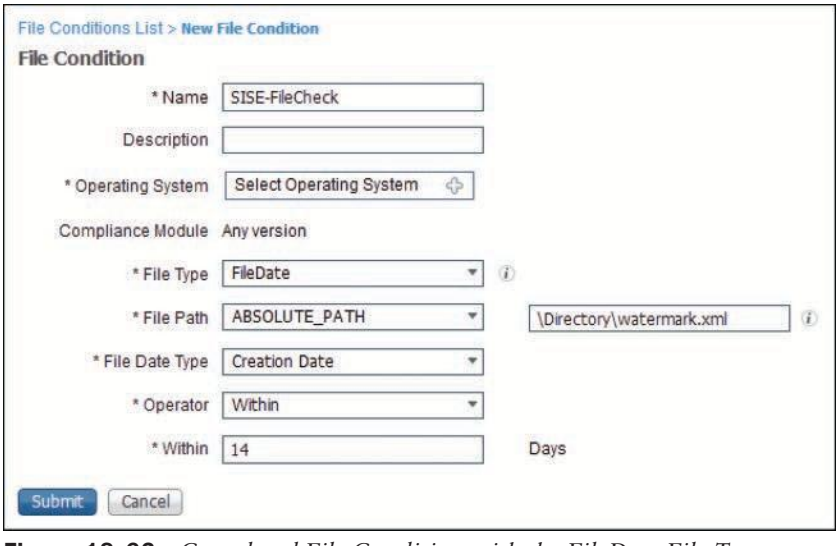


Figure 18-63 *Completed File Condition with the FileDate File Type*

- **FileExistence:** This file type looks to see if a file exists where it is supposed to—and that is all. With this option selected, there is no concern at all for validating the file dates, hashes, and so on. Figure 18-64 shows an example of using the FileExistence option.

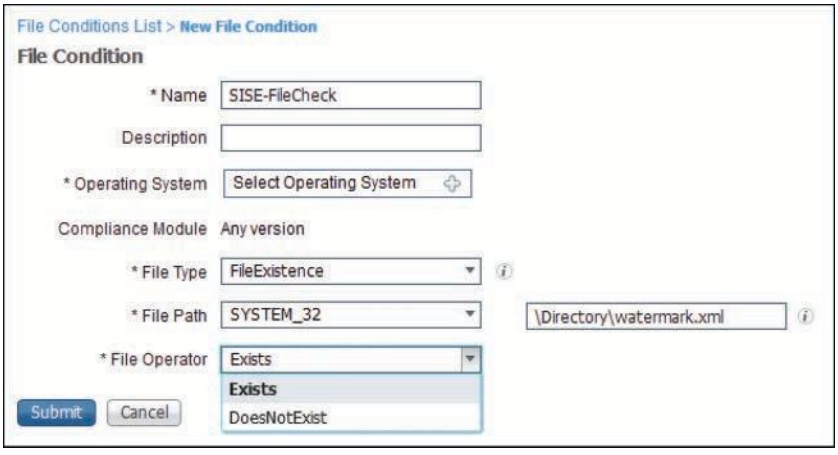


Figure 18-64 *Completed File Condition with the FileExistence File Type*

- **FileVersion:** This file type checks for the existence of a file and looks to see if the version of that file meets the specified criteria, as shown in Figure 18-65. The version can be set to EarlierThan, LaterThan, or EqualTo.

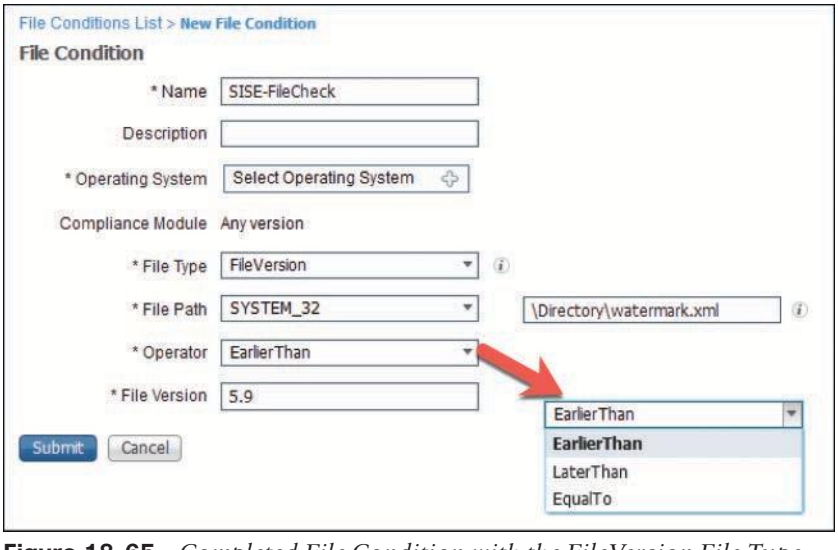


Figure 18-65 *Completed File Condition with the FileVersion File Type*

- **CRC32:** This file type validates the integrity of a file by using the checksum function, which is a hash of the file. Any changes to the file at all result in a different checksum value.
- **SHA-256:** Just like the CRC32 option, this option validates the hash of the file, but it uses a SHA-256 hash function, which has a much lower chance of collision (that is, two different files generating exactly the same hash).

Select a condition that you would like to use later in your practice and study and then click Submit to save the file condition for later use.

Patch Management

A patch management condition checks for the installation and state of software related to device managers that control the endpoint and ensure that the patches are up to date. This condition is very popular with enterprises as it allows the access control solution (ISE) to tie in to the chosen device management platform (such as SCCM for Windows or Jamf Pro for macOS).

This type of condition is easiest to understand by interacting with the UI. Click Add to create a new patch management condition, as shown in Figure 18-66.

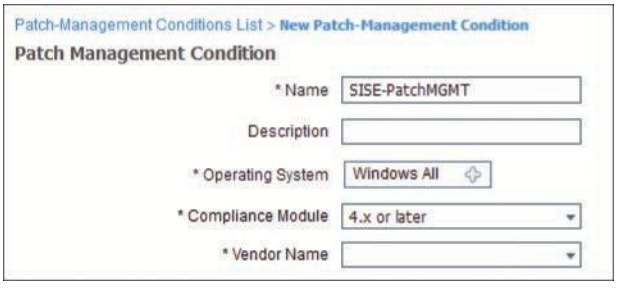


Figure 18-66 *New Patch Management Condition*

If you select Windows All from the Operating System dropdown, when you click the Vendor Name dropdown, you can see the large number of systems that the posture module can detect and work with (see Figure 18-67).

Select Microsoft Corporation from the Vendor Name dropdown, and the list of Microsoft products for patch management appears in the panel at the bottom of the screen, as shown in Figure 18-68. These are the three main choices:

- **Microsoft Intune Client:** Intune is the device manager that Microsoft offers through Microsoft Azure. It is capable of managing device policies and software installation on multiple operating systems (not just Windows).
- **System Center Configuration Manager (SCCM):** SCCM is the traditional on-premises device manager from Microsoft that the vast majority of enterprises use to maintain software and patches for their Active Directory–controlled Windows workstations.
- **Windows Update Agent:** This is the built-in Windows update tool that is used by default to keep Windows patched when another device manager is not leveraged. It uses a public Microsoft update service on the Internet and is very commonly used, even in enterprises that have BYOD policies.

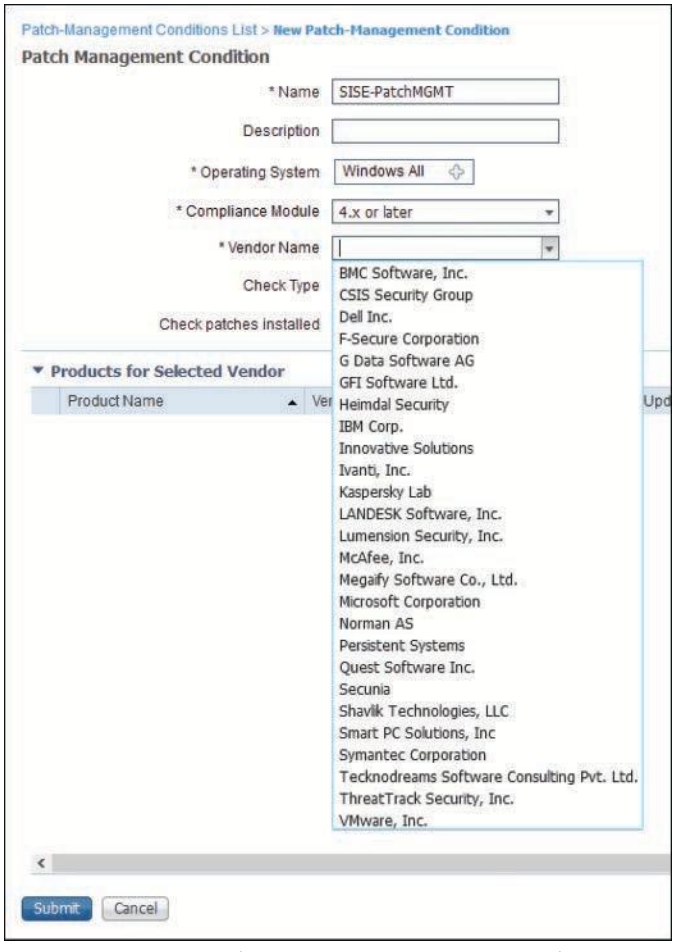


Figure 18-67 *Windows Patch Management Vendors List*

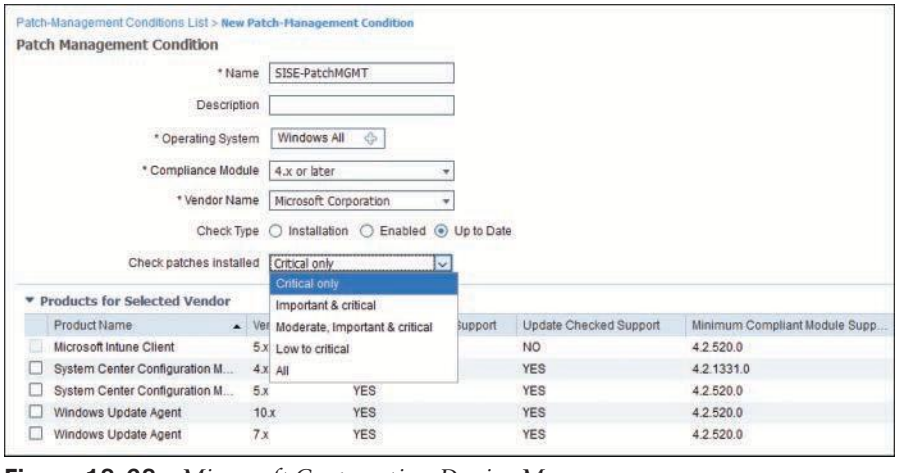


Figure 18-68 *Microsoft Corporation Device Managers*

When choosing any patch management vendor, you can build a posture condition to verify that the patch management agent is installed or is installed and enabled, or you can check whether the endpoint is up to date. Verifying that the device manager lists the endpoint as being up to date is what most businesses think they want, although verifying that the agent is installed and enabled is also common.

Figure 18-68 shows the patch management condition being created with Microsoft Corporation as the vendor, and Up to Date selected as the Check Type option. Notice that with Up to Date selected, the Intune client is grayed out. This means that installation can be verified, but the status of the patches cannot be verified.

As you can see in Figure 18-68, you can select from the Check Patches Installed dropdown to configure the condition to validate patches based on these categories:

- Critical Only
- Important & Critical
- Moderate, Important & Critical
- Low to Critical
- All

Registry

Just like file checks, Registry checks are very common and popular with ISE customers. These conditions help identify watermark Registry keys, which are keys or values that are buried deep in the Windows Registry and used to identify a system as being a managed endpoint of the organization.

As this book went to press, 191 preconfigured Registry conditions could be used to identify installed patches and installed software.

Click Add to create a custom Registry condition. As with all the other conditions, you can specify the operating system. In this case, the OS can only be a Windows version, but you could build the condition to apply only for a specific version of Windows, such as Window 10 or Windows 7.

For a Registry condition, you can select one of three options from the Registry Type dropdown:

- **RegistryKey**: You can select this option to validate whether the key exists.
- **RegistryValue**: You can select this option to indicate whether values can be numbers, strings, or versions.
- **RegistryValueDefault**: You can select this option to check the value of the Registry key and look to see if it is a default setting for the key.

Figure 18-69 shows an example of a Registry condition checking for the existence of a watermark key. Figure 18-70 shows an example using the RegistryValueDefault type.

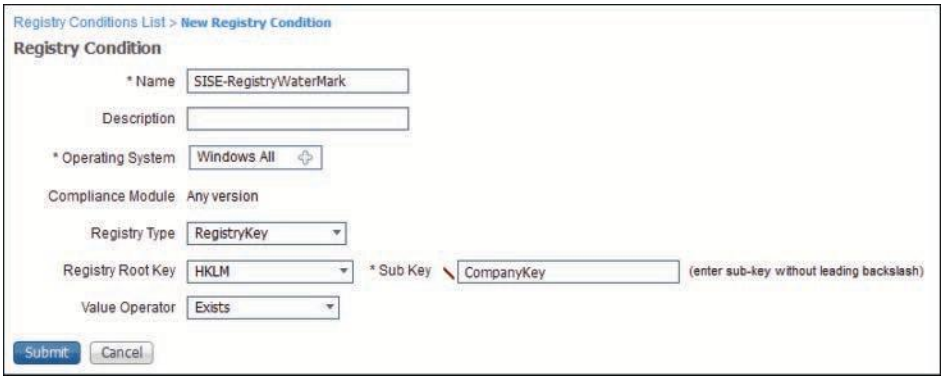


Figure 18-69 *RegistryKey Condition Example*

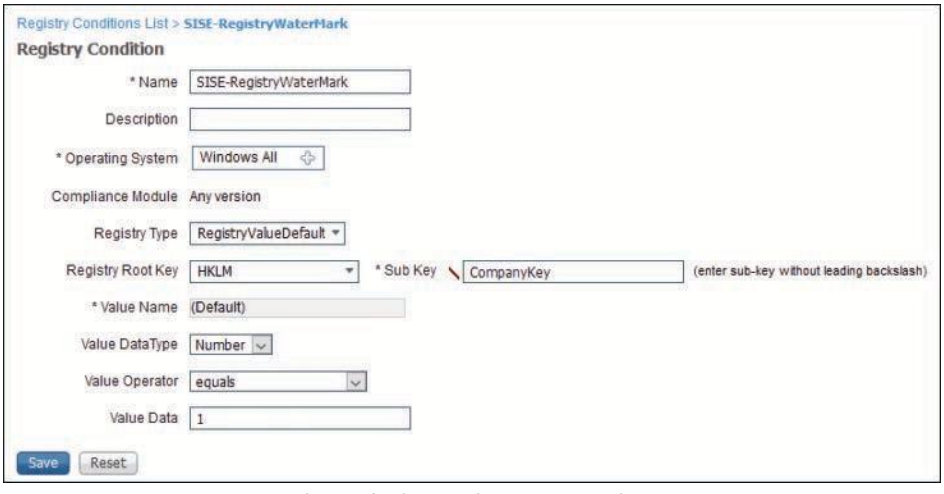


Figure 18-70 *RegistryValueDefault Condition Example*

Service

The service condition verifies that a Windows service or macOS daemon is running or ensures that a service or daemon is not running. Figure 18-71 shows the three preconfigured service conditions, including the long-dead Cisco Security Agent (CSA) service. Figure 18-72 shows a new custom service condition.

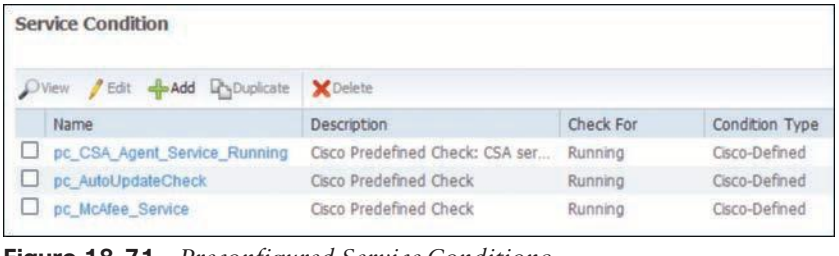


Figure 18-71 *Preconfigured Service Conditions*

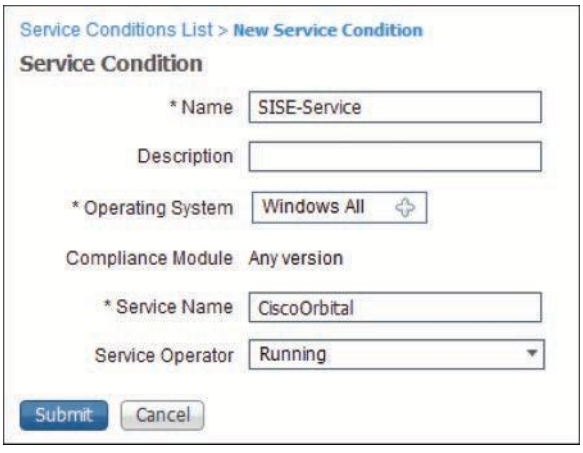


Figure 18-72 *Windows Service Condition Example*

Compound

A compound condition is made up of one or more simple conditions. Those simple conditions can be nearly any of the other conditions that exist. Cisco has created 40 prebuilt compound conditions. Figure 18-73 shows one of the preconfigured compound conditions.

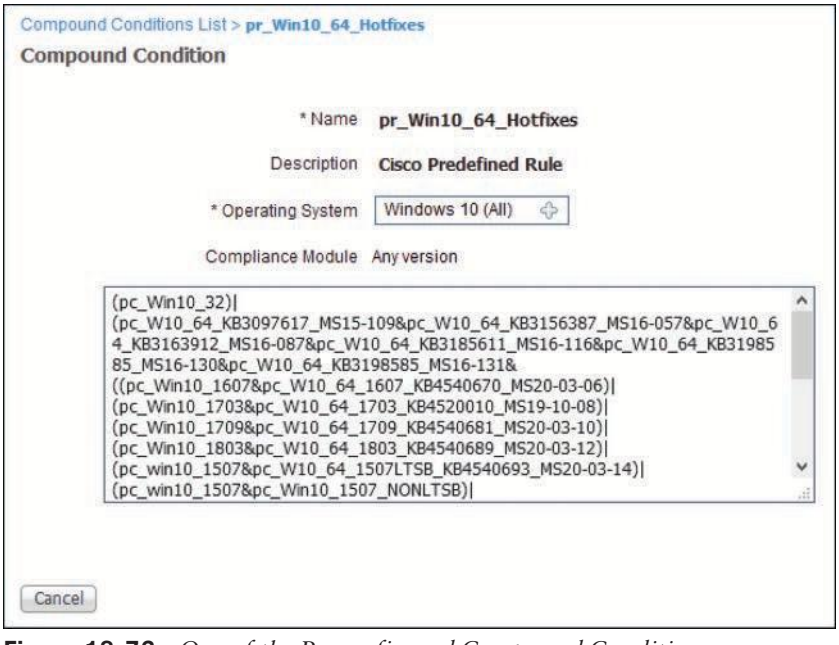


Figure 18-73 *One of the Preconfigured Compound Conditions*

The preconfigured compound conditions are basically just like mathematical formulas with a full order of operations, which can be illustrated as follows:

(Condition1 | Condition2) & Condition3

This translates to Condition1 OR Condition2 AND Condition3.

Figure 18-74 shows a new compound condition being created, with the expression editor at the bottom. You can add parentheses, an exclamation point to represent a "not" condition, an ampersand for an AND operator, and a pipes (|) for an OR operator. Figure 18-74 also shows the simple conditions picker, from which you can choose file, Registry, service, and application conditions.

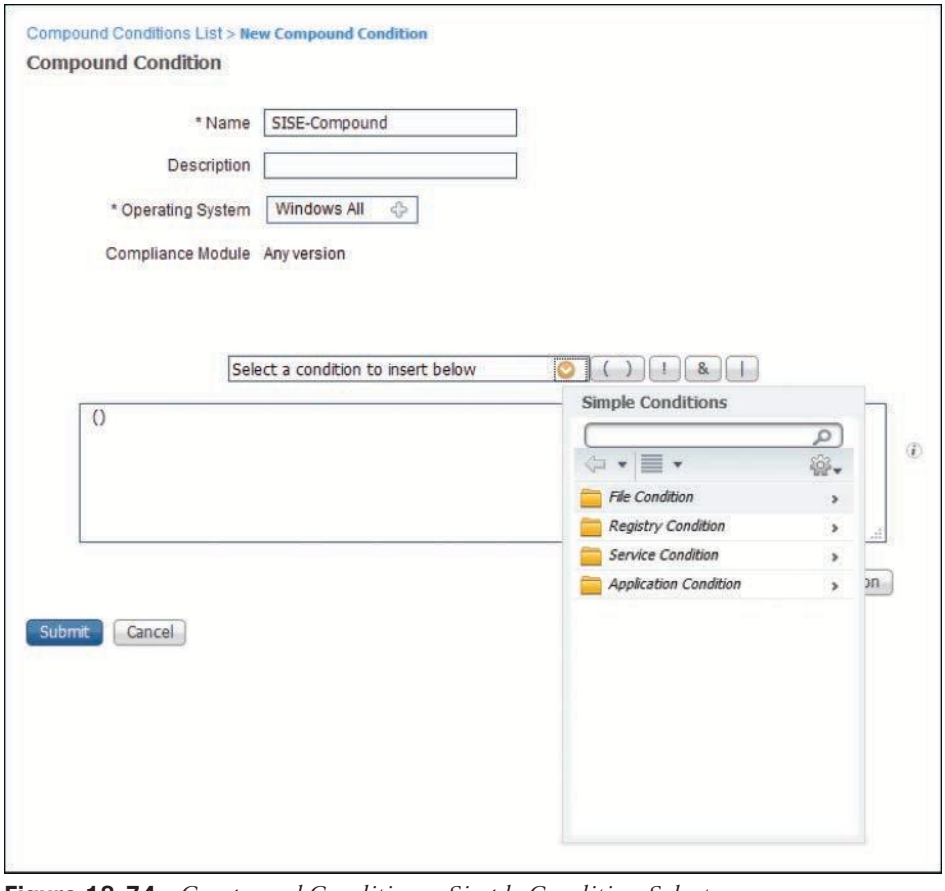


Figure 18-74 *Compound Condition – Simple Condition Selector*

Figure 18-75 shows a custom compound condition that will be true if either the file condition you created to look for watermark.xml is true or the Registry key you configured exists and the service (CiscoOrbital) is running.

You can click Validate Expression to ensure that the logic of a new compound condition works.

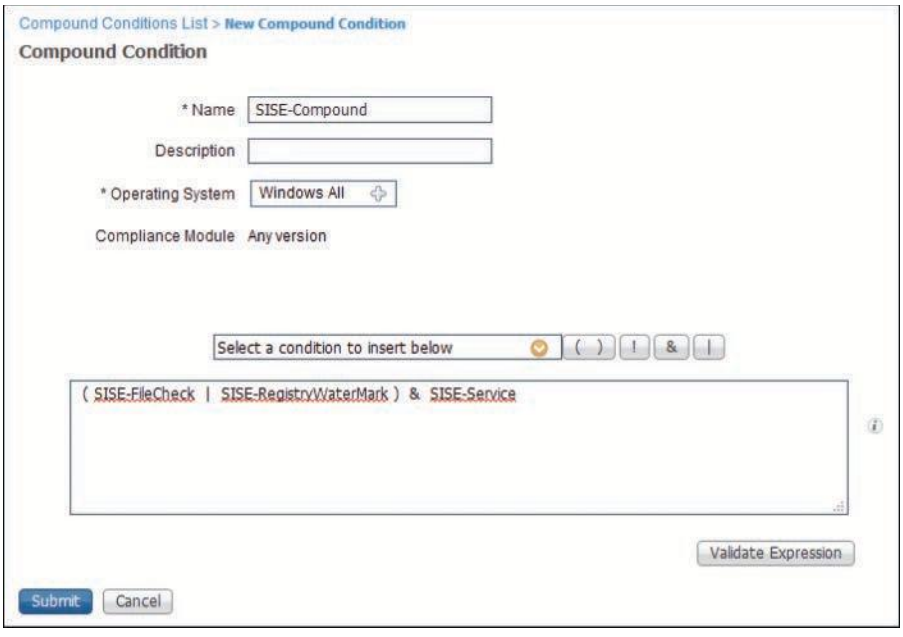


Figure 18-75 *Custom Compound Condition*

USB

The USB condition is the final condition we need to examine, and there is nothing to configure or add for this one. It is simply a preconfigured check to determine what mass storage is connected to the endpoint for use in the posture policy. Figure 18-76 shows the USB_Check condition.



Figure 18-76 *USB_Check Condition*

Remediations

It could be considered kind of humorous, but when discussing posture assessment and access control, many professionals only seem to focus on the access prevention capabilities. In other words, if an endpoint does not meet the minimum requirements, they deny access. You know what another term for this behavior is? Denial of service! If an employee is denied access to the network or some other resource that the person requires to accomplish day-to-day activities that keep the business running, that is a denial of service. At the very least, it will increase calls to your help desk, which makes the implementation more of a problem than a solution. Many NAC projects have been considered failures and ripped out of the network because the wrong person was denied access.

A much better experience is to automatically fix whatever is wrong or help the end user fix it. This is remediation.

ISE provides a number of remediation choices, each with a different use case or purpose (see Figure 18-77).



Figure 18-77 *ISE Remediations*

Application Remediations

The application of remediation is specifically for the application condition. It allows AnyConnect to uninstall the offending application or to simply kill the process to shut down the application.

Click on Application Remediations and click Add.

In the example shown in Figure 18-78, the remediation is set to Automatic, which means AnyConnect will run this action in the context of its system process. (Remember that AnyConnect does not run in the user context.) The remediation is set to uninstall the offending application, and the offending application is configured to be VMware Player Version 9.x.

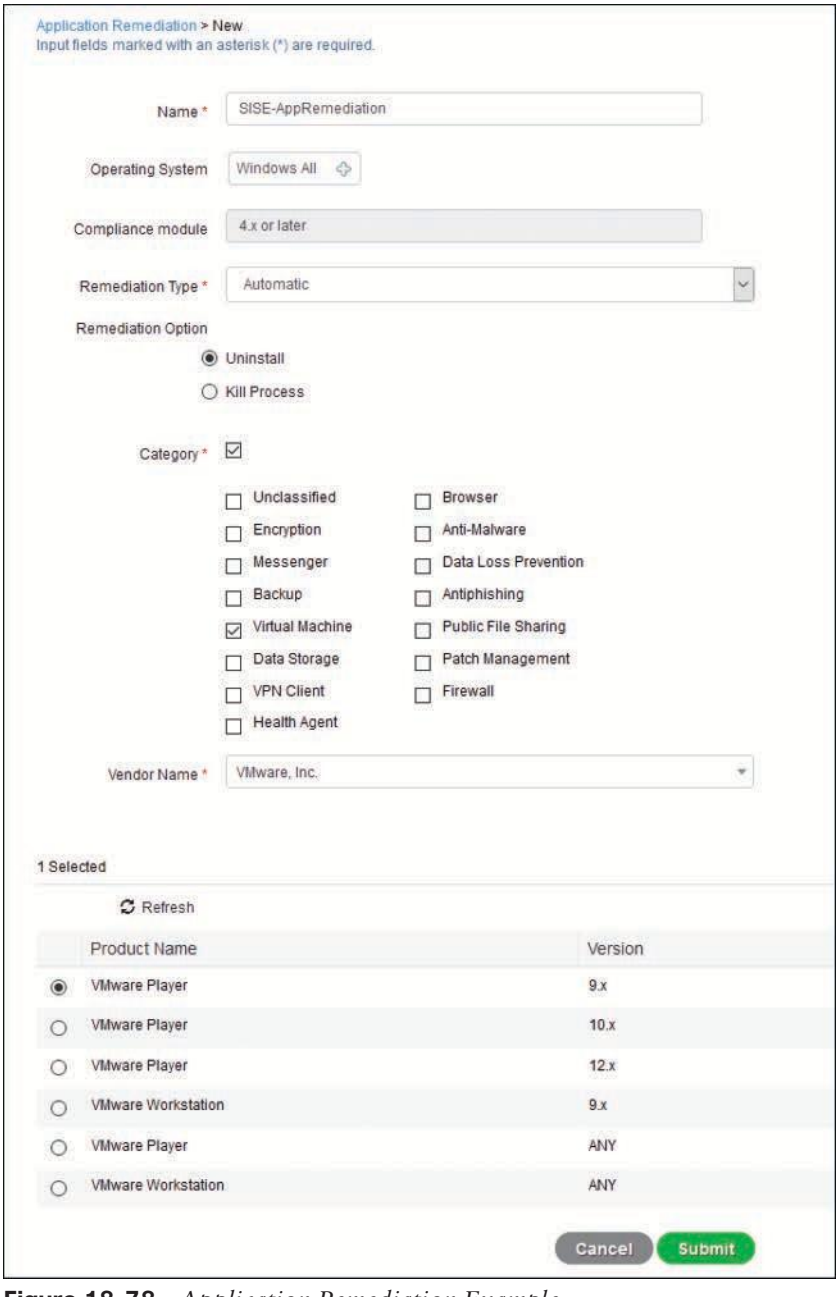


Figure 18-78 *Application Remediation Example*

Anti-Malware, Anti-Spyware, Anti-Virus Remediations

Cisco offers an anti-malware remediation that automatically launches each vendor's update process for any out-of-date products. This is a critical feature for remediation since it does not require you to manually identify the update mechanism for each of the endpoint

security suites that might exist. Each product is managed and configured by its own enterprise management application, which determines where and how to get the updates. These preconfigured remediation actions do not need to know what that configuration is; they simply have to be aware of how to execute the product's built-in update executable and allow the software configuration to handle the rest.

Figure 18-79 shows one example of a preconfigured remediation.

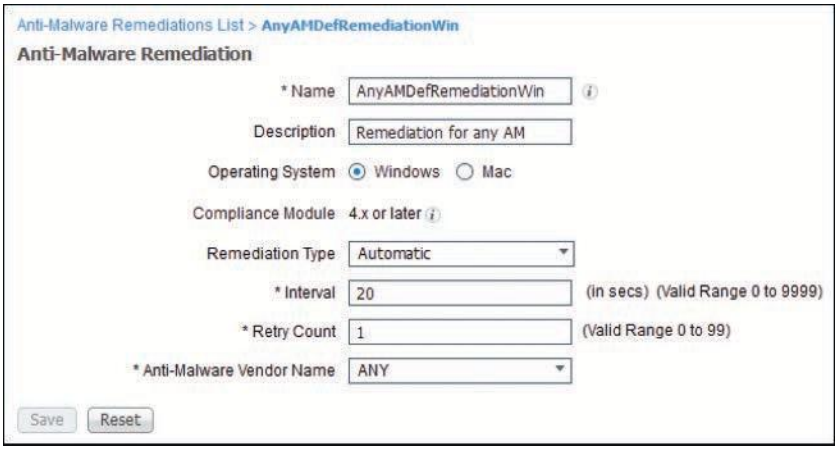


Figure 18-79 *Anti-Malware Remediation Example*

File Remediations

A file remediation allows an end user to download a required file in order to bring the endpoint up to compliant status. This file could be a watermark file or any other required file. Figure 18-80 shows an example file remediation where the file is named watermark.xml.

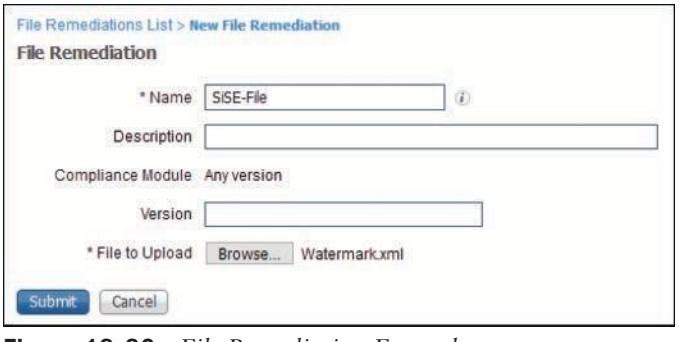


Figure 18-80 *File Remediation Example*

Firewall Remediations

A firewall remediation is designed to enable a host-based firewall when it is not enabled. Cisco provides two default remediations for firewalls: one for Windows and one for macOS. Figure 18-81 shows the default firewall remediation for macOS.

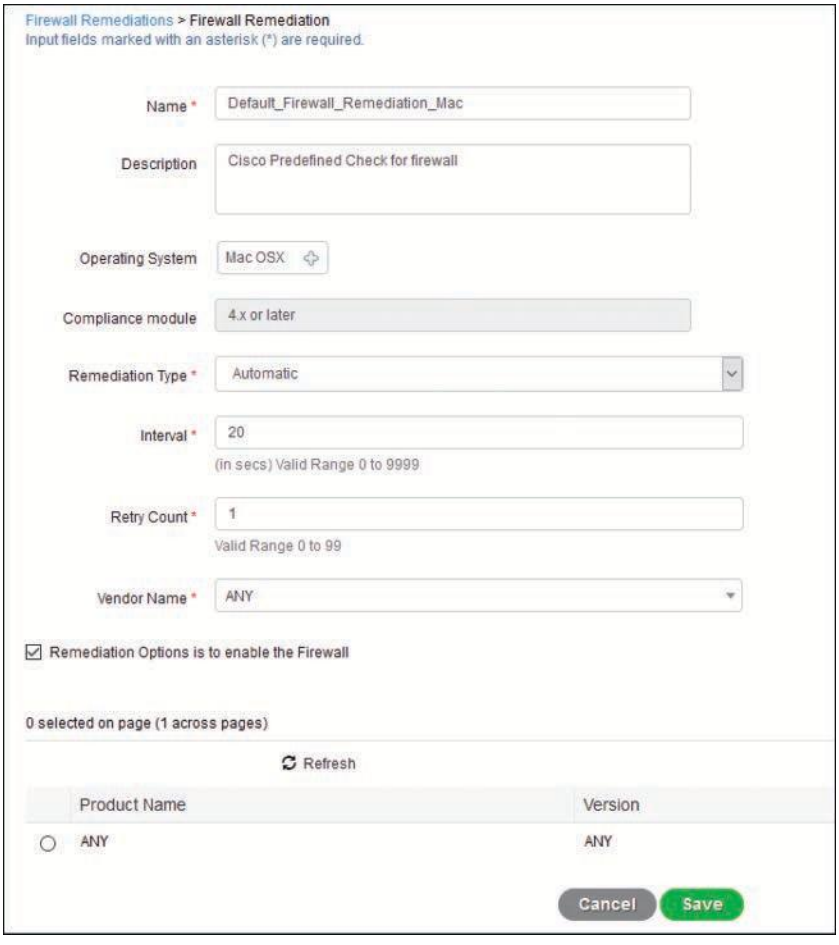


Figure 18-81 *Default Firewall Remediation*

Launch Program

Launching a specific application is a very common remediation type leveraged by large organizations all over the world. When an organization has a mature endpoint management team, it may have a number of custom tools and processes for updating systems, as well as logon scripts and other scripts to run these custom processes.

There are security implications related to allowing applications to launch scripts—and there is no remediation option for launching scripts because of those implications. However, you can convert a script into an executable, and that executable can be run with the launch program remediation. There is one catch to the use of this remediation: The executable must be signed by a trusted authority.

Figure 18-82 shows an example of a launch program remediation.

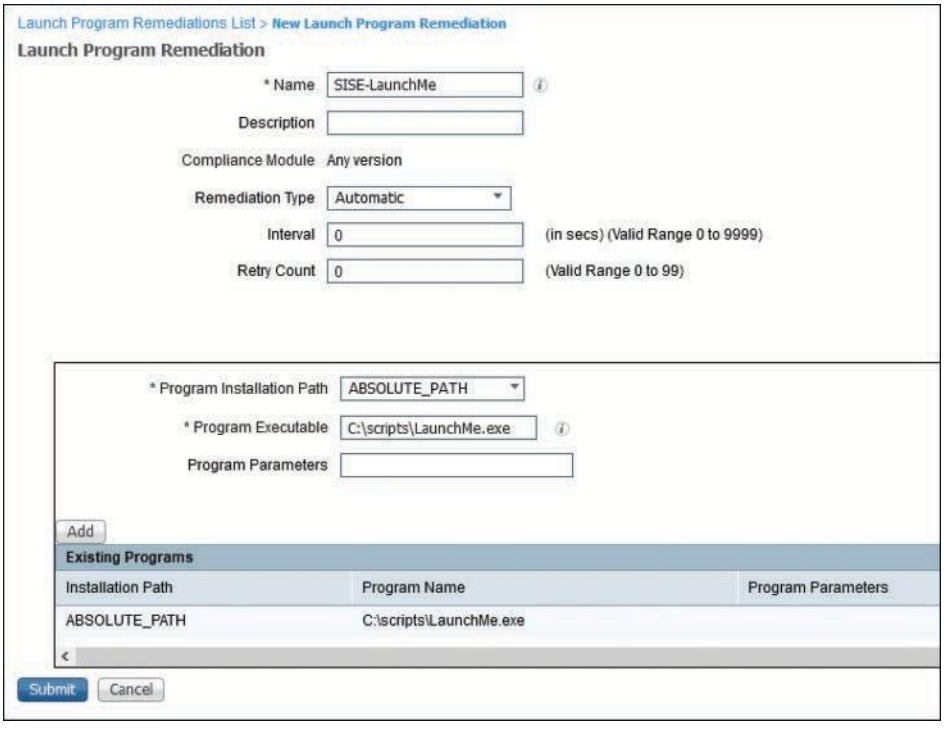


Figure 18-82 *Launch Program Remediation Example*

Link

A link remediation is used to display a clickable URL that directs the end user to a web page of your choosing. Figure 18-83 shows an example of a link remediation.

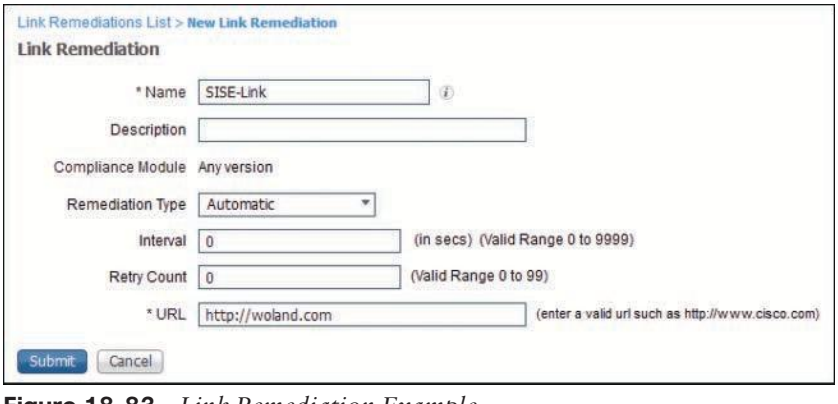


Figure 18-83 *Link Remediation Example*

Patch Management Remediations

A patch management remediation is among the most highly requested remediations from ISE customers. Integration with device management tools is valued highly—not only from a posture check perspective but from a remediation perspective.

Figure 18-84 shows an example of a patch management remediation. In this example, you can see that an action for SCCM to install missing critical patches is configured.

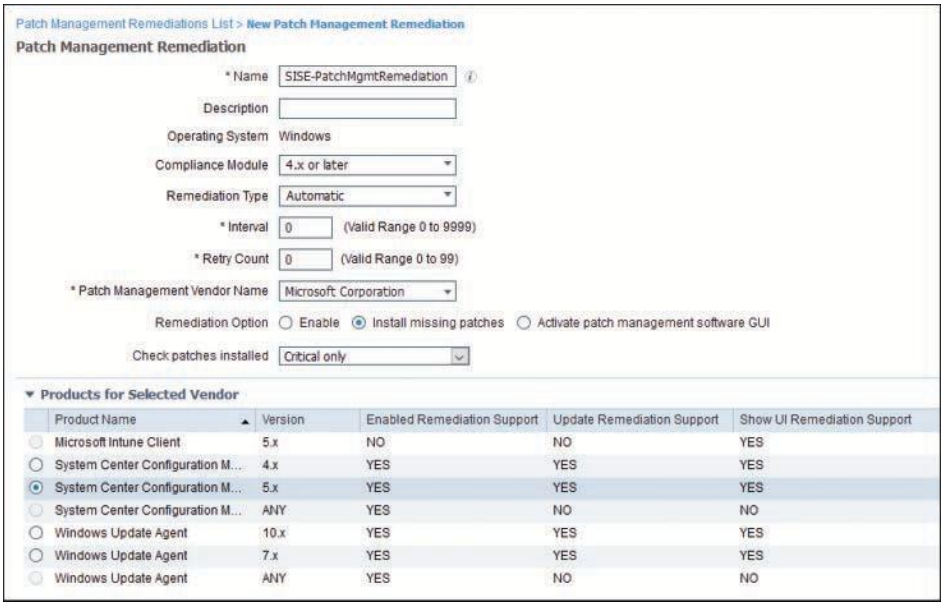


Figure 18-84 *SCCM Remediation Example*

Windows Server Update Services

Windows Server Update Services (WSUS) is an on-premises service that Windows administrators can implement on their own networks. It allows an administrator to control the distribution of all updates and hotfixes from Microsoft products to the managed Windows endpoints; without it, those endpoints must use the built-in Windows Update to retrieve software updates from Microsoft directly.

Figure 18-85 shows an example of a WSUS remediation.

Windows Update

The Windows Update remediation uses the built-in Windows service to download and install software patches and hotfixes directly from Microsoft. This remediation action also permits the changing of the Windows Update setting to overwrite the user's configured settings. Figure 18-86 shows an example of a Windows Update remediation.

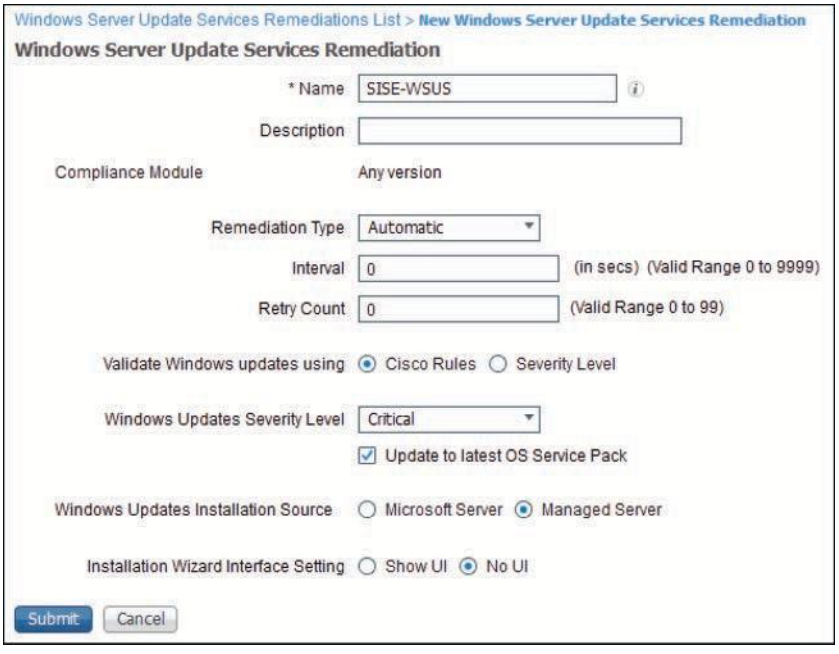


Figure 18-85 *WSUS Remediation Example*

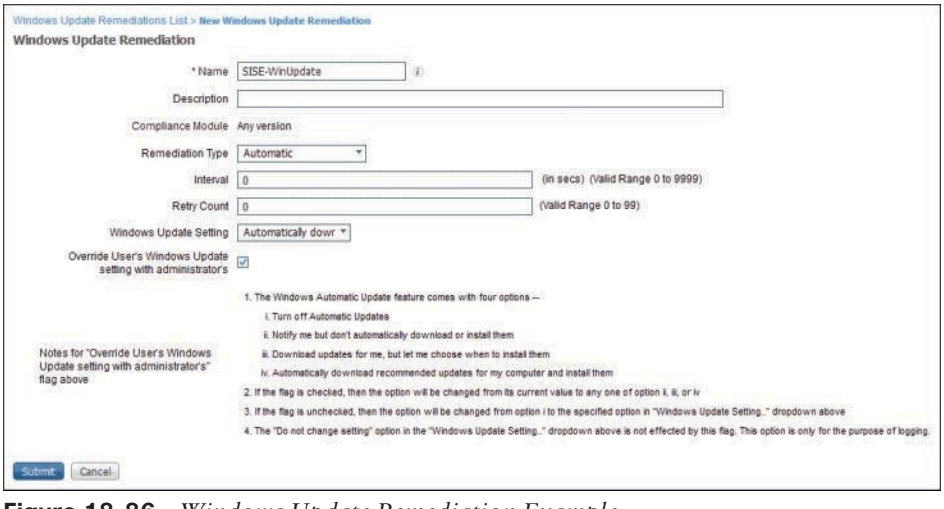


Figure 18-86 *Windows Update Remediation Example*

USB Remediations

A USB posture condition collects information about any attached USB devices, and a USB remediation is used to block the use of all USB mass storage devices. There is no granularity of control as of ISE Version 2.7 and AnyConnect Version 4.8; this remediation offers only a policy to block all USB storage usage, but it has met the needs of many customers.

Figure 18-87 shows the preconfigured USB remediation.

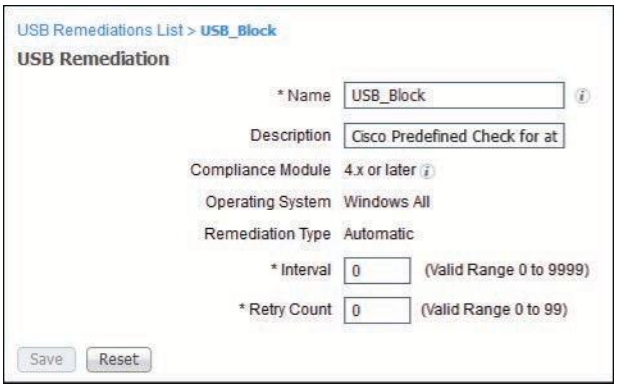


Figure 18-87 *Preconfigured USB Remediation*

Requirements

You have now seen posture conditions and remediations. Both of those policy elements are combined into ***posture requirements***, which are used in a posture policy.

NOTE If this section gets confusing, don't worry. You are not alone! Refer to Figure 18-33 any time you start to feel lost.

A requirement is a form of policy in which the remediation is the result of a set of conditions. As shown in Figure 18-88, the elements are listed in the Requirements screen in order of dependence, from left to right, meaning you need to select each policy element in that order:

- **Operating system:** Windows or All or Mac OSX
- **Compliance module:** 3.x, 4.x, or Any
- **Posture type:** AnyConnect, AnyConnect Stealth, or Temporal Agent
- **Conditions:** Compliance modules and agents (which become available after you select the OS)
- **Remediation actions:** Remediations that become available for selection after all the other conditions have been chosen

Figure 18-88 shows the default list of requirements. (To see this list more clearly, examine it within your own user interface, if possible.)

Unlike with most other areas of configuration in the ISE user interface, there is no Add button for creating a new requirement. You must either duplicate an existing requirement or insert a new one.

For the posture example we are following in this chapter, this section shows how to insert two new rules that each leverages a condition and a remediation created for Windows. (To get practice, follow along and create a new rule or two for the conditions that you built.)

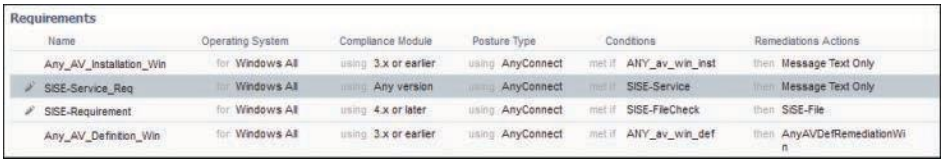


Requirements

Name	Operating System	Compliance Module	Posture Type	Conditions	Remediations Actions
Any_AV_Installation_Win	for Windows All	using 3.x or earlier	using AnyConnect	met if ANY_av_win_inst	then Message Text Only
Any_AV_Definition_Win	for Windows All	using 3.x or earlier	using AnyConnect	met if ANY_av_win_def	then AnyAVDefRemediationWin
Any_AS_Installation_Win	for Windows All	using 3.x or earlier	using AnyConnect	met if ANY_as_win_inst	then Message Text Only
Any_AS_Definition_Win	for Windows All	using 3.x or earlier	using AnyConnect	met if ANY_as_win_def	then AnyASDefRemediationWin
Any_AV_Installation_Mac	for Mac OSX	using 3.x or earlier	using AnyConnect	met if ANY_av_mac_inst	then Message Text Only
Any_AV_Definition_Mac	for Mac OSX	using 3.x or earlier	using AnyConnect	met if ANY_av_mac_def	then AnyAVDefRemediationMac
Any_AS_Installation_Mac	for Mac OSX	using 3.x or earlier	using AnyConnect	met if ANY_as_mac_inst	then Message Text Only
Any_AS_Definition_Mac	for Mac OSX	using 3.x or earlier	using AnyConnect	met if ANY_as_mac_def	then AnyASDefRemediationMac
Any_AM_Installation_Win	for Windows All	using 4.x or later	using AnyConnect	met if ANY_am_win_inst	then Message Text Only
Any_AM_Definition_Win	for Windows All	using 4.x or later	using AnyConnect	met if ANY_am_win_def	then AnyAMDefRemediationWin
Any_AM_Installation_Mac	for Mac OSX	using 4.x or later	using AnyConnect	met if ANY_am_mac_inst	then Message Text Only
Any_AM_Definition_Mac	for Mac OSX	using 4.x or later	using AnyConnect	met if ANY_am_mac_def	then AnyAMDefRemediationMac
USB_Block	for Windows All	using 4.x or later	using AnyConnect	met if USB_Check	then USB_Block
Default_AppVis_Requirement_Win	for Windows All	using 4.x or later	using AnyConnect	met if Default_AppVis_Condition_Win	then Select Remediations
Default_AppVis_Requirement_Mac	for Mac OSX	using 4.x or later	using AnyConnect	met if Default_AppVis_Condition_Mac	then Select Remediations
Default_Hardware_Attributes_Requirement_Win	for Windows All	using 4.x or later	using AnyConnect	met if Hardware_Attributes_Check	then Select Remediations
Default_Hardware_Attributes_Re	for Mac OSX	using 4.x or later	using AnyConnect	met if Hardware_Attributes_Ch	then Select Remediations

Figure 18-88 *Preconfigured Posture Requirements*

Figure 18-89 shows two inserted rules. The first rule has Compliance Module set to Any Version and Posture Type set to AnyConnect, it is checking for the Orbital service to be running (SISE-Service), and Remediation Action is set to Message Text Only. The second rule specifies the 4.x or later compliance module with AnyConnect and looks for the existence of the file specified in the SISE-FileCheck condition; it applies the file remediation created earlier in the chapter so the end user can download it.



Requirements

Name	Operating System	Compliance Module	Posture Type	Conditions	Remediations Actions
Any_AV_Installation_Win	for Windows All	using 3.x or earlier	using AnyConnect	met if ANY_av_win_inst	then Message Text Only
SISE-Service_Req	for Windows All	using Any version	using AnyConnect	met if SISE-Service	then Message Text Only
SISE-Requirement	for Windows All	using 4.x or later	using AnyConnect	met if SISE-FileCheck	then SISE-File
Any_AV_Definition_Win	for Windows All	using 3.x or earlier	using AnyConnect	met if ANY_av_win_def	then AnyAVDefRemediationWin

Figure 18-89 *Two Added Requirements*

NOTE As you build a requirement, it is possible to combine more than one condition together, but if you do, for each combination you see only one remediation option. The recommended approach is to use compound conditions when you plan to use more than one condition in a requirement. This way, you are still maintaining the ratio of one condition per remediation in the requirements table unless the conditions share the exact same remediation, such as using a patch management system to rectify the situation.

Configure Posture Policies

Back in the day of NAC Appliance, there were configurations referred to as "roles." The roles were based on user attributes, such as AD group membership. The important point to know is that a role defined which posture requirements were applied to members of that role.

Key Topic

With ISE, Cisco likes to refer to everything as a policy, and the posture policies perform the function that NAC Appliance used roles to perform. The posture policies determine which requirements are mandatory and for whom. It is critical to understand that all the clients connecting to your network must meet mandatory requirements during posture evaluation to become posture compliant.

Key Topic

It is also very important to understand that posture is just another data point when it comes to the authorization policy. A device may be noncompliant for posture all day long, but if the authorization policy is not using the posture attribute as a condition, then there is nothing to prevent the access at all.

If you navigate to the Posture Policy tab in the Posture work center, you see the default posture policy, as shown in Figure 18-90.



Figure 18-90 *Default Posture Policy*

Notice that in the default posture policy, all the rules are disabled by default. These preexisting policies are there to leverage the preexisting conditions, remediations, and requirements that ISE downloads with posture updates from Cisco.com. Yes, it is a lot to look at, and many people see this screen and just turn and run. However, it really does make life easier to have all these prebuilt options, as you can simply enable a few of them and get a tremendous amount of value without having to configure anything.

This section shows how to enable the default hardware attributes and application visibility policies for Windows and macOS, using both AnyConnect and temporal agents. In addition, you will see how to create two new policy rules for the custom file and Registry requirements created earlier. All the posture policies will apply to members of the Employees AD group only. Figure 18-91 shows this posture policy.

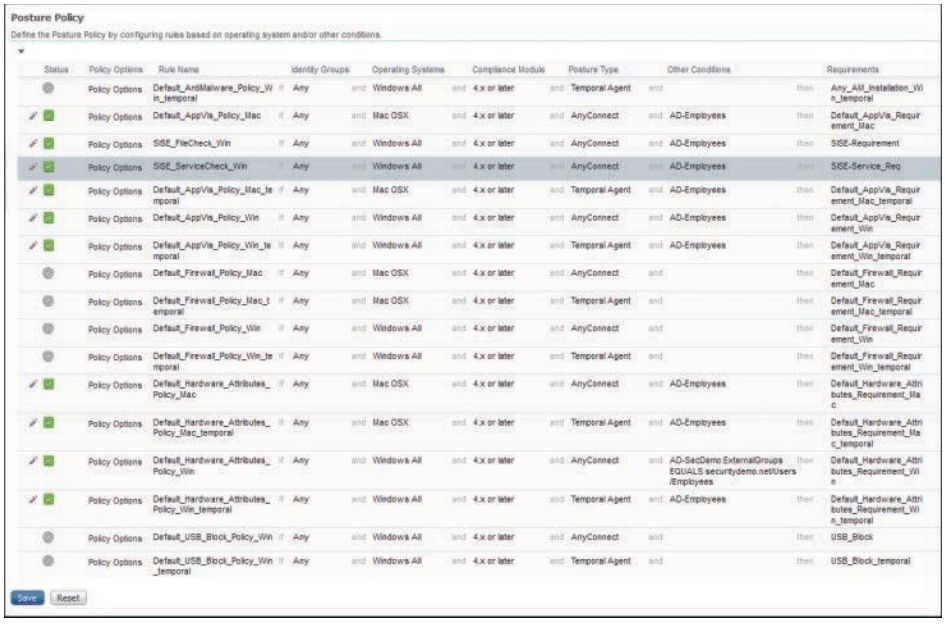


Figure 18-91 *Posture Policy with Enabled Rules*

Other Important Posture Settings

It is important to understand a number of settings in the Posture work center. In this section, we look at the most important ones.

Let's start with the Posture General Settings section, which is shown in Figure 18-92.

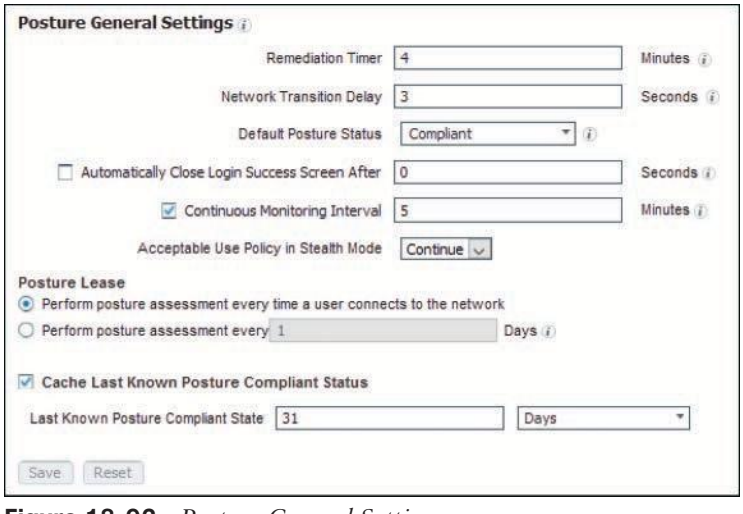


Figure 18-92 *Posture General Settings*

The important settings in the Posture General Settings section are as follows:

- **Remediation Timer:** This setting defines the amount of time a client has to correct a failed posture condition. There is also a remediation timer in the AnyConnect configuration; this timer is for ISE, not AnyConnect.
- **Default Posture Status:** This setting provides the posture status for devices without the posture agent or operating systems that cannot run the temporal agent, such as Linux-based operating systems.
- **Continuous Monitoring Interval:** This setting applies to the application and hardware conditions that are taking inventory of the endpoint. The setting specifies how often AnyConnect should send the monitoring data.
- **Acceptable Use Policy in Stealth Mode:** The only two choices for this setting are to block or continue. Block prevents stealth mode AnyConnect clients from proceeding if the AUP has not been acknowledged. Continue allows the stealth mode client to proceed even without acknowledging the AUP (which is often the intent when using the stealth mode setting of AnyConnect).

Posture Lease

Originally ISE performed a posture assessment every time an endpoint connected to the network. Some customers did not like that end-user experience and asked Cisco to speed up the login process. To meet that customer requirement, ISE added the posture lease concept, which allows an administrator to configure ISE to perform posture assessment in specified intervals for endpoints using the AnyConnect agent for posture assessment.

When the posture lease is active, ISE uses the last known posture state and does not reach out to the endpoint to check for compliance.

When the posture lease expires, the endpoint stays in the same compliance state until the endpoint reauthenticates, at which point the posture is reassessed. If the endpoint is compliant, the lease starts over.

Cache Last Known Posture Compliant Status

The Cache Last Known Posture Compliance setting has nothing to do with the posture assessment or authorization of an endpoint. This setting is used for context visibility. The endpoint record in the Context Visibility dashboard shows the last known compliance state of the endpoint until the timer for this setting expires.

Reassessment Configurations

Posture reassessments are a critical component for the posture workflow. You saw how to configure the AnyConnect agent for posture reassessment in the "Posture Protocol" section, earlier in this chapter. The agent periodically checks in with the PSNs defined based on the timer in that configuration.

When a request reaches the PSN, the PSN determines whether a posture reassessment is needed, based on the ISE configuration for that endpoint's role. If the client passes the reassessment, the PSN maintains the endpoint's posture-compliant state, and the posture lease is reset. If the endpoint fails the reassessment, the posture status changes to noncompliant, and any posture lease that existed is removed.

NOTE Periodic posture reassessment can be done only for clients that are already successfully postured for compliance. Posture reassessment cannot occur if clients are not compliant on the network.

Figure 18-93 shows a posture reassessment configuration being added. There are some strict rules that go along with the configuration. For starters, multiple posture reassessments can be added, but each one must have a unique group or combination of groups. If you use the group Any, as Figure 18-93 shows, then no other posture reassessment configuration can be added or used.

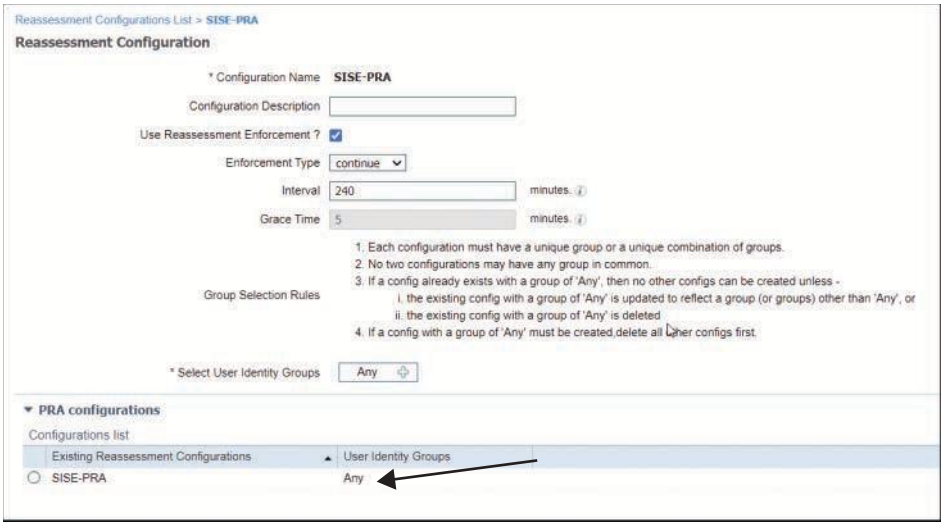


Figure 18-93 *Posture Reassessment*

In the posture reassessment configuration in Figure 18-93, notice the Use Reassessment Enforcement? checkbox. If this checkbox is not enabled, the posture data for an endpoint updates in ISE, but no action is taken until the next full access control event. When you enable the check box, the Enforcement Type dropdown comes alive, offering three choices:

- **Continue**: The user's session continues unaltered, regardless of the posture assessment.
- **Logoff**: If the posture assessment result is not compliant, the user is forced to log off from the network. When the user logs in again, the compliance status is set to unknown, and a full posture assessment takes place.
- **Remediate**: If the endpoint is not compliant, the agent waits for a specified time for the remediation to happen. Once the client has remediated, the agent sends the PRA report to the PSN. If the remediation is ignored on the client, the agent sends a logoff request to the PSN to force the client to log off from the network.

Authorization Rules

The authorization rules for posture assessment are much like the ones you create for guest access and BYOD. There must be a rule that redirects endpoints that do not have a posture-compliant state to the active PSN, and there must be another rule for permitting the endpoint on to the network if it is posture compliant.

The redirection is to the client provisioning portal; this is incredibly important to this process. If posture is not applicable to an endpoint, the endpoint is redirected to this portal and marked as not applicable. If posture is applicable but the agent is not yet on the endpoint, the CPP launches the network provisioning wizard on the endpoint to install the agent. Finally, if the client has the agent, the redirection sends the posture probes to the active PSN so the posture assessment can commence.

Create an Authorization Profile for Redirection

It should come as no surprise that you need an authorization profile to permit limited access and redirect the endpoint to the Client Provisioning portal. To set up an authorization profile, navigate to the Posture Elements tab in the Posture work center and click on Authorization Profiles.

As shown in Figure 19-94, you need to add a new authorization profile, include a dACL (in case you need it for wired access), select the Web Redirection checkbox, and select Client Provisioning (Posture). Use the ACL-WEBAUTH-REDIRECT ACL and select the default provisioning portal.

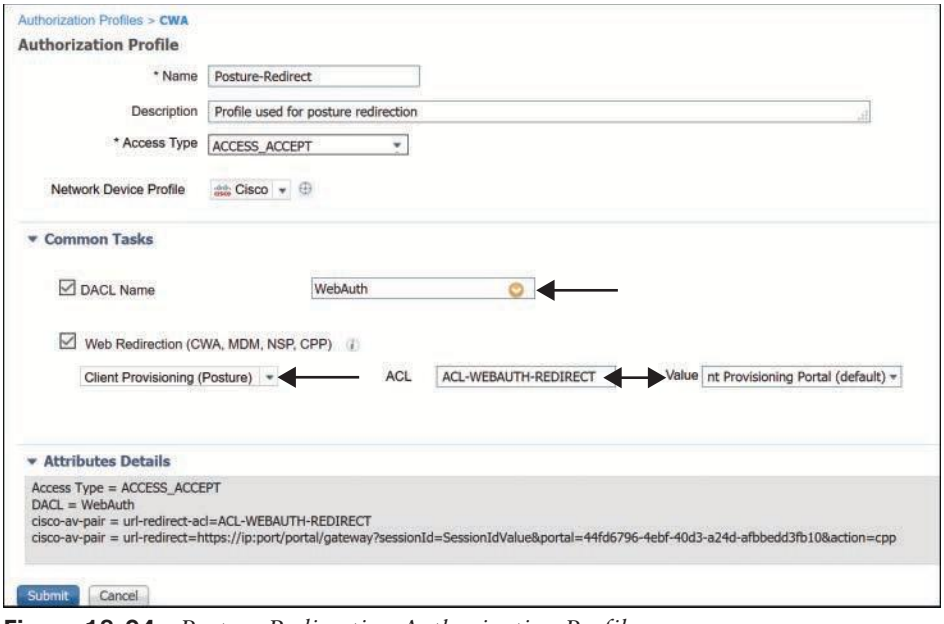


Figure 18-94 *Posture Redirection Authorization Profile*

Create the Authorization Rules

You can create your own authorization rules or modify prebuilt rules. In this section you will see how to do the latter so you can see the benefit of the prebuilt authorization rules.

There are three preconfigured authorization rules for posture:

- The first is configured to match when an authentication succeeds and a device's compliance is unknown.
- The second rule matches successful authentications with noncompliant endpoints.

NOTE Both of the first two rules have the same result, which is to use a preconfigured authorization profile that redirects the endpoint to the Client Provisioning portal.

- The final rule matches on successful authentication and posture-compliant endpoints and uses the prebuilt PermitAccess authorization profile.

Now, you might be thinking that the first two rules could be combined with an OR operator—and you would be correct. The conditions section of the combined rule would look similar to the following logical statement: Authentication Passed AND (Posture = unknown OR Posture = noncompliant).

Figure 18-95 shows the three preconfigured rules enabled, with a new condition added to each one to check for the user to be a member of the Employees group in AD. The noncompliant and unknown rules have the Posture-Redirect authorization profile created earlier; this profile uses the redirection ACL created in Chapter 11.



Figure 18-95 *Modified Default Authorization Rules for Posture*

Key Topic

The order of the authorization rules is very important. The rule table is top-down and first-match, so you should have the most specific rules toward the top. Figure 18-96 shows the order for your reference. Notice that the posture rules appear after the web authentication rules.



Figure 18-96 *Authorization Rule Order*

The Endpoint Experience

The purpose of this section of the chapter is to help you understand the full experience of posture from perspective of the most important people in the entire ISE ecosystem: the end users. Additionally, it is very important for the SISE exam that you understand the full experience from the end-user and endpoint perspective, not just from ISE's point of view.

Regardless of whether you're using a wired or wireless network—or even a remote-access VPN—the client connects to the network and receives the temporary access that redirects the endpoint, as shown in Example 18-1.

Example 18-1 *Showing the Redirected State*

```
3560-X# show authentication session int g0/23
            Interface:  GigabitEthernet0/23
          MAC Address:  0050.56a1.3e7f
           IP Address:  10.1.41.102
            User-Name:  associate1
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
        Authorized By:  Authentication Server
           Vlan Group:  N/A
              ACS ACL:  xACSACLx-IP-WebAuth-5e2a155e
     URL Redirect ACL:  ACL-WEBAUTH-REDIRECT
         URL Redirect:  https://atw-ise244.securitydemo.net:8443/portal/
gateway?sessionId=0A01283C000000056ABF55B4&portal=44fd6796-4ebf-40d3-a24d-
afbbedd3fb10&action=cpp&token=b7fb1e5cfdc2d62894b860c0619c1474
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  0A01283C000000056ABF55B4
      Acct Session ID:  0x00000011
               Handle:  0x6C000005


Runnable methods list:
       Method   State
       mab      Not run
       dot1x    Authc Success
```

Scenario 1: AnyConnect Not Installed on Endpoint Yet

In this scenario, a client does not have AnyConnect yet. The user opens a web browser and is redirected to the Client Provisioning portal. Because there is no agent installed, no posture discovery messages have been sent to the PSN, and the end user sees the message *Your computer requires security software to be installed before you can connect to the network*, as shown in Figure 18-97. The end user clicks Start.

The ISE PSN is waiting to hear from the AnyConnect System Scan module, and it still has not heard from it. The user clicks This Is My First Time Here, as shown in Figure 18-98. (Of course, all this text is configurable in the Client Provisioning portal, so you can make it more user friendly for your end users.)

The This Is My First Time Here section expands. It includes a link to download and install AnyConnect, as shown in Figure 18-99.

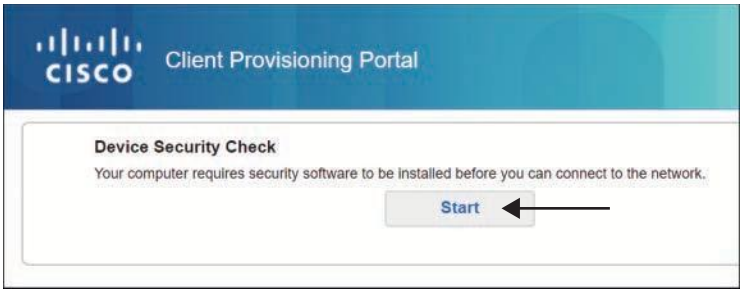


Figure 18-97 *Client Provisioning Portal with No Agent Installed*



Figure 18-98 *Client Provisioning Portal: This Is My First Time Here*

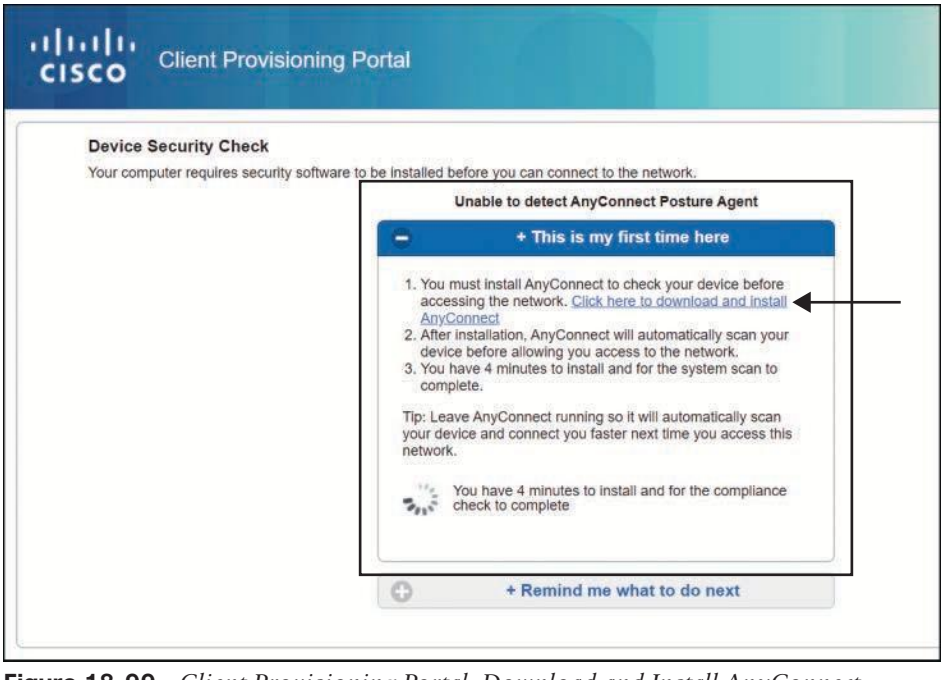


Figure 18-99 *Client Provisioning Portal: Download and Install AnyConnect*

The end user clicks this link, and the Network Setup Assistant (NSA) is launched. (This is the same NSA that you used for native supplicant provisioning in Chapter 16.) The NSA downloads the AnyConnect software that you uploaded to ISE along with the AnyConnect configurations you published there, and then AnyConnect installs the posture module, as shown in Figures 18-100 through 18-102.

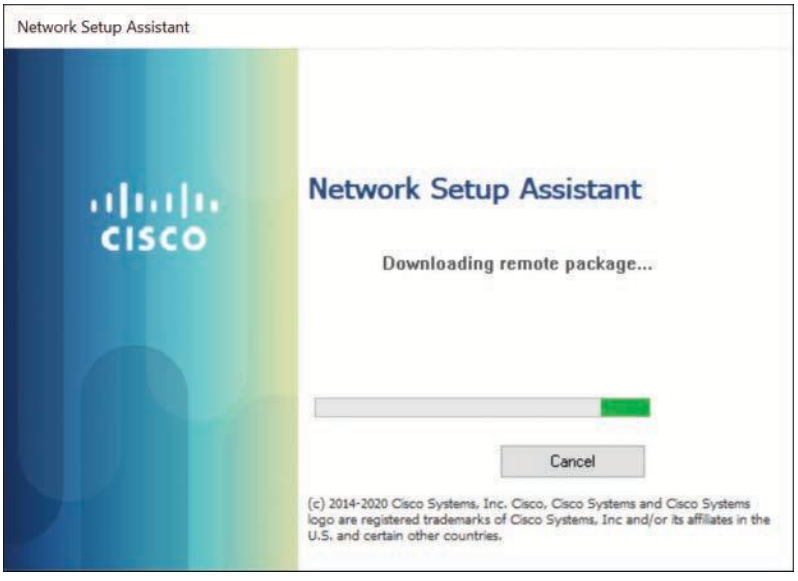


Figure 18-100 *Network Setup Assistant Downloading AnyConnect from ISE, Part 1*

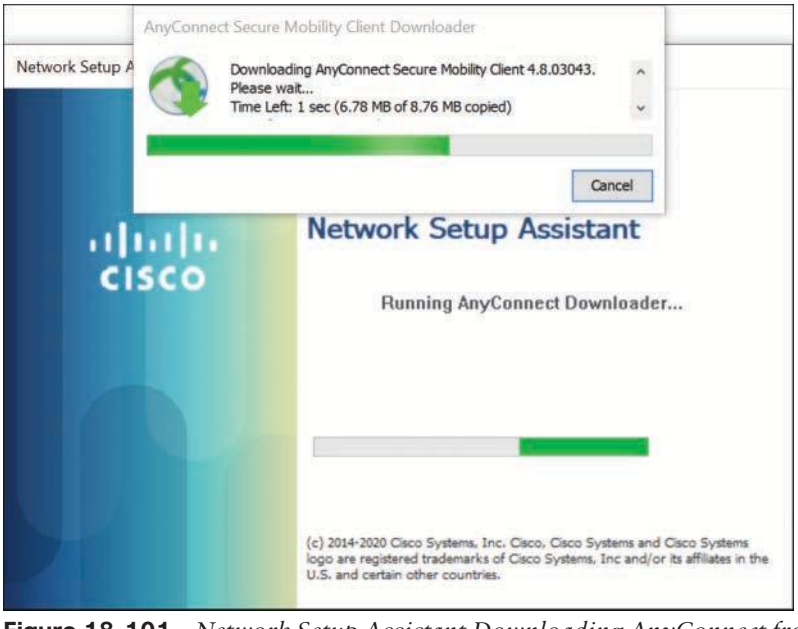


Figure 18-101 *Network Setup Assistant Downloading AnyConnect from ISE, Part 2*

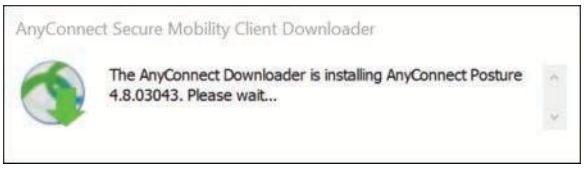


Figure 18-102 *AnyConnect Installing the Posture Module*

After the posture module is installed, the discovery messages go out to the active ISE PSN, and posture assessment is completed, as shown in Figures 18-103 and 18-104.

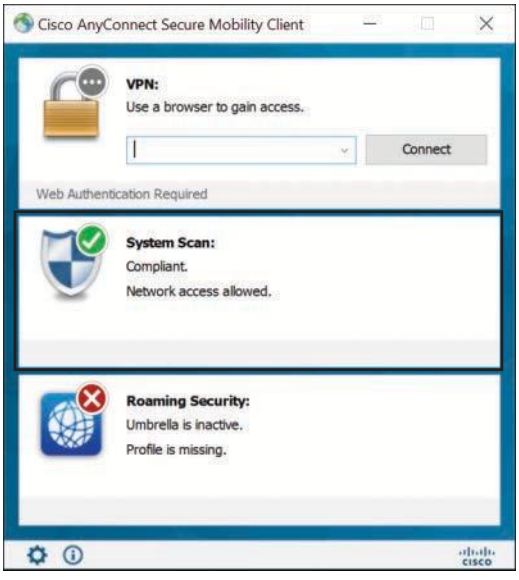


Figure 18-103 *Posture Is Compliant and Network Access Is Allowed*

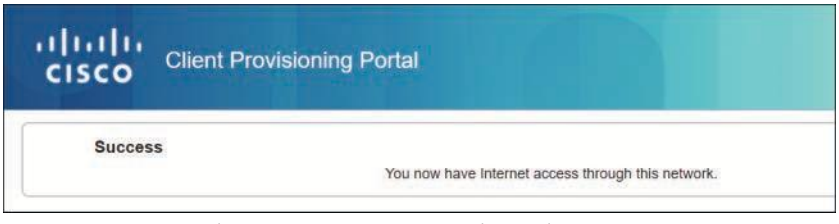


Figure 18-104 *Client Provisioning Portal Displaying a Success Message*

Let's look more closely at the discovery messages that are sent from AnyConnect. An endpoint session is created after the endpoint passes 802.1X authentication. The client agent then attempts to connect to the active ISE PSN by sending discovery packets using different methods, in the following order:

1. Via HTTP to port 80 on a Cisco ISE server (if configured)
2. Via HTTPS to port 8905 on a Cisco ISE server (if configured)

3. Via HTTP to port 80 on the default gateway
4. Via HTTPS to port 8905 to each previously contacted server
5. Via HTTP to port 80 to the FQDN enroll.cisco.com

The posture assessment begins when the AUP is accepted (if there is one). The Cisco ISE node issues a posture token to AnyConnect, and that token is used for the posture lease.

Scenario 2: AnyConnect Already Installed, Endpoint Not Compliant

This scenario considers the experience of a client that does have AnyConnect with the posture module installed already.

The 802.1X authentication has occurred, and the endpoint has temporary authorization with redirection to the Client Provisioning portal. At that point, the System Scan module starts the posture discovery process, and the posture process begins, as shown in Figure 18-105.



Figure 18-105 *Client Provisioning Portal*

The endpoint fails the compliance check, as the watermark file is missing, and the end user is given the four-minute default timer to finish the remediation. The remediation in this case is to click Start, as shown in Figure 18-106.

Once this posture requirement is passed, the next requirement is started. In this case, it is the first-time collection of hardware and application data, as shown in Figure 18-107.

After all the posture requirements are met, the network settings are refreshed to support a change of authorization, as shown in Figure 18-108.

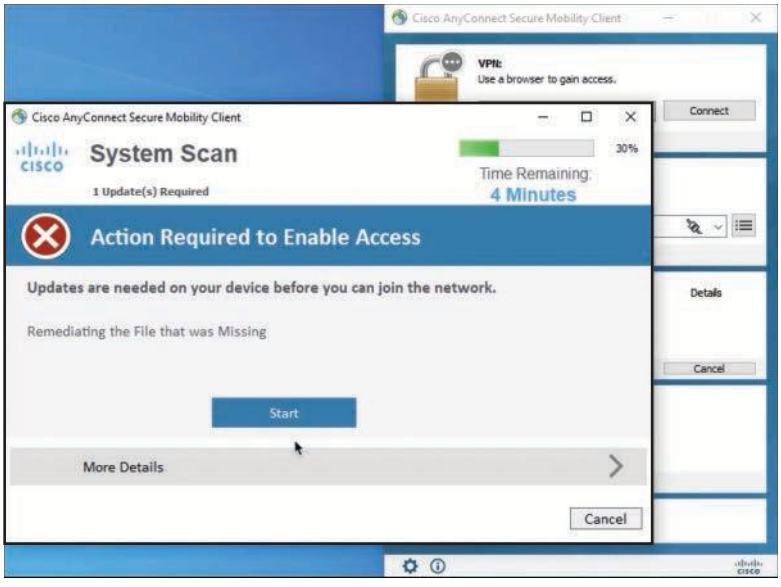


Figure 18-106 *Posture Requirement Failed, Start Remediation*



Figure 18-107 *Posture Gathering First-Time System Information*



Figure 18-108 *Updating Network Settings*

Finally, the endpoint is fully posture compliant and has received its posture token. The System Scan module displays the compliant status and shows that network access is allowed. A message pops up from AnyConnect, informing the end user that network access is allowed (see Figure 18-109).



Figure 18-109 *Endpoint Is Compliant and Network Access Is Allowed*

Scenario 3: Stealth Mode

Stealth mode is described earlier in this chapter, in the "Agent Behavior" section. Ultimately, AnyConnect can operate in clientless (stealth) mode or standard mode. When stealth mode is enabled, it runs as a service without any user interface at all. If it is the only AnyConnect module installed, the icon does not even appear in the system tray (Windows) or menu bar (macOS). The administrator even has the ability to prevent any messages from being displayed to the end user.

Figure 18-110 shows the AnyConnect user interface when stealth mode is enabled, and you can see that the System Scan module disappears. If there were no other modules installed, there would be no UI at all.

Because stealth mode is still AnyConnect, it has the full functionality of AnyConnect for posture, including posture leases, automatic remediation, posture discovery, and posture reassessment.



Figure 18-110 *AnyConnect in Stealth Mode for the Posture Module*

Figure 18-111 shows the Live Sessions screen, and Figure 18-112 shows the detailed posture report from the Posture work center. Both figures highlight the posture-compliant state of the endpoint, indicating that the endpoint still has the AnyConnect posture module.

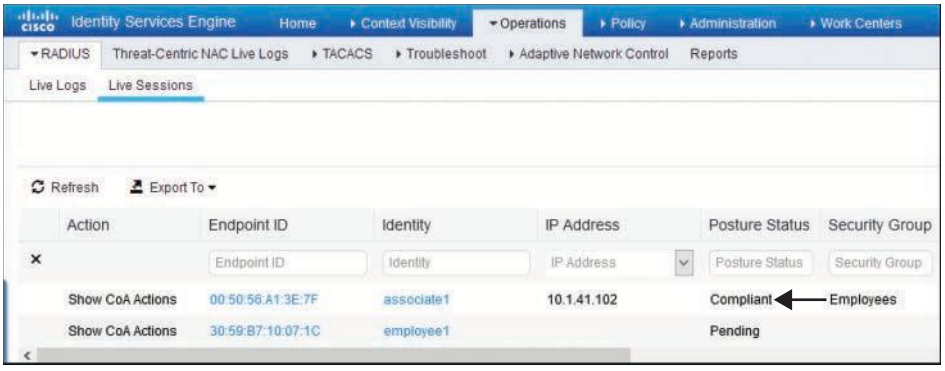


Figure 18-111 *Live Sessions View*

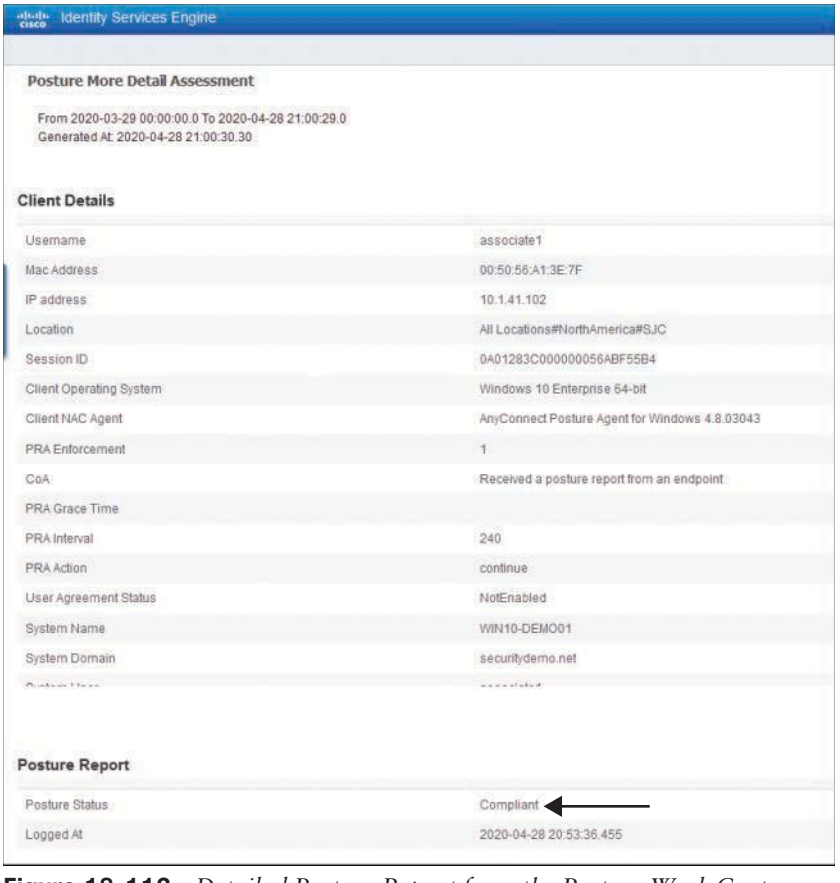


Figure 18-112 *Detailed Posture Report from the Posture Work Center*

Scenario 4: Temporal Agent and Posture Compliant

In this scenario, a contractor needs network access. The company's InfoSec team requires that all devices meet the endpoint compliance requirements. However, because the endpoint is not under the control of the company, forcing the AnyConnect agent to be installed for posture assessment may not be feasible. Another possibility is that the endpoint may lack the permissions to install a full agent. The temporal agent runs in the user space in the context of that user, and it dissolves away when the user closes the browser.

The user connects to the network and is redirected to the Client Provisioning portal because the endpoint posture is unknown. Figure 18-113 shows the contractor's browser redirected to the Client Provisioning portal, where the contractor clicks Start.



Figure 18-113 *Contractor's Browser Redirected to CPP*

The CPP tries to identify whether the temporal agent is already running on the endpoint, as shown in Figure 18-114.

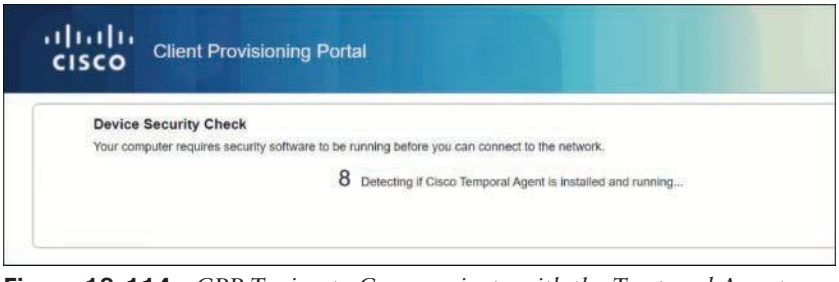


Figure 18-114 *CPP Trying to Communicate with the Temporal Agent*

Next, the contractor clicks the link to download and run the temporal agent, as shown in Figure 18-115.

The temporal agent runs and performs the posture assessment on the endpoint. After the agent determines that the posture is compliant, the correct network access is provided, as shown in Figures 18-116 and 18-117.

It is important to note that while the compliance modules are the same between AnyConnect and the temporal agent, automatic remediations do not apply because the temporal agent is running in the browser's memory space, in the context of the user.

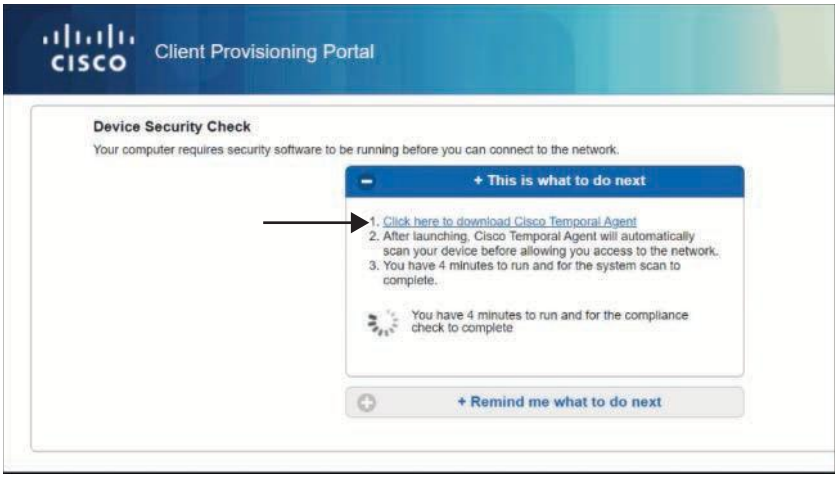


Figure 18-115 *Link to Download the Temporal Agent*

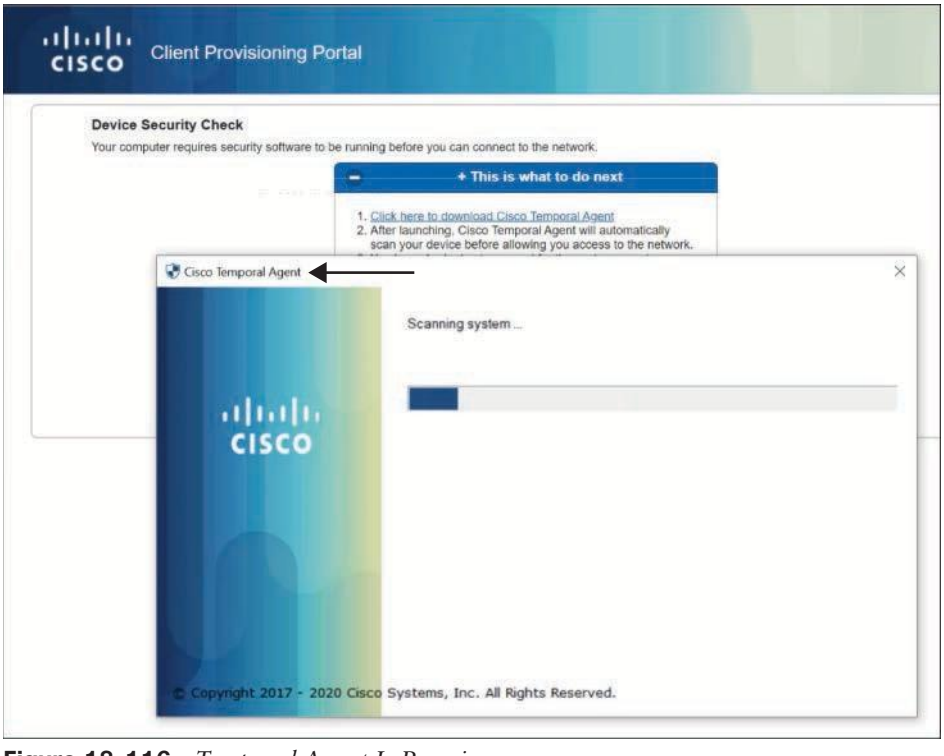


Figure 18-116 *Temporal Agent Is Running*

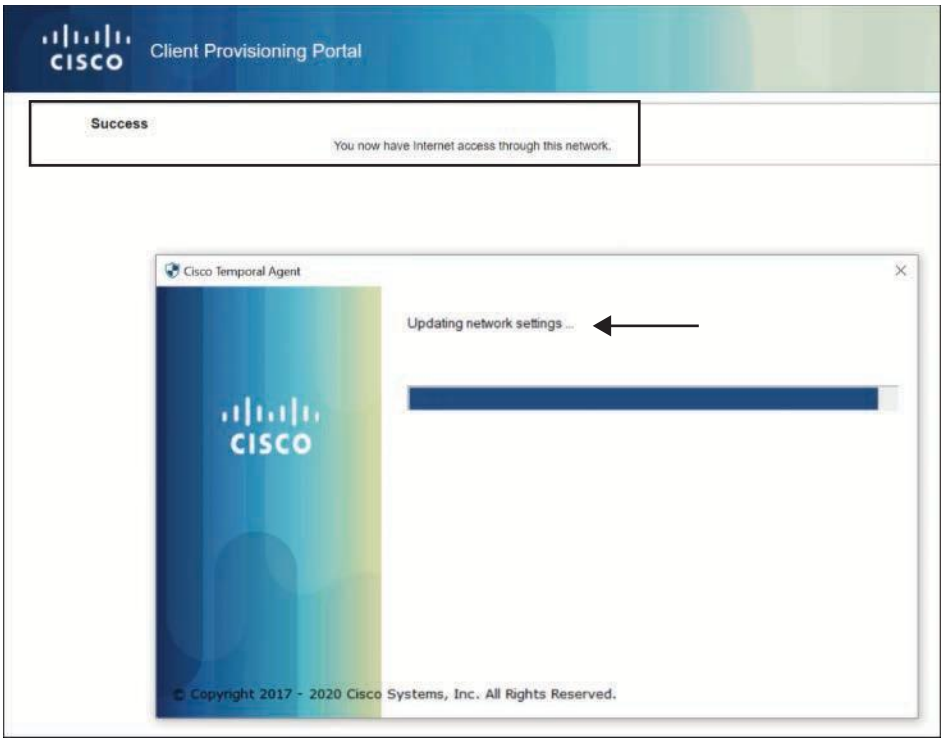


Figure 18-117 *Posture Is Compliant and Network Access Is Granted*

Mobile Posture

Everything you have seen so far in the Posture work center revolves around two operating systems: Windows and macOS. There are two more very common operating systems that end users will use on a daily basis: Apple iOS and Google Android OS. Posture works very differently with these mobile platforms. These platforms are inherently more secure than Windows and macOS, and they require a device manager to gain the level of visibility into the endpoint that AnyConnect is able to achieve with the desktop operating systems. ISE has links to one or more mobile device managers for BYOD onboarding and also to provide posture assessment on the endpoint.

There are many different device managers in the world, and each one has its own API and would therefore require most companies to write custom integrations for every mobile device manager it supports. Cisco has the most commonly deployed NAC solution in the world with ISE, so Cisco was able to take a very different approach: Cisco published a standard API for all mobile device managers, and if someone wants a mobile device manager to integrate with ISE, that product must implement the standard API. Then the mobile device manager vendor works with Cisco's Solution Partner program for integration testing and support. It's a really beautiful model that shows the value of being the 800-pound gorilla in certain markets.

The model of asking the mobile device manager vendor to support the Cisco API minimizes the development, TAC case load, and testing effort on Cisco ISE; it also expands the list of supported vendors to any mobile device manager that wants to join the ecosystem. An

ancillary benefit is that it permits the ISE user interface to remain very simple and identical across all mobile device manager vendors, thereby improving the user's experience.

At the time this book was written, every major device manager in the market was supported for integration, including Cisco Meraki Systems Manager (Meraki SM), Jamf Pro, VMware Workspace One (formerly Airwatch), Mobile Iron UEM, IBM MaaS360, Citrix XenMobile, BlackBerry (formerly Good Technology), and Microsoft Intune and SCCM.

The Cisco MDM API can retrieve specific attributes from the mobile device managers that are supported. You can see the list of supported mobile device managers when you create authorization conditions that use the MDM library. To see how this works, bring up the condition editor by navigating to Work Centers > Network Access > Policy Elements > Conditions > Library Conditions, as shown in Figure 18-118.

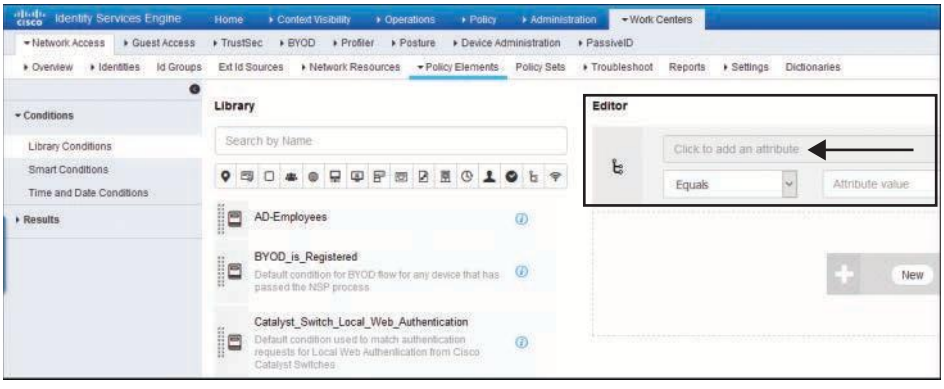


Figure 18-118 *Library Conditions*

In the Editor, click on the text Click to Add an Attribute and select the MDM dictionary from the dropdown, as shown in Figure 18-119.

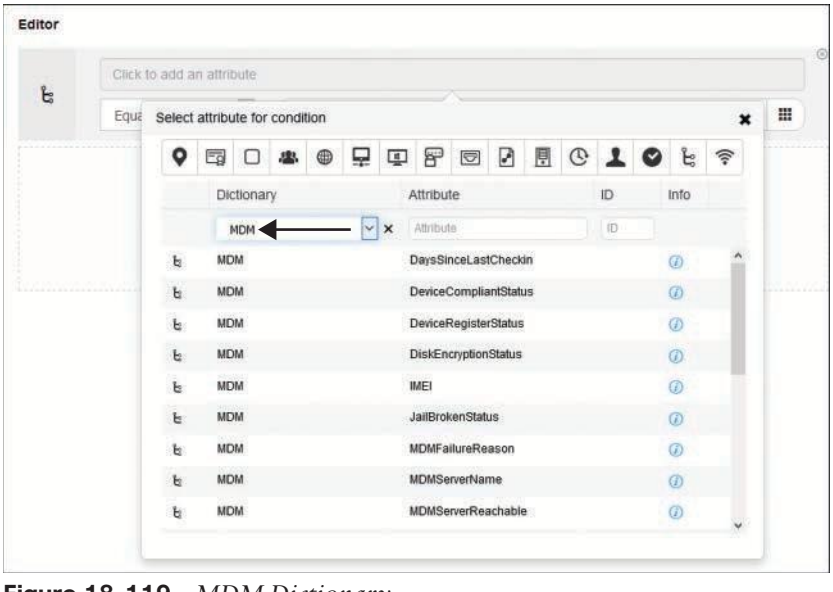


Figure 18-119 *MDM Dictionary*

The Editor now displays the attributes in the MDM dictionary, which is disabled and hidden until you configure a link to a mobile device manager under Administration > Network Resources > External MDM (refer to Chapter 16).

Some of these attributes come from the mobile device manager directly, such as DeviceCompliantStatus, which describes the mobile device manager's determination of the endpoint compared to the policies in the mobile device manager (aka posture). Other attributes are generated by ISE, such as MDMServerReachable, which describes the state of the mobile device manager as ISE tries to communicate with it and can be used for fail-open or fail-closed policies.

Create Mobile Posture Authorization Conditions

From the Editor screen shown in Figure 18-119, create and save a condition named MDM_Compliant, which has the following conditions: MDMServerReachable = Reachable AND DeviceCompliantStatus = Compliant (see Figure 18-120).



Figure 18-120 *MDM_Compliant Condition*

Now you can use the MDM_Compliant condition in authorization rules for allowing access to the corporate network. To do so, create and save another condition named MDM_NonCompliant with the conditions shown in Figure 18-121.



Figure 18-121 *MDM_NonCompliant Condition*

Now you can use the MDM_NonCompliant condition in authorization rules for explicitly denying access to the corporate network. To do so, create and save another condition named MDM_Detailed_Posture with the conditions shown in Figure 18-122.

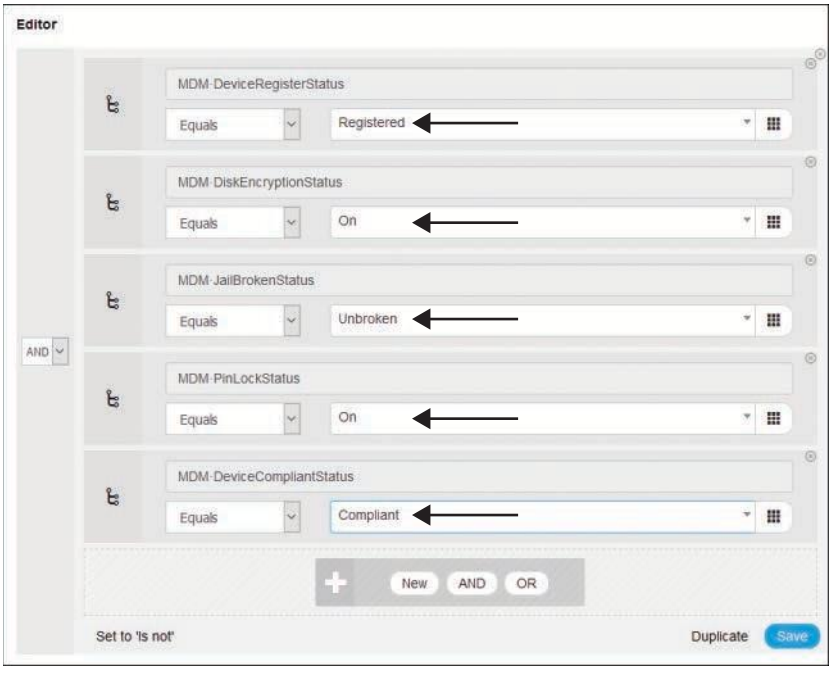


Figure 18-122 *MDM_Detailed_Posture Condition*

The MDM_Detailed_Posture condition shows you what is possible: You can use the individual attributes returned from the mobile device manager when that much control in the authorization is required, or you can use the simple compliant result in your authorization rules.

Create Mobile Posture Authorization Rules

As you've read many times throughout this book, you should always place the most specific rules toward the top of an authorization policy.

This section shows how to create three authorization rules to represent mobile posture: MDM_NotReachable, MDM_Compliant, and MDM_NonCompliant (see Figure 18-123). The following list describes these rules:

- **MDM_NotReachable:** This authorization rule is built to use five conditions. If all five conditions are true, then the endpoint is granted access to the Internet (and only to the Internet). These are the five conditions (in order, as shown in Figure 18-123):

  1. The MDM server or service is not reachable.
  2. The network access method is wireless 802.1X.
  3. The 802.1X authentication succeeded.
  4. The user is a member of the Employees group in Active Directory.
  5. The device is classified as part of the Mobile_Endpoint logical group.

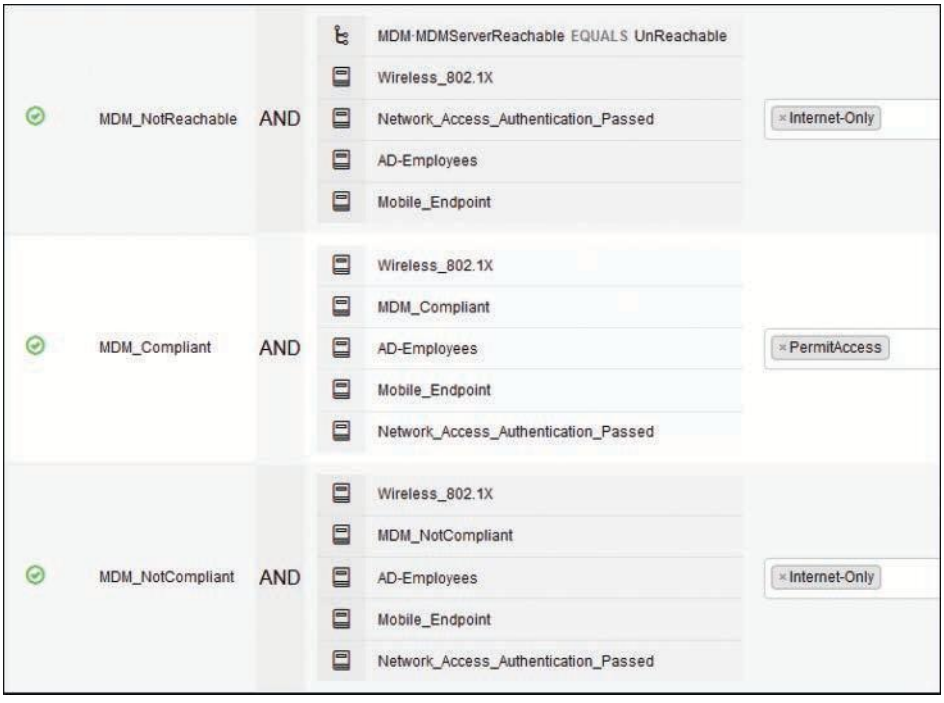


Figure 18-123 *MDM Posture Authorization Rules*

- **MDM_Compliant**: This authorization rule also has five conditions. When all five conditions are met, the endpoint is granted full access to the network. These are the conditions:
  1. The network access method is wireless 802.1X.
  2. The mobile device manager is reachable and responds that the mobile device is compliant.
  3. The user is a member of the Employees group in Active Directory.
  4. The device is classified as part of the Mobile_Endpoint logical group.
  5. The 802.1X authentication succeeded.
- **MDM_NonCompliant**: This authorization rule also has five conditions. When all five conditions are met the endpoint is granted Internet-only access. These are the conditions:
  1. The network access method is wireless 802.1X.
  2. The mobile device manager is reachable and responds that the mobile device is not compliant.
  3. The user is a member of the Employees group in Active Directory.
  4. The device is classified as part of the Mobile_Endpoint logical group.
  5. The 802.1X authentication succeeded.

Figures 18-124 and 18-125 show portions of the detailed authentication report. You can see that the endpoint that authenticated is a Samsung Galaxy S10 Android phone that is being used by the employee1 user.



Figure 18-124 *MDM Posture Authorization*

Figure 18-124 shows the Overview section of the detailed authentication report, accessed from Live Log. The authorization policy is MDM_Compliant, and the authorization results are the PermitAccess profile and the BYOD TrustSec tag.

Figure 18-125 shows a portion of the Authentication Details section of the detailed authentication report. The MDMServerReachable attribute is true, and the DeviceCompliantStatus is also true.



Figure 18-125 *MDM Posture Authorization*

Exam Preparation Tasks

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 18-2 lists these key topics and the page number on which each is found.

Key Topic

Table 18-2 Key Topics

Key Topic Element	Description	Page
Paragraph	Keeping the compliance modules up to date on ISE and AnyConnect	637
Paragraph	ISE as a headend and using the CPP to deploy AnyConnect and its modules	648
Paragraph	Mandatory requirements for posture evaluation	689
Paragraph	Posture assessment versus authorization policy	689
Paragraph	Order of authorization rules	694

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

posture, Cisco Network Admission Control (NAC), network access control, posture condition, posture remediation, posture requirement

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is the difference between an anti-malware condition and an anti-virus condition?
2. What is the difference between a dictionary compound condition and a compound condition?
3. What is a patch management condition?
4. What is a remediation?
5. What is the difference between stealth mode and the temporal agent?



CHAPTER 19

Deploying Safely

This chapter covers the following topic:

802.1X phasing (monitor mode, low impact, closed mode)

Throughout this book, you have seen quite a bit of configuration detail about ISE and about network access devices. You have explored the technical merit of policy creation, guest life-cycle management, posture assessment, and much more. There is obviously a great deal to consider when you deploy a system such as this one; it is not something you can just enable overnight by flipping a switch.

This chapter focuses on the recommended approach to deploying not only ISE but also a full secure network access system. It reviews some of the challenges that were common in the past and why certain technologies were enhanced to provide a more prescriptive approach to deployment.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 19-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 19-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Why Use a Phased Approach?	3, 9
Comparing authentication open to Standard 802.1X	4, 5
Prepare ISE for a Staged Deployment	6
Monitor Mode	1
Low-Impact Mode	2
Closed Mode	10
Transitioning from Monitor Mode to Your End State	8
Wireless Networks	7

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What is monitor mode?
   a. An authentication open interface configuration command on 802.1X-enabled interfaces
   b. A setting in ISE to record actions but not take them
   c. A method for identifying what device would have failed authentication and correcting the root cause before the failure occurs
   d. A method for alerting the administrator of failed authentications so the end user can be called and manually granted network access
2. What is low-impact mode?
   a. One of the two end states of authentication that limits access but still uses the **authentication open** interface configuration command
   b. One of the two end states of authentication that limits access but is less secure than closed mode
   c. A method to ensure that authentications occur but the authorizations are ignored to avoid causing denial of service
   d. A method for identifying what device would have failed authentication and correcting the root cause before the failure occurs
3. What is the primary benefit of a phased deployment approach?
   a. It allows an endpoint to go through multiple phases of authentication prior to gaining network access, including dual-factor authentication.
   b. It permits you to use Cisco-proprietary technology and therefore increase Cisco's stock value.
   c. It enables additional security protocols to extend authentication, such as the use of smart cards.
   d. It ensures that a port, switch, or location is fully ready to be successful before enabling enforcement and specific authorization results.
4. Which of the following statements about the **authentication open** command is true?
   a. The **authentication open** command performs EAP authentications but ignores authorization results.
   b. The **authentication open** command performs EAP authentications but ignores the Access-Reject response.
   c. The **authentication open** command ignores EAP authentications and simply allows devices on the network to be monitored.
   d. The **authentication open** command is used to configure networks without Wi-Fi Protected Setup.

5. Which of the following is true about the **authentication open** command?
   a. The **authentication open** command allows all traffic to pass through the switch port before the authentication result is received from the AAA server.
   b. The **authentication open** command is a key part of closed mode.
   c. The **authentication open** command blocks all traffic except EAP traffic before the authentication response is returned from the AAA server.
   d. The **authentication open** command is used to send TACACS authentication requests to any open and responding server.
6. What ISE configuration allows different groups of authentication and authorization policies?
   a. Policy groupings
   b. Policy sets
   c. Service selection rules
   d. Service sets
7. Where is monitor mode configured for wireless LANs?
   a. It is configured on the WLC, under the security properties for the WLAN.
   b. It is configured in the wireless monitor mode policy set in ISE.
   c. It is configured by enabling wireless monitor mode under the system settings in ISE.
   d. Monitor mode is not possible with wireless LANs.
8. When using policy sets, how can a switch be transitioned from monitor mode to one of the end-state modes?
   a. Move the NAD from the monitor mode NDG to the final state NDG.
   b. Remove the **authentication open** command from the switch interface.
   c. Enter the **low-impact** or **closed** keyword for the RADIUS server definition in the switch.
   d. Enable enforcement mode on the client supplicants.
9. Which of the following statements is true?
   a. A wired port must have a single configuration that supports authentication of supplicants, guests, and non-authenticating devices.
   b. A supplicant requests access to a specific network, and a wired port responds to the requested network.
   c. A wired port informs the RADIUS server what SAML protocol the supplicant used.
   d. A wired port must be preconfigured for each device type that might be connecting to it. Separate configurations are needed for supplicants, guests, and non-authenticating devices.
10. Which of the modes is most closely related to the default behavior of 802.1X?
    a. Closed mode
    b. Monitor mode
    c. Low-impact mode
    d. Cisco enhanced security mode

Foundation Topics

Why Use a Phased Approach?

Back in the early 2000s, a new technology was emerging that would revolutionize networking as we knew it. This technology, IEEE 802.1X, enabled authentication of a network access port prior to allowing devices onto the network. The concept was simple, and it was predicted that within five years, there would not be any "hot ports" in the world that wouldn't first authenticate the user, and no unauthorized users would be getting onto networks anymore.

802.1X was originally created to be very binary in nature so that a device would either be authenticated and get access or fail authentication and be denied. Figure 19-1 illustrates the intended behavior of 802.1X.

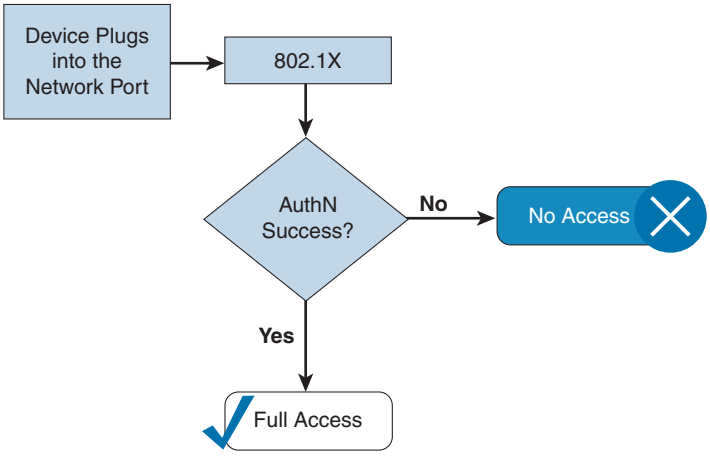


Figure 19-1 *802.1X Intended Behavior*

However, there are many different moving parts to this authentication process that must all be aligned properly in order to prevent denial of service to the user population. Supplicants must be configured on devices; lists of MAC addresses must be created in order to properly prepare for MAC Authentication Bypass (MAB) and allow non-authenticating devices to access the network; profiling probes must be enabled and must collect data about endpoints to help build that list of MAC addresses; certificates must be trusted; guest accounts must be created; and much more.

If you were to simply "flip the switch" and enable 802.1X on all access-layer switch ports at once, a swarm of angry users would likely converge on the IT department, screaming bloody murder and demanding that people lose their jobs. Such a move is often called a career limiting event, or CLE for short.

A financial organization implemented 802.1X with 2000 switch ports in a campus building. Due to an audit requirement, network authentication had to be enabled by a certain date, or the company would be subject to fines. The mandate came down from above, the project

received its funding, and away went the deployment. The company lab tested everything and proved it would all work using Cisco Catalyst 6500 Series switches and the native Windows XP (Service Pack 3) supplicant configured for EAP-TLS machine authentication with the Active Directory–issued machine certificate.

It was beautiful. Everything was working perfectly on the test systems in the lab. The desktop team assured everyone that the Group Policy Object (GPO) was sent out properly, and all the Windows XP systems were ready to authenticate. The network team should have (theoretically) been able to simply turn on the authentication on the switch ports.

The advice was still to deploy in monitor mode first and then change over to closed mode (the end state). This meant that the **authentication open** command needed to be applied to the switch port, but it would allow the team to validate that authentications would all be successful before the team truly enforced access to the network.

The security oversight committee nixed the idea immediately because the word *open* was in the command. Based on the short time requirements and politics, the team was not allowed to use the command. Not ever. Never mind that all 2000 ports were wide open.

So, the big day arrived. The change control window was set for 10 p.m. on a Sunday night. It was time to run the scripts and enable closed mode authentication across 2000 switch ports in a matter of minutes. Of those 2000 ports, a whopping 10 were authenticating successfully, and the team had accomplished exactly what it had warned about: It created a denial of service for all other systems.

Why did this happen? The policies were all correct. The certificates had all been pushed out to the desktops. The supplicants were configured. What no one realized is that the supplicant configuration would not take effect prior to rebooting the Windows systems! Of course, the team did not know this was the root cause until the next afternoon, when the desktop team researched it some more, but at that point, it no longer mattered. The team created a nice problem for all the users the morning after the change window. The rollback plan was to disable 802.1X network-wide in an emergency change window.

The story does have a happy ending, though. After the desktop team pushed out a job to reboot all the desktop systems, the team reenabled authentication at the next change control window and had 99% of the systems authenticating successfully.

However, not all deployments are so lucky or so well planned in advance. Therefore, a phased approach to deploying identity solutions is always a good idea. By using a phased deployment approach, you can start off in monitor mode and gradually transition into the end state of either low-impact mode or closed mode. By doing so, you can avoid denial of service. With a monitoring phase, you have time to build a list of endpoints with profiling, you can manually import the MAC addresses that will be using MAB without profiling, and you can ensure that you know exactly what will happen before it happens. Then you can gradually move into a final state, enforcement. Figure 19-2 illustrates the phased deployment concept.

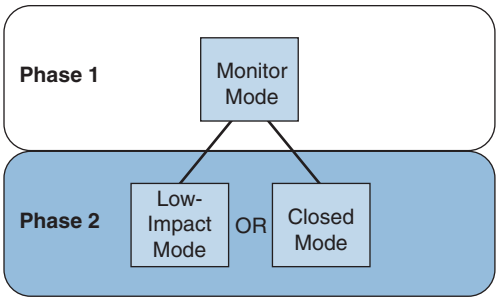


Figure 19-2 *Phased Deployment*

Comparing authentication open to Standard 802.1X

Per the IEEE 802.1X standard, a port that is protected with 802.1X does not allow network traffic to flow without successful authentication. Figure 19-3 illustrates that an 802.1X-controlled port normally only allows EAP, CDP, and LLDP traffic to enter the port, and no other traffic is permitted. (EAP, CDP, and LLDP are all Layer 2 protocols.) When 802.1X is enabled on a port, it is said to be a "supplicant authenticator," which is a complex way of saying that the port communicates with EAP at Layer 2, and the switch brokers that authentication to the RADIUS server.

Figure 19-3 illustrates the default port behavior of an 802.1X-enabled switch port.

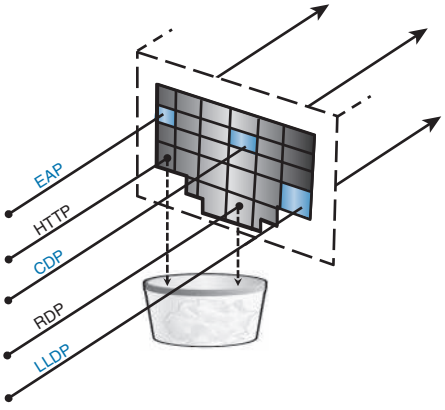


Figure 19-3 *Default Port Behavior with 802.1X*

Cisco created an enhancement to 802.1X ports that allows a port to be a supplicant authenticator and also allows all traffic to flow normally through the switch port, even without authentication occurring. This allows the supplicant to process an authentication correctly if it's configured, but if the device does not have a supplicant configured or if the switch received an Access-Reject message from the RADIUS server, the Access-Reject message is ignored.

Figure 19-4 illustrates that regardless of authentication, the switch port allows all traffic to flow, but it also authenticates the supplicant and performs MAB just like a standard 802.1X-enabled switch port.

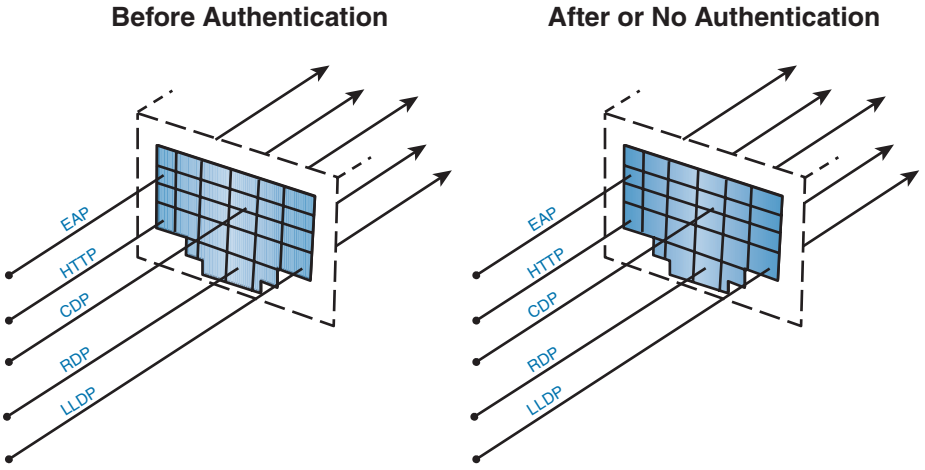


Figure 19-4 *Port Behavior with Open Authentication*

The creation of this authenticator enhancement made monitor mode possible. It is, of course, not the only necessary component of monitor mode, but it was certainly the catalyst.

Prepare ISE for a Staged Deployment

One of the primary ways to differentiate modes within ISE policies is with the use of network device groups (NDGs) and policy sets. You can configure the NDGs to have a top-level group named Stage and create subgroups for monitor mode, low-impact mode, and closed mode. You can then configure the authorization policies to have different policies for each of the network devices that are assigned to particular stages of deployment. While you could do all this in a single authorization policy, using a separate policy set for each stage of deployment keeps the policies clean, small, and organized.

To create the network device group named "Stage," follow these steps from the ISE user interface:

Step 1. Navigate to **Work Centers > Network Access > Network Resources > Device Groups**.

Step 2. Click **Add**.

Step 3. Name the network device group **Stage**.

Step 4. Select **Add as Root Group** from the Parent Group dropdown.

Step 5. Click **Save**.

Figure 19-5 shows the creation of the top-level NDG.

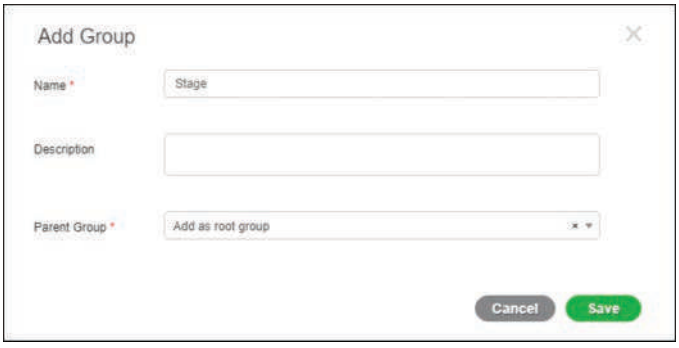


Figure 19-5 *Adding the Stage Top-Level Group*

After the root group named Stage has been created, you can follow these steps to create an NDG for monitor mode:

Step 1. Click **Add**.

Step 2. Name the network device group **Monitor Mode**.

Step 3. Select **Stage** from the Parent Group dropdown, as shown in Figure 19-6.

Step 4. Click **Save**.

Step 5. Repeat steps 1 through 4 and create a new subgroup for low-impact and closed modes.

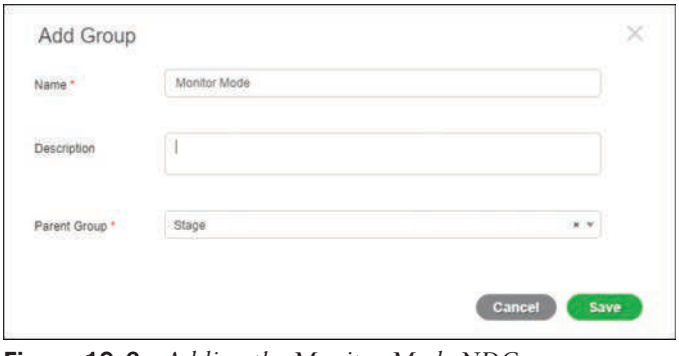


Figure 19-6 *Adding the Monitor Mode NDG*

Figure 19-7 shows the result creating these NDGs.

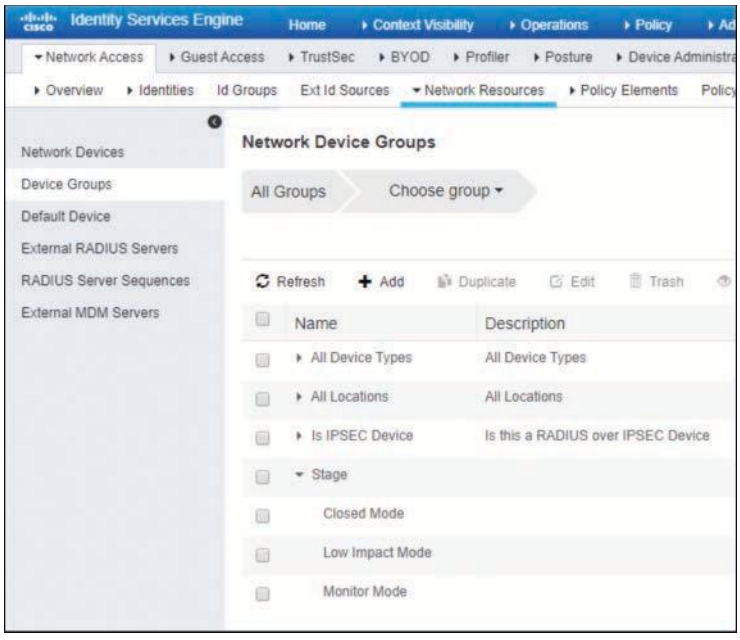


Figure 19-7 *Final Network Device Groups*

Monitor Mode

Monitor mode is a process, not just a command on a switch. The process involves enabling authentication (with **authentication open**), seeing exactly which devices fail and which ones succeed, and correcting the failed authentications before they cause any problems.

The high-level flow diagram in Figure 19-8 describes monitor mode.

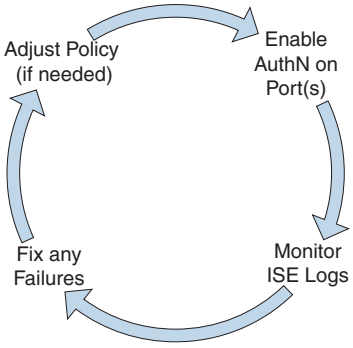


Figure 19-8 *Monitor Mode Operational Flow*

One key point to understand about monitor mode is that it is applicable to wired environments only. If you have ever configured a device to connect to a wireless network, you will recognize that configuring a network manager (which is considered a client and contains a supplicant) is expected behavior when using Wi-Fi. You must tell the Wi-Fi-capable endpoint which network to connect to and provide credentials for that network. It's common, it's expected, and it's well known.

On a wired network, however, there is no such concept of an SSID, so there is no popup on the endpoint that asks which network you would like to connect to; it's just assumed that you are physically connected, and therefore you are attached to the correct network. With wireless, if you don't have a supplicant, you do not connect. Wired environments are expected to always work—with or without a supplicant.

Key Topic

The wired port must be able to handle a device that has a supplicant (802.1X), a corporate device that doesn't have a supplicant but does belong on the network (IP phone or printer), and guest users. So there is quite a bit to audit when in monitor mode.

Key Topic

Another very important thing to understand about monitor mode is that authorization results from the RADIUS server are honored. (Access-Reject is the only command that is ignored.) So, if your authorization result from ISE includes VLAN assignment or dACLs, they are honored and applied to the port.

For a phased deployment approach, it is highly recommended to use NDGs in ISE. By using NDGs, you can build specific policies that send only the basic authorization results (Access-Accept and Access-Reject) to switches that are part of a monitor mode NDG.

The high-level flow diagram in Figure 19-9 describes monitor mode.

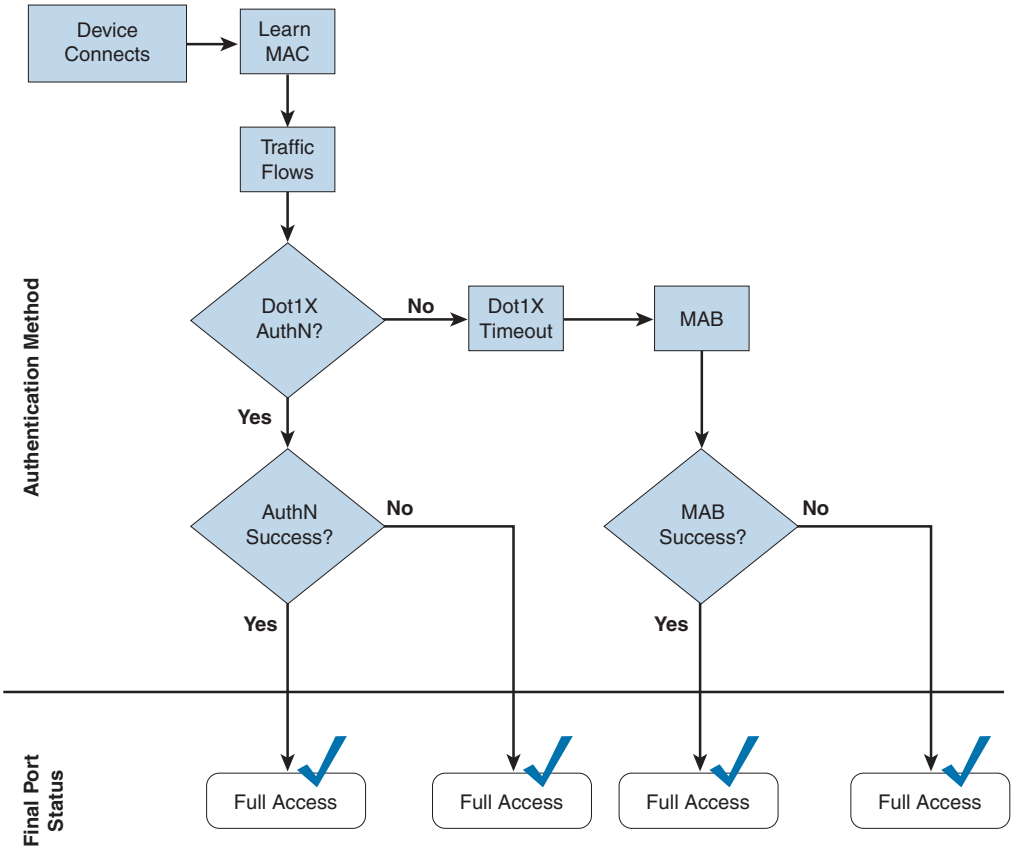


Figure 19-9 *Monitor Mode Flow Diagram*

Once you have NDGs configured for the different stages of the deployment, you can move on to creating the policy sets. To create a policy set for monitor mode, follow these steps:

Step 1. Navigate to **Work Centers > Network Access > Policy Sets**.

Step 2. Click the **+** sign above the default policy set, as shown in Figure 19-10.

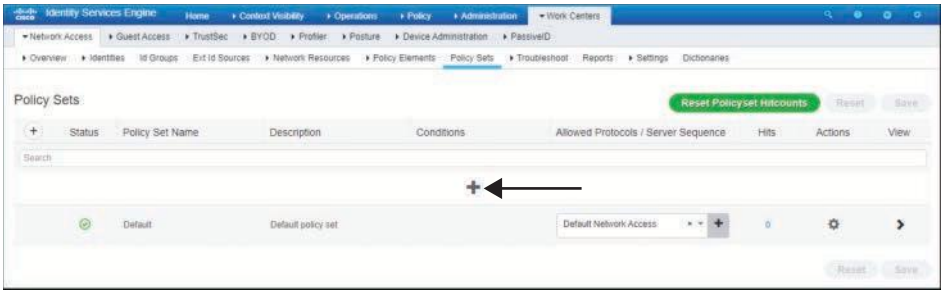


Figure 19-10 *Creating a New Policy Set Above the Default*

Step 3. Name the new policy set **Monitor Mode** and provide a description.

Step 4. Select **DEVICE > Stage > Starts With > Monitor Mode**, as shown in Figure 19-11.

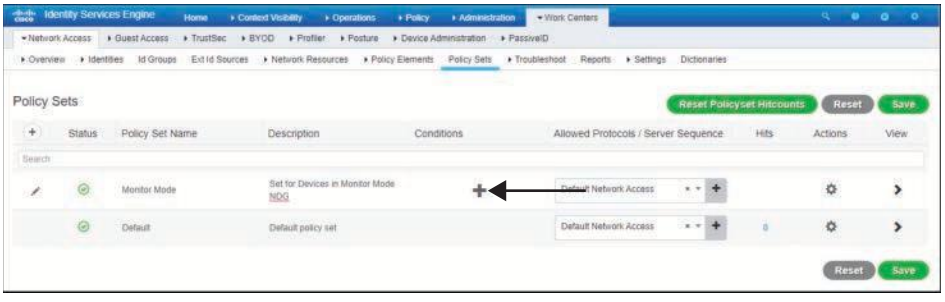


Figure 19-11 *Monitor Mode Condition*

Step 5. Click **Save**.

Figure 19-12 shows the policy set selection rule for monitor mode.

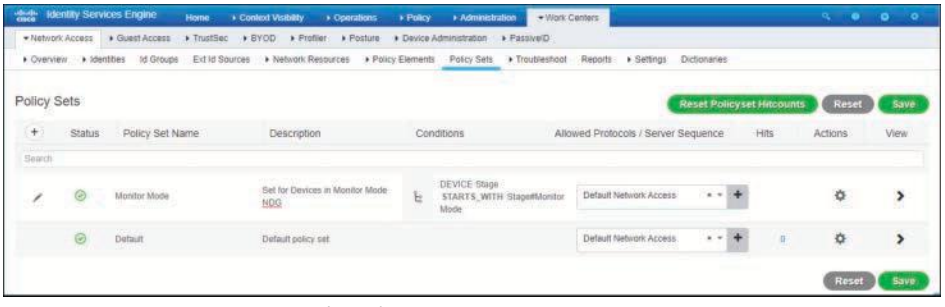


Figure 19-12 *Monitor Mode Policy Set*

At this point, any network device that is a member of the NDG named Monitor Mode uses this policy set. All authentication and authorizations with those network devices occur with this set of policies, and all other network devices continue to use the default policy set.

It is up to you, the administrator of ISE policies, to ensure that the authorization results for monitor mode switches are only Access-Accept and Access-Reject. Always remember that other authorization results will be accepted and applied to the switch port, so you must ensure that Web Authentication, access lists, and VLAN assignments do not occur for these switches. The only exception to this rule is the AV pair for IP phones that grants a phone permission to use the voice VLAN.

Low-Impact Mode

As described previously in this chapter, low-impact mode is one of the end-state choices for a deployment. Closed mode is the other final stage. There is no specific best practice for which mode is better to deploy; it depends on the organization and its needs.

A number of large organizations use a variety of technologies to reimage desktop systems that make use of the Preboot Execution Environment (PXE) to boot into a pseudo-OS and then connect to an imaging server that reimages the company computer. Those PXE environments may be time sensitive and may not have the ability to authenticate to the network infrastructure, but they must be able to seamlessly boot, connect to the reimaging server, and update the desktop to the latest corporate image—without any additional user interaction. Low-impact mode was the only way to make this work feasibly in some environments.

Another example is a retail organization that uses thin clients in its retail stores. These thin clients must be able to boot using PXE, gain limited access to the network, and download their operating systems from the local store server—and they need to have that access before the DHCP timers expire. Once the OS is loaded into memory and takes over the system, its supplicant sends an EAPoL start message to the network and authenticates with 802.1X. Low-impact mode allows the thin client to boot automatically and have the appropriate levels of access to the store server to download the OS.

Low-impact mode adds security on top of the framework that was built in monitor mode. It continues to use the **authentication open** capabilities of the switch port, allowing traffic to enter the switch prior to an authorization result. This permits the DHCP clients to be assigned IP addresses before their DHCP timers run out.

Key Topic

With low-impact mode, you add security right from the start by putting a port-based ACL (pACL) on the switch interface. This is a traffic-filtering ACL that gets applied to the port as part of the switch configuration and is then overridden by the dACL sent down from ISE.

Figure 19-13 shows the operational flow intended for low-impact mode. Because this mode is one of the two possible end states (closed mode being the second), very specific access may be provided per user, device, or other condition that you wish to use in the ISE authorization policies. Remember that the goal of low-impact mode is to provide very limited network access to devices without authentication and then provide very specific access to those that have been authorized; this is the least-privilege security principle.

As with any other security solution, tuning the authorization results can take a lot of operational person-hours. It is therefore recommended to deploy authorization results in stages. For example, begin with a policy that permits full access to any device that has

authenticated successfully. Ensure that the environment is fully functional and then begin to "ratchet down" the security, making the dACLs more specific and so on.

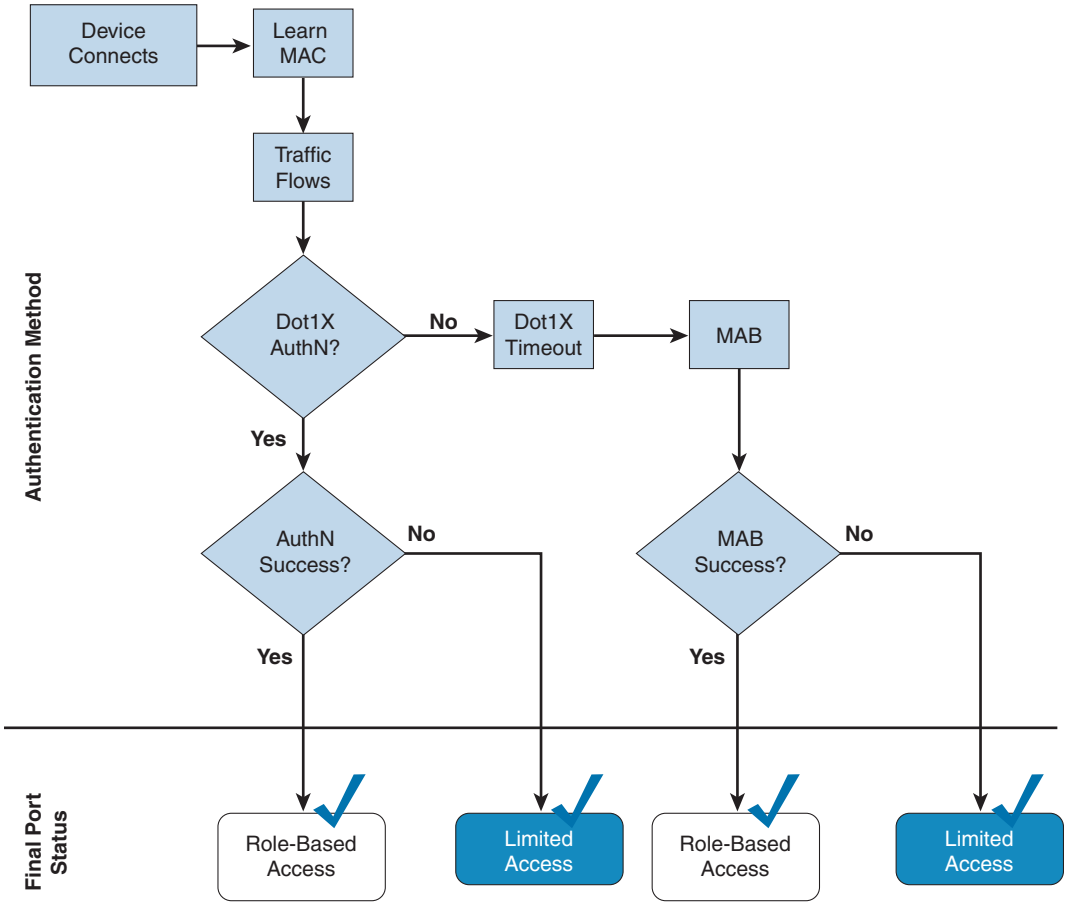


Figure 19-13 *Low-Impact Mode Operational Flow*

Figure 19-14 shows that the pACL is applied prior to authentication in order to allow only specific traffic into the port. Once the authentication occurs, the authorization needs to include a dACL that selectively permits and/or denies traffic. Other authorization results may also be applied at the port, such as URL redirection, VLAN assignment, 802.1AE MACsec encryption, and Security Group Tags.

Key Topic

Keep in mind that VLAN assignment should only be used on devices that use supplicants. Without a supplicant, a device will most likely not be able to identify the VLAN change and may end up with the wrong IP address for its final VLAN assignment.

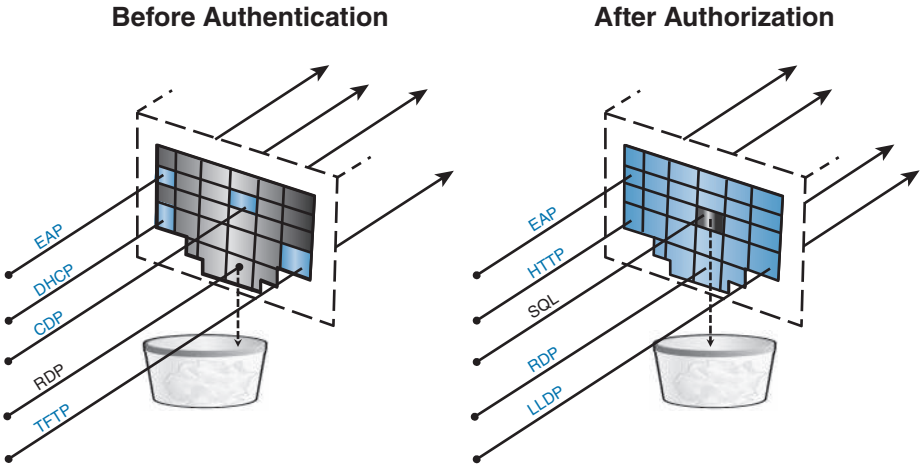


Figure 19-14 *Low-Impact Mode Port Behavior*

To create the policy set for low-impact mode, follow these steps:

Step 1. Navigate to **Work Centers > Network Access > Policy Sets**.

Step 2. Click the **+** sign in the upper-left corner, as shown in Figure 19-15.

Step 3. Name the new policy set **Low Impact Mode** and provide a description.

Step 4. Select **DEVICE > Stage > Starts With > Low Impact Mode**, as shown in Figure 19-15.

Step 5. Click **Save**.

Figure 19-15 shows the policy set selection rule for low-impact mode.

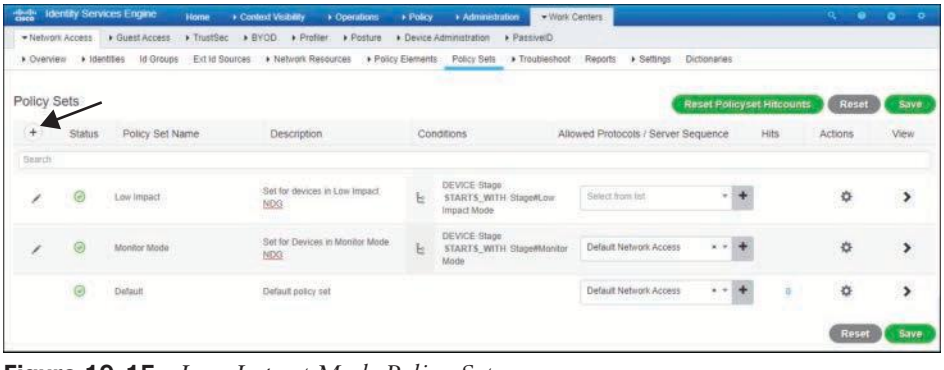


Figure 19-15 *Low-Impact Mode Policy Set*

Closed Mode

Closed mode is similar to the default behavior of 802.1X. As you can see in Figure 19-3, the port does not allow any traffic before the authentication (except for EAP, CDP, and LLDP), and then the port is assigned to specific authorization results after the authentication.

NOTE Closed mode was once called high-security mode but was renamed due to the perception that it was more secure than low-impact mode. In truth, the two modes are equally secure. The security level of either end state depends on the configuration of the devices and the policies on ISE, not on the mode of operation. In other words, an administrator can make closed mode very insecure or very secure, depending on the implementation.

As you can see in Figure 19-1, the operational model of 802.1X was designed to deny access to any device that does not authenticate successfully. This is a perfectly understandable model for wireless network access, where a human is required to interact with the device and configure a wireless client (supplicant) to connect to a specific SSID with specific credentials.

However, in a wired world, many devices require network access without any user interaction (for example, IP cameras, IP phones, printers, fax machines, badge readers). So, MAB had to be added to the process flow.

The concept of completely denying access to the network if authentication fails or if a supplicant is not configured proved to have operational difficulties. Some level of access was needed. Originally the switch itself would have a "failed authentication VLAN," where it would make a local decision to authorize access to a specific VLAN when a device failed authentication. In addition, if authentication timed out (meaning there was no supplicant on the endpoint), the switch would authorize access to a locally configured guest VLAN.

One of the problems with that original logic is the lack of centralized knowledge and control. As far as the policy server was concerned, the access was denied. Yet the device was still on the network because the NAD made a local decision despite what the policy server said.

Figure 19-16 shows the operational flow of closed mode. Notice that it is nearly the same as the flow for low-impact mode. All the same authorization results are available for use, but, unlike low-impact mode, closed mode does not allow any of the PXE-type traffic into the port prior to the authorization result.

Figure 19-17 shows the port behavior in closed mode. Virtually zero traffic is allowed into the port before the authentication. Once the session is authorized, very specific authorization results may be applied to the port, such as VLAN assignment, dACL, URL redirection, MACsec encryption, and Security Group Tags.

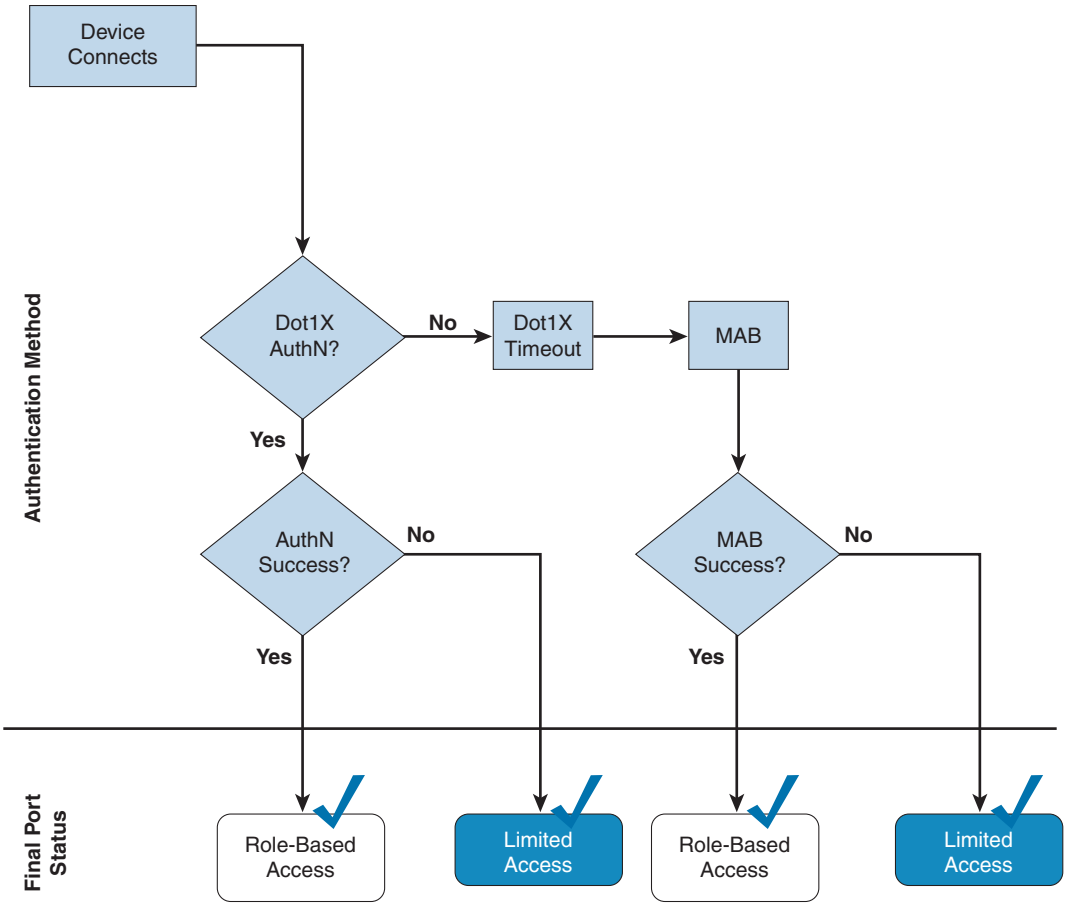


Figure 19-16 *Closed Mode Flow*

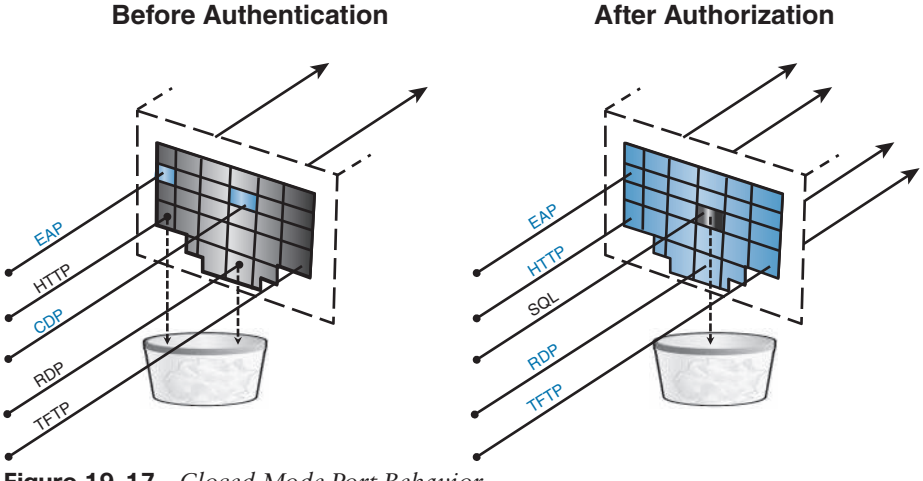


Figure 19-17 *Closed Mode Port Behavior*

To create the closed mode device group, follow these steps from the ISE user interface:

Step 1. Navigate to **Work Centers > Network Access > Policy Sets**.

Step 2. Click the **+** sign in the upper-left corner, as shown in Figure 19-18.

Step 3. Name the new policy set **Closed Mode** and provide a description.

Step 4. Select **DEVICE > Stage > Starts With > Closed Mode**, as shown in Figure 19-18.

Step 5. Click **Save**.

Figure 19-18 shows the policy set selection rule for closed mode.

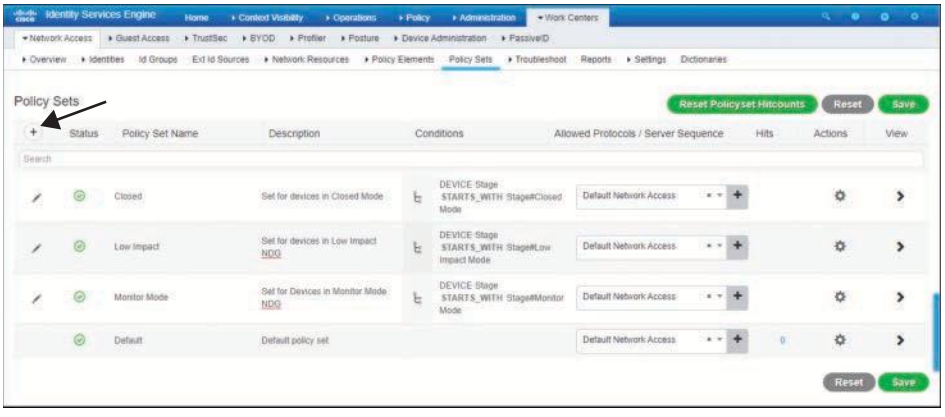


Figure 19-18 *Closed Mode Policy Set*

Transitioning from Monitor Mode to Your End State

The key to successfully using a phased deployment approach is understanding how to transition from monitor mode to the end state chosen (low-impact mode or closed mode). This is why you build out NDGs and policy sets, as shown in this chapter.

With monitor mode, you must ensure that only Access-Accept and Access-Reject authorizations are used. With low-impact and closed modes, you can send the other authorization results, such as sending a URL redirection for Centralized Web Authentication (CWA).

The purpose of monitor mode is to ensure that the endpoints are all authenticating correctly either via 802.1X with their supplicants or via MAB with profiling or even statically. If you get the first pilot switch ready, and all the devices are prepared for authentication and everything looks good, you flip the switch and change the default authorization policy to send a CWA result instead of just the basic accept or reject message. That first switch is fine, all the devices work correctly, and life looks easy.

However, you wouldn't want to push the CWA result to the switch port if you had not fully prepared the supplicants and educated the users about a possible change of experience when logging in to the network. That would be another CLE. That is why you use NDGs. You

ensure with the network access device's membership of the Stage NDG that you are sending the correct results to the correct network devices.

Imagine rolling out ISE to thousands of branch locations. You prepare a branch by putting it into monitor mode. When you are certain the branch is fully ready and all the devices are recognized and authenticating successfully, you can move the switch from the monitor mode NDG to the end-state NDG and make a few command modifications to the switches.

Wireless Networks

Wireless networks behave differently than wired networks. With the creation of a WLAN, an administrator must define the security for that LAN. When using 802.1X, you set the security to use WPA2 for key management. This setting cannot be mixed with an open authentication, and there are no fallback options.

For guest authentication, the guest needs to connect to a different SSID. This is fundamentally a much different model from a wired network, which needs to handle all different user, device, and authentication types on a single switch port configuration.

Even though wireless behaves differently, the authorization results in ISE may be configured to send the responses to wired devices and wireless devices, providing a unified access strategy. This permits wireless networks to be managed as part of your low-impact mode or closed mode deployments.

Some ISE installations use a separate policy set for wireless authentications and authorizations, and others will assign wireless NADs to the low-impact mode or closed mode NDG.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 19-2 lists these key topics and the page number on which each is found.

Key Topic

Table 19-2 Key Topics

Key Topic Element	Description	Page Number
Paragraph	Wired ports handling all types of devices with a single configuration	723
Paragraph	**authentication open**	723
Paragraph	Security added to low-impact mode	725
Paragraph	To not change VLANs on devices without a supplicant	726

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What command enables an enhancement to the 802.1X standard and allows traffic to flow before an authentication result is received?
2. What is the only response from the RADIUS server that is ignored by a switch port in monitor mode?
3. What should be used to differentiate switches in monitor mode from those in low-impact mode?
4. What network type cannot leverage monitor mode?
5. What deployment strategy should be employed to ensure a successful and safe deployment without causing outages?

This page intentionally left blank



CHAPTER 20

ISE Scale and High Availability

This chapter covers the following topics:

- Configuring ISE nodes in a distributed environment
- Understanding the HA options available
- Using load balancers
- Maintaining ISE deployments

While ISE can run with all services on a single physical or virtual appliance, scale is still an issue. A single appliance cannot scale to large enterprise levels, nor would such a model be highly resilient to failure.

ISE cubes can be designed with each persona running on separate physical or virtual appliances in a single location or in multiple locations. In addition, features within ISE and the network access devices provide redundancy for authentication.

The uptime and availability of a solution should be a full strategy, not just a product feature. Planning for patching and backups of a solution are critical to ensuring availability.

This chapter build up the foundation for ISE architecture that was built in Chapter 6, "Cisco Identity Services Engine Architecture."

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 20-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 20-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Configuring ISE Nodes in a Distributed Environment	1
Understanding the High Availability Options Available	3–5
Using Load Balancers	6, 7
Maintaining ISE Deployments	2, 8–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. How does a PSN join an ISE cube?
   a. From the Deployment screen on the secondary nodes, select Join Cube and enter the FQDN and credentials of the cube controller.
   b. From the Deployment screen on the PAN, click Create Cube and then click Register and add the FQDN and credentials for the other nodes.
   c. From the Deployment screen on the PAN, click Register and add the FQDN and credentials for the other nodes.
   d. PSNs are standalone nodes and do not join an ISE cube.
2. You must create a scheduled backup of ISE for recovering the ISE deployment in case of a disaster. Which of the following options are possible? (Choose two.)
   a. The Policy Administration node can be backed up to an FTP site.
   b. A single backup that includes the Policy Administration node and the MnT node can be scheduled to be stored on an SFTP server.
   c. A backup of the MnT node's operational database can be scheduled to be stored on an SFTP server.
   d. A backup of the Policy Administration node's configuration database can be scheduled to be stored on an HTTP server.
3. What three pieces of information are needed for an ISE license?
   a. The output from the **show license** CLI command, the email receipt for the purchase, and serial number
   b. The unique device ID (UDID), version number (VPID), and serial number
   c. The product ID (SPID), unique device ID (UDID), and serial number
   d. The product ID (SPID), version number (VPID), and serial number
4. How does high availability work for an ISE Policy Administration node?
   a. Gigabit Ethernet 4 is used for stateful heartbeat. When the primary no longer responds, the secondary takes over.
   b. The secondary is manually promoted from the secondary's GUI.
   c. The secondary is manually promoted from the primary's GUI.
   d. There is no HA for the Policy Administration node.
5. How does the Monitoring persona's high availability work?
   a. ISE uses TCP syslog, and if the primary node does not respond, the other nodes send logs to the secondary.
   b. Gigabit Ethernet 4 is used for stateful heartbeat. When the primary no longer responds, the secondary takes over.

c. The Monitoring persona does not have a high availability function.

d. Logs are sent to both MnT nodes automatically. If one MnT node goes down, the other node still receives logs.

6. What is the purpose of a node group?

a. Node groups are used for stateful synchronization between PSNs. If one PSN goes down, another PSN from the node group assumes its sessions automatically.

b. Node groups are used for a multicast heartbeat between PANs. If one PAN goes down, another PAN from the node group takes over.

c. Node groups are used for a heartbeat between PSNs. If one PSN goes down, another PSN from the node group sends a Change of Authorization (CoA) for establishing sessions of the fallen node.

d. Node groups are used for a multicast heartbeat between MnT nodes. If one MnT node goes down, another MnT node from the node group takes over.

7. Which of the following options best describes how NADs should be configured when using a load balancer with ISE PSNs?

a. Each NAD must have a separate RADIUS server entry for each PSN in the ISE cube and an additional one for the VIP address.

b. Each NAD should have a RADIUS server entry for the VIP address and list each of the ISE PSNs as a separate dynamic author client entry for CoA.

c. Each NAD should have one entry for each PSN, but all the entries should use the VIP address instead of the PSN's IP address.

d. The NAD is not configured any differently when a load balancer is in use.

8. How are patches applied to Cisco ISE?

a. Patches are downloaded and applied automatically using GitHub.

b. Patches are downloaded from Cisco.com and applied through the GUI.

c. Patches are downloaded but not applied automatically. They are downloaded from GitHub.

d. Patches are downloaded and applied automatically as part of the ISE feed service.

9. How do you verify the status of an ISE backup?

a. The status may only be viewed from the CLI.

b. The status of a restore is available in the GUI, but backup status is not.

c. The status is not viewable in ISE Version 1.2.

d. The status of a backup can be viewed from the GUI under Administration > System > Backup & Restore.

10. Where do you set the order for patching ISE nodes?

a. You set the order under Administration > System > Settings > Patch Management.

b. You set the order on the Administration > System > Maintenance > Patch Management page.

c. It is not configurable; all nodes are patched simultaneously.

d. It is not configurable; all nodes are patched, in alphabetical order.

Foundation Topics

Configuring ISE Nodes in a Distributed Environment

Recall from Chapter 6 that ISE has multiple personas or functions. Table 20-2 lists the personas and the function each persona provides. The third column in Table 20-2 provides the friendly name for an ISE node that is running that service or persona.

Table 20-2 ISE Persona Types

ISE Persona	Description	Node Name
Administration	Is responsible for providing policy and database synchronization for the entire ISE deployment (ISE cube). Also serves as the single point of interactive GUI administration for the ISE cube.	Policy Administration node (PAN)
Policy Services	Provides the RADIUS communication, performs the authentications and authorizations, and hosts all user-facing web portals.	Policy Services node (PSN)
Monitoring	Receives all the logs from the ISE cube members, maintains the global session directory, and publishes the main topic data to the pxGrid controller.	Monitoring and Troubleshooting (MnT) node
pxGrid	Controls the publish/subscribe (pub/sub) bus for scalable security context sharing to and from the ISE cube.	pxGrid

All ISE nodes are installed in a standalone mode by default. When in a standalone mode, the ISE node is configured to run all personas by default. This means that the standalone node runs Administration, Monitoring, and Policy Services personas. Also, each ISE standalone node is configured as its own root certificate authority (CA).

It is up to you, the ISE administrator, to promote the first node to be a primary Administration node and then join the additional nodes to this new deployment. At the time of joining, you also determine which services to run on which nodes; in other words, you determine which persona each node will have.

You can join more than one ISE node together to create a multi-node deployment, known commonly in the field as an *ISE cube*. It is important to understand that before any ISE nodes can be joined together, they must trust each other's administrative certificates. If that trust is absent, you receive the communication error "node was unreachable"; the root cause of such an error is the lack of trust.

Similarly, if you try to connect to a secure website that is not using a trusted certificate, you see an SSL error in your web browser. This is very much like the error just described, except that it is based on Transport Layer Security (TLS).

If you are still using the default self-signed certificates in ISE, you are required to import the public certificate of each ISE node into each of the other ISE nodes through the Administration > System > Certificates > Trusted Certificates screen. This is necessary because they are all self-signed (untrusted) certificates, and each ISE node needs to trust the primary node; in addition, the primary needs to trust each of the other nodes.

Instead of dealing with all the nonsense related to importing public keys for these self-signed certificates, the best practice is to always use certificates issued from the same trusted source. In that case, you just need to add the root certificates to the trusted certificates list.

Make the First Node a Primary Device

Because all ISE nodes are standalone by default, you must promote the ISE node that will become the primary Policy Administration node (PPAN) to be a primary device instead of a standalone. To do so, follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > System > Deployment**.

Step 2. Select the ISE node (there should only be one at this point), as shown in Figure 20-1.

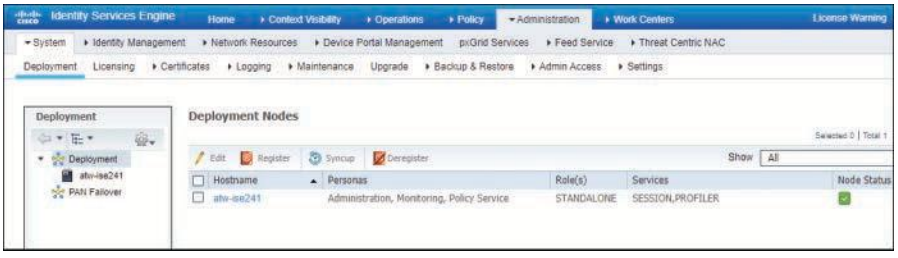


Figure 20-1 *Deployment Screen*

Step 3. Click the **Make Primary** button, as shown in Figure 20-2.

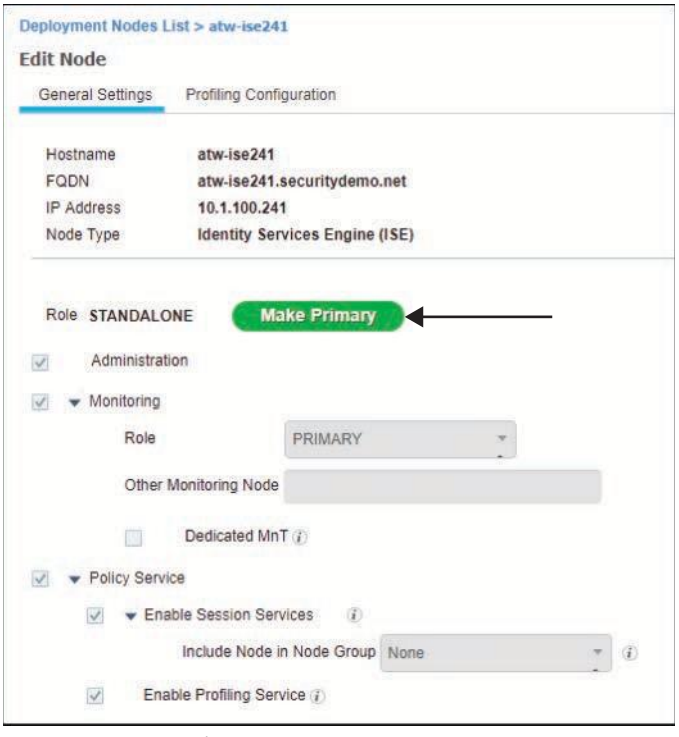


Figure 20-2 *Make Primary Button*

At this point, the Monitoring and Policy Service checkboxes have become selectable. If the primary node will not also be providing any of these services, uncheck them now. (You can always come back and make changes later.)

Step 4. Click **Save**.

Registering an ISE Node to the Deployment

Once there is a primary PAN, you can finally have a multi-node deployment. From the GUI on the PPAN, you register and assign personas to all ISE nodes.

NOTE In order to join another ISE node to the deployment, both nodes must have mutual trust. In other words, all nodes must trust each other's certificates. The easiest way to accomplish this is to use signed certificates on all nodes prior to joining them to a single cube. See Chapter 8, "Initial Configuration of Cisco ISE," for a detailed walkthrough of replacing the self-signed certificates with a signed certificate.

Follow these steps to register a new node to the primary and form your ISE cube:

Step 1. Navigate to **Administration > System > Deployment**.

Step 2. Click **Register**.

NOTE As with all other operations with ISE, DNS is a critical component.

Step 3. Enter the DNS name or IP address of the first ISE node you will be joining to the cube.

Step 4. Enter the administrator name (which is admin by default) and password. Figure 20-3 shows the registration page, with the hostname and credentials entered.

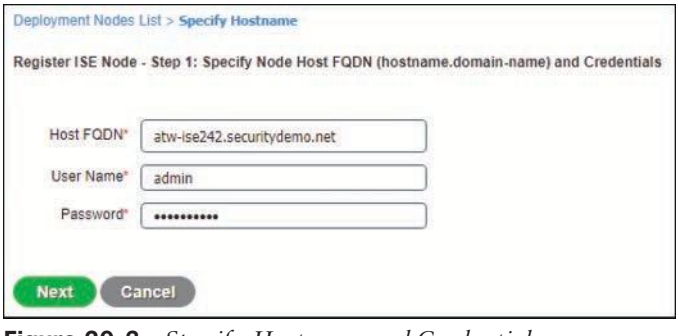


Figure 20-3 *Specify Hostname and Credentials*

Step 5. Click **Next**. The Configure Node screen, shown in Figure 20-4, appears. On this screen, you choose the main persona of the ISE node, including enabling profiling services; however, you cannot configure the individual profiling probes at this point.

Step 6. Choose the personas to run on the newly registered node.

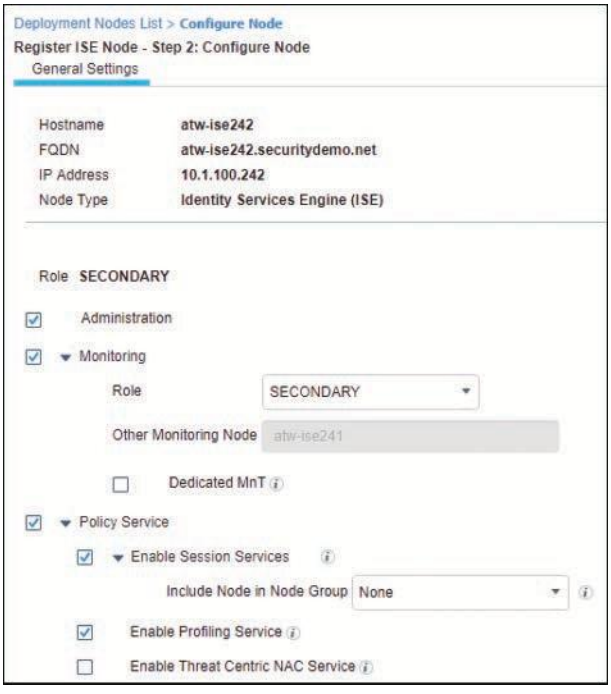


Figure 20-4 *Configuring a Node*

Notice in Figure 20-4 that the Administration and Monitoring personas are automatically configured to be secondary, leaving the existing node as the primary node in the ISE cube.

Step 7. Click **Submit**. The processing message shown in Figure 20-5 appears while the credentials and certificates are tested, and then the primary PAN syncs the entire configuration and the databases to the newly joined ISE node, as shown in Figure 20-6.

Step 8. Repeat steps 2 through 7 for all the ISE nodes that should be joined to the same cube.

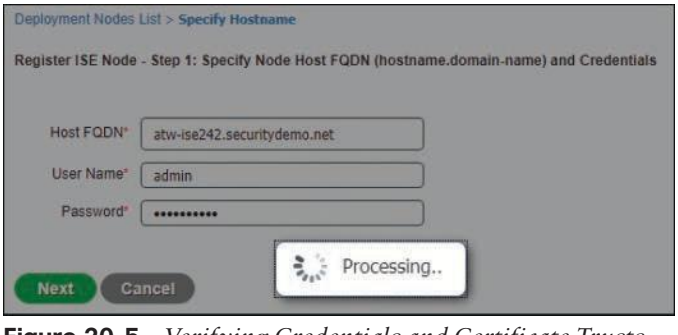


Figure 20-5 *Verifying Credentials and Certificate Trusts*

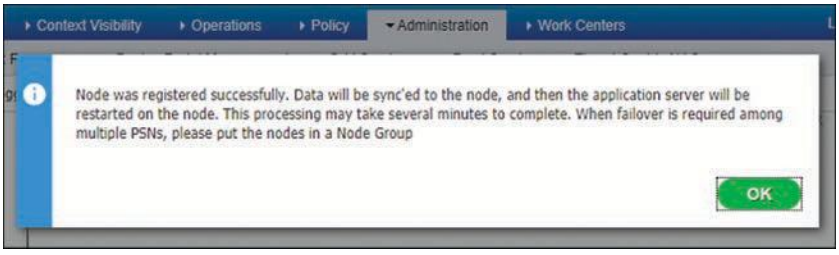


Figure 20-6 *Sync Initiated*

NOTE You can monitor the status of the ISE node that has been joined by using the **show application status ise** command from the command-line interface through either the console or an SSH session to the ISE node (see Example 20-1). When the application server state changes from initializing to running, you know that ISE is fully up and running.

Example 20-1 *show application status ise Output*

```
atw-ise242/admin# show application status ise

ISE PROCESS NAME                       STATE            PROCESS ID
------------------------------------------------------------------
Database Listener                      running          22027
Database Server                        running          96 PROCESSES
Application Server                     initializing
Profiler Database                      running          14397
ISE Indexing Engine                    running          18827
AD Connector                           running          10277
M&T Session Database                   running          16126
M&T Log Processor                      running          19403
Certificate Authority Service          running          11926
EST Service                            running          17631
SXP Engine Service                     disabled
Docker Daemon                          running          23498
TC-NAC Service                         disabled

Wifi Setup Helper Container            disabled
pxGrid Infrastructure Service          disabled
pxGrid Publisher Subscriber Service    disabled
pxGrid Connection Manager              disabled
pxGrid Controller                      disabled
PassiveID WMI Service                  disabled
PassiveID Syslog Service               disabled
PassiveID API Service                  disabled
PassiveID Agent Service                disabled
```

```
PassiveID Endpoint Service              disabled
PassiveID SPAN Service                  disabled
DHCP Server (dhcpd)                     disabled
DNS Server (named)                      disabled
ISE Messaging Service                   running          18433
Segmentation Policy Service             disabled
SSE Connector                           disabled

atw-ise242/admin#
```

NOTE ISE Version 2.6 added the new UI option Dedicated MnT. This option improves MnT node performance by disabling all other personas and services enabled on that node—even though you could do that yourself manually if you choose to. With ISE scale and performance numbers, the maximum scale requires dedicated nodes for each persona.

Ensure That the Persona of Each Node Is Accurate

When all your ISE nodes are joined to the deployment, you can ensure that the correct personas are assigned to the appropriate ISE nodes. Table 20-3 shows the ISE nodes in the sample deployment, as well as the associated persona that will be assigned. Figure 20-7 shows the final Deployment screen, after the synchronization has completed for all nodes. (The green checkmark icon indicates that a node is healthy and in sync.)

NOTE This is a good time to double-check that all the desired probes and services are enabled on the PSNs.

Table 20-3 ISE Nodes and Personas

ISE Node	Persona	Comments
atw-ise241	Administration, Monitoring	Primary PAN, primary MnT
atw-ise242	Administration, Monitoring	Secondary PAN, secondary MnT
atw-ise243	Policy Services	Session and profiler services
atw-ise244	Policy Services	Session and profiler services

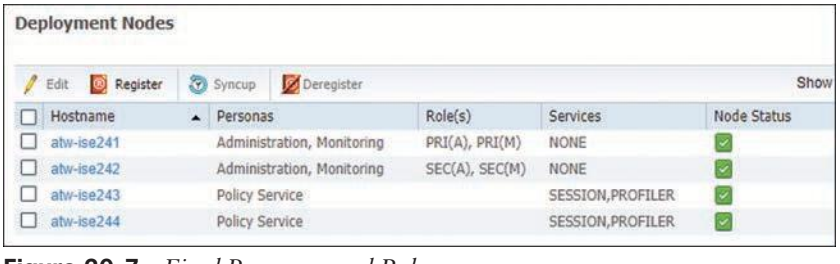


Figure 20-7 *Final Personas and Roles*

Understanding the High Availability Options Available

There are many different items to note when it comes to high availability (HA) in an ISE deployment. There are concerns about communication between the PANs and the other ISE nodes for database replication and synchronization, as well as about communication between the PSNs and MnT nodes for logging.

The ISE nodes themselves are not the only issue. There is also the issue of authentication sessions from the NAD reaching the PSNs in the event of a WAN outage. In addition, a NAD might recognize that a PSN may no longer be active and might send authentication requests to the active PSN instead.

Primary and Secondary Nodes

PANs and MnT nodes both have the concepts of primary and secondary nodes, but they operate very differently. Let's start with the easier type first: MnT nodes.

Monitoring and Troubleshooting Nodes

An MnT node is responsible for the logging and reporting functions of ISE. All PSNs send their logging data to the MnT node as syslog messages (UDP/20514).

When there are two MnT nodes in an ISE deployment, all ISE nodes send their audit data to both MnT nodes at the same time. Figure 20-8 illustrates the logging flows between the ISE nodes.

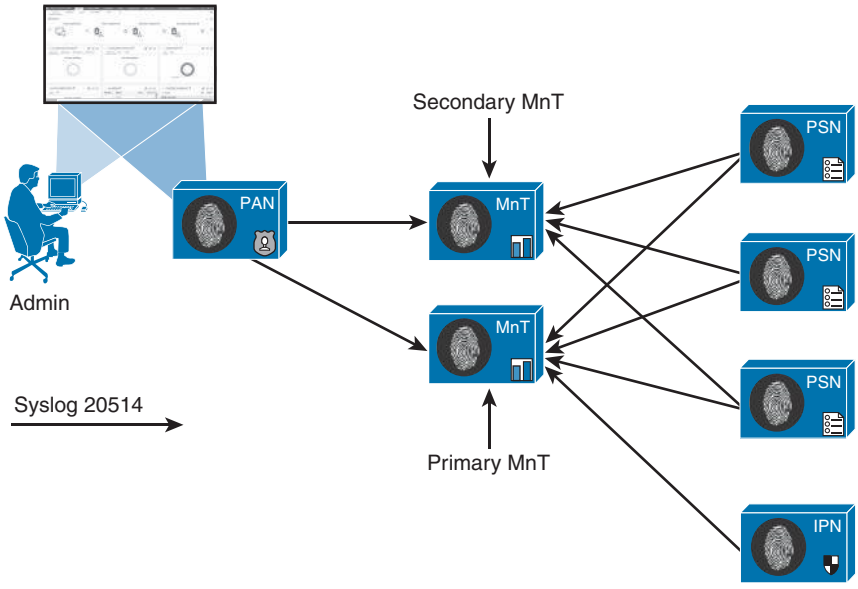
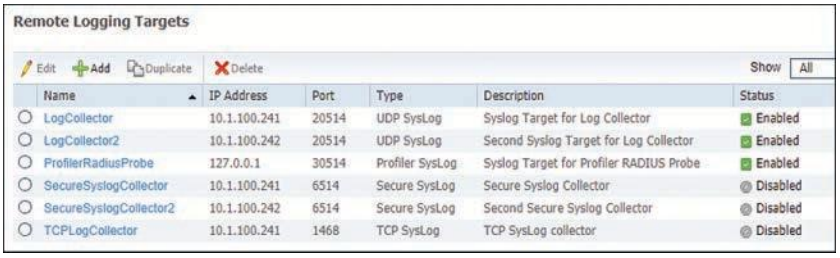


Figure 20-8 *Logging Flows*

The active/active nature of the MnT nodes can be viewed easily in the administrative console under Administration > System > Logging > Remote Logging Targets. In this screen you see the two MnT nodes defined as LogCollector and LogCollector2, as shown in Figure 20-9.

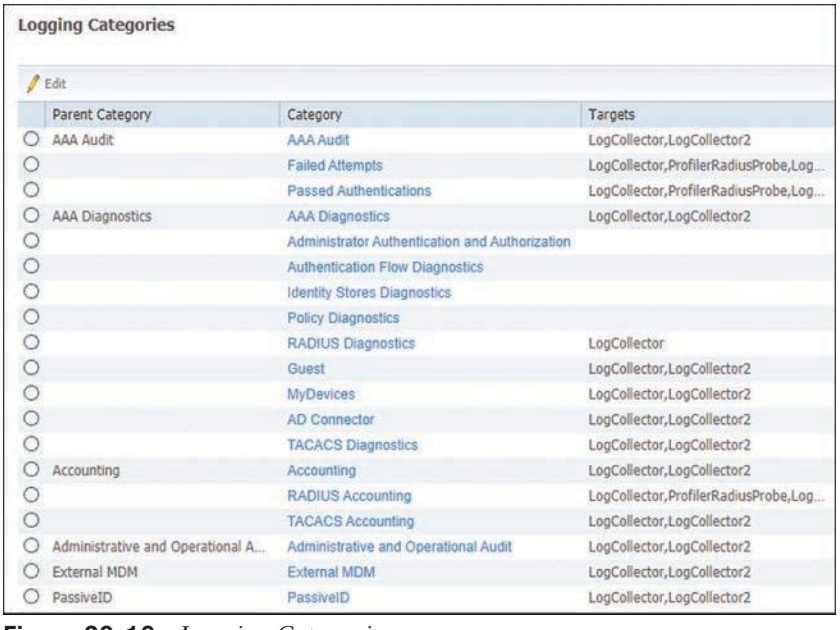


Name	IP Address	Port	Type	Description	Status
LogCollector	10.1.100.241	20514	UDP SysLog	Syslog Target for Log Collector	Enabled
LogCollector2	10.1.100.242	20514	UDP SysLog	Second Syslog Target for Log Collector	Enabled
ProfilerRadiusProbe	127.0.0.1	30514	Profiler SysLog	Syslog Target for Profiler RADIUS Probe	Enabled
SecureSyslogCollector	10.1.100.241	6514	Secure SysLog	Secure Syslog Collector	Disabled
SecureSyslogCollector2	10.1.100.242	6514	Secure SysLog	Second Secure Syslog Collector	Disabled
TCPLogCollector	10.1.100.241	1468	TCP SysLog	TCP SysLog collector	Disabled

Figure 20-9 *Logging Targets*

Key Topic

Notice under Administration > System > Logging > Logging Categories that both LogCollector and LogCollector2 are automatically listed as targets for the logs (see Figure 20-10). This ensures that all ISE nodes in the cube send logs to both MnT nodes.



Parent Category	Category	Targets
AAA Audit	AAA Audit	LogCollector,LogCollector2
	Failed Attempts	LogCollector,ProfilerRadiusProbe,Log...
	Passed Authentications	LogCollector,ProfilerRadiusProbe,Log...
AAA Diagnostics	AAA Diagnostics	LogCollector,LogCollector2
	Administrator Authentication and Authorization	
	Authentication Flow Diagnostics	
	Identity Stores Diagnostics	
	Policy Diagnostics	
	RADIUS Diagnostics	LogCollector
	Guest	LogCollector,LogCollector2
	MyDevices	LogCollector,LogCollector2
	AD Connector	LogCollector,LogCollector2
	TACACS Diagnostics	LogCollector,LogCollector2
Accounting	Accounting	LogCollector,LogCollector2
	RADIUS Accounting	LogCollector,ProfilerRadiusProbe,Log...
	TACACS Accounting	LogCollector,LogCollector2
Administrative and Operational A...	Administrative and Operational Audit	LogCollector,LogCollector2
External MDM	External MDM	LogCollector,LogCollector2
PassiveID	PassiveID	LogCollector,LogCollector2

Figure 20-10 *Logging Categories*

Upon an MnT failure, all nodes continue to send logs to the remaining MnT node so that no logs are lost. The PAN retrieves all logs and reports data from the secondary MnT node, so there is no administrative function loss, either. However, the log database is not synchronized between the primary and secondary MnT nodes; therefore, when the MnT node returns to service, a backup and restore of the monitoring node is required to keep the two MnT nodes in complete synchronization.

NOTE The best practice for logging is to also send logging data to a security information management (SIM) tool for long-term data archival and reporting.

Policy Administration Nodes

A PAN is responsible not only for providing an administrative GUI for ISE but also the mission-critical function of database synchronization of all ISE nodes. Every ISE node maintains a full copy of the database, and the master database exists on the primary PAN.

A PSN might receive data about a guest user, and when that occurs, it must sync that data to the primary PAN. The primary PAN then synchronizes that data out to all the ISE nodes in the deployment.

Because the functionality is so arduous, and having only a single source of truth for the data in the database is so critical, failing over to the secondary PAN is usually a manual process. If the primary PAN goes offline, no synchronizations occur until the secondary PAN is promoted to primary. Once it becomes the primary, it takes over all synchronization responsibility. This is sometimes referred to as a "warm spare" type of HA.

Promoting the Secondary PAN to Primary

To promote the secondary PAN to primary, you need to connect to the GUI on the secondary PAN. You should immediately notice that the UI is much more limited than the UI of a standalone or primary PAN. From this limited UI, perform the following steps:

Step 1. Choose **Administration > System > Deployment**.

Step 2. Click **Promote to Primary**. Figure 20-11 shows the Promote to Primary option available on the secondary node.

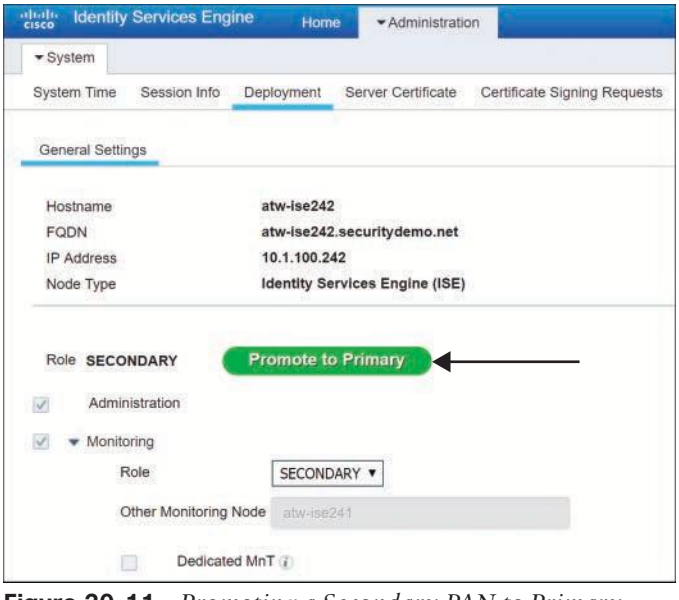


Figure 20-11 *Promoting a Secondary PAN to Primary*

Auto PAN Switchover

An automated promotion function was added to ISE beginning with Version 1.4. There must be two PANs and at least one other type of node in the deployment.

The node that is not a PAN uses a health check function for the PAN(s) and probes the PAN at specified intervals. The health check promotes the secondary PAN when the primary fails a configurable number of probes. Once the original secondary node is promoted, it is probed. Figure 20-12 illustrates the process.

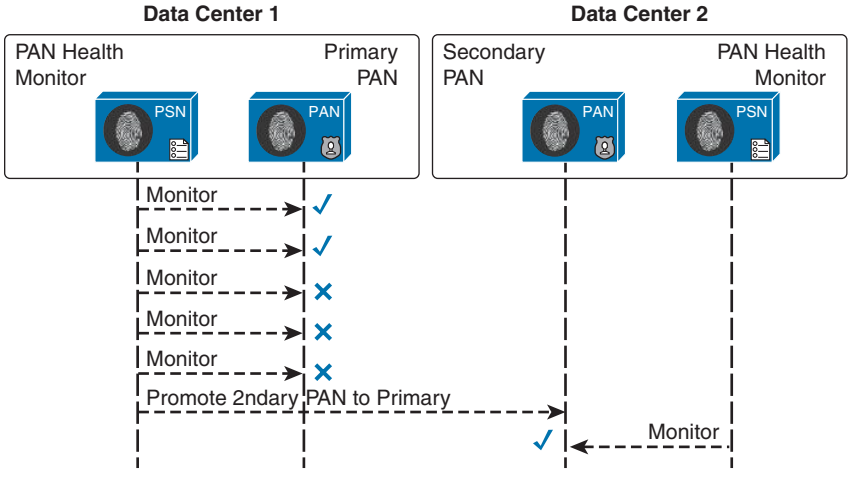


Figure 20-12 *Promoting a Secondary PAN to Primary*

NOTE As of ISE Version 2.7, there is no ability to automatically sync the original primary PAN back into the ISE cube. That is still a manual process.

Configuring Automatic Failover for the Primary PAN

In order for the configuration to be available, there must be two PANs and at least one node that is not a PAN in the deployment.

Follow these steps in the UI of the primary PAN:

Step 1. Navigate to **Administration > System > Deployment**.

Step 2. Click on **PAN Failover** on the left-hand side.

Step 3. Select the **Enable PAN Auto Failover** checkbox.

Step 4. Select the health check nodes from the dropdowns. Notice in Figure 20-13 that the UI lists the primary PAN and secondary PAN to the right of the selected health check nodes.

Step 5 Set the polling interval. The interval is in seconds and can be set between 30 and 300 seconds (5 minutes).

Step 6. Enter the number of failed probes that can occur before failover is initiated. The valid range is anywhere from 2 to 60 consecutive failed probes.

Step 7. Click **Save**.

Figure 20-13 shows the PAN configured for automatic failover.

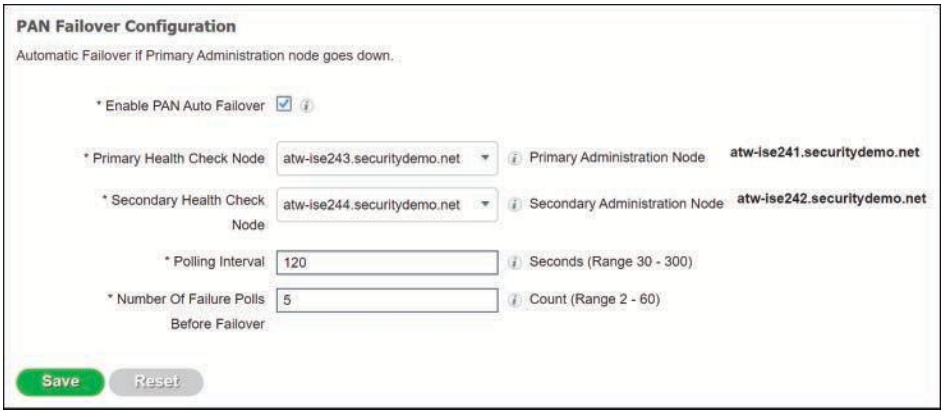


Figure 20-13 *PAN Failover*

Licensing in a Multi-Node ISE Cube

ISE uses a centralized pool of licenses, and all nodes within an ISE cube share a single set of licenses.

For redundancy to work properly, you need to ensure that licenses are issued with both the primary and secondary PANs listed. When requesting licenses, you should supply three pieces of information about each PAN: the product ID (SPID), version number (VPID), and serial number.

Example 20-2 shows how you log in to an ISE node and use the **show udi** command to obtain this information from the primary PAN (atw-ise241). Example 20-3 repeats the same command on the secondary PAN (atw-ise242).

Example 20-2 *show udi Command and Output from the Primary PAN*

```
atw-ise241/admin# sho udi

SPID: ISE-VM-K9
VPID: V01
Serial: AKGHJ7I48DQ

atw-ise241/admin#
```

Example 20-3 *show udi Command and Output from the Secondary PAN*

```
atw-ise242/admin# sho udi

SPID: ISE-VM-K9
VPID: V01
Serial: A59HJHJBK9N

atw-ise242/admin#
```

Ensuring that the license is created with both PANs in the license key helps to ensure there are no issues in the event of a disaster involving the loss of the primary PAN.

Node Groups

PSNs do not necessarily need to have a high availability configuration. Every ISE node maintains a full copy of the database, and the NADs have their own detection of a "dead" RADIUS server, which triggers the NAD to send AAA communication to the next RADIUS server in the list.



However, ISE has the concept of a ***node group***. A node group is made up of PSNs, where the PSNs maintain a heartbeat with each other. Beginning with ISE Version 1.3, the PSNs might be in different subnets, or they might be Layer 2 adjacent. In older versions, the PSNs required the use of multicast, but starting in Version 1.3, they used direct encrypted TCP-based communication instead:

- **TCP/7800**: Used for peer communication
- **TCP/7802**: Used for failure detection

If a PSN goes down and orphans a URL redirection session, one of the other PSNs in the node group sends a Change of Authorization (CoA) to the NAD so that the endpoint can restart the session establishment with a new PSN.

Node groups have another function, which is entirely related to data replication. ISE used a serial replication model in ISE 1.0, 1.1, and 1.1.x. This means that all data had to go through the primary PAN, which sent the data objects to each other node and waited for an acknowledgment for each piece of data before sending the next one in line.

In ISE Version 1.2 and later, ISE uses a common replication framework known as JGroups. One of the benefits of JGroups is that local replications are possible with local peers, and node groups are used to define them.

So, when a member of a node group learns endpoint attributes (through profiling), it is able to send the information directly to the other members of the node group. However, when that data needs to be replicated globally (to all PSNs), the JGroups communication must still go through the primary PAN, which replicates the communication to all the other PSNs.

Node groups are most commonly used when deploying the PSNs behind a load balancer; however, there is no reason node groups could not be used with regionally located PSNs. You would not want to use a node group with PSNs that are geographically and logically separate.

To configure a node group in the ISE GUI, perform the following steps:

Step 1. Choose **Administration > System > Deployment**.

Step 2. In the Deployment pane on the left side of the screen, click the cog icon and choose **Create Node Group**, as shown in Figure 20-14.

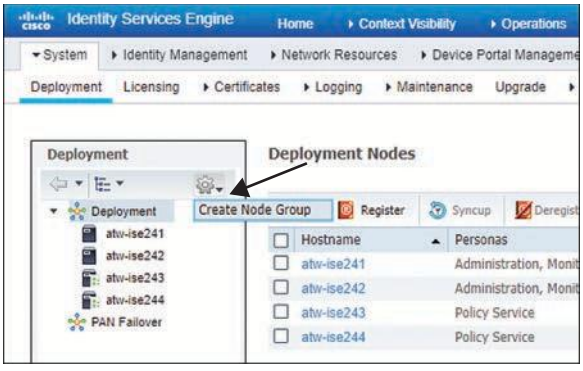


Figure 20-14 *Choosing to Create a Node Group*

Step 3. In the Create Node Group screen, shown in Figure 20-15, enter a name for the node group in the Node Group Name field. Ensure that the name describes the location of the group.

Step 4. In the Description field, enter a more detailed description that helps identify exactly where the node group is (for example, NodeGroup for Data Center 1).

Step 5. In the Multicast Address field, enter a multicast address for keepalive communication. As indicated below the field, you need to make sure you do not enter a multicast address that is already in use.

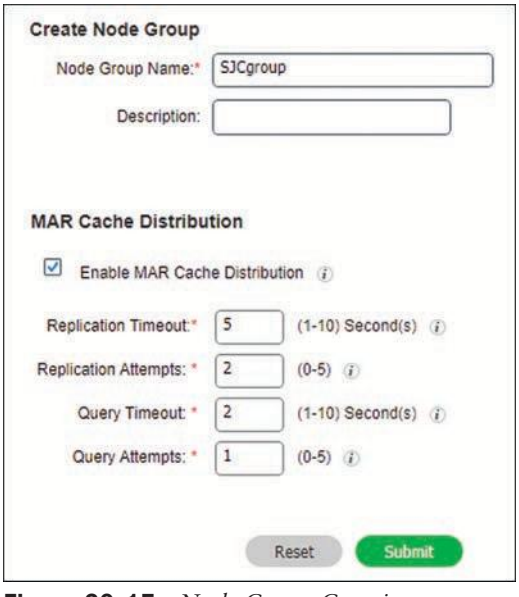


Figure 20-15 *Node Group Creation*

Step 6. Click **OK** in the success popup window, as shown in Figure 20-16. Also notice the appearance of the node group in the left pane.

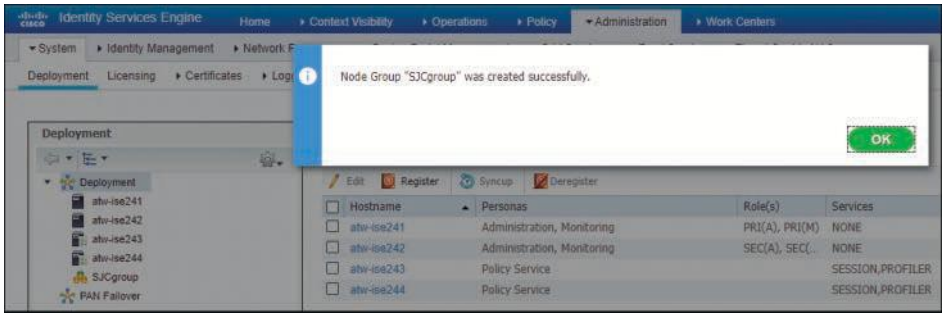


Figure 20-16 *Success Popup*

Add the Policy Services Nodes to the Node Group

To add the PSNs to the node group, in the ISE GUI, perform the following steps:

Step 1. Choose **Administration > System > Deployment**.

Step 2. Select one of the PSNs to add to the node group.

Step 3. From the Include Node in Node Group dropdown, select the newly created group, as shown in Figure 20-17.

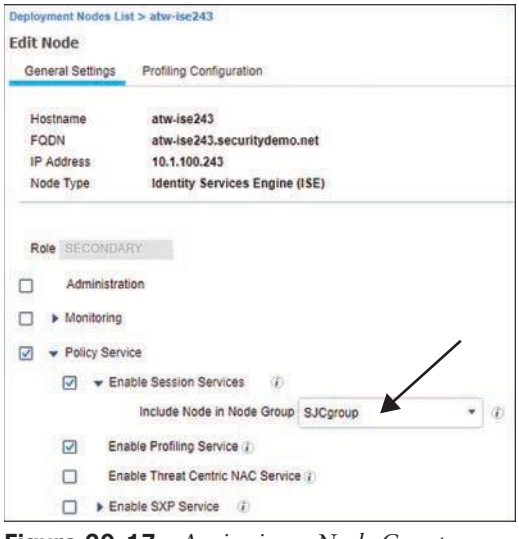


Figure 20-17 *Assigning a Node Group*

Step 4. Click **Save**.

Step 5. Repeat steps 1 through 4 for each PSN that should be part of the node group.

Figure 20-18 shows the reorganization of the PSNs within the node group in the navigation pane on the left-hand side of the Deployment screen.

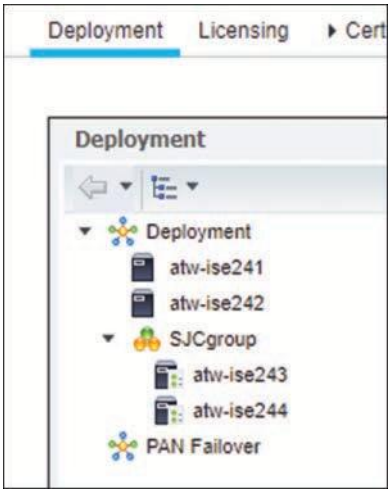


Figure 20-18 *Reorganized Deployment Screen Navigation Pane*

Using Load Balancers

One high-availability option that has skyrocketed in popularity for Cisco ISE deployments is the use of load balancers. The use of load balancers with ISE deployments has increased in part because it can greatly simplify larger deployments.

Figure 20-19 illustrates how the NADs only have to be configured with one IP address per set of ISE PSNs; this removes a lot of the complexity in the NAD configuration. The load balancer itself takes care of monitoring the ISE PSNs and removing them from service if they are down and allows you to scale more nodes behind the virtual IP address without ever touching the network device configuration again.

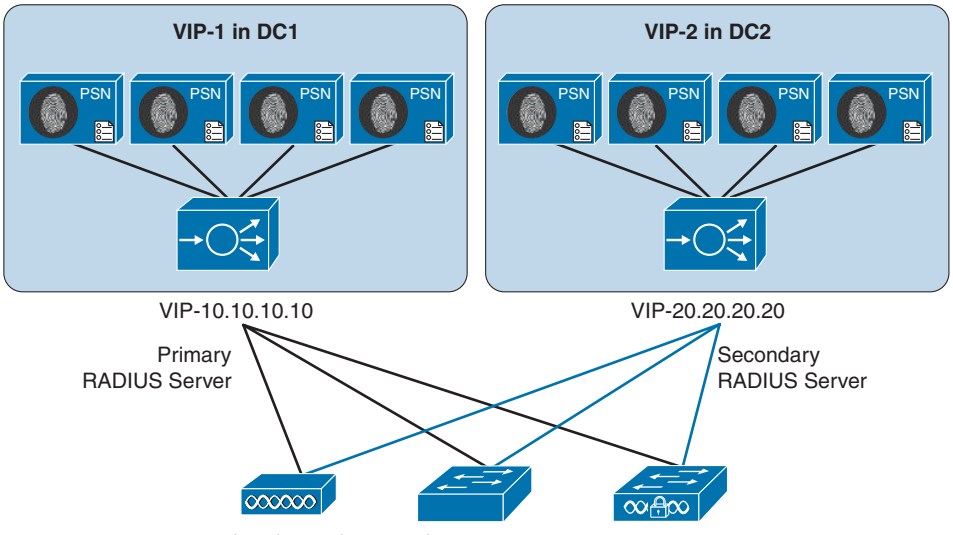


Figure 20-19 *Load-Balanced PSN Clusters*

NOTE While Craig Hyps was a principal technical marketing engineer at Cisco, he wrote what is considered to be the definitive guide on load balancing with ISE. The guide was written using F5 load balancers, but the principles are the same, no matter which load balancer you choose to implement. Instead of replicating that entire large and detailed guide in this chapter, the following sections focus on the basic principles that must be followed when using ISE with load balancers. If you want to read the complete document, simply search the Internet for "How To: Cisco and F5 Deployment Guide-ISE Load Balancing Using BIG-IP."

General Guidelines

ISE deployments are not exactly like typical RADIUS server deployments. With traditional RADIUS deployments, the traffic only flows from the NAD to a RADIUS server, and a CoA message might be sent from the RADIUS server to the NAD.

ISE is sort of like a next-generation RADIUS server that offers a more powerful and complex model. ISE requires the endpoint to also be able to reach the PSN for purposes of posture communication and Centralized Web Authentication. Therefore, if you want to deploy with load balancers, you must ensure the following:

- Each PSN must be reachable by the PAN/MnT node directly, without having to go through Network Address Translation (NAT). This sometimes is referred to as *routed mode* or *pass-through mode*.
- Each PSN must also be reachable directly from the endpoint.
  - When the PSN sends a URL redirection to the NAD, it uses the fully qualified domain name from the configuration, not the virtual IP (VIP) address.
  - You might want to use subject alternative names (SANs) in the certificate to include the FQDN of the load balancer's VIP address.
- The same PSN is used for the entire session. User persistence ("stickiness") needs to be based on the calling station ID.
- The VIP address gets listed as the RADIUS server of each NAD for all 802.1X-related AAA.
  - AAA includes both authentication and accounting packets.
  - Some load balancers use a separate VIP address for each protocol type.
- The list of RADIUS servers allowed to perform dynamic authorizations (also known as CoAs) on the NAD should use the real IP addresses of the PSNs, not the VIP address. The VIP address could be used for the CoAs if the load balancer is performing Source NAT (SNAT) for the CoAs sent from the PSNs.
- The load balancers are listed as NADs in ISE so that their test authentications can be answered. ISE uses a device's Layer 3 address to identify the NAD rather than the NAS IP address in the RADIUS packet; this is another reason to avoid SNAT for incoming RADIUS requests.

- Load balancers should be configured to use test probes to ensure that the PSNs are still alive and well.
  - A probe should be configured to ensure that RADIUS is responding.
  - HTTPS should also be checked.
  - If either probe fails, the PSN should be taken out of service.
  - A PSN must be marked dead and taken out of service in the load balancer before the NAD's built-in failover occurs.

Failure Scenarios

If a single PSN fails, the load balancer takes that PSN out of service and spreads the load over the remaining PSNs. When the failed PSN is returned to service, the load balancer adds it back to the rotation. By using node groups along with a load balancer, another member of the node group issues a CoA reauthorization for any sessions that were being established. This CoA causes the session to begin again. At that point, the load balancer directs the new authentication to a PSN that is still alive and still in rotation.

NADs have some built-in capabilities to detect when the configured RADIUS server is dead and to automatically fail over to the next RADIUS server configured. When using a load balancer, the RADIUS server IP address is actually the VIP address. So, if the entire VIP address is unreachable (for example, if the load balancer has died), the NAD should quickly fail over to the next RADIUS server in the list. That defined RADIUS server could be another VIP address in a second data center or another backup RADIUS server; the options are quite flexible.

Anycast High Availability for ISE PSNs

This section exists thanks to one of the most talented and gifted technologists roaming the earth today. E. Peter Karelis, CCIE #8068, designed the high availability solution discussed in this section for a small ISE deployment with two data centers. The network architecture is illustrated in Figure 20-20.

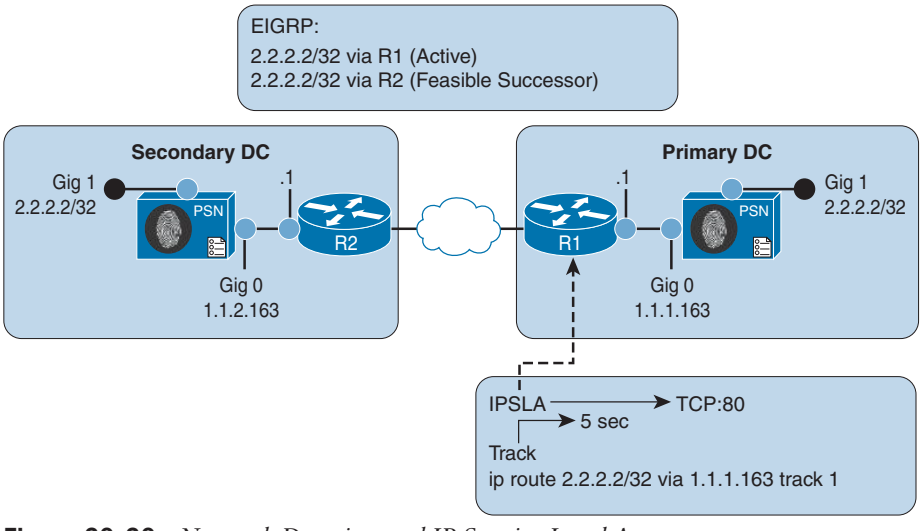


Figure 20-20 *Network Drawing and IP Service-Level Agreement*

Anycast is a networking technique in which exactly the same IP address exists in multiple places in the network. In this case, the same IP address (2.2.2.2) is assigned to the Gig1 interfaces on all the PSNs. Each default gateway (router) in each data center is configured with a static route to 2.2.2.2/32, with the Gig0 IP address of the PSN as the next hop. Those static routes are redistributed into the routing protocol—in this case, EIGRP. Anycast relies on the routing protocols to ensure that traffic destined to the Anycast address (2.2.2.2) is sent to the closest instance of that IP address.

When Anycast is set up to route 2.2.2.2 to the ISE PSN, EIGRP metrics can ensure that all routes prefer the primary data center, with the secondary data center route listed as the feasible successor (FS). With EIGRP, there is less than a one-second delay when a route (the successor) is replaced with the backup route (the feasible successor).

How do you make the successor route drop from the routing table when the ISE node goes down? You can configure an IP service-level agreement (SLA) on the router that checks the status of the HTTP service on the ISE PSN in the data center every five seconds. If the HTTP service stops responding on the active ISE PSN, the route is removed, and the feasible successor takes over, causing all the traffic for 2.2.2.2 to be sent to the PSN in the secondary data center. Figure 20-21 illustrates the IP SLA function; when it occurs, the only route left in the routing table is to the router at the secondary data center.

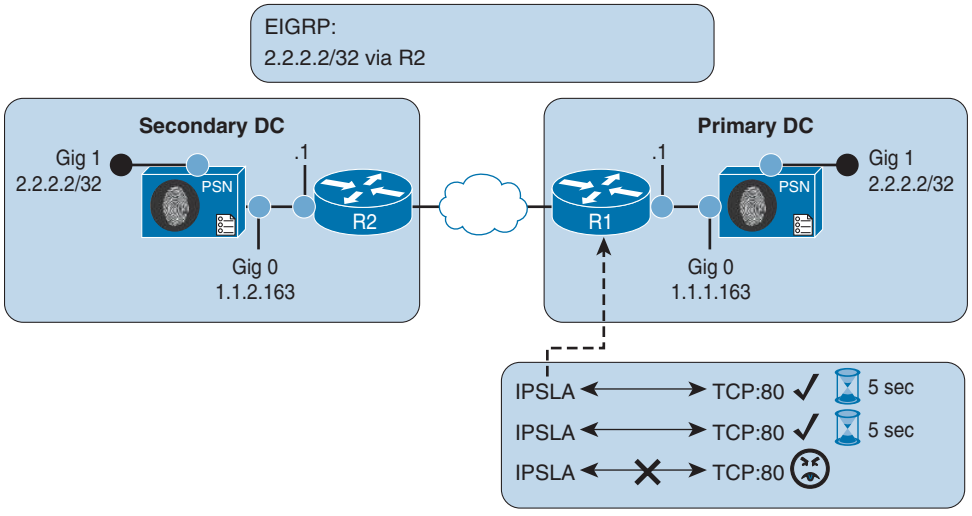


Figure 20-21 *IP SLA in Action*

All network devices are configured to use the Anycast address (2.2.2.2) as the only RADIUS server in their configuration. The RADIUS requests are always sent to whichever ISE node is active.

Example 20-4 shows the interface configuration on the ISE PSN. The Gig0 interface is the actual routable IP address of the PSN, and Gig1 is in a VLAN to nowhere, using the Anycast IP address.

Example 20-4 *ISE Interface Configuration*

```
interface gig 0
  !Actual  IP of Node
  ip address 1.1.1.1 255.255.255.0
interface gig 1
  !Anycast VIP assigned to all PSN nodes on G1
  ip address 2.2.2.2 255.255.255.255

ip default-gateway [Real Gateway for Gig0]
!note no static routes needed.
```

Example 20-5 shows the IP SLA configuration on the router to test port 80 on the PSN every 5 seconds and to time out after 1000 msec. When that timeout occurs, the router is removed.

Example 20-5 *IP SLA Configuration*

```
ip sla 1
  !Test TCP to port 80 to the actual IP of the node.
  !"control disable" is necessary, since you are connecting to an
  !actual host instead of an SLA responder

  tcp-connect 1.1.1.1 80 control disable
  ! Consider the SLA as down if response takes longer than 1000msec

    threshold 1000
    ! Timeout after 1000 msec.
    timeout 1000
    !Test every 5 Seconds:
    frequency 5

ip sla schedule 1 life forever start-time now
track 1 ip sla 1
ip route 2.2.2.2 255.255.255.255 1.1.1.1 track 1
```

Example 20-6 shows the route redistribution configuration where the EIGRP metrics are applied. Pete was able to use the metrics that he chose specifically because he was very familiar with his network. His warning to others attempting the same thing is to be familiar with your network or to test thoroughly when identifying the metrics that might work for you.

Example 20-6 *Route Redistribution*

```
router eigrp [Autonomous-System-Number]
  redistribute static route-map STATIC-TO-EIGRP

route-map STATIC-TO-EIGRP permit 20
  match ip address prefix-list ISE_VIP
  !Set metrics correctly
  set metric 1000000 1 255 1 1500

ip prefix-list ISE_VIP seq 5 permit 2.2.2.2/32
```

Anycast and IP SLAs are not magical features of ISE. These are normal networking tools and techniques used with Cisco networks all over the world. However, this section shows a very useful way to leverage the existing technology and capability of the Cisco infrastructure to enable a very useful form of high availability in a small deployment. Other deployments have used Anycast with the VIP addresses of load balancers physically located in different data centers to accomplish the same goal.

IOS Load Balancing

Cisco network devices have a lot of intelligence built into them to aid in an intelligent access layer for policy and policy enforcement. One such intelligence level is the ability to perform local load balancing of RADIUS servers. This does not mean using a Cisco switch as a server load balancer instead of a dedicated appliance. Rather, it refers to the ability of the access layer switch to balance the outbound authentication requests for endpoints that are authenticated to the switch, as illustrated in Figure 20-22.

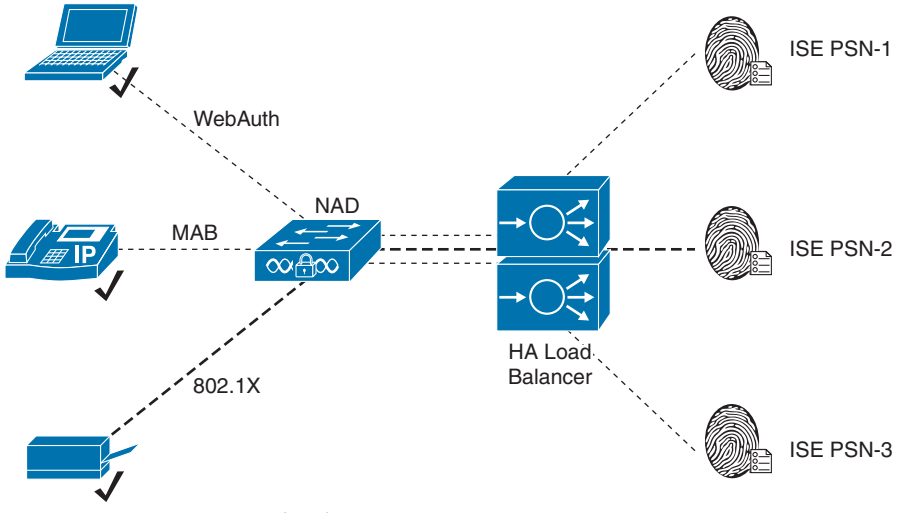


Figure 20-22 *IOS Load Balancing*

Enabling IOS RADIUS server load balancing requires only one additional command. After all the PSNs are defined as AAA servers in the switch, you use the **radius-server load-balance** global configuration command to enable RADIUS server load balancing.

Example 20-7 shows the use of a **show** command to verify that multiple ISE servers are configured.

Example 20-7 *Verifying That All ISE PSNs Are Configured on a Switch*

```
3750-X# show aaa server | incl host
RADIUS: id 4, priority 1, host 10.1.100.232, auth-port 1812, acct-port 1813
RADIUS: id 5, priority 2, host 10.1.100.233, auth-port 1812, acct-port 1813
RADIUS: id 6, priority 3, host 10.1.100.234, auth-port 1812, acct-port 1813
```

Example 20-8 shows how to enable IOS load balancing on a switch.

Example 20-8 *Enabling IOS Load Balancing*

```
3750-X(config)# radius-server load-balance method least-outstanding batch-size 5
```

Maintaining ISE Deployments

While having a distributed deployment and load-balanced architecture are certainly critical items when scaling a deployment and ensuring high availability, there are critical basic maintenance items that should always be considered to ensure the most uptime and stability, including patching and backup and restore strategies.

Patching ISE

ISE patches are released roughly once per month. These patches contain bug fixes as well as security fixes when necessary. (Think about the Heartbleed and Poodle vulnerabilities that were discovered with SSL.) To ensure that bug fixes are applied, security vulnerabilities are plugged, and a solution will work as seamlessly as possible, it is important to have a planned patching strategy.

You can download patches from Cisco.com by going to **Downloads > Products > Security > Access Control and Policy > Identity Services Engine > Identity Services Engine Software**, as shown in Figure 20-23.

NOTE This book was written using ISE Version 2.7, which at the time this book was written did not have any patches yet. Therefore, Figures 20-23 and 20-24 show ISE Version 2.6, and the example in Figures 20-25 through 20-27 use ISE Version 2.1.

Search the list of software available for your specific version of ISE. Figure 20-24 illustrates the naming convention for ISE patches. Cisco ISE patches are normally cumulative, so, for example, Version 1.2 patch 12 includes all the fixes in patches 1 through 11 as well.

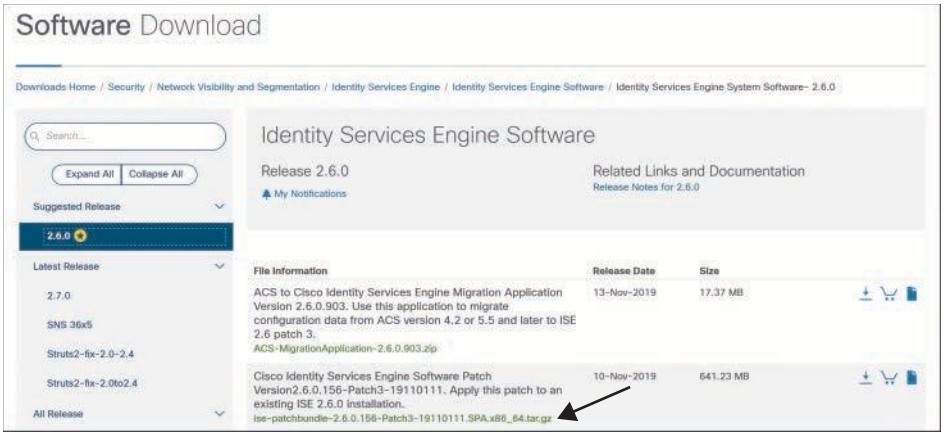


Figure 20-23 *ISE Downloads Page*

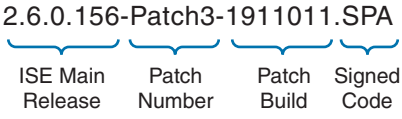


Figure 20-24 *Anatomy of ISE Patch Nomenclature*

To install patches in an ISE cube, follow these steps:

Step 1. Download the required patch.

Step 2. In the ISE GUI, navigate to **Administration > System > Maintenance > Patch Management**.

Step 3. Click the **Install** button, as shown in Figure 20-25.

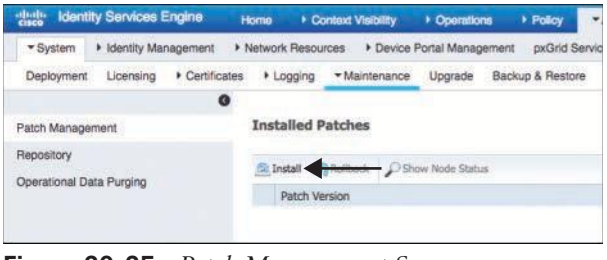


Figure 20-25 *Patch Management Screen*

Step 4. Click **Browse**, select the downloaded patch, and click **Install**, as shown in Figure 20-25.

The patch bundle is then uploaded through the web UI. As the patch is installed on the PAN, you are logged out of the GUI, and the patch is distributed from the PAN to all nodes in the ISE cube. The patch is applied to each node in the cube, one at a time, in alphabetic order.

You can log back in to the PAN when it's finished restarting services or rebooting. Use the Show Node Status button shown in Figure 20-26 to verify the progress of the patching. Figure 20-27 shows the resulting status of each node's progress for the patch installation.

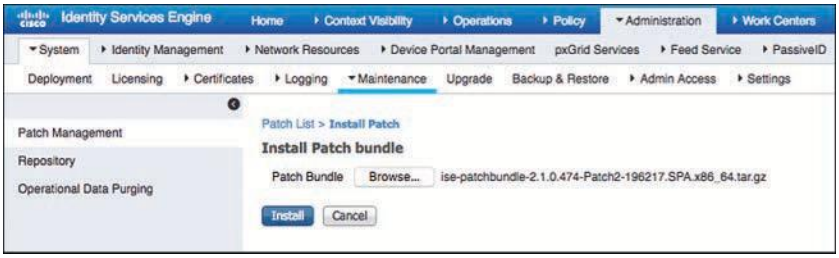


Figure 20-26 *Installing the Selected Patch*

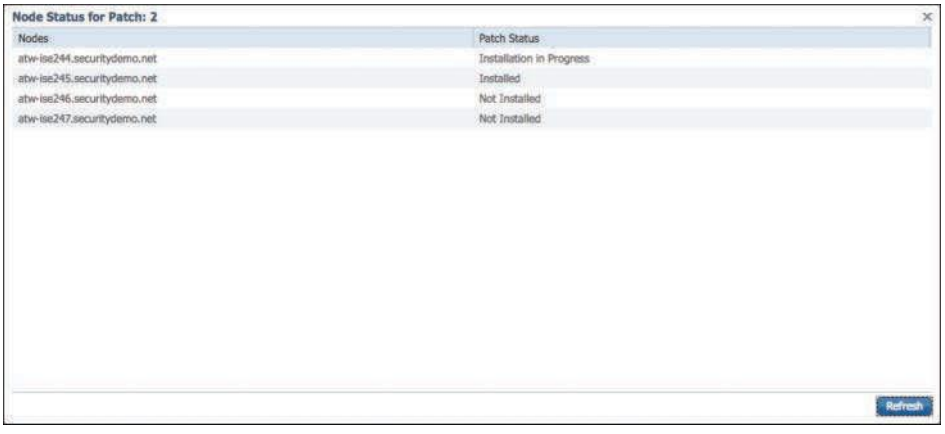


Figure 20-27 *Node Status*

NOTE Even though ISE may be deployed with high availability options, it is always recommended to save patching for an approved change window.

Backup and Restore

Another key strategy in assuring the availability of ISE in the environment is having a solid backup strategy. There are two types of ISE backups: configuration and operational backups. These two types are most easily related to backing up the product databases (configuration) from the PAN and backing up of the MnT data (operational).

Figure 20-28 shows the backup screen in ISE, which you reach by navigating to Administration > System > Backup & Restore.

As you can see in Figure 20-28, the backups are stored in a repository. They can be restored from this repository. Backups can be scheduled or run on demand. The status of a backup may be viewed from the GUI or the CLI, but the status of a restore may only be viewed from the CLI.

Figure 20-29 shows the creation of a configuration backup stored in a repository named FTP, and Figure 20-30 shows the backup progress in the UI.

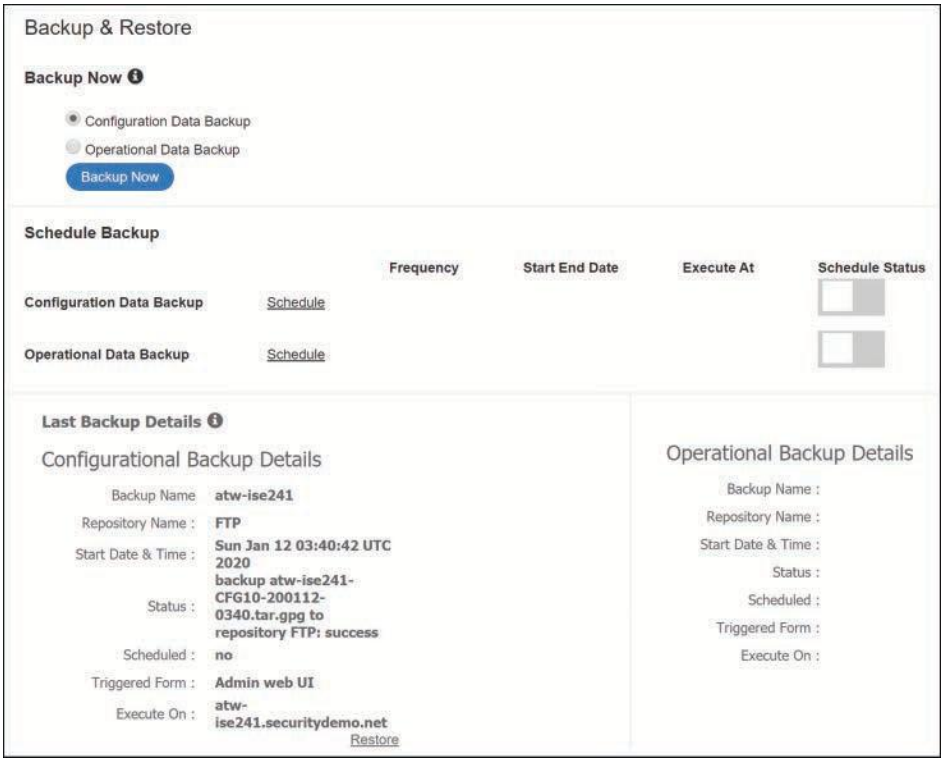


Figure 20-28 *Backup & Restore Screen*

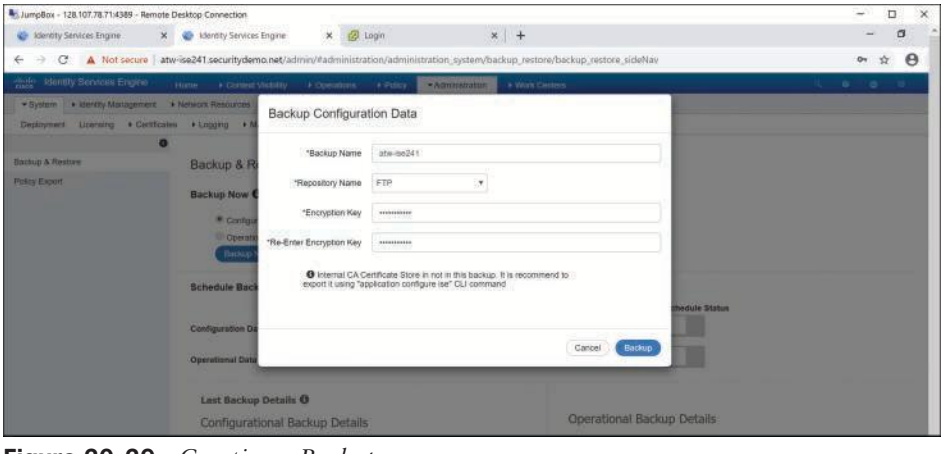


Figure 20-29 *Creating a Backup*

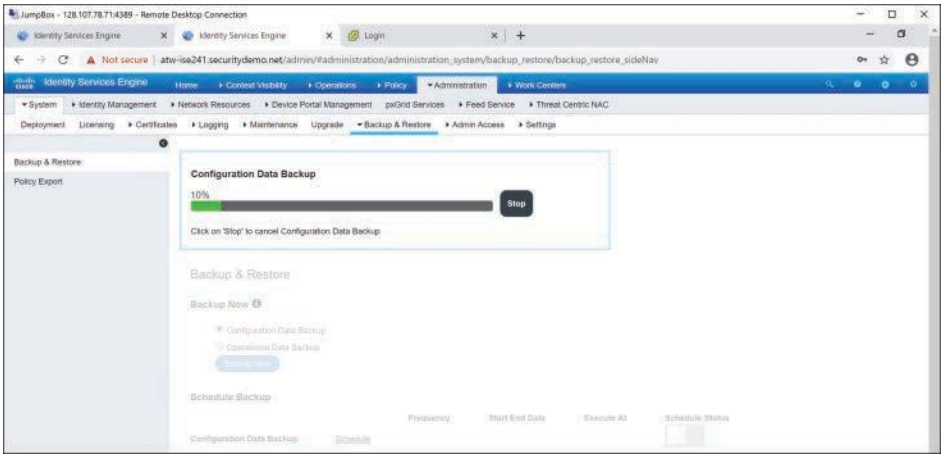


Figure 20-30 *Backup Progress*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 20-4 lists these key topics and the page number on which each is found.

Key Topic

Table 20-4 Key Topics

Key Topic Element	Description	Page
Paragraph	Ensuring that all ISE nodes in the cube send logs to multiple MnT nodes	744
Paragraph	Node groups	748

Define Key Term

Define the following key term from this chapter and check your answers in the glossary:

node group

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is the name of the technique in which a single IP address is owned by multiple network endpoints?
2. What common load-balancing technique should be avoided when using load balancers with Cisco ISE PSNs?
3. What are the two backup types in ISE?
4. Which node's UDI information should be used when generating an ISE license to ensure that failover occurs smoothly?
5. What is the prerequisite for joining an ISE node to an ISE cube?

This page intentionally left blank



CHAPTER 21

Troubleshooting Tools

This chapter covers the following topics:

- Logging
- Diagnostic tools
- Troubleshooting methodology
- Troubleshooting outside of ISE

Throughout this book, you have been verifying activities by using tools, such as Live Log, that are commonly used for troubleshooting. Such tools provide immeasurable value to an ISE administrator when attempting to identify and remediate the cause of an issue.

This chapter takes a deeper dive into the troubleshooting of secure access control and examines some of the tools that are built into ISE to aid you in your troubleshooting efforts.

Troubleshooting a product can sometimes be complex. Troubleshooting a solution made up of multiple products is bound to be downright difficult. The biggest tip we can provide for you is to stay calm and think through the flows. When you are comfortable with the solution and how the parts work together, you can more easily find and fix problems.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 21-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 21-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Logging	3–6
Diagnostic Tools	1, 2
Troubleshooting Methodology	7
Troubleshooting Outside of ISE	8–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What ISE diagnostic tool can be used to find misconfigurations in a Cisco NAD?
   - **a.** TCP Dump
   - **b.** Live Sessions
   - **c.** RADIUS Authentication Troubleshooting tool
   - **d.** Evaluate Configuration Validator
2. What ISE diagnostic tool can be used to examine different aspects of a session and provide additional details that may not have been available in the detailed authentication report?
   - **a.** TCP Dump
   - **b.** Live Sessions
   - **c.** RADIUS Authentication Troubleshooting tool
   - **d.** Evaluate Configuration Validator
3. Which of the following can be included in an ISE support bundle? (Choose three.)
   - **a.** Full configuration database
   - **b.** Policy configuration
   - **c.** Saved TCP Dump files
   - **d.** Screenshot of Live Log
   - **e.** Debug logs
4. What ISE tool displays a single-line summary view of authentications, Change of Authorizations (CoAs), and state changes of an endpoint through its lifecycle on a network?
   - **a.** Live Log
   - **b.** Live Sessions
   - **c.** RADIUS Authentication Troubleshooting tool
   - **d.** Evaluate Configuration Validator
5. Which ISE tool displays a near-real-time view of passed and failed authentications?
   - **a.** Live Log
   - **b.** Live Sessions
   - **c.** RADIUS Authentication Troubleshooting tool
   - **d.** Evaluate Configuration Validator

6. Which of the following best describes how external syslog servers receive logs from ISE?
   a. Each PSN must be configured locally to send syslog messages to all sources.
   b. It is not possible to configure ISE to log to external logging servers.
   c. The MnT node is configured to forward all received syslog messages to the external recipients.
   d. Each PSN sends syslog messages to the MnT nodes and the external syslog receivers at the same time.
7. Where does an ISE administrator disable all event de-duplication?
   a. Administration > System > Logging > Message Catalog
   b. Administration > System > Protocols > RADIUS
   c. Administration > System > Logging > Remote Logging Targets
   d. Administration > System > Protocols > IEEE 802.1X
8. Which tool gathers all the important log files and combines them into a single bundle for Cisco TAC?
   a. Cisco AnyConnect Network Access Manager (NAM)
   b. Cisco AnyConnect Diagnostic and Reporting Tool (DART)
   c. Cisco NAC Agent
   d. Cisco ISE Agent
9. What are the three main locations to troubleshoot network access authentication?
   a. ISE, firewall, NAD
   b. ISE, endpoint, firewall
   c. ISE, endpoint, NAD
   d. Endpoint, firewall, NAD
10. Which **debug** command provides the best detail to identify why a URL redirection may not be working?
    a. **debug authentication**
    b. **debug epm all**
    c. **debug dot1x all**
    d. **debug aaa all**

Foundation Topics

Logging

The number-one way to troubleshoot any product or solution is to examine the logs. There are numerous ways to examine logs with ISE, including using Live Log, Live Sessions, and individual log files that may be downloaded from the GUI, viewed from the CLI, or combined in a support bundle. The following sections cover these options in detail.

Live Log

ISE has a fantastic tool that is officially called the Live Authentications Log and commonly referred to as Live Log.

Live Log can display 23 columns, and you can hide, display, and reorder the columns as you like. Customizing Live Log to your specific preference will really aid in your ability to use it quickly and efficiently.

Figure 21-1 shows Live Log. The numerals in this figure correspond to the numerals in the following list:

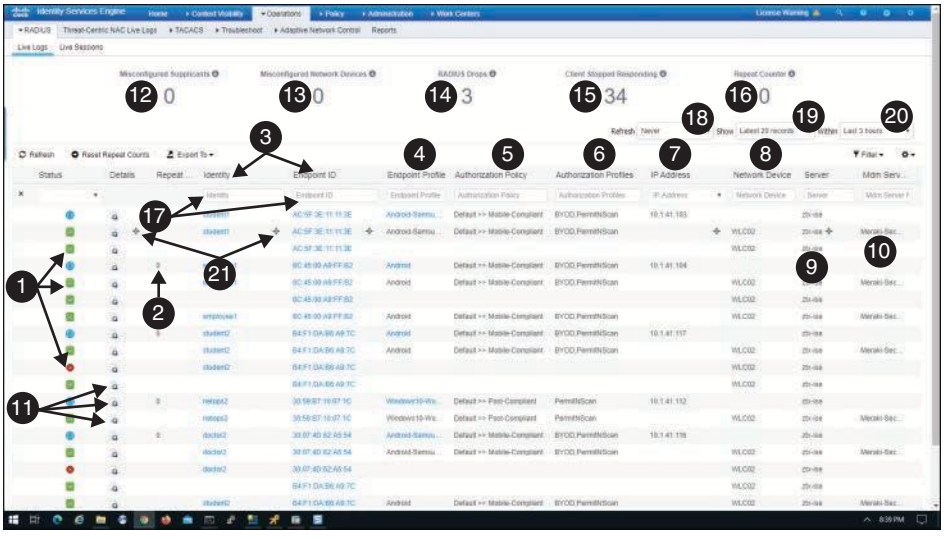


Figure 21-1 *Live Log*

1. **Status**: Live Log has three status types:
   - **Informational**: This status is indicated with a blue icon that indicates that client suppression is enabled (which is on by default). This status increments a repeat count for each pass/fail record that is suppressed.
   - **Pass**: This status is indicated with a green icon that indicates log entries for successful authentications.
   - **Fail**: This status is indicated with a red icon that indicates log entries for failed authentications.
2. **Repeat count**: This number increments only when client suppression is enabled (which is on by default). For each pass/fail record that is suppressed, the counter increments. Only informational record types have a repeat counter.
3. **Identity and Endpoint ID**: These columns show the username from the RADIUS request or WebAuth portal or the MAC address in the case of MAB. The Endpoint ID column is the same as the Calling-Station-Id field from the RADIUS request.
4. **Endpoint Profile**: This column lists the profile assigned to the endpoint involved in the authentication.
5. **Authorization Policy**: This column displays the authorization policy matched during network access. It is displayed in tree form, where the policy set is displayed followed by an arrow (>>) and then the authorization rule (for example, Policy Set >> AuthZ Rule).

6. **Authorization Profiles**: This column shows a comma-separated list of the authorization profiles and SGTs that are assigned from the authorization policy.
7. **IP Address**: This column displays the IP address assigned to the endpoint. The IP address is most often not known until after the Access-Accept response is sent back to the network access device (NAD). Remember that 802.1X occurs at Layer 2 with EAP, which means there is usually not a Layer 3 address assigned until after the link is fully up. ISE learns about that Layer 3 address through the RADIUS accounting message from the NAD, which tells ISE that the session is fully established. This is why the IP address is only visible with the blue informational status types.
8. **Network Device**: This column displays the NAD that performed the network access for the user or endpoint.
9. **Server**: This column shows the ISE PSN that performed the authentication.
10. **MDM Server**: This column lists the mobile device management (MDM) server for the endpoint when the endpoint is registered to one.
11. **Details**: You can click this icon to launch the authentication details report, which contains a tremendous amount of information about the authentication. See the "Authentication Details Report" section of this chapter for more information.
12. **Misconfigured Supplicants**: This counter is used with event suppression, and each time an event is discarded due to a misconfigured endpoint, the counter increases. You can click this counter to bring up a list of dropped events.
13. **Misconfigured Network Devices**: This counter is used to point out when a NAD is configured with a RADIUS accounting update frequency that is considered to be too often. You can click this counter to bring up a list of dropped events.
14. **RADIUS Drops**: When the RADIUS packet is dropped before it is even processed, this counter increments. For example, the packet may be dropped when a NAD is configured with an incorrect shared secret or when a NAD is not defined in the Network Devices section of ISE. You can click this counter to bring up a list of dropped events like the one shown in Figure 21-2.

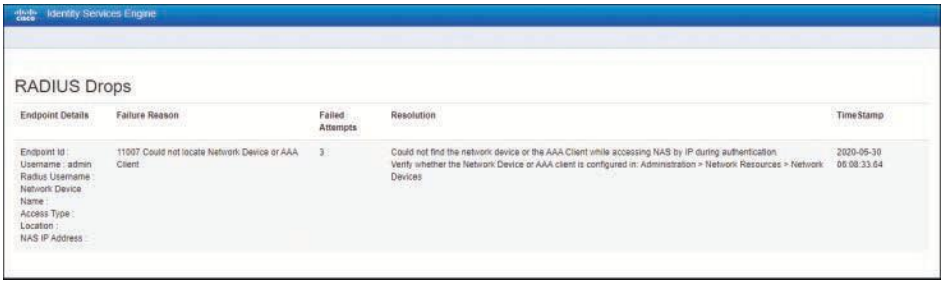


Figure 21-2 *RADIUS Drops*

15. **Client Stopped Responding**: An endpoint may start an authentication but fail to finish it. A client commonly stops responding when it does not trust the EAP certificate presented by ISE. A client may also stop responding if a wireless endpoint does not have a solid connection to the Wi-Fi network. Any time a client starts an EAP authentication but does not complete it, this counter increments. You can click this counter to bring up the list of dropped events like the one shown in Figure 21-3.

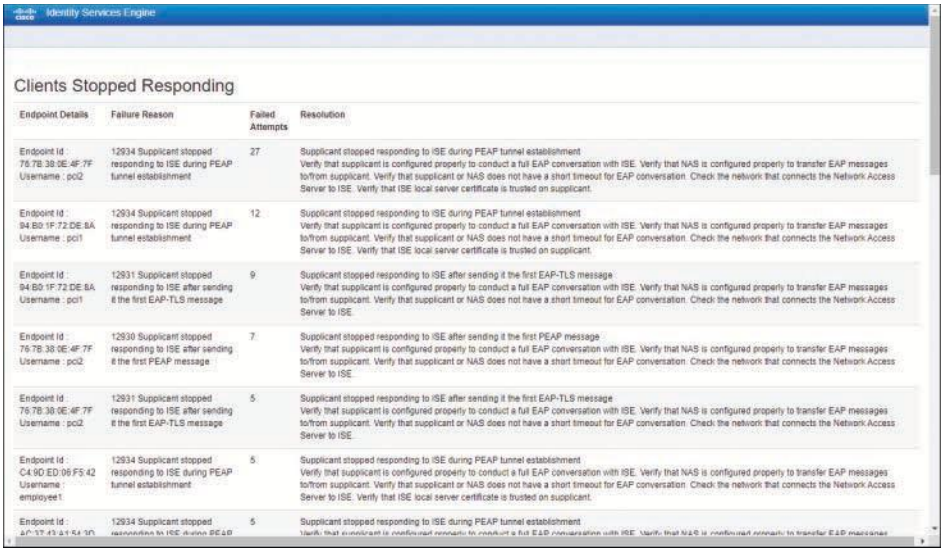


Figure 21-3 *Clients Stopped Responding*

16. **Repeat Counter:** This counter increments each time an event is suppressed and acts as an overall counter for all the individual repeat counts with the informational entries of Live Log (see item 2 in this list). Clicking this counter brings up a consolidated list of dropped events.

17. **Quick Filters:** Like nearly all the other tables in ISE, this table includes a quick filter at the top of every column. By typing a string in any of these columns, you can filter the table for matches on that string. Figure 21-4 shows the use of the quick filter with the string Androi to filter the Live Log entries to those who have Android in the profile name. You can click the X pointed out in the figure to clear the filter.

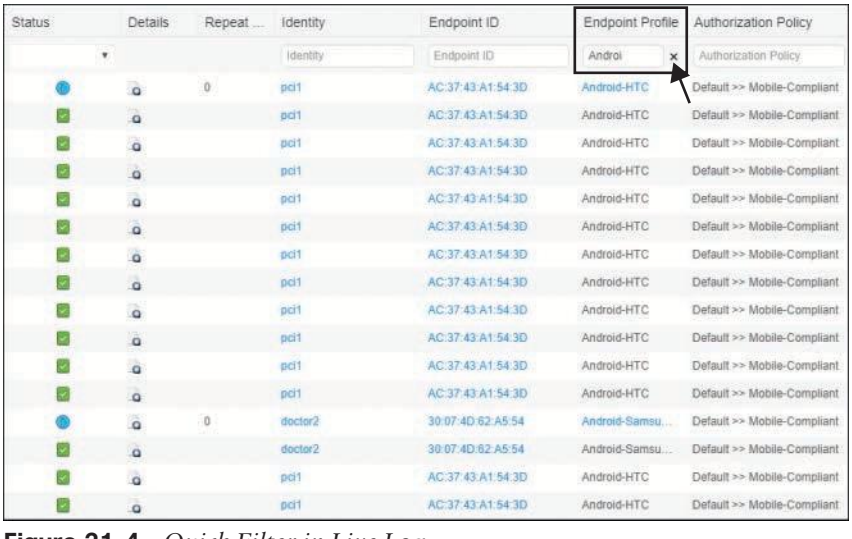


Figure 21-4 *Quick Filter in Live Log*

18. **Refresh**: Live Log is designed to be able to show authentication events in near real time. However, the number of entries can be daunting when many events are coming in each second (perhaps hundreds of events per second). ISE provides multiple refresh rate options, as shown in Figure 21-5. Experienced ISE administrators often use the Never option and manually click the Refresh button, shown in the figure, when they want the list to refresh.



Figure 21-5 *Manual Refresh*

19. **Record Selector**: Live Log provides a selector to limit the number of records displayed to keep it easier on the ISE administrator, as shown in Figure 21-6.

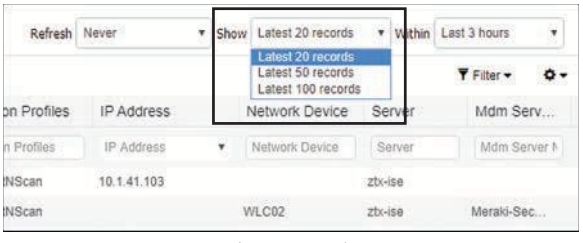


Figure 21-6 *Record Show Selector*

20. **Time Selector**: Live Log is designed for active troubleshooting, not long-term reporting. Therefore, it displays records from only the most recent 24 hours. Reporting is used for any logs that are older. In order to keep the record set more manageable, you can select different time ranges, as shown in Figure 21-7.

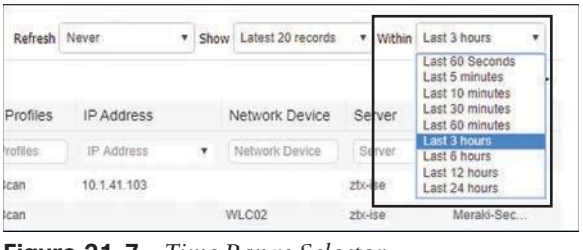


Figure 21-7 *Time Range Selector*

21. **Actions**: When you highlight a row, small target icons appear near key fields. You can click on such an icon to bring up an Actions menu related to those fields (see Figure 21-8).

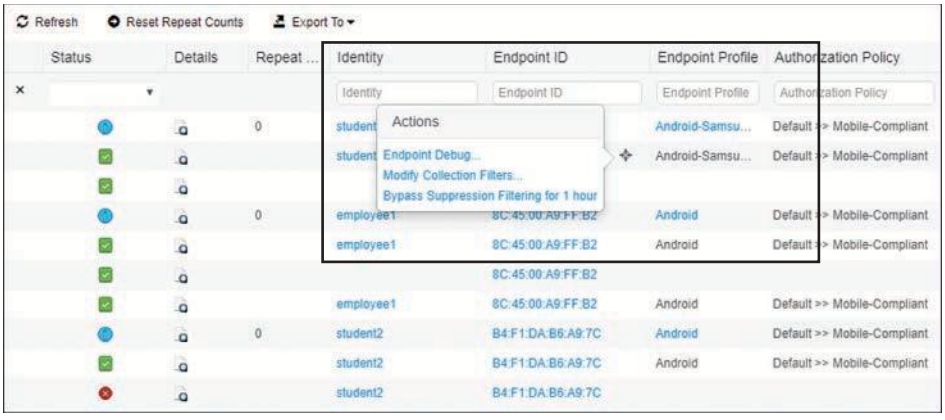


Figure 21-8 *Actions Shortcut Menu*

Advanced Filtering

The quick filter is displayed by default, but you can use an advanced filter by clicking Filter > Advanced Filter, as shown in Figure 21-9. The advanced filter allows you to create complex matching rules, as shown in Figure 21-10.



Figure 21-9 *Filter > Advanced Filter*



Figure 21-10 *Advanced Filter*

Authentication Details Report

The authentication details report is one of the single most useful tools to ever grace the network access control world. You might feel that like that is an opinion, but if that is how you feel, you probably have never had to manage a large NAC or 802.1X deployment.

Jesse Dubois is a genius-level ISE expert who leads the Cisco TAC escalations team for ISE. The very first thing he asks for any time there is an issue is the authentication details report because it provides answers to many problems.

Figures 21-11 through 21-14 show different sections of a single authentication details report for a failed event in Live Log.

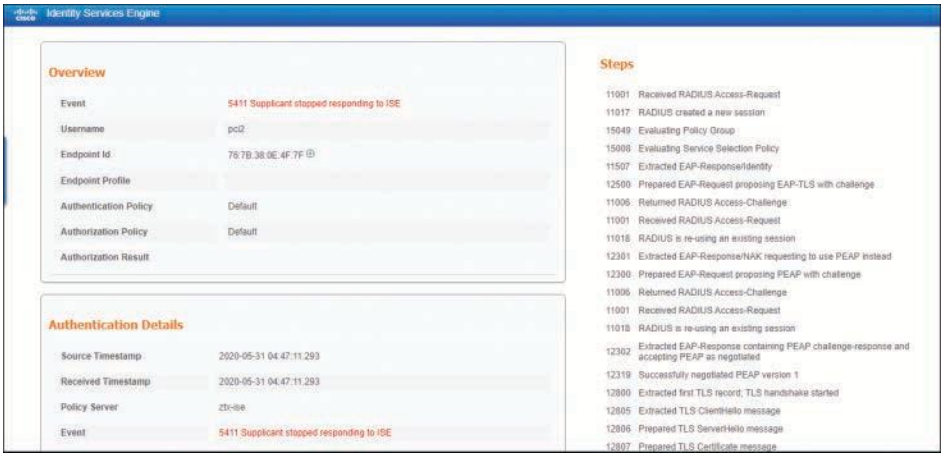


Figure 21-11 *Overview and Steps*

As you see in Figure 21-11, the first section of the report is the Overview section. It provides some of the most important fields for an ISE admin to see quickly, including the event itself, the username and endpoint identity, and the authentication and authorization results. The event displayed in this figure is event 5411, where a supplicant stopped responding to ISE.

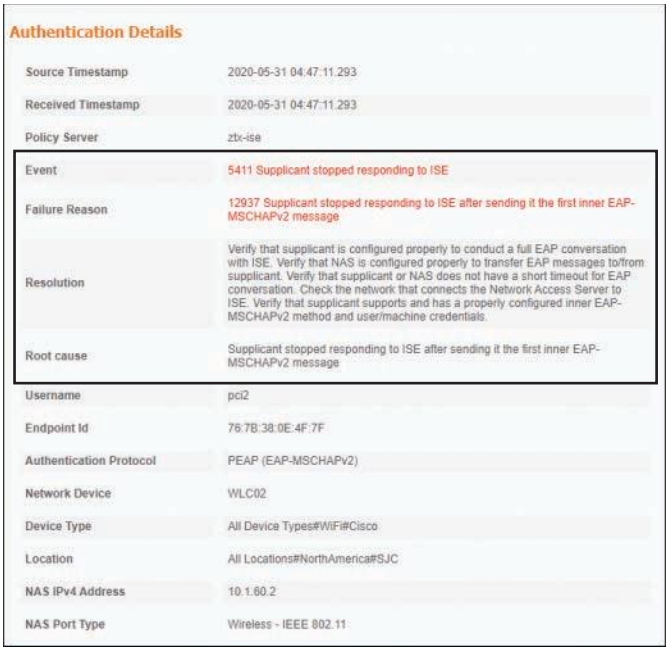
On the right in Figure 21-11, you can see the beginning of the Steps section. Experienced troubleshooters often focus on this long list the most. It shows every single step that ISE goes through during the authentication of a user or an endpoint.

Figure 21-12 shows the second major section of the report: the Authentication Details section.

Figure 21-12 shows the 5411 event again. It also shows a failure reason, a suggested resolution, and a root cause. Here you can see that the 5411 event occurred because an Android device temporarily had a poor wireless signal. When the device got a stronger signal, it authenticated again.

Figure 21-13 shows the final section of the report: Other Attributes. This section contains a handful of attributes from the authentication and the internal processes within ISE that can sometimes shed light on a situation.

Figure 12-14 shows the end of the Steps section for this report. The final steps of the event note the failure and show that a failed event is sent to Live Log.



Authentication Details

Source Timestamp	2020-05-31 04:47:11.293
Received Timestamp	2020-05-31 04:47:11.293
Policy Server	ztx-ise
Event	5411 Supplicant stopped responding to ISE
Failure Reason	12937 Supplicant stopped responding to ISE after sending it the first inner EAP-MSCHAPv2 message
Resolution	Verify that supplicant is configured properly to conduct a full EAP conversation with ISE. Verify that NAS is configured properly to transfer EAP messages to/from supplicant. Verify that supplicant or NAS does not have a short timeout for EAP conversation. Check the network that connects the Network Access Server to ISE. Verify that supplicant supports and has a properly configured inner EAP-MSCHAPv2 method and user/machine credentials.
Root cause	Supplicant stopped responding to ISE after sending it the first inner EAP-MSCHAPv2 message
Username	pci2
Endpoint Id	76:7B:38:0E:4F:7F
Authentication Protocol	PEAP (EAP-MSCHAPv2)
Network Device	WLC02
Device Type	All Device Types#WiFi#Cisco
Location	All Locations#NorthAmerica#SJC
NAS IPv4 Address	10.1.60.2
NAS Port Type	Wireless - IEEE 802.11

Figure 21-12 *Authentication Details Section*



Other Attributes

ConfigVersionId	81
RadiusPacketType	Drop
AcsSessionID	ztx-ise/378945942/56070
NAS-Port	2
CPMSessionID	0a013c02000008bb5ed1d098
EndPointMACAddress	76-7B-38-0E-4F-7F
ISEPolicySetName	Default
StepLatency	82=120001
TLSCipher	ECDHE-RSA-AES256-GCM-SHA384
TLSVersion	TLSv1.2
DTLSSupport	Unknown
Network Device Profile	Cisco
Location	Location#All Locations#NorthAmerica#SJC
Device Type	Device Type#All Device Types#WiFi#Cisco
IPSEC	IPSEC#Is IPSEC Device#No
Stage	Stage#Stage#Closed Mode
Device IP Address	10.1.60.2

Figure 21-13 *Other Attributes*

12319	Successfully negotiated PEAP version 1
12812	Extracted TLS ClientKeyExchange message
12813	Extracted TLS CertificateVerify message
12804	Extracted TLS Finished message
12801	Prepared TLS ChangeCipherSpec message
12802	Prepared TLS Finished message
12816	TLS handshake succeeded
12310	PEAP full handshake finished successfully
12305	Prepared EAP-Request with another PEAP challenge
11006	Returned RADIUS Access-Challenge
11001	Received RADIUS Access-Request
11018	RADIUS is re-using an existing session
12304	Extracted EAP-Response containing PEAP challenge-response
12313	PEAP inner method started
11521	Prepared EAP-Request/Identity for inner EAP method
12305	Prepared EAP-Request with another PEAP challenge
11006	Returned RADIUS Access-Challenge
11001	Received RADIUS Access-Request
11018	RADIUS is re-using an existing session
12304	Extracted EAP-Response containing PEAP challenge-response
11522	Extracted EAP-Response/Identity for inner EAP method
11806	Prepared EAP-Request for inner method proposing EAP-MSCHAP with challenge
12305	Prepared EAP-Request with another PEAP challenge
11006	Returned RADIUS Access-Challenge
12937	Supplicant stopped responding to ISE after sending it the first inner EAP-MSCHAPv2 message (Step latency=120001 ms)
61025	Open secure connection with TLS peer
5411	Supplicant stopped responding to ISE

Figure 21-14 *Final Steps*

The Blank Lines

Live Log always includes some blank lines in the identity column, as shown in Figure 21-15.

Status	Details	Repeat ...	Identity	Endpoint ID	Endpoint Profile	Authorization Policy	Authorization Profiles
			Identity	76:7B:38:0E:4F:7F	Endpoint Profile	Authorization Policy	Authorization Profiles
		0	pci2	76:7B:38:0E:4F:7F	Android	Default >> Give Access Now	Guests,PermitNScan
			pci2	76:7B:38:0E:4F:7F	Android	Default >> Give Access Now	Guests,PermitNScan
			pci2	76:7B:38:0E:4F:7F	Android	Default >> Give Access Now	Guests,PermitNScan
				76:7B:38:0E:4F:7F			
			pci2	76:7B:38:0E:4F:7F	Android	Default >> Give Access Now	Guests,PermitNScan
				76:7B:38:0E:4F:7F			
			pci2	76:7B:38:0E:4F:7F	Android	Default >> Give Access Now	Guests,PermitNScan
				76:7B:38:0E:4F:7F			

Figure 21-15 *CoA Events in Live Log*

These blanks indicate Change of Authorization (CoA) events. You can view such an event by clicking the details icon and viewing the authentication details report, shown in Figures 21-16 and 21-17.

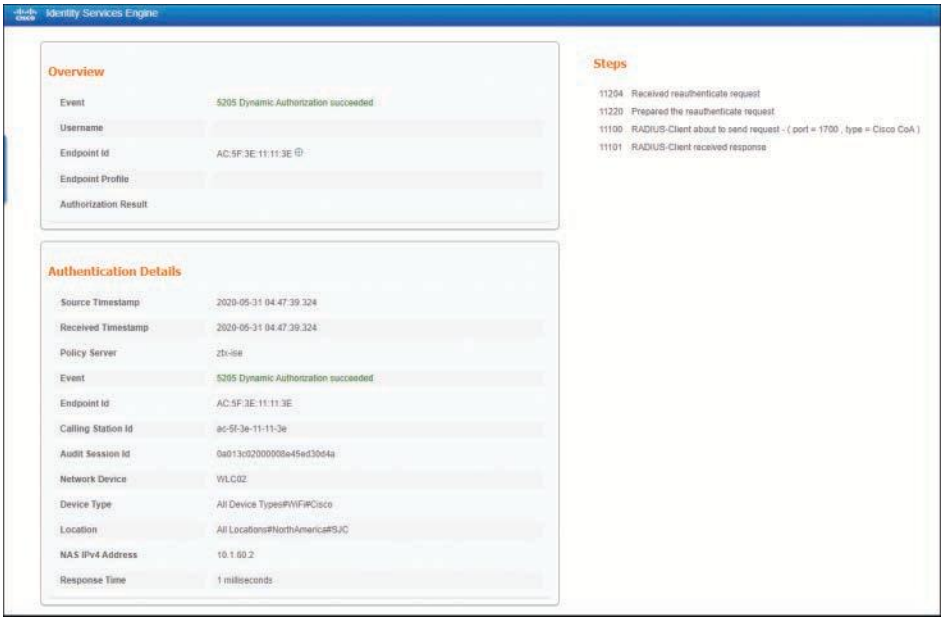


Figure 21-16 *Authentication Details Report for a CoA Event, Part 1*



Figure 21-17 *Authentication Details Report for a CoA Event, Part 2*

Live Sessions

Live Sessions differs from Live Log in that it shows activity related to the entire session, whereas Live Log is focused on the raw entries related to a passed or failed authentication. For example, recall from Chapter 16, "Bring Your Own Device," that a device may authenticate numerous times within a single session. This is true of nearly every advanced flow with ISE, including BYOD flows, guest flows, posture, and TC-NAC flows.

Key Topic

Live Sessions displays the current state of a session and allows you to take action on that session. Figure 21-18 shows a Live Sessions log with eight sessions being tracked.

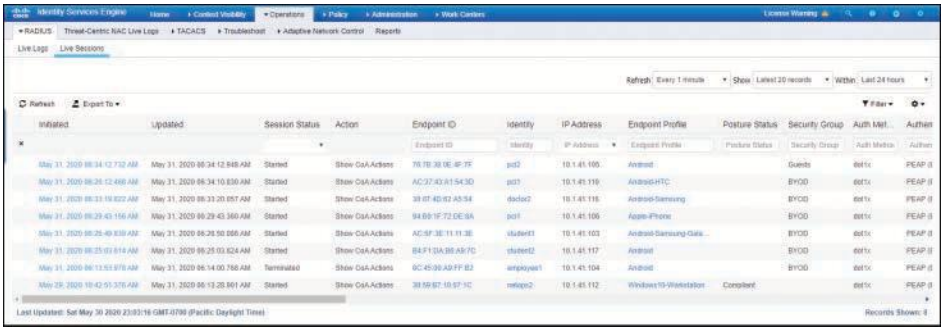


Figure 21-18 *Live Sessions*

In Figure 21-18 you can see two timestamp columns: Initiated and Updated. Initiated refers to the initial authentication, and Updated refers to the latest event that modified the session state. The status of the session is displayed as either Started or Terminated.

Key Topic

Live Sessions has a very useful function that allows you to send a CoA from the session listing by clicking on the target icon to bring up the Actions popup, as shown in Figure 21-19. This popup shows the available CoA types for the NAD being used in the session.

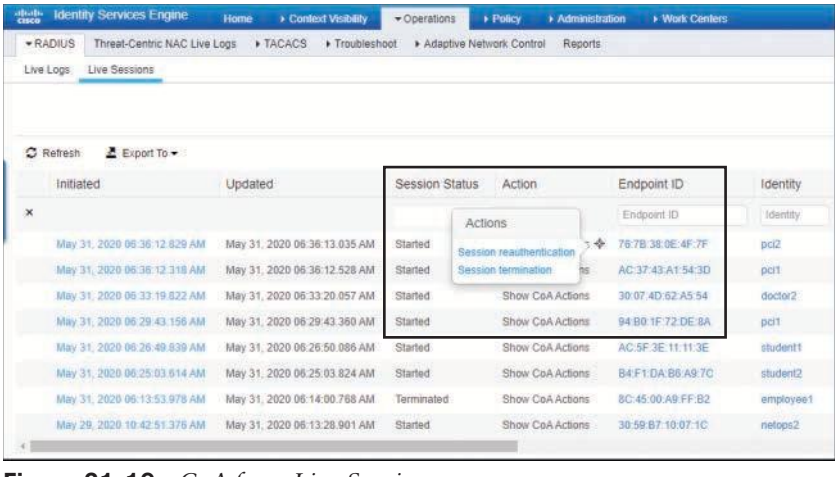


Figure 21-19 *CoA from Live Sessions*

Logging and Remote Logging

All nodes in an ISE deployment send information to the Monitoring and Troubleshooting (MnT) node(s) using syslog. The information that is sent includes system health, authentication, and authorization information, as well as accounting records, alarms, and any other system message. Traditional syslog transport uses UDP port 514, and ISE utilizes UDP port 20514.

Logging Targets

When a node joins an ISE cube, the policy is set for the nodes to send all logs to the defined LogCollector, which would be the primary monitoring persona.

To examine the configured LogCollector, navigate to Administration > System > Logging > Remote Logging Targets. As shown in Figure 21-20, this screen lists each syslog collector and its IP address, whether it is a TCP or UDP syslog, and the port number for the log collector. TCP syslog is used in certain environments to provide guaranteed log delivery.

Figure 21-20 shows the default log collector configuration for a standalone deployment, and Figure 21-21 shows a configuration with a distributed deployment and external syslog recipients.

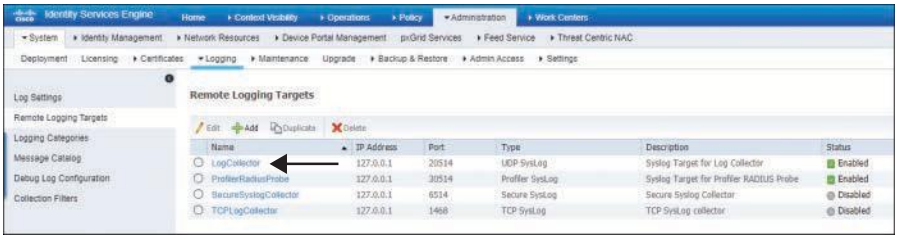


Figure 21-20 *Standalone Default Log Collector Configuration*

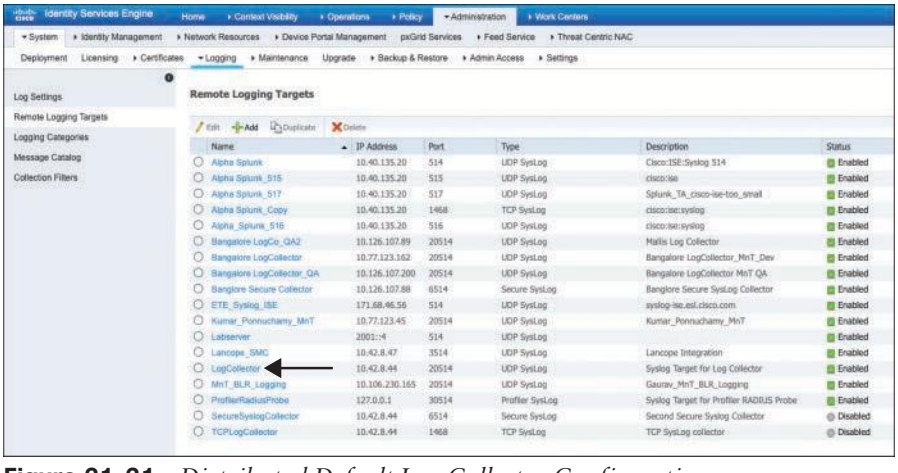


Figure 21-21 *Distributed Default Log Collector Configuration*

You can add a security information and event management (SIEM) or other syslog receiver to ISE simply by creating a new logging target. Figure 21-22 shows the addition of an external logging target.

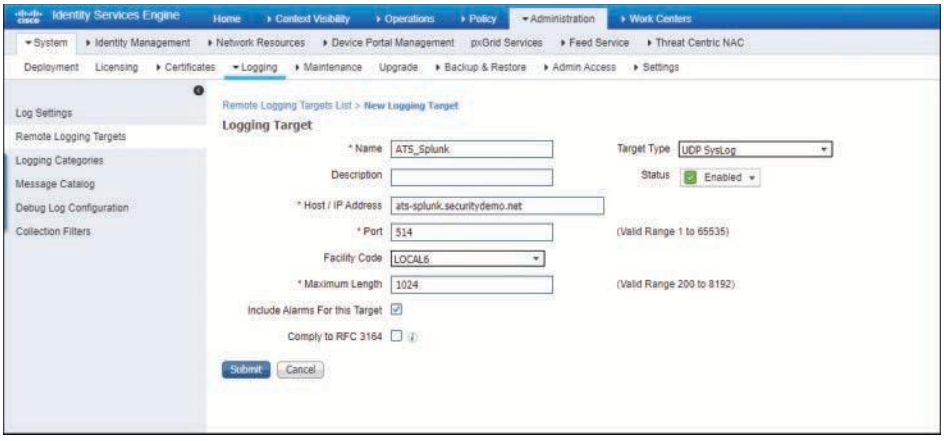


Figure 21-22 *Adding a Remote Logging Target*

In order to send logs to the newly added logging target, you need to add the target to the different logging categories.

Logging Categories

There are many different types of logs in ISE. You can see and configure the logging categories and the logging targets assigned to each of the categories by navigating to Administration > System > Logging > Logging Categories (see Figure 21-23).

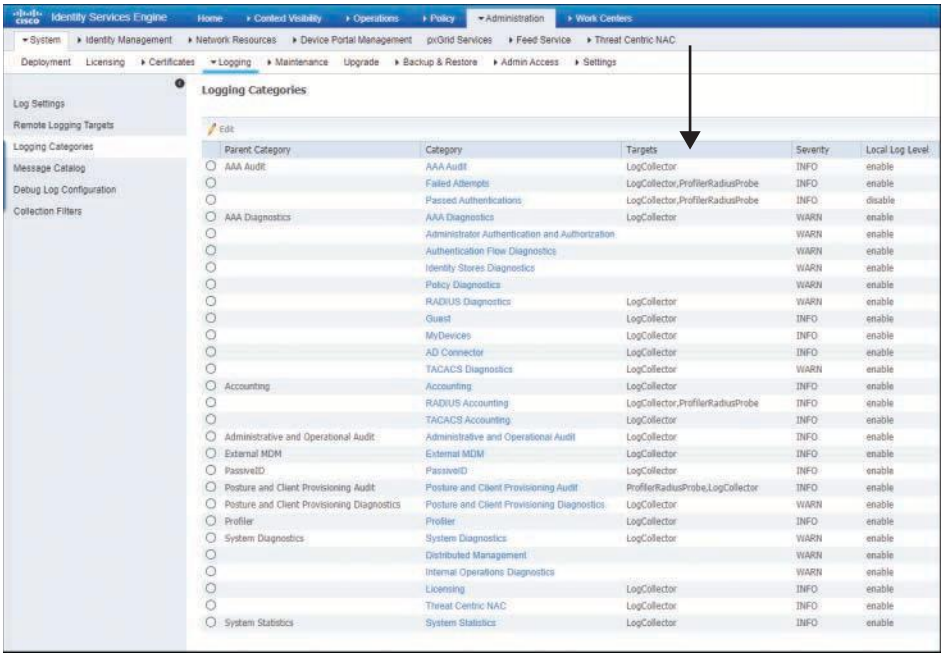


Figure 21-23 *Logging Categories*

For every category that you want an external SIEM or other collector to receive, you need to add that log collector to the individual category. The good news is that it takes effect for the entire ISE cube. To add a collector to a category, simply edit the category or click the category link. Then you need to ensure that the desired collectors are in the Selected column, as shown in Figure 21-24.

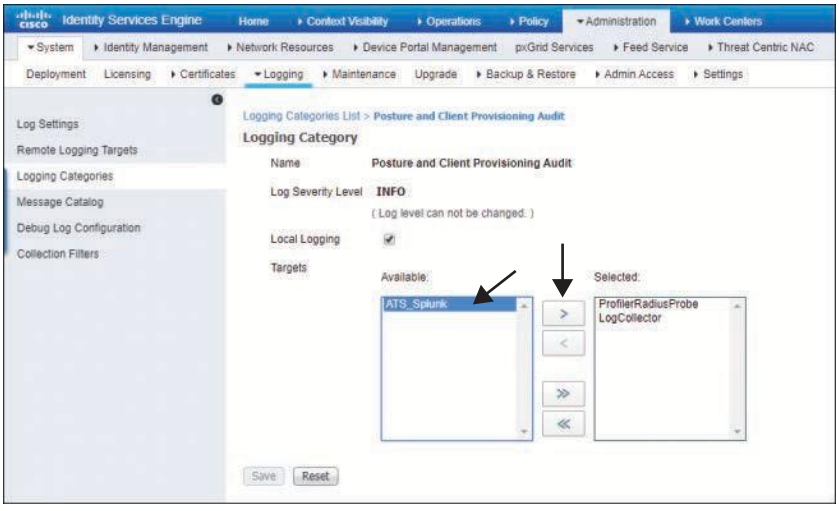


Figure 21-24 *Adding a Target to a Logging Category*

Debug Logs

With enterprise-class solutions, there is often a provision to provide very detailed levels of logging for troubleshooting and debugging. ISE provides the ability to set each component to log at different levels, including the debug level.

NOTE It is recommended that you not change any of the settings described in this section unless Cisco TAC directs you to do so. Setting components to debug level has an impact on system performance.

To change the logging level of a component, navigate to Administration > System > Logging > Debug Log Configuration. Click on the name of the ISE node that needs to have the log settings modified, and a list of components and the current log-level settings is displayed, as shown in Figure 21-25.

To modify a component's logging level, edit the component or double-click on the logging level. Select the new logging level from the dropdown and click Save, as shown in Figure 21-26.

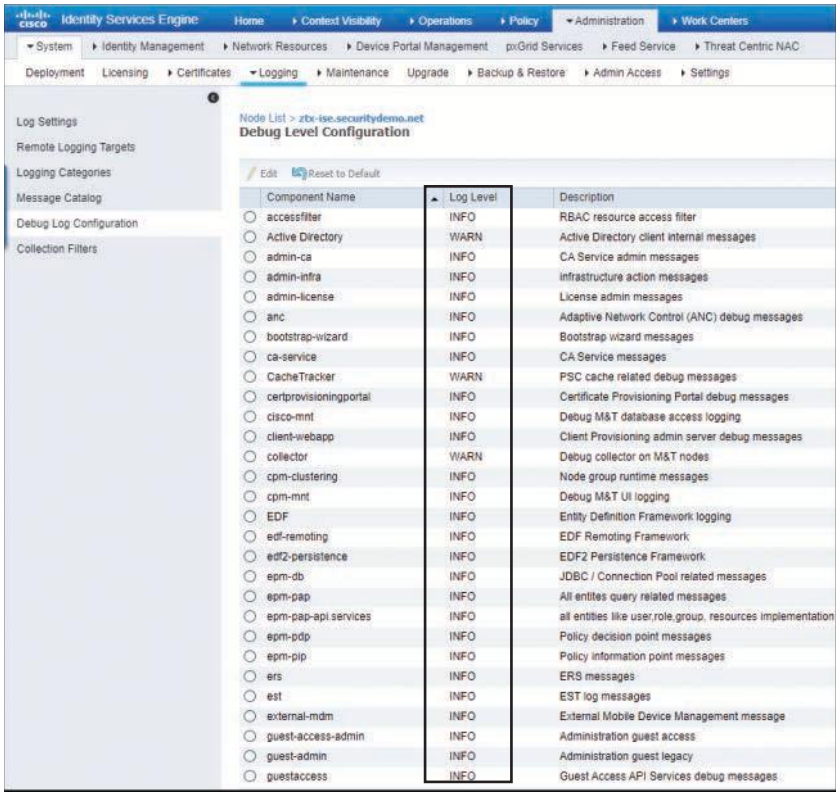


Figure 21-25 *Default Debug Log Configuration*



Figure 21-26 *Changing the Logging Level of a Category*

The resulting debug log file may be downloaded from the GUI, viewed from the CLI, or be included in a support bundle.

Downloading Debug Logs from the GUI

To debug log files from the GUI, navigate to Operations > Troubleshoot > Download Logs > Debug Logs. Then select the ISE node that contains the log files that need to be downloaded and click on the link for the specific log file, as shown in Figure 21-27.

Once the file is downloaded, you can view it in any standard text editor application.

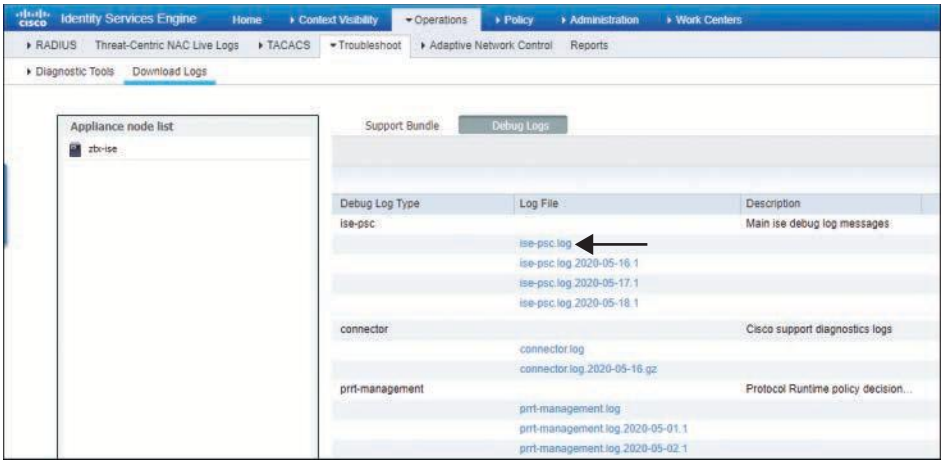


Figure 21-27 *Downloading a Debug Log File from the GUI*

Viewing Log Files from the CLI

Sometimes you might want to view logs on the local system through the command-line interface (CLI). ISE provides the ability to view and tail (view the last x number of lines of) the log files directly from the CLI by using the **show logging application** and **show logging system** commands. The **show logging application** command is used for the ISE application's logs, and the **show logging system** command is used for logs related to the underlying operating system.

To view a list of all the ISE log files, you use the **show logging application** command, as shown in Example 21-1.

Example 21-1 *show logging application Command Output*

```
ztx-ise/admin# show logging application
      53074 May 21 2020 16:02:03  ad_agent.log
          0 May 31 2020 00:00:02  appserver/catalina.log
      70582 May 31 2020 07:37:44  appserver/catalina.out
```

To view the contents of a specific log file, enter the **show logging application** command and the name of the log file, such as **ad_agent.log**, as shown in Example 21-2.

Example 21-2 *show logging application ad_agent.log Command Output*

```
ztx-ise/admin# show logging application ad_agent.log
2019-05-20 13:53:37,008 ALWAYS ,139899300931456,Logging started,LwSmSetLogger(),lwsm/
server/logger.c:489
2019-05-20 13:53:37,008 ALWAYS ,139899300931456,Locale: en_US.UTF-8,main(),lwsm/
server/main.c:303
2019-05-20 13:53:37,008 VERBOSE,139899300931456,LwSmCheckSaneShutdownMarker()
called. Checking sane shutdown marker /opt/pbis/var/.lwsmd_marker,LwSmCheckSaneShutd
ownMarker(),lwsm/server/main.c:1548
2019-05-20 13:53:37,008 VERBOSE,139899300931456,LwSmCheckSaneShutdownMarker() -
Marker file /opt/pbis/var/.lwsmd_marker doesn't exist. Process was shutdown
correctly. Creating now...,LwSmCheckSaneShutdownMarker(),lwsm/server/main.c:1558
```

```
2019-05-20 13:53:37,013 DEBUG  ,139898359117568,LwRtlSetThreadPriority: running
command: sudo /usr/bin/renice -n -10 -p 27364,LwRtlSetThreadPriorityEx(),lwbase/src/
threadpool-common.c:1202
2019-05-20 13:53:37,156 DEBUG  ,139898990667520,LwRtlSetThreadPriority: command
res=0; exit code=0,LwRtlSetThreadPriorityEx(),lwbase/src/threadpool-common.c:1219
2019-05-20 13:53:37,158 DEBUG  ,139898895988480,LwRtlSetThreadPriority: command
res=0; exit code=0,LwRtlSetThreadPriorityEx(),lwbase/src/threadpool-common.c:1219
2019-05-20 13:53:37,159 DEBUG  ,139899032631040,LwRtlSetThreadPriority: command
res=0; exit code=0,LwRtlSetThreadPriorityEx(),lwbase/src/threadpool-common.c:1219
2019-05-20 13:53:37,159 DEBUG  ,139898451371776,LwRtlSetThreadPriority: pid=27342
thread LWP=27404, set priority to -10,LwRtlSetThreadPriorityEx(),lwbase/src/
threadpool-common.c:1186
2019-05-20 13:53:37,159 INFO   ,139898451371776,AD connector service starting
up,Startup()
```

In addition to being able to view a log file, you can also perform standard operations on a file, such as **tail** and **grep**. You can see a list of the operations you can perform in the CLI by adding a pipe character (**|**) after the log name (see Example 21-3).

Example 21-3 *show logging application ad_agent.log | ? Command Output*

```
ztx-ise/admin# show logging application ad_agent.log | ?
Output modifier commands:
  begin    Begin with line that matches
  count    Count the number of lines in the output
  end      End with line that matches
  exclude  Exclude lines that match
  include  Include lines that match
  last     Display last few lines of the output
```

Support Bundles

The ultimate tool for Cisco TAC is the support bundle. You can think of a support bundle as being equivalent to the **show tech-support** command on a Cisco IOS device. A support bundle can contain all the logs across the entire ISE system, including all its functions, the configuration, and anything else a support engineer needs to troubleshoot the ISE installation and maybe even to re-create the environment in a lab, if necessary.

A support bundle is saved as a simple tar.gpg (GPG encrypted) file. The support bundle is automatically named with the date and timestamp, using the following format:

ise-support-bundle_ise-support-bundle-mm-dd-yyyy-hh-mm.tar.gpg.

You have the option to choose which logs you want to be part of a support bundle. For example, you can configure logs from a particular service to be part of a bundle.

The logs that you can download are categorized as follows:

- **Full Configuration Database**: If you choose this category, the ISE configuration database is saved into the support bundle and allows Cisco TAC to import this database configuration in another ISE node to re-create the scenario.

- **Debug Logs:** This category captures bootstrap, application configuration, runtime, deployment, monitoring and reporting, and public key infrastructure (PKI) information.
- **Local Logs:** This category contains log messages from the various processes that run but are not collected by the MnT node.
- **Core Files:** This category contains critical information that can help identify the cause of a crash. These logs are created if the application crashed and includes core dumps.
- **Monitoring and Reporting Logs:** This category contains information about the alarms and reports from the MnT node.
- **System Logs:** This category contains the underlying OS (Application Deployment Engine [ADE-OS]) logs. These logs are not directly part of the ISE application.
- **Policy Configuration:** This category includes the authentication and authorization policy.

There are two ways to create and download a support bundle: from the ISE GUI and from the CLI, as described in the sections that follow.

Creating a Support Bundle from the GUI

The more common way to create a support bundle is through the GUI, which provides more options than the CLI.

To create a support bundle by using the GUI, perform the following steps in the ISE GUI:

Step 1. Navigate to **Operations > Troubleshoot > Download Logs**.

Step 2. Select the ISE node in the left pane.

Step 3. Choose the categories of logs you want to include.

Step 4. Select a date range to limit the file size. Keep in mind that there is no need to send every file and spend time creating a large support bundle if the problem has been replicated in the last day or two.

Step 5. Select either **Public Key Encryption** or **Shared Key Encryption** and create a password. The Public Key Encryption option uses Cisco's PKI to encrypt the bundle, and only Cisco can decrypt it; this helps with Cisco TAC case automation. It is highly recommended to use this approach instead of using a password.

Step 6. Click **Create Support Bundle**, as shown in Figure 21-28.

Creating a Support Bundle from the CLI

Using the GUI is the more common way to create a support bundle, and it provides many options. However, the GUI is not always available, especially in disastrous situations, and for such situations, you need to be familiar with using the **backup-logs** CLI command. You cannot pick and choose which logs to capture from the CLI; the **backup-logs** command captures all logs, as shown in Example 21-4.

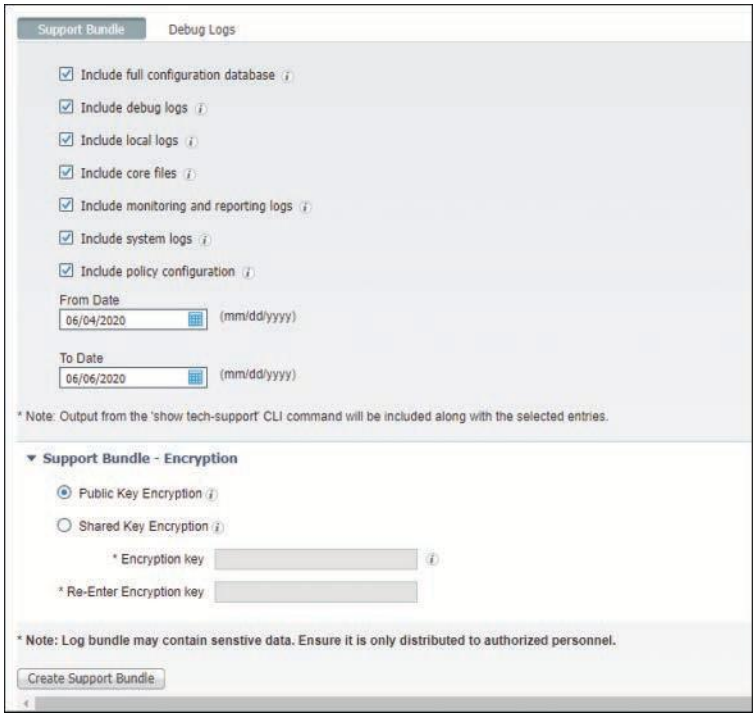


Figure 21-28 *Creating a Support Bundle from the GUI*

Key Topic

The support bundle must be uploaded to a repository, and you are given the option to use a password or to use the Cisco PKI to encrypt the bundle. Just as with the GUI, it is recommended to use the Cisco PKI instead of a custom password because the Cisco PKI helps with the Cisco TAC automation tools. If a Cisco ISE administrator creates a password for a bundle and forgets to tell Cisco TAC what that password is, in the event of a problem, there will be delays in the case resolution.

Example 21-4 shows an example of using the CLI to create a support bundle named SupportBundle. In this example, Cisco PKI is used to encrypt the file, and the bundle is saved to a repository named FTP'.

Example 21-4 *Creating a Support Bundle from the CLI*

```
ztx-ise/admin# backup-logs SupportBundle repository FTP public-key
% Creating log backup with timestamped filename: SupportBundle-pk-200607-0214.tar.
gpg
% supportbundle in progress: Copying database config files...10% completed
% supportbundle in progress: Copying debug logs...20% completed
% supportbundle in progress: Copying local logs...30% completed
ls: cannot access /localdisk/corefileanalysis: No such file or directory
% supportbundle in progress: Copying monitor logs...40% completed
% supportbundle in progress: Copying policy xml...50% completed
% supportbundle in progress: Copying system logs...60% completed
```

Diagnostic Tools

In addition to all its terrific logging capabilities, ISE has a number of built-in diagnostic tools that are found under Operations > Troubleshoot > Diagnostic Tools. The following sections provide a brief overview of the most commonly used tools.

RADIUS Authentication Troubleshooting Tool

The RADIUS Authentication Troubleshooting tool examines different aspects of a session and provides details that may not be available in the detailed authentication report. This tool also provides some suggestions for items to check next. To use this tool, follow these steps:

Step 1. Navigate to **Operations > Troubleshoot > Diagnostic Tools > General Tools > RADIUS Authentication Troubleshooting.**

Step 2. Limit your search by selecting specifics such as a username, failed or passed authentication status, and more. Click **Search.** Figure 21-29 shows a user named radius-test selected, along with any authentication status for the current day.

Step 3. Click **Troubleshoot** and select an authentication from the list that appears, as shown in Figure 21-29. A troubleshooting progress screen appears, as shown in Figure 21-30.

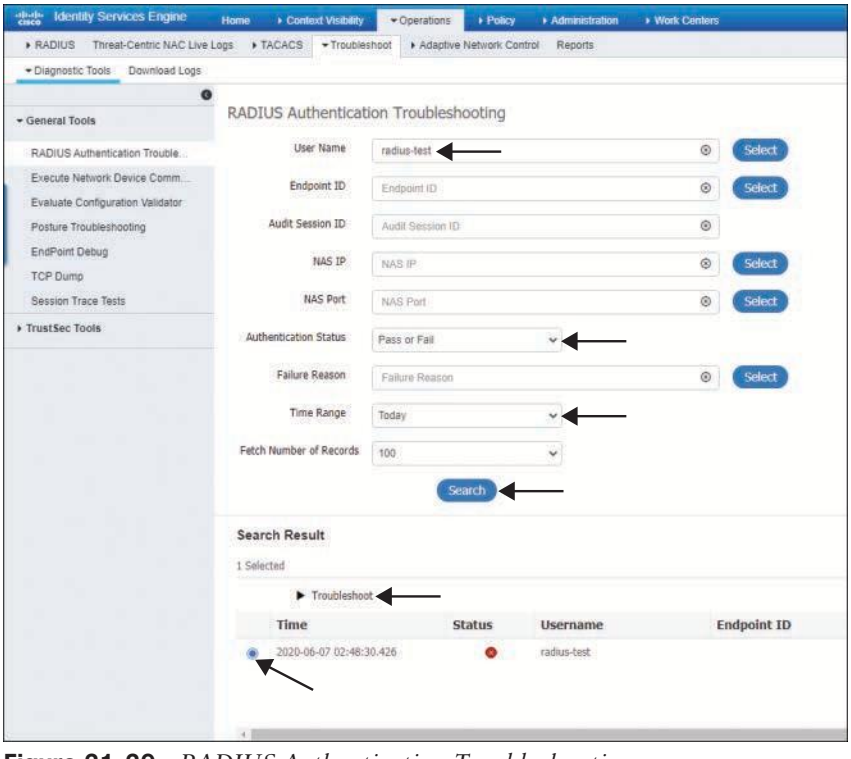


Figure 21-29 *RADIUS Authentication Troubleshooting*

Step 4. When the troubleshooting is complete, click **Show Result Summary**, shown in Figure 21-31, to see the diagnosis and suggested resolution.

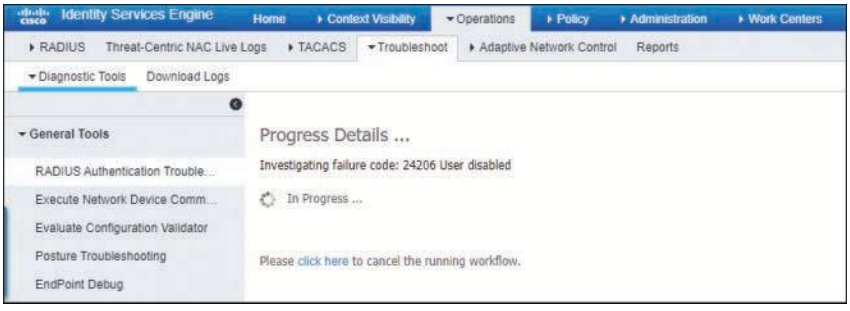


Figure 21-30 *RADIUS Authentication Troubleshooting in Progress*

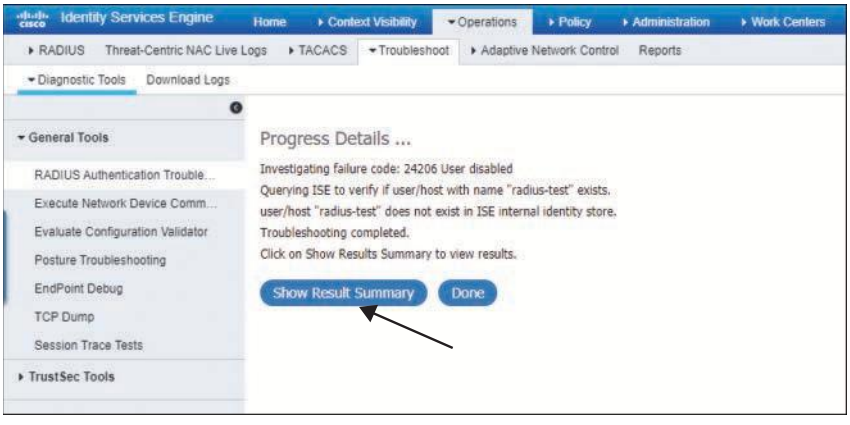


Figure 21-31 *RADIUS Authentication Troubleshooting Result Summary*

The Result Summary page in Figure 21-32 shows that there is no user or host named **radius-test.**

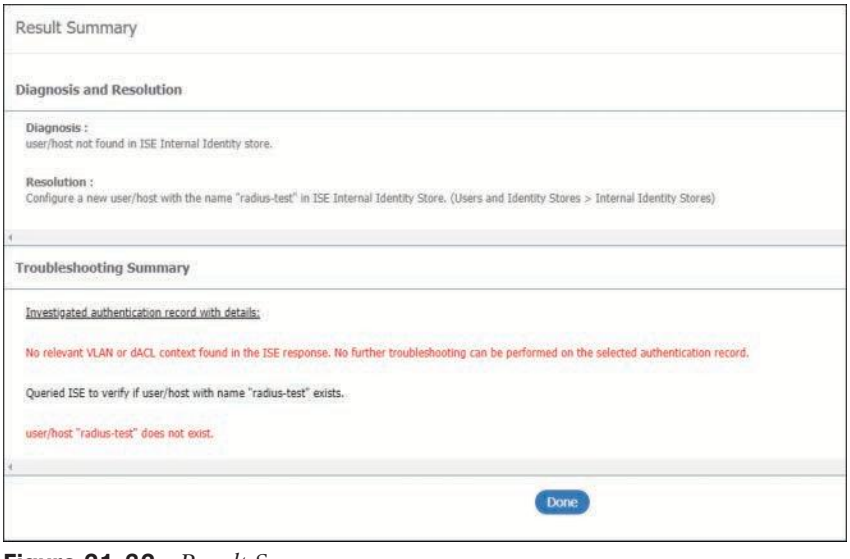


Figure 21-32 *Result Summary*

Execute Network Device Command

The Execute Network Device Command tool allows an ISE administrator to execute a command on a network device remotely. It seems like ISE should allow you to select one of the NADs that have already been configured within ISE, but it does not. As shown in Figure 21-33, you can enter an IP address for a network device and the command to run on that device.



Figure 21-33 *Execute Network Device Command*

After filling in the IP address and command, click **Run** to begin the remote login process. As shown in Figure 21-34, you are prompted to click a link to enter credentials.

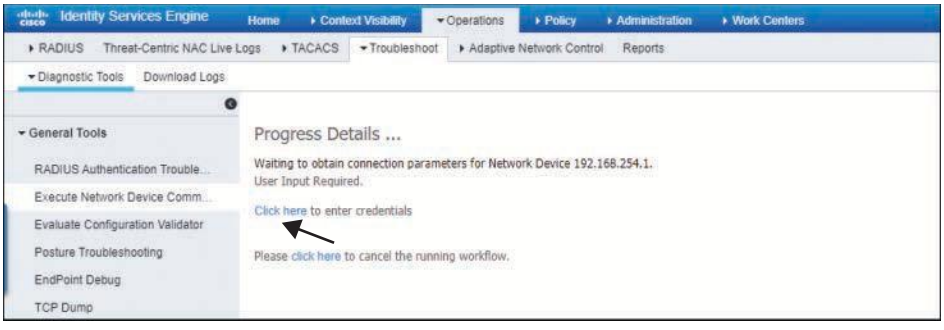


Figure 21-34 *Click to Enter Credentials*

Clicking the link launches a popup window where you enter the credentials for SSHv2 or Telnet. If a console server is in use, you can enter that IP address and port as well. Figure 21-35 shows the credentials window.

Click Submit, and ISE logs you in remotely and executes the command, as shown in Figure 21-36. The GUI displays the resulting output, as shown in Figure 21-37.

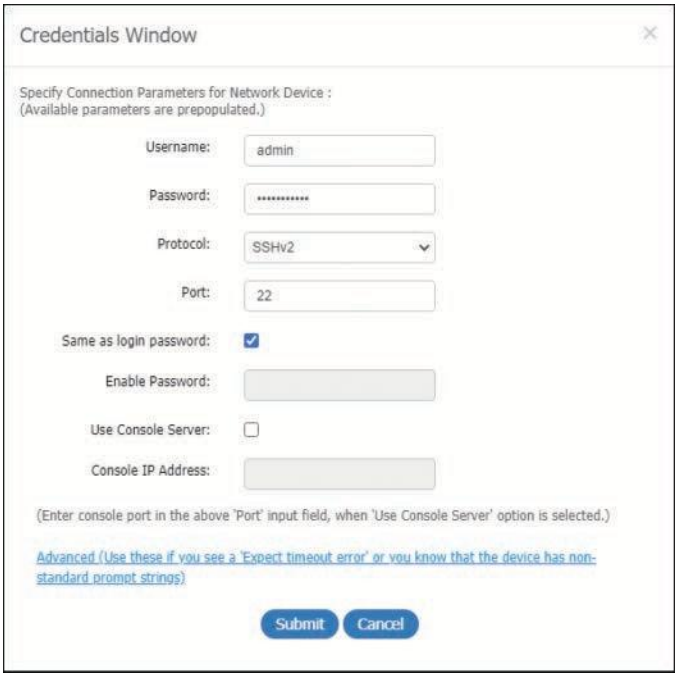


Figure 21-35 *Credentials Window*

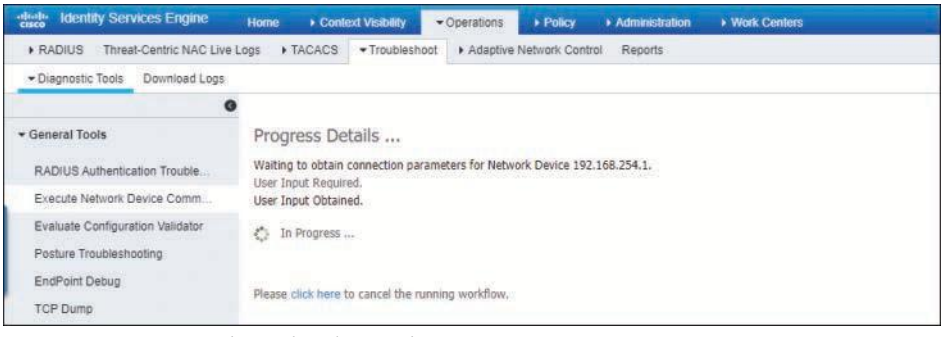


Figure 21-36 *Credentials Obtained*

Evaluate Configuration Validator

Evaluate Configuration Validator is a terrific tool with a less-than-terrific name. This tool can connect to a switch via Telnet, SSH, or even a console terminal server and examine a configuration. The tool compares the configuration to a template configuration that is built in to ISE and indicates any differences between the two configurations.

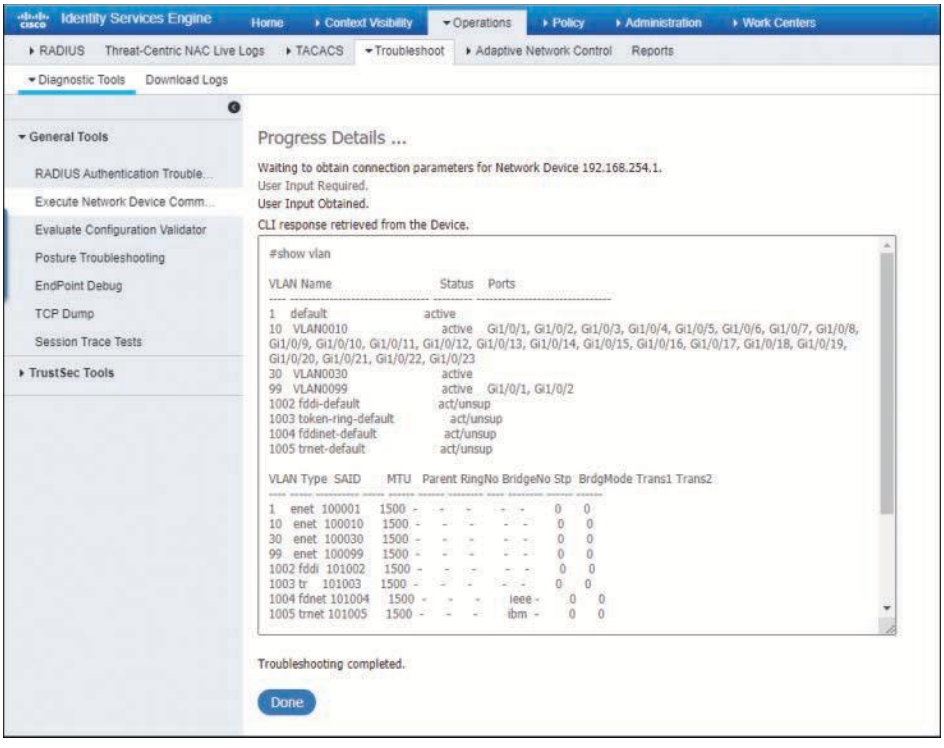


Figure 21-37 *Command Output Displayed*

The Evaluate Configuration Validator tool has been due for a major update since ISE Version 1.2, but it can still provide a lot of value, especially if you are using older switches. In fact, some ISE administrators have been known to point this tool to an unconfigured switch running Cisco IOS Version 12.2 and copy and paste the resulting commands in order to configure the switch.

The following are a few common configurations that may be misdiagnosed as missing or incorrect:

- The Evaluate Configuration Validator does not currently understand the device sensor and expects the use of SNMP for the SNMP probe(s).
- The Evaluate Configuration Validator does not recognize the active test options when defining a RADIUS server, and it may therefore think the RADIUS server definition is incorrect.
- Web Authentication (WebAuth) is Local WebAuth (LWA) only, and the tool does not recognize that MAB is used for Central WebAuth (CWA) instead.
- The Evaluate Configuration Validator does not support the Cisco Common Classification Policy Language (C3PL) commands available in IOS 15.2 and above, and it therefore does not support Cisco's flagship Catalyst 9000 Series switches.

Despite its limitations, Cisco TAC still suggests using the Evaluate Configuration Validator tool to quickly identify misconfigurations. In many cases, customers can avoid opening Cisco TAC cases by using it.

To run the Evaluate Configuration Validator tool, follow these steps:

Step 1. Navigate to **Operations > Troubleshoot > Diagnostic Tools > General Tools > Evaluate Configuration Validator.**

Step 2. Enter the IP address of the NAD and the options you would like the tool to examine, as shown in Figure 21-38.

NOTE In Figure 21-38, Local Web Authentication is unchecked because CWA is in use.

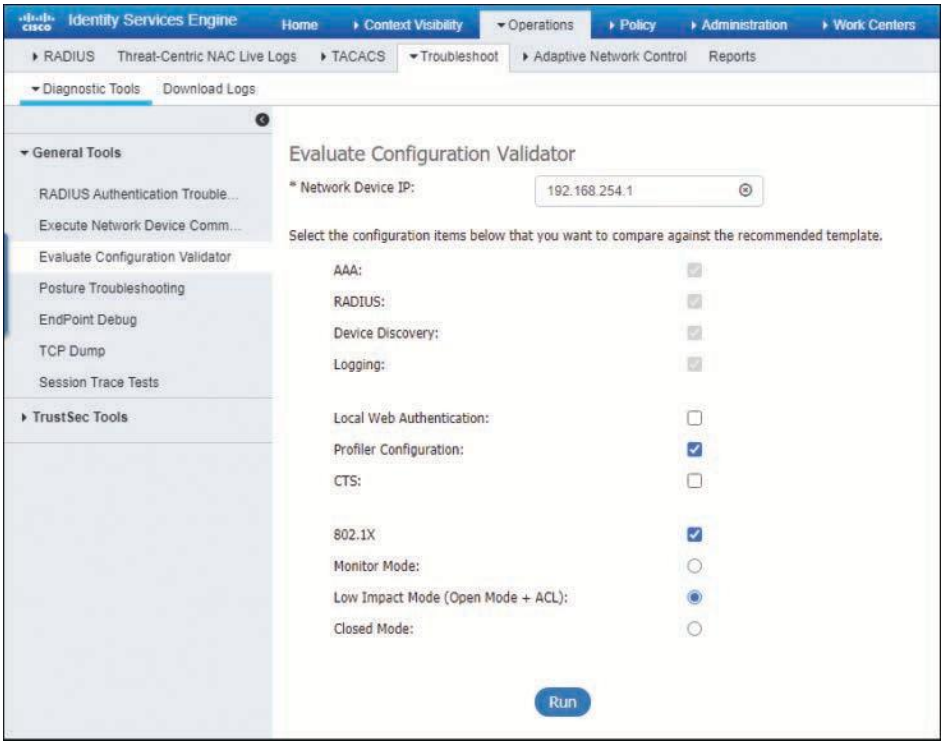


Figure 21-38 *Initializing the Evaluate Configuration Validator Tool*

Step 3. Click **Run** to begin the evaluation. ISE connects to the switch and prompts for your interaction if the switch asks for user authentication or other interactive prompts, as shown in Figure 21-39.

Step 4. When prompted to do so, select which interfaces you wish to compare, as shown in Figure 21-40. If you have been following the recommendations throughout this book, then nearly every interface will have exactly the same configuration, and you should select only one interface to compare.

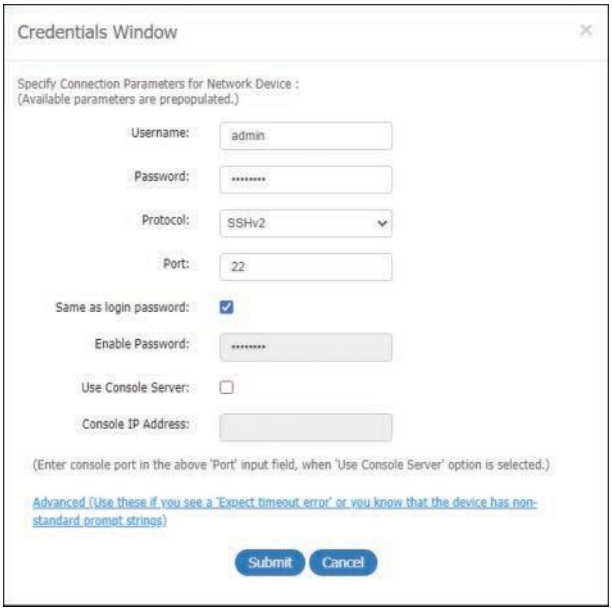


Figure 21-39 *Initializing the Evaluate Configuration Validator Tool*

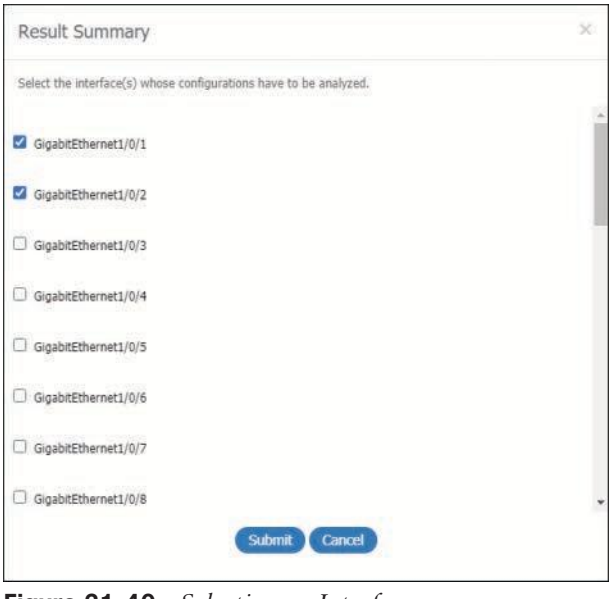


Figure 21-40 *Selecting an Interface*

Step 5. When you get a comparison similar to the one in Figure 21-41, click **Show Result Summary** to see the results. Figure 21-42 shows an example of a result summary.

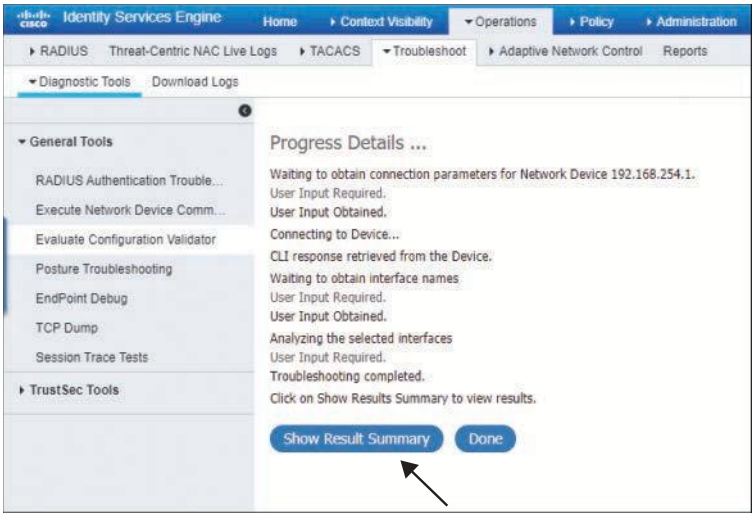


Figure 21-41 *Comparison Complete*

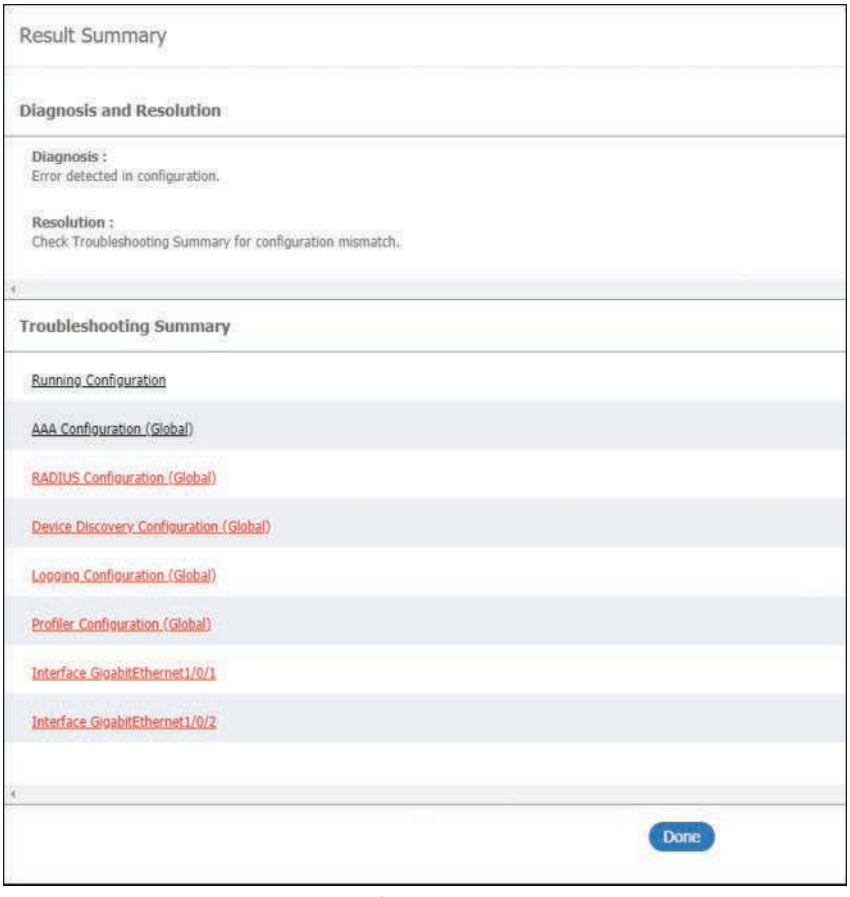


Figure 21-42 *Comparison Results*

As mentioned earlier in this chapter, the Evaluate Configuration Validator tool does not understand some of the newer commands that supersede the commands in the template used for the comparison. Figure 21-43 provides an example, in which the comparison result states that the **radius-server host** command is missing.

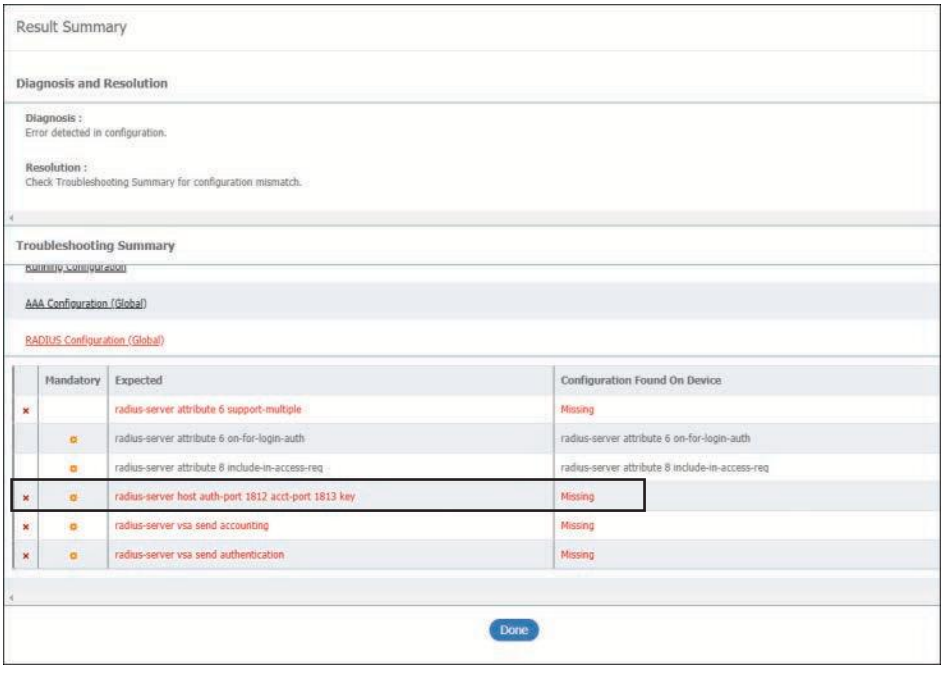


Figure 21-43 *RADIUS Configuration (Global)*

The command is truly missing, but the RADIUS server is actually configured (see Example 21-5). This situation exists because the switch being evaluated is running an IOS version newer than 12.2. Newer versions of IOS used a different approach to RADIUS server configuration in order to support both IPv4 and IPv6 addresses.

Example 21-5 *Ensuring That the RADIUS Server Is in the Configuration*

```
Cat9300# show run | begin radius server
radius server ZTX-ISE
 address ipv4 10.1.100.248 auth-port 1812 acct-port 1813
 key [redacted]
!
```

Therefore, it is up to you as an ISE administrator or consultant to use the Evaluate Configuration Validator tool carefully. This tool can provide a lot of value, but you should not take it at face value, and you should examine the switch configuration yourself to ensure that the necessary commands are in place.

Posture Troubleshooting

ISE includes the Posture Troubleshooting tool to help you identify why a user failed a posture check and was denied access to the network.

To run the Posture Troubleshooting tool, follow these steps:

Step 1. Navigate to **Operations > Troubleshoot > Diagnostic Tools > General Tools > Posture Troubleshooting.**

Step 2. Enter any of the search fields and click **Search**, as shown in Figure 21-44.

Step 3. Click **Troubleshoot** and select one of the failed posture assessments (see Figure 21-44).

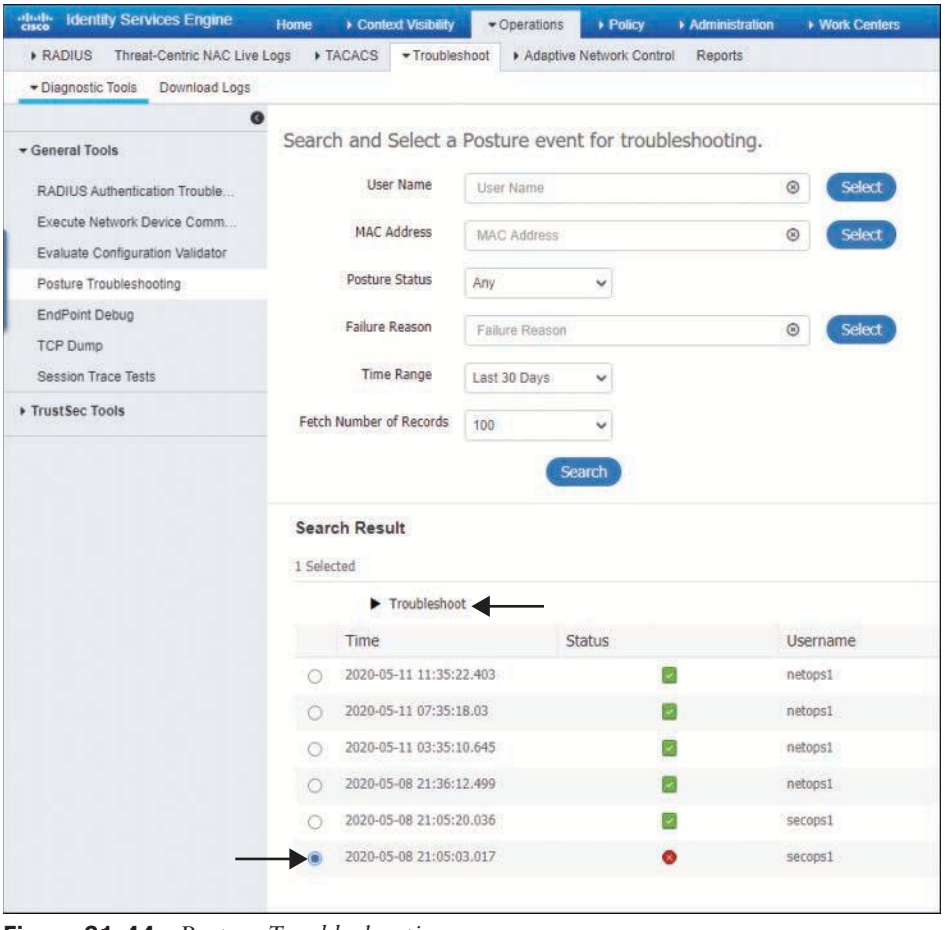


Figure 21-44 *Posture Troubleshooting*

Step 4. When the lookup is complete, click **Show Result Summary**, as shown Figure 21-45.

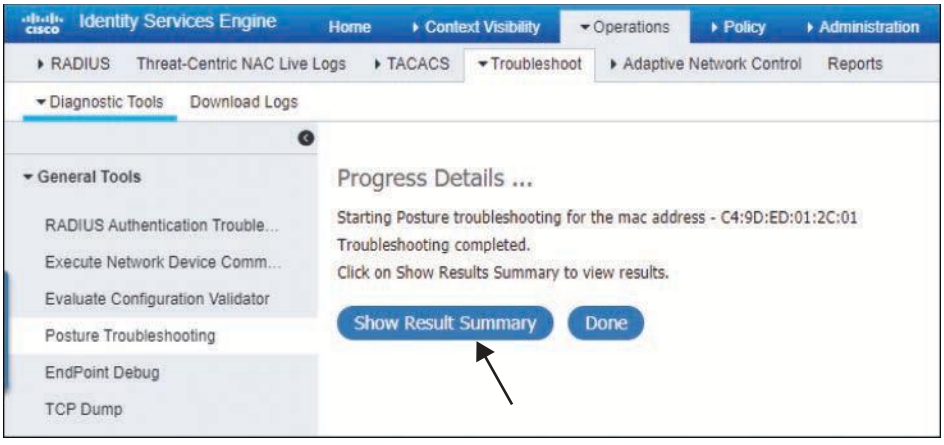


Figure 21-45 *Posture Troubleshooting Show Result Summary Button*

As shown in Figure 21-46, the tool shows you a report on which condition (or check) failed and why. In the example shown here, the failure occurred because the condition was not met; in this case, the file did not exist where the condition looked for it.

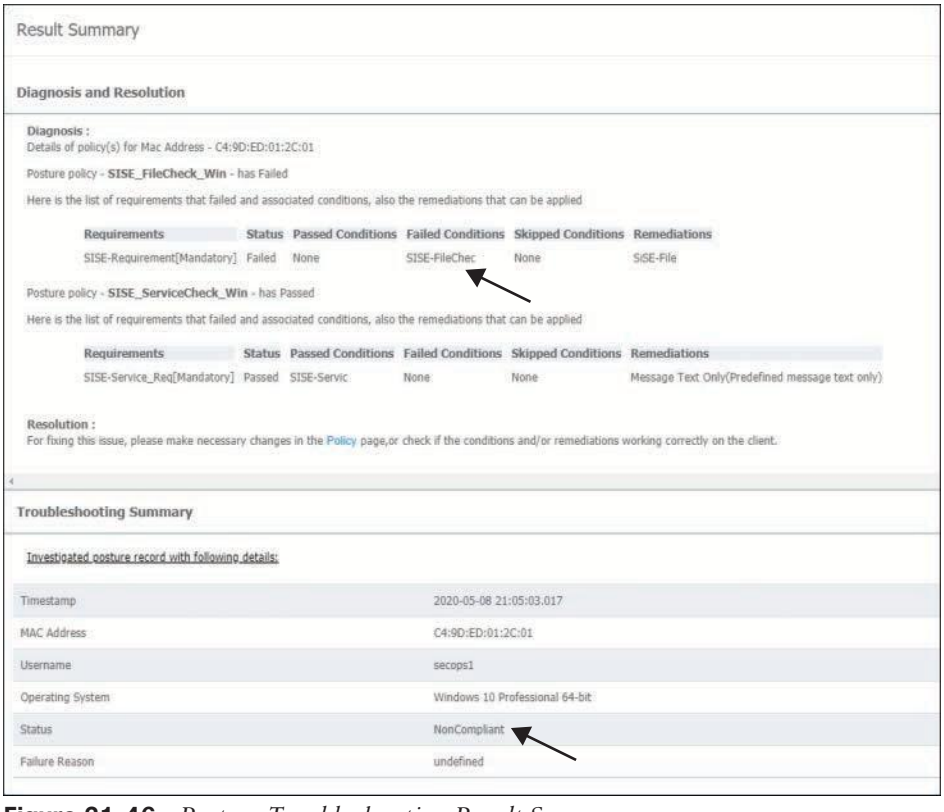


Figure 21-46 *Posture Troubleshooting Result Summary*

Endpoint Debug

An ISE deployment can be quite large and distributed in nature. Even with a single ISE node, many different components within ISE might be in use for any network session—for example, Guest for all web portal traffic as well as guest account creation and sponsorship, posture for checking the compliance of an endpoint, profiling to help identify the type of an endpoint, and RADIUS runtimes for the processing of incoming access requests.

As you are troubleshooting, there may come a time when you have to set the logging level of a particular component to debug and examine the debug logs within ISE. There are so many logs and so many logging categories in ISE that it is often quite burdensome to search through all those separate logs to find the entries that pertain to the endpoint, user, or session in question. This is where the Endpoint Debug tool comes into play. It was designed by Aaron Woland and a leader within Cisco's TAC named Jesse Dubois. The purpose of this tool is to allow you to easily and effectively troubleshoot an endpoint's activity with ISE in its entirety.



The Endpoint Debug tool provides a single debug file for all components for a specific endpoint across its entire session—across the entire deployment! So, if an endpoint is getting profiled in the east coast data center and the west coast data center at the same time, all that information shows up in the single, consolidated debug file. This tool prevents you from having to enable debugging on the components for all endpoints, and it focuses the debugging on a specific endpoint and its related session activity instead. This is incredibly elegant, and it helps advanced administrators and TAC engineers greatly reduce time to resolution when experiencing an issue.

There are two ways to launch the Endpoint Debug tool. The first (and more cumbersome) method is to follow these steps:

Step 1. **Navigate to Operations > Troubleshoot > Diagnostic Tools > General > EndPoint Debug** and select **MAC Address** or **IP**.

Step 2. Enter the endpoint's address.

Step 3. Click **Start**.

Step 4. Click **Stop** when it's time to examine the file.

Step 5. Click the link for the resulting file (which is named after the endpoint's MAC address) to download the resulting file.

Figure 21-47 shows the Endpoint Debug screen.

The second way to launch the Endpoint Debug tool is to click on the actions in Live Log and select Endpoint Debug from the list of available options, as shown in Figure 21-48.

When you launch the Endpoint Debug tool from Live Log, the tool automatically launches in a new window, prepopulated with the address for the endpoint.

Figure 21-49 shows debugging in progress, and Figure 21-50 shows the final result of the Endpoint Debug tool.

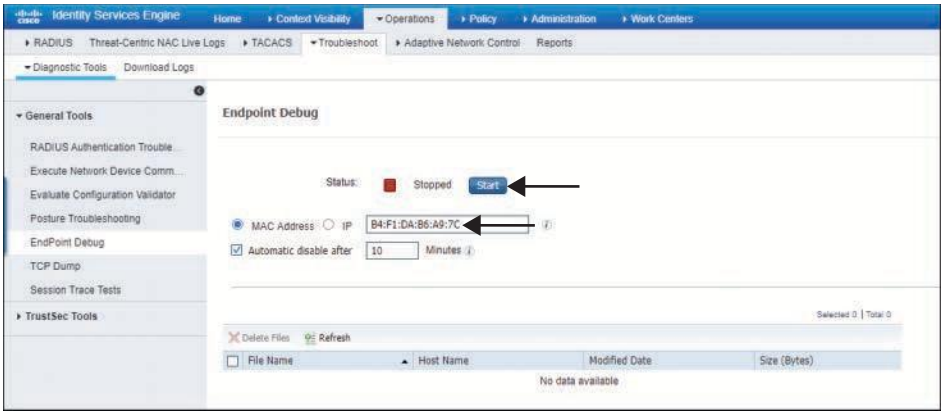


Figure 21-47 *Endpoint Debug Tool*

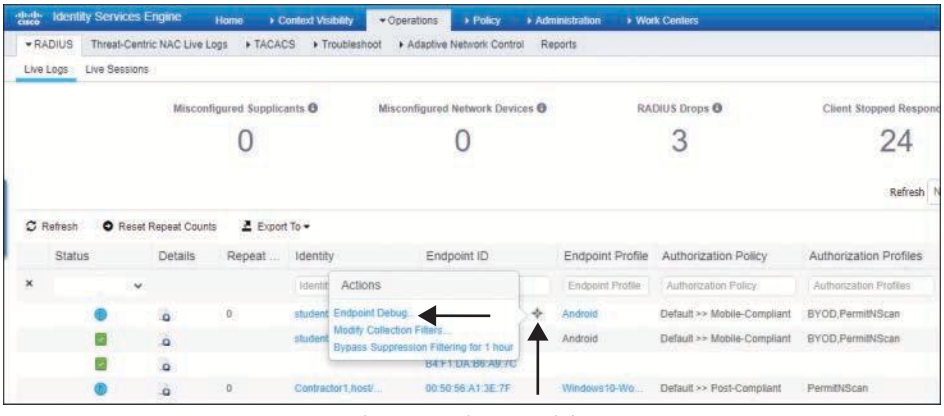


Figure 21-48 *Opening the Endpoint Debug Tool from Live Log*

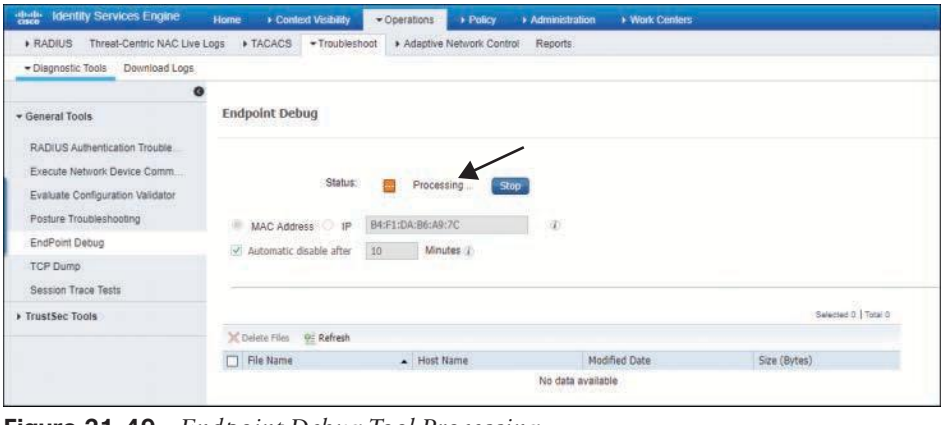


Figure 21-49 *Endpoint Debug Tool Processing*

When you run the Endpoint Debug tool, the tool creates a single file for the endpoint on each ISE node where that endpoint was active. All files are listed in the single GUI, as you can see in Figure 21-49. To view one of the files, click on the filename, as shown in Figure 21-50, and you are prompted to download just the single file or combine all the files from all the nodes into a single file for download and easier consumption.

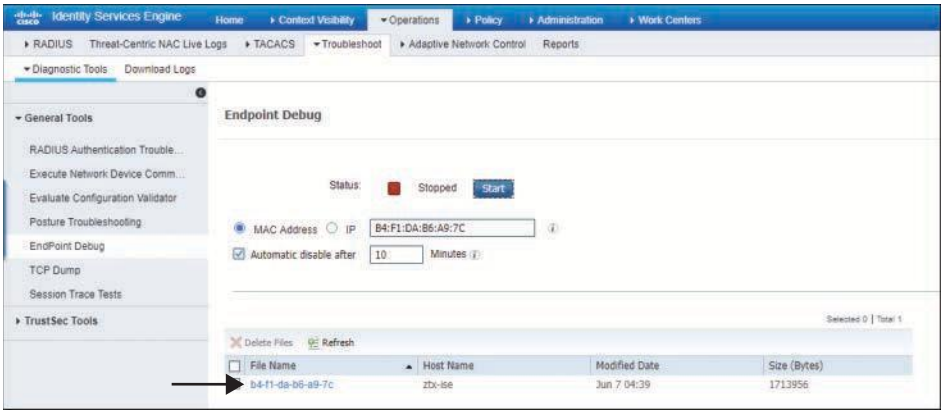


Figure 21-50 *Endpoint Debug Tool Resulting File*

If a certificate was used in the authentication, that captured certificate is listed with the same filename and a .cert extension. The certificate is downloaded in PEM format so you can more easily examine it.

TCP Dump

One of the most valuable troubleshooting options available to an ISE administrator is the ability to perform packet analysis on a RADIUS packet. Packet analysis provides visibility into the raw RADIUS attributes of an authentication, which can help you determine why a specific authentication rule was chosen or not chosen.

Cisco ISE provides the ability to perform a packet capture from any interface of any node in the entire ISE cube, download the resulting PCAP file, and perform that action from the central administrative console. You can then view the downloaded file by using your favorite packet analysis tool, such as Wireshark.

The TCP Dump tool is located under Operations > Troubleshoot > Diagnostic Tools > General Tools > TCP Dump, as shown in Figure 21-51.

With the TCP Dump tool, you can select any node of the ISE deployment and any interface of that node. You can also select the resulting packet capture format to be either the raw PCAP format or a human readable format.

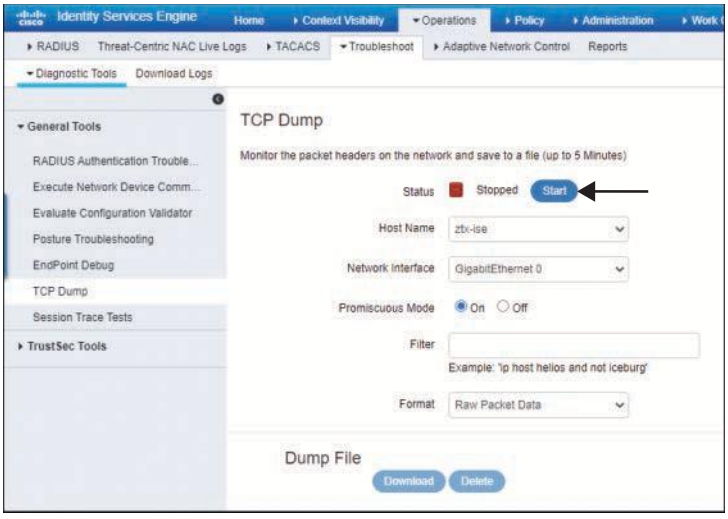


Figure 21-51 *TCP Dump*

To illustrate the value of the TCP Dump tool, this section walks through an example of using this tool to capture packets coming in to ISE and then using Wireshark to review the packet capture to examine the incoming RADIUS packet and look for the Called-Station-Id attribute. Follow these steps:

Step 1. Navigate to **Operations > Troubleshoot > Diagnostics Tool > General Tools > TCP Dump** (refer to Figure 21-51).

Step 2. Ensure that the correct ISE node (a PSN) and the proper interface (GigabitEthernet 0 for most installations) are selected.

Step 3. Click **Start**. The packet capture begins, and the approximate file size updates, as shown in Figure 21-52.

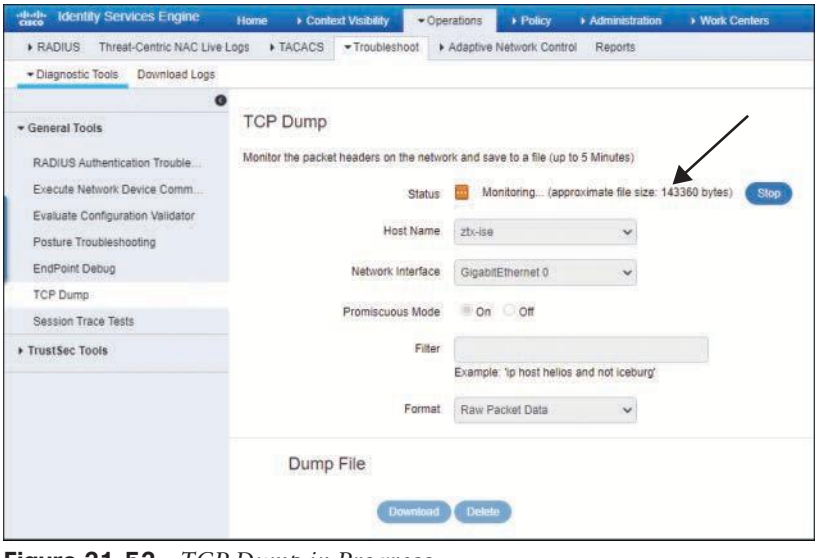


Figure 21-52 *TCP Dump in Progress*

Step 4. From an endpoint, connect to the wireless network.

Step 5. In ISE, click **Stop** when you are finished. The information in the Dump File section is updated on the screen, and the Download and Delete buttons are enabled, as shown in Figure 21-53.

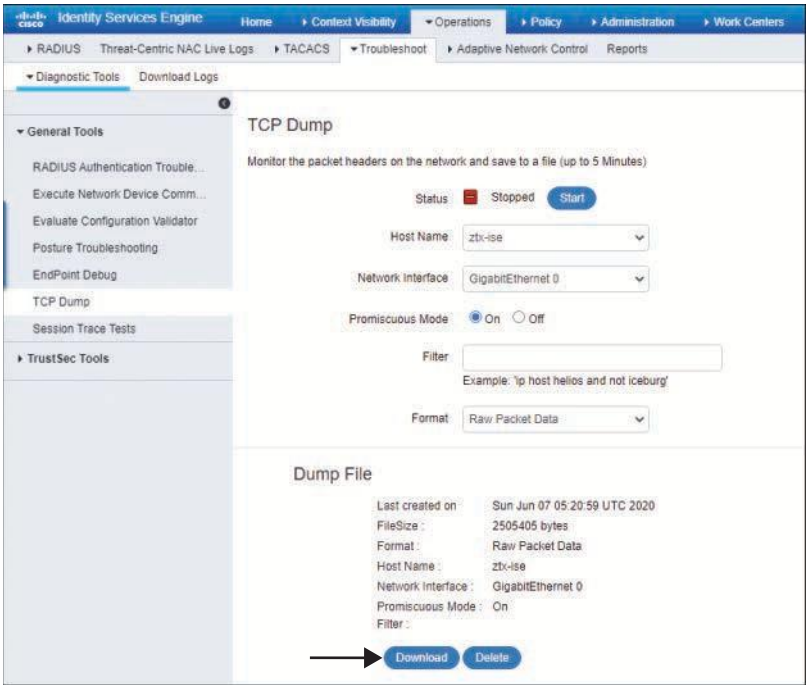


Figure 21-53 *TCP Dump Finished*

Step 6. Click **Download** to retrieve the PCAP file.

Step 7. Open the resulting PCAP file with Wireshark or another packet analysis tool and, if possible, narrow the results by applying a filter (such as "radius" to see only RADIUS packets), as shown in Figure 21-54.

In the results shown in Figure 21-54, you can see that the Called-Station-Id contains the MAC address of the wireless access point with the SSID string (ZTX-ISE) appended to it. This is a configurable setting in the wireless LAN controller, as shown in Figure 21-55.

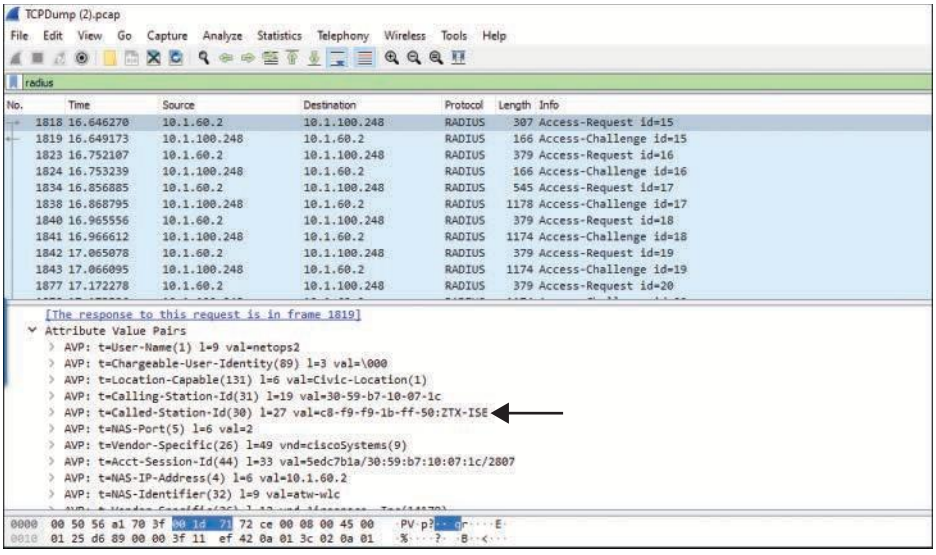


Figure 21-54 *Called-Station-Id Set to the AP's MAC Address and SSID*

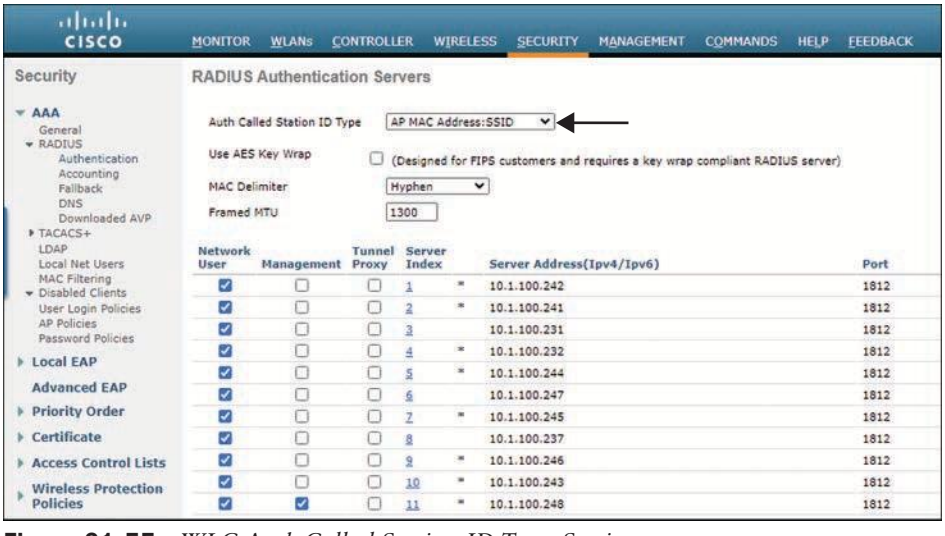


Figure 21-55 *WLC Auth Called Station ID Type Setting*

Session Trace Tests

The Session Trace tool was added in ISE Version 2.2, and you can think of it as a packet tracer for ISE. The packet tracer feature in the ASA firewall and ASDM GUI allows an administrator to test what the firewall would be expected to do with a packet matching

certain traits. Session Trace is a similar tool for ISE. The tool was designed by Aaron Woland, Jesse Dubois, and Vivek Santuka, as well as a slew of developers at Cisco, including Douglas Gash and Eyal Keren. Administrators, partners, and advocates for ISE had begged for such a tool since ISE Version 1.0.

Just as with the Endpoint Debug tool, there are two ways to launch the Session Trace tool. One way is to navigate directly to it at Operations > Troubleshoot > Diagnostics Tool > General Tools > Session Trace Tests and create a test by clicking Add and filling in all the attributes, as shown in Figure 21-56.

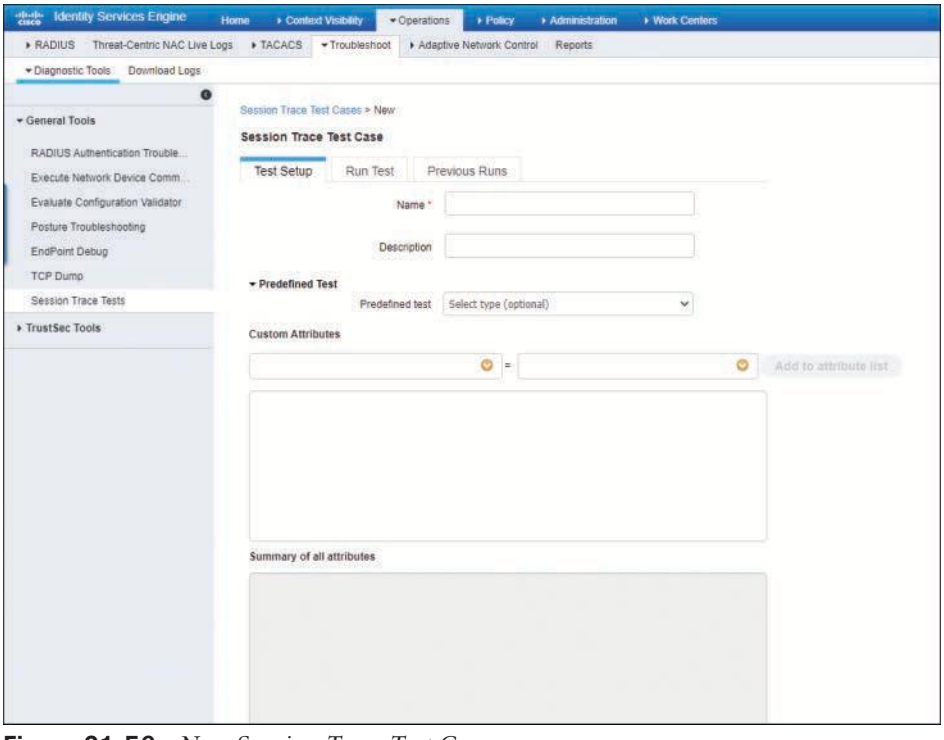


Figure 21-56 *New Session Trace Test Case*

However, it is much easier to start with an existing session. You can launch the Session Trace tool, with the test prepopulated with all the attributes of an existing session, from the Live Log screen by clicking on the target icon for details and then clicking on **Show Session Trace**, as shown in Figure 21-57.

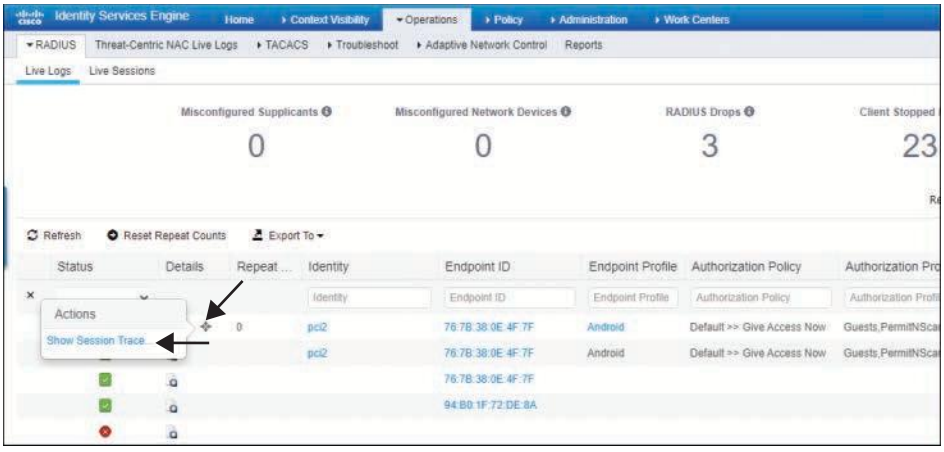


Figure 21-57 *Session Trace from Live Log*

Figure 21-58 shows the Session Trace tool prepopulated with the information from the session called from Live Log.

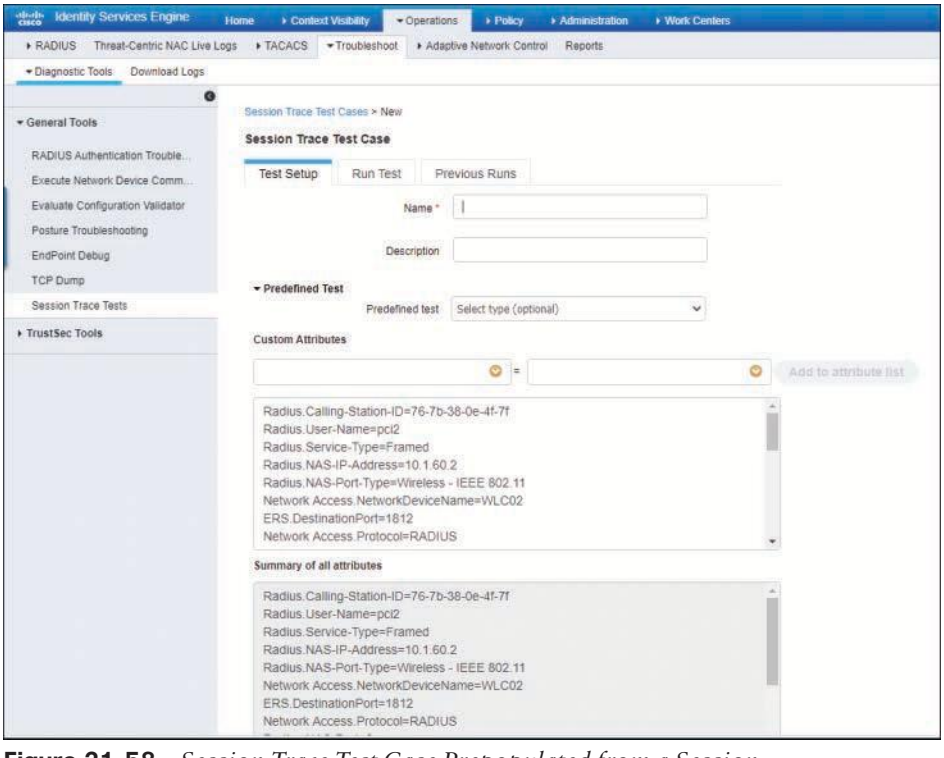


Figure 21-58 *Session Trace Test Case Prepopulated from a Session*

The first screen of the Session Trace tool is the Test Setup tab, shown Figure 21-58. Even though the tool is prepopulated, you can add other attributes by using the dropdowns under Custom Attributes.

You can modify any of the attributes from within the Test Setup tab. For example, you could try a different username or try changing the endpoint profile. This allows you to play around with what-if scenarios to see what authentication and authorization rules would be used, in which policy sets, and with what end results.

You can click Submit to save a test. Next, click on the Run Test tab of the tool, where you can select which ISE node to run the test against. The results are displayed right on the same page, as shown in Figure 21-59.

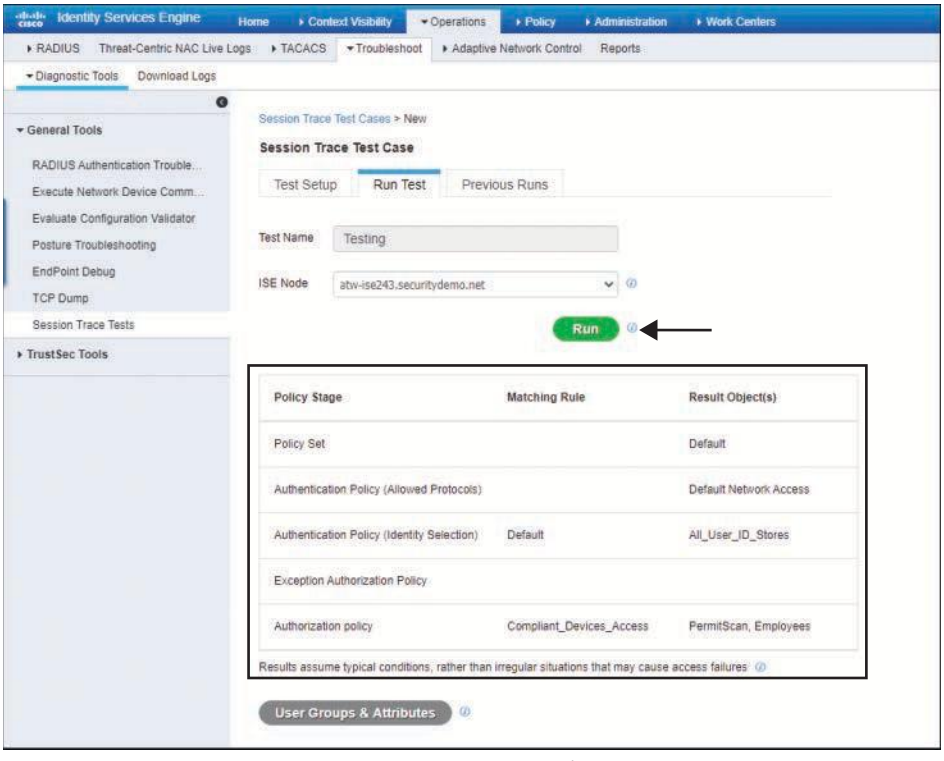


Figure 21-59 *Session Trace Test Case Run Test Tab*

Troubleshooting Methodology

As you read this section, keep in mind the tip from the beginning of the chapter: Stay calm, take your time, and think about how the solution works. Taking your time may sound counterproductive, but when you rush to fix a problem, you often end up taking much longer to reach a resolution. By staying calm, taking your time, and thinking through the flows, you may actually be able to come to a solution very quickly.

Figure 21-60 shows an example of authentication and authorization flows.

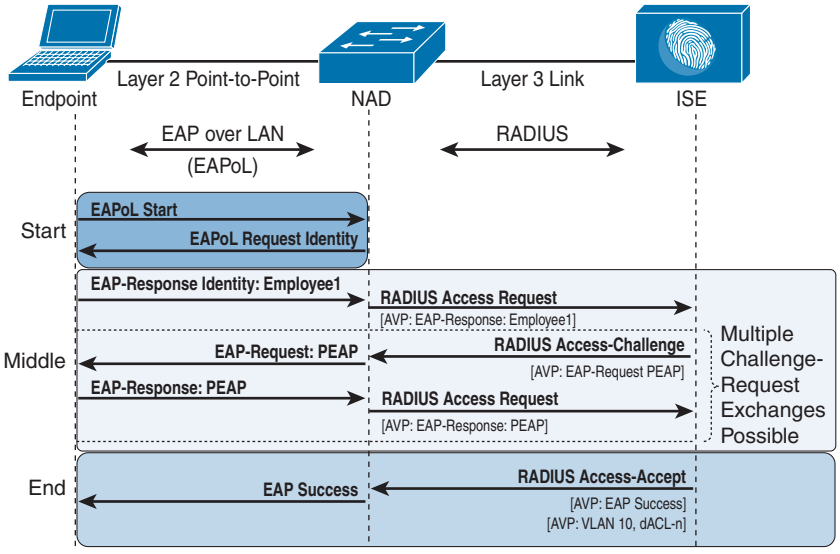


Figure 21-60 *Basic Authentication and Authorization Flows*

Log De-duplication

Prior to ISE Version 1.2, every authentication request created a 12 KB log record that needed to be stored. Therefore, when bad endpoint behavior causes millions of failed authentications a day, a lot of log data is being stored.

Beginning in ISE Version 1.2, ISE started to suppress anomalous clients by default; it stored only a single record and then logged each time that record was received. This saved a tremendous amount of processing effort and log storage, and it made better scalability possible.

Key Topic

De-duplication was a welcome change, but it did leave a few gaps to be addressed. Live Log is the first screen that you use when troubleshooting a login problem. However, if the entries are not showing up in Live Log because they are being suppressed, you are in a bad position with no visibility into what is going on.

There are two main ways to disable de-duplication so that you can effectively troubleshoot. The main way is to click on the actions in Live Log and select Bypass Suppression Filtering for 1 Hour, as shown in Figure 21-61. When you do this, you set a collection filter for the endpoint that bypasses the de-duplication for 60 minutes. After that hour, suppression is reenabled for that endpoint. In Figure 21-61 you can see the option Modify Collection Filters; this option allows you to change an existing filter to extend the time to disable the filter and more (see Figure 21-62).

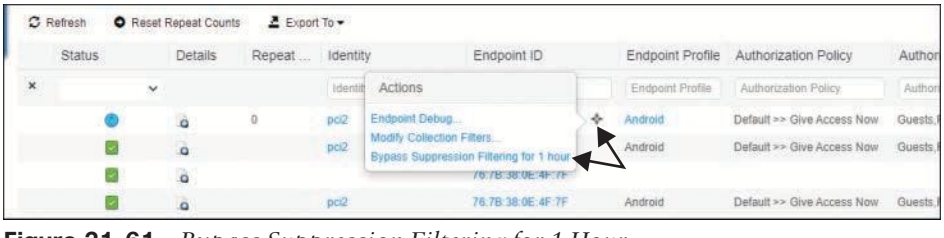


Figure 21-61 *Bypass Suppression Filtering for 1 Hour*

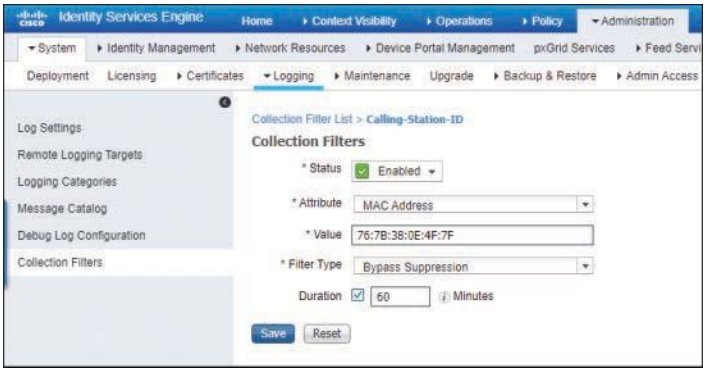


Figure 21-62 *Collection Filters*

De-duplication is a global setting that can be disabled. However, it is not recommended to disable this suppression in a production network unless directed to do so by Cisco TAC. The de-duplication setting is found under Administration > System > Settings > Protocols > RADIUS > Suppress Repeated Failed Clients, as shown in Figure 21-63.

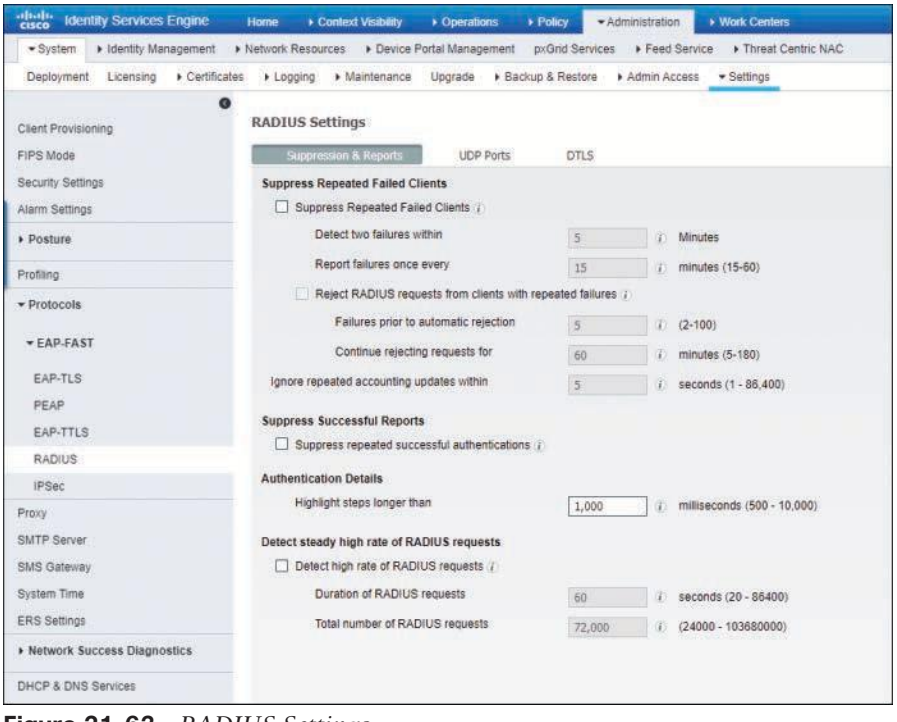


Figure 21-63 *RADIUS Settings*

Figure 21-63 shows that the RADIUS Settings screen offers numerous settings, including the following:

- **Detect two failures within**: This setting flags misbehaving supplicants when they fail authentication more than twice per interval.
- **Reporting failures once every**: This setting sends the alarm from the PSN to the MnT every specified number of minutes.
- **Reject RADIUS requests from clients with repeated failures**: With this setting, once the endpoint is in the reject interval, any requests with the same Calling-Station-Id (MAC address), NAD (NAS IP address), and failure reason are sent an Access-Reject, and the counter increments by 1 + timestamp. A log of these events is sent at the reporting interval specified on this screen.

NOTE A successful authentication clears all flags.

  - **Failures prior to automatic rejection**: This setting determines how many failed authentications before beginning the rejection of the clients.
  - **Continue rejecting requests for**: You can specify the length of time automatic rejections are sent before new attempts are allowed.
- **Ignore repeated accounting updates within**: With this setting, you can minimize the number of RADIUS accounting updates from the same source in order to prevent the session from updating too often. You can also de-duplicate successful authentications.
- **Suppress repeated successful authentications**: This setting applies de-duplication and suppresses the logs from the MnT node.

The USERNAME User

At some point in ISE's history, it was decided that ISE needed to hide the username field whenever an authentication fails. This was deemed the best security practice because it is possible for a user to accidently type a password into the username field during an attempted authentication—and it would be dangerous for passwords to show up in the Identity column. Therefore, the default setting is to show the string USERNAME in the Identity column for all failed authentications, as shown in Figure 21-64.

This makes troubleshooting very difficult because USERNAME does not tell you who is failing the authentication. The setting can be disabled permanently or for a specific time interval under the Administration > System > Settings > Security Settings, as shown in Figure 21-65.

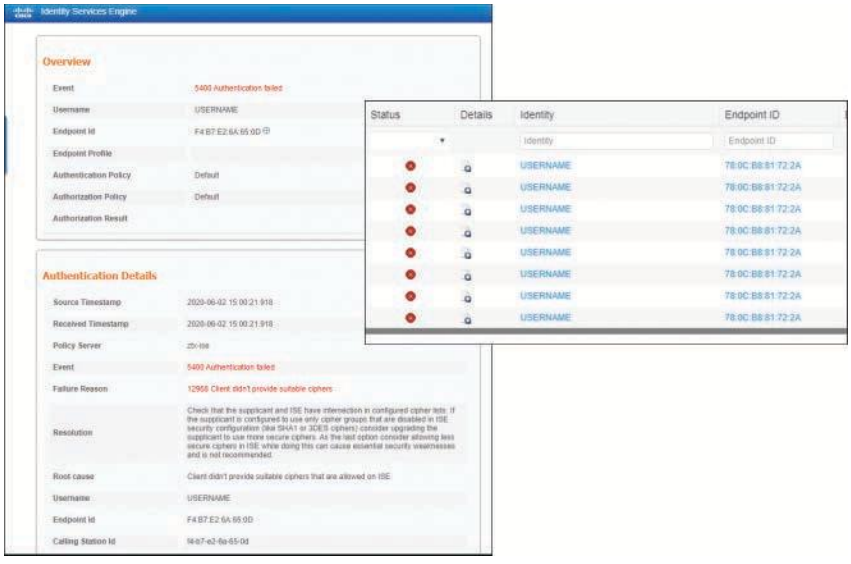


Figure 21-64 *USERNAME as the Identity*



Figure 21-65 *Security Settings*

Troubleshooting Outside of ISE

As described throughout this chapter and this book, ISE provides numerous tools that provide visibility into authentication. Although tools like Live Log, Live Sessions, and TCP Dump are very valuable when troubleshooting, they cannot give a complete picture of all the parts of authentication.

Think back to the basics of network authentication and recall that there are three "actors," or participants, in any authentication scenario: the endpoint, the NAD, and the authentication server. ISE is the authentication server, and it is very good at providing troubleshooting details. In the following sections look at the other two authentication participants: the endpoint and the NAD.

Endpoint Diagnostics

There are multiple aspects to troubleshooting from an endpoint. Naturally, there is the supplicant that is authenticating via 802.1X, but there are also other aspects that require visibility from time to time. DNS resolution is a critical aspect and may need to be validated from the client; in addition, you may need to perform activities such as BYOD onboarding with the Network Setup Assistant apps or the posture agent (NAC Agent).

Cisco AnyConnect Diagnostics and Reporting Tool (DART)

When troubleshooting 802.1X, you may need to be able to see exactly what is happening on the client side, not just what is happening on the server side. AnyConnect Network Access Manager (NAM) performs detailed logging and includes an optional module called the Diagnostics and Reporting Tool (DART).

You can use DART to create a support bundle that includes all relevant system log files and detailed communication from the supplicant; you can send this bundle to Cisco TAC. This ability can be incredibly valuable when troubleshooting 802.1X authentications.

To create a DART bundle on a Windows endpoint, follow these steps:

Step 1. Launch the DART tool by selecting **Start > Cisco > Cisco AnyConnect Secure Mobility Client > Cisco AnyConnect Diagnostics and Reporting Tool**, as shown in Figure 21-66.



Figure 21-66 *Launching DART from the Start Menu*

Step 2. Click **Next** to begin the bundle creation, as shown in Figure 21-67.



Figure 21-67 *Clicking Next*

Step 3. Choose the bundle creation option (the default is normally used) and click **Next**, as shown in Figure 21-68. The DART tool gathers all the necessary system and application logs, as shown in Figures 21-69 and 21-70.



Figure 21-68 *Choosing the Bundle Creation Option*

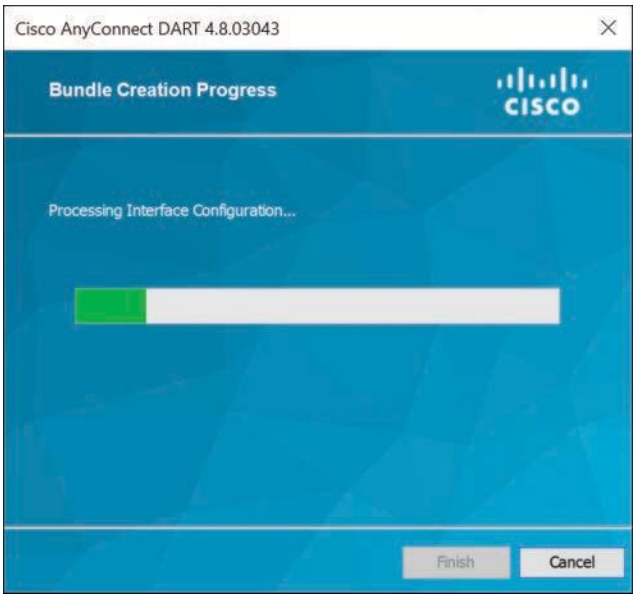


Figure 21-69 *Processing the Interface Configuration*

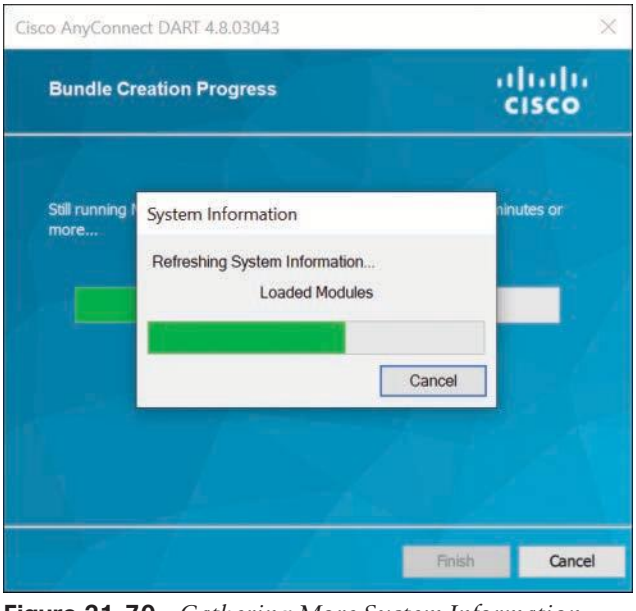


Figure 21-70 *Gathering More System Information*

Step 4. When the bundle creation is complete, find the resulting .zip file on the desktop, as seen in Figure 21-71, and email it to Cisco TAC.

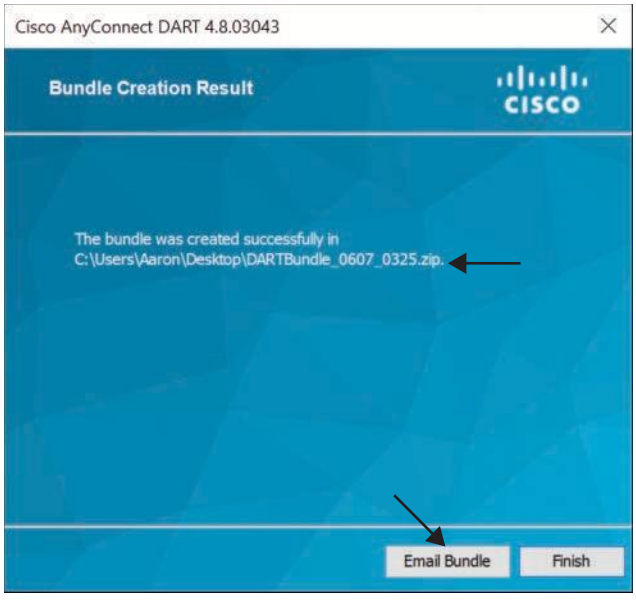


Figure 21-71 *Resulting Bundle .Zip File on the Desktop*

Supplicant Provisioning Logs

If an endpoint attempts a BYOD onboarding flow but is not able to complete the native supplicant provisioning process, examining the local endpoint logs may be the only way to find the root cause. Each operating system stores the logs in a different location, and you might even need special software to examine the logs.

The following are useful techniques that can help you troubleshoot client-side logging issues:

- Enter the **Log %temp%\spwProfileLog.txt** command to view the client-side logs for Microsoft Windows applications.
- Enter the **/sdcards/downloads/spw.log** command to view the client-side logs for Android applications.
- For MAC OSX, use the Console application and look for the SPW process.
- For Apple iOS, use the iPhone Configuration Utility (iPCU) to view messages. (Download iPCU from http://support.apple.com/downloads/.)

Network Device Troubleshooting

Seeing the authentication attempts and the authorization results from the point of view of the NAD (the enforcement device) is invaluable when troubleshooting.

Show Authentication Session Interface

Every ISE administrator should know how to use the **show authentication interface** command. As described in Chapter 11, "Implement Wired and Wireless Authentication," the

output of this command shows the current state of the authentication and the results sent back from ISE. As you see in the highlighted output in Example 21-6, this command is used to show that the WebAuth downloadable ACL (dACL) is applied to the session, the ACL-WEBAUTH-REDIRECT redirection ACL is applied, the session is being redirected to a CWA portal on the ISE node, it is assigned an SGT value of 8, and it is the result of MAC Authentication Bypass (MAB).

Example 21-6 *show authentication session interface Command Output*

```
3750-X# show authentication session interface g1/0/2
            Interface:  GigabitEthernet1/0/2
          MAC Address:  0050.5687.0004
           IP Address:  10.1.10.50
            User-Name:  00-50-56-87-00-04
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
        Authorized By:  Authentication Server
          Vlan Policy:  N/A
              ACS ACL:  xACSACLx-IP-WebAuth-53e3ffc0
     URL Redirect ACL:  ACL-WEBAUTH-REDIRECT
         URL Redirect:  https://atw-cp-ise02.ise.local:8443/guestportal/gateway?
sessionId=C0A8FE0100009884CAE356DA&action=cwa
                 SGT:  0008-0
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  C0A8FE0100009884CAE356DA
      Acct Session ID:  0x00013D98
               Handle:  0x0C0008C4

Runnable methods list:

       Method   State
       mab      Authc Success
       dot1x    Not run
```

Viewing Client Details on the WLC

A wireless LAN controller (WLC) has view that is similar to the show authentication details command on a Cisco Catalyst switch, but it is related to the client instead of the interface. You can view the client details by navigating to Monitor > Clients, as shown in Figure 21-72.

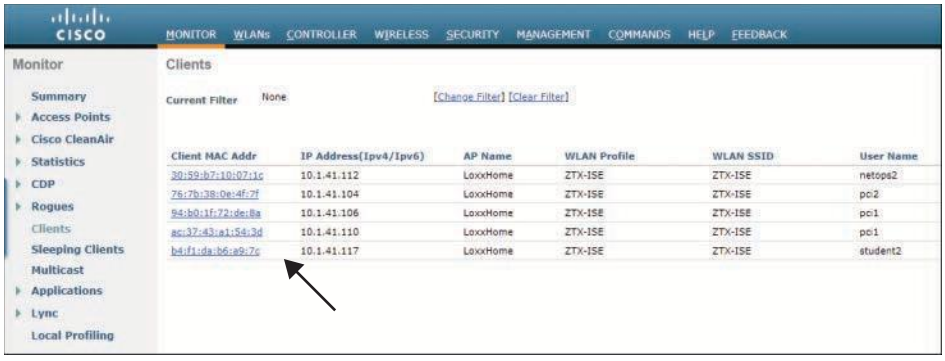


Figure 21-72 *Monitor > Clients*

You can select the client that you are interested in verifying by clicking its MAC address. This brings up the client details screen (see Figure 21-73). In the Security Information section in Figure 21-73, you can see that RADIUS NAC State is set to CENTRAL_WEB_AUTH, the ACL-WEBAUTH-REDIRECT redirection ACL is applied, and the session is being redirected to a client provisioning portal on the ISE node.

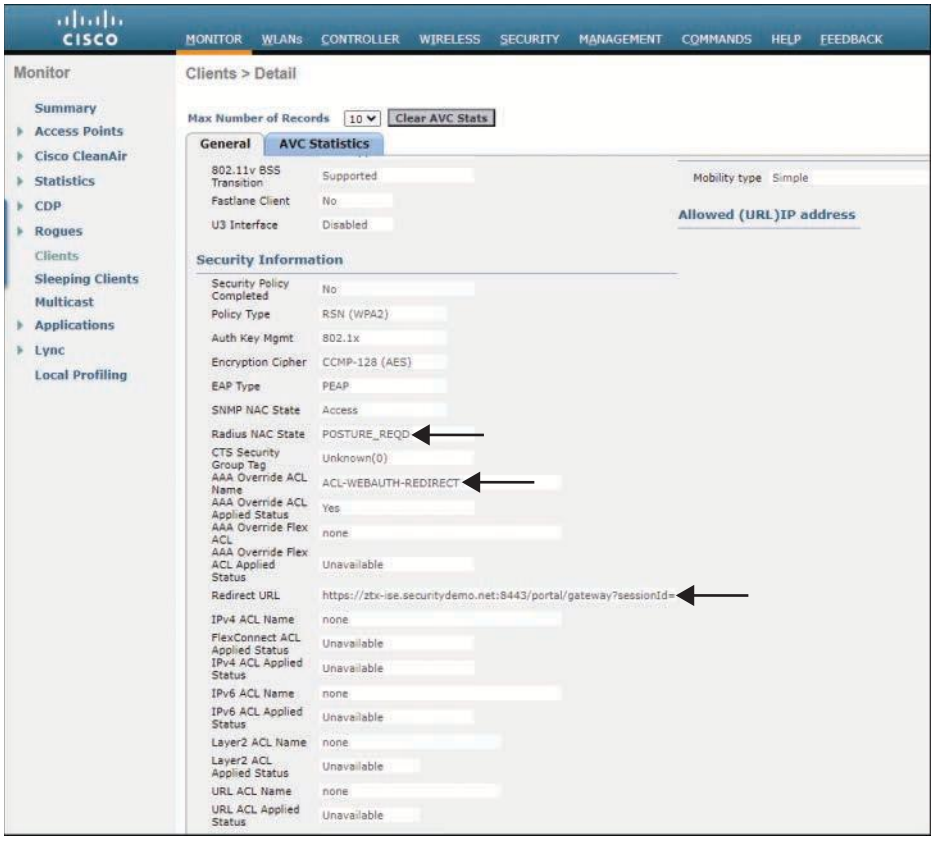


Figure 21-73 *Clients > Detail*

Debug Commands

A number of **debug** commands can be used on Cisco switches and WLCs to see authentication details. These commands are incredibly useful when you need to troubleshoot network devices.

The go-to **debug** command for a Cisco WLC is **debug client** *mac-address*. When you enter this command, all activity on the WLC related to the client is shown in the command output, and no other client traffic on the controller is affected.

Cisco switches do not offer a client debugging command, but you can constrain the **debug** command to a specific interface by using the syntax **debug interface** *interface-name*. You can use this command before entering one of the other **debug** commands to ensure that only the specified interface is affected by the debugging process.

The following **debug** commands are also useful for Cisco switches:

- **debug authentication** [**all** | **errors** | **events** | **feature** | **sync**]: Use this command to see all activity related to the authentication manager that handles 802.1X, MAB, WebAuth, and the related session activity.
- **debug dot1x** [**all** | **errors** | **events** | **packets** | **redundancy** | **registry** | **state-machine**]: Use this command to see all activity related to the 802.1X software in the switch.
- **debug epm** [**all** | **api** | **error** | **events** | **redirect**]: Use this **debug** command when troubleshooting URL redirection and application of DACLs.

NOTE For more on serviceability in ISE and troubleshooting, take a look at Aaron Woland's blog post "Troubleshooting Cisco's ISE without TAC," at https://www.networkworld.com/article/3053669/troubleshooting-ciscos-ise-without-tac.html.

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 21-2 lists these key topics and the page number on which each is found.

Key Topic

Table 21-2 Key Topics

Key Topic Element	Description	Page
Paragraph	Live Sessions and complex flows	776
Paragraph	Sending CoA from Live Sessions	776
Paragraph	Using the Cisco PKI option to help with TAC case automation	784
Paragraph	Endpoint debugging	796
Paragraph	De-duplication	805

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What does it mean when the Identity column in Live Log shows the string USERNAME?
2. What steps should you take if an expected authentication is not showing up in Live Log?
3. What does the Posture Troubleshooting tool do?
4. What is the function of the Endpoint Debug tool?
5. What is the function of the Diagnostic and Reporting Tool?

This page intentionally left blank



CHAPTER 22

ISE Context Sharing and Remediation

This chapter covers the following topics:

Platform Exchange Grid

Context-In

Rapid Threat Containment

Because ISE is positioned to know exactly who and what is on a network at any given time as well as assign different levels of access and context assignments with Security Group Tags, it is the perfect security tool to be at the center of a security ecosystem.

As a network administrator, you have many tools in your security toolbox, including firewalls, next-generation firewalls (NGFWs), intrusion prevention systems (IPSs), NG IPSs, security information and event management (SIEM) systems, secure web gateways, threat defense tools, vulnerability assessment scanners, mobile device managers, and more. Most of these tools do not know the identity of the user or the context of the endpoint. Instead, these tools may only be able to identify the IP address or MAC address of the endpoint. By being integrated into a security ecosystem with ISE, these tools become even more valuable. Wouldn't the reporting in a SIEM system provide more value if it were to show which user was involved in a security event instead of just an IP address? And when your IPS or threat defense solution identifies malicious activity on the network, wouldn't it be great if they could trigger something that would change the way the endpoint is treated on the network? With a single trigger, you could change the endpoint's level of network access, which would result in the Firepower Threat Defense firewall performing a deeper inspection on the endpoint's traffic, the Web Security Appliance (WSA) applying a different SSL decryption policy, and more. ISE can provide a single point of policy control for such threat containment and context sharing.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 22-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 22-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Integration Types in the ISE Ecosystem	1–3
pxGrid	4–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following does pxGrid not perform?
   a. Threat feed ingestion
   b. Context brokering
   c. Context sharing
   d. Rapid Threat Containment
2. Which actions could Endpoint Protection Service (EPS) take? (Choose three.)
   a. Quarantine
   b. Unquarantine
   c. Internet only
   d. Shutdown
   e. Profile
3. What does Adaptive Network Control (ANC) offer that EPS Version 1.0 did not?
   a. The ability to hard shut down ports
   b. The ability to change an endpoint's Security Group Tag (SGT)
   c. The ability to change the level of access for an endpoint
   d. The ability to create different custom labels to use in policies
4. What feature in pxGrid allows ISE to receive information for profiling purposes?
   a. Context-In
   b. Context sharing
   c. Rapid Threat Containment
   d. Context brokering
5. How many pxGrid nodes can you have in an ISE cube on ISE Version 2.4?
   a. 4
   b. 2
   c. 50
   d. 6

6. What are the basic requirements for adding a pxGrid participant? (Choose three.)
   a. The ISE certificate authority must be trusted.
   b. All the ISE node hostnames must be included in the pxGrid certificate SAN.
   c. A pxGrid certificate must be installed for the participant.
   d. The IP address or FQDN of the pxGrid controller must be configured.
   e. The exact pxGrid participant must have the certificate for each of other pxGrid participants.
7. Beginning in pxGrid Version 1, what was the original communication protocol?
   a. XMPP
   b. WebSocket
   c. TLS 1.3
   d. MQTT
8. With ISE Version 2.3, what was added to modernize the pxGrid interface?
   a. XMPP
   b. WebSocket
   c. TLS 1.3
   d. MQTT
9. In WSA Version 11.7 and later, Active Directory groups can be downloaded from ISE by using which of the following?
   a. pxGrid
   b. SSH
   c. ERS
   d. EPS
10. What pxGrid enhancement was added to Stealthwatch in Version 7.0?
   a. ERS integration
   b. EPS 1.0
   c. ANC
   d. SXP

Foundation Topics

Integration Types in the ISE Ecosystem

Key Topic

An *integration* might be ISE sharing data outbound or it might be ISE steering traffic toward another solution. The integration method could involve ISE receiving information inbound for use in ISE's own network access policies or ISE brokering data exchange between other members of the security system.

MDM Integration

In Chapter 16, "Bring Your Own Device," you read about BYOD and the integration between ISE and mobile device management (MDM). That integration is twofold: ISE

provides the redirection to the MDM service for onboarding, but the mobile device manager is also able to provide Context-In to ISE. In other words, the mobile device manager tells ISE about the mobile endpoints, compliance with the MDM endpoint policies, the status of encryption or PIN lock, and more.

This integration uses a specific bidirectional application programming interface (API) between ISE and the mobile device manager (cloud service or appliance). This API is unique and created just for MDM integration.

Rapid Threat Containment

MDM is one of the first and most common integration types for ISE. In true Cisco marketing fashion, the next integration type, ***Rapid Threat Containment***, has gone through many different names and marketing initiatives.

Key Topic

A feature called ***Endpoint Protection Service (EPS)*** was added to ISE Version 1.1. EPS provided an API to allow other applications to initiate three actions against an endpoint, based on IP address or MAC address:

- **Quarantine:** The Quarantine action set the binary flag on the endpoint record to true, added the endpoint to a list of quarantined endpoints, and allowed an administrator to create authorization policies that used that assignment to specify a different level of network access.
- **Unquarantine:** This action removed the endpoint from the list of quarantined endpoints and cleared the binary flag.
- **Shutdown:** This action was supposed to send a Change of Authorization (CoA) terminate to the network and shut down the port on the network switch. However, this option existed in the API but was not exposed to the policy and was therefore not usable.

Many of the first integrations with ISE used EPS including the original integration with Lancope Stealthwatch (now Cisco Stealthwatch). With this integration, an endpoint could be quarantined from the Stealthwatch user interface.

Figure 22-1 illustrates a flow with Stealthwatch initiating an EPS quarantine.

The flow illustrated in Figure 22-1 shows an endpoint being admitted to the network with full access. The Stealthwatch administrator initiates a quarantine, and Stealthwatch connects to ISE by using the EPS REST API. ISE quarantines the endpoint with the specific IP address. ISE then adds the endpoint to the EPS list, sets the flag on the endpoint object, and sends a CoA to the network.

When the new access request comes in, a rule created with the EPSStatus condition is matched. Figure 22-2 shows this condition.

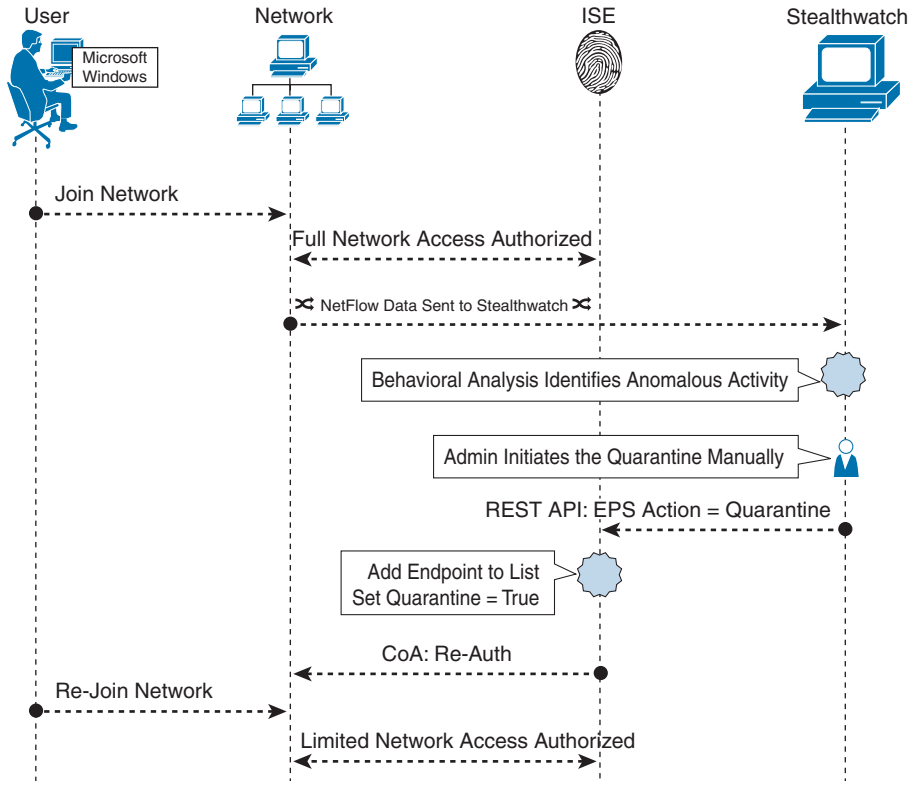


Figure 22-1 *Stealthwatch to ISE: EPS Quarantine*

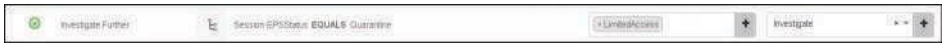


Figure 22-2 *The EPSStatus Authorization Rule*

The resulting network authorization may provide for limited access or even set a new Security Group Tag (SGT) that can be acted on differently at miscellaneous points in the network, such as Cisco Web Security Appliance.

Ultimately, EPS was too rigid as it provided for only a single actionable classification (quarantine). Users needed flexibility—not just to have many different options but also for integration into the newfangled context-sharing technology that Cisco was creating, named pxGrid. Therefore, it needed to evolve into EPS 2.0 or something like that.

ISE Version 1.3 introduced ***Adaptive Network Control (ANC)***, which was basically a renamed version of EPS. ISE Version 1.4 actually added new functionality to ANC. While it still supported the old EPS API calls for backward compatibility, ISE Version 1.4 also added a new API with different labels available, as well as the ability to create your own label.

ISE refers to these labels as ANC policies, but there is no policy involved here. It is just a tag or a label that gets assigned to an endpoint object. Those labels can be used as conditions in an authorization rule to invoke some action such as changing the authorization level and assigning a new SGT.

While you can add many different labels, there are only three choices for ANC types: Quarantine, Shut_Down, and Port_Bounce (which determines the CoA type used when the label is applied to the endpoint).

To create an ANC policy (that is, label), navigate to Operations > Adaptive Network Control > Policy List and click Add, as shown in Figure 22-3.

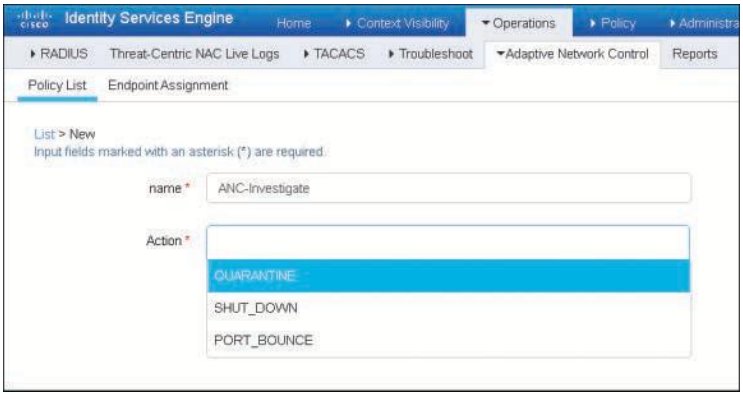


Figure 22-3 *Adding an ANC Policy*

You can create multiple ANC policies, and each policy can contain one or more actions. Each ANC policy can be associated with a different authorization result. For example, you might end up with ANC policies such as the following:

- Investigate
- Phasers on Stun
- Eradicate
- Nuke from Orbit

In addition to providing a much more flexible approach to classification (or, as Cisco's legendary Paul Forbes would say, "flexible name spaces"), ANC also integrates tightly with pxGrid. It allows pxGrid subscribers to trigger the ANC action in the pxGrid connection instead of in the point API, as it was done previously.

To recap, Endpoint Protection Services was renamed Adaptive Network Control. Then Adaptive Network Control got new functionality in ISE Version 1.4. Cisco then came up with a new naming convention to refer to the entire integrated security system where any Cisco security product takes action through another Cisco security product: Rapid Threat Containment. So, for example, you can use Rapid Threat Containment with Cisco Stealthwatch and Rapid Threat Containment with Cisco Firepower Management Center and Identity Services Engine.

While ISE is often the center of a security ecosystem, the Rapid Threat Containment portfolio includes more than just integrations with ISE. There are solutions like Rapid Threat Containment with Firepower Management Center and Cisco Stealthwatch, Firepower and Cisco Tetration, and many more. Actions taken using Cisco Threat Response are also part of the Rapid Threat Containment umbrella. Crystal clear, right?

Platform Exchange Grid

Platform Exchange Grid (pxGrid) is Cisco's premier publish/subscribe (pub/sub) communication bus. It was designed from the ground up to be a scalable secure data sharing system.

Like most other next-generation AAA solutions, ISE originally started sharing information through the use of APIs. People quickly recognized that point APIs would not scale to the level of data that needed to be shared and the scale at which it was requested.

Cisco went down the path of a pub/sub bus, similar to the way Call Manager works, and there is a controller that keeps track of all the topics that exist. A ***topic*** is a list of information that is available. A topic might be session data related to who and what is on the network, or it might be a list of vulnerable endpoints and list of vulnerabilities.

pxGrid participants can subscribe to any topic of interest and can be notified when there is data to be retrieved. Those participants are known as ***subscribers***. The true source of the data can be any other pxGrid participant; these participants are known as ***publishers***. A publisher registers a topic with a controller, which performs the authorization for each subscriber to retrieve the data from the miscellaneous publishers.

Figure 22-4 shows the standard Cisco drawing that is often used to explain pxGrid. In this illustration, you can see many different types of products, each with different information to publish and each needing information from one of the other products.

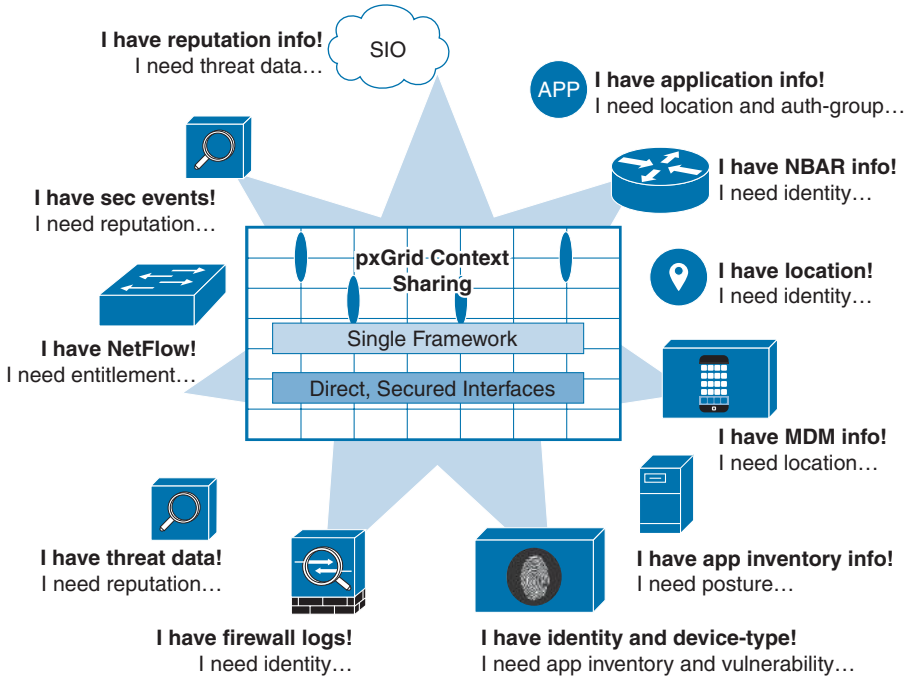


Figure 22-4 *Standard Cisco pxGrid Illustration*

pxGrid was initially added to ISE in Version 1.3, so it's been around for a while now and has an ecosystem of partner applications that continues to grow at a very rapid pace.

ISE Version 2.2 made strides in enhancing pxGrid, mostly in terms of ease of use in configuration and maintenance. ISE Version 2.2 also added more information into ISE's pxGrid topics for consumption by the subscribers. A specific example of additional information that was added to pxGrid topics is the list of groups of which each active user is a member. That list of groups was shared within the same topics used in previous ISE versions, thus seamlessly enabling backward compatibility.

pxGrid

Key Topic

pxGrid Version 1 was created by extending the Extensible Messaging and Presence Protocol (XMPP), which is also the communication protocol used by Jabber. In fact, the pxGrid controller is a modified Extensible Communication Platform (XCP) server. (For more on XMPP, see www.xmpp.org.)

The XCP needs a client that knows how to communicate with it. Cisco DevNet partners can create applications that use the pxGrid common library (GCL) to join the pxGrid controller without having to write their own client from scratch.

Beginning in ISE Version 2.3, ISE added a modernized WebSocket-based interface to pxGrid to make integration easier. The DevNet partners were no longer required to integrate a Java or C library into their applications. Instead, they could use common representational state transfer (REST) connections.

No matter what the version, always remember that pxGrid is made up of three main components: a controller, publishers, and subscribers (see Figure 22-5).

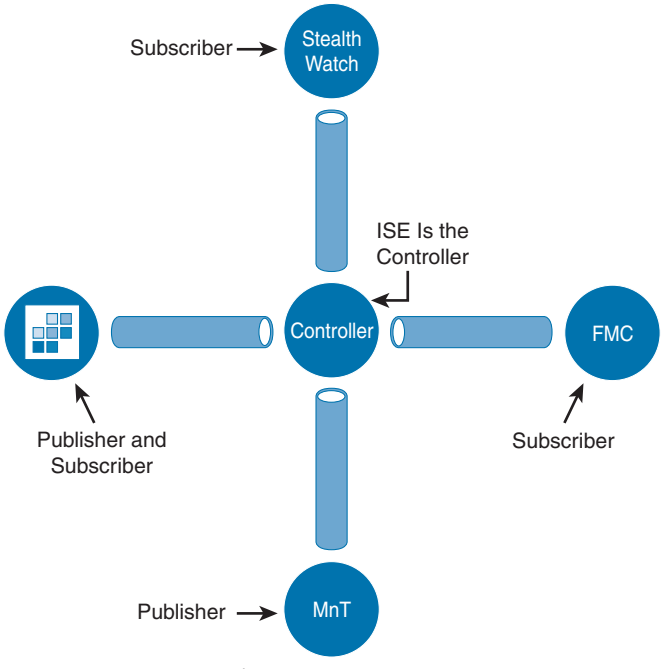


Figure 22-5 *pxGrid Components*

pxGrid in Action

pxGrid uses secure communication between participants. Therefore, certificates are very important to the success and ease of a deployment. Every pxGrid participant must trust the controller, and the controller must trust each of the participants.

The FMC needs to speak to the pxGrid controller to learn of the topics that exist and who has published those topics, and it also needs to speak directly to the MnT node to perform bulk downloads of the published session data. If the FMC were to trust the pxGrid controller's certificate but not the Monitoring and Troubleshooting (MnT) node's certificate, the communication would ultimately fail. Figure 22-6 illustrates this concept. As you can see in this figure, you end up needing a full mesh of trust between pxGrid participants. Each participant must trust the controller as well as the other participants.

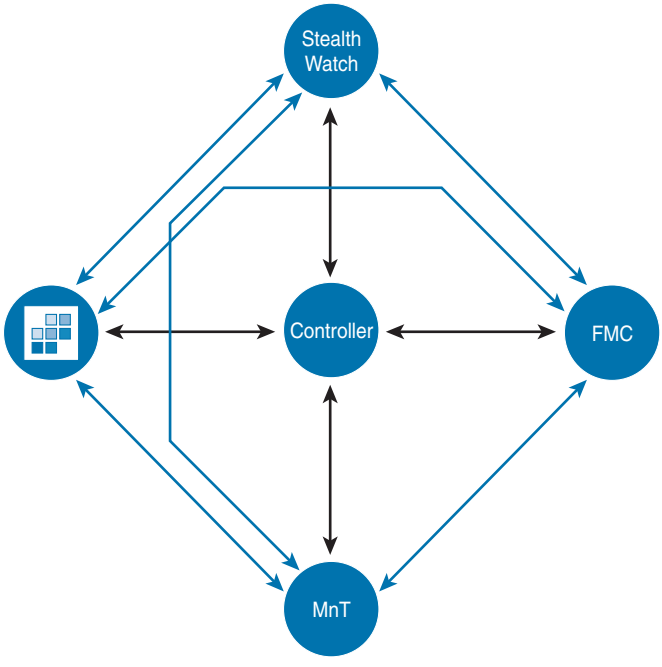


Figure 22-6 *Full Mesh of Trust*

Based on a lot of deployment experience, the best practice is to always use the same certificate authority (CA) to issue the pxGrid certificates to the participants. To make this easier, ISE's built-in CA was enhanced to issue pxGrid certificates in addition to endpoint certificates beginning in ISE Version 2.1. In addition to the enhancement to the CA, APIs were added to automate the certificate enrollment from a pxGrid ecosystem partner. In fact, Cisco's flagship DNA Center product uses exactly these APIs and this CA to integrate with ISE.

Figure 22-7 illustrates a single CA issuing the certificates to all the participants.

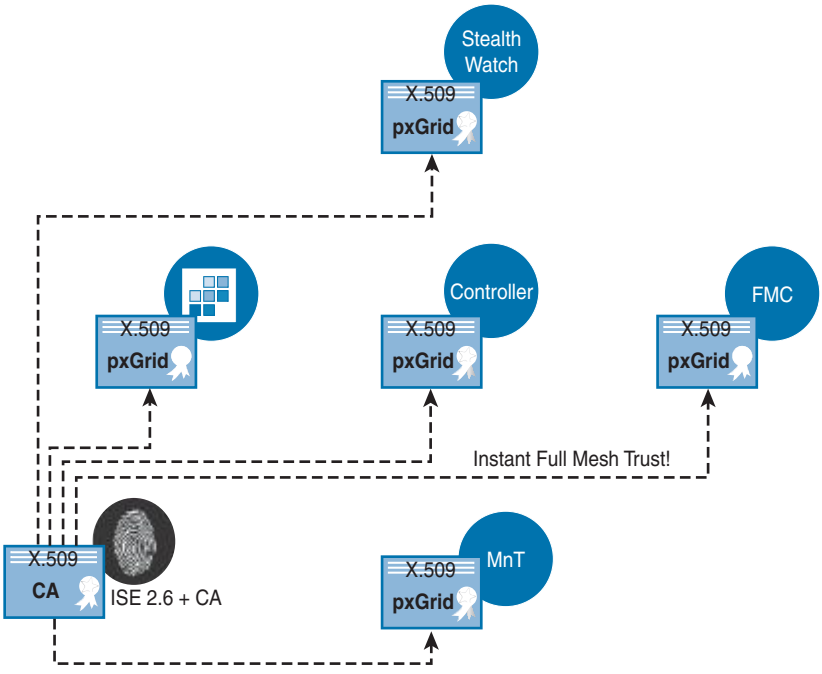


Figure 22-7 *ISE CA Issuing pxGrid Certificates to All Participants*

Context-In

pxGrid involves more than context being shared from ISE (referred to as ***Context-Out***); it is also used for sharing information between external systems. As of Version 2.4, ISE is also able to receive information through pxGrid to help with its own profiling policies. This is referred to as ***Context-In***.

In Chapter 14, "Profiling," you learned about profiling and the different probes that ISE can use. One of those probes, which was introduced in ISE Version 2.4, is the pxGrid probe, which is used to learn profiling data about endpoints through pxGrid Context-In.

The profiling probe was first used with the Cisco Industrial Network Director (IND), which communicates with industrial switches and IoT security devices and collects detailed information about the connected IoT devices. However, additional third-party ecosystem partners have been added since that time. IND Version 1.3 adds a pxGrid publisher interface to communicate to ISE IoT attributes that are leveraged in profiling (see Figure 22-8).

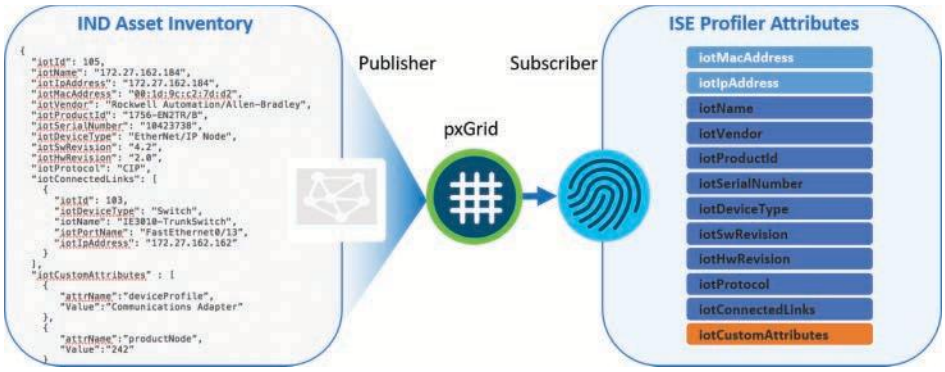


Figure 22-8 *Industrial Network Director Using ISE pxGrid Probe*

Configuring ISE for pxGrid

The pxGrid user interface is located under Administration > pxGrid Services. By default, the pxGrid services are not enabled on any ISE node, and the following message is displayed on the pxGrid Services page:

```
In order to navigate to the pxGrid Services page, pxGrid persona
must be enabled on at least one node in the ISE deployment. Please
click on this link to be redirected to the Deployment page.
```

You need to enable the pxGrid persona on at least one of the Policy Services nodes (PSNs) in your deployment. Before you enable pxGrid on any of the ISE nodes in a deployment, it is a good idea to ensure that all the nodes in the ISE cube have pxGrid certificates signed by the same certificate authority.

Beginning in ISE Version 2.2, each node's pxGrid certificate is signed automatically by the internal CA. Naturally, you can replace that certificate with one from an external CA of your choosing, but the default certificate uses the internal CA in an attempt to simplify the setup and follow best practices.

To check the pxGrid certificate in ISE:

Step 1. Navigate to **Administration > System > Certificates**.

Step 2. Select the pxGrid certificate of one of the nodes, as shown in Figure 22-9.

Step 3. Click **View**.

Step 4. Ensure that the root signer of the certificate is the primary PAN of the ISE cube (the root CA), as shown in Figure 22-10.

Once you've ensured that the certificates in use are all issued by the same PKI, it is time to enable them. Our experience-based recommendation is to have a pxGrid certificate on every single node in the ISE deployment, even if the node will not run the pxGrid controller function.

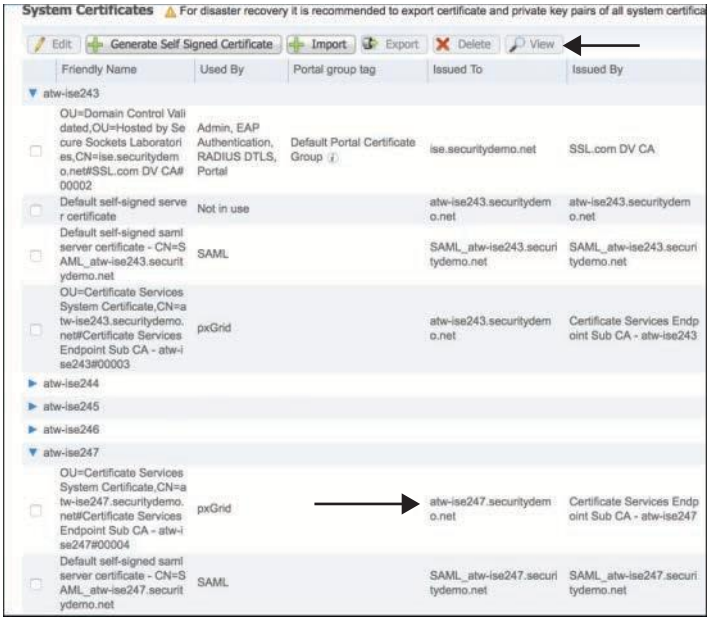


Figure 22-9 *Administration > System > Certificates*

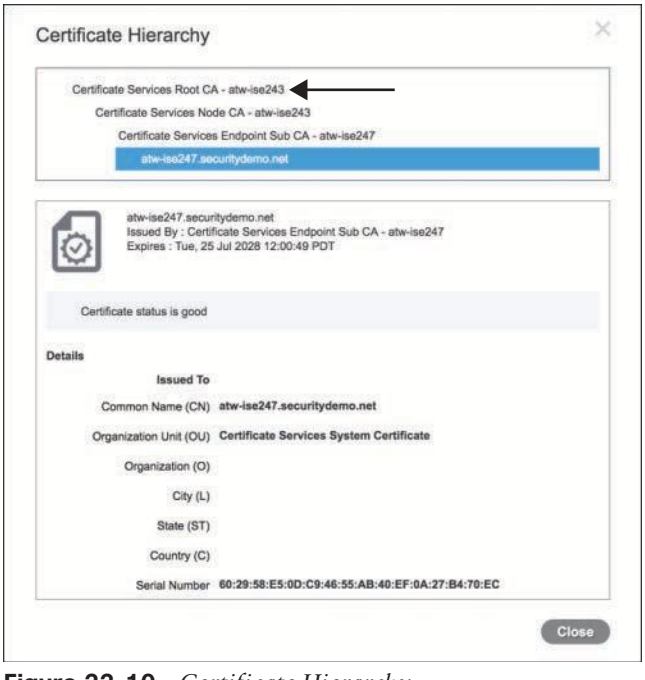


Figure 22-10 *Certificate Hierarchy*

IMPORTANT NOTE Beginning in ISE Version 2.2, all pxGrid communications occur within the secure pxGrid channel. In other words, all communication occurs by leveraging the pxGrid certificate of the ISE node. This differs from prior versions, in which all bulk downloads from the MnT node occurred using the Admin certificate rather than the pxGrid certificate. This caused a lot of confusion and thus needed to change. If you are implementing pxGrid on any ISE version prior to ISE 2.2, you must ensure that the participant trusts the CA that issued the Admin certificate, as well as the CA that issued the pxGrid certificate.

To enable the pxGrid persona on a PSN, follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > System > Deployment**.

Step 2. To ensure that the pxGrid controller function runs on a PSN, select one of the PSNs from the list.

Step 3. Select the checkbox for **pxGrid**, as shown in Figure 22-11.

Step 4. Click **Save**. This is all it takes to enable the pxGrid controller function on the PSN.

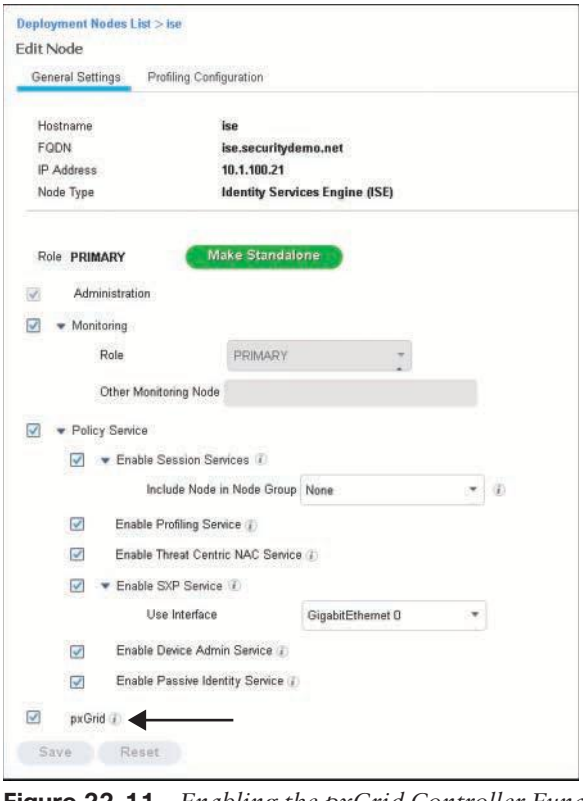


Figure 22-11 *Enabling the pxGrid Controller Function*

As of ISE Version 2.4, you can have up to four pxGrid controllers per ISE cube to provide redundancy. In previous ISE versions, only two pxGrid controllers per ISE cube were allowed.

When the pxGrid services are all up and running, the PAN and MnT node will automatically register and publish their respective topics into the grid, as shown in Figure 22-12.

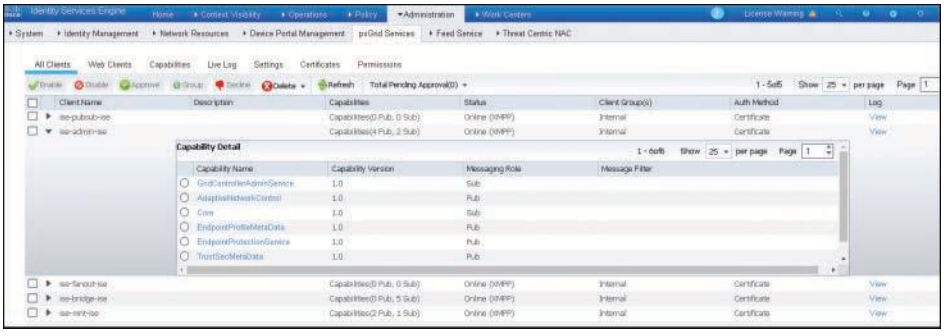


Figure 22-12 *Default pxGrid Services After Enabling the pxGrid Controller Function*

Notice in Figure 22-12 how the topics are listed under the pxGrid participant as well as the roles that each node plays with the topic (pub and sub).

By default, only ISE nodes are registered automatically. All other pxGrid participants either require manual approval or require you to enable auto-registration.

Configuring pxGrid Participants

Many different subscribers and publishers can participate in the ecosystem with pxGrid. Each one uses the information in its own way, and the integration UI is bound to be unique for each product. However, the basic requirements and configuration steps are always the same:

Step 1. Trust the ISE certificate authority.

Step 2. Install a certificate for pxGrid's own identity.

Step 3. Configure the IP address or FQDN of the pxGrid controller.

For the most part, this is all you really need to do on each participant. Some participants make things easier than others. The following sections show how to configure some of the main pxGrid participants: Cisco Firepower Management Center, Cisco Stealthwatch, and Cisco Web Security Appliance.

Configuring Firepower Management Center for Identity with pxGrid

The Cisco Firepower Management Center (FMC) is the enterprise-class device manager and security monitoring tool for Cisco's Firepower line of NGFWs and NGIPSs. The FMC has had pxGrid integration with ISE for a while, but Version 6.2 added an even better integration, with the ability to use the TrustSec data independently of user identities. FMC Version 6.5 advanced the pxGrid integration even further by allowing the FMC to subscribe to the SXP topic to receive the IP address-to-SGT bindings learned via SXP. The FMC can use context information provided by pxGrid—such as endpoint profiles, TrustSec tags, and both passive and active user identities—in its own policies.

NOTE Much like the FMC, the Firepower Device Management (FDM) solution is also capable of integrating with ISE using pxGrid, but this section is focused on the FMC integration.

Firepower Management Center leverages pxGrid to learn the context of who and what is on the network and determine the mapping of those devices to IP addresses. However, it leverages the LDAP-based realms to learn about what users and groups exist in Active Directory for the creation of access policy.

The following section shows how to configure the pxGrid integration, and the next one discusses realm configuration.

Configuring Firepower Management Center for pxGrid

Before configuring pxGrid on the FMC, you need to generate a pxGrid certificate for the FMC to use. Beginning with ISE Version 2.2, an administrator can download a CA's certificates and generate certificates directly from the pxGrid Services tab.

To generate a pxGrid certificate for the FMC, follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > pxGrid Services > Certificates**, as shown in Figure 22-13.

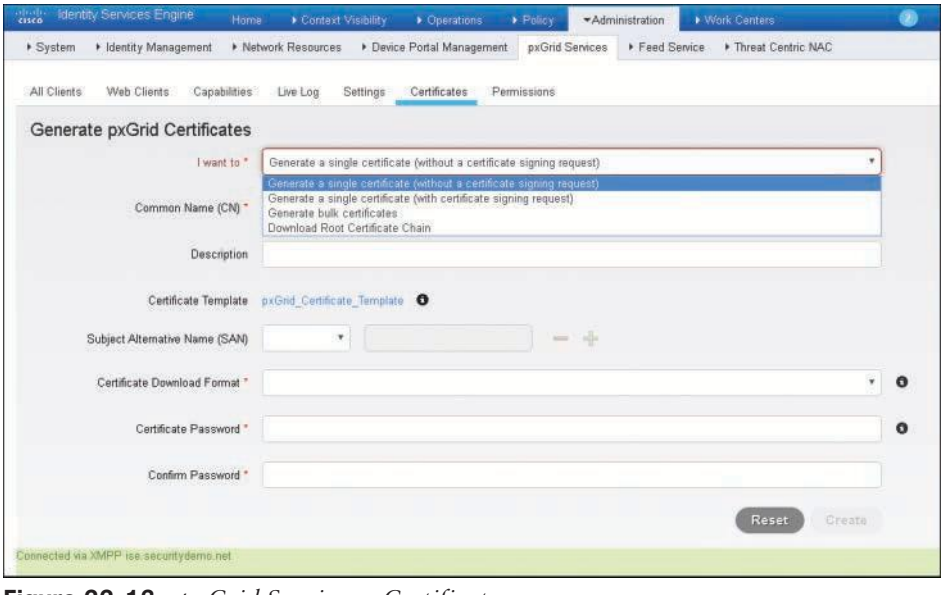


Figure 22-13 *pxGrid Services > Certificates*

As you can see in Figure 22-13, you can use the screen to generate a single certificate, sign a certificate signing request (CSR), generate bulk certificates from a CSV file, or download the certificate authority chain for import into the trust store of the pxGrid participant. For the FMC, you need to generate a certificate/key pair.

Step 2. Select **Generate a Single Certificate (Without a Certificate Signing Request)**.

Step 3. Enter a common name (CN) for the subject of your certificate. The CN is normally the FQDN of the host (for example, fmc.securitydemo.net). However, a common practice is to add a prefix to the CN, such as *pxGrid-*, which can help you avoid installation errors that sometimes occur when you try to install more than one certificate with the same FQDN.

Step 4. Add a SAN, if needed. If you use anything other than the true FQDN for the device, you need to use the Subject Alternative Name (SAN) field. Any time you use a SAN, it must also contain the CN. Add an entry for the FQDN of the host. Adding a SAN for the IP address is helpful in the event that one of the pxGrid peers is sent to the host via the IP address instead of the FQDN.

Step 5. Choose **Certificate in Privacy Enhanced Electronic Mail (PEM) Format, Key in PKCS8 PEM Format**. All options include the internal CA's certificates for the entire PKI hierarchy. There is also an option to download it as a PKCS12 chain file, where the public certificate, private key, and signing chain are all in a single file. For the FMC, however, the download format needs to be separate PEM files, not the PKCS12 chain.

Step 6. Add a password for the private key. (ISE never issues private keys without a password to encrypt the key.)

Step 7. Click **Create** to download the resulting .zip file.

Figure 22-14 shows the completed certificate form, and Figure 22-15 shows the contents of the .zip file.

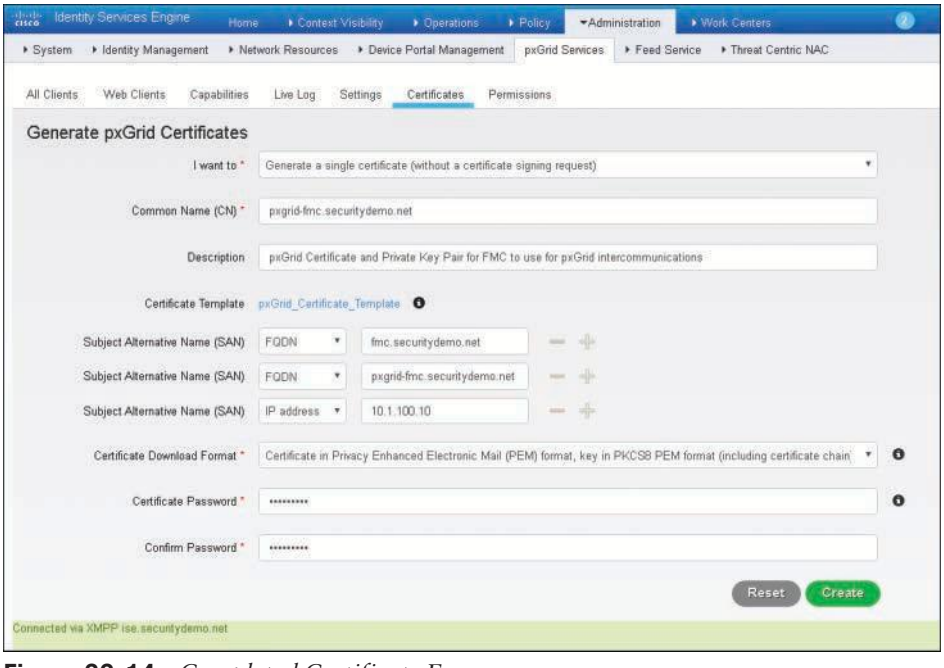


Figure 22-14 *Completed Certificate Form*

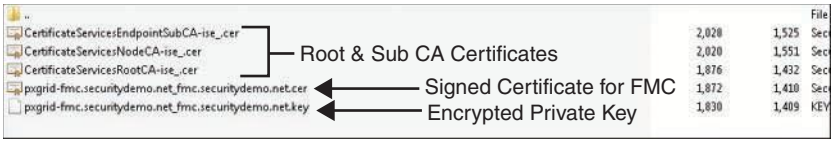


Figure 22-15 *Contents of the Resulting .zip File*

In Figure 22-15, you can see that the .zip file contains the signed certificate, the encrypted private key, and all the signing certificates in the PKI hierarchy for the issued certificate. In addition, the signing certificates in the PKI hierarchy for the Admin certificate are also included for good measure. Beginning with ISE Version 2.2, these certificates should not be required but are included in the .zip file anyway.

Now you have all the required certificates and the private key for the FMC.

To configure pxGrid in the FMC, follow these steps in the FMC GUI:

Step 1. Navigate to **System > Integration > Identity Sources** (see Figure 22-16).

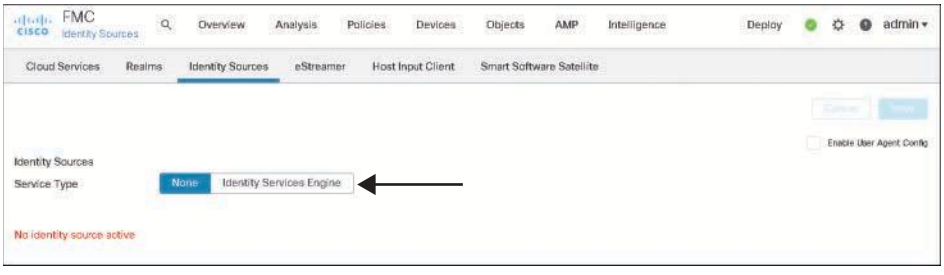


Figure 22-16 *FMC Identity Sources*

Step 2. Select **Identity Services Engine**.

Step 3. Enter the FQDN or IP address of the primary pxGrid controller.

Step 4. If there is a secondary controller, add its FQDN or IP address.

Step 5. Add the ISE root CA certificate by clicking the plus button for the pxGrid server CA. The root CA certificate is added to the list of trusted CAs in the FMC. Give the certificate a name such as ISErootCA, as shown in Figure 22-17.

Step 6. Click **Browse** and select the root CA certificate from the expanded .zip file you downloaded earlier.

Step 7. Click **Save**.

Step 8. Ensure that the newly imported root CA certificate is selected for both the pxGrid server CA and the MnT server CA, as shown in Figure 22-19.

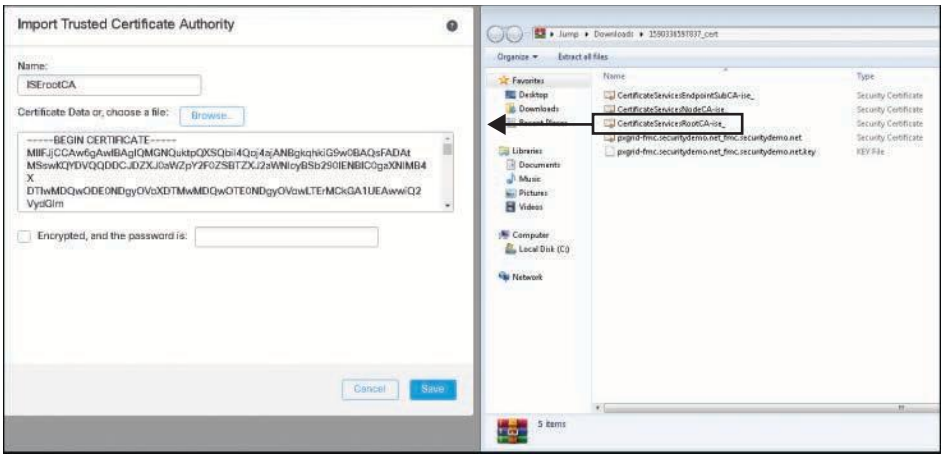


Figure 22-17 *Importing a Trusted Certificate Authority: ISE Root CA*

NOTE The separate MnT certificate is there just in case you are not using a single CA for all pxGrid clients. However, you now know that you should always use the same CA for all participants.

Step 9. Add the signed certificate and private key for the FMC by clicking the + button for the FMC server certificate. The PEM-encoded certificate that was signed by ISE's endpoint CA and the encrypted private key are now added to FMC. Give the internal certificate a name such as **pxGrid-FMC**, as shown in Figure 22-18.

Step 10. Click **Browse** next to Certificate Data and select the PEM certificate from the expanded .zip file you downloaded earlier (see Figure 22-18).

Step 11. Click **Browse** next to Key and select the PKCS8 key file from the expanded .zip file you downloaded earlier (see Figure 22-18).

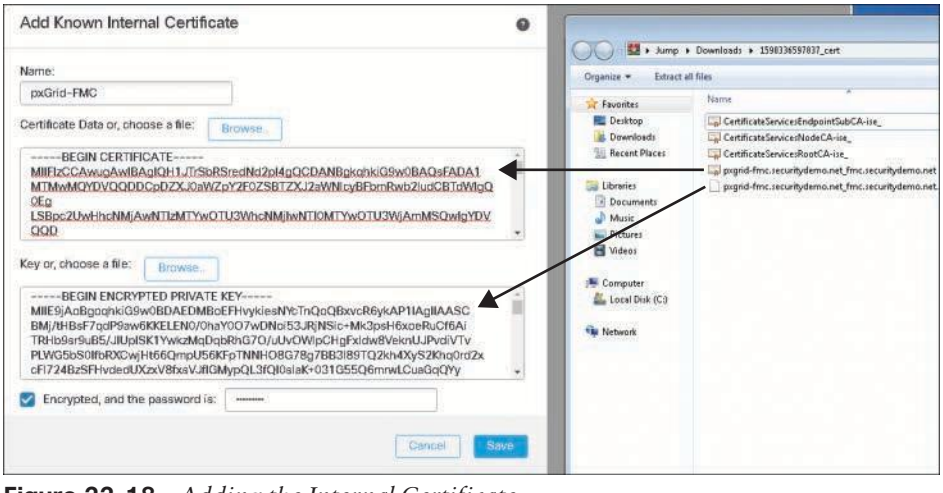


Figure 22-18 *Adding the Internal Certificate*

Step 12. To subscribe to the session topic, check the box next to **Session Directory Topic**, as shown in Figure 22-19.

Step 13. To subscribe to the SXP topic, check the box next to **SXP Topic**.

Step 14. Click **Save** in the upper-right corner of the screen. Figure 22-19 shows the completed form.

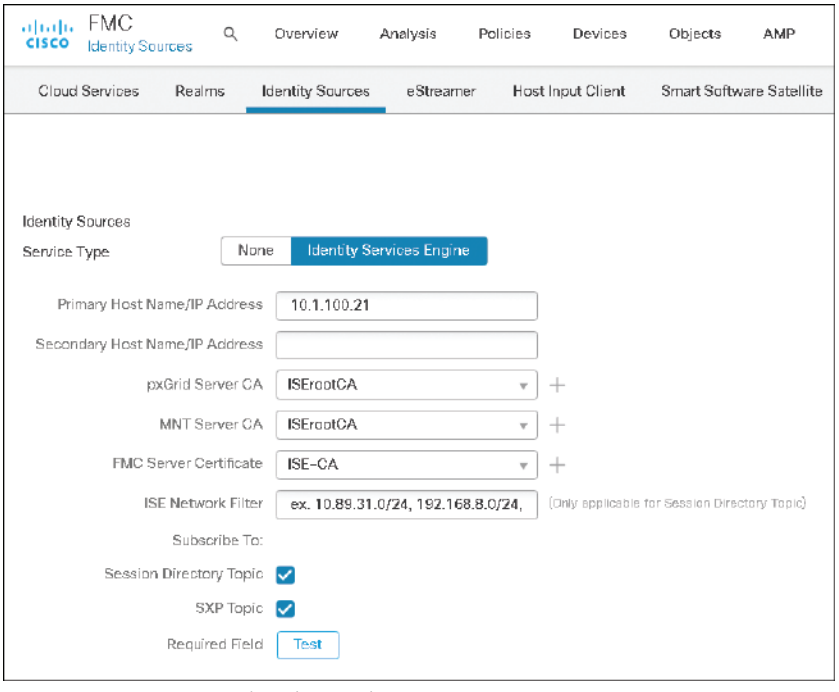


Figure 22-19 *Completed ISE Identity Source Form*

Step 15. Click **Test** to verify that the connection is successful. The test is likely to fail the first time you try unless ISE is configured to automatically approve new participants.

Step 16. In the ISE UI, navigate to **Administration > pxGrid Services > Clients**. If ISE is not configured to auto-approve participants, you need to accept the FMC's agent and test again.

Step 17. Select the **ia-fmc-** client for the FMC, as shown in Figure 22-20, and click **Approve**.

Step 18. Select the **st-fmc-** client and click **Approve**.

Step 19. Return to the FMC UI and attempt the test again. This test should be successful.

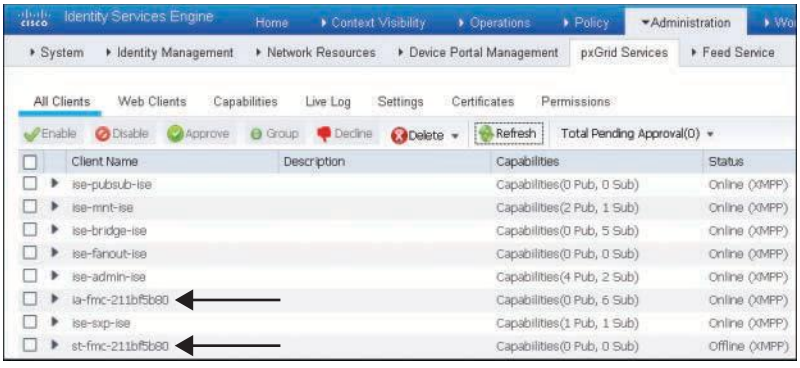


Figure 22-20 *pxGrid Clients*

Manually approving each and every pxGrid participant and test account can be time-consuming and somewhat confusing. Alternatively, you can enable automatic approval of certificate-based accounts in the pxGrid Settings screen, as shown in Figure 22-21. Just remember to disable it again after you are finished.

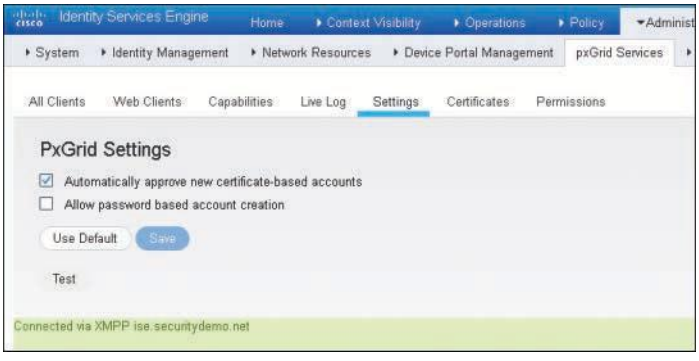


Figure 22-21 *pxGrid Settings*

NOTE In the pxGrid Settings screen, you have an option to allow password-based account creation. This is an alternative to the certificate-based accounts that you are seeing in this chapter. With password-based account creation, a password is leveraged instead and then tokens are assigned for authorization. At press time, there were not any pxGrid client applications leveraging this account method. Also, in the Settings screen you can click the Test button to verify that pxGrid is working as expected in ISE. This is very useful for checking that ISE trusts its own certificates.

Configuring Realms for Identity in Access Rules

The FMC may download all the users and IP address bindings, but none of the data that is downloaded is used in the policy until there is a realm configured to determine which groups and users to use in the firewall policies.

Realms leverage LDAP or LDAPS to communicate to query the data from Active Directory.

To configure realms in the FMC, follow these steps:

Step 1. Navigate to **System > Integration > Realms**.

Step 2. Click **New Realm**.

Step 3. In the Add New Realm dialog that appears (see Figure 22-22), provide a name for the new realm.

Step 4. For AD Primary Domain, enter the Active Directory domain name.

Step 5. Provide a UPN (user principal name) such as **administrator@securitydemo.net** for an AD user with enough permissions to join FMC to the domain.

Step 6. Enter the password for the AD user.

Step 7. Provide a UPN such as **administrator@securitydemo.net** for an AD user account to perform the LDAP queries.

Step 8. Enter a base DN (distinguished name) such as **ou=users,dc=securitydemo,dc=net** to begin the user account LDAP queries.

Step 9. Enter a base DN (distinguished name) such as **ou=groups,dc=securitydemo,dc=net** to begin the group LDAP queries.

HINT If you aren't getting the result you want, you can try backing up in the DN an extra level, such as **dc=securitydemo,dc-net**, so you can examine all organizational units (OUs).

Step 10. Click **OK**.

Figure 22-22 shows the completed Add New Realm dialog.

After the realm has been created, you are in the Realm Directory configuration. You need to add a directory, which is another way of saying that you need to add an LDAP server to perform the queries against.

To configure the directory, follow these steps:

Step 1. In the Realm Directory configuration screen, click **Add Directory**.

Step 2. Enter the IP address for the AD domain controller that FMC should use for LDAP queries.

Step 3. Enter the port for LDAP; **389** is the default port for unencrypted LDAP.

Step 4. If you are using secure LDAP, choose the method and certificate to trust.

Step 5. Click **OK**.

Figure 22-23 shows the completed Add Directory dialog.

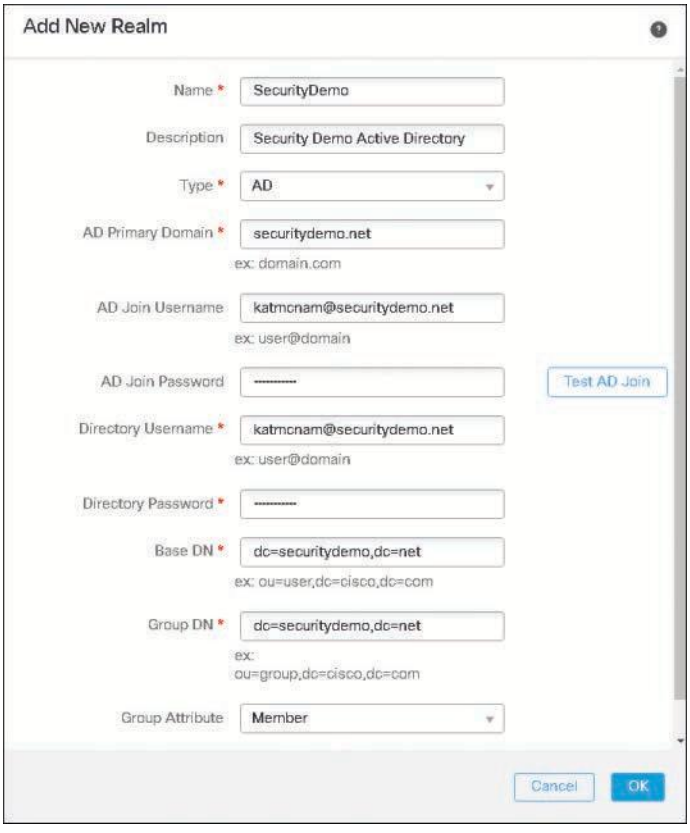


Figure 22-22 *Completed Add New Realm Form*

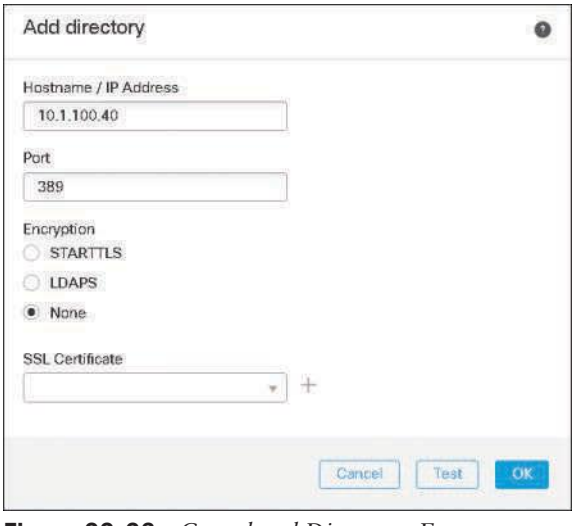


Figure 22-23 *Completed Directory Entry*

Now that the realm is configured, along with an LDAP server, it is time to download users and groups for use in the policies. Follow these steps:

Step 1. Click on the **User Download** tab.

Step 2. Enable the **Download Users and Groups** checkbox.

Step 3. Select the interesting groups from the Available Groups list and assign them for inclusion or exclusion from use in Firepower policies, as shown in Figure 22-24.

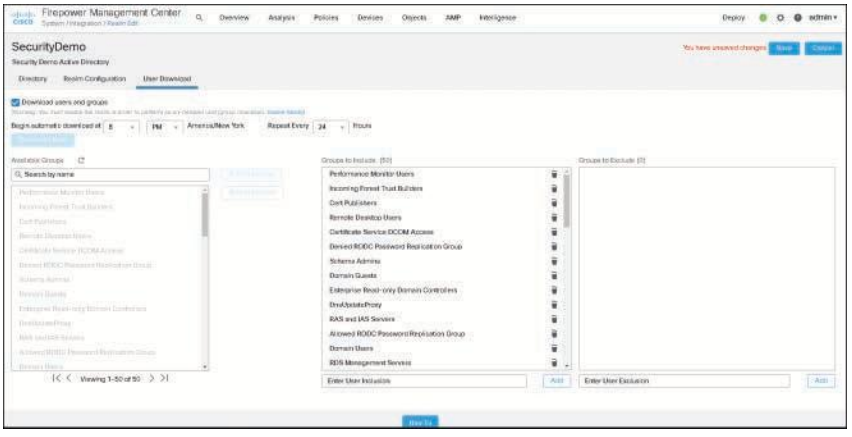


Figure 22-24 *Included and Excluded Groups*

NOTE Selective inclusion of AD groups is important for performance, as AD may have thousands of groups, most of which are not relevant for identity policies in the firewalls. In addition, the FMC would not perform very well if all groups were candidates for identity rules.

Step 4. Click **Save**.

Step 5. Enable the realm, as shown in Figure 22-25.



Figure 22-25 *Enabled Realm*

The realm is now fully configured for rule creation along with pxGrid integration for learning what IP addresses belong to which users and devices. Now you are ready to add identity information to the access policy rules in the FMC.

Configuring Firepower Access Rules with Context from pxGrid

Before you can add user identities or groups to the access policy rule, you must first create an identity rule by following these steps:

Step 1. Navigate to **Policies > Access Control > Identity**.

Step 2. Click **New Policy**.

Step 3. Enter a name, such as **ID_Policy** and, optionally, a description.

Step 4. Click **Save**.

Figure 22-26 shows the new identity policy being created.

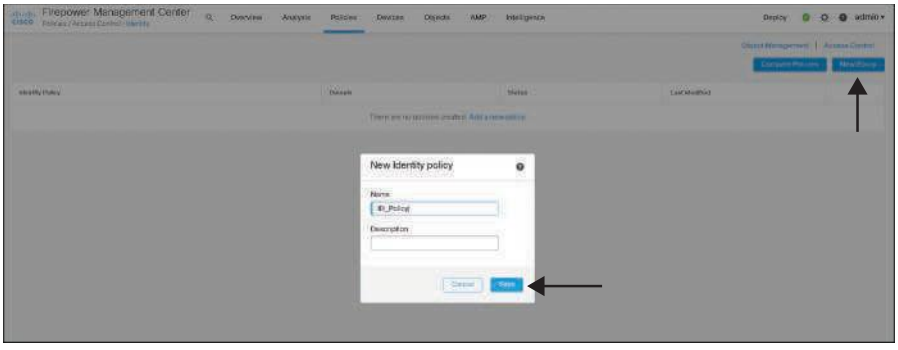


Figure 22-26 *New Identity Policy*

Step 5. Click **Add Rule** to configure an Identity rule.

Step 6. Enter a name, such as **ID-Rule**.

Step 7. Keep the **Enabled** checkbox checked.

Step 8. In the Action dropdown, select **Passive Authentication**.

Step 9. Click **Realm & Settings**.

Step 10. From the Realm dropdown, select your AD realm.

Step 11. Click **Add**.

Figure 22-27 shows the new rule being added to the identity policy.

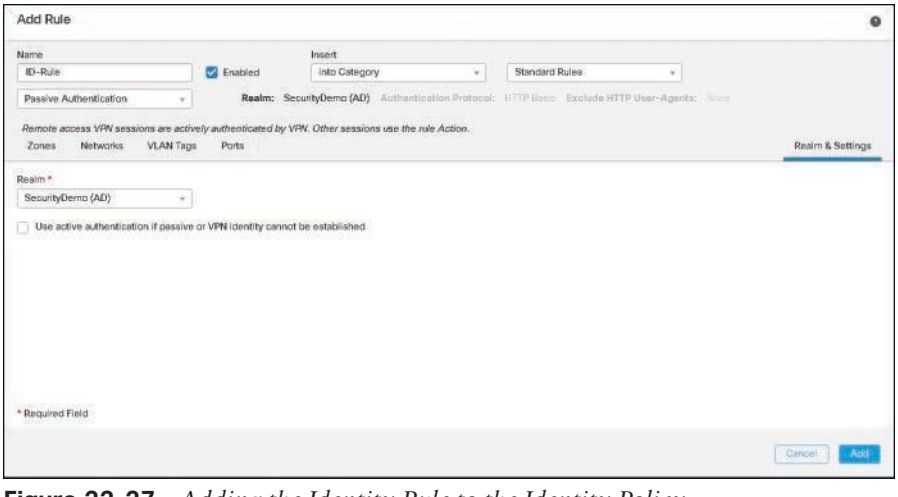


Figure 22-27 *Adding the Identity Rule to the Identity Policy*

Now that an identity policy has been created, it can be attached to the access policy.

Step 12. Navigate to **Policies > Access Policy > Access Policy**.

Step 13. Click on the **Identity Policy** link.

Step 14. Choose your identity policy from the dropdown.

Step 15. Click **OK**.

Figure 22-28 shows the identity policy being selected in the access policy.

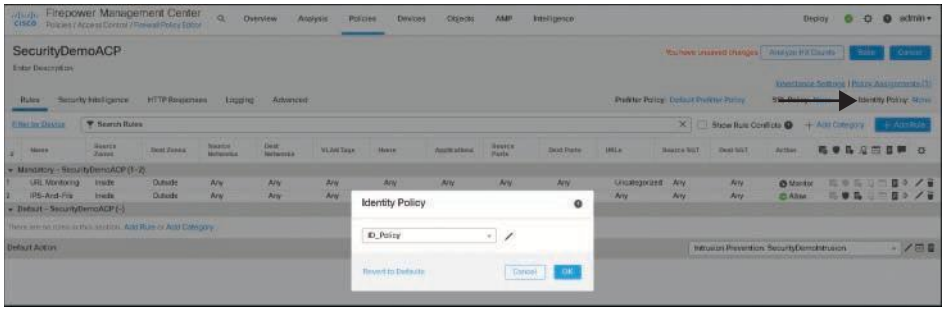


Figure 22-28 *Selecting the Identity Policy in the Access Policy*

Now that an identity policy has been attached to the access policy, identities can be added to the access rule.

Step 16. Navigate to **Policies Access Policy > Access Policy**.

Step 17. Either click **Add Policy** to create a new one or click **Edit** to edit an existing policy.

Step 18. Click the **Add Rule** button to create an access policy rule.

Step 19. Click the **Users** tab.

Step 20. Click on the realm you created (in the leftmost column).

Step 21. When users and groups populate in the middle column, select the groups or users to match in this access rule.

Step 22. Click **Add to Rule**.

Figure 22-29 shows the user group Employee being added to the access rule.

Since Firepower Management Center is integrated with ISE, you have access to other bits of contextual data to build a policy on, such as endpoint profiles and TrustSec tags (also known as Scalable Group Tags or Security Group Tags [SGTs]).

Step 23. Click on the **SGT/ISE Attributes** tab.

Step 24. In the leftmost column, select **Security Group Tag** from the Available Metadata dropdown.

Step 25. Select one of the SGTs from ISE and click **Add to Rule**.

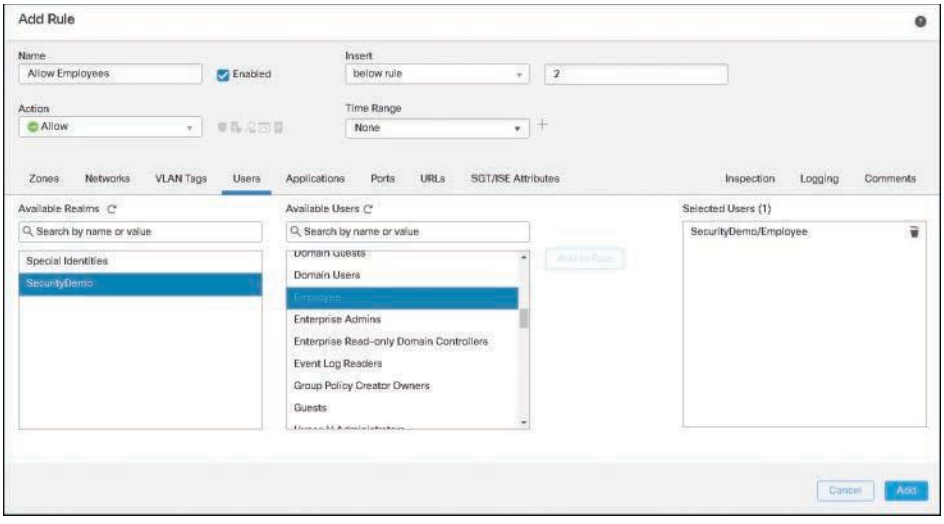


Figure 22-29 *Adding AD Groups to an Access Policy Rule*

Prior to FMC Version 6.5, FMC access control rules could only be configured with source metadata from ISE. This changed due to an enhancement in FMC Version 6.5 that allows the FMC to subscribe to the SXP topic to learn IP address-to-SGT mappings.

Figure 22-30 shows the source SGT named Employees being added to the access rule.

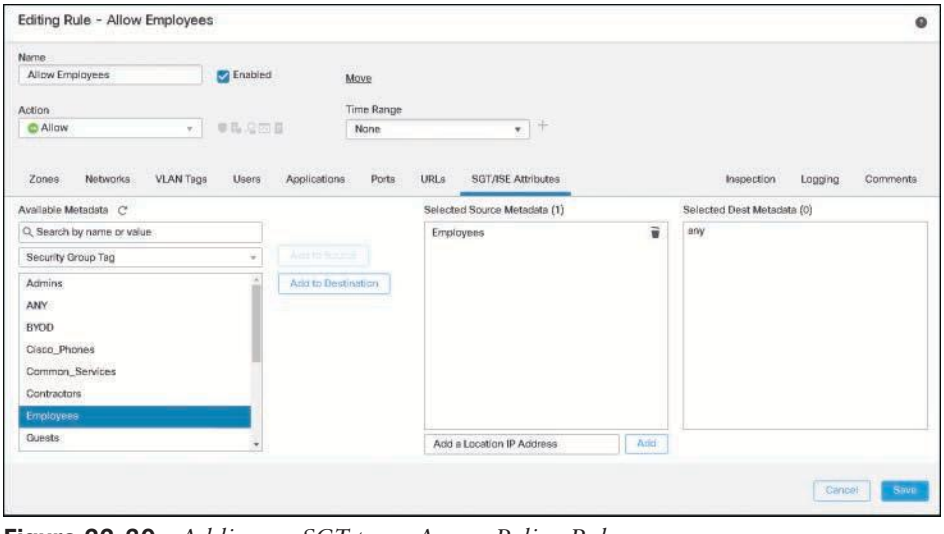


Figure 22-30 *Adding an SGT to an Access Policy Rule*

Step 26. In the leftmost column, select **Device Type** from the Available Metadata dropdown.

Step 27. Choose the appropriate endpoint profiles and add them to the policy.

Step 28. Click **Add** to save the Access Policy rule to the policy.

Step 29. Click **Save** to save the policy.

Figure 22-31 shows the Active Directory user group Employees being added to an access rule that allows access to a specific network segment.

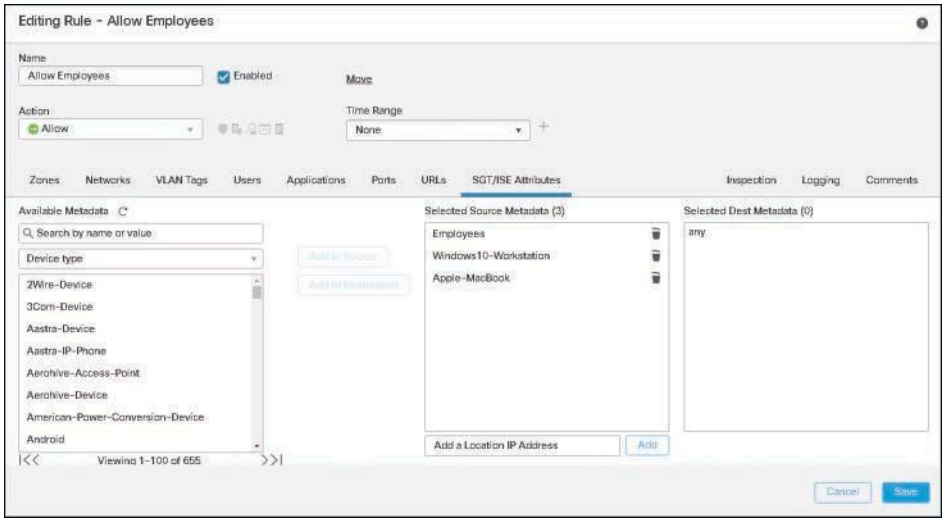


Figure 22-31 *Adding Endpoint Profiles to an Access Policy Rule*

Viewing Active Users

You've completed all the configuration steps for the identity integration with FMC and ISE, but how do you know that the FMC is learning about the active and passive online users and devices?

Navigate to Analysis > Users > Active Sessions, and you should start seeing domain logons, such as those shown in Figure 22-32. This figure shows the online users that ISE has learned about through either active or passive identity mappings.

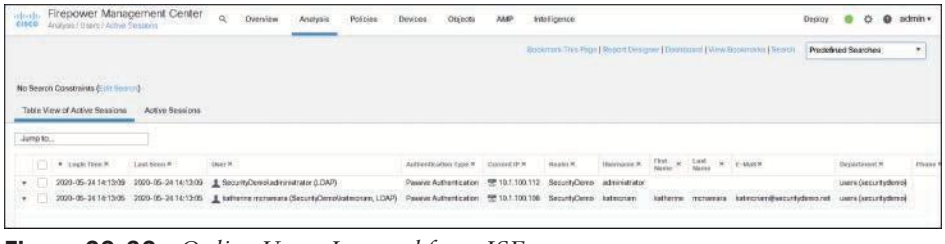


Figure 22-32 *Online Users Learned from ISE*

For those who enjoy using the CLI, there is also a great way to see the user identities from the expert mode command line in the FMC: **adi_cli session** (see Example 22-1).

Example 22-1 *Viewing Online Users from the FMC CLI*

```
admin@atw-fmc:~$ sudo adi_cli session | more
input 'q' to quit
received realm information: operation REALM_DELETE_ALL, Null realm info
received realm information: operation REALM_ADD, realm name securitydemo.net, sh
ort name SECURITYDEMO, id 3
ADI is connected
received security group operation: DELETE ALL
received security group operation: ADD id: 92bb1950-8c01-11e6-996c-525400b48521
name: ANY fullyQualifiedName: Any Security Group tag: 65535
received security group operation: ADD id: 934557f0-8c01-11e6-996c-525400b48521
name: Auditors fullyQualifiedName: Auditor Security Group tag: 9
received security group operation: ADD id: 935d4cc0-8c01-11e6-996c-525400b48521
name: BYOD fullyQualifiedName: BYOD Security Group tag: 15
received security group operation: ADD id: 9370d4c0-8c01-11e6-996c-525400b48521
name: Contractors fullyQualifiedName: Contractor Security Group tag: 5
received security group operation: ADD id: 93837260-8c01-11e6-996c-525400b48521
name: Developers fullyQualifiedName: Developer Security Group tag: 8
received security group operation: ADD id: 9396d350-8c01-11e6-996c-525400b48521
```

Configuring Rapid Threat Containment with Firepower and ISE

Learning about the online users and endpoints is only one of the use cases when integrating the FMC with ISE. Another common use case of the integration is to act when a malicious activity has occurred, as you learned earlier in this chapter.

Figure 22-33 illustrates how the FMC works with correlation rules and remediation modules to aid your understanding of how all the pieces fit together.

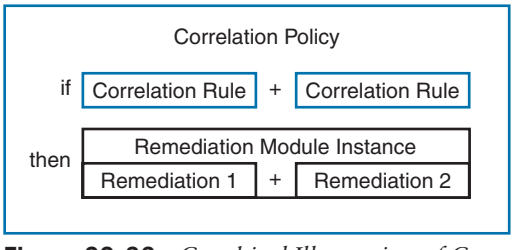


Figure 22-33 *Graphical Illustration of Correlation Policies and Components*

The parts that make up the response are as follows:

- **Correlation policy**: The policy construct that is made up of correlation rules and configured remediations.
- **Correlation rule**: An individual rule housed inside a correlation policy that is configured to look for one or more security events. There can be one or many correlation rules in each correlation policy.

- **Remediation module:** Modules of FMC that understand how to communicate to an external system. For example, the pxGrid module knows how to use EPS on ISE to quarantine endpoints.
- **Remediation instance:** A specific instance of a remediation module, as there can be many instances, each with a different configuration.
- **Remediation:** A specific action that is configured, such as quarantine. There can be many remediations in each instance of the remediation module.

The pxGrid mitigation module is built in to the FMC, and this module can be used to take an EPS quarantine action when a correlation rule is triggered.

Follow these steps in the FMC GUI to configure the built-in pxGrid mitigation module:

Step 1. Navigate to **Policies > Actions > Modules** (see Figure 22-34).



Figure 22-34 *Remediation Modules*

Step 2. Click the magnifying glass icon on the right-hand side for the **pxGrid Mitigation** module.

Step 3. Click **Add** to create a new instance of the module.

Step 4. Provide a name and, optionally, a description for the instance, as shown in Figure 22-35.

Step 5. Click **Create**.

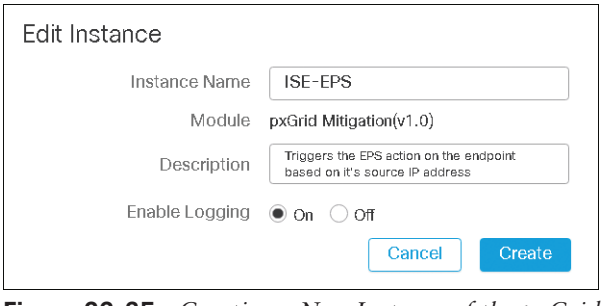


Figure 22-35 *Creating a New Instance of the pxGrid Mitigation Module*

Step 6. Choose **Mitigate Source** in the Configured Remediations dropdown.

Step 7. Click **Add**.

After clicking **Create**, you are brought to the window where you create a remediation action for the module (see Figure 22-36).

Step 8. Provide the remediation a name and, optionally, a description.

Step 9. Select the **quarantine** mitigation action.

Step 10. Click **Create**.

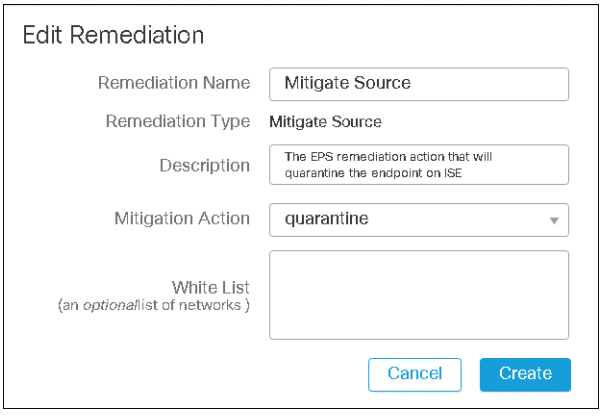


Figure 22-36 *Selecting Mitigate Source*

Step 11. Click **Save** (see Figure 22-37).

Step 12. Click **Done**.

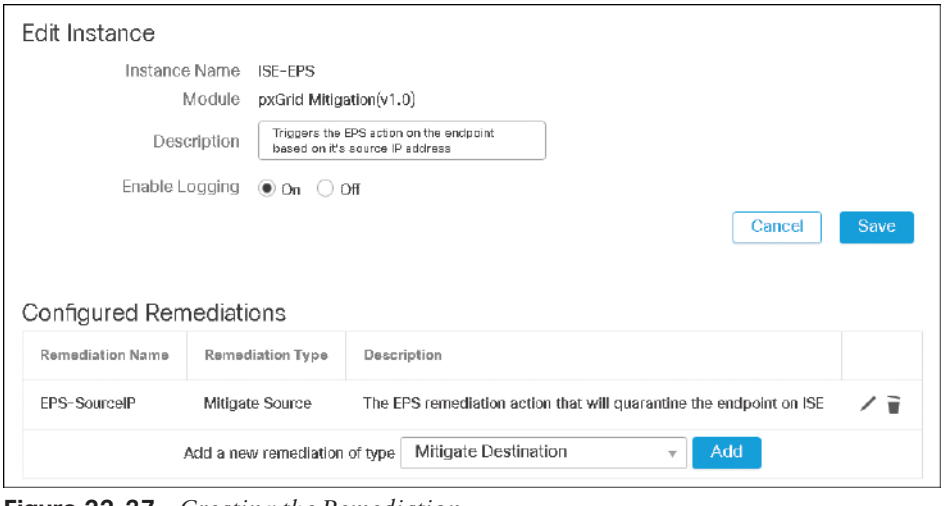


Figure 22-37 *Creating the Remediation*

Step 13. Click **Save** to save the module instance.

Figure 22-38 shows the completed instance of the pxGrid mitigation module.

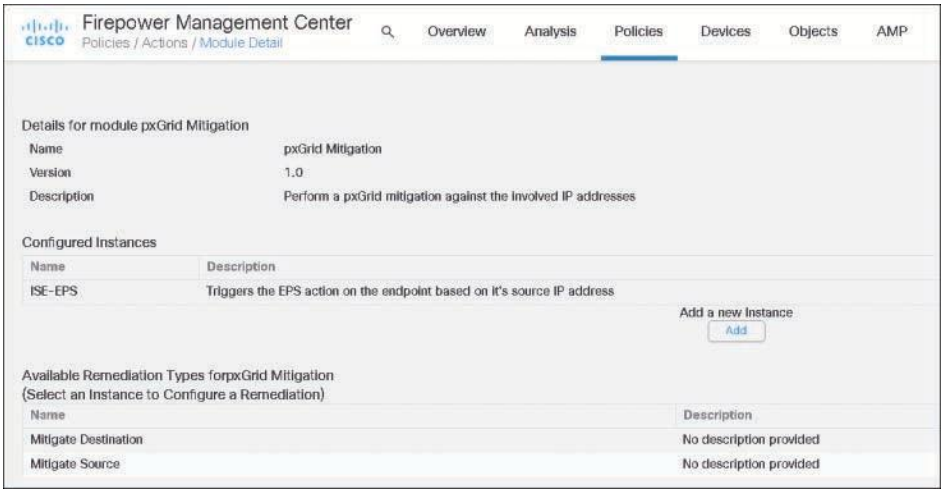


Figure 22-38 *Completed pxGrid Mitigation Module*

The remediation module is ready for use, so now you need to create a correlation rule that will use the remediation module whenever that rule is matched.

Step 14. Navigate to **Policies > Correlation > Rule Management**.

Step 15. Click **Create Rule**.

Figure 22-39 shows a completed correlation rule that is looking for an AMP for Endpoints event where malware is unable to be quarantined. Now that this rule exists, you can add it to the correlation policy.

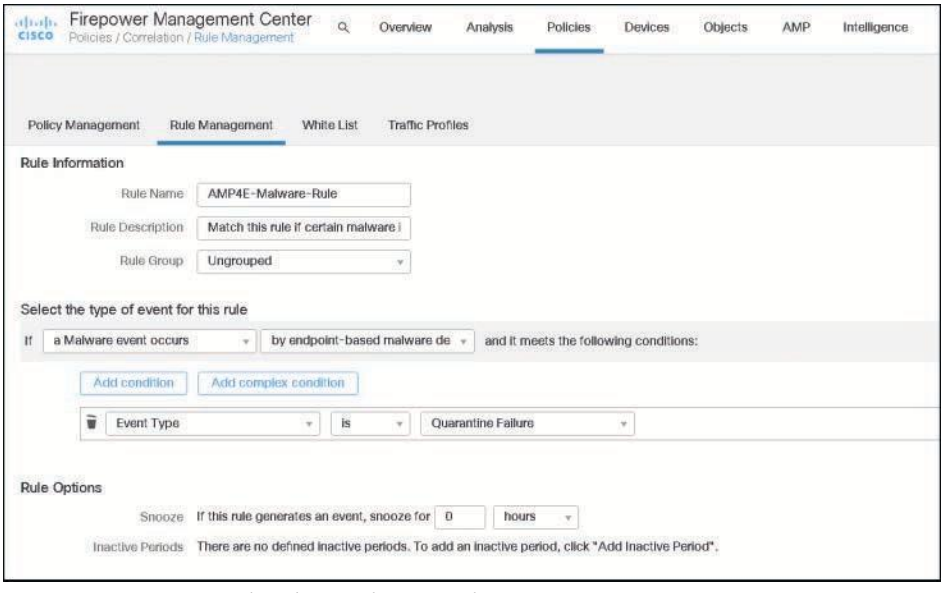


Figure 22-39 *Completed Correlation Rule*

Step 16. Navigate to the **Policy Management** tab.

Step 17. Click **Create Policy**.

Step 18. Provide a policy name and, optionally, a description.

Step 19. Click **Add Rules**.

Step 20. Select the correlation rule you created.

Figure 22-40 shows a correlation policy with the correlation rule added. However, there is no remediation action configured yet, and you can create one next.

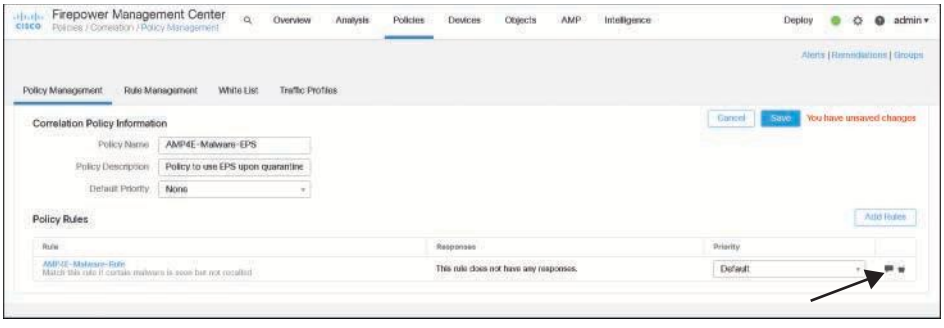


Figure 22-40 *Correlation Policy Without a Remediation Action*

Step 21. Click the response icon, which is highlighted in Figure 22-40.

Step 22. Assign the remediation action you created, as shown in Figure 22-41.

Step 23. Click **Update**.

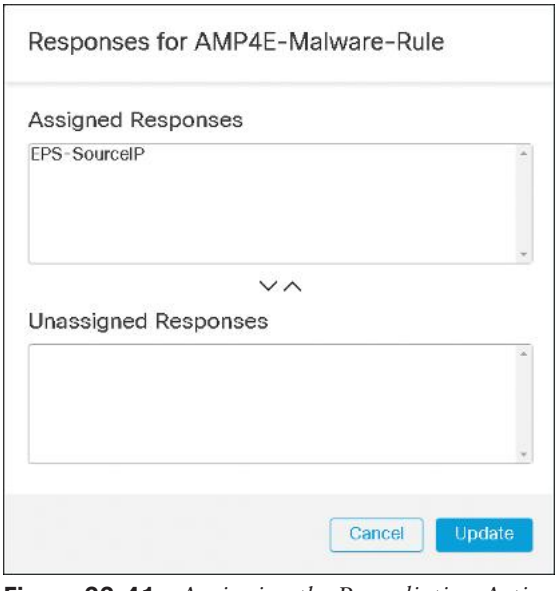


Figure 22-41 *Assigning the Remediation Action*

Step 24. Click **Save**.

Step 25. Enable the policy, as shown in Figure 22-42.



Figure 22-42 *Enabling the Final Policy*

Configuring the Web Security Appliance

The Cisco Web Security Appliance (WSA) was one of the first pxGrid partner applications in the security ecosystem. The WSA can use pxGrid to ascertain both passive and active user identities as well as TrustSec tags. Prior to WSA Version 11.7, the WSA was unable to combine Active Directory group membership with the identity information gathered from pxGrid, which meant TrustSec tagging was realistically the only scalable approach when using pxGrid in prior versions. Starting with WSA Version 11.7, the WSA now can transparently identify user-identity information, including usernames and Active Directory groups, by using pxGrid.

Integrating the WSA and ISE with pxGrid

With the WSA, you can create a certificate/key pair, just as you did for the FMC. Go ahead and create that certificate/key pair and then follow along with this section.

To configure pxGrid on WSA Version 11.7 or later, you need to enable ERS on ISE. Follow these steps in the ISE GUI:

Step 1. Navigate to **Administration > System > Admin Access > Administrators > Admin Users**.

Step 2. Click **Add > Create an Admin User** to create a new ERS admin user.

Step 3. Name this new admin account **ers-admin**, give it a password, and, optionally, give it a description.

Step 4. Add the admin account to the **ERS Admin** group and click **Submit** (see Figure 22-43).

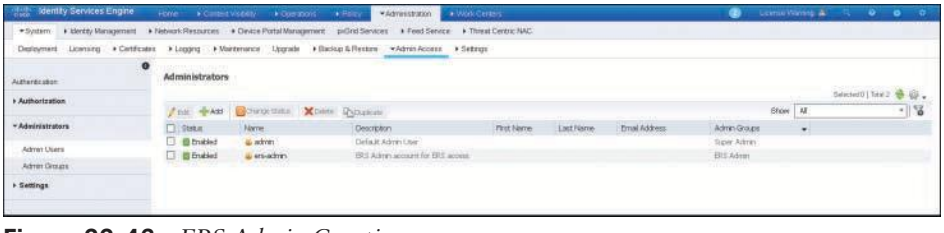


Figure 22-43 *ERS Admin Creation*

Step 5. Navigate to **Administration > System > Settings > ERS Settings**.

Step 6. Choose the **Enable ERS for Read/Write** setting if the Primary Administration node will be the pxGrid node.

Step 7. Choose the **Enable ERS for Read** setting for all other nodes.

Step 8. Click **Save** (see Figure 22-44).

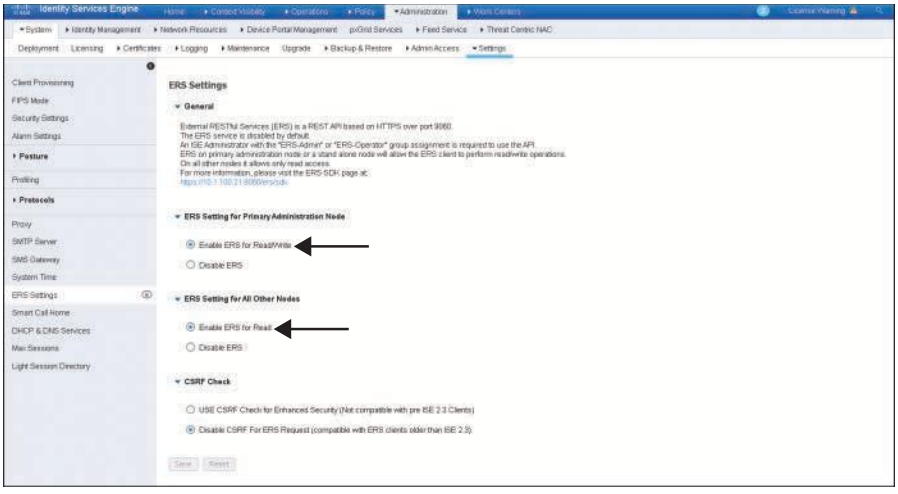


Figure 22-44 *Enabling ERS for Read*

To configure pxGrid on the WSA, you need to add the ISE root certificates to the trusted certificate store.

Step 9. In the WSA GUI, navigate to **Network > Certificate Management**.

Step 10. Click **Manage Trusted Root Certificates**, as shown in Figure 22-45.

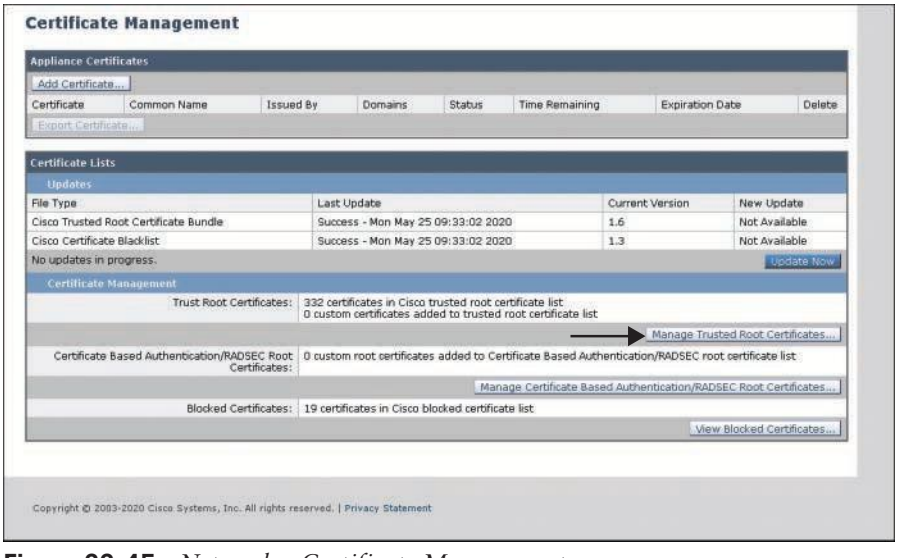


Figure 22-45 *Network > Certificate Management*

Step 11. Click **Import**.

Step 12. Browse for each of the ISE CA certificates (Root, Node, and Endpoint) and click **Submit** after selecting each one.

Step 13. When all the signing certificates are uploaded, as shown in Figure 22-46, click **Submit**.



Figure 22-46 *Manage Trusted Root Certificates*

Step 14. Click **Commit Changes** to save the WSA configuration.

Now that the ISE root certificates will be trusted, it is time to configure the WSA for pxGrid:

Step 15. Navigate to **Network > Identification Services > Identity Services Engine**.

Step 16. Click **Enable and Edit Settings**, as shown in Figure 22-47.

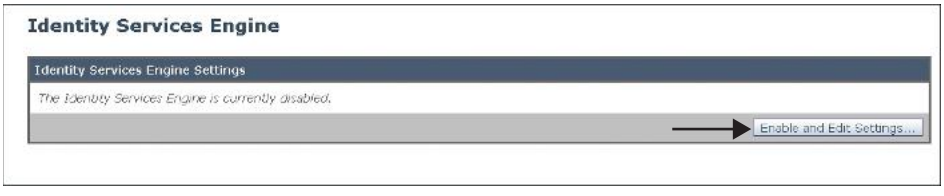


Figure 22-47 *Identification Servers > Identity Services Engine*

Step 17. In the Primary ISE pxGrid Node section, enter the FQDN for the primary pxGrid controller.

Step 18. Click **Choose File** and select the ISE root CA certificate.

Step 19. Click **Upload File**.

Figure 22-48 shows the completed Primary ISE pxGrid Node section.



Figure 22-48 *Primary ISE pxGrid Node*

Step 20. In the WSA Client Certificate section, select **Use Uploaded Certificate and Key**.

Step 21. Click **Choose File** next to Certificate and select the WSA's certificate from the ISE CA.

Step 22. Click **Choose File** next to Key and select the WSA's private key from the ISE CA.

Step 23. Enable the **Key Is Encrypted** checkbox.

Step 24. Enter the password that you used to encrypt the key.

Step 25. Click **Upload Files**.

Figure 22-49 shows the WSA certificate and key selected and ready for upload.



Figure 22-49 *WSA Client Certificate Section*

Next, you need to enable Active Directory group support via the ISE External Restful Service (ERS).

Step 26. To enable ERS, check the box next to **Enable ISE External Restful Service (ERS)**. In WSA Version 11.7 and later, the WSA can retrieve Active Directory groups and local ISE groups by using the ISE ERS. This feature provides the ability to configure the WSA's policies using AD groups or in combination with SGTs.

Step 27. Add your **ers-admin** credentials in the ERS Administrator Credentials section.

Step 28. If your ERS servers are the same ISE nodes as your ISE pxGrid node, check the **Server Name Same as ISE pxGrid Node** box. Otherwise, manually enter the hostname or IPv4 addresses for the primary and secondary servers.

Figure 22-50 shows the completed ISE ERS settings section.

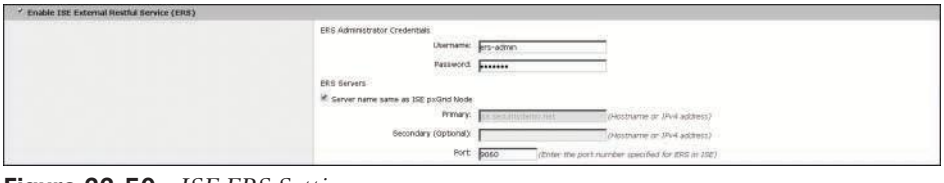


Figure 22-50 *ISE ERS Settings*

Step 29. Click **Submit** to complete the configuration.

Step 30. Click **Commit Changes** twice.

Step 31. To test the connection, click **Edit Settings**.

Step 32. Click **Start Test** at the bottom of the screen, as shown in Figure 22-51. If auto approval is enabled, the test should be successful. If it is not enabled, the test fails until you manually approve the two WSA accounts on the pxGrid controller.

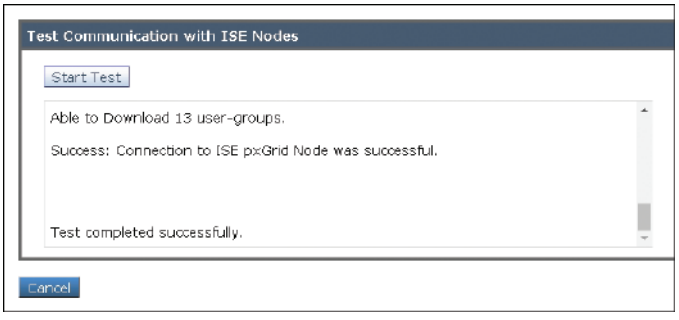


Figure 22-51 *Testing Communication with ISE Nodes*

Example 22-2 shows an example of the test output.

Example 22-2 *Example of Output for Testing Communication with ISE Nodes*

```
Checking DNS resolution of ISE pxGrid Node hostname(s) ...
Success: Resolved 'ise.securitydemo.net' address: 10.1.100.21

Validating WSA client certificate ...
Success: Certificate validation successful

Validating ISE pxGrid Node certificate(s) ...
Success: Certificate validation successful

Checking connection to ISE pxGrid Node(s) ...
Trying primary PxGrid server...
Preparing TLS connection...

Completed TLS handshake with PxGrid successfully.

Trying download user-session from (https://ise.securitydemo.net:8910)...

Able to Download 2 user-sessions.

Trying download SGT from (https://ise.securitydemo.net:8910)...
```

```
Able to Download 16 SGTs.

Trying connecting to primary ERS service...

Trying download user-groups...

Able to Download 13 user-groups.

Success: Connection to ISE pxGrid Node was successful.




Test completed successfully.
```

Configuring WSA Policies That Leverage the Data from ISE

Now that you have configured the WSA to work with ISE and to subscribe to the interesting pxGrid topics, it is time to configure policies. The first step is to create an identification profile. To do so, follow these steps in the ISE GUI:

Step 1. Navigate to **Web Security Manager > Identification Profiles**.

Step 2. Click **Add Identification Profile**.

Step 3. Provide a name for the profile.

Step 4. Under User Identification Method, from the Identification and Authentication dropdown, select **Transparently Identify Users with ISE**.

Step 5. Click **Submit**.

Step 6. Click **Commit Changes** to save the WSA configuration.

Figure 22-52 shows the completed identification profile.

Next, you need to add an Access Policy that leverages SGTs or Active Directory groups.

Step 7. In the ISE GUI, navigate to **Web Security Manager > Access Policies**.

Step 8. Click **Add Policy**.

Step 9. Provide a name for the policy, as shown in Figure 22-53.

Step 10. In the Identification Profiles and Users section, choose **Select One or More Identification Profiles** from the dropdown.

Step 11. Choose the configured ID profile from the dropdown.

Step 12. Enable the **Selected Groups and Users** radio button.

Step 13. If you would like to create the policy based on SGTs, select SGTs in the **ISE Secure Group Tags** area.

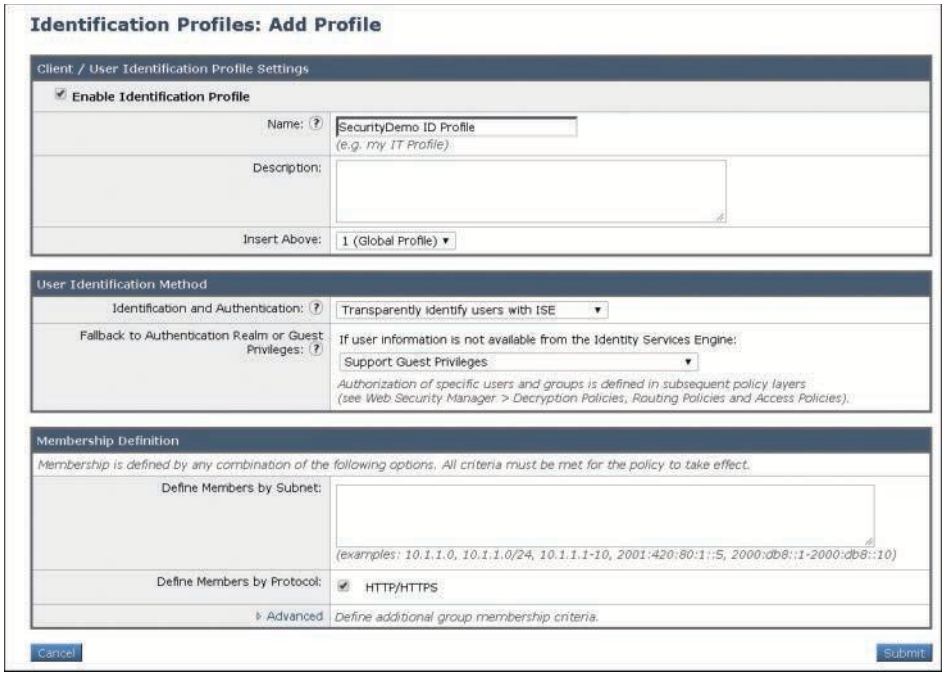


Figure 22-52 *Identification Profile*

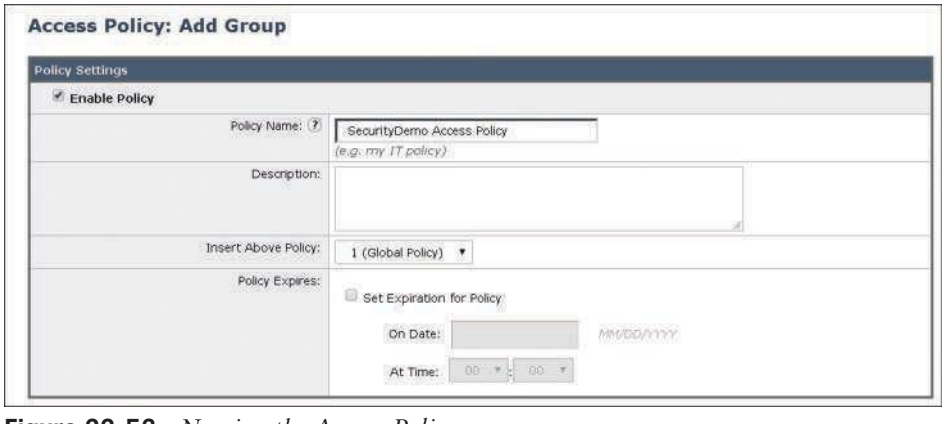


Figure 22-53 *Naming the Access Policy*

Step 14. If you would like to create the policy based on Active Directory groups learned from ISE, select **ISE Groups** in the ISE Groups area.

Step 15. Click **Submit**.

Figure 22-54 shows the completed access policy that will apply to all users with the Employee SGT assigned and the Employee Active Directory Group.



Figure 22-54 *Access Policy with Employee SGT and Employee AD Group*

Next, you need to add a decryption policy that will decrypt SSL traffic from users with a specific SGT.

Step 16. Navigate to **Web Security Manager > Decryption Policies**.

Step 17. Click **Add Policy**.

Step 18. Provide a name for the policy.

Step 19. In the Identification Profiles and Users section, choose **Select One or More Identification Profiles** from the dropdown.

Step 20. Choose the configured ID profile from the dropdown.

Step 21. Enable the **Selected Groups and Users** radio button.

Step 22. Select SGTs in the **ISE Secure Group Tags** area.

Step 23. Select Active Directory groups in the **ISE Groups** area.

Step 24. Click **Submit**.

Figure 22-55 shows the completed access policy that will apply to all users with the Employee SGT assigned and the Employee Active Directory group.



Figure 22-55 *Decryption Policy*

Integrating Stealthwatch and ISE

For years, Cisco had a proven solution known as Cyber Threat Defense (CTD), whose main components were Cisco ISE and a product called StealthWatch from Lancope. Cisco acquired Lancope in 2016 and proceeded to make the giant leap of rebranding the product from Lancope StealthWatch to Cisco Stealthwatch.

Regardless of what the product is called, Stealthwatch is phenomenal at security analytics and visibility. It works primarily by analyzing NetFlow records from the network and providing analytics related to the traffic, hosts, and other telemetry used to decorate the flows.

Why Integrate Stealthwatch and ISE?

Flow analysis itself is incredibly useful for analytics before and after an attack. Figure 22-56 shows a basic host report for a client PC in Stealthwatch before it is integrated with ISE. This report provides just a small taste of what Stealthwatch is able to provide to your security organization and security operation center (SOC) for incident response and alerting.

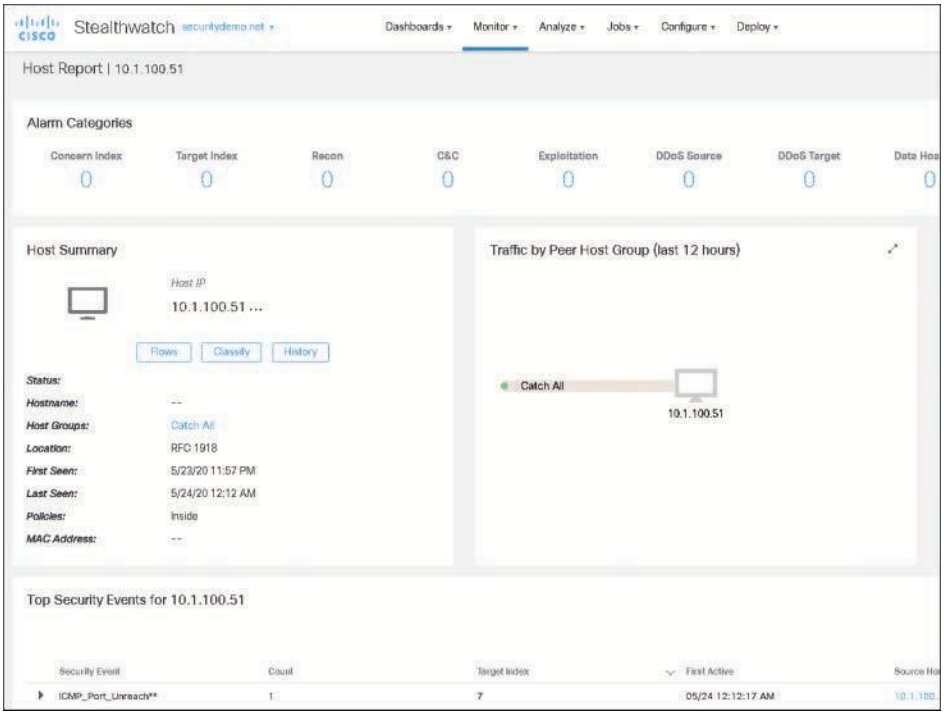


Figure 22-56 *Host Report: Before ISE Integration*

Beginning with Version 6.9, Cisco Stealthwatch uses ISE as the primary source for learning passive and active user identities to merge into the flow records used for security analytics. The mechanisms used are exactly the same whether it is full ISE or the ISE Passive Identity Connector (ISE-PIC), which provides only passive identities. In Version 7.0, Cisco Stealthwatch transitioned from pxGrid 1.0 (EPS) to pxGrid 2.0 (ANC), which meant that more than one custom policy could be created and used.

Just as with the WSA, the context provided to Stealthwatch can be much richer with full ISE, and it can therefore provide more value if you add an endpoint profile and TrustSec data.

After you integrate Stealthwatch and ISE, the flows contain much more context about the hosts, including the logged in user data. Figure 22-57 shows the populated Users & Sessions table after ISE integration.

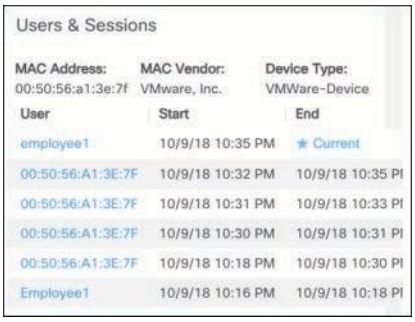


Figure 22-57 *Host Report: After ISE Integration*

Preparing Stealthwatch for pxGrid

To start configuring Stealthwatch for pxGrid, you can generate an additional TLS identity for the Stealthwatch Management Center (SMC)—that is, you will get a pxGrid certificate from ISE and install it on the SMC.

Unlike the FMC and the WSA, Stealthwatch uses the PKCS12 chain files instead of individual certificates. In other words, it requires the private key, a signed certificate, and all the signing root certificates in a single encrypted file.

NOTE All steps in this book are for Cisco Stealthwatch Version 7.0. To see the integration with Version 6.x, check out the Cisco Press book *Cisco ISE for BYOD and Secure Unified Access*.

Step 1. Click the settings cog in the upper-right corner and select **Central Management**, as shown in Figure 22-58. The Stealthwatch Central Management tab or window opens.

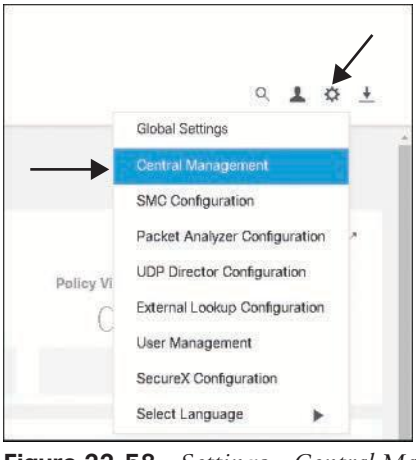


Figure 22-58 *Settings > Central Management*

Step 2. Click the actions bubble next to your Stealthwatch Management Center and click **Edit Appliance Configuration**, as shown in Figure 22-59.

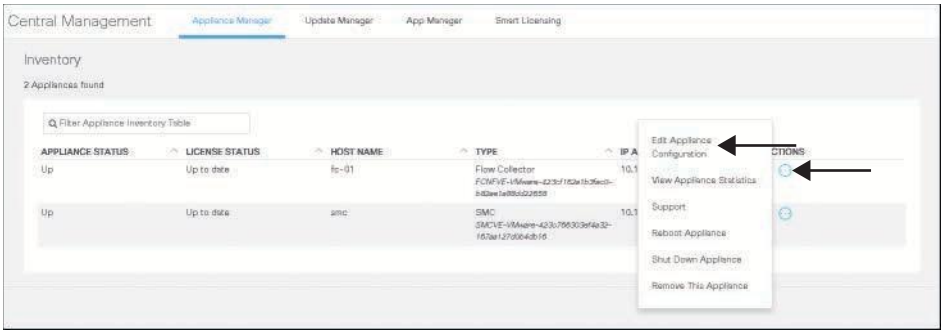


Figure 22-59 *Editing the Appliance Configuration*

Step 3. Scroll down to the section titled Additional SSL/TLS Client Identities and click **Add New**.

Step 4. Fill out the fields for the certificate signing request.

Step 5. Click **Generate CSR** again, as shown in Figure 22-60.

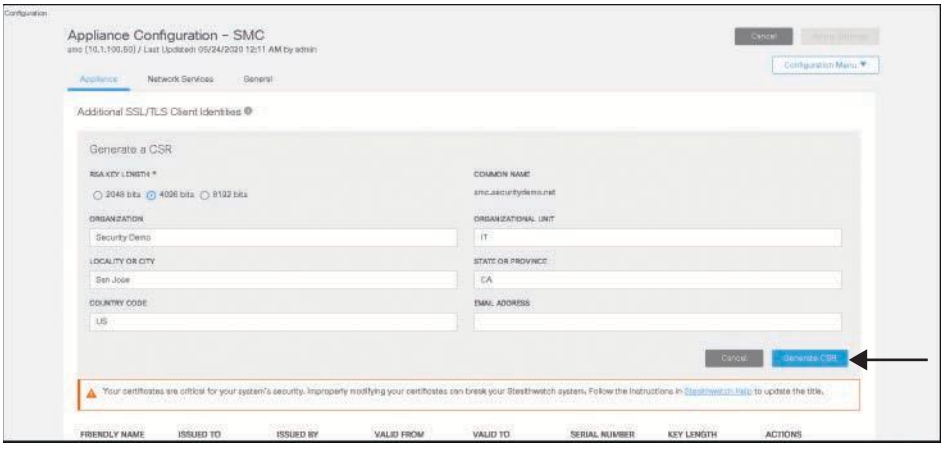


Figure 22-60 *Generating the CSR*

Step 6. Save the resulting .csr file to a location from which you can easily retrieve it.

Step 7. Open the CSR in your favorite text editor.

Step 8. Copy the contents of the CSR to your clipboard.

Step 9. In the ISE GUI, navigate to **Administration > pxGrid Services > Certificates**.

Step 10. Select **Generate a Single Certificate (with Certificate Signing Request)**.

Step 11. Choose the **PKCS12** format (including certificate chain—one file for both the certificate chain and key) from the dropdown.

Step 12. Enter and confirm a password for the encrypted file.

Step 13. Click **Create**.

Step 14. Save the resulting p12 file to a location from which you can easily retrieve it.

Figure 22-61 shows the completed certificate generation screen in ISE.

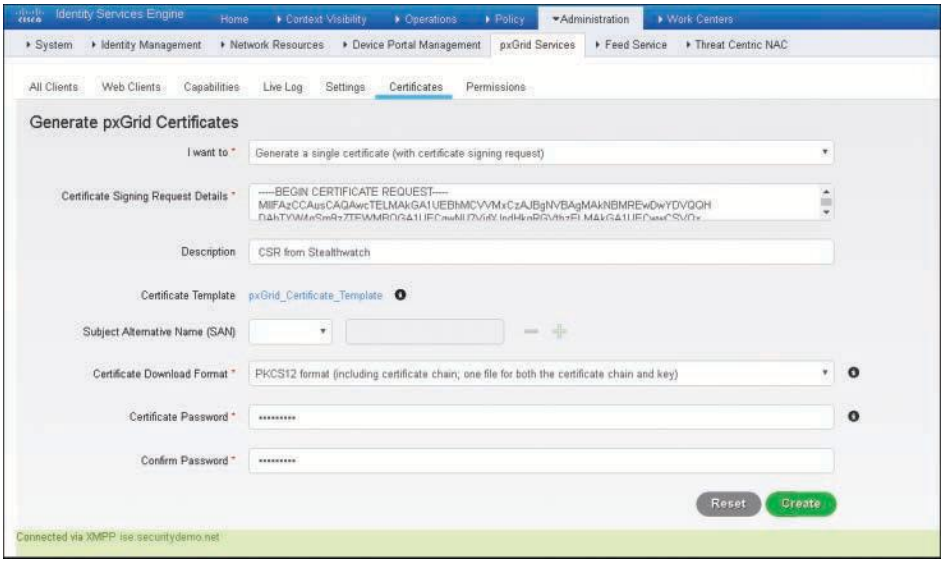


Figure 22-61 *Generating the Certificate Chain for Stealthwatch*

Step 15. In the Stealthwatch GUI, enter a friendly name for the identity certificate.

Step 16. Click **Choose File** and select the downloaded p12 chain file. After the UI recognizes the chain file, the bundle password field appears.

Step 17. Enter and confirm the bundle password.

Step 18. Click **Add Client Identity**.

Figure 22-62 shows the import of the PKCS certificate chain into Stealthwatch.

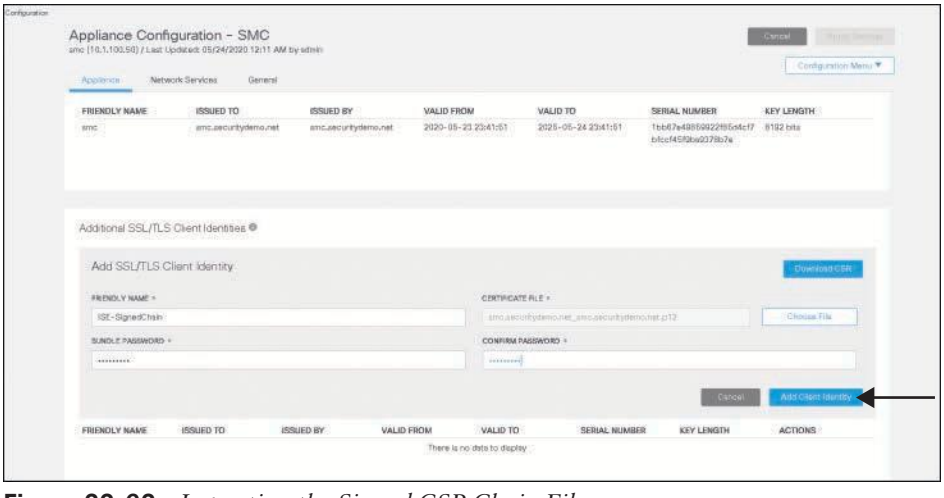


Figure 22-62 *Importing the Signed CSR Chain File*

Step 19. In order for the SMC to be able to use this new identity certificate, it needs to be added to the SMC's trusted CA store, so click on the **General** tab in the SMC's settings and scroll down to the **Trust Store** section (see Figure 22-63).

Step 20. Click **Add New**.

Step 21. Add the chain file to the Trust Store and click **Add Certificate**.

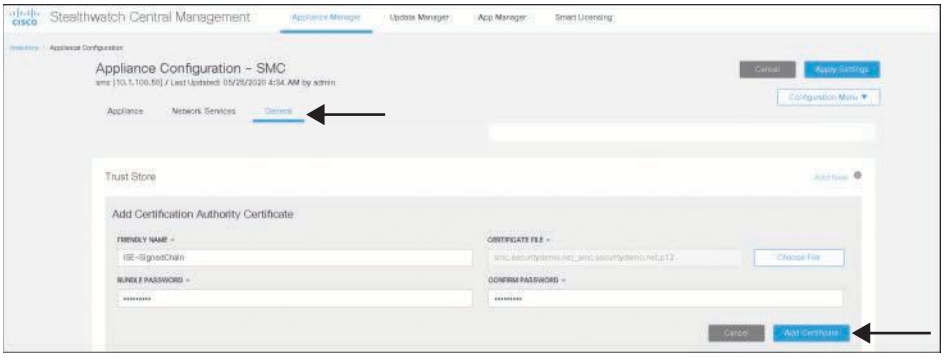


Figure 22-63 *Importing the Signed CSR Chain File into the Trust Store*

Step 22. Click **Apply Changes** to save the new identity certificate.

Configuring Stealthwatch for ISE

Now that the pxGrid client identity certificate has been imported into Stealthwatch, it is time to configure the ISE integration. In the Stealthwatch GUI, follow these steps:

Step 1. Navigate to **Deploy > Cisco ISE Configuration**, as shown in Figure 22-64.



Figure 22-64 *Deploy > Cisco ISE Configuration*

Step 2. Click **Add New Configuration**.

Step 3. Enter a friendly name for the ISE cube.

Step 4. Select the pxGrid certificate from the Certificate dropdown.

Step 5. Enter the IP addresses for the primary and secondary pxGrid controllers.

Step 6. Create a username to uniquely identify Stealthwatch in the ISE pxGrid UI.

Step 7. Check the box next to **Adaptive Network Control** to allow Stealthwatch to add or remove ANC policies.

Step 8. Check the box next to **Static SGT Classifications** to allow Stealthwatch to learn the static IP address-to-SGT bindings configured locally on ISE or learned via SXP.

Step 9. Check the box next to **Sessions** to allow Stealthwatch to learn of user session updates. Optionally, you can also check the **Track Sessions Derived from Machine Authentications** option if you want Stealthwatch to receive machine authentication updates and user updates.

Step 10. Click **Save**.

Figure 22-65 shows the completed ISE Configuration screen.

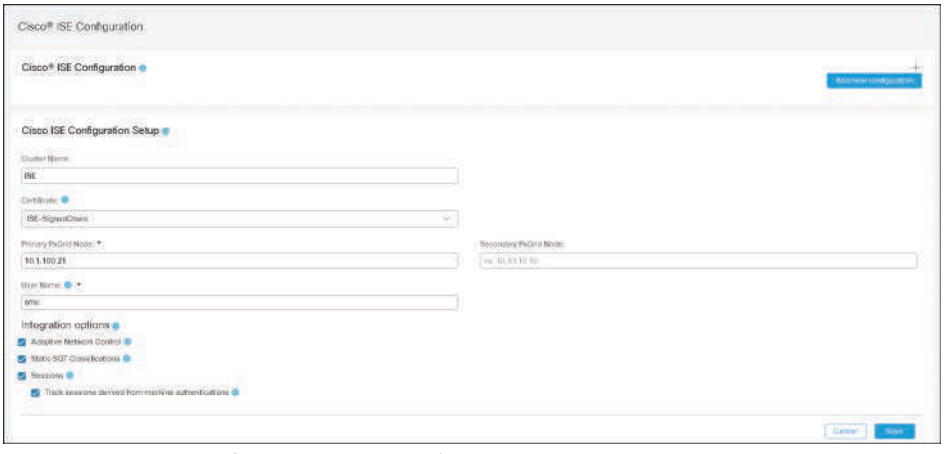


Figure 22-65 *Configuring the pxGrid Connection*

After a little bit of time, the status light for the pxGrid connection (highlighted in Figure 22-66) changes from yellow to green to indicate that the connection to pxGrid is up and running.



Figure 22-66 *Status Is Green*

Figure 22-67 shows the final pxGrid clients screen, where you can see the FMC, WSA, and the Stealthwatch clients in the list.

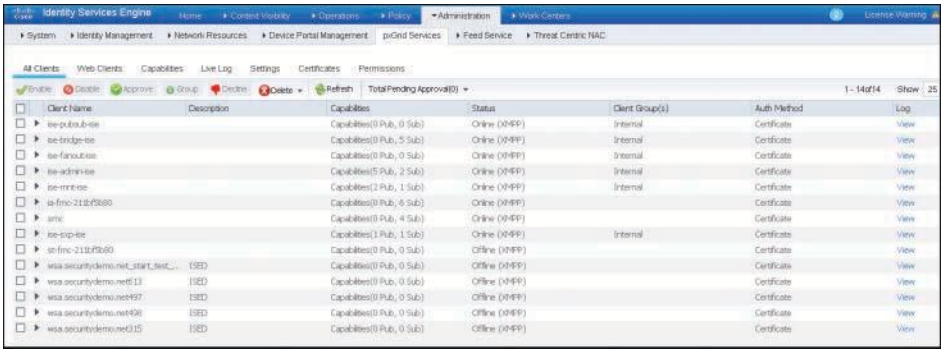


Figure 22-67 *Final pxGrid Clients Screen*

The Stealthwatch/ISE integration is not only for providing telemetry to Stealthwatch. You can also act during an investigation in Stealthwatch for enforcement through ISE. Stealthwatch Version 7.0 uses ANC, and previous versions use only EPS.

Unlike EPS, which had only two options (Quarantine and Unquarantine), ANC allows you to create many different labels of your choosing for a variety of actions. To create labels, follow these steps in the ISE GUI:

Step 1. Navigate to **Operations > Adaptive Network Control > Policy List**.

Step 2. Click **Add** to create a new label (called a policy).

Step 3. Give it a name, such as **Investigate**.

Step 4. Select the action to indicate the type of CoA ISE will issue: Shut Down, Port Bounce, or Re-Auth Quarantine.

Step 5. Click **Save**.

Figure 22-68 shows two configured ANC labels.

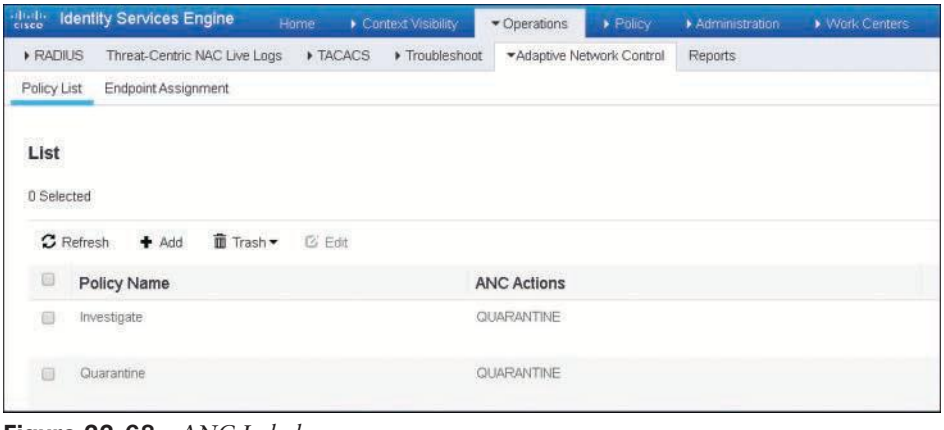


Figure 22-68 *ANC Labels*

Now that these labels exist, you can include them as conditions in your authorization rules, as shown in Figure 22-69.



Figure 22-69 *ANC Labels in the Authorization Policy*

Now when something looks awry during an incident response, you can assign an ANC label to a host right in the Stealthwatch interface and have ISE take action. To set this up, follow these steps in the Stealthwatch user interface:

Step 1. Click **Edit** next to ISE ANC Policy, as shown in Figure 22-70.

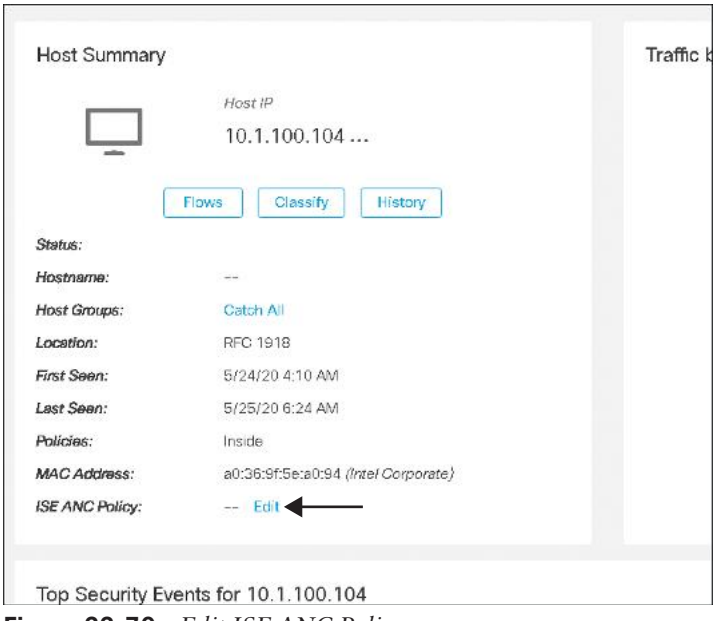


Figure 22-70 *Edit ISE ANC Policy*

Step 2. In the Applying ANC Policy screen, select the label you want from the drop-down, as shown in Figure 22-71.

Step 3. Click **Save**.

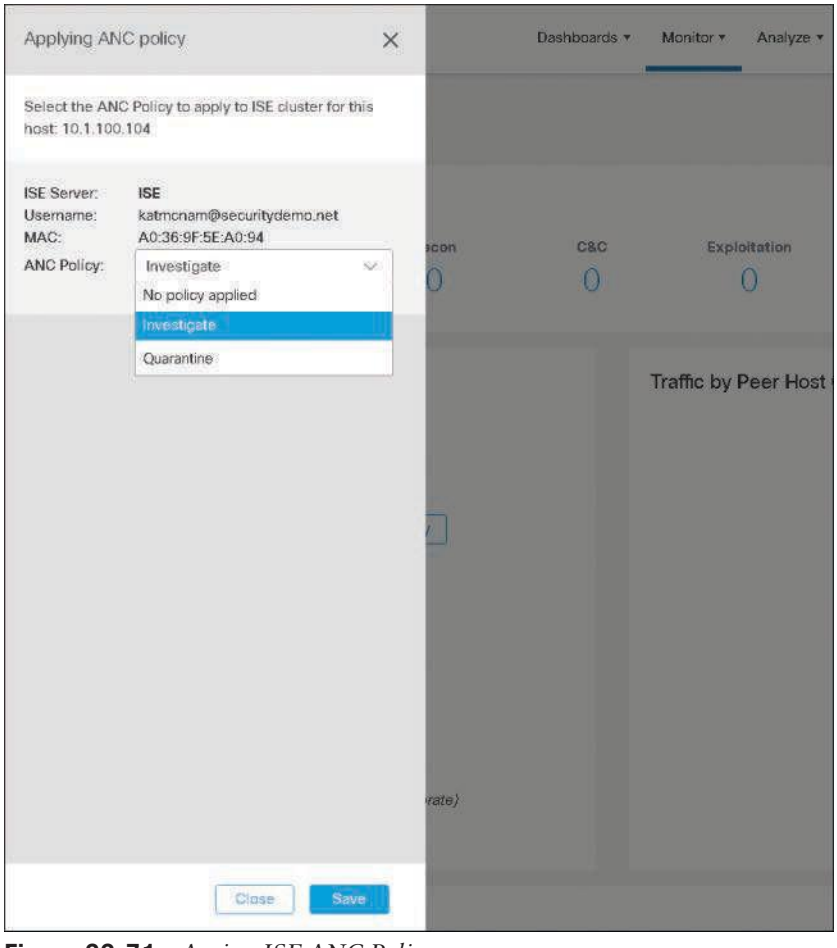


Figure 22-71 *Assign ISE ANC Policy*

Figure 22-72 shows the ISE RADIUS Live Log, where you can see ANC triggering a CoA.



Figure 22-72 *RADIUS Live Log*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 22-2 lists these key topics and the page number on which each is found.

Key Topic

Table 22-2 Key Topics

Key Topic Element	Description	Page Number
Paragraph	Integrations in ISE	820
Section	Rapid Threat Containment and the actions that can be taken	821
Section	The architecture and abilities of pxGrid	825

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

integration, Endpoint Protection Service (EPS), Adaptive Network Control (ANC), topic, subscriber, publisher, Context-Out, Context-In, Rapid Threat Containment

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is the goal of pxGrid?
2. What is the goal of Rapid Threat Containment?
3. What did Endpoint Protection Service (EPS) provide?
4. What is the value of Adaptive Network Control?
5. What is the difference between Context-In and Context-Out?



CHAPTER 23

Threat Centric NAC

This chapter covers the following topics:

Integrating ISE with vulnerability assessment tools

Integrating ISE with sources for threat intelligence

When you examine breach reports from companies like Cisco and Verizon, a common theme appears. Many breaches are caused by older vulnerabilities for which a patch or fix has been available for a long time. Consider the EternalBlue exploit as an example.

The National Security Agency developed the EternalBlue exploit, which was leaked to the world by the Shadow Brokers hacker group on April 14, 2017 (see https://blog.talosintelligence.com/2017/04/shadow-brokers.html). EternalBlue exploited a vulnerability in the Microsoft Server Message Block (SMB) protocol, which was fixed by Microsoft through a patch for the Microsoft Security Bulletin MS17-010 one month before Shadow Brokers released the exploit to the world.

On May 12, 2017, the WannaCry ransomware wreaked havoc all over the world, and that ransomware leveraged the EternalBlue exploit. Basically, two months after Microsoft released the patch, WannaCry wreaked havoc using a vulnerability for which there was already a patch. WannaCry worked because a large portion of the world had not yet deployed the patch for MS17-010.

In Chapter 18, "Posture Assessment," you learned all about posture assessment of endpoints, a foundational pillar of network access control. The idea behind NAC is to ensure that a device meets the security criteria of the company before the device is allowed to have access. That's the purpose of extending ISE with Threat Centric Network Access Control (TC-NAC). In addition to using traditional posture assessment and enforcement, you can examine data from vulnerability assessment tools and threat intelligence sources. If those sources find vulnerabilities or threats with an endpoint, you can allow ISE to take action on the endpoint.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 23-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 23-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Vulnerabilities and Threats, Oh My!	1, 3
Integrating Vulnerability Assessment Sources	5–8
Integrating with Threat Sources	2, 4, 9, 10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following is true of Threat Centric NAC (TC-NAC)?
   a. It makes it possible for integrated solutions to assign policies (or labels) that quarantine an endpoint that is behaving suspiciously.
   b. It is Cisco's next-generation access control system, which combines network access control, posture, and profiling.
   c. It is the publish/subscribe (pub/sub) communication bus used to share security intelligence between products.
   d. It provides the ability to use vulnerability assessment and threat intelligence as conditions in authorization.
2. Which of the following statements about the AMP connector for TC-NAC is true?
   a. AMP indicators and incidents are added as authorization conditions for access control.
   b. AMP does not integrate with TC-NAC.
   c. Indicators refer to indicators of compromise, and incidents are AMP events.
   d. AMP sends CVSS scores for vulnerabilities found on endpoints, and those CVSS scores are leveraged in as authorization conditions for access.
3. Which of the following statements describes a vulnerability?
   a. It is a language used to share cyber threat intelligence.
   b. It is a weakness in computer software or hardware.
   c. It is a weaponized software package that takes advantage of a security hole.
   d. It is a virus.

4. What kind of proxy does the AMP adapter for TC-NAC require?
   a. The AMP adapter does not support a proxy for Internet access.
   b. The AMP adapter requires the use of the Socks proxy.
   c. It must be a Cisco Web Security Appliance (WSA).
   d. Any standard HTTP/HTTPS web proxy may be used.
5. Which of the following statements regarding TC-NAC operational flows is most accurate?
   a. Policies can be created to limit network access or provide full access until a result from the vulnerability scanner is returned.
   b. ISE must provide limited access to the endpoint and then wait for a response from the vulnerability scanner before providing full network access.
   c. The AMP agent communicates to AnyConnect, and during the posture assessment process, ISE controls access based on the AMP agent's reported vulnerabilities.
   d. ISE must provide full access to the endpoint and then, when a response from the vulnerability scanner is returned, authorization may be changed.
6. Which of the following statements regarding the TC-NAC Live Log is most true?
   a. The TC-NAC Live Log shows all the transitions for an endpoint, from initial network access to the final threat-related Change of Authorization.
   b. There is no TC-NAC Live Log; there is only RADIUS Live Log.
   c. An entry appears in the TC-NAC Live Log only if a CoA changed authorization for an endpoint due to a vulnerability or compromise.
   d. An entry appears in the TC-NAC Live Log when an authorization result includes a vulnerability scan. Another entry appears when the results from the vulnerability scan are returned.
7. How does ISE integrate with Cisco Cognitive Threat Analytics (CTA)?
   a. ISE uses STIX for threat intelligence and TAXII for transport.
   b. ISE leverages STIX for threat intelligence and pxGrid for transport.
   c. CTA uses CVSS for threat intelligence and TAXII for transport.
   d. CTA uses CVE for threat intelligence and pxGrid for transport.
8. Where should an ISE administrator go to verify the CVSS score sent by the vulnerability scanner?
   a. TACACS Live Log
   b. RADIUS Live Log
   c. Vulnerability Assessment report
   d. Threat-Events report

9. Which of the following statements regarding the AMP adapter is most accurate?
   a. ISE integrates with all three public cloud instances of AMP, as well as on-premise private cloud installations.
   b. You must copy and paste the API username and password from the AMP console into the ISE UI.
   c. ISE gets authorized with an OAuth token to communicate with the AMP cloud and receive events related to a selected group of computers.
   d. ISE gets authorized with an OAuth token to communicate with the AMP cloud and receive events related to all endpoints.
10. Which of the following statements is most accurate regarding the CTA adapter?
   a. You must copy and paste the STIX/TAXII username and password from the CTA dashboard into the ISE UI.
   b. ISE integrates with the public and on-premises private cloud implementations of CTA.
   c. ISE gets authorized with an OAuth token to communicate with the CTA cloud and receive events related to a selected group of computers.
   d. You enter an administrative username and password for ISE to log in to CTA and retrieve the data on demand.

Foundation Topics

Vulnerabilities and Threats, Oh My!

There are lots of endpoints on a corporate network, but only a subset of these endpoints, representing a small portion of the ISE installation base, can go through posture assessment. In addition, ISE is often operated and controlled by the network security team, which can be different from the information security or operations team (InfoSec or SecOps). It's very rare for an organization to have a security team that does not invest heavily in vulnerability assessment tools. It would be nothing short of a crime if a policy engine, such as ISE, did not support these tools and bring added value to the investments that were made in both network security and the InfoSec team.

Key Topic

In ISE, Threat Centric Network Access Control (TC-NAC) enables you to create authorization policies based on the indicators of compromise (IoCs) and vulnerability attributes received from the threat and vulnerability products with which ISE can be integrated.

Threat severity levels and vulnerability assessment results can be used to dynamically control the access level assigned to an endpoint or a user. You can configure the vulnerability and threat adapters to send high-fidelity IoCs, detected threat events, and CVSS scores to Cisco ISE so that TC-NAC access policies can be created to change the privilege and context of an endpoint. Figure 23-1 illustrates how TC-NAC adds threats and vulnerabilities to the access control policy options so ISE can flag compromised and vulnerable hosts and limit access for remediation.

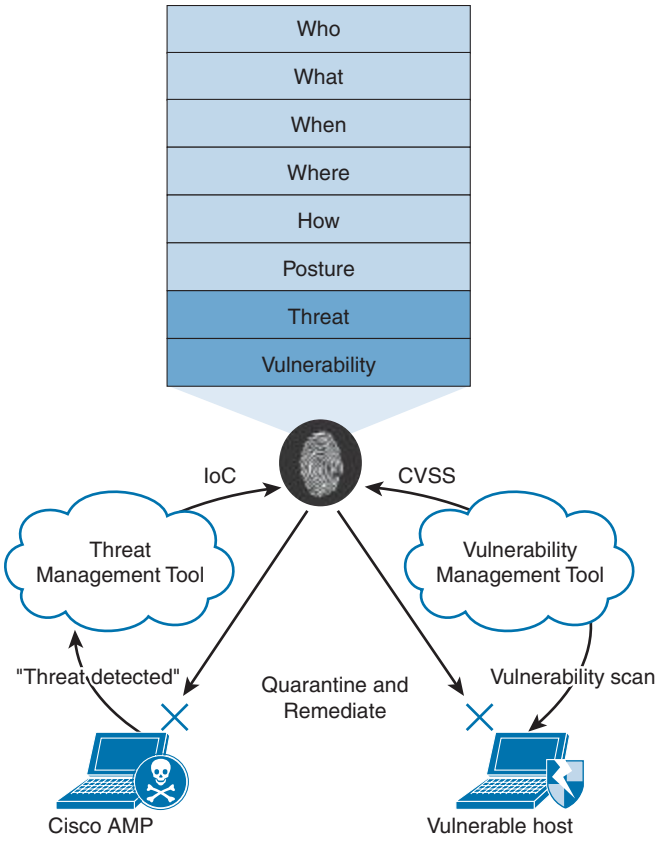


Figure 23-1 *TC-NAC Design*

ISE classifies vulnerabilities and threats differently, and different products integrate with ISE for both vulnerabilities and threats.

Integrating Vulnerability Assessment Sources

Key Topic

As of Version 2.7, ISE integrates with three vulnerability assessment tools: Qualys Enterprise Edition, Tenable Security Center, and Rapid7 Nexpose. This chapter shows examples of integration with Tenable Security Center (Tenable.SC) using a Tenable Nessus vulnerability scanner.

Before we dive into the flows, we need to consider the definitions of some important terms:

- A ***vulnerability*** (or *security hole*) is a weakness in computer software or hardware. Sometimes the shortened term *vuln* is used instead of *vulnerability*.
- An ***exploit*** is a weaponized software package that takes advantage of a vulnerability to accomplish a task that is not meant to be permitted. A vulnerable means there is theoretically a way to exploit something (that is, a vulnerability exists), and an exploit means there is a definite path to doing so.

- ***Common Vulnerabilities and Exposures (CVE)*** is a system to track, monitor, and describe publicly known security vulnerabilities.
- ***Common Vulnerability Scoring System (CVSS)*** is an open standard for assessing the severity of security vulnerabilities. CVSS makes it possible to assign severity scores from 0 to 10 to vulnerabilities so that incident responders can prioritize their responses.

Vulnerability scanners provide CVSS scores to ISE as part of TC-NAC, and those scores can be used in authorization rules.

TC-NAC Flows

An ISE authorization policy can be configured to trigger a vulnerability scan when an endpoint joins the network. This is what everyone thinks they want when they consider implementing network access control. They think it is a way to ensure that every endpoint attaching to the company network is posture compliant and vulnerability free before it is permitted access. Figure 23-2 shows an example of such a flow using vulnerability scanners.

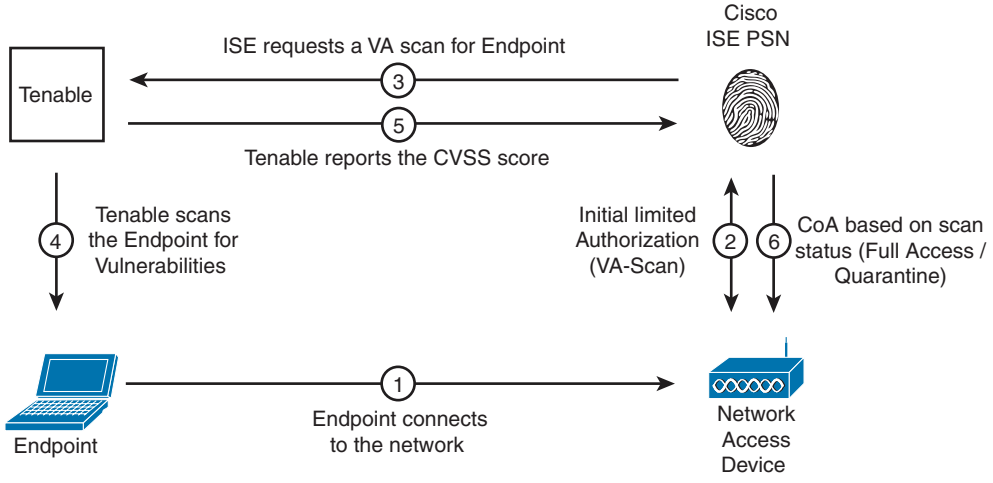


Figure 23-2 *Vulnerability-Based If Access Control*

Figure 23-2 illustrates six steps:

Step 1. An endpoint connects to the network.

Step 2. Limited access is granted.

Step 3. Part of the authorization result is to trigger a scan with the vulnerability scanner.

Step 4. The vulnerability scanner scans the endpoint.

Step 5. The resulting vulnerabilities and their CVSS score are reported to ISE.

Step 6. A CoA is issued, and full network access is granted.

The flow in Figure 23-2 illustrates a "scan before entry" model, but that is not always what a real deployment looks like. Vulnerability scanners can often be backed up, with a queue of scans waiting to be performed. Depending on how in depth a scan is, it might add significant time to the network access, and that is very undesirable behavior. To prevent such a situation, the flow could be reversed so that the flow triggers a new scan and allows the endpoint into the network. If the scan is returned with a high CVSS score, a CoA can be issued, and the endpoint's access can be changed. This more common approach to TC-NAC is illustrated in Figure 23-3.

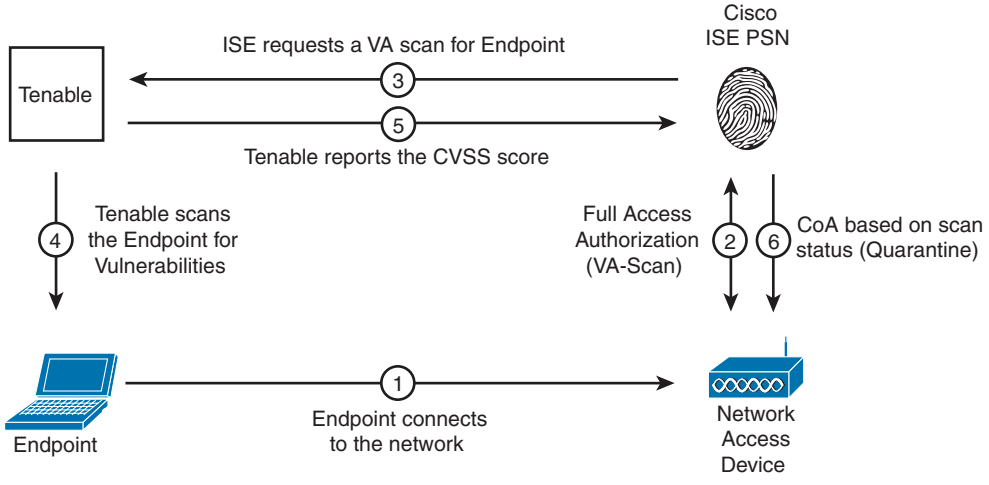


Figure 23-3 *Allowing Access and Changing If Vulnerable*

Figure 23-3 illustrates six steps:

Step 1. The endpoint connects to the network.

Step 2. Normal (full) access is granted.

Step 3. Part of the authorization result is to trigger a scan with the vulnerability scanner.

Step 4. The vulnerability scanner scans the endpoint.

Step 5. The resulting vulnerabilities and their CVSS score are reported to ISE.

Step 6. If the CVSS scores reported exceed the threshold, a CoA is issued to provide limited access for the endpoint to be remediated.

Enable TC-NAC

Before you can configure TC-NAC, it must first be enabled on a Policy Services node (PSN). Keep in mind that TC-NAC can be enabled on exactly one PSN, and you receive an error message if you try to enable it on a second PSN in a deployment. You enable TC-NAC in the Administration > Deployment screen in the ISE GUI (see Figure 23-4).

When you enable the TC-NAC service on a PSN, a few things happen. Example 23-1 shows the output of the **show application status ise** command on the PSN before you enable the TC-NAC service. You can see in this output that the TC-NAC service is disabled.

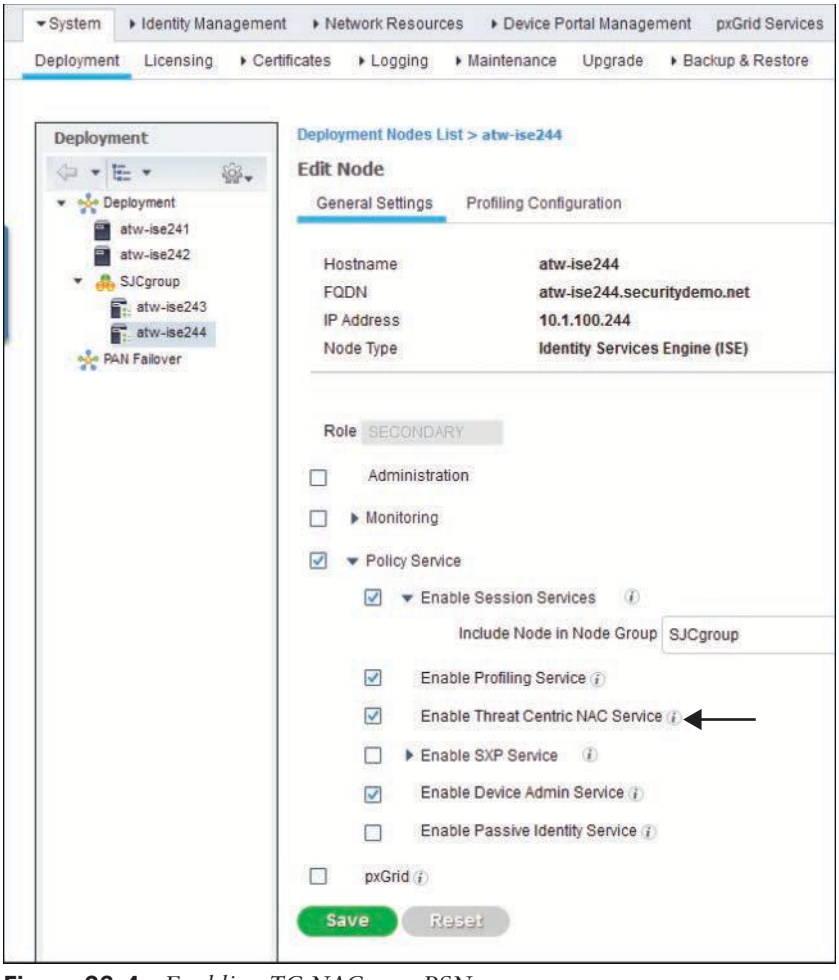


Figure 23-4 *Enabling TC-NAC on a PSN*

Example 23-1 *PSN Services Prior to Enabling TC-NAC*

```
atw-ise244/admin# show application status ise

ISE PROCESS NAME                       STATE            PROCESS ID
--------------------------------------------------------------------
Database Listener                      running          21719
Database Server                        running          73 PROCESSES
Application Server                     running          32747
Profiler Database                      running          29389
ISE Indexing Engine                    running          2540
```

```
AD Connector                          running          10435
M&T Session Database                  running          29195
M&T Log Processor                     running          461
Certificate Authority Service         running          7198
EST Service                           running          1311
SXP Engine Service                    running          7771
Docker Daemon                         running          23282
TC-NAC Service                        disabled

Wifi Setup Helper Container           disabled
pxGrid Infrastructure Service         running          1517
pxGrid Publisher Subscriber Service   running          1820
pxGrid Connection Manager             running          1746
pxGrid Controller                     running          1912
PassiveID WMI Service                 running          8256
PassiveID Syslog Service              running          8621
PassiveID API Service                 running          9022
PassiveID Agent Service               running          9416
PassiveID Endpoint Service            running          9752
PassiveID SPAN Service                running          10292
DHCP Server (dhcpd)                   disabled
DNS Server (named)                    disabled
ISE Messaging Service                 running          25880
Segmentation Policy Service           disabled
SSE Connector                         disabled
```

When you enable TC-NAC, the TC-NAC service spins up four different Docker containers and then starts those microservices. Example 23-2 shows that the TC-NAC service displayed in Example 23-1 has not transitioned to four different services.

Example 23-2 *Four TC-NAC Containers in the Running State*

```
atw-ise244/admin# show application status ise

ISE PROCESS NAME                      STATE            PROCESS ID
--------------------------------------------------------------------
Database Listener                     running          21719
Database Server                       running          87 PROCESSES
Application Server                    running          32747
Profiler Database                     running          29389
ISE Indexing Engine                   running          2540
AD Connector                          running          10435
M&T Session Database                  running          29195
M&T Log Processor                     running          461
Certificate Authority Service         running          7198
```

```
EST Service                          running          1311
SXP Engine Service                   running          7771
Docker Daemon                        running          23282
TC-NAC MongoDB Container             not running
TC-NAC Core Engine Container         not running
VA Database                          not running
VA Service                           not running

Wifi Setup Helper Container          disabled
pxGrid Infrastructure Service        running          1517
pxGrid Publisher Subscriber Service  running          1820
pxGrid Connection Manager            running          1746
pxGrid Controller                    running          1912
PassiveID WMI Service                running          8256
PassiveID Syslog Service             running          8621
PassiveID API Service                running          9022
PassiveID Agent Service              running          9416
PassiveID Endpoint Service           running          9752
PassiveID SPAN Service               running          10292
DHCP Server (dhcpd)                  disabled
DNS Server (named)                   disabled
ISE Messaging Service                running          25880
Segmentation Policy Service          disabled
SSE Connector                        disabled
```

Example 23-3 shows the status after the four TC-NAC containers are in the running state.

Example 23-3 *Four TC-NAC Containers in Running State*

```
atw-ise244/admin# show application status ise

ISE PROCESS NAME                     STATE            PROCESS ID
--------------------------------------------------------------------
Database Listener                    running          21719
Database Server                      running          90 PROCESSES
Application Server                   running          32747
Profiler Database                    running          29389
ISE Indexing Engine                  running          2540
AD Connector                         running          10435
M&T Session Database                 running          29195
M&T Log Processor                    running          461
Certificate Authority Service        running          7198
EST Service                          running          1311
SXP Engine Service                   running          7771
Docker Daemon                        running          20941
TC-NAC MongoDB Container             running          2736
```

```
TC-NAC Core Engine Container          running          3416
VA Database                           running          3860
VA Service                            running          4036

Wifi Setup Helper Container           disabled
pxGrid Infrastructure Service         running          1517
pxGrid Publisher Subscriber Service   running          1820
pxGrid Connection Manager             running          1746
pxGrid Controller                     running          1912
PassiveID WMI Service                 running          8256
PassiveID Syslog Service              running          8621
PassiveID API Service                 running          9022
PassiveID Agent Service               running          9416
PassiveID Endpoint Service            running          9752
PassiveID SPAN Service                running          10292
DHCP Server (dhcpd)                   disabled
DNS Server (named)                    disabled
ISE Messaging Service                 running          25880
Segmentation Policy Service           disabled
SSE Connector                         disabled
```

Configure the Integration with a Vulnerability Assessment Vendor

After TC-NAC has been enabled, you can add and configure TC-NAC integrations through Administration > Threat Centric NAC > Third Party Vendors, as shown in Figure 23-5.

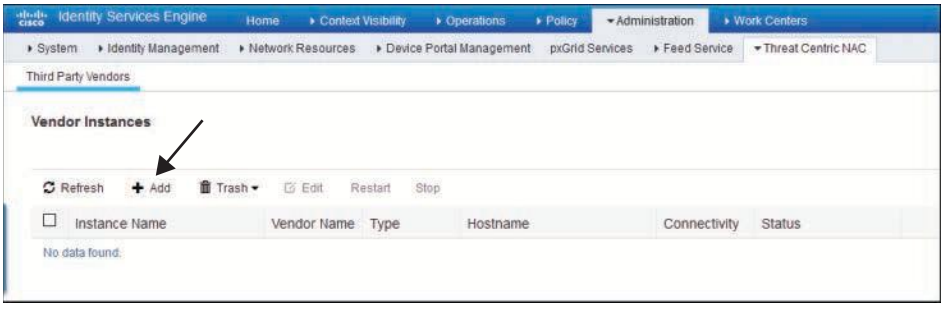


Figure 23-5 *TC-NAC Vendor Instances*

Click Add to create your first integration. Figure 23-6 shows the list that is presented when you click the Vendor dropdown. Each vendor that is supported also has an attribute presented next to its name: either VA (vulnerability assessment) or THREAT.

To add the Tenable vulnerability assessment vendor, select Tenable Security Center : VA, as shown in Figure 23-7, and click Save.

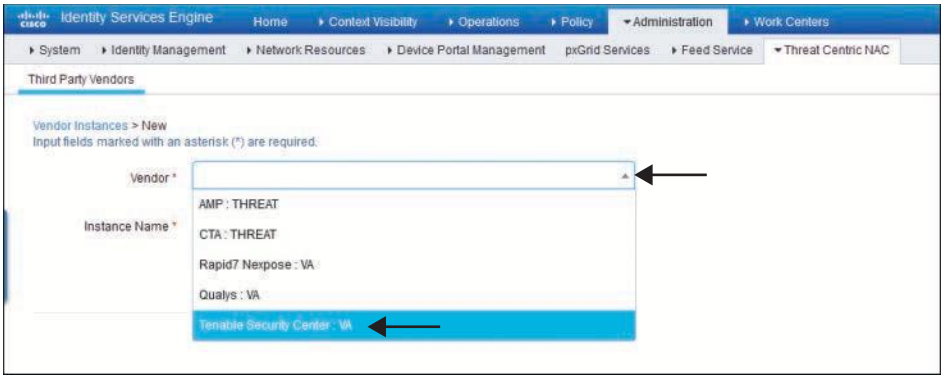


Figure 23-6 *TC-NAC Vendors, Part 1*

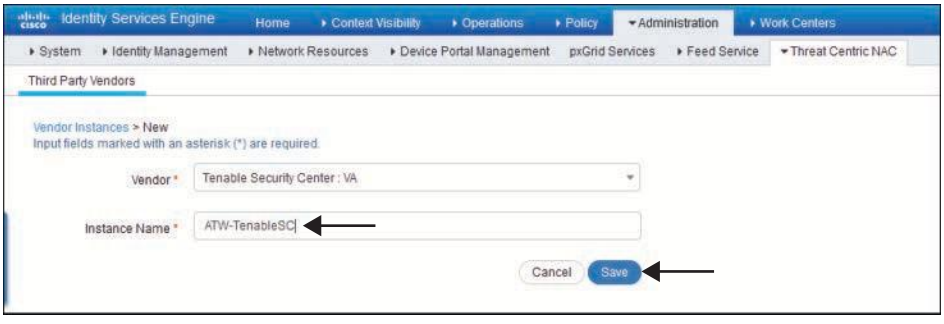


Figure 23-7 *TC-NAC Vendors, Part 2*

It is possible to have multiple instances of the same vendor type. Perhaps you want to integrate with multiple Tenable Security Centers (Tenable.SC), such as one in the United States and another in Europe. In this case, you provide a name for each instance that you are connecting to and click Save, as shown in Figure 23-7.

After you click Save, the TC-NAC instance is prepared in the background, and the status appears as Ready to Configure when that instance is ready to be set up, as shown in Figure 23-8.

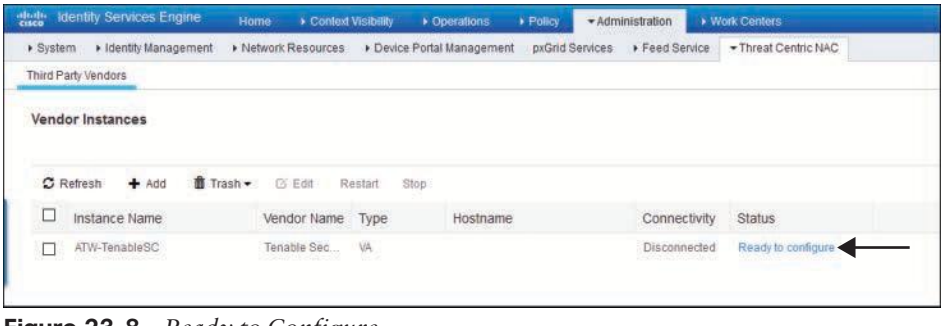


Figure 23-8 *Ready to Configure*

Click Ready to Configure to begin setting up the new instance. To start the configuration, you need to provide the basic connectivity information required to access Tenable.SC, such as the hostname and port, as well as the username and password of a Tenable.SC user, as shown in Figure 23-9.

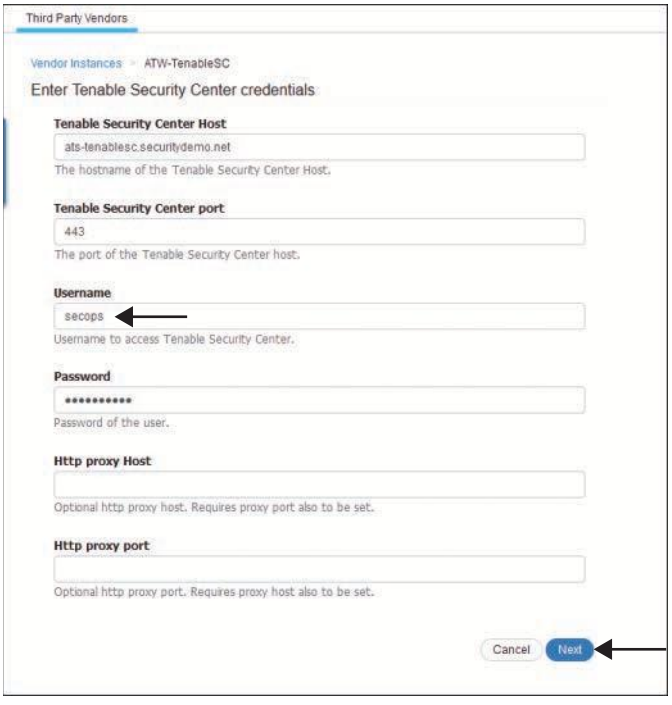


Figure 23-9 *Basic Setup of a Vulnerability Assessment Vendor*

The Tenable.SC user must be a member of the Security Manager role and not the Admin user. The Admin user has access to the Tenable Core application but not the actual scanning and security settings required. Figure 23-10 shows the Tenable.SC users, which are all in the Security Manager role.

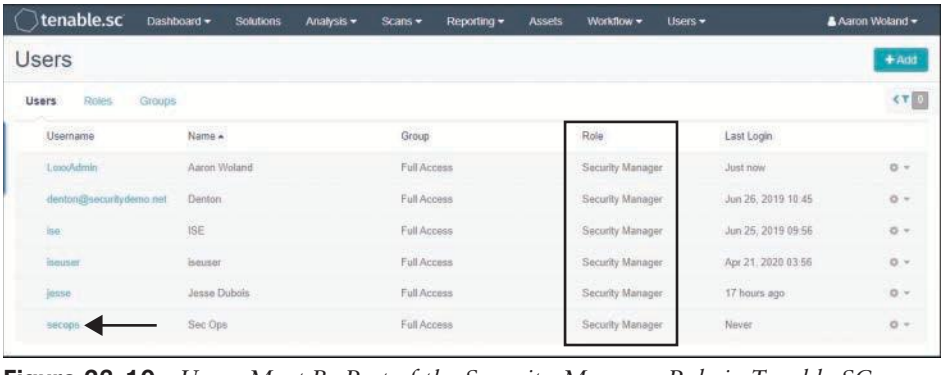


Figure 23-10 *Users Must Be Part of the Security Manager Role in Tenable.SC*

After filling in the basic connectivity information for the integration, click Next to enter the Advanced Settings for the connection, as shown in Figure 23-11. Each vendor has a unique set of advanced settings, based on the way that vendor's product works. The example used in this book is Tenable.SC.

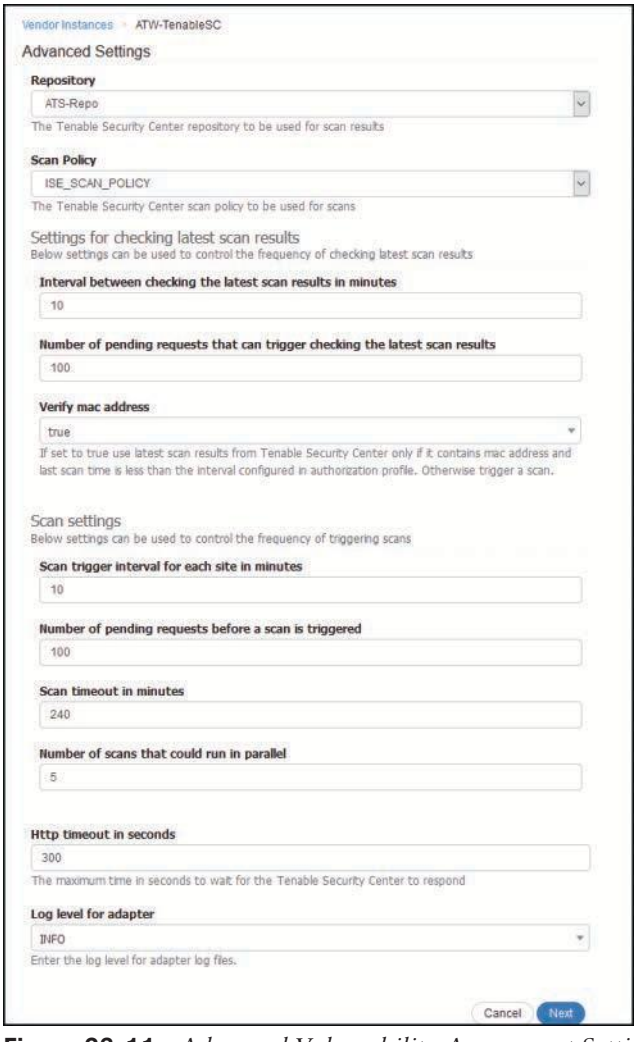


Figure 23-11 *Advanced Vulnerability Assessment Settings*

As shown in Figure 23-11, you can configure a fairly large number of advanced vulnerability assessment settings:

- **Repository:** This is a mandatory field in which you define where the scan results are stored within Tenable.SC. The repository is created within the Tenable.SC user interface, and the available repositories appear in the Repository dropdown in the ISE Advanced Settings screen. Figure 23-12 shows the Repositories in the Tenable.SC user interface.



Figure 23-12 *Tenable.SC Repositories*

- **Scan Policy:** Many scan policies can be created in a vulnerability assessment tool to define what to scan and how deeply to scan an endpoint. The scans are created in the Tenable.SC user interface, and the available repositories appear in the Scan Policy dropdown. Figure 23-13 shows the Policies screen in the Tenable.SC user interface.

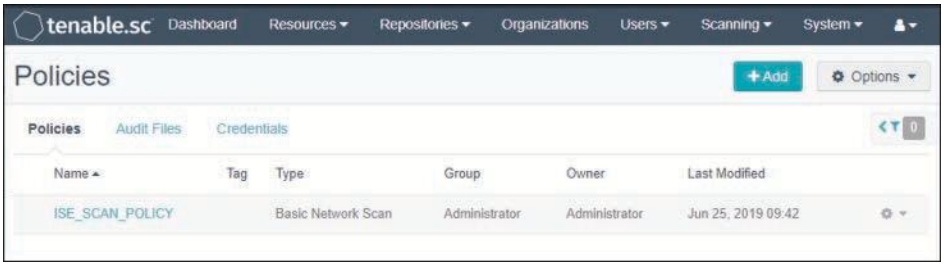


Figure 23-13 *Tenable.SC Scans*

In many cases, the default settings in the ISE Advanced Settings screen suffice. The following settings need to be changed only if you are fine-tuning your environment:

- **Interval Between Checking the Latest Scan Results in Minutes:** This setting controls how often ISE checks with Tenable.SC.
- **Number of Pending Requests That Can Trigger Checking the Latest Scan Results:** If ISE has queued this number of requests without response from Tenable.SC, this triggers an automatic poll to check for results.
- **Verify MAC Address:** When this option is set to true, ISE does not trigger a scan if the latest scan results contain the MAC address of the endpoint and the last scan time is lower than the interval defined in the authorization profile itself.
- **Scan Trigger Interface for Each Site in Minutes:** You can configure ISE to queue the scan requests and send the requests every 10 minutes.
- **Number of Pending Requests Before a Scan Is Triggered:** You can configure ISE to immediately trigger a scan when the queue has grown to this configured number.
- **Can Time Out in Minutes:** You can configure how long to wait before timing out the scan and removing it from the queue.

- **Number of Scans That Could Run in Parallel:** You can configure how many scans to trigger simultaneously.
- **HTTP Timeout in Seconds:** You can configure how many seconds to wait for Tenable.SC to respond to requests.
- **Log Level for Adapter:** You can set this to ERROR, INFO, DEBUG, or TRACE, and this setting should be changed only under the direction of Cisco TAC for trouble-shooting a TC-NAC integration.

Click Next and then click Finish, as shown in Figure 23-14. Figure 23-15 shows a configured vendor instance.

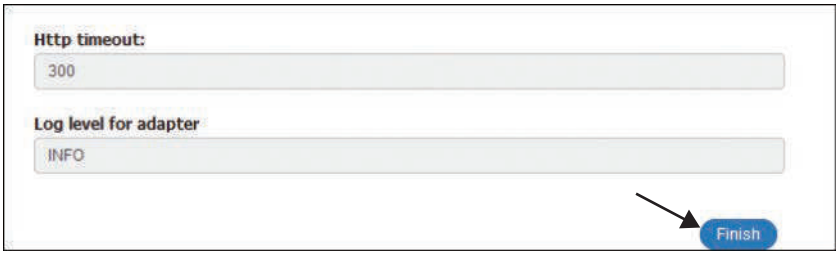


Figure 23-14 *Clicking Finish to Complete the Instance Configuration*

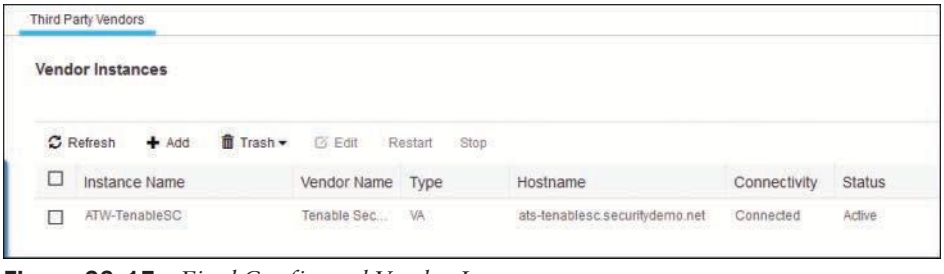


Figure 23-15 *Final Configured Vendor Instance*

When a configuration is finished, it must be reconfigured in order to make changes. To do this, you select an instance and click Edit, as shown in Figure 23-16.

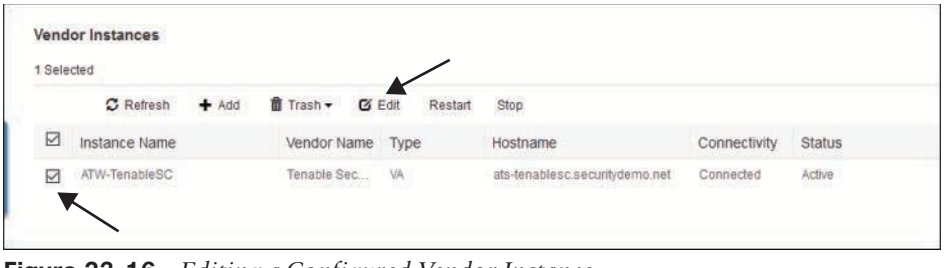


Figure 23-16 *Editing a Configured Vendor Instance*

To edit a configured instance, you must click Reconfigure in order to change any settings, as shown in Figure 23-17.

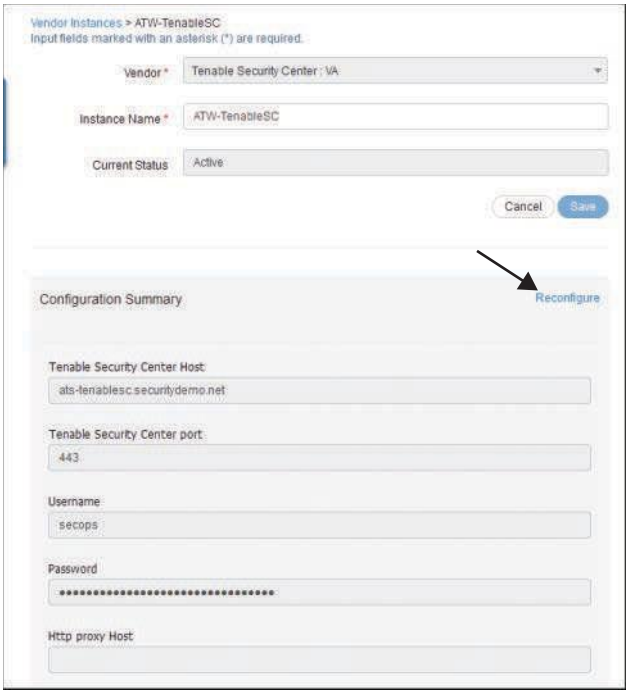


Figure 23-17 *Reconfiguring an Instance*

Authorization Profile and Authorization Rules

Configuring an integration is naturally a crucial part of configuring TC-NAC, but there is no network access control in TC-NAC without the appropriate authorization rules and profiles.

To permit access and assess vulnerabilities using the configured vulnerability assessment instance, navigate to Work Centers > Network Access > Policy Elements > Authorization Profiles and add a new Authorization Profile named PermitScan, as shown in Figure 23-18.

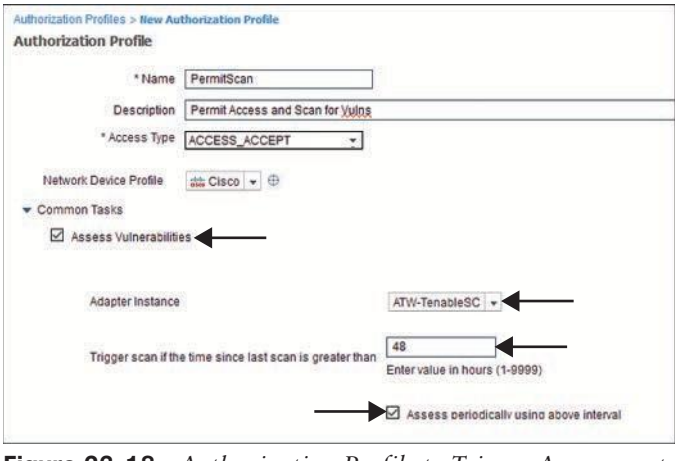


Figure 23-18 *Authorization Profile to Trigger Assessment*

Creating authorization rules can be a very subjective process, and it is best to identify what the organization really needs when it comes to TC-NAC.

The "most paranoid" thought process is to add a vulnerability check to each and every authorization to the ISE network. This way, ISE is tracking each and every endpoint with the vulnerability assessment tool(s). Doing this also creates the most load on the entire system, which isn't always desirable.

The "not paranoid enough" thought process is to add vulnerability checks only to the most secure authorizations granted, sch as those permitted to access PCI or PII data.

Most experienced implementors will tell you that the best path is usually somewhere in the middle of these two thought processes.

Figure 23-19 shows how to add the PermitScan result to all rules that are currently configured for PermitAccess.

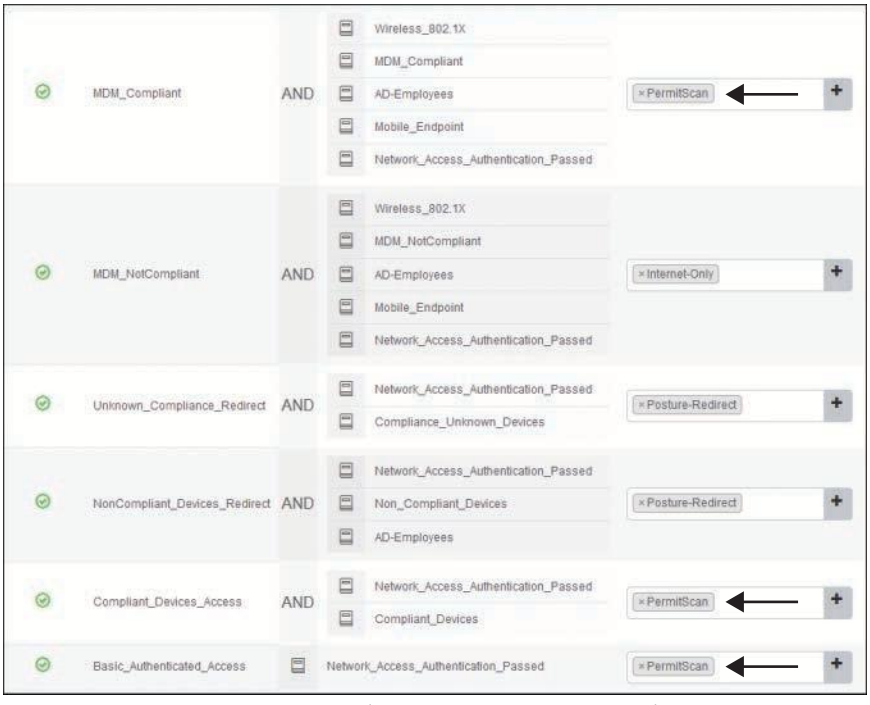


Figure 23-19 *PermitScan Results in an Authorization Policy*

Now you have rules that will trigger the scan and tracking of the endpoint. What happens when the endpoint is vulnerable? That is where you insert a rule near the top of the authorization policy that locks down access to any devices that have a CVSS score higher than the threshold defined by your security team. This would normally be a high-risk CVSS, but for learning purposes, in this chapter shows this configured as anything with a CVSS score higher than 4.

Figure 23-20 shows all the conditions available from the Threat dictionary in the conditions studio. Since our example is using Tenable.SC, the condition to select is Tenable Security Center-CVSS_Base_Score, as shown in Figure 23-21.

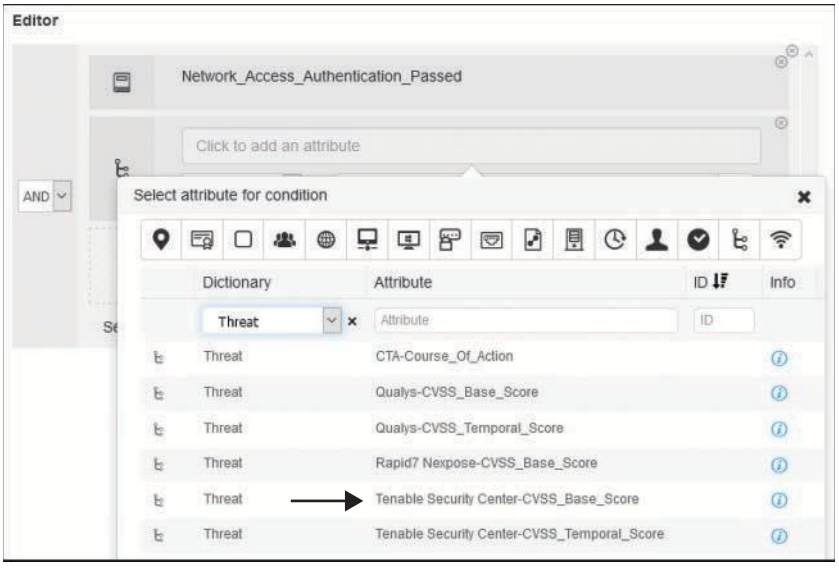


Figure 23-20 *Threat Dictionary Conditions*

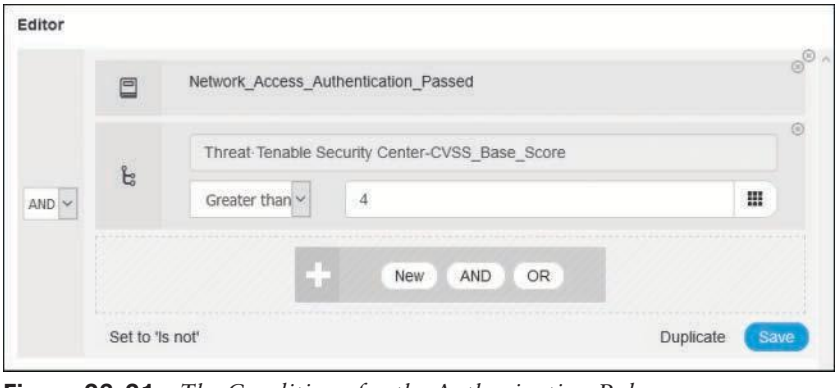


Figure 23-21 *The Conditions for the Authorization Rule*

Figure 23-22 shows the Vulnerable-LimitAccess authorization rule at the top of the authorization policy, with the two conditions in the rule providing Internet-only access and assigning the Security Group Tag (SGT) named Vulnerable.



Figure 23-22 *Vulnerable-LimitAccess Authorization Rule*

Seeing TC-NAC with Vulnerability Scanners in Action

Figure 23-23 shows a composite of the filtered Live Log, illustrating the transition of the endpoint's status on the network and the TC-NAC process. The numerals in the figure correspond to the numerals in the following list:

Details	Identity	Endpoint Profile	Authorization Policy	Authorization Profiles	Posture Status
	secops1 ×				
5	secops1	Windows10-Workstation	Default >> Vulnerable - Limit Access	Vulnerable,Internet-Only	
4	secops1	Windows10-Workstation	Default >> Vulnerable - Limit Access	Vulnerable,Internet-Only	
3	secops1	Microsoft-Workstation	Default >> Compliant_Devices_Access	PermitScan,Employees	Compliant
2	secops1	Microsoft-Workstation	Default >> NonCompliant_Devices_Redirect	Posture-Redirect	NonCompliant
1	secops1	Microsoft-Workstation	Default >> Unknown_Compliance_Redirect	Posture-Redirect	Pending

Figure 23-23 *Endpoint Transition on a Network with TC-NAC*

1. A new endpoint joins the network.
   - The user is secops.
   - The endpoint is initially profiled as a Windows workstation.
   - The posture status is pending.
   - The authorization result is Posture-Redirect.
2. The endpoint's posture status is updated.
   - The posture status is NonCompliant.
   - The authorization result remains Posture-Redirect.
3. The endpoint's posture status is updated again.
   - The posture status is Compliant.
   - A CoA was sent.
   - The new authorization result is PermitScan, and the SGT is Employees.
   - PermitScan triggers Tenable.SC to scan the endpoint.
4. The scan results are returned to ISE.
   - The results must have included a base CVSS score greater than 4.
   - A CoA occurred.
   - The Vulnerable-Limit Access policy was matched.
   - The Internet-only authorization profile was applied.
   - The SGT Vulnerable was assigned to the session.
5. This is an informational entry in Live Log that shows the current state of the session.

Verifying What Happened

Live Log shows the activity from the viewpoint of the authorization rules, but it does not do much to show the result of the vulnerability scan or any of the interactions with Tenable.SC. To see those interactions, navigate to Operations > Reports > Threat Centric NAC. Figures 23-24 through 23-28 show these reports.

Figure 23-24 shows the Threat-Events report, where you can see that the endpoint received a base CVSS score of 6.4 from the vulnerability scanner, which is certainly greater than the threshold of 4 that you set in the authorization policy.

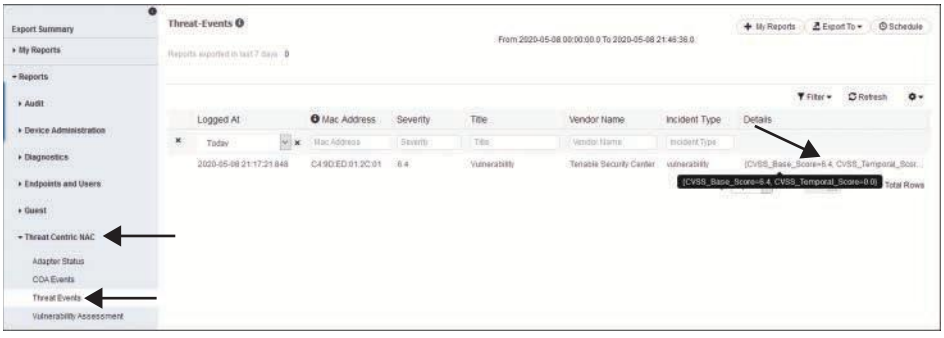


Figure 23-24 *Threat-Events Report*

Figure 23-25 shows the Vulnerability Assessment report. (Don't let the bug in the report title fool you: This report is not just for failures.) In this report, you can see that each request to the vulnerability assessment gets a unique adapter instance ID, which is a unique identifier used to track the individual requests and responses between the disjointed systems.

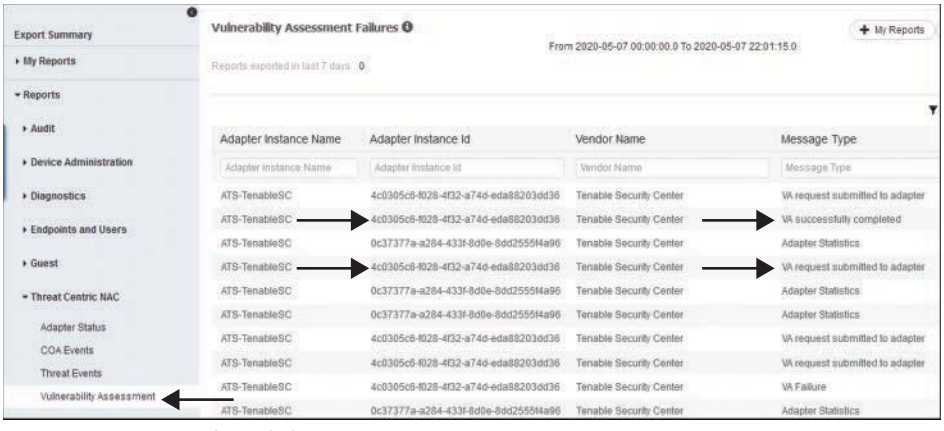


Figure 23-25 *Vulnerability Assessment Report*

When a vulnerability assessment occurs and the results are sent back to ISE, it is possible that a CoA needs to occur. When that CoA occurs, it is recorded in the Coa-Events report (shown in Figure 23-26), and if a change of access based on the TC-NAC event occurs, it is displayed in the TC-NAC Live Log (shown in Figure 23-27).

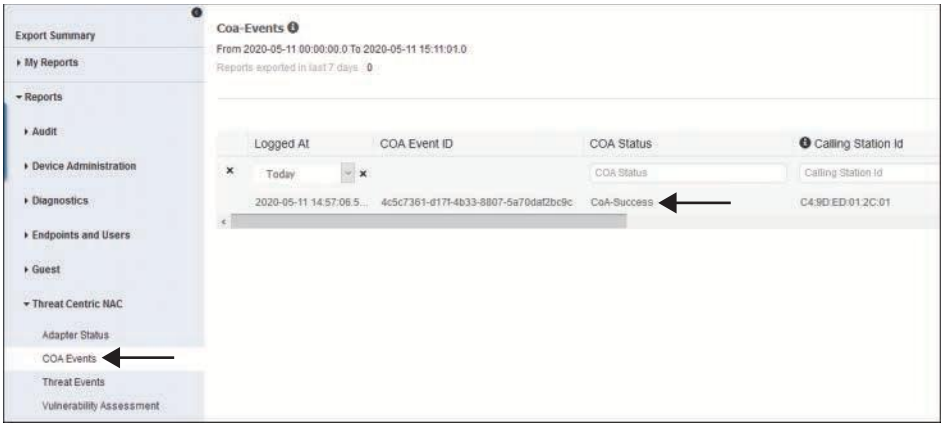


Figure 23-26 *Coa-Events Report*

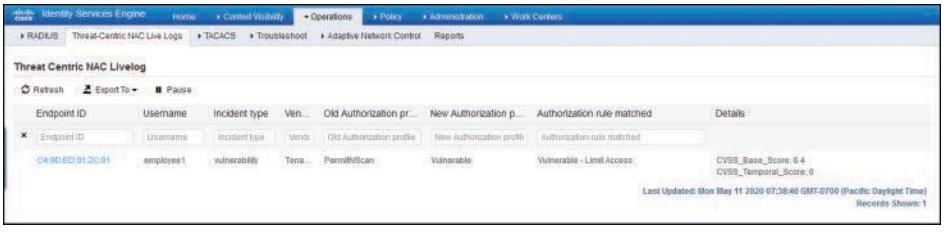


Figure 23-27 *TC-NAC Live Log*

Key Topic

It's very important to note that a successful CoA is required for any data to show up in either the Coa-Events Report or in the TC-NAC Live Log.

There are two more locations you can leverage to validate the TC-NAC behavior. When vulnerability assessment results are returned to ISE, the results are also tracked as part of the Context Visibility function, in the Vulnerable Endpoints section (see Figure 23-28). The second location is in a vulnerability assessment tool such as the Scan Results section in Tenable.SC, shown in Figure 23-29.

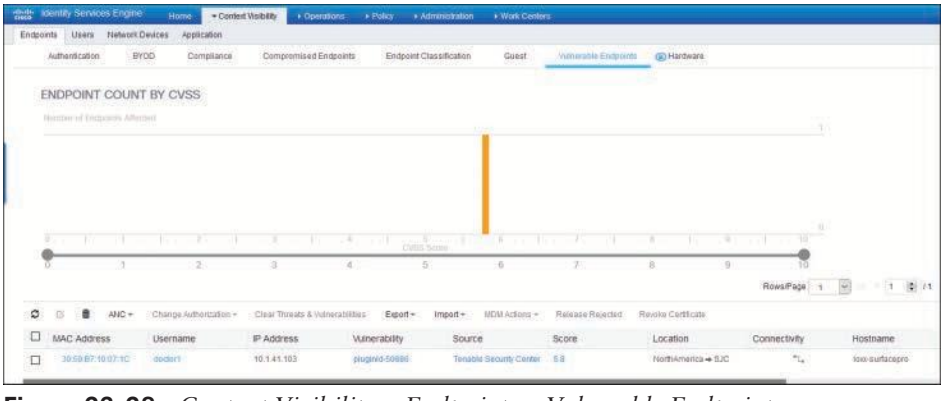


Figure 23-28 *Context Visibility > Endpoints > Vulnerable Endpoints*

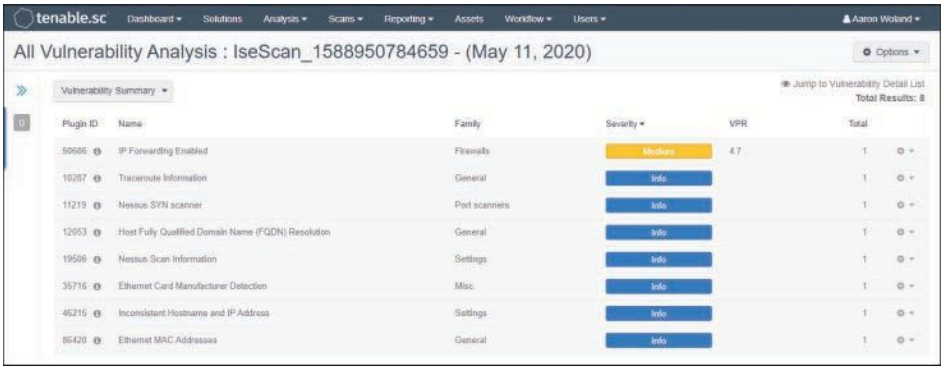


Figure 23-29 *Tenable.SC Scan Results*

Integrating with Threat Sources

Key Topic

In addition to vulnerability scanners, ISE also integrates with sources for threat intelligence. Specifically, ISE integrates with the Cisco next-generation endpoint security solution Advanced Malware Protection (AMP) and Cisco Cognitive Threat Analytics (CTA), which has since been renamed Cisco Cognitive Intelligence. These solutions are both described in this chapter.

Before we dive into the individual threat solutions with which ISE can integrate, we need to consider the definitions of some important terms:

- ***Cyber threat intelligence (CTI)*** is information about computer- and networking-related threats. Basically, CTI is data about an existing threat.
- ***Structured Threat Information eXpression (STIX)*** is a language used to share CTI. STIX is a format, not a transport protocol. It requires a protocol, such as TAXII or pxGrid, to carry it between consumers and producers of the STIX data.
- ***Trusted Automated eXchange of Intelligence Information (TAXII)*** is a protocol used to exchange CTI over secure communication (HTTPS). TAXII is designed specifically to carry STIX CTI. TAXII follows a publish/subscribe (pub/sub) model, similar to pxGrid.

NOTE Both STIX and TAXII are open standards created and maintained by the Organization for the Advancement of Structured Information Standards (OASIS).

ISE normalizes the threat intelligence from AMP into the STIX data model, and CTA shares its CTI via the STIX model, using TAXII natively.

Cognitive Threat Analytics (CTA)

Cisco's Cognitive Intelligence service, also known as CTA, is a threat analytics service that used to be sold as a separate product from Cisco's Security Business Group. The service offers advanced machine learning (ML) and artificial intelligence (AI)—which is a very "marketing way" to describe very advanced algorithms performing analytics on large data sets in order to identify threats.

The data provided by Cognitive was so valuable to the participating Cisco security services that the SBG leadership decided to make it a common service for all of the security portfolio and stop charging for it as a separate service. For example, Cisco AMP for Endpoints and the Cisco Stealthwatch analytics platform both rely on CTA as a core intelligence engine, and Figure 23-30 shows that CTA is a dashboard item in the AMP for Endpoints console.

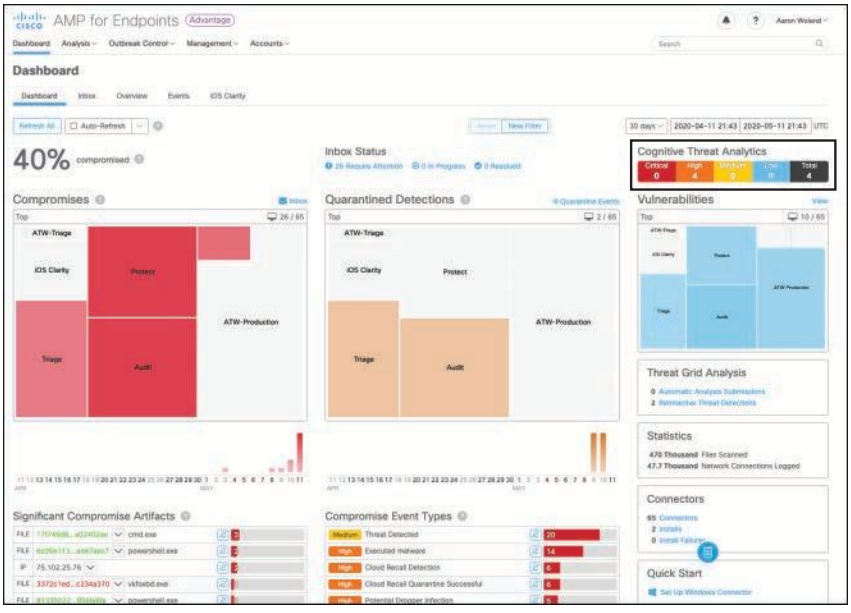


Figure 23-30 *Cognitive Dashboard Widget in the AMP for Endpoints Console*

You can click on the CTA widget on the AMP dashboard to go to the Cognitive Intelligence dashboard (see Figure 23-31). You can navigate to the dashboard directly.

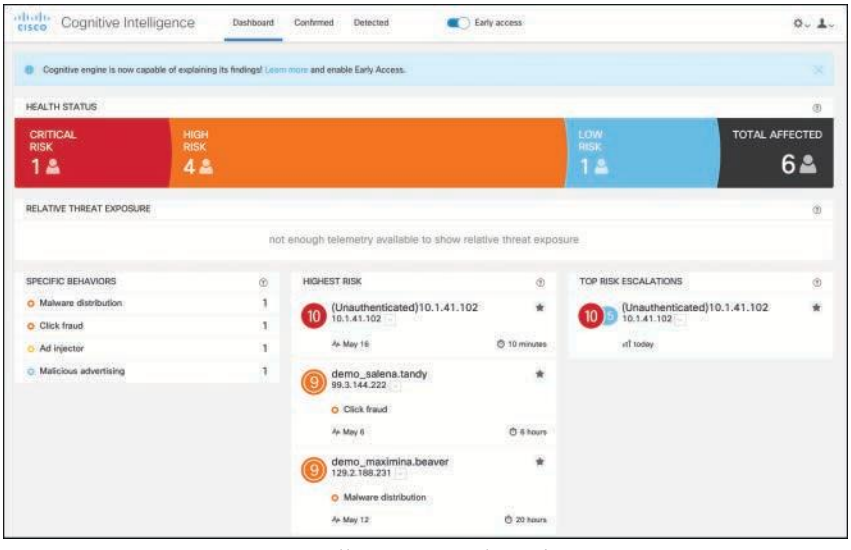


Figure 23-31 *Cognitive Intelligence Dashboard*

Notice in Figure 23-31 that CTA assigns a risk score (1–10) to each of the confirmed threats. You can click on one of the risks to bring up details about the curated threat information that lead to the curated threat intelligence and risk score, as shown in Figure 23-32.

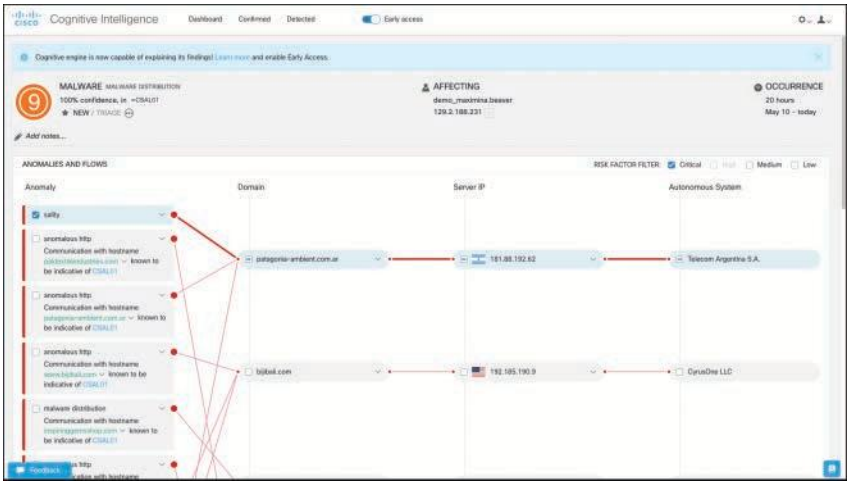


Figure 23-32 *Cognitive Intelligence Incident*

Now that you have seen an overview of what Cognitive Intelligence is and what it can provide, let's integrate it with ISE and see how TC-NAC can use this intelligence as an attribute for access control and context visibility.

Create a CTA STIX/TAXII API Account

The first step for integrating ISE and CTA together is to create a CTA STIX/TAXII API account for ISE to use for subscribing to the TAXII feed.

On the CTA dashboard, click on the gear icon in the upper-right corner and select CTA STIX/TAXII API, as shown in Figure 23-33.

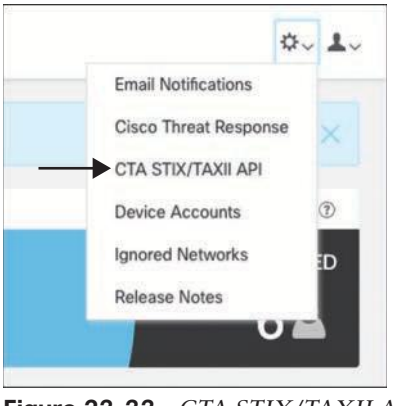


Figure 23-33 *CTA STIX/TAXII API Menu Item*

Selecting this menu option brings you to the TAXII provisioning page shown in Figure 23-34.

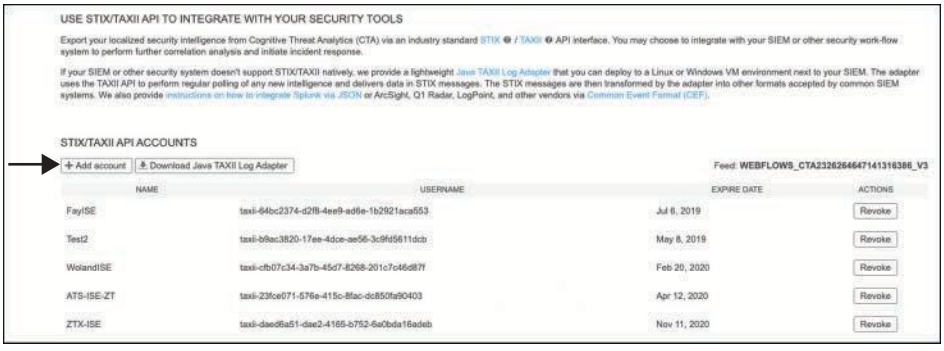


Figure 23-34 *CTA TAXII Provisioning Page*

Click Add Account, provide a name for the new account that describes what product will be using this new API account, and click Add Account, as shown in Figure 23-35.

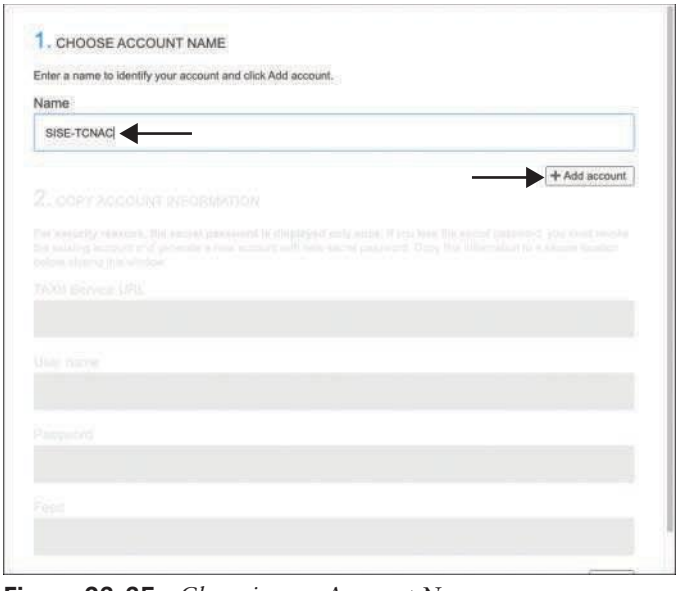


Figure 23-35 *Choosing an Account Name*

Now the account information is displayed as shown in Figure 23-36. The password shown on this screen will be displayed only this one time. If you misplace this password, you need to revoke the account and create a new one. Each piece of information you see in Figure 23-36 needs to be copied and pasted into the ISE UI in the upcoming configuration steps.

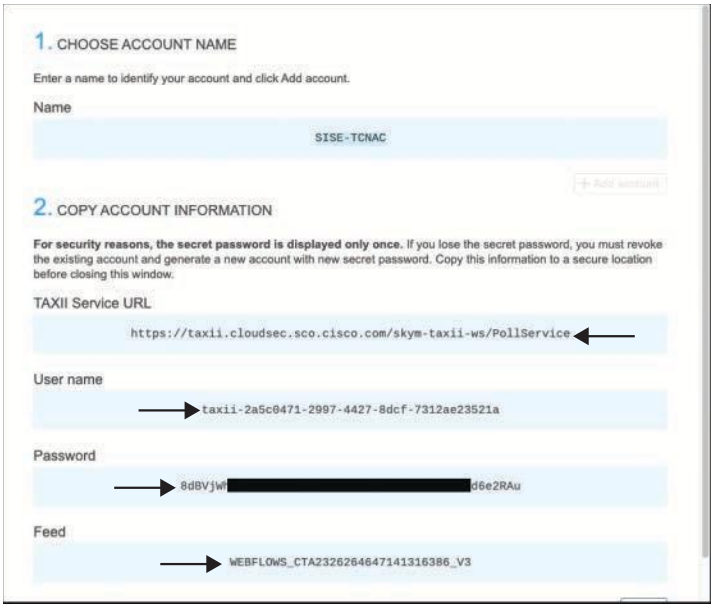


Figure 23-36 *The TAXII API Account Is Created and Displayed One Time*

Create a CTA Integration for TC-NAC

Navigate to Administration > Threat Centric NAC > Third Party Vendors, as shown in Figure 23-37, and click Add to add a new connector. Then, in the window shown in Figure 23-38, set Vendor to CTA: THREAT, name the instance, and click Save.

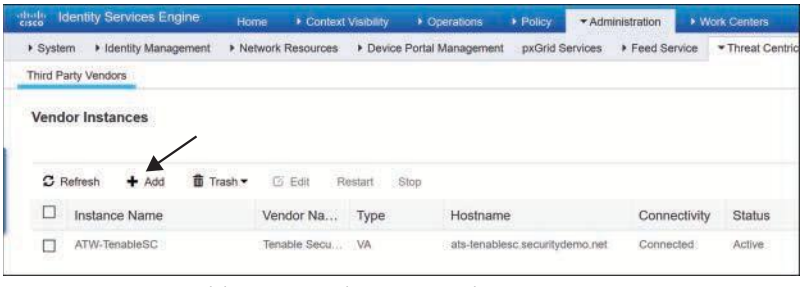


Figure 23-37 *Adding a Third-Party Vendor*

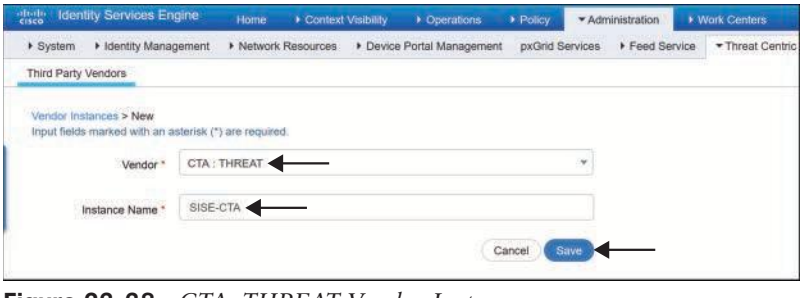


Figure 23-38 *CTA: THREAT Vendor Instance*

When the instance is ready, click Ready to Configure, as shown in Figure 23-39.

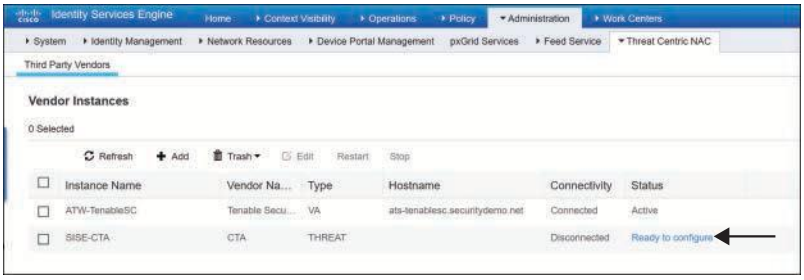


Figure 23-39 *Clicking Ready to Configure*

As shown in Figure 23-40, fill in the CTA STIX/TAXII fields with the strings copied from the Cognitive Intelligence UI in Figure 23-36.

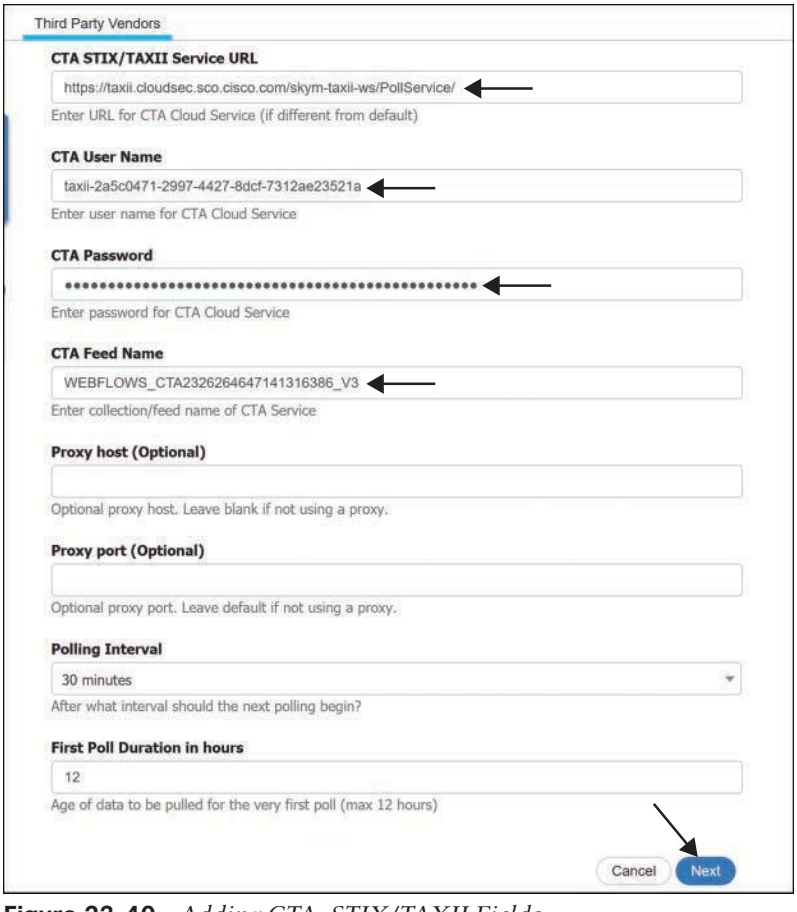


Figure 23-40 *Adding CTA: STIX/TAXII Fields*

Click Next and then Finish, as shown in Figure 23-41.

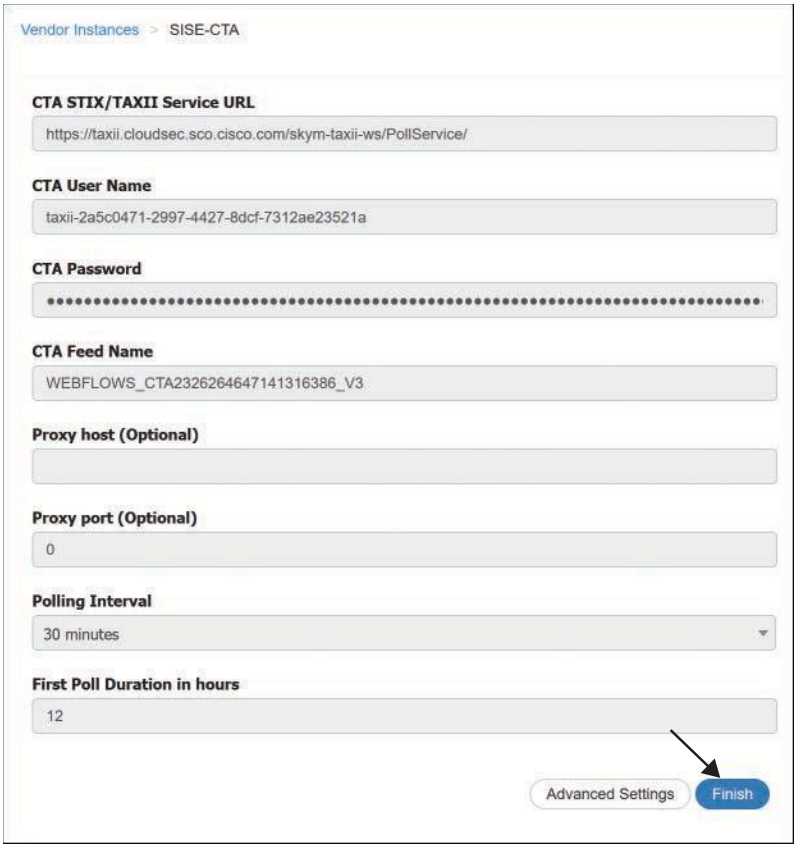


Figure 23-41 *Adding CTA: Finish*

Using CTA with Authorization

With every finding, CTA assigns a course of action. The course of action is a suggestion from the threat intelligence source about what to do with the endpoint that was identified as having a risk associated with it. Figures 23-42 and 23-43 show the three course of action options in CTA and how they can be used as authorization conditions.

The idea behind course of action is to alert the InfoSec team about the recommended remediation based on the severity of the compromise:

- **Eradication:** The compromise is so bad that the endpoint needs to be completely wiped and installed clean.
- **Internal Blocking:** The endpoint should be isolated, or quarantined, until it can be remediated.
- **Monitoring:** It is necessary to increase the inspection on this endpoint, as it is acting suspicious. This might trigger a vulnerability scan as well.

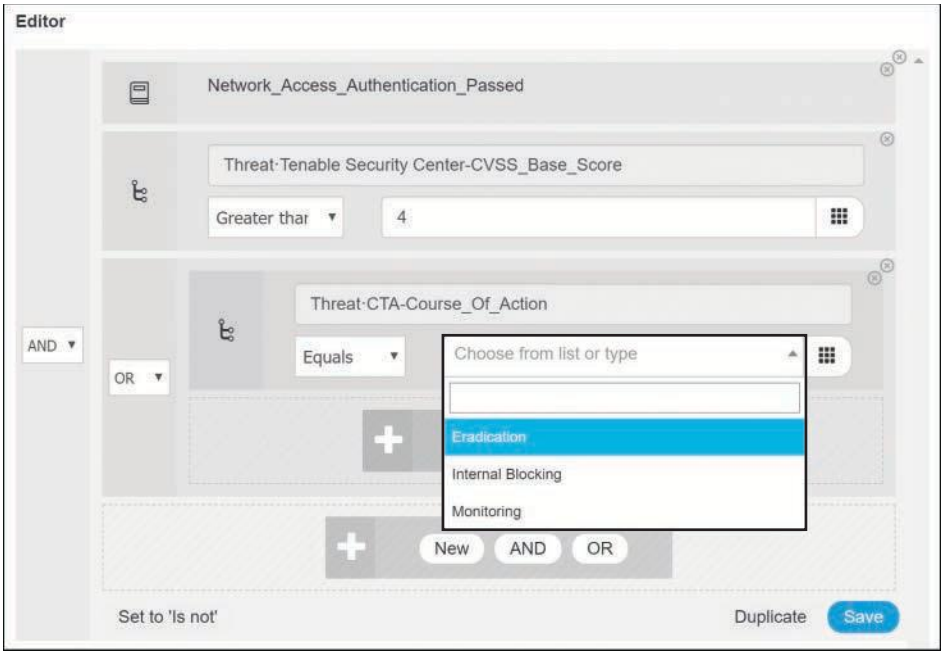


Figure 23-42 *CTA Course of Action Choices*

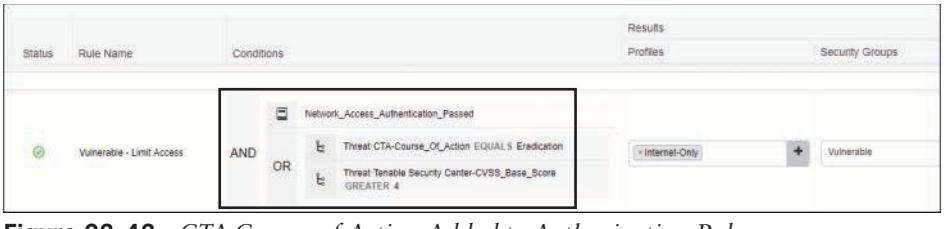


Figure 23-43 *CTA Course of Action Added to Authorization Rule*

As you know, the authorization rules in ISE can be very flexible and powerful. You can mix many different conditions to elicit a specific response. Using the CTA monitoring course of action is a perfect example of this flexibility. An authorization rule can be created in which you tie together adaptive network control policies with CTA course of action recommendations and vulnerability CVSS scores. When all three of these exist, you might take a strong action, but when just one exists, you might take nearly any action. You could assign a specific SGT that causes traffic to route differently, through packet-capturing tools, and crack open SSL traffic at the web proxy, for example.

AMP for Endpoints

Cisco Advanced Malware Protection (AMP) for Endpoints provides cloud-delivered next-generation antivirus, an endpoint protection platform (EPP), and advanced endpoint detection and response (EDR) capabilities. AMP has visibility into every file that is moved, copied, or executed. It has the ability to track the destination of each process that makes a network connection, the scripts that run on a system, and the contents of a system's memory.

In other words, AMP for Endpoints easily detects and prevents malicious behavior on an endpoint and can communicate that information to ISE through TC-NAC.

Key Topic

TC-NAC is supposed to be all about protecting a network from vulnerable or compromised endpoints. You've seen how CVSS scores and courses of action can be used in authorization rules. However, AMP is different. As of ISE Version 2.7, there is still no way to use the threat intelligence from AMP in an authorization policy.

NOTE If you are ever in the mood for a good laugh, watch the video-on-demand of Aaron Woland presenting on this topic at Cisco Live by going to www.ciscolive.com and clicking on the On-Demand Library. Search for "Woland Advanced Security Integration" and choose the session from Cisco Live! Orlando 2018 that is shown in Figure 23-44.

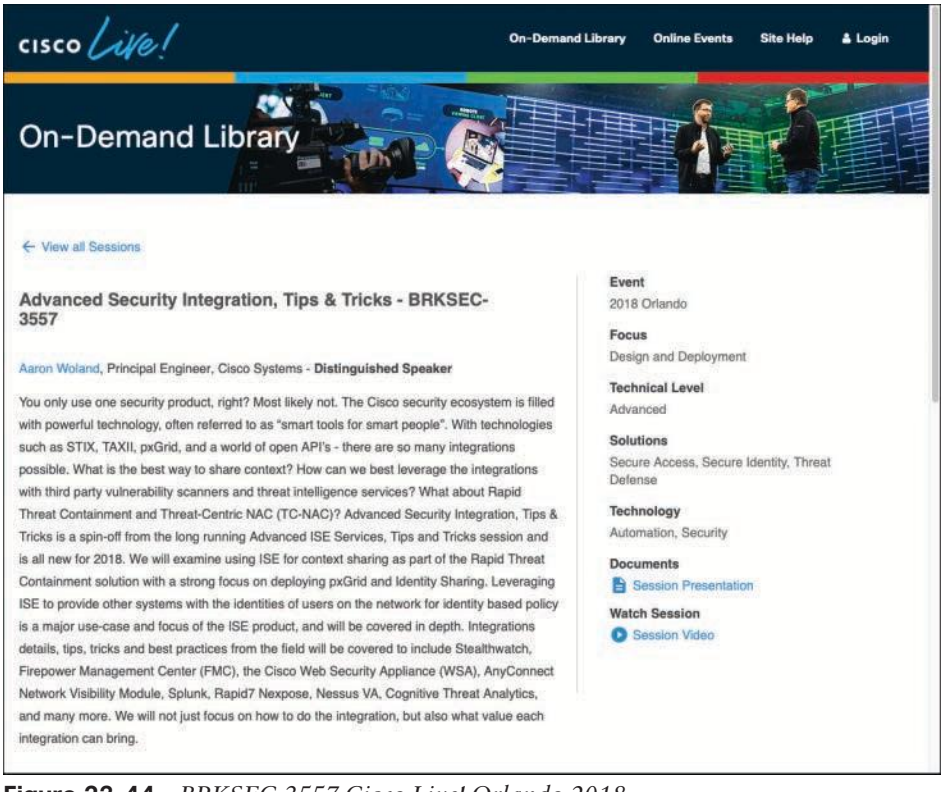


Figure 23-44 *BRKSEC-3557 Cisco Live! Orlando 2018*

Even though there is no automated action when integrating AMP with ISE via TC-NAC, there is still great value in doing so. The endpoint data in the Context Visibility store is updated with compromise data, and ISE administrators can send CoAs right from the Context Visibility screens.

Normalized Events

The threat intelligence from AMP is organized into indicators and incidents:

- An ***indicator of compromise (IoC)*** shows that there is an artifact or series of artifacts discovered on an endpoint indicating a breach with high confidence.
- An ***AMP incident*** is an event such as a detected threat or a vulnerable application.

Indicators

An IoC could be viewed as a type of signature. There is even an open source framework called OpenIOC for IoCs. This framework leverages an XML-based format for exchanging threat intelligence in a machine-readable way.

IoCs can be used like signatures to compare an endpoint against a signature and look for matches. AMP uses IoCs in this way and can trigger alarms that classify an endpoint as being in a compromised state. This classification is conveyed to ISE and added to the endpoint record within the Context Visibility store.

AMP assigns indicators to the likely impact levels low, medium, and high. This normalization helps ISE group the indicators and allows an administrator to take appropriate actions.

Incidents

An incident can be described as one of the many AMP events that can occur in the AMP console. Example 23-4 shows just a sampling of the events that exist in the AMP console.

NOTE You wouldn't want to expose each and every possible event to ISE, as that would be too much data.

Example 23-4 *Sampling of Events from the AMP for Endpoints Console*

```
| 1090519081 | Rootkit Detection                          |
| 1090519084 | DFC Threat Detected                        |
| 1090519089 | Endpoint IOC Scan Detection Summary        |
| 1090524040 | APK Threat Detected                        |
| 1090524041 | APK Custom Threat Detected                 |
| 1091567628 | Scan Completed With Detections             |
| 1107296257 | Multiple Infected Files                    |
| 1107296258 | Potential Dropper Infection                |
| 1107296260 | Java compromise                            |
| 1107296261 | Adobe Reader compromise                    |
| 1107296262 | Microsoft Word compromise                  |
| 1107296263 | Microsoft Excel compromise                 |
| 1107296264 | Microsoft PowerPoint compromise            |
| 1107296265 | Java launched a shell                      |
| 1107296266 | Adobe Reader launched a shell              |
| 1107296267 | Microsoft Word launched a shell            |
| 1107296268 | Microsoft Excel launched a shell           |
| 1107296269 | Microsoft PowerPoint launched a shell      |
```

```
| 1107296270 | Apple QuickTime compromise                    |
| 1107296271 | Apple QuickTime launched a shell              |
| 1107296273 | Suspected botnet connection                   |
| 1107296275 | Microsoft Calculator compromise               |
| 1107296276 | Microsoft Notepad compromise                  |
| 1107296277 | Connection to suspicious domain               |
| 1107296280 | Suspicious Download                           |
| 1107296281 | Microsoft CHM Compromise                      |
| 1107296282 | Suspicious Cscript Launch                     |
| 1107296283 | Possible Webshell                             |
| 1107296284 | Potential Ransomware                          |

| 1090519054 | Threat Detected                               |
| 1107296279 | Vulnerable Application Detected
```

AMP classifies incidents into five different impact levels: insignificant, distracting, painful, damaging, and catastrophic. This normalization helps ISE group the incidents and allows an administrator to take appropriate actions.

Configuring the AMP Adapter

Just as with Tenable and CTA, to add integration with AMP, you start by navigating to Administration > Threat Centric NAC > Third Party Vendors and click Add. Then you select AMP:THREAT from the Vendor dropdown, provide a name for the instance, and click Save, as shown in Figure 23-45.

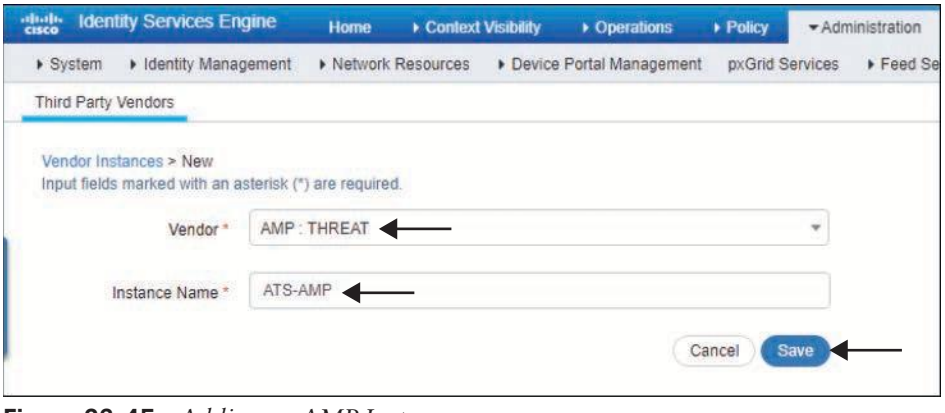


Figure 23-45 *Adding an AMP Instance*

Click Ready to Configure to begin the configuration process. In the first configuration screen, shown in Figure 23-46, you enter optional Socks proxy information. If your ISE nodes require a proxy server to communicate with the Internet, you need to add the information here.

NOTE It must be a Socks proxy; a standard HTTP proxy does not work for the AMP connector.

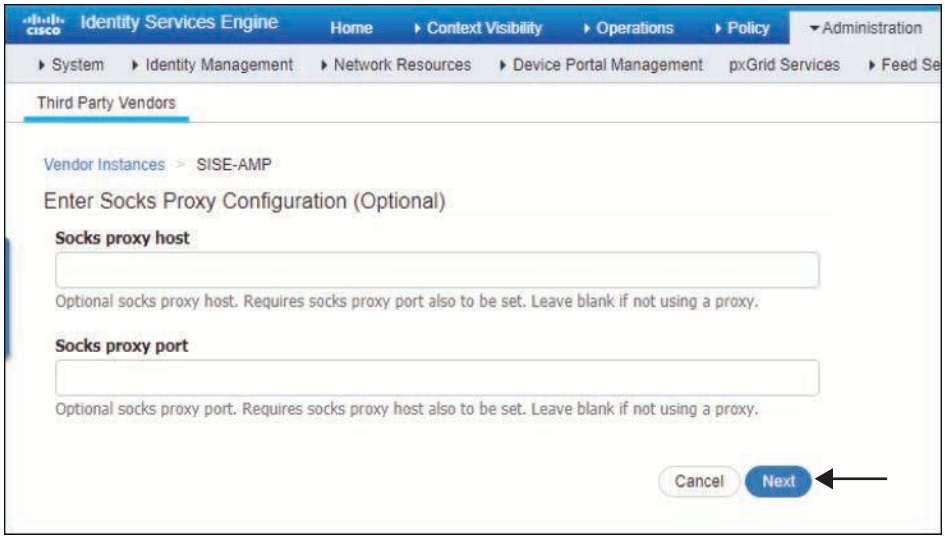


Figure 23-46 *Socks Proxy Configuration Screen*

If your ISE nodes require a proxy to reach the Internet, enter the information here and click Next. Otherwise, just click Next to continue the configuration.

Cisco's AMP solution has three regional clouds, and the next configuration screen allows you to select the AMP cloud instance for your business and click Next, as shown in Figure 23-47.

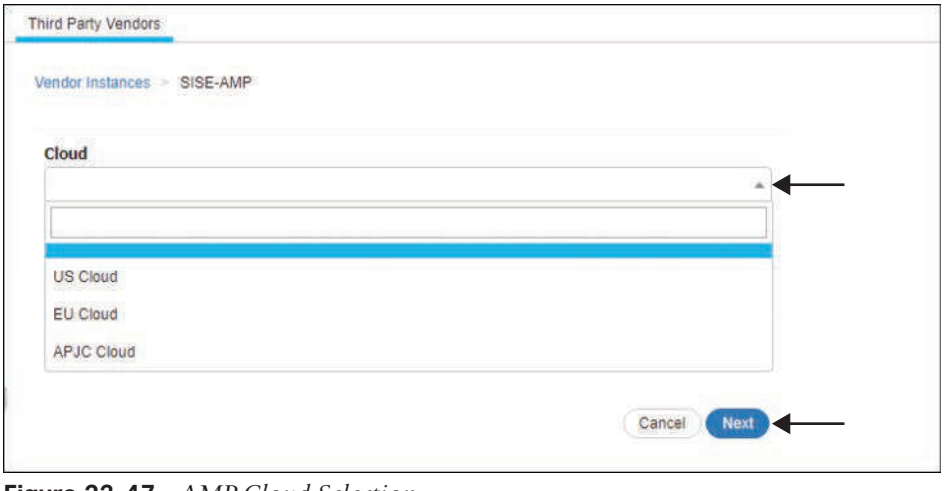


Figure 23-47 *AMP Cloud Selection*

NOTE As of ISE Version 2.7, only the AMP public cloud is supported. A customer who has deployed the on-premises version, known as AMP Private Cloud, cannot, unfortunately, integrate it with ISE.

The next step is to authenticate to the AMP service, which uses Open Authentication (OAuth) to generate a token for further access between the web services. Figure 23-48 shows the link you click for the authentication. When you click this link, you are redirected to the login screen for the AMP console, shown in Figure 23-49.

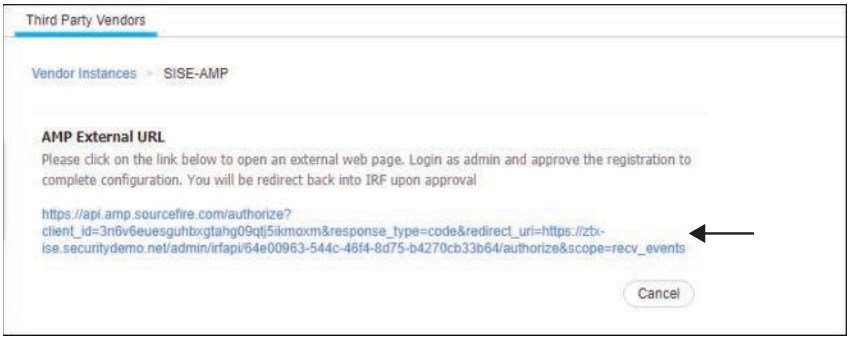


Figure 23-48 *Clicking the Link to Authenticate to the AMP Console*

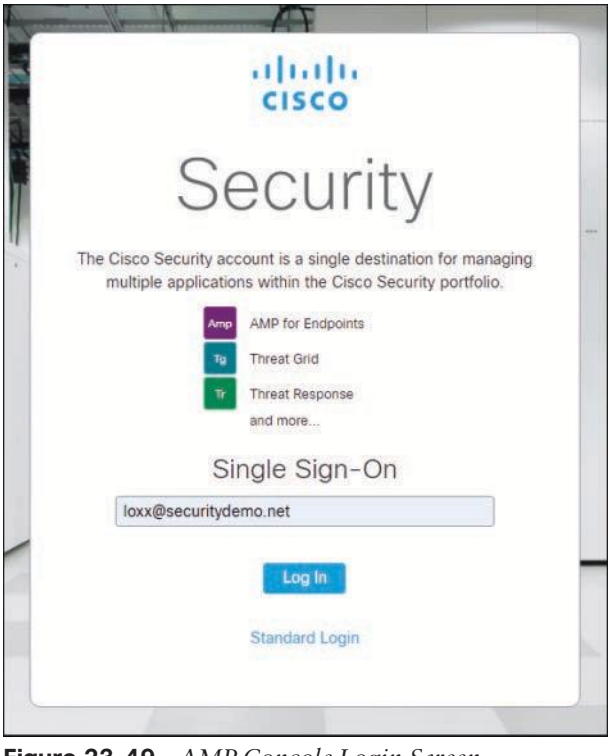


Figure 23-49 *AMP Console Login Screen*

When you are logged in to the AMP console, you can authorize the specific groups for which ISE will receive events and click Allow, as shown in Figure 23-50. You are sent back to the ISE configuration screen, where you can click Advanced Settings and select the individual events that ISE should subscribe to, as shown in the partial list in Figure 23-51.

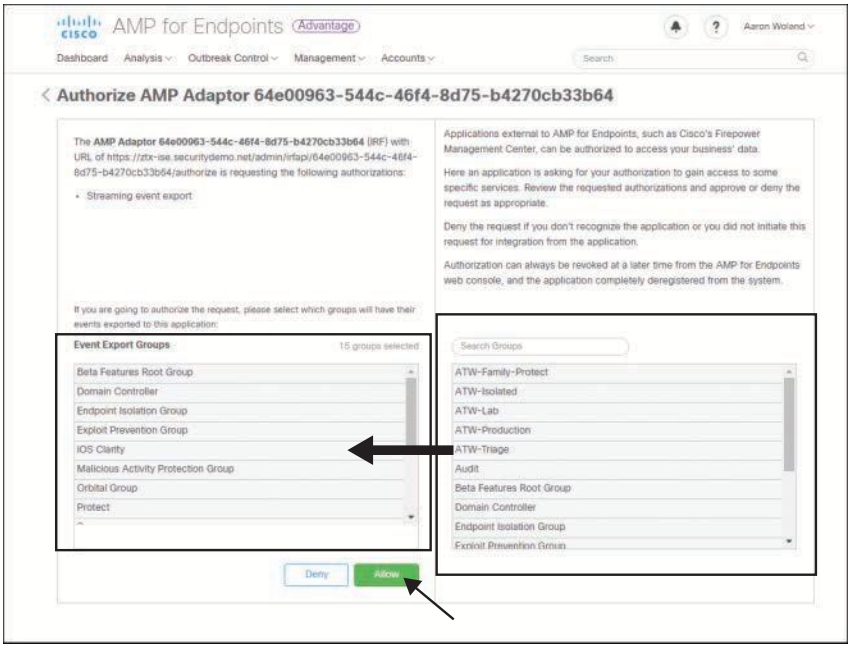


Figure 23-50 *Selecting the Group Events to Authorize*

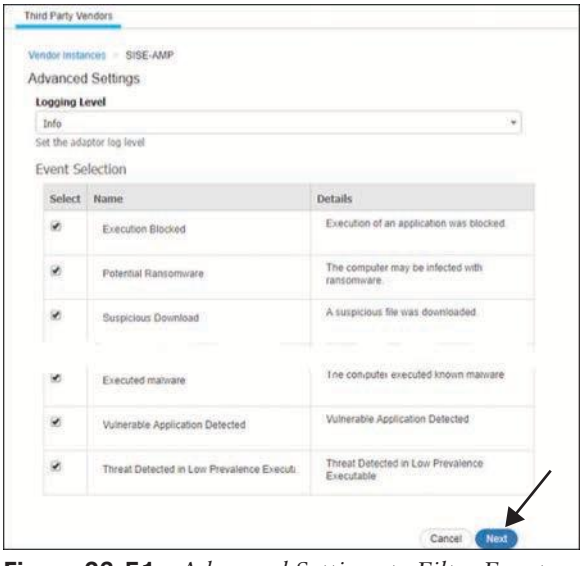


Figure 23-51 *Advanced Settings to Filter Events*

Figure 23-52 shows the Compromised Endpoints page in the Context Visibility store, and Figure 23-53 shows how an administrator can change the authorization of an endpoint in the list.

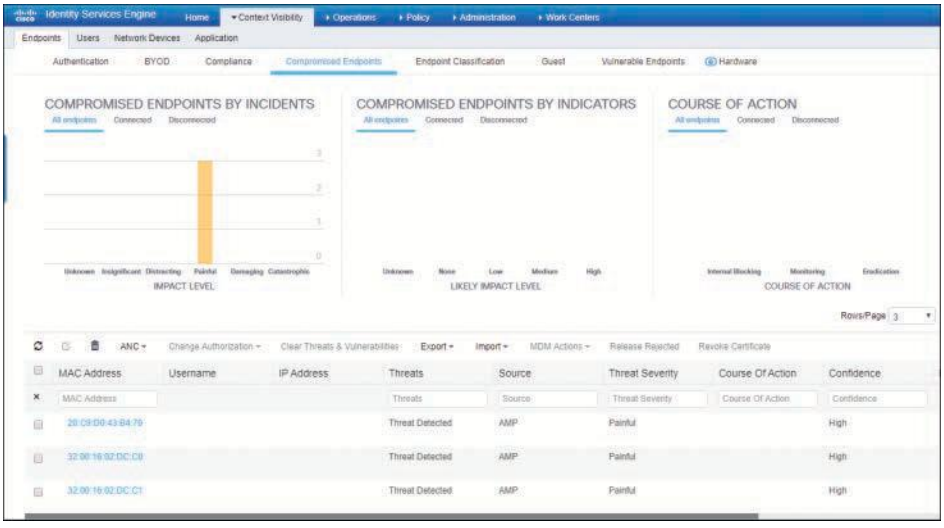


Figure 23-52 *Context Visibility > Endpoints > Compromised Endpoints*

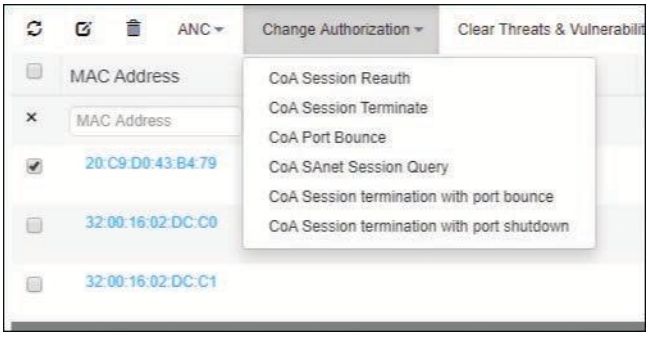


Figure 23-53 *Changing Authorization of Compromised Endpoints*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 23-2 lists these key topics and the page number on which each is found.

Key Topic

Table 23-2 Key Topics

Key Topic Element	Description	Page Number
Paragraph	TC-NAC and policy creation around threats and vulnerabilities	871
Paragraph	Tenable, Qualys, and Rapid7	872
Paragraph	The CoA report and TC-NAC Live Log	889
Paragraph	ISE integration with AMP and Cognitive for threat sources	890
Paragraph	AMP threat intelligence and authorization policies	898

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

vulnerability, exploit, Common Vulnerabilities and Exposures (CVE), Common Vulnerability Scoring System (CVSS), cyber threat intelligence (CTI), Structured Threat Information eXpression (STIX), Trusted Automated eXchange of Intelligence Information (TAXII), indicator of compromise (IoC), AMP incident

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What protocols are used between ISE and CTA?
2. What is the difference between an incident and an indicator in TC-NAC?
3. Can TC-NAC be applied before providing full access to an endpoint or after?
4. What is required for an entry to be added to the TC-NAC Live Log?
5. What is the difference between a vulnerability and an exploit?



CHAPTER 24

Device Administration AAA with ISE

This chapter covers the following topics:

- Configure personas
- Describe deployment options
- Describe identity store options
- Configure network access devices
- Configure policies including authentication and authorization profiles
- Compare AAA protocols
- Configure TACACS+ device administration and command authorization

Cisco ISE is a policy server. It exists to provide authentication, authorization, and accounting (AAA) for purposes of access control.

With ISE Version 1.0, there was only one AAA protocol that could be used with ISE: RADIUS. As discussed in Chapter 1, "Fundamentals of AAA," RADIUS is the primary AAA protocol used for controlling network access. ISE Version 1.x was designed as a next-generation network access control solution—period. Does this mean that an implementation could not overload the original policy sets with one created for the purposes of controlling administrative access into devices such as routers, switches, security appliances, or even Linux servers? Of course not, and the authors of this book are well aware of many implementations of ISE where RADIUS is being used for exactly that purpose.

In ISE Version 2.0, a very big addition to the product was the inclusion of the *Device Administration work center*, which assumes the use of the TACACS+ AAA protocol but creates policy sets designed for device administration and not network access.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 24-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 24-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Device Administration AAA Refresher	1
Device Administration in ISE	2, 4
Device Administration Global Settings	5
Device Administration Work Center	3, 6–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What makes TACACS+ the preferred AAA protocol for device administration?
   a. The TACACS+ protocol has built-in communication and devAdmin attribute/value pairs for device administration.
   b. TACACS+ is able to authorize many times from a single authentication.
   c. TACACS+ performs authentication and authorization in a single autonomous connection.
   d. TACACS+ was written by Cisco and therefore is the best protocol.
2. What is needed to enable the Device Administration work center in the ISE UI?
   a. You need one device administration license per PSN.
   b. Nothing. The Device Administration work center is always enabled, but the PSNs cannot use TACACS+ without a license.
   c. Nothing. The Device Administration work center is always enabled, as it is included with a base license.
   d. You need one or more device administration licenses.
3. What is the difference between the network devices listed in the Network Access work center and the network devices listed in the Device Administration work center?
   a. Network resources for device administration must use TACACS+, and that configuration is only available through the Device Administration work center.
   b. There are no network resources for device administration; there are only network resources for network access.
   c. Nothing at all. They are the same devices, and these two work centers display the same pages.
   d. Network resources for device administration must use RADIUS, and that configuration is only available through the Device Administration work center.

4. Which of the following is true when designing device administration for a large enterprise?
   a. For very large deployments, you should maintain a separate ISE cube for RADIUS (network access) and TACACS (device administration) functions.
   b. There is no difference when it comes to the size of an ISE deployment.
   c. TACACS must be enabled on the PAN as well as on the PSNs.
   d. You should design for a single ISE cube but ensure that there are separate MnT nodes for device administration.
5. From what page in the ISE UI can you configure TACACS to be used on a nonstandard TCP port?
   a. There is no page for this in the UI; the port can only be changed from the CLI.
   b. There is no way to change the port that TACACS+ is using. It must be TCP port 47.
   c. There is no way to change the port that TACACS+ is using. It must be TCP port 49.
   d. Navigate to Work Centers > Device Administration > Deployment.
6. Where do you configure the strings used in the password change prompts?
   a. There is no way to change passwords in an TACACS+ session. That is a RADIUS capability only.
   b. Navigate to Work Centers > Device Administration > Global Settings > Password Change Control.
   c. Passwords can be changed within a TACACS+ session, but there is no way to change the default strings used for the prompts.
   d. Navigate to Administration > Settings > TACACS > Password Change Control.
7. What policy element is used to dictate which commands a network administrator is permitted to use and which ones are restricted?
   a. TACACS command sets are used for this.
   b. It is not possible to restrict the commands from ISE. Control of the commands is strictly configured at the network device itself, through the use of privilege-level assignment.
   c. Device administration command sets are used for this.
   d. It is not possible to restrict the commands from ISE. Control of the commands is strictly configured at the network device itself, through the use of TrustSec role assignment.
8. Where does an ISE administrator configure the policy to deny device administration access to any user who has not first passed the network access authentication?
   a. Configure the policy at Work Center > Device Administration > Connection Settings.
   b. Configure the policy at Administration > Settings > TACACS.
   c. This configuration is available only through pxGrid.
   d. This is not possible today. The network access and device administration functions of ISE are completely separate.

9. Which of the following protocols can be used for device administration with ISE?
   a. RADIUS and TACACS+
   b. TACACS+ only
   c. TACACS+ and DIAMETER
   d. RADIUS and SAML
10. How can you configure different policy sets for IOS devices and WLCs?
   a. Policy sets are not used with device administration; they are used only with network access.
   b. The device type is conveyed in the TACACS+ communication, which allows ISE to select the correct policy set.
   c. Use network device groups (NDGs) as the condition for the policy set.
   d. Use vendor-specific attributes (VSAs) to determine the device type.

Foundation Topics

Device Administration AAA Refresher

In Chapter 1, you were introduced to the concepts of network access AAA and device administration AAA. While both are focused on identifying who is allowed to perform what action, they are vastly different in terms of purpose.

NOTE The purposes of network access AAA and device administration AAA are so different, in fact, that Aaron Woland once published a blog post on Network World about why both types shouldn't be in ISE. He contended that ISE should stick to network access AAA and leave Cisco's Access Control Server (ACS) for device administration AAA. You can read that opinion piece at http://www.networkworld.com/article/2838882/radius-versus-tacacs.html.

Key Topic

Device administration AAA exists to control access to network device consoles, Telnet sessions, Secure Shell (SSH) sessions, and other methods of gaining administrative control over a network device. The purpose is not only to control the access but to provide centralized visibility of all those actions. For example, if an administrator makes a change that causes a network outage, the organization needs an audit trail as proof of who caused it and what the person did.

Figure 24-1 illustrates device administration graphically.

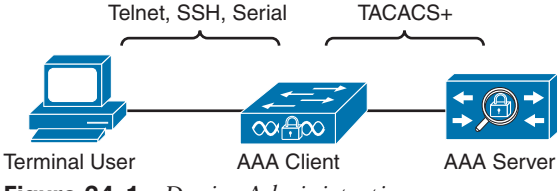


Figure 24-1 *Device Administration*

Device administration AAA is very interactive in nature—or at least it can be. Policies can be written so that a user is authorized a single time for an entire CLI session, or they can

be configured to authorize each and every command entered. Due to the extremely interactive nature of command authorization, TACACS+ is a perfect AAA protocol for device administration.

Figure 24-2 illustrates how TACACS+ can perform authentication once and authorize many times. This is key for device administration. Think about it: You can authenticate a user, then authorize the user to access the CLI of a router, and you may still have to authorize each of the commands entered from that CLI for granular control.

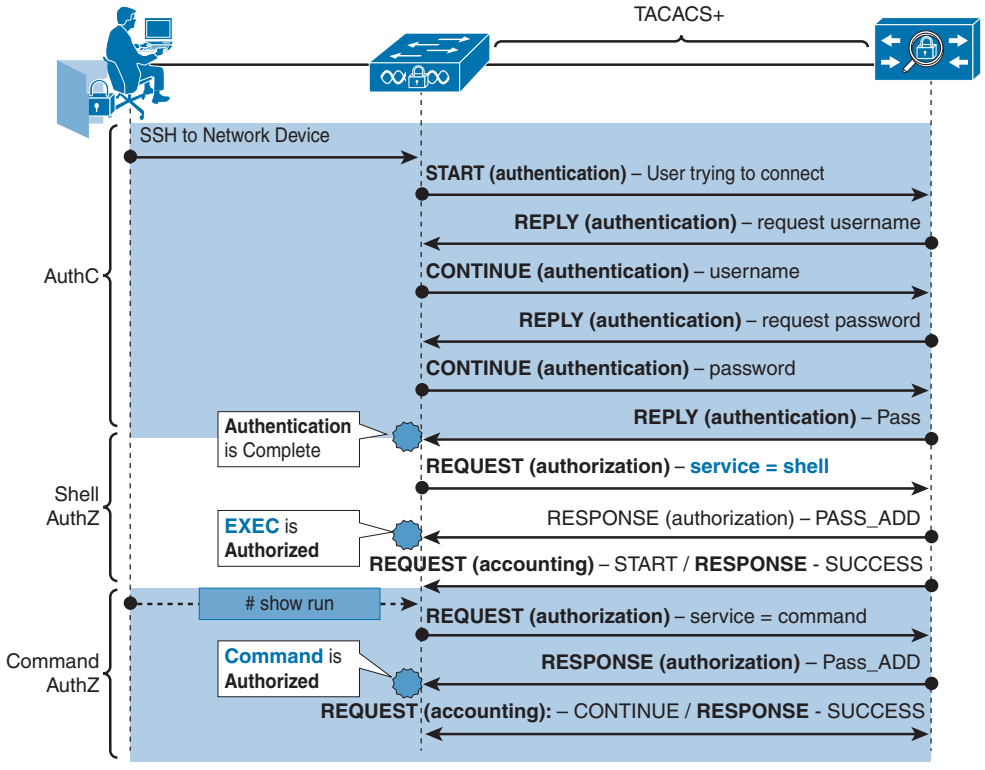


Figure 24-2 *Authenticate Once, Authorize Many*

NOTE Review Chapter 1 for more on what device administration AAA is and for a detailed look at TACACS+.

Device Administration in ISE

When it comes to ISE, device administration is synonymous with TACACS+. Everything to do with TACACS+ exists within the Device Administration work center.

In fact, to enable TACACS+, you need a single license, named Device Admin, as shown in Figure 24-3. Unlike the Base, Plus, and Apex licenses, the Device Admin license is based on the number of connections. It is a single license that is counted per PSN with TACACS+ enabled. This differs a bit from ISE's predecessor, Cisco Access Control Server (ACS), which had a base license and a large deployment license.

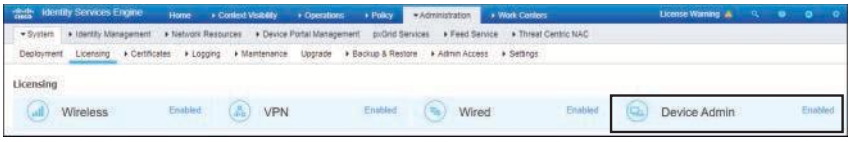


Figure 24-3 *Administration > System > Licensing*

Key Topic

The Device Administration work center is displayed only if you have at least one Device Admin license.

Device Administration Design

There are a few schools of thought when it comes to designing ISE for device administration:

- Separate ISE cubes, with one for RADIUS and another for TACACS+. This option provides dedicated PSNs and a dedicated MnT node for each function.
- A mixed ISE cube with separate PSNs. This option provides dedicated PSNs for each AAA protocol, but the cube shares a single MnT node for logging and reporting.
- A mixed ISE cube where the PSNs are not dedicated to any protocol.

NOTE Aaron Woland has publicly stated his belief that device administration AAA has no business being in a network access AAA product. There is too much disparity in the uses of the two AAA types, their traffic patterns, and the load level they put on the AAA and monitoring servers. However, Aaron lost that fight. Device administration AAA is squarely embedded in ISE, and ACS has been retired permanently.

Figure 24-4 shows a table created by Douglas Gash, the father of TACACS+ at Cisco. Doug is a brilliant Englishman who is the chief architect of TACACS+ for ACS and ISE at Cisco, as well as a major contributor to the TACACS+ standard in the IETF.

The table in Figure 24-4 helps you decide which deployment model to follow, based on what is most important to your organization. Is it scale of TACACS+ and RADIUS services? Are the audit trail and log retention most important? Perhaps you need something for the political separation of duties in your organization?

Priorities - According to Policy		**Separate Deployment** TACACS RADIUS	**Separate PSN Mode** RADIUS TACACS	**Mixed PSN Mode** RADIUS/TACACS
Separation of Configuration	Yes: Specialization for TACACS+	✓		
	No: Avoid Duplication of Shared Items Avoid cost of duplicate PAN/PSN		✓	✓
Separation of Logging Store	Yes: Optimize Log Retention VM	✓		
	No: Centralized Monitoring		✓	✓
Independent Scaling of Services	Yes: Scale as Needed Avoid NAC/Device Admin Load	✓	✓	
	No: Avoid underutilized PSNs			✓

Figure 24-4 *Doug Gash's Options for Deploying Device Administration*

For simplicity, in the following sections we separate the deployment recommendations into categories based on deployment size: large, medium, and small.

Large Deployments

For a large deployment, it is best to have separate ISE cubes for network access and device administration. This way, you are assured that one will never negatively impact the other. After all, you would not want to accidently prevent your CEO from getting on the wireless network just because a script is actively making a lot of changes across your network infrastructure. You also would not want the audit logs for network device changes to be kept for less time because the MnT disk space is shared between network access and device administration.

The MnT node plays a key role in network access functions beyond just logging and reporting. As of ISE Version 2.2, the MnT node is still the holder of critical functions such as the centralized session directory, the pxGrid topic publishing for session data, and merging of passive ID and active authentication data into the session directory.

If the MnT node must also process tremendous amounts of logging for all the command authorization accounting, logging entries for every single command entered on every single network device in the entire organization…well, you can see why this becomes a point of concern.

Figure 24-5 illustrates two different ISE cubes. Cube 1 is dedicated to TACACS+, and Cube 2 is dedicated to RADIUS. Therefore, the PAN and MnT nodes are also dedicated—not just the PSNs. While this figure does show virtual IP (VIP) addresses for the PSNs, a load balancer is not required. The setup is only illustrated this way to show a common practice.

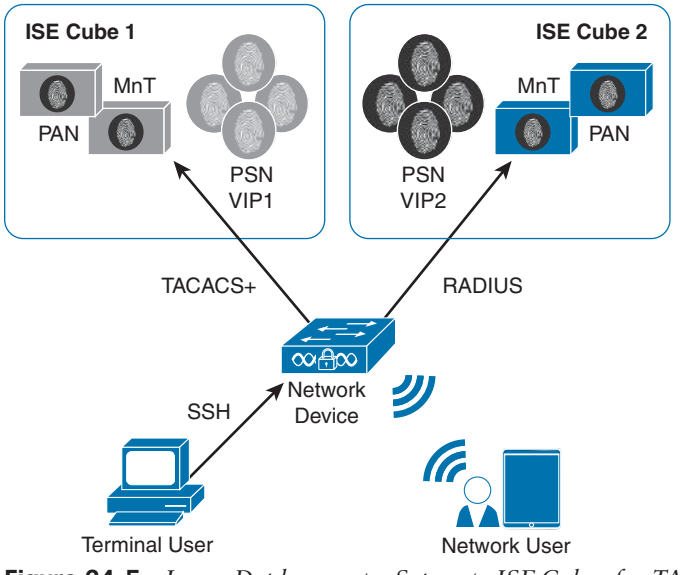


Figure 24-5 *Large Deployments: Separate ISE Cubes for TACACS+ and RADIUS*

Medium Deployments

With medium-size deployments, you might want to have a single ISE cube, but it is best to have separate PSNs for network access and device administration. This way, you are assured that one will never affect the other, and you will not have to maintain separate cubes.

A set of PSNs would primarily be responsible for all the RADIUS traffic, while another set of PSNs would have primary responsibility for the TACACS+ traffic. For redundancy, you might choose to send the RADIUS traffic to the TACACS+ PSNs, but only in the case of a disaster. The PSNs are still primarily dedicated to either RADIUS or TACACS+.

Figure 24-6 illustrates a single ISE cube with dedicated PSNs per function. While the illustration does show VIP addresses for the PSNs, a load balancer is not required. The setup is only illustrated this way to show a common practice.

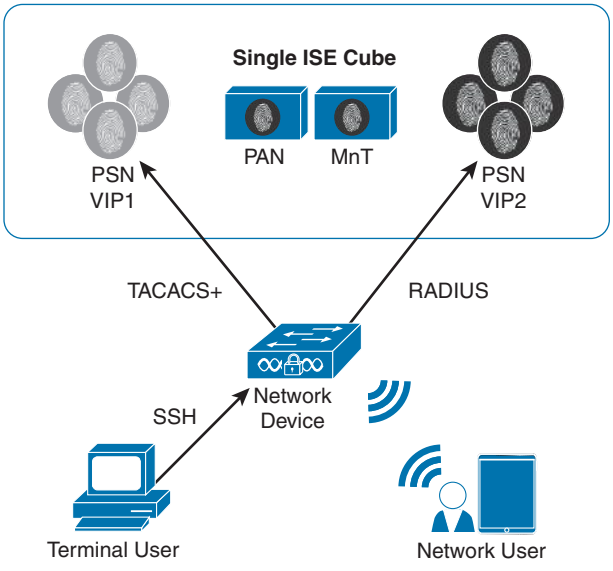


Figure 24-6 *Medium Deployments: One ISE Cube, Separate PSNs for TACACS+ and RADIUS*

Small Deployments

A small deployment could certainly be just a single ISE cube. It might even be a standalone ISE node that is performing all functions. In such instances, there are no dedicated nodes at all; all PSNs handle equal amounts of TACACS+ and RADIUS traffic.

Figure 24-7 illustrates a single ISE cube with dedicated PSNs per function. While the illustration does show VIP addresses for the PSNs, a load balancer is not required. The setup is only illustrated this way to show a common practice.

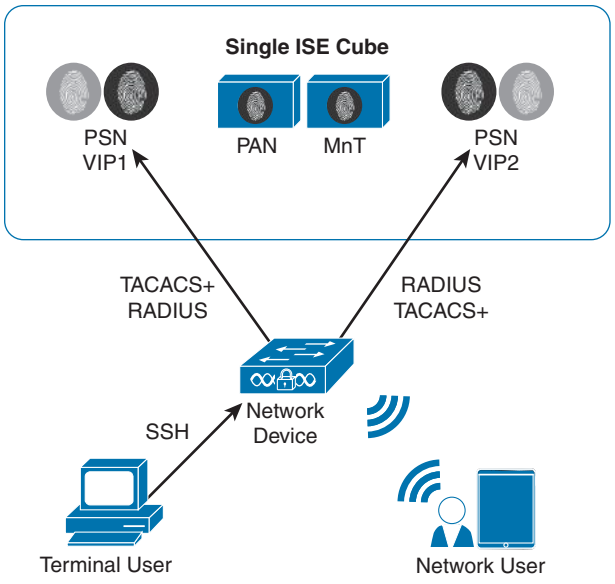


Figure 24-7 *Small Deployments: One ISE Cube, No Separation of TACACS+ and RADIUS*

Enabling TACACS+ in ISE

Installing the Device Admin license is not enough to start accepting TACACS+ communications. You also must enable it in the deployment settings.

It is important to know your intended design before you enable the TACACS+ functionality, and you should only enable the Device Admin service on the PSNs that will handle TACACS+ and leave it disabled on any PSNs that are supposed to be dedicated for RADIUS (and vice versa). You should keep the remainder of the session services disabled on the dedicated TACACS+ PSNs.

You can use the standard system deployment screen located at Administration > System > Deployment, or you can use the preferred method: navigate to Work Centers > Device Administration > Overview > Deployment.

When you use the standard deployment screen, device administration must be enabled per PSN, as shown in Figure 24-8.

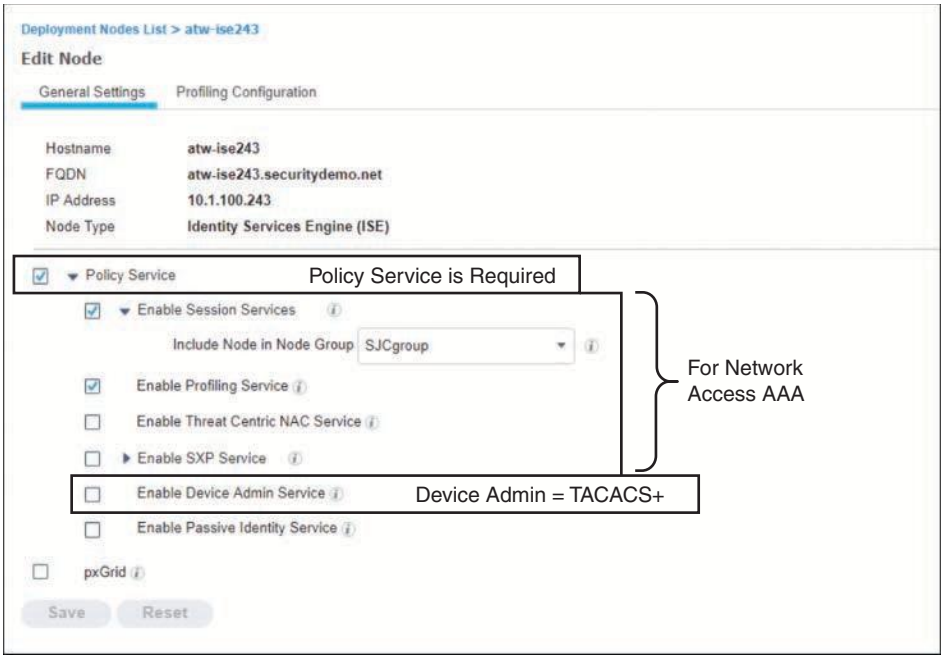


Figure 24-8 *Deployment Screen*

A better way to access the deployment settings is to navigate to Work Centers > Device Administration > Overview > Deployment, as shown in Figure 24-9. This work center allows you to activate device administration on all PSNs, deactivate it on all PSNs, or select individual PSNs. Figure 24-10 shows device administration being enabled on specific PSNs only. The deployment screen in the work center also permits an ISE administrator to modify the TCP port that TACACS+ will use; 49 is the default port.

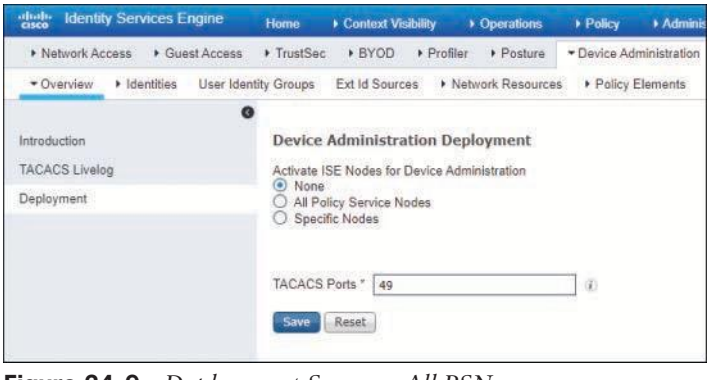


Figure 24-9 *Deployment Screen – All PSNs*



Figure 24-10 *Deployment Screen – Specific PSNs*

Network Devices

We could discuss ISE cubes and TACACS+ versus RADIUS PSNs all day long. However, none of it will matter if a network access device (NAD) is not set up correctly. A NAD object in ISE must be configured with the TACACS+ shared secret and the connection type.

The same network device group (NDG) guidelines for network access apply for device administration. The more applicable the NDG design is for your organization, the more it aids in your policy creation. Device type, location, and line of business are all very useful types of NAD groupings.

Figure 24-11 shows the TACACS+ portion of the NAD definition in the ISE UI.

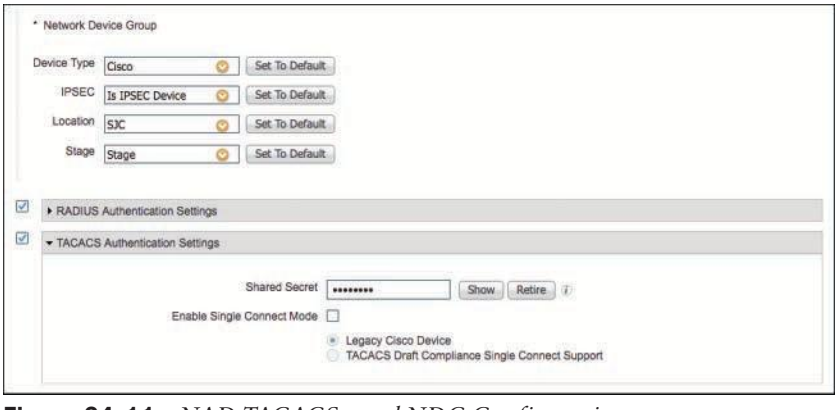


Figure 24-11 *NAD TACACS+ and NDG Configuration*

In Figure 24-11 you can see the NDG assignments, such as device type. In this instance, the device type assignment is leveraged to identify this switch as a Cisco access-layer switch,

and you effectively use that assignment to point the TACACS+ requests to a policy set that is created for IOS devices.

Also in Figure 24-11, you can see how the TACACS+ shared secret gets configured; the configuration is just the same as with RADIUS. The shared secret is the password used between the NAD and ISE to validate the source of the TACACS+ communication and to provide a seed value for the encryption.

There is also an option to enable single connect mode. The normal communication mode for TACACS+ is to open a new TCP session for each and every TACACS communication—that is, for each authentication request, for each authorization request, and for each accounting request. Enabling this setting allows the NAD to maintain a single open connection between itself and the TACACS+ service on ISE. This type of connection is efficient because it allows the service to handle a larger number of TACACS operations, but it does require the NAD to also be configured for single connect.

Next, you should notice a button labeled Retire. This allows you to retire the existing shared secret and configure a new one. During the retirement period, ISE accepts both the old and new shared secrets for the device, which gives you some buffer time and makes the updating of shared secrets more operationally feasible. The retirement time leverages the timer configured under Work Centers > Device Administration > Settings > Connection Settings, as shown in Figure 24-12.

Device Administration Global Settings

The retire option for shared secrets is just one of the many global settings for TACACS+ in ISE. To see the globally applicable settings, navigate to Work Centers > Device Administration > Settings, as shown in Figure 24-12.

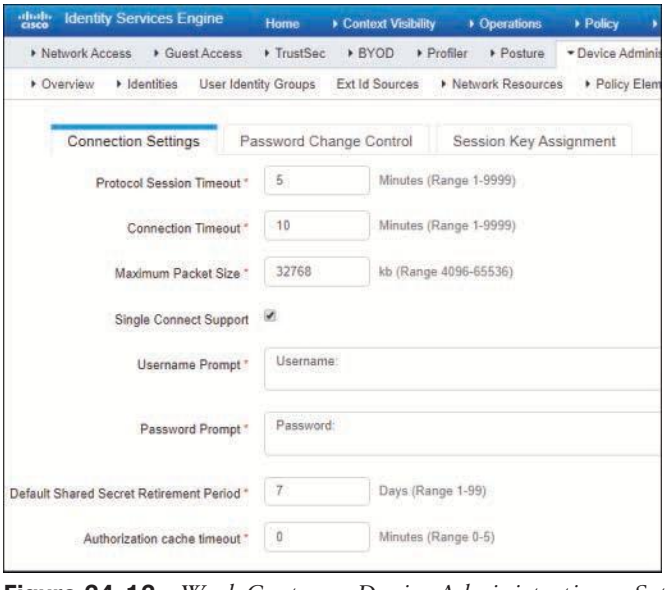


Figure 24-12 *Work Centers > Device Administration > Settings*

Connection Settings

Figure 24-12 shows that the default tab is the Connection Settings tab. The shared secret retirement period, which can range from 1 to 99 days, is configured on this tab. In addition to the retirement period, the connection and session timers are configured on this tab, as are the strings used for username and password prompts and whether single-connect mode should be enabled at all.

Password Change Control

If you navigate to Work Centers > Device Administration > Settings > Password Change Control, you can globally enable or disable the ability to change passwords through the TACACS+ session and the strings to be used in the prompts (see Figure 24-13).

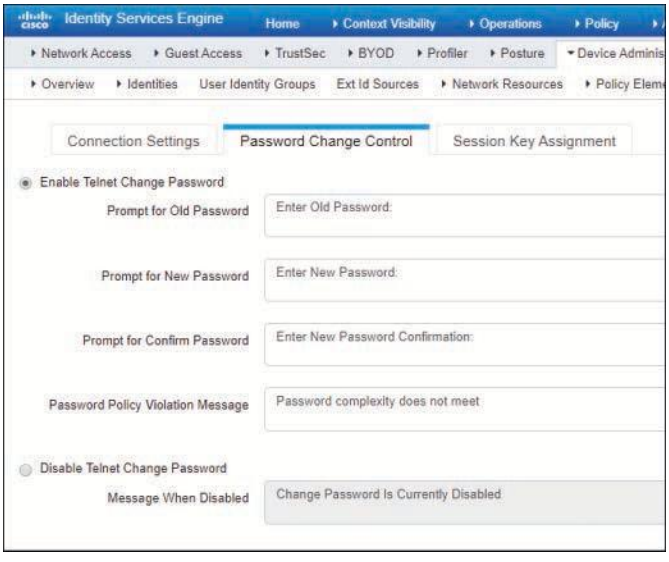


Figure 24-13 *Password Change Control Tab*

Session Key Assignment

The last of the global settings is found by navigating to Work Centers > Device Administration > Settings > Session Key Assignment (see Figure 24-14). On this tab you see the different values from the TACACS packets that are available for use in identifying when the packets are part of the same session. Because of these settings, ISE can track a full session throughout the single authentication and multiple authorization and accounting requests that may come in. There is rarely a need to change anything in this tab, unless you are directed by Cisco TAC to do so.

The real purpose of the Session Key Assignment tab is to help you determine the number of sessions when you configure a limit for the maximum sessions allowed.

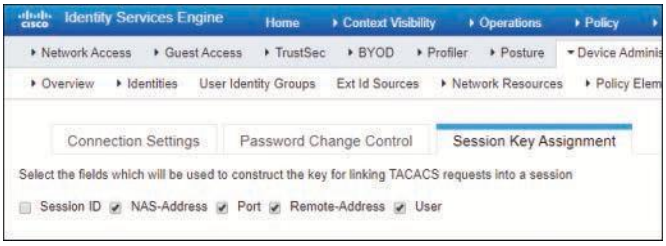


Figure 24-14 *Session Key Assignment Tab*

Device Administration Work Center

You have already been exposed to parts of the Device Administration work center. This section walks you through the remaining portions of this work center. As with the other work centers in ISE, the Device Administration work center is designed to provide you with access to just about everything you need to accomplish the administrative tasks associated with device administration AAA.

Figure 24-15 shows the ISE UI navigation for the Device Administration work center.

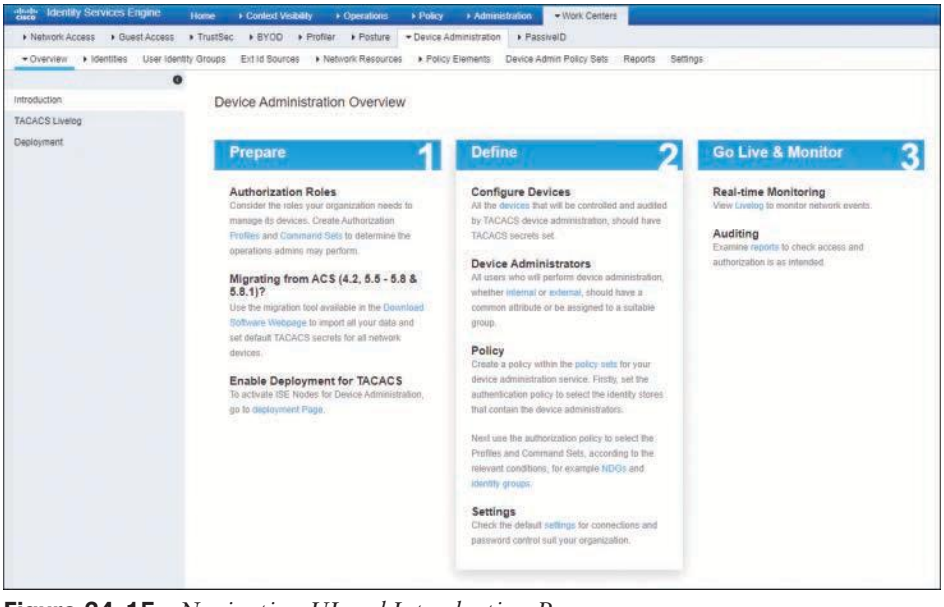


Figure 24-15 *Navigation UI and Introduction Page*

Working from the left to right, the first UI area is Overview. It has the standard work center introduction page, including a very broad review of the activities within the work center, as well as some helpful links. In this introduction page, there is a link to the ACS-to-ISE migration tool.

Next in the Overview section is TACACS Live Log, a dedicated near-real-time log screen for TACACS only. Much like the RADIUS Live Log, the TACACS Live Log is your go-to page for an overview of operations and for troubleshooting. We will come back to this page in the next few chapters.

The final part of the Overview section is a TACACS deployment page that is designed to allow you to quickly enable TACACS+ on any number of PSNs from a single screen instead of having to enable it one PSN at a time (see Figure 24-12).

Identities

The next three tabs in the Device Administration work center are related to user identities. These identities are not unique to the Device Administration work center at all. In fact, they are the same as the identities you find in the other work centers. The Identities tab houses the internal users created within ISE's internal user database. Figure 24-16 shows the Identities tab with a local user.

NOTE In Figure 24-16, you can see that when you select the Identities tab, the page that appears is labeled Network Access Users, but this is simply a legacy name, as these user accounts can be leveraged for device administration as well.

Remember that internal users can be configured to leverage internal passwords, or an internal user can be a local account with the same name as an AD user account that leverages the AD password. This is commonly the case with device administration as an additional authorization condition. In other words, you might not only leverage an account or group membership from AD but also rely on the need for the TACACS+ administrator to create that local account in ISE.



Figure 24-16 *Identities Tab*

The User Identity Groups tab of the Device Administration work center is a shortcut of sorts to the existing users, as shown in Figure 24-17. The Ext Id Sources tab is shown in Figure 24-18. These are exactly the same tabs as in the other areas of the UI, which makes the work centers easier to use and more complete.



Figure 24-17 *User Identity Groups Tab*

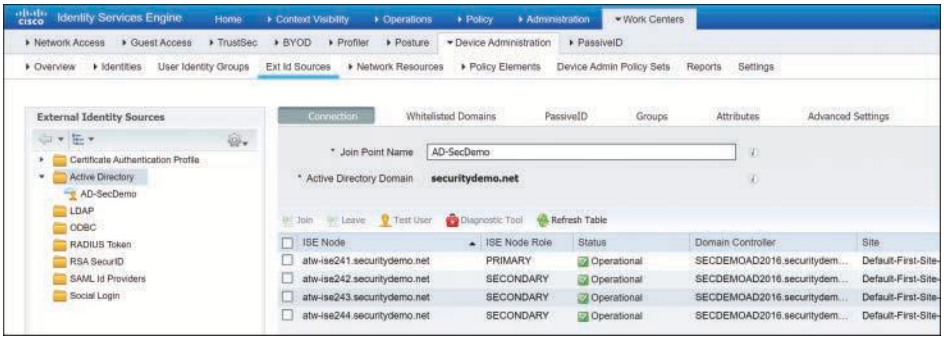


Figure 24-18 *Ext Id Sources Tab*

Network Resources

The Network Resources tab of the Device Administration work center is a shortcut to the same network resources available all throughout the UI and in other work centers. A very noticeable difference is the inclusion of TACACS External Servers and TACACS Server Sequence sections, as shown in Figure 24-19.

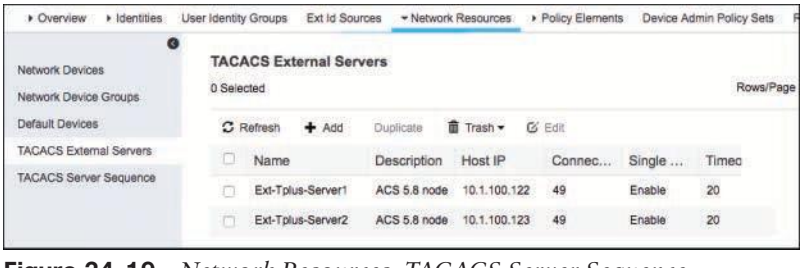


Figure 24-19 *Network Resources: TACACS Server Sequence*

Just as with external RADIUS servers and RADIUS server sequences, there are many reasons you might need to forward TACACS+ requests to other servers. It could be a requirement for proof of concept or a phase in migrating away from another solution. The external TACACS servers get added to one or more TACACS server sequences, which are then leveraged in a proxy sequence instead of a standard authentication being done within the policy set.

Policy Elements

The Policy Elements tab of the Device Administration work center is a shortcut to the conditions and results. The following sections focus on the results. The two main results that are used with device administration are TACACS command sets and TACACS profiles.

TACACS Command Sets

Key Topic

As you will learn in detail in Chapter 25, "Configuring Device Administration AAA with Cisco IOS," the purpose of command sets is to limit which commands a user can or cannot use. A command set is assigned to a session as one of the authorization results.

To work with the TACACS command sets, navigate to Work Centers > Device Administration > Policy Elements > Results > TACACS Command Sets, as shown in Figure 24-20.

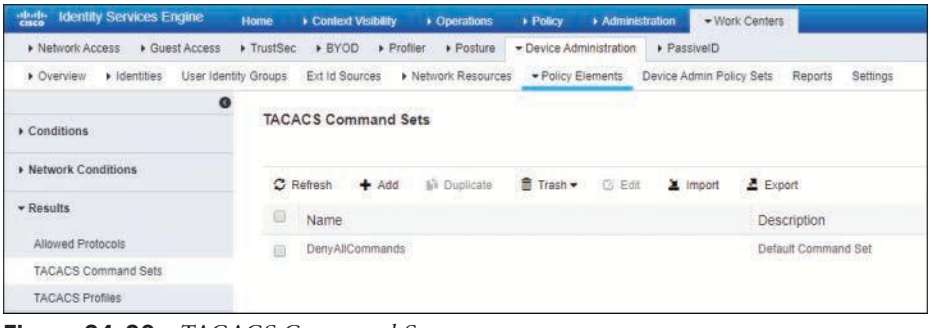


Figure 24-20 *TACACS Command Sets*

A command set is a group of commands that should be permitted, commands that should be denied, or any combination of the two. A command set provides you with the ability to use simple or complex combinations of permit and deny statements.

Click Add to see the available choices for a command set, as shown in Figure 24-21.

Figure 24-21 keeps things very simple to illustrate a point about command sets. For this particular command set, the checkbox Permit Any Command That Is Not Listed Below is enabled, and only one command is listed: The **configure** command with an asterisk (*) for the command's arguments is denied always.

The choices for each command are Permit, Deny, and Deny Always. The difference between Deny and Deny Always has to do with stacking command sets—that is, combining multiple command sets together in an authorization result. Permit always trumps Deny, but Deny Always wins every time.

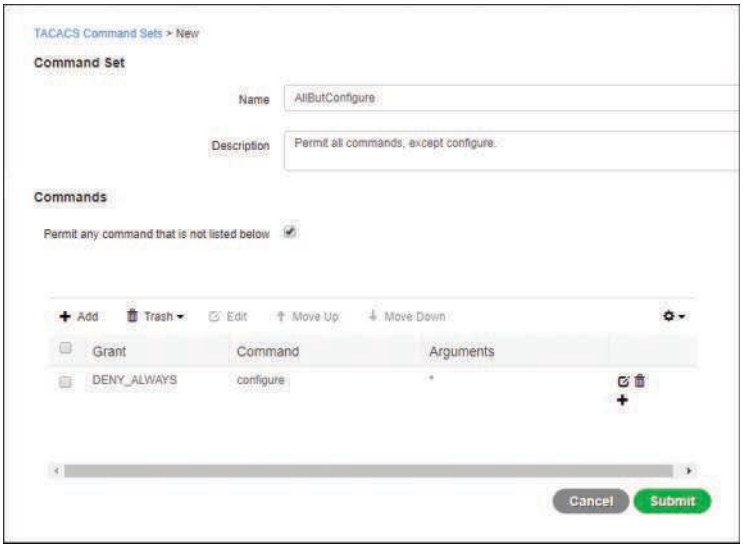


Figure 24-21 *A Look at a Command Set*

NOTE You will work more with command sets in the following chapters, when you configure device administration with TACACS+.

TACACS Profiles

Key Topic

TACACS profiles are aptly compared to authorization profiles with network access (RADIUS). They are the basic form of authorization result for device administration, and their function is similar to allowing access or denying access, but their results are less granular than are available with command sets.

To work with TACACS profiles, you navigate to Work Centers > Device Administration > Policy Elements > Results > TACACS Profiles, as shown in Figure 24-22. As you can see, there are some predefined profiles. In addition, there are different profile types: Shell (IOS), WLC, Nexus, and Generic. These types make it even easier to work with well-known devices.

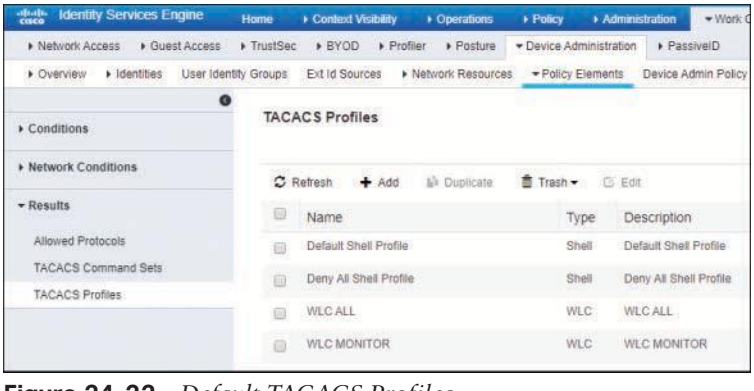


Figure 24-22 *Default TACACS Profiles*

Figure 24-23 shows an example of a TACACS profile for an IOS device (which is a shell type). Figure 24-24 shows the Raw View tab for the same profile. The Raw View tab is a live view, and the Task View tab provides a friendly view of the raw data. You can make changes in either the Task Attribute View tab or the Raw View tab.

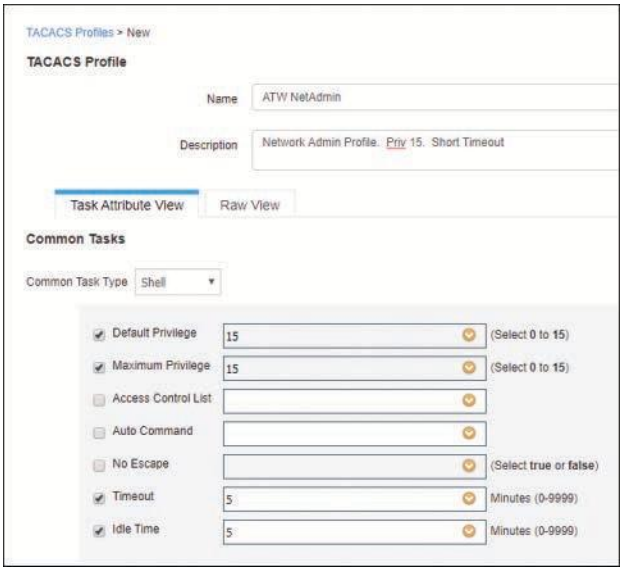


Figure 24-23 *TACACS Profile: Task Attribute View Tab*

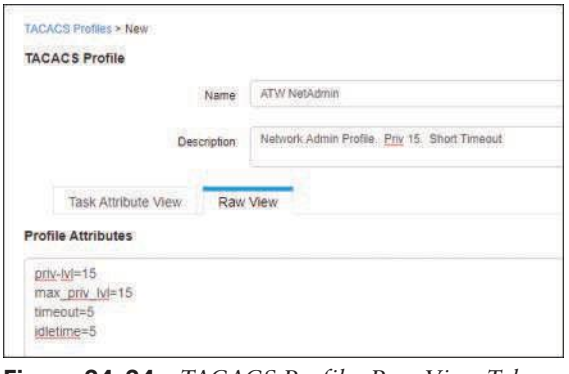


Figure 24-24 *TACACS Profile: Raw View Tab*

NOTE You will see the different profile types in subsequent chapters, when you configure TACACS+ for IOS, WLC, and Nexus devices.

Policy Sets

Policy sets work the same way with device administration as they do with network access. They are combinations of authentication and authorization rules.

You can and should create policy sets that work with your organization's needs, perhaps based on location, job role, or line of business.

For the purposes of this chapter and the chapters that follow, here we walk through the configuration of two policy sets, one per device type that we will work with in these chapters: Cisco Catalyst switches (IOS) and the Cisco Wireless LAN Controller (WLC).

To create the Catalyst policy set, follow these steps:

Step 1. Navigate to **Work Centers > Device Administration > Device Admin Policy Sets.**

Step 2. Click on the **+** sign to create a new policy set, as highlighted in Figure 24-25.

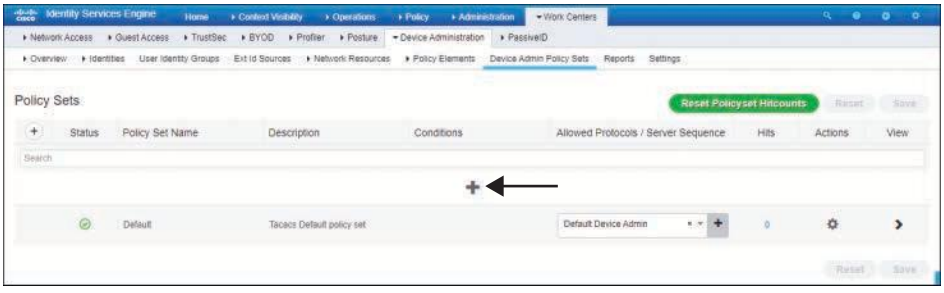


Figure 24-25 *Device Admin Policy Sets*

Step 3. Name the policy set **Catalyst Switches** and provide a description.

Step 4. Select **Default Device Admin** from the Allowed Protocols/Server Sequence dropdown, as shown in Figure 24-26.

Step 5. Click the **+** sign in the Conditions column to display the Conditions studio.

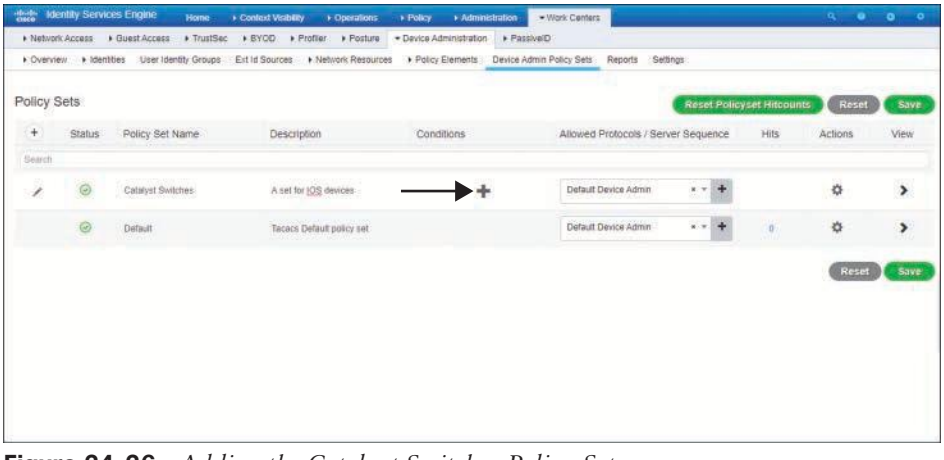


Figure 24-26 *Adding the Catalyst Switches Policy Set*

NOTE These policy sets use the NDGs created in earlier chapters.

Step 6. Create a condition in the Conditions studio that matches a device type that starts with **All Device Types#Switches#Access-Layer#Cisco**, as shown in Figure 24-27.

Step 7. Click **Use**.

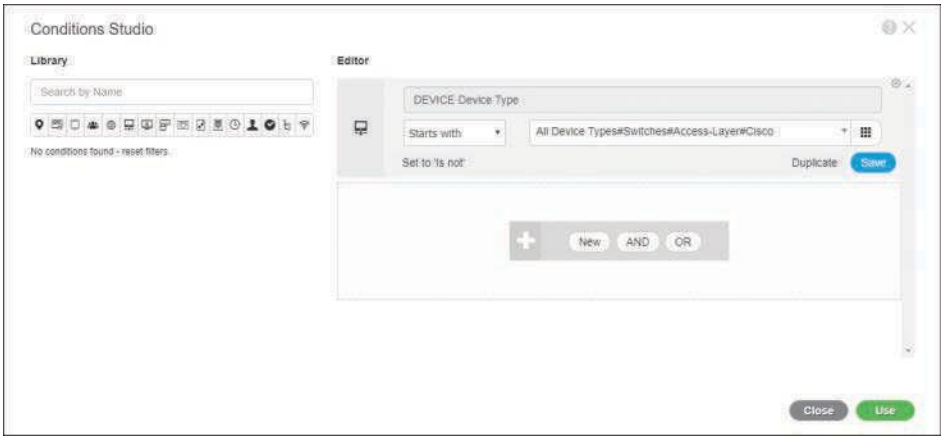


Figure 24-27 *Catalyst Switch Policy Set*

Next, create a WLC policy set by following these steps:

Step 1. Navigate to **Work Centers > Device Administration > Device Admin Policy Sets**.

Step 2. Click on the **+** sign to create a new policy set.

Step 3. Name the policy set **Wireless Controllers** and provide a description.

Step 4. Select **Default Device Admin** from the Allowed Protocols/Server Sequence dropdown.

Step 5. Choose the condition **DEVICE:Device Type STARTS WITH All Device Types#WiFi#Cisco**.

Step 6. Click **Use**.

Step 7. Click **Save**.

Figure 24-28 shows a summary of the added policy sets.

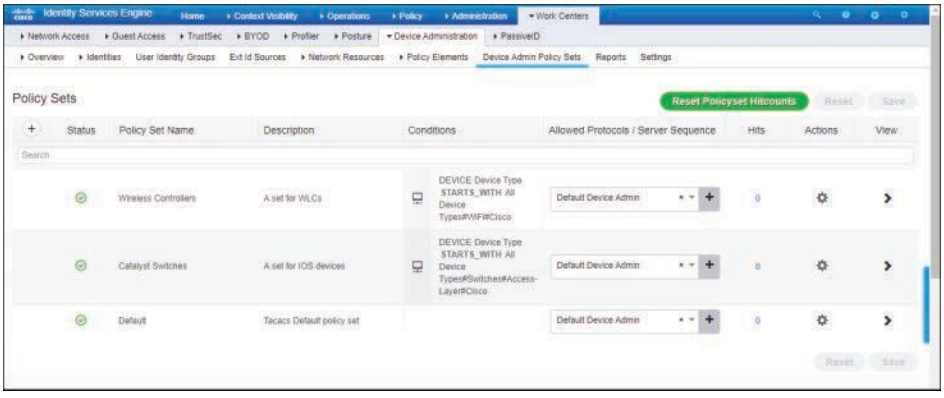


Figure 24-28 *Summary of Policies*

NOTE It is important to note that TACACS+ is not the only protocol that can be used for device administration. Some devices may use RADIUS, even though RADIUS is better suited for network access AAA. If you need a device administration AAA policy set for RADIUS, that set cannot exist within the Device Administration work center; it needs to be created in the Network Access work center because ISE logically separates them based on the AAA protocol.

Reports

The reports related to device administration are packaged up quite neatly within the Device Administration work center, to avoid your having to leave the work center. By navigating to Work Centers > Device Administration > Reports > Device Administration Reports, you can see the default reports, as shown in Figure 24-29.

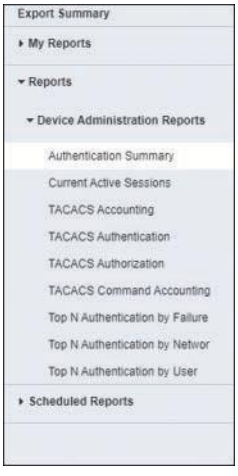


Figure 24-29 *Reports*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 24-2 lists these key topics and the page number on which each is found.

Key Topic

Table 24-2 Key Topics

Key Topic Element	Description	Page
Paragraph	Device administration AAA	909
Paragraph	The Device Admin license	911
Section	TACACS+ commands permitted for use	922
Section	TACACS profiles	923

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What is the name of the authorization result for a device administration rule?
2. What port does TACACS+ communicate on by default?
3. What are TACACS command sets used for?
4. When an administrator needs to refresh the shared secret between many network devices and ISE, what should the administrator do?
5. What network resource is used to forward requests to an external TACACS+ server in order?

This page intentionally left blank



CHAPTER 25

Configuring Device Administration AAA with Cisco IOS

This chapter covers the following topics:

- Configuring network access devices
- Configuring policies, including authentication and authorization profiles
- Compare AAA protocols
- Configure TACACS+ device administration and command authorization

Chapter 24, "Device Administration AAA with ISE," explains how Cisco ISE is a policy server that provides authentication, authorization, and accounting (AAA) for the purpose of access control and device administration.

This chapter explores device administration on Cisco IOS. With IOS devices, you can control the level of access to the shell as well as provide granular per-command authorization. This chapter covers both of the steps involved in enabling ISE to manage an IOS device: (1) configuring ISE for device administration and (2) configuring the IOS device for TACACS+.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 25-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 25-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Overview of IOS Device Administration AAA	1–4, 6
Configure ISE and an IOS Device for Device Administration AAA	5, 7, 9
Testing and Troubleshooting	8, 10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you incorrectly guess skews your self-assessment results and might provide you with a false sense of security.

1. How many privilege levels are available in Cisco IOS?
   a. 15
   b. 16
   c. 9
   d. 17
2. Which of the following best describes the way ISE provides device administration for IOS?
   a. ISE provides role-based access control (RBAC) by authenticating the user, authorizing the commands and privilege level, and providing full command accounting.
   b. ISE provides role-based access control (RBAC) by authenticating the user and authorizing the shell privilege level.
   c. ISE provides role-based access control (RBAC) by authenticating the user, authorizing the shell privilege level, and providing command accounting.
   d. ISE provides role-based access control (RBAC) by authenticating the user, authorizing the commands, and providing full command accounting.
3. Which IOS privilege level provides basic monitoring commands and has a prompt that features the symbol > after the device hostname?
   a. Privilege level 15
   b. Privilege level 0
   c. Privilege level 1
   d. Privilege level 7
4. Which of the following privilege levels is not defined by default and available for customization?
   a. Privilege level 0
   b. Privilege level 1
   c. Privilege level 15
   d. Privilege level 2
5. What is the purpose of configuring fallback in a AAA method list?
   a. To provide a fallback ISE PSN in the event that the first one fails
   b. To provide a fallback password to use as a TACACS shared secret
   c. To provide a fallback AAA method in the event that the first method is unavailable
   d. To provide a fallback TACACS configuration for troubleshooting
6. What is the common task type for an IOS device?
   a. Shell
   b. WLC
   c. Nexus
   d. Generic

7. What command would you use to define an authentication list for users moving from user exec mode to privileged exec mode?
   a. **aaa authentication login**
   b. **aaa authentication enable**
   c. **aaa authentication password-promp**
   d. **aaa authentication local-override**
8. Which **debug** command allows you to troubleshoot the transport protocol between an IOS device and ISE?
   a. **debug tacacs**
   b. **debug aaa**
   c. **debug ip tcp transactions**
   d. **debug tacacs packets**
9. Which command should be used to authenticate incoming SSH sessions?
   a. **aaa authentication ssh**
   b. **aaa authentication enable**
   c. **aaa authentication login**
   d. **aaa authentication password-prompt**
10. Which IOS **debug** command can you use to debug TACACS+ packet traces?
    a. **debug tacacs aaa**
    b. **debug tacacs packet**
    c. **debug tacacs**
    d. **debug tacacs events**

Foundation Topics

Overview of IOS Device Administration AAA

With Cisco IOS devices, administrative access to the network device console, Telnet sessions, and Secure Shell (SSH) sessions can be controlled by using RADIUS or TACACS+. While RADIUS may be used for device administration, there are limitations due to authentication and authorization being combined into the same step. With RADIUS, you do not have the ability to perform granular per-command authorization as you do with TACACS+. RADIUS is typically better suited for network access control than for device administration. For this reason, we mainly focus on TACACS+ in this chapter.

With TACACS, you can control access to IOS devices by using two elements: the TACACS profile and command sets.

TACACS Profile

Key Topic

The Cisco IOS shell is a command-line interface that has different ***privilege levels***. A privilege level defines the level of access a user has in an IOS shell session. Cisco IOS provides 16 levels of access privileges. By default, the higher the level, the more commands and access a user has available. Only three privileges are defined by default:

- **Privilege level 0**: No real access at all. This permits only the **disable**, **enable**, **exit**, **help**, and **logout** commands. At level 0, all these commands are available to all users.

- **Privilege level 1**: This is user exec mode, which is the default level for a user after login. The shell prompt is the device hostname followed by the > symbol (for example, Switch>).
- **Privilege level 15**: This is privileged exec mode. This level is reached after the **enable** command is entered. The shell prompt for privileged mode is the device hostname followed by the # sign (for example, Switch#).

The remaining privilege levels, 2 through 14, are available for customization in individual IOS devices. If you were to use a customized privilege level, it would need to be configured on an IOS device first before it could be called as a TACACS+ result. However, defining custom privilege levels is beyond the scope of the SISE exam, so this chapter focuses only on configuration using the defined privilege modes.

A TACACS profile in ISE controls the initial login session and the highest privilege level a user can access in the shell session. A TACACS profile can be viewed and created by navigating to Work Centers > Device Administration > Policy Elements > Results > TACACS Profiles. A TACACS profile enables you to configure privilege levels, timers, and other details that are sent as attribute/value pairs (AV pairs) to a network device.

Table 25-2 lists the options that are configurable in a TACACS profile in ISE.

Table 25-2 Configurable Options in a TACACS Profile in ISE

Option	Description
Common Task Type	From this dropdown, you can select the common tasks associated with creating this profile. There are four preset common task types: ■ **Shell**: This is used for an IOS or IOS XE device. ■ **WLC**: This is used for a wireless LAN controller. ■ **Nexus**: This is used for a Nexus device. ■ **Generic**: This common task type is for use with non-Cisco devices. You can add custom attributes for this task type.
Default Privilege	This option defines the privilege level for the administrator upon initial authorization. The values available, 0 through 15, correspond with the possible shell privilege levels.
Maximum Privilege	This option defines the maximum privilege level for which the user is authorized in the shell. The values available are 0 through 15.
Access Control List	If there is a configured access control list on the IOS NAD to permit or deny certain IP addresses, it can be called with this option.
Auto Command	This option can be used to specify a command that will be executed at exec mode startup.
No Escape	This option, which can be defined as true or false, can prevent a user from using the escape character.
Timeout	This option defines the number of minutes before an exec session disconnects. If set to 0, there will be no timeout.
Idle Timeout	This option defines the number of minutes after which an idle session is terminated.
Custom Attributes	An ISE administrator can add custom attributes in addition to the common tasks.

As shown in Figure 25-1, the Task Attribute View tab shows these attributes and options in an easy-to-understand format.

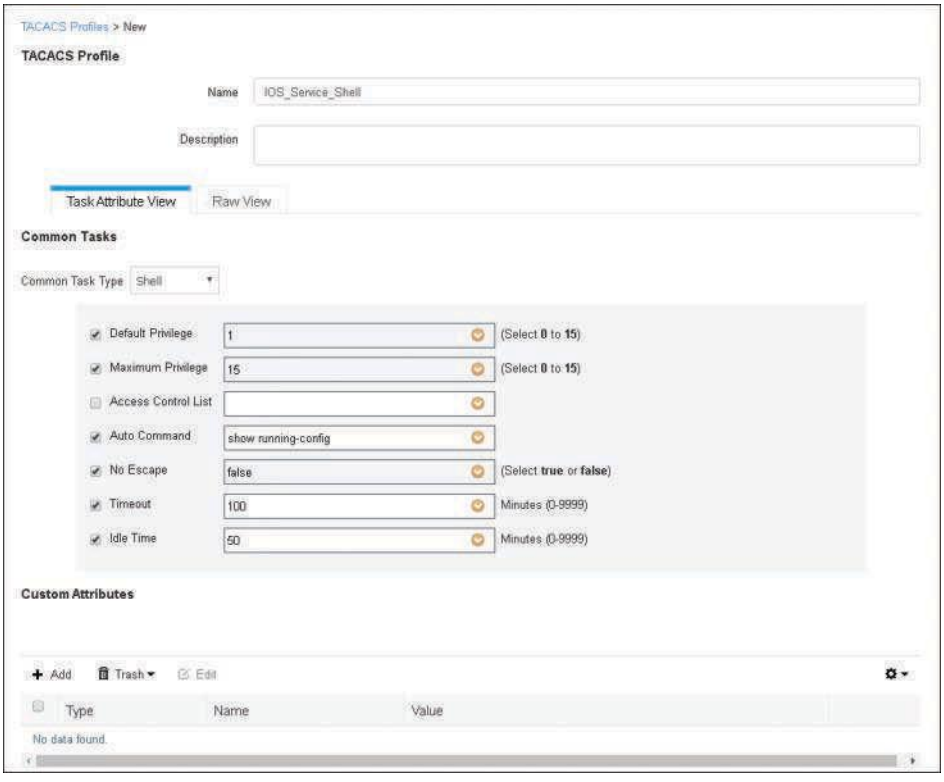


Figure 25-1 *TACACS Profile Task Attribute View*

If you click on the Raw View tab, shown in Figure 25-2, you can see the raw AV pairs that will be sent to the network access device.

TACACS+ Command Sets

The use case for defining a custom IOS privilege level could be to configure a customizable set of commands in order to achieve role-based access control (RBAC). However, these customizable privilege levels would have to be configured on every single network access device in order to be usable. If a command needs to be added or removed from a defined privilege level, a configuration change would be needed on each IOS device. As you can imagine, customizing privilege levels does not scale well.

Command authorization with TACACS+ is a much more eloquent solution than customizing privilege levels. With command authorization, you can allow a user to authenticate and authorize that user to privilege level 15 but allow the user to run only a subset of the commands that are usually available in that privilege level. As discussed in Chapter 24, TACACS+ allows you to authorize, in real time, every command entered on a network access device, as illustrated in Figure 25-3.

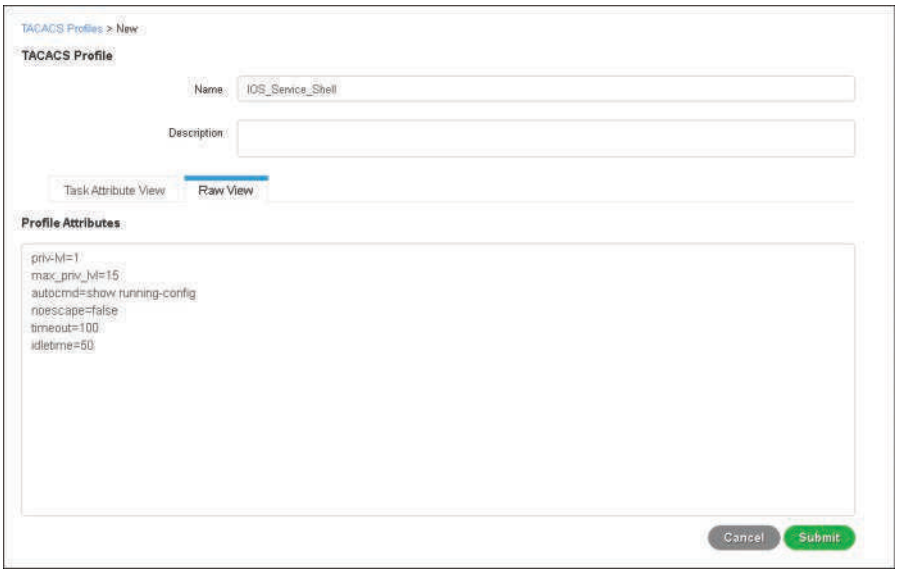


Figure 25-2 *TACACS Profile Raw View*

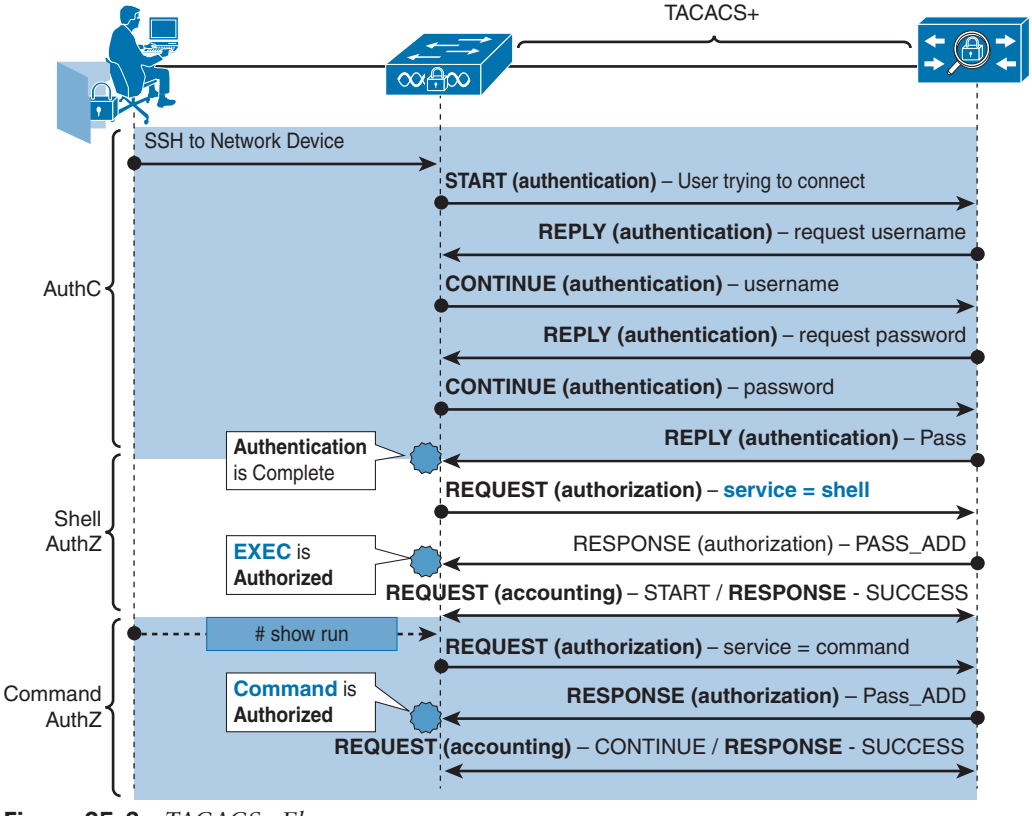


Figure 25-3 *TACACS+ Flow*

Key Topic

TACACS *command sets* are part of the configurable authorization results in ISE that allow you to define specific lists of commands that can be executed. You can configure TACACS+ command sets in ISE by navigating to Work Centers > Device Administration > Policy Elements > Results > TACACS Command Sets.

In a TACACS command set configuration, you list the commands that you would like to allow or deny. If you want to allow all commands except for a certain few, you can list the commands and check the Permit Any Command That Is Not Listed Below box.

As shown in Figure 25-4, there are three aspects to configure within a command set:

- **Grant**: You can choose Permit, Deny, or Deny Always. Except with Deny Always, ISE permits a command based on the first match that is permitted. The Deny Always setting always has a higher precedence over a match than does Permit.
- **Command**: This is where you define the command. You can use wildcards, and the command is not case-sensitive.
- **Arguments**: This might be one or more parameters that follow the command. You can use regular expressions with these arguments.

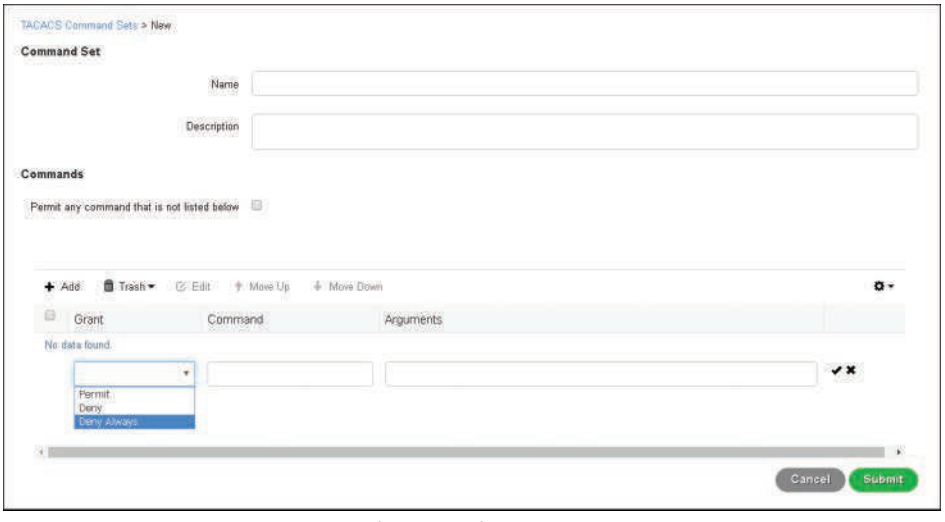


Figure 25-4 *TACACS Command Set Configuration*

Configure ISE and an IOS Device for Device Administration AAA

The steps in configuring device administration AAA are very similar on all types of network devices. You need to perform some configuration on ISE to ensure that ISE is ready to receive the TACACS+ requests from the NAD, and you must configure the NAD to send those TACACS+ requests to ISE.

Prepare ISE for IOS Device Administration AAA

In the ISE GUI, you need to ensure that network device objects are configured correctly and assigned to appropriate network device groups (NDGs). The authorization results should then be created to assign privileges. Finally, you create a policy set to tie the authorization result to conditions.

Ensure That the Device Administration Service Is Enabled

As discussed in Chapter 24, you need to enable the Device Administration service on a Policy Services node (PSN). To do so, follow these steps in the ISE GUI:

Step 1. Navigate to **Work Centers > Device Administration > Overview > Deployment** and activate the service on the applicable PSNs (see Figure 25-5).

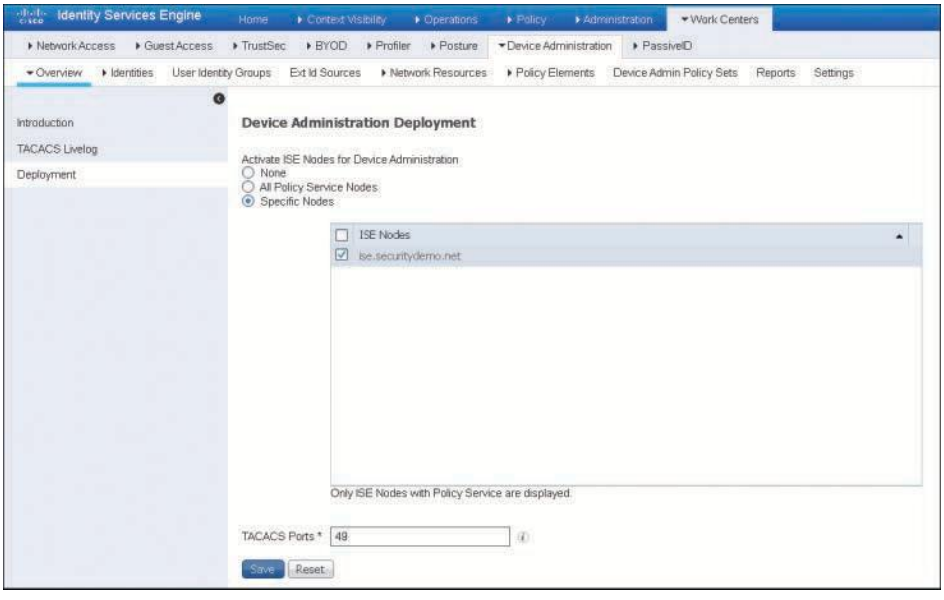


Figure 25-5 *Enabling the Device Administration Service*

Step 2. Click **Save**.

Prepare the Network Device

To prepare the network device, follow these steps in the ISE GUI:

Step 1. Navigate to **Work Centers > Device Administration > Network Resources > Network Device Groups**. You should already have a detailed set of hierarchical NDGs similar to what you see in Figure 25-6. Ensure that you have the appropriate groups to go along with the policies you want to create.

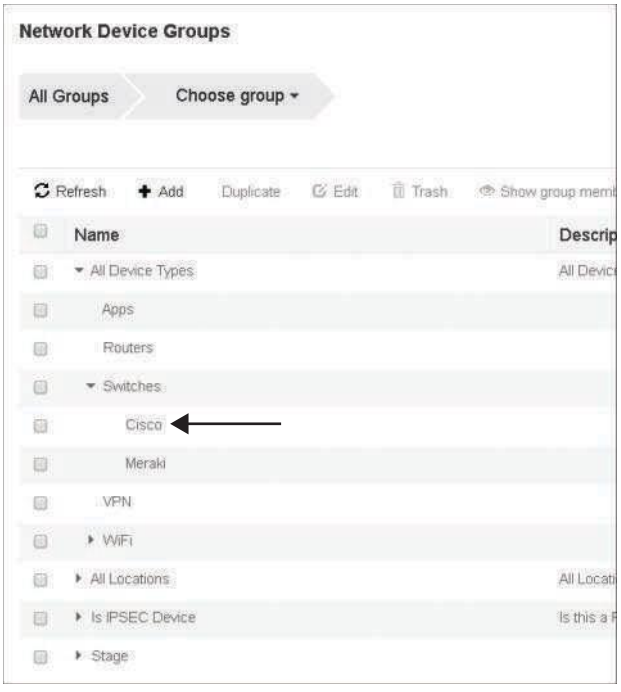


Figure 25-6 *NDG Hierarchy Example*

At the very least, ensure that these IOS devices have their own group for a device type. By using network device groups as conditions, you can ensure that a device administration policy set is built just for these IOS devices.

Step 2. Ensure that the Cisco IOS network access device itself has a network device object in ISE with the TACACS+ shared secret configured. In this example, you will be adding an IOS switch as your network device object, so navigate to **Work Centers > Device Administration > Network Resources > Network Devices.**

Step 3. Click on the NAD or click **Add** to create a new object.

Step 4. Ensure that the network device groups are assigned properly and that the TACACS+ shared secret is configured correctly.

Step 5. Check the box next to **TACACS Authentication Settings** and enter a TACACS+ shared secret.

Optionally, you may also check the **Enable Single Connect Mode** box if the IOS device supports this mode. Single Connect Mode optimizes TACACS+ by keeping a single TCP connection for the entirety of the session, and all the

transactions will be made on this connection. It is by no means a requirement and many network access devices do not support it.

Figure 25-7 shows the NAD settings.

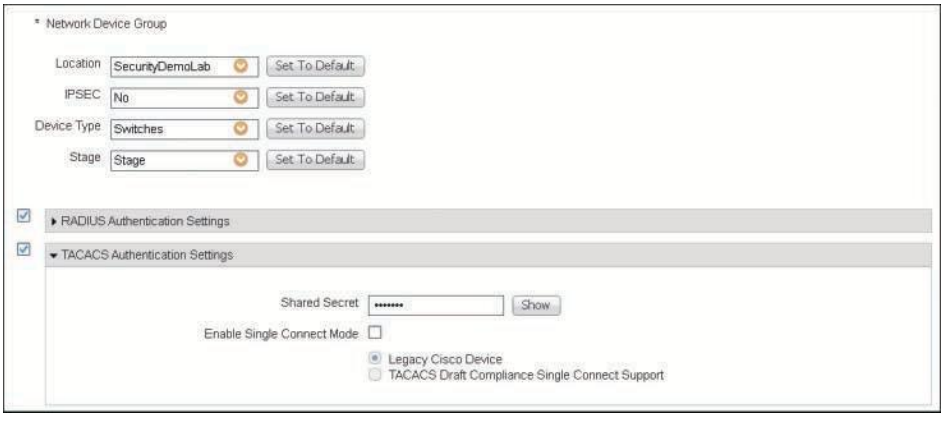


Figure 25-7 *IOS Network Device Object*

Step 6. Click **Save**.

Prepare the Policy

Based on your configurations so far, ISE can accept TACACS+ requests from the IOS device, but you have no policies and no authorizations to send down to the NAD yet. This section shows how to create TACACS+ shell profiles and command sets. Specifically, it shows how to create three different roles for use in a policy:

- **HelpDesk**: This role will be for help desk analysts who will perform **debug**, **undebug**, **traceroute**, **ping**, and **show** commands.
- **Security Administrators**: This role will be for security administrators who may need to shut down or un-shutdown certain interfaces. However, you do not want them to shut down the uplink interface and lose access to the switch, so we will prevent them from doing so.
- **Network Administrator**: This role will be for network administrators who should have access to the full range of commands on the IOS device.

Configure the TACACS Profiles

This section shows how to define the TACACS profiles that will be used for the three roles described in the preceding section. First, you have to navigate to Work Centers > Device Administration > Policy Elements > Results > TACACS Profiles. Figure 25-8 shows the TACACS Profiles page.

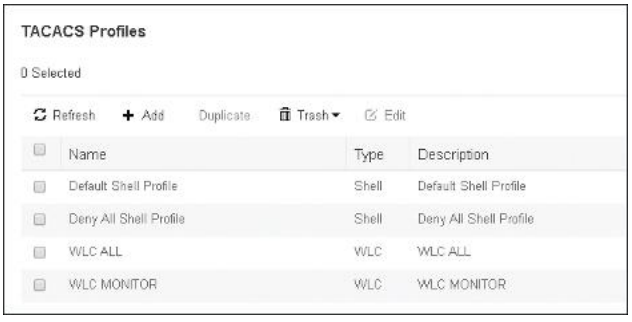


Figure 25-8 *TACACS Profiles*

There are some predefined profiles available, but in this case, you need to create some new profiles by clicking **Add**.

Create the HelpDesk Profile

Next, you need to create a new TACACS profile to use when a help desk analyst logs in to the IOS device:

Step 1. Name the profile **IOS_HelpDesk_Privilege**.

Step 2. Under Common Tasks, set Default Privilege to **1** and Maximum Privilege to **15**.

Step 3. Click **Save** to save the profile.

Figure 25-9 shows the completed IOS_HelpDesk_Privilege profile.



Figure 25-9 *IOS_HelpDesk_Privilege TACACS Profile*

Create the Administration Profile

Now you can create a new TACACS profile that you will use with both the Security Administrator and Network Administrator roles:

Step 1. Click **Add** to create another TACACS profile.

Step 2. Name the profile **IOS_Admin_Privilege**.

Step 3. Under Common Tasks, set Default Privilege to **15** and Maximum Privilege to **15** (see Figure 25-10).

Step 4. Click **Save** to save the profile.



Figure 25-10 *IOS_Admin_Privilege TACACS Profile*

Configure the TACACS Command Sets

Now you can define the TACACS command set for each of the roles to enforce granular command authorization, regardless of the privilege level a role is granted in the shell.

To define a command set, navigate to Work Centers > Device Administration > Policy Elements > Results > TACACS Command Sets. Figure 25-11 shows the TACACS Command Sets page.

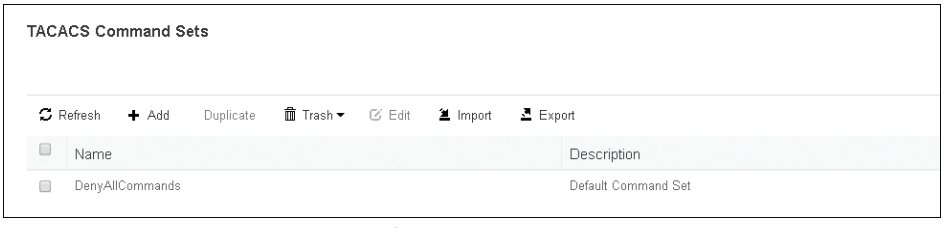


Figure 25-11 *TACACS Command Sets*

The only predefined command set is DenyAllCommands, so you need to add new command sets.

Follow these steps to create the HelpDesk command set:

Step 1. Click **Add** to create a new command set.

Step 2. Name the new command set **IOS_HelpDesk_CS**.

Step 3. In the Commands section, add the commands listed in Table 25-3.

Table 25-3 HelpDesk Command Set

Grant	Command	Argument
Permit	**debug**	
Permit	**undebug**	
Permit	**traceroute**	
Permit	**ping**	
Permit	**show**	

Step 4. Click the checkmark at the end of each entry to keep that line.

Step 5. Click **Save**.

Figure 25-12 shows the completed IOS_HelpDesk_CS configuration.

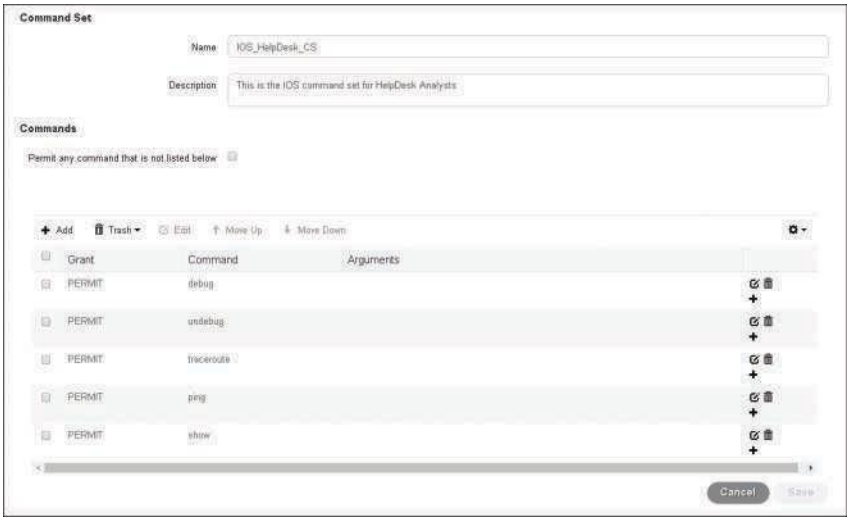


Figure 25-12 *IOS_HelpDesk_CS Command Set*

Follow these steps to **Create the Security Administrator Command Set:**

Step 1. Click **Add** to create a new command set.

Step 2. Name the new command set **IOS_Security_CS**.

Step 3. In the Commands section, add the commands listed in Table 25-4.

Table 25-4 Security Administrator Command Set

Grant	Command	Argument
Permit	**config***	
Deny Always	**interface**	GigabitEthernet 1 0 48
Permit	**interface**	
Permit	**shutdown**	
Permit	**no**	shutdown

Step 4. Click the checkmark at the end of each entry to keep that line.

Step 5. Click **Save**.

Figure 25-13 shows the completed IOS_Security_CS configuration.

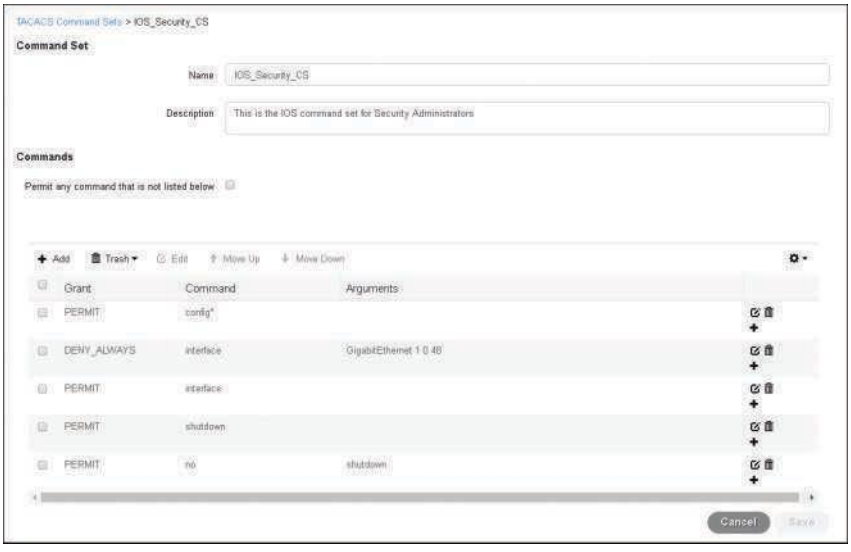


Figure 25-13 *IOS_Security_CS Command Set*

Follow these steps to **Create the Network Administrator Command Set:**

Step 1. Click **Add** to create a new command set.

Step 2. Name the new command set **IOS_ Network _CS.**

Step 3. Check the **Permit Any Command That Is Not Listed Below** box and do not add any commands below. By doing so, you will permit all commands that are available.

Step 4. Click **Save.**

Figure 25-14 shows the completed IOS_Network_CS configuration.

Configure the Policy Set

When you have all the role types you need, it's time to configure the device administration policy set for IOS devices, as described in the sections that follow.

Create the Policy Set

Step 1. Navigate to **Work Centers> Device Administration> Device Administration Policy Sets**

Step 2. Click the plus (+) sign to add a new policy set.

Step 3. Name the Policy Set as **IOS Device Administration.**

Step 4. Set the Policy Set condition as **DEVICE: Device Type EQUALS All Device Types#Switches#Cisco.**



Figure 25-14 *IOS_Network_CS Command Set*

NOTE If you have other IOS device types than switches, you may add the network device group with an OR operator.

Step 5. From the Allowed Protocols list, select **Default Device Administration**.

NOTE This is a built-in allowed protocol list that allows TACACS+ with PAP/ASCII, CHAP, and MS-CHAPv1. You may alternatively create a custom allowed protocol list under **Work Centers > Device Administration > Policy Elements > Results > Allowed Protocols**.

Step 6. Click **Save** to save the new policy set.

Figure 25-15 shows the newly created Device Administration Policy Set

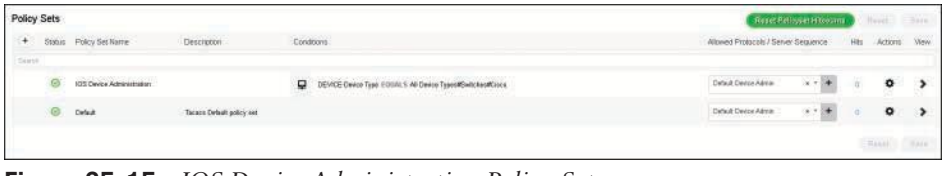


Figure 25-15 *IOS Device Administration Policy Set*

Configure the Authentication Policy

To configure the authentication policy, expand the newly created policy set and expand the authentication policy. For the authentication policy, you can easily use the default rule to specify an identity source sequence or a specific identity store. In this case, you can use the default rule and keep the All_User_ID_Store identity source sequence as the chosen identity store. Figure 25-16 shows the authentication policy.

Alternatively, you can create authentication rules with more specific conditions. For device administration, the dictionaries are limited to device, network access, and TACACS conditions.



Figure 25-16 *IOS Device Administration Authentication Policy*

Configure the Authorization Policy

This section uses the following Active Directory groups:

- **sec-admin**: Active Directory group membership for the security administrators
- **net-admin**: Active Directory group membership for the network administrators
- **helpdesk**: Active Directory group membership for the help desk analysts

Follow these steps to configure the authorization policy:

Step 1. Expand the authorization policy in the policy set.

Step 2. Click the + button to create a new authorization rule.

Step 3. Configure the following authorization rule:

Rule Name	Network Administrators
Conditions	ad1:ExternalGroups EQUALS net-admin
Command Sets	IOS_Network_CS
Shell Profiles	IOS_Admin_Privilege

Step 4. Add another authorization rule below the preceding rule and configure it as follows:

Rule Name	Security Administrators
Conditions	ad1:ExternalGroups EQUALS sec-admin
Command Sets	IOS_Security_CS
Shell Profiles	IOS_Admin_Privilege

Step 5. Add another authorization rule below the preceding rule and configure it as follows:

Rule Name	Helpdesk Analysts
Conditions	ad1:ExternalGroups EQUALS helpdesk
Command Sets	IOS_HelpDesk_CS
Shell Profiles	IOS_HelpDesk_Privilege

Step 6. Click **Save** to save the authorization policy.

Figure 25-17 shows the completed authorization policy.



Figure 25-17 *IOS Device Administration Authentication Policy*

IOS Configuration for TACACS+

Now that the policy set is ready, it is time to configure the IOS device. There are three main parts of the TACACS+ configuration on IOS devices:

- **Configuration of TACACS+ authentication and fallback:** With this configuration, you specify the TACACS+ server to authenticate to and a fallback method, such as a locally configured username and password on the IOS device. When a user authenticates to the device after this is configured, the TACACS+ shell profile configured in the authorization policy is returned. ISE supports both IPv4 and IPv6 for TACACS+.
- **Configuration of TACACS+ command authorization:** Command authorization allows the IOS network access device to send requests to ISE when a command is executed. ISE matches the request from the NAD to the list of commands in the TACACS+ command sets and allows the command to be executed, if that command is permitted.
- **Configuration of TACACS+ command accounting:** Command accounting sends information about each command that can be executed, including the command, date, and username.

Configure TACACS+ Authentication and Fallback

When you configure TACACS+ on an IOS device, there is a possibility of being locked out of the device if the commands are not entered correctly or in the correct order. To mitigate this possibility, network administrators often issue the **reload in** *min* command prior to starting the configuration. If the reload is not canceled with the **reload cancel** command within the configured number of minutes, the IOS device reloads. As long as you have not saved the running configuration, this should reset the device to your pre-TACACS+ configuration.



To configure TACACS+ authentication and fallback, follow these steps:

Step 1. Configure the switch to support SSH:

```
Switch(config)# hostname CAT9300
CAT9300(config)# ip domain name securitydemo.net
CAT9300(config)# crypto key generate rsa modulus 2048
% The key modulus size is 2048 bits
% Generating 2048 bit RSA keys, keys will be
non-exportable...
```

```
[OK] (elapsed time was 1 seconds)
CAT9300(config)# ip ssh version 2
```

Step 2. Configure an enable secret and local username and password on the switch:

```
CAT9300(config)# enable secret ISEc0ld
CAT9300(config)# username admin privilege 15 secret
ISEc0ld
```

Step 3. Enable AAA:

```
CAT9300(config)# aaa new-model
```

Step 4. Configure the TACACS+ servers on the IOS device:

```
CAT9300(config)# tacacs server ise-1
CAT9300(config)# address ipv4 10.1.100.21
CAT9300(config)# key ISEc0ld
```

Step 5. Add the TACACS server to a TACACS+ server group:

```
CAT9300(config)# aaa group server tacacs+ ISE-TACACS
CAT9300(config)# server name ise-1
```

Step 6. Configure the AAA method lists by using the TACACS+ server group and falling back to locally configured accounts if the TACACS+ servers are unreachable:

```
CAT9300(config)# aaa authentication login VTY group
ISE-TACACS local
CAT9300(config)# aaa authentication login CON none
CAT9300(config)# aaa authentication enable default group
ISE-TACACS enable
```

These method lists specify the following:

- The authentication method list named VTY attempts to authenticate to the ISE-TACACS server group and falls back to local login if it is unreachable.
- The authentication method list named CON uses no authentication.
- The default authentication for enable first tries the server group ISE-TACACS and then falls back to the enable password if there are no TACACS+ servers available.

Step 7. Apply a method list to the console access:

```
CAT9300(config)# line con 0
CAT9300(config-line)# exec-timeout 0 0
CAT9300(config-line)# login authentication CON
CAT9300(config-line)# logging synchronous
```

This method list allows access to the console without TACACS+. Alternatively, you can configure console access with the TACACS+ server group if you want to have every user authenticate.

Step 8. Apply the method list to the vty lines:

```
CAT9300(config)# line vty 0 15
CAT9300(config-line)# transport input ssh
CAT9300(config-line)# login authentication VTY
CAT9300(config-line)# logging synchronous
```

Step 9. If you want to further limit access to the vty lines, you can optionally configure an access control list with only the IP addresses or subnets that should be accessing it for device administration and apply it to the vty lines:

```
CAT9300(config)# ip access-list extended vtyAccess
CAT9300(config)# permit tcp 10.1.100.0 255.255.255.0 any
eq 22
CAT9300(config)# line vty 0 15
CAT9300(config-line)# access-class vtyAccess in
```

Step 10. To test the SSH access, open a new Putty window and log in with a variety of different accounts to verify that the security administrators, network administrators, and help desk analysts are able to connect via SSH.

Step 11. In ISE, navigate to **Operations > TACACS > Live Logs** to view the authentications. Figure 25-18 shows the TACACS+ authentications.

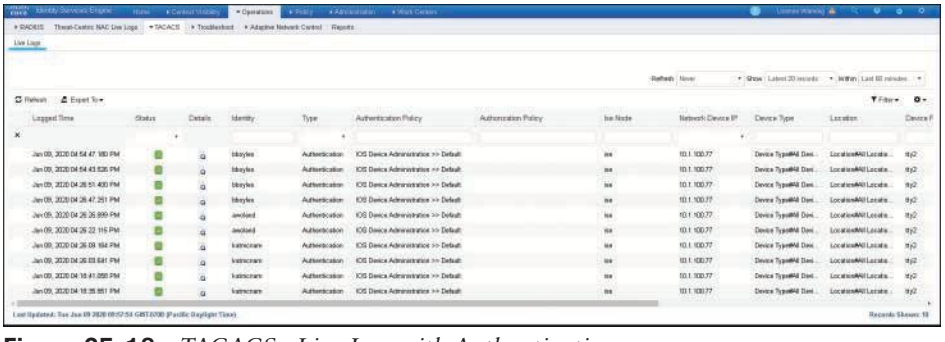


Figure 25-18 *TACACS+ Live Log with Authentications*

Configure TACACS+ Command Authorization

To configure TACACS+ command authorization, follow these steps:

Key Topic

Step 1. Create AAA authorization method lists for exec command authorization (which happens right after the user logs in to the device):

```
CAT9300(config)# aaa authorization exec CON none
CAT9300(config)# aaa authorization console
CAT9300(config)# aaa authorization exec VTY group
ISE-TACACS if-authenticated
```

These commands perform the following actions:

- Create an AAA authorization exec method list named CON that requires no authorization
- Enable authorization at the console level
- Create an AAA authorization exec method list named VTY that performs authorization against the ISE-TACACS server group if the user is authenticated

Step 2. Apply the AAA authorization method list to the console line:

```
CAT9300(config)# line con 0
CAT9300(config-line)# authorization exec CON
```

Step 3. Apply the AAA authorization method list to the vty lines:

```
CAT9300(config)# line vty 0 15
CAT9300(config-line)# authorization exec VTY
```

With this method list applied to the vty lines, the shell profiles with the default privilege attribute are applied for new SSH sessions.

Step 4. Ensure that the IOS device command authorization is in configuration mode and the privilege levels are enabled:

```
CAT9300(config)# aaa authorization config-commands
CAT9300(config)# aaa authorization commands 1 VTY group
ISE-TACACS local if-authenticated
CAT9300(config)# aaa authorization commands 15 VTY group
ISE-TACACS local if-authenticated
```

These commands perform the following actions:

- Enable authorization for config commands
- Perform authorization for commands at privilege level 1 by using the ISE-TACACS server group and fall back to the local login for all authenticated users
- Perform authorization for commands at privilege level 15 by using the ISE-TACACS server group and fall back to the local login for all authenticated users

Step 5. The command authorization is not enforced until the method lists are applied to the vty lines, so apply them to the VTY lines:

```
CAT9300(config)# line vty 0 15
CAT9300(config-line)# authorization commands 1 VTY
CAT9300(config-line)# authorization commands 15 VTY
```

Immediately after these commands are applied, you may find that you are no longer able to enter new commands in the console.

Step 6. Open a new SSH window and try using a variety of accounts. From the help desk account, attempt to write the memory to the IOS device. The command authorization should fail.

Step 7. In ISE, navigate to **Operations > TACACS > Live Logs** to view the passed and failed authorizations. Figure 25-19 shows the TACACS+ authorizations.

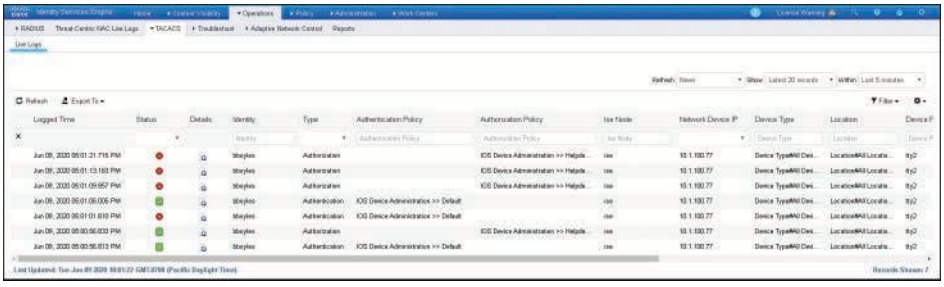


Figure 25-19 *TACACS+ Live Log with authorizations*

Click on **More Details** for a failed authorization to see more information about what failed. Figure 25-20 shows details on a failed authorization.



Figure 25-20 *TACACS+ Authorization Failure Details*

Configure TACACS+ Command Accountings

Follow these steps to configure TACACS+ command accountings:

Key Topic

Step 1. Configure the AAA accounting method lists on the IOS device:

```
CAT9300(config)# aaa accounting exec default start-stop
group ISE-TACACS
CAT9300(config)# aaa accounting commands 1 default
start-stop group ISE-TACACS
CAT9300(config)# aaa accounting commands 15 default
start-stop group ISE-TACACS
```

These method lists immediately start sending accounting packets to the ISE-TACACS server group. No further configuration is required on the IOS device.

Step 2. Try using a variety of commands on the network device to trigger failures and permits for command authorization.

Step 3. In ISE, navigate to **Operations > Reports > Reports > Device Administration > TACACS Command Accounting**. Figure 25-21 shows the TACACS Command Accounting report.

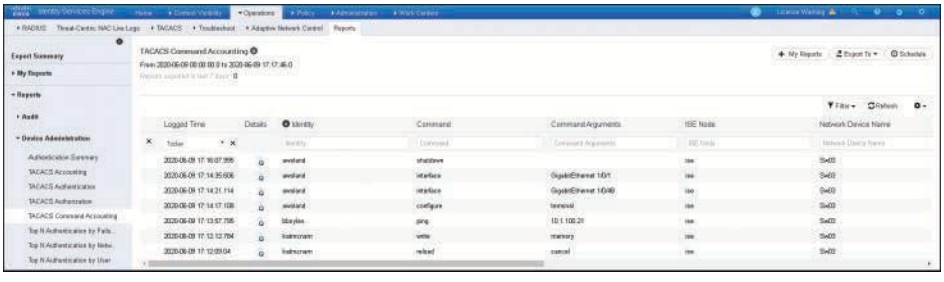


Figure 25-21 *TACACS+ Command Accounting Report*

Testing and Troubleshooting

Before testing, it is highly recommended that you change a security setting that was added in ISE Version 2.4 patch 3 and that remains a default security setting today. Once upon a time, someone performed a security audit on ISE and determined that it was an offensive security vulnerability to show the username of failed authentications in ISE Live Log because the idea was that it's possible that a user might accidently type his or her password into the username field, and that password would then show up in the log, violating security. Any failed authentications now appear as USERNAME in Live Log if the username is not found in any of the available identity stores until you enable the Disclose Invalid Usernames setting under Administration > Settings > Security Settings. Figure 25-22 shows this setting. Otherwise, a valid user account that fails authentication (such as using a wrong password) will still show up in the logs.



Figure 25-22 *Disabling the Disclose Invalid Usernames Setting*

Testing and Troubleshooting in ISE

It is typically easy to test logins and command authorization for an IOS device by initiating an SSH session to the device. Follow these steps:

Step 1. Use SSH to remotely access the IOS device and attempt to log in using the wrong password. Where do you look first to identify why authentication failed? You look in the TACACS Live Log in ISE.

Step 2. Navigate to **Work Centers > Device Administration > Overview > TACACS Live Logs**. Figure 25-23 shows the failed login from the bboyles account.

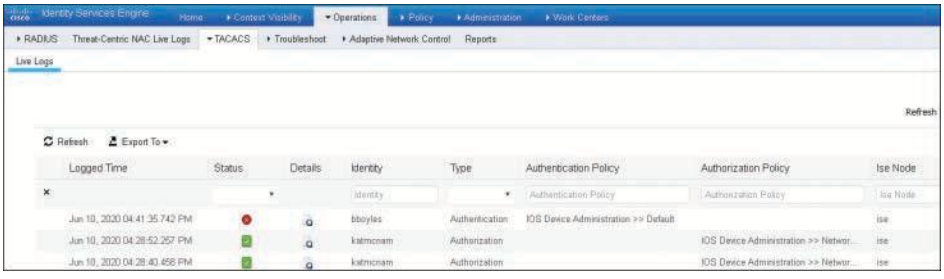
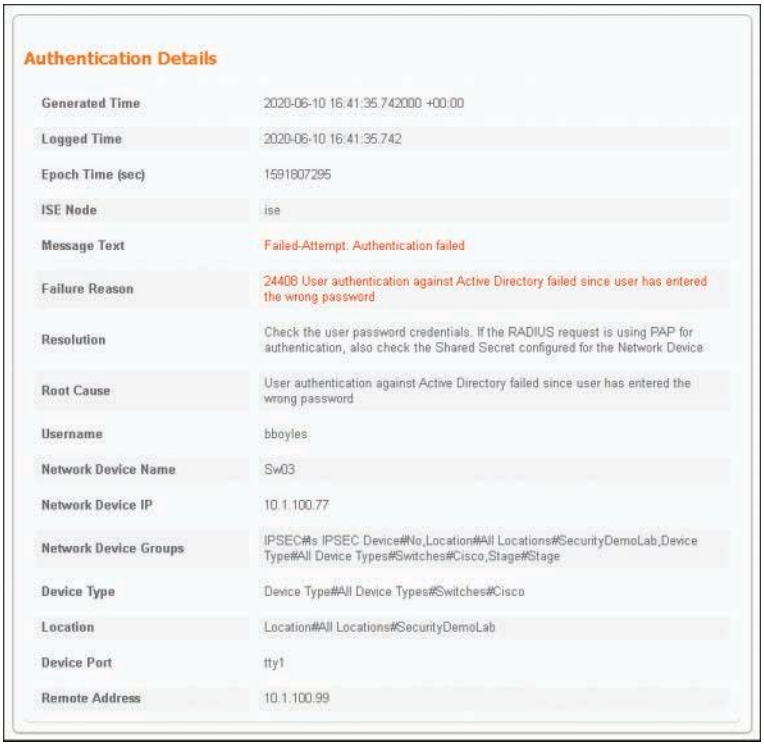


Figure 25-23 *TACACS Live Log – Failed Login*

Step 3. Click on the details icon to find the failure reason. In Figure 25-24, which shows the details screen that appears, you can confirm that the authentication failed because the password was incorrect.

Step 4. Log in again but this time use the correct password. You should now see a successful authentication completed in the TACACS Live Log, as shown in Figure 25-25.

Step 5. Enter the **show running-config** command as the bboyles user. This user is a help desk analyst, and the HelpDesk role is limited in the number of commands that can be issued. Issuing this command should be permitted, as shown in Figure 25-26.



Authentication Details

Generated Time	2020-06-10 16:41:35.742000 +00:00
Logged Time	2020-06-10 16:41:35.742
Epoch Time (sec)	1591807295
ISE Node	ise
Message Text	Failed-Attempt: Authentication failed
Failure Reason	24408 User authentication against Active Directory failed since user has entered the wrong password
Resolution	Check the user password credentials. If the RADIUS request is using PAP for authentication, also check the Shared Secret configured for the Network Device
Root Cause	User authentication against Active Directory failed since user has entered the wrong password
Username	bboyles
Network Device Name	Sw03
Network Device IP	10.1.100.77
Network Device Groups	IPSEC#Is IPSEC Device#No,Location#All Locations#SecurityDemoLab,Device Type#All Device Types#Switches#Cisco,Stage#Stage
Device Type	Device Type#All Device Types#Switches#Cisco
Location	Location#All Locations#SecurityDemoLab
Device Port	tty1
Remote Address	10.1.100.99

Figure 25-24 *TACACS Details*

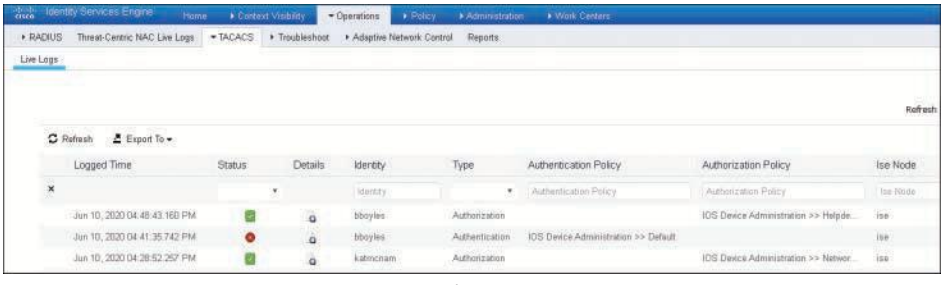


Figure 25-25 *TACACS Live Log Passed Login*

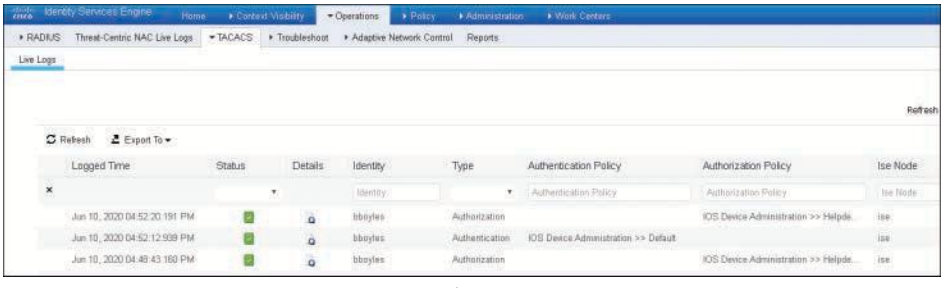


Figure 25-26 *TACACS Live Log Passed Authorization*

Step 6. Issue **config t** at the command line. The HelpDesk role should not be able to move into configuration mode, and you should see a failed authorization, as shown in Figure 25-27.

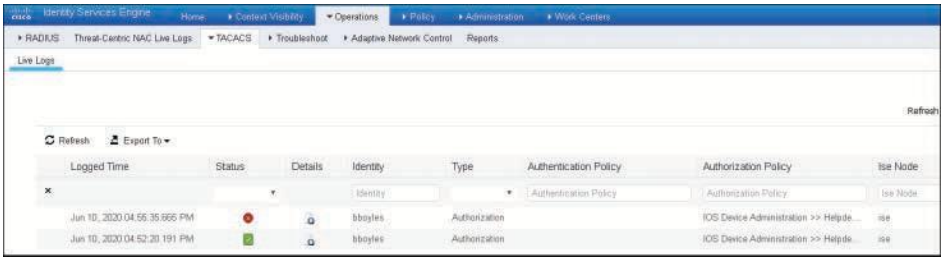


Figure 25-27 *TACACS Live Log Failed Authorization*

Step 7. Click the details icon to see more information about the failed authorization. As shown in Figure 25-28, the command failed to match a permit rule in the command set.

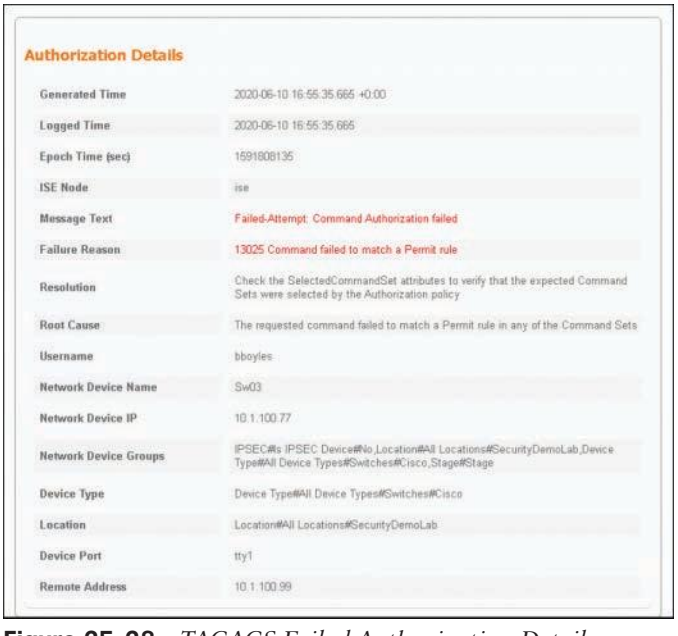


Figure 25-28 *TACACS Failed Authorization Details*

Troubleshooting at the IOS Command Line

While you can glean a lot of information from the ISE GUI, there are times when you need to troubleshoot from the network access device command line to gain more visibility. The following steps walk you through the process:

Step 1. Use SSH to remotely access the IOS device and log in to it by using one of the three roles.

Step 2. Verify the existing user connections:

```
CAT9300# show users
    Line      User       Host(s)     Idle      Location
   1 vty 0    bboyles    idle        00:00:19  10.1.100.99
*  2 vty 1    katmcnam   idle        00:00:00  10.1.100.99
  Interface   User            Mode   Idle      Peer Address
CAT9300#
```

Step 3. Use the **debug aaa authentication** command to troubleshoot AAA authentication issues on the IOS device:

NOTE The output of the **debug aaa authentication** command includes both RADIUS and TACACS+ authentications, so it may be rather chatty on a busy production switch using ISE for network access control.

```
CAT9300# debug aaa authentication
*Jun 10 08:49:51.212: AAA/BIND(00000FE7): Bind i/f
*Jun 10 08:49:51.212: AAA/AUTHEN/LOGIN (00000FE7): Pick
method list 'VTY'
*Jun 10 08:49:52.982: %SEC_LOGIN-5-LOGIN_SUCCESS: Login
Success [user: bboyles] [Source: 10.1.100.99] [localport:
22] at 08:49:52 UTC Wed Jun 10 2020
```

This output shows the successful authentication of the user bboyles.

Step 4. Use the **debug aaa authorization** command to troubleshoot AAA authorization issues on the IOS device:

NOTE The output of the **debug aaa authorization** command includes both RADIUS and TACACS+ authentications, so it may be rather chatty on a busy production switch using ISE for network access control.

```
CAT9300# debug aaa authorization
AAA Authorization debugging is on


*Jun 10 07:55:43.238: AAA/AUTHOR: auth_need : user=
'bboyles' ruser= 'CAT9300'rem_addr= '10.1.100.99' priv=
15 list= 'VTY' AUTHOR-TYPE= 'commands'
*Jun 10 07:55:43.238: AAA: parse name=tty1 idb type=-1
tty=-1
*Jun 10 07:55:43.238: AAA: name=tty1 flags=0x11 type=5
shelf=0 slot=0 adapter=0 port=1 channel=0
*Jun 10 07:55:43.238: AAA/MEMORY: create_user
(0x7FCBFD432410) user='bboyles' ruser='CAT9300' ds0=0
```

```
port='tty1' rem_addr='10.1.100.99' authen_type=ASCII
service=NONE priv=15 initial_task_id='0', vrf= (id=0)
key=D138FCA5
*Jun 10 07:55:43.238: tty1 AAA/AUTHOR/CMD (2095255949):
Port='tty1' list='VTY' service=CMD
*Jun 10 07:55:43.238: AAA/AUTHOR/CMD: tty1 (2095255949)
user='bboyles'
*Jun 10 07:55:43.238: tty1 AAA/AUTHOR/CMD (2095255949):
send AV service=shell
*Jun 10 07:55:43.238: tty1 AAA/AUTHOR/CMD (2095255949):
send AV cmd=show
*Jun 10 07:55:43.238: tty1 AAA/AUTHOR/CMD (2095255949):
send AV cmd-arg=running-config
*Jun 10 07:55:43.238: tty1 AAA/AUTHOR/CMD (2095255949):
send AV cmd-arg=<cr>
*Jun 10 07:55:43.238: tty1 AAA/AUTHOR/CMD(2095255949):
found list "VTY"
*Jun 10 07:55:43.238: tty1 AAA/AUTHOR/CMD (2095255949):
Method=ISE-TACACS (tacacs+)
*Jun 10 07:55:43.238: AAA/AUTHOR/TAC+: (2095255949):
user=bboyles
*Jun 10 07:55:43.238: AAA/AUTHOR/TAC+: (2095255949): send
AV service=shell
*Jun 10 07:55:43.238: AAA/AUTHOR/TAC+: (2095255949): send
AV cmd=show
CAT9300#
*Jun 10 07:55:43.238: AAA/AUTHOR/TAC+: (2095255949): send
AV cmd-arg=running-config
*Jun 10 07:55:43.238: AAA/AUTHOR/TAC+: (2095255949): send
AV cmd-arg=<cr>
*Jun 10 07:55:43.439: AAA/AUTHOR (2095255949): Post
authorization status = PASS_ADD
*Jun 10 07:55:43.439: AAA/MEMORY: free_user
(0x7FCBFD432410) user='bboyles' ruser='CAT9300'
port='tty1' rem_addr='10.1.100.99' authen_type=ASCII
service=NONE priv=15 vrf= (id=0)
CAT9300#
*Jun 10 07:56:05.009: AAA/AUTHOR (0xFCB): Pick method
list ' pvt_author_0'
*Jun 10 07:56:05.041: AAA/AUTHOR (0xFC9): Pick method
list ' pvt_author_0'
```

```
CAT9300#
*Jun 10 07:56:07.093: AAA/AUTHOR (0xFB7): Pick method
list ' pvt_author_0'
CAT9300#
*Jun 10 07:56:10.044: AAA/AUTHOR (0xFCA): Pick method
list ' pvt_author_0'
```

This output shows the AAA authorization process after user bboyles issues and is authorized to run the **show running-config** command. The output displays information on the AAA process, which details the AAA authorization method list being used for the request and the server group the request was sent to. The request was returned by the TACACS server with the authorization response PASS_ADD, which means the command was authorized.

Now, let's take a look at a command authorization that fails:

```
*Jun 10 07:58:41.109: AAA/AUTHOR: auth_need : user=
'bboyles' ruser= 'CAT9300'rem_addr= '10.1.100.99' priv=
15 list= 'VTY' AUTHOR-TYPE= 'commands'
*Jun 10 07:58:41.109: AAA: parse name=tty1 idb type=-1
tty=-1
*Jun 10 07:58:41.109: AAA: name=tty1 flags=0x11 type=5
shelf=0 slot=0 adapter=0 port=1 channel=0
*Jun 10 07:58:41.109: AAA/MEMORY: create_user
(0x7FCBFD432410) user='bboyles' ruser='CAT9300' ds0=0
port='tty1' rem_addr='10.1.100.99' authen_type=ASCII
service=NONE priv=15 initial_task_id='0', vrf= (id=0)
key=333EB499
*Jun 10 07:58:41.109: tty1 AAA/AUTHOR/CMD (3897019072):
Port='tty1' list='VTY' service=CMD
*Jun 10 07:58:41.109: AAA/AUTHOR/CMD: tty1 (3897019072)
user='bboyles'
*Jun 10 07:58:41.109: tty1 AAA/AUTHOR/CMD (3897019072):
send AV service=shell
*Jun 10 07:58:41.109: tty1 AAA/AUTHOR/CMD (3897019072):
send AV cmd=configure
*Jun 10 07:58:41.109: tty1 AAA/AUTHOR/CMD (3897019072):
send AV cmd-arg=terminal
*Jun 10 07:58:41.109: tty1 AAA/AUTHOR/CMD (3897019072):
send AV cmd-arg=<cr>
*Jun 10 07:58:41.109: tty1 AAA/AUTHOR/CMD(3897019072):
found list "VTY"
*Jun 10 07:58:41.109: tty1 AAA/AUTHOR/CMD (3897019072):
Method=ISE-TACACS (tacacs+)
```

```
*Jun 10 07:58:41.109: AAA/AUTHOR/TAC+: (3897019072):
user=bboyles
*Jun 10 07:58:41.109: AAA/AUTHOR/TAC+: (3897019072): send
AV service=shell
CAT9300#
*Jun 10 07:58:41.109: AAA/AUTHOR/TAC+: (3897019072): send
AV cmd=configure
*Jun 10 07:58:41.109: AAA/AUTHOR/TAC+: (3897019072): send
AV cmd-arg=terminal
*Jun 10 07:58:41.109: AAA/AUTHOR/TAC+: (3897019072): send
AV cmd-arg=<cr>
*Jun 10 07:58:41.310: AAA/AUTHOR (3897019072): Post
authorization status = FAIL
*Jun 10 07:58:41.310: AAA/MEMORY: free_user
(0x7FCBFD432410) user='bboyles' ruser='CAT9300'
port='tty1' rem_addr='10.1.100.99' authen_type=ASCII
service=NONE priv=15 vrf= (id=0)
```

In this case, instead of returning a PASS_ADD response, the TACACS+ server responded with FAIL, and the command was not authorized.

Step 5. To troubleshoot TACACS+ packet traces, use the **debug tacacs** command. The output of this command displays detailed information related to the interaction between the AAA client (IOS network device) and the AAA server (ISE). Let's look at some examples of output from this command.

The following output is from an SSH session initiated by user bboyles, who passes authentication and is placed into default privilege level 1:

```
CAT9300# debug tacacs


*Jun 10 08:31:13.006: TPLUS: Queuing AAA Authentication
request 4069 for processing
*Jun 10 08:31:13.006: TPLUS(00000FE5) login timer started
1020 sec timeout
*Jun 10 08:31:13.006: TPLUS: processing authentication
start request id 4069
*Jun 10 08:31:13.006: TPLUS: Authentication start packet
created for 4069(bboyles)
*Jun 10 08:31:13.006: TPLUS: Using server 10.1.100.21
*Jun 10 08:31:13.006: TPLUS(00000FE5)/0/NB_WAIT/
7FCC090E8D38: Started 5 sec timeout
*Jun 10 08:31:13.007: TPLUS(00000FE5)/0/NB_WAIT: socket
event 2
```

```
*Jun 10 08:31:13.007: TPLUS(00000FE5)/0/NB_WAIT: wrote
entire 42 bytes request
*Jun 10 08:31:13.007: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:13.007: TPLUS(00000FE5)/0/READ: Would block
while reading
*Jun 10 08:31:13.011: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:13.011: TPLUS(00000FE5)/0/READ: read entire
12 header bytes (expect 15 bytes data)
*Jun 10 08:31:13.012: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:13.012: TPLUS(00000FE5)/0/READ: read entire
27 bytes response
*Jun 10 08:31:13.012: TPLUS(00000FE5)/0/7FCC090E8D38:
Processing the reply packet
*Jun 10 08:31:13.012: TPLUS: Received authen response
status GET_PASSWORD (8)
*Jun 10 08:31:15.105: TPLUS: Queuing AAA Authentication
request 4069 for processing
*Jun 10 08:31:15.105: TPLUS(00000FE5) login timer started
1020 sec timeout
*Jun 10 08:31:15.105: TPLUS: processing authentication
continue request id 4069
*Jun 10 08:31:15.105: TPLUS: Authentication continue
packet generated for 4069
*Jun 10 08:31:15.105: TPLUS(00000FE5)/0/
WRITE/7FCC08A5A2C0: Started 5 sec timeout
*Jun 10 08:31:15.105: TPLUS(00000FE5)/0/WRITE: wrote
entire 24 bytes request
*Jun 10 08:31:15.141: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.141: TPLUS(00000FE5)/0/READ: read entire
12 header bytes (expect 6 bytes data)
*Jun 10 08:31:15.141: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.141: TPLUS(00000FE5)/0/READ: read entire
18 bytes response
*Jun 10 08:31:15.141: TPLUS(00000FE5)/0/7FCC08A5A2C0:
Processing the reply packet
```

```
*Jun 10 08:31:15.141: TPLUS: Received authen response
status PASS (2)
*Jun 10 08:31:15.142: %SEC_LOGIN-5-LOGIN_SUCCESS: Login
Success [user: bboyles] [Source: 10.1.100.99] [localport:
22] at 08:31:15 UTC Wed Jun 10 2020
*Jun 10 08:31:15.144: TPLUS: Queuing AAA Authorization
request 4069 for processing
*Jun 10 08:31:15.145: TPLUS(00000FE5) login timer started
1020 sec timeout
*Jun 10 08:31:15.145: TPLUS: processing authorization
request id 4069
*Jun 10 08:31:15.145: TPLUS: Protocol set to None .....
Skipping
*Jun 10 08:31:15.145: TPLUS: Sending AV service=shell
*Jun 10 08:31:15.145: TPLUS: Sending AV cmd*
*Jun 10 08:31:15.145: TPLUS: Authorization request
created for 4069(bboyles)
*Jun 10 08:31:15.145: TPLUS: using previously set server
10.1.100.21 from group ISE-TACACS
*Jun 10 08:31:15.145: TPLUS(00000FE5)/0/NB_
WAIT/7FCC08A5A2C0: Started 5 sec timeout
*Jun 10 08:31:15.145: TPLUS(00000FE5)/0/NB_WAIT: socket
event 2
*Jun 10 08:31:15.145: TPLUS(00000FE5)/0/NB_WAIT: wrote
entire 61 bytes request
*Jun 10 08:31:15.145: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.146: TPLUS(00000FE5)/0/READ: Would block
while reading
*Jun 10 08:31:15.163: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.163: TPLUS(00000FE5)/0/READ: read entire
12 header bytes (expect 17 bytes data)
*Jun 10 08:31:15.163: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.163: TPLUS(00000FE5)/0/READ: read entire
29 bytes response
*Jun 10 08:31:15.163: TPLUS(00000FE5)/0/7FCC08A5A2C0:
Processing the reply packet
*Jun 10 08:31:15.163: TPLUS: Processed AV priv-lvl=1
```

```
*Jun 10 08:31:15.163: TPLUS: received authorization
response for 4069: PASS
*Jun 10 08:31:15.163: TPLUS: Queuing AAA Accounting
request 4069 for processing
*Jun 10 08:31:15.164: TPLUS: processing accounting
request id 4069
*Jun 10 08:31:15.164: TPLUS: Sending AV task_id=4080
*Jun 10 08:31:15.164: TPLUS: Sending AV timezone=UTC
*Jun 10 08:31:15.164: TPLUS: Sending AV service=shell
*Jun 10 08:31:15.164: TPLUS: Accounting request created
for 4069(bboyles)
*Jun 10 08:31:15.164: TPLUS: using previously set server
10.1.100.21 from group ISE-TACACS
*Jun 10 08:31:15.164: TPLUS(00000FE5)/0/NB_WAIT/
7FCC08A5A2C0: Started 5 sec timeout
*Jun 10 08:31:15.164: TPLUS(00000FE5)/0/NB_WAIT: socket
event 2
*Jun 10 08:31:15.164: TPLUS(00000FE5)/0/NB_WAIT: wrote
entire 83 bytes request
*Jun 10 08:31:15.164: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.164: TPLUS(00000FE5)/0/READ: Would block
while reading
*Jun 10 08:31:15.167: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.167: TPLUS(00000FE5)/0/READ: read entire
12 header bytes (expect 5 bytes data)
*Jun 10 08:31:15.167: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:31:15.167: TPLUS(00000FE5)/0/READ: read entire
17 bytes response
*Jun 10 08:31:15.168: TPLUS(00000FE5)/0/7FCC08A5A2C0:
Processing the reply packet
*Jun 10 08:31:15.168: TPLUS: Received accounting response
with status PASS
```

The following output results when user bboyles issues the **show running-config** command and it passes command authorization:

```
*Jun 10 08:36:10.131: TAC+: using previously set server
10.1.100.21 from group ISE-TACACS
*Jun 10 08:36:10.131: TAC+: Opening TCP/IP to
10.1.100.21/49 timeout=5
```

```
*Jun 10 08:36:10.132: TAC+: Opened TCP/IP handle
0x7FCC0B33DF10 to 10.1.100.21/49 using source 0.0.0.0
*Jun 10 08:36:10.132: TAC+: Opened 10.1.100.21 index=1
*Jun 10 08:36:10.132: TAC+: 10.1.100.21
(18446744072565125362) AUTHOR/START queued
*Jun 10 08:36:10.332: TAC+: (18446744072565125362)
AUTHOR/START processed
*Jun 10 08:36:10.332: TAC+: (-1144426254): received
author response status = PASS_ADD
*Jun 10 08:36:10.332: TAC+: Closing TCP/IP 0x7FCC0B33DF10
connection to 10.1.100.21/49
*Jun 10 08:36:10.333: TPLUS: Queuing AAA Accounting
request 4069 for processing
*Jun 10 08:36:10.333: TPLUS: processing accounting
request id 4069
*Jun 10 08:36:10.333: TPLUS: Sending AV task_id=4080
*Jun 10 08:36:10.333: TPLUS: Sending AV timezone=UTC
*Jun 10 08:36:10.333: TPLUS: Sending AV service=shell
*Jun 10 08:36:10.333: TPLUS: Sending AV priv-lvl=15
*Jun 10 08:36:10.333: TPLUS: Sending AV cmd=show
running-config <cr>
*Jun 10 08:36:10.334: TPLUS: Accounting request created
for 4069(bboyles)
*Jun 10 08:36:10.334: TPLUS: using previously set server
10.1.100.21 from group ISE-TACACS
*Jun 10 08:36:10.334: TPLUS(00000FE5)/0/NB_
WAIT/7FCC0B330FE8: Started 5 sec timeout
*Jun 10 08:36:10.339: TPLUS(00000FE5)/0/NB_WAIT: socket
event 2
*Jun 10 08:36:10.340: TPLUS(00000FE5)/0/NB_WAIT: wrote
entire 124 bytes request
*Jun 10 08:36:10.340: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:36:10.340: TPLUS(00000FE5)/0/READ: Would block
while reading
*Jun 10 08:36:10.344: TPLUS(00000FE5)/0/READ: socket
event 1
*Jun 10 08:36:10.344: TPLUS(00000FE5)/0/READ: read entire
12 header bytes (expect 5 bytes data)
*Jun 10 08:36:10.344: TPLUS(00000FE5)/0/READ: socket
event 1
```

```
*Jun 10 08:36:10.344: TPLUS(00000FE5)/0/READ: read entire
17 bytes response
*Jun 10 08:36:10.344: TPLUS(00000FE5)/0/7FCC0B330FE8:
Processing the reply packet
*Jun 10 08:36:10.344: TPLUS: Received accounting response
with status PASS
```

The following output results when user bboyles issues the **configure terminal** command and the command fails authorization:

```
*Jun 10 08:38:03.800: TAC+: using previously set server
10.1.100.21 from group ISE-TACACS
*Jun 10 08:38:03.800: TAC+: Opening TCP/IP to
10.1.100.21/49 timeout=5
*Jun 10 08:38:03.801: TAC+: Opened TCP/IP handle
0x7FCC0B006078 to 10.1.100.21/49 using source 0.0.0.0
*Jun 10 08:38:03.801: TAC+: Opened 10.1.100.21 index=1
*Jun 10 08:38:03.801: TAC+: 10.1.100.21
(18446744073154346637) AUTHOR/START queued
*Jun 10 08:38:04.002: TAC+: (18446744073154346637)
AUTHOR/START processed
*Jun 10 08:38:04.002: TAC+: (-555204979): received author
response status = FAIL
*Jun 10 08:38:04.002: TAC+: Closing TCP/IP 0x7FCC0B006078
connection to 10.1.100.21/49
```

Step 6. To debug TACACS+ packets, issue the **debug tacacs packet** command. This example shows the output when user bboyles authenticates to the IOS device:

```
CAT9300# debug tacacs packets


*Jun 10 09:17:06.988: T+: Version 192 (0xC0), type 1,
seq 1, encryption 1, SC 0
*Jun 10 09:17:06.988: T+: session_id 1243645340
(0x4A20859C), dlen 30 (0x1E)
*Jun 10 09:17:06.988: T+: type:AUTHEN/START, priv_lvl:1
action:LOGIN ascii
*Jun 10 09:17:06.988: T+: svc:LOGIN user_len:7 port_len:4
(0x4) raddr_len:11 (0xB) data_len:0
*Jun 10 09:17:06.988: T+: user:  bboyles
*Jun 10 09:17:06.988: T+: port:  tty1
*Jun 10 09:17:06.988: T+: rem_addr:  10.1.100.99
*Jun 10 09:17:06.988: T+: data:
*Jun 10 09:17:06.988: T+: End Packet
```

```
*Jun 10 09:17:06.993: T+: Version 192 (0xC0), type 1,
seq 2, encryption 1, SC 0
*Jun 10 09:17:06.993: T+: session_id 1243645340
(0x4A20859C), dlen 15 (0xF)
*Jun 10 09:17:06.993: T+: AUTHEN/REPLY status:5 flags:0x1
msg_len:9, data_len:0
*Jun 10 09:17:06.993: T+: msg:  Password:
*Jun 10 09:17:06.993: T+: data:
*Jun 10 09:17:06.993: T+: End Packet
*Jun 10 09:17:13.915: T+: Version 192 (0xC0), type 1,
seq 3, encryption 1, SC 0
*Jun 10 09:17:13.915: T+: session_id 1243645340
(0x4A20859C), dlen 12 (0xC)
*Jun 10 09:17:13.915: T+: AUTHEN/CONT msg_len:7 (0x7),
data_len:0 (0x0) flags:0x0
*Jun 10 09:17:13.915: T+: User msg: <elided>
*Jun 10 09:17:13.915: T+: User data:
*Jun 10 09:17:13.915: T+: End Packet
*Jun 10 09:17:13.958: T+: Version 192 (0xC0), type 1,
seq 4, encryption 1, SC 0
*Jun 10 09:17:13.958: T+: session_id 1243645340
(0x4A20859C), dlen 6 (0x6)
*Jun 10 09:17:13.958: T+: AUTHEN/REPLY status:1 flags:0x0
msg_len:0, data_len:0
*Jun 10 09:17:13.958: T+: msg:
*Jun 10 09:17:13.958: T+: data:
*Jun 10 09:17:13.958: T+: End Packet
*Jun 10 09:17:13.958: %SEC_LOGIN-5-LOGIN_SUCCESS: Login
Success [user: bboyles] [Source: 10.1.100.99] [localport:
22] at 09:17:13 UTC Wed Jun 10 2020
*Jun 10 09:17:13.961: T+: Version 192 (0xC0), type 2,
seq 1, encryption 1, SC 0
*Jun 10 09:17:13.961: T+: session_id 1148614991
(0x4476794F), dlen 49 (0x31)
*Jun 10 09:17:13.961: T+: AUTHOR, priv_lvl:1, authen:1
method:tacacs+
*Jun 10 09:17:13.961: T+: svc:1 user_len:7 port_len:4
rem_addr_len:11 arg_cnt:2
*Jun 10 09:17:13.961: T+: user:  bboyles
*Jun 10 09:17:13.961: T+: port:  tty1
*Jun 10 09:17:13.961: T+: rem_addr:  10.1.100.99
```

```
*Jun 10 09:17:13.961: T+: arg[0]: size:13 service=shell
*Jun 10 09:17:13.961: T+: arg[1]: size:4 cmd*
*Jun 10 09:17:13.962: T+: End Packet
*Jun 10 09:17:13.975: T+: Version 192 (0xC0), type 2,
seq 2, encryption 1, SC 0
*Jun 10 09:17:13.975: T+: session_id 1148614991
(0x4476794F), dlen 17 (0x11)
*Jun 10 09:17:13.975: T+: AUTHOR/REPLY status:1
msg_len:0, data_len:0 arg_cnt:1
*Jun 10 09:17:13.975: T+: msg:
*Jun 10 09:17:13.975: T+: data:
*Jun 10 09:17:13.976: T+: arg[0] size:10
*Jun 10 09:17:13.976: T+:  priv-lvl=1
*Jun 10 09:17:13.976: T+: End Packet
*Jun 10 09:17:13.977: T+: Version 192 (0xC0), type 3,
seq 1, encryption 1, SC 0
*Jun 10 09:17:13.977: T+: session_id 2881630759
(0xABC23227), dlen 71 (0x47)
*Jun 10 09:17:13.977: T+: ACCT, flags:0x2 method:6
priv_lvl:1
*Jun 10 09:17:13.977: T+: type:1 svc:1 user_len:7 port_
len:4 rem_addr_len:11
*Jun 10 09:17:13.977: T+: arg_cnt:3
*Jun 10 09:17:13.977: T+: user:  bboyles
*Jun 10 09:17:13.977: T+: port:  tty1
*Jun 10 09:17:13.977: T+: rem_addr:  10.1.100.99
*Jun 10 09:17:13.977: T+: arg[0]: size:12 task_id=4100
*Jun 10 09:17:13.977: T+: arg[1]: size:12 timezone=UTC
*Jun 10 09:17:13.977: T+: arg[2]: size:13 service=shell
*Jun 10 09:17:13.977: T+: End Packet
*Jun 10 09:17:13.980: T+: Version 192 (0xC0), type 3,
seq 2, encryption 1, SC 0
*Jun 10 09:17:13.980: T+: session_id 2881630759
(0xABC23227), dlen 5 (0x5)
*Jun 10 09:17:13.980: T+: ACCT/REPLY status:1 msg_len:0
data_len:0
*Jun 10 09:17:13.980: T+: msg:
*Jun 10 09:17:13.980: T+: data:
*Jun 10 09:17:13.980: T+: End Packet
```

Step 7. To troubleshoot the TCP connection between the IOS device and ISE, issue the **debug ip tcp transactions** command. The following is a sample of the debug output showing the TCP connection between the switch and ISE:

NOTE The TACACS+ protocol uses TCP as the transport protocol with destination port 49.

```
CAT9300# debug ip tcp transactions


*Jun 10 09:26:25.691: TCP0: state was ESTAB -> CLOSEWAIT
[13709 -> 10.1.100.21(49)]
*Jun 10 09:26:25.691: TCP0: state was CLOSEWAIT ->
LASTACK [13709 -> 10.1.100.21(49)]
*Jun 10 09:26:25.691: TCP0: sending FIN
*Jun 10 09:26:25.692: %SEC_LOGIN-5-LOGIN_SUCCESS: Login
Success [user: bboyles] [Source: 10.1.100.99] [localport:
22] at 09:26:25 UTC Wed Jun 10 2020
*Jun 10 09:26:25.692: TCP0: Got ACK for our FIN
*Jun 10 09:26:25.692: TCP0: state was LASTACK -> CLOSED
[13709 -> 10.1.100.21(49)]
```

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 25-5 lists these key topics and the page number on which each is found.

Key Topic

Table 25-5 Key Topics

Key Topic Element	Description	Page Number
Section	Default shell privilege levels	932
Section	TACACS command sets	936
List	Configuration of TACACS+ authentication on a switch	946
List	Configuration of TACACS+ command authorization on a switch	948
List	Configuration of TACACS+ command accounting on a switch	951

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

privilege level, command set

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What ISE TACACS+ policy element is used to assign the shell privilege level?
2. What types of TACACS+ allowed protocols can be defined?
3. When is Deny Always used in a command set?
4. What are the three main parts of configuring TACACS+ on IOS?
5. What is the advantage of using command sets instead of defining custom privilege levels?



CHAPTER 26

Configuring Device Admin AAA with the Cisco WLC

This chapter covers the following topics:

- Configure network access devices
- Configure policies including authentication and authorization profiles
- Compare AAA protocols
- Configure TACACS+ device administration and command authorization

In Chapter 25, "Configuring Device Administration AAA with Cisco IOS," you learned about a network device that has very granular control all the way down to individual commands. Keep in mind that not all network devices have the capability to control access at a per-command level. Many provide a more broad-stroke access control, such as the assignment of roles.

Cisco Wireless LAN Controller (WLC) provides a broad approach to role-based access control by controlling access at a menu level.

NOTE This chapter focuses on the Cisco WLC, which manages and controls Cisco's lightweight access points. Cisco's autonomous access points run Cisco IOS and follow the same device administration as Cisco Catalyst switches and Cisco routers, covered in Chapter 25.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the "Exam Preparation Tasks" section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 26-1 lists the major headings in this chapter and their corresponding "Do I Know This Already?" quiz questions. You can find the answers in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections."

Table 26-1 "Do I Know This Already?" Section-to-Question Mapping

Foundation Topics Section	Questions
Overview of WLC Device Administration AAA	1, 3, 8, 10
Configure ISE and the WLC for Device Administration AAA	4–6, 9
Testing and Troubleshooting	2, 7

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following best describes the way Cisco WLC provides device admin AAA?
   a. WLC provides role-based access control (RBAC) on a per-menu basis. Only the menus that an administrator has access to are displayed, and all others are hidden.
   b. WLC provides role-based access control (RBAC) on a per-menu basis. A WLC administrator sees all configuration items in that menu, but any denied settings are grayed out.
   c. WLC provides role-based access control (RBAC) on a per-menu basis. A WLC administrator sees all configuration items in that menu, but authorization is performed before changes are applied.
   d. WLC provides role-based access control (RBAC) on a per-command basis. A WLC administrator sees all menus, but any commands that are denied are not displayed in the menu.
2. Which of the following statements about device administration AAA and Cisco WLC is most accurate?
   a. The priority order applies to access to the WLC web interface and the SSH interface simultaneously.
   b. You must configure a different access class for each method of access: HTTP, SSH, and console.
   c. You enable TACACS for AAA by assigning it in priority order, and then you must apply the priority for each access method (HTTP, SSH, and console).
   d. Device administration AAA is only applicable to the web interface and not to SSH or console access.
3. Which of the following role assignments is valid?
   a. WLAN, SECURITY, MONITOR
   b. WLAN, MONITOR, MANAGEMENT
   c. ALL
   d. WLAN, LOBBY
4. Which of the following statements about device administration policy sets is true?
   a. Command sets limit which options in WLC work for users.
   b. TACACS profiles are used for limiting the commands that are authorized.
   c. RADIUS profiles are used for controlling access to WLC.
   d. Command sets are not used with WLC.

5. Which of the following statements about TACACS profiles is accurate?
   a. Switching between the Raw view and the Task Attribute view requires you to reenter the values.
   b. You can switch seamlessly between the Task Attribute view and the Raw view without losing your configuration.
   c. TACACS profiles are used for per-command authorization.
   d. RADIUS profiles are used for per-command authorization.
6. Which of the following statements about device administration policy sets is true?
   a. Policy sets can be disabled under Work Centers > Device Administration > Settings.
   b. RADIUS device administration requires its own device administration policy set.
   c. Shell (IOS) and WLC rules can be in the same authorization policy.
   d. Cisco WLC cannot be used with a device administration policy set.
7. When troubleshooting, the Identity column in TACACS Live Log displays only the string USERNAME. Where do you correct this behavior?
   a. Administration > System > Settings > Security Settings > Hide Failed Usernames
   b. Work Centers > Device Administration > Settings > Hide Failed Usernames
   c. Work Centers > Device Administration > Settings > Disclose Invalid Usernames
   d. Administration > System > Settings > Security Settings > Disclose Invalid Usernames
8. Which of the following is true of the MONITOR role type?
   a. All menus and settings are visible to the user, but no changes are authorized.
   b. You can see only see the Dashboard page and nothing from the Advanced UI.
   c. All menus and settings are visible to the user, and actions can be taken on the client's page to disconnect endpoints from the Wi-Fi, but no other configuration changes are authorized.
   d. No UI access is granted; only SNMP is allowed.
9. Where must ISE be added to the WLC for device administration AAA?
   a. Security > Device Access AAA: authentication, authorization, and accounting lists
   b. Security > AAA > TACACS+: authentication, authorization, and accounting lists
   c. Security > AAA > RADIUS: authentication, authorization, and accounting lists
   d. Administration > AAA > TACACS+: authentication, authorization, and accounting lists
10. Which statement best describes the LOBBY role type?
    a. The role is granted access to the wireless dashboard with read-only access to all settings.
    b. It is a reserved role that is assigned to publicly accessible APs, such as those installed in a lobby.
    c. It is a role that provides access to WLC for a user to select which WLAN to connect to, much like a building lobby with elevator access.
    d. It is a role that provides access to the lobby ambassador functions, where the WLC performs its own guest lifecycle management.

Foundation Topics

Overview of WLC Device Administration AAA

WLC provides role-based access control (RBAC) on a per-menu basis. However, WLC does not prevent access to those menus by locking the user out of those sections of the GUI; instead, WLC prevents changes from being saved when a user is inside a menu section that is not authorized.

Figure 26-1 shows the WLC user interface. The top-level menus, highlighted in the figure, are where the RBAC is applicable.

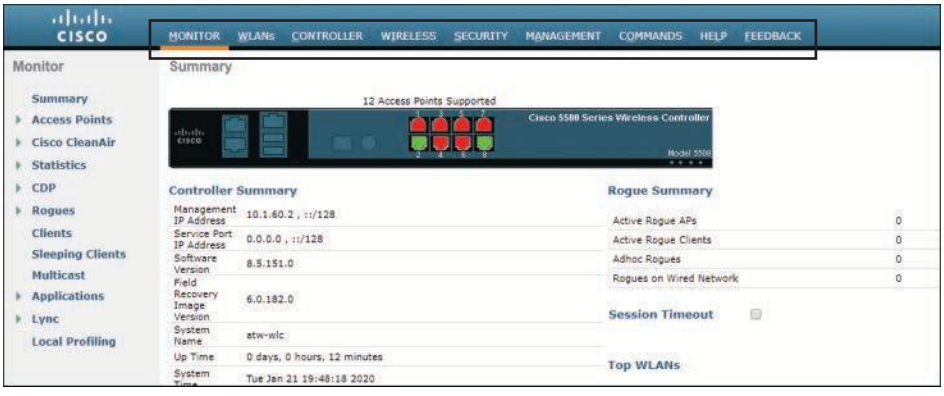


Figure 26-1 *Cisco WLC Top-Level Menus*

WLC expects the authorization results from ISE to include a list of the menus in which the user is authorized to make changes. WLC refers to those results as *roles*.

The AAA server can send back a single role or multiple roles that will be additive, as long as each role is designated with a unique role number. For example, to allow a user to make changes only within the WLANs, Wireless, and Security menus, the TACACS+ authorization result should include the following text: role1=WLANS role2=SECURITY role3=WIRELESS. The good news is that the order does not matter.

Key Topic

Six different roles exist, and they correspond directly to the menu system in the WLC user interface:

- **WLAN**: Permits write access to the WLAN menu structure.
- **CONTROLLER**: Permits write access to the Controller menu structure.
- **WIRELESS**: Permits write access to the Wireless menu structure.
- **SECURITY**: Permits write access to the Security menu structure.
- **MANAGEMENT**: Permits write access to the Management menu structure.
- **COMMANDS**: Permits write access to the Commands menu structure.

In addition, there are three special roles that cannot be mixed with any other roles:

- **LOBBY**: Permits access to the lobby ambassador functions, which remain from when WLC performed its own guest lifecycle management. The lobby ambassador functions on WLC have since been superseded by the guest capabilities in ISE.
- **MONITOR**: Provides a read-only style of access to the WLC. All pages are visible to the user, but no changes are authorized.
- **ALL**: Just as it sounds, permits access to all menus.

Configure ISE and the WLC for Device Administration AAA

The steps required to configure device administration AAA are very similar on different network devices. You need to perform some configuration on ISE to ensure that it is ready to receive the TACACS+ requests from the NAD, and you must configure the NAD to send those TACACS+ requests to ISE.

Prepare ISE for WLC Device Administration AAA

In the ISE GUI, you need to ensure that the network device object is configured correctly and assigned to appropriate network device groups (NDGs). Then the authorization results should be created, and then the policy set.

Prepare the Network Device

Navigate to Work Centers > Device Administration > Network Resources > Network Device Groups.

You should already have a detailed set of hierarchical NDGs, similar to what you see in Figure 26-2. Ensure that you have the appropriate groups to go along with the policies you want to create.

At the very least, ensure that these Cisco WLCs have their own group for a device type. With that group, you can ensure that a device administration policy set is built just for these devices and their unique way of doing device administration.

Next, you should ensure that Cisco WLC has a network device object in ISE with the TACACS+ shared secret configured.

Navigate to Work Centers > Device Administration > Network Resources > Network Devices and click on the WLC or click Add to create a new object.

Ensure that the network device groups are assigned properly and that the TACACS+ shared secret is configured correctly. Figure 26-3 shows the NAD settings.

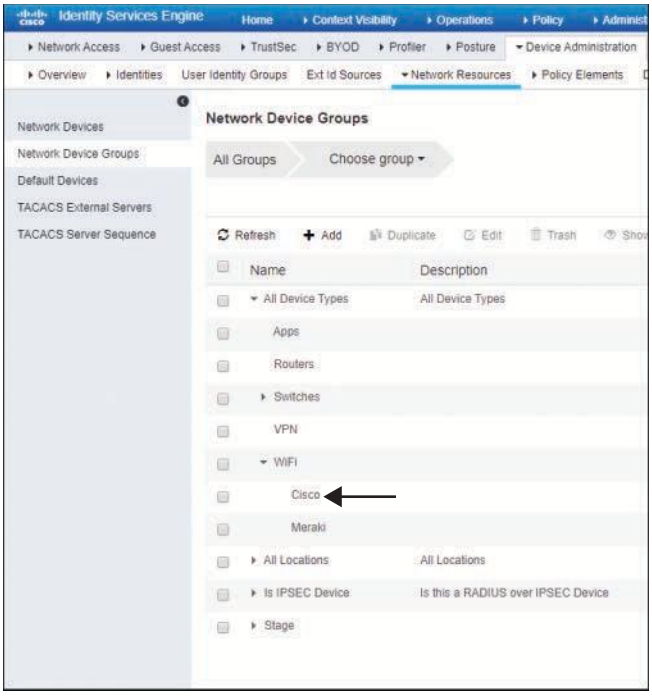


Figure 26-2 *NDG Hierarchy Example*

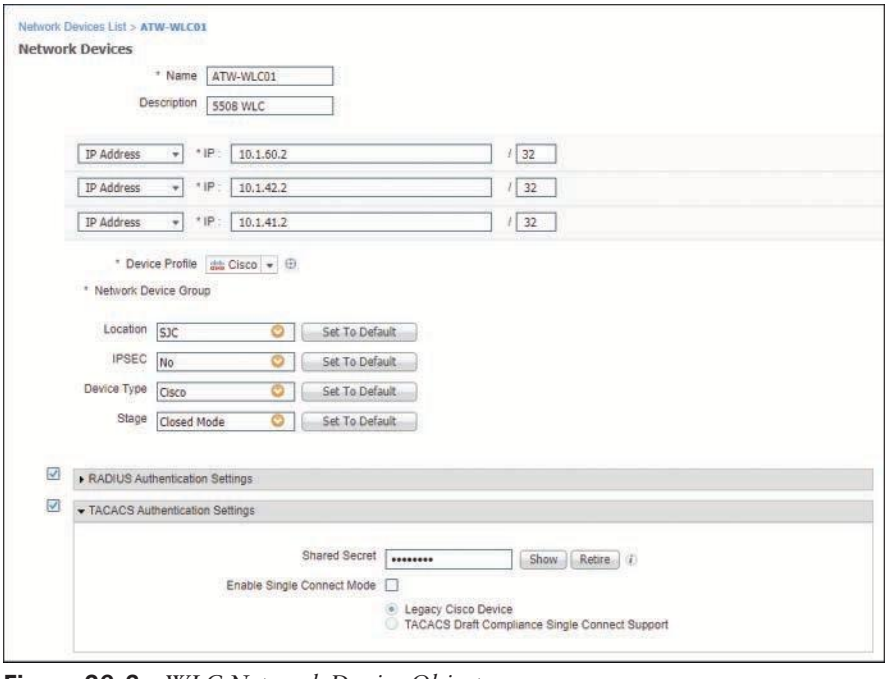


Figure 26-3 *WLC Network Device Object*

Prepare the Policy Results

If you have followed along, device administration AAA is ready on ISE, but you have no policies and no authorizations to send to WLC yet. Now you're going to add one TACACS profile for each role that you plan to use with WLC.

Here you will create the same three roles for WLC that were created in Chapter 25: NetAdmin, SecAdmin, and Helpdesk. Finally, you will create a Lobby Ambassador result to apply for any other employee who attempts to log in.

Navigate to Work Centers > Device Administration > Policy Elements > Results > TACACS Profiles (see Figure 26-4). Notice that there are predefined TACACS profiles: WLC ALL and WLC MONITOR.

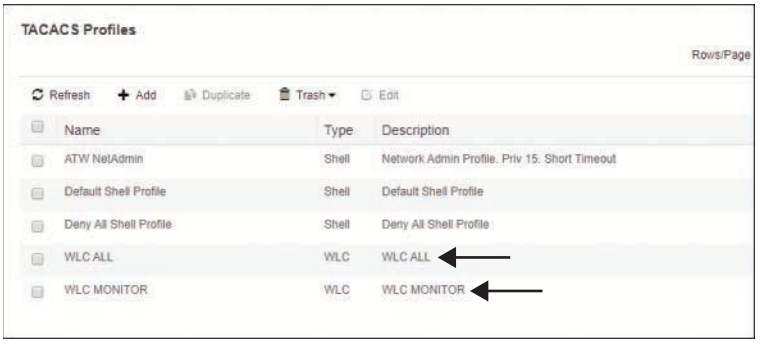


Figure 26-4 *Predefined TACACS Profiles for the WLC*

Create the NetAdmin Profile

Follow these steps to create a new TACACS profile to use as the result when a network administrator logs in to the WLC:

Step 1. Click **Add**.

Step 2. Name the profile **WLC NetAdmin**. It is always best to prefix the profile name with the device on which you will use it. Following this convention makes it easy to know which result is for WLC and which result is for IOS devices.

Step 3. Set Common Task Type to **WLC**. Notice that the UI changes to align with the UI of the WLC menu structure. You have a choice to select from the three special roles (All, Monitor, Lobby) or to select the individual menus. (Thank you, Doug Gash, for this very cool UI enhancement, which helps the experience of configuring the WLC TACACS profiles make sense to someone who is familiar with the WLC!)

Step 4. Select **All**. As you can see in Figure 26-5, the text "The configured options give a mgmtRole Debug value of: 0xfffffff8" appears. This text is calling out a mngtRole field that will appear in debug logs in WLC to help you troubleshoot if needed. (This is another wonderful user experience that we can thank Doug Gash for inventing.)

Step 5. Click **Submit**.

Figure 26-5 shows the completed WLC NetAdmin profile.

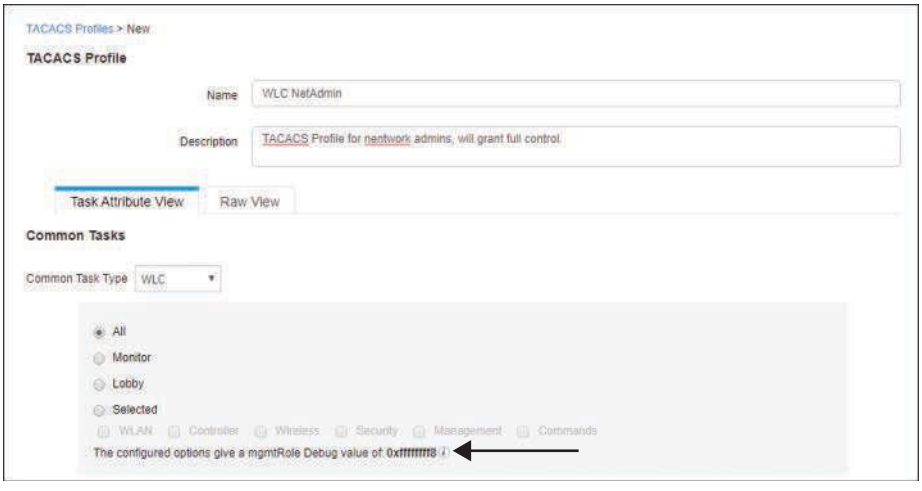


Figure 26-5 *Completed NetAdmin TACACS Profile*

Create the SecAdmin Profile

Follow these steps to create a new TACACS profile to use as the result when a security administrator logs in to the WLC and should only have access to the WLANs and Security menus:

Step 1. Click **Add**.

Step 2. Name the profile **WLC SecAdmin**.

Step 3. Set Common Task Type to **WLC**.

Step 4. Select **WLAN**.

Step 5. Select **Security**.

Step 6. Click **Submit**.

Figure 26-6 shows the completed WLC SecAdmin profile.

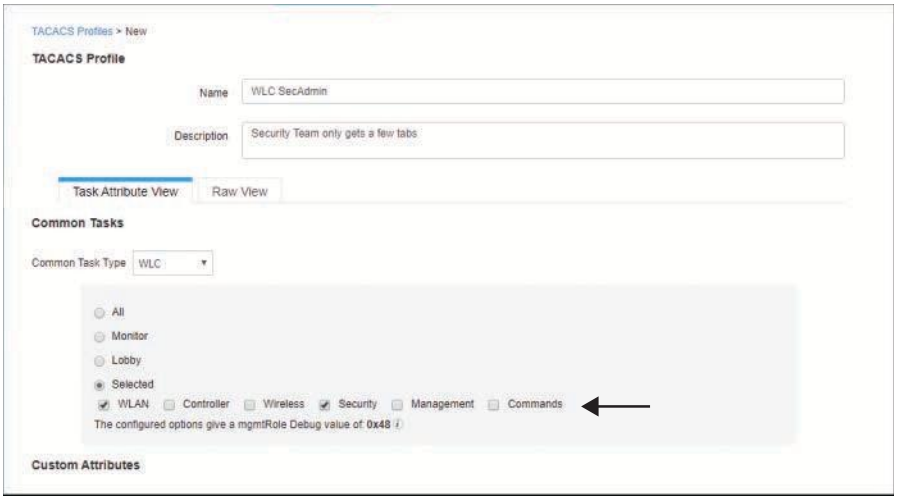


Figure 26-6 *Completed SecAdmin TACACS Profile*

Create the Helpdesk Profile

Follow these steps to create a profile for the members of the Helpdesk group in AD, which will be limited to the monitor result only:

Step 1. Click **Add**.

Step 2. Name the profile **WLC Helpdesk**.

Step 3. Set Common Task Type to **WLC**.

Step 4. Select **Monitor**.

Step 5. Click **Submit**.

Figure 26-7 shows the completed WLC Helpdesk profile.

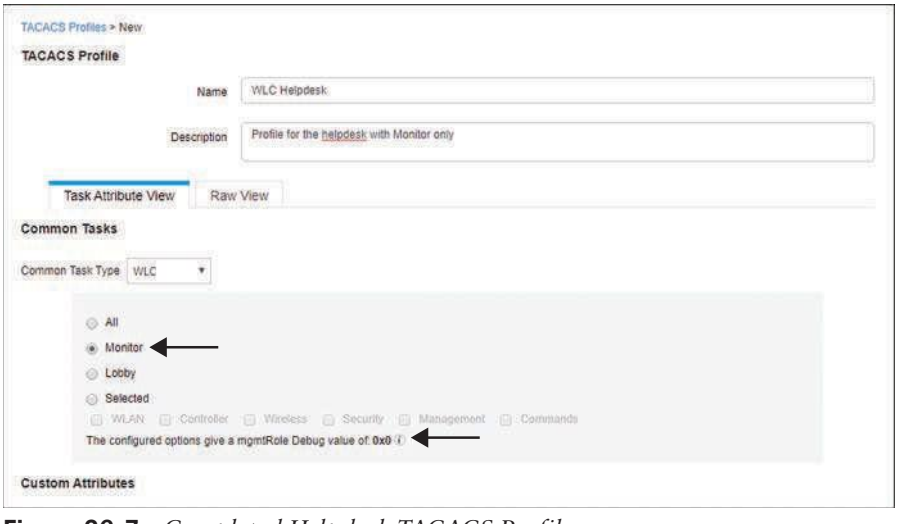


Figure 26-7 *Completed Helpdesk TACACS Profile*

Create the Employee Profile

Finally, you can create a profile for any employee of the company who doesn't match one of the previous groups. Those employees will be given lobby ambassador privileges on the WLC for creating guest accounts. This is not very relevant for this book or the Implementing and Configuring Cisco Identity Services Engine SISE 300-715 exam, since ISE is the guest server in this case, but we are simply using it as an example. Follow these steps:

Step 1. Click **Add**.

Step 2. Name the profile **WLC Employees**.

Step 3. Set Common Task Type to **WLC**.

Step 4. Select **Lobby**.

Step 5. Click **Submit**.

Figure 26-8 shows the completed WLC Employees profile.



Figure 26-8 *Completed Employees TACACS Profile*

Configure the Policy Set

When you have all the role types you need, it's time to configure the device administration policy set for the WLCs.

To add the authorization rule for network administrators, follow these steps:

Step 1. Navigate to **Work Centers > Device Administration > Device Admin Policy Sets**.

Step 2. Select the previously created **Wireless Controllers** policy set.

Step 3. Expand the **Authorization Policy** section.

Step 4. Insert a new authorization rule above the default rule by clicking the + sign and name it **NetAdmin**.

Step 5. Set the condition to be an external group from AD, as shown in Figure 26-9.

Step 6. Because there are no command sets for the WLC, ignore the Command Sets option.

Step 7. For the shell profile, select **WLC NetAdmin**.

Step 8. Click **Save**.

Figure 26-9 shows the completed NetAdmin WLC authorization rule.

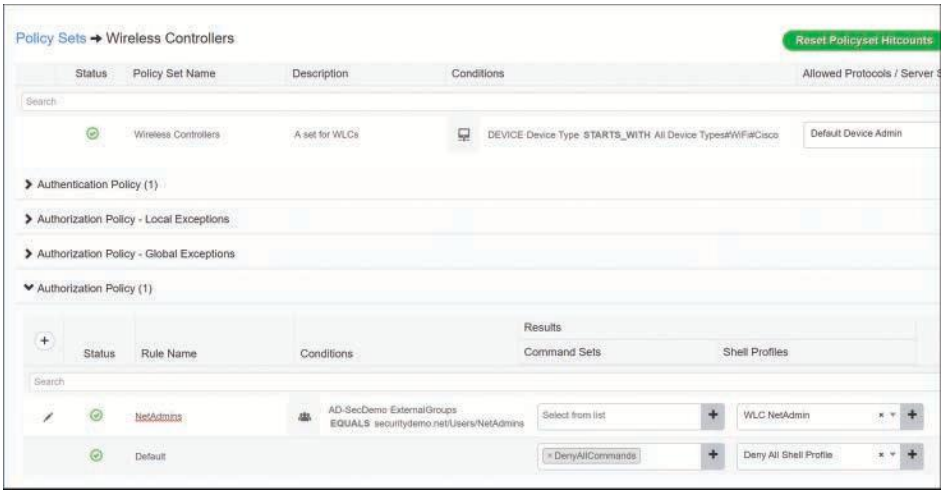


Figure 26-9 *NetAdmin WLC Authorization Rule*

To add the authorization rule for security operators, follow these steps:

Step 1. Insert a rule above the default rule and name it **SecOps**.

Step 2. Set the condition to be an external group from AD.

Step 3. For the shell profile, select **WLC SecAdmin**.

To add the authorization rule for the helpdesk, follow these steps:

Step 1. Insert a rule above the default rule and name it **Helpdesk**.

Step 2. Set the condition to be an external group from AD.

Step 3. For the shell profile, select **WLC Helpdesk**.

To add the authorization rule for the remainder of the employees to get LobbyAmbassador access, follow these steps:

Step 1. Insert a rule above the default rule and name it **Lobby Ambassador**.

Step 2. Set the condition to be an external group from AD.

Step 3. For the shell profile, select **WLC Employees**.

Step 4. Click **Done**.

Step 5. Click **Save**.

Figure 26-10 shows the completed wireless controllers policy set.

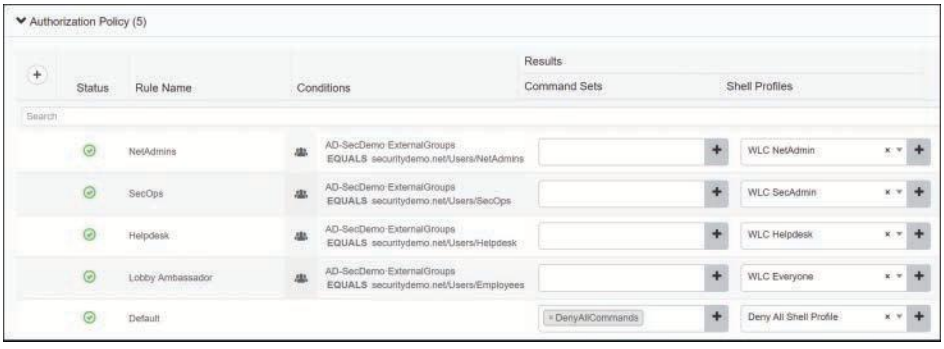


Figure 26-10 *Completed Wireless Controllers Policy Set*

26

Adding ISE to the WLC TACACS+ Servers

When the policy set is ready, it's time to start configuring WLC to authenticate and authorize interactive users. ISE needs to be added to WLC as a TACACS+ authentication server, authorization server, and accounting server.

The TACACS+ servers are configured in the WLC user interface under Security > AAA > TACACS+. Follow these steps:

Step 1. Navigate to **Security > AAA > TACACS+ > Authentication**.

Step 2. Click **New**.

Step 3. Fill in the Server IP Address and the Shared Secret text boxes as shown in Figure 26-11.

Step 4. Click **Apply**.

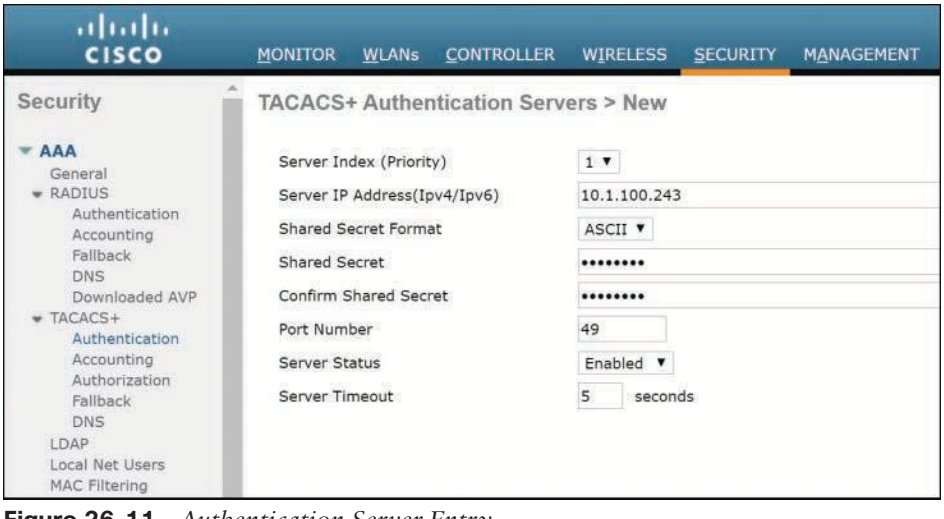


Figure 26-11 *Authentication Server Entry*

The WLC automatically creates an entry for authorization servers when you add an authentication server, so there is no need to add the authorization server separately.

Step 5. Navigate to **Security > AAA > TACACS+ > Accounting**.

Step 6. Click **New**.

Step 7. Fill in the Server IP Address and the Shared Secret text boxes as shown in Figure 26-12.

Step 8. Click **Apply**.

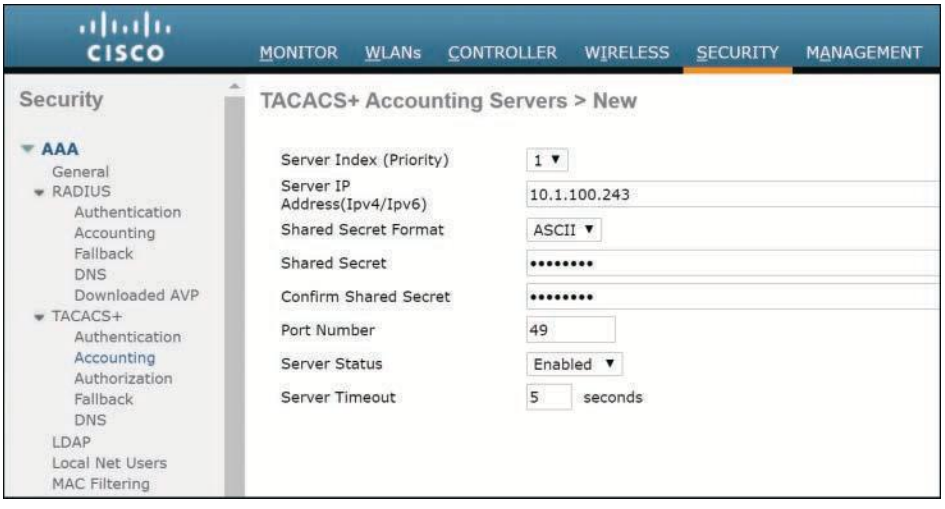


Figure 26-12 *Accounting Server Entry*

Repeat this process for each of the ISE PSNs running TACACS+.

After the ISE PSNs have been added to WLC for use with TACACS+, the authorization policy has been built, and the TACACS+ servers are defined in WLC, the next step is to configure WLC to use TACACS+ for administrative access. Follow these steps:

Step 1. Navigate to **Security > Priority Order > Management User**.

Step 2. Ensure that TACACS+ is at the top of the list titled Order Used for Authentication, as shown in Figure 26-13.

Step 3. Click **Apply**.

Key Topic

By putting TACACS+ at the top of the list and leaving LOCAL below it, you are providing yourself with a failsafe mechanism. WLC will use the local accounts (such as the built-in administrative account) in the event that the TACACS+ servers are not reachable.

WLC is now configured to use TACACS+ for logins, and you can test it.

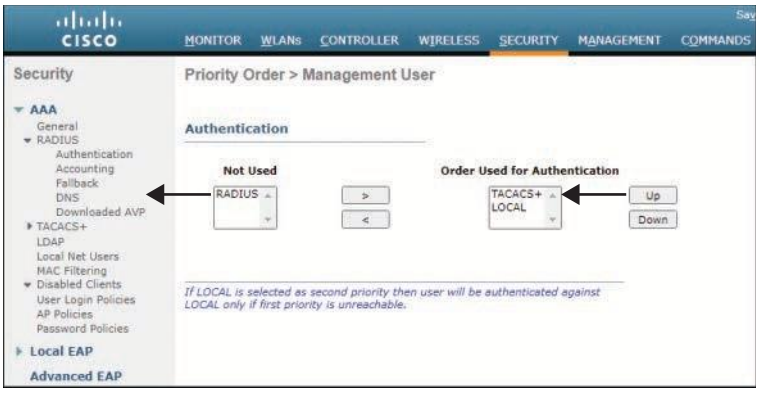


Figure 26-13 *Management Order*

26

Testing and Troubleshooting

Before testing WLC, it is highly recommended that you change a security setting that was added in ISE Version 2.4 patch 3, which remains a default security setting today. Once upon a time, someone performed a security audit on ISE and determined that it was an offensive security vulnerability to show the username of failed authentications in ISE's Live Log—because it's possible that a user might accidently type his or her password into the username field, so therefore displaying that text in the log would be a violation of security. Therefore, by default, any failed authentications will show USERNAME in the Identity column Live Log unless you enable the Disclose Invalid Usernames setting under Administration > Settings > Security Settings. Figure 26-14 shows this setting.

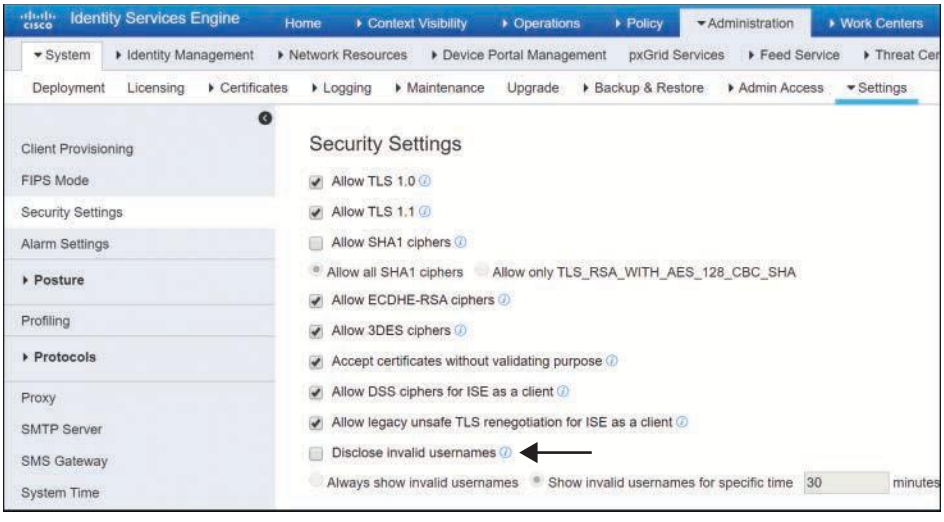


Figure 26-14 *The Setting That Should Not Be Set*

It should be very easy to test the logins from WLC: Just try to log in from a browser. Up to this point, you have been authenticating to WLC by using a LOCAL user account (admin). However, you configured the authentication order to use TACACS and then LOCAL, and the admin user account won't exist in Active Directory or ISE.

Open a browser window, bring up the WLC UI, and follow these steps:

Step 1. Log in as admin. Figure 26-15 shows an attempted login using the same admin account used to configure WLC up to this point.

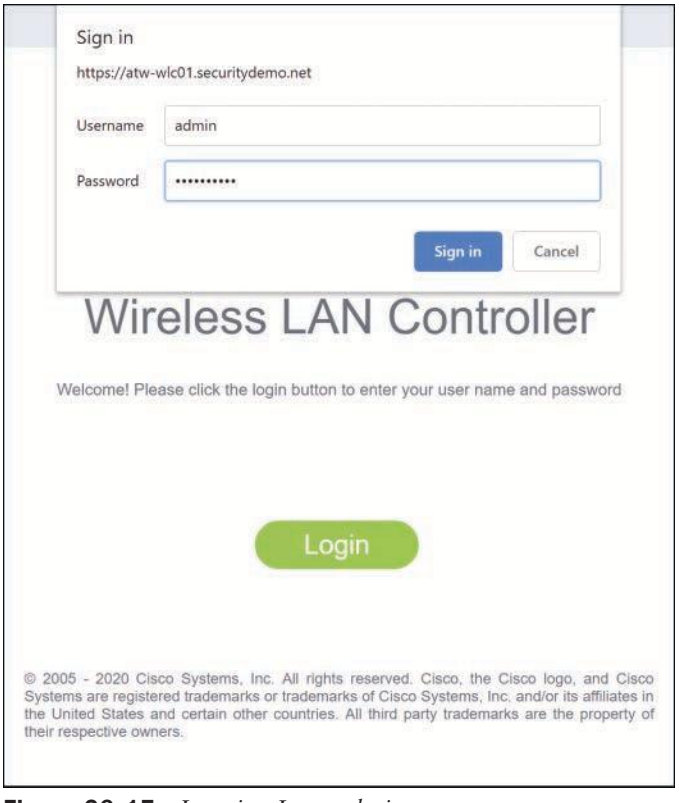


Figure 26-15 *Logging In as admin*

The authentication fails. So, where do you look first to identify why an authentication has failed? That's right: TACACS Live Log.

Step 2. Navigate to **Work Centers > Device Administration > Overview > TACACS Live Log**. Figure 26-16 shows the failed admin login.

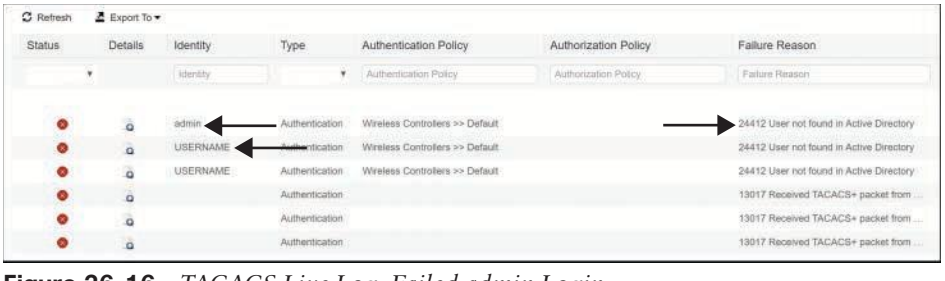


Figure 26-16 *TACACS Live Log: Failed admin Login*

In Figure 26-16, notice the Identity column entries that say USERNAME and admin. The USERNAME entries were recorded before the Disclose Invalid Usernames setting was enabled, and the latest entry, admin, was the attempt after that setting was enabled.

You can click on the details icon to see the authentication details report, which explains the failure reason as "22056 Subject not found in the applicable identity store(s)." This makes sense. ISE is configured to check with the backend identity stores, such as its own internal user database and active directory. The admin account is a local account that exists only on WLC. WLC does not attempt to authenticate local accounts while the TACACS+ server responds. This provides a fallback in case there is a network disruption and ISE is not reachable by WLC.

Step 3. Log in again, but this time as a network administrator, as shown in Figure 26-17. This time the authentication succeeds.

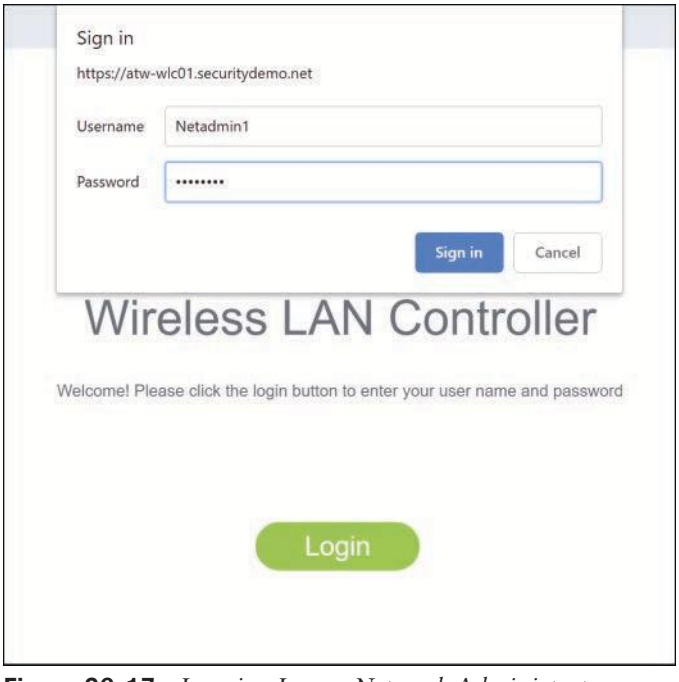


Figure 26-17 *Logging In as a Network Administrator*

Step 4. Examine Live Log, as shown in Figure 26-18. As you can see in the figure, the authentication succeeds, as does the subsequent authorization where the Net-Admin WLC authorization result is applied.

Step 5. To test what happens when you make a configuration change, navigate to a menu that the other roles do not have access to, such as **Controller**.

Step 6. Make a change to that configuration page, such as adding a DNS server, as shown in Figure 26-19.



Figure 26-18 *Successful Login, NetAdmin WLC Result*

Step 7. Click **Apply**.

Step 8. Click **Save Configuration**. As you can see in Figures 26-19 and 26-20, the change is applied successfully, and the configuration is saved successfully.

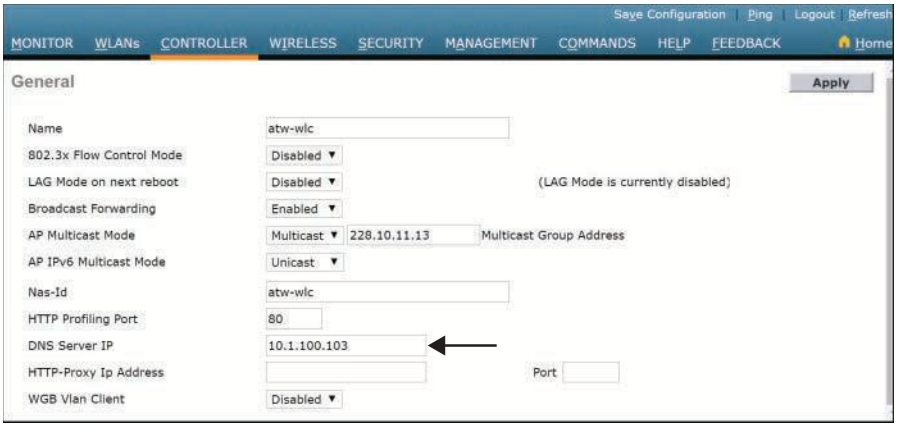


Figure 26-19 *Making a Change to the Controller Page*

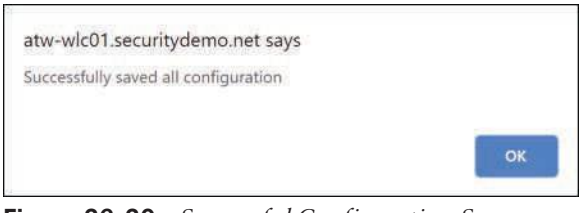


Figure 26-20 *Successful Configuration Save*

Step 9. Use SSH to access WLC.

Step 10. Log in as admin. The login fails.

Step 11. Log in as NetAdmin1. The login succeeds.

Now you can see that the CLI is being controlled by the same RBAC as the GUI. Figure 26-21 shows the login attempts from the CLI.

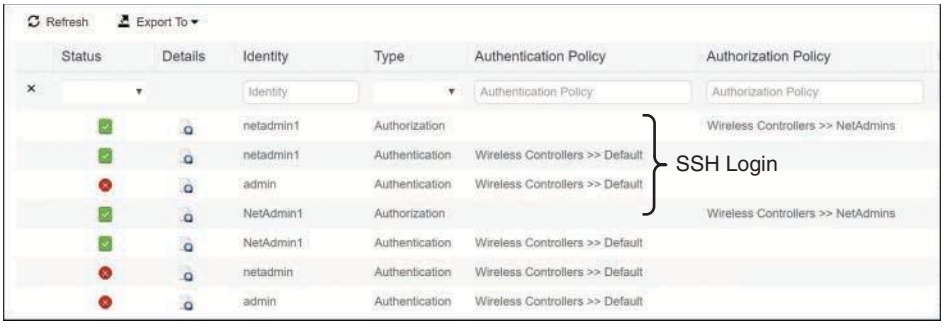


Figure 26-21 *CLI Login Attempts in Live Log*

Step 12. At the WLC CLI, type **debug aaa tacacs enable.**

Step 13. Log out of the GUI.

Step 14. Log in to the GUI as a user from the SecOps group. The login as SecOps1 succeeds. Figure 26-22 shows the success and the assigned SecAdmin WLC profile in Live Log. In addition, the CLI shows a tremendous amount of debug activity happening. One of the messages here read "User has the following mgmtRole 48." This number matches the value shown in ISE's GUI for the SecAdmin WLC TACACS profile.



Figure 26-22 *SecOps1 Login Shown in Live Log*

Step 15. To test what happens when you make a configuration change, navigate to a menu that the SecAdmin role should not have access to, such as **Controller**.

Step 16. Make a change to that configuration page, such as adding a DNS server. The change is accepted.

Step 17. Click **Apply**. As you can see in Figure 26-23, the user fails to apply the change due to insufficient privileges. This provides more proof of how RBAC works with the WLC. The menus are available to everyone, but only authorized changes are applied.



Figure 26-23 *Authorization Failed*

Step 18. Navigate to a menu that the SecAdmin role does have access to, such as **Security**.

Step 19. Make a change to that configuration page, such as adding a new ACL. The change is accepted.

Step 20. Click **Apply**.

Step 21. Click **Save Configuration**. The change is applied successfully, and the configuration is saved successfully.

Step 22. Log out of the GUI.

Step 23. Log in to the GUI as Helpdesk1. The login as Helpdesk1 succeeds. All the UI pages are available, but no changes are applied. The helpdesk users are assigned the MONITOR role, which is basically equivalent to "read everything, write nothing."

Step 24. Log out of the GUI.

Step 25. Log in to the GUI as Employee1. The login as Employee1 succeeds.

As you can see in Figure 26-24, it's a different GUI. The LOBBY role that is assigned to any employee not matching one of the more specific groups in the configuration provides access to a guest management user interface only.



Figure 26-24 *Lobby Ambassador UI*

Exam Preparation Tasks

As mentioned in the section "How to Use This Book" in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 27, "Final Preparation," and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topics icon in the outer margin of the page. Table 26-2 lists these key topics and the page number on which each is found.

Key Topic

Table 26-2 Key Topics

Key Topic Element	Description	Page
List	Six different roles	971
Paragraph	Local login as a backup if the TACACS server is unavailable	980

Q&A

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep software online.

1. What view provides the true TACACS results rather than a GUI representation of them?
2. Which role provides full read/write access to every section of the WLC UI?
3. Which role is used to provide access only to the guest services of the WLC?
4. What setting must be enabled in ISE to see the usernames of failed authentications?
5. Where does an ISE administrator go to begin troubleshooting device administration AAA authentications or authorizations?



CHAPTER 27

Final Preparation

The first 26 chapters of this book cover the technologies, protocols, design concepts, and considerations required you need to understand to pass the Cisco Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) exam.

If you are pursuing the CCNP Security certification, you need to pass the Implementing and Operating Cisco Security Core Technologies (SCORE 350-701) exam, as well as a concentration exam such as SISE 300-715.

Chapters 1 through 26 of this book cover the information that is necessary to pass the exam. However, most people need more preparation than simply reading the first 26 chapters of this book. This chapter, along with the Introduction of the book, suggests hands-on activities and a study plan that will help you complete your preparation for the exam.

Hands-on Activities

As mentioned, you should not expect to pass the SISE 300-715 exam by just reading this book. The SISE 300-715 exam requires hands-on experience with many of the Cisco technologies, tools, and techniques discussed in this book, including Cisco switches, Cisco Threat Response, Cisco ISE, Cisco pxGrid, Cisco AnyConnect Secure Mobility Client, and the different APIs supported by those products. Building your own lab, breaking it, and fixing it is the most effective way to learn the skills necessary to pass the exam.

Suggested Plan for Final Review and Study

This section lists a suggested study plan from the point at which you finish reading this book through Chapter 26 until you take the Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) exam. You can ignore this four-step plan, use it as is, or modify it to better meet your needs:

Step 1. **Review key topics and DIKTA questions:** You can use the table at the end of each chapter that lists the key topics in each chapter or just flip the pages looking for key topics. Also, reviewing the "Do I Know This Already?" (DIKTA) questions at the beginning of each chapter can be helpful for review. Do not simply memorize these questions and answers; instead, use them to determine which topics you need to spend more time on.

Step 2. **Review the exam blueprint:** Cisco maintains a list of testable content known as the Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) exam blueprint. Review it and make sure you are familiar with every item listed. You can download a copy from https://www.cisco.com/c/dam/en_us/training-events/le31/le46/cln/marketing/exam-topics/300-715-SISE.pdf.

Step 3. **Study the "Q&A" sections:** Go through the review questions at the end of each chapter to identify areas in which you need more study.

Step 4. **Use the Pearson Cert Practice Test engine to practice:** The Pearson Test Prep software provides a bank of unique exam-realistic questions available only with this book.

Summary

The tools and suggestions listed in this chapter have been designed with one goal in mind: to help you develop the skills required to pass the Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) exam to achieve the CCNP Security or the Cisco Certified Specialist certifications.

This book has been developed from the beginning both to present you with a collection of facts and to help you learn how to apply those facts. Regardless of your experience level before reading this book, it is our hope that the broad range of preparation tools and the structure of the book will help you pass the exam with ease.



GLOSSARY OF KEY TERMS

A

Active Directory (AD) An external identity store from Microsoft that can be used with ISE authentication and authorization policies.

Adaptive Network Control (ANC) The successor to EPS, which added different custom labels that can be used in authorization policies.

AMP incident An AMP event that has occurred, such as a detected threat or vulnerable application.

Anomalous Behavior Detection An optional feature in ISE for mitigating some of the attempts at MAC address spoofing.

AnyConnect DART (Diagnostics and Reporting Tool) A module of Cisco's AnyConnect client that is used to collect detailed logs related to all aspects of AnyConnect, including the NAM supplicant.

AnyConnect NAM (Network Access Manager) Cisco's enterprise-class supplicant, which is an evolution of a Meeting House acquisition that became the Cisco Security Services Client (CSSC) and is now a module in AnyConnect.

authentication server A device that validates EAP credentials from a supplicant and makes authentication and authorization decisions.

authenticator A device that is communicating with a supplicant using the EAP over LAN protocol. It is the device responsible for sending the EAP authentication request to the authentication server via RADIUS. Network access devices are authenticators.

C

CDP Second Port Disconnect A Cisco-proprietary method by which an IP Phone uses Cisco Discovery Protocol (CDP) to communicate when the endpoint connected behind the phone disconnects from the network. This method works for MAB, dot1x, and WebAuth.

Centralized Web Authentication (CWA) Web Authentication in which pages are hosted in a central location, such as on ISE. With CWA, the authenticator is only aware of a MAB authentication, and the authentication server (ISE) houses the web portal and maintains the synchronization between the user's credentials and the MAB authentication session.

certificate An X.509 certificate, which is an electronic document that has been signed for authenticity.

certificate authentication profile (CAP) A configuration that is used to extract the principle identity from a field in the certificate. When using the EAP-TLS protocol for user authentication with certificates, you need to configure CAP to define what certificate field to use as the username attribute for the ISE authentication process.

certificate authority (CA) The top-level entity in a PKI deployment that is responsible for issuing digital certificates. The CA could be a local enterprise resource in the company's infrastructure or could be a third-party CA provider.

certificate revocation list (CRL) A signed list of revoked certificates' serial numbers. The CA publishes the CRL via a website that can be retrieved by a device or an application that needs to validate whether a digital certificate has been revoked by the certificate authority administrator.

certificate signing request (CSR) A request sent from an applicant to a certificate authority in order to apply for a digital identity certificate.

Change of Authorization (CoA) A message that a RADIUS server (or a policy server) sends when it sends any change of the authorization status of a user or an endpoint. This can be accomplished via a port bounce, a reauthentication, or even the push of an entirely new policy.

Cisco AnyConnect ISE agent A persistent posture agent that can be installed on Windows or macOS clients to monitor and enforce posture compliance. It can be installed with the user interface visible to the user or in stealth mode, where it runs without any client interaction.

command-line interface (CLI) A text-based method used to access ISE or another device. When using a CLI, most interaction with the device occurs using only the keyboard.

command set Part of the configurable authorization results in ISE that allows you to define a specific list of commands that can be executed.

Common Access Card (CAC) An embedded integrated circuit that can be used for two-factor authentication. Also called a smart card. A CAC stores a set of X.509-encrypted digital certificates used for authentication processes. An employee's identification badge might be a CAC.

Common Vulnerabilities and Exposures (CVE) A system to track, monitor, and describe publicly known security vulnerabilities.

Common Vulnerability Scoring System (CVSS) An open standard for assessing the severity of security vulnerabilities. CVSS assigns severity scores of 0 to 10 to vulnerabilities to enable incident responders to prioritize their responses.

compound condition A policy object that contains multiple attributes along with an operator such as AND or OR.

context The end result of all attributes associated to an endpoint or user, including (but not limited to) the user identity, endpoint profile, posture compliance, location, time, and any other attribute related to an authentication request.

context brokering ISE brokering in which pxGrid participants share data exchanged between other members of the security system.

Context-In A feature that enables ISE to receive contextual information through pxGrid to help ISE with its own profiling policies.

Context-Out A feature that enables the sharing of contextual information about an endpoint or user to external pxGrid participants.

cyber threat intelligence (CTI) Information about threats related to the computer and networking world. Basically, CTI is data about existing threats.

D

device administration AAA The form of authentication, authorization, and accounting (AAA) used for controlling login and configuration command and control enforcement for network devices. It is most often used with TACACS+.

DNS ACL An ACL that is based on snooping performed by a wireless AP that updates the controller on which destination IP addresses to permit access to.

dual SSID onboarding BYOD onboarding in which a user joins an open SSID first, and the credentials from the CWA are used to authenticate and authorize the end user for onboarding. After the device is onboarded, the closed SSID is used with the provisioned supplicant.

E

EAP over LAN (EAPoL) IEEE 802.1X, a standard for port-based network access control for local-area and metropolitan-area networks.

EAPoL-Proxy-Logoff A standards-based method an IP Phone uses to send an EAP logoff message on behalf of the 802.1X authenticated endpoint that is connected to the phone. This method does not work for non-802.1X-authenticated endpoints.

Easy Connect An authentication method with which Microsoft Active Directory logins are used to passively map user information onto existing network sessions that were initiated with MAB.

Endpoint Protection Services (EPS) A Cisco tool that provided an API that allowed other applications to initiate a quarantine action against an endpoint based on IP address or MAC address.

environment data Information that a TrustSec device receives from ISE about the environment, including server lists, device SGTs, defined SGTs, and expiration timeout.

exploit A weaponized software package that takes advantage of a vulnerability to accomplish a task that was not meant to be permitted. So, a vulnerability means there is theoretically a way to exploit something (that is, a vulnerability exists), and an exploit means that there is a definite path to doing so.

Extensible Authentication Protocol (EAP) An authentication framework that defines the transport and use of identity credentials.

external identity store An identity store that is configured to use an external database server for authentication and/or authorization policies. Examples include Active Directory, LDAP, RADIUS token servers, RSA SecureID, and certificate authentication profiles.

F–G

FlexAuth (Flexible Authentication) A capability of a Cisco switch interface that allows a network administrator to set an authentication order and priority on a switch port, thereby allowing the port to attempt 802.1X, MAB, and WebAuth, in this order. All these functions are provided while maintaining the same configuration on all access ports, so this is a simpler operational model for customers than traditional 802.1X deployments.

global CoA A form of CoA that is sent automatically when a device transitions from unknown to any known profile.

graphical user interface (GUI) A graphic-based method used to access ISE or another device. Using a GUI usually requires a mouse to point and click on the various configuration or administrative components of the device. The keyboard is used to populate text-based fields within the GUI.

guest A user or device that needs network access who is not an employee or a corporate-owned device.

Guest Flow A special state leveraged in an ISE authorization policy that indicates that a user's credentials were entered through a Centralized Web Authentication portal.

guest locations A mechanism that simplifies time zone selection for a guest account.

I

identity source An individual identity store.

Identity Source Sequence (ISS) A list of multiple identity sources combined in a top-to-bottom sequence list for ISE authentication or authorization processes.

identity store An internal and/or external database that can be used to authenticate a user's or an endpoint's credentials.

indicator of compromise (IoC) An indicator that an artifact or a series of artifacts has been discovered on an endpoint that indicates, with high confidence, that a breach occurred.

integration A process in which ISE shares data outbound, receives information inbound, and brokers data exchange between other members of the ecosystem.

internal endpoint database An ISE database that is used to store information about all the endpoint devices that have connected to the ISE infrastructure. This database stores the endpoints' attributes, including MAC address, IP address, and various other attributes learned using the ISE profiling probes.

internal user database An ISE database that is used for the creation of internal username and password accounts (which are typically referred to as internal user accounts). This database can also be used for ISE authentication and authorization policies. An internal user database may be used, for example, for ISE guest services.

ISE CA (certificate authority) ISE's built-in certificate authority, which may be used to issue certificates to BYOD endpoints.

ISE cube Another name used to refer to an ISE deployment.

ISE profiling An advance subscription license feature of ISE that is used to identify endpoints based on network data obtained from a number of enabled probes. With profiling, you can build an authorization policy that combines a user's identity with the classification result and invoke specific authorization results.

K–L

key pair The combination of a public key and a private key that is used for asymmetric encryption.

Lightweight Directory Access Protocol (LDAP) A TCP/IP client-model protocol that is used to connect and retrieve information from an X.500 directory service database using TCP port 389. A secure version of LDAP is also supported, using Secure Sockets Layer (TLS/SSL) with TCP port 636. LDAP, which is defined in RFC 2251, can be used as an external identity store.

Local Web Authentication (LWA) A form of authentication in which WebAuth pages are hosted on a local device, such as a switch, a wireless controller, or even a firewall. With LWA, the authenticator receives the credentials from the web portal and sends those credentials to the authentication server through a RADIUS Access-Request. The portal may be hosted on the authenticator itself or on an external server.

logical profiles Groupings of individual profiles.

M

MAC Authentication Bypass (MAB) A method by which a NAD sends a RADIUS authentication on behalf of an endpoint, using the endpoint's MAC address as the identity credential.

MACsec IEEE 802.1AE, a standard for authenticating and encrypting packets between two MACsec-capable devices.

mobile device manager A security software system that allows an administrator to configure and secure a mobile device, regardless of where the device is located. Mobile device managers can provide application provisioning, security posturing, and remote wipe capabilities.

Monitoring and Troubleshooting (MnT) node An ISE node that provides monitoring and troubleshooting functions for a Cisco ISE deployment.

Multi-Auth A mode that is an extension to MDA and permits a virtually unlimited number of MAC addresses per interface, with a separate authentication session for each MAC address.

Multi-Hosts A mode that is an extension to MDA and permits a virtually unlimited number of MAC addresses per interface. All endpoints attached to the interface utilize the authorization results of the first MAC address to authenticate successfully. No other hosts are able to use 802.1X authentication.

Multiple-Domain Authentication (MDA) An enhancement to the default mode of 802.1X that allows for two MAC addresses, one in the voice domain and the other in the Data domain.

N

Native-EAP Specific methods of Extensible Authentication Protocol, including EAP-MD5, EAP-MSCHAPv2,and EAP-TLS.

network access AAA The form of authentication, authorization, and accounting (AAA) used for controlling login and access to networks, controlling who and what is allowed to communicate on a network.

network access control (NAC) The industry-standard term for the posture assessment as part of access control.

network access device (NAD) An access-layer device that acts as the 802.1X authenticator and policy enforcement point for a secure access solution. A switch, firewall, or other device that provides access control enforcement.

Network Device Admission Control (NDAC) An option that provides the ability to authenticate network access devices that are joining a TrustSec domain. The purpose is to create a chain of trust within a TrustSec domain and allow only trusted network devices.

network device group (NDG) A hierarchical group that can be used to logically group network devices based on various criteria, such as geographic location, device type, or the relative place in the network (access layer, data center, and so on).

Network Time Protocol (NTP) A time synchronization networking protocol which ensures that network-based computer systems have the correct date and time for the operating system. NTP is critical when it comes to issuing and using digital certificates in a PKI customer deployment. NTP uses User Datagram Protocol (UDP) port 123.

node A physical or virtual ISE appliance that runs one or more personas.

node groups Groupings of PSNs, where the PSNs maintain a heartbeat with each other. If a PSN goes down while a session is being authenticated, one of the other PSNs in the node group sends a CoA to the NAD so the endpoint can restart the session establishment with a new PSN. Local replication of profiling data also occurs within the node group.

non-seed device A device that acts as a supplicant in a TrustSec domain. Such a device does not initially require direct IP connectivity to ISE and enrolls and authenticates to the seed device.

O

one-time password (OTP) An example of a two-factor authentication process in which the user's account password is generated for each individual authentication access process. OTPs come in various forms, including software-based and hardware-based token generators. The account user inputs a secret PIN number, and a unique one-time password is generated for the user's authentication login process.

Online Certificate Status Protocol (OCSP) A protocol that allows a device or an application to send real-time requests, typically using HTTP. The OSCP process is a client request/server response process. OCSP is documented in IETF RFC 6960.

P

personas The different functions in Cisco ISE that are required for proper operation of the platform.

Policy Administration node (PAN) The administrative persona for Cisco ISE. All policies on ISE are configured and then pushed to all other ISE nodes/personas from this node.

policy matrix A common view in ISE that lists the source and destination tags. This is where filtering using SGACLs may be configured and deployed to network devices to ensure a consistent policy across the TrustSec domain.

Policy Services node (PSN) The workhorse node of an ISE deployment that is responsible for actively servicing RADIUS authentication and authorization requests.

port security A feature included in the Catalyst Integrated Security Features (CISF) of Cisco Catalyst switches that limits the number of MAC addresses permitted on a single switch interface. Port security and 802.1X are mutually exclusive, meaning they cannot both be configured on the same switch interface without causing unexpected negative behavior.

portal A website that is set up to allow the user to accomplish tasks.

posture A process by which an application (such as Cisco AnyConnect ISE agent) running on an endpoint or a temporary executable file (such as a Cisco temporal agent) is run on the client and provides critical information about the software or operating system that is actively running on the device.

posture conditions Definitions of the items that a posture module examines on an endpoint. A condition may be looking for the existence of an item (such as a file), the absence of that item, or details about an item (such as the file date and hash). Failing a mandatory condition results in noncompliance of the endpoint. Also known as a check.

posture remediation When an endpoint fails a posture condition, actions taken to correct the failure so the endpoint can be moved from a noncompliant state to a compliant state.

posture requirements Policy elements where conditions are bound to the remediations. The requirement elements go into the posture policy.

principle username X.509 attribute The identity used for the user credential that has been extracted from a field in a certificate.

private key The half of the key pair that should never be shared with any other entity. Data encrypted with the private key can only be decrypted by using the public key.

privilege level The level of access a user has in an IOS shell session.

profiling The process by which an authentication server makes educated observations about what a device is, based on various attributes of the endpoint.

Protected Access Credential (PAC) A unique shared credential between a server (ISE) and a client (network access device). The PAC allows the server and the client to authenticate via a secure tunnel and enables the client to download control data such as the SGACL policy and environment data.

public key The half of the key pair that is shared with other entities for purposes of authentication and/or encryption. Data encrypted with the public key can only be decrypted by using the private key.

Public Key Infrastructure (PKI) A set of hardware, software, and security policies that enables the enrollment, storage, distribution, and management processes for digital certificates. PKI is based on the X.509 standard.

publisher The source of the data that subscribers are subscribed to. Any other pxGrid participant may be a publisher.

pxGrid A scalable pub/sub communication bus that is used for the sharing of large amounts of security data at scale.

R

Rapid Threat Containment A trigger from a pxGrid participant that can give an endpoint a label in ISE and change the level of access the endpoint has, based on a corresponding authorization policy.

registration authority (RA) A second-level entity in a PKI deployment that can process and verify certificate-based requests. The RA helps offload the processing of certificate processing from the CA. This also enables the PKI deployment to be more scalable for large corporate PKI environments.

Remote Authentication Dial-In User Service (RADIUS) An application-layer network protocol used for authentication, authorization, and accounting. RADIUS is a client/server protocol that uses User Datagram Protocol (UDP) port 1645/1646 (IETF Draft) or UDP port 1812/1813 (IETF Revised Standard).

S

SAML assertion A packet of security info that contains the identity and attributes of a user and the user's authorization for a service that is passed from the IdP to the SP.

SAML identity provider (IdP) A policy engine that determines authorization and issues SAML assertions. The IdP is the identity provider against which ISE authenticates.

SAML service provider (SP) An application or service that is being accessed and that requires AuthN/AuthZ before it can be accessed. The ISE web portal is the SAML SP in this case.

security group access control list (SGACL) An access list that allows the control of access and permissions based on the SGTs assigned.

Security Group Exchange Protocol (SXP) A peering protocol that allows network devices to communicate their databases of IP address-to-SGT mappings to one another. This provides support for retagging a frame after it reenters a TrustSec domain.

security group firewall (SGFW) A firewall or router that can filter traffic based on source and destination SGTs.

Security Group Tag (SGT) A 16-bit value that ISE assigns to a user's or an endpoint's session upon login. Sometimes referred to as a Scalable Group Tag.

seed device The authenticator to network devices for the TrustSec domain. This device is configured manually and has connectivity to ISE initially.

Simple Certificate Enrollment Protocol (SCEP) A protocol used to broker the provisioning of a certificate to an endpoint.

simple condition A policy object that is defined with a single attribute.

Single-Host The default mode of an 802.1X-enabled port, which allows only a single MAC address per switch interface.

single SSID onboarding BYOD onboarding in which a username and password are used to authenticate and authorize the end user. After a user is onboarded, the same SSID is used but with a certificate.

sponsor An employee who introduces and supports a guest.

Structured Threat Information Expression (STIX) A language format used to share cyber threat intelligence (CTI) (that is, threat data). STIX is a format, not a transport protocol. It requires a protocol, such as TAXII or pxGrid, to carry it between consumers and producers of the STIX data.

subscriber A pxGrid participant that subscribes to a topic of interest and is notified when that data is retrieved or updated.

supplicant The software on an endpoint that understands how to communicate with EAP over LAN (802.1X).

T

temporal agent A temporary executable file that is run on a client to check compliance status. The Cisco temporal agent is run at the time of connection to the network and is removed after the login session is terminated.

topic A list of information that is available for pxGrid subscribers, which might include session data or a list of vulnerable endpoints.

trusted authority An entity whose signature is considered valid.

Trusted Automated Exchange of Intelligence Information (TAXII) A protocol used to exchange CTI over secure communication (HTTPS). TAXII is designed specifically to carry STIX CTI. TAXII follows a pub/sub model, much like pxGrid.

TrustSec A Cisco solution that simplifies the provisioning and management of secure access to network services and applications. By classifying traffic based on the contextual identity of the endpoint versus its IP address, Cisco TrustSec enables more flexible access controls for

dynamic networking environments and data centers. MACsec is also grouped into TrustSec as an option to encrypt at Layer 2, hop by hop.

TrustSec domain A chain of trust in which the Security Group Tags and security group ACLs are all downloaded from the same trusted source, well-known authentication methods are used, and SGTs are propagated throughout.

tunneled EAP An EAP method that involves building a secure tunnel between the supplicant and the authentication server. Native EAP communication occurs within the secure tunnel.

two-factor authentication A more secure method of providing authentication credentials that requires two separate pieces of information for the user authentication to be successful. Typically this is something you know and something you have to prove your identity. An example of two-factor authentication is using an ATM machine to withdraw money from your checking account and needing your physical bank ATM card plus your secret PIN code to prove your identity.

U–V–W

user identity group A group that has associated internal user accounts that sets the type of internal account being created.

vendor-specific attributes (VSAs) Attributes that developers and various network access device vendors use to extend RADIUS to perform other functions and carry information within RADIUS.

vulnerability A weakness in computer software or hardware. Also known as a security hole. The term *vuln* is a shortened form of *vulnerability*.

Web Authentication (WebAuth) A process in which an end user submits credentials to the authentication server through an interactive web page. WebAuth is used when stronger authentication (such as 802.1X) is unavailable.

This page intentionally left blank



APPENDIX A

Answers to the "Do I Know This Already?" Quizzes and Q&A Sections

Answers to the "Do I Know This Already?" Quizzes

Chapter 1

1. d. **Explanation:** Authentication involves validating someone's identity credentials. Authorization involves determining what is allowed or not allowed, based on those credentials.
2. a and d. **Explanation:** The two forms of authentication, authorization, and accounting that are relevant to the SISE exam are network access and device administration.
3. b. **Explanation:** TACACS+ is best suited for granular command-level control due to its ability to separate authentication and authorization.
4. c. **Explanation:** RADIUS is best suited for network access AAA due to its ability to work with numerous authentication protocols, such CHAP and MS-CHAPv2, and its use in 802.1X authentications and change authorization.
5. c. **Explanation:** Both TACACS+ and RADIUS may be used to provide device administration AAA services, but TACACS+ offers command-level authorization, and RADIUS does not. SAML and OAuth are used for AAA with modern web applications.
6. d. **Explanation:** TACACS+ for device administration was added to ISE beginning with ISE Version 2.0. Prior to that, Cisco Secure Access Control Server (ACS) was the product to use for device administration, although ACS is now end of sale. There is no ICS or TCS product. Cisco Centri was an old Cisco firewall product that is end of life.
7. **d. Explanation:** The majority of the authentication protocols used (for example, EAP, CHAP, MS-CHAPv2, PAP) are Layer 2 protocols, and they are meant to be topology independent. RADIUS and TACACS+ are used to connect an end user to an authentication server, even when they are not on the same LAN segment.
8. a. **Explanation:** TACACS+ clients send only two message types: START and CONTINUE. REPLY is sent from the AAA server to the AAA client.
9. b. **Explanation:** The Service-Type field value tells the RADIUS server what is being performed. For example, the value Call-Check informs the AAA server that the client is performing a MAB request.
10. a. **Explanation:** A RADIUS server may assign an attribute to an authentication session, such was with a VLAN. In this case, the VLAN placeholder is the attribute, and the assigned VLAN number is the value for that placeholder; together the attribute and value form a pair.

Chapter 2

1. b and c. **Explanation:** The two different types of identities used with Cisco ISE are endpoints and usernames. MAC addresses are used to identify the endpoints within a data store. An IP address is an ephemeral address that can be assigned to an endpoint for a time.
2. b and c. **Explanation:** ISE can leverage both internal and external identity sources. Temporal and permanent are not ISE identity sources.
3. a and d. **Explanation:** The two different identity types that exist with ISE are endpoints and users. Each identity store contains either user accounts or a list of endpoints.
4. a and d. **Explanation:** Time and accounting are two different attribute types, but they are not identities. User accounts and machine accounts can both be authenticated through ISE.
5. c. **Explanation:** An identity source sequence (ISS) is a list of identity sources that can be checked for authentication, in order.
6. c. **Explanation:** An identity source sequence (ISS) is a list of identity stores that are processed in top-down order.
7. a, c, and d. **Explanation:** Each of these are valid identity stores, but DIAMETER is not a supported protocol with ISE; it is used almost exclusively in the service provider world.
8. b and d. **Explanation:** MAC Authentication Bypass (MAB) compares the MAC address of an endpoint to the endpoint's internal identity store. Guest accounts are created and stored in an internal user's data store. Neither SSID nor time is an identity store.
9. a and b. **Explanation:** The two internal identity stores for ISE are for users and endpoints.
10. c and d. **Explanation:** Internal identity stores are great for testing, but using them is not realistic for large-scale operations. External identity stores such as Microsoft Active Directory centralize the identities and provide for much greater scalability and ease of management.

Chapter 3

1. b. **Explanation:** EAP communication occurs between the supplicant and the authentication server. The authenticator acts as a middleman and encapsulates the unmodified EAP frames within the RADIUS communication to the authentication server.
2. b. **Explanation:** For purposes of the SISE 300-715 exam, only Cisco AnyConnect NAM 3.1 and newer are capable of running EAP chaining. However, in December 2019, Microsoft released EAP chaining support as part of its initial TEAP offering, in conjunction with Cisco's ISE Version 2.7. However, the SISE 300-715 exam predates this release.
3. c. **Explanation:** The outer identity provides a mechanism to authenticate the identity of the endpoint during the tunnel establishment phase.

4. a. **Explanation:** IEEE 802.1X must use RADIUS or DIAMETER (but DIAMETER is beyond the scope of the SISE 300-715 exam).
5. d. **Explanation:** Supplicants have the option of not authenticating the server certificate. In addition, EAP-FAST offers the ability to use PAC files instead of certificates for tunnel establishment. DNS would not be a requirement for EAP since EAP communication occurs at Layer 2. While the supplicant *could* be using a certificate that was issued by the authentication server, a client certificate is not required and could be issued by any certificate authority.
6. d. **Explanation:** A Protected Access Credential (PAC) is a type of secure cookie that can be used instead of or in addition to a certificate.
7. a. **Explanation:** MS-CHAPv2 cannot be used with ISE when the identity store is LDAP.
8. b. **Explanation:** The tunnel mechanism is unrelated to the ability to do machine authentication. The requirement is simply that EAP-MS-CHAPv2 or EAP-TLS must be the authentication method.
9. c. **Explanation:** The three main components of 802.X are the authentication server, the supplicant, and the authenticator.
10. b and c. **Explanation:** Extensible Authentication Protocol (EAP) is an authentication framework that transports credentials. Traditional, or native, EAP types include EAP-MS-CHAPv2, EAP-TLS, and EAP-GTC. A tunneled EAP type is able to use native EAP types as its inner method. Examples of tunneled EAP types are PEAP, EAP-FAST, EAP-TTLS, and TEAP. There is no such thing as machine EAP, and EAP chains bind multiple EAP types together.

Chapter 4

1. b. **Explanation:** The available option for non-authenticating endpoints is MAC Authentication Bypass (MAB). For headless endpoints, Web Authentication and EasyConnect would not be appropriate because there would be no user interaction or authentication.
2. b. **Explanation:** With non-authenticating endpoints, the authenticator (a switch, for example) may be configured to send the MAC address of the endpoint to the authentication server in a RADIUS Access-Request message. This process is known as MAC Authentication Bypass (MAB).
3. d. **Explanation:** With MAB, it is not recommended to use VLAN assignment, but MAB authorizations do not limit the authorization results.
4. b. **Explanation:** The benefit of using EasyConnect is that no PKI or 802.1X supplicant is required. The Active Directory login is passively mapped to a MAB network session. This provides an ISE administrator with a way to easily roll out basic network access with user authentication.
5. a and c. **Explanation:** While EasyConnect can be auto-configured using a domain administrator account, it does not require one; it also does not require an agent to be installed on the domain controller. The domain account that ISE uses to access Active Directory does require permission to access WMI remotely and access to read the security event log of the domain controller.

6. a. **Explanation:** The four main non-802.1X authentication use cases are WebAuth (CWA and LWA), MAB, EasyConnect, and remote-access VPNs.
7. b. **Explanation:** When gathering additional data from probes, ISE can further identify an endpoint by using profiling.
8. c and e. **Explanation:** Centralized Web Authentication uses a web portal that is hosted on ISE to receive the user's credentials. The authenticator sends a MAB request to ISE, and ISE responds with a RADIUS Access-Accept and a URL redirection, and often a dACL limits the access to the network. After the credentials are received through the web portal, ISE sends a Change of Authorization (CoA) to the authenticator, causing reauthentication. The reauthentication maintains the same session ID, and ISE is able to tie the user's credentials to the MAB request and send the final authorization results for the end user. With EasyConnect, ISE passively maps an Active Directory login to a MAB network session.
9. b. **Explanation:** There are many different "headless" endpoints in an organization, such as IP phones, IP cameras, printers, badge readers, IV pumps, and medical imaging systems. Some of these endpoints do not have supplicants. For those that do, the enablement and configuration of supplicants on the disparate endpoints could be very complicated or operationally difficult for the company. Many such devices do not have a central management platform that is capable of configuring each supplicant across large numbers of devices deployed at scale. Therefore, MAB is chosen to provide network access to such headless devices.
10. a. **Explanation:** In ISE Version 2.1, DCHP and DNS services were introduced for third-party NADs that did not support CoA with URL redirection. Instead, an authentication VLAN would be utilized, and ISE would deliver a DHCP address to the endpoint. Subsequent DNS requests would redirect the endpoint's browser to ISE's centralized portal for authentication.

Chapter 5

1. b. **Explanation:** A RADIUS CoA allows an authentication server to trigger a reauthorization. This provides an opportunity for the server to update a user's level of network access as the server learns additional information about an endpoint, such as endpoint posture information.
2. c. **Explanation:** In a situation where a CoA is warranted, an authentication server can perform a number of actions, including no COA (that is, do nothing), port bounce (that is, whether to shut the relevant access port), or reauthenticate (that is, force the endpoint to reauthenticate in cases where multiple endpoints are present on a single access medium). Supported CoA actions may vary depending on the authentication server.
3. c. **Explanation:** Devices that don't have an 802.1X supplicant available use MAC Authentication Bypass. Without the supplicant, a device does not recognize EAP messages and, therefore, EAP authentication techniques are not available. In the absence of EAP, a device uses its MAC address as its unique identifier to authenticate to the network.
4. d. **Explanation:** The first three octets of a MAC address are the organizationally unique identifier (OUI), which indicates the vendor that manufactured the device. This

information can be useful, at times, in determining the function of the device (for instance, an IP phone or printer).

5. a. **Explanation:** Oftentimes, "dumb" network devices lack 802.1X supplicants. From this list, a printer would be the most common device to lack 802.1X support. Other examples include IP phones, IP cameras, and badge readers.
6. c. **Explanation:** Prior to the introduction of MAB, there was not a mechanism to authenticate a device based strictly on the device's MAC address. The switch port would therefore be configured without port security or any level of end-user or device authentication. This would allow any device—either the intended device or an unintended rogue device that was plugged in to that switchport—to have unfettered access to the network.
7. d. **Explanation:** Through posture checking, an endpoint can be checked for file conditions (existence, date, and version), registry conditions (whether a registry entry is present), and service condition (whether a service is running). Posture checking also can confirm the presence, absence, and status of antivirus and antispyware programs running on the endpoint.
8. d. **Explanation:** When using posture assessment as a condition for authorization policy, the value of the PostureStatus condition can be Compliant, NonCompliant, or Unknown. Different levels of network access or remediation can be authorized based on the status of this variable.
9. b. **Explanation:** To remediate a noncompliant endpoint, a redirect ACL must be defined on the switch, and the redirect destination must be set to the remediation portal.
10. d. **Explanation:** A mobile device manager is a third-party appliance or service that provides advanced posture assessment for mobile endpoints. A mobile device manager can determine the type of mobile device, the level of the operating system on the endpoint, the presence or absence of a PIN lock, and whether encryption is being used, and it can provide remote security services such as device lock and secure wipe. Depending on the mobile device management (MDM) vendor chosen, additional services may be available.

Chapter 6

1. c. **Explanation:** Cisco Identity Services Engine (ISE) is a network security and policy platform. Using Cisco ISE, a network administrator can maintain and serve security policy to all network devices from a central location.
2. a, d, e, and g. **Explanation:** Cisco ISE has four personas: Administration, Monitoring, Policy Services, and pxGrid.
3. a. **Explanation:** Cisco ISE's Administration persona is the instance of Cisco ISE where policy configuration actually happens. This persona distributes the policy to all other nodes.
4. d. **Explanation:** The Cisco ISE Monitoring persona provides a platform for logging and reporting data from the Cisco ISE deployment. As a user or device authenticates and authorizes to the network, the ability to monitor and log those AAA events is the responsibility of the Monitoring persona.

5. c. **Explanation:** The Cisco ISE Policy Services persona provides policy decision making. As a user or an endpoint attempts to authenticate to the network, the PSN is responsible for making the AAA decisions based on the policy as downloaded from the Cisco ISE Policy Administration node (PAN).
6. a. **Explanation:** The NAD sends a RADIUS access-request to ISE policy services nodes and the RADIUS response will contain attribute/value pairs (AV pairs) that contain the security policies such as downloadable ACLs, VLAN assignments, and security group tags (SGTs) to perform the enforcement.
7. b and c. **Explanation:** If you choose to deploy ISE as a virtual appliance, it is paramount that you allocate the appropriate virtual resources to best emulate the equivalent SNS-36xx physical appliance. Also, you should reserve 100% of these resources to ensure that other virtualized network functions do not starve ISE of the resources.
8. d. **Explanation:** In a single-node Cisco ISE deployment, all ISE personas reside on a single appliance. In this type of deployment, there are no options for redundancy. For instance, if the Policy Services persona fails, or if the physical appliance fails, RADIUS authentications and authorizations fail until the issue can be resolved.
9. f. **Explanation:** In a hybrid Cisco ISE deployment, the Administration and Monitoring personas are combined on two of the appliances; each of them acts as primary on one appliance and secondary on the other appliance. On the remaining two appliances, only the Policy Services persona is configured.
10. f. **Explanation:** In a fully distributed ISE deployment, the ISE Administration and Monitoring personas each reside on a separate appliance (or a separate pair of appliances if redundancy is required). The Administration and Monitoring appliances will each be an SNS-3695 appliance (or equivalent virtual appliance). With these Administration and Monitoring functions distributed, up to 50 PSNs can be deployed, with a maximum of 2,000,000 endpoints.

Chapter 7

1. d. **Explanation:** The best way to ensure a secure connection is by encrypting the communications between ISE and the device that is being used for the administrative portal. If HTTP were used, any device in the network flow between the administrative device and ISE could eavesdrop or play "man-in-the middle" on the communications, either compromising the administrative credentials or surreptitiously injecting a different security policy. To prevent this from happening, ISE leverages HTTPS, which encrypts all traffic between the administrative device and ISE, to ensure that the traffic sent from the administrative device arrives securely, without compromise. SSH and SCP are not protocols that are typically used for GUI-based portals.
2. b. **Explanation:** Before the initial secure connection with ISE can occur, ISE generates a self-signed certificate. This certificate is not signed by a trusted certificate authority (CA) (either a local CA or a third-party public CA), so you might see a security warning in the web browser that is being used for administrative access. If you are confident that a man-in-the-middle or other nefarious device is not presenting this certificate, you can permanently accept the certificate in the web browser to suppress these security warnings in the future. Ideally, it is best to install a certificate from a trusted CA (a CA that already exists in the browser store, either a local CA or

a third-party public CA) in ISE. Doing so also prevents these security warnings from popping up in the future.

3. a. **Explanation:** The Operations tab of Cisco ISE allows an administrator to monitor, report, and troubleshoot active authentication and authorization sessions.
4. c. **Explanation:** The Policy tab of the Cisco ISE GUI allows an administrator to configure authentication, authorization, profiling, posture, client provisioning, and security group access.
5. b and d. **Explanation:** The Administration component of Cisco ISE can be used to configure all setup-type functions of ISE. These functions are often set up one time and rarely modified thereafter. Certificates and network devices are two items that are configured under the Administration component and are rarely modified after their initial configuration.
6. b, c, and e. **Explanation:** When adding a new NAD to Cisco ISE, you must provide a device name and a device IP address. If you intend to use Cisco ISE as a RADIUS server for authentication and authorization (which is the usual purpose of Cisco ISE in a network deployment), you also need to add a shared secret key for RADIUS. The RADIUS server IP address is configured on the NAD pointing to Cisco ISE. Mobile device managers and SGA AAA servers are unrelated to the network device configuration.
7. c and e. **Explanation:** When an endpoint attempts to access the network, it automatically sends a number of different packets onto the network; this is normal communication for a networked device. The information contained within these packets can often be leveraged by ISE to determine the type of device that is sending the information. The MAC address of the endpoint, learned either via EAP or via MAC Authentication Bypass (MAB) on the NAD, is forwarded to ISE via RADIUS. The endpoint's DHCP requests to get an IP address can also be sent to ISE, allowing ISE to extract key identifying information from this DHCP process. Finally, HTTP(S) communications between the endpoint and ISE portals can be used to further identify the type of device that is accessing the network. Using RADIUS, DHCP, and HTTP (and other protocols), ISE can make a pretty good determination as to the type of device that is accessing the network. ISE currently does not support the use of SSH or FTP for profiling an endpoint.
8. a. **Explanation:** During the client provisioning process, the necessary credentials and configurations are deployed to the endpoint, allowing the endpoint to automatically join the network in the future with little or no interaction from the user.
9. c. **Explanation:** pxGrid is an API that integrates ISE with third-party ISE ecosystem partners.
10. c. **Explanation:** Quarantine in ANC means nothing without a corresponding ISE security policy. That security policy can provide access that is as restrictive or liberal as an ISE administrator wants it to be. It can grant full access to the network or completely remove access for the quarantined endpoint.

Chapter 8

1. c. **Explanation:** The permissions needed to join ISE to AD are Search Active Directory (to see if the ISE machine account already exists) and Add Workstation to Domain (if it does not already exist). An optional permission is Set Attributes on the New Machine Account (OS type and version).

2. b. **Explanation:** The **show application status ise** command list all the ISE processes and the status of each one.
3. c. **Explanation:** In both HTTPS and TLS connections, certificates are used to authenticate the server to the client and act as the basis for the encrypted transport between the client and the server.
4. b and c. **Explanation:** Currently, ISE is supported only on virtual machines and physical appliances.
5. b. **Explanation:** The **show application status ise** command lists all the ISE processes and the status of each of them.
6. a. **Explanation:** Use an NDG as the condition by which to build different policy sets for the staged deployment of ISE.
7. b. **Explanation:** The Admin certificate usage is only for the ISE Admin portal that the administrator logs in to.
8. a, c, and d. **Explanation:** There are three default top-level network device groups, but you can create additional customized NDGs.
9. c. **Explanation:** A certificate authentication profile serves as the identity source for certificate authentications and defines the field of a certificate whose data will be extracted and used as the principal identity for the authorization process.
10. a, d, and e. **Explanation:** The AD account to join ISE to the domain does not require Domain Admin permissions. It simply needs to be able to search Active Directory to see if the ISE machine account already exists, add a workstation to the domain if it doesn't already exist, and set OS and version attributes on the new machine account it creates.

Chapter 9

1. b. **Explanation:** The RADIUS packet must have Service-Type, which dictates the method of authentication, set to Call-Check. The Calling-Station-Id field must be populated with the MAC address of the endpoint.
2. b. **Explanation:** As this book went to press, only EAP-FAST and TEAP (RFC 7170) had EAP chaining capabilities.
3. c. **Explanation:** An authentication policy is meant to drop traffic that isn't allowed. If traffic is using an authentication protocol that is configured, the policy routes authentication requests to the correct identity store to validate the identity and passes successful authentications over to the authorization policy.
4. b. **Explanation:** Only the Process Host Lookup checkbox must be selected in the allowed protocols list for Cisco MAB to work.
5. a, b, and d. **Explanation:** Reusable conditions are saved in the library for reuse.
6. d. **Explanation:** Create one rule for each EAP-Type under the policy set's authentication policy that points to the appropriate identity store for each rule.
7. d. **Explanation:** The Called-Station-ID attribute is used to match the source SSID.
8. d. **Explanation:** The Calling-Station-ID attribute contains the MAC address of the endpoint.

9. c. **Explanation:** The continue option is used to send an authentication to the authorization policy, even if the authentication was not successful.
10. a. **Explanation:** The Conditions Studio allows you to create and edit conditions. It is not where the result of an authentication rule is set or modified.

Chapter 10

1. d. **Explanation:** An authorization profile is a required authorization result that is made up of multiple RADIUS attributes. These RADIUS results affect the ultimate security policy deployed to the NAD on behalf of the endpoint.
2. b. **Explanation:** It contains the RADIUS response Access-Accept or Access-Reject along with the additional authorization attributes to be sent to the network device for enforcement.
3. d. **Explanation:** dACL name, web redirection, Local WebAuth, and auto smart port are common tasks that are sent to a NAD for secure policy enforcement of the endpoint.
4. a. **Explanation:** An authorization policy contains authorization rules, and each rule has at least one authorization profile.
5. a. **Explanation:** Condition attributes can be saved into a library for future use and improved readability.
6. b. **Explanation:** It contains the voice domain permission (cisco-av-pair = device-traffic-class = voice), which authorizes the endpoint to access the voice VLAN assigned to the interface.
7. c. **Explanation:** A simple condition contains only one attribute. A compound condition contains multiple attributes along with an operator such as AND or OR.
8. a. **Explanation:** A hierarchical compound condition can contain a mixture of simple and complex conditions to meet the demands of a security requirement.
9. d. **Explanation:** The end goal of a Cisco Secure Access deployment is to provide very specific permissions to any authorization to provide defense-in-depth while meeting the goals of the organization's security policy. A printer, for example, should not have unfettered access to the network but should have only what is needed (such as reaching the print servers).
10. c. **Explanation:** Cisco dACLs are created entirely on the RADIUS server, and a full ACL is sent down to the network device within RADIUS AV pairs. Non-Cisco network devices must create ACLs on individual local network devices. This allows a Cisco administrator to create and maintain the access lists in a central place and have any changes applied nearly instantly.

Chapter 11

1. c. **Explanation:** 802.1X requires global-level configuration for AAA servers, such as enabling 802.1X on the system, configuring Change of Authorization, and enabling VSAs. In addition, each interface that will be performing authentication requires interface-level commands.
2. b and d. **Explanation:** When interacting with an advanced RADIUS server, such as Cisco ISE, Cisco WLCs require that the same ISE PSN be configured as the authentication and accounting server for the WLAN. In addition, RADIUS NAC must be enabled on the Advanced tab for the WLAN configuration.

3. b. **Explanation:** Cisco switches can be configured to send syslog messages to the MnT node, where the data is correlated as part of the authentication reports.
4. a. **Explanation:** The switch sends periodic test authentication messages to the RADIUS server (Cisco ISE), looking for a RADIUS response from the server (either an Access-Accept or Access-Reject). The username and password in the automated test must exist in the configuration.
5. b. **Explanation:** Device tracking must be enabled in order to capture the IP address of an endpoint.
6. c. **Explanation:** FlexAuth allows a network administrator to set an authentication order and priority on a switch port, thereby allowing the port to attempt 802.1X, MAC Authentication Bypass, and WebAuth, in this order. All these functions are provided while maintaining the same configuration on all access ports, thereby providing a much simpler operational model for customers than traditional 802.1X deployments.
7. d. **Explanation:** Multi-Hosts mode is not commonly used, but it is a valid option. Much like Multi-Auth mode, Multi-Hosts mode is an extension to MDA. There is one authentication on the voice domain, and there is one authentication on the data domain. All other hosts on the data domain are allowed onto the network using the first successful authentication. It's an "authenticate one, allow the rest" model.
8. a. **Explanation:** The **authentication port-control auto** command enables authentication on a port and allows the authorization result to be sent from the RADIUS server.
9. c. **Explanation:** The **show aaa servers** command gives you a quick and simple way to see the current status of the ISE server from the switch's perspective.
10. d. **Explanation:** The **show authentication session interface** command shows that authentications are being attempted, shows which ones are successful, shows what authorization results have been assigned, and much more. Some of the information that is quickly provided by this command's output include the endpoint's MAC address, the authentication method used, any assigned redirect URL, access control lists, and other RADIUS AV pairs provided via the authentication and authorization process.

Chapter 12

1. d. **Explanation:** The Cisco switch needs the HTTPS server to be enabled in order to redirect HTTPS traffic. Before that service can be enabled, the switch needs a certificate. A hostname and domain name must be configured to provide the switch a fully qualified domain name.
2. b. **Explanation:** A traffic-filtering ACL may be downloaded from ISE to a switch as a dACL, but the redirection ACL must preexist on the switch and is called using a RADIUS AV pair. The AirespaceOS-based Cisco WLCs support only locally configured ACLs, which are also applied through the use of a RADIUS AV pair.
3. c. **Explanation:** ISE NAC is a critical setting for a WLAN that enables URL redirection and the pre-RUN states. Without this setting, CWA is not possible. On some WLC versions, the ISE NAC may still be displayed as a RADIUS NAC in the UI.
4. b. **Explanation:** CWA is controlled by the authorization policy. Even an unknown MAC address needs to "continue" out of the authentication policy, so the appropriate response may be sent to the NAD, including the URL redirection to the portal.

5. b and d. **Explanation:** The first rule should match if no more specific authorization rule is used, and it should redirect the user to the CWA portal. The second rule type should exist above the redirection rule, and it should allow access to the user after the user successfully authenticates to the CWA portal.
6. d. **Explanation:** Web Authentication is used mainly when a device is not capable of performing a stronger user authentication, such as when a device does not have an 802.1X supplicant. If a device does not have a supplicant, you *should not* change VLANs on it. Without a supplicant, the device does not normally have the ability to recognize the change and request a new IP address. However, you *can* still use VLAN changes. There are tricks of the trade, such as implementing very short DHCP lease times to help enable that VLAN change on devices that do not have supplicants. TrustSec tags (SGTs) are always a great idea for microsegmentation, and they are not mutually exclusive with any segmentation techniques.
7. c. **Explanation:** With CWA, the NAD sends a MAB to ISE, and ISE responds with an Access-Accept that redirects traffic to a web portal hosted on the active ISE node. The user enters credentials into that web portal, and ISE sends a CoA-ReAuth down to the NAD. At that point, a new MAB authentication with the same SessionID is sent to ISE, and it binds the Web Authentication and the MAB together and then sends the final authorization result back to the NAD in response to the MAB authentication. The configuration of the NAD is different for wired and wireless, although the ISE configuration can be the same, and the overall process is also the same. Different policy sets for wired and wireless could be leveraged, but this is not a requirement.
8. d. **Explanation:** The **show authentication session interface** [*interface-name*] command is like the Swiss Army knife of **show** commands for authentications. Its output shows the MAC address, IP address, the dACL (listed as an ACS ACL), the URL-Redirect ACL, and the URL the end user is being redirected to.
9. b. **Explanation:** Cisco ISE has a phenomenally useful built-in tool that is commonly called Live Log. Live Log provides a near-real-time view of every incoming authentication, Change of Authorization (CoA), and more. You can also leverage Operations > RADIUS > Live Sessions or reach Live Log through Work Centers > Network Access > Overview > RADIUS Livelog.
10. a. **Explanation:** The CoA is a key function. Specifically, it is a CoA-Reauth and causes a switch to reauthenticate the endpoint without starting a new session. The switch sends another MAB request to ISE, which is able to tie the guest authentication from the centralized portal to the MAB request from the switch and assign the appropriate permission.

Chapter 13

1. d. **Explanation:** When a guest logs in, endpoints are added to groups to make it easy to assign guest access to an endpoint when it joins a network and not negatively impact the guest's user experience.
2. b. **Explanation:** A guest is a user or device that is not a corporate employee or corporate-owned endpoint. This user type could be a visitor, contractor, consultant, or customer. Guest accounts are created to allow these users to authenticate and gain Internet access.

3. a, c, and d. **Explanation:** There are three default guest portals with ISE: hotspot, self-registered and sponsored guest portals. SAML and social login are options for portal authentication.
4. a, c, e, and f. **Explanation:** While ISE allows you to create your own guest types, four default guest types are preconfigured: daily, weekly, contractor, and social. Each guest type has unique attributes, such as the lifetime for which the guest account may exist, hours the guest may be active on the network, and the number of devices the guest may use.
5. b. **Explanation:** A guest who connects to the network with an unknown device is redirected to the guest portal. Which portal a guest is redirected to is based on the configuration of ISE. Often, a special wireless SSID is used as the source condition. After the guest is registered, the device is added to the endpoint identity group so the guest is not redirected in subsequent connections.
6. d. **Explanation:** Members of GROUP_ACCOUNTS can edit, delete, reinstate, and approve all guests created by sponsors who are members of the same group. They cannot, however, work with guest accounts created by other groups.
7. b. **Explanation:** Sponsor groups exist to assign common permissions to multiple sponsors. When a sponsor is a member of more than one group, the permissions are additive, and the sponsor receives a combination of the least-restrictive permission from all the groups.
8. c. **Explanation:** While a sponsor could connect directly to the PSN with the correct port for the sponsor portal listener, this is not a very user-friendly process for the end users who are sponsoring guests. The sponsor portal has a field to configure an FQDN that acts as a "friendly name" to make it easy for the end user, and the PSN uses that friendly name to identify that the traffic is destined to the sponsor portal. In addition, when a guest registers, an approval notification can be sent to the sponsor, and the sponsor can respond to approve the request.
9. a. **Explanation:** The guest portal and the sponsor portal are fully customizable, and there are built-in templates to make customization easy. Advanced customizations can be performed offline and imported.
10. c. **Explanation:** SAML authentication replaces the identity source sequence (ISS) for a guest or sponsor portal, as ISS's and SAML authentications are mutually exclusive. SAML is available for guest and sponsor portals to leverage the organization's single sign-on solution for a better end-user experience.

Chapter 14

1. a. **Explanation:** Profiling probes are enabled under the PSNs in a deployment.
2. a. b, and d. **Explanation:** There is no such thing as an EndPointProfile attribute. While an OS scan is used as a condition to determine an endpoint's profile, it cannot be used directly in an authorization policy. An authorization policy may use identity groups (which contain lists of MAC addresses), the EndPoint Policy attribute (which is the actual endpoint profile), and logical profiles (groups of profiles).
3. a and d. **Explanation:** The SNMPQUERY probe periodically queries all the NADs configured with SNMP strings, but it is also a reactive probe. The SNMPQUERY

probe reactively queries a NAD when the RADIUS probe receives an accounting START message or when an SNMP trap is received.

4. b and c. **Explanation:** A WLC can send DHCP information to ISE either using Device Sensor, which sends the information via the RADIUS accounting packets, or by bridging DHCP to the LAN, which then has an IP helper address send that information to ISE. A WLC cannot send DHCP information to ISE by using the DHCP proxy.
5. a. **Explanation:** Cisco no longer includes profile updates within ISE version updates or patches. All new profiles are included and downloaded as part of Cisco Profiler Feed Service.
6. c. **Explanation:** Profiling relies on the certainty value. Each profile has a minimum certainty value, and matching the conditions increases the certainty value. The higher the certainty value of any profile, the more likely it is to be assigned.
7. a. **Explanation:** There are several places in the ISE GUI where this can be found but probably the easiest place is under Context Visibility > Endpoints > Endpoint Classification.
8. b and d. **Explanation:** HTTP User-Agent strings can be gleaned through SPAN monitoring, VACLS, and directly from the ISE web portals. Wired switches do not currently have an HTTP Device Sensor probe, but wireless controllers do.
9. b and c. **Explanation:** With IOS Device Sensor, ISE gets the same information it would get from DHCP without having to place an IP helper address on the Layer 3 interface and the SNMPQUERY probe.
10. c. **Explanation:** Profiles are classified as Cisco Provided, Administratively Modified, or Administrator Created. Only Cisco Provided profiles are overwritten.

Chapter 15

1. b. **Explanation:** The signing CA's public key must be imported into ISE's certificate store, and it must be stored under Administration > System > Certificates > Trusted Certificates, for the Trust for Client Authentication use case.
2. d. **Explanation:** It's very important to understand that the Valid-From field is just as important as the Valid-To field. A certificate is rejected if it is issued for a date and time after the current date and time. NTP is therefore critical for PKI.
3. b and d. **Explanation:** ISE supports both CRL and Online Certificate Status Protocol (OCSP). OCSP is the preferred method for scalability and security reasons. There is no such thing as Secure Authentication Mechanism Language; SAML stands for Security Assertion Markup Language. OAuth is not used for certificate status verification.
4. c. **Explanation:** ISE only leverages the CRL distribution point configured within the trusted certificate store for that signing CA and ignores the field that is in the client's certificate.
5. a. **Explanation:** ISE sends some "throw-away data" to the client that is encrypted with the combination of ISE's private key and the client's public key (the certificate sent for authentication). Then the endpoint decrypts the data with the combination of its private key and the server's public key, proving the client has the full key pair and not just a copy of a public key.
6. c. **Explanation:** A certificate issued by Active Directory Certificate Services is still just an X.509 certificate. It goes through all the same authentication validation as any

other certificate, regardless of the fact that the CA was integrated into AD. The CAP extracts the user's identity from the fields in the certificate for the authorization with AD.

7. b. **Explanation:** While both EAP-TLS and EAP-GTC are native EAP types capable of performing certificate-based authentication, EAP-TLS is more common. EAP-TTLS and EAP-FAST are tunneled EAP types, both of which are capable of having EAP-TLS as an inner method.
8. d. **Explanation:** Allowed protocols, CAP for an identity store, and trusting the signing CA for client authentication are all that is required. Certificate revocation checking and the authorization rule are both optional.
9. c. **Explanation:** Many certificate authorities have websites that permit administrators to download their public certificates and even the full certificate chains. This chapter provides an example of downloading the key from a Microsoft CA. Navigating to this web page and downloading the certificate is how an ISE administrator might obtain the public certificate of the signing CA to trust for client authentications. However, it is not recommended to use PKCS chain files unless there is no other option. Always use Base64-encoded files instead of DER-encoded files.
10. b. **Explanation:** While I'm flattered that you might want to call me to fix your problems, C is definitely not the correct answer. The first question that you would be asked is "What does it say for Failure Reason in the Authentication Details Report?" and to do this, you click the details icon. There is no report named Failed Authentications, and even if there were, it would not exist under Report. Please do not email Aaron directly to ask him why things are not working.

Chapter 16

1. b. **Explanation:** One of the business issues with a bring your own device (BYOD) model is walking end users through the process of configuring their network supplicants to meet corporate policies. The onboarding process helps end users perform those actions themselves, without requiring interaction from the IT department.
2. c. **Explanation:** In order to maintain a seamless experience for the end user, a CoA-Reauth message is used to keep the endpoint connected to the network and simply cause the supplicant to send credentials again. At this point, the supplicant uses the new certificate-based credentials to authenticate. The end user is completely unaware of the actions. A CoA-DM (disconnect message) would drop the endpoint from the network and would provide a poor user experience. Waiting for a reauth interval or a disconnect/reconnect to the network would not be an optimal user experience either.
3. a. **Explanation:** The software is hard-coded to prevent guest users from entering the flows. There is no configuration possible to allow guest users to enter the provisioning process through the dual SSID onboarding flows.
4. d. **Explanation:** In order to onboard, Android needs to download the Cisco Network Setup Assistant. The reason for this is due to Android not allowing unknown apps to be downloaded and executed by default. Instead, the provisioning app must be downloaded through the app store.
5. b. **Explanation:** ISE authenticates any endpoint that has been configured to authenticate to the network, regardless of the onboarding status. The policy can be configured

to send an Access-Reject message or to leave the user in the redirected state to receive a message explaining the need to configure the device independently or call their IT department for assistance.

6. b. **Explanation:** Apple iOS does not use an app to perform the provisioning. Instead, it leverages the native over-the-air (OTA) provisioning that is built in to the OS to handle the certificate signing requests and downloading of a network profile.
7. b. **Explanation:** While an ISE administrator might be able to remove devices from ISE in the Endpoints dashboard, end users log in to their My Devices portal, where they can add or remove devices at leisure.
8. a. **Explanation:** The live authentication log does not show any information about the registration or the NSA app. It does show all the authentications and the Changes of Authorization.
9. d. **Explanation:** ISE enables the use of a native .exe for Windows and a .dmg for macOS.
10. b. **Explanation:** The client provisioning policy determines which NAC agent and NSA wizard to use and which NSP to send to an endpoint. The policy is capable of using the operating system as one of many conditions to determine which result to provide for an endpoint.

Chapter 17

1. c. **Explanation**: A Security Group Tag (SGT) is a 16-bit value that ISE assigns to a user or to an endpoint's session upon login. The SGT can represent the context of the user and device and may be carried in the Layer 2 frame or communicated through SXP. The tag is assigned at ingress and enforced upon egress.
2. b. **Explanation:** Security Group Tags (SGTs) are considered an authorization result in the ISE administrative GUI. They are defined within the TrustSec Work Centers section of the GUI, under Components.
3. a, b, and d. **Explanation:** In order to use a Security Group Tag, the tag needs to be assigned in a process known as classification. Tags can be assigned dynamically and downloaded as a result of an ISE authorization, be assigned manually at the port level, and mapped to IP addresses and downloaded to SGT-capable devices.
4. a. **Explanation:** SXP Version 2 introduced the IPv6 binding and the ability for SXP to automatically negotiate the SXP version with its peer.
5. d. **Explanation**: Cisco has developed a peering protocol (similar to BGP) to allow devices to communicate their databases of mappings of IP addresses to SGTs to one another. This peering protocol is called SGT Exchange Protocol (SXP).
6. a and c. **Explanation:** Every SXP peer session has a speaker and a listener. A speaker sends the mappings of IP addresses to SGTs. The listener receives those updates and records them. A peer can be configured to be both a speaker and a listener for the same peer if both support it. It may have numerous peers as well.
7. a. **Explanation:** Native tagging of SGTs includes the 16-bit tag as a portion of the Cisco MetaData (CMD) field of the Layer 2 Ethernet frame. It may also be included as part of an IPsec link.
8. b. **Explanation:** An SGT can be encrypted within a MACsec-encrypted link between network infrastructure devices or even an IPsec connection. The endpoint is never

aware of the tag that has been assigned, so enabling downlink MACsec between the endpoint and the switch does not help.

9. a and c. **Explanation:** SGTs can be enforced with security group ACLs, which are egress ACLs that use Source and Destination tags as the conditions on which to invoke the egress ACL. In addition, the ASA, ASR, Firepower Threat Defense, and ISR can act as security group firewalls using the Source and/or Destination tags as ACL conditions.
10. d. **Explanation:** Uplink MACsec defines the encrypted connection between network infrastructure components, and downlink MACsec defines the encrypted connection between the access-layer device and the endpoint. While uplink MACsec and downlink MACsec use different keying mechanisms today, both still use the same encryption algorithm, AES-GCM-128.

Chapter 18

1. d. **Explanation:** The three possible posture outcomes following the initial connection with the network are compliant, noncompliant, and unknown. Compliant implies that the endpoint fully adheres to the company's security policy, as configured in ISE. Noncompliant implies that there is at least one deviation from the company security policy. Unknown implies that there is not an agent present on the device and, therefore, the endpoint is unable to report its posture to ISE.
2. d. **Explanation:** The file condition for posture can check for the existence and check the date and version of a file on the client. This can be very useful in determining if a particular endpoint is vulnerable to a new virus or if a specific software package is present on the endpoint.
3. c. **Explanation:** The Stealth Mode Agent is a special configuration of AnyConnect that hides the AnyConnect user interface and runs AnyConnect as a service for posture assessment. If system scan posture is the only AnyConnect service that is running, then the AnyConnect icon will not even be displayed in the system tray, and allows posture to appear to be operating in an "agentless" fashion.
4. a. **Explanation:** Remediation is the process by which an endpoint that is not compliant with security policy becomes compliant. This may include downloading the latest virus definitions, installing a service pack, or launching an application.
5. d. **Explanation:** A posture lease allows an administrator to configure ISE to perform posture assessment in specified intervals for endpoints using the AnyConnect agent for posture assessment instead of performing an assessment at every login. When posture is compliant, a token is issued to the posture module, and the lease begins. When the posture lease is active, ISE uses the last known posture state and does not reach out to the endpoint to check for compliance.
6. b. **Explanation:** The ISE admin must configure ISE to download posture updates which contain the latest definition files for what endpoint protection software suites exist, what their latest versions are, the latest data on hotfixes, and a tremendous number of preconfigured objects from Cisco. Many TAC cases are opened due to endpoints failing posture checks, and the remedy of those cases is simply to update the ISE and AnyConnect compliance modules.
7. c. **Explanation:** Mobile platforms are inherently more locked down than desktop operating systems and require a device manager to gain the level of visibility into the endpoint that AnyConnect is able to achieve with desktop operating systems. ISE has links to one or more mobile device managers for BYOD onboarding and also to provide the

posture assessment of the endpoint. AnyConnect is not used for posture assessment on mobile platforms.

8. a. **Explanation:** This condition type is for bulk attribute collection. This check is preconfigured by Cisco to collect details about the endpoint's hardware to aid in the endpoint context visibility aspects of ISE, which can be viewed in the Context Visibility > Endpoints > Hardware dashboard. This condition cannot be used for access control decisions or posture compliance.
9. c. **Explanation:** The posture elements contain conditions, remediations, and requirements. Conditions and remediations are linked together into requirements. Requirements are used in the posture policy.
10. a. **Explanation:** AnyConnect is a lightweight client and requires a headend to provide the configurations and modules to the endpoint. ISE is one of the possible headends, and Cisco ASA and Umbrella are two additional headend possibilities. Any of these headends can update the AnyConnect client and install modules and provide the configuration. Only ISE and the ASA are capable of installing AnyConnect on the endpoint.

Chapter 19

1. c. **Explanation:** Monitor mode is a process, not just a command on a switch. The process involves enabling authentication (with **authentication open**), seeing exactly which devices fail and which ones succeed, and correcting the failed authentications before they cause any problems.
2. a. **Explanation:** Low-impact mode uses **authentication open** but adds security on top of the framework was built in monitor mode. It uses a PACL on the switch port to permit critical traffic of certain endpoints, such as thin clients, to function prior to an attempted authentication. After the authentication, the authorization should provide specific access, unlike monitor mode, which is the same before and after authentication.
3. d. **Explanation:** By using a phased deployment approach, you are can start off in monitor mode and gradually transition into the end state of either low-impact mode or closed mode. By doing so, you can avoid the denial of service that can so often happen with 802.1X deployments.
4. b. **Explanation: authentication open** ignores RADIUS Access-Reject responses but honors and enforces all other authorization results.
5. a. **Explanation: authentication open** allows traffic to flow with or without authentication. When an authorization result is sent back from the authentication server, the switch ignores RADIUS Access-Reject responses but honors and enforces all other authorization results.
6. b. **Explanation:** Policy sets are groupings of authentication and authorization policies. The use of policy sets makes for a nice clean way to differentiate separate rules for each stage of the deployment.
7. d. **Explanation:** Wireless LANs cannot have a mixture of authentication and non-authentication. A WLAN must either use Wi-Fi Protected Access (which facilitates the 802.1X authentication) or open; it cannot be both.

8. a. **Explanation:** The NDG assignment of the NAD is used to determine which policy set ISE uses for incoming authentication. It involves moving the NAD from the monitor mode NDG to either the low-impact or closed mode NDGs.
9. a. **Explanation:** Wired clients do not get to pick their network; there is no SSID as there is for wireless. Therefore, all the different types of possible authentication mechanisms must work within a single port configuration. Otherwise, an administrator would have to change the port configuration for each type of device that needs to access the network, which would be extremely operationally expensive.
10. a. **Explanation:** Just like the default behavior of the original IEEE 802.1X, closed mode does not allow any traffic into the switch port until after a result has been received for the attempted authentication or a timeout occurs.

Chapter 20

1. c. **Explanation:** After a standalone node has been promoted to primary on the Deployment screen, you click Register and enter the FQDN and the credentials for any other node that you want to join the new primary and form an ISE cube.
2. a and c. **Explanation:** Two backup types exist: a configuration (Policy Administration node) backup and an operational (MnT node) backup. These backups are separate and must be scheduled separately. An HTTP repository is a read-only repository and cannot be used for uploading backup files.
3. d. **Explanation:** The **show udi** CLI command and the GUI provide the three required items: SPID, VPID, and serial number.
4. b. **Explanation:** There is no automatic failover, but there is a manual promotion from the secondary's GUI.
5. d. **Explanation:** There is no automatic failover, but the ISE nodes are configured to send logging to both the primary and secondary MnT nodes automatically. If one fails, the other still receives the logs.
6. c. **Explanation:** Node groups are made up of PSNs, and the PSNs maintain a heartbeat with each other. In the event that a PSN goes down while a session is being authenticated, one of the other PSNs in the node group sends a CoA to the NAD so the endpoint can restart the session establishment with a new PSN.
7. b. **Explanation:** Cisco ISE is commonly deployed with load balancers. There are caveats to pay attention to, such as not using Source NAT (SNAT). The NAD does not need to know how many PSNs are in use. It is simply configured for a single RADIUS server per VIP address. The exception to the rule is with the CoA configuration. It is often the case that each PSN must be entered separately as a dynamic author client (CoA client).
8. b. **Explanation:** Patches are downloaded from Cisco.com and applied to the PAN under Administration > System > Maintenance > Patch Management. The PAN pushes a patch to the other nodes in the deployment.
9. d. **Explanation:** The status of a backup can be viewed from the GUI or the CLI, but the status of a restore can only be viewed from the CLI.
10. d. **Explanation:** It is not configurable; rather, all nodes are patched, in alphabetical order. The PAN is patched first and pushes the patch to all other nodes.

Chapter 21

1. d. **Explanation:** The Evaluate Configuration Validator tool compares a switch configuration to a template configuration built into ISE and points out any differences between these configurations.
2. c. **Explanation:** The RADIUS Authentication Troubleshooting tool examines different aspects of a session and provides details that may not have been available in the detailed authentication report, as well as suggestions for items to check next.
3. a, b, and e. **Explanation:** A support bundle can include the full configuration database, debug logs, local logs, core files, MnT reporting logs, system logs, and the policy configuration. These options are selectable in the GUI, and all of them are included automatically from the CLI.
4. b. **Explanation:** Live Sessions correlates activity related to the entire session, not just the raw entries related to a passed or failed authentication.
5. a. **Explanation:** Live Log displays events related to the raw syslog messages sent from the PSN to the MnT node, focused on passed or failed authentications.
6. d. **Explanation:** Logging targets are configured centrally, and the settings are pushed down to each PSN. Each PSN is configured to send syslog messages to all configured logging targets concurrently.
7. b. **Explanation:** An administrator uses the Suppress Anomalous Clients setting under Administration > System > Protocols > RADIUS to enable log de-duplication.
8. b. **Explanation:** Cisco AnyConnect DART is the module used to collect all log files from the endpoint along with other important information. DART combines all this information into a single condensed file for analysis by Cisco TAC.
9. c. **Explanation:** While a firewall can sometimes be a good place to determine why communication is not successful, the three main locations to troubleshoot network access are ISE, the endpoint, and the NAD.
10. b. **Explanation: debug epm** is the go-to debugging command for URL redirection, applied dACLs, assigned SGTs, and other activity related to advanced capabilities of authentication sessions.

Chapter 22

1. a. **Explanation**: pxGrid can integrate with third-party systems to allow for context sharing, Context-In, context brokering, and rapid threat containment. However, it cannot ingest threat feeds.
2. a, b, and d. **Explanation**: EPS, which is considered pxGrid 1.0, provided three static actions/labels: Quarantine, Unquarantine, and Shutdown.
3. d. **Explanation**: EPS always had the ability to change an endpoint's SGT or level of access with a corresponding authorization policy. However, EPS had only the actions Quarantine and Unquarantine. This lack of diverse labels gave ISE administrators limited options if they wanted to have multiple policies to choose from. ANC allows ISE administrators to create multiple labels to use with ANC and authorization policies.
4. a. **Explanation**: Thanks to the Context-In feature, ISE can receive contextual information through pxGrid to help ISE with its own profiling policies.

5. a. **Explanation:** Starting in ISE Version 2.4, the maximum number of pxGrid nodes increased from 2 to 4.
6. a, c, and d. **Explanation:** In order to add a participant to pxGrid, that user must trust the ISE certificate authority, have installed a pxGrid certificate that was either issued by ISE or the CA that issued ISE's pxGrid certificate, and have the IP address or FQDN of the pxGrid controller (ISE) configured.
7. a. **Explanation:** pxGrid Version 1 was designed by extending the Extensible Messaging and Presence Protocol (XMPP), which is also the communication protocol used by Jabber.
8. b. **Explanation:** Beginning in ISE Version 2.3, ISE added a modernized WebSocket-based interface to pxGrid to make integration easier. The DevNet partners were no longer required to integrate a Java or C library into their applications. Instead, they could use common representational state transfer (REST) connections.
9. c. **Explanation:** In WSA Version 11.7 and later, the Web Security Appliance can retrieve Active Directory groups and local ISE groups by using the ISE External Restful Service (ERS). This feature provides the ability to configure the WSA's policies using AD groups.
10. c. **Explanation:** ANC was added to create multiple ANC policies that can be assigned to an endpoint.

Chapter 23

1. d. **Explanation:** Threat Centric Network Access Control (TC-NAC) enables you to create authorization policies based on indicators of compromise and vulnerability attributes received from the threat and vulnerability products that ISE can be integrated with. Assigning policies that quarantine an endpoint is part of Cisco Rapid Threat Containment, not TC-NAC.
2. c. **Explanation:** AMP does not send CVSS scores; vulnerability assessment tools do. AMP sends threat intelligence as indicators and incidents, and they are not exposed to the authorization policy at all as of ISE Version 2.7.
3. b. **Explanation:** A vulnerability is a weakness in computer software or hardware that is also referred to as a security hole. An exploit is a weaponized software package that takes advantage of that security hole. A virus is a type of exploit.
4. b. **Explanation:** The AMP adapter does not require a proxy, but it does require the ability to reach the AMP cloud on the Internet. If a web proxy is required for Internet access, the AMP adapter requires it to be a Socks proxy, not an HTTP/HTTPS proxy. Cisco's WSA does offer Socks proxy capabilities, but any Socks proxy can work—not just the WSA proxy.
5. a. **Explanation:** Policies are flexible enough to provide nearly any flow an organization wants. Waiting for the vulnerability scanner to respond could create a very slow login process for end users, which would be very undesirable. The AMP agent does not communicate to AnyConnect for posture assessment. ISE communicates to the AMP cloud for threat intelligence.
6. c. **Explanation:** When a vulnerability assessment occurs and the results are sent back to ISE, it is possible that a CoA needs to occur. When a CoA occurs, it is recorded in

the Coa-Events report, and if a change of access based on the TC-NAC event occurs, it is displayed in the TC-NAC Live Log.

7. a. **Explanation:** CTA uses the standard threat intelligence data structure known as STIX to communicate cyber threat intelligence. TAXII is used to transport the STIX data in a secure manner.
8. d. **Explanation:** The Threat-Events report indicates the endpoint's base CVSS score from the vulnerability scanner. The Vulnerability Assessment report does not show scores.
9. c. **Explanation:** An administrator follows a link to the AMP public cloud and authorizes the ISE adapter software to subscribe to events from the selected endpoint groups. The authorization uses OAuth for software-to-software authorization. ISE does not support on-premises AMP private cloud installations.
10. a. **Explanation:** From the CTA dashboard, you create a STIX/TAXII API account and then copy the details into the ISE UI for future use. There is no on-premises version of CTA.

Chapter 24

1. b. **Explanation:** Device administration is very interactive in nature, and the ability to authorize each command is a critical capability because it means you will not be forced to authorize one time only. There are no devAdmin attribute/value pairs; that is made up. While Cisco did extend TACACS to TACACS+, that is not an acceptable answer here.
2. d. **Explanation:** One device administration license enables the Device Administration work center in the ISE GUI. The first license enables the work center, but to be compliant with Cisco's licensing, one license per PSN that enables TACACS+ is required.
3. c. **Explanation:** Each work center in the ISE user interface provides a common place to perform all activities related to that work center. The network resources are the routers, switches, and other devices that interact with ISE, and they are the same devices for the Network Access and Device Administration work centers.
4. a. **Explanation:** When a deployment is expected to be very large, it is recommended to maintain separate ISE cubes (deployments) for network access and device administration. TACACS is never required to be enabled on the PAN. It is not possible to have dedicated MnTs per function within a single ISE cube.
5. d. **Explanation:** The deployment page in the Device Administration work center is special. Unlike the standard deployment page, the work center allows you to enable or disable TACACS+ on one or all of the PSNs together as well as set the TCP port that TACACS+ will use. The default port is 49. There is no CLI to configure the TACACS+ port.
6. b. **Explanation:** Work Centers > Device Administration > Settings > Password Change Control is where you can globally enable or disable the ability to change passwords through the TACACS+ session and configure the strings to be used in the prompts.
7. a. **Explanation:** TACACS command sets are designed to permit or deny commands and their associated arguments. There is no such thing as device administration command sets. TrustSec tags have nothing to do with device administration.

8. d. **Explanation:** It is not possible to configure a policy to deny device administration access to any user who has not first passed the network access authentication. While a few customers have requested this feature, the network access sessions and the device administration sessions are completely separate.
9. a. **Explanation:** The Device Administration work center and its policies only relate to TACACS+, but RADIUS can be used for device administration AAA; however, those policies have to be configured in the network access policy sets.
10. c. **Explanation:** Using a condition such as device type to differentiate the incoming TACACS+ authentication and authorizations is a common practice.

Chapter 25

1. b. **Explanation:** A privilege level defines the level of access a user has in the IOS shell session. Cisco IOS provides 16 levels of access privilege. By default, the higher the level, the more commands and access one has available to them.
2. a. **Explanation:** With Cisco IOS devices, administrative access to the network device console, Telnet sessions, and Secure Shell (SSH) sessions can be controlled to authenticate, authorize shell privilege level, and perform granular per-command authorization and accounting.
3. c. **Explanation:** Privilege level 1, known as user exec mode, is the default level for a user after login. The shell prompt is the device hostname followed by the > symbol (for example, Switch>).
4. d. **Explanation:** Privilege levels 0, 1, and 15 are defined by default in Cisco IOS. Privilege levels 2 through 14 are available for customization.
5. c. **Explanation:** When creating AAA method lists, you can define the primary AAA method as well as fallback methods. For example, you might configure a AAA authentication method list to first attempt to use the TACACS+ server group and, if it is unreachable, use the local account store on the IOS device.
6. a. **Explanation:** When configuring a TACACS+ profile, IOS devices use the shell common task type.
7. b. **Explanation:** The **aaa authentication enable** command defines what methods to use to authenticate from user exec mode to privileged exec mode.
8. c. **Explanation:** TACACS+ uses TCP port 49 as the transport protocol to connect from the TACACS+ client to the TACACS+ server. The easiest way to troubleshot that TCP connection is to debug the IP TCP transactions on the switch.
9. c. **Explanation:** The **aaa authentication login** method list can be applied to the vty lines to evaluate all incoming SSH sessions.
10. c. **Explanation:** The **debug tacacs** command debugs just the TACACS+ packet traces.

Chapter 26

1. c. **Explanation:** WLC provides role-based access control (RBAC) on a per-menu basis. A WLC administrator sees all configuration items in that menu, but authorization is performed before changes are applied. So if a user named admin1 does not have access to the CONTROLLER menu, admin1 can see the menu and all its items and make changes on the web page, but when admin1 clicks Apply, an authorization failure message appears.

2. a. **Explanation:** WLC applies TACACS access control to the SSH session and to the HTTP sessions simultaneously. Cisco WLC does not use access classes, but IOS does.
3. c. **Explanation:** An AAA server can respond with a single role or multiple roles (which are additive). The exceptions to this rule are the three special roles that cannot be used with any other roles: LOBBY, MONITOR, and ALL.
4. d. **Explanation:** TACACS profiles are used with WLC, and there is a specific WLC type of TACACS profile. Command sets are not used by WLC; its RBAC capabilities are related to entire menus in the WLC UI.
5. b. **Explanation:** The Task Attribute view is a graphical front end for creating the raw TACACS values. You can create a result in either view and seamlessly switch between the views without losing your work.
6. c. **Explanation:** IOS and WLC rules can be in the same policy set, although it is strongly recommended to separate them with the use of NDGs in the top-level condition. RADIUS is not used with device administration policy sets; only TACACS is. RADIUS may only be used with network access policy sets. However, you can use a network access policy set to perform RADIUS-based device administration AAA.
7. d. **Explanation:** Due to a default "security setting" for ISE, any failed authentications show USERNAME in Live Log until you enable the Disclose Invalid Usernames setting under Administration > Settings > Security Settings.
8. a. **Explanation:** MONITOR provides a read-only style of access to the WLC. All pages are visible to the user, but no changes are authorized.
9. b. **Explanation:** The ISE PSN(s) must be listed in the authentication, authorization, and accounting lists under Security > AAA > TACACS+ configuration of the UI. When you add a TACAC+ authentication server, the entry in the authorization server list is added automatically.
10. d. **Explanation:** The LOBBY role permits access to the lobby ambassador functions, which remain from when WLC performed its own guest lifecycle management. Lobby ambassador on the WLC has since been superseded by the guest capabilities in ISE.

Answers to the Q&A Sections

Chapter 1

1. TACACS+
2. RADIUS
3. TACACS+
4. Attribute/value (AV) pair
5. Change of Authorization (CoA)

Chapter 2

1. Something you know, something you have, and something you are
2. SAML identity provider (IdP) and social login
3. Identity source sequence
4. Active Directory
5. Identity

Chapter 3

1. EAP-FASTv2 and TEAP
2. A supplicant
3. AnyConnect Diagnostics and Reporting Tool (DART)
4. Windows, when it is a member of an Active Directory domain
5. AnyConnect Network Authentication Module (NAM)

Chapter 4

1. Many organizations have a large percentage of endpoints that are IoT devices or otherwise headless devices, guest endpoints, and VPN users.
2. EasyConnect is easy to deploy quickly, does not require you to configure supplicants, and is an alternative to the headaches of 802.1X.
3. Authentication is based on MAC address, which is easily spoofed. This is why a defense-in-depth approach should be utilized.
4. CWA is fully compatible with CoA, which offers more flexibility for authentication and authorization. It also allows for customized portals.
5. PAP or MS-CHAPv2

Chapter 5

1. A standards-based mechanism to change the AAA attributes of a session after it has been authenticated
2. Profiling
3. It gathers a series of attributes about an endpoint to help match it to a profiling policy and identify the endpoint type.
4. It allows ISE administrators to ensure a minimum security posture of an endpoint based on the software or operating system running on it.
5. Support for most mobile devices and phones, remote wipe, and device backup are a few.

Chapter 6

1. Admin, Monitoring, Policy Services, and pxGrid
2. 2,000,000 endpoints
3. Administration and Monitoring personas running on a primary and secondary node, with up to five dedicated policy services nodes
4. Policy Services
5. Policy Services

Chapter 7

1. Some of the criteria that can be used for a global endpoint search in ISE include username, user type, MAC address, IP address, authorization profile, endpoint profile, failure reason, identity group, identity store, network device name, network device type, operating system, posture status, location, and security group.
2. The Context Visibility dashboards are a series of prebuilt contextual dashboards with data on the endpoints. ISE administrators can quickly filter endpoints to customize their view and export the data.

3. The RADIUS Live Log dashboard provides real-time information about RADIUS authentication and gives an ISE administrator the ability to drill into a single authentication for more details when troubleshooting.
4. Dictionaries, conditions, and results
5. Authentication, authorization, profiling, posture, client provisioning, and TrustSec

Chapter 8

1. IP address, net mask, default gateway, DNS domain, name server, NTP server, time zone, and configuring the initial admin username and password
2. DNS name, IP address, uniform resource name, and directory name
3. A single certificate can be shared across multiple hosts in an organization, as long as they have the same domain name. The disadvantage is that it can't be used for EAP authentication, so it's limited to other certificate usages in ISE.
4. Network device groups can be used as a condition when creating network access control policies. Depending on the network device group that a NAD is in, a policy can be built to authenticate and authorize endpoints differently, depending on which NADs they connect to.
5. Check the permissions on the account that you're attempting to join with, make sure that the time settings on the domain controller and the ISE node are within 5 minutes of each other, and ensure that the required ports are open between ISE and the domain controller

Chapter 9

1. Authentication is the validation of a credential or an identity.
2. Authorization is the validation of what access should be granted after an identity has been validated.
3. Accept only allowed protocols, select the correct identity store, validate the identity, and then pass the request to the authorization policy.
4. A logical container for an authentication policy and an authorization policy that allows an ISE administrator to segment and organize policies into policy blocks
5. It has Service-Type set to Call-Check for MAB.

Chapter 10

1. To examine conditions in authorization rules in order to ultimately send an authorization result to a network access device (NAD)
2. Policies can be built using a myriad of attributes, such as location, time, whether a device was registered, whether a mobile device has been jail broken, or nearly any other attribute imaginable. Even the authentication can be used as an attribute: Was authentication successful? Which authentication protocol was used? What is the content of specific fields of the certificate that was used?
3. If certain conditions are met, assign the following permissions.
4. No, authentication alone does not guarantee segmentation. In order to provide a true defense-in-depth approach, the access granted must be least-privilege access.

5. Hierarchical compound conditions can contain a mixture of simple and complex conditions to meet the demands of a security requirement. Such a condition can also minimize the number of authorization rules needed in the policy.

Chapter 11

1. They define the method lists for RADIUS authentication, authorization, and accounting. The RADIUS requests are sent and accepted from the servers defined in these lists.
2. The authentication method that was chosen, the IP address of the endpoint, the service-type, and the class
3. It enabled the NAD to fall back to another authentication method if the first one failed or timed out. Instead of having to configure the switch port for just MAB or just 802.1X, ISE admins could instead have both configured on a port.
4. Single-Host, MDA, Multi-Auth, and Multi-Hosts
5. authentication port-control auto

Chapter 12

1. RUN state
2. URL redirection
3. MAB or, more specifically, the mab Authc Success method
4. On the NAD (switch or WLC)
5. The switch uses separate ACLs to control access through the switch and to control what is redirected to the web portal. The WLC uses a single ACL to perform both functions.

Chapter 13

1. The guest portal is used for Centralized Web Authentication (CWA) and can make it easier for an employee who is going through the CWA process.
2. Different guest types can have different timelines for account life, different active hours, and different permissions.
3. Hotspot portals are designed for use when guest authentication is not necessary, such as at a coffee shop. The user just clicks to accept an AUP, and the MAC address is recorded to prevent the user from having to revisit the AUP the next time.
4. There is a setting for endpoint purge, which is how frequently endpoints should be removed from the endpoint group.
5. The portals are exactly the same except for the settings that are preconfigured. The portals can be configured to work exactly the same way if you change the settings.

Chapter 14

1. ISE Profiler is the part of ISE that is responsible for endpoint detection and classification. It uses probes to collect attributes about the connected endpoints.
2. Anomalous Behaviour Detection provides the ability to discover some MAC spoofing through the connection type, DHCP Class-Identifier, and changes to the endpoint policy.
3. NETFLOW, DHCP, DHCPSPAN, HTTP, RADIUS, NMAP, DNS, SNMPQUERY, SNMPTRAP, Active Directory, and pxGrid

4. The three different types of CoA action are Reauth, no CoA, and Port Bounce.
5. This allows an ISE administrator to change a setting for a specific profile that needs a different CoA type in order to function. For example, certain non-802.1x devices might need a Port Bounce in order to trigger DHCP.

Chapter 15

1. Use Base64 encoding, also called PEM encoding, and avoid DER encoding when possible.
2. Add the public certificate from each of the possible issuing CAs and each CA in the issuing chain. In this scenario, the root, node, and signing CAs must all be imported and trusted.
3. http://[fqdn-of-CA]/certsrv
4. The private key should never be exported. The public certificate is meant to be shared.
5. A certificate authentication profile (CAP)

Chapter 16

1. Employees want to use their own devices on an organization's network because those devices help them be productive, thanks to the native applications running on those devices, which provide the user experience they are used to and love. In order to grant this experience, companies need to be able to identify a user, a device, the device's compliance, and more.
2. MDM policies ensure that devices have a minimum level of security enabled, such as encryption, remote wipe capabilities, and screen lock (PIN lock).
3. Con: Some manual supplicant configuration is required by the end user. Pro: A single SSID is used for onboarding and connection.
4. Pro: There is no manual configuration of the NAD supplicant required of the end user. Con: A user has to connect to a different SSID, either automatically or manually.
5. Certificates can be issued by ISE's internal certificate authority and managed locally, or ISE can act as a registration authority to an external certificate authority using SCEP. Alternatively, MDM providers can deliver their own certificates to endpoints through the MDM onboarding process.

Chapter 17

1. To communicate IP address-to-SGT mappings from one device to another so that routers, switches, and firewalls can learn identity information from access devices
2. To assign a tag to a user's or device's traffic at the point of ingress and enforce the access elsewhere in the network. This provides flexible enforcement of east–west and north–south traffic.
3. *Classification* is the process of assigning a tag to a user or device. *Transport* means propagating the tag throughout the network, either inline or through SXP. *Enforcement* is taking action on a SGACL policy or SGFW to either allow or deny the traffic, based on the source and destination tags.
4. Dynamic assignment via ISE authorization policy using normal authentication methods (802.1X, MAB, WebAuth) and static assignment, using the Layer 2 port, Layer 3 port, VLAN, or IP address, hostname, or subnet.

5. The switch must know the source SGT, destination SGT, SGACL policy, and environment data.

Chapter 18

1. Anti-malware conditions replaced both anti-virus and anti-spyware conditions in compliance module 4.x and later. Anti-virus and anti-spyware conditions are only valid in compliance module 3.x and below.
2. Dictionary compound conditions are made up of simple conditions that you created as an ISE administrator. Compound conditions are preconfigured conditions from Cisco.
3. Patch management systems are widely used and deployed in businesses to keep endpoints up to date, patched, and posture compliant. The patch management condition integrates with those tools and uses them for remediation; it also ensures that the patch management agents are running and up to date.
4. When an endpoint fails a posture check (a condition), the configured remediation action is used to correct that failure and allow the endpoint to become compliant. The requirement binds the condition to the remediation action.
5. The temporal agent is not installed on the endpoint; it runs in the user's context. It gets launched from the web browser and runs in memory with the web browser. It is not capable of automatic remediation at all. The stealth mode agent is actually the AnyConnect System Scan module with a configuration option that hides the GUI. If no other modules are enabled, the icon doesn't even display in the system tray (Windows) or menu bar (macOS). All user notification can be suppressed. The stealth mode agent runs as a service, not in the user space, and it can therefore perform auto-remediation.

Chapter 19

1. authentication open
2. Access-Reject
3. Network device groups (NDGs)
4. Wi-Fi/wireless networks
5. A phased approach

Chapter 20

1. Anycast
2. Source NAT (SNAT)
3. Configuration and operational
4. Primary PAN and secondary PAN
5. The primary PAN must trust the administrative certificate on the additional nodes, and those nodes must trust the administrative certificate on the primary PAN.

Chapter 21

1. When the string USERNAME shows up in the Identity column, it means a user has failed authentication. The column displays USERNAME when the default security setting to not disclose invalid usernames is set.
2. Disable the log de-duplication to ensure that all entries appear in Live Log.
3. The Posture Troubleshooting tool examines the detailed posture report and displays reasons that an endpoint might have failed a posture condition.

4. The Endpoint Debug tool sets all logs to debug for the endpoint and its session only and combines the entries into a single file for download.
5. DART collects any and all related settings and logs from an endpoint with AnyConnect to be sent to TAC.

Chapter 22

1. The ultimate goal of pxGrid is to integrate various Cisco security products and third-party ecosystem partners to form a security architecture to share information and remediate threats.
2. The goal of Rapid Threat Containment is to utilize information or triggers from other Cisco security products or ecosystem partners to respond to threats more quickly and to dynamically change an endpoint's access.
3. EPS provided ISE a way to assign a policy (label) on an endpoint that could be used in authorization policies. The labels available were static in nature: Quarantine and Unquarantine. This left little room for additional authorization policies.
4. ANC picked up where EPS left off but also provides the ability to create multiple custom policies/labels, which can be used in multiple authorization policies. This added flexibility and options to Rapid Threat Containment, where an ISE administrator no longer has to assign a single policy (quarantine) and have a single authorization policy tied to it.
5. Context-In is a feature of pxGrid that enables endpoint attributes and data to be shared from an ecosystem partner and used for profiling purposes in ISE. Context-Out enables ISE to share contextual information with pxGrid participants to enrich their database.

Chapter 23

1. STIX for the threat intelligence, and TAXII is used for transport.
2. An indictor shows that an endpoint matched an IoC, and an incident is an AMP event.
3. Both. A vulnerability scan can be triggered with any authorization profile and allows flexibility.
4. An entry appears in TC-NAC Live Log only if a CoA was directly related to a TC-NAC adapter.
5. A vulnerability is a weakness in computer hardware or software. An exploit is software created to take advantage of a vulnerability.

Chapter 24

1. TACACS profile
2. TCP port 49
3. TACACS command sets are used for dictating what commands and arguments are permitted or denied as part of the authorization result.
4. Retire the secret. It will remain valid for a period of time, and the new and old secrets will both be usable for some time.
5. TACACS server sequence

Chapter 25

1. TACACS+ profile
2. PAP/ASCII, CHAP, and MS-CHAPv1
3. Deny Always should be used in a command set when there is a command that should be denied and should have higher precedence over a match than any Permit in that command set.
4. Configuration of TACACS+ authentication and fallback, configuration of TACACS+ command authorization, and configuration of TACACS+ command accounting
5. Custom privilege levels need to be configured on each and every network device that will be used for device administration. While commands can be defined for the custom privilege level, this customization adds administrative burden in terms of managing the configuration of the privilege level when it needs to be modified.

Chapter 26

1. Raw view
2. ALL
3. LOBBY
4. Disclose invalid usernames
5. TACACS Live Log



APPENDIX B

CCNP Security Implementing and Configuring Cisco Identity Services Engine (SISE 300-715) Exam Updates

Over time, reader feedback allows Pearson to gauge which topics give our readers the most problems when taking the exams. To assist readers with those topics, the authors create new materials clarifying and expanding on those troublesome exam topics. As mentioned in the Introduction, the additional content about the exam is contained in a PDF on this book's companion website, at http://www.ciscopress.com/title/9780136642947.

This appendix is intended to provide you with updated information if Cisco makes minor modifications to the exam on which this book is based. When Cisco releases an entirely new exam, the changes are usually too extensive to provide in a simple update appendix. In those cases, you might need to consult the new edition of the book for the updated content. This appendix attempts to fill the void that occurs with any print book. In particular, this appendix does the following:

- Mentions technical items that might not have been mentioned elsewhere in the book
- Covers new topics if Cisco adds new content to the exam over time
- Provides a way to get up-to-the-minute information about content for the exam

Always Get the Latest at the Book's Product Page

You are reading the version of this appendix that was available when your book was printed. However, given that the main purpose of this appendix is to be a living, changing document, it is important that you look for the latest version online at the book's companion website. To do so, follow these steps:

Step 1. Browse to www.ciscopress.com/title/9780136642947.

Step 2. Click the **Updates** tab.

Step 3. If there is a new Appendix B document on the page, download the latest Appendix B document.

NOTE The downloaded document has a version number. When you compare the version of the print Appendix B (Version 1.0) with the latest online version of this appendix, you should do the following:

- **Same version:** Ignore the PDF that you downloaded from the companion website.
- **Website has a later version:** Ignore this Appendix B in your book and read only the latest version that you downloaded from the companion website.

Technical Content

The current Version 1.0 of this appendix does not contain additional technical coverage.



APPENDIX C

Sample Switch Configurations

This Appendix includes some full sample configurations of various device types with multiple IOS Versions, all designed to follow the guidelines and practices laid out throughout the different chapters of this book.

Catalyst 3000 Series, 12.2(55)SE

With this model and IOS version.

```
3560-X#sho run
Building configuration...


Current configuration : 22928 bytes
!
version 12.2
hostname 3560-X
logging monitor informational
username radius-test password 0 Cisco123
!
aaa new-model
!
!
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
aaa accounting update newinfo periodic 1440
!
!
aaa server radius dynamic-author
 client 10.1.103.231 server-key Cisco123
 client 10.1.103.4 server-key Cisco123
```

```
!
aaa session-id common
authentication mac-move permit
ip routing
!
ip domain-name ise.local
ip name-server 10.1.100.100
ip device tracking
!
!
crypto pki trustpoint TP-self-signed-4076357888
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-4076357888
 revocation-check none
 rsakeypair TP-self-signed-4076357888
!
!
crypto pki certificate chain TP-self-signed-4076357888
 certificate self-signed 01
quit
!
dot1x system-auth-control
!
interface Loopback0
 ip address 192.168.254.60 255.255.255.255
!
interface range <ALL EDGE PORTS>
 switchport access vlan 41
 switchport mode access
 switchport voice vlan 99
```

```
 ip access-group ACL-ALLOW in
 authentication event fail action next-method
 authentication event server dead action authorize vlan 41
 authentication event server dead action authorize voice
 authentication event server alive action reinitialize
 authentication host-mode multi-auth
 authentication open
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 10
 spanning-tree portfast
!
interface Vlan1
 no ip address
!
interface Vlan40
 ip address 10.1.40.60 255.255.255.0
!
!
ip http server
ip http secure-server
!
ip access-list extended ACL-AGENT-REDIRECT
 remark explicitly prevent DNS from being redirected to address a
bug
 deny   udp any any eq domain
 remark redirect HTTP traffic only
```

```
 permit tcp any any eq www
 remark all other traffic will be implicitly denied from the
redirection
ip access-list extended ACL-ALLOW
 permit ip any any
ip access-list extended ACL-DEFAULT
 remark DHCP
 permit udp any eq bootpc any eq bootps
 remark DNS
 permit udp any any eq domain
 remark Ping
 permit icmp any any
 remark PXE / TFTP
 permit udp any any eq tftp
 remark Drop all the rest
 deny   ip any any log
ip access-list extended ACL-WEBAUTH-REDIRECT
 remark explicitly prevent DNS from being redirected to
accommodate certain switches
 deny   udp any any eq domain
 remark redirect all applicable traffic to the ISE Server
 permit tcp any any eq www
 permit tcp any any eq 443
 remark all other traffic will be implicitly denied from the
redirection
!
ip radius source-interface Loopback0
logging origin-id ip
logging source-interface Loopback0
logging host 10.1.103.4 transport udp port 20514
!
snmp-server community CiscoPressRO RO
```

```
snmp-server trap-source Loopback0
snmp-server source-interface informs Loopback0
snmp-server host 10.1.103.231 version 2c CiscoPressRO
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 dead-criteria time 5 tries 3
radius-server host 10.1.103.231 auth-port 1812 acct-port 1813 key
Cisco123
radius-server host 10.1.103.4 auth-port 1812 acct-port 1813 key
Cisco123
radius-server vsa send accounting
radius-server vsa send authentication
!
end
```

Catalyst 3000 Series, 15.0(2)SE

```
C3750X#sho run brief
Building configuration...

Current configuration : 18936 bytes
!
version 15.0
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname C3750X
!
boot-start-marker
boot-end-marker
```

```
!
logging monitor informational
!
username radius-test password 0 Cisco123
aaa new-model
!
!
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
aaa accounting update newinfo periodic 1440
!
!
aaa server radius dynamic-author
 client 10.1.103.231 server-key Cisco123
 client 10.1.103.4 server-key Cisco123
!
aaa session-id common
clock timezone EDT -1 0
authentication mac-move permit
ip routing
!
!
ip dhcp snooping vlan 10-13
ip dhcp snooping
ip domain-name ise.local
ip device tracking
!
!
device-sensor filter-list cdp list my_cdp_list
 tlv name device-name
```

```
 tlv name platform-type
!
device-sensor filter-list lldp list my_lldp_list
 tlv name port-id
 tlv name system-name
 tlv name system-description
!
device-sensor filter-list dhcp list my_dhcp_list
 option name host-name
 option name class-identifier
 option name client-identifier
device-sensor filter-spec dhcp include list my_dhcp_list
device-sensor filter-spec lldp include list my_lldp_list
device-sensor filter-spec cdp include list my_cdp_list
device-sensor accounting
device-sensor notify all-changes
!
epm logging
!
crypto pki trustpoint TP-self-signed-254914560
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-254914560
 revocation-check none
 rsakeypair TP-self-signed-254914560
!
!
crypto pki certificate chain TP-self-signed-254914560
 certificate self-signed 01
cts role-based enforcement
!
dot1x system-auth-control
```

```
!
interface Loopback0
 ip address 192.168.254.1 255.255.255.255
!
interface <ALL EDGE PORTS>
 switchport access vlan 10
 switchport mode access
 switchport voice vlan 99
 ip access-group ACL-ALLOW in
 authentication event fail action next-method
 authentication event server dead action authorize vlan 10
 authentication event server dead action authorize voice
 authentication event server alive action reinitialize
 authentication host-mode multi-auth
 authentication open
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 10
 spanning-tree portfast
 ip dhcp snooping information option allow-untrusted
!
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan10
```

```
 ip address 10.1.10.1 255.255.255.0
!
interface Vlan20
 ip address 10.1.20.1 255.255.255.0
!
interface Vlan30
 ip address 10.1.30.1 255.255.255.0
!
interface Vlan99
 ip address 10.1.99.1 255.255.255.0
!
!
ip http server
ip http secure-server
!
!
ip access-list extended ACL-AGENT-REDIRECT
 remark explicitly prevent DNS from being redirected
 deny   udp any any eq domain
 remark redirect HTTP traffic only
 permit tcp any any eq www
 remark all other traffic will be implicitly denied from the
redirection
ip access-list extended ACL-ALLOW
 permit ip any any
ip access-list extended ACL-DEFAULT
 remark DHCP
 permit udp any eq bootpc any eq bootps
 remark DNS
 permit udp any any eq domain
 remark Ping
```

```
 permit icmp any any
 remark PXE / TFTP
 permit udp any any eq tftp
 remark Drop all the rest
 deny   ip any any log
ip access-list extended ACL-WEBAUTH-REDIRECT
 remark explicitly prevent DNS from being redirected to address
 deny   udp any any eq domain
 remark redirect all applicable traffic to the ISE Server
 permit tcp any any eq www
 permit tcp any any eq 443
 remark all other traffic will be implicitly denied from the
redirection
ip access-list extended AGENT-REDIRECT
 remark explicitly prevent DNS from being redirected to address
 deny   udp any any eq domain
 remark redirect HTTP traffic only
 permit tcp any any eq www
 remark all other traffic will be implicitly denied from the
redirection
!
ip radius source-interface Loopback0
ip sla enable reaction-alerts
logging origin-id ip
logging source-interface Loopback0
logging host 10.1.103.4 transport udp port 20514
!
snmp-server community Cisco123 RO
snmp-server community TrustSecRO RO
snmp-server trap-source Loopback0
snmp-server source-interface informs Loopback0
snmp-server host 10.1.103.4 version 2c Cisco123  mac-notification
```

```
!
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 dead-criteria time 5 tries 3
radius-server vsa send accounting
radius-server vsa send authentication
!
radius server CP-VIP
 address ipv4 10.1.103.231 auth-port 1812 acct-port 1813
 automate-tester username radius-test
 key Cisco123
!
radius server CP-04
 address ipv4 10.1.103.4 auth-port 1812 acct-port 1813
 automate-tester username radius-test
 key Cisco123
!
end
```

Catalyst 9000 Series, 16.9.5

```
9300#sho run brief
Building configuration...
Current configuration : 29669 bytes
!
! Last configuration change at 20:20:09 UTC Fri Jul 24 2020 by
admin
!
version 16.9
no service pad
service timestamps debug datetime msec
```

```
service timestamps log datetime msec
service password-encryption
service compress-config
! Call-home is enabled by Smart-Licensing.
service call-home
no platform punt-keepalive disable-kernel-core
!
aaa authentication dot1x default group ise-group
aaa authorization network default group ise-grou
aaa accounting update periodic 1440
aaa accounting dot1x default start-stop group ise-group
aaa accounting system default start-stop group ise-group
!
!
!
!
!
aaa server radius dynamic-author
 client 10.1.100.21 server-key Cisco123
 auth-type any
!
aaa session-id common
!
!
!
!
ip routing
!
!
!
!
```

```
!
ip name-server 10.1.100.40
ip domain name ise.local
!
device-tracking tracking auto-source
device-tracking policy TRACKING-POLICY
tracking enable
no ip cef optimize neighbor resolution
!
!
!
ip dhcp snooping vlan 100
no ip dhcp snooping information option
ip dhcp snooping
!
!
!
!
!
!
device-sensor filter-list lldp list TLV-LLDP
 tlv name port-id
 tlv name system-name
 tlv name system-description
 tlv name system-capabilities
!
device-sensor filter-list cdp list TLV-CDP
 tlv name device-name
 tlv name platform-type
!
```

```
device-sensor filter-list dhcp list TLV-DHCP
 option name host-name
 option name domain-name
 option name parameter-request-list
 option name class-identifier
 option name client-identifier
 option name client-fqdn
!
device-sensor filter-spec dhcp include list TLV-DHCP
device-sensor filter-spec lldp include list TLV-LLDP
device-sensor filter-spec cdp include list TLV-CDP
device-sensor accounting
device-sensor notify all-changes
vtp domain ise.local
vtp mode transparent
epm logging
authentication mac-move permit
access-session template monitor
access-session acl default passthrough
ipv6 neighbor tracking auto-source override
!
!
!
!
crypto pki trustpoint TP-self-signed-1546261629
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-1546261629
 revocation-check none
 rsakeypair TP-self-signed-1546261629
!
```

```
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
!
!
crypto pki certificate chain TP-self-signed-1546261629
 certificate self-signed 01
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
!
dot1x system-auth-control
license boot level ipbasek9
!
!
diagnostic bootup level minimal
!
spanning-tree mode pvst
spanning-tree extend system-id
!
!
username admin privilege 15 password Cisco123
!
redundancy
 mode sso
!
!
!
!
!
!
transceiver type all
```

```
 monitoring
hw-switch switch 1 logging onboard message
!
vlan 100
 name DATA
lldp run
!
!
interface <ALL EDGE PORTS>
switchport access vlan 41
switchport mode access
switchport voice vlan 99
ip access-group ACL-ALLOW in
device-tracking attach-policy TRACKING-POLICY
authentication event fail action next-method
authentication event server dead action reinitialize vlan 100
authentication event server dead action authorize voice
authentication event server alive action reinitialize
authentication host-mode multi-auth
authentication open
authentication order dot1x mab
authentication priority dot1x mab
authentication port-control auto
authentication periodic
authentication timer reauthenticate server
authentication timer inactivity server
authentication violation restrict
mab
snmp trap mac-notification change added
snmp trap mac-notification change removed
dot1x pae authenticator
```

```
dot1x timeout tx-period 10
spanning-tree portfast
spanning-tree bpduguard enable
!
interface Vlan100
ip address 10.1.100.75 255.255.255.0
ip helper-address 10.1.100.21
!
!
ip default-gateway 10.1.100.254
ip forward-protocol nd
ip http server
ip http authentication local
ip http secure-server
ip http secure-active-session-modules none
ip http active-session-modules none
ip ftp source-interface Vlan100
ip tftp source-interface Vlan100
ip route 0.0.0.0 0.0.0.0 10.1.100.254
ip ssh authentication-retries 2
ip ssh version 2
!
ip access-list extended ACL-AGENT-REDIRECT
 remark explicitly prevent DNS from being redirected to address
 deny   udp any any eq domain
 remark redirect HTTP traffic only
 permit tcp any any eq www
 remark all other traffic will be implicitly denied from the
redirection
ip access-list extended ACL-ALLOW
 permit ip any any
```

```
ip access-list extended ACL-DEFAULT
 remark DHCP
 permit udp any eq bootpc any eq bootps
 remark DNS
 permit udp any any eq domain
 remark Ping
 permit icmp any any
 remark PXE / TFTP
 permit udp any any eq tftp
 remark Drop all the rest
 deny   ip any any log
ip access-list extended ACL-WEBAUTH-REDIRECT
 remark explicitly prevent DNS from being redirected to address
 deny   udp any any eq domain
 remark redirect all applicable traffic to the ISE Server
 permit tcp any any eq www
 permit tcp any any eq 443
 remark all other traffic will be implicitly denied from the
redirection
!
ip radius source-interface Vlan100
ip sla enable reaction-alerts
logging trap debugging
logging origin-id ip
logging source-interface Vlan100
logging host 10.1.100.21 transport udp port 20514
!
!
ipv6 neighbor tracking auto-source override
!
!
```

```
snmp-server community ISEc0ld RO
snmp-server trap-source Vlan100
snmp-server source-interface informs Vlan100
snmp-server enable traps snmp authentication linkdown linkup
coldstart warmstart
snmp-server enable traps mac-notification change move threshold
snmp-server host 10.1.100.21 ISEc0ld
!
radius-server attribute 6 on-for-login-auth
radius-server attribute 6 support-multiple
radius-server attribute 8 include-in-access-req
radius-server attribute 25 access-request include
radius-server attribute 31 mac format ietf upper-case
radius-server attribute 31 send nas-port-detail
radius-server dead-criteria time 5 tries 3
radius-server deadtime 30
!
radius server ise
 address ipv4 10.1.100.21 auth-port 1812 acct-port 1813
 key ISEc0ld
!
line con 0
 stopbits 1
line aux 0
 stopbits 1
line vty 5 15
!
ntp source Vlan100
ntp server 10.1.100.40
mac address-table notification change
!
end
```

Catalyst 4500 Series, IOS-XE 3.3.0 / 15.1(1)SG


```
4503#show run brief
Building configuration...
Current configuration : 35699 bytes
!
!
version 15.1
!
hostname 4503
!
!
username radius-test password 0 Cisco123
aaa new-model
!
!
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
aaa accounting update newinfo periodic 1440
!
!
aaa server radius dynamic-author
 client 10.1.103.231 server-key Cisco123
 client 10.1.103.4 server-key Cisco123
!
aaa session-id common
clock timezone EDT -1 0
!
ip domain-name ise.local
```

```
!
ip device tracking
!
device-sensor filter-list cdp list my_cdp_list
 tlv name device-name
 tlv name platform-type
!
device-sensor filter-list lldp list my_lldp_list
 tlv name port-id
 tlv name system-name
 tlv name system-description
!
device-sensor filter-list dhcp list my_dhcp_list
 option name host-name
 option name class-identifier
 option name client-identifier
device-sensor filter-spec dhcp include list my_dhcp_list
device-sensor filter-spec lldp include list my_lldp_list
device-sensor filter-spec cdp include list my_cdp_list
device-sensor accounting
device-sensor notify all-changes
epm logging
!
!
crypto pki trustpoint CISCO_IDEVID_SUDI
 revocation-check none
 rsakeypair CISCO_IDEVID_SUDI
!
crypto pki trustpoint CISCO_IDEVID_SUDI0
 revocation-check none
!
```

```
!
crypto pki certificate chain CISCO_IDEVID_SUDI
 certificate 238FC0E90000002BFCA1
 certificate ca 6A6967B3000000000003
crypto pki certificate chain CISCO_IDEVID_SUDI0
 certificate ca 5FF87B282B54DC8D42A315B568C9ADFF
!
dot1x system-auth-control
!
!
vlan 40
 name jump
!
vlan 41
 name data
!
vlan 99
 name voice
!
interface <ALL EDGE PORTS>
 switchport access vlan 41
 switchport mode access
 switchport voice vlan 99
 ip access-group ACL-ALLOW in
 authentication event fail action next-method
 authentication event server dead action authorize vlan 41
 authentication event server dead action authorize voice
 authentication event server alive action reinitialize
 authentication host-mode multi-auth
 authentication open
 authentication order dot1x mab
```

```
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 10
 spanning-tree portfast
 ip dhcp snooping information option allow-untrusted
!
interface Vlan1
 no ip address
!
interface Vlan40
 ip address 10.1.40.2 255.255.255.0
!
ip http server
ip http secure-server
ip route 0.0.0.0 0.0.0.0 10.1.40.1
!
ip access-list extended ACL-AGENT-REDIRECT
 remark explicitly prevent DNS from being redirected to address
 deny   udp any any eq domain
 remark redirect HTTP traffic only
 permit tcp any any eq www
 remark all other traffic will be implicitly denied from the
redirection
ip access-list extended ACL-ALLOW
 permit ip any any
ip access-list extended ACL-DEFAULT
 remark DHCP
 permit udp any eq bootpc any eq bootps
```

```
 remark DNS
 permit udp any any eq domain
 remark Ping
 permit icmp any any
 remark PXE / TFTP
 permit udp any any eq tftp
 remark Drop all the rest
 deny   ip any any log
ip access-list extended ACL-WEBAUTH-REDIRECT
 remark explicitly prevent DNS from being redirected to address
 deny   udp any any eq domain
 remark redirect all applicable traffic to the ISE Server
 permit tcp any any eq www
 permit tcp any any eq 443
 remark all other traffic will be implicitly denied from the
redirection
!
logging 10.1.103.4
!
snmp-server community Cisco123 RO
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 dead-criteria time 5 tries 3
radius-server host 10.1.103.231 auth-port 1812 acct-port 1813 test
username radius-test key Cisco123
radius-server host 10.1.103.4 auth-port 1812 acct-port 1813 test
username radius-test key Cisco123
radius-server vsa send accounting
radius-server vsa send authentication
!
end
```

Catalyst 6500 Series, 12.2(33)SXJ

```
hostname 6503
logging monitor informational
username radius-test password 0 Cisco123
!
aaa new-model
!
!
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
aaa accounting update newinfo periodic 1440
!
!
aaa server radius dynamic-author
 client 10.1.103.231 server-key Cisco123
 client 10.1.103.4 server-key Cisco123
!
aaa session-id common
authentication mac-move permit
ip routing
!
ip domain-name ise.local
ip name-server 10.1.100.100
ip device tracking
!
!
crypto pki trustpoint TP-self-signed-4076357888
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-4076357888
 revocation-check none
```

```
 rsakeypair TP-self-signed-4076357888
!
!
crypto pki certificate chain TP-self-signed-4076357888
 certificate self-signed 01
quit
!
dot1x system-auth-control
!
interface Loopback0
 ip address 192.168.254.1 255.255.255.255
!
interface <ALL EDGE PORTS>
 switchport access vlan 10
 switchport mode access
 switchport voice vlan 99
 ip access-group ACL-ALLOW in
 authentication event fail action next-method
 authentication event server dead action authorize vlan 10
! - Note:  No Critical Voice VLAN Support on 6500's yet
 authentication event server alive action reinitialize
 authentication host-mode multi-auth
 authentication open
 authentication order dot1x mab
 authentication priority dot1x mab
 authentication port-control auto
 authentication violation restrict
 mab
 dot1x pae authenticator
 dot1x timeout tx-period 10
 spanning-tree portfast
```

```
!
interface Vlan1
 no ip address
!
interface Vlan40
 ip address 10.1.40.1 255.255.255.0
!
!
ip http server
ip http secure-server
!
ip access-list extended ACL-AGENT-REDIRECT
 remark explicitly prevent DNS from being redirected to address a
bug
 deny   udp any any eq domain
 remark redirect HTTP traffic only
 permit tcp any any eq www
 remark all other traffic will be implicitly denied from the
redirection
 deny ip any any
ip access-list extended ACL-ALLOW
 permit ip any any
ip access-list extended ACL-DEFAULT
 remark DHCP
 permit udp any eq bootpc any eq bootps
 remark DNS
 permit udp any any eq domain
 remark Ping
 permit icmp any any
 remark PXE / TFTP
 permit udp any any eq tftp
 remark Drop all the rest
```

```
 deny   ip any any log
ip access-list extended ACL-WEBAUTH-REDIRECT
 remark explicitly prevent DNS from being redirected
 deny   udp any any eq domain
 remark redirect all applicable traffic to the ISE Server
 permit tcp any any eq www
 permit tcp any any eq 443
 deny ip any any
!
ip radius source-interface Loopback0
logging origin-id ip
logging source-interface Loopback0
logging host 10.1.103.4 transport udp port 20514
!
snmp-server community CiscoPressRO RO
snmp-server trap-source Loopback0
snmp-server source-interface informs Loopback0
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 dead-criteria time 5 tries 3
radius-server host 10.1.103.231 auth-port 1812 acct-port 1813 key
Cisco123
radius-server host 10.1.103.4 auth-port 1812 acct-port 1813 key
Cisco123
radius-server vsa send accounting
radius-server vsa send authentication
!
end
```



Index

SYMBOLS

*** (asterisk), 494, 922**

2FA (two-factor authentication), 26, 1000

802.1Q, 284

802.1X, 7. ***See also*** **EAP (Extensible Authentication Protocol)**

authentication. *See* authentication

Cisco AnyConnect NAM supplicant, 59–73

AnyConnect NAM profiles, 71–72

Authentication Policy view, 60, 62

Client Policy view, 60, 61–62

EAP chaining, 73

Network Groups view, 60, 71

Networks view, 60, 62–70

overview of, 59–60

definition of, 993

FlexAuth (Flexible Authentication)

configuration, 269–272

definition of, 994

FAST (Flexible Authentication via Secure Tunneling), 45, 48–49, 215

global 802.1X commands, 266–267

history of, 38

identity stores. *See* identity stores

NADs (network access devices). *See* NADs (network access devices)

overview of, 41–42, 719–720

phased deployment

advantages of, 717–718

closed mode, 728–730

default port behavior, 719

low-impact mode, 725–727

monitor mode, 722–725, 730–731

preparation for, 720–721

transitioning to end state, 730–731

wireless networks, 731

SGT (security group tag) assignment with, 577

supplicants. *See* supplicants

A

AAA (authentication, authorization, and accounting). ***See*** **accounting; authentication; authorization**

aaa commands

aaa accounting commands 1 default start-stop group ISE-TACACS, 951

aaa accounting commands 15 default start-stop group ISE-TACACS, 951

aaa accounting dot1x default start-stop group ise-group, 562

aaa accounting dot1x default start-stop group radius, 262, 567

aaa accounting exec default start-stop group ISE-TACACS, 951

aaa accounting system default start-stop group ise-group, 562

aaa accounting system default start-stop group radius, 567

aaa authentication dot1x default group ise-group, 562

aaa authentication dot1x default group radius, 262, 567

aaa authentication enable default group ISE-TACACS enable, 947

aaa authentication login, 947
aaa authorization commands, 949
aaa authorization config-commands, 949
aaa authorization console, 948
aaa authorization exec, 948
aaa authorization network cts-list group ise-group, 562
aaa authorization network cts-list group radius, 567
aaa group server radius ise-group, 561
aaa group server tacacs+ ISE-TACACS, 947
aaa new-model, 261, 561, 567, 947
aaa server radius dynamic-author, 265, 563
AAA Identity Management Security, 6
AAA Servers tab
corporate WLAN, 293–294
guest WLAN, 289–290
ABSOLUTE_PATH file path option, 669
ACCEPT message, 9
acceptable use policy (AUP) page settings
hotspot guest portals, 354–356
sponsored guest portals, 386
Acceptable Use Policy in Stealth Mode setting (Posture General Settings), 691
access control entries (ACEs), 553
Access Control List option (TACACS profile), 933
access control lists. *See* ACLs (access control lists)
Access Control Server (ACS), 8, 630, 909, 910
access levels, Admin group roles, 127–132
access policy
for FMC (Firepower Management Center), 840–844
for WSA (Web Security Appliance), 855–856
Access-Accept message, 13
Access-Challenge message, 13
access-layer devices. *See* NADs (network access devices)
Access-Reject message, 13
Access-Request message, 13, 88
Account Expiration Notification settings, 346
accounting
configuration
RADIUS accounting servers, 278–279
RADIUS fallback, 279–280
device administration AAA with Cisco IOS, 951
messages
RADIUS, 13
TACACS+, 11
Accounting-Request message, 14
Accounting-Response message, 14
accounts
CTA STIX/TAXII API, 892–893
guest
contractors, 344–346
daily, 344–346
overview of, 341, 343
provisioning from sponsor portals, 389–394
social, 348
weekly, 347
sponsor groups, 381–382
ACEs (access control entries), 553
ACL-MDM-REDIRECT, 539–540
ACLs (access control lists)
ACEs (access control entries), 553
ACL-MDM-REDIRECT, 539–540
ACL-WEBAUTH-REDIRECT, 280–282, 314–315, 330, 492, 693
configuration for BYOD onboarding, 492–495
DNS ACLs, 494
NSP ACL, 493, 495
dACLs (downloadable access control lists), 236, 237, 553
configuration for pre-WebAuth authorization, 319–320
creating, 246–249

DNS (Domain Name System), 494, 993
ingress, 553–554
local, 268–269
named, 244
pACLs (port-based ACLs), 725
SGACLs (security group ACLs), 597–604
definition of, 998
Deny_All SGACL, 601–602
east-west deployment of, 598–599
egress policy, 597–598, 600–601
north-south deployment of, 598–599
Permit_HTTP_HTTPS SGACL, 601–602
Permit_ICMP SGACL, 602–603
Permit_Mgmt SGACL, 601
Permit_SRC_HTTP_HTTPS SGACL, 603–604
Permit_WEB_RDP SGACL contents, 598
syntax, 599–600
URL-redirect
configuration for Cisco switch, 314–315
configuration for WLC, 316
wireless authentication, 280–284
Google URLs for ACL Bypass, 282–283
Posture Agent Redirection ACL, 283–284
Web Authentication Redirection ACL, 280–282

ACL-WEBAUTH-REDIRECT access list, 280–282, 314–315, 330, 492, 693

ACS (Access Control Server), 8, 630, 909, 910

Active Directory. *See* AD (Active Directory)

active users, viewing in FMC (Firepower Management Center), 844–845

ActiveX, 85–86, 153, 322, 356

AD (Active Directory), 23, 24–25, 196–202
Active Directory probe, 422
CA (certificate authority), 469
definition of, 991
domains, 24
forests, 24
joining, 197–202
MDM support for, 101
objects, 24
prerequisites for, 196

Adaptive Network Control (ANC), 148–149, 156, 822–823, 864–866

Adaptive Security Appliances. *See* ASA (Adaptive Security Appliance)

Adaptive Security Device Manager. *See* ASDM (Adaptive Security Device Manager)

Add Dashlet(s) command, 134

Add Directory dialog, 838

Add New Dashboard option, 134

Add New Realm dialog, 838

Add New Rule command, 282, 283

Add-Remove URL command, 283

address ipv4 command, 264

Address Resolution Protocol (ARP), 646

addresses, MAC (Media Access Control). *See* MAC (Media Access Control) addresses

AD-Host-Exists attribute (Active Directory), 422

adi_cli session command, 844–845

AD-Join-Point attribute (Active Directory), 422

Admin Access tab (System), 160–161

Administration persona, 109, 737

Administration portal, 137–142
Admin group roles, 127–132
global search for endpoints, 139–140
Help menu, 138, 140–141
initial login, 125–126
ISE Setup Wizards, 141
Settings menu, 142
tabs, 137–139

Administration profile, 940–941

Administration screen, 142, 155–170
Device Portal Management tab, 166–169
Blacklist portal, 166
Certificate Provisioning portal, 166–167
Client Provisioning portal, 166–167, 650–651
Custom Portal Files portal, 168
Mobile Device Management portal, 168
My Devices portal, 168
Settings, 169
Feed Service tab, 169
Identity Management tab, 161–163
External Identity Sources tab, 162
Groups tab, 162
Identities tab, 161
Identity Source Sequences tab, 163
Settings tab, 163
Network Resources tab, 163–166
External MDM tab, 165
External RADIUS Servers tab, 165
Location Services tab, 166
NAC Managers tab, 165
Network Device Groups tab, 164–165
Network Device Profiles tab, 165
Network Devices tab, 163–164
RADIUS Server Sequences tab, 165
PassiveID Setup option, 138
pxGrid Services tab, 169
Search icon, 138
System Activities option, 139
System tab, 155–161
Admin Access tab, 160–161
Backup & Restore tab, 160
Certificates tab, 158
Deployment tab, 155
Licensing tab, 155–158
Logging tab, 159
Maintenance tab, 159
Settings tab, 161
Upgrade tab, 160
Threat Centric NAC tab, 170
Visibility Setup option, 138
Wireless Setup (BETA) option, 139
Administration tab (Guest Access work center), 340–341
AD-Operating-System attribute (Active Directory), 422
AD-OS-Version attribute (Active Directory), 422
AD-Service-Pack attribute (Active Directory), 422
Advance Filter, 771
Advanced Malware Protection. *See* AMP (Advanced Malware Protection) for Endpoints
Advanced Settings tab, Windows native supplicant, 57
Advanced tab
corporate WLAN, 294–295
guest WLAN, 289–290
agents, temporal, 999
Aggregation Services Router. *See* ASR (Aggregation Services Router)
Airwatch, 708
Alarms dashlet, 134
ALL role, 972
ALL_ACCOUNTS sponsor group, 381
All_User_ID_Stores, 472–474
alternative ID stores based on EAP type, 224–227
EapAuthentication equals EAP-TLS, 225
EapTunnel equals EAP-FAST, 226–227
EapTunnel equals PEAP, 226
AMP (Advanced Malware Protection) for Endpoints, 897–904
adapter configuration, 900–904
capabilities of, 897–898
incidents, 899–900, 991
indicators, 899
AMP Enabler profile, 642

ANC (Adaptive Network Control), 148–149, 156, 822–823, 864–866, 991
AND operator, 252–256
Android devices. *See also* BYOD (bring your own device) onboarding; MDM (mobile device management)
BYOD (bring your own device) onboarding, 526–530
device provisioning, 529–530
device registration, 526–528
NSP app download, 528–529
fingerprint technology, 27
mobile posture, 707–712
authorization conditions, 709–710
authorization rules, 710–712
supported device managers, 707–709
Anomalous Behaviour Detection, 406–408, 991
anti-malware, 100, 661–663, 681–682
anti-spyware, 100, 663, 681–682
anti-virus, 100, 663, 681–682, 753–756
Anycast high availability
IP SLAs (service-level agreements), 754–756
network architecture, 753–754
route redistribution, 755–756
AnyConnect Diagnostics and Reporting Tool. *See* DART (AnyConnect Diagnostics and Reporting Tool)
AnyConnect ISE agent, 992
AnyConnect ISE Posture Module, 99–100
AnyConnect NAM (Network Access Manager) supplicant, 59–73, 809
AnyConnect NAM profiles, 71–72
Authentication Policy view, 60, 62
Client Policy view, 60, 61–62
EAP chaining, 73
Network Groups view, 60, 71
Networks view, 60, 62–70
Certificates tab, 67
Connection Type tab, 66
Credentials tab, 68–70
Machine Auth tab, 66
PAC Files tab, 67–68
Security Level tab, 64–66
User Auth tab, 68–70
overview of, 59–60
profiles, 71–72
AnyConnect posture assessment endpoint scenarios
AnyConnect already installed, endpoint not compliant, 700–702
AnyConnect not installed on endpoint yet, 696–700
stealth mode, 645, 703
temporal agent and posture compliant, 705
AnyConnect Secure Mobility Client, 640–649
AnyConnect configuration, 648–649
AnyConnect posture profile configuration, 644–648
client provisioning resource configuration, 640–642
downloading, 640
headend deployment packages, uploading to ISE, 642–644
AnyConnect System Scan. *See* posture assessment
Apex license packages, 156
APIs (application programming interfaces)
license packages, 156
MDM (mobile device management), 821
Apple iOS. *See also* MDM (mobile device management)
BYOD (bring your own device), 523–526. *See also* BYOD (bring your own device) onboarding
device enrollment, 523–524
device provisioning, 526–527
device registration, 523–524
faceprint technology, 27
iPad, 482
iPCU (iPhone Configuration Utility), 812
mobile posture, 707–712
authorization conditions, 709–710

authorization rules, 710–712
supported device managers, 707–709
Push Notification, 101
Touch ID, 310
application conditions, 100, 655–660
application provisioning, 101
application remediations, 680
Application tab (Context Visibility screen), 143
architecture (ISE). *See* ISE (Identity Services Engine) architecture
ARP (Address Resolution Protocol), 646
ASA (Adaptive Security Appliance)
ASDM (Adaptive Security Device Manager), 592–593
configuration for TrustSec, 564–565
DAP (Dynamic Access Policy), 629
SGFW (security group firewall) on, 612
SXP (SGT Exchange Protocol) configuration on, 591–592
ASDM (Adaptive Security Device Manager), 592–593
Ask a Question option (Help menu), 141
ASR (Aggregation Services Router), 613
assertions, SAML (Security Assertion Markup Language), 35, 395, 998
assetConnectedLinks attribute (pxGrid), 423
assetCustomAttributes attribute (pxGrid), 423
assetDeviceType attribute (pxGrid), 423
assetHwRevision attribute (pxGrid), 423
assetId attribute (pxGrid), 423
assetIpAddress attribute (pxGrid), 423
assetMacAddress attribute (pxGrid), 423
assetName attribute (pxGrid), 423
assetProductId attribute (pxGrid), 423
assetProtocol attribute (pxGrid), 423
assetSerialNumber attribute (pxGrid), 423
assetSwRevision attribute (pxGrid), 423
assetVendor attribute (pxGrid), 423
assignment
SGTs (security group tags)
dynamically, 577
manually, 577–578
VLAN (virtual LAN), 551–553
asterisk (*), 494, 922
attribute/value (AV) pairs, 13, 15, 107
Audit reports, 150
AUP (acceptable use policy) page settings
hotspot guest portals, 354–356
sponsored guest portals, 386
authentication, 13–14. *See also* certificate-based authentication; CWA (Centralized Web Authentication); posture assessment
2FA (two-factor authentication), 26, 1000
authentication open versus 802.1X, 719–720
authentication servers, 41, 991
authenticators, 41, 991
authorization compared to, 209–210, 235
communications in, 42
definition of, 206
device administration AAA with Cisco IOS
device administration AAA with Cisco IOS, 946–948
TACACS+ command accounting, 951
TACACS+ command authorization, 948–950
devices without a supplicant, 79–80
EasyConnect, 89–90, 993
FlexAuth (Flexible Authentication)
configuration, 269–272
definition of, 994
FAST (Flexible Authentication via Secure Tunneling), 45, 48–49, 215
flows for, 804–805
importance of, 5
MAB (MAC Authentication Bypass)
authentication with, 80–82, 227–228
configuration, 265, 270–274, 318
definition of, 717, 995

MAB rule flowchart, 217
overview of, 96–99
profiling, 96–99
role-specific authorization rules, 241
wireless, 489
messages
RADIUS, 13–14
TACACS+, 9
MFA (multifactor authentication), 26–29
Multi-Auth (Multiauthentication), 995
policies and policy sets, 151
allowed protocols, 210, 213–216
for alternative ID stores based on EAP type, 224–227
authorization compared to, 209–210, 235
for certificate-based authentication, 472–474
conditions, 217–219
default, 216–217
definition of, 171
for device administration AAA with Cisco IOS, 944
goals of, 206–207, 210–211
identity stores, 210–211, 219–220
identity validation, 211
for MAB (MAC Authentication Bypass), 227–228
options, 220
for remote access VPN, 223–224
restoring, 229
for wireless SSID, 220–223
remote access connections, 88–89
SAML (Security Assertion Markup Language), 394–400
assertions, 395, 998
guest portal logins, 368, 394–400
IdPs (identity providers), 35, 394–400, 998
SPs (service providers), 394, 998
support for, 23
sponsored guest portals, 342, 380–381
AUP (acceptable use policy) page settings, 386
configuration flowchart, 380–381
default sponsor portal, 384
login settings, 386
other settings, 387
portal settings, 385–386
provisioning guest accounts from, 389–394
timers, 275
Web Authentication. *See* WebAuth (Web Authentication)
Windows native supplicant
machine authentication, 58–59
user authentication, 58
wired, 261–276
Cisco ISE verification, 302–303
endpoint supplicant verification, 295–296
global configuration AAA commands, 261–262
global configuration RADIUS commands, 262–269
interface configuration settings, 269–276
NAC (network access device) verification, 296–302
wireless
AAA servers, 276–280
airespace ACLs (access control lists), 280–284
Cisco ISE verification, 302–303
dynamic interfaces for client VLANs, 284–286
endpoint supplicant verification, 295–296
NAC (network access device) verification, 296–302
wireless LANs, 286–295
authentication event fail action command, 271

authentication event server alive action reinitialize command, 272
authentication event server dead action authorize vlan command, 272
authentication event server dead action authorize voice command, 272
authentication event server dead action reinitialize vlan command, 272
authentication host-mode multi-auth command, 273
authentication linksec policy command, 617
authentication open command, 275, 718, 719–720, 722, 725
authentication order dot1x mab command, 271
Authentication Policy view, Cisco AnyConnect NAM supplicant, 60, 62
authentication port-control auto command, 276
authentication priority dot1x mab command, 271
authentication servers, 41, 991
authentication success settings, hotspot guest portals, 357–358
Authentication tab, Windows native supplicant, 53, 56–57
authentication timers, 275
authentication violation restrict command, 273
authentication VLAN, 87–88
Authentications dashlet, 134
authenticators, 41, 991
authorization. *See also* ACLs (access control lists)
authentication compared to, 209–210, 235
for certificate-based authentication, 474–475
Cisco CTA (Cognitive Threat Analytics), 896–897
CoA (Change of Authorization)
CWA (Centralized Web Authentication) and, 85–86, 311
definition of, 16, 95–96, 992
enabling, 265, 277
ISE Profiler and, 442–444
messages, 110, 748
definition of, 209, 232
for device administration AAA with Cisco IOS, 948–950
flows for, 804–805
messages
RADIUS, 13
TACACS+, 10–11
mobile posture
authorization conditions, 709–710
authorization rules, 710–712
policies and policy sets, 151
authentication compared to, 209–210, 235
Blackhole_Wireless_Access, 240–241
for certificate-based authentication, 474–475
Cisco_IP_Phones, 237–241
compound conditions, 239, 251–256, 992
condition blocks, 252–256
configuration of, 241–249
default, 236–241
definition of, 171
for device administration AAA with Cisco IOS, 945–946
goals of, 235–241
for guest portals, 348–351
for MDM (mobile device management), 536–537
organization of, 216, 236
profile assignment in, 450–453
role-specific authorization rules, 241
rule processing for, 236–241
saving conditions for reuse in, 249–251
simple conditions, 239, 251, 999
profiles, 450–453
assignment of, 450–452, 453

for BYOD (bring your own device) onboarding, 516
configuration, 320–322
Downlink MACsec, 616
Employee Full Access, 241–243
Employee_Limited, 246–249
for hotspot guest portals, 362–364
Internet_Only, 243–246
MDM Onboard, 539–540
for posture assessment, 693
for self-registered guest portals, 373–380
for TC-NAC (Threat Centric Network Access Control), 884–886
rules, 236–241, 693–694
AND/OR operators in, 252–256
for BYOD (bring your own device) onboarding, 517, 518
for device administration AAA with Cisco WLC (Wireless LAN Controller), 977–979
Employee and CorpMachine, 242–243
employee full access, 241–243
employee limited access, 246–249
Internet-only access, 243–246
IT Users Access, 252–256
for MDM onboarding, 539–542
PERMIT_ALL_IPV4_TRAFFIC, 241–243
role-specific, 241
for self-registered guest portals, 373–380
for TC-NAC (Threat Centric Network Access Control), 884–886
Wireless Black List Default, 239
security context, 232, 235
TrustSec, 559
Authorization Policy column, Live Log, 767
Authorization Profiles column, Live Log, 768
Authy, 29
Auto Command option (TACACS profile), 933
auto PAN switchover, 745–746
automate-tester username command, 264
automatic failover, 746
auto-source command, 267
AV (attribute/value) pairs, 13, 15, 107

B

Backup & Restore tab (System), 160
backup and restore, 101, 160, 759–761
backup interface GigabitEthernet 3 command, 352
backup-logs command, 783–784
Base license packages, 156
Base64-encoded files, 477
BlackBerry, 508–509, 708
Blackhole_Wireless_Access authorization profile, 240–241
Blacklist portal, 166
blocks, condition, 252–256
bootstrapping ISE (Identity Services Engine), 177–180
bring your own device. *See* BYOD (bring your own device) onboarding
browsers
requirements for, 125
Softerra LDAP, 26
support for, 122–123
BYOD (bring your own device) onboarding. *See also* MDM (mobile device management)
Android onboarding flow, 526–530
device provisioning, 529–530
device registration, 526–528
NSP app download, 528–529
challenges of, 485–487
definition of, 487
dual SSID versus single SSID, 487–488, 993, 999
end-user experience

Blackberry example, 508–509
dual SSID with Android example, 503–508
single SSID with Apple iOS example, 496–503
history of, 482
iOS onboarding flow, 523–526
device enrollment, 523–524
device provisioning, 526–527
device registration, 523–524
ISE configuration for, 495–496
authorization profiles, 516
authorization rules for EAP-TLS authentications, 518
authorization rules for onboarding, 517
client provisioning policy configuration, 512–514
default unavailable client provisioning policy action, 515
ISE as certificate authority, 519–520, 521–523, 994
native supplicant profile, 510–512
SCEP (Simple Certificate Enrollment Protocol), 520–521, 999
WebAuth configuration, 514–515
overview of, 487
self-registered guest portal settings, 372
verification of BYOD flows, 534–535
endpoint identity groups database, 535
RADIUS Live Logs, 534
reports, 534–535
Windows and Mac onboarding flow, 531–533
device provisioning, 532–533
device registration, 531
WLC (Wireless LAN Controller) configuration, 489–495
ACLs (access control lists), 492–495
WLAN configuration, 490–491
BYOD Endpoints dashlet, 134

C

C (Country) field, 184
C3PL (Common Classification Policy Language), 789
CA_SERVICE_Certificate_Template, 520
Cache Last Known Posture Compliance setting, 691
CACs (Common Access Cards), 992
Call Home List setting (AnyConnect posture profile), 647
Call-Check for Service-Type, 228
Called-Station-Id attribute (RADIUS), 414
Calling-Station-Id attribute (RADIUS), 414
Calling-Station-Id field (RADIUS), 97
CAPs (certificate authentication profiles), 23, 202, 469, 471–472, 991
cards, smart, 29
career limiting events (CLEs), 717
CAs (certificate authorities). ***See also*** **certificate-based authentication**
AD (Active Directory), 469
CA-signed certificates, 182–191
characteristics of, 30–33
CSR (certificate signing requests), 182–191
binding certificates, 189–191
certificate subject fields, 183–184
downloading and importing certificates, 188–191
exporting certificates, 191
PEM files, 186–187
submitting, 186–187
wildcard certificates, 184
definition of, 992
ISE as, 519–520, 521–523, 994
signatures, 31–32
trusted, 31–32, 475–479
certificate status verification, 478–479
public certificates, 476–477
role in authentication process, 463–465

Catalyst 3000 Series, 12.2(55)SE switch configuration, 1034–1038
Catalyst 3000 Series, 15.0(2)SE switchconfiguration, 1038–1044
Catalyst 4500 Series, IOS-XE 3.3.0 / 15.1(1)SG switchconfiguration, 1053–1057
Catalyst 6500 Series, 12.2(33)SXJ switchconfiguration, 1058–1061
Catalyst 9000 Series, 16.9.5 switchconfiguration, 1044–1052
categories, logging, 778–779
CCP (client provisioning policy), 515
CDA (Context Directory Agent), 469
CDP (Cisco Discovery Protocol), 98, 273, 418, 991
centralized portal, LWA (Local Web Authentication) with, 84–85
Centralized Web Authentication. *See* CWA (Centralized Web Authentication)
certificate authentication profiles. *See* CAPs (certificate authentication profiles)
certificate authorities. *See* CAs (certificate authorities)
Certificate Authority Certificates menu (ISE Certificate Authority), 520
Certificate Provisioning portal, 166–167
certificate revocation lists (CRLs), 33, 466, 992
certificate signing requests. *See* CSRs (certificate signing requests)
Certificate Templates menu (ISE Certificate Authority), 520
certificate-based authentication, 158. *See also* PKI (public key infrastructure)
CAPs (certificate authentication profiles), 23, 202, 469, 471–472, 991
CAs. *See* CAs (certificate authorities)
concept of, 30–31
CRLs (certificate revocation lists), 466, 992
CSRs (certificate signing requests), 182–191, 521–523
binding certificates, 189–191
definition of, 992
downloading and importing certificates from, 188–191
exporting certificates, 191
ISE (Identity Services Engine), 521–523
PEM files, 186–187
submitting, 186–187
wildcard certificates, 184
CWA (Centralized Web Authentication) configuration, 313
definition of, 463
EAP-TLS (EAP Transport Layer Security), 470
expired certificates, 32–33, 465
guest services, 340–341
ISE configuration for, 181–191, 470–479
authentication policies, 472–474
authorization policies, 474–475
CAPs (certificate authentication profiles), 471–472
certificate status verification, 478–479
overview of, 470
principal username X.509 attribute, 470
protocols verification, 470–471
public certificates, importing, 476–477
trusted CAs (certificate authorities), 475–479
popularity of, 460
public certificates, 476–477
purpose of, 181
pxGrid certificates. *See* pxGrid (Platform Exchange Grid)
RAs (registration authorities), 998
revoked certificates, 33
checking for, 466–467
CRLS (certificate revocation lists), 33
CRLs (certificate revocation lists), 33, 466

OCSP (Online Certificate Status Protocol), 33, 466, 996
validity period, 467
self-signed certificates, 181–182
trusted certificates, 537–538
Certificates tab
Cisco AnyConnect NAM supplicant, 67
System, 158
chaining, EAP, 73, 216
Challenge Handshake Authentication Protocol. *See* CHAP (Challenge Handshake Authentication Protocol)
Change of Authorization. *See* CoA (Change of Authorization)
CHAP (Challenge Handshake Authentication Protocol), 7, 46, 214, 215
chip cards, 29
Chrome, support for, 122
Cisco Access Control Server (ACS), 909, 910
Cisco Access Registrar, 8
Cisco AnyConnect Diagnostics and Reporting Tool. *See* DART (AnyConnect Diagnostics and Reporting Tool)
Cisco AnyConnect ISE agent, 992
Cisco AnyConnect ISE Posture Module. *See* AnyConnect ISE Posture Module
Cisco AnyConnect Network Access Manager. *See* AnyConnect NAM (Network Access Manager) supplicant
Cisco Cognitive Threat Analytics. *See* CTA (Cognitive Threat Analytics)
Cisco Context Directory Agent (CDA), 469
Cisco Discovery Protocol (CDP), 98, 418
Cisco Duo Security, 27–29
Cisco Firepower Management Center. *See* FMC (Firepower Management Center) configuration
Cisco Identity Services Engine architecture. *See* ISE (Identity Services Engine) architecture
Cisco Industrial Network Director (IND), 827
Cisco IOS. *See* IOS (Internetwork Operating System)
Cisco ISE for BYOD and Secure Unified Access (Woland and Heary), 859
Cisco ISE (Identity Services Engine) architecture. *See* ISE (Identity Services Engine) architecture
Cisco Meraki Systems Manager (Meraki SM), 708
Cisco Meta Data (CMD) field, 593, 616
Cisco NAC (Network Admission Control), 626, 630
Cisco Network Setup Assistant app, 492, 507
Cisco Platform Exchange Grid. *See* pxGrid (Platform Exchange Grid)
Cisco Secure Access Control Server (ACS), 8
Cisco Secure Network Server (SNS), 177
Cisco Security Agent (CSA) service, 676
Cisco Software-Defined Access (SD-Access), 613–614
Cisco Stealthwatch. *See* Stealthwatch
Cisco Supplicant Provisioning Wizard, 513
Cisco Temporal Agent, 99–100
Cisco Umbrella, 640
Cisco Wireless LAN Controller. *See* WLC (Wireless LAN Controller)
Cisco_IP_Phones authorization profile, 237–241
cisco-av-pair command, 576
CiscoPress SSID policy set, 518
CiscoPress-TLS, 513–514
Citrix XenMobile, 708
Class (RADIUS attribute 25) VSA, 266
Class-Identifier attribute (DHCP), 411
Clean Machines, 631
CLEs (career limiting events), 717
CLI (command-line interface), 992. *See also specific commands*
CLI Setup Wizard, 178–179
Client Messaging servers, 101

Client Policy view, Cisco AnyConnect NAM supplicant, 60, 61–62
client provisioning policy. *See* CPP (client provisioning policy)
Client Provisioning portal, 166–167, 650–651
Client Provisioning tab (Policy page), 153
Client Stopped Responding counter, 768
Client Supplicant Not Capable of MACsec policy, 617
client VLANs, dynamic interfaces for, 284–286
employee dynamic interface, 284–285
guest dynamic interface, 285–286
Client-FQDN attribute (DHCP), 411
client/server communication, 7–8
closed mode, 728–730
CMD (Cisco Meta Data) field, 593, 616
CN (common name), 183, 833
CoA (Change of Authorization)
CWA (Centralized Web Authentication) and, 85–86, 311
definition of, 16, 992
enabling, 265, 277
ISE Profiler and, 442–444
global CoA, 442–443, 994
per-profile CoA, 443–444
messages, 110, 748
CoA Type setting, hotspot guest portals, 354
Coa-Events report, 888–889
Cognitive Intelligence. *See* CTA (Cognitive Threat Analytics)
Cognitive Threat Analytics. *See* CTA (Cognitive Threat Analytics)
command sets
definition of, 992
device administration AAA with Cisco IOS, 934–936
TACACS+, 934–936, 941–943
command-line interface (CLI), 992. *See also specific commands*
COMMANDS role, 971
comma-separated values (CSV) files, 194
Common Access Cards (CACs), 992
Common Classification Policy Language (C3PL), 789
common name (CN), 183, 833
Common Port option, Nmap, 416
Common Task Type option, TACACS profile, 933
Common Vulnerabilities and Exposures (CVE), 873, 992
Common Vulnerability Scoring System (CVSS), 873, 992
compliance. *See* posture assessment
compliance modules, updating, 637–638
compound conditions, 239, 251–256, 664–665, 677–678, 992
Compromised Endpoints Over Time dashlet, 137
conditions
authentication policy, 217–219
authorization policy
blocks, 252–256
compound, 239, 251–256, 664–665, 677–678, 992
definition of, 235
EmployeeFullEAPChain, 249–250
mobile posture, 709–710
saving for reuse, 249–251
simple, 239, 251, 999
posture policy, 654–679
anti-malware, 661–663
anti-spyware, 663
anti-virus, 663
application, 655–660
compound, 677–678
dictionary compound, 664–665
dictionary simple, 663–664
disk encryption, 665–666
file, 667–673
firewall, 660–661
hardware attributes, 655

patch management, 673–675
Registry, 675
USB, 679
Wired_802.1X, 242, 254
Wired_MAB, 242
Conditions Studio, 217–219
Conditions tab (Policy Elements), 154
conf t command, 621
configure command, 922
configure terminal command, 963
Connection Settings (Device Administration), 918
Connection Type tab (Cisco AnyConnect NAM supplicant), 66
Console application, 812
context
brokering, 992
definition of, 992
sharing
EPS (Endpoint Protection Services), 821–822
MDM integration, 820–821
need for, 818
pxGrid. See pxGrid (Platform Exchange Grid)
Rapid Threat Containment, 821–823, 993, 998
Context Directory Agent (CDA), 469
Context Visibility screen, 137, 142, 143
Context-In, 827, 992
Context-Out, 827, 993
CONTINUE message, 9, 11
Continue option
authentication policy, 220
posture reassessment, 692
Continuous Monitoring Interval setting (Posture General Settings), 691
contractors, 344–346
CONTROLLER role, 971
Core Files support bundles, 783
Corporate Wipe option, 543
corporate WLAN configuration, 291–295
AAA Servers tab, 293–294
Advanced tab, 294–295
General tab, 292
Layer 2 Security tab, 292–293
Layer 3 Security tab, 293
correlation policy, 845–847
correlation rules, 845–847
Country (C) field, 184
CPP (client provisioning policy), 172, 637–638
AnyConnect Secure Mobility Client, 640–649
AnyConnect configuration, building, 648–649
AnyConnect posture profile, 644–648
client provisioning resource configuration, 640–642
headend deployment packages, uploading to ISE, 642–644
BYOD (bring your own device) onboarding
client provisioning policies, 512–514
default unavailable client provisioning policy action, 515
Client Provisioning portal, 153, 166–167, 650–651
default client provisioning policy, 652
order of operations, 637–638
rules, creating, 652–653
CRC32 file type, 672
credentials, 20
Credentials tab (Cisco AnyConnect NAM supplicant), 68
CRLs (certificate revocation lists), 33, 466, 992
Cropped Portal Page Customization screen, 358
crypto key generate rsa general-keys mod 2048 command, 313
crypto key generate rsa modulus 2048 command, 946

CSA (Cisco Security Agent) service, 676
CSRs (certificate signing requests), 182–191, 521–523
binding certificates, 189–191
definition of, 992
downloading and importing certificates from, 188–191
exporting certificates, 191
ISE (Identity Services Engine), 521–523
PEM files, 186–187
submitting, 186–187
wildcard certificates, 184
CSV (comma-separated values) files, 194
CTA (Cognitive Threat Analytics), 890–897
authorization with, 896–897
CTA STIX/TAXII API account creation, 892–893
dashboard, 890–892
integration for TC-NAC, 894–896
CTD (Cyber Threat Defense), 857
CTI (cyber threat intelligence), 890, 993
cts authorization list cts-list command, 562
cts credentials id Sw01 password Cisco123 command, 561
cts manual command, 595, 621
cts manualN7K-DIST command, 578
cts role-based enforcement command, 562, 595, 596
cts role-based sgt-map command, 580
cts role-based sgt-map vlan-list command, 580
cts sxp connection peer command, 588, 589
cts sxp default password command, 588
cts sxp default source-ip command, 588
cts sxp enable command, 588–589
cubes, ISE
definition of, 109, 737, 995
joining, 182
licensing in, 747–748
Current, Chip, 337
Custom Attribute option (TACACS profile), 933
Custom Portal Files portal, 168
Custom Ports option (Nmap), 416
custom profiling attributes, 445–448
Customization admin role, 127
CVE (Common Vulnerabilities and Exposures), 873, 992
CVSS (Common Vulnerability Scoring System), 873, 992
CWA (Centralized Web Authentication), 730–731. ***See also*** **sponsored guest portals**
authentication process, 85–87
authorization policies, 322–324
custom authorization rules, 323–324
Guest Flow attribute, 323–324, 994
preconfigured authorization rules, 322
Cisco switch configuration, 313–315
certificates, 313
HTTP/HTTPS server, 314
URL-redirect ACL, 314–315
CoA (Change of Authorization) and, 311
definition of, 991
dual SSID onboarding and, 496
ISE (Identity Services Engine) configuration, 317–322
authorization profiles, 320–322
dACLs (downloadable access control lists), 319–320
Guest_Portal_Sequence ISS (identity source sequence), 319
MAB (MAC Authentication Bypass), 96–99, 318
services supported by, 311
with third-party network device support, 87–88
URL-redirected MAC authentication bypass, 311–313
verification from client, 324–326
verification on NAD (network access device), 327–331

client details, viewing on WLC, 329–331
show commands on wired switch, 328–329
WLC (Wireless LAN Controller) configuration, 98, 315–316, 329–331
ISE NAC feature, 315–316
MAC Filtering option, 315
URL-redirect ACL, 316
Cyber Threat Defense (CTD), 857
cyber threat intelligence (CTI), 890, 993

D

dACLs (downloadable access control lists), 13, 236, 237, 246–249, 319–320, 548, 553
daily guest accounts, 344–346
DAP (Dynamic Access Policy), 629
DART (AnyConnect Diagnostics and Reporting Tool), 59, 809–811, 991
dashboard, 132–137
Dashboard Settings menu, 134
Endpoints tab, 134–135
Guests tab, 135–136
Summary tab, 134
Threat tab, 137
verifying profiling with, 454
Vulnerability tab, 136
Dashboard Settings menu, 134
dashlets
Alarms, 134
Authentications, 134
BYOD Endpoints, 134
Compromised Endpoints Over Time, 137
Endpoint Categories, 135, 454
Endpoint Categories dashlet, 135
Endpoints, 134, 135
Failure Reason, 136
Guest Status, 136
Guest Type, 136
Location, 136
Metrics, 134
Network Devices, 134, 135
Status, 134
System Summary, 134
Threat Watchlist, 137
Top Threats, 137
Top Vulnerability, 136
Total Compromised Endpoints, 137
Total Vulnerable Endpoints, 136
Vulnerability Watchlist, 136
Vulnerable Endpoints Over Time, 136
databases, 994
endpoint identity groups, 535
endpoints, 455–456
internal endpoint, 22, 994
internal user, 994
Datagram Transport Layer Security (DTLS), 190
DaysSinceLastCheckIn attribute, 537
debug aaa authentication command, 955
debug aaa authorization command, 955–958
debug aaa tacacs enable command, 985
debug authentication command, 815
debug client command, 302, 815
debug commands, 815
debug dot1x command, 302, 815
debug epm command, 815
debug interface command, 815
debug ip tcp transactions command, 966
debug logs, 779–784
configuration, 779–780
downloading from GUI, 780
support bundles, 782–784
categories of, 782–783
creating from CLI, 783–784
creating from GUI, 783
definition of, 782
viewing from CLI, 781–782
Debug Logs support bundles, 783

debug tacacs command, 958–961
debug tacacs packet command, 963–965
decryption policy, 857
Dedicated MNT, 742
de-duplication of logs, 805–807
default authorization policies, 236–241
default client provisioning policy, 652
default policy sets, 211
Default Posture Status setting (Posture General Settings), 691
Default Privilege option (TACACS profile), 933
default sponsor portals, 384
defense-in-depth, 241
Delete option (endpoint management), 543
deny statement, 280, 316
Deny_All SGACL, 601–602
DenyAllCommands command, 941
deployment
AnyConnect headends
AnyConnect configuration, building, 648–649
package upload to ISE, 642–644
posture profile configuration, 644–648
device administration AAA, 911–913
large deployments, 912
medium deployments, 913
small deployments, 913
high availability, 743
Anycast high availability, 753–756
backup and restore, 759–761
failure scenarios, 753
general guidelines for, 752–753
licensing in multi-mode ISE cube, 747–748
load balancing, 751–752, 756–757
MnT (Monitoring and Troubleshooting) nodes, 109–110, 743–744
node groups, 748–750
PAN (Policy Administration node), 109, 743–744
patches, 757–759
ISE (Identity Services Engine)
distributed, 116–119
hybrid, 116–117
overview of, 113
physical versus virtual, 111–113
scale limits for, 118–119
single-node, 113
two-node, 114–116
ISE nodes in distributed environment, 737–742
ISE cubes, 737
ISE persona types, 737
node personas, 742
PPAN (Policy Administration Node), 738–739
primary devices, 738–739
registration of ISE nodes, 739–742
phased approach to
advantages of, 717–718
authentication open versus 802.1X, 719–720
closed mode, 728–730
low-impact mode, 725–727
monitor mode, 722–725, 730–731
preparation for, 720–721
transitioning to end state, 730–731
wireless networks, 731
Deployment tab (Administration), 155
design
for device administration, 911–913
large deployments, 912
medium deployments, 913
small deployments, 913
SXP (SGT Exchange Protocol), 582–583
Desktop Preview (portal page customization), 358
destination tree view (TrustSec policy matrix), 553

Details icon, Live Log, 768
device administration AAA, 478–479
concept of, 6
configuring with Cisco IOS, 930
device administration service, enabling, 937
network device preparation, 937–939
overview of, 932
policy preparation, 939–946
privilege levels, 932–933, 997
TACACS profiles, 932–934
TACACS+ authentication and fallback, 946–948
TACACS+ command accounting, 951
TACACS+ command authorization, 948–950
TACACS+ command sets, 934–936, 992
testing and troubleshooting in ISE, 952–954
troubleshooting at IOS command line, 954–966
configuring with Cisco WLC (Wireless LAN Controller)
ISE configuration on WLC TACACS+ servers, 979–980
network device preparation, 972
policy results preparation, 974–977
policy sets, 977–979
testing and troubleshooting, 981–986
top-level menus, 971–972
definition of, 2, 993
design and deployment, 911–913
large deployments, 912
medium deployments, 913
small deployments, 913
device administration service, enabling, 937
device backup, 101
device enrollment, 523–524
device tracking in IOS Xe 16.x and later, 267
devices without a supplicant, 79–80
graphical illustration of, 6, 909
interactive nature of, 909–910
license packages, 157
MDM (mobile device management), 108, 820–821. *See also* BYOD (bring your own device) onboarding
definition of, 995
overview of, 101–102
posture assessment with, 108
supported features, 101–102
vendors, 101–102
NAD (network access device) configuration, 917
connection settings, 918
identities, 920
network resources, 921–922
overview of, 916–917
password change control, 918
policy elements, 922–924
policy sets, 925–927
reports, 927
session key assignment, 918
UI navigation for, 919–920
network device troubleshooting, 812–815
client details, viewing on WLC, 813–814
debug commands, 815
show authentication interface command, 812–813
policies, 922–924, 939–946
policy sets, 943–946
roles, 939
TACACS command sets, 922–923, 941–943
TACACS profiles, 923–924, 939–941
purpose of, 909–910
RADIUS. *See* RADIUS (Remote Access Dial-In User Service)

reports, 150
resources for, 6
security context, 232, 235
smart devices
employee limited access, 246–249
Internet-only access for, 243–246
TACACS+. *See* TACACS+ (Terminal Access Controller Access Control System)
Device Administration Policy Set, 943–944
Device Portal Management tab (Administration screen), 166–169
Blacklist portal, 166
Certificate Provisioning portal, 166–167
Client Provisioning portal, 166–167, 650–651
Custom Portal Files portal, 168
Mobile Device Management portal, 168
My Devices portal, 168
Settings, 169
device provisioning
Android onboarding flow, 529–530
iOS onboarding flow, 526–527
Windows and Mac onboarding flow, 532–533
device registration
Android onboarding flow, 526–528
iOS onboarding flow, 523–524
Windows and Mac onboarding flow, 531
Device Sensor, 98, 267, 426–427, 457–458
DeviceComplianceStatus attribute, 536
DeviceCompliantStatus attribute, 712
DeviceRegisterStatus attribute, 536
device-sensor accounting command, 427
device-sensor filter-list cdp list command, 427
device-sensor filter-list dhcp list command, 426
device-sensor filter-list lldp list command, 427
device-sensor filter-spec cdp include list command, 427
device-sensor filter-spec dhcp include list command, 427
device-sensor filter-spec lldp include list command, 427
device-sensor notify all-changes command, 427
device-tracking attach-policy command, 267
device-tracking policy command, 267
device-tracking tracking auto-source command, 267
DHCP (Dynamic Host Configuration Protocol)
DHCP helper, 424–427
DHCP probe, 411–414
DHCPSPAN probe, 411–414
profiling, 97, 98
diagnostic tools
Endpoint Debug, 796–798
endpoint diagnostics, 809–812
Cisco AnyConnect Diagnostics and Reporting Tool (DART), 59, 809–811
supplicant provisioning logs, 812
Evaluate Configuration Validator, 788–793
Execute Network Device Command, 787–789
network device troubleshooting, 812–815
client details, viewing on WLC, 813–814
debug commands, 815
show authentication interface command, 812–813
Posture Troubleshooting, 794–795
RADIUS Authentication Troubleshooting tool, 785–786
Session Trace, 801–804
TCP Dump, 798–801
troubleshooting methodology, 804–808
authentication and authorization flows, 804–805
log de-duplication, 805–807
USERNAME user, 807

Diagnostics and Reporting Tool. *See* DART (AnyConnect Diagnostics and Reporting Tool)
Diagnostics reports, 150
diagrams, monitor mode flow, 7, 723–724
dial-up networking, 88
Dictionaries tab (Policy Elements), 154
dictionary compound conditions, 664–665
dictionary simple conditions, 100, 663–664
digital certificates. *See* certificate-based authentication
disable command, 932
Disable UAC Prompt setting (AnyConnect posture profile), 646
Discovery host setting (AnyConnect posture profile), 647
disk encryption conditions, posture policy, 665–666
DiskEncryptionStatus attribute, 536
Display Language setting, hotspot guest portals, 354
distributed ISE deployment, node configuration in, 116–119, 737–742
ISE cubes, 737
ISE persona types, 737
node personas, 742
PPAN (Policy Administration Node), 738–739
primary devices, 738–739
registration of ISE nodes, 739–742
Distribution Points, CRL, 467
DNS (Domain Name System), 196
ACLs (access control lists)
for BYOD onboarding, 494
definition of, 993
DNS probe, 417
snooping, 494
Domain Name System (DNS), 196, 417
Domain-Name attribute (DHCP), 411
domains
AD (Active Directory), 24
joining, 197–202
prerequisites for joining, 196
TrustSec, 557, 1000
Dot1x. *See* 802.1X
dot1x pae authenticator command, 275
dot1x system-auth-control command, 266, 563
dot1x timeout tx-period 10 command, 275
Downlink MACsec, 616–619
authorization profile, 616
ISE (Identity Services Engine) configuration, 619
policies, 616–618
switch configuration modes, 618–619
downloadable access control lists (dACLs), 13, 236, 237, 246–249, 319–320, 553
downloadable ACLs (dACLs), 548
downloading
AnyConnect Secure Mobility Client, 640
debug logs, 780
Drop option, authentication policy, 220
DTLS (Datagram Transport Layer Security), 190
dual SSID onboarding, 487–488, 993
definition of, 993
ISE (Identity Services Engine) configuration, 495–496, 510–523
authorization profiles, 516
authorization rules for EAP-TLS authentications, 518
authorization rules for onboarding, 517
client provisioning policy configuration, 512–514
default unavailable client provisioning policy action, 515
dual SSID with Android example, 503–508
ISE as certificate authority, 519–520, 521–523, 994
native supplicant profile, 510–512
SCEP (Simple Certificate Enrollment Protocol), 520–521, 999
WebAuth configuration, 514–515

verification of BYOD flows, 534–535
endpoint identity groups database, 535
RADIUS Live Logs, 534
reports, 534–535
WLC (Wireless LAN Controller) configuration, 489–495
ACLs (access control lists), 492–495
WLAN configuration, 490–491
Dubois, Jesse, 772
"dumb" devices, 96
Duo Security, 27–29
Dynamic Access Policy (DAP), 629
Dynamic Host Configuration Protocol. *See* DHCP (Dynamic Host Configuration Protocol)
dynamically assigning SGTs (security group tags), 577

E

EAP (Extensible Authentication Protocol), 7, 42–49, 73, 214, 545–546
alternative ID stores based on, 224–227
comparison of, 47–49
definition of, 993
EAP-FAST, 45, 48–49, 215
EAP-GTC, 43, 45, 215
EAP-MD5, 43, 46, 214
EAP-MS-CHAPv2, 43, 45, 46, 215
EAPoL (EAP over LAN). *See* 802.1X
EAPoL-Proxy-Logoff, 273, 993
EAPoL-Start (EAP over LAN Start), 270
EAP-TLS, 43, 45, 214, 215
EAP-TTLS, 45–46, 48–49
identity store comparison, 49
native types, 43–44, 996
PEAP, 44–45, 48–49, 53–54, 55–56, 108, 215
TEAP, 46–47, 48–49, 73
tunneled, 44–49, 1000
EAP MSCHAPv2 Properties dialog, 54
EAP_Authentication_Certificate_Template, 520
EAPChainingResult, 546
EAP-Identity-Request messages, 270
east-west segmentation, 554–555
east-west SGACL deployment, 598–599
EasyConnect, 89–90, 993
Edge, support for, 122
editors, Conditions Studio, 218
EDR (endpoint detection and response), 661, 897–898
eduroam initiative, 46
egress enforcement, of SGTs (security group tags), 555–556
EKU (extended key usage), 211
Elevated system admin role, 131
Employee and CorpMachine authorization rule, 242–243
employee dynamic interface, 284–285
employee full access rule, 241–243
employee limited access, 246–249
Employee profile, 977
Employee_Limited dACL, 246–249
EmployeeFullEAPChain condition, 249–250
Enable agent IP refresh setting (AnyConnect posture profile), 646
enable command, 932
Enable notifications in stealth mode setting (AnyConnect posture profile), 645
Enable Rescan Button setting (AnyConnect posture profile), 645
enable secret ISEc0ld command, 947
Enable signature check setting (AnyConnect posture profile), 645
encryption, posture policy conditions for, 665–666
end state, transitioning to, 730–731
Endpoint Assignment tab (Adaptive Network Control), 149
endpoint attribute filtering, 444–445
Endpoint Categories dashlet, 135, 454

Endpoint Debug tool, 796–798
Endpoint ID column (Live Log), 767
endpoint identity groups, 22, 354, 450–452
endpoint identity groups database, 535
Endpoint list, 454
Endpoint Profile column (Live Log), 767
Endpoint Profile Policies, 431–441
endpoint protection platform (EPP), 661, 897–898
Endpoint Protection Services (EPS), 821–822, 993
EndPointPolicy, 453
endpoints
diagnostics, 809–812
Cisco AnyConnect Diagnostics and Reporting Tool (DART), 809–811
supplicant provisioning logs, 812
EDR (endpoint detection and response), 661, 897–898
endpoint attribute filtering, 444–445
endpoint diagnostics, 809–812
Cisco AnyConnect Diagnostics and Reporting Tool (DART), 59, 809–811
supplicant provisioning logs, 812
endpoint identity groups, 22, 354, 450–452
endpoint identity groups database, 535
EndPointPolicy, 453
endpoints database, 455–456
EPP (endpoint protection platform), 661, 897–898
EPS (Endpoint Protection Services), 993
global search for, 139–140
guest access, 338
health of. *See* posture assessment
internal endpoint database, 22, 994
local endpoint groups, 195
MDM (mobile device management), 542–543
onboarding, 153
posture assessment, 695–705
AnyConnect already installed, endpoint not compliant, 700–702
AnyConnect not installed on endpoint yet, 696–700
redirected state, 695–696
stealth mode, 645, 703
temporal agent and posture compliant, 705
profile policies for, 431–441
purge policies for, 345
supplicant verification, 295–296
Endpoints and Users reports, 150
Endpoints dashlet, 134, 135
Endpoints tab
Context Visibility screen, 143
Home page, 134–135
Endpoints widget, 454
enforcement, traffic
SGACLs (security group ACLs), 597–604, 998
Deny_All SGACL, 601–602
east-west deployment of, 598–599
egress policy, 597–598, 600–601
north-south deployment of, 598–599
Permit_HTTP_HTTPS SGACL, 601–602
Permit_ICMP SGACL, 602–603
Permit_Mgmt SGACL, 601
Permit_SRC_HTTP_HTTPS SGACL, 603–604
Permit_WEB_RDP SGACL contents, 598
syntax, 599–600
SGFWs (security group firewalls), 611–613
on ASA (Adaptive Security Appliances), 612
on ASR (Aggregation Services Router), 613
definition of, 999
on Firepower, 612–613
on ISR (Integrated Services Router), 613

TrustSec policy matrix, 604–609
configuration of, 605–609
views of, 604–605
environment data, 558, 993
epm logging command, 299
EPP (endpoint protection platform), 661, 897–898
EPS (Endpoint Protection Services), 821–822, 993
EPSStatus condition, 821–822
eradication, 896
ERROR message, 9, 11
ERS (External RESTful Services), 131, 850–851
Evaluate Configuration Validator, 788–793
Evaluation license, 155, 157
exam preparation, 988–989
exam updates, 1032–1033
final study and review, 988–989
hands-on activities, 988–989
Execute Network Device Command tool, 787–789
Executive SGT (Security Group Tag), 555
ExernalBlue, 868
exit command, 932
expanding policy sets, 211
expired certificates, 32–33, 465
exploits, definition of, 872, 993
Export command (Dashboard settings), 134
Ext Id Sources tab (Guest Access work center), 340
extended key usage (EKU), 211
Extensible Authentication Protocol. ***See*** **EAP (Extensible Authentication Protocol)**
Extensible Communication Platform (XCP) server, 825
Extensible Messaging and Presence Protocol (XMPP), 825
External CA Settings, 520
external data source condition, 100
External Identity Sources tab (Identity Management), 162
external identity stores, 23–33, 993. ***See also*** **certificate-based authentication**
AD (Active Directory), 24–25
configuration, 196
definition of, 23, 993
LDAP (Lightweight Directory Access Protocol), 25–26
MFA (multifactor authentication), 26–29
OTP (one-time password) services, 29
smart cards, 29
External MDM tab (Network Resources), 165
External RADIUS Servers tab (Network Resources), 165
External RESTful Services (ERS), 131, 850–853

F

faceprint technology, 27
Fail Live Log status, 767
FAIL message, 10
Failure Reason dashlet, 136
fallback, TACACS+, 946–948
FAST (Flexible Authentication via Secure Tunneling)
EAP chaining, 216
EAP-FAS, 215
EAP-FAST, 45, 48–49
FDM (Firepower Device Management), 832
Feed Service tab (Administration screen), 169
FIDO2, 310
files
CRC32, 672
file conditions, 100, 667–673
FileDate, 669–671
FileExistence, 671–678
FileVersion, 672
ise-2.7.0.356.SPA.x86_64.iso, 112

ISE-2.7.0.356-virtual-SNS3615-SNS3655–300.ova, 112
ISE-2.7.0.356-virtual-SNS3655-SNS3695–1200.ova, 112
ISE-2.7.0.356-virtual-SNS3695–2400.ova, 112
paths for, 669–670
policy for, 612
remediations for, 682
SHA-256, 672
FileVersion file type, 672
filtering
endpoint attributes, 444–445
Live Log, 771
Finance SGT (Security Group Tag), 555
fingerprint technology, 27
Firefox, support for, 122
Firepower
FDM (Firepower Device Management), 832
FTD (Firepower Threat Defense), 629
SGFW (security group firewall) on, 612–613
firewalls, 611
conditions for, 100, 660–661
posture policy conditions, 660–661
remediations, 682
SGFWs (security group firewalls), 611–613
on ASA (Adaptive Security Appliances), 612
on ASR (Aggregation Services Router), 613
definition of, 999
on Firepower, 612–613
on ISR (Integrated Services Router), 613
FlexAuth (Flexible Authentication)
configuration, 269–272
definition of, 994
FAST (Flexible Authentication via Secure Tunneling)
EAP chaining, 216
EAP-FAS, 215
EAP-FAST, 45, 48–49
FMC (Firepower Management Center) configuration, 831–850
access rules, 840–844
active users, viewing, 844–845
FDM (Firepower Device Management), 832
pxGrid integration, 832–837
Rapid Threat Containment, 845–850
realms, 837–840
FOLLOW message, 11
Forbes, Paul, 633
forests, AD (Active Directory), 24
form factors
ISE (Identity Services Engine), 177
SNS (Secure Network Server) appliances, 111–112
FQDN (fully qualified domain name), 416, 833
Framed-IP-Address attribute (RADIUS), 265, 414
FTD (Firepower Threat Defense), 629
Full Configuration Database support bundles, 782
Full Wipe option (endpoint management), 543
fully qualified domain name (FQDN), 416, 833
Funk Software, 45

G

gateway providers, SMS (Short Message Service), 388–389
GCL (pxGrid common library), 825
General tab
corporate WLAN, 292
guest WLAN, 287–288
Generic Token Card (GTC), 43, 45, 215
global CoA (Change of Authorization), 442–443, 994

global configuration AAA commands, 261–262
global configuration RADIUS commands, 262–269
device tracking in IOS Xe 16.x and later, 267
global 802.1X commands, 266–267
IOS 12.2.x, 262–263, 264–266
IOS 15.x, 263–266
IOS XE, 263–266
local ACL (access control list) creation, 268–269
Global Page Customizations (portal page customization), 361
global profiler settings, 444–445
global search for endpoints, 139–140
Global Search tool, verifying profiling with, 454–455
Good Technology, 708
Google Chrome, 122
Google Client Messaging servers, 101
Google Play Store, 282, 492–493, 506
Google URLs for ACL Bypass, 282–283
graphical user interface. *See* GUI (graphical user interface)
GROUP_ACCOUNTS sponsor group, 381
groups
Admin group roles, 127–132
endpoint, 22
endpoint identity, 22, 354, 450–452, 535
local endpoint, 195
local user identity, 194
NDGs (network device groups), 720–721, 996
node, 748–750, 996
user identity, 22, 339–340, 840–844, 1000
Groups tab (Identity Management), 162
GRTC (Generic Token Card), 43, 45
GTC (Generic Token Card), 43, 215
Guest Access work center
Administration tab, 340–341
Ext Id Sources tab, 340
Identities tab
endpoints, 338
ISS (identity source sequence), 339
user identity groups, 339–340
overview page, 337–338
Portals & Components tab, 341
guest portals, 341–342
guests, 341, 343–348
overview of, 341
sponsor portals, 341
guest dynamic interface, 285–286
Guest Exceeds Limit setting, contractor accounts, 345
Guest Flow attribute, 323–324, 994
guest location setting, 369–371, 994
Guest management license packages, 156
guest portals, 341–342
authorization policies for, 348–351
definition of, 341–342
hotspot, 351–358
AUP (acceptable use policy) page settings, 354–356
authentication success settings, 357
authorization rule configuration, 362–365
configuration flowchart for, 351
definition of, 342
interface bonding, 352–353
portal page customization, 358–362
portal settings, 352–354
post-access banner page settings, 355–356
support information page settings, 357–358
VLAN DHCP release page settings, 355–356
self-registered
authorization rule configuration, 373–380
BYOD settings, 372
configuration flowchart, 365–366
definition of, 342

guest change password settings, 371–372
guest device compliance settings, 373
guest device registration settings, 371–372
guest location setting, 369–371, 994
login page settings, 367–368
portal settings, 366–367
registration form settings, 368–371
self-registration success, 371
sponsored, 380–381
AUP (acceptable use policy) page settings, 386
configuration flowchart, 380–381
default sponsor portal, 384
definition of, 342
login settings, 386
other settings, 387
portal settings, 385–386
provisioning guest accounts from, 389–394
Guest reports, 150
guest services, 337–341
administration, 340–341
authorization policies for, 348–351
guests, 343–348
contractors, 344–346
daily, 346–347
definition of, 341
overview of, 343
provisioning from sponsor portals, 389–394
social, 348
sponsor, 341
weekly, 347
hotspot guest portals, 351–365
AUP (acceptable use policy) page settings, 354–356
authentication success settings, 357
authorization rule configuration, 362–365
configuration flowchart for, 351
definition of, 341–342
interface bonding, 352–353
portal page customization, 358–362
portal settings, 352–354
post-access banner page settings, 355–356
support information page settings, 357–358
VLAN DHCP release page settings, 355–356
identities, 338–340
importance of, 334
notification services, 388–389
SMS gateway providers, 388–389
SMTP servers, 388
overview of, 337–341
SAML (Security Assertion Markup Language) authentication, 394–400
self-registered guest portals
authorization rule configuration, 373–380
BYOD settings, 372
configuration flowchart, 365–366
definition of, 342
guest change password settings, 371–372
guest device compliance settings, 373
guest device registration settings, 371–372
guest location setting, 369–371, 994
login page settings, 367–368
portal settings, 366–367
registration form settings, 368–371
self-registration success, 371
sponsored guest portals, 380–381
AUP (acceptable use policy) page settings, 386
configuration flowchart, 380–381
default sponsor portal, 384
definition of, 342

login settings, 386
other settings, 387
portal settings, 385–386
provisioning guest accounts from, 389–394
sponsors
definition of, 381, 999
sponsor groups, 381–382
Guest Status dashlet, 136
Guest Type dashlet, 136
guest WLAN configuration, 287–290
AAA Servers tab, 289–290
Advanced tab, 290
General tab, 287–288
Layer 2 Security tab, 288
Layer 3 Security tab, 289
Guest_Portal_Sequence ISS, 319, 339
guests, 343–348
contractors, 344–346
daily, 344–346
definition of, 341, 994
overview of, 343
provisioning from sponsor portals, 389–394
social, 348
weekly, 347
Guests tab (Home page), 135–136
GUI (graphical user interface)
Admin group roles, 127–132
Administration portal, 137–142
global search for endpoints, 139–140
Help menu, 138, 140–141
ISE Setup Wizards, 141
Settings menu, 142
tabs, 137–139
Administration screen, 142, 155–170
Device Portal Management tab, 166–169
Feed Service tab, 169
Identity Management tab, 161–163
Network Resources tab, 163–166
PassiveID Setup option, 138
pxGrid Services tab, 169
Search icon, 138
System Activities option, 139
System tab, 155–161
Threat Centric NAC tab, 170
Visibility Setup option, 138
Wireless Setup (BETA) option, 139
browser requirements for, 125
Context Visibility screen, 137, 142, 143
definition of, 994
downloading, 780
Home dashboards, 132–137
initial login, 125–126
Operations screen, 142, 143–150
ANC (Adaptive Network Control) component, 147–148, 991
RADIUS tab, 144–146
Reports tab, 150
TACACS Live Log tab, 147
Threat-Centric NAC Live Logs tab, 146
Troubleshoot tab, 147–148
Policy page, 138, 142, 150–154
Client Provisioning tab, 153
Policy Elements tab, 154
Policy Sets tab, 150–151
Posture tab, 152
Profiling tab, 152
support bundles, 783
supported browsers, 122–123
Work Centers screen, 142, 170–171

H

HA (high availability)
configuration, 269–272
RADIUS fallback, 279–280
hardware attributes condition, 100, 655
headends, AnyConnect
configuration of, 640–642

deployment packages
AnyConnect configuration, building, 648–649
posture profile configuration, 644–648
uploading to ISE, 642–644
number of, 640
Hello (Windows), 27
help command, 932
Help menu, Administration portal, 138, 140–141
Helpdesk admin role, 127
HelpDesk command set, 941–942
helpdesk group, 945
HelpDesk profile, 939, 940, 976
high availability (HA), 743
Anycast, 753–756
IP SLAs (service-level agreements), 754–756
network architecture, 753–754
route redistribution, 755–756
backup and restore, 759–761
configuration, 269–272
failure scenarios, 753
general guidelines for, 752–753
licensing in multi-mode ISE cube, 747–748
load balancing
IOS (Internetwork Operating System), 756–757
PSNs (Public Services Networks), 751–752
MnT (Monitoring and Troubleshooting) nodes, 109–110, 743–744
logging categories, 744
logging flows, 743
logging targets, 743–744
node groups, 748–750
PAN (Policy Administration node), 109, 743–744
patches, 757–759
RADIUS fallback, 279–280
high-security mode. *See* closed mode
Home dashboards, 132–137
Dashboard Settings menu, 134
Endpoints tab, 134–135
Guests tab, 135–136
Summary tab, 134
Threat tab, 137
Vulnerability tab, 136
hop-by-hop encryption, 548
host mode, switch port, 272–274
Host-Name attribute (DHCP), 411
hostname command, 946
HostScan, 629
hotspot guest portals, 342, 351–358
AUP (acceptable use policy) page settings, 354–356
authentication success settings, 357
authorization rule configuration, 362–365
configuration flowchart for, 351
interface bonding, 352–353
portal page customization, 358–362
portal settings, 352–354
post-access banner page settings, 355–356
support information page settings, 357–358
VLAN DHCP release page settings, 355–356
"HowTo: Cisco and F5 Deployment Guide-ISE Load Balancing Using BIG-IP" (Hyps), 752
HR SGT (Security Group Tag), 555
hrDeviceDescr option (Nmap), 416
HTTP (Hypertext Transfer Protocol)
HTTP probe, 420–421
HTTP/HTTPS servers, enabling, 314
HTTPS (HTTP Secure), 126, 180
POST method, 84–85
hybrid ISE deployment, 116–117
Hyper-V, 113
Hyps, Craig, 633, 752

I

IBM MaaS360, 708
IBM Tivoli Identity Manager (TIM), 25
identification profiles, WSA (Web Security Appliance), 855–857
Identities tab
Device Administration, 920
Guest Access work center
endpoints, 338
ISS (identity source sequence), 339
user identity groups, 339–340
Identity Management, 161
Identity admin role, 127
Identity column, Live Log, 767
Identity Management tab (Administration screen), 161–163
External Identity Sources tab, 162
Groups tab, 162
Identities tab, 161
Identity Source Sequences tab, 163
Settings tab, 163
identity providers. *See* IdPs (identity providers)
Identity Services Engine. *See* ISE (Identity Services Engine) architecture
Identity Source Sequences tab (Identity Management), 163
identity sources, 34, 35
identity stores, 21–33, 993
authentication policy for, 219–220
CAPs (certificate authentication profiles), 202
definition of, 20–21
EAP (Extensible Authentication Protocol). *See* EAP (Extensible Authentication Protocol)
external
Active Directory. See AD (Active Directory)
certificates. See certificate-based authentication
configuration, 196
definition of, 23, 993
LDAP (Lightweight Directory Access Protocol), 25–26
MFA (multifactor authentication), 26–29
OTP (one-time password) services, 29
smart cards, 29
identity sources, 21, 994
identity validation, 211
internal, 21–22
endpoint groups, 22
user identity groups, 22, 1000
ISS (identity source sequence), 34, 202–204, 218, 319, 339, 472, 994
local endpoint groups, 195
local user identity groups, 194
local users, 195–196
selection of, 210–211
Idle Timeout option (TACACS profile), 933
IdPs (identity providers), 35, 394–400, 998
IEEE 802.1X. *See* 802.1X
IETF (Internet Engineering Task Force), 12
NEA (Network Endpoint Assessment), 626
RADIUS. *See* RADIUS (Remote Access Dial-In User Service)
IF.THEN policy rules, 154
in authentication policies, 216
in authorization policies, 236
IMEI attribute, 537
importing public certificates, 476–477
incidents, 899–900
Include Service Version Information option (Nmap), 416
IND (Industrial Network Director), 827
indicators of compromise (IoCs), 899, 994
Industrial Network Director (IND), 827
Informational Live Log status, 767
infrastructure configuration
DHCP helper, 424–427
IOS Device Sensor, 426–427
SPAN (Switched Port Analyzer), 424–425

VACLs (VLAN Access Control Lists), 425–426
VMware vSwitches, 427
ingress access controls
ACLs (access control lists), 553–554
east-west segmentation, 554–555
VLAN assignment, 551–553
int eth1/3N7K-DIST command, 578
int GigabitEthernet 2 command, 352
Integrated Services Router. *See* ISR (Integrated Services Router)
integration
definition of, 820, 994
MDM (mobile device management), 820–821
administrative management, 545
configuration, 537–538
endpoint management, 542–543
integration points, 536–537
onboarding rules, 539–542
self-management, 543–544
Rapid Threat Containment, 821–823
ANC (Adaptive Network Control), 822–823
definition of, 998
EPS (Endpoint Protection Services), 821–822, 993
interface bonding, hotspot guest portals, 352–353
interface configuration, 269–276
application of initial ACL to port, 275–276
authentication settings, 274–275
authentication timers, 275
configuration of interfaces as switch ports, 269
FlexAuth (Flexible Authentication), 269–272
HA (high availability), 269–272
host mode of switch port, 272–274
interface g1/0/1 command, 267
interface range command, 269
intermediate CAs (certificate authorities), 521–523
internal blocking, 896
Internal CA Settings menu (ISE Certificate Authority), 520
internal endpoint database, 22, 994
internal identity stores, 21–22
endpoint groups, 22
user identity groups, 22, 1000
internal user database, 994
Internet Engineering Task Force. *See* IETF (Internet Engineering Task Force)
Internet Explorer, 122
Internet-only access for smart devices, 243–246
intrusion prevention systems (IPSs), 661, 818
IoCs (indicators of compromise), 899, 994
iOS. *See* Apple iOS
IOS (Internetwork Operating System)
device administration AAA with, 930
device administration service, enabling, 937
network device preparation, 937–939
overview of, 932
policy preparation, 939–946
privilege levels, 932–933, 997
TACACS profiles, 932–934
TACACS+ authentication and fallback, 946–948
TACACS+ command accounting, 951
TACACS+ command authorization, 948–950
TACACS+ command sets, 934–936, 992
testing and troubleshooting in ISE, 952–954
troubleshooting at IOS command line, 954–966

Device Sensor feature, 426–427
global configuration RADIUS commands
device tracking in IOS Xe 16.x and later, 267
IOS 12.2.x, 262–263, 264–266
IOS 15.x, 263–266
IOS XE, 263–266
local ACL (access control list) creation, 268–269
IOS XE switches
configuration for TrustSec, 560–563
global configuration RADIUS commands, 263–266
manual SGT (security group tag) propagation on, 595–597
IOS-based NADs, 495
load balancing, 756–757
SXP (SGT Exchange Protocol) configuration on, 588–590
IOS_ Network _CS, 943
IOS_Admin_Privilege profile, 940–941
IOS_HelpDesk_CS, 941–942
IOS_HelpDesk_Privilege profile, 940
IOS_Security_CS, 942–943
IoT SGT (Security Group Tag), 555
ip access-group ACL-ALLOW command, 276
ip access-list ext ACL-AGENT-REDIRECT command, 269
ip access-list ext ACL-DEFAULT command, 268
ip access-list ext ACL-WEBAUTH-REDIRECT command, 268, 315
ip access-list extended ACL-ALLOW command, 268
ip access-list extended command, 425, 948
IP Address column, Live Log, 768
ip device tracking command, 266, 298
ip device tracking probe auto-source command, 267
ip domain-name command, 313, 946
ip helper-address command, 412–413, 424
ip http active-session-modules none command, 314
ip http secure-server command, 314
ip http server command, 314
ip radius source-interface command, 266
IP SLAs (service-level agreements), 754–756
ip ssh version 2 command, 947
iPad, 482. ***See also*** **Apple iOS**
iPCU (iPhone Configuration Utility), 812
iPhone, 482. ***See also*** **Apple iOS**
iPhone Configuration Utility (iPCU), 812
IPSs (intrusion prevention systems), 612, 661, 818
ISE (Identity Services Engine) architecture, 6. ***See also*** **profiling**
Anomalous Behaviour Detection, 406–408, 991
configuration for BYOD onboarding, 495–496, 510–523
authorization profiles, 516
authorization rules for EAP-TLS authentications, 518
authorization rules for onboarding, 517
Blackberry example, 508–509
client provisioning policy configuration, 512–514
default unavailable client provisioning policy action, 515
dual SSID with Android example, 503–508
ISE as certificate authority, 519–520, 521–523, 994
native supplicant profile, 510–512
SCEP (Simple Certificate Enrollment Protocol), 520–521, 999
single SSID with Apple iOS example, 496–503
WebAuth configuration, 514–515
configuration for CWA (Centralized Web Authentication), 317–322, 327
authorization profiles, 320–322
dACLs (downloadable access control lists), 319–320

Guest_Portal_Sequence ISS (identity source sequence), 319
MAB (MAC Authentication Bypass), 96–99, 318
configuration for MACsec, 619
configuration for pxGrid, 828–831
cubes, 182
definition of, 109, 737, 995
licensing in, 747–748
deployment
distributed, 116–119, 737–742
hybrid, 116–117
overview of, 113
physical versus virtual, 111–113
scale limits for, 118–119
single-node, 113
two-node, 114–116
device administration AAA with Cisco IOS
device administration service, enabling, 937
network device preparation, 937–939
policy preparation, 939–946
TACACS+ authentication and fallback, 946–948
TACACS+ command accounting, 951
TACACS+ command authorization, 948–950
Endpoint Profile Policies, 431–441
form factors, 177
GUI (graphical user interface). *See* GUI (graphical user interface)
high availability, 743
Anycast high availability, 753–756
backup and restore, 759–761
failure scenarios, 753
general guidelines for, 752–753
licensing in multi-mode ISE cube, 747–748
load balancing, 751–752, 756–757
MnT (Monitoring and Troubleshooting) node, 109–110, 743–744
node groups, 748–750
PAN (Policy Administration node), 109, 743–744
patches, 757–759
initial configuration, 174
AD (Active Directory), 196–202
bootstrapping, 177–180
CAPs (certificate authentication profiles), 202
certificates, 181–191
external identity stores, 196, 993
form factors, 177
installation guides, 177
ISS (identity source sequence), 202–204, 994
local endpoint groups, 195
local user identity groups, 194
local users, 195–196
NADs (network access devices), 192
NDGs (network device groups), 192–194
SSL (Secure Sockets Layers), 181
TLS (Transport Layer Security), 181–182
ISE NAC feature, 315–316
ISE nodes in distributed environment, 737–742
ISE cubes, 737
ISE persona types, 737
node personas, 742
PPAN (Policy Administration Node), 738–739
primary devices, 738–739
registration of ISE nodes, 739–742
licensing
in multi-mode ISE cube, 747–748
packages, 155–158
Live Log, 327
NADs (network access devices), 113
network probes, 409–423
Active Directory, 422
configuration, 409–411

DHCP and DHCPSPAN, 411–414
DNS, 417
HTTP, 420–421
NETFLOW, 419–420
Nmap (network scan), 415–417
publishing endpoint probe data on, 450
pxGrid, 423
RADIUS, 414–415
SNMPQUERY and SNMPTRAP, 417–419
nodes
configuration in distributed environment, 737–742
definition of, 109
MnT (Monitoring and Troubleshooting), 109–110, 743–744
node groups, 748–750
PAN (Policy Administration node), 109, 745–748
PSN (Policy Services node), 110
single-node ISE deployment, 113
two-node ISE deployment, 114–116
overview of, 106–108
personas
Administration, 109, 737
definition of, 108–109
Monitoring, 109–110, 737
Policy Services, 110, 737
pxGrid, 111, 737
types of, 737
posture assessment. *See* posture assessment
Profiler Feed Service, 404–406, 429–430
CoA (Change of Authorization) with, 442–444
configuration, 429
verification of, 429–430
SXP (SGT Exchange Protocol) configuration on, 584–586
troubleshooting. *See* troubleshooting tools
verification of, 302–303
Live Sessions, 303
RADIUS Live Log, 302–303
version scalability, 118–119
virtual appliances, 177
WLC device administration AAA
ISE configuration on WLC TACACS+ servers, 979–980
network device preparation, 972
policy results preparation, 974–977
policy sets, 977–979
ISE Community Page option (Help menu), 141
ISE Documentation option (Help menu), 141
ISE on Cisco.com option (Help menu), 141
ISE Partner Ecosystem option (Help menu), 141
ISE Passive Identity Connector (ISE-PIC), 157, 858
ISE Portal Builder option (Help menu), 141
ISE Profiler. *See* profiling
ISE Setup Wizards, 141
ISE Software Downloads option (Help menu), 141
ISE YouTube Channel option (Help menu), 141
ise-2.7.0.356.SPA.x86_64.iso file, 112
ISE-2.7.0.356-virtual-SNS3615-SNS3655–300.ova file, 112
ISE-2.7.0.356-virtual-SNS3655-SNS3695–1200.ova file, 112
ISE-2.7.0.356-virtual-SNS3695–2400.ova file, 112
ISE-PIC (ISE Passive Identity Connector), 157, 858
ISR (Integrated Services Router), 613
ISS (identity source sequence), 34, 202–204, 218, 319, 339, 472, 994
Issued Certificates menu (ISE Certificate Authority), 520
IT Users Access authorization rule, 252–256

J

Jabber, 825
JailBrokenStatus attribute, 536
Jamf Pro, 673, 708
Java applets, 85–86, 153, 322, 356
JGroups, 748
Jobs, Steve, 482
joining AD (Active Directory) domains, 197–202
Juniper, 45

K

Karelis, E. Peter, 753
Kerberos, 196, 469
key command, 264
keys. ***See also*** **PKI (public key infrastructure)**
EKU (extended key usage), 211
key pairs, 468–469, 995
MKA (MAC Security Key Agreement), 616
private, 468–469
public, 468–469, 998
Kpasswd, 196

L

L (Locality) field, 183
Lancope Stealthwatch. ***See*** **Stealthwatch**
LANs (local area networks)
EAP over LAN. *See* 802.1X
VLANs (virtual LANs), 726
assignment of, 551–553, 726
authentication VLAN, 87–88
dynamic interfaces for, 284–286
mapping to SGTs (security group tags), 580
segmentation with, 322
VACLs (VLAN Access Control Lists), 424, 425–426
Launch Page Level Help option (Help menu), 140
launch program remediations, 683
Layer 2 Security tab
corporate WLAN, 292–293
guest WLAN, 288
Layer 3 Security tab
corporate WLAN, 293
guest WLAN, 289
Layout Template command (Dashboard settings), 134
LDAP (Lightweight Directory Access Protocol), 23, 25–26, 196, 995
lease, posture, 691
left-zero-padded keyword, 621
library, Conditions Studio, 218
library conditions, 708
licensing, 155–158, 747–748
Lightweight Directory Access Protocol. ***See*** **LDAP (Lightweight Directory Access Protocol)**
Lightweight Directory Access Protocol (LDAP), 23, 25–26, 196, 995
Line-of-Business-1 SGT (Security Group Tag), 555
Line-of-Business-2 SGT (Security Group Tag), 555
Link encryption (MACsec), 156
Link Layer Discovery Protocol (LLDP), 98, 418
link remediations, 684
Linux KVM, 113
lists
ACLs (access control lists). *See* ACLs (access control lists)
CRLs (certificate revocation lists), 33, 466
Live Log, 766–771
Actions menu, 770
advanced filtering, 771
authentication details report, 771–774
Authorization Policy column, 767
Authorization Profiles column, 768

blank lines in, 774–775
Cisco ISE verification with, 302–303
Client Stopped Responding counter, 768
CWA (Centralized Web Authentication) verification from, 327
Details icon, 768
device administration AAA with Cisco IOS, 952–954
Endpoint ID column, 767
Endpoint Profile column, 767
Identity column, 767
IP Address column, 768
MDM Server column, 768
Misconfigured Network Devices counter, 768
Misconfigured Supplicants counter, 768
Network Device column, 768
Quick Filters counter, 769
RADIUS (Remote Access Dial-In User Service), 144–146
RADIUS Drops counter, 768
Record Selector, 770
Refresh counter, 770
Repeat counter, 767, 769
Server column, 768
status types, 767
TACACS+, 147, 982–983
Threat-Centric NAC, 146
Time Selector, 770
verification of BYOD flows with, 534

Live Sessions, 776
Cisco ISE verification with, 303
RADIUS (Remote Access Dial-In User Service), 145–146

LLDP (Link Layer Discovery Protocol), 98, 418

load balancing
IOS (Internetwork Operating System), 756–757
PSNs (Public Services Networks), 751–752

LOBBY role, 972

local ACLs (access control lists), 268–269

local endpoint groups, 195

Local Logs support bundles, 783

local user identity groups, 194

local users, 195–196

Local Web Authentication. *See* LWA (Local Web Authentication)

Locality (L) field, 183

Location dashlet, 136

Location Services tab (Network Resources), 166

Log %temp%\spwProfileLog.txt command, 812

logging categories, 778–779

logging host command, 299

logging in to ISE (Identity Services Engine), 125–132
Admin group roles, 127–132
Administration portal, 137–142
global search for endpoints, 139–140
Help menu, 138, 140–141
ISE Setup Wizards, 141
Settings menu, 142
tabs, 137–139
daily guest accounts, 347
Home dashboards, 132–137
initial login, 125–126
self-registered guest portals, 367–368
social, 35
sponsored guest portals, 386

logging monitor informational command, 299

logging origin-id ip command, 299

logging source-interface command, 299

logging synchronous command, 947

Logging tab (System), 159

logging targets, 743–744, 777–778

logical profiles, 441–442, 995

login authentication command, 947

Login Page Settings (Client Provisioning portal), 651–652

Logoff option, posture reassessment, 692
logout command, 932
logs, 159, 766–785
debug, 779–784
configuration, 779–780
downloading from GUI, 780
support bundles, 782–784
viewing from CLI, 781–782
de-duplication of, 805–807
Live Log. *See* Live Log
Live Sessions, 776
logging categories, 744, 778–779
logging flows for, 743
logging targets for, 743–744, 777–778
Lost option (endpoint management), 543
low-impact mode, 725–727
LWA (Local Web Authentication), 310–311
with centralized portal, 84–85
definition of, 995
when to use, 84

M

MaaS360, 708
MAB (MAC Authentication Bypass)
authentication with, 80–82, 227–228
configuration, 265, 270–274, 318
definition of, 717, 995
MAB rule flowchart, 217
overview of, 96–99
profiling, 96–99
DHCP (Dynamic Host Configuration Protocol), 97
RADIUS (Remote Access Dial-In User Service), 97
role-specific authorization rules, 241
wireless, 489
MAC (Media Access Control) addresses
identity management with, 20
MAB (MAC Authentication Bypass), 80–82, 995
authentication with, 227–228
configuration, 265, 318
definition of, 717
MAB rule flowchart, 217
overview of, 96–99
profiling, 96–99
role-specific authorization rules, 241
wireless, 489
Mac filtering on WLCs, 315
MACsec, 995
MAM (MAC Address Management) model, 451
MKA (MAC Security Key Agreement), 616
URL-redirected MAC authentication bypass, 311–313
MAC Filtering option, WLC (Wireless LAN Controller), 315
Mac OSX
Console application, 812
onboarding flow, 531–533
device provisioning, 532–533
device registration, 531
Machine Access Restriction (MAR), 472
Machine Authentication, 66, 545–546
macro-segmentation, 613
MACsec, 995
definition of, 548
Downlink MACsec, 616–619
authorization profile for, 616
ISE (Identity Services Engine) configuration, 619
policies, 616–618
switch configuration modes, 618–619
history of, 614–616
Multi-Hosts mode, 273, 298, 618, 995
Uplink MACsec, 619–623
Maintenance tab (System), 159
Make Primary button, 738

malware
anti-malware posture policy conditions, 661–663
anti-malware remediations, 681–682
MAM (MAC Address Management) model, 451
Manage Dashboards command (Dashboard settings), 134
MANAGEMENT role, 971
manual SGT (security group tag) propagation, 595–597
manually assigning SGTs (security group tags), 577–578
Manufacturer attribute, 536
mapping
subnets to SGTs (security group tags), 580
VLANs to SGTs (security group tags), 580
MAR (Machine Access Restriction), 472
matrix view (TrustSec policy matrix), 553
Maximum Access Time setting
contractor accounts, 344
daily guest accounts, 347
weekly guest accounts, 347
Maximum Privilege option (TACACS profile), 933
Maximum Simultaneous Logins setting, 345
MD5 algorithm, 43, 46, 214
MDA (Multidomain Authentication), 273, 298, 618, 995
MDM (mobile device management), 820–821. *See also* BYOD (bring your own device) onboarding
definition of, 995
dictionary, 708–709
onboarding
administrative management, 545
advantages of, 535–536
attributes in authorization policies, 536–537
definition of, 487
endpoint management, 542–543
integration configuration, 537–538
onboarding rule configuration, 539–542
self-management, 543–544
overview of, 101–102
posture assessment with, 108
supported features, 101–102
vendors, 101–102
MDM Server column (Live Log), 768
MDM_Compliant authorization rule, 709, 711
MDM_Detailed_Posture condition, 710
MDM_NotCompliant authorization rule, 711
MDM_NotCompliant condition, 709–710
MDM_NotReachable authorization rule, 710
MDMFailureReason attribute, 537
MDMServerName attribute, 537
MDMServerReachable attribute, 537, 712
Media Access Control. *See* MAC (Media Access Control) addresses
MEID attribute, 537
Meraki Systems Manager (Meraki SM), 708
meshed SXP (SGT Exchange Protocol), 582–583
messages
CoA (Change of Authorization), 110, 311
RADIUS, 13–14
accounting messages, 13
authentication messages, 13
authorization messages, 13
syslog, 299–300
TACACS+
accounting messages, 11
authentication messages, 9
authorization messages, 10–11
communication flows, 12
Metrics dashlet, 134
MFA (multifactor authentication), 26–29
micro-segmentation, 614

Microsoft Active Director. *See* AD (Active Directory)
Microsoft CHAP. *See* MS-CHAP (Microsoft CHAP)
Microsoft Edge, 122
Microsoft Hyper-V, 113
Microsoft Internet Explorer, 122
Microsoft Intune, 673, 708
Microsoft Remote Procedure Call (MSRPC), 196
Miller, Darrin, 633
Misconfigured Network Devices counter, 768
Misconfigured Supplicants counter, 768
MKA (MAC Security Key Agreement), 616
MnT (Monitoring and Troubleshooting) admin role, 127
MnT (Monitoring and Troubleshooting) nodes, 109–110, 743–744, 912
Dedicated MNT, 742
definition of, 995
logging categories, 744
logging flows, 743
logging targets, 743–744
sending syslog messages to, 299
mobile device management. *See* MDM (mobile device management)
Mobile Device Management portal, 168
Mobile Iron UEM, 708
mobile posture, 707–712
authorization conditions, 709–710
authorization rules, 710–712
supported device managers, 707–709
Mobile Preview pane (portal page customization), 358
Mobility license packages, 156
Mobility Upgrade license packages, 157
Model attribute, 537
modes
closed, 728–730
low-impact, 725–727
monitor, 722–725
flow diagram, 723–724
operational flow, 722–723
policy set creation, 724–725
transitioning to end state, 730–731
modify device capabilities, 101
modules, compliance, 637–638
monitor mode, 722–725
flow diagram, 723–724
operational flow, 722–723
policy set creation, 724–725
transitioning to end state, 730–731
MONITOR role, 972
monitor session command, 425
Monitoring and Reporting Logs support bundles, 783
Monitoring and Troubleshooting nodes. *See* MnT (Monitoring and Troubleshooting) nodes
Monitoring persona, 109–110, 737
Mozilla Firefox, 122
MS-CHAP (Microsoft CHAP), 43, 45, 46
MSE integration, 156
MSRPC (Microsoft Remote Procedure Call), 196
Multi-Auth (Multiauthentication), 273, 298, 618, 995
Multidomain Authentication (MDA), 273, 298, 618, 995
multifactor authentication (MFA), 26–29
multi-hop, 582
Multi-Hosts mode (MACsec), 273, 298, 618, 995
multi-node deployment. *See* cubes, ISE
Multiple Hosts (Multi-Hosts), 273, 298
must-not-secure policy (MACsec), 616, 618
must-secure policy (MACsec), 616, 618
My Devices portal, 168, 543–544
MyDevices_Portal_Sequence, 544

N

NAC (network access control) projects, 258. ***See also*** **AAA (authentication, authorization, and accounting)**

NAC (Network Admission Control), 626, 630

NAC Appliance, 631–633

"NAC Framework" solution, 630

NAC Managers tab (Network Resources), 165

NADs (network access devices), 7, 107. ***See also*** **authorization**

configuration, 917

ACLs (access control lists), 492–495

for BYOD onboarding, 489–495

connection settings, 918

identities, 920

network resources, 921–922

overview of, 916–917

password change control, 918

policy elements, 922–924

policy sets, 925–927

reports, 927

Session Key Assignment, 918

UI navigation for, 919–920

WLAN configuration, 490–491

CWA (Centralized Web Authentication) verification, 327–331

client details, viewing on WLC, 329–331

show commands on wired switch, 328–329

definition of, 113, 996

enforcement types for, 50

ISE initial configuration, 192

role of, 49–50

TrustSec-enabled

defining TrustSec settings for, 559

PACs (Protected Access Credentials), 558–559

verification of, 296–302

with Cisco switches, 296–299

with Cisco WLC (Wireless LAN Controller), 300–302

with syslog messages, 299–300

NAM (Network Access Manager), 991

named ACLs, 244

NAP (Network Access Protection), 626

NAS-Port-Id attribute (RADIUS), 415

NAS-Port-Type attribute (RADIUS), 415

National Security Agency, 868

native EAP (Extensible Authentication Protocol), 43–44, 996

native SGT (security group tag) propagation, 593–597

native supplicant profile, 510–512, 642

native tagging, 593–597

configuration on IOS XE switches, 595–597

Layer 2 frame format with, 593–594

pervasive tagging, 594

NDAC (Network Device Admission Control), 566, 996

NDES (Network Device Enrollment Services), 520

NDGs (network device groups), 192–194, 720–721, 996

NDS (Novell Directory Services), 25

NEA (Network Endpoint Assessment), 626

net-admin group, 945

NetAdmin profile, 974–975

NETFLOW probe, 419–420

NetIQ eDirectory. ***See*** **AD (Active Directory)**

netsh ras set tracing * enable command, 296

network access AAA, 3, 7, 996

network access control (NAC) projects, 258. ***See also*** **AAA (authentication, authorization, and accounting)**

network access devices. ***See*** **NADs (network access devices)**

Network Access Manager (NAM), 809, 991

Network Access Protection (NAP), 626

network access users, 21
Network Administrator command set, 943
Network Administrator role, 939
Network Admission Control. *See* NAC (Network Admission Control)
Network device admin role, 128
Network Device Admission Control (NDAC), 566, 996
Network Device column (Live Log), 768
Network Device Enrollment Services (NDES), 520
network device groups (NDGs), 720–721, 996
Network Device Groups tab (Network Resources), 164–165
Network Device Profiles tab (Network Resources), 165
network device troubleshooting, 812–815. *See also* device administration AAA
client details, viewing on WLC, 813–814
debug commands, 815
show authentication interface command, 812–813
Network Devices dashlet, 134, 135
Network Devices tab, 143, 163–164
Network Devices widget, 454
Network Endpoint Assessment (NEA), 626
Network Groups view, Cisco AnyConnect NAM supplicant, 60, 71
Network Infrastructure SGT (Security Group Tag), 555
network interface cards (NICs), 97
network probes. *See* probes
Network Resources tab (Administration screen), 163–166
External MDM tab, 165
External RADIUS Servers tab, 165
Location Services tab, 166
NAC Managers tab, 165
Network Device Groups tab, 164–165
Network Device Profiles tab, 165
Network Devices tab, 163–164
RADIUS Server Sequences tab, 165
Network Resources tab (Device Administration), 921–922
network scan (Nmap), 415–417
Network Services Platform. *See* NSP (Network Services Platform)
Network Services SGT (Security Group Tag), 555
Network Setup Assistant app, 282, 507
Network Time Protocol (NTP), 32, 196, 465, 996
Networks view, Cisco AnyConnect NAM supplicant, 60, 62–70
Certificates tab, 67
Connection Type tab, 66
Credentials tab, 68–70
Machine Auth tab, 66
PAC Files tab, 67–68
Security Level tab, 64–66
User Auth tab, 68–69
NGFWs (next-generation firewalls), 818
NICs (network interface cards), 97
Nmap (network scan), 415–417
No Escape option (TACACS profile), 933
nodes
configuration in distributed environment, 737–742
ISE cubes, 737
ISE persona types, 737
node personas, 742
PPAN (Policy Administration Node), 738–739
primary devices, 738–739
registration of ISE nodes, 739–742
definition of, 109, 996
MnT (Monitoring and Troubleshooting), 109–110, 743–744
Dedicated MNT, 742
definition of, 995
logging categories, 744

logging flows, 743
logging targets, 743–744
sending syslog messages to, 299
node groups, 748–750, 996
PAN (Policy Administration node), 109, 745–748
auto PAN switchover, 745–746
automatic failover for, 746
definition of, 997
promoting to primary, 745
PSN (Policy Services node), 110, 633, 874–876
single-node ISE deployment, 113
two-node ISE deployment, 114–116
non-802.1X authentication
devices without a supplicant, 79–80
EasyConnect, 89–90, 993
MAB (MAC Authentication Bypass). *See* MAB (MAC Authentication Bypass)
need for, 76
remote access connections, 88–89
Web Authentication. *See* WebAuth (Web Authentication)
non-seed devices, 567–571, 996
non-tunneled EAP (Extensible Authentication Protocol), 43–44
north-south SGACL deployment, 598–599
Not MACsec capable, 617–618
notification services, 388–389
SMS gateway providers, 388–389
SMTP servers, 388
Novell Directory Services (NDS), 25
NSP (Network Services Platform)
NSP ACL (access control list), 493, 495
NSP app download, 528–529
NTP (Network Time Protocol), 32, 196, 465, 996

O

O (Organization) field, 183
objects, AD (Active Directory), 24
OCSP (Online Certificate Status Protocol), 33, 466, 519, 996
onboarding. *See* BYOD (bring your own device) onboarding; MDM (mobile device management)
one-time password (OTP), 29, 88, 996
Online Certificate Status Protocol (OCSP), 33, 466, 519, 996
open virtual appliances (OVAs), 112
Operate on non-802.1X wireless networks setting (AnyConnect posture profile), 645
Operations screen, 142, 143–150
ANC (Adaptive Network Control) component, 147–148, 991
RADIUS tab, 144–146
Reports tab, 150
TACACS Live Log tab, 147
Threat-Centric NAC Live Logs tab, 146
Troubleshoot tab, 147–148
operators
AND, 252–256
OR, 252–256
option name host-name command, 426
OR operator, 252–256
Oracle Identity Manager, 25
Organization (O) field, 183
organizational units (OUs), 25, 183
organizationally unique identifiers (OUIs), 97
OS option (Nmap), 416
OSCP (Online Certificate Status Protocol), 519
OSVersion attribute, 537
OTA (over-the-air) provisioning, 492
OTP (one-time password), 29, 88, 996
OUIs (organizationally unique identifiers), 97
OUs (organizational units), 25, 183
OVAs (open virtual appliances), 112
over-the-air (OTA) provisioning, 492
Overview menu (ISE Certificate Authority), 520
OWN_ACCOUNTS sponsor group, 382

P

PAC Files tab (Cisco AnyConnect NAM supplicant), 67–68
packages
AnyConnect headend deployment
AnyConnect configuration, building, 648–649
posture profile configuration, 644–648
uploading to ISE, 642–644
licensing, 155–158
pACLs (port-based ACLs), 725
PACs (Protected Access Credentials), 45, 546, 558–559, 997
PAN (Policy Administration node), 109, 519, 738–739, 745–748
auto PAN switchover, 745–746
automatic failover for, 746
definition of, 997
promoting to primary, 745
PAP (Password Authentication Protocol), 7, 46, 214
Parameter-Request-List attribute (DHCP), 411
Pass Live Log status, 767
PASS_ADD message, 10–11
PASS_REPL message, 11
passive identity service, 110
PassiveID Setup option (Admin portal), 138
pass-through mode, 752
Password Authentication Protocol (PAP), 46, 214
Password Change Control (Device Administration), 918
passwords
guest change password settings, 371–372
OTP (one-time password), 29, 88, 996
patch management, 685, 757–759
conditions for, 100, 673–675
remediations, 685
PCI (Payment Card Industry), 551
PEAP (Protected EAP), 44–45, 48–49, 53–54, 55–56, 108, 215
Pearson Test Prep software, 989
peering, 582
PEM (Privacy Enhanced Electronic Mail) format, 186–187, 833
Perfigo, 631
permit statement, 280, 314, 316
PERMIT_ALL_IPV4_TRAFFIC authorization rule, 241–243
Permit_HTTP_HTTPS SGACL, 601–602, 607
Permit_ICMP SGACL, 602–603
Permit_Mgmt SGACL, 601, 607
Permit_SRC_HTTP_HTTPS SGACL, 603–604
Permit_WEB_RDP SGACL, 598
per-profile CoA (Change of Authorization), 443–444
personas
Administration, 109, 737
definition of, 108–109, 997
ISE nodes and, 742
Monitoring, 109–110, 737
Policy Services, 110, 737
pxGrid, 111, 737
types of, 737
pervasive tagging, 594
phased deployment, 717–718
authentication open versus 802.1X, 719–720
closed mode, 728–730
low-impact mode, 725–727
monitor mode, 722–725
flow diagram, 723–724
operational flow, 722–723
policy set creation, 724–725
transitioning to end state, 730–731
preparation for, 720–721
transitioning to end state, 730–731
wireless networks, 731

PhoneNumber attribute, 537
physical ISE appliances, 111–113
form factors for, 111–112
scalability of, 112–113
PIN Lock option (endpoint management), 543
ping, 646
PinLockStatus attribute, 536
PKI (public key infrastructure), 519. *See also* certificate-based authentication
definition of, 463, 998
encryption with, 468–469
key pairs, 468–469, 995
prevalence of, 460
public versus private keys in, 468–469
RAs (registration authorities), 998
signatures, 31–32
smart cards, 29
Platform Exchange Grid. *See* pxGrid (Platform Exchange Grid)
Plus license packages, 156
Point-to-Point Protocol. *See* PPP (Point-to-Point Protocol)
Point-to-Point Protocol (PPP), 12
policies and policy sets. *See also* posture assessment; profiles
ANC (Adaptive Network Control), 148–149, 156, 822–823, 864–866
AUP (acceptable use policy)
hotspot guest portals, 354–356
sponsored guest portals, 386
authentication, 151
allowed protocols, 210, 213–216
for alternative ID stores based on EAP type, 224–227
authorization compared to, 209–210, 235
for certificate-based authentication, 472–474
conditions, 217–219
default, 216–217
definition of, 171
device administration, 944
for device administration AAA with Cisco IOS, 944
goals of, 206–207, 210–211
identity stores, 210–211, 219–220
identity validation, 211
for MAB (MAC Authentication Bypass), 227–228
options, 220
for remote access VPN, 223–224
restoring, 229
for wireless SSID, 220–223
authorization, 151
authentication compared to, 209–210, 235
Blackhole_Wireless_Access, 240–241
for certificate-based authentication, 474–475
Cisco_IP_Phones, 237–241
compound conditions, 239, 251–256, 992
condition blocks, 252–256
configuration of, 241–249
for CWA (Centralized Web Authentication), 322–324
default, 236–241
definition of, 171
device administration, 945–946
for device administration AAA with Cisco IOS, 945–946
goals of, 235–241
for guest portals, 348–351
for MDM (mobile device management), 536–537
organization of, 216, 236
profile assignment in, 450–453
role-specific authorization rules, 241
rule processing for, 236–241
saving conditions for reuse in, 249–251
simple conditions, 239, 251, 999

BYOD (bring your own device) onboarding
client provisioning policy, 512–514
default unavailable client provisioning policy, 515
Cisco AnyConnect NAM supplicant
Authentication Policy view, 60, 62
Client Policy view, 60, 61–62
Cisco WLC (Wireless LAN Controller), 974–979
CiscoPress SSID, 518
closed mode, 730
correlation, 845–847
CPP (client provisioning policy), 637–638
AnyConnect Secure Mobility Client, 640–649
Client Provisioning portal, 153, 166–167, 650–651
default client provisioning policy, 652
definition of, 172
for ISE for BYOD onboarding, 512–514, 515
order of operations, 637–638
rules, creating, 652–653
default, 211
device administration, 922–924, 925–927, 939–946
policy sets, 943–946
roles, 939
TACACS command sets, 922–923, 941–943
TACACS profiles, 923–924, 939–941
Endpoint Profile Policies, 431–441
endpoint purge, 345
expanding, 211
for hotspot guest portals, 362–364
low-impact mode, 727
MACsec, 616–618
monitor mode, 724–725
for self-registered guest portals, 373–380
SGTs (security group tags), 577, 612
"smart default", 318
top-level rules of, 212
TrustSec
egress policy, 597–598
policy matrix, 600–601, 604–609
WSA (Web Security Appliance), 855–857
access policy, 855–856
decryption policy, 857
Policy admin role, 128
Policy Administration node. *See* PAN (Policy Administration node)
policy configuration support bundles, 783
Policy List tab, 149
Policy page, 138, 142, 150–154
Client Provisioning tab, 153
Policy Elements tab, 154, 922–924
Policy Sets tab, 150–151
Posture tab, 152
Profiling tab, 152
Policy Service nodes (PSNs), 110, 210, 519, 633
Policy Services persona, 110, 737
Policy Sets tab, 150–151, 925–927
policy static sgt command, 578, 595, 596
portals
Administration, 137–142
global search for endpoints, 139–140
Help menu, 138, 140–141
ISE Setup Wizards, 141
Settings menu, 142
tabs, 137–139
authorization policies for, 348–351
definition of, 997
guest types, 341–342
contractors, 344–346
daily, 344–346
overview of, 343
social, 348
weekly, 347
hotspot, 341–342, 351–358

AUP (acceptable use policy) page settings, 354–356
authentication success settings, 357
authorization rule configuration, 362–365
configuration flowchart for, 351
definition of, 342
portal page customization, 358–362
portal settings, 352–354
post-access banner page settings, 355–356
support information page settings, 357–358
VLAN DHCP release page settings, 355–356
overview of, 341
self-registered, 366–367
authorization rule configuration, 373–380
BYOD settings, 372
configuration flowchart, 365–366
definition of, 342
guest change password settings, 371–372
guest device compliance settings, 373
guest device registration settings, 371–372
guest location setting, 369–371, 994
login page settings, 367–368
portal settings, 366–367
registration form settings, 368–371
self-registration success, 371
sponsored, 380–381, 385–386
AUP (acceptable use policy) page settings, 386
configuration flowchart, 380–381
default sponsor portal, 384
definition of, 342
login settings, 386
other settings, 387
portal settings, 385–386
provisioning guest accounts from, 389–394

Portals & Components tab (Guest Access work center). ***See*** **portals**

port-control command, 276

ports
802.1X default port behavior, 719
application of initial ACL to, 275–276
assigning SGTs (security group tags) to, 577–578
configuring interfaces as, 269
host mode, 272–274
pACLs (port-based ACLs), 725
ports, 273
security, 273, 997
switch
configuring interfaces as, 269
host mode, 272–274
TCP
port 49, 8
port 389, 25
port 636, 25
port 64999, 580
UDP, 32

POST method, 84–85

post-access banner page settings, 355–356

Posture Agent Redirection ACL, 283–284

posture assessment. ***See also*** **posture policy**
authorization rules, 693–694
centralization of, 629
compliance module updates, 637–638
conditions for, 99–100, 997
CPP (client provisioning policy) configuration, 637–638
AnyConnect Secure Mobility Client, 640–649
Client Provisioning portal, 650–651
default client provisioning policy, 652
order of operations, 637–638
rules, creating, 652–653

definition of, 172, 626, 997
endpoint experience, 695–705
AnyConnect already installed, endpoint not compliant, 700–702
AnyConnect not installed on endpoint yet, 696–700
redirected state, 695–696
stealth mode, 645, 703
temporal agent and posture compliant, 705
history of, 629–633
ISE posture flows, 107–110, 633–636
license packages, 156
mobile posture, 707–712
authorization conditions, 709–710
authorization rules, 710–712
supported device managers, 707–709
overview of, 99–101
Posture General Settings, 690–691
posture lease concept, 691
posture requirements, 997
Posture Troubleshooting tool, 794–795
Posture work center
Cache Last Known Posture Compliance setting, 691
Posture General Settings, 690–691
posture lease, 691
posture reassessment, 691–692
profiling versus, 402
remediation, 997

posture policy, 653–688
conditions, 654–679
anti-malware, 661–663
anti-spyware, 663
anti-virus, 663
application, 655–660
compound, 677–678
dictionary compound, 664–665
dictionary simple, 663–664
disk encryption, 665–666
file, 667–673
firewall, 660–661
hardware attributes, 655
patch management, 673–675
Registry, 675
USB, 679
configuration, 688–689
relationships between elements of, 653
remediations, 679–686
anti-malware, 681–682
anti-spyware, 681–682
anti-virus, 681–682
application, 680
definition of, 679–680
file, 682
firewall, 682
launch program, 683
link, 684
patch management, 685
USB, 686
WSUS (Windows Server Update Services), 685
requirements, 687–688

posture profiles, 644–648
Posture tab (Policy page), 152
PoV (proof of value) service, 138
PPAN (Policy Administration Node), 738–739
PPP (Point-to-Point Protocol), 12
PRA retransmission time setting, 647
Preboot Execution Environment (PXE), 725
primary devices, 738–739
primary PAN (Policy Administration node), 745
principle username X.509 attribute, 470, 997
Privacy Enhanced Electronic Mail (PEM) format, 186–187, 833
private keys, 468–469, 997
privilege levels, 932–933, 997
probe delay command, 267

probes, 409–423
Active Directory, 422
configuration, 409–411
DHCP and DHCPSPAN, 411–414
DNS, 417
HTTP, 420–421
NETFLOW, 419–420
Nmap, 415–417
publishing endpoint probe data on pxGrid, 450
pxGrid, 423
RADIUS, 414–415
SNMPQUERY and SNMPTRAP, 417–419
Process Host Lookup, 214
Profiled Cisco IP Phones rule, 237
Profiler Feed Service, 429–430
configuration, 429
verification of, 429–430
profiles, 884–886, 923–924. ***See also*** **profiling**
AMP Enabler, 642
AnyConnect NAM, 71–72
assignment in authorization policies, 450
endpoint identity groups, 450–452
EndPointPolicy, 453
BYOD (bring your own device) onboarding, 510–512, 516
CAPs (certificate authentication profiles), 23, 202, 469, 471–472, 991
Cisco WLC (Wireless LAN Controller)
Employee, 977
Helpdesk, 976
NetAdmin, 974–975
predefined TACACS profiles, 974
SecAdmin, 975–976
configuration, 320–322
device administration AAA with Cisco IOS, 932–934, 939–941
Administration profile, 940–941
HelpDesk profile, 940
Downlink MACsec, 616
Employee Full Access, 241–243
Employee_Limited, 246–249
for hotspot guest portals, 362–364
Internet_Only, 243–246
license packages, 156
logical, 441–442, 995
MDM Onboard, 539–540
native supplicant, 642
for posture assessment, 693
for self-registered guest portals, 373–380
verification of
dashboard, 454
Device Sensor show commands, 457–458
endpoints database, 455–456
Global Search tool, 454–455
WSA (Web Security Appliance), 855–857
profiling, 107, 404–406. ***See also*** **profiles**
Anomalous Behaviour Detection, 406–408
Cisco ISE probes, 409–423
Active Directory, 422
configuration, 409–411
DHCP and DHCPSPAN, 411–414
DNS, 417
HTTP, 420–421
NETFLOW, 419–420
Nmap, 415–417
publishing endpoint probe data on pxGrid, 450
pxGrid, 423
RADIUS, 414–415
SNMPQUERY and SNMPTRAP, 417–419
CoA (Change of Authorization) and, 442–444
global CoA, 442–443, 994
per-profile CoA, 443–444
custom attributes, 445–448
definition of, 172, 995
DHCP (Dynamic Host Configuration Protocol), 98

evolution of, 404–406
global settings for, 444–445
endpoint attribute filtering, 444–445
SNMP, 444
infrastructure configuration, 424–427
DHCP helper, 424
IOS Device Sensor, 426–427
SPAN (Switched Port Analyzer), 424–425
VACLs (VLAN Access Control Lists), 425–426
VMware vSwitches, 427
MAB (MAC Authentication Bypass), 80–82, 96–99
DHCP (Dynamic Host Configuration Protocol), 97
RADIUS (Remote Access Dial-In User Service), 97
posture versus, 402
Profiler Feed Service, 429–430
configuration, 429
verification of, 429–430
Profiling tab (Policy page), 152
Promote to Primary option, 745
proof of value (PoV) service, 138
Protected Access Credentials (PACs), 45, 546, 558–559, 997
Protected EAP (PEAP), 44–45, 48–49, 53–54, 55–56, 215
provisioning. *See also* CPP (client provisioning policy)
PACs (Protected Access Credentials), 558–559
supplicant provisioning reports, 534–535
PSNs (Policy Service nodes), 110, 210, 519, 633, 874–876
PSNs (Public Services Networks), 210, 519
for large deployments, 912
load balancing, 751–752
for medium deployments, 913
for small deployments, 913
public certificates, importing, 476–477
public key infrastructure. *See* PKI (public key infrastructure)
public keys, 468–469, 998
Public Services Networks. *See* PSNs (Public Services Networks)
publish and subscribe (pub/sub) communication bus, 111
publishers, 824, 998
purge policies, endpoint, 345
Push Notification (Apple), 101
PXE (Preboot Execution Environment), 725
pxGrid (Platform Exchange Grid), 169, 190
communication between participants, 826–827
components of, 824, 825
Context-In, 827, 992
Context-Out, 827, 993
definition of, 998
FMC (Firepower Management Center) configuration, 831–850
access rules, 840–844
active users, viewing, 844–845
FDM (Firepower Device Management), 832
pxGrid integration, 832–837
Rapid Threat Containment, 845–850
realms, 837–840
GCL (pxGrid common library), 825
ISE (Identity Services Engine) configuration, 828–831
license packages, 156
overview of, 824–825
password-based account creation, 837
persona, 111, 737
publishing endpoint probe data on, 450
pxGrid probe, 423
Stealthwatch, 857–866
capabilities of, 858
ISE integration, 862–866
pxGrid client identity certificate, importing, 859–862

version history, 825
WSA (Web Security Appliance) configuration, 850–857
access policy, 855–856
decryption policy, 857
ERS (External RESTful Services), 850–853
identification profiles, 855–857
integration with pxGrid and ISE, 850–855
policies, 855–857
pxGrid common library (GCL), 825
pxGrid Services tab (Administration screen), 169
pxGrid_Certificate_Template, 520

Q

Quarantine action (EPS), 821
Quick Filters counter, 769

R

RADIUS (Remote Access Dial-In User Service), 12–14
AV (attribute/value) pairs, 15
CoA (Change of Authorization)
CWA (Centralized Web Authentication) and, 311
definition of, 16, 95–96
messages, 110
definition of, 998
for device administration, 927
Drops counter, 768
fallback, 279–280
global configuration commands, 262–269
device tracking in IOS Xe 16.x and later, 267
global 802.1X commands, 266–267
IOS 12.2.x, 262–263, 264–266
IOS 15.x, 263–266
IOS XE, 263–266
local ACL (access control list) creation, 268–269
Layer 2 EAP communication with, 12–13
Live Log. *See* Live Log
Live Sessions, 145–146, 303
message types, 13–14
profiling, 97
RADIUS Authentication Troubleshooting tool, 785–786
RADIUS over Datagram Transport Layer Security (DTLS), 190
RADIUS probe, 414–415
with remote-access VPN, 88
role of, 107
server configuration
accounting servers, 278–279
authentication servers, 277–278
service-type values, 13
TACACS+ compared to, 13, 16
token servers, 23
radius server command, 263, 561
RADIUS Server Sequences tab (Network Resources), 165
RADIUS tab (Operations screen), 144–146
radius-server attribute command, 264, 265
radius-server dead-criteria time 5 tries 3 command, 264
radius-server host command, 263
radius-server load-balance command, 757
radius-server vsa send accounting command, 265, 562, 567
radius-server vsa send authentication command, 265, 562, 567
ransomware, WannaCry, 554–555, 868
Rapid Threat Containment, 821–823
ANC (Adaptive Network Control), 822–823
configuration, 845–850
definition of, 998
EPS (Endpoint Protection Services), 821–822, 993

RAs (registration authorities), 998
RBAC (role-based access control), 128, 934, 971. ***See also*** **WLC (Wireless LAN Controller)**
Read-only admin role, 129
Realm Directory configuration, 838
realms, configuration in FMC (Firepower Management Center), 837–840
reassessment, posture, 691–692
Record Selector, 770
redirected state, 695–696
redistribute static route-map STATIC-TO-EIGRP command, 756
Refresh counter, 770
Registrar, 8
registration authorities (RAs), 998
registration form, self-registered guest portals, 368–371
Registry conditions, 100, 675
RegistryKey option, 675–676
RegistryValue option, 675
RegistryValueDefault option, 675–676
Reinstate option (endpoint management), 543
REJECT message, 9
Reject option, authentication policy, 220
reload command, 946
remediation
posture, 679–686, 692, 997
anti-malware, 681–682
anti-spyware, 681–682
anti-virus, 681–682
application, 680
definition of, 679–680
file, 682
firewall, 682
launch program, 683
link, 684
patch management, 685
USB, 686
WSUS (Windows Server Update Services), 685
Rapid Threat Containment, 845–850
ANC (Adaptive Network Control), 822–823
EPS (Endpoint Protection Services), 821–822, 993
remediation timer, 645, 691
remote access connections, 88–89
Remote Access Dial-In User Service. ***See*** **RADIUS (Remote Access Dial-In User Service)**
remote access VPN (virtual private network), 223–224
Repeat counter, 767, 769
replication, 748
REPLY message, 9
reports
authentication details, 771–774
Device Administration, 927
supplicant provisioning, 534–535
TC-NAC (Threat Centric Network Access Control)
Coa-Events, 888–889
TC-NAC Live Log, 888–889
Threat-Events, 888
Vulnerability Assessment, 888
Reports tab (Operations screen), 150
repositories, Tenable.SC, 881–882
REQUEST message, 10, 11
requests for comments. ***See*** **RFCs (requests for comments)**
RESPONSE message, 10, 11
restore, 229, 759–761
Results tab (Policy Elements), 154
reuse, saving conditions for, 249–251
revoked certificates, 33
checking for, 466–467
CRLs (certificate revocation lists), 33, 466
OCSP (Online Certificate Status Protocol), 33, 466, 996
validity period, 467
RFCs (requests for comments)
RFC 5176, 106

RFC 5281, 48
RFC 6238, 29
RFC 7170, 46–47, 48
role-based access control (RBAC), 128, 934, 971. ***See also*** **WLC (Wireless LAN Controller)**
roles
Admin, 127–132
Cisco WLC (Wireless LAN Controller), 971–972
device administration AAA with Cisco IOS, 939
role-specific authorization rules, 241
route redistribution, 755–756
routed mode, 752
RSA SecurID, 23
rules, 239. ***See also*** **policies and policy sets**
authentication, 151
alternative ID stores based on EAP type, 224–227
MAB rule flowchart, 217
organization of, 216
policy sets, 211–212
remote access VPN, 223–224
wireless SSIDs (service set identifiers), 220–223
authorization, 151
AND/OR operators in, 252–256
for BYOD (bring your own device) onboarding, 517, 518
for device administration AAA with Cisco WLC, 977–979
Employee and CorpMachine, 242–243
employee full access, 241–243
employee limited access, 246–249
for guest portals, 348–358
Internet-only access, 243–246
IT Users Access, 252–256
MDM On-boarding, 539–542
mobile posture, 710–712
PERMIT_ALL_IPV4_TRAFFIC, 241–243
for posture assessment, 693–694
role-specific, 241
for self-registered guest portals, 373–380
TC-NAC (Threat Centric Network Access Control), 884–886
Wireless Black List Default, 239
correlation, 845–847
CPP (client provisioning policy), 652–653
CWA (Centralized Web Authentication)
custom authorization rules, 323–324
Guest Flow attribute, 323–324, 994
preconfigured authorization rules, 322
Profiled Cisco IP Phones, 237
Run SMB Discovery Script option (Nmap), 416
Russell, Paul, 337

S

Sales SGT (Security Group Tag), 555
SAML (Security Assertion Markup Language)
assertions, 395, 998
guest portal logins, 368, 394–400
IdPs (identity providers), 35, 394–400, 998
SPs (service providers), 394, 998
support for, 23
SANs (subject alternative names), 184–185, 752
SAP (Security Association Protocol), 616
sap mode-list no-encap command, 566, 568
sap pmk 26 mode-list gcm-encrypt command, 621
sap pmk pairwise-master-key mode-list gcm-encrypt command, 621
scalability
ISE (Identity Services Engine), 118–119, 737–742
ISE cubes, 737
ISE persona types, 737

node personas, 742
PPAN (Policy Administration Node), 738–739
primary devices, 738–739
registration of ISE nodes, 739–742
SNS (Secure Network Server) appliances, 112–113
scans, Tenable.SC, 882–883
SCCM (System Center Configuration Manager), 673, 708
SCEP (Simple Certificate Enrollment Protocol), 500, 520–521, 999
SD-Access (Software-Defined Access), 613–614
/sdcards/downloads/spw.log command, 812
Search icon (Admin portal), 138
sec-admin group, 945
SecAdmin profile, 975–976
Second Port Disconnect, CDP, 991
secondary PAN (Policy Administration node)
auto PAN switchover, 745–746
automatic failover for, 746
promoting to primary, 745
Secure Network Server. *See* SNS (Secure Network Server) appliances
Secure Shell. *See* SSH (Secure Shell)
Secure Sockets Layer. *See* SSL (Secure Sockets Layers)
secure web gateways, 818
Security Administrator command set, 942–943
Security Administrators role, 939
Security Assertion Markup Language. *See* SAML (Security Assertion Markup Language)
Security Association Protocol (SAP), 616
security context, 232, 235
Security Group Access. *See* TrustSec
security group ACLs. *See* SGACLs (security group ACLs)
security group firewalls. *See* SGFWs (security group firewalls)
security group tags. *See* SGTs (security group tags)
security holes. *See* vulnerability assessment
security information and event management (SIEM) systems, 818
security information management (SIM), 744
Security Level tab (Cisco AnyConnect NAM supplicant), 64–66
security posturing, 101
SECURITY role, 971
seed devices, 566, 999
self-management, MDM (mobile device management) onboarding, 543–544
self-registered guest portals, 342, 515
authorization rule configuration, 373–380
BYOD settings, 372
configuration flowchart, 365–366
guest change password settings, 371–372
guest device compliance settings, 373
guest device registration settings, 371–372
guest location setting, 369–371, 994
login page settings, 367–368
portal settings, 366–367
registration form settings, 368–371
self-registration success, 371
Self-Registration Success page, 371
self-signed certificates, 181–182
serial replication, 748
SerialNumber attribute, 537
Server column (Live Log), 768
Server name rules setting (AnyConnect posture profile), 647
servers, 276–280
authentication, 41, 277–278, 991
Cisco Access Control Server (ACS), 909, 910
HTTP/HTTPS, 314
RADIUS
accounting servers, 278–279
authentication servers, 277–278

fallback, 279–280
token servers, 23
SMTP (Simple Mail Transfer Protocol), 388
SNS (Secure Network Server) appliances, 177
Sun Directory Server, 25
XCP (Extensible Communication Platform), 825
service providers, 35
service set identifiers. *See* SSIDs (service set identifiers)
service-level agreements (SLAs), 754–756
Service-Type attribute (RADIUS), 96, 265, 415
service-type values, 13
session IDs, 97
Session Key Assignment, 918
Session Trace tool, 801–804
Settings tab, 161, 163
setup command, 178
SGA (Security Group Access). *See* TrustSec
SGACLs (security group ACLs), 597–604
definition of, 998
Deny_All SGACL, 601–602
east-west deployment of, 598–599
egress policy, 597–598, 600–601
north-south deployment of, 598–599
Permit_HTTP_HTTPS SGACL, 601–602
Permit_ICMP SGACL, 602–603
Permit_Mgmt SGACL, 601
Permit_SRC_HTTP_HTTPS SGACL, 603–604
Permit_WEB_RDP SGACL contents, 598
syntax, 599–600
SGFWs (security group firewalls), 611–613
on ASA (Adaptive Security Appliances), 612
on ASR (Aggregation Services Router), 613
definition of, 999
on Firepower, 612–613
on ISR (Integrated Services Router), 613
SGTs (security group tags), 13, 236, 822
access-layer devices that do not support, 580
classification of, 575–577
configuration, 572–574
definition of, 172, 556
dynamically assigning via 802.1X, 577
egress enforcement with, 555–556
manually assigning via 802.1X, 577–578
mapping subnets to, 580
mapping VLANs to, 580
native tagging, 593–597
configuration on IOS XE switches, 595–597
Layer 2 frame format with, 593–594
pervasive tagging, 594
SXP (SGT Exchange Protocol), 110, 303, 581–593
configuration on Cisco ASA, 591–592
configuration on IOS devices, 588–590
configuration on ISE (Identity Services Engine), 584–586
configuration on WLCs (wireless LAN controllers), 590–591
definition of, 998
design, 582–583
verification in ASDM (Adaptive Security Device Manager), 592–593
SHA-256 file type, 672
Shadow Brokers, 868
sho cts interface command, 596
sho run int command, 621
Short Message Service. *See* SMS (Short Message Service)
should-secure policy (MACsec), 616, 618
show commands, 147, 457–458
show aaa server | incl host, 757
show aaa servers, 296–297
show application status ise, 741–742

show authentication interface, 812–813
show authentication session int g0/23, 696
show authentication session interface, 297–299, 328–329
show cts environment-data, 562, 568
show cts interface, 571, 622–623
show cts pac, 568
show cts policy peer, 570
show cts rbacl, 609
show cts role-based permissions, 562, 609
show cts sxp connections brief, 589, 591
show cts sxp sgt-map brief, 590
show device-sensor cache all, 457–458
show device-tracking database details, 267
show ip access-list ACL-WEBAUTH-REDIRECT, 314
show ip device tracking all, 267
show logging application, 781–782
show logging system, 781–782
show run, 11, 620
show running-config, 952, 961–963
show running-config aaa, 569
show running-config all | inc ip, 266
show running-config all | inc radius-server, 265
show udi, 747–748
show users, 955

Shutdown action (EPS), 821

SIEM (security information and event management) systems,818

signatures
AnyConnect posture profile, 645
CAs (certificate authorities), 31–32

SIM (security information management), 744

Simple Certificate Enrollment Protocol (SCEP), 500, 520–521, 999

simple conditions, 239, 251, 663–664, 999

Simple Mail Transfer Protocol. *See* SMTP (Simple Mail Transfer Protocol) servers

Simple Network Management Protocol. *See* SNMP (Simple Network Management Protocol)

single sign-on. *See* SSL (Secure Sockets Layers)

single SSID onboarding, 487–488
Android onboarding flow, 526–530
device provisioning, 529–530
device registration, 526–528
NSP app download, 528–529
definition of, 999
iOS onboarding flow, 523–526
device enrollment, 523–524
device provisioning, 526–527
device registration, 523–524
ISE configuration, 495–496, 510–523
Apple iOS example, 496–503
authorization profiles, 516
authorization rules for EAP-TLS authentications, 518
authorization rules for onboarding, 517
Blackberry example, 508–509
client provisioning policy configuration, 512–514
default unavailable client provisioning policy action, 515
ISE as certificate authority, 519–520, 521–523, 994
native supplicant profile, 510–512
SCEP (Simple Certificate Enrollment Protocol), 520–521, 999
WebAuth configuration, 514–515
verification of BYOD flows, 534–535
endpoint identity groups database, 535
RADIUS Live Logs, 534
reports, 534–535
Windows and Mac onboarding flow, 531–533
device provisioning, 532–533
device registration, 531
WLC (Wireless LAN Controller) configuration, 489–495

ACLs (access control lists), 492–495
WLAN configuration, 490–491
Single-Host mode, 272, 298, 618
single-node ISE deployment, 113
SISE 300–715 exam preparation, 988–989
final study and review, 988–989
hands-on activities, 988–989
Skip NMAP Host Discovery option (Nmap), 416
SLAs (service-level agreements), 754–756
smart cards, 29
smart default policies, 318
smart devices
employee limited access, 246–249
Internet-only access, 243–246
SmartDevice logical profile, 245–246
SMS (Short Message Service), 388–389
SMTP (Simple Mail Transfer Protocol) servers, 388
SNAT (Source NAT), 752
SNMP (Simple Network Management Protocol)
global probe settings, 444
SNMP Ports option (Nmap), 416
SNMPQUERY and SNMPTRAP probes, 417–419
snmp trap mac-notification change added command, 418
snmp trap mac-notification change removed command, 418
SNMPQUERY probe, 417–419
snmp-server community command, 419
snmp-server source-interface informs command, 266
snmp-server trap-source command, 266
SNMPTRAP probe, 417–419
snooping, DNS, 494
SNS (Secure Network Server) appliances, 177
form factors, 111–112
scalability, 112–113
social guest accounts, 348
social login, 23, 35
Softerra LDAP browser, 26
Software-Defined Access (SD-Access), 613–614
Source NAT (SNAT), 752
source SGT, 559
source tree view (TrustSec policy matrix), 552
sources, identity, 34, 35
SPAN (Switched Port Analyzer)
configuration, 424–425
DHCPSPAN probe, 411–414
HTTP SPAN design, 420–421
Sponsor Groups settings
contractor accounts, 346
weekly guest accounts, 348
sponsored guest portals, 342, 380–381
AUP (acceptable use policy) page settings, 386
configuration flowchart, 380–381
default sponsor portal, 384
login settings, 386
other settings, 387
portal settings, 385–386
provisioning guest accounts from, 389–394
sponsors
definition of, 341, 381, 999
sponsor groups, 381–382
SPs (service providers), 35, 394, 998
spyware
anti-spyware posture policy conditions, 663
anti-spyware remediations, 681–682
SSH (Secure Shell), 6
SSIDs (service set identifiers), 220–223, 487–488, 730–731, 999. ***See also*** **dual SSID onboarding; single SSID onboarding**
SSL (Secure Sockets Layers), 25, 181
SSO (single sign-on), 35

ST (State) field, 184
START message, 9, 11
statements
deny, 280, 316
permit, 280, 314, 316
status, Live Log, 767
Status dashlet, 134
stealth mode, 645, 703
Stealthwatch
capabilities of, 858
configuration, 857–866
ISE integration, 862–866
pxGrid client identity certificate, importing, 859–862
STIX (Structured Threat Information eXpression), 890, 892–893, 999
STOP message, 11
subject alternative names (SANs), 184–185, 752
subnets, mapping to SGTs (security group tags), 580
subordinate CAs (certificate authorities), 521–523
subscribers, 824, 999
SUCCESS message, 11
Summary tab (Home page), 134
Sun Directory Server, 25
Super admin role, 130
supplicants
authenticators, 719
Cisco AnyConnect NAM, 59–73
AnyConnect NAM profiles, 71–72
Authentication Policy view, 60, 62
Client Policy view, 60, 61–62
EAP chaining, 73
Network Groups view, 60, 71
Networks view, 60, 62–70
overview of, 59–60
definition of, 41, 50, 999
devices without, 79–80
endpoint supplicant verification, 295–296
native supplicant profile, 510–512
policy for, 617
supplicant provisioning reports, 534–535
Windows native, 50–59
machine authentication, 58–59
user authentication, 58
Wired AutoConfig service, 50–57
support bundles, 782–784
categories of, 782–783
creating from CLI, 783–784
creating from GUI, 783
definition of, 782
support information page settings, hotspot guest portals, 357–358
Switched Port Analyzer. *See* SPAN (Switched Port Analyzer)
switches. *See also* ports; WLC (Wireless LAN Controller)
authentication on, 261–276
AAA servers, 276–280
Cisco ISE verification, 302–303
endpoint supplicant verification, 295–296
global configuration AAA commands, 261–262
global configuration RADIUS commands, 262–269
interface configuration settings, 269–276
NAC (network access device) verification, 296–302
IOS XE
configuration for TrustSec, 560–563
manual SGT (security group tag) propagation on, 595–597
MACsec, 618–619
sample configurations, 1034–1061
Catalyst 3000 Series, 12.2(55)SE, 1034–1038
Catalyst 3000 Series, 15.0(2)SE, 1038–1044
Catalyst 4500 Series, IOS-XE 3.3.0 / 15.1(1)SG, 1053–1057

Catalyst 6500 Series, 12.2(33)SXJ, 1058–1061
Catalyst 9000 Series, 16.9.5, 1044–1052
verifying authentications with, 296–299
show aaa servers command, 296–297
show authentication session interface command, 297–299
test aaa command, 297
WebAuth, 313–315
certificates, 313
HTTP/HTTPS server, 314
URL-redirect ACL, 314–315
switchport command, 269
SXP (SGT Exchange Protocol), 110, 303, 581–593
configuration on Cisco ASA, 591–592
configuration on IOS devices, 588–590
configuration on ISE (Identity Services Engine), 584–586
configuration on WLCs (wireless LAN controllers), 590–591
definition of, 998
design, 582–583
overview of, 581–582
verification in ASDM (Adaptive Security Device Manager), 592–593
sysContact option (Nmap), 416
sysDescr option (Nmap), 416
sysLocation option (Nmap), 416
syslog messages, 299–300
sysName option (Nmap), 416
System Activities option (Admin portal), 139
System admin role, 130
System Center Configuration Manager (SCCM), 673, 708
System Certificates settings, hotspot guest portals, 353
System logs support bundles, 783
System Scan module. *See* posture assessment
System Summary dashlet, 134
System tab (Administration screen), 155–161
Admin Access tab, 160–161
Backup & Restore tab, 160
Certificates tab, 158
Deployment tab, 155
Licensing tab, 155–158
Logging tab, 159
Maintenance tab, 159
Settings tab, 161
Upgrade tab, 160
SYSTEM_32 file path option, 669
SYSTEM_DRIVE file path option, 669
SYSTEM_PROGRAMS file path options, 669
SYSTEM_ROOT file path option, 669

T

TACACS Live Log tab (Operations screen), 147
TACACS+ (Terminal Access Controller Access Control System)
accounting messages, 11
admin role, 131
authentication messages, 9
authorization messages, 10–11
Cisco WLC (Wireless LAN Controller) profiles
Employee, 977
Helpdesk, 976
ISE configuration on WLC TACACS+ servers, 979–980
NetAdmin, 974–975
predefined, 974
SecAdmin, 975–976
client/server communication, 7–8
command sets, 922–923
communication flows, 12
device administration AAA with Cisco IOS
policy sets, 943–946

TACACS command sets, 941–943
TACACS profiles, 932–934, 939–941
TACACS+ authentication and fallback, 946–948
TACACS+ command accounting, 951
TACACS+ command authorization, 948–950
TACACS+ command sets, 934–936, 992
enabling, 910–911, 914–915
Live Log. *See* Live Log
profiles, 923–924
RADIUS compared to, 13, 16
support for, 8
tags, security group. *See* SGTs (security group tags)
targets, logging, 743–744, 777–778
TAXII (Trusted Automated eXchange of Intelligence Information), 890, 892–893, 999
TCAM (Ternary CAM), 554
TC-NAC (Threat Centric Network Access Control), 110, 886
AMP (Advanced Malware Protection) for Endpoints, 897–904
adapter configuration, 900–904
capabilities of, 897–898
incidents, 899–900
indicators, 899
authorization profiles/rules, 884–886
capabilities of, 871–873
Coa-Events report, 888–889
CTA (Cognitive Threat Analytics), 890–897
authorization with, 896–897
CTA STIX/TAXII API account creation, 892–893
dashboard, 890–892
integration for TC-NAC, 894–896
CVE (Common Vulnerabilities and Exposures), 873, 992
CVSS (Common Vulnerability Scoring System), 873, 992
enabling, 874–878
endpoint transition on network with TC-NAC, 887
exploits, definition of, 872, 993
flows for, 873–874
goals of, 868
integration with threat sources, 890
integration with vulnerability assessment vendor, 878–883
advanced settings, 881
basic setup, 880
configured vendor instances, 883
Tenable.SC repositories, 881–882
Tenable.SC scans, 882–883
users, 880
license packages, 156
TC-NAC Live Log, 888–889
Threat-Events report, 888
Vulnerability Assessment report, 888
vulnerability-based access control, 873–874
TCP (Transmission Control Protocol)
port 49, 8
port 389, 25
port 636, 25
port 64999, 580
TCP/7800, 748
TCP/7802, 748
TCPDump, 327, 798–801
TEAP (Tunnel EAP), 46–47, 48–49, 73, 216
temporal agents, 99–100, 999
Tenable.SC repositories, 881–882
Tenable.SC scans, 882–883
Terminal Access Controller Access Control System. *See* TACACS+ (Terminal Access Controller Access Control System)
Ternary CAM (TCAM), 554
test aaa command, 297
testing
device administration AAA with Cisco IOS
at IOS command line, 954–966

in ISE (Identity Services Engine), 952–954
device administration AAA with Cisco WLC, 981–986
themes, hotspot guest portals, 361
Threat Centric NAC reports, 150
Threat Centric NAC tab (Administration screen), 170
Threat Centric Network Access Control. ***See*** **TC-NAC (Threat Centric Network Access Control)**
threat sources, TC-NAC integration with, 890
Threat tab (Home page), 137
Threat Watchlist dashlet, 137
Threat-Centric NAC Live Logs tab (Operations screen), 146
Threat-Events report, 888
TIM (IBM Tivoli Identity Manager), 25
Time Selector, 770
Timeout option (TACACS profile), 933
timers
authentication, 275
remediation, 645, 691
Tivoli Identity Manager (TIM), 25
TLS (Transport Layer Security), 43, 45, 181–182, 214, 215, 235, 470, 737
token servers, 23
Top Threats dashlet, 137
Top Vulnerability dashlet, 136
topics (pxGrid), 824, 999
Total Compromised Endpoints dashlet, 137
Total Vulnerable Endpoints dashlet, 136
Touch ID, 310
tracking enable command, 267
Transmission Control Protocol. ***See*** **TCP (Transmission Control Protocol)**
transport
native tagging, 593–597
SXP (SGT Exchange Protocol), 303, 581–593
configuration on Cisco ASA, 591–592
configuration on IOS devices, 588–590
configuration on ISE (Identity Services Engine), 584–586
configuration on WLCs (wireless LAN controllers), 590–591
design, 582–583
overview of, 581–582
verification in ASDM (Adaptive Security Device Manager), 592–593
Transport Layer Security. ***See*** **TLS (Transport Layer Security)**
Transportation Security Administration (TSA), 2
Troubleshoot tab (Operations screen), 147–148
"Troubleshooting Cisco's ISE without TAC" (Woland), 815
troubleshooting tools
for device administration, 981–986
at IOS command line, 954–966
in ISE (Identity Services Engine), 952–954
Endpoint Debug, 796–798
endpoint diagnostics, 809–812
Cisco AnyConnect Diagnostics and Reporting Tool (DART), 59, 809–811
supplicant provisioning logs, 812
Evaluate Configuration Validator, 788–793
Execute Network Device Command, 787–789
logs, 766–785
debug logs, 779–784
de-duplication of, 805–807
Live Log, 327, 766–775
Live Sessions, 776
logging categories, 744, 778–779
logging flows, 743
logging targets, 743–744, 777–778

network device troubleshooting, 812–815
client details, viewing on WLC, 813–814
debug commands, 815
show authentication interface command, 812–813
Posture Troubleshooting, 794–795
RADIUS Authentication Troubleshooting tool, 785–786
Session Trace, 801–804
support bundles, 782–784
categories of, 782–783
creating from CLI, 783–784
creating from GUI, 783
definition of, 782
TCP Dump, 798–801
troubleshooting methodology, 804–808
authentication and authorization flows, 804–805
log de-duplication, 805–807
USERNAME user, 807
trusted authorities, 999
Trusted Automated eXchange of Intelligence Information (TAXII), 890, 892–893, 999
trusted CAs (certificate authorities), 31–32, 475–479
certificate status verification, 478–479
public certificates, 476–477
role in authentication process, 463–465
trusted certificates, 188, 537–538
Trusted for Client Auth field, 463
TrustSec
architecture of, 557–558
Cisco Software-Defined Access (SD-Access), 613–614
configuration
ASA (Adaptive Security Appliances), 564–565
IOS XE switches, 560–563
NAD (network access devices) settings, 559
PACs (Protected Access Credentials), 558–559
definition of, 172, 548, 555–556
domains, 557, 1000
environment data, 558, 993
goals of, 555
license packages, 156
native tagging, 593–597
configuration on IOS XE switches, 595–597
Layer 2 frame format with, 593–594
pervasive tagging, 594
NDAC (Network Device Admission Control), 566, 996
non-seed devices, 567–571, 996
seed devices, 566, 999
policy matrix, 604–609
configuration of, 605–609
views of, 604–605
reports, 150
SGACLs (security group ACLs), 597–604, 998
Deny_All SGACL, 601–602
east-west deployment of, 598–599
egress policy, 597–598, 600–601
north-south deployment of, 598–599
Permit_HTTP_HTTPS SGACL, 601–602
Permit_ICMP SGACL, 602–603
Permit_Mgmt SGACL, 601
Permit_SRC_HTTP_HTTPS SGACL, 603–604
Permit_WEB_RDP SGACL contents, 598
syntax, 599–600
SGFWs (security group firewalls), 611–613
on ASA (Adaptive Security Appliances), 612
on ASR (Aggregation Services Router), 613
definition of, 999
on Firepower, 612–613

on ISR (Integrated Services Router), 613

SGTs (security group tags)

access-layer devices that do not support, 580

classification of, 575–577

defining, 572–574

definition of, 556

dynamically assigning via 802.1X, 577

egress enforcement with, 555–556

manually assigning via 802.1X, 577–578

mapping subnets to, 580

mapping VLANs to, 580

native tagging, 593–597

SXP (SGT Exchange Protocol), 581–593

TrustSec-enabled NADs (network access devices)

defining TrustSec settings for, 559

PACs (Protected Access Credentials), 558–559

TSA (Transportation Security Administration), 2

TTLS (Tunneled Transport Layer Security)

EAP-TTLS, 45–46, 48–49, 216

TEAP, 216

Tunnel EAP (Tunnel EAP), 46–47, 48–49, 73, 216

tunneled EAP (Extensible Authentication Protocol) types, 44–49, 214–216, 1000

EAP chaining, 216

EAP-FAS, 215

EAP-FAST, 48–49

EAP-GTC, 215

EAP-MS-CHAPv2, 215

EAP-TLS, 215

EAP-TTLS, 48–49, 216

PEAP, 44–45, 48–49, 215

TEAP, 46–47, 73, 216

Tunneled Transport Layer Security (TTLS), 45–46, 48–49, 216

two-factor authentication (2FA), 26, 1000

two-node ISE deployment, 114–116

U

UDID attribute, 537

UDP (User Datagram Protocol) port 123, 32

Umbrella, 640

Unquarantine action (EPS), 821

unsupported devices, onboarding, 508–509

updates

to Cisco Identity Services Engine (SISE 300–715) Exam, 1034–1061

Catalyst 3000 Series, 12.2(55)SE, 1034–1038

Catalyst 3000 Series, 15.0(2)SE, 1038–1044

Catalyst 4500 Series, IOS-XE 3.3.0 / 15.1(1)SG, 1053–1057

Catalyst 6500 Series, 12.2(33)SXJ, 1058–1061

Catalyst 9000 Series, 16.9.5, 1044–1052

to compliance modules, 637–638

Upgrade tab (System), 160

upgrades, 160

Uplink MACsec, 619–623

uploading AnyConnect deployment packages, 642–644

URL policy, 612

URL redirection, 236

URL-redirect ACL

configuration for Cisco switch, 314–315

configuration for WLC, 316

URL-redirected MAC authentication bypass, 311–313

USB condition, 100, 679

USB remediations, 686

user agents (SAML), 35

User Auth tab (Cisco AnyConnect NAM supplicant), 68–69
user authentication. *See* authentication
User Datagram Protocol (UDP) port 123, 32
user identities, 920, 921–922
user identity groups, 22
definition of, 1000
FDM (Firepower Device Management), 840–844
Guest Access work center, 339–340
user interface, Device Administration work center, 919–920
user persistence, 752
USER_DESKTOP file path option, 669
USER_PROFILE file path option, 670
user-experience (UX), 337. *See also* guest services
username command, 263
username cts-user privilege command, 561
UserNotified attribute, 537
users, 807
internal, 21–22, 994
local, 195–196
network access, 21
TC-NAC (Threat Centric Network Access Control), 880
viewing in FMC (Firepower Management Center), 844–845
Users tab (Context Visibility screen), 143
UX (user-experience), 337. *See also* guest services

V

VACLs (VLAN Access Control Lists), 424, 425–426
Validate Server Certificate checkbox (Windows native supplicant), 53
validity period, for certificate, 467
vendors
MDM (mobile device management), 101–102
TC-NAC integration with, 878–883
advanced settings, 881
basic setup, 880
configured vendor instances, 883
Tenable.SC repositories, 881–882
Tenable.SC scans, 882–883
users, 880
VSAs (vendor-specific attributes), 265–266, 1000
virtual ISE appliances, 111–113, 177
virtual private networks. *See* VPNs (virtual private networks)
viruses
anti-virus posture policy conditions, 663
anti-virus remediations, 681–682
Visibility Setup option (Admin portal), 138
vlan access-map command, 426
VLAN detection interval setting (AnyConnect posture profile), 646
VLAN DHCP release page settings, hotspot guest portals, 355–356
VLANs (virtual LANs)
assignment of, 551–553, 726
authentication VLAN, 87–88
dynamic interfaces for, 284–286
employee dynamic interface, 284–285
guest dynamic interface, 285–286
mapping to SGTs (security group tags), 580
segmentation with, 322
VACLs (VLAN Access Control Lists), 424, 425–426
VMware
ISE support for, 113
vSwitches, 427
Workspace One, 708
VPNs (virtual private networks), 7
provisioning, 101
remote access, 88, 223–224
VRF instances, 548
VSAs (vendor-specific attributes), 265–266, 1000

vSwitches, 427
vulnerabilities, definition of, 872, 1000
vulnerability assessment
exploits
CVE (Common Vulnerabilities and Exposures), 873, 992
CVSS (Common Vulnerability Scoring System), 873, 992
definition of, 872, 993
TC-NAC (Threat Centric Network Access Control)
capabilities of, 871–873
Coa-Events report, 888–889
enabling, 874–878
endpoint transition on network with TC-NAC, 887
flows for, 873–874
integration with threat sources, 890
integration with vulnerability assessment vendor, 878–883
TC-NAC Live Log, 888–889
Threat-Events report, 888
Vulnerability Assessment report, 888
vulnerability-based access control, 873–874
Vulnerability Assessment report, 888
Vulnerability tab (Home page), 136
Vulnerability Watchlist dashlet, 136
Vulnerable Endpoints Over Time dashlet, 136
Vulnerable-LimitAccess authorization rule, 886

W

WannaCry ransomware, 554–555
Web Authentication. *See* WebAuth (Web Authentication)
Web Authentication Redirection ACL, 280–282
WebAuth (Web Authentication)
configuration for BYOD onboarding, 514–515
CWA (Centralized Web Authentication), 730–731. *See also* sponsored guest portals
authentication process, 85–87
authorization policies, 322–324
Cisco switch configuration, 313–315
CoA (Change of Authorization) and, 311
definition of, 991
dual SSID onboarding and, 496
ISE (Identity Services Engine) configuration, 317–322
services supported by, 311
with third-party network device support, 87–88
URL-redirected MAC authentication bypass, 311–313
verification from client, 324–326
verification from ISE UI, 327
verification on NAD (network access device), 327–331
WLC (Wireless LAN Controller) configuration, 98, 315–316, 329–331
definition of, 306, 1000
guests, 309
LWA (Local Web Authentication), 310–311
with centralized portal, 84–85
definition of, 995
when to use, 84
overview of, 83
scenarios for, 309–310
WebAuthN, 310
weekly guest accounts, 347
WEP (Wired Equivalency Protection), 614
Wi-Fi Protected Access (WPA/WPA2), 615
wildcard certificates, 184

Windows BYOD (bring your own device) onboarding, 531–533
device provisioning, 532–533
device registration, 531
Windows Hello, 27, 310
Windows native supplicant, 50–59
machine authentication, 58–59
user authentication, 58
Wired AutoConfig service
Advanced Settings tab, 57
Authentication tab, 53, 56–57
EAP MSCHAPv2 Properties dialog, 54
local area connection properties, 52–53
Protected EAP Properties page, 53–54, 55–56
Windows services control applet, 51–52
Windows Server Update Services (WSUS) remediations, 685
Windows Update Agent, 673
WinSPWizard, 513
wired authentication, 261–276
Cisco ISE verification, 302–303
Live Sessions, 303
RADIUS Live Log, 302–303
endpoint supplicant verification, 295–296
global configuration AAA commands, 261–262
global configuration RADIUS commands, 262–269
device tracking in IOS Xe 16.x and later, 267
global 802.1X commands, 266–267
IOS 12.2.x, 262–263, 264–266
IOS 15.x, 263–266
IOS XE, 263–266
local ACL (access control list) creation, 268–269
interface configuration settings, 269–276
application of initial ACL to port, 275–276
authentication settings, 274–275
authentication timers, 275
configuration of interfaces as switch ports, 269
FlexAuth (Flexible Authentication), 269–272
HA (high availability), 269–272
host mode of switch port, 272–274
NAC (network access device) verification, 296–302
with Cisco switches, 296–299
with Cisco WLC (Wireless LAN Controller), 300–302
with syslog messages, 299–300
Wired AutoConfig service
Advanced Settings tab, 57
Authentication tab, 53, 56–57
EAP MSCHAPv2 Properties dialog, 54
local area connection properties, 52–53
Protected EAP Properties page, 53–54, 55–56
Windows services control applet, 51–52
Wired Equivalency Protection (WEP), 614
Wired_802.1X condition, 242, 254
Wired_MAB condition, 242
wireless authentication, 731
airespace ACLs (access control lists), 280–284
Google URLs for ACL Bypass, 282–283
Posture Agent Redirection ACL, 283–284
Web Authentication Redirection ACL, 280–282
RADIUS, 276–280
accounting servers, 278–279
authentication servers, 277–278
fallback, 279–280
Wireless Black List Default rule, 239

Wireless LAN Controller. *See* WLC (Wireless LAN Controller)
wireless LANs. *See* WLAN (wireless LAN) configuration
wireless MAB, 489
WIRELESS role, 971
Wireless Setup (BETA) option (Admin portal), 139
wireless SSIDs, authentication policy for, 220–223
AD identity source, 223
allowed protocols, 221–222
completed authentication rule, 223
SSID name, 221
wizards, 179
Cisco Supplicant Provisioning Wizard, 513
CiscoPress-TLS, 513
CLI Setup Wizard, 178–179
ISE Setup Wizards, 141
WinSPWizard, 513
WLAN (wireless LAN) configuration, 286–295, 971
for BYOD onboarding, 490–491
corporate WLAN, 291–295
guest WLAN, 287–290
WLC (Wireless LAN Controller), 329–331
authentication configuration on
AAA servers, 276–280
airespace ACLs (access control lists), 280–284
Cisco ISE verification, 302–303
dynamic interfaces for client VLANs, 284–286
endpoint supplicant verification, 295–296
NAC (network access device) verification, 296–302
wireless LANs, 286–295
configuration, 315–316, 329–331
for BYOD onboarding, 489–495
ISE NAC feature, 315–316
MAC Filtering option, 315
URL-redirect ACL, 316
device administration AAA configuration
ISE configuration on WLC TACACS+ servers, 979–980
network device preparation, 972
policy results preparation, 974–977
policy sets, 977–979
testing and troubleshooting, 981–986
top-level menus, 971–972
Device Sensor feature, 98
DHCP probes with, 413
HTTP POST method, 84–85
SXP (SGT Exchange Protocol) configuration on, 590–591
verifying authentications with, 300–302
current clients, 300–302
debug clients, 302
viewing client details on, 813–814
Woland, Aaron, 482, 633
Work Centers, Device Administration, 918
Identities tab, 920
Network Resources tab, 921–922
Password Change Control settings, 918
Policy Elements tab, 922–924
Policy Sets tab, 925–927
Reports, 927
Session Key Assignment settings, 918
UI navigation for, 919–920
Work Centers screen, 142, 170–171
Workspace One, 708
WPA (Wi-Fi Protected Access), 548, 615
WSA (Web Security Appliance) configuration, 850–857
access policy, 855–856
decryption policy, 857
ERS (External RESTful Services), 850–853

identification profiles, 855–857
integration with pxGrid and ISE, 850–855
policies, 855–857
WSUS (Windows Server Update Services) remediations, 685

X

X.509 certificates. *See* certificate-based authentication
XCP (Extensible Communication Platform) server, 825
XenMobile, 708
XMPP (Extensible Messaging and Presence Protocol), 825

Y-Z

YubiKeys, 310
ZBF (zone-based firewall), 611

This page intentionally left blank

Exclusive Offer ☐40% OFF

Cisco Press Video Training

livelessons

ciscopress.com/video

Use coupon code **CPVIDEO40** during checkout.



Video Instruction from Technology Experts



Advance Your Skills

Get started with fundamentals, become an expert, or get certified.



Train Anywhere

Train anywhere, at your own pace, on any device.



Learn

Learn from trusted author trainers published by Cisco Press.

Try Our Popular Video Training for FREE!

ciscopress.com/video

Explore hundreds of **FREE** video lessons from our growing library of Complete Video Courses, LiveLessons, networking talks, and workshops.

Cisco Press ciscopress.com/video

ALWAYS LEARNING PEARSON



Photo courtesy of Cisco

Register Your Product at ciscopress.com/register

Access additional benefits and **save 35%** on your next purchase

- Automatically receive a coupon for 35% off your next purchase, valid for 30 days. Look for your code in your Cisco Press cart or the Manage Codes section of your account page.
- Download available product updates.
- Access bonus material if available.*
- Check the box to hear from us and receive exclusive offers on new editions and related products.

**Registration benefits vary by product. Benefits will be listed on your account page under Registered Products.*

Learning Solutions for Self-Paced Study, Enterprise, and the Classroom

Cisco Press is the Cisco Systems authorized book publisher of Cisco networking technology, Cisco certification self-study, and Cisco Networking Academy Program materials.

At ciscopress.com, you can:

- Shop our books, eBooks, practice tests, software, and video courses
- Sign up to receive special offers
- Access thousands of free chapters and video lessons

Visit ciscopress.com/community to connect with Cisco Press

Pearson

Addison-Wesley • Adobe Press • Cisco Press • Microsoft Press • Pearson IT Certification • Que • Sams • Peachpit Press

Appendix E

Study Planner

Practice Test	Reading	Task

Element	Task	Goal Date	First Date Completed	Second Date Completed (Optional)	Notes
Introduction	Read Introduction				
1. Fundamentals of AAA	Read Foundation Topics				
1. Fundamentals of AAA	Review Key Topics				
1. Fundamentals of AAA	Define Key Terms				
1. Fundamentals of AAA	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 1 in practice test software				
2. Identity Management	Read Foundation Topics				
2. Identity Management	Review Key Topics				
2. Identity Management	Define Key Terms				
2. Identity Management	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 2 in practice test software				
3. Extensible Authentication Protocol (EAP) over LAN: 802.1X	Read Foundation Topics				
3. Extensible Authentication Protocol (EAP) over LAN: 802.1X	Review Key Topics				
3. Extensible Authentication Protocol (EAP) over LAN: 802.1X	Define Key Terms				
3. Extensible Authentication Protocol (EAP) over LAN: 802.1X	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 3 in practice test software				
4.Non-802.1X Authentications	Read Foundation Topics				
4.Non-802.1X Authentications	Review Key Topics				

4.Non-802.1X Authentications	Define Key Terms				
4.Non-802.1X Authentications	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 4 in practice test software				
5. Introduction to Advanced Concepts	Read Foundation Topics				
5. Introduction to Advanced Concepts	Review Key Topics				
5. Introduction to Advanced Concepts	Define Key Terms				
5. Introduction to Advanced Concepts	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 5 in practice test software				
6. Cisco Identity Services Engine Architecture	Read Foundation Topics				
6. Cisco Identity Services Engine Architecture	Review Key Topics				
6. Cisco Identity Services Engine Architecture	Define Key Terms				
6. Cisco Identity Services Engine Architecture	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 6 in practice test software				
7. A Guided Tour of the Cisco ISE Graphical User Interface (GUI)					
7. A Guided Tour of the Cisco ISE Graphical User Interface (GUI)	Review Key Topics				
7. A Guided Tour of the Cisco ISE Graphical User Interface (GUI)	Define Key Terms				
7. A Guided Tour of the Cisco ISE Graphical User Interface (GUI)	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 7 in practice test software				
8. Initial Configuration of Cisco ISE	Read Foundation Topics				
8. Initial Configuration of Cisco ISE	Review Key Topics				
8. Initial Configuration of Cisco ISE	Define Key Terms				
8. Initial Configuration of Cisco ISE	Complete Q&A section				

Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 8 in practice test software				
9. Authentication Policies					
9. Authentication Policies	Review Key Topics				
9. Authentication Policies	Define Key Terms				
9. Authentication Policies	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 9 in practice test software				
10. Authorization Policies	Read Foundation Topics				
10. Authorization Policies	Review Key Topics				
10. Authorization Policies	Define Key Terms				
10. Authorization Policies	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 10 in practice test software				
11. Implement Wired and Wireless Authentication	Read Foundation Topics				
11. Implement Wired and Wireless Authentication	Review Key Topics				
11. Implement Wired and Wireless Authentication	Define Key Terms				
11. Implement Wired and Wireless Authentication	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 11 in practice test software				
12. Web Authentication	Read Foundation Topics				
12. Web Authentication	Review Key Topics				
12. Web Authentication	Define Key Terms				
12. Web Authentication	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 12 in practice test software				
13. Guest Services	Read Foundation Topics				
13. Guest Services	Review Key Topics				
13. Guest Services	Define Key Terms				
13. Guest Services	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 13 in practice test software				
14. Profiling	Read Foundation Topics				
14. Profiling	Review Key Topics				
14. Profiling	Define Key Terms				
14. Profiling	Complete Q&A section				

Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 14 in practice test software				
15. Certificate-Based Authentication	Read Foundation Topics				
15. Certificate-Based Authentication	Review Key Topics				
15. Certificate-Based Authentication	Define Key Terms				
15. Certificate-Based Authentication	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 15 in practice test software				
16. Bring-Your-Own-Device	Read Foundation Topics				
16. Bring-Your-Own-Device	Review Key Topics				
16. Bring-Your-Own-Device	Define Key Terms				
16. Bring-Your-Own-Device	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 16 in practice test software				
17. TrustSec and MACsec	Read Foundation Topics				
17. TrustSec and MACsec	Review Key Topics				
17. TrustSec and MACsec	Define Key Terms				
17. TrustSec and MACsec	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 17 in practice test software				
18. Posture Assessment	Read Foundation Topics				
18. Posture Assessment	Review Key Topics				
18. Posture Assessment	Define Key Terms				
18. Posture Assessment	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 18 in practice test software				
19. Deploying Safely	Read Foundation Topics				
19. Deploying Safely	Review Key Topics				
19. Deploying Safely	Define Key Terms				
19. Deploying Safely	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 19 in practice test software				
20. ISE Scale and High Availability	Read Foundation Topics				
20. ISE Scale and High Availability	Review Key Topics				
20. ISE Scale and High Availability	Define Key Terms				

20. ISE Scale and High Availability	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 20 in practice test software				
21. Troubleshooting Tools	Read Foundation Topics				
21. Troubleshooting Tools	Review Key Topics				
21. Troubleshooting Tools	Define Key Terms				
21. Troubleshooting Tools	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 21 in practice test software				
22. ISE Context Sharing and Remediation	Read Foundation Topics				
22. ISE Context Sharing and Remediation	Review Key Topics				
22. ISE Context Sharing and Remediation	Define Key Terms				
22. ISE Context Sharing and Remediation	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 22 in practice test software				
23. Threat Centric NAC	Read Foundation Topics				
23. Threat Centric NAC	Review Key Topics				
23. Threat Centric NAC	Define Key Terms				
23. Threat Centric NAC	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 23 in practice test software				
24. Device Administration AAA with ISE	Read Foundation Topics				
24. Device Administration AAA with ISE	Review Key Topics				
24. Device Administration AAA with ISE	Define Key Terms				
24. Device Administration AAA with ISE	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 24 in practice test software				
25. Configuring Device Admin AAA with Cisco IOS	Read Foundation Topics				
25. Configuring Device Admin AAA with Cisco IOS	Review Key Topics				
25. Configuring Device Admin AAA with Cisco IOS	Define Key Terms				
25. Configuring Device Admin AAA with Cisco IOS	Complete Q&A section				

Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 25 in practice test software				
26. Configuring Device Admin AAA with the Cisco WLC	Read Foundation Topics				
26. Configuring Device Admin AAA with the Cisco WLC	Review Key Topics				
26. Configuring Device Admin AAA with the Cisco WLC	Define Key Terms				
26. Configuring Device Admin AAA with the Cisco WLC	Complete Q&A section				
Practice Test	Take practice test in study mode using Exam Bank 1 questions for Chapter 26 in practice test software				
27. Final Preparation					
27. Final Preparation	Take practice test in study mode for all book questions in practice test software				
27. Final Preparation	Review all Key Topics in all chapters				
27. Final Preparation	Take practice test in practice exam mode using Exam Bank 1 questions for all chapters				
27. Final Preparation	Take practice test in practice exam mode using Exam Bank 2 questions for all chapters				

CCNP Security Identity Management SISE 300-715 Official Cert Guide Companion Website

Access interactive study tools on this book's companion website, including practice test software, memory tables, review exercises, Key Term flash card application, study planner, and more!

To access the companion website, simply follow these steps:

1. Go to **www.ciscopress.com/register**.
2. Enter the print book ISBN: **9780136642947**.
3. Answer the security question to validate your purchase.
4. Go to your account page.
5. Click on the **Registered Products** tab.
6. Under the book listing, click on the **Access Bonus Content** link.

If you have any issues accessing the companion website, you can contact our support team by going to **pearsonitp.echelp.org.**

Pearson

Where are the companion content files?



Thank you for purchasing this Premium Edition version of

CCNP Security Identity Management SISE 300-715 Official Cert Guide

This product comes with companion content. You have access to these files by following the steps below:

1. Go to **ciscopress.com/account** and log in.
2. Click on the "Access Bonus Content" link in the Registered Products section of your account page for this product, to be taken to the page where your downloadable content is available.

Please note that many of our companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following the steps at left, please visit **ciscopress.com/support** and select the chat, phone, or web ticket options to get help from a tech support representative.

The Professional and Personal Technology Brands of Pearson

Addison Wesley · Cisco Press · informIT · PEARSON IT Certification · que · SAMS