CISCO

Official Cert Guide

Learn, prepare, and practice for exam success



CCNP Security SISAS 300-208

AARON T. WOLAND, CCIE NO. 20113
KEVIN REDMON

ciscopress.com

https://t.me/learningnets

About This eBook

ePUB is an open, industry-standard format for eBooks. However, support of ePUB and its many features varies across reading devices and applications. Use your device or app settings to customize the presentation to your liking. Settings that you can customize often include font, font size, single or double column, landscape or portrait mode, and figures that you can click or tap to enlarge. For additional information about the settings and features on your reading device or app, visit the device manufacturer's Web site.

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a "Click here to view code image" link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.

https://t.me/learningnets

CCNP Security SISAS 300-208

Official Cert Guide

Aaron T. Woland, CCIE No. 20113
Kevin Redmon

Cisco Press
800 East 96th Street
Indianapolis, IN 46240

https://t.me/learningnets

CCNP Security SISAS 300-208 Official Cert Guide

Aaron T. Woland
Kevin Redmon

Copyright © 2015 Cisco Systems, Inc.

Published by:
Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

First Printing April 2015

Library of Congress Control Number: 2015936634

ISBN-13: 978-1-58714-426-4

ISBN-10: 1-58714-426-3

Warning and Disclaimer

This book is designed to provide information about network security. 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 authors 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.

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: U.S. Corporate and Government Sales 1-800-382-3419 corpsales@pearsontechgroup.com

For sales outside of the U.S. please contact: International Sales international@pearsoned.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.

https://t.me/learningnets

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 e-mail at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Development Editor: Eleanor C. Bru

Managing Editor: Sandra Schroeder

Project Editor: Seth Kerney

Editorial Assistant: Vanessa Evans

Cover Designer: Mark Shirar

Composition: Bumpy Design

Business Operation Manager, Cisco Press: Jan Cornelssen

Executive Editor: Mary Beth Ray

Copy Editor: Megan Wade-Taxter

Technical Editors: Tim Abbott, Konrad Reszka

Proofreader: Jess DeGabriele

Indexer: Tim Wright



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose. CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel:+65 6317 7777
Fax:+65 6317 7799

Europe Headquarters
Cisco Systems International BV

https://t.me/learningnets

Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel:+31 0 800 020 0791
Fax:+31 0 203 571 100

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.**

©2007 Cisco Systems, Inc. All rights reserved. CCVR the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live. Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar. Ainonet, BPX, Catalyst, CCDA, CCDP CCIE, CCIP CCNA, CCNP CCSP Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems. Cisco Systems Capital, the Cisco Systems logo. Cisco Unity, Enterprise/Solver. EtherChannel. EtherFast, EtherSwitoh, Fast Step, Follow Me Browsing, FormShare, GigaDrive. GigaStack HomeLink Internet Quotient, IOS, IP/TV iQ Expertise, the iQ logo iQ Net Readiness Scorecard, iQuick Study. LightStream, Linksys, MeetingPlace. MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)

https://t.me/learningnets

About the Authors

Aaron T. Woland, CCIE No. 20113, is a principal engineer within Cisco's technical marketing organization and works with Cisco's largest customers all over the world. His primary job responsibilities include secure access and identity deployments with ISE, solution enhancements, standards development, 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. His areas of expertise include network and host security architecture and implementation, regulatory compliance, virtualization, as well as route-switch and wireless. Technology is certainly his passion, and Aaron currently has two patents in pending status with the United States Patent and Trade Office.

Aaron is the author of the Cisco ISE for BYOD and Secure Unified Access book (Cisco Press) and many published whitepapers and design guides. Aaron is one of the first six members of the Hall of Fame for Distinguished Speakers at Cisco Live and is a security columnist for Network World, where he blogs on all things related to identity. In addition to being a proud holder of a CCIE-Security, his other certifications include GCIH, GSEC, CEH, MCSE, VCP, CCSP, CCNP, CCDP, and many other industry certifications.

Kevin Redmon is the youngest of 12 siblings and was born in Marion, Ohio. Since joining Cisco in October 2000, Kevin has worked closely with several Cisco design organizations; as a firewall/VPN customer support engineer with the Cisco Technical Assistant Center; as a systems test engineer in BYOD Smart Solutions Group; and now as a systems test engineer in the IoT Vertical Solutions Group in RTP, NC with a focus on the connected transportation systems.

Besides co-authoring this book with Aaron Woland, Kevin is also the author of the Cisco Press Video Series titled Cisco Bring Your Own Device (BYOD) Networking LiveLessons. He has a bachelor of science in computer engineering from Case Western Reserve University and a master of science in information security from East Carolina University, as well as several Cisco certifications. Kevin enjoys presenting on network security-related topics and Cisco's latest solutions. He has presented several times at Cisco Live, focusing on network security-related topics and has achieved the honor of Distinguished Speaker.

Kevin enjoys innovating new ideas to keep his mind fresh and currently has a patent listed with the United States Patent and Trade Office. He spends his free time relaxing with his wife, Sonya, and little girl, Melody, in Durham, North Carolina.

https://t.me/learningnets

About the Technical Reviewers

Tim Abbott is a technical marketing engineer at Cisco Systems who works with Cisco customers all over the world. He holds a bachelor's degree from the University of Texas at San Antonio. His primary responsibilities at Cisco include ISE deployment design and writing solution guides for Cisco customers and partners. Tim has held CCNA and CCNP certifications and was also named Distinguished Speaker at Cisco Live. He has more than 10 years of IT experience in areas such as network security, routing and switching, remote access, and data center technologies.

Konrad Reszka is a software engineer at Cisco Systems specializing in designing and validating end-to-end solutions. He has contributed to many architectures and design guides spanning multiple technologies, including data center, security, wireless, and Carrier Ethernet. He is a distinguished speaker at Cisco Live, where you can catch him giving talks on the Internet of Everything, BYOD, and MPLS VPNs. Konrad holds a degree in computer science from the University of North Carolina at Chapel Hill.

https://t.me/learningnets

Dedications

Aaron Woland: First and foremost, this book is dedicated to my amazing best friend, fellow adventurer, and wife, Suzanne. This book would surely not exist without your continued love, support, guidance, wisdom, encouragement, and patience, as well as the occasional reminder that I need to "get it done." Thank you for putting up with all the long nights and weekends I had to be writing. I doubt that I could be as patient and understanding with the bright laptop and the typing next to me while I tried to sleep. You are amazing.

To Mom and Pop. You have always believed in me and supported me in absolutely everything I've ever pursued, showed pride in my accomplishments no matter how small, encouraged me to never stop learning, and engrained in me the value of hard work and to strive for a career in a field that I love. I hope I can continue to fill your lives with pride and happiness, and if I succeed, it will still only be a fraction of what you deserve.

To my two awesome and brilliant children, Eden and Nyah: You girls are my inspiration, my pride and joy, and continue to make me want to be a better man. Eden, when I look at you and your accomplishments over your 16 years of life, I swell with pride. You are so intelligent, kind, and hard-working. You will make a brilliant engineer one day, or if you change your mind, I know you will be brilliant in whatever career you find yourself pursuing (perhaps a dolphin trainer). Nyah, you are my morning star, my princess. You have the biggest heart, the kindest soul, and a brilliant mind. You excel at everything you put your mind to, and I look forward to watching you grow and using that power to change the world. Maybe that power will be used within marine biology, or maybe you will follow in my footsteps. I can't wait to see it for myself.

To my brother, Dr. Bradley Woland: Thank you for being so ambitious, so driven. It forced my competitive nature to always want more. As I rambled on in the 12-minute wedding speech, you not only succeed at everything you try, you crush it! If you were a bum, I would never have pushed myself to the levels that I have. To Bradley's beautiful wife, Claire: I am so happy that you are a member of my family now; your kindness, intelligence, and wit certainly keep my brother in check and keep us all smiling.

My sister, Anna. If I hadn't always had to compete with you for our parents' attention and to keep my things during our "garage sales," I would probably have grown up very naive and vulnerable. You drove me to think outside the box and find new ways to accomplish the things I wanted to do. Seeing you not just succeed in life and in school truly had a profound effect on my life. Thank you for marrying Eddie, my brilliant brother-in-law. Eddie convinced me that I could actually have a career in this technology stuff, and without his influence I certainly would not be where I am today.

Lastly, to my grandparents: Jack, Lola, Herb, and Ida. You have taught me what it means to be alive and the true definition of courage, survival, perseverance, hard work, and never giving up.

—Aaron

Kevin Redmon: There are a number of people who, without them, my coauthoring this book would not be possible.

To my lovely wife, Sonya, and daughter, Melody: You both demonstrated an amazing amount of love, patience, and support throughout this book process, allowing me to spend numerous weekends and late nights in isolation to write. Sonya, you are my all, and I love you. I'm am the luckiest man alive to have you as my co-pilot in life. Melody, thank you for being the beautiful princess that you are—Daddy loves you so much! Now that this book is done, my time again belongs to you both! Thank you both—with big hugs and kisses! I love you with all of my heart!

To my mom, Helen, and my brother, Jeffrey: Through the years, you both have provided me the tools, confidence, and financial support to achieve my dreams and go to college, enabling me to achieve my long career at Cisco and to, eventually, write this book. You have always been there to remind me that I can do whatever I put my mind to and to never quit—and, when I doubted that, you kept me in check. You both deserve all the riches that this world can give you, and then some. I love you, Mom! I love you, Bro!

https://t.me/learningnets

To Adam Meiggs: You have been an inspiration, a rock, and an amazing friend. You helped me get over stage fright, allowing me to get in front of people, and to never say “I can’t!” Thanks for being there for me, Kid! I miss you, and there is rarely a day that goes by that I don’t think of you!

To Mr. Rick Heavner: Thank you for taking me under your wing in 4th grade and instilling in me humility and a love for computers. This was truly a turning point in my personal and, eventually, professional development. From the bottom of my heart, THANK YOU!!!

To Mrs. Joyce Johnston: Thank you for being you and helping me to recognize the intellectual gifts that I have been given. You helped me see my untapped talent and that I can achieve excellence with a little bit of hard work. From your Algebra King, thanks!

To Mr. Donald Wolfe: Thank you for being such a great friend and driving me to my scholarship interview in Columbus during my senior year. I didn’t get the scholarship, but that rejection gave me the fire in my belly to fight, kick, and scream through my undergrad at CWRU. Defeat was never an option. From one Baldy to another, thank you!

To my teachers from Glenwood Elementary, Edison Middle School, and Marion Harding High School in Marion, Ohio: I know that being a teacher can be a thankless career at times, but I do want to change that and say THANK YOU!!! Because of your dedication to teaching, I was able to achieve more than a man of my humble beginnings could ever dream of! Thank you for helping me achieve these dreams; without you, this would not have been possible.

To all of my friends: Thank you for being there through the years to support me. I know it was a tough job at times. Most of all, thank you for helping to make me who I am.

https://t.me/learningnets

Acknowledgments

Aaron Woland:

There are so many to acknowledge. This feels like a speech at the Academy Awards, and I'm afraid I will leave out too many people.

Thomas Howard and Allan Bolding from Cisco, for their continued support, encouragement, and guidance. Most importantly, for believing in me even though I can be difficult at times. I could not have done any of it without you.

Craig Hyps, a senior technical marketing engineer at Cisco. "Senior" doesn't do you justice, my friend. You are a machine. You possess such deep technical knowledge on absolutely everything (not just pop culture). Your constant references to pop culture keep me laughing, and your influence can be found on content all throughout the book and this industry. "Can you dig it?"

Christopher Heffner, an engineer at Cisco, for convincing me to step up and take a swing at being an author and for twisting my arm to put "pen to paper" a second time. Without your encouragement and enthusiasm, this book would not exist.

I am honored to work with so many brilliant and talented people every day. Among those: Jesse Dubois, Vivek Santuka, Christopher Murray, Doug Gash, Chad Mitchell, Jamie Sanbower, Louis Roggo, Kyle King, Tim Snow, Chad Sullivan, and Brad Spencer. You guys truly amaze me.

Chip Current and Paul Forbes: You guys continue to show the world what it means to be a real product owner and not just a PM. I have learned so much from you both, and I'm not referring only to vocabulary words.

To my world-class TME team: Hosuk Won, Tim Abbott, Hsing-Tsu Lai, Imran Bashir, Ziad Sarieddine, John Eppich, Fay-Ann Lee, Jason Kunst, Paul Carco, and Aruna Yerragudi. World-class is not a strong enough word to describe this team. You are beyond inspirational, and I am proud to be a member of this team.

Darrin Miller, Nancy Cam-Winget, and Jamey Heary, distinguished engineers who set the bar so incredibly high. You are truly inspirational; people to look up to and aspire to be like, and I appreciate all the guidance you have given me.

Jonny Rabinowitz, Mehdi Bouzouina, and Christopher Murray: You three guys continue to set a high bar and somehow move that bar higher all the time. All three of you have a fight in you to never lose, and it's completely infectious. Chris, your constant enthusiasm, energy, brilliance, and expertise impresses me and inspires me.

Lisa Lorenzin, Cliff Cahn, Scott Pope, Steve Hannah, and Steve Venema: What an amazing cast of people who are changing the world one standard at a time. It has been an honor and a privilege to work with you.

To the Original Cast Members of the one and only SSU, especially: Jason Halpern, Danelle Au, Mitsunori Sagae, Fay-Ann Lee, Pat Calhoun, Jay Bhansali, AJ Shipley, Joseph Salowey, Thomas Howard, Darrin Miller, Ron Tisinger, Brian Gonsalves, and Tien Do.

Max Pritkin, I think you have forgotten more about certificates and PKI than most experts will ever know. You have taught me so much, and I look forward to learning more from your vast knowledge and unique way of making complex technology seem easy.

To the world's greatest engineering team, and of course I mean the people who spend their days writing and testing the code that makes up Cisco's ISE. You guys continue to show the world what it means to be "world-class."

My colleagues: Naasief Edross, Andrae Middleton, Russell Rice, Dalton Hamilton, Tom Foucha, Matt Robertson, Brian Ford, Paul Russell, Brendan O'Connell, Jeremy Hyman, Kevin Sullivan, Mason Harris, David Anderson, Luc Billot, Dave White Jr., Nevin Absher, Ned Zaldivar, Mark Kassem, Greg Tillett, Chuck Parker, Jason Frazier, Shelly Cadora, Ralph Schmieder, Corey Elinburg, Scott Kenewell, Larry Boggis, Chad Sullivan, Dave Klein, Nelson Figueroa, Kevin Redmon, Konrad Reszka, and so many more! The

https://t.me/learningnets

contributions you make to this industry inspire me.

Kevin Redmon:

First and foremost, I would like to give my utmost respect and recognition to my coauthor, Aaron Woland. When it comes to Cisco Identity Services Engine (ISE) and Cisco Secure Access, Aaron has been an indispensable resource. Without his expertise and support, the Cisco ISE community and the networking security industry at-large would be devoid of a huge knowledge base. To be in the same audience with a well-respected network security expert such as Aaron is truly an amazing feeling. Thank you for allowing me the honor to coauthor this book with you.

Special acknowledgements go to my former BYOD colleagues. During the two and a half years we shared on BYOD, I learned so much from each of you. By working closely with some of the brightest minds in solutions test and networking, I was able to learn so much in such a short time, giving me the knowledge, confidence, contacts, and tools to coauthor this book. Thank you for letting some random "security guy" wreck the ranks and become a part of the team. You guys are truly the best team that I've ever had the pleasure to work with!

I want to give a special shout-out to Nelson Figueroa and Konrad Reszka. You guys are just awesome—both as friends and colleagues. You both have become my brothers, and it's always a blast to collaborate with you both. I hope the Three Musketeers can continue to shake up the networking industry, one pint at a time.

I would also like to thank our two technical editors, Tim Abbott and Konrad Reszka. Writing a book is hard, but writing a good book would be impossible without some of the best technical editors around. Both of these guys are truly gifted network engineers in their own right. These guys help to keep me honest when I randomly drop words or overlook a key detail. Also, when my schedule slips, these guys help to make up for the lost time. Thanks guys—your help is truly appreciated!

https://t.me/learningnets

Contents at a Glance

Part I The CCNP Certification

Chapter 1 CCNP Security Certification

Part II "The Triple A" (Authentication, Authorization, and Accounting)

Chapter 2 Fundamentals of AAA

Chapter 3 Identity Management

Chapter 4 EAP Over LAN (Also Known As 802.1X)

Chapter 5 Non-802.1X Authentications

Chapter 6 Introduction to Advanced Concepts

Part III Cisco Identity Services Engine

Chapter 7 Cisco Identity Services Engine Architecture

Chapter 8 A Guided Tour of the Cisco ISE Graphical User Interface

Chapter 9 Initial Configuration of the Cisco ISE

Chapter 10 Authentication Policies

Chapter 11 Authorization Policies

Part IV Implementing Secure Network Access

Chapter 12 Implement Wired and Wireless Authentication

Chapter 13 Web Authentication

Chapter 14 Deploying Guest Services

Chapter 15 Profiling

Part V Advanced Secure Network Access

Chapter 16 Certificate-Based User Authentications

Chapter 17 Bring Your Own Device

Chapter 18 TrustSec and MACSec

Chapter 19 Posture Assessment

Part VI Safely Deploying in the Enterprise

Chapter 20 Deploying Safely

Chapter 21 ISE Scale and High Availability

https://t.me/learningnets

Chapter 22 Troubleshooting Tools

Part VII Final Preparation

Chapter 23 Final Preparation

Part VIII Appendixes

Appendix A Answers to the “Do I Know This Already?” Quizzes

Appendix B Configuring the Microsoft CA for BYOD

Appendix C Using the Dogtag CA for BYOD

Appendix D Sample Switch Configurations

Glossary

Index

https://t.me/learningnets

Contents

Introduction

Part I The CCNP Certification

Chapter 1 CCNP Security Certification

CCNP Security Certification Overview

Contents of the CCNP-Security SISAS Exam

How to Take the SISAS Exam

Who Should Take This Exam and Read This Book?

Format of the CCNP-Security SISAS Exam

CCNP-Security SISAS 300-208 Official Certification Guide

Book Features and Exam Preparation Methods

Part II "The Triple A" (Authentication, Authorization, and Accounting)

Chapter 2 Fundamentals of AAA

"Do I Know This Already?" Quiz

Foundation Topics

Triple-A

Compare and Select AAA Options

Device Administration

Network Access

TACACS+

TACACS+ Authentication Messages

TACACS+ Authorization and Accounting Messages

RADIUS

AV-Pairs

Change of Authorization

Comparing RADIUS and TACACS+

Exam Preparation Tasks

Review All Key Topics

Define Key Terms

Chapter 3 Identity Management

"Do I Know This Already?" Quiz

Foundation Topics

What Is an Identity?

Identity Stores

Internal Identity Stores

https://t.me/learningnets

External Identity Stores
Active Directory
LDAP
Two-Factor Authentication
One-Time Password Services
Smart Cards
Certificate Authorities
Has the Certificate Expired?
Has the Certificate Been Revoked?
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Chapter 4 EAP Over LAN (Also Known As 802.1X)
"Do I Know This Already?" Quiz
Foundation Topics
Extensible Authentication Protocol
EAP over LAN (802.1X)
EAP Types
Native EAP Types (Nontunneled EAP)
Tunneled EAP Types
Summary of EAP Authentication Types
EAP Authentication Type Identity Store Comparison Chart
Network Access Devices
Supplicant Options
Windows Native Supplicant
Cisco AnyConnect NAM Supplicant
EAP Chaining
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Chapter 5 Non-802.1X Authentications
"Do I Know This Already?" Quiz
Foundation Topics
Devices Without a Supplicant
MAC Authentication Bypass
Web Authentication
Local Web Authentication
Local Web Authentication with a Centralized Portal
Centralized Web Authentication

https://t.me/learningnets

Remote Access Connections
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Chapter 6 Introduction to Advanced Concepts
"Do I Know This Already?" Quiz
Foundation Topics
Change of Authorization
Automating MAC Authentication Bypass
Posture Assessments
Mobile Device Managers
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Part III Cisco Identity Services Engine

Chapter 7 Cisco Identity Services Engine Architecture
"Do I Know This Already?" Quiz
Foundation Topics
What Is Cisco ISE?
Personas
Administration Node
Policy Service Node
Monitoring and Troubleshooting Node
Inline Posture Node
Physical or Virtual Appliance
ISE Deployment Scenarios
Single-Node Deployment
Two-Node Deployment
Four-Node Deployment
Fully Distributed Deployment
Communication Between Nodes
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Chapter 8 A Guided Tour of the Cisco ISE Graphical User Interface
"Do I Know This Already?" Quiz
Foundation Topics
Logging In to ISE

https://t.me/learningnets

- Initial Login
- Administration Dashboard
- Administration Home Page
- Server Information
- Setup Assistant
- Help

Organization of the ISE GUI

- Operations
- Authentications
- Reports
- Endpoint Protection Service
- Troubleshoot
- Policy
- Authentication
- Authorization
- Profiling
- Posture
- Client Provisioning
- Security Group Access
- Policy Elements
- Administration
- System
- Identity Management
- Network Resources
- Web Portal Management
- Feed Service

Type of Policies in ISE

- Authentication
- Authorization
- Profiling
- Posture
- Client Provisioning
- Security Group Access

Exam Preparation Tasks

Review All Key Topics

Define Key Terms

Chapter 9 Initial Configuration of Cisco ISE

"Do I Know This Already?" Quiz

Foundation Topics

Cisco Identity Services Engine Form Factors

https://t.me/learningnets

Bootstrapping Cisco ISE
Where Are Certificates Used with the Cisco Identity Services Engine?
Self-Signed Certificates
CA-Signed Certificates
Network Devices
Network Device Groups
Network Access Devices
Local User Identity Groups
Local Endpoint Groups
Local Users
External Identity Stores
Active Directory
Prerequisites for Joining an Active Directory Domain
Joining an Active Directory Domain
Certificate Authentication Profile
Identity Source Sequences
Exam Preparation Tasks
Review All Key Topics

Chapter 10 Authentication Policies

"Do I Know This Already?" Quiz
Foundation Topics
The Relationship Between Authentication and Authorization
Authentication Policy
Goals of an Authentication Policy
Goal 1—Accept Only Allowed Protocols
Goal 2—Select the Correct Identity Store
Goal 3—Validate the Identity
Goal 4—Pass the Request to the Authorization Policy
Understanding Authentication Policies
Conditions
Allowed Protocols
Extensible Authentication Protocol Types
Tunneled EAP Types
Identity Store
Options
Common Authentication Policy Examples
Using the Wireless SSID
Remote Access VPN
Alternative ID Stores Based on EAP Type
More on MAB

https://t.me/learningnets

Restore the Authentication Policy
Exam Preparation Tasks
Review All Key Topics

Chapter 11 Authorization Policies
"Do I Know This Already?" Quiz
Foundation Topics
Authentication Versus Authorization
Authorization Policies
Goals of Authorization Policies
Understanding Authorization Policies
Role-specific Authorization Rules
Authorization Policy Example
Employee Full Access Rule
Internet Only for Smart Devices
Employee Limited Access Rule
Saving Conditions for Reuse
Combining AND with OR Operators
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Part IV Implementing Secure Network Access

Chapter 12 Implement Wired and Wireless Authentication
"Do I Know This Already?" Quiz
Foundation Topics
Authentication Configuration on Wired Switches
Global Configuration AAA Commands
Global Configuration RADIUS Commands
IOS 12.2.X
IOS 15.X
Both IOS 12.2.X and 15.X
Global 802.1X Commands
Creating Local Access Control Lists
Interface Configuration Settings for All Cisco Switches
Configuring Interfaces as Switchports
Configuring Flexible Authentication and High Availability
Host Mode of the Switchport
Configuring Authentication Settings
Configuring Authentication Timers
Applying the Initial ACL to the Port and Enabling Authentication

https://t.me/learningnets

Authentication Configuration on WLCs
Configuring the AAA Servers
Adding the RADIUS Authentication Servers
Adding the RADIUS Accounting Servers
Configuring RADIUS Fallback (High-Availability)
Configuring the Airespace ACLs
Creating the Web Authentication Redirection ACL
Creating the Posture Agent Redirection ACL
Creating the Dynamic Interfaces for the Client VLANs
Creating the Guest Dynamic Interface
Creating the Wireless LANs
Creating the Guest WLAN
Creating the Corporate SSID
Verifying Dot1X and MAB
Endpoint Supplicant Verification
Network Access Device Verification
Verifying Authentications with Cisco Switches
Sending Syslog to ISE
Verifying Authentications with Cisco WLCs
Cisco ISE Verification
Live Authentications Log
Live Sessions Log
Looking Forward
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Chapter 13 Web Authentication

"Do I Know This Already?" Quiz
Foundation Topics
Web Authentication Scenarios
Local Web Authentication
Centralized Web Authentication
Device Registration WebAuth
Configuring Centralized Web Authentication
Cisco Switch Configuration
Configuring Certificates on the Switch
Enabling the Switch HTTP/HTTPS Server
Verifying the URL-Redirection ACL
Cisco WLC Configuration
Validating That MAC Filtering Is Enabled on the WLAN

https://t.me/learningnets

Validating That Radius NAC Is Enabled on the WLAN
Validate That the URL-Redirection ACL Is Configured
Captive Portal Bypass
Configuring ISE for Centralized Web Authentication
Configuring MAB for the Authentication
Configuring the Web Authentication Identity Source Sequence
Configuring a dACL for Pre-WebAuth Authorization
Configuring an Authorization Profile
Building CWA Authorization Policies
Creating the Rule to Redirect to CWA
Creating the Rules to Authorize Users Who Authenticate via CWA
Creating the Guest Rule
Creating the Employee Rule
Configuring Device Registration Web Authentication
Creating the Endpoint Identity Group
Creating the DRW Portal
Creating the Authorization Profile
Creating the Rule to Redirect to DRW
Creating the Rule to Authorize DRW-Registered Endpoints
Verifying Centralized Web Authentication
Checking the Experience from the Client
Checking on ISE
Checking the Live Log
Checking the Endpoint Identity Group
Checking the NAD
show Commands on the Wired Switch
Viewing the Client Details on the WLC
Exam Preparation Tasks
Review All Key Topics

Chapter 14 Deploying Guest Services

"Do I Know This Already?" Quiz
Foundation Topics
Guest Services Overview
Guest Services and WebAuth
Portal Types
Configuring the Web Portal Settings
Port Numbers
Interfaces
Friendly Names
Configuring the Sponsor Portal Policies

https://t.me/learningnets

Sponsor Types
Mapping Groups
Guest User Types
Managing Guest Portals
Portal Types
Building Guest Authorization Policies
Provisioning Guest Accounts from a Sponsor Portal
Individual
Random
Import
Verifying Guest Access on the WLC/Switch
WLC
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Chapter 15 Profiling

"Do I Know This Already?" Quiz
Foundation Topics
ISE Profiler
Cisco ISE Probes
Probe Configuration
DHCP and DHCPSPAN
RADIUS
Network Scan
DNS
SNMPQUERY and SNMPTRAP
NETFLOW
HTTP Probe
HTTP Profiling Without Probes
Infrastructure Configuration
DHCP Helper
SPAN Configuration
VLAN Access Control Lists
Device Sensor
VMware Configurations to Allow Promiscuous Mode
Profiling Policies
Profiler Feed Service
Configuring the Profiler Feed Service
Verifying the Profiler Feed Service
Endpoint Profile Policies

https://t.me/learningnets

Logical Profiles
ISE Profiler and CoA
Global CoA
Per-profile CoA
Global Profiler Settings
Endpoint Attribute Filtering
Profiles in Authorization Policies
Endpoint Identity Groups
EndPointPolicy
Verify Profiling
The Dashboard
Endpoints Drill-down
Global Search
Endpoint Identities
Device Sensor Show Commands
Exam Preparation Tasks
Review All Key Topics

Part V Advanced Secure Network Access

Chapter 16 Certificate-Based User Authentications
"Do I Know This Already?" Quiz
Foundation Topics
Certificate Authentication Primer
Determine Whether a Trusted Authority Has Signed the Digital Certificate
Examine Both the Start and End Dates to Determine Whether the Certificate Has Expired
Verify Whether the Certificate Has Been Revoked
Validate That the Client Has Provided Proof of Possession
A Common Misconception About Active Directory
EAP-TLS
Configuring ISE for Certificate-Based Authentications
Validate Allowed Protocols
Certificate Authentication Profile
Verify That the Authentication Policy Is Using CAP
Authorization Policies
Ensuring the Client Certificates Are Trusted
Importing the Certificate Authority's Public Certificate
Configuring Certificate Status Verification (optional)
Verifying Certificate Authentications
Exam Preparation Tasks
Review All Key Topics

https://t.me/learningnets

Define Key Terms

Chapter 17 Bring Your Own Device

"Do I Know This Already?" Quiz

Foundation Topics

BYOD Challenges

Onboarding Process

BYOD Onboarding

Dual SSID

Single SSID

Configuring NADs for Onboarding

Configuring the WLC for Dual-SSID Onboarding

Reviewing the WLAN Configuration

Verifying the Required ACLs

ISE Configuration for Onboarding

The End User Experience

Single-SSID with Apple iOS Example

Dual SSID with Android Example

Unsupported Mobile Device—Blackberry Example

Configuring ISE for Onboarding

Creating the Native Supplicant Profile

Configuring the Client Provisioning Policy

Configuring the WebAuth

Verifying Default Unavailable Client Provisioning Policy Action

Creating the Authorization Profiles

Creating the Authorization Rules for Onboarding

Creating the Authorization Rules for the EAP-TLS Authentications

Configuring SCEP

BYOD Onboarding Process Detailed

iOS Onboarding Flow

Phase 1: Device Registration

Phase 2: Device Enrollment

Phase 3: Device Provisioning

Android Flow

Phase 1: Device Registration

Phase 2: Download SPW

Phase 3: Device Provisioning

Windows and Mac OSX Flow

Phase 1: Device Registration

Phase 2: Device Provisioning

Verifying BYOD Flows

https://t.me/learningnets

- Live Log
- Reports
- Identities
- MDM Onboarding
  - Integration Points
  - Configuring MDM Integration
  - Configuring MDM Onboarding Rules
  - Creating the Authorization Profile
  - Creating the Authorization Rules
- Managing Endpoints
  - Self Management
  - Administrative Management
- The Opposite of BYOD: Identify Corporate Systems
- Exam Preparation Tasks
- Review All Key Topics
- Define Key Terms

Chapter 18 TrustSec and MACSec

- "Do I Know This Already?" Quiz
- Foundation Topics
- Ingress Access Control Challenges
  - VLAN Assignment
  - Ingress Access Control Lists
- What Is TrustSec?
- What Is a Security Group Tag?
- Defining the SGTs
- Classification
  - Dynamically Assigning SGT via 802.1X
  - Manually Assigning SGT at the Port
  - Manually Binding IP Addresses to SGTs
  - Access Layer Devices That Do Not Support SGTs
  - Mapping a Subnet to an SGT
  - Mapping a VLAN to an SGT
- Transport: Security Group Exchange Protocol
  - SXP Design
  - Configuring SXP on IOS Devices
  - Configuring SXP on Wireless LAN Controllers
  - Configuring SXP on Cisco ASA
  - Verifying SXP Connections in ASDM
- Transport: Native Tagging
  - Configuring Native SGT Propagation (Tagging)

https://t.me/learningnets

Configuring SGT Propagation on Cisco IOS Switches
Configuring SGT Propagation on a Catalyst 6500
Configuring SGT Propagation on a Nexus Series Switch
Enforcement
SGACL
Security Group Firewalls
Security Group Firewall on the ASA
Security Group Firewall on the ISR and ASR
MACSec
Downlink MACSec
Switch Configuration Modes
ISE Configuration
Uplink MACSec
Manually Configuring Uplink MACSec
Verifying the Manual Configuration
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

Chapter 19 Posture Assessment
"Do I Know This Already?" Quiz
Foundation Topics
Posture Service Overview
Posture Flow
Agent Types
Posture Conditions
CoA with Posture
Configuring Posture
Downloading CPP Resources
Client Provisioning Policy
Posture Policy Building Blocks
Condition
Remediation
Requirement
Modifying the Authorization Policy for CPP
Modifying the Authorization Policy for Compliance
Verifying Posture and Redirect
Exam Preparation Tasks
Review All Key Topics
Define Key Terms

https://t.me/learningnets

Part VI Safely Deploying in the Enterprise

Chapter 20 Deploying Safely

- "Do I Know This Already?" Quiz
- Foundation Topics
- Why Use a Phased Approach?
- A Phased Approach
  - Comparing Authentication Open to Standard 802.1X
  - Preparing ISE for a Staged Deployment
  - Monitor Mode
  - Low-Impact Mode
  - Closed Mode
- Transitioning from Monitor Mode to Your End State
- Wireless Networks
- Exam Preparation Tasks
- Review All Key Topics

Chapter 21 ISE Scale and High Availability

- "Do I Know This Already?" Quiz
- Foundation Topics
- Configuring ISE Nodes in a Distributed Environment
- Making the First Node a Primary Device
- Registering an ISE Node to the Deployment
  - Ensuring the Personas of All Nodes Are Accurate
- Licensing in a Multinode ISE Cube
- Understanding the HA Options Available
  - Primary and Secondary Nodes
  - Monitoring and Troubleshooting Nodes
  - Policy Administration Nodes
  - Node Groups
- Using Load Balancers
  - General Guidelines
  - Failure Scenarios
- IOS Load Balancing
- Maintaining ISE Deployments
  - Patching ISE
  - Backup and Restore
- Exam Preparation Tasks
- Review All Key Topics
- Define Key Terms

Chapter 22 Troubleshooting Tools

https://t.me/learningnets

- “Do I Know This Already?” Quiz
- Foundation Topics
- Logging
  - Live Log
  - Live Sessions Log
  - Logging and Remote Logging
  - Logging Targets
  - Logging Categories
  - Debug Logs
  - Downloading Debug Logs from the GUI
  - Viewing Log Files from the CLI
  - Support Bundles
- Diagnostics Tools
  - Evaluate Configuration Validator
  - RADIUS Authentication Troubleshooting Tool
  - TCP Dump
  - Ensuring Live Log Displays All Events (Bypassing Suppression)
  - Disabling Suppression
- Troubleshooting Outside of ISE
  - Endpoint Diagnostics
  - AnyConnect Diagnostics and Reporting Tool
  - AnyConnect NAM Extended Logging
  - Microsoft Native Supplicant
  - Supplicant Provisioning Logs
  - Network Device Troubleshooting
  - The Go-To: show authentication session interface
  - Viewing Client Details on the WLC
  - Debug Commands
- Exam Preparation Tasks
- Review All Key Topics

Part VII Final Preparation

Chapter 23 Final Preparation

- Advice About the Exam Event
  - Learning the Question Types Using the Cisco Certification Exam Tutorial
  - Thinking About Your Time Budget Versus Number of Questions
  - A Suggested Time-Check Method
  - Miscellaneous Pre-Exam Suggestions
  - Exam-Day Advice
- Exam Review
  - Taking Practice Exams

https://t.me/learningnets

Practicing Taking the SISAS Exam
Advice on How to Answer Exam Questions
Taking Other Practice Exams
Finding Knowledge Gaps Through Question Review
Other Study Tasks
Final Thoughts

Part VIII Appendixes

Appendix A Answers to the "Do I Know This Already?" Quizzes

Appendix B Configuring the Microsoft CA for BYOD

CA Requirements
Other Useful Information
Microsoft Hotfixes
AD Account Roles
Configuration Steps
Installing the CA
Adding the Remaining Roles
Configuring the Certificate Template
Publishing the Certificate Template
Editing the Registry
Useful Links

Appendix C Using the Dogtag CA for BYOD

What Is Dogtag, and Why Use It?
Prerequisites
Installing 32-bit Fedora 15
Configuring Networking
Installing Packages with yum
Configuring Proxy (if Needed)
Updating System Packages with yum
Installing and Configuring the NTP Service
Installing the LDAP Server
Installing the PHP Services
Installing and Configuring Dogtag
Modifying the Firewall Rules (iptables)
Creating a New CA Instance
Enabling and Configuring SCEP
Preparing Apache
Configuring ISE to Use the New Dogtag CA
Adding Dogtag to the SCEP RA Profiles

https://t.me/learningnets

Appendix D Sample Switch Configurations

Catalyst 2960/3560/3750 Series, 12.2(55)SE

Catalyst 3560/3750 Series, 15.0(2)SE

Catalyst 4500 Series, IOS-XE 3.3.0/15.1(1)SG

Catalyst 6500 Series, 12.2(33)SXJ

Glossary

Index

https://t.me/learningnets

Icons



https://t.me/learningnets

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).
- Italics indicate arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [ ] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.

https://t.me/learningnets

Introduction

Welcome to the world of Cisco Career Certifications and the CCNP-Security. Moreover, welcome to the world of access control. Technology continues to evolve the way we do business, the types of devices that we use, the new threat vectors, and how we protect our valued assets. Through all these changes, organizations need intelligent solutions to enforce corporate policy in the access technologies that are deployed.

This book is designed to help you prepare for the Cisco CCNP Security 300-208 SISAS (Implementing Cisco Secure Access Solutions) certification exam, which is one of the four required exams to achieve the Cisco CCNP Security.

Goals and Methods

This book will help the reader understand, design, and deploy Cisco's Secure Unified Access system. This system will combine 802.1X, profiling, posture assessments, device onboarding, and guest lifecycle management.

The reader will learn all the items that make up the SISAS 300-208 exam blueprint in a realistic method using building blocks of information. Each chapter builds on the knowledge learned in the previous chapters.

How This Book Is Organized

Although you could read this book cover-to-cover, it is designed to be flexible and allow you to easily move between chapters and sections of chapters to cover only the material you need. If you do intend to read them all, the order in which they are presented is an excellent sequence.

Chapters 1–23 cover the following topics:

- **Chapter 1, "CCNP Security Certification,"** discusses the CCNP security certification with an overview and the contents of the SISAS 300-208 exam. It includes a discussion on how to take the SISAS exam and the exam's format. Additionally, features of the book and exam preparation methods are covered.
- **Chapter 2, "Fundamentals of AAA,"** builds a strong foundation for the concepts of authentication, authorization, and accounting (AAA). Comparisons and examples of the current AAA technologies and purposes are provided.
- **Chapter 3, "Identity Management,"** covers the many identity sources and how they work as related to secure network access.
- **Chapter 4, "EAP over LAN (also Known as 802.1X),"** discusses the IEEE standard for port-based network access control, its history, its progression, and the current state of the art.
- **Chapter 5, "Non-802.1X Authentications,"** details MAC authentication bypass (MAB) and the various types of web authentications. This chapter strengthens the foundation built in the first four chapters and is reinforced by Chapters 6, 12, and 13.
- **Chapter 6, "Introduction to Advanced Concepts,"** builds on the strong foundation and starts to expand the reader's knowledge base with an introduction into technologies such as profiling, posture, and BYOD.
- **Chapter 7, "Cisco Identity Services Engine Architecture,"** discusses the design of Cisco ISE, personas, and general deployment model.
- **Chapter 8, "A Guided Tour of the Cisco ISE Graphical User Interface,"** walks the reader through the many screens that make up the Cisco ISE graphical user interface.
- **Chapter 9, "Initial Configuration of Cisco ISE,"** guides the reader step-by-step through the bootstrapping and initial setup of Cisco ISE.

https://t.me/learningnets

- **Chapter 10, "Authentication Policies,"** discusses the aspects of authentication policies, authentication methods, protocols, conditions, and results. The reader will learn about accessing the identity sources described in Chapter 3 to verify and validate the identity of the user or device attempting network access.
- **Chapter 11, "Authorization Policies,"** discusses the aspects of authorization policies, attribute sources, conditions, and results. The reader will learn about leveraging the identity learned in Chapter 11, accessing attributes of that identity, and utilizing those attributes to form the access control decision.
- **Chapter 12, "Implement Wired and Wireless Authentication,"** discusses the enabling of 802.1X and non-dot1x authentication and configuring the authorization policy to send the appropriate results.
- **Chapter 13, "Web Authentication,"** builds on the knowledge obtained in Chapter 5; this chapter puts the various web authentication mechanisms into play in the network access policies.
- **Chapter 14, "Deploying Guest Services,"** discusses extending the authentication and authorization policies with guest lifecycle services, including sponsored and self-registering guests.
- **Chapter 15, "Profiling,"** discusses the network configuration and ISE configuration related to profiling and profile data collection. Additionally, the chapter focuses on the profiling feed service and profile policies themselves.
- **Chapter 16, "Certificate-Based Authentications,"** discusses the use of end-entity certificates for authentication with EAP-Transport Layer Security (EAP-TLS). X.509 certificates, the signing of certificates, as well as the authentication process are examined in detail.
- **Chapter 17, "Bring Your Own Device,"** discusses the use of personal devices on the corporate network, differentiating between corporate and personal devices, and the onboarding of devices with Native Supplicant Provisioning (NSP). The ISE policies as well as the network device configuration are examined in detail.
- **Chapter 18, "TrustSec and MACSec,"** discusses the concepts and use of security group tags (SGTs), as well as the classification, propagation, and enforcement of those SGTs.
- **Chapter 19, "Posture Assessment,"** discusses endpoint compliance checking, the agents, and provisioning of the agents. The chapter dives into the posture policies themselves and integrating posture to the authorization policy.
- **Chapter 20, "Deploying Safely,"** examines a phased deployment approach that enables the administrator to implement ISE in the network environment in a safe and staged method using Monitor-Mode before moving a switch or location into Low-Impact Mode or Closed Mode.
- **Chapter 21, "ISE Scale and High Availability,"** describes how to configure ISE nodes in a distributed environment, installing ISE patches, using node groups, promotion of secondary to primary roles, and an introduction to the load-balancing of ISE PSNs.
- **Chapter 22, "Troubleshooting Tools,"** extends the validation and troubleshooting lessons learned throughout the book by describing and discussing the many troubleshooting tools within ISE and the network devices themselves.
- **Chapter 23, "Final Preparation,"** discusses the ways in which to prepare for the exam, from study methods to what to expect on exam day.

https://t.me/learningnets

Part I: The CCNP Certification

https://t.me/learningnets

Chapter 1. CCNP Security Certification

This chapter covers the following topics:

- CCNP Security Certification Overview
- Contents of the CCNP-Security SISAS Exam
- How to Take the SISAS Exam
- Who Should Take This Exam and Read This Book
- Format of the CCNP-Security SISAS Exam
- CCNP-Security SISAS 300-208 Official Certification Guide

The Cisco Certified Network Professional (CCNP) certification program has several technology tracks including Security, Routing and Switching, Data Center, Service Provider, Service Provider Operations, Voice, and—last but not least—Wireless. This book focuses on one of the four exams required to achieve your CCNP-Security certification: Implementing Cisco Secure Access Solutions (300-208 SISAS).

You might already have other Cisco certifications in other networking technologies, or this may be your first foray into the Cisco certification process. You might instead be reading this book to enrich your skill set 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 base 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 result in costly network downtime.

The CCNP Security certification track provides a solid basis in four key Cisco security technologies—Access Solutions, Firewall Solutions (IOS and ASA), Virtual Private Network (VPN) Solutions, and Threat Control Solutions. As highlighted previously, the focus of this book is on the implementation of Secure Access Solutions (Cisco Certification 300-208 SISAS). Table 1-1 lists the four exams required to receive the CCNP-Security Certification.

Required Exam(s)	Recommended Training
300-206 SENSS	Implementing Cisco Edge Network Security Solutions (SENSS)
300-207 SITCS	Implementing Cisco Threat Control Solutions (SITCS)
300-208 SISAS	Implementing Cisco Secure Access Solutions (SISAS)
300-209 SIMOS	Implementing Cisco Secure Mobility Solutions (SIMOS)

Table 1-1 CCNP-Security Exam Technologies

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 SISAS Exam

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 covered on the exam. By diving deeper into each of these blueprint topics, you will become better prepared for your certification exam.

https://t.me/learningnets

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. If you would like to jump directly to the CCNP-Security track, you can leverage the former name of the CCNP-Security—CCSP (Cisco Certified Security Professional). The link to go directly to the CCNP-Security certification track is http://www.cisco.com/go/ccsp.

To drill down specifically to the SISAS exam blueprint, click the link under Exams and Recommended Training corresponding to the SISAS exam. On this page, you will find a number of tabs that provide high-level descriptions of the SISAS exam, exam topics, recommended training, as well as additional resources. As you review the blueprint (under Exam Topics) and other content pertaining to the SISAS exam, you might find that many topics overlap with other Cisco certifications—namely, CCNA and CCNA-Security certifications. You can 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 SISAS exam are provided in Table 1-2.

https://t.me/learningnets

Certification Guide Chapter	Exam Topic
Identity Management and Secure Access	
8, 9	Implement device administration
3, 10	Describe identity management
4, 12	Implement wired/wireless 802.1X
5, 6	Implement MAB
11	Implement network authorization enforcement
5, 13	Implement Central Web Authentication (CWA)
15	Implement profiling
14	Implement guest services
6, 19	Implement posture services
17	Implement BYOD access
Threat Defense	
18	Describe TrustSec Architecture
Troubleshooting, Monitoring, and Reporting Tools	
22	Troubleshoot identity management solutions
Threat Defense Architectures	
12	Design highly secure wireless solution with ISE
Design Identity Management Architectures	
7, 8, 9, 20, 21	Device Administration
3, 10, 16	Identity Management
15	Profiling
14	Guest Services
6, 19	Posturing Services
17	BYOD Access

Table 1-2 CCNP-Security SISAS Exam (300-208) Topics

Besides the training resources provided on the SISAS exam page, you also can find study resources at the links provided in Table 1-3. Other unofficial texts, video, and online training resources can also be found via your favorite online search engine.

https://t.me/learningnets

Resource	URL
The Cisco Learning Network SISAS page	https://learningnetwork.cisco.com/community/certifications/ccnpsecurity/sisas
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-9781587143250
BYOD Networking LiveLessons	http://www.ciscopress.com/store/cisco-bring-your-own-device-byod-networking-livelessons-9781587144219

Table 1-3 Additional Training Resources

How to Take the SISAS Exam

To take the CCNP-Security SISAS Exam, browse to https://www.cisco.com/go/ccsp and click the link for the SISAS certification. You will find information about the exam including the languages in which the exam is offered, the duration of the exam, and a link to register for the exam. At the time of publication of this book, the only approved testing vendor for the SISAS exam is Pearson VUE (www.vue.com). To register, click the Pearson VUE link, create an account, and register for the 300-208 SISAS 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 SISAS 300-208 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 be used either as the totality of the study material or to supplement other study resources (other texts, videos, instructor-led training, online training, and so on). Whether you are participating in formalized training for the SISAS 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 jobs or their intended career paths. Sometimes, the CCNP-level exams are the pinnacle of an individual's intended training—once his 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 minors 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 Ph.D of the specific technology. See Figure 1-1 and Table 1-4.

https://t.me/learningnets

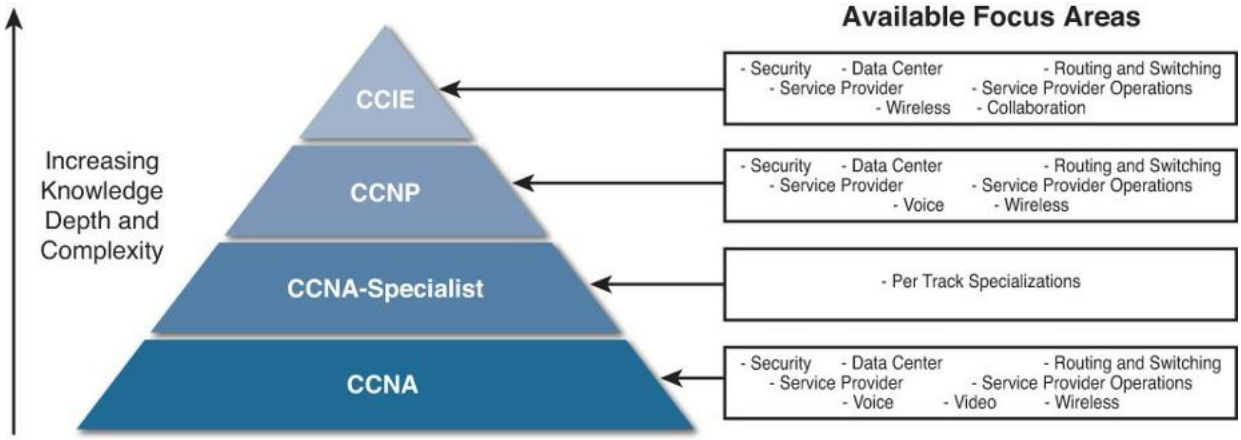


Figure 1-1 Cisco certification hierarchy.

Certification Name	Years of Experience	Job Role	Product/ Technology	Number of Exams	Prerequisites
CCNA Security	1–3	Network Security Specialist, Security Administrator, Network Security Support Engineer	Security Management, Cisco routers, Cisco IOS	1	Cisco CCENT, CCNA Routing and Switching, or any CCIE certification can act as a prerequisite.
CCNP Security	3–5	Network Security Engineer	Cisco IOS, Cisco routers and switches, ASA, security management, CSM, and firewall for IPS management	4	CCNA Security or any CCIE certification can act as a prerequisite.
CCIE Security	7+	Network Security Engineer	Cisco Adaptive Security Appliance (ASA) Firewall, Cisco Intrusion Prevention System (IPS), Cisco IOS Firewall, Cisco IOS IPS, Cisco Easy VPN, Cisco SSL VPN, Cisco Secure Desktop, Cisco Network Admission Control (NAC) Appliance, Cisco Security Agent, Cisco Secure ACS, Cisco Security Monitoring, Analysis and Response System, Cisco Catalyst 6500 Series Security	2	None.

https://t.me/learningnets

			Services Modules, Cisco Adaptive Security Device Manager, Cisco Router and Security Device Manager, Cisco Security Manager, Cisco IPS Device Manager, and Cisco IPS Manager Express		
Security Specialists					
Cisco ASA Specialist*	1–3	Security Engineers, Consultants, and Architects	Adaptive Security Appliance (ASA) technologies	2	Cisco CCNA Security certification.
Cisco Cybersecurity Specialist	1–3	Network Security Analyst	Cisco IOS, network security operations center (SOC) monitor security events, configure and tune security event detection and alarming, analyze traffic for security threats, respond appropriately to security incidents	1	None.
Cisco Firewall Security Specialist*	1–3	Network Security Engineers	Cisco ASA adaptive security appliance and Cisco IOS Software router- and switch-based networks	2	Cisco CCNA Security certification.

https://t.me/learningnets

Cisco IOS Security Specialist*	1–3	Network Security Engineers	Cisco IOS Software router- and switch-based networks	1	Cisco CCNA Security certification.
Cisco IPS Specialist*	1–3	Network Security Engineers	Cisco IPS sensor, Cisco IPS data management and analysis solution, and IPS policy virtualization	1	Cisco CCNA Security certification.
Cisco Network Admission Control Specialist	1–3	Network Security Engineers	Cisco Secure ACS; Cisco IOS routers, switches, and access points; and Trend Micro policy server	1	Cisco CCNA Security certification.
Cisco VPN Security Specialist*	1–3	Network Security Engineers	Cisco IOS Software Threat Control and Containment, Site-to-Site Transmission Security, and deploying Secure Remote Access	2	Cisco CCNA Security certification.

Table 1-4 Security Certification Comparison Chart

Format of the CCNP-Security SISAS Exam

If you have taken other Cisco certification exams, this exam format will not be much different. After registering for the SISAS exam, you will have a date and location for taking 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 form 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 cubicles to encourage privacy and 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). Often, 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, this will likely be easy to navigate.

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

- Multiple-choice
  - Single-answer
  - Multiple-answer
- Drag-and-drop
- Fill-in-the-blank
- Testlet
- Simlet

https://t.me/learningnets

- Simulated lab

Multiple-choice questions can take on one of two formats—single-answer and multiple-answer. With the single-answer, multiple-choice questions, you are given a question with several options for the correct answer. You are asked to select only one of the options using a round radio button to the left of the chosen answer—point your mouse icon at the radio button and left-click the mouse. For the multiple-answer, multiple-choice questions, you still are given a question with several options for the correct answer. However, you usually are asked to select a prescribed number of correct answers—for instance, "Choose 3." You select these using a square radio button to the left of the chosen answers. If you attempt to choose too many answers, you are prompted to choose only the prescribed amount.

Drag-and-drop questions test your ability to match or put into order a number of words/concepts. You 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 must match 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 expected to input the correct answer into the fill-in-the-blank box.

A testlet is a question in which a scenario is given. You are given multiple options to choose from to address the given scenario.

The simlet questions provide a simulated scenario. You are 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 Cisco devices interconnected in a simulated network. As part of this simulated lab question type, you are asked to configure the relevant network devices. You 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 must use the GUI to change the configuration and behavior of the affected device. If you normally use the command-line interface (CLI) to configure a device, the CLI can be the best way to configure the device during your exam. In this simulated lab environment, not all commands are available and the standard context-sensitive help available on Cisco routers and switches (the ? button) or tab-completion for commands might 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 similar to the format of CCNA-level tests. Examples of the question formats are 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 SISAS 300-208 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 you might 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.

The first section of the book is what you are reading right now; this is an overview of the CCNP-Security SISAS exam and everything that goes into it. Hopefully, you have a pretty good understanding as to what to expect as you schedule your exam.

In the second section of the book, the focus is on identity management and secure access. In this section, we discuss how to manage the users as well as how to allow them secure access to the network. This section presents the basis of authentication, authorization, and accounting (AAA). We cover the management of users, leveraging the internal user database of Cisco's Identity Services Engine (ISE), as well as third-

https://t.me/learningnets

party enterprise databases. The verification of the user via one of these databases—internal or external—is called authentication.

You can use a number of methods to authenticate a user when she is joining the network. We cover several of these authentication methods and the underlying protocols in this section of this book. We discuss how to authenticate a wired and wireless user using 802.1X, MAC Authentication Bypass, as well as nonstandard flows such as local and centralized web authentication.

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

We round out this section of the book by discussing some advanced concepts and diving deeper into some of the details of how ISE and the supporting network infrastructure accomplishes what it needs to accomplish. By the end of this first section, you should have a good overview of the end-to-end AAA process.

The third section of the book focuses on Cisco's Identity Services Engine (ISE) and its configuration. We discuss the specific roles that each persona plays in the ISE architecture and several common deployment scenarios. After this overview of ISE architecture, we 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 start to affect the policy on ISE for both authentication and authorization. We walk you step-by-step through how ISE is configured for authentication and authorization policies—highlighting all 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 you enforce the level of access can change. For instance, the enforcement mechanisms (VLANs, access control lists, Security Group Access, and so on) can 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, you can leverage ISE to granularly apply differentiated access to each endpoint individually.

The fourth section of this book moves 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 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 finish off the fourth section of this book 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, by providing either basic Internet access to employees or access to vendors and visitors. Profiling is a process whereby ISE can make an intelligent guess as to which type of device is joining the network, making granular authorization decisions based on device type. By the end of this fourth section, you should have a 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 fifth section of this book, we 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 this section, we incorporate certificate-based user authentication—authenticating a user based on an X.509 certificate, issued by either ISE or 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 covered in this section of the book. BYOD is a process and security infrastructure that allows a user to bring his 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. A number of special portals and configurations 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 several key

https://t.me/learningnets

factors.

The next advanced topic reviewed in this section is TrustSec and MACSec. We provide 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 we address in this section 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 the type of device that is trying to join the network, without ever actively probing the device. Posturing is a little more entrenched at the client/endpoint level. Posturing leverages information contained deep in the configuration of the endpoint, requiring a posturing agent to be run on the endpoint. After key information is read from the endpoint via this agent, the ISE makes 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.

The sixth section of this book is geared toward the operational aspects of having ISE. As part of this chapter, we 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 a log is thrown if the user doesn't match an available policy. This enables network administrators to fully discover and understand the endpoints on their networks without having an adverse effect on users. After 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 the sixth section is ISE scale and high availability. This section highlights how to configure and deploy a distributed ISE architecture to accommodate additional load, demand, and possible additional features. Each instance of ISE has an upper limit based on the platform and particular software on which it is running. 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 the sixth section we provide the reader 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 the necessary information to isolate any quality or network issues.

In the final section, section seven, we describe the steps that you'll need to take to prepare for the CCNP-Security SISAS.

Book Features and Exam Preparation 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. Therefore, this book does not try to help you pass the exams only by memorization, but by truly learning and understanding the topics.

The book includes many features that provide different ways to study to be ready for the exam. If you understand a topic when you read it but do not study it any further, you will probably not be ready to pass the exam with confidence. The features included in this book give you tools that help you determine what you know, review what you know, better learn what you don't know, and be well prepared for the exam. These tools include

- **"Do I Know This Already?" Quizzes**—Each chapter begins with a quiz that helps you determine the amount of time you need to spend studying that chapter.
- **Foundation Topics**—These are the core sections of each chapter. They explain the protocols, concepts, and configuration for the topics in that chapter.
- **Exam Preparation Tasks**—The Exam Preparation Tasks section lists a series of study activities that should be done after reading the Foundation Topics section. Each chapter has the activities that make the most sense for studying the topics in that chapter. The activities include
  - **Planning Tables**—The SISAS exam topics include some perspectives on how an engineer plans

https://t.me/learningnets

for various tasks. The idea is that the CCNP-level engineer in particular takes the design from another engineer, plans the implementation, and plans the verification steps, handing off the actual tasks to engineers working during change-window hours. Because the engineer plans the tasks but might not be at the keyboard when implementing a feature, that engineer must master the configuration and verification commands so that the planned commands work for the engineer making the changes off-shift. The planning tables at the end of the chapter give you the chance to take the details in the Foundation Topics core of the chapter and think about them as if you were writing the planning documents.

- **Key Topics Review**—The Key Topic icon is shown next to the most important items in the Foundation Topics section of the chapter. The Key Topics Review activity lists the key topics from the chapter and their corresponding page numbers. Although the contents of the entire chapter could be on the exam, you should know the information listed in each key topic. Review these topics carefully.
- **Memory Tables**—To help you exercise your memory and memorize some lists of facts, many of the more important lists and tables from the chapter are included in a document on the CD. This document lists only partial information, allowing you to complete the table or list. CD-only Appendix D holds the incomplete tables, and Appendix E includes the completed tables from which you can check your work.
- **Definition of Key Terms**—Although Cisco exams might be unlikely to ask a question such as "Define this term," the SISAS exam requires that you learn and know a lot of networking terminology. This section lists some of the most important terms from the chapter, asking you to write a short definition and compare your answer to the Glossary on the enclosed CD.

- **CD-based practice exam**—The companion CD contains an exam engine, including a bank of multiple-choice questions. Chapter 23, "Final Preparation," gives two suggestions on how to use these questions: either as study questions or to simulate the SISAS exam.
- **Companion website**—The website (http://www.ciscopress.com/title/) posts up-to-the-minute materials that further clarify complex exam topics. Check this site regularly for new and updated postings written by the author that provide further insight into the more troublesome topics on the exam.

https://t.me/learningnets

Part II: “The Triple A” (Authentication, Authorization, and Accounting)

https://t.me/learningnets

Chapter 2. Fundamentals of AAA

This chapter covers the following exam topics:

- **Compare and Select AAA Options**
- **TACACS+**
- **RADIUS**
- **Describe Features and Functionality of Authentication and Authorization**
- **Describe RADIUS Flows**
- **AV Pairs**
- **Device Administration**

In the world of security, we can only be as secure as our controls permit us to be. There are laws in the United States defining what a passenger of an airplane is permitted to bring onboard. If the 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), then how would the FAA ever really enforce those policies?

With technology, we are faced with the same challenges. We need to have controls in place to ensure that only the correct entities are using our technological "gadgets" while getting in the way of the user as little as possible. The same concepts can be applied to many use cases, including human interaction with a computer, a computer's interaction with a network, and even an application's interaction with data.



This security principle is known as authentication, authorization, and accounting (AAA), often referred to as Triple-A.

Before allowing an entity to perform certain actions, you must ensure you know who that entity actually is (authentication) and whether the entity is authorized to perform that action (authorization). Additionally, you need to ensure that accurate records are maintained showing that the action has occurred, so you should keep a security log of the events (accounting).

The concepts of AAA can be applied to many aspects of a technology lifecycle. However, this exam, and therefore this book, will focus on the two main aspects of AAA related to networking:

1. **Device administration**—Controlling access to who can log in to a network device console, telnet session, secure shell (SSH) session, or other method is one form of AAA that you should be aware of. This is AAA for device administration, and although it can often seem similar to network access AAA, it has a completely different purpose and requires different policy constructs.
2. **Network access**—Securing network access can provide the identity of the device or user before permitting the entity to communicate with the network. This is AAA for network access and is the type of AAA that is most focused on in this exam and book.

"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."

https://t.me/learningnets

Foundation Topics Section	Questions
Compare and Select AAA Options	1-8
TACACS+	3, 6-8
RADIUS	4, 5, 7
Describe Features and Functionality of Authentication and Authorization	1-2
Describe RADIUS Flows	9
AV Pairs	10
Device Administration	2

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

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, whereas an authentication determines what devices the user can interact with.
   - **c.** Authentication is used with both network access and device administration, whereas authorization applies only to device administration.
   - **d.** Authentication validates the user's identity, whereas authorization determines what that user is permitted to do.
2. Which of the following are types of AAA as related to the topics of this exam? (Select 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+

https://t.me/learningnets

b. CHAP

c. RADIUS

d. MS-CHAPv2

5. True or False? RADIUS can be used for device administration AAA.

a. True

b. False

6. Which of the following Cisco products should be used for device administration with TACACS+?

a. Cisco Secure Access Control Server (ACS)

b. Cisco Identity Services Engine

c. Cisco TACACS+ Control Server (TCS)

d. Cisco Centri

7. Why is RADIUS or TACACS+ needed? Why can't the end user authenticate directly to the authentication server?

a. The added level of complexity helps Cisco and other vendors to sell more products.

b. Because the names sound so cool.

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 are TACACS+ messages sent from the AAA client to the AAA server? (Select all that apply.)

a. START

b. REPLY

c. CHALLENGE

d. REQUEST

9. When using RADIUS, what tells the AAA server which type of action is being authenticated?

a. The TACACS+ service.

b. The Service-Type field.

c. RADIUS does not distinguish between 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 likes to throw in terms to confuse the reader.

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) for audio and video traffic.

https://t.me/learningnets

Foundation Topics

Triple-A

Simply put, authentication is the validation of 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 networking 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 verify the identity of the card user is authentication. Ensuring that the identity matches the name printed on the card is an authorization. Think about it: What if Kevin Redmon were to go into a retail store and hand over a credit card to pay for the $10,000 of electronics he is purchasing? He passes his driver's license to the cashier, who verifies that he and the picture on the license match. It is certainly his identity. However, the name printed on the credit card is Aaron Woland. Should the transaction go through? Of course not.

Kevin's attempt to use Aaron's credit card is now 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 Cisco Secure Access Solutions exam to understand the differences between and purposes of all three A's in Triple-A.

Compare and Select AAA Options

As previously described in this chapter, there are two uses of AAA that you will focus on for this exam and this book: device administration and network access.

Device Administration



Device administration is a method of AAA for controlling access to a network device console, telnet session, SSH session, or other method of accessing the device's operating system itself for purposes of configuration. For example, imagine your company has an Active Directory group named Network Administrators, which should have full access (privilege level 15) to the Cisco Switches in a company's network. They should therefore be able to make changes to virtual local area networks (VLANs), see the entire running configuration of the device, and more.

There could be another group named Network Operators who should be allowed to only view the output of show commands and not allowed to configure anything in the device.

Device administration AAA allows you to get more granular. Cisco Secure Access Control Server (ACS) has a capability to provide command sets, which are listings of commands that are permitted or denied to be run by an authenticated user. In other words, a user can authenticate to the Cisco IOS shell, and ACS can permit or deny his execution of individual commands, if you choose.

Figure 2-1 illustrates device administration.

https://t.me/learningnets

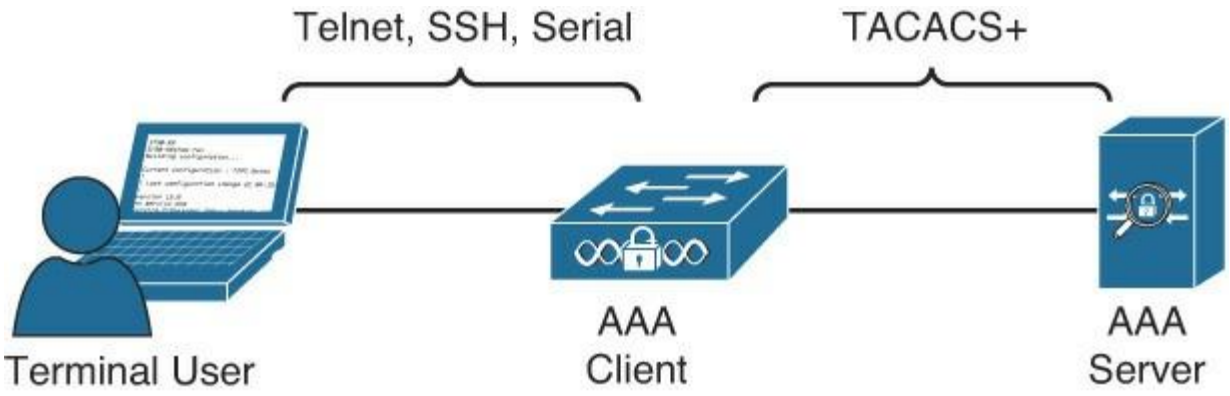


Figure 2-1 Device administration.

Device administration can be very interactive in nature, with the need to authenticate once but authorize many times during a single administrative session in the command line of a device. As such, it lends itself well to using the Terminal Access Controller Access Control System (TACACS) client/server protocol, more so than Remote Authentication Dial-in User Service (RADIUS).

As the name describes, TACACS was designed for device administration AAA to authenticate and authorize users for mainframes, Unix terminals, and other consoles.

Both the TACACS and RADIUS protocols are discussed in more depth within this chapter; however, TACACS separates out the authorization portion of AAA, allowing for a single authentication and multiple authorizations within the same session. This is why it lends itself to device administration more than RADIUS. RADIUS does not provide the ability to control which commands can be executed.

To learn more about device administration AAA, beyond the knowledge necessary to pass the SISAS exam, take a look at the Cisco Press book AAA Identity Management Security (ISBN-10 1-58714-144-2).

Network Access



Secure network access is essentially all 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 in this exam and book. Network access AAA really took a strong hold back in the day of modems and dial-up networking with plain old telephone service (POTS). In those days companies provided network access to workers from outside the physical boundaries of the company's buildings with the use of modems. People gained Internet access by using dial-up to Internet service providers (ISPs) over their modems, as well. Basically, all that was needed was a modem and a phone line.

Allowing anyone to dial in to your network just by dialing your modem's phone number is not a secure idea. The user needs to be authenticated before allowing the connection. That is where RADIUS came into play originally, as is evident in the name of the protocol (Remote Access Dial In User Service). RADIUS was used between the network access device (NAD) and the authentication server. The authentication was normally Password Authentication Protocol (PAP), Challenge/Handshake Authentication Protocol (CHAP), or Microsoft CHAP (MS-CHAP).

Figure 2-2 illustrates dial-up remote access.

https://t.me/learningnets

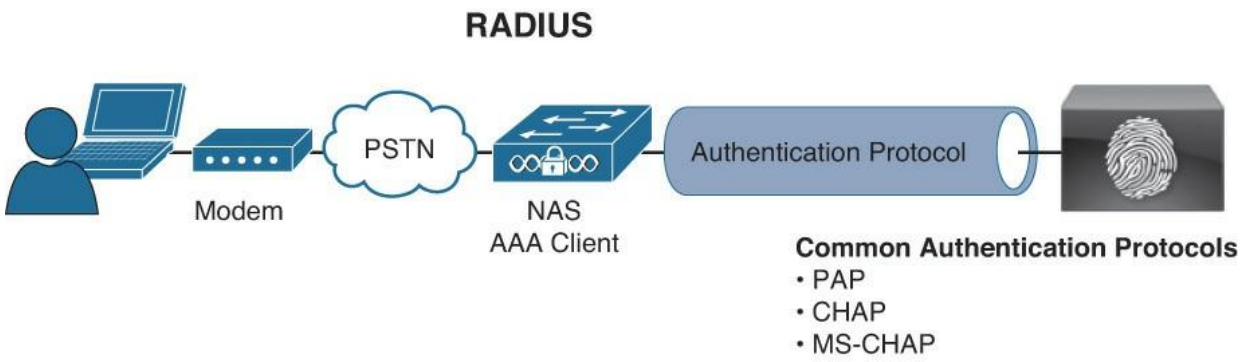


Figure 2-2 Dial-Up remote access.

As technology continued to evolve and Wi-Fi became prevalent, direct dial-in to a company was replaced by Remote Access virtual private networks (VPNs). The IEEE standardized on a method to use Extensible Authentication Protocol (EAP) over local area networks (IEEE 802.1X), and RADIUS was used as the protocol of choice 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, also can be used with 802.1X. However, it is mostly found in the service provider space and is not covered as part of this exam.

RADIUS is the protocol used almost exclusively with network access AAA.

TACACS+

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. Although TACACS+ is mainly used for device administration AAA, you can use it for some types of network access AAA.



It is important to understand that TACACS+ is not currently a supported protocol with Cisco Identity Services Engine (ISE) 1.2. The Cisco Secure Access Control Server (ACS) is Cisco's primary authentication server product for enterprises that have a need to use TACACS+ for device administration AAA.

> **Note**
>
> Other products from Cisco support TACACS+, such as the Cisco Access Registrar solution. However, those solutions are geared toward service providers and are not covered in the exam or this book.

TACACS+ uses Transmission Control Protocol (TCP) port 49 to communicate between the TACACS+ client and the TACACS+ server. An example is a Cisco switch authenticating and authorizing administrative access to the switch's IOS CLI. The switch is the TACACS+ client, and Cisco Secure ACS is the server, as illustrated in Figure 2-3.

https://t.me/learningnets

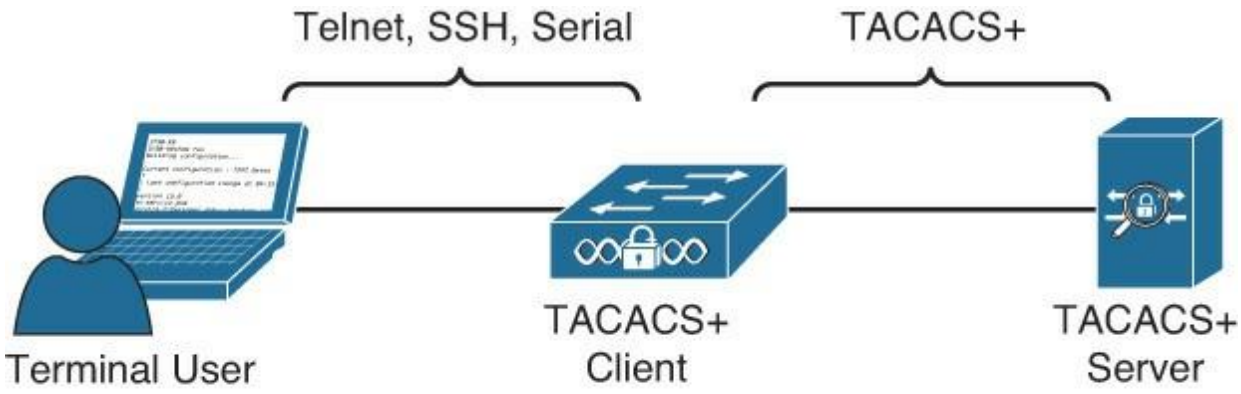


Figure 2-3 TACACS+ client server communication.

One of the key differentiators of TACACS+ is its capability 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 interactive in nature, with 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 might need to authorize a user's activity on a per-command basis. TACACS+ is designed to accommodate that type of authorization. As the name describes, 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 can be used for authentication than are used for authorization and accounting. Another interesting point to know, especially for the purposes of the SISAS exam, is that TACACS+ communication will encrypt the entire payload.

TACACS+ Authentication Messages

When using TACACS+ for authentication, only three types of packets are 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**—Messages sent from the AAA server to the AAA client.
- **CONTINUE**—Messages from the AAA client used to respond to the AAA server requests for username and password.

The following paragraphs describe the authentication flow process and the messages used.

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) will be a REPLY during authentication. The server sends a REPLY message asking the client to retrieve the username. The username is sent to the server within a CONTINUE message.

After 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 another CONTINUE message. The server then sends a final REPLY message with the pass or fail status of the authentication request.

The possible values returned from the AAA server to the AAA client within the final REPLY message are

- **ACCEPT**—The user authentication succeeded and the authorization process may begin, if the AAA client is configured for authorization.
- **REJECT**—The user authentication has failed. The login will be denied or the end user will be prompted to try again, depending on the configuration of the AAA client.

https://t.me/learningnets

- **ERROR**—An error occurred at some point during the authentication. AAA clients will 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 is not to be confused with the CONTINUE message sent from the AAA client to the AAA server. This value is sent from the AAA server within a REPLY message, indicating that more information is required.

Figure 2-4 illustrates the authentication messages between the client and server.

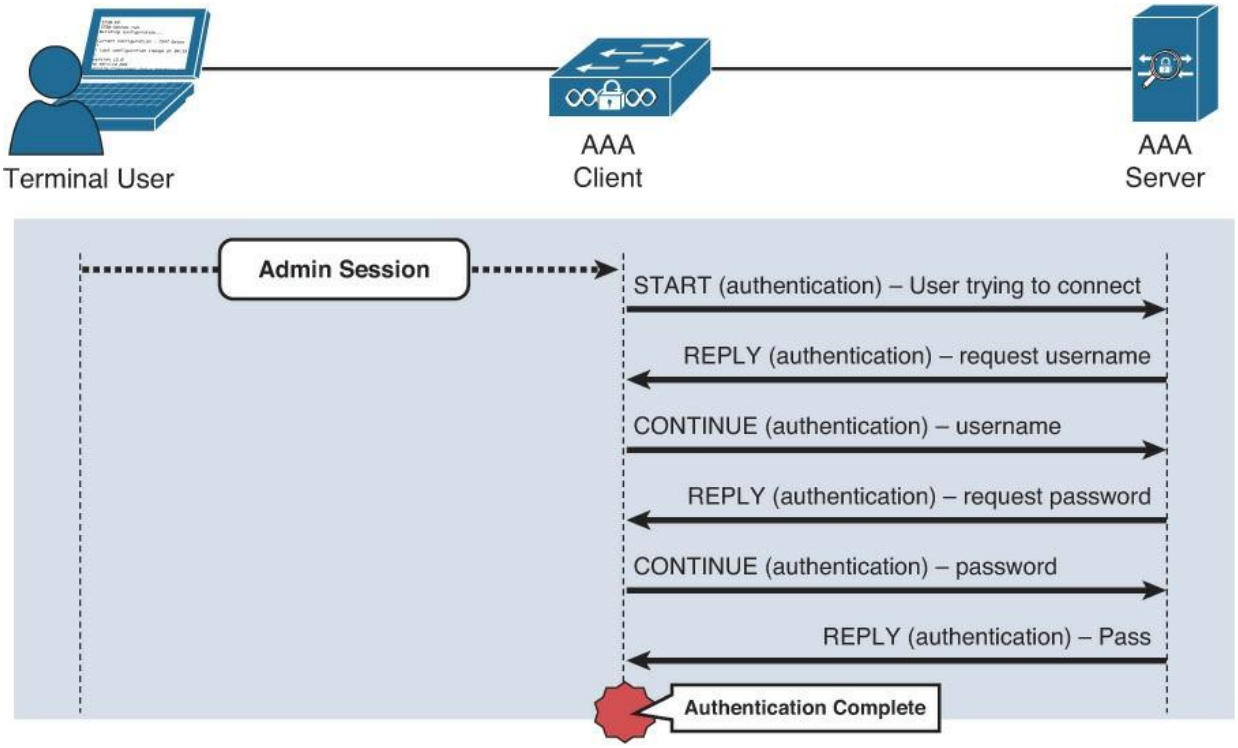


Figure 2-4 TACACS+ authentication communication flows.

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 an authorization. The authorization can 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. You can find more about specific TACACS+ AV pairs at http://bit.ly/1mF27aT.
- **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. RESPONSE messages can contain one of five replies:
  - **FAIL**—This response indicates the user should be denied access to the requested service.
  - **PASS_ADD**—This response indicates a 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, the request is simply authorized as-is.
  - **PASS_REPL**—Indicates a successful authorization, but the server has chosen to ignore the REQUEST and is replacing it with the information sent back in the RESPONSE.
  - **FOLLOW**—This reply indicates that the AAA server wants the AAA client to send the

https://t.me/learningnets

authorization request to a different server. The new server information will be listed in the RESPONSE packet. The AAA client can use that new server or treat the response as a FAIL.

- **ERROR**—A reply of ERROR indicates a problem occurring on the AAA server, and 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 like authorization messages, only two message types are used in accounting:

- **REQUEST**—This message is sent from the AAA client to the AAA server to indicate a notification of activity. Three values can be included with the REQUEST:
  - **START**—A start record indicates that a service has begun.
  - **STOP**—The stop record indicates that the service has ended.
  - **CONTINUE**—The continue record is also sometimes referred to as a Watchdog or UPDATE record. This is sent when a service has already started and is in progress, but there is updated information to provide in relation to the service.
- **RESPONSE**—This message is sent from the AAA server back to the AAA client with the result of the accounting REQUEST and can contain one of three replies:
  - **SUCCESS**—Indicates that the server received the record from the client
  - **ERROR**—Indicates an error on the server and that the record was not stored
  - **FOLLOW**—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

Figure 2-5 illustrates an end user being authorized to access the IOS exec CLI. The figure is a direct continuation of Figure 2-4 where the authentication occurred. In this illustration, the end user gets authorized to enter the IOS exec and is authorized to run the show run command. The command the user is requesting to use is contained within the REQUEST (authorization) message.

https://t.me/learningnets

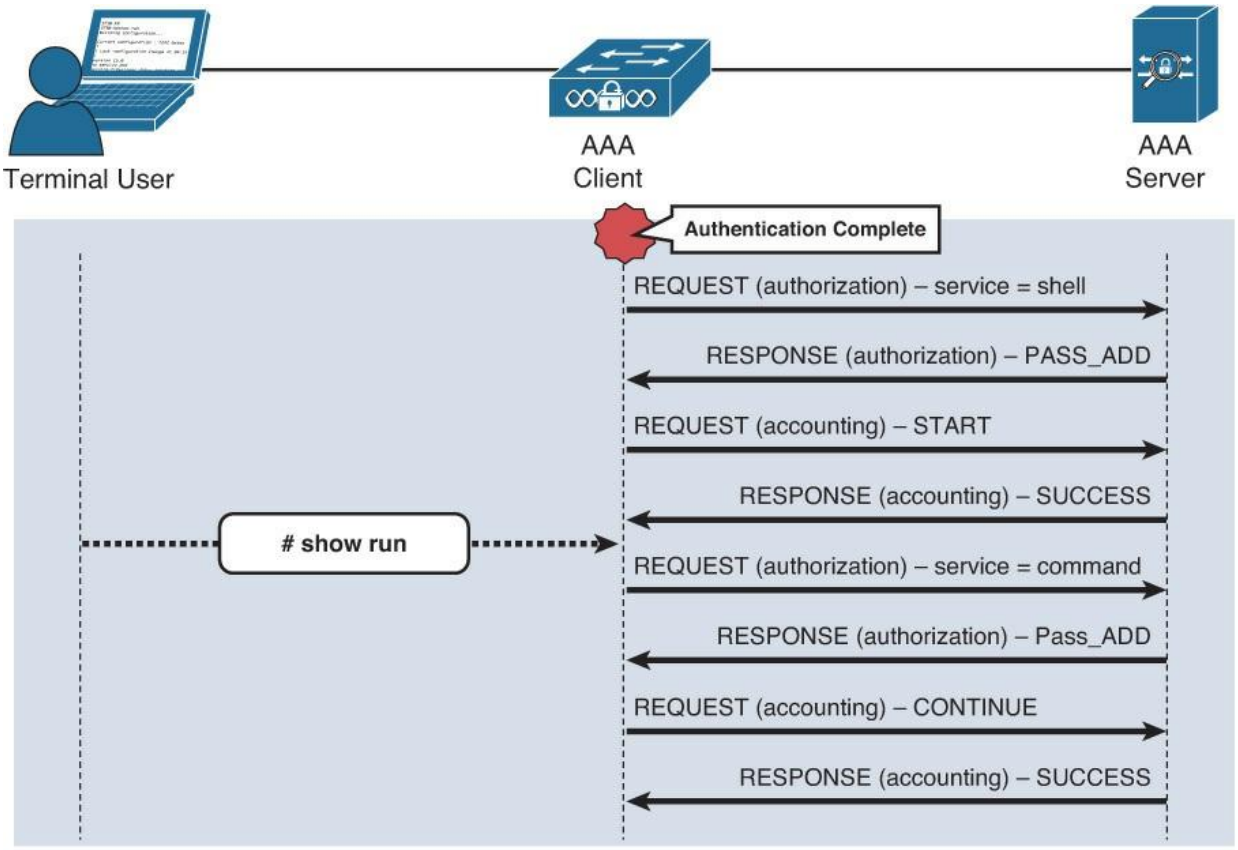


Figure 2-5 TACACS+ authorization and accounting communication flows.

This concludes the material related to TACACS+ for the purposes of the SISAS exam. For more information on TACACS+, take a look at the Cisco Press book AAA Identity Management Security (ISBN-10 1-58714-1442).

RADIUS



RADIUS is an IETF standard for AAA. As with TACACS+, RADIUS follows a client/server model in which the client initiates the requests to the server. RADIUS is the protocol of choice for network access AAA, and it's time to get very familiar with RADIUS. If you connect to a secure wireless network regularly, RADIUS is most likely being used between the wireless device and the AAA server. Why? Because RADIUS is the transport protocol for EAP, along with many other authentication protocols.

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

As described 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 2-6.

https://t.me/learningnets

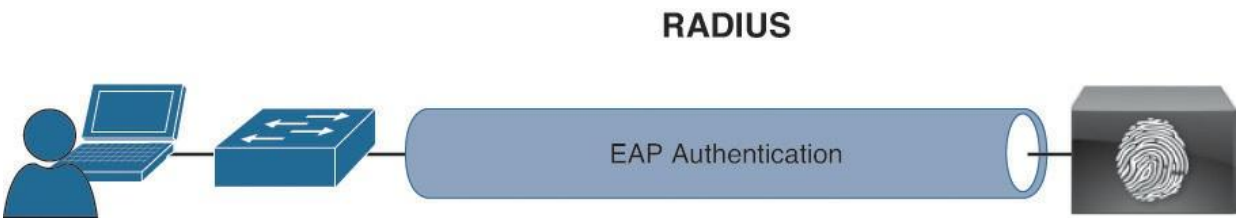


Figure 2-6 RADIUS carries the Layer-2 EAP communication.

There are many differences between RADIUS and TACACS+. One such difference is that authentication and authorization are not separated in a RADIUS transaction. When the authentication request is sent to an 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 an authentication and authorization. The request could be for network access or for device shell access—RADIUS does not discriminate. The function requested is known as a service type. For example, the service type might be "framed" for an IEEE 802.1X authentication. Some common RADIUS service types are shown in Table 2-2.

Value	Service Type Name*	Commonly Used for
1	Login	Login request; often used with web authentications for non-Cisco network equipment
2	Framed	IEEE 802.1X
5	Outbound	Used for local web authentication
10	Call-Check	Used for MAC Authentication Bypass (MAB)

*You can find a more complete listing of service-types at http://bit.ly/1CGDE8Y.

Table 2-2 Common RADIUS Service Types

- **Access-Accept**—Sent from the AAA server to the AAA client signaling a passed authentication. The authorization result will be included as AV-pairs. The AV-pairs can include items such as the assigned VLAN, a downloadable access control list (dACL), a security group tag (SGT), and much more.
- **Access-Reject**—Sent from the AAA server to the AAA client signaling the authentication failure. The failed authentication also signifies that no authorization has been granted.
- **Access-Challenge**—This optional message can 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 2-7 illustrates a sample RADIUS flow.

https://t.me/learningnets

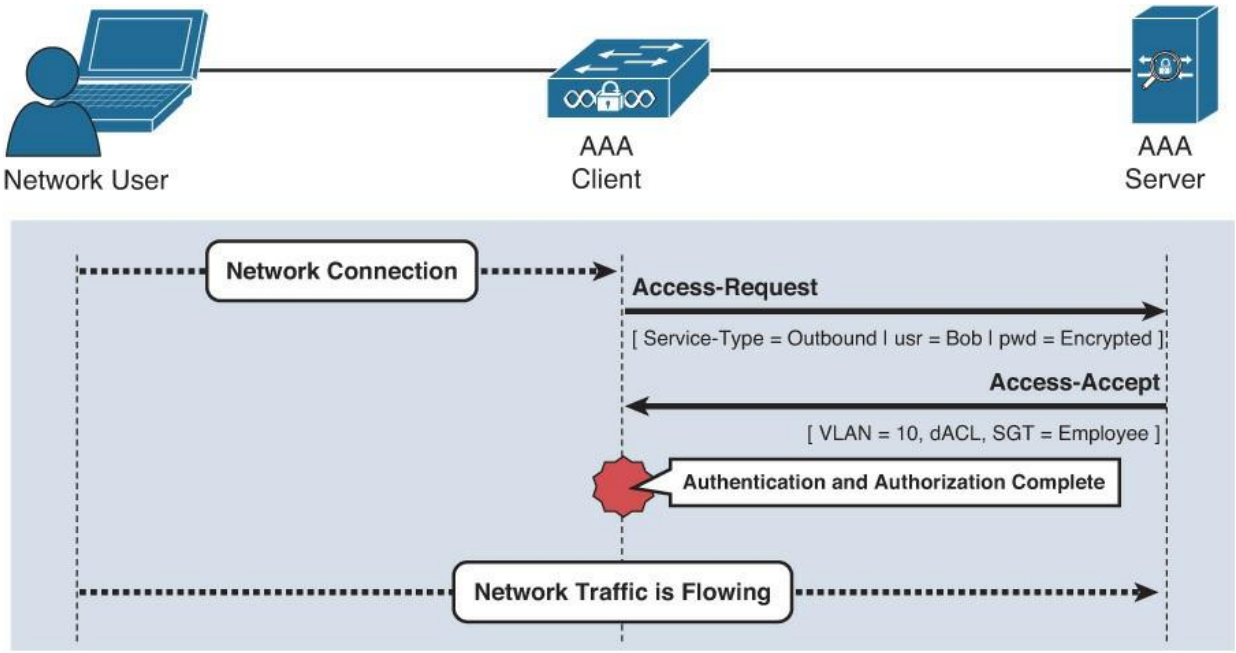


Figure 2-7 RADIUS authentication and authorization communication flows.

When looking at Figure 2-7, keep in mind that authentication and authorization are combined with RADIUS. The Access-Accept message includes the AV-pairs defining what the user is authorized to do.

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.

Only two message types are used in accounting:

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

Figure 2-8 illustrates a sample RADIUS accounting flow. The figure is a direct continuation of Figure 2-7 where the authentication and authorization occurred.

https://t.me/learningnets

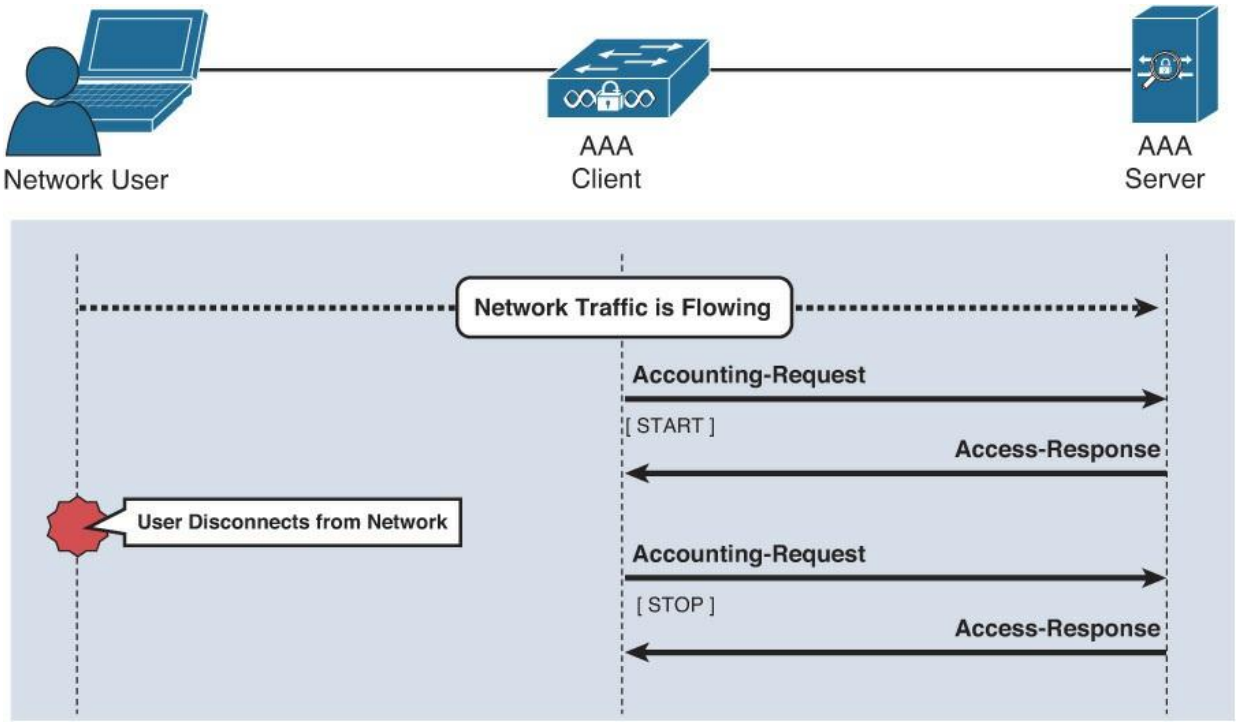


Figure 2-8 RADIUS accounting communication flows.

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

AV-Pairs



There are references to something called an attribute value pair (AV-pair) all through the TACACS+ and RADIUS sections. When communicating with an AAA protocol, many attributes can be referenced to clearly dictate answers or results. The RADIUS server might be assigning an attribute to the authentication session, such as a VLAN. The VLAN placeholder is the attribute, and the actual assigned VLAN number is the value for that placeholder. The placeholder and its value are paired together and referred to as AV-pairs.

Change of Authorization

Because RADIUS was always defined to be a client server architecture, with the client always initiating the conversation, it became challenging for the AAA server to take action. As RADIUS was defined, the AAA server could only assign an authorization as a result to an authentication request.

As technology advanced, many new demands appeared, including the ability for the network to kick out misbehaving clients, to quarantine them, or basically to just change their access.

How can that happen when the network access is using a RADIUS control plane and the AAA client must always initiate the RADIUS conversations? That is where RFC 3576 and its successor RFC 5176 come in. These RFCs define a new enhancement to RADIUS known as Dynamic Authorization Extensions to RADIUS or, as it is more commonly called, change of authorization (CoA).

https://t.me/learningnets



CoA is what enables a RADIUS server to initiate a conversation to the 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 also see how critically important CoA is. For more information regarding CoA, refer to Chapter 6, "Introduction to Advanced Concepts."

Comparing RADIUS and TACACS+

Table 2-3 is included as a summary to the chapter and to provide you with a reference table for your studies.

	RADIUS	TACACS+
Protocol and Port(s) Used	UDP: 1812 and 1813 -or- UDP: 1645 and 1646	TCP: 49
Encryption	Encrypts only the Password field	Encrypts the entire payload
Authentication and Authorization	Combines authentication and authorization	Separates authentication and authorization
Primary Use	Network access	Device administration

Table 2-3 Summary of the Two Main AAA Protocols: RADIUS and TACACS+

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 2-4 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	The security principle known as AAA	17
Section	Device administration	21
Section	Network access	22
Paragraphs	TACACS+	24
Section	RADIUS	28
Paragraph	AAA and AV-pairs	31
Paragraph	CoA and RADIUS	32

Table 2-4 Key Topics for Chapter 2

https://t.me/learningnets

Define Key Terms

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

device administration

network access

https://t.me/learningnets

Chapter 3. Identity Management

This chapter covers the following exam topics:

- **Identity management/secure access**
- **Describe identity management**
- **Describe identity store options (such as LDAP, AD, PKI, OTP, Smart Card, local)**
- **Describe native AD and LDAP**

Do you know what an identity is? You have an identity! I have an identity! What is the definition of an identity? How do we use identities in the design of your Identity Service Engine (ISE) infrastructure?

This chapter will answer these questions and introduce several key concepts regarding the use of ISE identity management, including key terminology such as identity store, identity store sequence, and internal and external identity sources.

ISE has multiple authentication configuration options that we will look at in this chapter, including Active Directory, LDAP, one-time token passwords, smart cards, and the internal ISE user database.

Additionally, we will have an in-depth discuss on 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 enables 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."

Foundation Topics Section	Questions
What is an identity?	1
What is an identity store?	2-6
Identity sources supported by ISE	7
Internal identity stores	8-9
External identity stores	10

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

> **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 are two types of identities used in Cisco Identity Service Engine?
   - **a.** SSID
   - **b.** MAC address
   - **c.** Username

https://t.me/learningnets

d. IP address

2. What are the two general types of identity stores used by Cisco ISE?

a. Temporary

b. External

c. Internal

d. Permanent

3. Cisco ISE internal identity stores are used to authentication which two of the following?

a. Endpoints

b. AD security groups

c. RADIUS

d. Users

4. Which identity store attributes can be used in an ISE authorization policy? (Choose two.)

a. User

b. Time

c. Accounting

d. Machine

5. What is an individual identity store called?

a. Authentication source

b. Identity database

c. Identity source

d. Authentication database

6. How is an identity source sequence processed?

a. Bottom to top

b. Left to right

c. Top to bottom

d. No particular order

7. Which of the following identity stores are supported by ISE for authentication? (Choose three.)

a. LDAP

b. TACACS

c. Microsoft Active Directory

d. RADIUS servers

8. Which of the following can be used with an internal identity store?

a. SSID

b. Guest login

c. Administration

d. MAB

9. What are the two types of internal identity stores used in ISE?

a. User database

b. Endpoint database

https://t.me/learningnets

c. System database

d. Admin database

10. What are the two primary reasons for using external identity stores?

a. Performance

b. Monitoring

c. Scalability

d. Management

Foundation Topics

What Is an Identity?

According to the Oxford English dictionary, an identity is "serving to establish who the holder, owner, or wearer is by bearing their name and often other details such as a signature or photograph."

Someone could walk up to you and say that her name is "Tammy," but how would you know that identity is true? You would need some type of "trusted" evidence such as a third-party form of identification—say, a driver's license or passport—to verify her credentials.

One of the first and foremost uses of identity in the world of secure network access control is the identity of a person or device. Typically, many factors 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 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 can be used to authenticate a user's or an endpoint's credentials. The identity store could be an internal database that resides on the Authenticaton Authorization and Accounting (AAA) server or even external database(s) housing the identities.

Identity stores can also be used for the retrieval of user or machine attributes used in authorization policies, which are covered in more detail in Chapter 11, "Authorization Policies."



Each individual identity store is also referred to as an identity source. You can combine multiple identity sources into an identity source sequence that can then be processed for user and endpoint authentications and authorizations. The identity source sequence will then process the identity stores in the order they are defined from top-to-bottom succession. See Figure 3-1 for an example of an ISE identity source sequence configuration.

https://t.me/learningnets

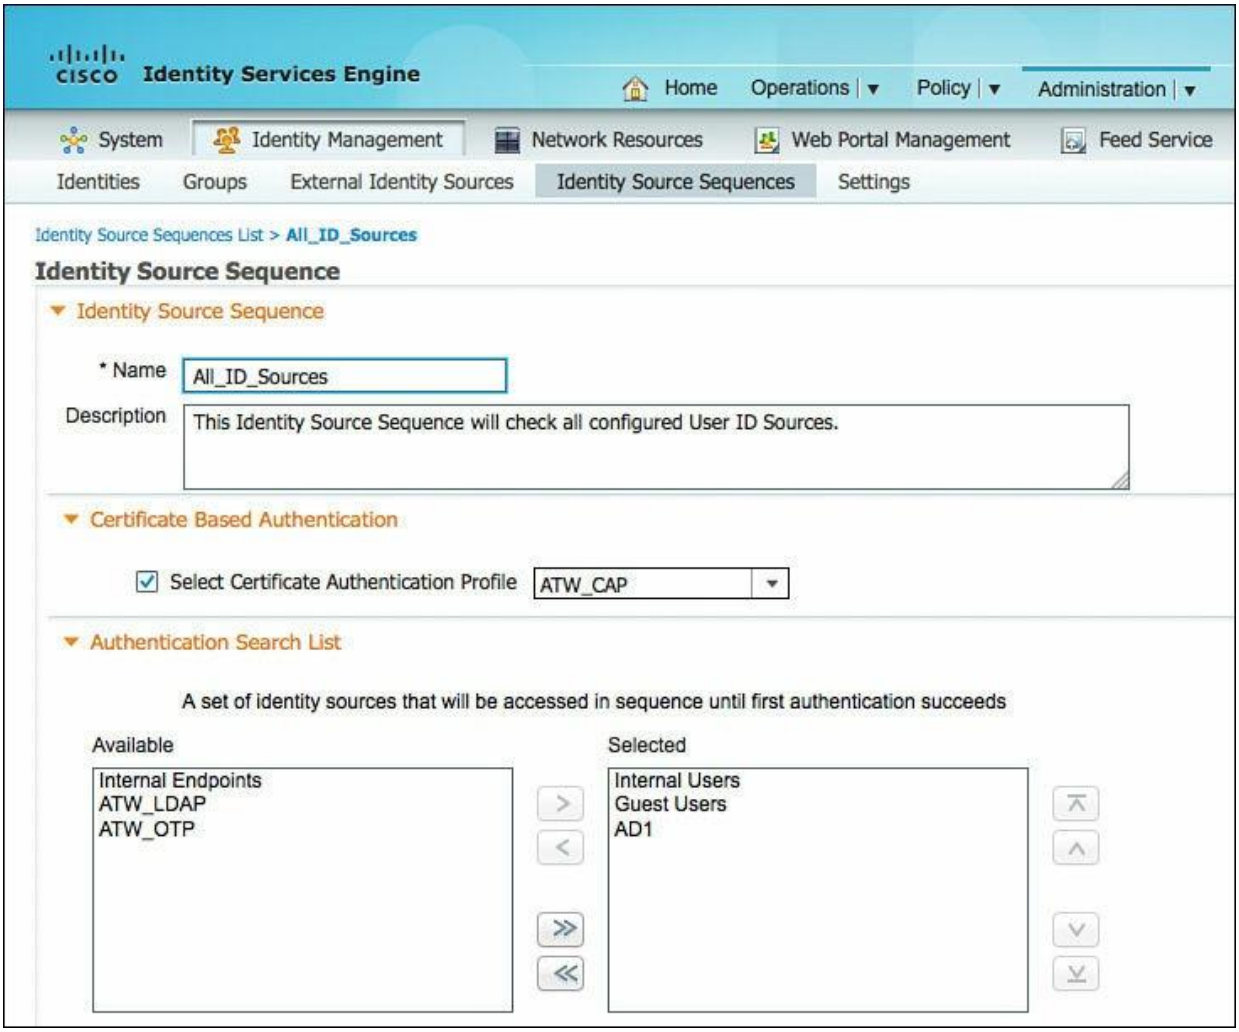


Figure 3-1 ISE identity source sequence configuration.

Both the Cisco Identity Services Engine and the Cisco Secure Access Control Server are capable of using identity sources inclusive of an internal identity database(s) as well as external identity databases. External identity sources include Microsoft Active Directory (AD), Lightweight Directory Access Protocol (LDAP), one-time password (OTP) servers, smart cards, and certificate authorities (CAs).

Internal Identity Stores

Cisco ISE has an internal user database that can be used for internal username and password accounts. These user accounts stored in the internal user database are referred to as internal users. The internal user database can be used as an internal identity store for local authentication and authorization policies.

One example for use of this internal identity store would be to authenticate and authorize a sponsored ISE guest accounts to access your corporation's guest wireless network. Cisco ISE maintains a separate guest identity store. Another example is the internal identity source for administrative accounts, which is another separate identity store.

A third example would be to create internal username accounts that could be used as sponsor accounts for the creation of ISE guest accounts. See Figure 3-2 for an example of an ISE internal user account configuration.

https://t.me/learningnets

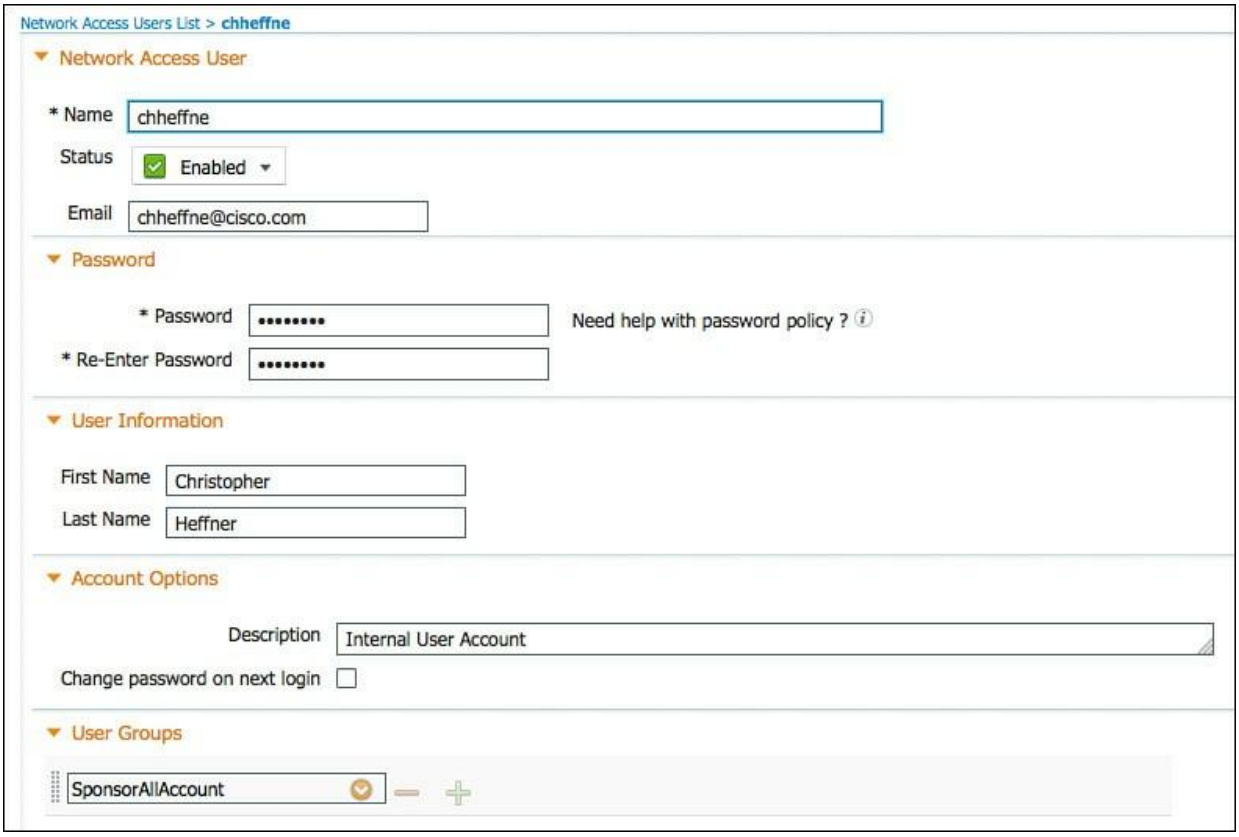


Figure 3-2 ISE internal user account configuration.

Any internal user accounts configured then are associated with a defined user identity group that sets the type of internal account being created. Several preconfigured user identity groups can be used when creating internal user accounts, including the following: Guest, Activated Guest, Employee, SponsorAllAccount, SponsorGroupAccounts, and SponsorOwnAccounts. See Figure 3-3 for an example of ISE user identity group options.

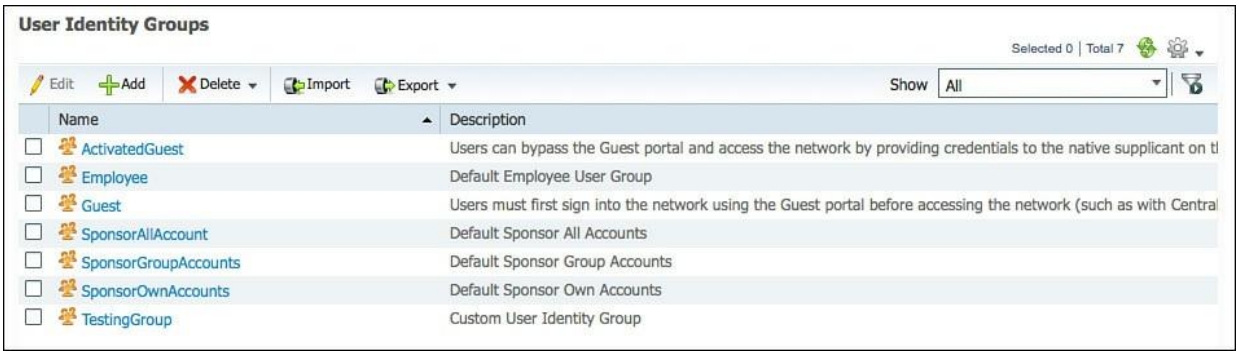


Figure 3-3 ISE user identity groups.

If you are using internal user account, make sure that the internal user's identity store is part of the identity source sequence. The identity source sequence can then be used in the ISE authentication and authorization policies.

The second type of internal identity store used in both Cisco ISE and Cisco ACS is called the internal endpoint database. It is used to store information about all endpoint devices that connect to the ISE infrastructure. See Figure 3-4 for an example of the ISE internal endpoint database.

https://t.me/learningnets

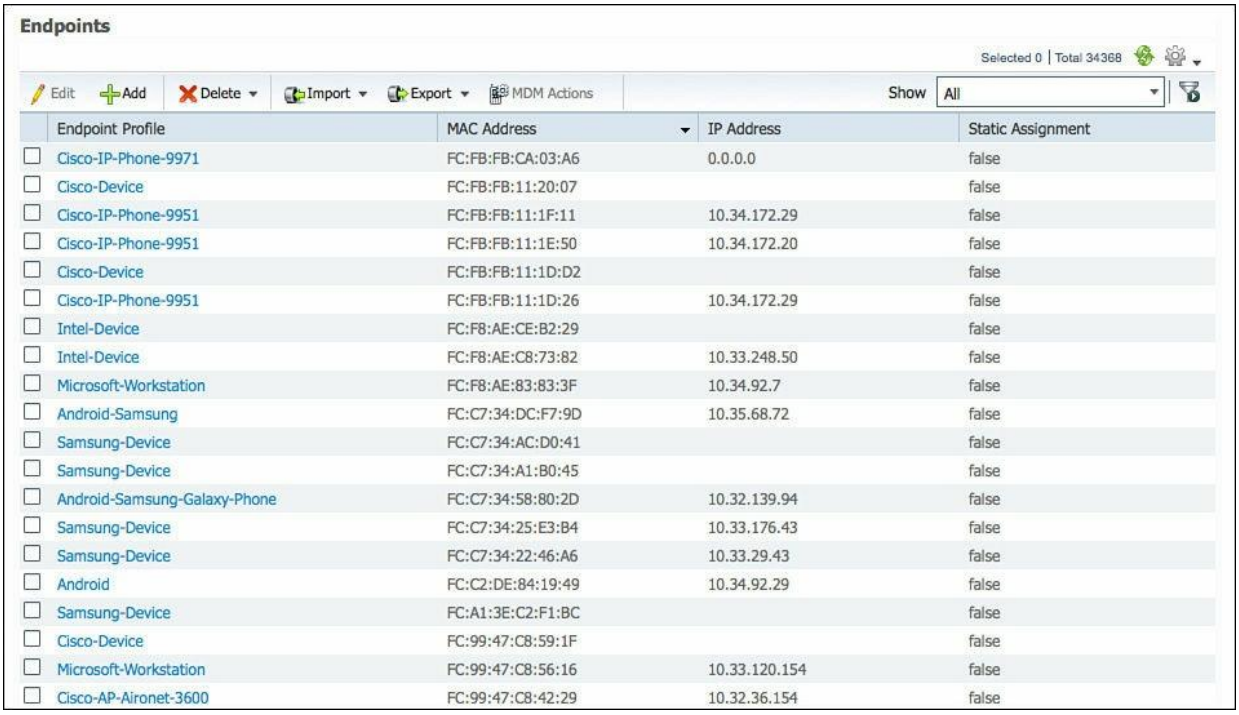


Figure 3-4 ISE internal endpoint database.

External Identity Stores

External identity stores are used to integrate your Cisco ISE deployment with your company's existing authentication infrastructure. By using the company's existing external authentication database, you can eliminate duplication of account creation and management. Because user and machine accounts are created, modified, or deleted in the external identity store, you do not have to worry about that process being duplicated for the Cisco ISE deployment. External identity stores allow for better scalability and management for ISE authentication and authorization processes.

The Cisco ISE deployment will connect to the external identity source to verify username and password credentials for authentication policies. Certificate information can also be retrieved from the external identity store. Cisco ISE uses authentication protocols to secure communications with the external identity stores. These authentication protocols are discussed more in depth in Chapter 4, "EAP over LAN (also Known as 802.1X)."



Examples of external identity stores include certificate authentication profiles, Active Directory, LDAP, RADIUS Token, and RSA SecureID. See Figure 3-5 for an example of external identity stores available for configuration in ISE.

https://t.me/learningnets



Figure 3-5 External identity stores.

Active Directory

One of the most popular external identity store supported with Cisco ISE is the integration with Microsoft Active Directory for user and machine authentication and authorization. Additional support includes authorization using AD security group membership and various other AD attributes.

Cisco ISE supports Active Directory natively by joining the Microsoft Active Directory domain. Support for multiple AD domains can be also supported by leveraging the use of two-way trust relationships between the domains.



Cisco ISE 1.2 supports AD integration as an external identity store using Microsoft Active Directory Servers 2003, 2008, 2008 R2, and 2012. Microsoft Active Directory Server 2000 is not support by Cisco ISE 1.2. See Figure 3-6 for an example of ISE Active Directory identity source configuration.

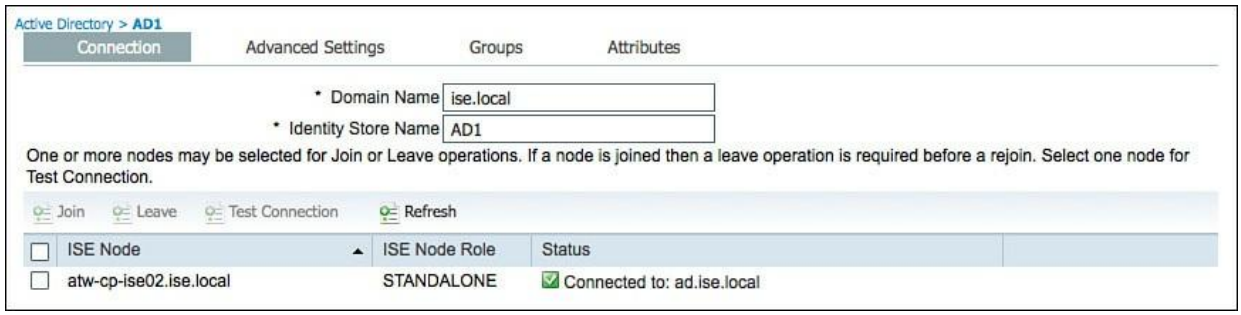


Figure 3-6 ISE Active Directory identity source configuration.

LDAP

The Lightweight Directory Access Protocol (LDAP) uses TCP/IP networking protocol as defined by RFC 2251. LDAP uses a client model to connect and retrieve information from X.500 directory service databases using TCP port 389. See Figure 3-7 for an example of LDAP identity source configuration.

https://t.me/learningnets

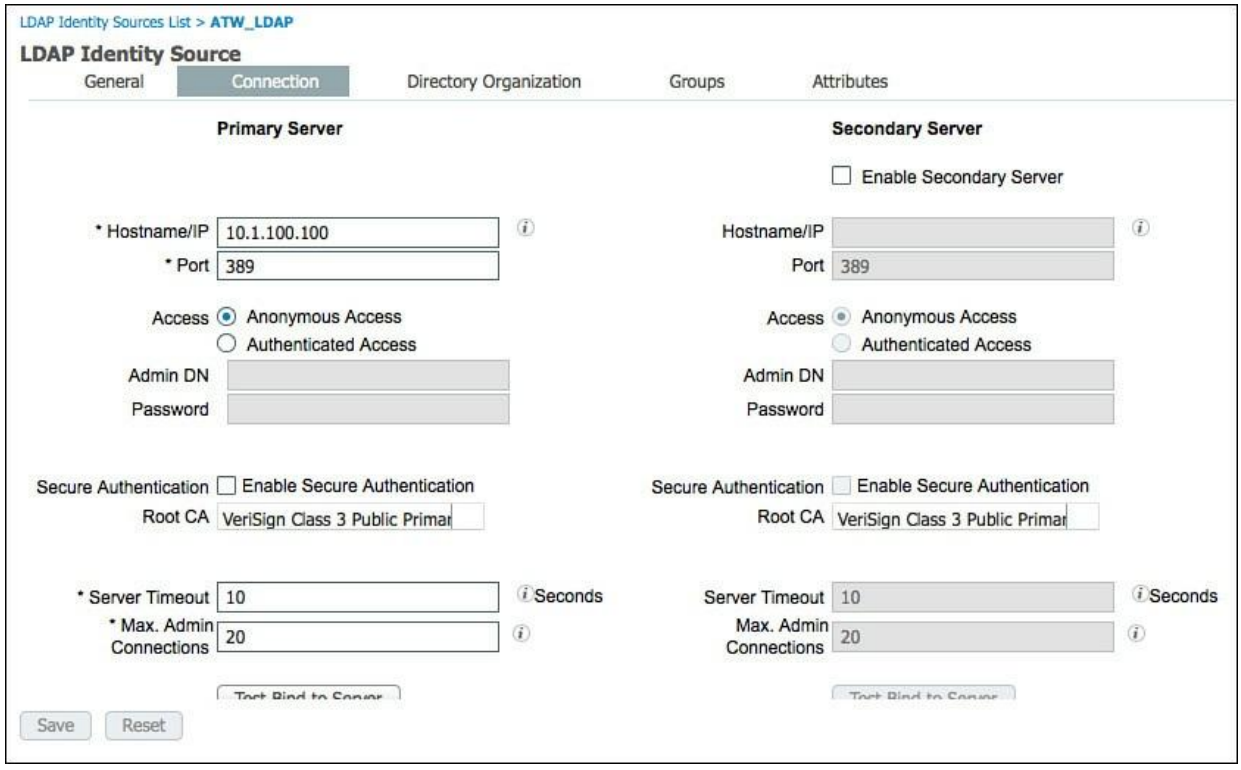


Figure 3-7 ISE LDAP identity source configuration.

Using LDAP as an external identity store can be supported for user authentication for Cisco ISE deployments. The username and passwords are sent in the clear by default when using LDAP. To prevent clear-text credentials from going across the wire, it is highly recommended to secure the LDAP communications using Secure Sockets Layer (SSL). LDAP over TLS/SSL uses TCP port 636.

LDAP also supports the retrieval of group membership information that can be used in the Cisco ISE authentication and authorization policies.

The LDAP external identity store supports configuring both a primary and secondary IP address for failover redundancy. You can also configure more than one Cisco ISE external identity store to connect to multiple LDAP X.500 database instances if needed.

The LDAP protocol is used in a variety of today's Enterprise Identity Management (IDM) solutions, including Active Directory, Sun Directory Server, Novell eDirectory (formerly Novell Network Directory Services [NDS]), Oracle Identity Manager, and IBM Tivoli Identity Manager (TIM) to name a few.

Two-Factor Authentication

In secure enterprise networks, typical user identities are based on a username and some type of static password that hopefully has a required security policy of changing the user's password every 30, 60, or even 90 days.

One of the main disadvantages of using usernames with static passwords is that the credentials could be packet captured and then later used for a replay attack to break into the company's networks.

When dealing with sensitive data access or remote access, security administrators might want the password used to access the corporate network or virtual private network (VPN) systems to include more security by requiring two forms of identity.

This process is known as two-factor authentication.

A typical example of using two-factor authentication system is when you go to your bank's ATM system to withdraw money from your checking or savings account. To make the withdrawal, you need two separate items.

The first item is the actual ATM debit/credit card issued by your bank. This is known as the something "you

https://t.me/learningnets

have."

The second item that is required is the actual personal identification number (PIN) that you selected when you set up your bank account. This is known as the something "you know." Your account PIN is typically a four-digit number to maintain your secure access to your account.

If someone were to have your PIN but not your actual debit/credit card, he would not be able to process an ATM withdrawal from your account.

So the two factors for authenticating the user identity that will come into play are the required something "you have" and the separate required something "you know."

One-Time Password Services

Another version of two-factor authentication is called one-time password (OTP) services. You are still required to "have something" and to "know something."

In this case, the "have something" comes in the form of a hard or soft token card.

Hard tokens typically come in various forms, including credit card, USB token key, or key FOB to name a few. Soft tokens, also known as virtual tokens, serve the same role as a hard token but run as software on your computer instead of you having to carry around a separate physical token card. See Figure 3-8 for an example of various hard tokens types.



Figure 3-8 Hard token types.

There are two main advantages of using a soft token over hard token cards. The first benefit is the extra cost that is associated with having to have an actual hard token card per each user. The second advantage to using a soft token is you do not have to worry about misplacing or losing the hard token card itself anymore.

OTP systems are based on algorithms that make use of pseudo randomness to generate unique security tokens. Two possible types of algorithms are usually used by OTP systems to generate these unique tokens. The first algorithm uses a time synchronization algorithm and the second method uses a mathematical algorithm.

The time synchronization method is used to guarantee the uniqueness of the one-time password. The typical hard token device (credit card, key FOB, or USB token) maintains an accurate clock that is in sync with the OTP authentication server. The time synchronization is critical to maintain the expected next unique token (password) between the end user's hard token and the OTP authentication system. Typically, there is a +1 to –1 ratio of acceptable passwords. This means that the authentication system will accept the last generated password, the current password, and the next password in sequence. This helps account for a minor drift in

https://t.me/learningnets

time synchronization between the token card and the token authentication server.

The mathematical algorithm uses the previous OTP to generate the next OTP or by using a challenge (also called a PIN) that must be entered to generate the next OTP. The second method is what you typically see when using a soft token (virtual token) system. The software is run, and then the user is prompted for her secret challenge phrase/PIN. Upon successful input of the secure PIN, the soft token software program generates a unique OTP for the user to enter as her password with her network user ID.

The OTP that is generated is valid for one user login session only. By requiring users to use OTPs for logging in to either enterprise secure systems or remote VPN networks, you are enforcing a security policy that requires a unique password for each unique user login. The primary advantage of using OTP systems is that they are not subject to the replay attacks we discussed in the previous section for typical network username/password login systems.

You must also still have the "know something," which in this case is the PIN associated with the hard token card or soft token software. The key difference when using an OTP system is that your PIN is actually used as input to the token software to generate an unique OTP that will then be entered as your password with your username to access your company's secure networks or remote VPN networks.

RSA SecurID and SafeWord are just two examples of OTP systems that can be implemented with Cisco ISE.

Smart Cards

Another type of identity store supported by Cisco ISE is with the use of smart cards. Smart cards use embedded integrated circuits that then can be used as forms of identification. That identification can be used as a method authentication for secure network access control systems.

Common access card (CAC) is another term used to refer to smart cards that allow users to use their identification badges to authenticate. The end user inserts his CAC into a card reader and then enters his PIN. Again, you have to "have something" and "know something" for the authentication to be valid.

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

Today's smart cards are used in variety of formats including bank credit cards, drivers' licenses, and hospital patient identity cards to name a few examples.

Certificate Authorities

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

To help understand the topic of certificates, it is often easiest to relate it to something most of us are familiar with, which is your driver's license. When you take your driver's exam and pass the test, you are provided your driver's license documentation. This document 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 will use various methods to sign that license to indicate it came from a trusted source and is valid and trusted. One example that is typically used is a hologram embedded on the driver's license, which makes it difficult to duplicate or forge.

An X.509 certificate is very much like a driver's license. It is used to identify the user, device, application, or whatever that entity might be. The certificate authority (CA) vouches for the identity contained in the X.509 certificate by "signing" the X.509 information.

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 driver is asked to provide his driver's license, which the police officer will then use to perform the following tasks:

https://t.me/learningnets

1. Validate that the driver's license appears to be a real document and not a forgery.
2. Ensure the picture does in fact look like the person in the driver's seat.
3. Ensure that the driver's license is still valid and not expired.
4. The police officer will then verify via computer query that the driver's license has not been suspended or revoked based on past offenses.



X.509 certificates work in a 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):

1. Has the digital certificate been issued and properly signed by a known and trusted certificate authority?
2. Is the certificate expired? (Must check both the issue and expiration dates!)
3. Has the certificate been revoked by the certificate authority?
4. Has the client provided proof of possession?

Has the Digital Certificate Been Signed by a Trusted CA?

As mentioned previously, a CA vouches for the identity that's contained within the certificate. The owner of the certificate (the client) and the access device (the server) must both mutually trust the CA. When an identity certificate is presented, you must be able to trust the signing authority. Without the established trust first, the authentication server has no way to validate the identity certificate being presented.

The signing of a certificate really has two requirements. The first requirement is that the certificate must be presented in the proper format as based on the X.509 standards. If the certificate is not formatted properly, the authentication server will discard the identity certificate being presented and the authentication process will then be immediately terminated.

The second requirement is that the trusted signing certificate authority's public key must be stored in the Trusted Certificates database of the authentication server and that certificate must be trusted for purposes of authentication processes.

Using Cisco ISE as an example, the trusted certificate will need to have the Trust for Client Authentication use case selected, as shown in Figure 3-9.

https://t.me/learningnets

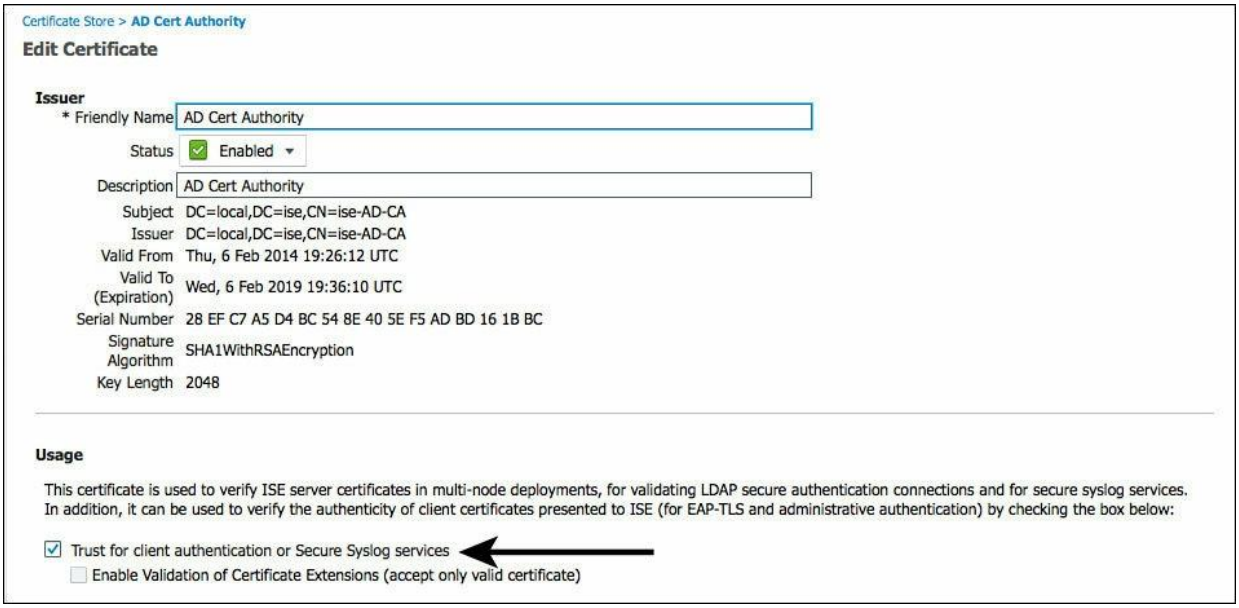


Figure 3-9 Trust for Client Authentication.

So not only does Cisco ISE "trust" the certificates that have been signed by the certificate authority, but it also must trust those certificates for the specific required use case, such as client authentication in this example. If a client presents a certificate and the chosen CA did not authorize the certificate for client authentication, the authentication immediately fails. This authentication process is comparable to someone entering the wrong password for a user account—or, better yet, a user attempting to use a movie ticket for the wrong movie.

Has the Certificate Expired?

Just like a driver's license or a passport that has an issued date and an expiration date, a certificate has validity dates. The certificate list the Issued On date, which is the starting date that the certificate is first valid. It also has an Expires On date, which is the date on which the certificate will be no longer valid. Figure 3-10 illustrates the valid-from and valid-to attributes within a certificate.

https://t.me/learningnets

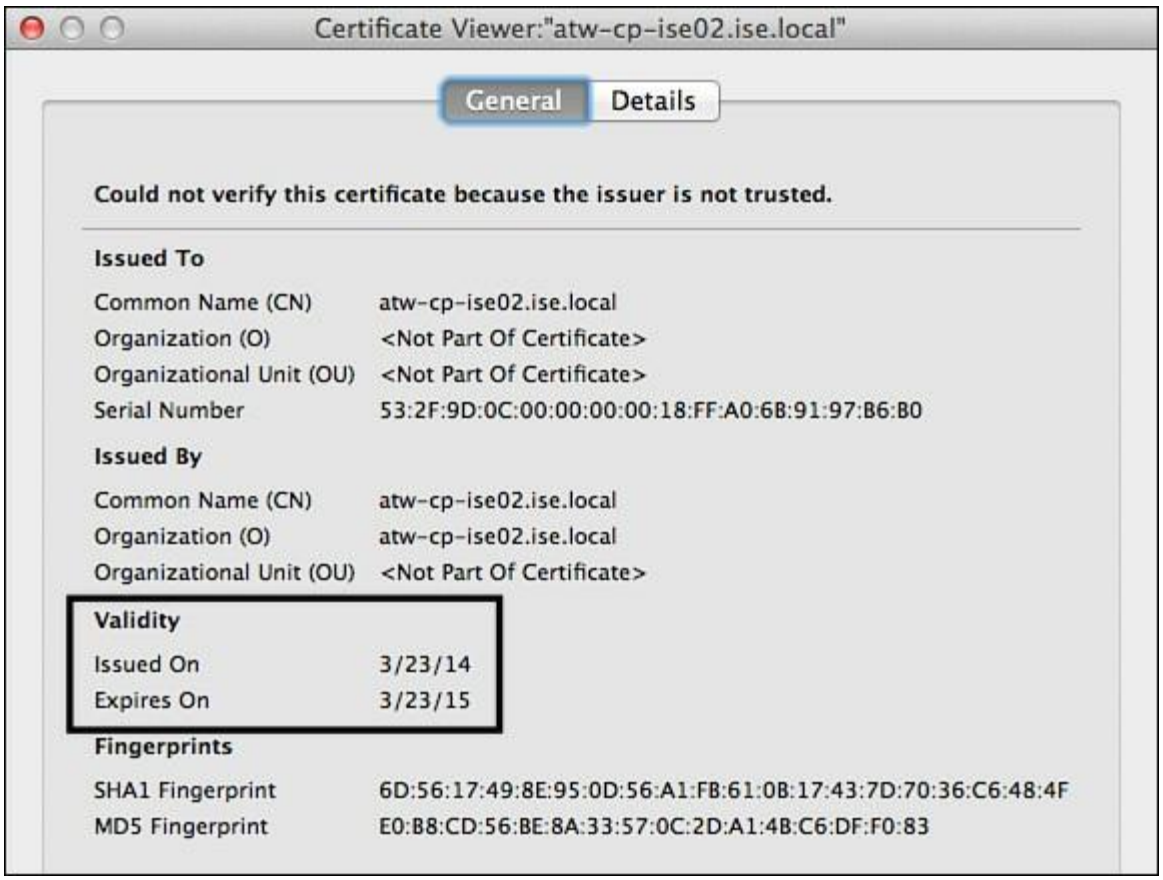


Figure 3-10 Certificate validity.

After the certificate has expired, it will have to be renewed with the trusted certificate authority. An authentication server will examine these dates to ensure 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 your CA, when creating new certificates along with your authentication servers when processing certificates used for authentication. The best practice is to use the Network Time Protocol (NTP) (UDP port 123) to a trusted and authenticated network time server. This NTP server acts as a centralized timekeeper for the network.

Many of the issues involved with creating, implementing, and processing certificates concern time being out of sync. A typical example is when the certificate Issued On date and time stamp is not valid yet. The certificate might be presented for authentication on September 9, 2014, at 10:13 a.m., but its Issued On value in the certificate itself is listed as starting on September 10, 2014, at 5:13 p.m.

Characteristically this is because the certificate authority was NOT in sync when the certificate was originally generated.

Has the Certificate Been Revoked?

Now if we go back to the police officer analogy that was described earlier, remember that the officer can look at your driver's license and verify whether the expiration date is still valid. The question still remains: How does the police officer know if your driver's license has been revoked?

To verify that your driver's license is not revoked, the police officer will have to perform additional checks of your driver's license number against the DMV's database. If the police officer runs your driver's license number and finds out that you have many speeding tickets and in fact your driver's license has been revoked, you would be immediately arrested for driving on a revoked license.

https://t.me/learningnets



When presented with a certificate for authentication, the authentication server has the same capability to verify that the trusted certificate authority has not revoked the identity certificate. Every CA should also have a service to publish a list of certificates that have been revoked. There are two primary ways to check for certificates that have been revoked: a certificate revocation list (CRL) and the Online Certificate Status Protocol (OCSP). These are discussed next.

Certificate Revocation List

This is basically a signed list of the revoked certificate serial numbers that the CA publishes on a website that can be read by any device or application needing to confirm a certificate's validity. This file is periodically downloaded and stored locally on the authentication server. When a certificate is then presented for authentication, the authentication server examines the CRL to see whether the client's certificate serial number is on the list indicating that the certificate has been revoked.

The use of CRL option does lead to some security concerns that you need to be aware of.

How often does the CA update the CRL? If the CRL is updated only once a week, then there is still the possibility that the CRL will be outdated if presented with a recently revoked certificate. For instance, if the CRL was just cached yesterday at noon on the authentication server with a caching time of one week, then the certificate that is revoked as of yesterday at 12:01 p.m. can still be used until next week when the CRL would be scheduled for the next update.

How often does the authentication server download the CRL from the CA? 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.

To deal with these security concerns, let's take a look at how OCSP works as compared to the CRL method.

Online Certificate Status Protocol

This 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 another device and checking the status of the certificate right then and there. OCSP could be compared to the police officer using the computer in her squad car and doing a live lookup on the DMV's database.

If the certificate has been revoked, then access is denied.

Figure 3-11 is an example of the configuration screen for a CA in ISE. The administrator has the ability to configure a location to check for OCSP and/or CRL when this particular CA signs a certificate or its subordinates.

https://t.me/learningnets

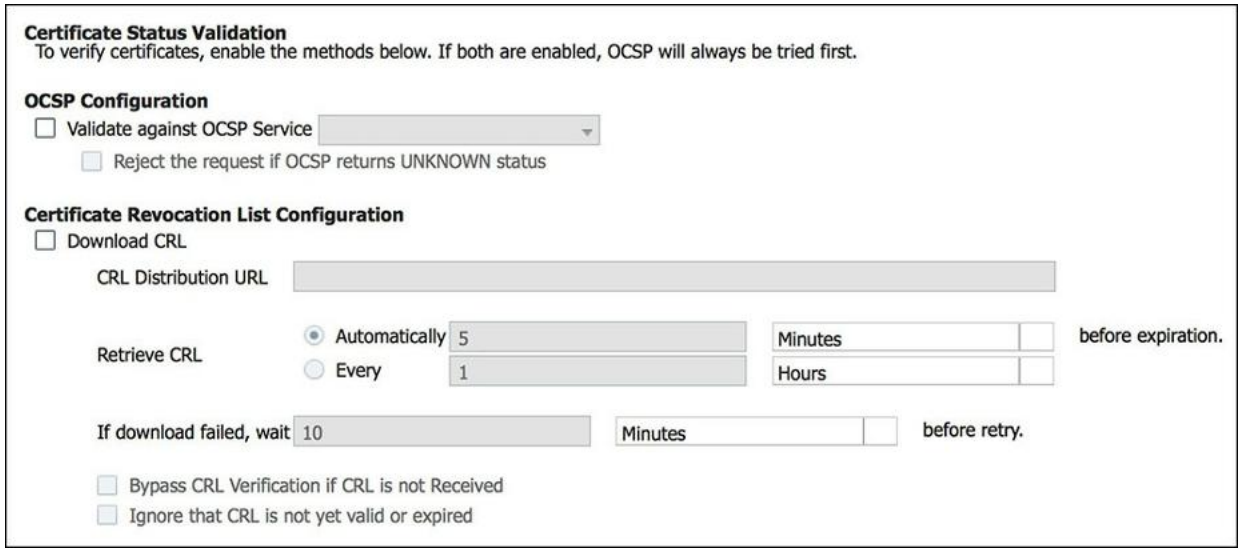


Figure 3-11 Certificate CRL and OSCP configuration.

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 23, "Final Preparation"; and the exam simulation questions on the CD-ROM.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. Table 3-2 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	Identity stores, identity sources, and identity source sequences	38
Paragraph	External identity sources	41
Paragraph	Supported Active Directory versions	42
Step List	X.509 certificates	46
Paragraph	Certificate revocation	49

Table 3-2 Key Topics for Chapter 3

Define Key Terms

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

identity store

identity source

identity source sequence

internal user database

https://t.me/learningnets

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)

two-factor authentication

one-time password (OTP)

common access card (CAC)

public key infrastructure (PKI)

certificate authority (CA)

registration authority (RA)

Network Time Protocol (NTP)

certificate revocation list (CRL)

Online Certificate Status Protocol (OCSP)

https://t.me/learningnets

Chapter 4. EAP Over LAN (Also Known As 802.1X)

This chapter covers the following exam topics:

- **EAP Types**
- **Describe Supplicant, Authenticator, Authentication Server**
- **Supplicant Options**
- **Network Access Devices**

Back in the early 2000s the IEEE standardized a solution for port-based network access control, known as 802.1X. It was predicted to revolutionize networking as we knew it, and no device would be able to plug in and communicate on a network without the user identifying himself and being authorized again. Well, here we are a decade later and 802.1X is really starting to catch on.

The three fundamental components of 802.1X are the supplicant, the authenticator, and the authentication server. This chapter explains those components, along with critical elements to an 802.1X solution, such as the various EAP types that can be used. A perfect companion topic for EAP types is the exploration of supplicants, which wraps up our chapter.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Network Access Devices	4
EAP over LAN (802.1X)	1-3, 5-6, 10
Describe Supplicant, Authenticator, Authentication Server	4, 9
EAP Chaining	2
Supplicant Options	2, 5-6, 8
Compare and Select AAA Options	7

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

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 the 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 whether the supplicant is allowed on the network.

   b. The EAP communication occurs between the supplicant and the authentication server.

https://t.me/learningnets

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(s) 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 an outer identity?

a. The outer identity is used for dual-factor authentications such as a username/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, whereas the inner identity represents the user during EAP chaining.

4. True or False? IEEE 802.1X may use TACACS+ to communicate the EAP identity to the authentication server.

a. True

b. False

5. True or False? The supplicant is required to trust the certificate of the authentication server before it will form the TLS tunnel within which the EAP transaction will occur.

a. True

b. False

6. What is the name of the "secure cookie" used with EAP-FAST that can be used in lieu of a certificate, or even 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. True or False? MSCHAPv2 may be used to perform machine authentication with an LDAP connection to Active Directory.

a. True

b. False

8. True or False? A machine authentication may use EAP-FAST.

a. True

b. False

9. What are the three main components of IEEE 802.1X?

a. Agent, broker, authentication server

b. Supplicant, authorizer, authorization server

c. Authentication server, supplicant, authenticator

d. EAP, RADIUS, TLS

https://t.me/learningnets

10. True or False? A tunneled EAP type is able to use native EAP types as its inner method.

   a. True

   b. False

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 password (OTPs), 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 authentication methods including VPN, 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. The three main components to 802.1X are the supplicant, the authenticator, and the authentication server; these are illustrated in Figure 4-1 and described in Table 4-2.

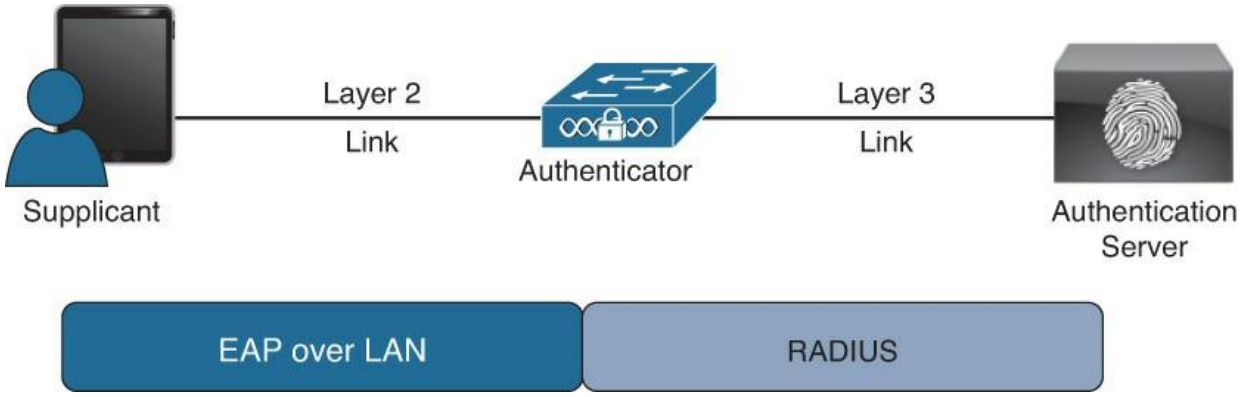


Figure 4-1 Components of 802.1X.

Key Topic

https://t.me/learningnets

Component Name	Description
Supplicant	Software on the endpoint (also called a *peer* in IETF) 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 taking Layer-2 EAP communication from the supplicant and encapsulating it 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 device (NADs).
Authentication Server	The server that is performing the actual 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 Identity Services Engine is an authentication server.

Table 4-2 802.1X Components

While reviewing Figure 4-1 and Table 4-2, keep in mind that the authenticator (switch or WLC) acts as the middleman or proxy. The actual EAP identity exchange and authentication is occurring between the supplicant and the authentication server. The switch or WLC has no idea of which EAP type is in use or whether the user's credentials are valid. It simply takes the unmodified EAP frame and encapsulates it within the RADIUS packet sent to the authentication server and authorizes the port if the authentication server tells it to. Therefore, the EAP authentication itself is completely transparent to the authenticator. Figure 4-2 displays the communication through a successful authentication.

https://t.me/learningnets

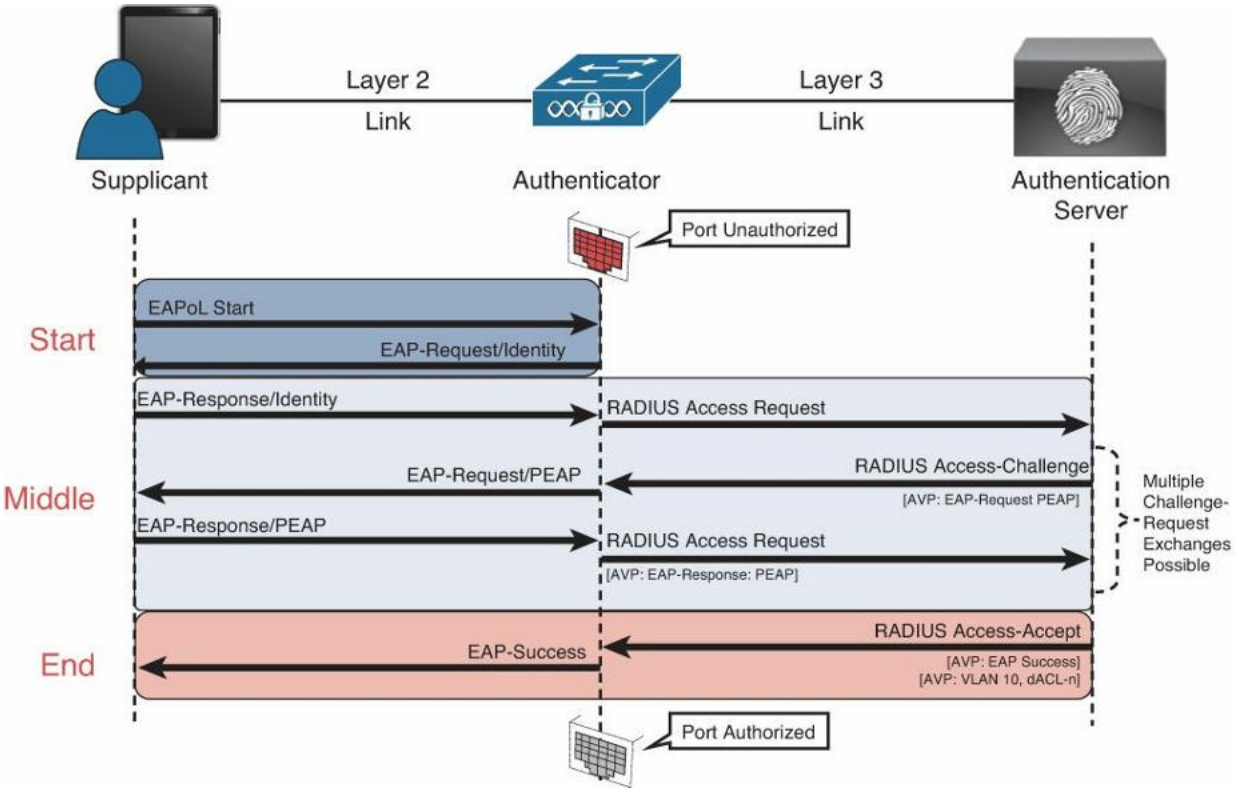


Figure 4-2 Authentication communication.

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 EAP types, and each one has its own benefit and downside. The EAP type defines the authentication mechanism to be used with EAP, which is usually self-evident in its name. Most of EAP types are not discussed in this book, due to lack of adoption or lack of inclusion in the exam blueprint, such as 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 nontunneled EAP inside a Transport Layer Security (TLS) tunnel between the supplicant and the authenticator. See Figures 4-3 and 4-4 for a graphical representation of native and tunneled EAPs.

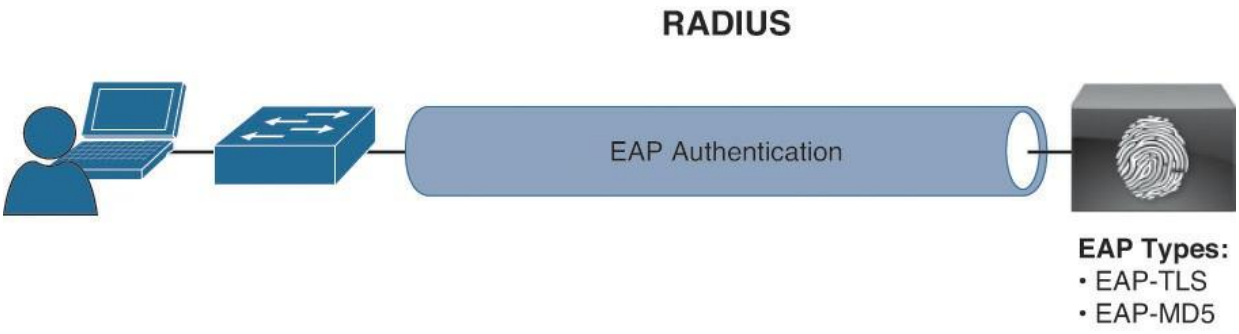


Figure 4-3 Native EAP methods.

https://t.me/learningnets

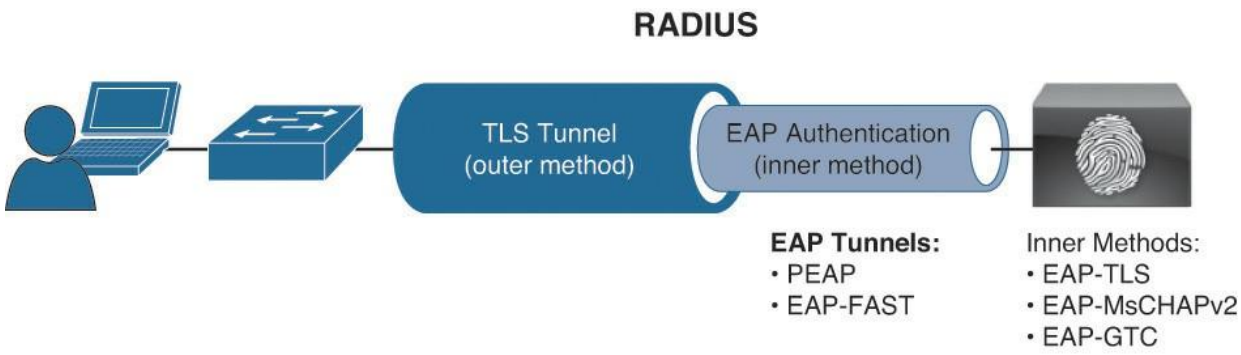


Figure 4-4 Tunneled EAP methods.

Native EAP Types (Nontunneled EAP)

Native EAP types include the following:

- **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 whether the credentials were accurate. However, EAP-MD5 does not have a mechanism for mutual authentication. That means the server is validating the client, but the client does not authenticate the server (that is, it does not check to see whether it should trust the server). EAP-MD5 is common on IP phones, and it is also possible that some switches will send MAC authentication bypass (MAB) requests using EAP-MD5.
- **EAP-TLS**—An EAP type that uses TLS to provide the secure identity transaction. This is similar to SSL and the way encryption is formed between your web browser and a secure website. EAP-TLS has the benefit of being an open IETF standard and is considered to be 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 because 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-MSCHAPv2**—An EAP type in which the client's credentials are sent to the server encrypted within an MSCHAPv2 session. This allows for simple transmission of username and password, or even computer name and computer passwords, to the RADIUS server, which in turn authenticates them to Active Directory (AD).

> **Note**
>
> Cisco ISE does not currently support EAP-MSCHAPv2 as a native EAP, only within a tunneled EAP.

- **EAP-GTC**—EAP-Generic Token Card (GTC) was created by Cisco as an alternative to MSCHAPv2 that allows generic authentications to virtually any identity store, including OTP token servers, LDAP, Novell E-Directory, and more.

> **Note**
>
> Cisco ISE does not currently support EAP-GTC as a native EAP, only when it is inside a tunneled EAP.

https://t.me/learningnets

Tunneled EAP Types

The previously mentioned native EAP types send their credentials immediately, while tunneled EAP types form an encrypted tunnel first and then transmit the credentials within that tunnel.



> **Note**
>
> It is important to understand that tunneled EAP methods work 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); then the user enters her banking login credentials within that secure tunnel.

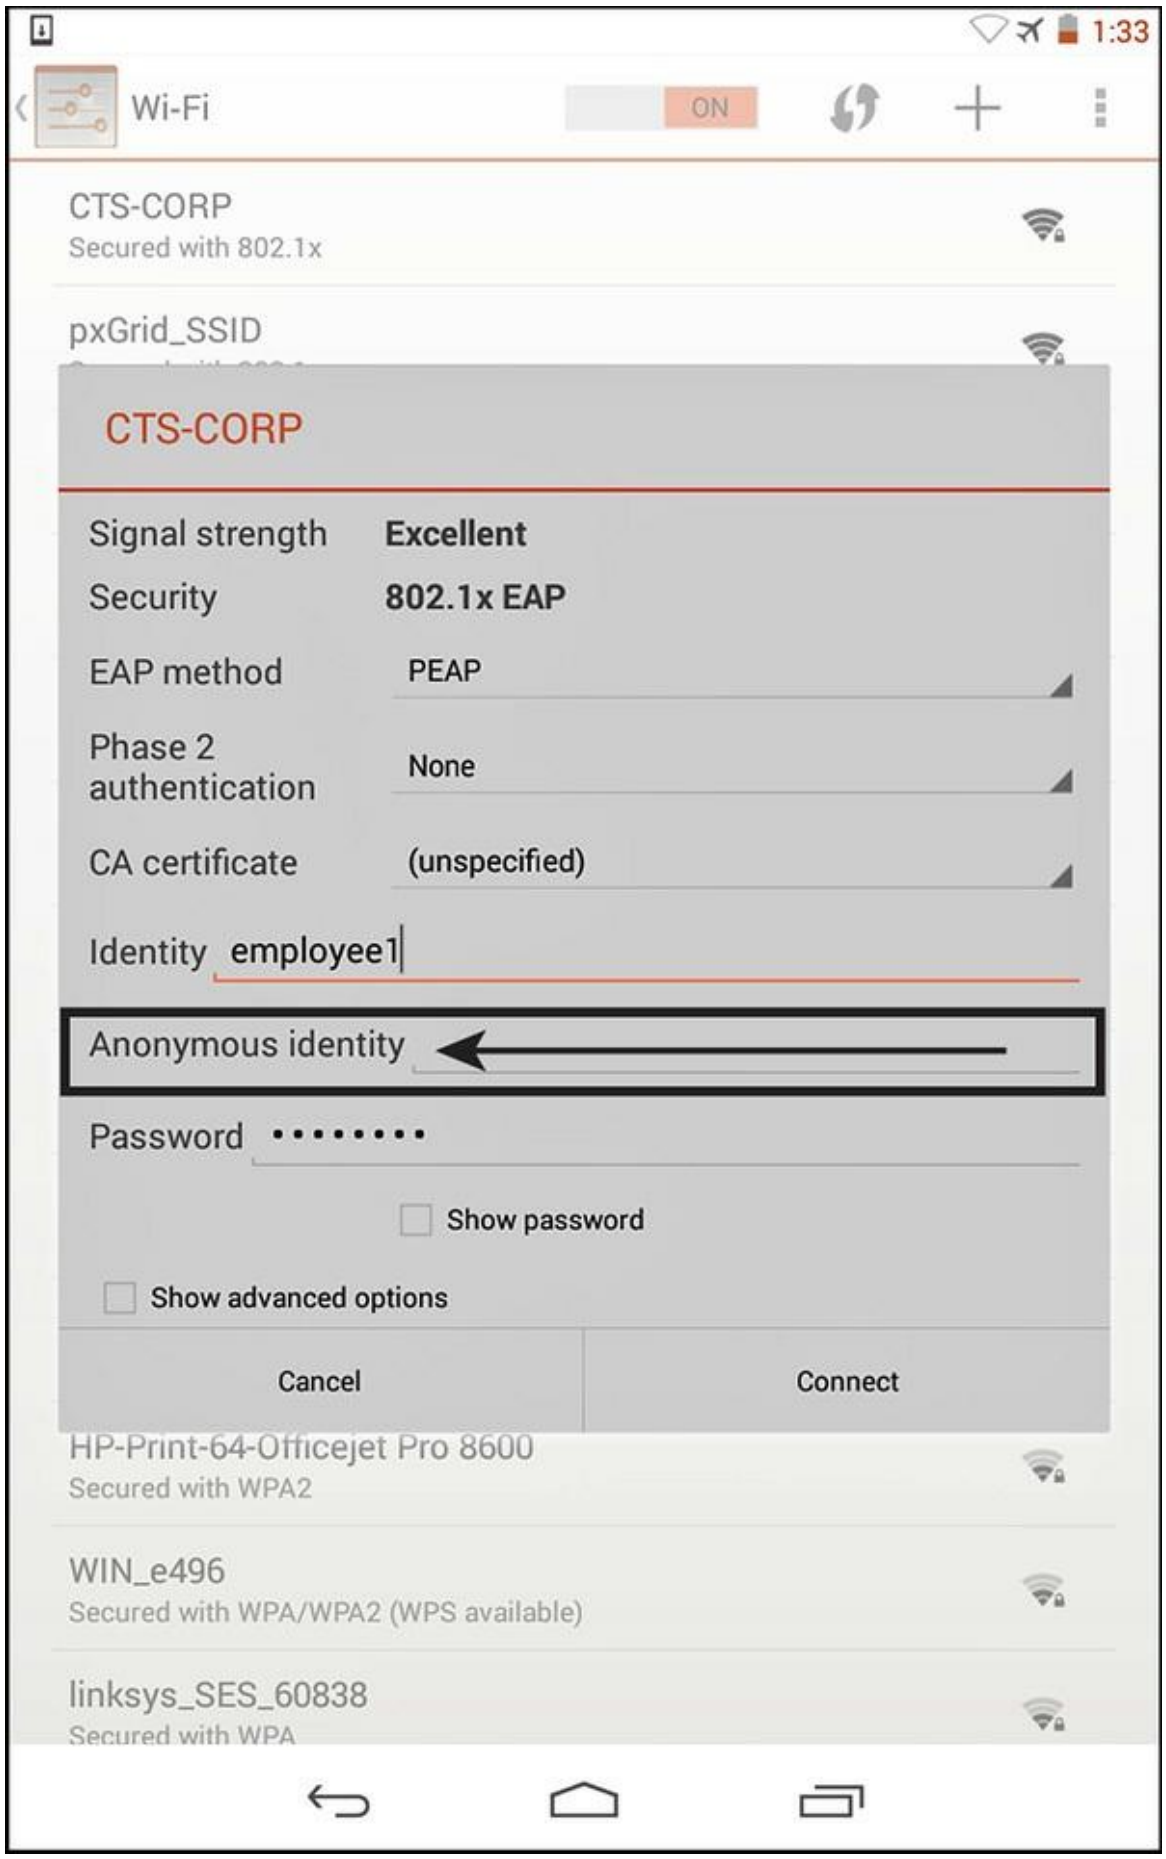
- **PEAP**—Protected EAP (PEAP) was originally proposed by Microsoft and is an EAP tunnel type that has quickly become the most popular and widely deployed EAP method in the world. PEAP forms a potentially encrypted TLS tunnel between the client and server using the x.509 certificate on the server in much the same way 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.
  - **EAP-MSCHAPv2**—Using this inner method, the client's credentials are sent to the server encrypted within an MSCHAPv2 session. This is the most common inner method because 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**—This inner method was created by Cisco as an alternative to MSCHAPv2 and enables generic authentications to virtually any identity store, including OTP token servers, LDAP, Novell E-Directory, and more.
  - **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 Tunnel (FAST) is similar to PEAP. FAST was created by Cisco Systems as an alternative to PEAP that enables faster reauthentications 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 the ability to use protected access credentials (PACs). A PAC can be thought of like a secure "cookie," stored locally on the host as proof of a successful authentication.
  - **EAP-MSCHAPv2**—Using this inner method, the client's credentials are sent to the server encrypted within an MSCHAPv2 session. This is the most common inner method because 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**—This inner method was created by Cisco as an alternative to MSCHAPv2 that enables generic authentications to virtually any identity store, including OTP token servers, LDAP, Novell E-Directory, and more.
  - **EAP-TLS**—EAP-FAST is capable of using EAP-TLS as an inner method. This has become quite popular with EAP chaining.

With tunneled EAPs, there is a concept of inner and outer identities. The inner identity is easiest to explain. It is the user's or device's actual identity credentials sent with the native EAP protocol. The outer identity is typically set to "anonymous." It's the identity that is used between the supplicant and the authentication server for the initial TLS tunnel setup.

https://t.me/learningnets

Cisco ISE is able to read that outer identity and use it to help make identity store selection decisions. Put simply, that outer identity can contain information (such as domain name) that tells Cisco ISE to submit the credentials to Active Directory, 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, as shown in Figure 4-5. One point of humor to this author is that the Android supplicant refers to the outer identity as the "Anonymous Identity," which is an oxymoron.

https://t.me/learningnets



https://t.me/learningnets

Figure 4-5 Android supplicant exposes the outer identity.

Summary of EAP Authentication Types

Table 4-3 shows a comparison of common EAP authentication types.



	EAP-PEAP	EAP-TLS	EAP-FAST
Client-side certs required	No	Yes.	No.
Server-side certs required	Yes	Yes.	No.
Security	High ■ Protects EAP methods: ■ TLS tunnel creation for EAP packets ■ Message authentication ■ Message encryption ■ Authentication of server to client	High ■ Authentication occurs automatically, no intervention by user. ■ No dependency on user password. ■ Digital certs provide strong authentication. ■ Message exchange is protected with public key encryption.	High ■ Crypto-binding between inner and outer methods prevents MITM attacks. ■ Immunity to passive dictionary attacks. ■ Use of preshared keys (PAC files) to construct tunnel through use of symmetric cryptography (optional).

Table 4-3 EAP Comparison

EAP Authentication Type Identity Store Comparison Chart

Selecting the appropriate EAP type is dependent on the operating system, 802.1X supplicant, and supported back-end credential database or identity store. The comparison is displayed in Table 4-4.

Identity Store	EAP/ TLS	PEAP/ TLS	PEAP/ MSCHAPv2	PEAP/ GTC	EAP-FAST/ MSCHAPv2	EAP-FAST/ GTC	EAP-FAST/ TLS
MS-AD	Yes*	Yes*	Yes	Yes	Yes	Yes	Yes*
LDAP	Yes*	Yes*	No	Yes	No	Yes	Yes*
Token Server	No	No	No	Yes	No	No	No
RSA	No	No	No	Yes	No	Yes	No

*TLS authentication is validating only the certificate; however, the authorization function may use these identity stores.

Table 4-4 EAP Authentication Type Identity Store Comparison

https://t.me/learningnets

Network Access Devices

Cisco ISE refers to the authenticator role as a network access device (NAD). The NAD serves multiple roles. Yes, it is an authenticator for 802.1X and will proxy EAP communications from a supplicant to the authentication server. The NAD is also what is commonly referred to as a policy enforcement point (PEP). The NAD is responsible for enforcing whatever authorization result it receives from the authentication server (Cisco ISE in this case).

In simple terms, a NAD is the Access Layer device but can be any device that is going to send RADIUS authentication requests to Cisco ISE. Common NAD types include

- Wired Ethernet switches
- Wireless LAN Controllers
- Cisco adaptive security appliances (ASAa)

Enforcement types are covered in more detail, but to provide a little color to this subject, common enforcement types include the following:

- Dynamic VLAN (dVLAN) assignment
- Downloadable access control lists (dACLs)
- Security group tags (SGTs)
- Airespace ACL names (for use with Wireless LAN Controllers)
- URL redirections

Network access devices are an incredibly important part within the authentication system known as secure access. It is doing so much more than simply passing frames at Layer-2. These devices are the security enforcement devices; they apply the policy to the end user. These devices are doing functionality that used to be available only in overlay appliances, such as URL redirection and enabling agent discovery of the active posture server. These topics are covered in much more depth; however, it is important to stress just how key these devices actually are 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 itself. That would never be complete without discussing the client side of the authentication flow, the supplicant. As described earlier in the 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 with the user's username and password.

Windows Native Supplicant

The single 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. Besides that, 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 the supplicant appear very attractive to desktop managers of corporations all over the world.

> **Note**
>
> In the authors' experience, companies tend to start off their 802.1X projects utilizing this supplicant. However, the first time any major troubleshooting is required, they quickly start looking at more premium supplicants such as Cisco's AnyConnect NAM or the Odyssey client.

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. This will need to be changed to Automatic for the

https://t.me/learningnets

supplicant to be enabled at each boot. There is also a WLAN AutoConfig service for the wireless supplicant, which is set to automatic by default.

At this point in the chapter, we will take a guided tour of the Windows native supplicant. This exercise will hopefully drive home some of the points made previously, surrounding EAP types and inner versus outer methods. If you have a Windows 7 or Windows 8 device at your disposal, feel free to follow along.

From the Windows services control applet, do the following:

Step 1. Locate **Wired AutoConfig** from the list of services as shown in Figure 4-6, and double-click it to enter the service properties.

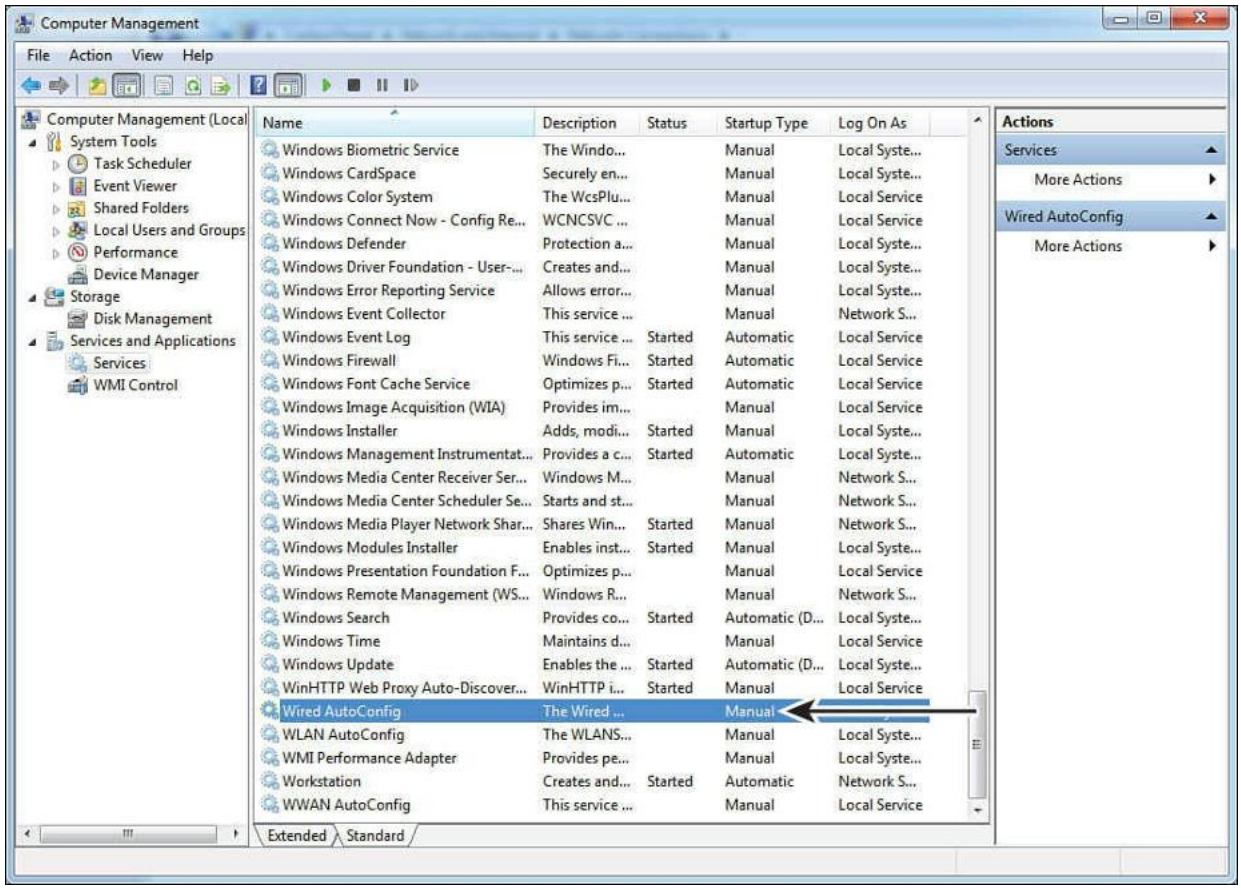


Figure 4-6 Windows service control applet.

Step 2. Change the startup type value to **Automatic**, as shown in Figure 4-7.

https://t.me/learningnets

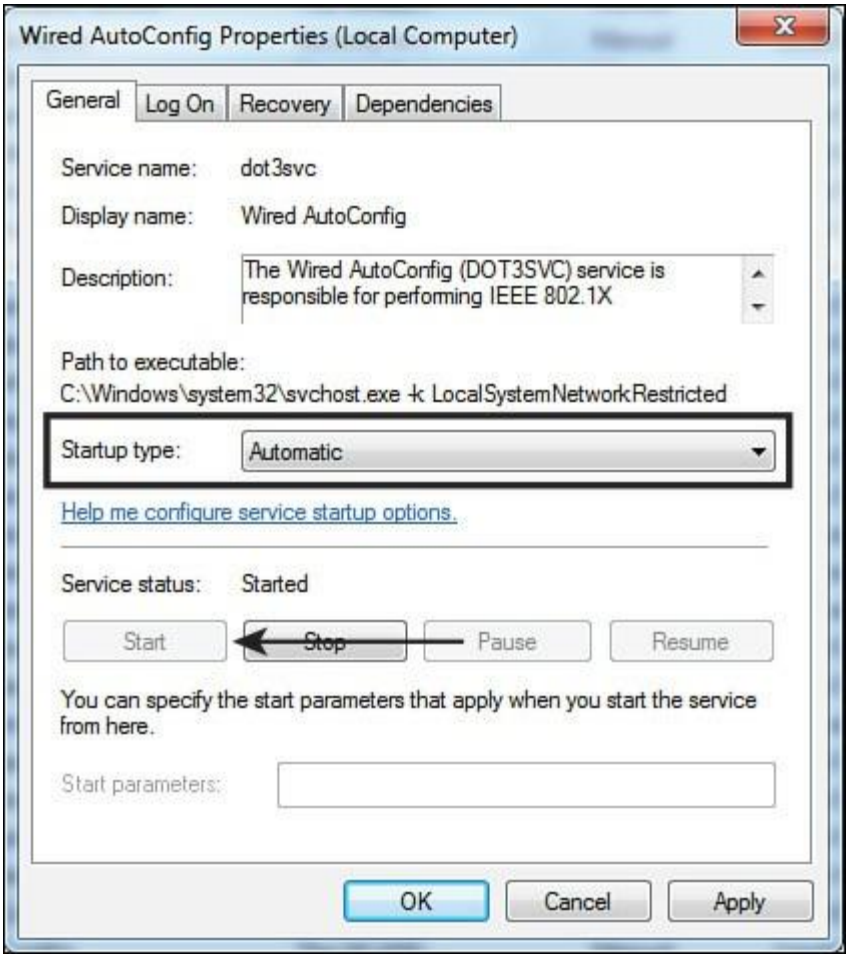


Figure 4-7 Wired AutoConfig service properties.

Step 3. Click **Start** to enable the supplicant. This will also populate the wired NIC properties page with an authentication tab, as shown in Figure 4-8.

https://t.me/learningnets

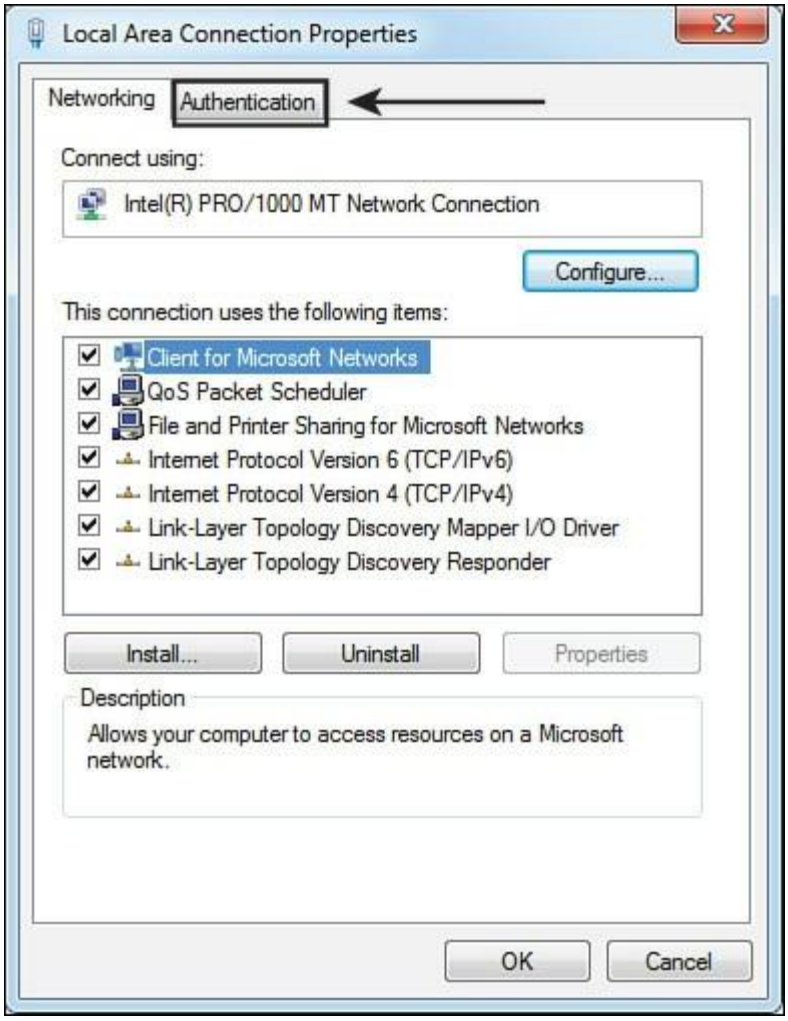


Figure 4-8 Wired AutoConfig service properties.

Navigate to **Control Panel > Network and Internet > Network Connections**.

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

https://t.me/learningnets

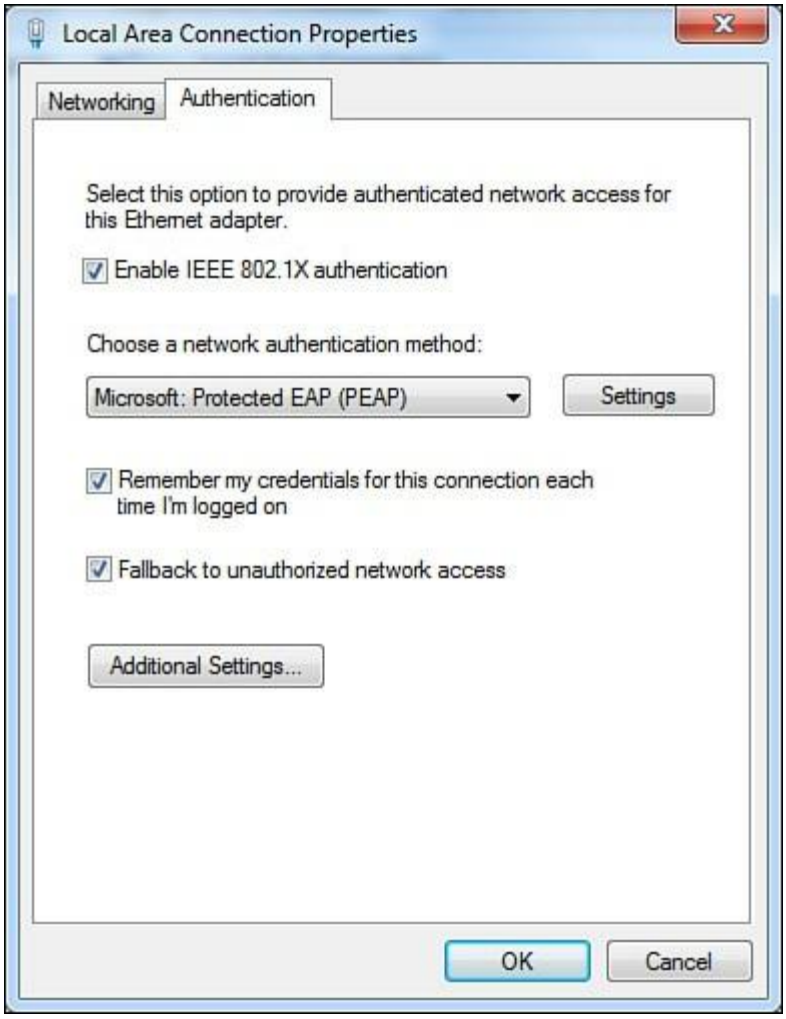


Figure 4-9 Authentication tab.

As shown in Figure 4-9, the Authentication tab has a check box to enable or disable IEEE 802.1X authentication. It also has a check box to allow the supplicant to store (remember) the credentials, thus not prompting the end user for his credentials at each network login.

The last check box is to Fallback to Unauthorized Network Access. This setting basically states that if the network device is not an authenticator (that is, it does not send EAP identity requests), the network connection should still work. If this setting was unchecked and the user plugs into a non–802.1X-enabled switch, the connection would be treated as if the authentication failed and the user would not have network access. It's a good idea to leave this setting enabled to ensure positive user experience.

There are two more areas for additional settings. The first is the **Settings** button for the selected network authentication method. As shown in Figure 4-10, the default network authentication method is to use PEAP. Do the following:

https://t.me/learningnets

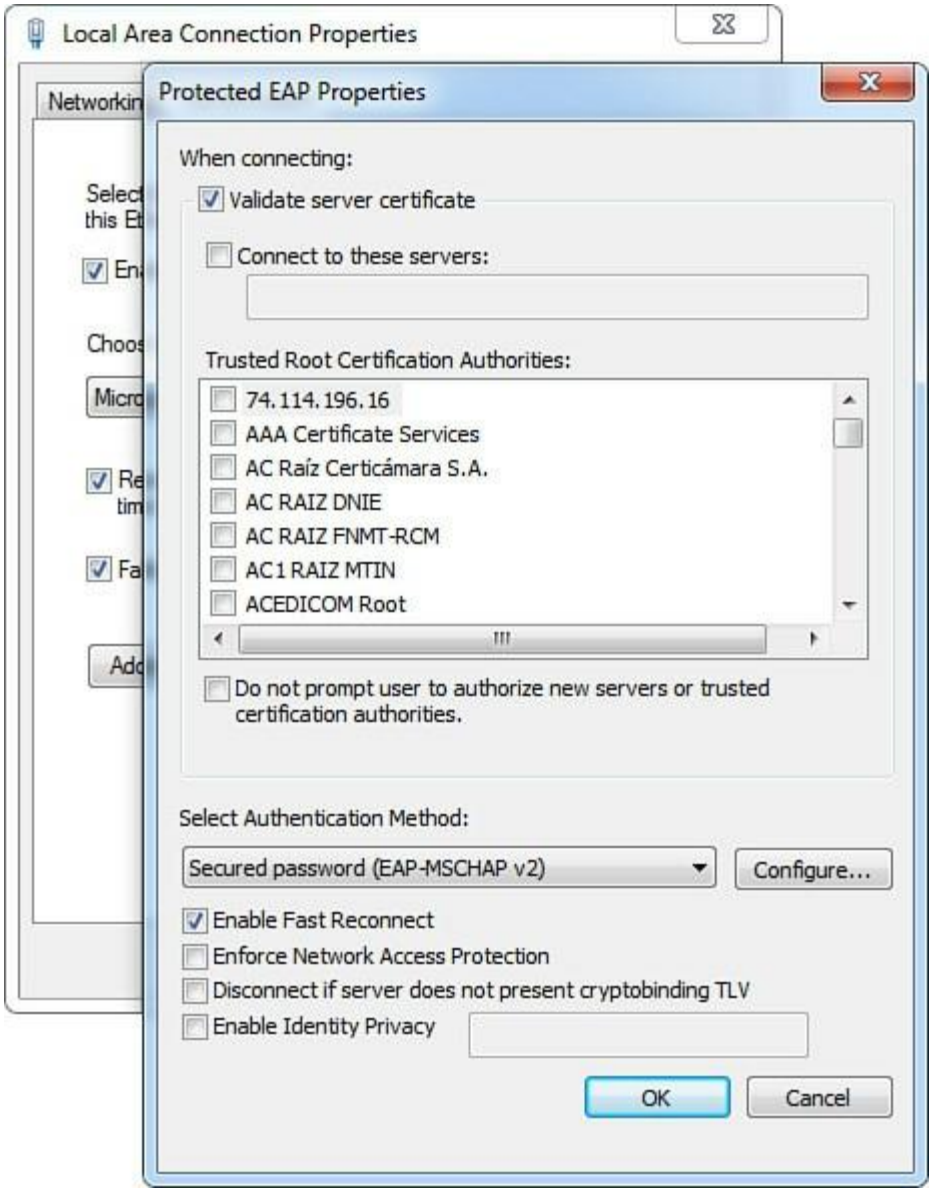


Figure 4-10 Protected EAP properties.

Step 1. Click the **Settings** button. The Protected EAP Properties page opens, as shown in Figure 4-10.

The first thing in the Protected EAP Properties page that should stand out is the check box for Validate Server Certificate, as shown in Figure 4-10. With this setting enabled, the supplicant is required to trust the certificate from the RADIUS server (802.1X authentication server) before it will form the secure TLS tunnel. Disabling this setting turns off that level of authentication and the supplicant would form a tunnel and send EAP credentials to any RADIUS server.

An option to specify the servers that are allowed is available, as well as which certificates to trust. This setting provides 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 customers.

In the Trusted Root Certificate Authorities area, the administrator is able to select which specific root CAs are trusted for authentication. The next check box is to not prompt a user to authorize new servers or trusted certificate authorities. By leaving this check box disabled, a user will not be prompted to trust a certificate that is not explicitly trusted in the list above the check box.

In the Select Authentication Method area of the properties page, the administrator is provided

https://t.me/learningnets

with an option to select the inner method for PEAP, which can be either Secured Password (EAP-MSCHAPv2) or Smart Card or Other Certificate (which will use EAP-TLS as the inner method for certificate authentication).

Step 2. Click **Configure** for Secured Password (EAP-MSCHAPv2) to bring up the properties page, as shown in Figure 4-11.



Figure 4-11 EAP MSCHAPv2 properties.

The only configurable option for this inner method is to automatically use the Windows logon credentials, also called single sign-on.

Step 3. Click **OK** to close the EAP-MSCHAPv2 Properties window.

Step 4. Change the Select Authentication Method to **Smart Card or Other Certificate**, and click the **Configure** button to bring up the properties window.

The properties window shown in Figure 4-12 give the administrator the option to select the use of a smart card or a certificate that is stored on the local computer. Along with that choice is a selection for simple certificate selection. A windows machine might be storing many different certificates that have multiple purposes. By enabling simple certificate selection, the list of certificates will be filtered down to only identity certificates, thereby simplifying the experience of choosing the correct certificate.

https://t.me/learningnets

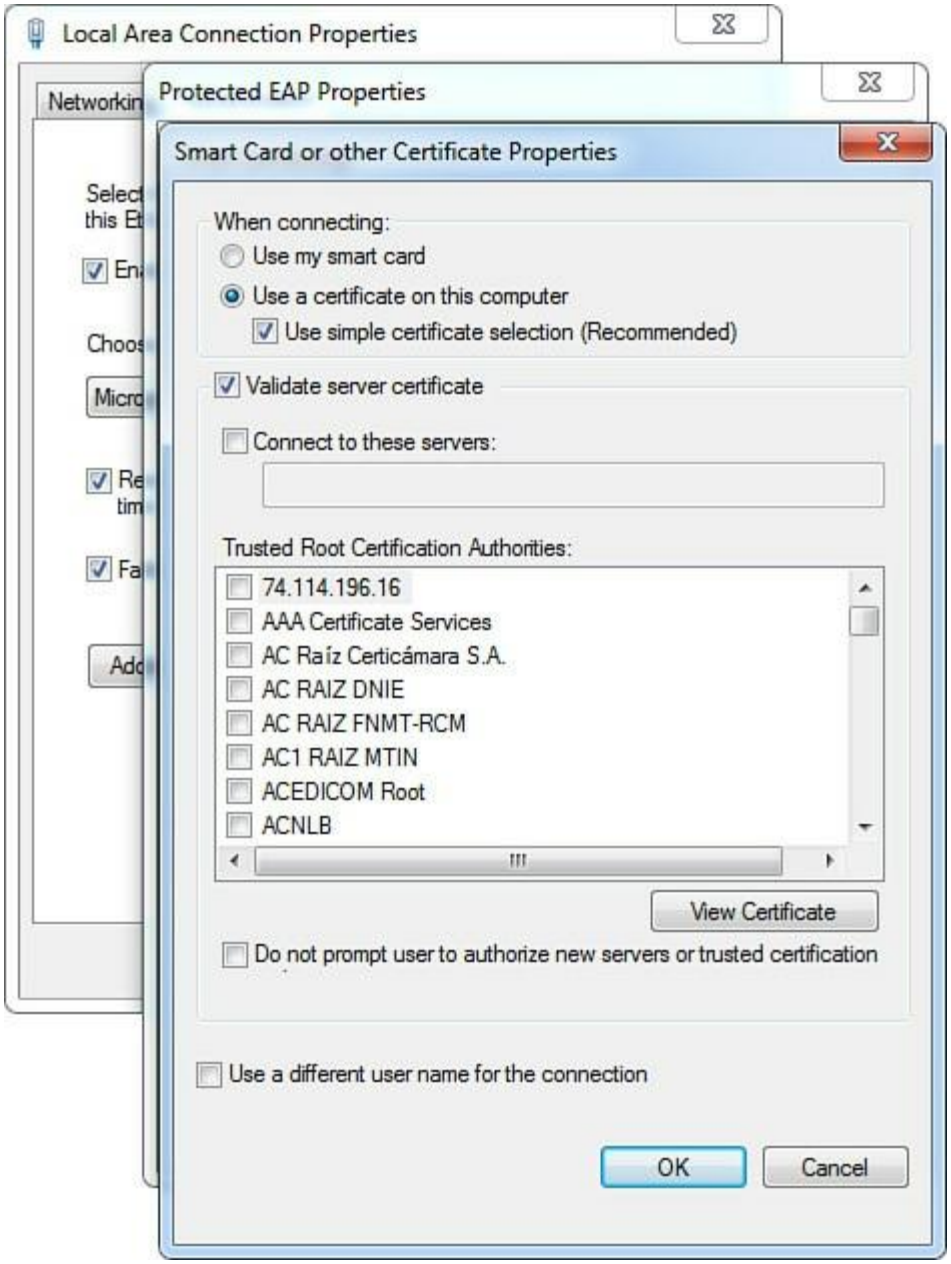


Figure 4-12 EAP MSCHAPv2 properties.

Just like the outer method of PEAP, this inner method of (EAP-TLS) has the options to validate server certificates, specify specific servers to trust, allow the administrator to select trusted root certificate authorities, and control the prompting to trust new servers. There is even an ability to specify a different username for this connection (inner identity).

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

Step 6. Return the Select Authentication Method to the default of Secured Password, as shown in Figure 4-13.

https://t.me/learningnets

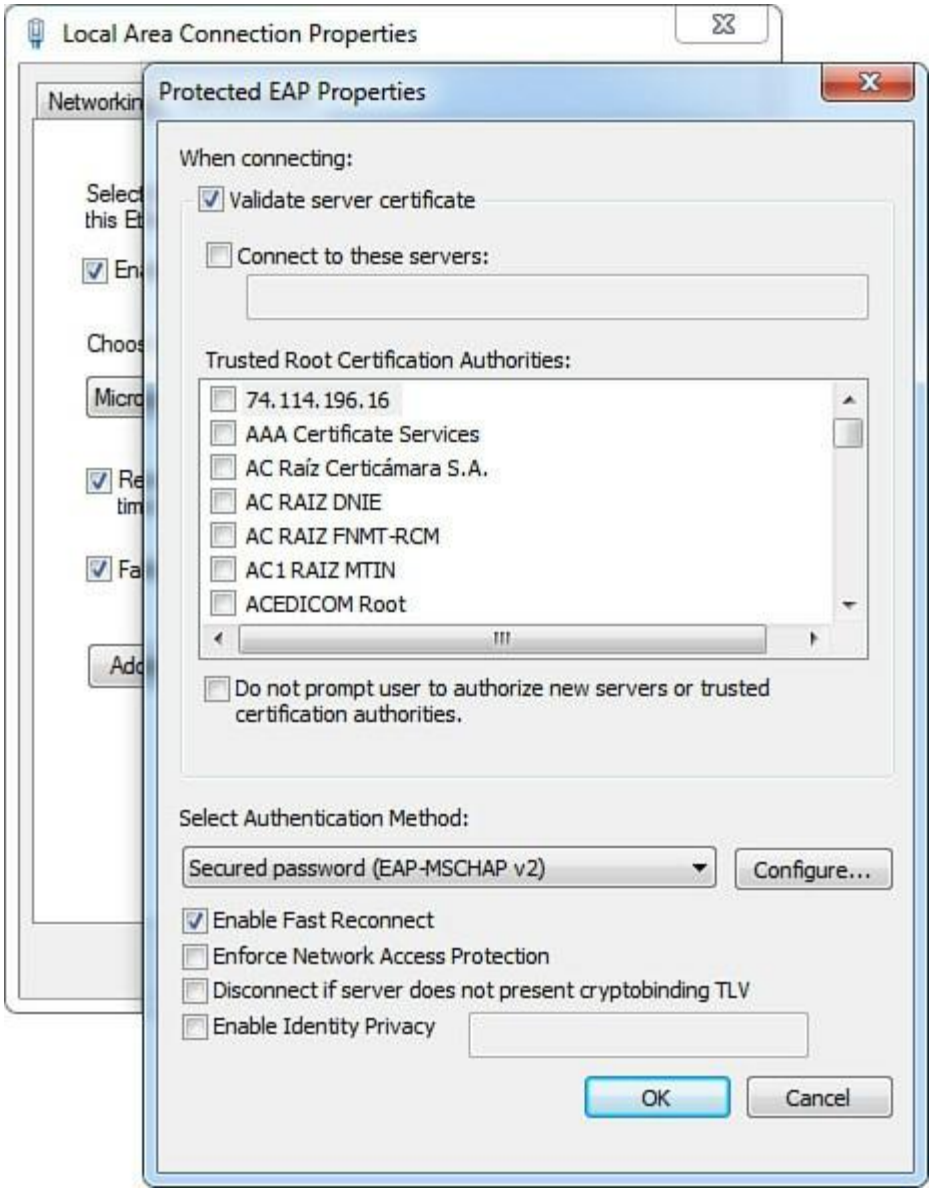


Figure 4-13 Back to the protected EAP properties.

The only other relevant setting in the Protected EAP properties page is the Enable Identity Privacy check box. This option allows the administrator to set the outer identity. Leaving this unchecked causes the outer identity to be set to anonymous.

Step 7. Click **Cancel** to return to the Authentication tab, as shown in Figure 4-14.

https://t.me/learningnets

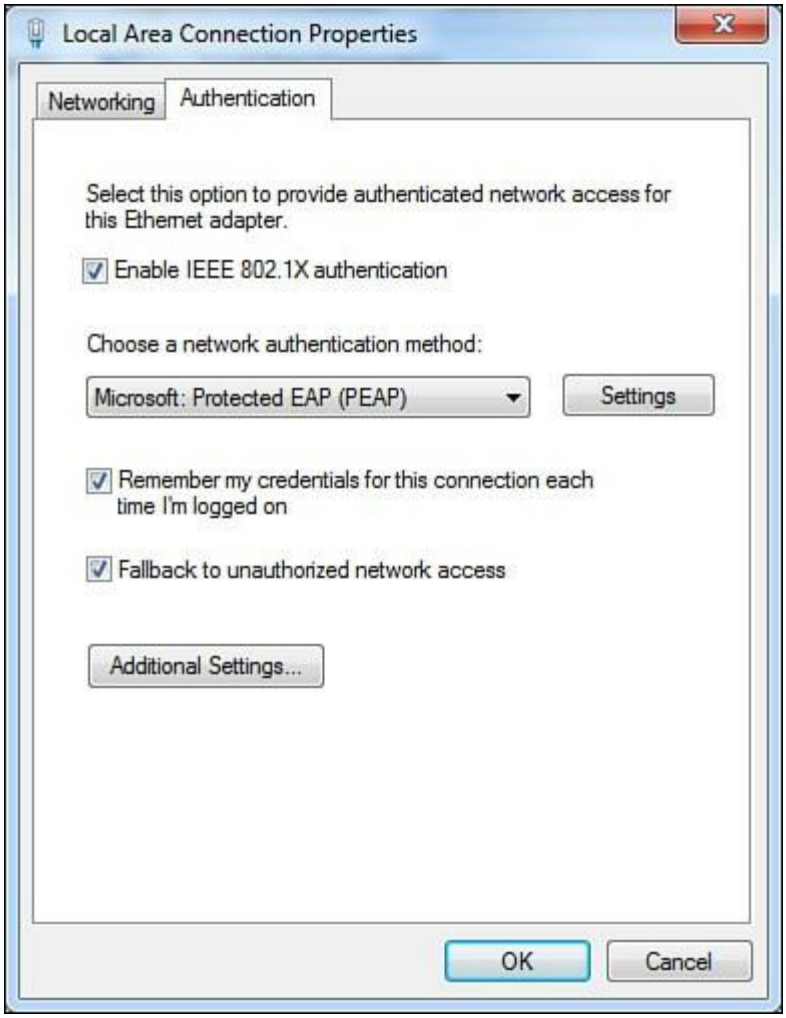


Figure 4-14 Back to the Authentication tab.

The only remaining option to cover is the **Additional Settings** button. This is where the supplicant has some additional authentication options related to mode and sign-on timing.

Step 8. Click the **Additional Settings** button. The advanced settings are displayed as shown in Figure 4-15.

https://t.me/learningnets

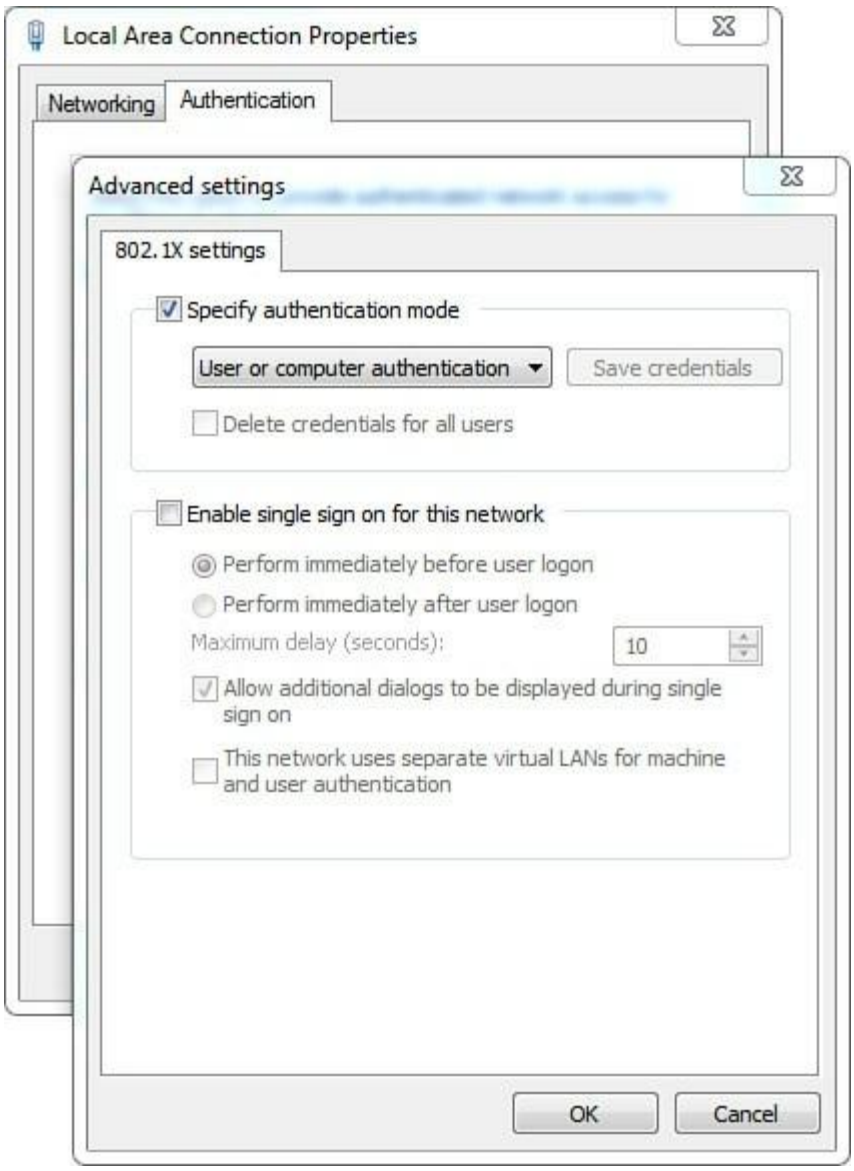


Figure 4-15 Advanced Settings tab.

Let us cover these settings in reverse order, starting with Enable Single Sign On for This Network first. This setting enables the administrator to allow or disallow the logging in to the workstation before the 802.1X authentication occurs. Selecting **Perform Immediately Before User Logon** sends the EAPoL_Start message and allows 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 logging the user off.

The Perform Immediate After User Logon option enables 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. One last choice on the bottom half of this tab is This Network Uses Separate Virtual LANs for Machine and User Authentication. This option will make more sense after you have read the next section, but ultimately 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, which is the Specify Authentication Mode check box. The options under this check box are User or Computer Authentication, Computer Authentication, User Authentication, and Guest Authentication (see Figure 4-16).

https://t.me/learningnets

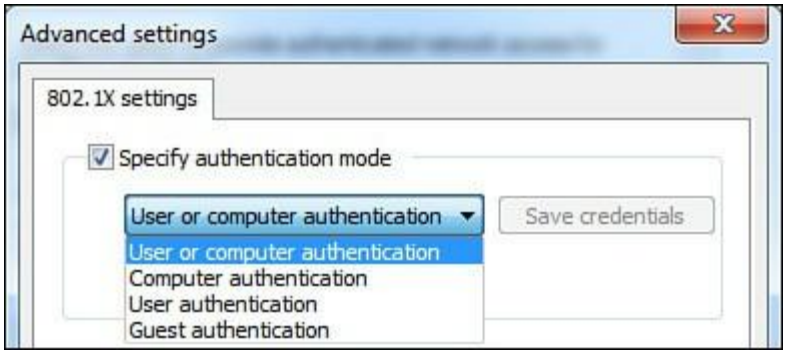


Figure 4-16 Authentication mode choices.

This concludes our guided tour through the Windows native supplicant, and it is the perfect time to introduce a new topic that is specific to Microsoft Windows operating systems: user authentication versus computer authentication.

User Authentication

This is the normal authentication that one thinks of when discussing 802.1X and network access control. It provides the identity credentials of the user to the authentication server, allowing for role-based access control to the network. The buzz on the street with the 802.1X standard was all about knowing who was accessing your network before allowing them onto the network.

User authentication on a Microsoft Windows device is able to use a username, password, user-issued certificate, or even a smart card. With these Windows devices there is a separate certificate store for user-issued certificates. Each user who logs in to a Windows workstation will have her own certificate store, and therefore authenticate differently to the network.

Machine Authentication (a.k.a. Computer Authentication)



Microsoft Windows workstations have a very powerful management system, known as Active Directory. Maybe you've heard of it? Well, AD needs to remain in contact with the computers it manages for policy updates and other management tasks. The IEEE 802.1X standard came out, and with no user logged in to the computer, no network access was granted. This concept broke the communication between the AD managed computer, as well as not allowing the user's Kerberos authentication to reach Active Directory, which didn't allow the credentials to be passed down to the supplicant itself—especially in the instances where the user's password was expired.

This should sound an awful lot like a denial of service (DoS) because it truly was. Microsoft (quite brilliantly) designed a way to solve this problem and prevent this chicken-and-egg scenario. They created multiple states for their supplicant: a machine state and a user state.

Whenever there was no interactive user logged in to the workstation (that is, no one had pressed CTRL-ALT-DEL 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 into the network, triggering a whole new authentication of the user's credentials.

Active Directory negotiates a password with a Windows workstation when it joins AD, which the machine-state supplicant is able to use as the credential. Additionally, the computer maintains its own system store for certificates that is separate from the users' certificate stores. Therefore, machine authentication (also called computer authentication) is capable of using a computer name and password (PEAP-MSCHAPv2) or a machine certificate (EAP-TLS, PEAP-EAP-TLS).

Figure 4-17 illustrates the computer starting with machine authentication, and then a user logging in to the system, thereby causing a new authentication with user authentication.

https://t.me/learningnets

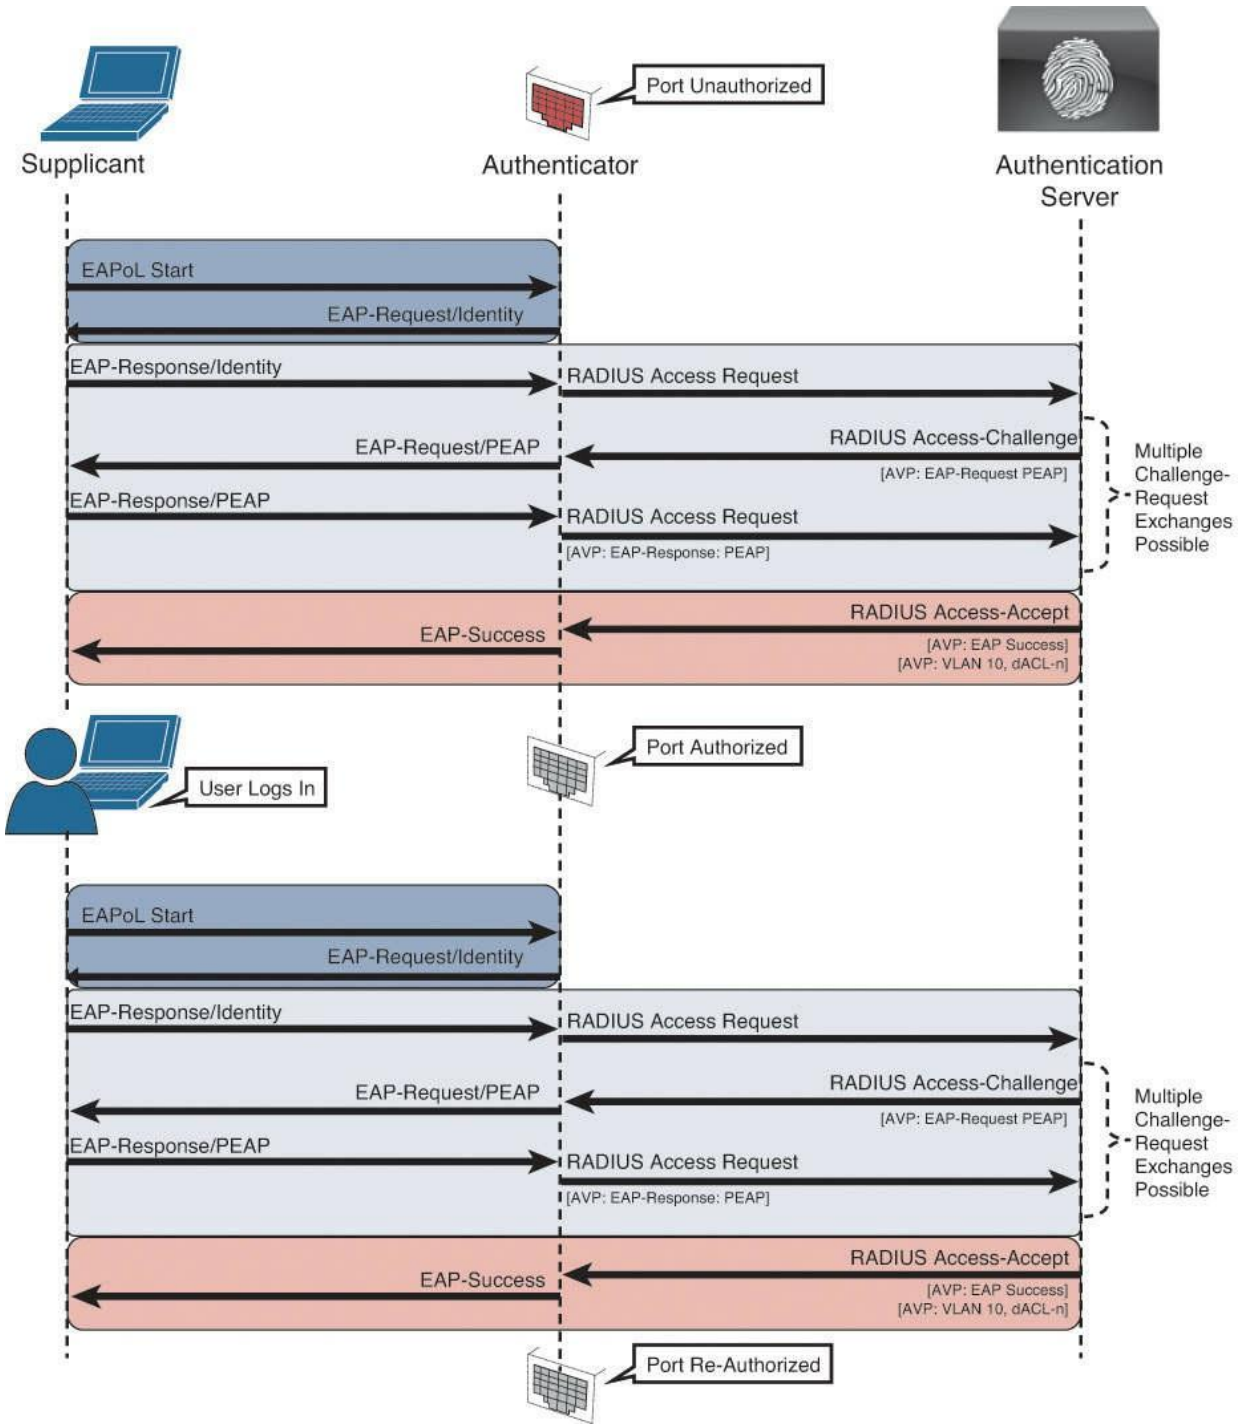


Figure 4-17 Multiple supplicant states.

This concludes the section on the Windows native supplicant, as well as the explanation of user and machine authentications.

https://t.me/learningnets

Cisco AnyConnect NAM Supplicant



The Cisco AnyConnect Secure Mobility Client (also called AnyConnect) is another one of those products that 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 itself is modularized with several modules in existence today that 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 4-18 shows the module choices when installing AnyConnect on a Windows system.

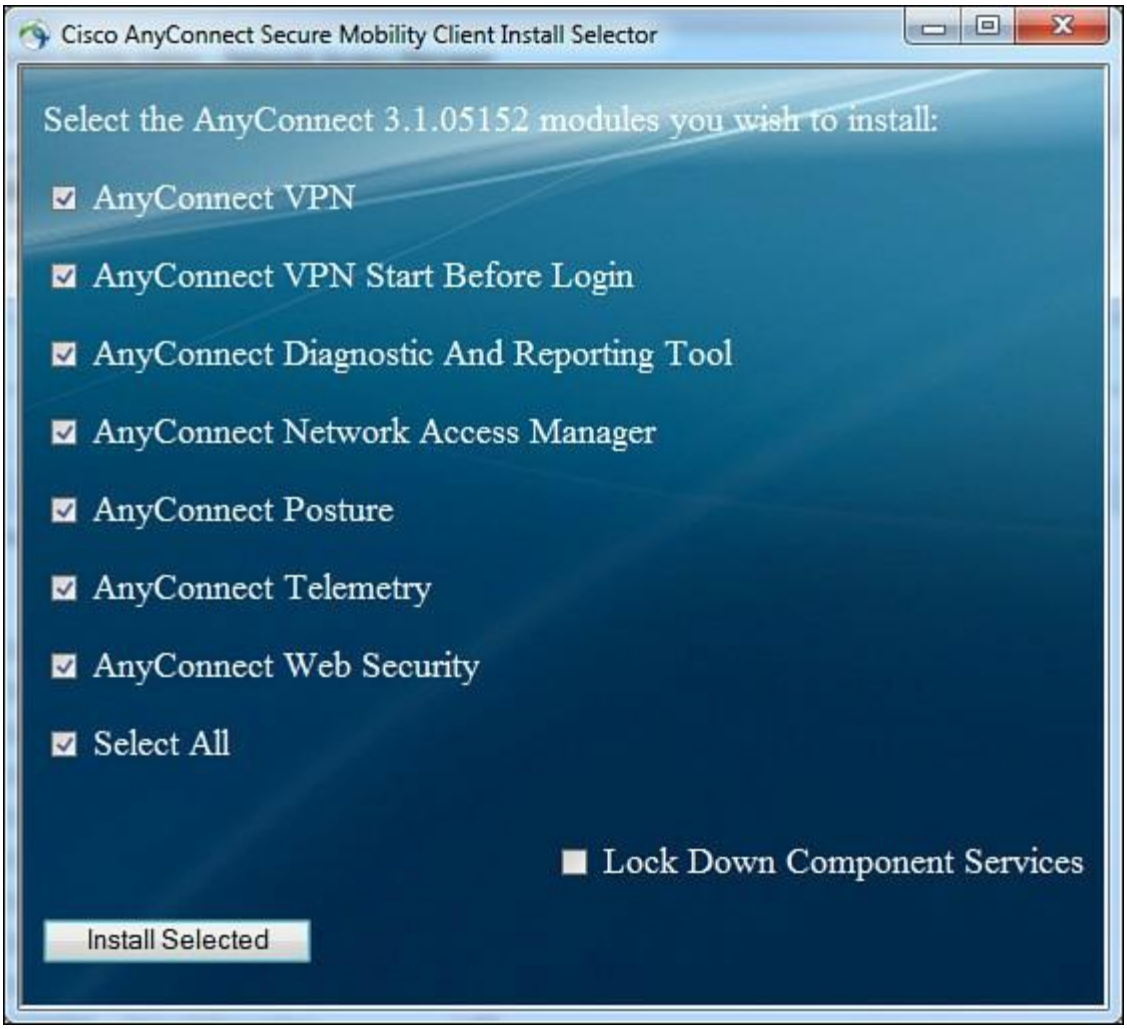


Figure 4-18 Cisco AnyConnect Secure mobility client modules.

The key modules that are of interest to Cisco ISE and this exam in particular are the AnyConnect Network Access Manager (NAM) and the AnyConnect Diagnostics and Reporting Tool (DART). To provide you with a little bit of history, many years ago in a galaxy far, far away there was a product known as the Cisco Trust Agent. This product has since been announced as end of life, so don't waste any time memorizing the name. The point of interest of this product is that it contained a supplicant in it that was OEM'd from a company named Meeting House. Since that time, Cisco has acquired Meeting House and packaged its supplicant as the Cisco Secure Services Client (CSSC).

Why have you bothered to read that little tidbit of history that only this author finds interesting? Because that 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

https://t.me/learningnets

supplicant got integrated into the AnyConnect client as the Network Access Module and is licensed at no charge.

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 Cisco's ASA itself.

Let's walk through the configuration of AnyConnect NAM. The Standalone Profiler Editor enables administrators to build configurations; the end user will never see these screens. The editor contains the following views:

- **Client Policy**—This view is for configuration options for connection, wired/wireless/3G mobile broadband, and end user and administrative settings. Many of the items in this section will look familiar as many of the options were 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 which 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 would 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.

At this point in the chapter, we will take a guided tour of the Standalone Profiler Editor. If you have a Windows 7 or Windows 8 device at your disposal, feel free to download the standalone editor from www.cisco.com and follow along.

Client Policy

This view is for configuration options for connection, wired/wireless/3G mobile broadband, and end user and administrative settings. Many of the items in this section will look familiar as many of the options were also in the Windows native supplicant. We will examine the options in the Client Policy view, as visible in Figure 4-19.

https://t.me/learningnets

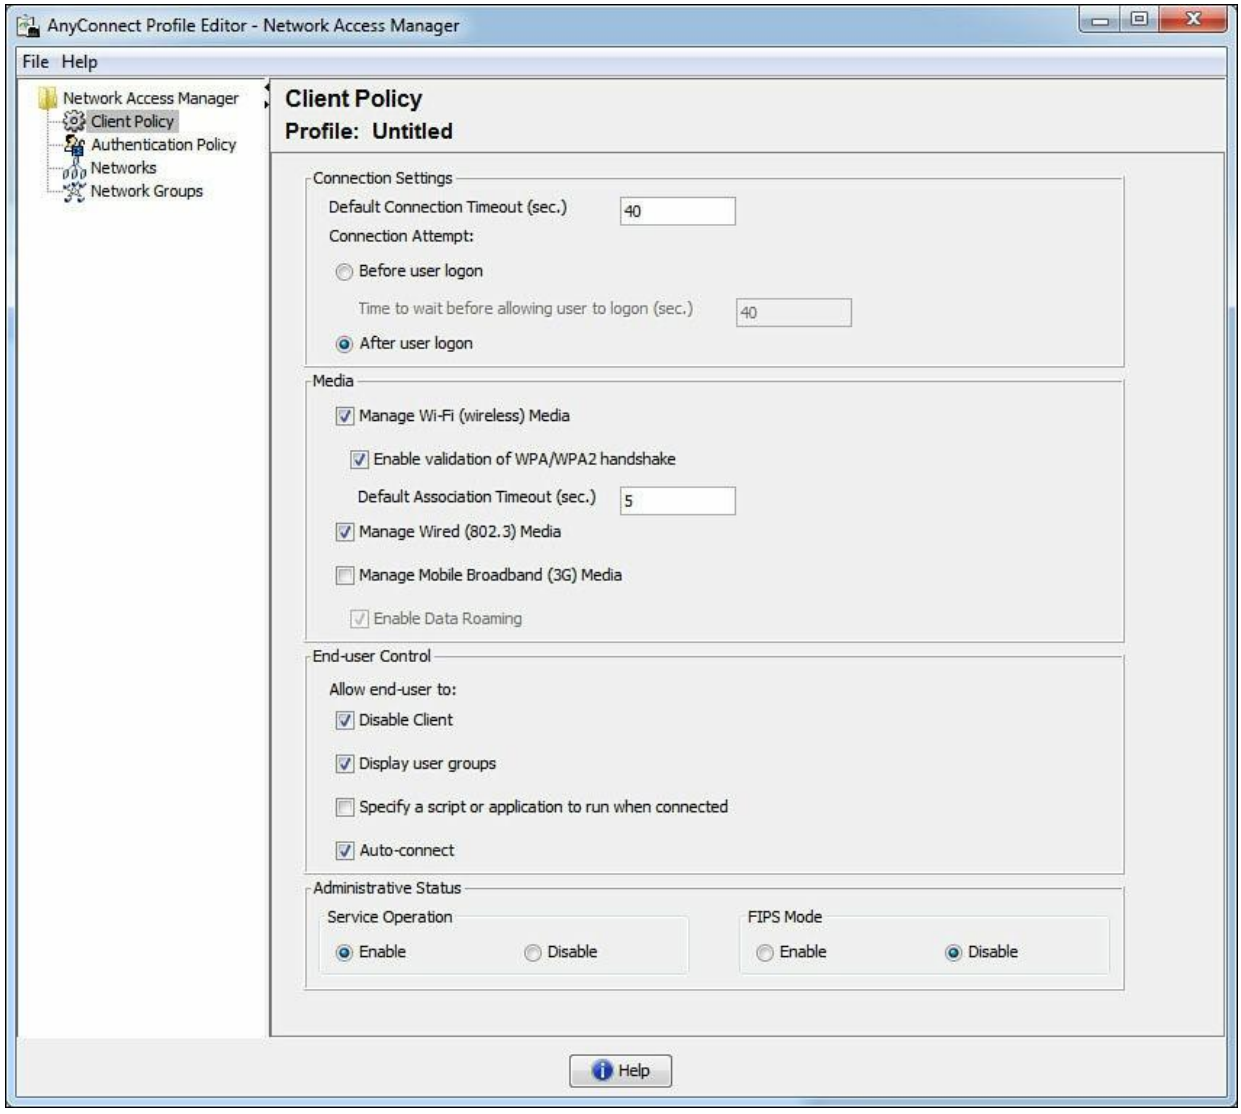


Figure 4-19 Client Policy view.

The first section, Connection Settings, is similar to the single sign-on settings with the native supplicant, where the administrator defines whether the supplicant performs authentication before or after allowing the user to interact with the desktop. You can compare these setting with those in Figure 4-15.

The next section, Media, defines if end users should be able to use AnyConnect NAM and join wired and wireless networks. Note that 3G broadband adapter cards must be Windows 7 logo certified; otherwise, the 3G cards will be incompatible.

With the End User Control section, the administrator defines if the end user will be allowed to 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 a connection. Even though they do not correspond to administrator-defined groups, they appear under the local group. User-created networks defined in AnyConnect NAM will appear under here as well.

Specify a Script or Application to Run When Connected enables users to 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) refresh from Active Directory or to map network drives.

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

https://t.me/learningnets

Administrative Status has two sub-settings. If you disable Service Operation, 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 enable FIPS mode, the NAM performs cryptographic operations in a way that meets the government requirements.

> **Note**
>
> There are many more complexities to FIPS mode; simply enabling this setting is not enough. Please see the official Cisco AnyConnect administrative documentation for more information on FIPS mode.

Authentication Policy

This view is where an administrator can restrict which types of networks the end user is permitted to connect to when not at the corporate location. Some organizations are very specific about which security levels are required for their corporate owned assets, and this is one way of controlling that. As you will see in Figure 4-20, the settings on this tab let the administrator get very specific on the types of wireless key management (WPA/WPA2) types to which an end user can connect, as well as very specific on the EAP types that may be employed on those networks.

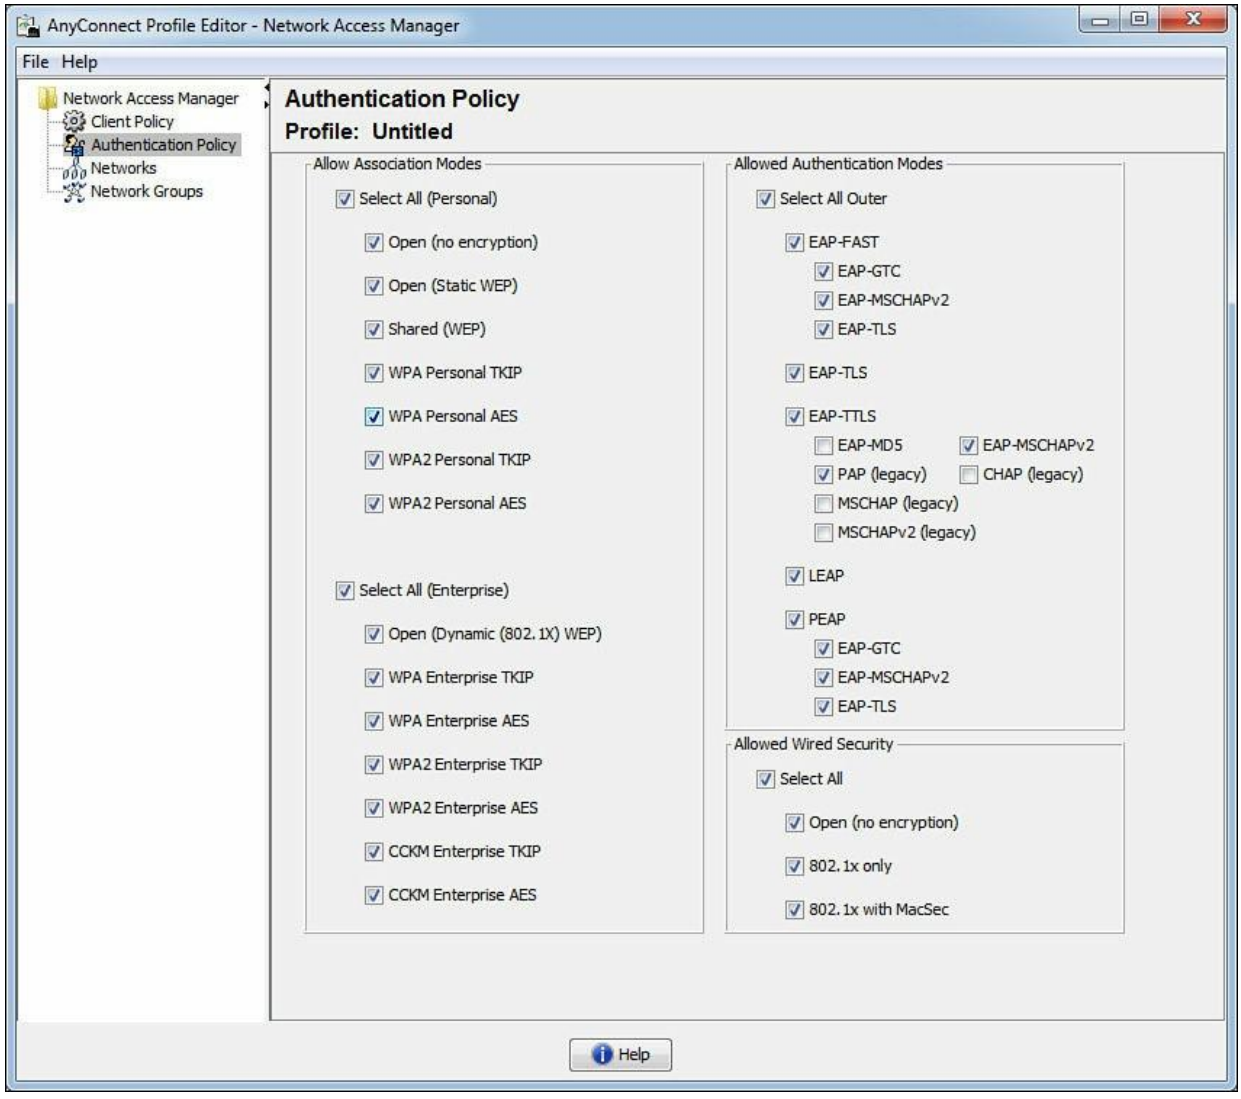


Figure 4-20 Authentication policy.

https://t.me/learningnets

Networks

This view is where the administrator defines the corporate network and the security to use on the corporate wired and wireless networks. As shown in Figure 4-21, there is 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, nonauthenticating wired network.

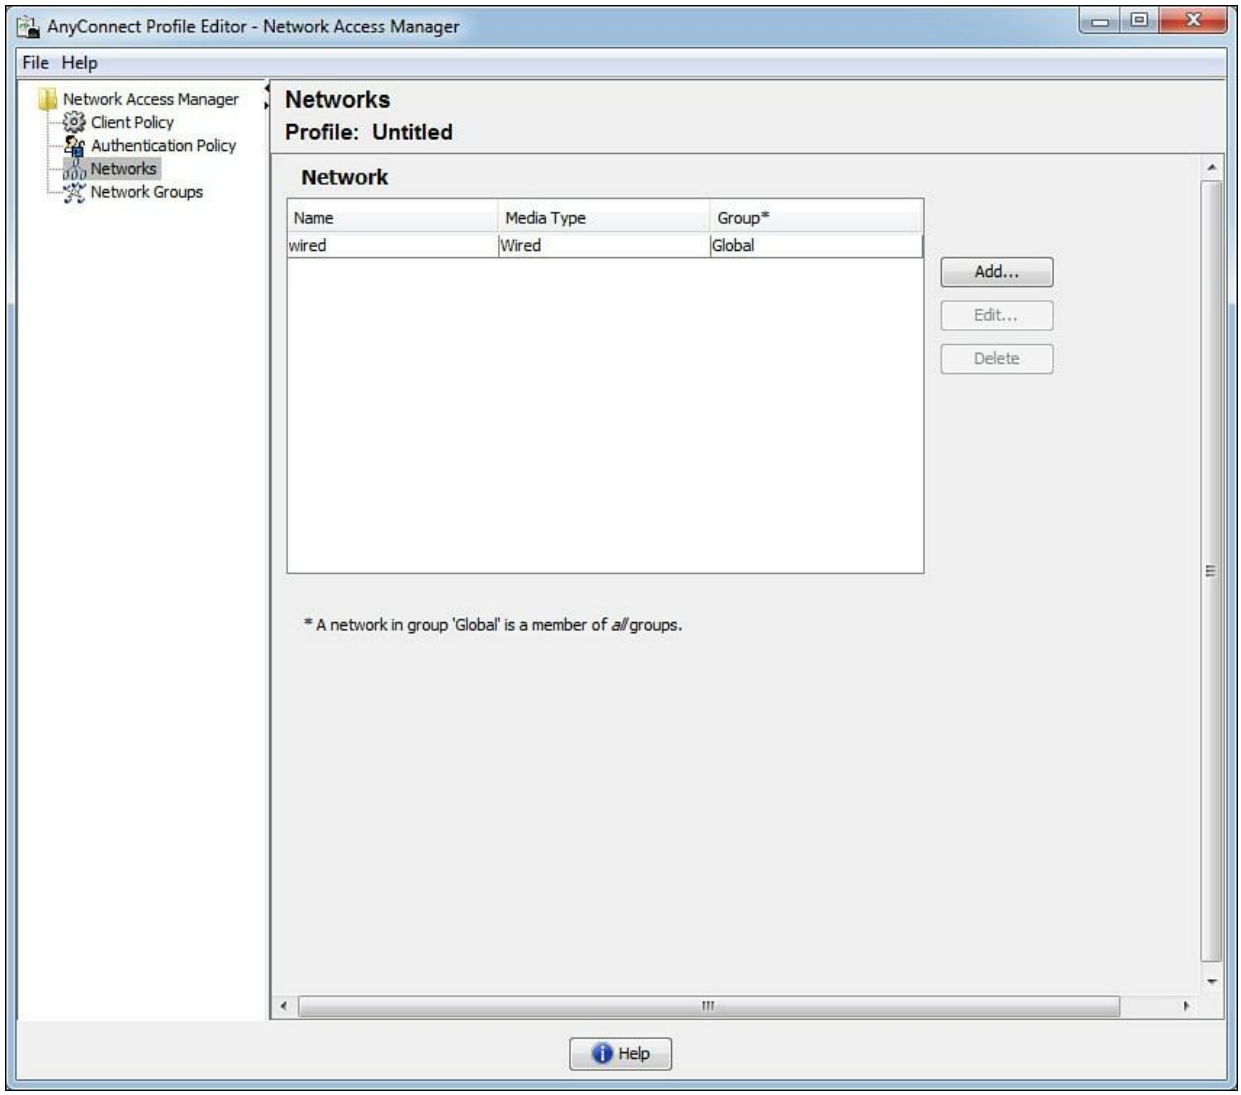


Figure 4-21 Networks tab.

Where things get interesting is when we click Add. Clicking this button starts a wizard that walks the administrator through creating administratively defined networks, as shown in Figure 4-22.

https://t.me/learningnets

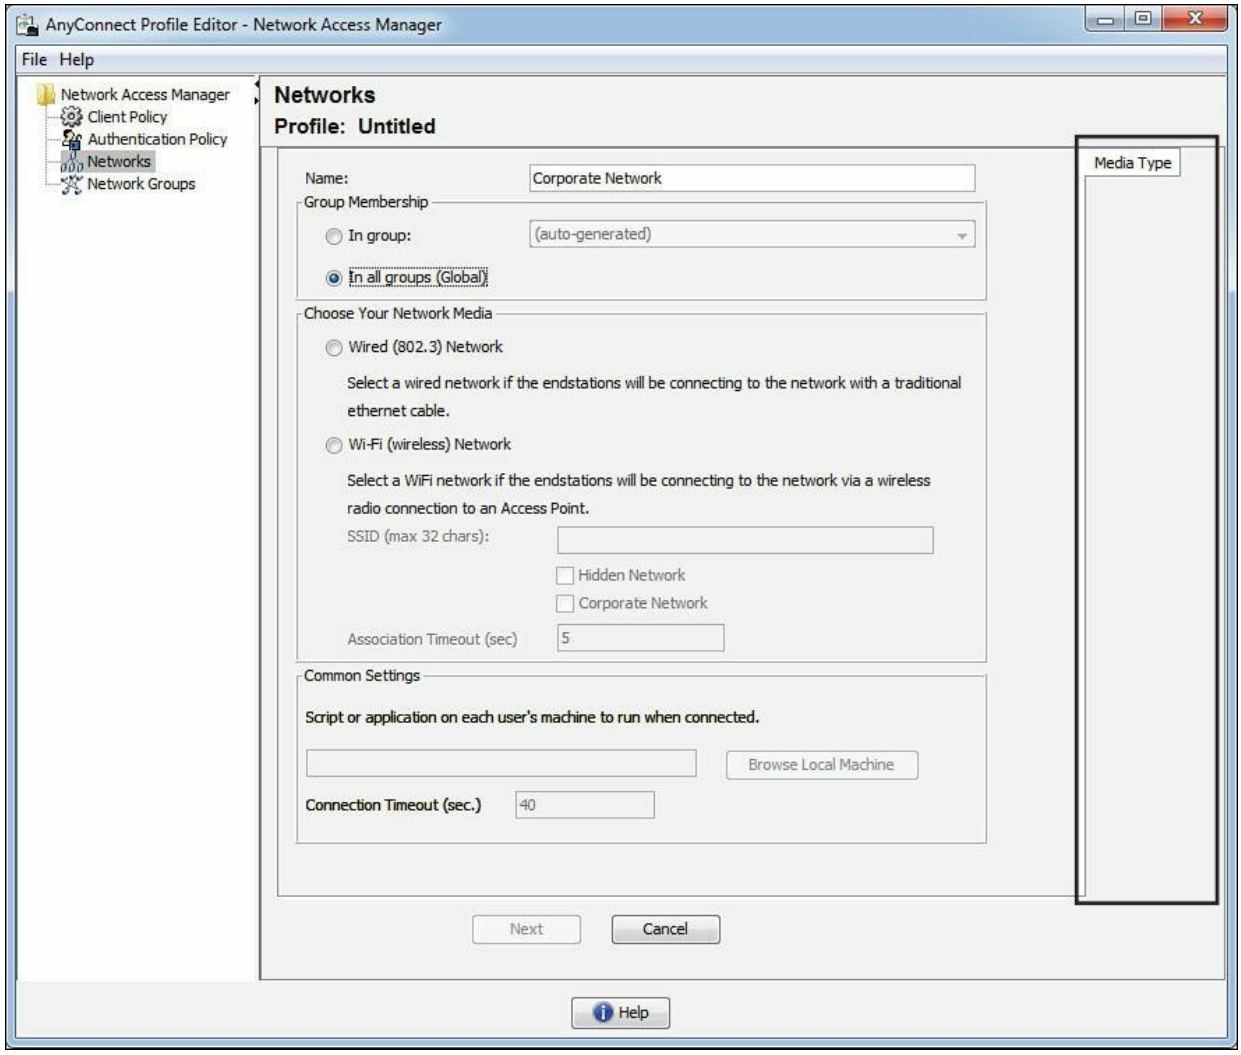


Figure 4-22 Adding a new network.

As an administrator builds the network profile, the wizard dynamically adds new tabs to the right side, which is highlighted in Figure 4-22. Selecting a wired (802.3) network adds a security level tab, as shown in Figure 4-23.

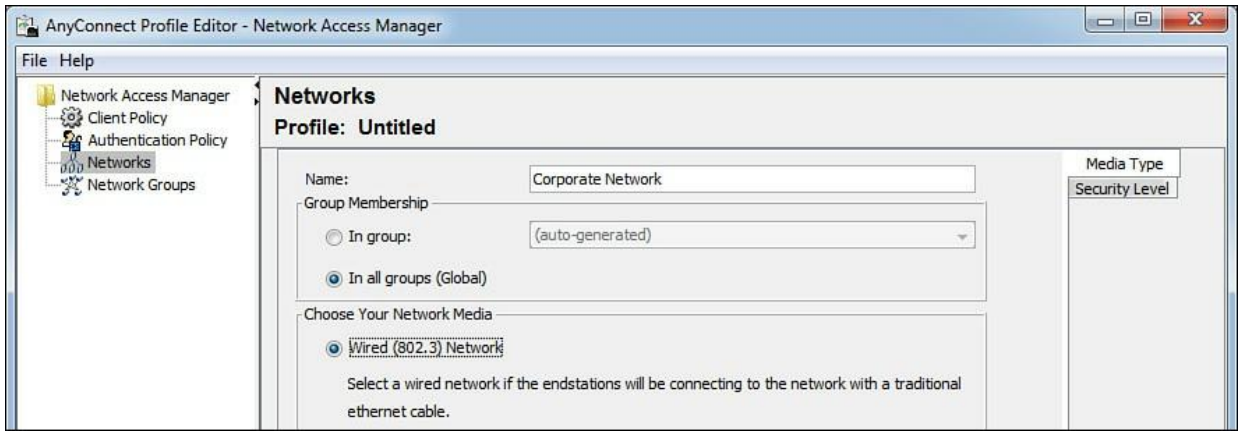


Figure 4-23 Adding the Security Level tab.

Clicking **Next** will bring you to that next tab, or you can click the tab directly. The result is shown in Figure 4-24.

https://t.me/learningnets

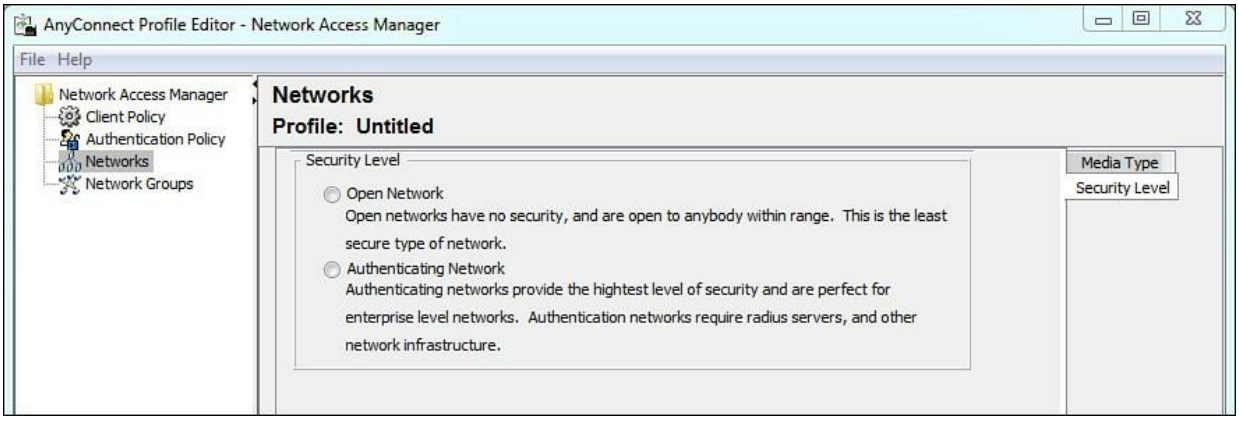


Figure 4-24 Security Level tab.

Because we are talking about corporate networks and this is an exam about secure network access, we will certainly be choosing the authenticating network option. As this option is selected, notice the Connection Type tab that appears in the right margins. Figure 4-25 shows the Security Level settings now that the profile is configured for an authenticating network.

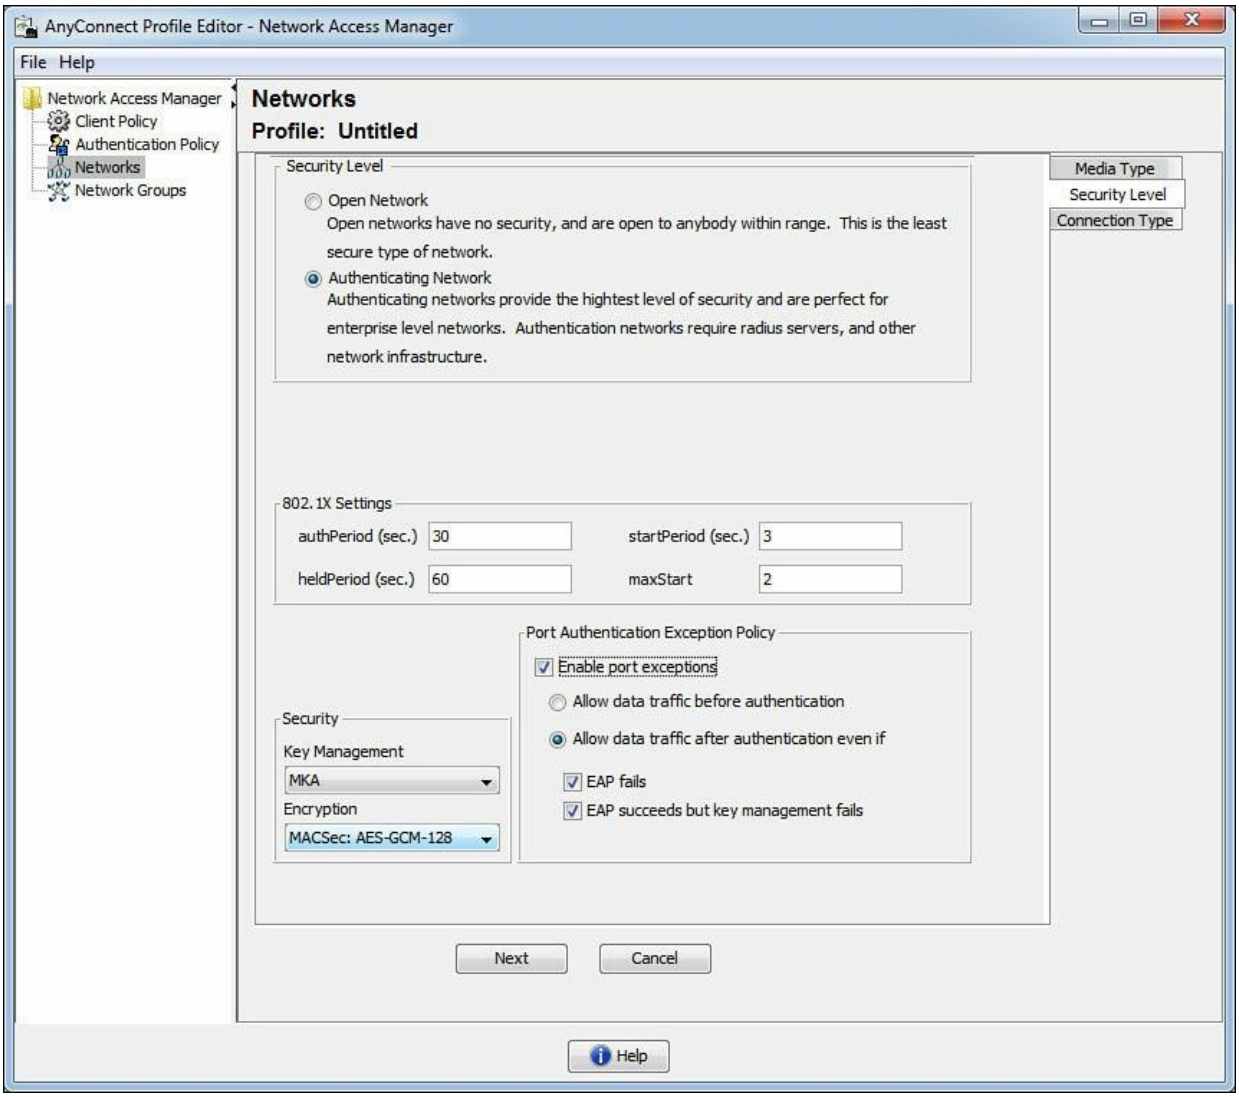


Figure 4-25 Security Level tab with an authenticating network.

https://t.me/learningnets



Notice the 802.1X timers that are available for the administrator to tune. Also pay attention to the Security section where Key Management can be turned on for Layer-2 MACSec encryption between the supplicant and the switch. This provides AES-GCM-128-bit encryption over the wire. Lastly, when enabled, the Port Authentication Exception Policy 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, like certain government institutions, require such strict controls that they demand the ability to deny traffic if it cannot be encrypted.

Clicking Next brings us to the Network Connection Type. The settings here should seem familiar becuse they are asking for machine authentication, user authentication, or both. Selecting Machine and User Connection adds four new tabs to the right side. Clicking Next opens the Machine Auth tab.

Selecting the EAP method populates the lower portion of the tab, as shown in Figure 4-26. In the figure, we have chosen EAP-FAST for the tunneled EAP method, and you can see just how much control is given to the administrator for customizing the behavior of the supplicant. The inner method is fully selectable for EAP-MSCHAPv2 or EAP-GTC, or EAP-TLS (authenticate using a certificate).

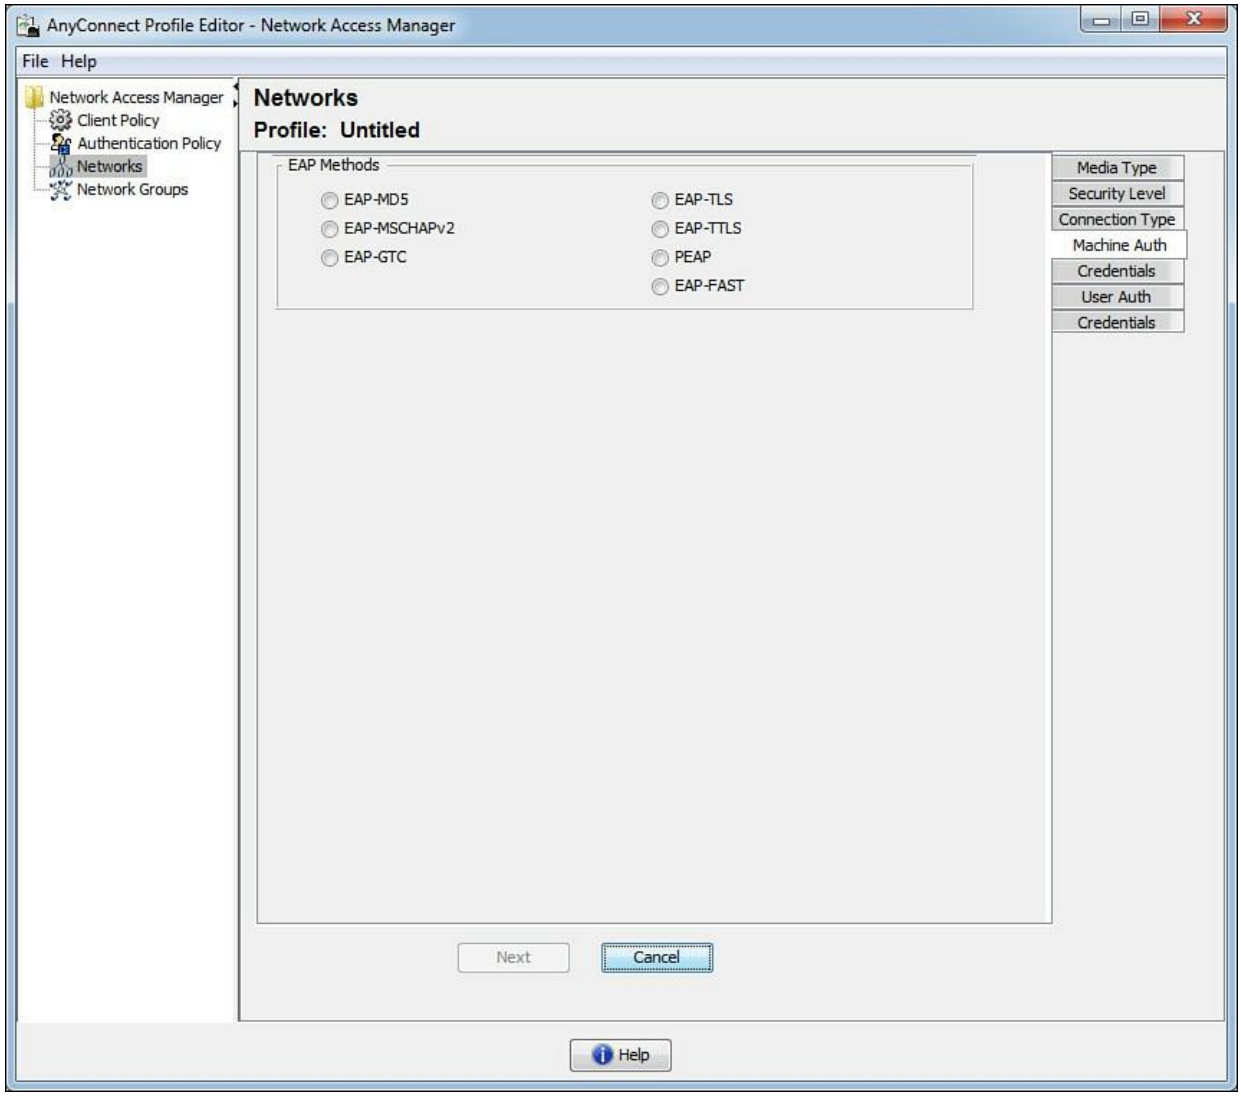


Figure 4-26 Machine Auth tab.

Clicking Next opens the Certificates tab, as shown in Figure 4-27. Why the Certificates tab? We chose a

https://t.me/learningnets

password based authentication method, didn't we? This tab is defining the certificates to trust for the outer method, the forming of the TLS tunnel. AnyConnect NAM provides powerful and flexible rules to filter which certificates to trust, which is beyond the scope of this exam and book. It also provides the choice to the administrator to trust any Root CA already trusted by the operating system or to specify the trusted Root CA in the profile.

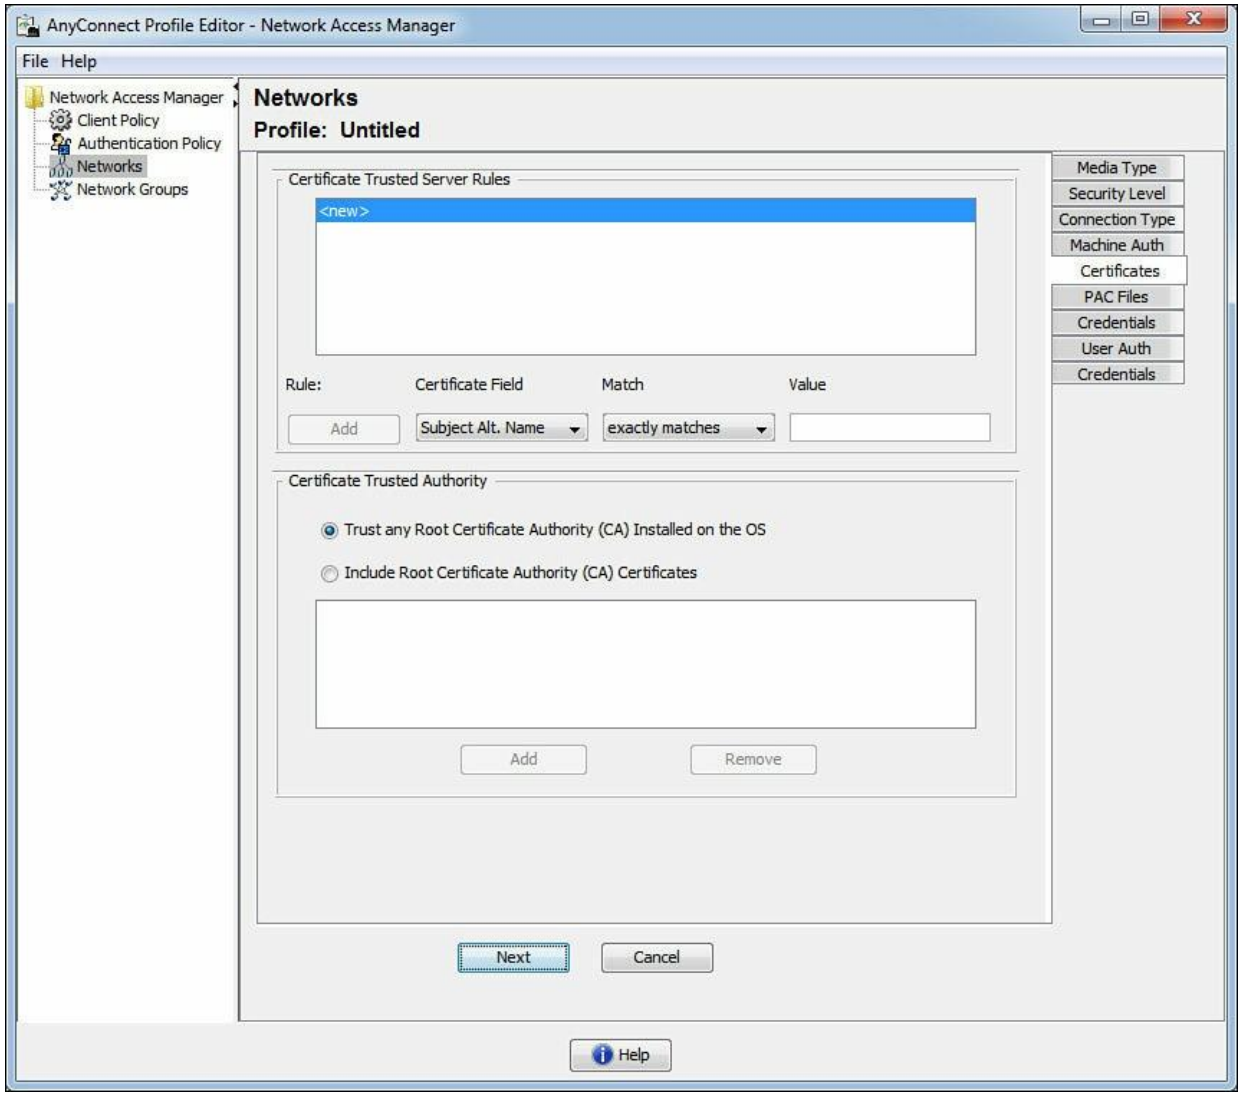


Figure 4-27 Machine Auth: Certificates tab.

Clicking Next opens the PAC files tab, as shown in Figure 4-28. This is specific to EAP-FAST, which has the option of using PAC files. A PAC file can be viewed a lot like a secure cookie. It's 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 provides the administrator a way to distribute the server's PAC file in the NAM profile for secure tunnel establishment without the need for certificates.

https://t.me/learningnets

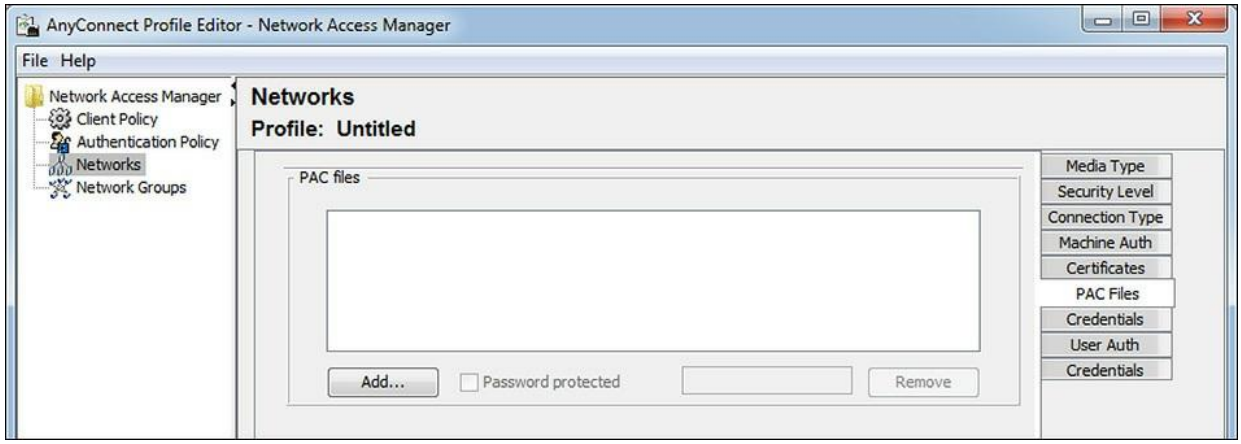


Figure 4-28 Machine Auth: PAC Files tab.

Click Next to get to the Credentials tab, as shown in Figure 4-29. This tab should also reinforce the concept of the inner and outer identity with a tunneled EAP. Because this is EAP-FAST, there is a notion of both. The Unprotected Identity is the outer identity. This is a machine authentication, so the outer identity begins with `host/`, and you can see that the outer identity is set to anonymous by default. An administrator 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, a rule could be configured 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 actual Active Directory machine account name. With the machine credentials, the default is to use the machine's password that was autogenerated when the computer joined AD. However, if something special was required, the administrator has the ability to enter a static password.

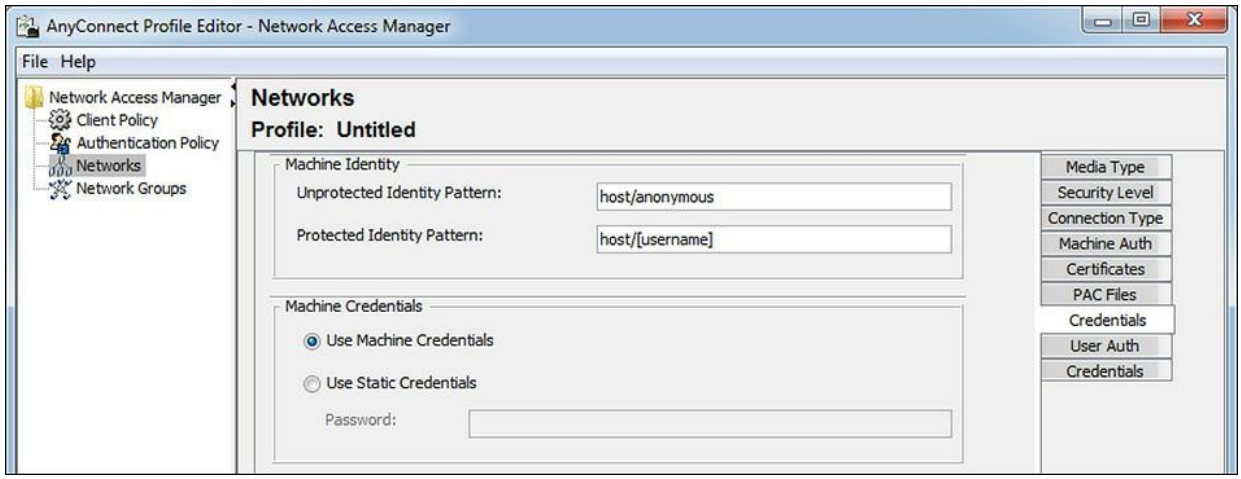


Figure 4-29 Machine Auth: Credentials tab.

Clicking Next opens the first tab of the user authentication, as shown in Figure 4-30. Take notice of the concept that a different EAP method could be selected for the user authentication. It does not have to be the same as the machine authentication. If you select EAP-FAST, you will notice all the same choices are available as were with the machine auth, 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.

https://t.me/learningnets

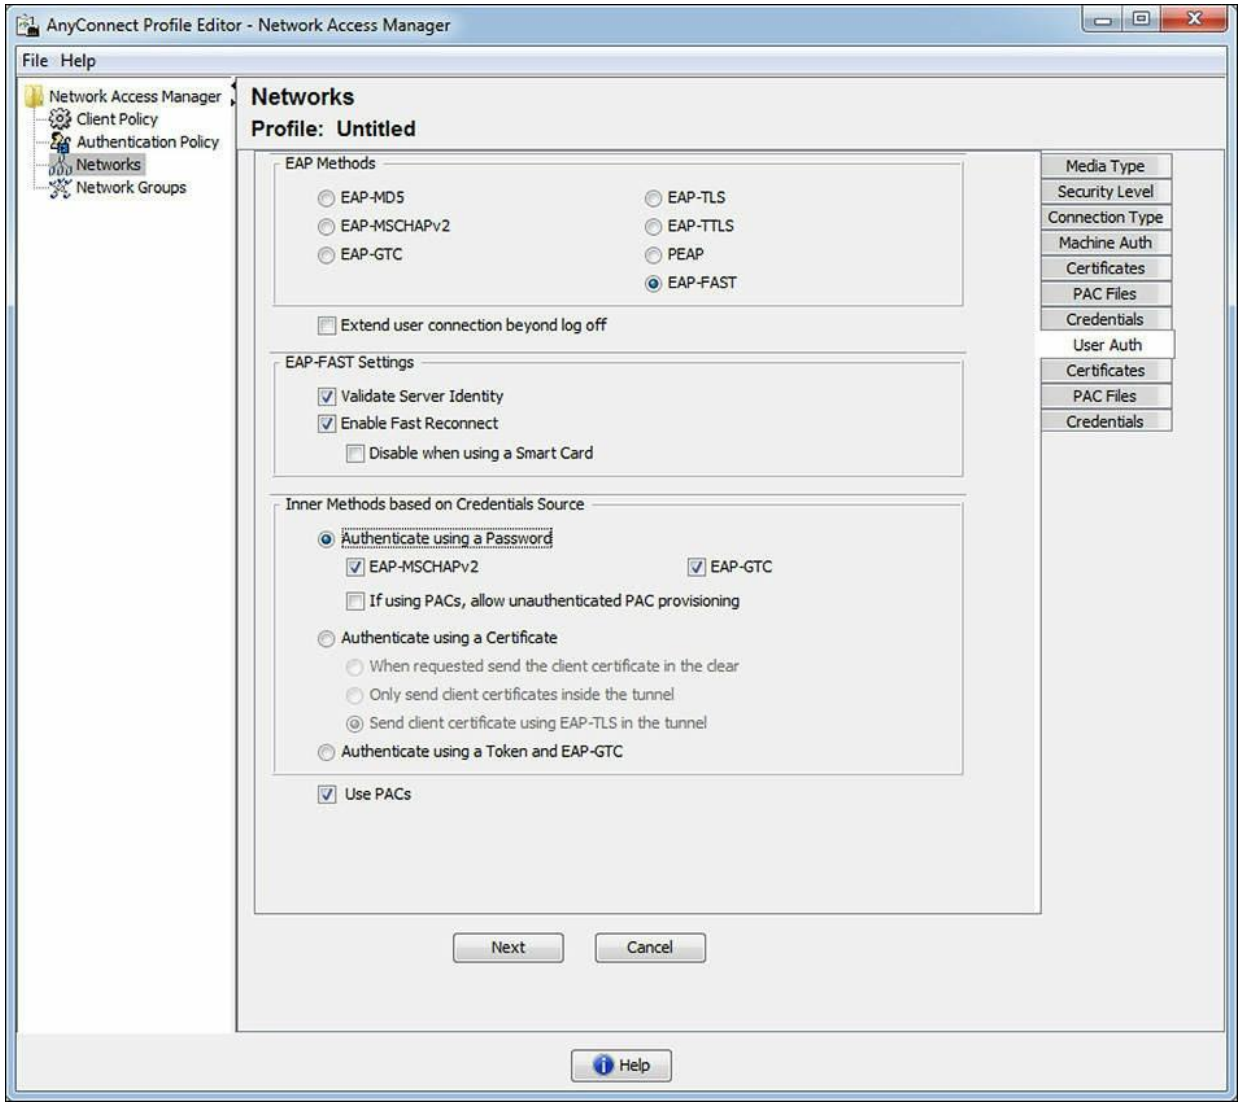


Figure 4-30 User Auth tab.

Clicking Next opens the Certificates tab for the user certificates. Just like the machine certificates, the administrator is provided with options to specify certificates, include them with the profile, and more. Clicking Next again displays the PAC files tab for users, which again is exactly like that of the Machine Auth PAC files tab.

Clicking Next opens our final tab for the Network wizard—the User Credentials tab, which is displayed in Figure 4-31. As with the Machine Auth Credentials tab, there are fields for an unprotected (outer) identity and a protected (inner) identity. There are also options for using single-sign-on credentials, prompting for the user's credentials, and for using static credentials.

https://t.me/learningnets

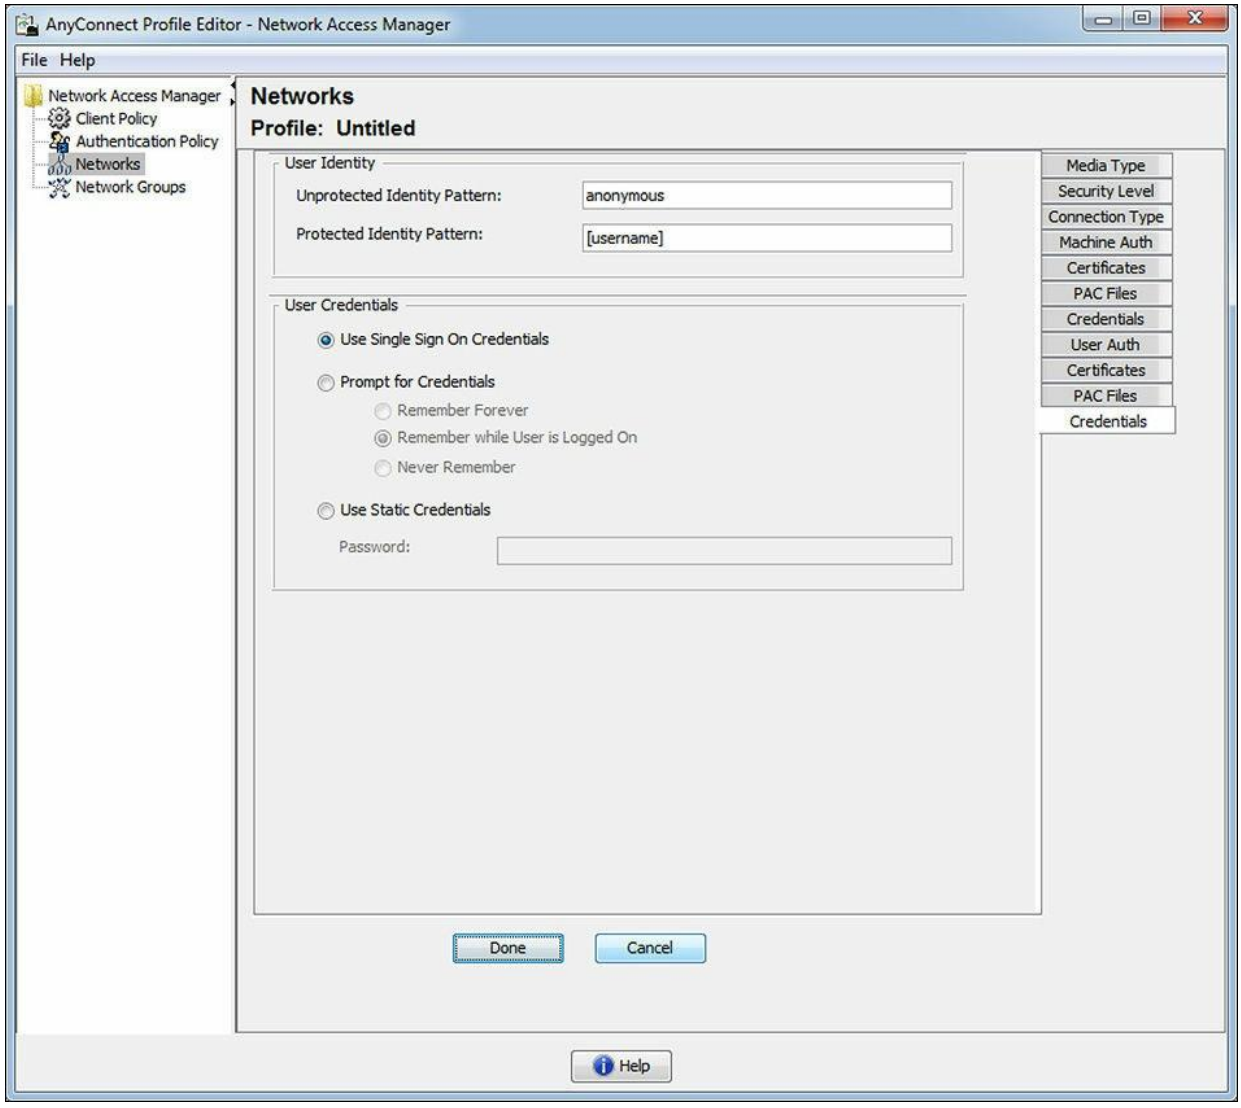


Figure 4-31 User Auth: Credentials tab.

Click **Done** and notice the new network configuration has been added to the list, as shown in Figure 4-32.

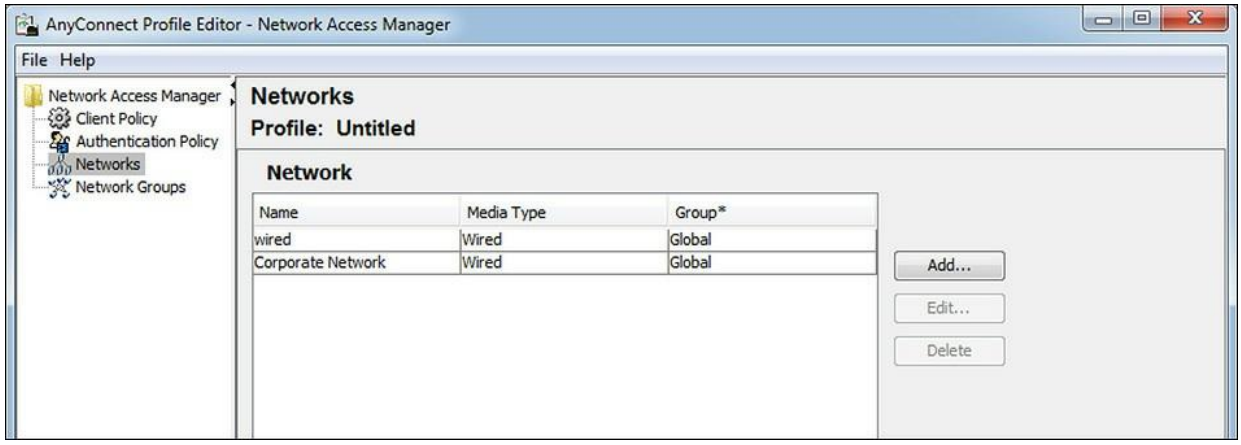


Figure 4-32 Network added.

Network Groups

As shown in Figure 4-33, the Network Groups view enables the administrator to organize network profiles into logical groupings for easier use.

https://t.me/learningnets

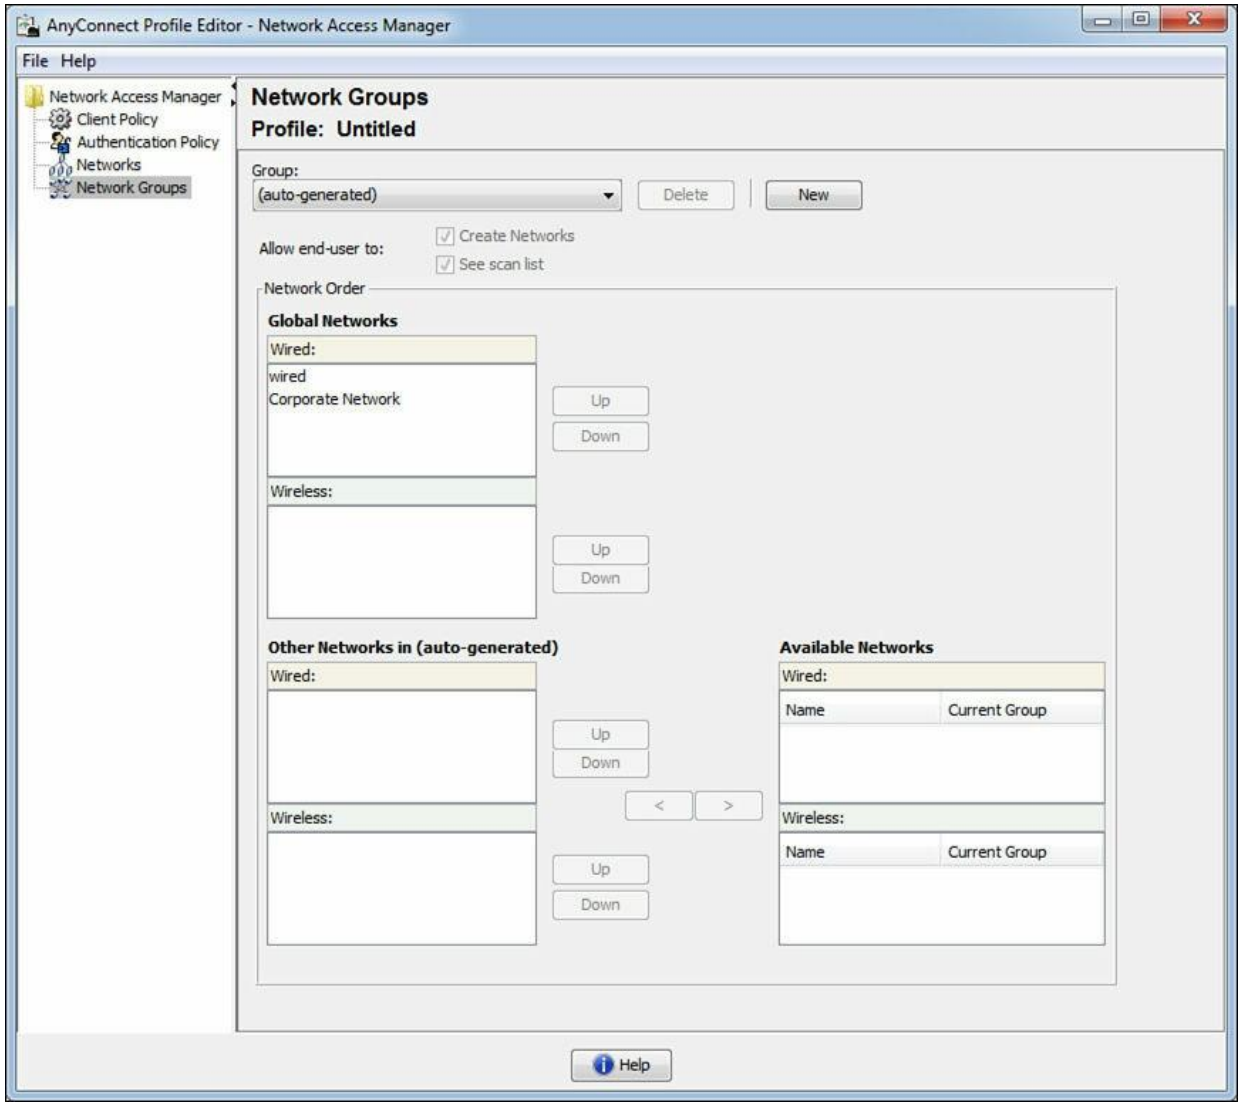


Figure 4-33 Network groups.

Implementing AnyConnect NAM Profiles

Now that 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 will ensure 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, IBM Big Fix, or another software management solution. Another option is to publish the configuration on the Cisco ASA, where the profiles get updated the next time a user VPNs into the company.

The file must be saved with a filename of configuration.xml and needs to be store in the %SystemDrive%\ProgramData\Application Data\Cisco\Cisco AnyConnect Secure Mobility Client\Network Access Manager\newConfigFiles directory, as shown in Figure 4-34. 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. The example in this paragraph will work for Windows 7 and Windows 8 with AnyConnect version 3.1.

https://t.me/learningnets

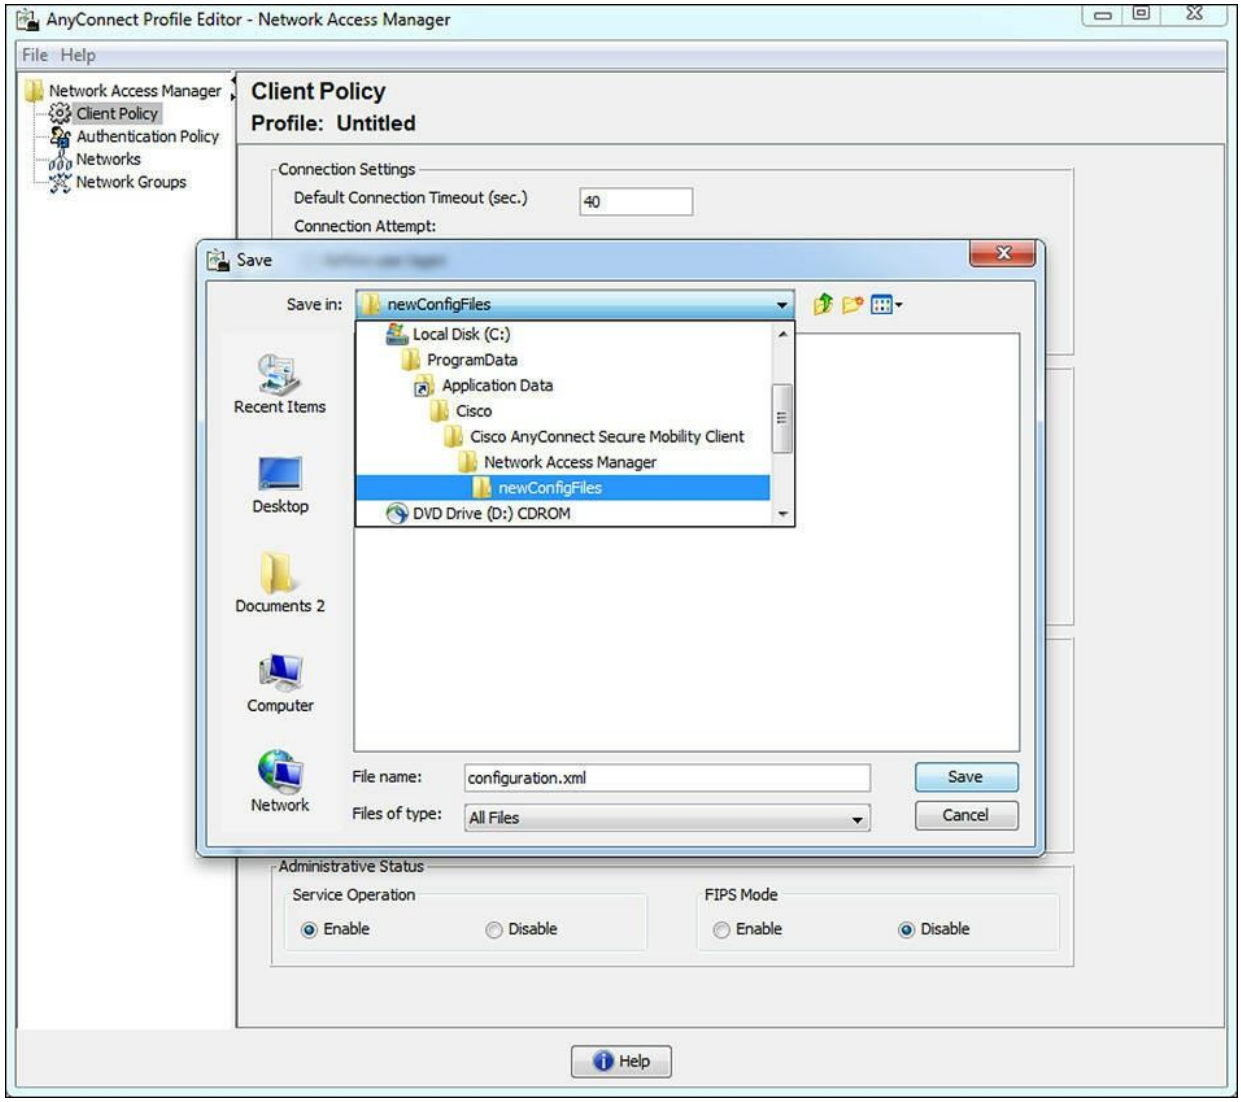


Figure 4-34 Saving configuration.xml.

The end user experience is clean and simple. The GUI with which the end user interacts is shown in Figure 4-35.

https://t.me/learningnets

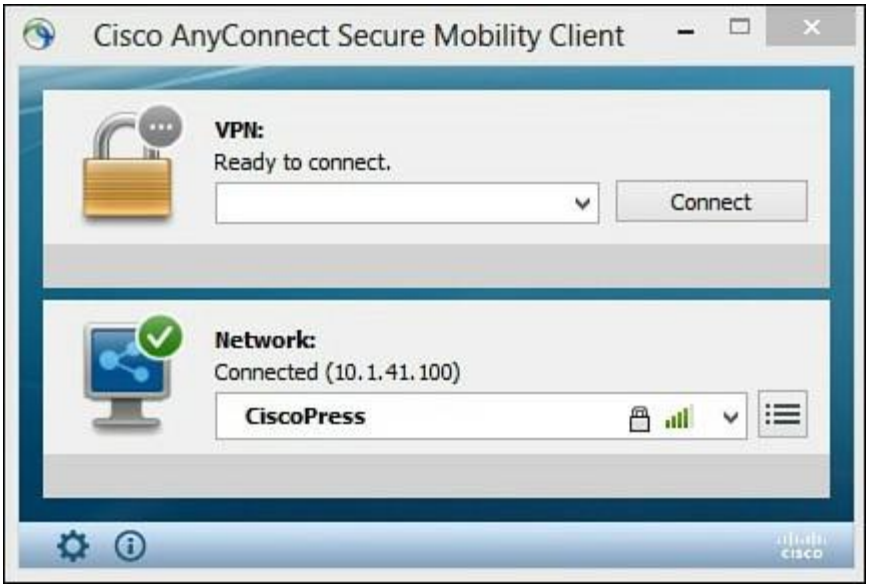


Figure 4-35 End user GUI.

EAP Chaining



Having a machine and user authentication sounds like a great thing. If a machine is able to successfully authenticate to AD, then it must be a company asset and therefore managed by the company (company official software packages and so on). Then knowing the user is an authorized user to be on the company network is of course a desirable state to be in. That is the entire point of 802.1X to begin with.

However, the two states are completely separate. Each EAP session is completely unaware of the other EAP session. There have been enhancements to RADIUS servers over the years to attempt to track the machine and user auth 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 the ability to authorize not only the user, but also the machine, and ensure that the two were authorized together.

Cisco went a step further than what is available with industry standards and created an enhancement to EAP-FAST. Now with EAP-FASTv2, a differentiation was made to have a user PAC and a machine PAC. After a successful machine authentication, ISE will issue a machine PAC to the client. Then when processing a user authentication, ISE will request 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. The IETF is creating a new open standard based on EAP-FASTv2; at the time this book was published, it was to be referred to as EAP-TEAP (tunneled EAP), which should eventually be supported by all major vendors. This topic is revisited in Chapter 17, "Bring Your Own Device."

> **Note**
>
> As of the time this book was published, the only supplicant to support EAP chaining is Cisco AnyConnect NAM 3.1 and newer. The only RADIUS server to support EAP chaining is Cisco ISE 1.1.1 and newer.

https://t.me/learningnets

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 4-5 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Table 4-2	802.1X components	56
Note	Tunneled and native EAP types	60
Table 4-3	Summary of EAP authentication types	62
Paragraph	Knowing that Windows supplicants have a machine and user state and how that machine state works	73
Paragraph	Describing the AnyConnect NAM supplicant	75
Paragraph	Describing 802.1AE (MACSec)	82
Paragraph	EAP chaining	89

Table 4-5 Key Topics Found in Chapter 4

Define Key Terms

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

Extension Authentication Protocol (EAP)

EAP over LAN

native EAP

tunneled EAP

supplicant

authenticator

authentication server

network access device (NAD)

AnyConnect NAM

AnyConnect DART

https://t.me/learningnets

Chapter 5. Non-802.1X Authentications

This chapter covers the following exam topics:

- **Describe Supplicant, Authenticator, Authentication Server**
- **MAC Authentication Bypass**
- **Describe the MAB Process Within an 802.1X Framework**
- **Centralized Web Authentication**

As described in Chapter 4, "EAP over LAN (also Known as 802.1X)," the IEEE standardized a solution for port-based network access control in the early 2000s, known as IEEE 802.1X. It was predicted that no device would be able to access a network without the user identifying himself and being authorized. Well, here we are a decade later and 802.1X is really starting to catch on.

In practice, the belief that all ports would be authenticating with 802.1X did not happen, and to this day does not happen as much as one might think. Ultimately, there can be complications managing a large number of devices that are not managed PCs. Think about it—who manages all the printers in a company, and what management is available? As for the team that manages the printers, do they 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.

The solution could be to list which switch ports have printers 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, digital signage, and so on?

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Describe Supplicant, Authenticator, Authentication Server	6
Describe the MAB Process Within an 802.1X Framework	1, 6
MAC Authentication Bypass	2-3, 6-7, 9
Web Authentication	4-8, 10

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

> **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. True or False? To allow endpoints without configured supplicants to connect to a network where IEEE 802.1X has been enabled, the administrator must disable 802.1X on the endpoints' switch port.

https://t.me/learningnets

a. True

b. False

2. Which of the following is true?

a. With nonauthenticating endpoints, the authenticator takes over the EAP communication instead of the endpoint.

b. With nonauthenticating endpoints, the authenticator can be configured to send the MAC address of the endpoint 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 when 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 can assign all authorization results, except for VLAN assignment.

c. An administrator can 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 can include dACL, VLAN Assignment, SGT, and others.

4. True or False? With centralized web authentication (CWA), ISE sends the username and password to the authenticator.

a. True

b. False

5. Which of following accurately describes local web authentication (LWA)?

a. With LWA, the authenticator redirects the end user's web traffic to a centralized portal hosted on the authentication server, which is then returned to the local device (authenticator).

b. With LWA, the authenticator hosts a local web portal, which is coded to send an HTTP POST to the authentication server containing the credentials of the end user. The authentication server returns an HTTP POST with the Access-Accept or Access-Reject.

c. With LWA, the authenticator receives the credentials from the end user through a locally hosted web portal, and it is the authenticator that sends the credentials to the authentication server through a RADIUS Access-Request.

d. With LWA, the authenticator receives the credentials from the end user through a locally hosted web portal, and the authenticator sends the credentials to the authentication server through a TACACS+ Access-Request.

6. Which of the following lists are non-802.1X authentications?

a. WebAuth, MAB, RA VPN

b. Remote Access, WebAuth, EAP-MSChapV2

c. PAP, LWA, RA VPN

d. WebAuth, EAP-GTC, HTTP POST

7. True or False? Cisco recommends changing the VLAN for a guest user after that visitor has authenticated through Web Authentication to put that guest user into an isolated "guest network."

a. True

b. False

https://t.me/learningnets

8. Which non-802.1X authentication method uses specialized authorization results to connect a user's credentials to a MAB session?
   a. Remote access
   b. Local web authentication with a centralized portal
   c. Centralized web authentication (CWA)
   d. Local web authentication
9. What is one of the main reasons that MAB is used in modern-day networks?
   a. Most endpoints, such as printers and IP phones, do not have supplicants and therefore cannot use 802.1X.
   b. The endpoints can have a supplicant, but the enablement and configuration of that supplicant could be overcomplicated or operationally difficult for the company. Therefore, the company opts to use MAB instead.
   c. The endpoints mostly do have supplicants, but those are not compatible with Cisco networks.
   d. MAB is equally as secure as 802.1X and therefore is chosen often to save the company the operational difficulties of configuring the supplicants on such disparate endpoints.
10. True or False? Web authentication can be used for guest users as well as internal employees.
   a. True
   b. False

Foundation Topics

Devices Without a Supplicant

As described in Chapter 4 and the introduction to this chapter, the IEEE standardized 802.1X for a port-based network access control solution. The original predictions for how quickly 802.1X would catch on and be universally deployed were greatly exaggerated. More than a decade later, we are now seeing 802.1X truly take hold and become the de-facto solution deployed everywhere on both wired and wired networks.

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

In practice, the belief that all ports would be authenticating with 802.1X or no access would be granted 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 devices when thinking about network authentication. However, organizations tend to possess a plethora of devices beyond those. Think about the printers, IP cameras, IP phones, thin-client terminals, building automation, and other "headless" devices that exist in a modern network.

Take Cisco IP phones, for example. These are devices that do have a supplicant, which can be configured individually at the keypad and 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 can 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 capable of configuring each supplicant across large numbers of devices deployed at scale?

One of the original "solutions" to deal with these nonauthenticating devices was to not configure 802.1X on the individual switch ports where the nonauthenticating endpoint would be plugged into the network. Take a moment and think about that for a minute or two—it would enable anyone to simply unplug the nonauthenticating endpoint and plug his laptop into the port. Voila! The device would have full network

https://t.me/learningnets

access without any challenge whatsoever.

What about when a device (such as a printer) needs to be moved? That might require the network team to be involved for the move and enable 802.1X on the old switch port while disabling 802.1X on the new switch port, which poses significant management burden on the IT department. It's just is not a sustainable business model.

Next, 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 bandages needed to be created to deal with all these variations and to deal with them in a sustainable way (where possible).

MAC Authentication Bypass



The first bandage to help with non-authenticating policies 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 a MAC Authentication Bypass (MAB).

Figure 5-1 illustrates the process of MAB.

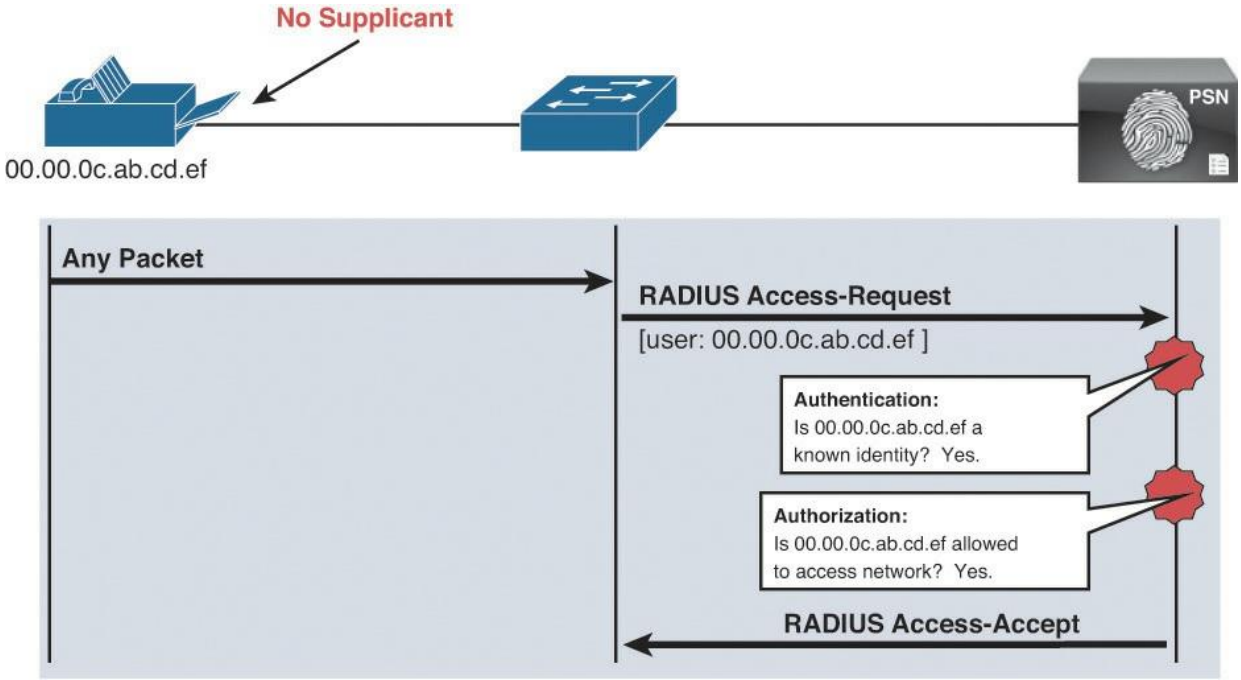


Figure 5-1 MAB.

Examining Figure 5-1, there is a nonauthenticating endpoint (a printer) with a MAC address of 00.00.0c.ab.cd.ef. There is also a switch, which is the authenticator, and an ISE server acting as the authentication server. Let's take a look at the steps that occurred:

Step 1. Because the printer does not have a supplicant, the authenticator crafted a RADIUS Access-

https://t.me/learningnets

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, which determines whether it 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 also important to note that although 802.1X is a standard, MAB is not. MAB is something that each vendor could implement differently if they so choose, just as long as 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 3 timeouts—a period of 90 seconds by default—it is accepted that the endpoint must not have a supplicant. As with most Cisco switch features, timers are adjustable. Figure 5-2 shows the timeouts occurring 3 times before MAB begins.

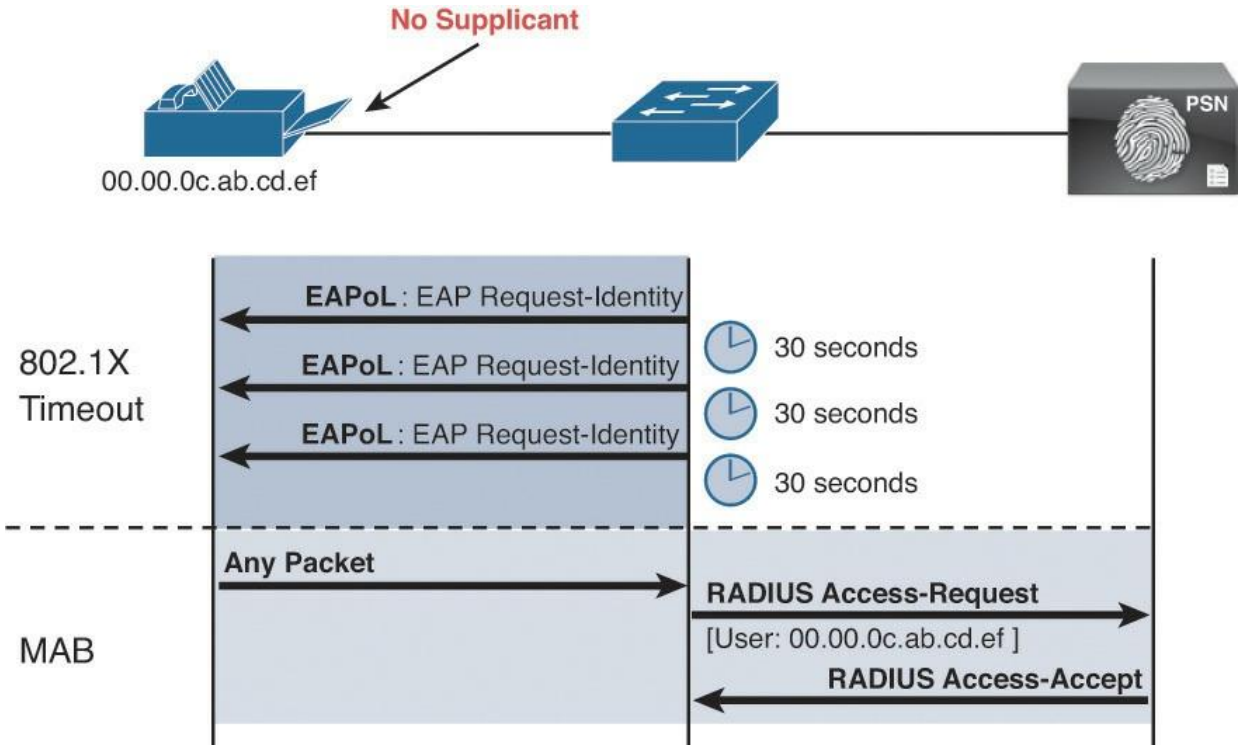


Figure 5-2 EAP Timeout for MAB.

Keep in mind that MAB is inherently not a secure technology. When implementing MAB, you are bypassing the stronger security of 802.1X by allowing specific MAC addresses to gain access without authentication. MAC addresses are easily spoofed, meaning 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'd device for network access, the endpoint should be granted access only to the networks and services that device is required to speak to.

Key Topic

In other words, don't provide full access to devices that have been MAB'd; 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 to the type of

https://t.me/learningnets

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 tag (SGT)
- Smart port macros

Keep in mind that if an endpoint does not have a supplicant, it is not recommended to ever change its VLAN. When changing 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, a supplicant exists on the client to do the VLAN change detection (is gateway reachable, and so on) and trigger the DHCP renewal.

If you still choose to change the VLAN on open networks, then you have only a few choices (none are considered a best practice). You can set the DHCP lease time to something very low, so it will renew the address frequently. You also can use an ActiveX or Java Applet on the portal that will do the VLAN change detection in lieu of a supplicant.

The topics of 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 11, "Authorization Policies."

Web Authentication

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

Enter web authentication, commonly referred to as just WebAuth. An authenticator would be able to send a user to a locally hosted web page—in other words a web page hosted on the local device itself, such as the switch, the wireless controller, or even the firewall or VPN concentrator. It is a simple thing, really; it's just a very basic web page where a user can submit her username and password.

The username and password that are submitted to the web portal are then sent from the authenticator to the authentication server in a standard RADIUS Access-Request packet. So, in a similar fashion to what occurs with MAB, the switch is sending the request for the endpoint because the endpoint is not authenticating to the switch. Figure 5-3 illustrates the WebAuth concept.

https://t.me/learningnets

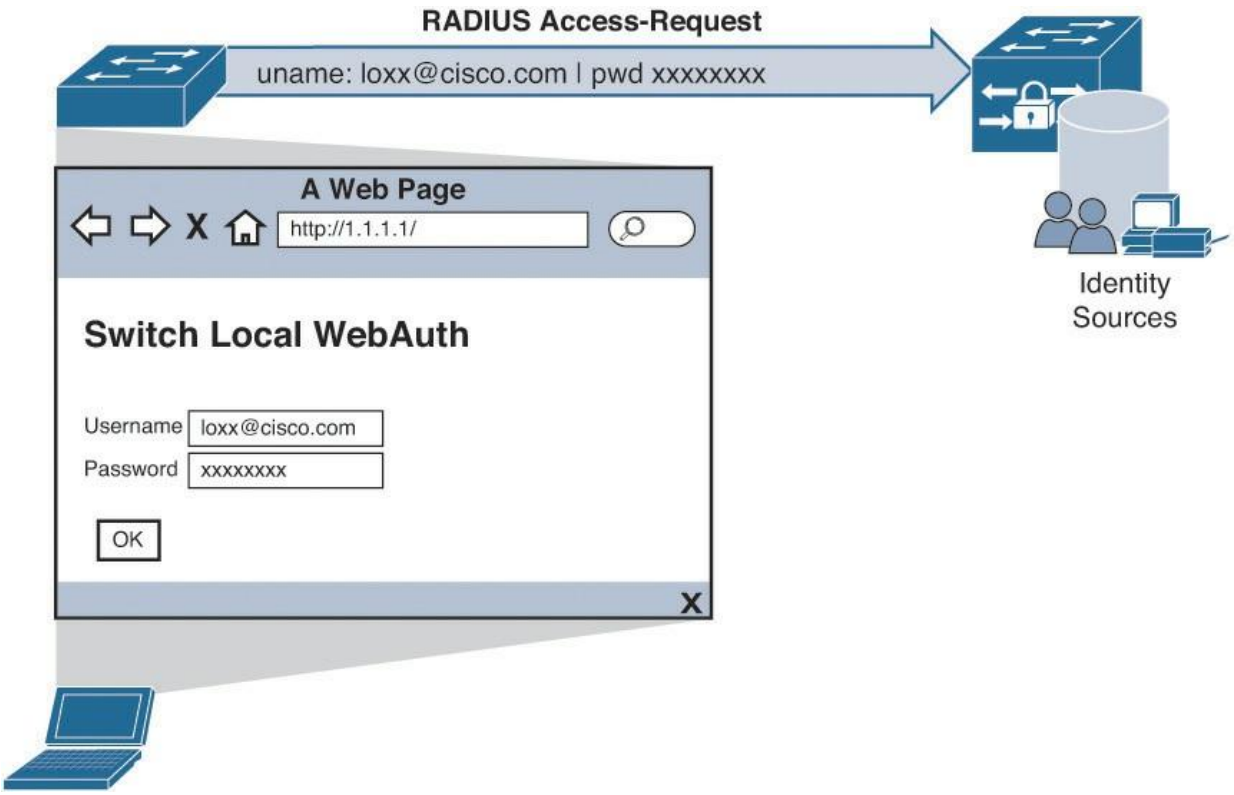


Figure 5-3 Web authentication.

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

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

As with MAB, WebAuth is not a standard. There are multiple ways to perform WebAuth, with benefits and downsides to each one.

Local Web Authentication

Local web authentication (LWA) is the original WebAuth. As described in the preceding paragraphs, the authenticator will redirect web traffic (HTTP and/or HTTPS) to a locally hosted web portal where a user can enter her username and password.



The credentials are submitted through the web portal, and the authenticator (switch, wireless controller, and so on) 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, it is considered LWA.

Storing the web authentication portal on the local authenticator can often limit the customization capabilities of the web pages themselves. On a Cisco switch, the pages are virtually not customizable at all. Some 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.

Additionally, when using LWA with Cisco switches, there is no native support for advanced services such as

https://t.me/learningnets

Acceptable Use Policy acceptance pages, client provisioning, password changing capabilities, self registration, or device registration. For those advanced capabilities, a company needs to consider using centralized web authentication.

Cisco switches as well as a variety of other 802.1X-compliant switches have a configuration option that assigns a special VLAN to endpoints when the authentication timer expires, meaning they don't have a supplicant. This is known as the Guest VLAN. It is an option that was available before the more powerful policy servers—such as Cisco ISE—existed. Many production deployments of 802.1X today still use this Guest VLAN to provide wired guests access to the Internet. However, it is important to note that once the switch makes a local decision, such as assigning the Guest VLAN, LWA is no longer an option.

In addition to LWA and Guest VLAN being mutually exclusive, there are some other interesting bits of information about LWA. 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 other need to quarantine the endpoint.

Local Web Authentication with a Centralized Portal

Many modern authenticators provide an option to redirect the LWA to a centralized web portal. Utilizing a centralized portal enables the organization to customize the portal with corporate branding and provide the right look and feel for their organization.

With Cisco switches, the locally stored web pages can contain the redirection string that sends the users' 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 i-frame.

Figure 5-4 illustrates the process of the user entering his credentials into the centrally located portal (in this case it's hosted on ISE), and Figure 5-5 shows the HTML POST occurring from the browser to the network device (authenticator).

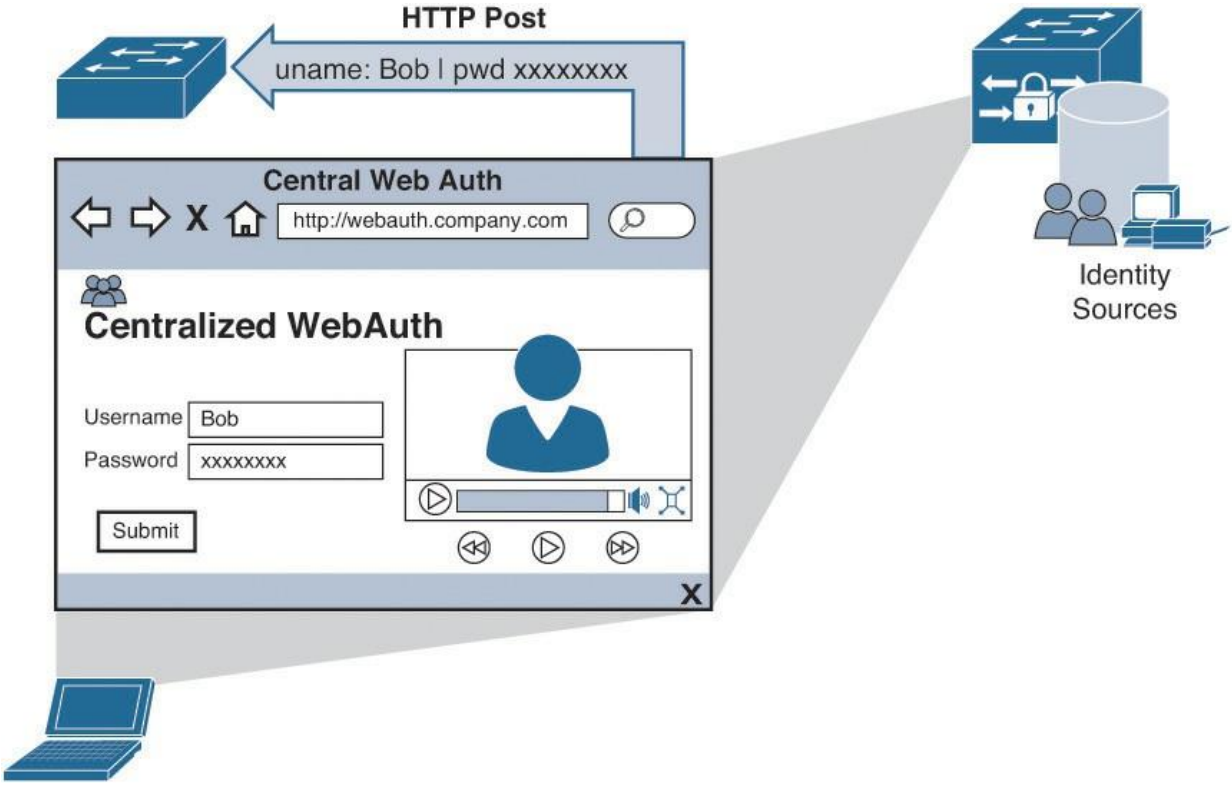


Figure 5-4 Centrally hosted portal for LWA.

https://t.me/learningnets

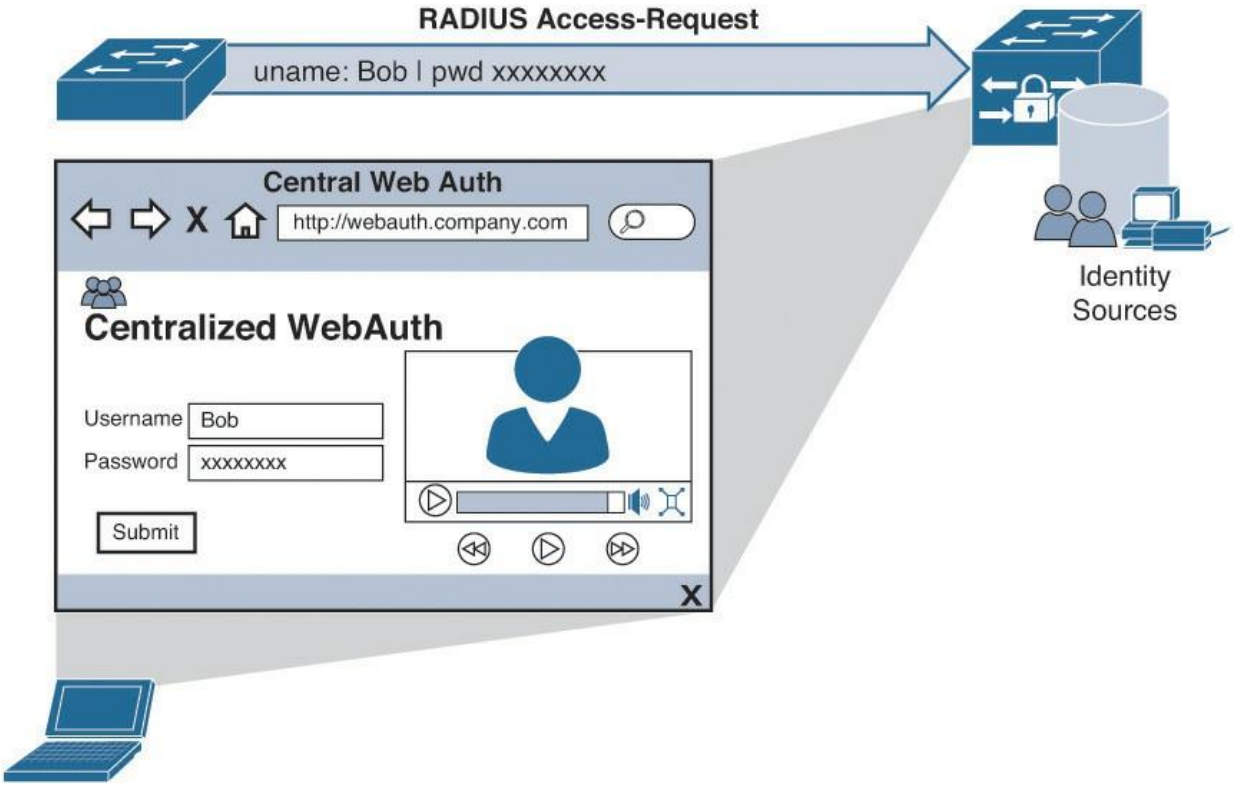


Figure 5-5 HTTP POST occurs from browser to authenticator.

The POST method and the i-frame method both have pros and cons. The i-frame method does not work with newer browsers such as Internet Explorer 9 and newer. The POST method does not work with some mobile browsers, and it is very unsecure 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.

The Cisco Wireless LAN Controller (WLC) version 7.0 required the use of either POST or i-frame method, to support LWA with a centralized portal on ISE. Traditional LWA would always still be available.

Even though the portal is centralized, the same restrictions associated with traditional LWA are still in effect. CoA will not function, nor will client provisioning and other advanced functionality.

Centralized Web Authentication

Centralized web authentication (CWA) is what Cisco ISE uses throughout the Secure Access solution. Although Cisco ISE is still capable of supporting LWA methods, those methods are typically reserved for non-Cisco network devices.

Just like LWA, CWA is only for interactive users that have a web browser, where the user will manually enter the username and password, and just as before, the WebAuth and Guest VLAN functions remain mutually exclusive.

CoA works fully with CWA, which leads to the support for all the authorization results, such as ACL and VLAN authorization. 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 normally is 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.

CWA also supports all the advanced services, such as client provisioning, posture assessments, acceptable use policies (AUPs), password changing, self registration, and device registration

https://t.me/learningnets



Now that you've read all the things it 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 5-6 shows the MAB occurring with a redirection to the centralized portal, and Figure 5-7 illustrates that the switch still sees only a MAB request while ISE is maintaining the user authentication.

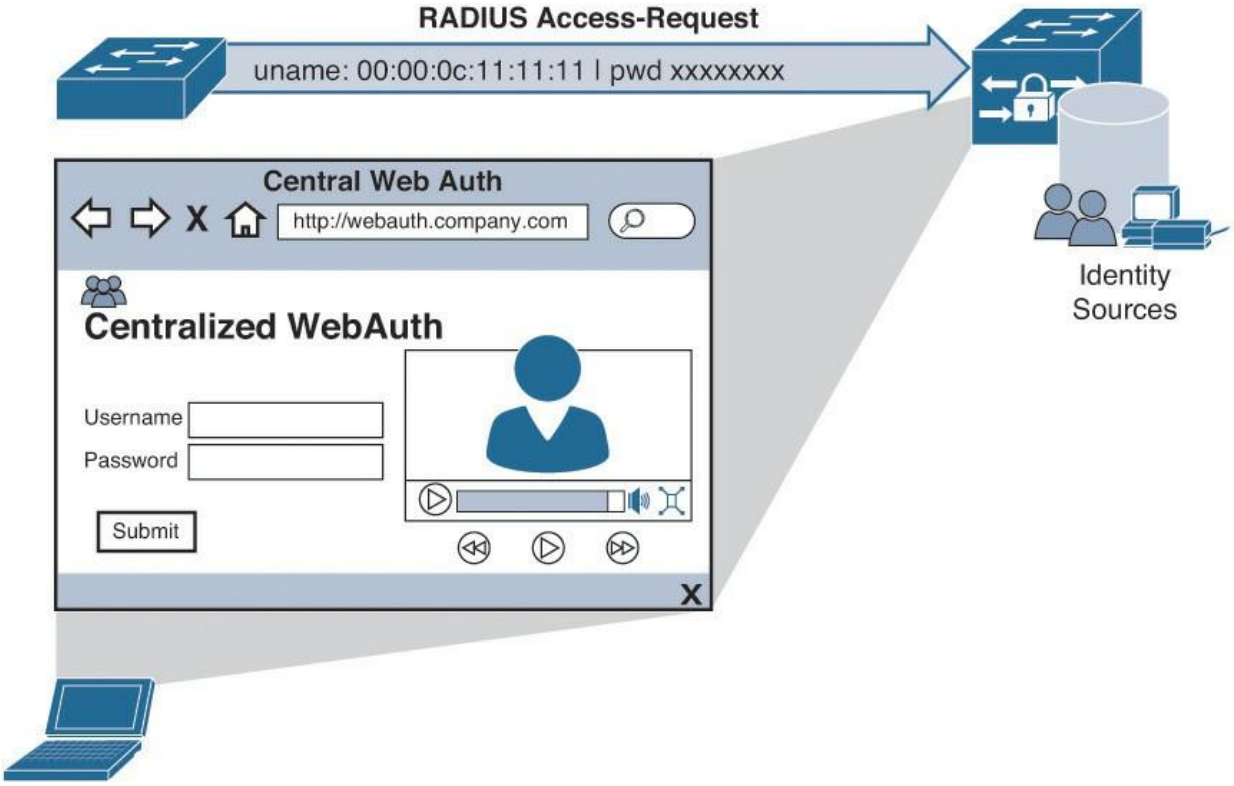


Figure 5-6 URL-redirected MAB.

https://t.me/learningnets

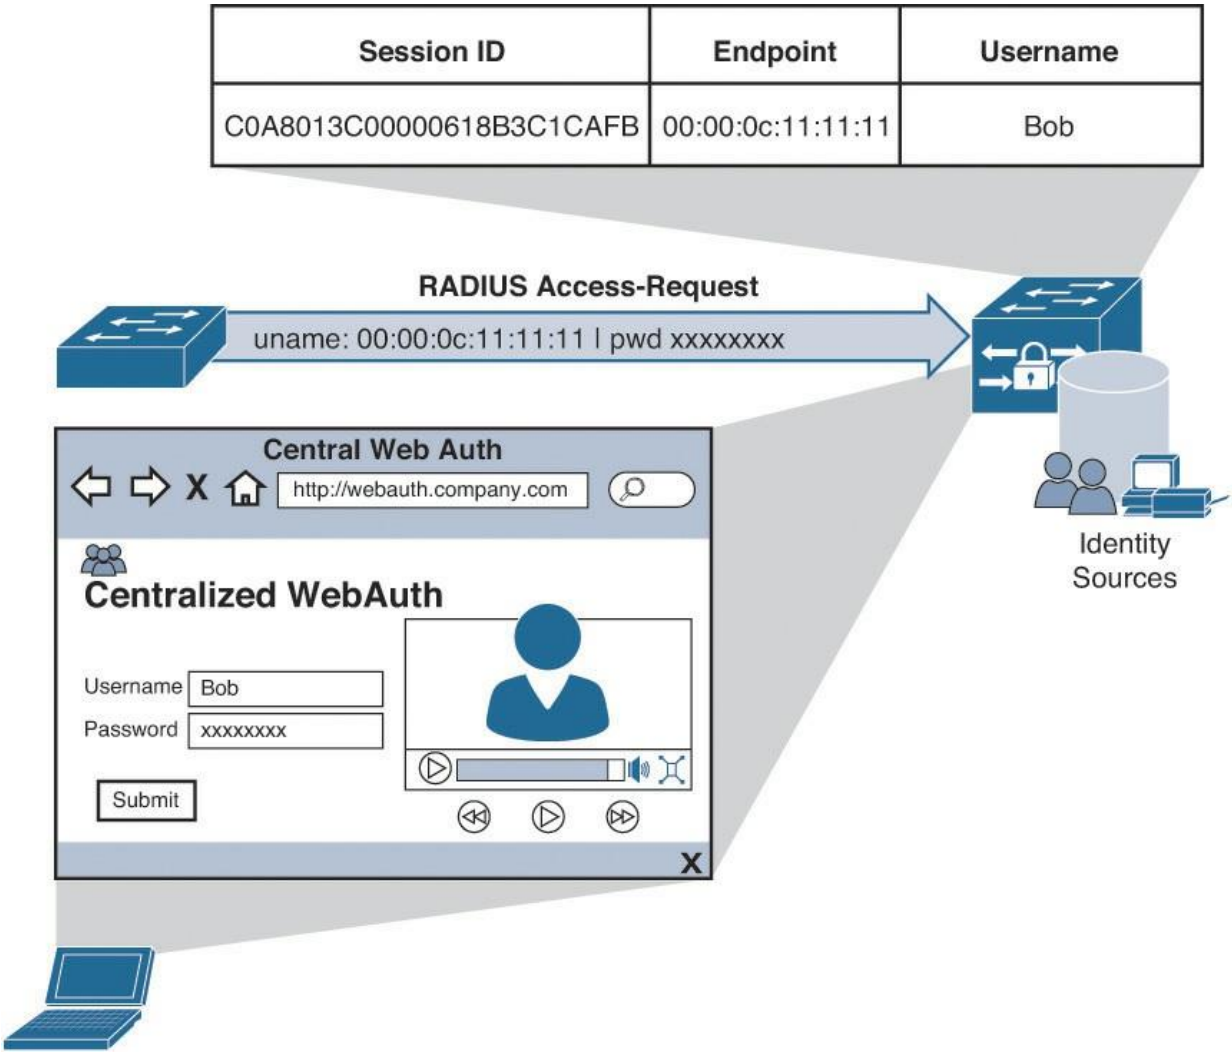


Figure 5-7 The credentials are never sent to authenticator.

The following steps detail what is occurring in Figures 5-6 and 5-7:

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

Step 2. The authenticator performs a 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 itself.

Step 4. The end user enters his credentials into the centralized portal. Unlike the LWA options, the credentials are never sent to the switch; 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, which is processed.

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

CWA is covered in more detail in Chapter 13, "Web Authentication," where the configuration of CWA and the policies involved are examined in depth.

https://t.me/learningnets

Remote Access Connections

So far this chapter has been focused on wired and wireless connections where the endpoint was not using 802.1X. However, there is a third network access use case that is one of the "holy trinity." That trinity is made up of wired, wireless, and remote access. More commonly nowadays, it's referred to as wired, wireless, and VPN.

When connecting an endpoint to a corporate network from another location, it is known as remote access (RA). This used to be handled mainly by dial-up networking and in fact is where RADIUS got its name. RADIUS stands for remote access dial up service. In today's world, using a virtual private network (VPN) to join the endpoint to the company's network through a public Internet is much more common.

With this type of connection, the VPN concentrator (usually a Cisco ASA) is the authenticator and ISE is still 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 can be formed using IPSec or SSL, and the user's credentials are passed inside that secure tunnel.

The VPN head sends the RADIUS Access-Request to ISE, which directs the identity lookup to the appropriate identity store, most often a one-time password (OTP) server. The user's credentials with remote access VPN using the ASA will be sent through either PAP or MSCHAPv2 from the ASA to ISE.

> **Note**
>
> No 802.1X is involved with remote access VPNs.

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 a reference of these key topics and the page numbers on which each is found.



https://t.me/learningnets

Key Topic Element	Description	Page
Paragraph	With MAB, the authenticator will craft a RADIUS Access-Request message and send it to the authentication server. The authenticator uses the endpoint's MAC address as the identity. If that MAC address is to be given access to the network, a RADIUS Access-Accept message is sent back from the authentication server.	98
Paragraph	Do not provide full access to devices that have been MAB'd; instead provide them with an authorization that is more limited.	100
Paragraph	With LWA, credentials are submitted through the web portal, and the authenticator (switch, wireless controller, and so on) sends the RADIUS Access-Request to the authentication server. It is key to remember that any time the switch is sending the credentials for the user, it is considered LWA.	101
Paragraph	With CWA, the authenticator sees only a MAB authentication type. All the rest is handled on the authentication server (ISE).	104

Table 5-2 Key Topics for Chapter 5

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 WebAuth (LWA)

Centralized Web Authentication (CWA)

Do not provide full access to devices that have been MAB'd; instead provide them with an authorization that is more limited.

https://t.me/learningnets

Chapter 6. Introduction to Advanced Concepts

This chapter covers the following exam topics:

- **Change of authorization**
- **Automating MAB**
- **Posture Assessments**
- **Mobile Device Managers (MDMs)**

The features and capabilities available within an authentication server are designed to accommodate a wide variety of deployments—from the simple to the more complex. This can enable network administrators to appropriately stage their deployments, implementing only the basic authentication and authorization features that are required in their networks, and grow into the more advanced features. The advanced features provide for a more comprehensive and often more granular deployment, providing a 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 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."

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

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

> **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. A RADIUS change of authorization enables an authentication server to do which of the following?
   - **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 of all network devices
   - **d.** Take over the world
2. Three possible options for change of authorization actions are which of the following?
   - **a.** IKEv1, IKEv2, SSL
   - **b.** HTTP, FTP, Telnet

https://t.me/learningnets

c. No COA, Port Bounce, Reauth

d. User mode, privileged mode, configuration mode

3. MAC Authentication Bypass 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, of which the first three octets are which of the following?

a. A duplicate of the IP address subnet in hexadecimal format

b. Always the same across all network devices

c. Assigned dynamically upon connection to the network

d. An organizationally unique identifier (OUI) that indicates the device's vendor

e. All F's—that is, FF:FF:FF

5. Which devices often lack an 802.1X supplicant?

a. Printers

b. Laptops

c. Cell phones

d. All of the above

6. Prior to MAB, a switchport with a non-802.1x client would be configured without 802.1x. This presented issues because of which of the following?

a. A broadcast storm would be created as the endpoint device was plugged into 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 switchport to go into error disable.

7. Posture assessment can check for which of the following?

a. File conditions including existence, date, and/or version

b. Registry condition, whether a registry entry is or is not present, on Windows-based endpoints

c. Service condition, whether a service is or is not running, on Windows-based endpoints

d. A and B

e. B and C

f. A, B, and C

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

https://t.me/learningnets

redirection must be destined to _____ 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. A mobile device manager is which of the following?

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 software system or service that provides advanced posture assessment for mobile endpoints

Foundation Topics

Change of Authorization

To have a secure network, it is paramount that no individual or endpoint is allowed access to the network or to network resources until they have been sufficiently authenticated. As we discussed in Chapters 2–5, you can use a variety of identity stores and authentication methods to authenticate a user or device. This authentication establishes who the user and/or endpoint is. For instance, via authentication, you can determine whether the user is part of the network administrators group of users or whether the user is simply a basic user.

Following—and often related to or based on—the authentication, the user or endpoint will be authorized to access the network at some level. Again, as mentioned in earlier chapters, authorization is the process by which you determine to which resources an authenticated user can have access. This level of authorization could allow the user unfettered access to all network resources, access to none of the network resources, or a sliding scale of access between these two extremes. Being able to provide a differentiated level of access to a network administrator versus a basic user, whereby the network administrator is given the greatest level of access and the basic user is given elementary level of access, is critical for network security. It gives users the least level of access that they require to accomplish their duties.

However, it's not always that clear cut. Sometimes a user, an endpoint, or the network can change 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 these changes and provide an updated level of access, or authorization, to this endpoint.



As highlighted in Chapter 2, "Fundamentals of AAA," this server-initiated authorization update process happens via a change of authorization (CoA). As an endpoint's level of authorization changes, the authentication server can elect to do nothing (that is, do not issue a CoA), perform a port bounce (such as shut/no-shut the port), or instruct the network access device to reauthenticate the user.

By using this CoA process, we, as network administrators, have an opportunity via the authentication server to rethink or reexecute the authorization of an endpoint as the network or device changes, enabling us to implement the advanced concepts that are the key topics of this chapter.

https://t.me/learningnets

Automating MAC Authentication Bypass

The first advanced concept we will cover is the automation of MAC Authentication Bypass (MAB). As discussed in Chapter 5, "Non-802.1X Authentications," 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 as both the identity and password fields with a RADIUS Service-Type of Call-Check. This is an indication 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 devices.

Often, the authentication server is configured to trigger this RADIUS Service-Type of Call-Check, authenticating 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 so on—and provided different levels of network access upon a successful authorization, depending on the chosen authentication server.



As discussed at length in Chapter 5, MAB is not a secure technology and it is best to leverage MAB as a secondary "authentication" method instead of as 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, to name just a few. Managing these IP phones, IP cameras, and printers manually within 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 must be added to an internal endpoints database on the authentication server, sometimes via "MAB Devices" list(s). If a user gets a new IP phone because he spilled coffee on his last one or 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 to automate this process for you via device profiling. Profiling is the process by which the authentication server can make educated observations as to what the device is based on its network activity and the contents of these network exchanges. We'll discuss some of the profiling process now and leave a more in-depth discussion of profiling in Chapter 15, "Profiling."

As a device connects to the network, the network access device (NAD) is configured for 802.1X, often with MAB as a backup authentication method—rarely with just MAB (remember: 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 within the RADIUS packets provides the MAC address of the endpoint requesting access to the network, while the Session ID is a unique identifier associated with that RADIUS session (and, therefore, uniquely associated to that MAC address).

As the device continues its network onboarding, it goes through its normal network processes, requesting/receiving an IP address, discovering/advertising its neighbors, and so on. At each step along the way, these various network exchanges can be sent to the authentication server and associated with the MAC address of the device.

To be most effective at profiling the endpoint, a number of profiling options must be enabled on the authentication server to facilitate and gather the information that is being provided by this endpoint. The two key profiling options that will prove most useful as we automate 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

https://t.me/learningnets

the way. From the very first RADIUS packet received on the RADIUS server for the endpoint, observations are being made and recorded. With each RADIUS exchange, the RADIUS Calling-Station-ID and/or Session ID is provided, making it easy for the RADIUS server to know which endpoint the newly found information is about.

This MAC address, alone, can reveal quite a bit of information about an endpoint. First of all, a MAC address is divided into two main portions; the first three octets (00:1D:71:9B:06:40 as shown in Figure 6-1) are the organizationally unique identifier (OUI), and the last three octets (00:1D:71:9B:06:40) is the network interface controller (NIC) specific portion. Where a governing body such as IEEE issues the OUI, the NIC specific portion is unique per device and acts as a serial number that is assigned by the device manufacturer. With the OUI portion of the MAC address, some observations can be made immediately. 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. Otherwise, the authentication server maintains a database of the OUIs associated with certain device vendors as well as the specific device types associated with these vendors, enabling the RADIUS server to dynamically associate that device with a particular network function or device group.

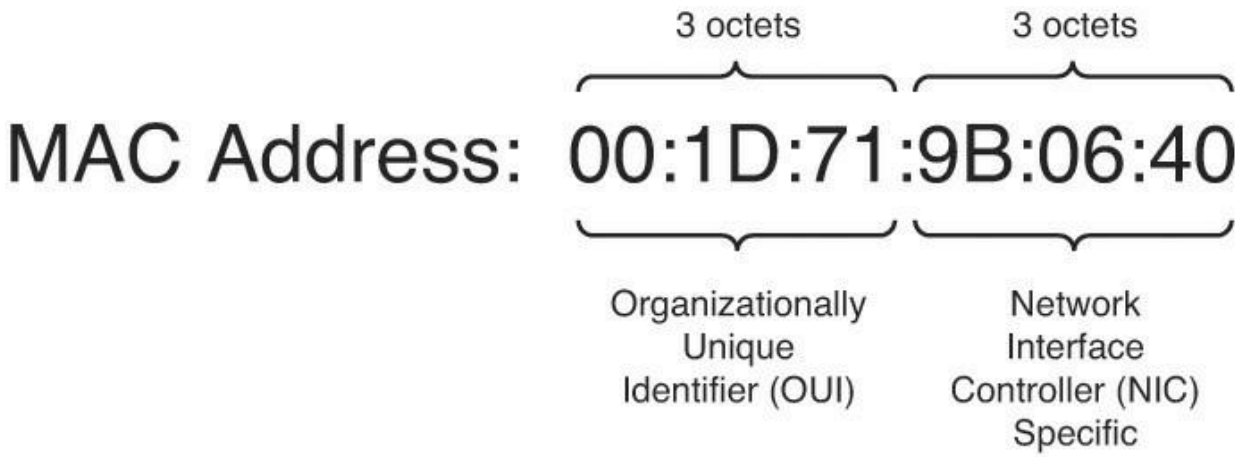


Figure 6-1 MAC address.

For those vendors that produce a large number of device types under a single OUI, such as Cisco (who makes routers, switches, IP phones, IP cameras, and so on), whereby differentiating the device type is not practical, additional information from the RADIUS exchange is often leveraged. Other devices, such as switches and Wireless LAN Controllers (WLCs), within the network can sometimes assist the authentication server in its profiling process. These devices are called device sensors. With these device sensors, the authentication server can gather additional information from the device that is not natively available to the server. "Neighbor" protocols, such as Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP), provide some useful information to the authentication server. These protocols are often confined to the Layer-2 domain; if the authentication server is not on the same Layer-2 domain, another device (such as the device sensors) must tunnel, port mirror, or otherwise relay this information to the authentication server. Within these two protocols, the device capabilities, a description of the device, the hostname, the platform type, and other useful fields can be learned. 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/or 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 the 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 "MAB Devices" list(s) and assigned the appropriate level of authorization.

DHCP profiling can also be a useful mechanism to develop an endpoints dossier. As most devices join a

https://t.me/learningnets

network, they will request an IP address using DHCP. This DHCP process can easily be replicated to the authentication server by simply adding an additional "IP helper" address to the interface on the switch, pointing to the profiling server. As the authentication server receives the DHCP packets from the endpoint, the server will leverage fields from the DHCP packets (which also, subsequently, contains the MAC address of the endpoint) to determine which type of device the endpoint could be. Some of the fields within DHCP that the authentication server can leverage include device dhcp-client-identifier (or MAC address, which contains the device OUI), the dhcp-requested-address (allowing the authentication server to further correlate MAC-to-IP mapping), and dhcp-parameter-request-list (which can also act as a device-type fingerprint).

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

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

To trigger this new level of authorization, we will leverage the CoA. The device was given basic access to the network for the authentication server to gather the required RADIUS and DHCP fields from the device. Now that the authentication server has a more updated view of the device, the authentication server can now, via the NAD, provide a more open or a stricter level of access to the endpoint as necessary. To invoke this reauthorization process, it sends a CoA to the NAD with the relevant Session ID for the endpoint. The NAD then forces the device to reauthenticate, and subsequently reauthorize, its level of access to the network.

Posture Assessments

Authentication is often focused on the user of an endpoint. As the user brings her device onto the network, we, as network administrators, often look at who is using device and do not put 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 her laptop without IT intervention. This is important for allowing the user to add home Wi-Fi parameters, installing productivity applications, as well as installing device drivers when the end user is working remotely from a home office or from a hotspot.

How do network administrators ensure that the changes to the laptop made by the end user in the sanctity of her home are not going to compromise the security of the corporate network? How can we ensure that the device continues to comply with the company security policy? How do we react to a device that is out of compliance? This is where device posture assessments come in.

A posture assessment, or simply posturing, is a process by which an operating system—or, as in this case—an application (for example, a Cisco NAC Agent) running on an endpoint provides critical information about the software that is actively running on the device. Several conditions can be checked as part of a posture assessment. The availability of these conditions or their exact functions 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.
- **Registry Condition**—Checks the existence or value of a registry key on the endpoint.
- **Application Condition**—Checks whether an application is running on the endpoint.
- **Service Condition**—Checks whether a service is running on the endpoint.

https://t.me/learningnets

- **Dictionary Simple Condition**—Checks an attribute that is associated to a user and value.
- **Windows Update Verification**—Confirms that the endpoint has the appropriate Windows service pack installed.
- **Virus Application Verification**—Confirms that the endpoint has an antivirus solution installed. If specific software is required, this can also be confirmed.
- **Virus Definition Verification**—Confirms that the endpoint has the virus definition file present and that they are newer than a particular date.
- **Spyware Application Verification**—Confirms that the endpoint has an antispyware solution installed.
- **Windows screen saver password verification**—Verifies that the user has a Windows screensaver password defined.

This information about the endpoint is forwarded to the authentication or posture server. The server compares this data to the security policy as configured, giving the device an assessment of Compliant, NonCompliant, or Unknown. This assessment outcome can then be used as part of the authorization for the endpoint, or what this endpoint will have access to. If the endpoint is compliant with the configured security policy, the device can be given unfettered access to the network via the selected authorization policy. If the endpoint is NonCompliant or Unknown, the resulting authorization policy could redirect the endpoint to a remediation portal by using a redirect ACL defined on the switch. If the posture status is Unknown, the posture agent will be installed on the endpoint so a proper posture assessment can be completed.

After the 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 the posturing configuration is provided in Chapter 19, "Posture Assessment."

Mobile Device Managers

A mobile device manager (MDM) is a security software system that enables an administrator to configure and secure a mobile device, regardless of where the device is located. As an endpoint onboards with an MDM, the owner of the endpoint gives permission to the MDM software to access the device and gather any required information. This is accomplished by leveraging an MDM agent on the endpoint—either a separate application or an agent that is built into the OS—and a server. After the required permissions are issued, the MDM software performs a number of "assessments" on the device, identifying the current state of security, software, and networking, amongst other things. 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. Some additional supporting servers reside in the cloud and supplement the operation of the MDM; these include Apple's Push Notification Servers and Google Client Messaging Servers.

A number of MDM vendors exist. Each of these vendors provides Comprehensive Device Provisioning, Detailed User and Device Context, and Increased Device and Application Security. Depending on your user base, cost, preferred deployment model (on premises or in the cloud), and security policy, you might choose one MDM vendor over the other.

Many of the features that are supported almost ubiquitously across all MDM vendors include these:

- Support for most mobile and laptop platforms—This includes Android, Apple iOS, Mac OSX, and some Windows.
- Application provisioning—Pushing required applications to the device.
- Security posturing—Includes checking whether Pin-lock is enabled, whether encryption is enabled, whether the device is jail broken, and so on.
- Modify device capabilities—Includes enabling and disabling Wi-Fi, GPS, Bluetooth, and the camera.
- Remote wipe—Whenever a device is lost or stolen, the employee or administrator can perform a partial or full wipe on the endpoint.
- Active Directory support—Includes the ability to authenticate MDM users based on AD credentials.

https://t.me/learningnets

- VPN provisioning—Deploying VPN configurations so the endpoint can securely connect to the corporate network.
- Device backup—Ensuring the device has been backed up recently.

Many of the fields discussed are not directly available to the authentication server's posture assessment process. Depending on the authentication server, some fields might be explicitly communicated to the authentication server. Otherwise, in absence of these explicit fields, the MDM server can provide a final, summary assessment of Compliant or NonCompliant back to the authentication server.

MDMs, either via an application on the endpoint or via software features embedded natively in the operating system, can see a greater amount of detail as described previously. MDMs also provide a greater level of security remotely, for user and administrators alike. This enables the endpoint to be locked remotely and wiped securely from a portal available within the MDM or by proxy via the authentication server.

To take advantage of the MDM features, the authentication server's authorization policy can query the status of the endpoint as Compliant/Noncompliant as seen by the MDM as well as check the status of several other explicit endpoint statuses from the MDM (for example, Jailbroken or Pin-lock). If a device needs to be redirected to register with the MDM or for noncompliance, the authentication server's authorization policy can perform a redirect, sending the endpoint to an MDM portal page for further mitigation and remediation information.

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 devices' security remain paramount. An MDM acts as just 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 6-2 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	Change of authorization (CoA)	113
Paragraph	MAC authentication bypass	114

Table 6-2 Key Topics for Chapter 6

Define Key Terms

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

change of authorization (CoA)

Cisco NAC agent

mobile device manager (MDM)

https://t.me/learningnets

Part III: Cisco Identity Services Engine

https://t.me/learningnets

Chapter 7. Cisco Identity Services Engine Architecture

This chapter covers the following exam topics:

- **What Is Cisco ISE?**
- **Personas**
- **Physical or virtual appliance**
- **Cisco ISE deployment scenarios**

Up to this point, we've looked generically at several key components of a network security architecture, highlighting the fundamentals of AAA, identity management, EAP/802.1X, and advanced concepts. In this chapter, we take our first look at Cisco's Identity Services Engine specifically. We provide a general description of Cisco ISE, the various personas that ISE will assume, physical and virtual instantiations of ISE, as well as 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 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."

Foundation Topics Section	Questions
What is Cisco Identity Services Engine (ISE)?	1
Cisco ISE personas	2-6
Cisco ISE physical or virtual appliance	7
Cisco ISE deployment scenarios	8-10

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

> **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 which of the following?
   - **a.** A switch that provides authenticated access to the network
   - **b.** A network management platform
   - **c.** A network security and policy platform
   - **d.** A unified computing system that incorporates virtualization of endpoints
2. The four key personas of Cisco ISE are which of the following? (Select four.)
   - **a.** Administration
   - **b.** Authentication Server

https://t.me/learningnets

c. File Download

d. Monitoring and Troubleshooting

e. Policy Services Node

f. Identity Management

g. Inline Posture Node

3. The Cisco ISE Administration Node persona is which of the following?

a. The node where policy configuration changes are made

b. The network management platform for the network

c. The engine where policy decisions are made

d. Responsible for logging and reporting data

4. The Cisco ISE Monitoring and Troubleshooting Node persona is which of the following?

a. The node where policy configuration changes are made

b. The network management platform for the network

c. The engine where policy decisions are made

d. Responsible for logging and reporting data

5. The Cisco ISE Policy Service Node persona is which of the following?

a. The node where policy configuration changes are made

b. The network management platform for the network

c. The engine where policy decisions are made

d. Responsible for logging and reporting data

6. Which of the following is true about the Cisco ISE Inline Posture Node persona?

a. A gatekeeper that enforces access policies and handles CoA requests, specifically for those that cannot process CoA requests

b. Is an ergonomic tool included within Cisco ISE to ensure that network administrators are not slouching on the job

c. Allows users to always bypass authentication and authorization, giving them unfettered access to the network.

d. Sniffs all the packets sent from an endpoint, inline, making sure that the endpoint is not distributing viruses and malware onto the network.

7. A virtual ISE appliance should do which of the following?

a. Be kept as small as possible for speed and agility

b. Be appropriately sized to match the equivalent physical appliance

c. Reserve the appropriate resources to ensure that other virtualized applications do not cannibalize the ISE resources

d. A and B

e. B and C

f. A, B, and C

8. In a single-node/standalone deployment of ISE which of the following is true?

a. Each ISE appliance services a single network access device.

b. Each ISE appliance services only a single ISE persona.

c. All endpoints bypass authentication.

https://t.me/learningnets

d. All core ISE personas reside on a single ISE appliance.

9. In a four-node deployment of Cisco ISE, the ____ and ____ personas are combined on two of the appliances, while the ____ persona is by itself on each of the other two appliances.

   a. PAN, PSN, MNT

   b. PAN, IPN, MNT

   c. PSN, MNT, IPN

   d. PSN, PAN, MNT

   e. PAN, MNT, IPN

   f. PAN, MNT, PSN

10. The maximum number of PSNs supported with ISE 1.2 in a fully distributed deployment model is ____, resulting in a maximum number of supported endpoints of _____.

   a. 5; 5,000

   b. 5; 10,000

   c. 5; 50,000

   d. 40; 5,000

   e. 40; 20,000

   f. 40; 250,000

Foundation Topics

What Is Cisco ISE?

Cisco's Identity Services Engine (ISE) is a security policy management platform and a key component of the Cisco secure access architecture. Its role within the network enables it to be the enforcement point for network security.

Network security is often delegated to singular devices within the network. For instance, you might allow unfettered access for all endpoints within the core of your corporate network and enforce the access policy at the edge firewall. For your wireless users, you might choose to enforce a singular policy for all users allowing every wireless user access to HTTP, HTTPS, SSH, and Telnet and implementing this policy at the access point (autonomous mode) or at the Wireless LAN Controller (lightweight mode). This "one-size-fits-all" approach is not the ideal way to implement network security.

Cisco ISE helps facilitate this paradigm shift of making the network—not a singular device—the enforcement point for network security. This is made possible via RADIUS. Remote Authentication Dial-In User Service (RADIUS) provides authentication, authorization, and accounting (AAA). Cisco ISE acts as a centralized, network security policy platform and RADIUS server, extending this AAA functionality 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 to the NAD with the resulting security policy to be applied to the user or endpoint by using RADIUS Attribute-Value Pairs (AVPs).

The authorization policy sent to the NAD is implemented on a per-user basis. By tightly coupling the authorization of the user with the authentication, a network administrator can 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 also might change whether you are accessing the network using wireless versus wired or whether you are accessing the network from a remote branch location or from the campus network. The granular policy building blocks available in Cisco ISE ensure that the network administrator can implement the best network security policy either on a per-user or per-group basis.

https://t.me/learningnets

Besides the basic functions provided as a RADIUS server, Cisco ISE also provides several advanced functions. Natively integrated within Cisco ISE, profiling 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), as well as other network-level assessments of the endpoint. The result of this profiling operation can also be used as one more criteria in the authorization policy that is applied to the endpoint.

Another advanced function available within Cisco ISE is posturing. Posturing can go slightly deeper than what is provided via profiling by using a slightly different approach. Where 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 will ensure 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 gaining 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 managers (MDMs), as described in Chapter 6, "Introduction to Advanced Concepts." Depending on the chosen MDM, additional industry compliance requirements (including HIPAA and PCI-DSS) can be leveraged by Cisco ISE.

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

Trusted level access can be implemented in a number of ways, either by dynamically onboarding the endpoint into the Cisco ISE database (often 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 future chapters.

Guest-level access is also available natively within ISE, leveraging built-in guest portals. The guest portals available within ISE allow a trusted employee to create temporary guest credentials for visitors, contractors, or other temporary users. These guest credentials then are 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 by using network-level authentication such as PEAP) is used to authenticate the user, pushing the relevant guest user access policy to the NAD (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 the network administrators are able to monitor the level of access granted to users and 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 the play-by-play of a user's network authentication and authorization.

Cisco ISE is a one-stop shop for network security policy. By combining an industry-trusted protocol such as RADIUS 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 his network, ensuring that the security policy deployed to each user and endpoint is as secure as possible.

https://t.me/learningnets

Personas

As we begin to take a detailed look at Cisco ISE, we will realize that the Cisco ISE is designed with scalability in mind. The manner by which Cisco ISE achieves this is by dividing the roles it provides into separate personas. Each persona is a different function within Cisco ISE that is required for proper operation of the platform.



The three key personas of a Cisco ISE deployment are Administration Node, Policy Service Node, and Monitoring and Troubleshooting Node.

Administration Node

The first, and possibly the most important, persona is the Administration Node. This function of Cisco ISE is important because it is how the network administrator configures and administers the network security policy.

When a network administrator needs to make a change to the security policy—whether it is the authentication, authorization, profiling, posturing, or other policy—on Cisco ISE, the network administrator must access the administrative graphical user interface (GUI). This administrative portal is a Flash-based webpage. Through this Policy Administration Node (PAN, or simply Admin node), which is the central control center for Cisco ISE, all policies are configured and then pushed to other ISE nodes or personas. This Admin node also includes the licensing function for the platform, allowing the network administrator to enable advanced features or additional functionalities within Cisco ISE by installing purchased licenses.

Policy Service Node

The next persona within an ISE deployment is the Policy Service Node (PSN). This PSN persona is the workhorse of the ISE deployment. After making changes on the PAN, all policies are pushed to the PSNs. As a network administrator, each NAD on the network will point to one or more of the ISE PSNs. 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," amongst others.

When a user authenticates to the network, the NAD forwards the RADIUS Access-Request and subsequent packets to the configured PSN(s). The PSN, having the complete security policy as pushed from the PAN, can authenticate the user and provide the requisite authorization policy.

If the user endpoint is a candidate of any profiling or posturing that is done by ISE, the PSN will also be responsible for deploying a CoA to the endpoint. The NAD also must have the dynamic authorization configuration present, allowing 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.

Monitoring and Troubleshooting Node

The third persona within an ISE deployment is the Monitoring and Troubleshooting (MnT) Node. The function of this node is exactly as it sounds—to provide monitoring and troubleshooting functions for the Cisco ISE deployment. As an endpoint authenticates to the PSN, events are created to keep track of the authentication and authorization process. These events are forwarded to the MnT Node, which then consolidates and processes these events into a legible format. A network administrator requires reports to be created for whatever purpose, such as managerial slides and presentations, access reports, and so on, and this function is also provided by the MnT node.

The second function of the MnT Node is troubleshooting. With each event that is forwarded to the MnT node, this can be a success or failure of the authentication or authorization process. Due to the detailed event tracking that happens on the Cisco ISE PSNs, the content of the logs present on the MnT are also quite

https://t.me/learningnets

detailed. If a user or endpoint is having any level of issue in accessing the network, a network administrator can leverage the MnT Node to see, firsthand, the cause of the failure and track the user's or endpoint's authentication in a play-by-play manner.

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

Inline Posture Node

The fourth and final persona within an ISE deployment is the Inline Posture Node (IPN). The IPN can be considered a gatekeeper between a NAD and the rest of the network. The IPN can ensure that the endpoint is adhering to the required security policy before it is given access to the network. The IPN completes a posture assessment of the endpoint, checking antivirus, antispyware, operating system patch level, security parameters including passwords, and other critical system parameters. Upon determining compliance with the security policy, the IPN can then allow the endpoint to access the network at the appropriate level of authorization. If the endpoint fails its compliance check or its compliance status is unknown, the IPN provides appropriate remediation to get the endpoint in compliance. The IPN is also responsible for handling CoA requests for those NADs that are unable to process their own.

Keep in mind that the Inline Posture Node is optional in ISE deployments, required only when the NAD is incapable of handling CoA requests or when additional posture checking is desired.

Physical or Virtual Appliance

Now that you are aware of the various ISE personas, you are better equipped to start the deployment of Cisco ISE on your network. However, there are still a few more decisions you will need to make as you plan for your deployment. The first of these decisions is whether you will deploy physical ISE appliances or will virtualize the Cisco ISE on your network. This section discusses both options. Your decision of whether to use a physical and/or virtual appliance ISE deployment will depend on a number of factors. Several of the key factors impacting this decision include deployment size, budget, scalability, and current network topology.

As we start this discussion of physical versus virtual ISE deployments, we should first highlight what options are available for each approach. The physical appliances come in two form factors:

- Cisco Secure Network Server 3415 (SNS-3415)
- Cisco Secure Network Server 3495 (SNS-3495)

The SNS-3415 is intended for the smaller size of the deployment scale (5,000 endpoints per server), whereas the SNS-3495 is intended for larger deployments (20,000 endpoints per server in a fully distributed model). See Table 7-2 for the physical appliance specifications. The virtual appliance specifications are provided with specifications to emulate the equivalent hardware appliance. Please see Table 7-3 for that reference.

https://t.me/learningnets

Cisco ISE Appliance	Hardware Specifications	Diagrams
Cisco SNS- 3415-K9	■ Cisco UCS C220 M3. ■ Single socket Intel E5-2609 2.4GHz CPU 4 total cores, 4 total threads. ■ 16 GB RAM. ■ 1 x 600-GB disk. ■ Embedded Software RAID 0. ■ 4 GE network interfaces. ■ For physical, environmental, and power specifications, see Cisco SNS-3400 Series Server Specifications.	■ Cisco SNS 3415/3495 Front Panel ■ SNS 3415/3495 Rear Panel
Cisco SNS- 3495-K9	■ Cisco UCS C220 M3. ■ Dual socket Intel E5-2609 2.4GHz CPU 8 total cores, 8 total threads. ■ 32 GB RAM. ■ 2 x 600-GB disks. ■ RAID 0+1. ■ 4 GE network interfaces. ■ For physical, environmental, and power specifications, see Cisco SNS-3400 Series Server Specifications.	■ Cisco SNS 3415/3495 Front Panel ■ SNS 3415/3495 Rear Panel

Table 7-2 Cisco ISE Physical Appliance Specifications

https://t.me/learningnets

Requirement Type	Minimum Requirements
CPU	Single Quad-core; 2.0 GHz or faster.
Memory	4–32 GB RAM.
Hard disks	From 200 GB to 2 TB of disk storage (size depends on deployment and tasks). We recommend that your VM host server use hard disks with a minimum speed of 10,000 RPM. The Cisco ISE VM requires a minimum write bandwidth of 50 MB per second. This write bandwidth can be easily achieved if the hosting environment uses 10,000-RPM disks. **Note** When you create the VM for Cisco ISE, use a single virtual disk that meets the storage requirement. If you use more than one disk to meet the disk space requirement, the installer might not recognize all the disk space.
Storage	■ File System—VMFS We recommend that you use VMFS for storage. Other storage protocols are not tested and might result in some file system errors. ■ Internal Storage—SCSI/SAS ■ External Storage—iSCSI/SAN The use of NFS storage is not supported.
Disk controller	SCSI controller.
NIC	1 GB NIC interface required (two or more NICs are recommended). **Note** When creating network connections for any NICs that you configure, we recommend that you select E1000 from the Adapter drop-down list. Cisco ISE, Release 1.2, supports the E1000 and VMXNET3 adapters for all NICs. It does not support any other virtual NIC drivers.

Table 7-3 Cisco ISE Virtual Appliance Specifications

If your current network infrastructure is highly virtualized, you might be considering the virtual appliance approach and leveraging existing Cisco UCS infrastructure. If you choose the virtualized appliance approach for your Cisco ISE deployment, it is best to ensure reservations on the virtual machine (VM) for the entire amount of virtual processors, memory, and hard disk space (refer to Table 7-3). This will guarantee that other VMs are not impacting the proper operation of ISE. Failure to reserve the recommended resources can greatly impact the overall performance and scalability of the Cisco ISE.

Whether you choose to use the 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 upcoming section.

https://t.me/learningnets

ISE Deployment Scenarios



Cisco ISE is not a single physical appliance or virtual machine; it is a highly flexible and scalable security platform. This platform is designed for fault tolerance and to scale to all levels of businesses, from a small mom-and-pop shop to the largest of multinational, global enterprises. As part of this section, we highlight a number of common deployment scenarios, including

- Single Node
- 2 Node
- 4 Node
- Fully Distributed

As we discuss each of these deployment scenarios, we also need to consider how each Cisco ISE instance communicates with each other node. The term node is used to refer to a physical or virtual instance of the Cisco ISE appliance.

Single-Node Deployment

A single-node, or standalone, deployment of Cisco ISE is the most basic deployment of the Cisco ISE security platform. In this deployment, all ISE personas reside on a single appliance, whether physical or virtual. As we discussed, the various personas earlier in this chapter, whereas each persona has a specific role to serve in the Cisco ISE ecosystem, this single node ISE deployment is responsible for performing all functions on a single device (see Figure 7-1).



Figure 7-1 Single-node/standalone deployment persona assignment.

> **Note**
>
> The IPN cannot be included in single-node deployments because the IPN requires a dedicated node and cannot assume or share other personas.

Figure 7-2 shows a screen capture of ISE node configuration (**Menu > Administration > System > Deployment > <ISE-Host>**). If the deployment is standalone (by default), all personas should be selected.

https://t.me/learningnets

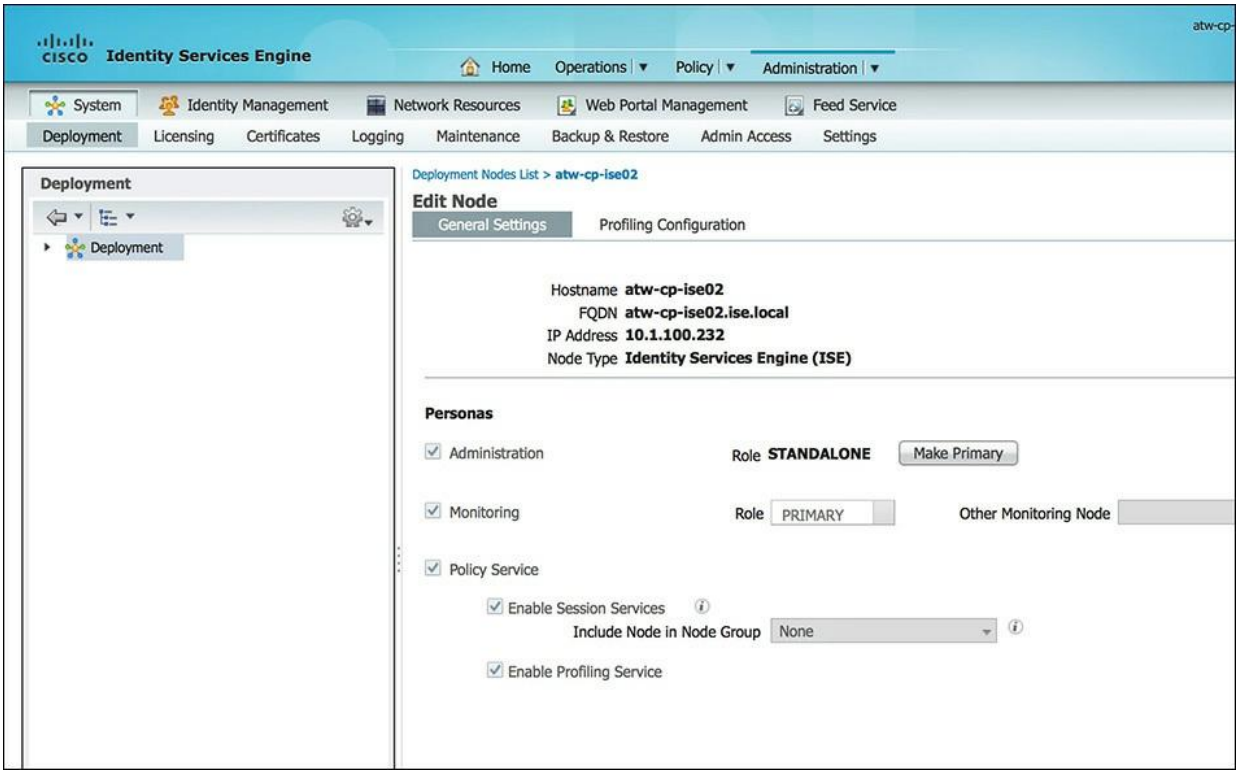


Figure 7-2 Single-node/standalone ISE configuration.

With this single-node, standalone 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 might not exist. The maximum number of supported endpoints in a single-node deployment is 5,000 for the SNS-3415 (or equivalent virtual appliance) and 10,000 for the SNS-3495 (or equivalent virtual appliance).

Because there is no redundancy in this solution and the availability of network resources can be impacted in case of an outage, this deployment is recommended only when testing a solution in a lab environment.

Two-Node Deployment

A two-node deployment of Cisco ISE is the first deployment approach that incorporates redundancy, which is highly recommended in a production network. In this deployment, the various ISE personas are distributed across two appliances. As in the single-node deployment, these appliances can be either physical or virtual appliances. A mixture of physical or virtual appliances also might be considered depending on your network topology and application. Each of the requisite personas—PAN, MNT, and PSN—must be assigned to one or more of the two available appliances. Whenever redundancy is available in this manner, it is sometime referred to as an ISE distributed deployment.

The first appliance, ISE-1, acts as the primary PAN and secondary MNT; the second appliance, ISE-2, acts as the secondary PAN and primary MNT (see Figure 7-3). Both the first and second appliances, ISE-1 and ISE-2, simultaneously act as PSNs. The PSN persona does not recognize a “primary” and “secondary” division of responsibilities, so a PSN is always “active” as soon as the service is enabled on the appliance.



Figure 7-3 Basic two-node deployment persona assignment.

https://t.me/learningnets

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 load of each ISE server is, for the most part, equally distributed because each of the nodes is actively performing a role in the ISE ecosystem. In case of a failure of a single persona, the secondary instance of the persona on the other node can take over that single role. If a failure of a single appliance occurs, every persona's role can be provided by the other node.



Because maintaining a single source of truth is paramount to the proper operation of Cisco ISE, failing over from a primary PAN to the secondary PAN is a manual process. 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. After it becomes primary, synchronization can again resume. This is sometimes referred to as a warm spare.

With regards to the PSN, because a PSN is always "active," each NAD on the network should point to each PSN. Each NAD will contain 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 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 might consider listing ISE-1 as the primary PSN and ISE-2 as the secondary PSN on 50% of the NADs on your network and ISE-2 as the primary PSN and ISE-1 as the secondary PSN on the other 50% of your NADs.

If the PSN persona (and, therefore, RADIUS processing) fails or a single appliance fails, the NAD itself will detect the failure of the RADIUS service on one of the PSNs and direct 100% of future RADIUS queries from that NAD to the remaining PSN. The availability, failover, and recovery detection criterion of the RADIUS service are configured on the NAD. If you want to ensure 100% redundancy of ISE services on your network in a two-node deployment, in case of a single logical or physical failure that would cause a single ISE node to take on 100% of ISE functionality, the maximum number of supported endpoints would be 5,000 for the SNS-3415 and 10,000 for the SNS-3495. If there are greater than the specified number of endpoints, the remaining operational ISE appliance might not be able to single-handedly manage the load.

As shown in Figure 7-2, each ISE host must have the appropriate roles and priorities selected by going to the ISE Node Configuration screen (**Menu > Administration > System > Deployment > <ISE-Host>**).

Four-Node Deployment

A four-node ISE deployment strategy allows for an equal amount of redundancy as the two-node deployment, with a greater level of persona distribution, further separating the function of each ISE appliance. With a four-node deployment, we will leverage the two-node approach for the PAN and MNT personas. The first appliance, ISE-1, acts as the primary PAN and secondary MNT; the second appliance, ISE-2, acts as the secondary PAN and primary MNT (see Figure 7-4).

https://t.me/learningnets

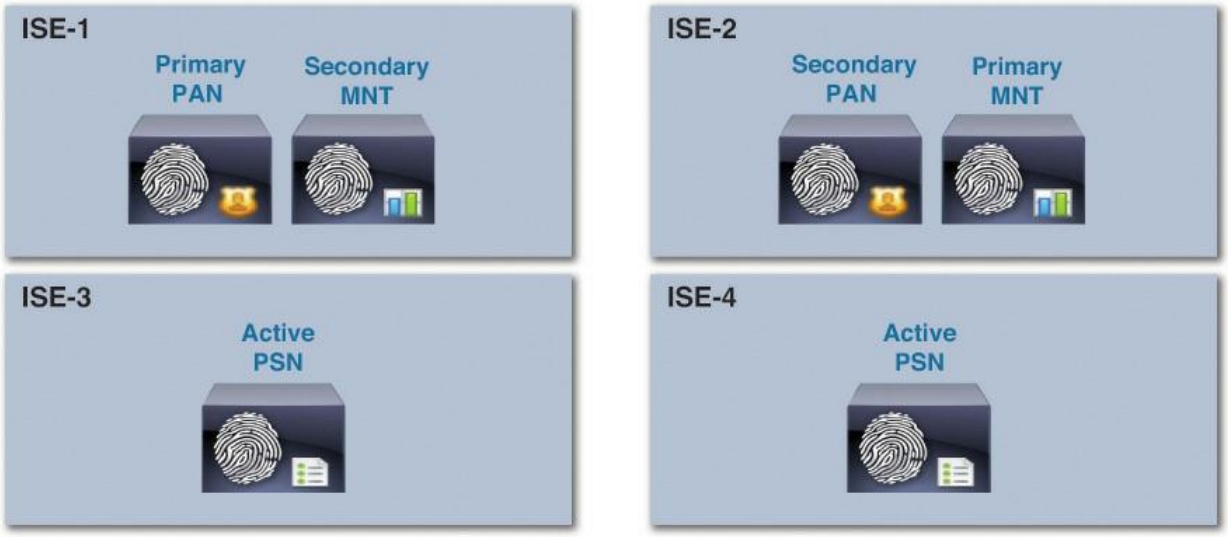


Figure 7-4 Basic four-node deployment persona assignment.

In the two-node deployment, the PSN resides on the each of the ISE nodes, sharing resources with the PAN and MNT personas. However, in the four-node deployment, we separate the administrative functions provided in the PAN and MNT personas from the PSN function. For this reason, ISE-3 and ISE-4 in this four-node deployment are used strictly for the PSN persona.

By separating the PSN persona from the PAN and MNT personas, RADIUS calls to a PSN will continue to function even if the PAN and MNT personas fail. Despite having all PSN functions dedicated to separate appliances, the number of supported endpoints remains the same as in the two-node deployment: 5,000 for the SNS-3415 and 10,000 for the SNS-3495. All other aspects of the four-node deployment will have the same operational features and characteristics as the two-node deployment.

Fully Distributed Deployment

A fully distributed deployment of Cisco ISE separates each persona and places each on a separate appliance. With this deployment model, we place the primary PAN on ISE-1 and the primary MNT on ISE-2, and (if redundancy were required) secondary PANs and MNTs would be deployed on ISE-3 and ISE-4, respectively.

In the four-node deployment, the maximum number of PSNs that can be associated to the combined PAN+MNT is five PSNs, making it a seven-node deployment. If you choose to add more PSNs (up to the maximum of 5), the maximum number of supported endpoints, however, does not increase; these stay at 5,000 and 10,000 endpoints (SNS-3415 and SNS-3495, respectively).

In the fully distributed deployment model, as shown in Figure 7-5, we see much greater scalability. By separating the PAN and MNT personas (using the SNS-3495 for each appliance), a fully distributed deployment can now support up to 250,000 endpoints as a maximum, sustaining approximately 5,000 endpoints per SNS-3415 PSN and approximately 20,000 endpoints per SNS-3495 PSN (an increase from the previous deployment models). The maximum number of PSNs in this deployment model also increases from 5 PSNs to a maximum of 40 PSNs.

https://t.me/learningnets

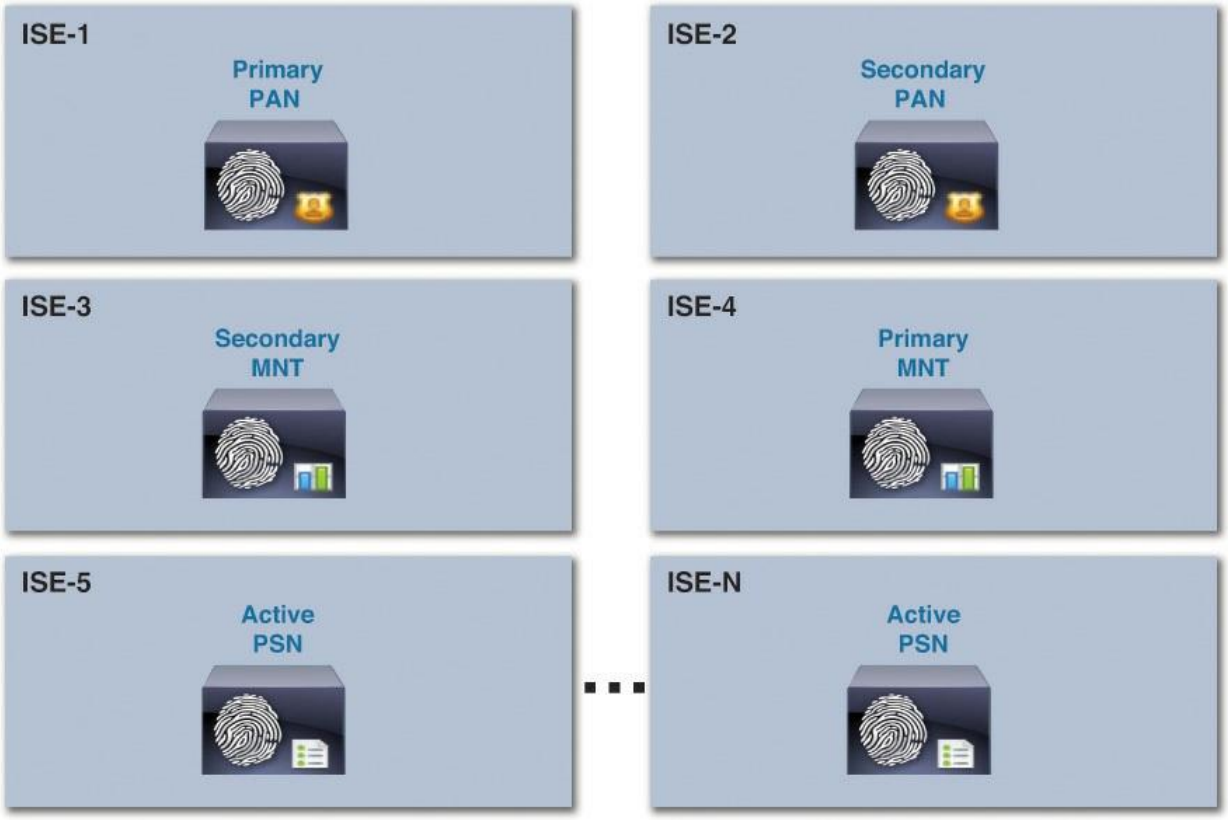


Figure 7-5 Fully distributed deployment persona assignment.

Communication Between Nodes

As mentioned previously, it is best to always deploy ISE in a redundant fashion when used in a production environment. This ensures that a single failure of ISE will not cause a permanent network outage. As you deploy the two-node, four-node, and fully distributed models as described previously, you must ensure that you allow the required communication and allocate the appropriate bandwidth between the individual nodes. If you fail to ensure proper communication between the nodes, the function and performance of the Cisco ISE ecosystem can be greatly impacted.

For all nodes, regardless of persona, it is imperative that you permit TCP/22 (SSH), TCP/80 (HTTP), and TCP/443 (HTTPS). These ports are used for administrative functions on the ISE nodes. Additionally, allowing port TCP/9060 to the PAN nodes will enable you to use the External RESTful Services (ERS) REST API.

There are also a number of ports that should be allowed for External Resources. These ports are used to allow access to LDAP (TCP/UDP/389, TCP/3268), SMB (TCP/445), NTP (UDP/123), DNS (TCP/UDP/53), and Kerberos (TCP/UDP/88 for KDC and TCP/464 for KPASS).

Additional ports are provided in Table 7-4. These enable communications amongst ISE personas and with other network resources and endpoints.

https://t.me/learningnets

Cisco ISE Node	Cisco ISE Service	Ports on Gigabit Ethernet 0	Ports on Gigabit Ethernet 1	Ports on Gigabit Ethernet 2	Ports on Gigabit Ethernet 3
Administration node	Administration	TCP: 22 (Secure Shell [SSH] server) TCP: 80[1] (HTTP) TCP: 443[1] (HTTPS) TCP: 9060 (External RESTful Services [ERS] REST API) Note: Port 80 is redirected to port 443 (not configurable). Note: Ports 80 and 443 support Admin web applications and are enabled by default.	Cisco ISE management is restricted to Gigabit Ethernet 0.	Cisco ISE management is restricted to Gigabit Ethernet 0.	Cisco ISE management is restricted to Gigabit Ethernet 0.
	Replication and Synchronization	TCP: 443 (HTTPS SOAP) TCP: 12001 Global (JGroups - Data synchronization and data replication)	—	—	—
	Monitoring	UDP: 161 (SNMP Query) Note This port is route table dependent.	—	—	—

https://t.me/learningnets

	Logging (Outbound)	UDP: 20514, TCP: 1468 (Syslog) TCP: 6514 (Secure Syslog) Note: Default ports are configurable for external logging. UDP: 162 (SNMP Traps)—			
	External Identity Stores and Resources	TCP: 389, 3268, UDP: 389 (LDAP) TCP: 445 (SMB) TCP: 88, UDP: 88 (KDC) TCP: 464 (KPASS) UDP: 123 (NTP) TCP: 53, UDP: 53 (DNS) (Admin user interface authentication)	—	—	—
	Guest	Guest account expiry email notification: SMTP: TCP/25			
Monitoring node	Administration	TCP: 22 (SSH server) TCP: 80[1] (HTTP) TCP: 443[1] (HTTPS)	—	—	—
	Replication and Synchronization	TCP: 443 (HTTPS SOAP) TCP: 1528 (Secure JDBC - Oracle DB Listener)	TCP: 1528 (Secure JDBC - Oracle DB Listener).	TCP: 1528 (Secure JDBC - Oracle DB Listener).	TCP: 1528 (Secure JDBC - Oracle DB Listener).

https://t.me/learningnets

	TCP: 12001 Global (JGroups - Data synchronization and data replication)			
Monitoring	UDP: 161 (SNMP) Note: This port is route table dependent.			
Logging	UDP: 20514, TCP: 1468 (Syslog) TCP: 6514 (Secure Syslog) Note: Default ports are configurable for external logging. TCP: 25 (SMTP) UDP: 162 (SNMP Traps)			
External Resources	TCP: 389, 3268, UDP: 389 (LDAP) TCP: 445 (SMB) TCP: 88, UDP: 88 (KDC) TCP: 464 (KPASS) UDP: 123 (NTP) TCP: 53, UDP: 53 (DNS) (Admin user interface authentication)	—	—	—

https://t.me/learningnets

Policy Service node	Administration	TCP: 22 (SSH server) TCP: 80[1] (HTTP) TCP: 443[1] (HTTPS)	—	—	—
	Replication and Synchronization	TCP: 443 (HTTPS SOAP) TCP: 12001 Global (JGroups - Data synchro-nization and data replication)	—	—	—
	Clustering (Node Group)	UDP: 45588, 45590 (Local JGroup) TCP: 7802 (Local JGroup failure detection)	—	—	—
	Monitoring	UDP: 161 (SNMP) Note: This port is route table dependent.	—	—	—
	Logging (Outbound)	UDP: 20514, TCP: 1468 (Syslog) TCP: 6514 (Secure Syslog) Note: Default ports are configurable for external logging. UDP: 162 (SNMP Traps)			

https://t.me/learningnets

Session	UDP:1645, 1812 (RADIUS Authentication) UDP:1646, 1813 (RADIUS Accounting) UDP: 1700 (RADIUS change of authorization Send) UDP: 1700, 3799 (RADIUS change of authorization Listen/Relay) Note: UDP port 3799 is not configurable.			
External Identity Stores and Resources	TCP: 389, 3268, (LDAP) TCP: 445 (SMB) TCP: 88 (KDC) TCP: 464 (KPASS) UDP: 123 (NTP) UDP: 53 (DNS) (Admin user interface authentication and endpoint authentication)	—	—	—
Web Portal Services: - Guest/Web Auth - Guest Sponsor portal - My Devices portal - Client Provisioning	HTTPS (Interface must be enabled for service in Cisco ISE.) TCP: 8000-8999 (Guest Portal and Client Provisioning. Default port is TCP: 8443.) TCP: 8000-8999 (Sponsor Portal. Default port is TCP: 8443.) TCP: 8000-8999 (My Devices Portal. Default port is TCP: 8443.) TCP: 8000-8999 (Blacklist Portal. Default port is TCP: 8444.) TCP: 25 (SMTP Notification.)			

https://t.me/learningnets

- BlackListing portal	
Posture	TCP: 80 (HTTP) Discovery - Client side
- Discovery	TCP: 8905 (HTTPS) Discovery - Client side
- Provisioning	Note: By default, TCP: 80 is redirected to TCP: 8443.
- Assessment/ Heartbeat	TCP: 8443, 8905 (HTTPS) Discovery - Policy Service node side
	URL Redirection—Provisioning.
	Active-X and Java Applet Install including IP refresh, Web Agent install, and launch NAC Agent install—Provisioning:
	TCP: 8443 Provisioning: NAC Agent Install
	UDP: 8905 (SWISS) Provisioning: NAC Agent update notification
	TCP: 8905 (HTTPS) Provisioning: NAC Agent and other package/ module updates
	TCP: 8905 (HTTPS) Assessment: Posture Negotiation and Agent Reports
	UDP: 8905 (SWISS) Assessment: PRA/Keep-alive
Bring Your Own Device (BYOD) / Network Service Protocol	URL Redirection—Provisioning.
	Active-X and Java Applet Install (includes the launch of Wizard Install)—Provisioning.
- Redirection	TCP: 8443 Provisioning: Wizard Install from Cisco ISE (Windows and Mac OS)
- Provisioning	TCP: 443 Provisioning: Wizard Install from Google Play (Android)
- SCEP	TCP: 8905 Provisioning: Supplicant Provisioning Process
	TCP: 80 or TCP: 443 SCEP Proxy to CA (Based on SCEP RA URL configuration)
Mobile Device Management (MDM) API Integration	URL Redirection
	API—Vendor-specific
	Agent Install and Device Registration—Vendor-specific

https://t.me/learningnets

	Profiling	UDP: 9996 (NetFlow) Note: This port is configurable. UDP: 67 (DHCP) Note: This port is configurable. UDP: 68 (DHCP SPAN) TCP: 80, 8080 (HTTP) NMAP uses ports 0-65535 [2] (outbound). UDP: 53 (DNS lookup) Note: This port is route table dependent. UDP: 161 (SNMP Query) Note: This port is route table dependent. UDP: 162 (SNMP Trap) Note: This port is configurable.			
Inline Posture node	Administration	TCP: 22 (SSH server) TCP: 8443 (HTTPS) Note: TCP: 8443 is used by the Administration node.	—	—	—

https://t.me/learningnets

Inline Posture	UDP: 1645, 1812 (RADIUS proxy for authentication) UDP: 1646, 1813 (RADIUS proxy for accounting) UDP: 1700, 3799 (RADIUS CoA) Note: UDP port 3799 is not configurable. TCP: 9090 (Redirect)	UDP: 1645, 1812 (RADIUS proxy for authentication) UDP: 1646, 1813 (RADIUS proxy for accounting) RADIUS CoA: Not Applicable TCP: 9090 (Redirect)	—	—
Logging	UDP: 20154 (Syslog) Note: This port is configurable.	UDP: 20154 (Syslog) Note: This port is configurable.	—	—
Note: Inline Posture node High Availability does not apply to any other Cisco ISE node types.				
Inline Posture node (continued)	High Availability	—	—	UDP: 694 (Heartbeat)	UDP: 694 (Heartbeat)

1. Because Inline Posture nodes do not support the Administration persona, they will not have access to this port.
2. NMAP OS Scan uses ports 0-65535 to detect endpoint operating system.

Table 7-4 Cisco ISE Services and Ports

The amount of communication exchanged between the various ISE personas and other network resources can be quite substantial, so it is critical that you allocate sufficient bandwidth across the network. LAN speed (1 Gbps) between nodes is recommended (see Table 7-5) to ensure that this communication is not interrupted. If your deployment is unable to provide this minimum bandwidth, please contact your Cisco Channel Partner or Cisco account representative for guidance.

Connection Between	Minimum Bandwidth
PAN and MNT	1 Gbps
Redundant MNT Pair	1 Gbps
PSN and PAN	1 Gbps
PSN and MNT	1 Gbps
Endpoint and PSN (Posture)	125 bps per endpoint

Table 7-5 Cisco ISE Bandwidth Requirements

https://t.me/learningnets

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 7-6 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	The three key personas of Cisco ISE	129
Paragraph	ISE can be deployed as standalone or as a fault-tolerant and scalable distributed solution.	133
Paragraph	Failing over from a primary PAN to a secondary PAN is a manual process	135

Table 7-6 Key Topics for Chapter 7

Define Key Terms

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

personas

Policy Admin Node (PAN)

Monitoring and Troubleshooting (MNT) Node

Policy Service Node (PSN)

Inline Posture Node

https://t.me/learningnets

Chapter 8. A Guided Tour of the Cisco ISE Graphical User Interface

This chapter covers the following exam topics:

- **Logging into ISE**
- **Organization of the ISE GUI**
- **Types of Policies in ISE**

In Chapter 7, "Cisco Identity Services Engine Architecture," we discussed the roles of each ISE persona and how ISE is deployed within the network. After you have decided the deployment model you plan to use on your network, you then need to begin the configuration of the ISE security policy. However, you must first become familiar with the ISE administration portal.

As you saw briefly in the deployment section of Chapter 7, ISE's administrative portal is a Flash-based, web interface—meaning the administration of ISE is done using a supported web browser that has Adobe Flash installed. The supported web browsers, Adobe Flash, and screen resolution requirements for ISE 1.2 are provided in Table 8-1.

Mozilla Firefox	5.x and later
Microsoft Internet Explorer	8.x and later
Adobe Flash	11.2.0.0 and later
Minimum Required Screen Resolution	1280×800 pixels

Table 8-1 Supported Web Browsers, Adobe Flash, and Screen Resolution Requirements

> **Note**
>
> The Cisco ISE user interface does not support using the Microsoft IE8 browser in IE7 compatibility mode. The Microsoft IE8 is supported in its IE8-only mode.

During this chapter, we will log in to ISE, provide a quick overview of the organization of the ISE GUI, and briefly discuss the various policies that are available within 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 8-2 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."

Foundation Topics Section	Questions
A Guided Tour of Cisco ISE Graphical User Interface	1
Logging In to Cisco ISE	2-3
Organization of the ISE GUI	4-7
Types of Policy in ISE	8-10

Table 8-2 "Do I Know This Already?" Section-to-Question Mapping

https://t.me/learningnets

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 is true of the Cisco ISE GUI?
   - **a.** Requires a separate application to access it
   - **b.** Uses a "standard," Adobe Flash-capable web-browser
   - **c.** Does not exist—ISE is only configurable via command-line interface (CLI)
   - **d.** Requires Cisco Network Assistant
2. To ensure the highest level of security, the ISE administrative GUI uses which of the following?
   - **a.** SSH
   - **b.** SCP
   - **c.** HTTP
   - **d.** HTTPS
3. The initial certificate presented by the ISE administrative GUI is typically which of the following?
   - **a.** Signed by a trusted, public certificate authority
   - **b.** A self-signed certificate automatically generated by ISE
   - **c.** Delivered in a separate envelope from the ISE appliance
   - **d.** Put in a frame and hung over your desk at work
4. Components within the Operations section of ISE allow an administrator to do which of the following?
   - **a.** Actively monitor, report, and troubleshoot active authentication and authorization sessions
   - **b.** Configure how ISE will operate on the network
   - **c.** Create the web portals for client provisioning
   - **d.** Modify the security policy of ISE
5. The Policy tab of the Cisco ISE GUI allows an administrator to configure all of the following EXCEPT which?
   - **a.** Authorization
   - **b.** Client provisioning
   - **c.** Web portals
   - **d.** Security group access
6. You can configure which of the following item(s) under the Administration tab of Cisco ISE?
   - **a.** Policy elements
   - **b.** Certificates
   - **c.** Dictionaries
   - **d.** Network devices
   - **e.** A, B, and C
   - **f.** B, C, and D

https://t.me/learningnets

g. B and D

7. When adding a network access device to Cisco ISE, which of the following details can be configured under the network device? (Select 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

8. An authentication policy within ISE is used to do which of the following?

   a. Determine what the endpoint will be given access to

   b. Identify the endpoint or the user of the endpoint as it connects to the network

   c. Determine the type of security software that is running on the endpoint

   d. Quarantine a user if the endpoint is on the Blacklist

9. Profiling policies within ISE can leverage all of the following protocols to determine the type of endpoint that is accessing the network EXCEPT which? (Select two.)

   a. DHCP

   b. RADIUS (by proxy)

   c. SSH

   d. HTTP(S)

   e. FTP

10. Client provisioning is a process whereby all necessary _____ and _____ are deployed to the endpoint, allowing the endpoint to more easily, maybe even automatically, join the network in the future.

    a. credentials, configurations

    b. regulations, policies

    c. IP addresses, ACLs

    d. protocols, processes

Foundation Topics

Logging In to ISE

Like many modern products, Cisco ISE uses a web-based interface for the GUI. As we discussed in the previous chapter summary, the Cisco ISE administrative portal is an Adobe Flash-based web interface.

Initial Login

To log in to Cisco ISE, you need to open a supported web browser running a supported version of Adobe Flash (refer to Table 8-1). If your web browser does not meet these requirements, your ability to access or configure the various components of Cisco ISE can be drastically affected.

After you have opened your supported, preferred web browser, you must access Cisco ISE. To maintain the highest level of security, the Cisco ISE GUI leverages HTTPS for administrative access. Therefore, the resulting URL to access your Cisco ISE is https://<ISE_FQDN>. If you have not added your ISE into your

https://t.me/learningnets

DNS server, either by hostname or FQDN, you might need to type the ISE's IP address explicitly, in place of ISE_FQDN, to access the administrative GUI.



As you first access Cisco ISE, you might receive a security alert from your browser. The default-generated X.509 certificate 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 into your web browser or load a new X.509 certificate that is "pre-trusted" by your web browser. The process by which this certificate is permanently accepted by the web browser is browser-specific and, therefore, beyond the scope of this chapter. The use of X.509 certificates within ISE are discussed further in Chapter 9, "Initial Configuration of Cisco ISE."

As you first browse to the provided URL, a login screen appears. This login screen will help to authenticate you as a valid administrative user. The username (and password) given here define the level of access you will be given as you access this administrative GUI. The initial administrative user (configured during the initial bootstrapping of Cisco ISE) is given complete access to all components of the administrative portal (the "Super Admin" role; refer to Table 8-3). In Figure 8-1, you will see that our configured initial administrative username is "admin".

https://t.me/learningnets

Admin Group Role	Access Level	Permissions	Restrictions
Helpdesk Admin	Query monitoring and troubleshooting operations.	Run all reports. Run all troubleshooting flows. View the Cisco ISE dashboard and livelogs. 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, livelogs, 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 all reports (run, create, and delete). Run all troubleshooting flows. View the Cisco ISE dashboard and livelogs.	Cannot perform any policy management, identity management, or system-level configuration tasks in Cisco ISE.

https://t.me/learningnets

		Manage alarms (create, update, view, and delete).	
Network Device Admin	Manage Cisco ISE network devices and 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, livelogs, alarms, and reports. Run all troubleshooting flows.	Cannot perform any policy management, 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, livelogs, alarms, and reports. Run all troubleshooting flows.	Cannot perform any identity management or system-level configuration tasks in Cisco ISE.

https://t.me/learningnets

RBAC Admin	All tasks under the Operations menu except for the Endpoint Protection Service, and partial access to some menu items under Administration.	View the authentication details. Enable or disable endpoint protection service. Create, edit, and delete alarms; generate and view reports; and use Cisco ISE to troubleshoot problems in your network. Read permissions on administrator account settings and admin group settings. View permissions on admin access and data access permissions along with the RBAC policy page. View the Cisco ISE dashboard, livelogs, alarms, and reports. Run all troubleshooting flows.	Cannot perform any identity management or system-level configuration tasks in Cisco ISE.

https://t.me/learningnets

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** 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.	—
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.	Cannot perform any policy management or system-level configuration tasks in Cisco ISE.

https://t.me/learningnets

		View the authentication details. Enable or disable endpoint protection service. Create, edit, and delete alarms; generate and view reports; and use Cisco ISE to troubleshoot problems in your network.	
External RESTful Services (ERS) Admin [1]	Full access to all ERS API requests such as GET, POST, DELETE, and PUT.	Create, read, update, and delete ERS API requests.	The role is meant only for ERS authorization supporting internal users, identity groups, endpoints, endpoint groups, and SGT.
External RESTful Services (ERS) Guest	Read-only access to ERS API, only GET.	Can only read ERS API requests.	The role is meant only for ERS authorization supporting internal users, identity groups, endpoints, endpoint groups, and SGT.

Table 8-3 Cisco ISE Admin Groups, Access Levels, Permissions, and Restrictions



Figure 8-1 Initial ISE administrative GUI login.

Additional administrative users can be created after the initial login. Each additional administrative user can

https://t.me/learningnets

be assigned a specific administrative role. As mentioned previously, depending on the username/password, and, therefore, the resulting administrative role, that is provided upon login, you will have access to those relevant functions within the administrative portal. These additional administrative roles are defined in Table 8-3.

To access these administrative roles within Cisco ISE, you must browse to **Administration > System > Admin Access > Administrators > Admin Groups** (see Figure 8-2). To add an administrative user with the given role, go to **Administration > System > Admin Access > Administrators > Admin Users** (see Figure 8-3).

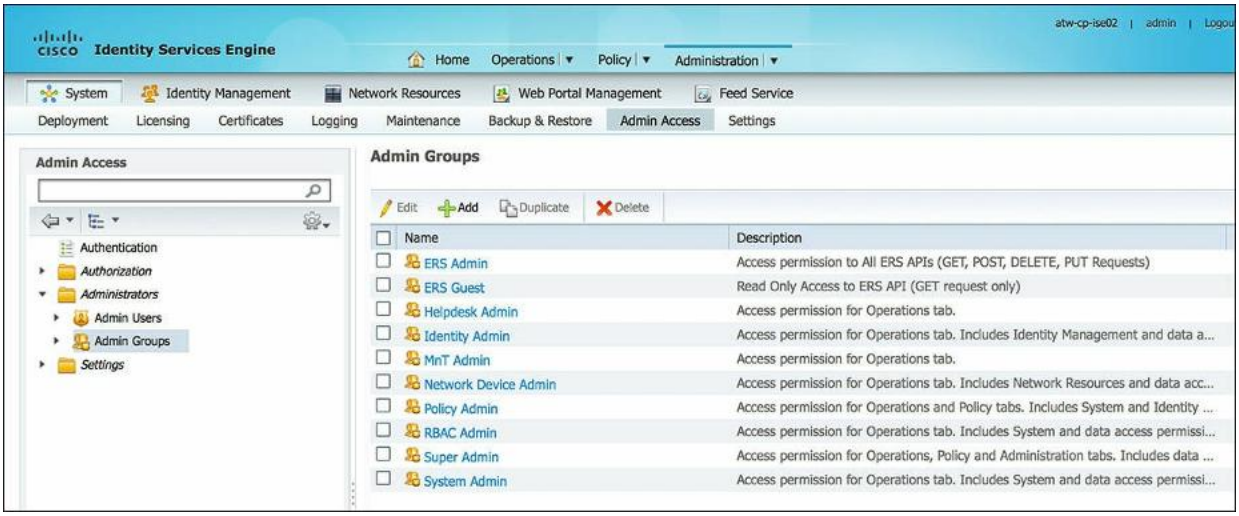


Figure 8-2 Additional administrative roles.

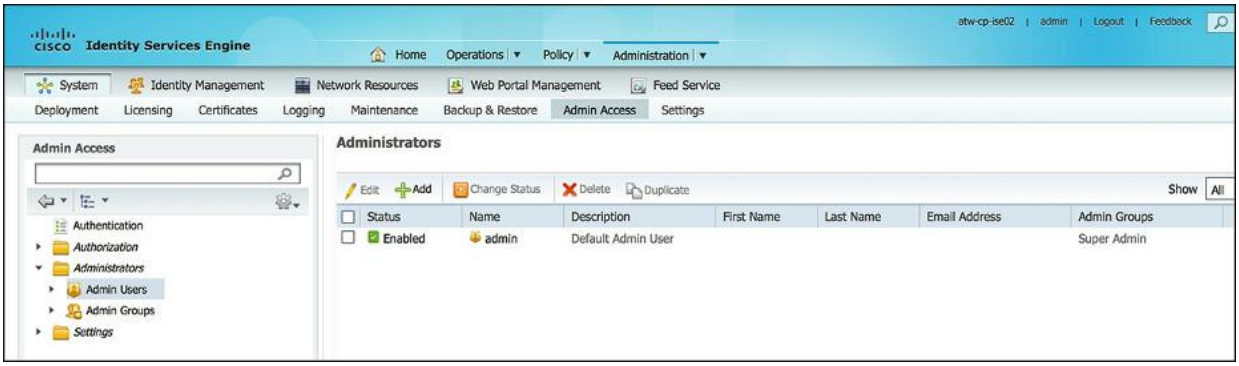


Figure 8-3 Adding administrative users.

Administration Dashboard

The first screen an administrative user sees is the Administration Dashboard (see Figure 8-4 and Table 8-4). This dashboard is analogous to the dashboard of your car because it provides a condensed overview of the status of ISE. As an administrator, this dashboard can be your first stop when checking the status of the entire ISE deployment.

https://t.me/learningnets

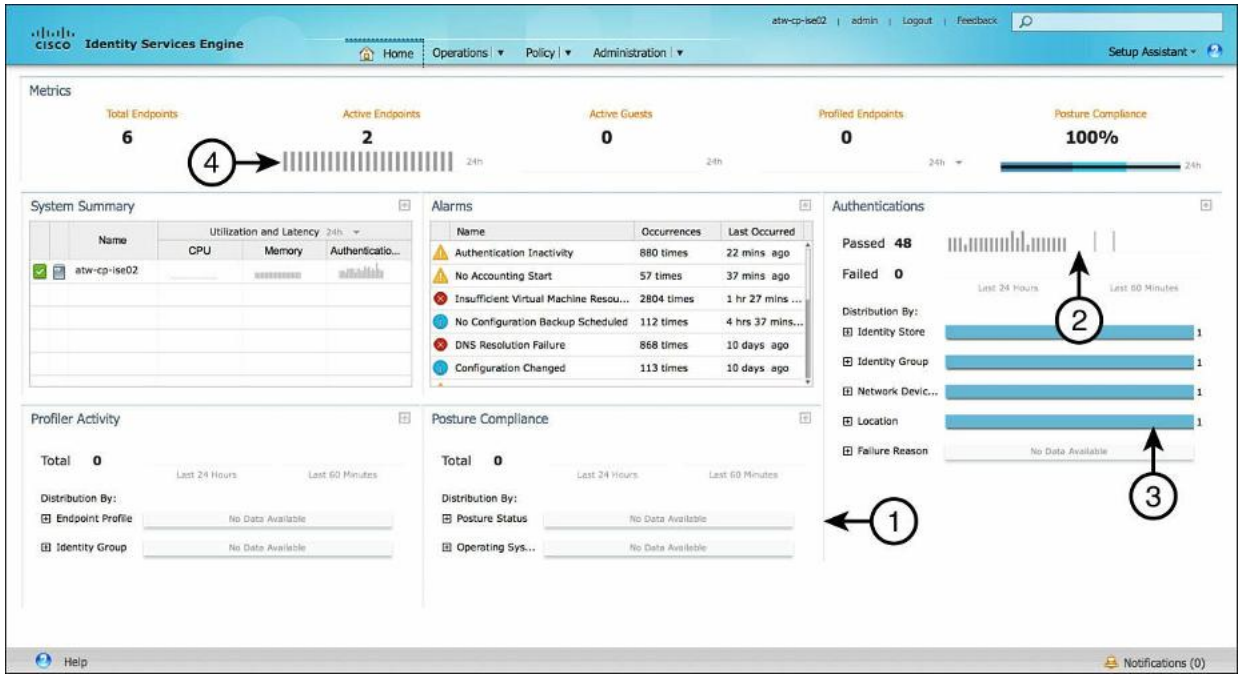


Figure 8-4 Administration dashboard.

1	Dashlets	Dashboard element that displays statistical summaries about the devices and user accessing the network. In some dashlets, colored icons are displayed prior to the device names to convey the system health: Green = Healthy Yellow = Warning Red = Critical Gray = No information
2	Sparklines	Depict trends over time.
3	Stacked bar charts	Display the distribution of parameters using color as the dividing element, so you can see where one parameter ends and another begins. Display is limited to the top 10 distributions. In general, stacked bar charts use color to mark the boundary points between one data measurement and another.
4	Metric meters	Summarize the most important statistics regarding the devices that are accessing the network. Metric meters provide an at-a-glance view of network health and performance. You can click the number displayed above the metrics meter to view more information about the devices.

Table 8-4 Administration Dashboard Components (refer to Figure 8-4)

The administration dashboard provides the following dashlets:

- **Metrics**—Shows a high-level summary of how many endpoints are active, how many active guests exist, the profiled endpoints, as well as posture compliance.
- **System Summary**—Provides a quick overview of all ISE appliances within the deployment. This summary includes the health status, CPU level, memory usage, and authentication latency for each ISE appliance.

https://t.me/learningnets

- **Alarms**—Provides any alarms or anomalous behaviors that have been seen by ISE. A few examples of alarms may be authentication inactivity, NTP sync issues, or insufficient virtual machine resources, just to name a few.
- **Authentications**—Provides a 24-hour and 60-minute summary of the ISE passed and failed authentications, providing a distribution between identity store, identity group, network device, location, and failure reason (if applicable).
- **Profiler Activity**—If profiling is enabled, this dashlet provides a 24-hour and 60-minute overview of which endpoint profile or identity group the endpoints were profiled.
- **Posture Compliance**—This pane provides the posture status of endpoints, highlighting the distribution based on profile status as well as operating system.

Based on the status of the various dashlets, further research, troubleshooting, and reporting might be required.

Administration Home Page

Besides the administration dashboard, other important information is available on the Home page.

Server Information

In the upper-right corner of the administration Home page, you will see the hostname of the ISE appliance that is currently being viewed. By hovering over this hostname, you will see some additional information about the ISE appliance. The information given in this Server Information pop-up is shown in the following list and Figure 8-5:

- **Personas**—The active personas on the current ISE appliance
- **Role**—Whether this ISE is currently in a standalone, primary, or secondary role
- **System Time**—The current time as seen by the current ISE appliance
- **FIPS Mode**—If the current ISE node is running in FIPS compliant mode, it is indicated here
- **Version**—The current version of ISE that is running
- **Patch Information**—If there is a system patch that has been installed, it is indicated here



Figure 8-5 Server information pop-up.

https://t.me/learningnets

Setup Assistant

The Setup Assistant link on the Administration Home page, located in the upper-right corner of the page, can be used to do an initial configuration of the Cisco ISE after the initial CLI bootstrapping of ISE. The Setup Assistant guides you through a series of questions to configure the basic functionality of Cisco ISE. Following the Setup Assistant Wizard, the input you provided is used to configure Cisco ISE directly. Some of the areas that are affected by the Setup Assistant include authentication, authorization, profiling, posture, client provisioning, guest services, and support for personal devices. Figure 8-6 shows you how to run the Setup Assistant.

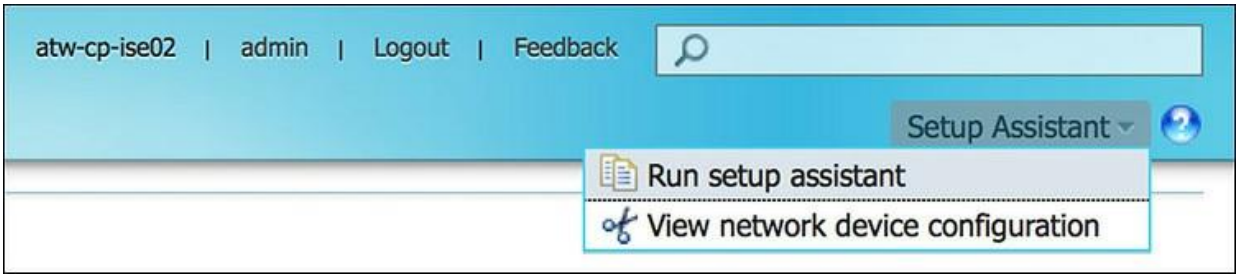


Figure 8-6 Setup Assistant drop-down.

Help

In the bottom-left corner of the Administration Home page, the Help link provides a number of useful resources to help you manage your Cisco ISE.

The first resource provided via the Help link is the Task Navigator. The Task Navigator is a set of wizards that will walk you through a number of common configuration tasks. These walkthroughs are given as a linear set of tasks—whereby the order of the tasks is the order in which the tasks must be configured. The links provided as part of this workflow take you directly to the relevant configuration page within ISE.

The second resource provided via the Help link is the Online Help repository (see Figure 8-7). This link provides the online documentation for Cisco ISE. When you select Online Help, an indexed and searchable guide appears, often in a new tab within your browser (you must allow for pop-ups from the ISE node). This is a local copy of the Cisco ISE User Guide for the version of ISE that is currently running. Also within this Online Help, you can add your favorite links for those tasks that you refer to often.

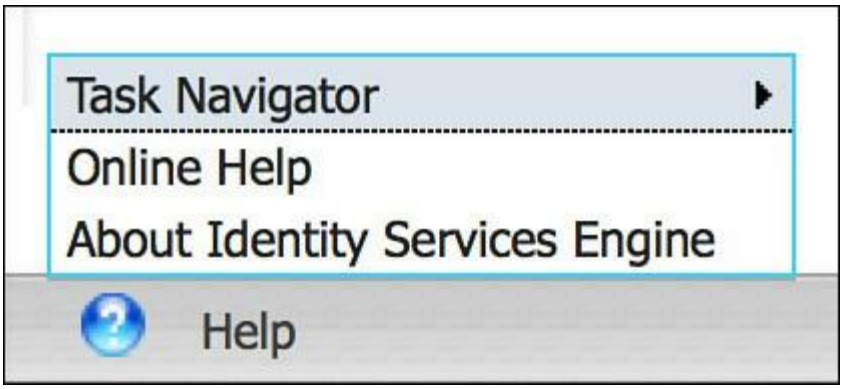


Figure 8-7 Help link.

The third resource provided via the Help link is About Cisco Identity Services Engine (see Figure 8-8). This resource provides the ISE and ADE software versions that are running on the current node. The product ID, version ID, and serial number are also provided via the About link; these can prove quite useful when requesting customer support from Cisco TAC, RMAs (in case of a defective ISE), and to place additional orders for similarly configured equipment. This information is also important for licensing purposes.

https://t.me/learningnets

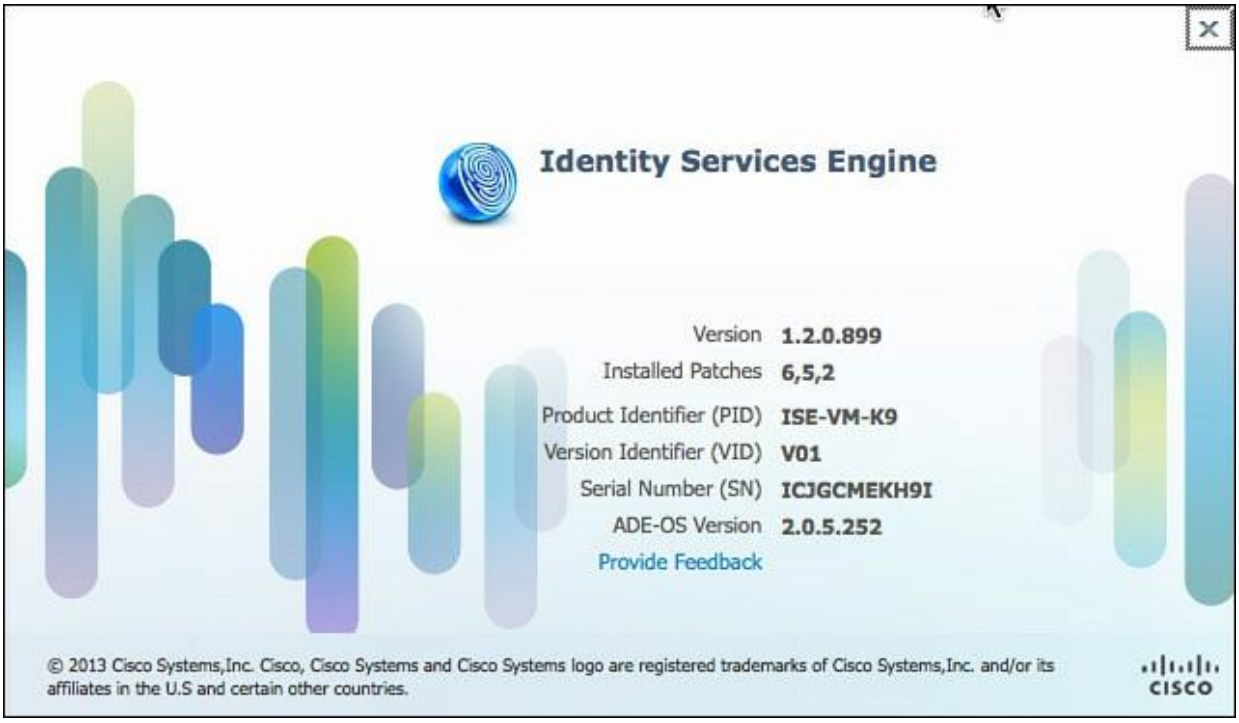


Figure 8-8 About Cisco Identity Services Engine.

Organization of the ISE GUI

Besides the Administrative Home screen, the Cisco ISE GUI is divided into three functional components—Operations, Policy, and Administration.

Operations are those components of ISE that enable the administrator to actively monitor, report, and troubleshoot ongoing authentication and authorization sessions. It is also a place where the administrator can monitor, report, and troubleshoot those network devices and policies that are already configured on ISE.

The Policy functions are those components of ISE that allow the administrator to configure the security policy. These policy functions include authentication, authorization, profiling, posture, client provisioning, and security group access policy—as well as the supplementary building blocks that are used within these policies. As a network device authenticates and authorizes to ISE, ISE processes the credentials provided by the NAD through this policy, providing the resulting authorization security policy back to the NAD.

The third major component of the Cisco ISE GUI is Administration. Administration focuses on the configuration of the ISE component itself—what, who, and how users and devices can access ISE. This configuration section of ISE enables the administrator to define how the ISE deployment behaves, which external identity resources are going to be used, which devices are allowed to use the ISE security policy, which services ISE will provide to the user base, and how often ISE will update its device databases.

We'll now dive a little deeper into each of these major functional components, describing how each, smaller part contributes to the complete ISE ecosystem.

Operations

As you look at the Operations tab, you will see four subcomponents—Authentications, Reports, Endpoint Protection Service, and Troubleshoot.

https://t.me/learningnets

Authentications

The authentication subcomponent of Operations takes on a very similar feel as the administration dashboard discussed earlier in this chapter. This authentication screen can be configured to update in a periodic fashion, providing a live view of the authentications that are occurring and the resulting security policies that are actively being issued to the endpoints. By monitoring this screen, an administrator can quickly discern whether users are successfully authenticating/authorizing or whether there is currently a problem with authentications or authorizations.

Within the Live Authentications Log, we can see those users (Identity), devices (Endpoints), or NADs (Network Device) as they attempt to authenticate to ISE. If the policy is successful, the resulting security policy (Authorization Profile) is sent to the NAD. Each authentication event is logged on this Live Authentications Log, as shown in Figure 8-9.

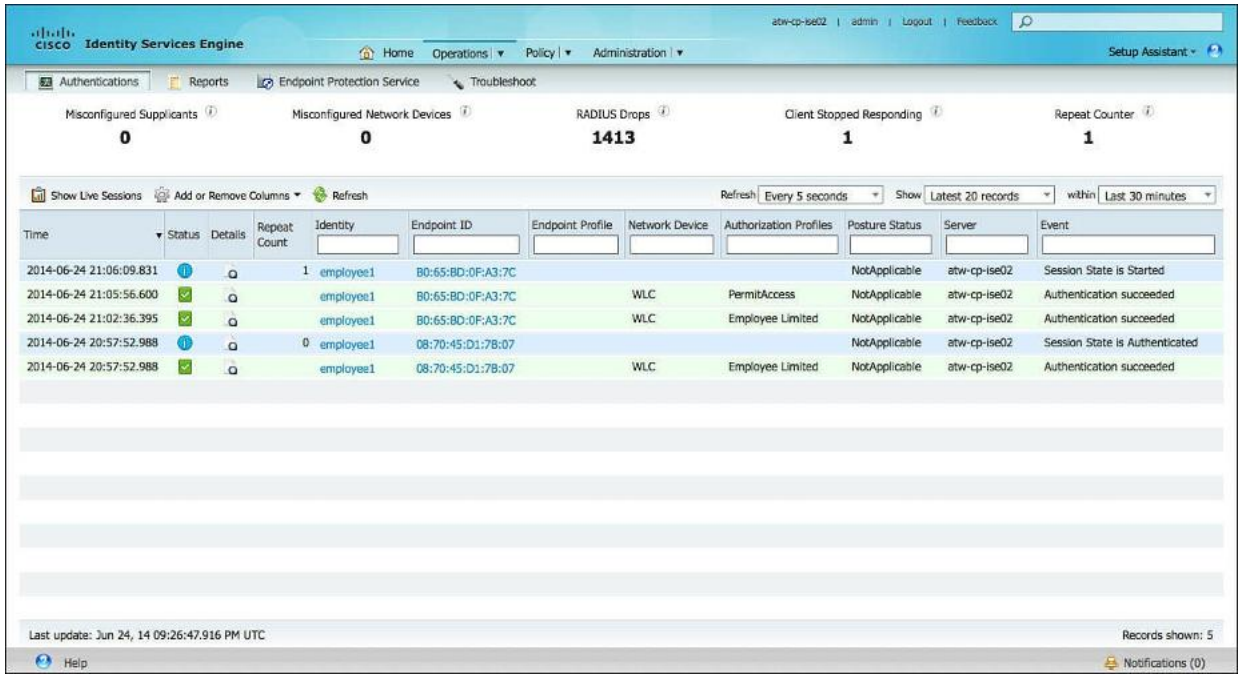
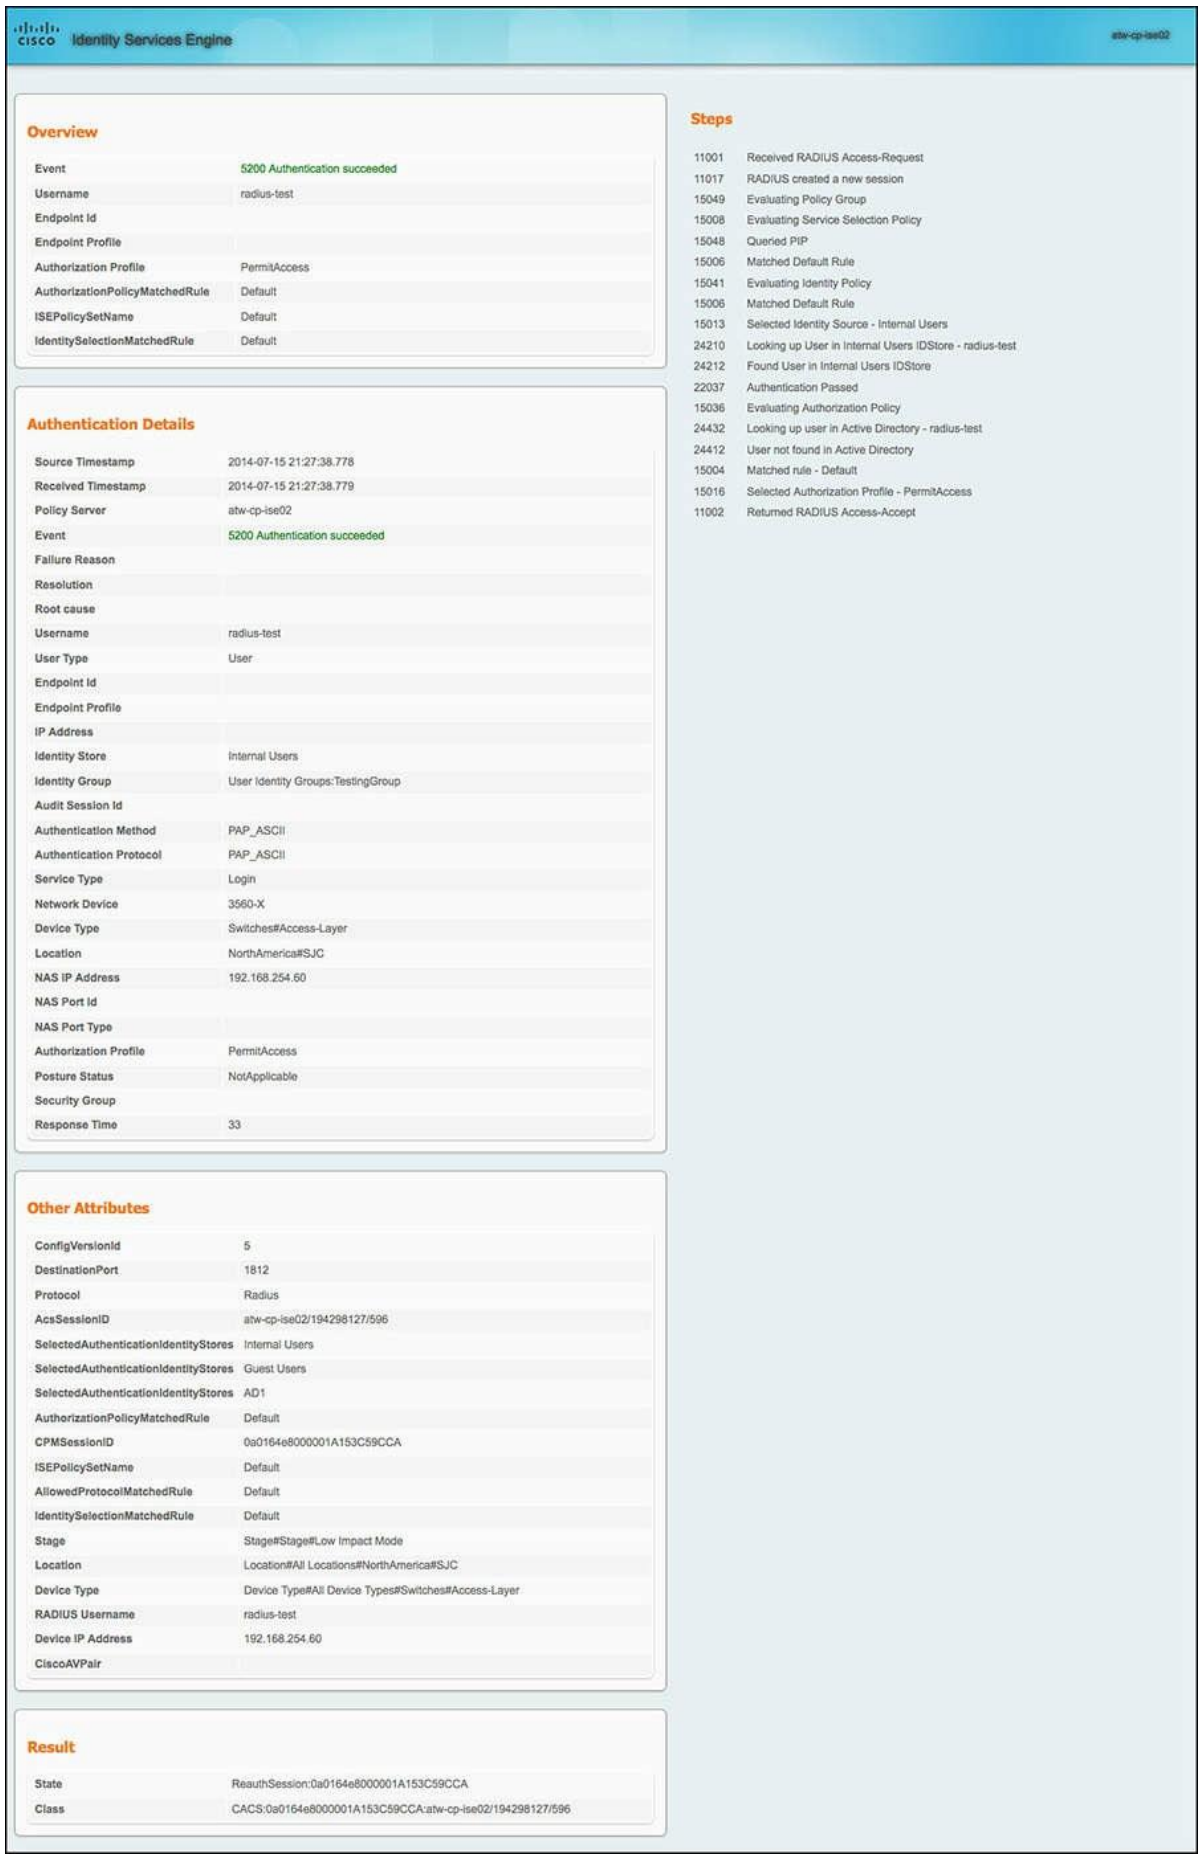


Figure 8-9 Live Authentication Log.

The Live Authentications Log provides only a summary of the authentication but not all of the details. However, the details of the authentication session are sometimes as important—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 will help an administrator determine the nature of the problem. To see the details of the authentication session, click the magnifying glass (in the Details column) for the relevant authentication event.

As you click the Live Log Details link (see Figure 8-10), you will see a large amount of information about the Authentication session—the details of the RADIUS exchange that resulted in the authentication event, the endpoint, the NAD, and the identity store—to name just a few. By reviewing portions of the Live Log Details and knowing the expected policy outcome, you should be able to determine the reason a particular authentication, authorization, posturing, or profiling policy was chosen, correctly or incorrectly, as the result of the authentication session.

https://t.me/learningnets



https://t.me/learningnets

Figure 8-10 Live Authentication Log details.

The alternative view of the Live logs—Live Sessions—can be seen by clicking the Show Live Sessions icon shown in Figure 8-11. Whereas the Live Authentications Log showed real-time authentications to ISE, the Live Sessions Log shows the status of those active sessions with ISE. This can be a session that has just started or a session that is ongoing.

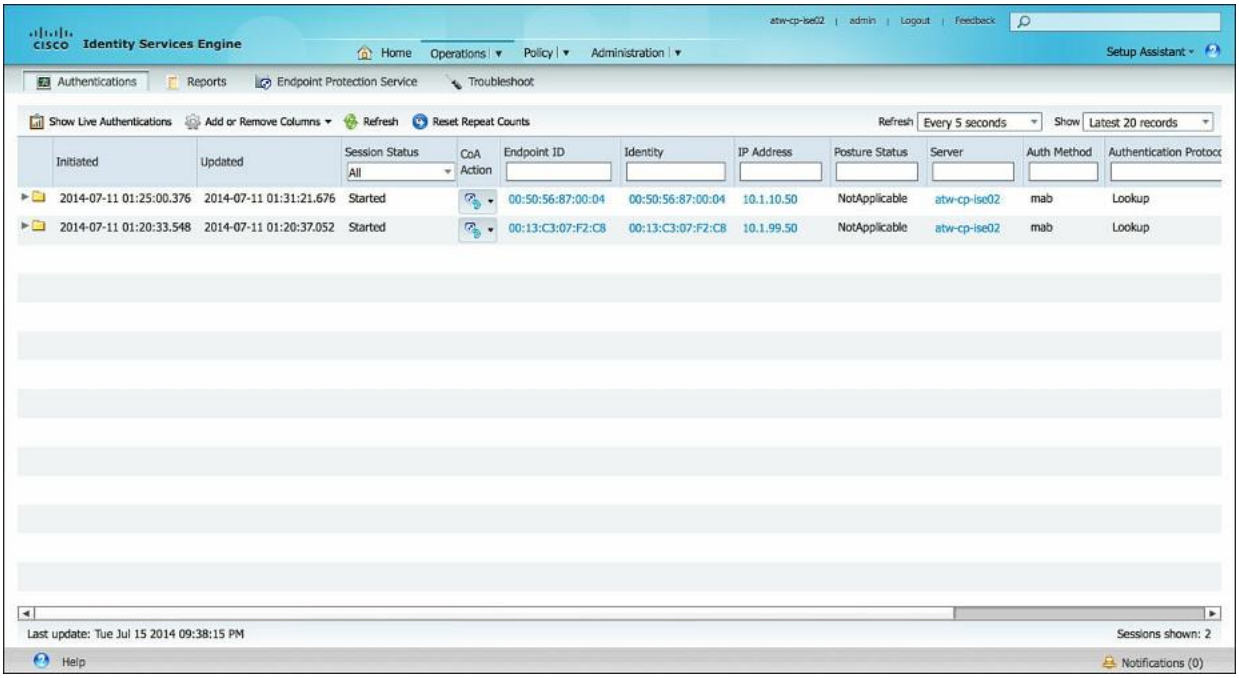


Figure 8-11 Live Sessions Log.

By expanding a session or filtering for the relevant session that you are interested in, you can see the play-by-play of the authentication session (see Figure 8-12). If a session is having issues, this view should provide a sufficient timeline as to when each step of the authentication took place as well as which ones succeeded or failed.

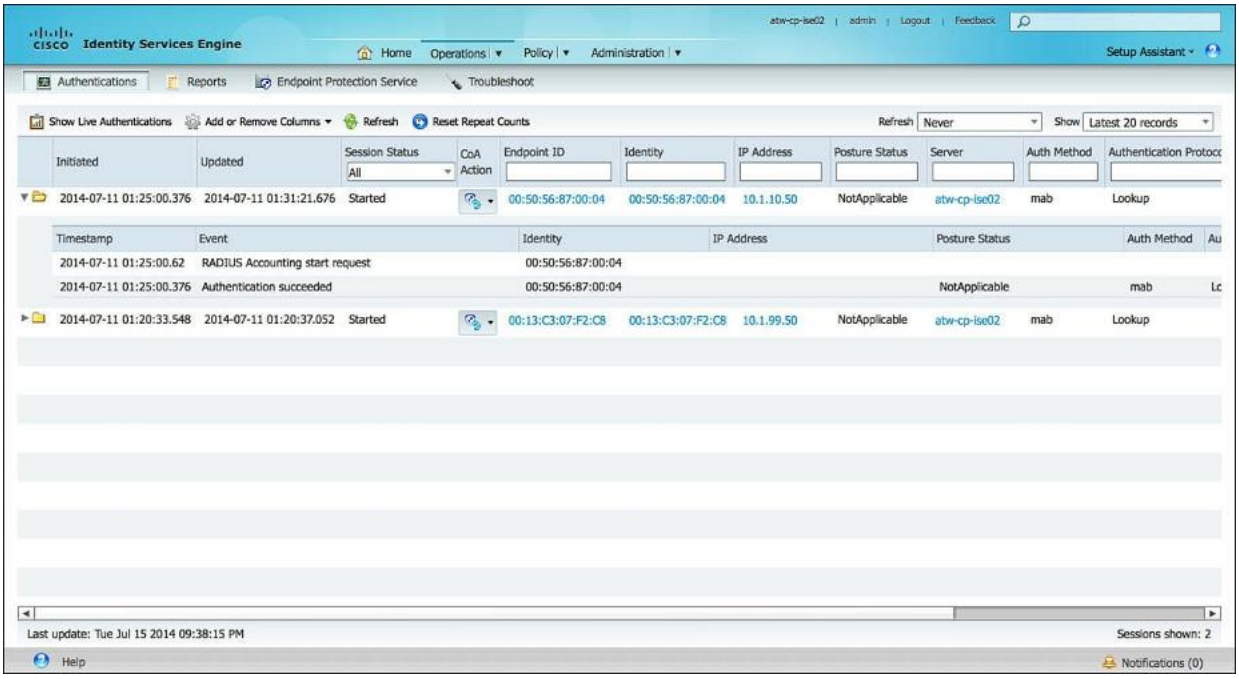


Figure 8-12 Live Sessions play-by-play.

https://t.me/learningnets

A further function of this Live Sessions Log view is the capability to force the termination of the session or to force the session to reauthenticate. This is accomplished by selecting the CoA Action icon for the relevant session, as shown in Figure 8-13.

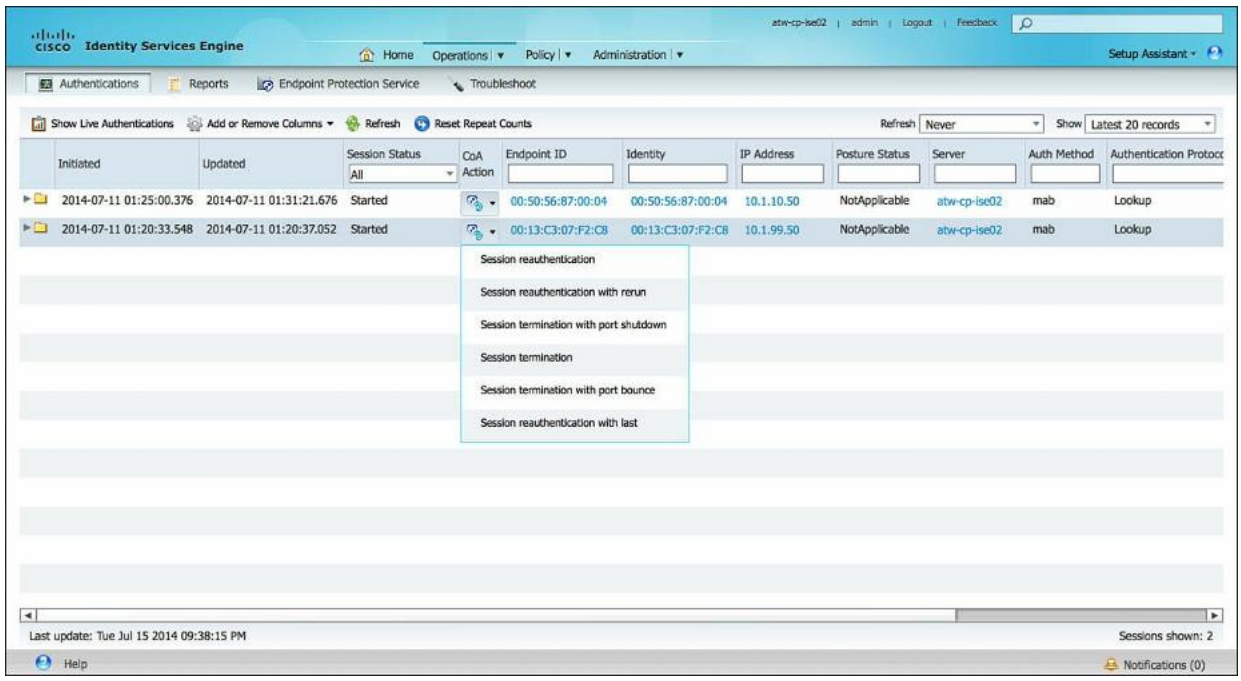


Figure 8-13 Live Sessions additional actions.

Reports

The reports subcomponent of Operations is exactly as it sounds—a place to generate reports for ISE functions and sessions. The reports are broken down into the following categories and also shown in Figure 8-14:

- **Auth Services Status**—Access reports about authentications, authorization, and accounting (also known as AAA services). These reports focus on tracking the transactions between users and Cisco ISE.
- **Deployment Status**—Access reports about the health of the Cisco ISE deployment. These reports focus on the Cisco ISE infrastructure and administration.
- **Endpoints and Users**—Access reports about user provisioning, endpoint profiling, and user management. These reports focus on Cisco ISE end users and endpoints.
- **Security Group Access**—Access reports for the Security Group Access (SGA) feature, which is available with the Cisco Advanced ISE license. These reports provide details about the integration between Cisco ISE and the network devices that support the SGA feature.

https://t.me/learningnets

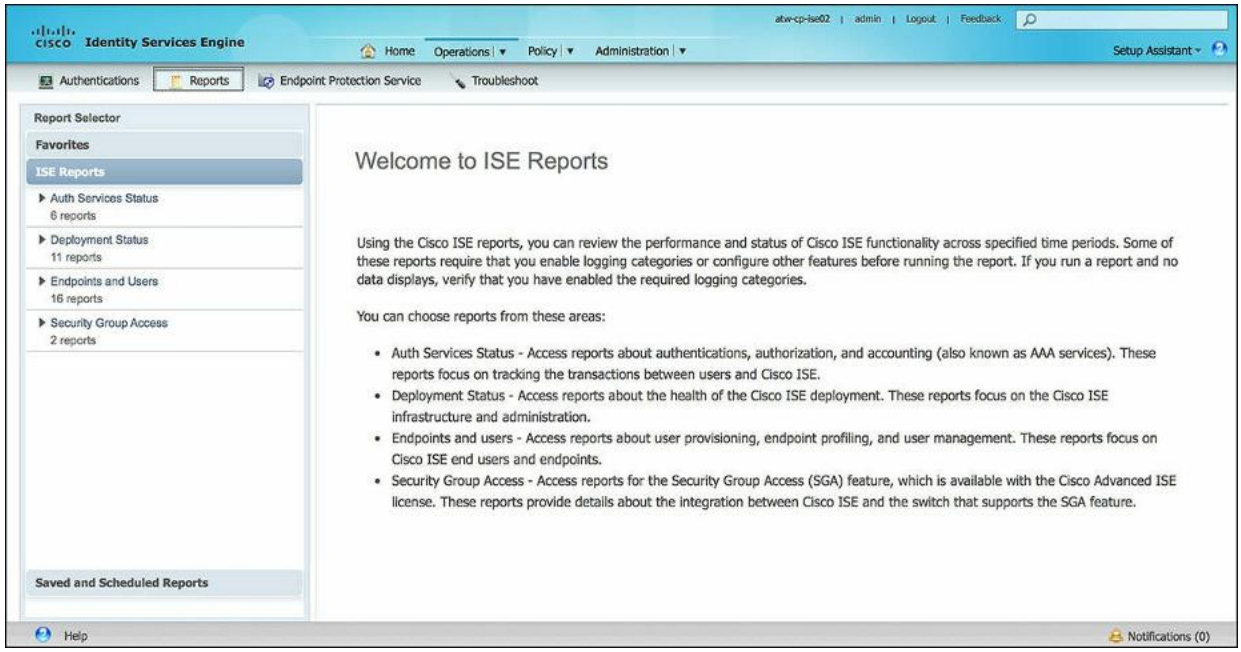


Figure 8-14 Available reports.

Endpoint Protection Service

The Endpoint Protection Service is a useful tool available to network administrators (see Figure 8-15). This function of ISE enables an administrator to manually refuse network access to an endpoint. Depending on the security policy that is configured, a network administrator can quarantine an endpoint using the endpoint's MAC address (most effective) or IP address.

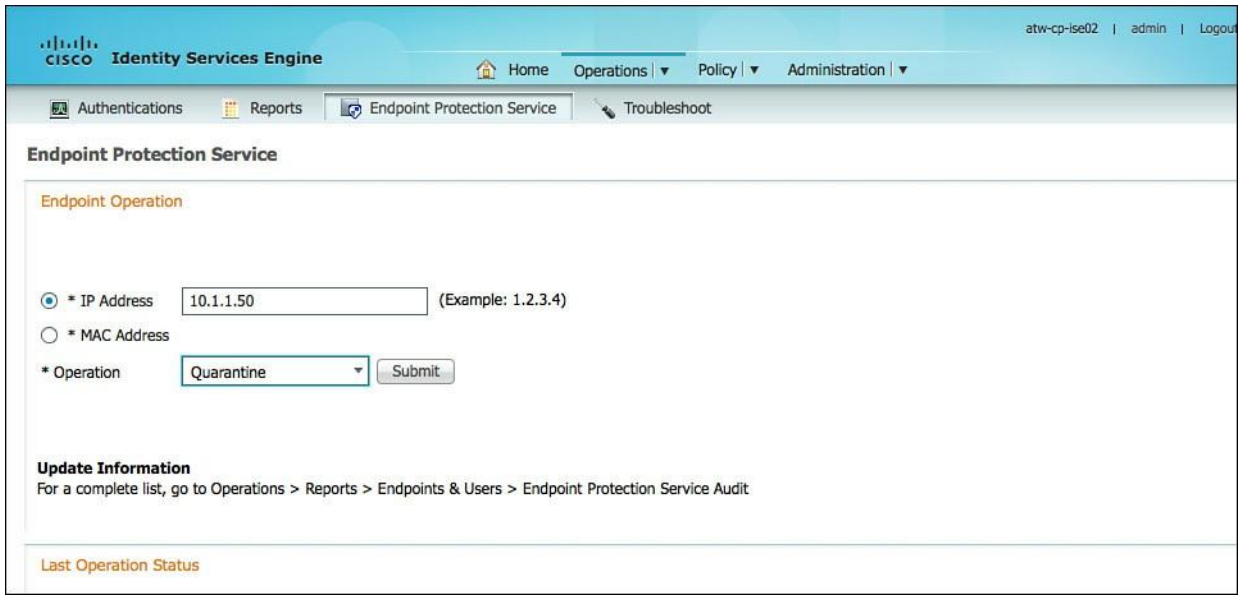


Figure 8-15 Endpoint Protection Services.

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 NAD. In the case of a wired endpoint, this is usually the physical connectivity to the switch and, with wireless, the connection between the endpoint and the Wireless LAN Controller (WLC) from where the RADIUS calls are originated.

This is a very strong-handed approach to keep a device off of the network and is often used to address the most egregious security offenses on the network. In this case, the user or endpoint is simply refused

https://t.me/learningnets

network access without being given any concrete feedback as to why (such as a redirect to a webpage with detailed feedback). A few scenarios where this can be appropriate are a device that is infected that is accessing the network, a rogue device, or a device that is otherwise DoSing the network.

Troubleshoot

The Troubleshoot subcomponent of Operations provides a number of tools for the network administrator to use in case of issues with ISE. These tools provide a number of useful functions including a configuration validation tool (Evaluate Configuration Validator), a RADIUS Authentication Troubleshooting tool, a packet capture utility (TCP Dump), and others, as shown in Figure 8-16.

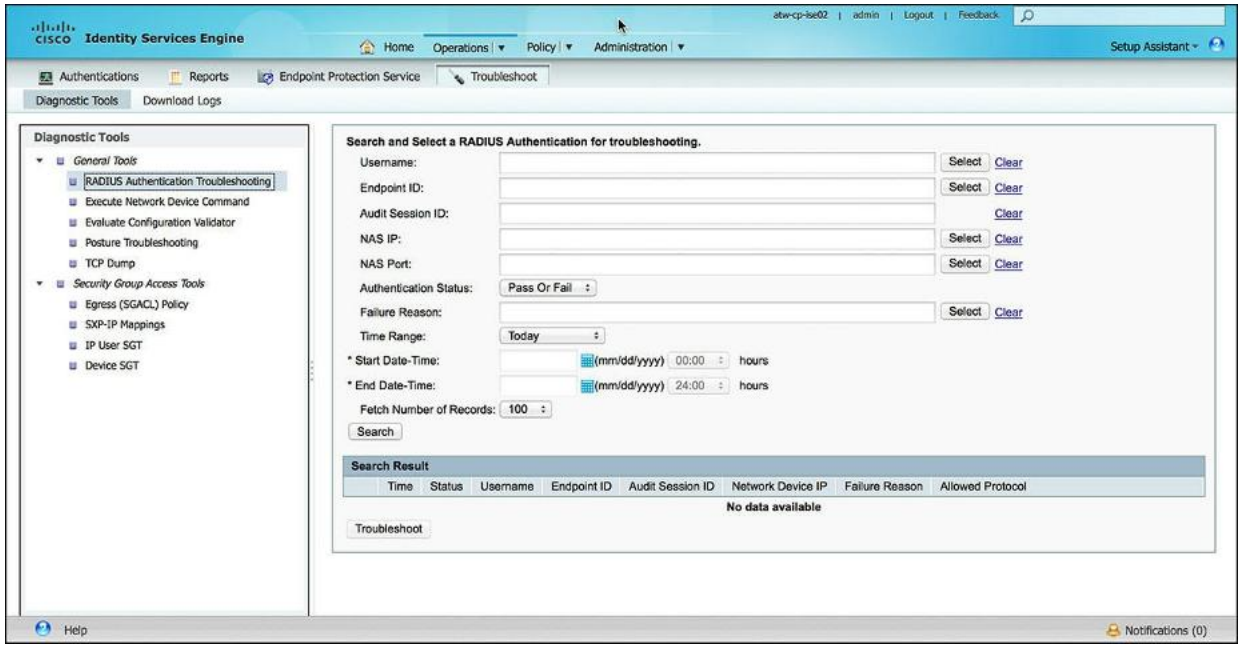


Figure 8-16 Troubleshooting tools.

In those cases where you cannot resolve an issue by monitoring the Live Logs or troubleshooting on your own, you might need to open a support case with Cisco TAC. To provide the customer support engineer with the information she needs, you might be asked to gather the logs from ISE at the time of the issue. These logs are available as a Support Bundle or as individual debug logs, both of which are under the Troubleshoot subcomponent, as shown in Figures 8-17 and 8-18.

https://t.me/learningnets

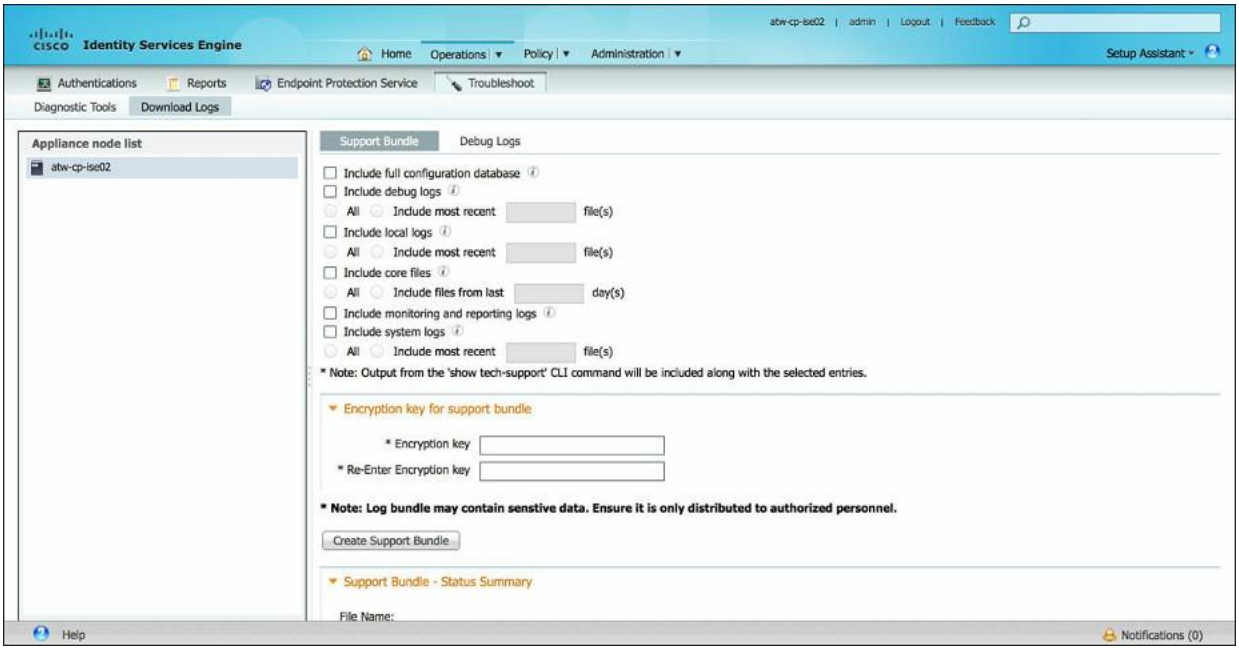


Figure 8-17 Support bundle.

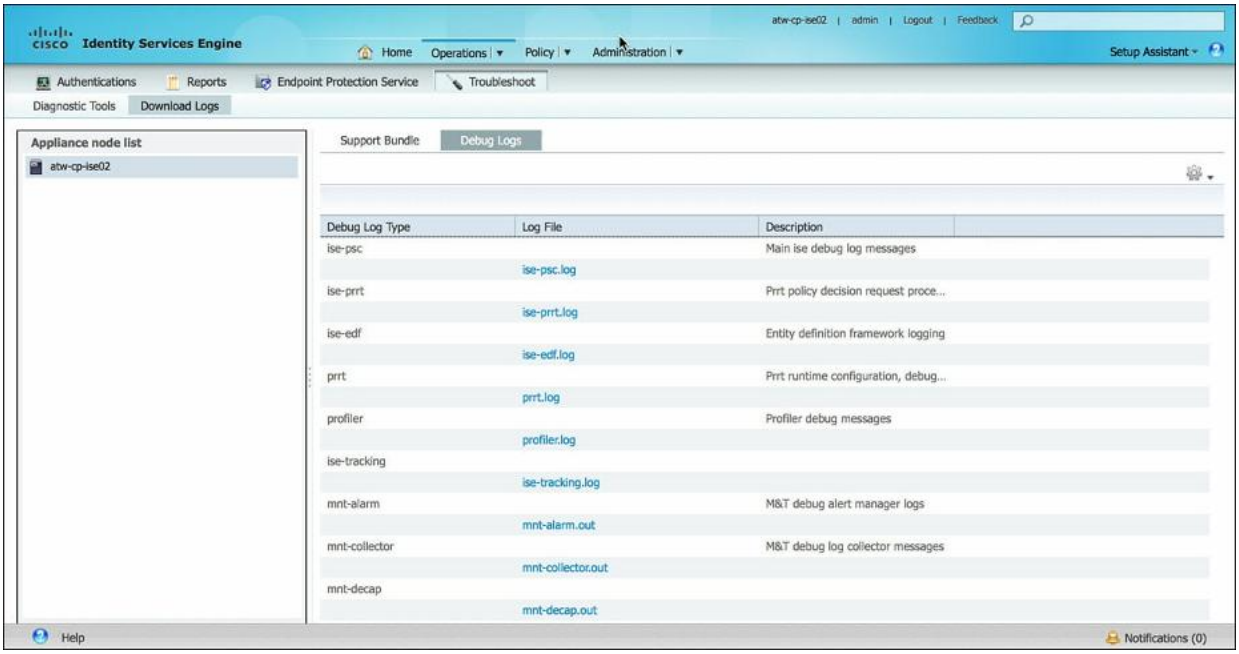


Figure 8-18 Debug logs.

Policy

The next major tab we'll discuss in the ISE GUI is the Policy tab. This component of ISE contains configuration for authentication, authorization, profiling, posturing, client provisioning, security group access, and policy elements.

https://t.me/learningnets

Authentication

The authentication subcomponent of Policy provides the interface to configure authentication policies on ISE. Authentication, as you learned in Chapter 2, "Fundamentals of AAA," is the process by which the identity of the 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. Figure 8-19 shows a sample authentication policy.

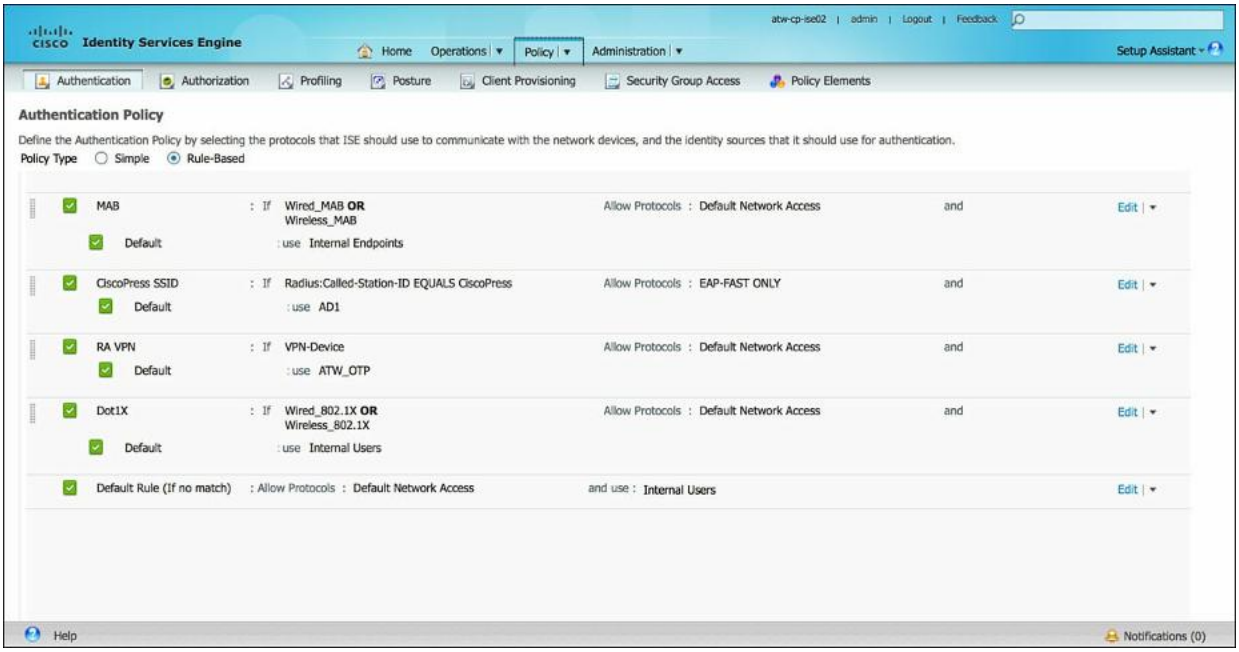


Figure 8-19 Authentication policy.

Authorization

Authorization is the process of determining what an endpoint device will have access to on the network after 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. Each policy has a name, conditions the endpoint must match, and the resulting permission. Figure 8-20 shows a sample set of authorization policies.

https://t.me/learningnets

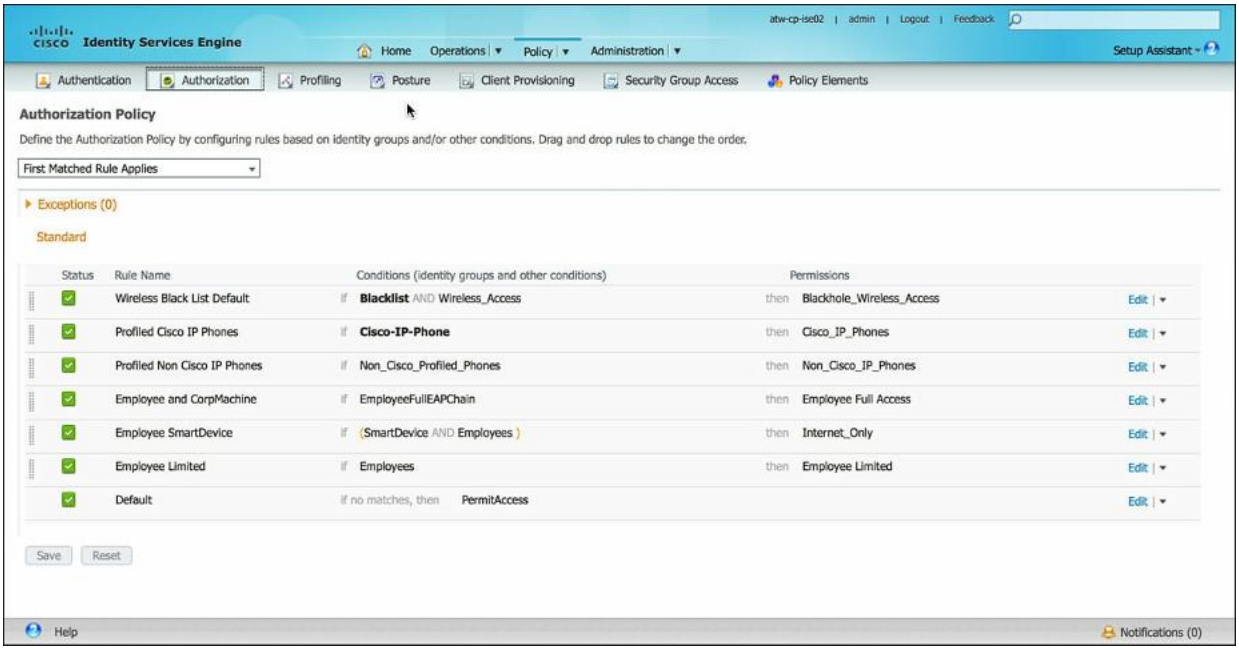


Figure 8-20 Authorization policy.

Profiling

ISE has the ability to determine, or make a very educated guess, about which type of endpoint is attempting to access the network. This process, called profiling, is accomplished by ISE's monitoring of those protocols sent by the endpoint onto the network—namely, DHCP, HTTP, and (by proxy) RADIUS, amongst a few others. The profiling done by ISE can sometimes be general in nature, identifying just the vendor or just the device type, or very specific, identifying the endpoint's vendor, device model, and function on the network (such as a printer, phone, and so on). A large number of profiling policies are included by default, a sample of which can be seen in Figure 8-21.

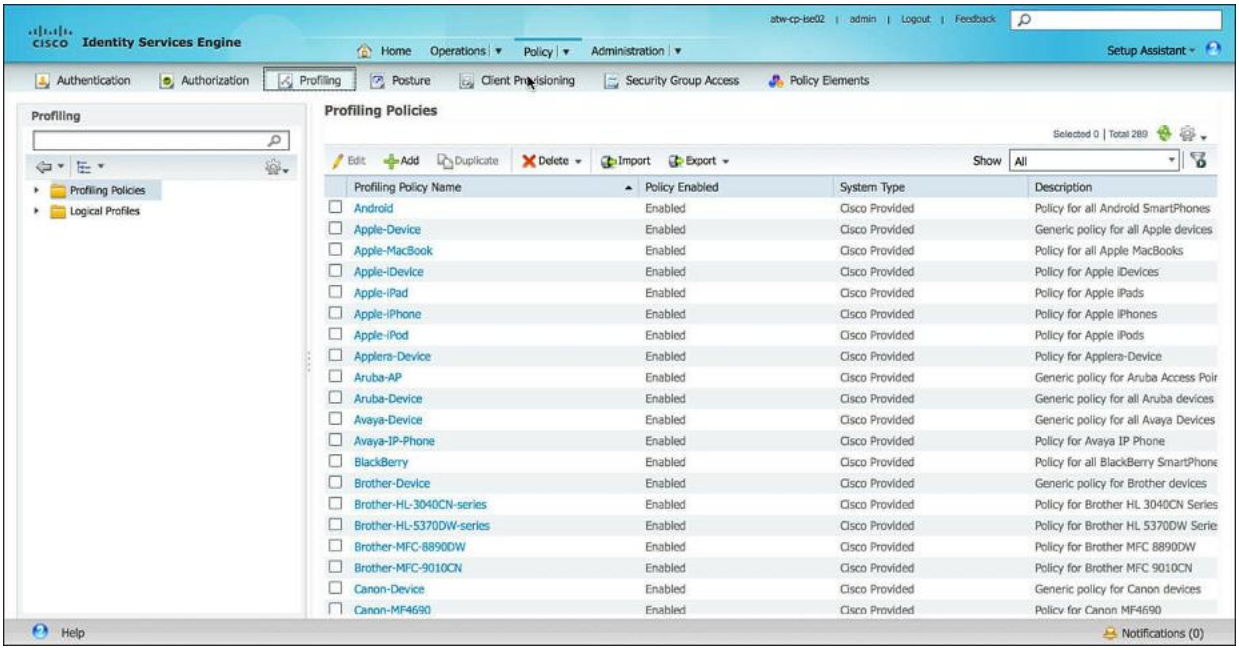


Figure 8-21 Profiling policy.

The profiling policy on ISE enables an administrator to leverage the Cisco preconfigured profiling policy to identify endpoints and/or create their own policy, defining the rules or criteria by which the type of endpoint that is accessing the network is determined.

https://t.me/learningnets

Posture

Posturing is a mechanism by which a software supplicant or agent on the endpoint provides ISE detailed information about the endpoint's software and hardware configuration. The information provided by the endpoint can include information about the operating system, antivirus/antispyware/antimalware software, as well as current OS patch levels. 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, comparing it to the configured posture policy on ISE, ISE might quarantine the endpoint with a noncompliant or unknown posture—possibly providing the endpoint reduced network access to remediate its posture. After taking all remediation steps needed to become compliant, the endpoint then can be allowed appropriate network access. The posture policy screen is shown in Figure 8-22.

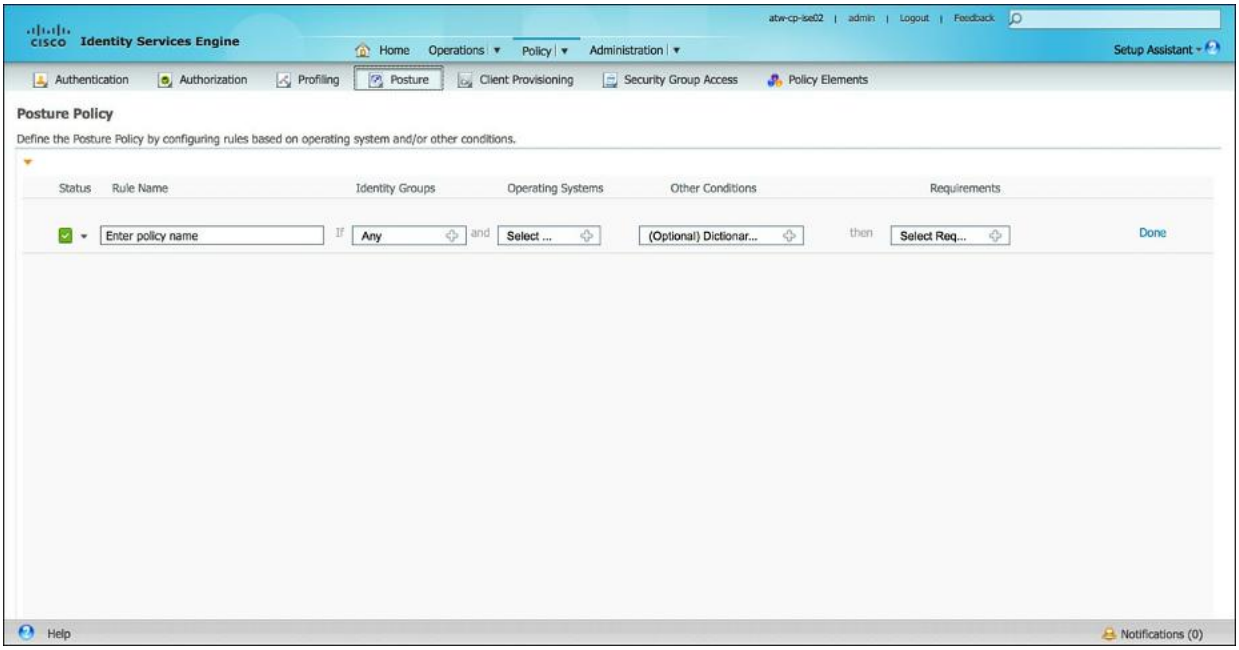


Figure 8-22 Posture policy.

Client Provisioning

In the Client Provisioning section of the Policy component, you can configure how each endpoint type will be provisioned to join the network. This can 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) must be downloaded onto the endpoint. After being onboarded, the user of the endpoint then can join the network in the future with minimal, if any, manual interaction.

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 8-23), based on the OS and other criteria provided by the endpoint and the NAD, ISE will present the required credentials (again, X.509 certificates and wireless profiles) to the endpoint to store for future use. This sometimes might require a supplicant built in to the OS or one that is automatically downloaded from ISE (often a Java- or ActiveX-based supplicant).

https://t.me/learningnets

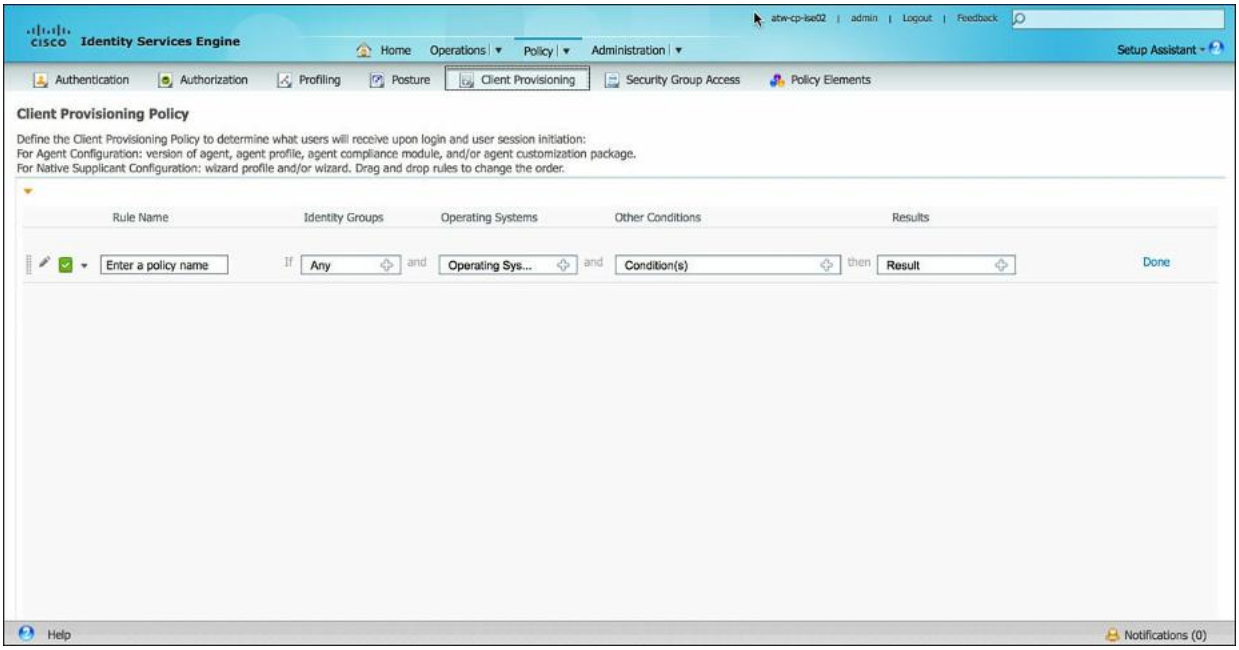


Figure 8-23 Client provisioning policy.

In a similar fashion, the client provisioning process can be used to deploy a Cisco NAC Agent onto the endpoint if posturing is being used.

Security Group Access

The Security Group Access (SGA) section of the Policy component is where ISE's TrustSec access policy is defined. After identifying a NAD as a TrustSec-enabled or -capable device (via the **Administration > Network Resources > Network Devices > Network Devices**, which will be discussed shortly), an administrator must configure the appropriate SGA policy. Specifically, he must configure what each Security Group Tag source or destination will be allowed to access. SGA policy configuration is discussed in more detail in Chapter 11, "Authorization Policies," and Chapter 18, "TrustSec and MACSec." Figure 8-24 shows a sample SGT policy matrix.

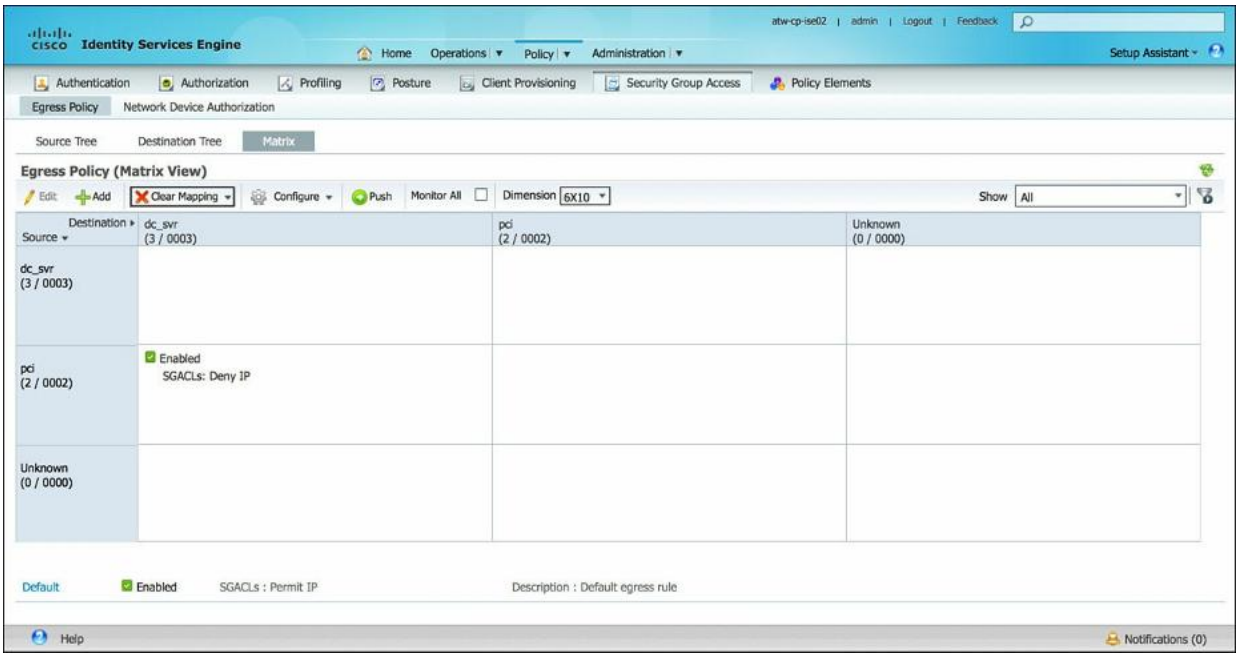


Figure 8-24 Security Group Access policy.

https://t.me/learningnets

Policy Elements

Policies are often comprised of numerous building blocks, or elements. These building blocks 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 section of the Policy component is the place within ISE where you can add and modify these reusable building blocks.

Within Policy Elements, there are subsections defined for Dictionaries, Conditions, and Results:

- **Dictionaries**—A list of system and user-defined elements and their available attributes. The individual dictionary elements leverage an Attribute-Value pair approach, similar to RADIUS, whereby each chosen attribute is assigned a value by the endpoint, the NAD, or ISE. These attributes and their assigned values can be used as the smallest building blocks within the ISE policy.
- **Conditions**—Many of the policy constructs within 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 the ultimate value of TRUE or FALSE. If the totality of this conditional statement is TRUE, everything following the THEN portion of the clause is executed. The Conditions subsection of Policy Elements is the location within ISE where the IF portions of ISE policy are defined.
- **Results**—As noted in the Conditions definition, many of the policy constructs within ISE leverage an IF-THEN programmatic approach. The Results subsection of Policy Elements is the location within ISE where the THEN portions of ISE policy are defined. Therefore, if the IF statement (the conditions) result in a TRUE outcome, everything following the THEN (the results) will be executed.

Administration

The final tab we'll discuss in the ISE GUI is the Administration tab. The Administration component is comprised of all setup-type functions of ISE. These functions are those that are often set up one time and rarely modified thereafter. These include System, Identity Management, Network Resources, Web Portal Management, and the Feed Service.

System

The System section of the Administration configuration component enables the administrator to configure those elements that affects ISE's behavior directly, including

- **Deployment**—Under the Deployment subsection, you can configure all the ISE nodes. Each node must be defined and registered within the Deployment configuration on the Admin node. As discussed in Chapter 7, each node is assigned at least one persona, role (primary, secondary, or standalone), as well as an IP address or a DNS-resolvable hostname.

  The Deployment subsection, shown in Figure 8-25, is also the place within ISE where you enable the profiling function on a per-node basis—that is, which nodes will actively participate in endpoint profiling.

https://t.me/learningnets

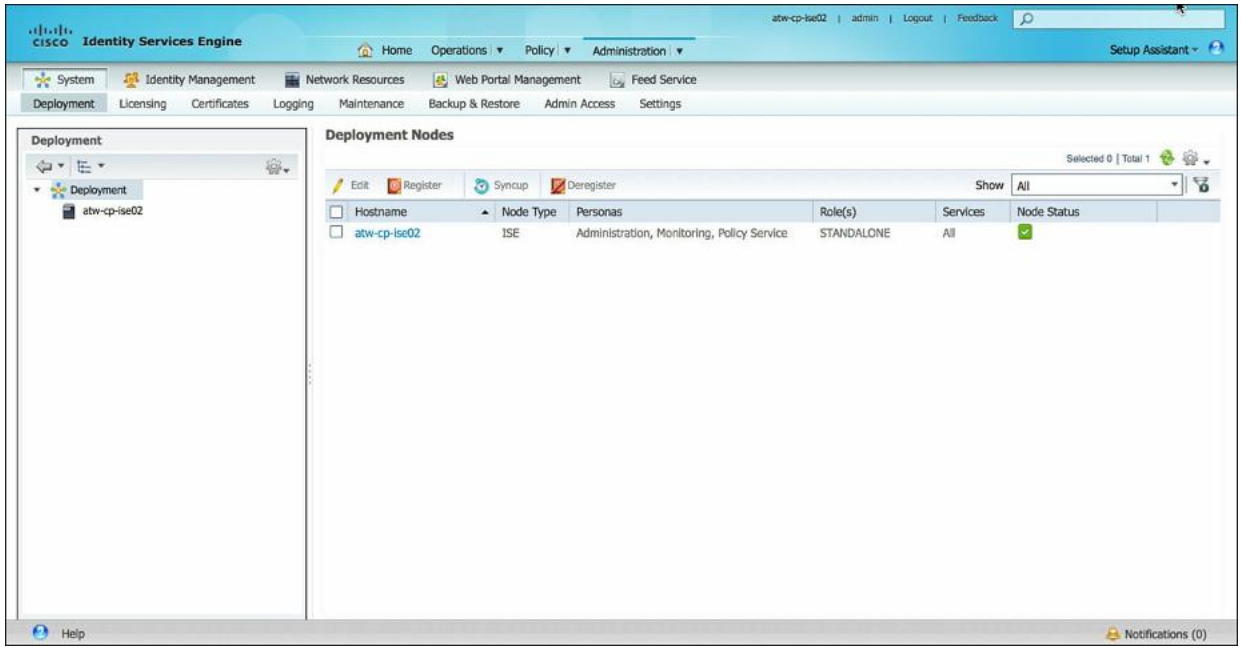


Figure 8-25 ISE deployment nodes.

- **Licensing**—Every feature within 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. Those features that require periodic updates from Cisco require an Advanced license. By default, ISE comes with a 90-day Evaluation license. A breakdown of what features are covered with each license is provided in Table 8-5, and Figure 8-26 shows the licensing screen.

https://t.me/learningnets

License Package	Perpetual or Subscription	ISE Functionality Covered	Notes
Base	Perpetual	Basic network access: AAA, IEEE-802.1X Guest management Link encryption (MACSec)	
Plus	Subscription (1, 3, or 5 years)	Bring Your Own Device (BYOD) Profiling Endpoint Protection Service (EPS) TrustSec SGT	Does not include Base services. A Base license is required for each Plus license.
Advanced	Subscription (1, 3, or 5 years)	Bring Your Own Device (BYOD) Profiling Endpoint Protection Service (EPS) TrustSec SGT Mobile Device Manager Health Compliance and Remediation Posture	Does not include Base services. A Base license is required for each Advanced license. The Advanced license includes all the functionality of the Plus license.
Wireless	Subscription (1, 3, or 5 years)	A Wireless license turns on the functionality of Base and Advanced licenses for wireless LAN deployments.	Cannot coexist on a Cisco Administration node with Base, Plus, or Advanced Licenses.
Wireless Upgrade	Subscription (1, 3, or 5 years)	A Wireless Upgrade license turns on the functionality of Base and Advanced licenses for all wireless and nonwireless client-access methods, including wired and VPN Concentrator access.	You can only install a Wireless Upgrade License on top of an existing Wireless license.
Evaluation	Temporary (90 days)	Full Cisco ISE functionality is provided for 100 endpoints.	Limited use of Cisco ISE product for presale customer evaluations. All Cisco ISE appliances are supplied with an Evaluation license.

Table 8-5 Cisco ISE License Packages

https://t.me/learningnets

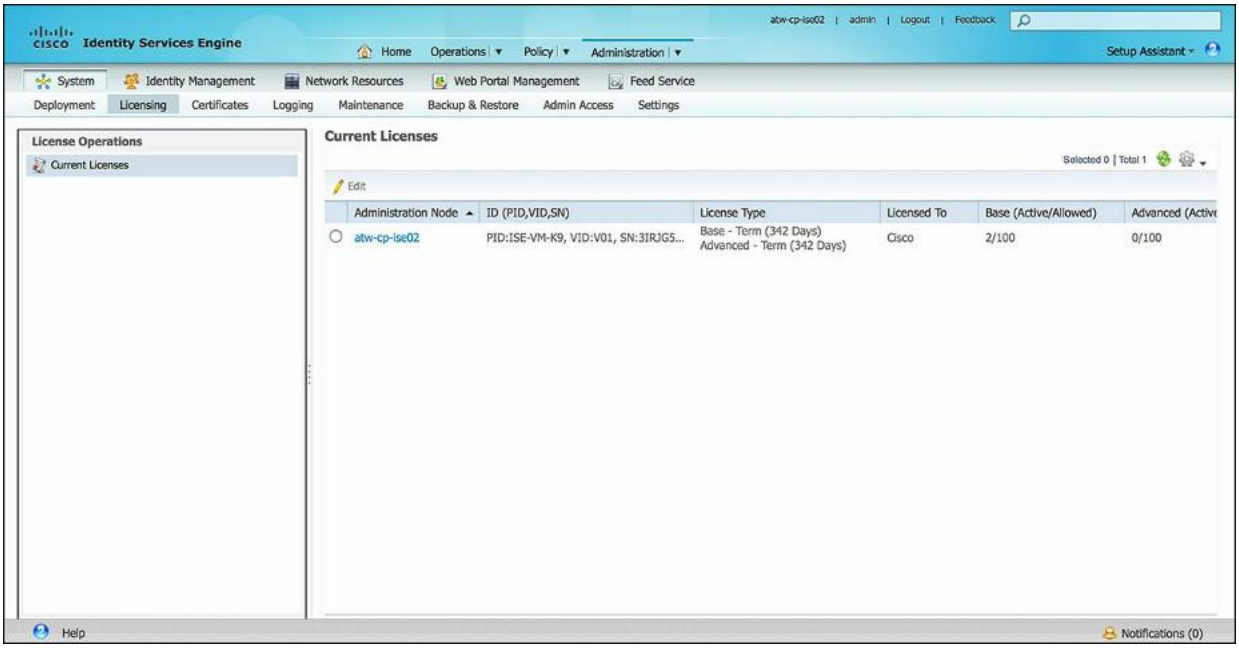


Figure 8-26 ISE licensing.

- **Certificates**—X.509 certificates (often referred to as simply certificates) are the key components to a Public Key Infrastructure (PKI), providing a scalable deployment option that marries each endpoint to the user as well as nonrepudiation. To ensure that ISE can properly authenticate endpoints and the endpoints can properly authenticate ISE, we must leverage the X.509 certificates (see Figure 8-27) for both server authentication as well as client authentication.

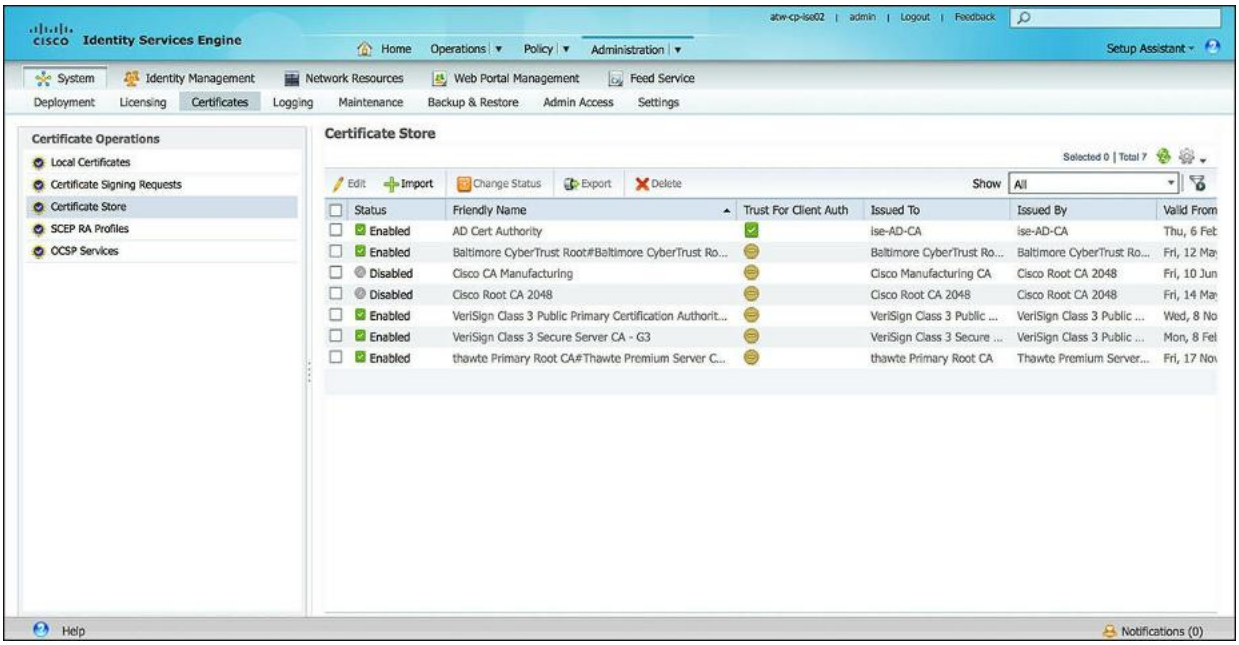


Figure 8-27 X.509 certificates.

The Certificates subsection enables an administrator to configure the certificates that will be used for the administration portal as well as any portal that is served to the client as part of the onboarding or remediation process. The requisite certificate authorities 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 9.

- **Logging**—The Logging subsection, shown in Figure 8-28, is where the administrator configures the specifics of how ISE creates and stores syslogs. In the logging configuration, the administrator can

https://t.me/learningnets

define how long logs will be kept, how detailed the logs will be, and where to send the logs. These logs might 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 that is chosen is too low, the amount of detail provided in these logs might not be sufficient to address any issues as they occur. If the logging level that is chosen is too high, then the amount of logs generated can adversely affect storage space on the ISE appliance and affect the overall performance of the appliance.

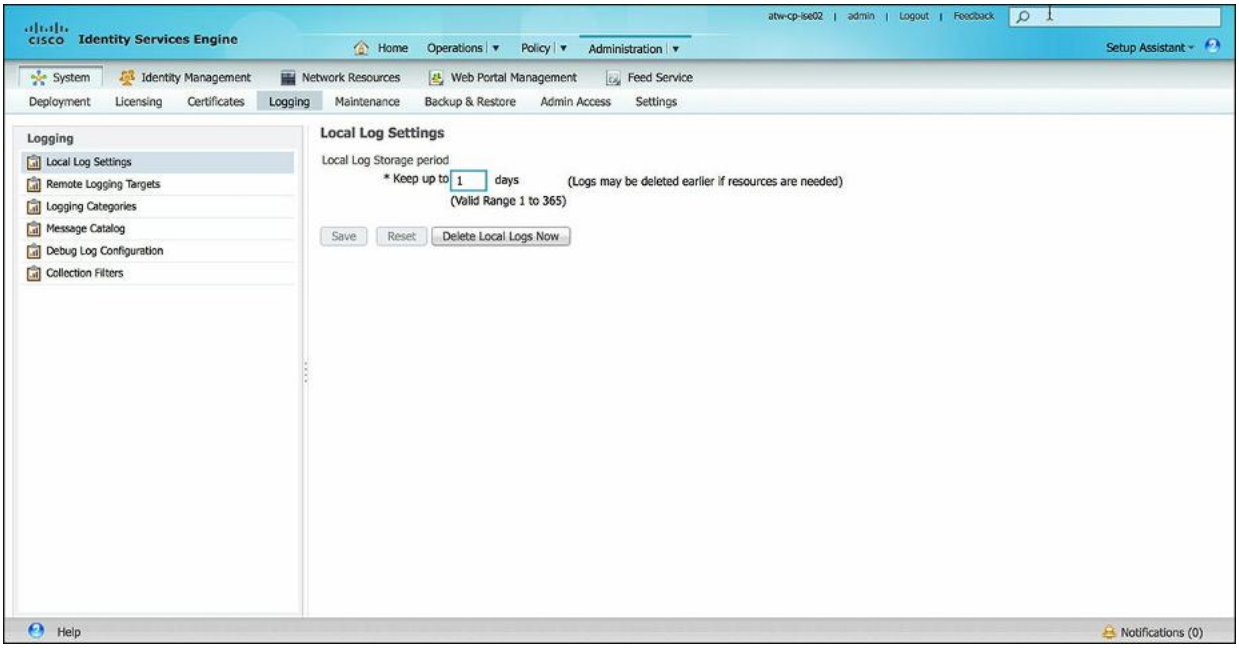


Figure 8-28 Logging configuration.

Key Topic

The level of logging 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 might request detailed logs at the time of the issue. If debug-level logs are requested, it is best to enable the debug level logs temporarily to gather the requested level of detail. As soon as the requested information is gathered with the debug-level logs, you should return the logging level to a more neutral logging level. Figure 8-28 shows the logging configuration screen.

- **Maintenance**—As with software systems from other vendors, Cisco periodically releases software updates to add features to ISE or to resolve any defects. On a periodic basis, Cisco compiles a patch to address a number of defects that have been found within previous ISE releases. To load these patches, manage download locations for software maintenance, or manage how often operations data is purged to maintain optimal system performance, you must access the Maintenance subsection of the System configuration.
- **Backup & Restore**—As part of any disaster recovery process, it is important that an administrator periodically perform a backup of all network configurations. ISE's configuration is no different. The Backup & Restore subsection enables 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. If disaster does strike, the configuration can be restored accordingly.
- **Admin Access**—The Admin Access subsection (see Figure 8-29) does exactly as the name implies—it defines how an administrator can access ISE. Some of the options available from the Admin Access section include usernames and passwords for administrative users, the administrative groups that are defined, the requisite password policy for administrative users, and from where and for

https://t.me/learningnets

how long administrators are allowed to access the administrative portal of ISE. Additional options are also available on the Admin Access section.

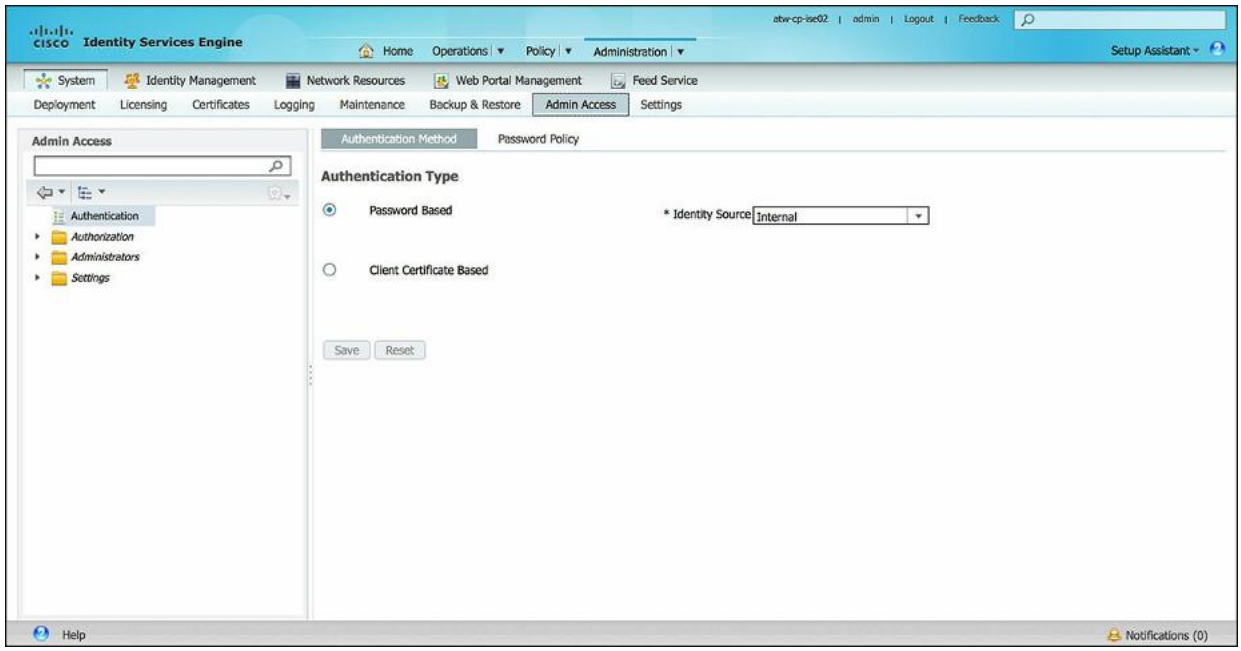


Figure 8-29 Admin access configuration.

- **Settings**—The Settings subsection (see Figure 8-30) of the System configuration enables an administrator to define how the system will behave. Some features configured via the Settings subsection include client provisioning support, Endpoint Protection Services, FIPS Mode, posture timers, and profiling CoA behavior. Proxy, SMTP, and NTP servers can also be defined here.

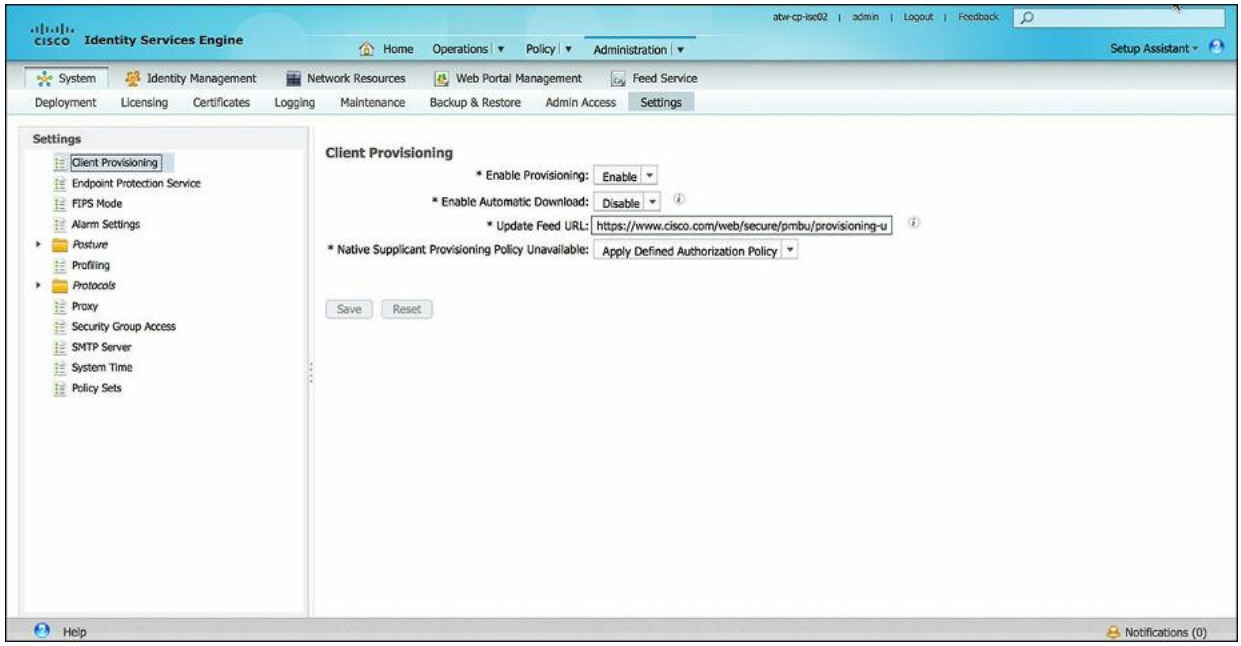


Figure 8-30 Admin settings configuration.

Identity Management

Identity management is a key construct within ISE. Nearly every step within the ISE authentication and authorization processes requires at least one identity as the subject of the operation. The Identity Management subsection of Administration enables an administrator to manage all ISE configurations that pertain to identities:

https://t.me/learningnets

- **Identities**—The Identities subsection (see Figure 8-31) is used to manage both users and endpoints. Whenever a user or an endpoint first authenticates or authorizes to ISE—as well as any subsequent reauthentications—information about the entity is updated here. If a user or an endpoint needs to be removed from ISE, the administrator will remove the user/endpoint from this Identities database.

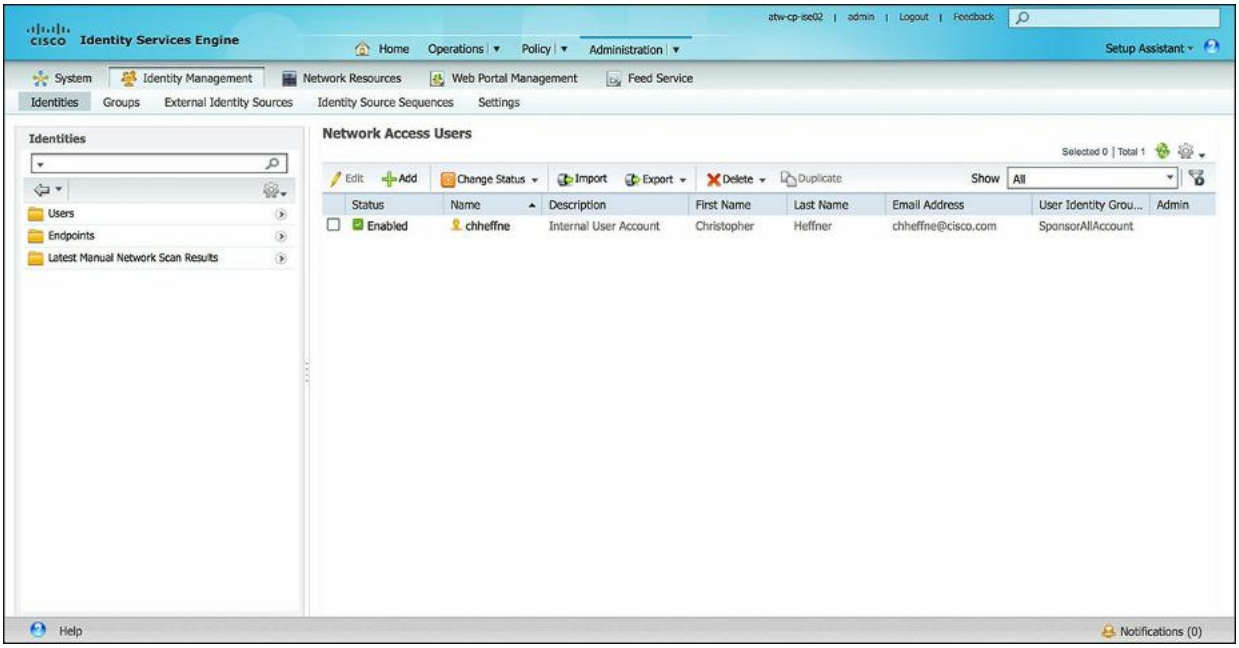


Figure 8-31 Identities configuration.

- **Groups**—Identity groups, shown in Figure 8-32, enable administrators to subdivide and manage users or endpoints based on common criteria. 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 could have one security policy while all laptop computers could have a slightly different security policy.

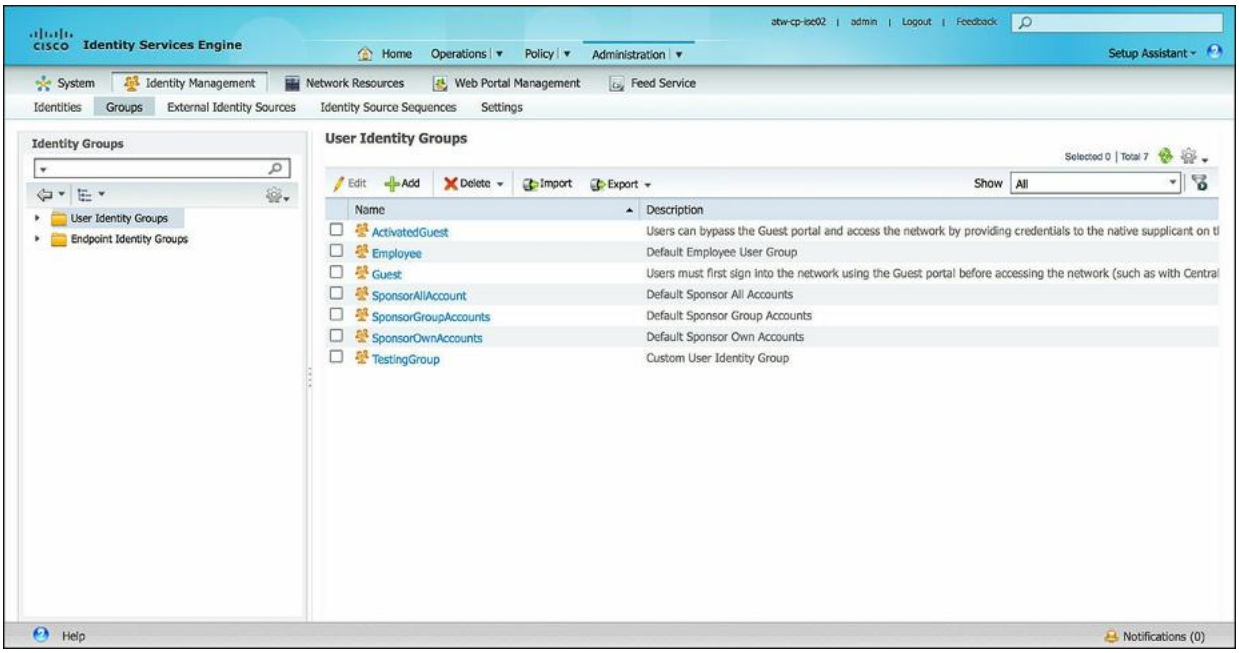


Figure 8-32 Identity group configuration.

- **External Identity Sources**—In many ISE deployment scenarios, ISE is being deployed into a network that has been running for a while. Often, the user identities and management are stored

https://t.me/learningnets

outside of ISE. Instead of importing each identity into ISE and, therefore, managing the user in two different places, ISE allows for external identity sources.

These external identity sources can include Active Directory (as illustrated in Figure 8-33), LDAP, or RSA SecurID (token-based, two-factor authentication). Depending on the authentication policy as configured on ISE, a user's or an endpoint's credentials may be authenticated to any one of these external identity sources. We discuss authentication more in Chapter 9, "Initial Configuration of Cisco ISE," and Chapter 10, "Authentication Policies."

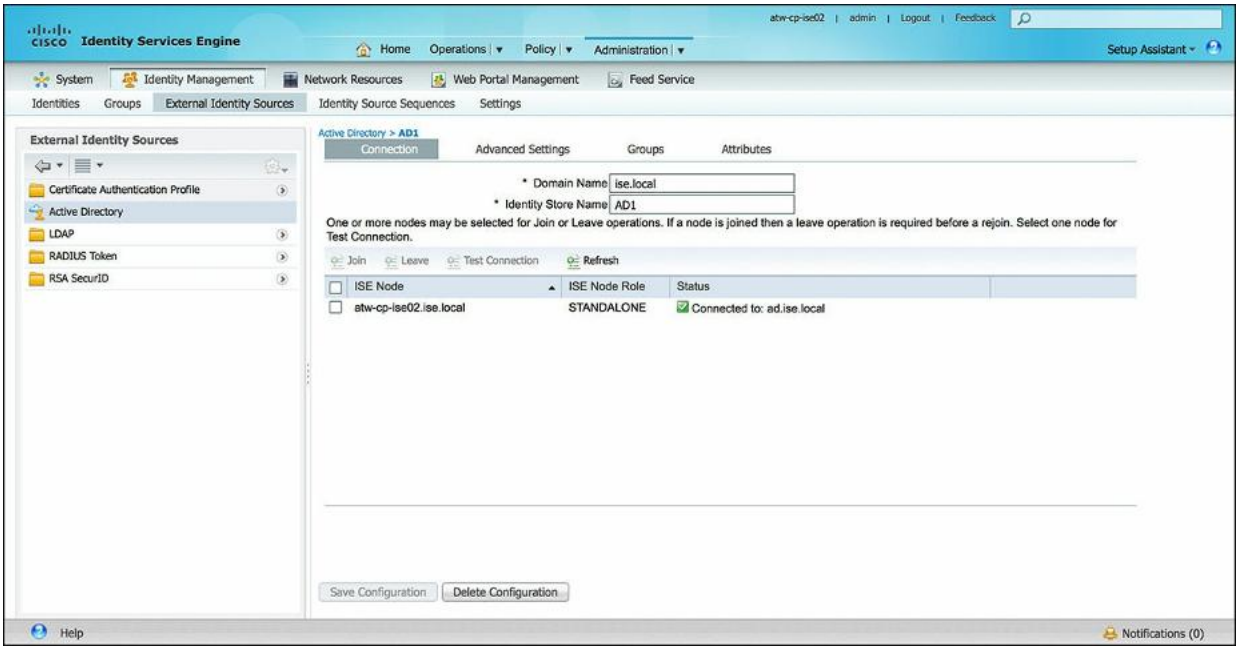


Figure 8-33 External identity sources configuration.

- **Identity Source Sequences**—Used in conjunction with the Local and External Identity Sources as mentioned previously, Identity Source Sequences as shown in Figure 8-34 allow a policy to sequentially check a number of Identity Sources for a successful authentication. This sequential search can be useful if users or endpoints can exist in multiple identity databases. A policy can then be configured to "check all identity stores sequentially" as a single rule versus searching each identity store separately. This approach can greatly simplify authentication policy. The authentication process ends upon the first successful authentication as based on the chosen Authentication policy. Again, we'll discuss authentication more in Chapter 9 and Chapter 10.

https://t.me/learningnets

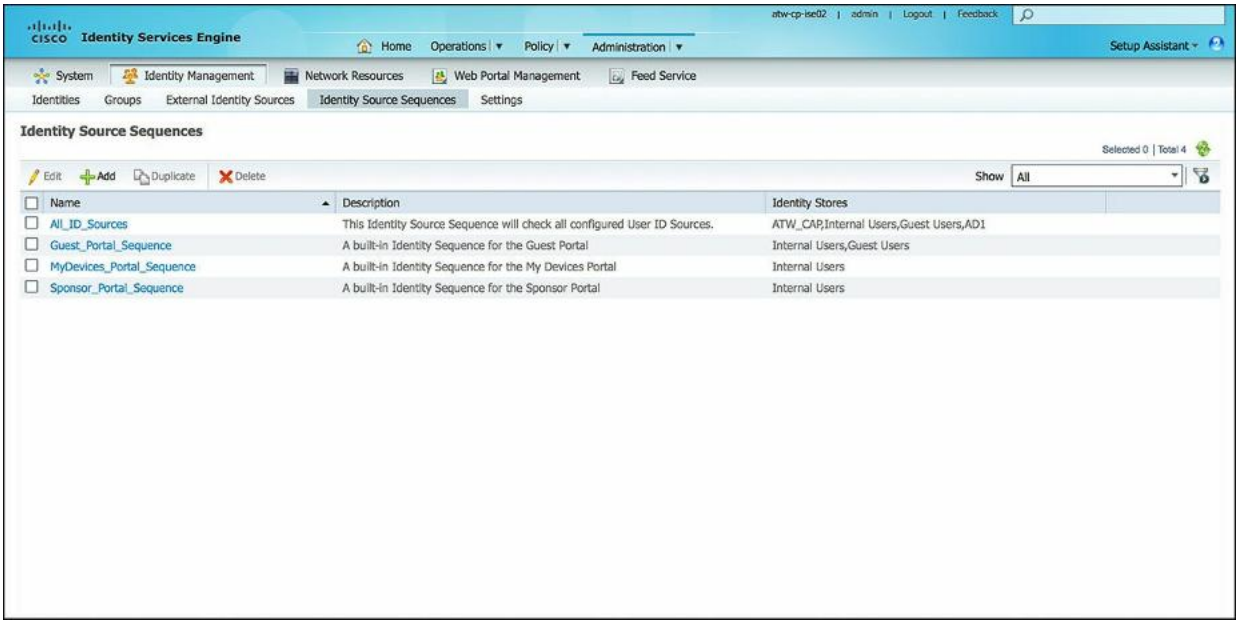


Figure 8-34 Identity source sequences configuration.

- **Settings**—The Identity Management Settings subsection, shown in Figure 8-35, can be used to define custom attributes that can be associated to each user, such as the user password policy.

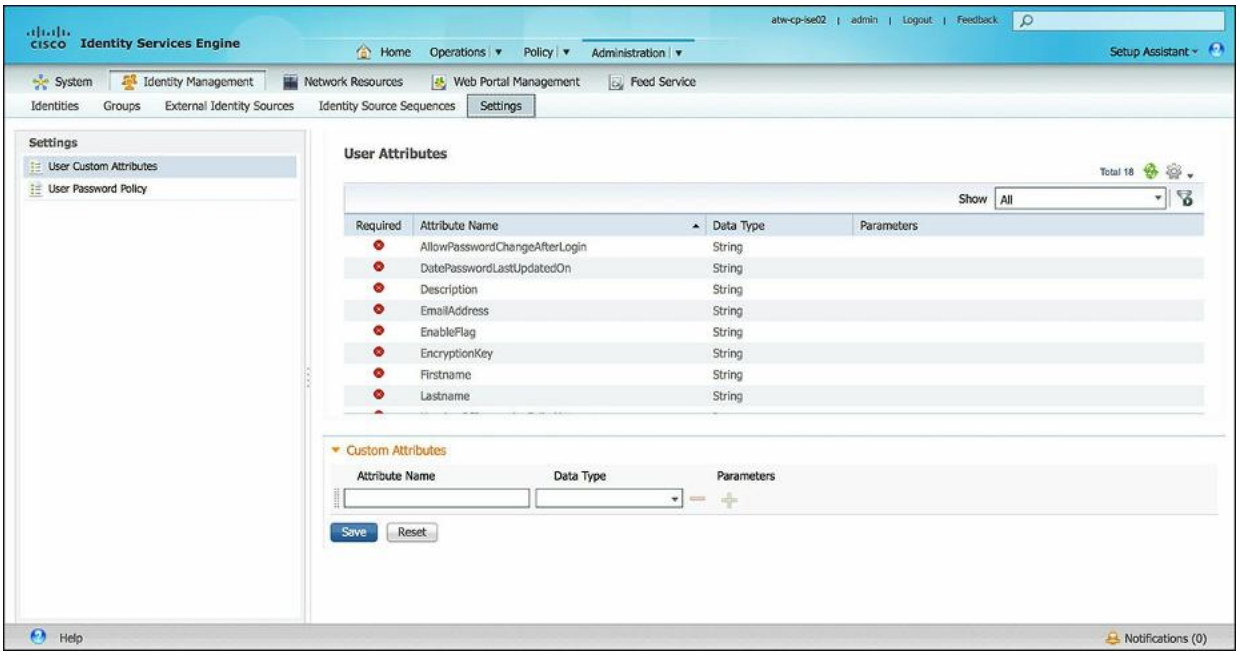


Figure 8-35 Identity management settings configuration.

Network Resources

The Network Resources section of the Administration component is where all external, supplementary resources are configured. These resources include all the NADs, external RADIUS servers, NAC managers, and MDMs:

- **Network Devices**—Whenever a new NAD, such as a switch, router, firewall, or Wireless LAN Controller (WLC), is added to the network, you might need to add this device's credentials to ISE. This will enable the NAD to authenticate endpoints to ISE. Without being added to the Network Devices, this NAD will not be able to use any of the AAA resources provided by ISE.

  For each NAD that is added to this Network Devices database (see Figure 8-36), additional

https://t.me/learningnets

information can be provided about the NAD, including the IP address(es), model name, software version, device location, and device type. Additionally, 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.

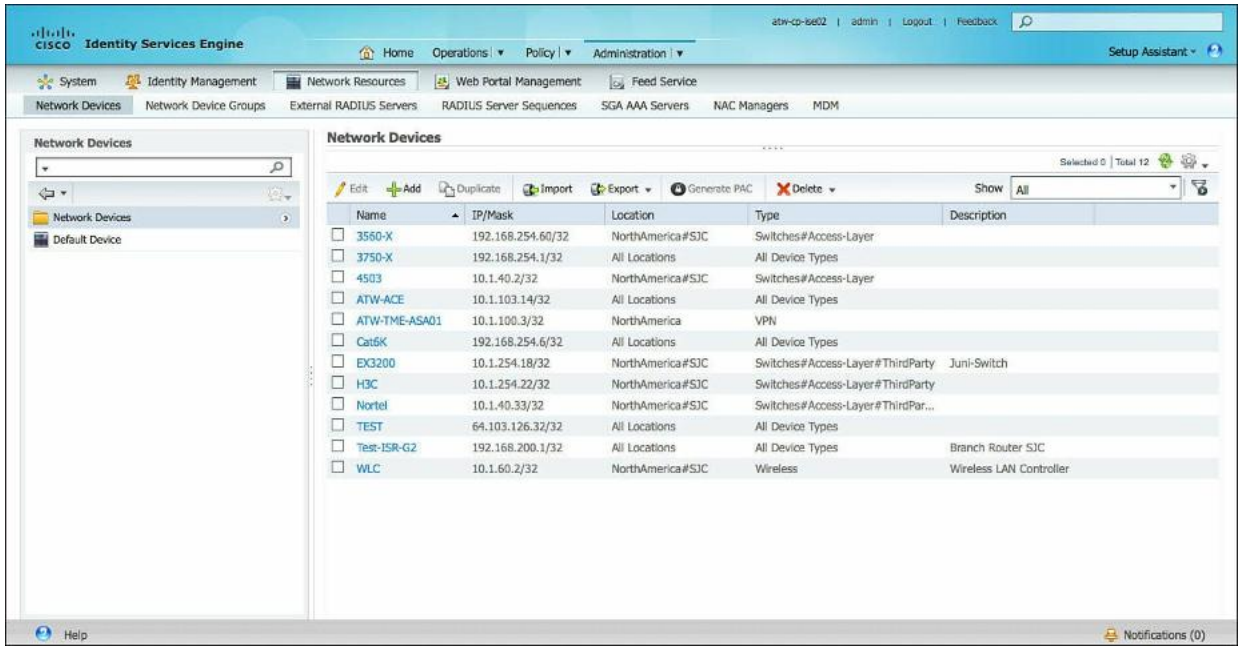


Figure 8-36 Network devices configuration.

- **Network Device Groups**—As the number of NADs on the network increases, the ability to manage these network devices can become incrementally more difficult. Furthermore, as the ISE security policy is configured, you can greatly simplify the policy rules by grouping the NADs according to certain criteria.

  By default, ISE enables each device to have a device type and device location component, though additional criteria can be defined. The device type can refer to the functional role of the device on the network—such as a router, switch, firewall, or WLC. The device location might refer to the geographical 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 want.

  These device groups can be used within the security policy to provide a differentiated level of access to the endpoints, 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. It will also allow the network administrator a level of flexibility in the policy that is pushed to the NAD. If certain NADs support a more feature-rich set of 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 leveraging the default device type group or adding your own custom device group). For an example of a custom device group, consider the “Stage” group in Figure 8-37—used to facilitate a tiered, rollout approach for ISE.

https://t.me/learningnets

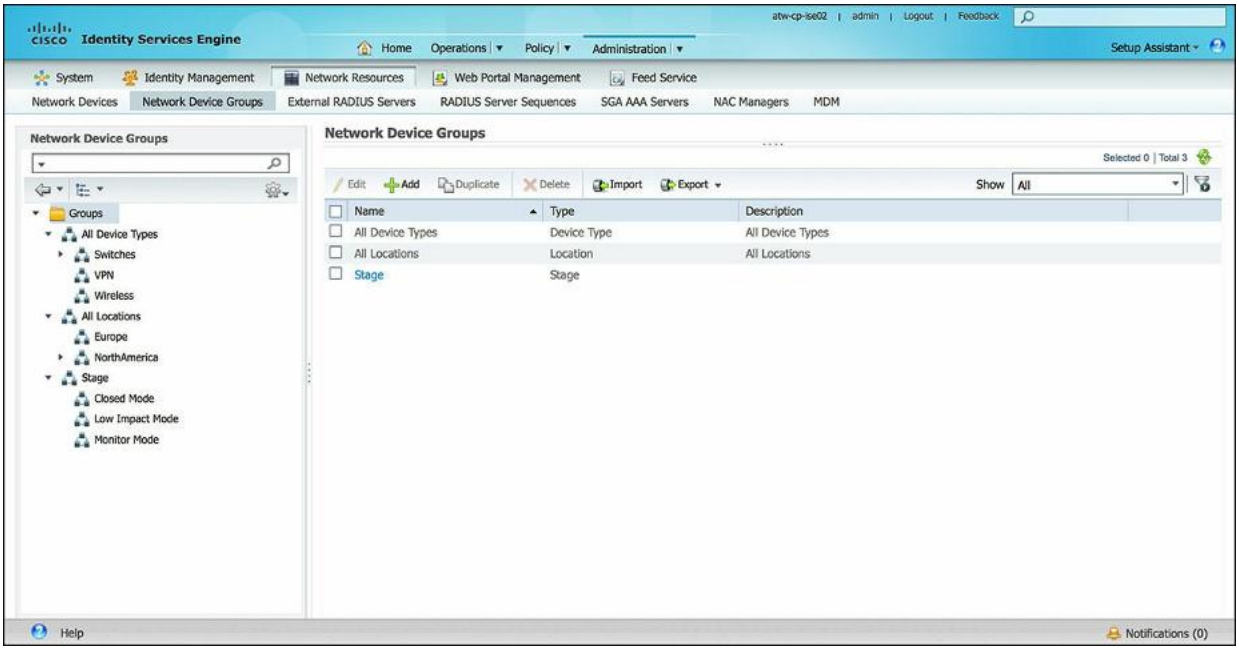


Figure 8-37 Network groups configuration.

We'll learn more about differentiating levels of access more in Chapter 10 and Chapter 11, "Authorization Policies."

- **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 deploy ISE. If an external RADIUS server is in use, ISE can receive the RADIUS requests and proxy them to the external server. The response from the external RADIUS server will also be proxied by ISE back to the original NAD.
- **RADIUS Server Sequences**—RADIUS server sequences leverages the External RADIUS Servers that we mentioned previously. After configuring the external RADIUS servers, you can select one or more of these servers to comprise a RADIUS server sequence. Furthermore, you can elect whether you would like to modify the RADIUS request before it is forwarded to the external RADIUS server and whether you want to modify the RADIUS response before it is forwarded to the NAD. Again, ISE is acting as a RADIUS proxy in both the forward and reverse direction.
- **SGA AAA Servers**—SGA AAA servers enable the network administrator to configure a list of ISE servers in your deployment in the AAA server list, allowing SGA devices to be authenticated against any of these servers. When you add ISE servers to this list, all these server details are downloaded to the SGA device. When an SGA device tries to authenticate, it would choose any ISE server from this list. If the first server is unavailable, the SGA device can choose another SGA AAA server from this list.
- **NAC Managers**—Network admission control (NAC) 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 prior to allowing users onto the network. The solution identifies whether machines are compliant with security policies and repairs vulnerabilities before permitting access to the network.
- **MDM**—Mobile device managers (MDMs) are third-party security software system that enables an administrator to configure and secure an endpoint, regardless of where the device is located. 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. ISE 1.2 allows for only one active MDM at a time. Figure 8-38 shows the MDM configuration screen.

https://t.me/learningnets

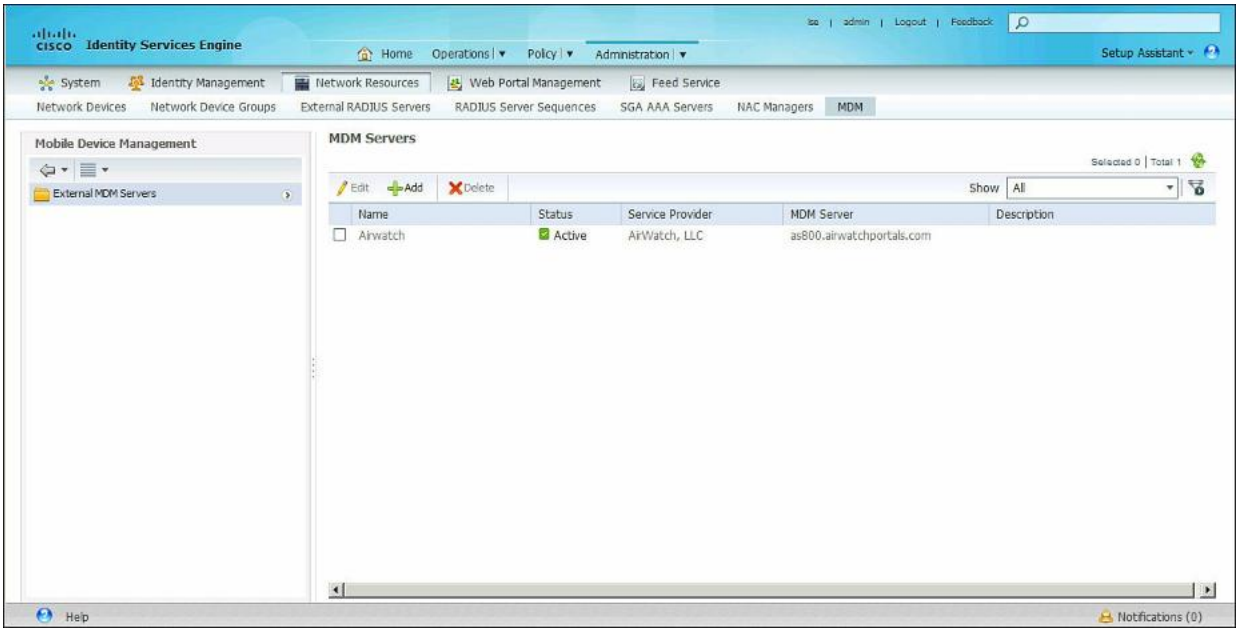


Figure 8-38 Mobile device manager configuration.

Web Portal Management

The Web Portal Management section of the Administration component is where all operational portals are configured. These portals are those that can be used to create guest accounts, allow for onboarding of endpoints, and manage any of your provisioned endpoints. The Web Portal Management enables the configuration of sponsor group policy, sponsor groups, and settings.

- **Sponsor Group Policy**—Sponsor group policy, as shown in Figure 8-39, provides the policy for network sponsors. A network sponsor is a user that can create temporary network accounts—often for guests or contractors. The sponsor group policy, like any other policy within ISE, is heavily based on IF-THEN programming construct. When defining the sponsor group policy, the policy can be based on the identity group of the authenticating sponsor and any other conditions that can be associated with the authenticating sponsor, such as the Active Directory group membership. Based on the outcome of the sponsor group policy, the sponsor will be limited accordingly as to which types of guest accounts that can be created, the length of that guest access, and other restrictions pertaining to guest account creation and access.

https://t.me/learningnets

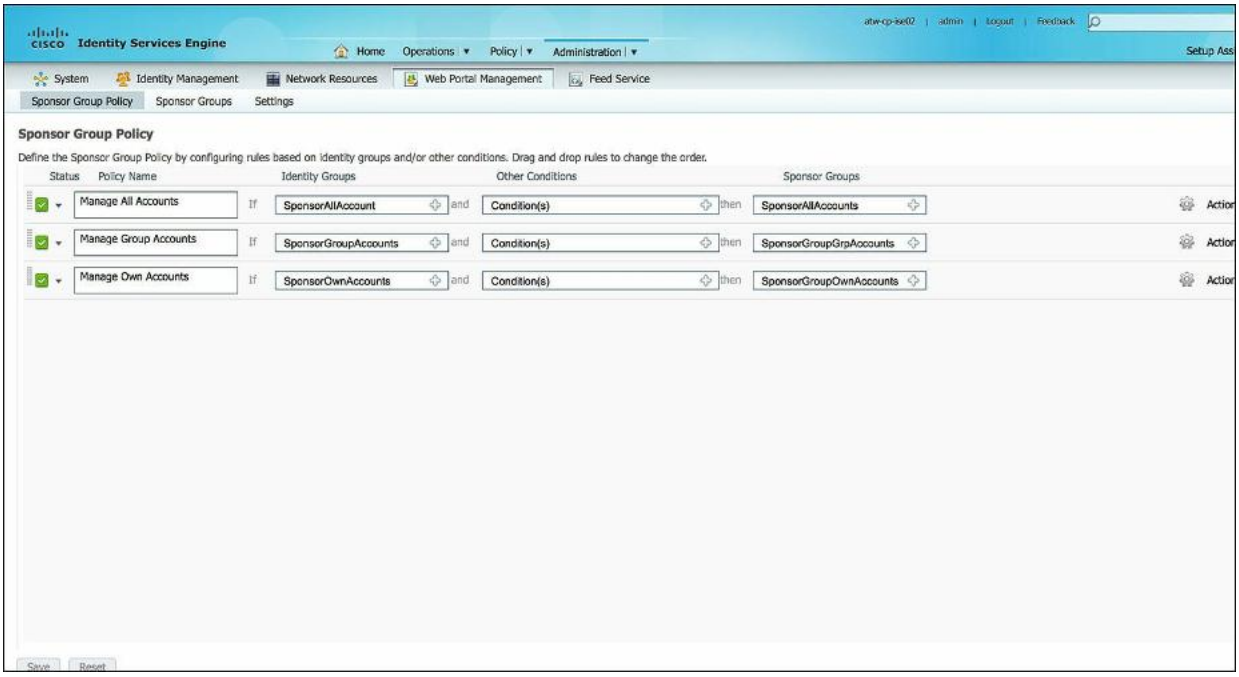


Figure 8-39 Sponsor group policy configuration.

- **Sponsor Groups**—Sponsor groups, as shown in Figure 8-40, describe the level of access that a sponsor will be given in creating guest accounts. Each sponsor group highlights the types of guest accounts that can be created, the length of access, as well as a number of other restrictions. These are the result of the sponsor group policy.

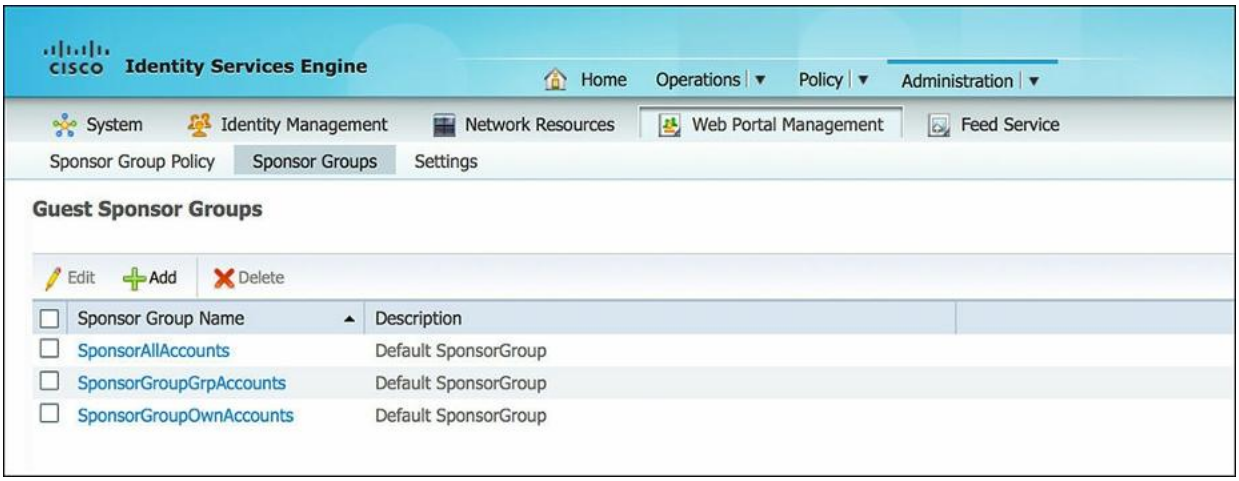


Figure 8-40 Sponsor groups configuration.

- **Settings**—The Settings subsection of Web Portal Management (see Figure 8-41) enables an administrator to customize the operational portals of ISE. These operational portals include the MyDevices portal (allowing users to manage their provisioned endpoints), guest portals (including those for guest access as well as self-provisioning operations), and the sponsor portal (where guest accounts can be created).

https://t.me/learningnets

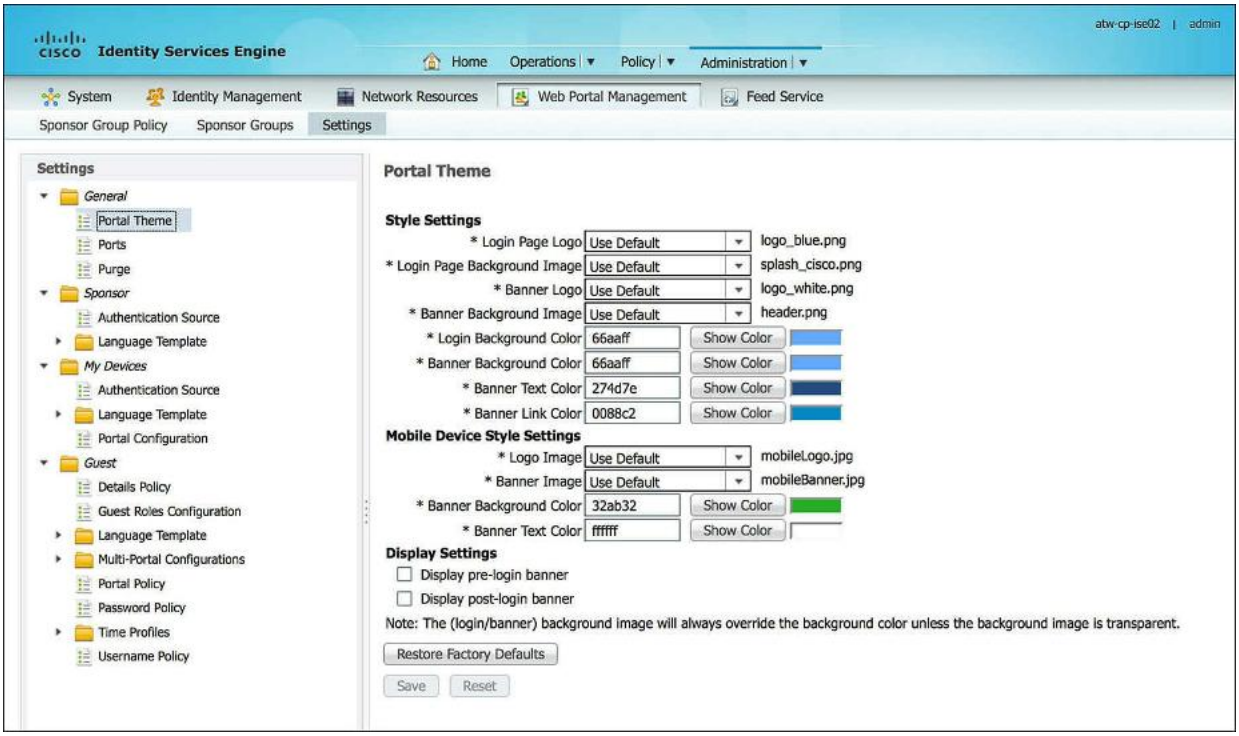


Figure 8-41 Web portal settings configuration.

Feed Service

The Feed Service provides for periodic updates to ISE's Profiler service. The Profiler service is the capability of ISE to identify the specifics of endpoints that are accessing the network. As was discussed in Chapter 6, "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. (We discuss this more in Chapter 15, "Profiling.") As endpoint vendors develop new endpoints, ISE can use the Feed Service (as shown in Figure 8-42) to download the updated database of the requisite criteria (MAC Address, DHCP parameters, and so on) that is needed to identify the new endpoint.

https://t.me/learningnets

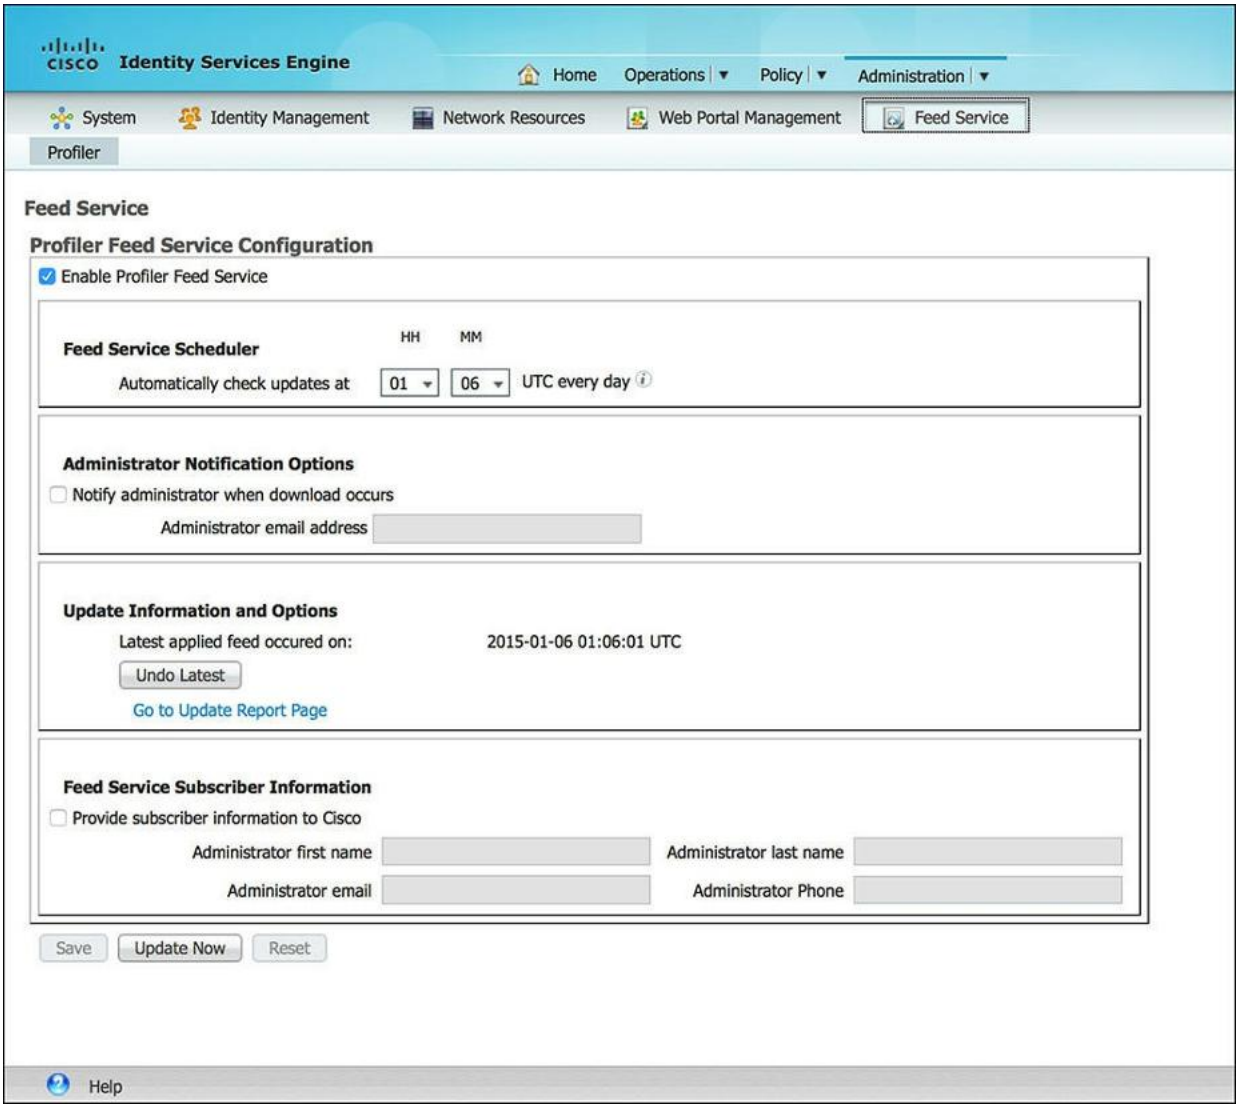


Figure 8-42 Feed Service.

Type 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 policies. Each of these policies serves a purpose to limit or appropriately manage the level of access given to these endpoints. Many of these policies take on the programmatic construct referred to as an IF-THEN statement.

Authentication

Authentication is the capability of ISE to identify the endpoint or the user of the endpoint as it connects to the network. By determining the identity of the endpoint or user, the network administrator can make additional decisions as to which level of access should be given to the endpoint or user. We discuss authentication policy in more detail in Chapter 10.

https://t.me/learningnets

Authorization

After ISE has determined the identity of the user or endpoint, ISE can use this identity to determine to what the endpoint will be given access. The granting of, or even the denial of, access is referred to as authorization. This authorization policy enables a network administrator to make the final determination based on any number of criteria—information provided via authentication, by the endpoint, or by the NAD from where the endpoint is accessing the network. We discuss authorization policy in more detail in Chapter 11.

Profiling

Profiling is the capability of ISE to determine the type of device that is accessing the network. This 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 (that is, DHCP, HTTP, and so on). We discuss profiling in more detail in Chapter 15.

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, antimalware, antispyware, and OS patches), which is 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. We discuss posturing in more detail in Chapter 19, "Posture Assessment."

Client Provisioning

Client provisioning is a process whereby all necessary credentials and configurations are deployed to the endpoint, allowing the endpoint to more easily, maybe even automatically, join the network in the future. Provisioning is sometimes referred to as onboarding. This onboarding process might defer depending 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 the decision tree as to which credentials and configurations are deployed to the endpoint in these scenarios. Client provisioning is also instrumental in deploying NAC agents to endpoints during posturing. Client provisioning is discussed in greater detail in several chapters as we move forward.

Security Group Access

Security group access (SGA) is a security solution whereby a security group tag (SGT) is assigned to an endpoint following authentication. The SGT plan is often designed to subdivide and easily distinguish different user scenarios—for instance, full access versus partial access, employee versus guest, and engineering versus marketing. Every packet that has the same SGT is given the same level of access on the network. The NAD inserts this SGT into each packet sent by the endpoint, allowing each SGA-enabled upstream device to know which "group" of users or endpoints was responsible for sending that packet and to enforce the appropriate security policy.

The SGA policy within ISE determines the appropriate SGT to assign to an endpoint based on the user of the endpoint, the type of endpoint, the method of access (wired versus wireless), as well as any other criteria that can be determined from the authentication process. SGA is discussed in detail in Chapter 18, "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 8-6 lists a reference of these key topics and the page numbers on which each is found.

https://t.me/learningnets



Key Topic Element	Description	Page
Paragraph	Cisco ISE X.509 certificate warning	155
Paragraph	Choosing an appropriate logging level	181

Table 8-6 Key Topics for Chapter 8

Define Key Terms

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

command-line interface (CLI)

graphical user interface (GUI)

https://t.me/learningnets

Chapter 9. Initial Configuration of Cisco ISE

This chapter covers the following exam 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 (ISE) appliance, you can take it out of the box and power on the appliance and it will be ready to control access to the network, right? Wrong. Of course not. Cisco ISE is a very powerful policy engine, and as such it requires a fair amount of preparation to teach it how to communicate in the company's environment.

We are not simply talking about the standard requirement of "this device requires a valid IP address before it can communicate on a network." Although that is absolutely a requirement for Cisco ISE, there is so much more.

Cisco ISE will need to be configured for working domain name resolution (DNS), for correct time synchronization (NTP), and to have a valid and resolvable fully qualified domain name (FQDN) that any network-attached device might use to resolve the address. Cisco ISE will need a certificate installed that has been signed by the company's root certificate authority or a public certificate authority. It will also need to be configured with the addresses and secret keys of all the network devices that will use it—and most likely be joined to an Active Directory (AD) infrastructure to use as one of the identity sources.

The blueprint for this exam focuses mainly on the state of Cisco ISE after the administrative graphical user interface (GUI) is running and available. However, critical steps during the initial installation can make the overall deployment successful or make it more difficult. Therefore, this chapter focuses on all of it.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Describe Native AD and LDAP	1
Describe Identity Store Options	1
Network Access Devices	5-6
Implement Guest Services	4
Troubleshoot Using Cisco ISE Diagnostic Tools	2
Identity Management	3-4, 7
Identity Management	3, 8

Table 9-1 "Do I Know This Already?" Section-to-Question Mapping

https://t.me/learningnets

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 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. Which CLI command lists all the ISE processes and their statuses?

   a. `show status ise`

   b. `show application status ise`

   c. `show application status`

   d. `show version`

3. Which two functions does a certificate fulfill when used with HTTPS and EAPoverLAN?

   a. Authenticates the server to the client, and the encryption method is embedded in the transform-set field within the certificate.

   b. Identifies the client to the NAD and is used as the basis for the encrypted transport between the client and the NAD.

   c. Authenticates the server to the client and is used as the basis for the encrypted transport between the client and server.

   d. Authenticates the client to the NAD, and the encryption method is embedded in the transform-set field within the certificate.

4. True or False? When submitting a certificate signing request (CSR), the CSR and the private key are sent to the signing certificate authority (CA), so the CA can sign the key-pair.

   a. True

   b. False

5. True or False? Settings such as RADIUS shared secret keys and SNMP strings can be set on a per Network Device Group (NDG) level.

   a. True

   b. False

6. What is a valid use of network device groups?

   a. Use NDG as the condition by which to build different policy sets for the staged deployment of ISE.

   b. Use the incoming authentication protocol type to route the authentication to a network device group that is able to process that authentication type.

   c. Use the NDG to determine to which ISE policy node to route the authentication request.

   d. The result of an authorization policy will allow the user to log in and control devices within the

https://t.me/learningnets

assigned network device group.

7. True or False? Local endpoint identity groups should be created per endpoint profile instead of using the attribute itself.
   a. True
   b. False
8. True or False? Cisco ISE 1.2 can join 1 Active Directory Forest and process authentications for any domain in the forest with 2-way trusts.
   a. True
   b. False
9. What is the purpose of a certificate authentication profile (CAP)?
   a. Defines which CA to use for revocation checking via either certificate revocation lists (CRLs) or online certificate status protocol (OCSP).
   b. Used with MSCHAPv2 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 principle identity for the authorization process.
   d. Used with EAP-FAST to allow for faster reauthentications and secure transport without the use of X.509 certificates.
10. True or False? It is critical to use Network Time Protocol (NTP) to ensure the time is synchronized correctly between Cisco ISE and Microsoft Active Directory.
    a. True
    b. False

Foundation Topics

Cisco Identity Services Engine Form Factors

Before the Cisco Identity Services Engine can be configured, it must first be installed. The installation of Cisco ISE is beyond the scope of the exam blueprint, and therefore is behind the scope of this book. There are two form factors for Cisco ISE. There is a physical appliance called the Cisco Secure Network Server (SNS). The SNS appliance is a multipurpose appliance and is used for not only Cisco ISE, but also Cisco Secure Access Control Server (ACS) and Cisco Network Admission Control (NAC) Appliance.

Cisco ISE is also available as a virtual appliance. It is important to recognize that a virtual appliance is different from a traditional virtual machine. How is it different? 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. That means the hardware cannot be oversubscribed and reservations for the memory and CPU must be configured at equal values to the physical SNS appliances.

Why is this so important? The authors of this book have both been involved in a number of escalation cases where the root of the problem that the customer ran into was the misconfiguration of the VMware virtual machine(s) itself. The main misconfigurations are no resource reservations were set for the memory and CPU; the VM storage was much slower than the physical drives of the SNS appliance; and delayed writes causing database corruption. Please use that as a lesson to ensure the virtual appliances are always configured with the correct reservations and storage in place.

For instructions on installation and configuration, see the Cisco Identity Services Engine Installation guides at http://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-guides-list.html.

https://t.me/learningnets

Bootstrapping Cisco ISE

As stated in the introduction, Cisco ISE is a powerful policy engine, and as such, it requires a fair amount of preparation to teach it how to communicate in the company's environment.

The SNS Appliance will come with Cisco ISE preinstalled and waiting at the setup screen (as shown in Figure 9-2). For the virtual appliance and for any time you are installing from a DVD (or .iso file), you will see the installation choices menu displayed in Figure 9-1.

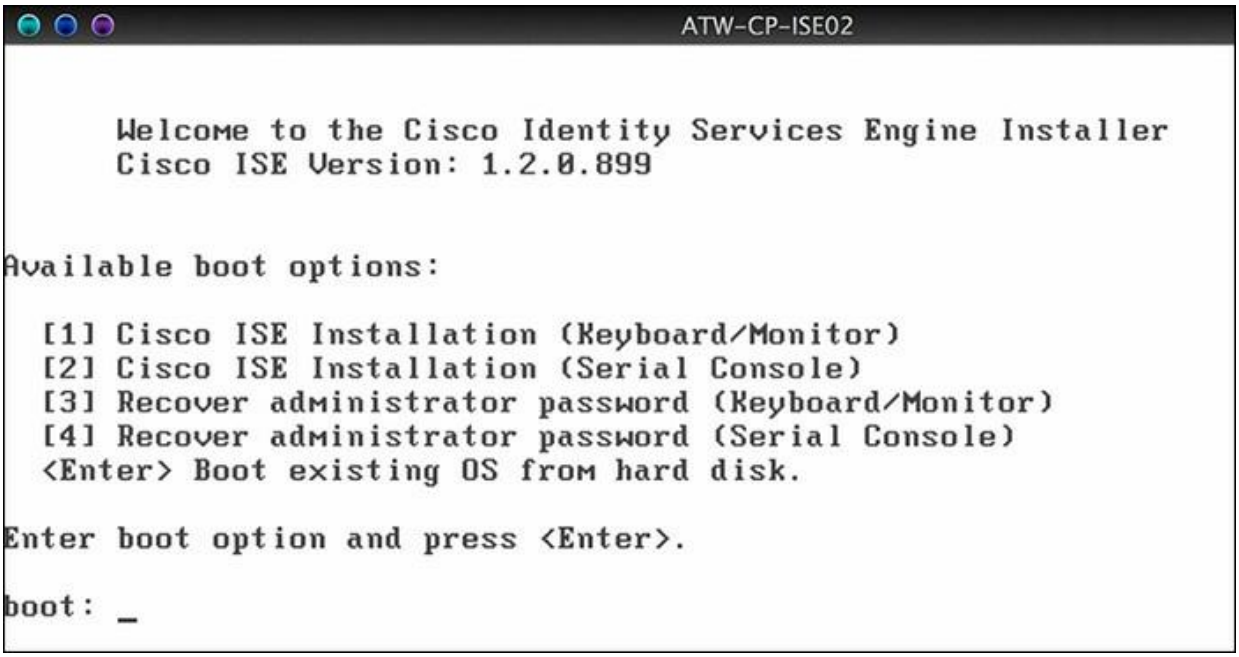


Figure 9-1 Installation choices.

The underlying operating system (OS) that will be installed is a customized Linux called the Cisco Application Development Environment Operating System (ADE-OS). After the OS has completed installing, the appliance will restart and boot into the setup prompt, as shown in Figure 9-2.

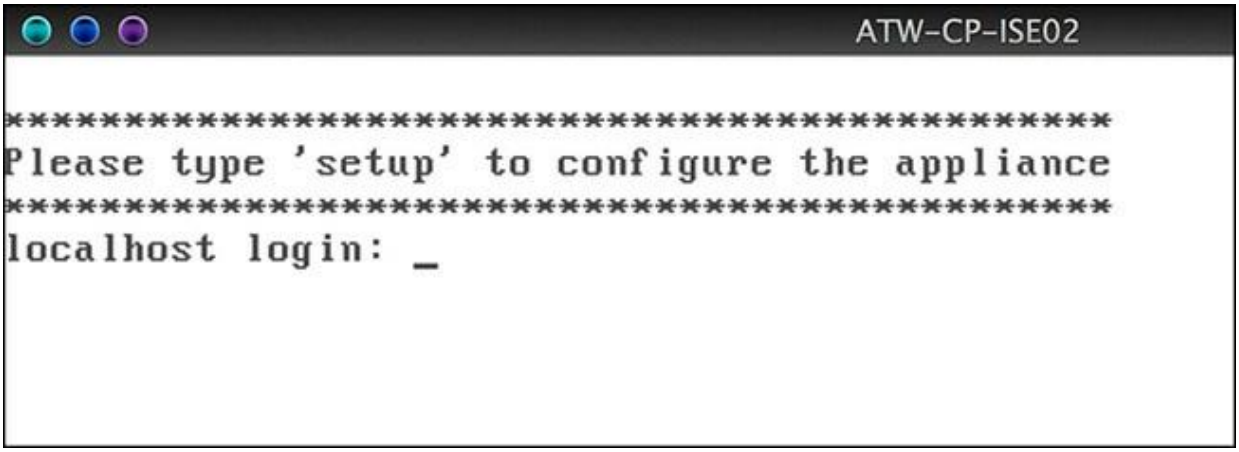


Figure 9-2 Setup prompt.

Typing **`setup`** will launch the command-line setup wizard. Following the prompts, enter all the bootstrapping information the setup wizard asks for, as shown in Figure 9-3.

https://t.me/learningnets

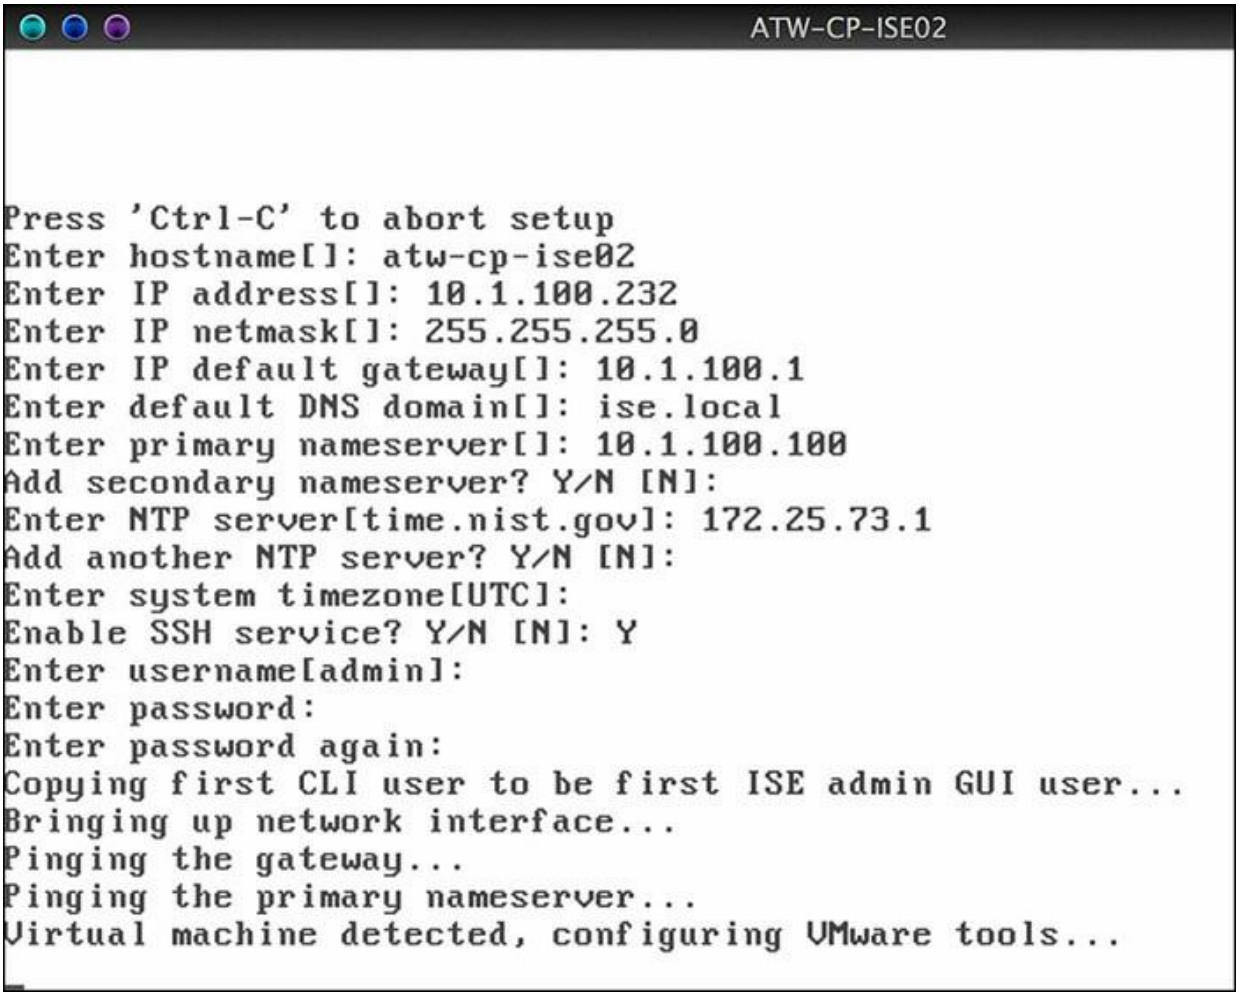


Figure 9-3 CLI Setup Wizard.

After the settings have been validated, the ISE application itself is installed. This will take quite a bit of time, so at this point you should go drink a few gallons of coffee. After the ISE application is fully installed, the device will reboot. You can 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 their statuses, as shown in Figure 9-4. When the application service is running, ISE is ready for you to connect to the administrative GUI.

https://t.me/learningnets

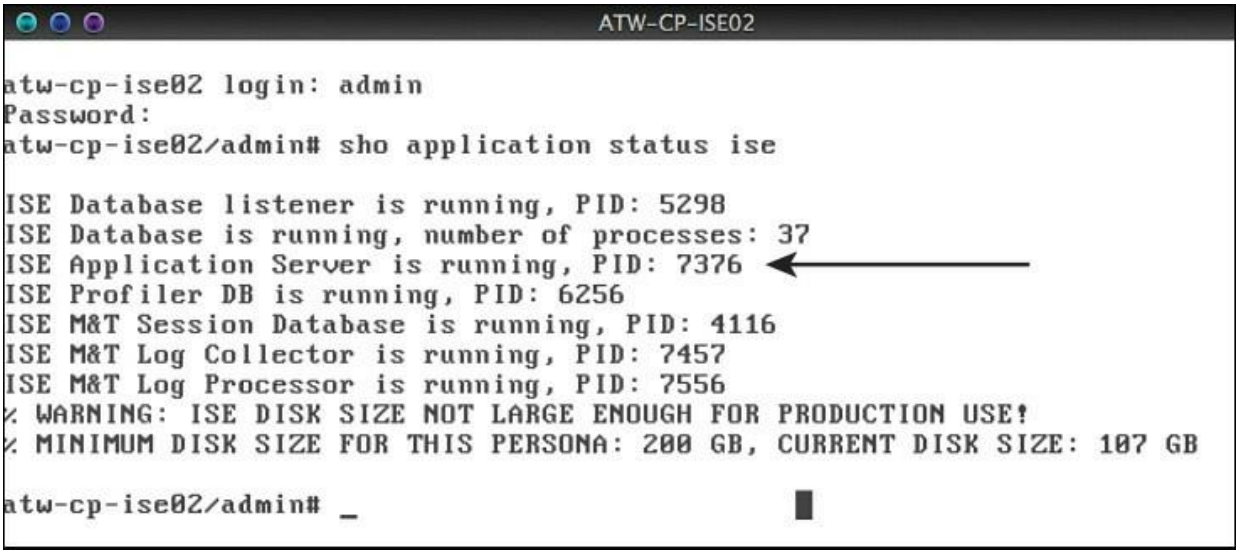


Figure 9-4 Output of the `show application status ise` command.

You connect to the Cisco ISE GUI using secure HTTP (HTTPS), using the following URL: https://<ise-fqdn>/admin/. Upon connecting to the ISE GUI the first time, your web browser will most likely present an error about the certificate being unknown or untrusted, as shown in Figure 9-5. This error message is due to the self-signed certificate that is created during installation.

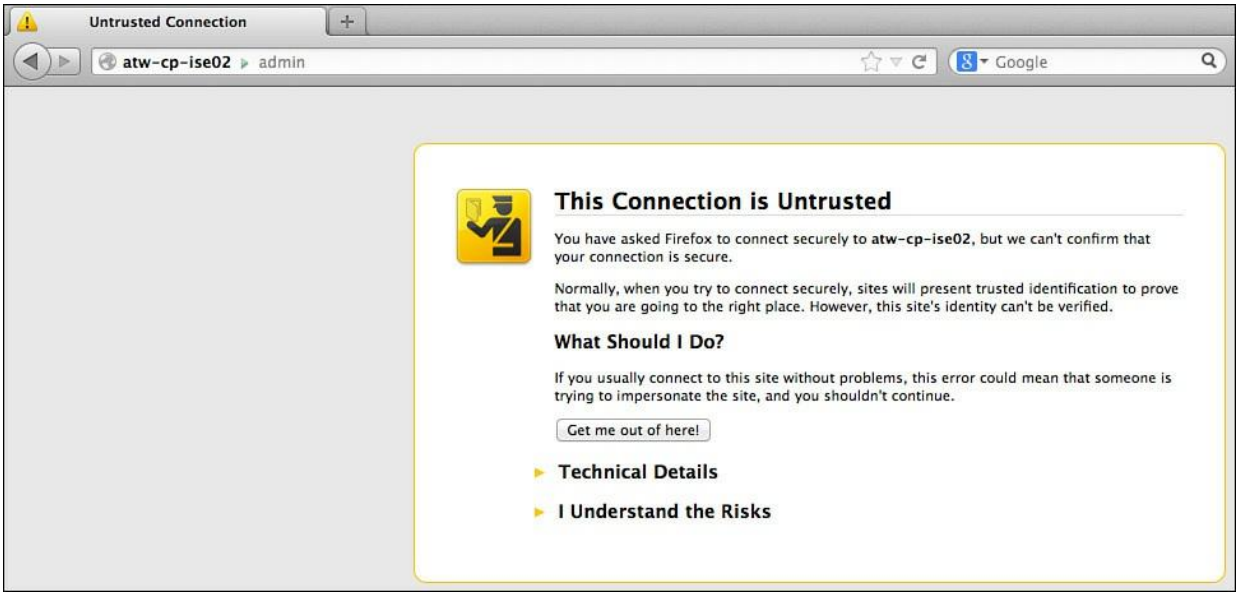


Figure 9-5 Untrusted certificate error.

Accept and trust this certificate to gain access to the administrative GUI login, as shown in Figure 9-6. Enter the username and password entered during the installation setup wizard.

https://t.me/learningnets

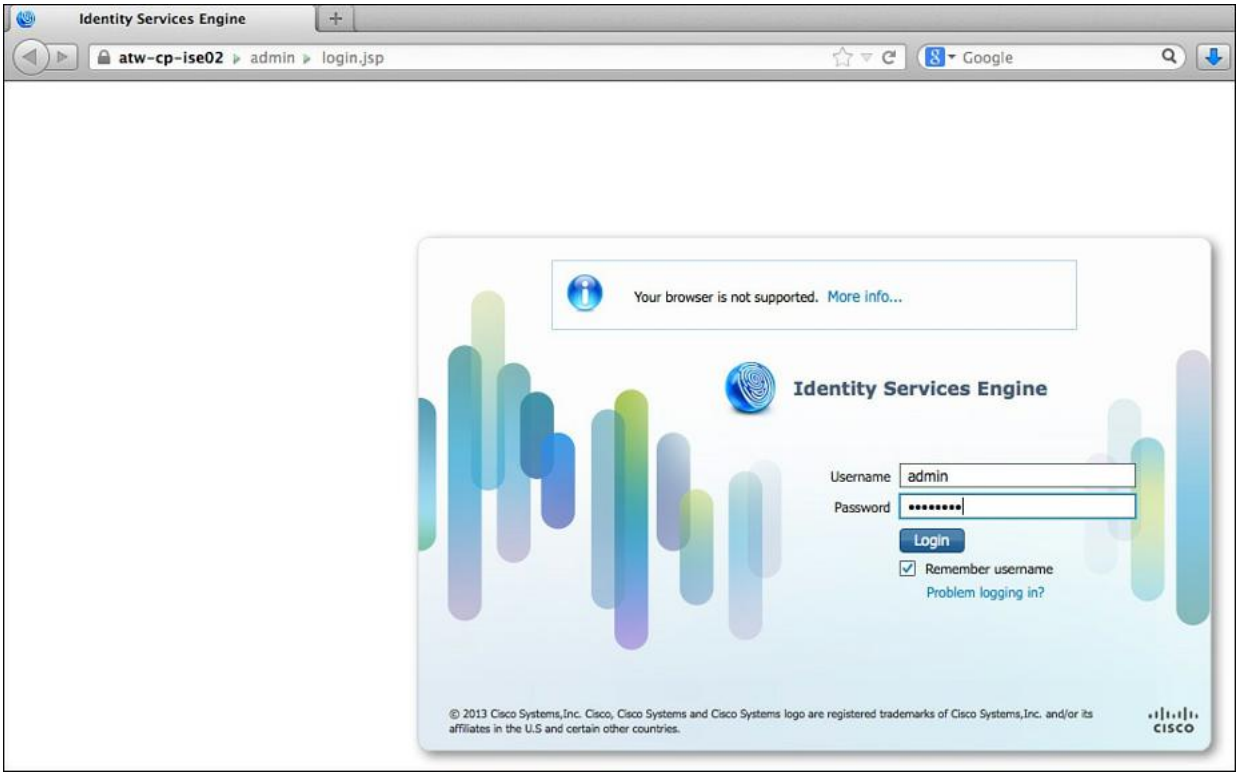


Figure 9-6 Administrative login.

Where Are Certificates Used with the Cisco Identity Services Engine?

When you connect to a website secured with Secure Sockets Layer (SSL, a.k.a.: HTTPS), the website will use an x.509 certificate to establish that secure communication. The certificate is used for more than one function:

1. To prove "identity" of the website
2. To serve as the basis of encryption between the client (web browser) and the website

Figure 9-7 illustrates the process of securing a website. Numerous web portals are hosted on ISE, including the administrative portal, which was displayed in Figure 9-6. In addition, there are many end-user–facing portals for web authentication, endpoint management, and other uses. All portals on ISE require encryption by default. In other words, all portals in ISE will use HTTPS, and a certificate on ISE will be used to secure that communication.

https://t.me/learningnets

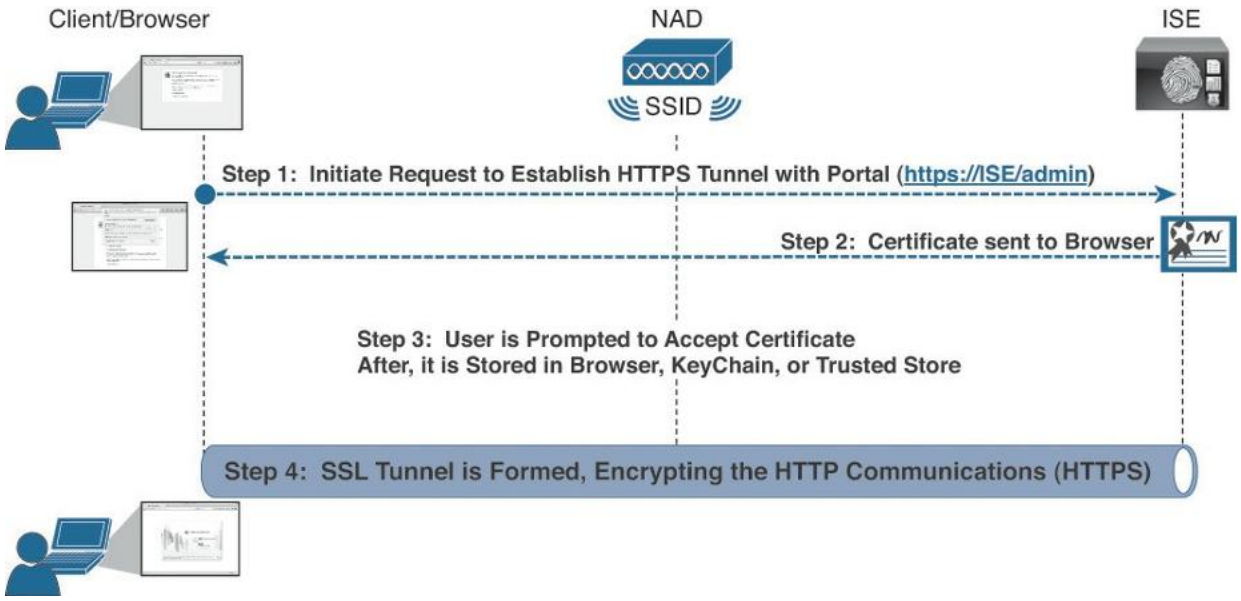


Figure 9-7 Securing Web communications with SSL.

Certificates are also used to secure Extensible Authentication Protocol communications, as described in great detail in Chapter 4, "EAP over LAN (also Known as 802.1X)." Much like the SSL communication with websites, certificates are used to form a Transport Layer Security (TLS) tunnel between the endpoint and the RADIUS server, within which EAP authentication will occur, as shown in Figure 9-8. Much like the secure-HTTPS example, the x.509 certificate has multiple functions with EAP as well. It is used to identify the RADIUS server to the client (supplicant) as well as form the basis of the encrypted communication.

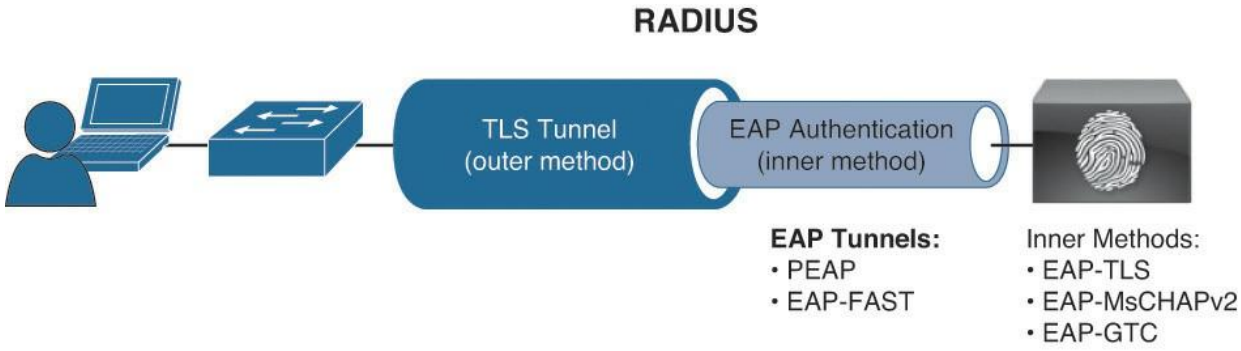


Figure 9-8 Securing EAP communications with TLS.

Self-Signed Certificates



When initially installing ISE, the application will generate a self-signed certificate, meaning no certificate authority (CA) was used to sign the certificate and vouch that it should indeed be trusted. By default, many browsers and operating systems alike come with a number of preinstalled CA certificates from the more popular CAs. Because a well-known and well-trusted CA was not used to sign this self-signed certificate, clients (browsers and supplicants alike) will not trust the certificate and warning messages similar to the message shown in Figure 9-5 will be displayed.

https://t.me/learningnets

> **Note**
>
> To 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.

Numerous complexities are involved with running a production deployment with self-signed certificates. The first that will present itself is that each ISE node will need to trust each other's ISE node (a problem that will increase itself exponentially as the deployment grows). For this and many other reasons, it is always recommended to use signed certificates. This does not mean the certificate must be signed by a public trusted root; a company's own internal CA would suffice nicely as well.

CA-Signed Certificates



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 could simply be imported into whatever device you want to use those keys. This is not normally what happens when first installing 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 9-9.

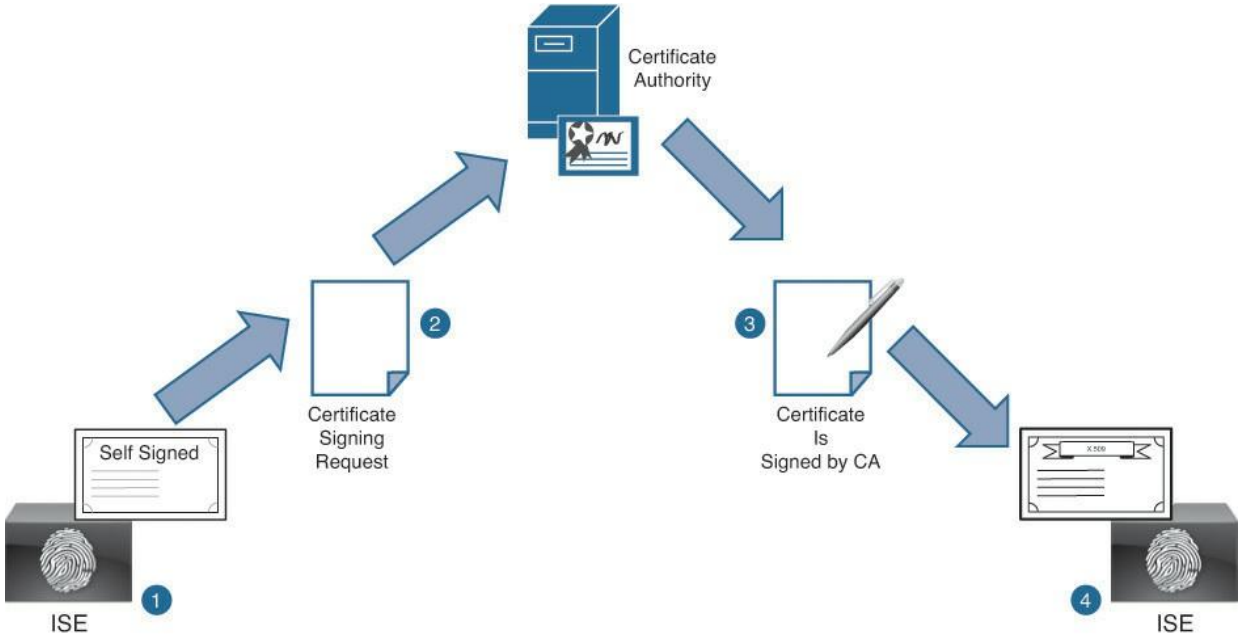


Figure 9-9 Stages of certificate signing overview.

For ISE, all certificate management is handled within the ISE GUI under **Administration > System > Certificates**. The initial screen displays the local certificates, which displays the default self-signed certificate, as shown in Figure 9-10.

https://t.me/learningnets

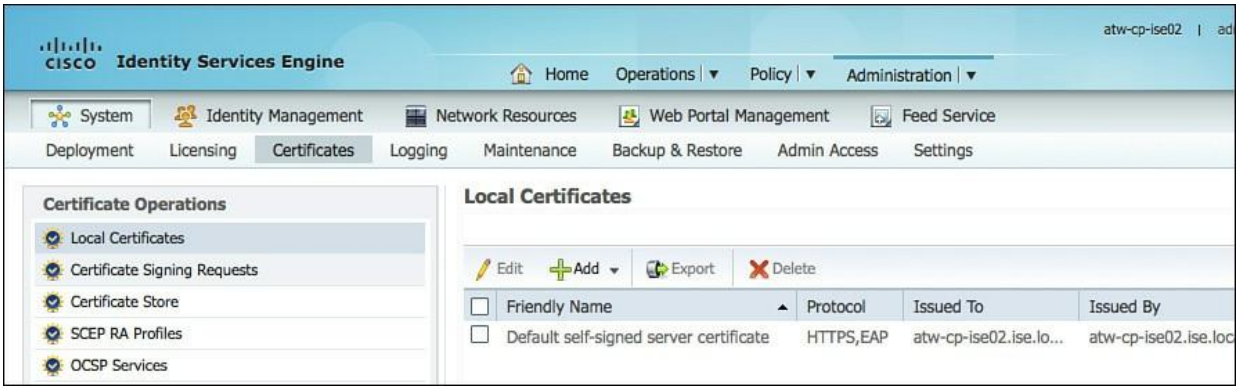


Figure 9-10 Administration > System > Certificates > Local Certificates.

Initiate the generation of a CSR by clicking **Add > Certificate Signing Request**, as shown in Figure 9-11.



Figure 9-11 Local Certificates > Add > Generate Certificate Signing Request.

Fill in the Certificate Subject field for the CSR after the Generate Certificate Signing Request page is displayed in the main window, as shown in Figure 9-12. A number of fields can make up the certificate subject and should be separated by commas. These include

- **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 an IP address of the ISE node. The standard would be to use the hostname and domain name configured during the initial setup of ISE (ex: atw-cp-ise02.ise.local).

> **Note**
>
> Pay close attention to capitalization because it will matter when binding the signed certificate.

- **Organization (O)**—This should be the organization the certificate is meant to represent; in other words, the name of the company that is implementing ISE.
- **Organizational Unit (OU)**—This field should represent the division of the company that is controlling the device for which the certificate is being used.
- **Locality (L)**—Another field for defining the entity, it is often used to denote the city where the entity resides.
- **State (ST)**—The state or province in which the entity resides.
- **Country (C)**—The country in which the entity resides.

https://t.me/learningnets

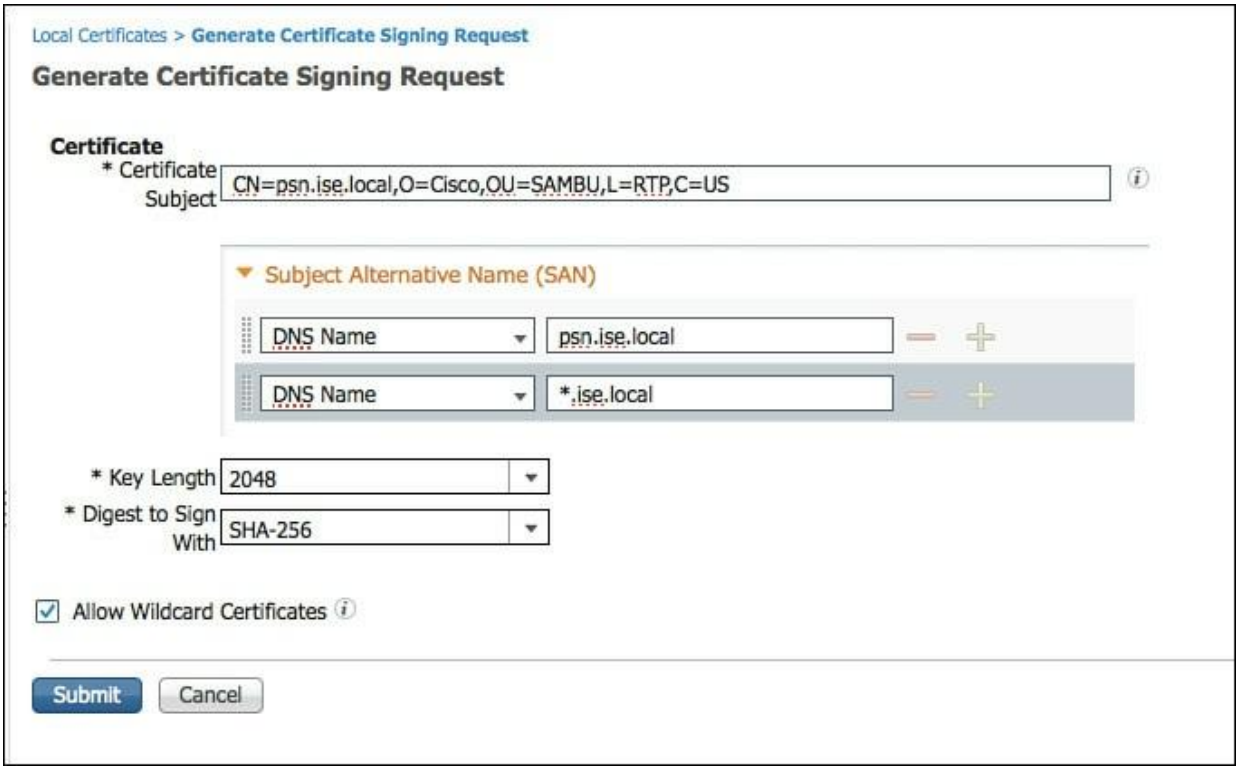


Figure 9-12 CSR form.

There are many more optional fields, or attributes. The only required one is the Common Name (CN). Another section of certificate attributes that the ISE GUI enables customization of is the Subject Alternative Name (SAN) field. The SAN field is meant to carry a number of attributes, each designed to aid in the identification of the certificate.



When presented with a certificate, a client will compare 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, it will cause a certificate name mismatch. For example, if a secure website is hosted at https://www.ciscopress.com and the subject of that certificate has a CN=www.ciscopress.com, the name and the CN will match. However, the FQDN www.ciscopress.com resolves to an IP address of 159.182.165.54. That same address also might resolve to the FQDN of www.pearsonpublishing.com. If you were to enter the URL of https://159.182.165.54 or https://www.pearsonpublishing.com into the web browser, a certificate mismatch error would occur.

This brings the conversation to one of wildcard certificates. A wildcard certificate is one that uses a wildcard notation (an asterisk and a period before the domain name) and enables the certificate to be shared across multiple hosts in an organization. A sample CN value for a wildcard certificate's Subject Name would look like the following:

*wolandheffner.com

If you configure a wildcard certificate to use *wolandheffner.com, that same certificate can be used to secure any host whose DNS name ends in "wolandheffner.com", such as:

aaa.wolandheffner.com

psn.wolandheffner.com

mydevices.wolandheffner.com

sponsor.wolandheffner.com

https://t.me/learningnets

Although a wildcard certificate would solve the issue for a different hostname in the CN, it would still require the domain name to 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 will not trust wildcard certificates for EAP communications.

As an alternative to wildcard certificates, focus can shift to the Subject Alternative Name (SAN) field, which was 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) defines the requirement for the FQDN from the CN field to also be listed as a DNSName in the SAN field.

As shown in Figure 9-12, a wildcard value also can be added as a DNSName entry in the SAN field. This provides a best of both worlds approach for matching multiple values. It enables the use of the same private and public key-pair to be installed across all the ISE nodes, which enables a much smoother user experience with mobile devices.

After clicking **Submit**, a public and private key pair is created. The private key will be kept secret and never distributed to any third party. The public key and the X.509 information are then combined to form the CSR. This CSR is stored under **Administration > System > Certificates > Certificate Signing Requests**, as shown in Figure 9-13.

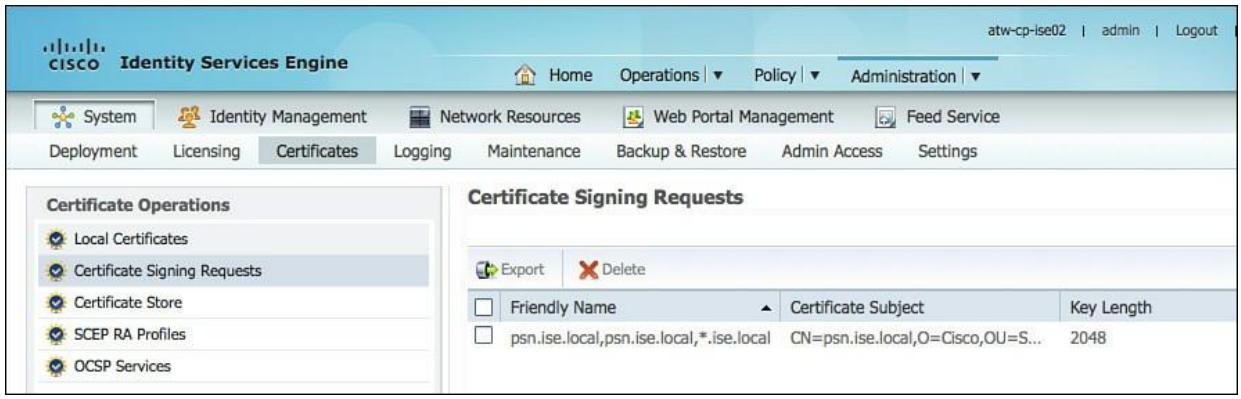


Figure 9-13 Administration > System > Certificates > Certificate Signing Requests.

From here, you will select the certificate signing request and click **Export**. This will prompt you to save the CSR to your local drive, as shown in Figure 9-14.

https://t.me/learningnets

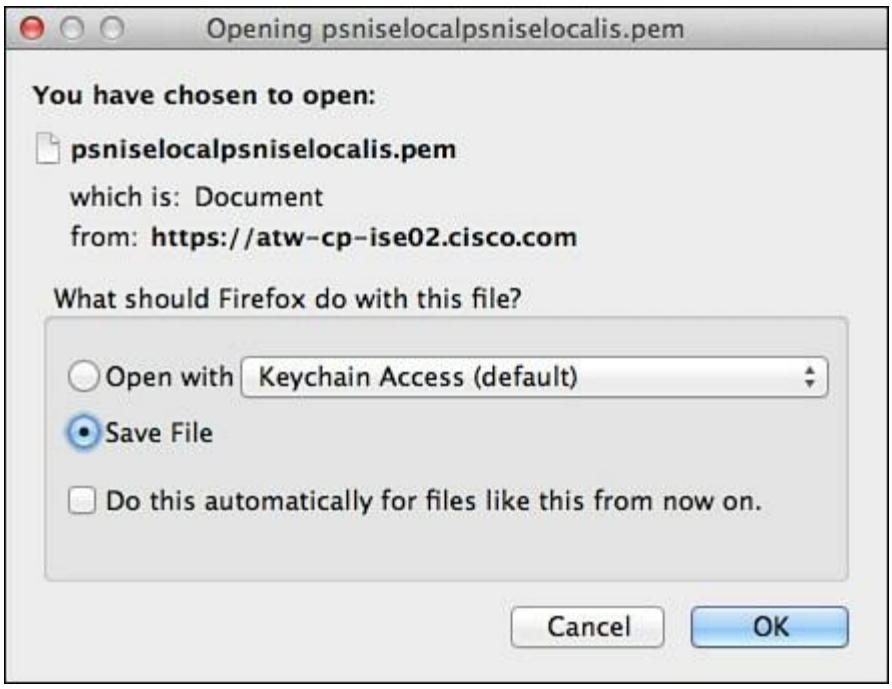


Figure 9-14 Downloading the CSR.

The CSR will be a Base-64 encoded, plain-text format file, known as a PEM file. This can be opened in a standard text editor, as shown in Figure 9-15. Copy all the text including the "-----BEGIN CERTIFICATE REQUEST-----" and the "-----END CERTIFICATE REQUEST-----", but do not include any extra spaces (if they exist).



Figure 9-15 Certificate authority.

Open a web browser and navigate to the CA. In this chapter we are using the CA built in to Active Directory—for example, http://atw-cp-ad.ise.local/certsrv, which is shown in Figure 9-15. If you are not using a CA server, a different process might be required.

https://t.me/learningnets

Next, click **Request a Certificate**, as shown in Figure 9-16.

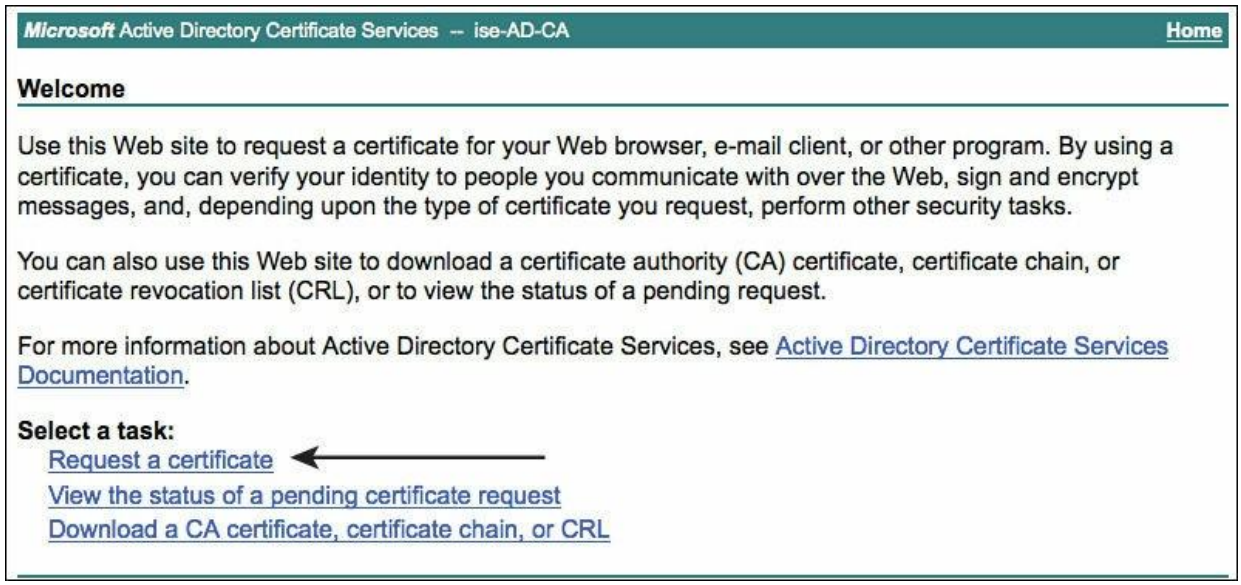


Figure 9-16 Request a certificate.

Select the **Advanced Certificate Request**, as shown in Figure 9-17.

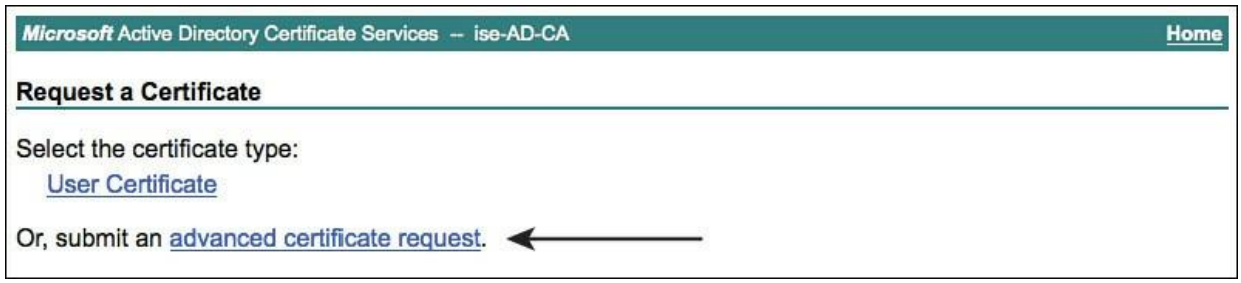


Figure 9-17 Advanced certificate request.

Paste the contents of the CSR into the form, ensuring there are no extra spaces. Select the Web Server certificate template from the drop–down list, and click **Submit**, as shown in Figure 9-18.

https://t.me/learningnets

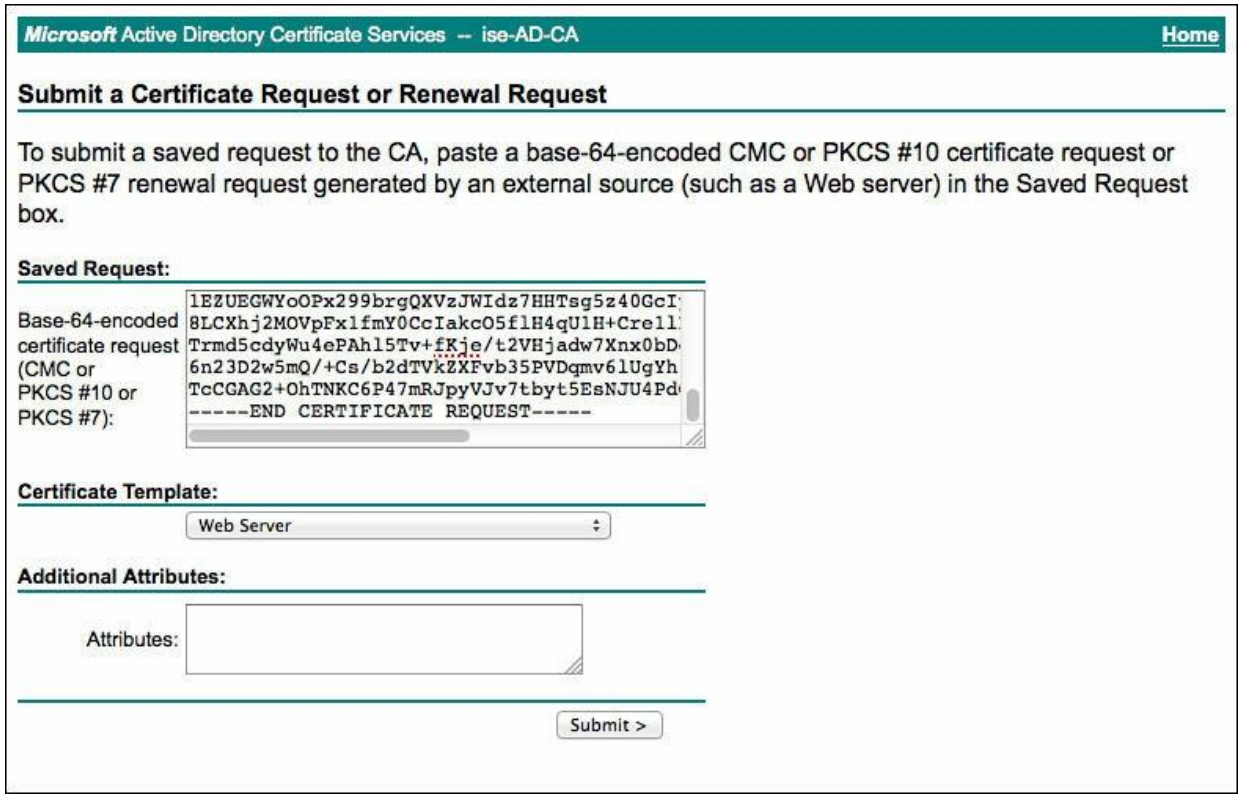


Figure 9-18 Submit the certificate signing request.

Select the **Base 64 encoded** option. This is the PEM format. Then click **Download Certificate**, as shown in Figure 9-19.

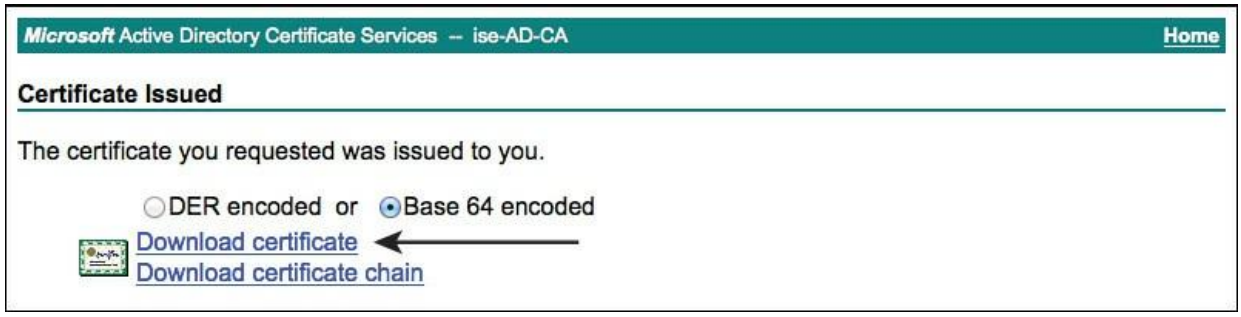


Figure 9-19 Download the signed certificate.

> **Note**
>
> Do not download the certificate chain. The best practice from Cisco is to import all certificates separately, instead of using certificate chains (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, and click **Download a CA Certificate, Certificate Chain, or CRL**. This is where you will download the public certificate of the CA itself.

Select **Base 64** and click **Download CA Certificate**, as shown in Figure 9-20.

https://t.me/learningnets

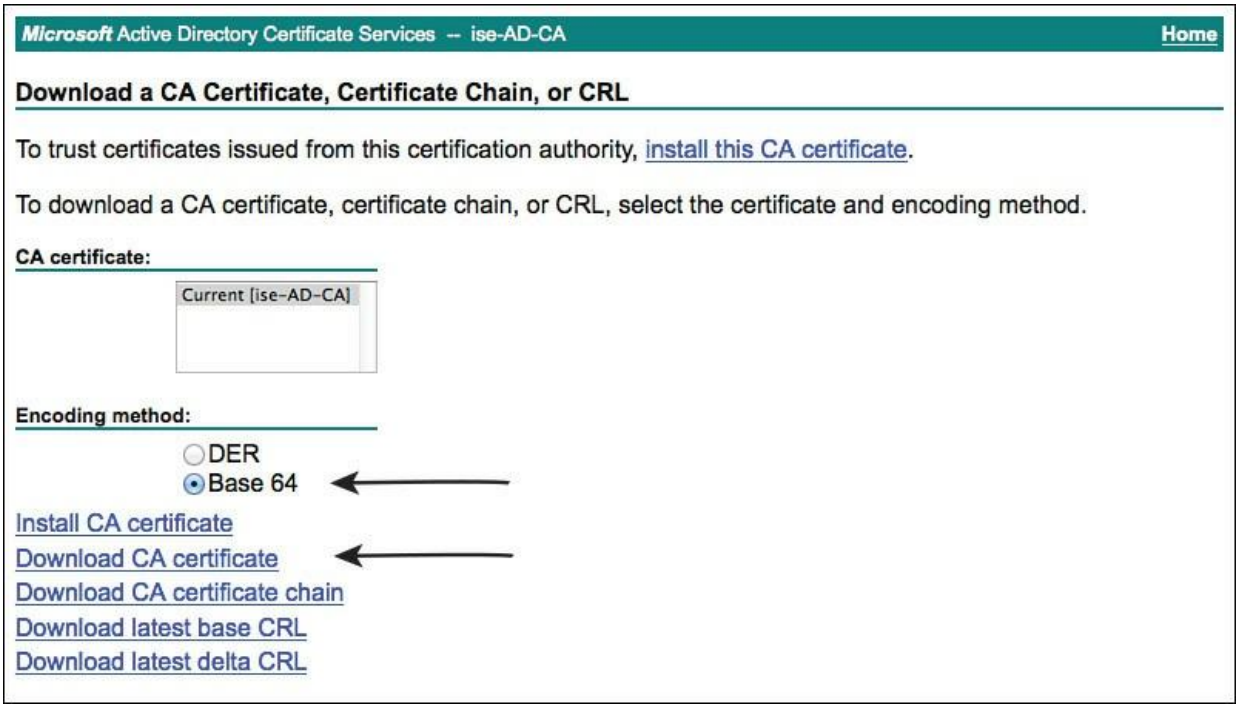


Figure 9-20 Download the CA's public certificate.

Now you will import the CA's certificate into the ISE trusted store by navigating to **Administration > System > Certificates > Certificate Store** in the ISE GUI, as shown in Figure 9-21.

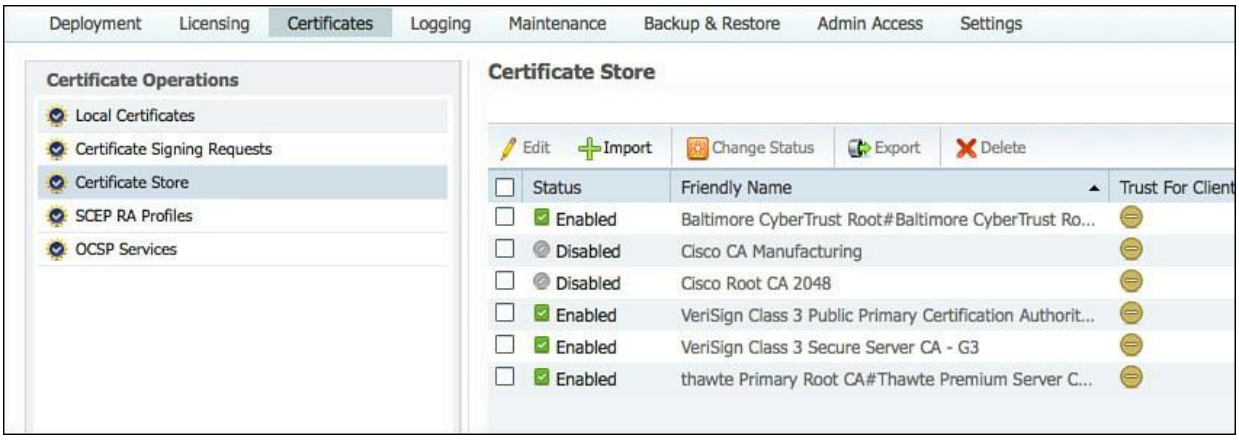


Figure 9-21 Administration > System > Certificates > Certificate Store.



This is the location of all certificates that ISE will trust, for which functions certificates from that CA are trusted, and where the revocation checking is defined.

Click **Browse** and select the CA's certificate. Provide a friendly name to help identify the certificate in the list, and click **Trust for Client Authentication or Secure Syslog Service**. Then click **Submit**, as shown in Figure 9-22.

https://t.me/learningnets



Figure 9-22 Importing the CA's certificate to be trusted.

Now that ISE will trust certificates signed by the Active Directory CA, it is time to bind the CA signed identity certificate to the CSR for ISE to now use the certificate. Navigate to **Administration > System > Certificates > Local Certificates.** Click **Add > Bind CA Signed Certificate**.

Click **Browse** and select the signed identity certificate. Provide a friendly name, and ensure the **Allow Wildcard Certificates** check box is checked if you provided a CN that is not the exact same as the ISE node's FQDN. Select both the **EAP** and **HTTPS** check boxes for the use of this certificate, and click **Submit**. Figure 9-23 shows the Bind CA Signed Certificate screen.



Figure 9-23 Bind CA signed certificate.

The ISE node will now restart all services to use the newly applied certificate. Upon your next login, you also can delete the CSR and the original self-signed certificate because they are no longer needed. When you connect after the reboot, you will also notice that the certificate used to protect and identify the ISE admin portal is now the signed certificate. Figure 9-24 illustrates the identity certificate in use.

https://t.me/learningnets



Figure 9-24 The signed certificate in use.

There is one more step that should be completed for various reasons. The public and private key-pair should be exported and stored in a secure location. Why? This is a recommended step to ensure that the same certificates can be used on any additional ISE nodes and to ensure there is a safe backup in case of an emergency that requires the ISE node to be rebuilt.

From the ISE GUI, navigate to **Administration > System > Certificates > Local Certificates**. Select the signed-certificate, and click **Export**. Select **Export Certificate and Private Key**, and provide a secure password that can be used to unencrypt the keys for future use, as shown in Figure 9-25.

https://t.me/learningnets



Figure 9-25 Exporting the private/public key pair.

This completes the certificate portion of the ISE installation and bootstrapping.

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 actual access control to the endpoint.

This device, or more specifically the device's capabilities, is a critical part of any secure network access strategy. The more capable the network access devices (switch, wireless controller, and VPN concentrator) are, the more flexibility and power the ISE admin will have to accomplish the business goals.

Network Device Groups

Before adding network access devices to ISE, it is highly recommended to create some logical network device groups (NDGs), located under **Administration > Network Resources > Network Device Groups**. These groups can be a powerful tool when used correctly. With the use of NDGs, policy can be created based on the type of network device, its location, or any other logical grouping that an organization might want.

ISE is prebuilt with two top-level groups named All Device Types and All Locations. An ISE administrator can create other top-level groups as required. Examples of other top-level groups could be Line of Business, Deployment Stage, and so on—basically anything that would be an organizational unit for your business or deployment structure.

Figure 9-26 displays an example NDG hierarchy.

https://t.me/learningnets

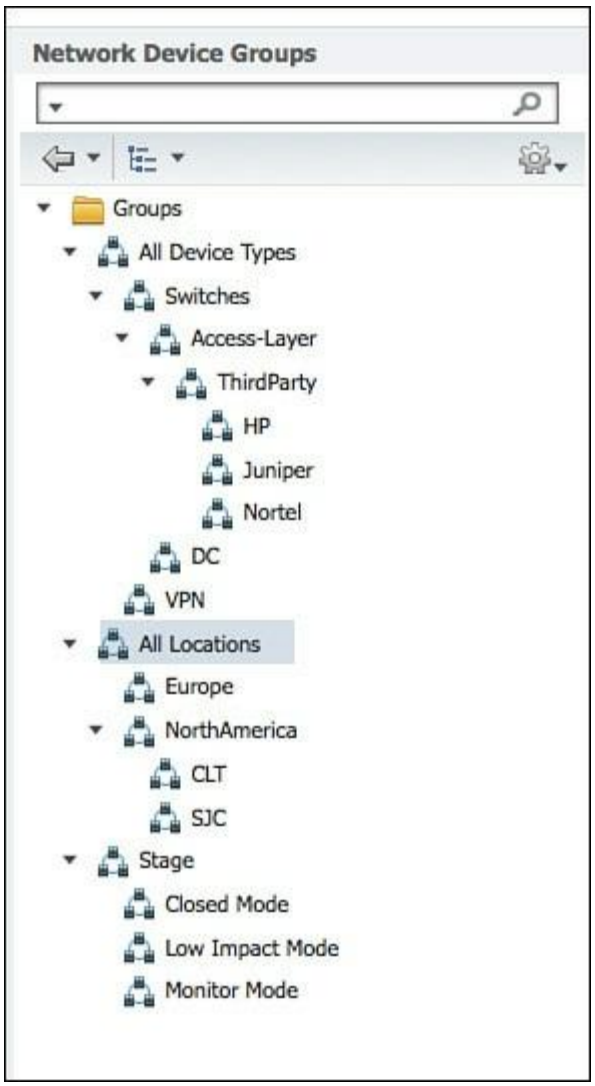


Figure 9-26 Sample NDG structure.

Network Access Devices

Now that the NDG structure is in place, it is time to add the network access devices themselves. When you add a NAD (switch, wireless LAN controller, or VPN) to ISE, you are teaching ISE about its IP address, configuring a shared secret key for the RADIUS communication, and maybe even configuring some SNMP strings to pull data from that device (think profiling data). This process is what 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, SNMP community strings, and TrustSec configuration. The shared secret key must match the configured key on the NAD exactly.

As shown in Figure 9-27, each NAD object must be configured with one or more IP address that ISE will use to communicate with the NAD, as well as identify which NAD is the source of an incoming request. An entire subnet can be associated with a single entry, isolating which NAD on the configured subnet had authenticated the endpoint; however, this might make troubleshooting difficult in the future. This screen is also where the NDGs will be selected.

https://t.me/learningnets

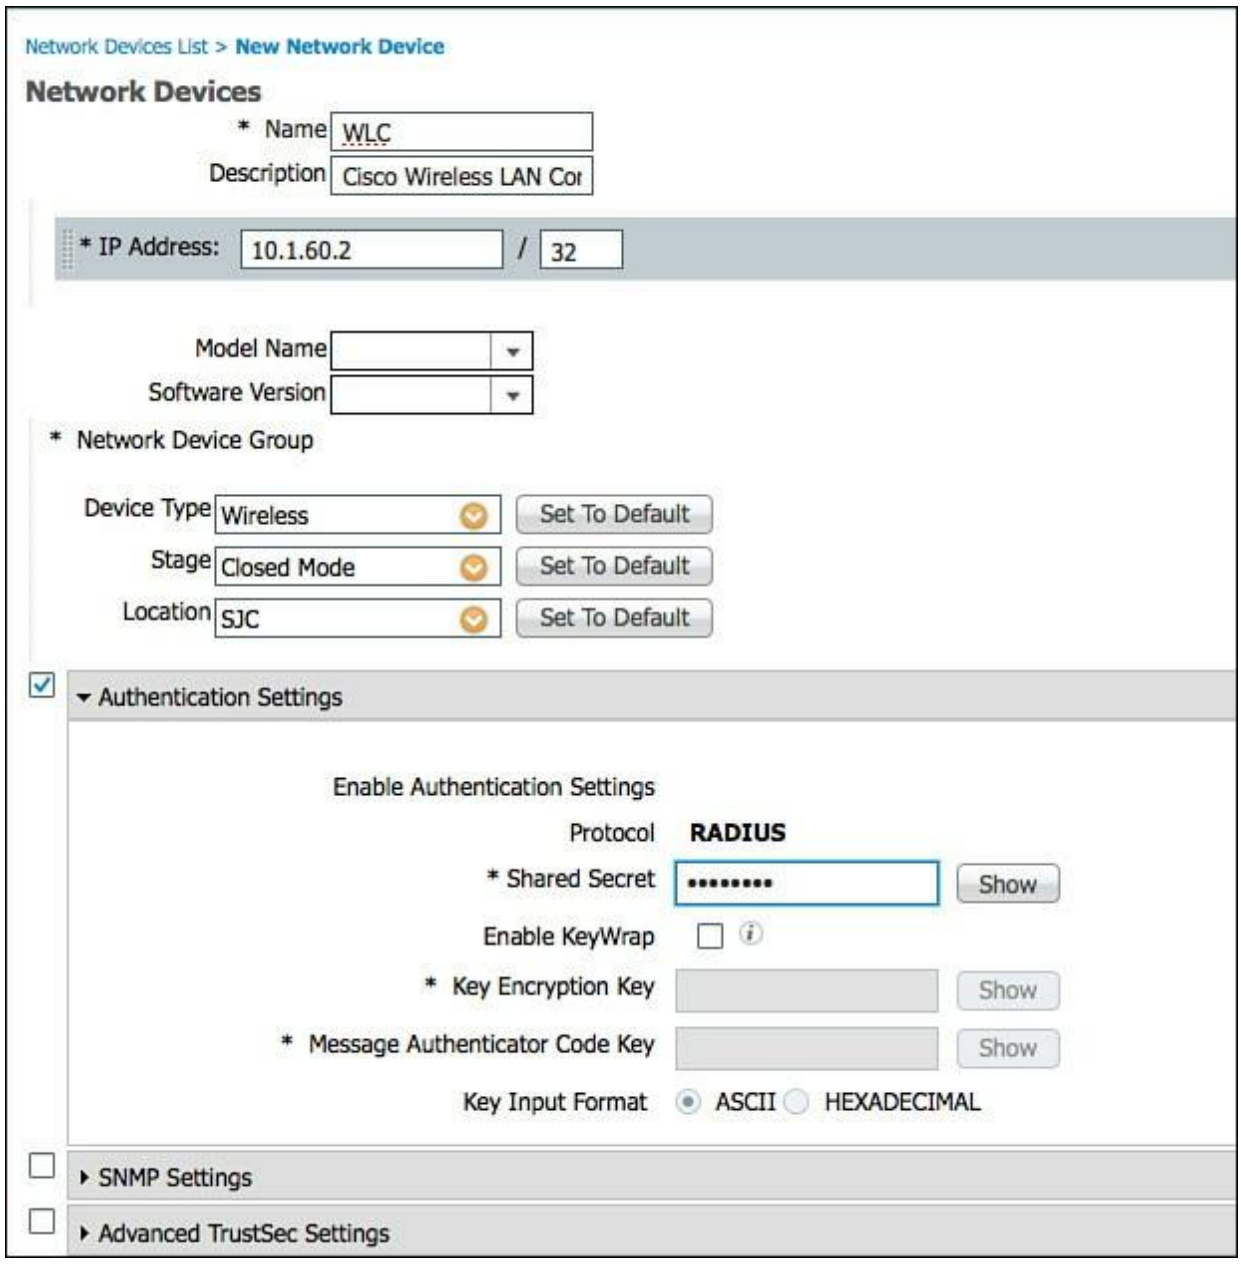


Figure 9-27 Adding a network device.

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 NDG, such as Stage.

The important section labeled Authentication Settings is where the RADIUS shared secret is configured. The SNMP Settings section configures 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. Only configure the SNMP settings for devices that do not use the Device Sensor to send ISE profiling data within RADIUS accounting packets. The Advanced TrustSec Settings is where security group tag (SGT) and network device admission control (NDAC) settings.

> **Note**
>
> It is possible and common to use a comma-separated value (CSV) file to do bulk imports of NDGs and NADs. However, these functions are beyond the scope of this book.

https://t.me/learningnets

Local User Identity Groups

Local users and local user groups are a form of identity store. Identity stores are covered in much more detail in Chapter 3, "Identity Management." However, it's important to cover these in this section for the instances where local users and groups are needed during initial setup for a small proof of concept or other small deployment.

Local user groups are used with local users, including guest users. The groups can be used to define specific levels of access and other attributes that should be applied to a group of local users instead of individuals. Sample uses of groups could be for local sponsors of guest users, guest user types, and more. These special groups are covered more in Chapter 14, "Deploying Guest Services."

Figure 9-28 displays the prebuilt groups as can be shown under **Administration > Identity Management > Groups > User Identity Groups**. These identity groups may be used directly in the authorization policies and even have a special location within the policy where a user or endpoint identity group can be used.

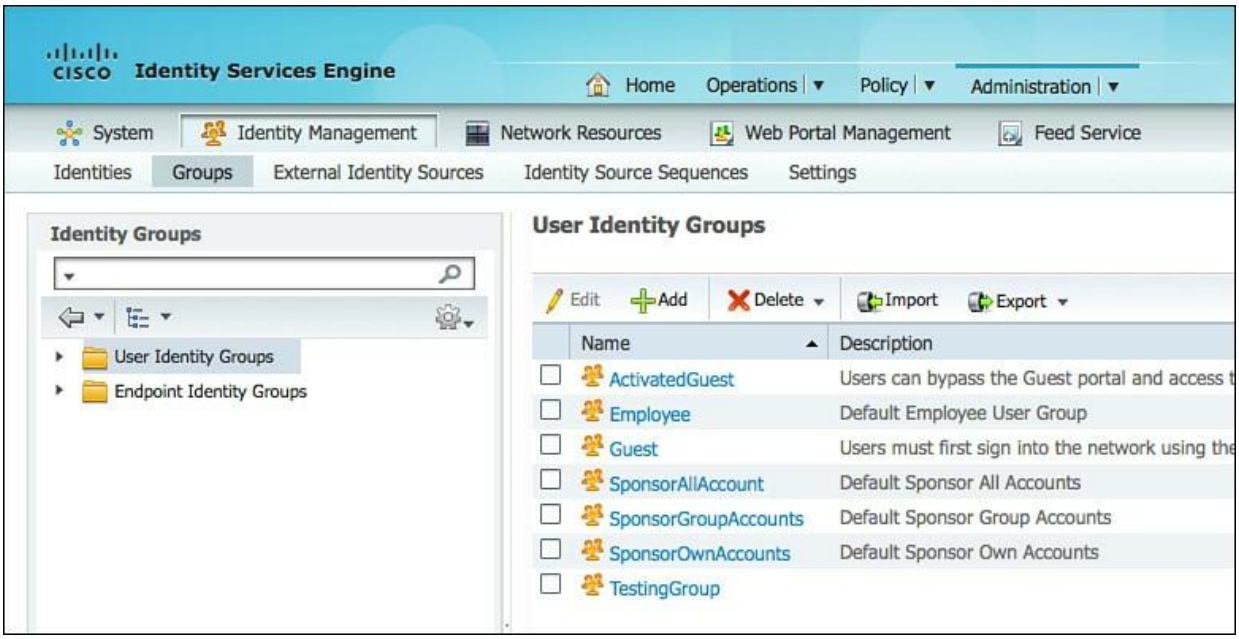


Figure 9-28 Prebuilt user identity groups.

> **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



Endpoints and endpoint identity groups are another form of identity store. Endpoint identities are covered in much more detail in Chapter 15, "Profiling." These groups are best used for what many organizations call MAC address management. For example, an endpoint identity group can be created with the name Corporate Asset, and a list of MAC addresses of corporate-owned tablets, laptops, desktops, and other devices can be added to the group.

https://t.me/learningnets

> **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 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 1.2 and above, the profile attribute itself is used in the Authorization Policy, negating the need to waste endpoint identity groups on a profile. Now they can be used for pure MAC address management.

Figure 9-29 shows the preconfigured endpoint identity groups located under **Administration > Identity Management > Groups > Endpoint Identity Groups**. Notice how Cisco-IP-Phone and Workstation are nested within the Profiled group. These are legacy groups that are left over from the pre-1.2 versions.



Figure 9-29 Preconfigured endpoint identity groups.

Local Users

Found within **Administration > Identity Management > Identities > Users**, the local users are another identity store where user accounts can be created for use with network logins, sponsor accounts for guest users, and sometimes to create test accounts for service probes. Although you can 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"—while leveraging locally configured users for administrative purposes.

For more on identity stores, refer to Chapter 3.

External Identity Stores

Even though it would be easy to continue to think of a policy server maintaining all the user accounts locally, this would not be practical, nor would it succeed in today's world. With large enterprises, there is usually a need for what is commonly called a single source of truth for identity. Often many different identity stores are used within an organization.

You will most likely need to configure one or more external identity sources during the configuration of your ISE deployment. The most common external identity sources and how to configure them are described in

https://t.me/learningnets

the following sections. For more on identity source types, refer to Chapter 3.

Active Directory

Microsoft Active Directory is the single most common external identity store shown with ISE deployments. Cisco ISE will join AD, creating a machine account and using native AD communications to query the identity store for user accounts and their attributes.



Cisco ISE 1.2 will join a single AD domain and can query other domains within the AD forest, as long as the trust relationships exist. For more on AD as an identity source, please refer to Chapter 3.

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 whether ISE machine account already exists).
- Add workstation to domain (if one does not already exist).
- Set attributes on the new machine account (OS type and version—optional).

In addition to an account with the appropriate rights and permissions, certain communication ports must be open between ISE and the domain controllers. Table 9-2 lists the ports that are required for Active Directory communication.

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	

Table 9-2 Active Directory Communication Ports

> **Note**
>
> It is also critical to ensure that good time synchronization is in place before joining Active Directory. Time synchronization issues are the number-one reason ISE fails to join AD.

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**.

The first procedure is to name the AD domain in the Connection tab and save it, as shown in Figure 9-30. The domain name should be in a DNS style format, such as ciscopress.com or ise.local. After the domain name is saved, the Connection tab will show that the domain has not been joined yet, as displayed in Figure 9-31. For the sake of this book, we'll retain the default identity store name of "AD1." This value will be referenced in future chapters as we define the ISE policy.

https://t.me/learningnets

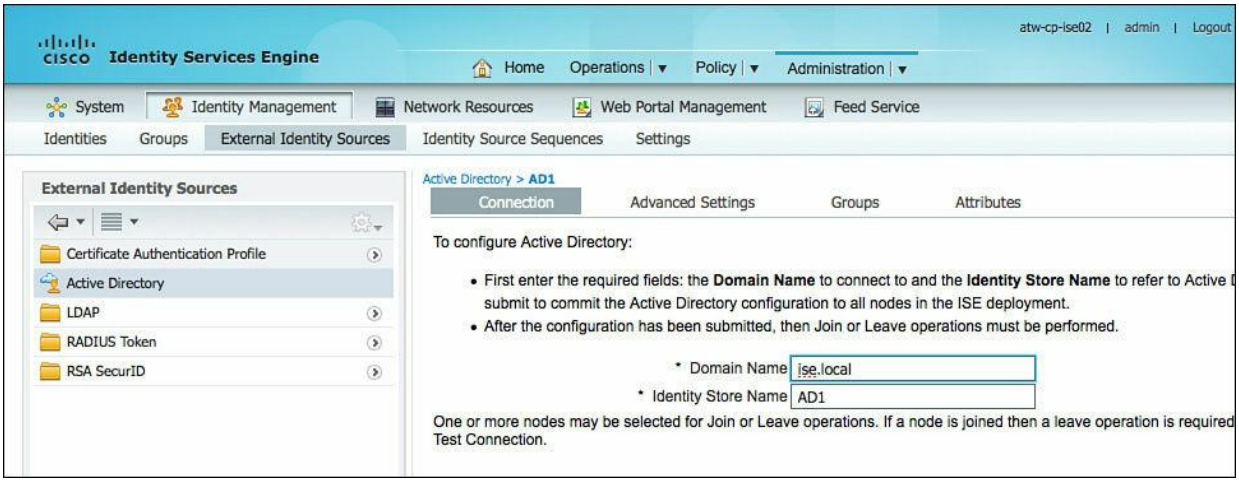


Figure 9-30 Active Directory > Connection tab.

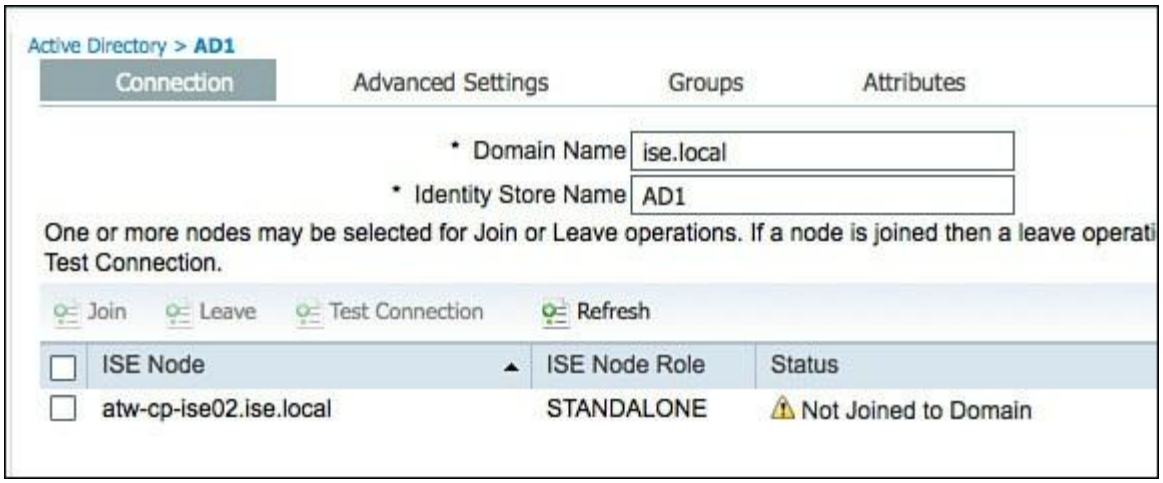


Figure 9-31 Domain not joined.

Although the cConnection tab in Figure 9-30 displays only a single ISE node, the screen will show the status of all ISE nodes in the ISE cube as it pertains to the AD connection(s). In other words, if the ISE deployment (also called the ISE cube) had 34 nodes in it, this tab would display all 34 nodes and the current status of each node's connection to the configured AD domain. This is necessary because every ISE node will join AD independently of the other nodes, although the entire configuration is handled from this centralized GUI.

Click the **Join** button, which will open a dialog box for the username and password of an account with the correct rights and permissions to join the machine to AD. Enter the correct username and password and click **OK**, as shown in Figure 9-32.

https://t.me/learningnets

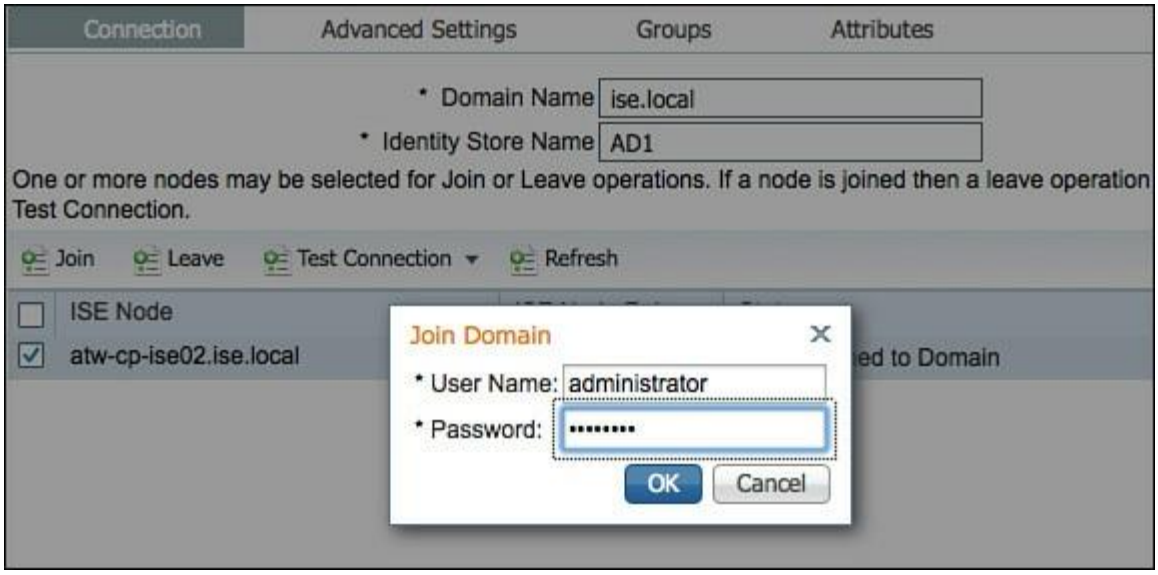


Figure 9-32 Join domain.

Figure 9-33 displays the Connection tab after the ISE node has been joined to the domain.

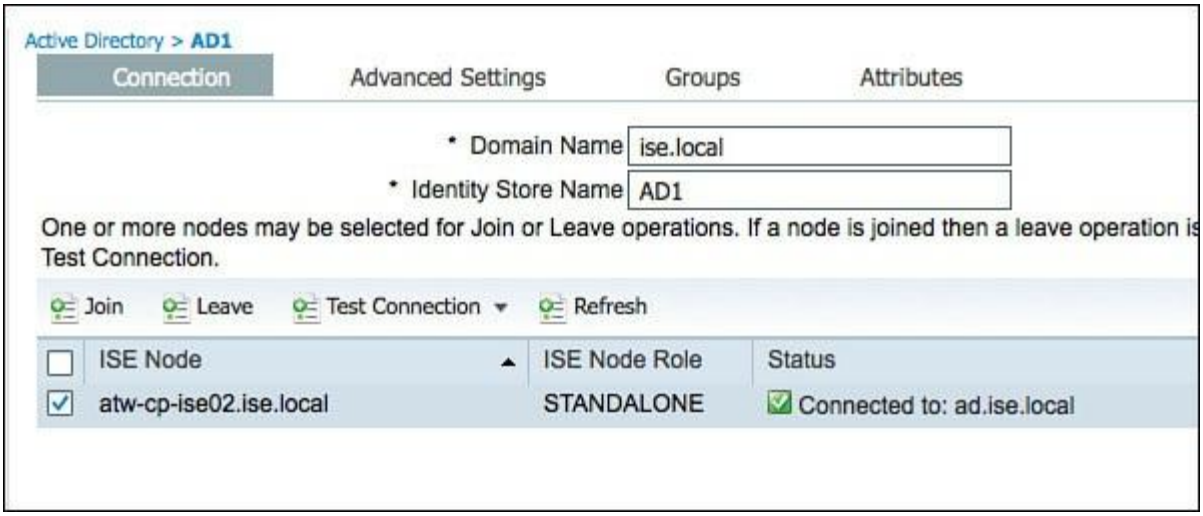


Figure 9-33 Active Directory > Connection tab (after domain join).

Now that ISE has been joined to Active Directory, there are other items that will make the overall deployment easier, such as preselecting the AD groups of interest or even preselecting specific account attributes that will be used in access policy.

To preselect the groups of interest, click the **Groups** tab from **Administration > Identity Management > External Identity Sources > Active Directory**. Click **Add > Select Groups from Directory**, and then the **Select Directory Groups** pop-up will open, as shown in Figure 9-34. After checking each of the groups of interest, click **OK** to accept the choices.

https://t.me/learningnets

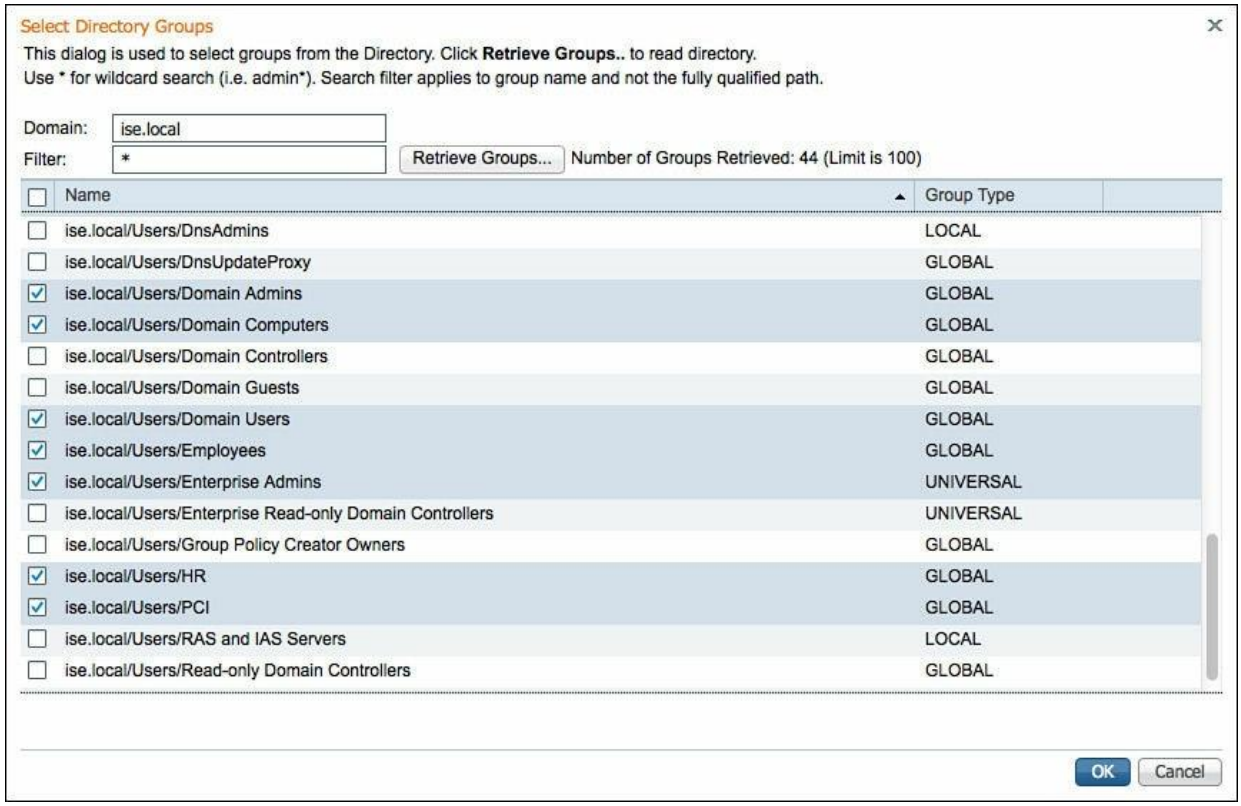


Figure 9-34 Select Directory groups.

Now, when you are building a policy that will use AD group membership as a condition, you will be presented with this list of groups instead of having to sort through the entire list of all AD groups. Figure 9-35 displays the final selection of groups.

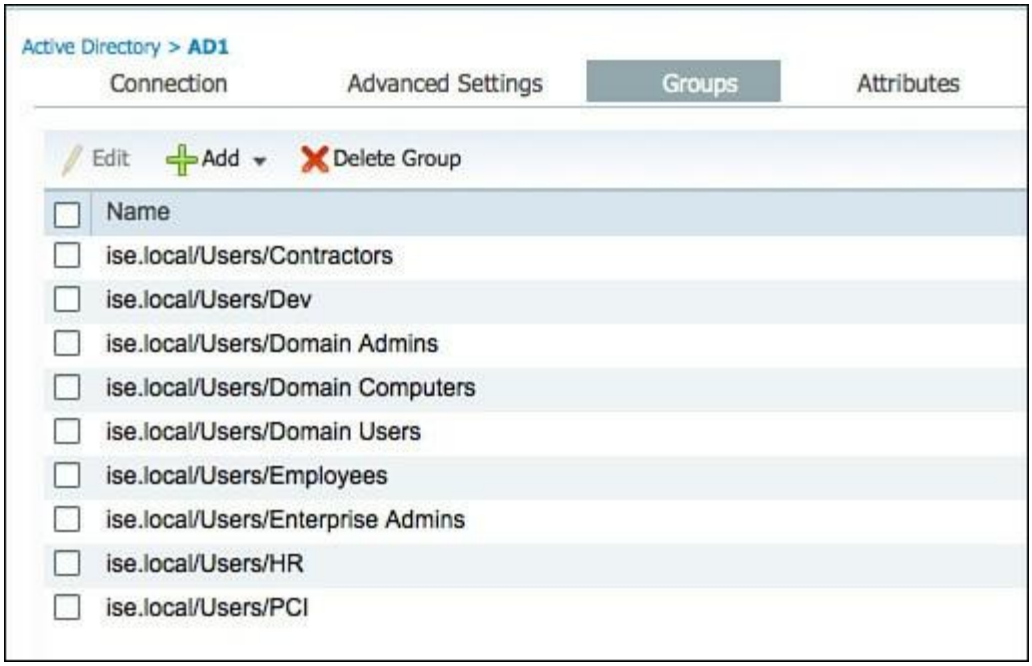


Figure 9-35 Selected Active Directory groups.

In addition to being members of groups, other attributes within Active Directory might be of interest for the creation of access policies. These attributes can be selected in the Attributes tab by clicking **Add** > **Select Attributes from Directory**. To make the administration easier, ISE will require the name of a sample user account. Then after you click the Retrieve Attributes button, all attributes of that sample user account will be

https://t.me/learningnets

displayed. Select the attributes of interest, and click. Don't forget to save the configuration when you are finished.

Figure 9-36 shows attributes being selected, using the employee1 account as the example, and Figure 9-37 displays the final selection of attributes.

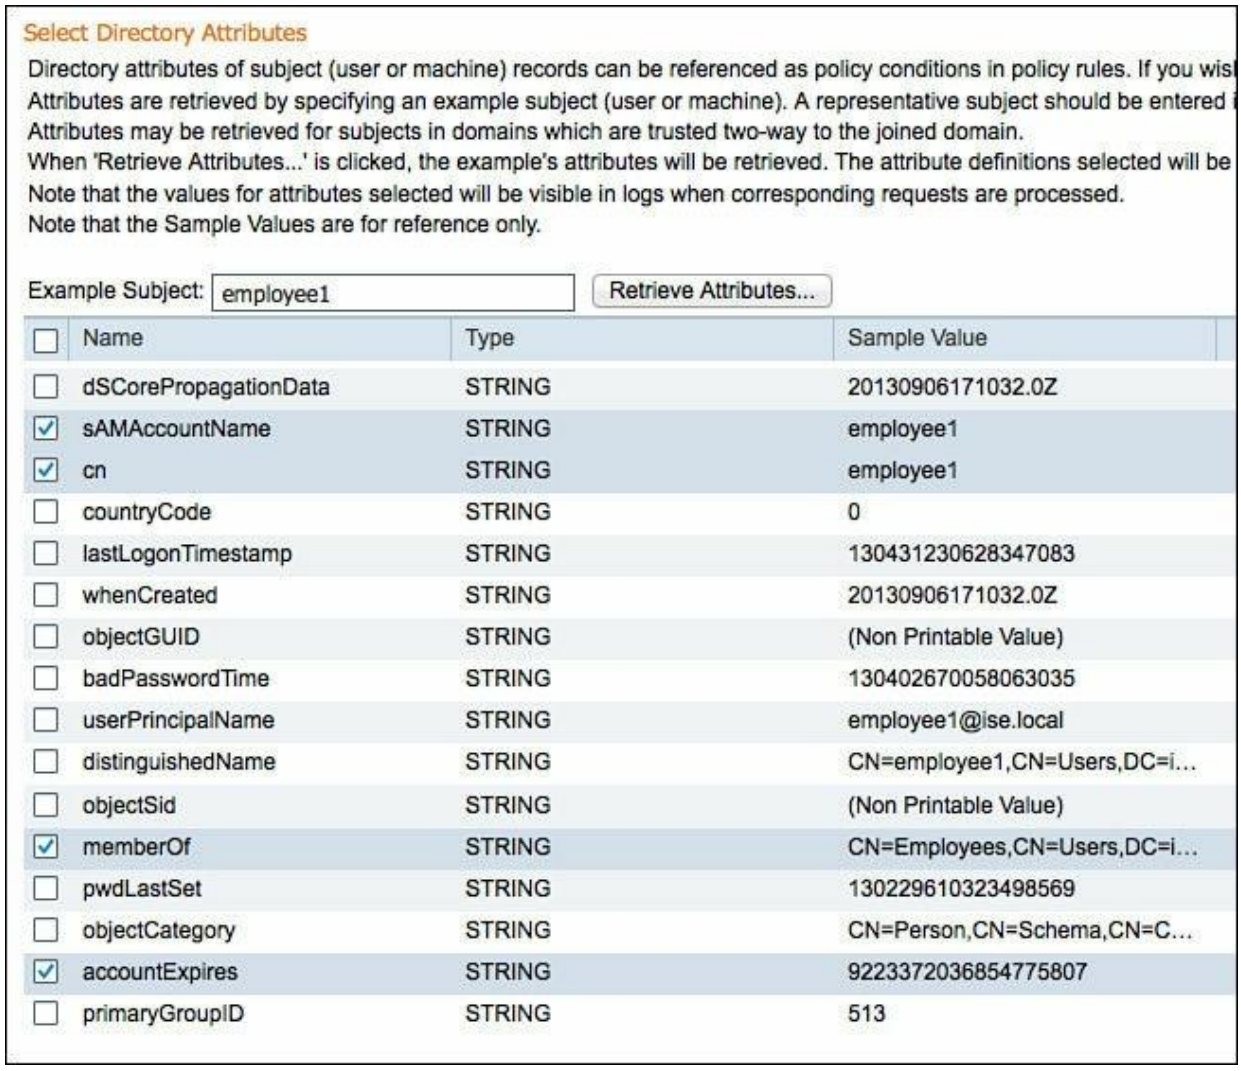


Figure 9-36 Selected Active Directory groups.

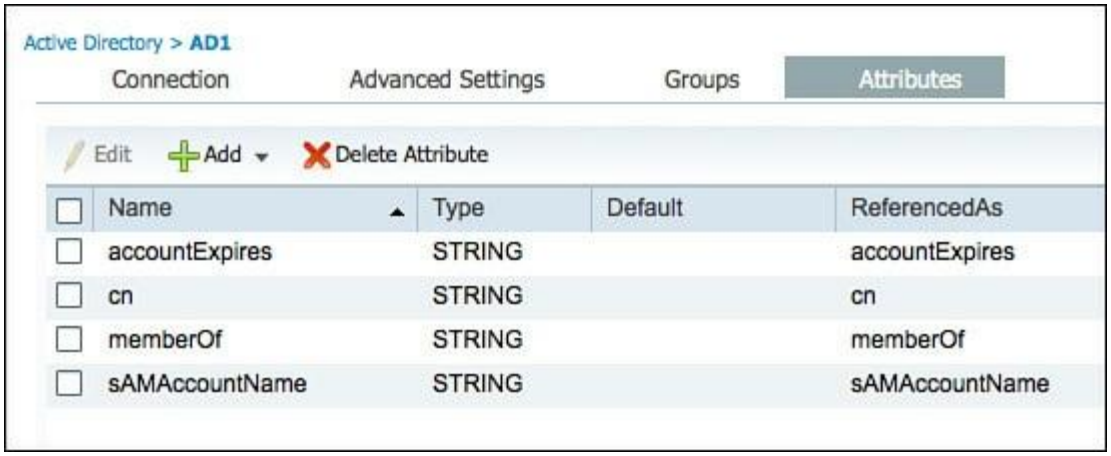


Figure 9-37 Final selection of AD attributes.

In addition to the configuration of the AD Connector, a Test Connection Tool includes both basic and

https://t.me/learningnets

detailed diagnostic tests for the connection. ISE administrators as well as Cisco TAC commonly use this tool to determine theroot cause of any issues with AD communication. Figure 9-38 shows the launching of the tool, and Figure 9-39 shows the output of the detailed test results.

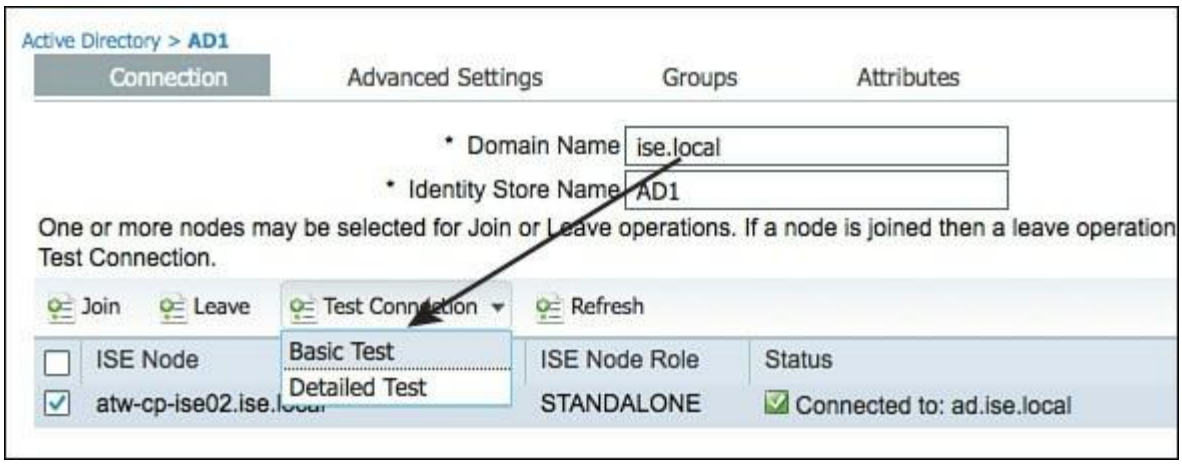


Figure 9-38 Test connection selection.

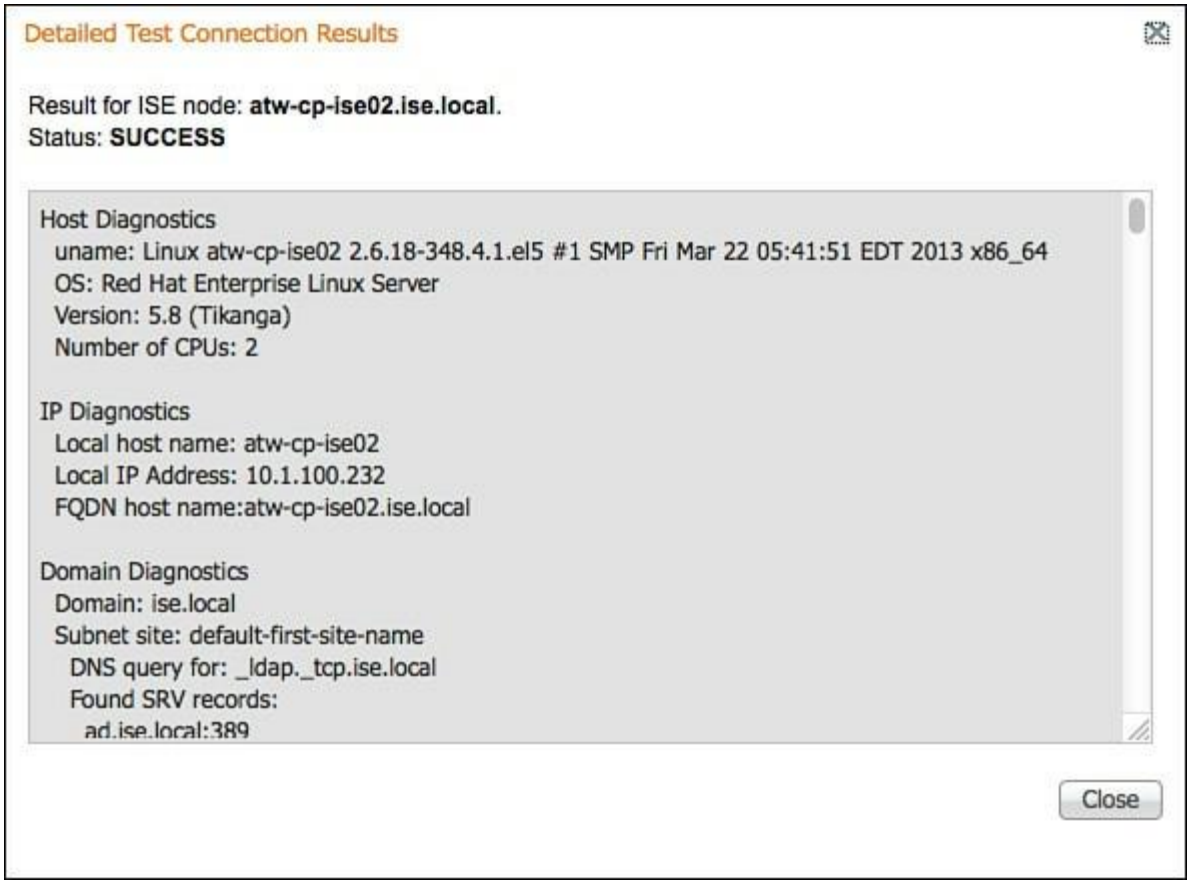


Figure 9-39 Detailed test output.

It's important to note that AD can be used for authentication as well as authorization. However, it is common to use AD for the authorization operation only, while another source is used for the authentication operation.

https://t.me/learningnets

Certificate Authentication Profile

As described in the previous section, many identity stores can be used for authentication, not only Active Directory. However, the authorization of the user or device might still require some information that is available in AD or another ID store. The use of digital certificates is one such example.

When you break it down, the authentication of a digital certificate is basically nothing more than ISE validating that the certificate was signed by a trusted authority and the certificate is still valid. However, to perform authorization for the 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 authentications is a certificate authentication profile (CAP). The CAP tells ISE which value from the certificate should be used as the principle identity for comparison during the authorization phase.

Figure 9-40 shows a common CAP where the certificate subject's common name (CN) field is being assigned as the principle identity.

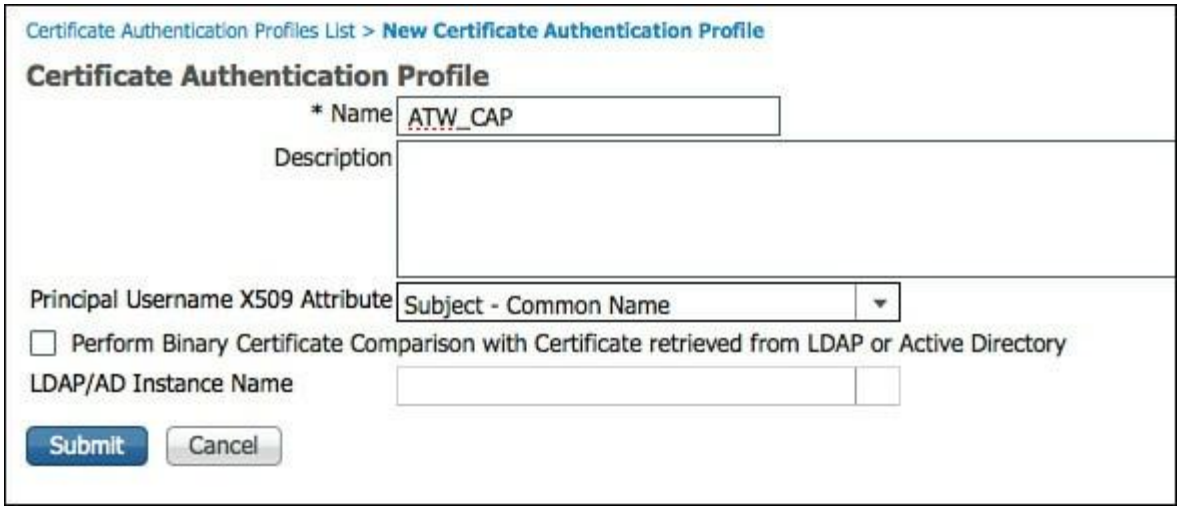


Figure 9-40 Certificate authentication profile.

The second major function of a CAP is to determine whether a binary comparison of the certificate should be performed, and if so, whicht 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 itself). This setting is configured within the CAP by checking 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 will check a series of identity sources from top to bottom (as configured), allowing a single authentication rule check a number of identity sources, instead of just one.

ISSes are configured within the ISE GUI under Administration > Identity Management > Identity Source Sequences. With ISE 1.2, three preexisting ISSes are displayed in Figure 9-41 and more can be created as needed.

https://t.me/learningnets

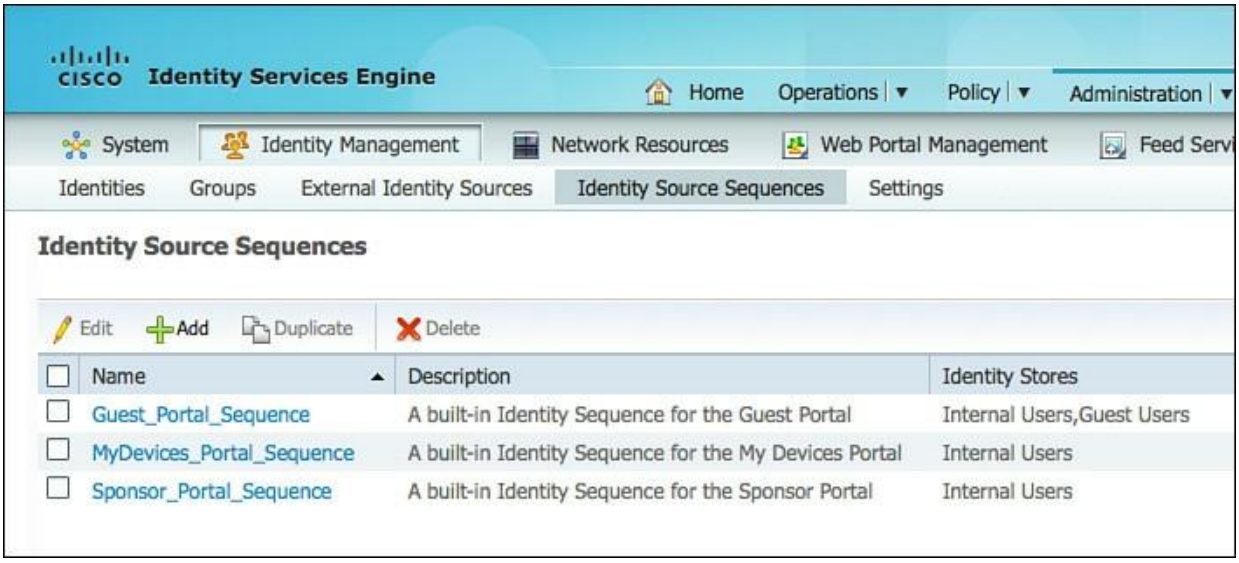


Figure 9-41 Identity source sequences.

To add a new ISS, click **Add**. Provide a name for the sequence, such as ALL_ID_STORES. As with all objects created within ISE, you should provide a good description to help document the configuration for any other administrators of the system. The description will even help you—in case you have to revisit the configuration after a break, solid descriptions will help you remember why this ISS was created.

A CAP can be selected to include certificate-based authentications, along with the other selected identity stores. Figure 9-42 shows part one of an ISS that will authenticate against a CAP followed by the Internal Users, Guest Users, and finally Active Directory.

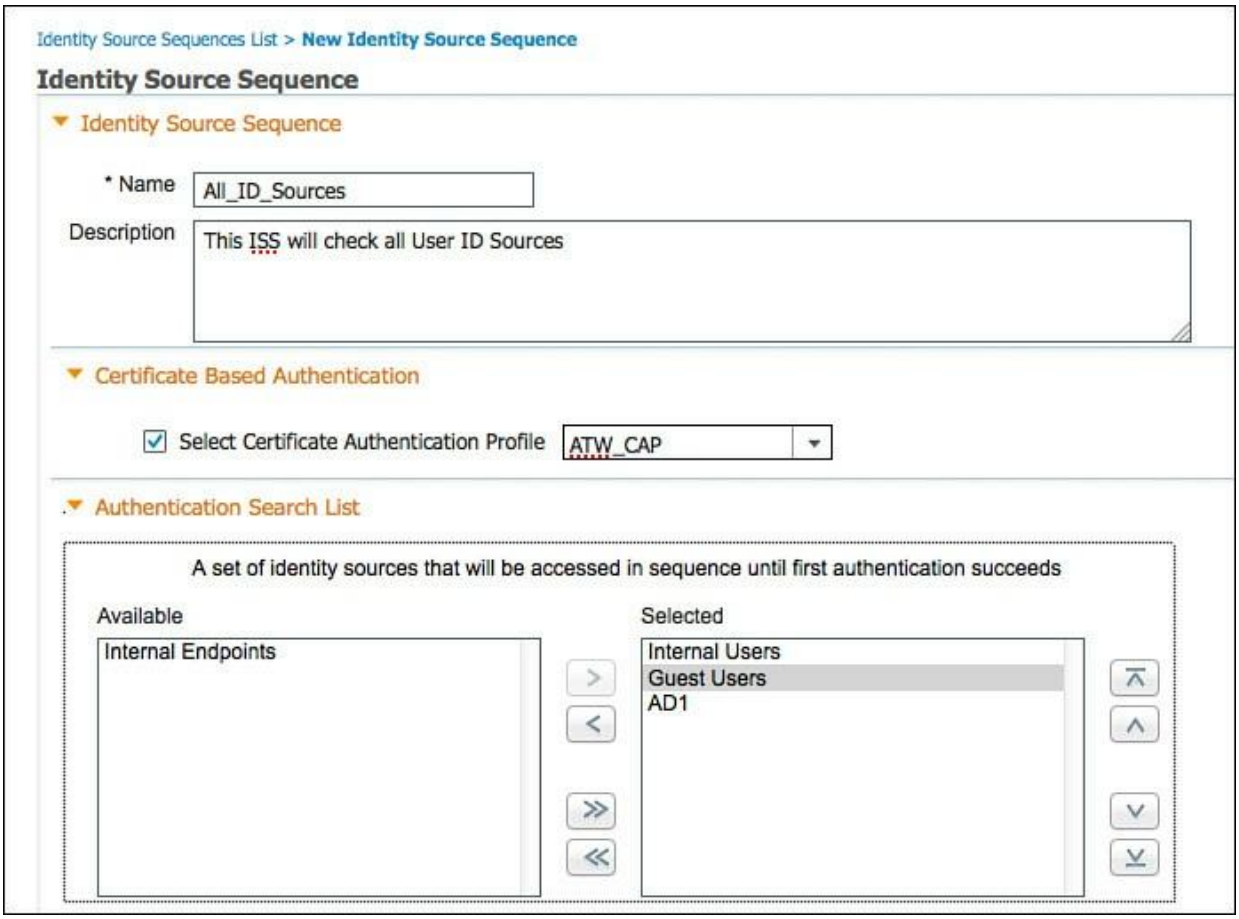


Figure 9-42 All identity stores ISS, part one.

https://t.me/learningnets

Figure 9-43 shows the second part of the ISS, where the decision on which action to take if a selected identity store cannot be accessed is configured. 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.

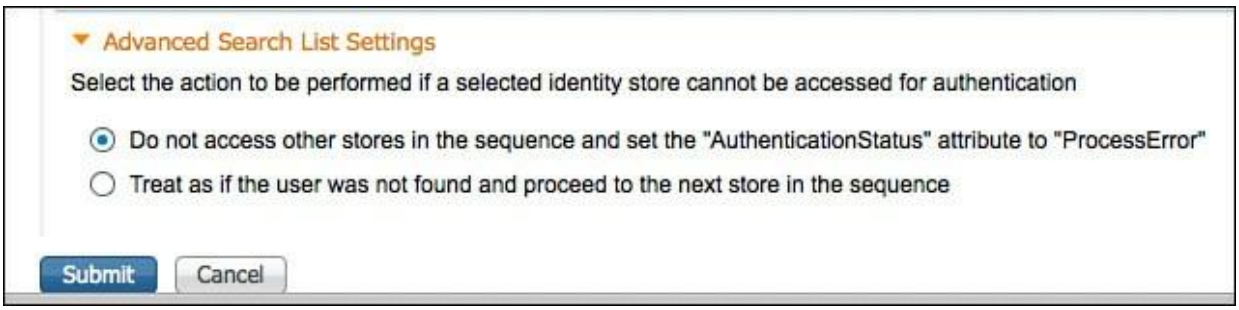


Figure 9-43 All identity stores ISS, part two.

This chapter discussed the initial configuration of ISE and the creation of identity sources for use in authentication policies. The next chapter, Chapter 10, "Authentication Policies," will detail the use of those identity sources in the authentication policies themselves.

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 9-3 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Key Topic Element	Description	Page
Section	Self-signed certificates	206
Section	CA-Signed certificates	206
Paragraph	Certificate name matching	208
Paragraph	Trusted certificates	213
Paragraph	Local endpoint groups	219
Paragraph	Active directory	221

Table 9-3 Key Topics for Chapter 9

https://t.me/learningnets

Chapter 10. Authentication Policies

This chapter covers the following exam topics:

- **Describe Identity Store Options (i.e., 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 validating of a credential. It is an important step in the process of performing any sort of secure network access control. When thinking about authentication, it often helps to relate the topic to something that occurs within your day-to-day life

Consider when a highway patrol officer has a driver pull his car over to the side of the road. The officer will walk up to the driver's window and ask for his driver's license and proof of insurance (at least that is what happens in the United States). The driver will hopefully hand over these documents for the officer to inspect.

The officer should examine the driver's license and determine whether it appears to be real. The hologram and watermarks in the driver's license are there, so it appears to be real. The picture on the license looks like the driver who handed over the license. The license hasn't expired. After going back to the squad car, the officer will perform a lookup into the Department of Motor Vehicles database to determine whether the license has 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 (the officer would immediately reject a library card because that is not an allowed form of ID for a driver). The policy will route authentication requests to the correct identity store (North Carolina DMV, or New York DMV, and so on and so on); validate the identity (was this a valid license for that driver); and finally "pass" successful authentications over to the authorization policy (was the driver allowed to exceed the speed limit and run other drivers off the road).

When thinking about authentication for network access, it often helps to relate the topic to an example such as this one, where it is something that occurs within your day-to-day life. Typically, the goals are similar, and it helps to understand the difference between authentication and authorization.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

https://t.me/learningnets

Foundation Topics Section	Questions
Describe Identity Store Options	6
Implement Wired/Wireless 802.1X	6-7
AV Pairs	7-8
EAP Types	2
Implement MAB	1, 4
Describe the MAB Process Within an 802.1X Framework	1
ISE Authentication/Authorization Policies	3, 5, 9-10

Table 10-1 "Do I Know This Already?" Section-to-Question Mapping

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 the `service-type` set to `login` and the `called-station-id` populated with the MAC address of the endpoint.
   - **b.** The RADIUS packet must have the `service-type` set to `Call-Check` and the `calling-station-id` populated with the MAC address of the endpoint.
   - **c.** The RADIUS packet must have the `service-type` set to `Call-Check` and the `called-station-id` populated with the MAC address of the endpoint
   - **d.** The RADIUS packet must have the `service-type` set to `login` and the `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 choices 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 dACL 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. True or False? You must select Detect PAP as Host Lookup to enable MAB requests for Cisco nNetwork devices.

https://t.me/learningnets

a. True

b. False

5. True or False? Policy conditions from attribute dictionaries can be saved as conditions inline while building authentication policies.

a. True

b. False

6. Which method will work effectively to allow a different Identity store to be selected for each EAP type used?

a. This is not possible because the first rule to match 802.1X will be used and no further rules can be used.

b. Create one authentication rule that matches a service type framed for each of the EAP protocols. Each authentication rule should have one subrule that matches the EapAuthentication (such as EAP-TLS, EAP-FAST, and so on).

c. This is only possible for the main EAP types. If there is an inner method of EAP-MSCHAPv2 with PEAP, it must be sent to the same identity store as the EAP-MSCHAPv2 inner method of EAP-FAST.

d. Create one sub-rule for each EAP type under the default 802.1X authentication rule that points to the appropriate identity store per rule.

7. Which 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 sub-rule 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 will send an authentication to the selected identity store.

10. True or False? The Drop option for an authentication rule will allow ISE to act as if it were not "alive" so the network device will no longer send authentication requests to that ISE server.

a. True

b. False

https://t.me/learningnets

Foundation Topics

The Relationship Between Authentication and Authorization

What is authentication, and what is authorization? Many IT professionals, especially those with wireless backgrounds (versus those with a security background), will tend to confuse these terms and what they actually do. Wireless was really the first place in the network where 802.1X took hold and is still the most prevalent use case of 802.1X authentication. With that in mind, the vast majority of wireless environments would provide a user with full network access as long as their usernames and passwords were correct (meaning that authentication was successful).



An authentication is simply the validating of credentials. If you were to go into a bank and request a withdrawal from an account, the teller would ask for your ID. You would pass your driver's license to the teller, and she would look the license over, going through a checklist of sorts:

- Is the license from a recognized authority (one of the United States or a military ID)?
- Does the picture on the ID look like the person in front of the teller?

Let's say for conversation's sake that you handed her a valid ID (authentication was successful). Does that mean you are entitled to the money you asked for?

The next step for the bank teller would be to check the account and ensure that the person requesting the withdrawal is entitled to complete that transaction. Perhaps you are allowed to withdraw up to $1,000 but no more. This is the process of authorization. Just having a successful authentication does not prove entitlement.

This is why most of the time expended working within a product like Cisco ISE is spent setting up and tuning the authorization policy. Authorization is where the bulk of the decisions are made.

Authentication Policy

Authentication policies are the first opportunity for Cisco ISE to interact with the RADIUS Access-Request coming from the network access device (NAD). The authentication policy has very specific goals, but ultimately the main goal is to process the authentication request quickly so it can be dropped (if invalid), denied immediately if the credentials were incorrect, or forwarded to be run through the authorization policies (if successful).

Goals of an Authentication Policy

Authentication policies have a few goals:

1. Drop traffic that isn't allowed and prevent it from taking up any more processing power.
2. Route authentication requests to the correct identity store—sometimes called a policy information point (PIP).
3. Validate the identity.
4. Pass successful authentications over to the authorization policy.

https://t.me/learningnets

Goal 1—Accept Only Allowed Protocols

By default, ISE will allow nearly all supported authentication protocols; however, it would behoove the organization to lock this down to only the ones that are expected and supported. This serves a few purposes: It keeps the load on the Policy Service Nodes down and uses the authentication protocol to help choose the right identity store. For example, think of a corporation that wants to support only 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 allow only EAP-TLS and not waste time processing PEAP requests from device that are not configured with a certificate.

Keep in mind that a company is best served to have its security policy dictate which authentication protocols meet the security requirements of the organization. This is where the less secure protocols can be disabled, ensuring that any protocol that is more easily compromised is shut off.

Allowing only certain protocols defines which set of protocols should be permitted as well as the specific tuning of those protocols. For example, EAP-FAST can be in the allowed protocol list, but it also configures the options for EAP-FAST such as whether to allow in-band PAC provisioning or to use EAP chaining.

Goal 2—Select the Correct Identity Store

After the authentication has been accepted, ISE must make an identity store selection decision; you can even consider it to be an identity routing decision. Based on the attributes of an incoming authentication, it must determine which identity store should be used. Obviously, if a certificate is being presented, ISE should not try to validate that certificate against the internal user database that is expecting usernames and passwords.

If your company has multiple lines of business, it can also have more than one Active Directory (AD) domain or more than one LDAP store. Using attributes in the authentication request, you can pick the correct domain or LDAP store.

Goal 3—Validate the Identity

After the correct identity store has been identified, ISE must make sure the credentials are valid. In the case of password-based authentications, it must determine whether:

- The username is valid.
- The user's password matches what is in the Directory Store.

For certificate-based authentications, it must determine whether:

- The digital certificate has been issued and signed by a trusted certificate authority (CA).
- The certificate has expired (checks both the start and end dates).
- The certificate has been revoked.
- The client has provided proof of possession.
- The certificate presented has the correct key usage, critical extensions, and extended key usage values present.

Goal 4—Pass the Request to the Authorization Policy

If the authentication failed, the policy can reject the request without wasting the CPU cycles comparing the request to the authorization policy. Also, if the request did not match any of the configured rules, should we send a reject message? However, when the request passes authentication, it is now time for the hand-off to the authorization policy.

Understanding Authentication Policies

Now that you understand the four main responsibilities of the authentication policy, it will be easier to understand why you are doing the things we are introducing in this section.

To understand authentication policies even more, we will now examine a few.

https://t.me/learningnets

From the ISE GUI, navigate to **Policy > Authentication.** You will notice the default rules as displayed in Figure 10-1. Basic authentication policy rules are logically organized in this manner:

Click here to view code image

```
IF conditions THEN ALLOW PROTCOLS IN LIST AllowedProtocolList
AND CHECK THE IDENTITY STORE IN LIST IdentityStore
```

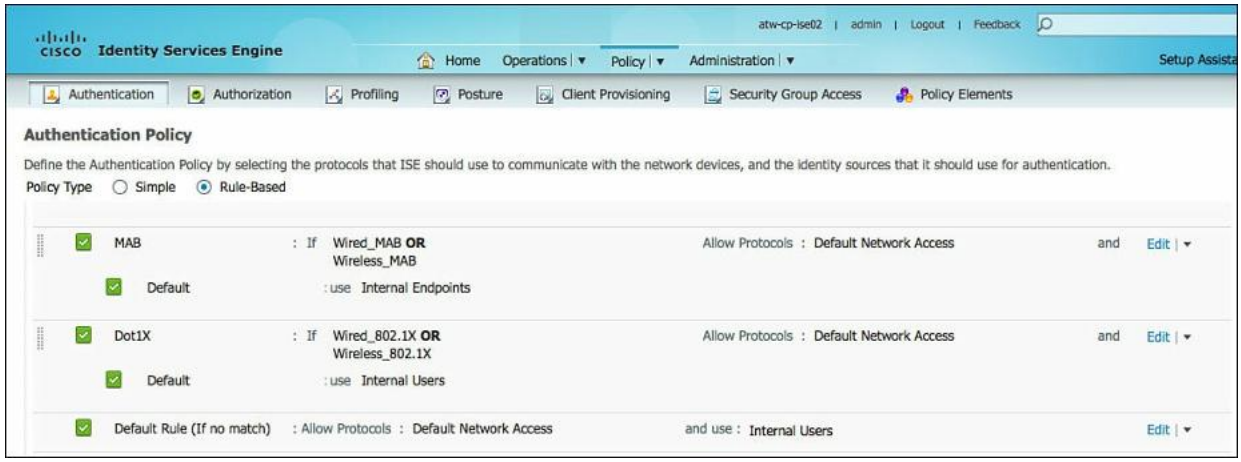


Figure 10-1 Default authentication policy.

Rules are processed in a top-down, first-match order, just like a firewall policy. So if the conditions do not match, the authentication will be compared to the next rule in the policy.

As shown in Figure 10-1, ISE is preconfigured with a default rule for MAC Authentication Bypass (MAB). MAB is used for a number of things, such as allowing nonauthenticating endpoints onto the network, guest access, BYOD, and more, that will be covered in further 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 can help to follow along with the text.

Figure 10-2 demonstrates the MAB rule in flow chart format.

https://t.me/learningnets

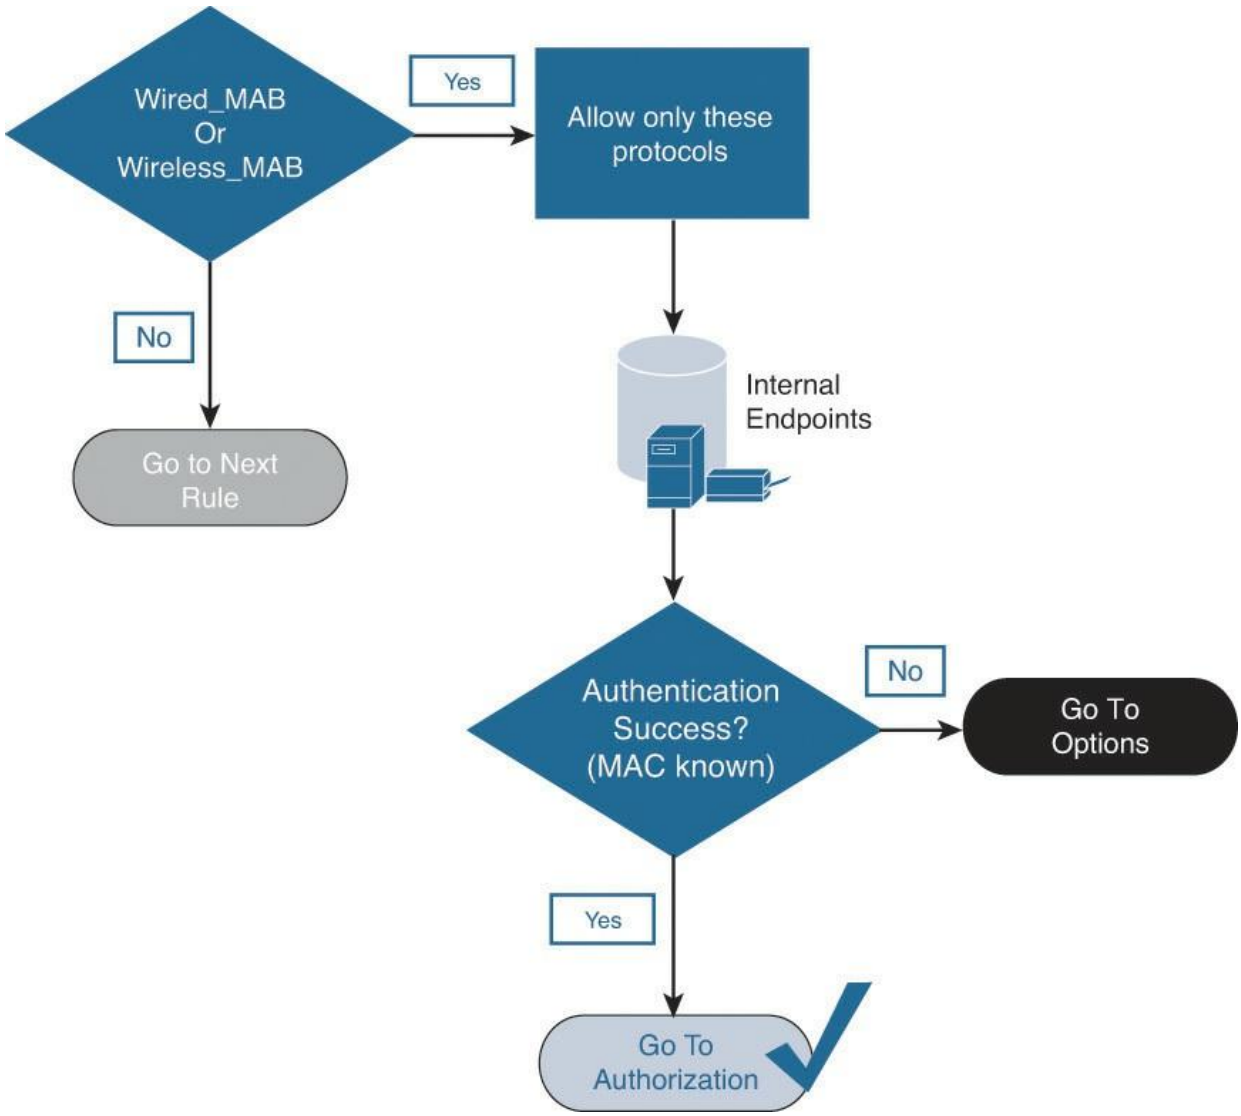


Figure 10-2 MAB rule flow chart.

Conditions

The conditions of this rule state, "If the authentication request is Wired_MAB or Wireless_MAB, it will match this rule." We can expand these conditions by mousing over the conditions and clicking the target icon that appears or by looking directly at the authentication conditions. Here's how:

Step 1. Navigate to **Policy > Policy Elements > Conditions > Authentication > Compound Conditions**.

Step 2. Select **Wired_MAB**.

As shown in Figure 10-3, Wired_MAB is looking for the RADIUS `Service-Type` to be `Call-Check` and the `NAS-Port-Type` to be `Ethernet`. This combination of attributes from the RADIUS authentication packet tells ISE that it is a MABs (Service-type = Call-check) request from a switch (NAS-Port-Type = Ethernet).

https://t.me/learningnets

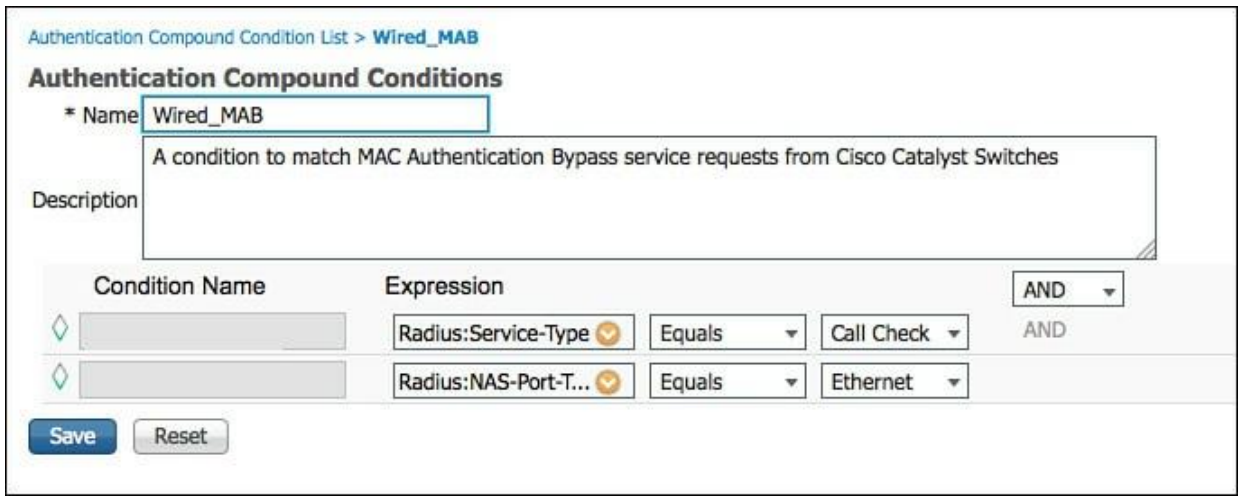


Figure 10-3 Wired_MAB condition.

Figure 10-4 highlights these key attributes in a packet capture of the MAB authentication request.

https://t.me/learningnets

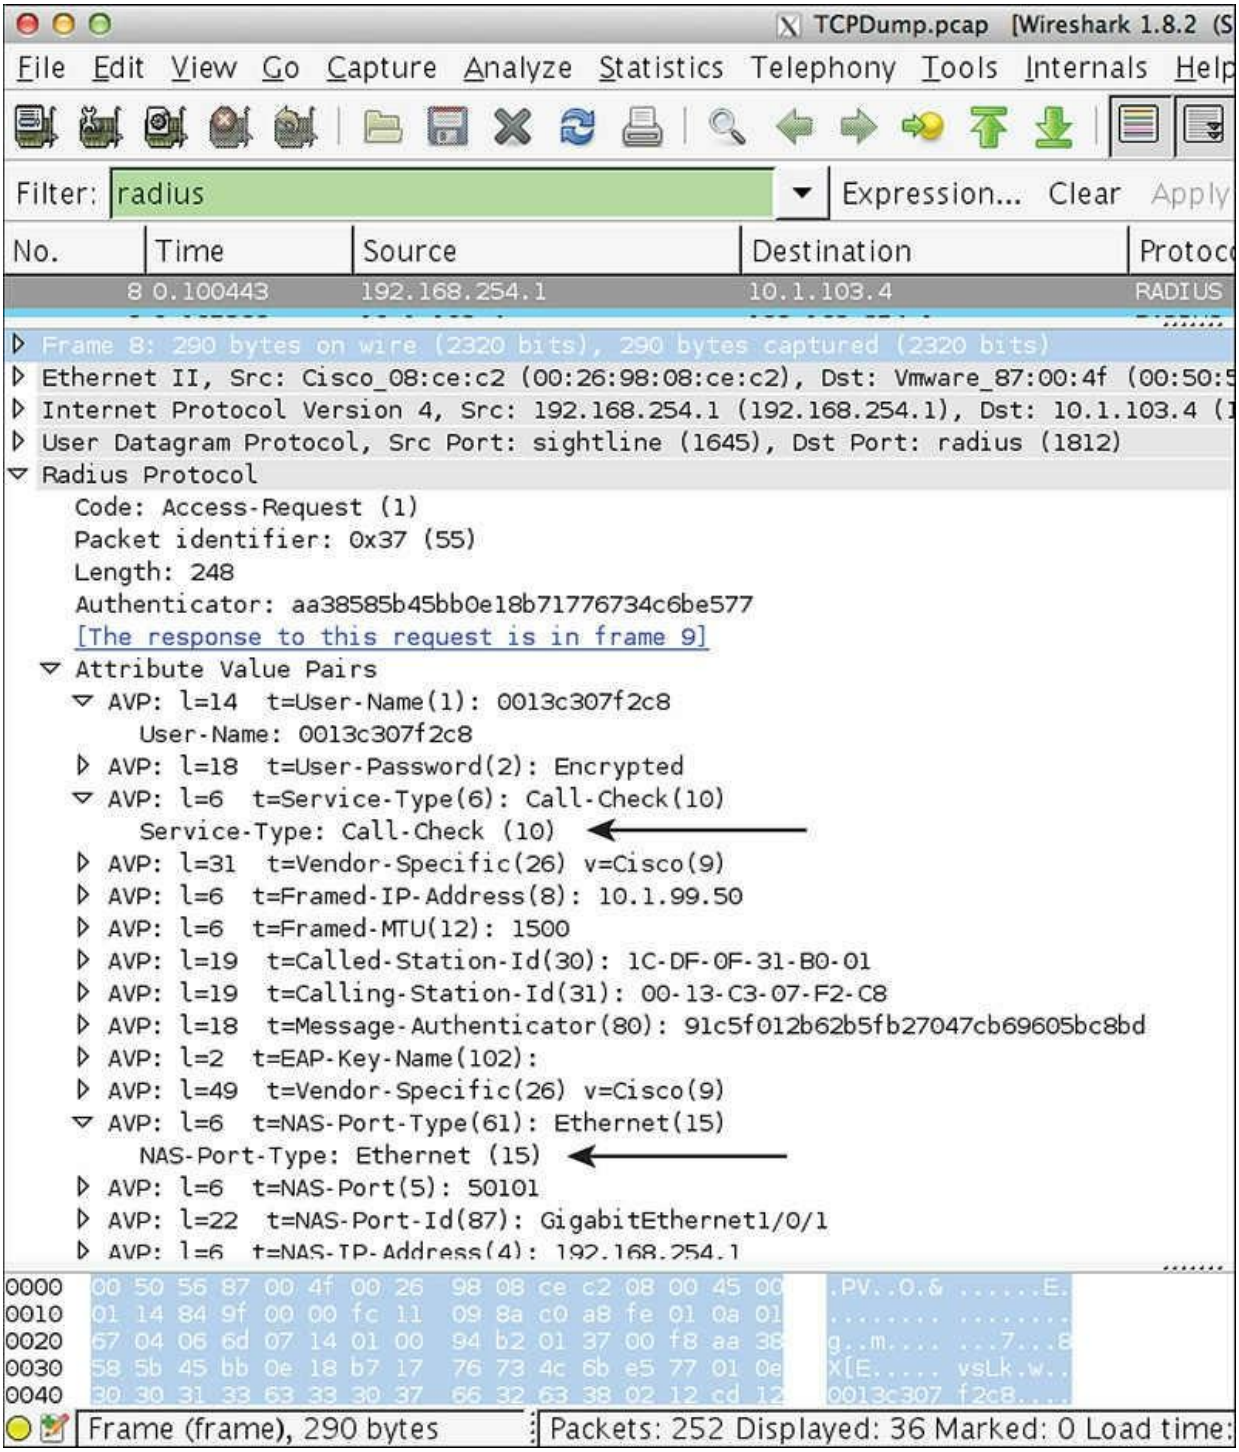


Figure 10-4 Packet capture of wired MAB.

Step 3. Navigate back to **Policy > Policy Elements > Conditions > Authentication > Compound Conditions**.

Step 4. Select **Wireless_MAB**.

As shown in Figure 10-5, Wireless MAB is similar. However, it uses a NAS-Port-Type of Wireless - IEEE 802.11. This combination of attributes from the RADIUS authentication packet tells ISE that it is a MAB request from a wireless device.

https://t.me/learningnets

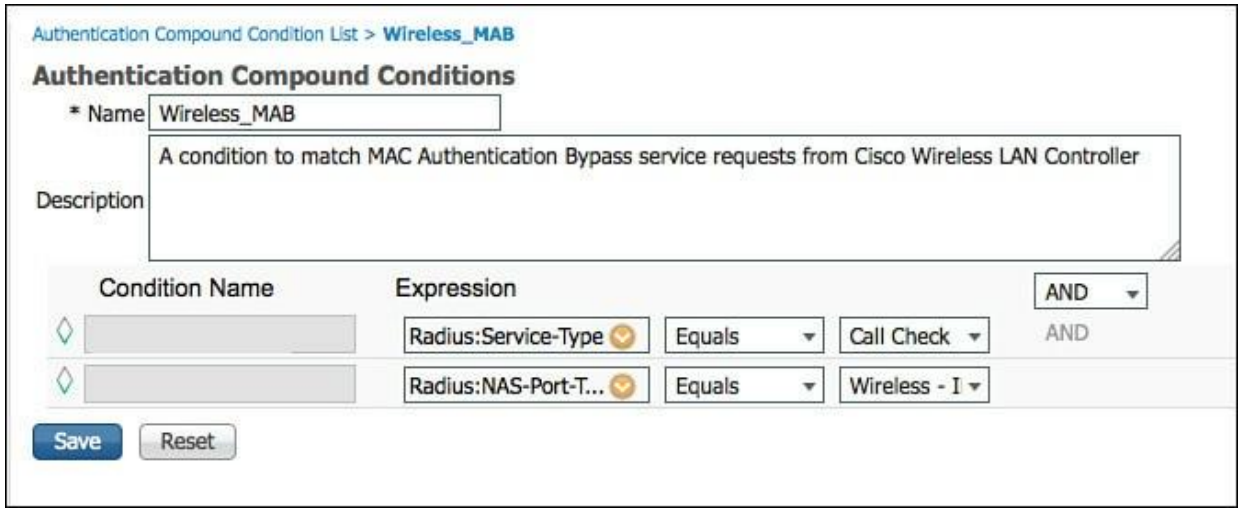


Figure 10-5 Wireless_MAB condition.

Allowed Protocols

After the conditions are matched, the rule now dictates which authentication protocols are permitted. Looking at the predefined MAB rule, this rule uses the Default Network Access list of allowed protocols (which is almost every supported authentication protocol). You can create multiple allowed protocols list, using a different one in each authentication policy rule.

Let's examine the default allowed protocols. From the ISE GUI, do the following:

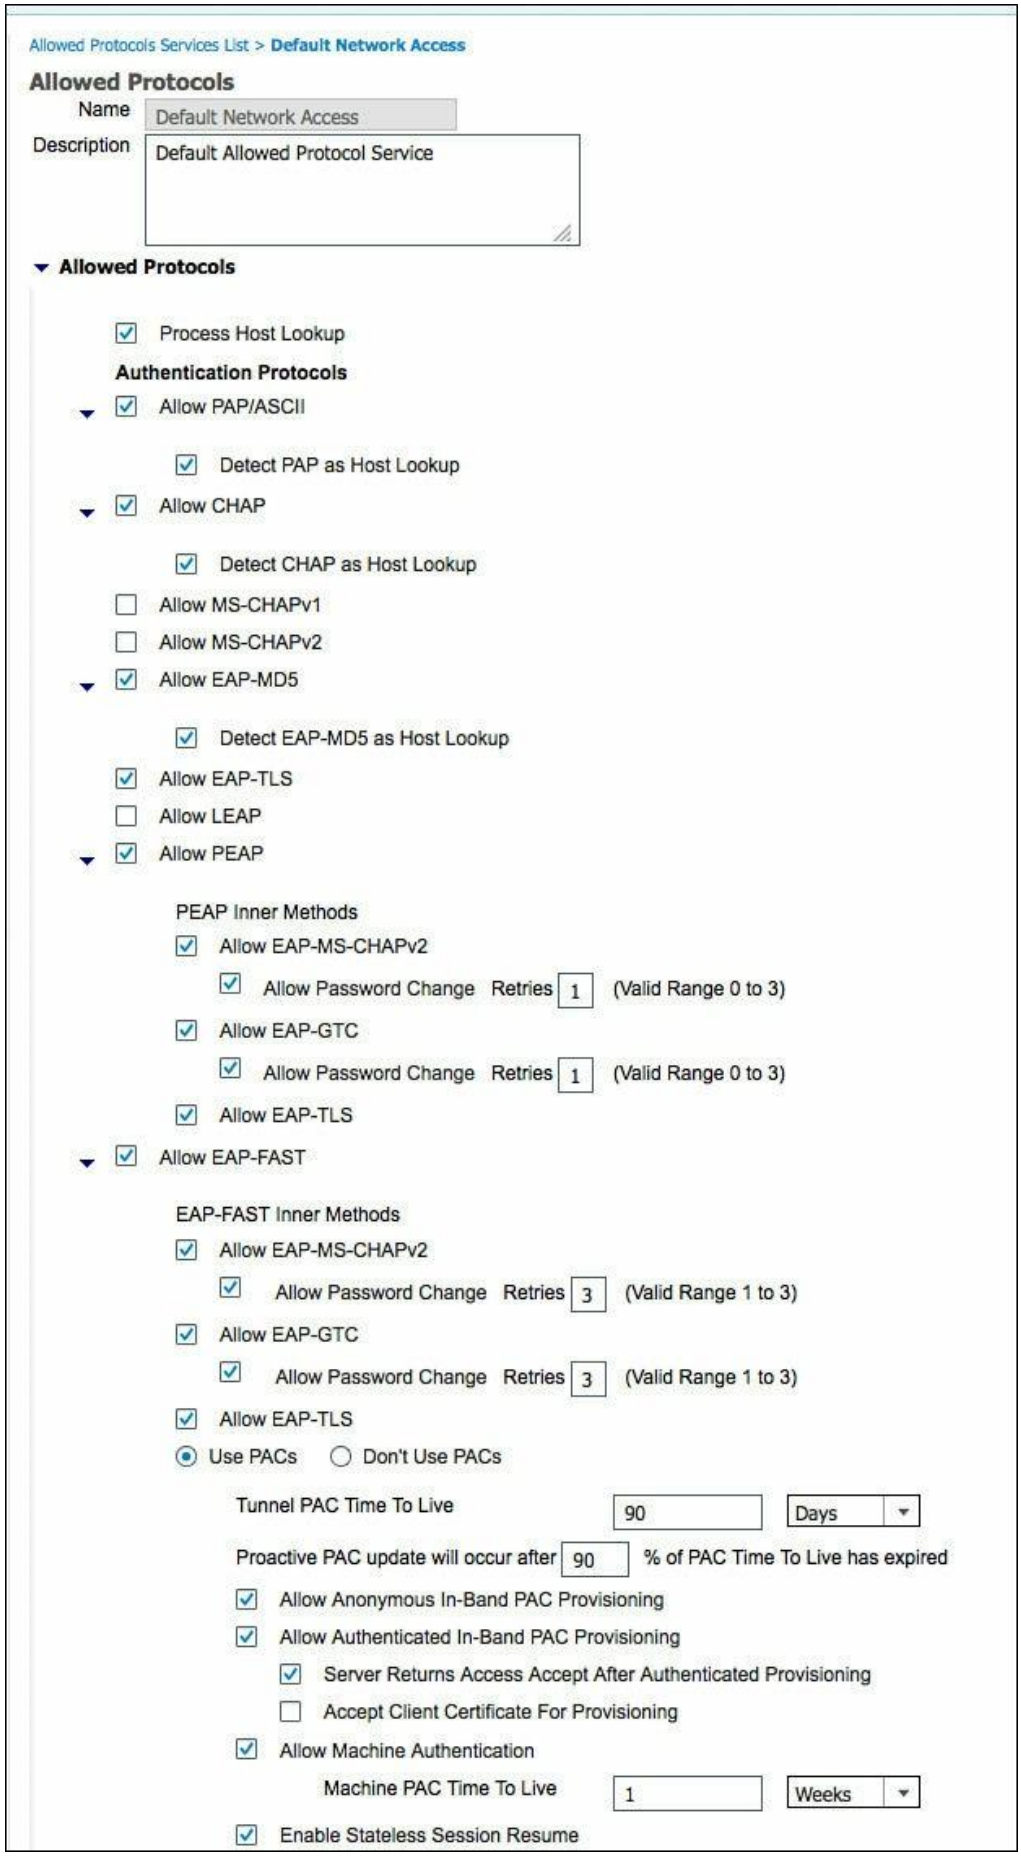
Step 1. Navigate to **Policy > Policy Elements > Results > Authentication > Allowed Protocols**.

Step 2. Select **Default Network Access**.



As Figure 10-6 shows, the list of supported protocols and their options is very extensive. This default list is inclusive with the intention of making deployments work easily for customers, but security best practice is to lock this down to only the protocols needed for that rule. Be sure to elect the protocols that are consistent with your corporate security policy, ensuring that the most secure protocol is chosen for each particular application.

https://t.me/learningnets



https://t.me/learningnets

Figure 10-6 Default network access.

Let's examine the main authentication (most common) protocols and their uses, so you will be able to create a more specific list of allowed protocols for your deployment. We will follow Figure 10-6, from top down:

- **PAP**—Password Authentication Protocol. Username is sent in the clear; password is optionally encrypted. PAP is normally used with MAB, and some devices use PAP for web authentications. We recommend you enable this for the MAB rule only and disable PAP for any authentication rules for real authentications. The check box for Detect PAP as Host Lookup enables PAP authentications to access the Internal Endpoints Database. Without this check box selected, MAB would not work.
- **CHAP**—Challenge Handshake Authentication Protocol. Usernames and passwords are encrypted using a challenge sent from the server. Challenge Handshake Authentication Protocol (CHAP) is not often used with network access; however, some vendors will send MAB using CHAP instead of PAP. The check box for Detect CHAP as Host Lookup enables CHAP authentications to access the Internal Endpoints Database. Without this check box selected, MAB would not work.

Extensible Authentication Protocol Types

Extensible Authentication Protocol (EAP) is an authentication framework providing for the transport and usage of identity credentials. EAP encapsulates the usernames, passwords, and certificates that a client is sending for purposes of authentication. There are many EAP types, each one with its own benefit and downside. As an interesting side note, 802.1X defines EAP over LAN. Here are the variations:

- **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 whether the credentials were accurate. However, EAP-MD5 does not have a mechanism for mutual authentication. That means the server is validating the client, but the client does not authenticate the server (that is, it does not check to see whether it should trust the server). EAP-MD5 is common on some IP phones, and it is also possible that some switches will send MAB requests within EAP-MD5.

  The check box for Detect EAP-MD5 as Host Lookup enables EAP-MD5 authentications to access the Internal Endpoints Database. Without this check box selected, MAB would not work.
- **EAP-TLS**—An EAP type that uses Transport Layer Security (TLS) to provide the secure identity transaction. This is similar to SSL and the way encryption is formed between your web browser and a secure website. EAP-TLS 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 because password capture is not an option; the endpoint must still have the private key. EAP-TLS is quickly becoming the EAP type of choice when supporting BYOD in the Enterprise.

Tunneled EAP Types

The previously mentioned EAP types transmit their credentials immediately. These next two EAP types form encrypted tunnels first and then transmit the credentials within the tunnel. Figure 10-7 illustrates the tunneled EAP.

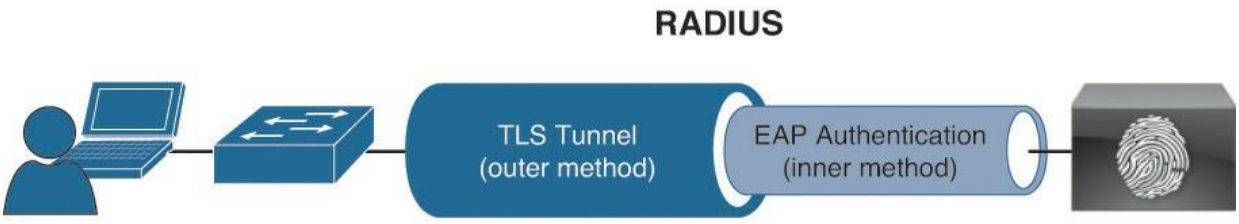


Figure 10-7 Tunneled EAP types (PEAP and FAST).

- **PEAP**—Protected EAP. 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

https://t.me/learningnets

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.

- **EAP-MSCHAPv2**—Using this inner method, the client's credentials are sent to the server encrypted within an MSCHAPv2 session. This is the most common inner method because it enables simple transmission of usernames and passwords, or even computer names and computer passwords to the RADIUS server, which in turn will authenticate them to Active Directory.
- **EAP-GTC**—EAP-Generic Token Card (GTC). This inner method was created by Cisco as an alternative to MSCHAPv2 that allows generic authentications to virtually any identity store, including one-time-password (OTP) token servers, LDAP, Novell E-Directory, and more.
- **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 Tunnel (FAST) is similar to PEAP. FAST was created by Cisco as an alternative to PEAP that allows for faster reauthentications and support for 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 the ability to use protected access credentials (PACs). A PAC can be thought of like a secure cookie, stored locally on the host as proof of a successful authentication.
  - **EAP-MSCHAPv2**—Using this inner method, the client's credentials are sent to the server encrypted within an MSCHAPv2 session. This is the most common inner method because it enables simple transmission of usernames and passwords, or even computer names and computer passwords to the RADIUS server, which in turn will authenticate them to Active Directory.
  - **EAP-GTC**—This inner method was created by Cisco as an alternative to MSCHAPv2 that allows generic authentications to virtually any identity store, including OTP token servers, LDAP, Novell E-Directory, 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**—As an enhancement to EAP-FAST, a differentiation was made to have a user PAC and a machine PAC. After a successful machine authentication, ISE will issue a machine PAC to the client. Then when processing a user authentication, ISE will request 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; it is known as EAP chaining.

The IETF has recently published RFC 7170, a new open standard for Tunnel Extensible Authentication Protocol (TEAP), which is based on EAP-FASTv2. At the time of this book publishing, the RFC was brand-new and no known vendors have adopted TEAP yet. It is expected to take the industry by storm, providing the dual authentication for enterprises.

Identity Store

After processing the allowed protocols, the authentication request is then authenticated against the chosen identity store, or in this case with MAB it is compared to the internal endpoints database (list of MAC addresses stored locally on ISE).

If the MAC Address is known, meaning it's present in the provided endpoint database, it is considered to be a successful MAB (notice this did not say successful "authentication"). MAB is exactly that—bypassing authentication—and it is not considered a secure authentication.

The selected identity source can also be an identity source sequence, which will try a series of identity

https://t.me/learningnets

stores in order. This is covered in more detail in Chapter 21, "ISE Scale and High Availability."

Options

Every authentication rule has a set of options that are 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. The options are Drop, Reject, and Continue:

- **Reject**—Send Access-Reject back to the NAD
- **Continue**—Continue to the authorization policy regardless of authentication pass/fail (used with web authentication)
- **Drop**—Do not respond to the NAD because NAD will treat as if RADIUS server is dead

Please see Chapters 20–23 for more on when to use these options.

Common Authentication Policy Examples

In this section, you will see a few quick examples of authentication policies based on common use case, or simply because they were interesting.

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 the policy is not difficult; what becomes challenging is the identification of the attribute to use because "Source-SSI" is not a field in a RADIUS packet. The attribute we need to use is `called-station-id`. That is the field that describes the wireless SSID name.



For this example, we will build a rule for an SSID named CiscoPress. This rule will be configured to:

- Only match authentications coming from that SSID
- Allow only EAP-FAST authentications
- Utilize EAP chaining
- Authenticate against Active Directory

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authentication**.

Step 2. Insert a new rule above the preconfigured Dot1X rule.

Step 3. Provide a name for the rule. In this case, we named it CiscoPress SSID.

Step 4. For the condition, select **RADIUS > Called-Station-ID**.

Step 5. Select **Contains**.

Step 6. Type the SSID name in the text box; Figure 10-8 shows the condition.

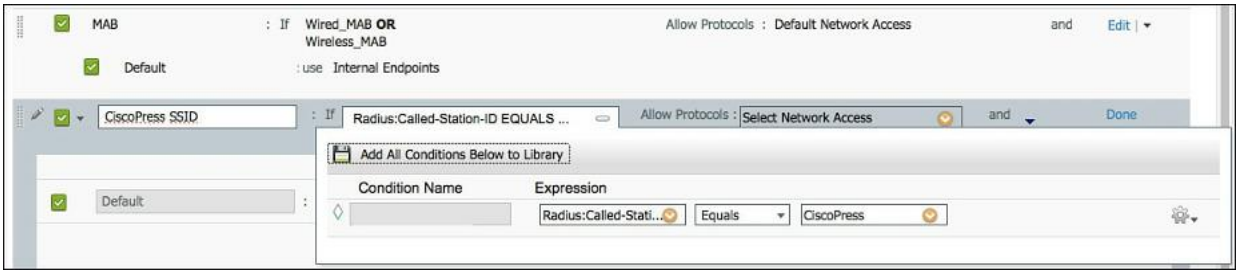


Figure 10-8 `Called-Station-ID` contains CiscoPress.

https://t.me/learningnets

Step 7. Next, we will create a new allowed protocol object that allows only EAP-FAST. Select the drop-down list for **Allowed Protocols**.

Step 8. Click the cog in the upper-right corner, and select **Create a new Allowed Protocol**, as shown in Figure 10-9.

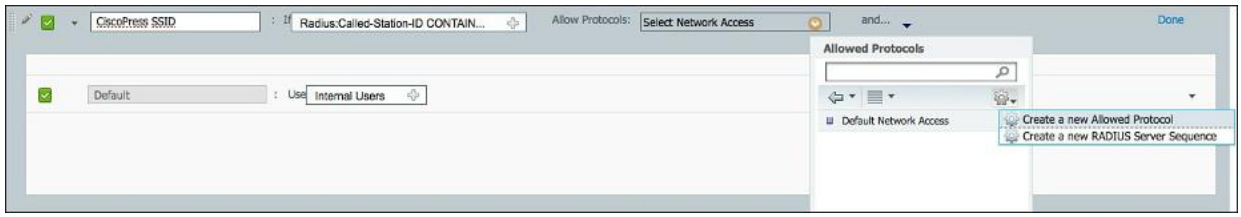


Figure 10-9 Create a new allowed protocol.

Step 9. Provide a name. In this case, we named it EAP-FAST ONLY.

Step 10. Optionally, provide a description.

Step 11. Working top-down, ensure all the check boxes are unchecked until you reach Allow EAP-FAST.

Step 12. Ensure Allow EAP-FAST is enabled.

Step 13. For ease of use, enable EAP-MS-CHAPv2, EAP-GTC, and EAP-TLS for inner methods.

Step 14. Select **Use PACs** as shown Figure 10-10 for faster session reestablishment and to allow EAP chaining.

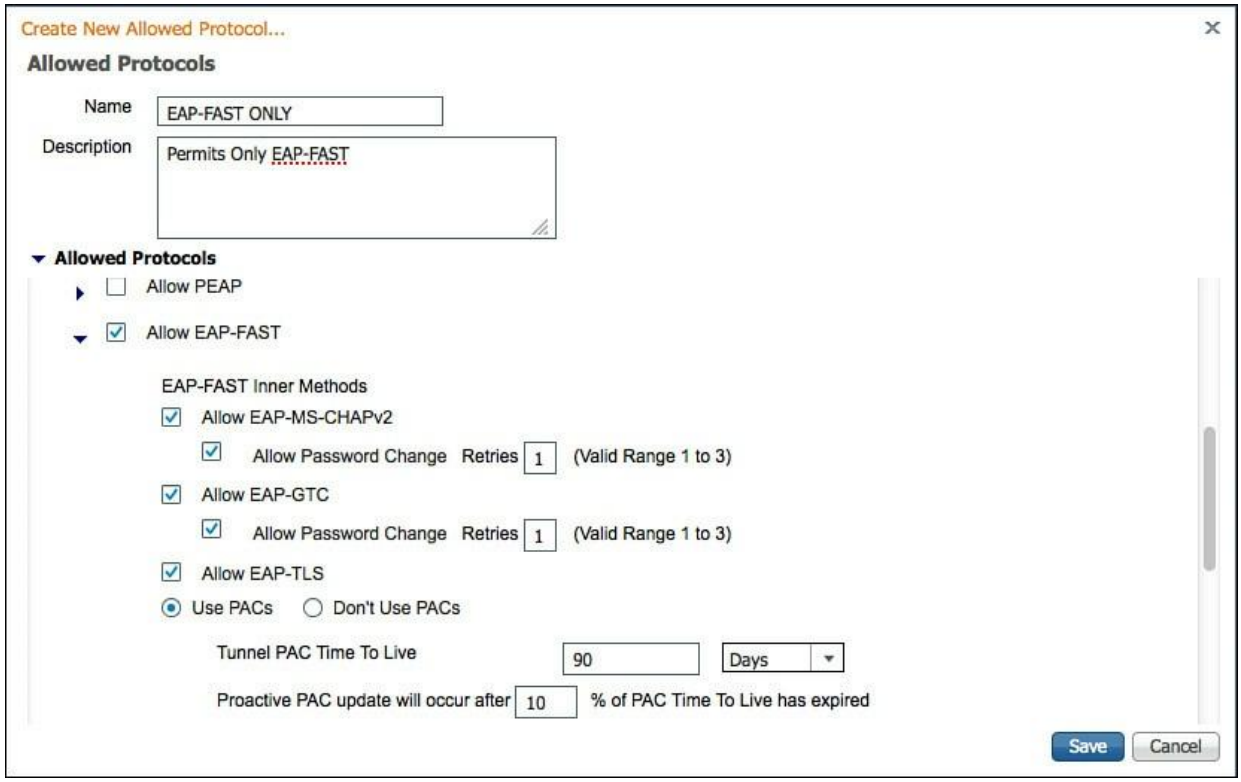


Figure 10-10 Allowed protocols.

Step 15. For ease of deployment, select **Allow Anonymous in-Band PAC Provisioning** and **Allow Authenticated in-Band PAC Provisioning**.

Step 16. Check the boxes for **Server Returns Access-Accept After Authenticated Provisioning** and **Accept Client Certificate for Provisioning**.

Step 17. Enable **Allow Machine Authentication**.

https://t.me/learningnets

Step 18. Select **Enable Stateless Session Resume**.

Step 19. Select **Enable EAP Chaining**.

Steps 15–19 are displayed in Figure 10-11.

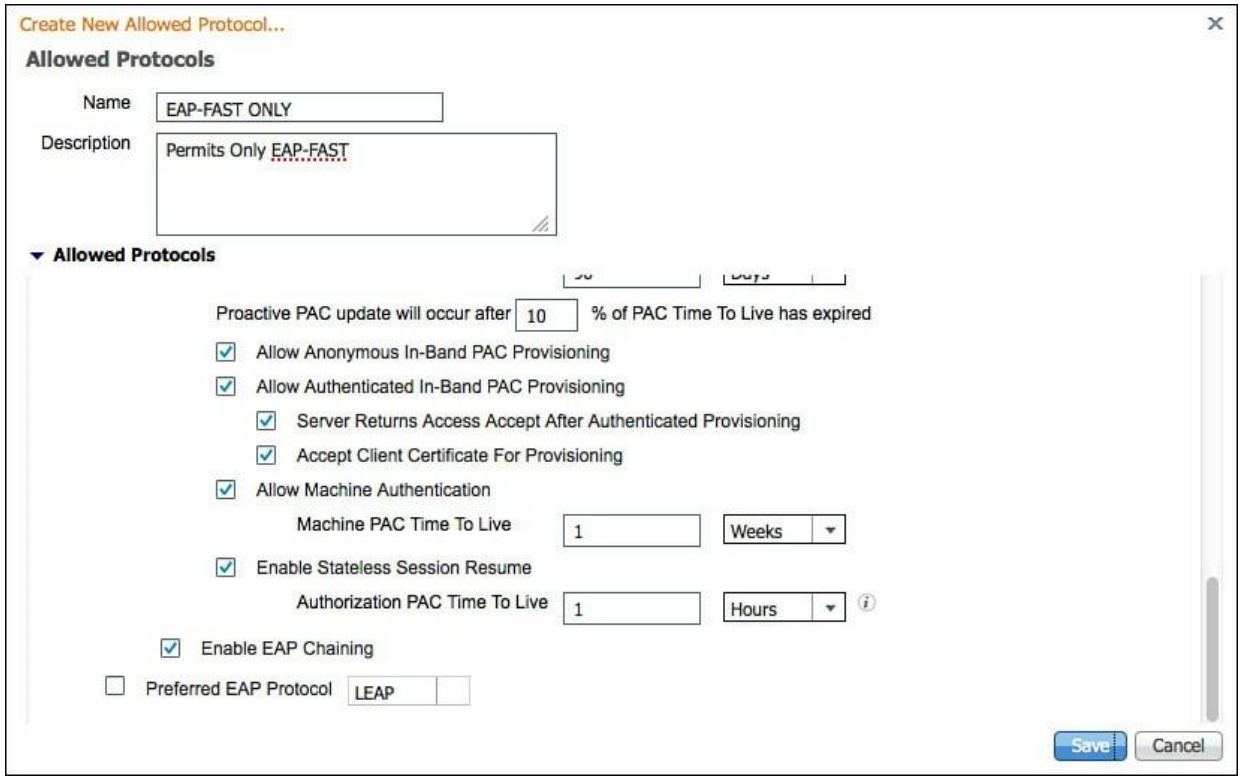


Figure 10-11 Allowed protocols, continued.

Step 20. Because we allowed only one protocol, there is no need to set a preferred EAP Protocol. Click **Save**.

Step 21. Even though we created the allowed protocol object for this specific authentication rule, it doesn't automatically select it from the drop-down list. Select the drop-down list for the identity source (currently set for Internal Users).

Step 22. Select your AD source; in this case, the name is **AD1**, as shown in Figure 10-12.

https://t.me/learningnets

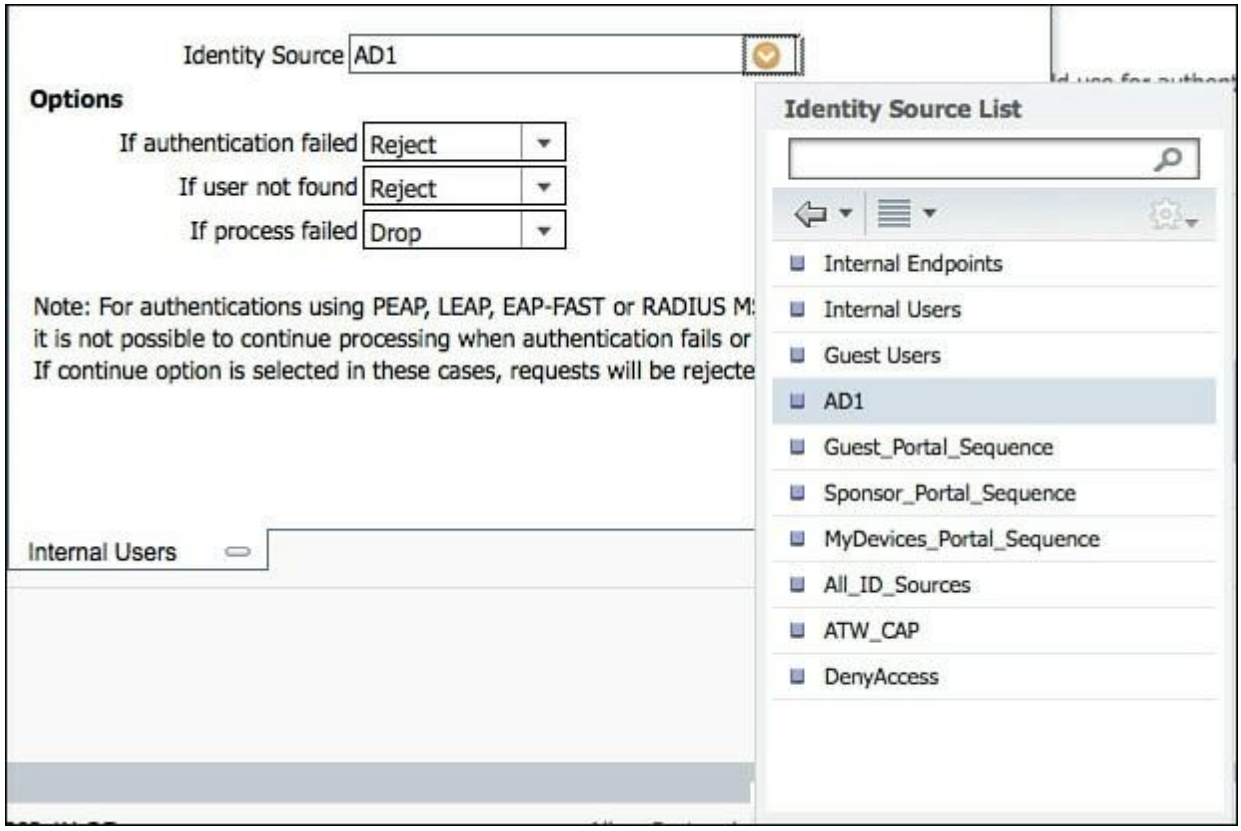


Figure 10-12 Selecting the AD identity source.

Step 23. Leave the default Options, and click **Done**.

Step 24. Click **Save**.

Figure 10-13 shows the completed authentication rule.

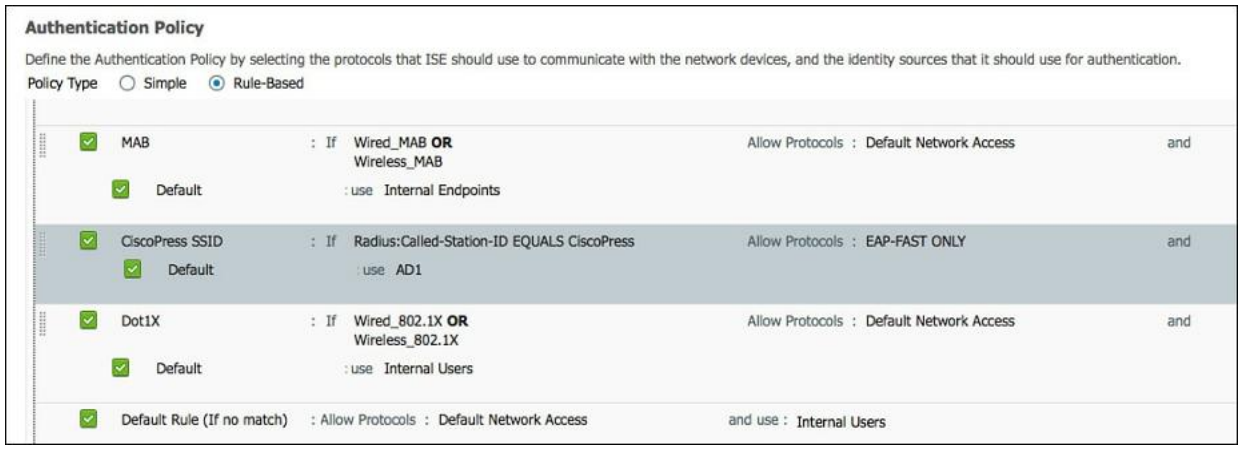


Figure 10-13 Completed authentication rule.

This completes the creation of the authentication rule. Determining which actions to take for the authentications that passed will be handled in the authorization policy.

Remote Access VPN

Often authentications for a remote access VPN connection get routed to an OTP server, such as RSAs SecureID. For this example, we will build a rule for remote access VPN authentications. This rule will be configured to:

- Only match authentications coming from the VPN device

https://t.me/learningnets

■ Route that authentication to an OTP server

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authentication**.

Step 2. Insert a new rule above the preconfigured Dot1X rule.

Step 3. Provide a name for the rule. In this case, we named it RA VPN.

Step 4. For the condition, select **Create New Condition (Advanced Option) > DEVICE > Device Type**.

Step 5. Set the operator to Equals.

Step 6. Select the Network Device Group VPN.

Step 7. Save the selection as a condition by clicking the cog on the right side and then selecting **Add Condition to Library**. Name the condition **VPN-Device**.

Step 8. Click the green check mark.

Figure 10-14 shows the completed condition.

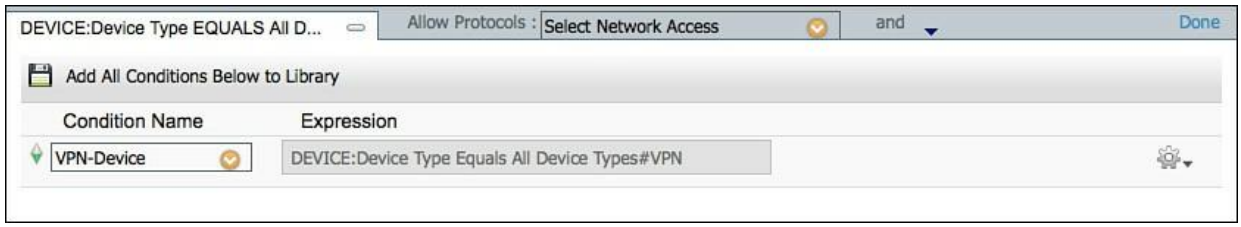


Figure 10-14 Device type equals VPN.

Step 9. For this example, we will just use the allowed protocol of Default Network Access. Click the drop-down list and select **Allowed Protocols** and then **Default Network Access**.

Step 10. For the identity store, we have selected the OTP server (previously configured) in **Administration > Identity Management > External Identity Sources > RADIUS Token (ATW_OTP)**.

Step 11. We are leaving the default options; click **Done**.

Step 12. Click **Save**.

Figure 10-15 shows the completed authentication rule.

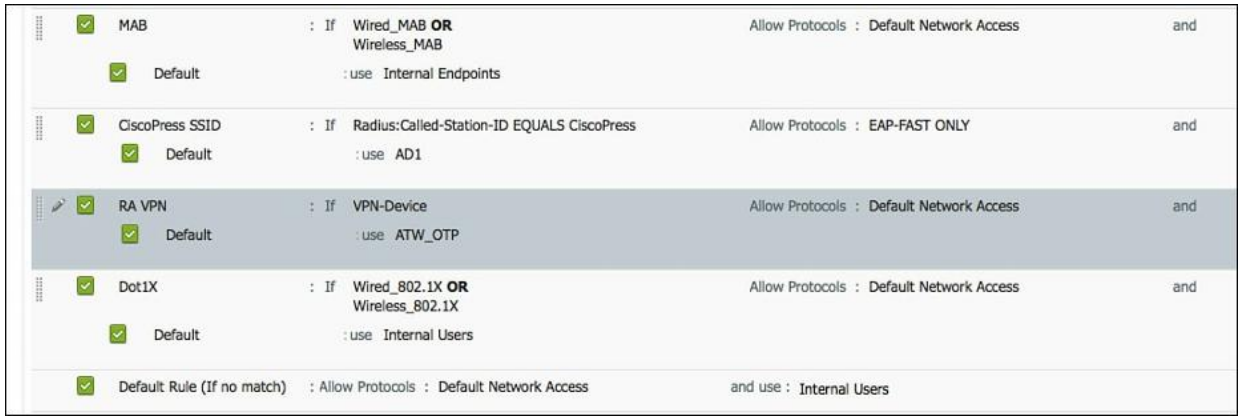


Figure 10-15 Completed authentication rule.

https://t.me/learningnets

Alternative ID Stores Based on EAP Type

In this modern day of BYOD and mobility, it is common to have multiple user and device types connecting to the same wireless SSID. In scenarios like this, often the users with corporate laptops authenticate using EAP-FAST with EAP chaining; while BYOD-type devices must use certificates and EAP-TLS. Anyone authenticating with PEAP would be recognized as a noncorporate and nonregistered asset and be sent to a device registration portal instead of being permitted network access.

For this example, we will modify the preconfigured Dot1X rule by creating sub-rules for each EAP type. This rule will be configured to:

- Match wired or wireless 802.1X
- Route EAP-TLS authentications to a certificate authentication profile (CAP)
- Route PEAP authentications to an LDAP server
- Route EAP-FAST to Active Directory
- Route EAP-MD5 to internal endpoints for host lookup as a MAB request

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authentication**.

Step 2. Edit the preconfigured Dot1X rule.

Step 3. Next, we will create a new allowed protocol object that allows only EAP authentications. Select the drop-down list for **Allowed Protocols**.

Step 4. Click the cog in the upper-right corner, and select **Create a New Allowed Protocol**.

Step 5. Provide a name. In this case, we named it All EAP Types.

Step 6. Optionally, provide a description.

Step 7. Working top-down, ensure all EAP types are enabled, except for LEAP (unless you need LEAP for backward compatibility).

Step 8. Enable EAP chaining, as we did previously in the wireless SSID exercise.

Step 9. Click **Save**.

Step 10. Select the All EAP Types allowed protocol object for this authentication rule.

Step 11. Insert a new sub-rule above the Default identity store sub-rule, and name it **EAP-TLS**.

Step 12. For the condition, select **Create a New Condition (Advanced Option) > Network Access > EapAuthentication Equals EAP-TLS** (as shows in Figure 10-16).

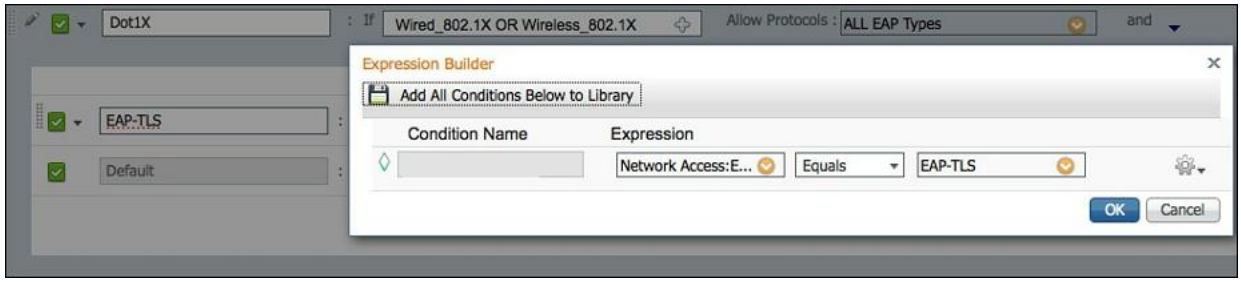


Figure 10-16 Network access: EapAuthentication equals EAP-TLS.

Step 13. For the identity source, we are choosing a preconfigured CAP. This was configured at **Administration > Identity Management > External Identity Sources > Certificate Authentication Profile**.

Step 14. Insert a new row above the EAP-TLS row, to insert EAP-FAST. We are placing EAP-FAST above EAP-TLS because EAP-TLS can be used as an inner method of EAP-FAST.

Step 15. Select **Network Access > EapTunnel Equals EAP-FAST** for the condition.

https://t.me/learningnets

Step 16. Select the **Active Directory Object** for the identity source.

Step 17. Insert a new row above the EAP-TLS row to insert PEAP.

Step 18. Select **Network Access > EapTunnel Equals PEAP** for the condition.

Step 19. Select the **LDAP object** for the identity source.

Step 20. Insert a new row below the EAP-TLS row to insert EAP-MD5.

Step 21. Select **Network Access > EapAuthentication Equals EAP-MD5** for the condition.

Step 22. Select **Internal Endpoints** for the identity source.

Step 23. Change the default identity store (bottom row) to be **Deny Access**.

Step 24. Click **Done**.

Step 25. Click **Save**.

Figure 10-17 shows the completed rule and sub-rules.

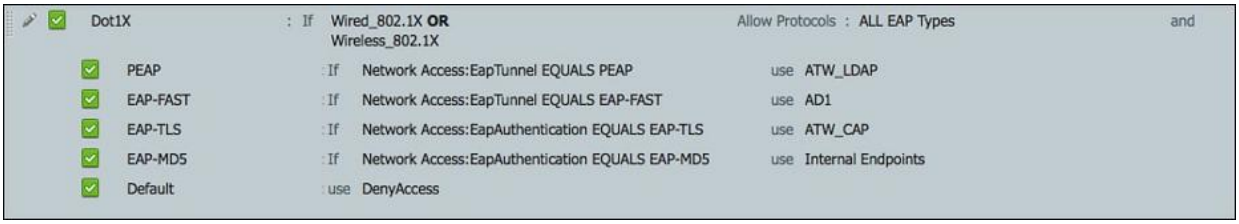


Figure 10-17 Completed authentication rule and Sub-rules.

More on MAB

One of the things that is often not understood, especially when looking to mix access device vendors, is MAB. There is no standard for MAB. Different vendors implement MAB in different ways. Ultimately, the goal is to allow the supplicant in the switch itself to run an authentication request for the endpoint because the endpoint obviously must not have a supplicant.

Some vendors send a RADIUS service-type of Login; some send a RADIUS service-type of Framed. Cisco uses a service-type of Call-Check for MAB. Why would Cisco use Call-Check if no other vendor does? Why does Cisco do MAB differently from everyone else? Quick answer: security.



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 and the same format. There was no database separation of user IDs from endpoint IDs (MAC addresses). As displayed in Figure 10-18, a malicious user could enter a MAC address into the username and password fields of a web authentication or maybe even into the endpoint supplicant and gain access to the network.

https://t.me/learningnets

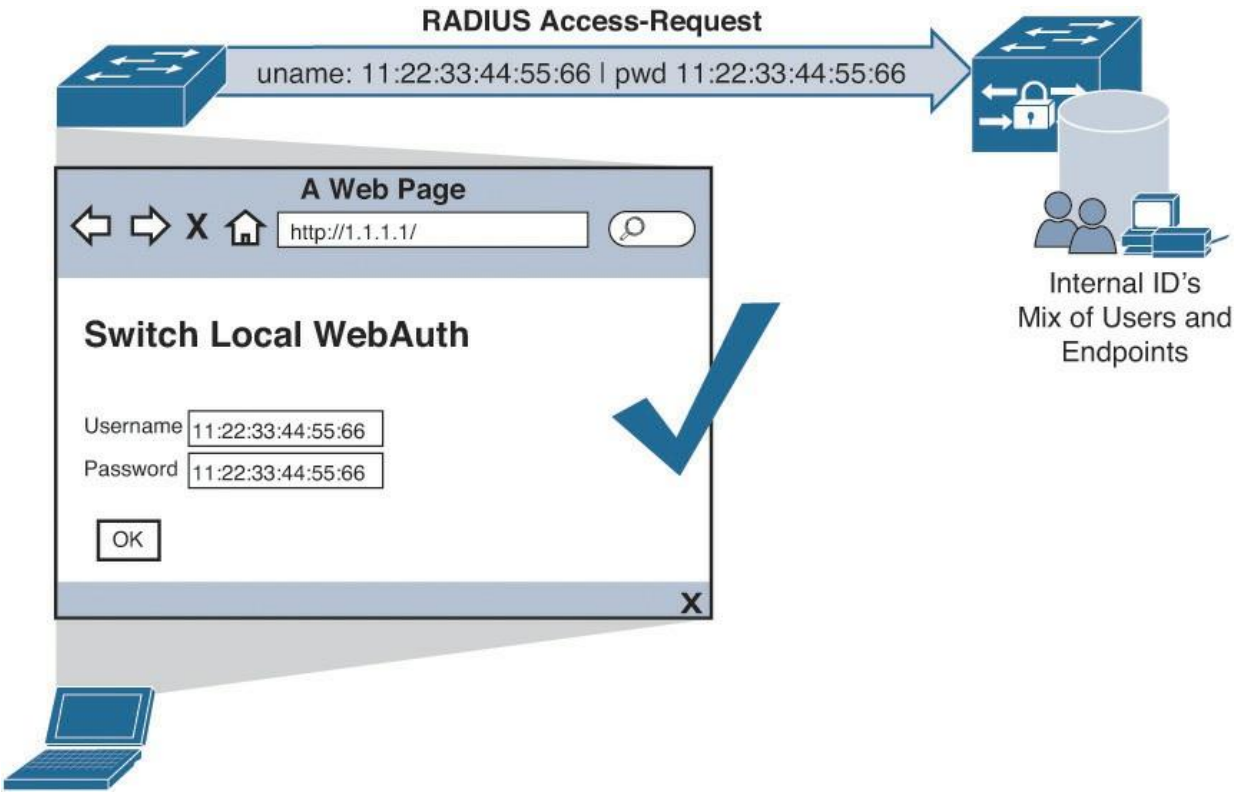


Figure 10-18 Web authentication with MAC address instead of username.

In an effort to close this security hole and make MAB a bit more secure, Cisco changed the way it does MAB. The key differences are listed here:

- For authentication requests to be processed as MAB (by default), the service type must be Call-Check.
- RADIUS servers (ACS and ISE) maintain a separate endpoint database.
- The `calling-station-id` is the value that will be compared to the endpoint database, ignoring the username and password fields of the MAB request.

All supported Cisco NADs use a service type of Call-Check for MAB requests. They also ensure the `calling-station-id` is populated with the MAC address of endpoint. Lastly, Cisco ISE uses a simple check box within the allowed protocols configuration as another method to permit or deny the access into the endpoint database for the MAB request, as shown in Figure 10-19.

https://t.me/learningnets

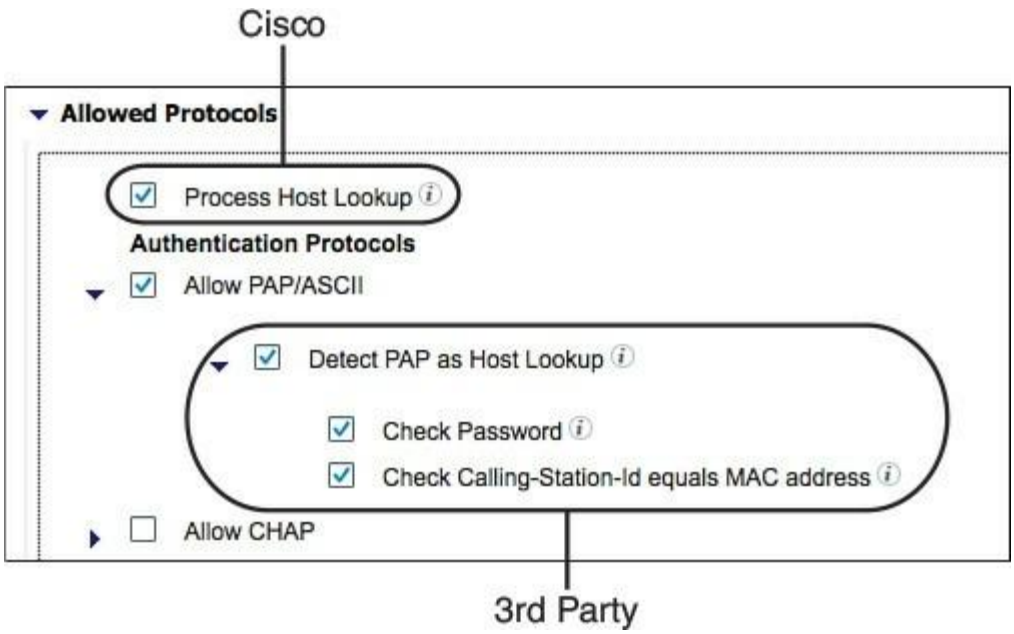


Figure 10-19 Process host lookup.

As Figure 10-19 shows, the top selection for Process Host Lookup is the one for Cisco network devices. That check box allows RADIUS authentications with a service type of Call-Check to have the RADIUS `calling-station-id` value compared with the contents of the endpoints database. The selection for Process Host Lookup also exists under each of the individual authentication protocols (such as PAP, CHAP, and EAP-MD5). These are there for third-party support and are the reason there are two other check boxes: Check Password and Check Calling-Station-Id Equals MAC Address.

These check boxes make an insecure mechanism such as MAB a bit more insecure, so it is recommended that you secure it as much as possible by only allowing the network devices that must use MAB in the less secure manner to use it in that manner. This topic is discussed further in the successful deployment strategies section(s) of this book.

Keep in mind that MAB is inherently not a secure technology. When implementing 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 to the networks and services that device is required to speak to only. In other words, don't provide full access to devices that have been MAB'd; instead provide them with an authorization that is more limited. This topic is covered in more detail in the next chapter where authorization policies are covered.

Restore the Authentication Policy

In this chapter, you have created a complex and specific authentication policy. This is useful for learning how authentication policies work, but it might make things a bit too complicated for you as you navigate through the future chapters.

To keep things simple, follow these steps to restore your authentication policy to something simple that will work for all use cases remaining in this book.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authentication**.

Step 2. Delete the rule named CiscoPress SSID.

Step 3. Edit the **r**ule named Dot1X.

Step 4. Delete the PEAP, EAP-FAST, EAP-TLS, and EAP-MD5 rules.

Step 5. Change the Dot1X > Default rule from Deny Access to **All ID_Sources**.

https://t.me/learningnets

Figure 10-20 shows the final authentication policy that will enable you to support the use cases in the remainder of the book in an easy way.

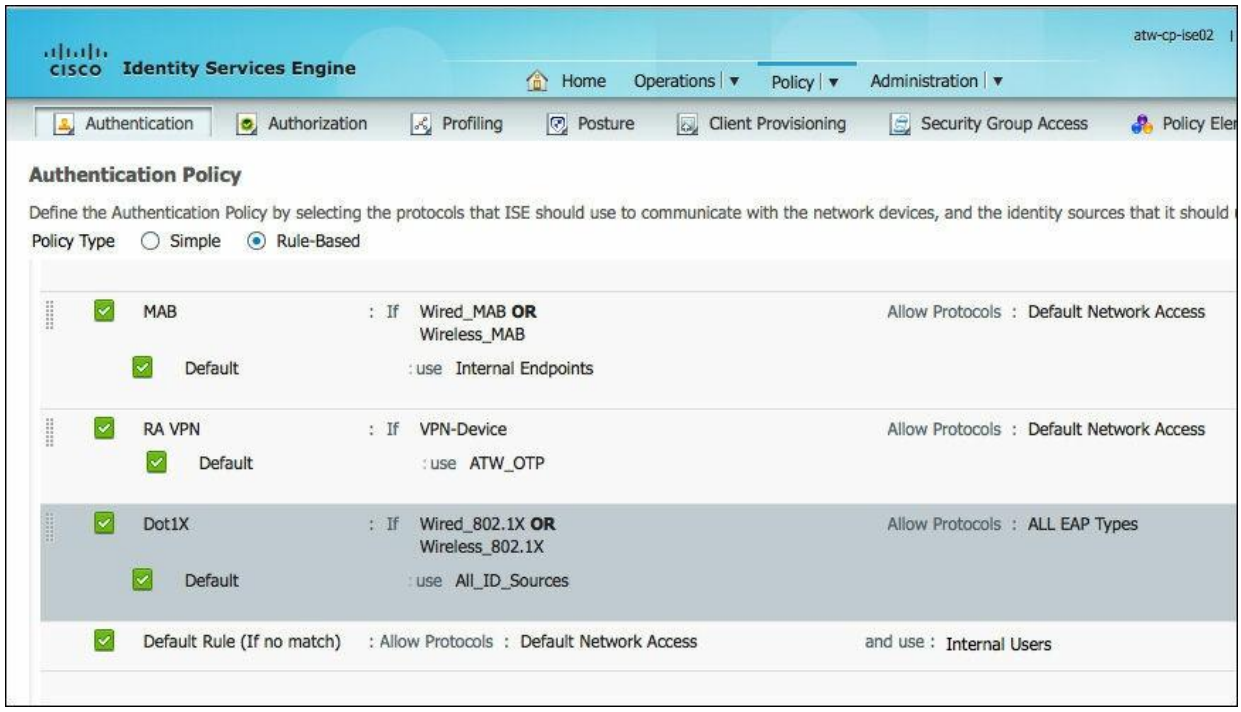


Figure 10-20 Simplified authentication policy.

This completes the authentication chapter. In the next chapter we take an in-depth look at authorization policies and common authorization rules.

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 10-2 lists a reference of these key topics and the page numbers on which each is found.

Key Topic

Key Topic Element	Description	Page
Paragraph	Authentication and authorization	237
List	Allowed protocols	243
Paragraph	Matching an SSID	248
Paragraph	MAC authentication bypass	255

Table 10-2 Key Topics for Chapter 10

https://t.me/learningnets

Chapter 11. Authorization Policies

This chapter covers the following exam topics:

- **Implement Wired/Wireless 802.1X**
- **AV Pairs**
- **Implement MAB**
- **ISE Authentication/Authorization Policies**
- **Implement Network Authorization Enforcement**
- **Network Authorization Enforcement**

As discussed in Chapter 10, "Authentication Policies," an authentication is the validation of a credential or an identity. The other major portion of secure network access control is the actual control portion. That is where authorization policies come in.

Consider the bank example: walk into a bank and provide a valid identification. It's a government-issued passport. It checks out completely—the seal is valid, the picture matches. This was a successful authentication. However, what does this successful authentication entitle you to withdraw from the bank? Does it allow you to withdraw money from any account in the bank?

This entitlement would also be known as the 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 also what Cisco calls the "Security Context" of the user or device.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
AV Pairs	2, 6
ISE Authentication/Authorization Policies	1, 3-5, 7-8
Implement Network Authorization Enforcement	1-4
Network Authorization Enforcement	1-3, 9

Table 11-1 "Do I Know This Already?" Section-to-Question Mapping

> **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?

https://t.me/learningnets

a. An authorization profile is a rule in the policy table that is formatted like “IF condition THEN result.”

b. An authorization profile is created to determine which identity store to validate the credentials with.

c. An authorization profile is a sequential list of identity stores to validate the credentials with.

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 options are part of the common tasks section of 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. Each rule will have at least one authorization profile.

b. An authorization rule contains authorization policies. Each policy will have at least one authorization profile.

c. An authentication policy contains authorization rules. Each rule must have an authentication result.

d. An authentication rule contains the authorization profiles. Each profile must contain one authentication result.

5. True or False? Condition attributes can be saved into a library for future use and improved readability.

a. True

b. False

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 a simple condition and compound condition?

a. Simple conditions are easier to use than compound conditions.

https://t.me/learningnets

b. Simple conditions are created on-the-fly within the expression builder, while compound conditions must be created separately.

c. Simple conditions contain only one attribute. Compound conditions contain multiple attributes along with an operator such as AND or OR.

d. Simple conditions and compound conditions can each contain multiple attributes, but compound conditions can mix operators such as AND or OR.

8. True or False? A compound condition can contain a mixture of simple conditions and raw attributes.

a. True

b. False

9. What should be the end goal of a 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, and provide limited access to any failed authentications

c. To secure the network by purchasing Cisco ISE, thereby increasing the stock value of the company

d. To provide very specific permissions to any authorization, providing 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, while non-Cisco network devices cannot apply ACLs.

b. Cisco downloadable ACLs 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 the full ACL is sent down to the network device within RADIUS AV pairs, while non-Cisco network devices must create the ACL on the individual local network device.

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 from the burden of traffic control.

Foundation Topics

Authentication Versus Authorization

At the beginning of the previous chapter, you examined the relationship between authentication and authorization.

An authentication is the validating of credentials. Consider the BYOD use case, in which a tablet connects to the network 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 prove entitlement.

There is still a process that can determine the level of access to assign to the tablet. That access can be based on virtually any attribute, such as the device type, the location, the business unit, the user identity extracted from the certificate, Active Directory group membership of that user, which certificate authority (CA) signed the certificate, the time of day, or the user's hair and eye color even. The list can go on and on. The conglomeration of attributes is often referred to as the context of the device. A different level of access can be assigned as long as an authorization rule in the policy matches the context.

These rich policies are part of the authorization policy. The flexibility of being able to assign a different authorization per possible context is why most of the time that is expended working within a product like

https://t.me/learningnets

Cisco ISE is spent setting up and tuning the authorization policy. Authorization is where the bulk of the decisions are made.

Authorization Policies

After an endpoint has been successfully authenticated, the authentication process is passed to the authorization policy, where the policy is evaluated to determine the final access results.

Goals of Authorization Policies



Authorization policies have one main goal: to examine conditions to send an authorization result to the network access device (NAD). Which 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 within ISE. Policies can be built using attributes such as location, time, whether a device was registered, whether a mobile device has been jail-broken—nearly any attribute imaginable. Even the authentication is an attribute: Was authentication successful, which authentication protocol was used, and what is the content of specific fields of the certificate that was used?

The policy compares these conditions with the explicit goal of providing an authorization result. The result can be a standard RADIUS Access-Accept or Access-Reject message, and it can include more advanced things such as VLAN assignment, dACLs, security group tag (SGT), URL redirection, and more.

The result will allow or deny access to the network, and when it is allowed, it can 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 the authorization policy, it will be easier to understand the exercises of this section. To understand authorization policies even more, we will now go through and examine a few.

Basic authorization policy rules are logically organized in this manor:

Click here to view code image

```
IF [conditions] THEN [AssignThesePermissions]
```

Just like the authentication policy, authorization policy rules are processed in a top-down, first-match order. So, if the conditions do not match, the authentication will be compared to the next rule in the policy.

ISE is preconfigured with a default rule for blacklisted devices, named Wireless Blacklist Default, Profiled Cisco IP Phones, and Profiled Non Cisco IP Phones. We are going to examine the Cisco IP phone and blacklist rules to dig into authorization rules and how they work. If you have a live ISE system, it can help to follow along with the text.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authorization**. You should notice an immediate difference between the authorization policy and the authentication policy examined earlier in this chapter (see Figure 11-1). The authorization policy attempts to display the rule logic in plain English. The bold text designates an identity group, while the standard font is a normal attribute. The operator is always “AND” when both Identity Group and Other Conditions are used in the same rule.

https://t.me/learningnets

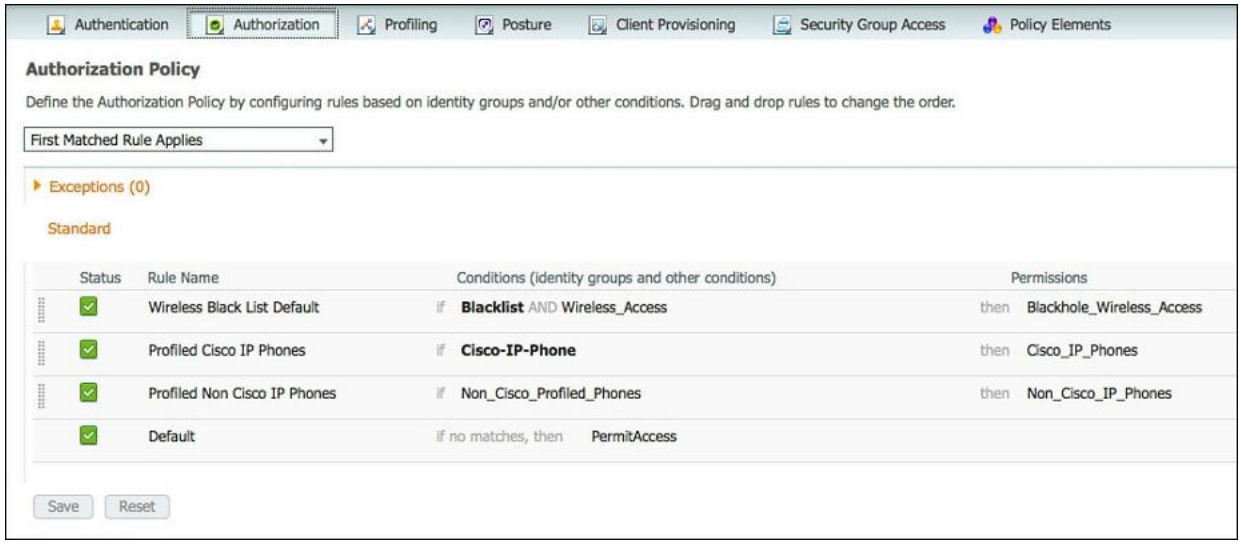


Figure 11-1 Default authorization policy.

Step 2. Edit the rule named Cisco IP Phones.

Notice the identity group is a separate list than the other conditions. In this rule, there is an identity group named Cisco-IP-Phones (see Figure 11-2). The next field is where other conditions would be selected.



Figure 11-2 Profiled Cisco IP phones.

This particular rule is a prebuilt rule that permits any device that was profiled as a Cisco IP phone, sending an Access-Accept that also sends an attribute value pair (AVP) that permits the phone into the voice VLAN.

Step 3. Let's examine the permissions (result) that are sent. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

Authorization profiles are a set of authorization results that should be sent together. Notice there are two other categories of authorization results: Downloadable ACLs and Inline Posture Node Profiles. Every authorization rule must have a result (see Figure 11-3).

https://t.me/learningnets



Figure 11-3 Default authorization profiles.

Step 4. Select the **Cisco_IP_Phones Authorization Profiles**.

The authorization result needs to be RADIUS attributes. To make that easier for the users of ISE, Cisco has included a Common Tasks section that presents the options in more of a plain-English format (see Figure 11-4). The Attributes Details section at the bottom displays the raw RADIUS result that will be sent.

https://t.me/learningnets

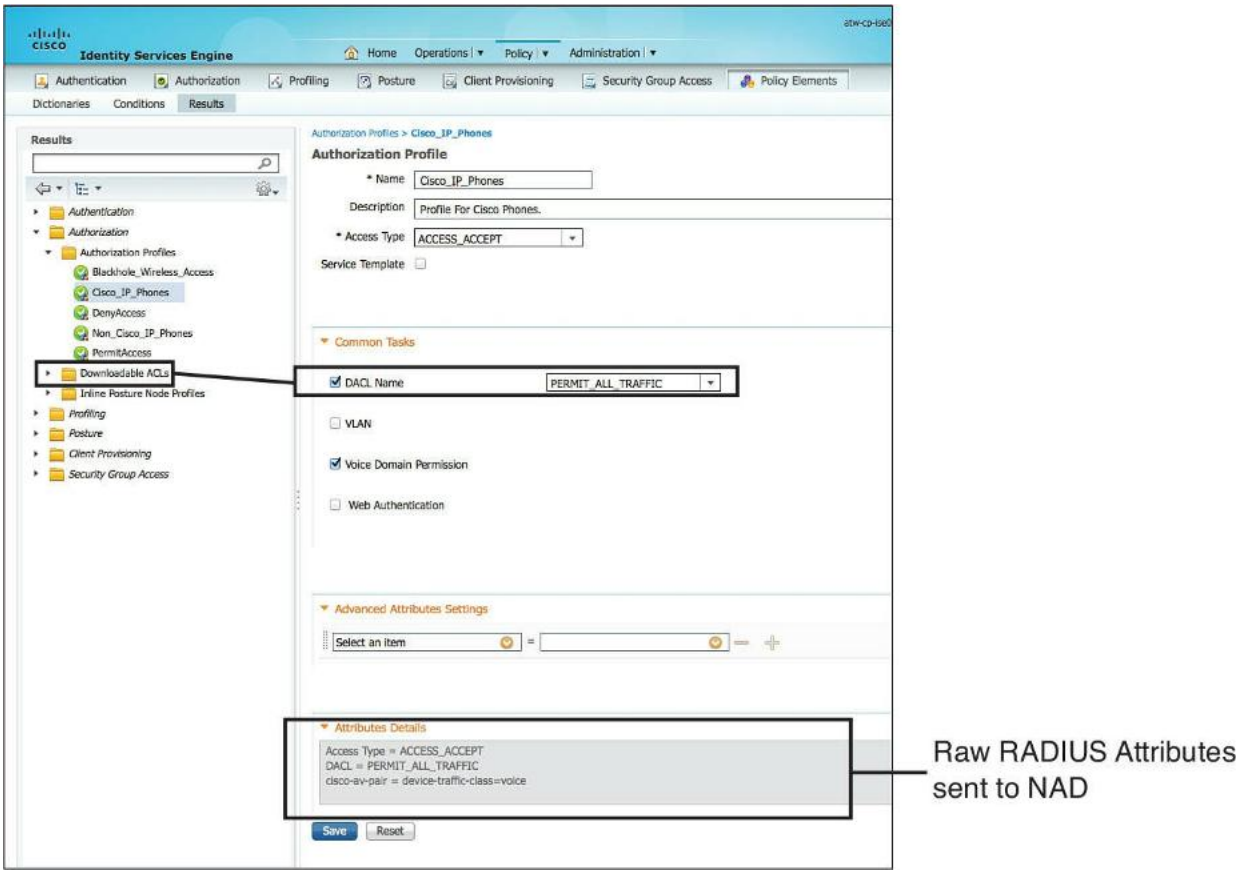


Figure 11-4 Cisco_IP_Phones authorization profile.

Take note in Figure 11-4 the DACL Name is a drop-down box where you will select a dACL that is created and stored in ISE. The Voice Domain Permission check box is required for the switch to allow the phone into the Voice VLAN on the switch.

Step 5. Notice in the Attributes Detail section that this authorization result will send a RADIUS result with an Access-Accept, a dACL that permits all traffic, and the voice domain VSA to permit the phone to join the Voice VLAN.

Next, let's examine the Wireless Black List Default Rule:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Edit the rule named Wireless Black List Default.

Notice the identity group is a separate list than the other conditions. In this rule, there is an identity group named Blacklist. The next field is populated with a prebuilt condition specifying wireless connections. This particular rule is built to prevent devices that have been marked lost or stolen from accessing the network.

Step 3. Let's examine the authorization condition being used. Navigate to **Policy > Policy Elements > Conditions > Authorization > Compound Conditions**.

Figure 11-5 shows the default list of compound conditions. A simple condition is just a single attribute that has been saved to the library for reuse. A compound condition is the joining of more than one condition with an AND or OR operator. This compound condition can also be saved to the library for reuse.

https://t.me/learningnets

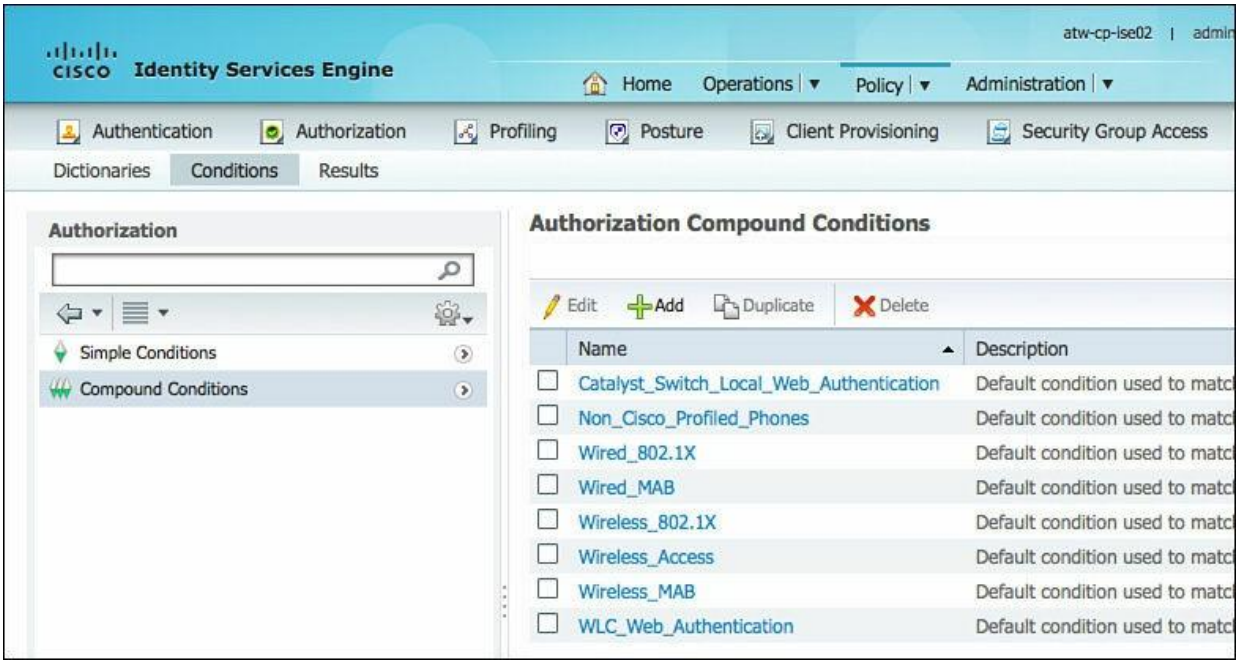


Figure 11-5 Prebuilt authorization compound conditions.

Step 4. Select **Wireless_Access**.

As shown in Figure 11-6, the Wireless_Access compound condition references the RADIUS attribute of NAS-Port-Type Equals Wireless – IEEE 802.11.

https://t.me/learningnets

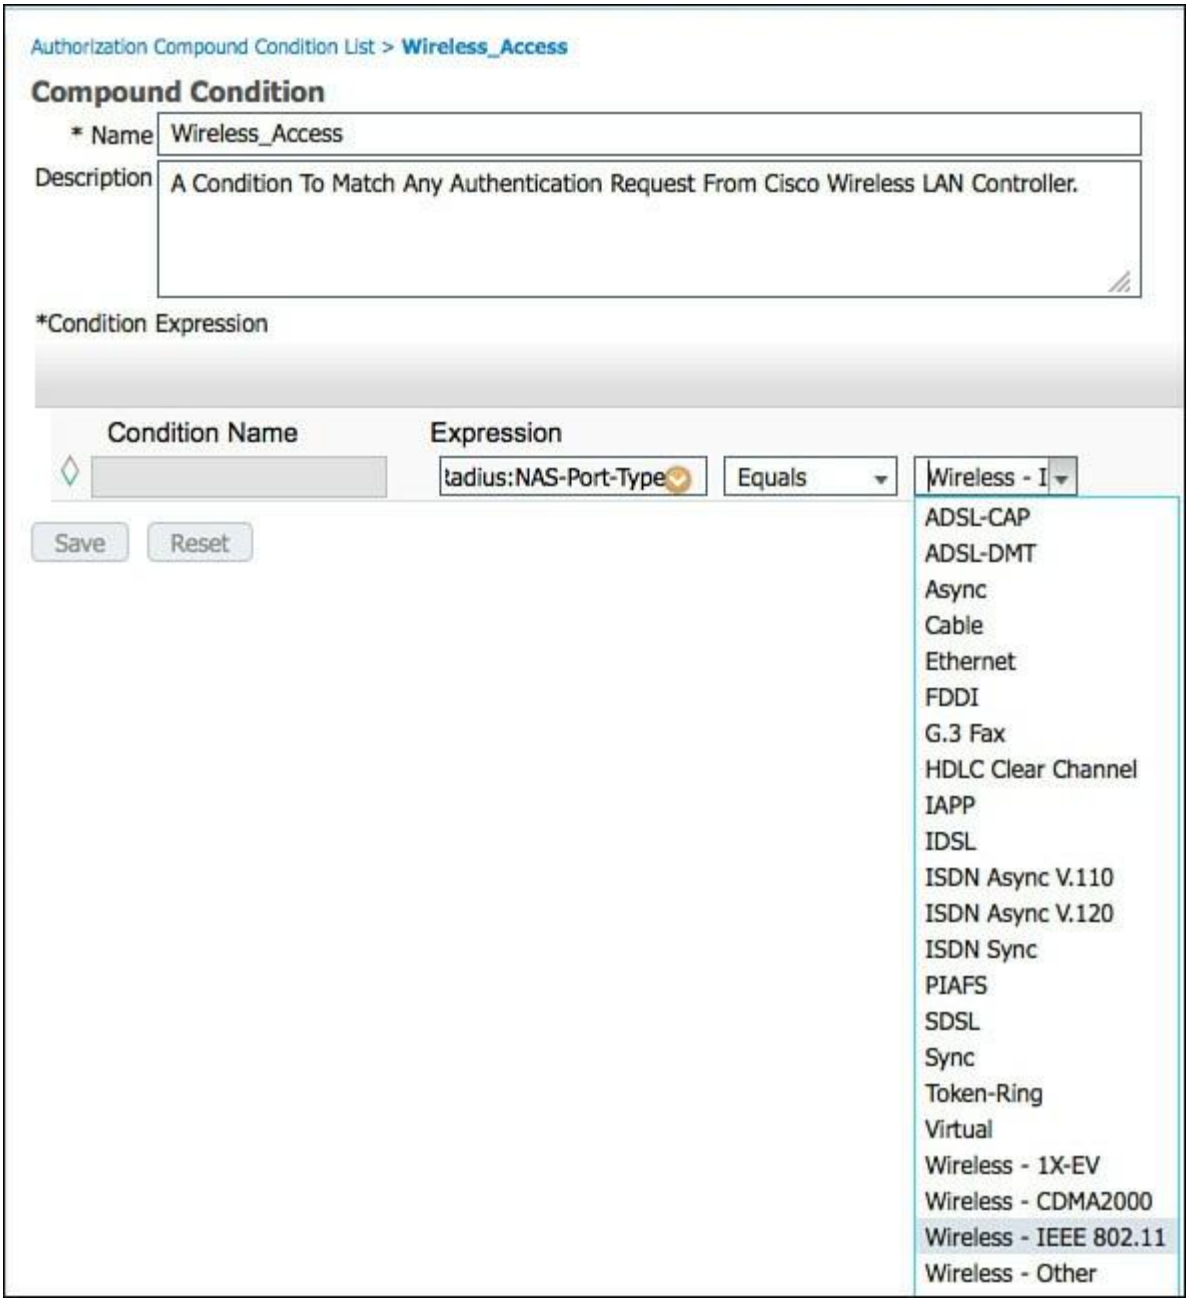


Figure 11-6 Wireless_Access compound condition.

Step 5. Now let's examine the authorization result that is being sent for this authorization rule. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

Step 6. Select **Blackhole_Wireless_Access**. As shown in Figure 11-7, the Blackhole_Wireless_Access authorization profile does not use any of the common tasks. Instead, it is using the Advanced Attribute Settings to send a URL-Redirect and URL-Redirect-ACL result to the wireless LAN controller (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 web page describing that the device was blacklisted.

https://t.me/learningnets

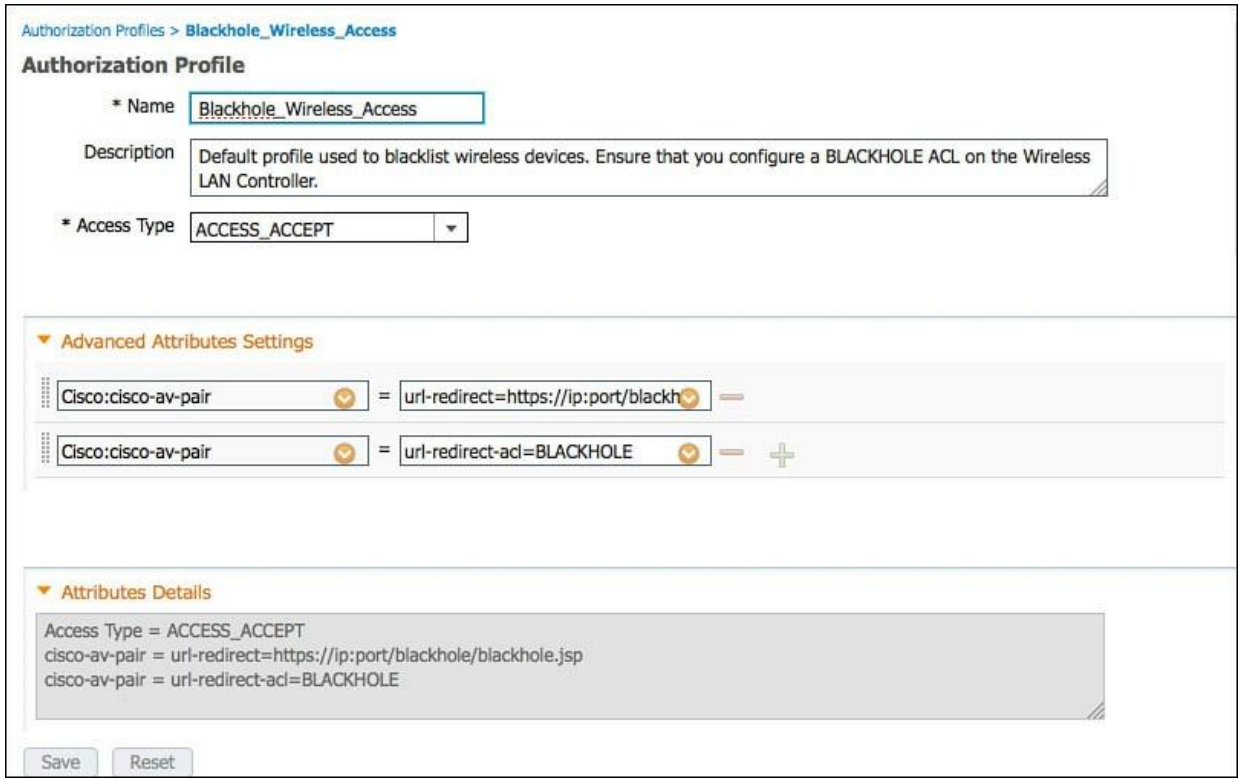


Figure 11-7 Blackhole_Wireless_Access authorization profile.

These two authorization rules demonstrate a variety of rules. Later in this chapter we will examine a few common authorization policies.

Role-specific Authorization Rules



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. Everything is dependent on the security policies built by you, the administrator.

For example, companies often use a combination of ISE profiling capabilities and MAC Authentication Bypass (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 NAD to send an authentication request containing the MAC address of the printer to ISE. From there, ISE should send down the Access-Accept result.

There was a successful MAB, and the printer needs to have access to the network, but 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, if a malicious user were to spoof the MAC address of the printer, he would not able to get to much on 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 to limit the impact to the end users.

Chapter 20, "Deployment Phasing," is dedicated to this phased approach.

https://t.me/learningnets

Authorization Policy Example

In this section, you will see an example of an authorization policy made up of numerous rules based on a common use case. This use case was selected to show multiple aspects of the authorization policy and help to solidify your working knowledge of the pieces of an authorization policy and the workflows associated with creating the policies.

For this example, we will configure three authorization rules. One assigns full access to an employee that authenticated successfully with EAP chaining followed by a rule that assigns more limited access to the same employee authenticating with a noncorporate machine. The last rule assigns Internet only access to the same employee authenticating on a mobile device.

Employee Full Access Rule

In this rule, we will assign full access permissions to an employee who is authenticating from a valid corporate asset. From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Insert a new rule above the Default rule.

Step 3. Name the new rule **`Employee and CorpMachine`**.

Step 4. For the other conditions drop-down list, click the plus sign (+) next to Conditions and select **Create New Condition**.

Step 5. Select **Network Access > EapChainingResult**.

Step 6. Select **Equals**.

Step 7. Select **User and Machine both Succeeded**.

Step 8. Click the cog on the right side and select **Add Attribute/Value**.

Step 9. Select **AD1 > External Groups Equals Employees** (or another AD group of your choosing).

Step 10. For the **AuthZ Profiles**, click the plus sign (+).

Step 11. Click the cog in the upper-right corner and select **Add New Standard Profile**.

Step 12. Name the new authorization profile **`Employee Full Access`** (see Figure 11-8).

https://t.me/learningnets

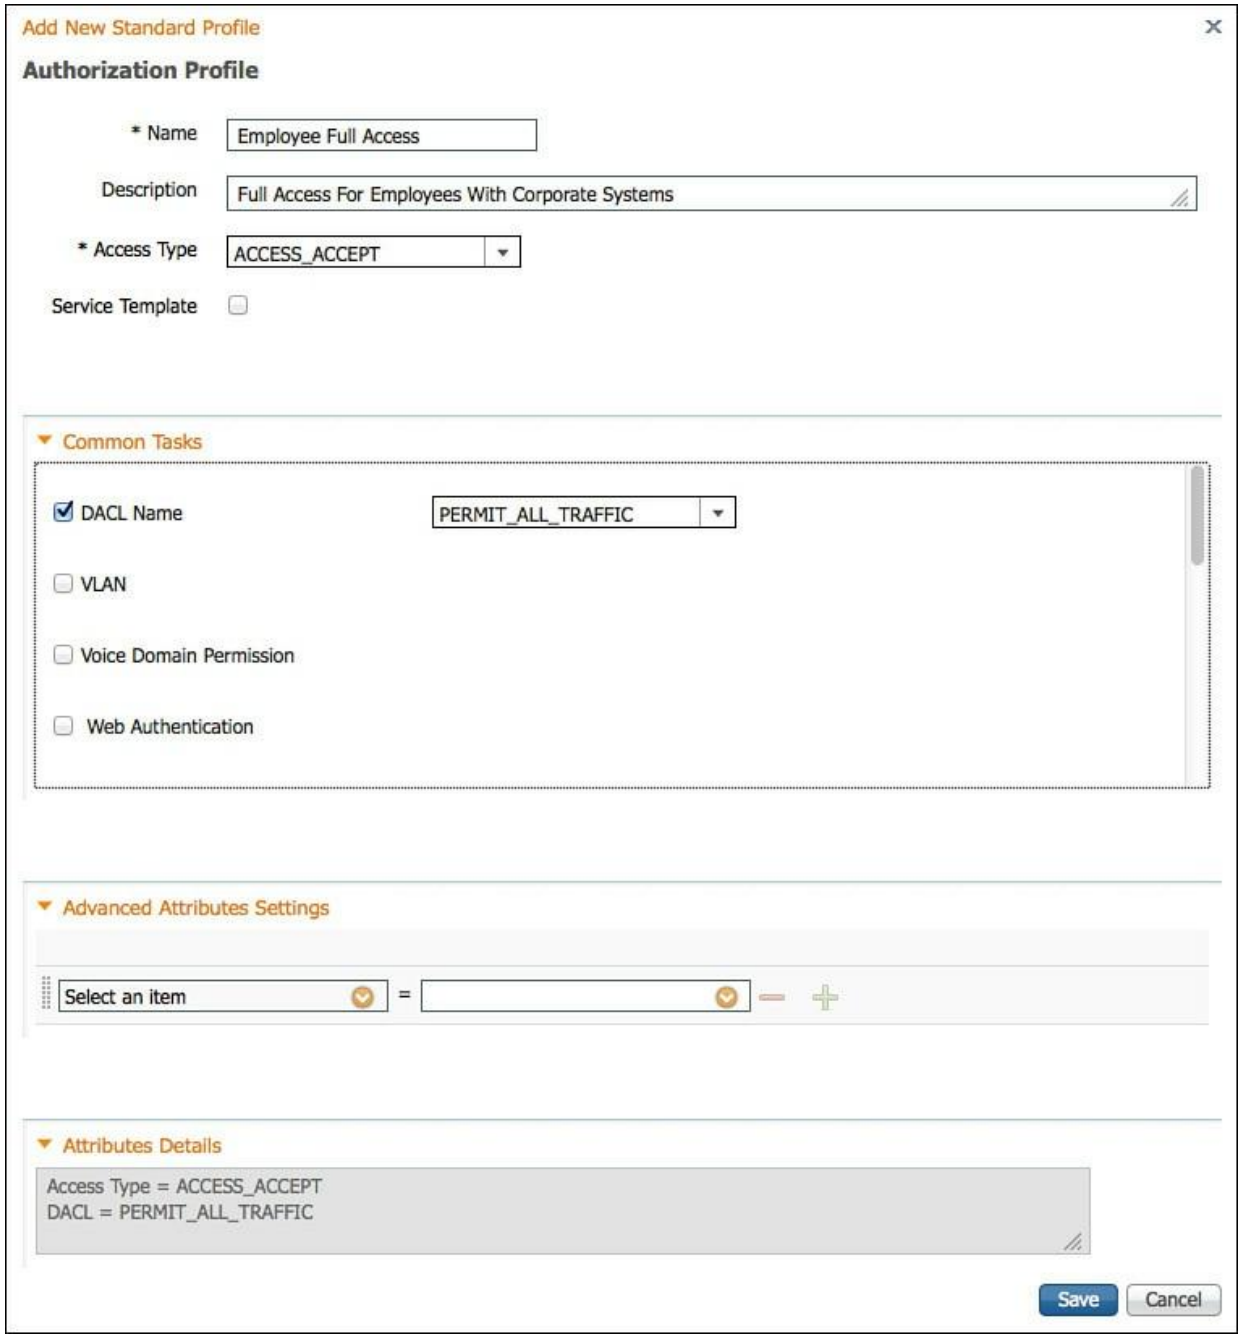


Figure 11-8 Employee full access authorization profile.

Step 13. Optionally add a description.

Step 14. Access Type = Access_Accept.

Step 15. Select **DACL Name > PERMIT_ALL_TRAFFIC**.

Step 16. Click **Save**.

Step 17. Click **Done** to finish editing the rule.

Step 18. Click **Save** to save the authorization policy.

Figure 11-9 shows the completed authorization rule.

https://t.me/learningnets

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
✓	Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
✓	Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
✓	Employee and CorpMachine	if (Network Access:EapChainingResult EQUALS User and machine both succeeded AND AD1:ExternalGroups EQUALS ise.local/Users/Employees)	then Employee Full Access
✓	Default	if no matches, then PermitAccess	

Figure 11-9 Completed employee and CorpMachine rule.

Internet Only for Smart Devices

Now the rule for employees with corporate devices has been created, we need to create the rule below it that provides Internet access only to employee authentications on mobile devices.

To begin this rule, we will first create a new dACL that will be applied to switches; then we will create the authorization result. Next, we will go back into the authorization policy and build the rule. Follow these steps:

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. Within dACL content, provide an ACL that would permit required traffic and deny traffic destined to the corporate network.

Step 6. Click **Submit**.

The example in Figure 11-10 is just an example; each organization might need its own customization.

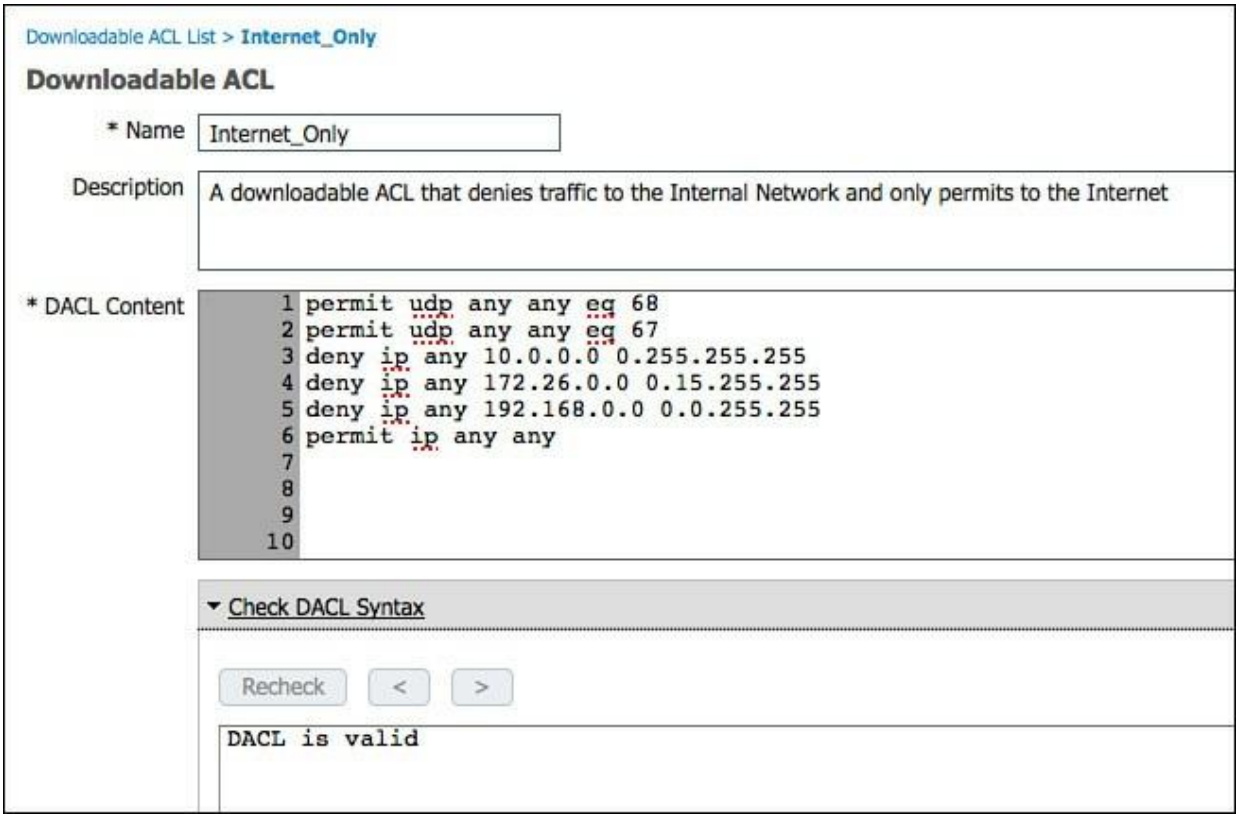


Figure 11-10 Internet-only dACL.

Now the dACL is created, it's time to create the authorization profile.

https://t.me/learningnets

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. The access type is ACCESS_ACCEPT.

Step 6. Select **DACL Name** and select **Internet Only**.

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 will be assigned VLAN on wireless but not wired. In either case, the VLAN must already exist on the NAD.

Step 8. Select **Airespace ACL Name** and fill in the name of the ACL on the controller that will provide Internet-only access. Because we are referencing this ACL by name only, this ACL must already exist on the WLC—we refer to this as a named ACL.

Step 9. Click **Submit**.

Figure 11-11 shows the completed authorization profile.

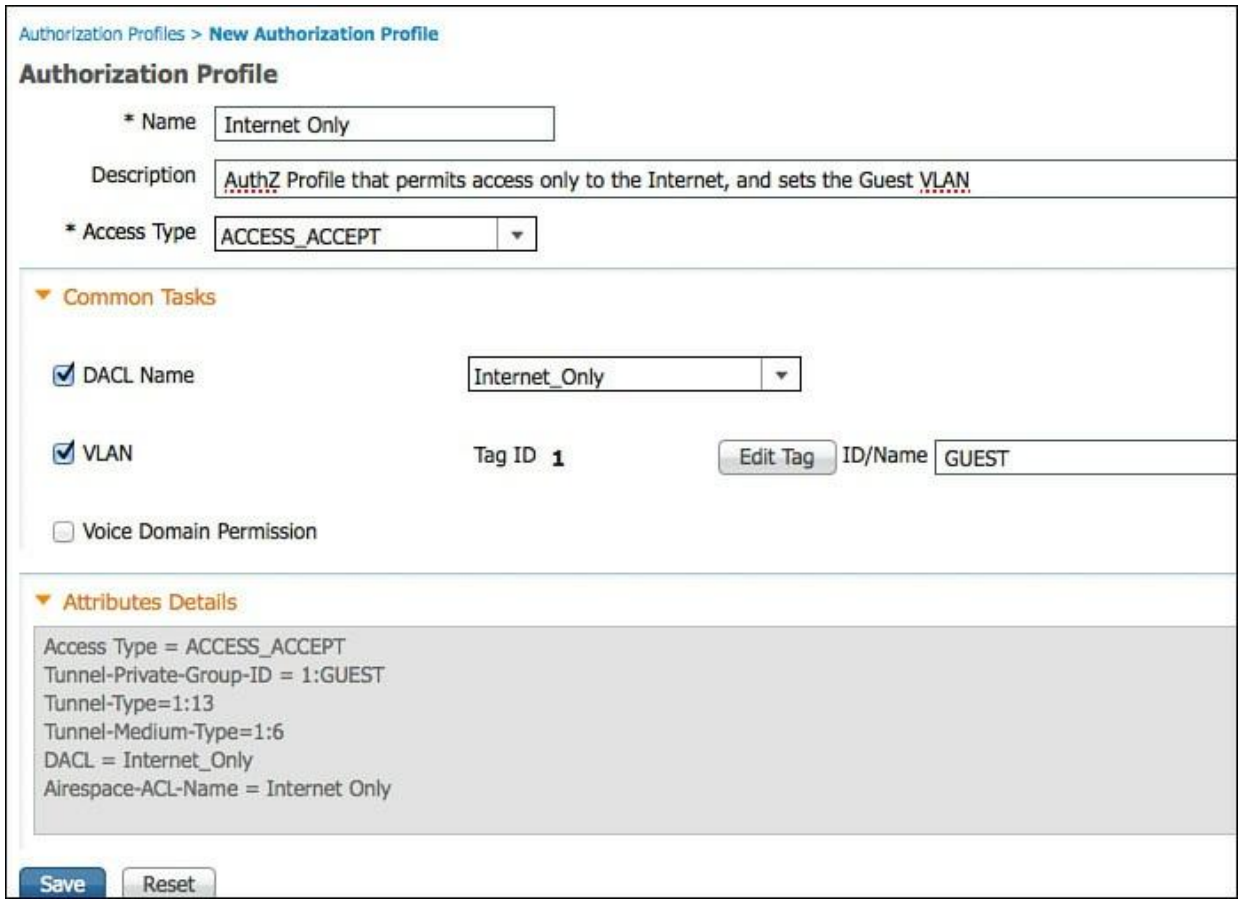


Figure 11-11 Internet-only authorization profile.

Before we build the authorization policy, we 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. Here's how:

Step 1. Navigate to **Policy > Profiling > Logical Profiles**.

Step 2. Click **Add**.

Step 3. Name the logical policy **`SmartDevice`** (see Figure 11-12).

https://t.me/learningnets



Figure 11-12 SmartDevice logical profile.

Step 4. Optionally provide a description.

Step 5. Select all the mobile platforms from the **Available Devices** side, and click the **>** to move them to the **Assigned Policies** side.

Step 6. Click **Submit**.

Lastly, it is now time to create the authorization rule:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Insert a new rule above the default rule.

Step 3. Name the rule **`Employee SmartDevices`**.

Step 4. Select the + sign for Condition(s), and select **Endpoints > LogicalProfile**.

Step 5. Select **Equals**.

Step 6. Select **SmartDevice**.

Step 7. Click the cog on the right side and select **Add Attribute/Value**.

Step 8. Select **AD1 > External Groups Equals "Employees"** (or another AD Group of your choosing).

Step 9. For the AuthZ Profiles, click the plus sign (+).

Step 10. Select **Standard > Internet Only**.

Step 11. Click **Done**.

Step 12. Click **Save**.

The completed authorization rule is displayed in Figure 11-13.

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
✓	Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
✓	Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
✓	Employee and CorpMachine	if (Network Access:EapChainingResult EQUALS User and machine both succeeded AND AD1:ExternalGroups EQUALS ise.local/Users/Employees)	then Employee Full Access
✓	Employee SmartDevice	if (EndPoints:LogicalProfile EQUALS SmartDevice AND AD1:ExternalGroups EQUALS ise.local/Users/Employees)	then Internet_Only
✓	Default	if no matches, then PermitAccess	

Figure 11-13 Employee SmartDevice authorization rule.

https://t.me/learningnets

Employee Limited Access Rule

Now that the rule for employees connecting with mobile devices has been created, we need to create the rule below it that provides limited access only to employee authentications on any other device.

To begin this rule, we will first create a new dACL that will be applied to switches. Then we will create the authorization result, and then go back into the authorization policy and build the rule. Here's how:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Downloadable ACLs**.

Step 2. Click **Add**.

Step 3. Name the ACL **`Employee_Limited`** (see Figure 11-14).

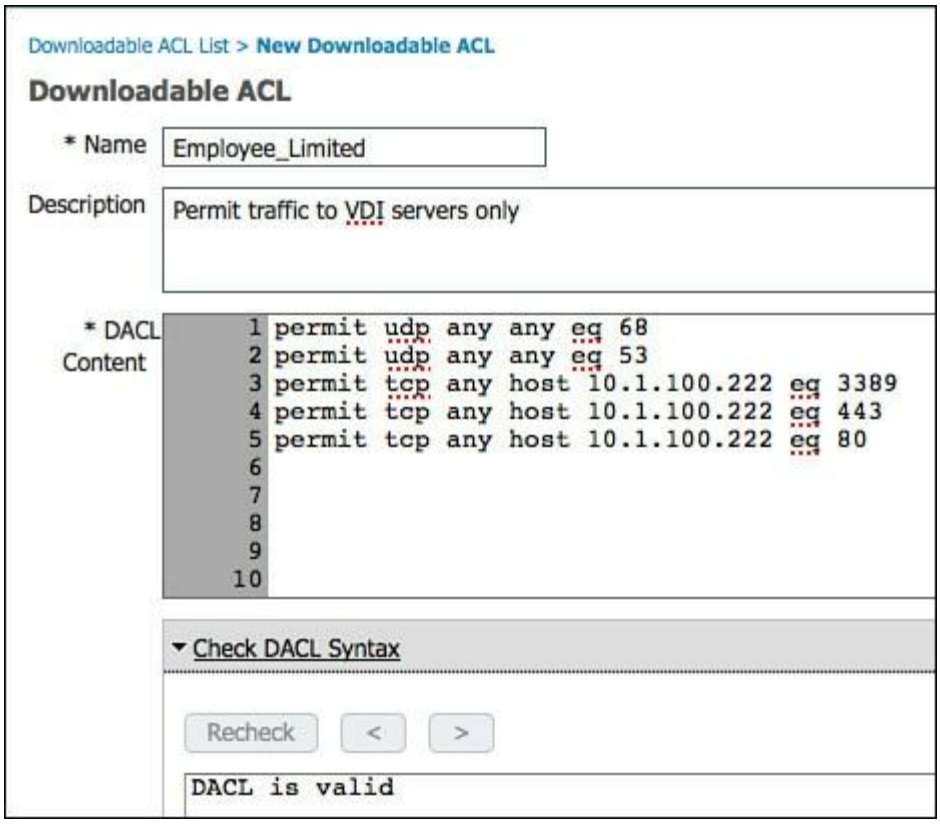


Figure 11-14 Employee limited dACL.

Step 4. Optionally provide a description.

Step 5. Within dACL content, provide an ACL that would permit required traffic and deny traffic destined to the corporate network. For our example, we will allow traffic to reach our virtual desktop infrastructure and essential services such as DNS only.

Step 6. Click **Submit**.

Now that the dACL is created, we will build the authorization policy to permit network access and apply that dACL. Follow these steps:

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. The access type is ACCESS_ACCEPT.

Step 6. Select **DACL Name** and select **Employee_Limited**.

Step 7. We will not assign a different VLAN for this authorization.

https://t.me/learningnets

Step 8. Select **Airespace ACL Name** and fill in the name of the ACL on the controller that will provide Internet-only access.

Step 9. Click **Submit**.

Figure 11-15 shows the completed authorization profile.



Figure 11-15 Employee limited authorization profile.

Now, create the authorization policy rule to assign that authorization profile by doing the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Insert a new rule above the default rule.

Step 3. Name the rule **`Employee Limited`**.

Step 4. Select the plus sign (+) for conditions.

Step 5. Select **AD1 > External Groups Equals "Employees"** (or another AD Group of your choosing).

Step 6. For the AuthZ Profiles, click the plus sign.

Step 7. Select **Standard > Employee Limited**.

Step 8. Click **Done**.

Step 9. Click **Save**.

https://t.me/learningnets



Figure 11-16 Employee limited authorization rule.

Saving Conditions for Reuse



ISE offers the ability to save conditions to the library to make it much easier to reuse them again in other policies. To show this, we will go back into our sample authorization policy and save a few of the conditions. Not only does it make it quicker to find the condition in the future, but it also helps clean up the policies and make them easier to read.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Edit the Employee and CorpMachine rule.

Step 3. Expand the Conditions.

Step 4. Click **Add All Conditions Below to Library**.

This is adding the full set of conditions, including the AND operator (see Figure 11-17).

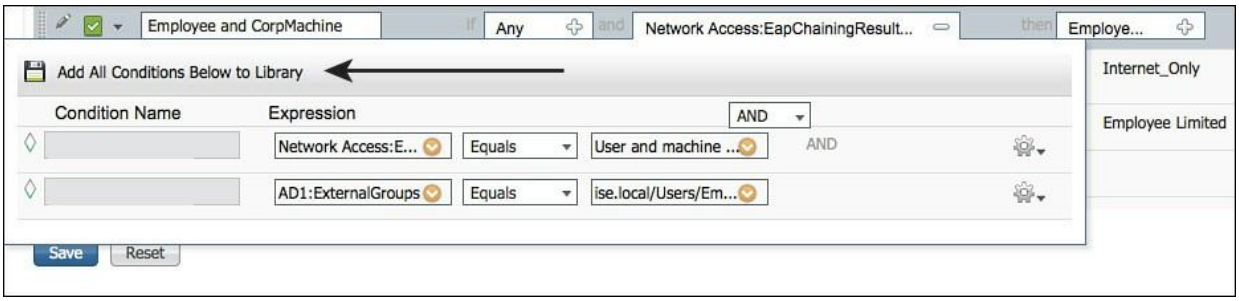


Figure 11-17 Add all conditions below to library.

Step 5. Provide a name for this new saved condition. We have called it EmployeeFullEAPChain.

Step 6. Finish editing the rule.

Step 7. Click **Save**.

As shown in Figure 11-18, the authorization policy text is simplified now with the name of the saved conditions instead of the raw attributes.



Figure 11-18 Authorization policy after saving conditions to library.

Conditions can be saved individually, unlike what is displayed in Figure 11-18. Individual conditions are often

https://t.me/learningnets

more desirable to save to the library because they can be mixed and matched with other library conditions or raw attributes, allowing for more flexibility. Next, we will save the Employees group for Active Directory as a condition because it is a common condition used in many rules, as shown in Figure 11-19. Follow these steps:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Edit the **Employee SmartDevice** rule.

Step 3. Expand the **Conditions**.

Step 4. Click the cog on the right side of the Employees line.

Step 5. Select **Add Condition to Library**.

Step 6. Name the condition **`Employees`**.

Step 7. Click the green check mark.

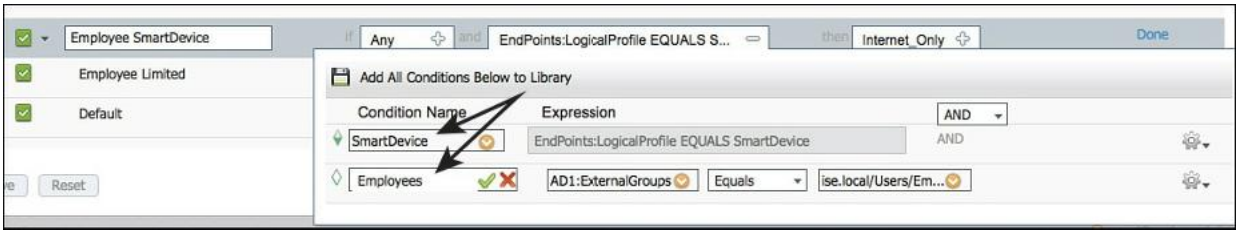


Figure 11-19 Saving employees to library.

Step 8. Click **Done** to finish editing the rule.

Step 9. Click **Save**.

Figure 11-20 shows the final authorization policy.

Looking at Figure 11-20, it is easy to see why saving the raw attributes as simple or compound conditions can clean up the readability of the policies. These reusable conditions exist for many policies within ISE, not only the authorization policy.

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
✓	Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
✓	Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
✓	Employee and CorpMachine	if EmployeeFullEAPChain	then Employee Full Access
✓	Employee SmartDevice	if (SmartDevice AND Employees)	then Internet_Only
✓	Employee Limited	if Employees	then Employee Limited
✓	Default	if no matches, then PermitAccess	

Figure 11-20 Final authorization policy.

Combining AND with OR Operators

Through the normal ISE GUI, it is not easy to see how to combine the AND operator with the OR operator. However, on some occasions a compound condition needs to blend the two operators. Cisco ISE does have the capability to create a complex compound condition made up of simple conditions and any mixture of operators. The catch is that raw attributes cannot be used—only simple conditions can.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Policy Elements > Conditions > Authorization > Simple Conditions**.

At this point, you should have at least two simple conditions listed: Employees and SmartDevice. Create some new simple conditions to be used in the complex compound condition.

https://t.me/learningnets

Step 2. Click **Add**.

Step 3. Name the new simple condition **`PostureCompliant`**.

Step 4. Select **Session > PostureStatus** under the *Attribute column.

Step 5. Select **Equals** under the *Operator column.

Step 6. Select **Compliant** under the *Value column.

Step 7. Click **Submit**.

Figure 11-21 shows the final values for the PostureCompliant simple condition.

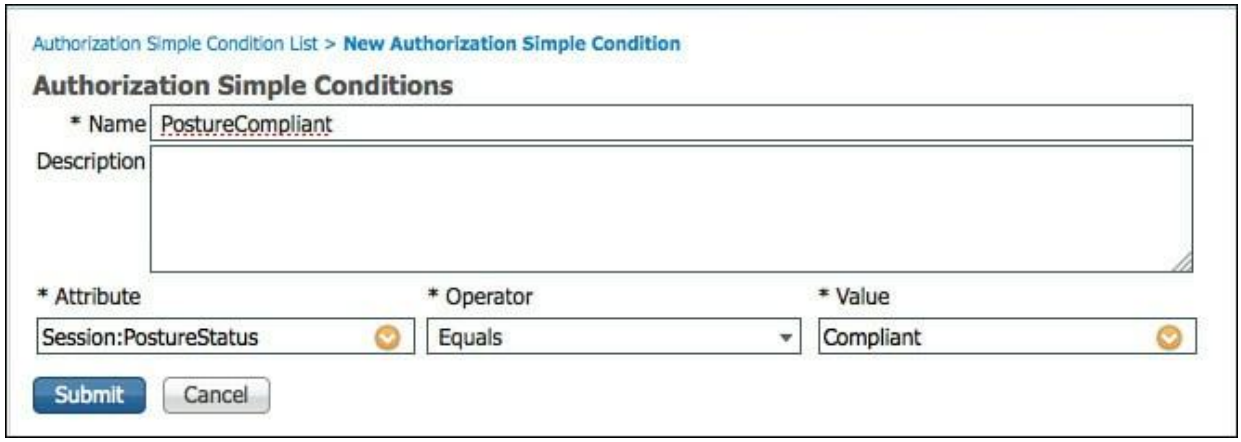


Figure 11-21 PostureCompliant simple condition.

Step 8. Click **Add**.

Step 9. Name the new simple condition **`PEAP`**.

Step 10. Select **Network Access > EapTunnel** under the *Attribute column.

Step 11. Select **Equals** under the *Operator column.

Step 12. Select **PEAP** under the *Value column.

Step 13. Click **Submit**.

Figure 11-22 shows the final values for the PEAP simple condition.



Figure 11-22 PEAP Simple Condition

Step 14. Click **Add**.

Step 15. Name the new simple condition **`EAP-TLS`**.

Step 16. Select **Network Access > EapAuthentication** under the *Attribute column.

Step 17. Select **Equals** under the *Operator column.

https://t.me/learningnets

Step 18. Select **EAP-TLS** under the *Value column.

Step 19. Click **Submit**.

Figure 11-23 shows the final values for the EAP-TLS simple condition.

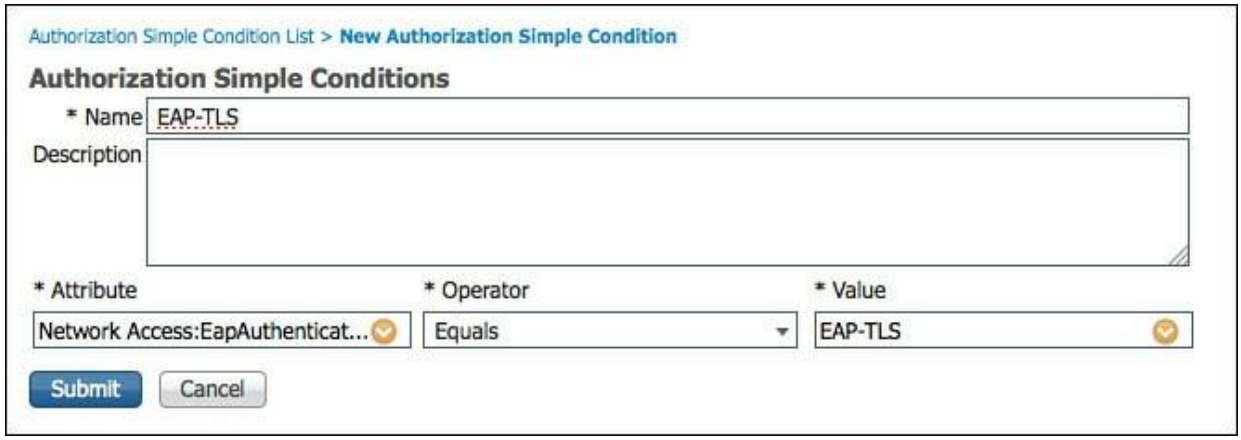


Figure 11-23 EAP-TLS simple condition.

Step 20. Click **Add**.

Step 21. Name the new simple condition **`EAPChainingSuccess`**.

Step 22. Select **Network Access > EapChainingResult** under the *Attribute column.

Step 23. Select **Equals** under the *Operator column.

Step 24. Select **User and Machine both succeeded** under the *Value column.

Step 25. Click **Submit**.

Figure 11-24 shows the final values for the EAPChainingSuccess simple condition.

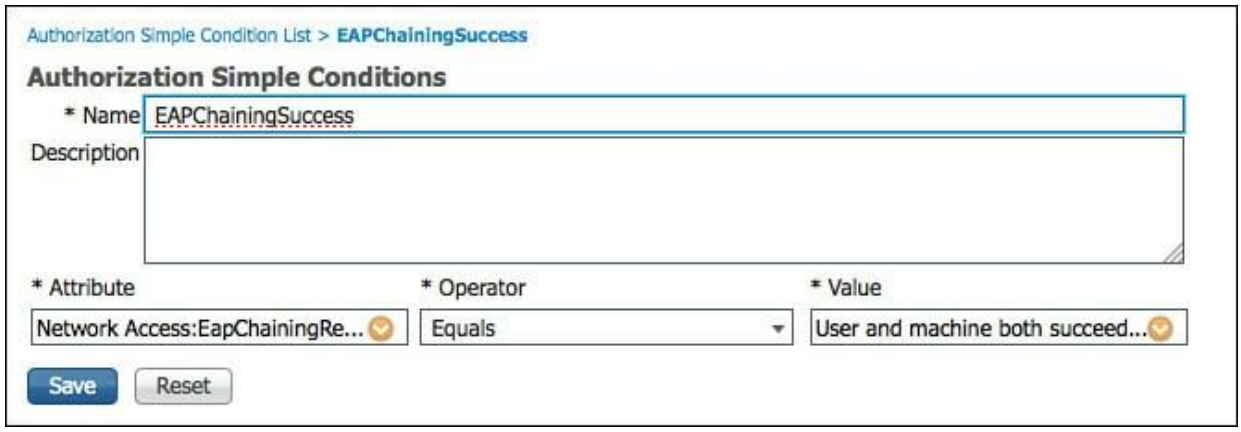


Figure 11-24 EAP chaining success simple condition.

Now, let's create a complex compound condition that uses the simple conditions created so far:

Step 1. Navigate to **Policy > Policy Elements > Conditions > Authorization > Compound Conditions**.

Step 2. Click **Add**.

Step 3. Name the new compound condition **`TestComplexCondition`**.

Step 4. Click the button labeled **Select Existing Condition from Library**, as shown in Figure 11-25.

https://t.me/learningnets

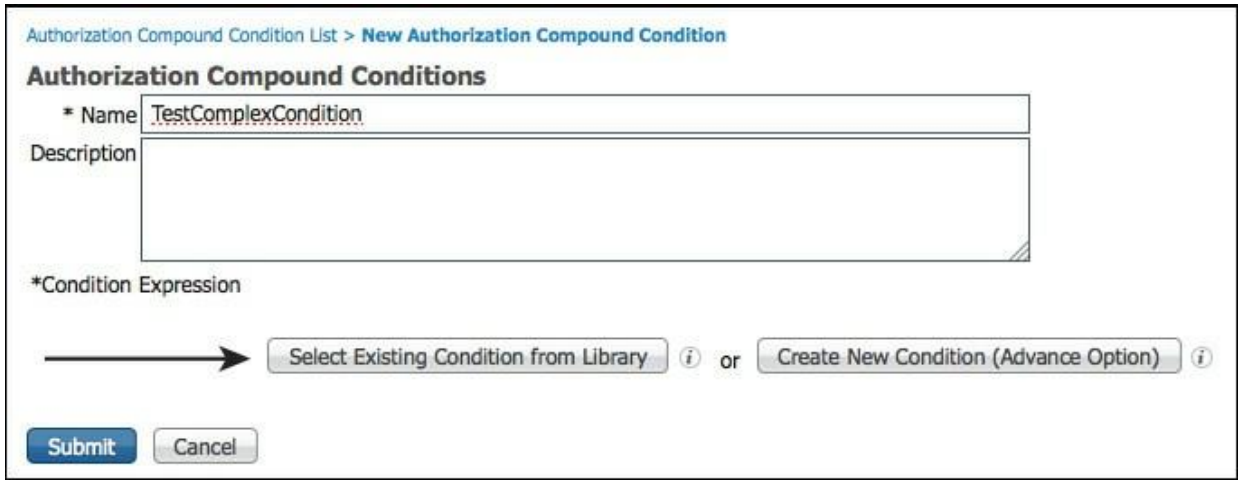


Figure 11-25 Select existing condition from library.

When you begin to build a compound condition using an existing condition from the library, two buttons appear in the upper-right corner. These are the Simple View and Advance View buttons.

Step 5. Click the **Advance View** button that looks a bit like a tic-tac-toe board, as shown in Figure 11-26.

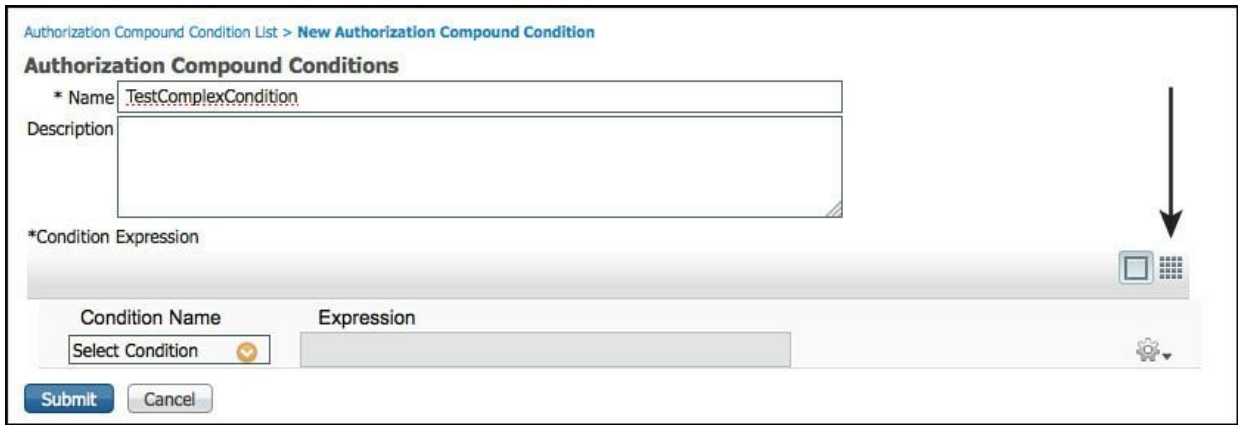


Figure 11-26 Advance View button.

After you click the Advance View button, the screen will reconfigure itself as shown in Figure 11-27. In this view you will be able to build entire expressions, similar to a mathematical formula.

https://t.me/learningnets

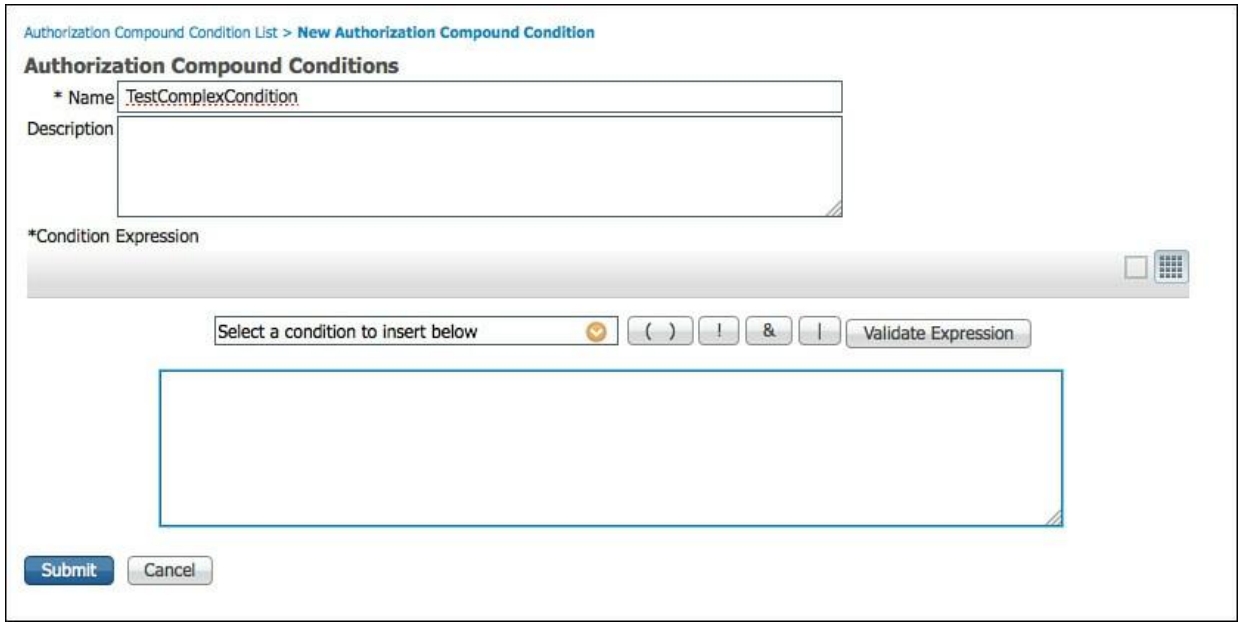


Figure 11-27 Advanced view.

Step 6. Click the parenthesis button **( )** to begin your first match in the expression.

Step 7. Select the **SmartDevice** condition from the drop-down list, so it appears within the parentheses like this: **(SmartDevice)**.

Step 8. Click the ampersand (&) button, which represents the AND operator, so it appears after SmartDevice like this: **(SmartDevice and)**.

Step 9. Select the EAP-TLS condition from the drop-down list, so your expression now looks like this: **(SmartDevice and EAP-TLS)**.

Step 10. Click the ampersand button, so it appears after EAP-TLS like this: **(SmartDevice and EAP-TLS and)**.

Step 11. Select the Employees condition from the drop-down list. The expression should now look like this: **(SmartDevice and EAP-TLS and Employees)**.

Step 12. Now move your cursor outside of the expression and click the pipe button (|), which represents the OR operator, followed by the parentheses button. The full expression now looks like this: **(SmartDevice and EAP-TLS and Employees) | ( )**.

Step 13. Inside the second set of parentheses, build the following expression: **(! SmartDevice and (EAPChainingSuccess | EAP-TLS) and PostureCompliant)**.

Figure 11-28 shows the final expression to be saved. The exclamation point (!) is the symbol for NOT.

https://t.me/learningnets

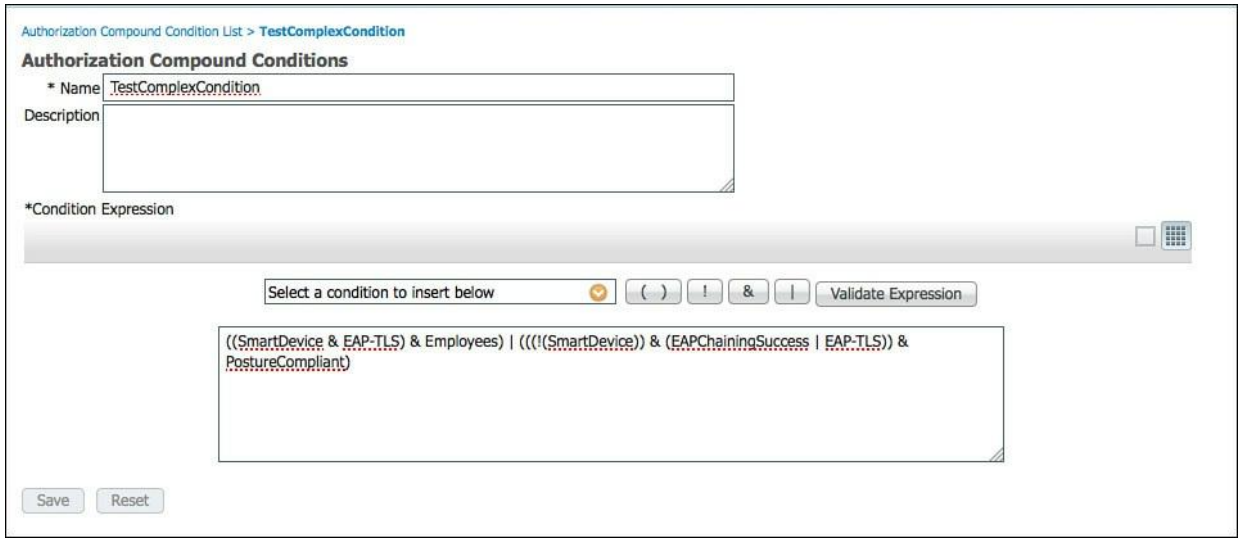


Figure 11-28 Final complex compound expression.

So, the final expression would read as If (endpoint is a smart device AND was authenticated with EAP-TLS AND the user is a member of the Employees AD group) OR (endpoint is not a smart device AND either EAPChaining or EAP-TLS authentication succeeded AND the endpoint's posture is compliant).

Step 14. Click **Submit** to save the complex compound condition.

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 11-2 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Section	Goals of authorization policies	265
Section	Role-specific authorization rules	271
Step list	Saving conditions for re-use from the ISE GUI	279

Table 11-2 Key Topics for Chapter 11

Define Key Terms

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

context

simple condition

compound condition

https://t.me/learningnets

Part IV: Implementing Secure Network Access

https://t.me/learningnets

Chapter 12. Implement Wired and Wireless Authentication

This chapter covers the following exam topics:

- **Implement Wired/Wireless 802.1X**
- **AV Pairs**
- **Implement MAB**
- **Implement Network Authorization Enforcement**
- **Troubleshooting, Monitoring, and Reporting Tools**

In Chapter 9, "Initial Configuration of Cisco ISE," during the initial configuration of Cisco ISE, you added a network access device (NAD) definition into ISE and were introduced to the concept of importing multiple NADs that are defined in a comma-separated value (CSV) file. In Chapter 10, "Authentication Policies," and Chapter 11, "Authorization Policies," some basic authentication and authorization policies were defined on ISE.

What is still missing is one of the most important aspects of secure network access—the configuration of the NADs themselves (switch, wireless controller, and so on), also commonly referred to as the Access Layer device.

The NAD is the point of access for an endpoint and is the first network device through which an endpoint will communicate. This makes it a critical part of any secure network architecture because the NAD is also the main policy enforcement point.

When a company begins the project planning for a secure network access project (most often called a network access control [NAC] project), the basic use cases are always considered, such as how will our corporate assets authenticate to the network and what will be our plan for guest users. A lab might be built, and technology tested with a proof of concept environment (usually at small scale), and things often seem perfect. The proof of concept is considered successful and everyone celebrates the win.

In truth, many important things are not considered until the project team truly starts drilling into the deployment at scale. For instance: How to handle WAN outages when the policy server is centrally located is often considered, but the devil is in the details? What about the desktop reimaging process? In other words, when a corporate workstation needs a software repair, do they use a preexecutable environment (PXE) and boot to the network where a clean image is put onto the workstation? What about using thin-client devices that don't have a local operating system and that boot to the network and download an operating system (OS) or even connect to a centrally located terminal server?

How will these devices get access to the network if they cannot authenticate, and how does the team ensure these devices can reach the appropriate location to download their OSes or the reimaging server, while restricting access to anything unnecessary? The policy server is obviously a lynch pin to a successful deployment. However, intelligence in the Access Layer device is crucial and will make or break the entire project.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

https://t.me/learningnets

Foundation Topics Section	Questions
Implement Wired/Wireless 802.1X	1-2, 4, 7-8
AV Pairs	1
Implement MAB	5-6
Implement Network Authorization Enforcement	1, 4, 8
Troubleshooting, Monitoring, and Reporting Tools	3, 9-10

Table 12-1 "Do I Know This Already?" Section-to-Question Mapping

> **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 which level of the configuration do the 802.1X-related commands exist?
   - **a.** Global configuration only.
   - **b.** Interface configuration only.
   - **c.** Both at global configuration level as well as per interface.
   - **d.** Enabling 802.1X changes the context to a dot1x subconfiguration mode, where all related commands are entered.
2. When configuring a Cisco Wireless LAN Controller (WLC) for communication with ISE, what must be configured for the wireless LAN (WLAN)? (Choose two.)
   - **a.** The authentication and authorization RADIUS servers can be pointed to different ISE PSNs, as long as those PSNs are part of a node group.
   - **b.** The authentication and authorization RADIUS servers can 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. True or False? Cisco switches should be configured in production to send syslog messages to the ISE MNT node.
   - **a.** True
   - **b.** False
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-Requests to the AAA servers to verify whether they are still alive. The username and password will be 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 local user database.
   - **d.** Without the local username and password in the configuration, an administrator can be locked

https://t.me/learningnets

out of the switch when the RADIUS server is unavailable.

5. True or False? 802.1X can be configured on all switch interfaces, including Layer-3 interfaces.
   a. True
   b. False
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. Flex-Auth
   d. Flex-Connect
7. Which host mode will permit 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 Mode
   b. MDA
   c. Multi-Auth
   d. Multi-Host
8. Which interface-level command is the equivalent of "turn authentication on"?
   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 will display 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 will validate that authentications are being attempted, which authentications are successful, and which 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 will be required. Some aspects must be configured from global configuration mode, and others are configured at the individual interfaces. There is an authentication subsystem that must be enabled across the entire switch (global configuration), and there are interface configurations that must be entered for authentication settings (timers, enabling 802.1X communication, and more).

https://t.me/learningnets

Global Configuration AAA Commands



By default, the AAA subsystem of the Cisco switch is disabled. Prior to enabling the AAA subsystem, none of the required commands will be available as a configuration option. Here's how you enable it:

Step 1. Enable Authentication, Authorization, and Accounting on the access switch(es):

```
C3560X(config)#aaa new-model
```

An authentication method is required to instruct the switch on which group of RADIUS servers to use for 802.1X authentication requests.

Step 2. Create an authentication method for 802.1X:

Click here to view code image

```
C3560X(config)#aaa authentication dot1x default group radius
```

The method created in Step 2 will enable 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 is what defines that the user or device is actually allowed to access the network and which level of access is actually permitted.

Step 3. Create an authorization method for 802.1X:

Click here to view code image

```
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 (Cisco ISE) knows the exact state of the switchport and endpoint. Without the accounting packets, Cisco 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. That field is populated with the IP address of the endpoint and updates the RADIUS server (ISE) so it has the capability 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 Device Sensor, the sensor data will be sent to ISE using the RADIUS accounting configuration.

Step 4. Create an accounting method for 802.1X:

Click here to view code image

```
C3560X(config)# aaa accounting dot1x default start-stop group radius
```

Global Configuration RADIUS Commands

A change in the way RADIUS servers are configured in IOS occurred between version 12.2.x and the newer 15.x IOS code. When the configuration differs in the following steps, it will be broken out into a separate section.

https://t.me/learningnets

IOS 12.2.X

Cisco switches provide a proactive method to check the availability of the RADIUS server. With this practice, the switch will send periodic test authentication messages to the RADIUS server (Cisco ISE). It is looking for a RADIUS response from the server. A success message is not necessary—a failed authentication will suffice because it shows that the server is alive.

If you want to receive an Access-Accept response, the username you are creating here must already exist in either ISE or another identity store, or it needs to be added to the local user database in Cisco ISE at a later step. This account will be used in a later step where you define the RADIUS server:

Step 1. Within global configuration mode, add a username and password for the RADIUS keepalive:

Click here to view code image

```
C3560X(config)#username radius-test password password
```

In Step 2 you are adding each Cisco ISE policy service node (PSN) to the switch configuration, using the test account we created in Step 1. The server will proactively be checked for responses once per hour, in addition to any authentications or authorizations occurring through normal processes. Repeat Step 2 for each PSN.



Step 2. Add the Cisco ISE servers to the RADIUS group:

Click here to view code image

```
C3560X(config)#radius-server host ise_ip_address auth-port 1812
acct-port 1813 test username radius-test key shared_secret
```

The switch has been configured to proactively check the Cisco ISE server for RADIUS responses. Now 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 they begin in the "Both IOS 12.2.X and 15.X" section.

IOS 15.X

Cisco switches provide a proactive method to check the availability of the RADIUS server. With this practice, the switch will send periodic test authentication messages to the RADIUS server (Cisco ISE). It is looking for a RADIUS response from the server. A success message is not necessary—a failed authentication will suffice because it shows that the server is alive.

If you want to receive an Access-Accept response, the username you are creating here must already exist in either ISE or another identity store, or it needs to be added to the local user database in Cisco ISE at a later step. This account will be used in a later step where you define the RADIUS server:

Step 1. Within global configuration mode, add a username and password for the RADIUS keepalive:

Click here to view code image

```
C3560X(config)#username radius-test password password
```

Next, you are adding each Cisco ISE PSN to the switch configuration.



Step 2. Add the Cisco ISE server:

Click here to view code image

https://t.me/learningnets

```
C3560X(config)#radius server [name]
```

This is where the configuration really differs between IOS 12.2.X and IOS 15.X. In Step 2, you created a new RADIUS server object. The configuration related to that server is 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:

Click here to view code image

```
C3560X(config-radius-server)# address ipv4 [ip-address] auth-port 1812
acct-port 1813
```

Step 4. Configure the shared secret:

Click here to view code image

```
C3560X(config-radius-server)#key [shared-secret]
```

This uses the test account we created in Step 1. The server will proactively be checked for responses once per hour, in addition to any authentications or authorizations occurring through normal processes. Repeat Step 2 for each PSN.

Step 5. Configure the IP address, authentication port, and accounting port:

Click here to view code image

```
C3560X(config-radius-server)# automate-tester username
[username-from-step-1]
```

Optionally: If you are using a Cisco 3850, 3650, or 5760 series switch, the device sends its **calling-station-id** in the format of xxxx.xxxx.xxxx by default, instead of xx-xx-xx-xx-xx-xx. This can be changed 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**:

Click here to view code image

```
3850(config)# radius-server attribute 31 mac format ietf lower-case
```

Both IOS 12.2.X and 15.X

The switch has been configured to proactively check the Cisco ISE server for RADIUS responses. Now configure the counters on the switch to determine whether the server is alive or dead. Your settings will configure the switch to wait 5 seconds for a response from the RADIUS server and attempt the test 3 times before marking the server dead. If that RADIUS server doesn't send a valid response within 15 seconds, it will be marked as dead. High availability is covered in more detail in Chapter 21, "ISE Scale and High Availability." Do the following:

Step 1. Set the dead criteria:

Click here to view code image

```
C3560X(config)#radius-server dead-criteria time 5 tries 3
```

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 change of authorization (CoA) (RFC 3576 / RFC 5176) operations in a different listing, also within global configuration mode.



https://t.me/learningnets

Step 2. Enable change of authorization:

Click here to view code image

```
C3560X(config)#aaa server radius dynamic-author
C3560X(config-locsvr-da-radius)#client ise_ip_address server-key
shared_secret
```

Here we configure the switch to send any defined vendor-specific attributes (VSAs) to Cisco ISE PSNs during authentication requests and accounting updates.

Step 3. Configure the switch to use the Cisco vendor-specific attributes:

Click here to view code image

```
C3560X(config)#radius-server vsa send authentication
C3560X(config)#radius-server vsa send accounting
```

The three VSAs you will enable in Step 4 are as follows:

- Service-Type (RADIUS attribute 6). This is to ensure that login requests include the Service-Type with them. This Service-Type attribute indicates to ISE which authentication method was chosen for the switchport and, ultimately, the endpoint. The options are
  - Service-Type:Framed = 802.1X
  - Service-Type:Call-Check = MAB
  - Service-Type:Login = Local WebAuth
- Framed-IP-Address (RADIUS attribute 8). This VSA ensures the IP address is included (if one exists). This Framed-IP-Address can be used by ISE to correlate the user with other ISE profiling techniques, which are discussed in detail in Chapter 15, "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 VPN).

Step 4. Next, enable the VSAs:

Click here to view code image

```
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
```

Switches can often have multiple IP addresses associated to them. Therefore, it is a best practice to always force any management communications to occur via a specific interface. This interface IP address must match the IP address defined in the Cisco ISE network device object.

The applicable traffic that can be sourced is RADIUS and SNMP. SNMP is optional and used mostly for sending SNMP traps or informs for endpoint profiling.

Step 5. Ensure the switch always sends traffic from the correct interface:

Click here to view code image

```
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
```

https://t.me/learningnets

Global 802.1X Commands



Enabling 802.1X globally on the switch does not actually enable authentication on any of the switchports. Authentication will be configured but not enabled until the interface configuration occurs. Do the following:

Step 1. Enable 802.1X globally on the switch:

Click here to view code image

```
C3560X(config)#dot1x system-auth-control
```

Downloadable access control lists (dACLs) are a common enforcement mechanism in Cisco ISE deployments. For dACLs to function properly on a switch, a function called **IP device tracking** must be enabled globally. This command allows for the **any** source in the provided dACL to be replaced with the IP address of the single device connected to the switchport.

Step 2. Enable dACLs to function:

Click here to view code image

```
C3560X(config)#ip device tracking
```

Creating Local Access Control Lists

Certain functions on the switch require the use of locally configured ACLs, such as URL redirection. Some of these ACLs created will be used immediately, and some might not be used until a later phase of your deployment. The goal of this section is to prepare the switches for all possible deployment models at one time and limit the operational expense of repeated switch configuration.

The details of these ACLs and why they are created this way will be discussed in Chapter 13, "Web Authentication," and Chapter 20, "Deploying Safely."

Step 1. Add the following ACL to be used on switchports in Monitor Mode:

Click here to view code image

```
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 switchports in Low-Impact Mode:

Click here to view code image

```
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:

Click here to view code image

```
C3560X(config)#ip access-list ext ACL-WEBAUTH-REDIRECT
C3560X(config-ext-nacl)#remark explicitly deny DNS from being
redirected to address a bug
```

https://t.me/learningnets

```
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:

Click here to view code image

```
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

You have just completed the global configuration settings of the access layer switches, including the RADIUS and AAA methods.

This section focuses on building a single port configuration that can be used across your entire secure unified access deployment, regardless of switch type, deployment stage, or deployment model you choose.

Configuring Interfaces as Switchports

One of the first things to do before configuring any of the authentication settings on the switchport is to ensure the switchport is configured as a Layer-2 port, not a Layer-3 port. This command is a simple, one-word command that you will run, and from that point the other commands you enter will all take effect:

Step 1. Enter interface configuration mode for the switchport range:

Click here to view code image

```
C3560X(config)#interface range first_interface - last_interface
```

Step 2. Ensure the ports are Layer-2 switchports:

Click here to view code image

```
C3560X(config-if-range)#switchport
```

The switch has some predefined macros that run a set of commands in order. The host macro automatically runs three commands for you: configuring the port to be an access port (non-trunk), disabling channel groups, and configuring spanning tree to be in portfast mode.

Step 3. Configure the port for Layer-2 edge using the host macro:

Click here to view code image

```
C3560X(config-if-range)#switchport host
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
```

https://t.me/learningnets

Configuring 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 an undesirable behavior in many customer deployments because it does not allow for guest access, nor does it 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/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 nonauthenticating devices. With the default behavior of 802.1X, an administrator would have to configure ports for printers and other devices that do not have supplicants differently from the ports where she planned to do authentication.



Therefore, Cisco created flexible authentication (Flex-Auth). Flex-Auth allows a network administrator to set an authentication order and priority on the switchport, thereby allowing the port to attempt 802.1X, MAB, and then WebAuth in 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.

As mentioned previously, there are multiple methods of authentication on a switchport: 802.1X (dot1x), (MAB), and web-based authentication (WebAuth). For the purpose of this wired deployment, we will 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. Additionally, the endpoint supplicant should send a periodic EAP over LAN Start (EAPoL-Start) message into the switchport to speed up authentication. If a device is not able to authenticate, it merely has to wait until the dot1x timeout occurs; then MAB occurs. Assuming the device's MAC address is in the correct database, it is then authorized to access the network. Figure 12-1 illustrates this concept.

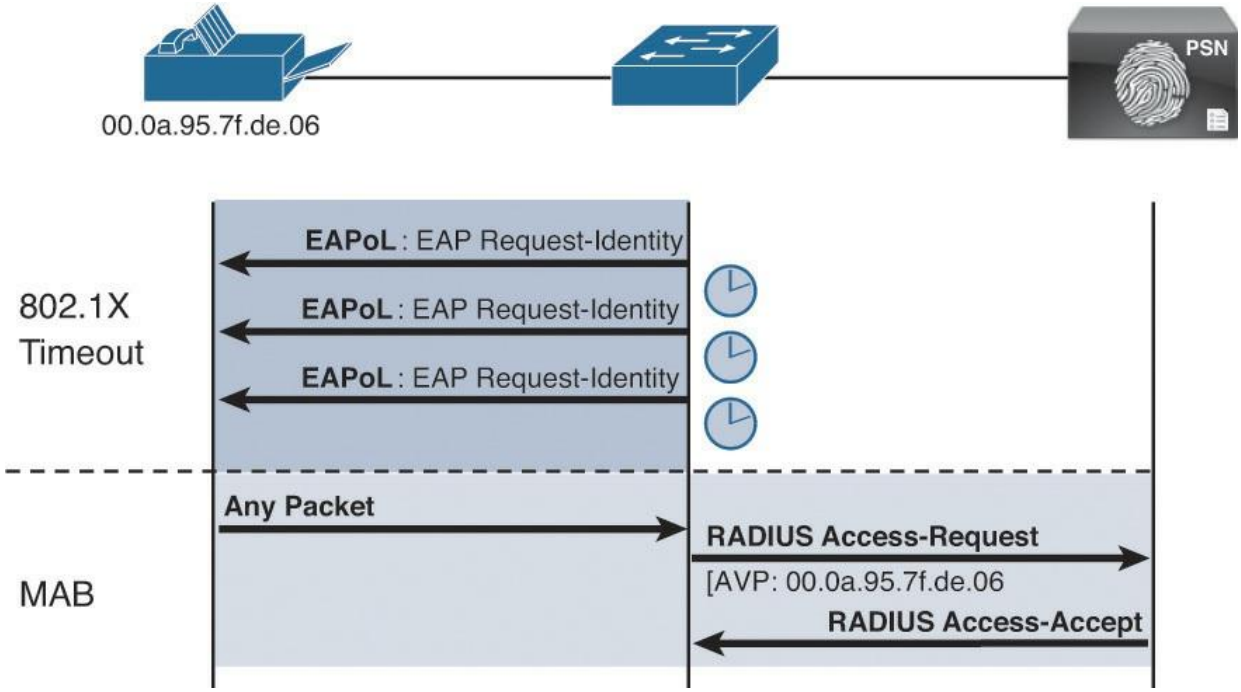


Figure 12-1 Flexible authentication.

The following steps walk you through the configuration of Flex-Auth and the configurable actions for authentication high availability.

With Flex-Auth, you are able to toggle whether dot1x or MAB has the higher priority. The best practice is to

https://t.me/learningnets

always prefer the stronger authentication method (dot1x). The dot1x method is also the default of all Cisco switches. Here's how:

Step 1. Configure the authentication method priority on the switchports:

Click here to view code image

```
C3560X(config-if-range)#authentication priority dot1x mab
```

In certain deployment methods MAB should occur before 802.1X authentication, such as an environment in which the majority of devices are nonauthenticating. For those corner cases, Flex-Auth does allow for a network administrator to set a user-definable authentication order. However, the best practice is to maintain the order of 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 in flight and immediately switch to an 802.1X transaction.

Step 2. Configure the authentication method order on the switchports:

Click here to view code image

```
C3560X(config-if-range)#authentication order dot1x mab
```

Just setting the order and priority alone is not what is actually providing the Flex-Auth technology. You still must enable a key command. This is the command that dictates how the switch should behave when the RADIUS server (Cisco ISE) sends an Access-Reject response. The **authentication event fail action** [action] determines what should occur in the event of an Access-Reject 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 is known as using an Auth-Fail VLAN. In the early days of dot1x, this was how guest access was handled. If the authentication failed, you would assign the Guest" VLAN. It's important to note that in a Cisco ISE secure access deployment, the fail action should be set to **next-method**. That is the third and final piece of the puzzle to make up the Flex-Auth technology.

Note

If the switch makes a local decision, such as assigning a VLAN instead of denying access, the switch will no longer accept CoA commands from the RADIUS server. A way to think about this is that Cisco ISE now thinks that the switchport is down, but the switch has actually permitted the endpoint to have access to the guest VLAN; therefore, the state is now out of sync.

Step 3. Configure the port to use Flex-Auth:

Click here to view code image

```
C3560X(config-if-range)#authentication event fail action next-method
```

In the global configuration section, you configured the RADIUS server entry to use a test account that will proactively alert the switch when Cisco ISE has stopped responding to RADIUS requests. Now we will configure the switchport to locally authorize the port when that server is found to be dead and reinitialize authentication when the server becomes alive again. The topic of high availability is covered in more detail in Chapter 21.

Step 4. Configure the port to use a local VLAN for voice and data when the RADIUS server is dead (when it stops responding):

Click here to view code image

https://t.me/learningnets

```
C3560X(config-if-range)#authentication event server dead action
authorize
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 is **up**) would get full access to network and the others (the new hosts) would not get any 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 a command called **authentication event server dead action reinitialize vlan** vlan-id. With this new feature/CLI, when new hosts try to access the network and the RADIUS server is down, that port is reinitialized immediately and all hosts (in this port) get the same VLAN.

Step 5. Configure the port to use a local VLAN when the RADIUS server is dead, and allow existing and new hosts (optional):

Click here to view code image

```
C3560X(config-if-range)#authentication event server dead action
reinitialize vlan vlan-id
```

Host Mode of the Switchport



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 ending of the authentication session. Since the inception of 802.1X, enhancements have occurred to allow one or more devices per port. The host modes available are listed here:

- **Single Host**—The default mode of 802.1X is Single Host. Only a single MAC address is permitted per interface. In the early days of 802.1X, the protocol expected to rely on link state for authentication control. When the link comes up, the authenticator (switch) starts sending the EAPoL-ID-Requests 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**—Multidomain Authentication (MDA) extends 802.1X to support IP telephony by modifying it to understand two separate domains: a data domain and a voice domain. A switchport 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 more than one device is on the switchport.

  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 to terminate any authentications for that endpoint. This will work for all authentication types (dot1x, MAB, and WebAuth).

  The 802.1X standard supports a different method known as EAPoL-Proxy-Logoff. This enables the IP phone to issue a logoff method for the endpoint supplicant that was authenticated behind the phone. However, that only works for dot1x authentications, and not for MAB or WebAuth.
- **Multiple Authentication**—Multiple Authentication (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 in the MDA section are still applicable to

https://t.me/learningnets

Multi-Auth mode.

- **Multiple Host**—Multiple Host (Multi-Host) mode is not commonly used but is still a valid option. Much like Multi-Auth mode, Multi-Host mode is an extension to MDA. One authentication exists on the voice domain, and one authentication exists 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 type of model. The IP Phone solutions described in the MDA section are still applicable to Mult-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 no denial of service occurs while deploying 802.1X. When the deployment moves into an enforcement phase, it is recommended to use MDA mode. In this sample configuration, we are assuming we are in the initial phases of our deployment.

Step 1. Set the host mode of the port:

Click here to view code image

```
C3560X(config-if-range)#authentication host-mode multi-auth
```

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 might seem to be a nice and secure behavior, it can create an accidental denial of service, especially during the initial phases of deployment. Therefore, you must set the action to be restrict. This mode of operation enables the first authenticated device to continue with its authorization and deny any additional devices.

Step 2. Configure the violation action:

Click here to view code image

```
C3560X(config-if-range)#authentication violation restrict
```

Configuring 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 environments (PXEs) or might have some thin clients that have to boot with DHCP and don't have any way to run a supplicant.

Additionally, when early adopters of 802.1X would deploy authentication companywide, there were repercussions. Many supplicants were misconfigured, and there were unknown devices that could not authenticate because of a lack of supplicant and other reasons.

Cisco created open authentication to aid with deployments. Open authentication allows all traffic to flow through the switchport, even without the port being authorized. This feature enables authentication to be configured across the entire organization but not deny access to any device.

Figure 12-2 depicts the difference between a switchport configured with the default behavior of 802.1X versus a switchport with open authentication configured. This is a key feature that enables the phased approach to deploying authentication, which is discussed in great length in Chapter 20.

https://t.me/learningnets

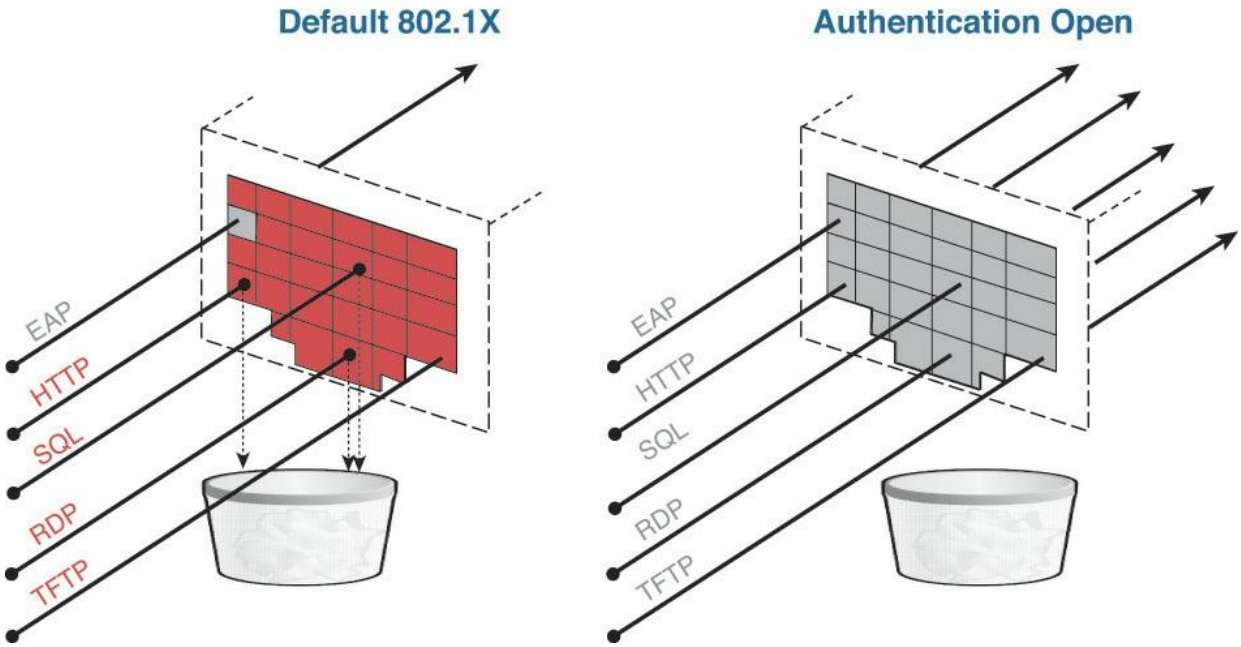


Figure 12-2 Default 802.1X versus authentication open.

Step 1. Set the port for open authentication:

Click here to view code image

```
C3560X(config-if-range)#authentication open
```

Throughout all the discussion in the text, and more specifically this chapter, surrounding MAC authentication bypass, there really is only a single interface command that enables it: **mab**.

Step 2. Enable MAB on the port:

```
C3560X(config-if-range)#mab
```

As previously mentioned, you have enabled dot1x globally on the switch, but no authentication will actually occur yet because the individual switchports have not been configured to do the 802.1X communication. To enable the 802.1X communication to and from a switchport, use the command **dot1x pae authenticator**.

Step 3. Enable the port to perform IEEE 802.1X authentication:

Click here to view code image

```
C3560X(config-if-range)#dot1x pae authenticator
```

Even with all the configuration commands that have been entered into the switch so far, you will still not be processing any authentication requests yet. A key command at the end of the switch configuration section is required first.

Configuring Authentication Timers



Many timers can be modified as needed in a deployment. Unless you are experiencing a specific problem where the adjusting of a timer can correct some unwanted behavior, it is recommended to leave all timers at their default values except for the 802.1X transmit timer (tx-period).

https://t.me/learningnets

The tx-period timer defaults to a value of 30 seconds. Leaving this value at 30 seconds provides a default wait of 90 seconds (3 × tx-period) before a switchport begins the next method of authentication and begins the MAB process for nonauthenticating devices.

Based on numerous deployment experiences, the recommendation is that you set the tx-period value to 10 seconds to provide the most optimal time for MAB devices. Setting the value below 10 seconds can result in unwanted behavior; setting the value greater than 10 seconds can result in DHCP timeouts.

Step 1. Configure the tx-period timer:

Click here to view code image

```
C3560X(config-if-range)#dot1x timeout tx-period 10
```

The recommendation of 10 seconds might need to be reduced or increased, depending on the needs of your specific deployment. However, 10 seconds has been a successful timer for a vast majority of deployments.

Applying the Initial ACL to the Port and Enabling Authentication

This step prepares the port for Monitor Mode by applying a default ACL on the port without denying any traffic.

Step 1. Apply the initial ACL (ACL-ALLOW):

Click here to view code image

```
C3560X(config-if-range)#ip access-group ACL-ALLOW in
```



If you want to enable authentication now, you can. A key interface command is needed to truly enable all these authentication commands: **authentication port-control auto**. Without this command, everything will appear to be working but no authentications will be sent to the RADIUS server.

Step 2. Turn on authentication:

Click here to view code image

```
C3560X(config-if-range)#authentication port-control auto
```

You have three options for the **port-control** command, detailed here:

- **auto**—Enables authentication on this port and allows the authorization result to be sent from the RADIUS server. Short answer: "Turn authentication on!"
- **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. Short answer: "Allow all endpoints connected to this switchport."
- **force-unauthorized**—Hard-codes the port to be in an unauthorized state. This basically disables the process of authentication and denies network access to all endpoints. Short answer: "Denial of service (DoS) the endpoints connected to this switchport."

Authentication Configuration on WLCs

This section reviews the configuration for the Cisco wireless LAN controller. It focuses on version 7.3 and above. With each version of the WLC software, new features are released to aid in BYOD solutions, such as DNS-based ACLs in version 7.6. All WLC screenshots in this chapter are from WLC version 7.6.

As with the Cisco catalyst switches, this chapter assumes you have established basic connectivity with the NAD and are now to the point of bootstrapping the WLC for use with ISE. Similarly, some configuration is

https://t.me/learningnets

globally applicable, meaning it applies to the entire system. Other configuration is per-wireless LAN (per-SSID), which is comparable to interface configurations on the wired NADs.

Configuring the AAA Servers

Much like the wired NADs, the first procedure in configuring the WLC is to add the ISE policy service nodes (PSNs) to the WLC as RADIUS authentication and accounting servers. This is a configuration that affects the entire WLC as a system, similar to the global configuration with a switch.

Adding the RADIUS Authentication Servers

The WLC separates the AAA server configuration into two sections. The authentication servers are configured in a separate section from the accounting servers. Start by adding the ISE PSN to the authentication servers. From the WLC GUI, do the following:

Step 1. Navigate to **Security** > **RADIUS** > **Authentication**.

You need to ensure that the endpoint's MAC address is populated in the RADIUS **calling-station-id** field in all RADIUS packets and that a hyphen is used for the delimiter. This is important because Cisco ISE uses the endpoint MAC address as the unique identifier for the endpoint's identity. This is used for with both 802.1X authentications and wireless MAB.

Step 2. Ensure that both the Acct Call Station ID Type and Auth Call Station ID Type drop-down boxes are set to **System MAC Address**.

Step 3. Ensure that the MAC Delimiter is set to **Hyphen**.

Figure 12-3 shows the final configuration for the Call Station ID and MAC Delimiter settings.

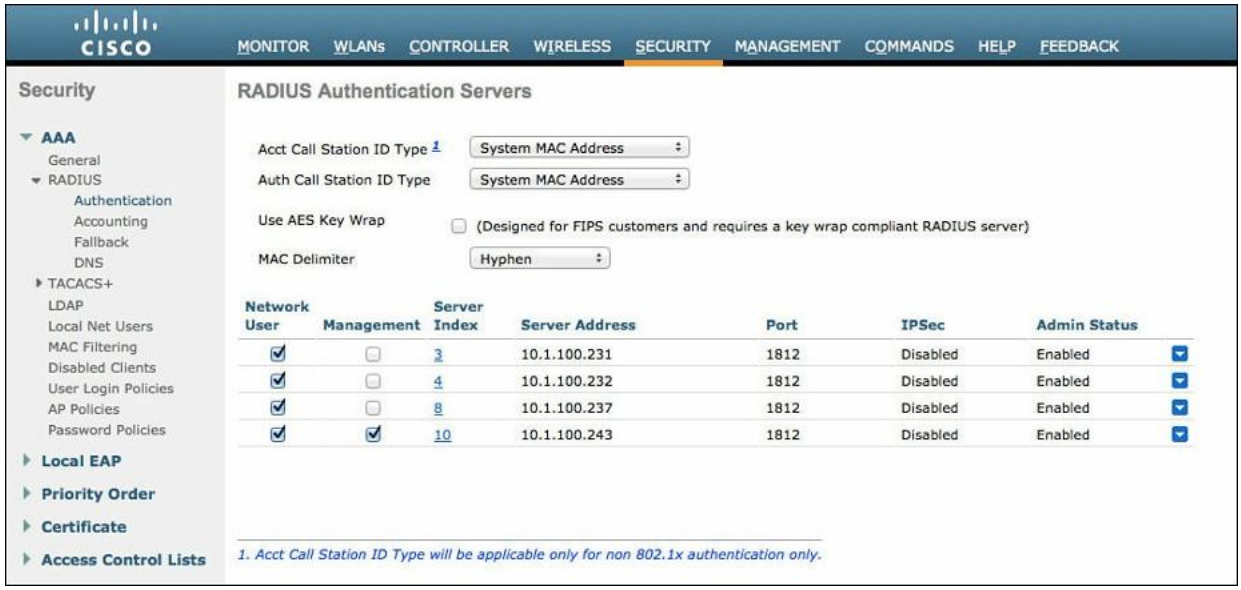


Figure 12-3 Security > RADIUS > Authentication.

Step 4. Click **New** to add the ISE PSN.

Step 5. Enter the IP address of the PSN (or the VIP, if using a load balancer).

Step 6. Fill in the **Shared Secret** field. This must match what is configured in the ISE network device object.

Step 7. The port must be 1812 for authentication.

Step 8. Ensure the server is enabled.

Step 9. Ensure that RFC 3576 (CoA) is enabled.

Step 10. Set the default server timeout to 2 seconds.

https://t.me/learningnets

Step 11. Ensure network user is enabled. This is simply stating that the RADIUS server can be used for network authentications.

Step 12. Click **Apply** in the upper-right corner.

Step 13. Save the configuration.

Figure 12-4 shows a completed server configuration.



Figure 12-4 Completed server configuration.

Repeat these steps for each PSN you need to add.

Adding the RADIUS Accounting Servers

As mentioned previously, the WLC separates the AAA server configuration into two sections. Now we will define the ISE PSN in the accounting server section (see Figure 12-5). From the WLC GUI, follow these steps:

Step 1. Navigate to **Security > RADIUS > Accounting**.

Step 2. Ensure the MAC Delimiter is set to **Hyphen**.



Figure 12-5 RADIUS accounting servers.

Step 3. Click **New** to add the ISE PSN.

Step 4. Enter the IP address of the PSN.

https://t.me/learningnets

Step 5. Enter the shared secret to match what is configured on ISE.

Step 6. Ensure the port is 1813.

Step 7. Ensure the server is enabled.

Step 8. Verify that the **Network User** check box is enabled.

Figure 12-6 shows a completed server entry.



Figure 12-6 Completed accounting server configuration.

Step 9. Click **Apply** in the upper-right corner.

Step 10. Save the configuration.

Repeat these steps for each PSN you need to add.

Configuring RADIUS Fallback (High-Availability)

The primary RADIUS server (the server with the lowest server index) is assumed to be the most preferable server for the 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 either the primary RADIUS server when it recovers and becomes responsive or a more preferable server from the available backup servers.

From the WLC GUI, do the following:

Step 1. Navigate to **Security** > **AAA** > **RADIUS** > **Fallback**.

Step 2. Set the Fallback Mode to **Active**.

Selecting Active causes the Cisco 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 the **Username,** enter the name to be sent in the inactive server probes.

We have been using radius-test as the username so far in the book. It's important to note, though, that you do not need to enter a password for this test user account because the system will simply look for a response from the RADIUS server; pass or fail does not matter.

Step 4. Enter a value for the **Interval in Sec** field.

The interval states the inactive time in passive mode and probe interval in active mode. The valid range is 180–3600 seconds, and the default value is 300 seconds. Figure 12-7 shows the completed RADIUS fallback settings.

https://t.me/learningnets

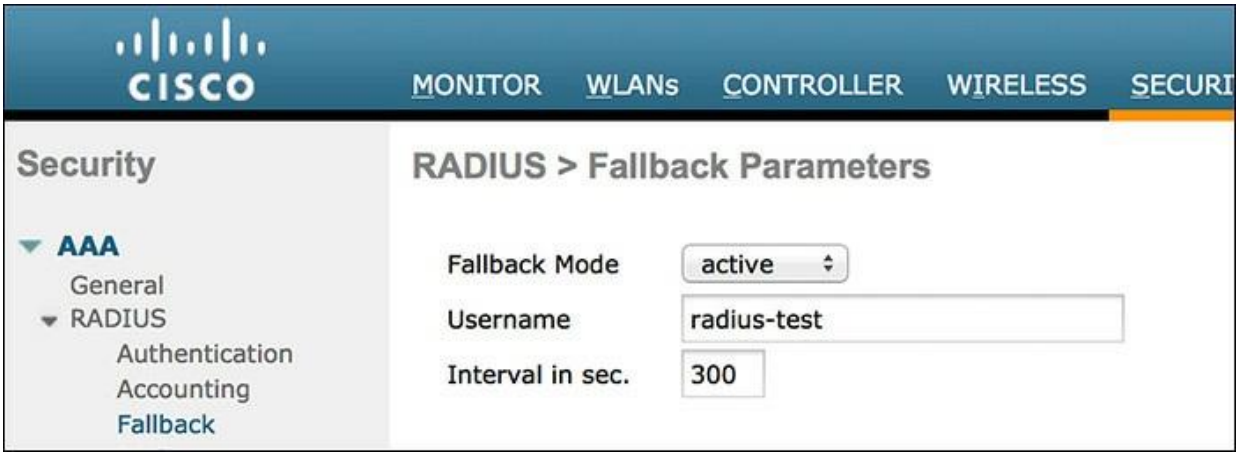


Figure 12-7 RADIUS fallback settings.

Configuring the Airespace ACLs

Just as we did with the Cisco catalyst switches, we will prestage the WLC with some access lists for:

- A web authentication ACL (ACL-WEBAUTH-REDIRECT)
- Posture agent redirection (ACL-AGENT-REDIRECT)

Notice that we are using the same ACL names as the Cisco switches. This is not required, but we do it this way to ensure consistency where possible.

Creating the Web Authentication Redirection ACL

As with the catalyst switches, we will need a local ACL on the wireless controller that will redirect web traffic to the centralized web authentication portal. However, with the 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 (permit or deny traffic flow).

With the Cisco WLC, there is a single access list, and it pulls double duty. It permits and denies traffic flow; at the same time, though, the traffic that is denied is redirected to the centralized web authentication portal.

From the WLC GUI, do the following:

Step 1. Navigate to **Security > Access Control Lists > Access Control Lists**.

Step 2. Click **New** to add a new ACL (see Figure 12-8).

https://t.me/learningnets



Figure 12-8 ACLs.

Step 3. Fill in the name **ACL-WEBAUTH-REDIRECT** (see Figure 12-9).

https://t.me/learningnets



Figure 12-9 Naming the ACL.

Step 4. Click **Apply**.

You will be returned to the main access list screen.

Step 5. Click the new entry, **ACL-WEBAUTH-REDIRECT**, as shown in Figure 12-10.

https://t.me/learningnets



Figure 12-10 Listing of ACLs.

Step 6. Click **Add New Rule** in the upper-right corner.

A rule in the WLC is the equivalent of an access control entry in the switch. It is a line in the access list.

Step 7. You need to 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, which will redirect the rest

Figure 12-11 shows an example of a completed ACL.

https://t.me/learningnets

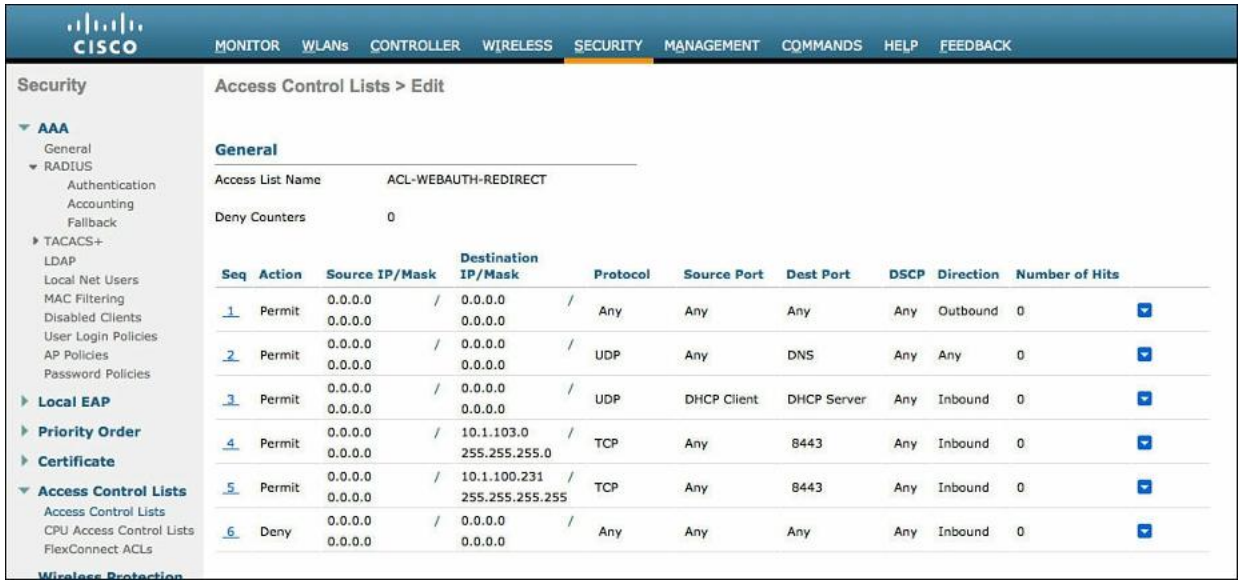


Figure 12-11 ACL-WEBAUTH-REDIRECT example.

Creating the Posture Agent Redirection ACL

Just as with the web authentication traffic, you should have an ACL for posture assessments. When a client successfully authenticates to the network but their 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, just like with web authentication.

The difference here is that we must open some more ports to allow even more traffic to flow than with web authentication. Technically, there is nothing to prevent you from using a single ACL for both web authentication and posture agent redirection if you are willing to open up all that traffic in both use cases.

To maintain consistency throughout our deployment examples, we will create a separate ACL here.

From the WLC GUI, do the following:

Step 1. Navigate to **Security > Access Control Lists > Access Control Lists**.

Step 2. Click **New** to add a new ACL.

Step 3. Enter the name **ACL-AGENT-REDIRECT**.

Step 4. Click **Apply**.

Step 5. Click 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 servers
- Permits TCP and UDP Port 8905 and 8909 to the ISE servers
- Permits applicable traffic to the remediation servers (the servers that will fix the client if it fails the posture assessment; could be patch management systems or even external websites)
- Denies all other traffic, which will redirect the rest

Figure 12-12 shows an example of a completed ACL

https://t.me/learningnets

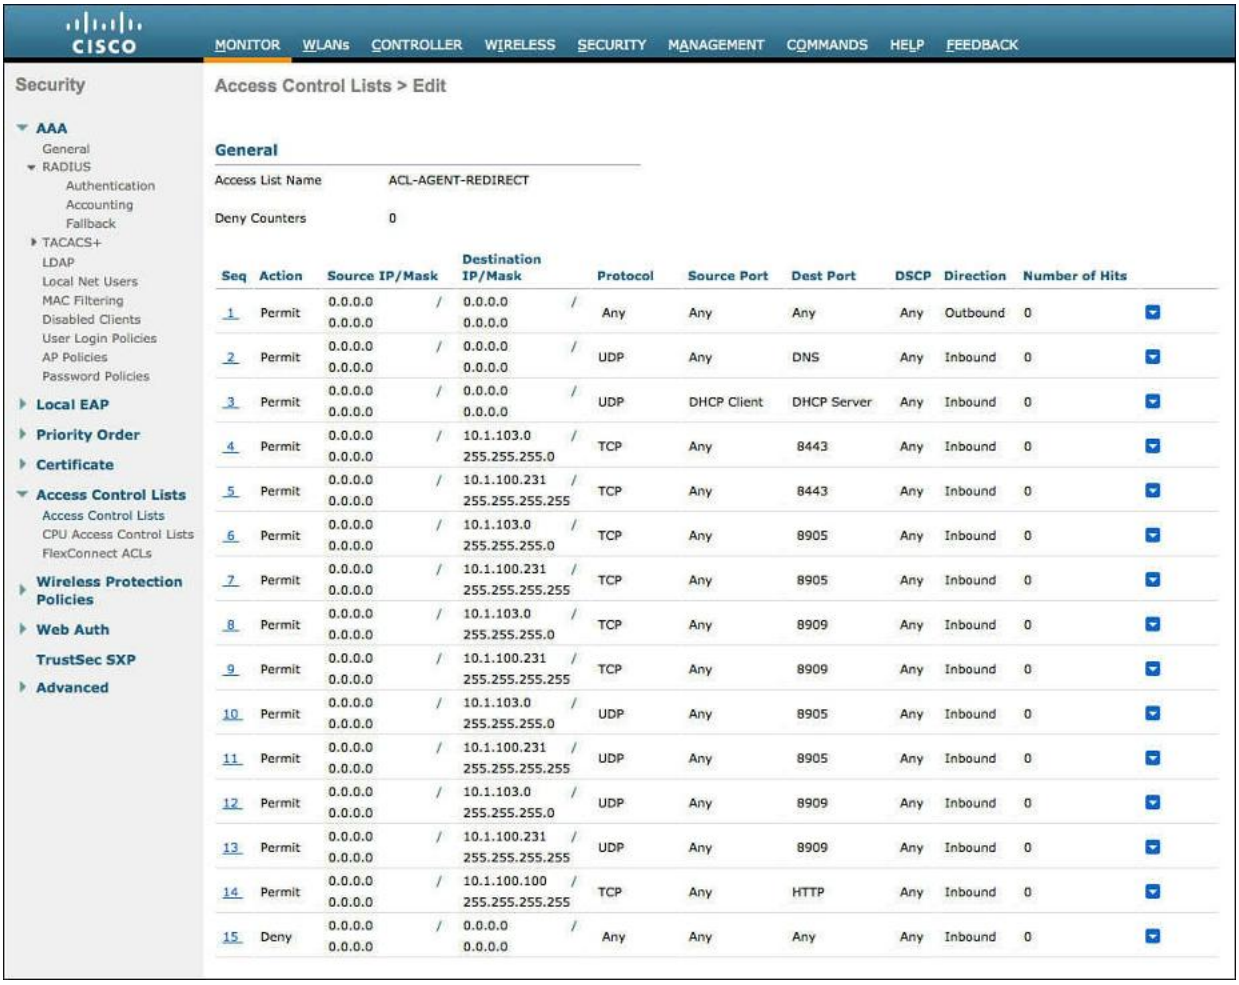


Figure 12-12 ACL-AGENT-REDIRECT example.

Creating the Dynamic Interfaces for the Client VLANs

When you want to assign a user or device to a VLAN on a catalyst switch, you simply assign the VLAN to the port, and the entire switch port is then 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 networks (Wi-Fi) to the physical wired network. It also must have the capability to assign different VLANs per authenticated session (if necessary). I know, you're thinking it just needs to be connected with a trunk, right? Well, yes, that's true.

The WLC will be configured to use 802.1Q to tag traffic for a specific VLAN as that traffic exits the controller. However, the controller will call this a dynamic interface because of the way it can assign a physical interface to traffic or assign a tag.

We will create two dynamic interfaces in this chapter—one for employee traffic and one for guest traffic.

Creating the Employee Dynamic Interface

This interface is used for all successful authentications to the corporate wireless LAN, providing full access to the entire network.

From the WLC GUI, do the folowing:

Step 1. Navigate to **Controller > Interfaces**.

Step 2. Click **New**.

Step 3. Name your Interface. We will use the name employee in our example.

https://t.me/learningnets

Step 4. Provide the VLAN ID to be used in the 802.1Q tag.

Step 5. Click **Apply**.

Step 6. Click the new interface named **employee**.

Step 7. You will most likely leave the settings alone until you reach the physical information section.

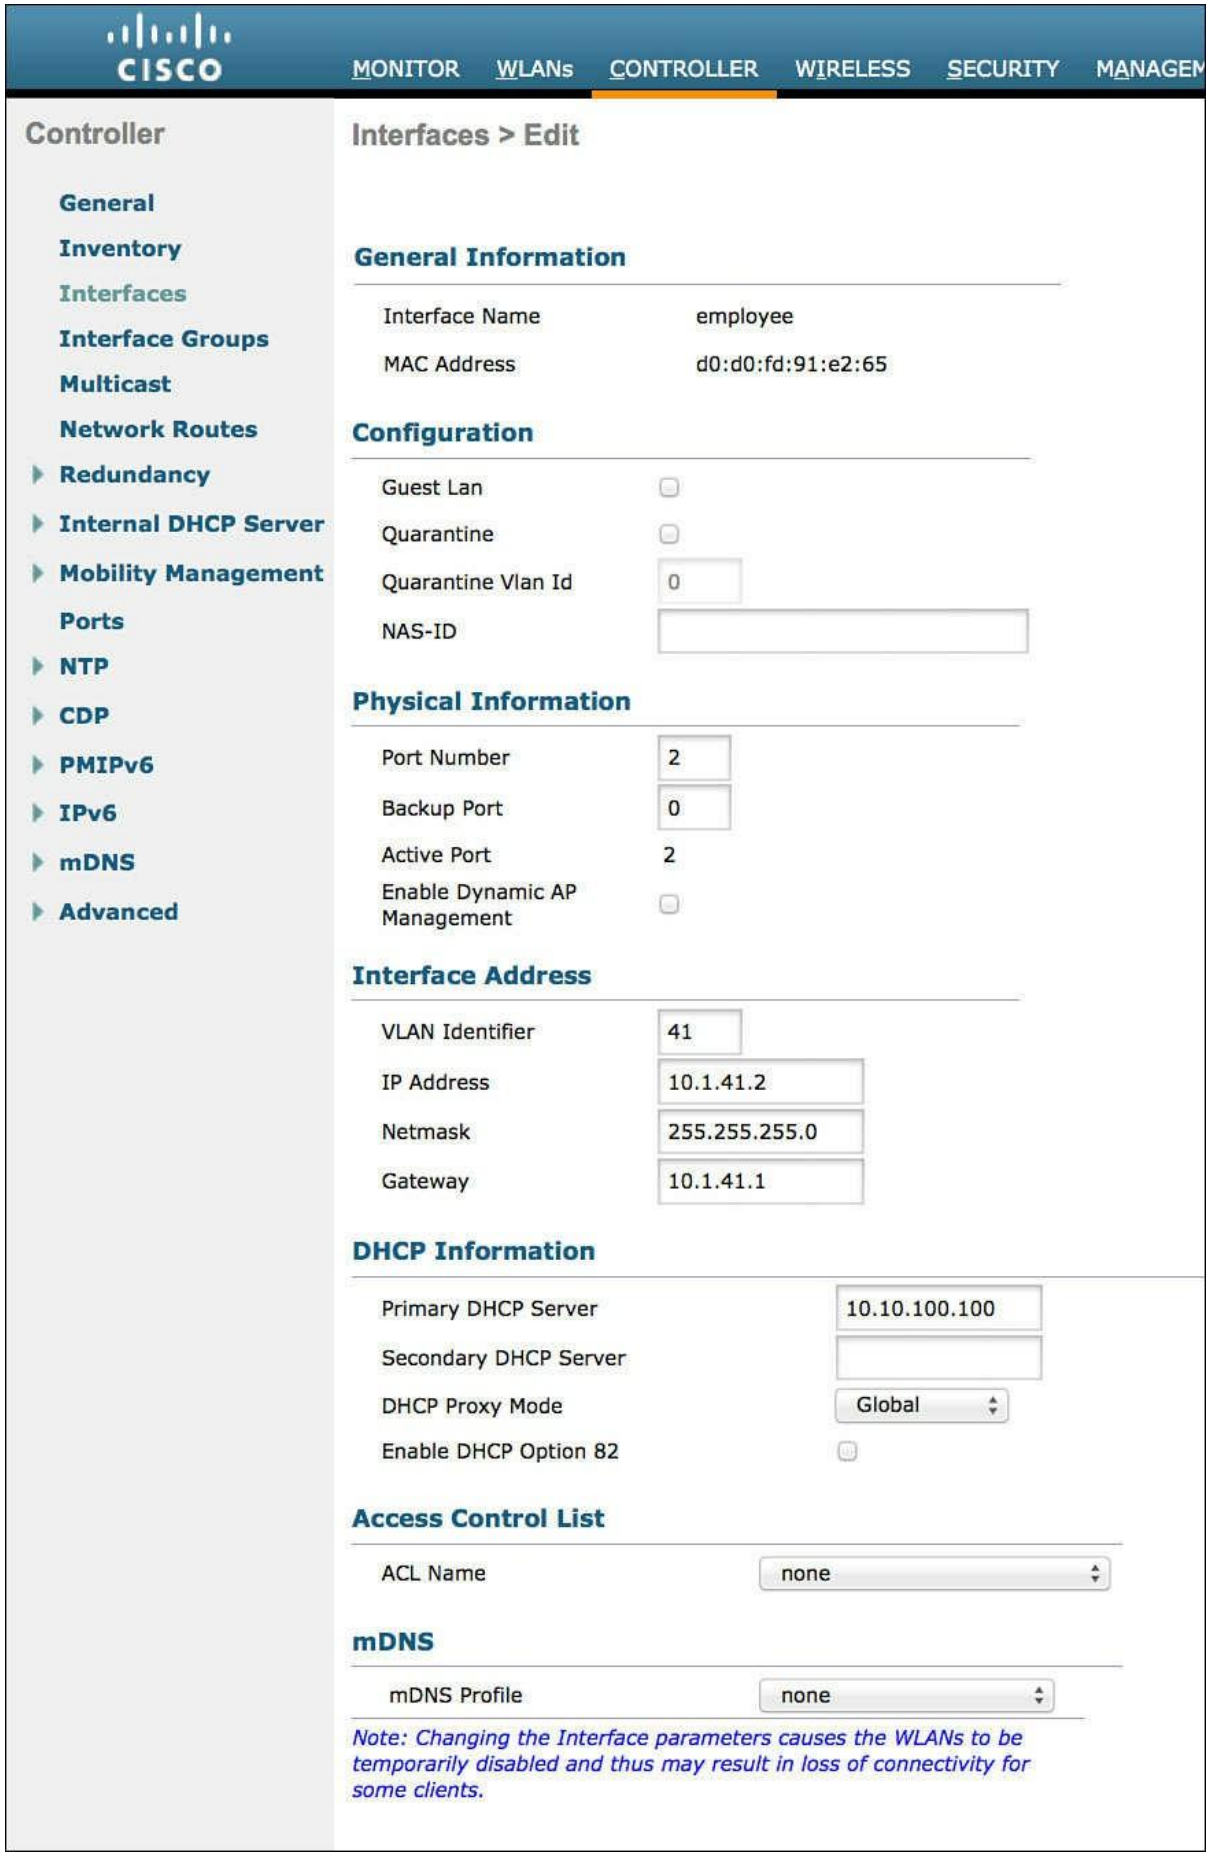
Step 8. Provide an IP address, netmask, and gateway for the VLAN.

Step 9. Provide the DHCP server address.

Step 10. Click **Apply**.

Figure 12-13 shows a sample employee dynamic interface configuration.

https://t.me/learningnets



https://t.me/learningnets

Figure 12-13 Employee dynamic interface.

Creating the Guest Dynamic Interface

This interface is used for all devices connecting to the GUEST WLAN, as well as unsuccessful or unauthorized authentications to the corporate wireless LAN. This interface will have Internet access only.

From the WLC GUI, do the folowing:

Step 1. Navigate to **Controller > Interfaces**.

Step 2. Click **New**.

Step 3. Name your Interface. We will use the name guest in our example.

Step 4. Provide the VLAN ID to be used in the 802.1Q tag.

Step 5. Click **Apply**.

Step 6. Click the new interface named **employee**.

Step 7. You will most likely leave the settings alone until you reach the physical information section.

There is a Guest LAN check box—do not enable it. This is not for GUEST WLANs; it is for providing guest access to directly connected wired LANs.

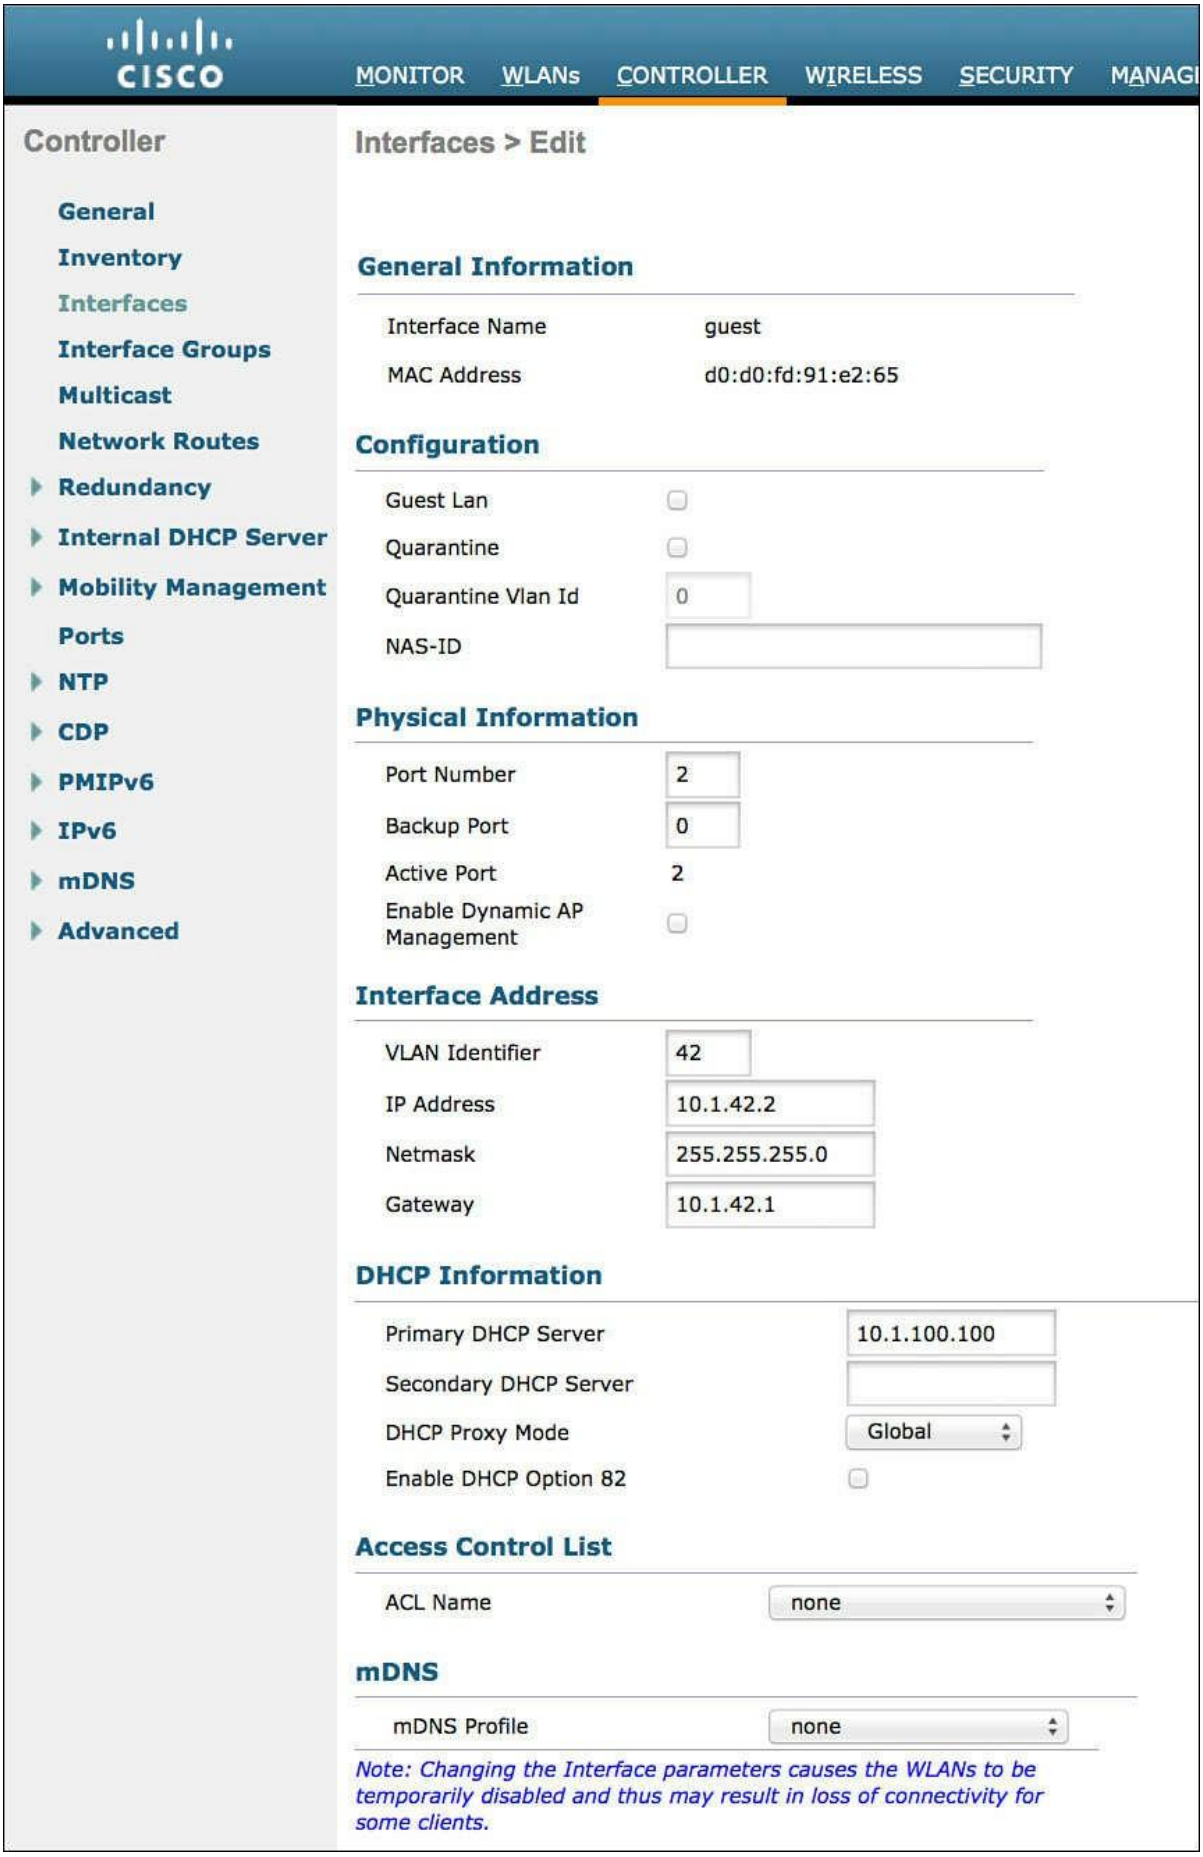
Step 8. Enter an IP address, netmask, and gateway for the VLAN.

Step 9. Enter the DHCP server address.

Step 10. Click **Apply**.

Figure 12-14 shows a sample guest dynamic interface configuration.

https://t.me/learningnets



https://t.me/learningnets

Figure 12-14 Guest dynamic interface.

Creating the Wireless LANs

Now that the RADIUS servers, ACLs, and dynamic interfaces are all created and configured, we will move on to creating two WLANs: a guest WLAN and a corporate WLAN. The guest WLAN will be an open WLAN, whereas the corporate WLAN will be configured to use 802.1X to authenticate devices.

Creating the Guest WLAN

The guest WLAN will be created as an open SSID but will send the endpoint MAC addresses to ISE for MAB, just like the wired networks.

From the WLC GUI, follow these steps (see Figure 12-15):

Step 1. Using the headers across the top, navigate to **WLAN**.

Step 2. Select **Create New** in the drop-down menu.

Step 3. Click **Go**.

Step 4. Select **WLAN**.

Step 5. Give the WLAN profile a name; we will use CP-Guest in this example.

Step 6. Enter an SSID name; we will use CiscoPress-Guest in this example.



Figure 12-15 Guest WLAN creation.

> **Note**
>
> WLAN ID is a unique identifier for each WLAN/SSID on a controller. The default should be fine in most cases. However, the WLAN ID can be used as a condition within your authorization policy. If you choose to use the WLAN ID in your policy, ensure that the WLAN ID is the same across all WLCs; otherwise, you might get unwanted results.

Step 7. Click **Apply**.

General Tab

The General tab is where a WLAN is enabled or disabled, and it is where the default interface is selected:

Step 1. If you are ready to put this SSID in an operative state, ensure the **Enabled** check box is selected.

Step 2. Next to **Interface/Interface Group (G)**, select the GUEST interface we created previously from the drop-down menu. Figure 12-16 shows the General tab.

https://t.me/learningnets



Figure 12-16 GUEST WLAN General tab.

Step 3. Click the **Security** tab.

Layer 2 Security Tab

The Layer 2 Security tab relates to the manner of protection for which devices associate to the wireless network and what encryption should be used (if any). It defines the use of Wi-Fi protected access (WPA), wired equivalent privacy (WEP), or even the ability to do MAC filtering. For an open guest network, you will select no Layer-2 security and will enable MAC filtering, which is how you configure wireless MAB. Here's how:

Step 1. Change the Layer 2 Security from the default (WPA) to **None**.

Step 2. Check the box for **MAC Filtering** (which is MAB). Figure 12-17 shows the completed tab.

https://t.me/learningnets

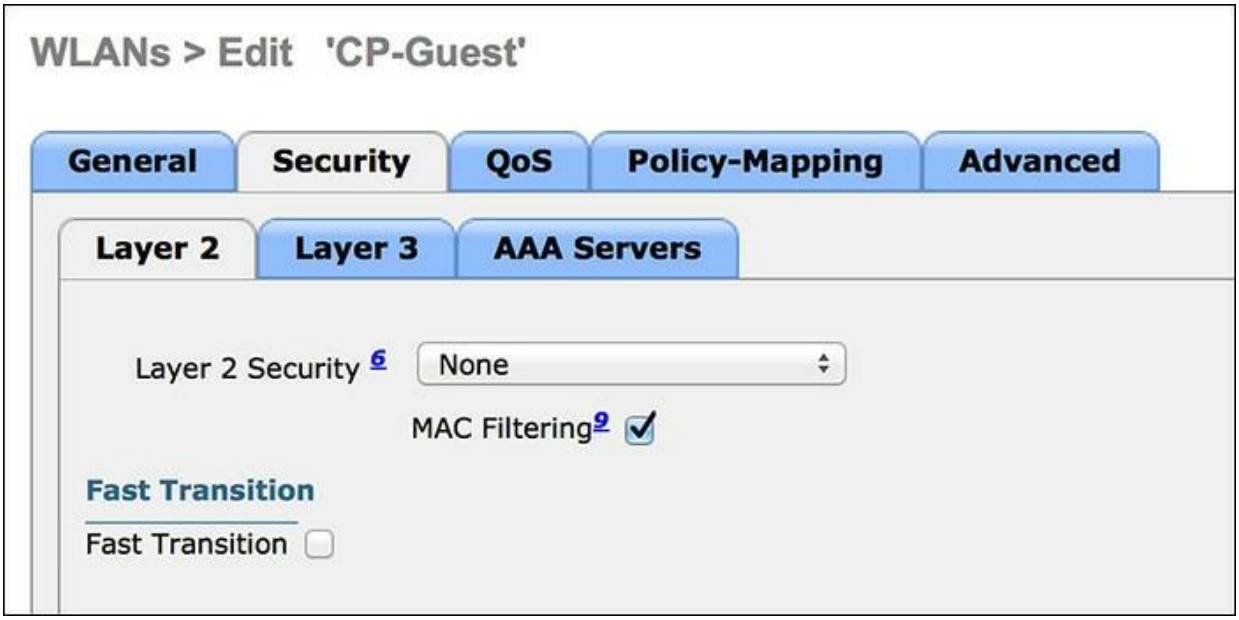


Figure 12-17 Layer 2 Security tab.

Step 3. Click the **Layer 3** tab.

Layer 3 Security Tab

The Layer-3 security options are used for local web authentication and should be set to None for an ISE deployment with centralized web authentication:

Step 1. Ensure the Layer-3 security option is set to **None**, as shown in Figure 12-18.



Figure 12-18 Layer 3 Security tab.

Step 2. Click the **AAA Servers** tab.

AAA Servers Tab

For the Cisco WLC to work correctly for all ISE services, the authentication and accounting servers must be set identically. While the GUI of the WLC permits them to be set for differing servers, that will break certain functionality related to ISE services:

Step 1. Select your ISE PSN(s) for both authentication and accounting.

Step 2. Click **Apply**. Figure 12-19 shows the completed tab.

https://t.me/learningnets

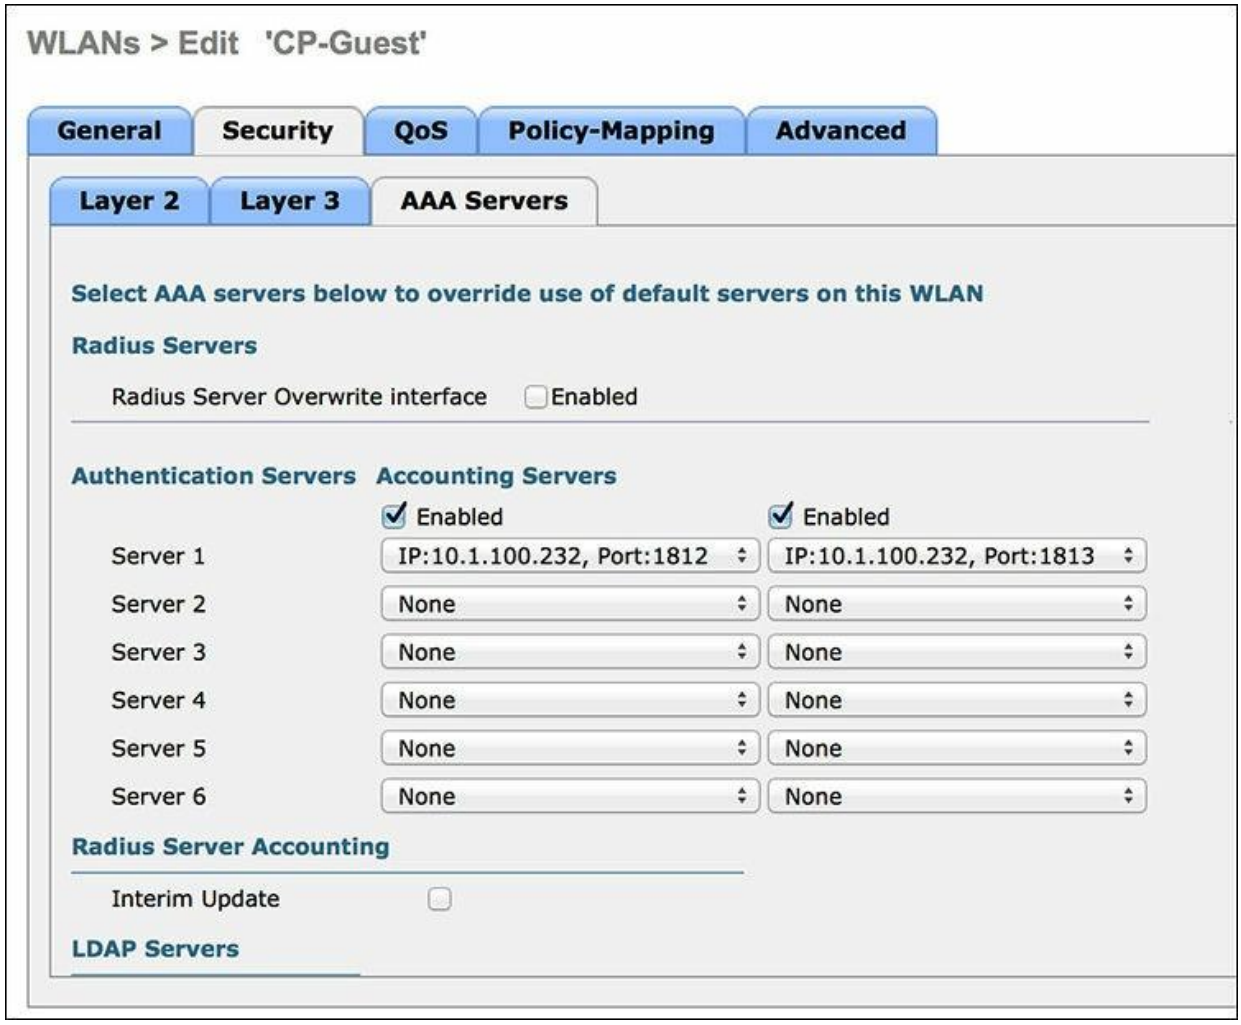


Figure 12-19 AAA Servers tab.

Step 3. Click the **Advanced** tab.

Advanced Tab

To allow ISE to override the assigned VLAN, interface, or even to assign an ACL, the Allow AAA Override setting must be enabled:

Step 1. Enable **Allow AAA Override**.

The WLC has built-in capabilities for the profiling of endpoints and sending that profiling data to ISE. The DHCP Addr. Assignment check box must be enabled for the DHCP Device Sensor built in to the WLC to function correctly.

Step 2. Check the box for **DHCP Addr. Assignment**.

A NAC State setting is available on the Advanced tab. 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 3. Change the NAC State to **RADIUS NAC**.

Step 4. Scroll down and enable both the **DHCP** and **HTTP RADIUS Client Profiling** options. This is revisited in Chapter 15, "Profiling," where endpoint profiling is covered in detail.

Step 5. Click **Apply**.

Step 6. Save the configuration.

Figures 12-20 and 12-21 display the completed Advanced tab.

https://t.me/learningnets

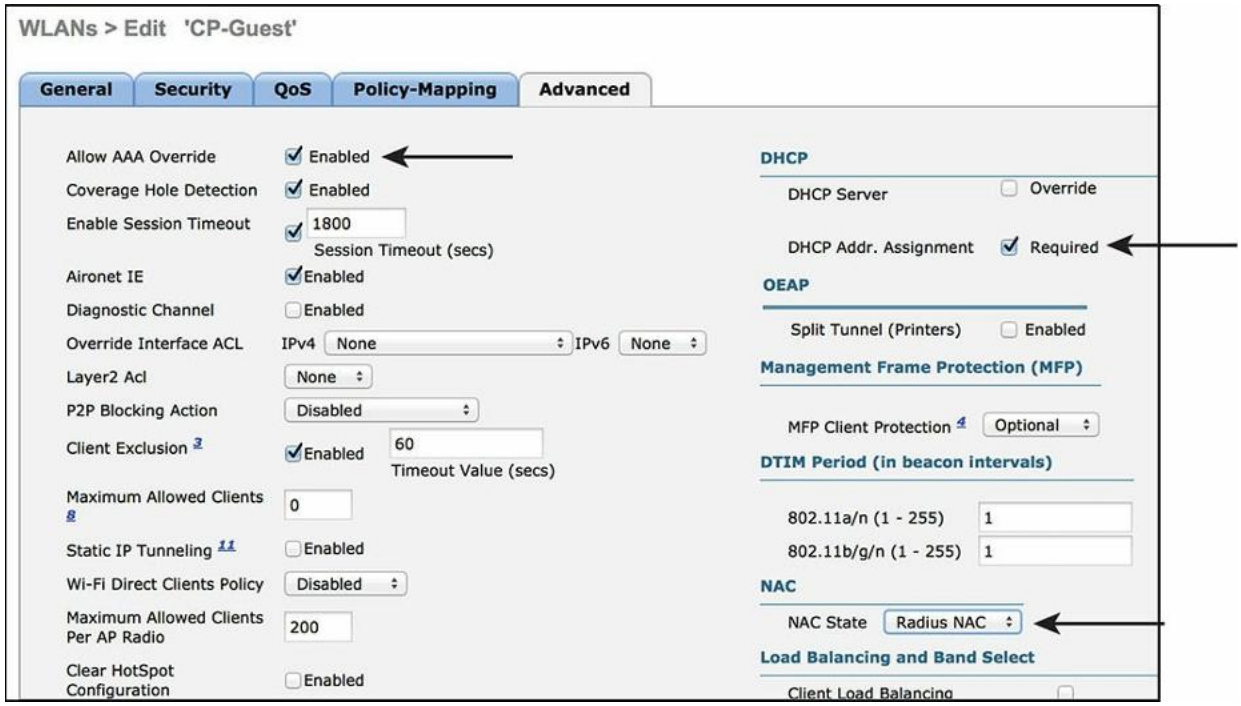


Figure 12-20 Advanced Tab part 1.

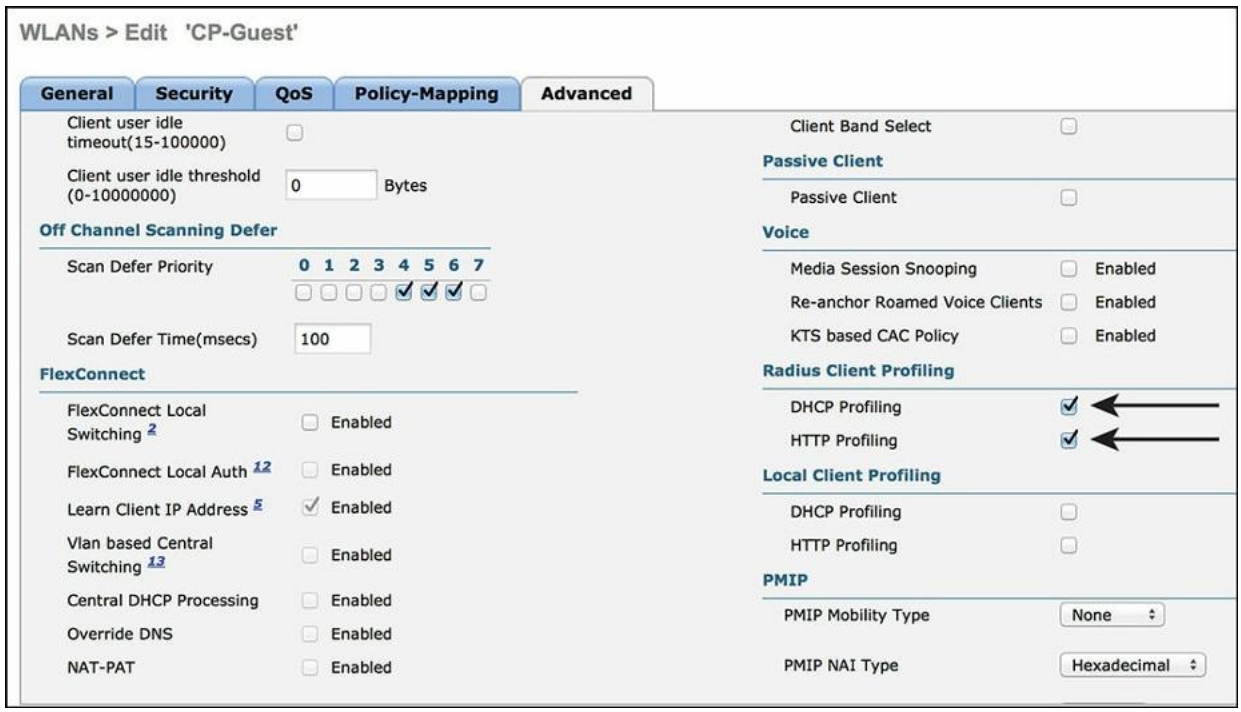


Figure 12-21 Advanced Tab part 2.

Creating the Corporate SSID

The corporate WLAN will be created as a closed SSID and will require 802.1X authentication for an endpoint to associate to the WLAN. Unlike wired networks, wireless networks have the added benefit of truly rejecting all access without a successful authentication. Users are attuned to the requirement of configuring software to connect to a wired network. The same is very much not true when it comes to wired networks.

From the WLC GUI, do the following:

https://t.me/learningnets

Step 1. Navigate to WLAN.

Step 2. Select **Create New**.

Step 3. Click **Go**.

Step 4. Select **WLAN**.

Step 5. Give the WLAN profile a name; we will use CiscoPress in this example.

Step 6. Enter an SSID name; we will use CiscoPress in this example.

Step 7. Click **Apply**.

Figure 12-22 shows the new WLAN being added.



Figure 12-22 CiscoPress WLAN.

General Tab

Step 1. If you are ready to put this SSID into an operative state, ensure the **Enabled** check box is selected.

Step 2. Next to the Interface/Interface Group (G), select the employee interface we created previously from the drop-down.

Step 3. Click the **Security** tab.

Figure 12-23 shows the final General tab.

https://t.me/learningnets

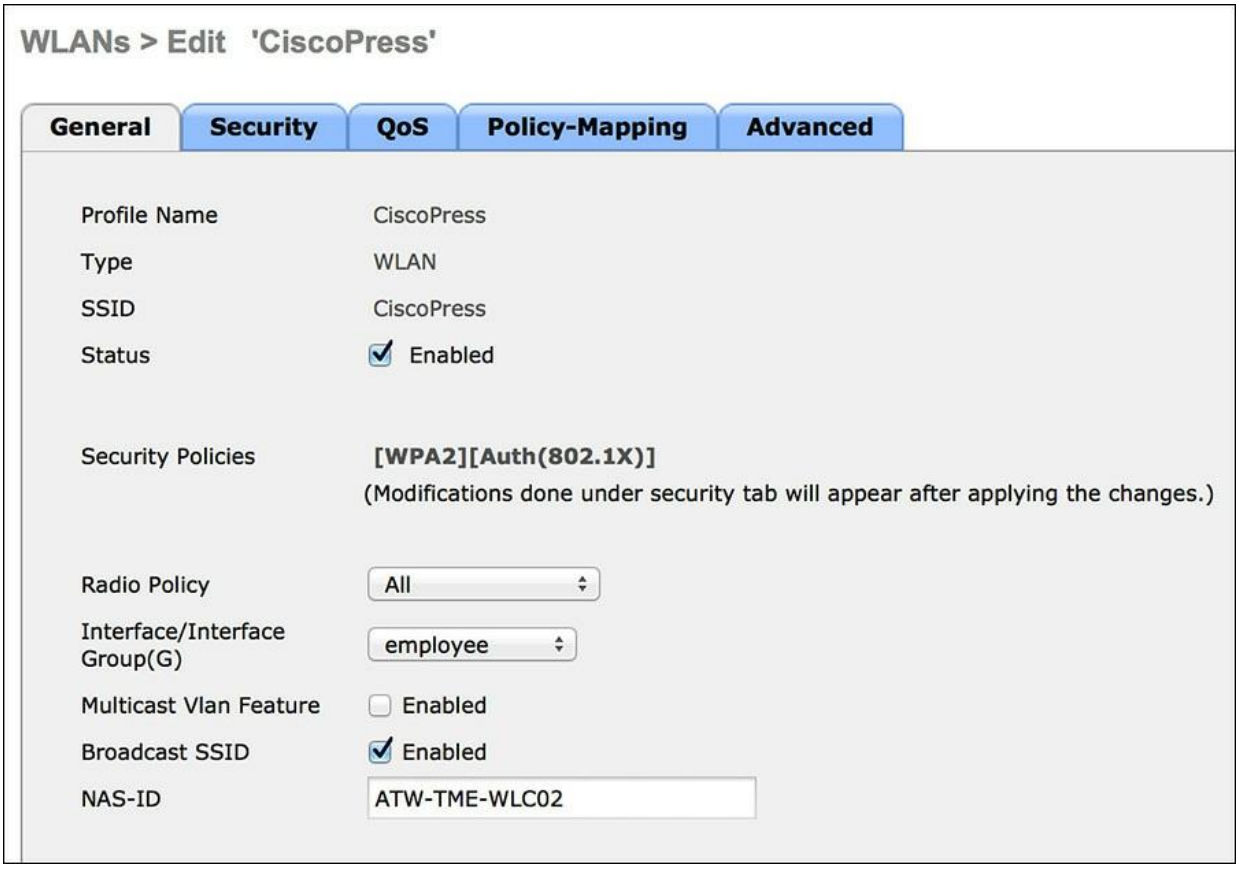


Figure 12-23 General tab.

Layer 2 Security Tab

The Layer 2 Security tab relates to the manner of protection for which devices associate to the wireless network and what encryption should be used (if any). It defines the use of WPA, WEP, or even the ability to perform MAC filtering. Do the following:

Step 1. The default setting of WPA+WPA2 is correct.

Step 2. You will not be enabling MAC filtering.

Step 3. Ensure that 802.1X is checked for **Authentication Key Management**.

Figure 12-24 shows the final Layer 2 Security tab.

https://t.me/learningnets

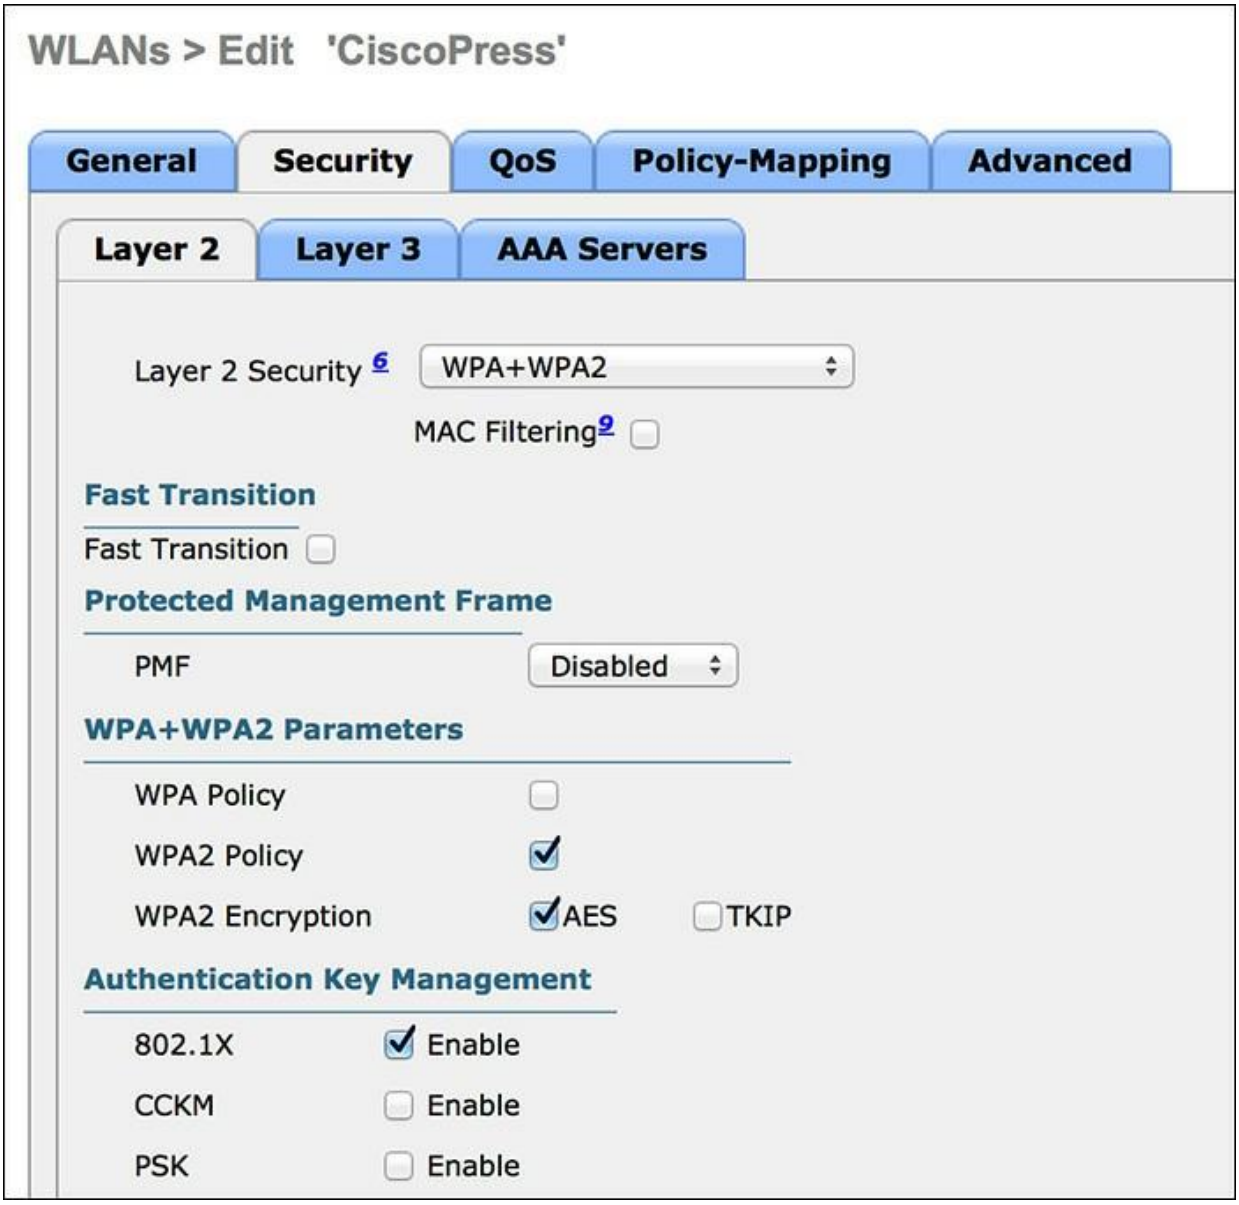


Figure 12-24 Layer 2 Security tab.

Step 4. Click the **Layer 3** tab.

Layer 3 Security Tab

Step 5. Ensure the Layer 3 security is set to "**None**" (see Figure 12-25).

https://t.me/learningnets



Figure 12-25 Layer 3 Security tab.

Step 6. Click the **AAA Servers** tab.

AAA Servers Tab

For the Cisco WLC to work correctly for all ISE services, the authentication and accounting servers must be set identically. Although the GUI of the WLC permits them to be set for differing servers, that will break certain functionality related to ISE services:

Step 1. Select your ISE PSN(s) for both authentication and accounting (see Figure 12-26).

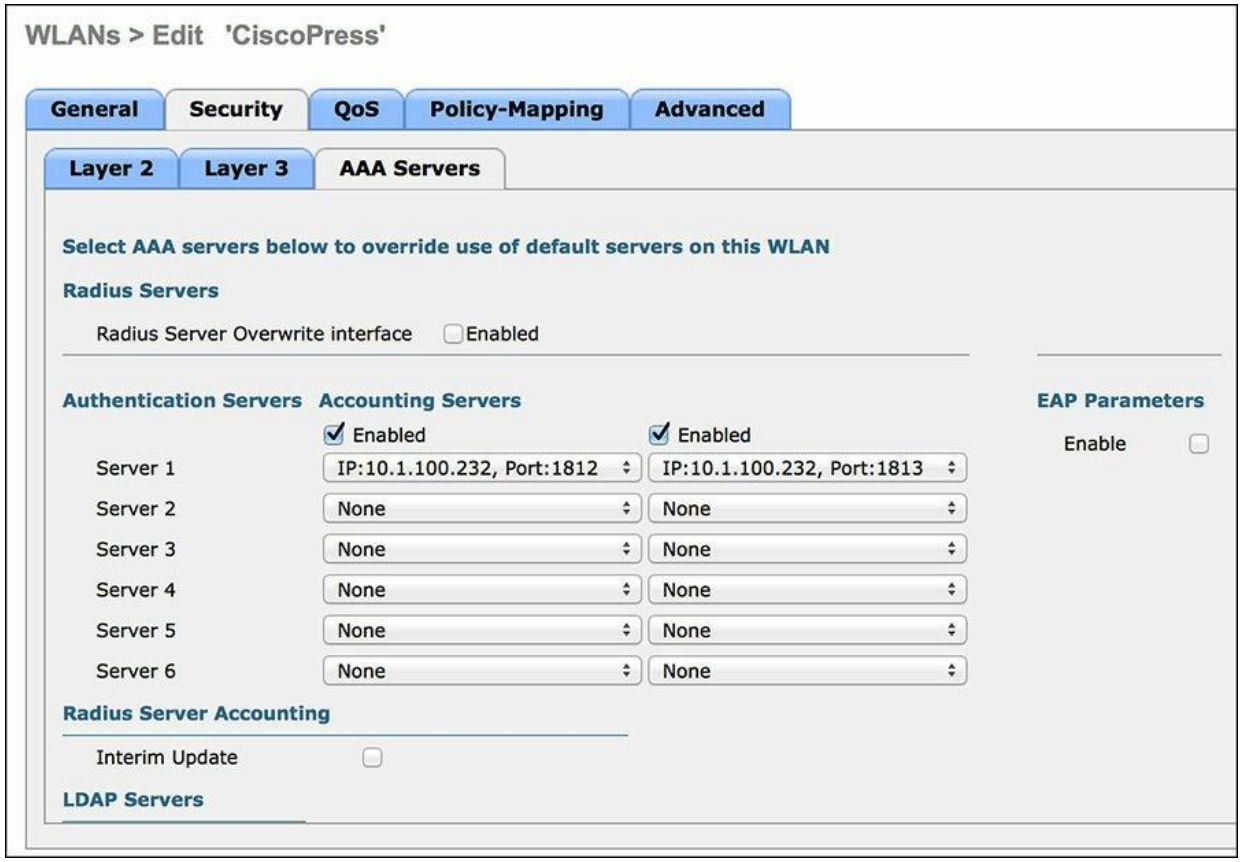


Figure 12-26 AAA Servers tab.

Step 2. Click **Apply**.

Step 3. Click the **Advanced** tab.

https://t.me/learningnets

Advanced Tab

To allow ISE to override the assigned VLAN or interface, or even to assign an ACL, the Allow AAA Override setting must be enabled:

Step 1. Enable **Allow AAA Override**.

The WLC has built-in capabilities for the profiling of endpoints and sending that profiling data to ISE. The DHCP Addr. Assignment check box will need to be enabled for the DHCP device sensor built in to the WLC to function correctly.

Step 2. Check the box for **DHCP Addr. Assignment**.

A NAC State setting exists on the Advanced tab. 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 3. Change the NAC State to **RADIUS NAC**.

Step 4. Scroll down and enable both the **DHCP** and **HTTP RADIUS Client Profiling** options. This is revisited in Chapter 15, where endpoint profiling is covered in detail.

Step 5. Click **Apply**.

Step 6. Save the configuration.

Step 7. Figures 12-27 and 12-28 display the completed Advanced tab.

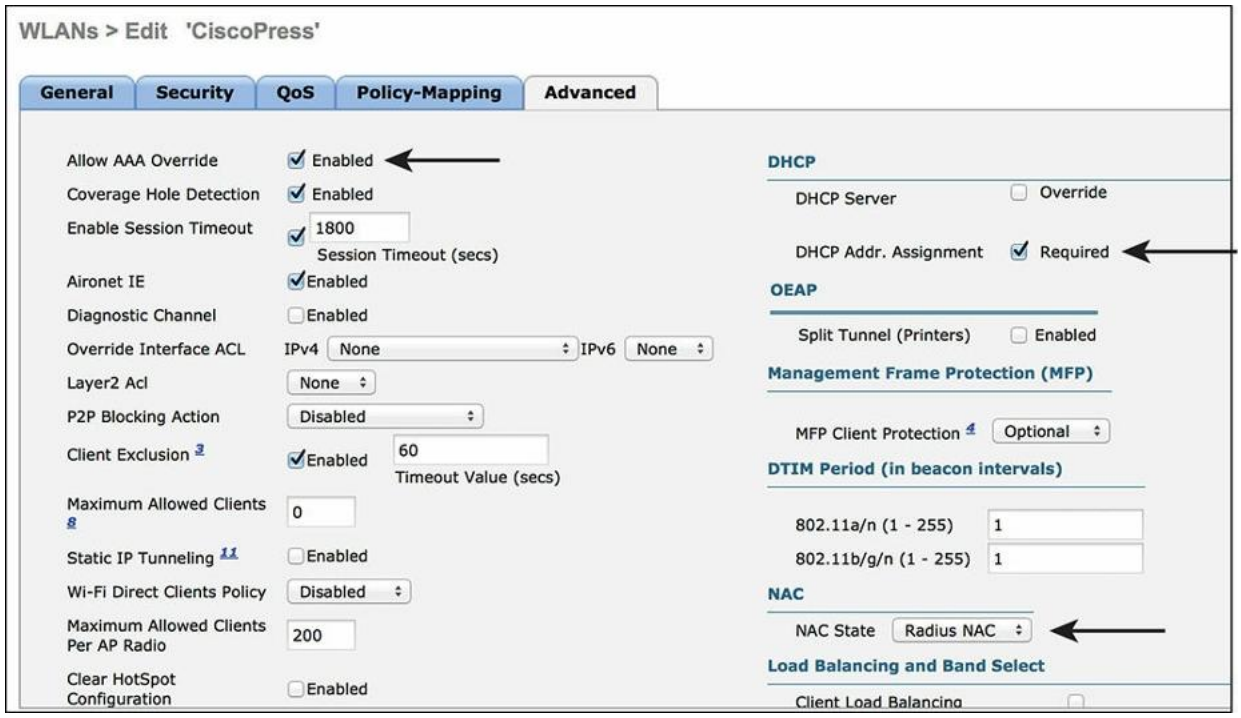


Figure 12-27 Advanced Tab part 1.

https://t.me/learningnets

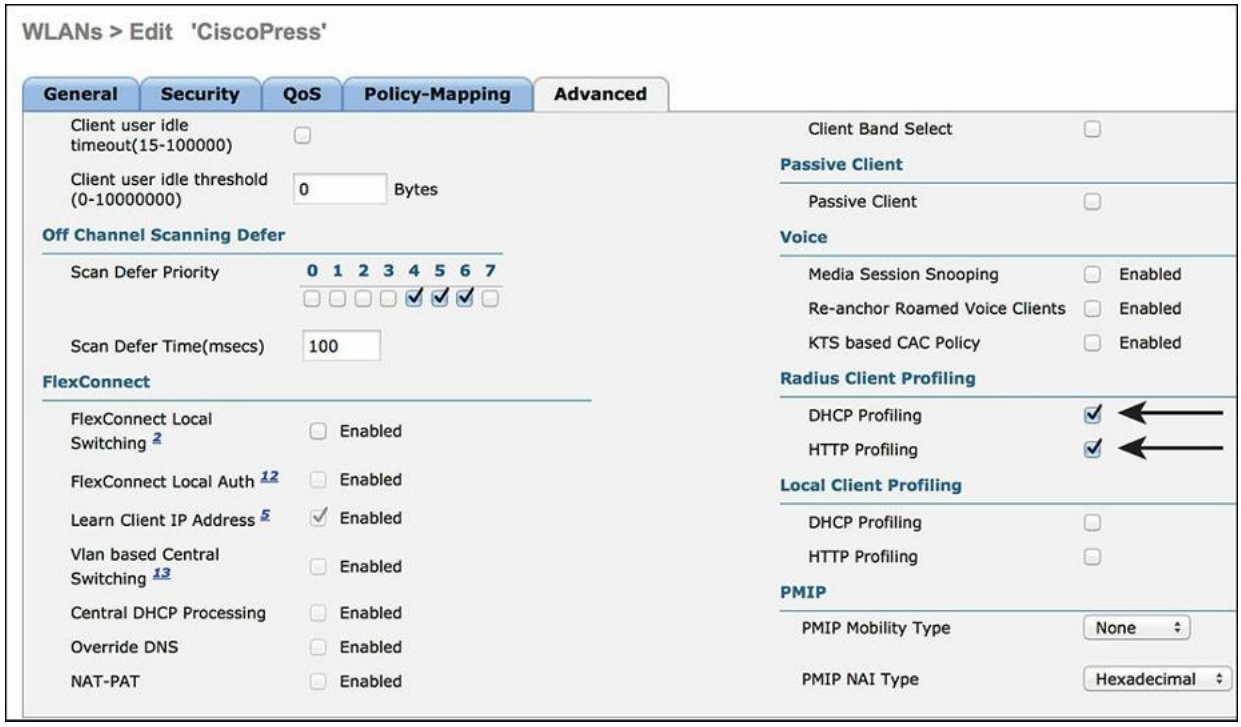


Figure 12-28 Advanced Tab part 2.

Verifying Dot1X and MAB

There are numerous ways to verify the authentication operations of switches and wireless controllers. Three locations always must be examined to validate a complete end-to-end transaction. Two of those three locations are much more common and easy to use. The three locations are:

- Endpoint supplicant—For 802.1X authentications
- Network access device—For all authentications
- Cisco ISE—For all authentications

Endpoint Supplicant Verification

Verifying the authentications from the supplicant is a bit outside of the exam blueprint, so this book will not focus on it much. With Cisco AnyConnect NAM as your supplicant, you can use the DART tool to get a detailed communication, even perform packet captures at the endpoint. If the supplicant is an Apple supplicant (MAC-OSX 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-line **netsh ras set tracing * enable** to enable the supplicant's logging capabilities. Once enabled, the logs are added to the **%systemroot%\tracing** folder.

Network Access Device Verification

Two NADs are focused on here: Cisco switches and Cisco wireless LAN controllers. Each is quite different in how authentications are verified and will therefore be discussed in two separate sections.

Verifying Authentications with Cisco Switches

Many items can be tested with a Cisco switch, with many tools being provided in Cisco IOS. The ones used most often are described in this section.

show AAA servers **Command**

https://t.me/learningnets

One of the first things to check with a Cisco switch is the status of the RADIUS server (ISE). 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 12-1 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 12-1, the current state is UP. Use this command to validate that the server is up. If it is down, communication to the RADIUS server will not occur.

Example 12-1 show aaa servers Command

Click here to view code image

```
3750-X#sho aaa servers

RADIUS: id 1, priority 1, host 10.1.100.232, 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
3750-X#
```

test AAA **Command**

Cisco switches have a built-in mechanism to send test authentications to the AAA servers they are configured to use. Using the **test aaa** command, you can verify that an authentication is successfully sent to and received by the RADIUS server. Example 12-2 shows the use of the **test aaa** command and the successful response. The **test** command will send an authentication request using PAP_ASCII and return 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 be sure the authentication is not being rejected for that reason.

Example 12-2 test aaa Command

Click here to view code image

https://t.me/learningnets

```
3750-X# test aaa group radius employee1 Cisco123 legacy
Attempting authentication test to server-group radius using radius
User was successfully authenticated.

3750-X#
```

show authentication session <interface> Command

One of the go-to commands that is in every implementer's bag of tools is the **show authentication session** <interface> command. Yes, the interface option is added to the base command of **show authentication session**, but that is to provide more detail. Example 12-3 shows the use of this command and the output, which displays a successful MAB authentication. Use this command to validate that the authentications are being attempted, which are successful, what authorization results have been assigned, and much more.

Example 12-3 show authentication session interface Command

Click here to view code image

```
3750-X# show authentication session int 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
     Session timeout: N/A
        Idle timeout: N/A
   Common Session ID: 0A013002000000110073D1F6
     Acct Session ID: 0x00000002
              Handle: 0xA9000012

Runnable methods list:
       Method   State
       mab      Authc Success
       dot1x    Not run
```



Many facets of the authentication session are displayed in this commands output. Because 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 controlling the authentication session.
- **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 which IP addresses are associated to the endpoints connected to the switch interface. This applies to both static IP addresses and DHCP assigned addresses. After the switch has learned the endpoint's IP address, it will be 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.

https://t.me/learningnets

- **Status**—This lists the status of the authentication session. Status can be 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-host modes, there are two domains: DATA and VOICE. Each authentication session can be assigned to one and only one of the domains.
- **Security Policy**—A better name for this would be MACSec Policy because that is exactly what the field is referring to. MACSec is the friendly name for IEEE 802.1AE, a Layer-2 encryption standard. The three options are Should Secure, Must secure, and Must Not Secure.
- **Security Status**—This field displays the current MACSec encryption applied. When secure, there is encryption. When unsecure, there is no encryption.
- **Oper host mode**—Lists the host-mode of the switch interface. Single-mode, multi-domain, multi-auth, and mutli-host are the available modes of operation.
- **Common Session ID**—The session ID is used to correlate authentication session information between the NAD and Cisco RADIUS server. When troubleshooting, it is often needed to compare this value with the one seen within Cisco ISE.
- **Runnable methods**—The available methods are mab, dot1x, and webauth. The possible states of the methods are Not Run, Running, Failed over, Authc Succeeded, and Authc Failed.

Sending Syslog to ISE

Syslog can be generated on Cisco IOS software in many events. Some of the syslog message can be sent to the ISE Monitoring Node (MNT) to be used in troubleshooting purposes. The ISE MNT node will correlate the syslog data with the RADIUS data and display both together in a report. This can be useful when checking to see whether dACLs have been applied correctly, as well as other important validations.



It is not recommended to enable the sending of syslog messages to ISE from all NADs at all times, but to enable it only when troubleshooting.

To ensure Cisco ISE is able to compile appropriate syslog messages from the switch, use the following commands:

Step 1. Enable syslog on the switch:

Click here to view code image

```
C3560X(config)#logging monitor informational
C3560X(config)#logging origin-id ip
C3560X(config)#logging source-interface <interface_id>
C3560X(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 dACL. Enabling EPM logging generates a syslog related to dACL authorization, and part of the log can be correlated inside Cisco ISE when such logs are sent to Cisco ISE.

Step 2. Set up standard logging functions on the switch to support possible troubleshooting and recording for Cisco ISE functions:

```
C3560X(config)#epm logging
```

Only the following NAD syslog messages are actually collected and used by Cisco ISE:

- AP-6-AUTH_PROXY_AUDIT_START

https://t.me/learningnets

- 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 WLC GUI, navigate to **Monitor** > **Clients**, as shown in Figure 12-29.

https://t.me/learningnets



Figure 12-29 Monitor menu.

As shown in Figure 12-30, the Clients screen displays all current clients associated to the WLC, along with very valuable information, such as:

- **IP Address**—The IP address of the endpoint, when known.
- **AP Name**—The name of the access point to which the endpoint is associated.
- **WLAN Profile**—The name of the WLAN profile created in the WLC.
- **WLAN SSID**—The name of the SSID for the WLAN profile, of 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 will be displayed.
- **Auth**—If the endpoint/supplicant has authenticated successfully, it will be listed here.
- **Device Type**—When the endpoint profile of the device is known, it will be displayed in this field.



Figure 12-30 Monitor > Clients.

Figure 12-30 shows the clients monitor screen.

For much more detail, click the endpoint's MAC address. This brings up the details related to the individual endpoint's wireless session. As shown in Figure 12-31, a key value to verify for authentication is that the 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.

https://t.me/learningnets

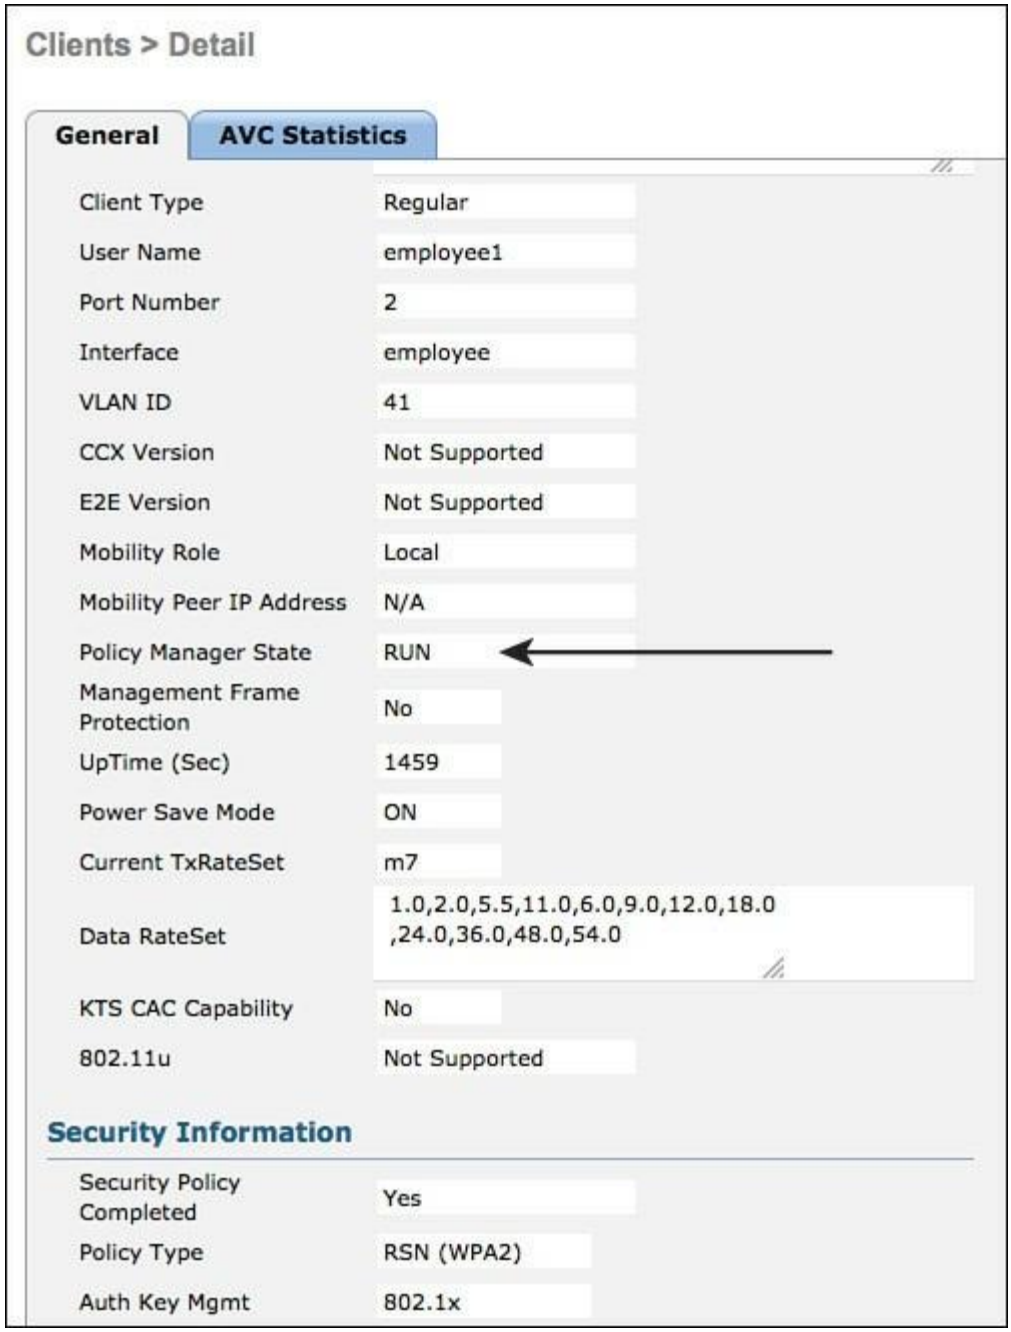


Figure 12-31 Monitor > Clients.

There are multiple states that an endpoint can be in. Those states are examined more in Chapter 13, "Web Authentication," and Chapter 17, "Bring Your Own Device."

Debug Client

The WLC provides a few useful debug commands: **debug dot1x** and **debug client** <mac-address>. **debug dot1x** can be a bit overwhelming on a live network, but the debug client command will show only events related to that specific endpoint. Example 12-4 shows an example of the debug client command.

Example 12-4 debug client <mac-address>

Click here to view code image

```
(Cisco Controller) >debug client 10:bf:48:d0:05:67
```

https://t.me/learningnets

```
*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 the authentication from the 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 the Live Authentications Log (commonly known as Live Log) and Live Sessions Log (commonly known as Live Sessions).

Live Authentications Log

To view the Live Authentications Log (Live Log), navigate to **Operations > Authentications**, as shown in Figure 12-32.

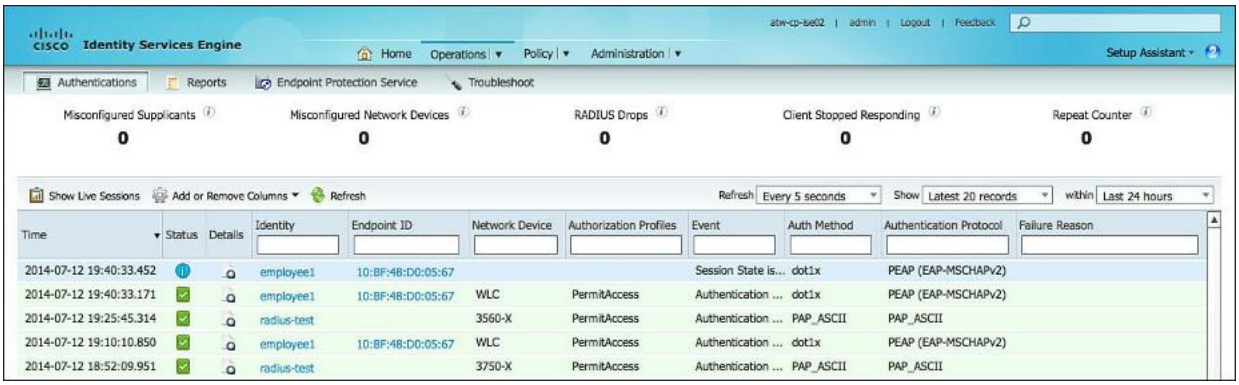


Figure 12-32 Live Authentications Log.

As shown in Figure 12-32, Live Log displays a near real-time display of authentication activity for the ISE Cube (also called ISE deployment). In Figure 12-32, you will see successful authentications from the radius-test account, which represents the automated testing performed by the wired switches. Additionally, the figure shows successful wireless authentications from employee1.

The Live Authentication Log is covered in greater detail throughout this book.

Live Sessions Log

To view the Live Sessions Log (Live Sessions), click the **Show Live Sessions** button in Live Log, as shown in Figure 12-33.

https://t.me/learningnets

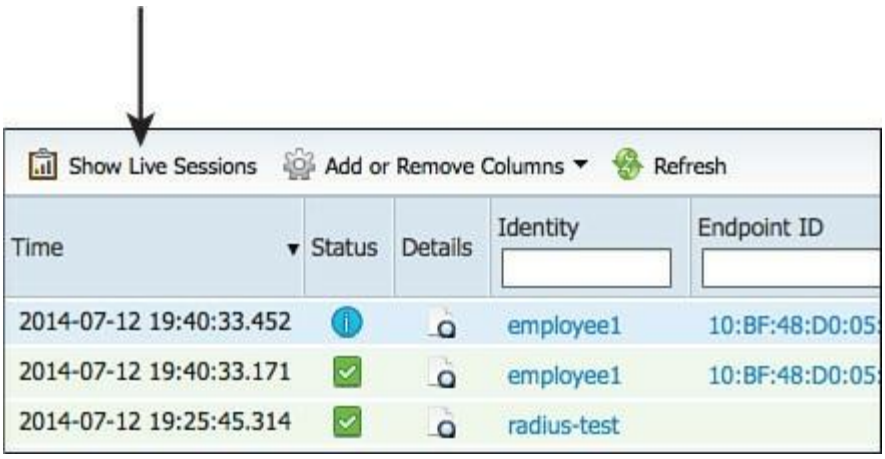


Figure 12-33 Live Authentications Log.

As shown in Figure 12-34, Live Log displays a correlated view of all activity related to an authentication session. If an endpoint changes state, that change will be listed in this view.

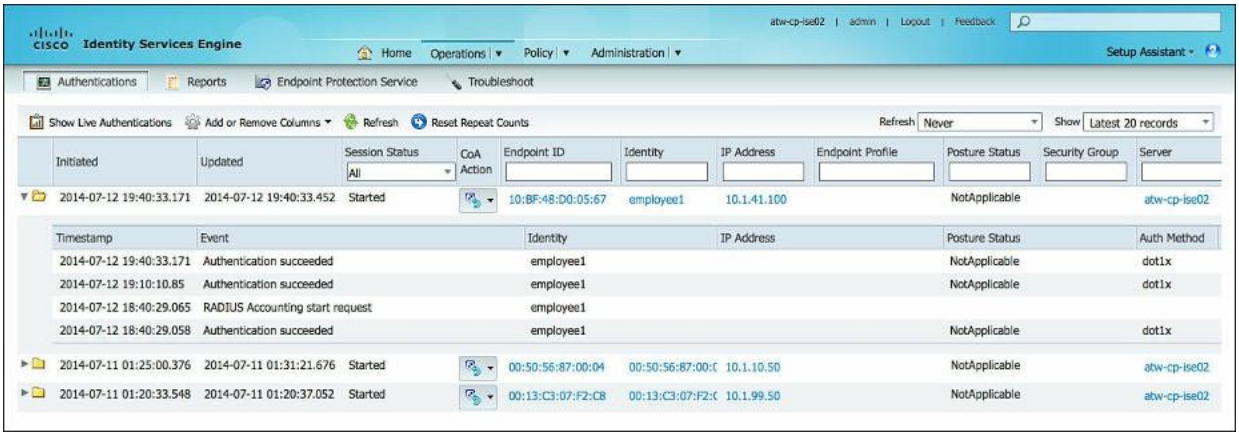


Figure 12-34 Live Sessions Log.

The Live Sessions Log is covered in greater detail throughout this book.

Looking Forward

While this chapter has covered the required configuration(s) for switches and wireless controllers to perform 802.1X and MAB authentications successfully with ISE, these NADs will be responsible for so much more than just basic authentication and authorization. Local ACLs will need to be deployed to define what traffic should be redirected for all use cases, possibly configuring Security-group eXchange Protocol (SXP) and other specific items. Those specific NAD configurations are covered in Chapter 13; Chapter 14, "Deploying Guest Services"; Chapter 17; and Chapter 18, "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 12-2 lists a reference of these key topics and the page numbers on which each is found.



https://t.me/learningnets

Key Topic Element	Description	Page
Step list	Creating global AAA methods	293
Step list	Adding ISE PSNs to the switch configuration with IOS 12.2.X	294
Step list	Adding ISE PSNs to the switch configuration with IOS 15.X	295
Step list	CoA configuration in IOS	296
Step list	Global 802.1X configuration	297
Paragraph	Understanding Flex-Auth	299
List	Host modes	302
Paragraph	Understanding Dot1X timers	305
Step List	Enabling authentication on the switchport	305
List	**show authentication session interface** command	332
Paragraph	Sending syslog from NADs to MNT	333

Table 12-2 Key Topics for Chapter 12

Define Key Terms

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

flexible authentication (Flex-Auth)

network access device (NAD)

change of authorization (CoA)

vendor-specific attributes (VSAs)

MAC authentication bypass (MAB)

single-host

multiple domain authentication (MDA)

multiple authentication (Multi-Auth)

multi-host

Cisco Discovery Protocol (CDP) Second port disconnect

EAPoL-Proxy-Logoff

port security

https://t.me/learningnets

Chapter 13. Web Authentication

This chapter covers the following exam topics:

- **Implement Wired/Wireless 802.1X**
- **AV Pairs**
- **Implement MAB**
- **Implement Network Authorization Enforcement**
- **Troubleshooting, Monitoring, and Reporting Tools**
- **Implement Central Web Authentication (CWA)**
- **Describe the Function of CoA to Support Web Authentication**
- **Configure Authentication Policy to Facilitate CWA**
- **URL Redirect Policy**
- **Redirect ACL**
- **Customize Web Portal**
- **Verify Central Web Authentication Operation**

As discussed in Chapter 5, "Non-802.1X Authentications," just because there is no configured supplicant on an endpoint does not mean that the user of that endpoint does not need to authenticate. Consider the use cases of guests or visitors, or maybe just a misconfiguration or expired credential for the end user. Based on who the user is, they still might require network access and be granted access to the network.

Enter web authentication, commonly referred to as just WebAuth. An authenticator would be able to send a user to a locally hosted web page—in other words, a web page hosted on the local device itself: the switch, wireless controller, or even the firewall or VPN concentrator. It is a simple thing, really, just a very basic web page where a user can submit his username and password.

Chapter 5 also discussed the multiple types of WebAuth and that Centralized WebAuth (CWA) is the type used with Cisco Secure Access and ISE. CWA is the focus of the exam and therefore the main focus of this book.

> **Note**
>
> The content of this chapter is written with the assumption that the switches and wireless LAN controllers (WLCs) have been configured as defined in Chapter 12, "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 12.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

https://t.me/learningnets

Foundation Topics Section	Questions
Implement Wired/Wireless 802.1X	1, 3
AV Pairs	2, 5, 7, 10
Implement MAB	3-4, 10
Implement Network Authorization Enforcement	2-3, 5, 7
Troubleshooting, Monitoring, and Reporting Tools	1, 8-10
Implement Central Web Authentication (CWA)	1-2, 4-7, 10
Describe the Function of CoA to Support Web Authentication	6, 10
Configure Authentication Policy to Facilitate CWA	4
URL Redirect Policy	1-2, 5, 7
Redirect ACL	2, 7
Verify central Web Authentication Operation	2, 8-10

Table 13-1 "Do I Know This Already?" Section-to-Question Mapping

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 will generate a self-signed certificate, which configuration is required?
   - **a.** The internal CA must be enabled.
   - **b.** An IPv6 address.
   - **c.** A Cisco switch cannot generate a self-signed certificate.
   - **d.** A domain name.
2. True or False? The URL redirection ACL can be downloaded from ISE to the NAD.
   - **a.** True
   - **b.** False
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.** RADIUS NAC
   - **d.** Fast Transition
4. For wired and wireless MAB, which option must be configured for unknown identities?
   - **a.** Drop
   - **b.** Continue
   - **c.** Reject

https://t.me/learningnets

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 that redirects the user to the CWA portal.

   c. An authentication rule must be created that permits access to users who have successfully authorized through the CWA portal.

   d. An authorization rule must be created that permits access to users who have successfully authenticated through the CWA portal.

   e. A WebAuth authentication rule must be created that redirects the end user to the CWA portal.

6. Which of the following capabilities exists for MyDevices portals in ISE 1.2 but not the Device-Registration portal?

   a. MyDevices provide a portal for the end user to manage his endpoints.

   b. MyDevices provides the ability to automatically populate the MAC address of the endpoint.

   c. MyDevices did not exist in ISE version 1.2.

   d. MyDevices is linked to the MDM and has the knowledge of which device belongs to a user.

7. True or False? CWA and DRW are using the same RADIUS attributes; the difference is in the actual URL sent down to the NAD.

   a. True

   b. False

8. Which command on the NAD will display information about the URL-redirected session, including the MAC address, IP address, dACL, URL-redirect ACL, and the URL to which the end user is being redirected?

   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 within the ISE GUI should you examine to validate that CWA is working? (Choose the best answer.)

   a. Policy > Policy Elements > Results > Authorization

   b. Operations > Authentications

   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 the MAB and CWA authentications together.

    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 packet of disconnect (PoD) to the NAD. ISE is then able to tie the original MAB session to the new web-authenticated session by correlating the MAC addresses from both authentication sessions.

https://t.me/learningnets

Foundation Topics

Web Authentication Scenarios

There are a number of reasons a company might choose to implement a WebAuth strategy. One of the most common reasons is to provide Internet access to visitors (also caled guests), which is detailed in Chapter 14, "Deploying Guest Services." Additionally, as newer versions of ISE are released, many companies are looking to add an interactive login to capture a user's username and password, as an additional credential to a certificate-based authentication. Think: two-factor authentication.

The end user is presented with a web portal to input his username and password. The credentials are then sent from the authenticator to ISE in a standard RADIUS Access-Request packet. So, in a similar fashion to what occurs with MAC address bypass (MAB), the switch is sending the request for the endpoint, because the endpoint itself is not participating in authentication. Figure 13-1 illustrates the WebAuth concept.

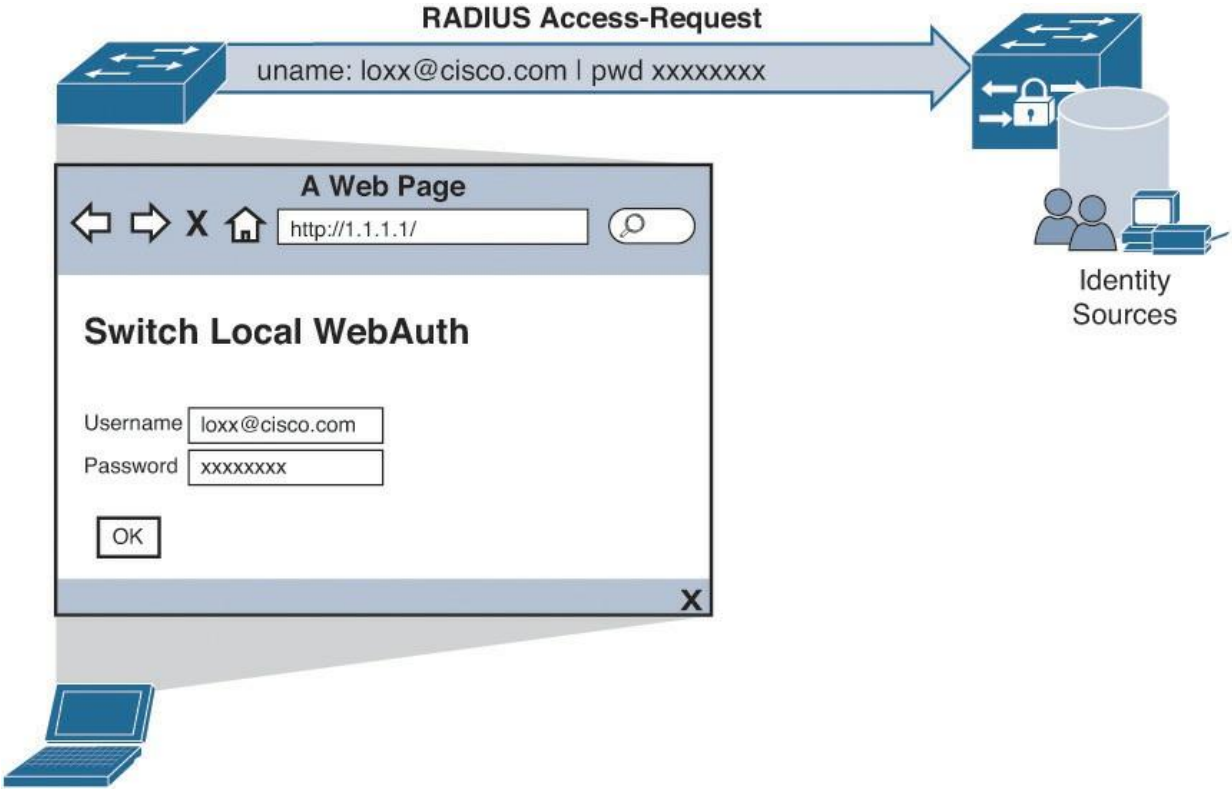


Figure 13-1 Web authentication.

The credential that gets submitted through the WebAuth page could be 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 identity type.

Keep in mind that WebAuth is an effective authentication method only for a device that has an interactive user. In other words, it would not make sense to try to use WebAuth for a printer. There is no user to interact with the web portal, enter her credentials, and then click Submit.

As with MAB, WebAuth is not a standard either. There are multiple ways to perform WebAuth, with benefits and downsides to each one.

Local Web Authentication

Local web authentication (LWA) is the original WebAuth. The authenticator will redirect web browser traffic to a locally hosted web portal where a user can enter her username and password.

https://t.me/learningnets



The credentials are submitted through the switch, or the wireless controller 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 switch, the locally hosted web pages are minimally customizable. Many companies 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 WebAuth with wired connections.

Additionally, when using LWA with Cisco switches, no native support exists for advanced services, including:

- Acceptable Use Policy acceptance pages
- Client provisioning
- Password changing capabilities
- Self-registration
- Device registration

For those advanced capabilities, a company truly needs to consider using centralized web authentication (CWA). For more details on LWA, see Chapter 5.

Centralized Web Authentication

CWA is what Cisco ISE uses throughout the Secure Access solution. Although Cisco ISE is still capable of supporting LWA methods, those methods are typically reserved for non-Cisco network devices.

As with all web authentications, CWA is only for interactive users who have a web browser, where the user will manually enter the username and password.

Change of authorization (CoA) works fully with CWA, which leads to the support for all the authorization results, such as ACL and VLAN authorization. Keep in mind that any time you change virtual local area networks (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 normally is no supplicant on the endpoint. Therefore, the DHCP scope length on the initial subnet 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.

CWA also supports all the advanced services, such as these:

- Client provisioning
- Posture assessments
- Acceptable use policies (AUPs)
- Password changing
- Self-registration
- Device registration



As described in Chapter 5, the switch or wireless controller sees only a MAB, and the rest is handled on the authentication server (ISE). Figure 13-2 shows the MAB occurring with a redirection to the centralized portal, and Figure 13-3 illustrates that the switch still sees only a MAB request while ISE is maintaining the user authentication.

https://t.me/learningnets

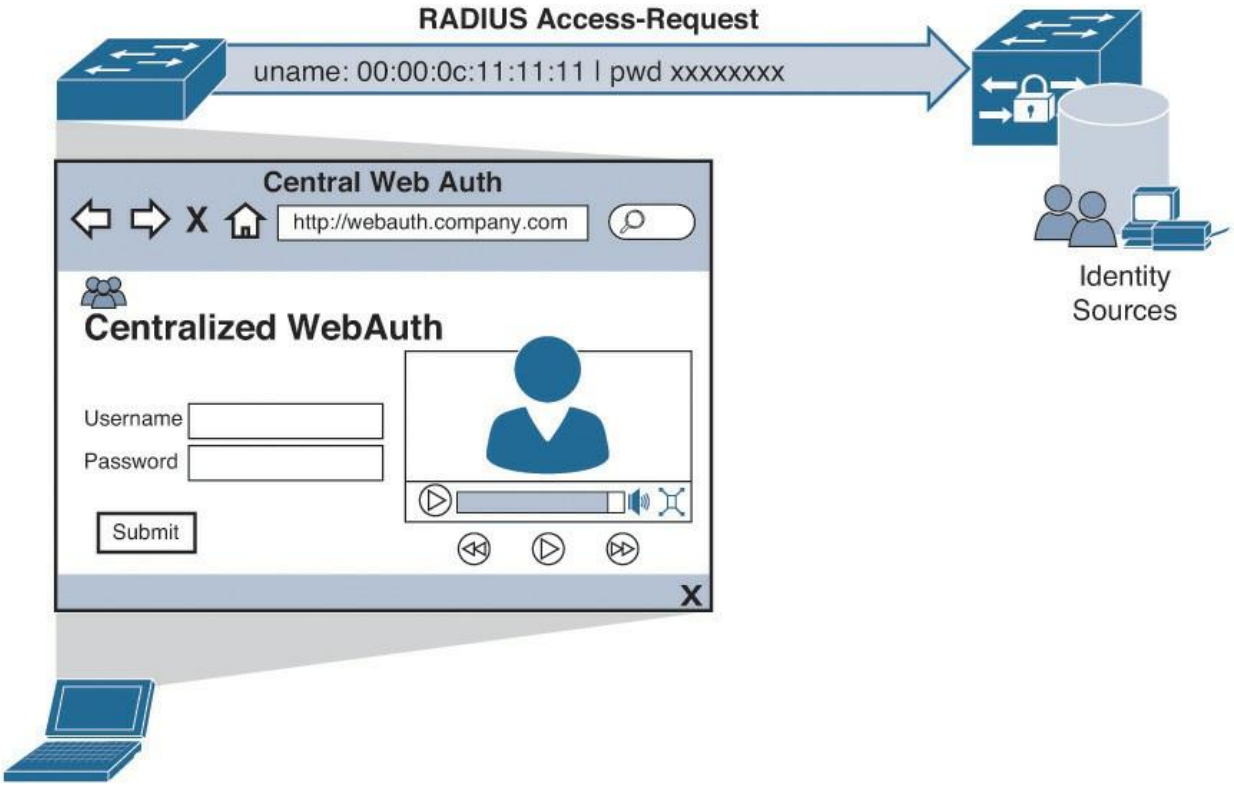


Figure 13-2 URL-redirected MAB.

https://t.me/learningnets

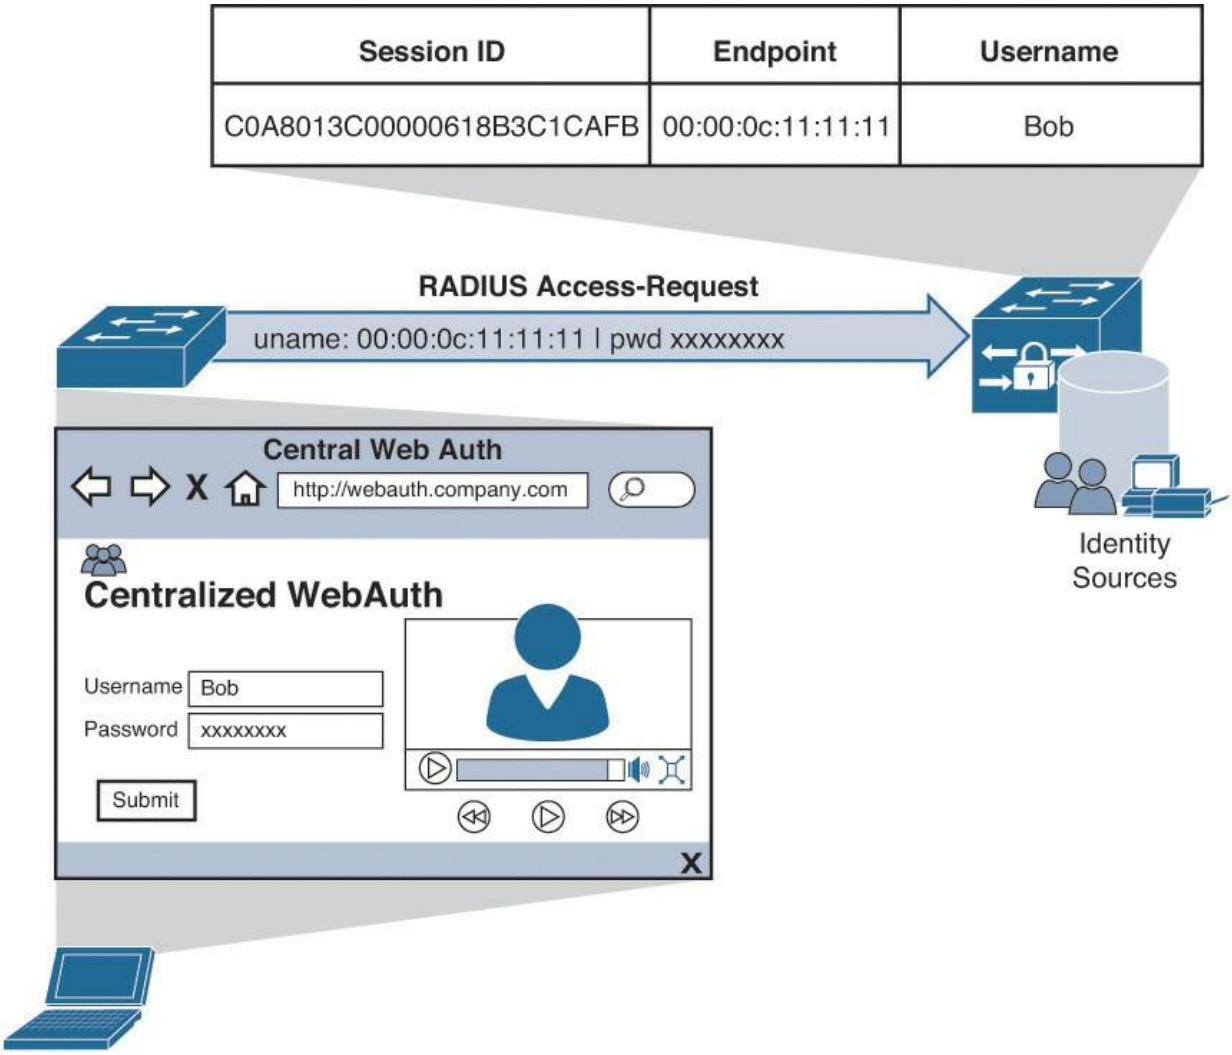


Figure 13-3 The credentials are never sent to the authenticator.

The following steps detail what is occurring in Figures 13-2 and 13-3:

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

Step 2. The authenticator performs a 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 itself.

Step 4. The end user enters his credentials into the centralized portal. Unlike the LWA options, the credentials are never sent to the switch; 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, which 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 is the basis for many of the other solutions within ISE, including device registration WebAuth, BYOD onboarding, MDM onboarding, posture assessments, and much more.

https://t.me/learningnets

Device Registration WebAuth

Device registration WebAuth (DRW) is a special form of WebAuth. Beginning with ISE 1.0 and continuing until ISE 1.2.x (which this exam is focused on), this special WebAuth was used to allow an interactive user to add her MAC address to an Endpoint Identity Group. An authorization rule that permits members of that Endpoint Identity Group to have network access (hopefully limited network access because this is not a strong method of authentication) is required.



It is important as both an implementer of ISE and a taker of this exam to understand the difference between DRW and the BYOD portal known as MyDevices. Both portals allow an interactive user to add her MAC address to the system. However, a device that is authorized via the DRW portal uses a base license, whereas a MyDevices endpoint uses an advanced license.

License consumption is not the only difference between the two processes. Table 13-2 displays a comparison of the DRW and MyDevices functions.

Feature or Function	DRW	MyDevices
Advanced license consumption	No	Yes
Provides a portal for end user to manage his registered Devices	No	Yes
End user can be automatically redirected to the registration portal	Yes	Yes
MAC address must be manually entered?	No	No
BYOD-Registration attribute is set on the endpoint object	No	Yes
Endpoint can be added to the registered device's Endpoint ID Group	Yes	Yes
Endpoint ID Group is configurable per portal	Yes	No
Ability to provision a configuration for the endpoint's native supplicant	No	Yes
Ability to provision a certificate to the endpoint	No	Yes

Table 13-2 Comparison of DRW and MyDevices

The MyDevices portal and BYOD processes are covered in much more detail in Chapter 17, "Bring Your Own Device."

Configuring Centralized Web Authentication

Multiple devices will require configuration to enable CWA. The network access device (NAD) requires some special configuration, such as a redirection ACL, and ISE will need authentication and authorization rules set up for CWA as well.

Let's start with the NADs.

Cisco Switch Configuration

Within Cisco's Secure Access system, the switch performs the URL redirection for web authentication as well as redirecting the discovery traffic from the posture agent to the ISE policy service node (PSN).

Performing URL redirection at the Layer-2 access (edge) device is a vast improvement over previous NAC

https://t.me/learningnets

solutions that required an appliance (such as the inline posture node) to capture web traffic and perform redirection to a web authentication page, simplifying the deployment for both web authentication and the posture agent discovery process.

Configuring Certificates on the Switch

To redirect HTTPS traffic, the switch must already have its own certificate.

Cisco IOS does not enable certificates, or even self-generated keys, to be created and installed without first defining a DNS domain name on the device.

From Global Configuration mode on the switch, do the following:

Step 1. 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, the keys can be generated.

Step 2. Generate keys to be used for HTTPS; type **crypto key generate rsa general-keys mod 2048** at the Global Configuration prompt.

Enabling the Switch HTTP/HTTPS Server

The embedded HTTP/S server in IOS will be used to grab HTTP traffic from the user and redirect that user's browser to the CWA portal, a device registration portal, or even the Mobile Device Management onboarding portal. This same function is used for redirecting the posture agent's traffic to the PSN:

Step 1. Enable the HTTP server in Global Configuration mode:

Click here to view code image

```
C3560X(config)# ip http server
```

Step 2. Enable the HTTP secure server:

Click here to view code image

```
C3560X(config)# ip http secure-server
```

Many organizations will want to ensure this redirection process that is using the switch's internal HTTP server is decoupled from the management of the switch itself. This can be accomplished by running the commands listed in Step 3, from Global Configuration mode:

Step 3. Disconnect the HTTP management process from the URL-Redirection process:

Click here to view code image

```
C3560X(config)# ip http active-session-modules none
C3560X(config)# ip http secure-active-session-modules none
```

Verifying the URL-Redirection ACL

In Chapter 12, you created an Access List named **ACL-WEBAUTH-REDIRECT**. This list is used to determine what traffic is redirected to the CWA portal with the **permit** statement. Any traffic that is denied will not be redirected.



Contrary to the way a WLC works, the URL-redirection ACL on a switch is used only to determine which traffic is redirected and which is not. If network traffic is denied from redirection, that does not deny it from traversing the network. The traffic-filtering capability will come from the downloadable ACL (dACL) that is sent down 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 in the WLC configuration within this chapter:

https://t.me/learningnets

Step 1. Validate whether the **ACL-WEBAUTH-REDIRECT** ACL is configured on the NAD:

Click here to view code image

```
C3560X# 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, follow Step 2.

Step 2. Add the following ACL to be used for URL redirection with web authentication:

Click here to view code image

```
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
```

Cisco WLC Configuration

Just as the Cisco switches are responsible for redirecting the web browser traffic to the centralized portal(s), the Cisco WLC must do the same thing. As stated in the introduction to this chapter, you are expected to have already configured the WLC as per Chapter 12.

In Chapter 12, you created an open WLAN with MAC filtering enabled and the NAC state configured for Radius NAC. Additionally, you created an Access List named **ACL-WEBAUTH-REDIRECT**.

> **Note**
>
> At the time of writing this book, the latest version of the Cisco WLC was 7.6 MR3. As of that version, the Cisco WLC could redirect only HTTP traffic. HTTPS redirections were not possible but were expected in a future release.

Validating That MAC Filtering Is Enabled on the WLAN

The MAC Filtering option for the open wireless network is configuring the WLAN for wireless MAB. This is necessary to ensure an authentication is sent from the WLC to the ISE, so ISE can return the URL-Redirection in the authorization result.

From the WLC GUI, do the following:

Step 1. Navigate to **WLANs**.

Step 2. Examine the list of WLANs, ensuring that MAC Filtering is listed under the Security Policies column, as shown for the CiscoPress-Guest SSID in Figure 13-4.

https://t.me/learningnets



Figure 13-4 MAC Filtering on an open SSID.

Validating That Radius NAC Is Enabled on the WLAN

Even though the Radius NAC feature is mislabeled with the wrong capitalization of the acronym RADIUS, it is still a valid and important setting. This setting is critical to allow for URL redirection, CWA, posture assessment, native supplicant provisioning, and more.

From the WLC GUI, do the following:

Step 1. Navigate to **WLANs >** and select your open SSID.

Step 2. Click the **Advanced** tab.

Step 3. Examine the setting for the NAC state and ensure it is configured for **Radius NAC**, as shown in Figure 13-5. While at this configuration screen, also ensure that **Allow AAA Override** is enabled, which is also depicted in Figure 13-5.

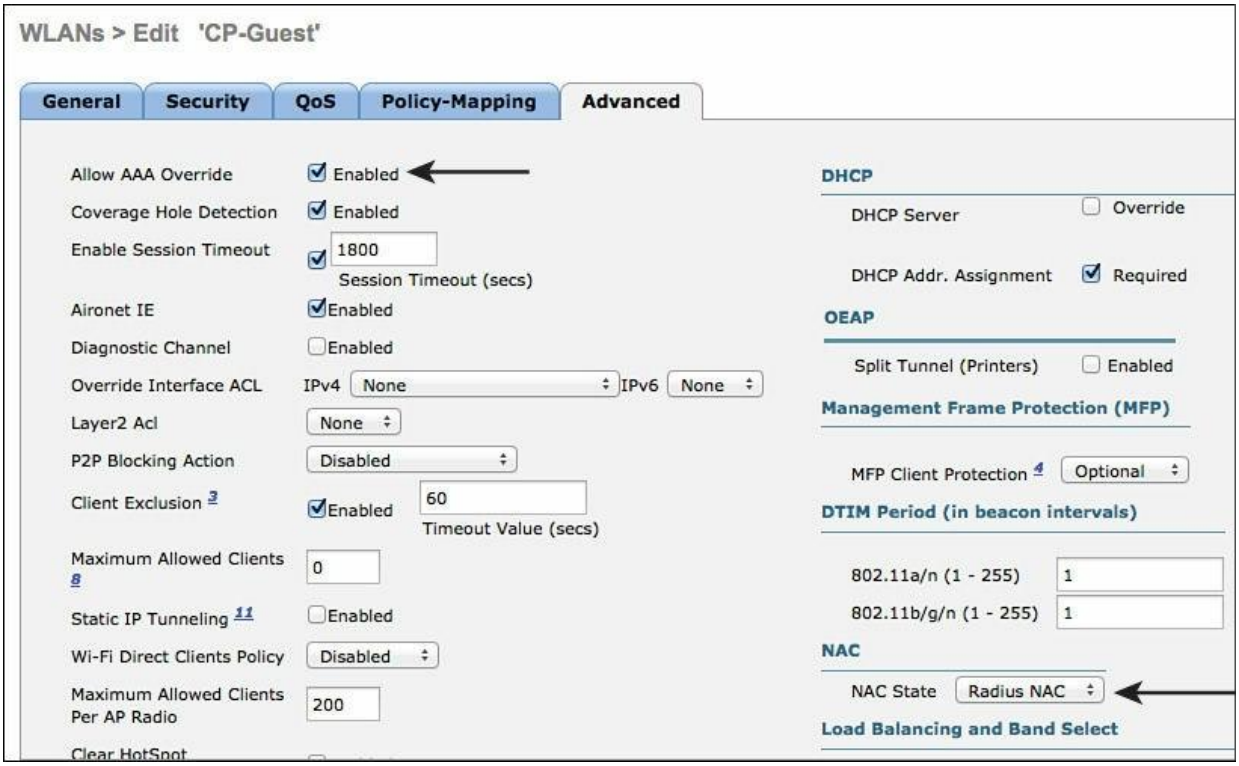


Figure 13-5 Radius NAC setting.

https://t.me/learningnets

Validate That the URL-Redirection ACL Is Configured

The last critical item to ensure exists in the WLC configuration is an ACL to use for UR -redirections. In Chapter 12, you created an ACL named **ACL-WEBAUTH-REDIRECT**. This list is used to determine which traffic is redirected to the CWA portal with the **deny** statement. Any traffic that is permitted will not be redirected.



Unlike IOS-based NADs, the 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 the single ACL. Therefore, the logistics of which traffic is redirected are not the same as with IOS-based devices. With these WLCs, a **deny** statement means that traffic should be redirected. A **permit** statement allows the traffic through the WLC and bypasses the redirection.

From the WLC GUI, do the following:

Step 1. Navigate to **Security > Access-Control-Lists > Access-Control Lists**. Ensure the **ACL-WEBAUTH-REDIRECT** ACL is in the list, as shown in Figure 13-6.

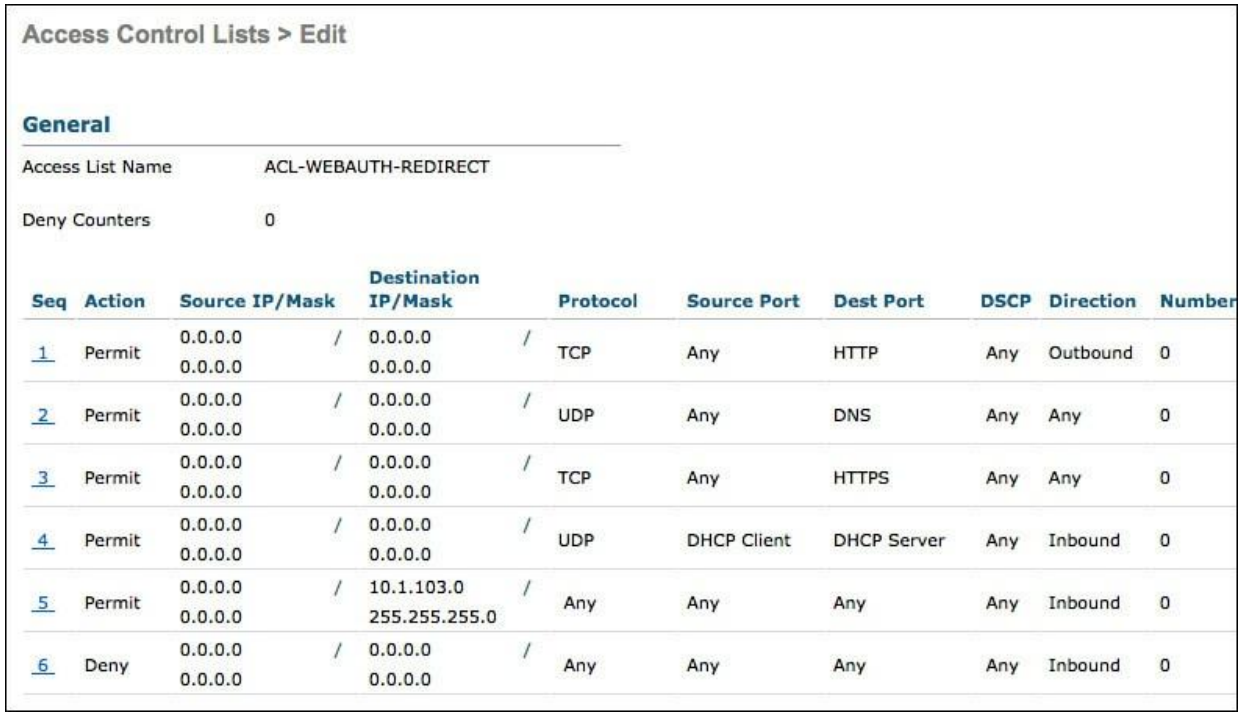


Figure 13-6 ACLs.

Step 2. Click the ACL, and ensure the entries for your environment are there.

Figure 13-7 shows the ACL as it is defined in the WLC.

https://t.me/learningnets



Access Control Lists > Edit

General

Access List Name ACL-WEBAUTH-REDIRECT

Deny Counters 0

Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port	DSCP	Direction	Number
1	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	TCP	Any	HTTP	Any	Outbound	0
2	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	UDP	Any	DNS	Any	Any	0
3	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	TCP	Any	HTTPS	Any	Any	0
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	0
5	Permit	0.0.0.0 / 0.0.0.0	10.1.103.0 / 255.255.255.0	Any	Any	Any	Any	Inbound	0
6	Deny	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Inbound	0

Figure 13-7 ACL-WEBAUTH-REDIRECT contents.

Captive Portal Bypass

One more local configuration is required for the WLC. You need to enable a feature called captive portal bypass. Apple iOS devices have a feature to help facilitate network access when their end users are on public hotspot-type networks. When joining an open wireless network, the iOS device attempts to reach a few web pages hosted on apple.com, and it attempts this in the background so the end user is not aware of the attempted communication.

If the communication is successful, nothing happens because (obviously) the end user must be on the Internet if that was successful. If the connection is not successful, iOS launches a pseudo-browser to facilitate the end user in entering credentials or accepting the public hotspot's acceptable use policy (or whatever other reason the redirection can be happening).

The pseudo-browser is not fully functioning and has problems with many types of redirections, including the ones to ISE's secure portals. Therefore, the pseudo-browser doesn't work for most CWA types or other onboarding. The icing on the cake is that if the end user closes that full-screen pseudo-browser, iOS will drop the connection to the network. That's right—instead of allowing the end user to launch a real browser, it actually disconnects him from the network. This causes a denial of service (DoS) against iOS-based endpoints with CWA or BYOD onboarding, and so on.

To get around that DoS, Cisco and other vendors have all created different ways to "spoof" a success message to the iOS endpoint. The following command needs to be run on all your Cisco WLCs, which will also require a reboot before it is active:

Click here to view code image

```
> config network web-auth captive-bypass enable
```

Configuring ISE for Centralized Web Authentication

Now that you're sure the key elements of the network devices are configured correctly, it's time to ensure ISE is configured correctly, too. A key change must be made in the authentication policy: an identity source sequence that will use all the appropriate identity stores and the configuration of the appropriate traffic filtering dACLS. Additionally, you must create the appropriate authorization rules for both pre- and post-web authentication.

https://t.me/learningnets

Configuring MAB for the Authentication

WebAuth is often used for guest access, which means the endpoint will most likely be unknown to ISE when the guest is attached to the network. As such, it is critical to set the identity options to continue when the MAC address is unknown.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authentication**.

Step 2. Edit the MAB rule, and click the plus (+) sign next to **Internal Endpoints**.

Step 3. Ensure the **If User Not Found** setting is set to **Continue**, as shown in Figure 13-8.

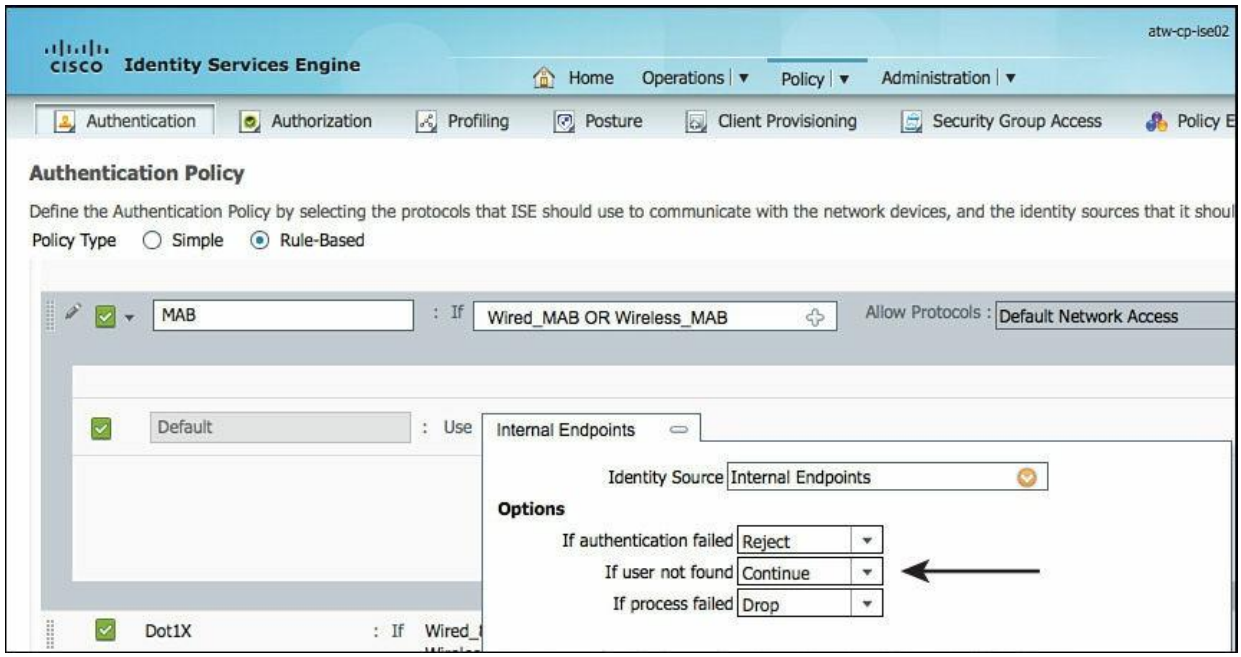


Figure 13-8 MAB, continued.

Step 4. Click **Done** and then click **Save**.

Configuring the Web Authentication Identity Source Sequence

There is a pre-onfigured 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 both the Internal Users and the Guest Users identity stores. Web authentication is not only used for guest access, so in this case we add Active Directory to the list of identity sources to query.

From the ISE GUI, follow these steps:

Step 1. Navigate to **Administration** > **Identity Management** > **Identity Source Sequences**, as shown in Figure 13-9.

https://t.me/learningnets

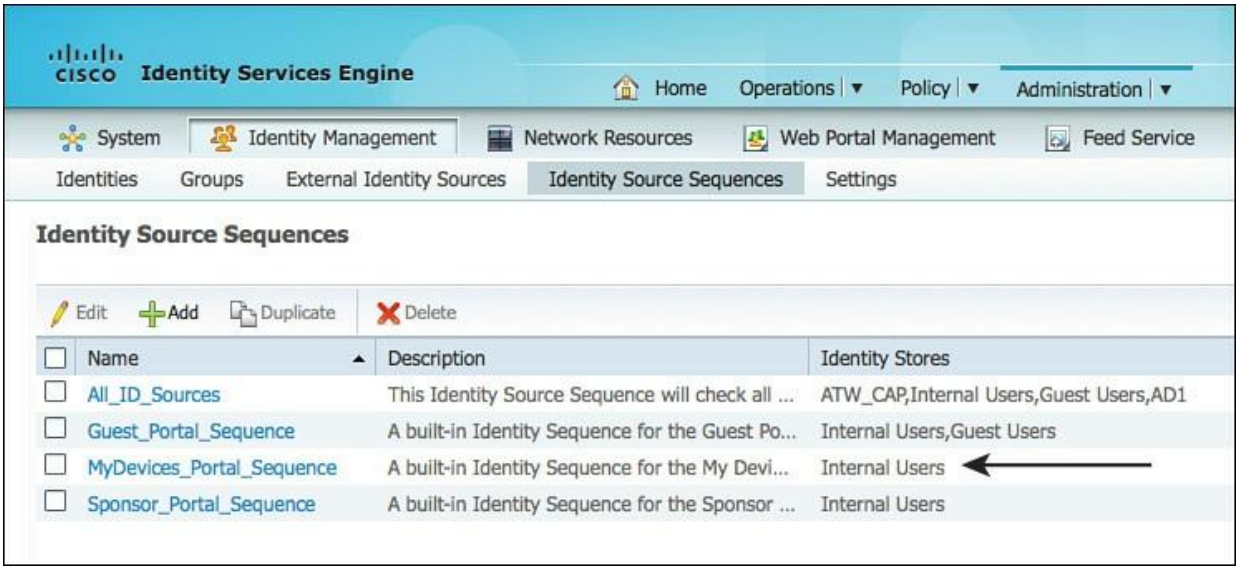


Figure 13-9 Identity source sequences.

Step 2. Click the ISS named **Guest_Portal_Sequence**.

Step 3. Add the Active Directory identity store to the sequence, as shown in Figure 13-10.

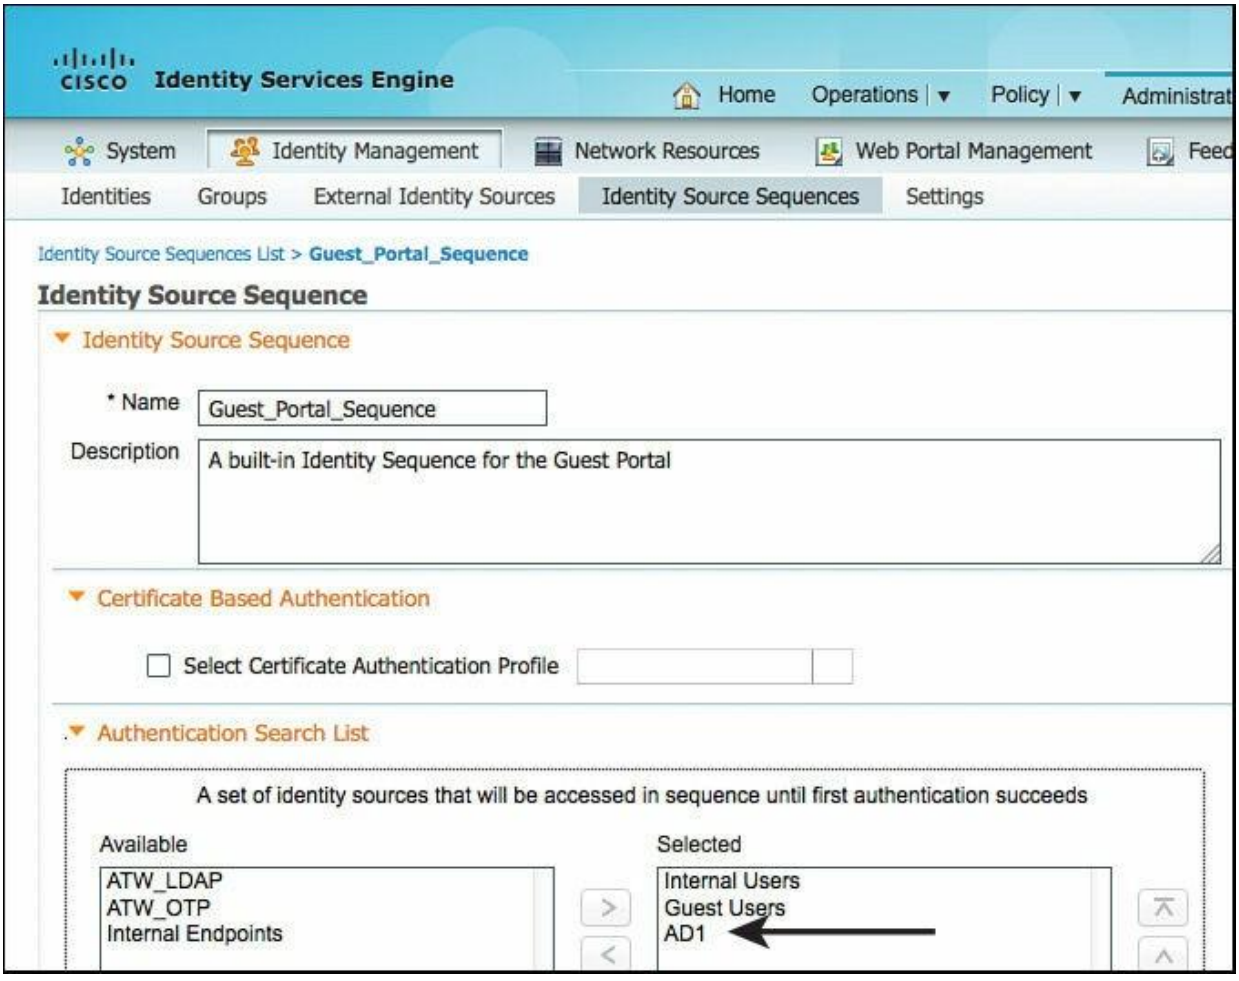


Figure 13-10 Guest_Portal_Sequence.

Step 4. Click **Save**.

https://t.me/learningnets

Configuring 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.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Downloadable ACLs**, as shown in Figure 13-11.

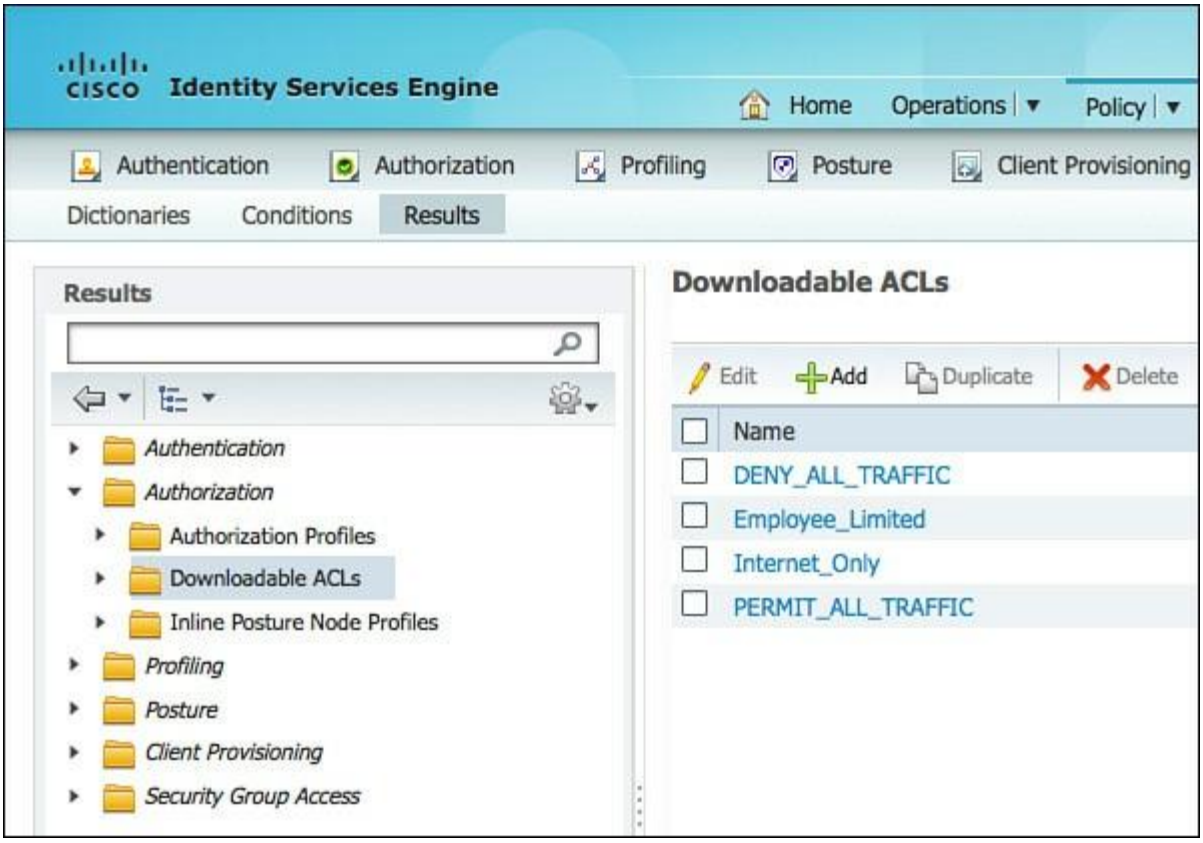


Figure 13-11 Downloadable ACLs.

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 PSNs, but deny access to the remainder of the internal network, similar to what is shown in Figure 13-12.

https://t.me/learningnets

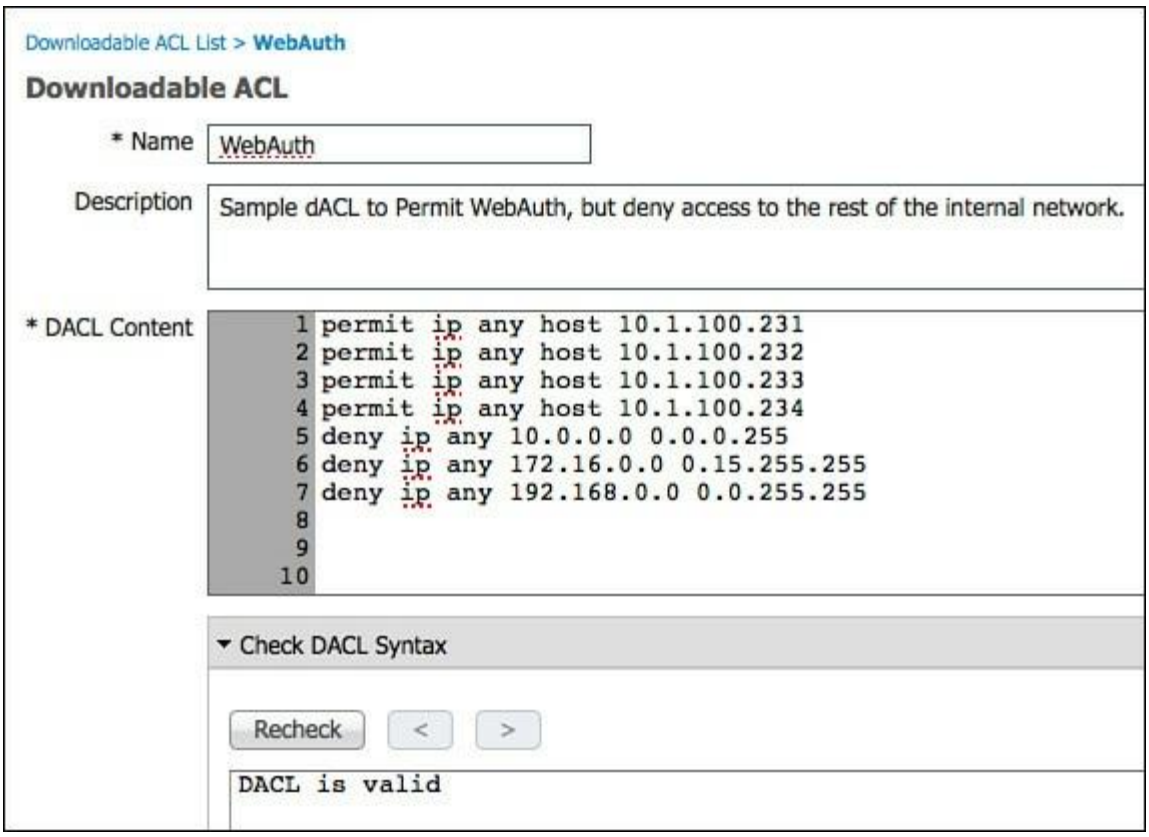


Figure 13-12 Sample WebAuth dACL.

Step 6. Click **Save**.

Configuring an Authorization Profile

At this point, everything is 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-redirection ACL, and apply the dACL to limit network traffic.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

Step 2. Click **Add**.

Step 3. Name the new authorization profile **CWA**.

Step 4. Select the **WebAuth** dACL.

Step 5. Select the check box for **Web Redirection**, choosing **Centralized Web Auth** from the first drop-down list. In the ACL text box, type **ACL-WEBAUTH-REDIRECT**. This ACL name must be typed exactly with proper capitalization. You are using the default WebAuth portal, so ensure that **Default** is selected for the redirect drop-down menu.

Figure 13-13 is a composite image, created to show all the key parts in one graphic, and it shows the complete authorization profile.

https://t.me/learningnets

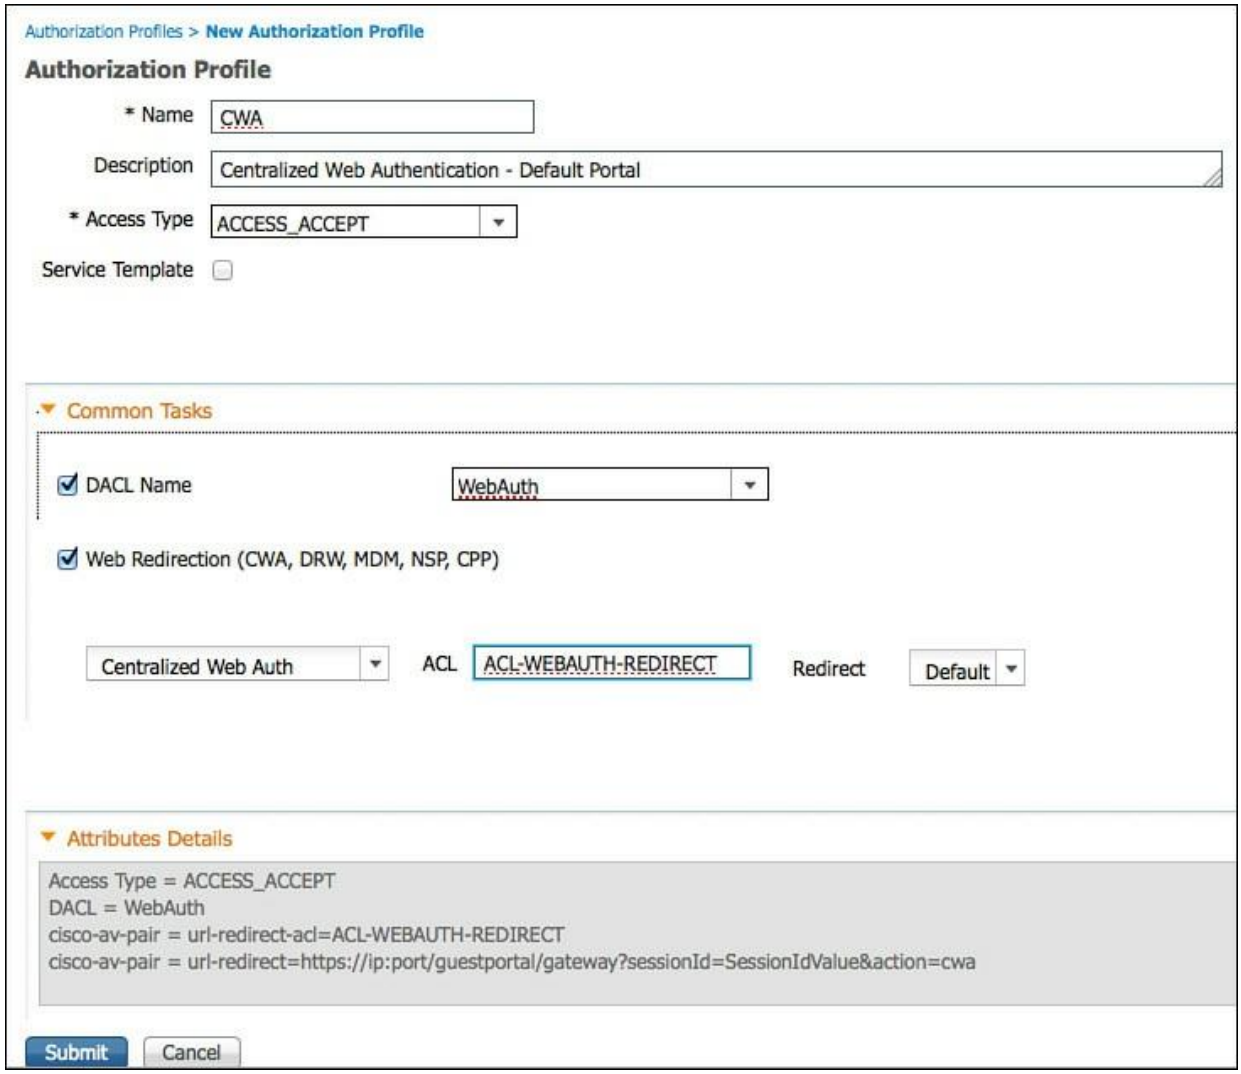


Figure 13-13 The complete CWA authorization profile.

Building CWA Authorization Policies

Two different authorization rule types must be created. The first rule should match if no more specific authorization rule is used and should redirect the user to the CWA portal. The second rule type should exist above the redirection rule and allow access to the user after she has successfully authenticated to the CWA portal.

Creating the Rule to Redirect to CWA

The first rule to create is one that will redirect to the CWA portal when no more specific rules are matched.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy** > **Authorization**.

Step 2. Add a new rule above the default rule at the bottom, as shown in Figure 13-14.

https://t.me/learningnets

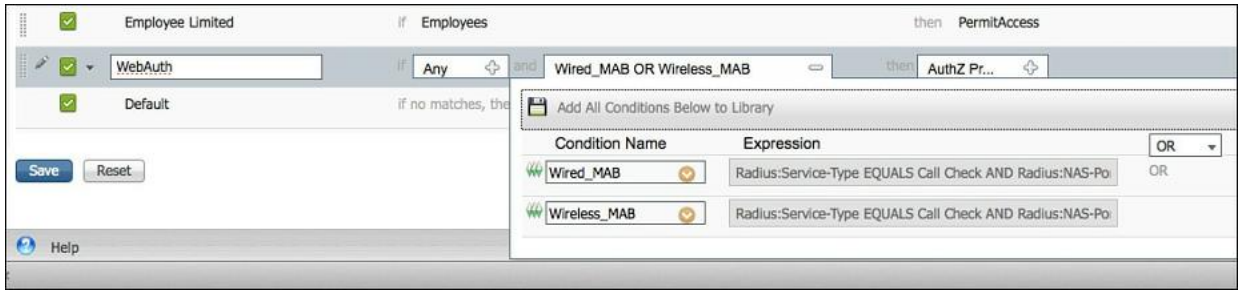


Figure 13-14 Creating the WebAuth authorization rule.

Step 3. Name the new rule **WebAuth**.

Step 4. For the conditions, select two existing compound conditions from the library: **Wired_MAB** and **Wireless MAB**. Ensure the **OR** operator is used with the conditions, as shown in Figure 13-14.

Figure 13-14 shows the compound conditions being added to the authorization rule.

Step 5. Use the **CWA** authorization profile you created previously for the result, as shown in Figure 13-15.

Rule Name	Conditions (identity groups and other conditions)	Permissions
Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
Employee and CorpMachine	if EmployeeFullEAPChain	then Employee Full Access
Employee SmartDevice	if (SmartDevice AND Employees)	then Internet_Only
Employee Limited	if Employees	then PermitAccess
WebAuth	if (Wired_MAB OR Wireless_MAB)	then CWA
Default	if no matches, then PermitAccess	

Figure 13-15 Completed WebAuth authorization rule.

Step 6. Click **Done** and then click **Save**.

Figure 13-15 shows the completed WebAuth authorization rule.

Creating the Rules to Authorize Users Who Authenticate via CWA

The second rule type is one that will allow the user(s) who authenticates via WebAuth to have her specific access to the network. The number of rules created will greatly depend on the needs of your organization. For the purposes of this chapter, you will create two authorization rules: one for employees and another for guests.

Creating the Guest Rule

You are constructing a new authorization rule that will allow guest users (users who belong to the User Identity Group named Guest) who have successfully authenticated through the web portal to have Internet-only network access. You will also use a dictionary item, called Guest Flow, in your rule. This dictionary object is used by ISE to identify when an authentication has occurred via an ISE web portal. Chapter 14 discusses guest access in much more detail.

From the ISE GUI, follow these steps:

https://t.me/learningnets

Step 1. Navigate to **Policy > Authorization**.

Step 2. Add a new rule above the first employee rule, as shown in Figure 13-16.



Figure 13-16 Creating the guest authorization rule.

Step 3. Name the new rule **Guest**.

Step 4. Select the **Guest** Internal User Identity Group, as shown in Figure 13-16.

Step 5. Create a new condition for **Network Access > Use Case EQUALS Guest Flow**, as shown in Figure 13-16.

Step 6. Save the condition as **GuestFlow** for reuse in other rules.

Figure 13-16 shows the guest authorization rule being constructed.

Step 7. Select the **Internet-Only** authorization profile you created earlier as the result.

Step 8. Click **Done** and then click **Save**.

Figure 13-17 shows the completed guest rule.

Rule Name	Conditions (identity groups and other conditions)	Permissions
Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
Guest	if **Guest** AND GuestFlow	then Internet_Only
Employee and CorpMachine	if EmployeeFullEAPChain	then Employee Full Access
Employee SmartDevice	if (SmartDevice AND Employees)	then Internet_Only
Employee Limited	if Employees	then PermitAccess
WebAuth	if (Wired_MAB OR Wireless_MAB)	then CWA
Default	if no matches, then PermitAccess	

Figure 13-17 The completed guest rule.

Creating the Employee Rule

Technically, you already have a rule that will work. 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 are constructing a new authorization rule that will allow 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.

From the ISE GUI, do the following:

https://t.me/learningnets

Step 1. Navigate to **Policy > Authorization**.

Step 2. Add a new rule below the guest rule, as shown in Figure 13-18.

Rule Name	Conditions (identity groups and other conditions)	Permissions
Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
Guest	if **Guest** AND GuestFlow	then Internet_Only
Employee CWA	if (Employees AND GuestFlow)	then Internet_Only
Employee and CorpMachine	if EmployeeFullEAPChain	then Employee Full Access
Employee SmartDevice	if (SmartDevice AND Employees)	then Internet_Only
Employee Limited	if Employees	then PermitAccess
WebAuth	if (Wired_MAB OR Wireless_MAB)	then CWA
Default	if no matches, then PermitAccess	

Figure 13-18 The completed employee CWA rule.

Step 3. Name the new rule **Employee CWA**.

Step 4. Create a new condition that uses the new **Guest Flow** condition from the library and the saved simple condition named **Employees** with the **AND** operator, as shown in Figure 13-18.

Step 5. Select the **Internet-Only** authorization profile you created earlier as the result.

Step 6. Click **Done** and then click **Save**.

Figure 13-18 shows the completed employee CWA rule.

Configuring Device Registration Web Authentication

In this section you will create a new DRW portal. DRW is used to allow an interactive user to add his MAC address to an Endpoint Identity Group, so that endpoint will be permitted back on the network repeatedly without having to go through another interactive authentication.

Creating the Endpoint Identity Group

The first procedure is to create an Identity Group for all the devices your users register. There is a preexisting endpoint identity group called Registered-Devices, but the best practice is to reserve that group for the BYOD devices that have been registered through MyDevices. You will create a new endpoint identity group named DRW-Endpoints to help differentiate them.

From the ISE GUI, follow these steps:

Step 1. Navigate to **Administration > Identity Management > Groups > Endpoint Identity Groups**.

Step 2. Click **Add**.

Step 3. Name the new group **DRW-Endpoints**.

Step 4. Click **Submit**.

Figure 13-19 shows the completed Endpoint Identity Group.

https://t.me/learningnets



Figure 13-19 The completed DRW-Endpoints Group.

Creating the DRW Portal

The next procedure is for you to create the actual DRW portal to which to redirect the end user. Because you used the default portal for the CWA procedures, this will be a new experience for you. The portal configuration is pretty well buried in a location that has proven itself difficult to find.

From the ISE GUI, follow these steps:

Step 1. Navigate to **Administration > Web Portal Management > Settings**.

Step 2. On the left side, select **Guest > Multi-Portal Configurations**.

Step 3. Click **Add**.

Step 4. Name the new portal **DRW**.

Step 5. Select **Device Web Authorization Portal**.

Step 6. Select **DRW-Endpoints** for the Endpoint Identity Group.

Step 7. Click **Submit**.

Figure 13-20 shows the completed DRW multi-portal.

https://t.me/learningnets

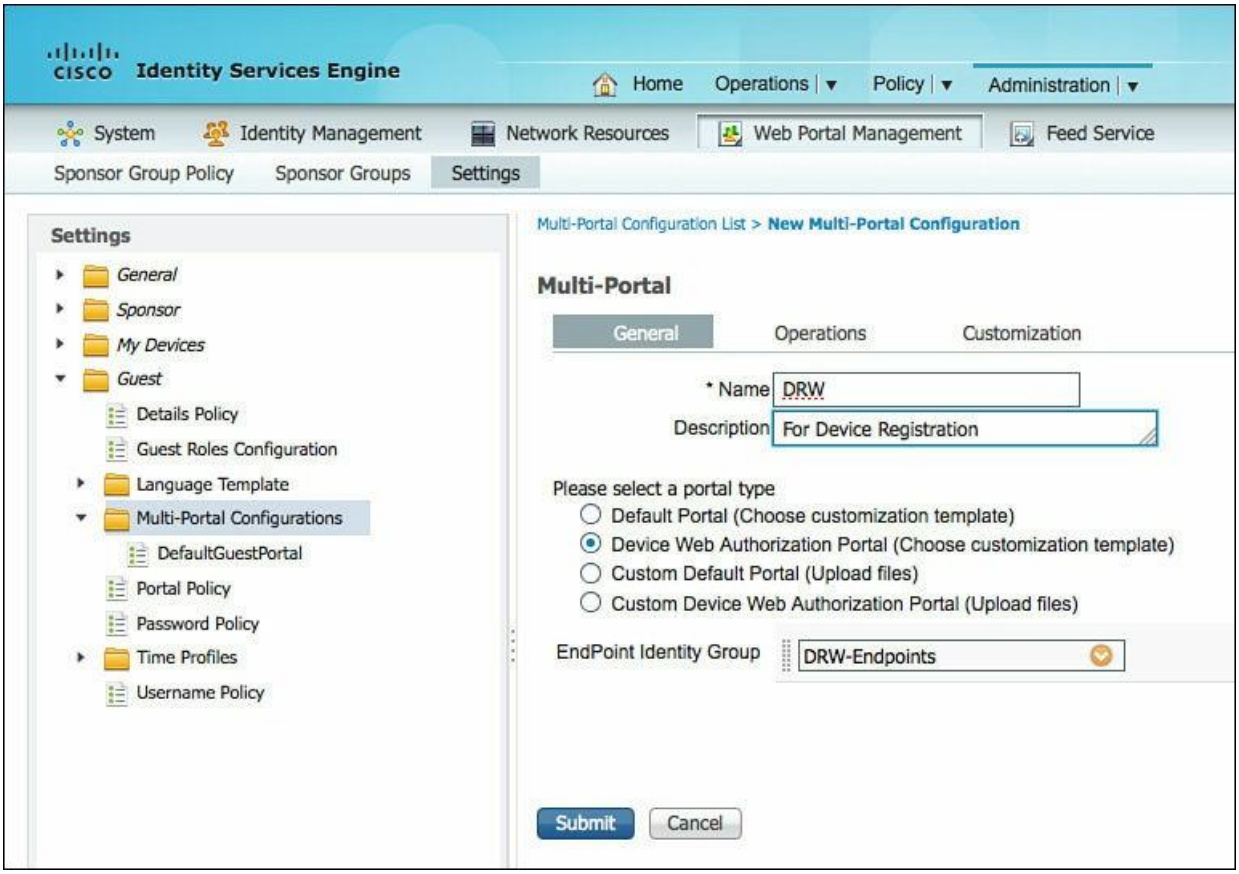


Figure 13-20 The completed DRW multi-portal.

Creating the Authorization Profile

Now that you have a DRW portal created, you need to create an authorization result that will redirect users to that portal.

From the ISE GUI, follow these steps:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

Step 2. Click **Add**.

Step 3. Name the new authorization profile **DRW**.

Step 4. Select the **WebAuth** dACL.

Step 5. Select the check box for **Web Redirection**, choosing **Device Registration Web Auth** from the first drop-down list. In the ACL text box, type **ACL-WEBAUTH-REDIRECT**.

Step 6. Select **DRW** from the Value drop-down list.

Figure 13-21 is a composite image, created to show all the key parts in one graphic, and it shows the complete authorization profile.

https://t.me/learningnets

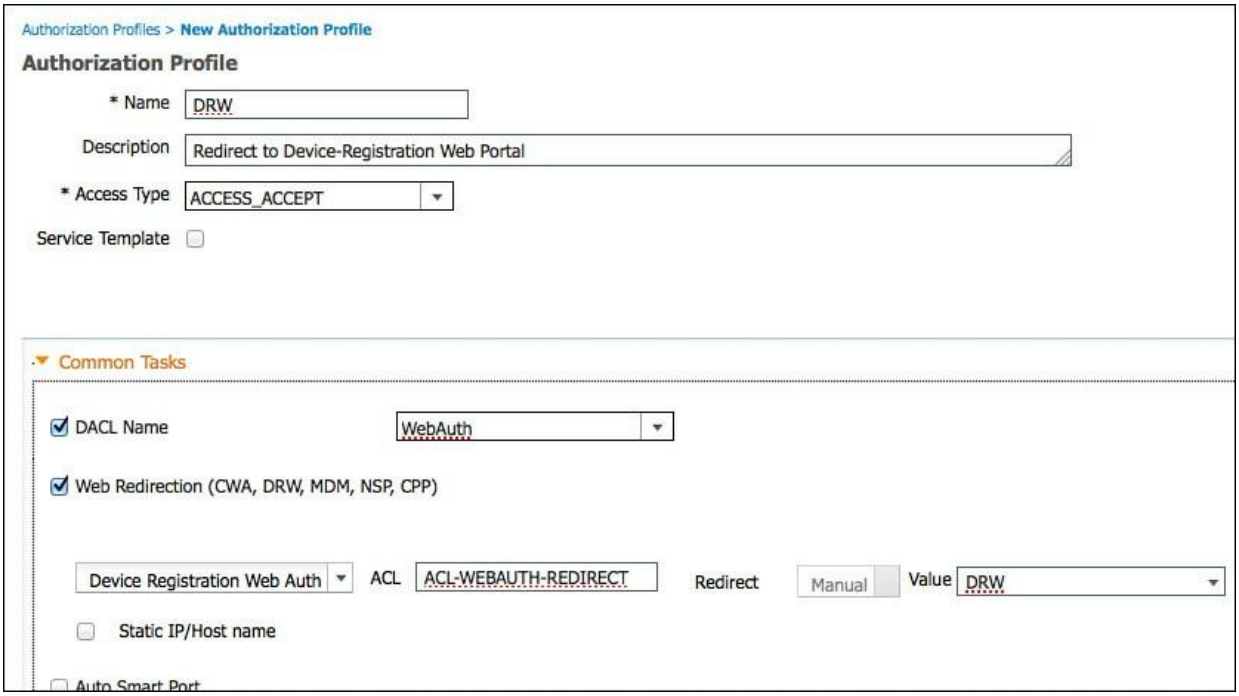


Figure 13-21 The complete DRW authorization profile.

This should have felt a lot like the procedure for creating a CWA authorization profile. Pay attention to how the Web Redirection field in the authorization profile is used for many different actions. When you select Device Registration Web Auth from the first drop-down list, it automatically grayed out the redirect drop-down menu and populated the Value drop-down list with all the DRW-type portals that had been created.

Figure 13-22 illustrates the difference in the raw RADIUS being used for both CWA and DRW. Notice the differences in the URL.

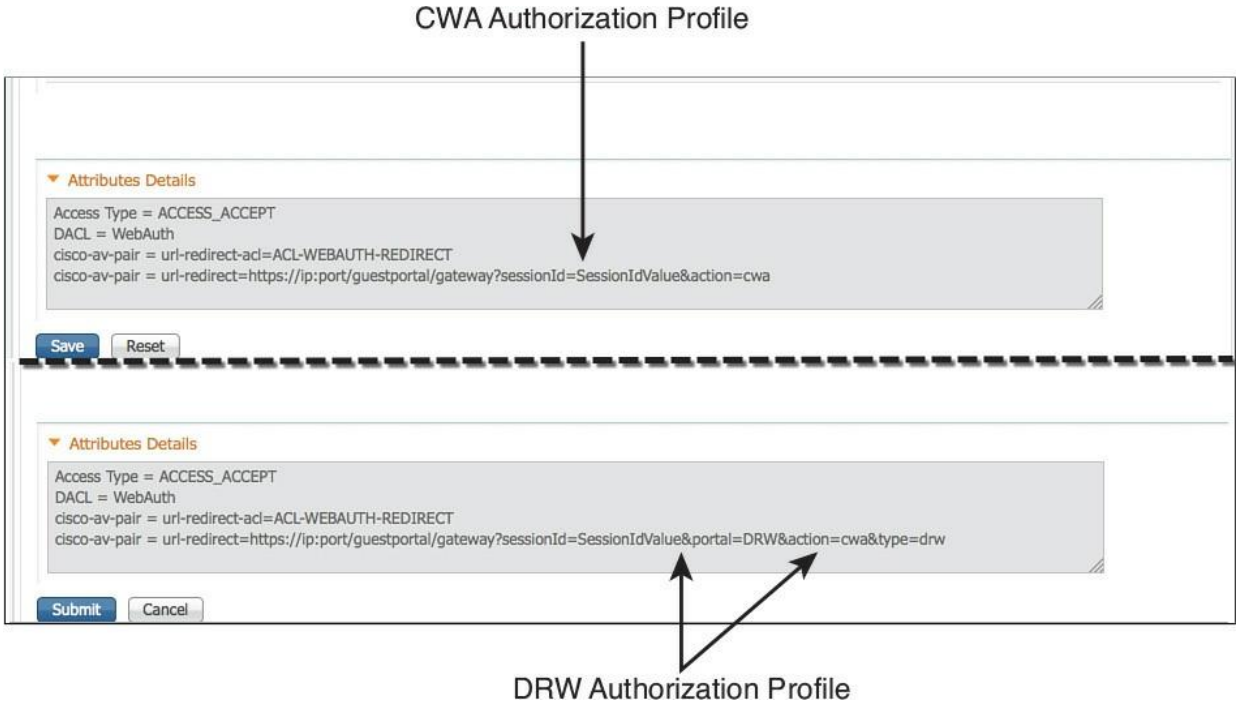


Figure 13-22 Comparing the CWA and DRW RADIUS details.

https://t.me/learningnets

Creating the Rule to Redirect to DRW

Much like the process for CWA, the DRW redirection rule will need some parameters on when to redirect a session to the DRW portal. For the purposes of this example, you will create a rule that redirects all MAB authenticated sourced from the CiscoPress-Guest SSID only. One major difference between CWA and DRW is that a user must only accept the AUP to register her device. No username or password is required.

From the ISE GUI, follow these steps:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Add a new rule above the WebAuth rule, as shown in Figure 13-23.

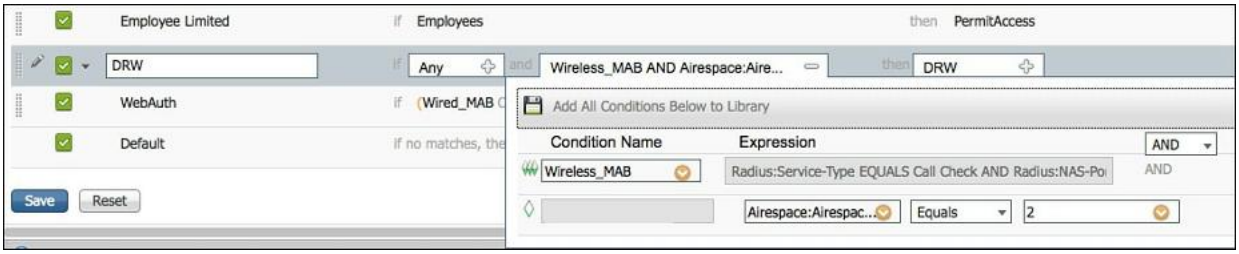


Figure 13-23 Creating the DRW authorization rule.

Step 3. Name the new rule **DRW**.

Step 4. For the first condition, select the existing compound conditions from the library: **Wireless MAB**.

Step 5. For the second condition, create a new one with **Airespace > Airespace-Wlan-Id EQUALS 2**.

In Step 5 you are leveraging a Cisco Airespace attribute to identify the WLAN. In our example, we used an Airespace-Wlan-Id of 2, but this can vary depending on your WLC configuration.

Step 6. Optionally save that condition for use again in other rules.

Step 7. Ensure the **AND** operator is used with the conditions, as shown in Figure 13-23.

Figure 13-23 shows the compound conditions being added to the authorization rule.

Step 8. Use the **DRW** authorization profile you created previously for the result, as shown in Figure 13-23.

Step 9. Click **Done** and then click **Save**.

Figure 13-24 shows the completed DRW authorization rule.

https://t.me/learningnets

Rule Name	Conditions (identity groups and other conditions)	Permissions
Guest	if **Guest** AND GuestFlow	then Internet_Only
Employee CWA	if (Employees AND GuestFlow)	then Internet_Only
Employee and CorpMachine	if EmployeeFullEAPChain	then Employee Full Access
Employee SmartDevice	if (SmartDevice AND Employees)	then Internet_Only
Employee Limited_Certificate	if (Employees AND x509_PKI)	then PermitAccess
Employee Limited	if Employees	then PermitAccess
DRW	if (Wireless_MAB AND Airespace:Airespace-Wlan-Id EQUALS 2)	then DRW
WebAuth	if (Wired_MAB OR Wireless_MAB)	then CWA
Default	if no matches, then PermitAccess	

Figure 13-24 The completed DRW authorization rule.

Creating the Rule to Authorize DRW-Registered Endpoints

Much like what you completed in the guest authorization rule, you will leverage the Identity Group field. Here you are creating a new authorization rule, above the DRW redirection rule that will allow any devices that are part of the DRW-Endpoints group to have Internet-only access.

From the ISE GUI, follow these steps:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Add a new rule above the guest rule, as shown in Figure 13-25.

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
✓	DRW Endpoints	if **DRW-Endpoints** ←	then Internet_Only
✓	Guest	if **Guest** AND GuestFlow	then Internet_Only
✓	Employee CWA	if (Employees AND GuestFlow)	then Internet_Only
✓	Employee and CorpMachine	if EmployeeFullEAPChain	then Employee Full Access
✓	Employee SmartDevice	if (SmartDevice AND Employees)	then Internet_Only
✓	Employee Limited_Certificate	if (Employees AND x509_PKI)	then PermitAccess
✓	Employee Limited	if Employees	then PermitAccess
✓	DRW	if (Wireless_MAB AND Airespace:Airespace-Wlan-Id EQUALS 2)	then DRW

Figure 13-25 The completed DRW-Endpoints rule.

Step 3. Name the new rule **DRW Endpoints**.

Step 4. Select the **DRW-Endpoints** Internal Endpoint Identity Group, as shown in Figure 13-25.

Step 5. Select the **Internet-Only** authorization profile you created earlier as the result.

Step 6. Click **Done** and then click **Save**.

Figure 13-25 shows the completed DRW-Endpoints rule.

https://t.me/learningnets

Verifying Centralized Web Authentication

Congratulations! You've gone through a lot of configuration this chapter. Now that you've gotten the CWA and the DRW configurations complete, it helps to see them in action. There are a number of locations to verify the action. You can examine the effect of the user experience on the client, check the Live Log in ISE, run **show** commands on the wired switch, or even examine the client details on the WLC.

Checking the Experience from the Client

Obviously, a quick way to see whether it's working is to try opening a web browser on the client machine and see if the browser is redirected to a portal.

Figure 13-26 shows the client experience on a wired Windows client being redirected to the CWA portal and the user entering his credentials in the login form.

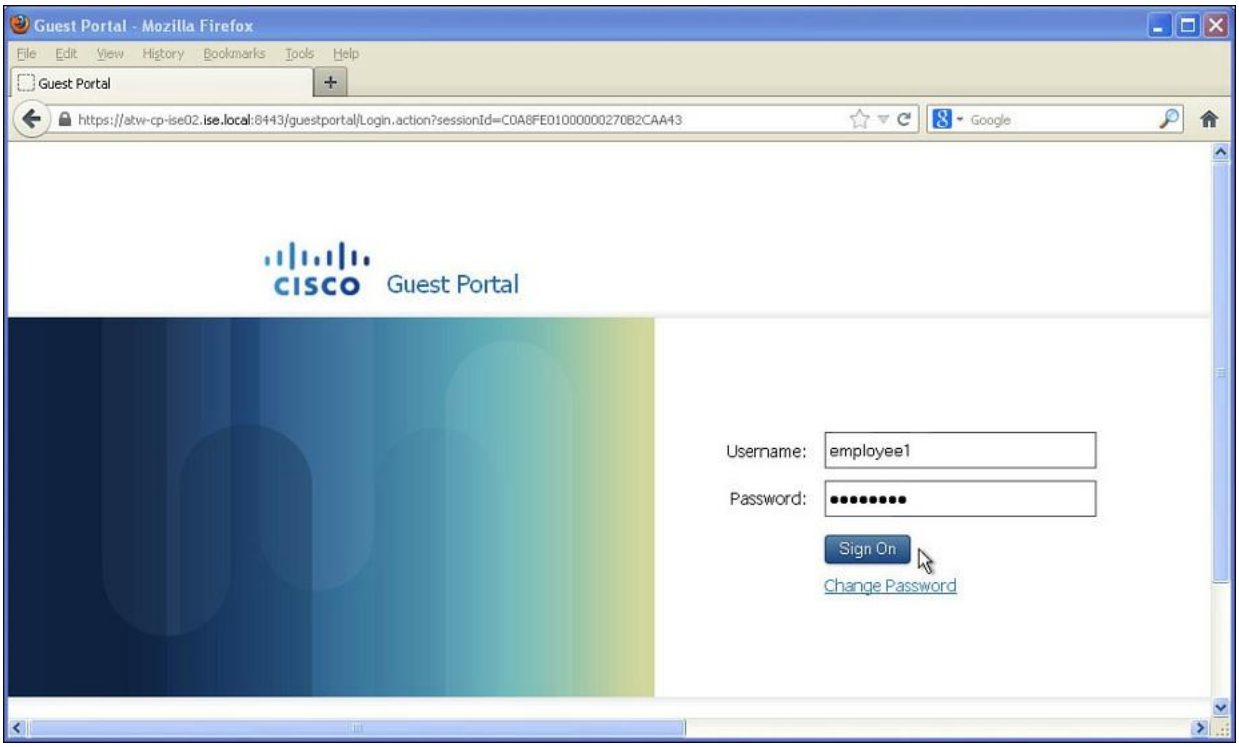


Figure 13-26 Browser redirected to CWA portal.

Figure 13-27 displays the client experience after the authentication was submitted. As the figure shows, the test client was able to connect to the Internet.

https://t.me/learningnets

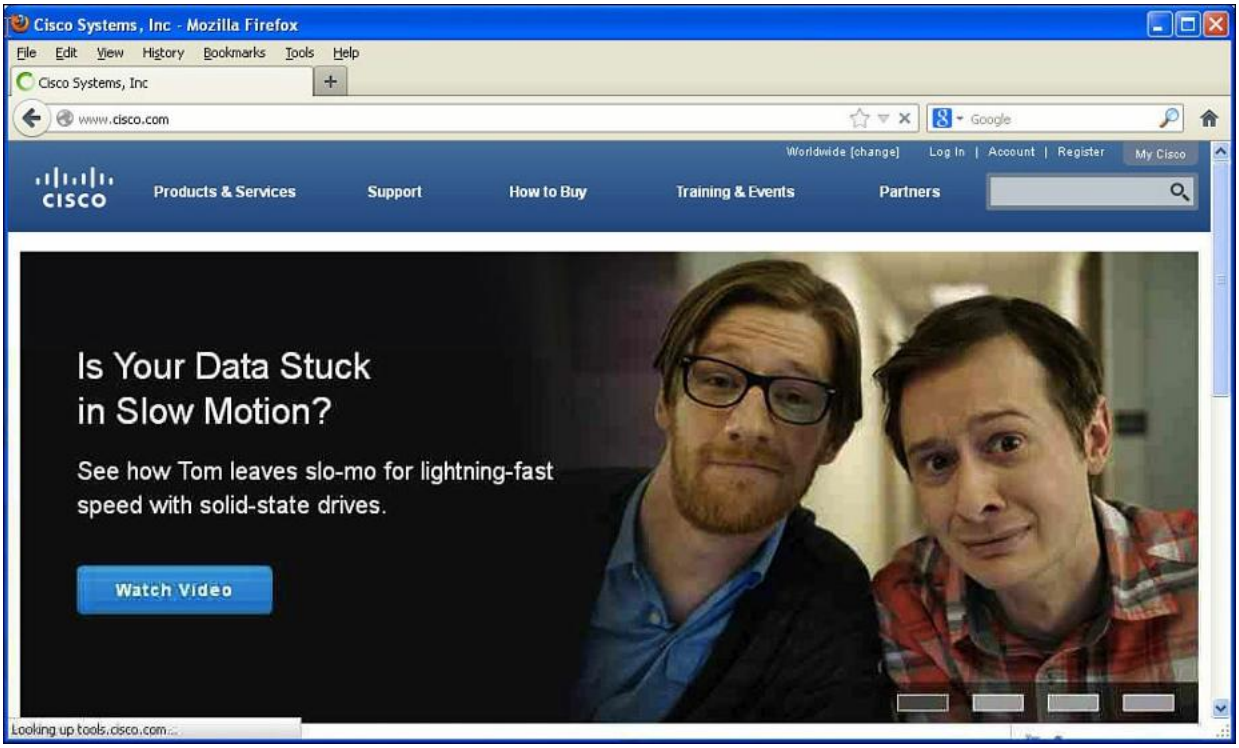


Figure 13-27 Successfully browsing the Internet.

Next, we will try the DRW experience from a wireless client. Figure 13-28 shows an iPad that has successfully associated to the CiscoPress-Guest wireless SSID and is redirected to the DRW portal.

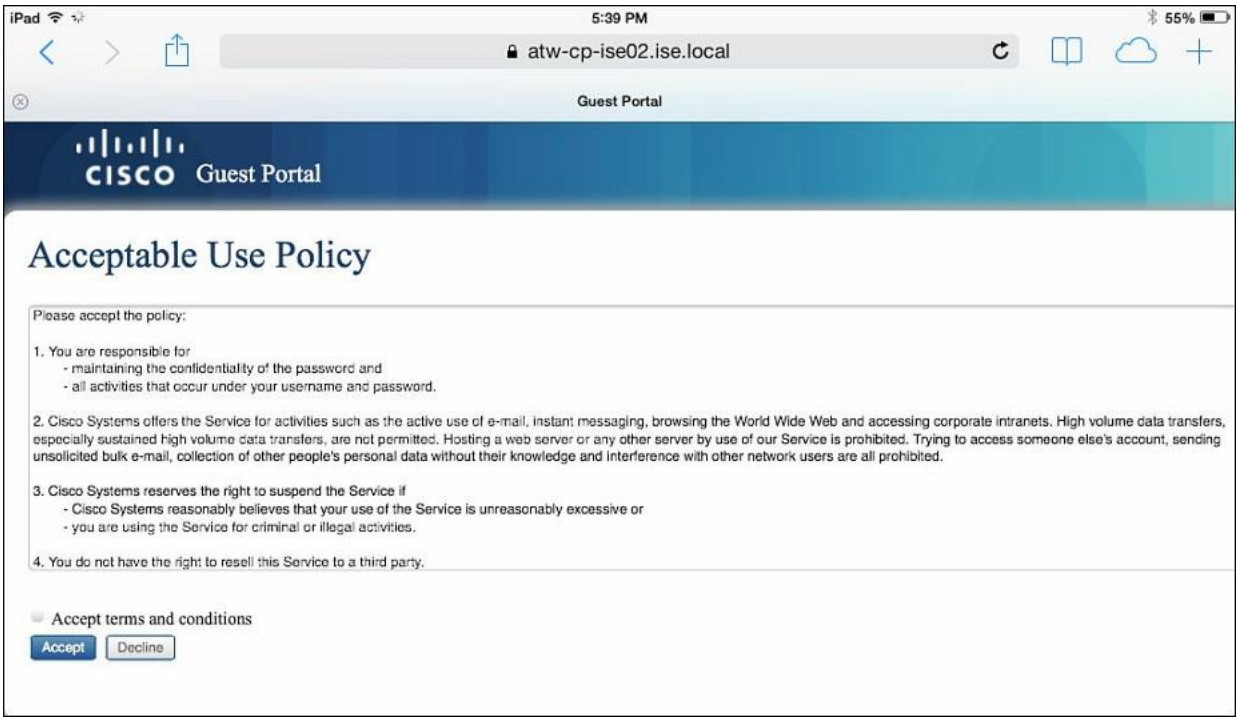


Figure 13-28 DRW's acceptable use policy.

After the user accepts the AUP, the MAC address is automatically added to the DRW-Endpoints identity group. A success message is displayed to the end user, as shown in Figure 13-29.

https://t.me/learningnets

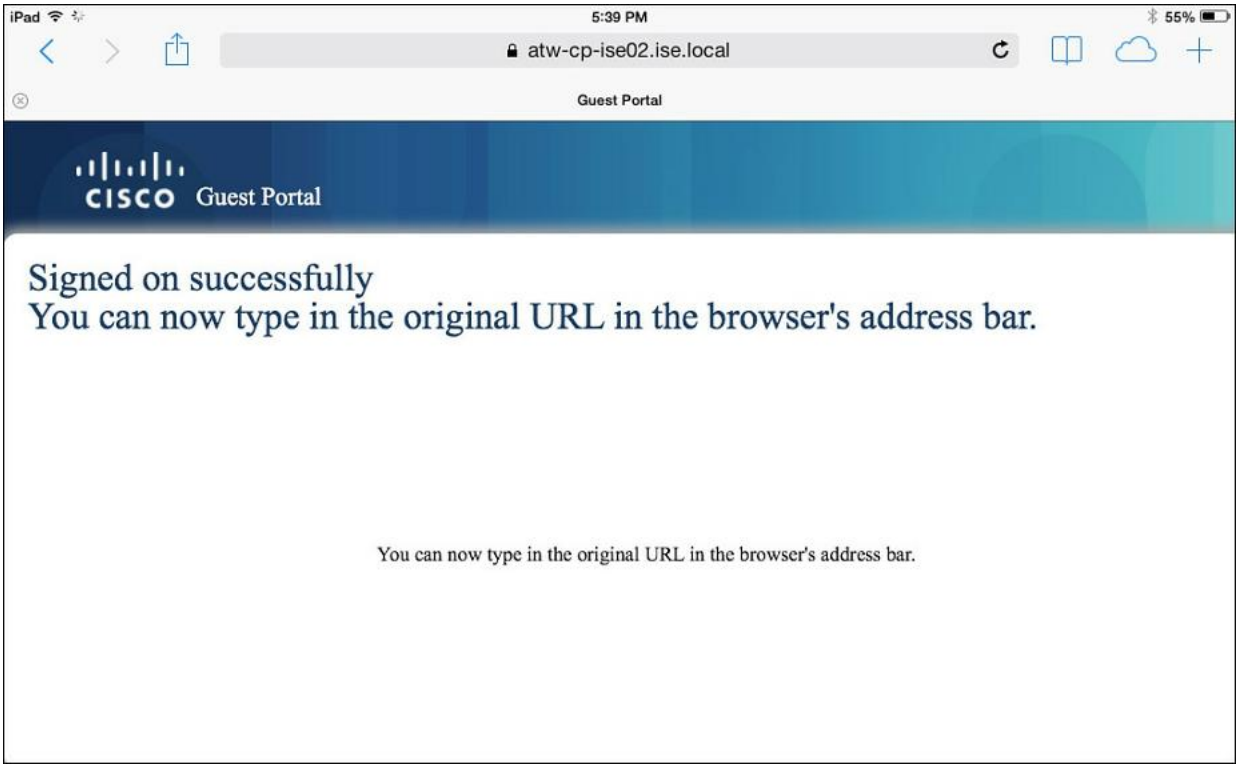


Figure 13-29 Success message.

Now that the device has been successfully registered, the end user is authorized for Internet-only access and is able to reach the Internet, as shown in Figure 13-30.

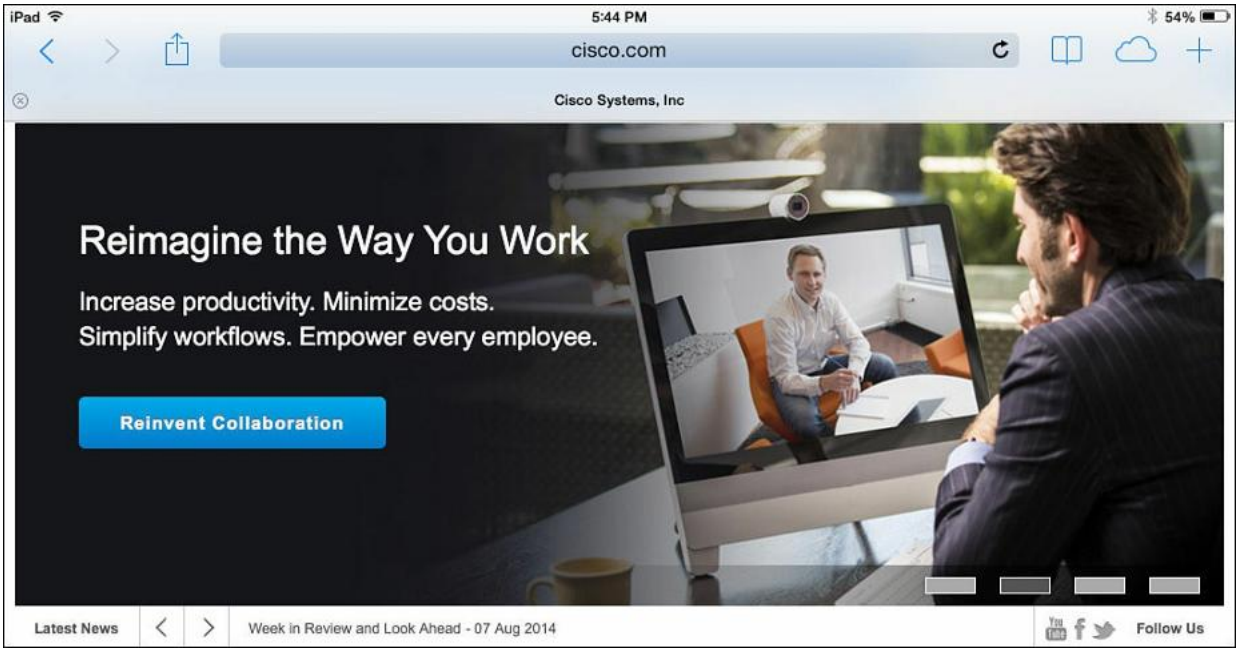


Figure 13-30 Internet access achieved.

Checking on ISE

The next most logical place to look is at the centralized policy server itself. ISE has a number of tools that can be used to verify the WebAuth. The most common is the Live Log. But in the case of DRW, you can also check the Identity Group itself. Many other tools can be used, such as TCPDump, but those tools are not covered in this chapter. For more on those other tools, see Chapter 22, "Troubleshooting Tools."

https://t.me/learningnets

Checking the Live Log

Cisco ISE has a phenomenally useful tool built in to it, commonly called Live Log. Live Log provides a near real-time view of all incoming authentications, CoA, and more. In this section, you will follow the client experience from the perspective of the ISE management console.

Figure 13-31 is highlighting the initial MAB, which has been assigned the CWA authorization result. Immediately following the successful authorization, you see the successful download of the dACL, which is also highlighted in Figure 13-31.

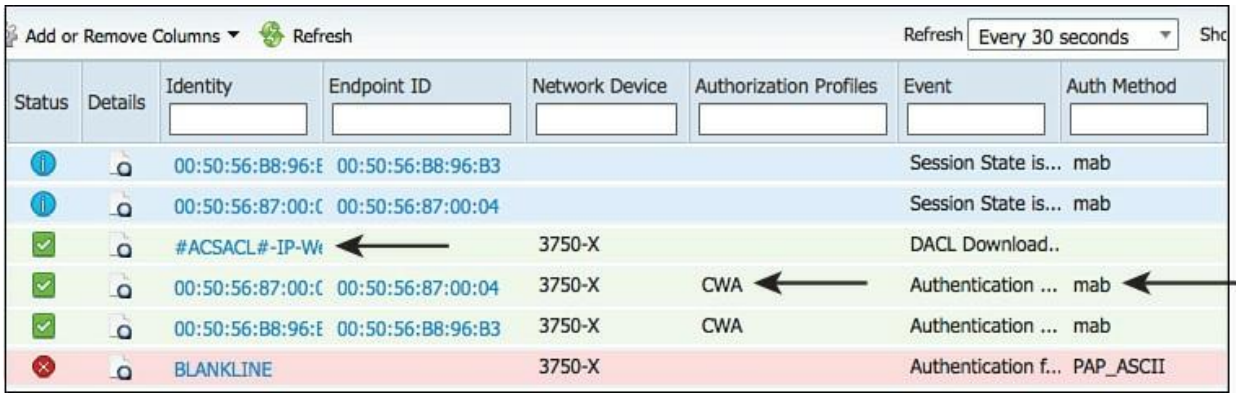


Figure 13-31 Live Log.



After the end user enters her credentials and clicks Submit, a guest authentication is recorded, followed by a CoA (displayed as a dynamic authorization). The CoA is a key function. Specifically, it is a CoA-Reauth and causes the switch to reauthentciate 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.

Due to the correlation of the centralized guest 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. All these steps are shown in Figure 13-32.

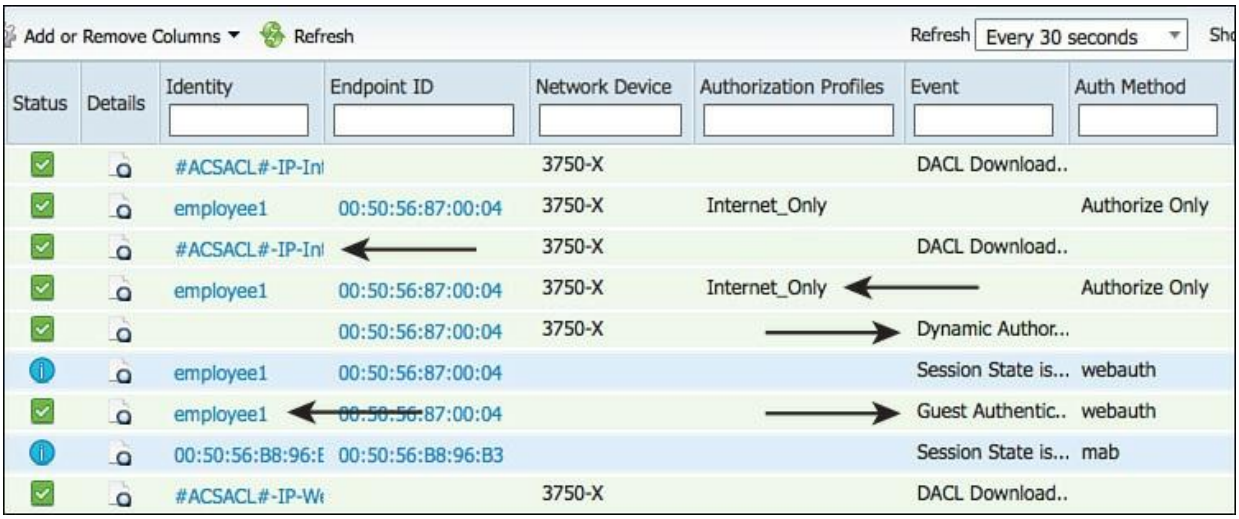


Figure 13-32 Live Log part 2.

https://t.me/learningnets

Checking the Endpoint Identity Group

With DRW, the endpoint is added to the configured Endpoint Identity Group. So, to ensure the redirection to the AUP occurred and the user accepted the AUP, you can check the membership of the Endpoint Identity Group, as shown in Figure 13-33.

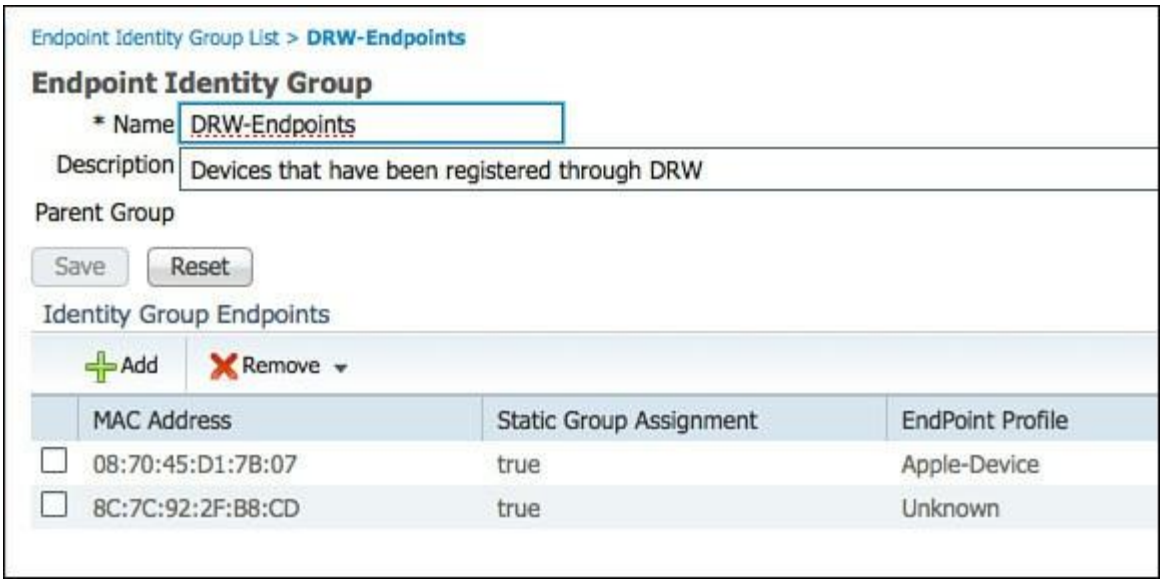


Figure 13-33 DRW-Endpoints Group membership.

Checking the NAD

Certainly, checking the device that is performing the enforcement should be a good way to validate that the CWA is working. In this section you will 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 command is like a Swiss Army knife with all the valuable information it provides. Example 13-1 displays the command and its output for the endpoint as it was being redirected to the CWA portal.

Highlighted in the example are the MAC address, IP address, dACL (listed as an ACS ACL), URL-redirect ACL, and URL to which the end user is being redirected.

Example 13-1 show authentication session interface g1/0/2

Click here to view code image

```
C3560X#  sho authen session int 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-PERMIT_ALL_TRAFFIC-51ef7db1
    URL Redirect ACL: ACL-WEBAUTH-REDIRECT
```

https://t.me/learningnets

```
          URL Redirect: https://atw-cp-ise02.ise.local:8443/guestportal/gateway?sess
ionId=C0A8FE01000000270B2CAA43&action=cwa
       Session timeout: N/A
          Idle timeout: N/A
     Common Session ID: C0A8FE01000000270B2CAA43
       Acct Session ID: 0x0000002D
               Handle: 0x68000028

Runnable methods list:
        Method   State
        mab      Authc Success
        dot1x    Not run
```

Viewing the Client Details on the WLC

With the WLC, navigating to **Monitor > Clients** will show you a list of all clients currently associated to that controller. Clicking the MAC address will bring up the details about the client and its association, including the authentication information such as the redirection and run state.

Figure 13-34 shows the details for that client, with a focus on the Security Information section. Pay attention to the RADIUS NAC state, which is set for **CENTRAL_WEB_AUTH**.

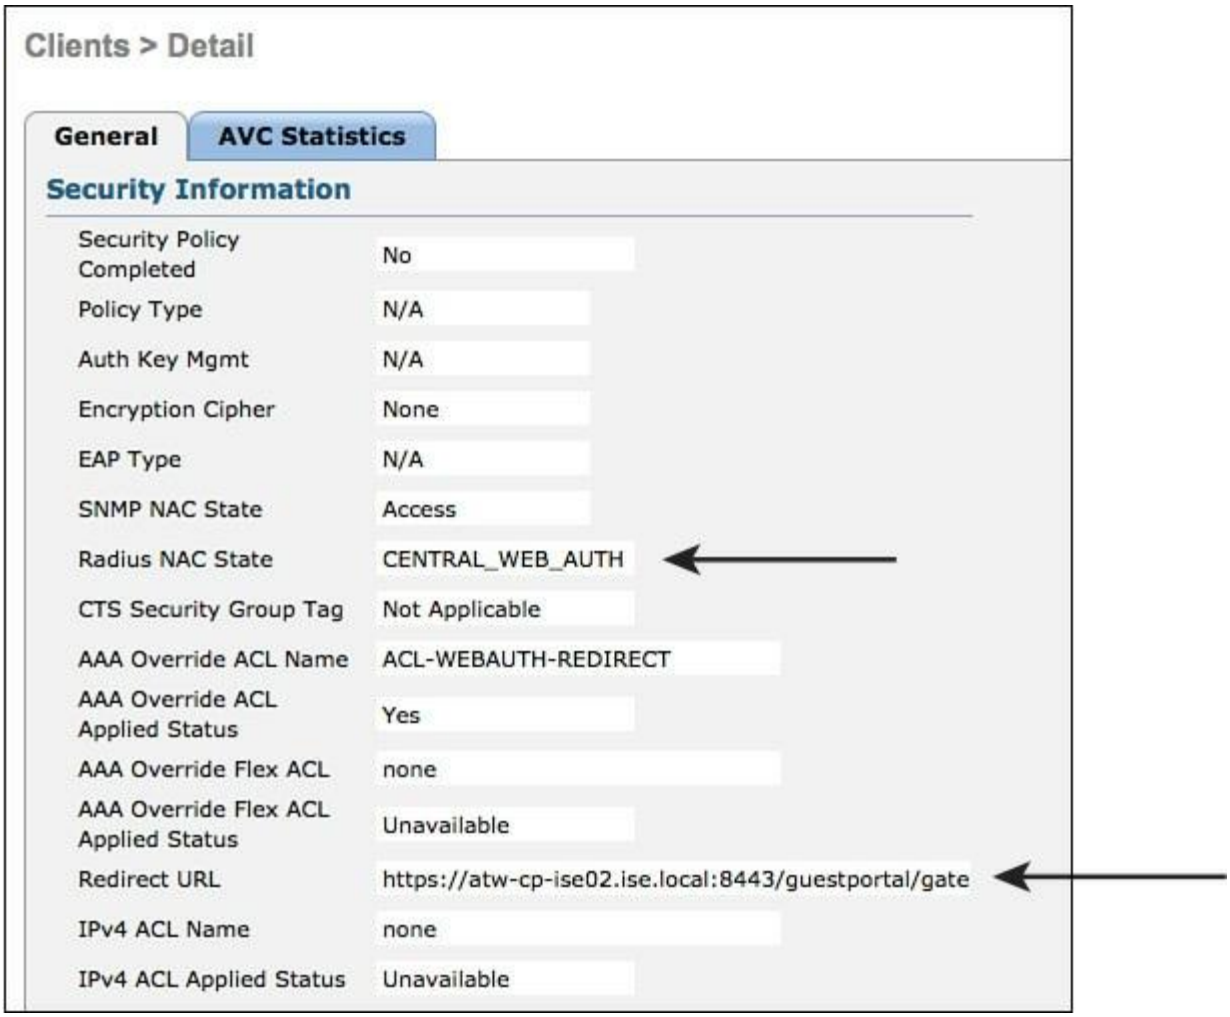


Figure 13-34 Security Information section of client details.

Figure 13-35 shows the same client details after the endpoint was registered. Notice the RADIUS NAC state is now RUN. This is a key setting; the redirection will never work if the state is RUN because this is the final state.

https://t.me/learningnets

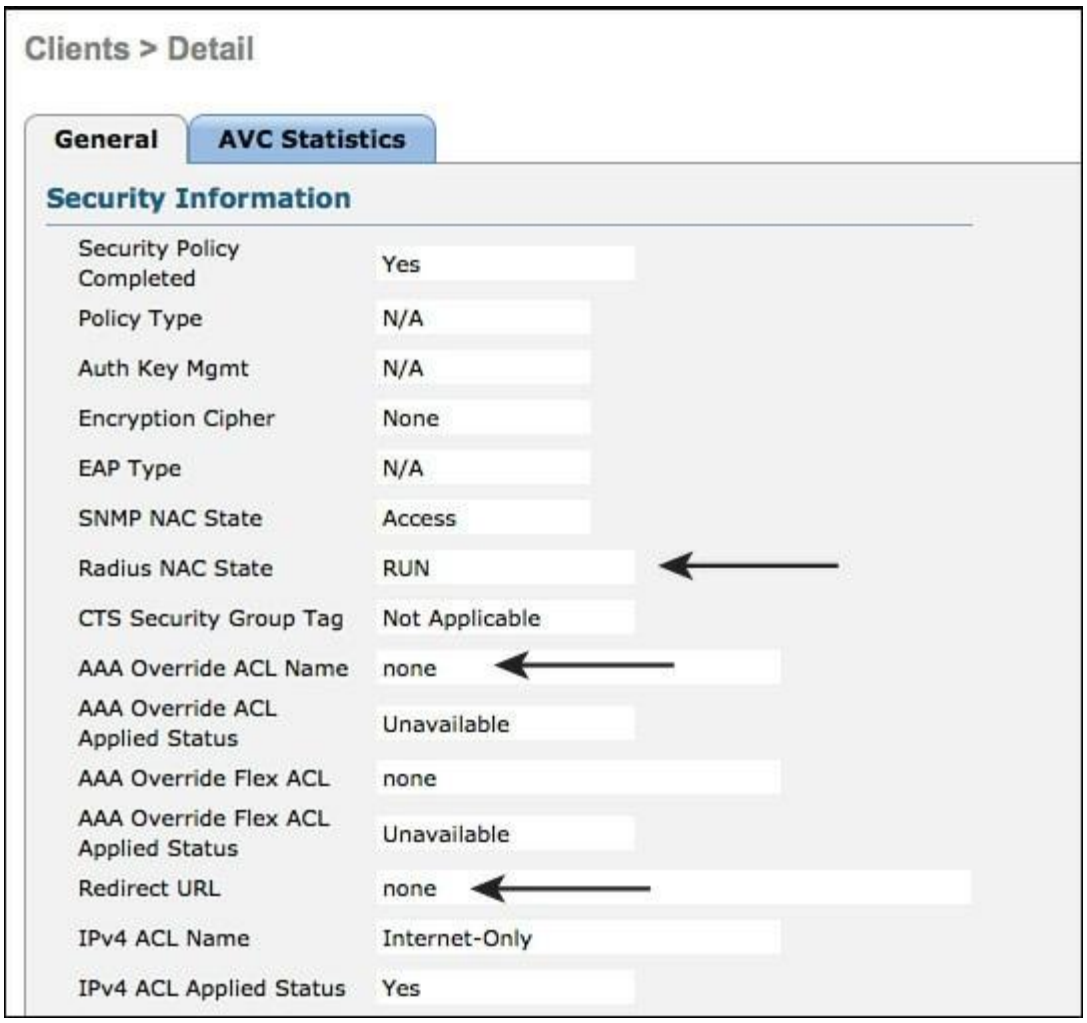


Figure 13-35 Security Information section of client details—run state.

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 13-3 lists a reference of these key topics and the page numbers on which each is found.



https://t.me/learningnets

Key Topic Element	Description	Page
Paragraph	Local Web Authentication (LWA)	346
Paragraph	Centralized Web Authentication (CWA)	347
Paragraph	DRW and MyDevices	349
Paragraph	IOS URL-Redirection ACL	351
Paragraph	WLC URL-Redirection ACL	353
Paragraph	CoA with CWA and DRW	373

Table 13-3 Key Topics for Chapter 13

https://t.me/learningnets

Chapter 14. Deploying Guest Services

This chapter covers the following exam topics:

- Guest Services Overview
- Configuring the Web Portal Settings
- Configuring the Sponsor Portal Policies
- Managing Guest Portals
- Building Guest Authorization Policies
- Provisioning Guest Accounts from the Sponsor Portal
- **Verifying Guest Access in ISE and the Network Access Device**

So far, much of our focus has been on the admission of employees and trusted users onto the network. However, ISE is capable of much more than that. What should we do when we need to provide a user temporary access to the network, such as a visitor, consultant, customer, or contractor? How do we control the ability of a user to create these temporary accounts, and how do we control the level of access that is given to these temporary users?

In this chapter, you will learn how Cisco ISE supports these temporary visitors to the network—a feature referred to as Guest Services. You will learn how to configure the various portals and policies within ISE that are used by the sponsors of these temporary users.

In a similar manner, you will also configure the guest portals—the webpages whereby the guest will authenticate—and build the authorization policy that manages the guest user's level of access that will follow this authentication.

To finish this chapter, you will provision a guest account using the configured sponsor portal and verify the guest's access on the network access device.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Guest Services Overview	1-2
Configuring the Web Portal Settings	3
Configuring the Sponsor Portal Policies	4-5
Managing Guest Portals	6
Build Guest Authorization	7
Provisioning Guest Accounts from Sponsor Portal	8
Verify Guest Access in ISE and Wwitch/WLC	9-10

Table 14-1 "Do I Know This Already?" Section-to-Question Mapping

https://t.me/learningnets

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. ISE Guest Services use which of the following approaches to authenticate a user?

   a. Badge

   b. WebAuth

   c. TACACS+

   d. SSH

2. The sponsor and guest portals can run on which of the following ISE personas?

   a. Admin

   b. MnT

   c. PSN

   d. a and b

   e. a and c

   f. b and c

3. True or False: A network administrator can customize the guest portals to run on any port greater than 1024.

   a. True

   b. False

4. Which default sponsor groups are available on ISE? (Select three.)

   a. SponsorAllAccounts

   b. SponsorADAccounts

   c. SponsorAdministrator

   d. SponsorGroupGrpAccounts

   e. SponsorAllUsers

   f. SponsorGroupOwnAccounts

5. When using Active Directory group membership as authentication and authorization for sponsors, which of the following must occur?

   a. ISE must be associated to the domain.

   b. The sponsor must create all guest accounts on the Active Directory Server.

   c. The Active Directory identity store must be part of the identity source sequence for the sponsor portal.

   d. a and b.

   e. b and c.

   f. a and c.

6. Under the Operations tab of the portal configuration page, which of the following items can be configured?

https://t.me/learningnets

a. Guest Device Registration

b. Allow or Require Guest to change password

c. Guest Self-Service

d. Acceptable Use Policy frequency

e. All of the above

7. What are the three configurable options for a sponsor group?

a. Authorization Levels, Guest Roles, Time Profiles

b. Access-List, VLAN, Security Group Tag

c. Switch, Router, Firewall

d. Centralized WebAuth, Network Supplicant Provisioning, Device Registration Webpage

8. Which of the following are options for provisioning guest accounts on Cisco ISE?

a. Guest, Contractor, Consultant

b. OneDay, OneWeek, OneMonth

c. Individual, Import, Random

d. Full, Basic, InternetOnly

9. Which security policy must be enabled on the Guest WLAN/SSID to facilitate WebAuth on a Cisco WLC?

a. WPA2 with 802.1X Key Management

b. WPA2 with 802.1X and CCKM Key Management

c. MAC Filtering and RADIUS NAC

d. Open

10. To verify a guest user's access policy on a Cisco switch, you should run which of the following commands?

a. show crypto ipsec sa

b. show aaa authorization <username> **details**

c. show authorization level guest interface <if_name>

d. show authentication sessions interface <if_name> **details**

Foundation Topics

Guest Services Overview

Cisco ISE can be flexible in its deployment, allowing network administrators to configure varying levels of access for a number of different user scenarios—such as for full-time employees and visitors alike. This section focuses on the configuration of Guest Services.

Guest Services and WebAuth

As you just learned, Guest Services is the ability for ISE to provide access to a temporary user. This temporary user could be one of the following:

- **Visitor**—Possibly the employee's sick child who came to work with his mother. The mother might choose to provide the user basic Internet access so the child can watch cat videos all day.
- **Consultant**—This user might need a slightly higher level of access than the child who is visiting his parent. However, we don't want to give the consultant an unfettered level of access. This consultant's access level may allow her to reach a few servers that are internal to the corporate

https://t.me/learningnets

network as well as access to the Internet.

- **Contractor**—This user, because he is acting as a temporary employee, can require the same amount of access as an employee—just for a shorter period of time.

This temporary user will be added to the network by a trusted user, such as a network administrator, an employee, or a manager—and often allowed access to a finite set of network resources based on the sample use cases given previously.

As we discuss Guest Services, it is imperative that we also discuss Web Authentication (WebAuth). WebAuth is exactly like it sounds—the end user will use a webpage to authenticate, requiring only a valid username and password. Because you typically have no control over the hardware or software capabilities of a guest's endpoint (the device might not support 802.1X or have no supplicant), WebAuth is well suited for guest services. How must we configure ISE and the network access device (NAD) to allow for this authentication? We discuss the specifics of the authentication and authorization process later in this chapter. In the meantime, we'll briefly discuss the theory behind the Guest Services operation and how it interacts with the WebAuth function of ISE.

When a user first authenticates to the network, ISE issues the endpoint a very restrictive authorization profile (see Figure 14-1). This restrictive level of access provides the endpoint basic network connectivity—just enough to allow it to access ISE's Guest Services web portal. As the user attempts to browse a website, he is redirected to ISE's WebAuth portal. This WebAuth portal allows the user to provide login credentials—a pre-provisioned username and password. After authenticating via this web portal, ISE issues a Change of Authorization (CoA) to the NAD. This CoA forces a reauthorization of the endpoint on the NAD. Via this reauthorization process, ISE deploys a different authorization profile that provides the endpoint its final level of access. This level of access—in both the length and the content—can be as permissive or restrictive as necessary, commensurate with the use cases as defined in the company's security policy and implemented on ISE.

https://t.me/learningnets

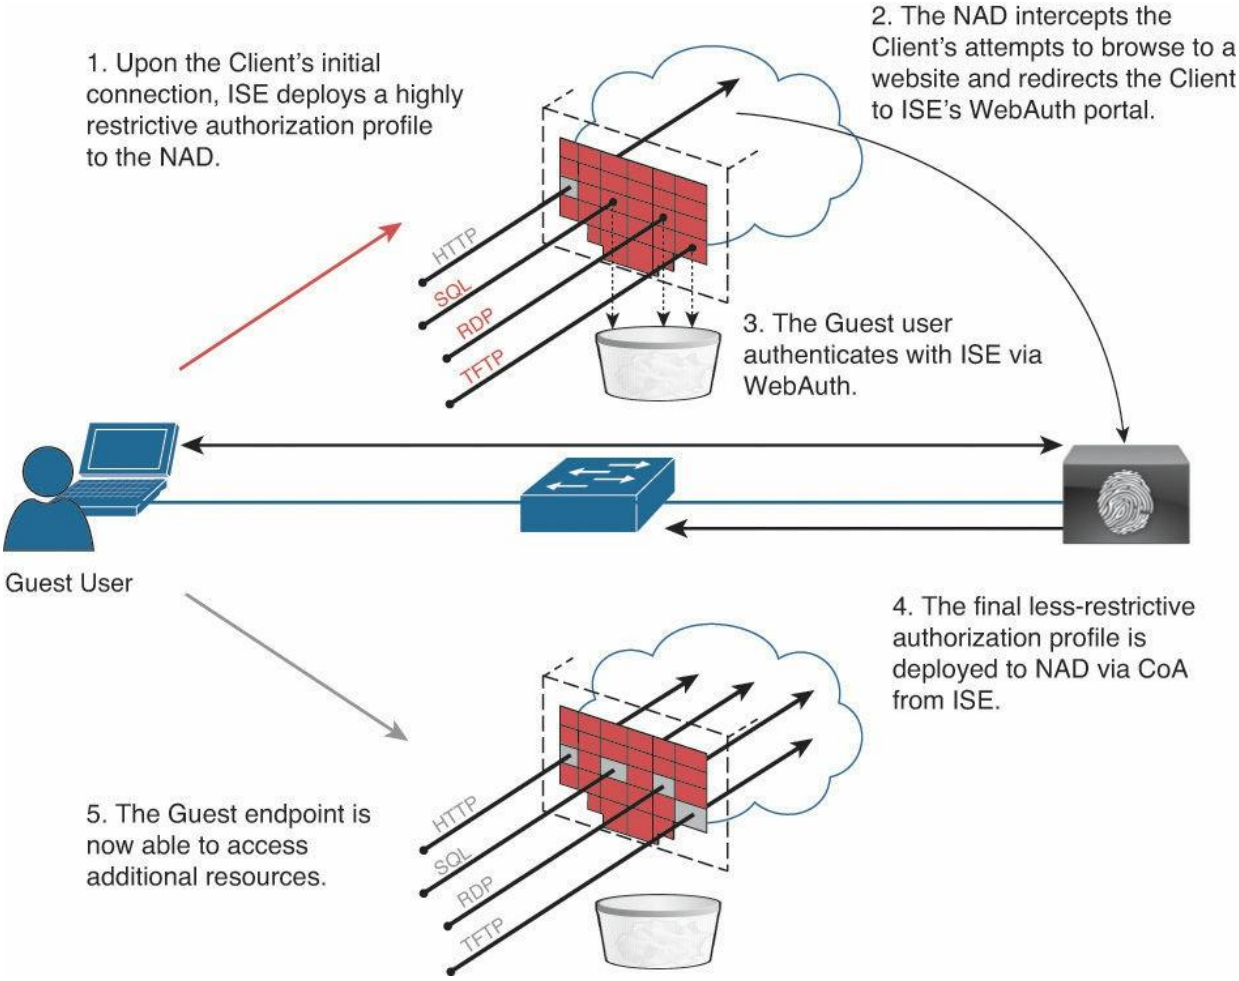


Figure 14-1 WebAuth process flow.

As we review the configuration of the Guest Services authorization policy later in this chapter, you will have a better understanding of what happens during these individual steps.

Portal Types

Cisco ISE has several portal types that we'll leverage for Guest Services. The three major types of application portals as they relate to guest services are

- **Admin portal**—Shown in Figure 14-2, it allows the network administrator to configure sponsor and guest user policy.

https://t.me/learningnets

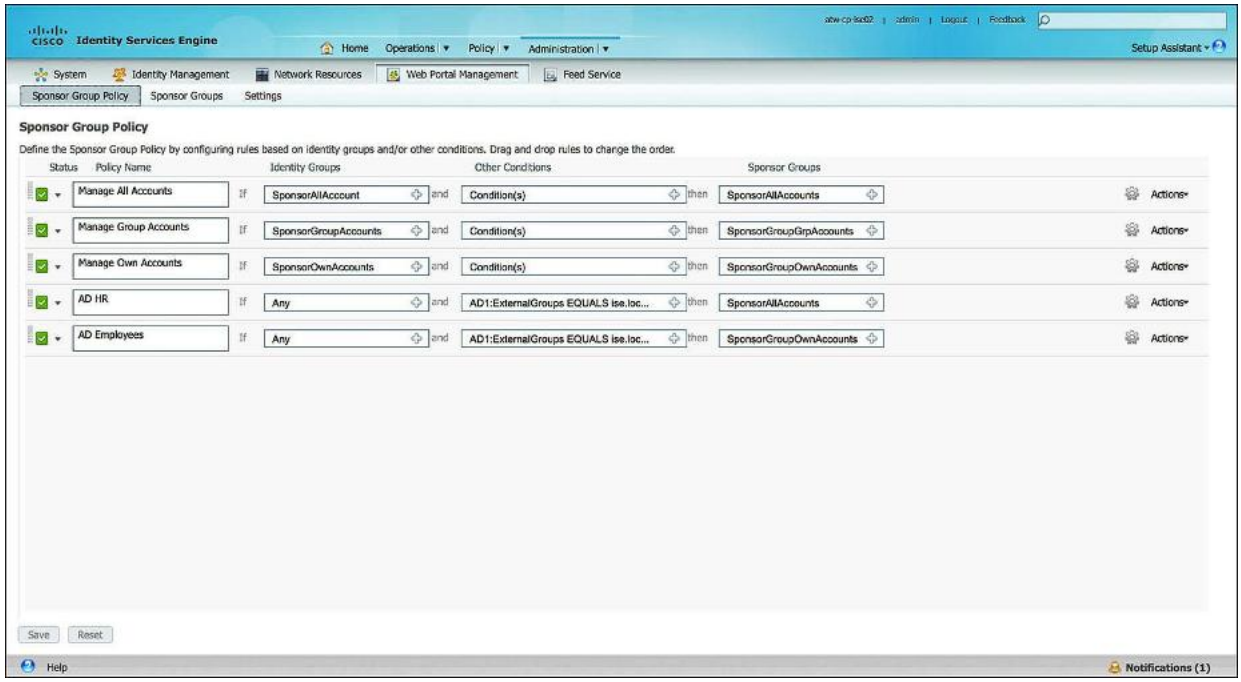


Figure 14-2 Admin portal.

■ **Sponsor portal**—Shown in Figure 14-3, it allows the sponsor to manage guest user accounts.

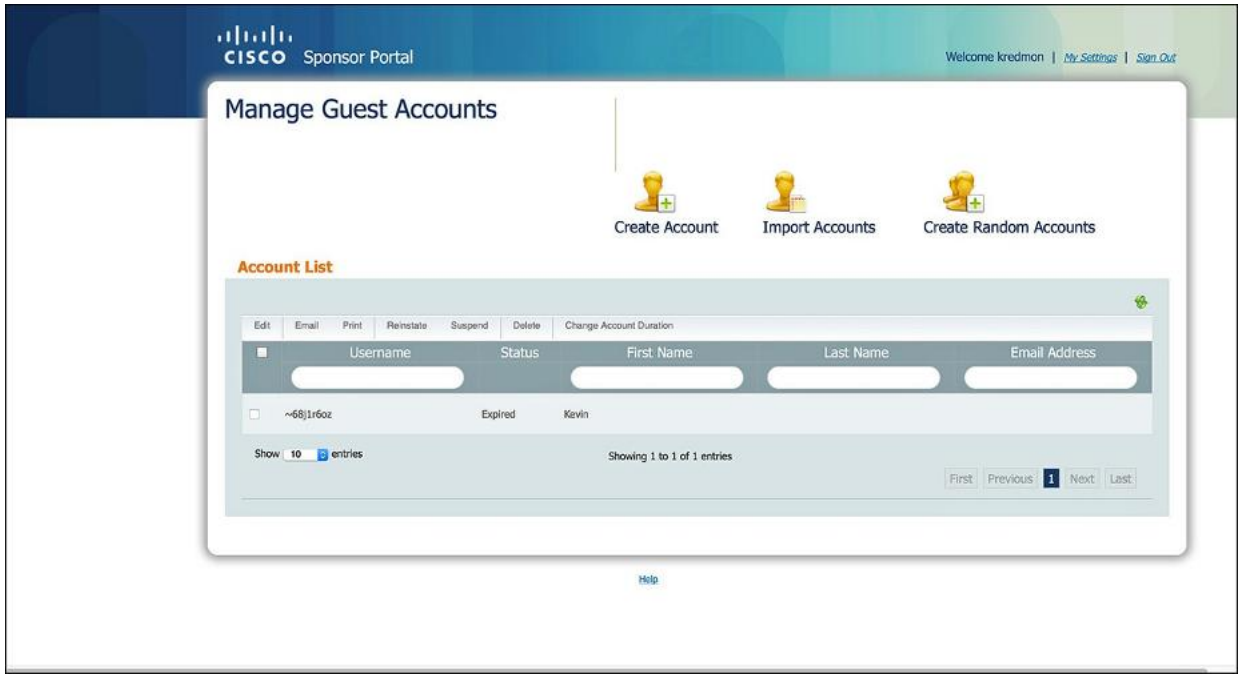


Figure 14-3 Sponsor portal.

■ **Guest user portal**—Shown in Figure 14-4, it allows the guest user to log in.

https://t.me/learningnets

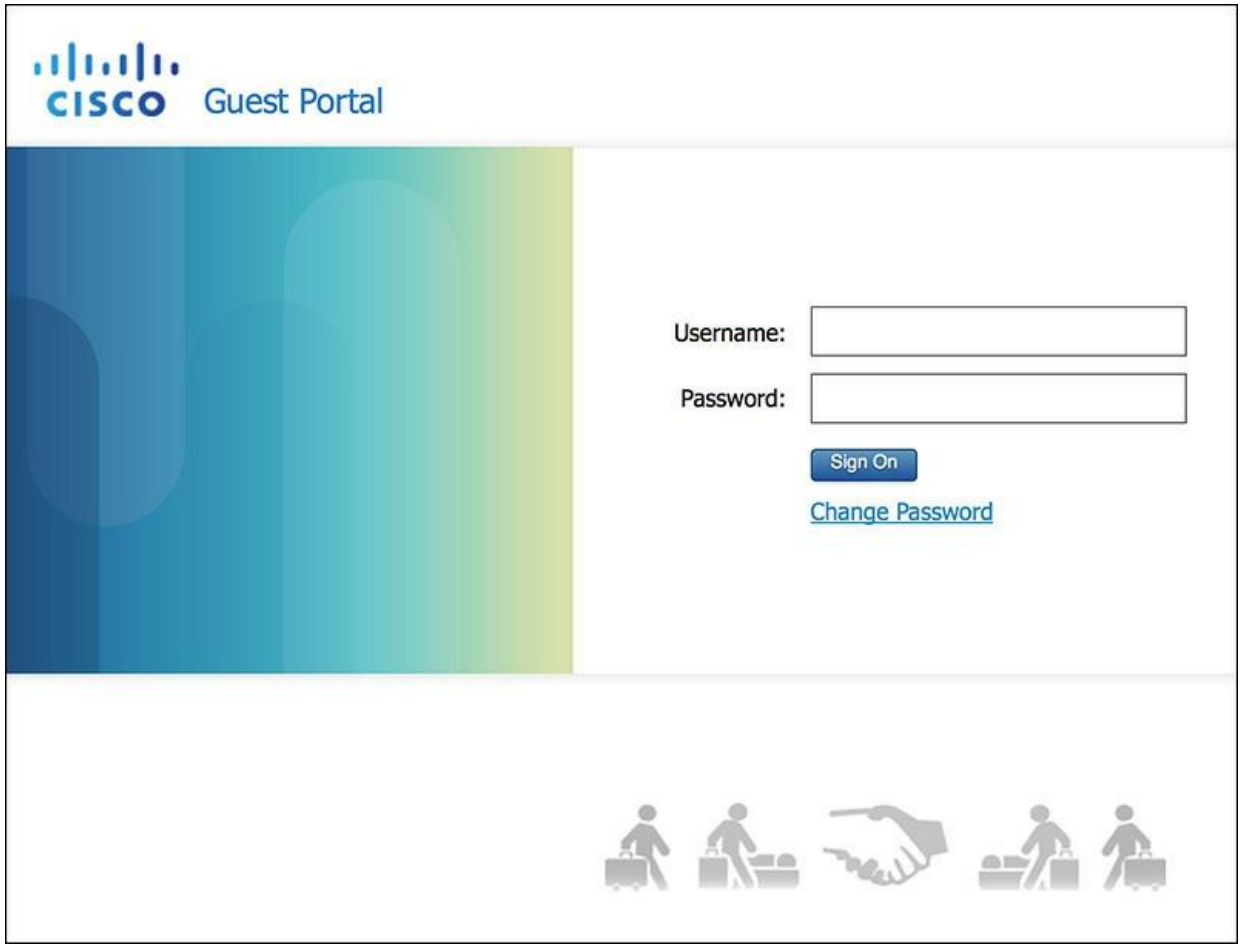


Figure 14-4 Guest user portal.

The guest user portal contains a number of elements, including these:

- **Guest User Login screen**—This screen within the guest portal prompts the user for a username and password field (see Figure 14-4). This is the authentication piece of WebAuth.
- **Acceptable Use Policy screen**—This screen enables administrators to provide the end user the rules, regulations, and guidelines—the Terms of Use—for using the network (see Figure 14-5). The user might get this screen upon the first login only or each time she logs in. This is configurable and can be customized with the requisite Acceptable Use Policy for your organization.

https://t.me/learningnets

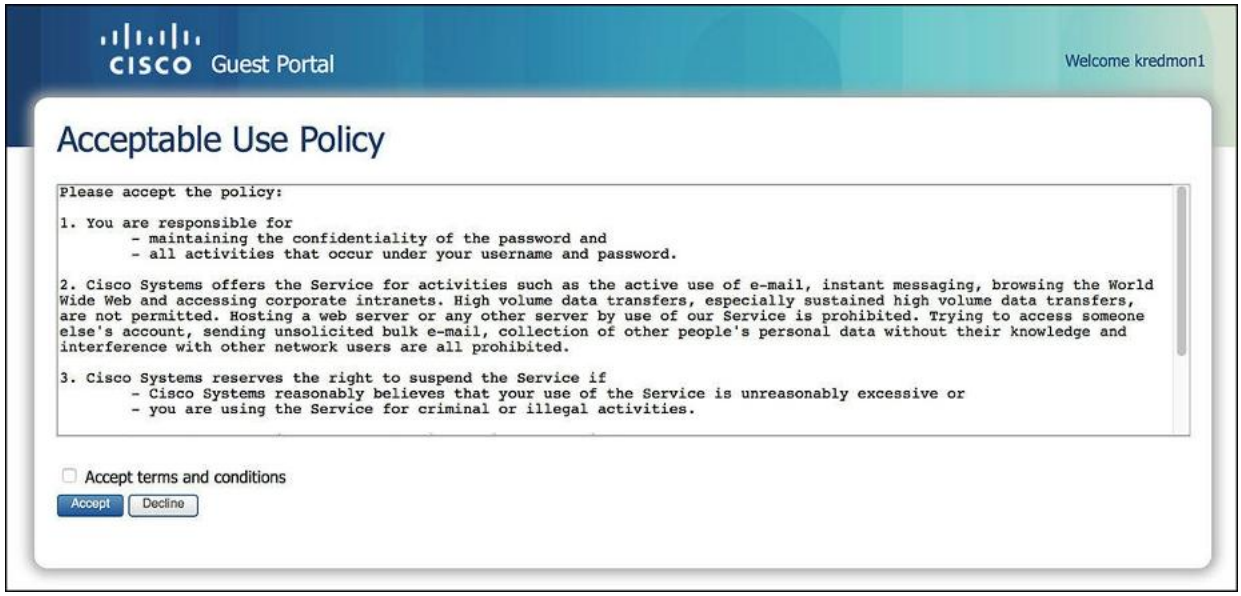


Figure 14-5 Acceptable Use Policy screen.

- **Required Password Change screen**—This screen requires the user to change the provided password. This can help ensure that the guest user can change the password to something that she knows as well as provide the guest user the option to change the password to something more private that only she knows—as the sponsor already knows the initial password that was issued. This is enabled by default.
- **Allow Password Change screen**—This screen allows the user to change the password if she chooses to (see Figure 14-6). If she clicks a Change Password link, this screen appears. This feature is enabled by default.

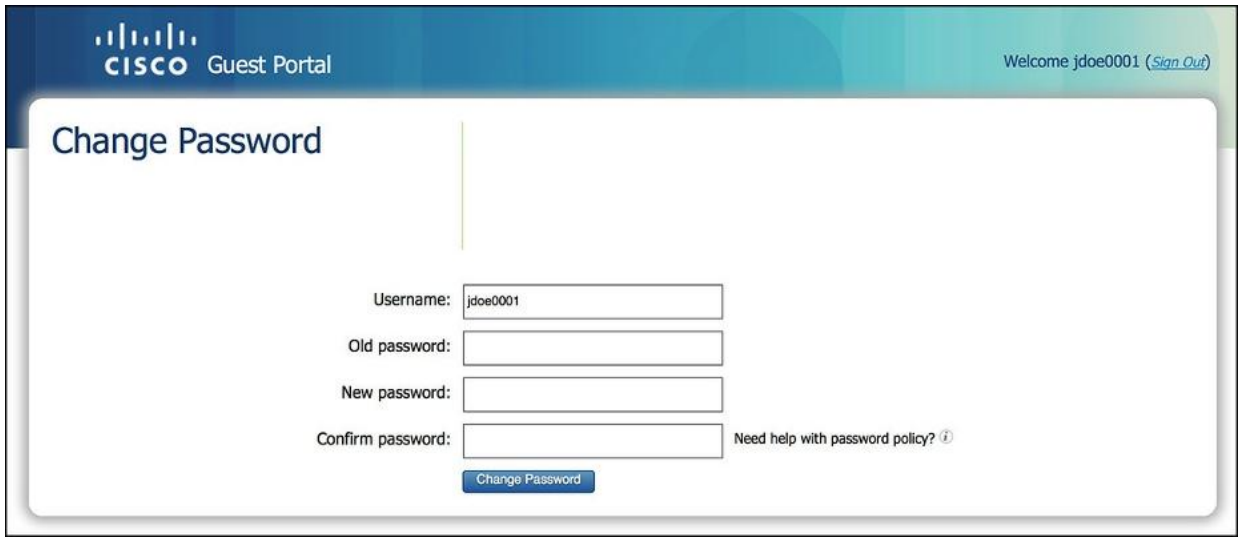


Figure 14-6 Required Password/Allow Password Change screen.

- **Self-Registration screen**—This screen is an optional screen that lets guests to set up their own accounts (see Figures 14-7 and 14-8).

https://t.me/learningnets

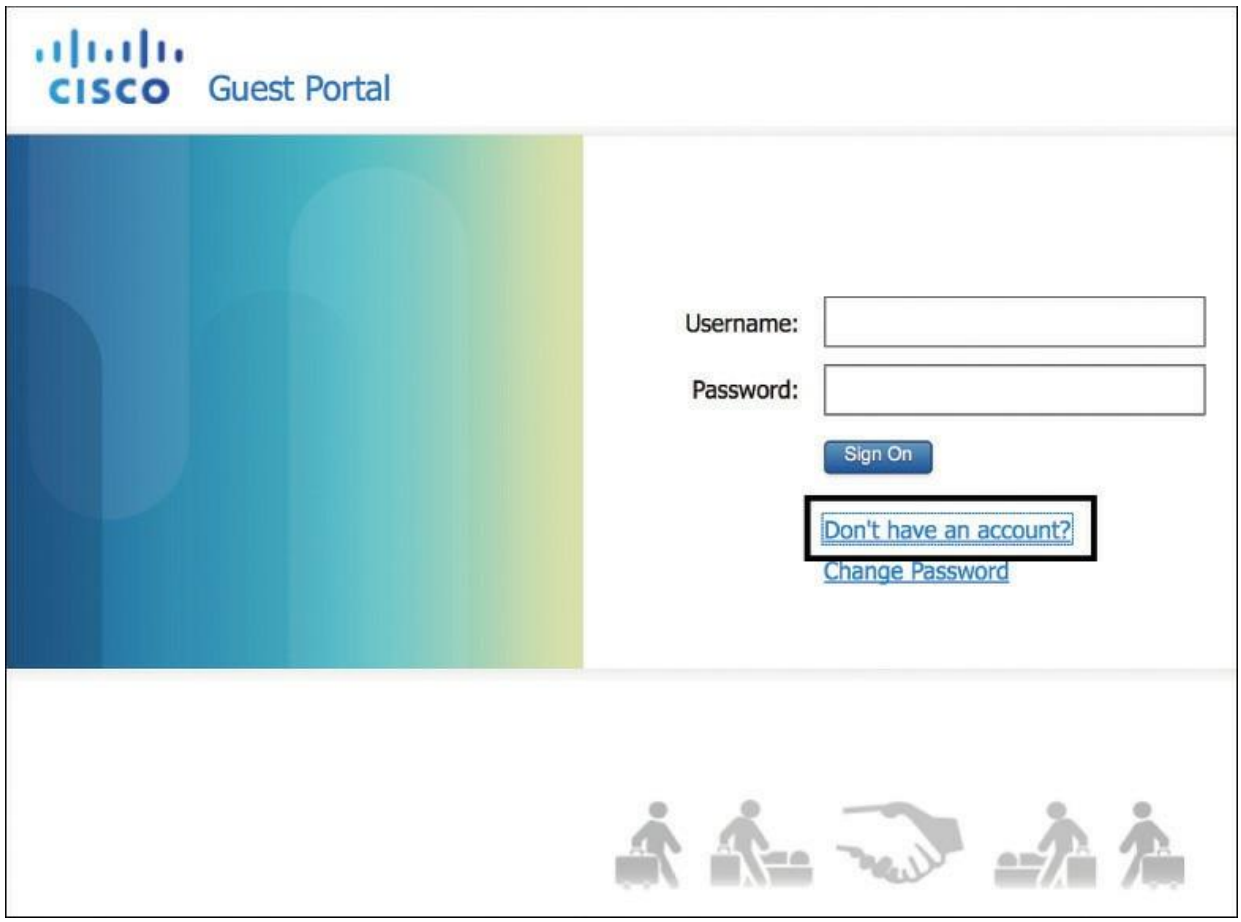


Figure 14-7 Self-Registration Login screen.

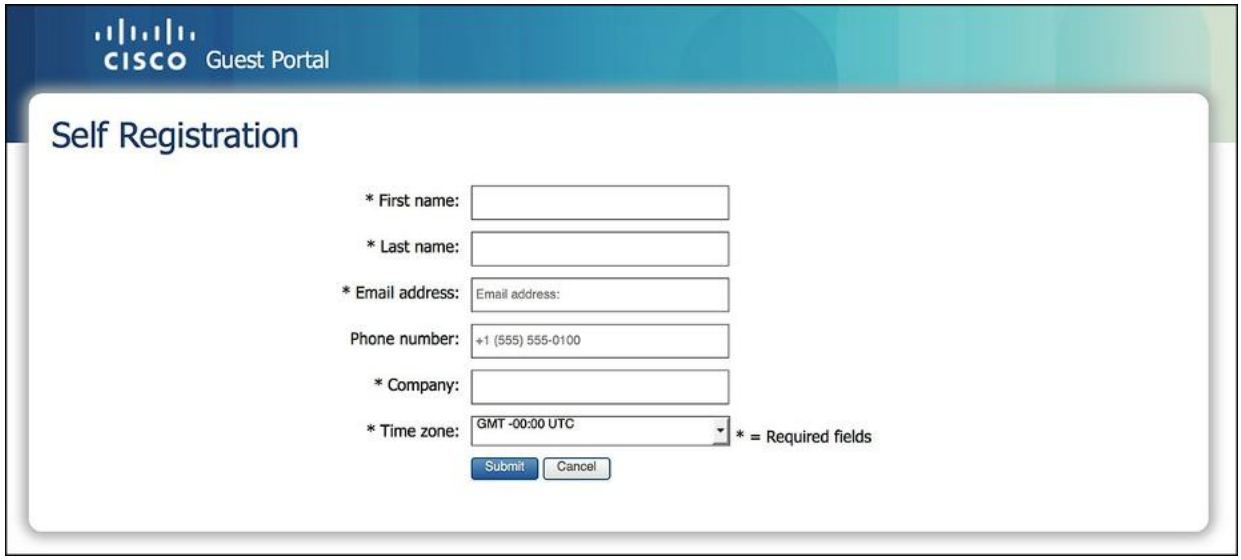


Figure 14-8 Self-Registration Creation screen.

To change whether the fields are optional, mandatory, or unused, you can go to **Administration > Web Portal Management > Settings, Guest > Details Policy** within the administration portal (see Figure 14-9). If you want to modify the Optional Data fields to be something more meaningful, you can do so by going to **Administration > Web Portal Management > Settings > Sponsor > Language Template** and choosing the appropriate language that should be modified. After selecting the appropriate language, you can now choose which sponsor portal page needs to be modified (for example, Configure Self Registration).

https://t.me/learningnets

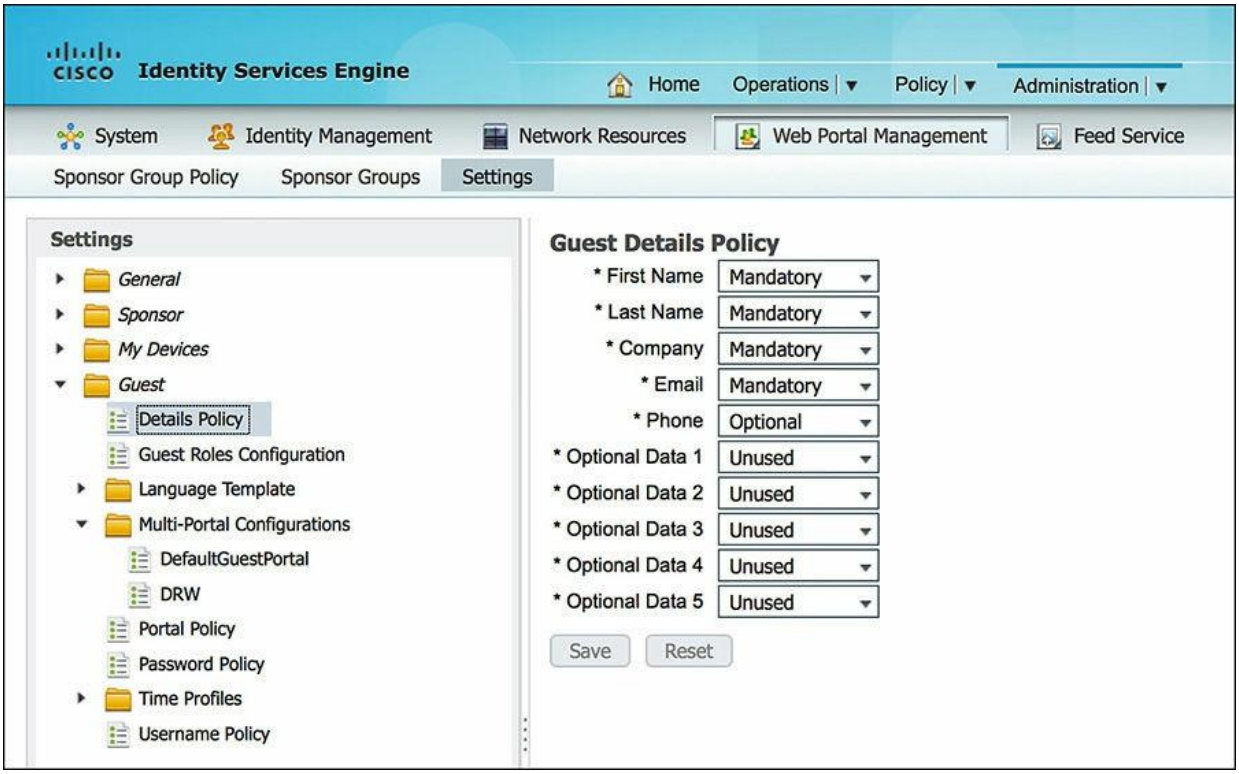


Figure 14-9 Self-Registration input field configuration.

- **Device Registration screen**—This screen enables a guest user to self-register a predefined number of endpoints by MAC address (see Figure 14-10). This registration results in static population of an internal endpoint identity database. The user requires valid credentials if she wants to register devices.

https://t.me/learningnets

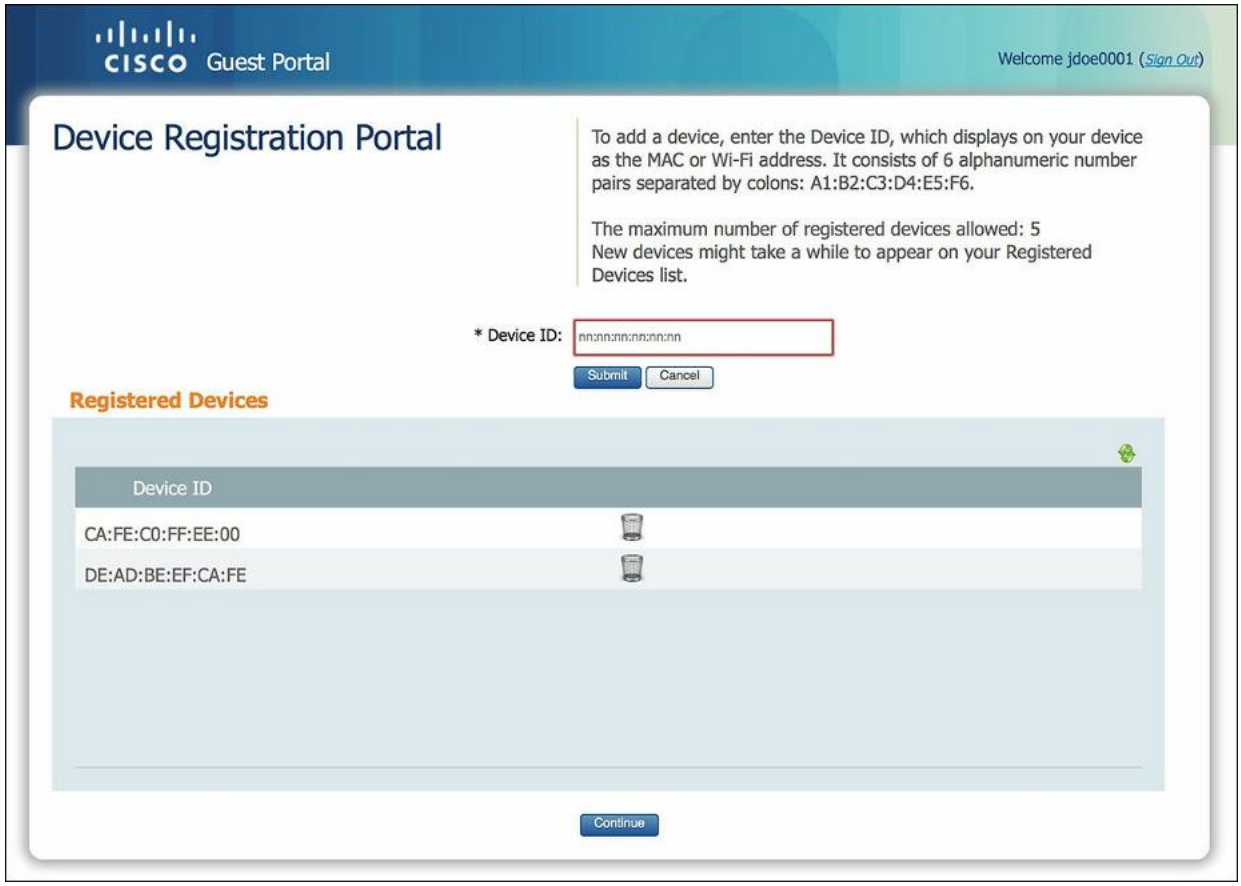


Figure 14-10 Device Registration screen.

Some of these elements can be optionally leveraged to customize the guest user's experience as she joins the network.

Key Topic

In a distributed ISE deployment, the Admin portal runs exclusively on the Admin node. The sponsor and guest user portals run on those policy service nodes (PSNs) that are configured to run session services.

We discuss each of these portal types in more detail as we move along through this chapter.

Configuring the Web Portal Settings

Now that we know what each portal is capable of doing, we'll walk through the configuration of many of these portals. The default configuration of the portals will often work fine for many deployments. However, we will provide a quick overview of several of the more common configuration options in case you choose to customize your deployment to match your company's security policy.

The base configuration location to change the General Web Portal Settings within ISE is **Administration > Web Portal Management > Settings > General**. Expanding the **General** folder, you will see options for **Portal Theme, Ports,** and **Purge** (see Figure 14-11). We'll be focusing mainly on the **Ports** configuration at this time.

https://t.me/learningnets

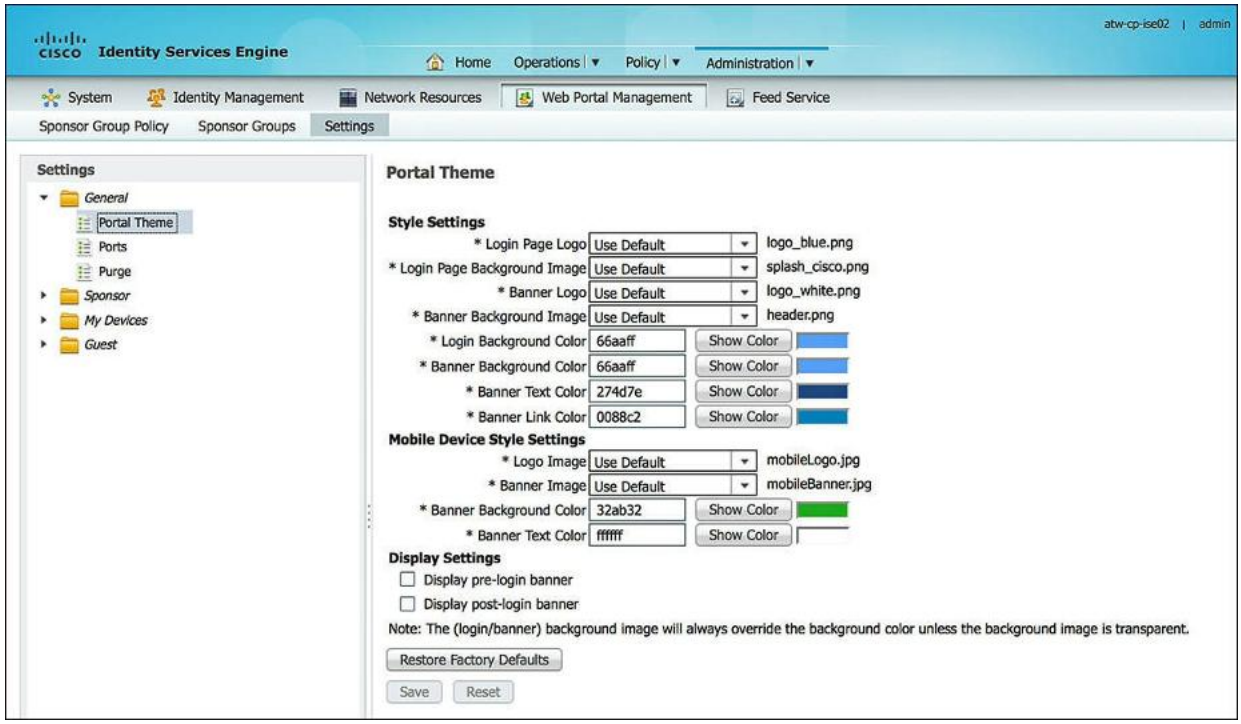


Figure 14-11 General web portal settings.

Port Numbers

As we discussed previously, ISE provides a number of portals to administrators, sponsors, and users alike, and many of these portals are central to the Guest Services Function. Starting with the base configuration location for the General Web Portal Settings, you will see an option for **Ports** (see Figure 14-12). These port settings enable you to modify the HTTPS ports that are used to access each portal on ISE. By default, these ports are all port 8443—except the Blacklist portal (which is port 8444). To change the port for a particular portal, simply change the port within the box to a value between 8000 and 8999 and then click **Save**.

https://t.me/learningnets

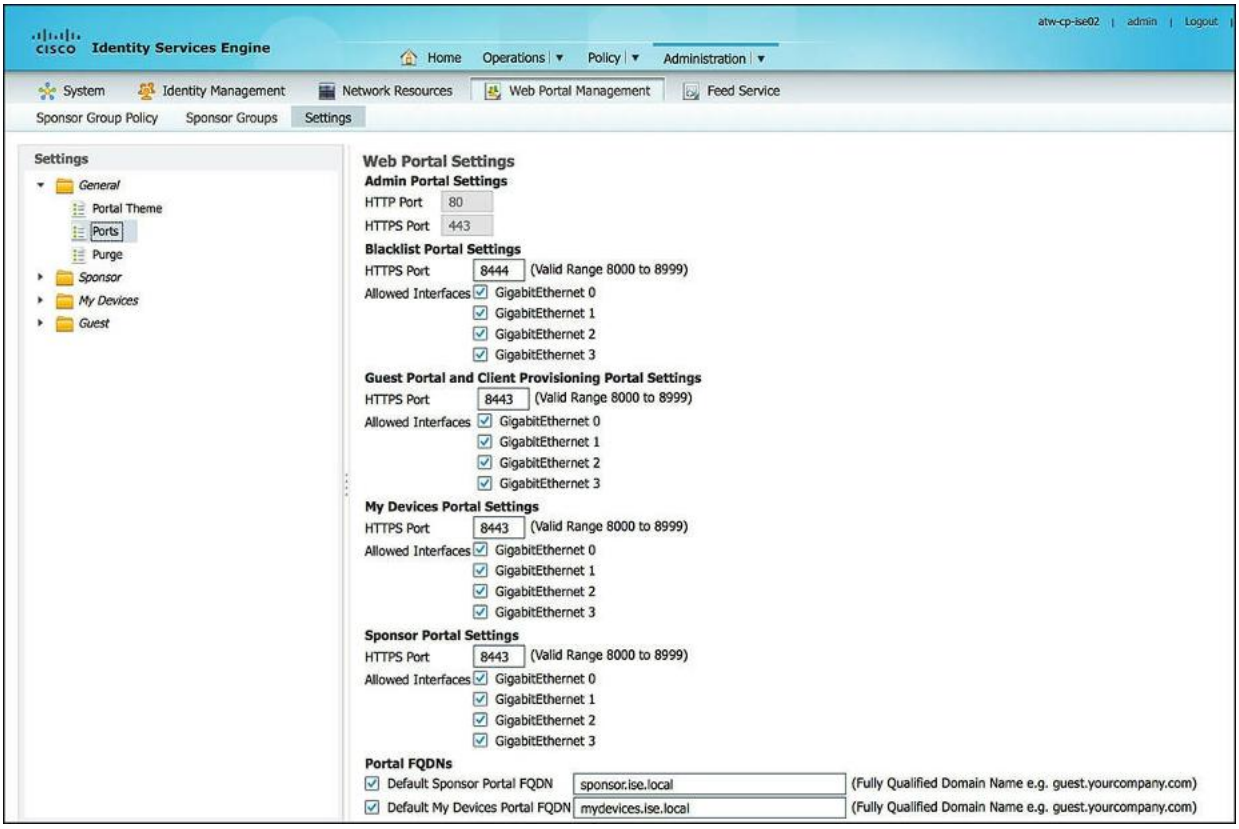


Figure 14-12 General Web Portal Settings – Ports.

Depending on your corporate security policy, you might choose to allow certain ports through your network by default or choose to block access to these ports from certain users or devices. By modifying the HTTPS port on which these portals reside, you have a greater amount of flexibility to provide certain users access to certain ports and portals while denying access to other users.

Again, the default configuration for these port settings will often work fine for many deployments.

Interfaces

While still under the **Ports** settings under the General Web Portal Settings (refer to Figure 14-12), you will also notice that you can enable or disable access to the individual portals on certain interfaces of the PSNs. Depending on the size and configuration of your ISE deployment, you can choose to allow access to the various portals on a number of interfaces. By enabling the portal on more interfaces, this can help you as you can now distribute the load across multiple interfaces.

From a security standpoint, for similar reasons to changing the HTTPS ports where the portals are hosted, you might choose to host a certain portal on a specific physical interface. For instance, you might elect to allow access to the sponsor portal out-of-band from the rest of your network. This would enable you to isolate the creation and management of guest accounts to computers on a management network.

Friendly Names

The Friendly Names configuration—again, still on the **Ports** configuration screen of the General Web Portal Settings (refer to Figure 14-12)—allows the administrator to simplify access to certain portals. As a sponsor, you will need to periodically access the sponsor portal to create guest user accounts. To make this account creation webpage easier to remember, you can configure a friendly name.

If you want to configure a Friendly Name for your sponsor portal, you can check the box beside Default Sponsor Portal FQDN and fill in the text box with the appropriate FQDN. For instance, a few good names for a sponsor portal might be sponsor.domain.com or guestaccounts.domain.com. This FQDN must be added to the relevant domain name service (DNS) servers for your organization, pointing to the appropriate

https://t.me/learningnets

interface Internet Protocol (IP) address(es) that are configured for the sponsor portal access. Also, because the sponsor portals are provided over HTTPS and leverage a X.509 certificate to authenticate the ISE portal server, it is imperative that you add the FQDN to the Subject Alternative Name of the portal's X.509 certificate to minimize security warning pop-ups on your browser. Otherwise, without a properly constructed X.509 certificate by a trusted certificate authority (local or well-known), you can continually get untrusted server warnings as you access the sponsor and other ISE portals.

Besides the sponsor portal, you can also enable the friendly name for the My Devices Portal of ISE. End users can manage their registered devices via this My Devices Portal. For the lay user, it can be difficult to remember a long URL. If the user has just lost his endpoint or had it stolen, he might need to browse to the portal via a kiosk or a co-worker's computer. An easy-to-remember URL will enable him to more quickly report his device as lost or stolen and minimize the risk of lost corporate intellectual property. A few examples of a My Devices FQDN would be mydevices.domain.com or byod.domain.com. The MyDevices portal is not one of the three portals required by Guest Services but is provided here for completeness.

The default URLs to access the common portals are listed in Table 14-2.



Portal	URL
Sponsor Portal	https://<PSN_IP_OR_FQDN>:8443/sponsorportal/
Guest Portal	https://<PSN_IP_OR_FQDN>:8443/guestportal/
My Devices Portal	https://<PSN_IP_OR_FQDN>:8443/mydevices/

Table 14-2 Default URLs to Access the Common Portals

Configuring the Sponsor Portal Policies

ISE provides a great deal of flexibility in determining who can be a sponsor of guest user accounts. As an ISE administrator, you can define which users, if any, on your network can create a guest user account. The users given the privilege of creating guest accounts will be referred to as sponsors. Depending on your corporate security policy, you can allow only certain users to create accounts for guests or can allow all full-time employees to create the guest accounts. As we'll discuss shortly, the level of access given to a sponsor can vary depending on any number of factors.

To configure the sponsor portal policies, start at **Administration > Web Portal Management > Settings > Sponsor**. Expanding the **Sponsor** folder, you will see **Authentication Source**. This authentication source is an identity source sequence that ISE uses to authenticate sponsors as they attempt to log in. To configure the identity source sequence, go to **Administration > Identity Management > Identity Source Sequences**. You can add a new identity source sequence or modify the existing Sponsor_Portal_Sequence (see Figure 14-13). Any sponsor that will be creating guest users must be authenticated using one of the identity stores within the chosen identity source sequence (i.e., Sponsor_Portal_Sequence). For more information or to review identity source sequences, please refer to Chapter 9, "Initial Configuration of Cisco ISE."

https://t.me/learningnets

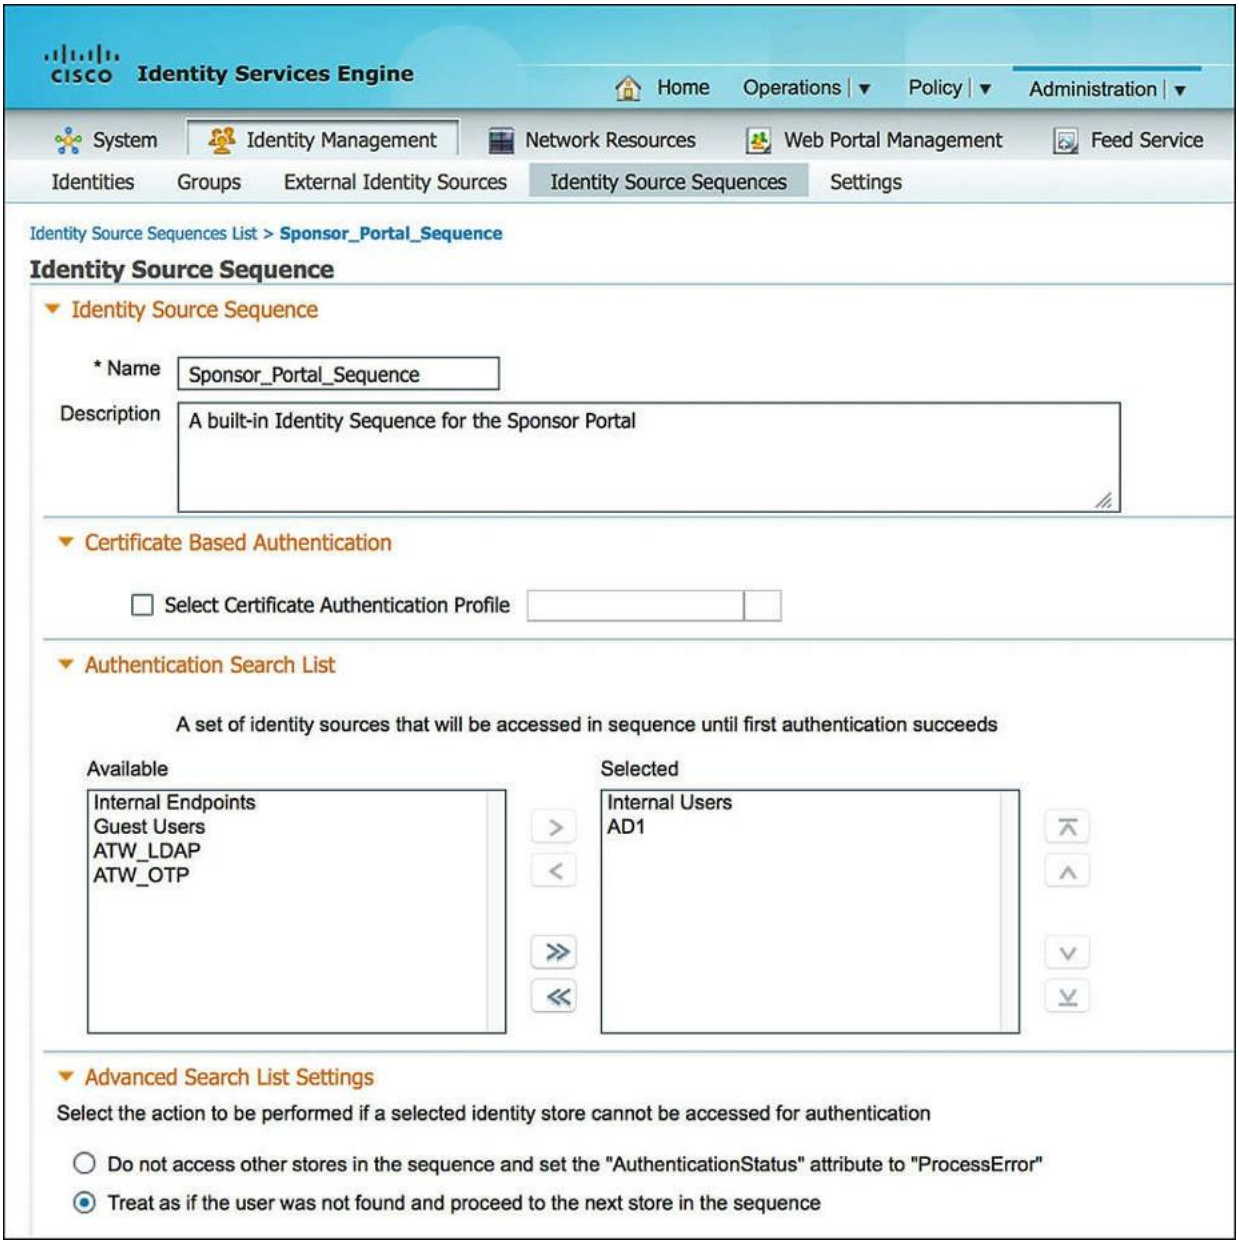


Figure 14-13 Sponsor_Portal_Sequence.

Sponsor Types

Sponsors can have varying levels of access, or capabilities, when it comes to creating guest users. ISE provides this flexibility in differentiating sponsors, as not all sponsors should be created equal. Earlier, we discussed that sponsors can be network administrators or full-time employees. Should all guest users created by either of these two groups of sponsors have the same level of access to network resources? Maybe, but probably not.

As a network administrator, you might have a better understanding of the level of security provided by the different guest user account types (which we'll discuss more shortly). Furthermore a full-time employee might not be well versed on networking and might have no idea what network security even means.

To further highlight the need for different levels of sponsors, the guest accounts created by the two sponsor groups can have different purposes. A guest account created by an employee might be to allow Internet access for a day, week, or month. However, a guest account created by a network administrator might need to be semi-permanent and have a much greater level of access. For instance, a network administrator may be responsible for providing long term network access to contract employees, giving them full access to all

https://t.me/learningnets

corporate resources for a year or more. This is not a responsibility that you would normally give to just any employee. By allowing your sponsors to create guest user accounts, they are allowing, potentially, unfamiliar, dangerous people onto your network. These guest users might be innocent or they might be nefarious hackers trying to steal all your corporate secrets and sell them to a competitor. A network administrator might have a few additional resources to verify his guest users or might receive additional training to help ensure that the guest user is legit and not a criminal.

Now that we have a background as to why we need differentiation between sponsors, let's take a look at how this is done. Within ISE, browse to **Administration > Web Portal Management > Settings > Sponsor Groups** (see Figure 14-14). Looking at the default configuration, you will notice that there are three different sponsor groups—SponsorAllAccounts, SponsorGroupGrpAccounts, and SponsorGroupOwnAccounts. These default groups and their associated policy provide the following sponsor access:

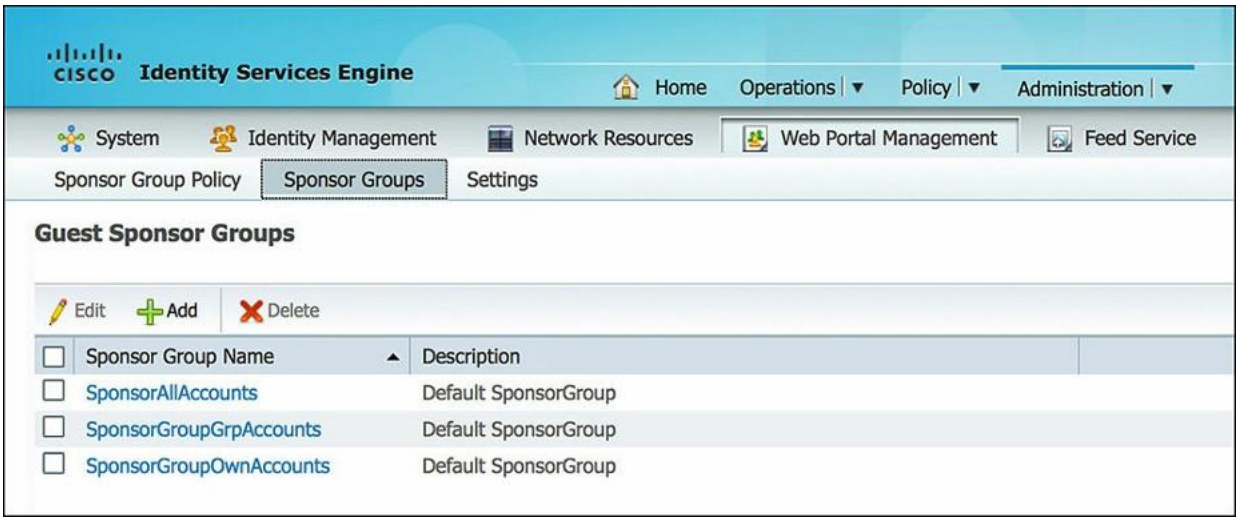


Figure 14-14 Sponsor groups.

- **SponsorAllAccounts**—These sponsors can manage all guest accounts on your ISE network.
- **SponsorGroupGrpAccounts**—These sponsors can manage only guest accounts created by sponsor users from the same sponsor group.
- **SponsorGroupOwnAccounts**—These sponsors can manage only those guest accounts that they have created.

Opening any of these three groups, you will see that there are several tabs to configure and many useful options to consider.

- **General**—This is the name and a brief description of the sponsor group.
- **Authorization Levels**—This tab highlights what the sponsor can and cannot do (see Figure 14-15). For instance, a sponsor might be allowed to create random accounts, create accounts from a CSV file, send emails/SMS, and view the guests' passwords, amongst other options.

https://t.me/learningnets

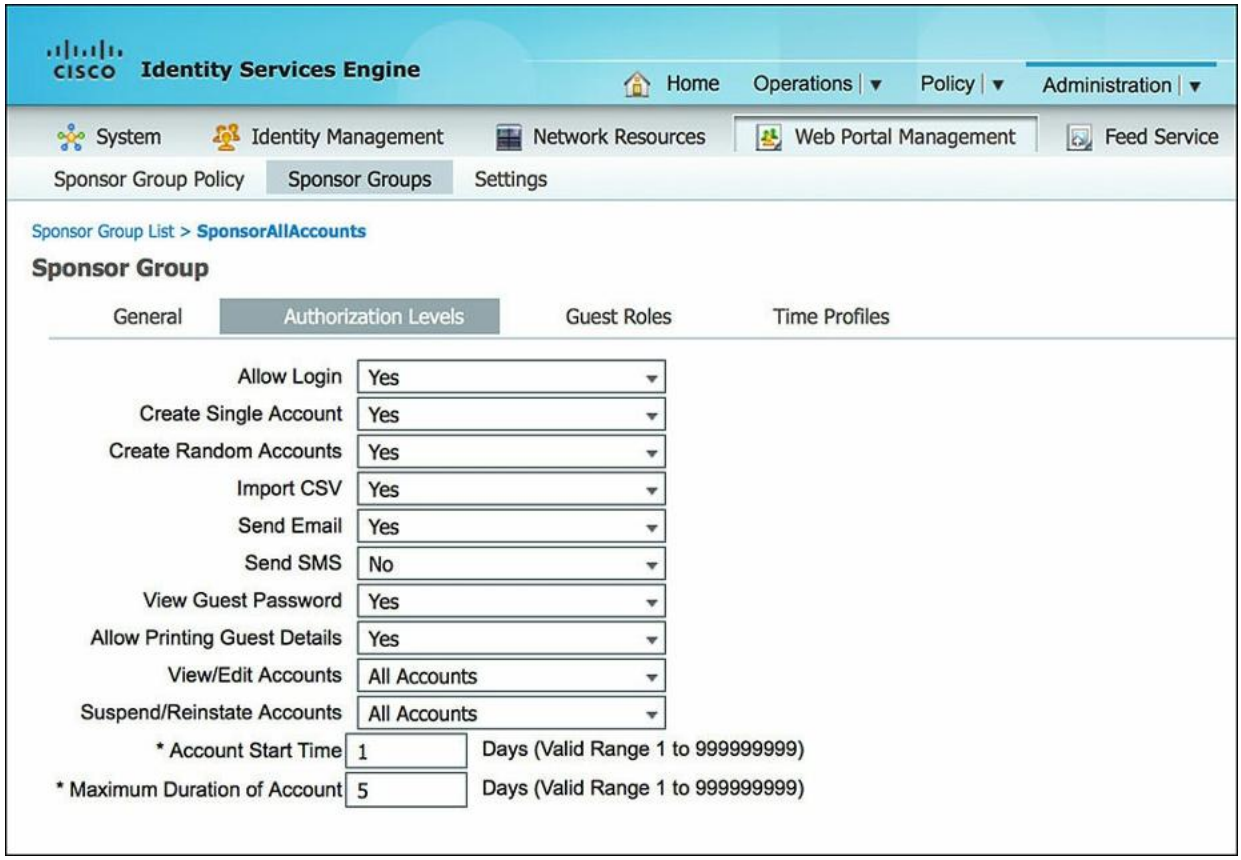


Figure 14-15 Authorization levels.

- **Guest Roles**—This tab shows what guest roles a sponsor can create (see Figure 14-16). These guest roles can eventually play a role in the level of access provided to the guest user. For instance, you might choose to create a guest role of Contractor (giving full access to the network on a temporary basis) or a Visitor (giving only Internet access to the guest user).



Figure 14-16 Guest roles.

- **Time Profiles**—This tab shows the length of time that a sponsor can allocate to a guest account (see Figure 14-17). These profiles are defined via **Administration > Web Portal Management > Settings > Guest > Time Profiles** (see Figure 14-18). Depending on the sponsor group, you can choose which time profiles are viable options.

https://t.me/learningnets

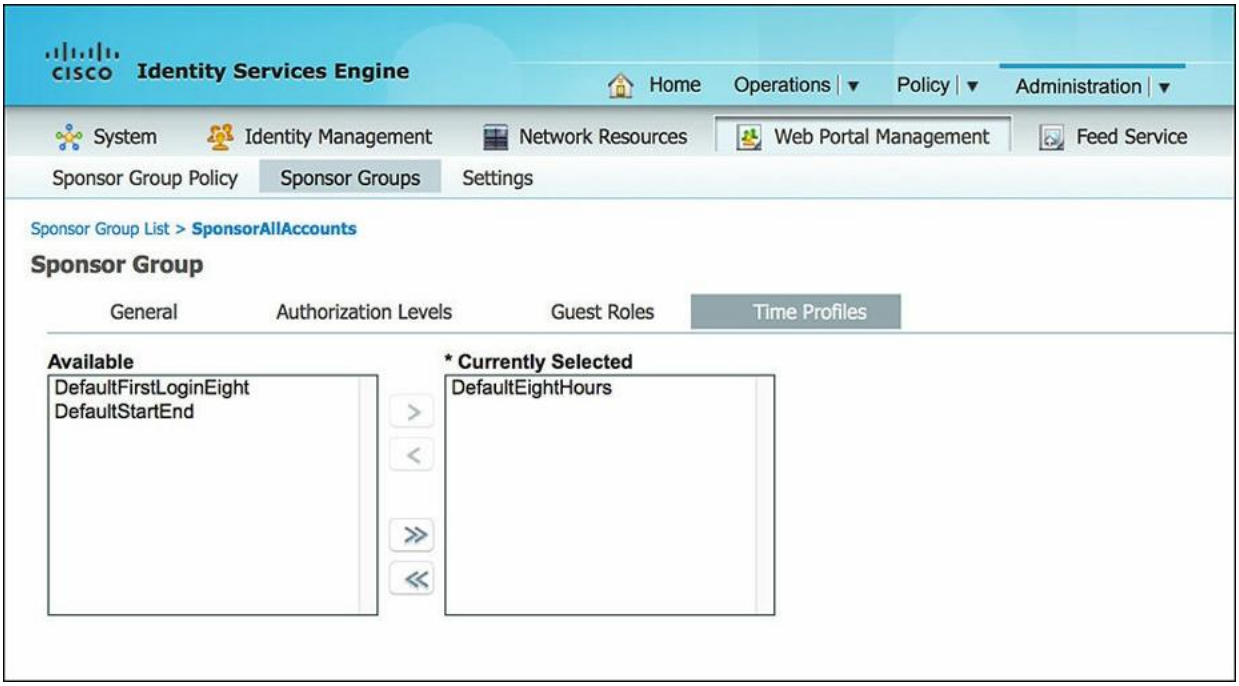


Figure 14-17 Time profiles.

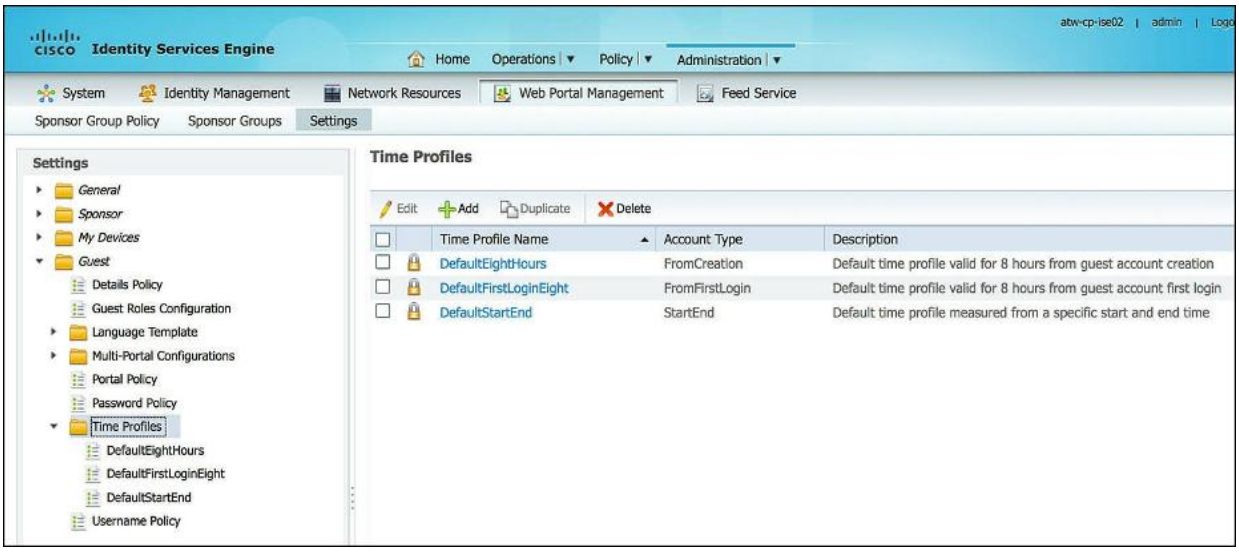


Figure 14-18 Time profiles definition.

Mapping Groups

Now that we have defined the sponsor groups and the groups' capabilities, we must map these groups to the sponsors who will create the guest user accounts. The manner in which this mapping happens is similar to the IF-THEN constructs we used in defining authentication and authorization policies.

To configure this sponsor-to-group mapping, go to **Administration > Web Portal Management > Settings > Sponsor Group Policy** (see Figure 14-19).

https://t.me/learningnets

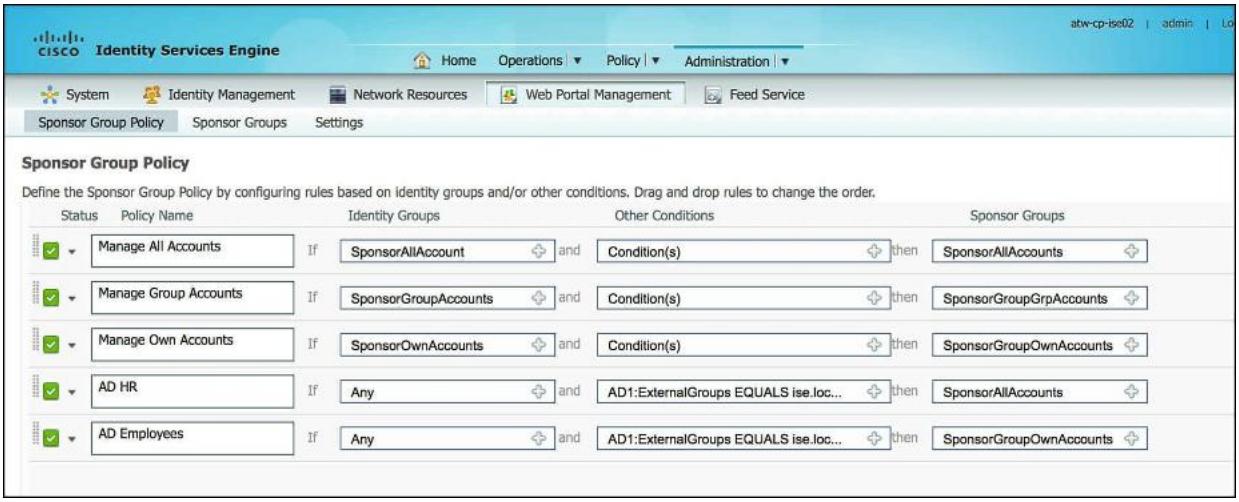


Figure 14-19 Sponsor group policy.

Looking at this policy, the first three rules are the default rules that are preconfigured on ISE. In each of these rules, you will see that the only condition governing the access provided to the sponsor is what ISE identity group that user is a part of. These three identity groups are also predefined on ISE. For any user to create user accounts via these first three rules, the usernames and passwords must be defined within ISE's User Identity store at **Administration > Identity Management > Identities > Users**. A user can be associated to a number of user groups.

Managing a separate sponsor database on ISE for every employee could quickly become cumbersome, especially if you already have an identity management system that is centralized for your corporation. For that reason, ISE can leverage an external identity management system, such as Microsoft Active Directory, to manage the users for you. The last two sponsor group policy rules have been added to demonstrate this concept.

The first added rule, named AD HR, has no ISE identity group assigned with the rule. The condition, however, references an externally defined AD group—AD1:ExternalGroups EQUALS ise.local/Users/HR. Because these users are members of our human resources (HR) department, we want to allow HR to create and manage all guest user accounts. This results in the HR user being assigned to the sponsor group SponsorAllAccounts. As we discussed previously, this high level of access could provide HR the ability to create or delete guest user accounts for contractors or others, regardless of which employee or network administrator created the account.

The second added rule, named AD Employees, similarly does not have an ISE identity group assigned with the rule. The condition does reference an externally defined AD group: AD1:ExternalGroups EQUALS ise.local/Users/Employees. As we don't want one employee to delete the user accounts of other employees, we have chosen to assign this group of users to the sponsor group SponsorGroupOwnAccounts. Therefore, the employee can edit or delete only those accounts that she personally created.

For these last two rules to be effective, allowing the designated employees to create the guest accounts, we'll need to ensure that AD1 is added to the Sponsor_Portal_Sequence, as demonstrated in Figure 14-20.

https://t.me/learningnets

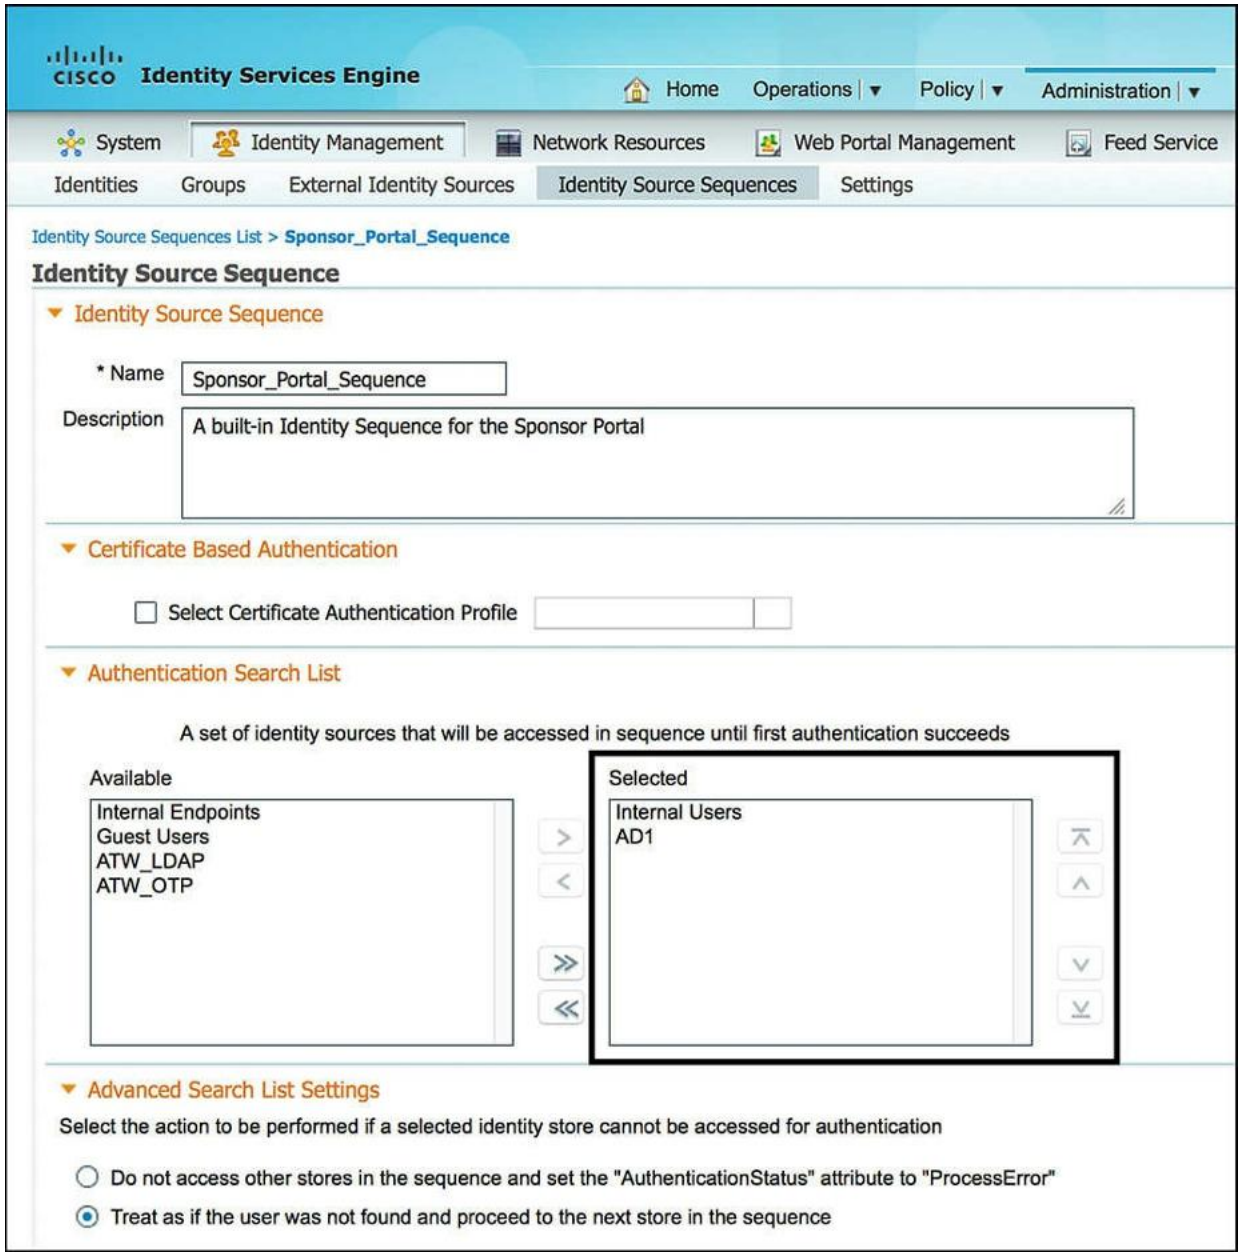


Figure 14-20 Add AD1 to Sponsor_Portal_Sequence.

Guest User Types

As a sponsor creates a guest user account, the user account can be assigned to any guest role or ISE identity group, as we previously mentioned in the Sponsor Types section. These guest roles, via the configuration **Administration > Web Portal Management > Settings > Guest > Guest Roles Configuration**, must be assigned to one of two guest user types—Guest or ActivatedGuest. By default, every ISE identity group is assigned to the guest user type of Guest.



What is the difference between Guest and ActivatedGuest? A guest user must authenticate via the Guest Services web portal. An ActivatedGuest can choose to authenticate to the network either via Guest Services web portal or use his credentials to authenticate to the network using 802.1X.

https://t.me/learningnets

Managing Guest Portals

As you leverage the Guests Services functionality of Cisco ISE, you might decide that you would like to customize the guest user experience, possibly adding company branding or modifying color schemes. If you want to add your own company branding, this can be done via the **Administration > Web Portal Management > Settings > General > Portal Theme**.

Portal Types

If you would rather customize a portal, you can do so by modifying an existing portal or by uploading complete HTML pages that are specific to your organization. This can be done via **Administration > Web Portal Management > Settings > Guest > Multi-Portal Configurations** (see Figure 14-21). As you browse to this configuration menu, you will see that there is a DefaultGuestPortal. The DefaultGuestPortal is the standard guest portal that comes predefined on ISE. You can choose to modify this guest portal or create a custom portal from scratch, modifying the configuration and pages as you wish.

https://t.me/learningnets

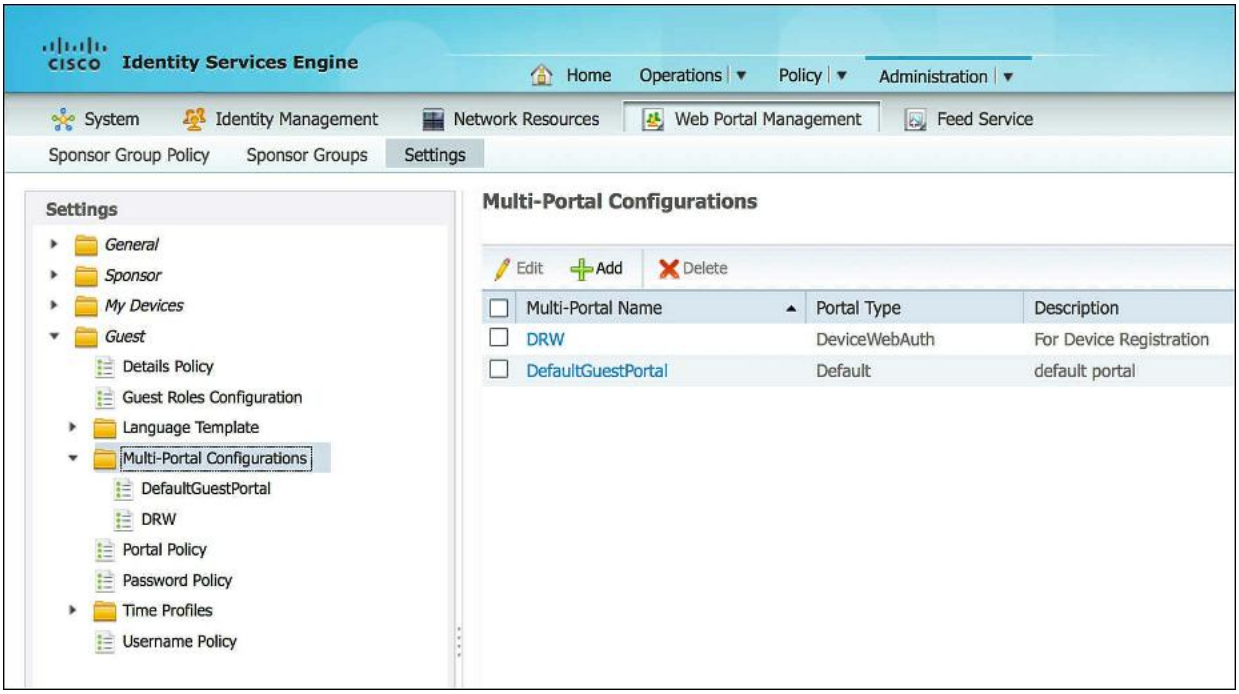


Figure 14-21 Multiportal configurations.

If you elect to create a custom portal from scratch, select **Add** from the **Multi-Portal Configurations** menu. As you click **Add**, you will be prompted to create a default portal or device web authorization portal —with each of these options allowing the administrator to select a customization template or to upload the HTML files. If you elect to upload HTML, you will be asked to map the HTML files with the specific ISE guest portal function, such as Login, Acceptable Use Policy (AUP), and so on.

Many of the customization options for a custom guest web portal are found on the Operations tab of the portal configuration. On this tab, you can choose whether the user will receive the Acceptable Use Policy (AUP) every time he logs in, only on the first login, or never receive the AUP at all. You can also allow a guest user to change his passwords as he wants, or you might choose to force a password change after the initial login. Equally important, a portal can allow a guest user to self-provision, self-service, and do device registration. Each of these options are unique to each custom portal you create, allowing a different visual experience as well as enabling only those features that are needed for each guest user scenario.



If you create a new guest portal, you will select this new portal as part of the authorization profile that is pushed to the endpoint upon successful authentication. We discuss the configuration of guest access policy later in this chapter.

Depending on your company's security policy, you also can choose to modify the guest user username or password policies. By going to **Administration > Web Portal Management > Settings > Guest > Password Policy** (see Figure 14-22) or **Administration > Web Portal Management > Settings > Guest > Username Policy** (see Figure 14-23) within the ISE admin portal, you will see that you can define minimum guest user password requirements or username requirements. These requirements would define the character sets and quantities of each character class that must be used in a password or username, accordingly.

https://t.me/learningnets

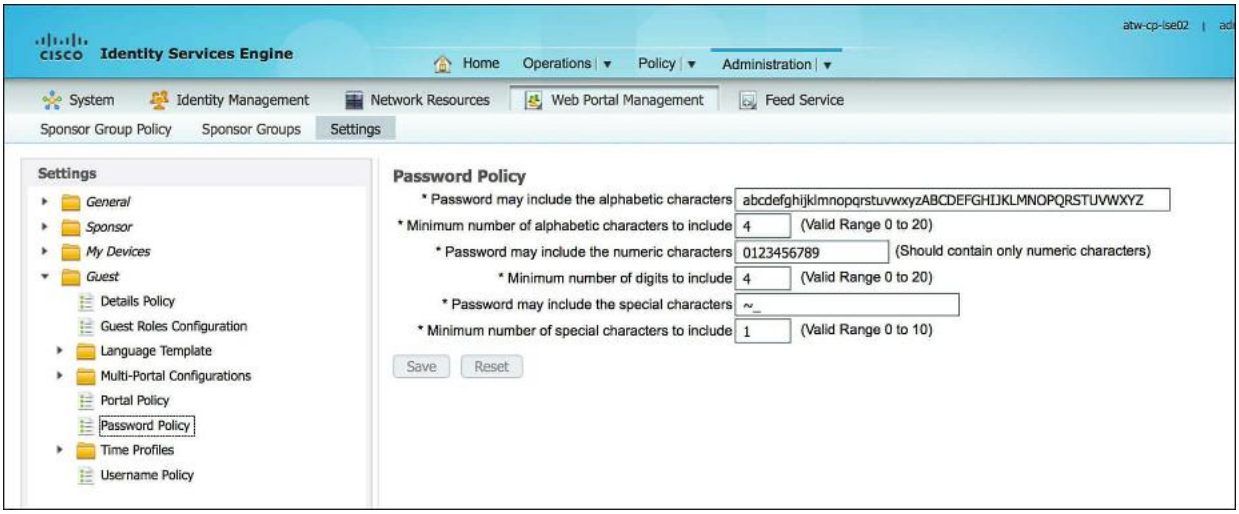


Figure 14-22 Guest user password policy.

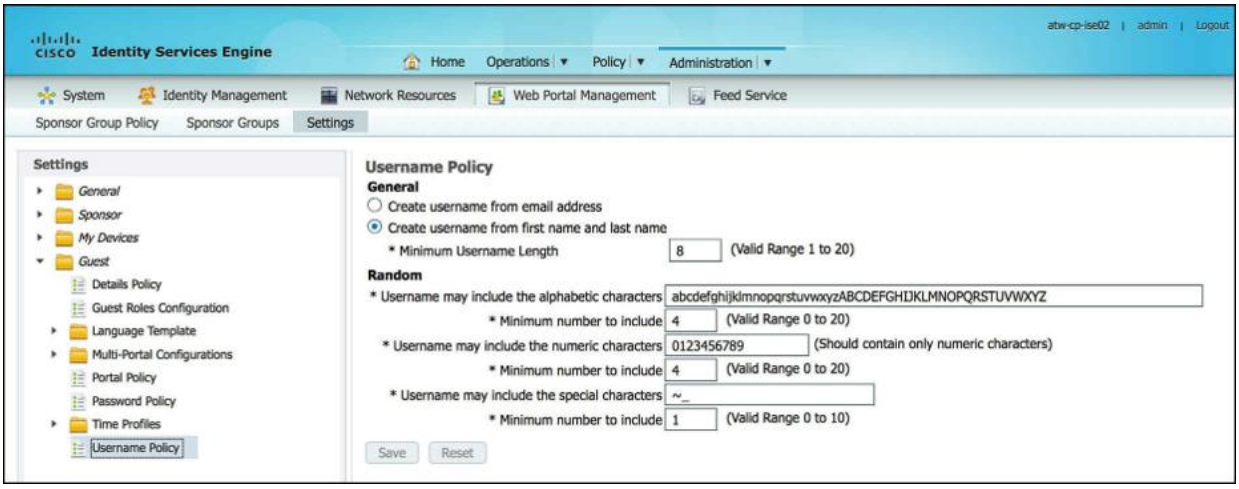


Figure 14-23 Guest user username policy.

Building Guest Authorization Policies

Before we configure guest authorization, let's take a look at what happens during the authentication process for Guest Services (see Figure 14-24). When we look at the authentication policy on ISE, we will see that the first authentication rule configured is named MAB—using the Internal Endpoints identity store. This MAB policy is leveraged to authenticate users for Guest Services.



Figure 14-24 Authentication policy.

This MAB authentication policy is used only if one of the following conditions is met:

https://t.me/learningnets

- Wired_MAB = Radius:Service-Type = Call Check AND Radius:NAS-Port-Type = Ethernet
- Wireless_MAB = Radius:Service-Type = Call Check AND Radius:NAS-Port-Type = Wireless - IEEE 802.11

Another authentication rule that deserves to be mentioned is the 802.1X authentication rule—Dot1X. The Dot1X uses the All_ID_Sources identity store. This Dot1X policy is met with either of the compound conditions:

- Wired_802.1X = Radius:Service-Type = Framed AND Radius:NAS-Port-Type = Ethernet)
- Wireless_802.1X = Radius:Service-Type = Framed AND Radius:NAS-Port-Type = Wireless—IEEE 802.11)

Finally, the Default authentication rule—chosen if none of the other policies match—uses the Internal Users identity store. Currently, the RA VPN policy does not have an impact on Guest Services.

As a NAD attempts to authenticate a user or an endpoint, there might be an authentication failure. With authentication, it is equally important for us to understand ISE's behavior during an authentication failure as it is to understand an authentication success. There are several ways that an authentication failure can occur. The authentication failures that ISE can explicitly handle are

- **Authentication Failed**—ISE received an explicit response that authentication has failed, for example due to bad credentials, a disabled user, and so on. The default course of action is reject.
- **User not found**—No such user was found in any of the identity databases. The default course of action is reject.
- **Process failed**—Unable to access the identity database or databases. The default course of action is to drop.

Within ISE, you have the ability to choose how to react to these authentication failures. This functionality enables you to continue to the authorization step, despite a failure during the authentication process. The possible courses of action that ISE can take in case of an authentication failure are

- **Reject**—A reject response is sent.
- **Drop**—No response is sent.
- **Continue**—Cisco ISE continues with the authorization policy.



Even when you choose the Continue option, there might be restrictions on the protocol whereby Cisco ISE cannot truly continue. However, Continue is a viable option for PAP/ASCII, EAP-TLS, and MAC Authentication Bypass (MAB). For all other authentication protocols—those that cannot truly continue—an ACCESS_REJECT response is sent if there is an authentication failure or if the user or host is not found in the identity database. Otherwise, it drops the request and sends no response if the process fails. We discuss how we will leverage the authentication failures within the WebAuth process in more detail shortly.

Now that we understand the current authentication configuration on ISE, let's take a look at WebAuth from the standpoint of the endpoint device. Because guest users can use several devices with different operating systems and capabilities, we cannot assume that all these guest users' devices can support 802.1X. However, it is reasonable to expect that all guest devices that need network access have a web browser that can be used for authentication.

As the end user attempts to access the network, the NAD, as we configured it in Chapter 12, "Implement Wired and Wireless Authentication," attempts to authenticate the user by trying 802.1X first and then MAB next and forwarding this information back to ISE. As the guest's endpoint does not have an 802.1X supplicant, credentials, or configuration to authenticate to the network, the NAD's attempts to authenticate the endpoint using 802.1X times out. The NAD then tries to authenticate the user and endpoint via MAB—the second method on the NAD's authentication list configuration.

https://t.me/learningnets

As we discussed in Chapter 6, "Introduction to Advanced Concepts," and Chapter 12, MAB requires that the endpoint's MAC address already exists in the endpoint identity store. Because this device is likely unknown to the network, this too will fail authentication. In this case, we are actually anticipating the authentication failure and will structure our authentication policy accordingly to handle this failure. To accommodate this anticipated failure and allow the authorization policy to provide the appropriate security policy, we will set the authentication failure reaction for **If user not found** and **If process failed** to Continue (see Figure 14-25).

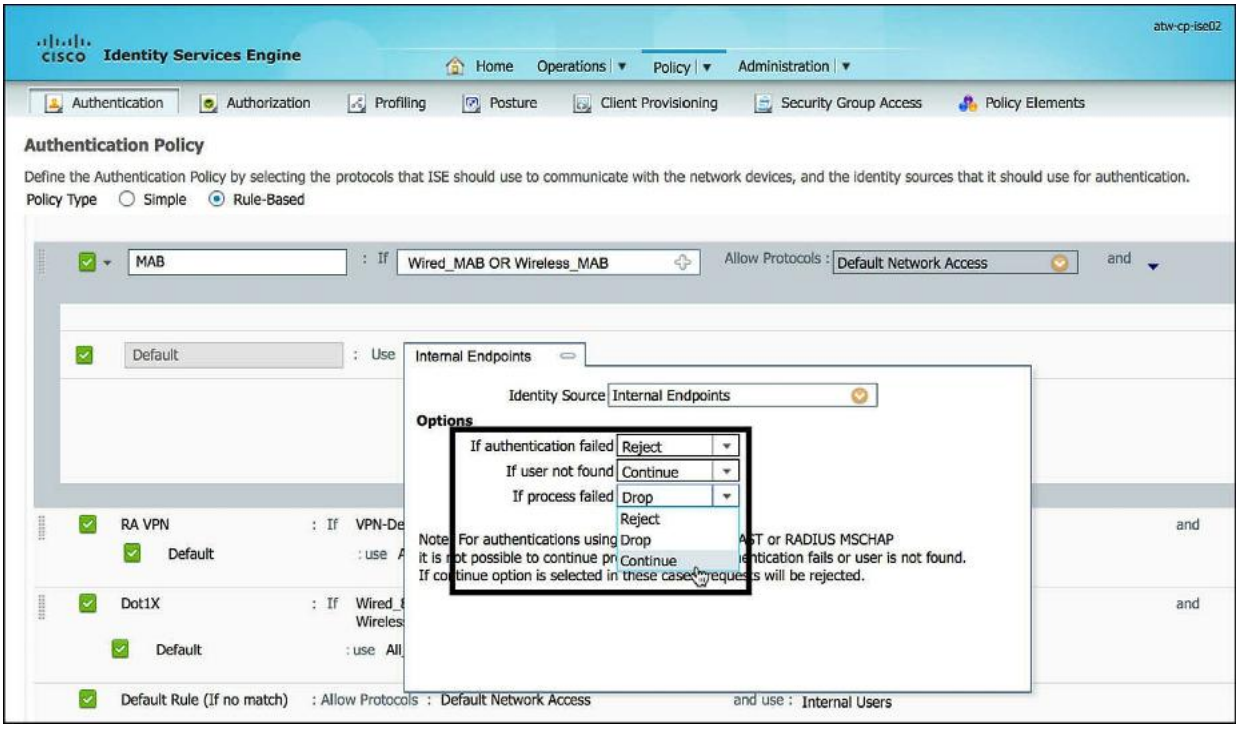


Figure 14-25 Changing authentication failure policies If User Not Found and If Process Failed to Continue.

Let's walk through a play-by-play of what will happen with MAB, considering the changes we made on ISE. With the policy configured the way it is, anytime ISE attempts to authenticate an endpoint using MAB, one of two scenarios will play out:

- The endpoint's MAC address is present in the Internal Endpoints or other identity store(s)—depending on the acceptable identity stores per the MAB authentication policy.
- The endpoint's MAC address is NOT present in any of the elected endpoint identity store(s).

In the first scenario, the endpoint's MAC address is sent as the username for the Radius:Service-Type = Call Check ACCESS_REQUEST. The endpoint can be a member of an explicit endpoint identity list, such as a printer, an IP phone, or a copier MAB list, causing the authentication to pass, moving onto the authorization policy. The resulting authorization policy for an explicit MAB authentication success can be structured to allow the printer, the IP phone, copier, or another device access to a finite list of specific resources. Therefore, the final level of access is determined by the authorization policy on ISE.

In the second scenario, the endpoint's MAC address is still sent as the username for the Radius:Service-Type = Call Check ACCESS_REQUEST. However, the endpoint will not be found in any of the configured endpoint identity stores. By default, this would result in an immediate authentication failure and an ACCESS_REJECT. With the **If user not found** being changed from Reject to Continue, the endpoint will still be allowed to continue with the authorization step. It is our responsibility, by our changing **If user not found** to Continue, to appropriately configure the authorization policy and only allow the device an appropriate level of access. We do not want this change to open the network to any unwanted network traffic.

Ultimately, with this new configuration, anytime the MAB rule is hit, ISE will "pass" authentication and all

https://t.me/learningnets

endpoints will move onto the authorization step. Theoretically, you could choose to configure the Default authentication rule's **If user not found** setting to be Continue and it could serve the same purpose as this MAB rule. However, by explicitly mentioning MAB within the authentication rules, it provides the network administrator a more granular approach to handle a larger number of scenarios.

It is a best practice to handle all expected scenarios explicitly and fail closed otherwise. For this reason, it is best to set the Default authentication rule to Deny, resulting in ACCESS_REJECT for authentication failures or user not found scenarios, dropping the authentication request if the process fails.

We've now established that we will always move onto the authorization step when using MAB to accomplish Guest Services. Within the authorization step, we can now differentiate those scenarios where the device was found in the endpoint identity stores and those scenarios where the device was not found. This is usually handled by appropriately managing the order of the authorization rules or by electing the conditions of the authorization rules to set apart one MAB authorization from the others.

As was the case with the authentication process, it is a best practice to handle all expected authorization scenarios explicitly and then have a default rule whereby this default rule is as secure as reasonably possible. Often, this final rule will be a RADIUS ACCESS_REJECT response, such as the DenyAccess authorization profile, that denies all traffic from the endpoint.

Figure 14-26 shows this approach. Within the authorization policy, the identity groups are provided in bold typeface. Whenever an identity group is listed, it can be accompanied with other authorization conditions or by itself. In Figure 14-26, you will see that the **DRW-Endpoints** will be given the authorization profile of Internet_Only AND GUEST. The only condition present in this authorization rule is an endpoint identity group, implying that MAB alone is being used to authenticate this endpoint. The goal of this rule is to check the identity list **DRW-Endpoints** and, if the endpoint's MAC address is present, provide the endpoint with the access given in the authorization profile Internet_Only and the SGT policy given by GUEST.

https://t.me/learningnets

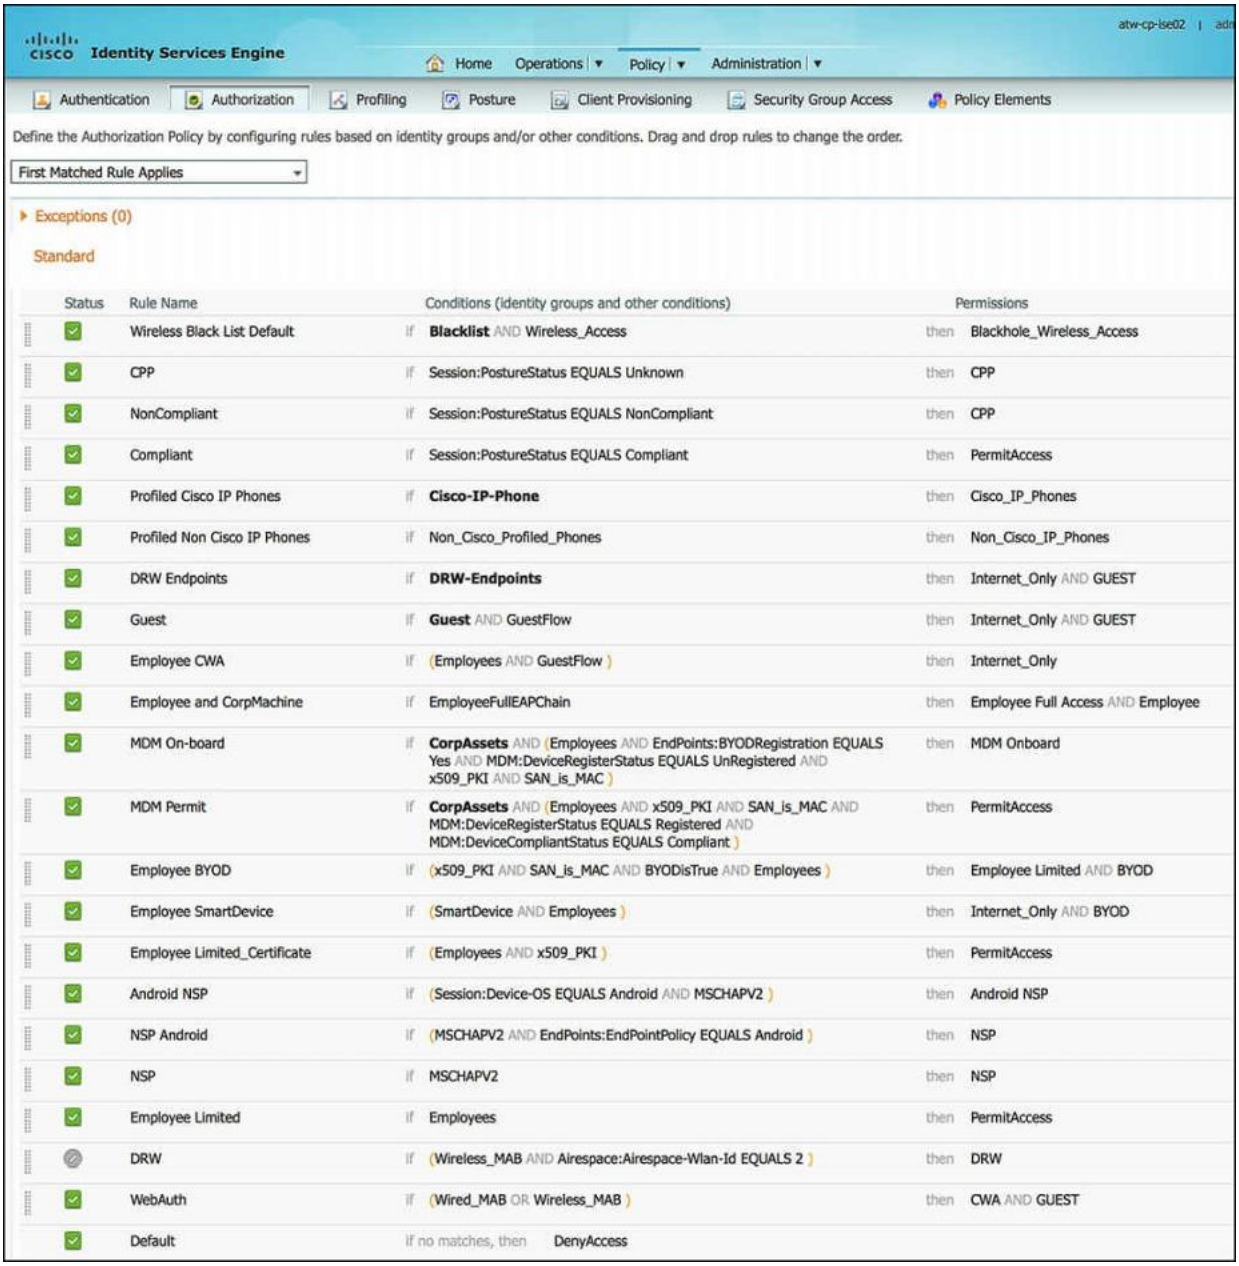


Figure 14-26 Proper ordering of authorization rules for MAB successes and failures.

As we look further down the authorization policy, we will see that there is an explicit mention of Wireless_MAB and Wired_MAB in the authorization rule WebAuth, just before the Default rule. Looking at this rule, the conditions column contains Wired_MAB OR Wireless_MAB (please notice the OR keyword). This rule will take effect only if

- All other rules before it don't match.
- The endpoint is using MAB to authenticate to the network (either wired or wireless).

Again, the Wired_MAB and Wireless_MAB conditions are authentication conditions that are configured as

- Wired_MAB = Radius:Service-Type = Call Check AND Radius:NAS-Port-Type = Ethernet
- Wireless_MAB = Radius:Service-Type = Call Check AND Radius:NAS-Port-Type = Wireless—IEEE 802.11

If either of these conditions is met, the endpoint will be granted the level of access contained in the authorization profile CWA and SGA Policy GUEST. This will provide an initial authorization that provides basic network access and allows the guest user an opportunity to authenticate via a web-based interface. We

https://t.me/learningnets

discuss these authorization profiles and how this web authentication process happens shortly.

A second option besides the one we just mentioned is to use a specific RADIUS condition within ISE designed for this process flow—Network Access:UseCase EQUALS Host Lookup. This rule implies that a MAB process is underway (that is, a host lookup has happened) and has failed to hit all previous and explicit MAB authorization rules within the authorization policy. It accomplishes exactly the same task as our original, configured rule. Depending on your company's security policy and required use cases, your policy and Guest Services approach can take a slightly different approach yet end up with the same result. ISE is extremely flexible and can accommodate several scenarios.

Usually, the Guest Services' failed MAB rule is placed just above the Default authorization rule. As mentioned previously, you can design your rule set, authentication and-authorization both, to allow a user to continue connecting to the network following a failed authentication and then eventually elect the Default authorization rule. This would be just one more way to accommodate Guest Services access. Again, ISE is extremely flexible and you can use several approaches to accomplish the same goal.



Please use extreme caution when having a final Default rule that does not provide a RADIUS ACCESS_REJECT response because this could provide a nefarious user access to your network that you might not completely intend.

Now that the user has been authorized to access the network, let's take a look at what this authorization means to the endpoint and end user. As an endpoint fails to match a MAB endpoint identity group and hits the WebAuth authorization rule, the endpoint—and therefore the user—will be assigned the authorization profile CWA and SGA Security Group GUEST. Let's first take a look at the CWA authorization profile (Figure 14-27 through Figure 14-30).

https://t.me/learningnets

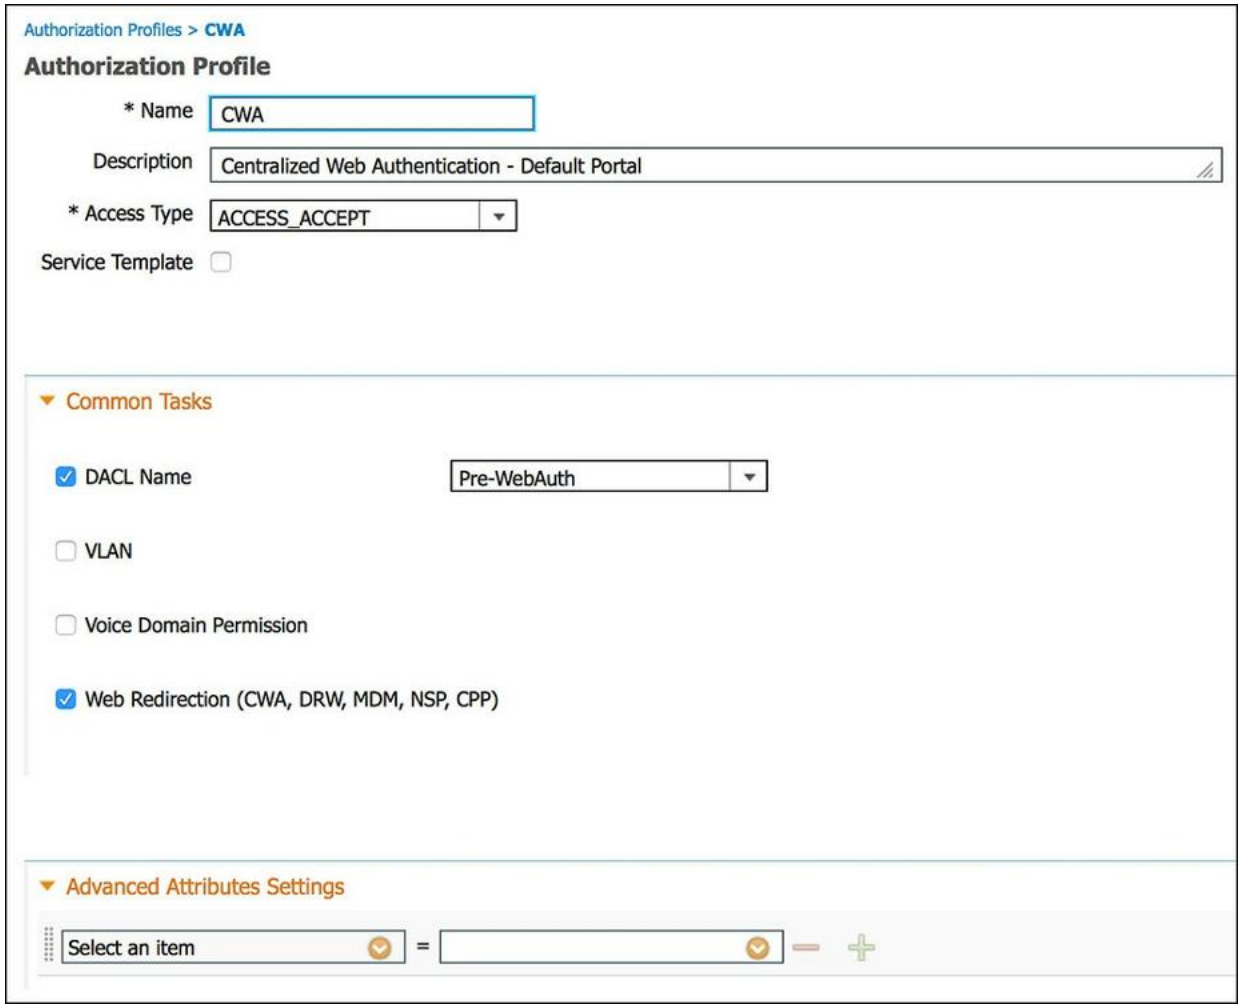


Figure 14-27 CWA authorization profile, part 1.

https://t.me/learningnets

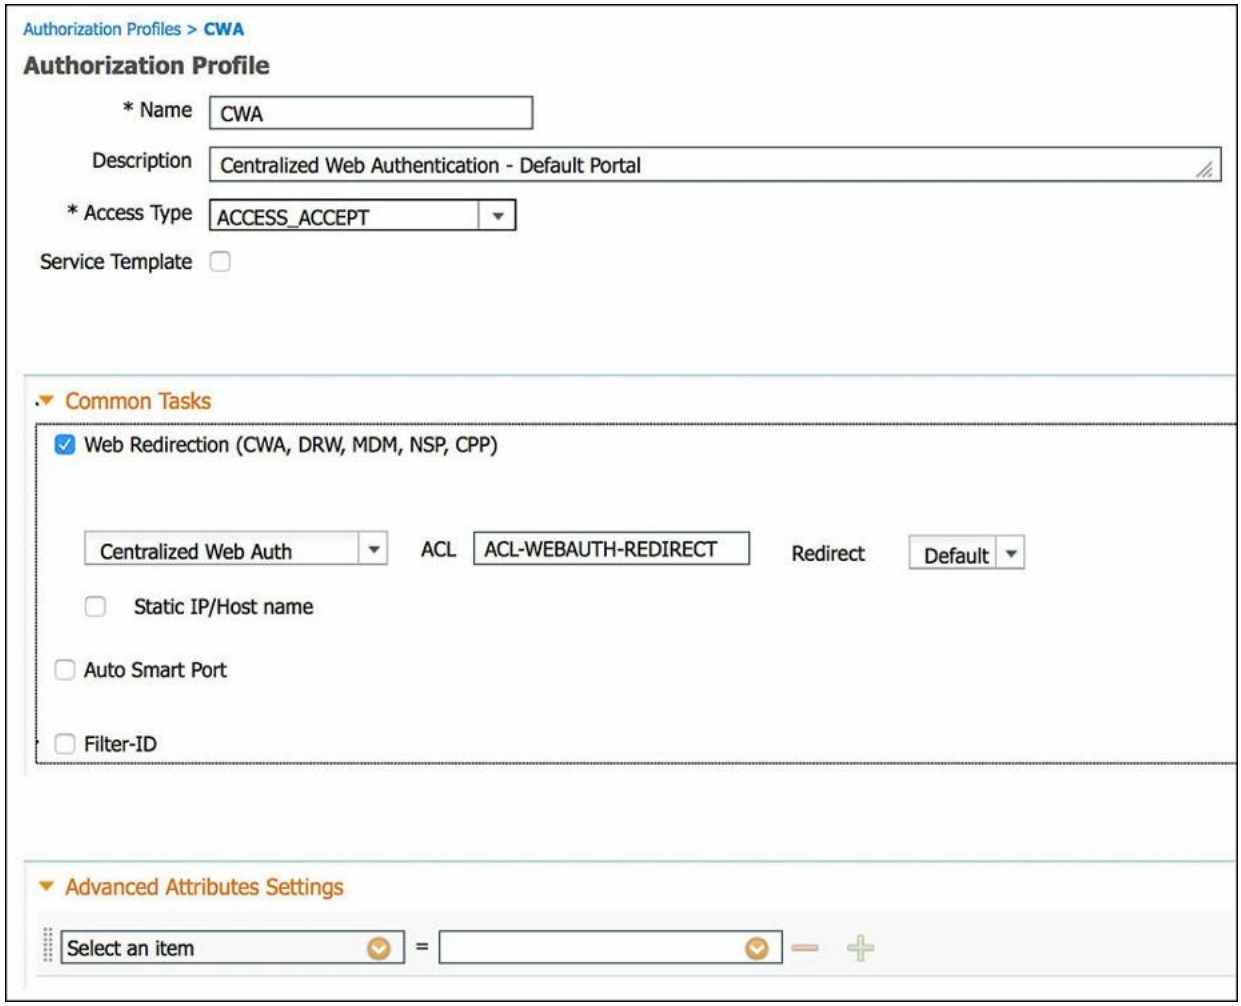


Figure 14-28 CWA authorization profile, part 2.

https://t.me/learningnets

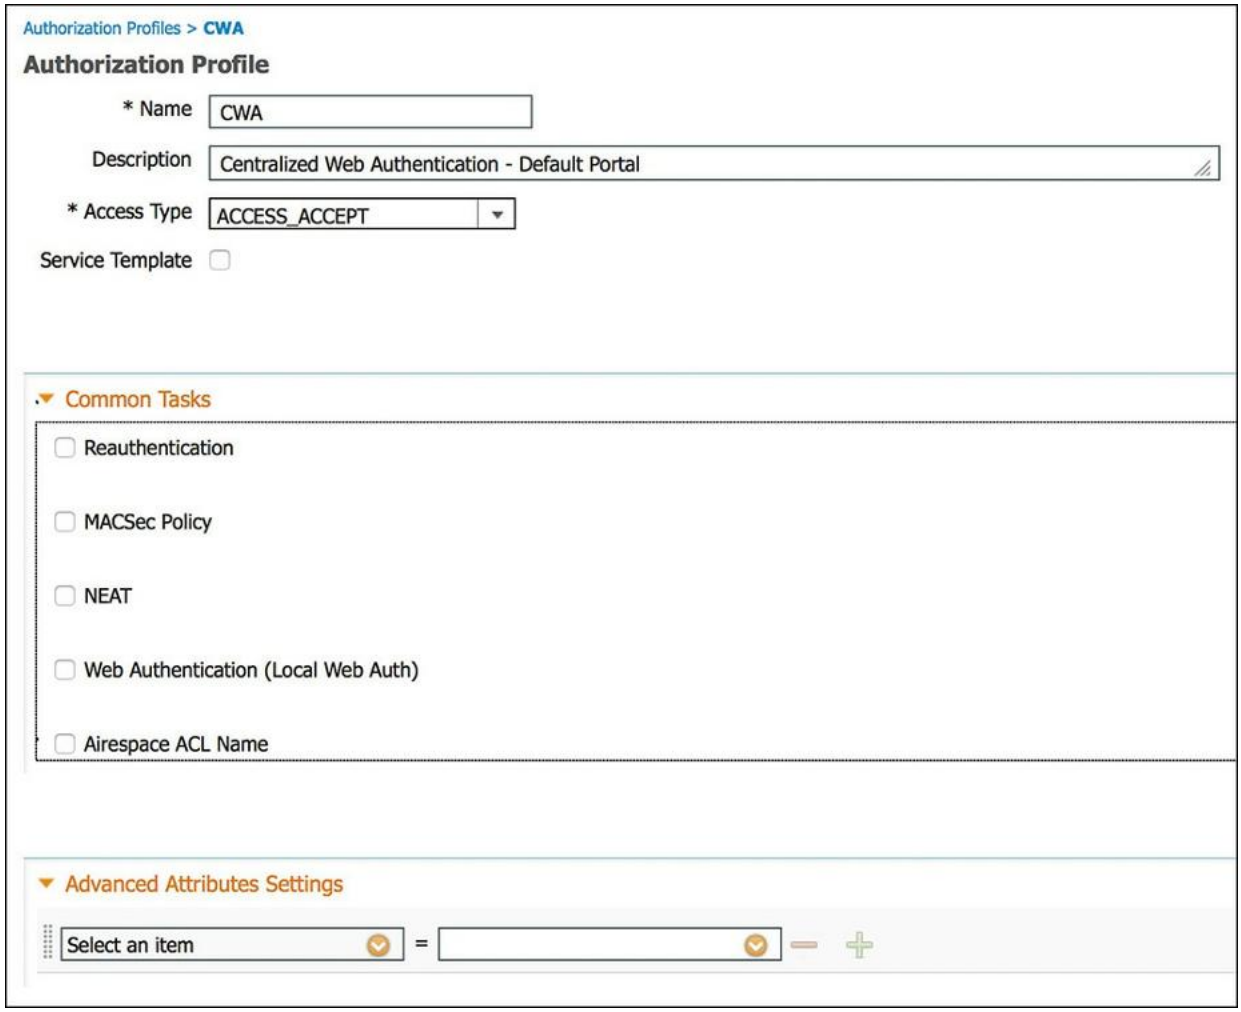


Figure 14-29 CWA authorization profile, part 3.

https://t.me/learningnets

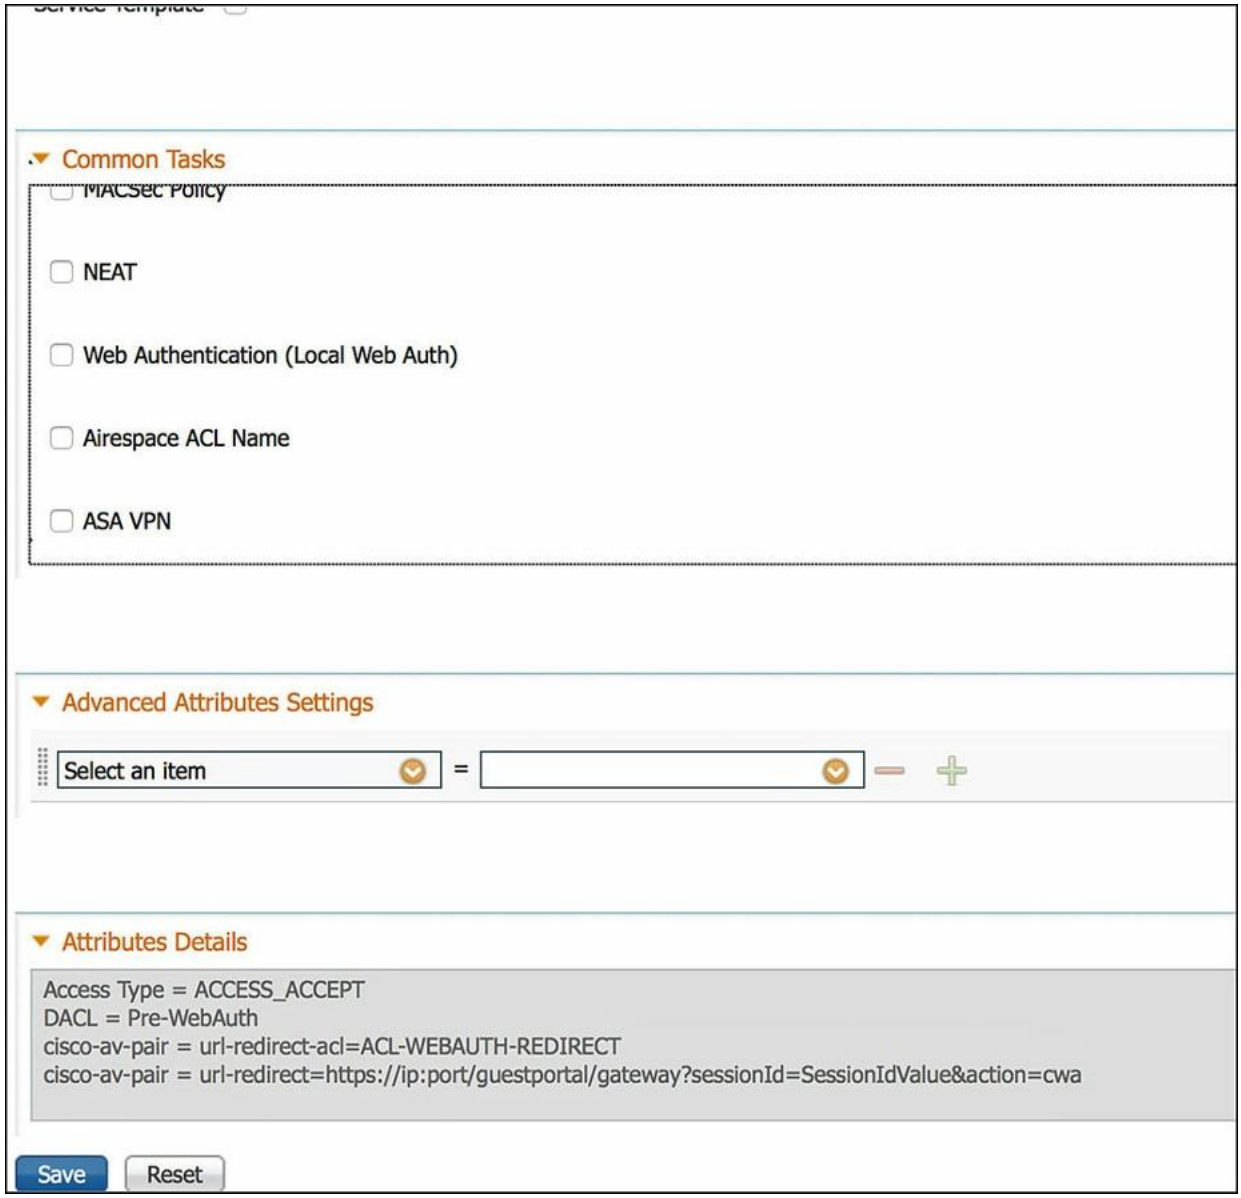


Figure 14-30 CWA authorization profile, part 4.

Looking at Figure 14-27, you will first notice the name and description followed by an access type of ACCESS_ACCEPT. This ACCESS_ACCEPT implies that some level of access will be provided to the endpoint. This access can be extremely open or extremely restrictive, but the result is access nonetheless.

Further down in Figure 14-27 is the Common Tasks section. This section, as you learned in earlier chapters, helps simplify the configuration for those security elements that are used often by network administrators. For the CWA authorization profile, we will leverage several of these common tasks.

The first common task we will leverage is the DACL name (see Figure 14-27). DACL, an acronym for Downloadable Access Control List, is an access list, most often used in a wired scenario, that is pushed from ISE down to the NAD. The DACL that we will push to the NAD for the CWA rule is named Pre-WebAuth. This ACL is defined on ISE via the **Policy > Policy Elements > Results > Authorization > Downloadable ACLs** (see Figure 14-31).

https://t.me/learningnets

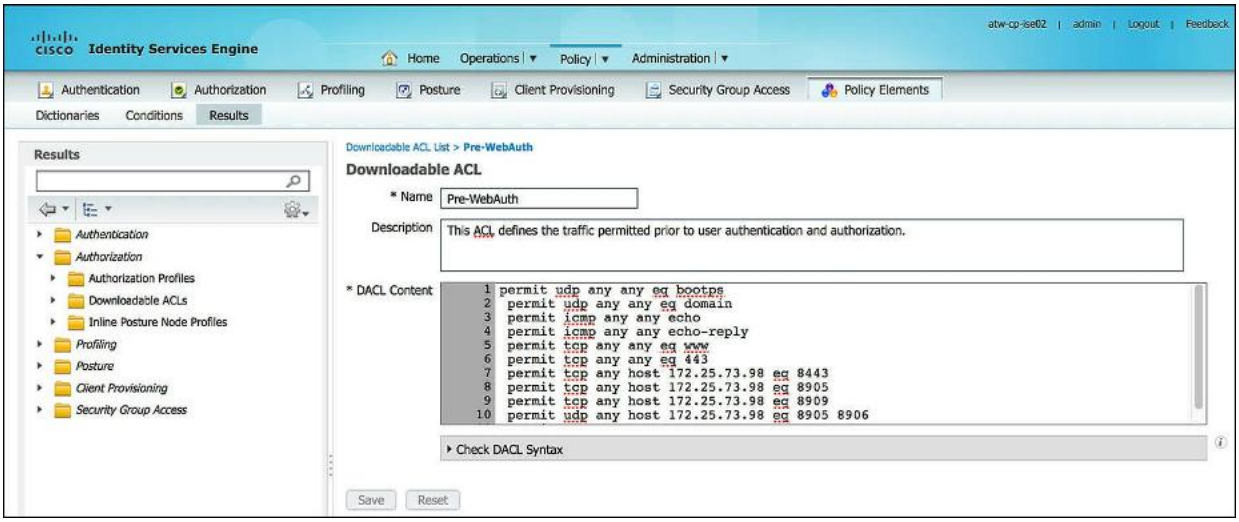


Figure 14-31 Pre-WebAuth DACL.

As the contents of this ACL continue beyond the bottom of the textbox on Figure 14-31, the complete contents of this Pre-WebAuth DACL, commented for your understanding, are shown in Example 14-1:

Example 14-1 Pre-WebAuth DACL

Click here to view code image

```
! Traffic to a DHCP server
permit udp any any eq bootps
! Traffic to a DNS server
permit udp any any eq domain
! Ping traffic and response for diagnostics purposes
permit icmp any any echo
permit icmp any any echo-reply
! Traffic to HTTP websites—must be permitted before being redirected
permit tcp any any eq www
! Traffic to HTTPS websites—must be permitted before being redirected
permit tcp any any eq 443
! Traffic to ISE portals, posturing, posture remediation, and provisioning
! functions
permit tcp any host 172.25.73.98 eq 8443
permit tcp any host 172.25.73.98 eq 8905
permit tcp any host 172.25.73.98 eq 8909
permit udp any host 172.25.73.98 eq 8905 8906
permit udp any host 172.25.73.98 eq 8909
```

As the user is assigned the CWA authorization profile, this Pre-WebAuth ACL is pushed to the NAD, line by line, in its entirety, using the RADIUS protocol. This ACL will be instantiated on the NAD and will restrict to which devices the endpoint will have access.

The next common task we'll leverage for the CWA authorization profile is Web Redirection (see Figure 14-28). Web Redirection can serve a number of purposes depending on which parameters are selected. In the case of Guest Services, we will be leveraging the Centralized Web Auth web redirection configuration. This is the reason we chose to name this authorization profile CWA—it contains a Centralized Web Authentication portal component. As the name suggests, this will redirect the user to a web portal that is hosted on ISE to authenticate the user.

One of the parameters required by the Centralized Web Auth function is an ACL via the ACL drop-down. The ACL we will choose here is a named ACL versus the downloadable ACL we just mentioned. A named ACL is statically defined on the NAD; ISE will simply invoke the ACL's use by reference alone. In this case, the ACL named ACL-WEBAUTH-REDIRECT is defined on the NAD and referenced as part of the Web

https://t.me/learningnets

Redirection configuration. The function of this ACL is to define what traffic will trigger the redirect. If traffic generated by the endpoint matches this ACL, the endpoint will be sent to the web portal that is selected—in this case, the Centralized Web Auth portal. The contents of this ACL were defined in Chapter 12.

The final parameter that we can configure as part of the Centralized Web Auth web redirection is the Redirect parameter. This Redirect drop-down enables the network administrator to select either the default Centralized Web Authentication portal or a custom portal, allowing the network administrator to customize the user experience as well as invoke specific portal rules based on the company's security policy. By default, the Default redirect is chosen—the default Centralized Web Authentication portal. However, if you have configured a custom portal, modifying any of the security parameters or portal behavior, you can redirect the end user to this custom portal by selecting the custom portal from the drop-down.

The rest of the common tasks are left as their defaults (unselected) for the CWA authorization profile (see Figure 14-29 and Figure 14-30). You can see the RADIUS Attribute-Value pairs for the totality of the configured parameters for the CWA authorization profile by looking at the box labeled Attributes Details in Figure 14-30. These attributes are sent to the NAD whenever the CWA authorization is hit.

We'll now take a look at the SGA Security Group GUEST. Looking at Figure 14-32, you will see that the GUEST SGA Security Group is assigned the Security Group Tag (SGT) of 8.

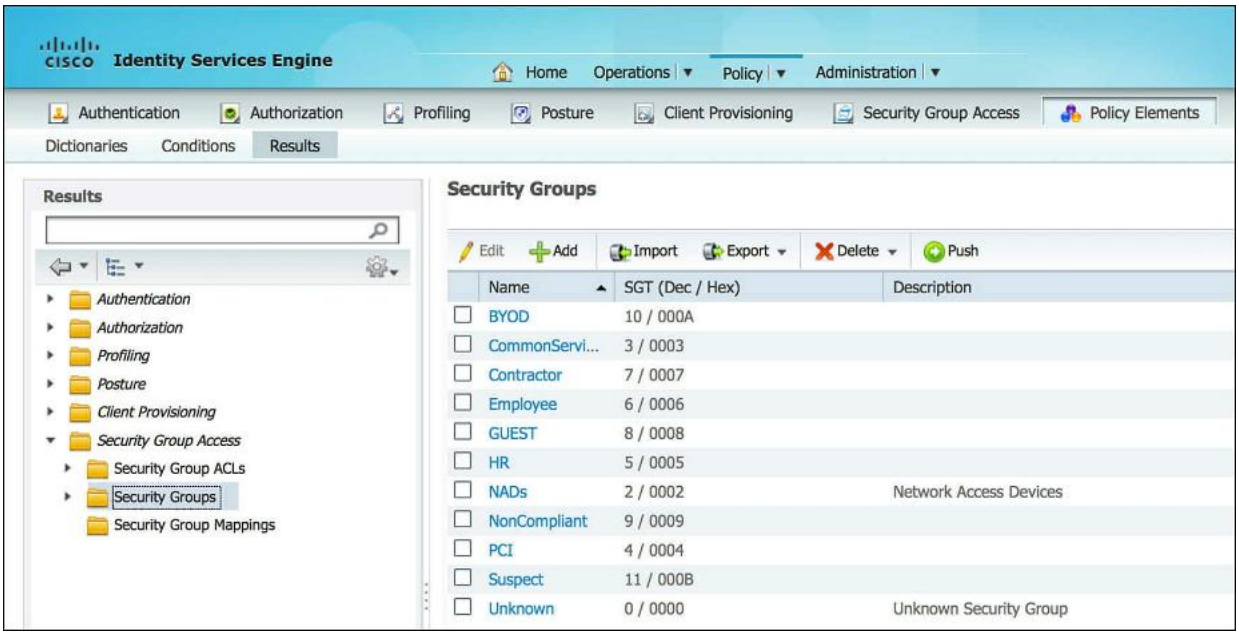


Figure 14-32 SGA security group.

With the SGT assignment of the GUEST policy, we really don't have the whole picture. We need to see the Security Group Access Egress Policy Matrix (see Figure 14-33). In this matrix, you can see that other SGTs are defined for other devices on the network. By looking at the SGT assigned to GUEST (that is, 8) on the left side of this table—the Source Tag—you can determine which destination tag(s) is a viable destination for the GUEST SGT. In Figure 14-33, you can see that SGT 8 is explicitly denied access to Contractors (SGT 7), Employees (SGT 6), Guests (SGT 8), and HR (SGT 5). Otherwise, by explicit configuration or by the default rule at the bottom, the GUEST SGT is allowed access to all other resources on the network.

https://t.me/learningnets

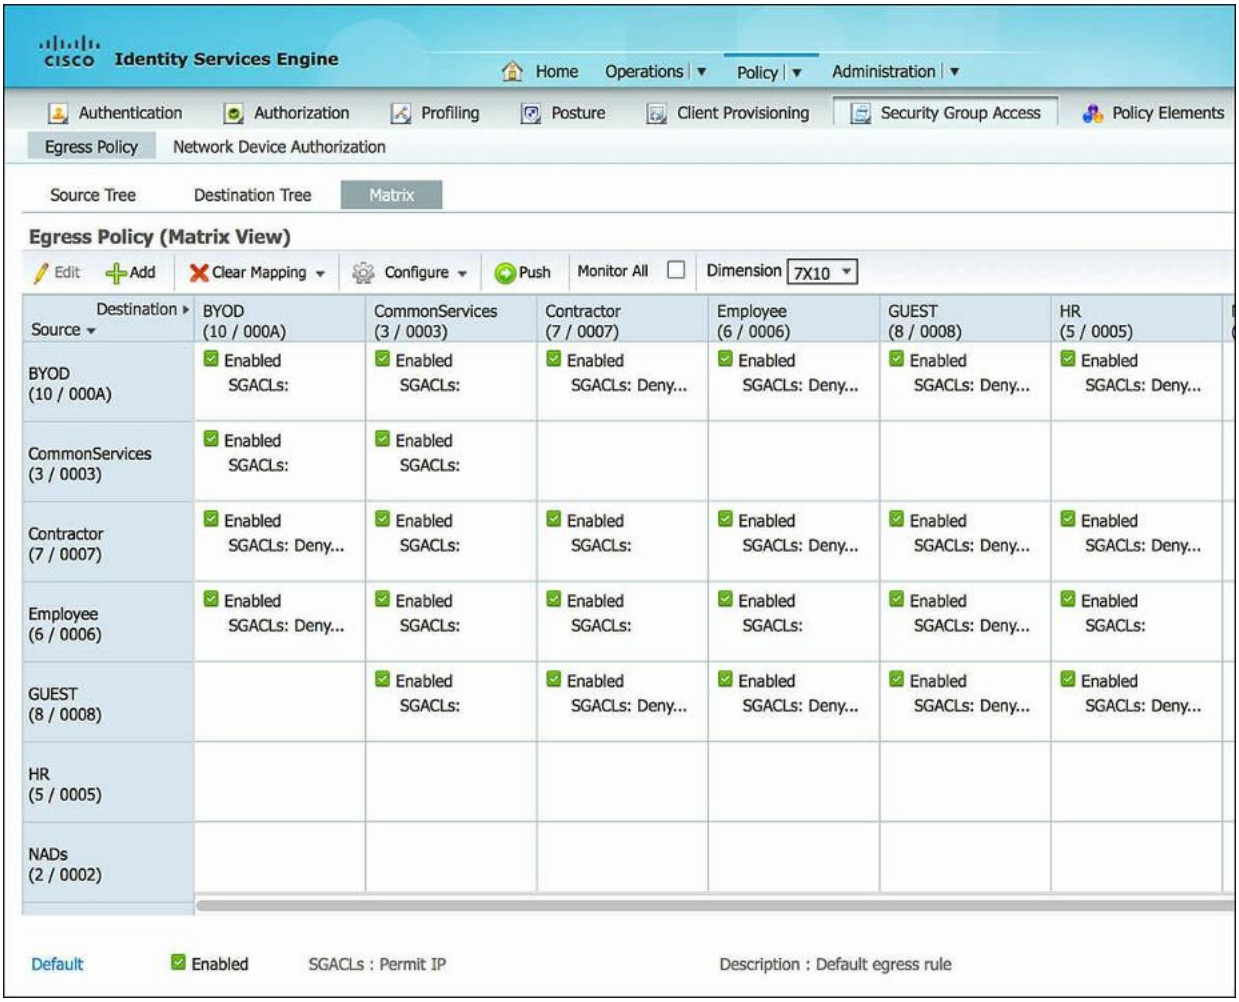


Figure 14-33 SGA egress policy matrix.

Now, the guest user's authorization on the NAD contains a downloadable ACL that limits what the user is allowed to access and a redirect URL (with corresponding redirect ACL). As the user attempts to access a network resource—any resource that matches the ACL-WEBAUTH-REDIRECT ACL—the NAD will redirect the unauthenticated user to the ISE for authentication.

Via the WebAuth portal, the client will be asked to provide his sponsor-generated username and password. If the guest user provides valid credentials, ISE will trigger a CoA to the NAD as we discussed at the beginning of this chapter, forcing a reauthorization. The user will again be processed through the authorization rule set.

Everything about the second authorization process will behave exactly as the first authorization process, with one exception—a flag will be set within ISE that says that the initial WebAuth was successful. This prevents ISE from getting into an endless loop for authentication and authorization whereby the client hits the catch-all MAB rule each time. To achieve a distinct experience during the reauthorization, we've created an Authorization rule named Guest.

This Guest rule will match if two conditions are met—the username is part of the user identity group called Guests and the RADIUS session is using GuestFlow. The user identity group Guests is the repository where the created guest users' credentials will be stored for use. GuestFlow is a simple condition whereby Network Access:UseCase Equals Guest Flow, meaning the guest user had already successfully authenticated to the WebAuth portal. If the initial WebAuth process was successful, ISE will link the successful authentication with the RADIUS session, making the Network Access:UseCase Equals Guest Flow equal TRUE. Because both conditions are true, the endpoint, and therefore the user, will be assigned the authorization profile of Internet_Only and SGA policy of GUEST. We've already discussed the SGA policy of GUEST, whereby the endpoint is assigned SGT 8 and explicitly denied access to Contractors (SGT 7), Employees (SGT 6), Guests (SGT 8), and HR (SGT 5). Let's briefly discuss the Internet_Only authorization

https://t.me/learningnets

profile.

Looking at the Internet_Only profile (see Figure 14-34 through Figure 14-36), as was the case with the CWA authorization profile, the Access Type is ACCESS_ACCEPT. Again, ACCESS_ACCEPT simply means that some level of access will be granted to the endpoint, however restrictive or permissive. You will also notice that the DACL Name is also selected as a common task (see Figure 14-34). However, the value of this DACL name is indeed different from the CWA profile. In this case, the DACL is named Internet_Only and is configured as shown in Figure 14-37 for wired deployments.

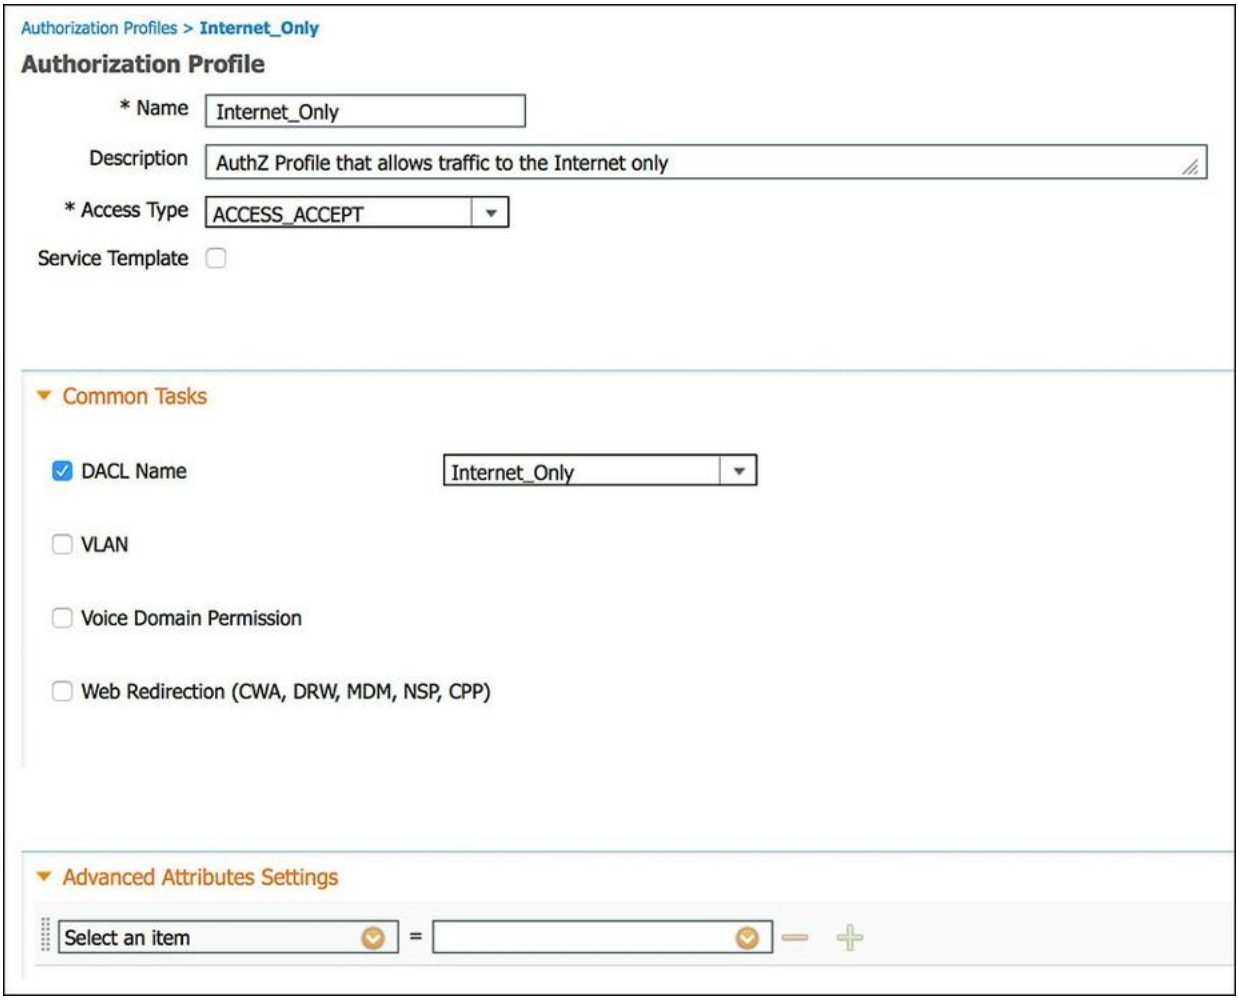


Figure 14-34 Internet_Only authorization profile, part 1.

https://t.me/learningnets

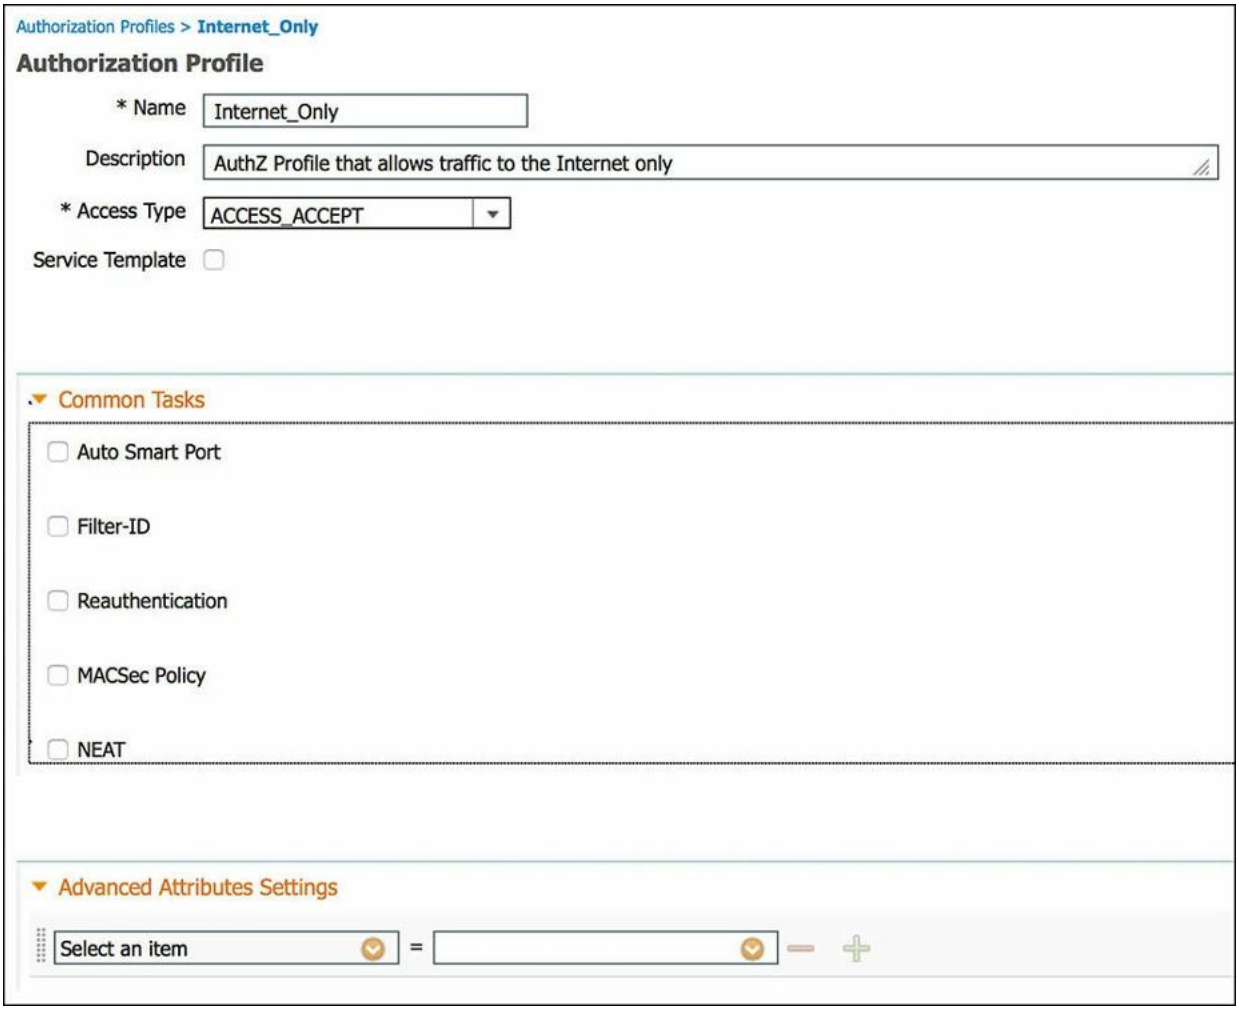


Figure 14-35 Internet_Only authorization profile, part 2.

https://t.me/learningnets

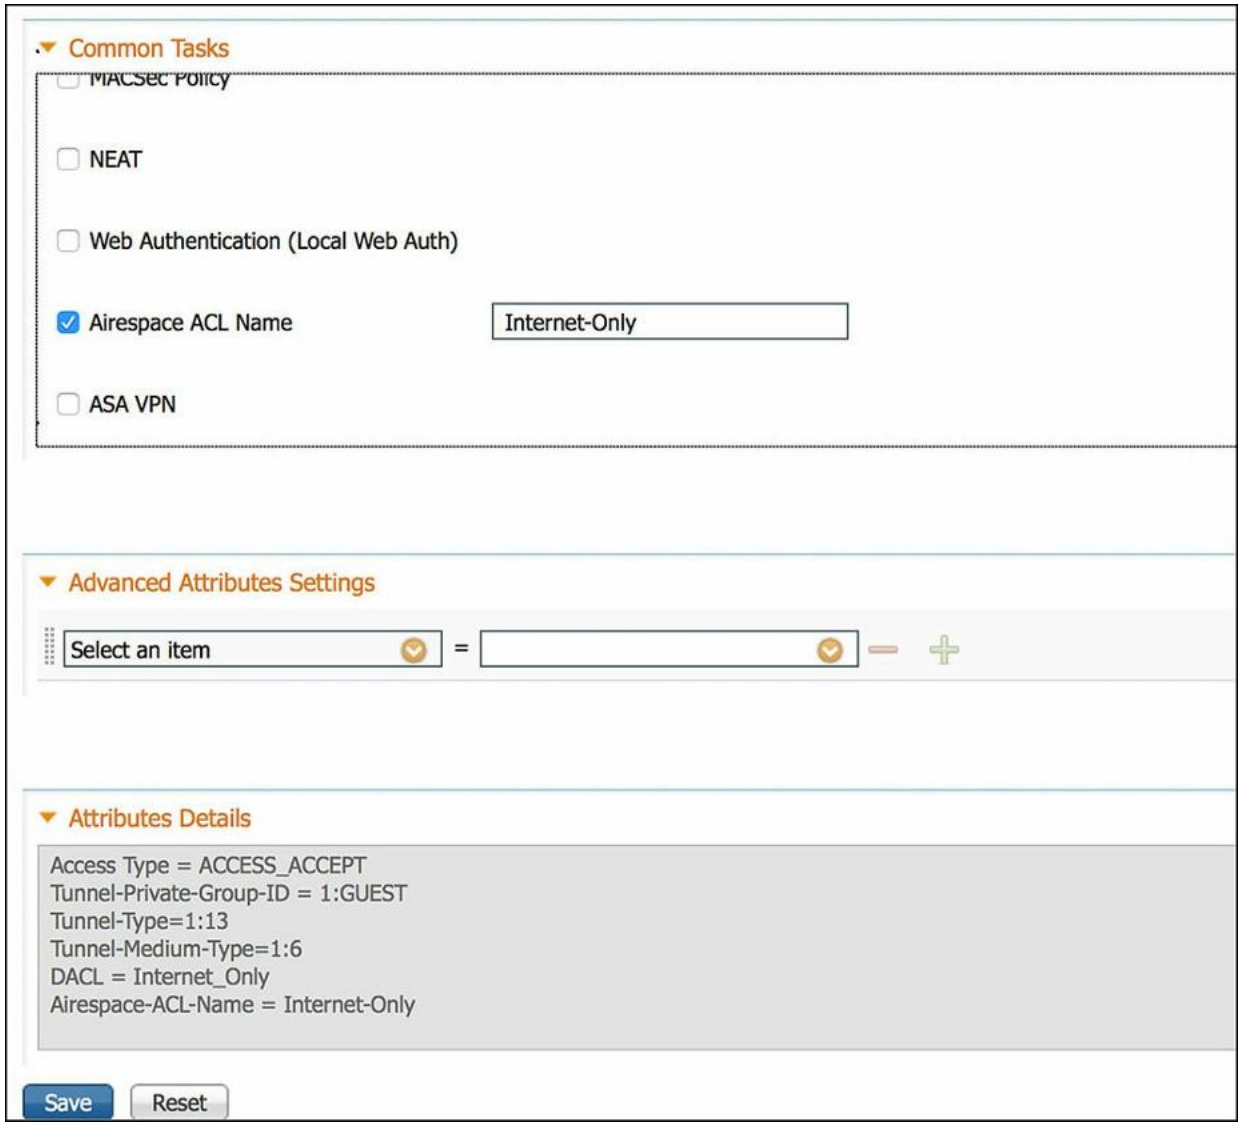


Figure 14-36 Internet_Only authorization profile, part 3.

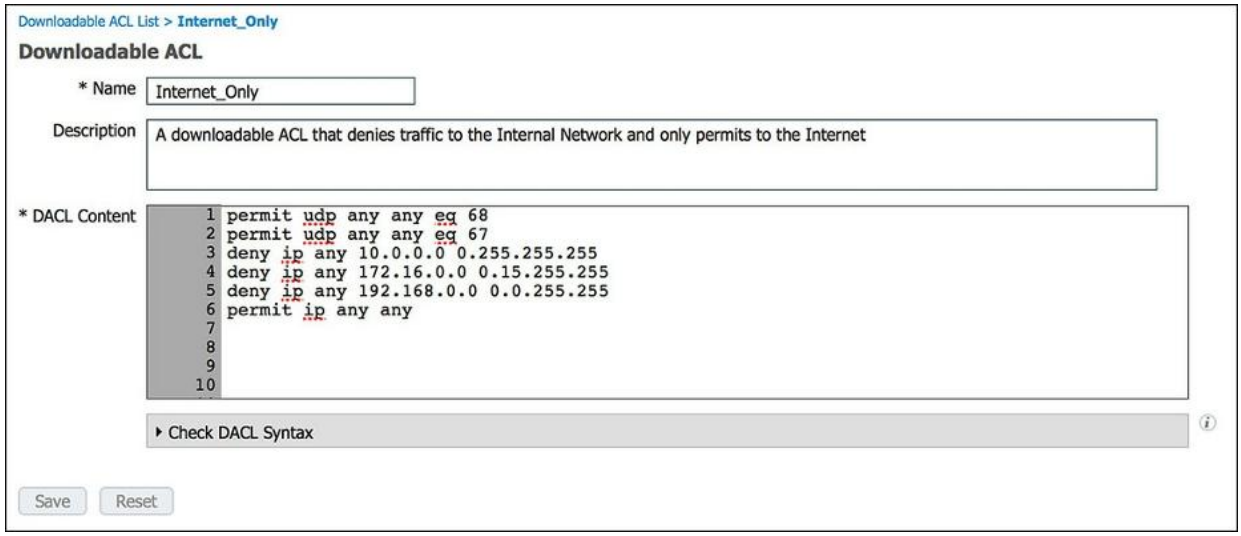


Figure 14-37 Internet_Only DACL.

The entries in the Internet_Only ACL is commented in Example 14-2 for your understanding:

https://t.me/learningnets

Example 14-2 Internet_Only ACL

Click here to view code image

```
! Traffic from DHCP client
permit udp any any eq 68
! Traffic to a DHCP server
permit udp any any eq 67
! Deny all traffic to RFC1918 private/internal IP ranges
deny ip any 10.0.0.0 0.255.255.255
deny ip any 172.16.0.0 0.15.255.255
deny ip any 192.168.0.0 0.0.255.255
! Allow all other traffic—i.e. Internet-bound traffic
permit ip any any
```

The other common task we'll leverage for the Internet_Only authorization profile is Airespace ACL Name. Airespace ACL Name is a RADIUS Attribute-Value Pair to send a named ACL to Cisco's wireless access points (APs), usually via a Wireless LAN Controller (WLC). Many of Cisco's APs do not allow more than 64 ACLs and no more than 64 access control entries per list. This restriction can limit the number of ACLs that can be used in the wireless domain. To support this approach while minimizing overhead, ISE leverages the named ACL approach, whereby the administrator will define ACLs explicitly on the WLC that is controlling the APs. ISE will then send a RADIUS Attribute-Value pair to the NAD instructing which ACL to associate with the endpoint/user. At that point, the endpoint's access will be appropriately limited depending on the named ACL that was referenced in the RADIUS exchange.

The Internet-Only named ACL for the WLC is shown in Figure 14-38.

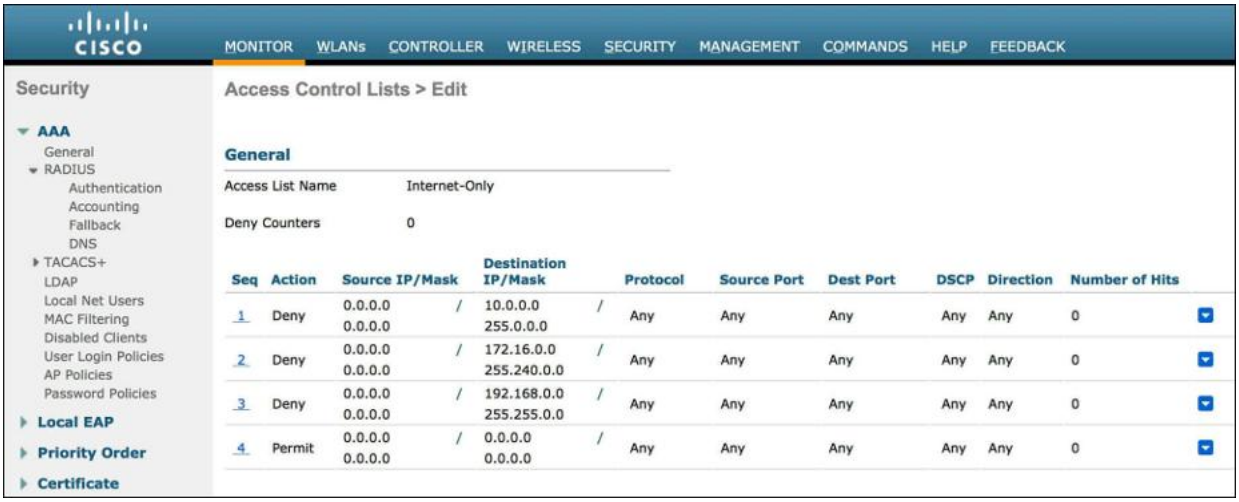


Figure 14-38 Internet-Only named ACL.

As you can see in this section, ISE can be extremely flexible in handling a number of scenarios for guest users—for both authentication as well as authorization. By breaking the overall process flow down into smaller, more easily consumable pieces, we were able to construct a very effective and quite restrictive guest access policy. Of course, depending on your own security policy, the level of access granted at each step along the way, as well as the conditions that trigger the access, might be completely different.

Provisioning Guest Accounts from a Sponsor Portal

The sponsor portal of ISE is mostly self-explanatory (see Figure 14-39). As you browse to the ISE sponsor portal (the default URL is https://<PSN_IP_OR_FQDN>:8443/sponsorportal), you will see that there are three key options—create a single account, create random accounts, and import accounts. Depending on which sponsor group you are a member of, you might or might not be able to use all three of these options. We'll now discuss each of these options in more detail.

https://t.me/learningnets

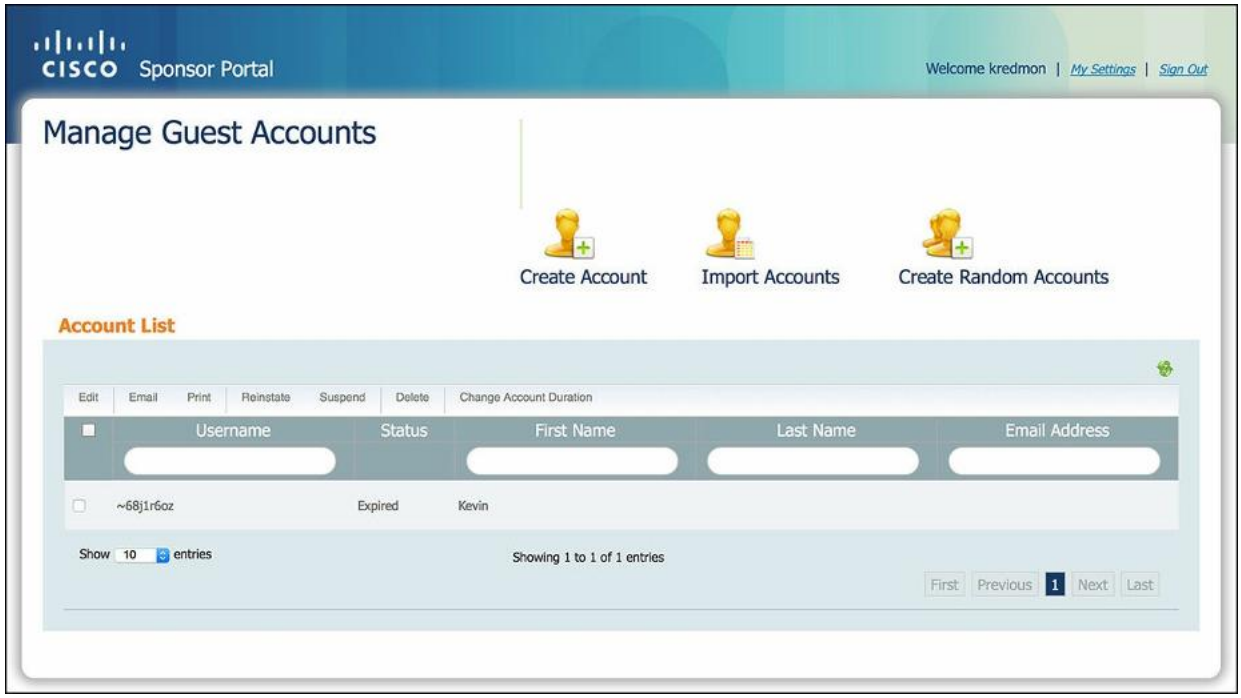


Figure 14-39 Sponsor portal.

Individual

To create an individual user account, such as a single guest user who is visiting the office for the day, click the man icon labeled Create Account. As you click this link, you will be presented with several options, as shown in Figure 14-40. Some of these fields will be limited to certain values, commensurate with your sponsor group.

https://t.me/learningnets

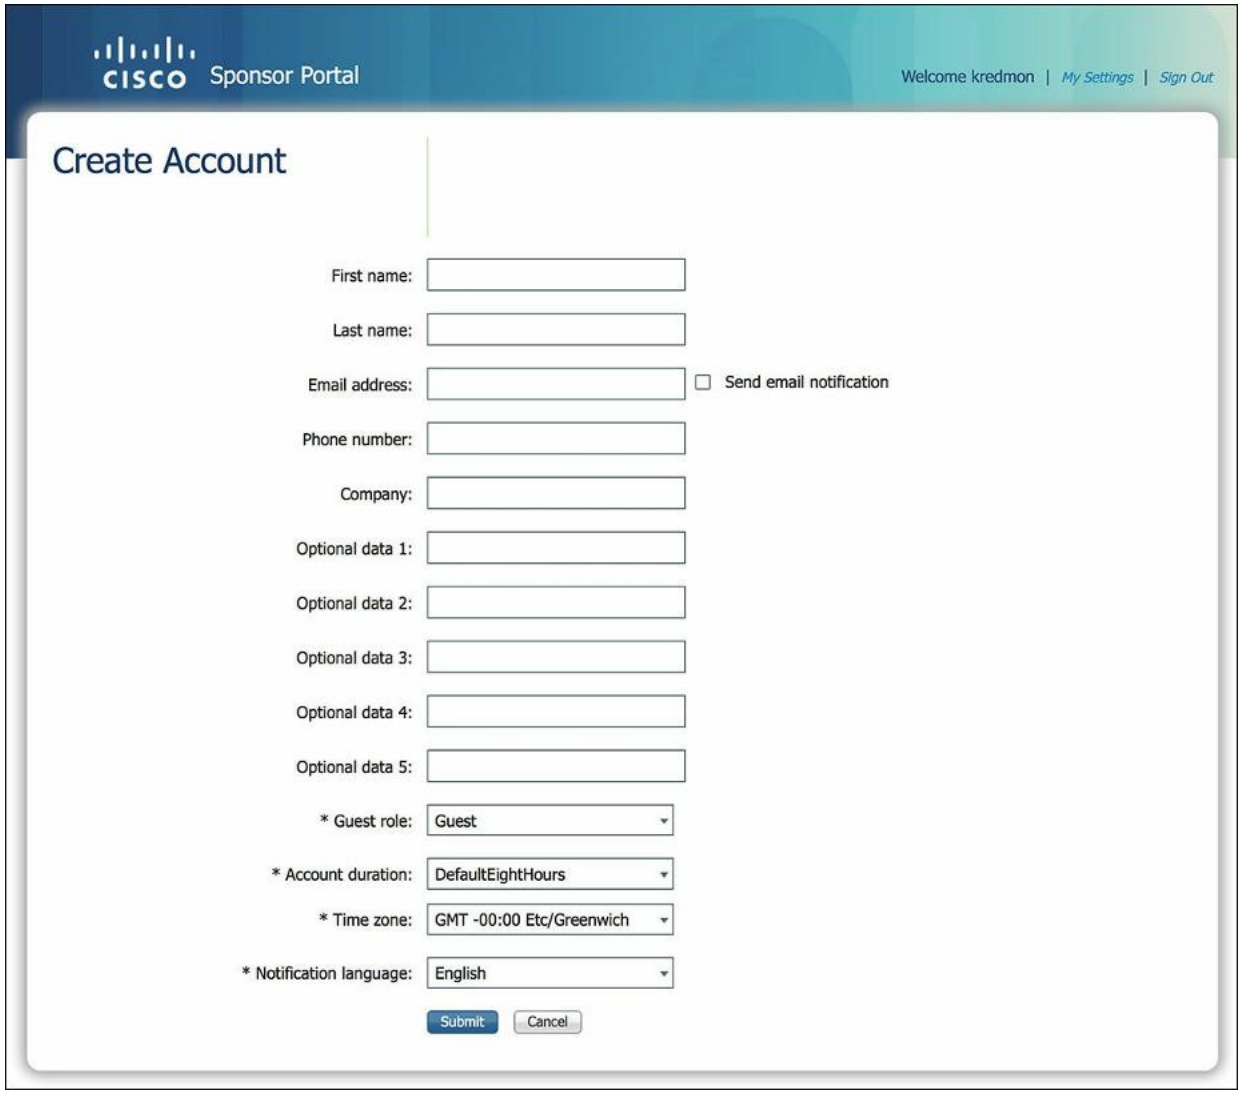


Figure 14-40 Sponsor portal individual account creation.

Key Topic

Depending on your security policy, some of these options might be required while others might be optional. You can configure which options are required by browsing to **Administration > Web Portal Management > Settings > Guest > Details Policy**. If you want to modify the Optional Data fields to be something more meaningful, you can do so by going to the **Administration > Web Portal Management > Settings > Sponsor > Language Template**. Select the appropriate language that should be modified. After selecting the appropriate language, you can now choose which sponsor portal page needs to be modified (such as Configure Create Single Guest Account).

Optional versus required fields are global across all guest portals.

Random

If you are hosting a conference with a large number of people attending or hosting a tour of your corporation, it might not always be possible to know the names of everyone attending well ahead of time. Furthermore, you may simply not want to give the guest users easily guessable usernames. Whichever the case, ISE provides the ability to create random user accounts. To create random user accounts, click the man icon labeled Create Random Accounts.

https://t.me/learningnets

After you click this link, you be asked to input the number of user accounts that will be required as well as a username prefix, which is optional (see Figure 14-41). This prefix will be prepended to the username of each guest user account that will be created. Again, some of the fields will be limited to certain values, commensurate with your sponsor group.

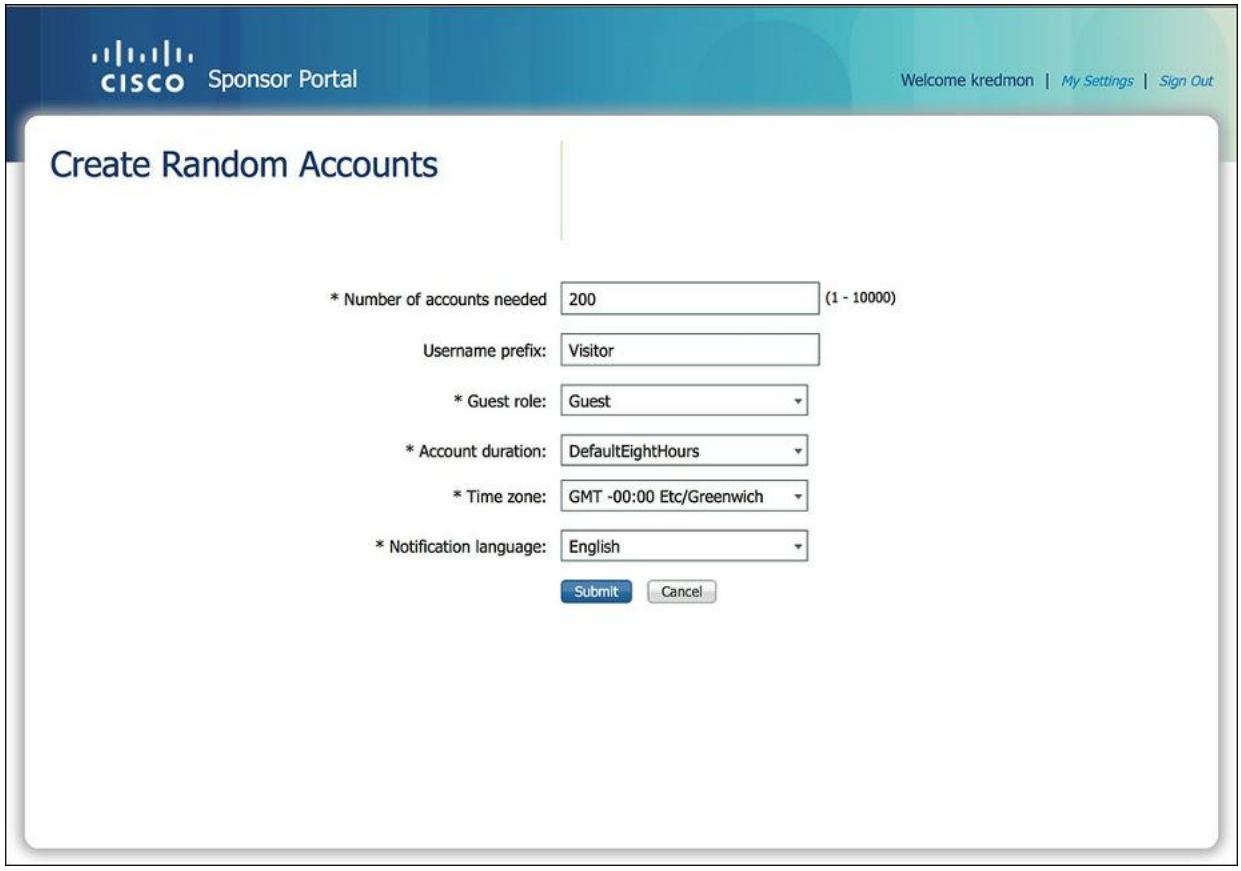


Figure 14-41 Sponsor portal random account creation.

Import

A third option to create guest user accounts is by importing a CSV that contains the requested data for each guest user account. To create guest user accounts using the import function, click the man icon labeled Import Accounts.

As you click the link, you can select the file that contains the user accounts (see Figure 14-42). If you are not confident as to the format of the file, you can download a template by clicking the Download Template button. This template will provide many of the same fields as requested in the Individual Account creation. Once you have populated the file with the required fields, select the file by clicking the Browse button; then click Upload. As the file uploads, you will see a list within the sponsor portal of all user accounts that were detected in the file. You are not done yet. If all the usernames look correct and were properly parsed, click Submit. After you click Submit, you will see a list of all usernames from the file and the created passwords.

https://t.me/learningnets

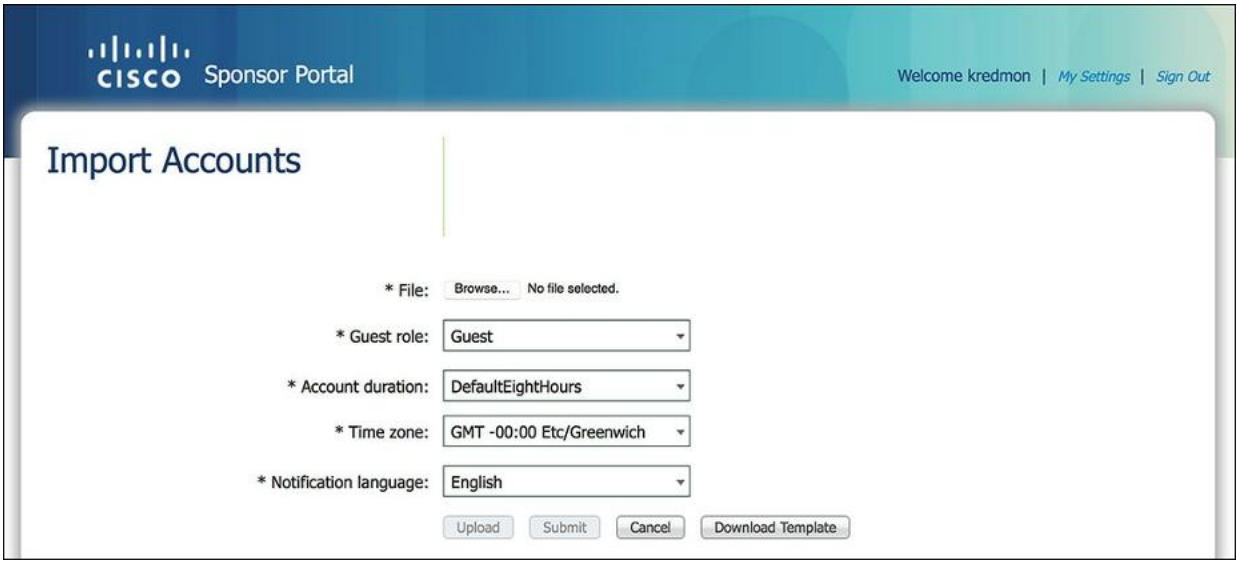


Figure 14-42 Sponsor portal import account creation.

Verifying Guest Access on the WLC/Switch

When a guest user connects to the network via Guest Services, several RADIUS exchanges must happen between the NAD (WLC or switch) and ISE for the guest's network access to be successful. We'll review the steps to verify that the guest access is going well from all angles—from the standpoint of the NAD (WLC and switch) and from ISE.

WLC

As a guest user attempts to connect to the network wirelessly, the endpoint must go through the following steps:

1. **The endpoint must be connected to the correct SSID**—Guest Services, we have determined that the endpoint must be authenticating to the network using MAB. On the WLC, you must enable MAC filtering to trigger a MAB authentication on the SSID in question. In this case, the SSID that you'll be associating to is CP-Guest, which does indeed have MAC filtering enabled under the column Security Policies. You can view a summary of the SSIDs/WLANs present on a WLC by clicking the WLANs link across the top bar (see Figure 14-43).



Figure 14-43 Wireless LAN controller SSIDs/WLANs.

2. **The initial authorization policy pushed to the endpoint must be CWA and SGA policy of GUEST**—As the endpoint first associates to the SSID, ISE will determine that the client is using MAB based on the RADIUS Service-Type sent by the WLC to ISE. However, as we discussed earlier, the MAC address is not present in the endpoint identity store. This will enable the endpoint to hit the WebAuth rule because WebAuth is the catchall MAB rule when the endpoint is not in the endpoint database. The resulting authorization profile should be CWA and an SGA Policy of GUEST.

https://t.me/learningnets

On ISE, you can easily see that we have hit the WebAuth rule, which we'll discuss shortly. However, on the WLC, it is not as straightforward. On the WLC, to ensure that the client has hit the correct rule, we must look at the status of the endpoint on the WLC (see Figure 14-44). To look at the status of the endpoint on the WLC, you must click Monitor > Clients and then click the MAC address of the client in question.

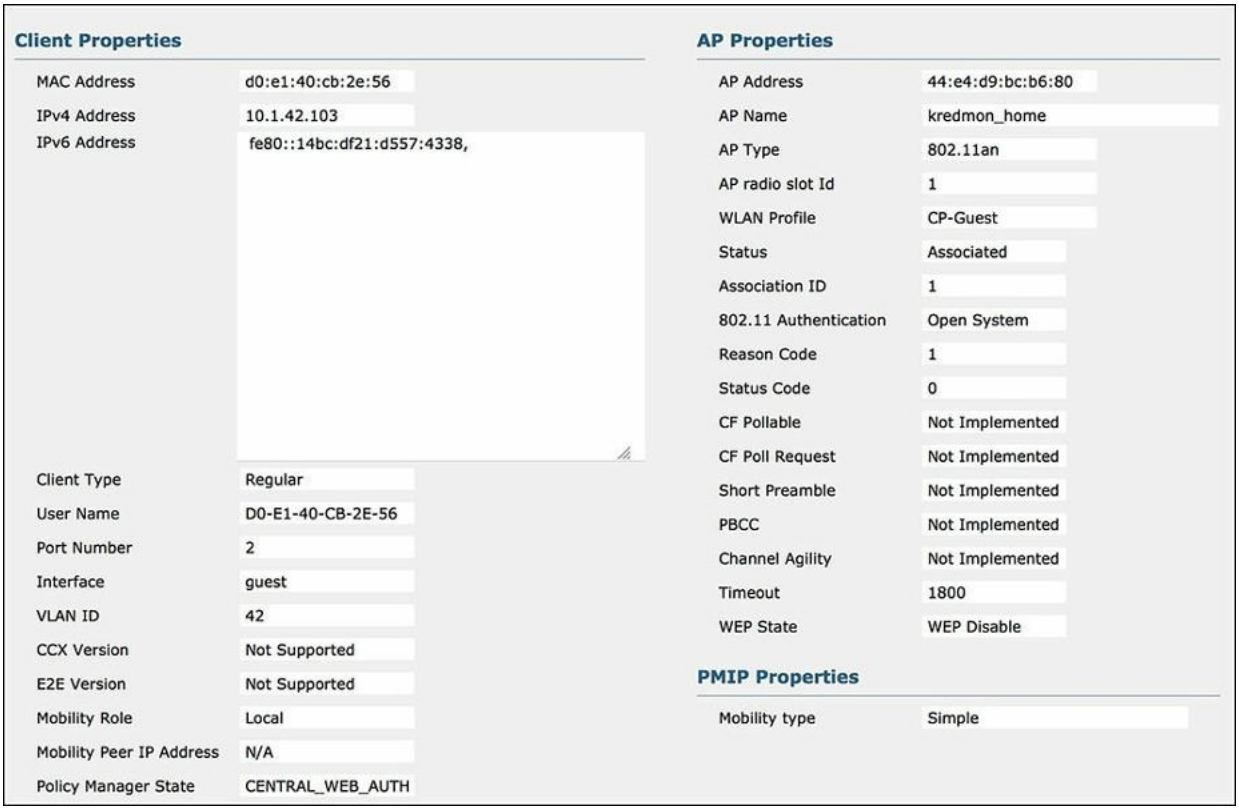


Figure 14-44 WLC endpoint status – WebAuth.

ISE should push the following information to the WLC on behalf of the endpoint when the CWA policy is hit:

- **MAC Address**—Confirm that the MAC address matches the endpoint in question.
- **WLAN Profile**—Although this isn't pushed to the WLC by ISE, it is still a good idea to confirm that the correct SSID is showing. In this case, the SSID should be CP-Guest, which it is.
- **Policy Manager State**—The Policy Manager state is an indication of the wireless connection status. In the case of CWA, the Policy Manager state should be CENTRAL_WEB_AUTH. This corresponds to the Centralized Web Auth web redirection we enabled on ISE CWA policy.
- **AAA Override ACL Name**—The AAA override ACL name is the ACL that is used for web redirection. This named ACL was defined as ACL-WEBAUTH-REDIRECT on the CWA policy on ISE.
- **Redirect URL**—The redirect URL is the URL for the WebAuth portal for the guest user. As the guest user attempts to browse to a website, this redirect will send the user to ISE to authenticate. This URL will point to one of the ISE PSNs.
- **CTS Security Group Tag**—The SGA policy of GUEST assigns the SGT of 8 to the endpoint. You can see that value here under the variable CTS security group tag.

3. **The final authorization policy pushed to the endpoint must be Internet_Only and an SGA policy of GUEST**—As the guest user authenticates and the NAD receives a CoA from ISE, the endpoint is sent back through the authorization process. Becuase a successful authentication occurs, the endpoint should now hit the Internet_Only authorization profile. This profile is a little less

https://t.me/learningnets

involved than the CWA policy.

ISE should push the following information to the WLC when the Internet_Only profile is hit (see Figure 14-45):

- **MAC Address**—Again, confirm that the MAC address matches the endpoint in question.
- **WLAN Profile**—Although this isn't pushed to the WLC by ISE, it is still a good idea to confirm that the correct SSID is showing. In this case, the SSID should still be CP-Guest, which it is.
- **Policy Manager State**—In the case of Internet_Only, the Policy Manager state should be RUN. This is the final connection state for the endpoint.
- **AAA Override ACL Name**—In the Internet_Only profile, there should NOT be a AAA override ACL name.
- **Redirect URL**—Because the guest user has authenticated, the endpoint should no longer need to be redirected. For this reason, the redirect URL should be blank.
- **IPv4 ACL Name**—In the case of (Wireless) CWA, there was no IPv4 ACL name. The AAA override ACL name was doing all necessary filtering of the traffic. Now that ISE has authenticated the guest user and, therefore, the endpoint, ISE will now push the Airespace ACL of Internet-Only to the WLC. As we discussed earlier in the chapter, an Airespace ACL is also a named ACL that is predefined on the WLC.
- **CTS Security Group Tag**—The SGA Policy of GUEST is exactly the same as the CWA policy. Looking at this field, you will again see that it is 8.

https://t.me/learningnets

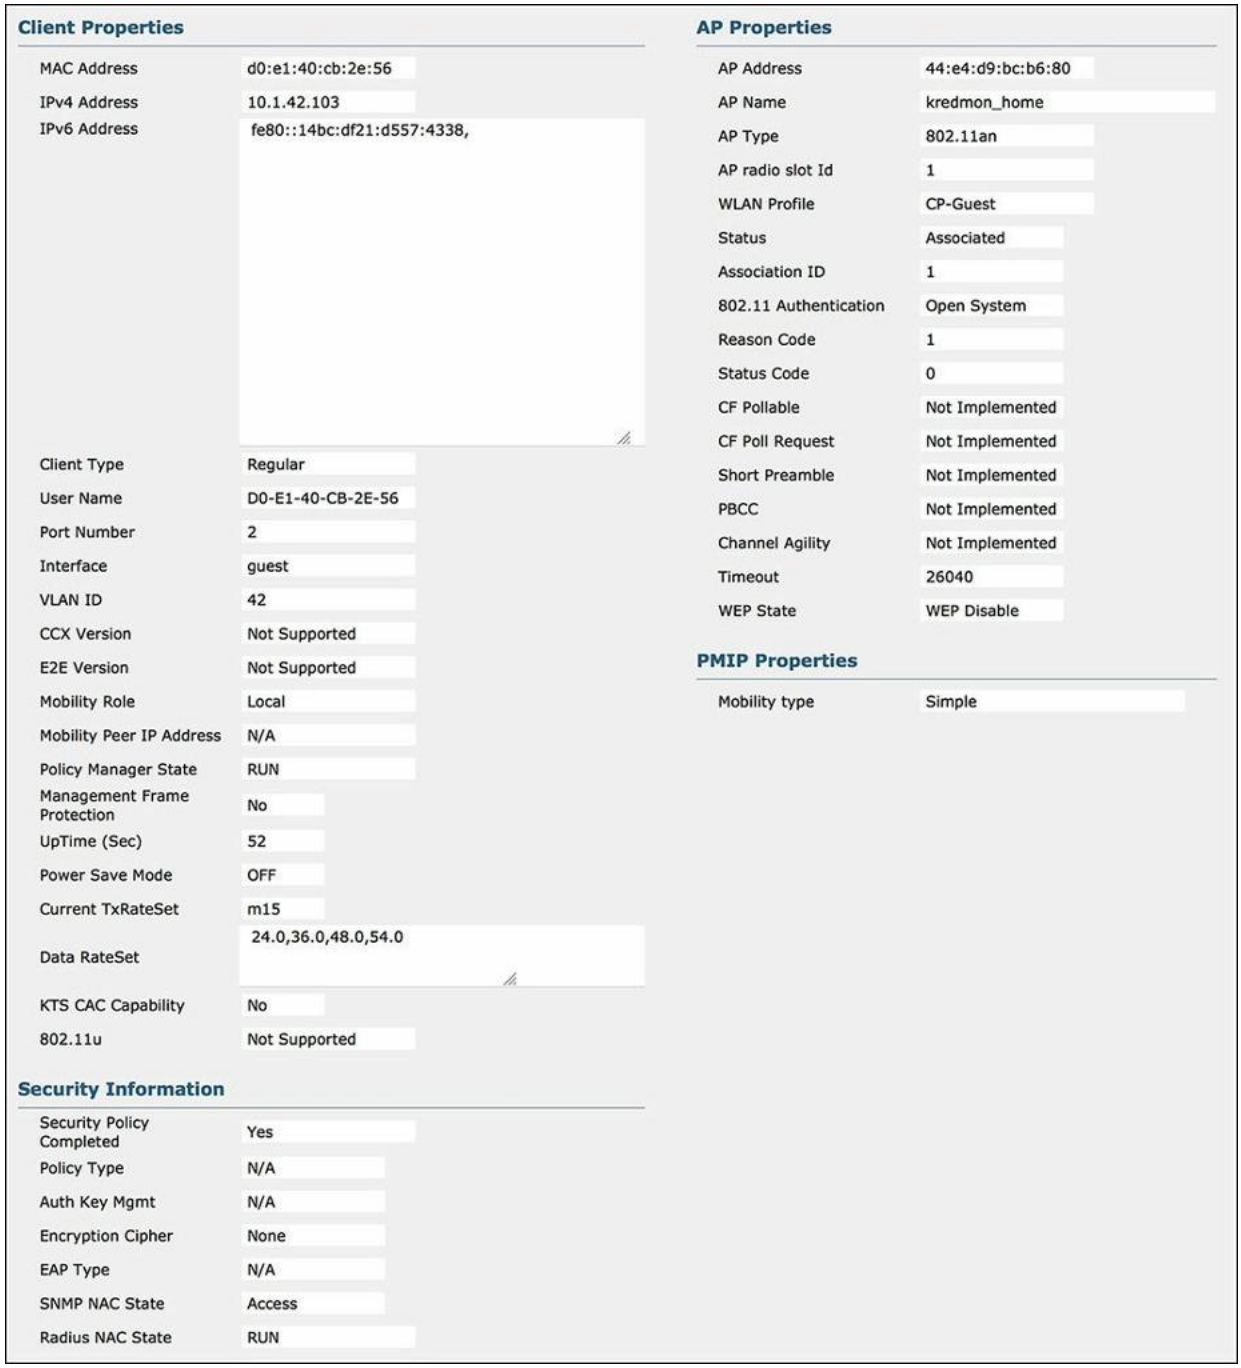


Figure 14-45 WLC Endpoint Status—Internet_Only.

Now that we have validated the connection from the standpoint of the WLC, we can take a look at ISE, highlighting the communications that are sent between the WLC and ISE and the return traffic. When we look at the ISE authentication details for the initial WebAuth authorization, and the final Internet_Only authorization, we will need to ensure that all authorization conditions are met from the WLC and that the ISE pushes the correct information back to the WLC in both cases.

4. **ISE gets the initial authentication request and sends the WLC the CWA authorization profile and SGA policy GUEST**—When the endpoint first joins the WLC, the WLC will send a MAB request to ISE. For the endpoint to hit the WebAuth policy, the WLC must meet the conditions of Wireless_MAB (or Wired_MAB, but we know that this endpoint is wireless). This initial authentication request should result in the following communications:

   ■ **From WLC to ISE**—Looking at the authentication details for the endpoint in question (the

https://t.me/learningnets

magnifying glass in the Details column of the Authentication Live Log), there are a number of sections in the output. For the communication from the NAD, you will need to look at the Overview and Authentication Details section of the output (see Figure 14-46).

- **Authorization Profile**—The Overview shows the condensed version of the authorization output. Looking at the authorization profile, you will see CWA and GUEST. These are the authorization profile and SGA policy you should be seeing.
- **AuthorizationPolicyMatchedRule**—This variable tells us the name of the authorization rule we hit. In the case of a guest authenticating to the network, the rule we should hit is indeed WebAuth.
- **Endpoint ID**—When looking at the authentication details for a particular endpoint, ensure that you are looking at the MAC address of the correct endpoint.
- **Authentication Method**—The Authentication method as seen on the authentication details should be MAB.
- **Service Type**—The service type for a MAB connection from a Cisco WLC should be Call Check.
- **NAS Port Type**—The NAS port type for a wireless connection should be Wireless—IEEE 802.11.

https://t.me/learningnets

Overview

Event	5200 Authentication succeeded
Username	D0:E1:40:CB:2E:56
Endpoint Id	D0:E1:40:CB:2E:56
Endpoint Profile	Apple-iPad
Authorization Profile	CWA,GUEST
AuthorizationPolicyMatchedRule	WebAuth
ISEPolicySetName	Default
IdentitySelectionMatchedRule	Default

Authentication Details

Source Timestamp	2015-01-30 23:58:34.876
Received Timestamp	2015-01-30 23:58:34.876
Policy Server	atw-cp-ise02
Event	5200 Authentication succeeded
Failure Reason	
Resolution	
Root cause	
Username	D0:E1:40:CB:2E:56
User Type	Host
Endpoint Id	D0:E1:40:CB:2E:56
Endpoint Profile	Apple-iPad
IP Address	
Identity Store	Internal Endpoints
Identity Group	Profiled
Audit Session Id	0a013c02000000b354cc1ac1
Authentication Method	mab
Authentication Protocol	Lookup
Service Type	Call Check
Network Device	WLC
Device Type	Wireless
Location	NorthAmerica#SJC
NAS IP Address	10.1.60.2
NAS Port Id	
NAS Port Type	Wireless - IEEE 802.11
Authorization Profile	CWA,GUEST
Posture Status	NotApplicable
Security Group	GUEST
Response Time	36

https://t.me/learningnets

Figure 14-46 ISE Authentication Details Live Log—Wireless WebAuth.

Because Wireless_MAB is defined by a service type of Call Check and a NAS port type of Wireless—IEEE 802.11, the conditions for the CWA authorization profile and SGA policy of GUEST have been met. ISE must now send the resulting authorization profile back to the WLC. The information in this authorization profile is provided in the Result section of the authentication details. Looking at this output (shown in Figure 14-47), we see:

- **From ISE to WLC**—The output in the Result section of the authentication live log is
  - **Redirect URL**—The redirect URL is contained within one of the cisco-av-pair outputs, specifically the one that begins with url-redirect=. As you look at the rest of the URL, you will notice that there is a parameter called action=cwa. As the endpoint is redirected to that URL, ISE recognizes this parameter as a CWA call and provides the endpoint the correct webpage.
  - **Redirect ACL**—The redirect ACL is also contained within one of the cisco-av-pair outputs in the Result section. In this case, the cisco-av-pair will begin with url-redirect-acl= (notice the acl at the end), specifically mentioning the ACL-WEBAUTH-REDIRECT ACL.
  - **Security Group Tag**—The SGT that was assigned to the endpoint can also be found in the Result section of the authentication details. Again, this value will be embedded in a cisco-av-pair—namely the AV pair of cts:security-group-tag=0008-0.



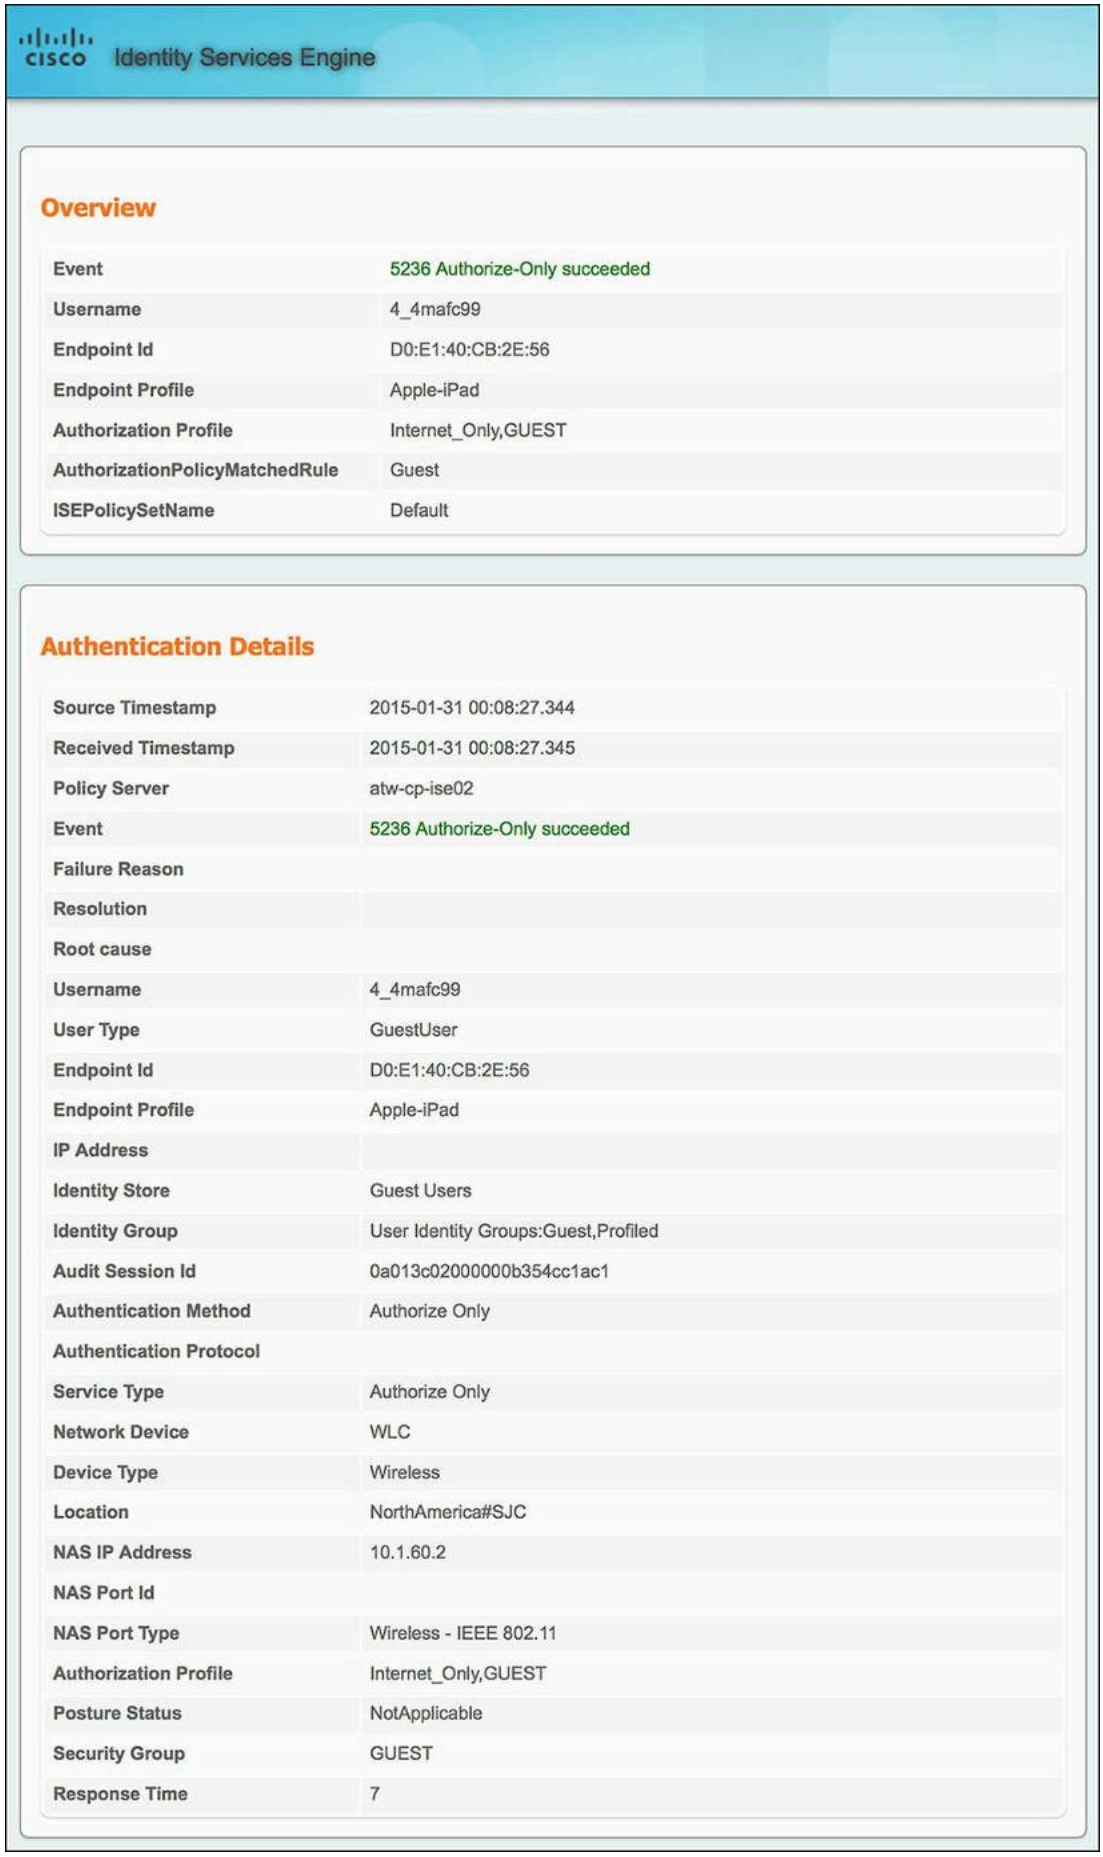
Result

UserName	D0:E1:40:CB:2E:56
User-Name	D0-E1-40-CB-2E-56
State	ReauthSession:0a013c02000000b354cc1ac1
Class	CACS:0a013c02000000b354cc1ac1:atw-cp-ise02/210675310/23037
cisco-av-pair	url-redirect-acl=ACL-WEBAUTH-REDIRECT
cisco-av-pair	url-redirect=https://atw-cp-ise02.ise.local:8443/guestportal /gateway?sessionId=0a013c02000000b354cc1ac1&action=cwa
cisco-av-pair	ACS:CiscoSecure-Defined-ACL=#ACSACL#-IP-Pre-WebAuth-54bb3896
cisco-av-pair	cts:security-group-tag=0008-0
cisco-av-pair	profile-name=Apple-iPad

Figure 14-47 ISE Authentication Details Live Log—Wireless WebAuth results.

5. **ISE gets the final authentication request and sends the WLC the Internet_Only authorization profile and SGA policy GUEST**—When the guest user authenticates to the web portal, a CoA will be sent to the WLC on his behalf, forcing the endpoint to reauthorize. If the guest user's web authentication is successful, the user should be given the authorization policy of Internet_Only and an SGA policy GUEST. The conditions that must match for this to happen are Guest user identity store and Network Access:UseCase = Guest Flow. This final authentication request should result in the following communications (see Figure 14-48):

https://t.me/learningnets



https://t.me/learningnets

Figure 14-48 ISE Authentication Details Live Log—Wireless Internet_Only.

- **From WLC to ISE**—Looking at the authentication details for the endpoint in question, we will focus on the Overview and Authentication Details section of the output.
  - **Authorization profile**—Under the Overview, you will get the condensed version of the authorization output. Looking at the authorization profile, you will see Internet_Only and GUEST. These are the authorization profile and SGA policy you should be seeing.
  - **AuthorizationPolicyMatchedRule**—In the case of a successful guest user web authentication and reauthorization, this rule should be Guest.
  - **Endpoint ID**—When looking at the authentication details for a particular endpoint, ensure that you are looking at the MAC address of the correct endpoint.
  - **Authentication method**—With a CoA, there will not be a need to authenticate. Therefore, this output should be Authorize Only.
  - **Identity group**—Because the guest user is a part of the Guest identity group, this identity group list should contain User Identity Groups:Guest.
  - **Use case**—The other condition, besides the **Guest** identity group, is GuestFlow, which is equal to Network Access:UseCase EQUALS Guest Flow. Under Other Attributes, you will see UseCase equals Guest Flow.

Because the Guest rule has the two conditions of Guest identity group and GuestFlow, and because both of these conditions are matched, the authorization profile of Internet_Only and SGA policy of GUEST must be pushed back to the WLC. The information in this authorization profile is provided in the Result section of the authentication details. This output (shown in Figure 14-49) is as follows:

- **From ISE to WLC**—The Result section of the authentication live log shows
  - **Airespace-ACL-Name**—The Airespace ACL name is a named ACL pushed to the NAD on behalf the endpoint. This named ACL, Internet-Only, will ensure that the endpoint can access only non-RFC1918 address space.
  - **Security Group Tag**—The SGT that was assigned to the endpoint can also be found in the Result section of the Authentication details. Again, this value will be embedded in a cisco-av-pair—namely the AV pair of cts:security-group-tag=0008-0. This is the same SGT value that was pushed in the WebAuth scenario.

https://t.me/learningnets



Result

User-Name	4_4mafc99
State	ReauthSession:0a013c02000000b354cc1ac1
Class	CACS:0a013c02000000b354cc1ac1:atw-cp-ise02/210675310/23046
Session-Timeout	26040
Termination-Action	Default
Tunnel-Type	(tag=1) VLAN
Tunnel-Medium-Type	(tag=1) 802
Tunnel-Private-Group-ID	(tag=1) GUEST
cisco-av-pair	ACS:CiscoSecure-Defined-ACL=#ACSACL#-IP-Internet_Only-54cb00fc
cisco-av-pair	cts:security-group-tag=0008-0
cisco-av-pair	profile-name=Apple-iPad
Airespace-ACL-Name	Internet-Only

Figure 14-49 ISE Authentication Details Live Log—Wireless Internet_Only results.

Switch

In a similar fashion as a wireless guest user, the wired guest user will also go through a number of distinct steps to authenticate to the network. The overall flow is mostly the same, except some of the RADIUS exchanges change from the wireless to the wired counterpart. Let's look at each of the steps in detail:

1. **The endpoint plugs into a switchport that is configured for MAB**—In the WLC, MAB was configured as MAC Filtering. On a switch, the switchport must be explicitly configured for MAB. In many cases, the port is configured for dot1x first and foremost, with a failover option of MAB. That is the case with this scenario.
2. **Because the port is configured for dot1x, we'll need to confirm that it fails over to MAB**—As mentioned previously, in most scenarios, dot1x is the preferred authentication method. Therefore, the wired guest user must wait until dot1x fails to receive a response. This will result in a successful MAB authentication. You can monitor the status of the authentication process by running the command **show authentication sessions interface <if_name> details**.
3. **The initial authorization policy pushed to the endpoint must be CWA and SGA policy of GUEST**—The output of **show authentication sessions interface <if_name> details** shows the RADIUS Attribute-Value pairs that are sent to the switch (see Example 14-3).

Example 14-3 show authentication sessions interface <if_name> **details**

Click here to view code image

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
      Oper host mode:  multi-auth
    Oper control dir:  both
```

https://t.me/learningnets

```
      Session timeout:  N/A
    Common Session ID:  0A742B860000003B0231777D
      Acct Session ID:  0x00000049
               Handle:  0x2B000029
       Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:
         URL Redirect:  https://atw-cp-ise02.ise.local:8443/guestportal/
gateway?sessionId=0A742B860000003B0231777D&action=cwa
     URL Redirect ACL:  ACL-WEBAUTH-REDIRECT
              ACS ACL:  xACSACLx-IP-Pre-WebAuth-54bb3896
            SGT Value:  8

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

Looking at this output, we can see the following information is pushed to the NAD:

- **MAC Address**—Again, confirm that the MAC address matches the endpoint in question.
- **Status**—The status of the endpoint should be Authorized. If it is not Authorized, you might have a communication or policy issue on the ISE server.
- **URL Redirect**—The redirect URL is the URL for the WebAuth portal for the guest user. Because the guest user attempts to browse to a website, this redirect will send the user to ISE to authenticate. This URL will point to one of the ISE PSNs.
- **URL Redirect ACL**—The URL redirect ACL is the ACL used for web redirection. This named ACL was defined as ACL-WEBAUTH-REDIRECT on the CWA policy on ISE.
- **ACS ACL**—The ACL Pre-WebAuth is a DACL pushed from ISE as part of the CWA policy. As a DACL is pushed from ISE, ISE will make a copy of the ACL for each endpoint, giving each DACL a unique identifier. In this case, the DACL was given the per-username xACSACLx-IP-Pre-WebAuth-54bb3896.
- **SGT Value**—The SGA policy of GUEST assigns the SGT of 8 to the endpoint. You can see that value here under the variable SGT value.
- **Method Status List**—The method status list is a list of the authentication methods that can be used to authenticate a client. This list should indicate a success (such as an Authc Success) for the MAB method. The other methods that might be available may be in a Stopped state. You can see that dot1x is in a Stopped state in this example.

4. **The final authorization policy pushed to the endpoint must be Internet_Only and SGA policy of GUEST**—Again, looking at the output of **show authentication sessions interface <if_name> details**, we see the following policy is pushed to the NAD (see Example 14-4).

Example 14-4 show authentication sessions interface <if_name> **details**

Click here to view code image

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
```

https://t.me/learningnets

```
            User-Name:  ~68j1r6oz
               Status:  Authorized
               Domain:  DATA
        Oper host mode:  multi-auth
      Oper control dir:  both
       Session timeout:  N/A
     Common Session ID:  0A742B860000003B0231777D
       Acct Session ID:  0x00000049
                Handle:  0x2B000029
        Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
       Security Policy:  Should Secure
       Security Status:  Link Unsecure

Server Policies:
              ACS ACL:  xACSACLx-IP-Internet_Only-54cb00fc
            SGT Value:  8

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

- **MAC Address**—Always check the MAC address!
- **Status**—As before, the status of the endpoint should be Authorized. If it is not Authorized, you might have a communication or policy issue on the ISE server.
- **URL Redirect**—The URL redirect will no longer be present after the endpoint has authenticated via WebAuth.
- **URL Redirect ACL**—As was the case with the URL redirect, the URL redirect will also no longer be needed after the endpoint has authenticated via WebAuth.
- **ACS ACL**—The ACL Internet_Only is a DACL pushed from ISE as part of the Internet_Only policy. As was the case with the Pre-WebAuth ACL, ISE will make a copy of the ACL for each endpoint, giving each instance of the DACL a unique identifier. In this case, the DACL was given the per-username xACSACLx-IP-Internet_Only-54cb00fc.
- **SGT Value**—The SGA policy of GUEST, as before, does not change. You can see that value here under the variable SGT value.
- **Method status list**—The method status list is a list of the authentication methods that can be used to authenticate a client. This list should indicate a success (such as an Authc Success) for the MAB method because we are still using MAB. The other methods that might be available may be in a Stopped state. You can see that dot1x is in a Stopped state in this example.

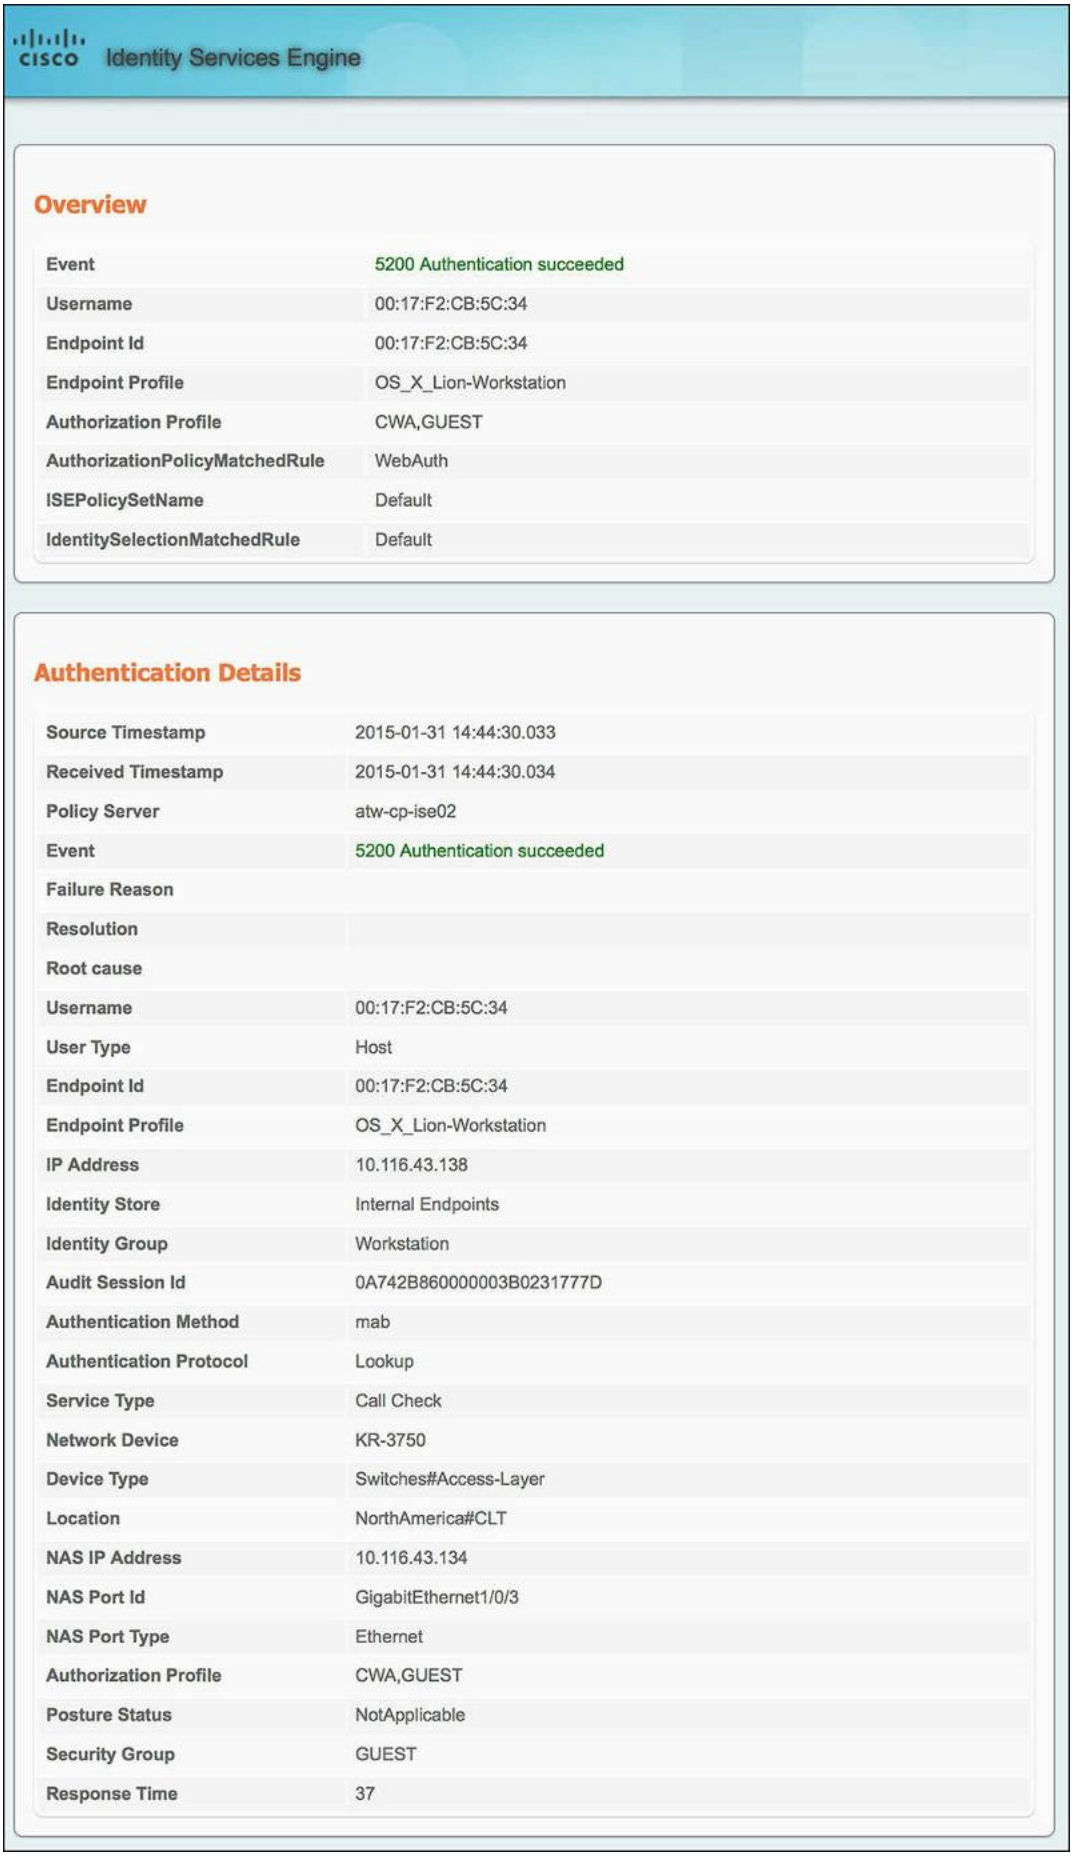
Now that we have validated the endpoint's connection from the standpoint of the switch, we can take a look at ISE, highlighting the communications sent between the switch and ISE and the return traffic. When we look at the ISE authentication details for the initial (WebAuth) authorization and the final (Internet_Only) authorization, we must ensure that all authorization conditions are met from the switch and that the ISE pushes the correct information back to the switch in both cases.

5. **ISE gets the initial authentication request and sends the switch the CWA authorization profile and SGA policy GUEST**—When the endpoint first plugs into the switch, the switch will send a MAB request to ISE. For the endpoint to hit the WebAuth policy, the switch must meet the conditions of Wired_MAB (or Wireless_MAB, but we know that this endpoint is wired). This initial authentication request should result in the following communications:
   - **From switch to ISE**—Looking at the authentication details for the endpoint in question (the magnifying glass in the Details column of the Authentication Live Log), there are a number of

https://t.me/learningnets

sections in the output. For the communication from the NAD, you will need to look at the Overview and Authentication Details section of the output (see Figure 14-50).

https://t.me/learningnets



https://t.me/learningnets

Figure 14-50 ISE Authentication Details Live Log—Wired WebAuth.

- **Authorization Profile**—The Overview shows the condensed version of the authorization output. Looking at the authorization profile, you will see CWA and GUEST. These are the authorization profile and SGA policy you should be seeing based on the WebAuth rule.
- **AuthorizationPolicyMatchedRule**—This variable tells us the name of the authorization rule we hit. In the case of a guest authenticating to the network, the rule we should hit is indeed WebAuth.
- **Endpoint ID**—When looking at the authentication details for a particular endpoint, ensure that you are looking at the MAC address of the correct endpoint.
- **Authentication Method**—The authentication method as shown on the authentication details should be MAB.
- **Service Type**—The service type for a MAB connection from a Cisco switch should be Call Check.
- **NAS Port Type**—The NAS port type for a wired connection should be Ethernet.

Because Wired_MAB is defined by a service type of Call Check and a NAS port type of Ethernet, the conditions for the CWA authorization profile and SGA policy of GUEST have been met. ISE must now send the resulting authorization profile back to the switch. The information in this authorization profile is provided in the Result section of the authentication details (see Figure 14-51). Looking at this output, we see:

Result

UserName	00:17:F2:CB:5C:34
User-Name	00-17-F2-CB-5C-34
State	ReauthSession:0A742B860000003B0231777D
Class	CACS:0A742B860000003B0231777D:atw-cp-ise02/210675310/23919
cisco-av-pair	url-redirect-acl=ACL-WEBAUTH-REDIRECT
cisco-av-pair	url-redirect=https://atw-cp-ise02.ise.local:8443/guestportal /gateway?sessionId=0A742B860000003B0231777D&action=cwa
cisco-av-pair	ACS:CiscoSecure-Defined-ACL=#ACSACL#-IP-Pre-WebAuth-54bb3896
cisco-av-pair	cts:security-group-tag=0008-0
cisco-av-pair	profile-name=OS_X_Lion-Workstation

Figure 14-51 ISE Authentication Details Live Log—Wired WebAuth results.

- **From ISE to switch**—The output in the Result section of the Authentication Live Log is as follows:
  - **URL Redirect**—The redirect URL is contained within one of the cisco-av-pair outputs, specifically the one that begins with url-redirect=. As you look at the rest of the URL, you will notice that there is a parameter called action=cwa. As the endpoint is redirected to that URL, ISE recognizes this parameter as a CWA call and provides the endpoint the correct webpage.
  - **URL Redirect ACL**—The redirect ACL is also contained within one of the cisco-av-pair outputs in the Result section. In this case, the cisco-av-pair will begin with url-redirect-acl= (notice the acl at the end), specifically mentioning the ACL-WEBAUTH-REDIRECT ACL.
  - **ACS ACL**—The DACL that is sent from ISE is also one of the cisco-av-pair outputs—the

https://t.me/learningnets

one with the attribute of ACS:CiscoSecure-Defined-ACL. The value, #ACSACL#-IP-Pre-WebAuth-54bb3896, is unique on a per-user basis.

- **Security Group Tag**—The SGT that was assigned to the endpoint can also be found in the Result section of the authentication details. Again, this value will be embedded in a cisco-av-pair—namely, the AV pair of cts:security-group-tag=0008-0.

Supplementing this authorization output, ISE will also create a separate log in the Authentication Live Log for the DACL that is pushed to the NAD on behalf of the endpoint. Looking at the authentication details for this DACL (see Figure 14-52), focusing on the Results section, you will see the individual lines of the ACL that is pushed to the NAD.

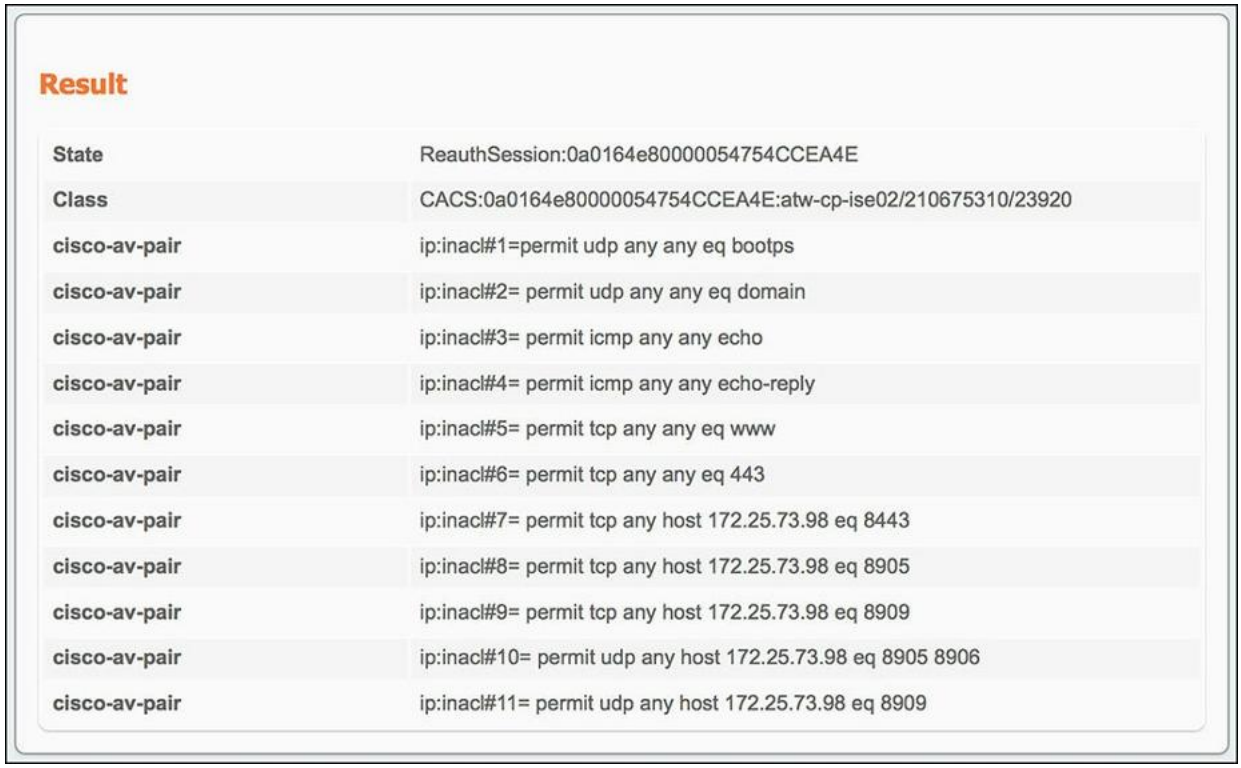


Figure 14-52 ISE Authentication Details Live Log—Wired WebAuth—Pre-WebAuth DACL.

6. **ISE gets the final authentication request and sends the switch the Internet_Only authorization profile and SGA policy GUEST**—When the guest user authenticates to the web portal, a CoA will be sent to the switch on her behalf, forcing the endpoint to reauthorize. If the guest user's web authentication is successful, the user should be given the authorization policy of Internet_Only and a SGA policy GUEST. The conditions that must match for this to happen are Guest user identity store and Network Access:UseCase = Guest Flow. This final authentication request should result in the following communications:

   - **From switch to ISE**—Looking at the authentication details (shown in Figure 14-53) for the endpoint in question, we will focus on the Overview and Authentication Details section of the output.

https://t.me/learningnets

Overview

Event	5236 Authorize-Only succeeded
Username	~68j1r6oz
Endpoint Id	00:17:F2:CB:5C:34
Endpoint Profile	OS_X_Lion-Workstation
Authorization Profile	Internet_Only,GUEST
AuthorizationPolicyMatchedRule	Guest
ISEPolicySetName	Default

Authentication Details

Source Timestamp	2015-01-31 14:58:19.322
Received Timestamp	2015-01-31 14:58:19.323
Policy Server	atw-cp-ise02
Event	5236 Authorize-Only succeeded
Failure Reason	
Resolution	
Root cause	
Username	~68j1r6oz
User Type	GuestUser
Endpoint Id	00:17:F2:CB:5C:34
Endpoint Profile	OS_X_Lion-Workstation
IP Address	10.116.43.138
Identity Store	Guest Users
Identity Group	User Identity Groups:Guest,Workstation
Audit Session Id	0A742B860000003B0231777D
Authentication Method	Authorize Only
Authentication Protocol	
Service Type	Authorize Only
Network Device	KR-3750
Device Type	Switches#Access-Layer
Location	NorthAmerica#CLT
NAS IP Address	10.116.43.134
NAS Port Id	GigabitEthernet1/0/3
NAS Port Type	Ethernet
Authorization Profile	Internet_Only,GUEST
Posture Status	NotApplicable
Security Group	GUEST
Response Time	13

https://t.me/learningnets

Figure 14-53 ISE Authentication Details Live Log—Wired Internet_Only.

- **Authorization Profile**—Under the Overview, you will get the condensed version of the authorization output. Looking at the Authorization Profile, you will see Internet_Only and GUEST. These are the authorization profile and SGA policy that we should be seeing.
- **AuthorizationPolicyMatchedRule**—In the case of a successful guest user web authentication and reauthorization, this rule should be Guest.
- **Endpoint ID**—When looking at the authentication details for a particular endpoint, ensure that you are looking at the MAC address of the correct endpoint.
- **Authentication Method**—With a CoA, there will not be a need to authenticate. Therefore, this output should be Authorize Only.
- **Identity Group**—As the guest user is a part of the Guest identity group, this list of identity groups should contain User Identity Groups:Guest.
- **Use Case**—The other condition, besides the Guest identity group, is GuestFlow, which is equal to Network Access:UseCase EQUALS Guest Flow. Under Other Attributes, you will see UseCase equals Guest Flow.

Because the Guest rule has the two conditions of Guest identity group and GuestFlow—and because both of these conditions are matched—the Authorization profile of Internet_Only and SGA policy of GUEST must be pushed back to the switch. The information in this authorization profile is provided in the Result section of the authentication details. Looking at this output (see Figure 14-54), we see:

- **From ISE to Switch**—Looking at the output in the Result section of the Authentication Live Log:
  - **ACS ACL**—The DACL that is sent from ISE is also one of the cisco-av-pair outputs—the one with the attribute of ACS:CiscoSecure-Defined-ACL. The value, #ACSACL#-IP-Internet_Only-54cb00fc, again, is unique on a per-user basis.
  - **Security Group Tag**—The SGT that was assigned to the endpoint can also be found in the Result section of the authentication details. Again, this value will be embedded in a cisco-av-pair—namely the AV pair of cts:security-group-tag=0008-0. This is the same SGT value that was pushed in the WebAuth scenario.

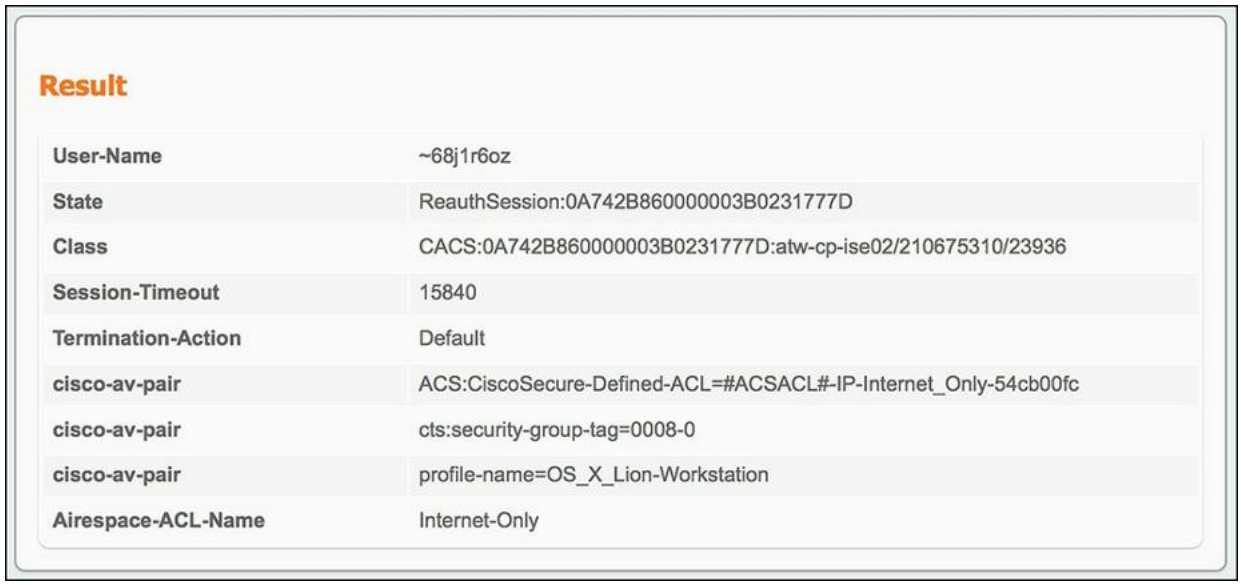


Figure 14-54 ISE Authentication Details Live Log—Wired Internet_Only Results.

Again, ISE will also create a separate log in the Authentication Live Log for the DACL that is pushed to the NAD on behalf of the endpoint. Looking at the authentication details for this DACL (see

https://t.me/learningnets

Figure 14-55), focusing on the Results section, you will see the individual lines of the ACL that is pushed to the NAD.

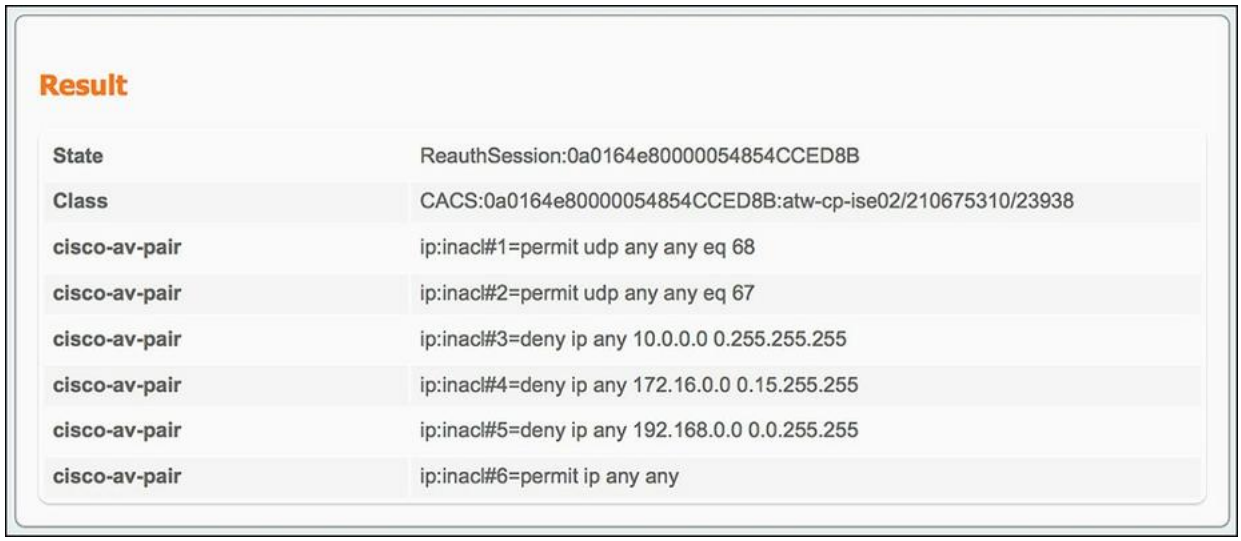
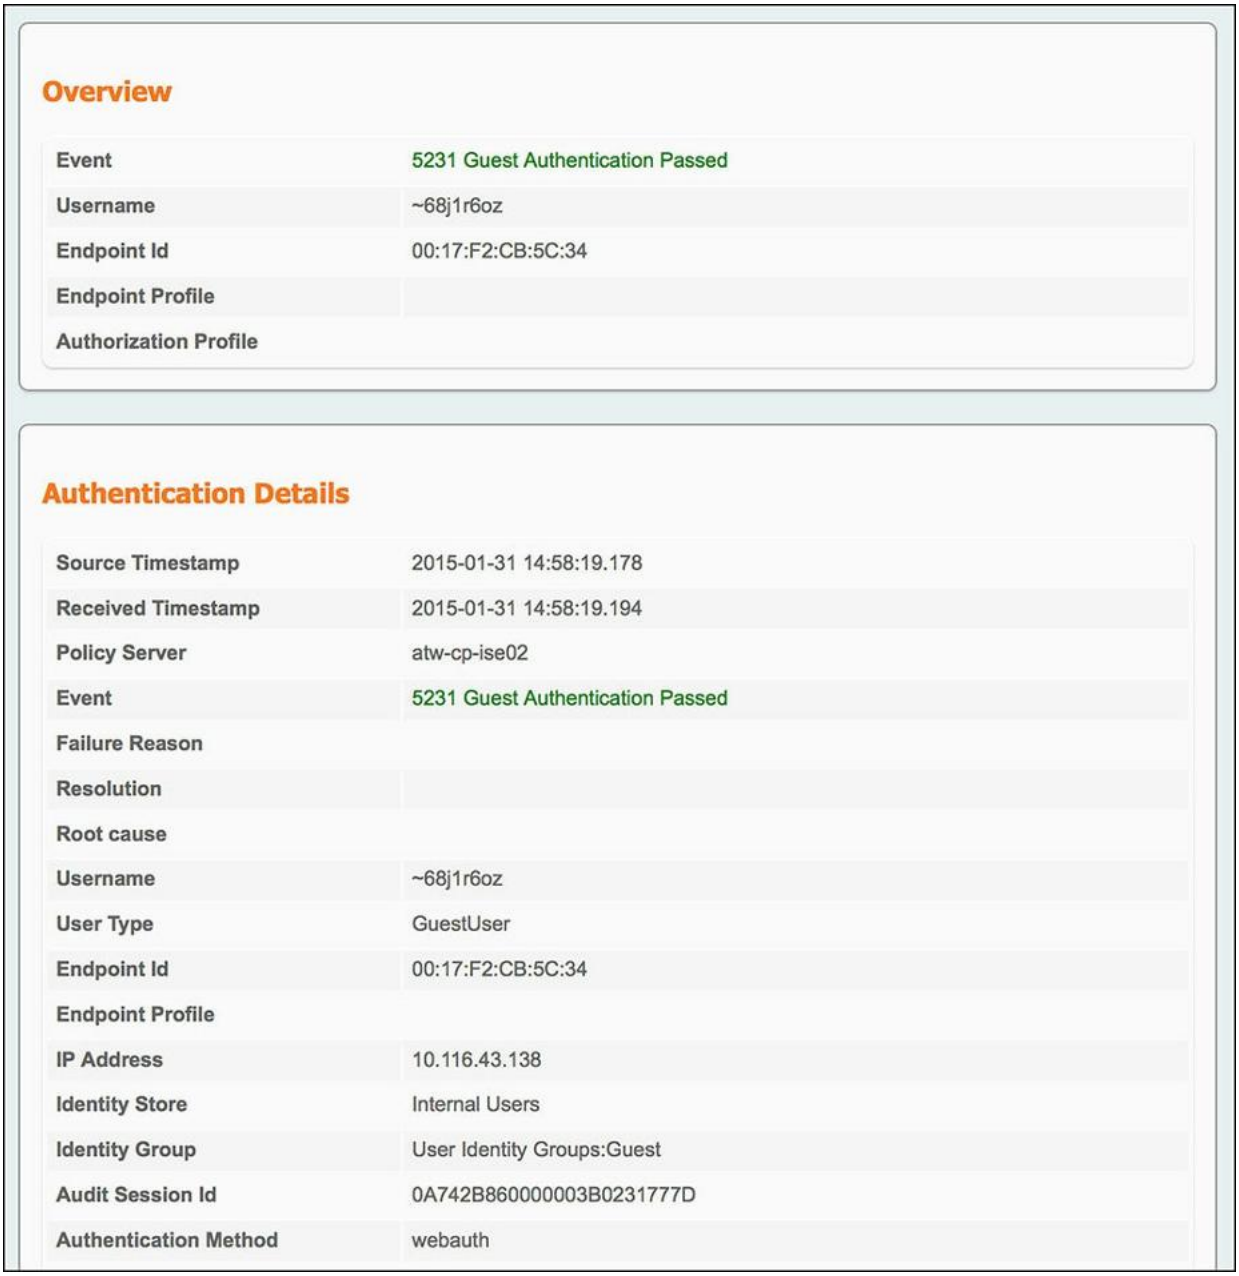


Figure 14-55 ISE Authentication Details Live Log—Wired Internet Only—Internet_Only DACL.

In both the WLC and switch use cases, the guest user was sent to a webpage to complete the WebAuth process. This WebAuth process will also create an Authentication Live Log. This allows the administrator to confirm if the user is typing the correct username and/or password during the login process. All the information within this Live Log is maintained locally and doesn't get sent explicitly to the NAD. Let's take a quick look at the key fields in the WebAuth Live Log (see Figure 14-56).

https://t.me/learningnets



Overview

Event	5231 Guest Authentication Passed
Username	~68j1r6oz
Endpoint Id	00:17:F2:CB:5C:34
Endpoint Profile	
Authorization Profile	

Authentication Details

Source Timestamp	2015-01-31 14:58:19.178
Received Timestamp	2015-01-31 14:58:19.194
Policy Server	atw-cp-ise02
Event	5231 Guest Authentication Passed
Failure Reason	
Resolution	
Root cause	
Username	~68j1r6oz
User Type	GuestUser
Endpoint Id	00:17:F2:CB:5C:34
Endpoint Profile	
IP Address	10.116.43.138
Identity Store	Internal Users
Identity Group	User Identity Groups:Guest
Audit Session Id	0A742B860000003B0231777D
Authentication Method	webauth

Figure 14-56 ISE Authentication Details Live Log—WebAuth.

- **Username**—Within the Overview section, you will see the guest user username. Confirm that this username was typed correctly by the user.
- **Endpoint Id**—As before, make sure that you are looking at the correct Live Log for the endpoint using the endpoints MAC address for confirmation.
- **Identity Group**—The Identity Group list in this case should contain User Identity Groups:Guest.
- **Authentication Method**—Because the guest user will be logging in directly to the webpage in this scenario, the Authentication Method should be webauth.
- **PortalName**—In the Other Attributes section of the authentication details, you will see which guest portal was used to process this authentication request.

The Authentication Live Log output is extremely useful when troubleshooting the communication between a NAD and ISE or a user and ISE. In both of these cases, comparing the conditions and result of the authorization policy that the user or endpoint DID hit versus the conditions and result of the authorization that the user or endpoint SHOULD HAVE hit will usually point you toward any policy issues on ISE.

https://t.me/learningnets

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 14-3 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	Portals and ISE personas	389
Table 14-2	Portal URLs	392
Paragraph	Guests versus ActivatedGuest	398
Paragraph	Using a custom web portal within an authorization profile	400
Paragraph	Restrictions when using the Continue option with authentication protocols	402
Paragraph	Use caution when using a Default rule that does not have an ACCESS_REJECT	406
Paragraph	Optional versus required user fields during guest account creation	416

Table 14-3 Key Topics for Chapter 14

Define Key Terms

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

sponsor

https://t.me/learningnets

Chapter 15. Profiling

This chapter covers the following exam 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 a lot of confusion about the difference between profiling and posture, and the terms can get mixed up. In Chapter 6, "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 the system and report the facts. Posture is covered more in Chapter 19, "Posture Assessment."

Profiling is used to build a database of endpoints that are on the network and assign them to a category based on their external attributes. Although profiling started out as a technology used to automate MAC Authentication Bypass (MAB), it has since morphed into an integral part of access control policy.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Enable the profiling services	1
Network probes	1, 3-4, 8
IOS Device Sensor	4, 8
Feed service	5-6, 10
Profiling policy rules	6, 9
Utilize profile assignment in authorization policies	2
Verify profiling operation	7

Table 15-1 "Do I Know This Already?" Section-to-Question Mapping

> **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.

https://t.me/learningnets

1. True or False? The profiling service is enabled by default on ISE policy service nodes.
   a. True
   b. False
2. Name three ways in which an endpoint profile can be used in an authorization policy rule?
   a. Logical profiles
   b. Endpoint identity groups
   c. NMAP OS-Scan result
   d. EndPointPolicy attribute
   e. EndPointProfile attribute
3. Which probe is used to trigger the SNMPQUERY probe to query a NAD?
   a. RADIUS
   b. SNMPQUERY
   c. HTTP
   d. SNMPTRAP
   e. Both A and D
   f. Both C and D
4. Which three probes exist with device sensor?
   a. CDP, DHCP, RADIUS
   b. HTTP, CDP, RADIUS
   c. CDP, DHCP, LLDP
   d. CDP, HTTP, SNMP
5. How are updated profiles distributed to customer ISE deployments?
   a. Cisco's Profiler Feed Service.
   b. Each new version of ISE or ISE patch includes new profile policies.
   c. The profiles are distributed together with the posture checks and compliance modules.
   d. Import the update packs that are downloaded from Cisco.com.
6. What determines when an endpoint is assigned to a profile?
   a. The profile that matches the most conditions will be 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 will identify the profile of an endpoint to ISE.
7. Which ISE tool enables an administrator to drill down in to the profiles that have been assigned to locate a specific endpoint with that profile?
   a. Endpoints Drill-down
   b. Cisco Endpoint Profiling Examination Tool (CEPET)
   c. Profiled Endpoints Counter
   d. Profiler Activity Window
8. What are two ways to collect HTTP user agent strings?
   a. Through the AnyConnect HTTP User Agent Reporting Tool
   b. SPAN port mirroring

https://t.me/learningnets

c. The Cisco WSA device sensor

d. Directly from ISE web portals

e. Device sensor in the switch

9. True or False? ISE deployments must wait for Feed Service updates for new profiles.

a. True

b. False

10. What will happen 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 admin will be prompted to choose to overwrite or ignore the profile update.

c. All nonconflicting profiles will be downloaded and installed. The conflicting profiles will be ignored.

d. The update will fail and an alarm will be triggered on the dashboard and in email.

Foundation Topics

ISE Profiler

The Cisco ISE Profiler is the component of the Cisco Identity Services Engine (ISE) platform that is responsible for endpoint detection and classification. It does so by using a probe or series of probes that collect attributes about an endpoint. The profiler then compares the collected attributes to predefined device signatures to locate a match.

In the early days of identity-based networks and 802.1X, countless hours were spent identifying all the devices that did not have supplicants—in other words, the devices that could not authenticate to the network using 802.1X, such as printers and fax machines. Your choices were to identify the port connected to the printer and configure it to either

- Not use 802.1X
- Use MAC Authentication Bypass (MAB)

MAB is an extension to 802.1X that enables the switch to send the device's MAC address to the authentication server. If that MAC address is on the approved list of devices, the authentication server sends an "accept" result, therefore enabling specific MAC addresses to bypass authentication.

Countless hours were spent collecting and maintaining this list of MAC addresses. A manual onboarding procedure would eventually be created, so that when a new printer (or other 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 and save all those hours of preparation and maintenance.

As described in Chapter 6, this is where profiling technology enters the picture. It enables you to collect attributes about devices from sources such as DHCP, NetFlow, HTTP User-Agent strings, and more. Those collected attributes are then compared to a set of signatures, similar to the way an intrusion prevention system (IPS) works. These signatures are commonly referred to as profiles.

An example is collecting a MAC address that belongs to Apple, seeing that the HTTP User-Agent string contains the word "iPad," and therefore assigning that device to the profile of "Apple-iPad."

Profiling technology has evolved, as technology tends to do. Now your authentication server (ISE) has the ability to use that profiling data for much more than just building the MAB list.

https://t.me/learningnets

Key Topic

Cisco's Identity Services Engine uses the resulting collection and classification data from the profiler as conditions in the authorization policy. Now you can build an authorization policy that looks at much more than your identity credentials. We can combine a user's identity with the classification result and invoke specific authorization results.

Figures 15-1 and 15-2 are examples of a differentiated authorization policy based on profiling.

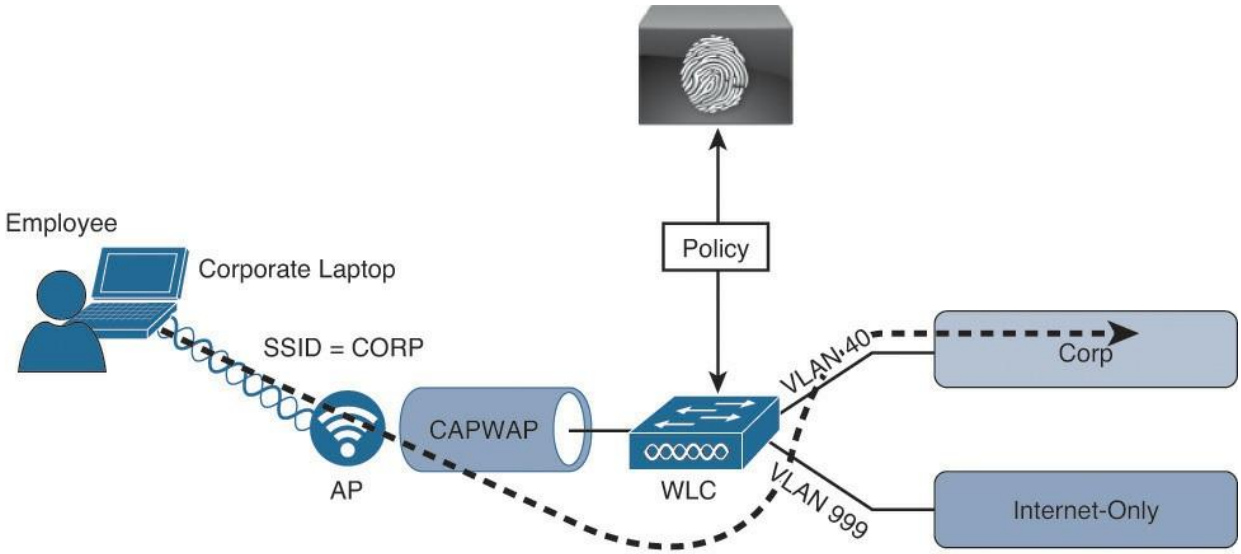


Figure 15-1 Employee using corporate laptop gains full access.

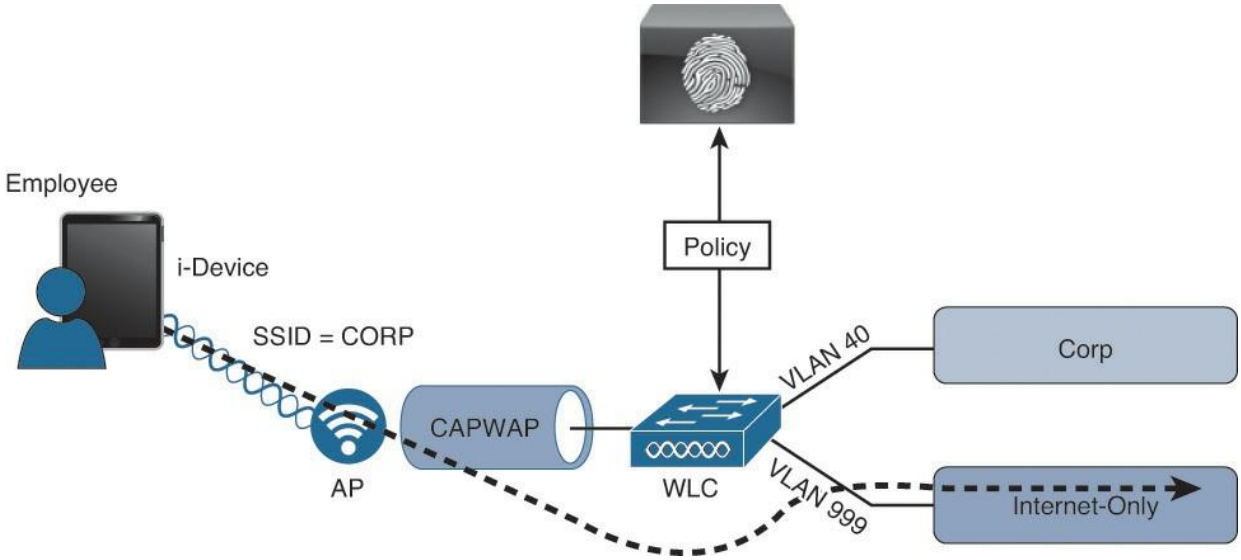


Figure 15-2 The same employee credentials on an i-device will get limited access.

Users who are using the same wireless SSID and the same credentials can be associated to different wired virtual local area network (VLAN) interfaces based on the device profile:

- Any employee using a corporate laptop with her Active Directory user ID is assigned to the corporate VLAN and given full access to the network.
- Any employee using an i-device with her same Active Directory user ID is assigned to a GUEST VLAN and provided Internet access only.

While it can be quite intuitive to visualize the types of network access policies you will be able to create

https://t.me/learningnets

based on the device's profile, the design of where and how the profiler collects the data about the endpoints requires thought and planning.



One of the first questions a security team might ask when discussing profiling with any network access control solutions is "Can we use this as an antispoofing solution?" Remember that MAC Authentication Bypass (MAB) 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 his laptop to use the same MAC address as the printer (spoofing), and gain access to the network.

It is key to always keep in mind that profiling is a technology that compares collected attributes about an endpoint to a set of signatures called profiling policies to make the best guess of what a device is. Can this type of technology be used to prevent spoofing? Sure. However, it is difficult to accomplish antispoofing with this type of technology. It would require a lot of tuning, trial and error, and constant adjustment, which makes it too operationally expensive and untenable.

A best-practice approach is to use a least-privilege strategy instead. If the previously mentioned malicious user is successful in spoofing the MAC address of the printer and gains network access, what level of network access should that device have? In other words, the authorization policy for printers should not provide full network access, but should provide a very limited subset of access instead. So, a printer should be permitted to communicate using only network ports critical to printer operations (such as TCP port 9100 or 9600).

Cisco ISE Probes

As described previously, the Cisco Identity Services Engine solution is capable of providing access policies on which the decisions can be made based: who, what, where, when, how, and more. Profiling is focused on the "what" elements of the policy. For the policy engine to know what the device is, you must first collect that data.

The Cisco Identity Services Engine solution uses a number of collection mechanisms known as probes. The probe is software designed to collect data to be used in a profiling decision. An example of this is the HTTP Probe, which captures HTTP traffic and enables the profiler to examine attributes from the traffic, such as HTTP User-Agent strings.

Probe Configuration

You enable the probes on each policy service node (PSN) where appropriate. In the Administration GUI of the PSN, navigate to **Administration > System > Deployment**. From here, select the PSN for which you are configuring the probes. You must repeat these steps for each PSN in your deployment:

Step 1. Select one of the PSNs, as shown in Figure 15-3. In this case, the node is standalone, which means it is a single node running all personas (Admin, Monitoring, and Policy Service).

https://t.me/learningnets

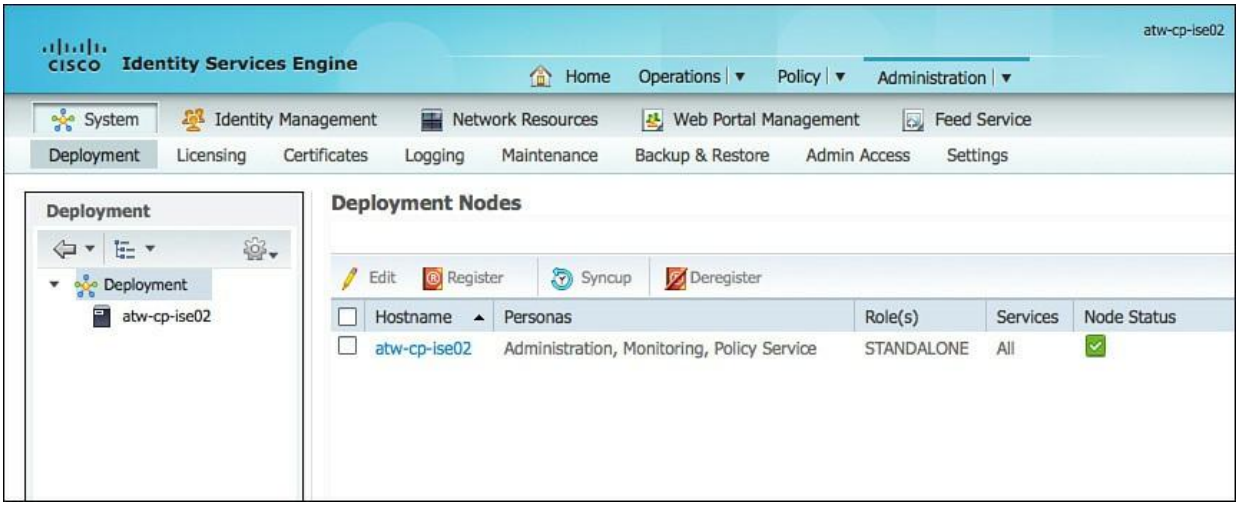


Figure 15-3 ISE Deployment screen.

Step 2. Ensure the **Enable Profiling Service** check box is selected. This service is enabled by default on all PSNs and is not configurable when in standalone mode, as shown in Figure 15-4.

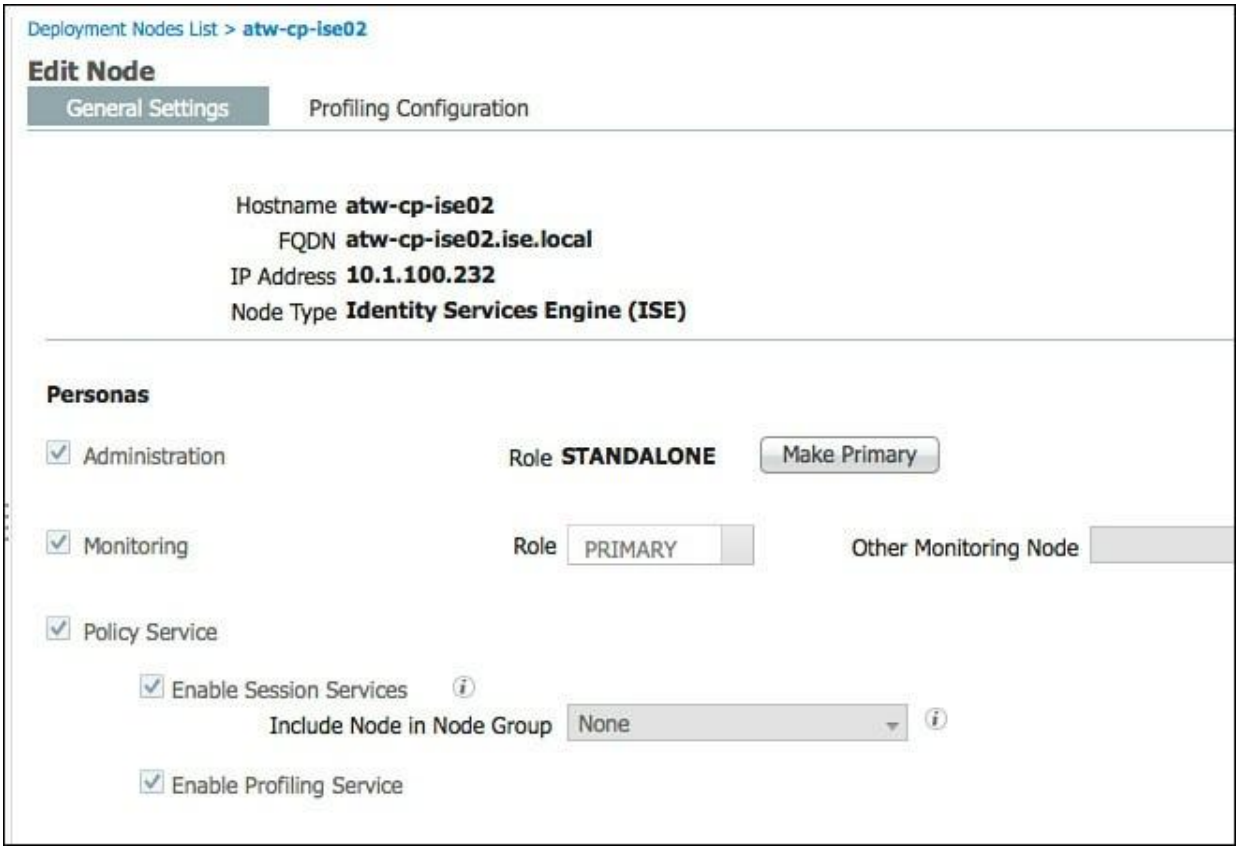


Figure 15-4 General settings.

Step 3. Select the Profiling Configuration tab, as shown in Figure 15-5.

https://t.me/learningnets

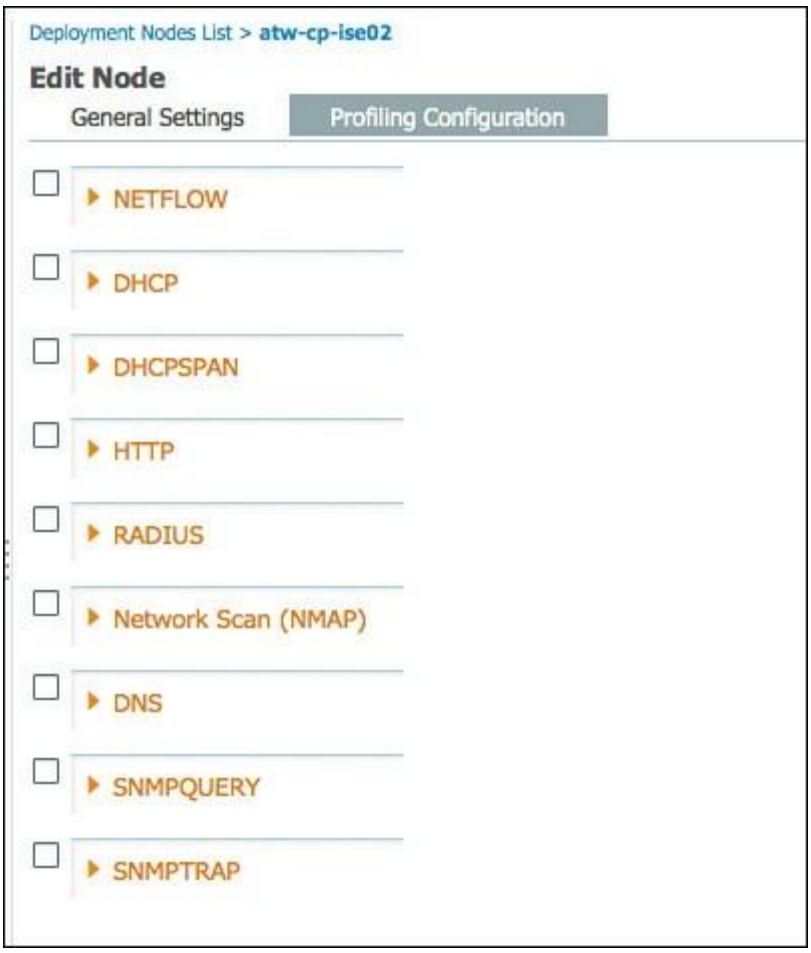


Figure 15-5 Profiling configuration.

There are nine probes on each PSN:

- NETFLOW
- DHCP
- DHCPSPAN
- HTTP
- RADIUS
- Network Scan (NMAP)
- DNS
- SNMPQUERY
- SNMPTRAP

Each probe will be examined in detail, but not in order.

DHCP and DHCPSPAN

DHCP can be one of the most useful data sources for an endpoint device. A primary use of DHCP in profiling is to capture the device MAC address; however, there are many other uses for the data. Much like HTTP, DHCP requests also carry a User-Agent field that helps to identify the operating system (OS) of the device. Some organizations have been known to use a custom DHCP User-Agent string, which helps to identify the device as a corporate asset.

Not only the populated fields from the DHCP client are useful in classification, but other attributes, such as requested DHCP options and the DHCP hostname, can also be helpful in further classifying the device.

https://t.me/learningnets

Unlike HTTP, there are two DHCP probes, each working in a slightly different way: DHCP and DHCP SPAN.

DHCP

The DHCP probe requires the DHCP requests to be sent directly to the ISE PSN(s). This is often done by using the **ip helper-address** interface configuration command to forward the request to both the DHCP server and the PSN, as illustrated in Figure 15-6.

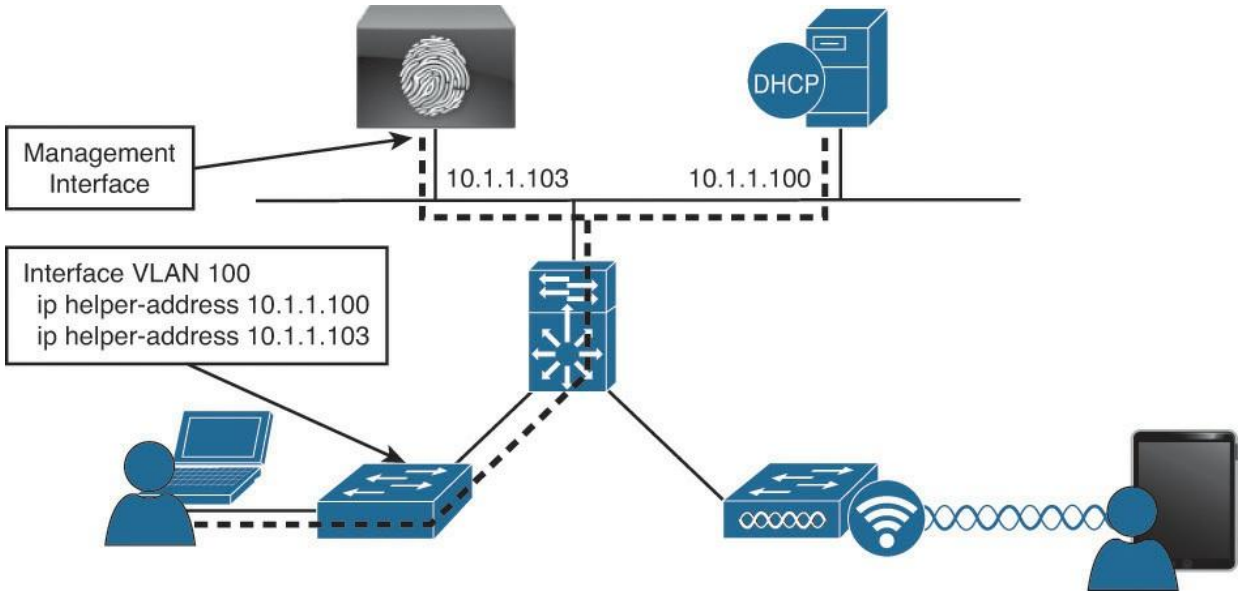


Figure 15-6 DHCP with ip helper-address logical design.

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 is sending the broadcast to all hosts on a specific subnet). Simply add the IP address of your PSN(s) to the list of helper addresses, and it will be copied on all DHCP requests. Not to worry, though—the Cisco Identity Services Engine server will never respond to that request. It will only use the incoming data to profile endpoints.

DHCPSPAN

Another way for ISE to glean the DHCP requests and even the DHCP responses is the use of Switched Port Analyzer (SPAN) session in true promiscuous mode. A SPAN session copies all traffic to/from a source interface on a switch to a destination interface, which would be one of ISE's interfaces assigned to the DHCPSPAN probe. The logical design of using SPAN is illustrated in Figure 15-7.

https://t.me/learningnets

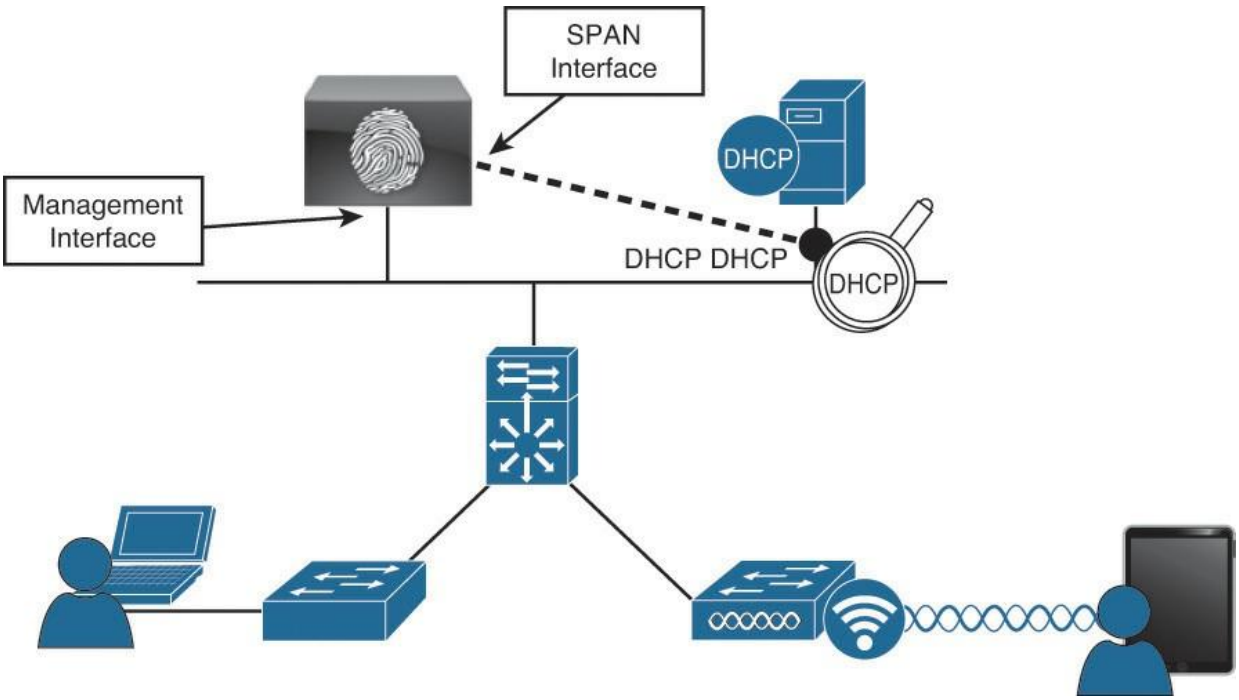


Figure 15-7 DHCP SPAN logical design.

When using the SPAN method, you need to consider where the best location is to create the SPAN session and gather the data. One recommended location is the DHCP server, where the DHCP probe will see both ends of the conversation (request and response). However, there are caveats to this method, such as: "What if the organization uses distributed DHCP servers?" This is why the non-SPAN method tends to be the most commonly deployed.

Considerations with the Cisco WLC

It is important to note that regardless of the SPAN or helper address methods of using the DHCP probe(s), when using a Wireless LAN Controller (WLC), the WLC has a default configuration of acting as a RADIUS proxy, which is its own form of a helper address where the WLC acts as a middle man for all DHCP transactions. Unfortunately, this behavior will have a negative effect on the DHCP probe and must be disabled on the WLC. Upon doing so, the DHCP requests from wireless endpoints will appear as broadcast messages on the VLAN, and an **ip helper-address** statement should be configured on the Layer-3 interface of that VLAN (the switch or router).

Probe Configuration

Minimal configuration is required on the ISE side to enable the DHCP probe(s). From the Profiling Configuration tab displayed in Figure 15-5, do the following:

Step 1. Click the check box next to the DHCP and/or DHCP SPAN probe to enable them.

Step 2. Select either the Gigabit 0 interface or all interfaces. Multiple interfaces cannot be individually selected. The choices will be a single interface or all interfaces.

Figure 15-8 shows the DHCP probes. You should never need to enable both probes for the same interface. That would cause double processing of DHCP packets and be somewhat wasteful of system resources.

https://t.me/learningnets

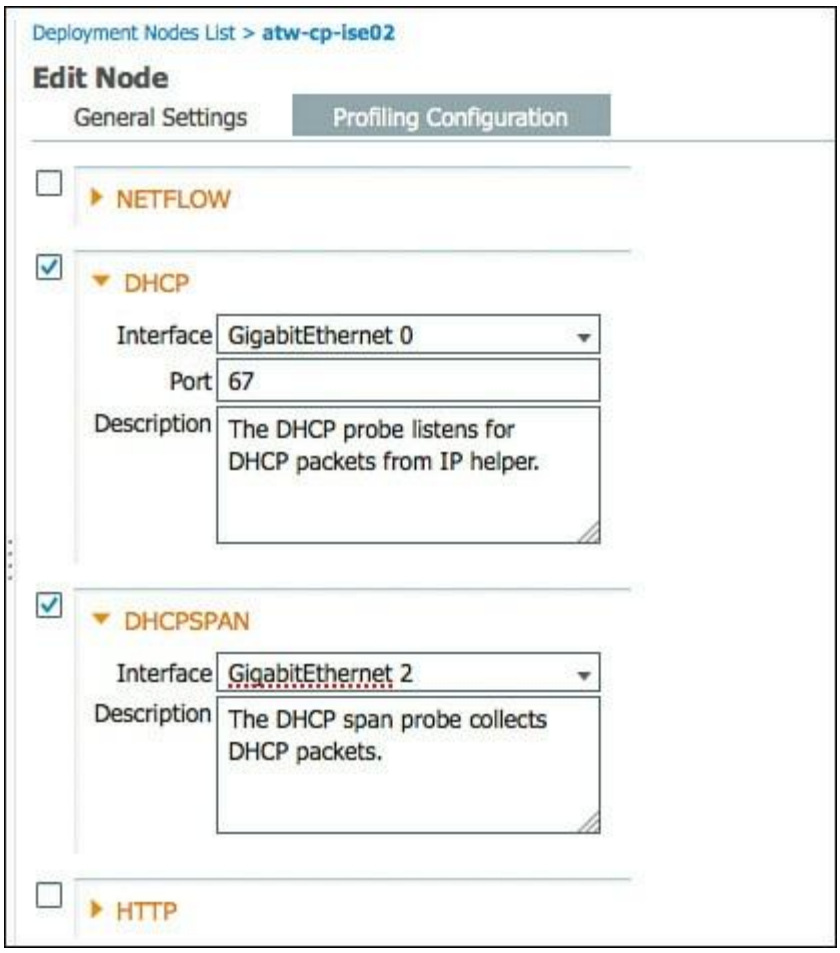


Figure 15-8 DHCP probes.

> **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). There is useful data to help you classify a device that exists within RADIUS communication.

Originally, this was focused on the MAC address and IP address of the device. By having this data conveyed in the RADIUS packet, ISE is capable of building the all-important MAC-to-IP address bindings. Because the endpoint database uses MAC addresses as the unique identifier for all endpoints, these bindings are absolutely critical. Without them, the Layer-3 probes such as HTTP and NMAP scanning would never work correctly.

The Calling-Station-ID field in the RADIUS packet provides the endpoint's MAC address, and the Framed-IP-Address field provides its IP address in the RADIUS accounting packet.

Additionally, the RADIUS probe triggers the SNMPQUERY probe to poll the NAD (see the following SNMP probe information).



https://t.me/learningnets

Most importantly, with the proliferation of device sensor-capable switches and wireless controllers, the RADIUS probe becomes even more critical. Device sensor is a feature in the switch or controller that collects endpoint attributes locally and then sends those attributes to ISE within RADIUS accounting packets.

By allowing the network device to proactively send the 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. Additionally, it 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 (see the section "SNMPQUERY").

Considerations with RADIUS Probe

All NADs in the Secure Unified Access Deployment should be configured to send RADIUS accounting packets. It is also important to note that the Cisco switch must learn the endpoint's IP address via DHCP snooping or **ip-device-tracking** to fill in the Framed-IP-Address field.

Probe Configuration

Minimal configuration is required on the ISE side to enable the RADIUS probe(s). From the Profiling Configuration tab displayed in Figure 15-5, just click the check box next to the RADIUS to enable the RADIUS probe, as shown in Figure 15-9.

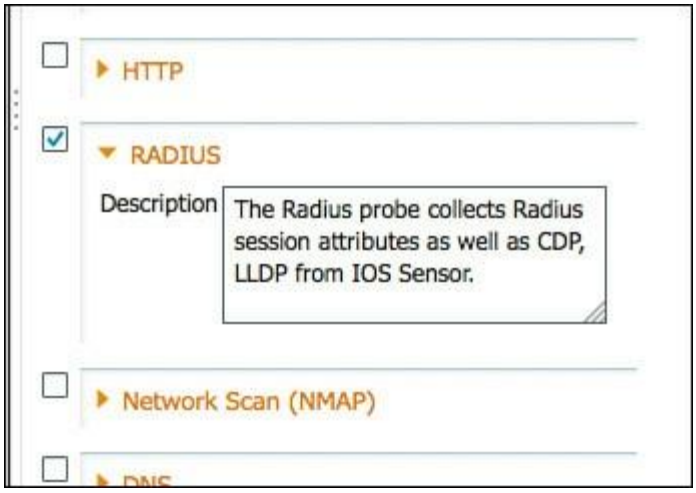


Figure 15-9 RADIUS probe.

Notice that configuration really isn't possible with this probe. However, it is one of the most useful probes, especially when combined with device sensor.

Network Scan

A very welcome improvement to ISE version 1.1 is the addition of the endpoint scanning (NMAP) probe. NMAP is a tool that uses port scans, SNMP, and other mechanisms to identity a device's OS or other attributes of the device. The NMAP probe can be manually run against a single IP address or subnet. More importantly, the profiler engine can be configured to react to a profiling event with a reactive NMAP probe.

For example, when an endpoint is discovered to be an "Apple-Device," ISE automatically launches an NMAP OS scan against that endpoint to determine whether it is running MAC OSX or iOS. From the results of that scan, ISE further classifies the device as a MAC device or an i-device.

Considerations with the NMAP Probe

The Endpoint Scanning (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. As such the Policy Services Node will rely on the MAC Address-to-IP Address binding to update an endpoint's attributes with the results of the NMAP scan. Therefore it is

https://t.me/learningnets

absolutely critical that the PSN had received valid information from the other probes.

The NMAP probe can be manually run against a single IP-Address or subnet, or (more commonly) an NMAP scan can be triggered as an action of a profile.

Probe Configuration:

As with all the other probes, only minimal configuration is needed in the ISE GUI. From the Profiling Configuration tab displayed in Figure 15-5, just click the check box next to the Network Scan (NMAP) probe to enable it. You can also run a manual scan from this screen against a single node, or an entire network, as shown in Figure 15-10.

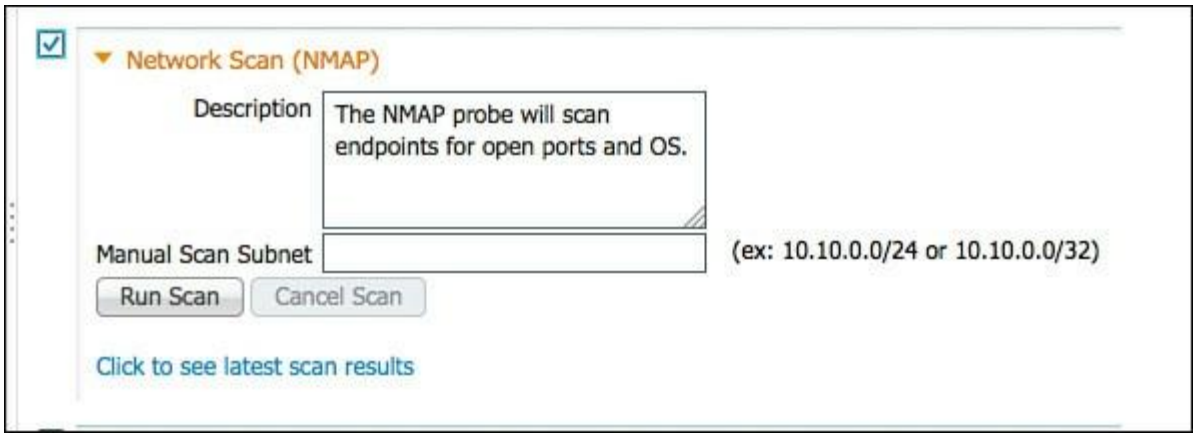


Figure 15-10 Network Scan (NMAP) probe.

DNS

The DNS probe is used to collect the fully qualified domain name (FQDN) of an endpoint using a reverse-lookup for the static or dynamic DNS registration of that endpoint. It is quite useful when looking for a specific DNS name format of corporate assets (Active Directory members).

A reverse DNS lookup will be completed only when an endpoint is detected by one of the DHCP, RADIUS, HTTP, and SNMP probes.

To enable the DNS probe, simply click the check box next to the DNS probe to enable it, as shown in Figure 15-11. This probe will use the name server configuration from the Identity Services Engine node itself.

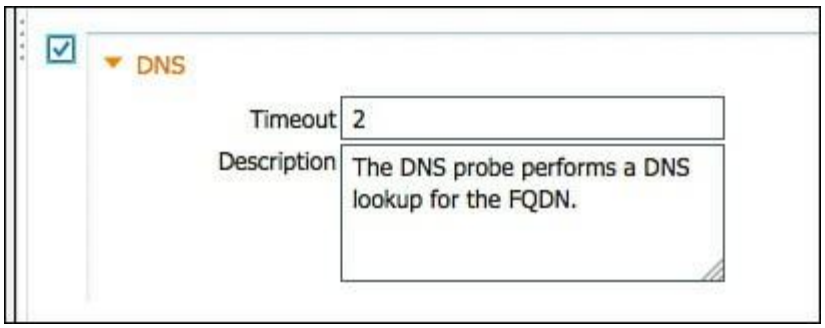


Figure 15-11 DNS probe.

SNMPQUERY and SNMPTRAP

SNMP is used to query NADs that do not yet support Cisco's device sensor. After enabling the SNMPQUERY probe, ISE will poll all the SNMP-enabled NADs at the configured polling interval.

https://t.me/learningnets

> **Note**
>
> It is recommended to remove SNMP settings from NADs that support IOS sensor to avoid double work and wasted processing.

There are two SNMP probes: SNMPQUERY and SNMPTRAP.

SNMPTRAP

The SNMP Trap receives information from the configured NAD(s) that support 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 enable the SNMPQUERY probe to reactively query a NAD instead of waiting for the periodic polling interval. Therefore, for SNMPTRAP to be functional, you must also enable the SNMPQUERY probe.

The SNMPTRAP probe receives information from the specific NAD(s) when the MAC address table changes or when link state changes on a switch port. To make this feature functional, you must configure the NAD to send SNMP traps or informs.

SNMPQUERY

SNMPQUERY does the bulk of the work. There are three kinds of SNMPQUERY probes:

- **System**—The System probe polls all NADs that are configured for SNMP at the configured interval.
- **Interface**—The Interface probe occurs in response to an SNMPTRAP or RADIUS Accounting start packet (only if the SNMPTRAP probe is enabled).
- **Network Scan**—The Network Scan (NMAP) probe triggers the SNMP walk of an endpoint.

When querying a NAD, ISE looks for interface data (which interface, which VLAN, and so on), session data (if the interface is Ethernet), Cisco Discovery Protocol (CDP) data, and Link Layer Discovery Protocol (LLDP) data. The CDP and LLDP data can be useful in identifying a device type by its registered capabilities and similar attributes.

> **Note**
>
> For distributed deployments, NAD polling is distributed among all PSNs enabled for SNMPQUERY probes.

Probe Configuration

When configuration to these probes is necessary, such as the trap types to examine and the SNMP port, it is recommended to leave these at their default settings unless directed otherwise by Cisco TAC. Do the following:

Step 1. Click the check box next to the SNMPQUERY and SNMPTRAP probes to enable them.

Step 2. For the SNMPTRAP probe, select either the Gigabit 0 interface or all interfaces. Multiple interfaces cannot be individually selected. The choices will be a single interface or all interfaces.

Figure 15-12 shows the enabling of the SNMP probes and their default settings.

https://t.me/learningnets

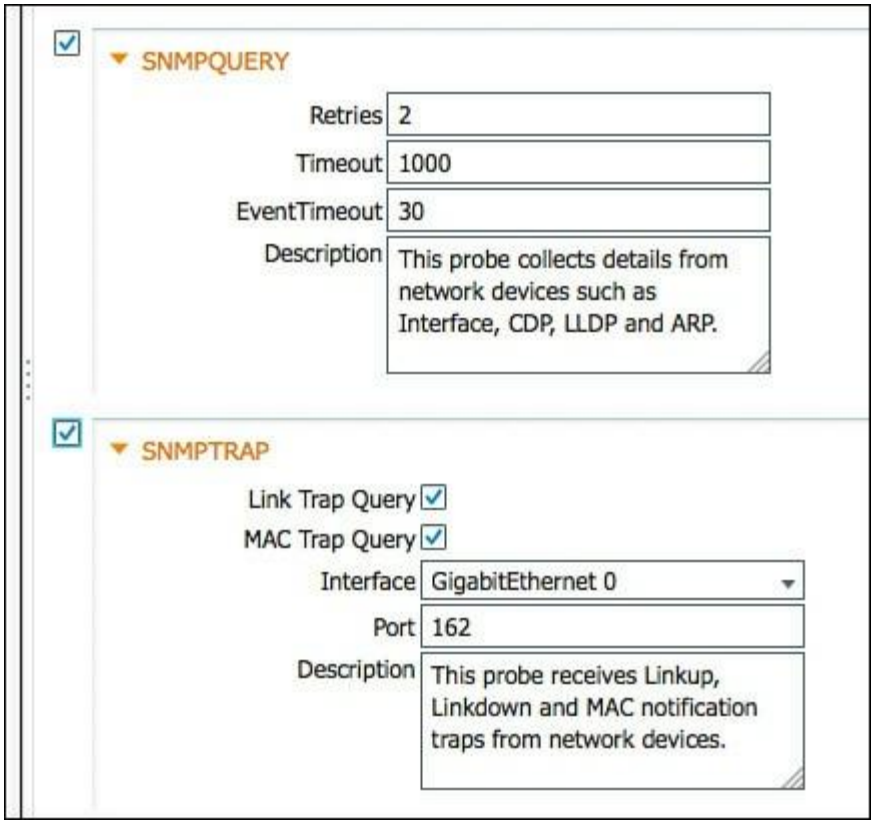


Figure 15-12 SNMP probes.

NETFLOW

NetFlow is an incredibly useful and undervalued security tool. Essentially, it is similar to a phone bill. A phone bill does not include recordings of all the conversations you have had in their entirety; instead it is a summary record of all calls sent and received.

Cisco routers and switches support NetFlow, sending a "record" of each packet that has been routed, including the ports and other useful information.

Just enabling NetFlow in your infrastructure and forwarding it all to the ISE can quickly oversubscribe your PSN. If you are planning to use the NetFlow probe, it is highly recommended that you have a third-party solution, such as LanCope StealthWatch, filter out any unnecessary data and only send what you truly need to ISE. For that reason, this probe is not focused on heavily in this exam or book, and it is recommended to perform extensive planning prior to its use.

Configuring the NetFlow probe is limited to enabling the check box next to the NetFlow probe and selecting either the Gigabit 0 interface or all interfaces. Figure 15-13 shows the enabled NetFLow probe.

https://t.me/learningnets

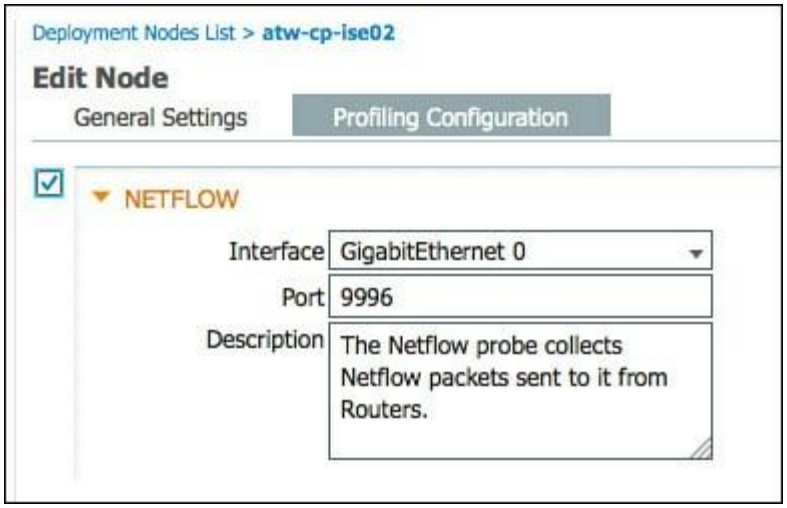


Figure 15-13 NetFlow probe.

HTTP Probe

When applications use HTTP, such as a web browser or even software such as Microsoft Outlook and Windows Update, it typically identifies itself and 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.

The Cisco Identity Services Engine uses the information in HTTP packets; especially the User-Agent field, to help match signatures of what profile a device belongs in.

The two primary mechanisms for the HTTP probe to collect the HTTP traffic are as follows (see Figure 15-14):

- Use a Switched Port Analyzer (SPAN) session in true promiscuous mode. When using the SPAN method, you need to consider where the best place is 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.
- Use a SPAN session in conjunction with a filter to limit the traffic visible to ISE. Another option to use with the SPAN design is the use of VLAN ACLs (VACLs) on a Catalyst 6500 or ACL-based SPAN sessions on a Nexus 7000. These options enable you to build an ACL that defines exactly what traffic you want to capture and send along to ISE—instead of a pure promiscuous SPAN, where the ISE interface will see all traffic. This is a better way to manage the resource utilization on your ISE server when available.

https://t.me/learningnets

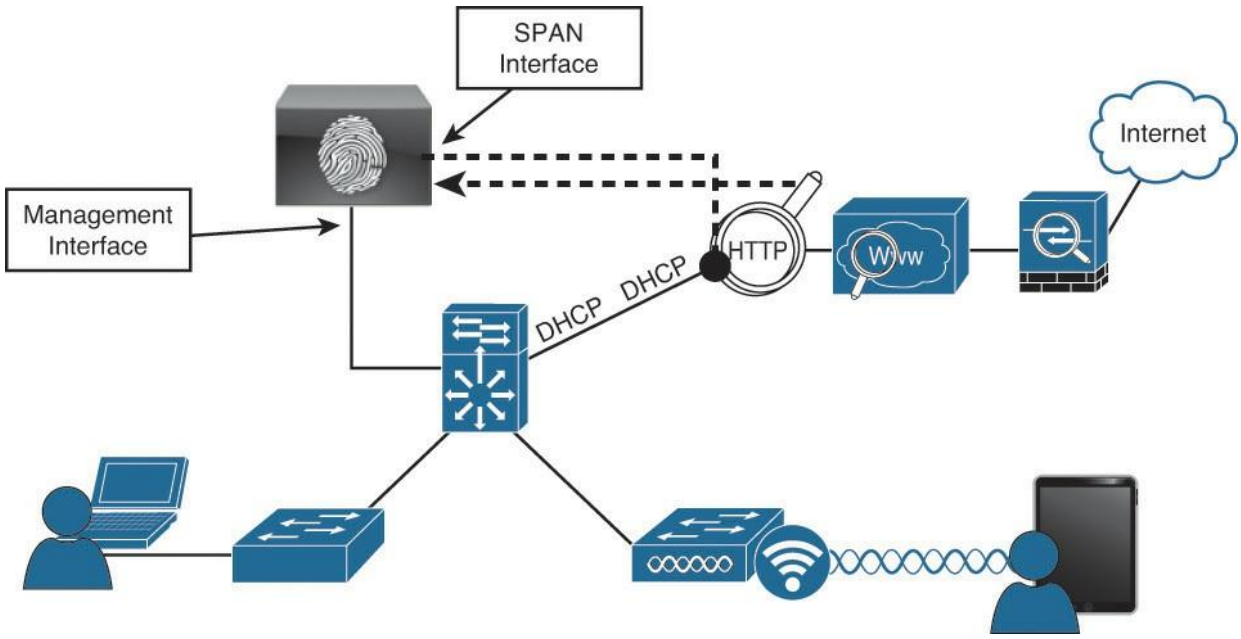


Figure 15-14 HTTP SPAN logical design.

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 with 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 check box next to the HTTP probe, as shown in Figure 10-9. Select either the Gigabit 0 interface or all interfaces.

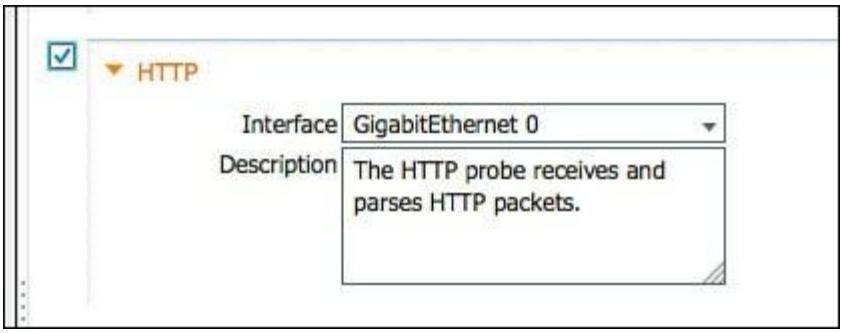


Figure 15-15 HTTP probe.

HTTP Profiling Without Probes

The Web Portal system within 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 profiling being enabled or not. The user agent is used to know which operating system is connecting and therefore which agent or client to send to the endpoint (in the cases of client and native supplicant provisioning).

When any portal collects that user agent, it is automatically passed over to the profiling engine within ISE, without requiring the HTTP probe to be enabled. It is a simple and efficient way to get the extremely valuable user agent string without having to rely on the computationally expensive SPAN methods.

Infrastructure Configuration

As an overall best practice, it is recommended to examine the cost benefit analysis of using processor-intensive probes or probe designs. For example, it is often recommended to use DHCP Helper instead of configuring a SPAN session and examining a multitude of traffic that might or might not be relevant.

Use HTTP traffic as an example. HTTP traffic is extremely useful for identifying the OS on a client

https://t.me/learningnets

endpoint; however, HTTP SPAN can consume a large amount of system resources on the PSN. Additionally, it might 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 VACL captures that can limit which traffic is sent to the SPAN session. Other deployments use the Authorization policy in ISE to send unknown devices to an ISE portal, enabling the portal to update the profiling data.

DHCP Helper

As previously shown in Figure 15-6, the **ip helper-address** commands are configured on the default gateway for each of the Access Layer VLANs. To configure the destination address to which to copy DHCP requests, enter interface configuration mode of the Layer-3 address for the VLAN and enter the **ip helper-address [ip-address]** command. You then must add the DHCP server and all applicable ISE PSNs to the list of helper address destinations.

Figure 15-16 shows sample output from a Layer-3 interface that is configured to send DHCP requests to the DHCP server in addition to two different ISE PSNs.

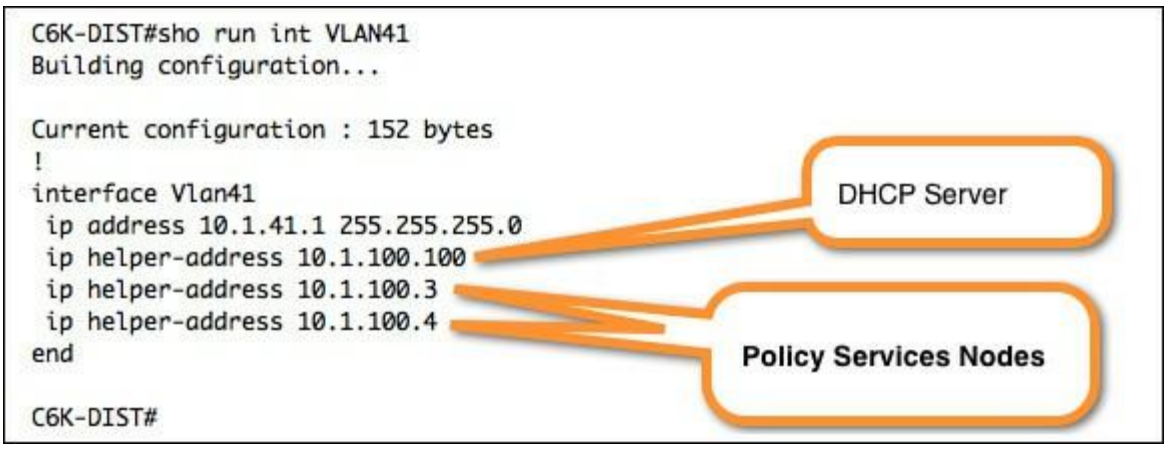


Figure 15-16 ip helper-address Settings.

SPAN Configuration

A monitor session is configured in global configuration mode and can be local (SPAN) or remote (RSPAN). Example 15-1 shows a SPAN configuration where an Internet-facing VLAN is the source of the session and an interface on the PSN is the destination. For more on SPAN configuration, see http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a008015c612.shtml.

Example 15-1 Monitor Session Command Input

Click here to view code image

```
DC-4948(config)#monitor session [1-4] source [interface | vlan] [rx | tx]
DC-4948(config)#monitor session [1-4] destination interface [interface_name]
```

Figure 15-17 shows the output of the **show monitor** command, where you can see the source and destination of the session.

https://t.me/learningnets

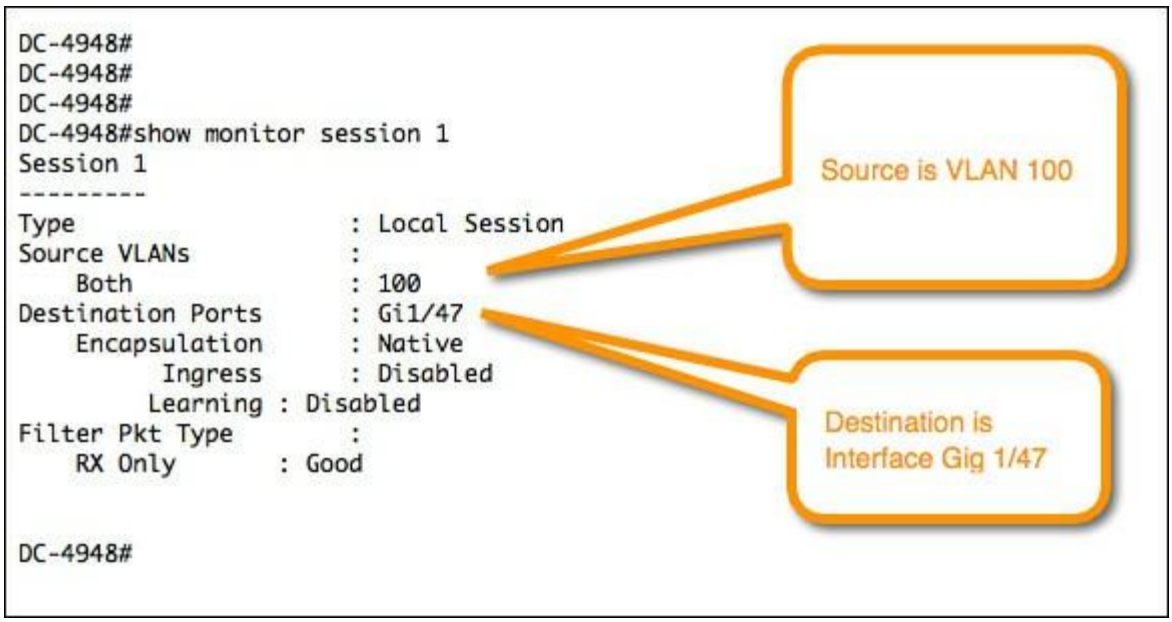


Figure 15-17 Sample monitor session (SPAN) configuration.

VLAN Access Control Lists

VACL configuration is a multistep process. It requires an access list to specify which traffic should be captured and another 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 finally you define the destination, like so:

Step 1. Build an access list to classify the traffic you want to capture, such as the one shown in Example 15-2.

Example 15-2 Creating an Access List to Match HTTP Traffic

Click here to view code image

```
C6K-DIST(config)#ip access-list extended HTTP_TRAFFIC
C6K-DIST(config-ext-nacl)#permit tcp any any eq www
```

Step 2. Build an access list for all the rest of the traffic, such as the one shown in Example 15-3.

Example 15-3 Creating an Access List to Match All Other Traffic

Click here to view code image

```
C6K-DIST(config)#ip access-list extended ALL_TRAFFIC
C6K-DIST(config-ext-nacl)#permit ip any any
```

Step 3. Create a VLAN access map sequence to capture HTTP traffic, such as what's shown in Example 15-4.

Example 15-4 Creating a VLAN Access Map

Click here to view code image

```
C6K-DIST(config)#vlan access-map HTTP_MAP 10
C6K-DIST(config-access-map)#match ip address HTTP_TRAFFIC
```

https://t.me/learningnets

```
C6K-DIST(config-access-map)#action forward capture
```

Step 4. Add a new sequence to the access map to forward all other traffic, such as what's shown in Example 15-5.

Example 15-5 Adding Another Sequence to the Access Map

Click here to view code image

```
C6K-DIST(config)#vlan access-map HTTP_MAP 20
C6K-DIST(config-access-map)#match ip address ALL_TRAFFIC
C6K-DIST(config-access-map)#action forward
```

Step 5. Apply the VLAN access map to the VLAN list, as shown in Example 15-6.

Example 15-6 Applying the VLAN Access Map to the VLAN(s)

Click here to view code image

```
C6K-DIST(config)#vlan filter HTTP_MAP vlan-list 41,42
```

Step 6. Configure the destination port for the PSN's SPAN interface, as shown in Example 15-7.

Example 15-7 Setting the Switchport Capture Destination

Click here to view code image

```
C6K-DIST(config)# interface gig1/0/1
C6K-DIST(config-if)#switchport capture allowed vlan 41
C6K-DIST(config-if)#switchport capture allowed vlan add 42
C6K-DIST(config-if)#switchport capture
```

Device Sensor

IOS device sensor requires a multipart configuration. The first part is to configure the device sensor filter lists. These lists inform device sensor of which items to consider important for each of the protocols.

The IOS device sensor supports three protocols: DHCP, CDP, and LLDP. Therefore, you must create one list for each protocol:

Step 1. Create a list for DHCP. You will need to configure three DHCP options for ISE: **host-name**, **class-identifier**, and **client-identifier**, as shown in Example 15-8.

Example 15-8 DHCP Filter List for Device Sensor

Click here to view code image

```
C3750X(config)#device-sensor filter-list dhcp list <dhcp_list_name>
C3750X(config-sensor-dhcplist)#option name host-name
C3750X(config-sensor-dhcplist)#option name class-identifier
C3750X(config-sensor-dhcplist)#option name client-identifier
```

Step 2. Create a list for CDP.

You will need to configure two CDP options for ISE: **device-name** and **platform-type**, as shown in Example 15-9.

https://t.me/learningnets

Example 15-9 XDP Filter List for Device Sensor

Click here to view code image

```
C3750X(config)#device-sensor filter-list cdp list <cdp_list_name>
C3750X(config-sensor-cdplist)#tlv name device-name
C3750X(config-sensor-cdplist)#tlv name platform-type
```

Step 3. Create a list for LLDP:

You will need to configure three LLDP options for ISE: **port-id**, **system-name**, and **system-description**, as shown in Example 15-10.

Example 15-10 LLDP Filter List for Device Sensor

Click here to view code image

```
C3750X(config)#device-sensor filter-list lldp list <lldp_list_name>
C3750X(config-sensor-lldplist)#tlv name port-id
C3750X(config-sensor-lldplist)#tlv name system-name
C3750X(config-sensor-lldplist)#tlv name system-description
```

Step 4. Include the lists created in Steps 1–3 in the device sensor.

In the previous steps you defined the options that device sensor should store. At this point you will configure device sensor to use those lists, as shown in Example 15-11.

Example 15-11 Including the Protocol Lists in the Device Sensor

Click here to view code image

```
C3750X(config)#device-sensor filter-spec dhcp include list <dhcp_list_name>
C3750X(config)#device-sensor filter-spec lldp include list <cdp_list_name>
C3750X(config)#device-sensor filter-spec cdp include list <lldp_list_name>
```

Step 5. Enable the device sensor.

The device sensor is now configured. Next, you will enable the device sensor service to run on the switch and configure when it will send its updates, as shown in Example 15-12.

Example 15-12 Enable the Device Sensor to Notify of Changes

Click here to view code image

```
C3750X(config)#device-sensor accounting
C3750X(config)#device-sensor notify all-changes
```

VMware Configurations to Allow Promiscuous Mode

As shown in Figure 15-18, VMware vSwitches will reject promiscuous mode by default. To use SPAN-type probes with ISE, you must configure the vSwitch to allow promiscuous connections.

https://t.me/learningnets

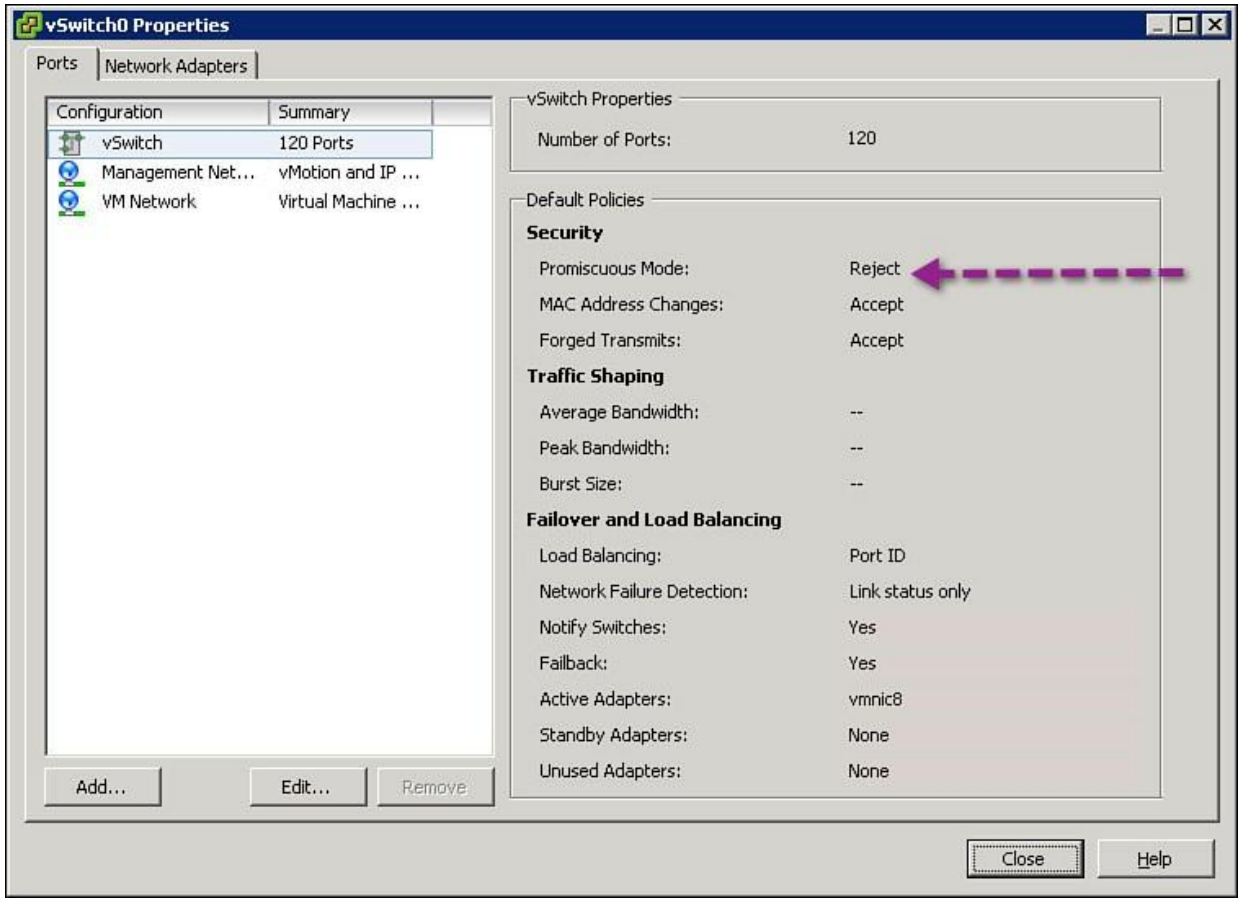


Figure 15-18 Default vSwitch configuration.

To modify the configuration and enable promiscuous traffic, do the following:

Step 1. Highlight the vSwitch and click **Edit**.

Step 2. Click on the **Security** tab.

Step 3. Change the Promiscuous Mode drop-down to **Accept**, as shown in Figure 15-19.

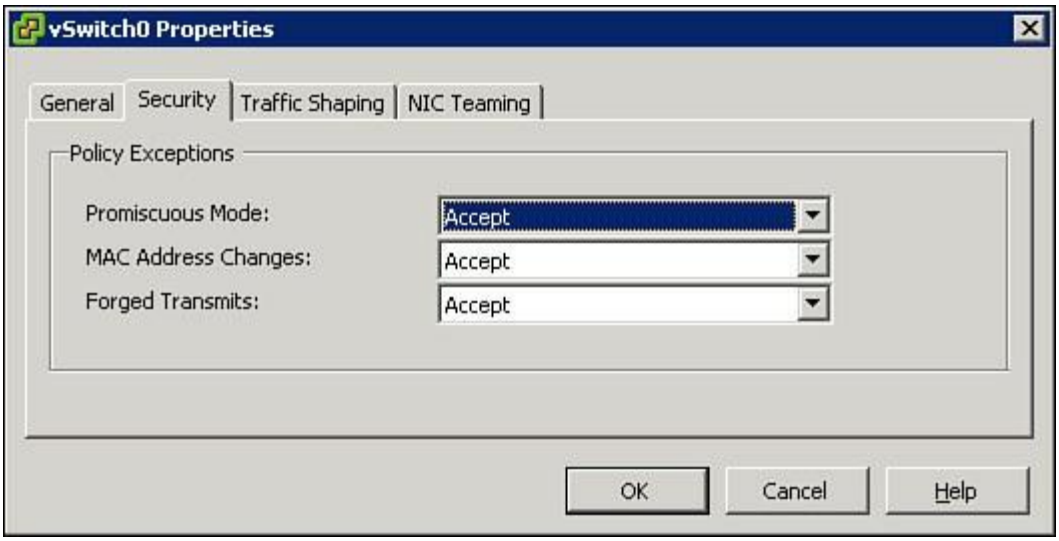


Figure 15-19 Promiscuous vSwitch setting.

https://t.me/learningnets

Profiling Policies

Collecting the data for profiling is only part of the solution. The other aspects are to have endpoint signatures and a policy engine to compare the collected attributes to those signatures; this will lead 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, which are designed to match when certain attributes exist. Additionally, much like an IDS system, an update service enables the engine to download new signatures.

Profiler Feed Service

Although ISE comes with a large and comprehensive list of signatures to classify endpoints (profiles), many devices are produced almost daily (think of the next smart phone or version of the phone's OS). In addition, a constant stream of new profiles is created by Cisco that should be shared to the ISE deployments of the world.

That's why Cisco created a profiler feed service. New devices are released to market, and Cisco creates profiles for them. New profiles are created by Cisco partners (manufacturers). Cisco also has a team that focuses on profile creation. The ISE profiler feed service is used to distribute these new profiles after the Quality Assurance team has approved them.

Configuring the Profiler Feed Service

Configuring the feed service is straightforward. Once enabled, it will reach out to cisco.com at the set time interval and download any published profiles. It 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 lastly a section to send your information to Cisco to help with understanding how many customers are using the feed service.

Figure 15-20 shows an enabled and configured Profiler Feed Service screen, which is located under **Administration > Feed Service > Profiler**.

https://t.me/learningnets

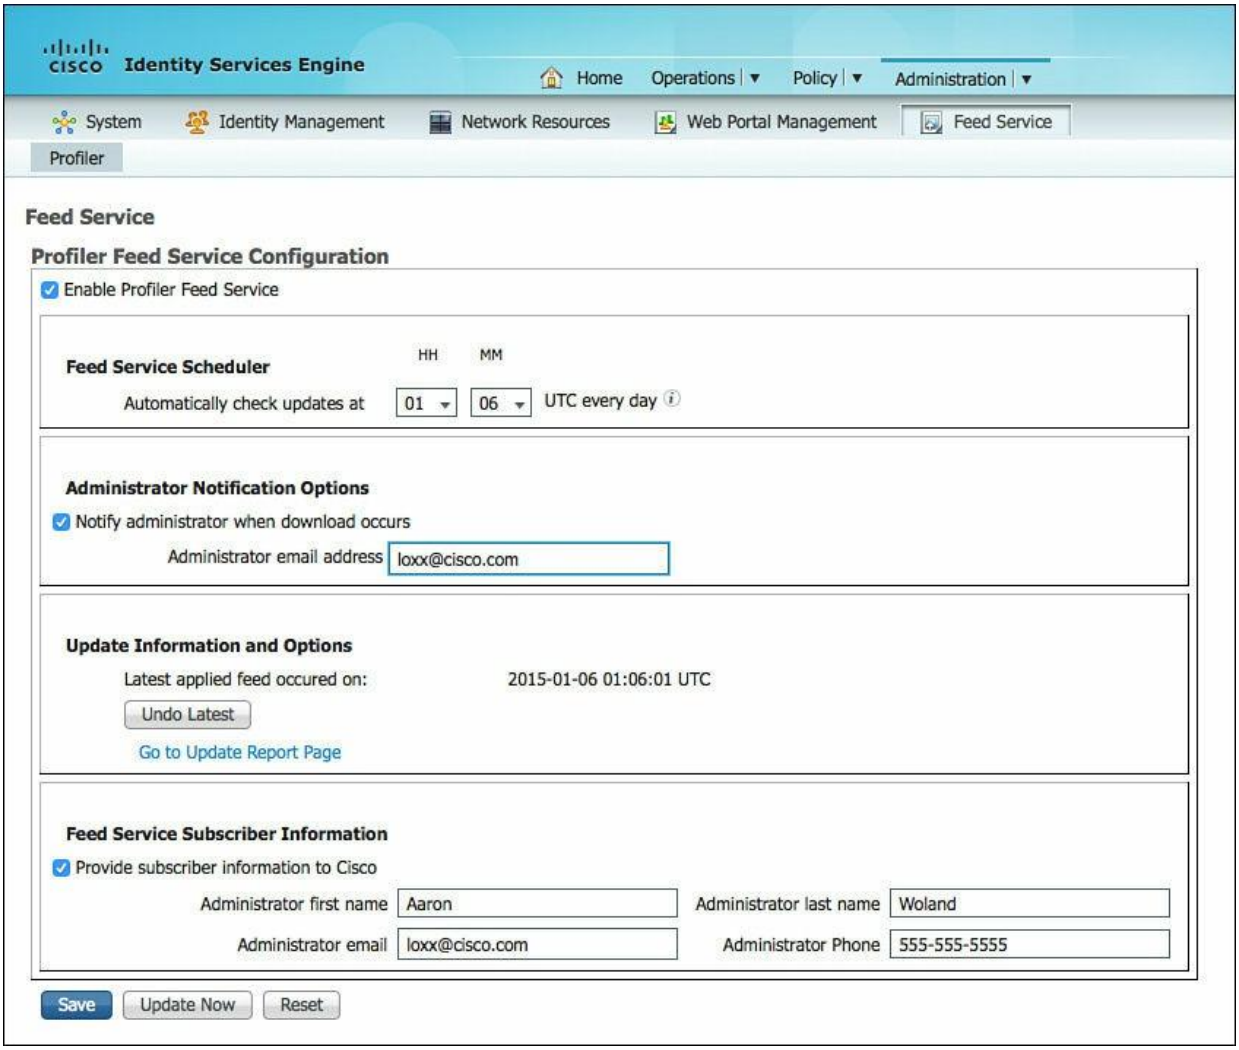


Figure 15-20 Configured profiling feed service.

When you don't want to wait for a configured interval for the feed service to run, you can click an Update Now button. Be cautious with manually updating the profiles during a production workday. When the profiles are updated, it will cause all the endpoints in the endpoint database to be compared against the new list of profiles. In other words, a complete re-profiling of endpoints will occur, which can be very processor intensive.

Verifying the Profiler Feed Service

The easiest way to verify whether the feed service is operational is to click the **Go to Update Report Page** link on the configuration screen, located at **Administration > Feed Service > Profiler**, as shown in Figure 15-21.

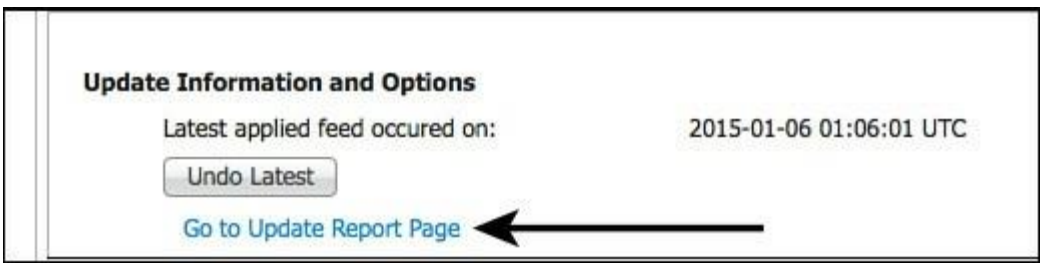
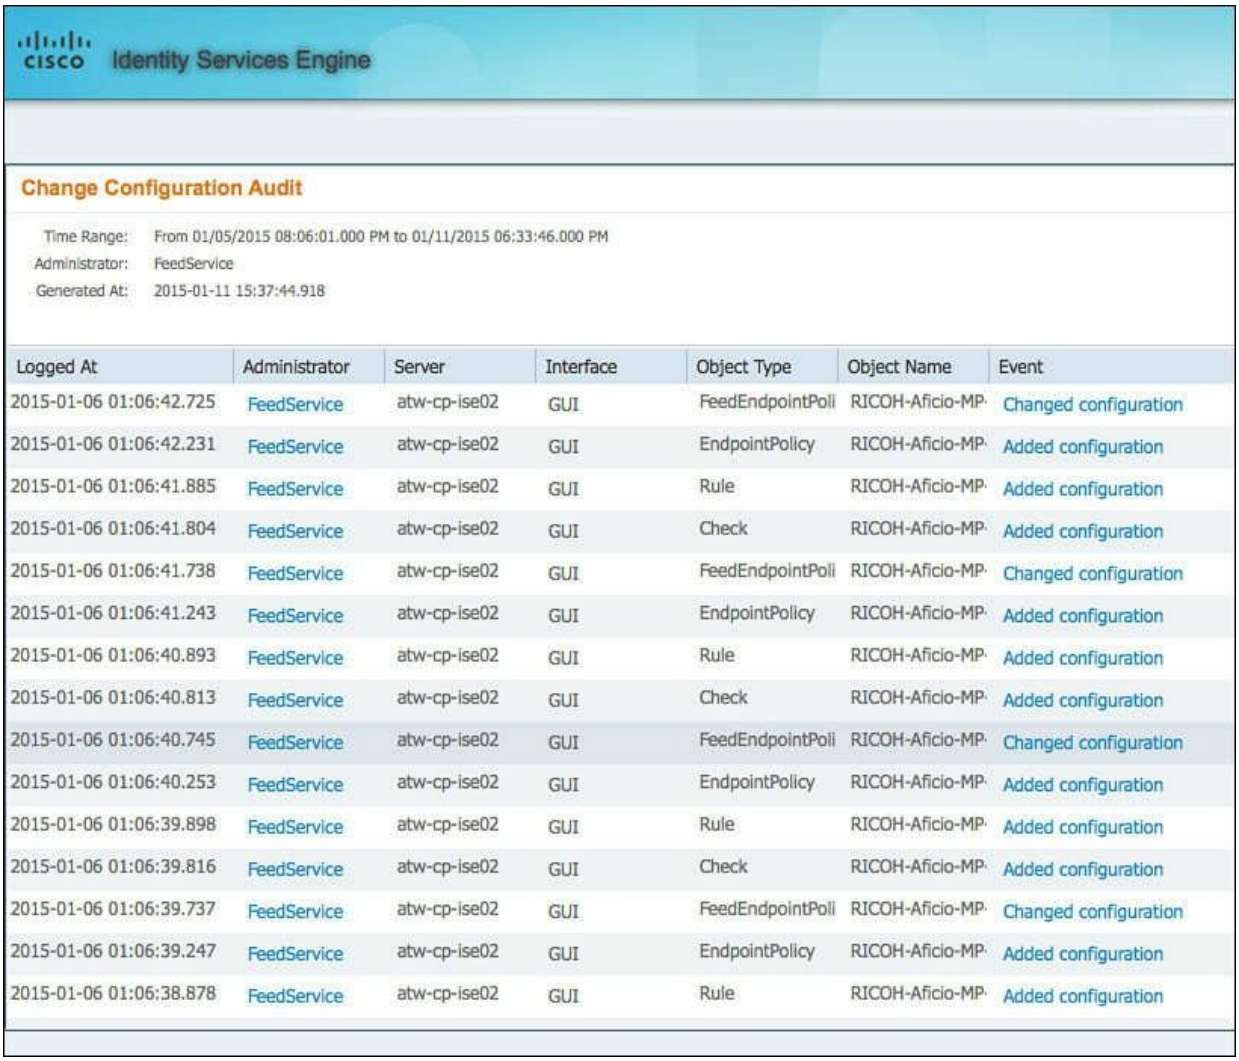


Figure 15-21 Go to update report page link.

Clicking that link opens another window with the Change Configuration Audit report prefiltered for the feed

https://t.me/learningnets

service related entries, as shown in Figure 15-22.



Logged At	Administrator	Server	Interface	Object Type	Object Name	Event
2015-01-06 01:06:42.725	FeedService	atw-cp-ise02	GUI	FeedEndpointPoli	RICOH-Aficio-MP	Changed configuration
2015-01-06 01:06:42.231	FeedService	atw-cp-ise02	GUI	EndpointPolicy	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:41.885	FeedService	atw-cp-ise02	GUI	Rule	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:41.804	FeedService	atw-cp-ise02	GUI	Check	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:41.738	FeedService	atw-cp-ise02	GUI	FeedEndpointPoli	RICOH-Aficio-MP	Changed configuration
2015-01-06 01:06:41.243	FeedService	atw-cp-ise02	GUI	EndpointPolicy	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:40.893	FeedService	atw-cp-ise02	GUI	Rule	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:40.813	FeedService	atw-cp-ise02	GUI	Check	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:40.745	FeedService	atw-cp-ise02	GUI	FeedEndpointPoli	RICOH-Aficio-MP	Changed configuration
2015-01-06 01:06:40.253	FeedService	atw-cp-ise02	GUI	EndpointPolicy	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:39.898	FeedService	atw-cp-ise02	GUI	Rule	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:39.816	FeedService	atw-cp-ise02	GUI	Check	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:39.737	FeedService	atw-cp-ise02	GUI	FeedEndpointPoli	RICOH-Aficio-MP	Changed configuration
2015-01-06 01:06:39.247	FeedService	atw-cp-ise02	GUI	EndpointPolicy	RICOH-Aficio-MP	Added configuration
2015-01-06 01:06:38.878	FeedService	atw-cp-ise02	GUI	Rule	RICOH-Aficio-MP	Added configuration

Figure 15-22 Change configuration audit for feed service.

Because ISE must be able to reach cisco.com for the profiling feed service to function, you might need to configure ISE to use a proxy server to reach the Internet. This optional configuration is located at **Administration > System > Settings > Proxy**, as shown in Figure 15-23.

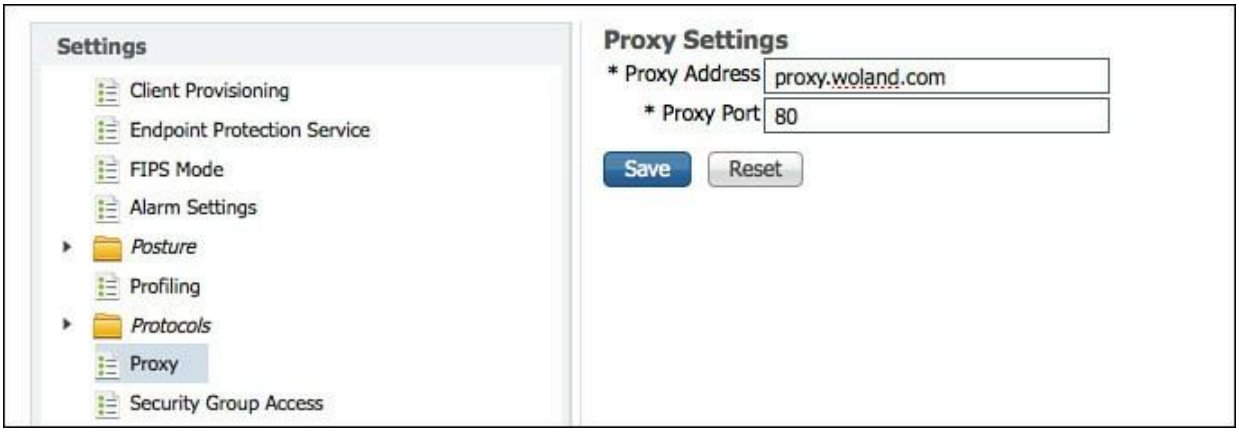


Figure 15-23 Proxy setting.

https://t.me/learningnets

Endpoint Profile Policies

The profiling service probes collect attributes of endpoints, while the profile policies are similar to signatures, which define the endpoint profiles themselves. For example, to match an Apple-Device profile, the endpoint must have a MAC address beginning 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 large number of predefined profile policies, and you have just read about the feed service that's used to update those policies and provide new ones.

The endpoint profile policies can be viewed by navigating to **Policy > Profiling**, as shown in Figure 15-24.

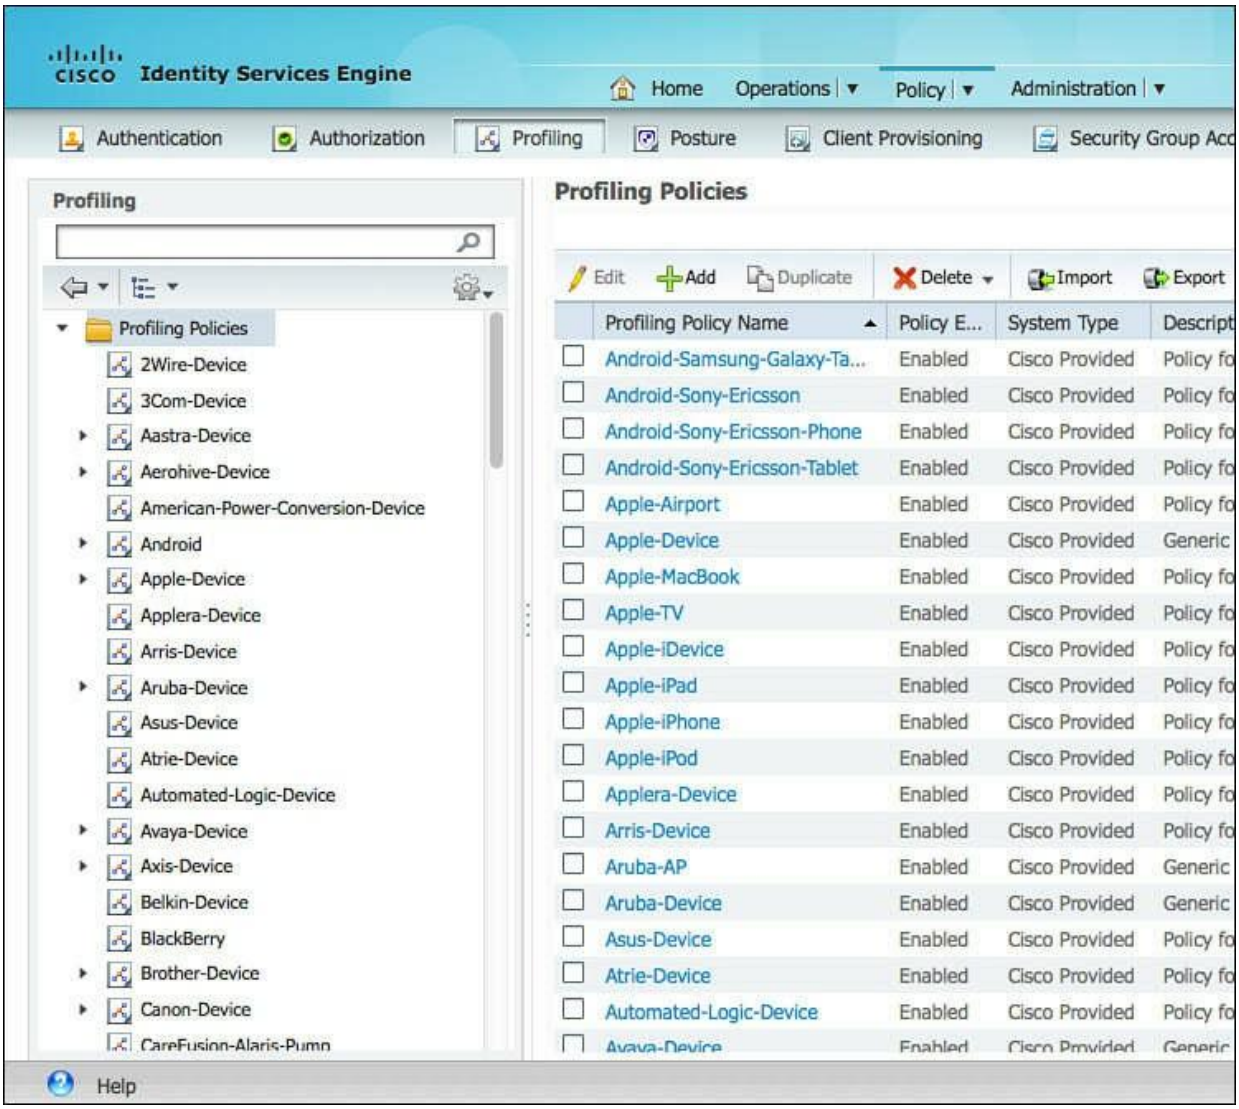


Figure 15-24 Policy > Profiling.

Each profile is listed as Cisco provided or Administrator Modified. This classification ensures that the feed service will not override one that has been changed by the administrator.

Profiles are hierarchical and inclusive in nature, and you can pick any level to use within your authorization policies, enabling you to be very specific or broad in your rules. Examining Figure 15-25, you see the existence of a parent policy named Cisco-Device with a child policy named Cisco-IP-Phone, and that has another child policy named Cisco-IP-Phone-7940 (Cisco-Device > Cisco-IP-Phone > Cisco-IP-Phone-7940).

https://t.me/learningnets

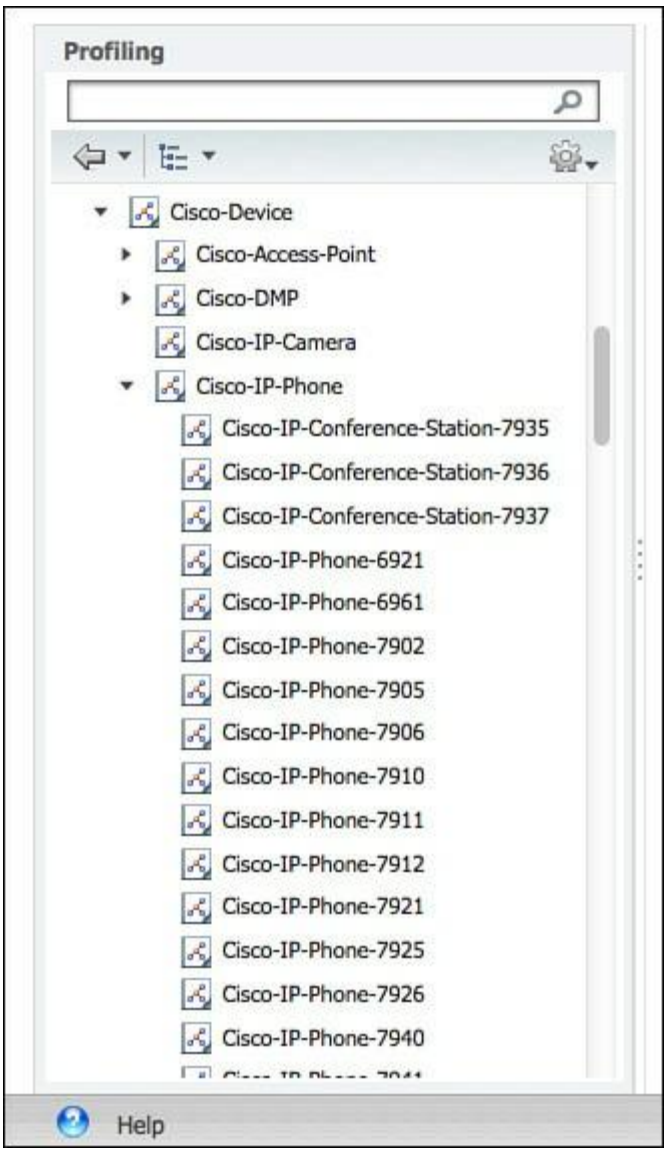


Figure 15-25 Cisco-Device parent profile.

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 15-26 graphically displays the way a profile hierarchy is built within ISE.

https://t.me/learningnets

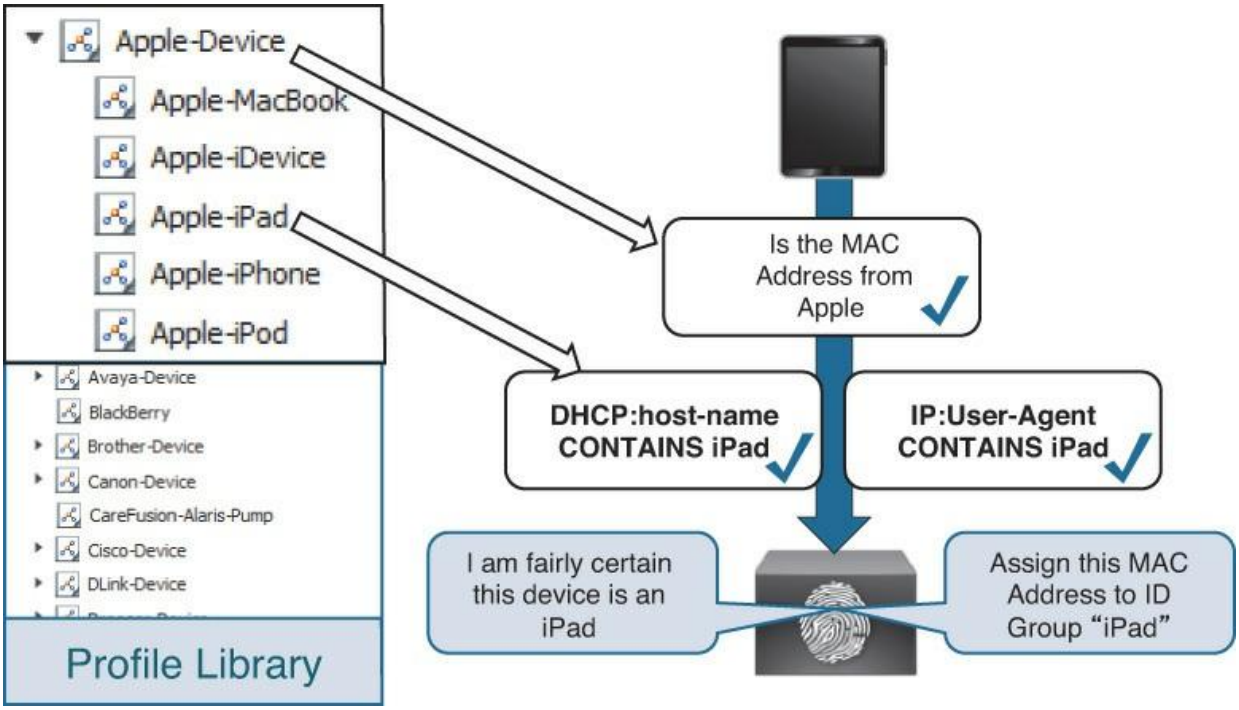


Figure 15-26 Profiling hierarchy illustrated.

To serve as further examples, there is a predefined authorization rule named Profiled Cisco IP Phones. This rule permits full access to the network and assigns permission to join the Voice VLAN for all devices that that are profiled as a Cisco-IP-Phone parent profile and any of the child profiles. Figure 15-27 shows this rule.

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
✓	Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
✓	Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones

Figure 15-27 Default Profiled Cisco IP Phones rule.

Continuing to examine policy hierarchy and how it works, let's dig into a specific example. In this sample network, there is a Cisco 7970 IP Phone, and that phone has been granted access from the Profiled Cisco IP Phones default authorization rule. This permits all devices matching a Cisco-IP-Phone profile Identity Group.

Start by examining the endpoint attributes and comparing them to the profiling policies.

Step 1. Navigate to **Administration > Identity Management > Identities > Endpoints**. Locate the 7970 phone, as shown in Figure 15-28.

https://t.me/learningnets

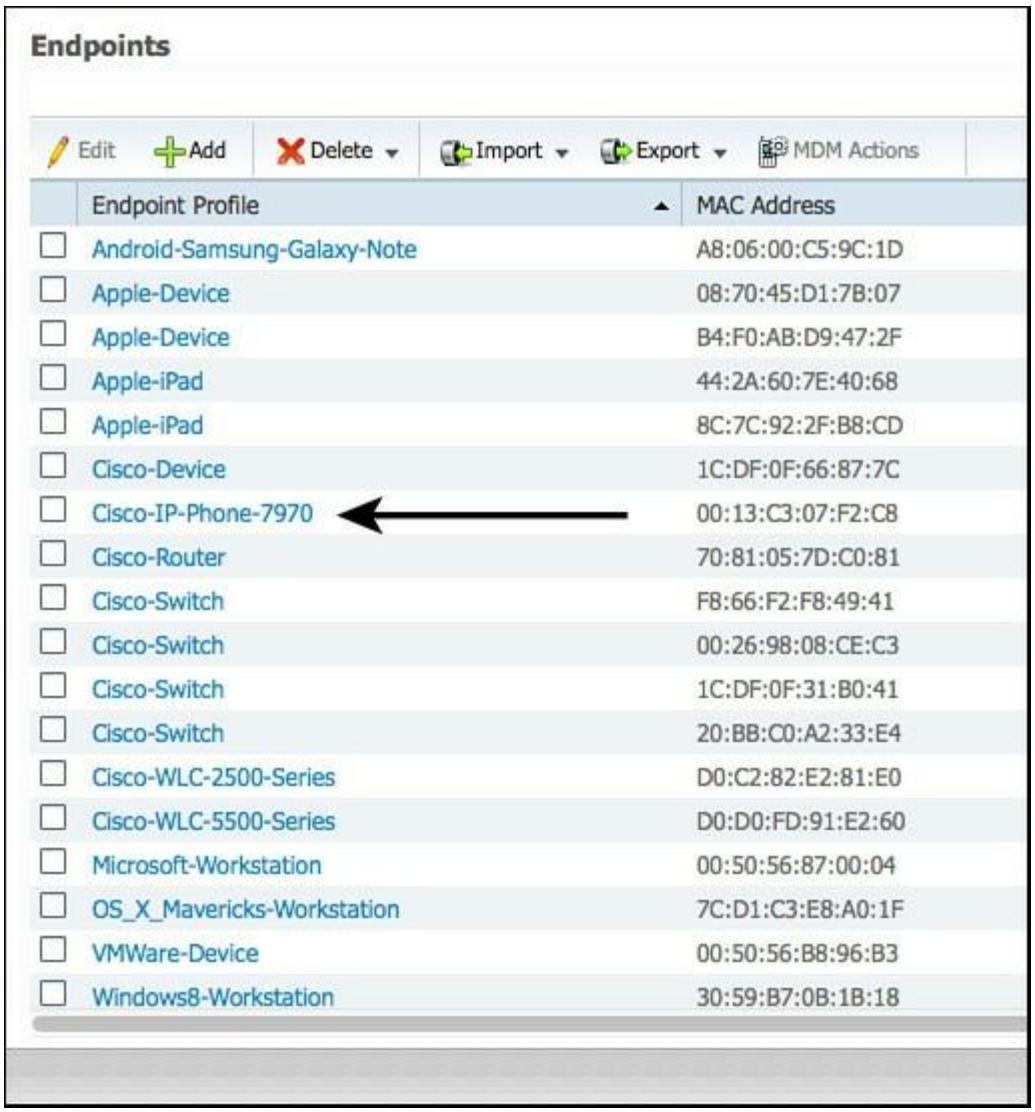


Figure 15-28 Filtered endpoint identities.

Step 2. Click that endpoint to enter into the endpoint details.

Step 3. You will immediately see that the Policy Assignment (profile of the device) is Cisco-IP-Phone-7970. However, the Identity Group Assignment is Cisco-IP-Phone. Additionally, the AuthorizationPolicyMatchedRule is Profiled Cisco IP Phones, which means the correct authorization rule is being matched, as shown in Figure 15-29.

https://t.me/learningnets



Figure 15-29 Endpoint details.

Step 4. Scroll downward in the endpoint details and see that the EndPointSource is the RADIUS probe, as shown in Figure 15-30. That means the RADIUS probe provided us with the most information needed to classify this device (which must have arrived from the device sensor on the NAD).

https://t.me/learningnets

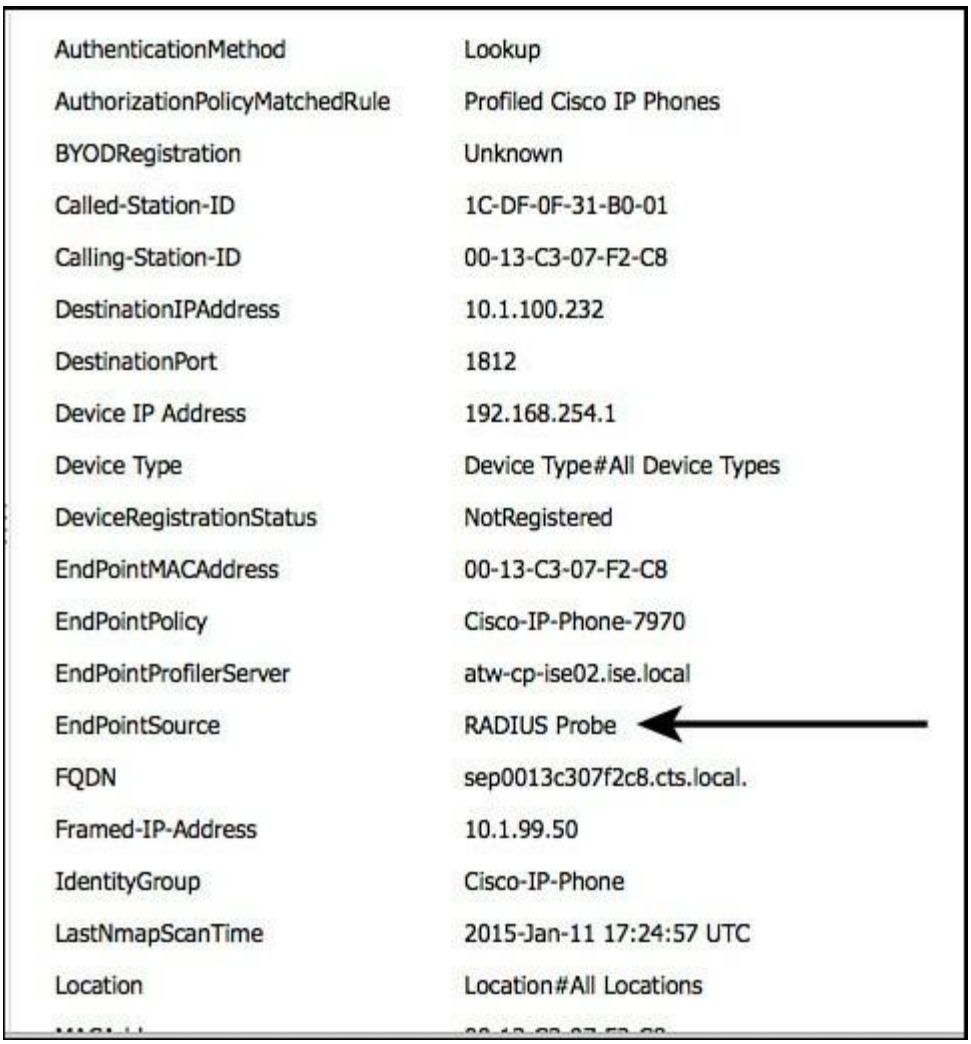


Figure 15-30 Endpoint details: EndPointSource.

Step 5. Continuing down the list of attributes, notice the CDP cached data from the switch that was sent to ISE via the RADIUS probe (see Figure 15-31).

https://t.me/learningnets

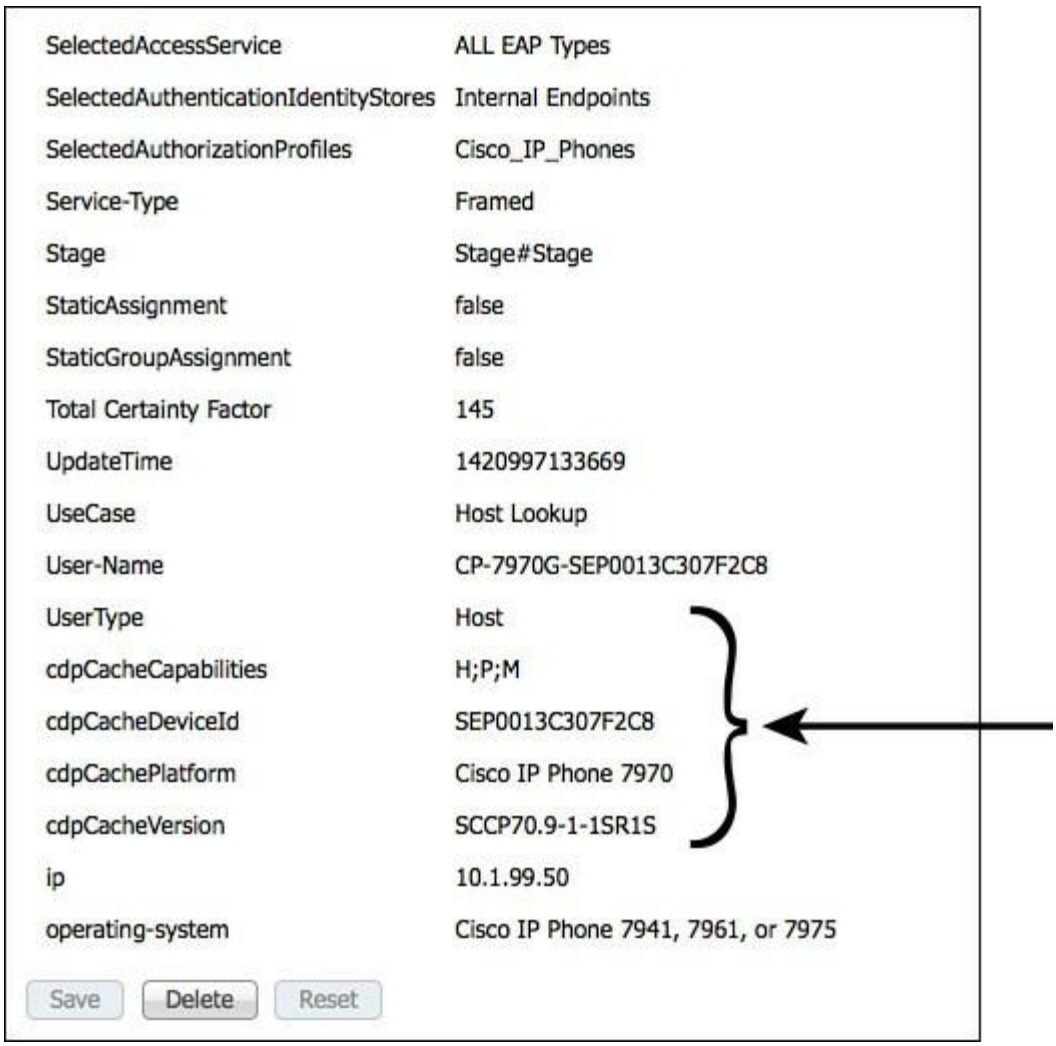


Figure 15-31 Endpoint details: CDP data.

You can see these attributes of the endpoint, but how is that used within the profiling policy? To answer that question, 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, as shown in Figure 15-32.

https://t.me/learningnets

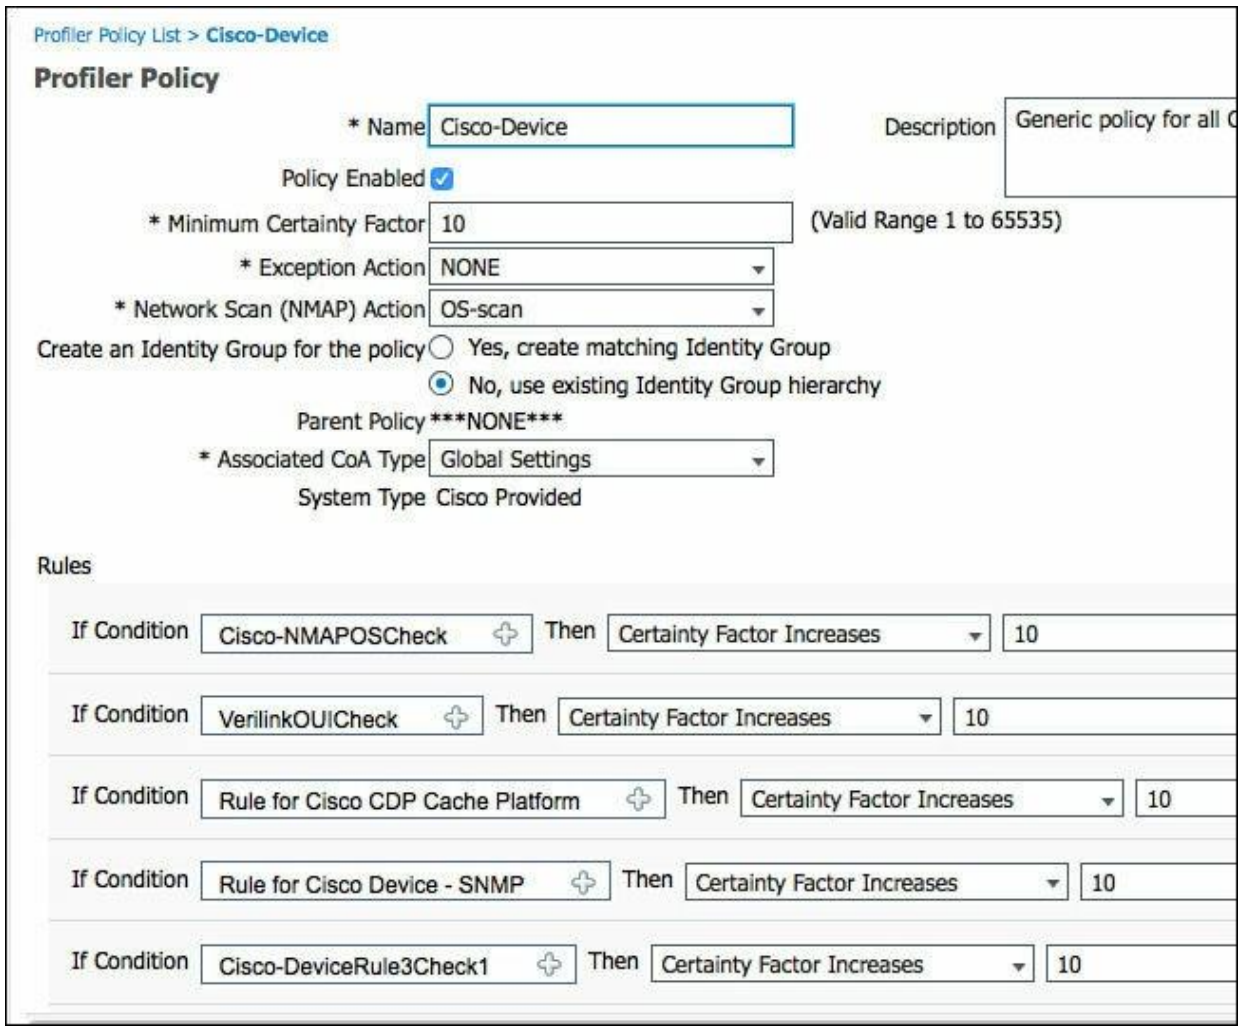
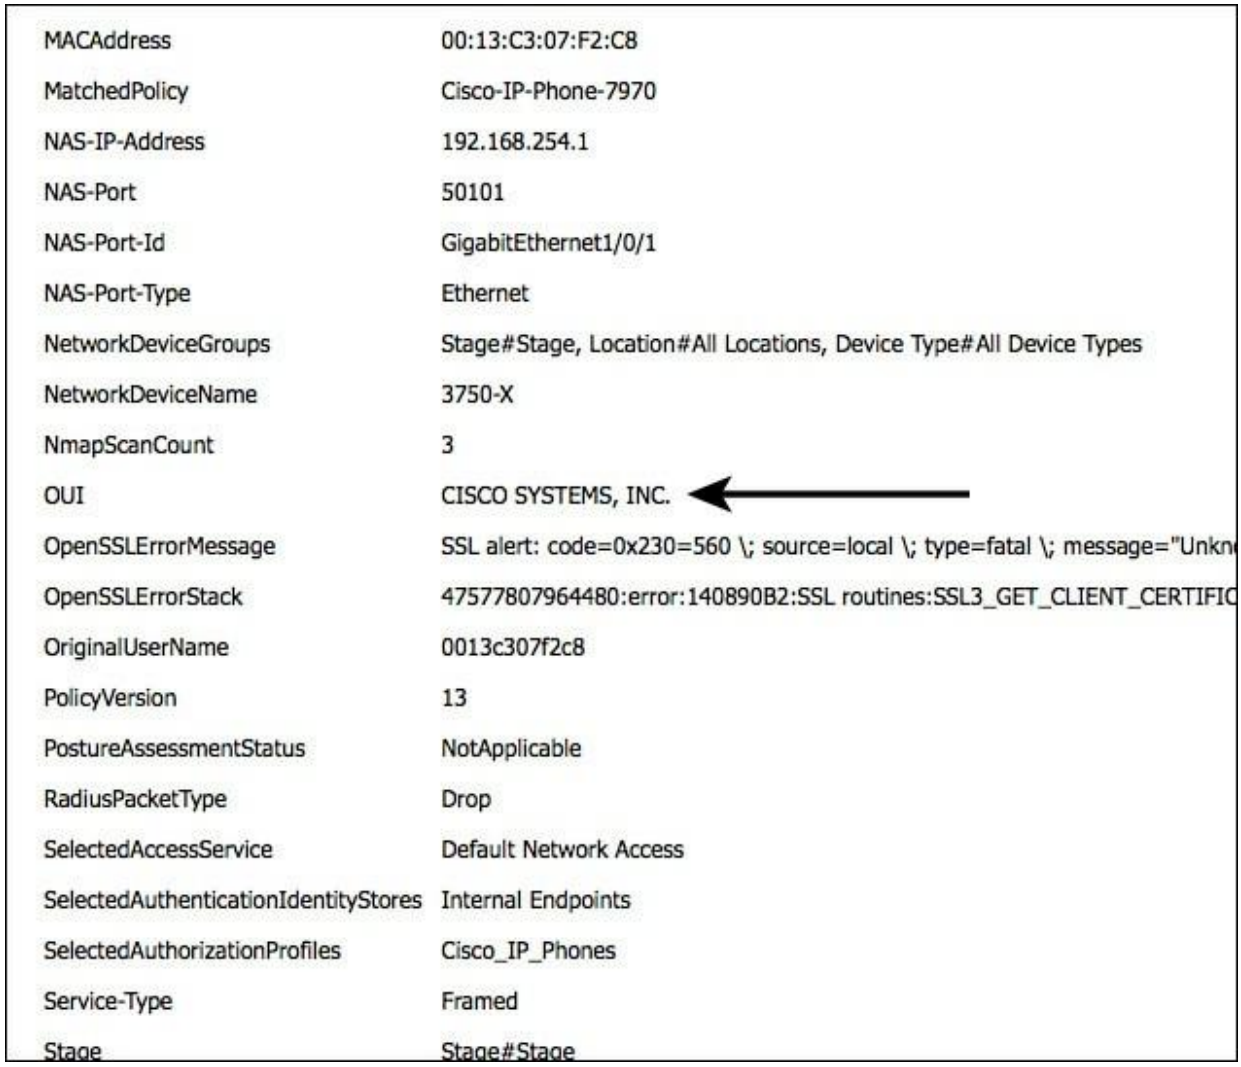


Figure 15-32 Cisco-Device profiling policy.

Using Figure 15-32 as a reference point, please 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. Each of the conditions at the bottom of the policy that are matched will add up to equal the endpoints certainty value. The higher the value, the more certain ISE profiler is that an endpoint matches the specific profile.
- Note that the OUI is CISCO SYSTEMS, INC., shown in Figure 15-33, and that will match the condition for Cisco-Device shown in Figure 15-34. This is one possible mapping that meets the minimum certainty value and should match the endpoint to this parent policy.

https://t.me/learningnets



MACAddress	00:13:C3:07:F2:C8
MatchedPolicy	Cisco-IP-Phone-7970
NAS-IP-Address	192.168.254.1
NAS-Port	50101
NAS-Port-Id	GigabitEthernet1/0/1
NAS-Port-Type	Ethernet
NetworkDeviceGroups	Stage#Stage, Location#All Locations, Device Type#All Device Types
NetworkDeviceName	3750-X
NmapScanCount	3
OUI	CISCO SYSTEMS, INC.
OpenSSLErrorMessage	SSL alert: code=0x230=560 \; source=local \; type=fatal \; message="Unkn
OpenSSLErrorStack	47577807964480:error:140890B2:SSL routines:SSL3_GET_CLIENT_CERTIFIC
OriginalUserName	0013c307f2c8
PolicyVersion	13
PostureAssessmentStatus	NotApplicable
RadiusPacketType	Drop
SelectedAccessService	Default Network Access
SelectedAuthenticationIdentityStores	Internal Endpoints
SelectedAuthorizationProfiles	Cisco_IP_Phones
Service-Type	Framed
Stage	Stage#Stage

Figure 15-33 OUI of the endpoint.

https://t.me/learningnets

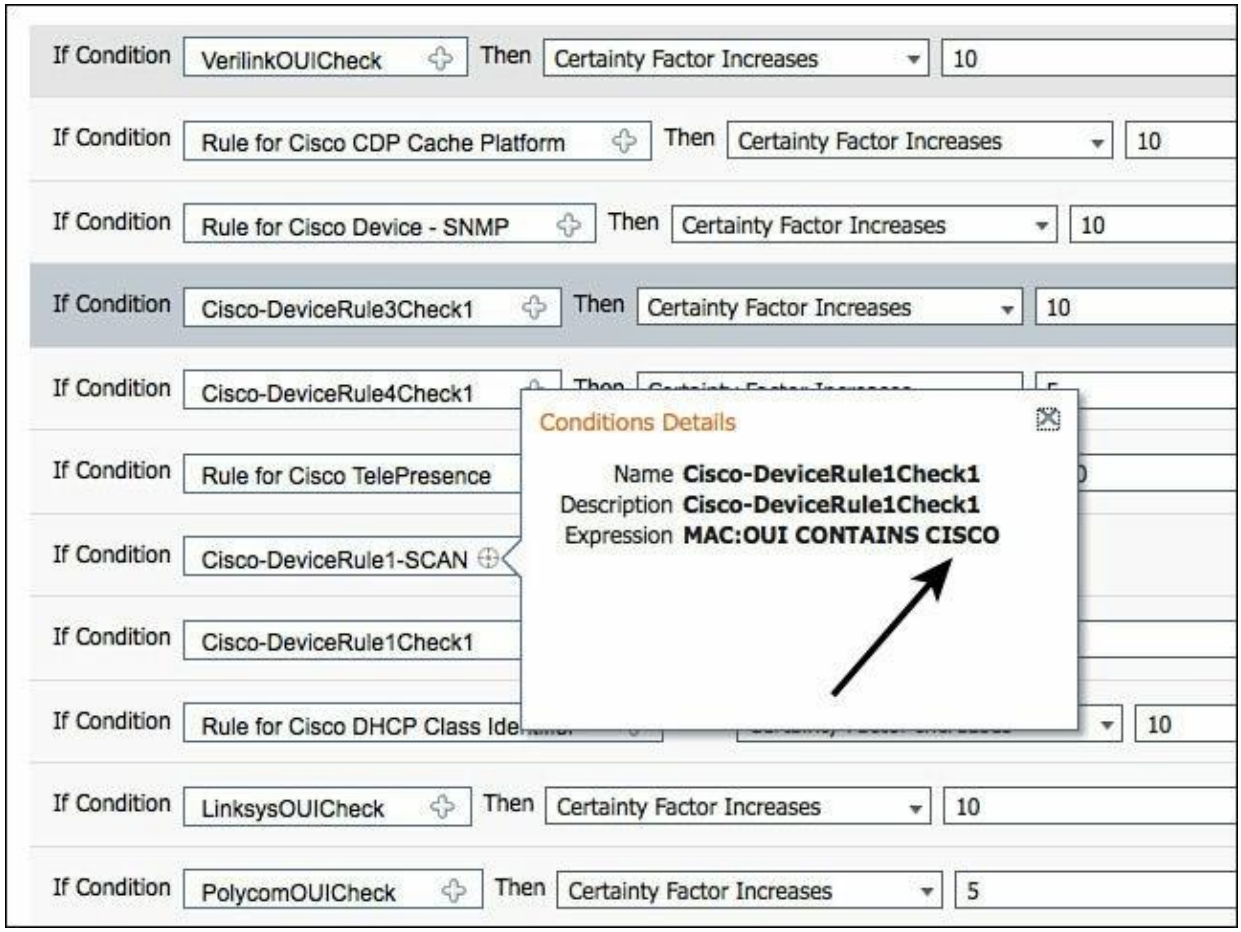


Figure 15-34 OUI condition in Cisco-Device policy.

- A Network Scan (NMAP) Action is set to OS-Scan. For this action to occur, there must be a condition in the profile that has a result to trigger the network scan. Figure 15-35 displays this mapping of the condition to the action. Figure 15-34 displays the contents of the condition, meaning that if the MAC OUI contains CISCO, then run the NMAP Scan type defined.

https://t.me/learningnets

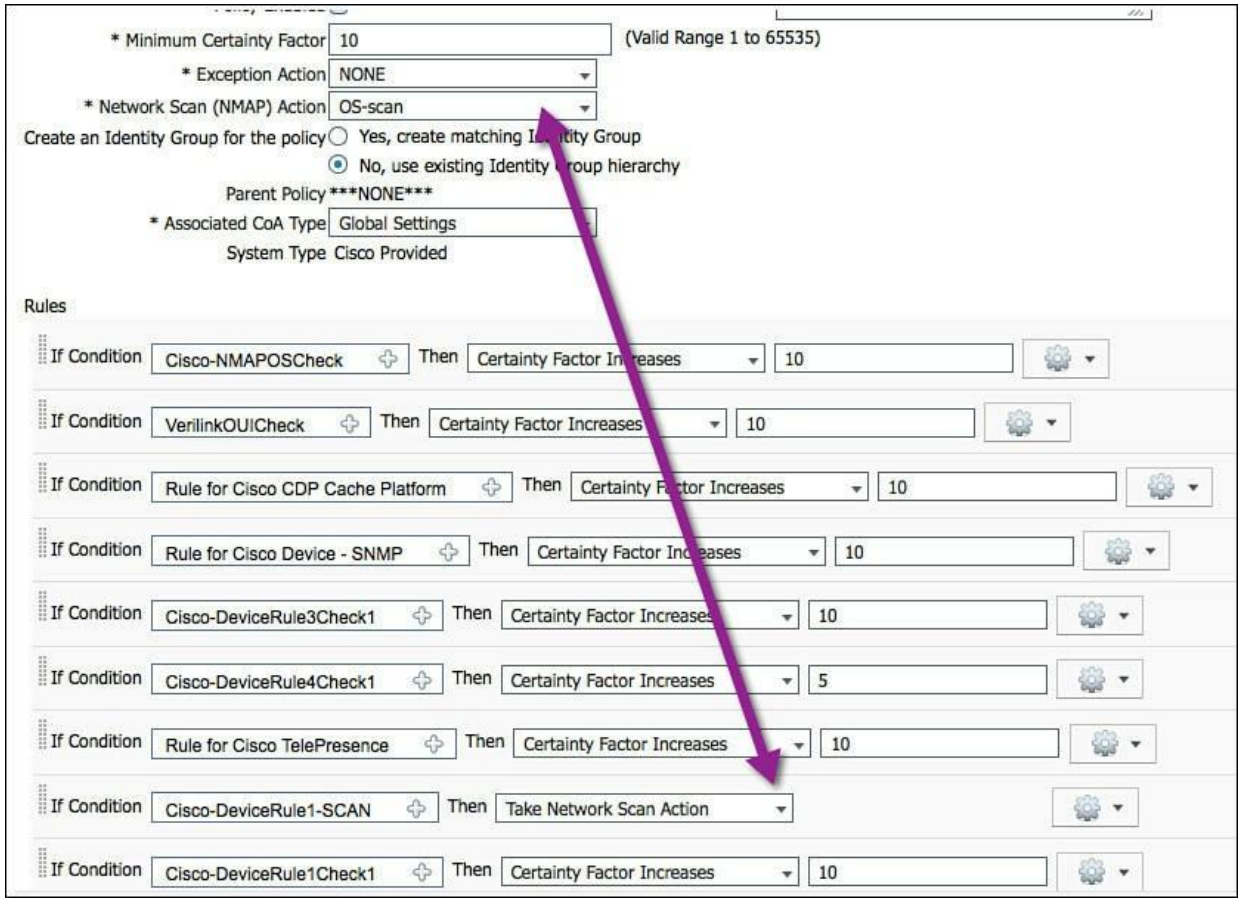


Figure 15-35 Network scan action and condition with scan result.

- This profiling policy has a tremendous number of conditions. Most of them add a certainty value of 5 or 10. The Certainty Value needs to be only a minimum of 10 to match the profile, so matching any one of these conditions will most likely equal a match.

Step 7. Navigate to **Policy > Profiling > Cisco-Device > Cisco-IP-Phone**, as shown in Figure 15-36, to examine the conditions used for this policy.

https://t.me/learningnets

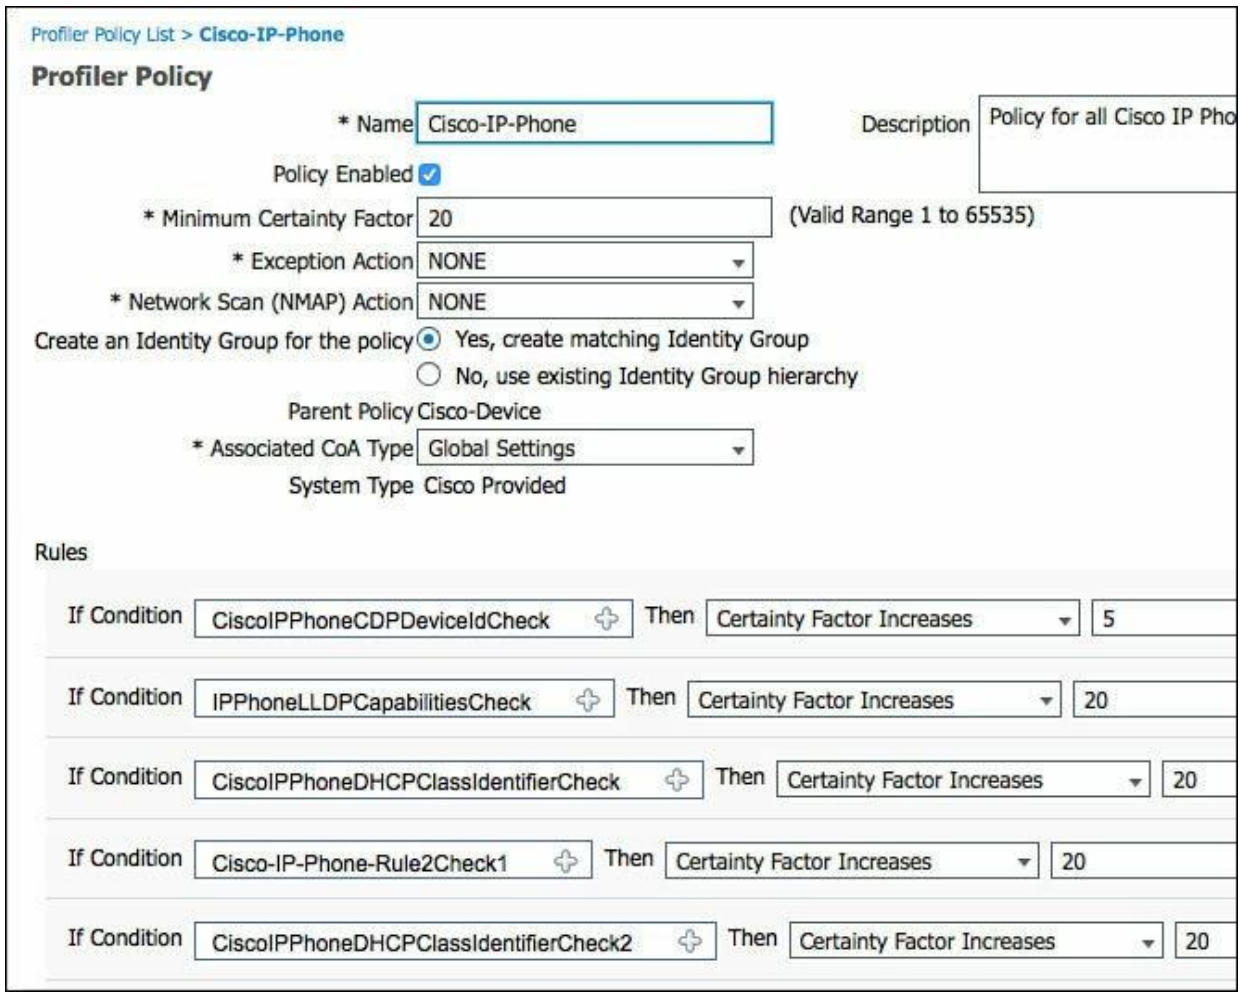


Figure 15-36 Cisco-IP-Phone profiling policy.

Using this profile as a reference point, please take note of the following details:

- 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 policy before these conditions would ever be examined.
- The minimum certainty factor of this profile is 20, as shown in Figure 15-36. Certainty factor is an aggregate value between 1 and 65535. Each of the conditions at the bottom of the policy can add more certainty to this profile, if they are matched.
- The CDP Cache shown in Figure 15-31 shows that the cdpCachePlatform attribute was sent as Cisco IP Phone 7970.
- Figure 15-37 shows that the Cisco IP Phone profile policy uses a condition looking for the cdpCachePlatform value to contain Cisco IP Phone and increase the certainty by 20.

Step 7 examined Cisco-IP-Phone, but what does it take to get one step further, to reach the final Cisco-IP-Phone-7970 profile?

https://t.me/learningnets

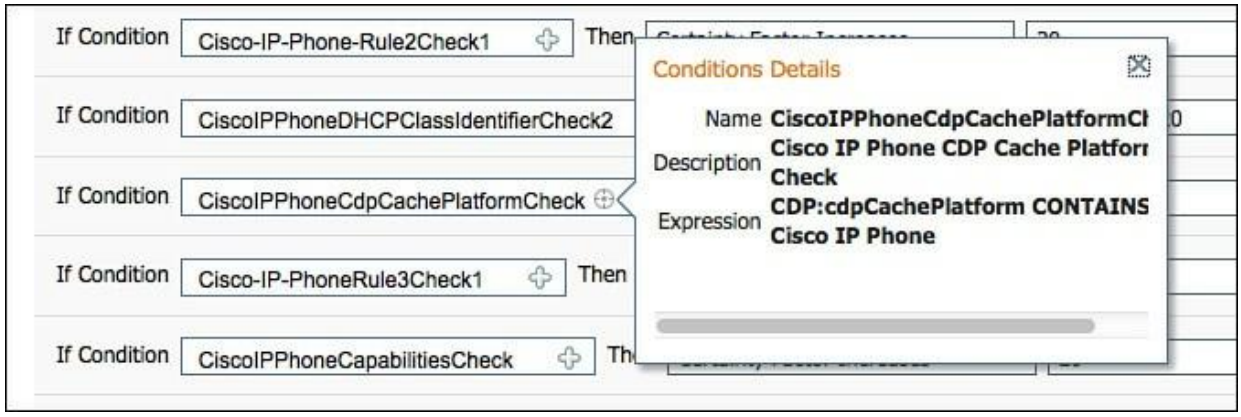


Figure 15-37 cdpCachePlatform condition.

Step 8. Navigate to **Policy > Profiling > Cisco-Device > Cisco-IP-Phone > Cisco-IP-Phone-7970**, and we will examine the conditions used for this policy (see Figure 15-38).

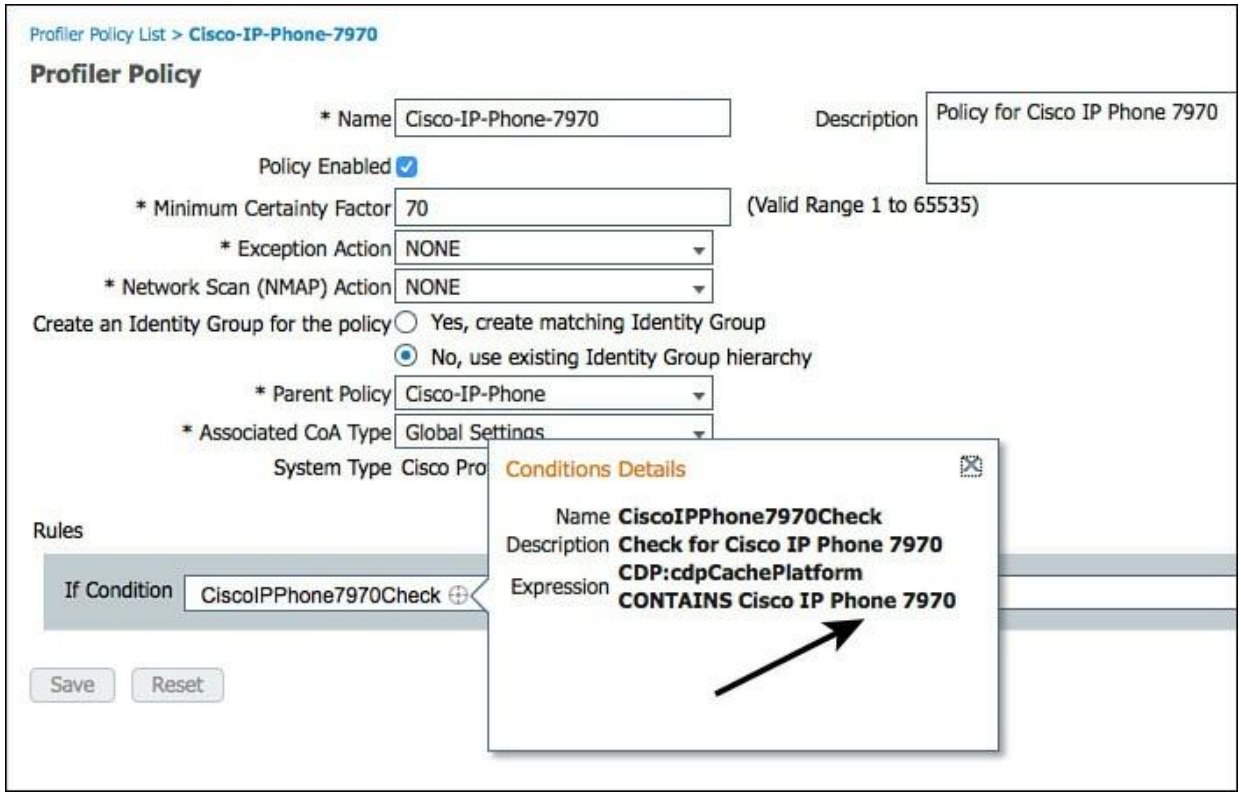


Figure 15-38 Cisco-IP-Phone-7970 profile.

Using Figure 15-38 as a reference point, please take note of the following details:

- 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 policy before these conditions would ever be looked at.
- The profile itself has only one condition, and that condition is that the cdpCachePlatform attribute is Cisco IP Phone 7970. Figure 15-31 has shown that it is.

Rarely would you build an authorization policy that is specific to the point of the model number of the Cisco Phone; instead you would just use the Cisco-IP-Phone parent policy in your authorization policies.

https://t.me/learningnets

Logical Profiles

When ISE 1.0 was first released, it was quickly requested by many customers to have a grouping of profiles that are not hierarchical. For example, many requested to create a profile group named IP-Phones that contains all the individual profiles of IP phones, Cisco and non-Cisco alike.



ISE 1.2 answered that request. With that release, Cisco introduced the new concept of logical profiles. These logical profiles are exactly what customers requested: a grouping of profiles.

To examine the only prebuilt logical profile in ISE, navigate to **Policy > Profiling > Logical Profiles**. Select **IP-Phones**, as shown in Figure 15-39.

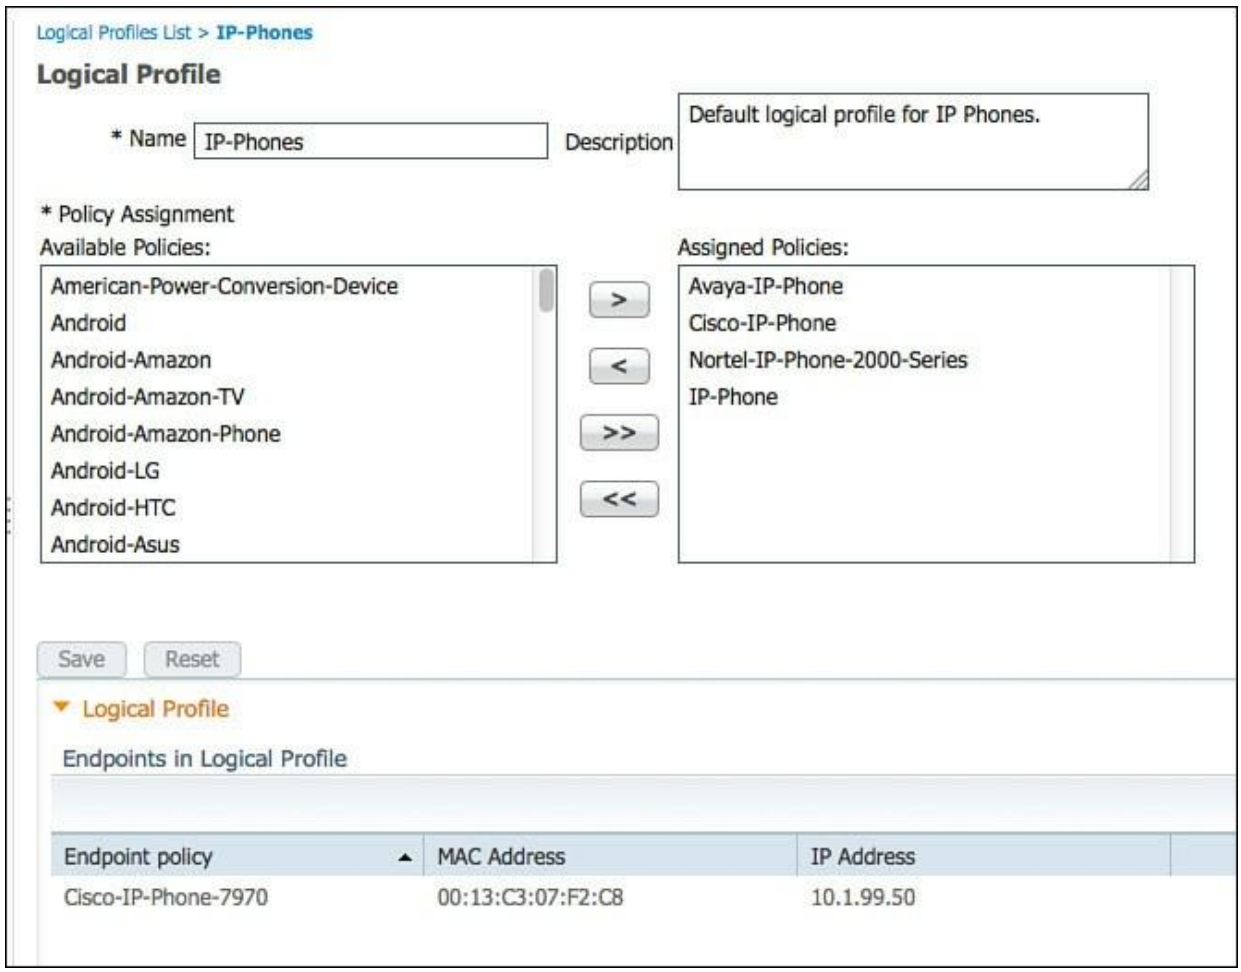


Figure 15-39 IP-Phones logical profile.

Notice in Figure 15-39 that the logical profile contains Avaya, Cisco, Nortel, and generic IP phone profiles.

ISE Profiler and CoA

When using an endpoint profile as an attribute within your authorization policy, you will be providing 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 devices is, yet you cannot find out what the device is without providing some level of access. The endpoint must be active on the network for ISE to identify the endpoint profile.

Enter the concept of change of authorization (CoA), which was introduced to you in Chapter 6. Without

https://t.me/learningnets

CoA, the only time a policy server such as ISE is permitted to send a command to the NAD is during a response to an authentication request. This created numerous issues because there would not be a way to disconnect a bad actor from the network or change the level of access an endpoint is permitted to have based on a newer data that has been learned at the policy engine. The current authorization to the network would have to sustain until the next time the endpoint has to authenticate.

Because the 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 nailed down further and the final result after the endpoint profile is definitely known, you cannot wait for the next authentication request each time. Instead, the profiling engine can use CoA to change the level for each state the endpoint goes through.

There are two main areas for configuring CoA with profiling. A global setting enables CoA for profiling in the ISE deployment, and a CoA can be configured 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 15-40.

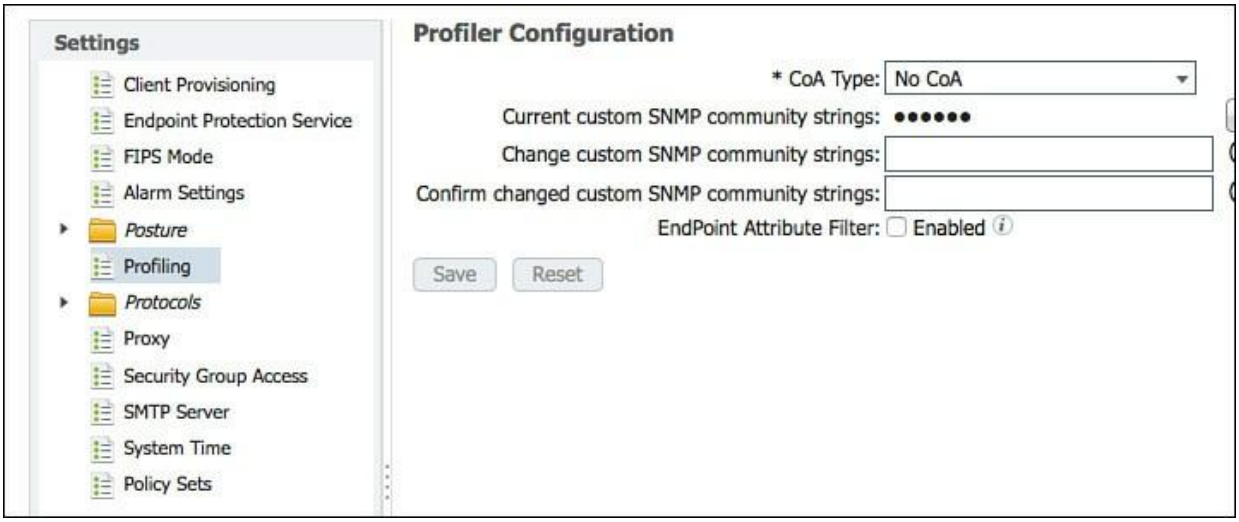


Figure 15-40 Profiler global settings.

As shown in Figure 15-40, the default setting is No CoA. Click the drop-down list, as shown in Figure 15-41, to see the choices: Port Bounce and Reauth.

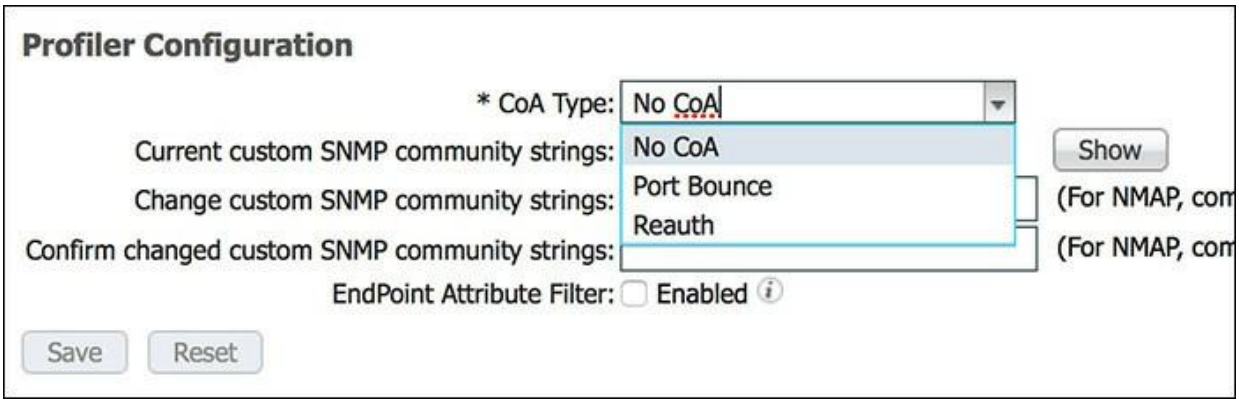


Figure 15-41 Profiler CoA types.

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, simulating the unplugging and plugging in of network cable. The benefit to this type of CoA is that many devices will try to renew their DHCP assigned IP addresses when the link state

https://t.me/learningnets

changes. Additionally, there is a built-in failsafe to never send a Port Bounce when more than one MAC address is seen on the switch port. That failsafe is in place to ensure there is no negative impact on IP telephony. When more than one MAC address exists on the switch port, a Reauth is sent instead.

The Reauth CoA instructs the NAD to initiate a new authentication to the endpoint, sending another EAPoL Start message to trigger the supplicant to send the credentials again. In the case of MAB, however, the NAD resends a RADIUS authentication with the endpoint MAC address as the identity credential. Either way there is 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.



Regardless of the CoA type used, ISE has now forced a new authentication attempt so that a different authorization result can be sent to the NAD, providing 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.

After 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

ISE 1.2 added a new setting to profiles that enable an administrator to control her own destiny with CoA. This came about with the need to send a Port Bounce CoA for certain devices only, while using the global Reauth CoA for the remaining endpoints.

This is known as the per-profile CoA. When a CoA type is configured for a profile, it is used when an endpoint is classified as that profile type. The setting is shown in Figure 15-42.

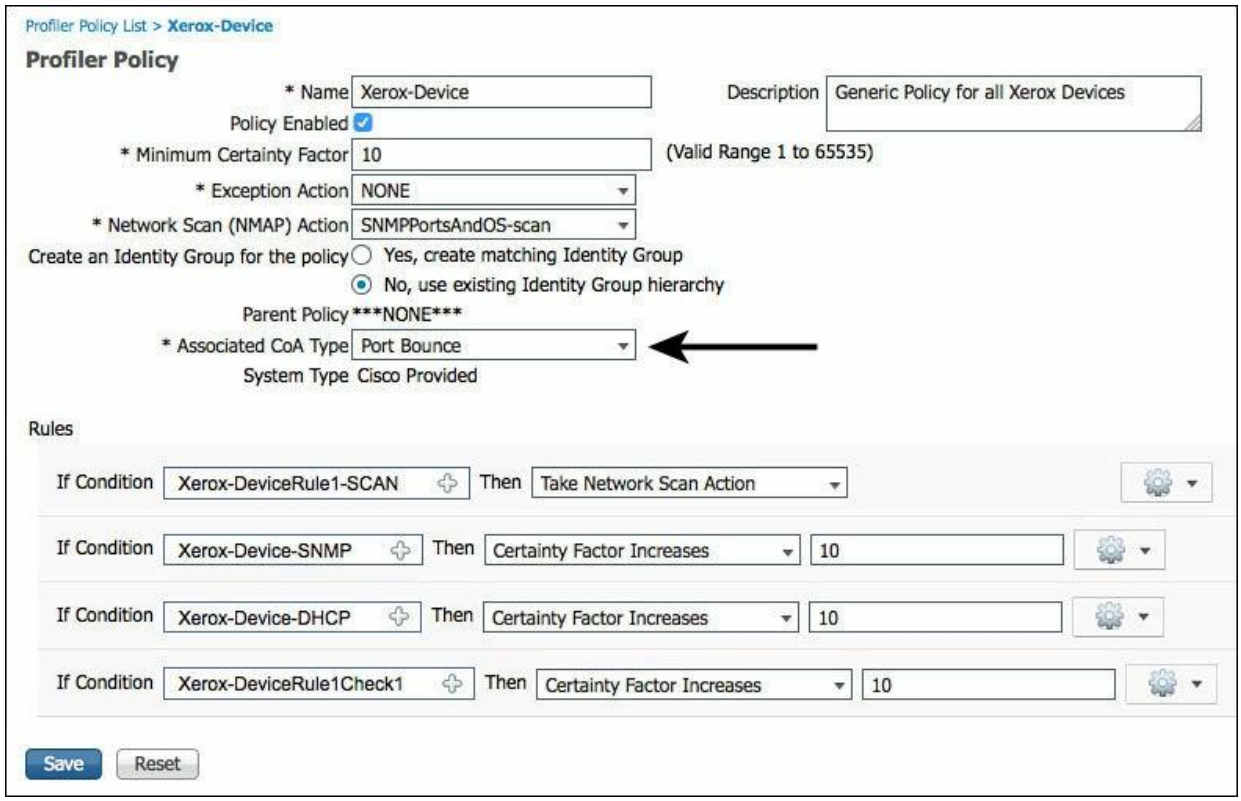


Figure 15-42 Per-profile CoA.

As shown in Figure 15-42, a device profiled as a Xerox-Device will now trigger a Port Bounce CoA, causing the link to go down and back up again. This in turn triggers the endpoint to request a new IP address from

https://t.me/learningnets

the DHCP server. This is useful when a device is using MAB and needs to be assigned to a different VLAN.

Global Profiler Settings

Additional settings related to profiling are set at the global (system-wide) level, not just the global CoA type. These include the SNMP community strings for NMAP SNMP walks and the enable setting for endpoint attribute filtering.

Configuring 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 in the Network Scan (NMAP) section, NMAP uses SNMP to examine endpoints. For this to function, ISE profiler must know which SNMP community strings to use. The community strings to use are configured within **Administration > System > Settings > Profiler** and are configured by listing each community string one-by-one with a comma separating each value. After they are saved, the two text boxes are erased; you then must click the **Show** button to see the configured strings, as displayed in Figure 15-43.

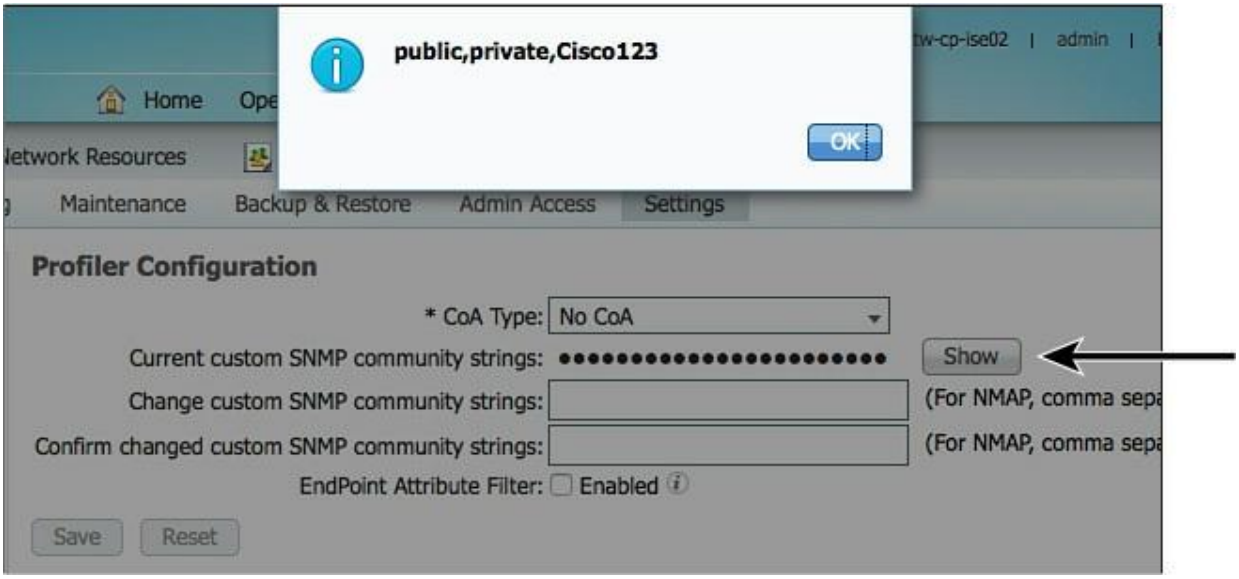


Figure 15-43 Global SNMP settings for NMAP probe.

Endpoint Attribute Filtering

ISE Profiler can and does collect a lot of data about endpoints. It stores all that data and replicates it to the other ISE nodes in the deployment. To help keep the replication traffic down, ISE has endpoint attribute filters, which are enabled at **Administration > System > Settings > Profiling**, as shown in Figure 15-44.

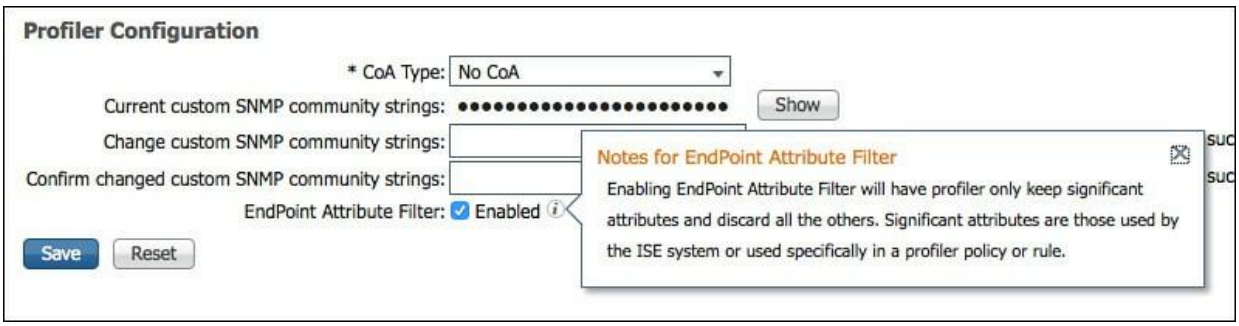


Figure 15-44 Endpoint attribute filter.

When the filtering is enabled, profiler builds a white list of attributes that are used in the existing profiler

https://t.me/learningnets

policies. In other words, the profiler examines every policy that is enabled and creates a list of attributes that are needed for all those policies. Only those attributes will now be collected and stored in the endpoint database.

Use of the endpoint attribute filter is highly recommended but only after a deployment has been up and running properly for an extended period of time. The reason it is not recommended for use immediately with an ISE install is to enable the administrator to have all the necessary attributes to build new profiles. If the list of collected attributes is filtered down, there might not be enough data available to create the necessary profiles.

Profiles in Authorization Policies

As you saw with the profiled Cisco IP phone authorization rule earlier in this chapter, the profile can be used as a condition of an authorization policy rule in the form of an identity group. Originally, ISE required an identity group to use any of the profiling policies in the rule; however, it has evolved into the ability to use the profile directly (called the EndPointPolicy).

Endpoint Identity Groups

Local identities within the ISE database can be in the form of user identities or endpoint identities. Some identity groups can contain multiple identities, although an identity (user or endpoint) can be a member of only one identity group at a time.

To create an identity group based on the profile, you simply select the labeled Yes, Create Matching Identity Group option on the profile, as displayed in Figure 15-45.



Figure 15-45 Create matching identity group.

If that option is selected, the matching identity group will be found under **Administration > Identity Management > Groups > Endpoint Identity Groups** (as shown in Figure 15-45).

https://t.me/learningnets



Figure 15-46 Endpoint identity groups.

Even though endpoint identity groups were the way to go with ISE 1.0, the use of them for profiling has been deprecated in favor of using the actual endpoint profile or logical profiles directly in the authorization policy.

Therefore, starting with ISE 1.2, endpoint identity groups are used for a different purpose. They are used for more of a MAC Address Management (MAM) model, where you can create a static list of MAC addresses to be authorized specifically. For example, you could create a list of all Apple iPads that are owned by the company so they can be differentiated from personally owned iPads.

The Blacklist identity group is a perfect example of identity group usage in this manner. If a user were to lose his personal device, he would be able to log in to the MyDevices portal and mark a device as lost, as shown in Figure 15-47. This would immediately add the endpoint to the Blacklist group, as shown in Figure 15-48. The device would then be denied network access by default, as shown in Figure 15-49.

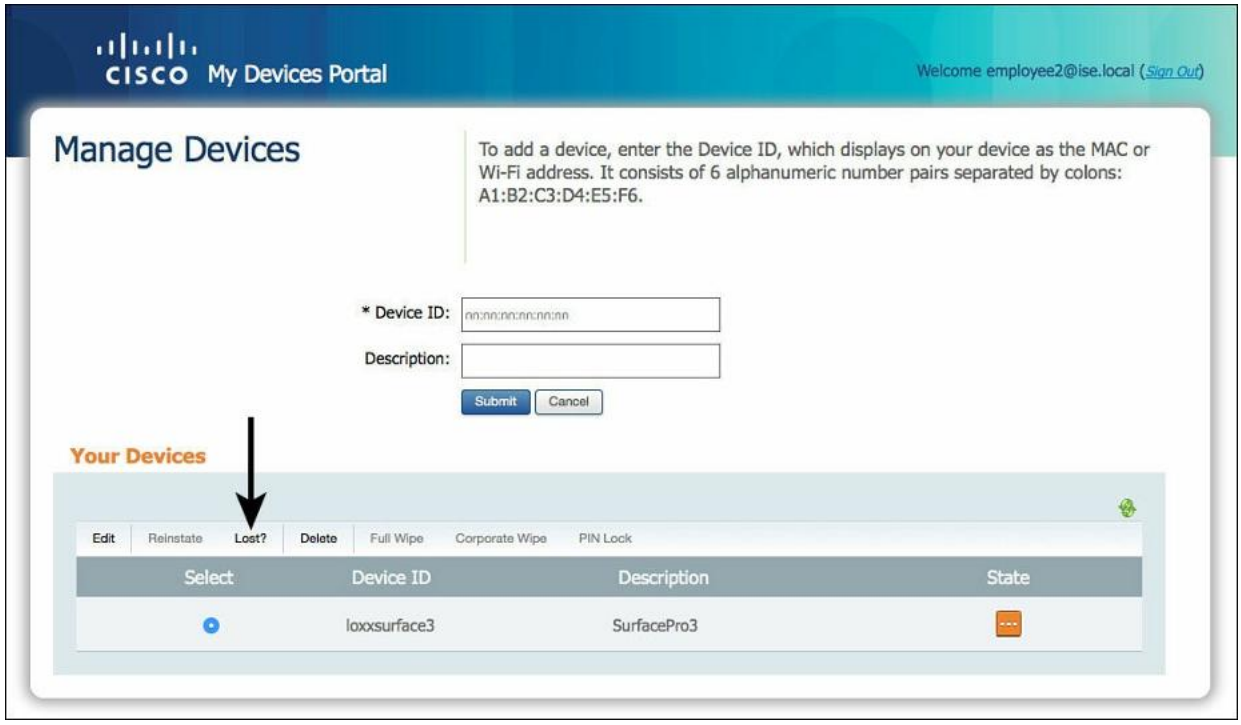


Figure 15-47 MyDevices: Mark endpoint as lost.

https://t.me/learningnets

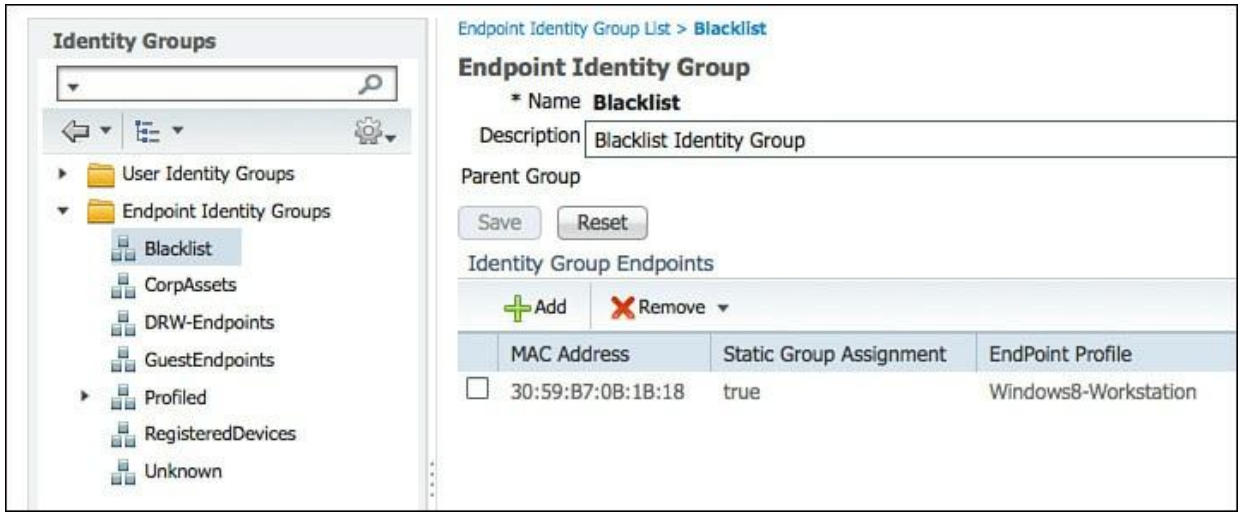


Figure 15-48 Blacklisted devices.

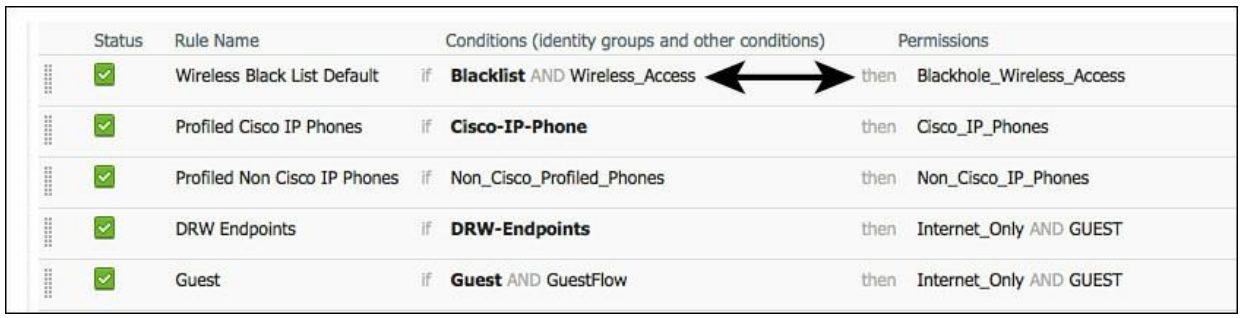


Figure 15-49 Default blacklist authorization rule.

From the MyDevices portal, selecting Reinstate moves the device from the Blacklist group to the RegisteredDevices group. Figure 15-50 shows the MyDevices portal, and Figure 15-51 shows the RegisteredDevices group.

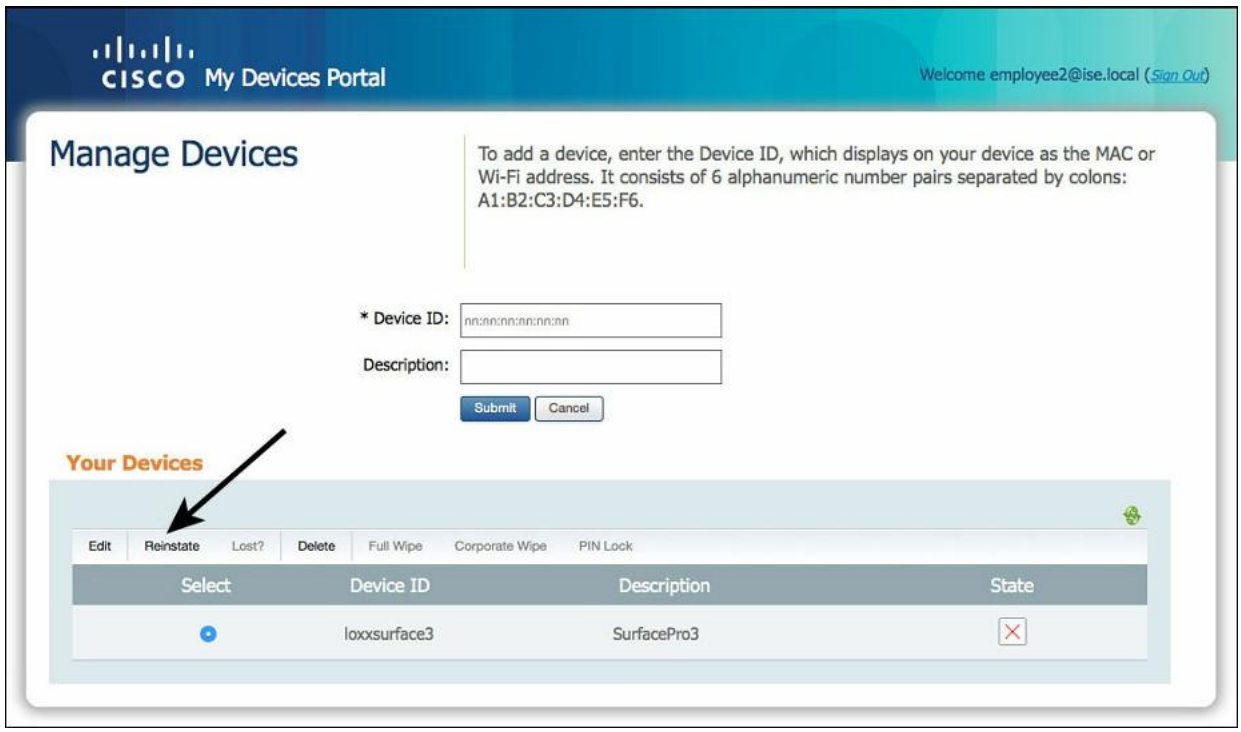


Figure 15-50 MyDevices: Reinstate the endpoint.

https://t.me/learningnets

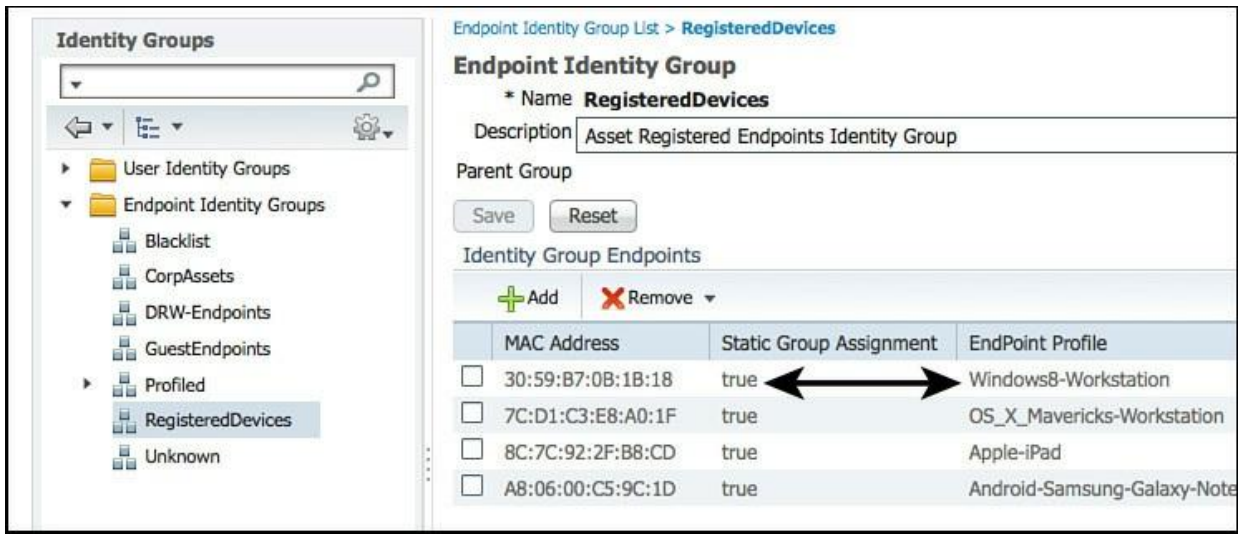


Figure 15-51 RegisteredDevices endpoint identity group.

EndPointPolicy

Beginning with ISE 1.2, identity groups are no longer the way to apply policy based on the endpoint's profile. Policy is now built with the actual profile through the use of the **Endpoints:EndPointPolicy** attribute.

To see the use of the profile in a policy, you will 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.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Duplicate the NSP rule.

Step 3. Add a new condition matching Android devices to the existing MSCHAP condition of **Endpoints > EndPointPolicy**.

Figure 15-52 shows the resulting authorization rule.

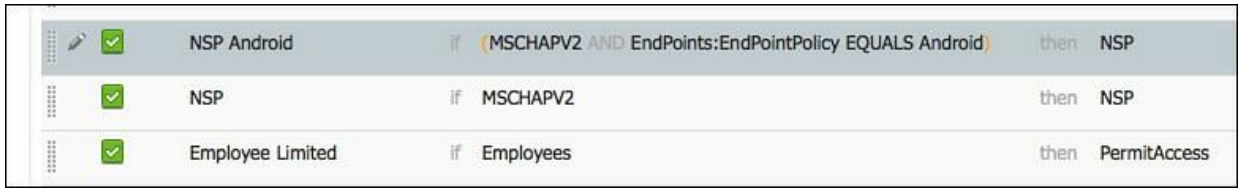


Figure 15-52 Authorization rule using the EndPointPolicy condition.

Throughout all of this, always keep in mind the simplicity of using logical profiles to streamline your authorization policies, and limit the number of individual **EndPointPolicies** that you need to add to your authorization rules.

Verify Profiling

There are a few key places to verify profiling operation. Most exist within the ISE GUI, and of course you should always examine the network device itself.

The Dashboard

The dashboard is always the first screen you see when you log in to the ISE GUI, as shown in Figure 15-53. Right on the dashboard are numerous places to look to see whether profiling is working.

https://t.me/learningnets

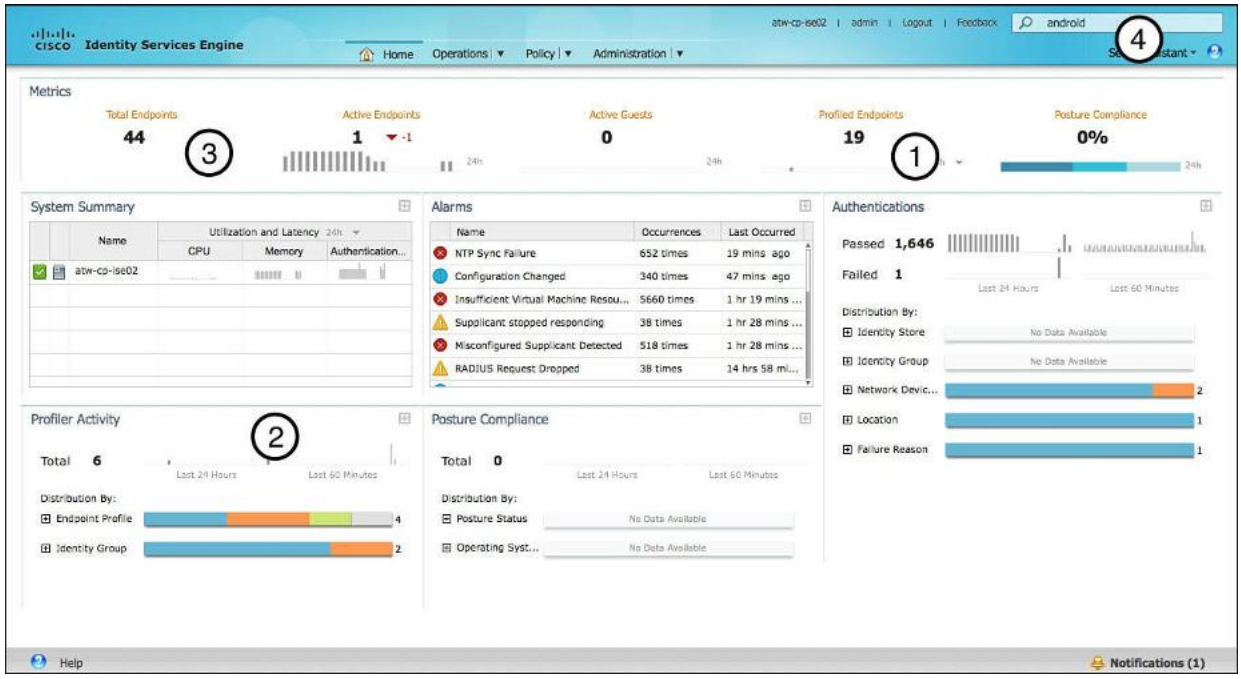


Figure 15-53 The Cisco ISE dashboard.

Examining Figure 15-53, four areas of the dashboard are called out:

- **Profiled Endpoints Counter (1)**—This counter shows the number of endpoints that have been profiled in the last 24 hours.
- **Profiler Activity Window (2)**—This shows a bit more detail than the Profiled Endpoints Counter; it shows two different counters: last 60 minutes and 24 hours. It also lists the individual profiles in a color-coded graph and can be expanded. The entire window can be expanded to see more detail.
- **The Total Endpoints Counter (3)**—This shows the total number of endpoints that have been examined by the MnT persona within the ISE cube. What makes this particular counter so interesting is what happens after you click it. This is examined in detail in the next section, "Endpoints Drill-down."
- **Global Search (4)**—This search bar enables you to search based on the profile, MAC address, or even username. This tool is covered in more detail in the section "Global Search."

Endpoints Drill-down

Often thought of as the best kept secret in ISE, the Endpoints Drill-down tool is frequently overlooked or unknown. Just like the rest of the dashboard, it is part of the MnT persona (not the PSN persona) and therefore displays a result based on the logs sent from the PSN to the MnT node and not the actual endpoint database. Although these can often contain the same data, they can sometimes not match identically, due to log purging.

Regardless of where the data is sourced from, the tool is an excellent way to look into the profiled endpoints and verify that the profiling service is working. Figure 15-54 shows the Endpoints Drill-down tool that appears when you click the total endpoints counter (#3 in Figure 15-53).

https://t.me/learningnets

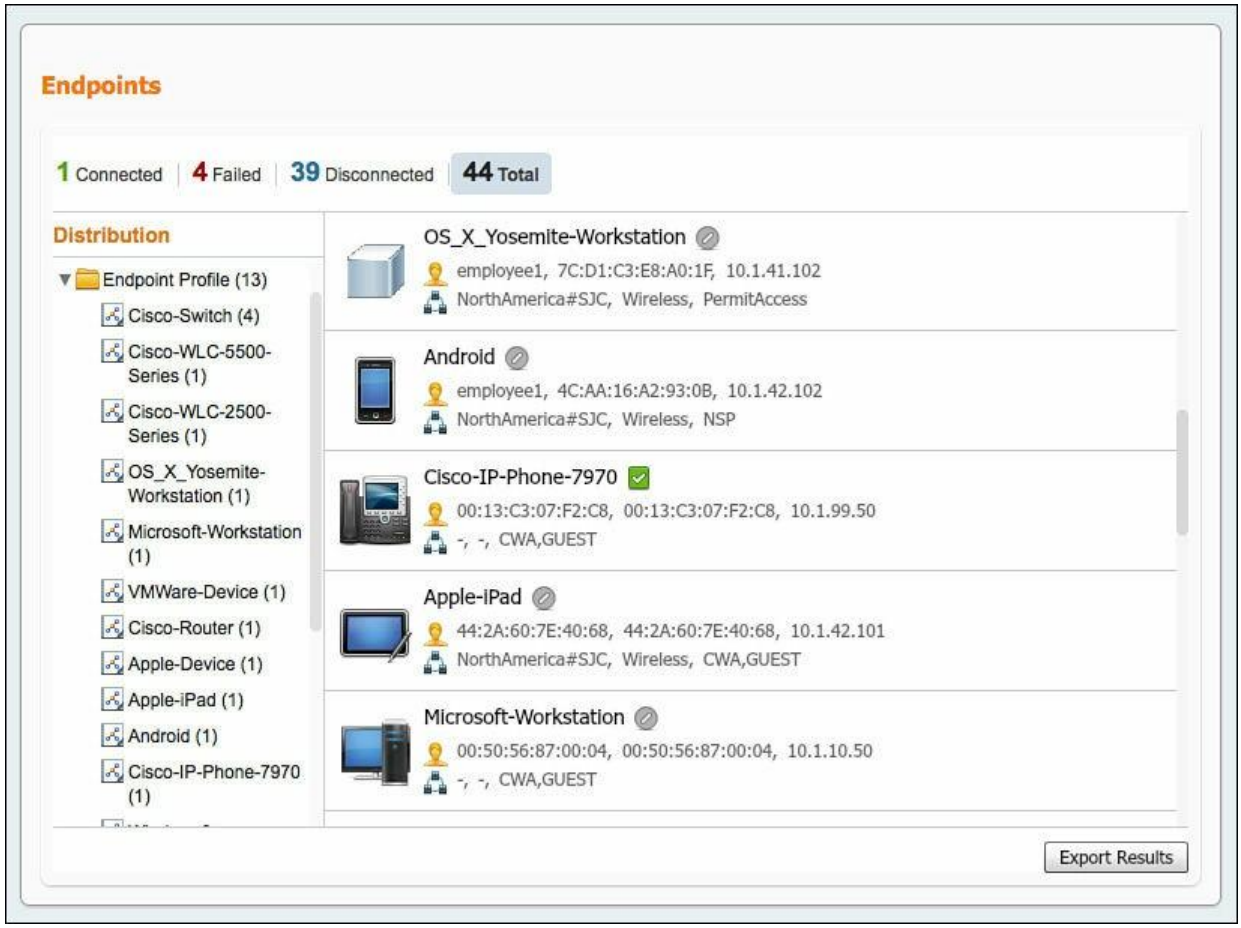


Figure 15-54 Endpoints Drill-down.

Selecting items from the left side of the tool filters the list of endpoints, and each endpoint can be drilled into further on the right side of the tool. In addition, the endpoint list can be exported with the Export button in the lower-right corner.

Global Search

The Global Search tool is available from any screen within the ISE GUI. Typing in a profile name displays results that can be drilled into further. Figure 15-55 shows the use of the Global Search tool to find profiles containing "Apple."

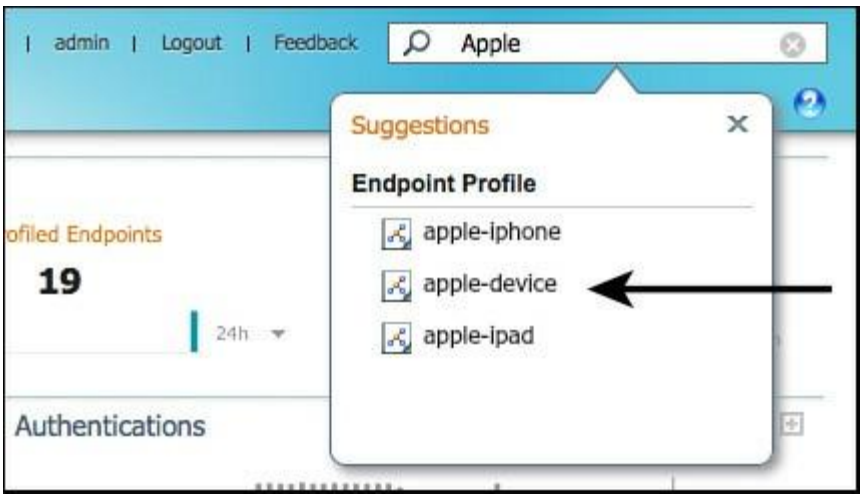


Figure 15-55 Global searching for Apple.

https://t.me/learningnets

Selecting Apple-Device from the resulting list displays a more focused list of endpoints matching that profile, as shown in Figure 15-56.

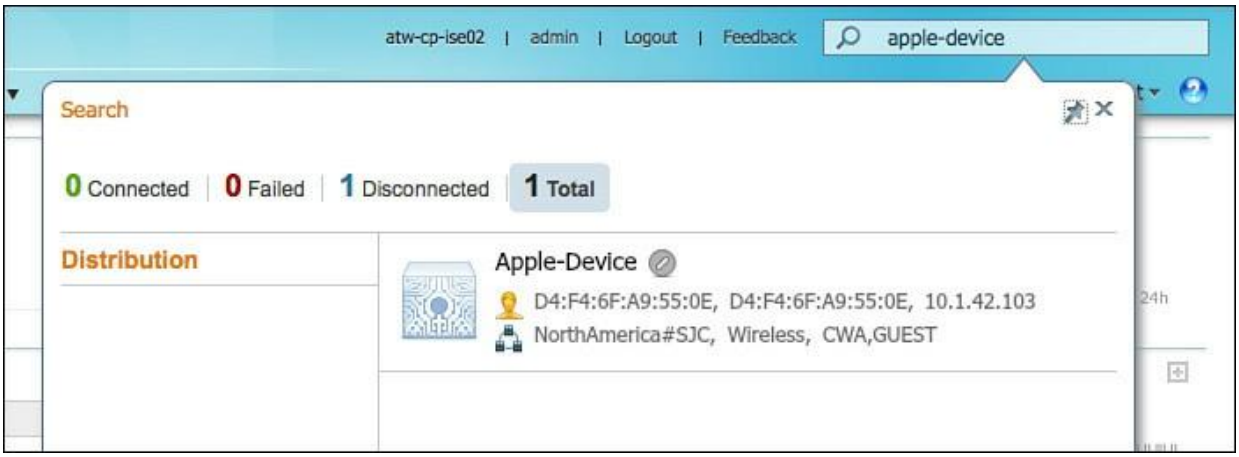


Figure 15-56 Apple-Device.

Endpoint Identities

The ultimate source of truth about endpoints is the actual endpoints database. The graphical view of that list of endpoints is located at **Administration > Identity Management > Identities > Endpoints**. This brings up the endpoints table, and the profile of the endpoint is visible in the first column.

https://t.me/learningnets

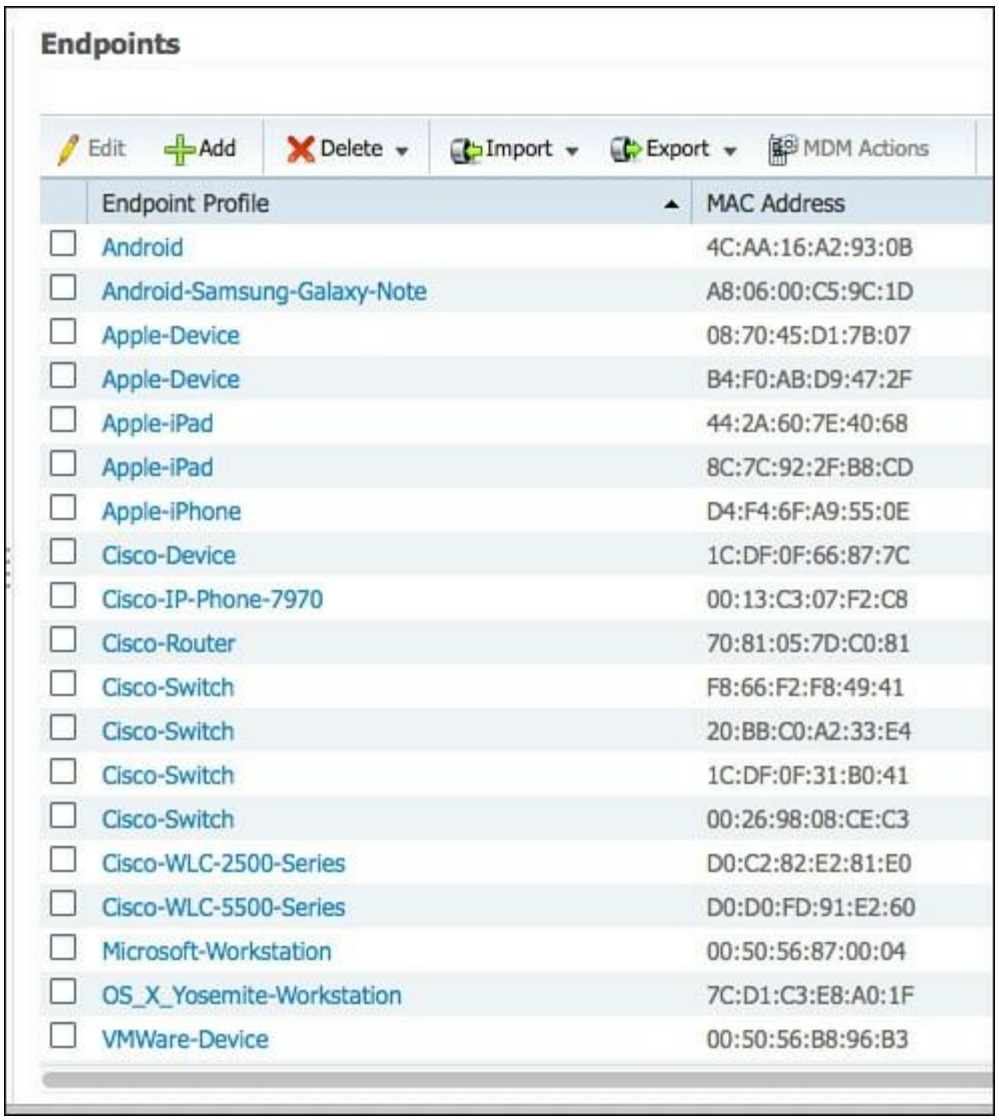


Figure 15-57 Endpoints table.

Clicking any of the endpoints in the table will bring up the endpoint details. The list of details maintained for each endpoint can be quite large, especially if the Endpoint Attribute Filter is not enabled. Figure 15-58 shows the endpoint details screen for an Apple iPhone, and Figure 15-59 shows some of the details further down the list, including the HTTP probe being given credit for the final profile assignment.

https://t.me/learningnets

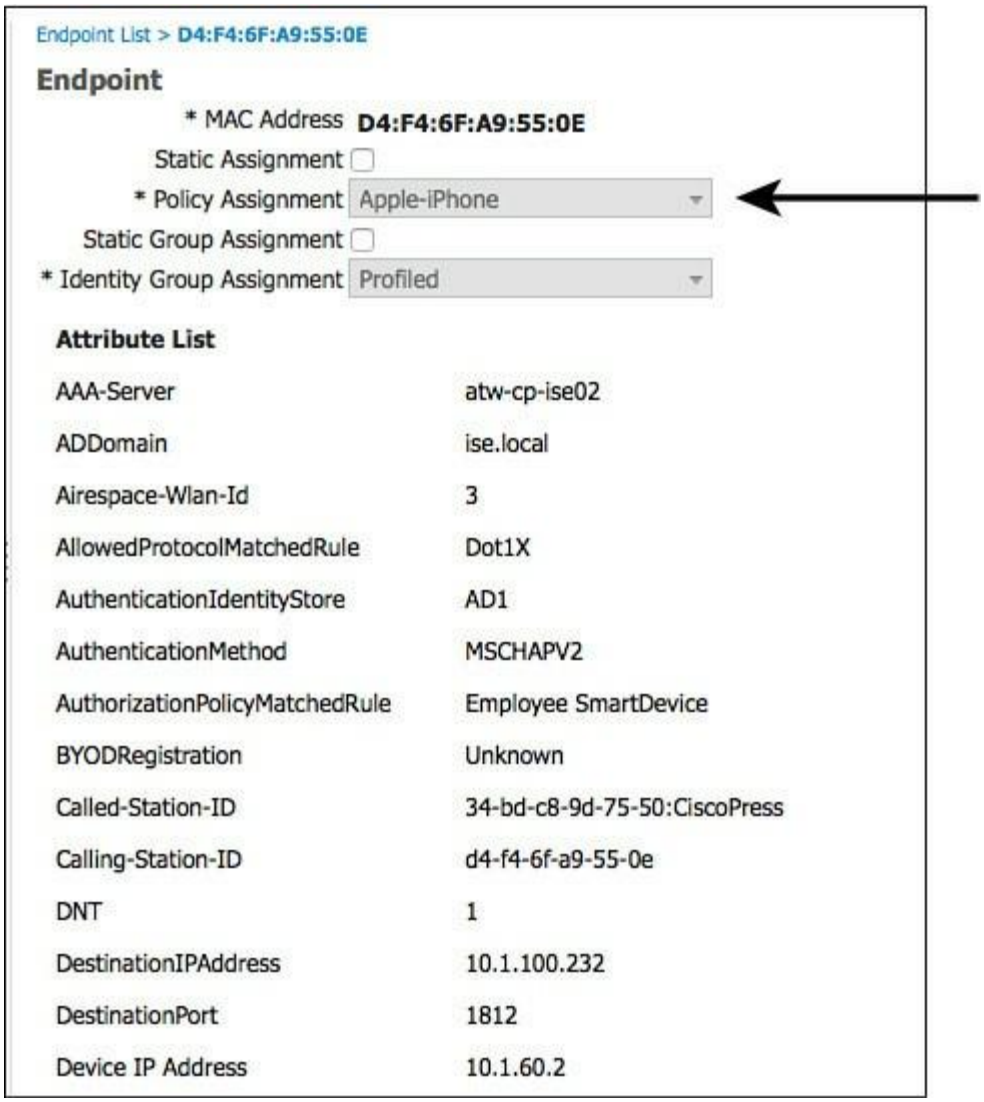
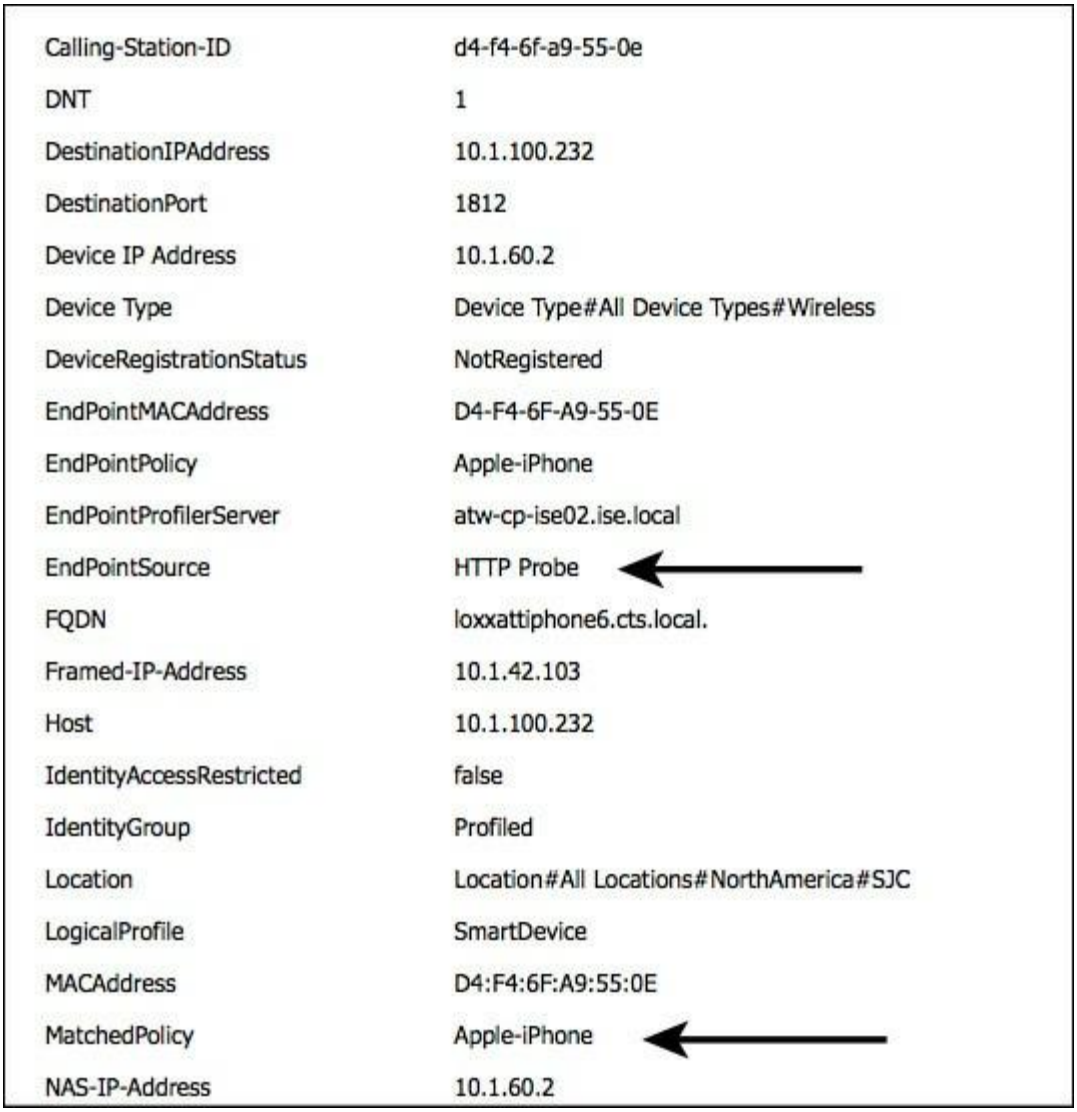


Figure 15-58 Endpoints details.

https://t.me/learningnets



Calling-Station-ID	d4-f4-6f-a9-55-0e
DNT	1
DestinationIPAddress	10.1.100.232
DestinationPort	1812
Device IP Address	10.1.60.2
Device Type	Device Type#All Device Types#Wireless
DeviceRegistrationStatus	NotRegistered
EndPointMACAddress	D4-F4-6F-A9-55-0E
EndPointPolicy	Apple-iPhone
EndPointProfilerServer	atw-cp-ise02.ise.local
EndPointSource	HTTP Probe
FQDN	loxxattiphone6.cts.local.
Framed-IP-Address	10.1.42.103
Host	10.1.100.232
IdentityAccessRestricted	false
IdentityGroup	Profiled
Location	Location#All Locations#NorthAmerica#SJC
LogicalProfile	SmartDevice
MACAddress	D4:F4:6F:A9:55:0E
MatchedPolicy	Apple-iPhone
NAS-IP-Address	10.1.60.2

Figure 15-59 Endpoint details, continued.

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 15-13 displays the output of the **show** command. While the values might not make a lot of sense to a human being, it does demonstrate that the device sensor is collecting and caching profiling data.

Example 15-13 show device-sensor cache all

Click here to view code image

```
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
```

https://t.me/learningnets

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 15-2 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
ISE profiling	With ISE, you can build an authorization policy that combines a user's identity with the classification result and invoke specific authorization results.	447
ISE profiling	Profiling is not an antispoofing solution. Use a defense-in-depth and least-privilege strategy.	453
Logical profiles	Logical profiles are groupings of individual profiles.	478
Global CoA	A CoA will be sent automatically when a device transitions from unknown to any known profile.	480

Table 15-2 Key Topics for Chapter 15

https://t.me/learningnets

Part V: Advanced Secure Network Access

https://t.me/learningnets

Chapter 16. Certificate-Based User Authentications

This chapter covers the following exam 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 that had large security operations (SecOps) teams or for closed-loop systems that hid the certificate usage from the admin, such as Active Directory. Many IT professionals hear the words certificate, X.509, or PKI and run away quietly, pretending they didn't hear.

Well, in today's world of mobile devices, bring-your-own-device IT models, and borderless networks, certificates are quickly becoming the norm. 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 enables 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."

Foundation Topics Section	Questions
Describe identity store options	1-7
EAP types	6-9
Network authorization enforcement	6
Identify issues using authentication event details in Cisco ISE	10
Describe features and functionality of authentication and authorization	1-6, 8-9

Table 16-1 "Do I Know This Already?" Section-to-Question Mapping

> **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 to ISE's Certificate Store.

https://t.me/learningnets

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. 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. Any time leading up to the date listed in the Revocation Date field of the X.509 certificate.

d. The time span between the dates listed in the Valid-From and Valid-To fields of the X.509 certificate.

3. True or False? Certificate Revocation List (CRL) is the only revocation status mechanism supported by ISE.

a. True

b. False

4. True or False? ISE will ignore the CRL distribution point listed in the X.509 client certificate.

a. True

b. False

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 screen shot 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 it cross-references it to the end user to which the certificate was issued, returning 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 principle 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 username and password for the Active Directory component.

7. Which 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, Signing CA's Public Certificate added to the

https://t.me/learningnets

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, 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, 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 an ISE administrator?

a. A form for the admin to fill out and request the CA administrator send its public key, including any intermediary CAs.

b. Configures the Windows client to provide the signer's public key during the authentication process, along with its own (hence, its certificate chain).

c. Downloads a PKCS file, which is a certificate chain file that will contain the public certificates for the CA and any intermediate CA in the hierarchy.

d. Redirects the admin to a new page where she can purchase the public key from the certificate authority.

10. Live Log provides a glance at a lot of information, including a brief failure reason. What should an admin do to find a more detailed explanation of the failed certificate authentication and a possible resolution?

a. From Live Log, navigate to **Operations > Reports > Failed Authentications**.

b. From Live Log, click the Details icon, which will launch the authentication details report.

c. Immediately email Aaron Woland of 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 authentication does not have to be scary; there are a few universal truths when discussing certificate-based authentications. Standard certificate authentication is the same, regardless of whether the certificate authority is built by Microsoft, Cisco, Symantec, Entrust, and so on.

Let's start with a quick definition of a certificate. A certificate is a signed document that represents an identity. When thinking of a certificate, try to relate it to a passport, a driver's license, or another personal identification card. That identification card is meant to represent you and prove you are who you say you are.

How does a business or security checkpoint know to trust that identification?



When presented with a certificate, an authentication server will perform the following checks (at a minimum):

1. Determine whether a **trusted authority** has signed the digital certificate.
2. Examine both the start and end dates to determine whether the certificate has expired.
3. Verify whether the certificate has been revoked (you could use OCSP or CRL check).
4. Validate that the client has provided proof of possession.

https://t.me/learningnets

Let's examine these four items one at a time:

Determine Whether a Trusted Authority Has Signed the Digital Certificate

The signing of the certificate really has two parts.

The first part is that the certificate must have been signed correctly (following the correct format and so on). If it is not, it will be discarded immediately. Next, the signing certificate authority's (CA'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 CA's public key must be stored at **Administration > System > Certificates > Certificate Store**, and it must have the Trust for client authentication use-case.

Figure 16-1 shows ISE's certificate store and the Trust for Client Auth field.

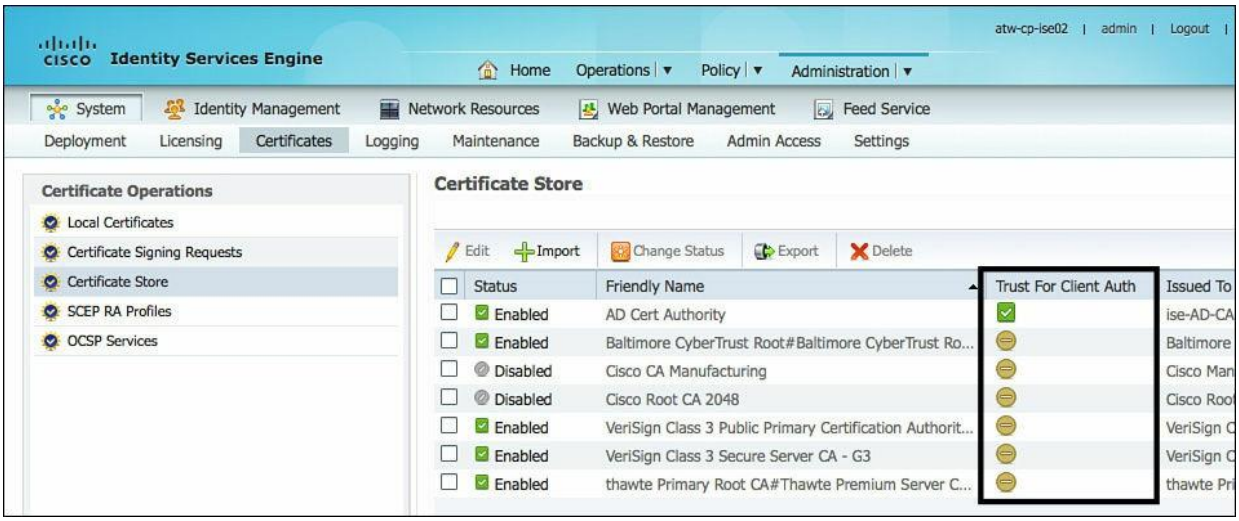


Figure 16-1 Administration > System > Certificates > Certificate Store.

Figure 16-2 shows the details of the public key belonging the Active Directory CA and has the Trust for Client Authentication check box enabled.

https://t.me/learningnets

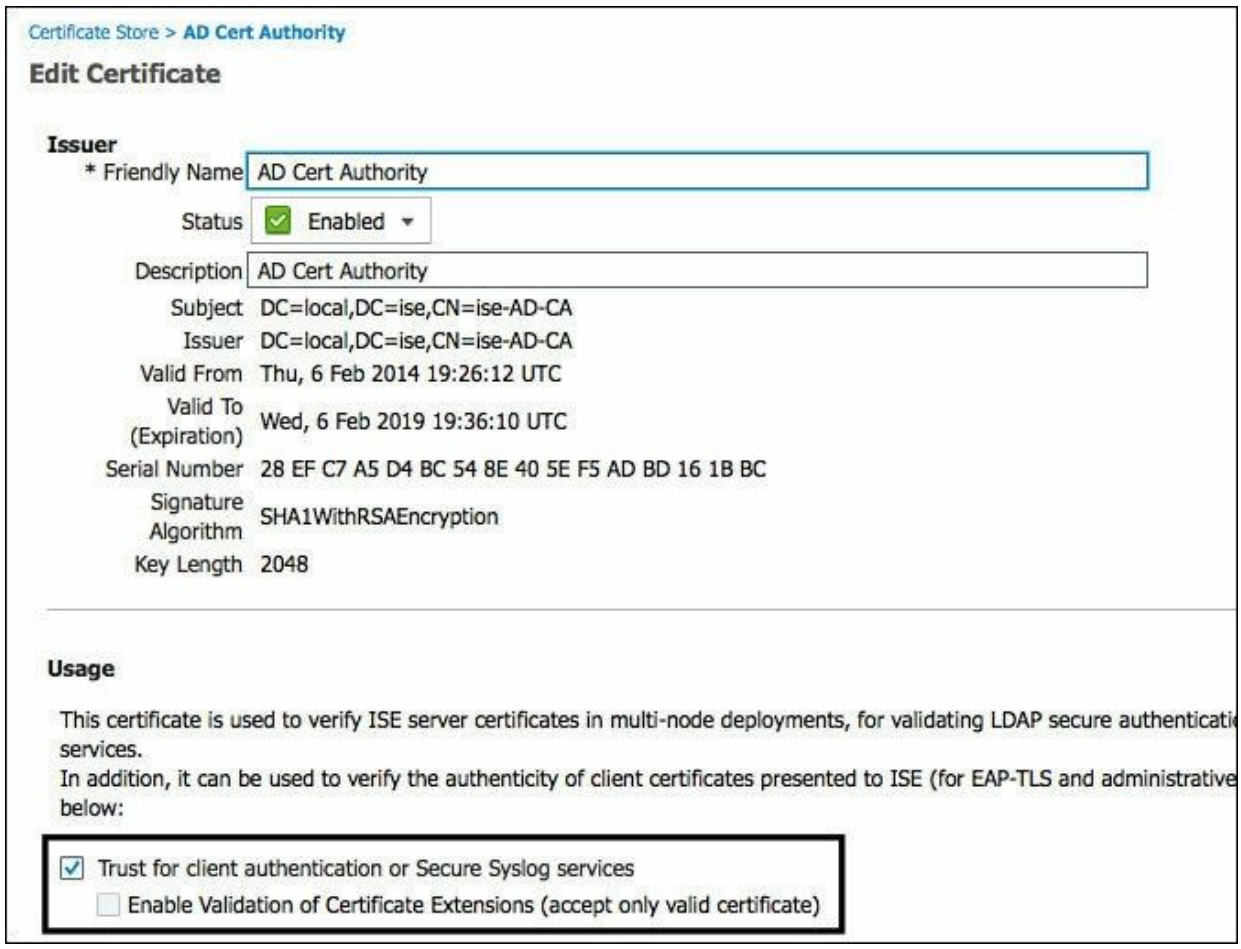


Figure 16-2 Active Directory CA certificate details.

So not only does ISE "trust" certificates that have been signed by this CA, but it also 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, then the authentication will fail. It's exactly like someone entering the wrong password or a bartender spotting a fake ID. Access rejected!

Examine Both the Start and End Dates to Determine Whether the Certificate Has Expired

Continuing to use the driver's license or passport as an example, two dates are usually listed on the document: the date 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 were trying to get through security at an airport and presented the TSA agent with an expired driver's license, he is not supposed to allow you through the gate.

The picture on the identity is still clearly you, and it was obviously a valid identification card once, so why would they ever reject it? Ultimately, because the signing authority (DMV) is no longer warrantying that the document is valid, that's why.

A certificate will also have two dates listed in it: a date issued and date it is valid until (when it expires). The authentication server performs a similar check to what the TSA agent does and should reject any certificate that is not within its validity range. In other words, it will check to see whether the certificate is valid for the date and time at which the authentication request comes in.

This is one reason Network Time Protocol (NTP) is so important when working with certificates. Many of us have seen problems where time was out of sync. For example, a certificate was presented on January 10, 2014, at 11:11 a.m., but its "valid-from" value started on January 10 at 11:30 a.m. In this example, the certificate was issued for a time 19 minutes in the future. This was because of a time sync issue where the CA thought it was 20 minutes later than the authentication server, so the brand-new certificate was not valid yet! This type of problem is so common you would laugh (or cry).

https://t.me/learningnets

Figure 16-3 shows a certificate and highlights the Valid-From and Valid-To attributes.

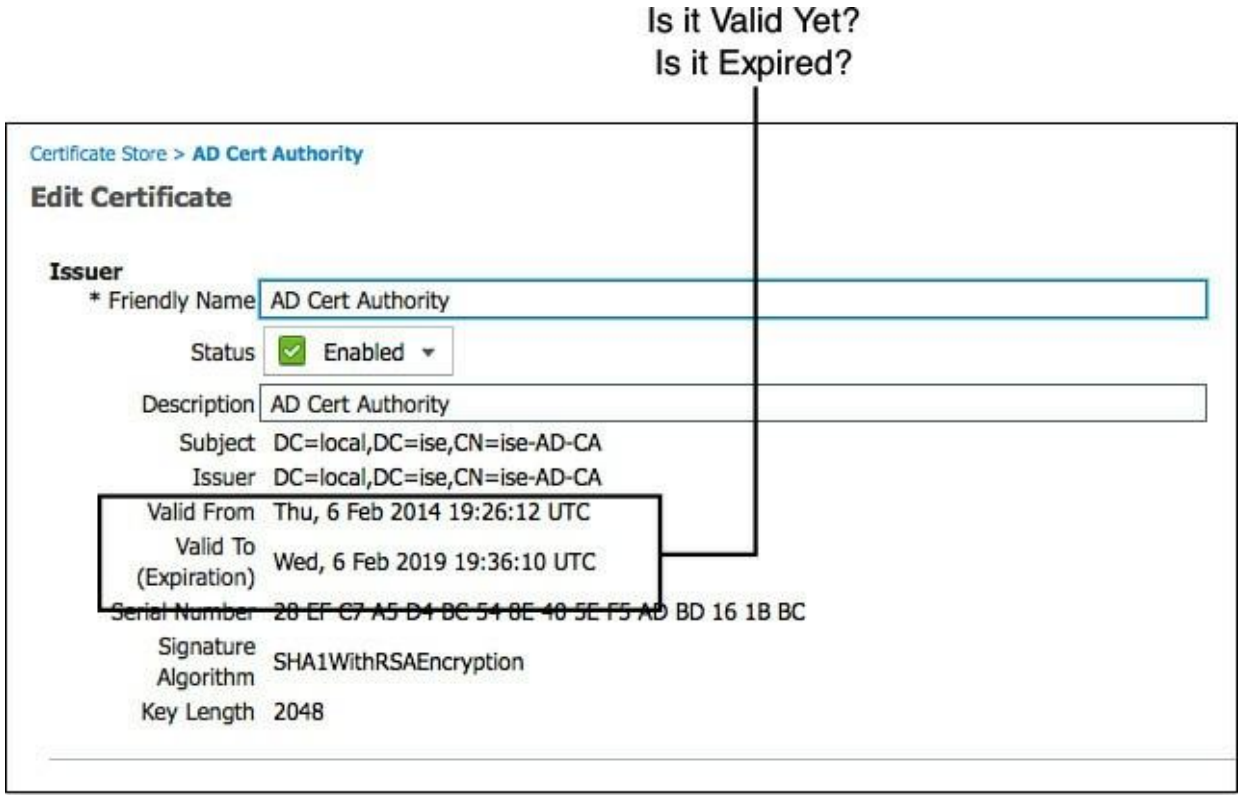


Figure 16-3 Certificate validity period.

Verify Whether the Certificate Has Been Revoked

Pretend for a moment that you are driving down the road and are pulled over by a police officer. The officer asks for your driver's license and proof of insurance. You hand him a driver's license, which is immediately checked for evidence of authenticity—that is, does it look like a valid driver's license or a forgery? Okay, it's not a forgery. Check. Next, expiration? It is not expired. Check. Now 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 (are you a registered owner of the car you were driving, and so on). Those are not important for this hypothetical scenario, though. What is important is that the officer must make sure your valid driver's license was not revoked by the Department of Motor Vehicles. A quick lookup on the computer into the DMV records shows that your driver's license was revoked for too many DWIs. The cold steel of the handcuffs and the rough shove into the back of the squad car as you are hauled off to jail make you reevaluate your life choices.



Certificate authentication has the same capability (not the handcuffs, I'm talking about the lookup into the DMV to check on revocation status). Every CA should also have a service to publish a list of certificates that have been revoked. The two main ways to do this today are

- Certificate Revocation List (CRL)—This is basically a signed list that the CA publishes on a website that can be read by authentication servers. The file is periodically downloaded and stored locally on the authentication server; when a certificate is being authenticated, the server examines the CRL to see whether the client's cert was revoked. A CRL could be compared to the police officer having a list of suspended drivers in his squad car. Much like that list, the CRL is downloaded and cached for a period of time and therefore might not be valid in some cases (such as a certificate that has been reinstated).

https://t.me/learningnets

- Online Certificate Status Protocol (OCSP)—This 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 police officer using the computer in the squad car and doing a lookup into the DMV's database.

If the certificate has been revoked, then access is denied. Enjoy the lights and sirens on the way to jail.

Figure 16-4 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 CRL on a per-trusted CA basis. When ISE is authenticating a certificate that has been signed by that particular root (or it's subordinates), ISE will use the configured service to lookup revocation status.

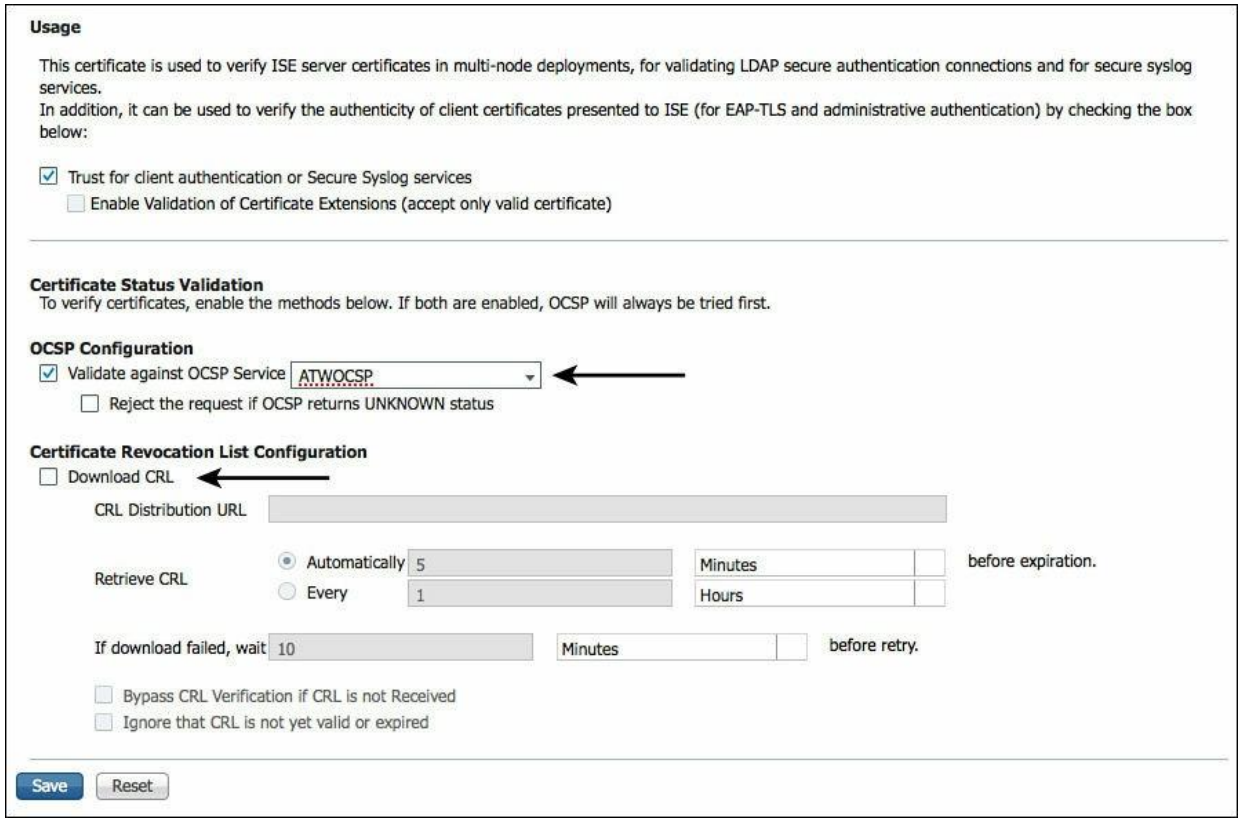


Figure 16-4 Configuration for certificate revocation checking.

It's very important for you to understand that checking for certificate revocation is an option. You'll notice that neither check box for CRL or OCSP are on by default; they require the admin to configure the URL or the service location. It is also critical to understand what behavior will happen if the service is not available or the status of the certificate is unknown—that is, how does the authentication policy handle exceptions? Is it configured to "fail-open" or "fail-closed"?

The client's certificate itself will have an extension called CRL Distribution Points, which can be populated with the URI where the authentication server may locate the CRL. Figure 16-5 illustrates this certificate attribute. However, it is important to note that Cisco ISE will ignore this field and use only the configured CRL point shown in Figure 16-4.

https://t.me/learningnets

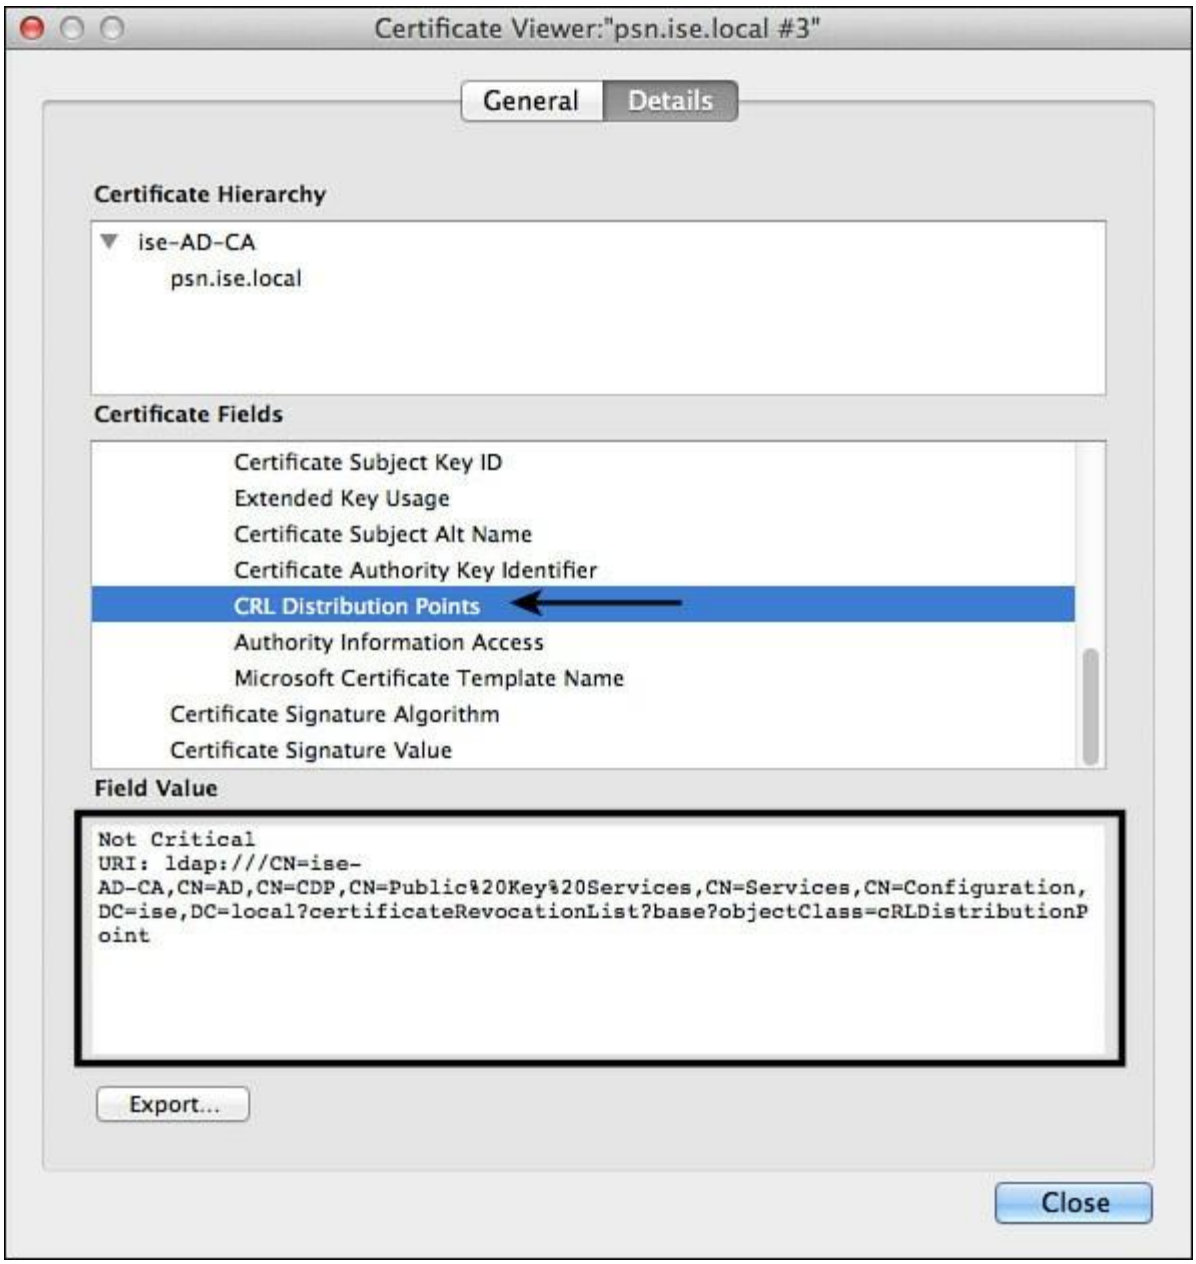


Figure 16-5 CRL distribution points attribute.

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 that 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 at manageable sizes.

Validate That the Client Has Provided Proof of Possession

Is there a way for an authentication server to be sure the client truly owns the certificate? Think of the example of being pulled over by a police officer for speeding (yes, again). You hand the officer a driver's license, and she evaluates it. The results of the evaluation are as follows:

- It is a valid driver's license, and has it been issued by a trusted signer (the state DMV)?
- It has not expired yet, so that is no problem.
- The police officer uses the computer to validate with the DMV that the driver's license has not been revoked.

https://t.me/learningnets

All three of those have passed their evaluation checks. However, the picture on the driver's license is a picture of a woman with long, flowing brown hair and hazel eyes, yet you yourself are a bald, elderly man. OOPS! This "valid" driver's license was not issued to you—the proof of possession has failed! Again—jail time!

Certificate authentications do something similar. There will 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 it is the proof of possession. This ensures that someone did not just grab the client's public key and try to present that as being his own. If the client cannot provide proof of possession, then the authentication will fail—Access-Reject.



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 this interesting is that something that has been encrypted using the private key can be decrypted only by using the public key. Additionally, if something was encrypted using the public key, it can be decrypted only by using the private key. These public and private keys are mathematically married. That is why the proof of possession works to ensure 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 (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 (AD) has a certificate authority built in to it that can be enabled as a full enterprise-wide 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. This is not the case.

What can often confuse people with regard to Authentication Authorization and Accounting (AAA) is the difference between authentication and authorization. They often blend so much. A certificate issued by Active Directory Certificate Services is still just an X.509 certificate. It will go through all the authentication validation of any other certificate, regardless of the fact that the CA was integrated into AD. Authenticating the certificate itself is not using Kerberos, such as AD authentications do; it doesn't retrieve any Active Directory attributes, such as group membership; and no AD permission/rights are assigned because of the certificate itself.

What is possible, however, is to examine a field of the certificate and then to do a separate lookup into AD based on that field during the authorization phase. For example, a certificate with a subject of 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 will take the certificate subject (Aaron) and do a lookup into AD for that username. This is where group membership and other policy conditions are examined and the specific authorization result issued.

Cisco ISE uses something called a Certificate Authentication Profile (CAP) to examine a specific field and map it to a username for authorization. Figure 16-6 shows a sample CAP.

https://t.me/learningnets

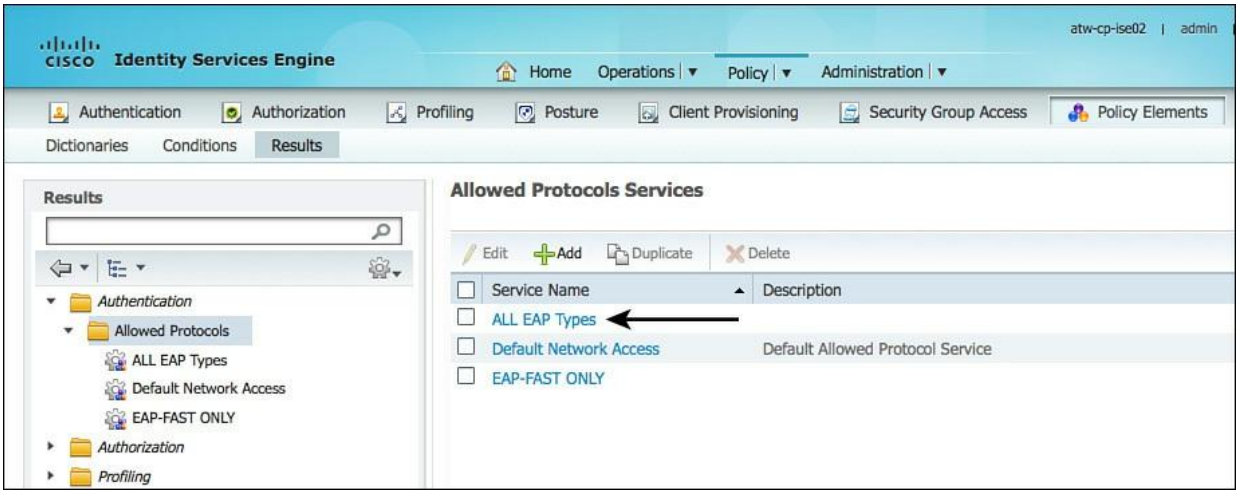


Figure 16-6 Allowed Protocols Services.

> **Note**
>
> Cisco ISE will also do a courtesy check to validate whether the machine or account has been disabled in AD. If the account were disabled in AD, then the authorization will be to deny access.

Please note, this is a very different process from an Active Directory authentication, which uses Kerberos, meaning AD logs will be recorded differently. There are solutions on the market that examine AD log files and use that information to help tie together usernames and IP addresses for single-sign-on to web proxy servers, identity-enabled firewalls, and other services (Cisco Context Directory Agent [CDA] is one such example).

If the authentication was a certificate-based authentication (EAP-TLS) but the user was authorized from an AD lookup, that process will most likely not 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 local area networks (LANs), EAP-TLS will be 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 use the bring-your-own-device (BYOD) onboarding functionality with ISE that will provision the native supplicant for Windows, MAC, iOS, and Android devices in addition to issuing that device a certificate for use with EAP-TLS.

For more on EAP methods, see Chapter 4, "EAP over LAN (also Known as 802.1X)." For more on device onboarding, see Chapter 17, "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 will use certificates:
  - A CAP must be created and used in the authentication rule:
    - The CAP should define which certificate field will be the Principle Username X.509 Attribute.
    - Optionally, you can add that CAP to an identity source sequence (ISS) that includes Active Directory for authorization purposes.

https://t.me/learningnets

- Cisco ISE must be configured to trust the client certificates:
  - Those certificates must be trusted for client authentication.
  - Optionally, you can configure OCSP or CRL checking for the CAs.
- Configure the authorization policy.

Validate Allowed Protocols

To process a certificate-based authentication, the Allowed Protocols definition will need to permit EAP-TLS. The definition can be configured to allow native EAP-TLS itself, within a tunneled EAP type such as PEAP or EAP-FAST, or possibly allowing all options regardless of using native or tunneled transport.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Policy Elements > Results > Authentication > Allowed Protocols**.

In Chapter 10, "Authentication Policies," you created an Allowed Protocols definition named ALL EAP Types, as shown in Figure 16-6. If you did not create this one already, feel free to skip Step 2 and follow along using whichever Allowed Protocols definition you will be using in your own policy.

Step 2. Click **All EAP Types**.

Step 3. Ensure that EAP-TLS is enabled natively, as well as within the PEAP and EAP-FAST tunnels, as shown in Figure 16-7.

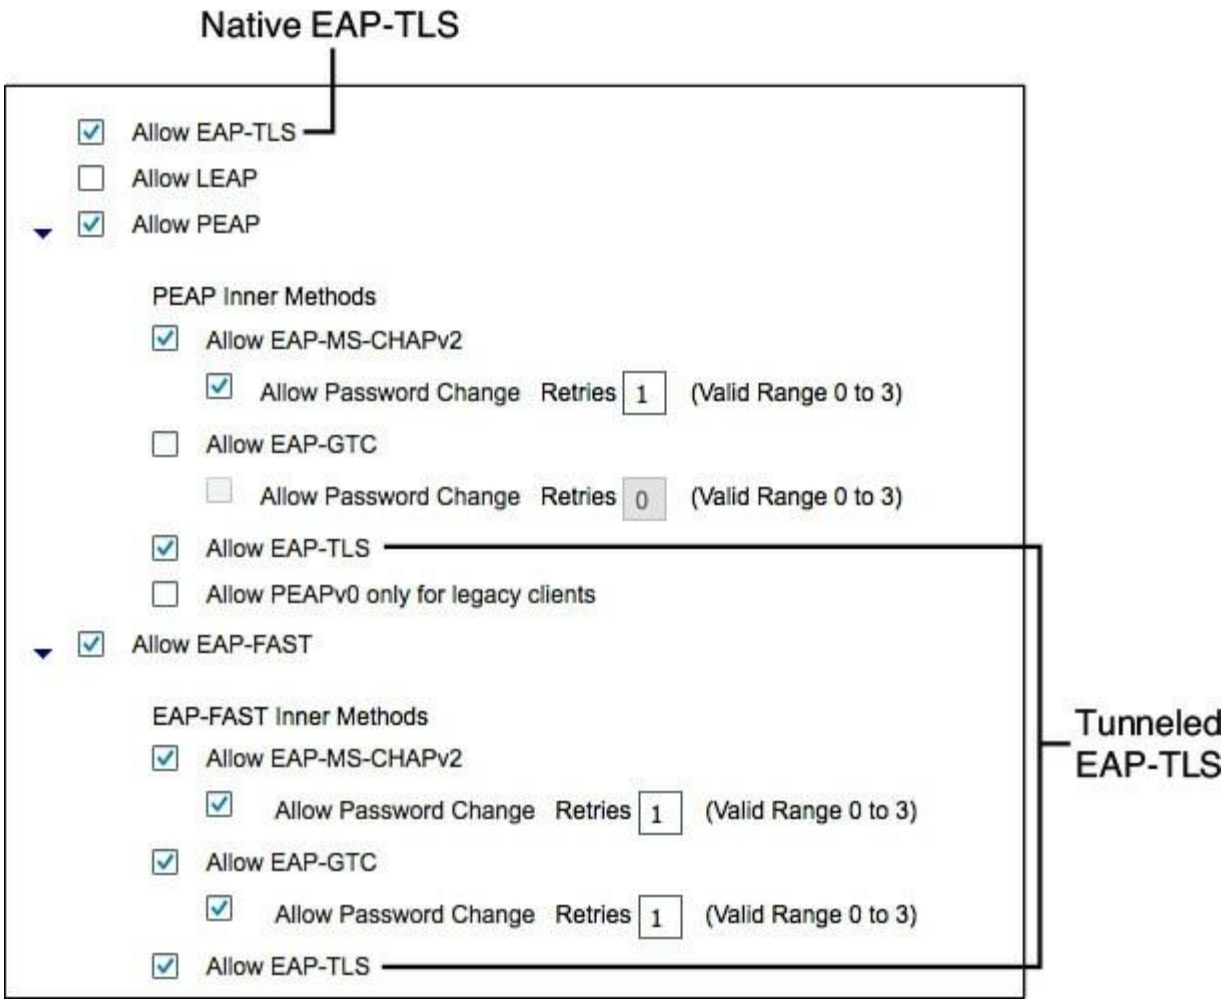


Figure 16-7 EAP-TLS enabled natively and within the EAP tunnels.

Now you are certain that EAP-TLS is allowed in this particular Allowed Protocols definition. Next, you will

https://t.me/learningnets

ensure there is a CAP for the Authentication Identity Store.

Certificate Authentication Profile

As described in Chapter 3, "Identity Management," and Chapter 10, ISE uses a certificate authentication profile as the identity source for certificate-based authentications. In reality, the authentication is simply validation of the certificate, as described in detail within this chapter.

Keep in mind that a certificate is really just a signed document and that the document 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 Principle Username X.509 Attribute.

To verify the CAP has been created correctly, take a look at the CAP you created in Chapter 10.

From the ISE GUI, do the following:

Step 1. Navigate to **Administration > Identity Management > External Identity Sources > Certificate Authentication Profile**.

Step 2. Click the CAP you created in Chapter 10 (ATW_CAP in the example). If you have not created one yet, click **Add**.

Figure 16-8 shows the sample CAP that was created in Chapter 10. Notice in the figure that the Principle Username X509 Attribute is configured to use the Common Name attribute from the certificate's subject field as the Principle Username.



Figure 16-8 Certificate authentication profile.

As described throughout this chapter, certificate-based authentications validate 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 are using certificate-based authentications for Windows AD member computers and want 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 Perform Binary Certificate Comparison with Certificate Retrieved from LDAP or Active Directory with the Active Directory as the source for comparison.

https://t.me/learningnets

Verify That the Authentication Policy Is Using CAP

You have just ensured a CAP exists to use with the authentication policy. Now you will examine the authentication policy and ensure the CAP is being used for EAP-TLS traffic.

In Chapter 10, you configured a detailed authentication policy that was very specific regarding which identity store to use with EAP protocol. This was designed to help you learn how the authentication policies work. At the end of the chapter, you modified the configuration to be simple but still work with all the use cases throughout this book. Now, you will verify that the policy is ready for certificate-based authentications.

You will ensure that there is an authentication rule for 802.1X configured to use the ALL EAP Types Allowed Protocols definition and that it will use the CAP as the identity source or within an SISS.

From the ISE GUI, follow these steps:

Step 1. Navigate to **Policy > Authentication**. The authentication policy should look similar to what is shown in Figure 16-9.

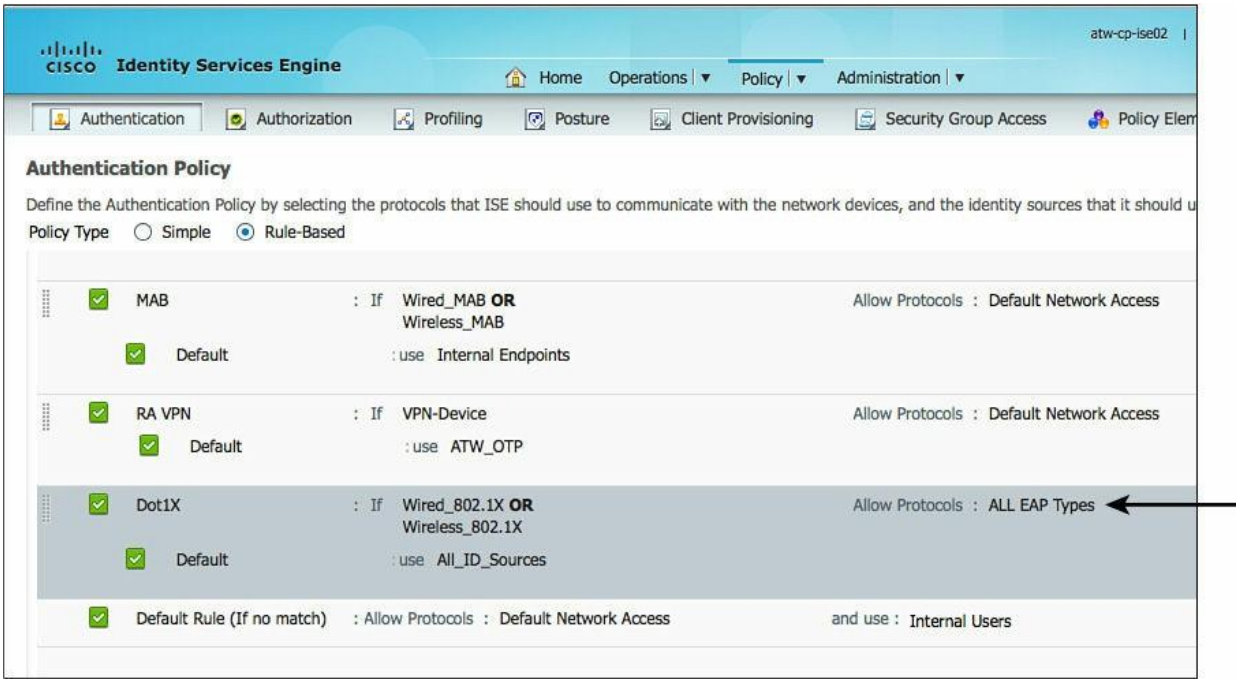


Figure 16-9 Simplified authentication policy from Chapter 10.

Step 2. Visually ensure that the Allowed Protocols is set for **ALL EAP Types**, as displayed in Figure 16-9.

Step 3. Edit the Dot1X rule.

Step 4. Hover over the identity source, which should be All_ID_Sources, and ensure that the Certificate Authentication Profile Name is there, such as ATW_CAP in this example (see Figure 16-10).

https://t.me/learningnets

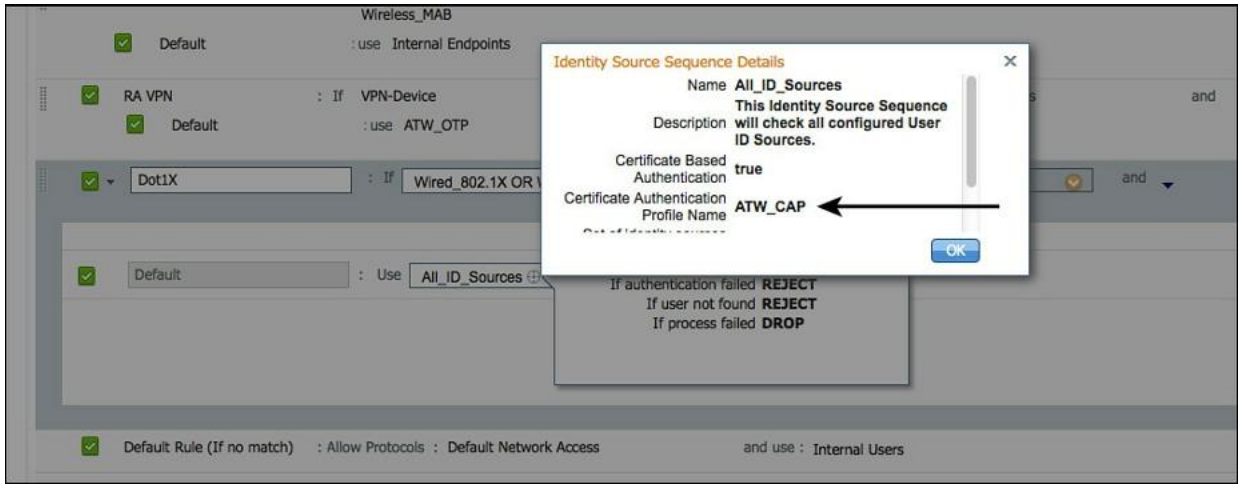


Figure 16-10 All_ID_Sources ISS contains the CAP.

> **Note**
>
> You could also examine this ISS by navigating to Administration > Identity Source Sequences and editing All_ID_Stores.

Step 5. Click **Done**.

You've now ensured that the Dot1X authentication rule will permit EAP-TLS and will authenticate to an ISS that includes the CAP. The authentication rule has been confirmed as ready, and you will now think about the authorization policy.

Authorization Policies

With the CAP, you specified which field to use for the principle identity. That identity can be used in the authorization policy, just as if a user typed in her ID and provided a password. You can examine AD group membership or other AD attributes for that identity, as well as compare fields in the certificate itself. You might choose to authorize the session based on the signing CA or perhaps assign a particular virtual LAN (VLAN) because the organizational unit (OU) of the certificate states it is for the HR department (for example).

To help foster a better understanding of certificate-based authorizations, let's create an authorization rule for employees to gain access to the network.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Duplicate the rule named Employee Limited by selecting **Duplicate Above** from the drop-down arrow next to **Edit**.

Step 3. Rename the duplicated rule **Employee Limited_Certificate**.

Step 4. For the conditions, select **Add a New Attribute / Value** from the condition Picker (the + sign).

Step 5. Select **Network Access > Authentication Method EQUALS x509_pki**, as shown in Figure 16-11.

https://t.me/learningnets

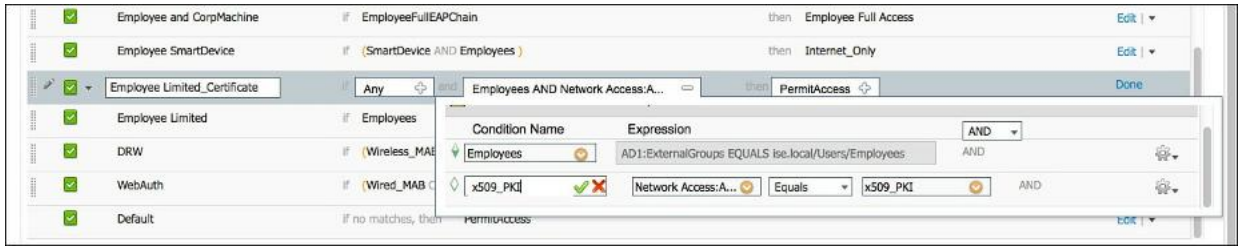


Figure 16-11 Employee Limited_Certificate conditions.

Step 6. Add another New Attribute / Value, and then select **Certificate**. You are doing this just to see which values are there; you are not adding another attribute at this time.

Step 7. Delete the unused attribute from the conditions list, so your conditions look like the ones in Figure 16-11.

Step 8. Add the condition to the Library for reuse in other rules, and name the new condition **x509_PKI**, as shown in Figure 16-11.

Step 9. Click **Done**.

Step 10. Click **Save**.

Figure 16-11 shows the Employee Limited_Certificate conditions. Your end result should look the same.

You've just created a policy that used a signed document and X.509 certificate and ran authorization against the AD membership of a user ID, using the Employees Group from Active Directory. This is due to the CAP extracting that identity from a field in that certificate and using that extracted identity for the authorization.

Additionally, you saw the options of using other information from the certificate to help make your authorization decisions.

You will examine these types of policies in more detail in Chapter 17.

Ensuring the Client Certificates Are Trusted

A step that is often overlooked is to ensure that ISE trusts the client certificate. For a certificate to be trusted for client authentication, you must trust the CA that has signed the certificate.



Keep in mind that X.509 provides for a hierarchy that enables scale. A certificate authority can actually belong to a full "tree" of CAs, all stemming from the root CA. Figure 16-12 illustrates this concept, showing a root CA with two subordinate CAs. Below the subordinate CA is the certificate holder (a tablet). When ISE is presented with the certificate, it should trust all the CAs in the path. Using Figure 16-12 as an example, ISE should trust the subordinate CA and the root CA.

https://t.me/learningnets

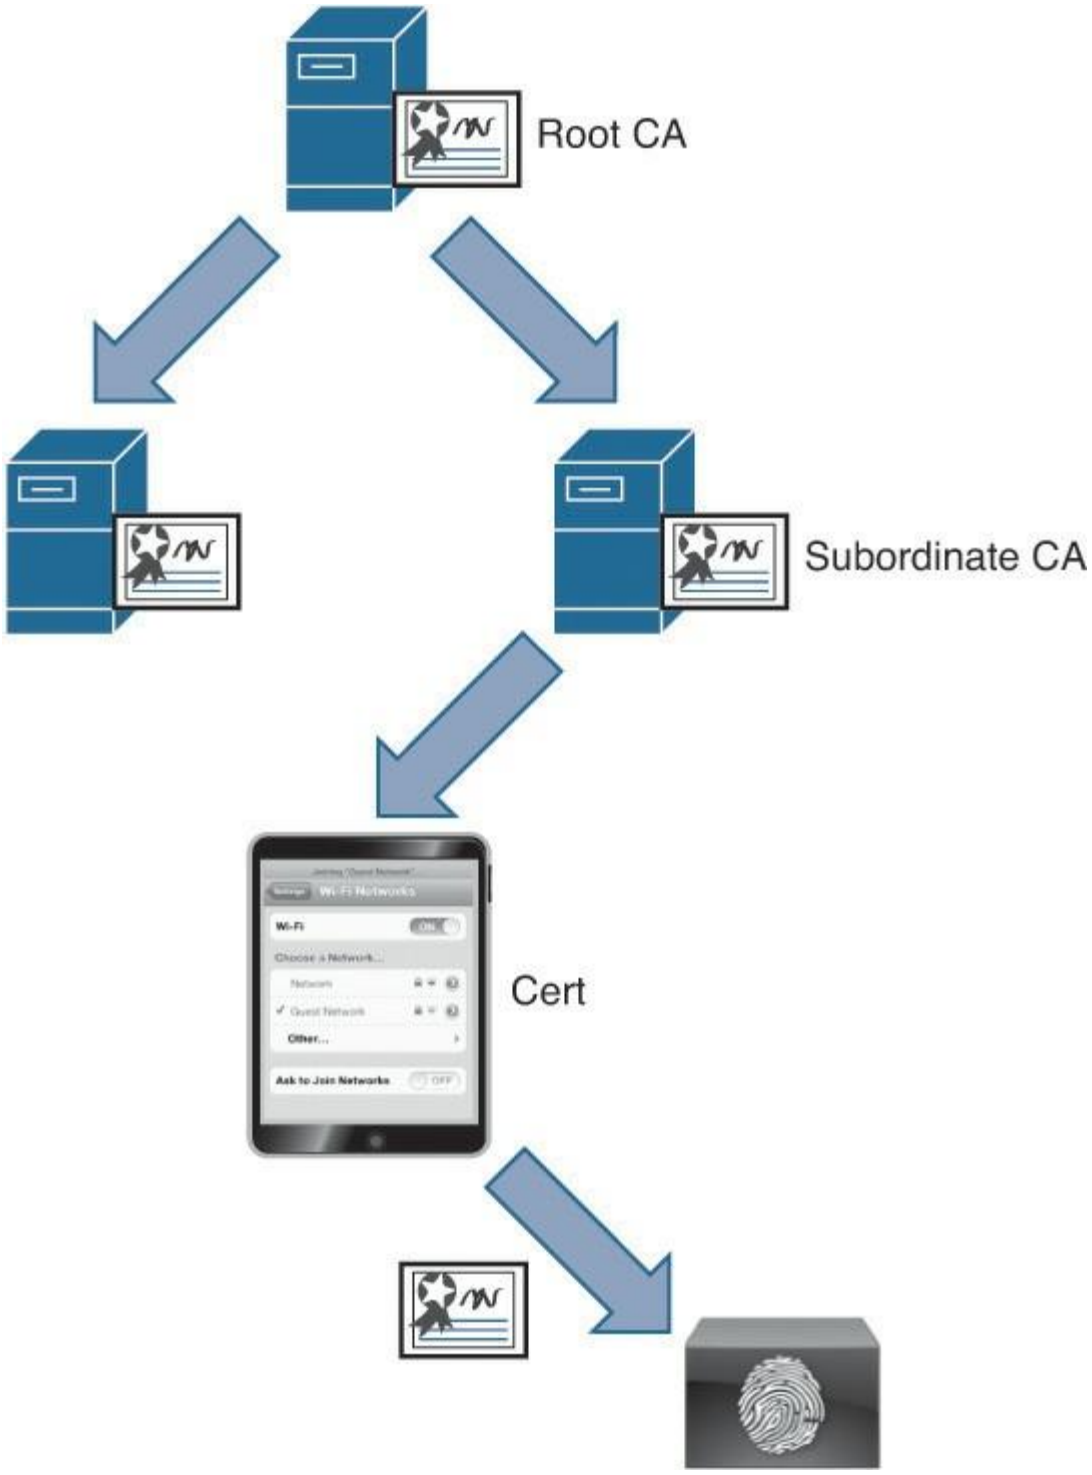


Figure 16-12 X.509 PKI hierarchy.

When trusting a certificate authority, keep in mind that each of the CAs that are in the signing hierarchy must also be trusted.

Importing 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 > Certificate Store**. This process was covered in detail within Chapter 9, "Initial Configuration of Cisco ISE."

Figure 16-13 shows a simple example of connecting to the website of your CA to begin the download

https://t.me/learningnets

process (the URL for a Windows CA is http://[fqdn-of-CA]/certsrv/).

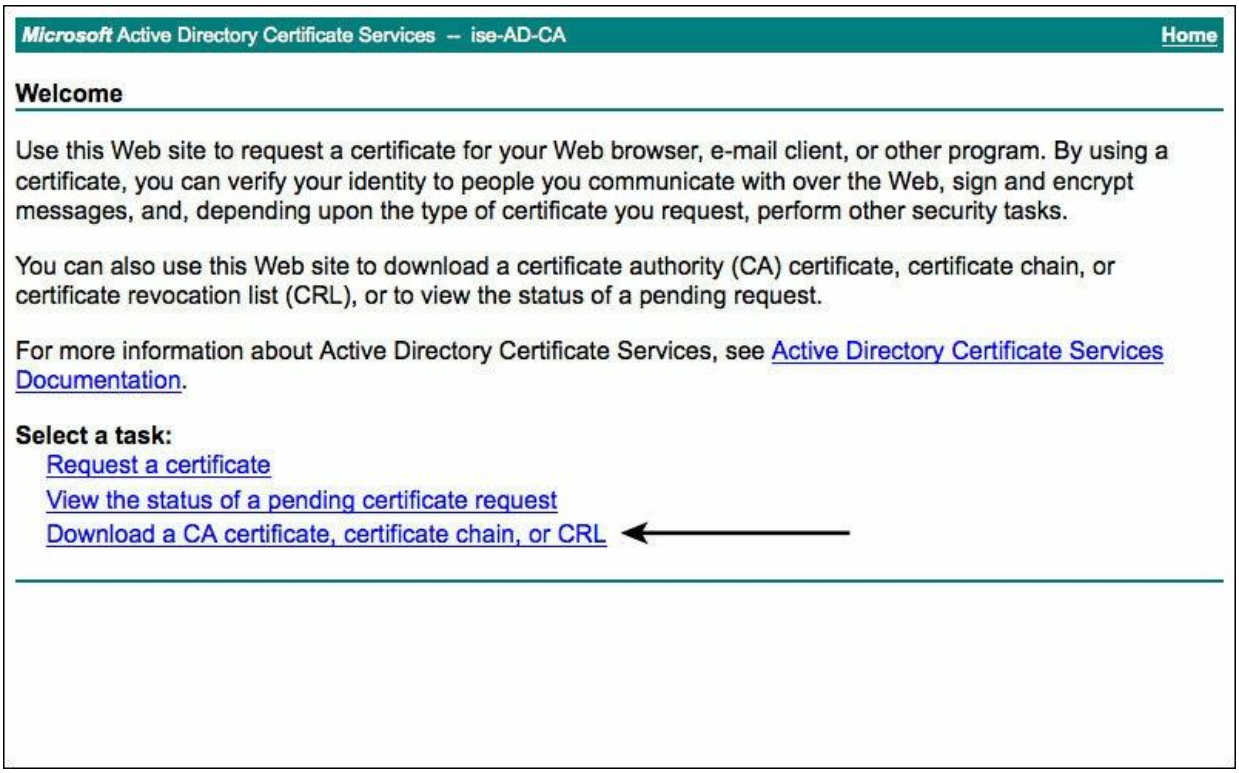


Figure 16-13 Connecting to a CA's website.

Figure 16-14 shows a simple example of downloading the public certificate from a Windows CA.

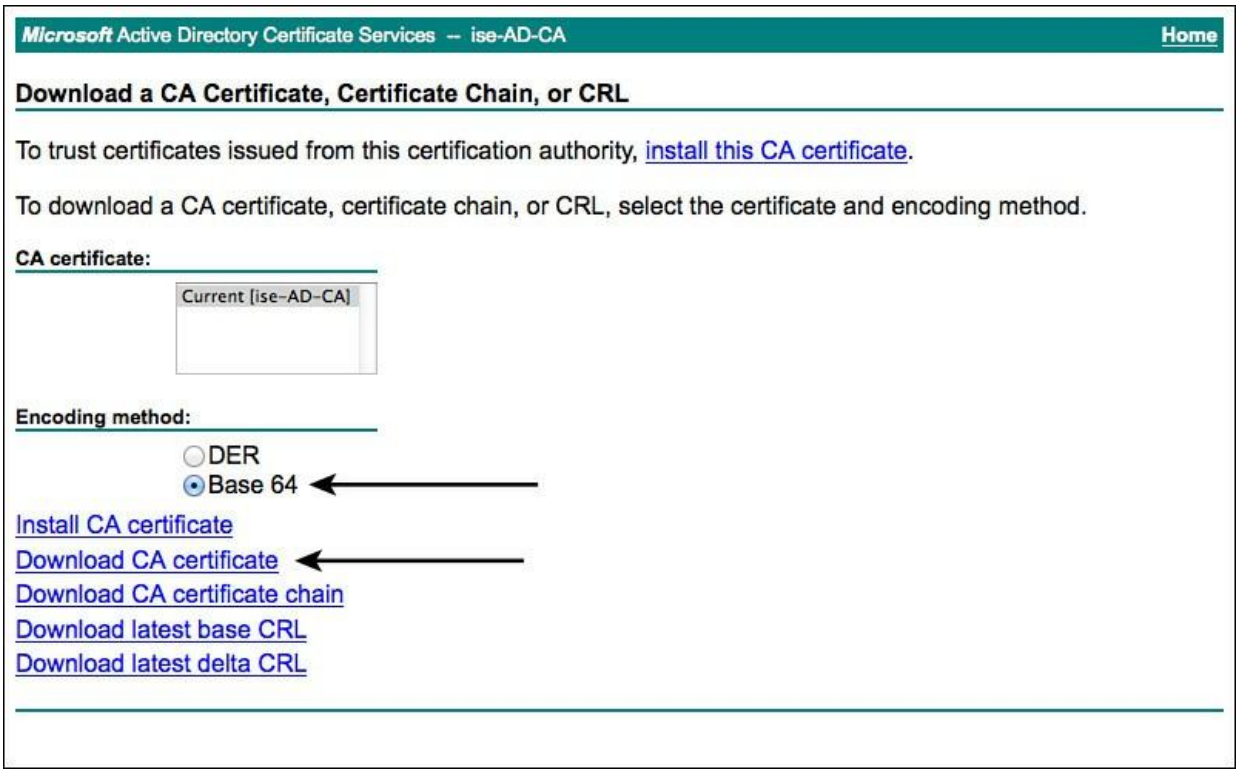


Figure 16-14 Downloading a public certificate.

When possible, avoid the use of certificate chain files (these have a .pkcs extension). Always use BASE-64–encoded files instead of DER-encoded files.

https://t.me/learningnets

When you have your root CA and all the intermediate CAs' public certificates downloaded, you can add them to ISE.

From the ISE GUI, repeat these steps for the root CA and every intermediate CA:

Step 1. Navigate to **Administration > System > Certificates > Certificate Store**.

Step 2. Click **Import**.

Step 3. Click **Browse**.

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 16-15 displays the importing of a public certificate from the root CA.

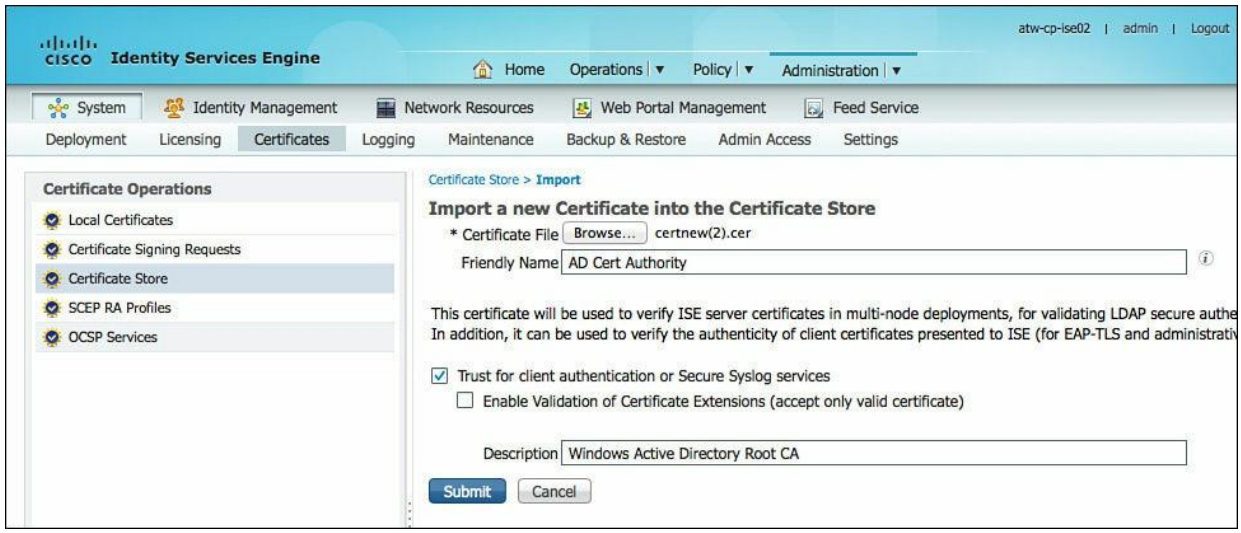


Figure 16-15 Importing a public certificate.

Configuring Certificate Status Verification (optional)

At this point, all certificates signed by that CA are trusted as a valid authentication. Commonly in production environments, it is not good enough to just trust the certificate—you must also know whether those certificates have been revoked. That requires the use of either CRLs or OCSP.

Within the certificate trust you just created for client authentication, you can specify the CRL or OCSP configuration.

From the ISE GUI, do the following:

Step 1. Navigate to **Administration > System > Certificates > Certificate Store**.

Step 2. Select the certificate you just added for client authentication.

Step 3. Scroll down to **Certificate Status Validation**.

Step 4. Select **Validate against OCSP Server**.

Step 5. Choose the correct OCSP provider from the drop-down. If you have not configured an OCSP provider yet, please see Chapter 9 for details.

Step 6. Click **Save**.

Figure 16-16 displays the importing of a public certificate from the root CA.

https://t.me/learningnets

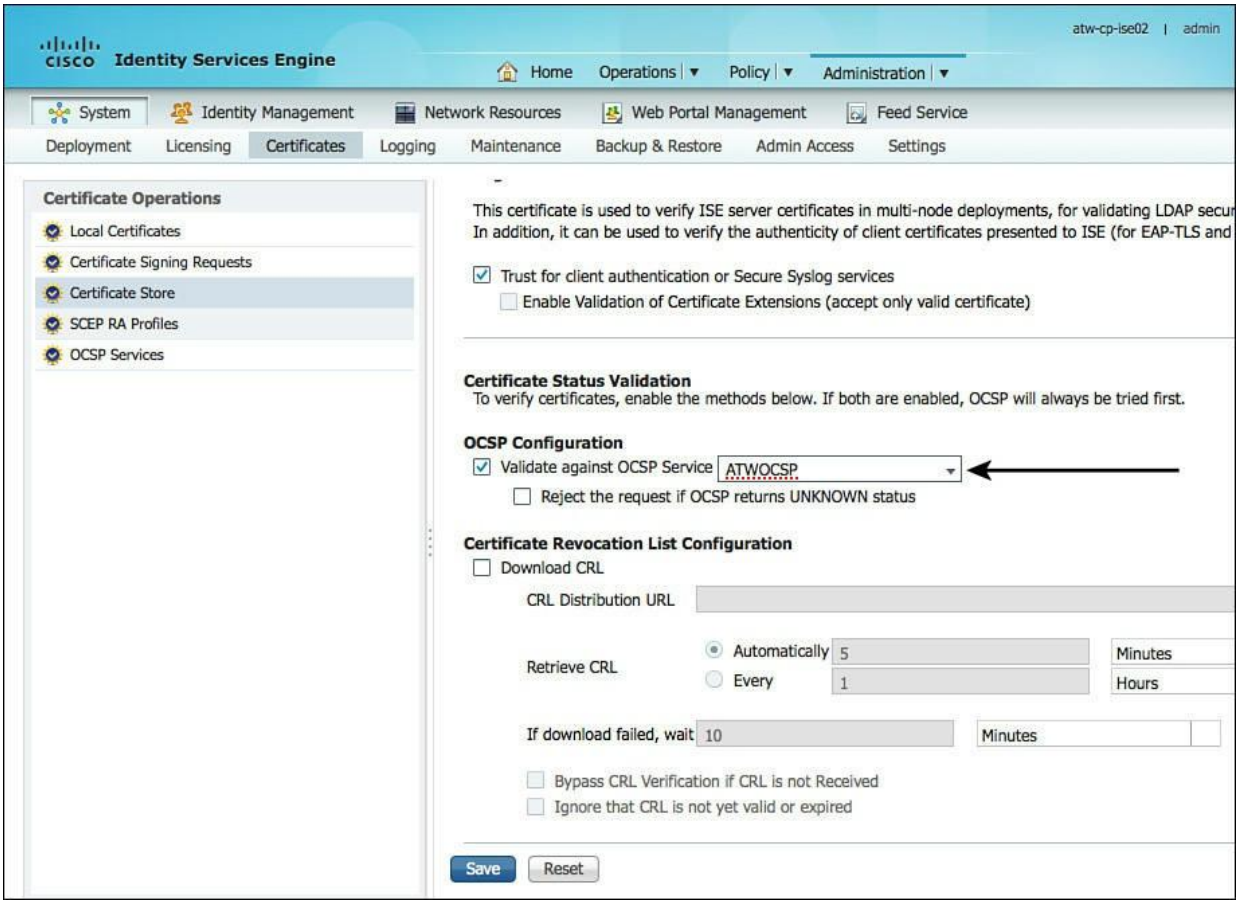


Figure 16-16 Importing a public certificate.

You are now all set to begin fully authenticating client certificates.

Verifying Certificate Authentications

To verify that a certificate authentication is working, you first need to have certificates provisioned. So, it's a bit of a chicken-and-egg issue for those of you following along in the book. However, assuming you have a client that is already configured to use a certificate for authentication, you can follow along with your own install. If you don't have clients configured yet, don't worry—you can still follow along.

Cisco ISE has a phenomenally useful tool built in to it, commonly called Live Log. Live Log provides a near real-time view of all incoming authentications, change of authorizations (CoAs), and more.

Identifying a failed authentication is rather straightforward: failed authentications are highlighted. Figure 16-17 shows Live Log with two failed authentications (the image was cropped to fit better in the print of this book). Notice in the figure that the authentication is using EAP-TLS for the protocol with a method of dot1x. If you so choose, you can filter in Live Log so that only the protocol you are looking for is displayed, which makes it easier to find your entry of interest. For example, you can enter TLS in the Authentication Protocol field of Live Log, and only those entries will be displayed.

https://t.me/learningnets

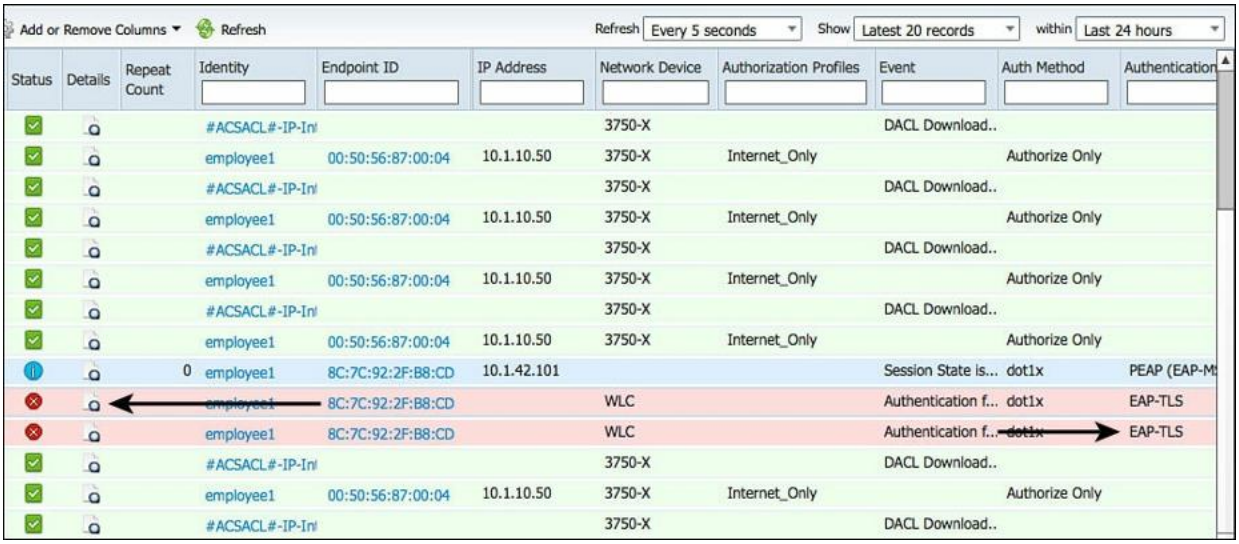


Figure 16-17 Live Log with failed authentication attempts.

To learn more about the attempted authentication, you click the Details icon. This icon looks a bit like a piece of paper with a magnifying glass over it, and it is in the Details column of Live Log. When you click that icon, it pops up another window with the Authentication Details report, which is displayed in Figure 16-18. Figure 16-18 is a cropped screen capture of the details report to make it more legible in this book. As you examine Figure 16-18, notice the text intended to point out the reason the authentication failed.

https://t.me/learningnets

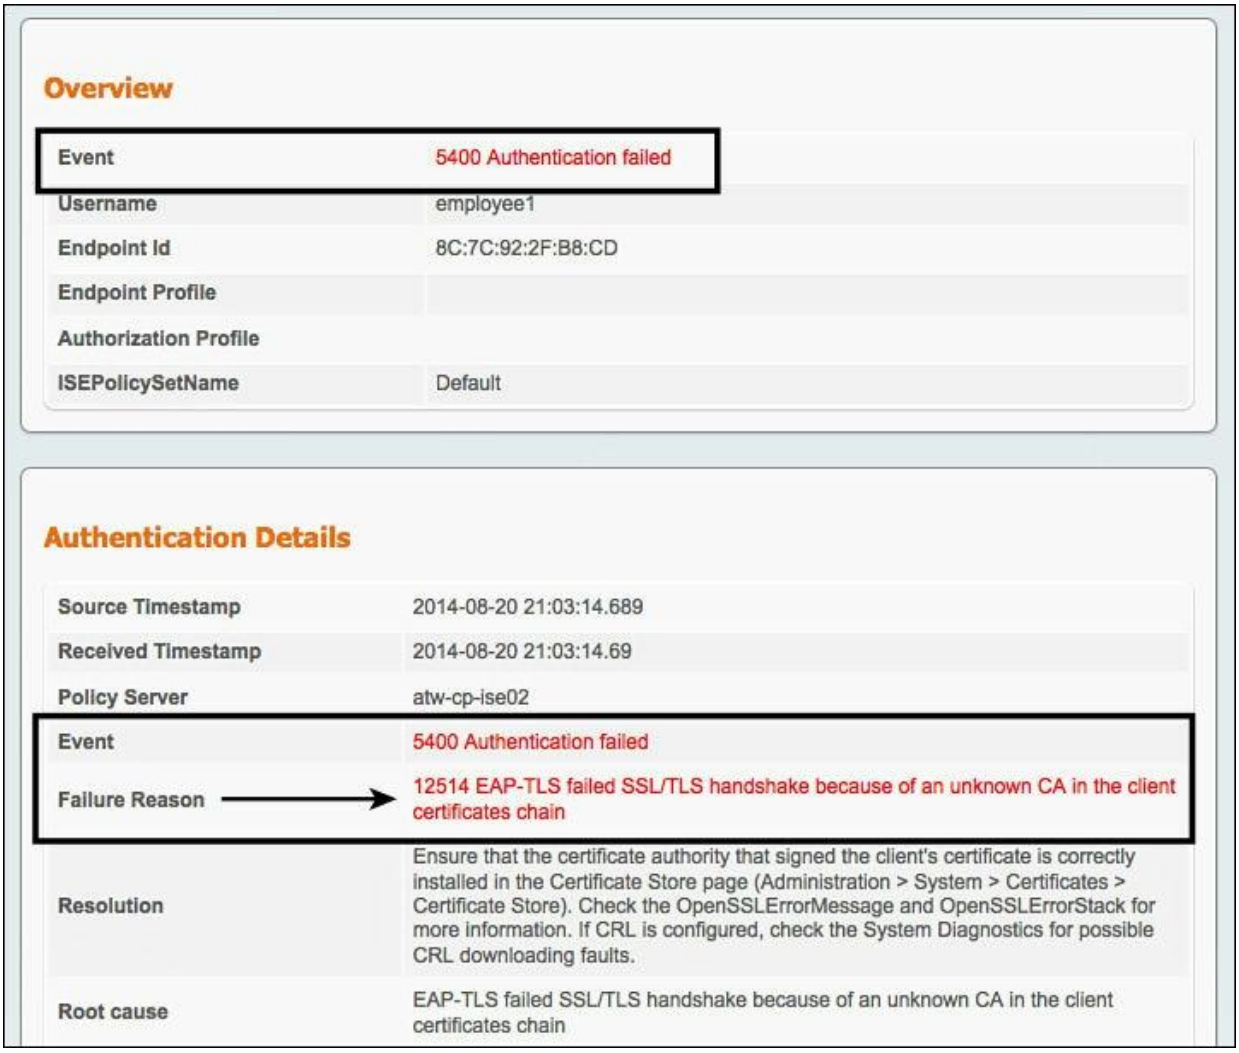


Figure 16-18 Excerpt from the authentication details report.

In fact, a field in the report named Failure Reason contains a fairly detailed message that reads: "12514 EAP-TLS failed SSL/TLS handshake because of an unknown CA in the client certificates chain." That message has just informed you that the signing CA, or one of the CAs in the chain, is not trusted by ISE. There is even a resolution field in the report that will help guide you in your journey to locate a fix for the issue.

The authentication details report is very thorough and is the top tool used to resolve the majority of TAC cases related to authentications. For an example of a full authentication details report, see Figure 16-20.

In this example, we have resolved the issue and trusted the signing CA, and now the same endpoint is authenticating successfully. Figure 16-19 shows a cropped screen capture of the Live Log with the successful EAP-TLS authentication that results in Internet-only access.

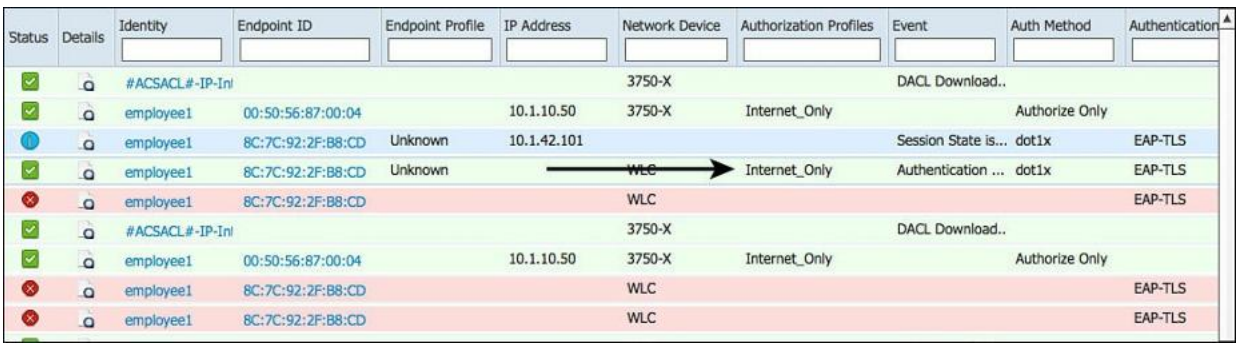


Figure 16-19 Live Log with successful authentication.

https://t.me/learningnets

Clicking the Details icon launches the authentication details report for the successful authentication. Figure 16-20 is too small for most to read, but it will give you a solid idea of the sheer amount of data that exists in the report.

https://t.me/learningnets



https://t.me/learningnets

Figure 16-20 Full authentication details report.

This completes the chapter on certificate-based authentications. The topic will be covered again in Chapter 17.

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 16-2 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
List	Minimal checking to authenticate a certificate	499
List	Revocation checking	502
Paragraph	Public/private key pairs and asymmetric encryption	505
Paragraph	X.509 PKI hierarchy	512

Table 16-2 Key Topics for Chapter 16

Define Key Terms

certificate

trusted authority

certificate authority

PKI

key-pair

public key

private key

CAP

Principle Username X.509 Attribute

https://t.me/learningnets

Chapter 17. Bring Your Own Device

This chapter covers the following exam topics:

- **Implement BYOD access**
- **Describe elements of a BYOD policy**
- **Device registration**
- **My Devices portal**
- **Describe supplicant provisioning**

Back in January 2007, Steve Jobs made a world-altering announcement when 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 multitouch screen interface, boasted a "real browser" instead of the cut-down versions that had been used on mobile devices up to that point, and arguably changed the game for the experience that users expected from their mobile devices from that point on. In January 2010, the iPad was released. In June of that same year, Aaron Woland was at Cisco Live in Las Vegas presenting a session on network access control (NAC).

In that NAC session, Mr. Woland asked the audience members a question: Which of their companies 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 the response was nearly 90% "no-way" and only 10% "yes."

At that same conference, Cisco announced the CIUS, which was designed to be a "corporate tablet"—a device to provide that wonderful multitouch user experience along with the security and assurance that IT departments required. Fast-forward 18 months, and Cisco announced the end-of-sale of the CIUS due to lack of adoption.

In June 2012, just two short years later, when Aaron asked the same question about allowing personal devices, the result was 90% affirmative and only 10% said their organizations would not allow personal devices. What a difference two years can make! Bring your own device (BYOD) has become an absolute reality.

As shown in Figure 17-1, we are moving into an era of BYOD, choose your own device (CYOD), and even into a bring your own app (BYOA) type of model.

https://t.me/learningnets

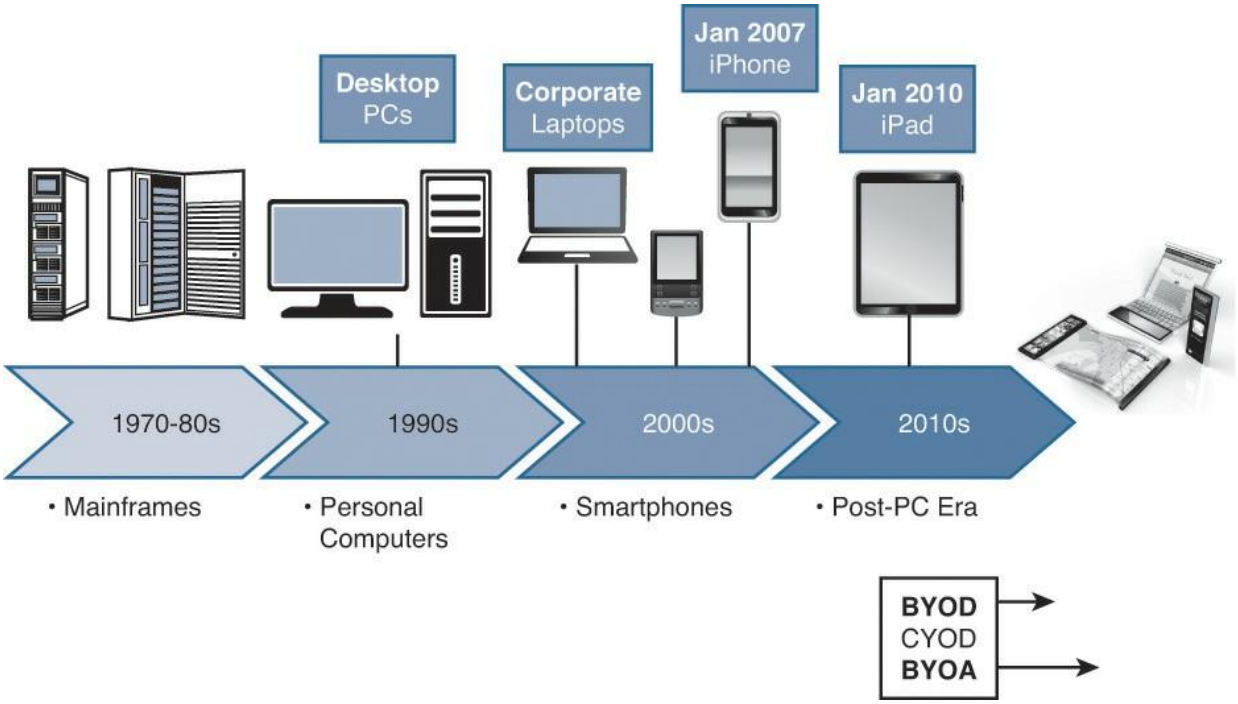


Figure 17-1 BYOD timeline.

Employees are demanding the use of the devices that make them most productive, with native applications running on those platforms that provide the user experience they are used to and love. This introduces a whole new paradigm for security, requiring the identification of the user, the device, the location of the user, and so much more.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Implement BYOD access	2-6
Describe elements of a BYOD policy	1-2, 4-5, 9-10
Device registration	1, 5
My Devices portal	7
Describe supplicant provisioning	1-6, 8-9

Table 17-1 "Do I Know This Already?" Section-to-Question Mapping

> **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.

https://t.me/learningnets

1. What is the process of onboarding as it relates to BYOD?
   a. It's a form of torture used in military interrogations.
   b. It prepares an endpoint for network access with supplicant configuration, and possibly even certificate provisioning.
   c. It's the process in which an IT department will prestage an endpoint for corporate use before issuing the endpoint to the end user.
   d. It prepares an endpoint for network access by preconfiguring an installation package that the end user runs with administrator privilege to configure the endpoint.
2. With a single-SSID model for BYOD onboarding, how does the supplicant begin using its new certificate-based credentials?
   a. The endpoint will continue to use the initial credentials until the next reauthentication interval.
   b. ISE will send a CoA-DM, causing a new authentication.
   c. ISE will send a CoA-Reauth, causing a new authentication.
   d. The endpoint will continue to use the initial credentials until the endpoint is deassociated 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. It's a configurable option, so nothing prevents guests from receiving the certificate and supplicant profile.
   c. It's a configurable option based on the authorization result given to the user.
   d. It's a configurable option in the client provisioning policies to permit guests to enter the provisioning flow.
4. The same ACL can be used for all endpoints to be onboarded. However, the security of the ACL needs to be relaxed for Androids. What is that reason?
   a. Google just feels that it is so special, so Androids require special access to keep up.
   b. Androids require access to the local app store in ISE.
   c. Because Android is inherently an insecure operating system, it therefore needs a less secure ACL.
   d. Androids require access to their app store to download and execute Cisco's Network Setup Assistant APP.
5. What are an ISE admin's options for dealing with endpoints that are not supported by the BYOD onboarding process?
   a. Cisco ISE will reject an authentication from any endpoint that cannot go through the onboarding process.
   b. The admin 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 will automatically permit access to any device that can't be onboarded.
   d. After the BYOD onboarding flow is enabled, every device must be onboarded. There are custom templates to be able to push profiles to any device that is not natively supported.
6. From where does an iOS-based device download the iOS Network Setup Assistant?
   a. From the Apple App Store.
   b. iOS uses the native OTA functionality.
   c. From ISE directly.

https://t.me/learningnets

d. From the Cisco App Store.

7. True or False? The ISE admin may log in to the MyDevices portal to manage all the registered devices.

   a. True

   b. False

8. Which of following lists most accurately describes the portions of BYOD onboarding that can be verified within Live Log?

   a. An entry will exist for the initial authentication, CoA, and final authentication.

   b. An entry will exist for the initial authentication, successful launch of the NSA app, and the final authentication.

   c. An entry will exist for the initial authentication, successful endpoint registration, download of the NSA app, and the final authentication.

   d. An entry will exist for the initial authentication, successful endpoint registration, CoA, and the final authentication.

9. As it relates to ISE 1.2, from where do Windows and Mac OSX endpoints download their Network Setup Assistant applications?

   a. Windows downloads the NSA app from the Microsoft App Store. Mac OSX uses the native OTA.

   b. Neither Windows nor Mac use NSA; they use native capabilities instead.

   c. Windows uses native capabilities, but the Mac will use a Java applet downloaded from the CPP.

   d. Windows and Mac will use a Java applet that is downloaded from the CPP hosted on ISE.

10. At which one of the following locations does an ISE admin determine which NSP 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, IT does not possess the ability to create and enforce a rigorous secure access policy. Although the user might be authorized, the device, location, time, and access medium can pose a company policy violation or even regulatory compliance violation that cannot be adequately detected or enforced.

In this chapter, you are focusing on the technical challenges of providing a secure BYOD access model. One of the most common challenges is what the industry refers to as onboarding. A user buys his new tablet or device and decides he wants to connect it to the corporate Wi-Fi and be productive on that consumer device. IT has the challenge of identifying the device as a noncorporate device and providing a limited set of access to the device.

Figure 17-2 illustrates what many companies used ISE to do, originally (back in the 1.0 days).

https://t.me/learningnets

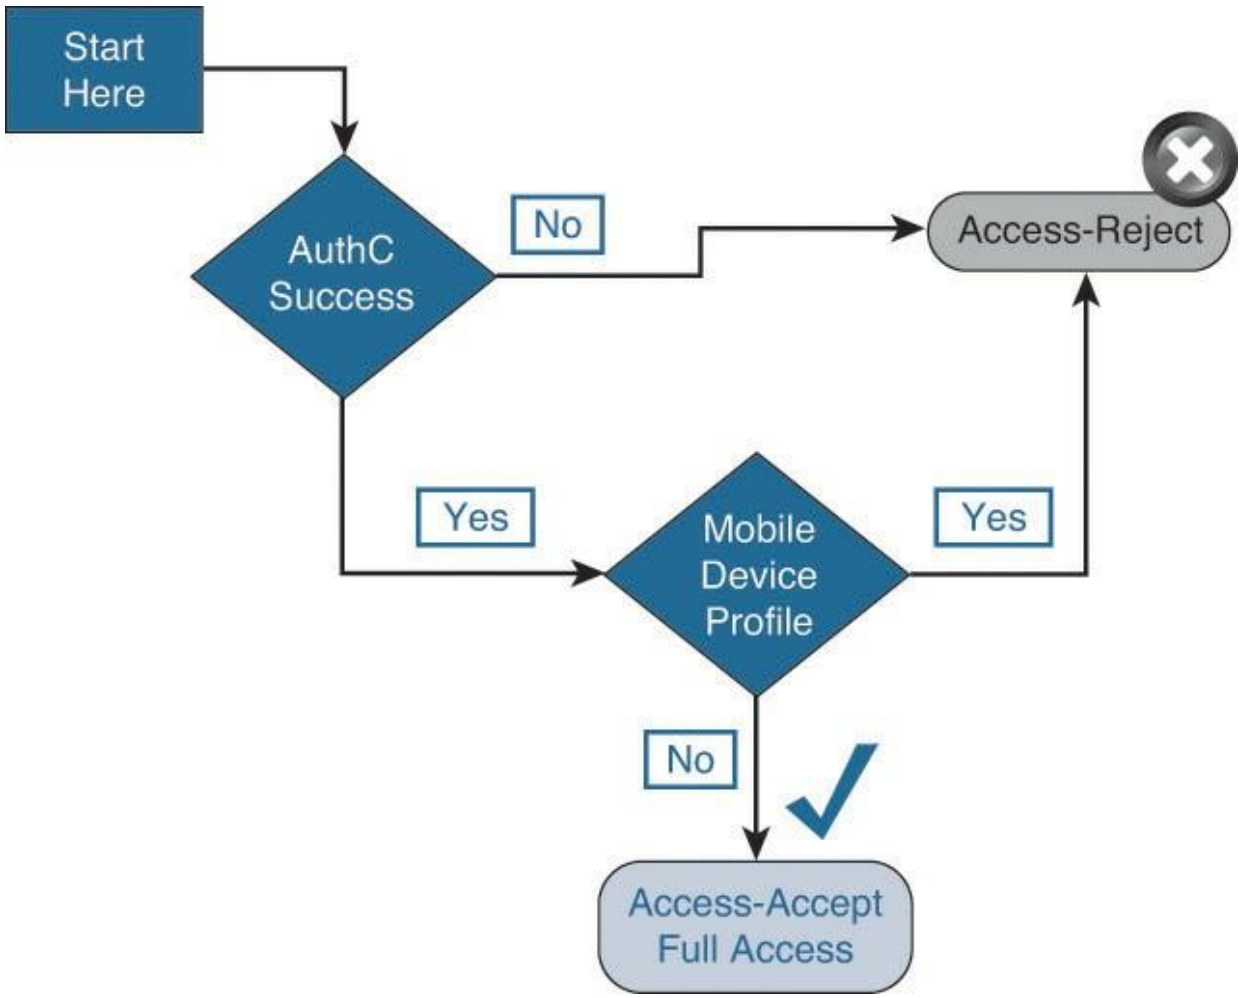


Figure 17-2 Old-style policy.

Then mobile device management (MDM) solutions came into play. These MDM solutions were able to manage these mobile platforms to some extent. MDM policies would ensure that devices had security enabled, such as encryption, remote wipe capabilities, screen lock (pin lock), and so forth. The MDM would even provision certificates down to the device as well as configurations for the device's supplicant to preconfigure the device to have network access. Then ISE would provide the correct level of access for each device based on the certificate the device had.

Figure 17-3 illustrates this type of policy logic.

https://t.me/learningnets

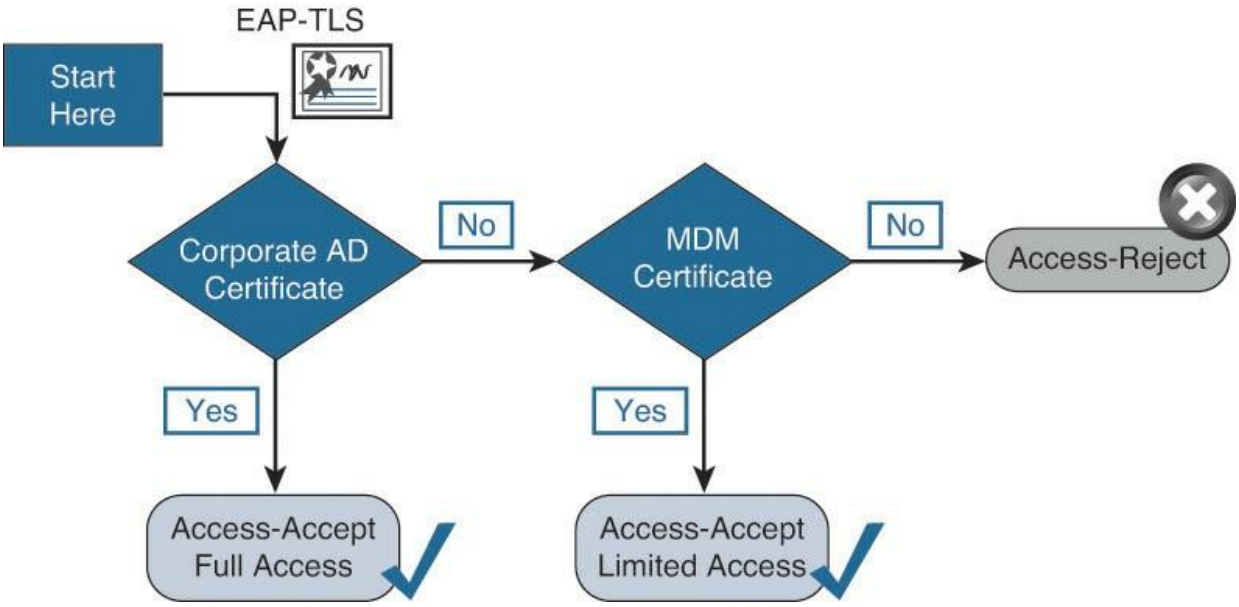


Figure 17-3 Using certificates to differentiate access.

MDMs typically cost money per device, and many companies were only looking for a good way to provision certificates and configure the device supplicant to use that certificate with 802.1X. The MDM cost was often prohibitive just to provision the certificate and get the device on the network. Cisco customers were looking for a much easier and cheaper way to accomplish the onboarding aspect of network access.

Onboarding Process

There are two types of onboarding that we will focus on. The first is what Cisco calls BYOD onboarding, which includes registering the device with ISE, provisioning the certificate to the device, and configuring the device's supplicant. This is using the native supplicant within the operating system, not installing a new one.

The second is MDM onboarding, which is the process of registering the device with the Mobile Device Manager, installing the MDM client software, and enforcing the security policy on that device. The key to successful onboarding within a company is to make it self-service and not require involvement of IT.

BYOD Onboarding

As introduced in Chapter 14, "Deploying Guest Services," ISE provides a My Devices portal that allows users to register devices and manage those devices they have registered. It's possible for a device to simply be registered, which can provide one level of authorization, such as Internet-only access. It's also possible for the 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, there can be a single-SSID or dual-SSID approach to the onboarding, plus a wired option (of course).

Figure 17-4 depicts the dual-SSID approach, while Figure 17-5 depicts the single-SSID approach. A quick comparison of the approaches follows.

https://t.me/learningnets

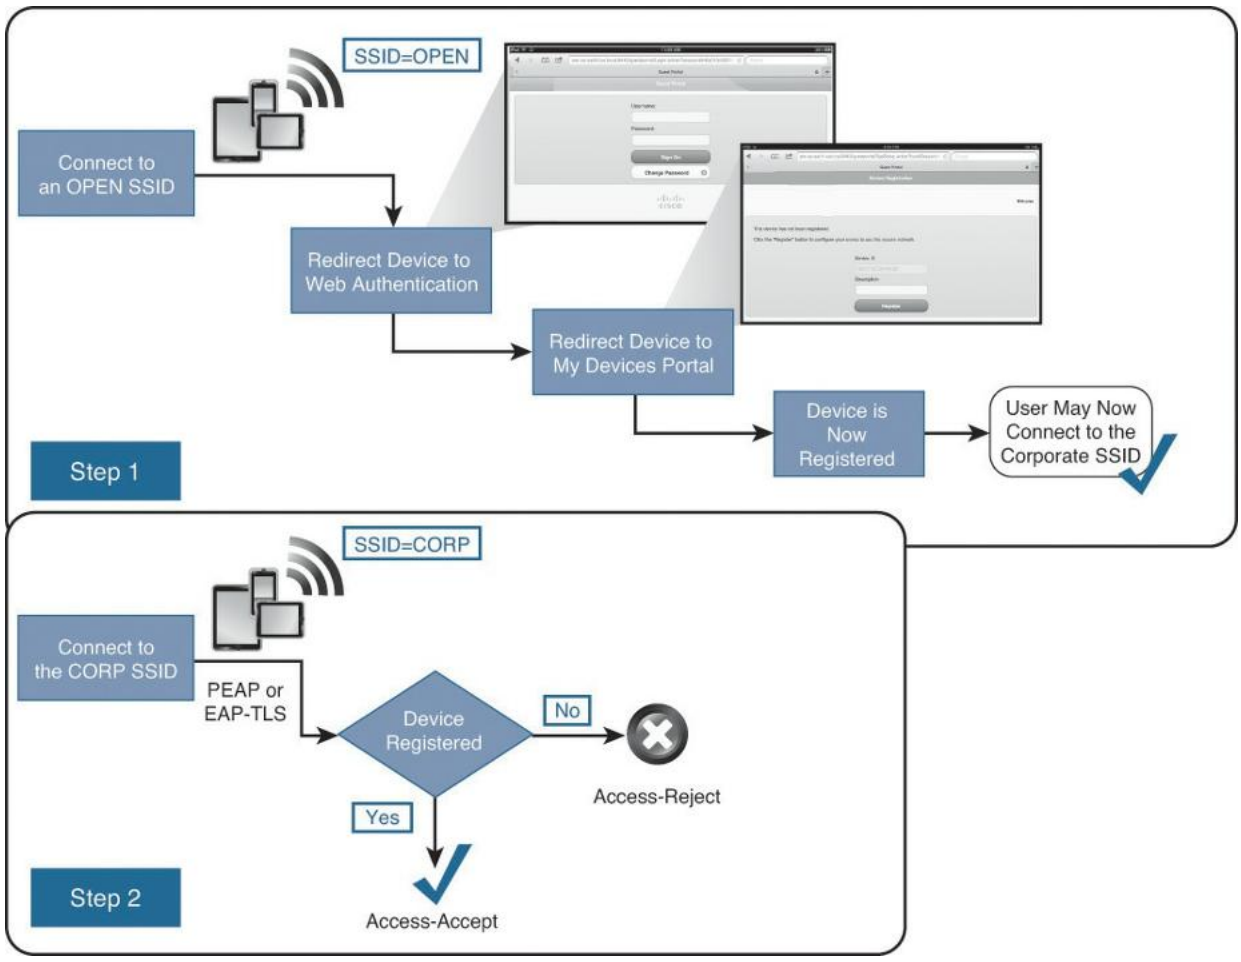


Figure 17-4 Dual-SSID flow.

https://t.me/learningnets

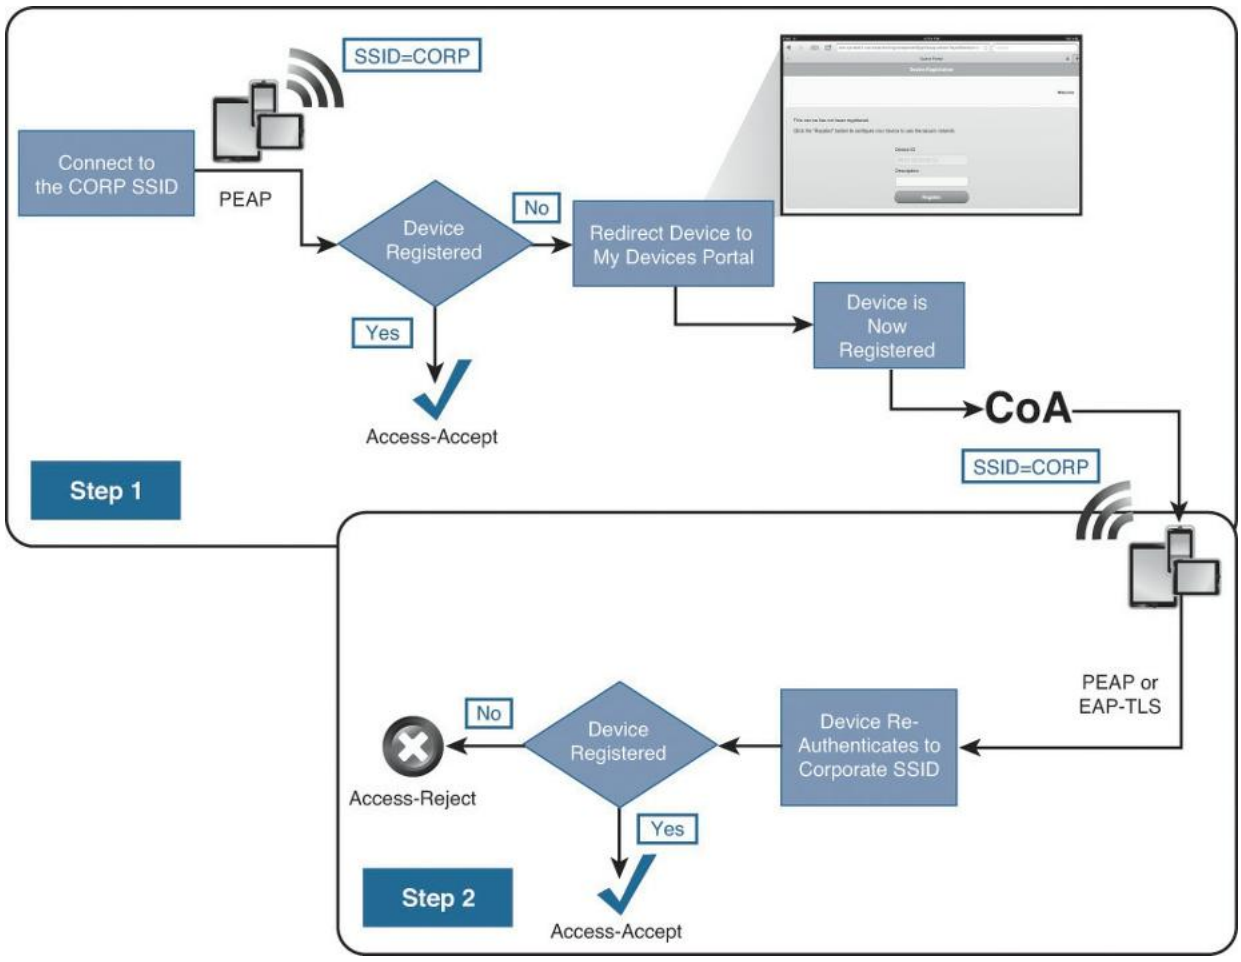


Figure 17-5 Single-SSID flow.

Dual SSID

With the dual-SSID approach to onboarding:

- The employee does not need to configure the supplicant on the device.
- The employee authenticates to a web form.
- The employee connects to the Open SSID before the provisioning process, and the employee must connect to the corporate SSID after the process (may be automated or manual).

Figure 17-4 depicts the dual-SSID approach.

Single SSID

With the single-SSID approach to onboarding:

- 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 17-5 depicts the single-SSID approach.

Configuring NADs for Onboarding

In this section you are configuring the NADs for both dual- and single-SSID onboarding, beginning with the dual-SSID flow.

https://t.me/learningnets

Configuring the WLC for Dual-SSID Onboarding

Dual-SSID onboarding uses an open WLAN configured for NAC RADIUS, and the security settings will only use Mac Filtering (also called Wireless MAB). This network should have been created already in Chapter 12, "Implement Wired and Wireless Authentication," and reviewed again in Chapter 13, "Web Authentication," but you will verify the WLC settings briefly in this section, to be certain.

Reviewing the WLAN Configuration

From the WLC GUI, edit the Open WLAN (named CP-Guest with an SSID of CiscoPress-Guest, in this example) and ensure that the RADIUS server definitions are configured for RFC 3576 (CoA).

The General tab of the WLAN should provide an SSID and profile name. CiscoPress-Guest is used in this example, as shown in Figure 17-6.



Figure 17-6 Open WLAN General tab.

Under **Security > Layer 2**, Layer 2 security should be set to **None**. The **Mac Filtering** check box should be enabled, as shown in Figure 17-7.

https://t.me/learningnets

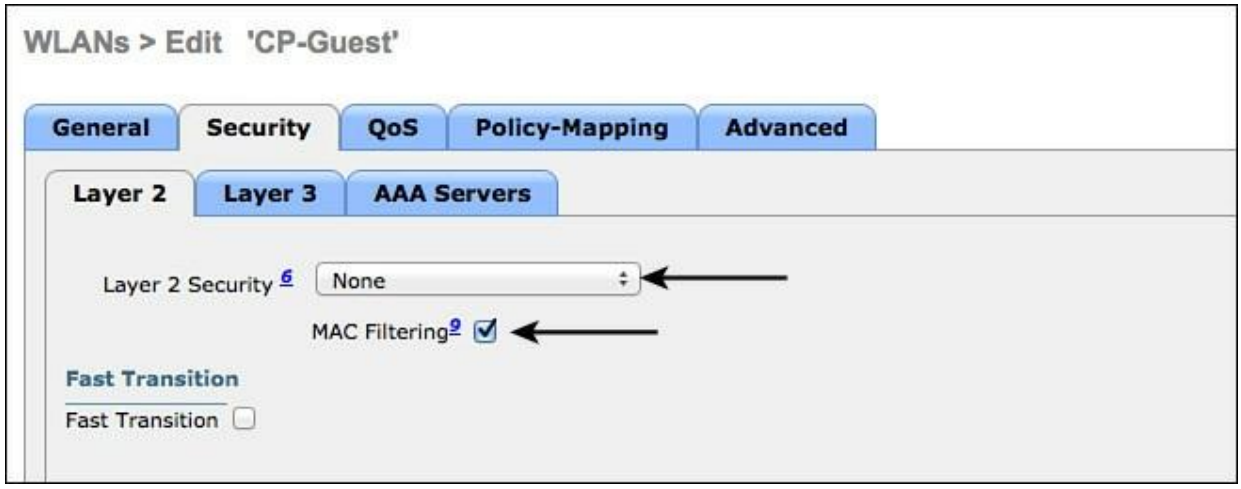


Figure 17-7 Layer 2 Security tab.

Under **Security > Layer 3**, Layer 3 security should be set to **None**, as shown in Figure 17-8.

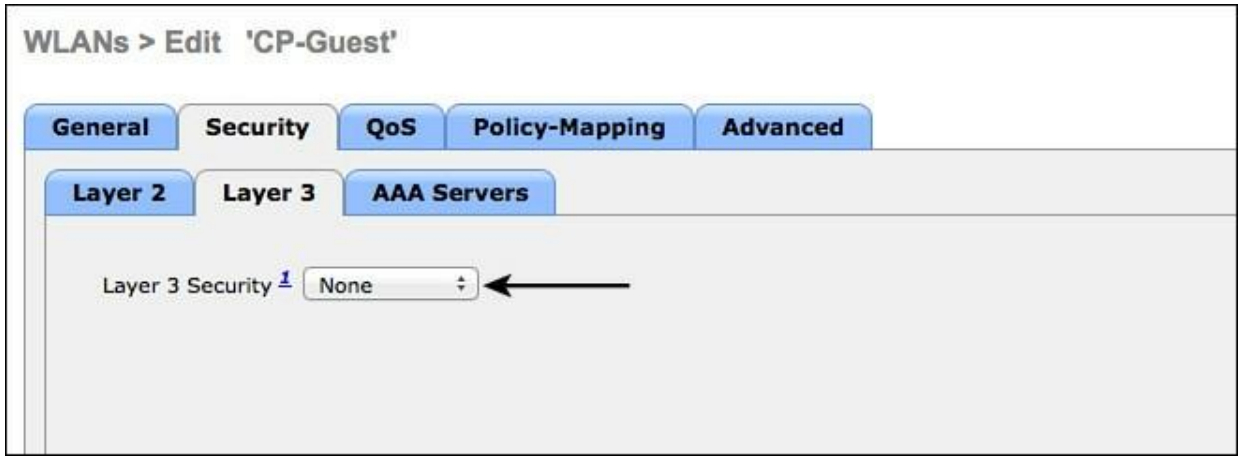


Figure 17-8 Layer 3 Security tab.

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 Service Node(s) must be selected for authentication and accounting servers, as shown in Figure 17-9.

https://t.me/learningnets

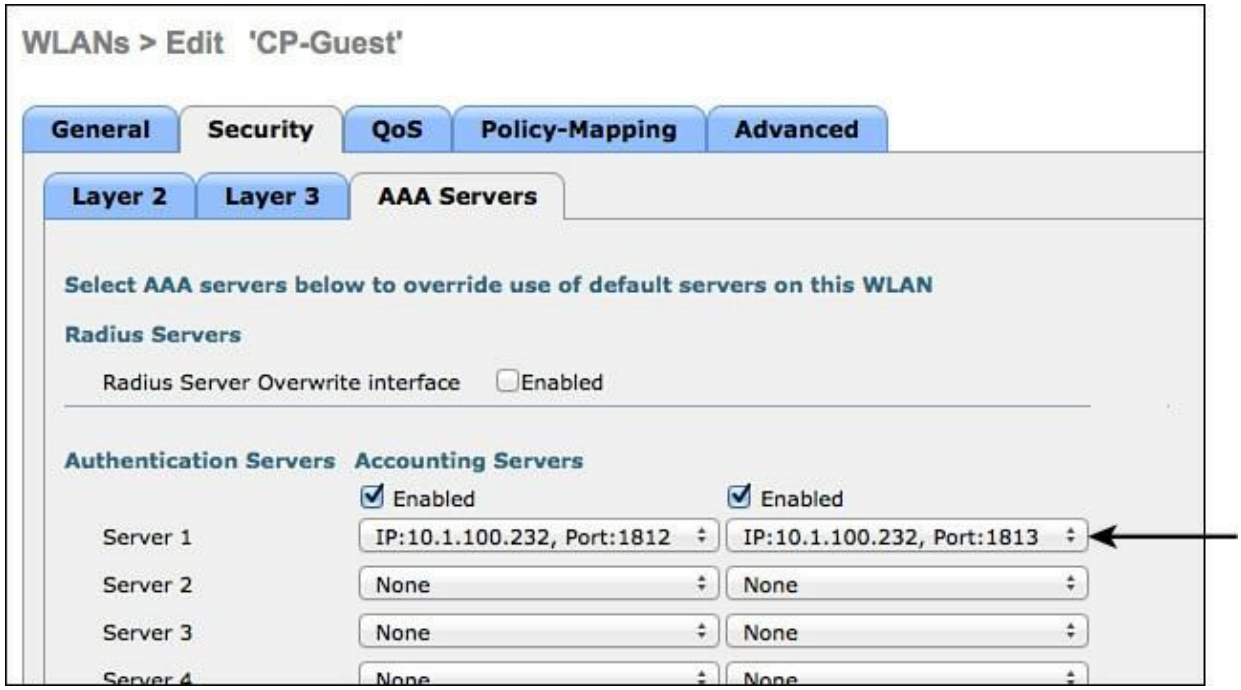


Figure 17-9 AAA Servers Security tab.

For the **Advanced** tab, the following should be configured: **Allow AAA Override** must be enabled and NAC State should be configured for **Radius NAC**, as shown in Figure 17-10.

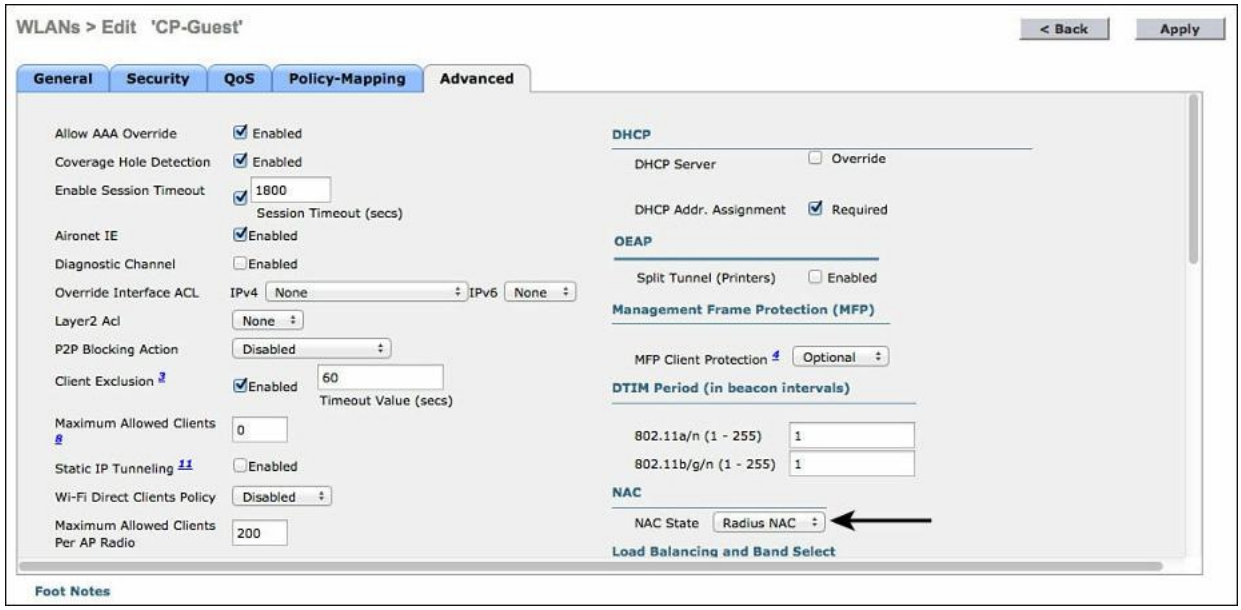


Figure 17-10 Advanced tabs.

You should also double-check that each RADIUS Server Definition under **Security > AAA > RADIUS > Authentication** is configured to allow CoA, which is the RFC 3576 drop-down shown in Figure 17-11.

https://t.me/learningnets

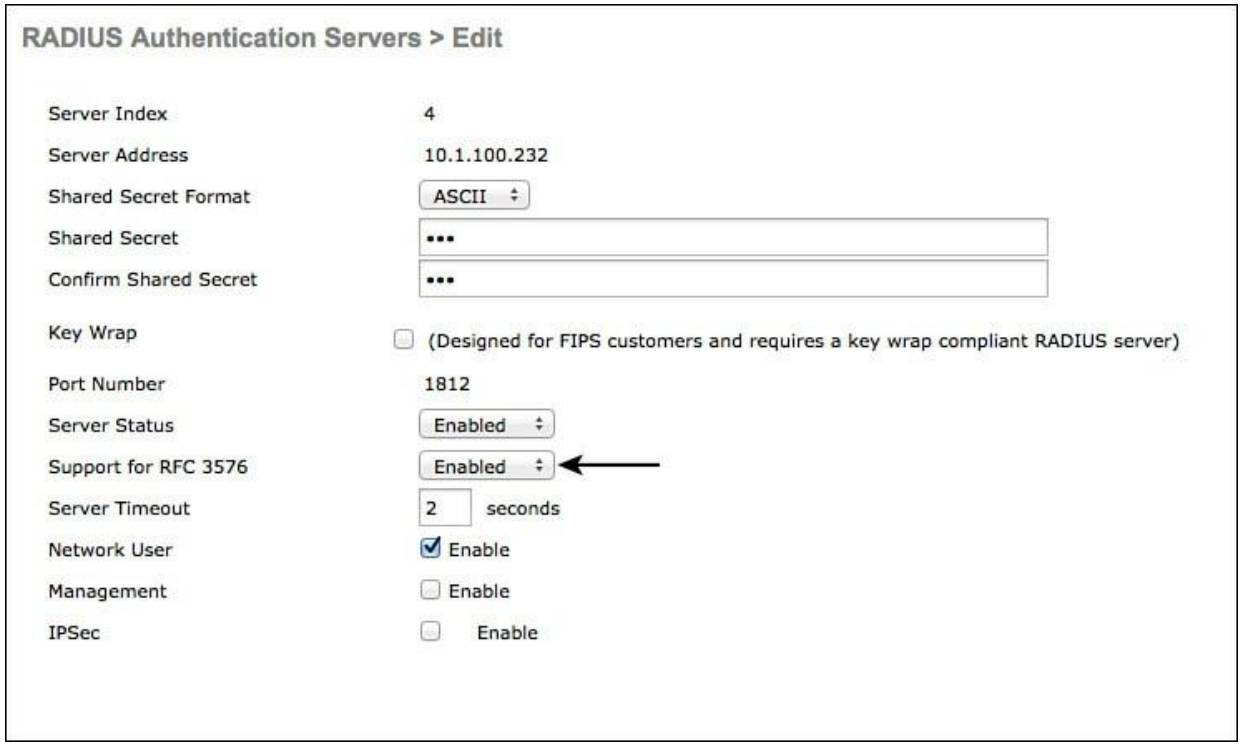


Figure 17-11 RADIUS Authentication Servers tab.

Verifying the Required ACLs

You should have an access control list (ACL) on the switches and the wireless controllers already named ACL-WEBAUTH-REDIRECT that permit DHCP, DNS, and traffic to ISE and denies most other traffic. We added that in Chapter 14.

When onboarding with iOS, Windows, and Mac OSX, the endpoint need only communicate with ISE. iOS uses its native Over the Air (OTA) provisioning process. Windows and Mac both use a Java-based wizard that is downloaded from ISE through the devices browser. Because the communication is limited to ISE only, the ACL-WEBAUTH-REDIRECT ACL is sufficient to be repurposed for the onboarding ACL as well.



However, Android is a different story altogether. Android devices inherently do not trust apps being installed from an app store other than those trusted during the factory install; therefore, ISE would not be allowed to host the app for Android devices by default. To keep the process simple for the end user, we will have to open up the ACL to allow access to a range of addresses for Google Play.

The Google Play app store (play.google.com) is a cloud service, and the addresses it uses can change regularly. This presents a bit of a challenge for us to permit access to those ranges. Depending on the version of Cisco WLC you are implementing, the solution might be to use a DNS-snooping ACL (7.6 and newer). For the WLC versions older than 7.6, the solution is to permit a series of blocks of addresses that are known to be used by the Android Marketplace:

- 74.125.0.0/16
- 173.194.0.0/16
- 173.227.0.0/16
- 206.111.0.0/16
- 8.35.0.0/16

These ACLs will be used for both single- and dual-SSID onboarding.

https://t.me/learningnets

> **Note**
>
> It is quite possible that these address blocks will change, or that additional blocks may need to be added. Google is very dynamic in its content location(s).

Creating the Android ACL on the Wireless LAN Controller

From the WLC GUI, under **Security > Access Control Lists**, add a new ACL named Android-Marketplace; then configure the ACL as the one in Figure 17-12.

Access Control Lists > Edit

General

Access List Name: Android-Marketplace

Deny Counters: 0

Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port	DSCP	Direction	Number
1	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Outbound	2062
2	Permit	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	UDP	Any	DNS	Any	Inbound	153
3	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
4	Permit	0.0.0.0 / 0.0.0.0	10.1.100.0 / 255.255.255.0	Any	Any	Any	Any	Inbound	1819
5	Permit	0.0.0.0 / 0.0.0.0	10.1.103.0 / 255.255.255.0	Any	Any	Any	Any	Inbound	0
6	Permit	0.0.0.0 / 0.0.0.0	74.125.0.0 / 255.255.0.0	TCP	Any	Any	Any	Inbound	28
7	Permit	0.0.0.0 / 0.0.0.0	173.194.0.0 / 255.255.0.0	TCP	Any	Any	Any	Inbound	54
8	Permit	0.0.0.0 / 0.0.0.0	173.227.0.0 / 255.255.0.0	TCP	Any	Any	Any	Inbound	0
9	Permit	0.0.0.0 / 0.0.0.0	206.111.0.0 / 255.255.0.0	TCP	Any	Any	Any	Inbound	0
10	Deny	0.0.0.0 / 0.0.0.0	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Inbound	3685

Figure 17-12 Android Marketplace ACL.

Our sample ACL is permitting all traffic from the inside network to speak to the client and is allowing the client to communicate into the network for DNS and DHCP, as well as TCP traffic destined to the ISE servers. Next will be the lines that permit TCP traffic to the Android Market IP address ranges.

Adding the DNS Values to the Android ACL on the WLAN Controller

The DNS ACLs added to the WLC in version 7.6 work in an interesting way. The wireless access point (WAP) itself performs DNS snooping to see the response that is sent to the endpoint. The WAP does not have to examine all DNS requests and responses; instead, it is configured to consider only certain domain names as "interesting."



The WAP updates the controller to allow the specific IP address through the ACL. Figure 17-13 illustrates

https://t.me/learningnets

the DNS snooping functionality.

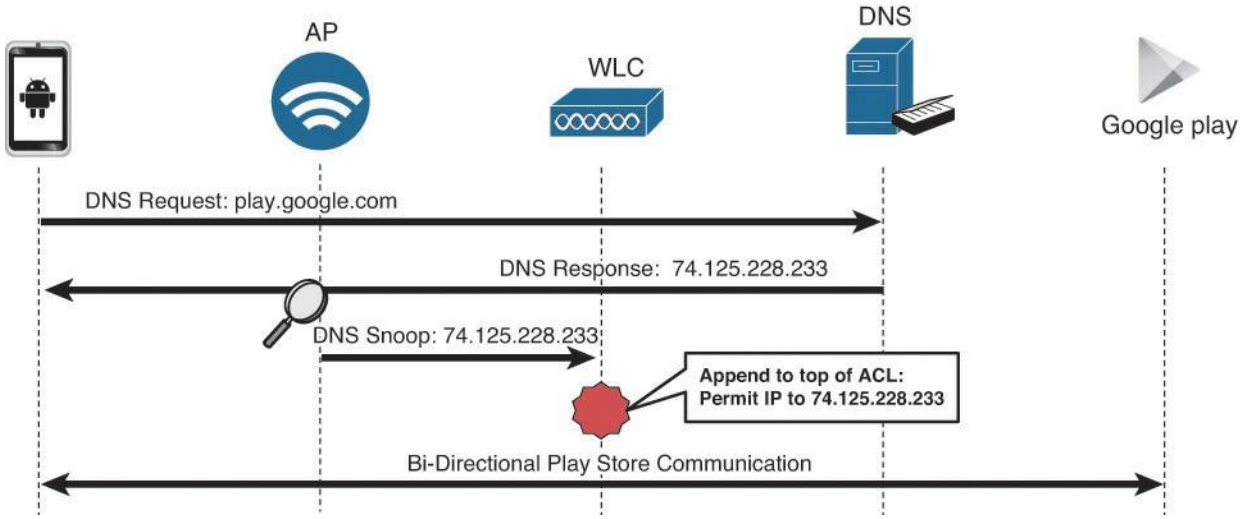


Figure 17-13 DNS snooping ACL.

Click the blue down arrow next to the ACL, and select **Add-Remove URL**, as shown in Figure 17-14.



Figure 17-14 Add-Remove URL.

Adding the URL strings to this automatically puts a wildcard (*) in front of the domain. So google.com will actually match *.google.com. Entering google.* is a way to ensure the DNS snooping will match for all country codes (that is, *.google.co.uk, and so on). The two DNS names you should have for the play store (at a minimum) are google.com and android.clients.google.com. If you are outside the United States, use an asterisk (*) instead of ".com".

Figure 17-15 shows the URL strings for the ACL.

https://t.me/learningnets



Figure 17-15 Adding URL strings to the ACL.

Creating the ACL on the IOS-based NADs

From Global-Configuration mode, add a new extended ACL named Android-Marketplace and configure the ACL as the one in Example 17-1. This ACL allows DNS to bypass redirection, along with all four of the known common Play Store address ranges; all other web traffic will be redirected to ISE.

Example 17-1 Android-Marketplace ACL for Switches

Click here to view code image

```
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 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

With the NADs prepared for the onboarding process, it's time to build the logic within the ISE Authorization Policy for both the dual- and single-SSID onboarding models.



The easier one to set up and understand first 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 an MSChapV2 inner method), then we know the device must not have been onboarded yet.

This way, when an employee shows up to work with her new mobile device and decides to try to connect to the corporate Wi-Fi, it prompts her for a username and password. The employee will enter her Active Directory credentials (as would be expected), and when she opens the browser on her mobile device, she will be redirected to the MyDevices portal where she can begin the onboarding process. ISE is using the single-sign-on credentials from the MSChapV2 authentication and does not need to prompt the end user for her credentials again during the onboarding. The process is simple, quick, and intuitive to most end users in the today's world.



https://t.me/learningnets

Dual-SSID onboarding takes a bit of a different approach. 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 a Centralized Web Authentication (CWA) portal, where she will enter her credentials. If the end user is a Guest, she will not be permitted to enter the provisioning process. This is hard-coded in ISE.

However, if the end user is not a Guest user, she will be prompted to register the device and begin the onboarding process. ISE will use the credentials entered into the CWA portal for the entire process, so the end user does not have to authenticate again.

The End User Experience

To fully understand the configuration of ISE, it is best that you experience the end user experience for both single- and dual-SSID onboarding. That will aid you in your understanding of each policy that must be created and each choice you will have to make. To demonstrate multiple user experiences, the following examples use Apple iOS for one and Android for the other. However, each onboarding method could be used with any of the supported clients (iOS, Android, Mac OSX, and Windows).

Single-SSID with Apple iOS Example

For this example, you will see an Apple iOS device (an iPad) connecting to a secured network with the SSID of "CiscoPress-Corp." This does not match the SSID 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. Instead, just try to see and understand the demonstrated end user experience and why it is happening.

Step 1. The user comes in with his iOS device. He opens Settings and connects to the corporate Wi-Fi, as shown in Figure 17-16.

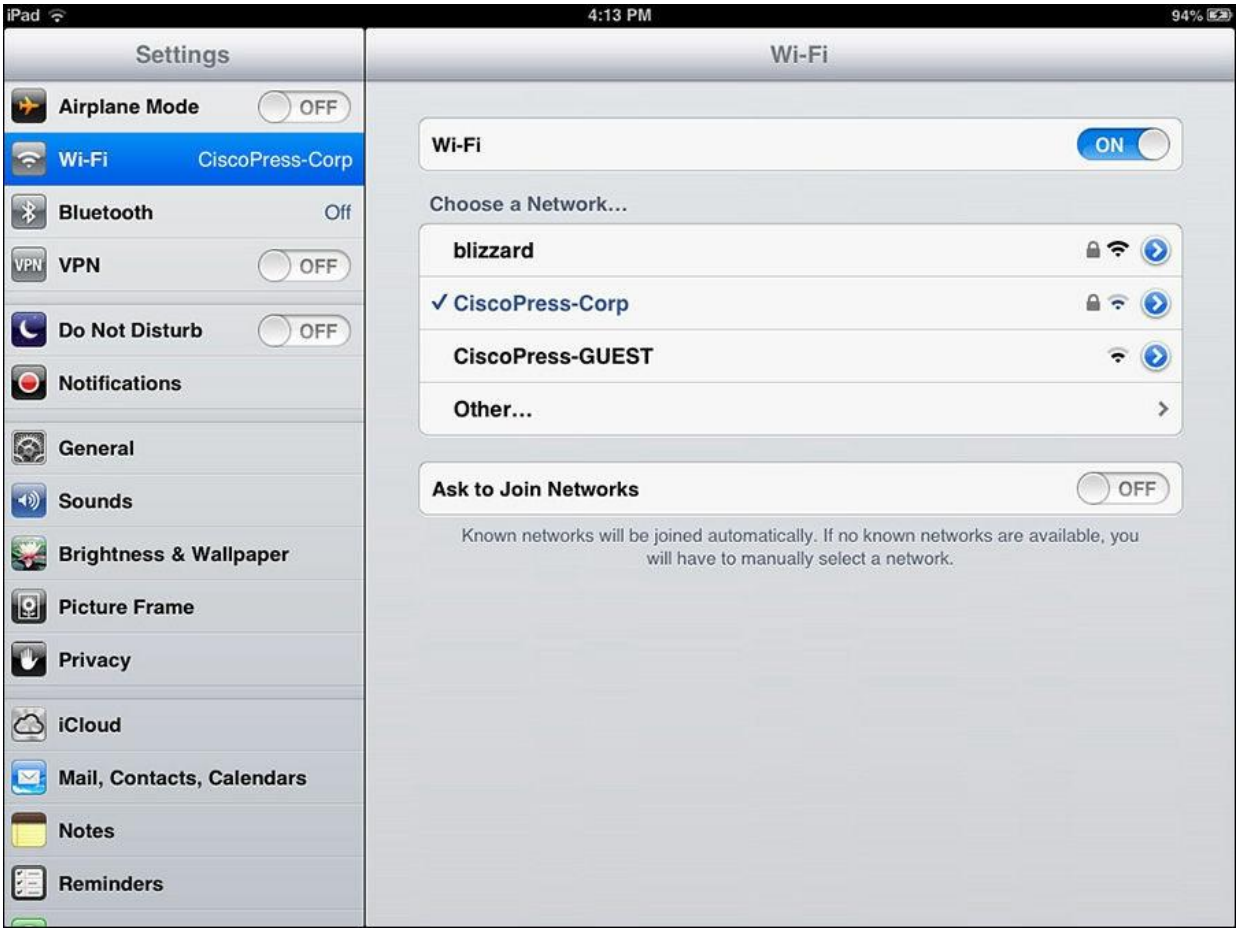


Figure 17-16 iOS: Choose a Wi-Fi network.

Step 2. The user is prompted to input a username and password. He will use his Active Directory

https://t.me/learningnets

username and password, just like in Figure 17-17.

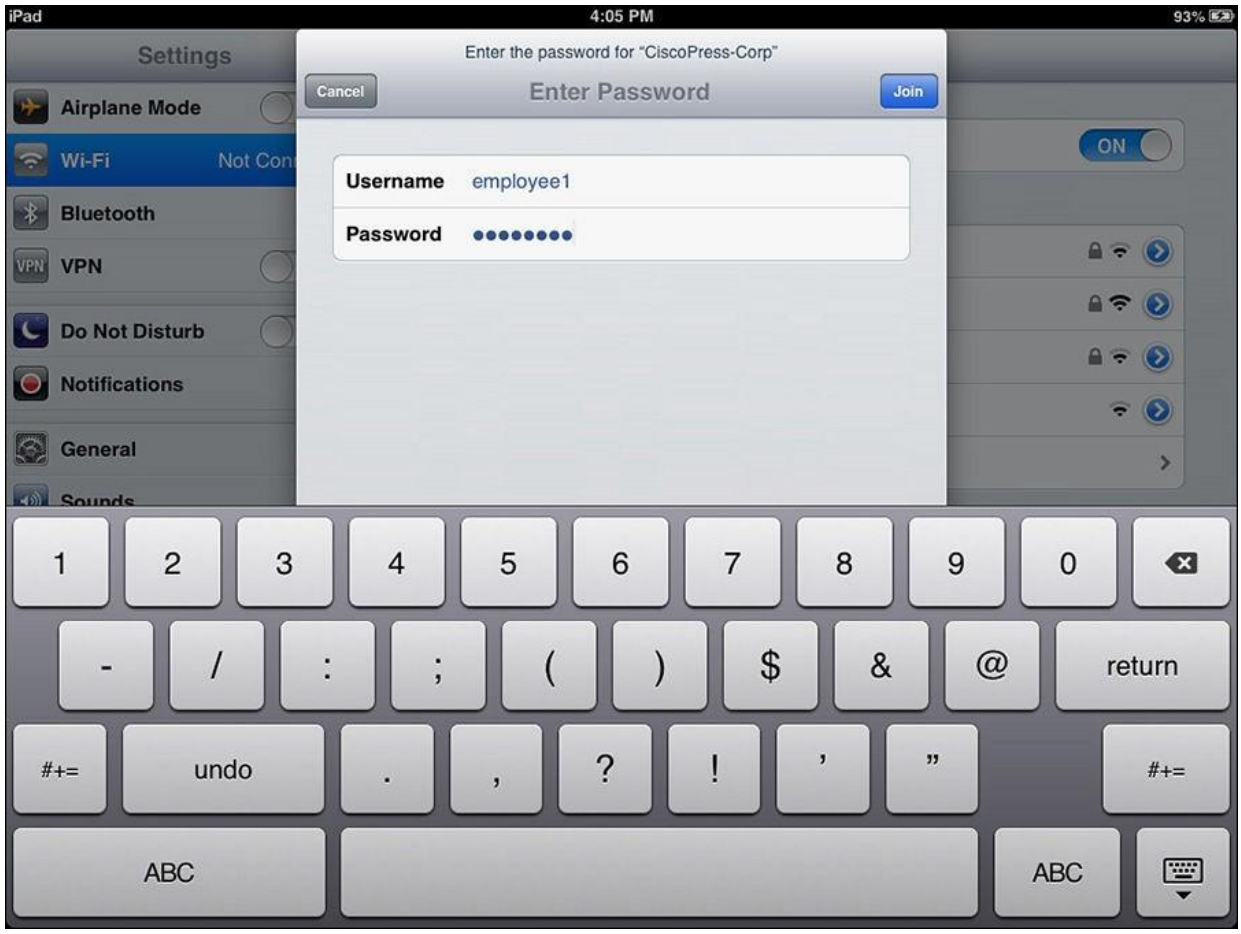


Figure 17-17 iOS: Enter credentials.

Step 3. If ISE's certificate is not signed by a trusted root, the user will be prompted to accept (trust) the certificate used by ISE, as shown in Figure 17-18.

https://t.me/learningnets

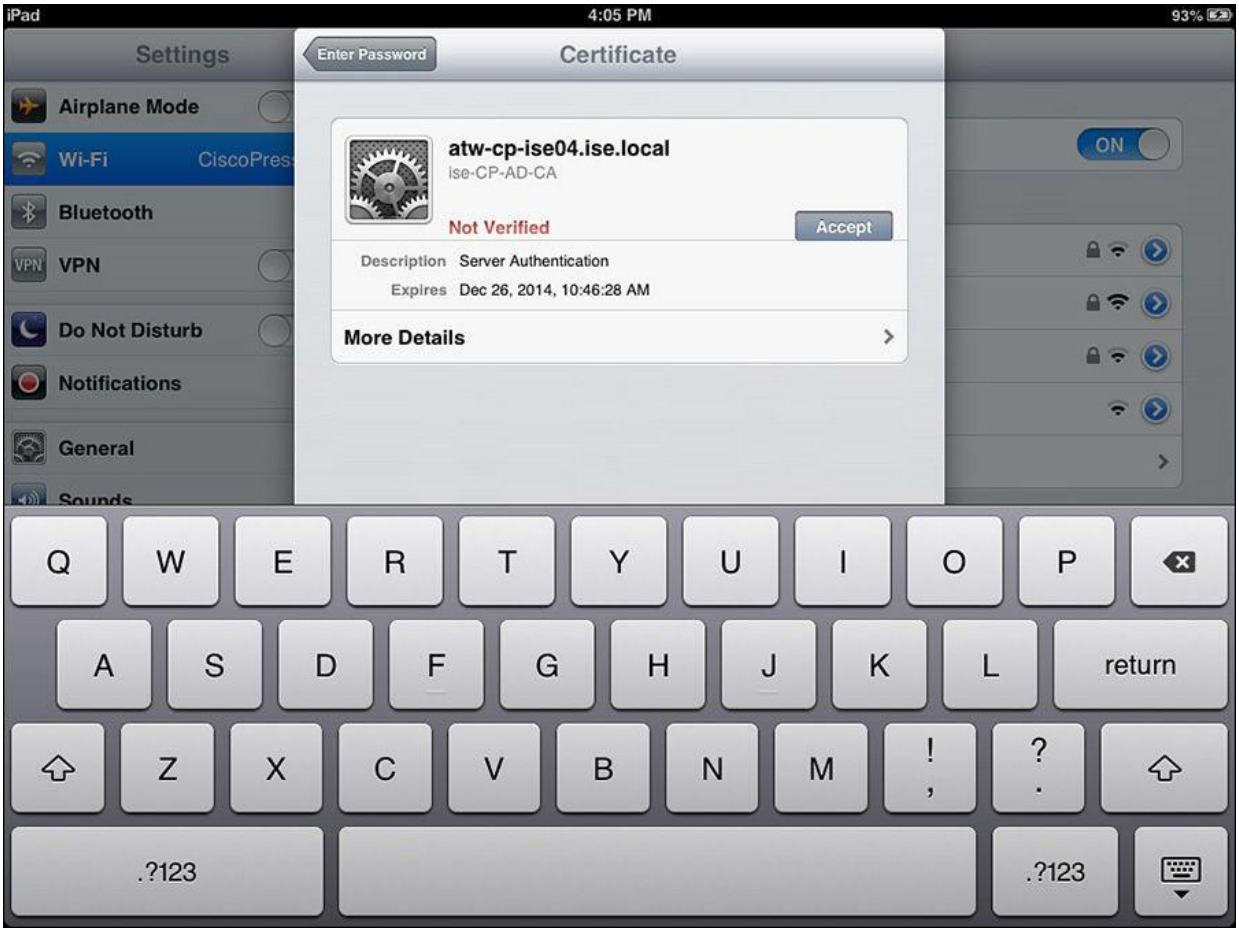


Figure 17-18 iOS: Trust the ISE certificate.

Step 4. The user will now see that he is successfully connected to the corporate network. He will not know that his access is actually limited, similar to what is shown in Figure 17-19.

https://t.me/learningnets

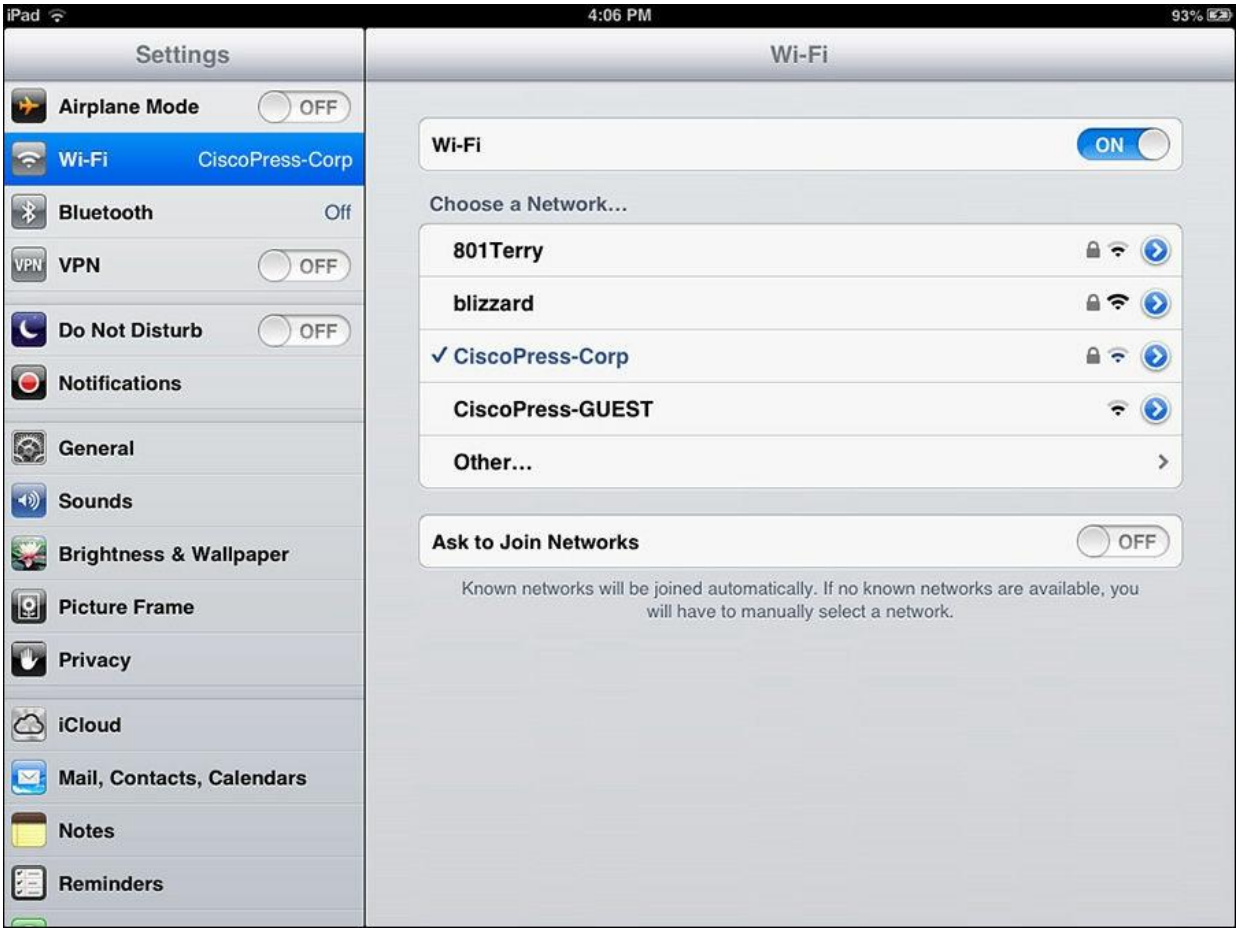


Figure 17-19 iOS: Connected to corporate Wi-Fi.

Step 5. When the end user opens a web browser, he is redirected to the Client Provisioning Portal, where 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, as shown in Figure 17-20.

https://t.me/learningnets

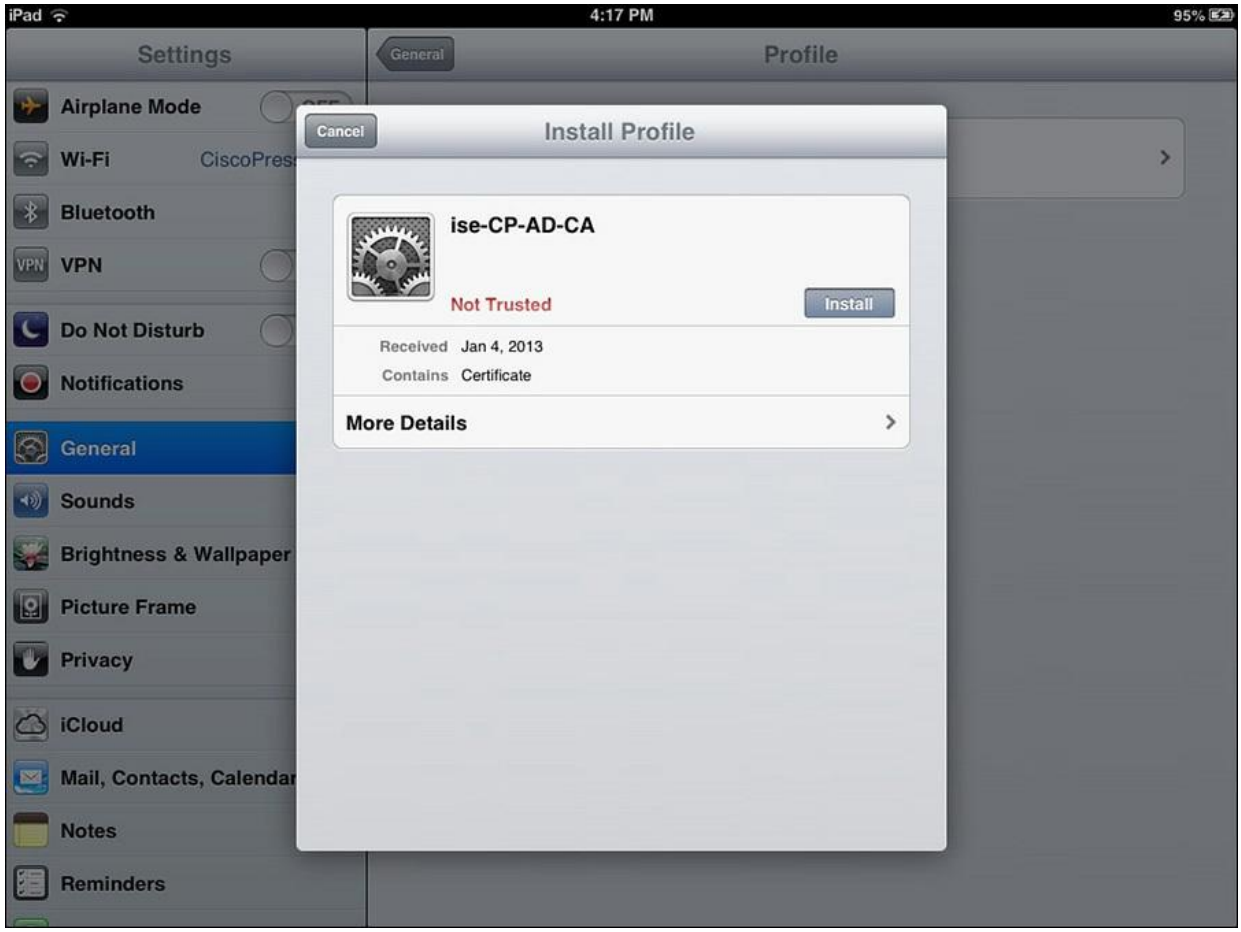


Figure 17-20 iOS: OTA trust the root CA.

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). This is displayed in Figure 17-21.

https://t.me/learningnets

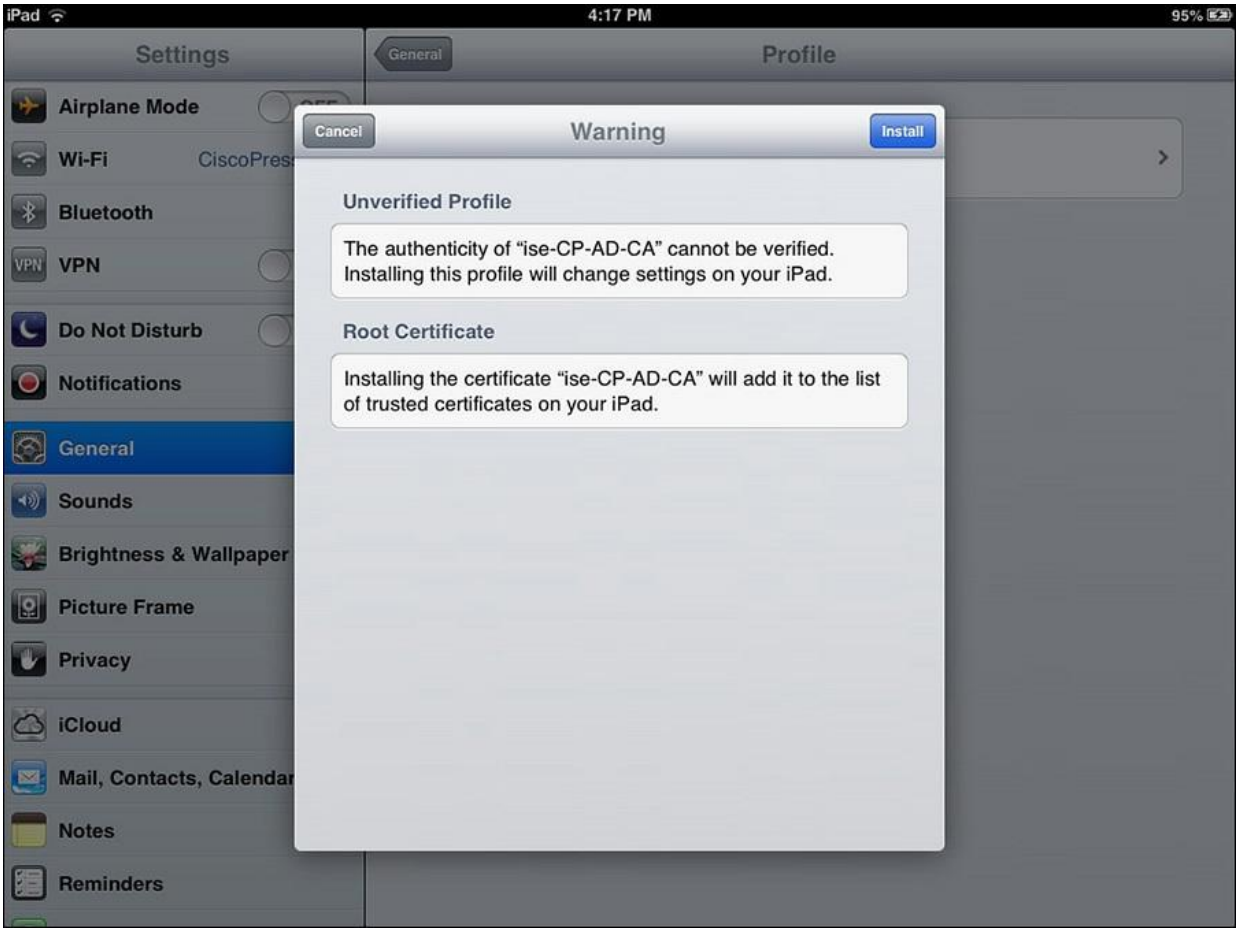


Figure 17-21 iOS: OTA warning message.

Step 8. The Certificate and Profile to allow OTA are successfully installed. The user clicks Done, as shown in Figure 17-22.

https://t.me/learningnets

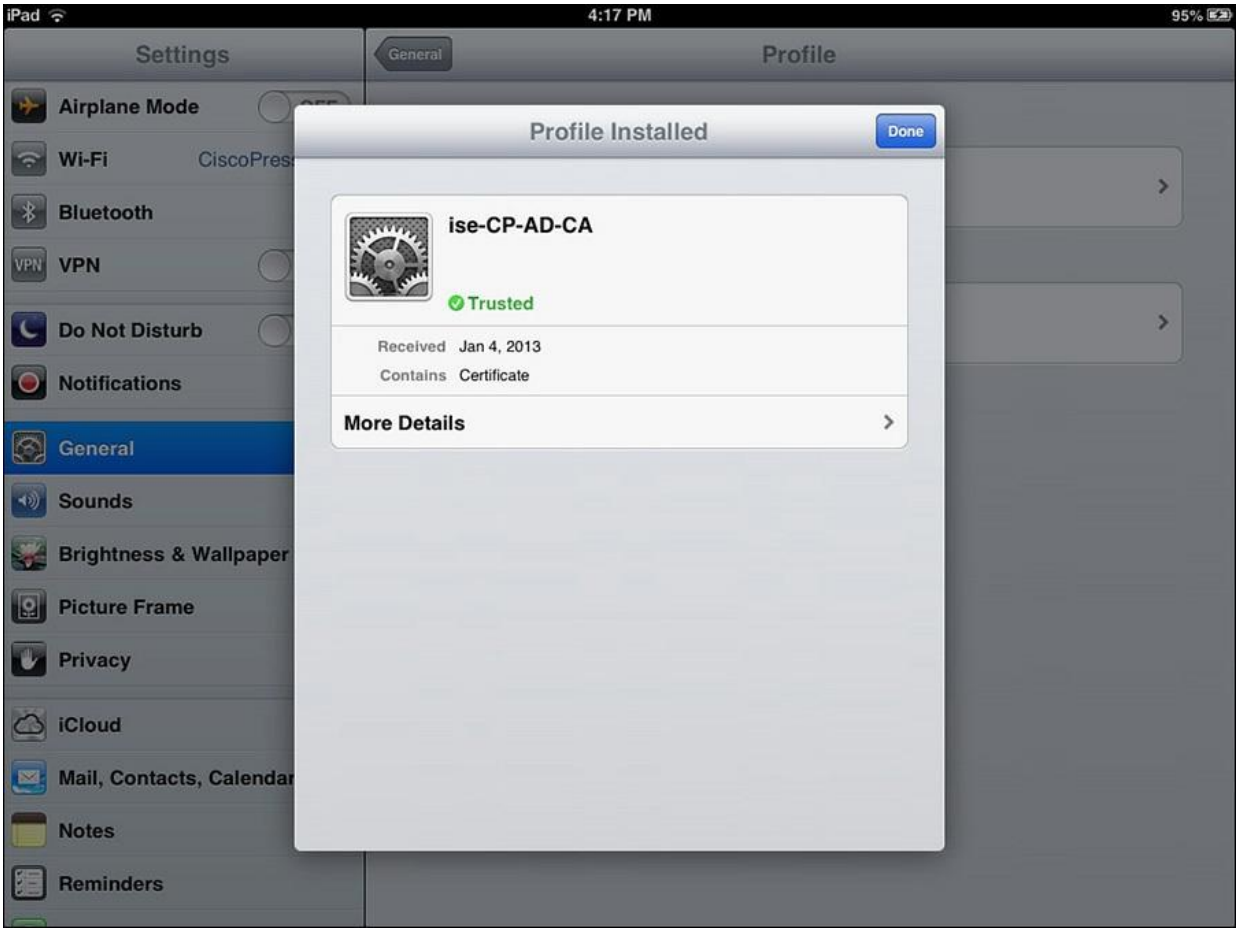


Figure 17-22 iOS: OTA profile installed.

Step 9. The user is returned to his browser window, which displays the Device Registration page within the MyDevices portal. The device's MAC address will be prepopulated and not editable. A description field will be displayed for the user to fill out, as shown in Figure 17-23.

https://t.me/learningnets

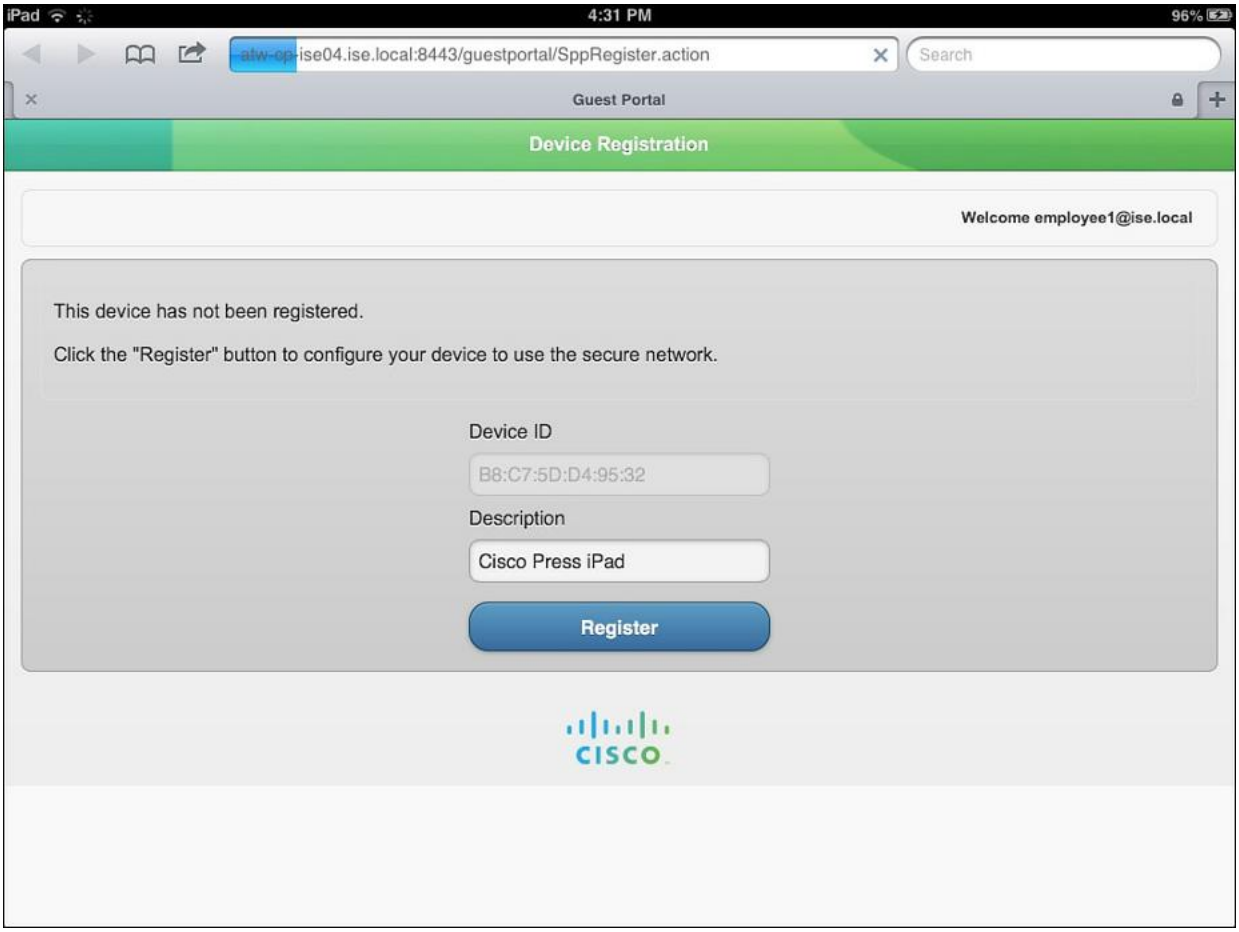


Figure 17-23 iOS: Device registration page.

Step 10. When Register is clicked, the screen immediately changes to the Install Profile page, as shown in Figure 17-24.

https://t.me/learningnets



Figure 17-24 iOS: Install the profile.

Step 11. The user is warned that clicking Install will actually install a profile. The user must click Install Now, as shown in Figure 17-25.

https://t.me/learningnets

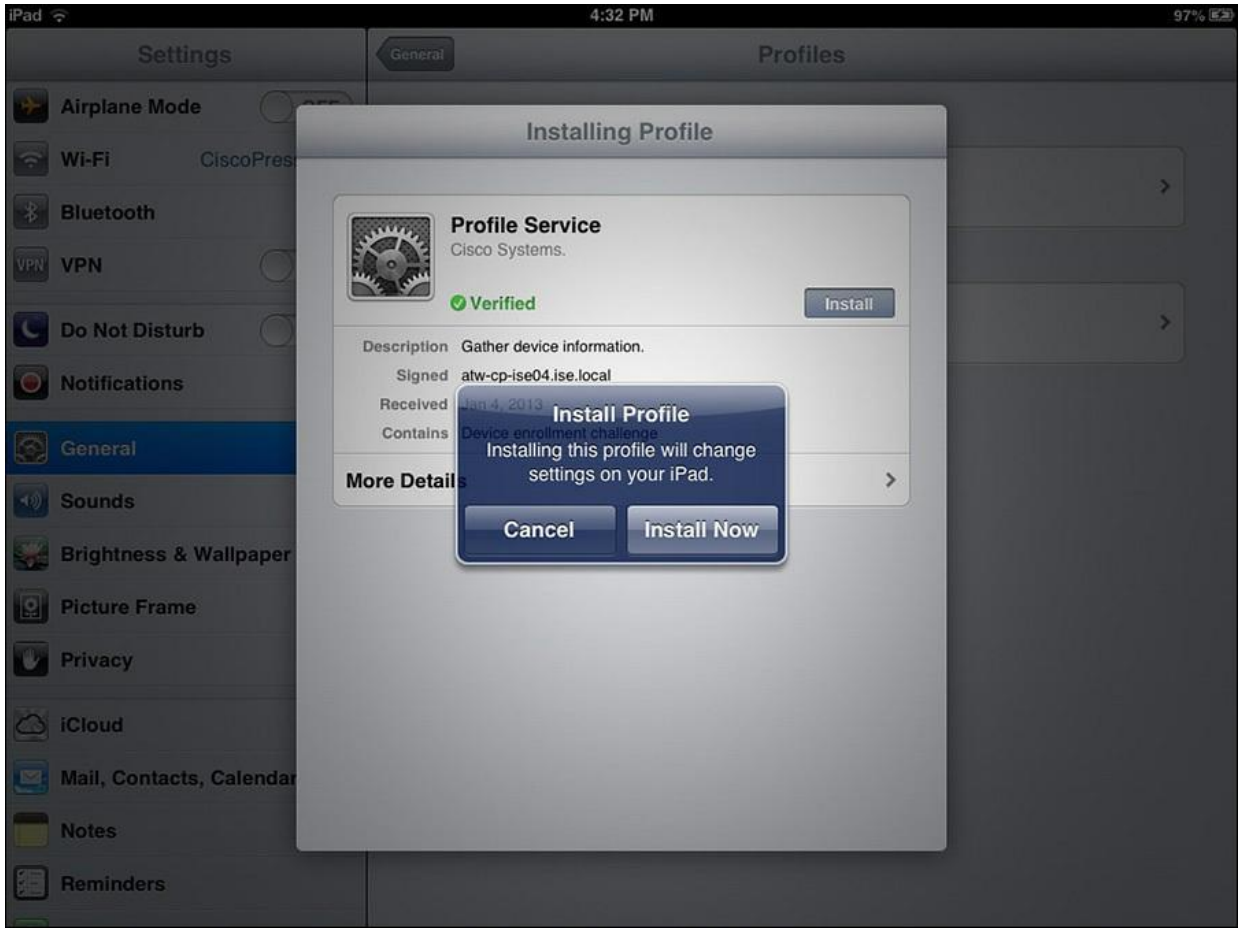


Figure 17-25 iOS: Redundant warning message.

Step 12. The profile begins to install, generates a certificate using SCEP, and prepares the device to be connected to the corporate SSID.

Figures 17-26 through 17-29 illustrate what the end user will see on the screen automatically. No user action is taken until the end user clicks Done in Step 13.

https://t.me/learningnets



Figure 17-26 iOS: Gathering device information.

https://t.me/learningnets



Figure 17-27 iOS: Generating a certificate request.

https://t.me/learningnets



Figure 17-28 iOS: Installing the profile.

https://t.me/learningnets

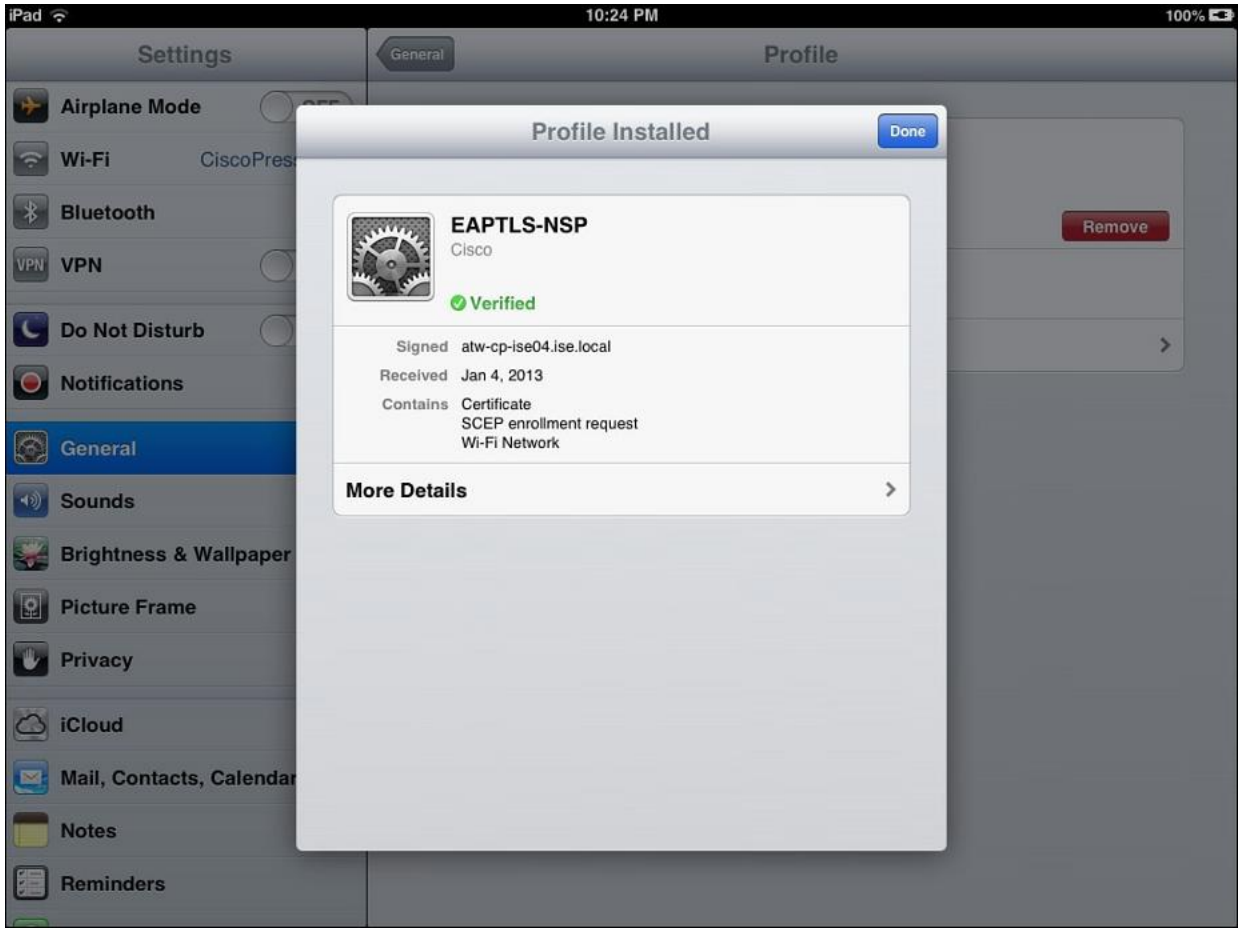


Figure 17-29 iOS: Profile is installed.

Step 13. When the profile is installed, the user must click Done. Then he is returned to his web browser where the success message is waiting for him, as shown in Figure 17-30.

https://t.me/learningnets

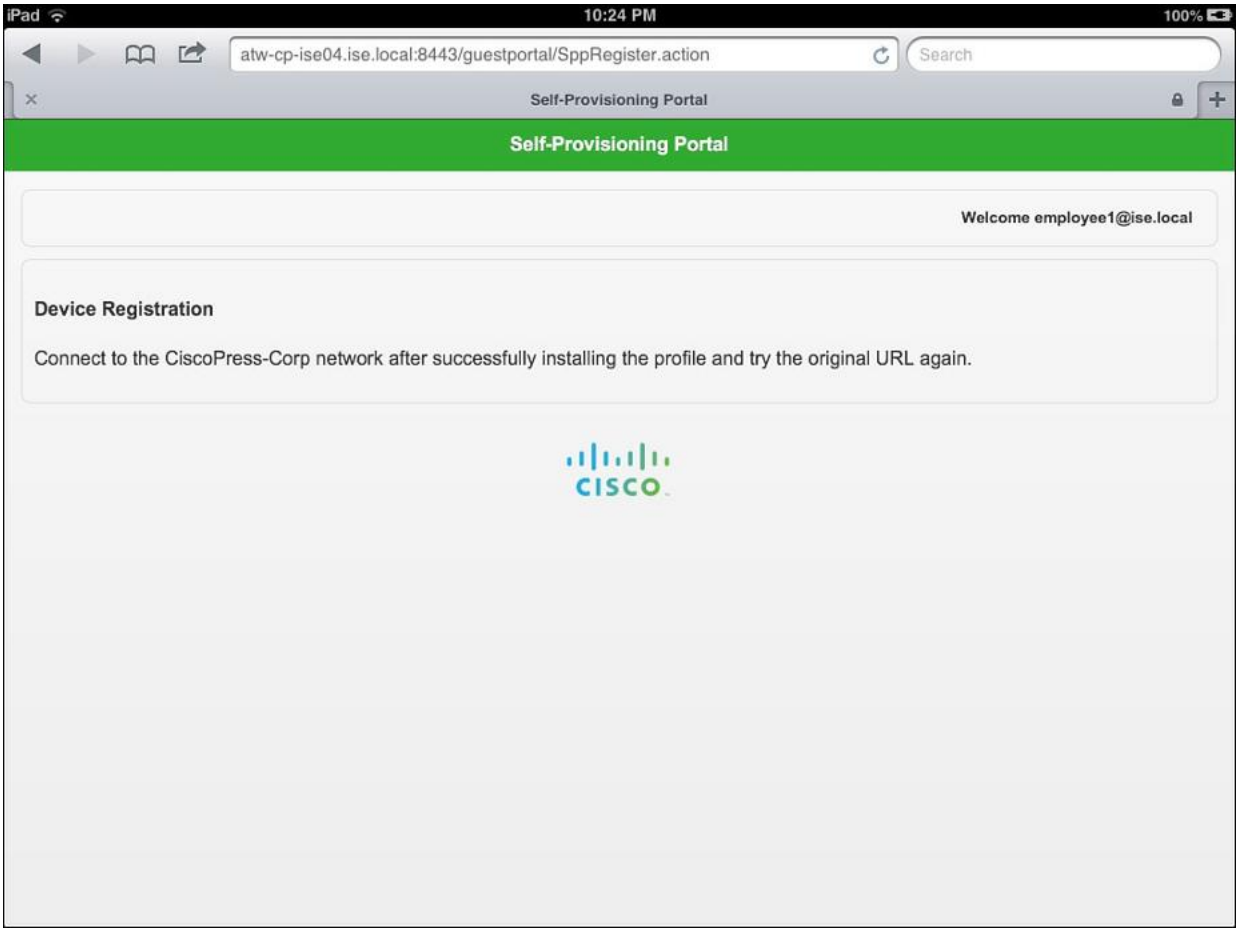


Figure 17-30 iOS: Success message.

Step 14. The user is now able to browse resources on the network, which is shown in Figure 17-31.

https://t.me/learningnets



Figure 17-31 iOS: Final network access.

That concludes the onboarding process for iOS with a single SSID. The user's iOS device is now ready to be used regularly on the corporate network. The onboarding was a one-time thing. Future attempts to access the network will happen automatically using the downloaded certificate as authentication.

Next, we will examine the user experience with dual SSID by using an Android example.

Dual SSID with Android Example

For this example, you will see an Android device connecting to secured network with the SSID of "CiscoPress-Corp." This does not match the SSID you have been using so far because you have not completed the entire configuration. For this example, don't worry about following along on your own; instead just try to see and understand the demonstrated end user experience and why it is happening.

Step 1. The user comes in with her Android device. She opens Settings and connects to the Guest Wi-Fi, such as what is shown in Figure 17-32.

https://t.me/learningnets

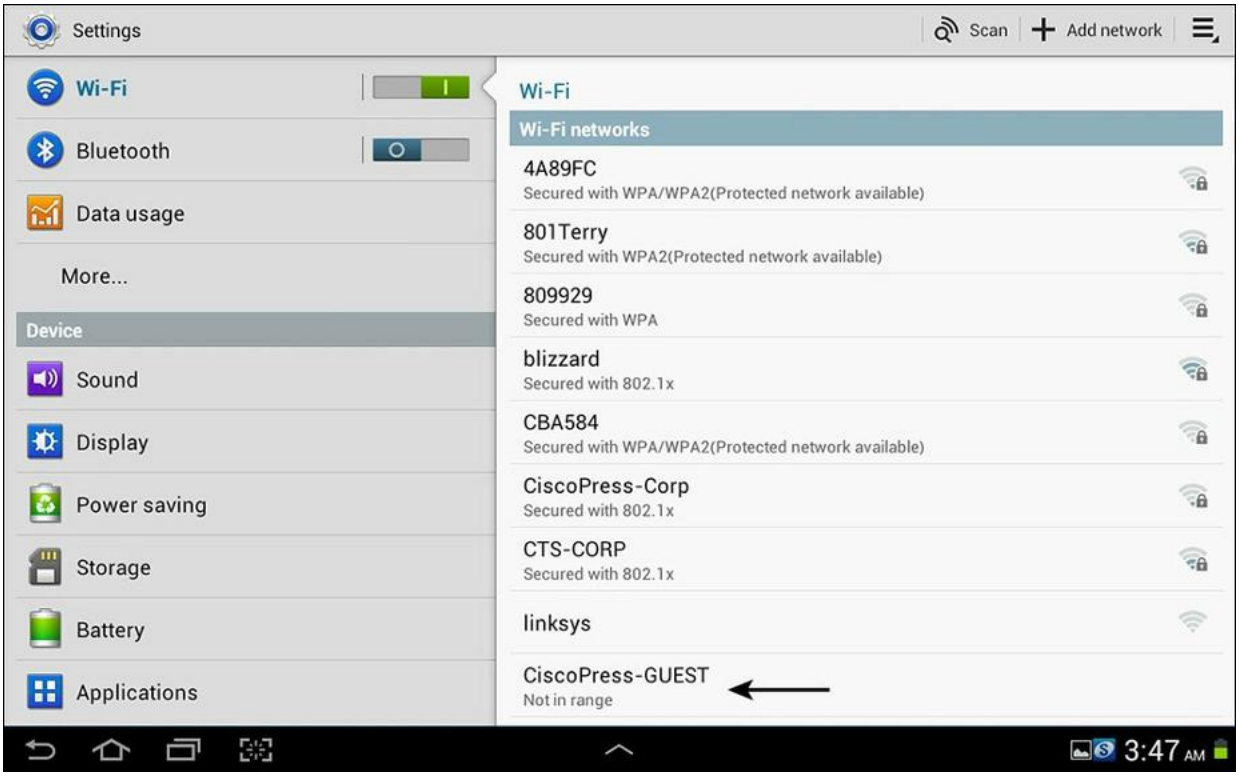


Figure 17-32 Android: Choose a Wi-Fi network.

Step 2. Because you have protected the Guest Wi-Fi by requiring a login (Guest or Active Directory), any attempts to browse the Internet will result in the user being redirected to the Web Authentication page, as is displayed in Figure 17-33.

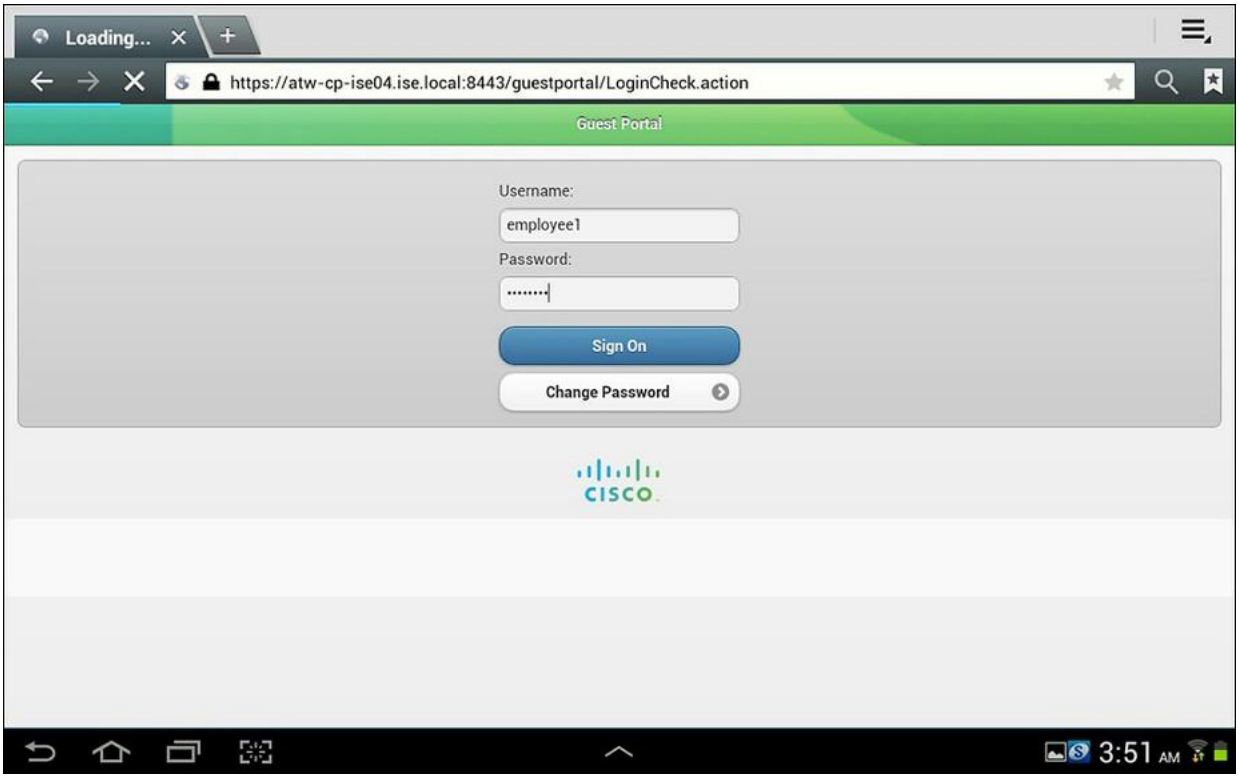


Figure 17-33 Android: Web Auth portal.

Step 3. After the user logs in to the WebAuth portal, she is redirected to the device registration page

https://t.me/learningnets

and the device ID is predefined, as shown in Figure 17-34.

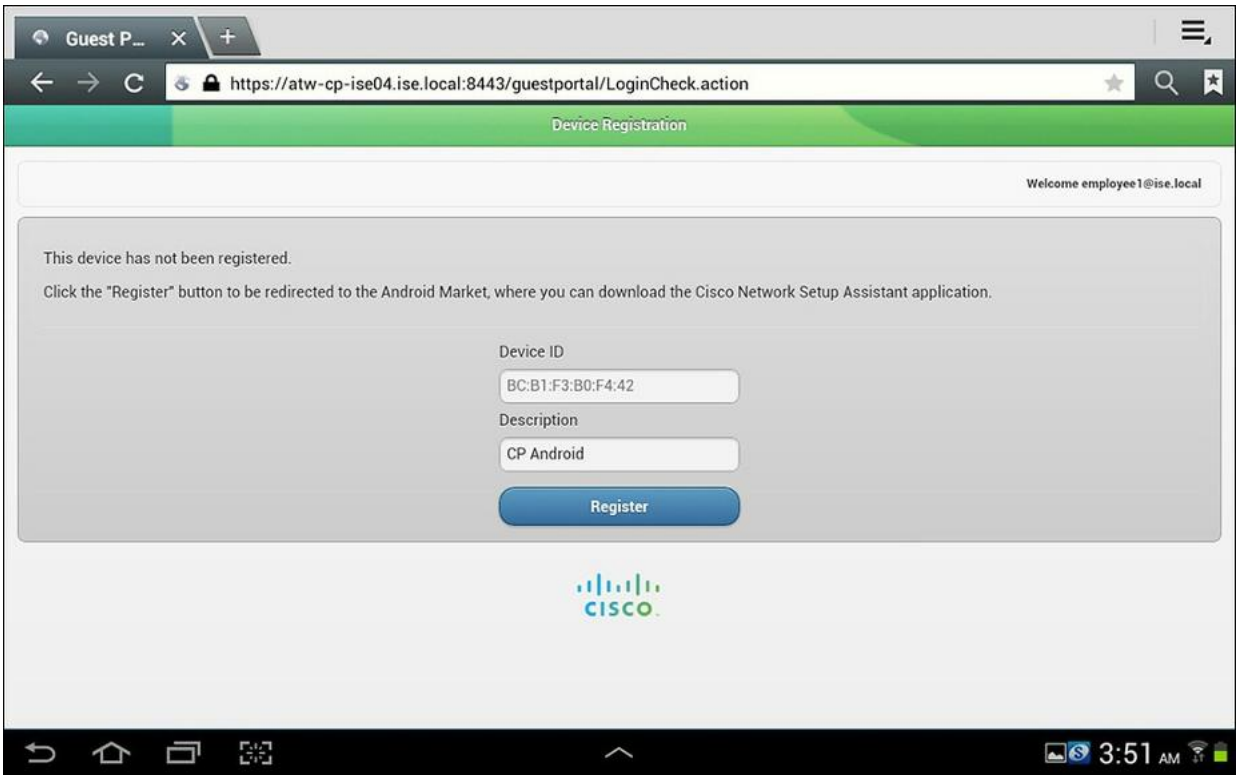


Figure 17-34 Android: MyDevice registration.

Step 4. When the user clicks Register, she is prompted to connect to the Android Marketplace (play.google.com) and is given the choice between the Internet and the Play Store App, which is shown in Figure 17-35.

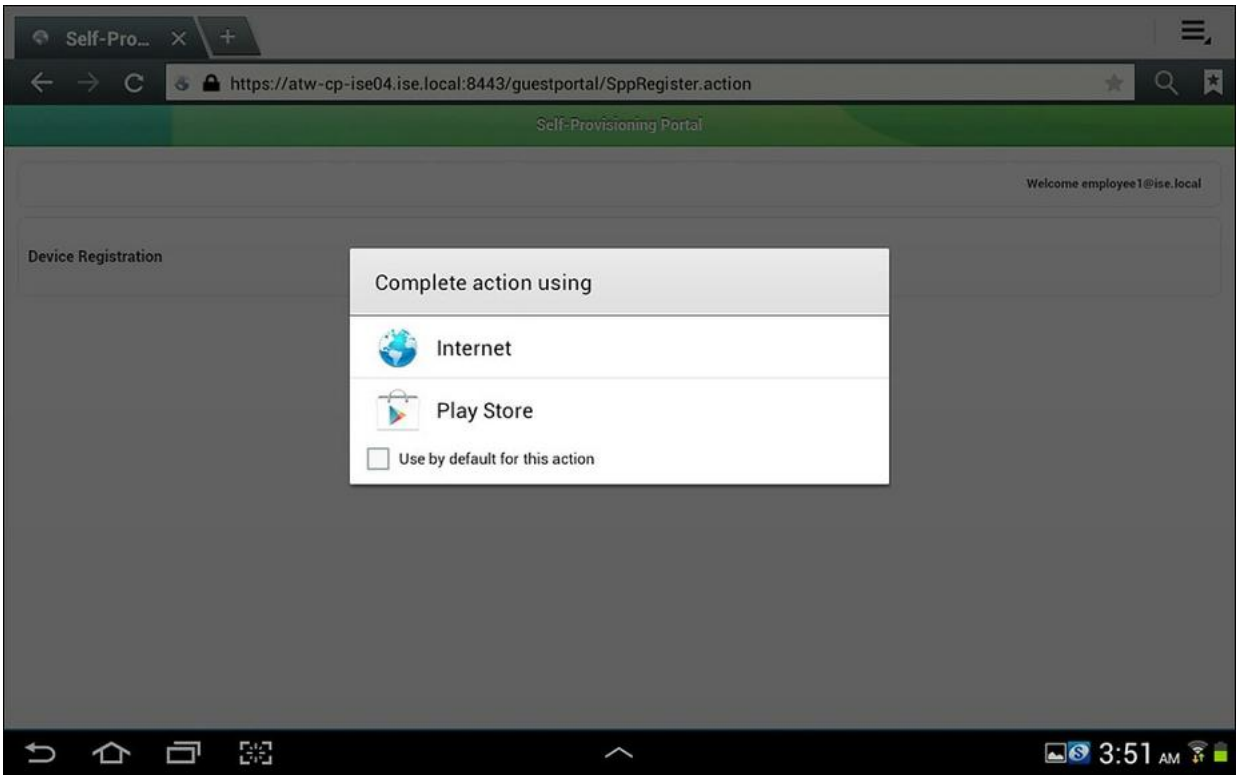


Figure 17-35 Android: Connect to Android Marketplace.

https://t.me/learningnets

Step 5. The user then downloads the app from the marketplace, similar to what is shown in Figure 17-36.

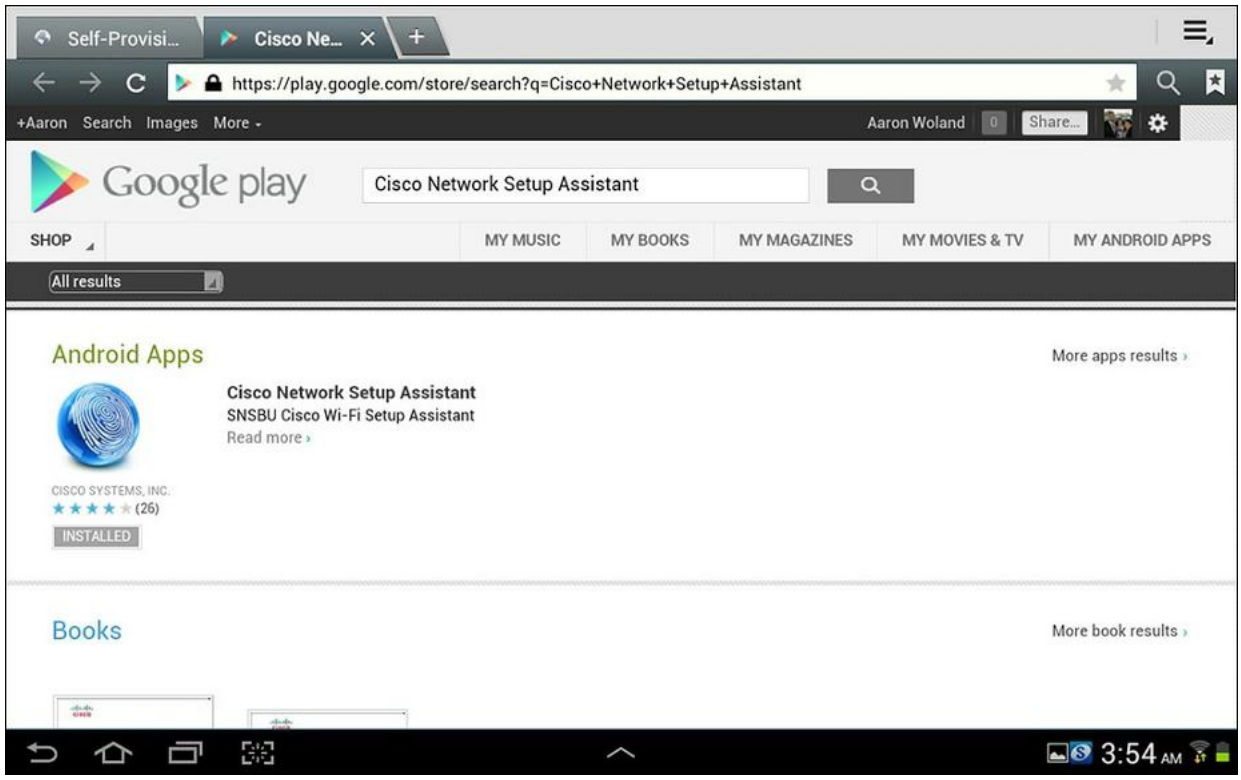


Figure 17-36 Android: Download the Cisco Network Setup Assistant.

Step 6. The user runs the app, and a screen like the one shown in Figure 17-37 appears, where the user must click Start.

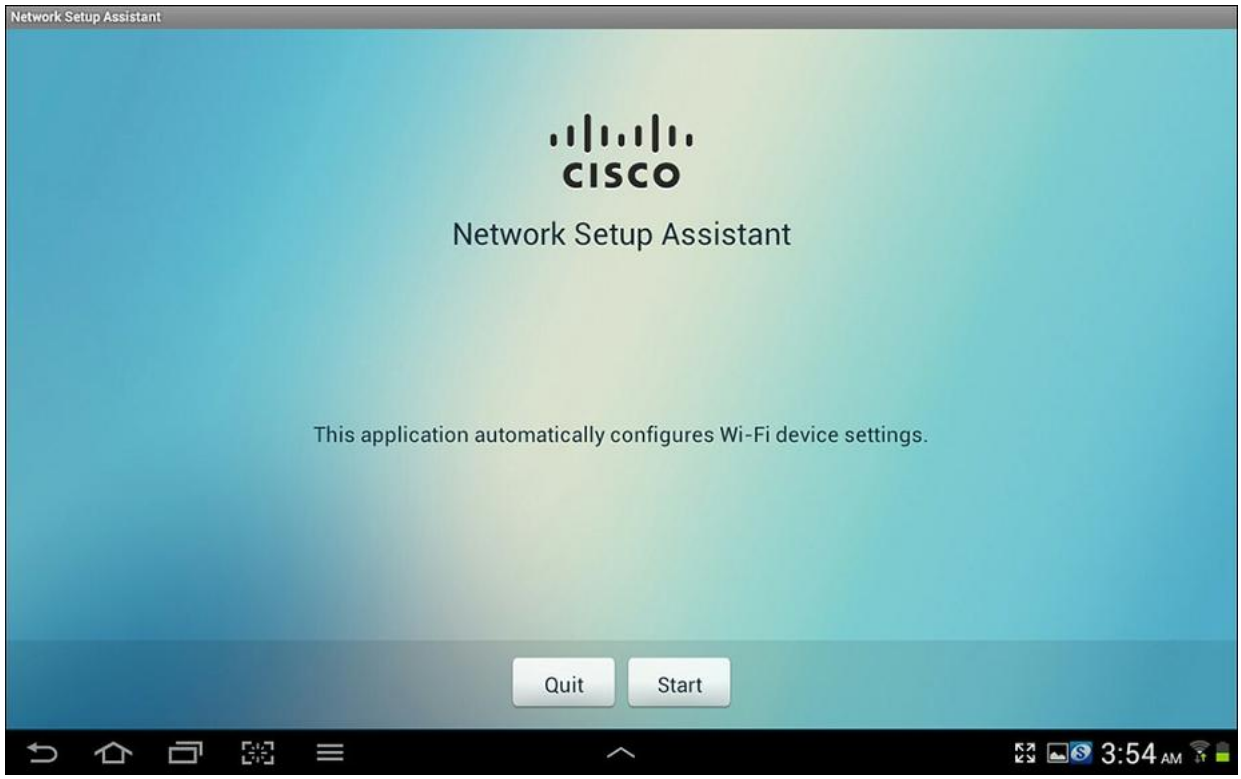


Figure 17-37 Android: Run the NSP app.

https://t.me/learningnets

Step 7. The NSP app downloads the profile from ISE, as shown in Figure 17-38.

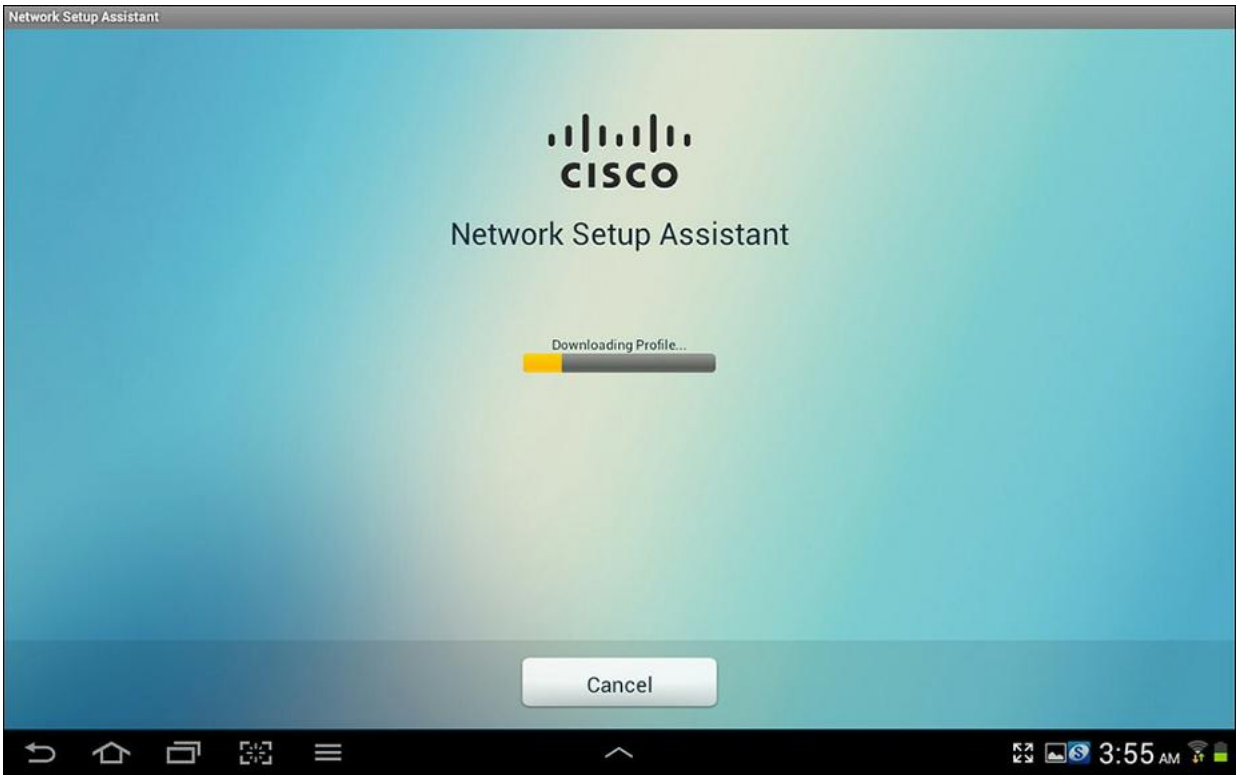


Figure 17-38 Android: NSP app downloading profile.

Step 8. As displayed in Figure 17-39, the user must name the user identity certificate.

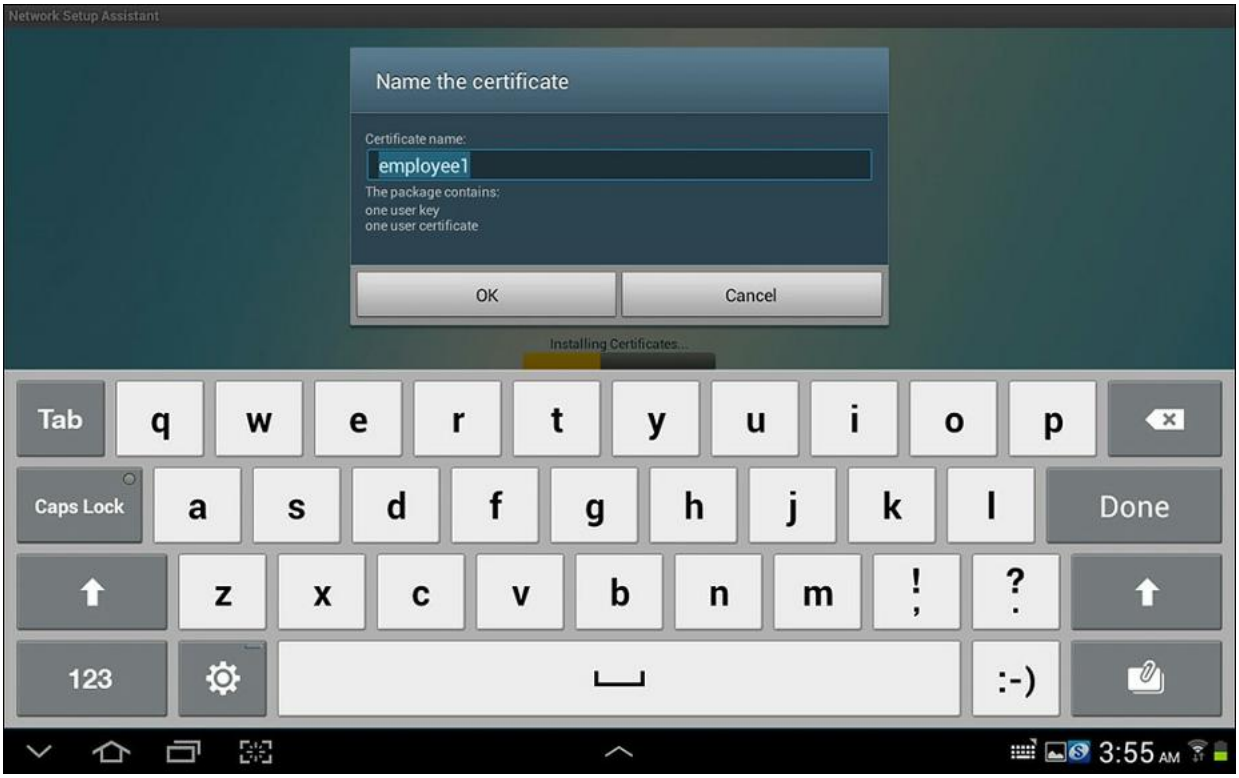


Figure 17-39 Android: Name the certificate.

Step 9. The user must now provide a name for the CA's public certificate, as shown in Figure 17-40.

https://t.me/learningnets

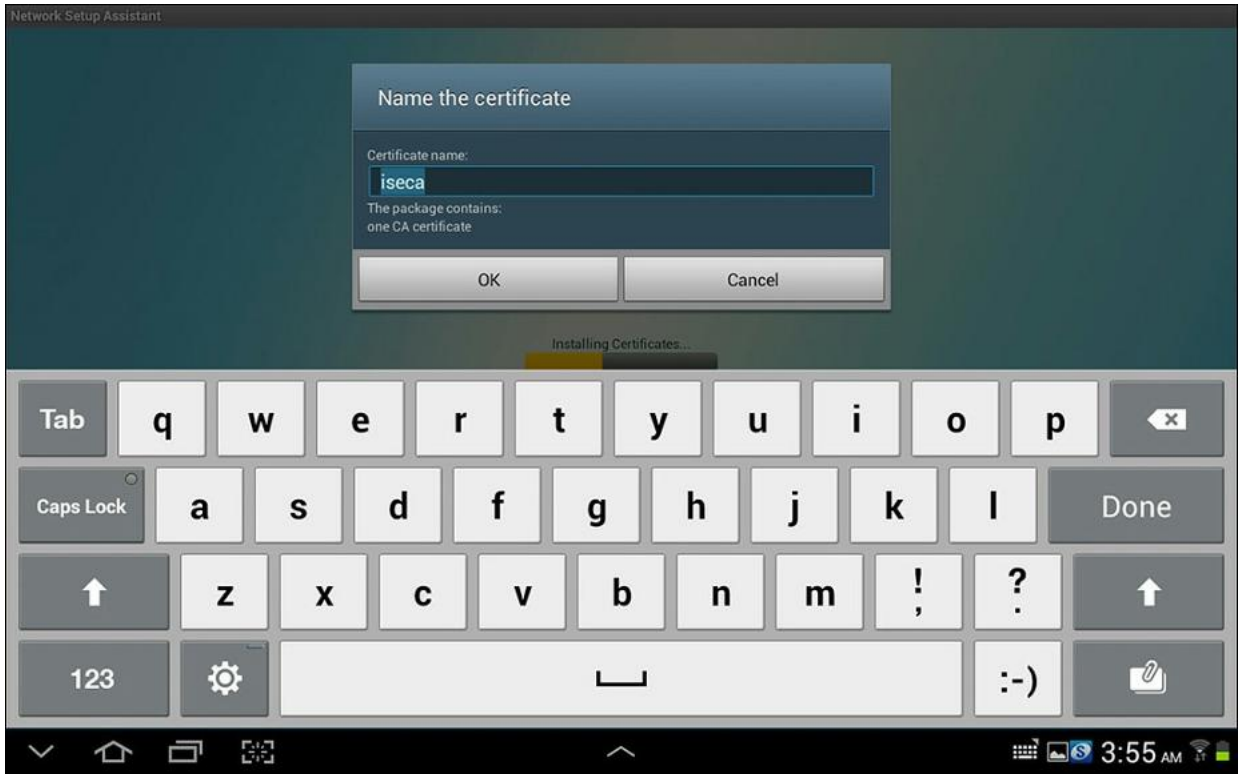


Figure 17-40 Android: Name the CA certificate.

Step 10. The NSP app automatically changes the network connection to the Corporate SSID and authenticates with the new certificate using EAP-TLS, as shown in Figure 17-41.

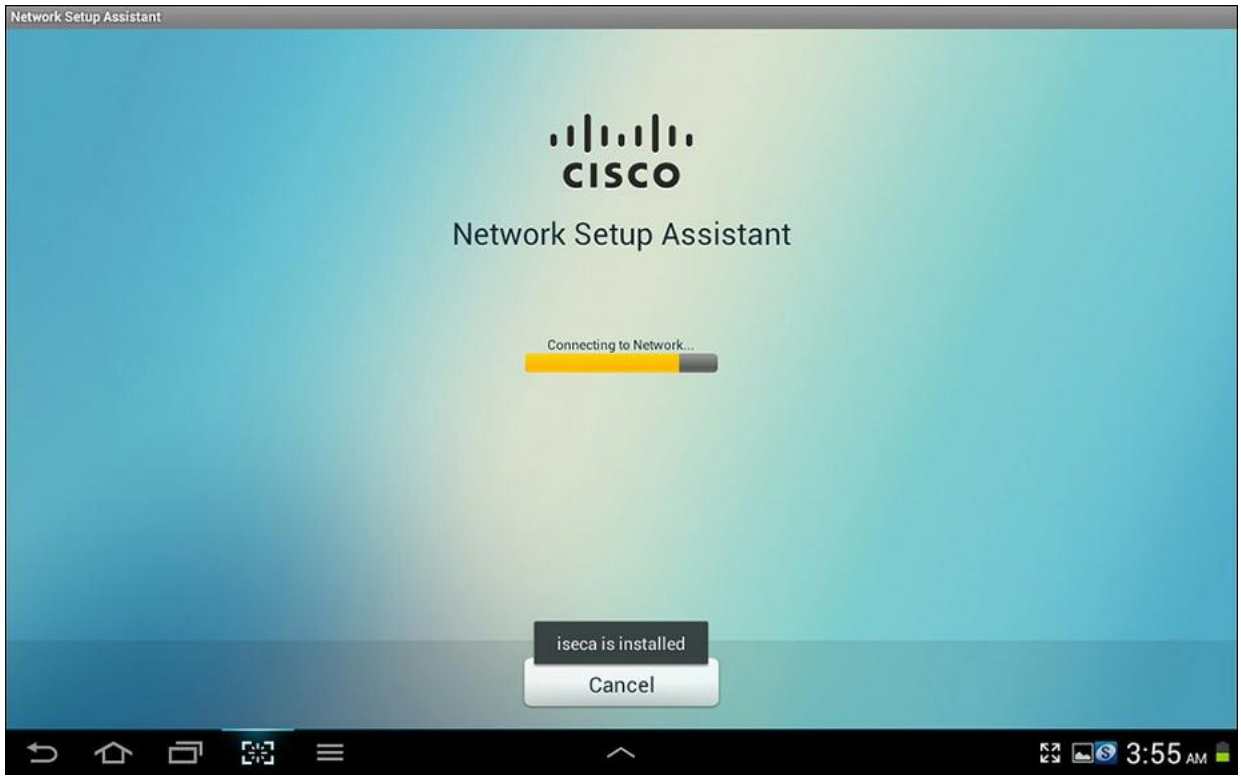


Figure 17-41 Android: Connect to the corporate SSID.

Step 11. As shown in Figure 17-42, the process is now complete. The user's Android device is now ready to be used regularly on the corporate network. The onboarding was a one-time thing.

https://t.me/learningnets

Future attempts to access the network will happen automatically using the downloaded certificate as authentication.

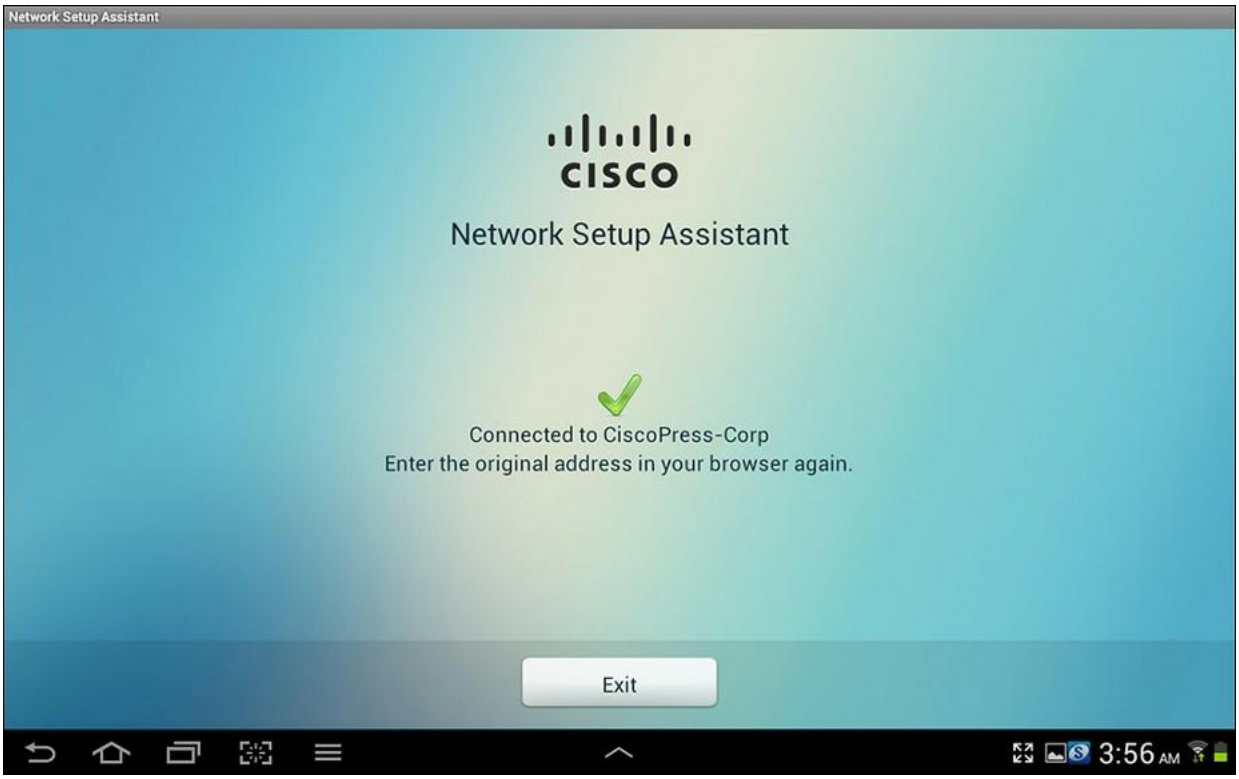


Figure 17-42 Android: Done.

Unsupported Mobile Device—Blackberry Example

You have now seen both single-SSID and dual-SSID onboarding with supported devices. You will now see a brief example of a device that is not supported. Figures 17-43 through 17-47 show screen captures from an unsupported BlackBerry device.

Step 1. A user brings his Blackberry mobile device to work and selects the GUEST Wireless network, as shown in Figure 17-43.



Figure 17-43 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 17-44.

https://t.me/learningnets



Figure 17-44 Blackberry: Web authentication.

Step 3. If the Global setting is Apply Defined Authorization Policy, the user receives the message shown in Figure 17-45 and is not able to gain network access (this is covered in more detail later in the chapter).



Figure 17-45 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 web portal (this is covered in more detail later in the chapter).

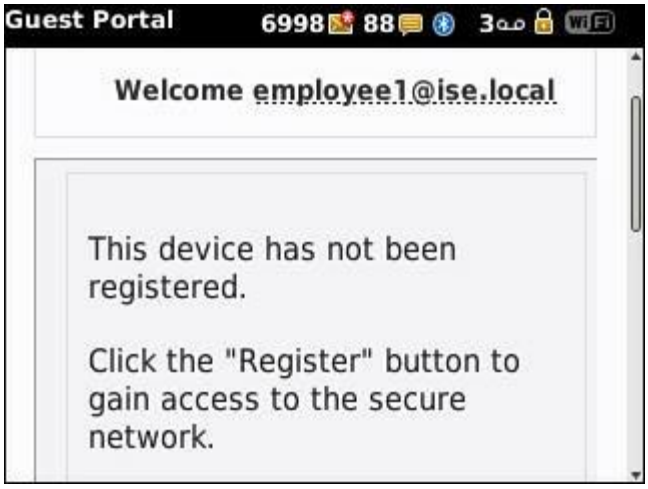


Figure 17-46 Blackberry: Registration permitted.

https://t.me/learningnets

Step 5. The user is now allowed to register the device and gain network access by manually configuring his supplicant, as shown in Figure 17-47.



Figure 17-47 Blackberry: Registering the BlackBerry.

Configuring ISE for Onboarding

The end user experience is designed to be very straightforward and easy for a typical user to be able to follow without any interaction with the IT department. To keep things so easy for the end user, there is quite a bit of up-front work you will need to do on the configuration side. We will cover that configuration in this section.

The first step is to create the Native Supplicant Profile (NSP) that will be sent to the end user's device. From there, you will configure the Client Provisioning Portal so the correct profiles are sent for the correct operating systems. Next, you will configure the default action that should be taken when an unsupported device attempts to connect and cannot be provisioned (such as the BlackBerry example you just saw).

Creating the Native Supplicant Profile

The NSP is a required object for BYOD onboarding. The NSP 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
- If the profile applies to wired, wireless or both

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Policy Elements > Results > Client Provisioning > Resources**.

Step 2. Click **Add > Agent Resources from Cisco site**.

Step 3. Select the latest versions of the clients and wizards.

Step 4. Click **Save**.

Figure 17-48 displays the screen to download agent resources from Cisco's site.

https://t.me/learningnets

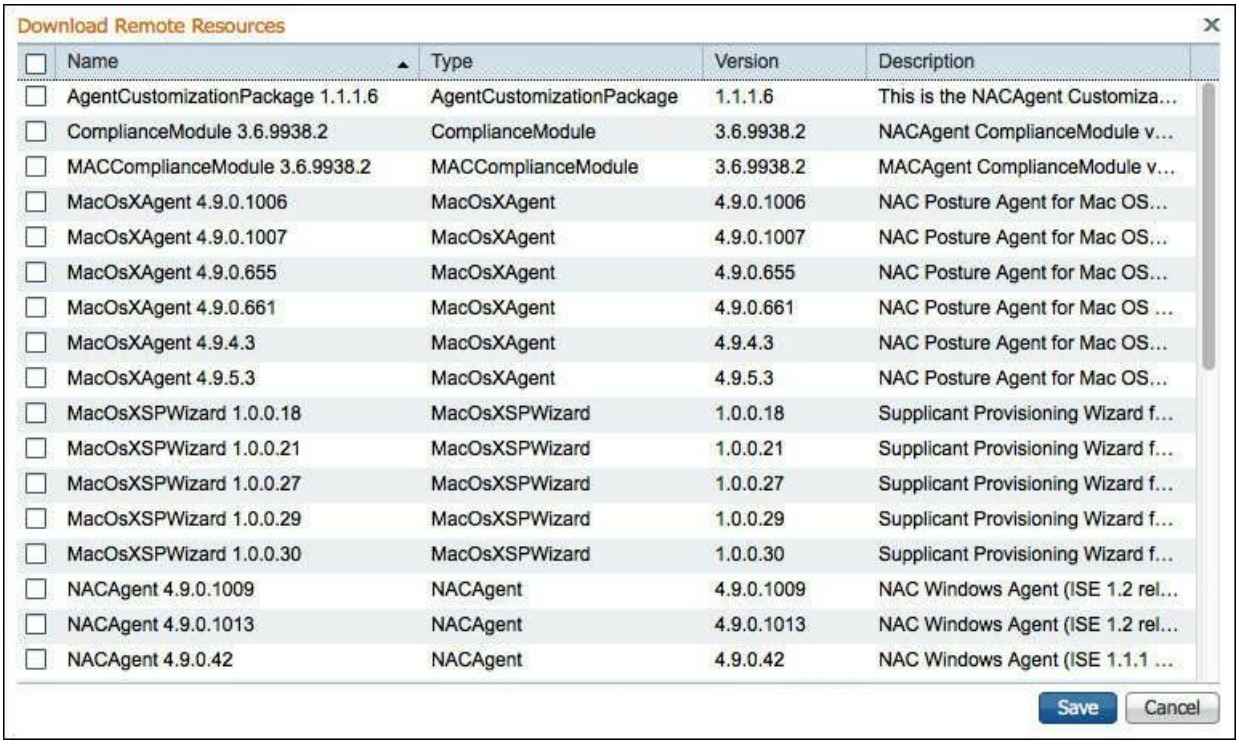


Figure 17-48 Agent resources from Cisco site.

Step 5. Click **Add > Native Supplicant Profile**.

Step 6. Name the NSP **CiscoPress-TLS**.

Step 7. The operating system can remain the default of ALL.

Step 8. Ensure that **Wireless** is checked.

Step 9. Wired is optional, so select it if you want.

Step 10. Provide the SSID for the corporate wireless network.

Step 11. Select the security level (such as **WPA2**).

Step 12. Select **TLS** for the Allowed Protocol.

Step 13. Select the certificate size (such as **2048**).

Step 14. Click **Submit**.

Figure 17-49 shows the completed NSP using EAP-TLS.

https://t.me/learningnets

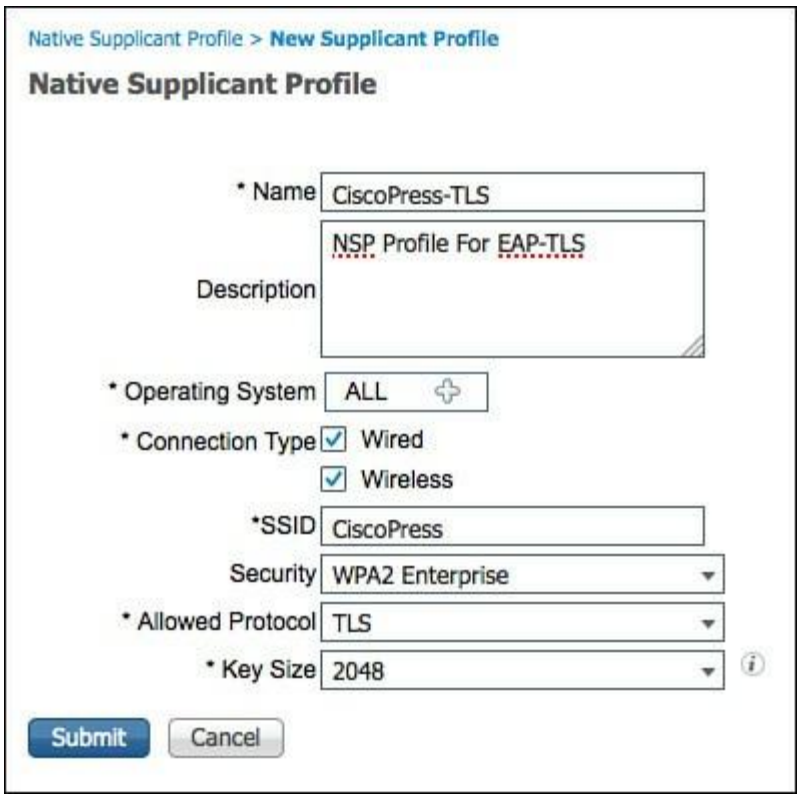


Figure 17-49 NSP for EAP-TLS.

Configuring the Client Provisioning Policy

You configure a client provisioning policy to dictate the software and profiles that 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 while configuring the CONTRACTOR-SSID for all vendors and contractors who are going through the provisioning process.

For our example, we will create one client provisioning policy per OS. Do the following:

Step 1. Navigate to **Policy > Client Provisioning**.

Step 2. Name a new rule **iOS**.

Step 3. Set the operating system as **Apple iOS All**.

Step 4. Set the result to be the **CiscoPress-TLS** supplicant profile from the drop-down.

> **Note**
>
> The ISE Client Provisioning Portal will automatically use the OTA provisioning process that is native to iOS for Apple iOS, so there is no need to specify that here.

Figure 17-50 is a composite image that has been cropped to display the completed client provisioning policy for iOS.

https://t.me/learningnets

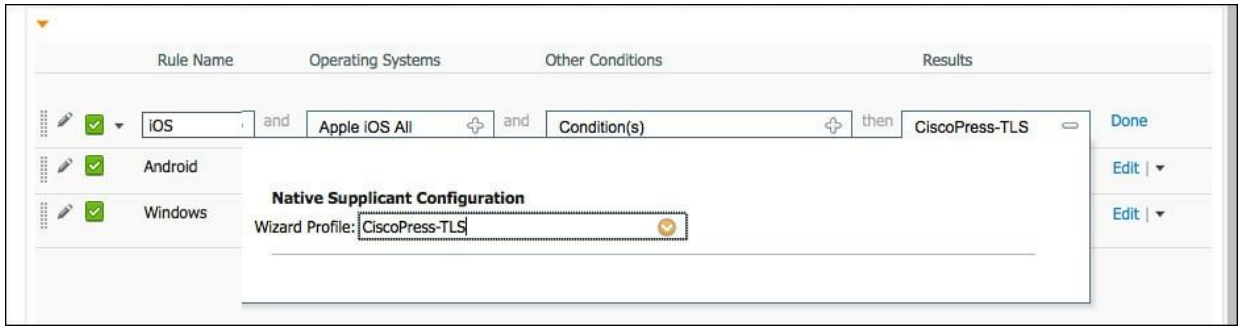


Figure 17-50 Client provisioning policy for iOS.

Step 5. Insert a new rule below the iOS rule, and name that rule **Android**.

Step 6. Set the operating system as **Android**.

Step 7. Set the result to be the **CiscoPress-TLS** supplicant profile from the drop-down.

> **Note**
>
> The ISE Client Provisioning Portal will automatically redirect 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 8. Insert a new rule below the Android rule, and name that rule **Windows**.

Step 9. Set the operating system as **Windows All**.

Step 10. Configure the results to use the **WinSPWizard** and **CiscoPress-TLS**.

The results drop-down provides many more possibilities for Windows. This is due to the ability to also provision the NAC Agent or Web Agent to the Windows operating system. Windows will use the Cisco Supplicant Provisioning Wizard (a Java applet) to do the provisioning, and that must be specified here.

Figure 17-51 is a composite image that has been cropped to display the completed Client Provisioning Policy for Windows.

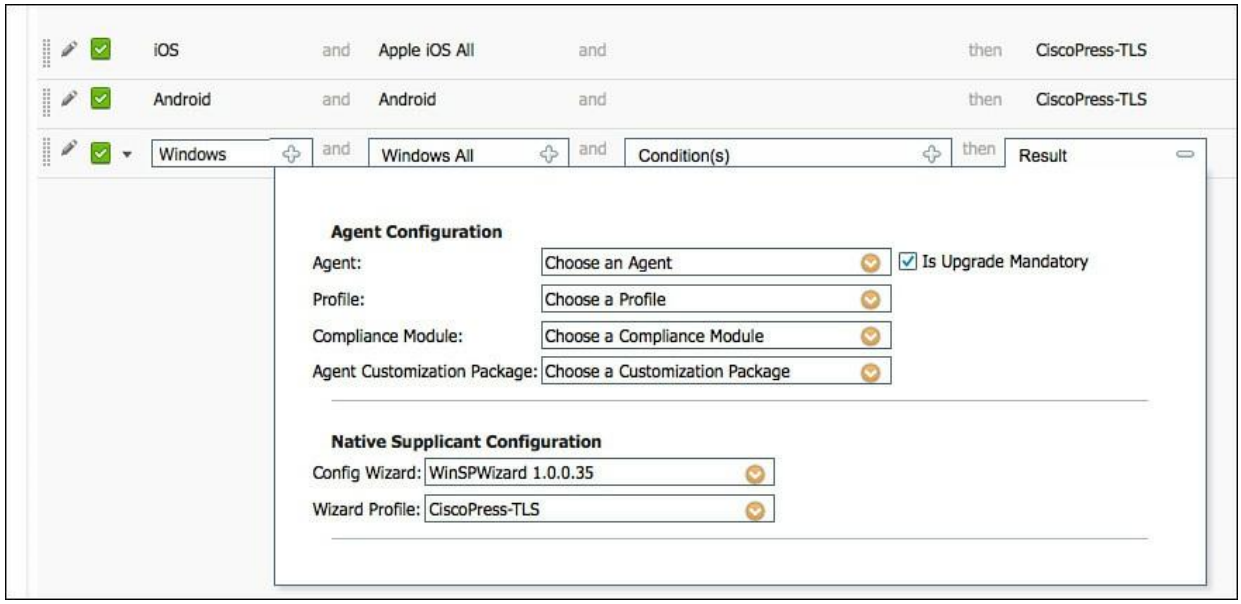


Figure 17-51 CPP results for Windows operating systems.

Step 11. Insert a new rule below the Windows rule, and name that rule **Mac OS**.

https://t.me/learningnets

Step 12. Set the operating system as **Mac OSX**.

Step 13. Configure the results to use the **MacOsXSPWizard** and **CiscoPress-TLS**.

Like Windows, the results choices for Mac OSX provide many more possibilities. This is due to the ability to also provision the NAC Agent to the Mac OSX operating system. Mac OSX will use the Cisco Supplicant Provisioning Wizard (a Java applet) to do the provisioning, and that must be specified here.

> **Note**
>
> In ISE 1.2 patch 11, the Java-based applet was replaced with a native application for both Windows and Mac OSX. However, it is not relevant for the purposes of the exam, and therefore is not included in this book.

Step 14. Click **Save**.

Figure 17-52 is a composite image that has been cropped to show the final Client Provisioning Policy, with all four OSes configured.

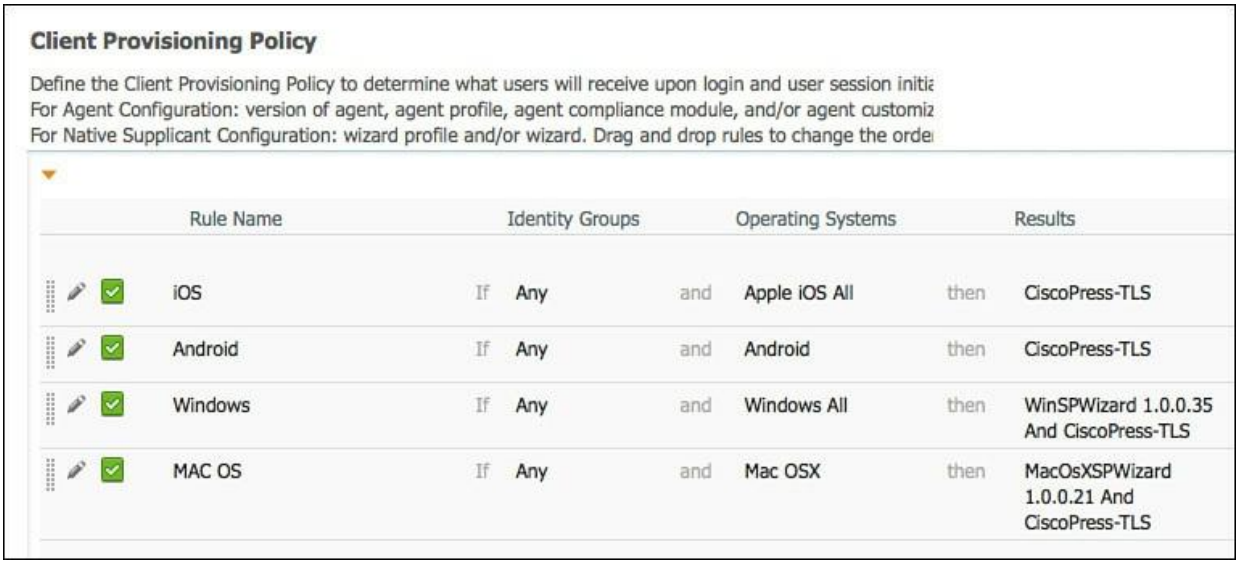


Figure 17-52 Client provisioning policy.

Configuring the WebAuth

The CPP is now created, and it will be used with both dual-SSID and single-SSID provisioning. However, we must ensure that the WebAuth Portal Page is ready for the Dual SSID flow.

There is a plethora of options when it comes to Web Authentication and supplicant provisioning. For instance, you can 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 chose.

For simplicity, we will use DefaultGuestPortal for our example. Do the following:

Step 1. Navigate to **Administration > Web Portal Management > Settings**.

Step 2. Select **Guest > Multi-Portal Configuration > DefaultGuestPortal**.

Step 3. Click the **Operations** tab.

Step 4. Ensure the **Enable Self-Provisioning Flow** check box is selected.

Step 5. Click **Save**.

Figure 17-53 shows the portal configuration with the Enable Self-Provisioning Flow check box enabled.

https://t.me/learningnets

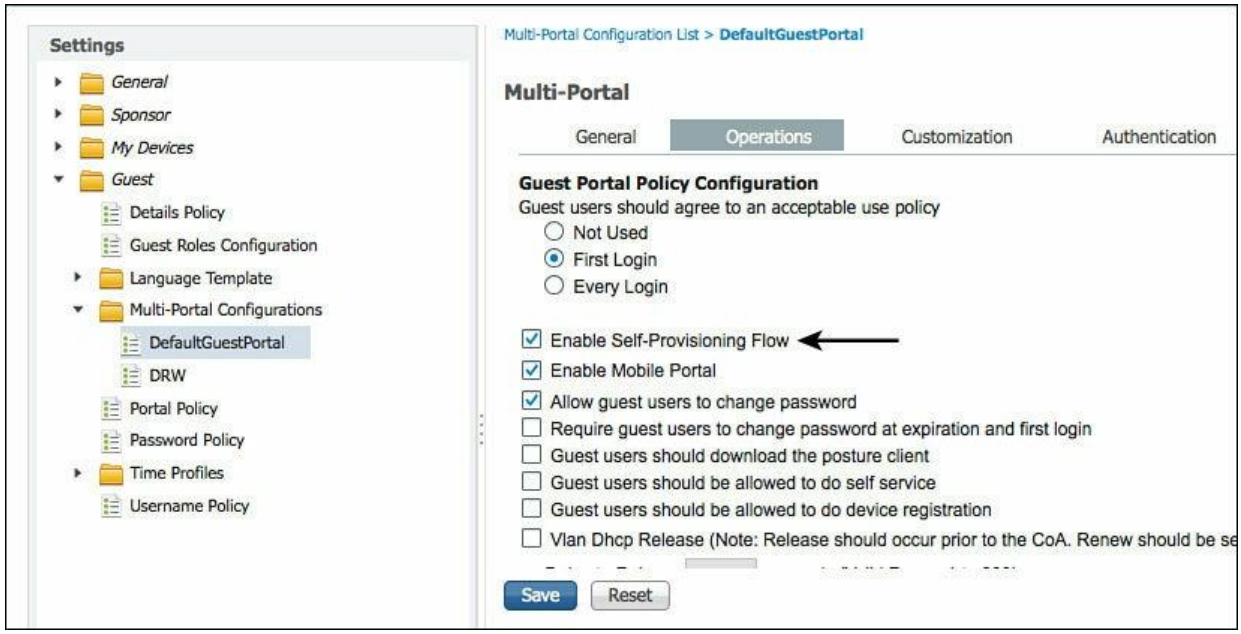


Figure 17-53 Web portal configuration.

Verifying Default Unavailable Client Provisioning Policy Action

ISE 1.2.x supports iOS, Android, Windows, and Mac OSX. However, it is quite possible for an end user to attempt access with a client that is not supported by ISE NSP (such as attempting with a BlackBerry or Windows mobile device). ISE offers two options for that situation:

- **Allow Network Access**—With this option, a user is allowed to register her device through the MyDevices Portal and gain network access without having to install and launch a native supplicant wizard. This assumes the user will have to interact and configure the supplicant in an out-of-band fashion. In other words, she must configure her supplicant herself. This option can be very attractive if the end users are capable of requesting and installing their own certificates.
- **Apply Defined Authorization Policy**—Basically, this option leave the client in the current state, which is a state of limited access. This is also the default setting.

Figure 17-54 shows the unavailable CPP action.

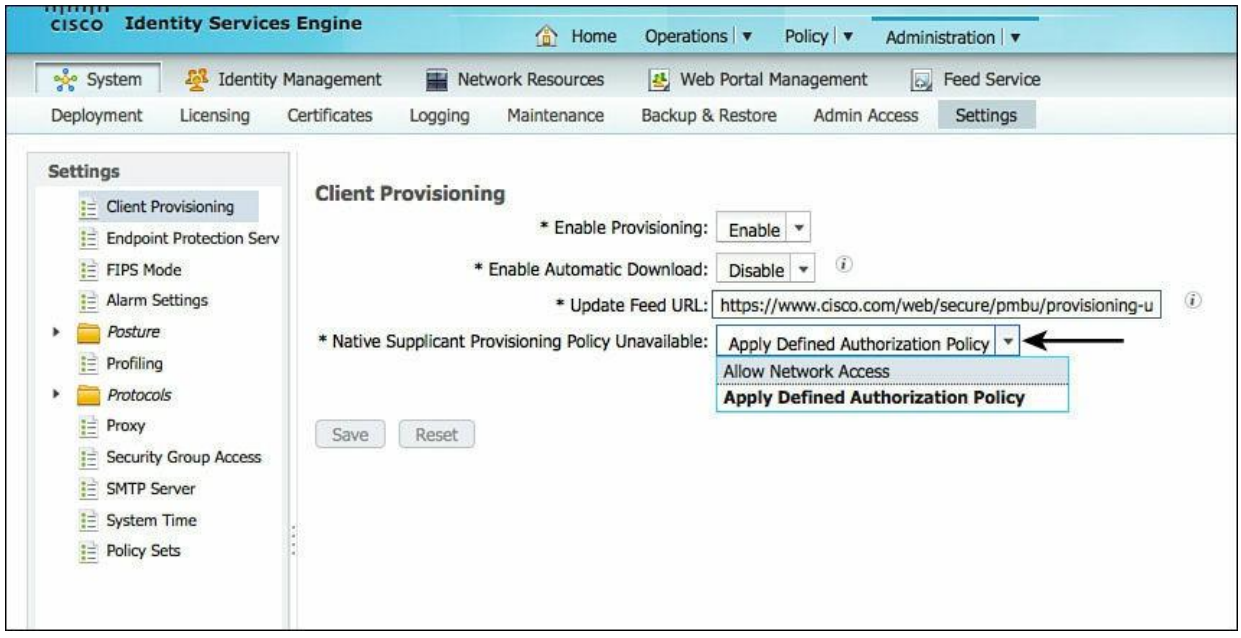


Figure 17-54 Default unavailable CPP action.

https://t.me/learningnets

Creating the Authorization Profiles

As described previously, you will create two authorization profiles. One of the authorization profiles is used for Android to accommodate a different ACL that permits Android devices to reach the Google Play Store. The second profile is used for all other applicable OSes that do not need to reach Google's Play Store. Do the following:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

Step 2. Add a new authorization profile, named **Android NSP**.

Step 3. Select **Web Redirection (CWA, DRW, MDM, NSP, CPP)**.

Step 4. From the drop-down, select **Native Supplicant Provisioning**.

Step 5. Enter the **Android-Marketplace** in the ACL field.

Step 6. Click **Submit**.

Figure 17-55 shows the completed Android NSP authorization profile.

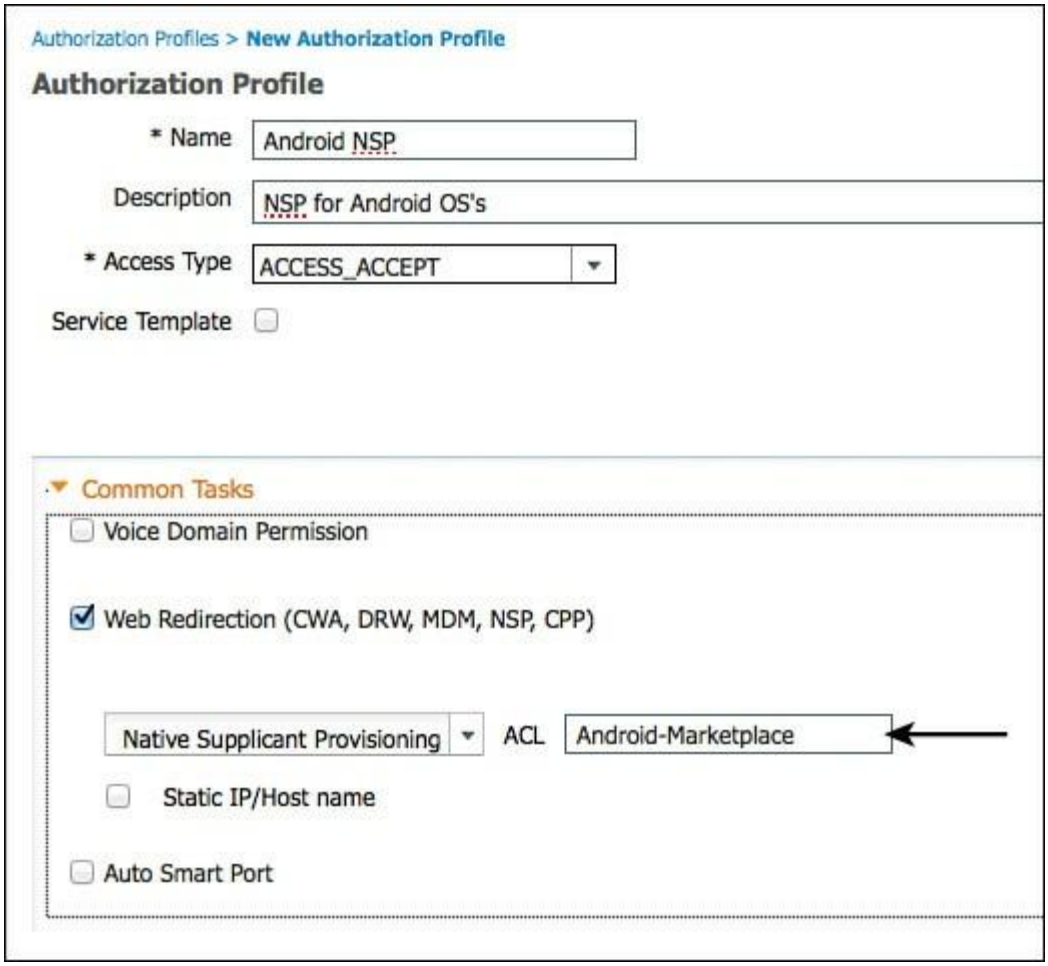


Figure 17-55 Android NSP authorization profile.

Step 7. Add a new authorization profile, named **NSP**.

Step 8. Select **Web Redirection (CWA, DRW, MDM, NSP, CPP)**.

Step 9. From the drop-down, select **Native Supplicant Provisioning**.

Step 10. Enter the **ACL-WEBAUTH-REDIRECT** in the ACL field.

Step 11. Click **Submit**.

Figure 17-56 shows the completed NSP authorization profile.

https://t.me/learningnets

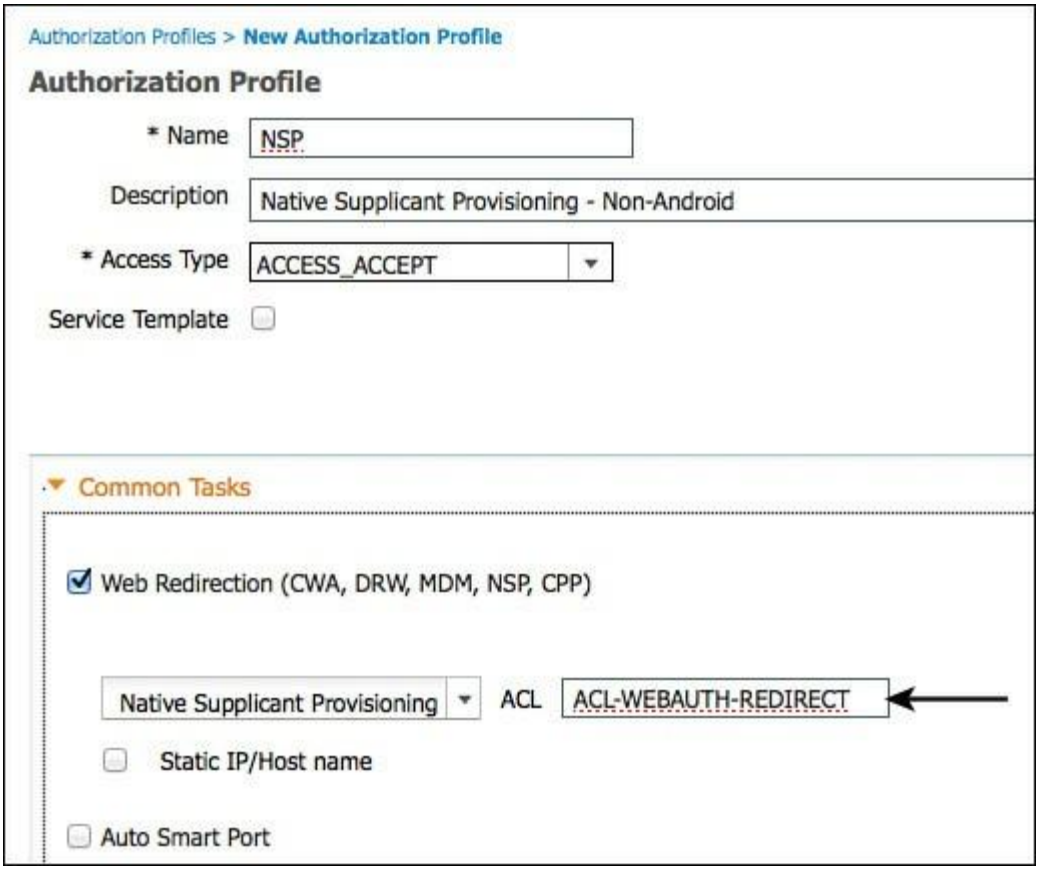


Figure 17-56 NSP authorization profile.

Creating the Authorization Rules for Onboarding

Now that the authorization profiles are ready, you will create two different authorization rules. The first rule is used for Android, and the second rule is used for the other devices. Do the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Add a new rule, just above the default rule.

Step 3. Name the rule **Android NSP**.

Step 4. Set the conditions as:

Click here to view code image

```
Network Access:AuthenticationMethod EQUALS MSCHAPV2
AND
Session:Device-OS EQUALS Android
```

Step 5. Set the authorization result to be the **Android NSP** created earlier.

Step 6. Click **Done**.

Step 7. Click Save.

Figure 17-57 shows the Android NSP conditions. You optionally can save the conditions individually or together into the library for reuse in other rules.

https://t.me/learningnets

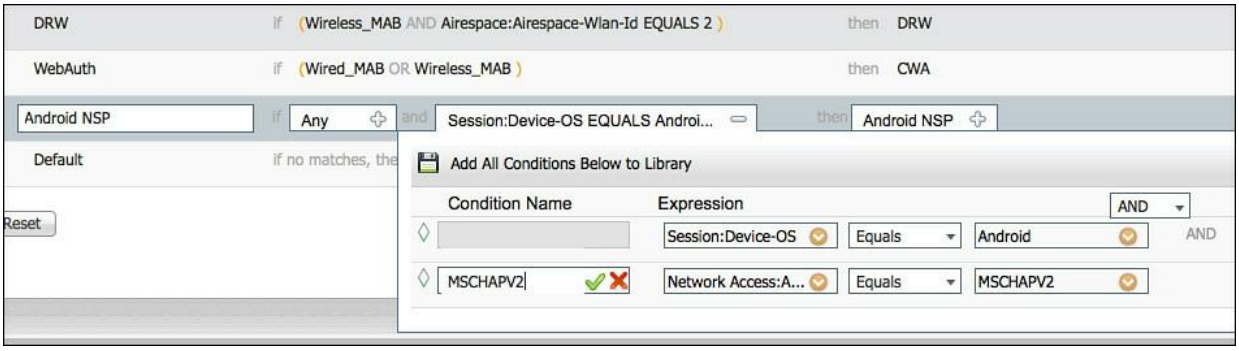


Figure 17-57 Android NSP authorization conditions.

Step 8. Add the NSP rule for non-Android.

Step 9. Add a new rule, just above the default rule and below the Android rule.

Step 10. Name the rule **NSP**.

Step 11. Set the conditions as:

Click here to view code image

```
Network Access:AuthenticationMethod EQUALS MSCHAPV2
```

Step 12. Set the authorization result to be the **NSP** authorization profile created earlier.

Step 13. Click **Done**.

Step 14. Click **Save**.

Figure 17-58 shows the NSP condition.

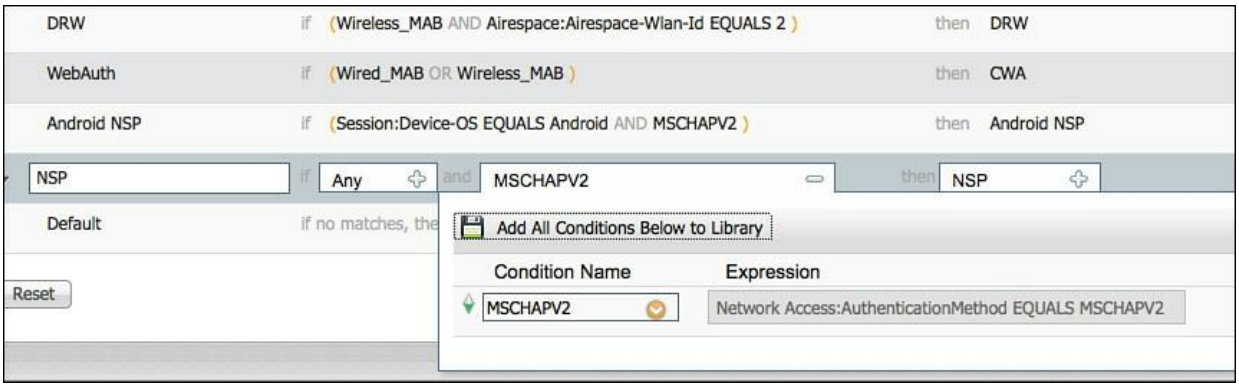


Figure 17-58 NSP authorization conditions.

Creating the Authorization Rules for the EAP-TLS Authentications

You have just created two different authorization rules that will result in the onboarding of the endpoints. After these endpoints are onboarded, they will be authenticating with EAP-TLS and certificates.

You created an authorization rule in Chapter 16, "Certificate-Based Authentications," that authorizes employees using certificates. That rule is named Employee Limited_Certificate. There is no reason this same rule cannot be used because it will certainly match any employee onboarded device that is authenticating with a certificate.

However, you are going to add a new rule that looks for certificate-based authentications and is also ensuring that the MAC address of the endpoint matches the MAC address burned into the certificate during the onboarding process. As an additional metric for the rule, it will also ensure that the endpoint has been through the BYOD registration process. This is a Cisco recommended configuration for BYOD with Cisco ISE.

From the ISE GUI, do the following:

https://t.me/learningnets

Step 1. Navigate to **Policy > Authorization**.

Step 2. Add a new rule, just above the **Employee SmartDevice** rule.

Step 3. Name the rule **Employee BYOD**.

Step 4. Set the conditions as:

Click here to view code image

```
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 ise.local/Users/Employees
```

Step 5. Set the authorization result to be **Employee Limited**.

Step 6. Click **Done**.

Step 7. Click **Save**.

Figure 17-59 shows the Employee BYOD conditions. As shown in that figure, you optionally can save the conditions individually or together into the library for reuse in other rules.

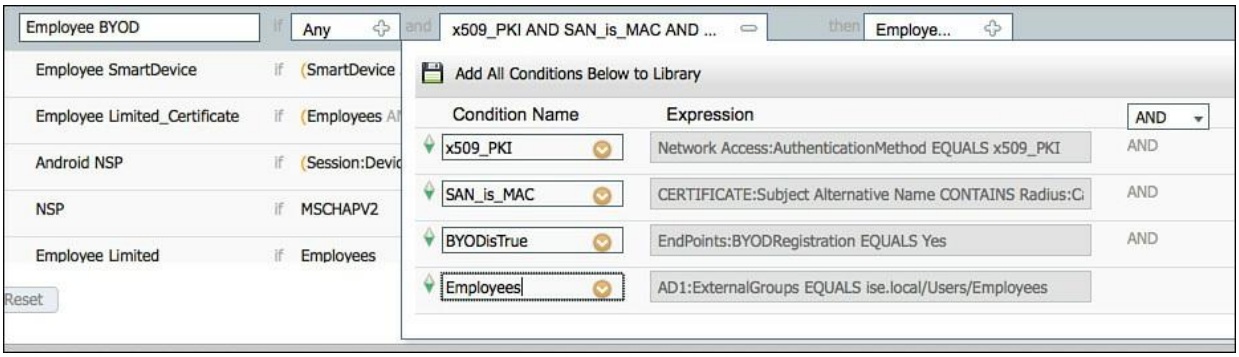


Figure 17-59 Employee BYOD conditions.

A composite screen shot of the authorization policy is shown in Figure 17-60.

https://t.me/learningnets

Rule Name	Conditions (identity groups and other conditions)	Permissions
Wireless Black List Default	if **Blacklist** AND Wireless_Access	then Blackhole_Wireless_Access
Profiled Cisco IP Phones	if **Cisco-IP-Phone**	then Cisco_IP_Phones
Profiled Non Cisco IP Phones	if Non_Cisco_Profiled_Phones	then Non_Cisco_IP_Phones
DRW Endpoints	if **DRW-Endpoints**	then Internet_Only
Guest	if **Guest** AND GuestFlow	then Internet_Only
Employee CWA	if (Employees AND GuestFlow)	then Internet_Only
Employee and CorpMachine	if EmployeeFullEAPChain	then Employee Full Access
Employee BYOD	if (x509_PKI AND SAN_is_MAC AND BYODisTrue AND Employees)	then Employee Limited
Employee SmartDevice	if (SmartDevice AND Employees)	then Internet_Only
Employee Limited_Certificate	if (Employees AND x509_PKI)	then PermitAccess
Android NSP	if (Session:Device-OS EQUALS Android AND MSCHAPV2)	then Android NSP
NSP	if MSCHAPV2	then NSP
Employee Limited	if Employees	then PermitAccess
DRW	if (Wireless_MAB AND Airespace:Airespace-Wlan-Id EQUALS 2)	then DRW
WebAuth	if (Wired_MAB OR Wireless_MAB)	then CWA
Default	if no matches, then PermitAccess	

Figure 17-60 Final authorization policy.

Configuring SCEP

ISE will act as a registration authority (RA) for the certificate enrollment and provisioning of the devices that are being onboarded. So, with ISE 1.2, ISE is not the certificate authority (CA); instead it is a "broker" of sorts that uses Simple Certificate Enrollment Protocol (SCEP) to request and provision a certificate to the client from the CA of your choice.



There are multiple CAs that could be used. The single most common CA deployed is the Microsoft Active Directory Certificate Authority. The Microsoft CA has support for SCEP and 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.

The authors have worked with companies who created a new Active Directory domain with only the one server in it, just for the CA and NDES functionality to be separate from their production AD. This is not necessarily recommended—we are just using the example to show the level of flexibility you might have, if needed. For more details on the configuration of a Microsoft CA, please see Appendix B, "Configuring the Microsoft CA for BYOD."

Some companies have a strict "no Microsoft" policy. The authors have worked with several companies that either have that policy or their AD team will not configure the CA services. Another fairly common CA to deploy for ISE is the Open Source Dogtag CA. There is also an Enterprise version of Dogtag, known as the Red Hat Enterprise CA. For more details on the configuration of a Dogtag CA, please see Appendix C, "Using the Dogtag CA for BYOD."

https://t.me/learningnets

> **Note**
>
> It's also good to note that ISE 1.3 provides a built-in CA. However, that is not applicable to this exam and therefore is out of scope for this book.

Although not all CAs have been tested, you theoretically can use any CA of your choosing, just as long as it meets these requirements:

- Must support SCEP
- Must support an automated or automatic issuing of the requested certificates
- Does not require a shared secret to be configured between the RA and the CA

The configuration of ISE to use the CA does not have too many steps. There is really just one setting to be made.

From the ISE GUI, do the following:

Step 1. Navigate to **Administration > System > Certificates**.

Step 2. Click **SCEP RA Profiles**.

Step 3. Click **Add**.

Step 4. Name the CA (our example is CP_AD_CA).

Step 5. Add the URL, such as http://ip-address-of-ca/certsrv/mscep/.

Step 6. Click **Test Connectivity**.

Step 7. Click **Submit**.

Figure 17-61 shows the completed SCEP RA profile.



Figure 17-61 Completed SCEP RA profile.

When you add a new SCEP profile to ISE, the CA's public key will automatically be imported into the Certificate Store and the Trust for Client Authentication or Secure Syslog Services setting will be enabled. This means that for any endpoints that are provisioned a certificate from the CA will be trusted. If you choose to use CRL or OCSP, you will need to edit this certificate and set that configuration. See Chapter 16 for more details around the Certificate Store and trust and revocation settings.

Figure 17-62 shows the certificate that was automatically imported and trusted.

https://t.me/learningnets

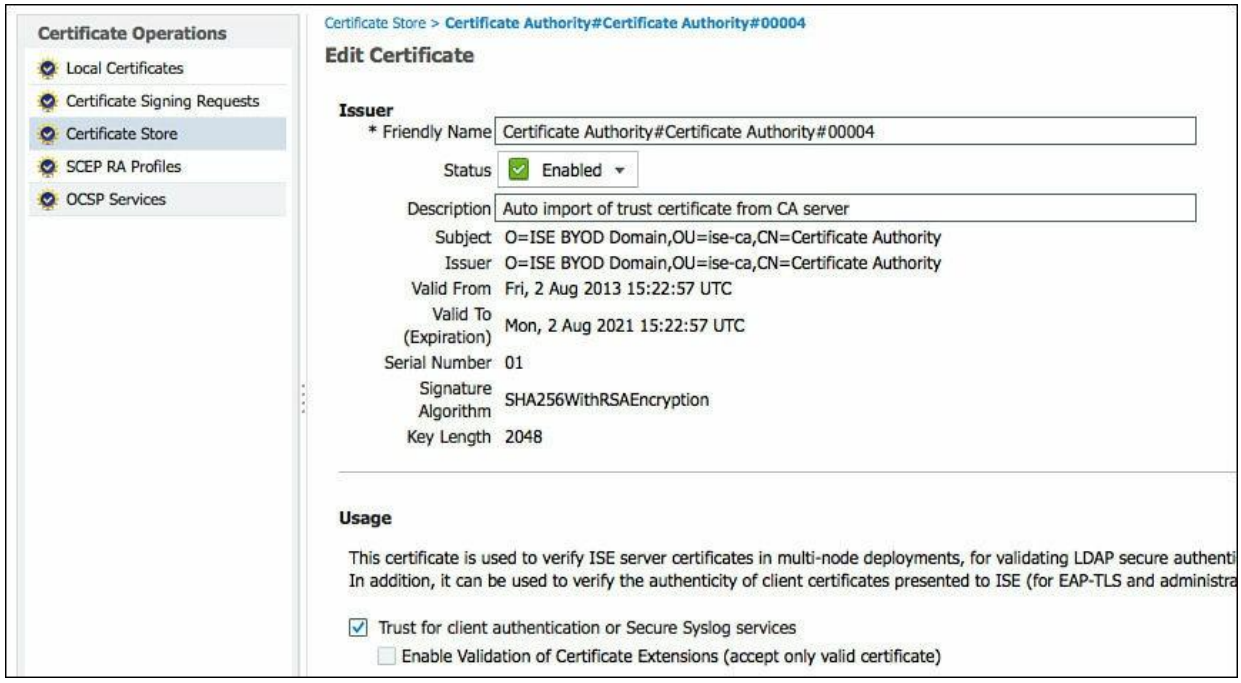


Figure 17-62 Automatically imported and trusted CA certificate.

Assuming the configuration on your CA was completed and accurate, your deployment is now ready to do BYOD onboarding. See Appendixes B and C for details on configuring Microsoft and Dogtag CAs, respectively. As you join the open network and authenticate or join the closed network and are directed to the onboarding process, you will experience first-hand the client experience outlined earlier in this chapter within the section titled "The End User Experience."

BYOD Onboarding Process Detailed

Yes, this chapter is getting very long. However, it is our hope that you will find all this information useful if you ever find yourself in a spot where you need to perform any troubleshooting of this process or need to describe it as part of this certification. You have seen that the user experience is quite simple and straightforward, but the process behind the scenes is a bit more complex.

iOS Onboarding Flow

We will examine, in detail, the experience with iOS endpoints and onboarding. To do so, we are focusing on a single-SSID onboarding experience.

Phase 1: Device Registration

The first phase of the onboarding is device registration. Remember that the overall process might have a lot of steps, but the end user should have to complete only four actions, as noted in Figure 17-63. However, we will take 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 his 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 will include a URL-Redirection to the NSP portal.

Step 5. The user opens his web browser, which is redirected to the NSP portal on ISE, displaying the device registration page with the Device ID field prepopulated with the MAC address (Action 3 for the end user experience).

https://t.me/learningnets

Step 6. The user clicks **Register** (Action 4 for the end user experience), which immediately 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 for OTA.

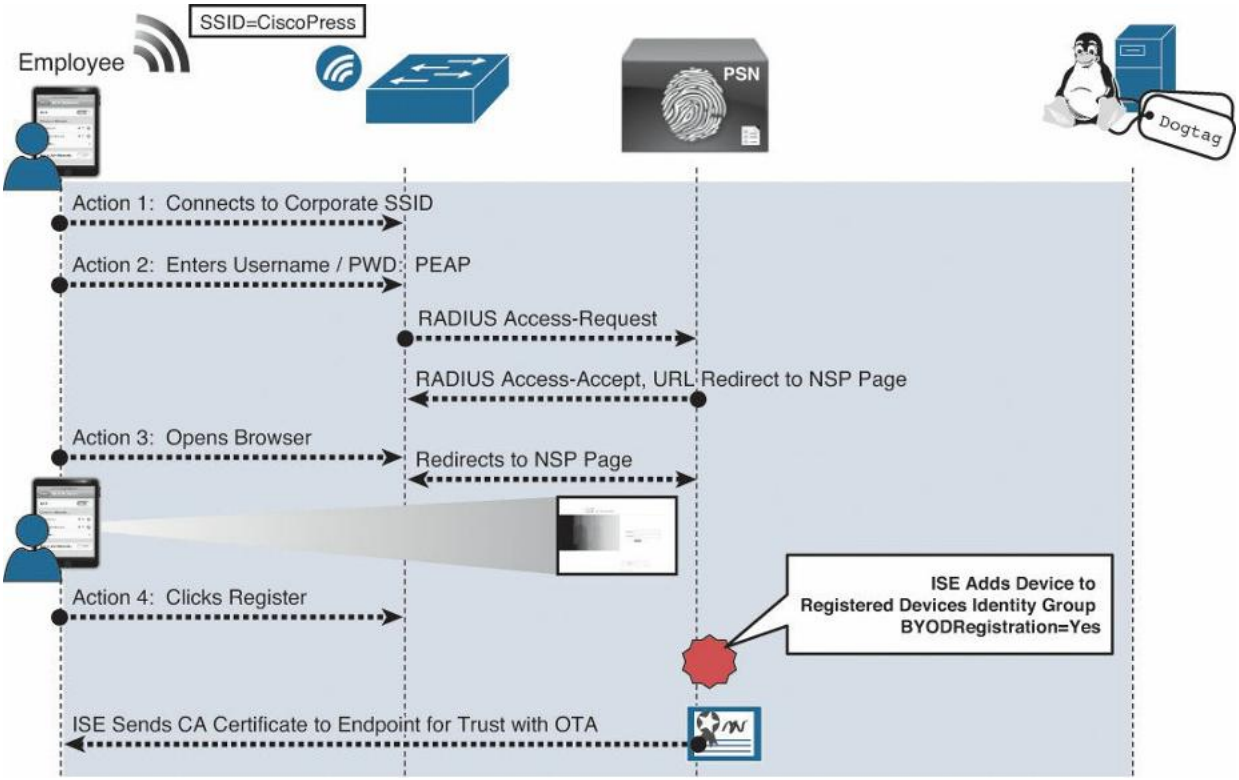


Figure 17-63 Phase 1: Device registration.

Figure 17-63 illustrates Phase 1.

Phase 2: Device Enrollment

This phase is where the device is enrolled in the BYOD system of ISE. The device itself will get its own certificate, although the authentication will only use the user certificate that is part of Phase 3.

Step 7. ISE sends a profile service to the iOS device via OTA.

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 uses SCEP to proxy the enrollment request to the CA.

Step 10. The CA automatically issues the certificate.

Step 11. The certificate is sent back to ISE, which sends it to the endpoint through the OTA service.

Figure 17-64 shows Phase 2.

https://t.me/learningnets

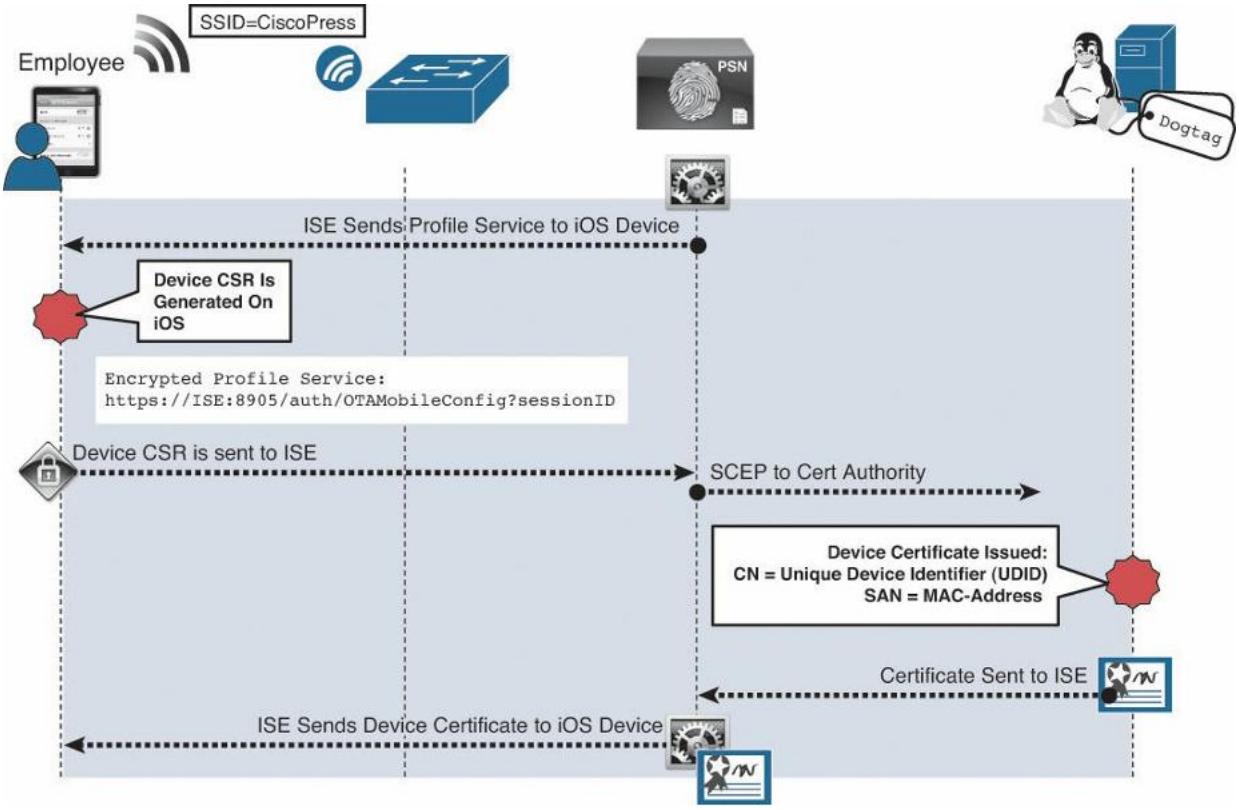


Figure 17-64 Phase 2: Device enrollment.

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. Additionally, this is the phase where the final authentication and authorization based on the user-certificate and EAP-TLS will occur.

Step 12. 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 SAN field:

- CN=Username
- SAN=MAC Address

Step 13. The user CSR is sent to ISE, which uses SCEP to proxy the certificate enrollment request to the certificate authority.

Step 14. The CA automatically issues the certificate.

Step 15. The certificate is sent back to ISE, which sends it to the device through the OTA service. Included in that profile is the Wi-Fi configuration, which details the SSID and to use EAP-TLS.

Step 16. ISE sends a CoA to the NAD of the type ReAuth, which causes a new authentication.

Step 17. The endpoint will now authenticate to the corporate SSID using the certificate via EAP-TLS.

Figure 17-65 shows Phase 3.

https://t.me/learningnets

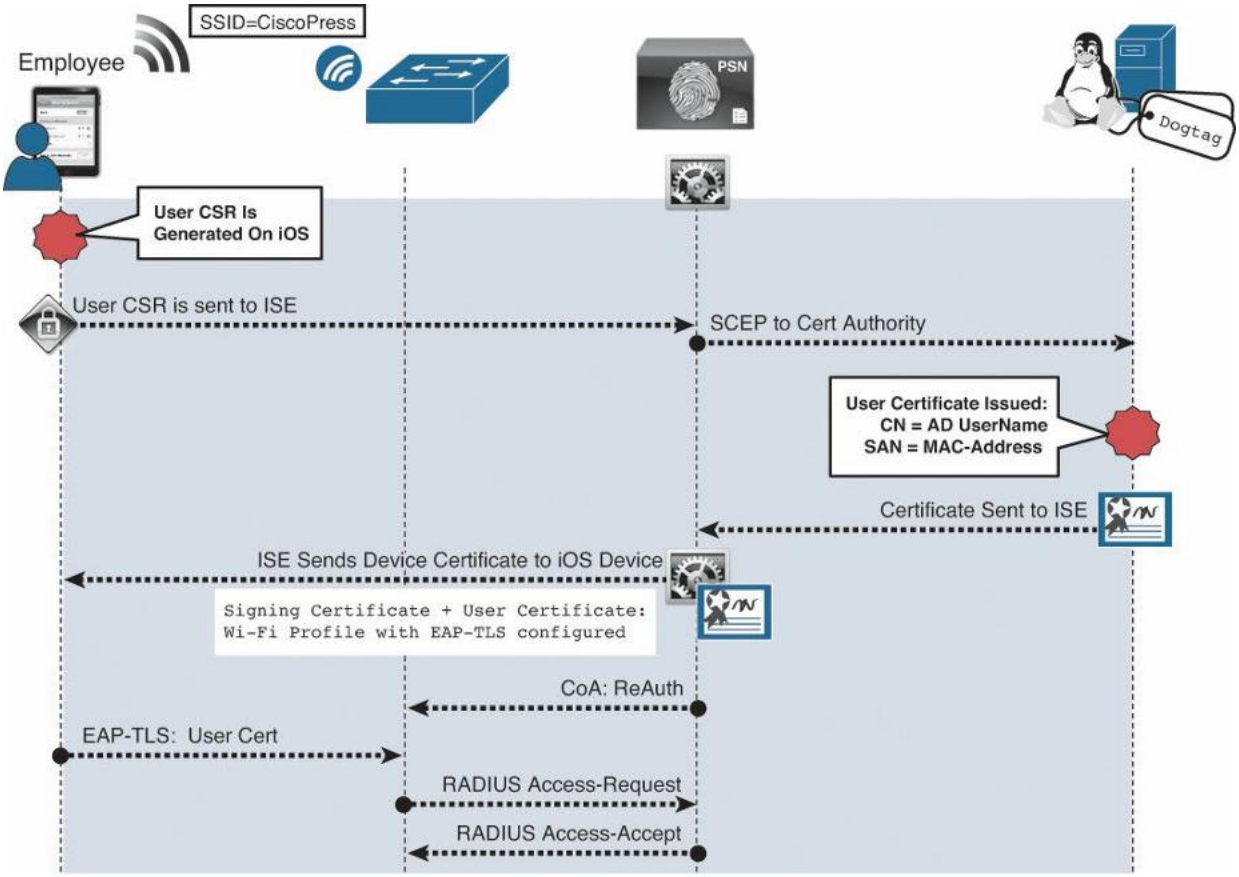


Figure 17-65 Phase 3: Device provisioning.

Android Flow

To detail the flow of onboarding with Android, we will use the dual-SSID approach. Android is certainly capable of doing 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.

Phase 1: Device Registration

The first phase of the onboarding is device registration. Remember that the overall process might have a lot of steps, but the end user should have to complete only six actions in total. There are four actions in this phase, as shown in Figure 17-66. More actions are required of the end user, but that is due to the Android requirement to use an app for the onboarding, versus the native OTA with iOS.

Step 1. 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.

Step 3. The authorization result from ISE will include 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 her AD username and password (Action 3 for the end user experience).

Step 6. The successful Web Auth triggers two events:

- The web page changes to the NSP portal.
- A CoA is sent to the WLC that 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.

https://t.me/learningnets

Step 8. The user clicks Register (Action 4 for the end user experience), which immediately 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 (this is a temporary attribute and only used for the provisioning process).
- Sends a CoA to the WLC to apply the correct ACL, allowing Google Marketplace access for the Android device.
- The web page sends the browser to the Google Play Store.

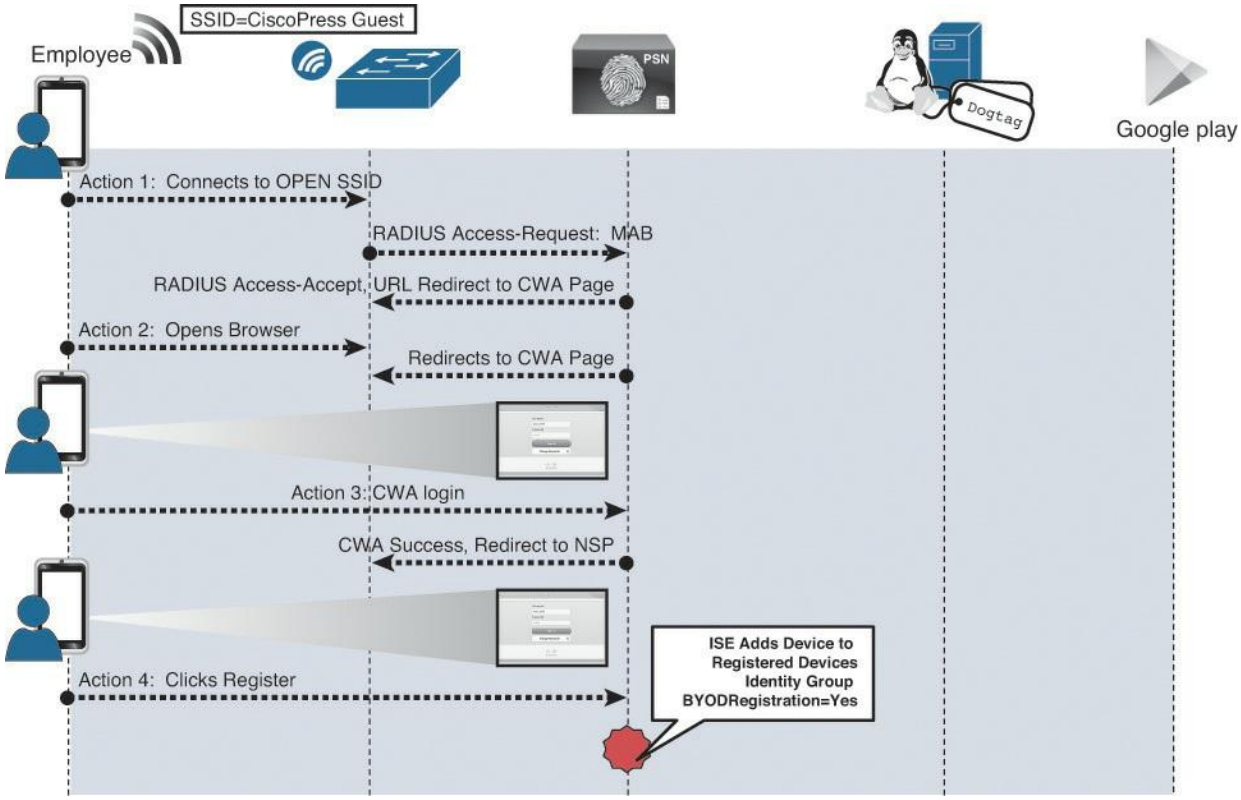


Figure 17-66 Phase 1: Device registration.

Figure 17-66 illustrates Phase 1 of the dual-SSID onboarding with Android.

Phase 2: Download SPW

This phase is where the end user is sent to the Google Play Store to download the NSA app. The device itself will get its own certificate, although the authentication will only use the user certificate that is part of Phase 3.

Step 9. The CoA from Phase 1 applies an ACL that permits traffic to the Google Play Store (the Android-Marketplace ACL).

Step 10. The browser was automatically sent to the Google Play Store, and the Android device prompts the user to choose the Internet browser or Play Store App to complete the request.

Step 11. The user may be prompted to log in to the Google Play Store.

Step 12. The user clicks to install the Cisco NSA app (Action 5 for the end user experience).

Figure 17-67 shows Phase 2 of the dual-SSID onboarding with Android.

https://t.me/learningnets

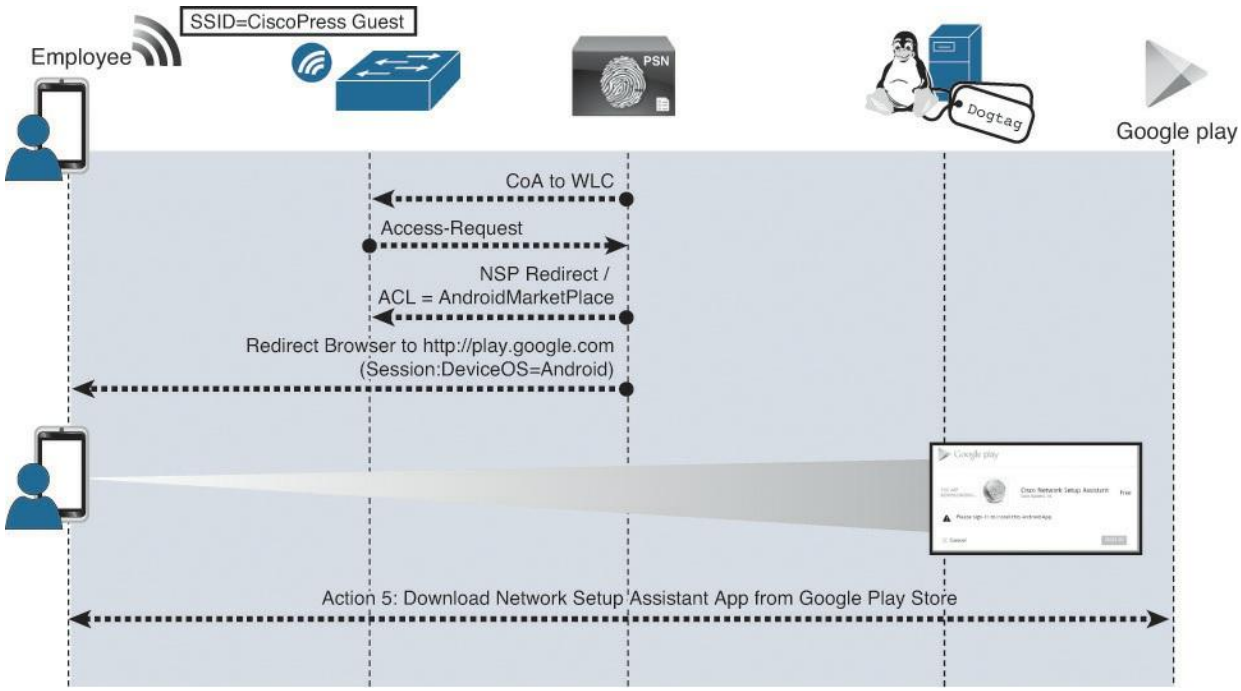


Figure 17-67 Phase 2: Download NSA.

Phase 3: Device Provisioning

This final phase is where the end user runs the NSA app, which handles the CSR generation and network profile installation, in a similar fashion to what iOS devices did natively. The sixth and final end user action occurs in this phase.

Step 13. The NSA app installs, and the user runs it (Action 6 for the end user experience).

Step 14. NSA sends a discovery message to http://default-gateway/auth/discovery.

Step 15. The WLC redirects that HTTP message to the ISE NSP portal based on the URL-REDIRECT result within the authorization from ISE.

Step 16. ISE sends the Android profile based on the CiscoPress-TLS supplicant NSP to the endpoint.

Step 17. NSA generates the CSR, using the employee's credentials as the certificate's subject and the MAC address as the SAN field:

- CN=Username
- SAN=MAC Address

Step 18. The CSR is sent to ISE, which uses SCEP to proxy the certificate enrollment request to the certificate authority.

Step 19. The CA automatically issues the certificate.

Step 20. The certificate is sent back to ISE, which sends it to NSA app. Included in that profile is the Wi-Fi configuration, which details the SSID and to use EAP-TLS.

Step 21. The NSA app connects the endpoint to the corporate SSID.

Step 22. The endpoint will now authenticate to the corporate SSID using the certificate via EAP-TLS.

Figure 17-68 shows the third and final phase of the dual-SSID onboarding with Android.

https://t.me/learningnets

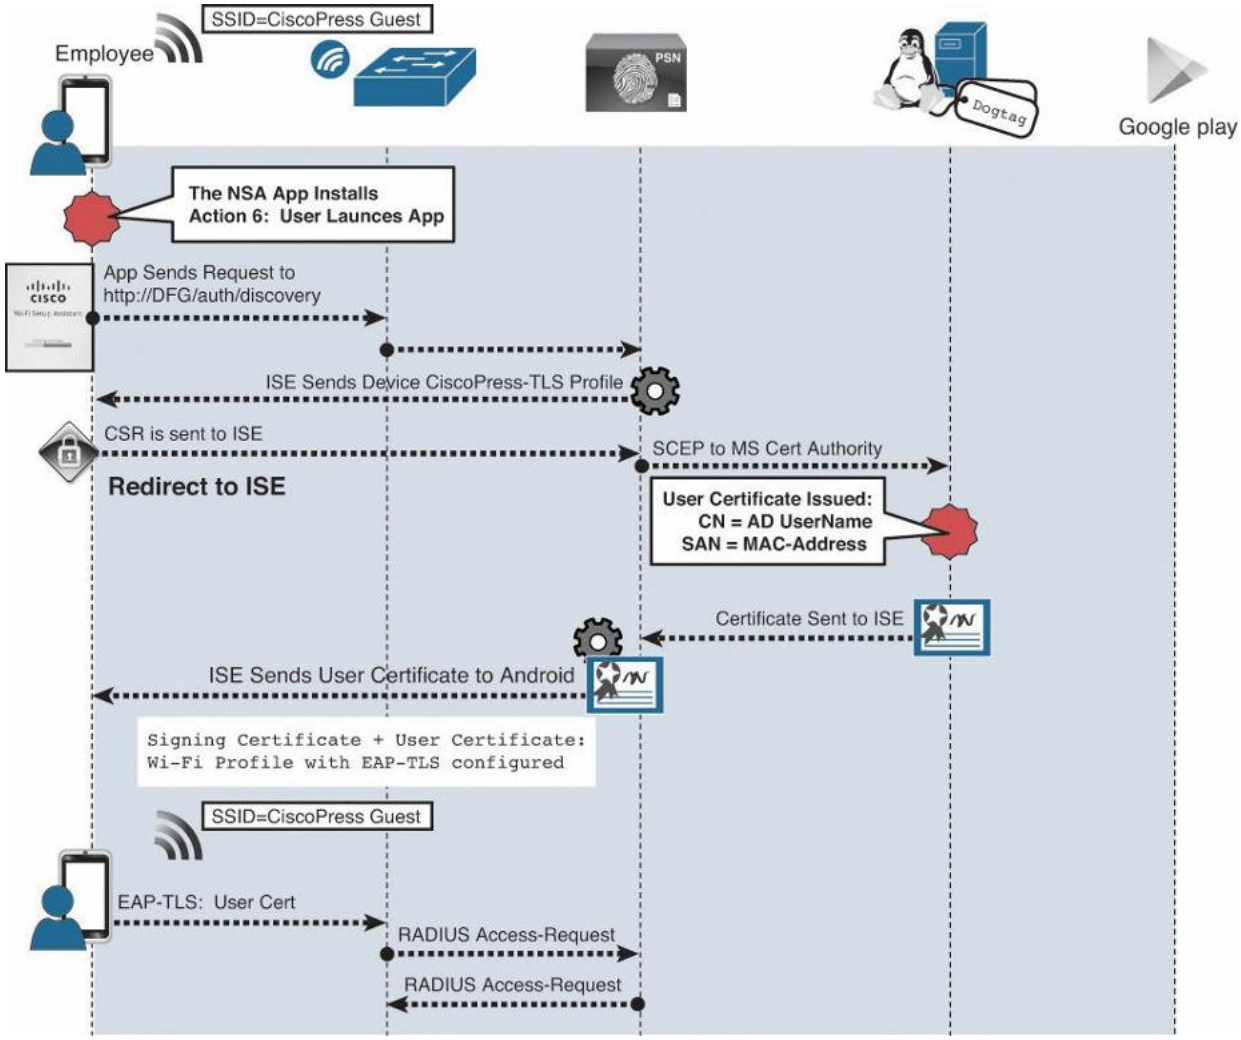


Figure 17-68 Phase 3: Device provisioning.

Windows and Mac OSX Flow

Mac OSX and Windows both use a wizard to accomplish the onboarding and provisioning. It is a Java-based applet called the Cisco Native Supplicant Provisioning Wizard. The wizard takes care of triggering the CSR from the operating system and installing the supplicant profile.

Unlike the mobile operating systems, this is only a two-phase process, whereas the mobile devices required three phases. We will use the single-SSID onboarding to demonstrate this flow.

There are only five actions for the end user, in a perfect world. However, we all know this is not a perfect world. Since the initial creation of the supplicant provisioning wizards with ISE 1.1.1, security vulnerabilities in client-side Java have been under attack. In response, the security levels in the client-side Java have been tightened, making the end user experience increasingly challenging, with many security warning prompts and workarounds.

In response, Cisco has created native applications for both Mac OSX and Windows operating systems that don't use the client-side Java. These clients became available with ISE 1.3 and were back-ported to ISE 1.2 patch 11. Neither of those versions is pertinent to this exam, which is why this book will still focus on the Java-based versions.

https://t.me/learningnets

Phase 1: Device Registration

Phase 1 is much like the same phase within the mobile device sections. The user joins the network in the same fashion, enters his credentials, and is redirected to download a piece of software that will do the onboarding for the operating system. It's very similar to the flows with Android, with one major difference: There is no requirement for access to an app store because the software is able to be hosted directly on ISE.

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 his username and password.

Step 2. The user enters his 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 his web browser (Action 3 for the end user experience), which is redirected to the NSP portal on ISE, displaying 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 immediately triggers two events:

- Sets the BYODRegistration flag for the endpoint identity to Yes.
- Adds the endpoint to the RegisteredDevices identity group.

Figure 17-69 illustrates the device registration phase.

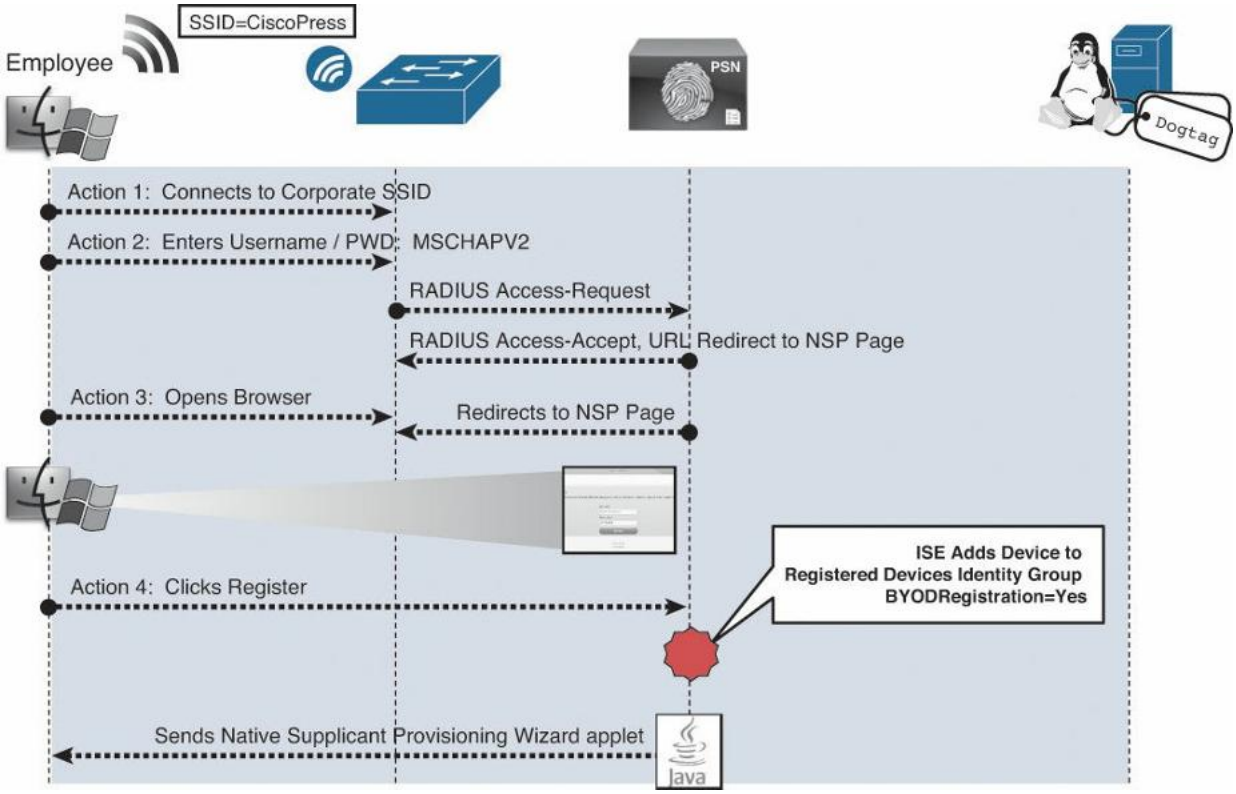


Figure 17-69 Phase 1: Device registration.

Phase 2: Device Provisioning

Phase 2 is the final phase for the Mac OSX and Windows onboarding and is much like the final phase within the Android sections. The Java applet is downloaded and executes, performing the CSR and provisioning for the endpoint. There should not be any user action.

Step 7. The Native Supplicant Provisioning Wizard is downloaded and runs.

https://t.me/learningnets

Step 8. The user clicks Start (Action 5 for the end user experience).

Step 9. The NSP Wizard sends a discovery message to http://default-gateway/auth/discovery.

Step 10. The WLC redirects that HTTP message to the ISE NSP portal based on the URL-REDIRECT result within the authorization from ISE.

Step 11. ISE sends the NSP profile based on the CiscoPress-TLS supplicant profile to the endpoint.

Step 12. NSP Wizard generates the CSR, using the employee's credentials as the certificate's subject and the MAC address as the SAN field:

- CN=Username
- SAN=MAC Address

Step 13. The CSR is sent to ISE, which uses SCEP to proxy the certificate enrollment request to the certificate authority.

Step 14. The CA automatically issues the certificate.

Step 15. The certificate is sent back to ISE and is sent down to the NSP Wizard. Included in that profile is the Wi-Fi configuration, which details the SSID and to use EAP-TLS.

Step 16. ISE sends a CoA to the NAD of the type ReAuth, which also causes a new authentication. (If this was a dual-SSID situation, the applet would automatically connect the endpoint to the corporate SSID at this point.)

Step 17. The endpoint will now authenticate to the corporate SSID using the certificate via EAP-TLS.

Figure 17-70 illustrates the device provisioning phase.

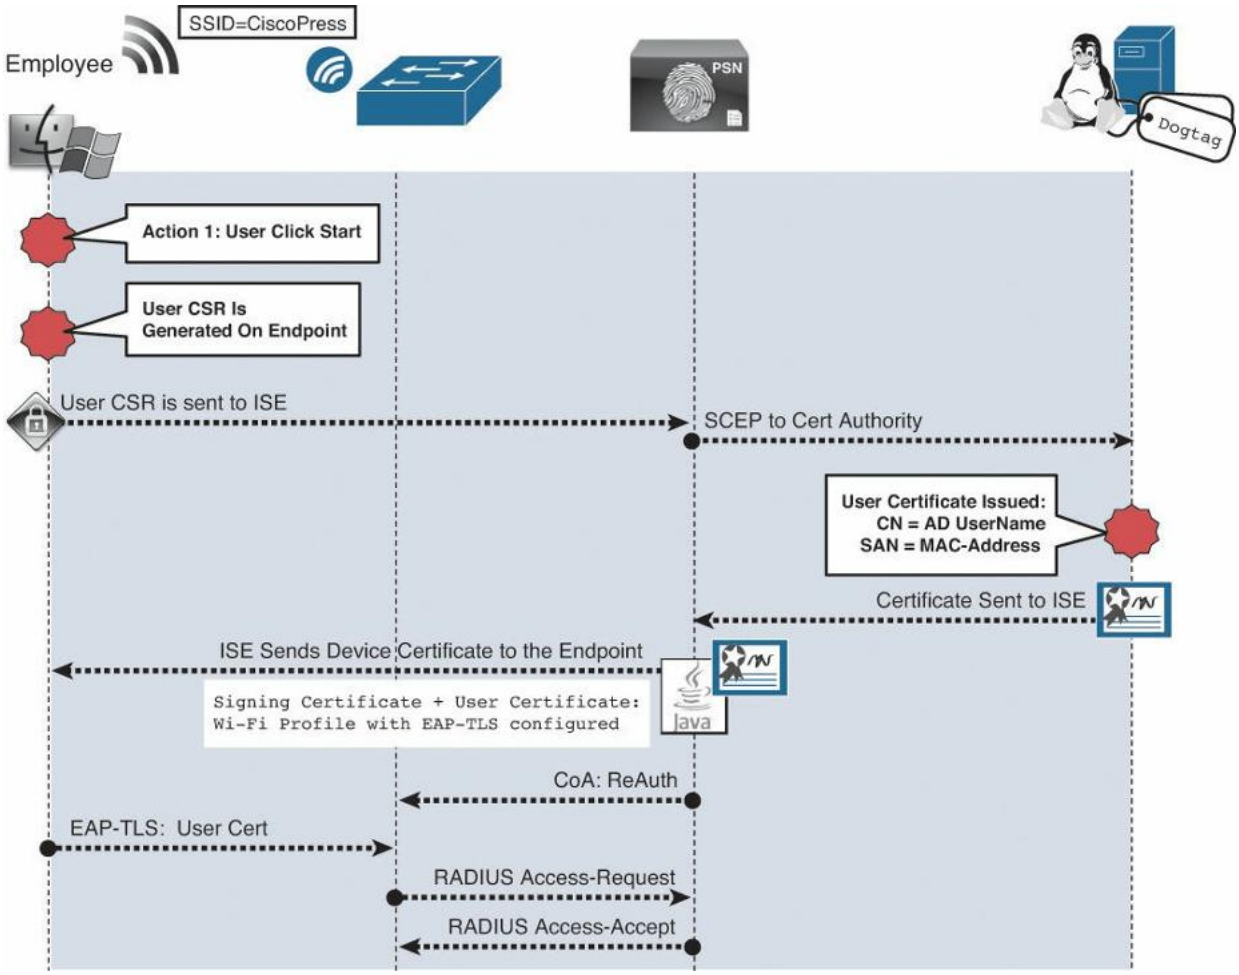


Figure 17-70 Phase 2: Device provisioning.

https://t.me/learningnets

Verifying BYOD Flows

In this section we will examine a few ways to verify that BYOD flows are operating successfully. This topic is covered in more depth in Chapter 22, "Troubleshooting Tools," but this section will get you started.

There are numerous locations that an ISE admin should look to see whether something is functioning correctly. The first is Live Log, but there are also reports to examine. Plus, you can take a look in the endpoint identities database or even MyDevices portal. The MyDevices portal is covered in more detail in the "Self Management" section of this chapter.

Live Log

As described numerous times throughout this book, the first location any ISE admin should always look to verify functionality is Live Log. As it relates to single-SSID onboarding, one can see at a glance the various authentications of an endpoint and the authorization result of each.

Figure 17-71 shows a screen capture of Live Log during a single-SSID onboarding flow.



Figure 17-71 Live Log showing the single-SSID.

Examining Figure 17-71, you should see three logged events that have been pointed out. From the earliest (bottom) to the latest (top):

- You see the initial authentication that uses PEAP(EAP-MSCHAPv2) for authentication and is given the NSP authorization profile. This is the authorization profile you created to redirect the user to the provisioning portal for device registration and provisioning.
- The next event is a Dynamic Authorization Success message. This is the successful CoA that tells the WLC to reauthenticate the session.
- Finally, 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 logical groupings such as Auth Services, Deployment Status, Endpoints and Users, and Security Group Access (also called TrustSec). There is even a dedicated report to supplicant provisioning.

From the ISE GUI, do the following:

Step 1. Navigate to **Operations > Reports > Endpoints and Users > Supplicant Provisioning**.

Step 2. Select a time range, such as Today or **Last 7 Days**.

Step 3. Click **Run**.

Figure 17-72 shows the output of this report.

https://t.me/learningnets

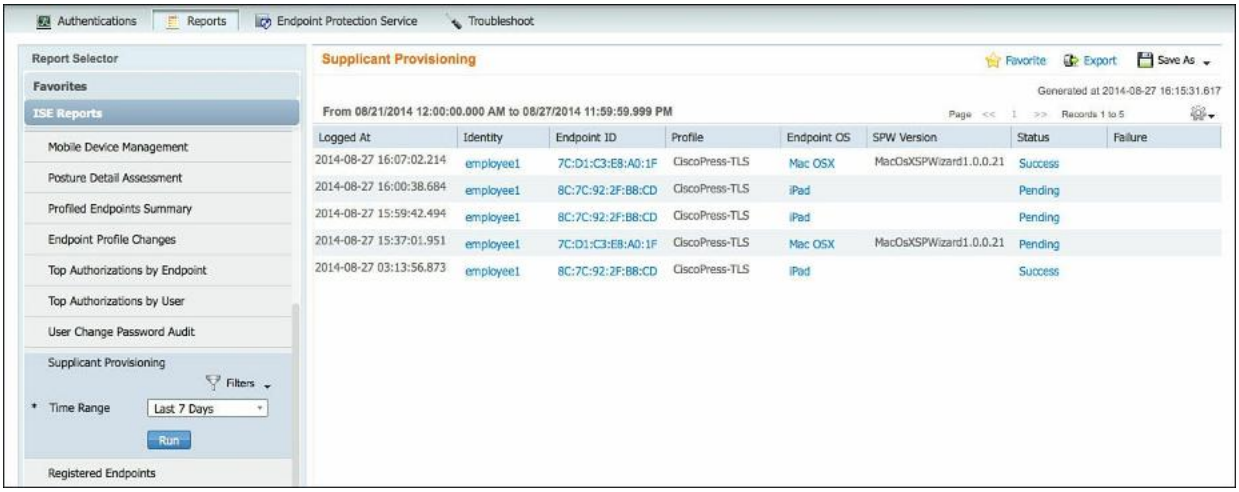


Figure 17-72 Supplicant provisioning report.

Identities

Certainly, another way to verify the onboarding is to examine the endpoint in the Endpoints database.

From the ISE GUI, do the following:

Step 1. Navigate to **Administration > Identity Management > Identities > Endpoints**.

Step 2. Optionally, apply a **Quick Filter** to make the list more manageable, as shown in Figure 17-73.

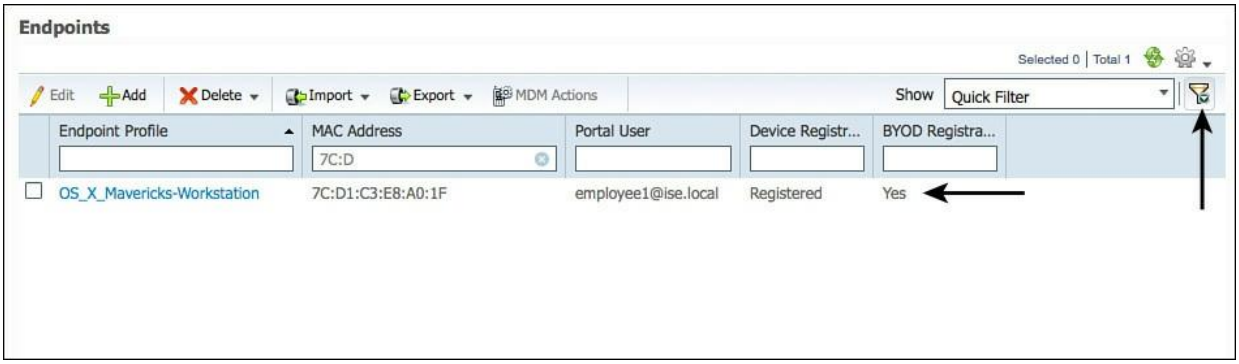


Figure 17-73 Identities > Endpoints.

Step 3. Examine the BYOD Registration Status and Portal User fields.

Figure 17-73 shows the endpoints database with a Quick Filter enabled.

MDM Onboarding

Many organizations use Mobile Device Management solutions. These solutions provide endpoint management for a plethora of devices. They help enforce specific security requirements such as endpoint encryption, PIN lock, jail-break detection, remote wipe capabilities, and more. Many MDMs will even provision supplicants and certificates to devices as part of their management package.

Typically, a user would bring in a mobile device and, to gain access to the network, she had to call the help desk to receive instructions on how to onboard the device with the MDM so she could gain access to the network. There were some significant downsides to this process, such as:

- 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 that solution.
- An MDM license was required for every device the organization would provision and allow to have network access, which could be cost prohibitive.

However, this presented a very beneficial and strategic opportunity for Cisco and these MDM vendors. The

https://t.me/learningnets

vendors had the mobile device management capabilities, and Cisco had the onboarding, network access policy, and enforcement mechanisms.

ISE 1.2 adds the integration of many of the industry's top MDM vendors. The list is inclusive of VMware AirWatch, Mobile Iron, Citrix ZenPrise, Good Technologies, Fiberlink, SAP Afaria, Symantec, and Tangoe—and the list continues to grow. These vendors have implemented an Application Programming Interface (API) written by Cisco to enable scalable bidirectional communication between their solution and ISE. By utilizing the same API for all vendors instead of a custom integration with each vendor, Cisco can ensure consistency and quality.

Integration Points

The API provides ISE with the ability to use MDM attributes in the authorization policies. The authorization can use a macro-level attribute stating that the device is in compliance with the MDM policy or a micro-level attributes such as jail-break status, PIN lock, or even endpoint encryption.

Table 17-2 lists the MDM attributes and what their possible values are.

https://t.me/learningnets

MDM Attribute	Definition	Possible Values
DeviceRegisterStatus	Is the device registered with the MDM?	■ Unregistered ■ Registered
DeviceComplianceStatus	Macro-level attribute stating whether the device meets the security policy of the MDM.	■ NonCompliant ■ Compliant
DiskEncryptionStatus	Is encryption enabled on the storage of the device?	■ On ■ Off
PinLockStatus	Does the device have an automatic lock, requiring a PIN or password to unlock the device?	■ On ■ Off
JailBrokenStatus	Has the device been jail-broken?	■ Unbroken ■ Broken
Manufacturer	Who is the manufacturer of the device?	Text field, or can be compared to attribute from AD/LDAP
Model	What is the model of the device?	Text field, or can be compared to attribute from AD/LDAP
IMEI	What is the unique ID?	Text field, or can be compared to attribute from AD/LDAP
SerialNumber	What is the device's serial number?	Text field, or can be compared to attribute from AD/LDAP
OSVersion	What is the version of the operating system?	Text field, or can be compared to attribute from AD/LDAP
PhoneNumber	What is thephone number?	Text field, or can be compared to attribute from AD/LDAP

Table 17-2 MDM Attributes and Values

Configuring MDM Integration

Before you can configure ISE to communicate with the MDM, ISE must trust the certificate of the MDM for the SSL-encrypted communications.

From the ISE GUI, do the following:

Step 1. Navigate to **Administration > System > Certificates**.

Step 2. Select **Certificate Store**.

Step 3. Import the certificate of the MDM as a trusted certificate.

https://t.me/learningnets

> **Note**
>
> If the MDM will be provisioning certificates to endpoints with their own CAs, instead of using the BYOD onboarding with ISE and your CA, then you will also need to trust the certificate for client authentication. Additionally, you might need to import ISE's public certificate into the MDM so the MDM will also trust ISE.

Now that the certificate is trusted, we can add the MDM Server to ISE. ISE 1.2 can be configured with the knowledge of many MDMs, but only one can be active at a time.

Step 4. Navigate to **Administration > Network Resources > MDM**.

Step 5. Click **Add**.

Step 6. Input a name for the connection to the MDM.

Step 7. Input the hostname of the server (should be DNS resolvable).

Step 8. The port should be 443, unless otherwise instructed by your MDM vendor.

Step 9. An instance name is used in some cases when the vendor is multitenant capable (such as a cloud-based MDM solution).

Step 10. Use the **Administrator Name** that the API will connect with to enroll all the mobile devices.

Step 11. Enter the **Password** for the API account.

Step 12. Add an optional description.

Step 13. Click **Test Connection**.

Step 14. Click **Submit**.

Figure 17-74 shows a sample MDM definition.

https://t.me/learningnets



Figure 17-74 Add MDM.

Configuring MDM Onboarding Rules

Now that the MDM has been added to ISE, the dictionaries for MDM attributes have been added for use in ISE authorization policies. So it is now 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, only this time 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 would permit final access. Some organizations would not want to send all devices to the MDM but would want only specific devices to be included. One such 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 option would be to use a logical profile that includes the relevant mobile device vendors and device types (that is, Apple or Android). In this second scenario, you can choose to send only mobile phones and/or tablet devices to the MDM for provisioning, allowing Windows and Mac OSX laptops to not leverage the MDM.

The example in this section will use those endpoint identity groups because it is something that the authors have seen used in production at a number of installs.

Creating the Authorization Profile

The first step is to create the authorization profile that redirects the endpoint to the MDM's provisioning portal for onboarding.

https://t.me/learningnets

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Policy Elements > Results > Authorization**.

Step 2. Select **Authorization Profiles**.

Step 3. Add a new authorization profile named **MDM Onboard**.

Step 4. The access type should be **ACCESS_ACCEPT**.

Step 5. Select **Web Redirection**.

Step 6. Select **MDM Redirect**.

Step 7. The Web Redirection ACL should reference an ACL that permits access to the MDM 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 8. Click **Submit**.

Figure 17-75 shows a sample authorization profile.

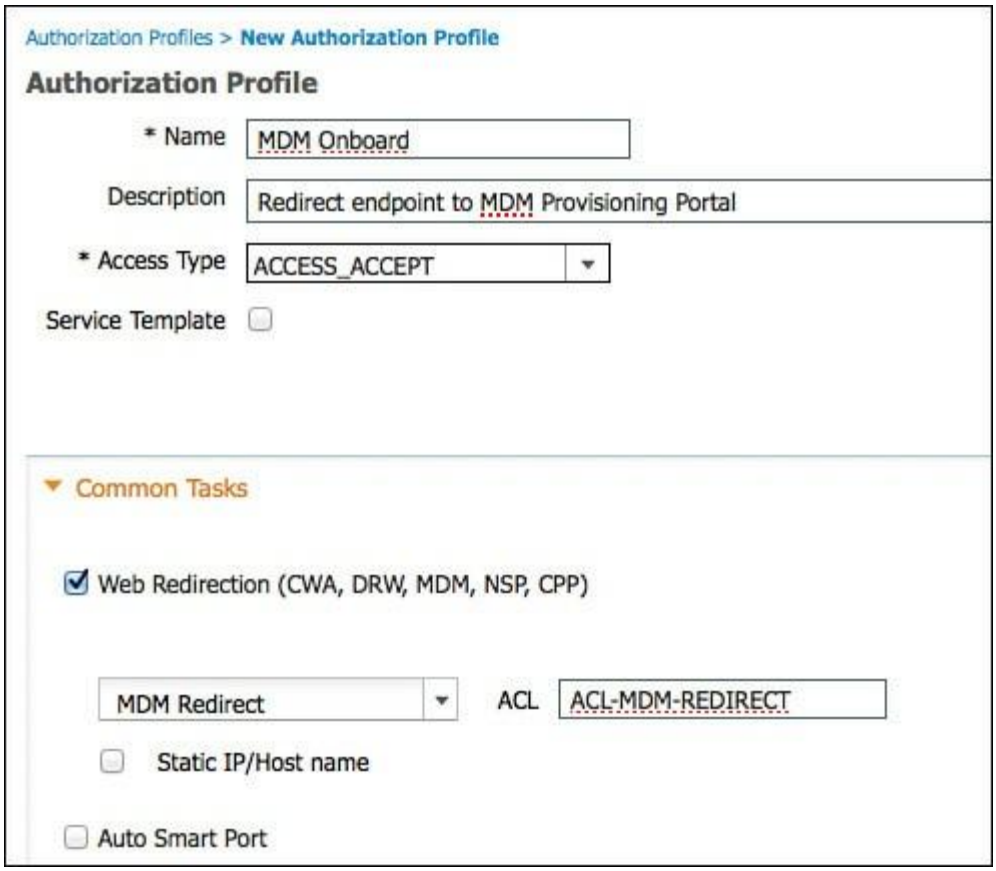


Figure 17-75 MDM onboard authorization profile.

Figures 17-76 and 17-77 show a sample ACL with DNS snooping.

https://t.me/learningnets

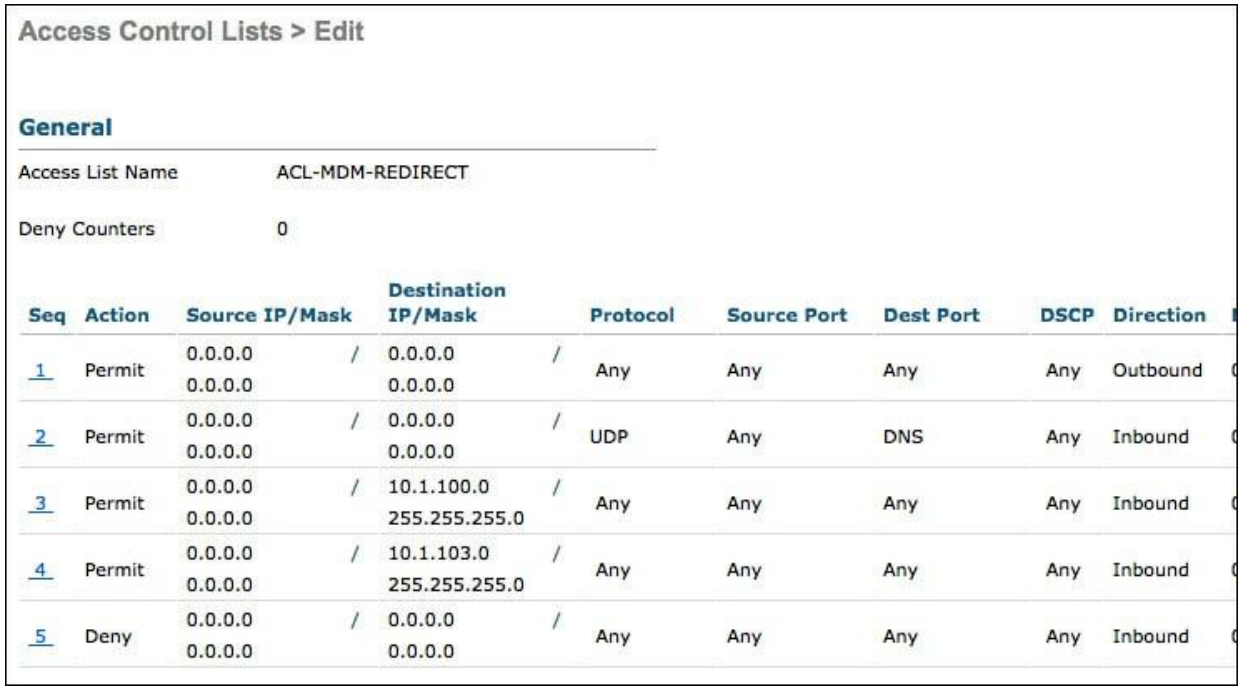


Figure 17-76 ACL-MDM-REDIRECT.



Figure 17-77 ACL-MDM-REDIRECT URLs for DNS snooping.

Creating the Authorization Rules

The authorization results are ready for the onboarding, so now you will create an authorization rule to send devices to the MDM vendor's provisioning portal for onboarding. As mentioned at the beginning of the "Configuring MDM Onboarding Rules" section, the placement of these rules is important. For this example, you will place the redirection rule above the Employee BYOD rule.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Insert the new rule above the Employee BYOD rule.

Step 3. Name the rule **MDM Onboard**.

Step 4. (Optional) Select an identity group that matches your corporate assets (**CorpAssets** in this example).

Step 5. Add the conditions as follows:

`Employees`

Click here to view code image

```
AND
x509_PKI
```

https://t.me/learningnets

```
AND
SAN_is_Mac
AND
Endpoints:BYODRegistration EQUALS Yes
AND
MDM:DeviceRegisterStatus EQUALS Unregistered
```

Step 6. Set the result to the **MDM Onboard** authorization profile.

Step 7. Click **Done**.

Figure 17-78 show the completed rule.

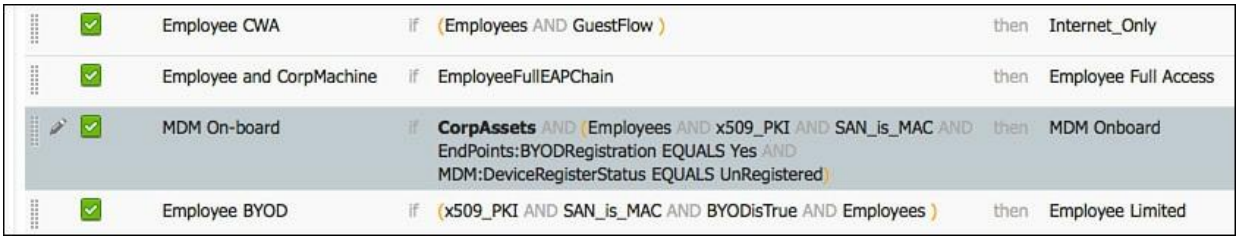


Figure 17-78 MDM onboard rule.

Add another rule below it that permits access to devices that are registered and meet the MDM compliance and that therefore should be provided with full access to the network.

Step 8. Navigate to **Policy > Authorization**.

Step 9. Click the down arrow for the MDM onboard rule and select **Duplicate Below**.

Step 10. Name the rule **MDM Permit**.

Step 11. Delete the BYODRegistration condition.

Step 12. Change the **MDM:DeviceRegisterStatus** to **EQUALS Registered**.

Step 13. Add a new condition for MDM:ComlianceStatus EQUALS Compliant:

The final condition set should look like this:

Click here to view code image

```
Employees
AND
x509_PKI
AND
SAN_is_Mac
AND
MDM:DeviceRegisterStatus EQUALS Registered
AND
MDM:DeviceComplianceStatus EQUAL Compliant
```

Step 14. Set the result to **Permit Access**.

Step 15. Click **Done**.

Step 16. Click **Save**.

Figure 17-79 shows the completed rule.

https://t.me/learningnets



Figure 17-79 MDM permit rule.

When using MDM attributes as part of the authorization policy, ISE will check with the MDM at every authorization. So, if the MDM is unavailable, the rule will never match. In addition, a bulk download of data from the MDM occurs every four hours, detailing endpoint status. If an endpoint is marked noncompliant during that download, a CoA is sent and the device is forced to reauthenticate, providing a different result (such as quarantine).

Managing Endpoints

Each user can manage the devices she has personally onboarded from the MyDevices portal, or the administrator can manage the devices from the endpoints screen. The options for the device are shown in Table 17-3.

Registered Device Option	Description
Lost	When marking a device as lost, it is added to the Blacklist Endpoint identity group and denied further access to the network until reinstated.
Reinstate	When the device is reinstated, it is moved from the Blacklist identity group to the RegisteredDevices identity group, so it can be permitted access to the network again.
Delete	The endpoint is removed from the Endpoint Directory.
Full Wipe	An MDM action to wipe the entire endpoint.
Corporate Wipe	An MDM action to remove only the containerized data on the endpoint considered “corporate”.
PIN Lock	An MDM action to initiate a remote lock on the device.

Table 17-3 MyDevices Registered Device Options

Self Management

The MyDevices portal is where end users can manage any endpoints that they have registered and register new endpoints.

An employee can log in to the MyDevices portal by navigating to either the default URL of https://[ISE_fqdn]:8443/mydevices/ or the friendly URL configured at Administration > Web Portal Management > Settings > General > Ports, as shown in Figure 17-80.

https://t.me/learningnets

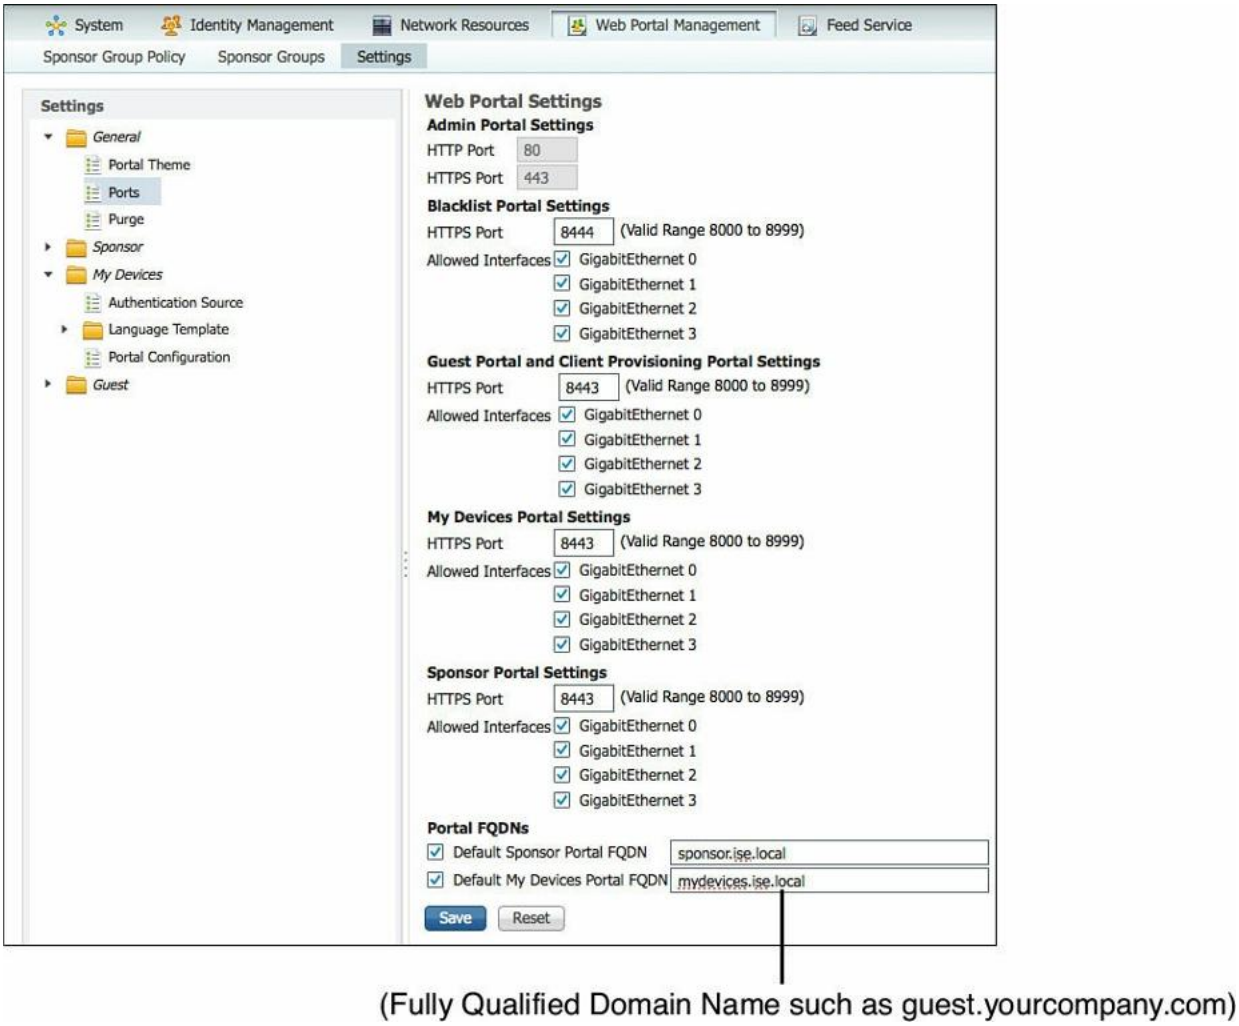


Figure 17-80 Settings > General > Ports.

Figure 17-80 shows the Ports page, where the friendly URL is configured.

By default, the MyDevices portal uses an identity source sequence (ISS) called MyDevices_Portal_Sequence, which includes only internal users by default in 1.2. To allow employees to log in successfully, you need to add Active Directory to the ISS.

From the ISE GUI, navigate to **Administration > Identity Management > Identity Source Sequences**. Select the **MyDevices_Portal_Sequence**, and add the AD1 identity source, as shown in Figure 17-81.

https://t.me/learningnets

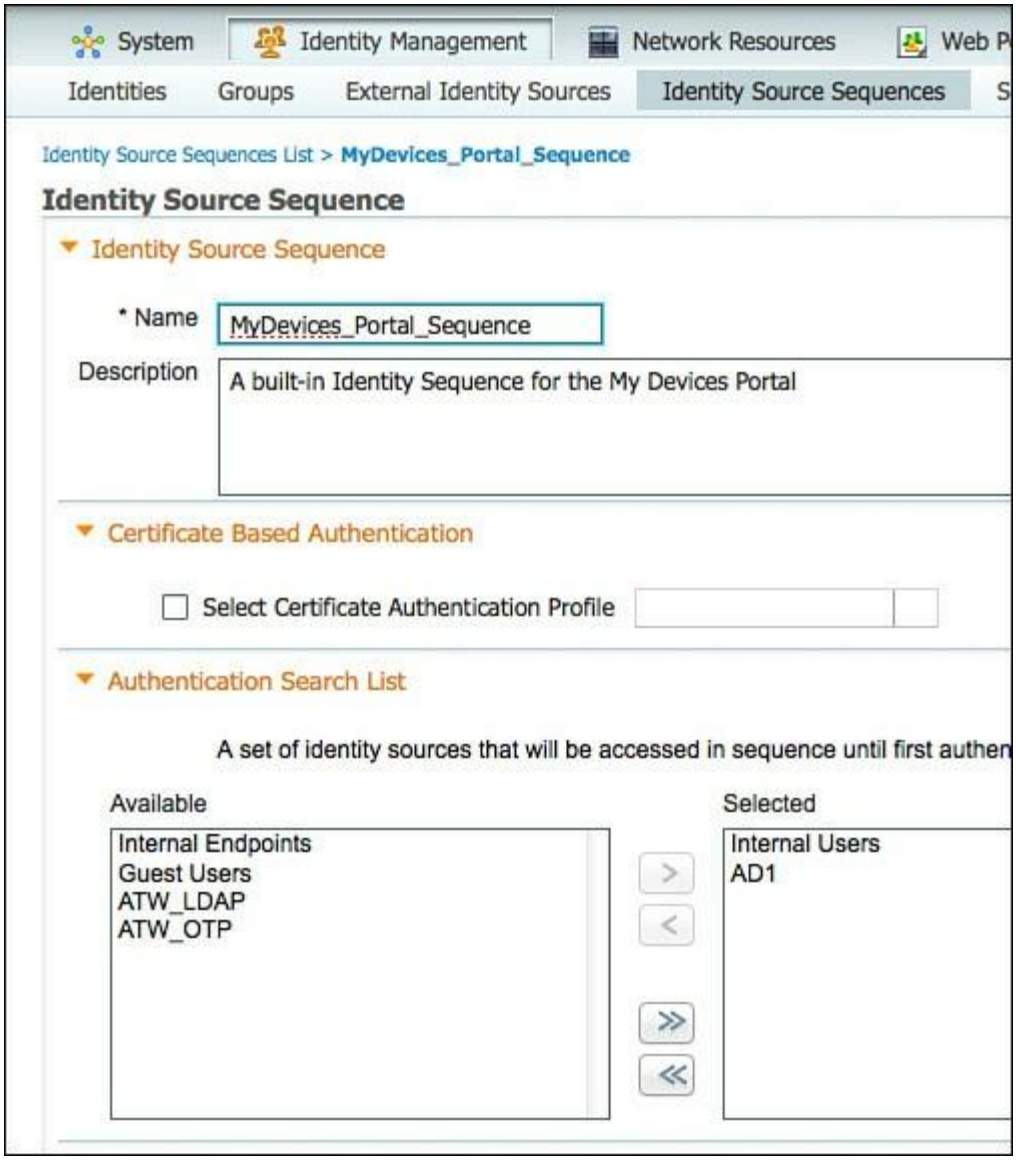


Figure 17-81 MyDevices_Portal_Sequence.

After logging in to the MyDevices portal, the list of registered devices is displayed. As displayed in Figure 17-82, if an MDM is being used, the user can select her device and initiate one of the options, such as Corporate Wipe.

https://t.me/learningnets

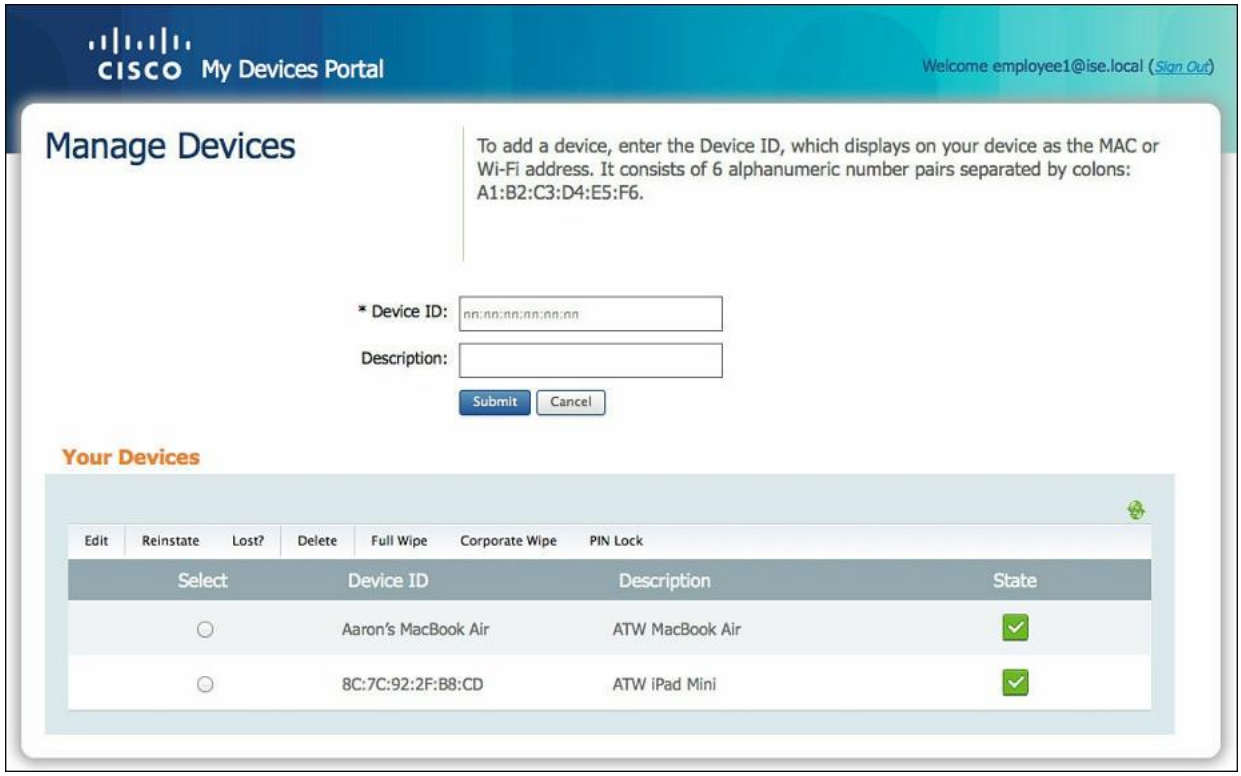


Figure 17-82 MyDevices portal.

Administrative Management

From an administrative perspective, BYOD registered endpoints are administered just like any other endpoint at Administration > Identity Management > Identities > Endpoints. From here, an administrator can initiate actions against the registered devices, as shown in Figure 17-83.

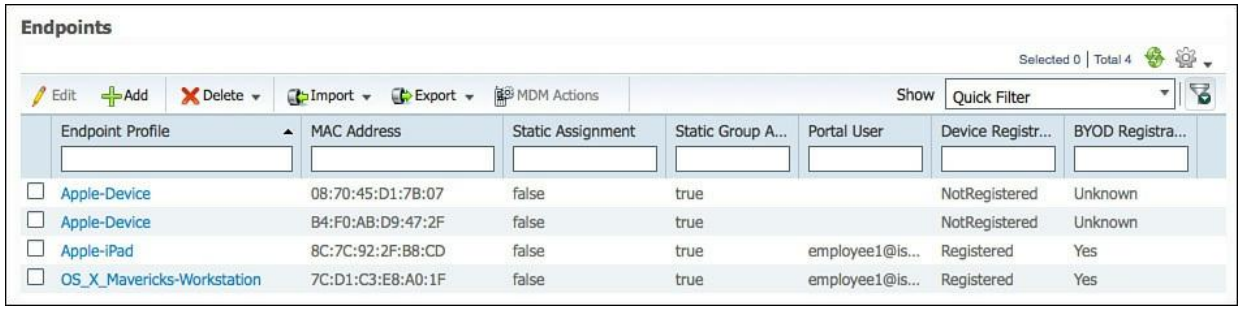


Figure 17-83 Endpoint identities.

The Opposite of BYOD: Identify Corporate Systems

For many years, Cisco has had customers explain their business need to identify the machine as an authorized asset, in addition to the user being an authorized user. Given that Microsoft Windows has both a user and a machine state, the device can be authenticated to the network with what is commonly known as machine auth, as well as have the ability to have the interactive user authenticated to the network.

The issue is that EAP was always designed to transport a single credential. The 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, it changes to a user state and issues the credentials associated to the user. With standard RADIUS and standard EAP, there was no way to join those authentications together.

To resolve the issue, Cisco enhanced EAP-FAST with the ability to do EAP Chaining. EAP Chaining is a capability to authenticate both the machine and the user within the same authentication session. EAP-

https://t.me/learningnets

FASTv2 has been standardized in the IETF and is known as EAP-TEAP (RFC7170). At the time this book was being written, no known vendor (OS or RADIUS server) has implemented TEAP, but the majority of vendors do have plans to implement it.

With EAP-FASTv2 and EAP-Chaining, both the machine and the user are issued a protected access credential (PAC), similar to a secure cookie. Thus, 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 options are

- No chaining.
- User and machine both failed.
- User and machine both succeeded.
- User failed and machine succeeded.
- User succeeded and machine failed.

With that level of flexibility, an authorization result can be provided that permits limited access to remediate a single failure, no access if neither succeeds, and full access if both succeed (for example).

A practical example from a real-world deployment was to use EAP Chaining to identify corporate-owned and -managed devices; the authorization rule acted like this:

Click here to view code image

```
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)
```

That authorization rule allowed only those devices to communicate to the servers housing credit card data.

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 17-4 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	Single-SSID versus dual-SSID onboarding	535
Paragraph	DNS ACLs on WLC	536
Paragraph	Single-SSID onboarding	538
Paragraph	Dual-SSID onboarding	538
Paragraph	CAs with SCEP	568

Table 17-4 Key Topics for Chapter 17

https://t.me/learningnets

Define Key Terms

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

onboarding

MDM mobile Device Management (MDM)

https://t.me/learningnets

Chapter 18. TrustSec and MACSec

This chapter covers the following exam 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 ways of controlling network access based on the context of a user and device. There is virtual local area network (VLAN) assignment, in which access is controlled at the Layer-3 edge or by isolating that VLAN into a segmented virtual network (VRFs). Additionally, there is access control list (ACL) assignment, which can be a local ACL, called into action by a RADIUS attribute, or a downloaded ACL (dACL). These ACLs are applied ingress at the switch port or at the virtual port in the case of the Wireless LAN Controller (WLC).

These are all very good access control methods, but controlling access only at the point of network ingress can leave room for a more desirable and scalable solution. In this chapter, you will learn about one such Cisco enhancement to make access control more scalable and powerful; it's called TrustSec. This is a very confusing name to give to a solution because Cisco has redefined what TrustSec means numerous times. However, you will learn about it in the context expected for the exam, and that definition of TrustSec is a policy-defined segmentation solution that provides advanced enforcement. This was formerly referred to as security group access (SGA).

TrustSec is a natural progression on top of the IEEE 802.1AE industry standard, commonly referred to as MACSec. This standard defines a Layer-2 encryption for Ethernet environments very similar to Wi-Fi Protected Access (WPA).

This chapter focuses on the fundamentals of TrustSec and MACSec as well as showing the configuration of the many devices available for use in a TrustSec environment. Basic use cases are presented for security group ACL, 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 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."

Foundation Topics Section	Questions
Describe SGA	1-3
Describe TrustSec Architecture	1-3, 7
SGT Classification—dynamic/static	3-4
SGT Transport—inline tagging and SXP	6-8
SGT Enforcement—SGACL and SGFW	9
MACsec	8, 10

Table 18-1 "Do I Know This Already?" Section-to-Question Mapping

https://t.me/learningnets

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 a security group tag?

   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. Policy > Policy Elements > Results > Security Group Access

   c. Policy > Policy Elements > Dictionaries > System > Security Group Access

   d. Policy > Firewall > Identity by TrustSec

3. What are three ways that an SGT can be assigned to network traffic?

   a. Manual binding of the IP address to an SGT

   b. Manually configured on the switch port

   c. Dynamically assigned by the network access device

   d. Dynamically assigned by the 802.1X authorization result

   e. Manually configured in the NAC agent profile

   f. Dynamically assigned by the AnyConnect network access manager

4. True or False? An SGT-capable device can automatically map traffic to an SGT based on the VLAN of that traffic.

   a. True

   b. False

5. Which peering protocol can be used to transmit a mapping of IP address 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—Intermediate System (IS-IS)

   c. Border Gateway Protocol (BGP)

   d. Security Group Exchange Protocol (SXP)

6. What are two modes of SXP peers?

   a. Speaker

   b. SGT-Reflector

   c. Listener

   d. SGT-Sender

7. How is the SGT transmitted when using native tagging?

   a. The SGT is included in the Cisco Metadata (CMD) portion of the Layer-2 Frame.

https://t.me/learningnets

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. What are two methods of enforcement with SGTs?

a. SG-ACLs on switches.

b. SG-ACLs on routers.

c. SG-Firewalls.

d. SG-Appliances.

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, whereas downlink MACSec is the encrypted traffic leaving the switch, destined to the endpoint itself.

b. There is no difference between uplink 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, whereas 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 (dACLs) 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. Let's take a look at each one of these standard use cases individually and discuss the challenges.

VLAN Assignment

VLAN assignment based on the context of a user or device is a common way to control access to the 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:

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 into the PCI-Allowed VLAN on the switch or WLC. To use that VLAN assignment to control access to the servers that house that PCI data, an ACL must be applied somewhere. Let's assume the ACL is applied at a firewall between the campus/branch networks and the data center.
4. 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 18-1 illustrates the concept of VLAN-based access control on a single switch.

https://t.me/learningnets

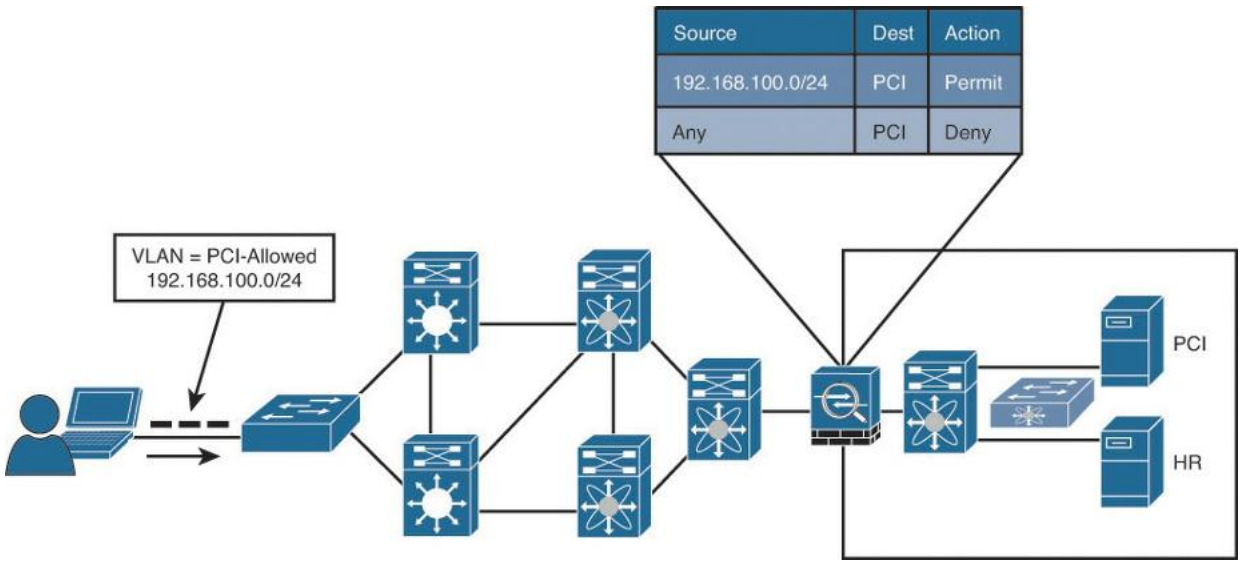


Figure 18-1 Controlling access with VLANs on a single switch.

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 must be built that assign the HR VLAN and another set of entries in the access list. Figure 18-2 illustrates the concept of controlling access with two VLANs on a single switch.

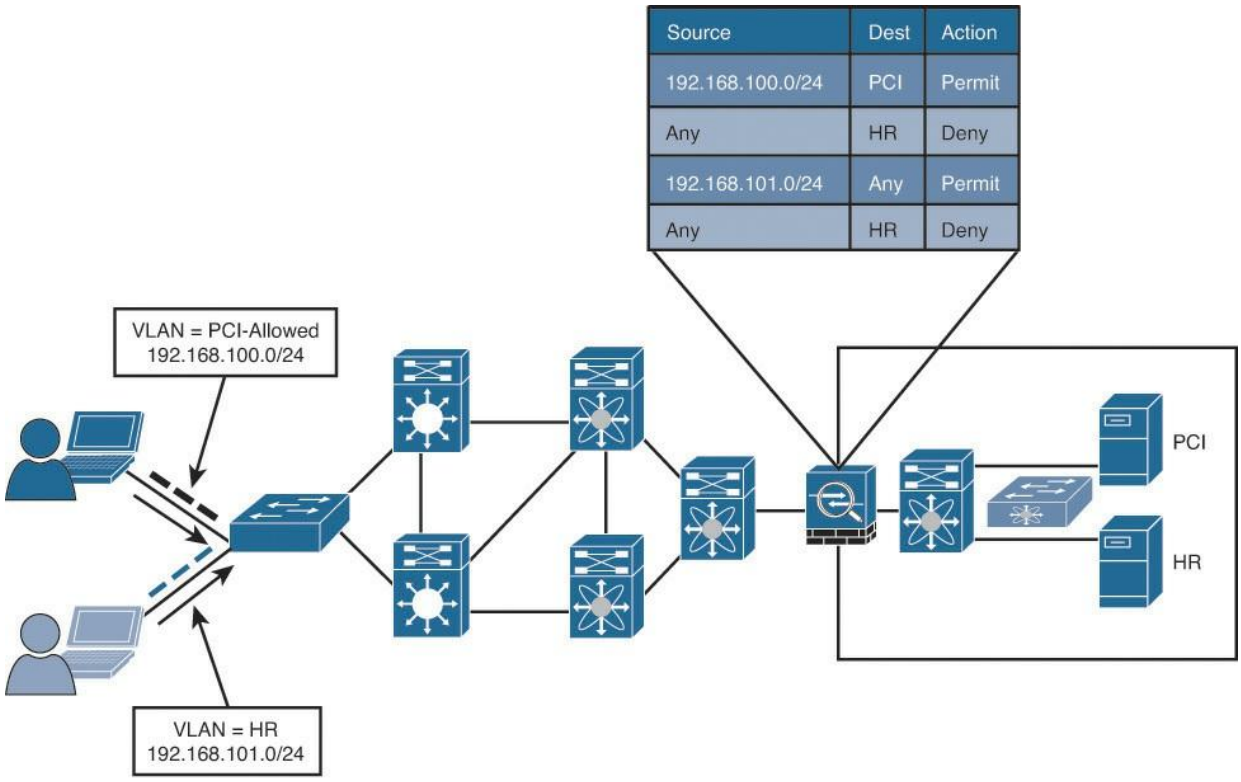


Figure 18-2 Controlling access with two VLANs on a single switch.

Now, consider how this can scale as we continue to add VLANs and add switches and WLCs to the equation. One customer that the author works with has more than 50,000 switches in their access layer. That is a tremendous number of VLANs to create and addresses to maintain in an access list on a firewall. That same customer had 15 full-time employees managing the firewall rules. They needed to find some better mechanism to control access that would lower their OPEX tremendously.

What if you had 100 remote sites? That would mean 100 new IP subnets, and that can easily modify your

https://t.me/learningnets

existing route summarization strategy. When that is the case, the route summarization alone can cause a network redesign, which will add even more operational cost. Figure 18-3 illustrates the growing complexity as the network grows.

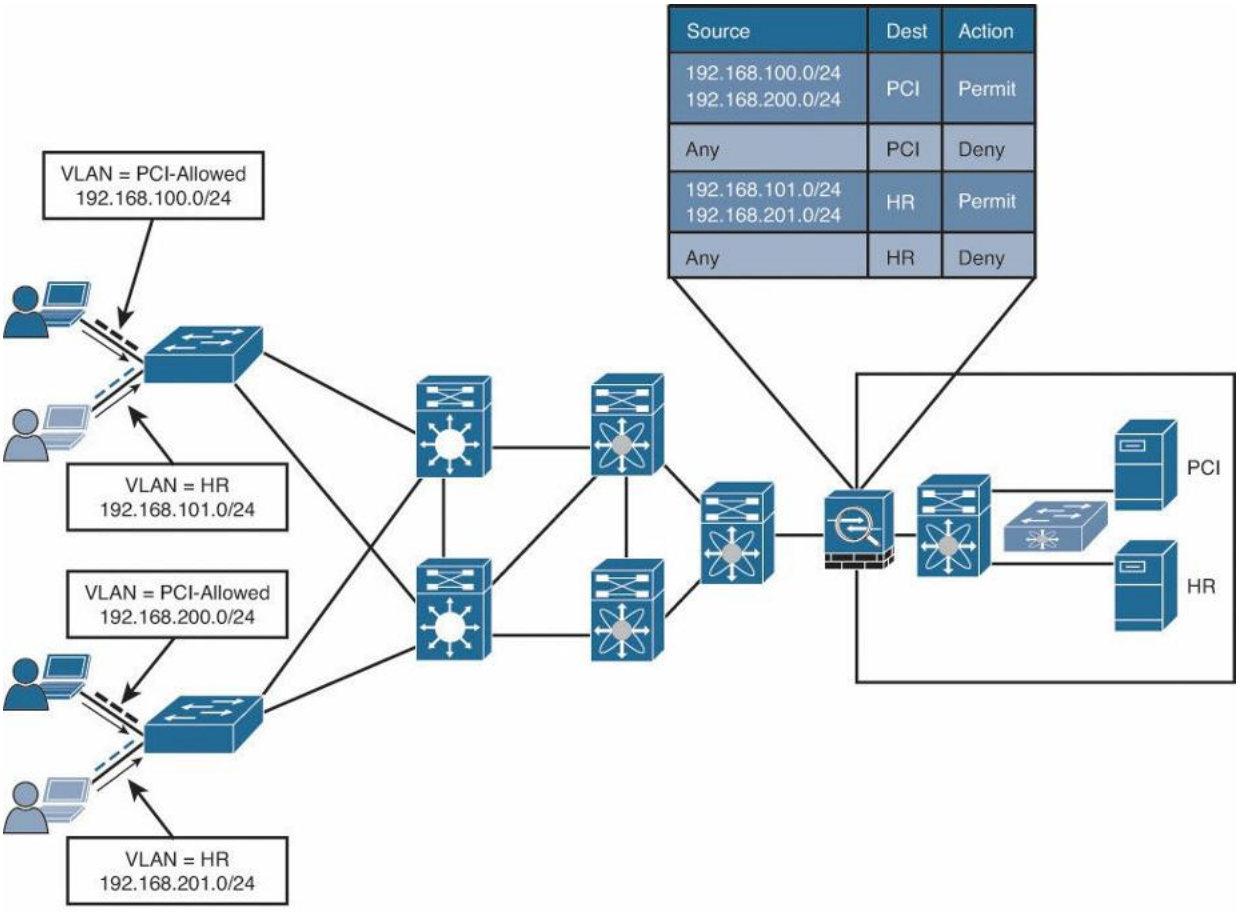


Figure 18-3 VLAN control can be operationally expensive.

There is a formula to determine the number of access control entries (ACEs) in an access control List (ACL). The formula takes the number of sources multiplied by the number of destinations multiplied by the permissions of the ACL:

(# of sources) * (# of destinations) * permissions = # of ACEs

So, with the environment depicted in Figure 18-3—only 4 sources * 2 destinations * 4 permissions—we would need 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

Another way to control access is to use ACLs applied ingress (inbound) at the port (or virtual port) that the user or device is using to access the network. This could be locally defined ACLs that are called by using the filter-id RADIUS attribute, or they could be dACLs where the entire ACL is defined on ISE and downloaded to the port. On the WLC, Airespace ACL is also a viable option.

Obviously, dACLs provide a better operational model because there is only one place to update an ACL when a change needs to be made. Additionally, the number of ACEs required is lower when applying the ACL to a switchport than it would be to apply the ACL to a centralized location. This is because the ACL is being applied at the point of ingress, so there would be only a single source IP address (theoretically). Cisco switches will 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 switchport.

https://t.me/learningnets

Using the same formula for 6 destinations and 4 permissions, we would have:

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 ACLs and the size of the ACLs.

If ACLs are being used to explicitly defend hosts, they must be updated regularly for all new destinations that get added to the network. This can cause an exorbitant amount of operational expense maintaining the lists and ensuring they get updated correctly. Additionally, there is a limited number of ACEs that a switch will be able to apply.

ACLs get loaded into and executed from Ternary CAM (TCAM). Access layer switches have a limited amount of TCAM, which 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.

Due to that limited amount of TCAM, ACLs cannot be overly large, especially when the access layer can be a mixture of different switches, with each switch having a different level of TCAM per ASIC. The best practice recommendation is to keep the ACEs at fewer than 64 per dACL. This might need to be adjusted for your specific environment, but it is a good place to start. On WLCs, the maximum configurable limit is 64 ACLs with a maximum of 64 ACEs per ACL. Figure 18-4 illustrates the concept of ingress ACLs.

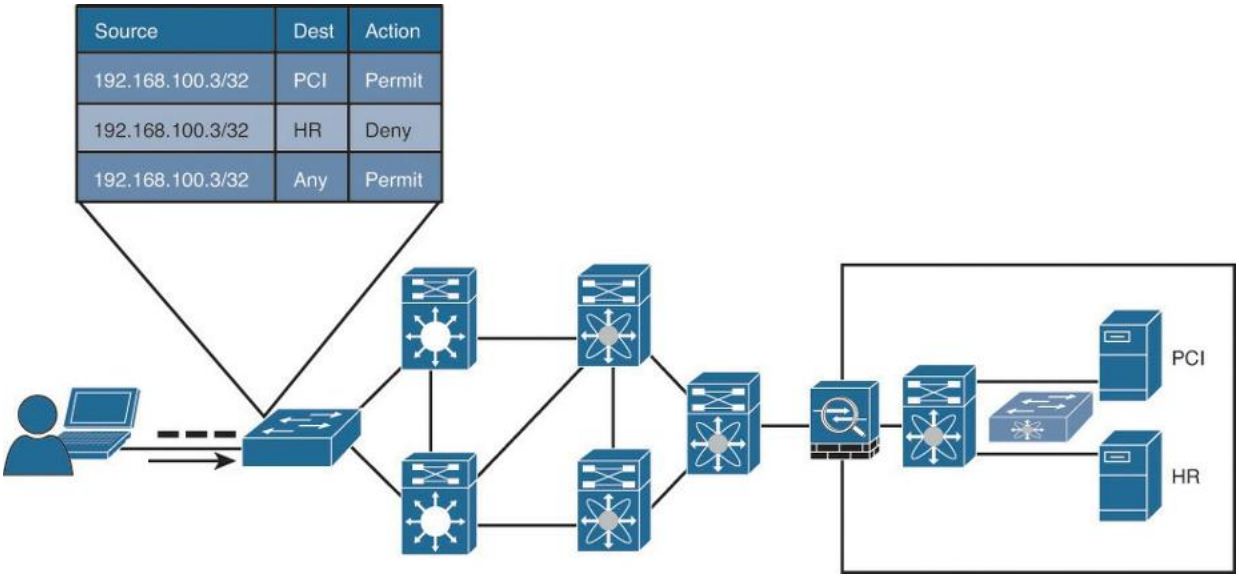


Figure 18-4 Ingress ACLs.

What Is TrustSec?

TrustSec, also 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 with maintaining firewall rules and access lists. SGA is a complimentary 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), to the user's or device's traffic at ingress (inbound into the network) and then enforce the access elsewhere in the infrastructure (in the data center, for example). So, SGA assigns a tag at login and enforces that tag elsewhere in the network (egress enforcement). To be clear, this enforcement happens at the egress of the first SGT-enabled device—not the egress point closest to the destination. Therefore, this approach should not create any unnecessary load on the upstream network for traffic that will be dropped.

https://t.me/learningnets

The SGT should be representative of some overarching role(s) within the company. For instance, an SGT can be assigned to a GUEST user so that GUEST traffic can be isolated from non-GUEST traffic throughout the infrastructure. Here is a list of some common security groups:

- **Network Infrastructure**—This SGT gets assigned to all the switches, routers, WLCs, and firewalls within the organization. This ensures that any management protocols are not impacted by the enforcement of SGT.
- **Network Services**—This SGT is assigned to the servers providing common services that most everyone should be able to reach (DNS, DHCP, NTP, and so on). As noted previously, this ensures that these critical protocols are not impacted by upstream SGT enforcement.
- **Executive**—Many organizations might classify their executives into their own SGT, simply to ensure that executives will never be denied access to anything.
- **Sales**—For the sales force because they are users who might need specific access to customer account management information.
- **Finance**—For the bean counters. These users might require access to payroll, asset management, and profit information—information that others do not require access to.
- **HR**—For the group of employees who have access to your personal and financial data.
- **Line-of-Business-1**—SGTs often are used when an umbrella company has many different lines of business and those lines of business cannot have access to each other's data.
- **Line-of-Business-2**—This is the same as the previous one.

The trick with SGTs is to use them for bulk access control and do your fine-grain access control within the application security itself. Additionally, each end user or end device may be assigned only a single SGT. So, you do not want to create too many roles; otherwise, you will spend too much time trying to figure out all the permutations and will lose the value of TrustSec.

What Is a Security Group Tag?

A security group tag is a 16-bit value that ISE assigns to the user's or endpoint's session upon login. The network infrastructure views the SGT as another attribute to assign to the session and will insert the Layer-2 tag to all traffic from that session. The SGT also can be manually assigned or mapped from the source VLAN. The SGT can represent the context of the user and device. Let's look at an example.

The following is one of my favorite examples from a customer I worked with directly. The customer is a retail organization and therefore accepts credit cards from their customers, which places them under the domain of PCI compliance. Access to any server housing credit card data must be protected as strictly as any technology will allow.

In this customer's case, we 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 that the machine's posture was compliant. If the user and machine met all these conditions, an SGT named PCI was assigned. No access was granted to PCI servers without the PCI SGT.

So, as you can see, SGTs 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.

Figure 18-5 and Figure 18-6 show the assignment of an SGT to an authentication session on a Cisco switch and Cisco WLC, respectively.

https://t.me/learningnets

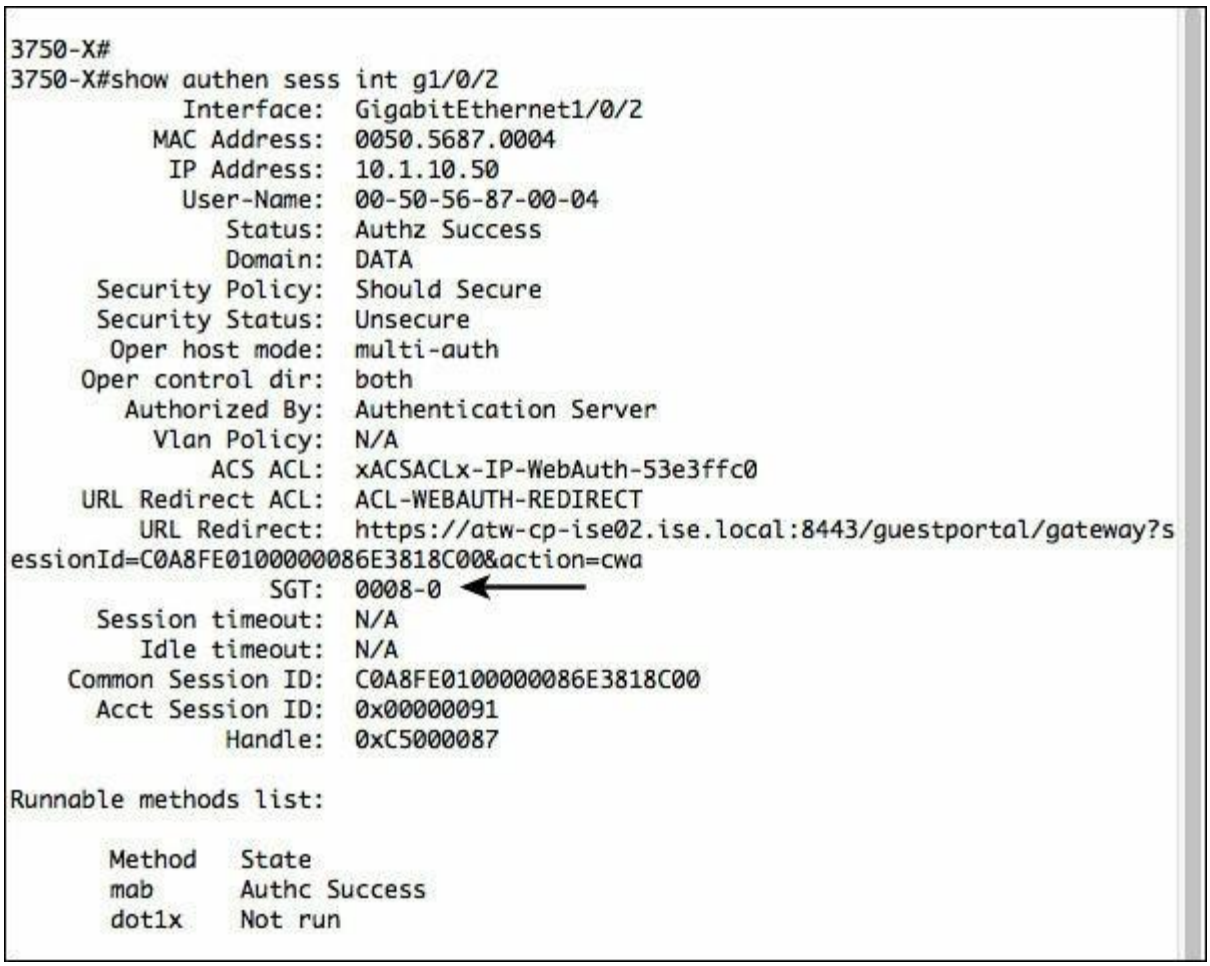


Figure 18-5 SGT applied to session on a switch.

https://t.me/learningnets

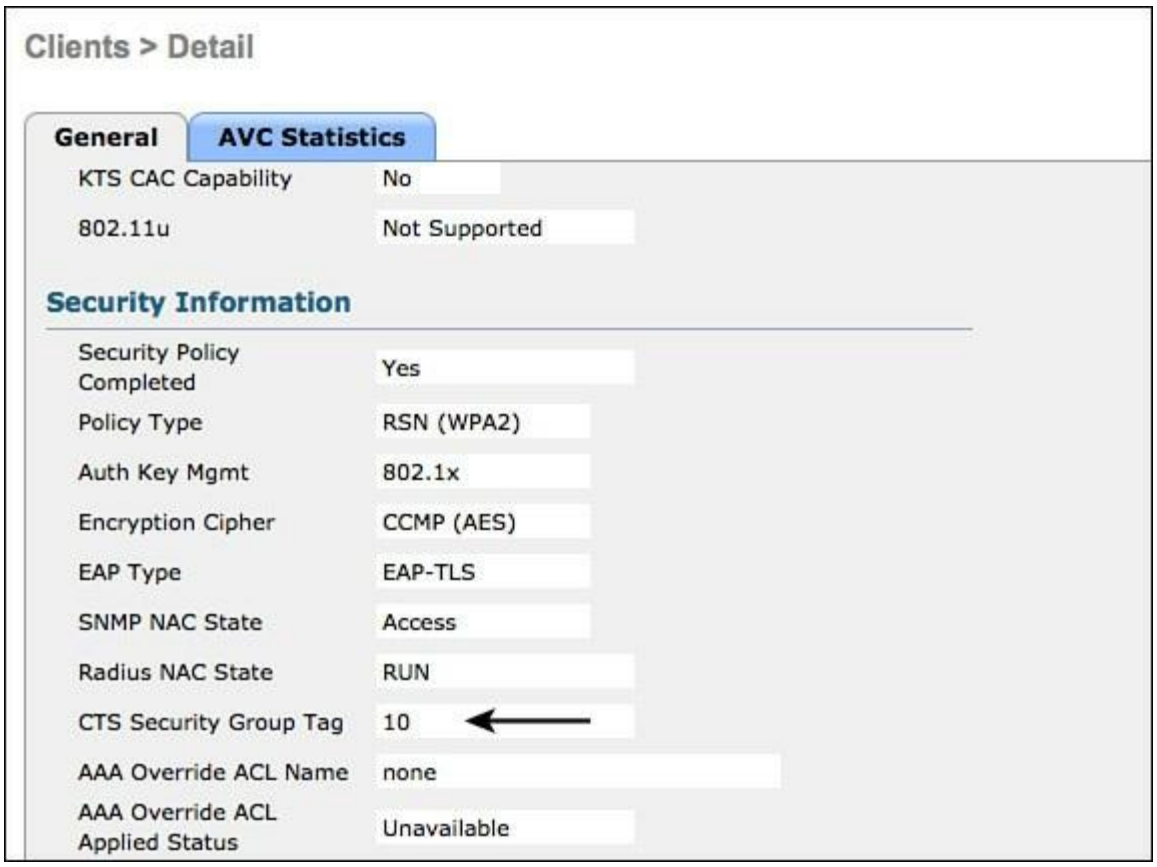


Figure 18-6 SGT applied to session on a WLC.

Defining the SGTs

ISE serves as the single-source-of-truth for which SGTs exist, and ISE considers an SGT a policy result. Therefore, you will create one SGT result for each SGT you want to define in the environment.

Create the SGTs in ISE by doing the following from within the ISE GUI:

Step 1. Navigate to **Policy > Policy Elements > Results**.

Step 2. Select **Security Group Access > Security Groups**.

Figure 18-7 shows the Security Groups list. Notice the default is a SGT of 0, "unknown." This is the tag that will be used if traffic arrives that is untagged. In other words, even the lack of an SGT can be used in the security policy.

https://t.me/learningnets

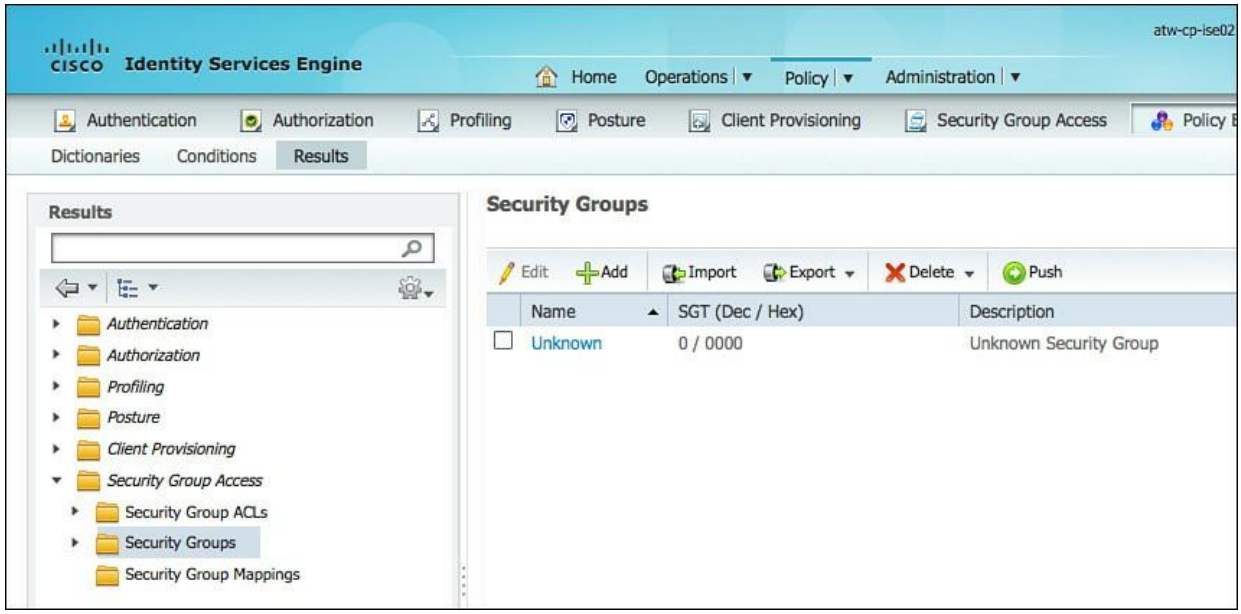


Figure 18-7 Security groups.

Now you are adding new security groups to be used in your network infrastructure. You begin by creating a security group for network access devices:

Step 3. Click **Add**.

Step 4. Name the new SGT **NADs**.

Figure 18-8 shows the adding of a new SGT for the NADs. Notice in the figure that the SGT value is predetermined. ISE 1.2 automatically assigns the value from 2 to 65535. Automatic assignment can be disabled for more manual control over tag values.



Figure 18-8 Adding a security group for NADs.

Step 5. Click **Submit** to save.

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.

Step 6. Click **Add**.

https://t.me/learningnets

Step 7. Name the new group **CommonServices**.

Step 8. Click **Submit** to save.

Step 9. Repeat Steps 6–8 until you have the security groups created for your organization's security policy.

Figure 18-9 shows a sample set of security groups.

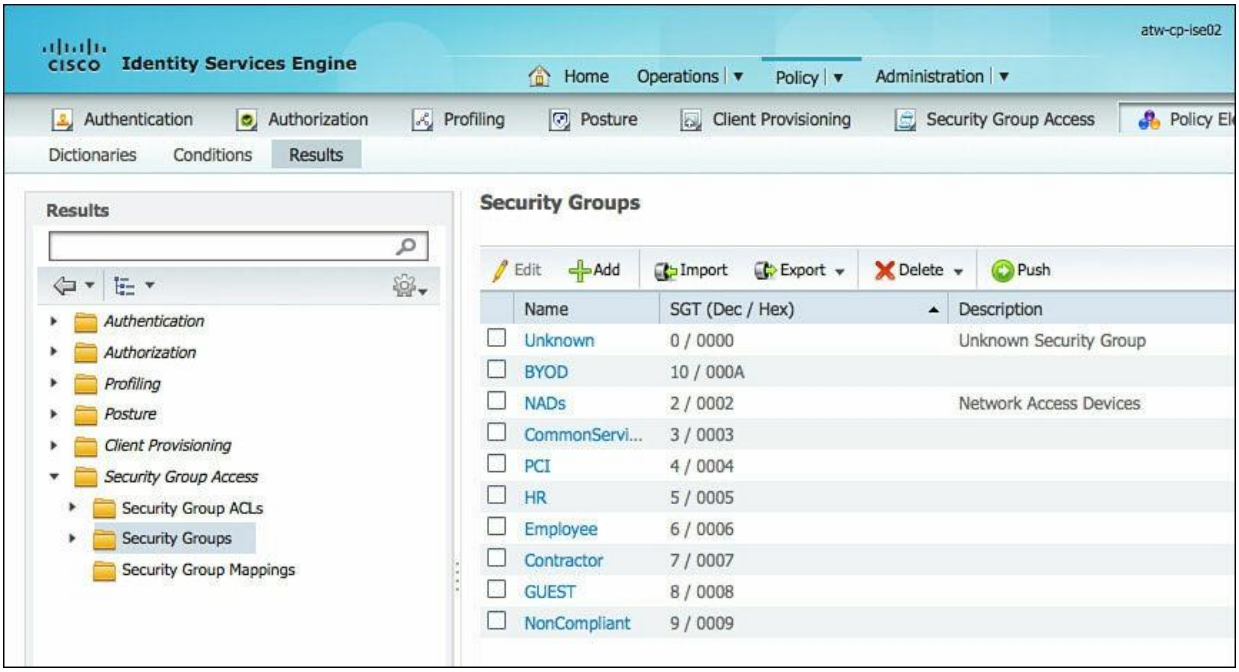


Figure 18-9 Sample security groups.

Classification

To use SGTs within your network infrastructure, your network devices must support SGTs. All Secure Access–supported Cisco switches and wireless controllers do support the assignment of the SGT. The assignment of an SGT to an authentication session is defined as classification. The process of communicating that assigned SGT upstream into the network can occur via either native tagging or a peering-protocol; this process is defined as Transport or Propagation.

Figure 18-10 shows an example of one access switch that has native tagging, with the packets getting tagged on the uplink port, and through the infrastructure. It also shows a switch that is not native tagging capable; it uses a peering protocol to update the upstream switch. In both cases, the upstream switch continues to tag the traffic throughout the infrastructure.

https://t.me/learningnets

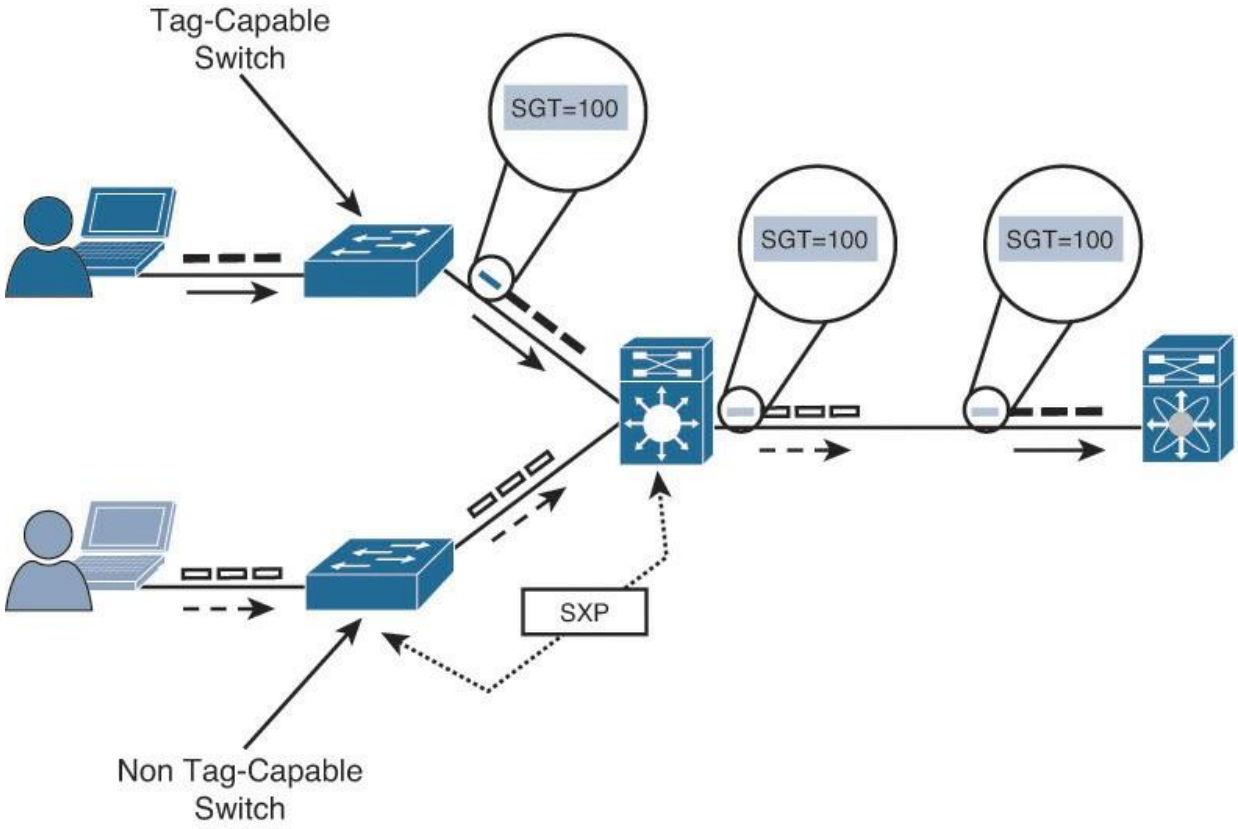


Figure 18-10 Security group tagging.

To use the SGT, the tag must be assigned (classification). This can happen dynamically and be downloaded as the result of an ISE authorization. Tags also can be assigned manually at the port level or even mapped to IP addresses and downloaded to SGT-capable devices.

Dynamically Assigning SGT via 802.1X

Assigning a tag is as simple as adding it as another permission or result of an authorization in an authorization policy. From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Authorization**.

Step 2. Edit an existing authorization rule.

Step 3. Click the plus (+) sign under Permissions.

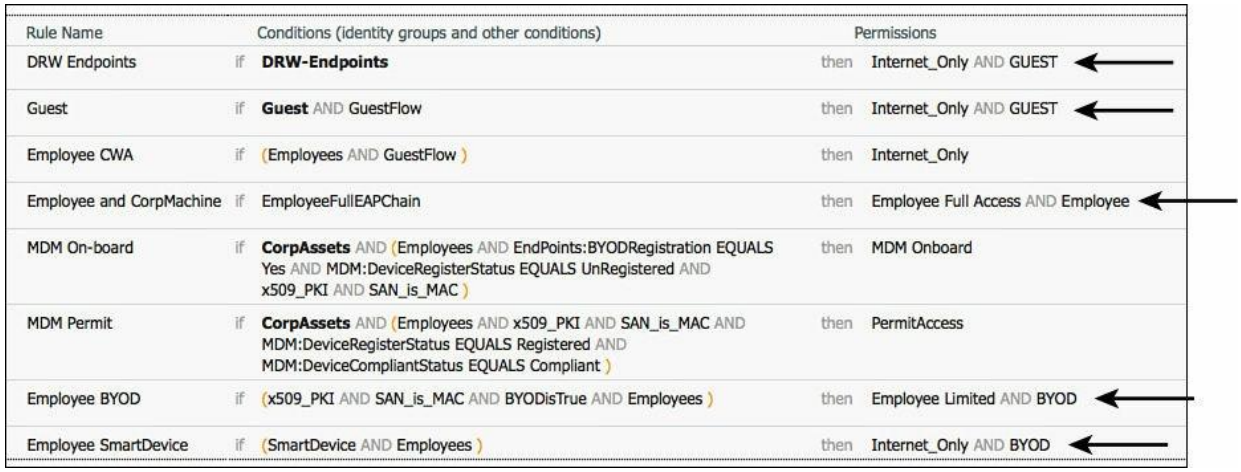
Step 4. Click the plus (+) sign next to the Authorization Profile.

Step 5. Select **Security Group**.

Step 6. Select the appropriate security group to apply.

Figure 18-11 shows some of the authorization rules using both a standard authorization profile result and a security group tag. The two results are concatenated with the AND operator. The tag is deployed only to devices that are SGA capable, as provided in the Administration -> Network Resources -> Network Devices configuration for the NAD. Otherwise, the standard authorization profile policy is deployed to the NAD on behalf of the endpoint.

https://t.me/learningnets



Rule Name		Conditions (identity groups and other conditions)		Permissions
DRW Endpoints	if	**DRW-Endpoints**	then	Internet_Only AND GUEST
Guest	if	**Guest** AND GuestFlow	then	Internet_Only AND GUEST
Employee CWA	if	(Employees AND GuestFlow)	then	Internet_Only
Employee and CorpMachine	if	EmployeeFullEAPChain	then	Employee Full Access AND Employee
MDM On-board	if	**CorpAssets** AND (Employees AND EndPoints:BYODRegistration EQUALS Yes AND MDM:DeviceRegisterStatus EQUALS UnRegistered AND x509_PKI AND SAN_is_MAC)	then	MDM Onboard
MDM Permit	if	**CorpAssets** AND (Employees AND x509_PKI AND SAN_is_MAC AND MDM:DeviceRegisterStatus EQUALS Registered AND MDM:DeviceCompliantStatus EQUALS Compliant)	then	PermitAccess
Employee BYOD	if	(x509_PKI AND SAN_is_MAC AND BYODisTrue AND Employees)	then	Employee Limited AND BYOD
Employee SmartDevice	if	(SmartDevice AND Employees)	then	Internet_Only AND BYOD

Figure 18-11 Authorization rules with SGT assignment.

Manually Assigning SGT at the Port

In most cases, 802.1X will not be used in the data center. Servers are not usually required to authenticate themselves to the data center switch because the data center is normally considered physically secure and no network access control is applied there. However, the servers themselves will be the destination of traffic coming from the campus and from within the data center itself.

Because 802.1X is not typically used in the data center, we need a manual way to apply the SGT. This is configured at the interface level of the Nexus configuration and is manually applied to the port itself. Example 18-1 shows the assignment of an SGT tag to a switch port.

Example 18-1 Static SGT Assignment

Click here to view code image

```
NX7K-DIST(config)# int eth1/3
NX7K-DIST(config-if)# cts manual
NX7K-DIST(config-if-cts-manual)# policy static sgt 0x3
```

Example 18-1 has manually assigned the SGT “3” to the port on the Nexus 7000. This is also available on the Nexus 5000 and 1000v.

Manually Binding IP Addresses to SGTs

As an alternative to assigning the SGT to the port itself, ISE has the ability to centrally configure a database of IP addresses and their corresponding SGTs, and then SGT-capable devices can download that list from ISE.

From the ISE GUI, do the following:

Step 1. Navigate to **Policy > Policy Elements > Results > Security Group Access > Security Group Mappings**.

Step 2. Click **Add**.

Step 3. Select the appropriate SGT from the drop-down list.

Step 4. Enter the IP address or DNS name of the host.

Step 5. Click **Submit** to save.

Figure 18-12 illustrates an example of mapping a server IP address to an SGT named PCI.

https://t.me/learningnets

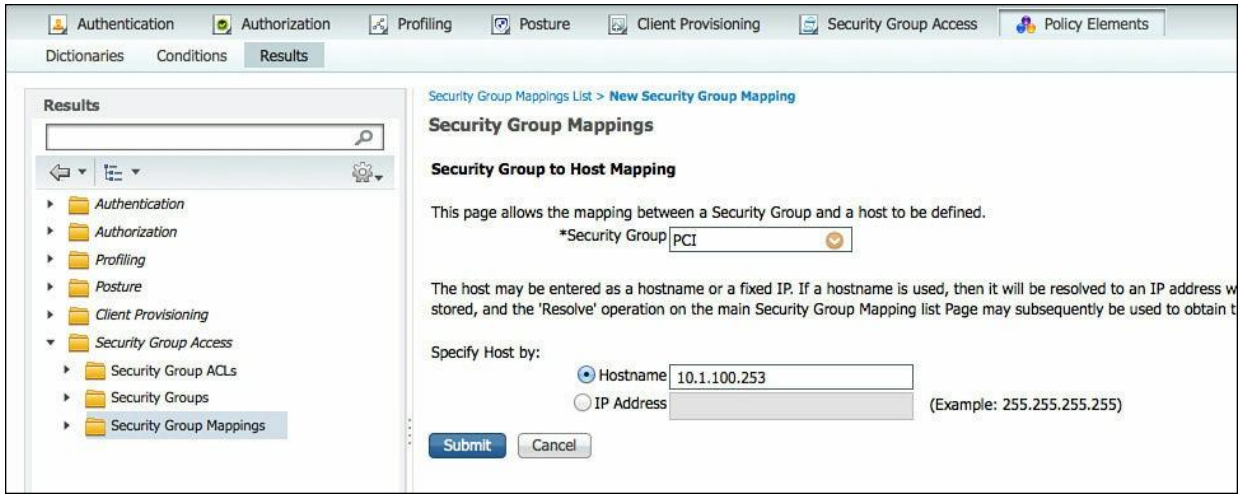


Figure 18-12 Mapping an SGT to an IP address in ISE.

Repeat Steps 2–5 for any other static IP-to-SGT mappings required for your organization. Figure 18-13 shows the security group mappings screen with multiple mappings.

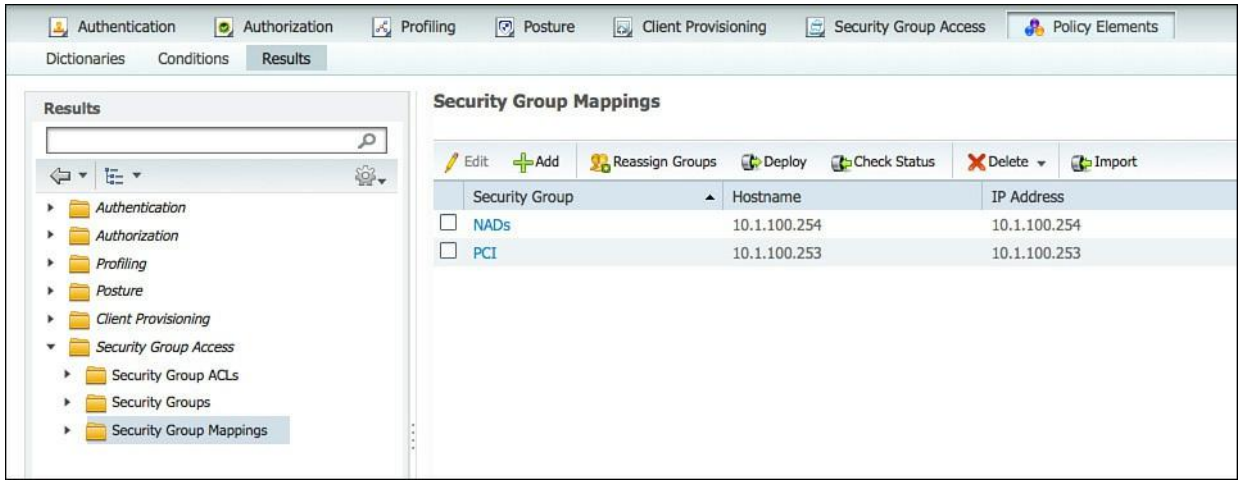


Figure 18-13 Mappings screen with multiple SGT-to-IP mappings.

Now that the mappings exist on ISE, other devices will be able to download them, such as a Nexus 7000 data center switch.

Access Layer Devices That Do Not Support SGTs

Because we do not live in a perfect world, and not all the equipment on the network will be the latest and greatest, we need another way to classify the endpoint traffic. One example is an older Cisco WLC (such as the 4400) that does not support version 7.2 or newer and therefore cannot accept the SGT classification from ISE, nor send the update via SXP.

Additionally, this could be a VPN concentrator or some third-party equipment that found its way into the deployment. Although that gear might not support the classification and transport natively, it might be capable of assigning different VLANs or IP addresses per authorization result.

With the many distribution layer devices, you have the ability to map subnets and VLANs and assign all source IP addresses from the subnet or VLAN to a specific tag.

Mapping a Subnet to an SGT

The **cts role-based sgt-map** [ipv4-subnet | ipv6-subnet] **sgt** tag-value command enables the subnet to SGT binding. When used, the IP device tracking feature in the switch identifies matches and assigns the SGT. Example 18-2 shows the use of the **cts role-based sgt-map** command.

https://t.me/learningnets

Example 18-2 cts role-based sgt-map for Subnets

Click here to view code image

```
C6K-DIST(config)#cts role-based sgt-map 192.168.26.0/24 sgt 4
```

Mapping a VLAN to an SGT

The **cts role-based sgt-map vlan-list** vlans **sgt** tag-value command enables the VLAN-to-SGT binding. When used, the IP device tracking feature in the switch identifies matches and assigns the SGT. Example 18-3 shows the mapping of a VLAN list to an SGT.

Example 18-3 cts role-based sgt-map for VLAN Lists

Click here to view code image

```
C6K-DIST(config)#cts role-based sgt-map vlan-list 40 sgt 4
```

Transport: Security Group Exchange Protocol

In a perfect world, all your access layer devices would support tagging the users' traffic natively. As we do not live in a perfect world, not all devices support native tagging. However, all supported Cisco switches support the assignment of the SGT to the authentication session (known as classification).

Cisco has developed a peering protocol (similar to BGP or LDP) to enable devices to communicate their database of IP address-to-SGT mappings to one another. This peering protocol is called Security Group Exchange Protocol (SXP). Because this is a peering protocol, it is possible to be very specific and deterministic as to which devices send updates and which ones receive updates.



An SXP peer can be a speaker or a listener. A speaker is a device that sends the IP address-to-SGT bindings. A listener is a device that receives the bindings.

SXP Design

SXP uses TCP as its transport, so the peers can be Layer-2 adjacent or multiple hops away. A network device can peer to an upstream switch for tag propagation or even peer directly to the enforcement device (the data center switch or security group firewall). Figure 18-14 illustrates the concept of SXP peering.

https://t.me/learningnets

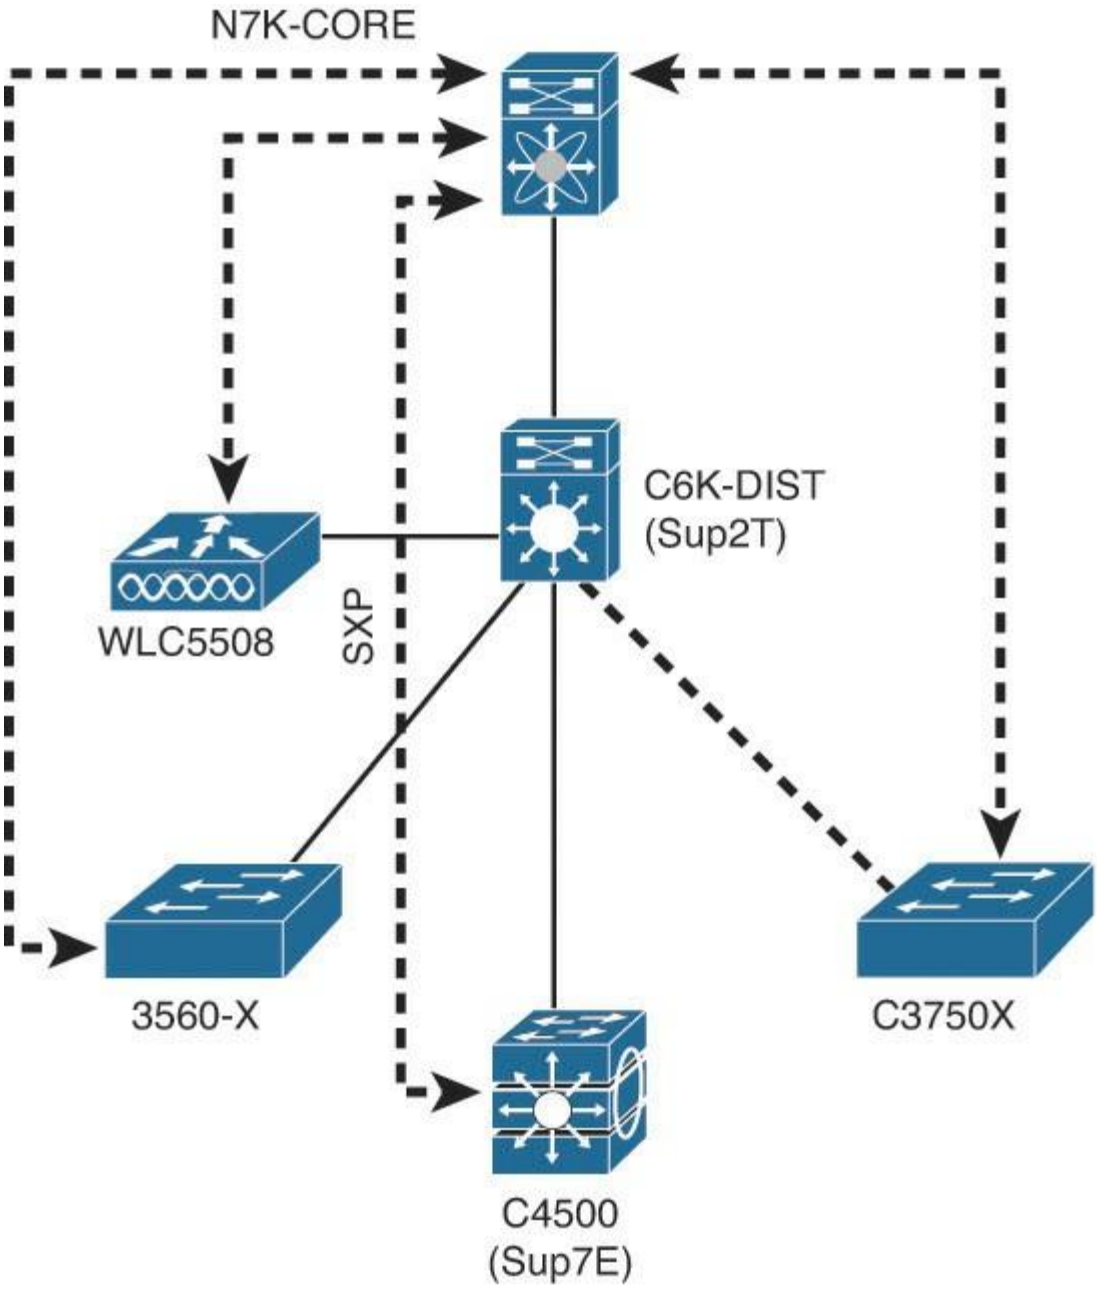


Figure 18-14 SXP peering.

Routing protocols have a limitation for the number of neighbors they can scale to, and so does SXP. Due to the limitations of scale for the number of peers, SXP design can be architected to be multi-hop, which allows for aggregation points. Devices such as the Catalyst 6500 with a Supervisor 2T engine or the Aggregation Services Router (ASR) are solid choices for SXP aggregation. Figure 18-15 illustrates the multi-hop SXP concept.

https://t.me/learningnets

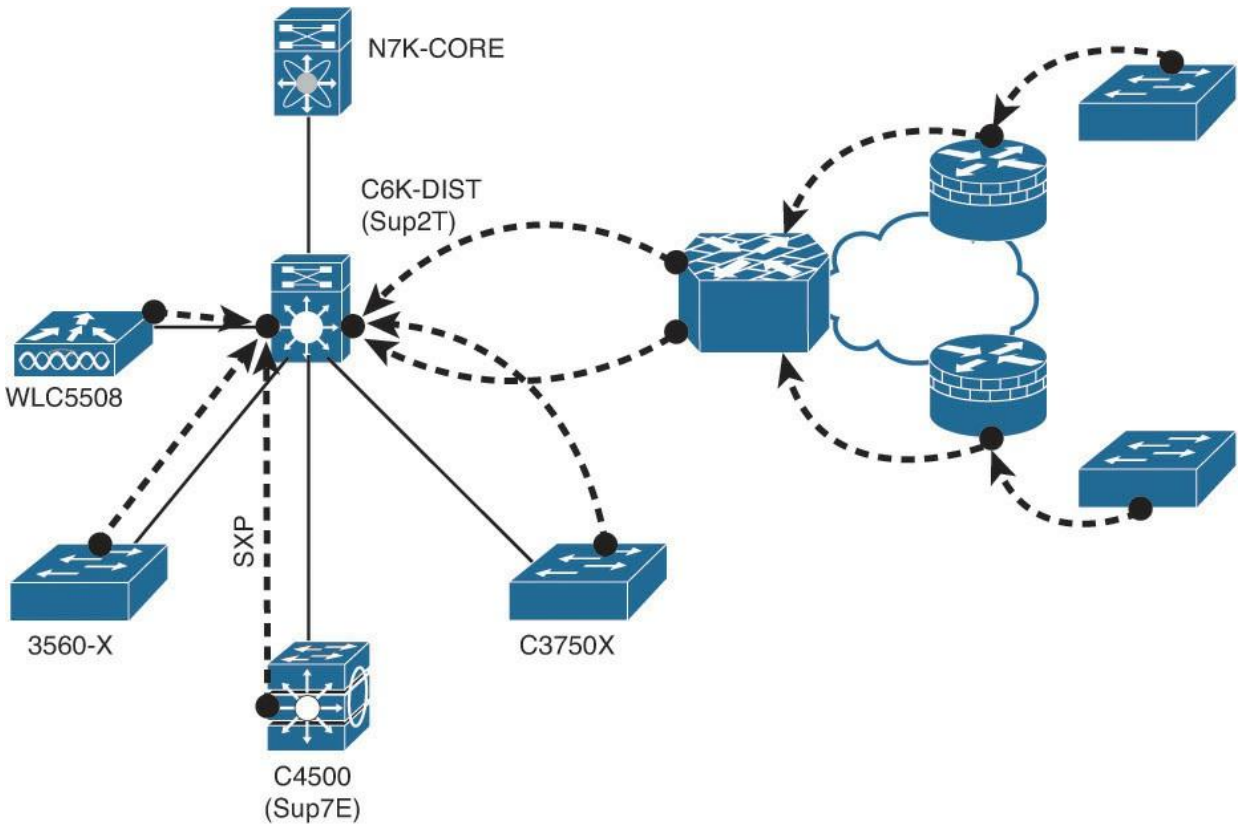


Figure 18-15 SXP multi-hop.

There are numerous benefits to this design. Not only does it not require SXP-aware infrastructure along every hop in the network path, but it also provides a deterministic scalable design.

Configuring SXP on IOS Devices

From global configuration, you must enable SXP globally. Each peer will need to be added individually, as well as have a global default SXP password set. Follow these steps:

Step 1. Enter **cts sxp enable**.

Step 2. Enter **cts sxp connection peer** [peer-ip-address] **password** [default | none] **mode** [local | peer] [listener | speaker].

This command is used to define the SXP peer. The options are as follows:

- **password default**—This states to use the password defined globally for all SXP connections. At the current time, it is not possible to have different SXP passwords per peer.
- **password none**—Do not use a password with this SXP peer.
- **mode local**—This states that the following SXP argument is defining the local side of the connection.
- **mode peer**—This states that the following SXP argument is defining the peer's side of the connection.
- **listener**—Defines that the specified device (local or peer) will receive SXP updates through this connection.
- **speaker**—Defines that the specified device (local or peer) will send SXP updates through this connection.

Step 3. (optional) **cts sxp default password password.**

https://t.me/learningnets

Step 3 is an optional step for when your connections will use the globally defined password instead of no password.

Examples 18-4 and 18-5 display the steps for setting up the SXP connection between a 4500 (an access layer device that does not support native tagging) and a 6500 with a Supervisor 2T (a distribution layer device that supports native tagging). Figure 18-16 illustrates the connection.

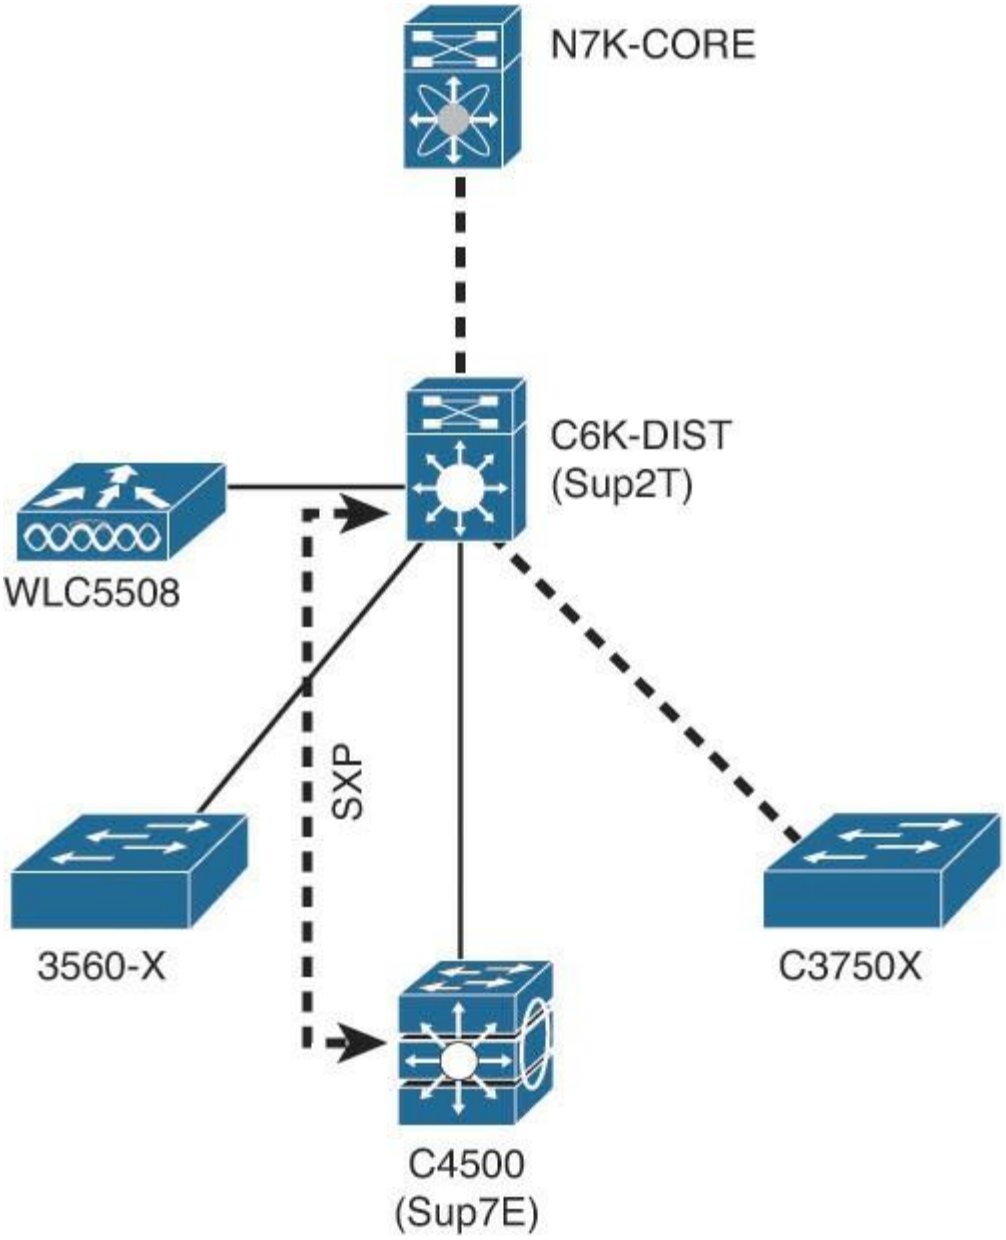


Figure 18-16 SXP between 4500 and 6500.

Example 18-4 Enabling SXP on the 4500

Click here to view code image

```
4503(config)#cts sxp enable
4503(config)#
*Aug  9 06:51:04.000: %CTS-5-SXP_STATE_CHANGE: CTS SXP enabled
4503(config)#cts sxp connection peer 10.1.40.1 password default mode peer listener
4503(config)#
```

https://t.me/learningnets

```
*Aug 10 09:15:15.564: %CTS-6-SXP_TIMER_START: Connection <0.0.0.0, 0.0.0.0> retry
open timer started.
*Aug 10 09:15:15.565: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.1,
10.1.40.2>-1 state changed from Off to Pending_On.
4503(config)#
*Aug 10 09:15:15.566: %CTS-3-SXP_CONN_STATE_CHG_OFF: Connection <10.1.40.1,
10.1.40.2>-1 state changed from Pending_On to Off.
4503(config)#cts sxp default password TrustSec123
4503(config)#
*Aug 10 09:17:20.936: %CTS-5-SXP_DFT_PASSWORD_CHANGE: CTS SXP password changed.
```

Example 18-5 Enabling SXP on the 6500

Click here to view code image

```
C6K-DIST(config)#cts sxp enable
C6K-DIST(config)#
Aug 10 16:16:25.719: %CTS-6-SXP_TIMER_START: Connection <0.0.0.0, 0.0.0.0> retry
open timer started.
C6K-DIST(config)#cts sxp default password TrustSec123
C6K-DIST(config)#cts sxp connection peer 10.1.40.2 password default mode peer
speaker
C6K-DIST(config)#
Aug 10 16:17:26.687: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.2, 10.1.40.1>-1
state changed from Off to Pending_On.
Aug 10 16:17:26.687: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.2, 10.1.40.1>-1
state changed from Pending_On to On.
```

Configuring SXP on Wireless LAN Controllers

Cisco's WLC added support for SGT classification and SXP transport beginning with the 7.2 release.

From the WLC user interface, do the following:

Step 1. Using the top menu navigation, select **Security**.

Step 2. Along the left side, select **TrustSec SXP** (second from the bottom).

Configure the settings on this page to be:

- SXP State = Enabled
- Default Password = the same default password you configured on the switches (all passwords in the SXP "domain" will need to be the same)

This has enabled SXP globally. Each peer will need to be added individually. To add a new SXP peer (a listener), follow these steps:

Step 3. Click **New** (in the upper-right corner).

Step 4. Enter the IP address of the listener peer.

Step 5. Click **Apply** (in the upper-right corner).

The added peers will be displayed on the TrustSec SXP page. Their statuses will be listed next to their IP addresses. After the peer is configured on the other side, the status should change from off to on.

Figure 18-17 displays a completed TrustSec SXP page.

https://t.me/learningnets



Figure 18-17 TrustSec SXP page—peer status.

You also can verify the SXP connection from the other side, as shown in Example 18-6.

Example 18-6 Verifying the Connection Between the WLC and 6500

Click here to view code image

```
C6K-DIST#sho cts sxp connections brief
 SXP              : Enabled
 Default Password : Set
 Default Source IP: Not Set
Connection retry open period: 120 secs
Reconcile period: 120 secs
Retry open timer is not running

------------------------------------------------------------------------------
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
```

https://t.me/learningnets

Configuring SXP on Cisco ASA

The Cisco Adaptive Security Appliance (ASA) has support for SGT enforcement (known commonly as SG firewall). Certain models of the ASA have added support for native tagging with version 9.3 and newer, but SXP is the most common method used for the transport of IP-to-TAG bindings, and it is the solution covered by the exam objectives.

It is important to note that the ASA has multiple functions. These functions include deep packet inspection firewalling and remote access VPN (among many others). ASA version 9.2.1 and newer enforce SGTs, receive SGTs (transport), and assign SGTs to remote access VPN users (classification).

From the ASA Device Manager (ASDM), do the following:

Step 1. Navigate to **Configuration > Firewall > Identity by TrustSec**.

Step 2. Globally enable SXP by checking the **Enable SGT Exchange Protocol (SXP)** check box in the upper-left corner.

Step 3. Click **Add** to add a new SXP peer.

Step 4. In the Add Connection Peer pop-up window, add the IP address of the remote peer.

Step 5. Select Default for the Password (unless you will not be using passwords).

Step 6. Set the mode to Peer.

Step 7. Set the role to Speaker.

Step 8. Click **OK**.

After clicking **OK**, you are returned to the main identity by TrustSec page. At this point, you will have SXP enabled and a single peer defined, but no default password yet.

Step 9. (optional) If you will be specifying the source IP address of the ASA, you can configure that source in the Default Source field.

Step 10. Enter the default password for your entire SXP deployment.

Figure 18-18 shows the SXP connections on the ASA via the ASDM management GUI.

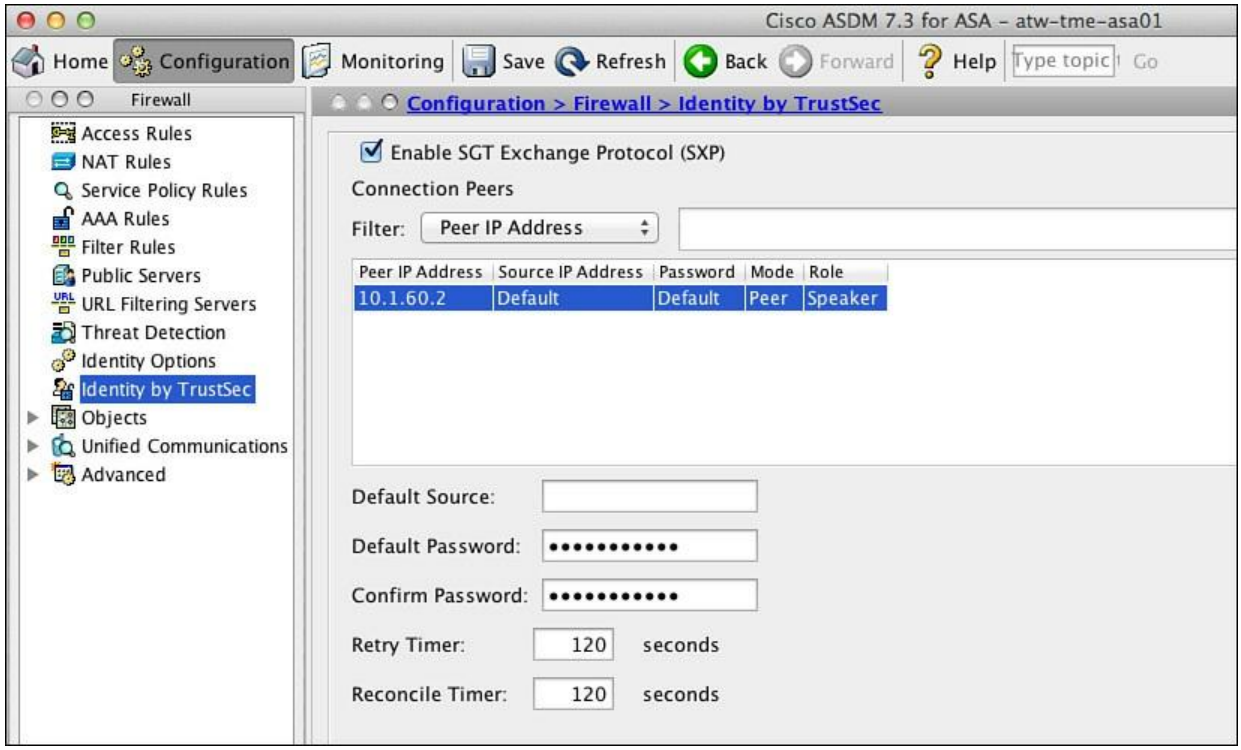


Figure 18-18 SXP connections in ASDM.

https://t.me/learningnets

Verifying SXP Connections in ASDM

Verifying SXP connections in ASDM is accomplished through **Monitoring > Properties > Identity by TrustSec pages**.

From ASDM, follow these steps:

Step 1. Navigate to **Monitoring > Properties > Identity by TrustSec**.

Step 2. Click **SXP Connections** to see the configured peers and their statuses, as shown in Figure 18-19.

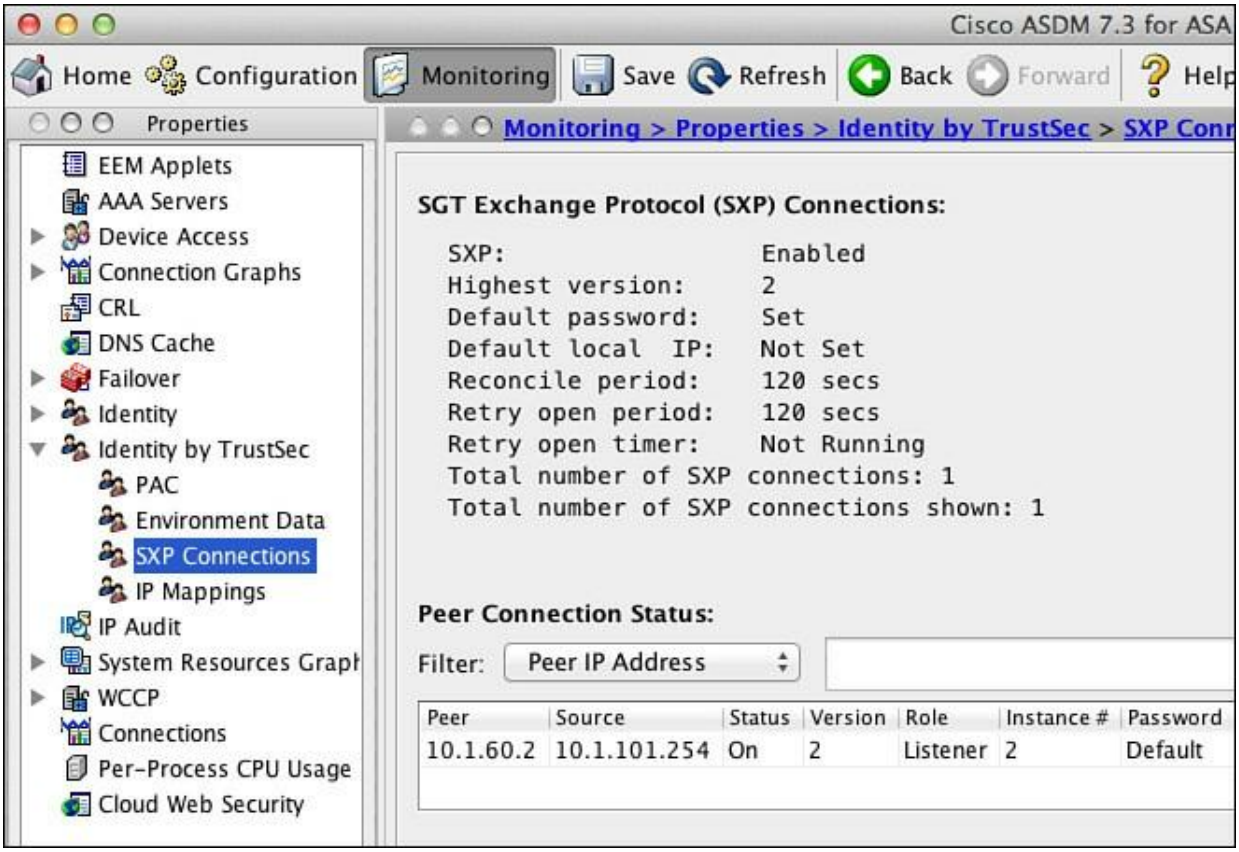


Figure 18-19 Monitoring SXP connections in ASDM.

Step 3. Click **IP Mappings** to see any IP address-to-SGT mappings that the ASA has learned about, as shown in Figure 18-20.

https://t.me/learningnets

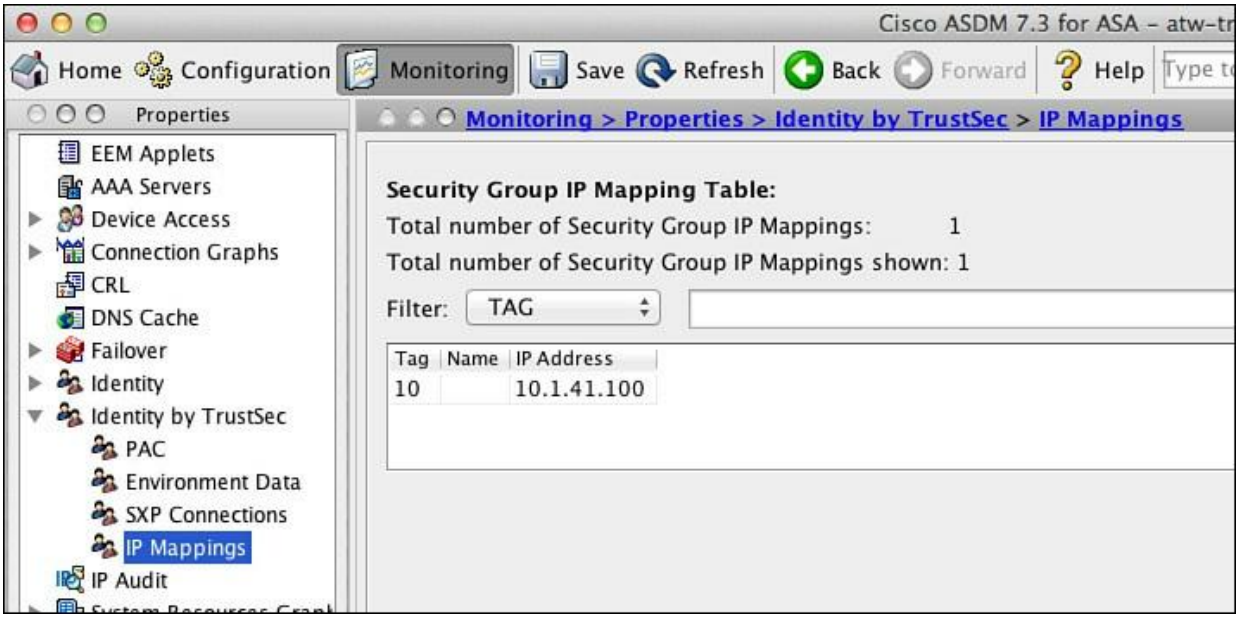


Figure 18-20 Monitoring IP address to SGT mappings in ASDM.

Transport: Native Tagging

Native tagging is the ultimate goal with any TrustSec deployment. With this approach, the access layer is capable of applying the SGT to the Layer-2 frame as it is sent across the wire to the upstream host. The upstream host will continue that and ensure the tag is applied as the traffic is sent out the egress interface toward the next hop and so on; this way, the tag is always present throughout the entire infrastructure.

Figure 18-21 shows the format of the Layer-2 frame, with the SGT value that exists within the Cisco CMD field circled.

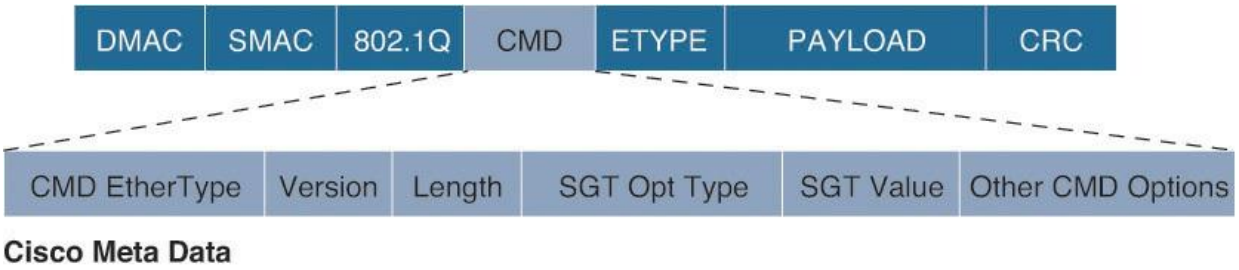


Figure 18-21 Layer-2 frame format with SGT.

Native tagging enables the technology to scale to virtually endless size, and it remains 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 18-22 illustrates the pervasive tagging concept.

https://t.me/learningnets

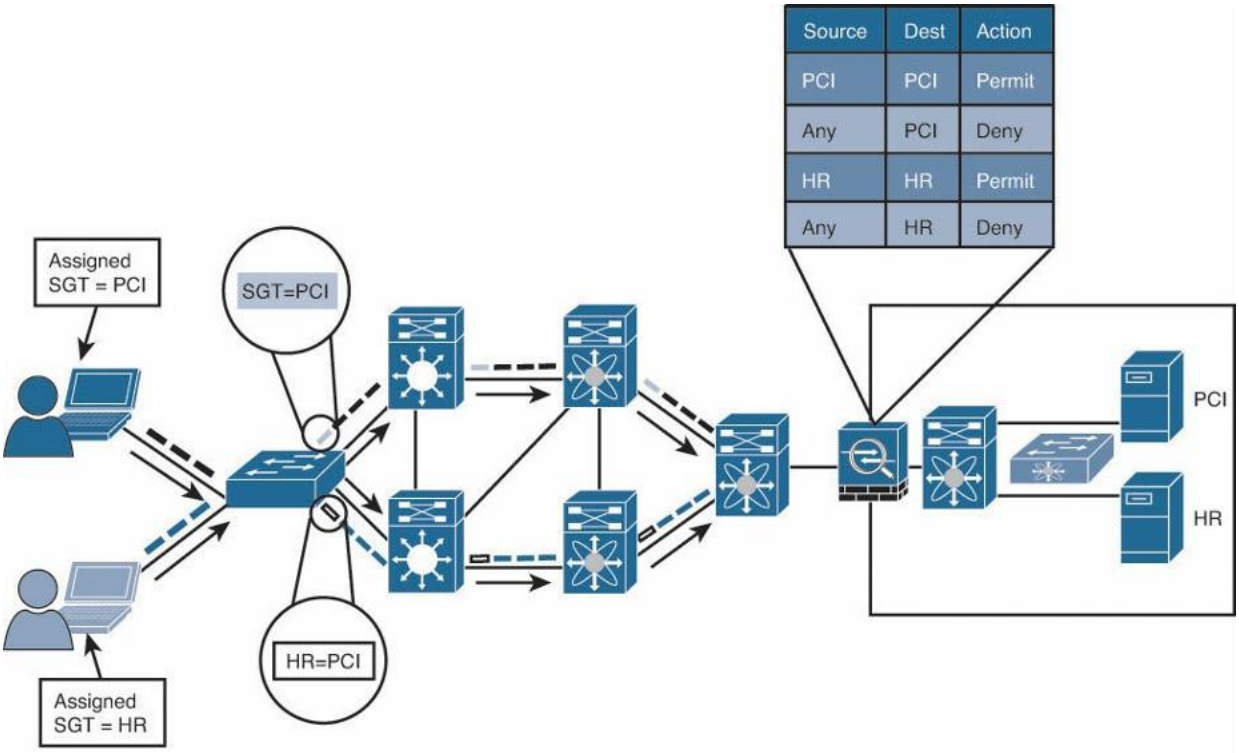


Figure 18-22 Pervasive tagging.

As you can see in Figure 18-22, when native tags are supported pervasively within the infrastructure, the SGT is communicated hop-by-hop. This provides for end-to-end segmentation and tremendous scale. With the tag being applied to the traffic at every Layer-2 link, you are able to enforce policy at any point in the infrastructure and there are no limitations to the size of an IP-to-SGT mapping database because the mapping database is not being used at all.

For added security, the tag can be encrypted with MACSec or IPSec and the network infrastructure can be authenticated prior to sending or receiving Tags.

Configuring Native SGT Propagation (Tagging)

The next few configuration exercises show the enabling of native security group tagging on three types of switches: a 3750 series access layer switch, a 6500 series distribution layer switch, and Nexus data center switches. Figure 18-23 shows the logical network layout used in the configuration examples to follow.

https://t.me/learningnets

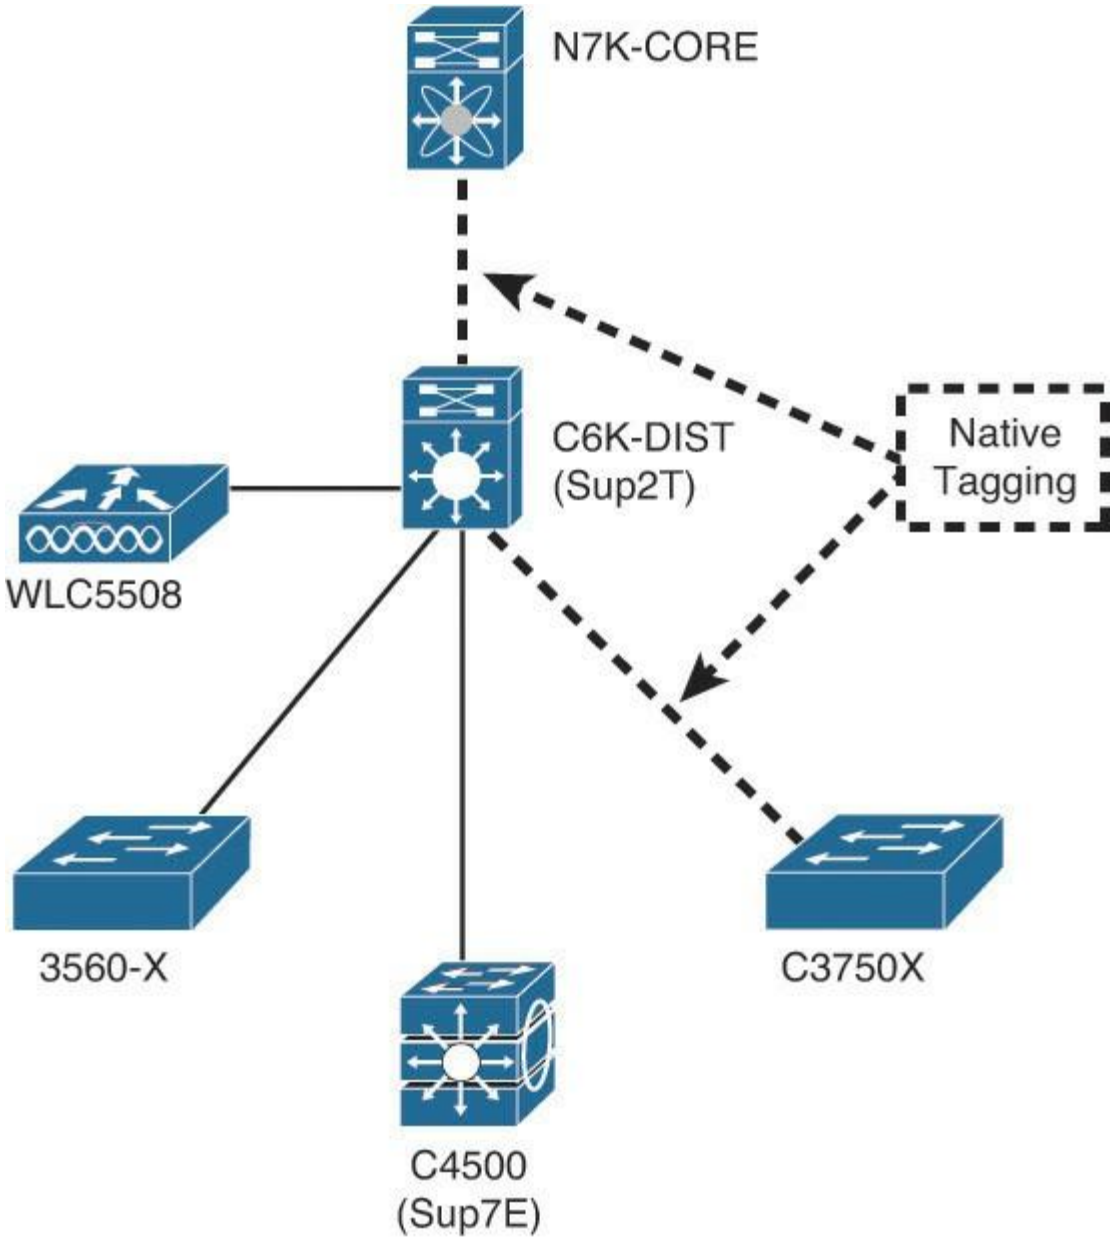


Figure 18-23 SGTs from access to distribution and distribution to data center.

Configuring SGT Propagation on Cisco IOS Switches

This section discusses the configuration of SGT propagation on access layer switches such as the 3560-X and 3750-X that have the capability to use native tags. The Catalyst 6500 and Nexus series switches are covered in the following sections.

When it comes to inserting the tag into Layer-2 traffic, there is a fundamental choice to make: to use encryption or not to use encryption. For simplicity, let's focus on the easy one: without encryption. Encrypting the tag is discussed more the section "MACSec."

From global configuration, do the following:

Step 1. Enter the **cts role-based enforcement** command.

This will globally enable the tagging of SGTs. It also enables the ability to enforce SGACLs (discussed in a future section). However, without this command in the global configuration, the switch will not tag the Layer-2 traffic.

https://t.me/learningnets

Step 2. Enter into interface configuration mode of the tagging-capable port by typing **interface** interface-name.

Step 3. Enter the **cts manual** command.

We are using **cts manual** because we are not utilizing NDAC at this point. The **cts manual** mode of operation enables us to apply the tag to the Layer-2 frame without needing to negotiate encryption or requiring a fully trusted domain of Cisco switches (such as we would need with NDAC).

Step 4. Enter the **policy static sgt** sgt-value **trusted** command.

When we created our security groups earlier in this chapter, we created a special group for NADs. We called that security group "NADs" and the value of that group was 2 (0x02). That is the value we are applying here with this **policy static sgt 2 trusted** command. The trusted keyword in this command ensures that no changes are made to the incoming tags because they are from a trusted source.

Example 18-7 shows the enabling of tagging on a 3750-X, and Example 18-8 shows how to verify the tagging.

Example 18-7 Enabling Tagging on a 3750-X Series Access Switch

Click here to view code image

```
C3750X(config)#cts role-based enforcement
C3750X(config)#interface Ten 1/1/1
C3750X(config-if)#cts manual
C3750X(config-if-cts-manual)#policy static sgt 2 trusted
```

Example 18-8 Verifying Tagging on a 3750-X Series Access Switch

Click here to view code image

```
C3750X#sho cts interface Ten 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: ""
    Authorization Status: SUCCEEDED
        Peer SGT:         2
        Peer SGT assignment: Trusted
    SAP Status:           NOT APPLICABLE
        Configured pairwise ciphers:
            gcm-encrypt
            null

        Replay protection:   enabled
        Replay protection mode: STRICT

        Selected cipher:

    Propagate SGT:        Enabled
    Cache Info:
        Cache applied to link : NONE

    Statistics:
        authc success:           0
        authc reject:            0
```

https://t.me/learningnets

```
        authc failure:           0
        authc no response:       0
        authc logoff:            0
        sap success:             0
        sap fail:                0
        authz success:           3
        authz fail:              0
        port auth fail:          0

    L3 IPM: disabled.
```

Configuring SGT Propagation on a Catalyst 6500

The Catalyst 6500 is a special case. This switch is sometimes used in the access layer, but it's most often used in the distribution layer or even in the data center. There are also a tremendous number of line cards possible for this chassis-based switch, some of which can support native tagging and some of which cannot. Due to the possibility of multiple locations and multiple line card possibilities, the Catalyst 6500 requires the administrator to set whether the switch should be used for egress (receiving the tag from other devices) or ingress, which would place it at the access layer. These modes are referred to as reflector modes.

> **Note**
>
> This switch is unable to be configured for both ingress and egress mode simultaneously.

Ingress reflector mode should be used only in the access layer. This mode allows the use of non–SGA-capable line cards along with an SGA-capable supervisor. (An example of this would be a Catalyst 6504-E chassis populated with a Supervisor 2T and a 6148 series line card.) With this mode, all packet forwarding occurs on the Supervisor 2T PFC. Line cards that use distributed forwarding are not supported in ingress reflector mode (such as the 6748-GE-TX).

With this mode of operation, ISE is able to assign an SGT to a device entering the access layer via any supported line card, but that tag is applied only to network traffic leaving one of the ports physically on the Supervisor 2T. In other words, the switch can apply the tag on an uplink port, but not any of the downlink ports. Additionally, the switch will not be able to read the incoming tag on any ports except the ones physically on the Supervisor 2T module itself.

Egress reflector mode is normally associated with the 6500 being deployed in the distribution layer or data center. With this mode, SGA propagation and encryption (MACSec) can be enabled on the Supervisor 2T and 6900 series line cards. These are the model of line card most often seen in the distribution layer, and as such, this provides for a nice SGA aggregation design. The switch will be able to read all incoming SGT tagged packets and apply that tag to the traffic leaving the switch as well. This is the model of security group tagging that one normally thinks of when discussing the topic. Additionally, if the Catalyst 6500 is a SXP peer, it is capable of applying the SGT to Layer-2 traffic based on the IP to SGT bindings learned via SXP.

From global configuration on the Catalyast 6500, do the following:

Step 1. Select the CTS reflector mode by typing **platform cts** {**egress** | **ingress**}.

Because this is a distribution layer deployment of the Catalyst 6500, choose egress mode. If this were an access layer deployment, where end users would be authenticated, then you would have chosen ingress mode.

Step 2. Enter **cts role-based enforcement**.

This will globally enable the tagging of SGTs. It also enables the ability to enforce SGACLs (discussed in a future section). However, without this command in the global configuration, the switch will not tag the Layer-2 traffic.

https://t.me/learningnets

Step 3. Enter into interface configuration mode of the tagging-capable port by typing **interface** interface-name.

Step 4. Enter **cts manual**.

We are using **cts manual** because we are not utilizing NDAC at this point. The **cts manual** mode of operation enables us to apply the tag to the Layer-2 frame, without needing to negotiate encryption or requiring a fully trusted domain of Cisco switches (such as we would need with NDAC).

Step 5. Enter **policy static sgt** sgt-value trusted.

When we created our security groups earlier in this chapter, we created a special group for network access devices. We called that security group "NADs," and the value of that group was 2 (0x02). That is the value we are applying here with this **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 18-9 Enabling Tagging on Catalyst 6500 Supervisor 2T

Click here to view code image

```
C6K-DIST(config)#platform cts egress
C6K-DIST(config)#cts role-based enforcement
C6K-DIST(config)#interface Ten1/5
C6K-DIST(config-if)#cts manual
C6K-DIST(config-if-cts-manual)#policy static sgt 2 trusted
```

Example 18-10 Verifying Tagging on the Catalyst 6500 Supervisor 2T

Click here to view code image

```
C6K-DIST#show cts interface Ten1/5
Global Dot1x feature is Enabled
Interface TenGigabitEthernet1/5:
    CTS is enabled, mode: MANUAL
    IFC state:            OPEN
    Authentication Status: NOT APPLICABLE
        Peer identity:    "unknown"
        Peer's advertised capabilities: ""
    Authorization Status: SUCCEEDED
        Peer SGT:         2
        Peer SGT assignment: Trusted
    SAP Status:           NOT APPLICABLE
        Configured pairwise ciphers:
            gcm-encrypt
            null

        Replay protection:   enabled
        Replay protection mode: STRICT

        Selected cipher:

    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
```

https://t.me/learningnets

```
        sap success:              0
        sap fail:                 0
        authz success:            1
        authz fail:               0
        port auth fail:           0

    L3 IPM: disabled.
```

Configuring SGT Propagation on a Nexus Series Switch

You've just configured the SGT propagation on the access and distribution layers. Now you will enable the propagation on the data center switch.

From global configuration on the Nexus Series switch, do the following:

Step 1. Type **feature dot1x** at global configuration mode.

The Nexus Series requires the feature dot1x to be enabled before enabling CTS features.

Step 2. Type **cts enable** at global configuration mode.

This command enables security group access and MACSec and NDAC features to be enabled and configured.

Step 3. Enter **cts role-based enforcement**.

This will globally enable the tagging of SGTs. It also enables the ability to enforce SGACLs (discussed in a future section). However, without this command in the global configuration, the switch will not tag the Layer-2 traffic.

Step 4. Enter into interface configuration mode of the tagging-capable port by typing **interface** interface-name

Step 5. Enter **cts manual**.

We are using **cts manual** because we are not utilizing NDAC at this point. The **cts manual** mode of operation enables us to apply the tag to the Layer-2 frame, without needing to negotiate encryption or requiring a fully trusted domain of Cisco switches (such as we would need with NDAC).

Step 6. Enter **policy static sgt** sgt-value **trusted**.

When we created our security groups earlier in this chapter, we created a special group for network access devices. We called that security group "NADs," and the value of that group was 2 (0x02). That is the value we are applying here with this **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 18-11 Enabling Tagging on Nexus 7000

Click here to view code image

```
NX7K-CORE(config)# feature dot1x
NX7K-CORE(config)# cts enable
NX7K-CORE(config)# cts role-based enforcement
NX7K-CORE(config)# int eth1/26
NX7K-CORE(config-if)# cts manual
NX7K-CORE(config-if-cts-manual)# policy static sgt 0x2 trusted
```

Enforcement

Now that we 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.

https://t.me/learningnets

There are multiple ways to enforce traffic based on the tag, but ultimately, we can summarize them into two major types:

- Enforcement on a switch (SGACL)
- Enforcement on a firewall (SG-FW)

SGACL

Historically, enforcement with SGACL was the only option available. It started with the Nexus 7000 Series and has expanded to the many additional switches. A major benefit to SGACL usage is the consolidation of access control entries (ACEs) and the operational savings involved with maintenance of those traditional access lists.

An SGACL can be visualized in a format similar to a spreadsheet. It is always based on a source tag to a destination tag. Figure 18-24 shows an example SGACL policy on ISE, which represents the SGACLs in a columns-and-rows presentation.

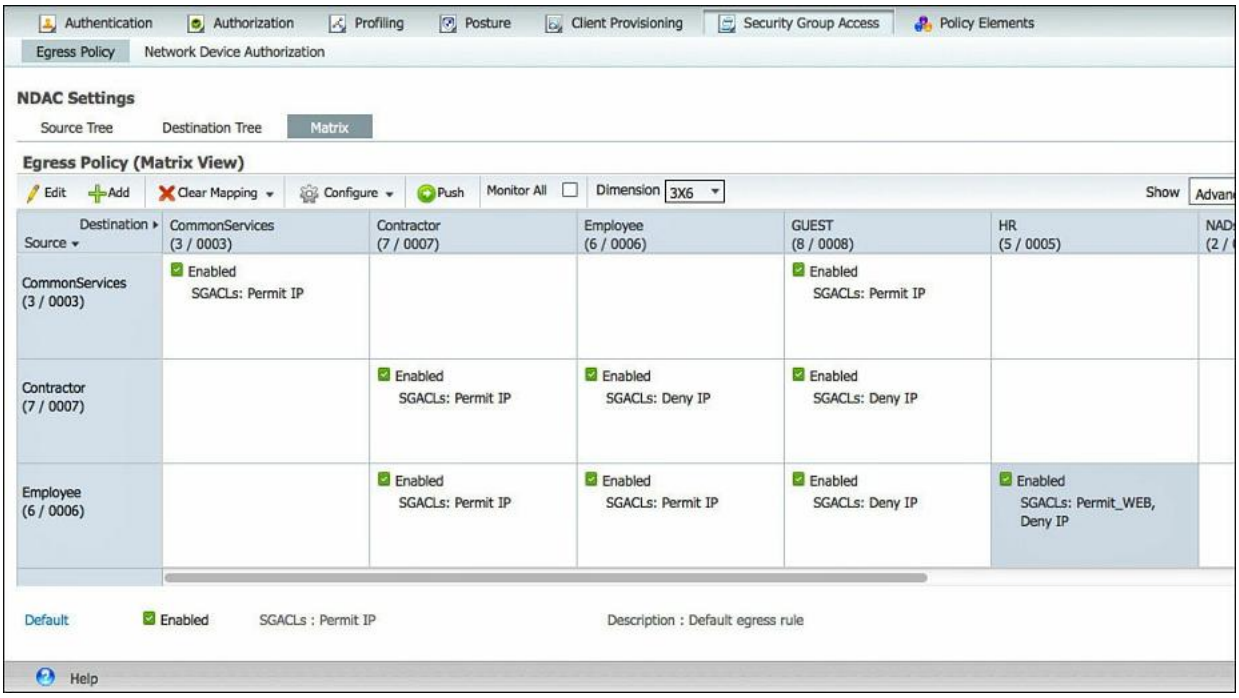


Figure 18-24 SGACL egress policy: matrix view.

The box that is highlighted in gray shows that any traffic with a source of the Employee SGT (6) attempts to reach a destination with the HR SGT (5), an SGACL named Permit_WEB will be applied along with a catch-all of Deny IP. The contents of the Permit_WEB ACL are displayed in Figure 18-25, where you can see that only HTTP and HTTPS are permitted.

https://t.me/learningnets

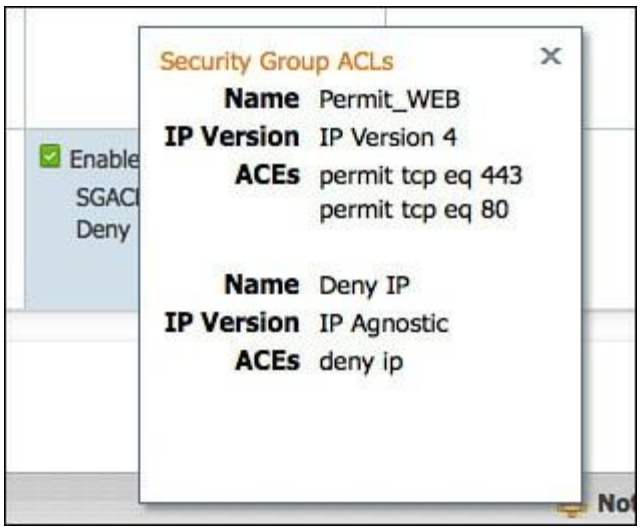


Figure 18-25 Permit_WEB SGACL contents.

As you can see with Figures 18-24 and 18-25, the resulting ACL would be to permit HTTP (tcp/80) and HTTP (tcp/443) and deny all other traffic. This traffic is applied egress of the switch where the SGACL is configured. In this case it is applied at the Nexus 7000 in the data center, as traffic attempts to reach the HR server.

This form of traffic enforcement can provide a tremendous savings on the complexity and number of ACEs to maintain. There is a general formula to see the savings:

(# of sources) * (# of destinations) * permissions = # of ACE's

With a traditional ACL on firewall:

4 VLANs (src) * 30 (dst) * 4 permission = 480 ACE's

Per source IP on port using dACL:

1 group (source) * 30 (dst) * 4 permission = 120 ACE's

With SGACLs, the number of ACEs is a magnitude smaller:

4 SGT (src) * 3 SGT (dst) * 4 permission = 48 ACE's

There are two main conceptual paths in which to deploy SG-ACLs: North-South and East-West. North-South refers to the use case of a user or device being classified at the access layer, but with 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 will be dropped if it tries to reach a server with financial data.

East-West refers to the use case of an SGACL protecting resources that exist on the same switch. For example, with a development server and a production server on the same Nexus 5000 Series switch in the data center, an SGACL can 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 can 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.

Figure 18-26 illustrates the concepts of North-South and East-West traffic flow enforcement.

https://t.me/learningnets

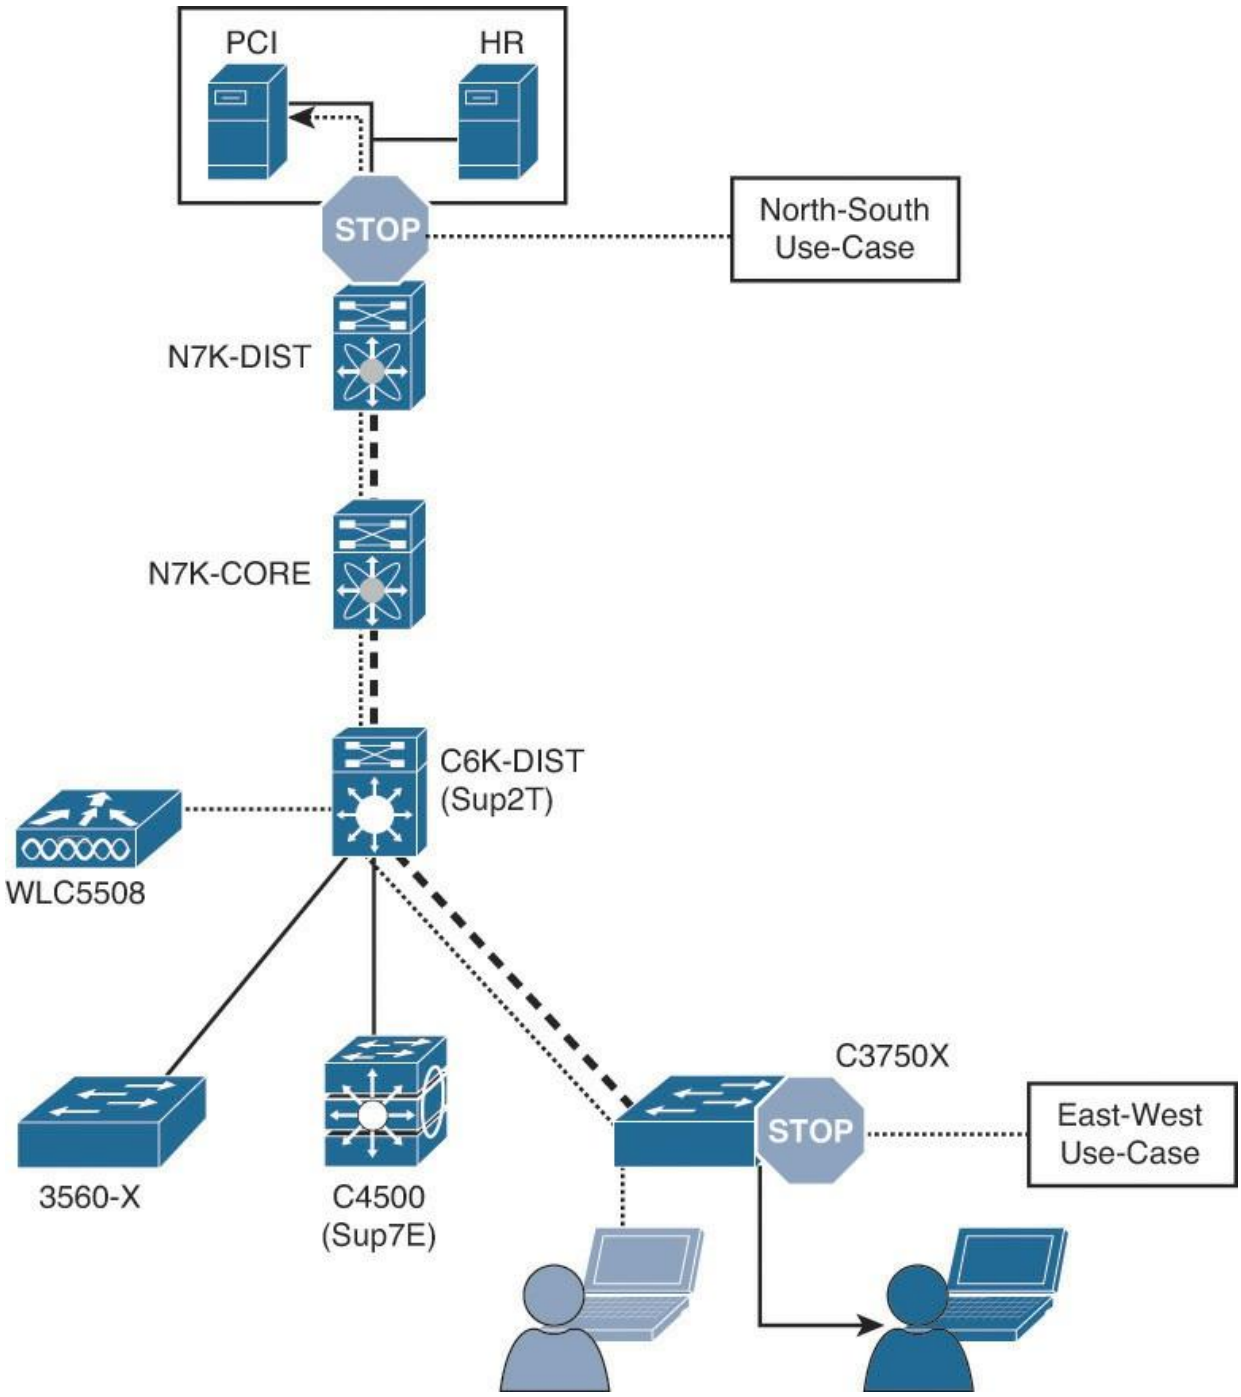


Figure 18-26 North-South versus East-West visually explained.

There is much that can be learned about security group ACLs and the integration of the switches with ISE. However, that would be beyond the scope of the topic for the certification exam, and this book is therefore stopping here with the conceptual overview.

Security Group Firewalls

Some organizations prefer to do the traffic enforcement on the switching infrastructure, and others prefer it to be on a device that was purpose-built to do traffic filtering: a firewall. Cisco has added the capability to enforce traffic on firewalls by implementing the security group firewall (SG-FW). Two types of security group firewalls exist—the ASA-based SG-FW and the router-based SG-FW. This makes sense because the routers use a zone-based firewall (ZBF) and the ASA does not.

https://t.me/learningnets

Security Group Firewall on the ASA

Beginning with ASA version 9.0, the ASA firewall added SG-FW functionality. ASDM supports the full configuration, and therefore the ASA is the only SG-FW that has a GUI (as of the writing of this book).

The SG-FW 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 into the decision. As you can see in Figure 18-27, there is a new "Security Group" column in the Source and Destination Criteria sections.



Figure 18-27 ASDM firewall 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 Generation and the Aggregation Services Router (ASR) have a very powerful ZBF capability.

The Cisco ISR's began support of SG-FW as of version 15.2(2)T, and the Cisco ASR 1000 added support of the SG-FW as of IOS-XE version 3.4. Both routers are capable of running the SXP, and the ASR has native tagging capabilities.

MACSec

In years long passed, when Wi-Fi was first being introduced into the consumer and corporate space, security concerns were raised. The concerns were around sensitive data being transmitted through the air, without any level of confidentiality. The temporary answer to this 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 as a wired network.

Well, the fun part is that 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. So, 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. What was the answer? Should we use IPSec on every endpoint to every other endpoint, encrypting the entire communication from end-to-end? If we did that, how do we provide strong levels of quality of service (QoS) if we cannot see the content of a packet, and how would we glean into the packet to ensure security with all the security tooling that we have invested in? We would be simply encrypting both good and bad traffic across the network if we used end-to-end IPSec.

The answer to provide Wireless Equivalency, and a viable alternative to end-to-end IPSec, was to layer on the confidentiality and integrity using IEEE 802.1AE (MACSec). MACSec provides Layer-2 encryption on

https://t.me/learningnets

the LAN between endpoints and the switch, as well as between the switches themselves. Figure 18-28 illustrates the concept of the hop-by-hop encryption.

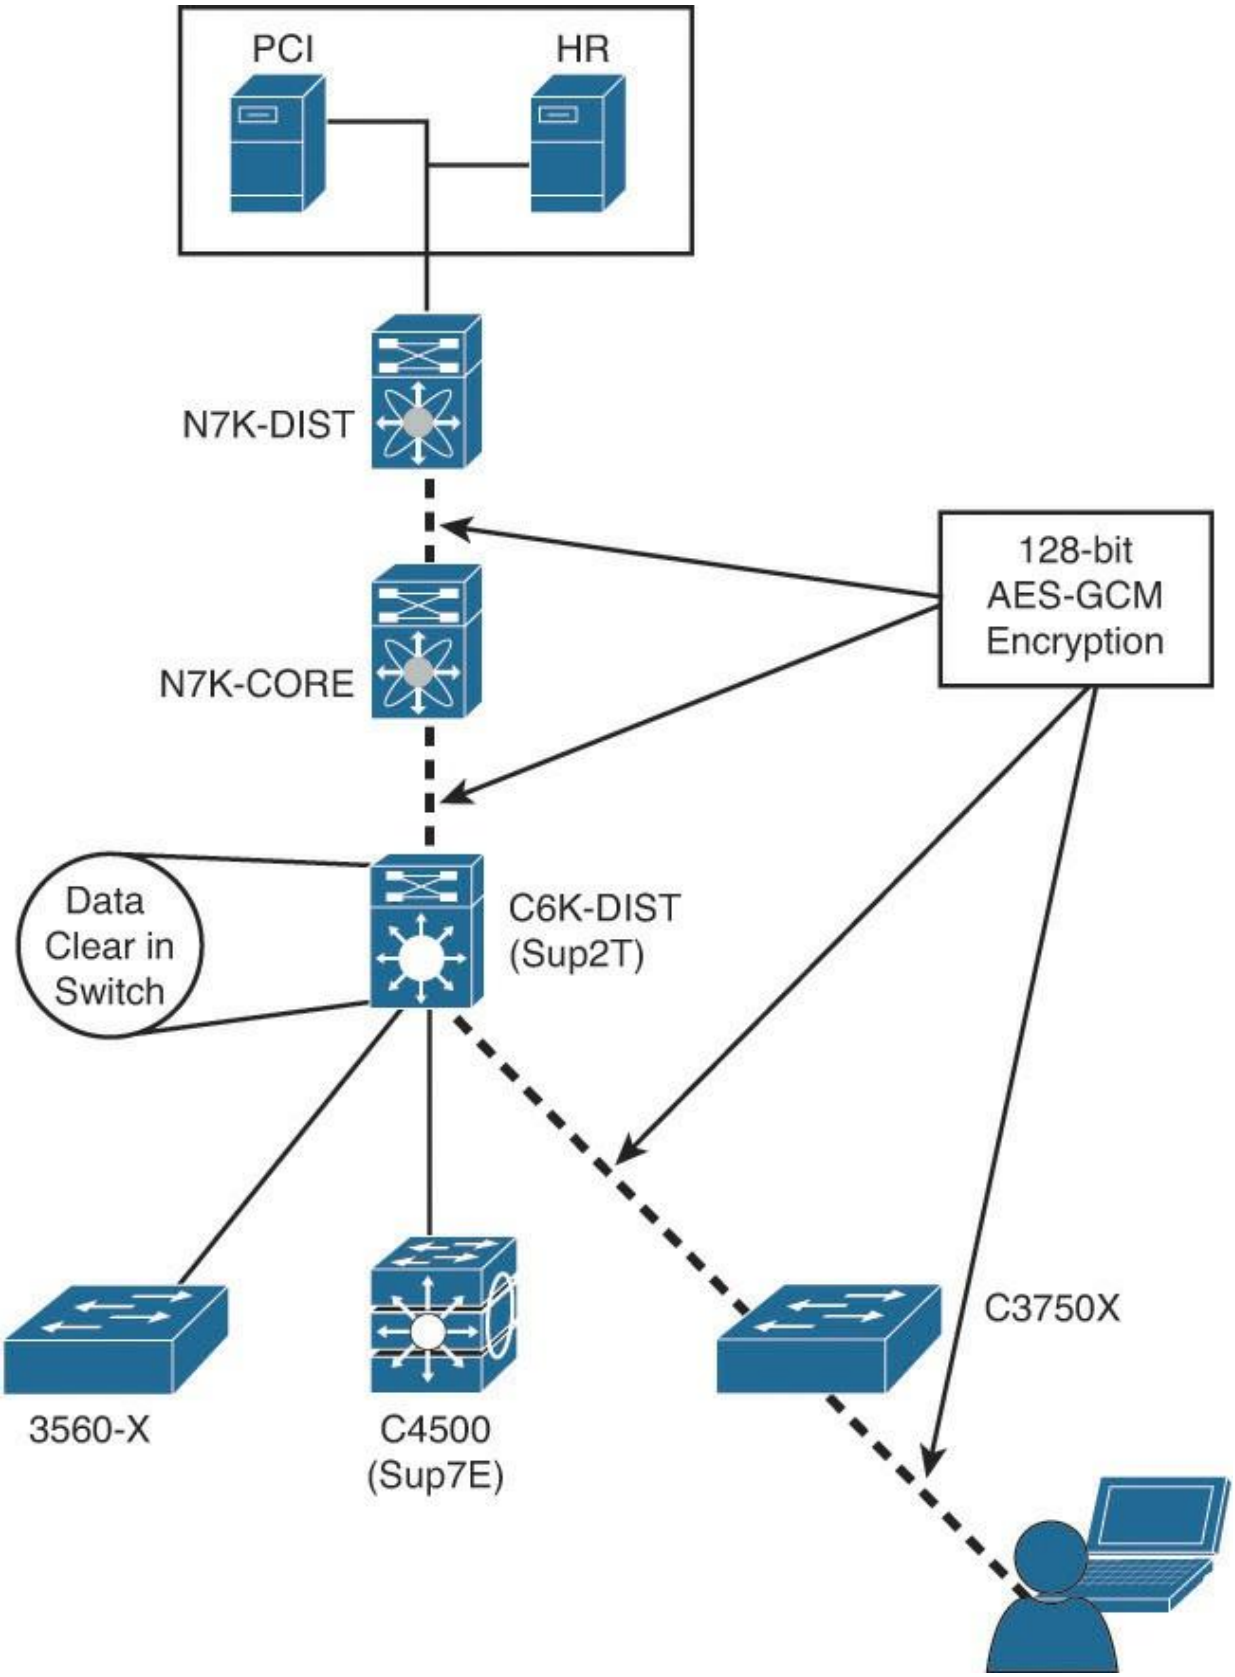


Figure 18-28 MACSec Layer-2 hop-by-hop encryption.

MACSec (IEEE 802.1AE) provides Layer-2 encryption on the LAN. The encryption also encapsulates and

https://t.me/learningnets

protects the CMD field, which carries the SGT as described the earlier in this chapter.

As of the time this book was published, there are two keying mechanisms available, both using 128-bit AES-GCM (Galois/Counter Mode) symmetric encryption that is capable of line-rate encryption and decryption for both 1Gb and 10Gb Ethernet interfaces and that provides replay attack protection of 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, whereas MKA is going to be the industry standard and is currently used between endpoints and Cisco switches.

Downlink MACSec

Downlink MACSec is the term used to describe the encrypted link between an endpoint and the switch. The encryption between the endpoint and the switch is handled by the MKA keying protocol. This requires a MACSec-capable switch (such as a Cisco Catalyst 3750-X) and a MACSec-capable supplicant on the endpoint (such as the 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.

The Cisco switch has the ability to force encryption, make it optional, or force nonencryption; that setting can be configured manually per port (not very common) or dynamically as an authorization result from ISE (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 18-29 shows the MACSec policy within an authorization profile on ISE. You can notice at the bottom, that the attribute sent to the switch is **cisco-av-pair=subscriber:linksec-policy**, followed by the policy itself. The choices are **Must-Secure**, **Should-Secure**, and **Must-Not-Secure**. Example 18-12 shows these options on the switch CLI, and Table 18-2 displays the resulting policy based on the supplicant policy and switch policy.

https://t.me/learningnets

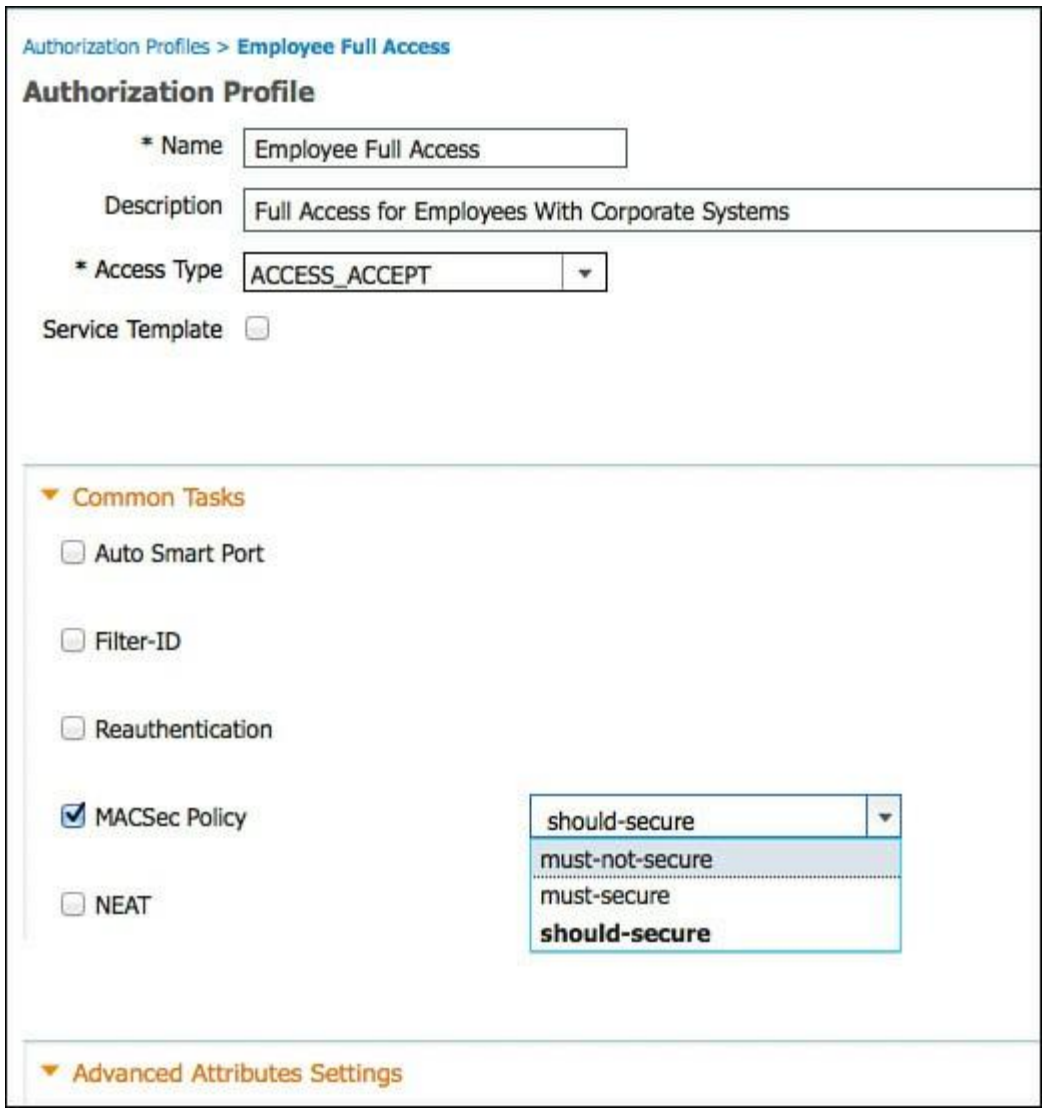


Figure 18-29 Authorization profile with MACSec policy.

Example 18-12 MACSec Policy Switch CLI

Click here to view code image

```
C3750X(config-if)#authentication linksec policy ?
  must-not-secure  Never secure sessions
  must-secure      Always secure sessions
  should-secure    OPTIONALLY secure sessions
```

https://t.me/learningnets

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
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

Table 18-2 Resulting MACSec Policies

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 of **Should-Secure**.

Switch Configuration Modes

A few configurations on the 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.

- **Single Mode**—MACsec is fully supported in single-host mode. In single-host mode, only a single MAC 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 can be on the Data domain, and another endpoint can be on the Voice domain. MACsec is fully supported in MDA host mode. If both endpoints are MACsec capable, each will be secured by its own independent MACsec session. If only one endpoint is MACsec capable, that endpoint can be secured while the other endpoint sends traffic in the clear.

https://t.me/learningnets

- **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.
- **Multihost Mode**—Although MACSec usage with this mode technically might be possible, it is not recommended. With multihost mode, the first endpoint on the port authenticates, and then any additional endpoints will be permitted onto the network via the first authorization. So, MACSec would work with the first connected host, but no other endpoint's traffic would actually pass because it would not be encrypted traffic.

Example 18-13 shows a switch interface configuration for MACSec-enabled endpoints. The example is using the default MACSec policy of **Should-Secure** and therefore the default setting is not displayed.

Example 18-13 Switch Interface Configuration for MACSec

Click here to view code image

```
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 attribute within the authorization profile (the result of an authorization). To add this result to an authorization profile, do the following:

Step 1. Navigate to **Policy > Policy Elements > Results**.

Step 2. Select **Authorization Profiles**.

Step 3. Edit an authorization profile to which you would like to add MACSec.

Step 4. Under **Common Tasks**, scroll down to **MACSec Policy**.

Step 5. Select **must-secure**, **should-secure**, or **must-not-secure**.

Step 6. Click **Submit** or **Save** to save the change.

Uplink MACSec

Uplink MACSec is the term used to describe encrypting the link between the switches with IEEE 802.1AE. At the time this book was written, the switch-to-switch encryption keying mechanism used Cisco's proprietary SAP instead of MKA, which is used with the downlink MACSec. The encryption is still the same AES-GCM-128 encryption used with both uplink and downlink MACSec.

https://t.me/learningnets

Uplink MACSec can be achieved manually or dynamically. Dynamic MACSec requires 802.1X between the switches and is covered in the NDAC section. For this section, we will focus on manual mode.

Manually Configuring Uplink MACSec

This method of MACSec is perfect to layer on top of the Manual SGTs configured earlier in this chapter. It will encrypt the inter-switch links. Start by reexamining the configuration of our uplink interface as we had it configured earlier in this chapter. Example 18-14 shows the uplink configuration.

Example 18-14 Uplink Configuration

Click here to view code image

```
C3750X# 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 18-13, our uplink between the 3750-X and the 6500-Sup2T is set up to use manual keying, without any encryption at all, and to apply SGTs to the frames. Now we will layer encryption on top of this to provide confidentiality and integrity of the SGTs and the data. Figure 18-30 depicts the relevant infrastructure configuration used for this example.

https://t.me/learningnets

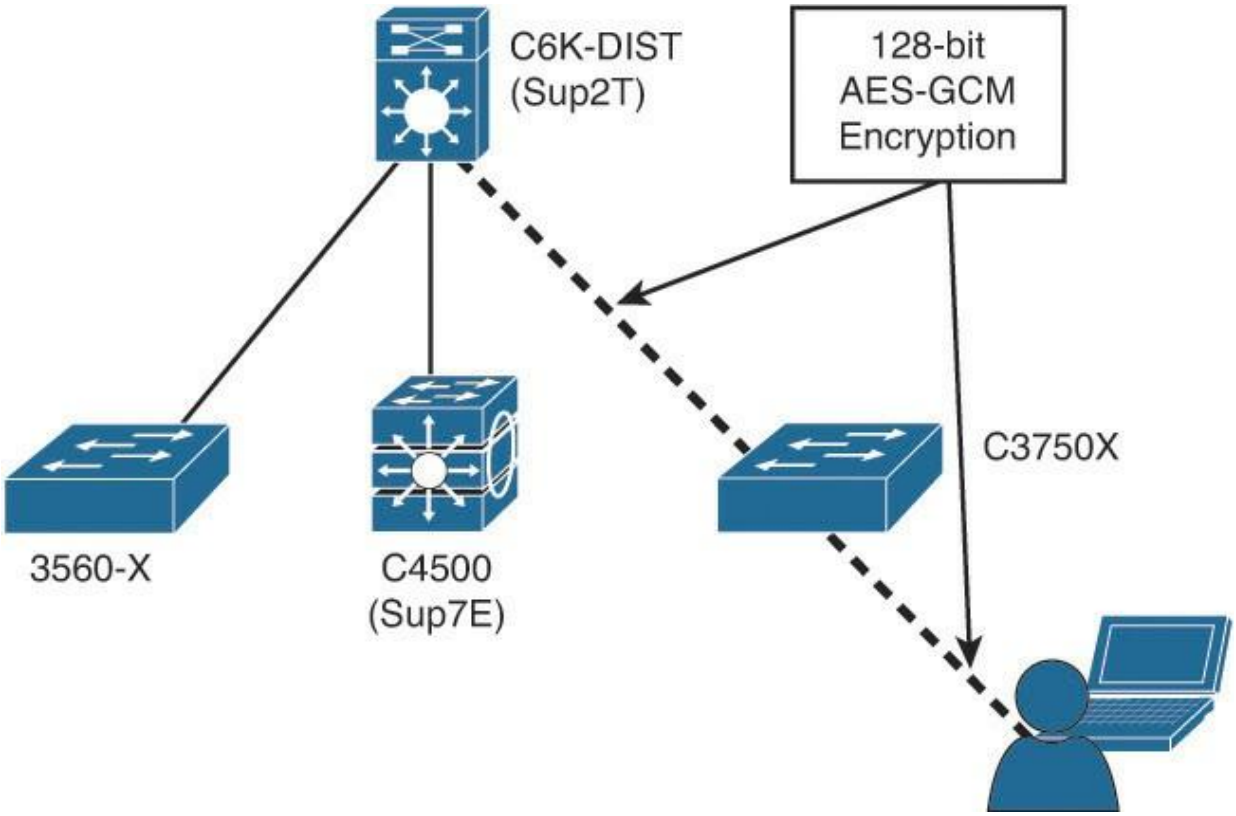


Figure 18-30 Adding MACSec to the uplink.

Step 1. From interface configuration mode, enter **cts manual**.

Step 2. Enable encryption with the **sap pmk** pairwise-master-key **mode-list gcm-encrypt** command.

The Pairwise Master Key (PMK) should be a hexadecimal value configured to be the same on both sides of the link. 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.

Step 3. Add the **sap pmk** pairwise-master-key **mode-list gcm-encrypt** command to the other side of the link.

Example 18-15 displays the sample configuration steps, and Example 18-16 shows the final configuration for the uplink port on the 3750-X.

Example 18-15 Adding Encryption to the Uplink Interface

Click here to view code image

```
C3750X#conf t
Enter configuration commands, one per line. End with CNTL/Z.
C3750X(config)#int Ten1/1/1
C3750X(config-if)#cts manual
C3750X(config-if-cts-manual)#sap pmk 26 mode-list gcm-encrypt
C3750X(config-if-cts-manual)#end
C3750X#
```

Example 18-16 Final Configuration for Uplink Interface

Click here to view code image

```
C3750X#sho run int ten1/1/1
Building configuration...
```

https://t.me/learningnets

```
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 00000000000000000000000000000000000000000000000000000000000000000000000026
mode-list gcm-encrypt
 no macro auto processing
end
```

Verifying the Manual Configuration

You can validate that the manual encryption on the uplink was successful with the **show cts interface** command, as shown in Example 18-17. SAP status is the status of the encryption, and the example shows that SAP succeeded, the pairwise cypher is using gcm-encrypt, and replay protection is enabled.

Example 18-17 Output of **show cts interface** Command

Click here to view code image

```
C3750X#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
```

https://t.me/learningnets

```
          port auth fail:          0

      L3 IPM: disabled.

C3750X#
```

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 18-3 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	Goal of TrustSec	605
Section	SXP modes (speaker and listener)	613

Table 18-3 Key Topics for Chapter 18

Define Key Terms

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

TrustSec

security group tag (SGT)

classification

transport

Security Group Exchange Protocol (SXP)

North-South

East-West

https://t.me/learningnets

Chapter 19. Posture Assessment

This chapter covers the following exam topics:

- **Posture Service Overview**
- **Posture Flow**
- **Agent Types**
- **Posture Conditions**
- **CoA with Posture**
- **Configuring Posture**

Over the last 18 chapters, we've highlighted a number of security technologies and policies that help to protect your network, leveraging Cisco's Identity Services Engine (ISE) as the cornerstone of our approach. Up to this point, most of ISE's decisions have been made based on the user or the endpoint—not the software that is running on the endpoint. What level of security do we truly have on our network if a virus or malware infects the legitimate endpoint of an authorized user? What if an authorized user is running unapproved software, such as a network sniffing tool or penetration tool on the network to look for network vulnerabilities? In both of these situations, the security of the network could be greatly compromised.

Posturing is ISE's solution to both of these situations. The posture of the endpoint is a function of what software is present on the endpoint. Whether it is an authorized user that is knowingly or unknowingly running a piece of software that could compromise network security, ISE has the ability to appropriately react to limit network access until the issue is remedied. Throughout this chapter, we discuss how to configure ISE to determine the posture of the endpoint.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Posture Service Overview	1
Posture Flow	2
Agent Types	3
Posture Conditions	4-6
CoA with Posture	7
Configuring Posture	8-10

Table 19-1 "Do I Know This Already?" Section-to-Question Mapping

https://t.me/learningnets

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. The Posture Service is comprised of which of the following functional components? (Select three.)
   - **a.** Profiling
   - **b.** Client provisioning
   - **c.** Authorization policy
   - **d.** Mobile device managers
   - **e.** Access lists
   - **f.** Guest Services
   - **g.** Posture Policy
2. What are the three possible posture outcomes following the initial connection to the network?
   - **a.** Location, Location, and Location
   - **b.** Routes, Translations, and Permissions
   - **c.** Authentication, Authorization, and Accounting
   - **d.** Compliant, Noncompliant, and Unknown
3. Which is a benefit of a NAC web agent versus a persistent agent?
   - **a.** The web agent provides enhanced remediation techniques.
   - **b.** The web agent does not require Administrator privileges to install.
   - **c.** The web agent provides additional firewall functionality for the endpoint.
   - **d.** The web agent can provide a greater number of Posture conditions.
4. True or False? The Process Check posture condition is supported on all NAC agent types.
   - **a.** True
   - **b.** False
5. The File condition for Posture does which of the following?
   - **a.** Checks 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 the above
6. True or False? Cisco offers periodic Posture Elements updates.
   - **a.** True
   - **b.** False
7. The CoA process is used for which of the following?
   - **a.** To force an endpoint to reauthorize following a change in status
   - **b.** Following a change of posture compliancy from the NAC agent
   - **c.** Only after a NAD has terminated an endpoint's connection

https://t.me/learningnets

d. a and b

e. b and c

f. a, b, and c

8. When configuring the Client Provisioning Policy, you can elect each of the following except which?

a. NAC Agent Configuration

b. Network Supplicant Provisioning

c. Access list

d. Profile

9. Remediation is a process by which of the following occurs?

a. An endpoint that is not compliant with security policy can become 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.

10. Which remediation type is available on a Macintosh OS X endpoint?

a. Automatic Launch Program Remediation

b. Manual Antispyware remediation

c. File Remediation

d. Manual Antivirus Remediation

Foundation Topics

Posture Service Overview

ISE only has the ability to "see" the limited amount of network traffic that the endpoint sends onto the network and is actively forwarded to ISE, usually with the assistance of other network devices (specifically, network access devices and device sensors). Using the information it receives, ISE can make an educated guess as to the type of the device (via profiling) and adjust the level of access given to the endpoint accordingly. However, the information sent to ISE will likely not include the latent virus or piece of malware running on the endpoint that is waiting to wreak havoc on the network at the right moment. ISE will likely not know about this attack on the network until it is too late.

ISE's posture service enables ISE to have a window into the software that is running on the device. Using a number of condition and policy constructs on ISE, a network administrator can define the requisite security policy for her corporation. As the endpoint attempts to join the network, ISE can confirm whether the endpoint is in compliance with the configured security policy. If the endpoint is compliant, the endpoint will be allowed the commensurate level of network access. However, if the endpoint is not compliant, ISE can immediately react and provide a limited amount of network access to allow the endpoint to remediate its noncompliance.

What approach does ISE use to determine whether an endpoint is compliant? To get the required insight as to what is running on the endpoint, ISE leverages a network admission control (NAC) agent that is installed on the endpoint. The NAC agent performs the necessary reconnaissance of the endpoint and does ISE's bidding to check the endpoint's compliance to the security policy. If the endpoint is not compliant, it is the job of the NAC agent to inform ISE of the issue. Without the NAC agent, ISE would be unable to do a proper posture assessment of the endpoint.



The posture service has three major functional components:

https://t.me/learningnets

- **Client Provisioning**—The Client Provisioning portion of the posture service installs the NAC agent onto the endpoint.
- **Posture Policy**—The Posture Policy defines what is compliant and what is noncompliant about the endpoint. This compliance could be based on the version of antispyware or antimalware software present on the endpoint and whether the anti-X definitions are up-to-date.
- **Authorization Policy**—Once we know whether the endpoint is compliant, we can determine the level of access that should be granted to the endpoint. Typically, a compliant endpoint is given an increased level of access while a noncompliant endpoint is cordoned off on the network until it can remediate its noncompliant stance.

Posture Flow

Much like the Profiler and Guest Services function of ISE, posturing is another optional component a network administrator can choose to deploy in a Cisco ISE network. If implemented correctly, the Posture service of ISE can greatly increase the security of a corporate network.

One of the key components of ISE that makes posturing possible is the change of authorization (CoA) function, which has been discussed numerous times in this book. To review, a CoA is the capability of ISE to instruct the network access device (NAD) to reauthorize the endpoint. This CoA normally takes place after a condition that may have changed the endpoint's network-worthiness or level of access. For instance, if an endpoint is determined to be a mobile device via a Profiling process, a more restrictive level of network access can be given to the endpoint.

How does posturing leverage the CoA function? As an endpoint first connects to the network, an initial network authentication and authorization begin. This provides at least a minimal amount of network access to the endpoint. This initial level of network access will be sufficient for the NAC agent that resides on the endpoint to download the security policy from ISE.



After downloading the security policy, the NAC agent can then proceed to run a posture assessment of the endpoint, reporting the result to ISE. This initial connection will return with one of three results:

- **Compliant**—After downloading the security policy from ISE, the NAC agent on the endpoint determined that the software assessment on the endpoint adheres to the security policy, reporting this information to ISE.
- **Noncompliant**—After downloading the security policy from ISE, the NAC agent determined that there is something on the endpoint that is in violation of the defined security policy, reporting this information to ISE.
- **Unknown**—The endpoint failed to report a posture assessment to ISE.

Whatever the outcome from the posture assessment—Compliant, Noncompliant, or Unknown—ISE can now make a determination as to the new authorization that should be granted to the endpoint. To effect this new level of access, ISE will send a CoA to the NAD on behalf of the endpoint. The NAD will then reauthorize the endpoint, implementing the updated security policy. Figure 19-1 shows a complete overview of the ISE authentication and authorization flow.

https://t.me/learningnets

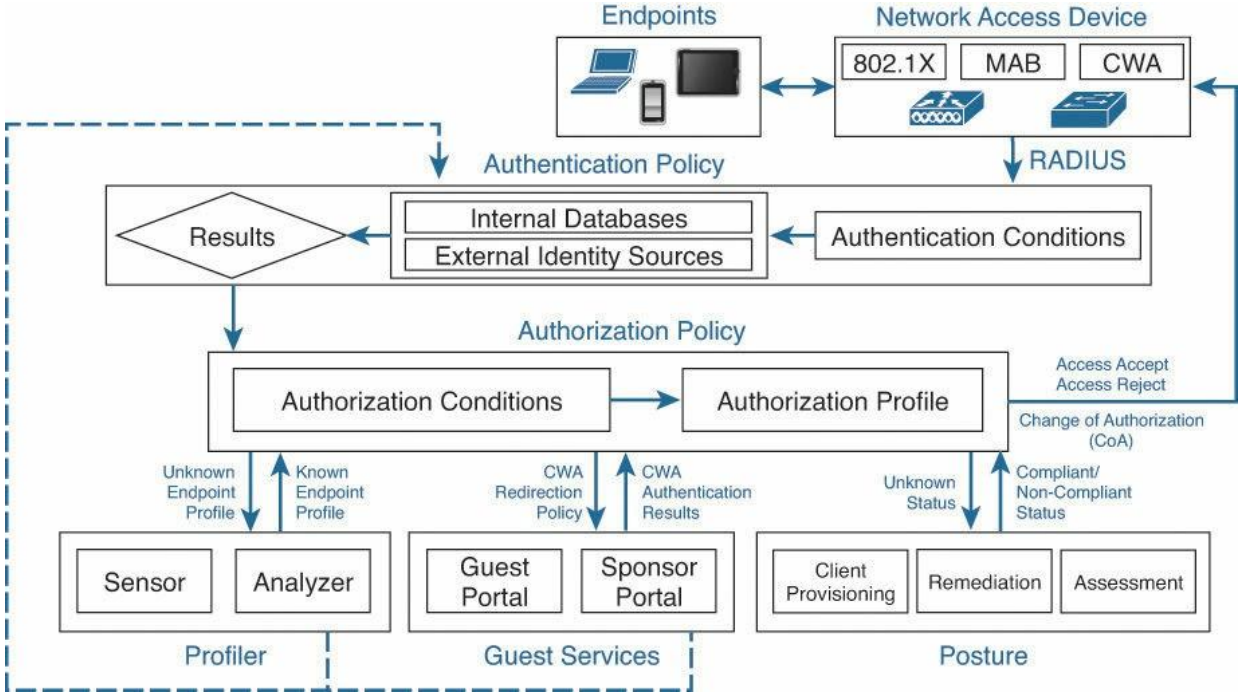


Figure 19-1 ISE authentication and authorization flow.

Typically, if an endpoint is found to be in compliance with the security policy, it will be able to access a larger amount of network resources. For instance, if the endpoint is compliant, it might be able to access all network resources, including file shares, web servers, email, and so on. If an endpoint is given a posture assessment of Noncompliant, the endpoint might be granted a minimal level of network access to remediate its noncompliance. This can allow the endpoint to access a webpage to download the latest antivirus and antimalware software and definitions or to download the appropriate corporate software packages. If an endpoint is found to be Unknown from a posture standpoint, this implies that a NAC agent is likely not installed or was unable to communicate a posture assessment back to ISE, with the former being the more likely scenario. In this case, ISE would likely force the endpoint to a web portal that would assist the user in installing the required NAC agent.

Agent Types

The NAC agent plays a critical role in the posture service on ISE. Without this agent and the assessments it provides, posturing would not be possible. However, not all NAC agents are created equal. There are two key types of NAC agents.

The first NAC agent type is the persistent agent, which is available for Mac and Windows computers. This agent type is typically used on managed endpoints—those endpoints that are maintained by a centralized IT department. This agent is typically installed from the web browser or via an MSI package and provides the greatest level of functionality of the NAC agents. However, to achieve this increased level of functionality, Administrator rights are required to install and update the agent software.

With a persistent agent, the user is guided through the posture service. The persistent agent also supports posture remediation and will inform the PSN about an IP address change that occurs while a session is active. The Cisco Identity Services Engine supports automatic NAC agent updates.

The second NAC agent type is the web agent. As the name implies, this agent is installed via a webpage, leveraging ActiveX or Java. As long as the endpoint has one or both of these software packages installed, the web agent should be able to function. Because this agent is temporary, it is well suited for unmanaged PCs and guests and does not require administrator access. However, without the increased level of operating system (OS) access provided with Administrator privileges, it offers fewer remediation options than the persistent agent. The web agent is also capable of refreshing any changes in Dynamic Host Configuration Protocol (DHCP) addresses, often seen where authorization policy enforces a change in a virtual local area

https://t.me/learningnets

network (VLAN).

The list of supported remediation approaches is found in Table 19-2.

Remediation Type	NAC Agent for Windows	Web Agent for Windows	NAC Agent for Macintosh
File Remediation	Supported	Supported	—
Link remediation (manual)	Supported	Supported	Supported
Link remediation (automatic)	Supported	Not supported	Not supported
Antivirus remediation (manual)	Supported	Not supported	Supported
Antivirus remediation (automatic)	Supported	Not supported	Not supported
Antispyware remediation (manual)	Supported	Not supported	Not supported
Antispyware remediation (automatic)	Supported	Not supported	Not supported
Launch Program remediation (manual)	Supported	Not supported	—
Launch Program remediation (automatic)	Supported	Not supported	—
Windows Update remediation (manual)	Supported	Not supported	—
Windows Update remediation (automatic)	Supported	Not supported	—
Windows Server Update Services remediation (manual)	Supported	Not supported	—
Windows Server Update Services remediation (automatic)	Supported	Not supported	—

Table 19-2 NAC Agent-Supported Remediation Types

Posture Conditions

The purpose of posturing is to determine which software or software configurations reside on an endpoint that might soon be given access to your network. As we just discussed, this posture assessment is the job of the NAC agent, but as a network administrator, you have to define the security policy the NAC agent will adhere to as it does the assessment. The manner by which a network administrator defines this policy is by using posture conditions.

Posture conditions are used to check specific attributes on the client system. Because Cisco has partnered with many vendors over the years in developing its NAC agents, they have developed a large array of automated rule sets for more than 350 partner applications, including Windows updates (auto updates, hotfixes, and service packs), antispyware, antivirus, Windows Registry, file, and service data. This helps to

https://t.me/learningnets

simplify the configuration of these conditions within ISE.

There are several categories for assessment options. These conditions can be used alone or in combination with other conditions:

- **File**—A condition that checks the existence of a file, the date of a file, and the versions of a file on the client.
- **Registry**—A condition that checks for the existence of a Registry key or the value of the Registry key on the client.
- **Application**—A condition that checks whether or not an application (process) is running on the client.
- **Service**—A condition that checks whether a service is running on the client.
- **Dictionary**—A condition that checks a dictionary attribute with a value.

See Table 19-3 for the assessment options that are supported for each NAC agent.

NAC Agent for Windows	Web Agent for Windows	NAC Agent for Macintosh OS X
Operating System/ Service Packs/Hotfixes	Operating System/Service Packs/Hotfixes	—
Process Check	Process Check	—
Registry Check	Registry Check	—
File Check	File Check	—
Application Check	Application Check	—
Antivirus Installation	Antivirus Installation	Antivirus Installation
Antivirus Version/Antivirus Definition Date	Antivirus Version/Antivirus Definition Date	Antivirus Version/Antivirus Definition Date
Antispyware Installation	Antispyware Installation	Antispyware Installation
Antispyware Version/ Antispyware Definition Date	Antispyware Version/ Antispyware Definition Date	Antispyware Version/ Antispyware Definition Date
Windows Update Running	Windows Update Running	—
Windows Update Configuration	Windows Update Configuration	—
WSUS Compliance Settings	WSUS Compliance Settings	—

Table 19-3 Posture Assessment Options

There are several ways that you can leverage the previously mentioned supported posture assessment options to keep your network secure. Several common approaches include

- **Windows update verification**—Microsoft frequently releases updates to its OS and Productivity Suites to patch any known vulnerabilities that were found by its millions of users. This posture condition can help ensure that the endpoint has the appropriate service packs installed and the appropriate patch level.
- **Antivirus application verification**—Maintaining the correct antivirus software can really help lower the risk profile of an endpoint. Many corporations spend millions of dollars a year to have the best antivirus solution on their endpoints. The cost of the software is miniscule compared to the

https://t.me/learningnets

public relations impact, manpower cost of a possible compromise, loss of data, and loss of time that can result if a virus were to take down even a small portion of their endpoints and network. Depending on the size of the company and the types of endpoints that might exist, looking for a single vendor's antivirus on the endpoint may not be practical. This posture condition can also look for any antivirus software, allowing a level of flexibility for the end user.

- **Virus definition verification**—Much like the Windows update verification, it is equally, if not more so, important to keep virus definitions up-to-date. With the quick exchange of information that occurs via social media and the proliferation of cyber-warfare and cyber-espionage, it can be a matter of seconds that a newly found vulnerability can reach the far edges of the Internet. If a new virus definition is available, it is very important to incorporate the updated file into your endpoint security profile as quickly as reasonably possible.
- **Windows screensaver password verification**—It is important to lock down access to your endpoints with a password. In absence of a password on your endpoint, a passerby, or possibly even a remote user, can access your computer and place any number of rogue software packages on your computer, including viruses, malware, key-loggers, and so on. With a password, this can completely prevent a user from gaining access to your computer or, at a minimum, act as a deterrent whereby the attacker will move on to an easier target.
- **Registry entry verification**—On Windows devices, the Registry can be like the Crown Jewels of the OS. It is important to protect the integrity and content of this file. If the content, integrity, or availability of this file are compromised in any way, the entire OS can be rendered inoperable, causing large amounts of downtime, lost productivity, and possible loss of intellectual property—not to mention any public relations impact that might result. Also, by verifying the content of certain Registry keys, you can confirm whether installed software is present and whether it is configured adequately to meet the requirements of the corporate security policy.

ISE supports periodic updates to its posture elements. The posture updates can be initiated manually or scheduled to run automatically on a periodic basis. To trigger an update or to configure the updates to run on a schedule, go to **Administration > System > Settings > Posture > Updates** (see Figure 19-2). These updates include predefined checks, rules, and support charts for antivirus and antispyware for both Windows and Macintosh operating systems and operating system information supported by Cisco.

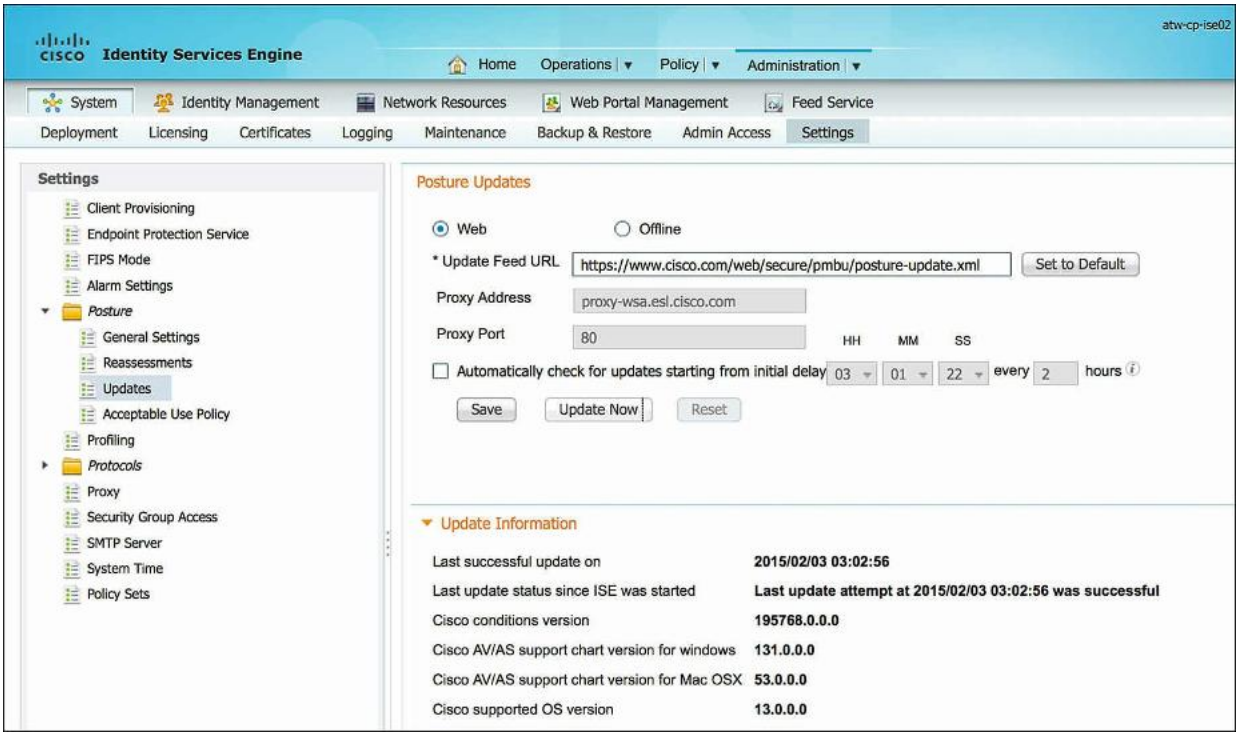


Figure 19-2 Posture updates.

https://t.me/learningnets

CoA with Posture

As mentioned previously in this chapter and as shown in Figure 19-1, ISE uses the CoA feature heavily to facilitate the posture service. When a client first attempts to join the network, it might not have a NAC agent available to perform a posture assessment. Without a NAC agent, ISE will have no mechanism to determine what software is running on the endpoint. Accordingly, the endpoint will not know how to communicate its posture assessment to ISE. For this reason, an endpoint without a posture agent will be assigned an Unknown posture status.

ISE will be configured to appropriately handle the Unknown status by adding a rule to the authorization policy. Within the authorization policy of ISE, we'll configure a rule that has a condition that matches a posture status of Unknown. The authorization profile that is attached to that rule will leverage the common task of redirection to send the endpoint to the Client Provisioning Portal (CPP). As the endpoint attempts to browse to a website, the redirect will trigger, sending the endpoint to the CPP where it can install an appropriate NAC agent.

Now that the previously Unknown endpoint has a NAC agent, a proper posture assessment can be completed. The NAC downloads the requisite posture security policy from ISE and completes the assessment, sending the assessment result to ISE. Whether the endpoint is compliant or noncompliant, ISE sends the CoA to the NAD to signal the endpoint that a new authorization profile will be applied. If the endpoint is compliant, this CoA will likely provide an increased level of access for the endpoint. If the endpoint is noncompliant, additional remediation steps can be required to access the network.

We'll discuss the actual configuration that must occur here shortly.

Configuring Posture



Prior to beginning the implementation of posturing on your network, it is paramount that you do a review of your network security policy. With this review, it is best to highlight all the required checks that must be done to determine whether an endpoint is compliant as well as the types of endpoints you will be assessing. Pay specific attention to the OS running on the endpoints and cross-reference these OSes with the compliancy policy.

As we alluded to previously, a number of conditions might not yet be supported with the current versions of the NAC agent on a particular OS in question (for instance, Linux operating systems do not currently have an agent installation policy within ISE). If a "required" condition is not supported, it might be prudent to come up with an acceptable approach to deal with the deviation. It is best to make sure that all required conditions are supported or understand what approach is going to be taken for unsupported conditions prior to beginning the configuration process. Brainstorm all possible scenarios so as to minimize any unintended outages during the implementation.

When you are comfortable that you have all the required information to configure a complete security policy on ISE, you can begin the configuration process. The first step of this configuration is to review the Posture general settings (see Figure 19-3) at **Administration > System > Settings > Posture > General Settings**. With these settings, you can configure such items as the default posture status as well as remediation timers, reassessments, and posture acceptable use policy (AUP). The configuration of the reassessments and Posture AUP is beyond the scope of the SISAS test and, therefore, is not included in this book.

https://t.me/learningnets

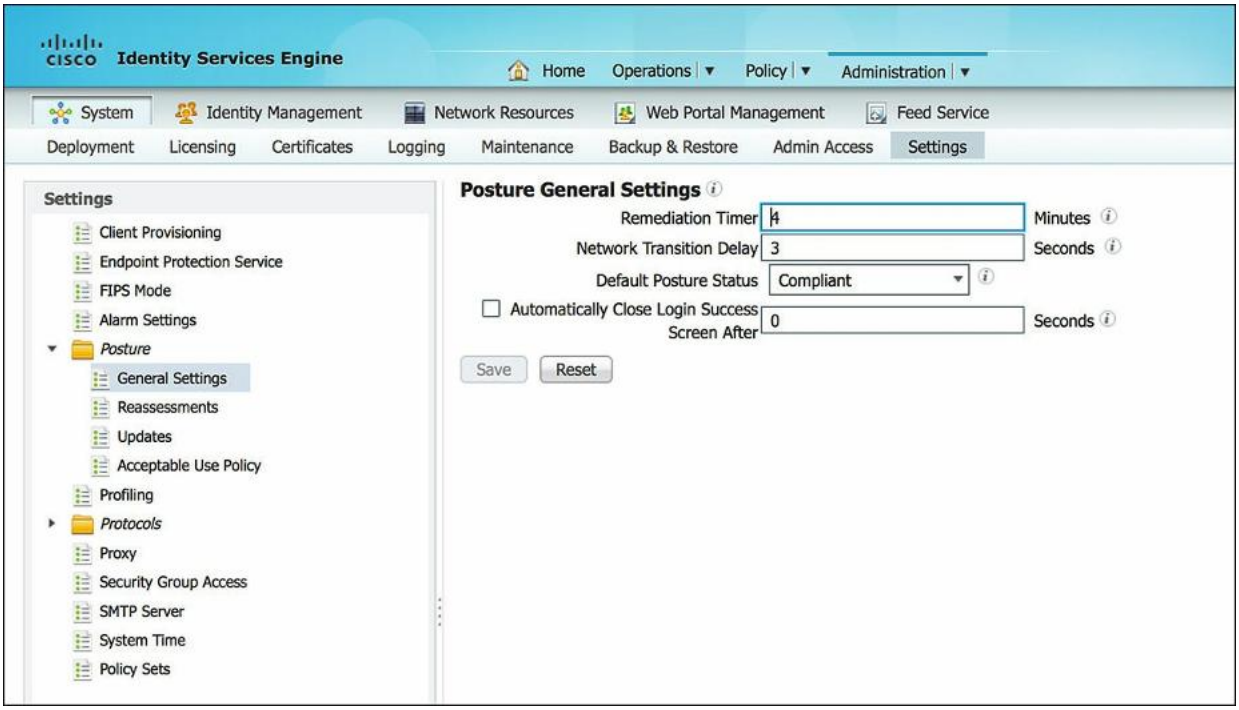


Figure 19-3 Posture general settings.

The general settings also provide, as we mentioned previously, the ability to update the Posture elements from www.cisco.com.

Downloading CPP Resources

Before we start configuring the policy, it is important that we download all the required CPP resources. These CPP resources include any agents or compliance modules that might be required on the endpoints. For posturing, our focus is going to be mostly on the NACAgent or WebAgent types. These resources can be downloaded (see Figure 19-4) by going to **Policy > Policy Elements > Results, Client Provisioning > Resources**. Click the **Add** button and select the required resource, such as **Agent Resources from Cisco site**. This is also the same location whereby you can download the Supplicant Provisioning Wizards that are required for BYOD, as discussed in Chapter 17, "Bring Your Own Device."

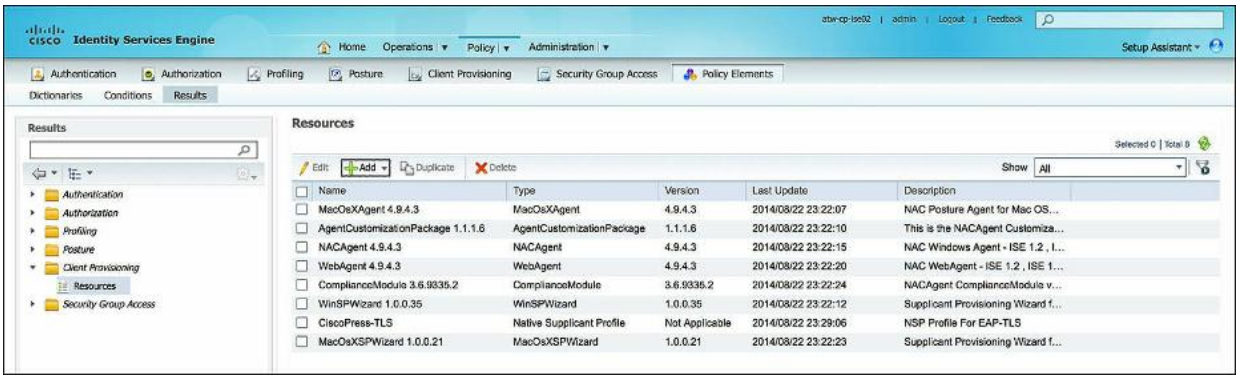


Figure 19-4 Download client provisioning resources.

Now that we have established our posture security policy, updated the Posture Elements, reviewed the Posture general settings, and updated our CPP resources, we are ready to begin our Posture service configuration. In this configuration, much like we did with our authentication, authorization, and Guest Services policies, we will lay some basic groundwork—building blocks if you will—before we start configuring the actual conditions and policies.

https://t.me/learningnets

Client Provisioning Policy

With the appropriate resources downloaded to ISE, we have to define how these resources will be deployed to the endpoints. For instance, as the network administrators, we need to decide which NAC agent we intend to use with each endpoint type and under what conditions. This task of deploying the appropriate NAC agent to the endpoint under the right conditions is defined using the Client Provisioning Policy. As we mentioned near the beginning of this chapter, we might not want to require posturing for guest users—or, if we do, we might choose a nonintrusive, web-based agent.

The Client Provisioning Policy is what we used in Chapter 17 to deploy the Supplicant Provisioning Wizards. This time, we'll define similar policies, or modify existing policies, to deploy the NAC agent to the endpoints. Even if you did not review the BYOD chapter, you will still likely be familiar with the intuitive constructs that were also used in the authentication and authorization policies—the IF-THEN programmatic approach. Each client provisioning rule will have a name, a set of conditions (including identity group, OS, and/or other conditions), and then the resulting NAC agent to deploy.

For the Operating Systems option, select from the predefined OSes, which include Android, Apple iOS, Mac OS X, and Microsoft Windows. Posturing of mobile endpoints are not currently supported using a NAC agent. If posturing of mobile endpoints is required, you might consider leveraging one of the Cisco's Mobile Device Manager (MDM) partners, as we briefly discussed in Chapter 6, "Introduction to Advanced Concepts."

If you want to cover all Windows varieties with one rule, you can select Windows All. If you require more specificity for Windows, you can also drill down into the folder and select a specific version of Windows: Windows 7, Windows 8, Windows 8.1, Windows Vista, or Windows XP.

For the Other Conditions column, you can leave this blank (triggering strictly off of the OS) or be as specific as you would like. For instance, you could trigger based on the authentication method (such as 802.1X, MAB, and so on), AD group membership, or the WLAN/SSID that is being used for access. These are just a few possible scenarios you might consider as you deploy NAC agents to your endpoints.

As you elect the NAC agent that is to be deployed to the endpoint, choosing the appropriate NAC agent type and release version, you can specify whether upgrading to the latest version is mandatory. The NAC agents available to choose from during this selection process are those that were downloaded from Cisco in the "Downloading CPP Resources" section, filtering the options based on the selected OS condition (see Figure 19-5). If you chose to not use the default posture timers under the Posture General Settings, you can choose the specific profile you configured to define various timers. You can also specify the compliance module and customization package. The finished policy is shown in Figure 19-6.

https://t.me/learningnets

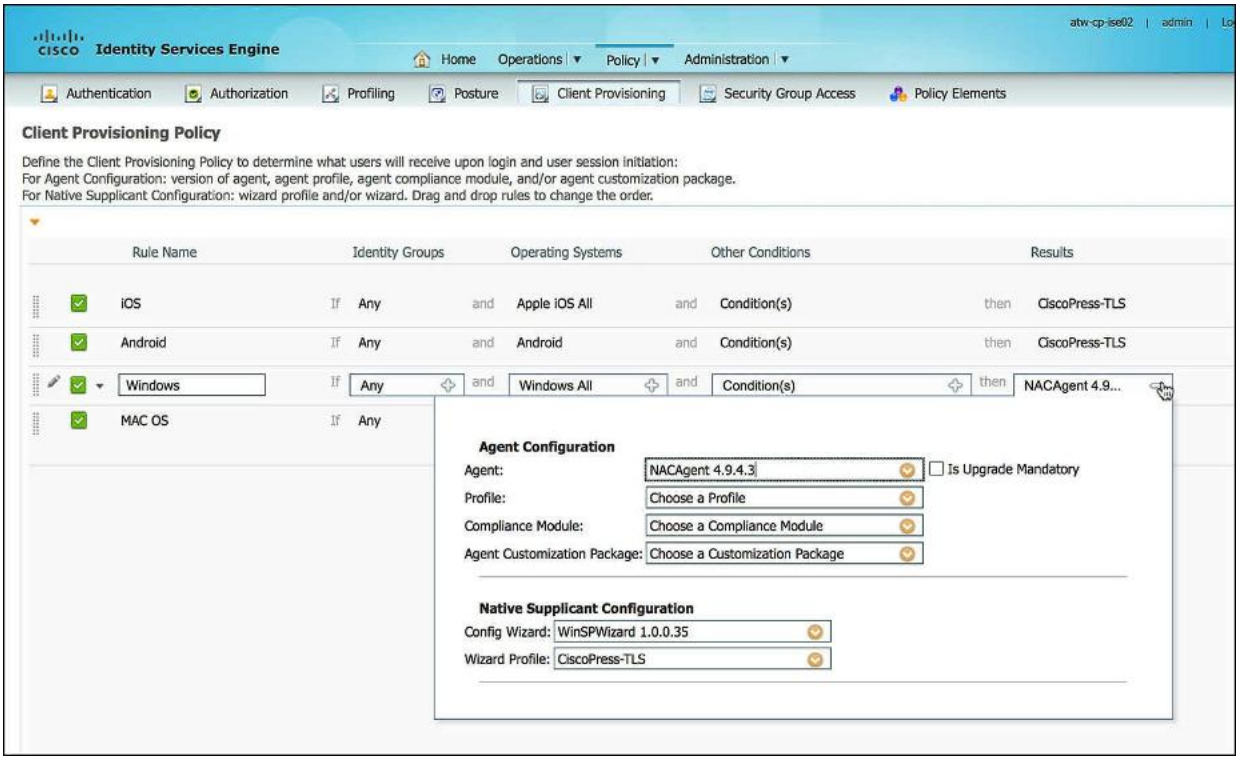


Figure 19-5 Selecting the NAC agent and policy.

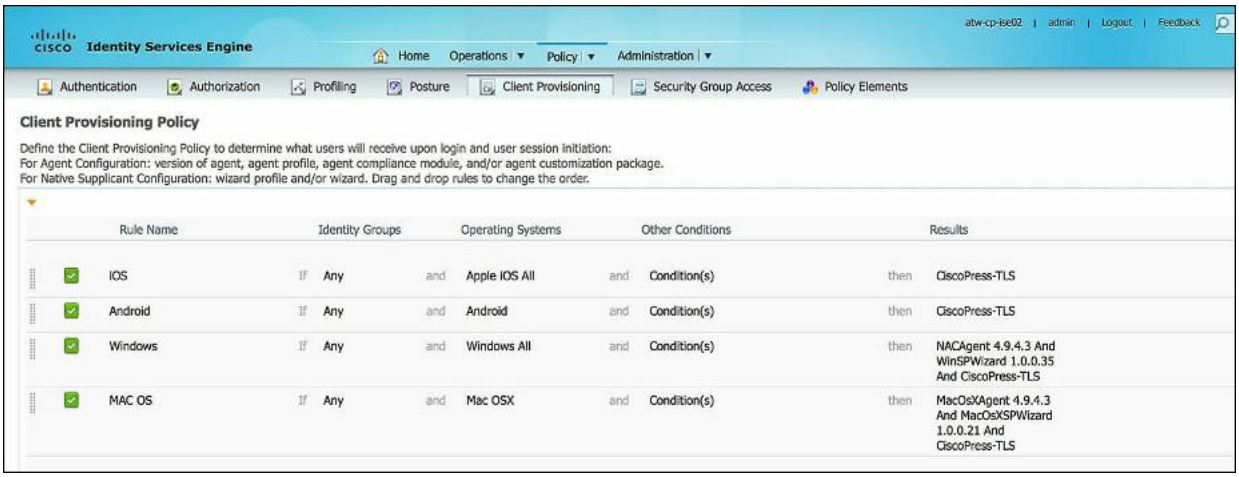


Figure 19-6 Client provisioning policy.

Posture Policy Building Blocks

Everything we have configured thus far with regards to posture has been to get a NAC agent onto the endpoint. We've not yet started the configuration of the actual posture policy—that is which software is or isn't running on the endpoint. These posture policies, similar to other ISE policies, have a name, conditions, and finally a result. The conditions can be a simple or compound condition, leveraging the file, Registry, service, application, or dictionary posture conditions we discussed earlier in this chapter. If these conditions match, we can then move onto the result; in this case, we refer to the result as a requirement. This requirement is somewhat different from other ISE policies because the requirement, too, is written as an expression.

This requirement expression can be read as "Requirement for 'OS' met IF 'posture condition' ELSE 'remediation.'" The remediation in this expression defines the action that must be performed. Earlier in this chapter, we discussed the capabilities of each NAC agent when it comes to remediation opportunities (refer to Table 19-2). Be sure that whichever remediation you select can actually be accomplished in the selected

https://t.me/learningnets

OS.

Let's now try our hand at defining a few posture policies.

Condition

The first policy building block we'll define is the Posture conditions. These conditions will be the test for compliance performed on the endpoint. This is why the NAC agent is required: to perform this reconnaissance on the endpoint and determine which one of these defined conditions holds true on the endpoint. This condition can be to see if particular software is or is not present, such as any antivirus or antispyware software, or a specific version of antivirus or antispyware. To configure Posture conditions, browse to **Policy > Policy Elements > Conditions> Posture**. Once here, you can elect a specific type of condition you want to define. Let's take a look at the AV Compound Condition option.

If you click AV Compound Condition (shown in Figure 19-7), you will see four predefined AV-related compound conditions. Two are for Windows and two are for Mac OS X. For each OS, a condition for the installation of an AV exists, as well as the definition file of the AV.

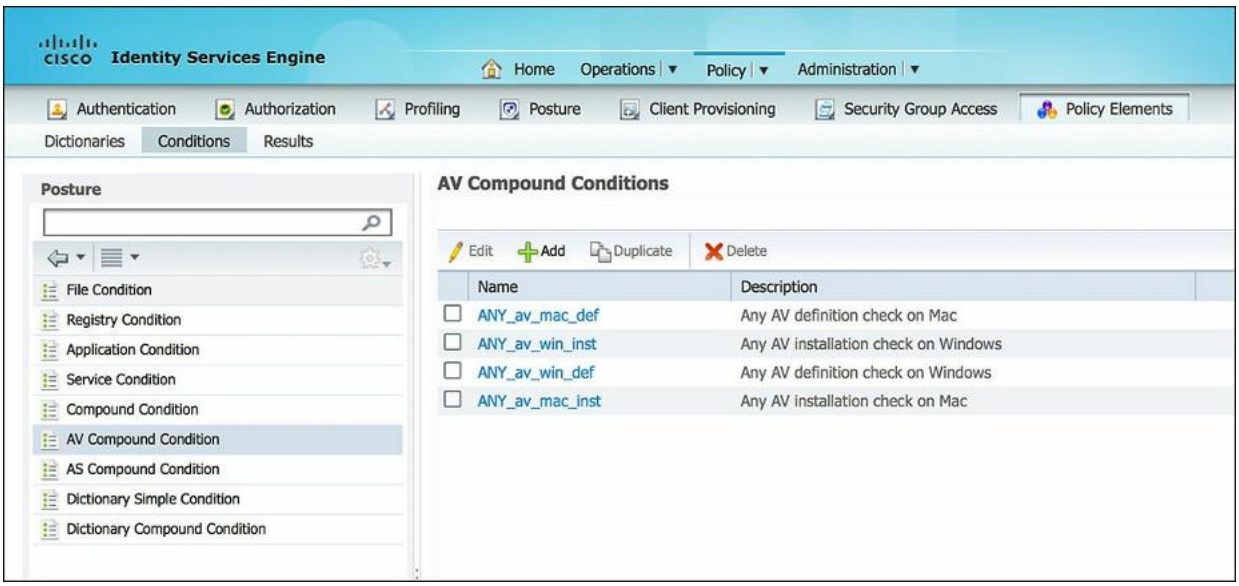


Figure 19-7 AV compound conditions.

If you click ANY_av_mac_def, you'll see that this rule has a requirement of Mac OSX as the OS where the Vendor option can be ANY with a Check Type of Definition (see Figure 19-8). This rule simply ensures that the Mac OS X endpoint has any antivirus running, but it must have a virus definition that is, at most, five days old.

https://t.me/learningnets

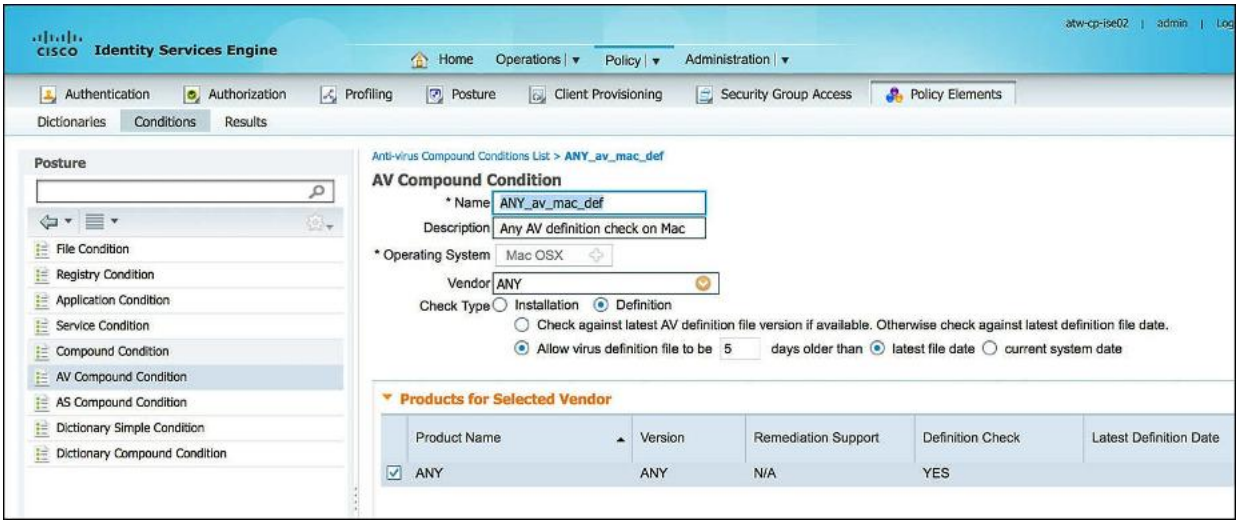


Figure 19-8 ANY_av_mac_def posture condition.

Now that we've looked at a simple condition, let's take a quick look at one that is a little more complex. Clicking Registry Condition in the left pane displays a long list of conditions that are predefined. Many of these conditions are looking for OS updates for known vulnerabilities or to determine which version of OS is running.

Let's define a Registry Condition of our own. Click the Add button. You are prompted to provide a name and description for your Registry condition. You can also select the RegistryType. This RegistryType allows for three options: the presence of the Registry key, a specific value being in the Registry key, or the default value being in the Registry key. You must then provide the "path" of the Registry key, selecting the Registry Root Key first and then providing the remainder of the path. Because you might want to confirm that a Registry key is NOT present, does NOT have a specific value, or is NOT the default value, you can change the Value Operator from Exists to DoesNotExist. Finally, you can select the relevant OS.

Figure 19-9 shows one example of a Registry condition that checks for the patch of the W32.Blaster.Worm (https://technet.microsoft.com/library/security/ms03-026) on Windows XP.



Figure 19-9 W32.Blaster.Worm Registry condition.

Remediation

With these conditions, we defined what we are checking for on the endpoint. This is the posture assessment. Next, we must define what we are going to do if the condition we found is or is not what we want it to be. If the result of the condition goes against our security policy, we'll need the user to remediate the discrepancy before gaining access to the network. To configure the remediation for posturing, go to **Policy > Policy Elements > Results > Posture > Remediation Actions.**

https://t.me/learningnets

With our previous examples, if the Mac OS X computer had a virus definition that was older than five days, we can require the user to install the latest antivirus definition before moving forward with his network access. For the W32.Blaster.Worm patch on Windows XP, we can force the user to install a Windows update. Again, some remediation types might not be available on all operating systems.

For this example, let's revisit the Mac OS X endpoint that is missing the latest antivirus definition. This is a huge risk to our network and our well-being as a company. For this reason, we'll require this endpoint to install the latest AV definition before it will be granted access to the network. To begin the configuration of the remediation, click AV Remediation.

Clicking AV Remediation shows that there are two predefined remediations: AnyAVDefRemediationWin and AnyAVDefRemediationMac. The latter remediation will be the relevant one for the Mac OS X endpoint that is lacking the latest definition. Opening this remediation policy, you will see a name, a description, that it is a manual remediation (meaning, you'll need to click a link), that the remediation is for a Mac OS X computer, and that it is for any AV vendor name.

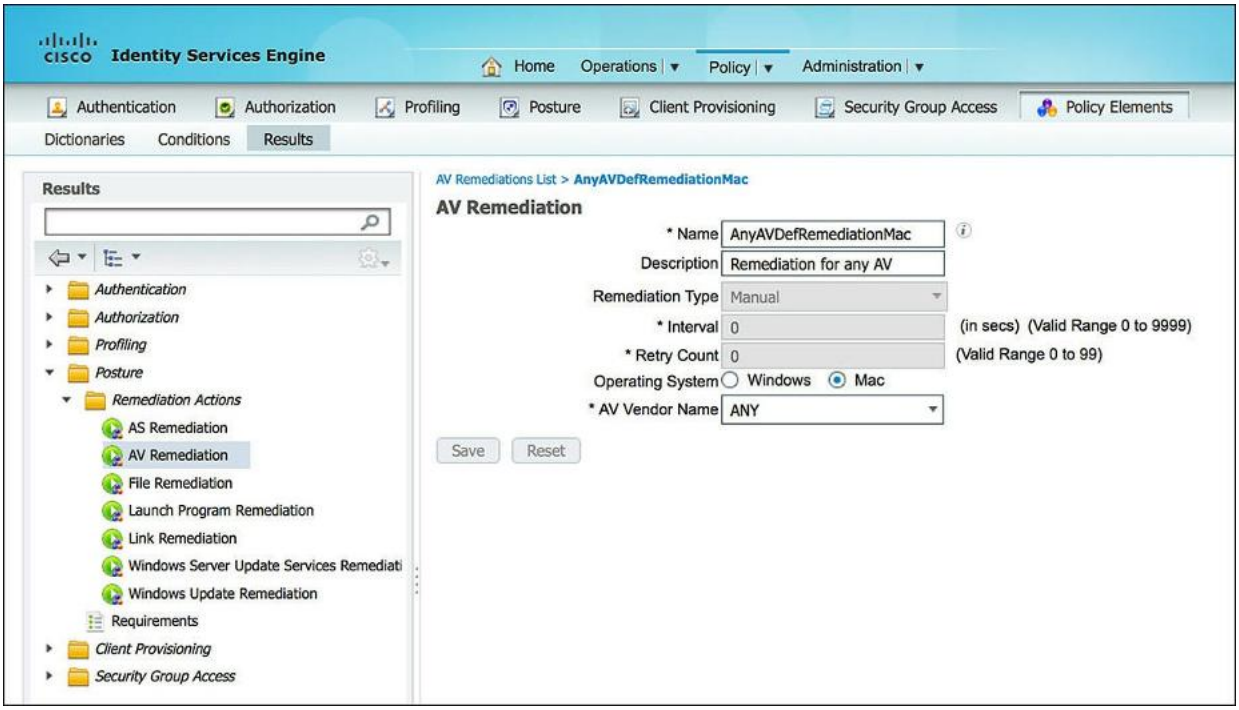


Figure 19-10 AV remediation for Mac OS X.

Requirement

The next step in our posturing process is to define what we plan to do about our noncompliance to the security policy. With the conditions, we defined what we want to check for on the endpoint. We also defined what we intend to do if the condition is not met by using the remediation construct. However, outside of our own insight, we have nothing else that links the condition to the remediation. The requirement function brings the two together.

To get to the requirement function, browse to **Policy > Policy Elements > Results > Posture > Requirements** (see Figure 19-11). Several requirements already are written to cover many of the antivirus and antispyware posture assessments on Windows and Mac OS X endpoints alike. Looking down the list, you can see that the relevant rule for our Mac OS X endpoint that doesn't have the latest antivirus definition is Any_AV_Definition_Mac.

https://t.me/learningnets

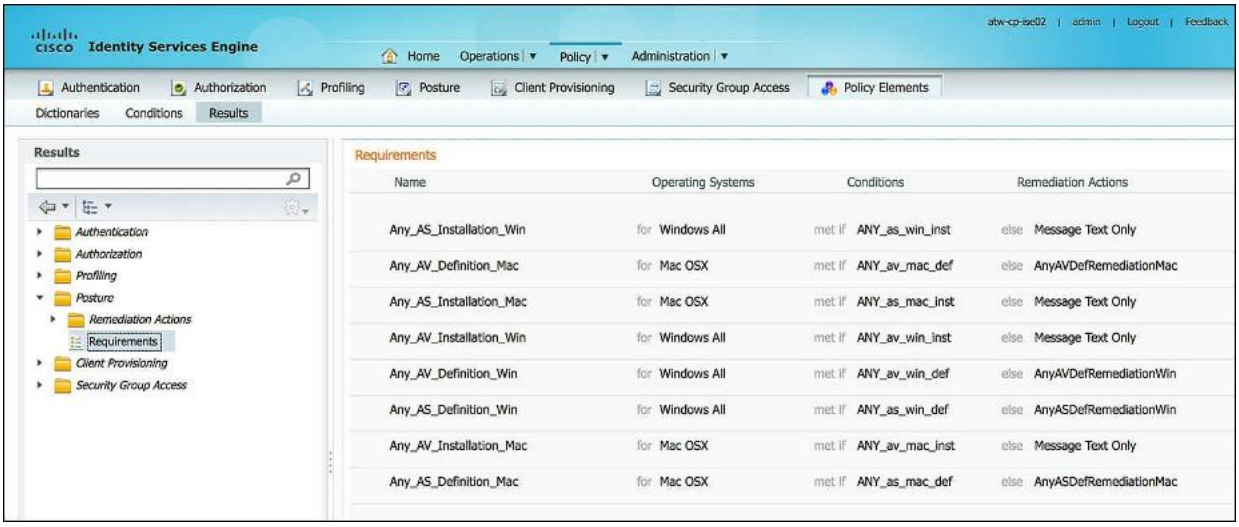


Figure 19-11 Posture requirements.

Looking at that Any_AV_Definition_Mac closely, you'll see it pertains to Mac OSX and is met if ANY_av_mac_def; if you recall, that condition states that my AV definition file must be less than five days old. If this condition is NOT met—that is, else—ISE will take the steps defined in the column Remediation Actions. In this case, the Remediation is AnyAVDefRemediationMac.

In this case, the conditions that were predefined were pretty simple—for a relatively simple test. However, if you wanted to have a few more conditions defined before you worried about upgrading the virus definition, you can define multiple conditions under the Conditions column (see Figure 19-12). As you combine more than one condition, you can define a few options on how those conditions are met. The three options that are available when you are combining multiple conditions are All Selected Conditions Succeed, Any Selected Condition Succeeds, and No Selected Condition Succeeds.

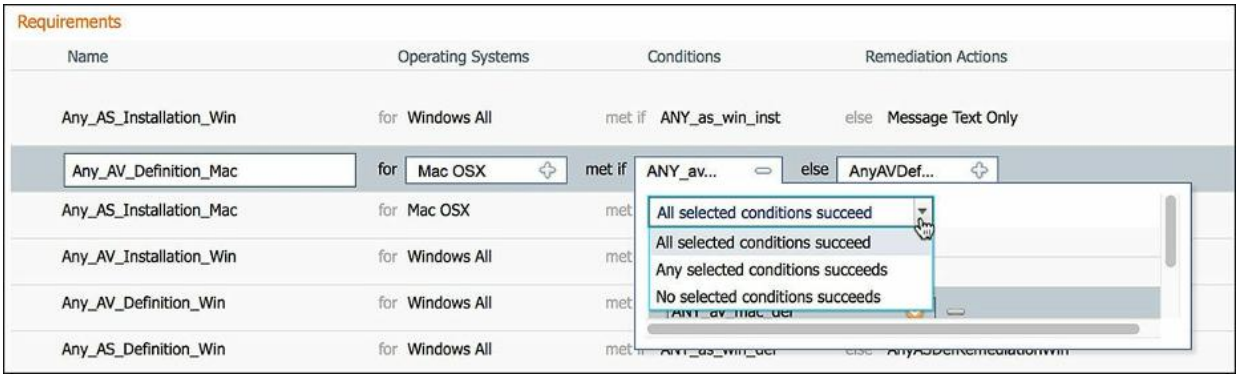


Figure 19-12 Multiple condition combining options.

Even yet, we still haven't defined the totality of conditions that each endpoint must meet. For instance, do all Mac OS X computer have to have the latest AV definitions or only those laptops that are taken off campus? What about the people in IT who refuse to install AV programs because it slows down their computers during intensive lunchtime gaming? With the posture policy, we will now lay out all the rules and define who has to meet which requirements. To get to the posture policy, go to **Policy > Posture**.

The Posture Policy screen displays a familiar IF-THEN construct. In this case, we will be defining which endpoints (based on identity groups, operating systems, and other conditions) must meet which requirements (see Figure 19-13). For our Mac OS X endpoint that is missing the latest antivirus definition, we'll give the posture policy the name MacOSX_OldAVDef. For the OS, we know that we only care about the Mac OS X virus definitions. For the other conditions, we could select a group that excludes the IT team, so the IT team would be allowed to play their lunchtime games without any issues. However, for the sake of simplicity, we'll enforce the rule for everyone who has a Mac OS X endpoint. Finally, on the last column, we have to select the requirements to which this user group must adhere to—that is,

https://t.me/learningnets

Any_AV_Definition_Mac. In summary, this posture policy rule states that any user with a Mac OS X endpoint must have an antivirus definition that is less than five days old.



Figure 19-13 Posture policy rule.

Modifying the Authorization Policy for CPP

As the Mac OS X endpoint—the one without the latest antivirus definition—attempts to connect to the network, it will still be required to go through the normal authentication and authorization process. Our configuration of "all things posture" hasn't changed that requirement. Actually, now that we have posturing in place, we can make the authorization process that much stricter. Let's take a look at what is required to accomplish that.

Up to this point, the authorization policy was based on username, AD user group, certificates, and possibly profiling. Nothing, yet, has looked at the OS or other software running on the endpoint. With the posturing configuration complete, we can now start taking a look at the software on the endpoint (specifically, the Mac OS X endpoint) and start making authorization and access decisions based on that outcome.

The first step in this journey is to make sure each endpoint has a NAC agent installed. How do we tell whether the endpoint lacks a NAC agent? The easiest way to tell whether the endpoint lacks a NAC agent is if the posture status is unknown. If the posture status is compliant or noncompliant, this implies that the NAC agent was present on the endpoint, made a posture assessment, and reported to ISE (either for the good or the bad). Without a NAC agent present, the endpoint will not be able to tell ISE anything about its compliance. Let's add an authorization policy that will leverage this unknown posture status and send the user off to the Client Provisioning Portal (CPP) to download the NAC agent (see Figure 19-14). The AV pair that stores the posture status is Session:PostureStatus EQUALS Unknown.

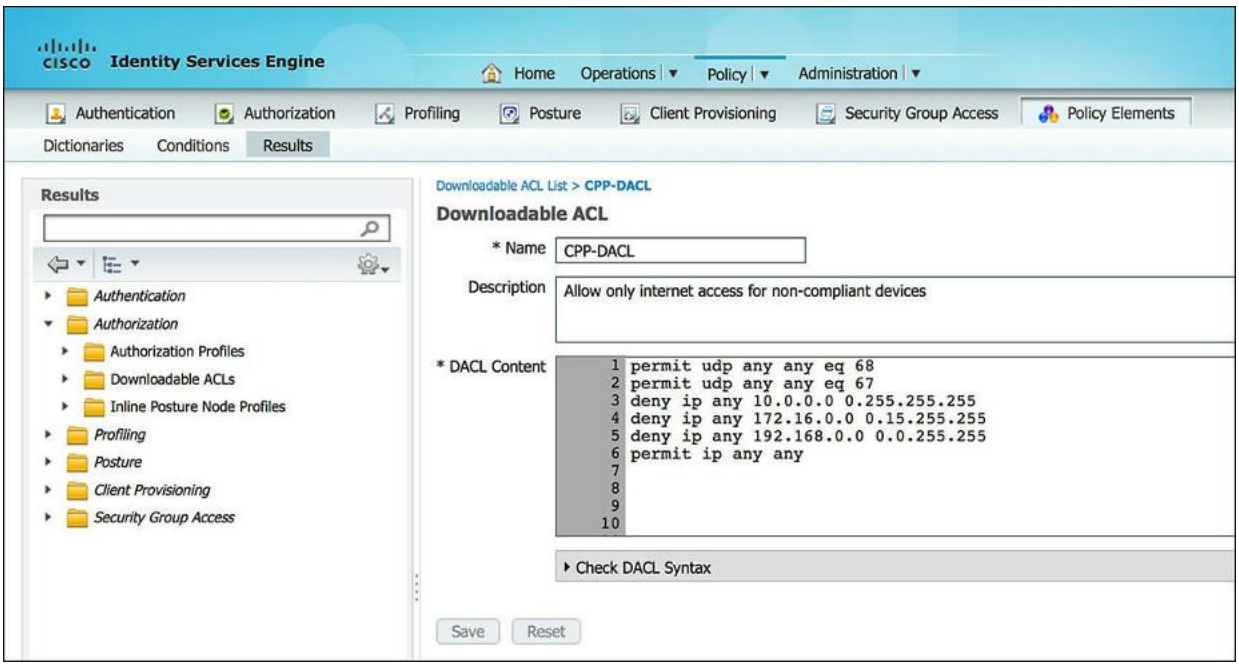


Figure 19-14 CPP downloadable ACL.

https://t.me/learningnets

To send the client to the CPP, we have to first configure an authorization profile that will trigger a redirect to the CPP. For this, we need to leverage two key common tasks:

- **DACL name**—In this case, we want to create a downloadable ACL that will allow the unknown user very limited access to the network.
- **Web redirection**—To send the endpoint to the correct web portal for CPP, we have to configure the web redirection and the associated redirect ACL to point to the CPP (see Figure 19-15). This redirect ACL is a named ACL, so we'll also need to define it on the NAD.

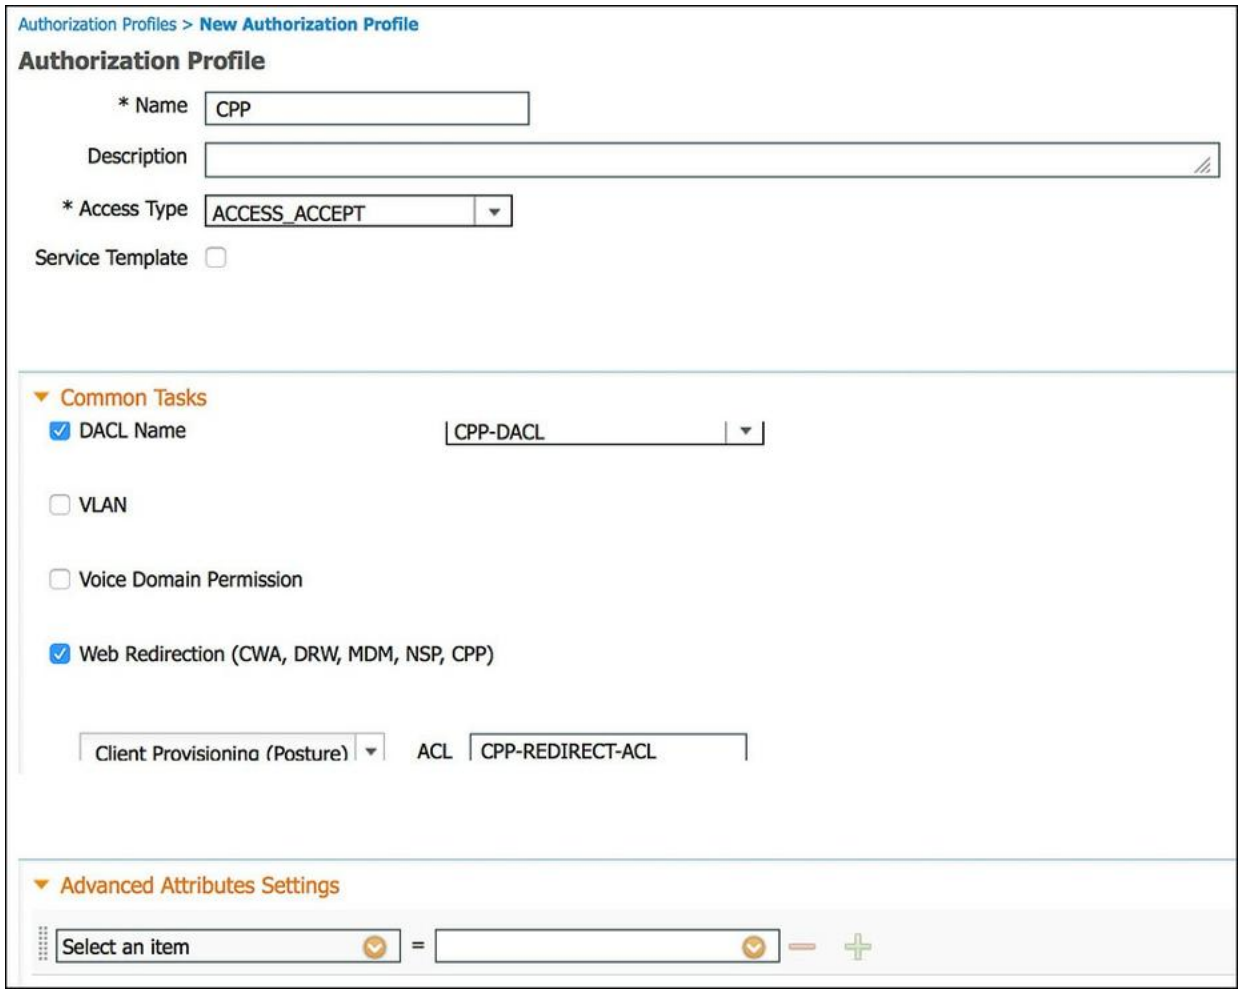


Figure 19-15 CPP authorization profile.

The redirect ACL we'll configure on the NAD will be CPP-REDIRECT-ACL and will match the ACL shown in Example 19-1.

Example 19-1 CPP-REDIRECT-ACL Contents

Click here to view code image

```
ip access-list extended CPP-REDIRECT-ACL
          ! Don't redirect DHCP traffic
          deny udp any any eq bootps
          ! Don't redirect DNS traffic
          deny udp any any eq domain
          ! Don't redirect traffic to ISE
          deny ip any host 10.1.100.232
          ! Redirect all other traffic
          permit ip any any
```

https://t.me/learningnets

With the authorization profile defined, we must now configure the appropriate authorization policy rule to send unknown clients to the CPP (see Figure 19-16).

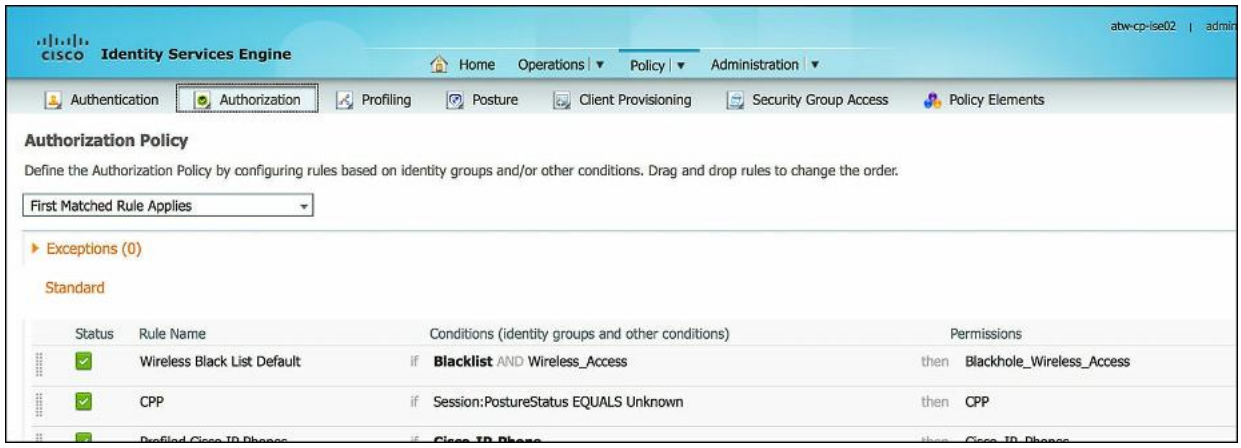


Figure 19-16 Authorization rule for CPP.

Modifying the Authorization Policy for Compliance

With the CPP rule, ISE was able to handle the situation in which the Mac OS X endpoint lacked a NAC agent altogether. What should be the behavior if I install a NAC agent—either out-of-band or via the CPP on ISE? If I have a NAC agent, my posture status should no longer reflect Unknown. My posture status should be updated to either Noncompliant or Compliant as, with a NAC agent installed, the NAC agent should report back to ISE with my correct posture.

Let's define the authorization profile that should be associated with a noncompliant endpoint (see Figure 19-17). Much like an unknown endpoint, we'll provide a noncompliant endpoint access to the Internet, allowing the endpoint access to whatever is needed to become compliant and/or enough "basic" access to surf the Internet. Because this policy will be the same as the policy within the CPP authorization profile (meaning the Redirect URL, Redirect ACL, and DACL will be exactly the same content), we'll simply create a new authorization rule named Noncompliant with a Session:PostureStatus EQUALS NonCompliant, yet point the authorization profile to CPP. We could also modify the CPP rule to contain Session:PostureStatus EQUALS Unknown OR Session:PostureStatus EQUALS NonCompliant, and this would accomplish the same task. But by keeping them separate, administrators have visibility into which endpoints do not have a NAC agent.

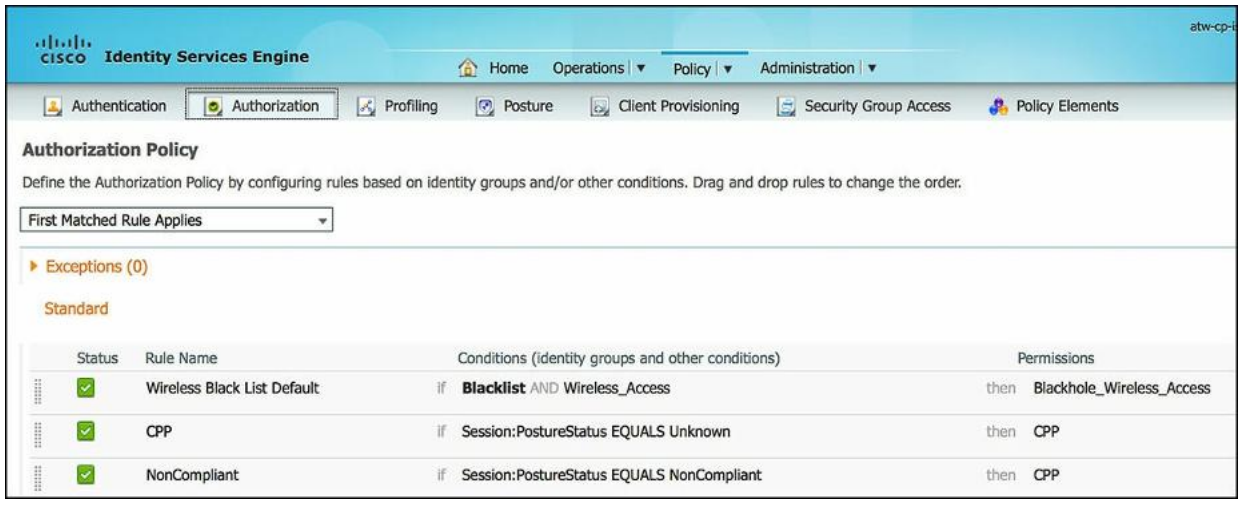


Figure 19-17 Authorization rule for Noncompliant.

If the endpoint is compliant, the rule we should hit is the Compliant rule, which has the condition of Session:PostureStatus EQUALS Compliant. To differentiate the level of access given to a compliant endpoint versus a noncompliant or unknown endpoint, I gave a compliant endpoint full access, giving the authorization profile of PermitAccess (see Figure 19-18). This PermitAccess authorization profile gives

https://t.me/learningnets

unfettered access by using the Access Type of ACCESS_ACCEPT without any further restrictions.

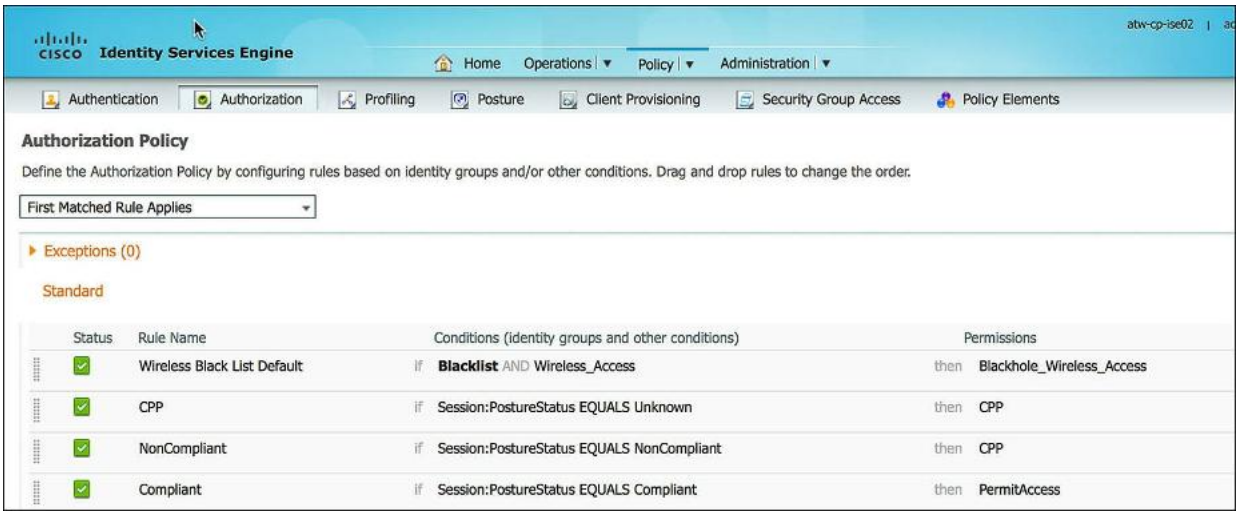


Figure 19-18 Authorization rule for Compliant.

Verifying Posture and Redirect

To verify that the posturing rules you have configured are correct and that you are hitting the correct policy, you can first try to browse the Internet from a Mac OS X endpoint that does not have a NAC agent installed. By not having the NAC agent installed, your endpoint will likely have the attribute of Session:PostureStatus EQUALS Unknown. Because the CPP authorization rule is at the top of the authorization policy and has this unknown as its sole condition, you should receive the authorization profile of CPP. This CPP profile should, upon an attempt to browse to a website, redirect you to the CPP portal whereby a NAC agent will be pushed down to your endpoint (see Figure 19-19). This is indeed what happened when I opened my browser; instead of reaching Cisco.com, I received the CPP portal with the NAC agent download button.

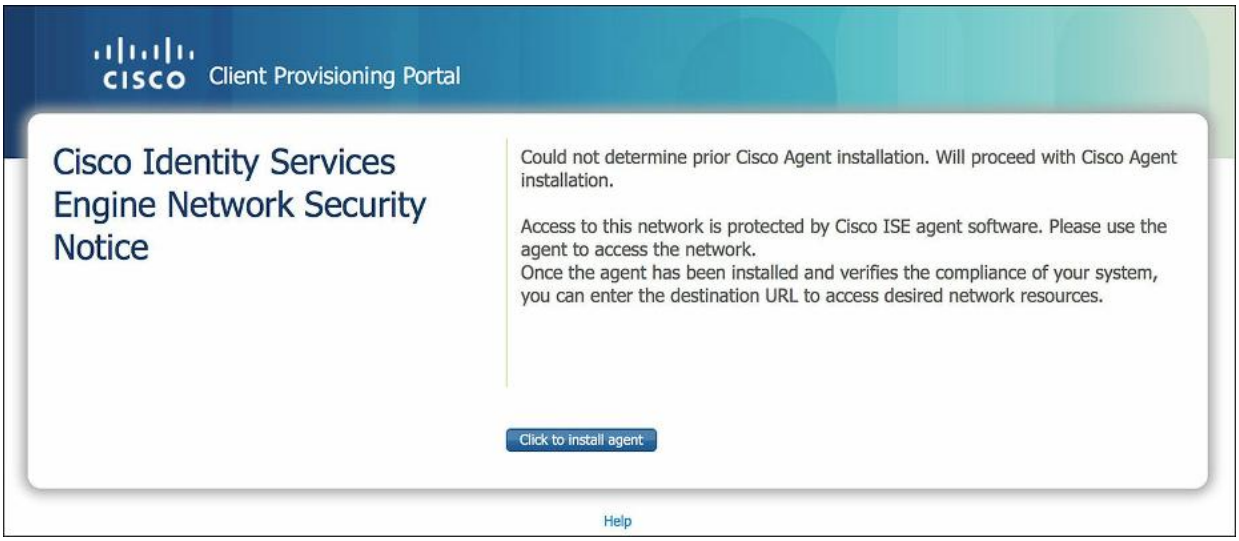


Figure 19-19 NAC agent download page.

To further indicate proper operation and confirm that you hit the CPP authorization profile, you can look at the output of **show authentication sessions interface <if_name>** details on the NAD, a Cisco 3750 Switch (see Example 19-2). If you are using a wireless client, the process to verify proper posture assessment on a WLC would be similar to the section "Verifying Guest Access on WLC/Switch" in Chapter 14, "Deploying Guest Services."

Example 19-2 show authentication sessions interface <if_name> **details**—Unknown Posture

https://t.me/learningnets

Click here to view code image

```
KR-3750#show authentication sessions interface gigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B860000004E1047F6F4
      Acct Session ID:  Unknown
               Handle:  0xCB00002B
       Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:
        URL Redirect:  https://atw-cp-ise02.ise.local:8443/guestportal/
gateway?sessionId=0A742B860000004E1047F6F4&action=cpp
     URL Redirect ACL:  CPP-REDIRECT-ACL
            ACS ACL:  xACSACLx-IP-CPP-DACL-54d081dd

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

Also note that the redirect URL, redirect ACL, and DACL as configured in the CPP authorization profile are present in the **show authentication sessions** output.

One final way to confirm that the endpoint has hit the correct policy due its unknown status is by looking at the Authentication Live Logs on ISE. Looking at the Live Log for the unknown posture, whereby I hit the CPP authorization profile, you can see in the Overview box that I indeed hit the CPP authorization profile and the CPP AuthorizationPolicyMatchedRule (see Figure 19-20). The Authentication Details box shows that my posture status remains at pending, which is synonymous with unknown, and hits the CPP authorization profile.

https://t.me/learningnets



https://t.me/learningnets

Figure 19-20 Authentication details, posture pending/unknown.

After clicking and loading the NAC agent on my Mac OS X client, I should no longer hit the unknown posture status. However, I still do not have any antivirus software. Therefore, I should now hit the noncompliant rule. As before, I attempt to browse to a website and hit the redirect, and I'm appropriately sent to the CPP with a posture status of pending. On my endpoint, I see the NAC agent do its posture assessment, and it pops up a dialog box that states that the endpoint is not in compliance as I'm lacking the AV definitions. I click Cancel on the NAC agent, refusing to comply with the updated AV definition and ensuring that my endpoint is noncompliant.

The output of **show authentications sessions interface** <if_name> **details** on the switch (see Example 19-3) shows that the policy hasn't really changed there. This is because my CPP rule for my unknown posture status resulted in a CPP authorization profile and the noncompliant rule for my noncompliant posture status also resulted in the same CPP authorization profile.

Example 19-3 show authentication sessions interface <if_name> **details**—Noncompliant Posture

Click here to view code image

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B860000005412E727CB
      Acct Session ID:  0x00000068
               Handle:  0x7A000030
       Current Policy:  POLICY_Gi1/0/3

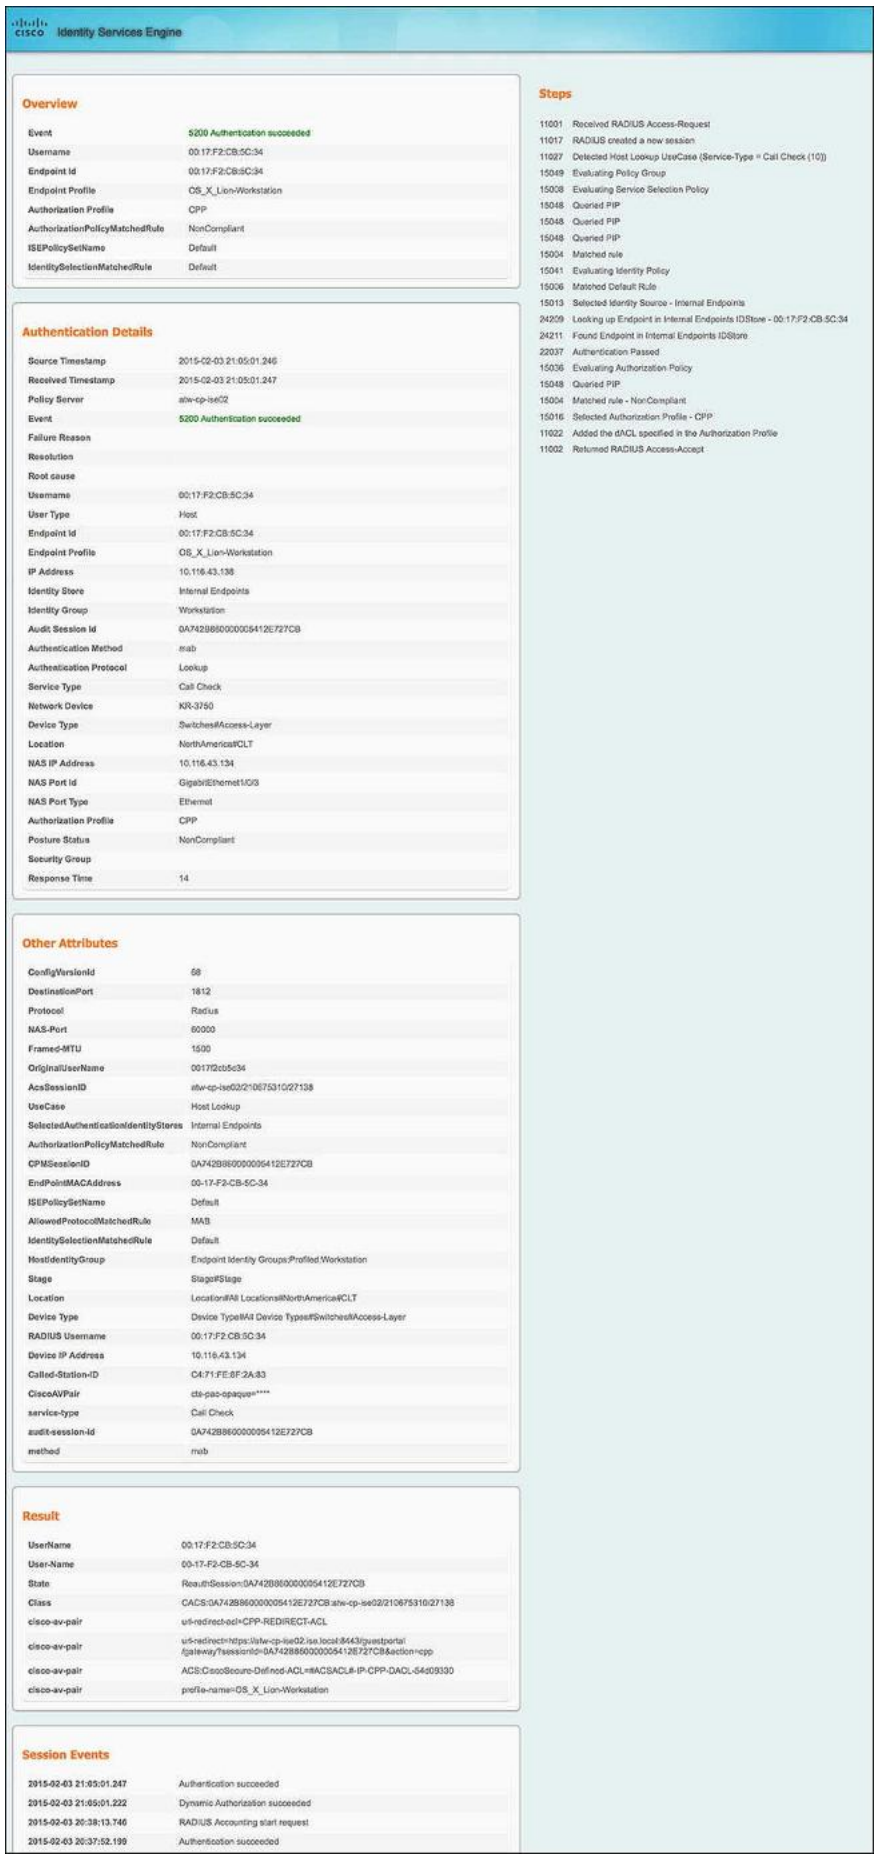
Local Policies:
           Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:
         URL Redirect:  https://atw-cp-ise02.ise.local:8443/guestportal/
gateway?sessionId=0A742B860000005412E727CB&action=cpp
     URL Redirect ACL:  CPP-REDIRECT-ACL
              ACS ACL:  xACSACLx-IP-CPP-DACL-54d09330

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

The Authentication Live Log for the NonCompliant use case (see Figure 19-21) shows that I've now hit the NonCompliant authorization rule, getting assigned the CPP authorization profile as we anticipated.

https://t.me/learningnets



https://t.me/learningnets

Figure 19-21 Authentication details, posture noncompliant.

If I choose to comply with the antivirus requirements, installing an antivirus solution and updating to a definition file that is less than five days old, my Mac OS X endpoint would then be compliant. After doing so, the output of **show authentication sessions** <if_name> **details** on the switch reveals an updated security policy (see Example 19-4). This indicates an unfettered level of access consistent with the authorization profile of PermitAccess.

Example 19-4 show authentication sessions interface <if_name> **details**—Compliant Posture

Click here to view code image

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B86000000581446F323
      Acct Session ID:  0x0000006F
               Handle:  0x5C000033
       Current Policy:  POLICY_Gi1/0/3

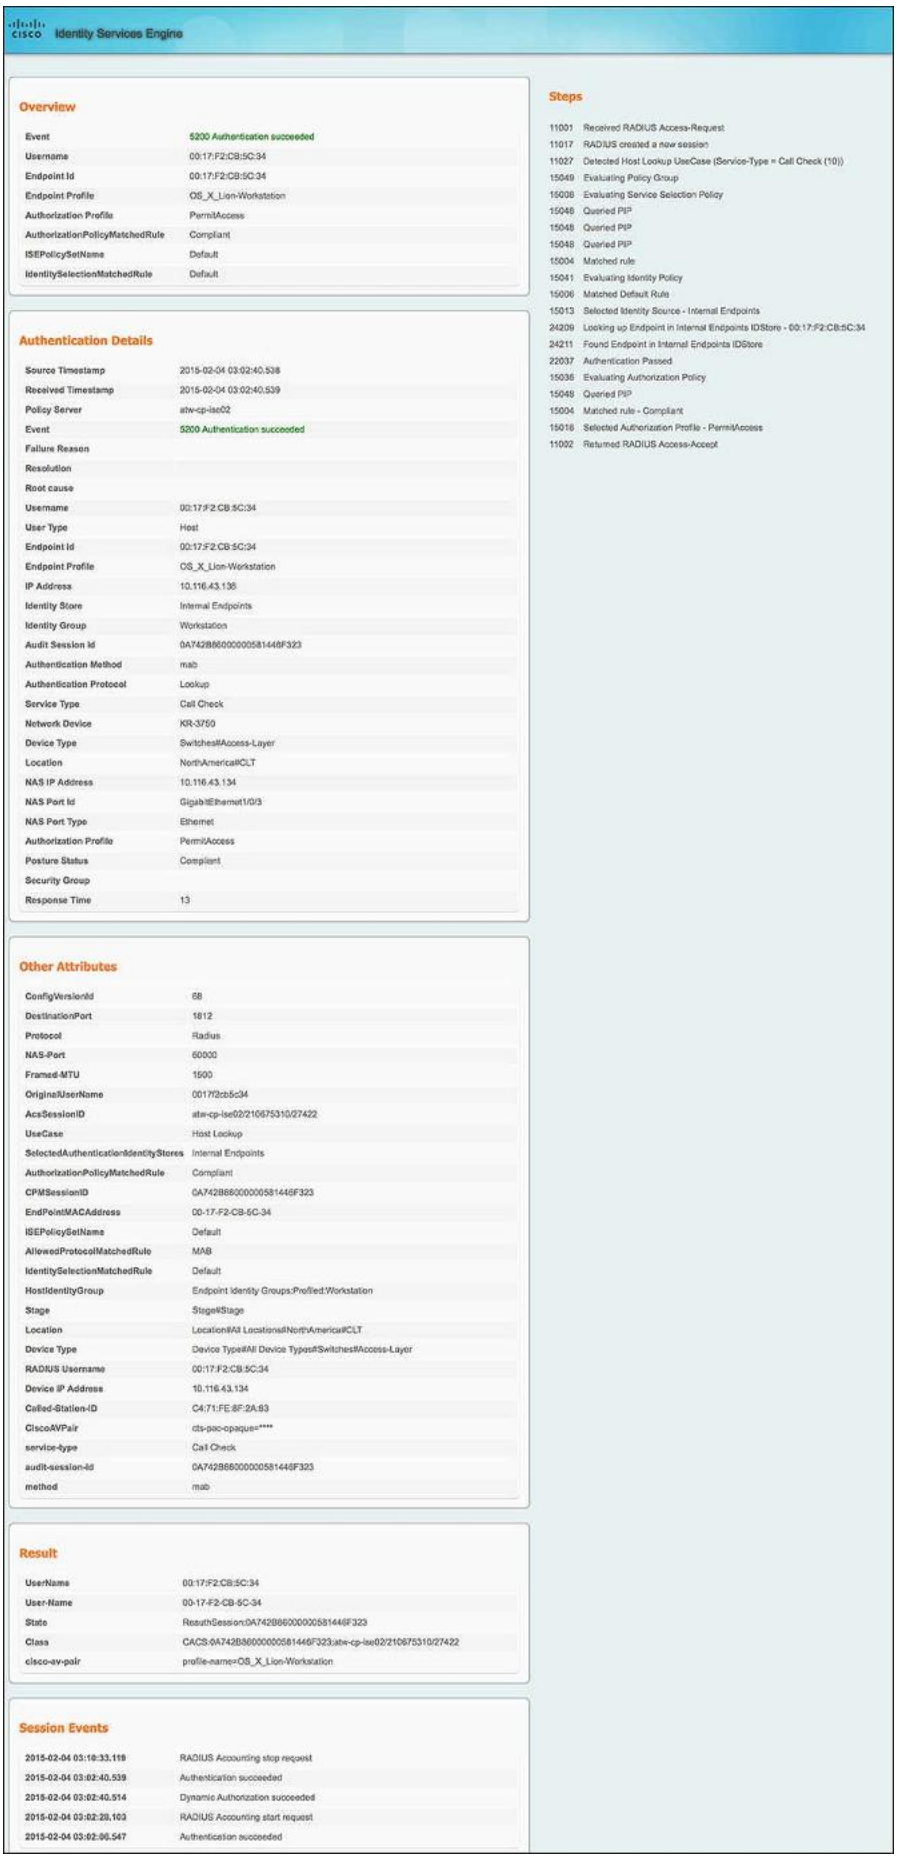
Local Policies:
           Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

Notice that there are no server policies present and no ACLs or redirects to hinder network access. From the Authentication Live Log on the ISE, you can see that we have indeed hit the Compliant rule and received the authorization profile of PermitAccess, again, the unfettered authorization profile (see Figure 19-22).

https://t.me/learningnets



https://t.me/learningnets

Figure 19-22 Authentication details, posture compliant.

During this chapter, we discussed the importance of knowing what software is running on an endpoint, a process known as posturing. Via the Posture Service on ISE, we are able to determine whether the software on an endpoint is in compliance with the network security policy. If an endpoint is not compliant with the security policy, ISE has the capability to restrict the endpoint's level of access, remediating the endpoint when possible while ensuring that its compromised state does not have a negative impact on the security of the network as a whole.

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 19-4 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
List	Posture has three major functional components.	648
List	There are three outcomes for a posture assessment.	649
Paragraph	Prior to implementing policy on your network, perform a comprehensive review of your security policy.	655

Table 19-4 Key Topics for Chapter 19

Define Key Terms

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

client provisioning portal (CPP)

https://t.me/learningnets

Part VI: Safely Deploying in the Enterprise

https://t.me/learningnets

Chapter 20. Deploying Safely

This chapter covers the following exam topics:

- **802.1X phasing (monitor mode, low impact, closed mode)**

Throughout this book, you have examined quite a bit of configuration detail about ISE and about the network access devices (NADs) themselves. You have explored the technical merit of policy creation, guest lifecycle 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 should just enable overnight with the flip of a switch.

This chapter focuses on the recommended approach to deploying the Secure Access system. We will review some of the challenges that were seen 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 enables 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."

Foundation Topics Section	Questions
802.1X phasing (monitor mode, low impact, closed mode)	1-10

Table 20-1 "Do I Know This Already?" Section-to-Question Mapping

> **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.** Using the 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 which device would have failed authentication and correcting the root cause prior to it taking effect
   - **d.** A method for alerting the administrator of failed authentications, so the end user may 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 authentications occur, but the authorizations are ignored, so as not to cause a denial of service
   - **d.** A method for identifying which device would have failed authentication and correcting the root

https://t.me/learningnets

cause prior to it taking effect

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 authentications, such as the use of smart cards.

   d. To ensure that a port, switch, or location is fully ready to be successful before enabling enforcement and specific authorization results.

4. True or False? The **authentication open** command performs EAP authentications but ignores authorization results.

   a. True

   b. False

5. True of False? **authentication open** allows all traffic to pass through the switch port before the authentication result is received from the AAA server.

   a. True

   b. False

6. What is the ISE configuration that will allow 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 within ISE.

   c. It is configured in ISE by enabling wireless monitor mode under the system settings.

   d. Monitor Mode is not possible with wireless LANs.

8. Using policy sets as described in this chapter, how would 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. True or False? A wired port must have a single configuration that supports authenticating supplicants, guests, and nonauthenticating devices.

   a. True

   b. False

10. Which of the modes is most closely related to the default of 802.1X?

    a. Closed Mode

    b. Monitor Mode

    c. Low-Impact Mode

https://t.me/learningnets

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 was IEEE 802.1X, which enabled authentication of a network access port prior to allowing devices on to 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. A device is either authenticated and gets access or fails authentication and is denied. Figure 20-1 illustrates the intended behavior of 802.1X.

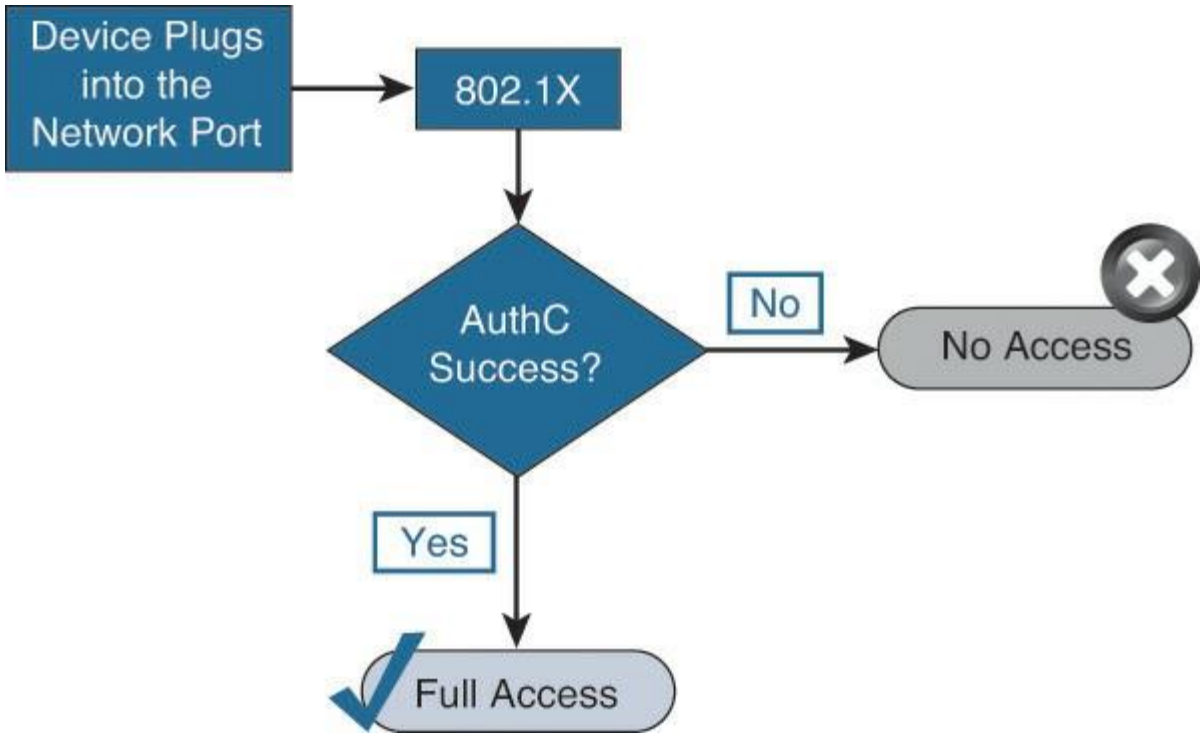


Figure 20-1 802.1X intended behavior.

However, there are many moving parts to this authentication process that must all be aligned properly to prevent causing a denial of service on the user population. Supplicants must be configured on devices; lists of MAC addresses must be created to allow MAB devices; profiling probes must be enabled and must be collecting 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 just flip the switch and enable 802.1X on all access-layer switch ports all at once, there would most likely be a swarm of angry users converging on the IT department demanding that people lose their jobs. This is often called a "career limiting event," or CLE for short.

Not long ago, there was an 802.1X implementation at a financial organization with 2,000 switch ports in their campus building. Due to an audit requirement, network authentication had to be enabled by a certain date; otherwise, they would be subject to fines. The mandate came down from above, the project received its funding, and away the deployment went. They lab tested everything and proved it all would work using Cisco Catalyst 6513 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 (GPO) was sent properly, and all the Windows XP systems were ready to

https://t.me/learningnets

authenticate. All the network team should have been required to do was turn on the authentication on the switch ports (theoretically).

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, which would enable the team to validate that authentications would all be successful before they 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 2,000 ports were wide open currently, and this was not making matters worse at all—they simply were not allowed to use that command.

So, the big day arrived; 10 p.m. on a Sunday night was the change control window. It was time to run the scripts and enable Closed Mode authentication across 2,000 switch ports in a matter of minutes. Of those 2,000 ports, a whopping 10 were authenticating successfully and the team had accomplished exactly what they were warned about: They had 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. However, no one realized that the supplicant configuration would not take effect prior to rebooting the Windows systems! Of course, the team did not know that 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 when they came into the office that Monday morning.

The story has a happy ending. Not only is there a fun story to write about in this book, but after the desktop team pushed out a job to reboot all the systems, the team reenabled authentication at the next change-control window and had 99% of the systems authenticating successfully.

However, not all deployments are that lucky, or that well planned-out in advance. This is why a phased approach to deploying identity solutions is always a good idea.

A Phased Approach

By using a phased approach to deployment, you are able to start 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 as illustrated in that story. With a monitoring phase, you have time to build your list of endpoints with profiling; you can manually import the MAC addresses that will be MAB'd 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 of enforcement. Figure 20-2 illustrates the phased deployment concept.

https://t.me/learningnets

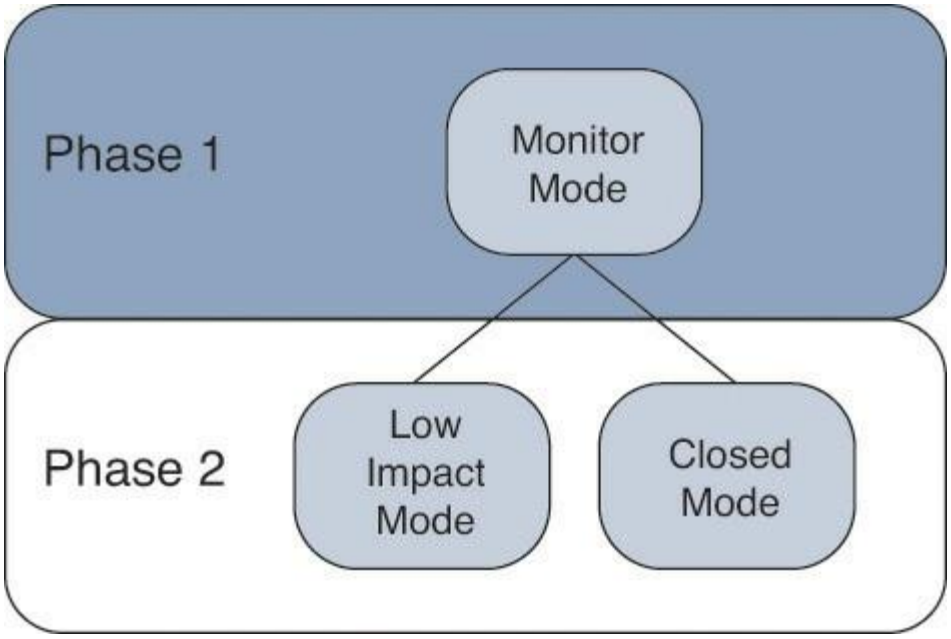


Figure 20-2 Phased deployments.

Comparing Authentication Open to Standard 802.1X

Per the IEEE standard, a port that is protected with 802.1X will not allow network traffic to flow without a successful authentication. Figure 20-3 illustrates that an 802.1X-controlled port normally allows only EAP, CDP, and LLDP traffic to enter the port (all three are Layer-2 protocols); no other traffic is permitted. When 802.1X is enabled on a port, it is said to be a "supplicant authenticator." That is a complex way of stating that the port will communicate with EAP at Layer-2 and the switch will broker that authentication to the RADIUS server.

https://t.me/learningnets

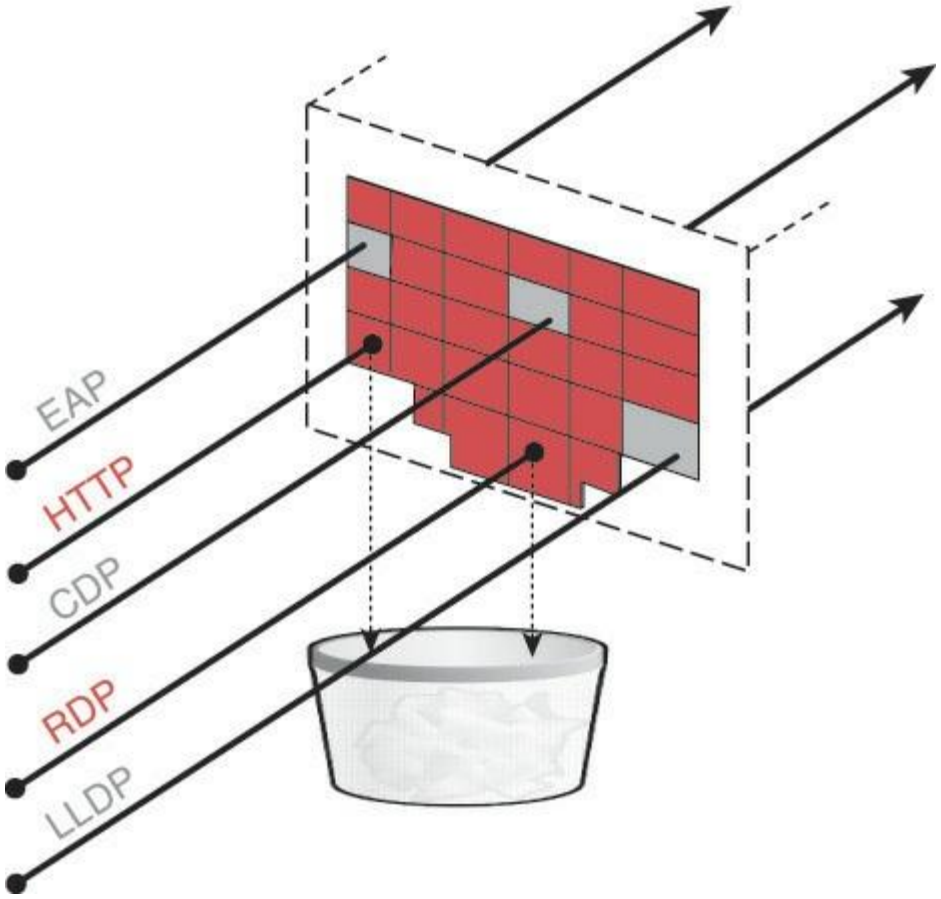


Figure 20-3 Default port behavior with 802.1X.

Figure 20-3 illustrates the default port behavior of an 802.1X-enabled switch port.

Cisco created an enhancement to standard 802.1X ports that enables the port to be a supplicant authenticator; however, it also allows all traffic to flow normally through the switch port even without an 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 the switch received an Access-Reject message from the RADIUS server, the Access-Reject message is ignored.

Figure 20-4 illustrates that regardless of authentication, the switch port will allow all traffic to flow, but it will also authenticate the supplicant and perform MAB just like a standard 802.1X-enabled switch port.

https://t.me/learningnets

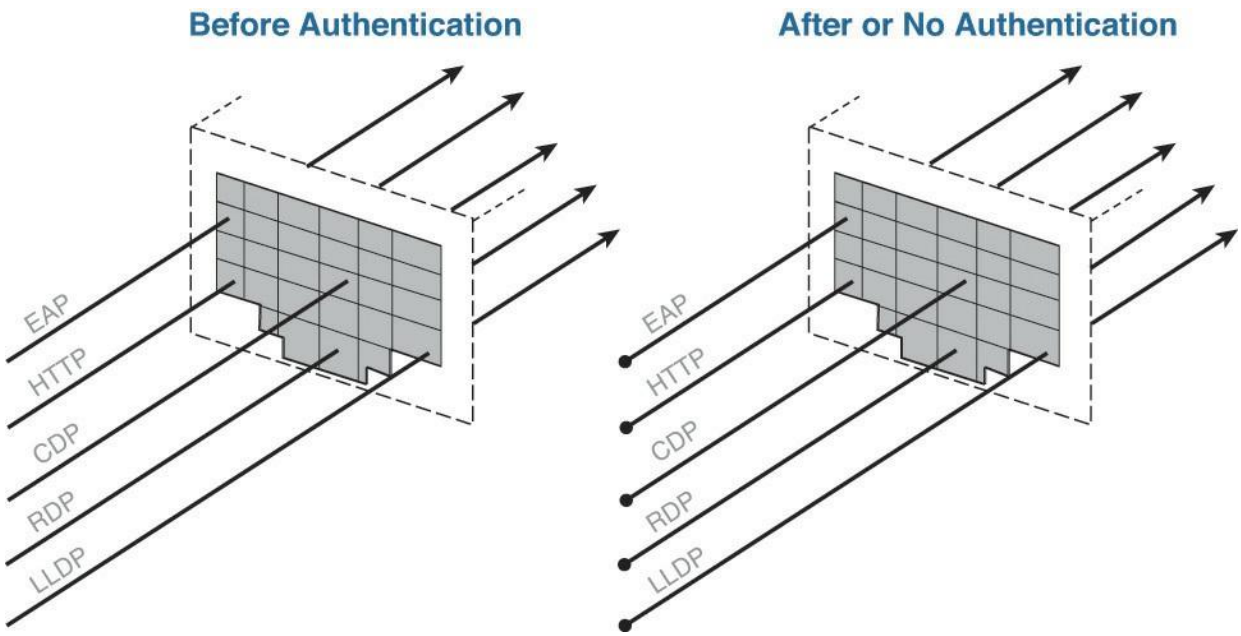


Figure 20-4 Port behavior with open authentication.

The creation of this authenticator enhancement is what truly made Monitor Mode possible. It is, of course, not the only necessary component of Monitor Mode, but it is certainly the catalyst (pardon the pun).

Preparing ISE for a Staged Deployment

One of the primary ways to differentiate modes within your ISE policies is with the use of network device groups (NDGs) and policy sets. Configure the NDGs to have a top-level group named Stage, and then configure subgroups for Monitor Mode, Low-Impact Mode, and Closed Mode. Then the authorization policies can be configured to have different policies for each of the network devices that are assigned to particular stages of deployment. Although you could do this all in a single authorization policy, using a separate policy set for each stage of deployment will keep the policies nice, clean, small, and organized.

Start by enabling policy sets under **Administration > System > Settings > Policy Sets**, as shown in Figure 20-5.

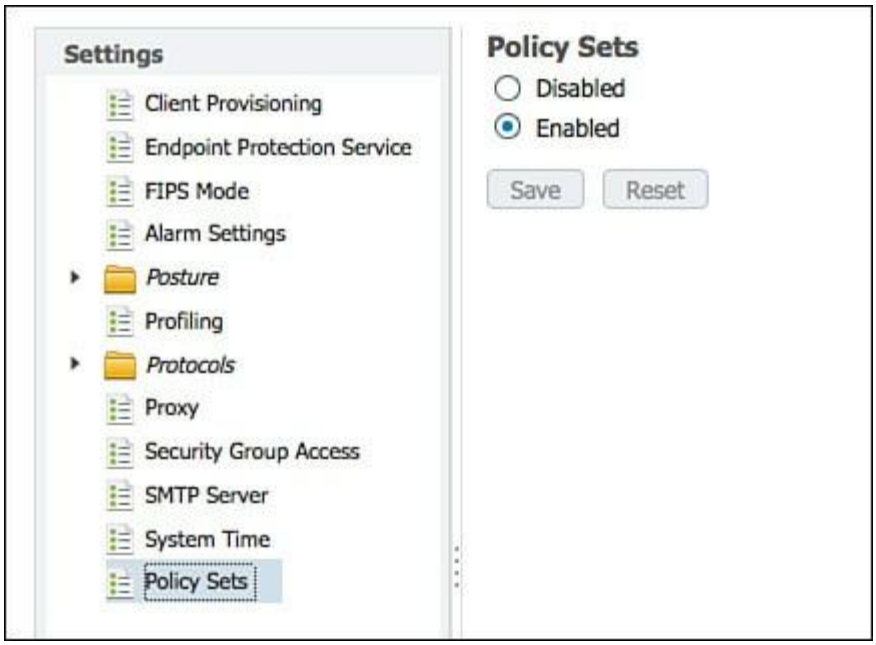


Figure 20-5 Enabling policy sets.

https://t.me/learningnets

All the existing authentication and authorization policies are preserved in the Default Policy Set, and all authentications and authorizations will proceed as they used to.

From within the ISE GUI, do the following:

Step 1. Navigate to **Administration > Network Resources > Network Device Groups**.

Step 2. Click **Add**.

Step 3. Name the NDG **Stage**.

Step 4. Set the Type to be **Stage**.

Step 5. Click **Submit**.

Figure 20-6 shows the creation of the top-level NDG.

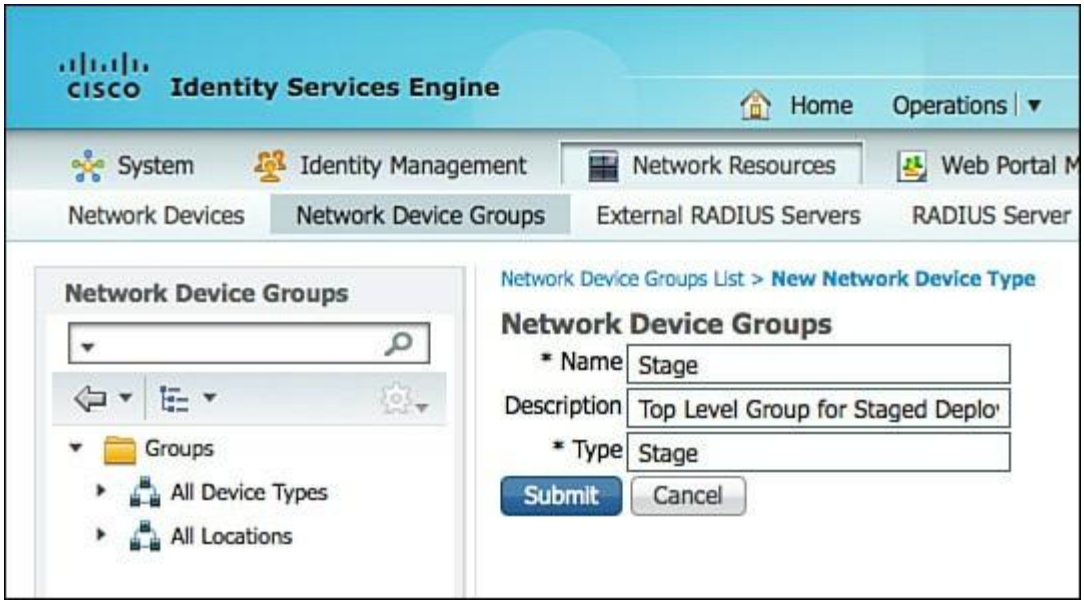


Figure 20-6 Add a stage NDG.

Step 6. Click the new Stage NDG on the left side.

Step 7. Repeat steps 2–5 and create a new subgroup for each of the deployment modes (as shown in Figure 20-7):

- Monitor Mode
- Low-Impact Mode
- Closed Mode

https://t.me/learningnets

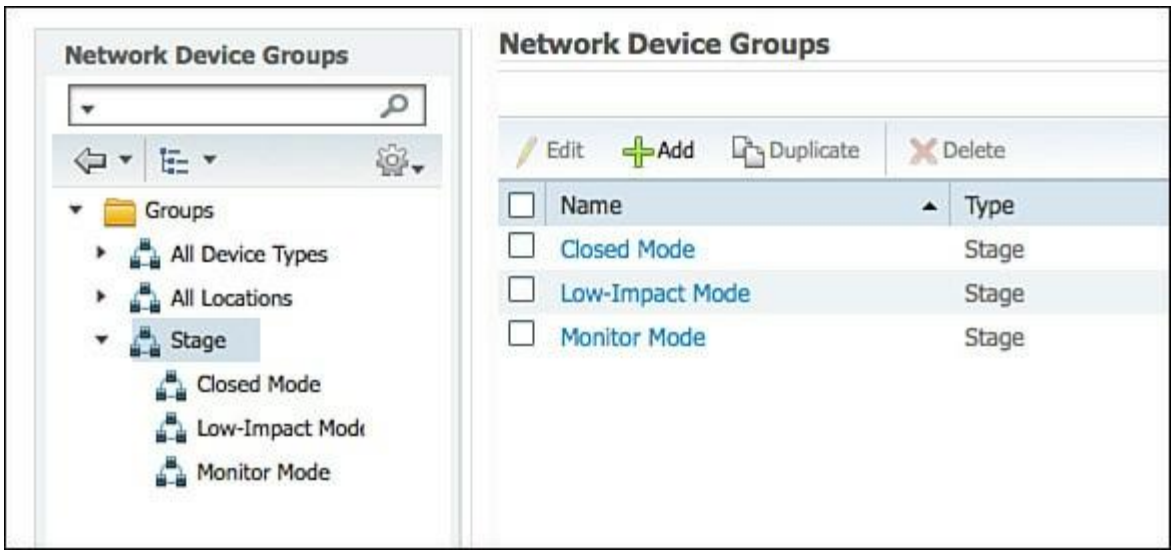


Figure 20-7 Final stage NDGs.

Monitor Mode

Monitor Mode is a process, not just a command on a switch. The process is to enable authentication (with authentication open), see exactly which devices fail and which ones succeed, and correct the failed authentications before they cause any problems.

Figure 20-8 shows a high-level flow diagram describing the Monitor Mode.

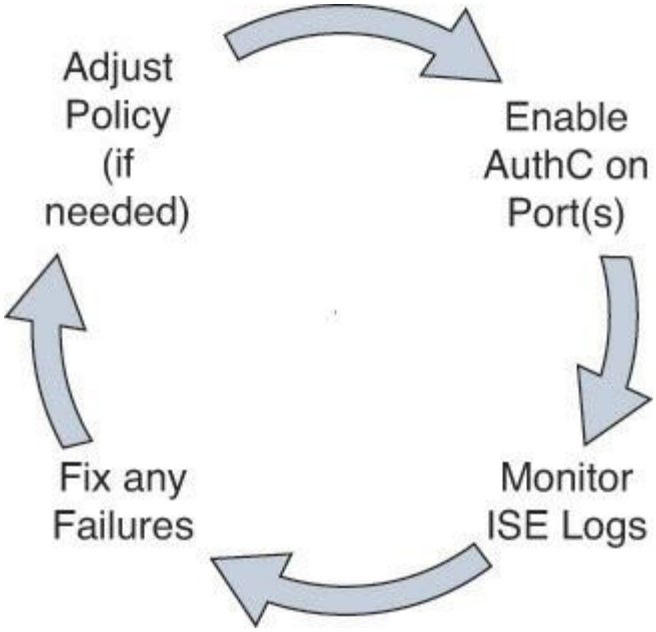


Figure 20-8 Monitor Mode operational flow.

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 to which network to connect 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 pop-up on the endpoint asking to which network you would like to connect. It's just assumed that you are physically connected and therefore 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, supplicant or not.

https://t.me/learningnets



The wired port must be able to handle the following:

- 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, printer, and so on)
- Guest users

Thus, there is quite a bit to audit when in Monitor Mode.



Another important thing to understand about Monitor Mode is that authorization results from the RADIUS server will absolutely be honored (Access-Reject is the only command that is ignored). So, if your authorization result from ISE includes dVLAN assignment or dACLs, those will be honored and applied to the port.

For a phased deployment approach, it is highly recommend to use NDGs in ISE. Using these NDGs, you can build specific policies that send the basic authorization results (Access-Accept and Access-Reject) only to switches that are part of a Monitor Mode NDG.

Figure 20-9 shows a high-level flow diagram describing the Monitor Mode.

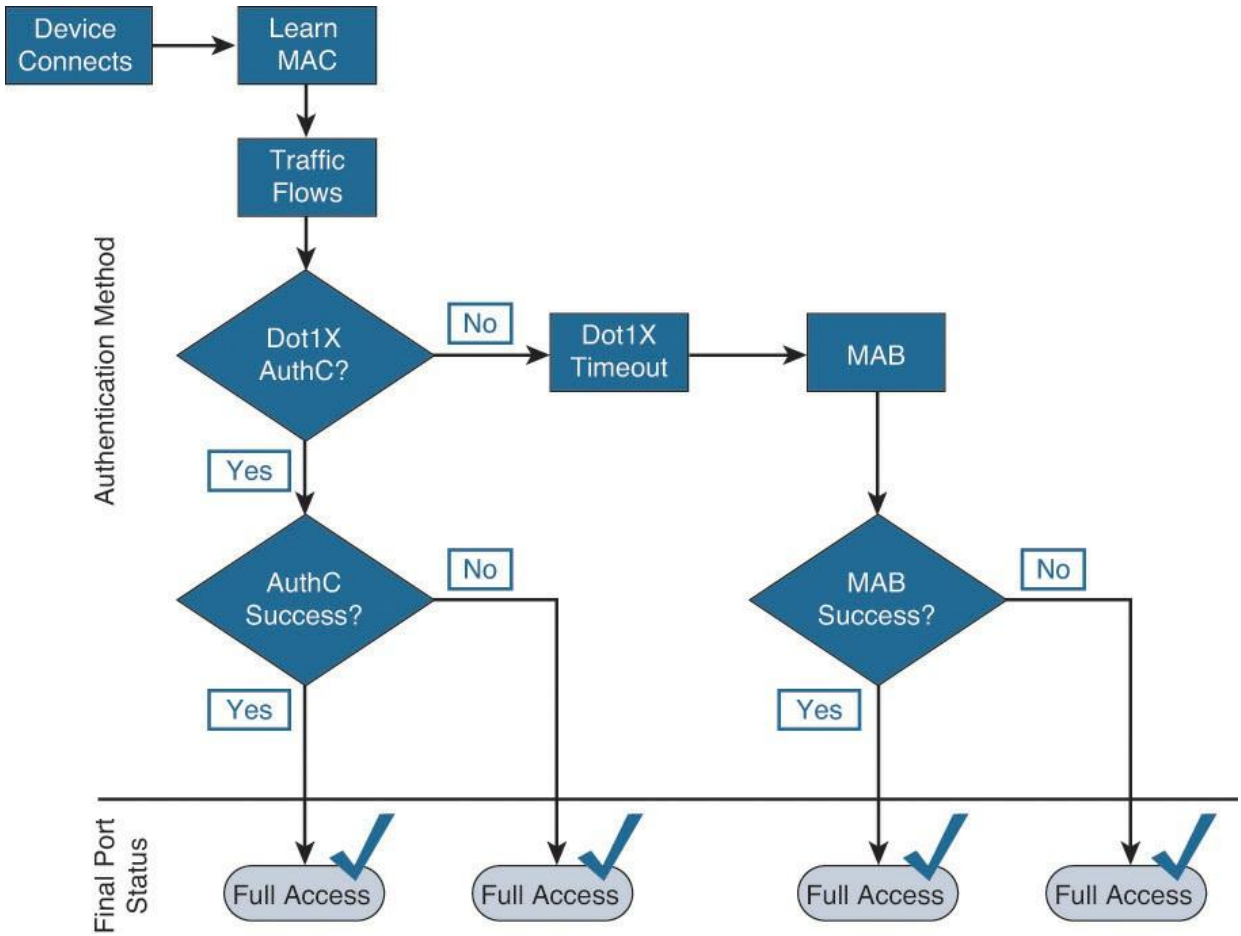


Figure 20-9 Monitor Mode flow diagram.

https://t.me/learningnets

Now that you have the NDGs configured for the different stages of the deployment, you can move on to creating the policy sets themselves.

To create the policy set for Monitor Mode, do the following:

Step 1. Navigate to **Policy > Policy Set**.

Step 2. Ensure that your default policy is selected on the left side (as shown in Figure 20-10), and then click the plus sign (+) in the upper-left corner.

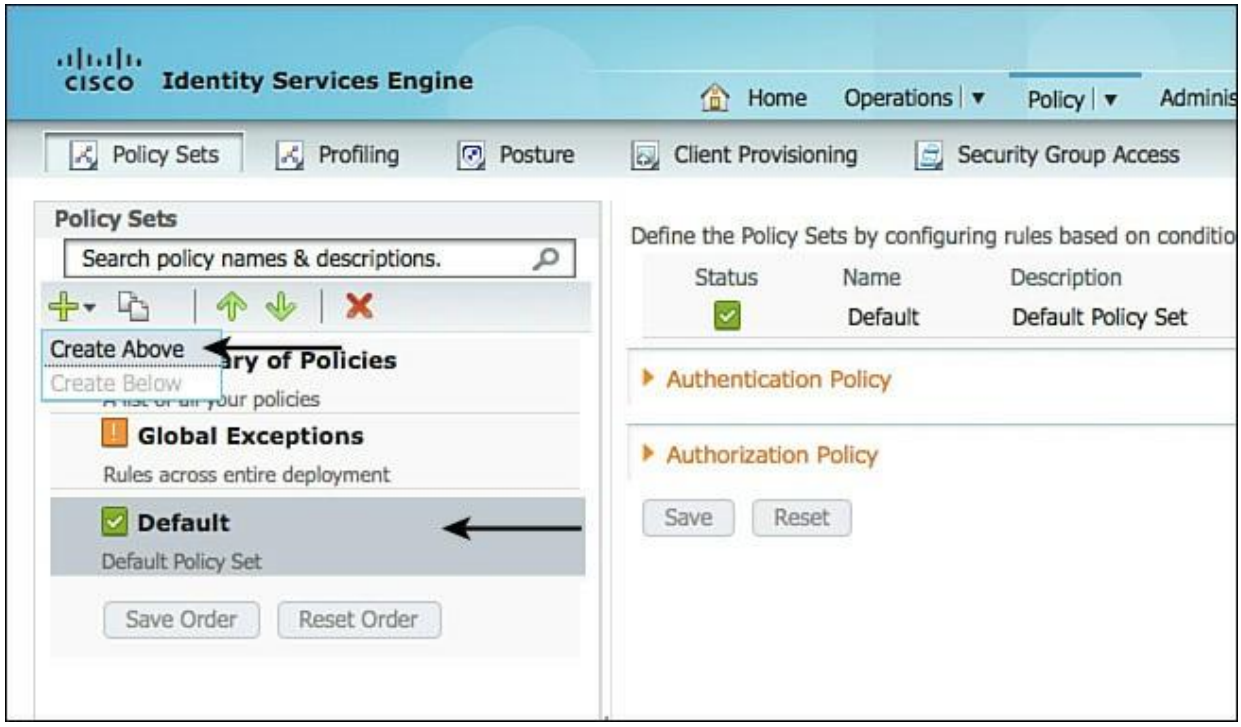


Figure 20-10 Default policy set selected, create above.

Step 3. Select **Create Above**.

Step 4. In the upper-right side of the policy window, double-click the text **Enter Policy Name**, as shown in Figure 20-11.

https://t.me/learningnets



Figure 20-11 Enter Policy Name.

Step 5. Name the new policy set **Monitor_Mode** and provide a description.

Step 6. Create a new condition of **DEVICE :: Stage Equals Stage#Monitor Mode**.

Step 7. Click **Done**.

Step 8. Click **Submit**.

Figure 20-12 shows the policy set selection rule for Monitor Mode.

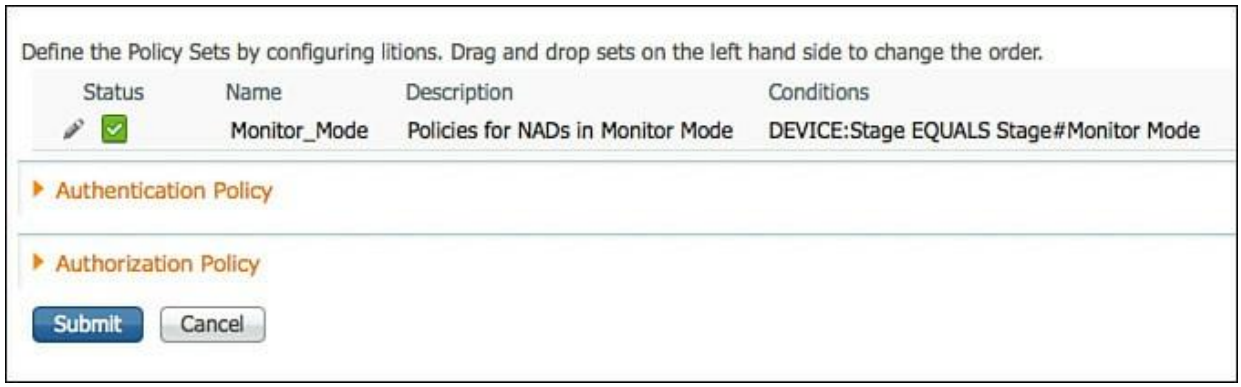


Figure 20-12 Monitor Mode policy set.

At this point, any network device that is a member of the NDG named Monitor Mode will use this policy set. All authentication and authorizations with those network devices will occur with this set of policies, and all other network devices will still use the default policy set.

It is up to you, the administrator of ISE policies, to ensure that any 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 that rule is the AV-pair for IP phones that grants a phone permission to use the voice VLAN.

https://t.me/learningnets

Low-Impact Mode

As described previously in this chapter, Low-Impact Mode is one of the end state choices for your deployment. Closed Mode is the other final stage. There is no specific best practice for which mode is better to deploy; it is entirely dependent on the organization and the needs.

A number of large organizations use a variety of technologies to reimage desktop systems that use the Pre-eXecutable Environment (PXE) to boot into a pseudo-OS and then connect to an imaging server that reimages the company computer. Those PXE environments can be time sensitive and might not have the ability to authenticate to the network infrastructure. However, it has to seamlessly be able to boot, connect to the reimaging server, and update the desktop to the latest corporate image—and do so without any additional user interaction. Low-Impact Mode was the only way to make this work feasibly in these 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, download their operating systems (OSes) from the local store server, and have that access before the DHCP timers expire. After that OS is loaded into memory and takes over the system, its supplicant would send an EAPoL start message into the network and authenticate with 802.1X. Low-Impact Mode allowed 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, which allows traffic to enter the switch prior to an authorization result. This permits the DHCP clients to be assigned an IP address before their DHCP timers run out (for example).



With Low-Impact Mode, you add security right from the get-go 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 downloadable ACL (dACL) sent down from ISE.

Figure 20-13 shows the operational flow intended for Low-Impact Mode. This mode is one of the two possible end states (closed mode being the second), and as such, very specific access can be provided per user, device, or any other condition that you want to use in the ISE authorization policies. Remember 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 who have been authorized (security principle of least privilege).

https://t.me/learningnets

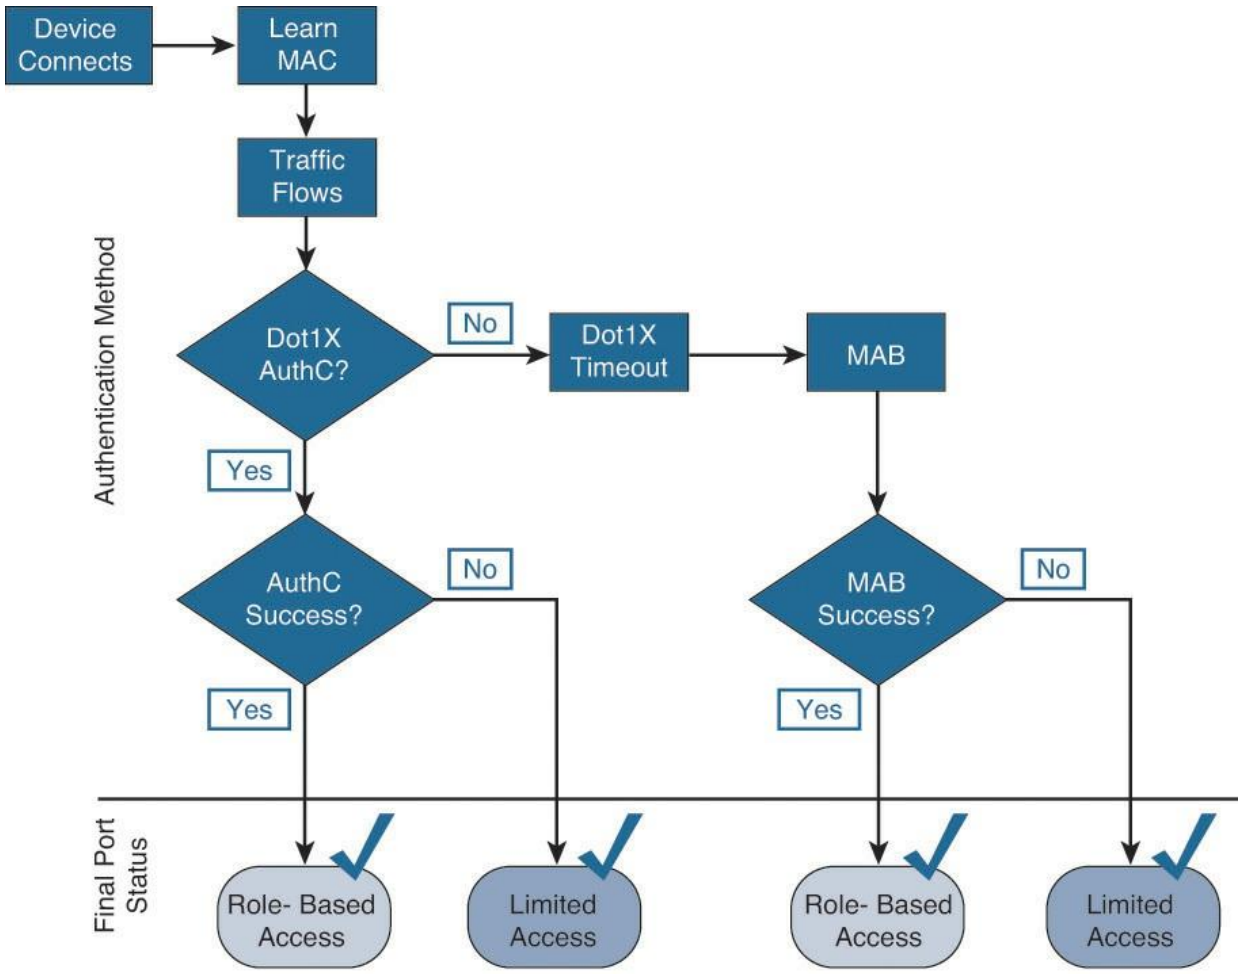


Figure 20-13 Low-Impact Mode operational flow.

As with any other security solution, tuning the authorization results is something that can take a lot of operational hours. Therefore, it is always recommended to deploy authorization results in stages as well. 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 decrease the security. Make the dACLs more specific, and so on.

Figure 20-14 shows that the pACL is applied prior to authentication, which allows only specific traffic into the port. After the authentication occurs, the authorization will need to include a dACL that selectively permits and/or denies traffic. Other authorization results also can be applied at the port, such as URL redirection, VLAN assignment, MACSec encryption, security group tags (SGTs), and more.

https://t.me/learningnets

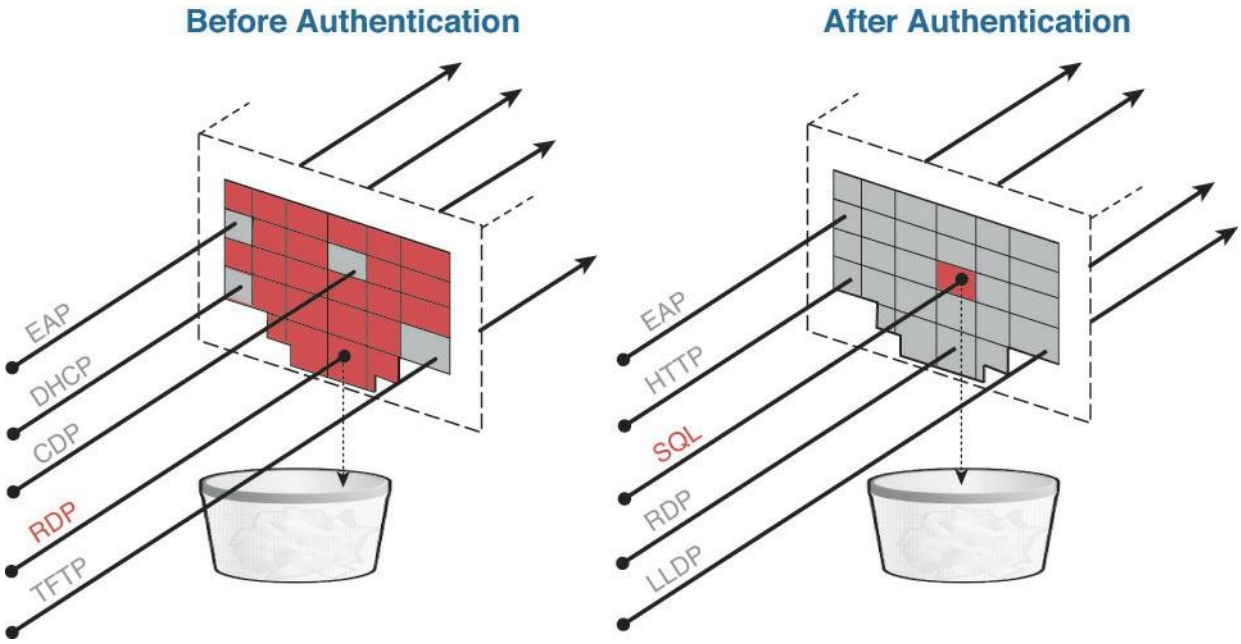


Figure 20-14 Low-Impact Mode port behavior.



Always keep in mind that VLAN assignment should be used only on devices that use supplicants. Without a supplicant, the device will most likely not be able to identify the VLAN change and might end up with the wrong IP address for its final VLAN assignment.

To create the policy set for Low-Impact Mode, do the following:

Step 1. Navigate to **Policy > Policy Set**.

Step 2. Select the default policy on the left side, and click the plus sign (+) in the upper-left corner.

Step 3. Select **Create Above**.

Step 4. In the upper-right side of the policy window, double-click the text **Enter Policy Name**, as you did in the Monitor Mode section.

Step 5. Name the new policy set **Low_Impact_Mode** and provide a description.

Step 6. Create a new condition of **DEVICE > Stage Equals Stage#Low-Impact Mode**.

Step 7. Click **Done**.

Step 8. Click **Submit**.

Figure 20-15 shows the policy set selection rule for Low-Impact Mode.

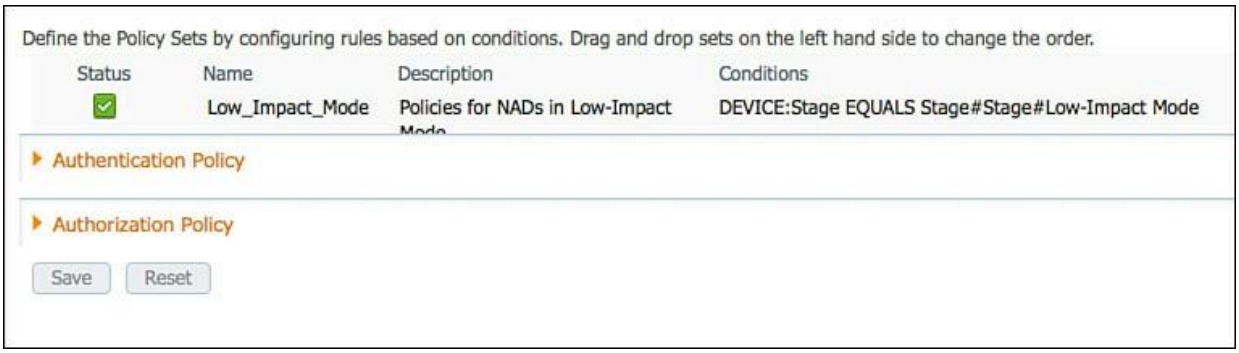


Figure 20-15 Low-Impact Mode policy set.

https://t.me/learningnets

Closed Mode

Closed Mode is similar to the default behavior of 802.1X. As shown in Figure 20-3, the port does not allow any traffic before the authentication (except for EAP, CDP, and LLDP), and then the port will be assigned to specific authorization results after the authentication.

> **Note**
>
> Closed Mode was once called High-Security Mode, but it was renamed due to the perception that it was more secure than Low-Impact Mode. In truth, both modes are equally secure. The security level of either end state is truly dependent on the configuration of the devices and the policies on ISE, not the mode of operation. In other words, an admin can make Closed Mode very insecure or very secure, depending on the implementation.

As shown previously in Figure 20-1, the operational model of 802.1X was always 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 there are so many devices that require network access without any user interaction. Consider devices such as IP cameras, IP phones, printers, fax machines, badge readers, and so much more. Therefore, MAC Authentication Bypass (MAB) had to be added to the process flow.

The concept of completely denying access to the network if authentication were to fail, or if a supplicant was 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. Additionally, if authentication were to time out (meaning there was no supplicant on the endpoint), then it 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 in spite of what the policy server said.

Figure 20-16 shows the operational flow of Closed Mode. Notice it is nearly exactly the same as Low-Impact Mode. All the same authorization results are available for use, but Closed Mode does not allow any of the PXE-type traffic into the port prior to the authorization result, unlike Low-Impact Mode.

https://t.me/learningnets

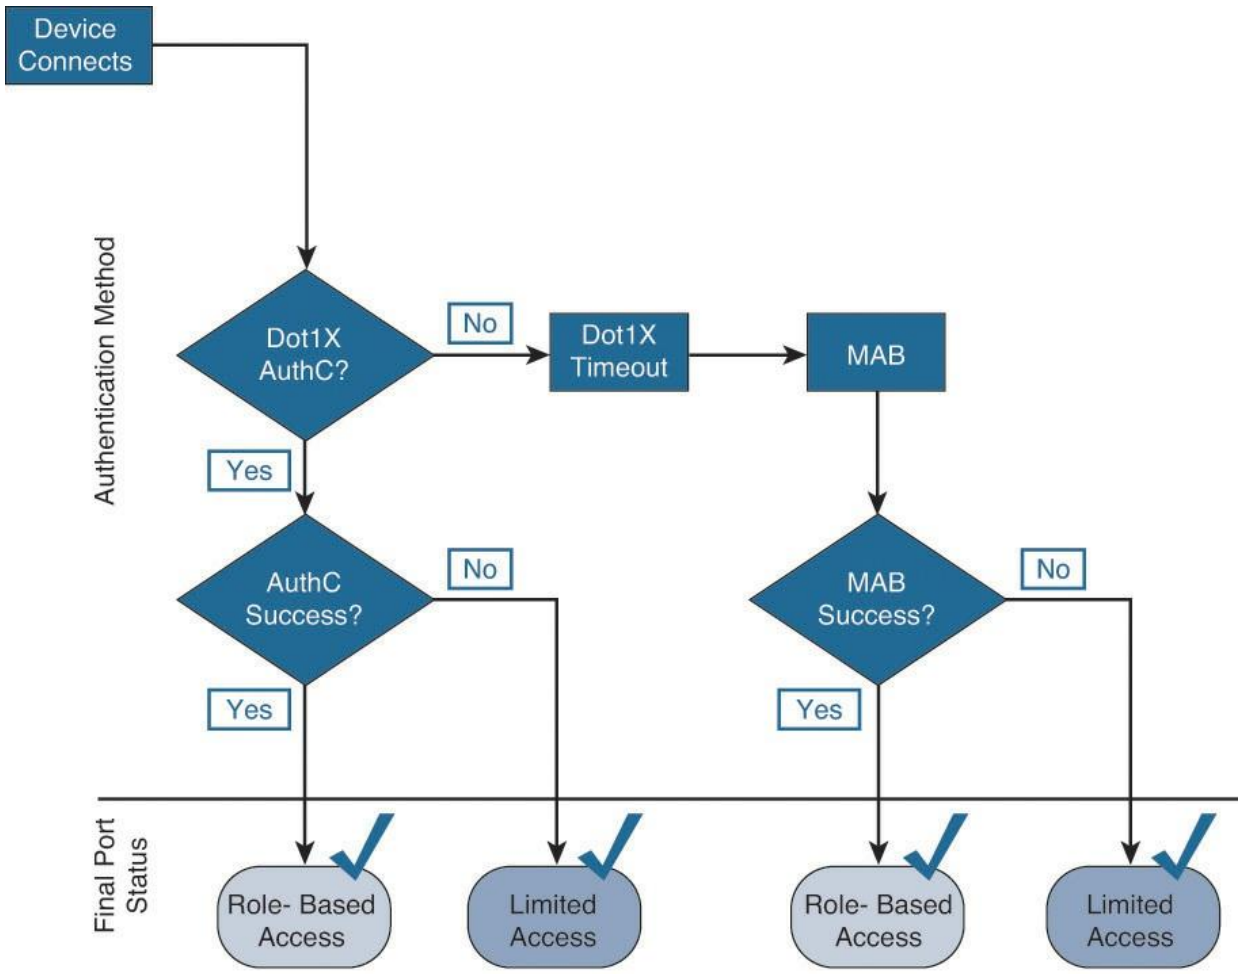


Figure 20-16 Closed Mode flow.

Figure 20-17 shows the port behavior in Closed Mode. Virtually zero traffic is allowed into the port before the authentication. After the session is authorized, very specific authorization results can be applied to the port, such as VLAN assignment, dACL, URL redirection, MACSec encryption, and SGTs.

https://t.me/learningnets

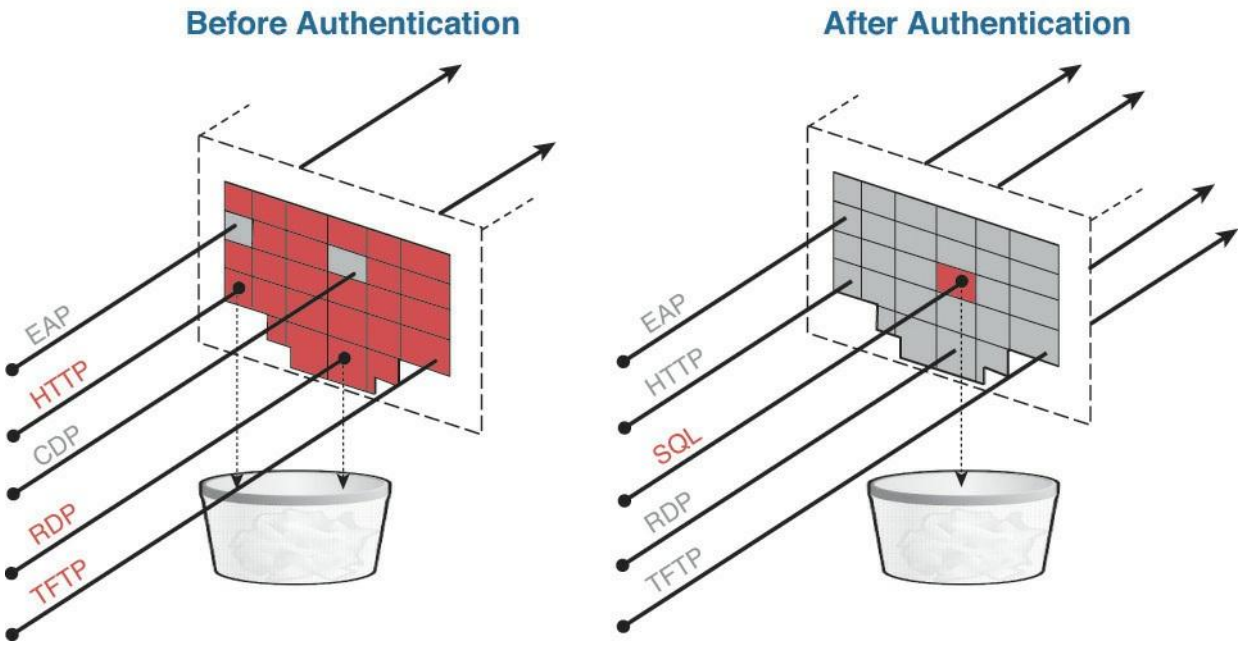


Figure 20-17 Closed Mode port behavior.

To create the policy set for Closed Mode, do the following:

Step 1. Navigate to **Policy > Policy Set**.

Step 2. Select the default policy on the left side, and click the plus sign (+) in the upper-left corner.

Step 3. Select **Create Above**.

Step 4. In the upper-right side of the policy window, double-click the text **Enter Policy Name**, as you did in the Monitor Mode and Low-Impact Mode sections.

Step 5. Name the new policy set **Closed_Mode** and provide a description.

Step 6. Create a new condition of **DEVICE :: Stage Equals Stage#Closed Mode**.

Step 7. Click **Done**.

Step 8. Click **Submit**.

Figure 20-18 shows the policy set selection rule for Closed Mode.

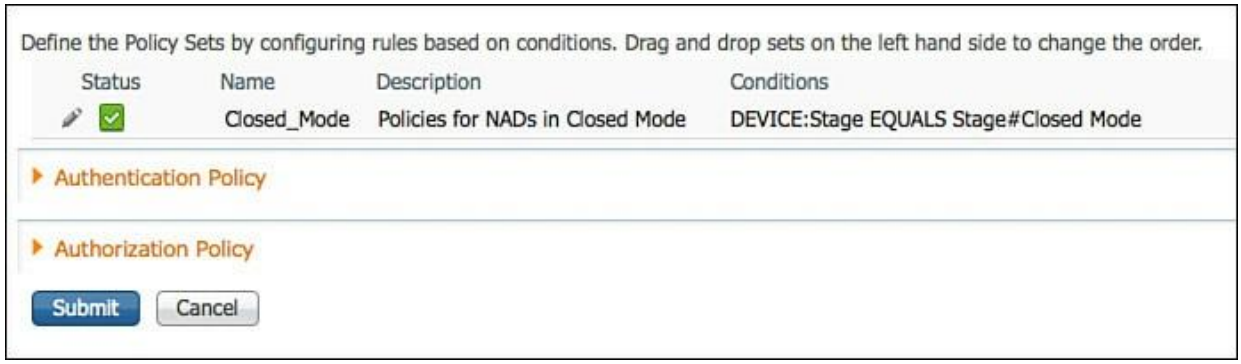


Figure 20-18 Closed Mode policy set.

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 or Closed Mode). This is why you built out NDGs and policy sets previously in this chapter.

With Monitor Mode, you must ensure that only Access-Accept and Access-Reject authorizations are used.

https://t.me/learningnets

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 the endpoints are all authenticating correctly either via 802.1X with their supplicants or via MAB with profiling or even statically. If we get the first pilot switch ready, and all the devices are prepared for authentication and everything looks good, we 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 will be fine, all the devices will be working correctly, and life looks easy.

However, we wouldn't want to push the CWA result to the switch port if we have not fully prepared the supplicants and educated the users on a possible change of experience when logging in to the network. That would be another career limiting event.

That is why we use NDGs. We ensure with the NAD's membership of the stage NDG that we 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 that branch is fully ready and all the devices are recognized and authenticating successfully, you can then 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 do. With the creation of a WLAN, the administrator must define the security for that LAN. When using 802.1X, you must 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 a guest authentication, the guest will need 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 can 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 will use a separate policy set for wireless authentications and authorizations, whereas others will assign their wireless NADs to the Low-Impact Mode or Closed Mode NDG.

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 20-2 lists a reference of these key topics and the page numbers on which each is found.



https://t.me/learningnets

Key Topic Element	Description	Page
List	Wired ports must be able to handle all types of devices with a single configuration.	686
Paragraph	**authentication open** will ignore Access-Reject, but all other authorization commands will be honored and applied.	686
Paragraph	Security is added to Low-Impact Mode by applying pACLs that are overridden by dACLs.	690

Table 20-2 Key Topics for Chapter 20

https://t.me/learningnets

Chapter 21. ISE Scale and High Availability

This chapter covers the following exam topics:

- **Configuring ISE nodes in a distributed environment**
- **Understanding the HA options available**
- **Using load balancers**
- **Maintaining ISE deployments**

> **Note**
>
> These topics are not specific exam objectives, but they are on the exam.

Although ISE can run with all services on a single physical or virtual appliance, there is still the issue of scale. A single appliance cannot scale to large enterprise levels, nor would that 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 multiple locations. Additionally, there are features within ISE and the network access devices (NADs) themselves that provide redundancy for authentication.

Availability of a solution should be a full strategy, not just a product feature. Planning for patching and backups of the solution are critical items to ensure availability.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Configuring ISE Nodes in a Distributed Environment	1-7
Understanding the HA options Available	1-7
Using Load Balancers	5-7
Maintaining ISE Deployments	8-10

Table 21-1 "Do I Know This Already?" Section-to-Question Mapping

> **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

https://t.me/learningnets

credentials of the cube controller.

b. From the Deployment screen on the PAN, click Create Cube. 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. They do not join an ISE cube.

2. True or False? When joining a node to an ISE cube, you specify which personas the node should have.

   a. True

   b. False

3. Which three pieces of information are needed for an ISE license?

   a. The output from the **show license** CLI command.

   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 HA 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, then the other nodes will 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. Monitoring persona does not have an HA function.

   d. Logs are sent to both MnT nodes automatically. If one MnT node goes down, the other node is still receiving logs.

6. What is the purpose of a node group?

   a. Node groups are used for stateful sync between PSNs. If one PSN goes down, another PSN from the node group will assume 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 will take over.

   c. Node groups are used for a multicast heartbeat between PSNs. If one PSN goes down, another PSN from the node group will send 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 goes down, another MnT from the node group will take over.

7. True or False? Cisco ISE cannot be used with load balancers.

   a. True

   b. False

8. How are patches applied to Cisco ISE?

https://t.me/learningnets

a. Patches are downloaded and applied automatically using Cisco 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 Cisco 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 can be viewed only from the CLI.

b. The status of a restore is available in the GUI, but not backup status.

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. This is configured under Administration > System > Settings > Patch Management.

b. It is configured on the Administration > System > Maintenance > Patch Management page.

c. It is not configurable and will patch all nodes simultaneously.

d. It is not configurable and will patch all nodes in alphabetical order.

Foundation Topics

Configuring ISE Nodes in a Distributed Environment

As you should recall from Chapter 7, "Cisco Identity Services Engine Architecture," ISE has multiple personas or functions. Table 21-2 contains a simple chart of personas and which function that persona provides. The third column in Table 21-2 provides the friendly name for an ISE node that is running that service or persona.

ISE Persona	Description	Node Name
Admin	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 Administrative Node (PAN)
Policy Services	Provides the RADIUS communication, performing the authentications and authorizations, and hosts all user-facing web portals.	Policy Services Node (PSN)
Monitoring	Provides monitoring, log collection, and troubleshooting functions for the deployment.	Monitoring and Troubleshooting Node (MnT)

Table 21-2 ISE Persona Types

All ISE nodes are installed in a standalone mode by default. When in a standalone mode, the ISE node will be configured to run all personas by default. That means that the standalone node will run administration, monitoring, and policy services.

It is up to you, the ISE administrator, to promote the first node to be a primary and then join the additional nodes to this new deployment (known as an ISE cube). At the time of joining, you will also determine which services will run on which nodes; in other words, you will determine which persona the node will have.

https://t.me/learningnets

Making the First Node a Primary Device

Because all ISE nodes are standalone by default, we must first promote the ISE node that will become the primary policy administration node to be a primary device, instead of a standalone.

From within the ISE GUI, do the following:

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 21-1.



Figure 21-1 Deployment screen.

Step 3. Click the **Make Primary** button, as shown in Figure 21-2.

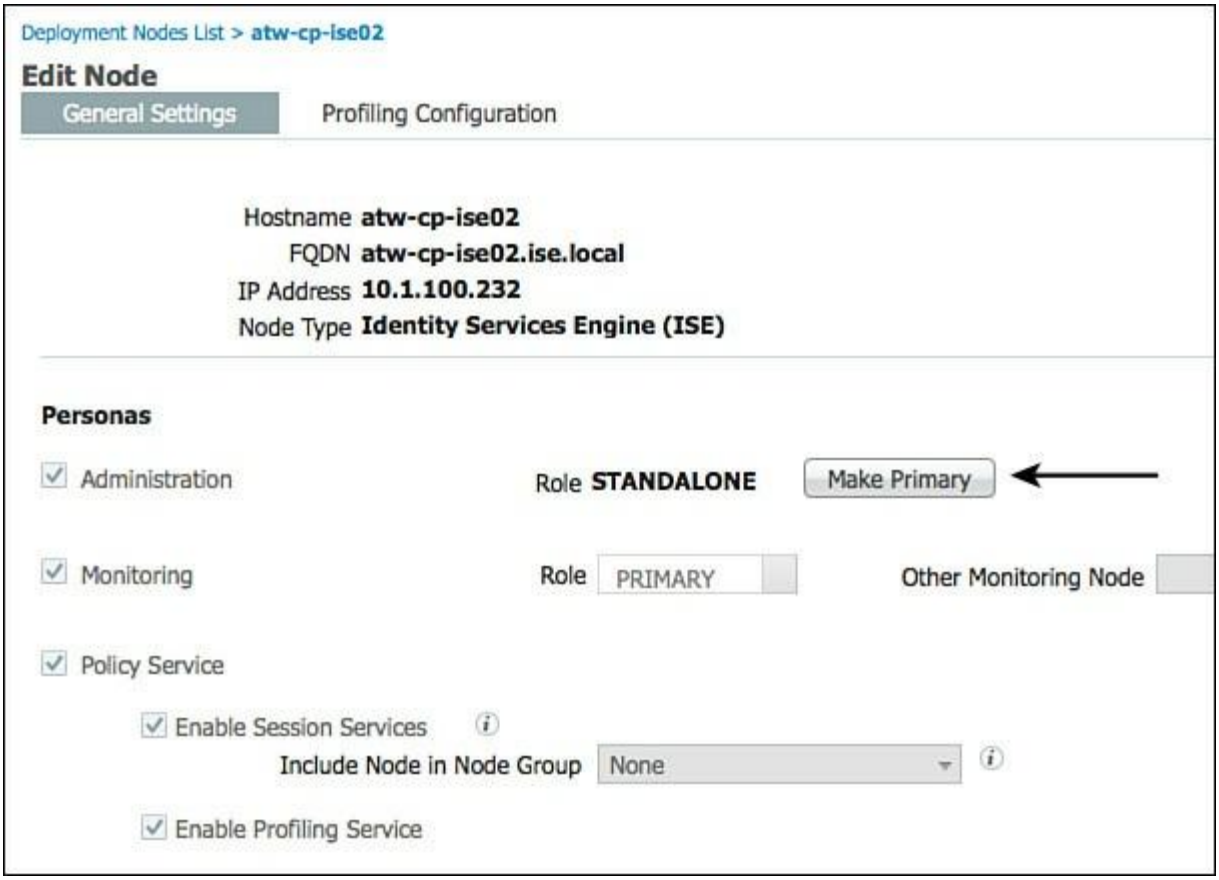


Figure 21-2 Make Primary button.

At this point, the Monitoring and Policy Service check boxes 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**.

https://t.me/learningnets

Registering an ISE Node to the Deployment

Now that there is a primary administrative node, you can finally have a multinode deployment. From the GUI on the primary admin node, you will register and assign personas to all ISE nodes.

> **Note**
>
> 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. Refer to Chapter 9, "Initial Configuration of Cisco ISE," for a detailed walkthrough of replacing the self-signed certificates with a signed certificate.

To register a new node to the primary and form your ISE cube, follow these steps:

Step 1. Navigate to **Administration > System > Deployment**.

Step 2. Click **Register > Register an ISE Node**, as shown in Figure 21-3.

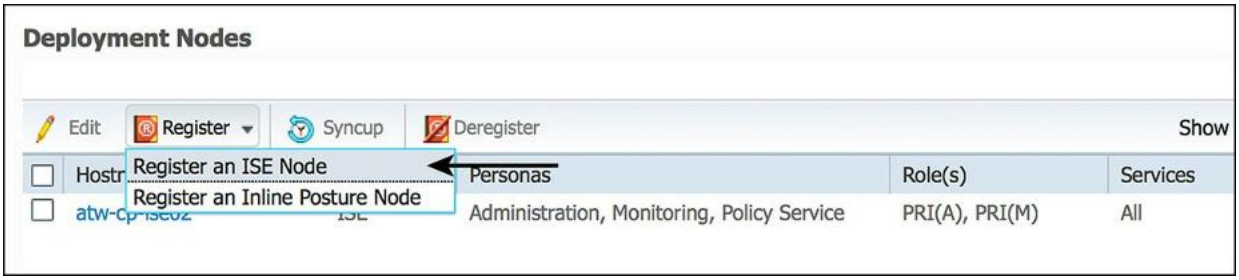


Figure 21-3 Register an ISE node.

> **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 (admin by default) and password. Figure 21-4 shows the registration page, with the hostname and credentials entered.

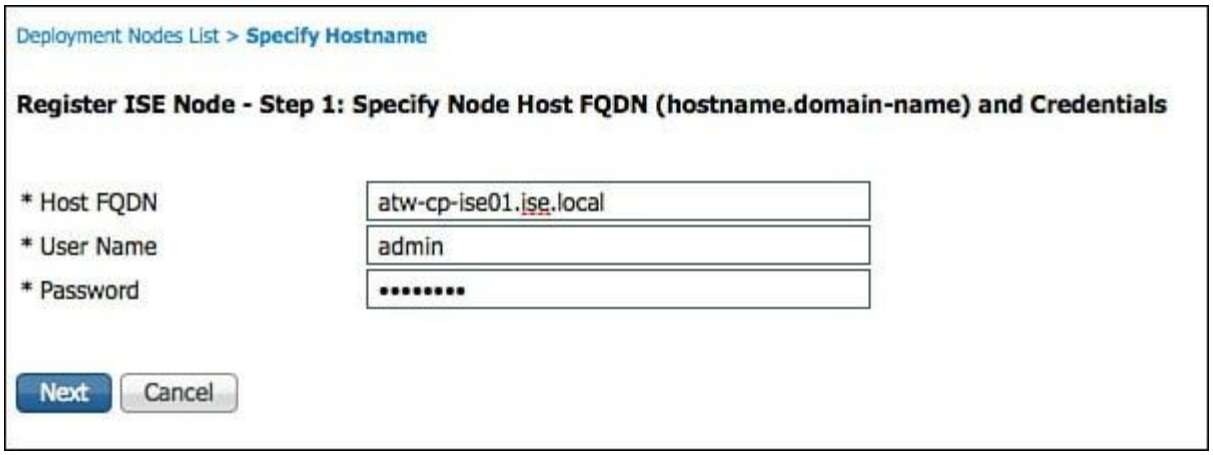


Figure 21-4 Specify the hostname and credentials.

Step 5. Click **Next**.

The next screen is the Configure Node screen, as shown in Figure 21-5. On this screen, you choose the main persona of the ISE node, including the enabling of profiling services. However, you cannot configure the individual profiling probes at this point.

https://t.me/learningnets

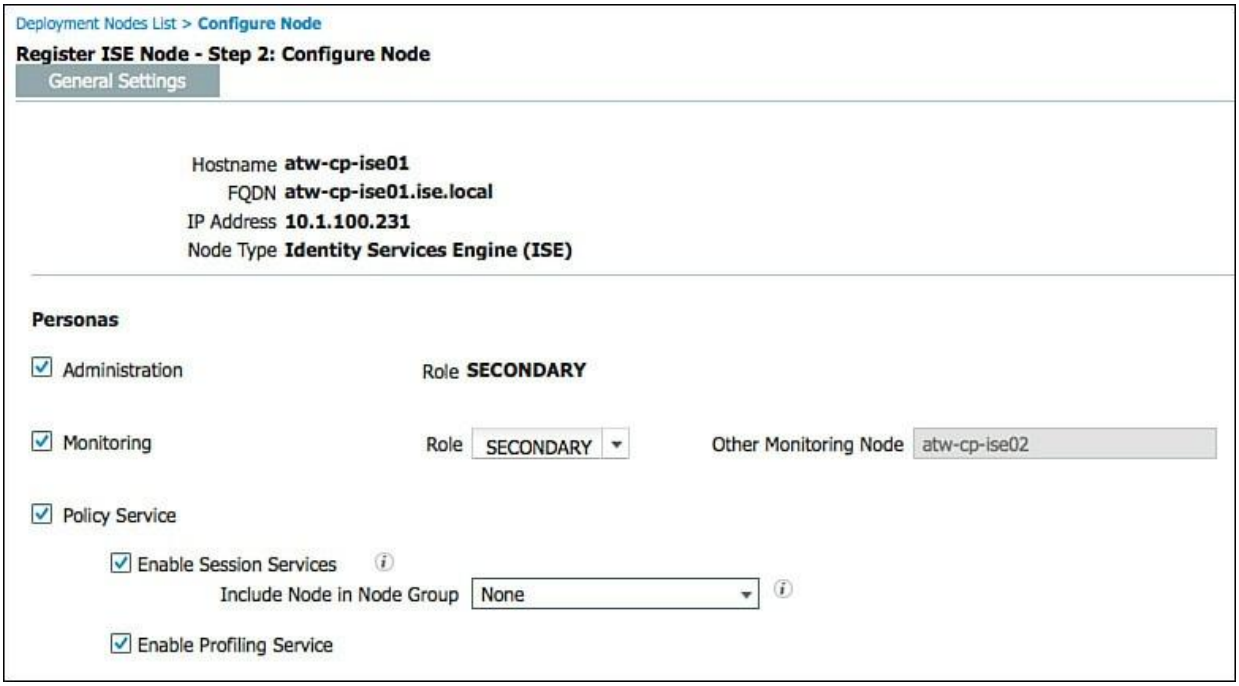


Figure 21-5 Configure node.

Step 6. Choose the personas to run on the newly registered node.

Notice in Figure 21-5 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 **Next**.

At this point, the primary PAN will sync the entire configuration and the databases to the newly joined ISE node. A success message will be displayed communicating the impending synchronization, as shown in Figure 21-6.



Figure 21-6 Sync initiated.

Step 8. Repeat steps 2–7 for all the ISE nodes that should be joined to the same cube.

Ensuring the Personas of All Nodes Are Accurate

Now that all your ISE nodes are joined to the deployment, you can ensure that the correct personas are assigned to the appropriate ISE nodes. Table 21-3 shows the ISE nodes in our sample deployment and the associated persona that will be assigned.

https://t.me/learningnets

ISE Node	Persona	Comments
atw-cp-ise01	Admin, Monitoring	Pri PAN, Pri MnT
atw-cp-ise02	Admin, Monitoring	Sec PAN, Sec MnT
atw-cp-ise03	Policy Services	Policy and profiling services
atw-cp-ise04	Policy Services	Policy and profiling services

Table 21-3 ISE Nodes and Personas



This is also a good time to double-check that all the desired probes are enabled on the PSNs. Many customers have called Cisco TAC complaining of "database not replicating" problems, but it actually turned out to be that the probes were not enabled. Remember that even though you enable profiling services when joining a PSN to the ISE cube, none of the actual profiler probes are enabled by default. For detail on enabling profiler probes, see Chapter 15, "Profiling."

Table 21-3 lists the ISE nodes used within this book's examples and their personas; Figure 21-7 shows the Deployment screen with all nodes joined to the ISE cube.



Figure 21-7 Final personas and roles.

Licensing in a Multinode ISE Cube

The good news is that ISE uses a centralized pool of licenses, where all nodes within an ISE cube share the single set of licenses (base and advanced). Therefore, there is no need to focus on how many licenses exist per PSN or anything of that nature.

However, you will need to ensure that licenses are issued with both the primary and secondary administrative nodes listed. When requesting licenses, you should supply three pieces of information about each administrative node: the product ID (SPID), version number (VPID), and serial number.

Example 21-1 shows logging in to an ISE node and using the **show udi** command to obtain this information from the primary administrative node (atw-cp-ise02). Example 21-2 repeats the same command on the secondary administrative node (atw-cp-ise01).

Example 21-1 show udi Command and Output from the Primary Administrative Node

```
atw-cp-ise02/admin# show udi

SPID: ISE-VM-K9
VPID: V01
```

https://t.me/learningnets

```
Serial: 3IRJG5IUQJC

atw-cp-ise02/admin#
```

Example 21-2 show udi Command and Output from the Secondary Administrative Node

```
atw-cp-ise01/admin# show udi

SPID: ISE-VM-K9
VPID: V01
Serial: JMQJIIKDGHK

atw-cp-ise01/admin#
```

Ensuring that the license is created with both administrative nodes in the license key helps to ensure there are no issues in the case of a disaster involving the loss of the primary policy administrative node.

Understanding the HA Options Available

When it comes to high availability within a secure access solution, you need to make note of several things. You need to consider the communication between the PAN and the other ISE nodes, for database replications and synchronization, as well as authentication sessions reaching the PSNs in the event of a WAN outage. You should also think about a NAD recognizing that a PSN might no longer be active and sending authentication requests to the active PSN instead.

Primary and Secondary Nodes

The policy administration nodes (PANs) and monitoring and troubleshooting nodes (MnTs) have a concept of primary and secondary nodes, but they operate very differently.

Monitoring and Troubleshooting Nodes

Let's start with the easiest one first, the MnT. As you know, the MnT node is responsible for the logging and reporting functions of ISE. All PSNs send their logging data to the MnT as Syslog messages (UDP/20514).

When two monitoring nodes exist in an ISE deployment, all ISE nodes send their audit data to both MnTs at the same time. The logging flows between the ISE nodes is illustrated in Figure 21-8.

https://t.me/learningnets

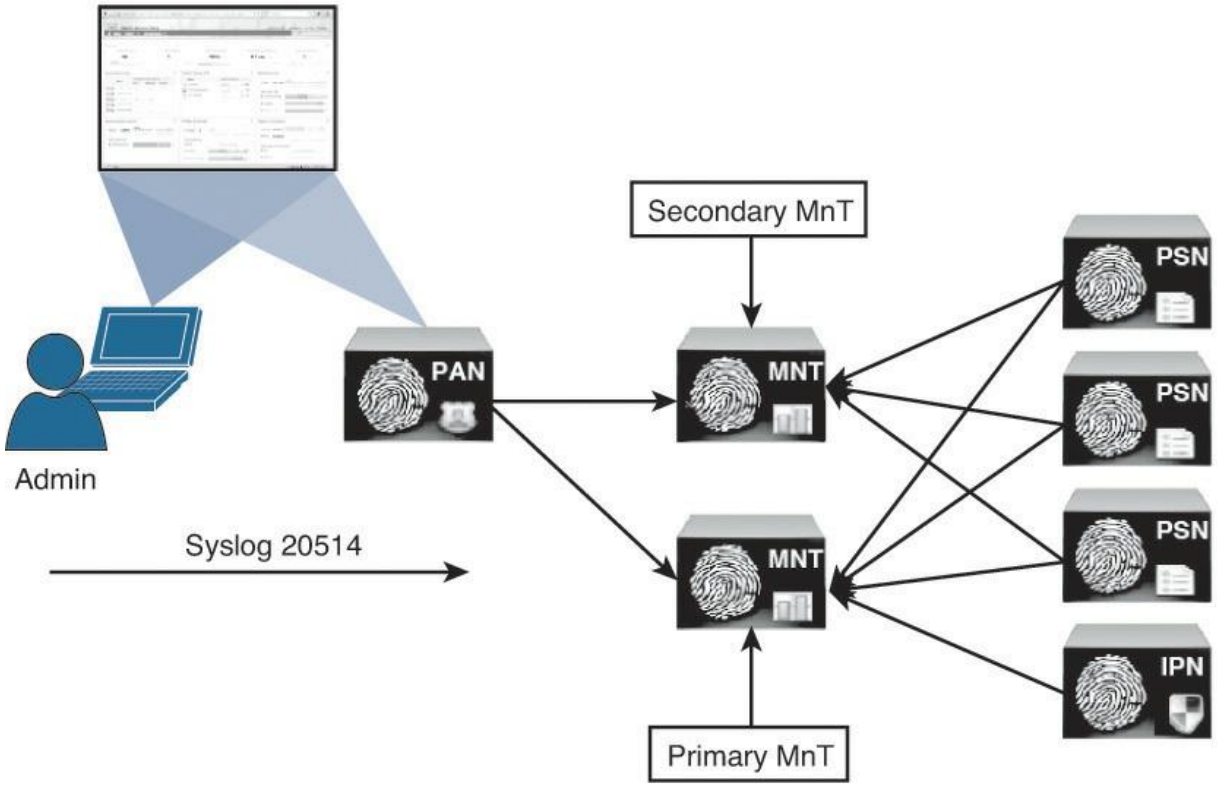


Figure 21-8 Logging flows.

> **Note**
>
> Inline posture nodes (IPNs) send syslog to only a single target.

The active/active nature of the MnT nodes can be viewed easily in the administrative console under **Administration > System > Logging > Remote Logging Targets**. Figure 21-9 shows the two MnTs being defined as LogCollector and LogCollector2.

Remote Logging Targets

Edit | Add | Duplicate | Delete | Show All

Name	IP Address	Port	Type	Description	Status
LogCollector	10.1.100.232	20514	UDP SysLog	Syslog Target for Log Collector	Enabled
LogCollector2	10.1.100.231	20514	SysLog	Second Syslog Target for Log Collector	Enabled
ProfilerRadiusProbe	127.0.0.1	30514	Profiler SysLog	Syslog Target for Profiler RADIUS Probe	Enabled
Splunk	10.1.100.98	514	UDP SysLog	SIEM solution	Enabled
TCPLogCollector	10.1.100.232	1468	TCP SysLog	TCP SysLog collector	Disabled

Figure 21-9 Logging targets.



You'll notice within **Administration > System > Logging > Logging Categories** that both LogCollector and LogCollector2 are automatically listed as targets for the logs, as shown in Figure 21-10. This ensures all ISE nodes in the cube send logs to both MnT nodes.

https://t.me/learningnets

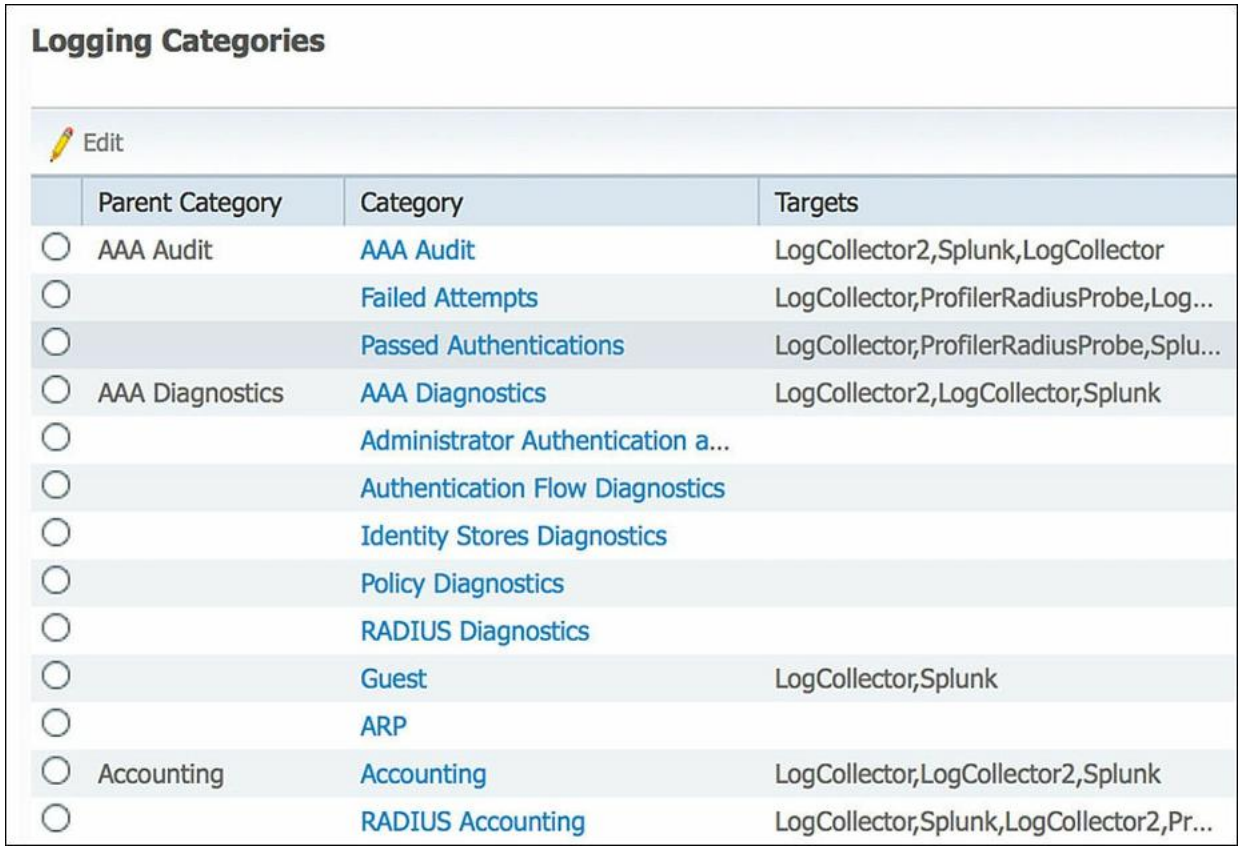


Figure 21-10 Logging categories.

Upon an MnT failure, all nodes will continue to send logs to the remaining MnT node. Therefore, no logs are lost. The PAN will retrieve all log and report data from the secondary MnT node, so no administrative function loss occurs, 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 sync.

> **Note**
>
> The best practice for logging is to also send logging data to a security information manager (SIM) tool for long-term data archiving and reporting.

Policy Administration Nodes

The policy administration node is not only responsible for providing an administrative GUI for ISE, but is also responsible for the critical function of database synchronization of all ISE nodes. All ISE nodes maintain a full copy of the database, with the master database existing on the primary PAN.

A policy services node can get new data about an endpoint, 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 PAN can have many updates to send to each node, those nodes must acknowledge the last update before a new update can be sent.

Because the functionality is so arduous and it is absolutely critical to have only a single source of truth for the data in the database, it is a manual process to fail over to the secondary PAN. In the event of the primary PAN going offline, no synchronizations will occur until the secondary PAN is promoted to primary. After it becomes the primary, it will take over all synchronization responsibility. Some might consider this a “warm spare” type of high availability (HA).

To promote the secondary PAN to the primary role, connect to the GUI on the secondary PAN by doing the following:

https://t.me/learningnets

Step 1. Navigate to **Administration > System > Deployment**.

Step 2. Click **Promote to Primary**, as shown in Figure 21-11.

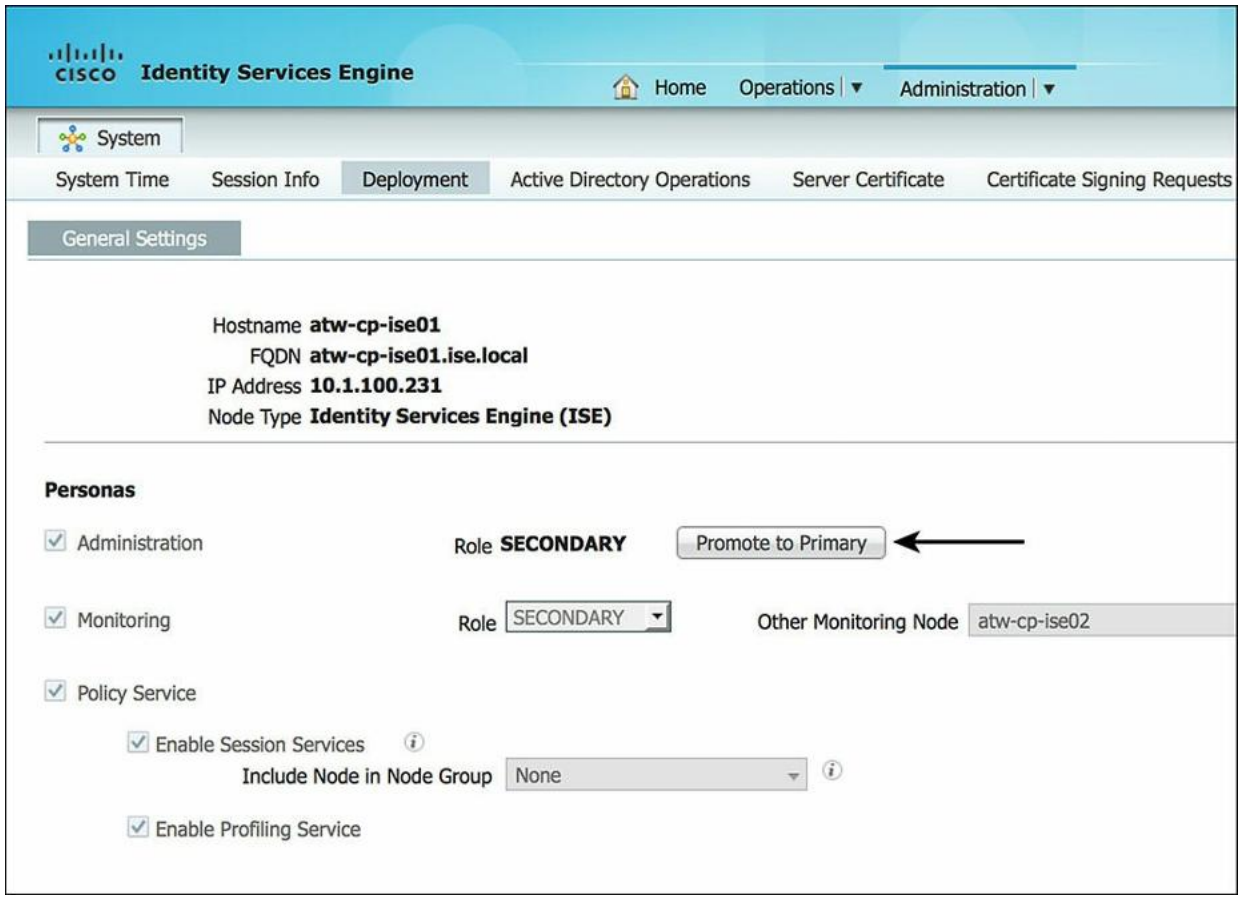


Figure 21-11 Promote to primary.

Node Groups

PSNs do not necessarily need to have an HA type of 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 does have a concept of a node group. Node groups are made up of Layer-2 adjacent (same VLAN) PSNs, where the nodes 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.

This is most commonly used when deploying the PSNs behind a load balancer. However, there is no requirement for node groups to be used solely with any Layer-2 adjacent PSNs.

To configure a node group, do the following:

Step 1. Navigate to **Administration > System > Deployment**.

Step 2. On the left side, click the cog and select Create Node Group, as shown in Figure 21-12.

https://t.me/learningnets

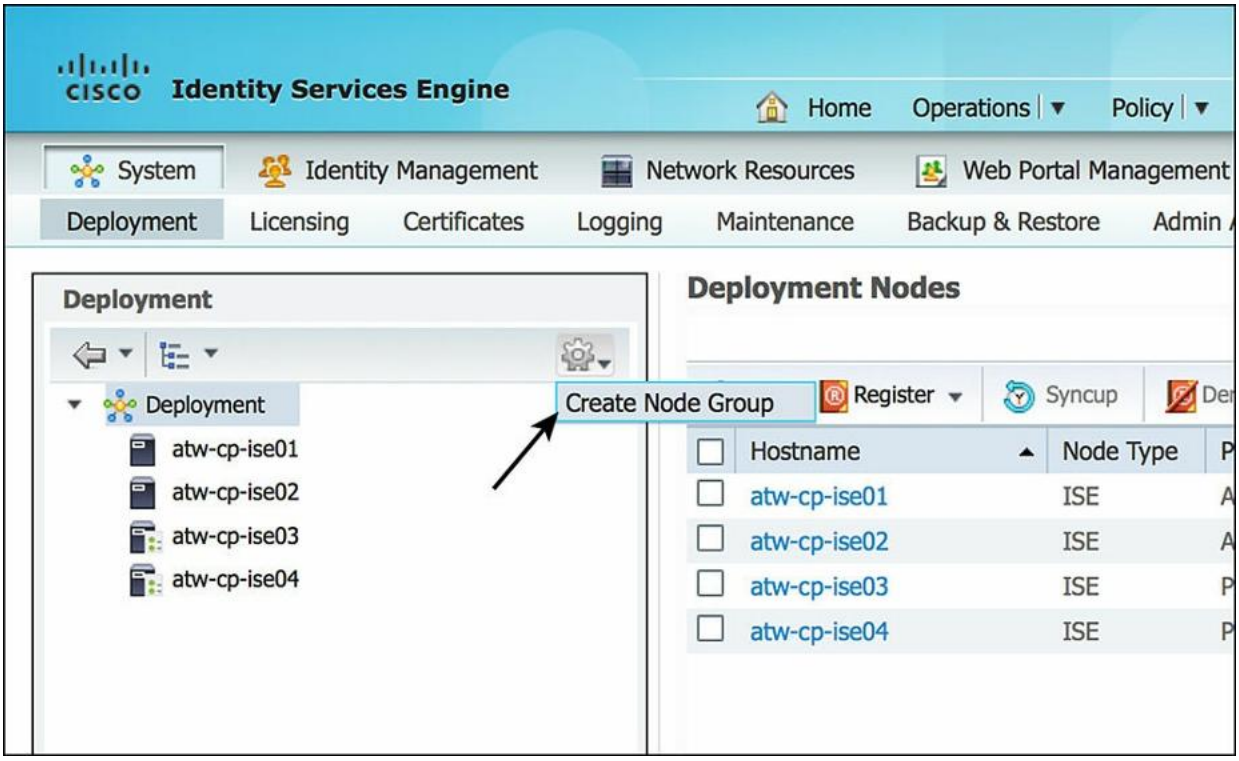


Figure 21-12 Create node group.

Step 3. Enter a name for the node group. Use a name that also helps describe the location of the group.

Step 4. Enter a more detailed description, helping to identify exactly where the node group is—for example, **DataCenter 1 Node Group**.

Step 5. Select a multicast address for keep-alive communication. As the screen dictates, ensure it is not a multicast address already in use, as shown in Figure 21-13.

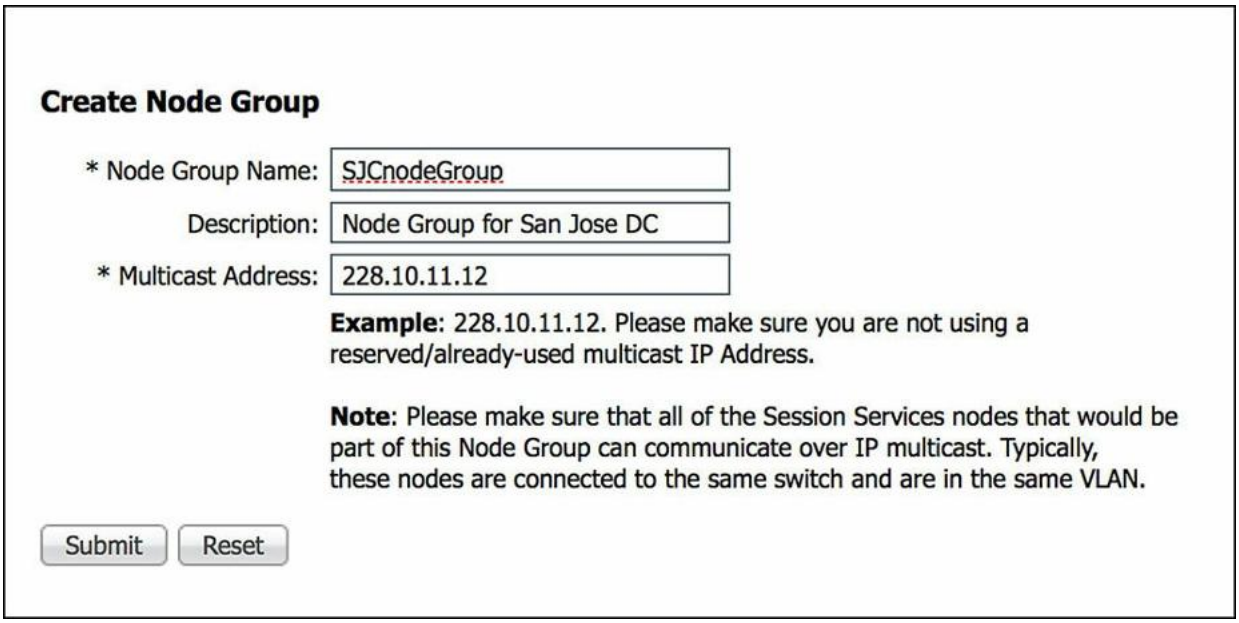


Figure 21-13 Node group creation.

Step 6. Click **OK** on the success pop-up window. You will also notice the appearance of the node group on the left side.

Now that the node group exists, it's time to add the PSNs to the node group.

https://t.me/learningnets

Step 7. Navigate to **Administration > System > Deployment**.

Step 8. Select one of the PSNs to add to the node group.

Step 9. Click the node group drop-down and select the newly created group, as shown in Figure 21-14.

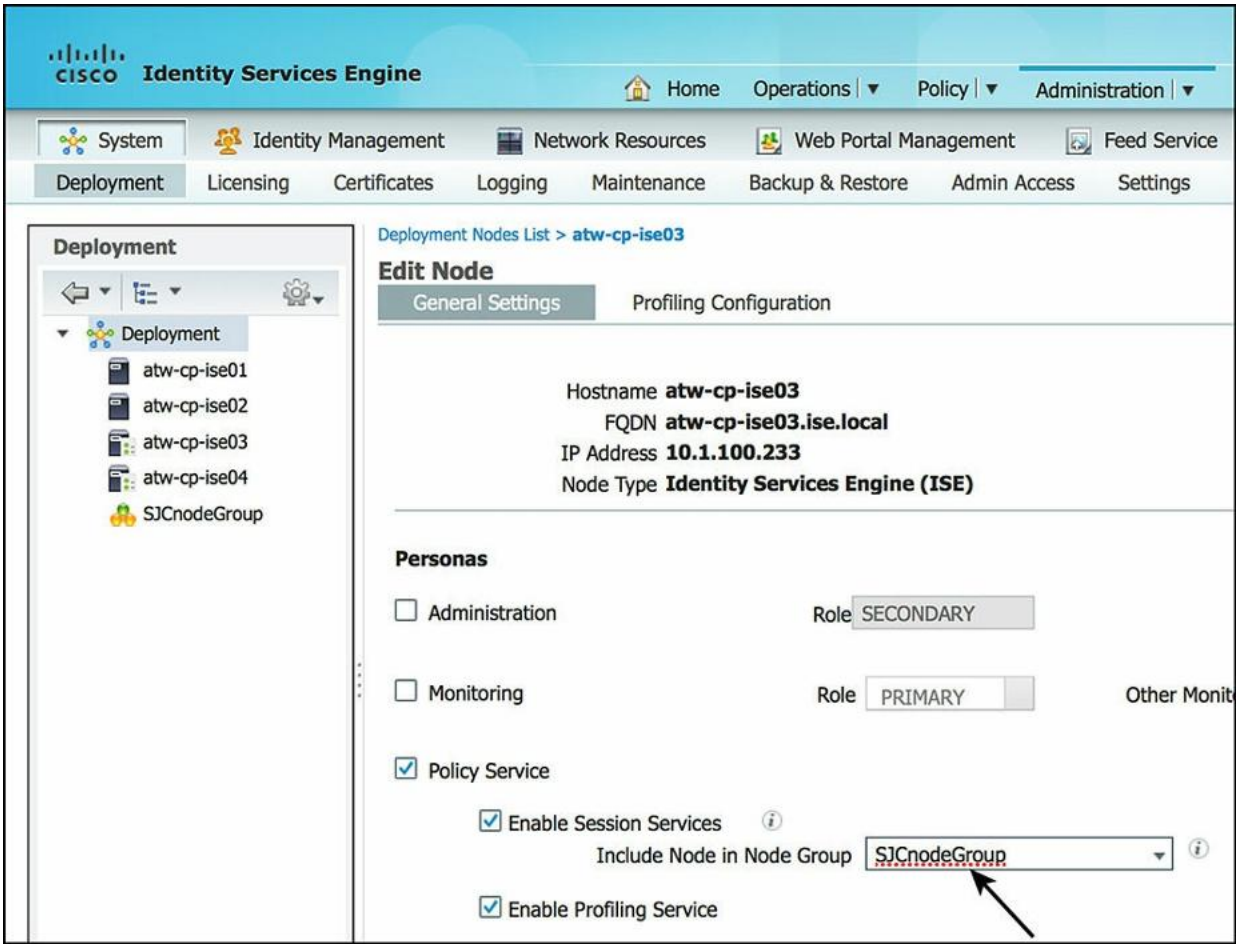


Figure 21-14 Assign node group.

Step 10. Click **Save**.

Step 11. Repeat steps 7–10 for each PSN that should be part of the node group. Figure 21-15 shows the final deployment tree after the two PSNs (atw-cp-ise03 and atw-cp-ise04) are added to the node group.

https://t.me/learningnets

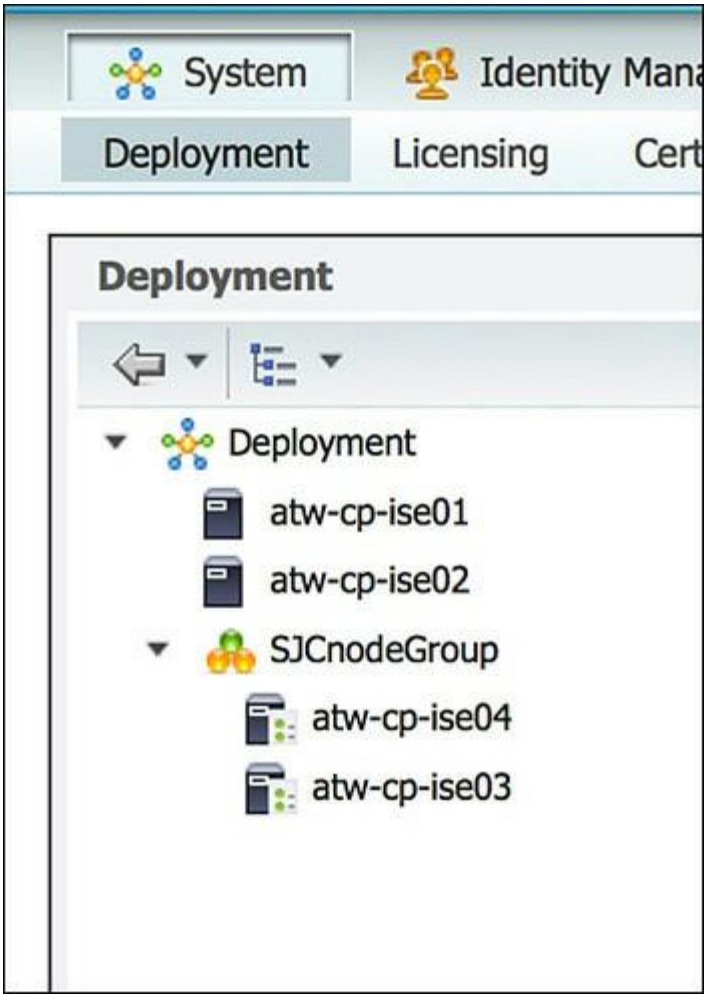


Figure 21-15 Deployment tree.

Using Load Balancers

A somewhat controversial and often debated topic is the use of load balancers with ISE policy services nodes. The reason that this becomes controversial is due to the server load balancer teams being used to installing using a certain architecture and the ISE solution requiring a different architecture.

The first question is usually why? Ultimately, the complexities with load balancing are related to the many types of traffic that must flow between the endpoint and the active PSN.

General Guidelines

There are a few general guidelines to follow when using a load balancer. Of course, a specific design can and will get more specific and more detailed than this list, but overall the rules of thumb are as follows:

- Each PSN must be reachable by the PAN/MnT directly, without having to go through NAT.
  - This can sometimes be referred to as Routed Mode or Pass-Through Mode.
- You cannot use Source NAT (SNAT). ISE uses the Layer-3 address to identify the NAD, not the NAS-IP address in the RADIUS packet.
  - SNAT would hide the real IP address of the switch and NAT it to a VIP address. Therefore, you cannot use Source NAT in the load balancer design.
- Each PSN must also be reachable directly from the endpoint.
  - When the PSN sends a URL redirection to the NAD, it will use the FQDN from the ISE node's configuration (that is, atw-cp-ise03.ise.local), not the IP address of the VIP.

https://t.me/learningnets

- The certificates used on the PSNs should include a DNS name in the subject alternative name (SAN) field that resolves to the IP address of the VIP.
  - For example, ise.company.com
- The load-balancing rules must ensure that the same PSN is used for the entire session.
  - Use persistence, sometimes called "sticky" based on Calling-Station-ID and Framed-IP-address.
  - Some load balancers have built-in intelligence that allows the same persistence value to persist across both RADIUS and DHCP. This should be used if the capability exists.
- The virtual IP (VIP) will be listed as the RADIUS server of each NAD for all 802.1X-related AAA.
  - Includes both authentication and accounting packets.
  - Some load balancers use a separate VIP for each protocol type.
- The list of RADIUS servers allowed to perform dynamic authorizations (CoA) on the NAD should use the real IP addresses of the PSNs.
  - If the load balancer is capable of using NAT for the PSN IP address for the CoA traffic destined to the network device, then only the VIP IP address would need to be included in the list of dynamic authorization (CoA).
  - The load balancer(s) will also get listed as NADs in ISE so their test authentications can be answered.
- Load balancers should be configured to use test probes to ensure the PSNs are still "alive and well."
  - A probe should be configured to ensure RADIUS is responding.
  - HTTPS should also be checked.
  - If either probe fails, the PSN should be taken out of service.

Failure Scenarios

If a single PSN fails, the load balancer should take that PSN out of service and spread the load over the remaining PSNs. When the failed PSN is returned to service, the load balancer will add it back into the rotation. By using node groups along with a load balancer, one of the other node group members will issue a CoA-reauth for any sessions that were establishing. This CoA will cause the session to begin again. At this point, the load balancer should direct the new authentication to a different PSN.



NADs have some built-in capabilities to detect when the configured RADIUS server is "dead" and 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 is unreachable (that is, the load balancer has died), the NAD should quickly fail over to the next RADIUS server in the list. That RADIUS server could be another VIP in a second datacenter or another backup RADIUS server.

IOS Load Balancing

Cisco network devices have a lot of intelligence built in to 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.

What it's referring to is the ability of the access layer switch to load-balance the outbound authentication requests for endpoints that are authenticated to the switch itself, as illustrated in Figure 21-16.

https://t.me/learningnets

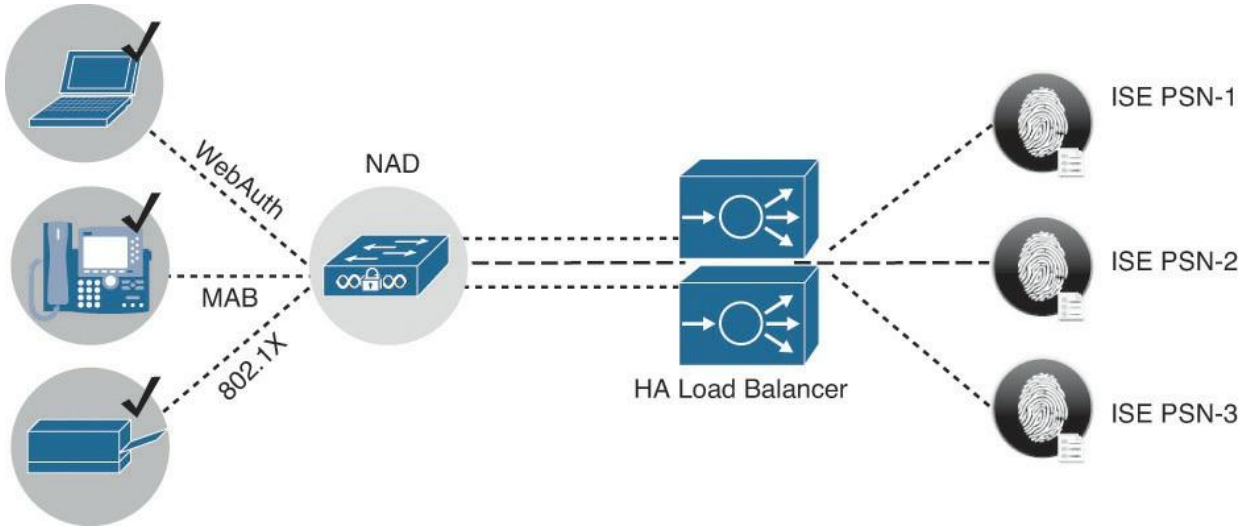


Figure 21-16 IOS load-balancing.

Enabling IOS RADIUS server load-balancing takes only one additional command. After all the PSNs are defined as AAA servers in the switch, use the **radius-server load-balance** global configuration command to enable it.

Example 21-3 shows use of a **show** command to verify that multiple ISE servers are configured. Example 21-4 shows how to enable the RADIUS server load-balancing.

Example 21-3 Verifying That All ISE PSNs Are Configured on Switch

Click here to view code image

```
3750-X#show aaa server | i 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 21-4 The Enabling of IOS Load-Balancing

Click here to view code image

```
Example 21-4 Enabling IOS Load Balancing
3750-X(config)# radius-server load-balance method least-outstanding batch-size 5
```

Maintaining ISE Deployments

Although having a distributed deployment and load-balanced architecture is certainly critical items to scaling the deployment and ensuring it is highly available, there are critical basic maintenance items that should always be considered to ensure the most uptime and stability.

That means having a patching as well as a backup and restore strategy.

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 recently discovered with SSL. To ensure that bug fixes are applied, security vulnerabilities are plugged, and the solution will work as seamlessly as possible, always have a planned patching strategy.

Patches are downloaded from cisco.com and are found under **Downloads > Products > Security > Access**

https://t.me/learningnets

Control and Policy > Identity Services Engine > Identity Services Engine Software, as shown at the top of Figure 21-17.

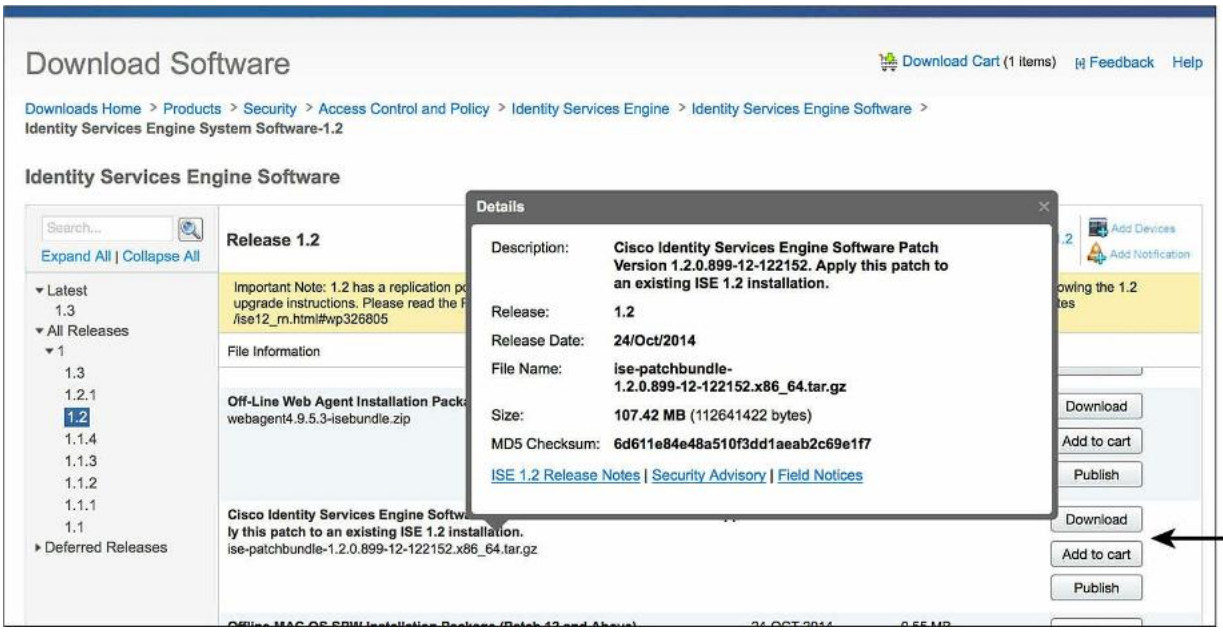


Figure 21-17 ISE Downloads page.

Search the list of software available for your specific version of ISE. Figure 21-18 illustrates the naming convention for ISE patches. Cisco ISE patches are normally cumulative, meaning that installing 1.2 patch 12 will include all the fixes in patches 1–11 as well. The following steps walk you through ISE patching.



Figure 21-18 Anatomy of ISE patch nomenclature.

Step 1. Download the required patch.

Step 2. From the ISE GUI, navigate to **Administration > System > Maintenance > Patch Management**.

Step 3. Click the **Install** button, as shown in Figure 21-19.

https://t.me/learningnets

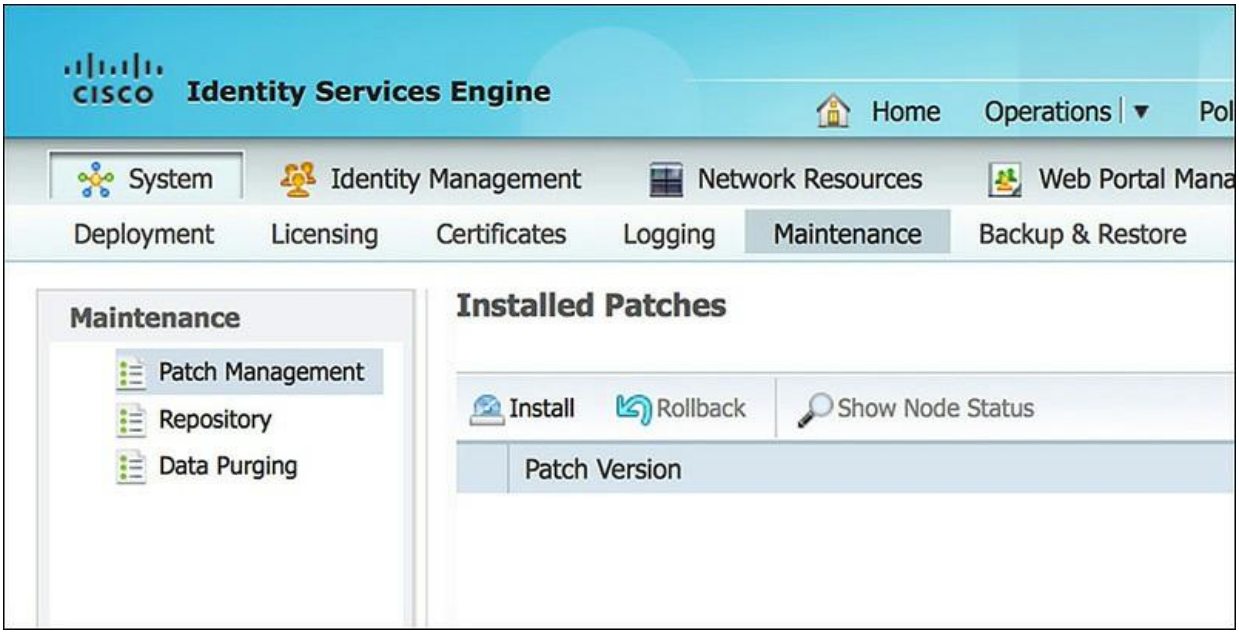


Figure 21-19 Patch Management screen.

Step 4. Click **Browse**, select the downloaded patch, and click **Install**, as shown in Figure 21-20.



Figure 21-20 Install the selected patch.

The patch bundle is now uploaded through the web UI, and you are prompted to confirm the MD5 hash of the file, as shown in Figure 21-21.



Figure 21-21 Verify MD5 hash.

Step 5. Click **OK** if the MD5 hash matches.

As the patch is installed on the PAN, you are logged out of the GUI and the patch is distributed from the

https://t.me/learningnets

PAN to all nodes in the ISE cube. The patch will be applied to all nodes in the cube one at a time, in alphabetical 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 21-19 to verify the progress of the patching. Figure 21-22 shows the resulting status of each nodes progress for the patch installation.

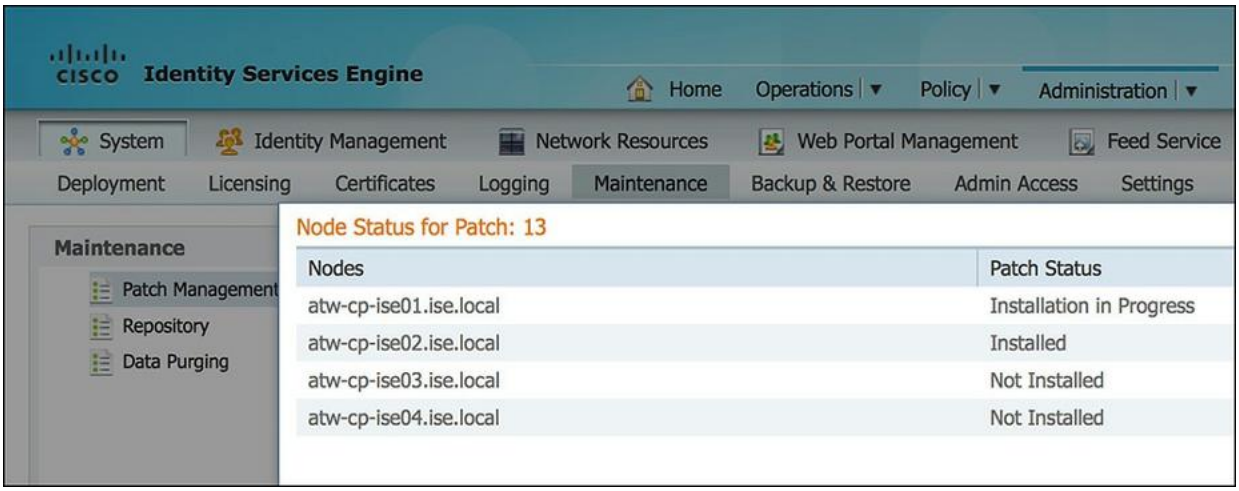


Figure 21-22 Node status.

Backup and Restore

Another key strategy to ensuring the availability of ISE in the environment is having a solid backup strategy. There are two types of ISE backups: configuration and operational backup. These two types are most easily related to backing up the product databases (configuration) and the MnT data (operational).

Figure 21-23 shows the backup screen in ISE, located at **Administration > System > Backup & Restore**.

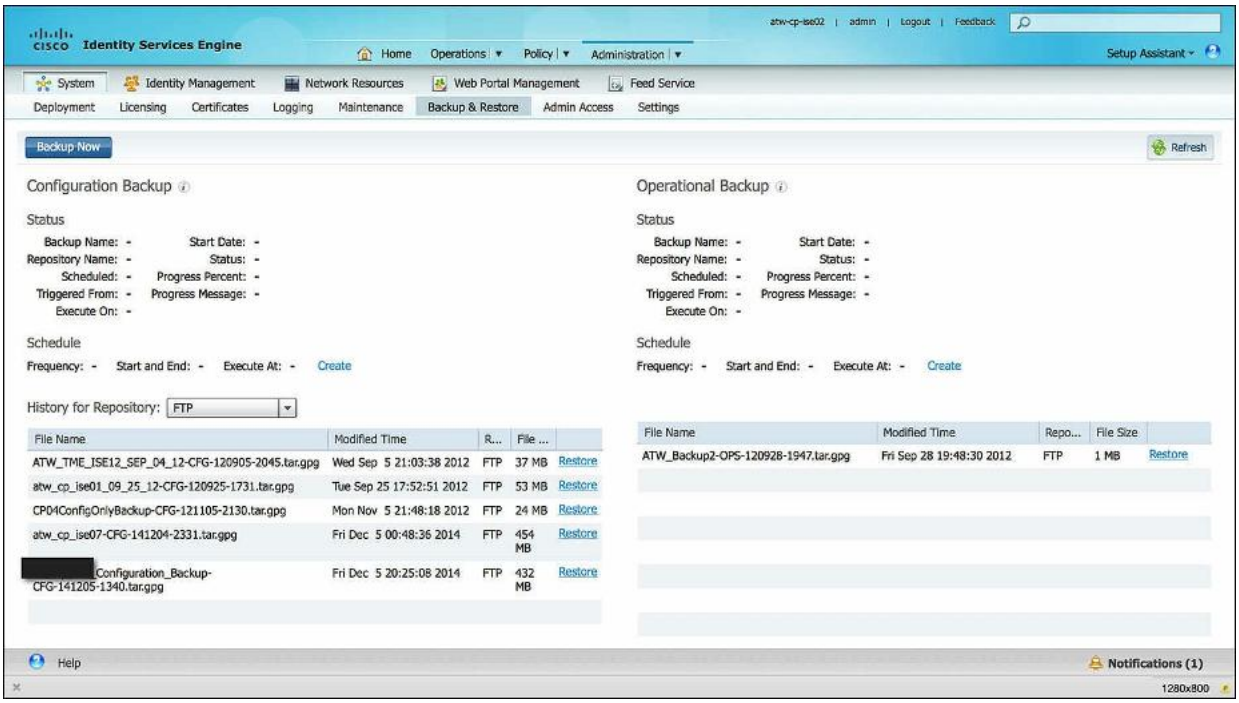


Figure 21-23 Backup and Restore screen.

As shown in Figure 21-23, the backups are stored in a repository and can be restored from the same repository. Backups can be scheduled or run on demand. The status of a backup can be viewed from the GUI or the CLI, but the status of a restore can be viewed only from the CLI.

https://t.me/learningnets

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 21-4 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Paragraph	Profiler probes configured per node.	706
Paragraph	All nodes send syslog to both MnTs automatically.	708
Paragraph	Node groups.	710
Paragraph	NADs have a built-in fail to next RADIUS server capability.	715

Table 21-4 Key Topics for Chapter 21

Define Key Terms

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

node groups

https://t.me/learningnets

Chapter 22. Troubleshooting Tools

This chapter covers the following exam topics:

- **Identify issues using authentication event details in Cisco ISE**
- **Troubleshoot using Cisco ISE diagnostic tools**
- **Troubleshoot endpoint issues**
- **Use debug commands to troubleshoot RADIUS and 802.1X on IOS switches and wireless controllers**

Throughout this book, you have been verifying activities using tools that are commonly used for troubleshooting. Tools such as Live Log provide immeasurable value to an ISE administrator when attempting to identify and remediate the cause of an issue.

In this chapter, you will take a deeper look into the troubleshooting of Secure Network Access and examine some of the tooling that has been built in to ISE 1.2 to aid you in your troubleshooting efforts.

Troubleshooting of product can sometimes get fairly complex. Troubleshooting a solution made up of multiple products is bound to get downright difficult. The biggest tip we can provide for you is this: Always stay calm and think through the flows. When you are comfortable with the Secure Access solution and how the parts work together, it really is not bad at all.

"Do I Know This Already?" Quiz

The "Do I Know This Already?" quiz enables 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."

Foundation Topics Section	Questions
Identify issues using authentication event details in Cisco ISE	1, 4
Troubleshoot using Cisco ISE diagnostic tools	1-7, 9
Troubleshoot endpoint issues	2-5, 9-10
Use debug commands to troubleshoot RADIUS and 802.1X on IOS switches and wireless controllers	10

Table 22-1 "Do I Know This Already?" Section-to-Question Mapping

> **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 ISE diagnostic tool can be used to find misconfigurations in a Cisco NAD?

   a. TCP Dump

   b. Live Sessions Log

https://t.me/learningnets

c. RADIUS Authentication Troubleshooting Tool

d. Evaluate Configuration Validator

2. Which ISE diagnostic tool can be used to examine different aspects of a session and provide some additional details that might not have been available in the detailed authentication report?

a. TCP Dump

b. Live Sessions Log

c. RADIUS Authentication Troubleshooting Tool

d. Evaluate Configuration Validator

3. True or False? Logging levels in ISE can be set to debug level only from the command-line interface.

a. True

b. False

4. Which ISE tool displays a correlated view of authentications, change of authorizations, and state changes of an endpoint through its lifecycle on a network?

a. Live Log

b. Live Sessions Log

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 Log

c. RADIUS Authentication Troubleshooting Tool

d. Evaluate Configuration Validator

6. Choose the option that best describes how external syslog servers can receive logs from ISE.

a. Each PSN must be configured locally to send syslog 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 to the external recipients.

d. Each PSN sends syslog to the MNT nodes, and the external syslog receivers at the same time.

7. Where does an ISE admin 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 will gather all the important log files and combine them into a single bundle for 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

https://t.me/learningnets

c. ISE, endpoint, NAD

d. Endpoint, firewall, NAD

10. Which debug command will provide the best detail to identify why a URL redirection might not be working?

a. debug authentication

b. debug epm all

c. debug dot1x all

d. debug aaa all

Foundation Topics

Logging

There are numerous ways to examine logs with ISE. You can use the Live Log; Live Sessions Log; and individual log files that can be downloaded from the GUI, viewed from the CLI, or combined in a support bundle.

Live Log

As discussed numerous times throughout this book, ISE has a fantastic tool: The Live Authentications Log (commonly referred to as Live Log). As an administrator, you can select which columns to display and which to hide. Figure 22-1 shows Live Log and attempts to point out some key areas of the screen. Each number is listed with a description of the item to which the number refers. Continually reference the figure as you read the list.

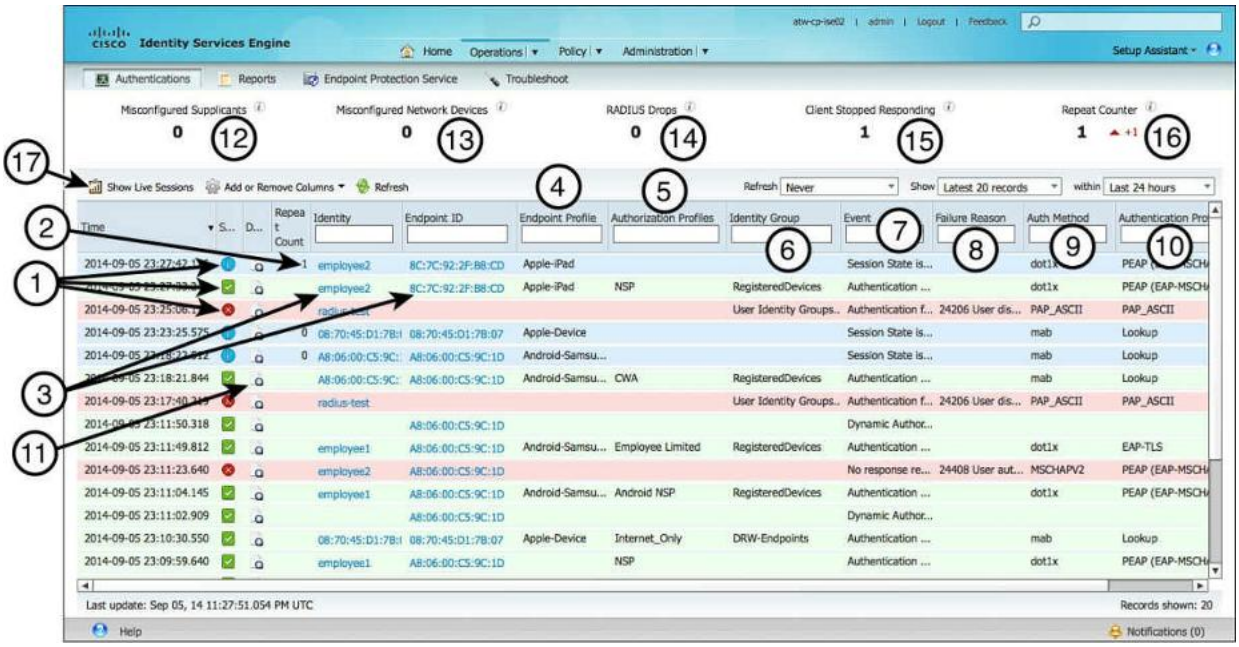


Figure 22-1 Live Log.

#1 – The status—Three status types exist in Live Log: Informational, Pass, and Fail. They are as follows:

- **Informational**—These records are blue in color and show up while client suppression is enabled (which is on by default). These increment a repeat count for each pass/fail record that is suppressed.
- **Pass**—These records are green in color and are the log entries of successful authentications.

https://t.me/learningnets

■ **Fail**—These records are red in color and are the 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. The figure is pointing out that employee2 has successfully authenticated two times because the repeat counter is set to 1.

#3 – Identity and Endpoint ID—These fields show the username from the RADIUS request, WebAuth portal, or even the MAC address in the case of MAB. The endpoint ID is the calling-station-id field from the RADIUS request. Beginning with 1.2.0 patch 5, both the identity and endpoint ID fields are right-clickable, providing an option to copy their value(s) to be used elsewhere.

#4 – Endpoint Profile—This lists the profile assigned to the endpoint involved in the authentication.

#5 – Authorization Profile—This lists the result that was sent during the authentication.

#6 – Identity Group—This lists the user or endpoint identity group, when applicable. For more on identity groups, see Chapter 3, "Identity Management," and Chapter 9, "Initial Configuration of Cisco ISE."

#7 – Event—This lists the main event ID and description for the success or failure. More information is available after clicking details (#11).

#8 – Failure Reason—When the record indicates a failed authentication, the main failure reason is listed here. More information is available after clicking details (#11).

#9 – Authentication Method—Lists the method used, such as MAB, dot1x, PAP_ASCII, MSCHAPv2, and so on.

#10 – Authentication Protocol—Displays the actual protocol that was used, which could be the tunneled EAP type, or even the lookup (for MAB endpoint lookups).

#11 – Details Icon—Clicking this icon launches the Authentication Details report, which contains a tremendous amount of information about the authentication.

#12 – Misconfigured Supplicants—This counter is used with event suppression, and each time an event is discarded due to a misconfigured endpoint, this counter increases. Clicking this counter brings up the list of dropped events.

#13 – Misconfigured Network Devices—This counter is used to point out when a network access device (NAD) is configured with a RADIUS accounting update frequency that is considered to be "too often." Clicking this counter brings up the list of dropped events.

#14 – RADIUS Drops—When the RADIUS packet is dropped before it is even processed, this counter increments. An example of a situation in which the packet will be dropped is when a NAD is configured with an incorrect shared secret. Clicking this counter brings up the list of dropped events.

#15 – Client Stopped Responding—Often an endpoint starts an authentication but does not finish. One such reason is the client doesn't trust the certificate presented by ISE. Any time that happens, this counter increments. Clicking this counter brings up the list of dropped events.

#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 the Live Log. Clicking this counter brings up the list of dropped events.

#17 – Show Live Sessions—Click this button to switch from Live Log to the Live Sessions Log, which is described in the next section.

https://t.me/learningnets

Live Sessions Log



The Live Sessions Log differs from the Live Log in that it correlates activity related to an entire session, whereas the Live Log is focused on the raw entries related to a passed or failed authentication. For example, recall from Chapter 17, "Bring Your Own Device," that a device can authenticate numerous times within a single session.

The Live Sessions Log correlates all the activities and the various states of the session. Figure 22-2 points out the single session of a device that must have gone through onboarding. It is listed with an authentication using PEAP(EAP-MSCHAPv2) and then another authentication using EAP-TLS. Also notice that the IP address is learned through an accounting message.

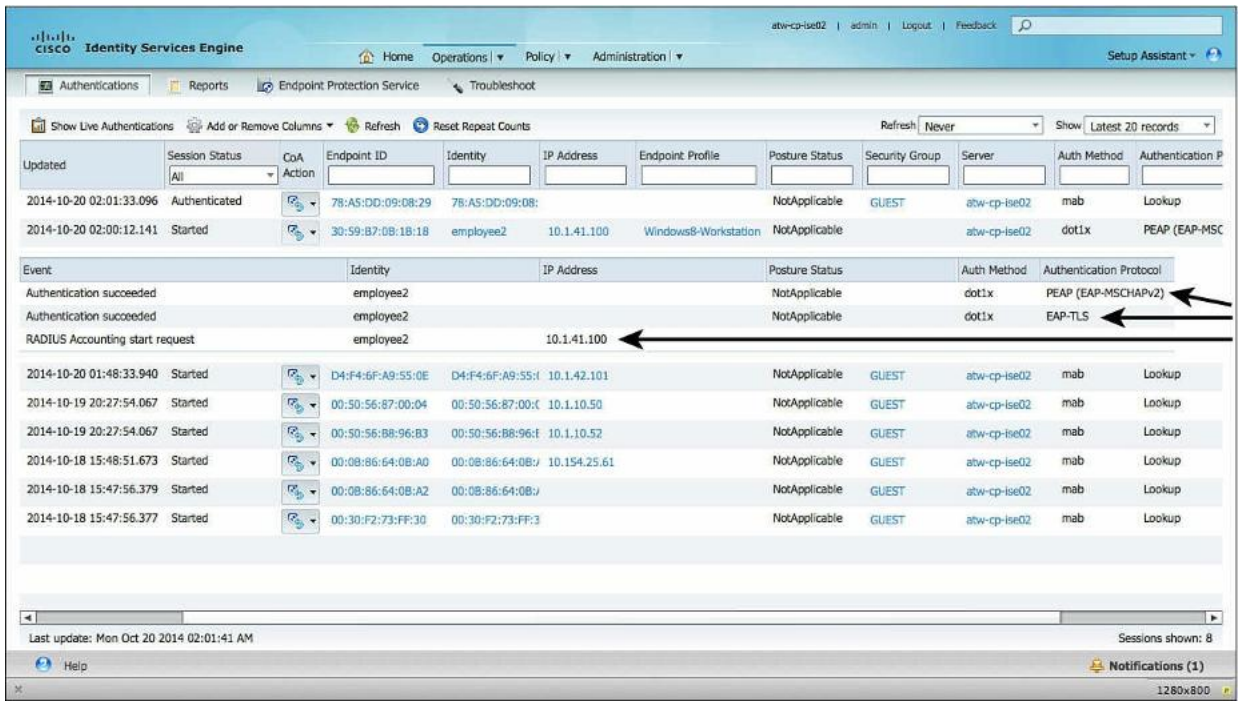


Figure 22-2 Live sessions.

The Live Sessions log has a useful function that enables you to send a change of authorization (CoA) from the session listing, as shown in Figure 22-3.

https://t.me/learningnets

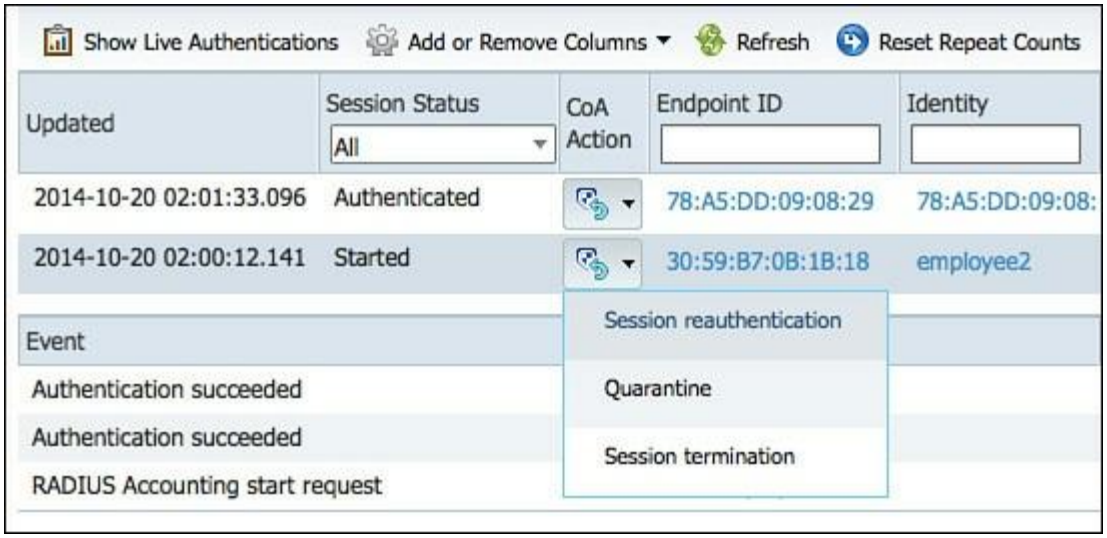


Figure 22-3 CoA from live sessions.

Logging and Remote Logging

All nodes in an ISE deployment send information to the Monitoring and Troubleshooting node(s) using syslog. The information sent includes system health; authentication, authorization, and accounting records; alarms; and any other system message. Traditional syslog transport uses UDP port 514, whereas 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 is the Primary Monitoring and Troubleshooting (MnT) persona.

To examine the configured LogCollector, navigate to **Administration > System > Logging > Remote Logging Targets**. Here a friendly name is given to a syslog collector and its definition—in other words, the IP Address, TCP or UDP syslog, and port number for the log collector, as shown in Figure 22-4. TCP syslog is used in certain environments to provide guaranteed log delivery.

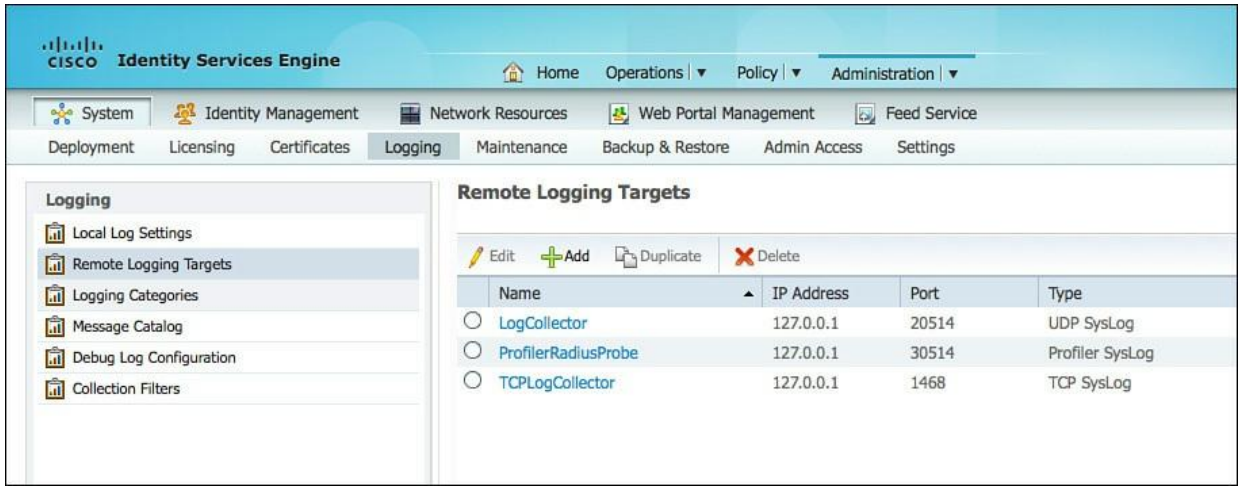


Figure 22-4 Default log collector configuration.

A Security Information and Event Management (SIEM) or other syslog receiver can be added to ISE simply

https://t.me/learningnets

by creating a new logging target, as shown in Figure 22-5.

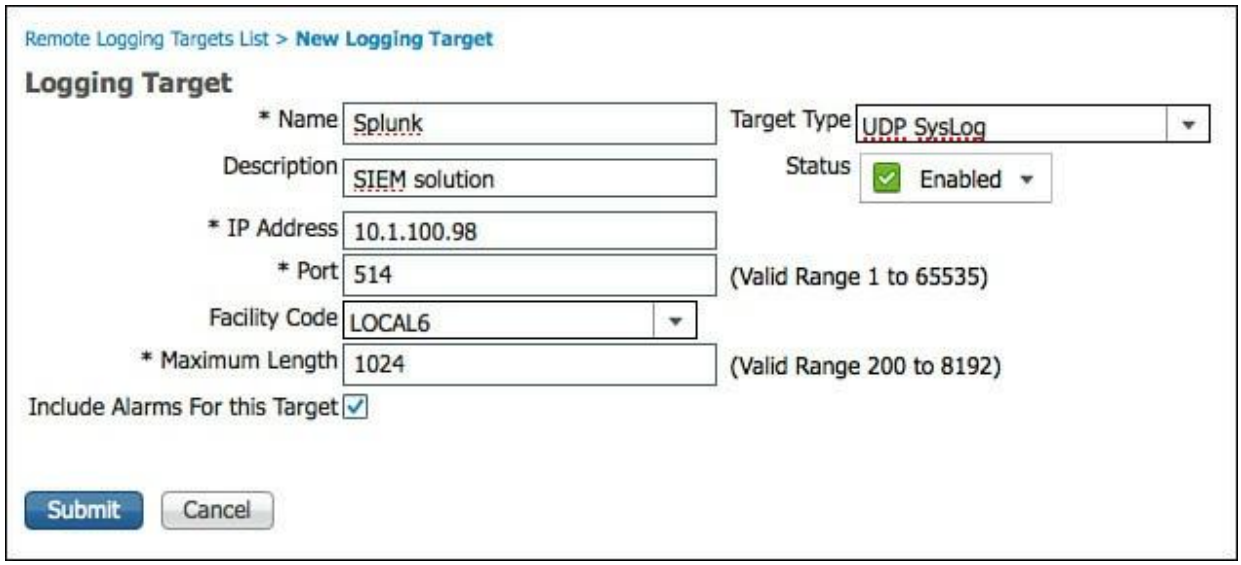


Figure 22-5 Adding a remote logging target.

To send logs to the newly added logging target, the target must be added to the various logging categories.

Logging Categories

There are many types of logs with a solution as in-depth as ISE. The logging categories and the logging targets assigned to each of the categories can be seen and configured by navigating to **Administration > System > Logging > Logging Categories**, as shown in Figure 22-6.

https://t.me/learningnets

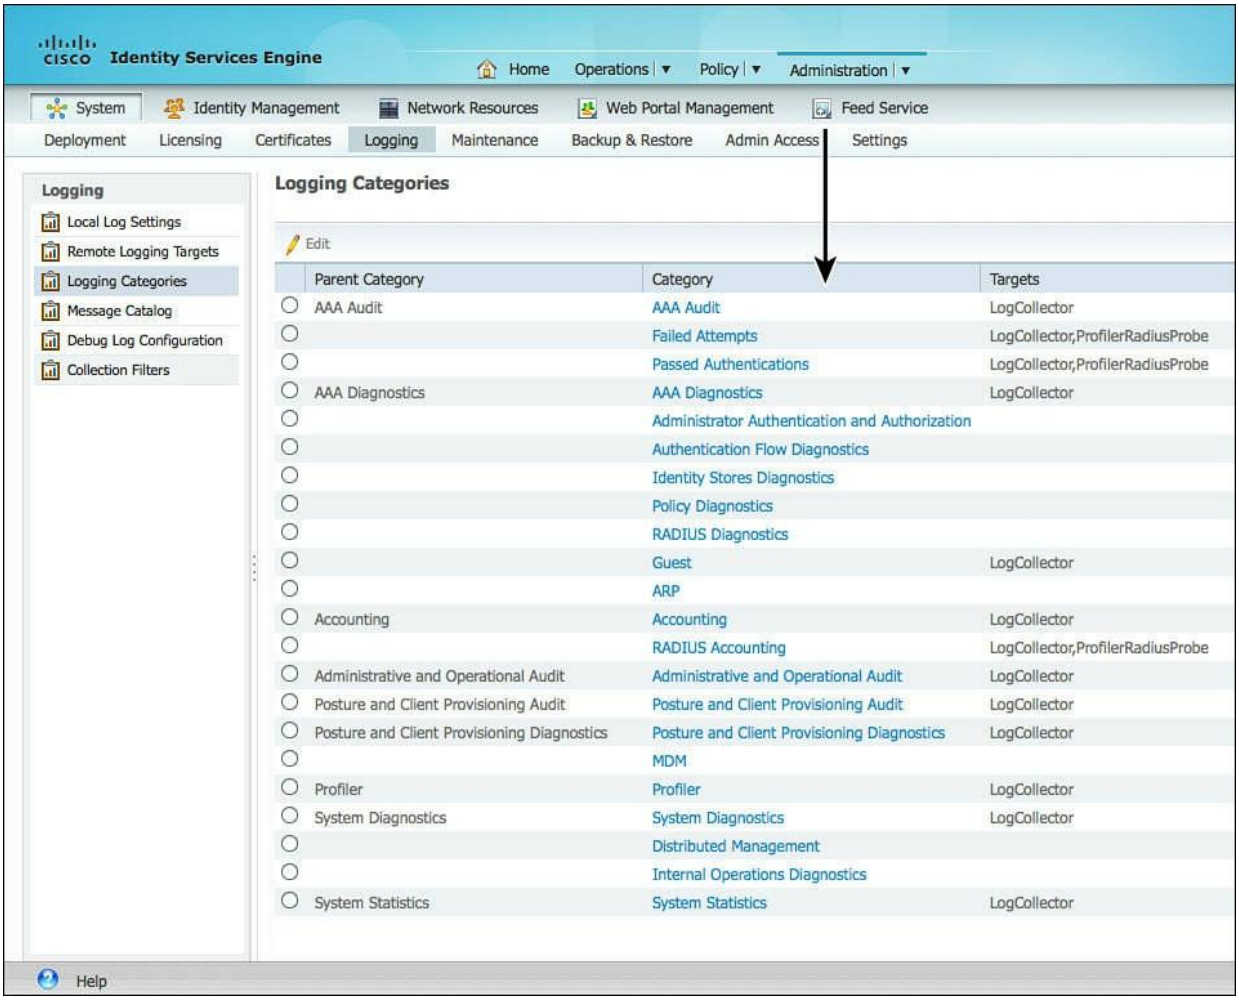


Figure 22-6 Logging categories.

For every category you want an external SIEM or other collector to receive, you must add that log collector to the individual category. The good news is that it takes effect for the entire ISE cube. To add that collector to a category, simply edit the category or click the category link. From there, ensure the desired collectors are in the Selected column, as shown in Figure 22-7.

https://t.me/learningnets

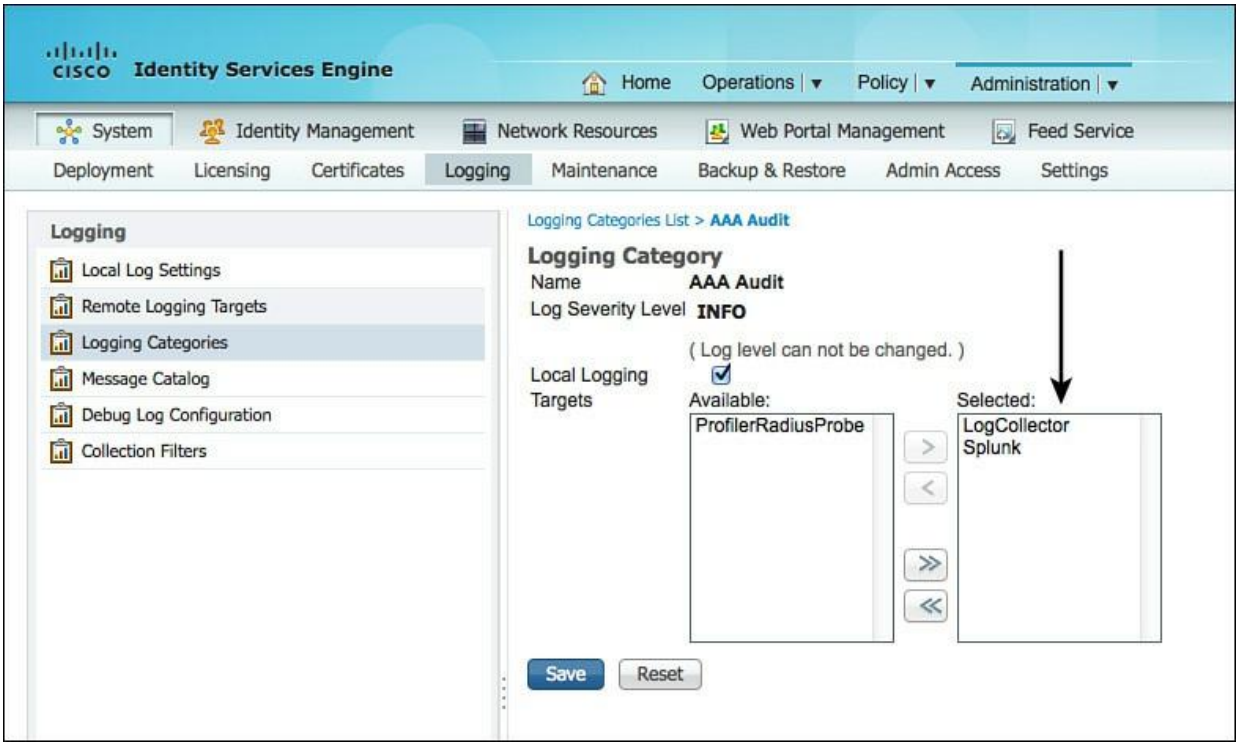


Figure 22-7 Adding a target to a logging category.

Debug Logs

As with most enterprise class solutions, a provision is often made to provide very detailed levels of logging for troubleshooting purposes known as debugging. ISE also provides the ability to set each component to log at different levels, including the debug level.

> **Note**
>
> It is never recommended to change any of these settings unless directed by Cisco TAC. Setting components to debug level will have an impact on system performance.

To change the logging level of a component, navigate to **Administration > System > Logging > Debug Log Configuration**. Click the name of the ISE node that needs to have the log settings modified; the list of components and the current log level settings are displayed, as shown in Figure 22-8.

https://t.me/learningnets

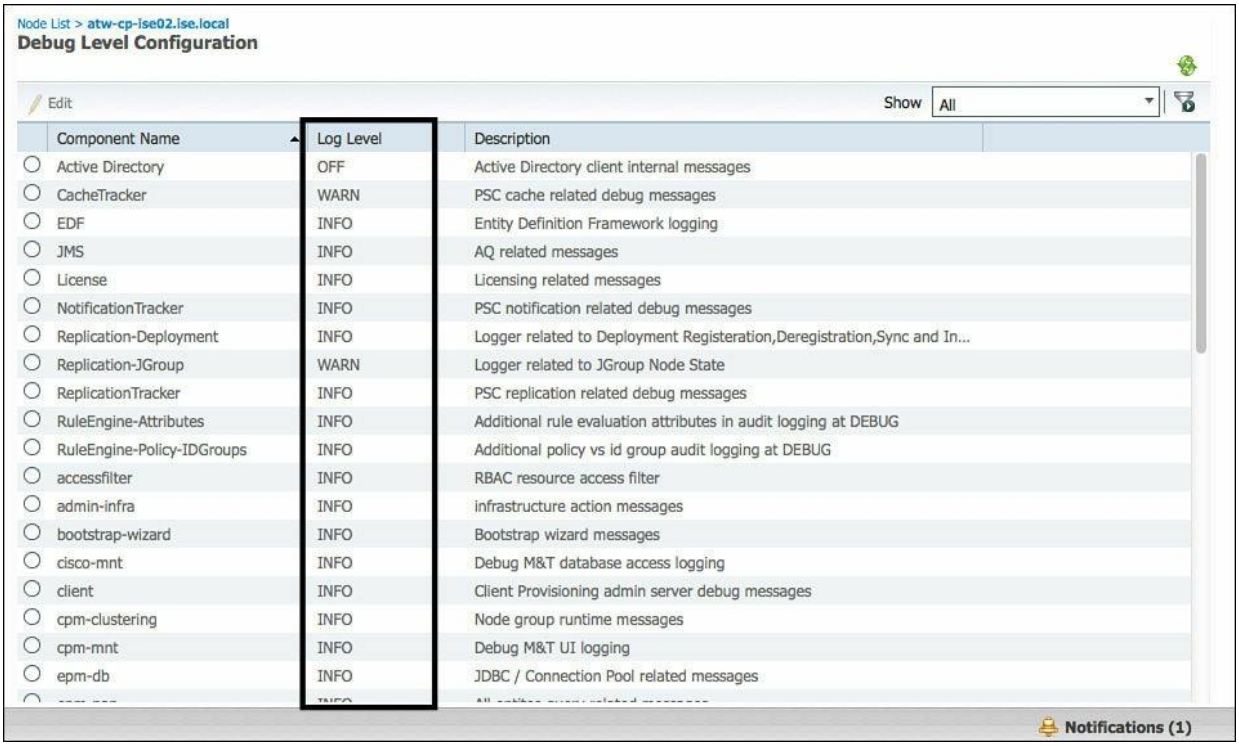


Figure 22-8 Default debug log configuration.

To modify a component's logging level, edit the component or double-click the logging level. Select the new logging level from the drop-down list, and click **Save**, as shown in Figure 22-9.

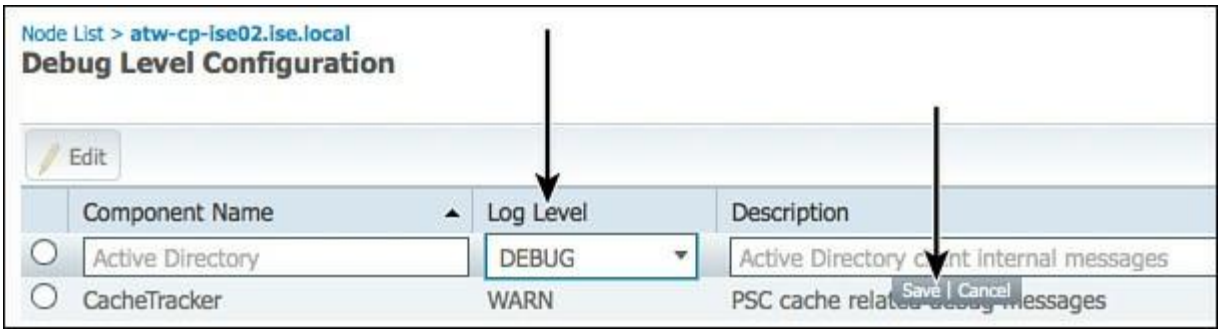


Figure 22-9 Changing the logging level of a category.

The resulting debug log file can be downloaded from the GUI, viewed from the CLI, or be included in a support bundle.

Downloading Debug Logs from the GUI

Debug log files can be downloaded from the GUI. Navigate to **Operations > Troubleshoot > Download Logs**. Select the ISE node that contains the log files that need to be downloaded, and click the link for the specific log file, as shown in Figure 22-10.

https://t.me/learningnets

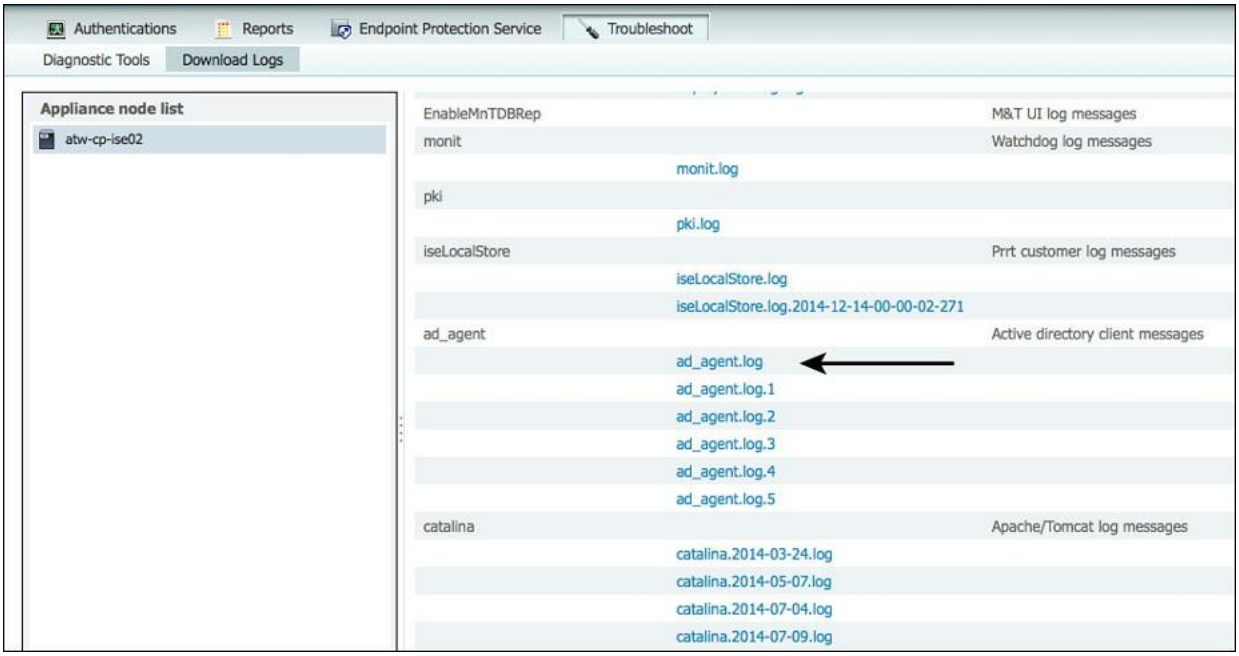


Figure 22-10 Downloading a debug log file from the GUI.

After the file is downloaded, it can be viewed in any standard text editor application.

Viewing Log Files from the CLI

Sometimes you might choose to view the logs on the local system through the command-line interface (CLI). ISE provides the ability to view and tail the log files directly from the CLI using the **show logging application** and **show logging system** commands. The application command is used for the ISE application's logs, and the system command is used for logs related to the underlying operating system (ADE-OS).

To view a list of all the ISE log files, simply type **show logging application** and press Enter, as shown in Example 22-1.

Example 22-1 show logging application Command

Click here to view code image

```
atw-cp-ise02/admin# show logging application
    5244288 Nov 03 2014 02:15:01  ad_agent.log.4
    5243450 Dec 05 2014 06:30:01  ad_agent.log.1
      22742 Dec 16 2014 16:33:31  mnt-report.log
```

To view the contents of a specific log file, simply type **show logging application** and the name of the log file, such as **ad_agent.log**, as shown in Example 22-2.

Example 22-2 show logging application ad_agent.log

Click here to view code image

```
atw-cp-ise02/admin# show logging application ad_agent.log
2014-12-05T06:37:51.031851+00:00 atw-cp-ise02 adclient[7795]: WARN  <bg:machinep
w> daemon.main computer password change failed: Last error was unexpected discon
nect (GC); deferring reconnect.
2014-12-05T06:38:42.005040+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Start trusted domain discovery
```

https://t.me/learningnets

```
2014-12-05T06:38:42.005074+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Delay /etc/krb5.conf update, Skipping trusted domains update beca
use the system is in disconnected mode
2014-12-05T06:48:42.004611+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Start trusted domain discovery
2014-12-05T06:48:42.004660+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Delay /etc/krb5.conf update, Skipping trusted domains update beca
use the system is in disconnected mode
2014-12-05T06:58:42.004231+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Start trusted domain discovery
```

In addition to the ability to view a log file, you can also perform certain standard UNIX operations on the file, such as tail and grep. This is exposed through the CLI by adding a pipe (|) after the log name and choosing the output modifier. The options are shown in Example 22-3.

Example 22-3 show logging application ad_agent.log | ?

Click here to view code image

```
atw-cp-ise02/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

Cisco routers, switches, and other IOS-based devices have a very useful command: **show tech-support**. This command builds the output of nearly every usable **show** command on the IOS device, so the administrator has to submit only a single file to the TAC case that provided the TAC engineer with nearly every answer to any possible question that could be asked about the device configuration and current state.



Cisco ISE has something similar, called a support bundle. The support bundle is an encrypted archive that contains the output to many CLI show commands, a backup of the databases and ISE configuration, copies of all the log files, and more. These bundles can be quite large in nature but typically provide Cisco TAC with everything needed to re-create an issue for the purposes of troubleshooting and resolving your case.

To create a support bundle, navigate to **Operations > Troubleshoot > Download Logs**. Select the ISE node to create the support bundle from. On the Support Bundle tab, select which items to include in the bundle and define a password for the encryption and decryption of the data, as shown in Figure 22-11.

https://t.me/learningnets

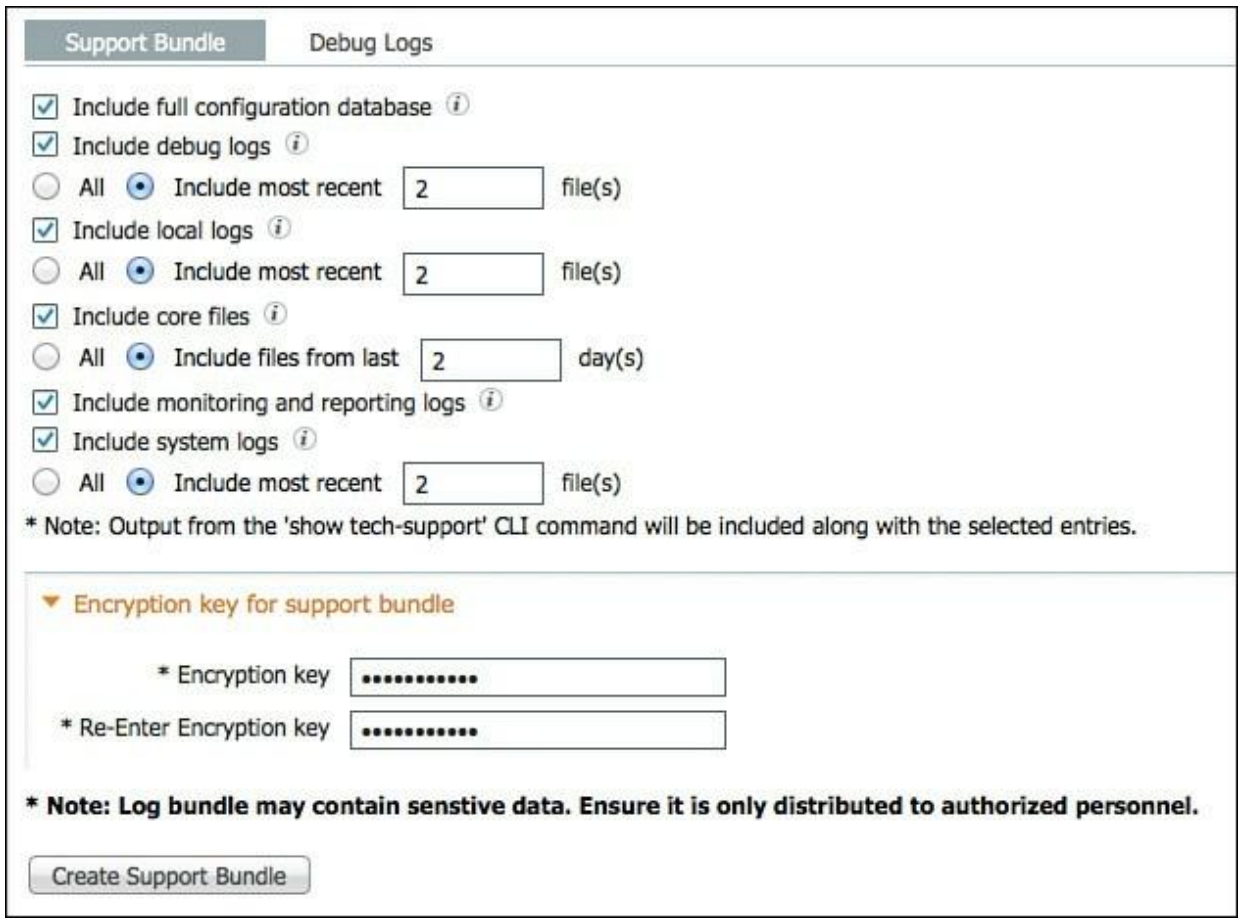


Figure 22-11 Creating a support bundle.

Diagnostics Tools

In addition to all the terrific logging capabilities in ISE, a number of built-in diagnostics tools are found under **Operations > Troubleshoot > Diagnostics Tools**. In this section, you will get a brief overview of the most commonly used tools.

Evaluate Configuration Validator

A terrific tool with a less-than-terrific name, the Evaluate Configuration Validator tool connects to a switch via telnet or SSH, or even through a console terminal server and examines the configuration. The configuration is then compared to a template configuration built in to ISE, and any differences between the configurations are called out. At the time this book was written, the tool was due for an update, but it still can provide a lot of value. Some administrators of ISE have been known to point this tool to an unconfigured switch and copy/paste the resulting commands to configure the switch.

Until it is updated, the following list displays a few of the common configurations that might be misdiagnosed as missing or incorrect:

- Does not currently understand device sensor and will be expecting the use of SNMP for the SNMP probe(s).
- Does not recognize the active test options when defining a RADIUS server, and therefore might think the RADIUS server definition is incorrect.
- WebAuth is Local WebAuth only, and the tool does not recognize that MAB is used for Central WebAuth instead.
- Does not support the Cisco Common Classification Policy Language (C3PL) commands available in IOS 15.2 and above (C3PL is out of scope for this exam and therefore this book).

https://t.me/learningnets

Even with the limitations listed, this tool is still recognized by Cisco TAC as quickly identifying a high number of misconfigurations, and in many cases it would have prevented a customer from opening the TAC case.

To run the Evaluate Configuration Validator tool, do the following:

Step 1. Navigate to **Operations > Troubleshoot > Diagnostics 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 22-12.

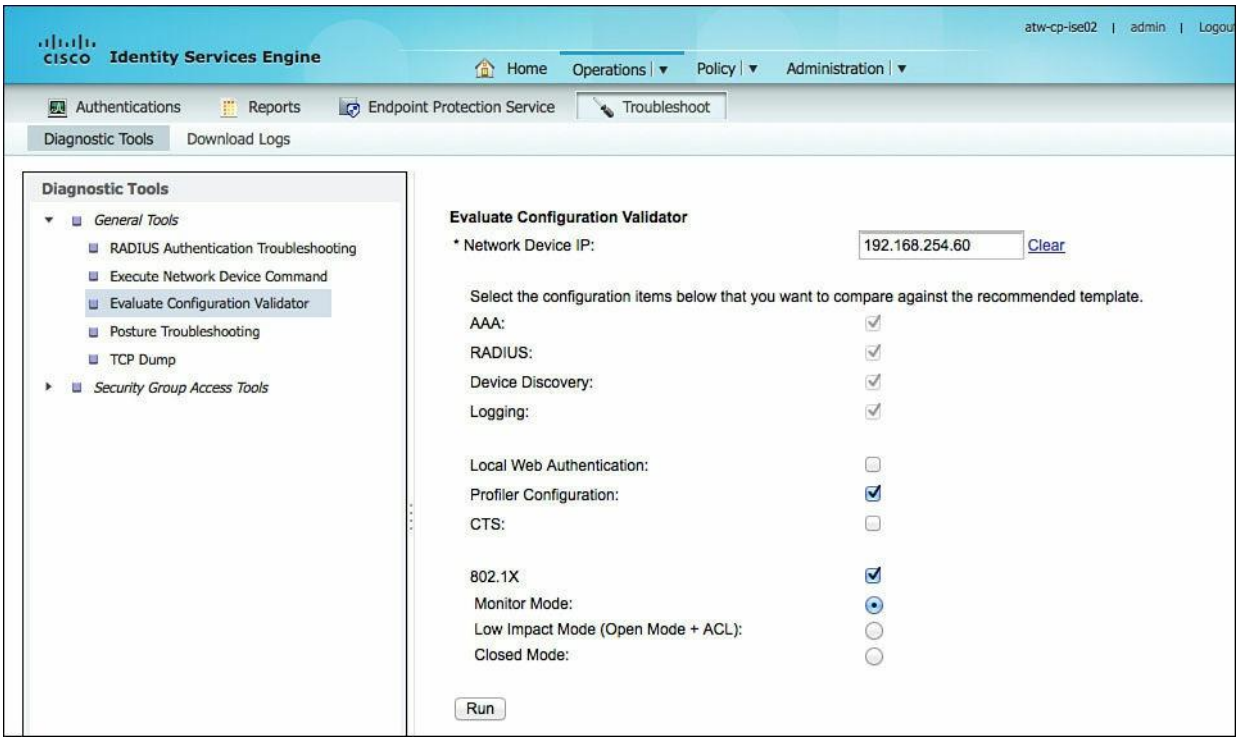


Figure 22-12 Initializing the Evaluate Configuration Validator tool.

> **Note**
>
> In Figure 22-12, Local Web Authentication has been unchecked because we are using Centralized Web Authentication (CWA).

Step 3. Click **Run** to begin the evaluation.

Step 4. ISE connects to the switch and prompts for your interaction if the switch asks for user authentication or other interactive prompts, as shown in Figures 22-13 and 22-14.

https://t.me/learningnets

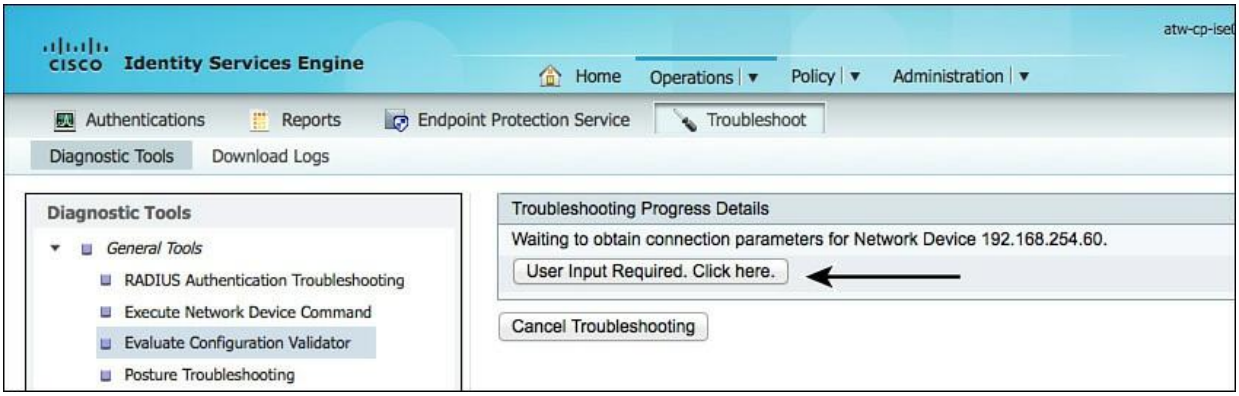


Figure 22-13 User input required.

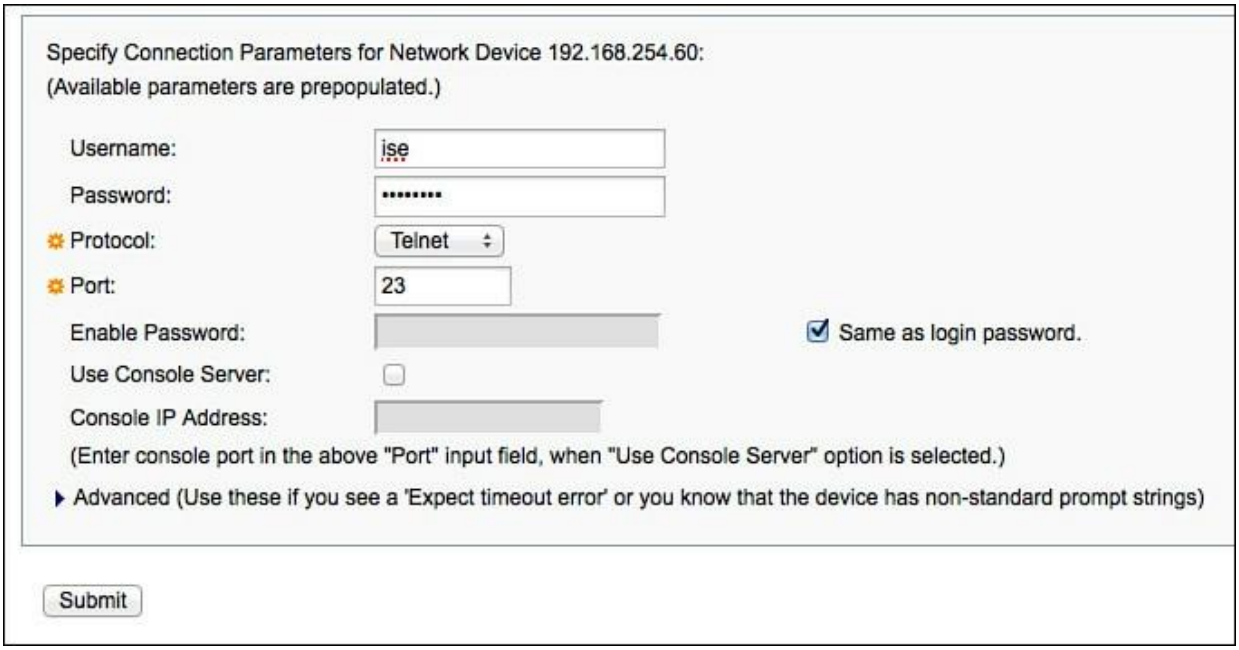


Figure 22-14 Entering credentials.

Step 5. Next, you are prompted to select which interfaces you want to compare (see Figure 22-15). If you are following the guidelines we set in this book, then nearly every interface will have the exact same configuration, and you should select only one interface to compare.

https://t.me/learningnets

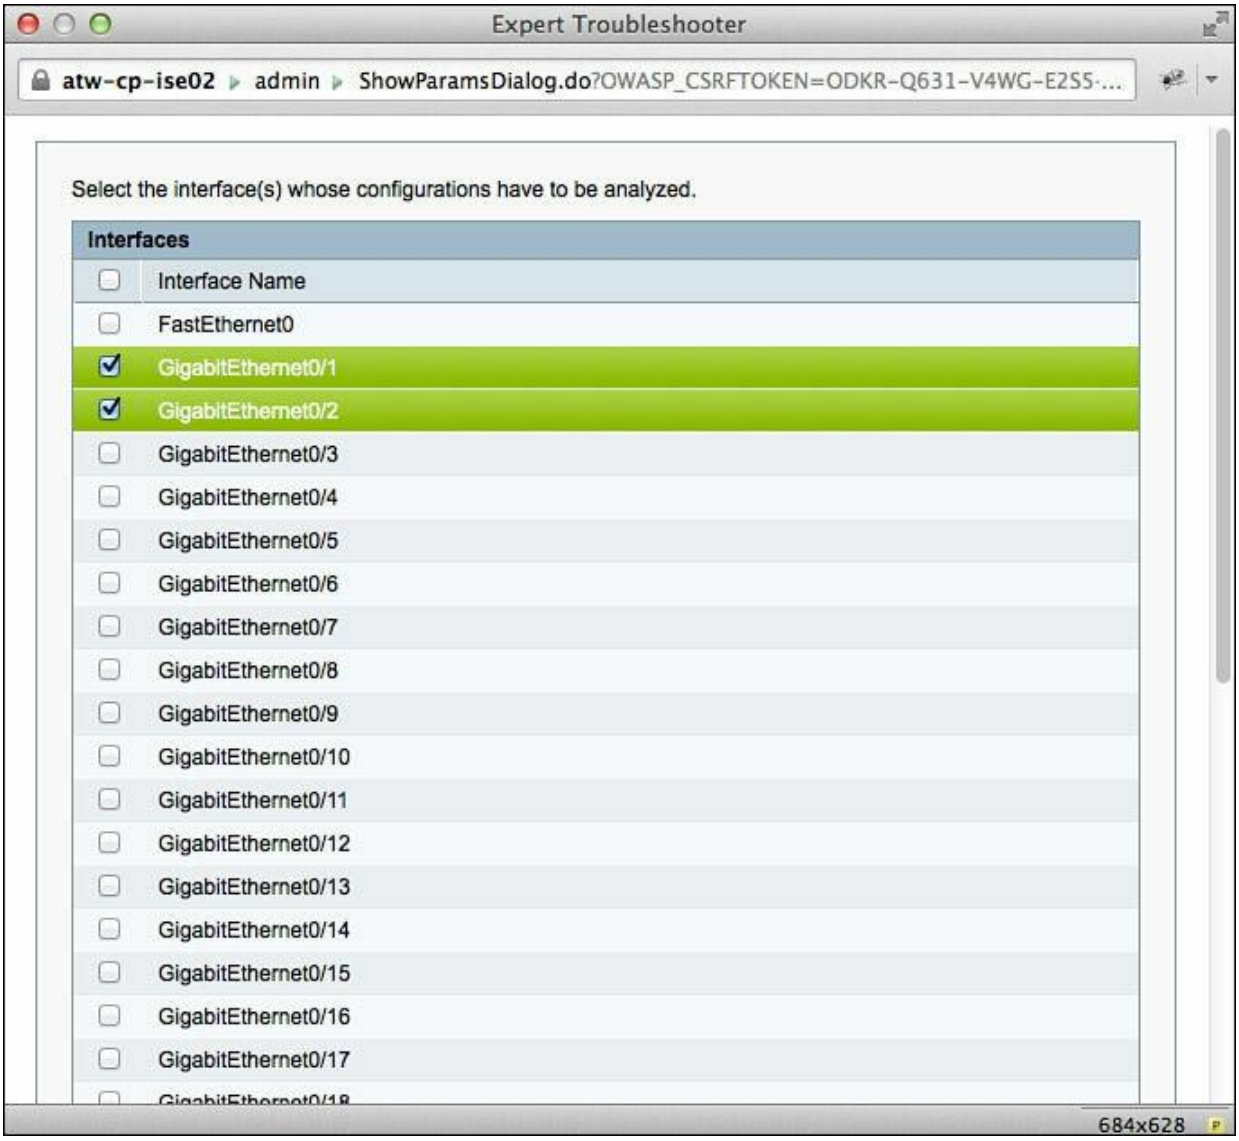


Figure 22-15 Select interface(s).

Step 6. The comparison should be complete, as shown in Figure 22-16. Click **Show Results Summary** to see the results, which are displayed in Figure 22-17.

https://t.me/learningnets

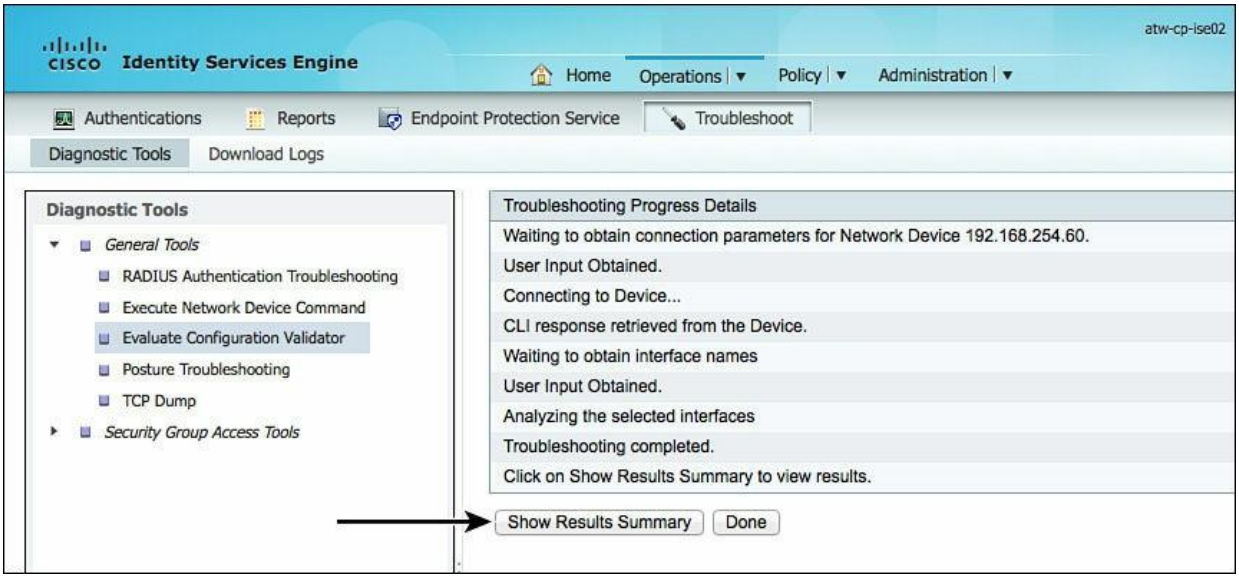


Figure 22-16 Comparison complete.

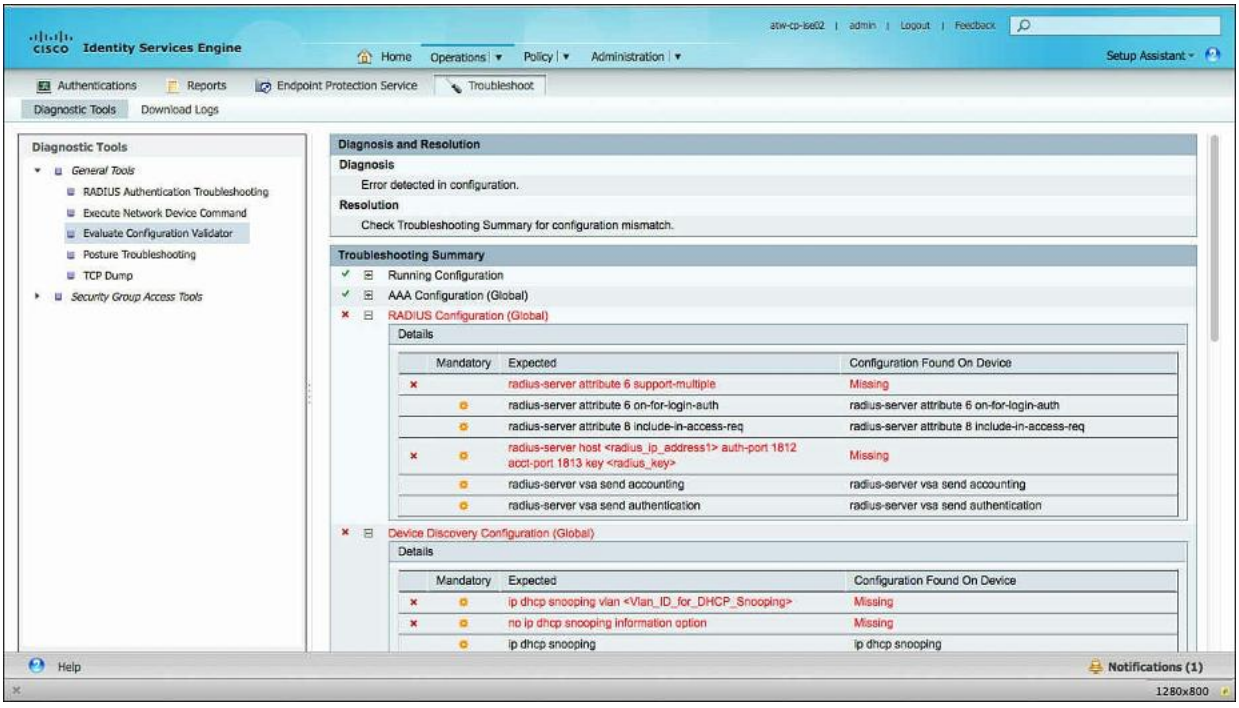


Figure 22-17 Comparison results.

As previously mentioned, the tool does not understand some of the newer commands that supersede the commands in the template it uses for the comparison. A perfect example is in Figure 22-17, where the comparison result states the **radius-server host** command is missing. The command is in fact on the switch, but it includes additional arguments for the proactive testing of the RADIUS server, as shown in Example 22-4. So it is up to you, the ISE administrator or consultant, to use the tool for the value that it does provide—not take it a face value—and examine the actual switch configuration to ensure the necessary commands are in place.

Example 22-4 radius-server host Command Is Indeed in the Configuration

Click here to view code image

```
3560-X#show run | include radius-server host
```

https://t.me/learningnets

```
radius-server host 10.1.100.232 auth-port 1812 acct-port 1813 test username
radius-test key Cisco123
3560-X#
```

RADIUS Authentication Troubleshooting Tool

The RADIUS Authentication Troubleshooting tool attempts to examine various aspects of a session and provide additional details that might not have been available in the detailed authentication report, as well as provide some suggestions for items to check next. Do the following:

Step 1. Navigate to **Operations > Troubleshoot > Diagnostic Tools > General Tools > RADIUS Authentication Troubleshooting**.

Step 2. From here you can select any number of specifics to limit your search, such as a specific username, failed or passed authentication status, and more. Click **Search**.

Step 3. Select one of the resulting authentications and click **Troubleshoot**, as shown in Figure 22-18.

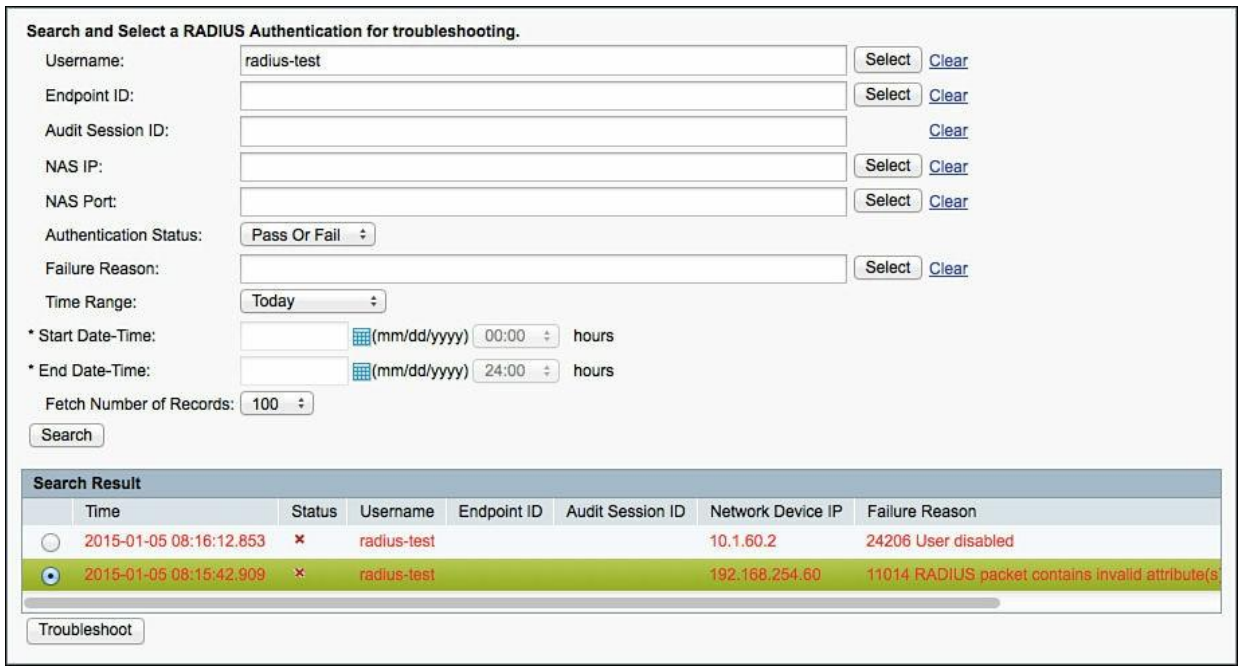


Figure 22-18 Troubleshooting the selected authentication.

Step 4. The troubleshooting progress is displayed, as shown in Figure 22-19.

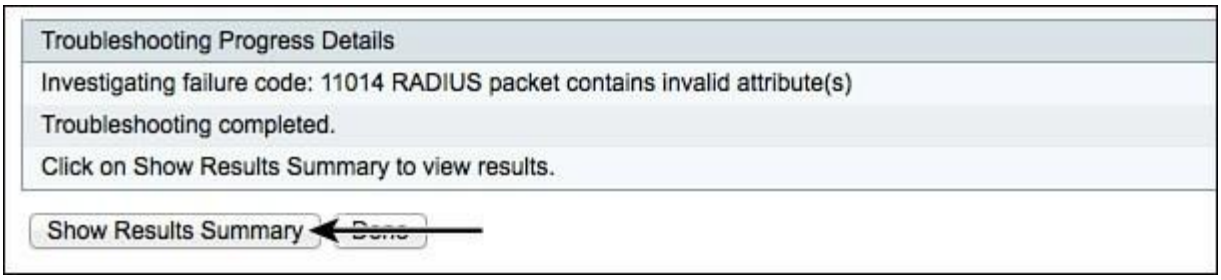


Figure 22-19 Troubleshooting progress.

Step 5. When the troubleshooting is complete, click **Show Results Summary** to see the diagnosis and suggested resolution, as displayed in Figure 22-20.

https://t.me/learningnets

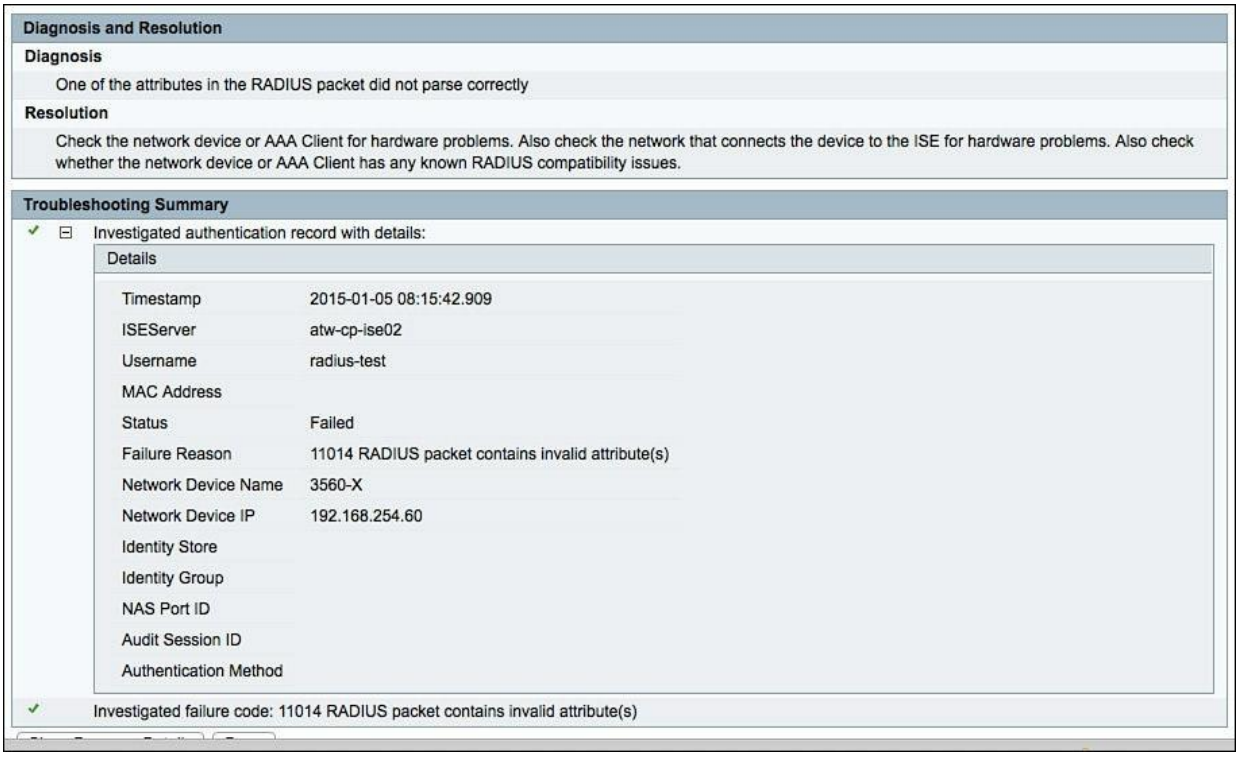


Figure 22-20 Diagnosis and resolution.

TCP Dump



One of the most valuable troubleshooting tools available to an ISE administrator is the ability to perform a packet analysis of the RADIUS packet. The packet analysis provides visibility into the raw RADIUS attributes of an authentication, so you can determine why a specific authentication rule was chosen or not chosen.

Cisco ISE provides the capability 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 admin console. The downloaded file can then be viewed using your favorite packet analysis tool, such as Wireshark.

The TCPDump tool is located under **Operations > Troubleshoot > Diagnostics Tools > General Tools > TCP Dump**, as shown in Figure 22-21.

https://t.me/learningnets

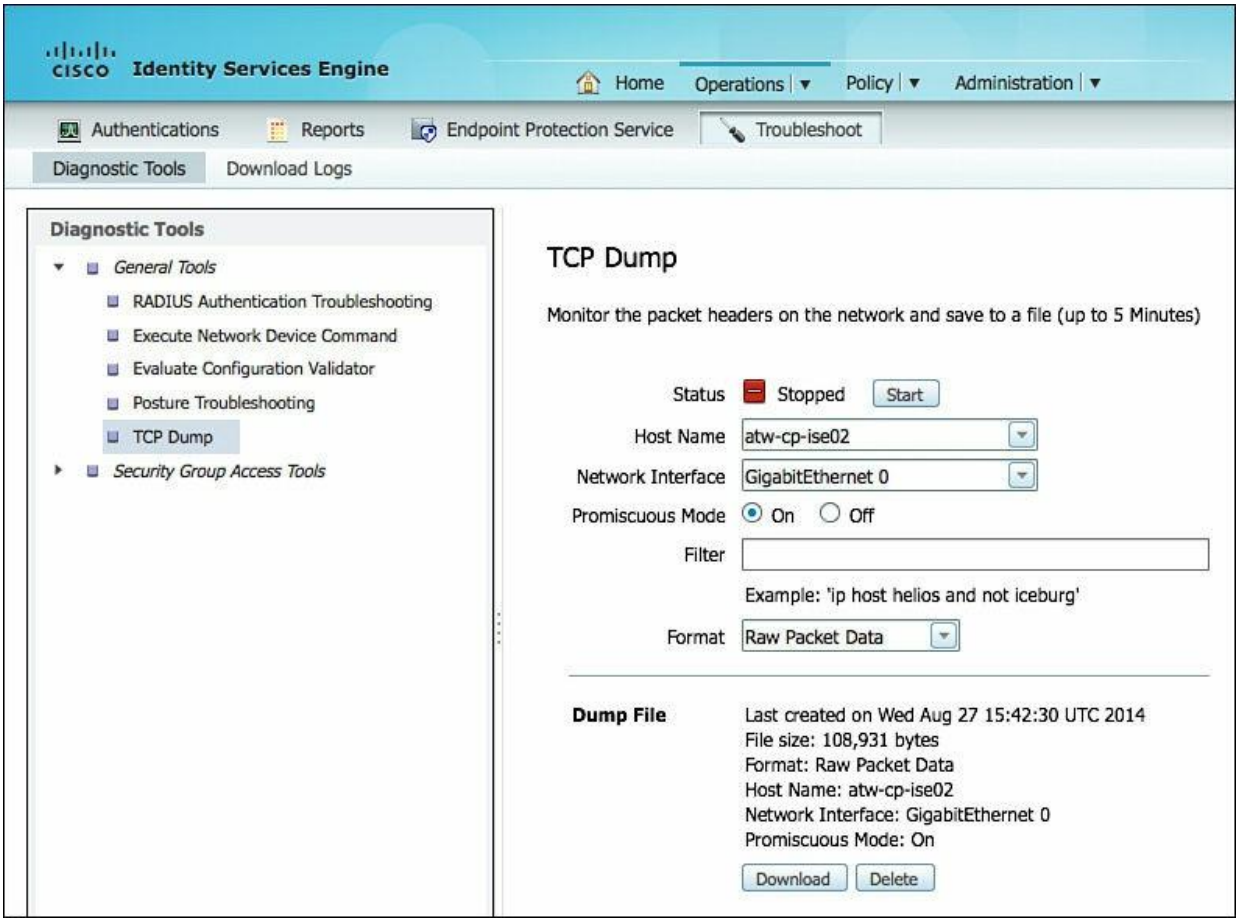


Figure 22-21 TCP Dump.

With the tool, you can select any node of the ISE deployment and any interface of that node. You are also able to select the resulting packet capture format to be either the raw PCAP format or in a human readable format.

To illustrate the value of this tool, this chapter walks through an example of checking an incoming wireless authentication. In Chapter 12, "Implement Wired and Wireless Authentication," you created an authorization rule that examined the wireless SSID an endpoint connected to, which used the RADIUS attribute Called-Station-ID. This example uses TCPDump to examine the incoming RADIUS packet and look for the Called-Station-ID attribute:

Step 1. Navigate to **Operations > Troubleshoot > Diagnostics Tools > General Tools > TCP Dump**, as shown previously in Figure 22-21.

Step 2. The correct ISE node (a PSN) and the proper interface are selected (should be GigabitEthernet 0 for most installs).

Step 3. Click **Start** to begin the TCP Dump. The packet capture begins and the approximate file size updates, as shown in Figure 22-22.

https://t.me/learningnets

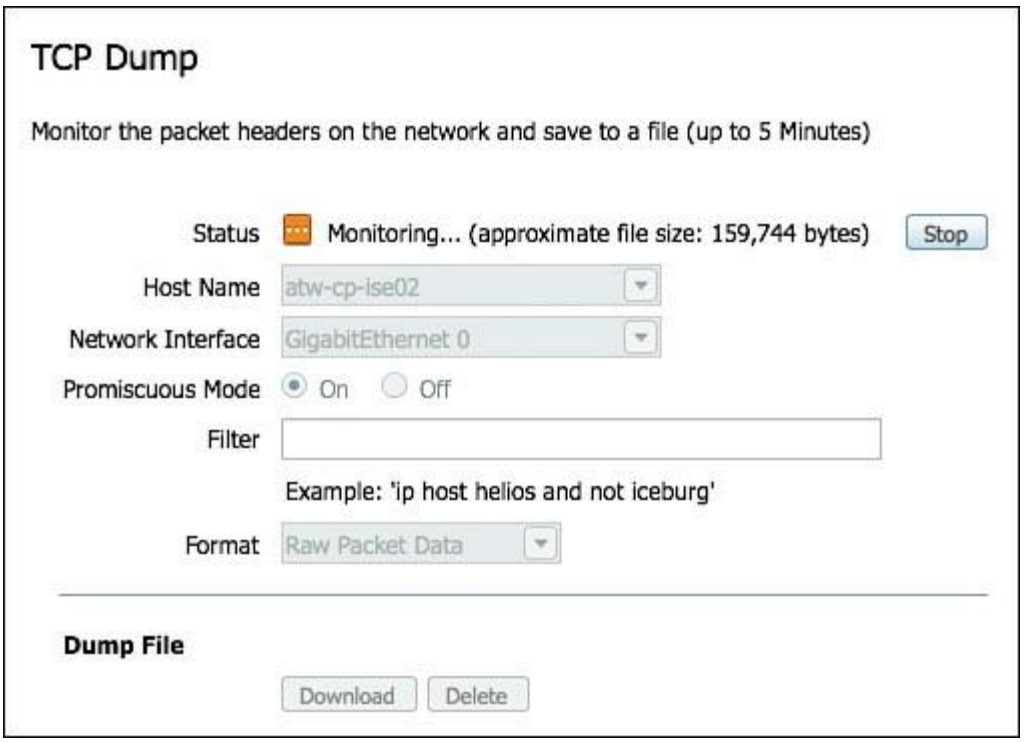


Figure 22-22 TCP Dump in progress.

Step 4. From an endpoint, connect to the wireless network (the CiscoPress SSID is being used in this example).

Step 5. From ISE, click **Stop** when you are finished. The dump file information will be updated on the screen, and the Download and Delete buttons will be enabled, as shown in Figure 22-23.

https://t.me/learningnets

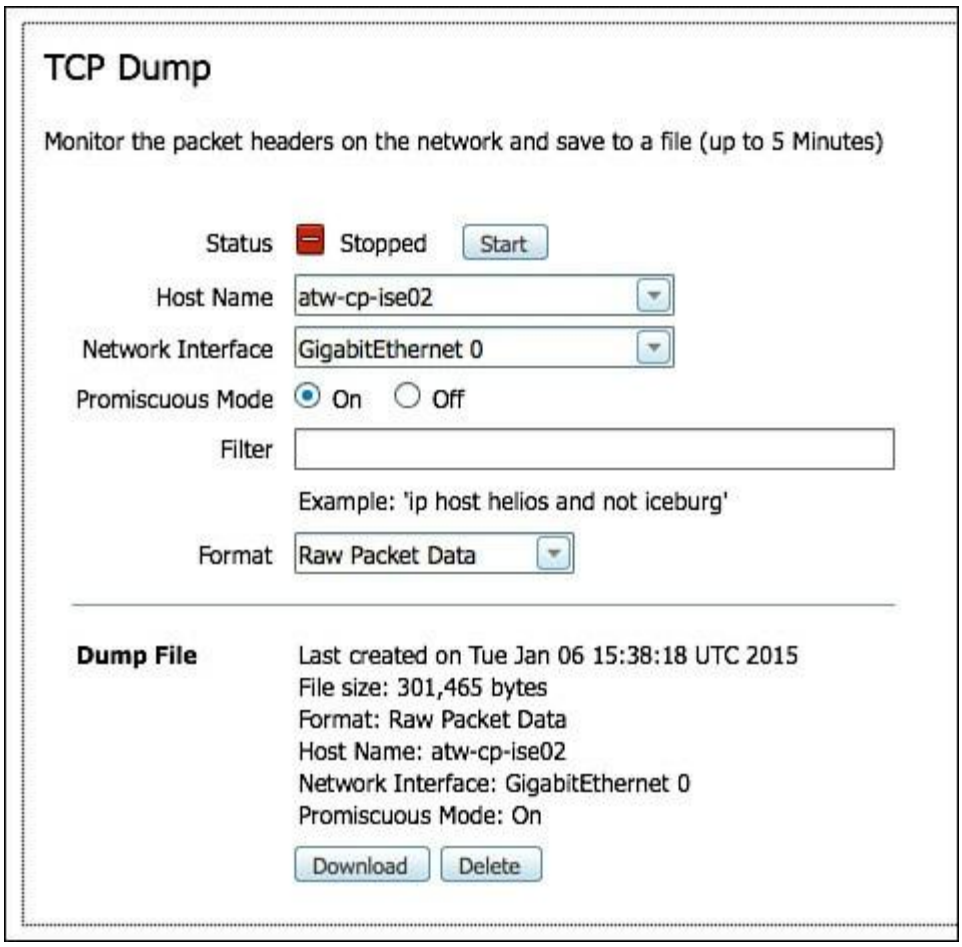


Figure 22-23 TCP Dump finished.

Step 6. Click **Download** to retrieve the PCAP file.

Step 7. Open the resulting PCAP file with your favorite packet analysis tool. Wireshark is used in this example, and it helps to apply a filter, such as "radius" to only see RADIUS packets, as shown in Figure 22-24.

https://t.me/learningnets

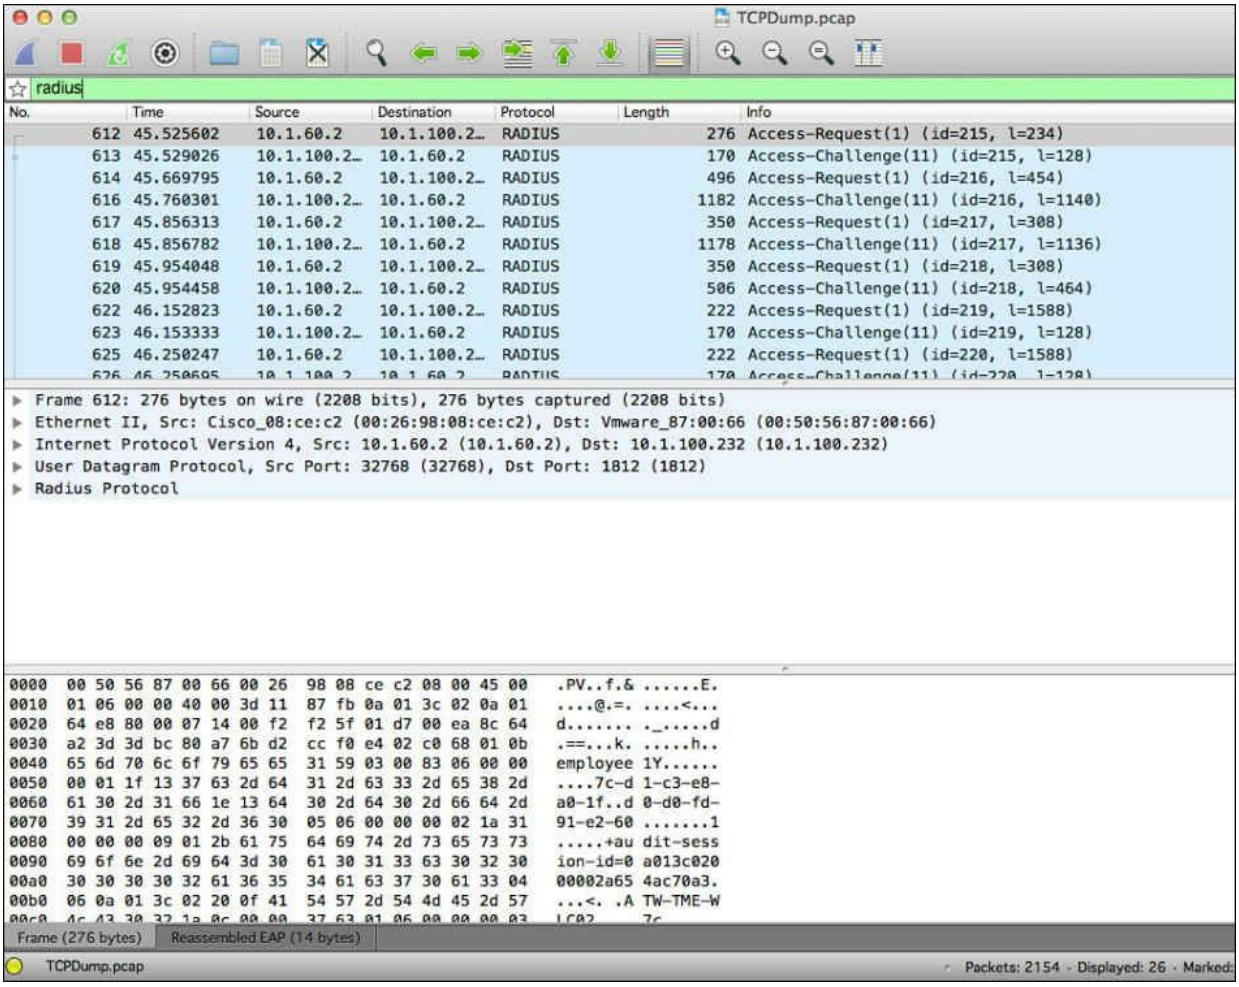


Figure 22-24 Opening the PCAP file in Wireshark.

Examining the RADIUS details in the packet, as shown in Figure 22-25, we see that the called-station-id did not contain the SSID; instead it contains the MAC address of the wireless access point. This is a configurable setting in the Wireless LAN Controller (WLC), and we will now change that setting beginning with Step 8.

https://t.me/learningnets

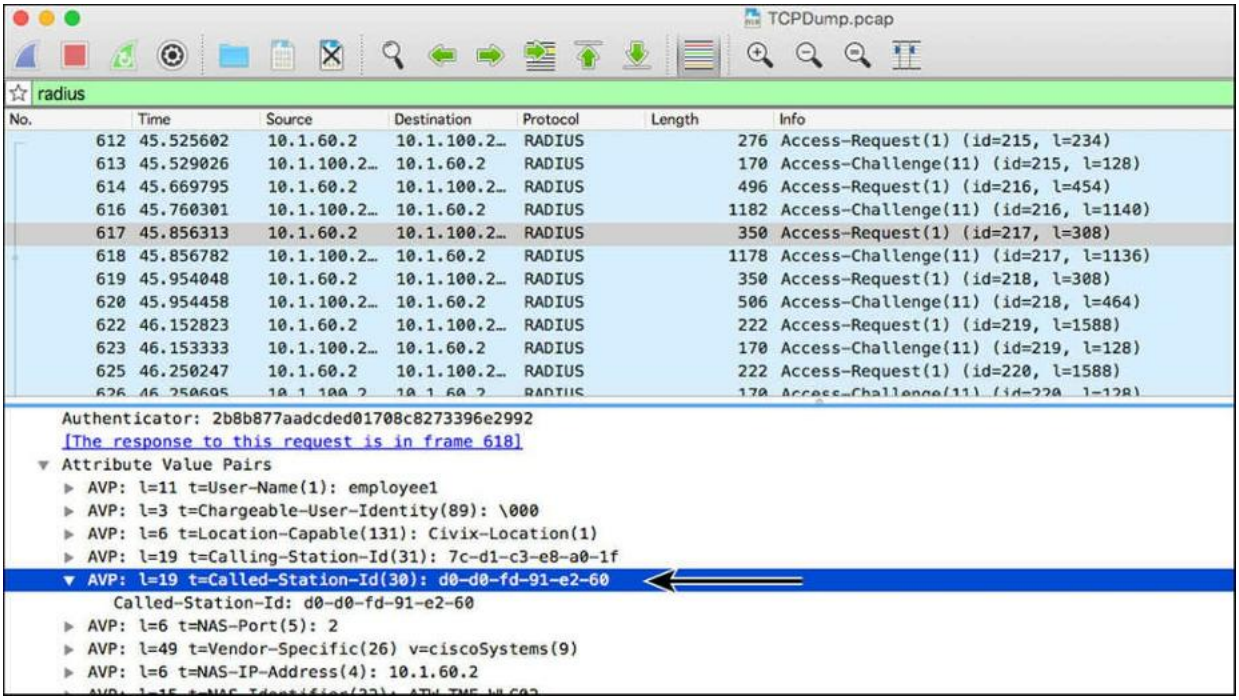


Figure 22-25 Called-Station-ID is AP's MAC address.

Step 8. Log in to the WLC Administrative GUI.

Step 9. Navigate to **Security > AAA > RADIUS > Authentication**.

Step 10. Set the **Auth Call Station ID Type** drop-down to **AP MAC Address:SSID**, as shown in Figure 22-26. This configures the RADIUS authentication requests to be sent with the SSID appended to the access point's MAC address.



Figure 22-26 Auth call station ID type.

Step 11. Click **Apply** and **Save Configuration**.

Step 12. Repeat Steps 1–6 to generate a new packet capture.

Step 13. Open the PCAP file in a packet analyzer, and the called-station-id (RADIUS attribute 30) will be the MAC address of the AP, followed by a colon (:) and then the SSID Name (CiscoPress), as shown in Figure 22-27.

https://t.me/learningnets

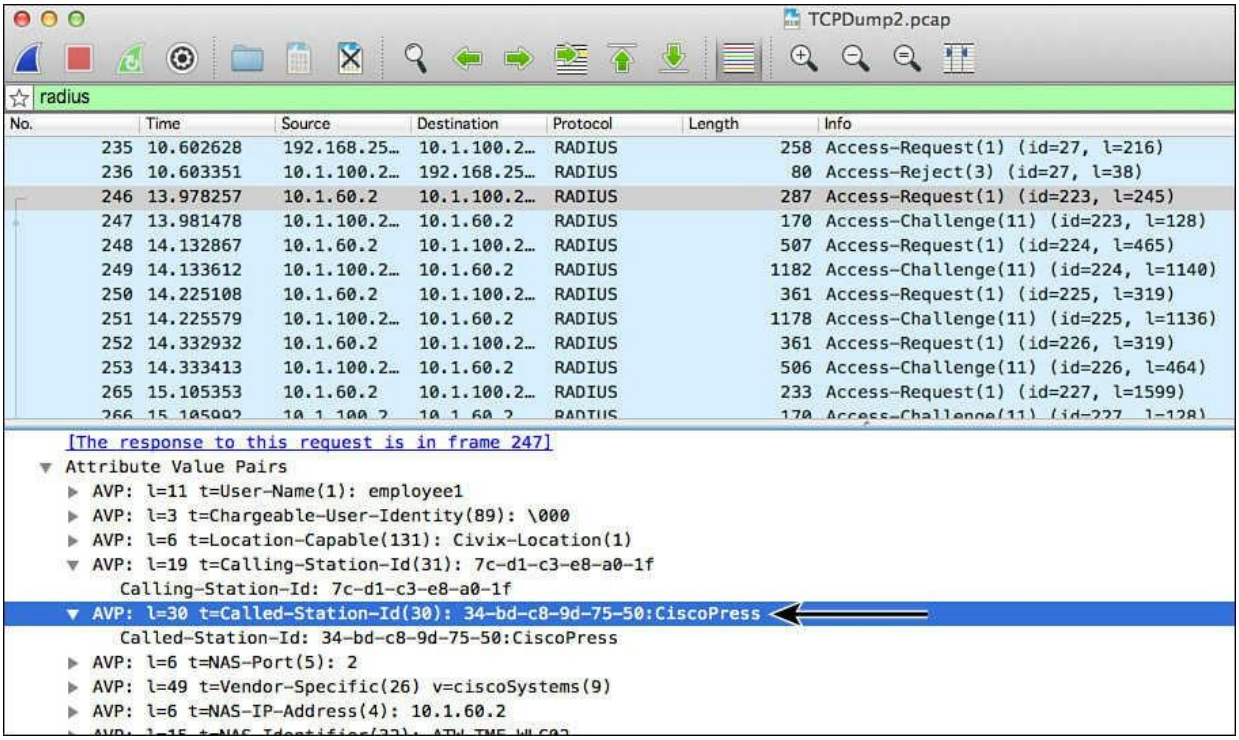


Figure 22-27 Called-Station-ID now contains the wireless SSID.

In the previous example, you have just seen what occurs when you configure a rule on ISE and then troubleshoot when that rule is not being matched by looking at the raw RADIUS packets. The root cause for the rule not being matched is identified as a misconfigured setting on the WLC, and the change was made.

Another common use for the TCP Dump tool is to verify that traffic from a NAD is actually reaching the ISE node and not relying solely on the Live Log.

Ensuring Live Log Displays All Events (Bypassing Suppression)

As networks continued to change and the bring your own device (BYOD) evolution continued to take hold, the number and types of endpoints grew and authentication traffic grew exponentially. Along with that shift in authentication traffic was a shift in the number of anomalous authentications, failures, and more.

To allow ISE to scale to larger environments, ISE 1.2 added new functionality to de-duplicate logs sent from the PSNs to the MNT. By default, repeated authentications with the same code are filtered out (de-duplicated) and not sent to the MnT node. Instead a summary is sent every 15 minutes with a count of the number of de-duplicated log entries.

Although this does wonders for the scalability of the solution, it also makes troubleshooting a tad more difficult because Live Log is not being updated with those suppressed events. There are two common ways to disable the suppression and enable you to troubleshoot more effectively. The first is to use a collection filter to bypass suppression. The last resort is to disable the suppression all together.

Collection Filters

Collection filters were created to prevent certain events from being logged to the MnT. For example, they can be used to prevent the logging of load balancer health probes or the proactive test probes from switches.

For troubleshooting purposes, the filters can be used to bypass suppression. To configure a collection filter, do the following:

Step 1. From the ISE GUI, navigate to **Administration > Logging > Collection Filters**. Any existing filters will be displayed in the right pane.

Step 2. Click **Add** to create a new filter.

https://t.me/learningnets

Step 3. Select an attribute to uniquely identify which traffic should bypass suppression.

The options are User Name, Policy Set Name, NAS IP Address, Device IP Address, and MAC Address, as shown in Figure 22-28. The MAC address is the most common attribute to use when troubleshooting.

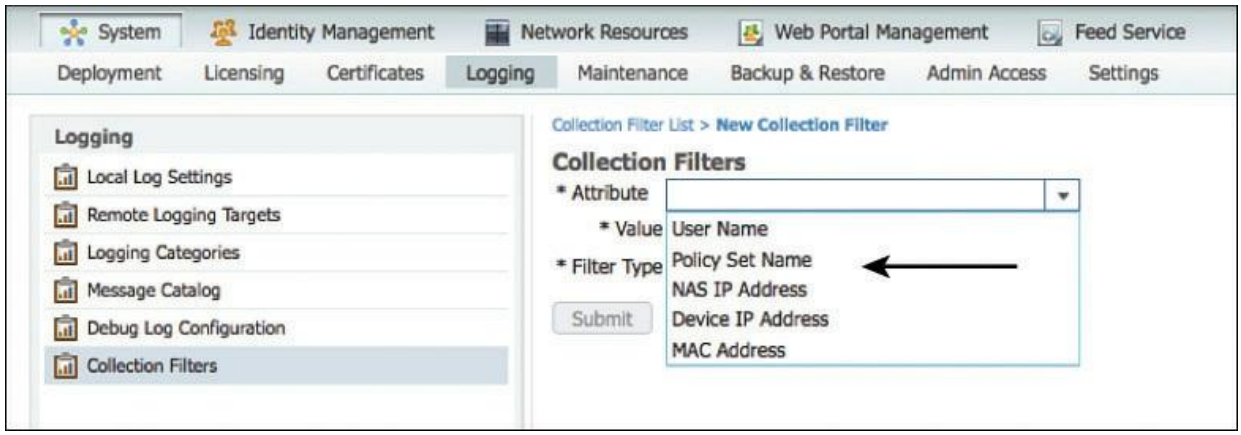


Figure 22-28 Choosing an attribute on which to filter.

Step 4. Enter the attribute value, such as the MAC address, as shown in Figure 22-29.

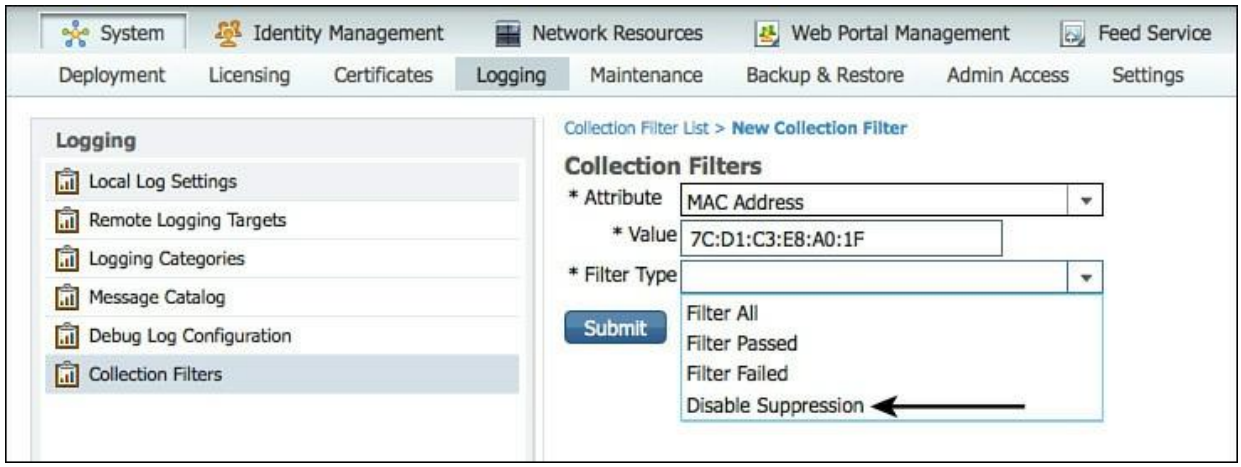


Figure 22-29 Enter an attribute value and select Disable Suppression.

Step 5. Select the Disable Suppression option from the Filter Type drop-down list, as shown in Figure 22-29.

Step 6. Click **Submit**.

Disabling Suppression

Bypassing suppression for a specific user or endpoint is a useful troubleshooting tool. However, that is not enough sometimes. Sometimes you must disable the suppression all together, like so:

Step 1. From the ISE GUI, navigate to **Administration > System > Settings > Protocols > RADIUS**, as shown in Figure 22-30.

https://t.me/learningnets

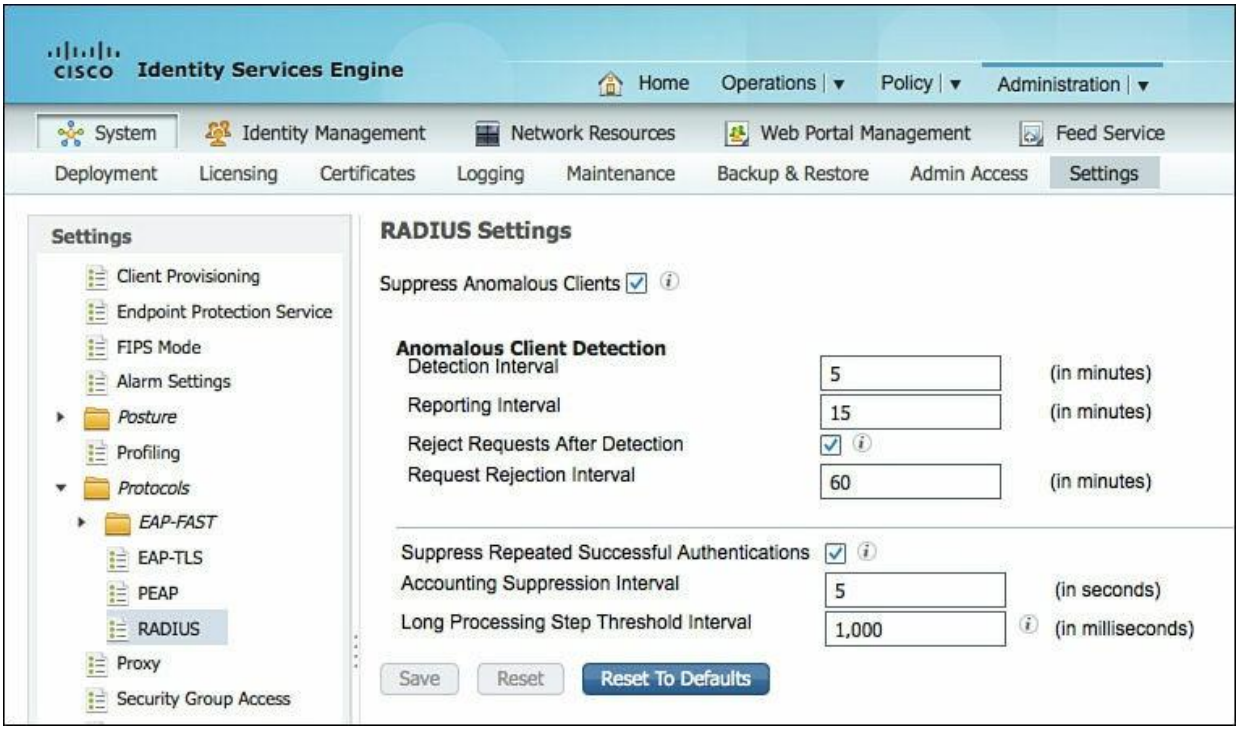


Figure 22-30 RADIUS settings.

Step 2. To disable suppression, uncheck the **Suppress Anomalous Clients** check box. Be sure to save afterwards.

Troubleshooting Outside of ISE

As described throughout this chapter and this book, ISE provides numerous tools that provide visibility to the authentication as it pertains to ISE's view of the network. Tools such as Live Log, Live Sessions Log, and TCP Dump very valuable when troubleshooting, but they cannot give a complete picture of all the parts of an authentication.

Think back to the basics of network authentication. There are three "actors" or participants in any authentication scenario: the endpoint, NAD, and authentication server. ISE is the authentication server and is world-class when it comes to providing troubleshooting details. However, there are still the other two participants in the authentication that you might need to look at: the endpoint and network device.

Endpoint Diagnostics

There are multiple aspects to troubleshooting from the 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 might need to be validated from the client; however, activities such as BYOD onboarding with the Network Setup Assistant apps or the posture agent (NAC Agent) also require validation.

AnyConnect Diagnostics and Reporting Tool

When troubleshooting 802.1X, it can often come down to seeing exactly what is happening at the client side, not only the server side. Unlike many of the supplicants out in the world, the AnyConnect Network Access Manager (NAM) performs detailed logging and includes an optional module called the Diagnostics and Reporting Tool (DART).

DART creates a support bundle that can be sent to Cisco TAC that includes all relevant system log files, as well as detailed communication from the supplicant. This is incredibly valuable when troubleshooting 802.1X authentications.

https://t.me/learningnets

To create a DART bundle on a Windows endpoint, do the following:

Step 1. Launch the DART tool from **Start > Cisco > Cisco AnyConnect Secure Mobility Client > Cisco AnyConnect Diagnostics and Reporting Tool**, as shown in Figure 22-31.

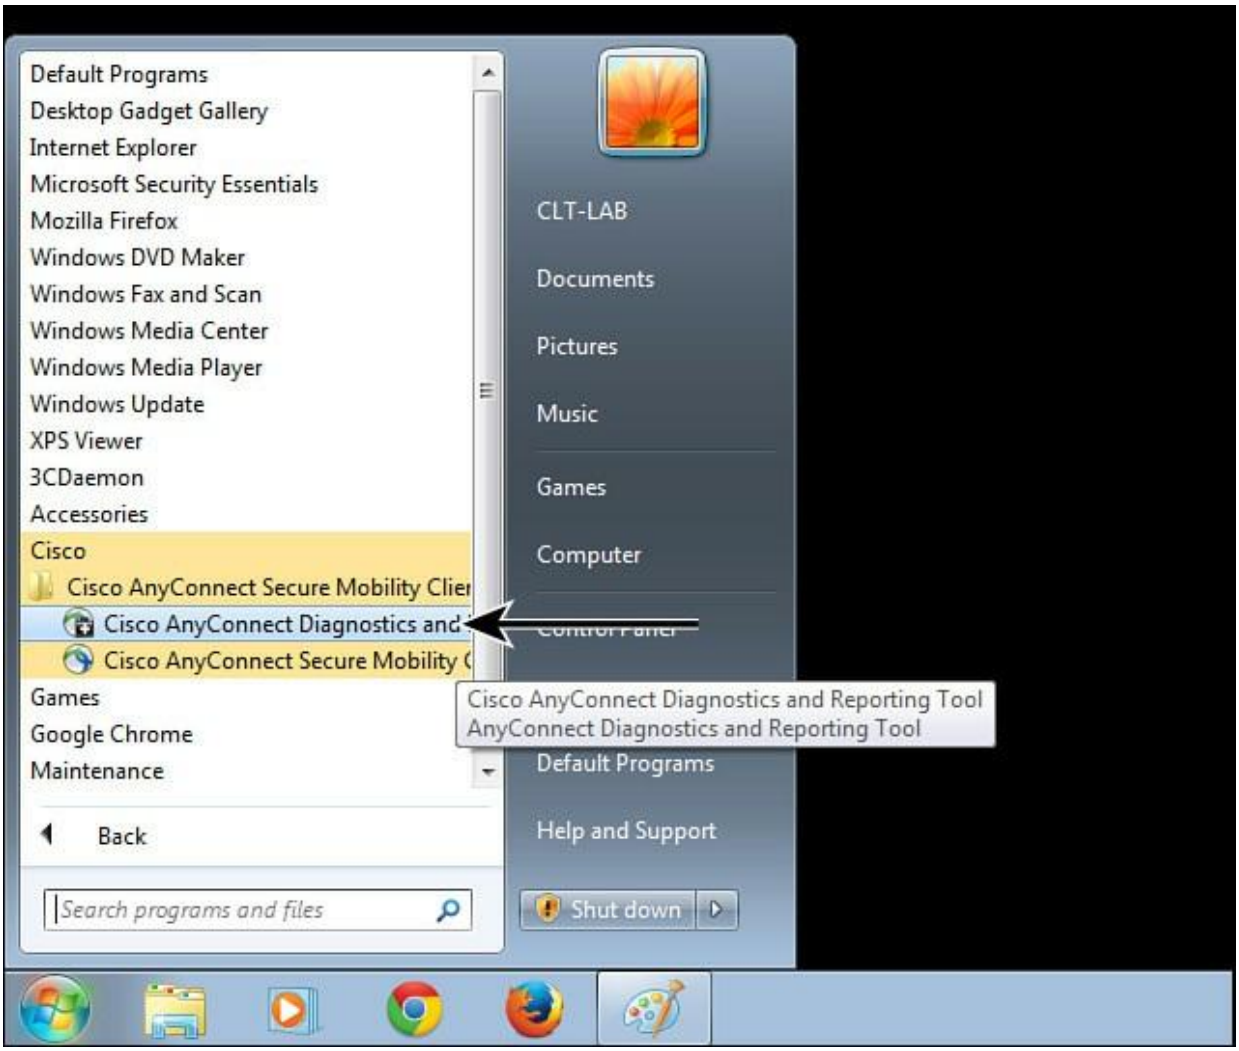


Figure 22-31 Launch DART from the Start menu.

Step 2. Click Next to begin the bundle creation, as shown in Figure 22-32.

https://t.me/learningnets



Figure 22-32 Click Next.

Step 3. Select the bundle creation option (default is normally used) and click **Next**, as shown in Figure 22-33. The DART tool will gather all necessary system and application logs, as shown in Figure 22-34.



Figure 22-33 The Bundle Creation option.

https://t.me/learningnets



Figure 22-34 The Bundle Creation option.

Step 4. When the bundle creation is complete, the resulting Zip file will be saved on the desktop, as shown in Figure 22-35.

https://t.me/learningnets

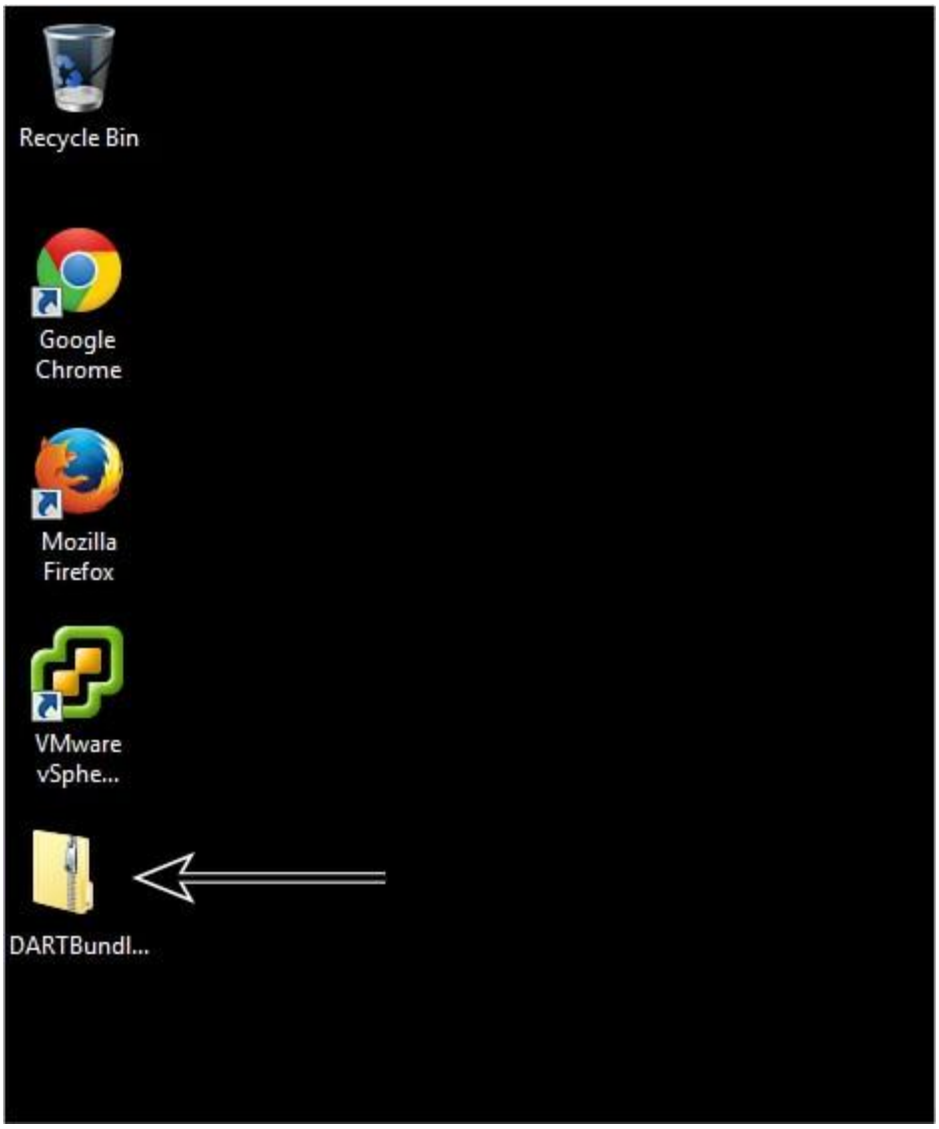


Figure 22-35 Resulting bundle Zip file on the desktop.

AnyConnect NAM Extended Logging

Even though AnyConnect NAM does have detailed logging capabilities enabled by default, it also has a special debug mode as well, known as extended logging. This mode can even provide packet captures from the client side for analysis, with a maximum file size of 10MB.

To enable extended logging, follow these steps:

Step 1. Type **ALT** + **SHIFT** + **L** at the desktop.

Step 2. Then right-click the **AnyConnect** icon in the system tray.

Step 3. Select **Extended Logging** from the context menu, as shown in Figure 22-36.

https://t.me/learningnets

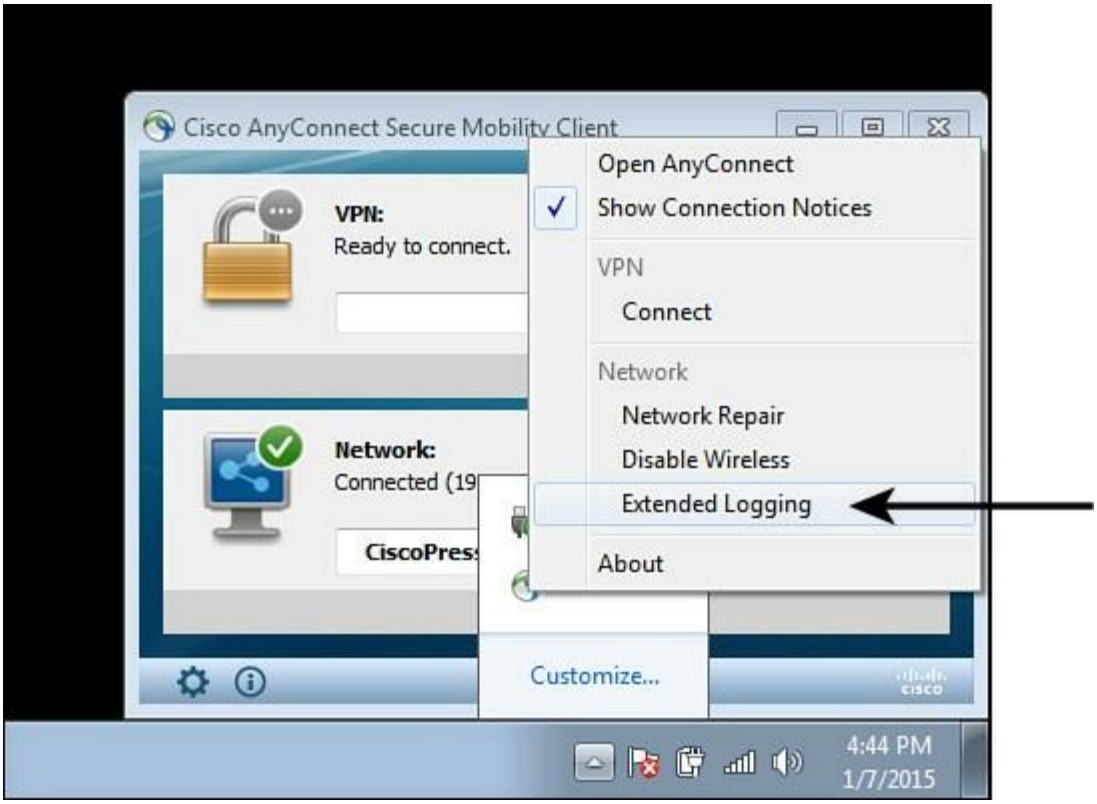


Figure 22-36 Enabling extended logging.

Microsoft Native Supplicant

There are many reasons that an organization might choose to deploy the supplicant that is native to the Windows OS, especially given the terrific centralized control offered by Active Directory. However, Cisco TAC has accumulated a list of eight hotfixes that an organization should ensure are installed if it is going to use the native supplicant:

- **KB 976373**—Windows 7 connected behind IP phones will not authenticate after waking up from sleep or hibernation.
- **KB 980295**—Windows 7 stops responding to 802.1X after first authentication fails.
- **KB 2481614**—Windows 7 selects a protocol different from what the GPO states.
- **KB 2491809**—Windows 7 does not prompt for 802.1X credentials to some users on a shared PC.
- **KB 2835595**—Windows 7 does not prompt for 802.1X credentials.
- **KB 2494172**—Windows 7 cannot authenticate if a valid certificate and an invalid certificate are present.
- **KB 2769121**—Windows 7 selects the wrong certificate for a machine migrated across two forests.
- **KB 2736878**—Windows 7 authentication fails intermittently.

Vivek Santuka, CCIE #17621 maintains a blog entry on the Cisco Support Forums with details on the subject. That blog is located at https://supportforums.cisco.com/blog/12256681/getting-past-intermittentunexplained-8021x-problems-windows-7.

Supplicant Provisioning Logs

If an endpoint attempts a BYOD onboarding flow but is not able to complete the native supplicant provisioning (NSP) process, examining the local endpoint logs might be the only way to find the root cause. Each operating system stores the logs in a different location, and some may even require special software to examine the logs.

https://t.me/learningnets

Here is a list of useful techniques that are used to troubleshoot client-side logging issues:

- **Microsoft Windows**—Enter the **Log %temp%\spwProfileLog.txt** command to view the client-side logs.
- **Android**—Enter the **/sdcards/downloads/spw.log** command to view the client-side logs.
- **Mac OSX**—Use the Console application, and look for the SPW process.
- **Apple iOS**—Use the iPhone Configuration Utility (iPCU) to view messages. It can be downloaded 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.

The Go-To: show authentication session interface

The command every ISE admin should have in her back pocket is the **show authentication interface** command. As described in Chapter 12, this command output shows the current state of the authentication and the results sent back from ISE. As shown in the bold text within Example 22-5, 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 was assigned an SGT value of 8, and it is the result of a MAB.

Example 22-5 show authentication session interface

Click here to view code image

```
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
```

https://t.me/learningnets

Viewing Client Details on the WLC

The WLC has a similar view, 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 22-37.

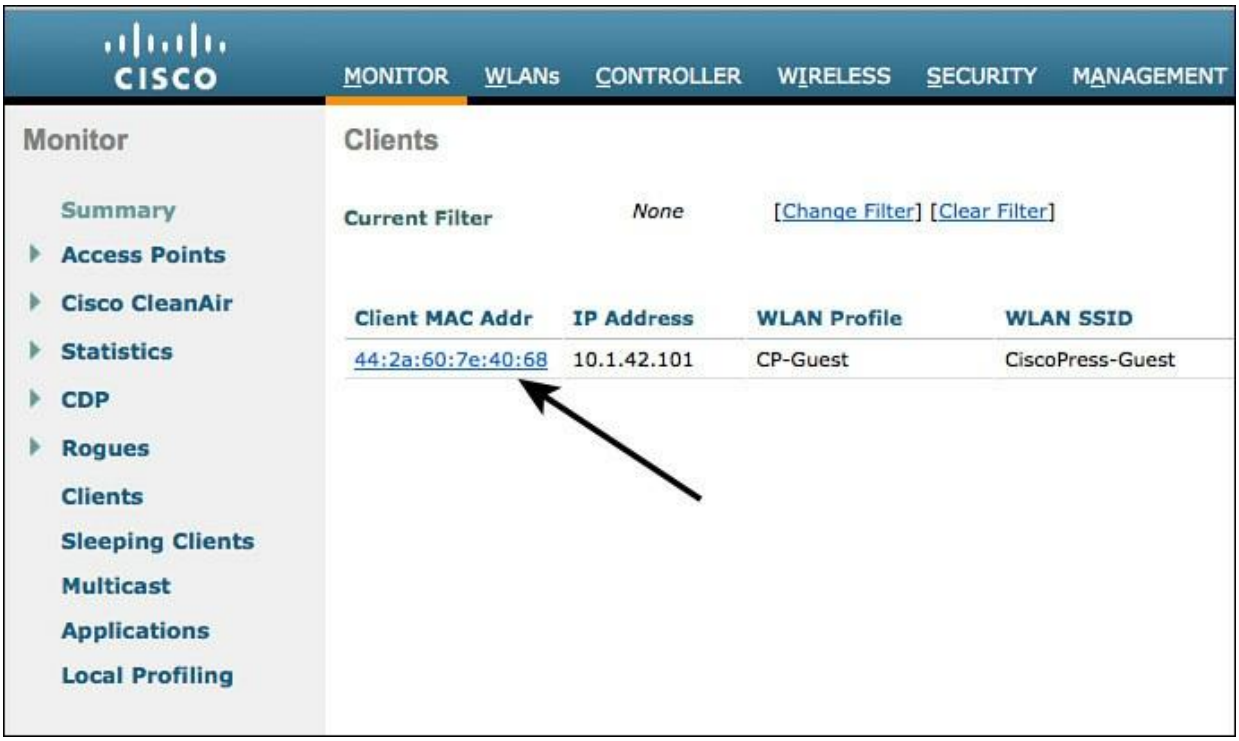


Figure 22-37 Monitor > Clients.

Select the client you are interested in verifying by clicking its MAC address. This brings up the client details screen shown in Figure 22-38. Scrolling down to the Security Information section, you can see RADIUS NAC State is set to CENTRAL_WEB_AUTH, the ACL-WEBAUTH-REDIRECT redirection ACL is applied, it was assigned an SGT value of 8, and the session is being redirected to a CWA portal on the ISE node.

https://t.me/learningnets

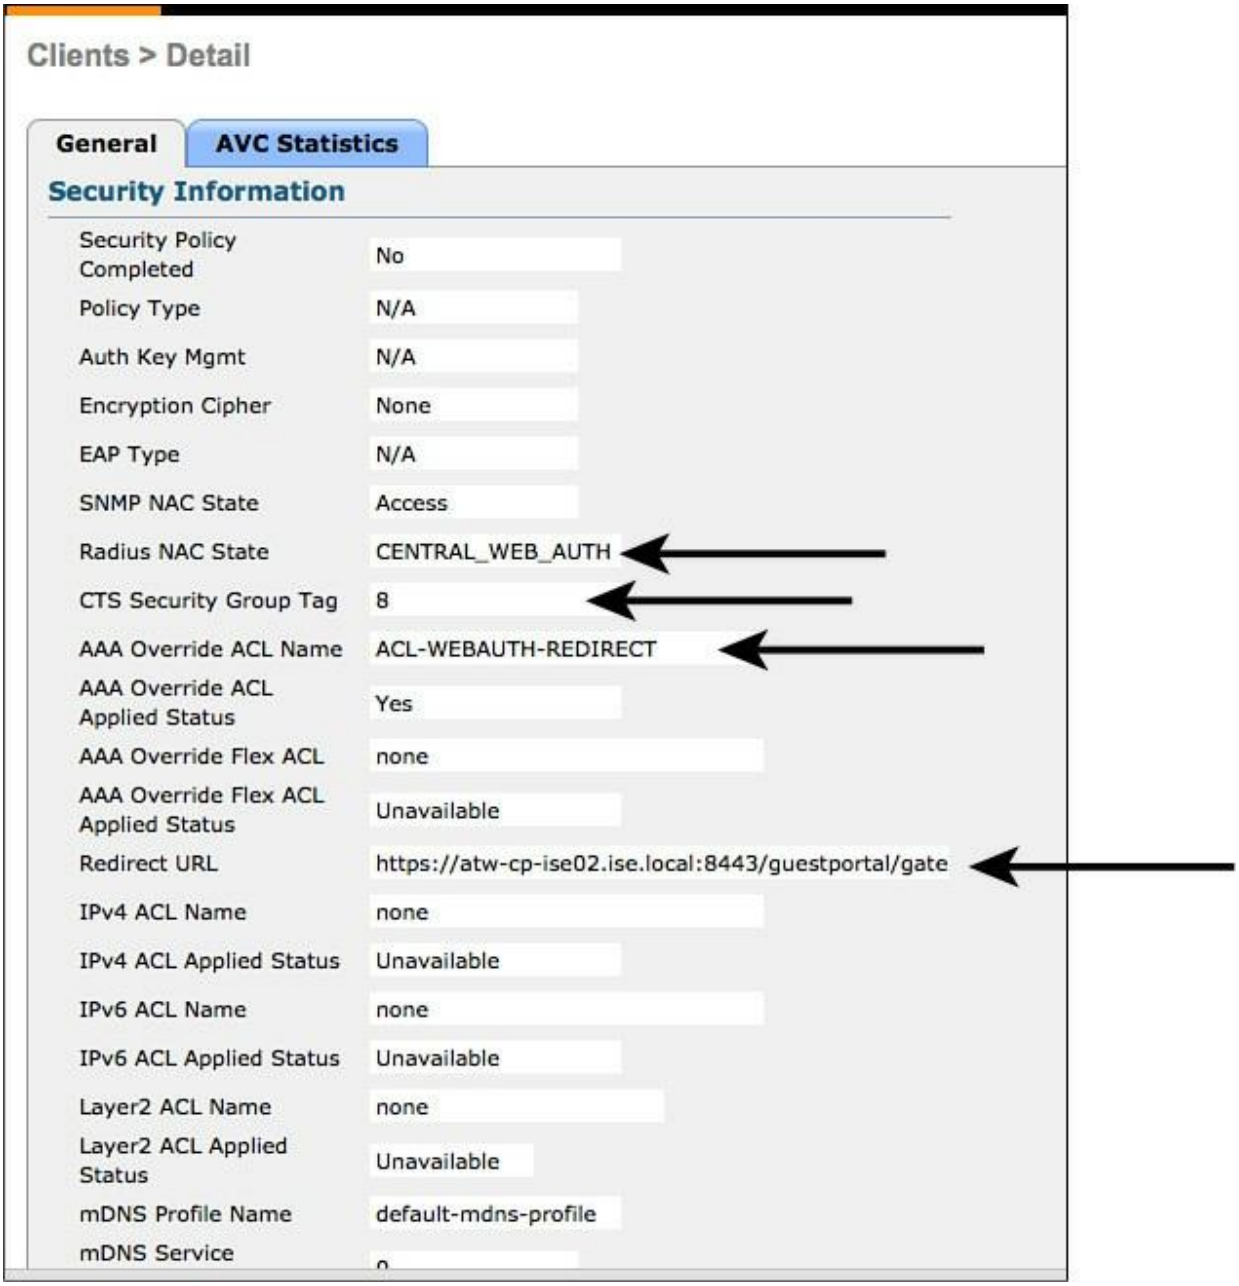


Figure 22-38 Clients > Detail.

Debug Commands

A number of useful debug commands can be used on Cisco switches and WLCs to see the details of an authentication. This type of visibility tends to be taken for granted, until such time as you need to troubleshoot.

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 that client is shown in the debug output, and no other client traffic on the controller is affected.

Cisco switches do not offer a debug client command, but they do have the ability to set the specific interface to constrain the debug command to **debug interface <interface name>**. Using this command before entering one of the other debug commands ensures that only the specified interface is affected with the debug process.

Other useful debug commands for Cisco switches include the following:

https://t.me/learningnets

- **debug authentication [all | errors | events | feature | sync]**—Use this 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.

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 22-2 lists a reference of these key topics and the page numbers on which each is found.



Key Topic Element	Description	Page
Section	Live Sessions Log	728
Section	Logging Targets	729
Section	Support Bundles	734
Section	TCP Dump	741

Table 22-2 Key Topics for Chapter 22

https://t.me/learningnets

Part VII: Final Preparation

https://t.me/learningnets

Chapter 23. Final Preparation

Congratulations! You made it through the book, and now it's time to finish getting ready for the exam. This chapter helps you get ready to take and pass the exam in two ways.

This chapter begins by talking about the exam itself. You know the content and topics. Now you need to think about what happens during the exam, and what you need to do in these last few weeks before taking the exam. At this point, everything you do should be focused on getting you ready to pass so that you can finish. You are very close to the finish line, so let's put forth one last effort. Everything is possible; you just need to believe it.

The second section of this chapter gives you some exam review tasks as your final preparation for your SISAS 300-208 exam.

Advice About the Exam Event

Now that you have finished most of this book, you could register for your Cisco SISAS 300-208 exam, show up, and take the exam. However, if you spend a little time thinking about the exam event itself, learning more about the user interface of the real Cisco exams and the environment at the Vue testing centers, you will be better prepared—particularly if this is your first Cisco exam. This first of three major sections in this chapter provides some advice about the Cisco exams and the exam event itself.

Learning the Question Types Using the Cisco Certification Exam Tutorial

In the weeks leading up to your exam, you should think more about the various types of exam questions and have a plan for how to approach those questions. One of the best ways to learn about the exam questions is to use the Cisco Exam Tutorial.

To find the Cisco Certification Exam Tutorial, go to www.cisco.com and search for "exam tutorial." The tutorial sits inside a webpage with a Flash presentation of the exam user interface. The tutorial even lets you take control as if you were taking the exam. When using the tutorial, make sure that you take control and try the following:

- Try to click Next on a multiple-choice, single-answer question without clicking an answer, and see that the testing software tells you that you have too few answers.
- On a multiple-choice, multiple-answer question, select too few answers and click Next to again see how the user interface responds.
- In a drag-and-drop question, drag the answers to the obvious answer locations, but then drag them back to the original location. (You might do this on the real exam if you change your mind when answering the question.)
- On a simulation question, first just make sure that you can get to the command-line interface (CLI) on one of the routers. To do so, you have to click the PC icon for a PC connected to the router console; the console cable appears as a dashed line, and network cables are solid lines. (Note that the exam tutorial uses the IOS CLI, not NX-OS, but it is similar enough to get the idea.)
- Still on the simulation question, make sure that you look at the scroll areas at the top, at the side, and in the terminal emulator window.
- Still on the simulation question, make sure that you can toggle between the topology window and the terminal emulator window by clicking Show Topology and Hide Topology.
- On a testlet question, answer one multiple-choice question, move to the second, answer it, move back to the first question, and confirm that inside a testlet, you can move around between questions.
- Again on the testlet question, click the Next button to see the pop-up window Cisco uses as a prompt to ask whether you want to move on. Testlets might actually allow you to give too few answers and still move on. After you click to move past the testlet, you cannot go back to change

https://t.me/learningnets

your answer for any of these questions.

Thinking About Your Time Budget Versus Number of Questions

On exam day, you need to keep an eye on your speed. Going too slowly hurts you because you might not have time to answer all the questions. Going too fast can hurt as well if your fast speed is because you are rushing and not taking the time to fully understand the questions. So, you need to be able to know whether you are moving quickly enough to answer all the questions while not rushing.

The exam user interface shows some useful time information—namely, a countdown timer as well as a question counter. The question counter shows a question number for the question you are answering, and it shows the total number of questions on your exam.

Unfortunately, treating each question equally does not give you an accurate time estimate. For example, if your exam allows 90 minutes and has 45 questions, you would have 2 minutes per question. After answering 20 questions, if you had taken 40 minutes, you would be right on time. However, several factors make that kind of estimate difficult.

First, Cisco does not tell us beforehand the exact number of questions for each exam. For example, Cisco.com (at the time this book was published) listed SISAS as a 90-minute exam with 65–75 questions. So, you only know a range of questions. But you do not know how many questions are on your exam until it begins, when you go through the screens that lead up to the point where you click Start Exam.

Next, some questions (call them time burners) clearly take a lot more time to answer:

- **Normal-time questions**—Multiple-choice and drag-and-drop take approximately 1 minute each.
- **Time burners**—Simulations, simlets, and testlets take approximately 4–5 minutes each.

Finally, the exam software counts each testlet and simlet question as one question in the question counter. For example, if a testlet question has four embedded multiple-choice questions, in the exam software's question counter that counts as one question.

You need a plan for how you will check your time, and use a plan that does not distract you from the exam. It might be worth taking a bit of a guess to keep things simple, like this:

> 60 questions, 90 minutes, are exactly 1:15 per question. Then just guess a little based on how many time-burner questions you have seen so far.

No matter how you plan to check your time, think about it before exam day. You can even use the method listed in the next section.

A Suggested Time-Check Method

The following math can be used to do your time check in a way that weights the time based on those time-burner questions. You do not have to use this method. But this math uses only addition of whole numbers to keep it simple. It gives you a pretty close time estimate, in my opinion.

The concept is simple. Just do a simple calculation that estimates the time you should have used so far. Basically, this process gives you 1 minute for normal questions and 7 minutes per time burner; here's the math:

> Number of Questions Answered So Far + 7 Per Time Burner

Then, check the timer to figure out how much time you have spent:

- You have used exactly that much time, or a little more: Your timing is perfect.
- You have used less time: You are ahead of schedule.
- You have used noticeably more time: You are behind schedule.

For example, if you have already finished 17 questions, 2 of which were time burners, your time estimate is 17 + 7 + 7 = 31 minutes. If your actual time is also 31 minutes, or maybe 32 or 33 minutes, you are right on schedule. If you have spent less than 31 minutes, you are ahead of schedule.

So, the math is pretty easy: Questions answered, plus 7 per time burner, is the guesstimate of how long you

https://t.me/learningnets

should have taken so far if you are right on time.

> **Note**
>
> This math is an estimate; we make no guarantees that the math will be an accurate predictor on every exam.

Miscellaneous Pre-Exam Suggestions

Here are just a few more suggestions for things to think about before exam day arrives:

- Get some earplugs. Testing centers often have some, but if you do not want to chance it, come prepared. The testing center is typically a room inside the space of a company that does something else as well, often a training center. So, people might be talking in nearby rooms and there might be other office noises. Earplugs can help. (Headphones, as electronic devices, are not allowed.)
- Some people like to spend the first minute of the exam writing down notes for reference. For example, maybe you want to write down the table of magic numbers for finding IPv4 subnet IDs, the critical steps required to accomplish a certain task within ISE, or the more common RADIUS Attribute-Value pairs. If you plan to do that, practice making those notes. Before each practice exam, transcribe those lists, just like you expect to do at the real exam.
- Plan your travel to the testing center with enough time so that you will not be rushing to make it just in time.
- If you tend to be nervous before exams, practice your favorite relaxation techniques for a few minutes before each practice exam, just to be ready to use them.

Exam-Day Advice

We hope the exam goes well for you. Certainly, the better prepared you are, the better chances you have on the exam. But these small tips can help you do your best on exam day:

- Rest the night before the exam, rather than staying up late to study. Clarity of thought is more important than one extra fact, especially because the exam requires so much analysis and thinking rather than just remembering facts.
- If you did not bring earplugs, ask the testing center for some, even if you cannot imagine you would use them. You never know whether it might help.
- You can bring personal effects into the building and the testing company's space, but not into the actual room in which you take the exam. So, take as little extra stuff with you as possible. If you have a safe place to leave briefcases, purses, electronics, and so on, leave them there. However, the testing center should have a place to store your things as well. Simply put, the less you bring, the less you have to worry about storing. Still, some testing centers actually insist that you leave cell phones in your car.
- The exam center will give you a laminated sheet and pen to take notes. (Test center personnel typically do not let you bring paper and pen into the room, even if supplied by the testing center.)
- Leave for the testing center with extra time, so you do not have to rush.
- Plan on finding a restroom before going into the testing center. If you cannot find one, of course you can use one in the testing center, and test personnel will direct you and give you time before your exam starts.
- Do not drink a 64-ounce drink on the trip to the testing center. After the exam starts, the exam timer will not stop while you go to the restroom.
- On exam day, use any relaxation techniques that you have practiced to help get your mind focused while you wait for the exam.

https://t.me/learningnets

Exam Review

This exam review completes the Study Plan materials as suggested by this book. At this point, you have read the other chapters of the book and have done the chapter review Exam Preparation Tasks. Now you need to do the final study and review activities before taking the exam, as detailed in this section.

The Exam Review section suggests some new activities, as well as repeating some old ones. However, whether new or old, the activities all focus on filling in your knowledge gaps, finishing your skills, and completing the study process. Although repeating some tasks you did at chapter review and part review can help, you need to be ready to take an exam, so the Exam Review asks you to spend a lot of time answering exam questions.

The Exam Review walks you through suggestions for several types of tasks and gives you some tracking tables for each activity. The main categories are as follows:

- Practicing for speed
- Taking practice exams
- Finding what you do not know well yet (knowledge gaps)
- Configuring and verifying functions from the CLI or GUI
- Repeating the chapter review tasks

Taking Practice Exams

One day soon, you will need to pass a real Cisco exam at a Vue testing center. So, it's time to practice the real event as much as possible.

A practice exam using the Pearson IT Certification Practice Test (PCPT) exam software lets you experience many of the same issues as when taking a real Cisco exam. The software gives you a number of questions, with a countdown timer shown in the window. After you answer a question, you cannot go back to it (yes, that's true on Cisco exams). If you run out of time, the questions you did not answer count as incorrect.

The process of taking the timed practice exams helps you prepare in three key ways:

- To practice the exam event itself, including time pressure, the need to read carefully, and a need to concentrate for long periods
- To build your analysis and critical thinking skills when examining the network scenario built in to many questions
- To discover the gaps in your networking knowledge so you can study those topics before the real exam

As much as possible, treat the practice exam events as if you were taking the real Cisco exam at a Vue testing center. The following list gives some advice on how to make your practice exam more meaningful, rather than as just one more thing to do before exam day rolls around:

- Set aside two hours for taking the 90-minute, timed practice exam.
- Make a list of what you expect to do for the 10 minutes before the real exam event. Then visualize yourself doing those things. Before taking each practice exam, practice those final 10 minutes before your exam timer starts. (The earlier section, "Exam-Day Advice," lists some suggestions about what to do in those last minutes.)
- You cannot bring anything with you into the Vue exam room, so remove all notes and help materials from your work area before taking a practice exam. You can use blank paper, a pen, and your brain only. Do not use calculators, notes, web browsers, or any other app on your computer.
- Real life can get in the way, but if at all possible, ask anyone around you to leave you alone for the time you will practice. If you must do your practice exam in a distracting environment, wear headphones or earplugs to reduce distractions.
- Do not guess, hoping to improve your score. Answer only when you have confidence in the answer. Then, if you get the question wrong, you can go back and think more about the question in a later

https://t.me/learningnets

study session.

Practicing Taking the SISAS Exam

To take a SISAS practice exam, you need to select one or both of the SISAS exams from PCPT. If you followed the study plan in this book, you will not have seen any of the questions in these two exam databases before now. After you select one of these two exams, you simply need to choose the Practice Exam option in the upper-right corner and start the exam.

If you are interested in purchasing more practice exams, check out the CCNP Security SISAS 300-208 Official Cert Guide Premium Edition eBook and Practice Test product at www.ciscopress.com/title/9780133888782, and be sure to use the 70% off coupon included in the CD sleeve of this book.

Table 23-1 gives you a checklist to record your practice exam events. Note that recording both the date and the score is helpful for some other work you will do, so note both. Also, in the Time Notes section, if you finish on time, note how much extra time you had; if you run out of time, note how many questions you did not have time to answer.

Exam	Date	Score	Time Notes
SISAS 300-208			
SISAS 300-208			
SISAS 300-208			
SISAS 300-208			

Table 23-1 SISAS 300-208 Practice Exam Checklist

Advice on How to Answer Exam Questions

Open a web browser. Yes, take a break and open a web browser on any device. Do a quick search on a fun topic. Then, before you click a link, get ready to think where your eyes go for the first 5–10 seconds after you click the link. Now, click a link and look at the page. Where did your eyes go?

Interestingly, web browsers and the content on those webpages have trained us all to scan. Webpage designers actually design content with the expectation that people will scan with different patterns. Regardless of the pattern, when reading a webpage, almost no one reads sequentially, and no one reads entire sentences. They scan for the interesting graphics and the big words, and then scan the space around those noticeable items.

Other parts of our electronic culture have also changed how the average person reads. For example, many of you grew up using texting and social media, sifting through hundreds or thousands of messages—but each message barely fills an entire sentence. (In fact, that previous sentence would not fit in a tweet because it's longer than 140 characters.)

Those everyday habits have changed how we all read and think in front of a screen. Unfortunately, those same habits often hurt our scores when taking computer-based exams.

If you scan exam questions like you read webpages, texts, and tweets, you will probably make some mistakes because you missed a key fact in the question, answer, or exhibits. It helps to start at the beginning and read all the words—a process that is amazingly unnatural for many people today.

When taking the practice exams and answering individual questions, let me make two suggestions. First, before the practice exam, think about your own personal strategy for how you will read a question. Make your approach to multiple-choice questions in particular be a conscious decision on your part. Second, if you want some suggestions on how to read an exam question, use the following strategy:

Step 1. Read the question itself, thoroughly, from start to finish.

Step 2. Scan any exhibit (usually command output) or figure.

https://t.me/learningnets

Step 3. Scan the answers to look for the types of information. (Numeric? Terms? Single words? Phrases?)

Step 4. Reread the question thoroughly, from start to finish, to ensure that you understand it.

Step 5. Read each answer thoroughly, while referring to the figure/exhibit as needed. After reading each answer, before reading the next answer:

A. If correct, select as correct.

B. If for sure it is incorrect, mentally rule it out.

C. If unsure, mentally note it as a possible correct answer.

> **Note**
>
> Cisco exams will tell you the number of correct answers. The exam software also helps you finish the question with the right number of answers noted. For example, the software prevents you from selecting too many answers. Also, if you try to move on to the next question but have too few answers noted, the exam software asks if you truly want to move on.

Use the practice exams as a place to practice your approach to reading. Every time you click to the next question, try to read the question following your approach. If you are feeling time pressure, that is the perfect time to keep practicing your approach to reduce and eliminate questions you miss because of scanning the question instead of reading thoroughly.

Taking Other Practice Exams

Many people use other practice exams and questions other than the questions that come with this book. Frankly, using other practice exams in addition to the questions that come with this book can be a good idea, for many reasons. The other exam questions can use different terms in different ways, emphasize different topics, and show different scenarios that make you rethink some topics.

No matter where you get additional exam questions, if you use the exam questions for a timed practice exam, it helps to take a few notes about the results. Table 23-2 gives you a place to take those notes. Also, take a guess at the percent of questions you have seen before taking the exam, and note whether you think the questions are less, more, or the same challenge level as the questions that come with this book. And as usual, note whether you ran out of time or had extra time left over at the end.

Exam Source	Other Exam Notes	% Questions Repeated (Estimate)	Challenging?	Date	Score	Time Notes

Table 23-2 Checklist for Practice Exams from Other Sources

Note that the publisher does sell products that include additional test questions. The CCNP Security SISAS 300-208 Official Cert Guide Premium Edition eBook and Practice Test product is basically the publisher's eBook version of this book. It includes a soft copy of the book, in formats you can read on your computer or on the most common book readers and tablets. The product includes all the content you would normally get with the DVD that comes with the print book, including all the question databases mentioned in this chapter. Additionally, this product includes two more SISAS exam databases for extra exams.

https://t.me/learningnets

Finding Knowledge Gaps Through Question Review

You just took a number of practice exams. You probably learned a lot, gained some exam-taking skills, and improved your networking knowledge and skills. But if you go back and look at all the questions you missed, you might be able to find a few small gaps in your knowledge.

One of the hardest things to find when doing your final exam preparation is to discover gaps in your knowledge and skills. In other words, what topics and skills do you need to know that you do not know? Or what topics do you think you know, but you misunderstand about some important fact? Finding gaps in your knowledge at this late stage requires more than just your gut feeling about your strengths and weaknesses.

This next task uses a feature of PCPT to help you find those gaps. The PCPT software tracks each practice exam you take, remembering your answer for every question and whether you got it wrong. You can view the results and move back and forth between seeing the question and seeing the results page. To find gaps in your knowledge, follow these steps:

Step 1. Pick and review one of your practice exams.

Step 2. Review each incorrect question until you are happy that you understand the question.

Step 3. When finished with your review for a question, mark the question.

Step 4. Review all incorrect questions from your exam until all are marked.

Step 5. Move on to the next practice exam.

Figure 23-1 shows a sample Question Review page, in which all the questions were answered incorrectly. The results list a Correct column, with no check mark meaning that the answer was incorrect.

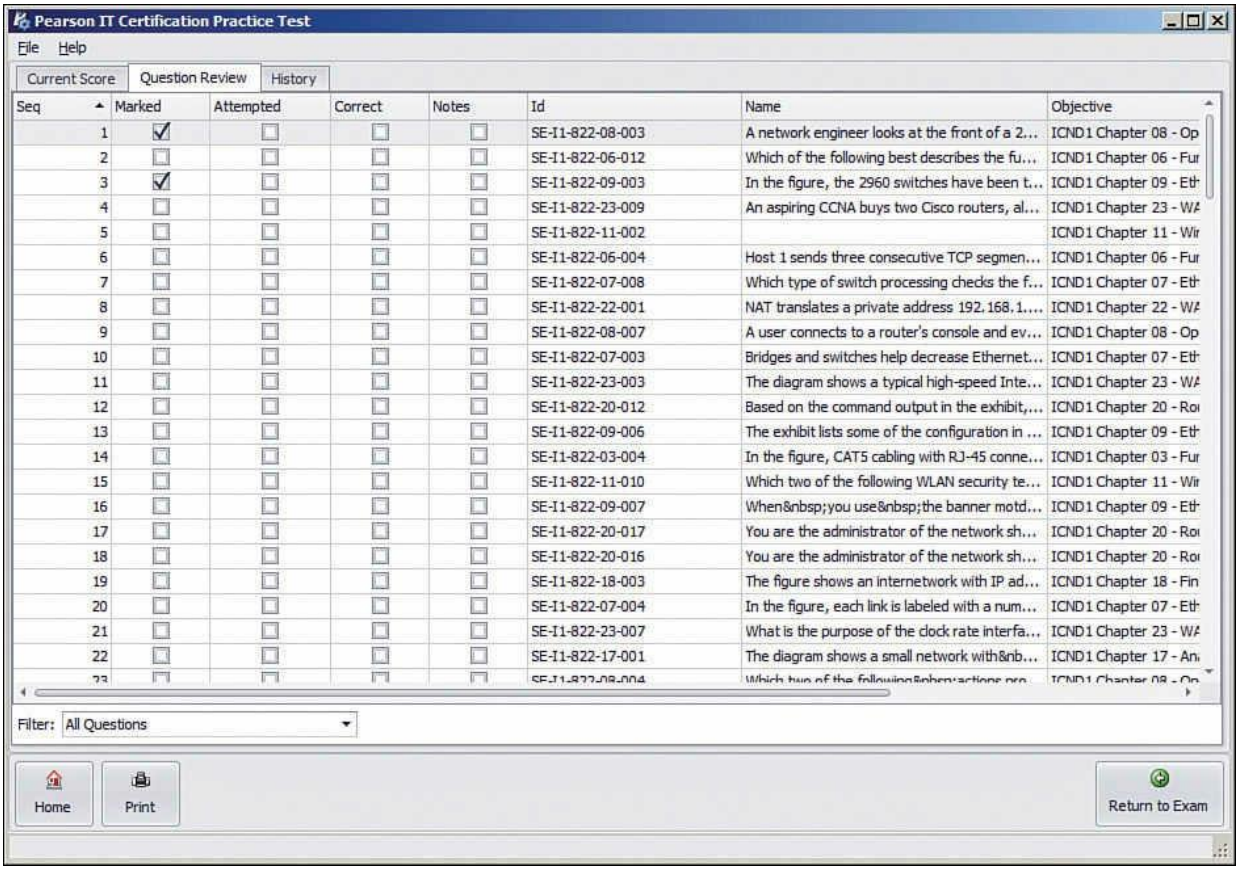


Figure 23-1 PCPT grading results page.

To perform the process of reviewing questions and marking them as complete, you can move between this Question Review page and the individual questions. Just double-click a question to move back to that question. From the question, you can click Grade Exam to move back to the grading results and to the

https://t.me/learningnets

Question Review page shown in Figure 23-1. The question window also shows the place to mark the question, in the upper left, as shown in Figure 23-2.

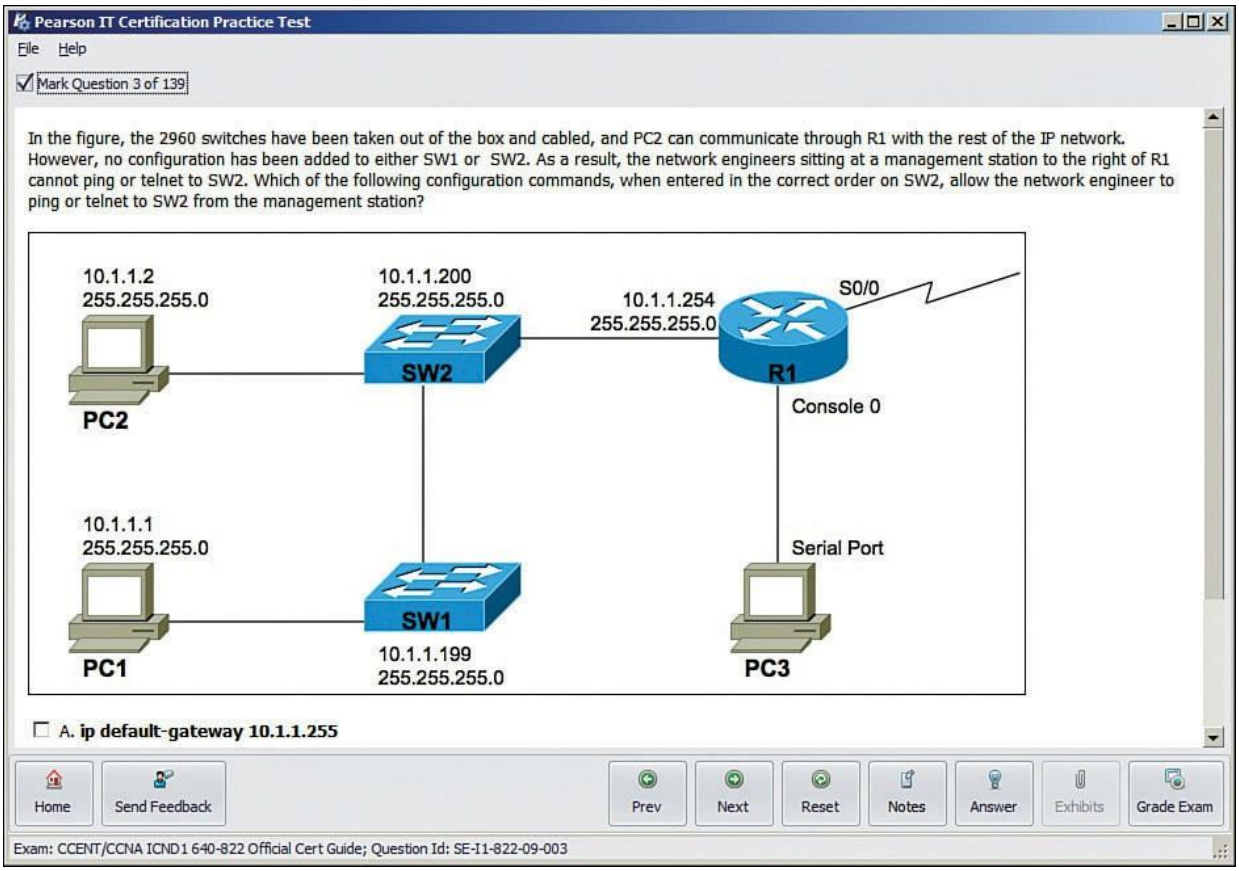


Figure 23-2 Reviewing a question with the mark feature in upper left.

If you want to come back later to look through the questions you missed from an earlier exam, start at the PCPT home screen. From there, instead of clicking the Start button to start a new exam, click the View Grade History button to see your earlier exam attempts and work through any missed questions.

Track your progress through your gap review in Table 23-5. PCPT lists your previous practice exams by date and score, so it helps to note those values in the table for comparison to the PCPT menu.

SISAS 300-208 Exam	Original Practice Exam Date	Original Exam Score	Date Gap Review Was Completed

Table 23-5 Tracking Checklist for Gap Review of Practice Exams

Other Study Tasks

If you get to this point and still feel the need to prepare some more, this last topic gives you three suggestions.

https://t.me/learningnets

First, the chapter review Exam Preparation Tasks sections give you some useful study tasks.

Second, take more exam questions from other sources. You can always get more questions in the Cisco Press Premium Edition eBook and Practice Test products, which include an eBook copy of this book plus additional questions in additional PCPT exam banks. You also can search the Internet for questions from many sources and review those questions as well.

> **Note**
>
> Some vendors claim to sell practice exams that contain the literal exam questions from the exam. These exams, called "brain dumps," go against the Cisco testing policies. Cisco strongly discourages using any such tools for study.

Finally, join in the discussions on the Cisco Learning Network. Try to answer questions asked by other learners; the process of answering makes you think much harder about the topic. When someone posts an answer with which you disagree, think about why and talk about it online. This is a great way to both learn more and build confidence.

Final Thoughts

You have studied quite a bit, worked very hard, and sacrificed time and money to be ready for the exam. We hope your exam goes well, that you pass, and that you pass because you really know your stuff and will do well in your IT and networking career.

We encourage you to celebrate when you pass and ask advice when you do not. Congratulations on achieving a major milestone in your career.

https://t.me/learningnets

Part VIII: Appendixes

https://t.me/learningnets

Appendix A. Answers to the "Do I Know This Already?" Quizzes

Chapter 2

1. D. Simply put, authentication is the validation of the identity credentials. Authorization is the determination of what is allowed or disallowed based on those credentials.
2. A, D. The two forms of authentication, authorization, and accounting that are relevant to the SISAS exam are network access and device administration.
3. B. TACACS+ is best suited for granular command-level control due to its ability to separate authentication and authorization.
4. C. RADIUS is best suited for network access AAA due to its capability to work with numerous authentication protocols, such as CHAP and MS-CHAPv2, but more importantly the dependency on RADIUS for 802.1X authenticationsand the enhancements to RADIUS for change of authorization.
5. A. Both TACACS+ and RADIUS can be used to provide device administration AAA services; however, TACACS+ offers command-level authorization and RADIUS does not.
6. A. Cisco ACS supports both RADIUS and TACACS+ and command sets, while Cisco ISE version 1.2 supports only RADIUS.
7. D. The majority of the authentication protocols used (EAP, CHAP, MS-CHAPv2, PAP) are Layer-2 protocols meant to be topology independent. RADIUS and TACACS+ are used to connect the end user to the authentication server, even when they are not on the same LAN segment.
8. A. TACACS+ clients send only two message types: START and CONTINUE. REPLY is sent from the AAA server to the AAA client.
9. B. The Service-Type value tells the RADIUS server what is being performed. For example, service-type of Call-Check informs the AAA server that the client is performing a MAB request.
10. A. The RADIUS server may be assigning an attribute to the authentication session, like a VLAN, for example. The VLAN place holder is the attribute, and the actual assigned VLAN number is the value for that place holder, as a pair.

Chapter 3

1. B, C. An identity is a representation of who a user or device is. Cisco ISE uses an endpoint's MAC address to uniquely identify that endpoint. A username is one method of uniquely identifying an end user. Although SSIDs and IP addresses can be used as conditions or attributes in ISE policies, they are not identities.
2. B, C. Cisco ISE can use identities stored in a database that resides as part of the ISE application itself; these are known as internal identity stores. Examples are the GUEST user identity store and the endpoints identity store. Identities can live outside of ISE, such as Active Directory, and these are known as external identity stores.
3. A, D. ISE has two different types of internal identity stores: users and endpoints. The user identity stores hold identities for interactive users, such as guests or employees. These have attributes such as passwords for the authentication of the user. Endpoints have a different kind of identity. Because they don't interact with an authentication in most cases, their identities can often just be their MAC addresses.
4. A, D. Either a user or a machine (endpoint) can be authorized for network access. Sometimes it is possible to authorize based on the identity or attributes of both the user and the machine.
5. C. The identity store is known as an identity source or an information source. The data contained in the identity store is used for authentication and authorization purposes.
6. C. An identity source sequence (ISS) is a list of identity stores. Much like an access control list

https://t.me/learningnets

(ACL), the ISS list is processed with from the top to the bottom, where the first entry that has the identity is used and the processing of the ISS ends.

7. A, C, D. Lightweight Directory Access Protocol is a standard directory type that allows vendors to use a common communication structure to provide authentications and information about identities. Microsoft's Active Directory is an LDAP-like directory source and is one of the most common identity sources in the modern world. In addition to querying an identity source directly, ISE is also able to proxy RADIUS authentications to a different RADIUS server.

8. B, D. Internal identity stores can be used to authenticate user accounts or endpoints. A guest is a type of internal user that ISE can authenticate. MAB is often used to "authenticate" endpoints against the internal endpoints identity store.

9. A, B. ISE has two different types of internal identity stores: users and endpoints. The user identity stores hold identities for interactive users, like guests or employees. These have attributes such as passwords for the authentication of the user. Endpoints have a different kind of identity. Because they don't interact with an authentication in most cases, their identities can often just be their MAC addresses.

10. C, D. External identity stores often exist already in an organization before ISE would be installed. By pointing to those identity sources, the management overhead is dramatically reduced because the accounts don't have to be created again in ISE's internal database(s). Additionally, this enables the organization to scale more effectively by having a single source of truth for identity.

Chapter 4

1. B. 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. Only Cisco AnyConnect NAM 3.1 and newer are capable of running EAP chaining as of the date this book was published.

3. C. The outer identity provides a mechanism to authenticate the identity of the endpoint during the tunnel establishment phase.

4. B. IEEE 802.1X must use RADIUS or DIAMETER. Note: DIAMETER is out of scope of the exam blueprint.

5. B. Supplicants have the option to not authenticate the server certificate. Additionally, EAP-FAST offers the ability to use PAC files instead of certificates for tunnel establishment.

6. D. Protected access credentials (PACs) are a type of "secure cookie" that can be used instead of or in addition to a certificate.

7. B. MSCHAPv2 may be used for user authentication against LDAP, but not machine authentication.

8. A. The actual tunnel mechanism is unrelated to the ability to do a machine authentication. The requirement is simply that it must be EAP-MSCHAPv2 for the authentication method.

9. C. The three main components of 802.X are the authentication server, supplicant, and authenticator.

10. A. A tunneled EAP type is able to use native EAP types as its inner method.

Chapter 5

1. B. The available options for nonauthenticating endpoints are MAC Authentication Bypass (MAB) and Web Authentication (WebAuth).

2. B. With nonauthenticating endpoints, the authenticator (a switch, for example) can 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. With MAB, it is not recommended to use VLAN assignment, but MAB authorizations do not limit

https://t.me/learningnets

the authorization results.

4. B. With CWA, the authenticator only recognizes a MAB, and ISE maintains administrative control of the entire session and the tracking of the user's credentials.

5. C. With LWA, the web portal is hosted within the authenticator, the end user enters her credentials into the web portal and the authenticator sends those credentials inside a RADIUS Access-Request message to the authentication server. The authentication server returns the Access-Accept or Access-Reject along with the full response.

6. A. The three main non-802.1X authentication use cases are WebAuth (CWA and LWA), MAB, and Remote Access VPN (RA VPN).

7. B. When changing a VLAN assigned to an endpoint, that endpoint must know (somehow) to renew the DHCP address. 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.

8. C. 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, a URL redirection, and often a dACL that 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 a reauthentication. The reauthentication maintains the same session ID, and ISE is able to tie the user's credentials to the MAB request, sending the final authorization results for the end user.

9. B. There are many different "headless" endpoints in an organization, such as IP phones, IP cameras, printers, badge readers, IV pumps, medical imaging systems, and so many more. Some do not have supplicants. For those that do, the enablement and configuration of supplicants on the disparate endpoints could be overcomplicated or operationally difficult for the company. Many of the 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 those headless devices.

10. A. Web authentication is used for any interactive login when a supplicant is not available, and sometimes it is even used as second authentication after 802.1X.

Chapter 6

1. B. 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. In a situation where a CoA is warranted, an authentication server can perform a number of actions: No COA (that is, do nothing), Port Bounce (i.e. shut/no shut the relevant access "port"), or Reauth (that is, force the endpoint to reauthenticate in cases where multiple endpoints are present on a single access medium.). Supported CoA actions can vary depending on the selected authentication server.

3. C. Those devices that don't have an 802.1X supplicant available use MAC Authentication Bypass. Without the supplicant, the device does not recognize EAP messages and, therefore, EAP authentication techniques are NOT available. In the absence of EAP, the device will use its MAC address as its unique identifier to authenticate to the network.

4. D. The first three octets of a MAC address are the organizationally unique identifier (OUI). This OUI indicates which vendor manufactured the device. This can be useful, at times, to also indicate the function of the device—for instance, an IP phone or printer.

5. A. Often, the "dumb" network devices are those that lack 802.1X supplicants. From this list, a printer would be the most common device to lack 802.1X support. Other examples would include an IP phone, IP cameras, and badge readers, amongst others.

6. C. Prior to MAB, there wasn't a mechanism to authenticate a device based strictly on the device's

https://t.me/learningnets

MAC address. For this reason, the switchport would 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 into that switchport, to have unfettered access to the network.

7. A, B, C. Via posture checking, the endpoint can be checked for file conditions (existence, date, and/or version), registry conditions (whether a registry entry is or is not present), and service condition (whether a service is or is not running), so all of the above are correct. posture checking also can confirm the presence, absence, and status of antivirus and antispyware programs running on the endpoint.

8. D. When using posture assessment as a condition for authorization policy, the values of the PostureStatus condition can be Compliant, NonCompliant, or Unknown. Different levels of network access and/or remediation can be authorized based on the status of this variable.

9. B. To remediate a noncompliant endpoint, a redirect ACL must be defined on the switch and the redirect destination must be set to remediation portal.

10. D. A mobile device manager is a software system or service that provides advanced posture assessment for mobile endpoints. The MDM can determine the type of mobile device, the level of operating system on the endpoint, the presence/absence of PIN lock, and whether encryption is being used, as well as provide remote security services such as device lock and secure wipe. Depending on the MDM vendor chosen, additional services also might be available.

Chapter 7

1. C. Cisco Identity Services Engine 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, G. Cisco ISE has four personas. These personas are Administration, Monitoring and Troubleshooting, Policy Services Node, and Inline Posture Node. Each of these personas is required at least once in an ISE deployment, with the exception of the Inline Posture Node. The function of each persona is discussed within the chapter.

3. A. Cisco ISE's Policy Administration Node (PAN) persona is the instance of Cisco ISE where policy configuration actually happens. This persona will then distribute this policy to all other nodes.

4. D. The Cisco ISE Monitoring and Troubleshooting (MnT) Node 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 will be the responsibility of the Monitoring and Troubleshooting Node.

5. C. The Cisco ISE Policy Service Node (PSN) persona provides policy decision-making. As a user or an endpoint attempts to authenticate to the network, the PSN will be responsible for making the AAA decisions based on the policy as downloaded from the Cisco ISE Policy Administration Node (PAN).

6. A. The Cisco ISE Inline Posture Node is responsible for enforcing access policies and handling the CoA requests for those network access devices that cannot process CoA requests. After an endpoint is authenticated, the Inline Posture Node will ensure that the posture of the endpoint adheres to the network security policy.

7. E. 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-3415 or SNS-3495 physical appliance. Also, you should reserve 100% of these resources to ensure that other virtualized network functions do not starve the ISE of the resources.

8. D. In a single-node deployment of ISE, all ISE personas (PAN, MNT, and PSN) reside on a single appliance. In this deployment, there are no options for redundancy. For instance, if the PSN persona fails, or if the physical appliance fails, RADIUS authentications and authorizations will fail until the

https://t.me/learningnets

issue can be resolved.

9. F. In a four-node ISE deployment, the PAN and MNT personas are combined on two of the appliances, with each acting as primary on one appliance and secondary on the other appliance. On the remaining two appliances, only the PSN persona is configured.

10. F. In a fully distributed ISE deployment, the ISE PAN and MNT personas each reside on a separate appliance (or a separate pair of appliances if redundancy is required). Each of the PAN and MNT appliances will be an SNS-3495 appliance (or equivalent virtual appliance). With these PAN and MNT functions distributed, up to 40 PSNs can be deployed. For each SNS-3415 PSN deployed, up to 5,000 endpoints can be supported. For each SNS-3495 PSN deployed, up to 20,000 endpoints can be supported. A limitation on the PAN/MNT nodes, however, will allow only up to 250,000 endpoints to be supported in a single fully distributed ISE 1.2 deployment.

Chapter 8

1. B. The Cisco ISE GUI is available via an Adobe Flash-capable web-browser. As of Cisco ISE 1.2, the two supported browsers are Mozilla Firefox and Microsoft Internet Explorer.

2. D. The best way to ensure a secure connection is by encrypting the communications between the ISE and the device being used for the administrative portal. If HTTP were to be 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, encrypting all traffic between the administrative device and ISE, and ensuring 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.

3. B. To establish the initial, secure connection with ISE, ISE will generate a self-signed certificate. Because a trusted certificate authority, either a local CA or a third-party, public CA, has not signed it, the certificate can cause a security warning within 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 this certificate within the web browser to prevent 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) onto ISE. This, too, will prevent these security warnings in the future.

4. A. The Operations tab of Cisco ISE allows an administrator to monitor, report, and troubleshoot active authentication and authorization sessions.

5. C. The Policy tab of the Cisco ISE GUI allows an administrator to configure authentication, authorization, profiling, posture, client provisioning, and security group access—amongst others. web portals, however, are configured under the Administration tab.

6. G. The Administration tab of Cisco ISE can be used to configure all "setup"-type functions of ISE. These functions are those that are often set up one time and rarely modified thereafter. In this case, certificates and network devices are two items that are configured under the Administration tab and are rarely modified after their initial configurations.

7. B, C, E. When adding a new network access device to Cisco ISE, you must provide a device name and a device IP address. If you intend to use a Cisco ISE RADIUS server for authentication and authorization (the usual purpose of Cisco ISE in a network deployment), you will 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.

8. B. Authentication is the process by which ISE identifies the endpoint or the user of the endpoint as it connects to the network. The authentication policy is used for this purpose.

9. C, E. When an endpoint attempts to access the network, it automatically sends a number of different

https://t.me/learningnets

packets onto the network—"normal" communication for a networked device. The information contained within these packets can often be leveraged by ISE to determine the type of device (profiling the device) that is sending the information. The MAC address of the endpoint—either learned via EAP or via MAC Authentication Bypass 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 as a vehicle for profiling an endpoint.

10. A. During the client provisioning process, the necessary credentials and configurations are deployed to the endpoint, allowing the endpoint to automatically join the network on the next attempt with little or no interaction from the user.

Chapter 9

1. C. The permissions needed to join ISE to AD are Search Active Directory (to see whether ISE machine account already exists), Add workstation to domain (if it does not already exist), and Set attributes on the new machine account (OS type and version—optional).
2. B. The `show application status ise` command lists all the ISE processes and their statuses.
3. C. In both HTTPS and TLS connections, certificates are used to authenticate the server to client and act as the basis for the encrypted transport between the client and the server.
4. B. Only the CSR is submitted to the signing CA. The private key should be backed up but never given out to a third party.
5. B. Settings such as RADIUS shared secret keys and SNMP strings can be set only on a per-NAD basis.
6. A. Use NDG to build different policy sets for the staged deployment of ISE.
7. B. False. It is a best practice to use endpoint identity groups only for MAC address management instead of profiles.
8. A. ISE 1.2 is capable of joining only a single AD domain.
9. 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 principle identity for the authorization process.
10. A. The Network Time Protocol is critical for all network interactions that require time-sensitive interactions, including the interaction between the Cisco ISE and the Active Directory. Endpoint identity certificates also require an NTP synchronized time on Cisco ISE.

Chapter 10

1. B. The RADIUS packet must have the `service-type` set to `Call-Check`. The `service-type` dictates the method of authentication. The `calling-station-id` field must be populated with the MAC address of the endpoint.
2. B. Only EAP-FAST and TEAP (RFC 7170) have EAP chaining capabilities as of the publishing of this book.
3. C. An authentication policy is meant to drop traffic that isn't allowed, meaning it is using an authentication protocol that is not configured, it will route authentication requests to the correct identity store to validate the identity, and "pass" successful authentications over to the authorization policy.
4. B. Only the Process Host Lookup check box must be select in the Allowed Protocols for Cisco MAB to work. Detecting another protocol as Host Lookup is only for non-Cisco network devices.

https://t.me/learningnets

5. A. Reusable conditions can be built on-the-fly while building the authentication policy, and they are saved as dictionary objects.
6. D. Create one sub-rule for each EAP type under the default 802.1X authentication rule that points to the appropriate identity store per rule.
7. D. The Called-Station-ID attribute is used to match the source SSID.
8. A. The Calling-Station-ID attribute contains the MAC address of the endpoint.
9. C. The continue option is used to send an authentication to the authorization policy even if the authentication was not successful.
10. A. The Drop option for an authentication rule will allow ISE to act as if it were not "alive" so the network device will no longer send authentication requests to that ISE server.

Chapter 11

1. D. An authorization profile is the required authorization result that is made up of multiple RADIUS attributes. These RADIUS results will affect the ultimate security policy deployed to the NAD on behalf of the endpoint.
2. 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.
3. D. DACL-Name, Web-Redirection, Local WebAuth, Auto Smart Port. These common tasks, as well as the others, are the most often used RADIUS AVPs that will be sent to the NAD for secure policy enforcement of the endpoint.
4. A. An authorization policy contains authorization rules. Each rule will have at least one authorization profile.
5. A. True. Condition attributes can be saved into a library for future use and improved readability.
6. B. 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. Simple conditions contain only one attribute. Compound conditions contain multiple attributes along with an operator such as AND or OR.
8. A. A compound condition can contain a mixture of simple conditions (which are saved dictionary attributes) and raw attributes themselves.
9. D. To provide very specific permissions to any authorization, providing defense-in-depth while meeting the goals of the company's security policy. A printer, for example, should not have unfettered access to the network; instead it should have only what is needed (such as reaching the print servers).
10. C. Cisco dACLs are created entirely on the RADIUS server, and the full ACL is sent down to the network device within RADIUS AV pairs, while non-Cisco network devices must create the ACL on the individual local network device. This allows the Cisco admin to create and maintain the access lists in a central place and have any changes applied nearly instantly.

Chapter 12

1. C. 802.1X requires global-level configuration for servers, enabling 802.1X on the system itself, configuring change of authorization, and enabling VSAs among others. Additionally, each interface that will be performing authentication will require interface-level commands.
2. B, D. 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. Additionally, RADIUS NAC must be enabled on the advanced tab of the WLAN configuration.
3. B. Cisco switches can be configured to send syslog to the MNT node, where the data will be correlated as part of the authentication reports. However, this should be configured only when

https://t.me/learningnets

performing active troubleshooting or during an initial pilot/PoC.

4. A. The switch will send periodic test authentication messages to the RADIUS server (Cisco ISE). It is looking for a RADIUS response from the server, either an Access-Accept or Access-Reject will suffice. The username and password used by the automated test must exist in the configuration.
5. B. Switch interfaces must be configured as Layer-2 access ports to run 802.1X (switchport).
6. C. Flex-Auth allows a network administrator to set an authentication order and priority on the switchport, thereby allowing the port to attempt 802.1X, MAC authentication bypass, and then WebAuth in order. All of 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. Multi-Host mode is not commonly used but is still a valid option. Much like Multi-Auth mode, Multi-Host 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 will be allowed onto the network using the first successful authentication. It's an "authenticate one, allow the rest" type of model.
8. A. The `authentication port-control auto` command will enable authentication on the port and allow the authorization result to be sent from the RADIUS server. Short answer: "Turn authentication on!"
9. C. 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.
10. D. The command will show that the authentications are being attempted, which are successful, which authorization results have been assigned, and much more. Some of the information that is quickly provided by this command output includes the endpoint's MAC address, the authentication method used, any assigned redirect URL, Access Control Lists, and other RADIUS AVPs that are provided via the authentication and authorization process.

Chapter 13

1. D. The Cisco switch will need the https server enabled to redirect https traffic. Before that service can be enabled, the switch needs a certificate. One of the prerequisites is a hostname and domain name, providing the switch a fully qualified domain name (FQDN). This FQDN will become the Subject Name of the self-signed certificate.
2. B. The traffic filtering ACL can be downloaded from ISE as a dACL, but the redirection ACL must preexist on the switch and is called by reference using a RADIUS AV-Pair. The AirespaceOS-based Cisco WLCs support only locally configured ACLs; therefore, all ACLs must be called by reference (also named ACLs).
3. C. RADIUS NAC is a critical setting for the WLAN that enables URL redirection and the pre-RUN states. Without this setting, CWA is not possible.
4. B. CWA is controlled by the Authorization Policy. Even an unknown MAC address needs to "continue" out of the Authentication Policy, so the appropriate response can be sent to the NAD, including the URL redirection to the portal.
5. B, D. The first rule should match if no more specific authorization rule is used and should redirect the user to the CWA portal. The second rule types should exist above the redirection rule and allow access to the user after she has successfully authenticated to the CWA portal. The authorization policy rules read like an ACL—from top down, whereby the first matched rule is applied.
6. A. DRW is an older method but uses a base license only. It does not provide a portal for the end user to manage his endpoints. When the end user accepts the AUP, the device's MAC address is automatically added to the configured Endpoint Identity Group.
7. A. The same URL-Redirect and URL-Redirect-ACL AV pairs are sent to the Cisco NADs regardless

https://t.me/learningnets

of the redirection type. The URL will be different for each portal type. When building the authorization profile, the common tasks area will provide a drop-down to select the type of URL redirection being used and to change the URL accordingly.

8. D. The **show authentication sessions interface [interface-name]** is like the Swiss Army knife of **show** commands for authentications. With the output, you see the MAC address, IP address, dACL (listed as an ACS ACL), URL-redirect ACL, and URL to which the end user is being redirected.

9. B. Cisco ISE has a phenomenally useful tool built in to it, commonly called Live Log. Live Log provides a near real-time view of all incoming authentications, change of authorizations (CoAs), and more.

10. A. The CoA is a key function. Specifically, it is a CoA-Reauth and causes the 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 14

1. B. When a guest connects to the network, they are given a web-redirect authorization policy. This web redirect will intercept any attempts to browse the Internet, forcing the guest user to a webpage where they will authenticate—that is, WebAuth.

2. C. The sponsor and guest portals can run on any PSN that has session services running.

3. B. Currently, the ISE guest portals can run only on those ports between 8000 and 8999.

4. A, D, F. The three default sponsor groups on ISE are SponsorAllAccounts, SponsorGroupGrpAccounts, and SponsorGroupOwnAccounts.

5. F. To use Active Directory group membership as the source of authentication and authorization for sponsors, ISE must first be associated to the domain. Furthermore, the AD identity store must also be a part of the identity source sequence that is in use for the sponsor portal. If you choose, you can provide a differentiated level of guest account creation based on the AD group membership as will be demonstrated in this chapter.

6. E. The Operations tab of the portal configuration page allows a network administrator to define the security policy for the portal. This page outlines how often the guest will be prompted to accept the Acceptable Use Policy, whether a guest can or must change their given password, whether the guest can perform device registration, or whether a user can create their own guest account. A few additional options are also available on the portal configuration page.

7. A. Under the sponsor group, the three settings that are configurable are the Authorization Levels, Guest Roles, and Time Profiles. With Authorization levels, the network administrator can configure which functions a sponsor user can configure for his guest. The Guest Roles option allows the sponsor to create guest users for specific Guest Roles—possibly allowing a differentiated level of access for each role. The final option, Time Profiles, defines the length of time for the guest accounts that can be created by the sponsor.

8. C. From the sponsor portal, when you are creating guest accounts, you have three options—Individual, Import, and Random. The Individual option creates a single guest user account, Import allows you to create multiple accounts using a spreadsheet template, and Random allows you to create a number of random guest accounts. The level of access and the length of the account also are configurable.

9. C. To trigger the WebAuth policy on Cisco ISE, the NAD must be using the MAB process. This MAB process, or RADIUS Service-Type of Call Check, is indicated by the security policy of MAC Filtering on the WLC. RADIUS NAC must also be configured as the NAC State on the Advanced tab of the SSID configuration.

10. D. The correct command to verify the level of access given to a guest user on a Cisco switch is **show authentication sessions interface <if_name> details.** This output will provide you with

https://t.me/learningnets

any ACLs or URL Redirects that have been deployed to the device from ISE.

Chapter 15

1. A. Profiler is enabled by default on all policy service nodes and standalone nodes. However, not a single probe is enabled by default in ISE 1.2.
2. A, B, D. There is no such thing as an EndPointProfile attribute. Although OS-Scan is used as a condition to determine the endpoint's profile, it cannot be used directly in an authorization policy. The authorization policy can use identity groups (which contain a list of MAC addresses), EndPoint Policy attribute (which is the actual endpoint profile), and logical profiles (a group of profiles).
3. E. The SNMPQUERY probe will periodically query all the NADs configured with SNMP strings, but it is also a reactive probe. The SNMPQUERY probe will reactively query a NAD when the RADIUS probe receives an accounting START message or when an SNMP trap is received.
4. C. The three probes that exist in device sensor on Cisco switches are CDP, DHCP, and LLDP. Wireless controllers have two probes: DHCP and HTTP.
5. A. Cisco no longer includes profile updates within the ISE version updates or patches. All new profiles are included and downloaded as part of the Cisco Profiler Feed Service.
6. C. Profiling is all about the certainty value. Each profile has a minimum certainty value, and matching the conditions will increase the certainty value. A higher the certainty value of any profile means it will be assigned.
7. A. The Endpoints Drill-down tool is an excellent way to look into the profiled endpoints and verify that the profiling service is working.
8. B, D. HTTP user agent strings could be gleaned through SPAN monitoring and 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. ISE provides the ability for administrators to create their own custom profiles using any of the attributes available to the profiling engine.
10. C. Profiles are classified as Cisco provided, administratively modified, or administrator created. Only Cisco-provided profiles will be overwritten.

Chapter 16

1. B. A copy of the signing CA's public key must be stored at Administration > System > Certificates > Certificate Store, and it needs to have the Trust for Client Authentication option selected.
2. D. It's vital to understand that the Valid-From field is just as important as the Valid-To field. A certificate will be rejected if it is issued for a date and time after the current date and time. This is why NTP is so critical for PKI.
3. B. ISE supports checking both CRL and Online Certificate Status Protocol (OCSP). OCSP is the preferred method for scalability and security reasons.
4. A. ISE will only leverage the CRL distribution point configured within the trusted certificate store for that signing CA and ignore the field that is in the client's certificate.
5. A. 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 must decrypt 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. A certificate issued by Active Directory Certificate Services is still just an X.509 certificate. It will go through all the authentication validation of 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.

https://t.me/learningnets

7. B. Although 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. 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. Many certificate authorities have a website where they permit the downloading of their public certificate and even the full certificate chain. In this chapter you see an example of downloading the key from a Microsoft CA. Navigating to this web page and downloading the certificate is how an ISE admin can 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. As a best practice, always use Base-64 encoded files instead of DER-encoded files.
10. B. Although I'm flattered that you might want to call me to fix your problems, C is definitely not the correct answer. The first question you would be asked is: "What does it say for Failure Reason in the Authentication Details Report?" which is the correct answer: B. There is no report named Failed Authentications, and besides it would not exist in the root of "reports."

Chapter 17

1. B. One of the business issues with a BYOD model is walking an end user through the process of configuring his network supplicant to meet corporate policies. Onboarding is used to help an end user perform those actions himself, without requiring interaction from the IT department.
2. C. To maintain a seamless experience for the end user, a CoA-Reauth message is used. This keeps the endpoint connected to the network and simply causes the supplicant to send credentials again. At this point, it will be using 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 be 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. The software is hard-coded to deny guest users from entering the flow. There is no configuration possible to allow guest users to enter the provisioning process through the dual-SSID onboarding flows.
4. D. While both C and D could be viewed as correct answers, only D is technically accurate.
5. B. ISE will authenticate 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 or to leave the user in the redirected state to receive a message explaining that she must configure her device on her own or call her IT department for assistance.
6. B. Apple iOS does not use an app to perform the provisioning; instead it leverages the native Over the Air (OTA) provisioning built in to the OS to handle the certificate signing requests and downloading of a network profile.
7. B. The admin may manage endpoints from the Endpoints Identity section within the ISE administrative GUI. The MyDevices portal is designed for an individual to perform self-service of registered devices.
8. A. Live authentications log does not show any information about the registration or the NSA app. It does show all the authentications and the change of authorizations.
9. D. With the ISE 1.2 versions pertinent to this exam, both Windows and Mac are still using a Java applet that is downloaded from ISE's Client Provisioning Portal (CPP). 1.2 patch 11 and 1.3 versions of ISE will enable the use of a native .exe for Windows and a .dmg for Mac OSX, but that is out of scope of this exam blueprint and therefore out of scope for this book.
10. B. The Client Provisioning Policy determines which NAC agent, NSA Wizard, and Native Supplicant Profile to send to an endpoint. The policy is capable of using the operating system as one of many

https://t.me/learningnets

conditions to determine which result to provide an endpoint.

Chapter 18

1. C. A security group tag (SGT) is a 16-bit value that ISE assigns to the user's or endpoint's session upon login. The SGT can represent the context of the user and device and can be carried in the Layer-2 frame or communicated through SXP. The SGT is assigned at ingress and enforced upon egress.
2. B. SGTs are considered an authorization result in the ISE administrative GUI. They are defined within the policy elements section of the GUI as an authorization result. They also can be defined from the **Policy > Security Group Access > Egress Policy** screens by clicking on **Configure > Create New Security Group**; however, that method was never discussed in the text of this chapter.
3. A, B, D. To use the SGT, the tag needs to be assigned (known as classification). This can happen dynamically and be downloaded as the result of an ISE authorization; they also can be assigned manually at the port level or even mapped to IP addresses and downloaded to SGT-capable devices.
4. A. Although that gear might not support the classification and transport natively, it might be capable of assigning different VLANs or IP addresses per authorization result. A distribution layer device may have the ability to map subnets and VLANs and assign all source IP addresses from the subnet or VLAN to a specific tag.
5. D. Cisco has developed a peering protocol (similar to BGP or LDP) to enable devices to communicate their database of IP-address-to-SGT mappings to one another. This peering protocol is called Security Group Exchange Protocol (SXP).
6. **A, C.** Every SXP peer session has a speaker and listener. A speaker sends the mappings of IP addresses to SGTs. The listener receives those updates and records them. A peer can be configured for both roles simultaneously and can have numerous peers.
7. A. Native tagging of SGTs includes the 16-bit tag as a portion of the Cisco Metadata field of the Layer-2 Ethernet frame. It also can be included as part of an IPSec link.
8. B. The tag 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 will not help.
9. A, C. SGTs can be enforced with security group ACLs, which are egress ACLs that use source and destination tags as the condition upon which to invoke the egress ACL. Additionally the ASA, ASR, and ISR can act as security group firewalls, using the source and/or destination tag as ACL conditions.
10. D. Uplink MACSec defines the encrypted connection between network infrastructure components, whereas downlink MACSec defines the encrypted connection between the access layer device and the endpoint. Although uplink and downlink MACSec use different keying mechanisms today, both are still using the same encryption algorithm of AES-GCM-128.

Chapter 19

1. B, C, G. The three major functional areas of the Posture Service are Client Provisioning, Posture Policy, and Authorization Policy. The first, Client Provisioning, is the process by which the NAC agent is installed on the endpoint. The second, Posture Policy, is the configuration of the Posture rules: what is compliant and what is not compliant within the security policy. The final functional area is Authorization Policy. After we have determined the compliance or noncompliance of the endpoint, what will the endpoint have access to.
2. D. The three possible posture outcomes following the initial connection to the network are Compliant, Noncompliant, and Unknown. Compliant implies that the endpoint fully adheres to the company's security policy as configured on 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

https://t.me/learningnets

the device and, therefore, the endpoint is unable to report its posture to ISE.

3. B. One benefit of the NAC web agent is that it does not require administrative privileges to install. Unfortunately, the web agent is lacking additional features that are standard in the persistent agent.

4. B. The Process Check posture condition is not supported on Macintosh operating systems.

5. D. The File condition for Posture can check the existence, date, and version of a file on the client. This can be very useful to determine if a particular endpoint is vulnerable to a new virus or if a specific software package is present on the endpoint. This feature is only supported on Windows PCs.

6. A. These Posture Elements can be updated manually or configured to update automatically on a fixed schedule.

7. D. The CoA process is used to force an endpoint to reauthorize following a change in status or following a change of posture compliance from the NAC agent.

8. C. When configuring the Client Provisioning Policy, a network administrator is responsible for defining what NAC agents or Network Supplicant Provisioning (NSP) client is getting pushed to what endpoints under which circumstances. The network administrator, besides specifying the elected NAC Agent and NSP client, can also specify the period of time between reassessments and whether or not an Acceptable Use Policy will be used.

9. A. Remediation is the process by which an endpoint that is not compliant with security policy can become compliant. This may include downloading the latest virus definitions, installing a service pack, or enabling a screen saver password.

10. D. The only remediation from this list that is available on a Macintosh OS X endpoint is Manual Antivirus Remediation. As an endpoint is found to be noncompliant due to a deviation in his antivirus signatures, the NAC agent will provide a link for the user to download the latest definition file. All other remediations provided in this list are not possible on the Macintosh NAC agent.

Chapter 20

1. C. Monitor Mode is a process, not just a command on a switch. The process is to enable authentication (with authentication open), see exactly what devices fail and which ones succeed, and correct the failed authentications before they cause any problems.

2. A. Low-Impact Mode uses authentication open, but adds security on top of the framework that was built in Monitor Mode. It uses a PACL on the switch port to permit critical traffic of certain endpoints, like 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 pre and post authentication.

3. D. By using a phased deployment approach, you are able to 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 often happen with 802.1X deployments.

4. B. **authentication open** will ignore RADIUS Access-Reject responses, but all other authorization results will be honored and enforced.

5. A. **authentication open** allows traffic to flow with our without an authentication. When an authorization result is sent back from the authentication server, the switch will ignore RADIUS Access-Reject responses, but all other authorization results will be honored and enforced.

6. B. Policy sets are groupings of authentication and authorization policies. The use of policy sets makes for a nice clean way to differentiate rules for each stage of the deployment.

7. D. Wireless LANs cannot have a mixture of authentication and nonauthentication. The WLAN must either be using Wi-Fi Protected Access (which facilitates the 802.1X authentication) or will be open; it cannot be both.

8. A. The NDG assignment of the NAD is used to determine which policy set ISE uses for the

https://t.me/learningnets

incoming authentications. To change the policy set being used, move the NAD from the Monitor Mode NDG to either the Low-Impact or Closed mode NDGs.

9. A. Wired clients do not get to pick their network; there is no SSID like there is for wireless. Therefore, all the various types of authentication mechanisms possible must work within a single port configuration. Without this, an admin 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.** 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 21

1. C. 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. When joining the node to the cube, you will specify the persona and whether it will be primary or secondary (Monitoring only).

3. D. The **show udi** CLI command and the GUI will provide the three required items: SPID, VPID, and serial number.

4. B. There is no automatic failover, but there is a manual promotion from the secondary's GUI.

5. D. There is no automatic failover, but the ISE nodes are configured to send logging to both primary and secondary MnT automatically. If one fails, the other is still receiving the logs.

6. C. Node groups are made up of Layer-2 adjacent (same VLAN) PSNs, where the PSNs maintain a heartbeat with each other. In the event that a PSN were to go down while a session was being authenticated, one of the other PSNs in the node group would send a CoA to the NAD so the endpoint could restart the session establishment with a new PSN.

7. B. Cisco ISE is commonly deployed with load balancers. There are caveats to pay attention to, such as not to use Source NAT (SNAT).

8. B. Patches are downloaded from cisco.com and applied to the PAN under Administration > System > Maintenance > Patch Management. The PAN will push the patch to the other nodes in the deployment.

9. D. 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. It is not configurable, and will patch all nodes in alphabetical order. The PAN is patched first, and will push the patch to all other nodes.

Chapter 22

1. D. The Evaluate Configuration Validator tool compares a switch configuration to a "template" configuration built in to ISE, and any differences between the configurations are pointed out.

2. C. The RADIUS Authentication Troubleshooting tool attempts to examine different aspects of a session and provide some additional details that might not have been available in the detailed authentication report, as well as provide some suggestions for items to check next.

3. B. Each ISE component can have its logging levels changed through the graphical user interface only.

4. B. The Live Sessions Log correlates activity related to the entire session, not just the raw entries related to a passed or failed authentication.

5. A. The Live Log displays events related to the raw syslog messages sent from the PSN to the MNT node, focused on passed or failed authentications.

https://t.me/learningnets

6. D. 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. The Suppress Anomalous Clients setting within Administration > System > Protocols > RADIUS is used to enable log de-duplication.
8. B. Cisco AnyConnect DART is the module used to collect all log files from the endpoint along with other important information, combining them all into a single Zip file for analysis by Cisco TAC.
9. C. Although a firewall can sometimes be a good place to troubleshoot why communication is not successful, the three main locations to troubleshoot network access are ISE, the endpoint, and the NAD.
10. **B. debug epm** is the go-to debug command for all activities related to URL-redirection, dACLs being applied, SGTs being assigned, and all other activity related to an authentication session advanced capabilities.

https://t.me/learningnets

Appendix B. Configuring the Microsoft CA for BYOD

With ISE 1.2 and previous versions, the vast majority of Cisco ISE deployments have been using the Microsoft Certificate Authority as their certificate authority (CA) for BYOD. Therefore, this appendix is included to provide some information on how to configure the Microsoft CA for use in the BYOD solution.

CA Requirements

For the Microsoft CA to provide all the functions necessary, it must meet the following requirements:

- Windows 2008 R2 Enterprise Server.
- Certificate Enrollment Web Service and Certificate Enrollment Policy Web Service must be installed.
- The Windows 2008 R2 Enterprise server must be joined to a domain.
- An Enterprise CA is required. The Certificate Enrollment Web Service cannot be configured to work with a standalone CA.

Other Useful Information

- The Certificate Enrollment Web Service can be configured to work with an Enterprise CA on the same server or a different server.
- The services can be installed on the same computer as the CA, Web Enrollment, Online Responder, and Network Device Enrollment Services (NDES) role services.
  - However, if you intend on using the NDES service for certificates issued to a Cisco IOS device (for example, Cisco CVO deployment), you will need to run Certificate Enrollment Web Services on a separate server from NDES.
  - This is due to an IIS incompatibility with SCEP-to-IOS routers when NDES and CWES are running on the same server. If you do run on separate servers, be sure to review the MS hotfixes listed in the next section.

Microsoft Hotfixes

- **http://support.microsoft.com/kb/2483564**—Renewal request for an SCEP certificate fails in Windows Server 2008 R2 if the certificate is managed by using NDES. This issue occurs because NDES does not support the GetCACaps operation.
- **http://support.microsoft.com/kb/2633200**—NDES does not submit certificate requests after the enterprise CA is restarted in Windows Server 2008 R2. You will see the message **The Network Device Enrollment Service cannot submit the certificate request (0x800706ba). The RPC server is unavailable** in the Event Viewer.

AD Account Roles

You will use two Active Directory accounts for the ongoing operation of the CA and NDES. Those two accounts will serve for the following roles:

1. SCEP Administrator—Used to install the NDES role service and must meet the following requirements:
   - Member of Domain Admins or Enterprise Admins
   - Enroll permissions on the Certificate Authority template (completed later)
2. SCEP Service Account—Used by the NDES application pool for the application to Run As. The account must meet the following requirements:
   - Member of the local IIS_IUSRS group

https://t.me/learningnets

- Request permission on the configured CA
- **Read** and **Enroll** permissions on configured device certificate templates

Configuration Steps

In this section you will install the CA with the critical services and then go back and add the remaining services in another configuration section.

Installing the CA

Begin by creating the service account:

Step 1. Add a new Active Directory user, such as SCEP_User.

Step 2. Ensure the user is added to the IIS_IUSRS local group.

Step 3. Install Active Directory Certificate Services.

Step 4. From Server Manager, select **Add Role**.

Step 5. Select **Active Directory Certificate Services**.

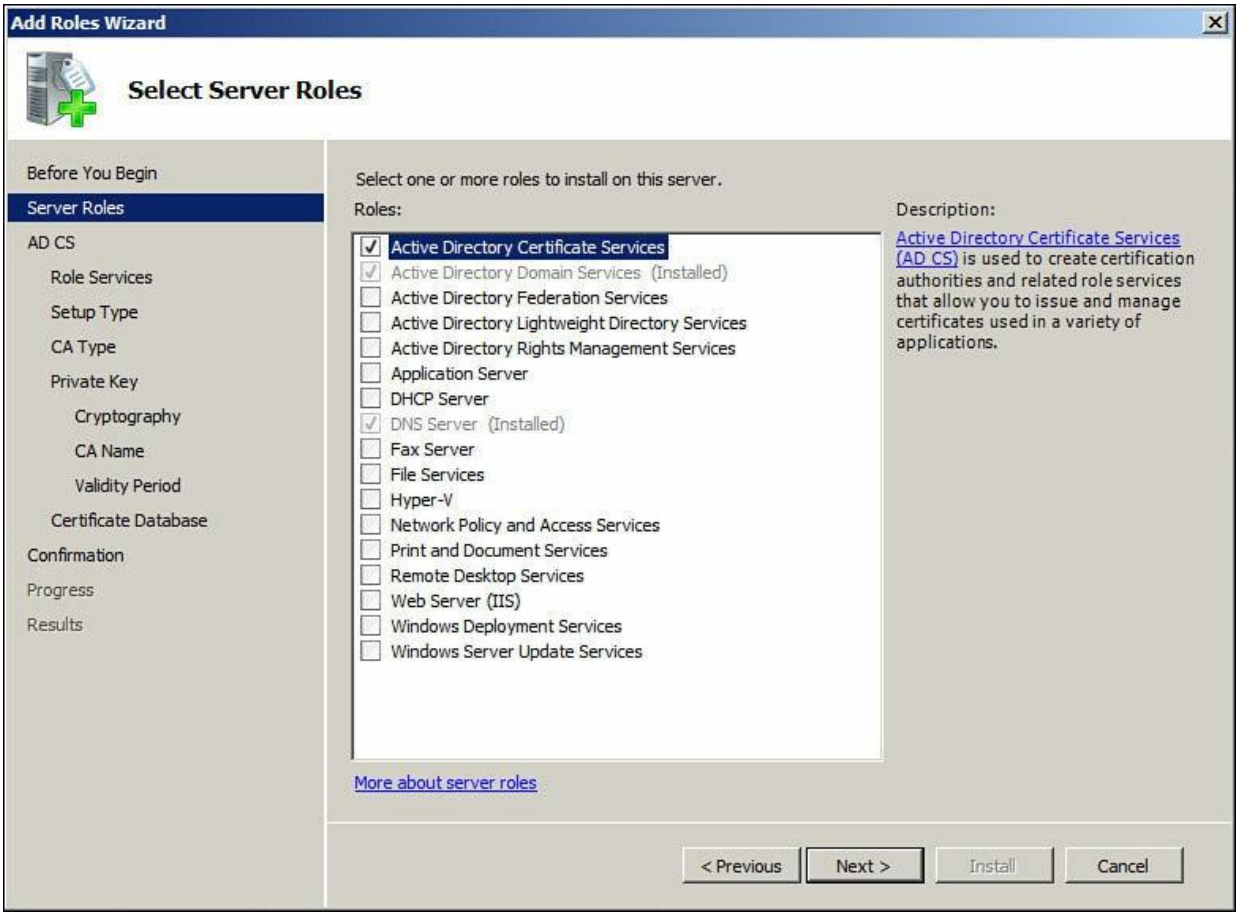


Figure B-1 Add Active Directory Certificate Services.

Step 6. Click **Next**.

Step 7. Select **Certification Authority, Certification Authority Web Enrollment, Online Responder, and Certificate Enrollment Policy Web Service.**

https://t.me/learningnets

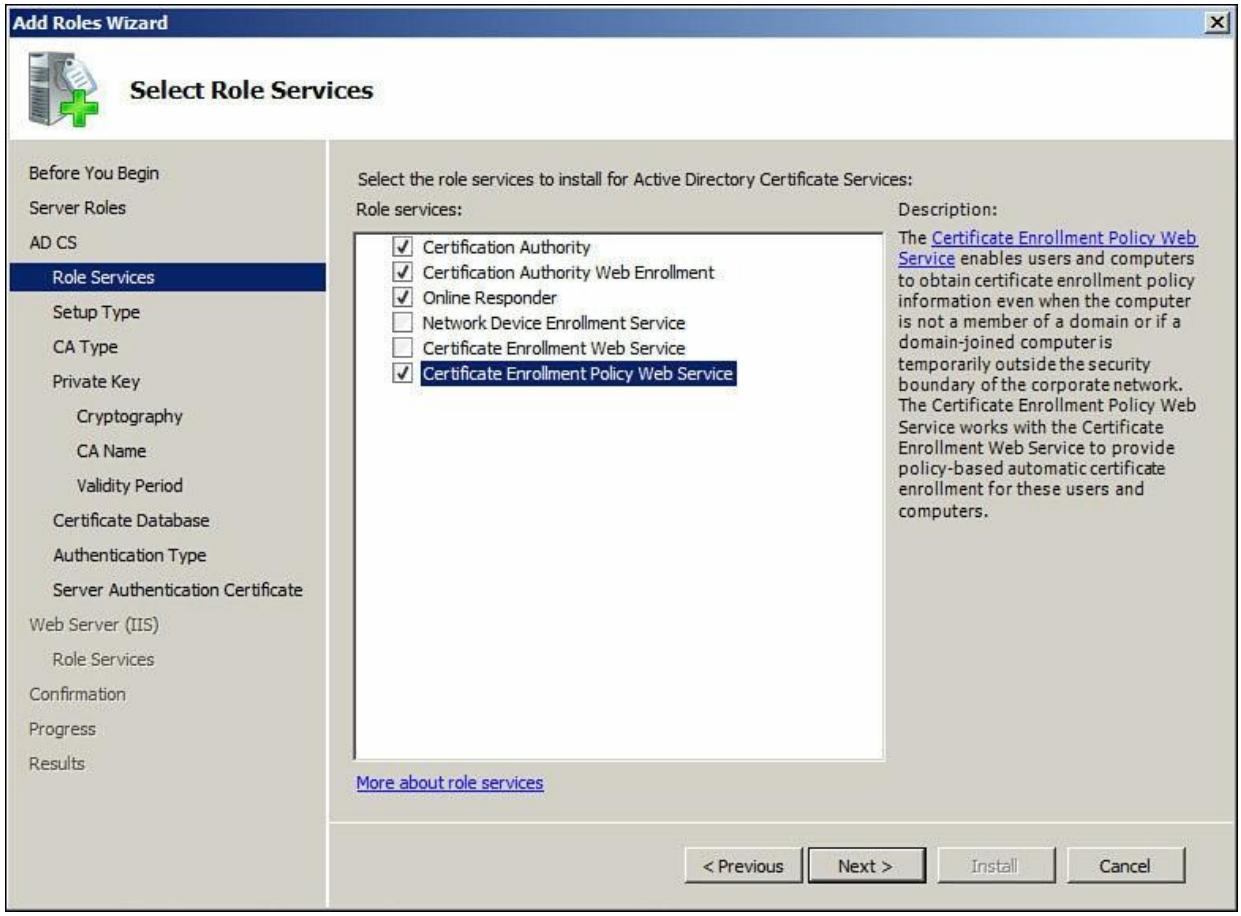


Figure B-2 Selected role services.

Step 8. Click **Next**.

Step 9. Select **Enterprise** to use Directory Services with the CA.

https://t.me/learningnets

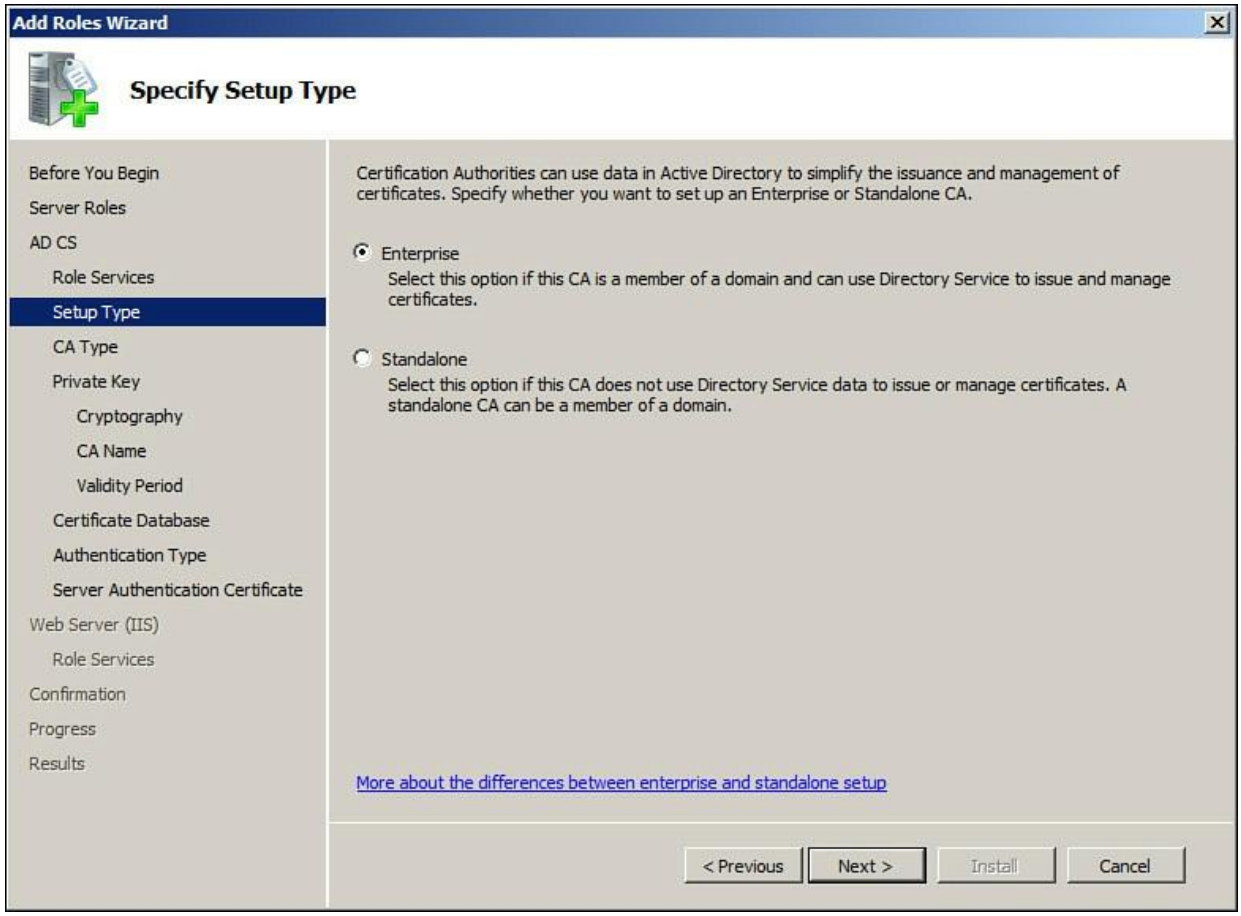


Figure B-3 Enterprise.

Step 10. Click **Next**.

Step 11. Specify the CA type. If this is a new CA, it should be root.

https://t.me/learningnets

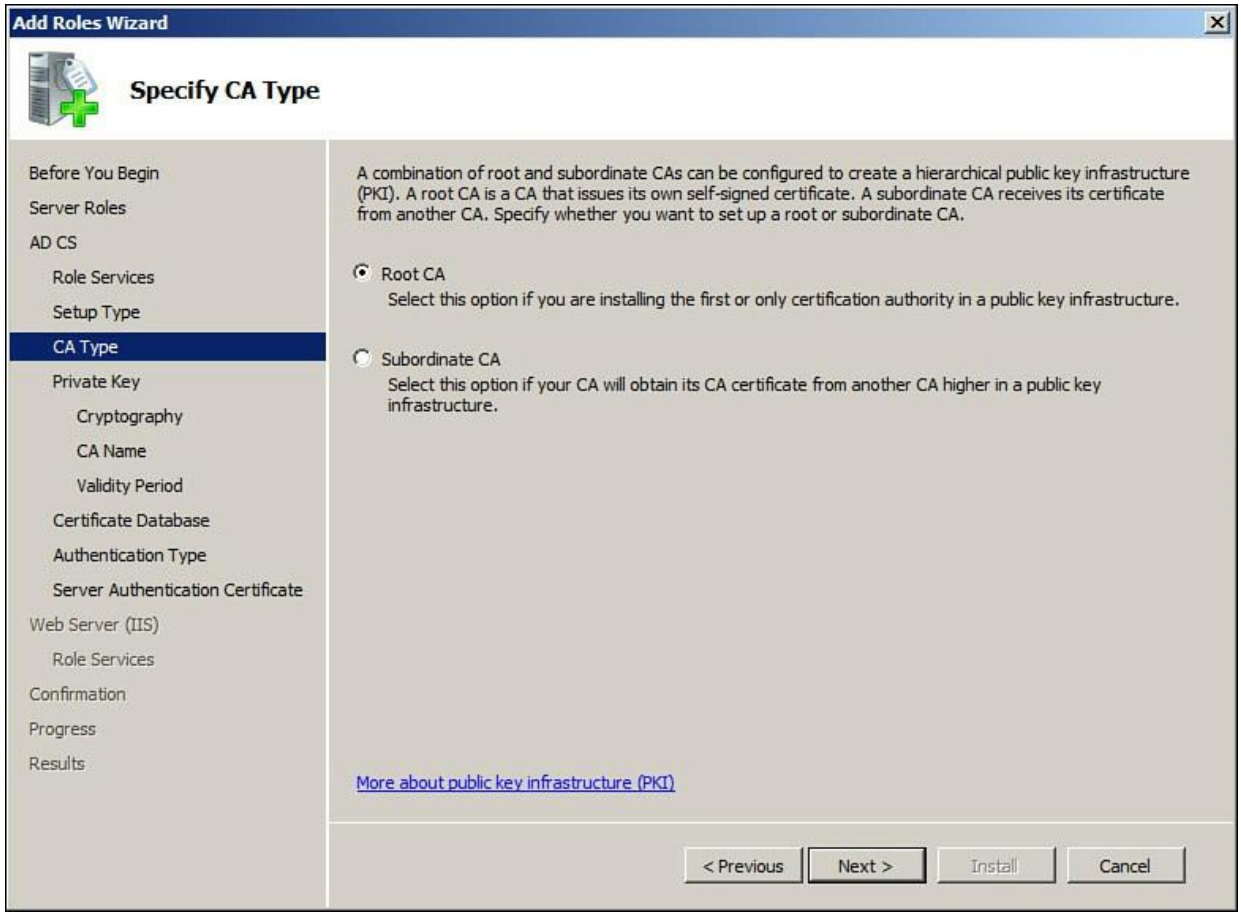


Figure B-4 Root CA type.

Step 12. Click **Next**.

Step 13. Create a new private key.

https://t.me/learningnets

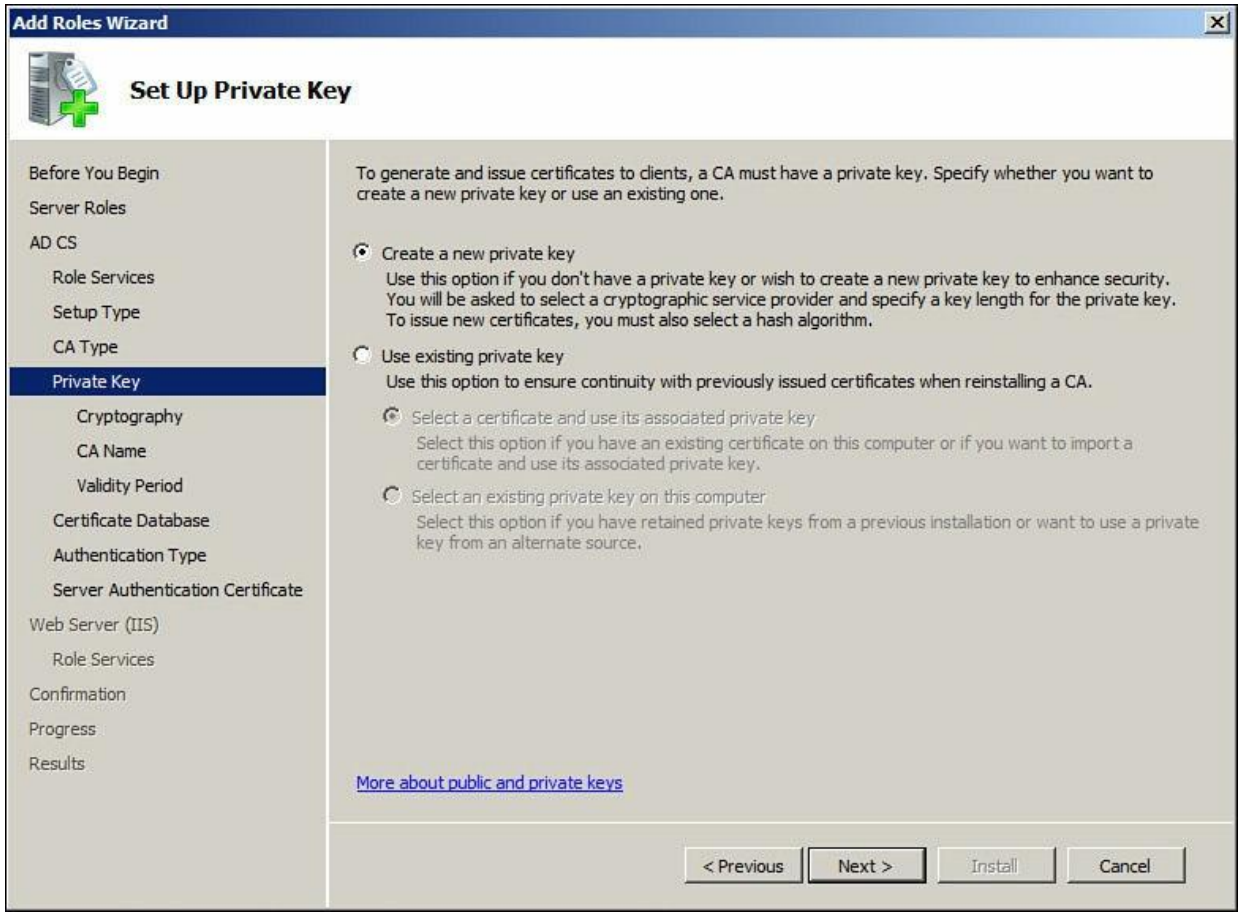


Figure B-5 Set up a private key.

Step 14. Click **Next**.

Step 15. Configure the cryptography for the CA.

Step 16. Most of the installations we have been involved with use the settings shown in Figure B-6.

https://t.me/learningnets

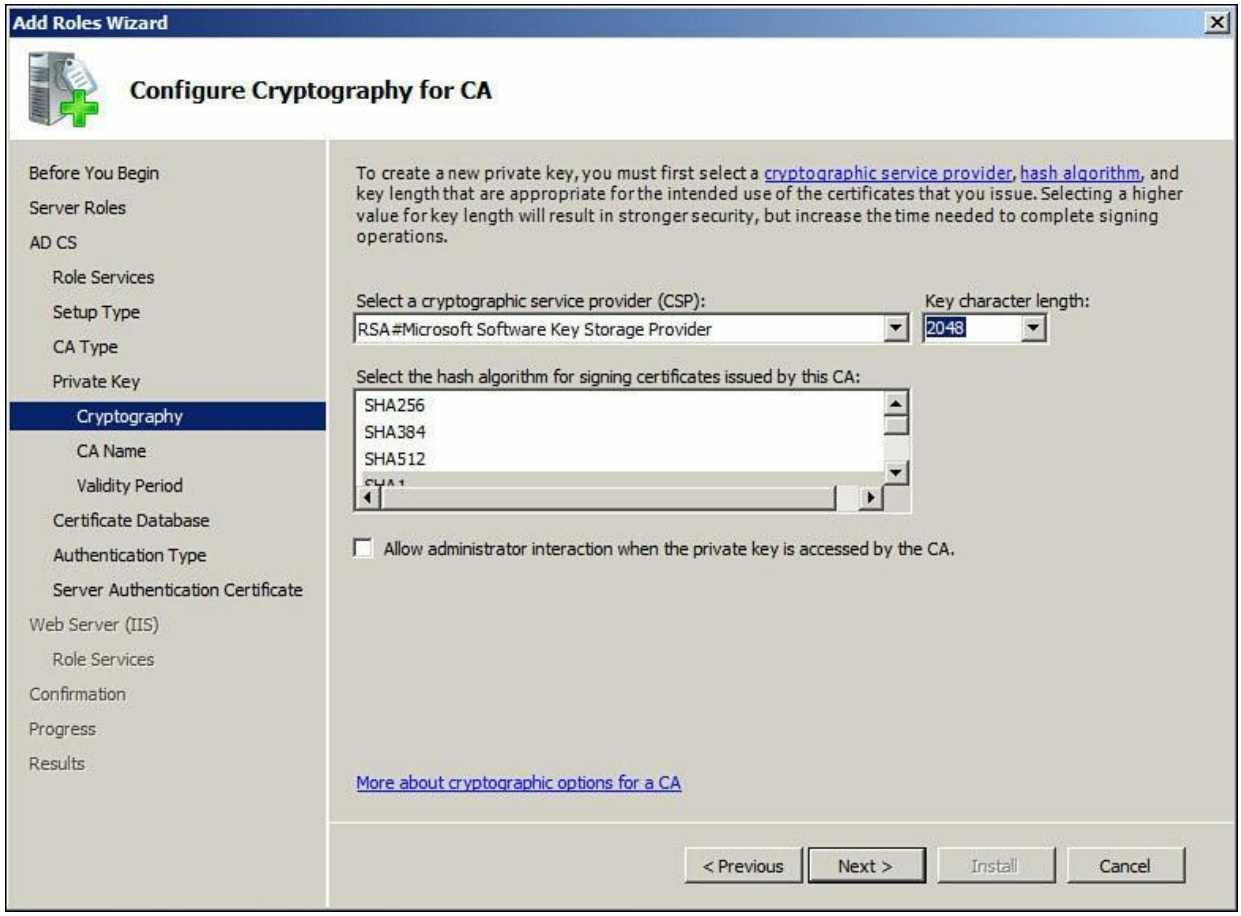


Figure B-6 Configure cryptography.

Step 17. Click **Next**.

Step 18. Configure the CA name.

https://t.me/learningnets

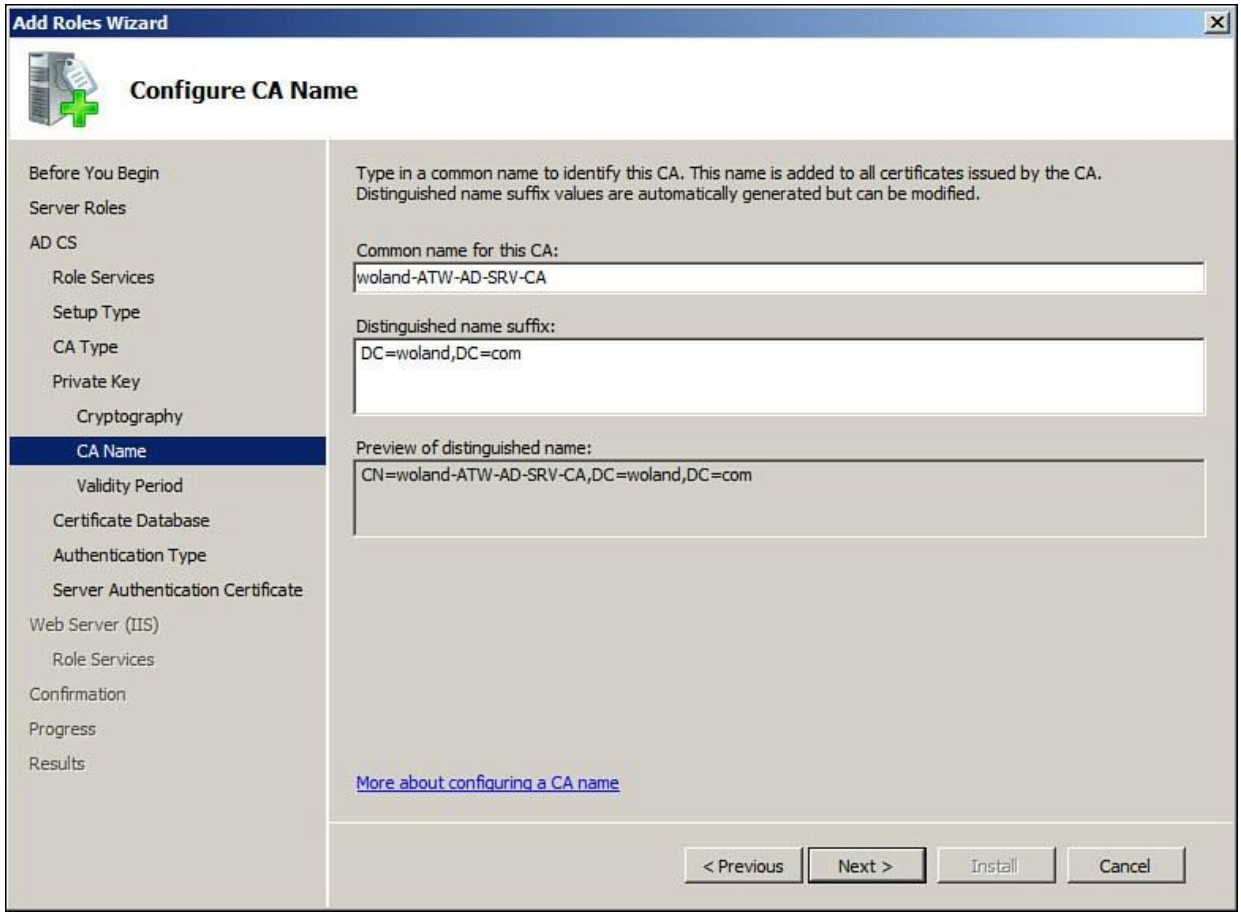


Figure B-7 Name the CA.

Step 19. Click **Next**.

Step 20. Set the validity period per your company's information security policy.

https://t.me/learningnets

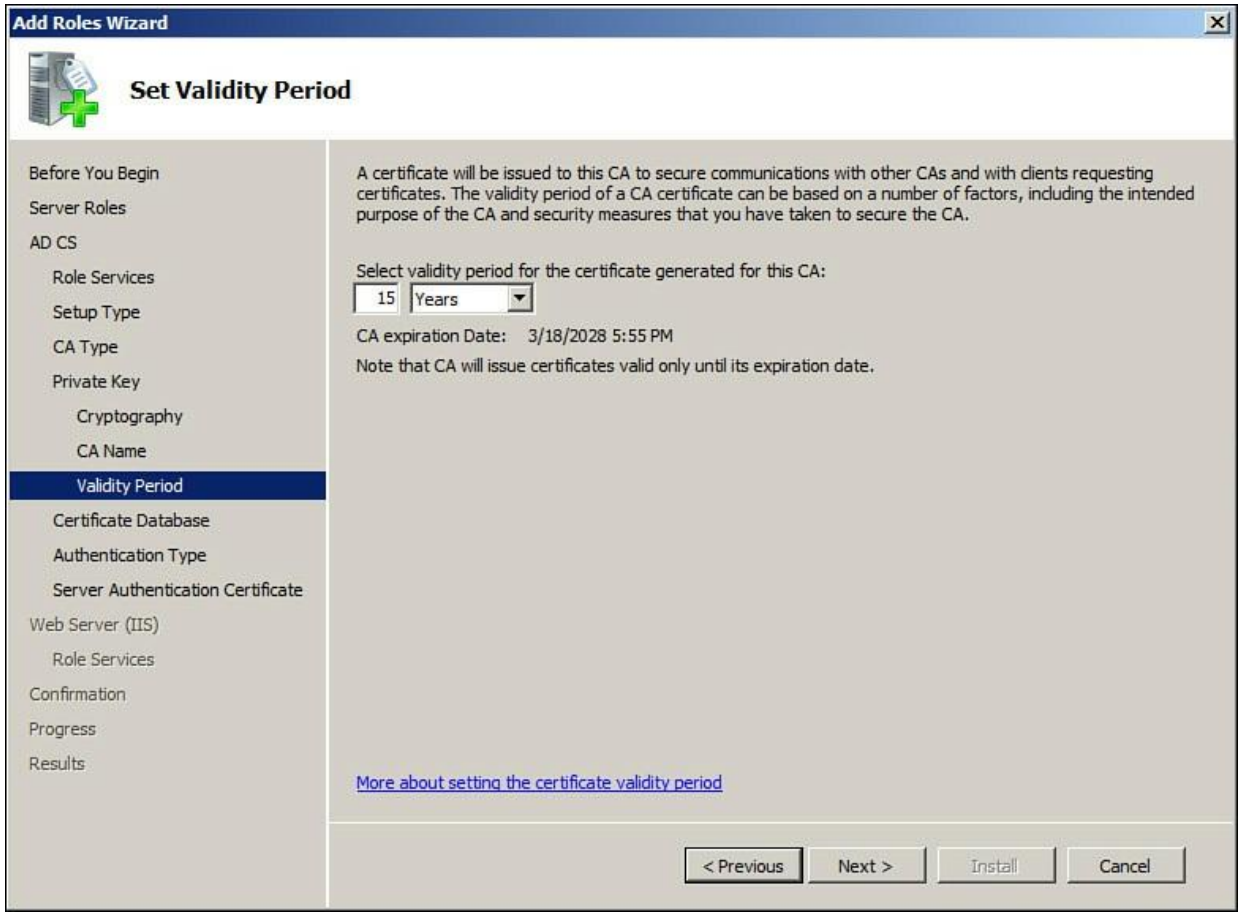


Figure B-8 Validity period.

Step 21. Click **Next**.

Step 22. Set the location of the CA database. The default is usually fine.

https://t.me/learningnets

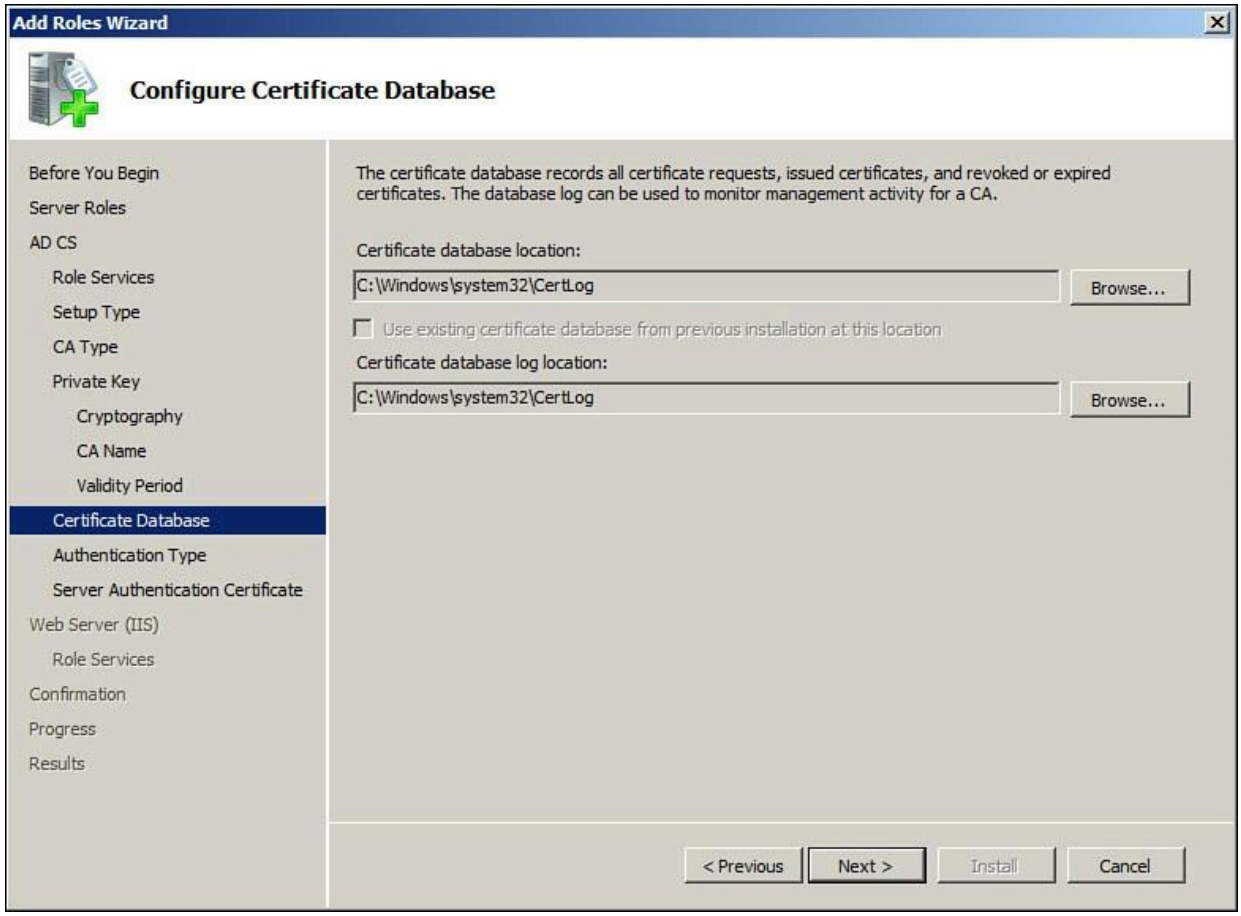


Figure B-9 Database location.

Step 23. Click **Next**.

Step 24. Set the authentication to use a username and password.

https://t.me/learningnets

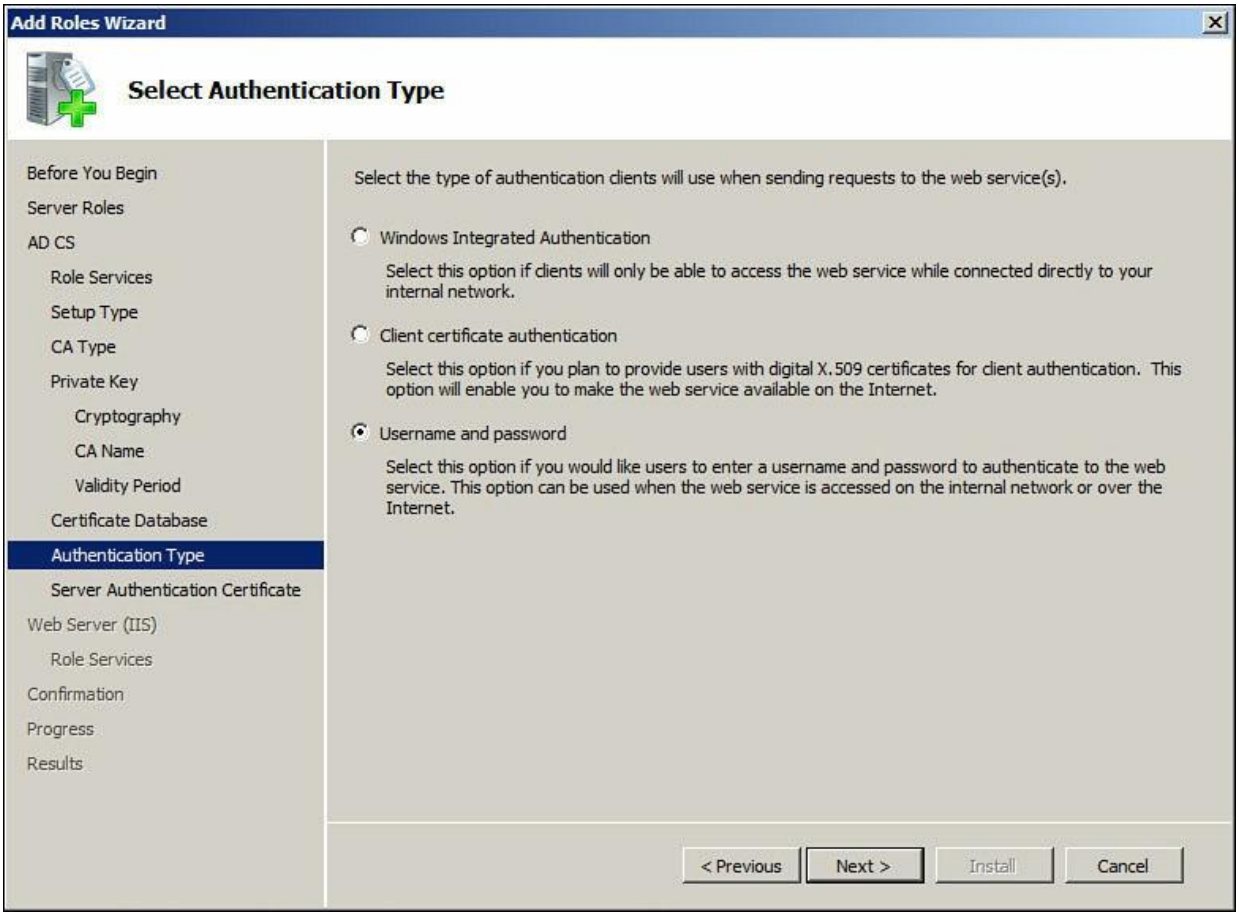


Figure B-10 Username and password.

Step 25. Click **Next**.

Step 26. Select the server certificate.

https://t.me/learningnets

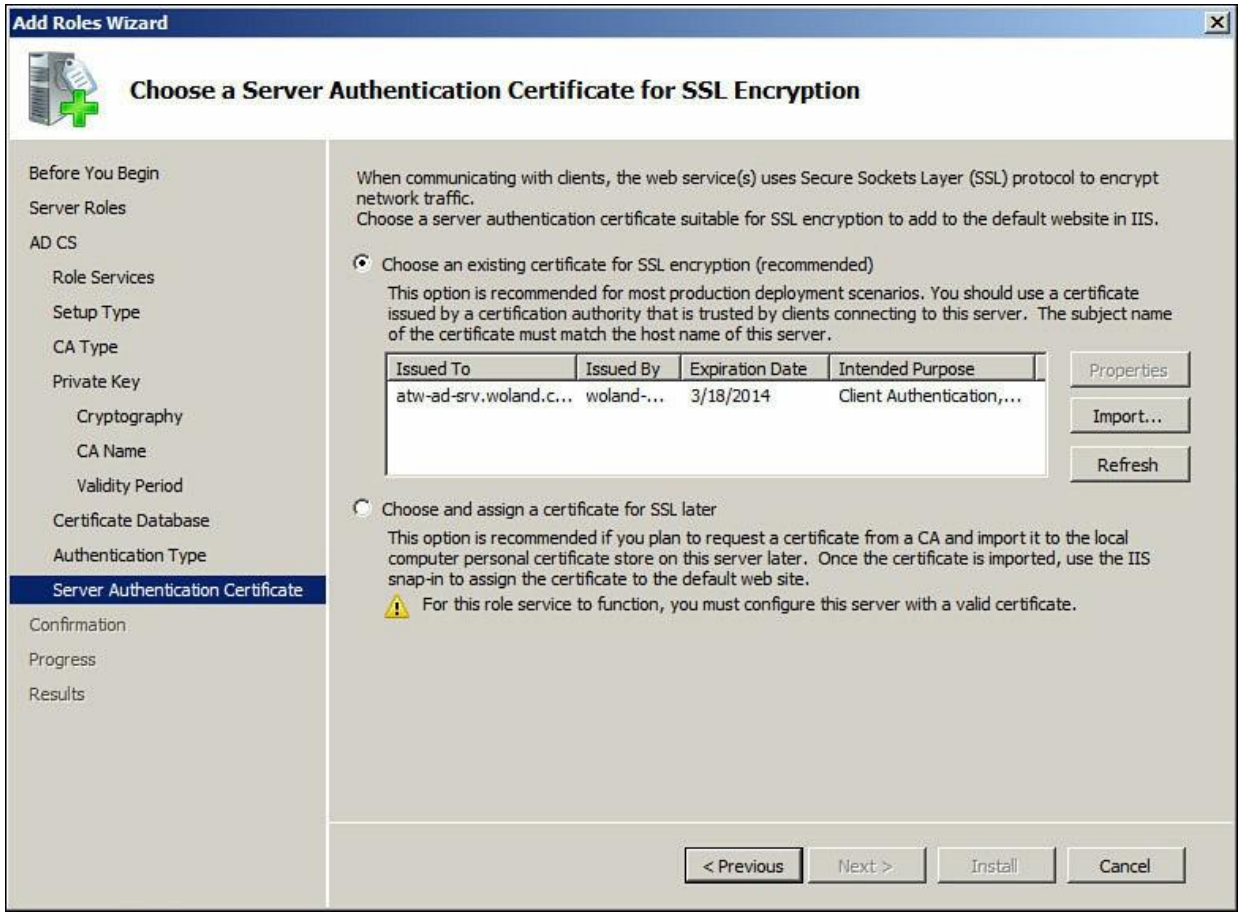


Figure B-11 Choose the server certificate.

Step 27. Click **Next**.

Step 28. Review your choices.

Step 29. Click **Install**.

Adding the Remaining Roles

Now that the CA is installed, you can go back in and add the remaining roles:

Step 1. From Server Manager, select **Add Role Services** for the CA.

https://t.me/learningnets

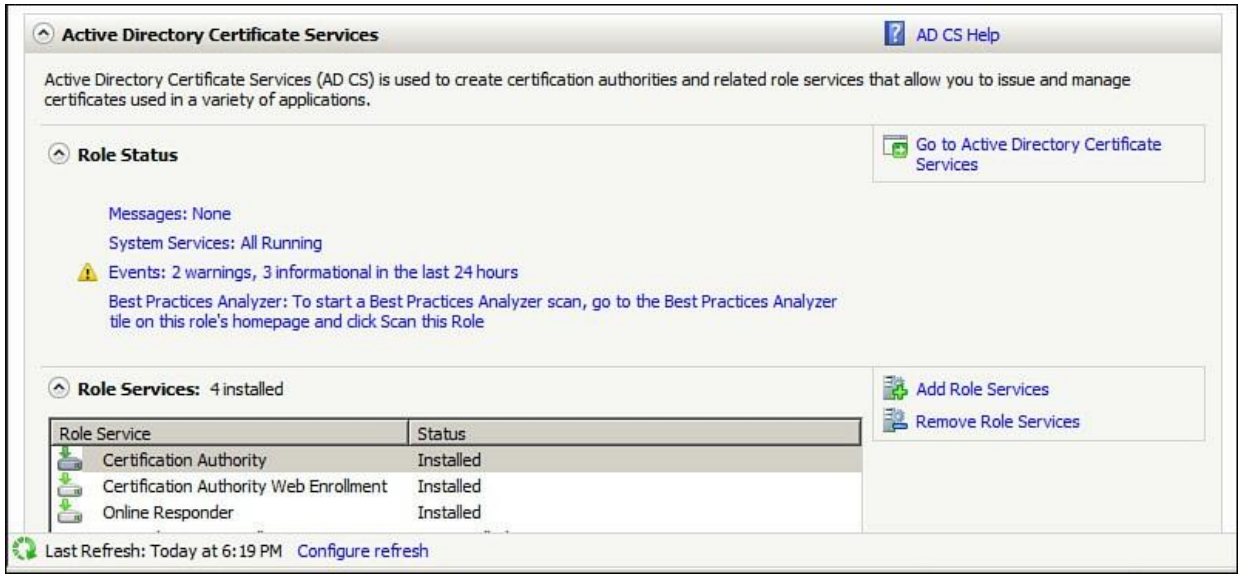


Figure B-12 Add a role service.

Step 2. Click **Next**.

Step 3. Select **Network Device Enrollment Service**, and then select **Certificate Enrollment Web Service**.

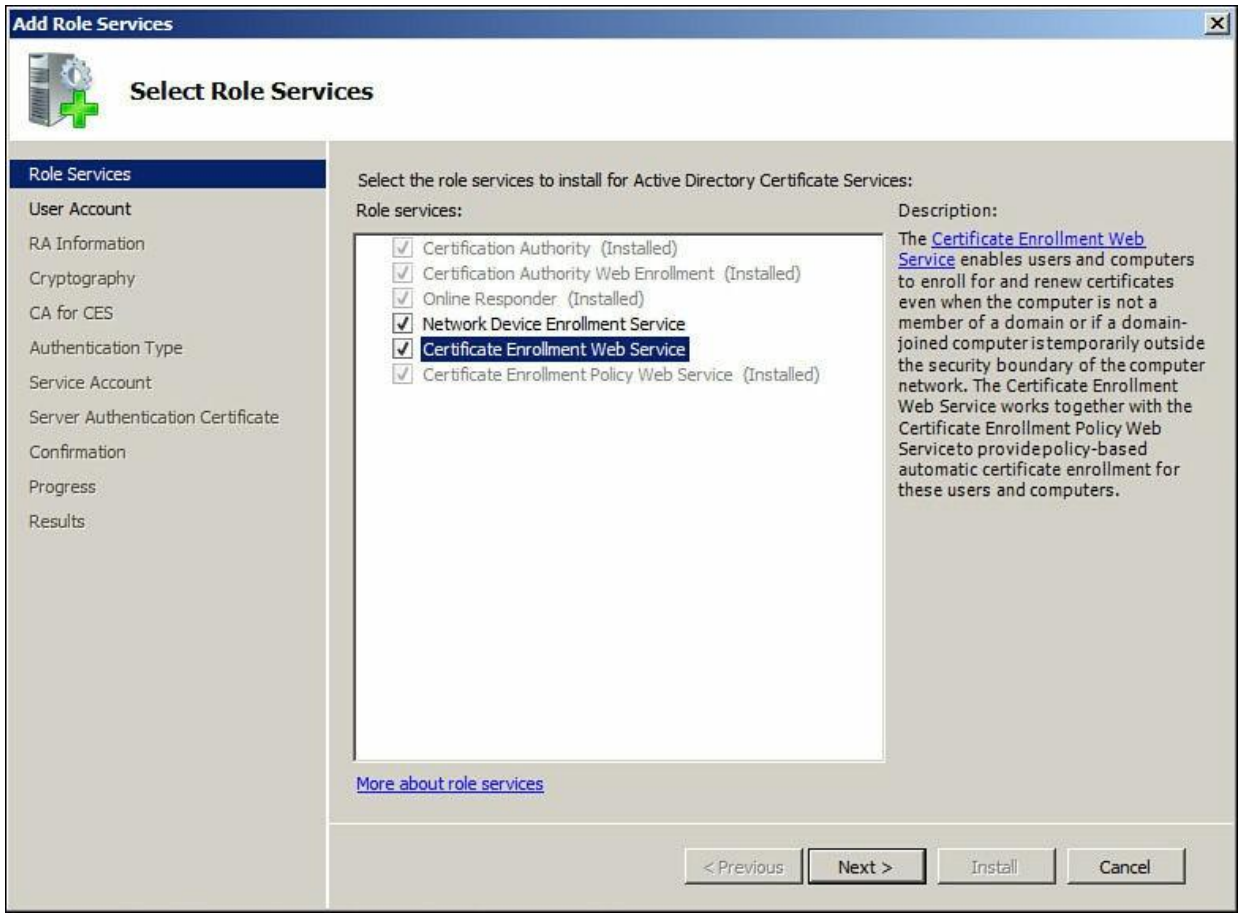


Figure B-13 Select the role services.

Step 4. Click **Next**.

Step 5. Specify the service account user created previously.

https://t.me/learningnets

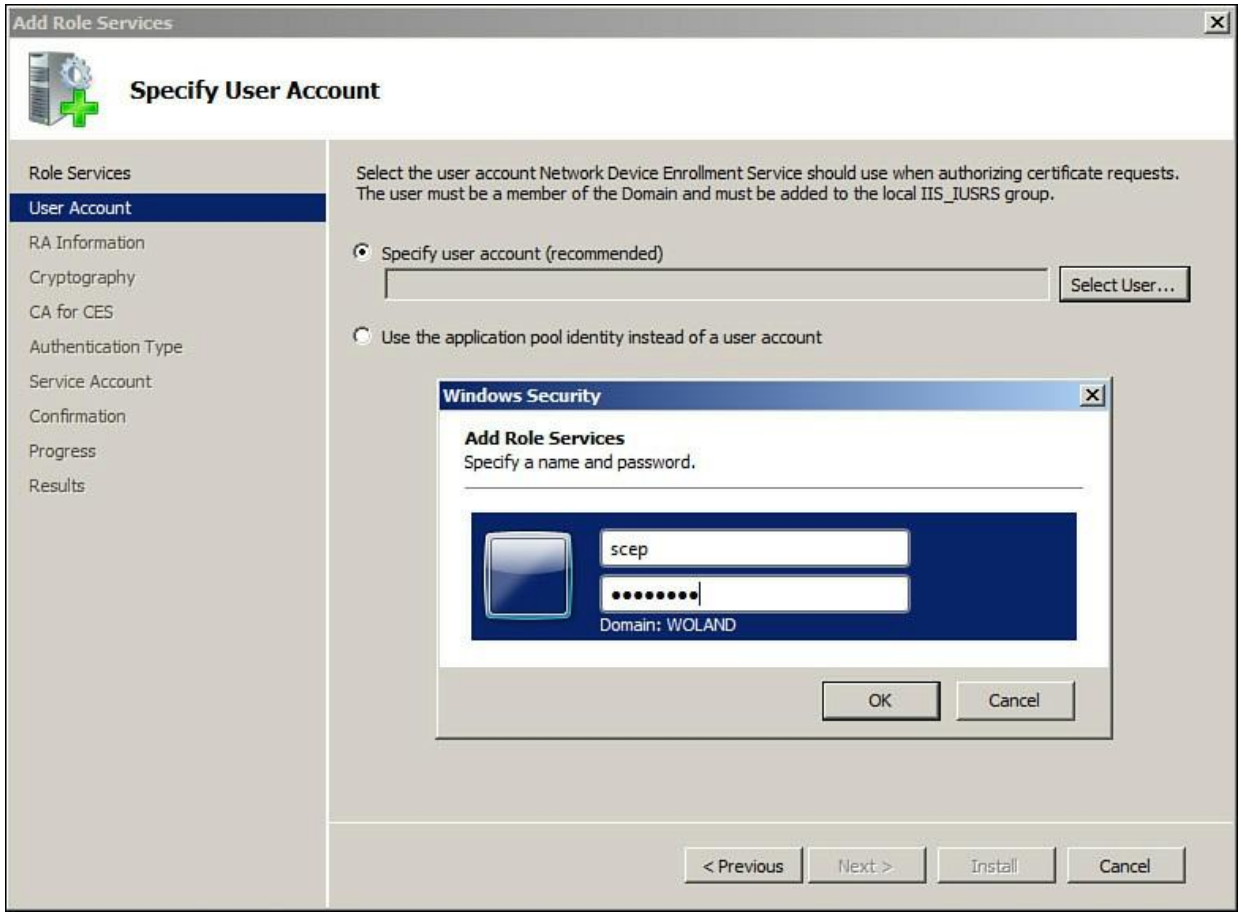


Figure B-14 The service account.

Step 6. Click **Next**.

Step 7. Fill in the **Registration Authority Data** field.

https://t.me/learningnets

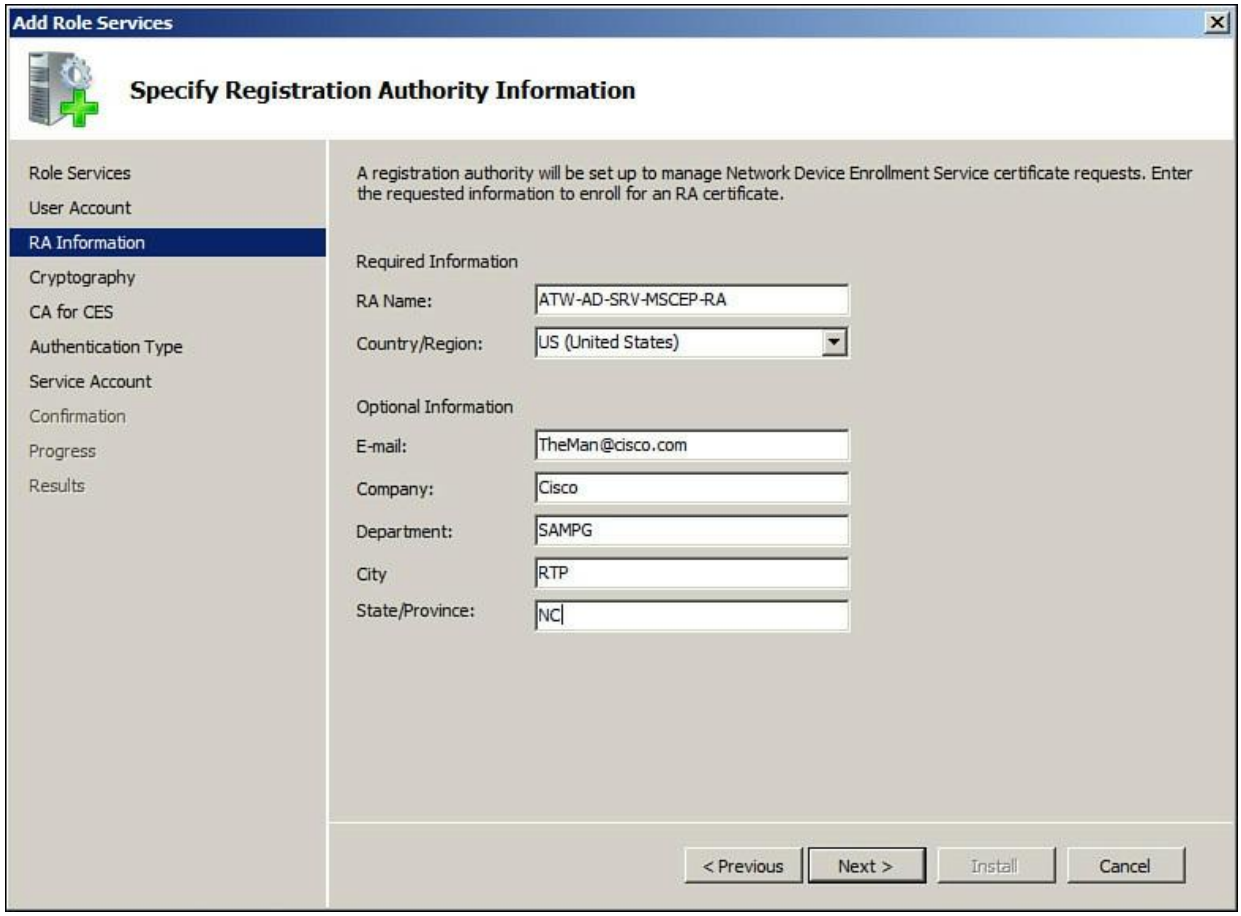


Figure B-15 RA data.

Step 8. Click **Next**.

Step 9. Set the RA cryptographic settings.

https://t.me/learningnets

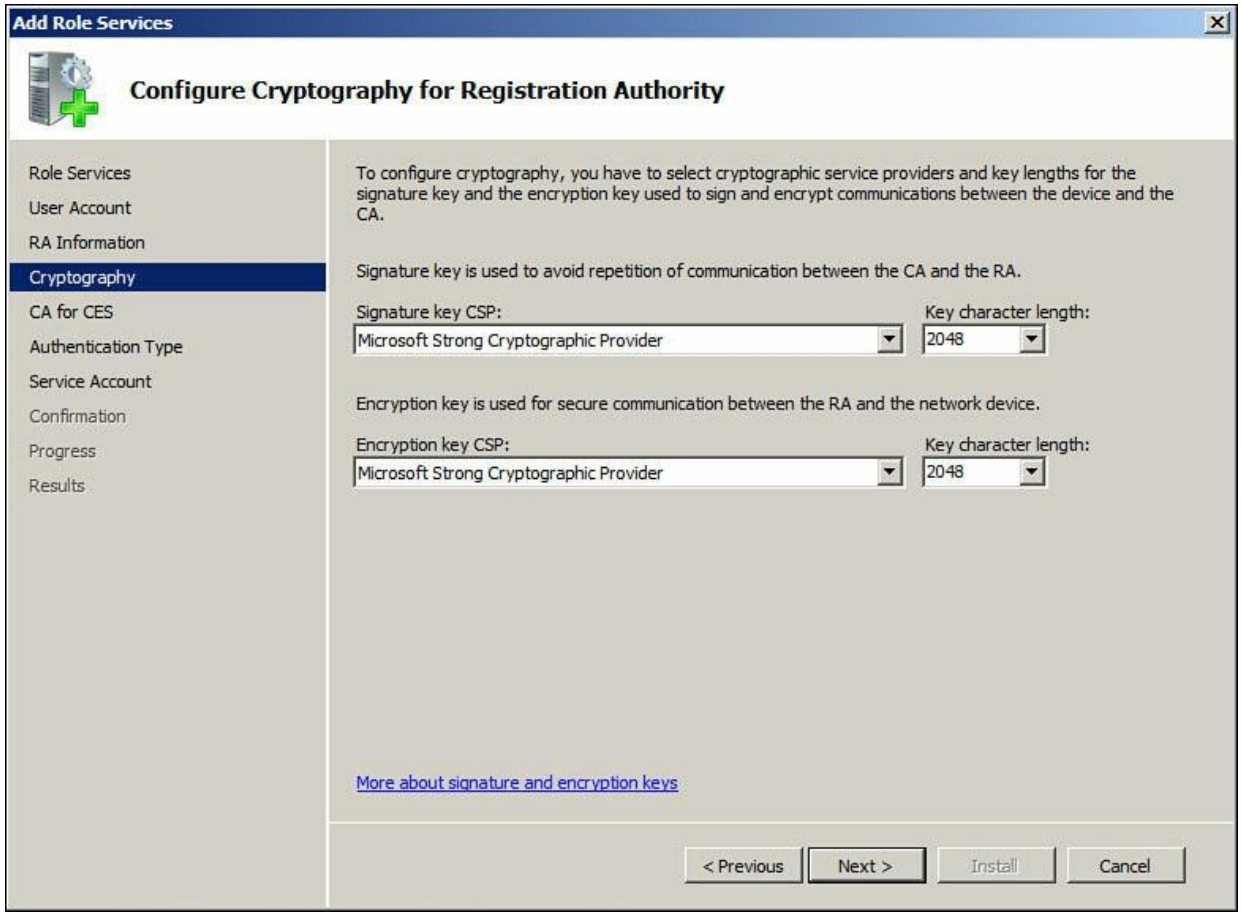


Figure B-16 RA cryptography.

Step 10. Click **Next**.

Step 11. Specify the CA for the Web Services.

https://t.me/learningnets

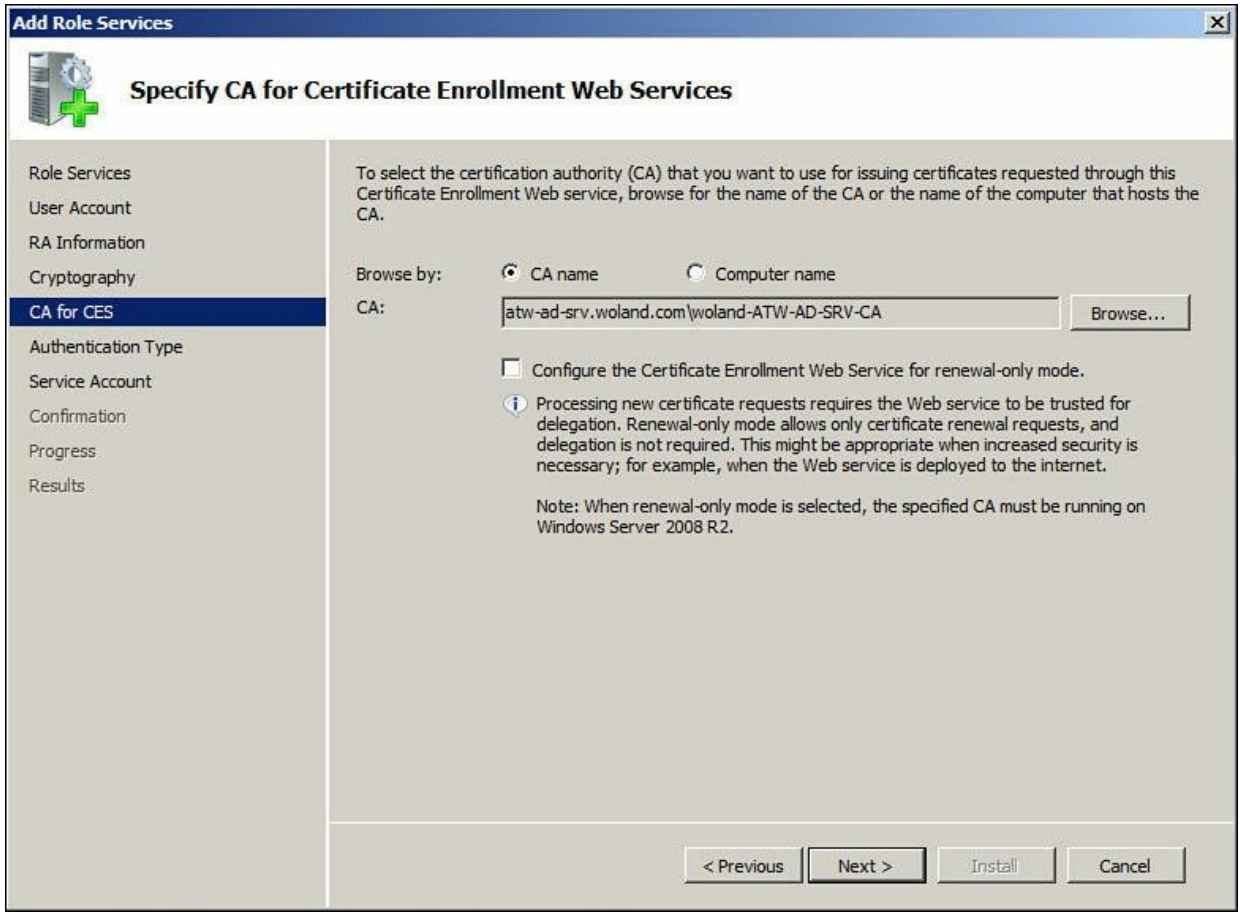


Figure B-17 Web Services CA.

Step 12. Click **Next**.

Step 13. Set the username and password as **Authentication Type**.

https://t.me/learningnets

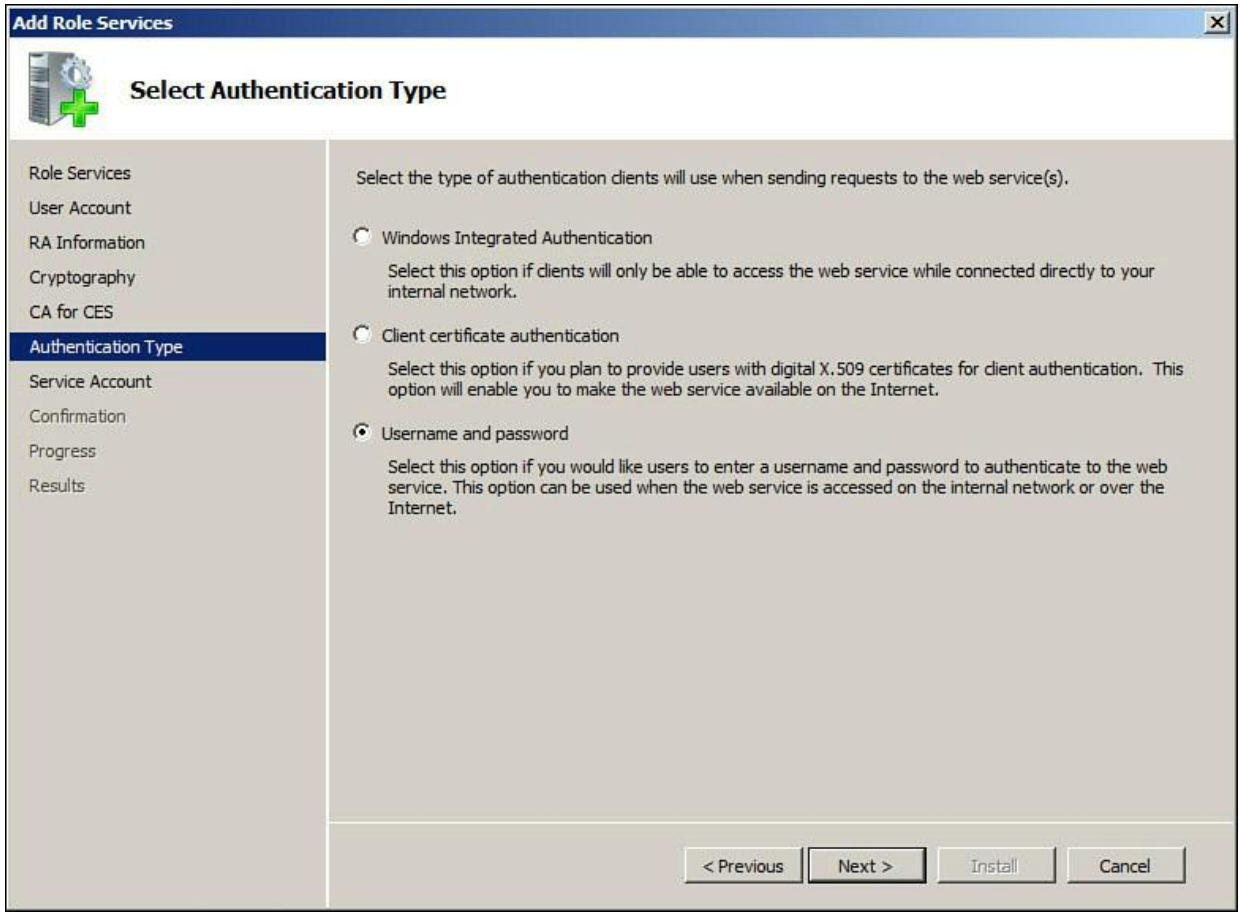


Figure B-18 Authentication type.

Step 14. Click **Next**.

Step 15. Specify the service account user again.

https://t.me/learningnets

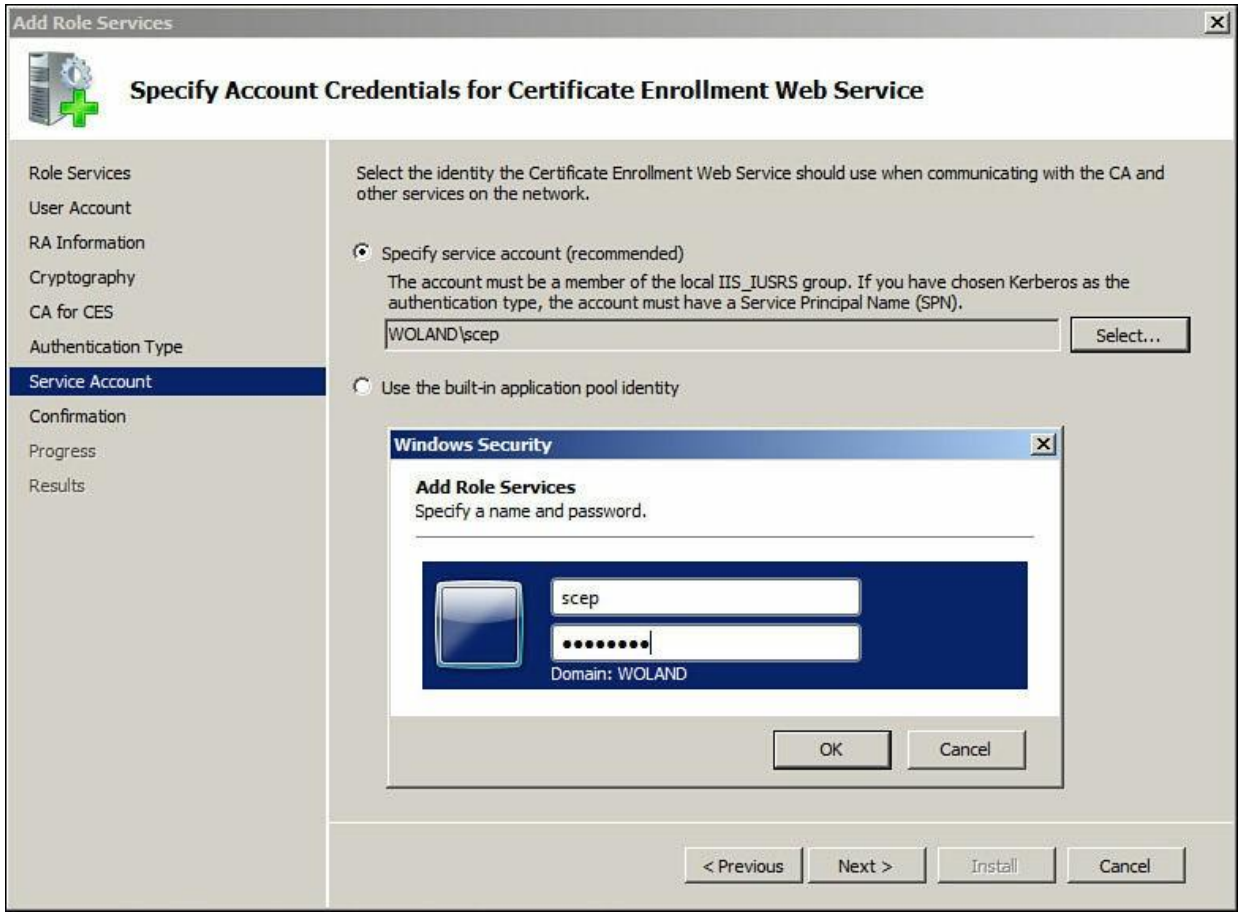


Figure B-19 The service account.

Step 16. Click **Install**.

Configuring the Certificate Template

A certificate template is used to define what fields and usages a certificate will have. The best certificate to base a BYOD certificate on is the user certificate. It will provide an identity certificate used for the end-entity.

Step 1. Navigate to **Server Manager > Roles > AD Certificate Authority > Certificate Templates**.

Step 2. Highlight the certificate template named **User**, and select **Duplicate**.

https://t.me/learningnets

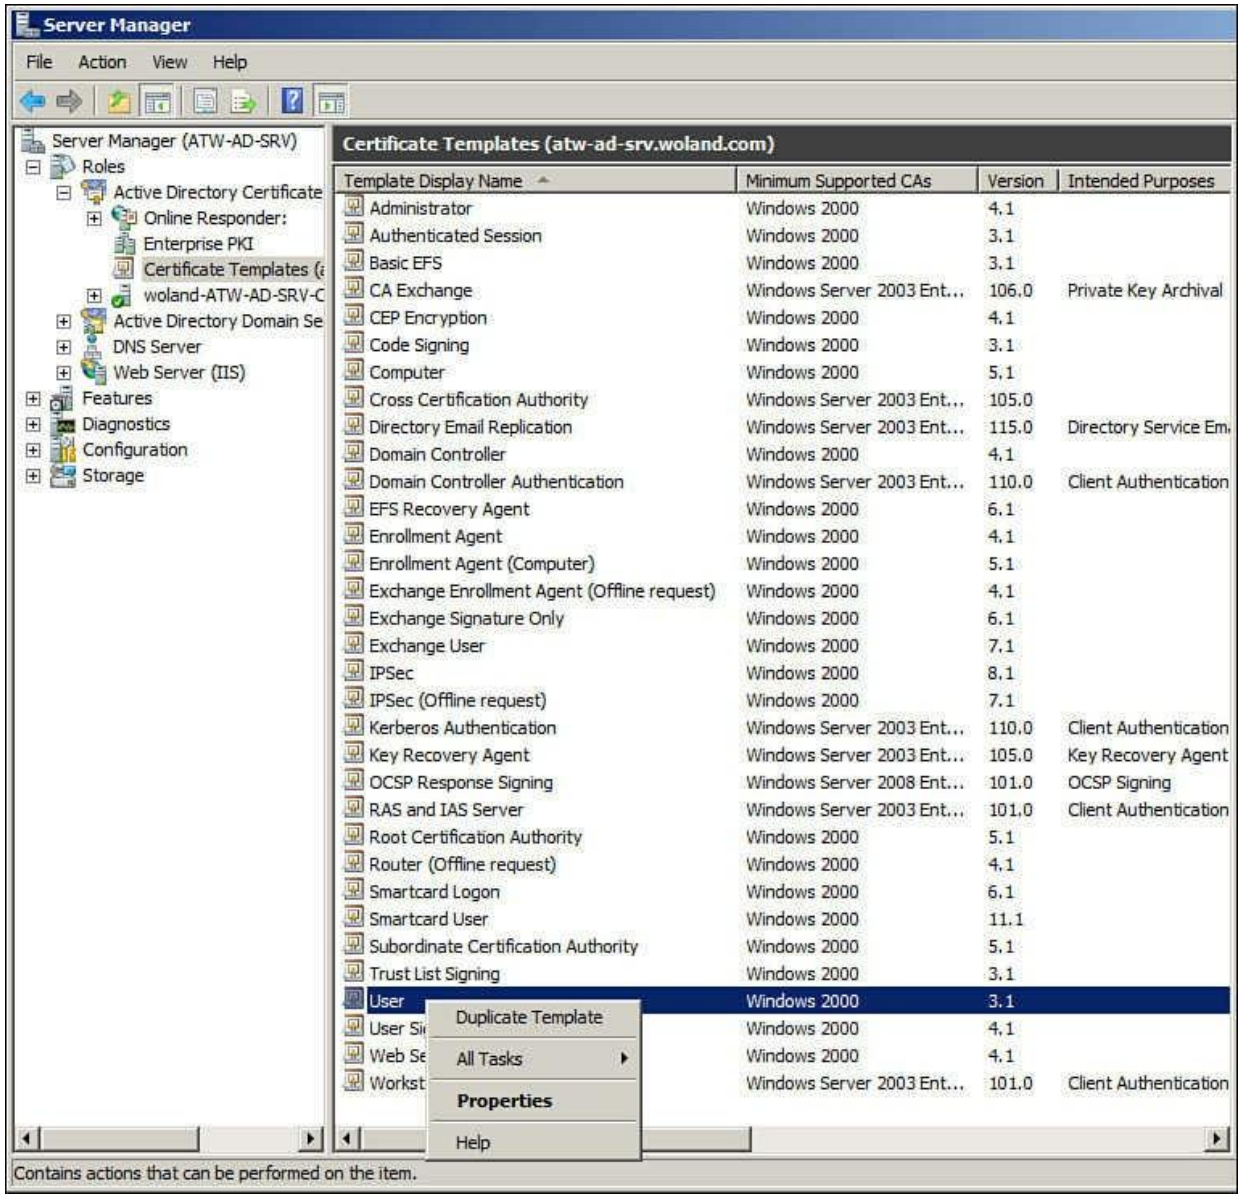


Figure B-20 Duplicate the user template.

Step 3. Click **Next**.

Step 4. Select the **Template Version**. Either will work; the sample screenshots are from a 2008 template version.

https://t.me/learningnets



Figure B-21 Template version.

Step 5. Click **OK**.

Step 6. Name the certificate template, and uncheck the **Publish in Active Directory** check box. This is an important check box, so be sure it is not checked to avoid storage issues.

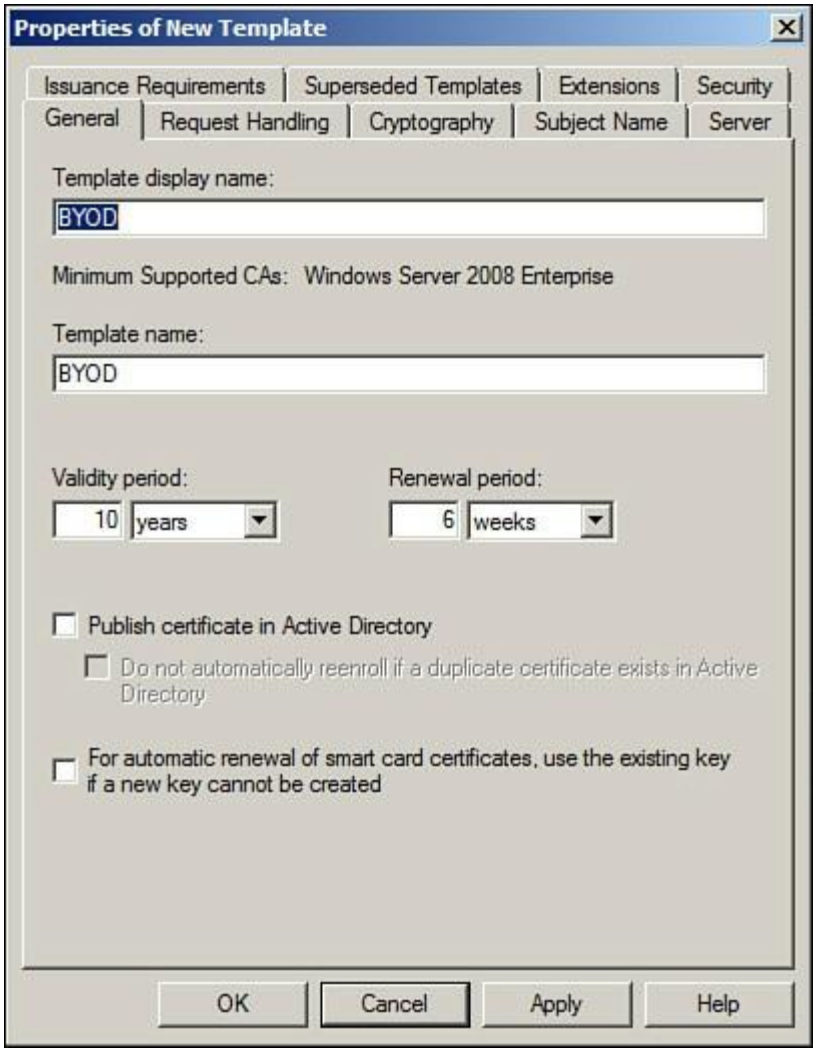


Figure B-22 General tab.

Step 7. Click the **Request Handling** tab:

https://t.me/learningnets

- Purpose—From the drop-down list, select **Signature and Encryption** so the certificate will be used for signing and encrypting.
- Uncheck the **Allow Private Key to Be Exported** if you want; this marks it as non-exportable.
- SCEP is an automated process, so be sure the **Enroll Subject Without Requiring Any User Input** radio button is selected.

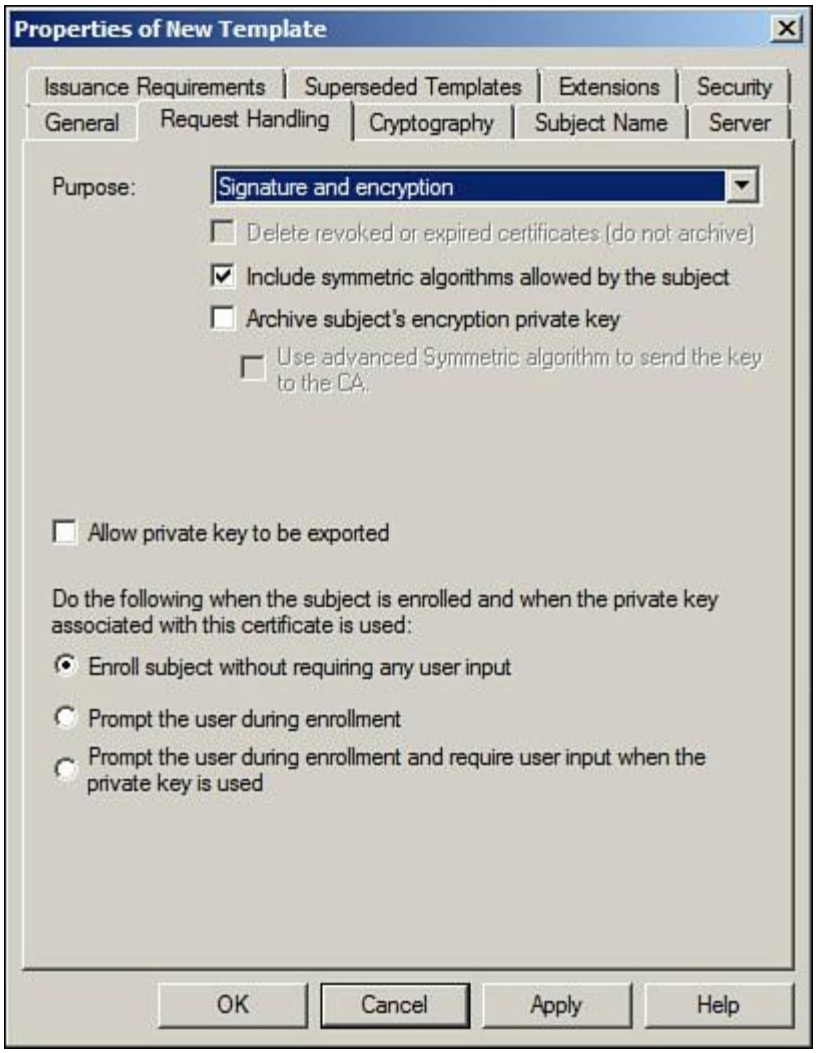


Figure B-23 Request Handling tab.

Step 8. Click the **Subject Name** tab.

Step 9. The BYOD process is pre-building the Certificate Signing Request, so ensure that the **Supply in the Request** option is selected.

https://t.me/learningnets

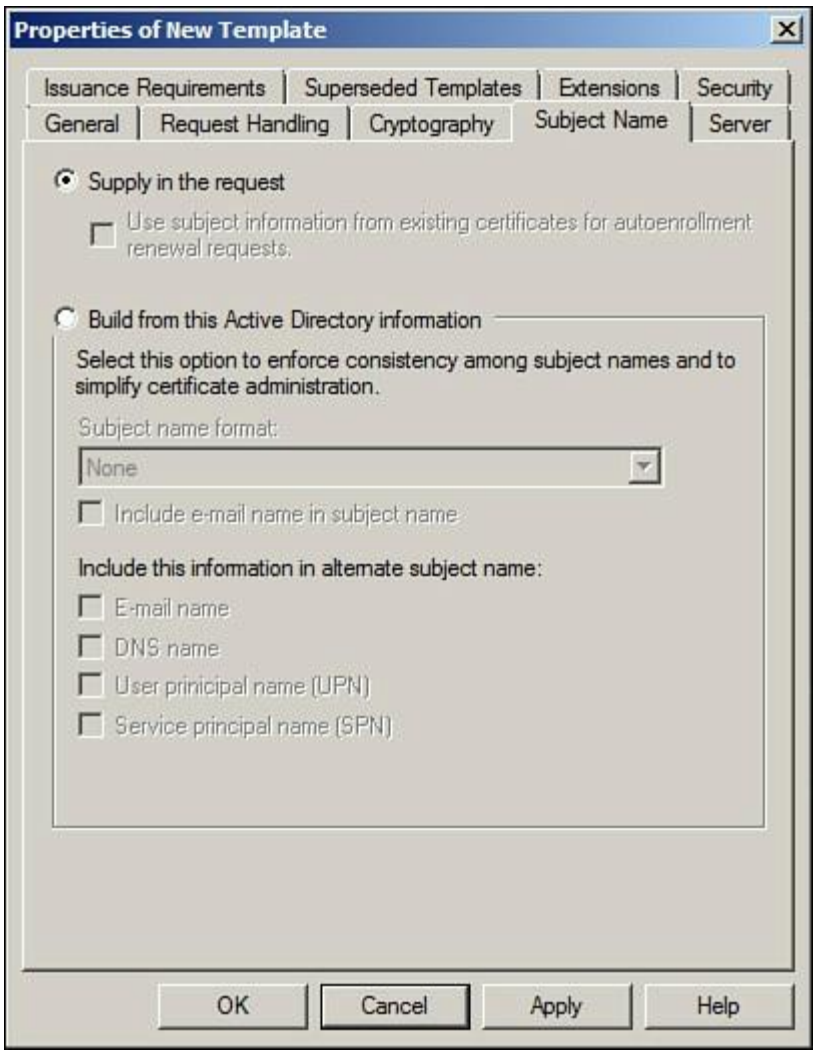


Figure B-24 Subject Name tab.

Step 10. Click the **Extensions** tab.

Step 11. Click on **Application Policies**, ensure that client authentication is listed.

https://t.me/learningnets

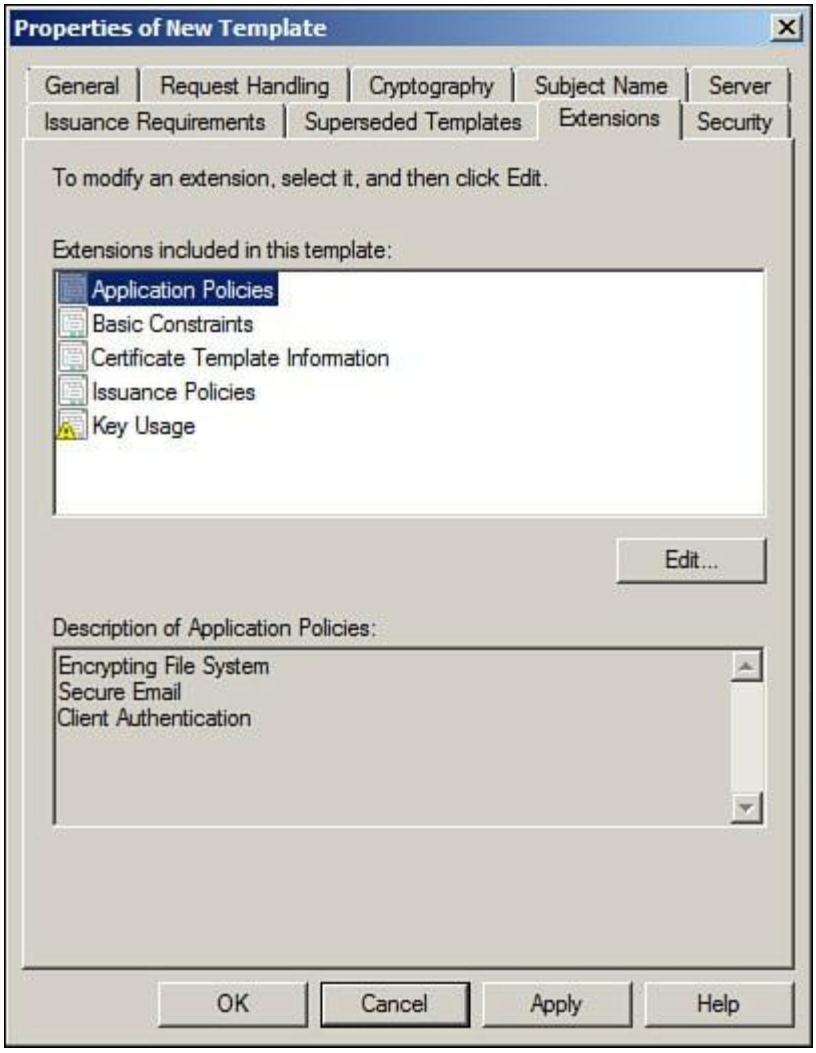


Figure B-25 Application Policies.

Step 12. Click **Issuance Policies**.

Step 13. Click **Edit**, and then click **Add**.

Step 14. Select **All Issuance Policies**. This is a critical step to ensure the certificate is issued to the endpoint.

https://t.me/learningnets

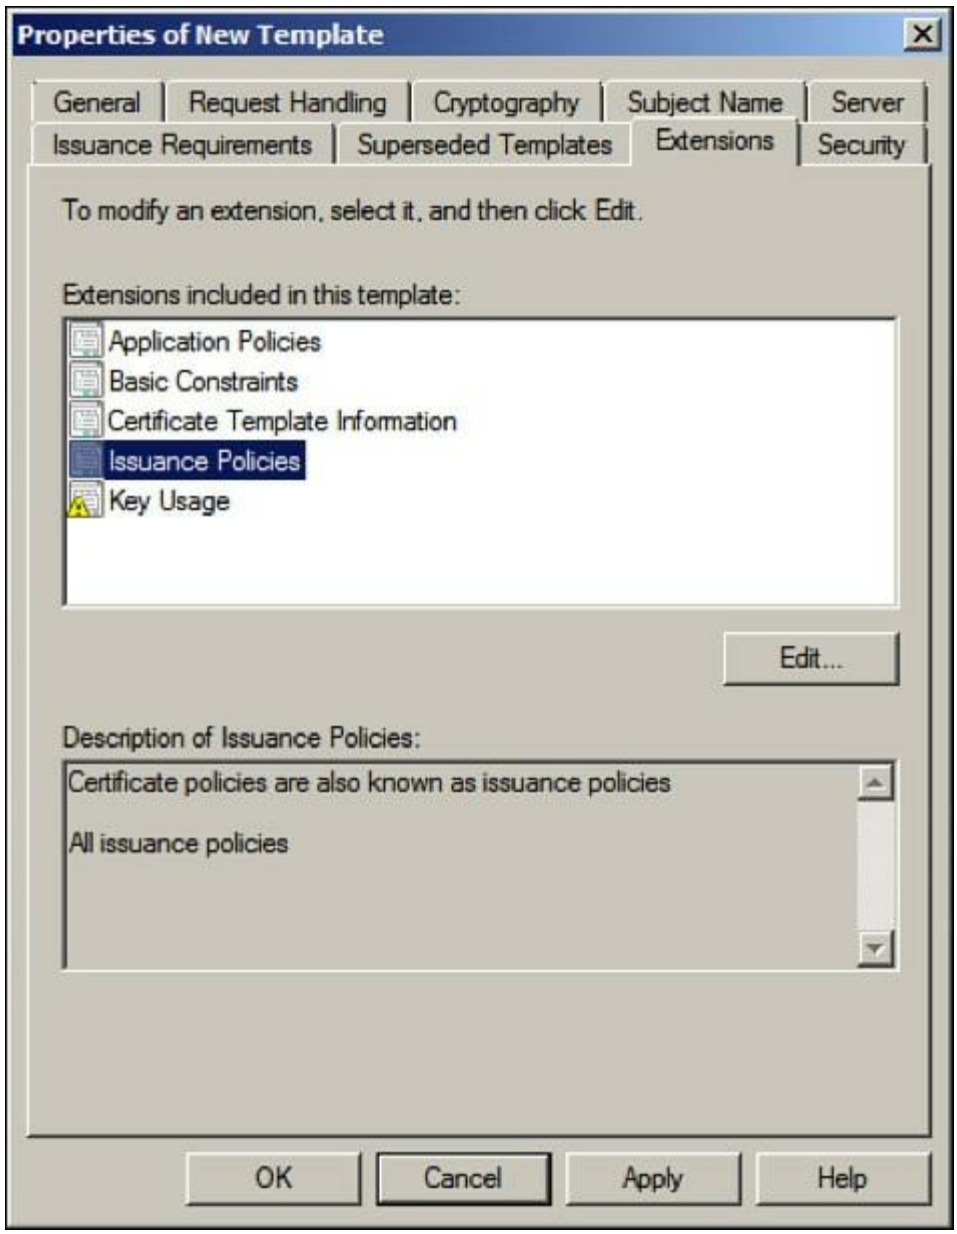


Figure B-26 Issuance policies.

Step 15. Click the **Security** tab.

Step 16. Add the service account user to have full control.

Step 17. Click **OK** to save the template.

Publishing the Certificate Template

The template is created, but we have to choose it as one to be issued:

Step 1. Select **Server Manager > Roles > AD Certificate Authority >** <your CA> **> Certificate Templates**.

Step 2. Right-click in the window.

Step 3. Select **New > Certificate Template to Issue**.

https://t.me/learningnets

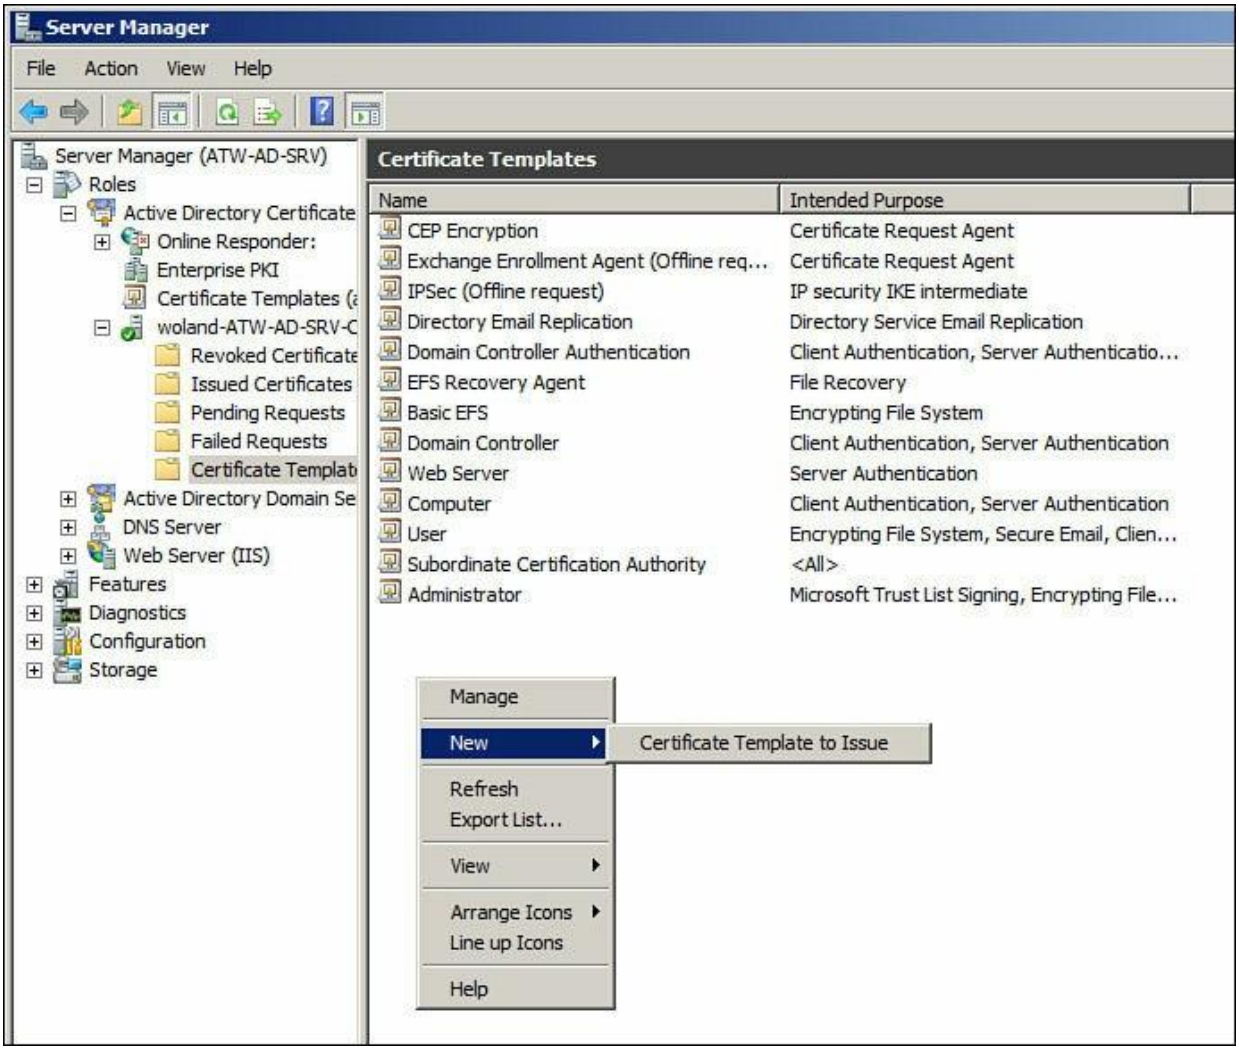


Figure B-27 Certificate template to issue.

Step 4. Choose your new certificate template.

https://t.me/learningnets



Figure B-28 BYOD certificate template.

Editing the Registry

The service account user must have full control of the MSCEP registry key. Do the following:

Step 1. Open the Regedit application.

Step 2. Select the MSCEP registry key from **HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\MSCEP**.

https://t.me/learningnets

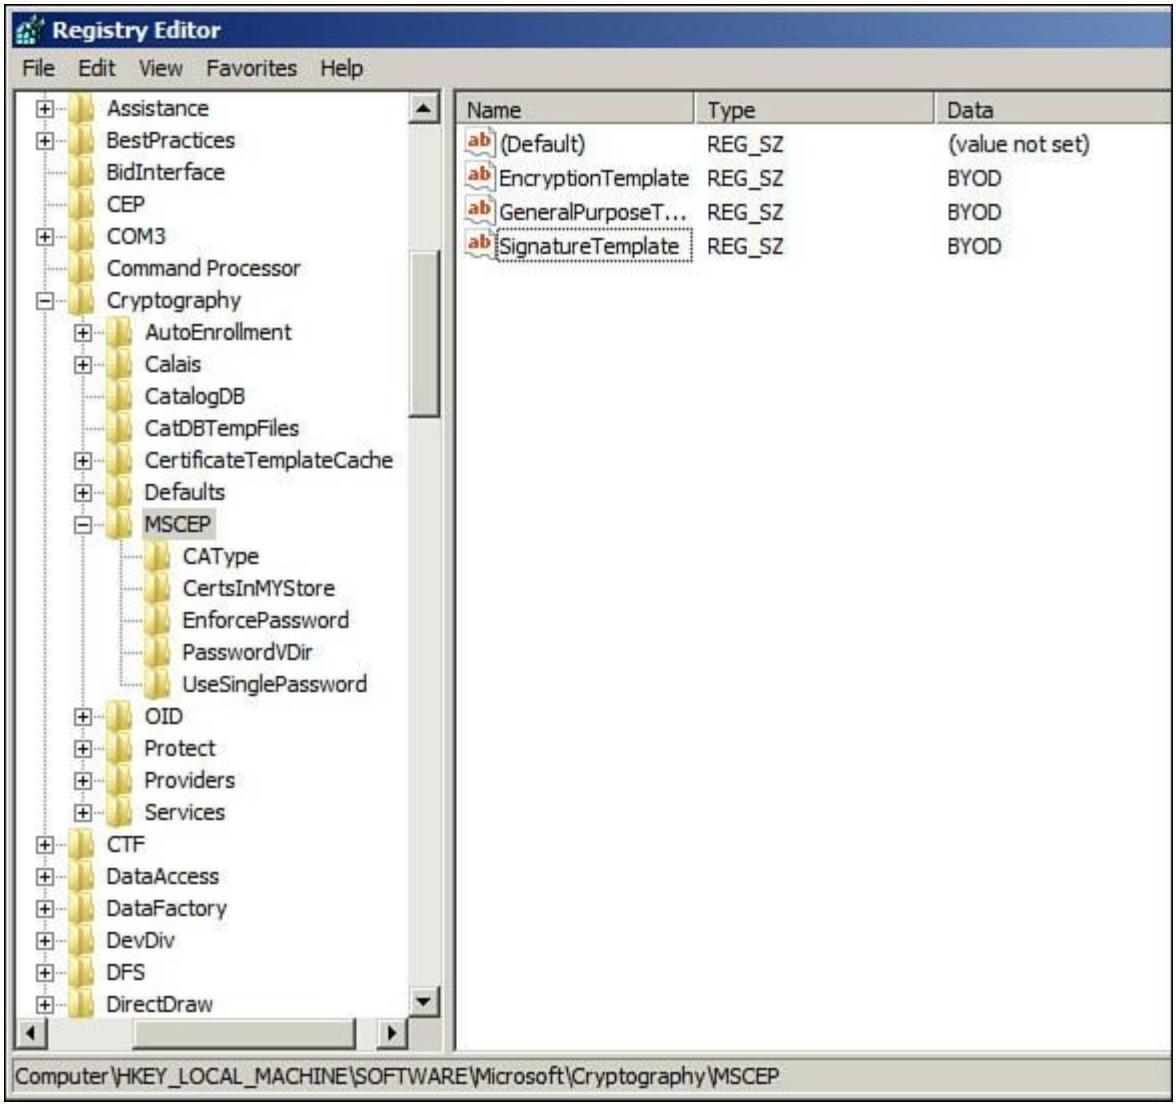


Figure B-29 Regedit.

Step 3. Right-click **MSCEP** > **Permissions**.

Step 4. Add the service account, and give it full control.

https://t.me/learningnets

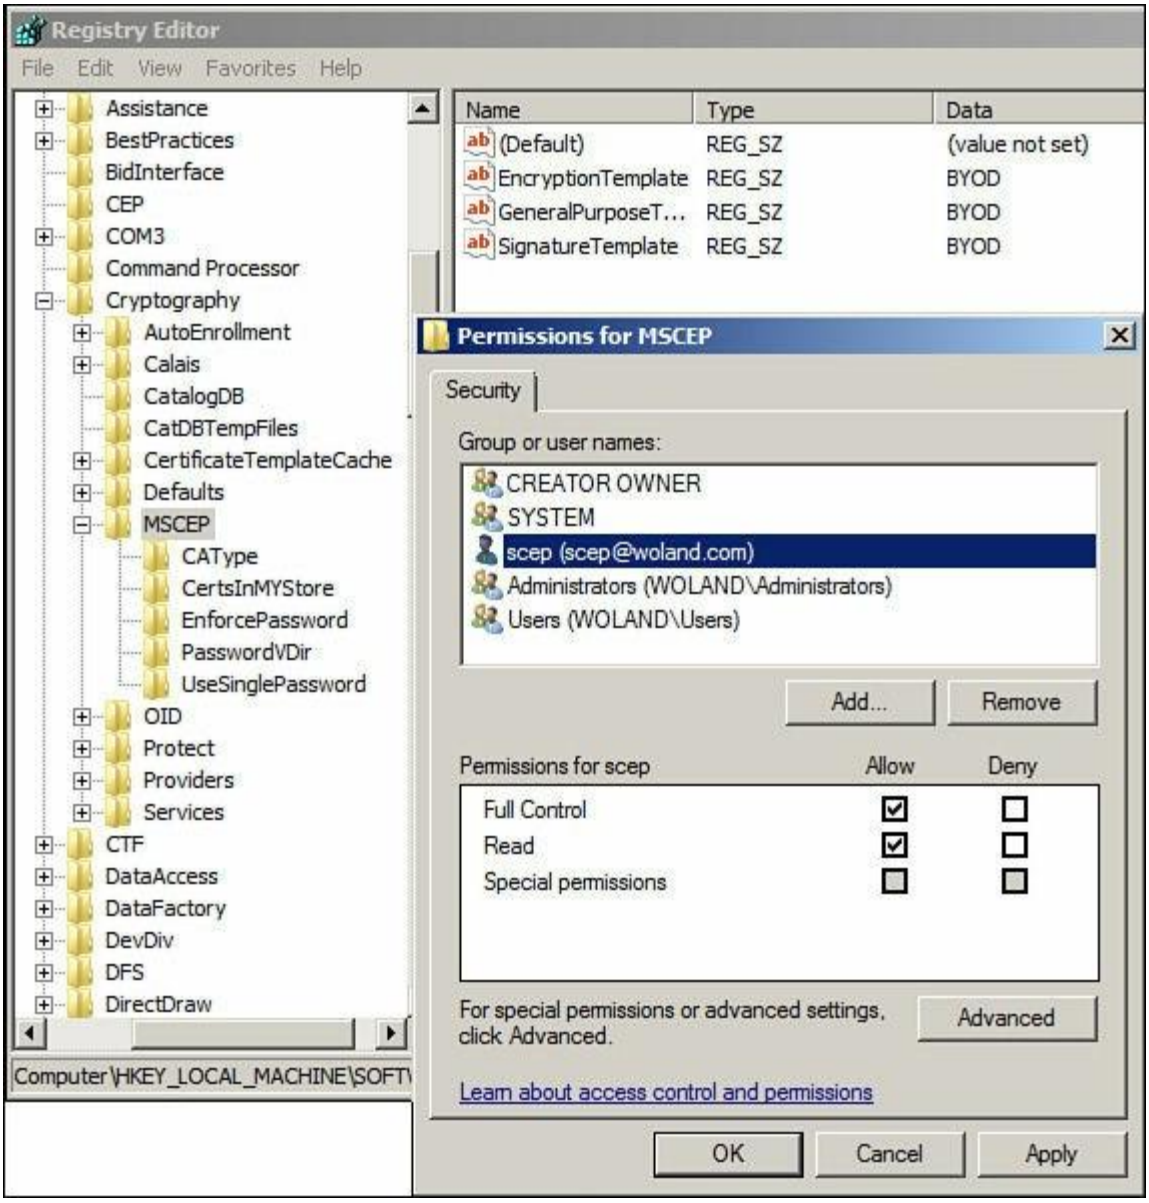


Figure B-30 Giving full control to the service account.

The default certificate template for SCEP to issue is an IPSec template. You must change this to use the new user template:

Step 1. From **HKLM > Software > Microsoft > Cryptography > MSCEP**, change the three Registry values to be the name of your newly created template, as shown in Figure B-29.

While in Regedit, you must disable the UseSinglePassword setting:

Step 1. From **HKLM > Software > Microsoft > Cryptography > MSCEP**, select the **UseSinglePassword** key.

Step 2. Change the value to **0.**

https://t.me/learningnets

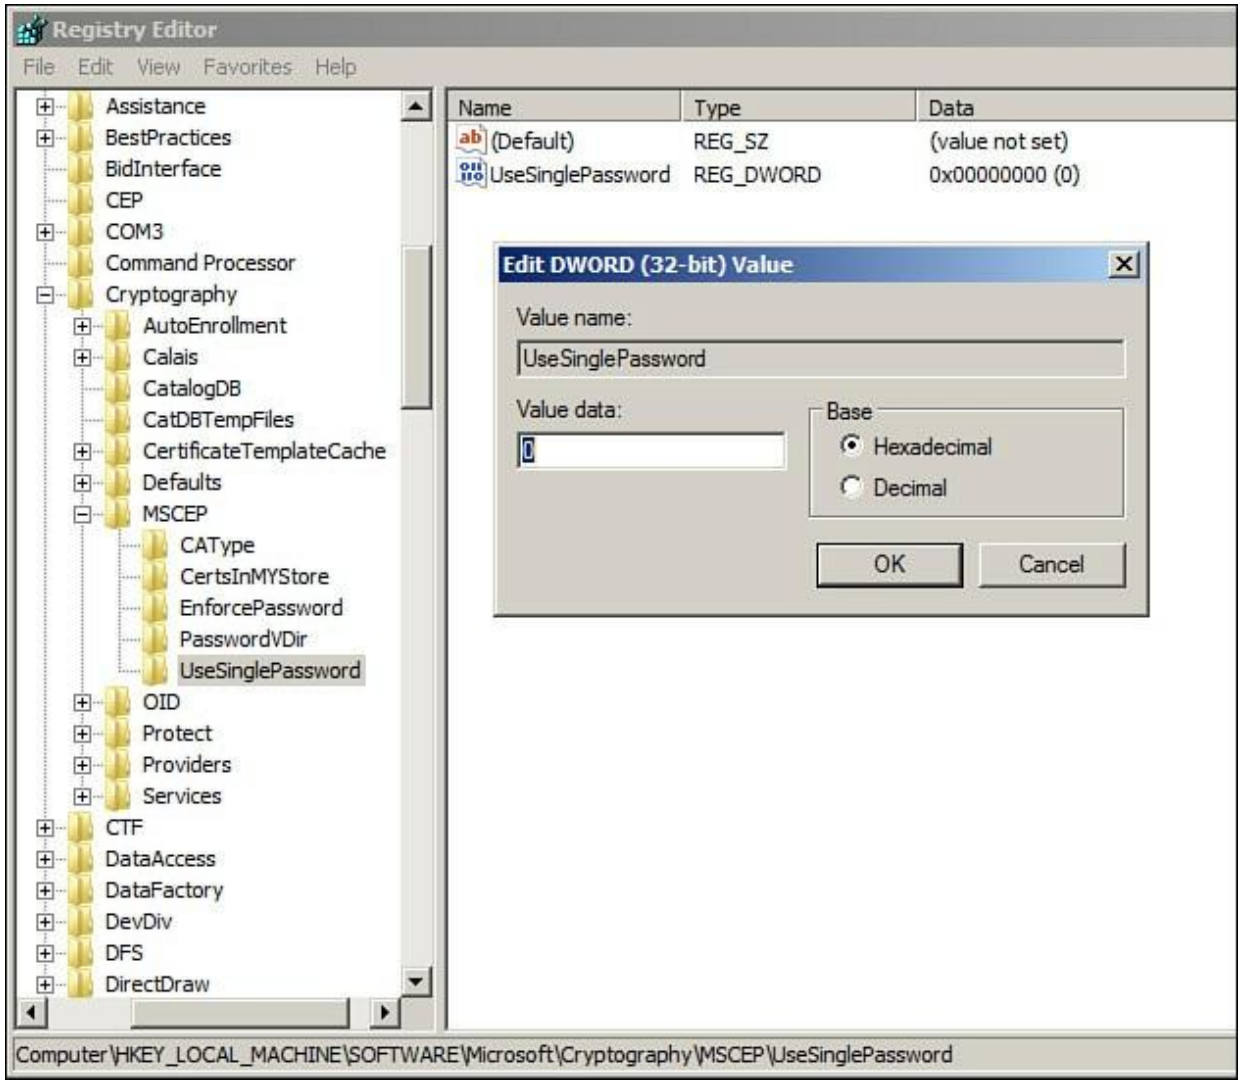


Figure B-31 The UseSinglePassword setting.

Step 3. Select the **EnforcePassword** key.

Step 4. Change the value to **0.**

https://t.me/learningnets

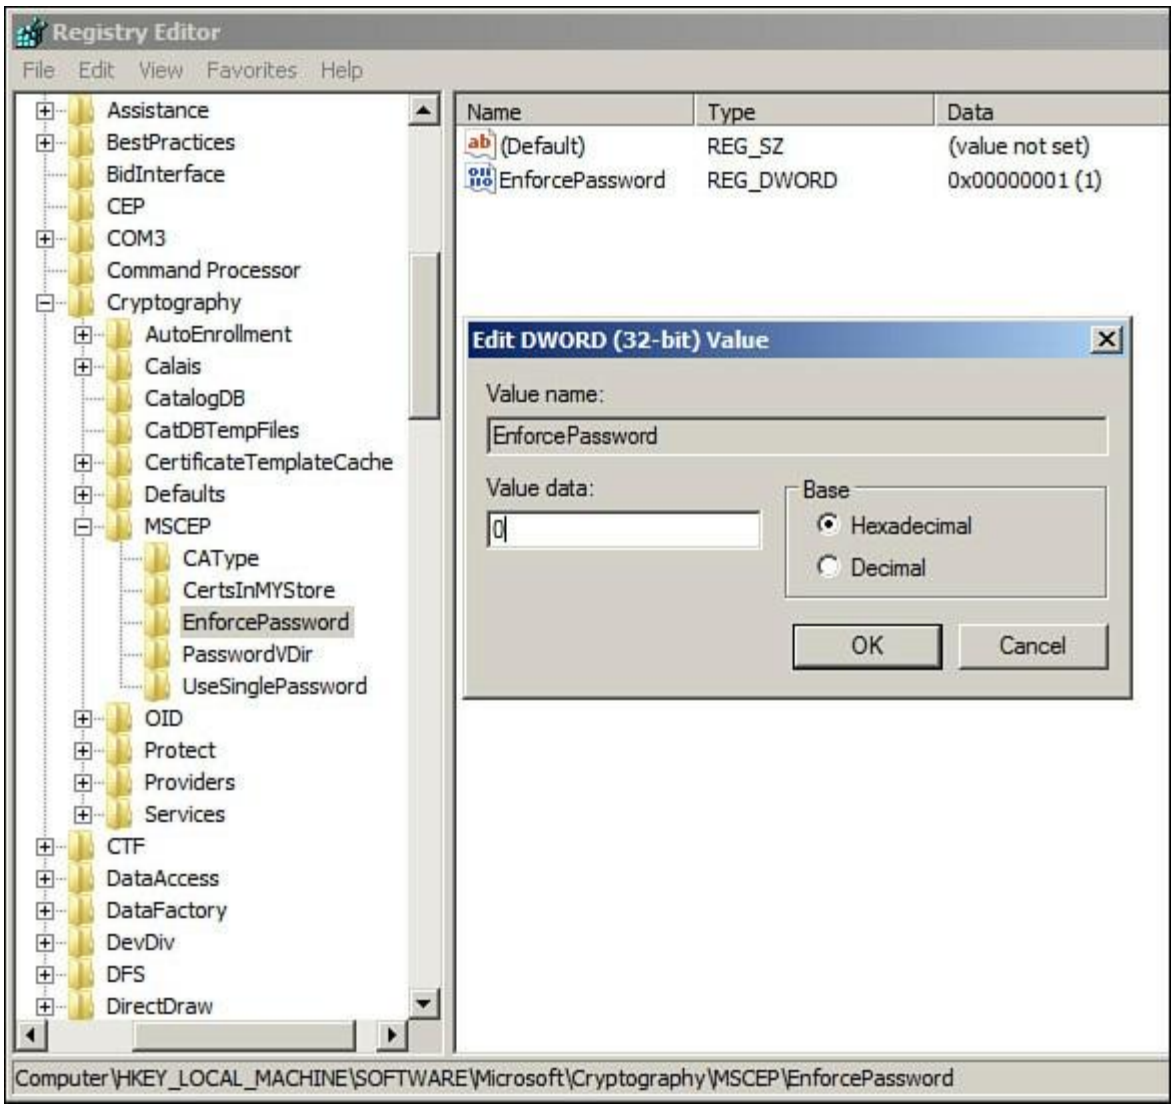


Figure B-32 The EnforcePassword setting.

You are finished! Reboot the server.

Useful Links

Microsoft Technet Article, "Certificate Enrollment Web Services in Active Directory Certificate Services," http://social.technet.microsoft.com/wiki/contents/articles/7734.certificate-enrollment-web-services-in-active-directory-certificate-services.aspx

Microsoft Technet Article, "AD CS: Deploying Network Device Enrollment Services," http://technet.microsoft.com/en-us/library/ff955646%28v=ws.10%29.aspx.

https://t.me/learningnets

Appendix C. Using the Dogtag CA for BYOD

As of ISE 1.2, no certificate authority (CA) is built in to ISE, and it requires an external CA. Whatever the reason, some installations require an alternative to the Microsoft CA. This appendix was created to show one such alternative to the MS Certificate Authority.

What Is Dogtag, and Why Use It?

Dogtag is an enterprise-class open source CA that Red Hat purchased from AOL back in 2004. Red Hat opened it up to the open source community in 2008. Dogtag supports all aspects of certificate lifecycle management, including key archival, OCSP and smartcard management, and much more.

> **Note**
>
> An enterprise-level version of Dogtag known as the Red Hat Certificate System also exists.

Prerequisites

Dogtag will run on most Red Hat variants. For the purposes of this appendix, we will focus on Fedora Core 15 (32-bit). This is the version that is known to work and has been tested with ISE 1.2. This version of Fedora can be installed with the minimum option and will leverage the Apache web server, PHP, and the open source directory server.

Installing 32-bit Fedora 15

Before you can install Dogtag, you have to install a base operating system and its components. As a prerequisite, you must ensure that a DNS entry exists for the host, so it can be reached by its fully qualified domain name (FQDN):

Step 1. Boot the machine with the 32-bit Fedora 15 ISO file or DVD available here:
https://mirrors.fedoraproject.org/publiclist/Fedora/15/

Step 2. Select **Install system with Basic Video Driver**, as shown in Figure C-1.

https://t.me/learningnets

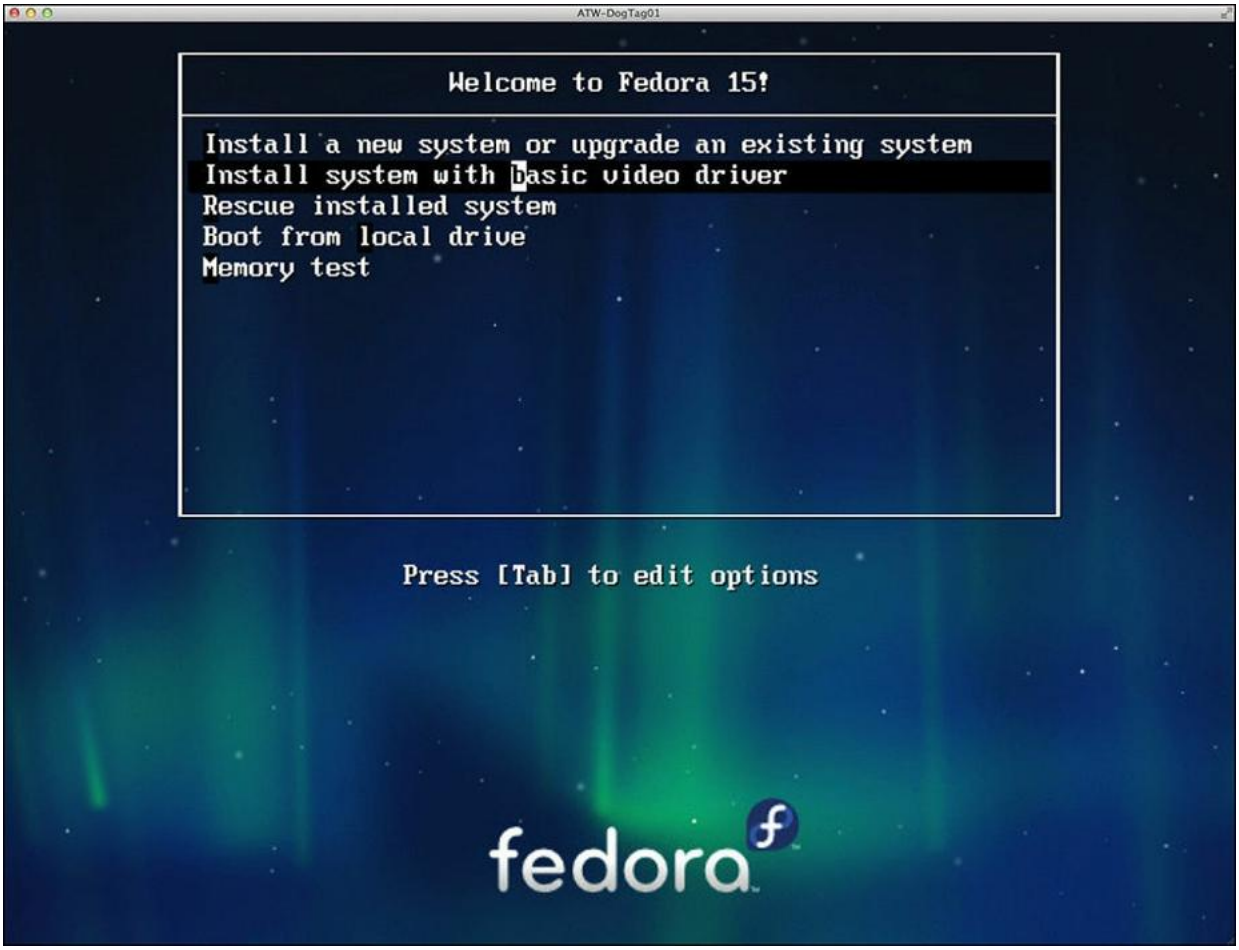


Figure C-1 Install the system with a basic video driver.

Step 3. The minimal installation type is all you need for this use case. So, select the **Minimal** software set, as shown in Figure C-2.

https://t.me/learningnets

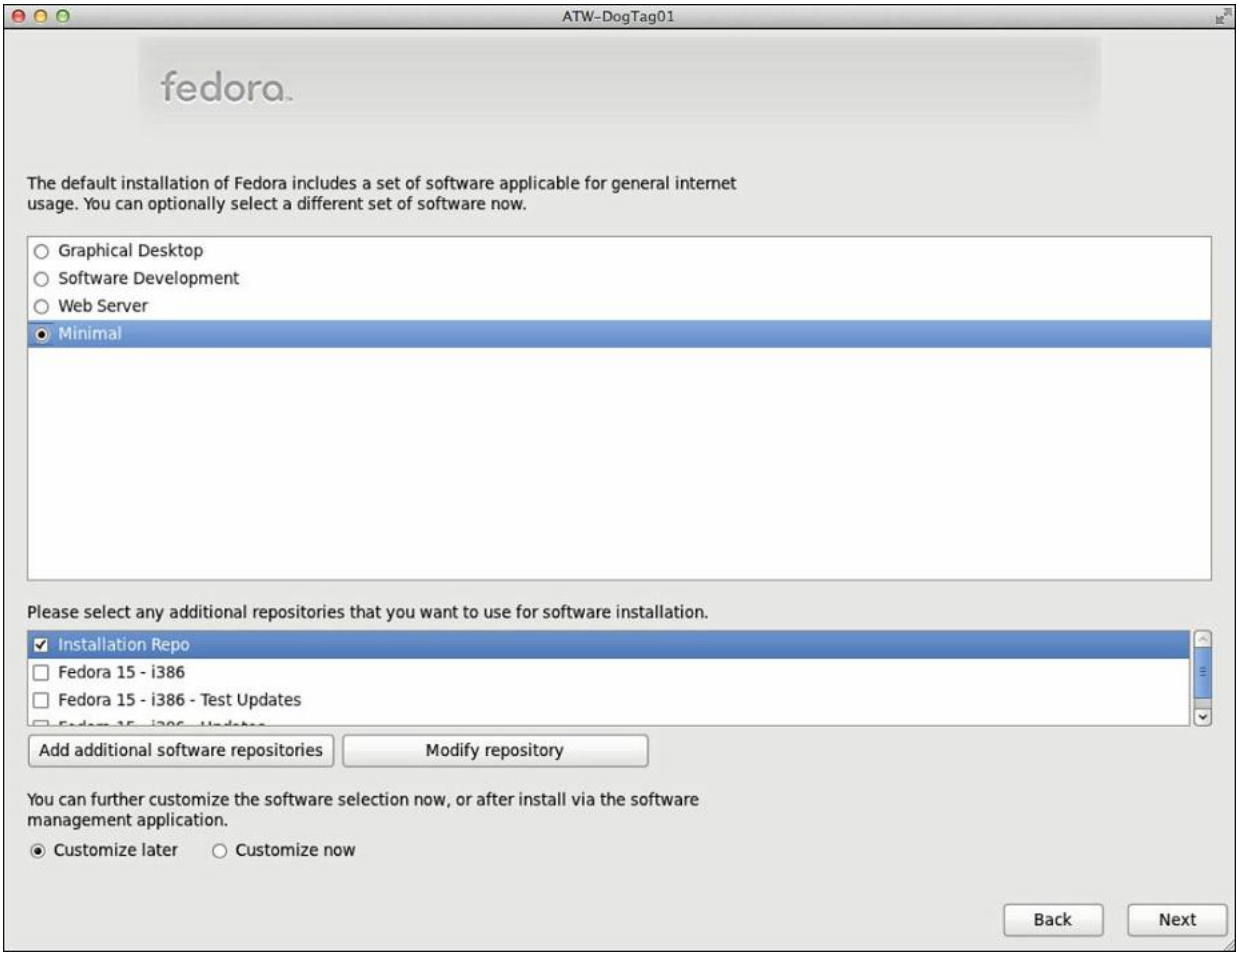


Figure C-2 Use the minimal installation type.

Step 4. Accept the default choices for the remainder of the installation.

Configuring Networking

The CA should have a static IP address to ensure that communication is always optimal. One component of the setup wizard enables you to configure the network prior to the installation finishing. However, with Fedora 15, the majority of the time those settings do not seem to be maintained, and when the Fedora operating system is fully installed, there is no assigned IP address, as shown in Figure C-3.

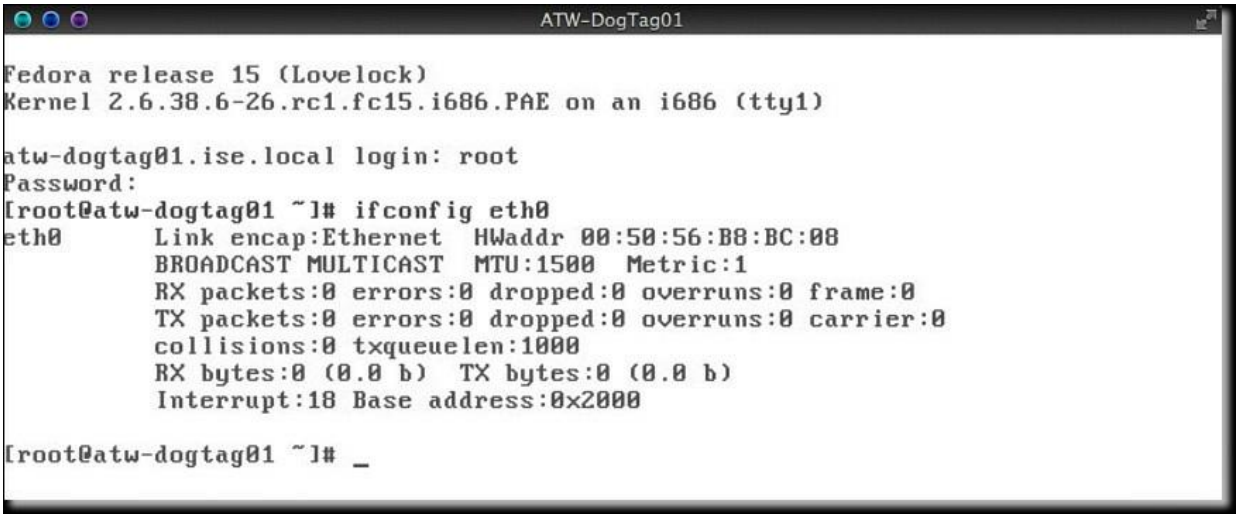


Figure C-3 No IP address after installation.

https://t.me/learningnets

> **Note**
>
> It is assumed that you are logged in as root to perform the activities in this document. If not, use the **su –** command to change your login context to the superuser (root).

Step 1. After the installation, verify whether an IP address exists. Use the **ifconfig eth0** command. Figure C-3 shows the result when no IP address has been configured.

Step 2. Using your favorite editor, edit the ifcfg-eth0 file to set up the network stack for the interface. This example is using the vi editor, as shown in Example C-1.

Example C-1 Editing the ifcfg-eth0 File

Click here to view code image

```
[root@atw-dogtag01 ~]# vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

Step 3. With the ifcfg-eth0 file open, ensure that the **ONBOOT** option is set to **yes**. This ensures that the interface will be on when the system reboots.

Step 4. Ensure the **BOOTPROTO** option is set to **none**. This configures the interface to use a static IP address.

Step 5. Set the **IPADDR** option to be the desired IP address of the server, and set the **NETMASK** to be the subnet mask for that IP address.

Step 6. You can use the **DNS1** and **DNS2** options to point the server to the correct DNS server(s).

Step 7. Use the **GATEWAY** option to specify the IP address of the default gateway. Example C-2 shows the details of a configured ifcfg-eth0 file.

Example C-2 Configured ifcfg-eth0 File

Click here to view code image

```
[root@atw-dogtag01 ~]# cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE="eth0"
HWADDR="00:50:56:B8:BC:08"
ONBOOT="yes"
NM_CONTROLLED="yes"
BOOTPROTO=none
IPADDR=10.1.100.229
NETMASK=255.255.255.0
USERCTL=yes
TYPE=Ethernet
DNS1=10.1.100.103
GATEWAY=10.1.100.1
```

Step 8. Ensure the network starts at boot with the **chkconfig network on** command.

Example C-3 Ensuring the Network Starts at Boot and Restarting the Service

Click here to view code image

```
[root@atw-dogtag01 ~]# chkconfig network on
[root@atw-dogtag01 ~]# service network restart
```

https://t.me/learningnets

Installing Packages with yum

Fedora uses a software package manager called yum to manage the installed packages within the operating system. yum provides the advantage of identifying dependencies and helping to manage the installation of the application and all of that application's dependencies. See http://fedoraproject.org/wiki/Yum for more on yum.

In this section you will use yum to update the Fedora 15 server to have the latest packages, as well as install needed applications such as NTP.

Configuring Proxy (if Needed)

The setup used to write this section required a proxy server to access the Internet. Therefore, this procedure was included. If your environment does not require a proxy to access the Internet, you can skip this section. Do the following:

Step 1. Use your favorite text editor to edit the yum configuration file located at /etc/yum.conf (see Example C-4).

Example C-4 Editing the yum Configuration File

Click here to view code image

```
[root@atw-dogtag01 ~]# vi /etc/yum.conf
```

Step 2. Add a line for with a field of **proxy=** followed by the URL and port for your proxy server. The completed file is shown in Example C-5.

Example C-5 The Complete yum.conf File

Click here to view code image

```
[root@atw-dogtag01 ~]# cat /etc/yum.conf
[main]
cachedir=/var/cache/yum/$basearch/$releasever
keepcache=0
debuglevel=2
logfile=/var/log/yum.log
exactarch=1
obsoletes=1
gpgcheck=1
plugins=1
installonly_limit=3
proxy=http://proxy.woland.com:8080
```

Updating System Packages with yum

With Linux systems, software applications are known as packages. yum is the package manager used to install and update the packages:

Step 1. Add a yum plug-in to choose the fastest location from which to download, as seen in Example C-6. This plug-in saved hours during the writing of this book.

Example C-6 Installing the Fastest Mirror Plug-in

Click here to view code image

```
[root@atw-dogtag01 ~]# yum install yum-plugin-fastestmirror
```

https://t.me/learningnets

Step 2. Update all the installed packages with the **yum update** command, as seen in Example C-7.

Example C-7 Updating All Installed Packages with yum

Click here to view code image

```
[root@atw-dogtag01 ~]# yum update
Loaded plugins: fastestmirror
Determining fastest mirrors

<<SNIP>>

Transaction Summary
================================================================================
Install       4 Package(s)
Upgrade     104 Package(s)

Total download size: 89 M
Is this ok [y/N]:
```

Installing and Configuring the NTP Service

Certificates require strict time synchronization. It's recommended that you use the network time protocol (NTP) to ensure the time is accurate on the CA. The NTP service (all called NTP daemon) is not installed by default with the minimal installation of Fedora 15, so we will use yum to install it:

Step 1. Install the NTP service with the **yum install ntp** command.

Step 2. Use the **chkconfig ntpd on** command to ensure the NTP daemon starts at boot.

Step 3. Use the **ntpdate ntp_server_ip_address** command to sync to an NTP source.

Step 4. Ensure the service is started with the **/etc/init.d/ntpd start** command.

The four steps are also shown in Example C-8.

Example C-8 Installing, Syncing, and Starting NTP

Click here to view code image

```
[root@atw-dogtag01 ~]# yum install ntp
[root@atw-dogtag01 ~]# chkconfig ntpd on
[root@atw-dogtag01 ~]# ntpdate [ip-address]
31 Jul 13:47:44 ntpdate[11361]: step time server [ip-address] offset 64.503042 sec
[root@atw-dogtag01 ~]# /etc/init.d/ntpd start
Starting ntpd (via systemctl):                             [  OK  ]
[root@atw-dogtag01 ~]#
```

Installing the LDAP Server

Dogtag uses an open source LDAP server called Directory Server to store its data. Before you can install Dogtag, Directory Server must be installed and prepared:

Step 1. Install the LDAP server package with the **yum install 389-ds** command.

Step 2. Create a new user named **ds389** to be used by the Directory Server.

Both steps are shown in Example C-9.

Example C-9 Installing Directory Server and Creating the Service Account

https://t.me/learningnets

Click here to view code image

```
[root@atw-dogtag01 ~]# yum install 389-ds
[root@atw-dogtag01 ~]# useradd ds389
```

Step 3. Launch the Directory Server configuration wizard using the **setup-ds.pl** script located in /usr/sbin/setup-ds.pl, as seen in Example C-10.

Example C-10 Launching the Setup Script

Click here to view code image

```
[root@atw-dogtag01 ~]#  /usr/sbin/setup-ds.pl
```

Step 4. Accept the defaults. When you reach the portion where the wizard is asking for a system user, you will need to change the default (nobody) to the ds389 user. Use the ds389 user for the group as well, as shown in Example C–11.

Example C-11 Setting the System User and Group to ds389

Click here to view code image

```
==============================================================================
The server must run as a specific user in a specific group.
It is strongly recommended that this user should have no privileges
on the computer (i.e. a non-root user). The setup procedure
will give this user/group some permissions in specific paths/files
to perform server-specific operations.

If you have not yet created a user and group for the server,
create this user and group using your native operating
system utilities.

System User [nobody]: ds389
System Group [nobody]: ds389
```

Step 5. Set the password for the Directory Manager, as shown in Example C-12.

Example C-12 Setting the Directory Manager password and Success Message

Click here to view code image

```
Directory Manager DN [cn=Directory Manager]:
Password:
Password (confirm):
Your new DS instance 'atw-dogtag01' was successfully created.
Exiting . . .
Log file is '/tmp/setupo0Vx6g.log'
```

Installing the PHP Services

Step 1. Use **yum** to install PHP, as shown in Example C-13.

Example C-13 Installing PHP with yum

Click here to view code image

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# yum install php
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package php.i686 0:5.3.13-1.fc15 will be installed
--> Processing Dependency: php-common(x86-32) = 5.3.13-1.fc15 for package: php-
5.3.13-1.fc15.i686
--> Processing Dependency: php-cli(x86-32) = 5.3.13-1.fc15 for package: php-
5.3.13-1.fc15.i686
--> Running transaction check
---> Package php-cli.i686 0:5.3.13-1.fc15 will be installed
---> Package php-common.i686 0:5.3.13-1.fc15 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

================================================================================
 Package            Arch         Version                   Repository       Size
================================================================================
Installing:
 php                i686         5.3.13-1.fc15             updates         1.1 M
Installing for dependencies:
 php-cli            i686         5.3.13-1.fc15             updates         2.2 M
 php-common         i686         5.3.13-1.fc15             updates         547 k

Transaction Summary
================================================================================
Install       3 Package(s)

Total download size: 3.9 M
Installed size: 13 M
Is this ok [y/N]: y
Downloading Packages:

Running Transaction
  Installing : php-common-5.3.13-1.fc15.i686                                1/3
  Installing : php-cli-5.3.13-1.fc15.i686                                   2/3
  Installing : php-5.3.13-1.fc15.i686                                       3/3

Installed:
  php.i686 0:5.3.13-1.fc15

Dependency Installed:
  php-cli.i686 0:5.3.13-1.fc15          php-common.i686 0:5.3.13-1.fc15

Complete!
[root@atw-dogtag01 ~]#
```

Step 2. Start the Apache (httpd) and Directory Server (dirsrv) services, and configure them to start on bootup, as shown in Example C-14.

Example C-14 Starting the Apache and Directory Server Services

Click here to view code image

```
[root@atw-dogtag01 ~]# service httpd start
Starting httpd (via systemctl):                            [  OK  ]
[root@atw-dogtag01 ~]# service dirsrv start
Starting dirsrv:
    atw-dogtag01... already running                        [  OK  ]
[root@atw-dogtag01 ~]# chkconfig dirsrv on
[root@atw-dogtag01 ~]# chkconfig httpd on
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]#
```

Installing and Configuring Dogtag

Now that all the perquisite services are installed and prepared, you are ready to install the Dogtag certificate authority.

Install the Dogtag CA with the **yum install pki-ca** command, as shown in Example C-15.

Example C-15 Installing Dogtag

Click here to view code image

```
[root@atw-dogtag01 ~]# yum install pki-ca
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package pki-ca.noarch 0:9.0.20-1.fc15 will be installed
--> Processing Dependency: pki-selinux = 9.0.20-1.fc15 for package:
pki-ca-9.0.20-1.fc15.noarch
--> Processing Dependency: pki-common = 9.0.20-1.fc15 for package: pki-ca-9.0.20-1.
fc15.noarch
--> Processing Dependency: pki-ca-theme >= 9.0.0 for package: pki-ca-9.0.20-1.fc15.
noarch
--> Running transaction check
---> Package dogtag-pki-ca-theme.noarch 0:9.0.11-1.fc15 will be installed
--> Processing Dependency: dogtag-pki-common-theme = 9.0.11-1.fc15 for package:
dogtag-pki-ca-theme-9.0.11-1.fc15.noarch
---> Package pki-common.noarch 0:9.0.20-1.fc15 will be installed
```

Modifying the Firewall Rules (iptables)

To connect to the Dogtag service on the ports used with Dogtag, you must modify the Linux server's host firewall (iptables) to allow the connections. Obviously, the best choice is to put specific rules into iptables that will allow the communication on the ports. However, to keep this simple, these instructions walk you through disabling iptables:

Step 1. Stop the firewall service with the **service iptables stop** command.

Step 2. Keep the firewall from starting when the server is booted with the **chkconfig iptables off** command.

Example C-16 Shutting Off the Firewall

Click here to view code image

```
[root@atw-dogtag01 ~]# service iptables stop
Stopping iptables (via systemctl):                         [  OK  ]
[root@atw-dogtag01 ~]# chkconfig iptables off
[root@atw-dogtag01 ~]#
```

Creating a New CA Instance

Now that Dogtag is installed, you need to create a new CA instance. The following uses ports that the author prefers to use. You can change any of the parameters in the following section to suit the needs of your organization:

Step 1. Create a PKI instance using the **pkicreate** command with the following options:

https://t.me/learningnets

- **pki_instance_root=/var/lib**—This sets the root location to store the pki instance. Based on the settings used in Example C-17, it will be placed in directory **/var/lib/ise-ca**.
- **pki_instance_name=ise-ca**—This names the new CA instance **ise-ca**. You can replace this with another name to suit the needs of your organization.
- **subsystem_type=ca**—This sets the subsystem to be a CA. Other possible subsystems are not applicable to this guide.
- **agent_secure_port=9443**—Agent services are where an administrator can see which certificate has been provisioned, revoke certificates, and so on.
- **ee_secure_port=9444**—This sets the SSL port for EnE-Entities web services.
- **ee_secure_client_auth_port=9446**—This sets the SSL port for EnE-Entities authentication.
- **admin_secure_port=9447**—This is the default port to use to access the CA Services Page as the administrator.
- **unsecure_port=9180**—This sets the regular port number. When not specified, it will be randomly generated.
- **tomcat_server_port=9701**
- **user=pkiuser**
- **group=pkiuser**
- **redirect conf=/etc/ise-ca**—This configures the configuration data to be stored in /etc/ise-ca.
- **redirect logs=/var/log/ise-ca**—This configures the logs to be in the /var/log/ise-ca directory.
- **verbose**—This sets the install to be in verbose mode, to provide you with as much detail as possible.

Example C-17 Creating the PKI Instance

Click here to view code image

```
pkicreate    -pki_instance_root=/var/lib -pki_instance_name=ise-ca -subsystem_
type=ca -agent_secure_port=9443 -ee_secure_port=9444 -ee_secure_client_auth_
port=9446 -admin_secure_port=9447 -unsecure_port=9180 -tomcat_server_port=9701
-user=pkiuser -group=pkiuser -redirect conf=/etc/ise-ca -redirect logs=/var/log/
ise-ca -verbose
```

Note

If you run the **pkicreate** command as provided here, are given an admin GUI URL that does NOT contain the FQDN, and are unable to connect to the GUI, you might be missing a DNS entry for your FQDN. It is imperative that this DNS entry exists, either in the DNS server given in the Linux configuration or as part of the /etc/hosts file. After remedying the absence of DNS entry, run the **pkiremove** command as provided here:

Click here to view code image

```
pkiremove -pki_instance_root=/var/lib -pki_instance_name=ise-ca
```

After removing the erroneous PKI instance, rerun the **pkicreate** command as provided in Example C-17:

https://t.me/learningnets

Step 2. Proceed with the graphical configuration of the Dogtag CA.

After the setup script completes running, a message is displayed with a unique URL to access the Dogtag GUI and complete the CA installation, as shown in Example C-18 in highlighted text.

Example C-18 Example of Unique URL to Dogtag GUI

Click here to view code image

```
Installation information recorded in /var/log/ise-ca-install.log.
[debug] run_command(/sbin/service pki-cad restart ise-ca)
Before proceeding with the configuration, make sure
the firewall settings of this machine permit proper
access to this subsystem.

Please start the configuration by accessing:

https://atw-dogtag01.ise.local:9447/ca/admin/console/config/login?pin=UUVMDHRvTojQ
rdeod91e

After configuration, the server can be operated by the command:

    /sbin/service pki-cad restart ise-ca
```

Step 3. Click **Next** from the Welcome screen, as shown in Figure C-4.



Figure C-4 The Dogtag certificate system Welcome screen.

Step 4. Create a new security domain. Name that security domain **ISE BYOD Domain** and click **Next**, as shown in Figure C-5.

https://t.me/learningnets

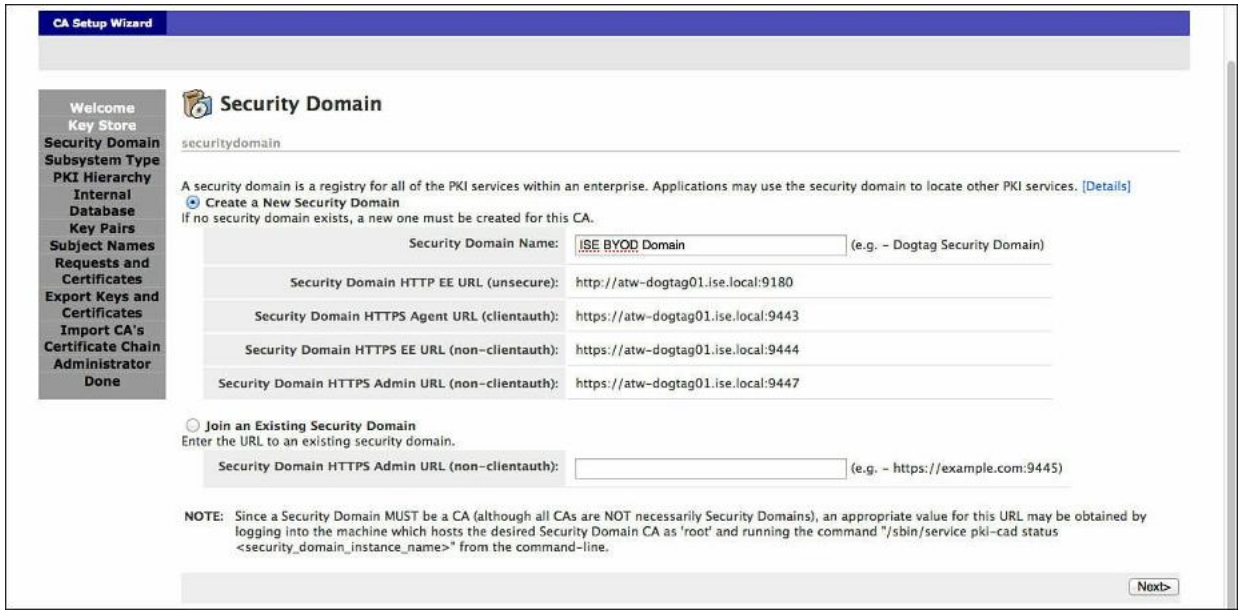


Figure C-5 Create the security domain.

Step 5. Name the subsystem **Certificate Authority** and click **Next**, as shown in Figure C-6.

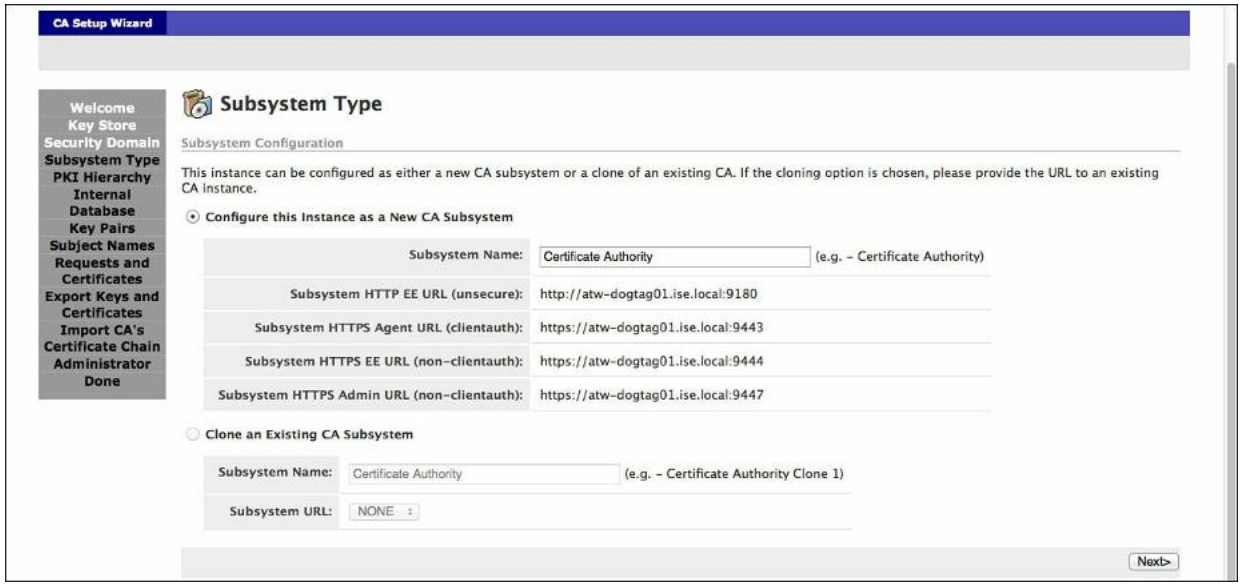


Figure C-6 The subsystem name is Certificate Authority.

Step 6. Select **Make This a Self-Signed Root CA Within This New PKI Hierarchy**, as shown in Figure C-7.

https://t.me/learningnets



Figure C-7 Make this a self-signed root CA.

Of course, this could become a subordinate CA of an existing CA. However, that is not the focus of this appendix.

Step 7. The internal database is the Directory Server (ds389) that we installed earlier. All the settings should be filled in correctly, as shown in Figure C-8. Add the Directory Manager password created earlier in Example C-12.

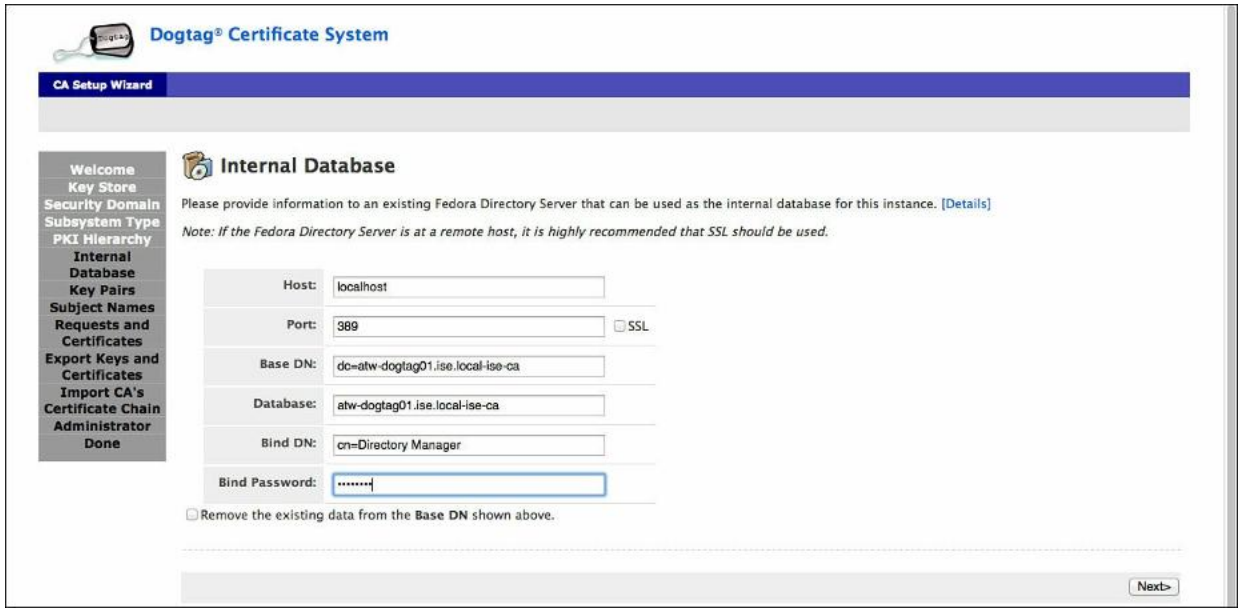


Figure C-8 Add the directory manager password.

Step 8. Generate the key pairs. The default of RSA with SHA256 and a key size of 2048 bits will work fine. Click **Next**, as shown in Figure C-9. As always, please reference your corporate security policy to ensure that these values are sufficient.

https://t.me/learningnets

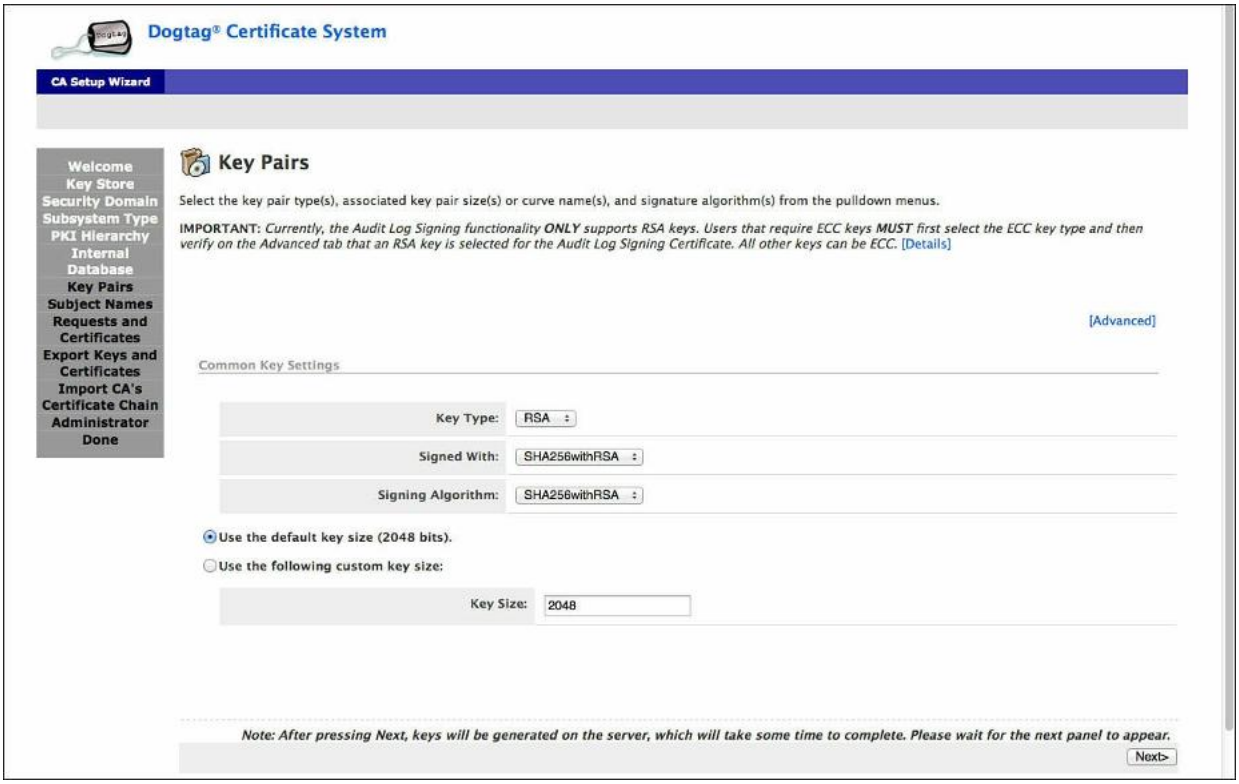


Figure C-9 Generate the key pairs.

Step 9. The certificate subject names can be left at their default values, as shown in Figure C-10. Click **Next**.

https://t.me/learningnets



Figure C-10 Default subject names.

Step 10. If an action is needed, it will appear in red. If no actions are required, click **Next**, as shown in Figure C-11.

https://t.me/learningnets

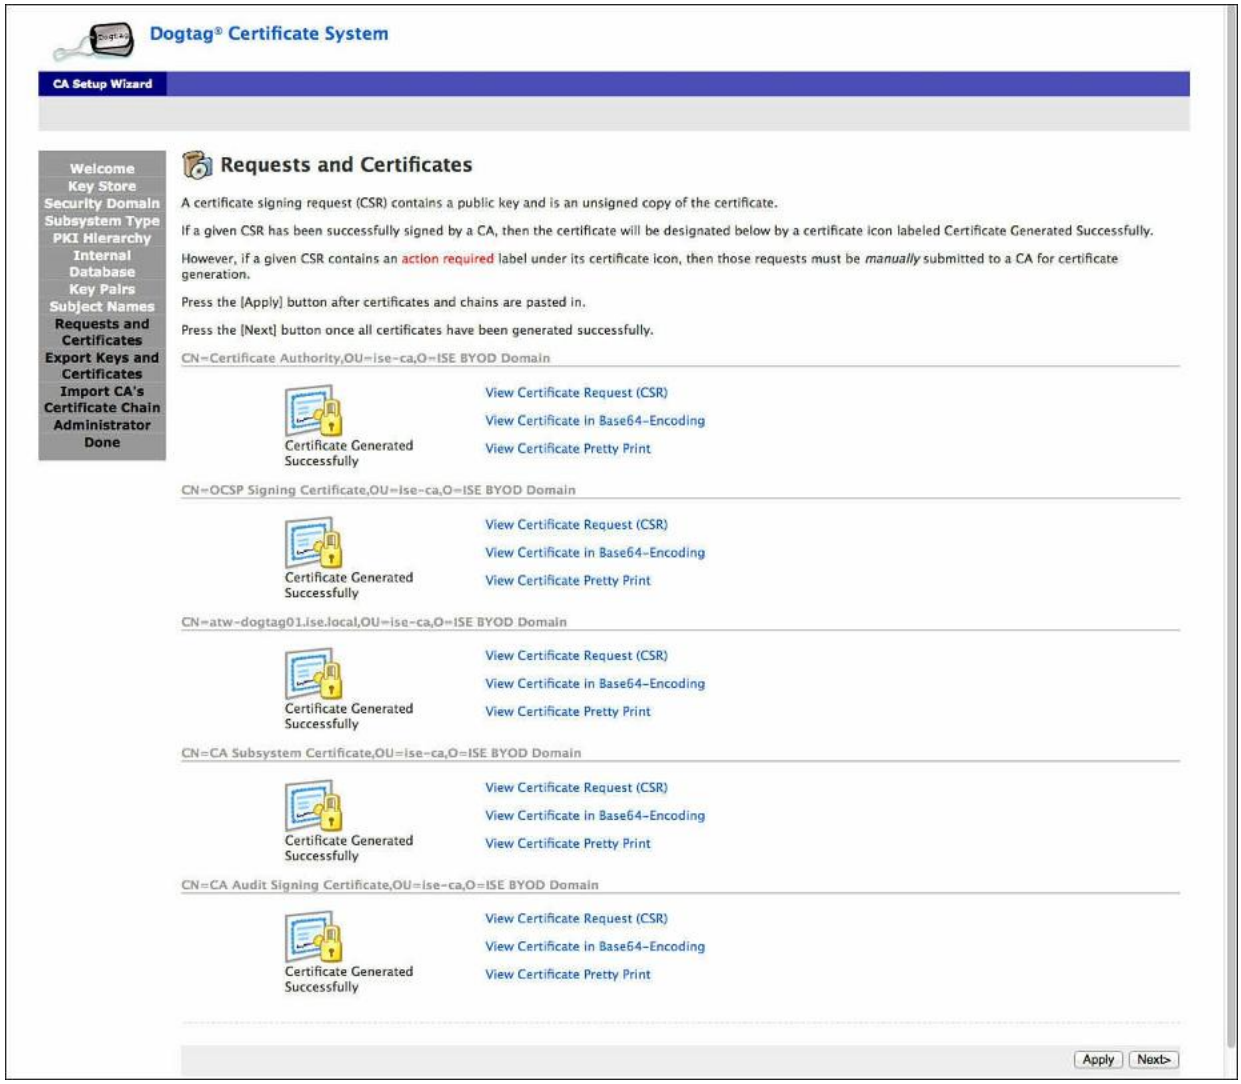


Figure C-11 Required actions will be in red.

Step 11. Provide a password and export the CA's key pair, as shown in Figure C-12. Store the key pair in a secure location. If you must restore this CA from a backup, this key pair will be required. Furthermore, compromise of this key pair would be equivalent to a compromise of this entire PKI infrastructure. Protect this key pair with the utmost care.

https://t.me/learningnets

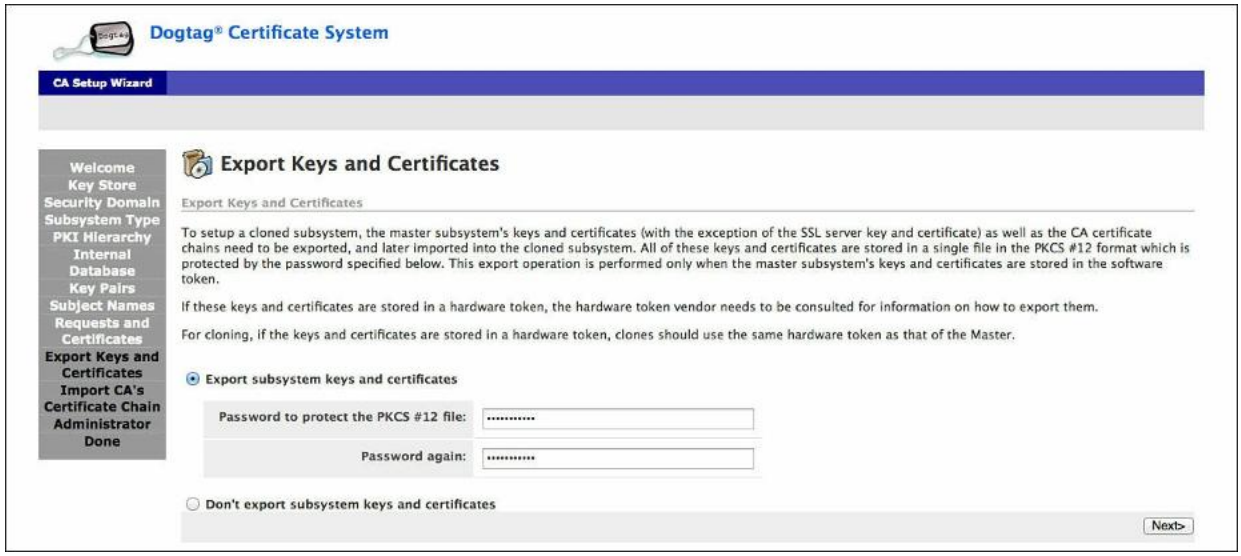


Figure C-12 Export and securely store the key pairs.

Step 12. The new root CA certificate will be imported into your browser or your system certificate store. You can select for which purposes the certificate should be trusted, as shown in Figure C-13.

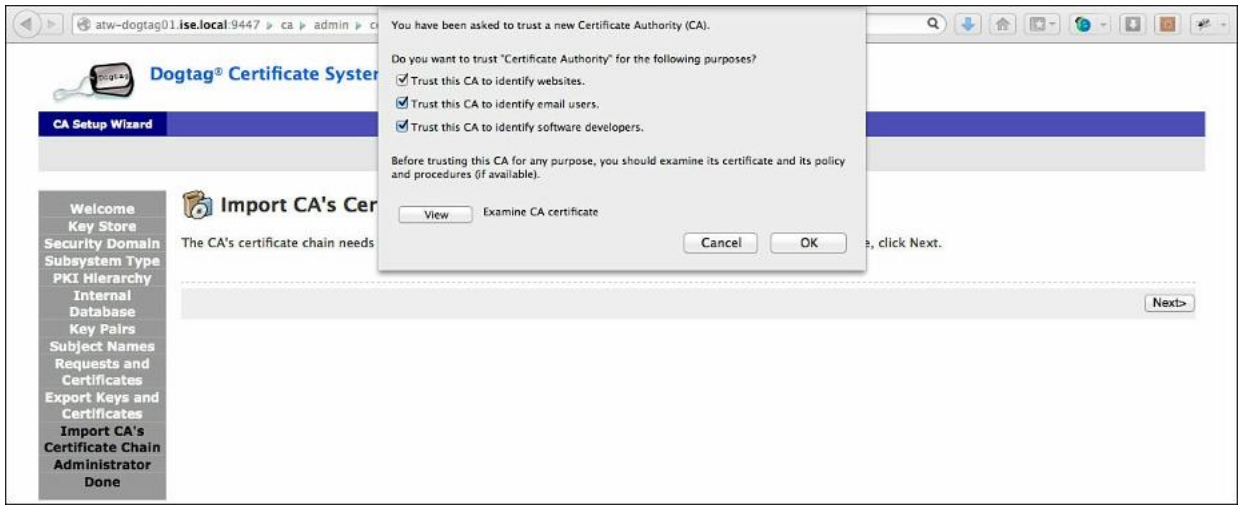


Figure C-13 Trust the certificate.

Step 13. Create an administrative certificate to identify you (the administrator) to the CA for administrative purposes, as shown in Figure C-14.

https://t.me/learningnets

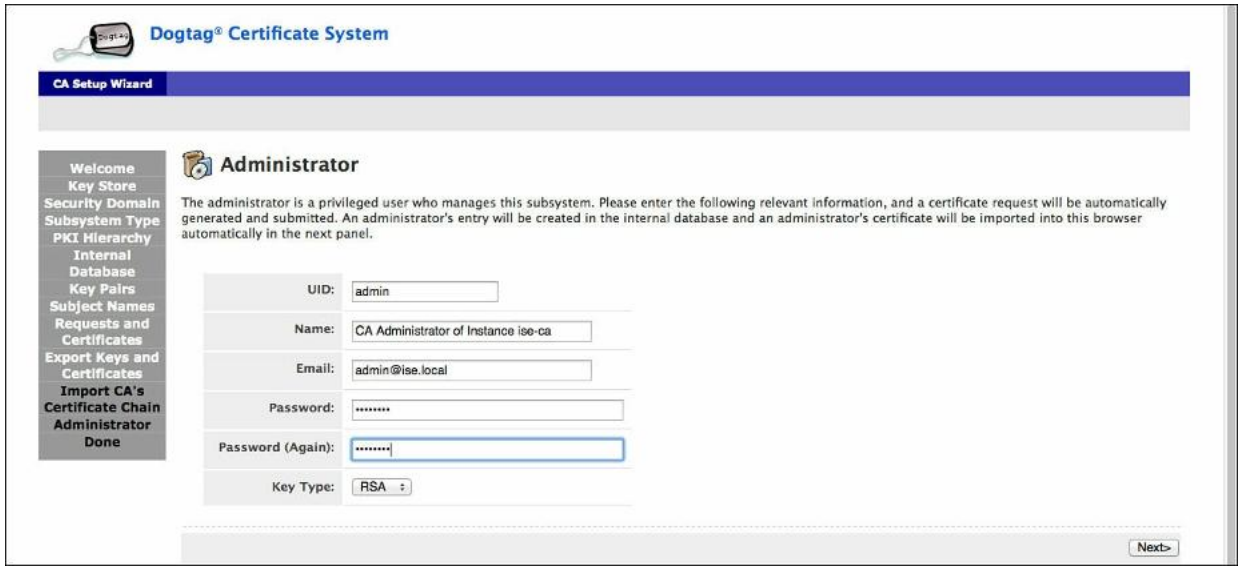


Figure C-14 Administrator certificate.

Step 14. You must now install that certificate into your browser, as shown in Figure C-15, so your browser will identify you when connecting to the CA. Be sure you back up and store this key in a secure location because you will not be able to administer the CA without this certificate.

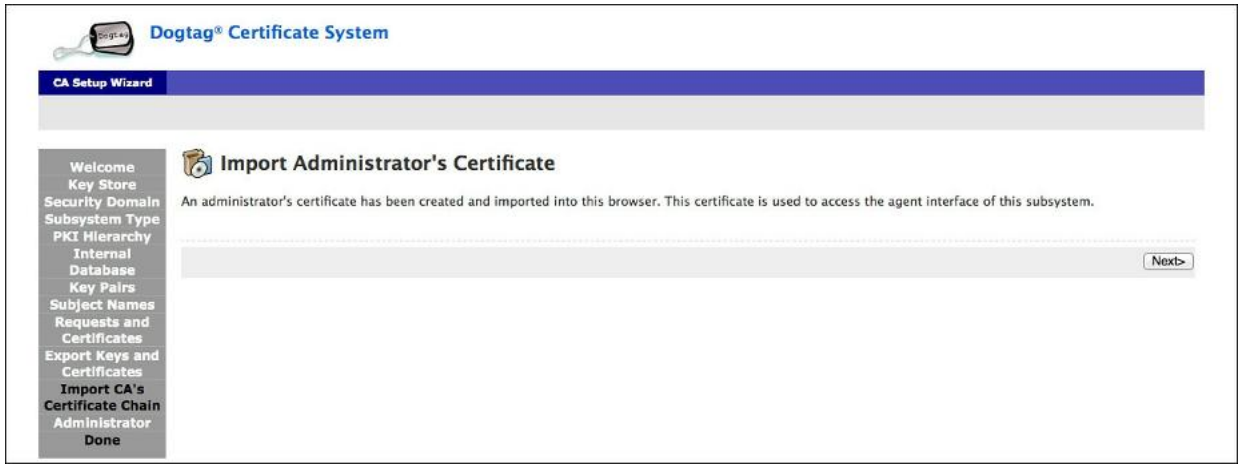


Figure C-15 Import the administrator certificate.

Step 15. You are finished with the GUI-based configuration, as shown in Figure C-16.

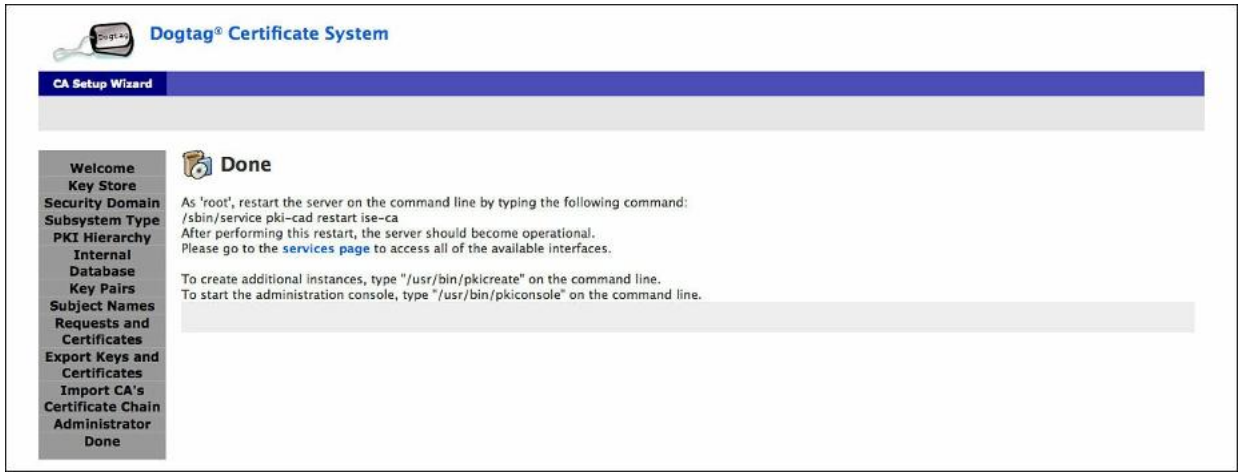


Figure C-16 Done.

https://t.me/learningnets

> **Note**
>
> Although the GUI configuration is complete, we are not ready to begin using the CA just yet. We still need to add a custom script and modify some more configuration files.

Step 16. After configuration, the server cannot be operated until the **service pki-cad restart ise-ca** command has been run, shown in Example C-19.

Example C-19 Restart the CA

Click here to view code image

```
[root@atw-dogtag01 ~]# /sbin/service pki-cad restart ise-ca
Stopping ise-ca:                                           [  OK  ]
Starting ise-ca:                                           [  OK  ]
[root@atw-dogtag01 ~]#
```

Step 17. You can access the administrative interface, but you're not ready to add it to ISE yet.

Enabling and Configuring SCEP

In this procedure, we will enable and configure Simple Certificate Enrollment Protocol (SCEP) by directly modifying the CS.cfg file:

Step 1. Back up the CS.cfg file before making any changes, as shown in Example C-20.

Example C-20 Backup of the CS.cfg File

Click here to view code image

```
[root@atw-dogtag01 ~]# cp /etc/ise-ca/CS.cfg /etc/ise-ca/CS.cfg.bak
```

Step 2. Open the CS.cfg file in a text editor, such as vi, as shown in Example C-21.

Example C-21 Edit the CS.cfg File

Click here to view code image

```
[root@atw-dogtag01 ~]# vi /etc/ise-ca/CS.cfg
```

Step 3. Add the following lines to the bottom of the CS.cfg file, and save the changes:

Click here to view code image

```
ca.scep.allowedEncryptionAlgorithms=DES3
ca.scep.allowedHashAlgorithms=SHA1,SHA256,SHA512
ca.scep.enable=true
ca.scep.encryptionAlgorithm=DES3
ca.scep.hashAlgorithm=SHA256
ca.scep.nonceSizeLimit=16
```

Step 4. Back up the caRouterCert.cfg file before making any changes, as shown in Example C-22.

Example C-22 Backing Up the caRouterCert.cfg File

Click here to view code image

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# cp /var/lib/ise-ca/profiles/ca/caRouterCert.cfg /var/lib/
ise-ca/profiles/ca/caRouterCert.cfg.bak
```

Step 5. Edit the caRouterCert.cfg file using a text editor. Delete the value for the variable **auth.instance_id**, and save your changes. Example C-23 shows the editing of the file, and the end result should look like Example C-24.

Example C-23 Edit the caRouterCert.cfg File

Click here to view code image

```
[root@atw-dogtag01 ~]# vi /var/lib/ise-ca/profiles/ca/caRouterCert.cfg
```

Example C-24 The Final Setting for auth.instance_id= field in the caRouterCert.cfg File

Click here to view code image

```
[root@atw-dogtag01 ise-ca]# cat /var/lib/ise-ca/profiles/ca/caRouterCert.cfg
desc=This certificate profile is for enrolling router certificates.
visible=false
enable=true
enableBy=admin
auth.instance_id=
name=One Time Pin Router Certificate Enrollment
input.list=i1,i2
input.i1.class_id=certReqInputImpl
input.i2.class_id=submitterInfoInputImpl
output.list=o1
<<SNIP>>
```

Step 6. Restart the CA services with the **service pki-cad restart** command, as shown in Example C-25.

Example C-25 Restart the CA Services

Click here to view code image

```
[root@atw-dogtag01 ise-ca]# service pki-cad restart
Stopping ise-ca:                                           [  OK  ]
Starting ise-ca:                                           [  OK  ]
```

Preparing Apache

Step 1. Move the Apache welcome.conf file to disable the default installation, as shown in Example C-26.

Example C-26 Move the welcome.conf File

Click here to view code image

```
[root@atw-dogtag01 ise-ca]# mv /etc/httpd/conf.d/welcome.conf /etc/httpd/conf.d/
welcome.conf.bak
```

Step 2. Create a new file called scepproxy.php at /var/www/html, as shown in Example C-27.

https://t.me/learningnets

Example C-27 Creating the scepproxy.php File

Click here to view code image

```
[root@atw-dogtag01 ise-ca]# vi /var/www/html/scepproxy.php
```

Step 3. Populate the file with the following PHP script, and save the file when completed. Example C-28 shows the full PHP script.

Example C-28 Contents for the scepproxy.php File

Click here to view code image

```
<?php
$ops = $_GET['operation'];
$msg= $_GET['message'];
$order   = array("\r\n", "\n", "\r", " ");
$msg = str_replace($order, "", $msg);
$msg = rawurldecode($msg);
$msg=urlencode($msg);
if ($ops == "GetCACaps")
{
echo "";
}
else
{
$url = "http://127.0.0.1/ca/cgi-bin/pkiclient.exe?operation=".$ops."&message=".$
msg;
$ch = curl_init();
curl_setopt($ch, CURLOPT_PORT, 9180);
curl_setopt($ch, CURLOPT_URL, $url);
curl_setopt($ch, CURLOPT_HEADER, 0);
curl_setopt($ch, CURLOPT_RETURNTRANSFER, 1);
//curl_setopt($ch, CURLOPT_POST, 1);
$body = curl_exec($ch);
curl_close($ch);
if ($ops=="PKIOperation")
{
header("Content-Type: application/x-pki-message");
}
else
{
header("Content-Type: application/x-x509-ca-cert");
}
echo $body;
}
?>
```

Step 4. Restart the Apache service to reflect your changes with the **service httpd restart** command, as shown in Example C-29.

Example C-29 Restarting the Web Server

Click here to view code image

```
[root@atw-dogtag01 ise-ca]# service httpd restart
Restarting httpd (via systemctl):                          [  OK  ]
[root@atw-dogtag01 ise-ca]#
```

https://t.me/learningnets

The Dogtag installation is complete. You are ready to add this CA to ISE for BYOD certificate provisioning.

Configuring ISE to Use the New Dogtag CA

This document is assuming that you already have your BYOD policies ready, or you will create them afterwards. In this section, we will focus on the simple task of adding the new Dogtag CA to ISE for purposes of SCEP provisioning the BYOD certificates.

Adding Dogtag to the SCEP RA Profiles

From the ISE administrative GUI, we will add the Dogtag server to the SCEP RA profiles:

Step 1. Navigate to **Administration > System > Certificates > SCEP RA Profiles**, and click **Add**.

Step 2. Name the RA **DogTag** and enter a description.

Step 3. Enter the Dogtag server URL of http://<server_name>/scepproxy.php, as shown in Figure C-17.



Figure C-17 SCEP profile configuration.

Step 4. Click **Test Connectivity**.

Step 5. Click **Submit**.

You are finished and ready to use Dogtag as the CA for BYOD onboarding.

https://t.me/learningnets

Appendix D. 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 chapters of this book.

Many commands will be identical with each platform because there are only minor variances between the switch platforms and the features they support. For specific explanations of each command, please see Chapter 12, "Implement Wired and Wireless Authentication."

Catalyst 2960/3560/3750 Series, 12.2(55)SE

With this model and IOS version, there is no device sensor. Therefore, the SNMP commands are there to allow ISE to do profiling via probing the switch with SNMP.

Click here to view code image

```
3560-X# show 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
!
!Dynamic Authorization is the official term for CoA, which allows a RADIUS server
to send commands to the NAD to change the authorization for the endpoint
aaa server radius dynamic-author
 client 10.1.103.232 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 is required in order to insert the client IP address into
downloadable ACLs (dACLs).
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
```

https://t.me/learningnets

```
!
interface Loopback0
 ip address 192.168.254.60 255.255.255.255
!
interface range <ALL EDGE PORTS>
switchport access vlan 41
! VLAN 41 is the data VLAN for this particular switch.  This will vary in your
environment.
 switchport mode access
 switchport voice vlan 99
! VLAN 99 is the voice VLAN in this particular switch. This will vary in your
environment
 ip access-group ACL-ALLOW in
 authentication event fail action next-method
 authentication event server dead action authorize
 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
!
```

https://t.me/learningnets

```
ip radius source-interface Loopback0
! The following logging commands should be used ONLY when troubleshooting,
or doing an initial roll-out.  Sending all the syslog to the MNT node can be
overwhelming and affect performance and storage.
See Chapter 22, "Troubleshooting Tools," for more information.
logging origin-id ip
logging source-interface Loopback0
logging host 10.1.103.4 transport udp port 20514
! Cisco ISE Monitoring uses a nontraditional port for Syslog.  It uses 20514
instead of 514.
!
! The following SNMP commands are used because this switch version does not have
Device Sensor capabilities.
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
!  The definitions of these attributes and their importance are covered in detail
in Chapter 12, "Implement Wired and Wireless Authentication."
radius-server dead-criteria time 5 tries 3
radius-server host 10.1.103.232 auth-port 1812 acct-port 1813 key Cisco123
radius-server vsa send accounting
radius-server vsa send authentication
!
end
```

Catalyst 3560/3750 Series, 15.0(2)SE

Depending on your hardware version of your 3750 or 3560 catalyst switch, you might be able to upgrade to 15.x code. With this model and IOS version, you will gain the device sensor capability described in detail in Chapter 15, "Profiling." Therefore, the SNMP commands will not be used, and the switch is using device sensor commands instead. Additionally, version 15.x support both IPv4 and IPv6, so the RADIUS server host definitions are very different from the 12.x version configuration.

Click here to view code image

```
C3750X# show 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
```

https://t.me/learningnets

```
!
!
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 Sensors enables ISE to gather raw endpoint information and send that
information to ISE in a RADIUS accounting packet.
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
!
dot1x system-auth-control
!
interface Loopback0
 ip address 192.168.254.1 255.255.255.255
!
interface <ALL EDGE PORTS>
 switchport access vlan 10
 ! VlAN 10 is the data VLAN for this particular switch.  It may differ in your
environment.
 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
```

https://t.me/learningnets

```
 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
!
! There are more VLANs on this switch than the 12.x version above, because it was
in a different part of the network with different requirements.  The author has
left the configuration this way to show that these are but sample configurations,
and you have many options.
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
```

https://t.me/learningnets

```
 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
!
!
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-02
 address ipv4 10.1.100.232 auth-port 1812 acct-port 1813
 automate-tester username radius-test
 key Cisco123
!
end
```

Catalyst 4500 Series, IOS-XE 3.3.0/15.1(1)SG

The Catalyst 4500 chassis-based switch has a supervisor engine that will run IOS-XE. With this model and IOS version, you gain the device sensor capability described in detail in Chapter 15. Therefore, the SNMP commands will not be used, and the switch is using device sensor commands instead.

Click here to view code image

```
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 10.1.103.232 server-key Cisco123
!
aaa session-id common
!
ip domain-name ise.local
!
ip device tracking
!
```

https://t.me/learningnets

```
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
 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
```

https://t.me/learningnets

```
 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
!
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 3radius-server vsa send accounting
radius-server vsa send authentication
!
!
radius server CP-02
 address ipv4 10.1.100.232 auth-port 1812 acct-port 1813
 automate-tester username radius-test
 key Cisco123
!


!
end
```

Catalyst 6500 Series, 12.2(33)SXJ

The Catalyst 6500 chassis-based switch has a supervisor engine that will run the IOS 12.2(33)SX versions. With this model and IOS version, you will need to use SNMP for profiling because there is no device sensor. Additionally, there is no support for critical voice.

Click here to view code image

```
hostname 6503
```

https://t.me/learningnets

```
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 10.1.103.232 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
!
```

https://t.me/learningnets

```
!
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 10.1.103.232 auth-port 1812 acct-port 1813 key Cisco123
radius-server vsa send accounting
radius-server vsa send authentication
!
end
```

https://t.me/learningnets

Glossary

A

Active Directory (AD) Microsoft Active Directory is an example of an external identity store that can be used with ISE authentication and authorization policies. ISE supports Microsoft Server 2003, 2008, 2008 R2, and 2012.

AnyConnect NAM AnyConnect Network Access Manager (NAM) is Cisco's Enterprise-class supplicant. It is an evolution from the acquisition of Meeting House, which became the Cisco Security Services Client (CSSC), and is now a module within AnyConnect.

AnyConnect DART AnyConnect Diagnostics and Reporting Tool (DART) is a module of Cisco's AnyConnect Client that is used to collect detailed logs related to all aspects of AnyConnect, including the NAM supplicant.

Authenticator The official name for the device that is communicating with the 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.

Authentication Serve The device that validates the EAP credentials from the supplicant and makes the authentication and authorization decisions.

C

CDP Second Port Disconnect A Cisco proprietary method of an IP phone using 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.

Certificate Meaning an X.509 certificate. It is an electronic document that has been signed for authenticity.

Certificate authentication profile (CAP) When using the EAP-TLS protocol for user authentication with certificates, you need to configure a certificate authentication profile. The CAP will define which certificate field to use as the username attribute for the ISE authentication process. Used to extract the principle identity from a field in the certificate.

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 certificate authority provider.

Certificate revocation list (CRL) A signed list of the revoked certificate serial number. The CA publishes the CRL via a website that can be retrieved by a device or application needing to validate whether a digital certificate has been revoked by the CA administrator.

Change of authorization (CoA) Described in RFC 3576 and its successor RFC 5176, change of authorization (CoA) is a process whereby a RADIUS server can disconnect or otherwise force a change in the authorization status of a user or endpoint. This can be accomplished via a port bounce or reauthentication. (CoA) is an enhancement to RADIUS, defined by RFC 3576 and superseded by RFC 5176, which allows the authentication server to initiate a command into the network access device and perform a change to the authentication session.

Client provisioning portal (CPP) A portal on ISE that enables an endpoint to install a NAC agent as well as any other required supplicants.

Cisco NAC agent The Cisco NAC agent is a software application that runs on an endpoint, providing detailed information about the endpoint, including the operating system, patch level, antivirus, and

https://t.me/learningnets

antispyware. This application is required for proper posture assessment.

Classification As it relates to TrustSec, classification is the process of assigning the SGT to a user's or endpoint's network session.

Command-line interface (CLI) A text-based method used to access ISE or other device. When using a CLI, most interaction with the device occurs using solely the keyboard.

Common access card (CAC) Common access card is another term used to refer to a smart card, which is an embedded integrated circuit that can be used for two-factor authentication. The smart cards store a set of X.509 encrypted digital certificates used for authentication processes. A typical example of a CAC is an employee's identification badge.

Compound condition Compound conditions are policy objects that contain 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.

CWA Centralized web authentication (CWA) is where the authenticator is only aware of a MAB authentication, while the authentication server (ISE) houses the web portal and maintains the sync between the user's credentials and the MAB authentication session.

D

Device administration Controlling access to who can log in to a network device console, telnet session, SSH session, or other method. Although it can often seem similar to network access AAA, it has a completely different purpose and requires different policy constructs.

E

EAP Extensible Authentication Protocol (EAP). An authentication framework that defines the transport and usage of identity credentials.

EAPoL-Proxy-Logoff A standards-based way an IP phone sends 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.

EAP over LAN Also known by the standard for EAP over LAN: IEEE 802.1X. This is a standard for port-based network access control for local area and metropolitan area networks.

East-West Refers to the enforcement of traffic between two or more endpoints in the same layer of the network.

External identity store An external identity store is an identity store configured to use an external database server for the authentication and/or authorization policies. Examples include Active Directory, LDAP, RADIUS token servers, RSA SecureID, and certificate authentication profiles.

F

Flex-Auth Flexible Authentication (Flex-Auth) is a capability of a Cisco switch interface that enables a network administrator to set an authentication order and priority on the switchport, thereby allowing the port to attempt 802.1X, MAB, and then WebAuth in 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.

https://t.me/learningnets

G

Graphical user interface (GUI) A graphic-based method used to access ISE or other device. Using a GUI usually requires a mouse to point and click the various configuration or administrative components of the device. The keyboard is still used to populate text-based fields within the GUI.

I

Identity source Each individual identity store is referred to as an identity source.

Identity source sequence When multiple identity sources are combined in a top-to-bottom sequence list for ISE authentication or authorization processes, this is referred to as an identity source sequence.

Identity store Internal and/or external database(s) that can be used to authenticate a user's or an endpoint's credentials.

Inline Posture Node A gatekeeper between the network access device and the rest of the network. This node will help determine the compliance of the endpoint to the company security policy by checking the status of antivirus, antispyware, patch level, and other critical system parameters.

Internal endpoint database The ISE internal endpoint database is used to store information about all the endpoint devices that have connected to the ISE infrastructure. The internal endpoint database will store the endpoints' attributes including MAC address, IP address, and various other attributes learned using the ISE profiling probes.

Internal User Database The ISE internal user database is used for the creation of internal username and password accounts. These accounts are typically referred to as internal user accounts. The internal user database can also be used for ISE authentication and authorization policies. ISE guest services is one example for the use of an internal user database.

K

Key-pair The combination of a public and private key used for asymmetric encryption.

L

Lightweight Directory Access Protocol (LDAP) LDAP is a TCP/IP client model protocol to connect and retrieve information from an X.500 directory service databases using TCP port 389. A secure version of LDAP is also supported using secure sockets layer (TLS/SSL) using TCP port 636. LDAP is defined in RFC 2251 and can be used as an external identity store.

LWA Local web authentication (LWA) is where the authenticator receives the credentials from the web portal and sends those credentials to the authentication server through a RADIUS Access-Request. The portal can be hosted on the authenticator itself or even on an external server.

M

MAB MAC authentication bypass (MAB) is the process of an authenticator sending the RADIUS Access-Request on behalf of the endpoint that doesn't have a supplicant. The Access-Request includes the MAC address of the endpoint as the credential.

MDA Multiple domain authentication (MDA) is an enhancement to the default mode of 802.1X allowing for two MAC addresses, one in the voice domain and the other in the data domain.

Mobile device manager (MDM) A security software system that enables an administrator to configure and secure a mobile device, regardless of where the device is located. MDMs can provide application provisioning, security posturing, as well as remote wipe capabilities. MDM solutions are used to configure,

https://t.me/learningnets

manage, secure, and control mobile platforms such as iPads, iPhones, Android devices, and more.

Monitoring and Troubleshooting (MNT) Node The ISE persona that provides monitoring and troubleshooting functions for the Cisco ISE deployment.

Multi-Auth Multiple authentication (Multi-Auth) mode is an extension to MDA that permits a virtually unlimited number of MAC addresses per interface, with a separate authentication session per MAC address.

Multi-Host Multi-Host mode is an extension to MDA that permits a virtually unlimited number of MAC addresses per interface. All endpoints attached to the interface will utilize the authorization results of the first MAC address to authenticate successfully. No other hosts will be able to use 802.1X authentications.

N

NAD Network access device (NAD). The network device where endpoints connect to enter the network. This is the 802.1X authenticator and is usually referring to a LAN switch, Wireless LAN Controller, or Cisco ASA VPN headend.

Native-EAP Specific methods of Extensible Authentication Protocol. Examples are EAP-MD5, EAP-MSCHAPv2, and EAP-TLS.

Network access Securing network access can provide the identity of the device or user before permitting the entity to communicate with the network. Network access is the type of AAA that is most focused on in this exam and book.

Network Time Protocol (NTP) Networking Time Protocol is a time synchronization networking protocol that ensures 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 UDP port 123.

Node groups Groupings of Layer-2 adjacent (same VLAN) PSNs, where the PSNs will 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.

North-South Refers the enforcement of traffic from the network into the data center.

O

Onboarding Refers to the ability to configure an endpoint for access to a network, including the Wi-Fi configuration, device registration, and possibly identity certificates.

One-time password (OTP) An OTP is an example of a two-factor authentication process. 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 her secret PIN, which will then generate a unique one-time password for the user's authentication login process.

Online Certificate Status Protocol (OCSP) Online Certificate Status Protocol allows a device or an application to send real-time requests typically using the HTTP protocol. The OSCP process is a client request/server response process. OCSP is documented in IETF RFC 6960.

P

Personas The various functions within Cisco ISE that are required for proper operation of the platform.

Policy Admin 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 Service Node (PSN) The workhorse persona of the ISE deployment. This node is responsible for

https://t.me/learningnets

actively servicing RADIUS authentication and authorization requests.

Port security Port security is part of the Catalyst Integrated Security Features (CISF) of Cisco Catalyst switches, it and 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.

Principle Username X.509 Attribute The identity to be used for the user credential that has been extracted from a field in the 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 be decrypted only by the public key, and vice versa.

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 be decrypted only by the private key.

Public key infrastructure (PKI) The set of hardware, software, and security policies to enable the enrollment, storage, distribution, and management processes of digital certificates. PKI is based on the X.509 standard.

R

Registration authority (RA) The 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 allows the PKI deployment to be more scalable for large corporate PKI environments.

Remote Authentication Dial In User Service (RADIUS) RADIUS is 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 1812/1813 (IETF Revised Standard).

S

Security Group Exchange Protocol (SXP) A peering protocol used to communicate IP address-to-SGT bindings between SXP speakers and SXP listeners.

Security group tag (SGT) A 16-bit value that is assigned to the user's or endpoint's session upon login. The SGT represents the role or context of the user and device.

Simple condition Simple conditions are policy objects that are defined with a single attribute.

Single-host Single-host mode is the default mode of an 802.1X-enabled port, allowing only a single MAC address per switch interface.

Sponsor A sponsor is the individual responsible for creating and managing a guest user account.

Supplicant The software on an endpoint that understands how to communicate with EAP over LAN (802.1X).

T

Tunneled EAP An EAP method that first builds a secure tunnel between the supplicant and the authentication server. Native EAP communication occurs within the secure tunnel.

Transport As it relates to TrustSec, transport is the propagation of the SGT through the network, either natively in the Layer-2 frame or through the use of application layer protocols.

Trusted authority An entity whose signature is considered valid.

https://t.me/learningnets

TrustSec Also known as security group access (SGA), it is a next-generation access control enforcement and segmentation technology.

Two-factor authentication Two-factor authentication is a more secure method of providing authentication credentials. It requires two separate pieces of information for the user authentication to be successful. Typically, these are something you know and something you have to prove your identity. An example of two-factor authentication process is when you use an ATM to withdraw money from your checking account. You must have your physical bank ATM card plus your secret PIN code to prove your identity.

U

User identity group Internal User accounts are associated with a defined User Identity Group that will set the type of internal account being created.

V

VSA RADIUS is an extensible standard, and vendor-specific attributes (VSAs) are the method in which developers and various network access device vendors are able to extend RADIUS to perform other functions and carry information within RADIUS.

W

WebAuth Web Authentication (WebAuth) is the process of using an interactive web page for the end user to submit her credentials to the authentication server.

https://t.me/learningnets

Index

Numerics

802.1X, 23, 56-58
- components, 56
- computer authentication, 73
- intended behavior of, 680-681
- NADs, 63
- supplicants, 63-89
  - Cisco AnyConnect NAM supplicant, 75-88
  - devices without supplicants, 97-98
  - Windows native supplicant, 64-72
- user authentication, 72-73

A

AAA (authentication, authorization, and accounting), 16, 21
- device administration, 16, 21-22
  - command sets, 22
  - TACACS, 22
- network access, 16, 22-32
  - RADIUS, 22, 28-32
  - TACACS+, 23-27

ACCEPT message (TACACS+), 25

access control
- ingress ACLs, 603-604
- VLAN assignment, 601-603

Access-Request message (RADIUS), 29

accounting, 21
- RADIUS accounting servers, adding to WLC, 308-309

accuracy of personas, ensuring, 706

ACLs (access control lists)
- applying
  - to switchports, 305-306
  - to WLCs, 310
- DACLs,
- ingress ACLs, 603-604
- local ACLs, creating, 297-298
- number of ACEs, calculating, 603
- SGACLs, 629-632
- VACLs, 461-462
- web authentication redirection ACLs, creating, 310-313

ACS (Cisco Secure Access Control Server), 22

ActivatedGuests, 398

AD (Active Directory), 42, 221-226

https://t.me/learningnets

CA, 505-506
computer authentication, 73
domains, joining, 221-226
admin groups, Cisco ISE, 156
admin portal, 384
Administration Dashboard (Cisco ISE), 161-162
Administration node (Cisco ISE), 129
Administration tab (Cisco ISE), 178-191
Feed Service subcomponent, 191
Identity Management subcomponent, 183-186
Network Resources subcomponent, 186-189
System subcomponent, 178-182
Web Portal Management subcomponent, 189-190
advanced functions, Cisco ISE, 127
posturing, 128
profiling, 127
Advanced license package (Cisco ISE), 178
agent types (NAC), 650-651
Alarms dashlet (Administration Dashboard), 162
alternative ID stores example, authentication policies, 253-254
AND operator, combining with OR operator, 281-286
Android device onboarding flow, 573-577
answering exam questions, 765-766
AnyConnect Diagnostics and Reporting tool, 748-750
AnyConnect Profile Editor, configuring Cisco AnyConnect NAM supplicant, 75-88
Authentication Policy view, 78
Client Policy view, 76-78
Network Groups view, 87
Networks view, 79-86
answers to "Do I Know This Already?" quizzes, 773-792
applying ACLs
to switchports, 305-306
to WLCs, 310
assessment options for posturing, 652-654
authentiation servers, 56
authentication, 21, 233-232. See also authentication policies; WebAuth
802.1X
computer authentication, 73
intended behavior, 680-681
NADs, 63
supplicants, 63-89
user authentication, 72-73
and authorization, 237
certificate authentication
configuring, 506-516
EAP-TLS, 506

https://t.me/learningnets

proof of possession, verifying, 504-505
revocation, verifying, 502-503
signing of certificates, verifying, 499-500
validity dates, verifying, 501
verifying, 516-519
devices without supplicants, managing
MAB, 98-100, 115-116
web authentication, 100-106
EAP
802.1X, 56
native types, 58-59
tunneled types, 59-61
failure types, 401
OTP services, 44-45
posturing, 117-118
two-factor authentication, 43-44
verifying
on Cisco switches, 329-334
on WLCs, 334-336
WebAuth, 340-341
CWA, 346-349
DRW, 349
LWA (local web authentication), 346
on wired switches, configuring
ACL, applying, 305-306
creating local ACLs, 297-298
Flex-Auth, 299-302
global 802.1X commands, 297
global configuration AAA commands, 293-294
global configuration RADIUS commands, 294-297
HA, 299-302
host mode of switchport, setting, 302-303
settings, 303-305
switchports, 299
timers, 305
on WLCs, configuring, 306-328
ACLs, applying, 310
corporate SSID, creating, 324-328
dynamic interfaces for client VLAN, creating, 315
guest dynamic interface, creating, 317
guest WLAN, creating, 319-323
posture agent redirection ACL, creating, 313-314
RADIUS accounting servers, adding, 308-309
RADIUS authentication servers, adding, 306-308
RADIUS fallback, 309-310
web authentication redirection ACL, creating, 310-313
Authentication Live Log, 436-438

https://t.me/learningnets

authentication policies
allowed protocols, 243-247
conditions, 241-243
examples
alternative ID stores example, 253-254
remote access VPN example, 251-252
wireless SSID example, 248-251
goals, 238-239
identity store, 247
MAB, 255-257
MAB rule flow chart, 240
options, 247
restoring, 257
Authentication Policy view (AnyConnect Profile Editor), 78
Authentication subcomponent (Cisco ISE), 173
Authentication subcomponent (Operations tab), 165-169
Authentications dashlet (Administration Dashboard), 162
authenticators, 56
authorization, 21. See also authorization policies
and authentication, 237
CoA, 31-32, 113
and Cisco ISE Profiler, 478-482
posturing, 654-655
guest authorization, configuring, 400-415
authorization policies, 265-279
building, 360-362
conditions, saving for reuse, 279-281
goals of, 265-266
profiles
endpoint identity groups, 483-485
EndPointPolicy, 486
rules, 266-279
examples, 272-279
role-specific authorization rules, 271
Authorization subcomponent (Cisco ISE), 173
AV-pairs, 31

B

backup and restore strategies, 718-719
bandwidth, requirements for Cisco ISE, 139
Base license package (Cisco ISE), 178
bootstrapping Cisco ISE, 201-215
CA-signed certificates, 206-215
certificates
locating, 204-205
self-signed certificates, 206

https://t.me/learningnets

building blocks for posturing, 658-659
BYOD (bring your own device)
challenges, 528-529
Dogtag, 821-825
configuring, 829-842
installing, 829-830
installing packages with yum, 825
NTP service, configuring, 826-827
PHP services, installing, 828-829
prerequisites, 821-825
EAP chaining, 593-594
LDAP server, installing, 827-828
Microsoft CA, configuring for BYOD, 796-819
requirements, 795-796
onboarding process, 529-530
Android onboarding flow, 573-577
configuring on Cisco ISE, 538-569
device enrollment, 571
device registration, 570
dual-SSID approach, 530
flows, verifying, 581-582
iOS onboarding flow, 570-573
Mac OSX onboarding flow, 577-580
NADs, configuring, 532-538
Windows device onboarding flow, 577-580

C

CACs (common access cards), 45-46
calculating number of ACEs per ACL, 603
CAP (certificate authentication profile), 226-227
captive portal bypass, 354-355
CAs (certificate authorities), 46
AD, 505-506
CRLs, 49
Dogtag, 821-825
configuring, 830-842
configuring ISE for use, 842-843
installing, 829-830
LDAP server, installing, 827-828
NTP service, configuring, 826-827
PHP services, installing, 828-829
Microsoft CA, configuring for BYOD, 796-819
requirements, 795-796
CA-signed certificates, 206-215
CCNP Security certification track, 2-4
CDP (Cisco Discovery Protocol), 116

https://t.me/learningnets

certificate authentication
configuring, 506-516
allowed protocols, validating, 507-508
authorization policies, 511-512
CAP, verifying, 508-511
ensuring trust, 512-513
public certificate, importing, 513-515
EAP-TLS, 506
proof of possession, verifying, 504-505
revocation, verifying, 502-503
signing of certificates, verifying, 499-500
validity dates, verifying, 501
verifying, 516-519
certificates, 180. See also certificate authentication
CAP, 226-227
CA-signed certificates, 206-215
CRLs, 49
locating, 204-205
OCSP, 49-50
self-signed certificates, 206
X.509 certificates, 46
revocation, 48-49
validity dates, 47-48
certification
CCNP Security certification track, 2-4
security certifications, comparing, 6
challenges to BYOD, 528-529
CHAP (Challenge/Handshake Authentication Protocol), 23
Cisco AnyConnect NAM supplicant, 87-88
AnyConnect NAM profiles, implementing, 87-88
AnyConnect Profile Editor
Authentication Policy view, 78
Client Policy view, 76-78
Network Groups view, 87
Networks view, 79-86
configuring, 75-88
EAP chaining, 89
Cisco Certification Exam Tutorial, 759-760
Cisco IOS 12.2.X, global configuration RADIUS commands, 294
Cisco IOS 15.X, global configuration AAA commands, 295-297
Cisco ISE (Identity Services Engine), 127-129
admin groups, 156
advanced functions, 127
authentication policies, 237-254
allowed protocols, 243-247
alternative ID stores example, 253-254

https://t.me/learningnets

conditions, 241-243
goals of, 238-239
identity store, 247
MAB, 255-257
options, 247
remote access VPN example, 251-252
restoring, 257
wireless SSID example, 248-251
backup and restore strategies, 718-719
bandwidth requirements, 139
bootstrapping, 201-215
CA-signed certificates, 206-215
locating certificates, 204-205
self-signed certificates, 206
certificate authentication, configuring, 506-516
allowed protocols, validating, 507-508
authorization policies, 511-512
CAP, verifying, 508-511
ensuring trust, 512-513
public certificate, importing, 513-515
verifying configuration, 516-519
CRLs, 49
deploying, 133
four-node deployment, 136-137
fully distributed deployment, 137
single-node deployment, 133-135
two-node deployment, 135-136
Dogtag, configuring, 842-843
external identity stores, 41-50
AD, 42
LDAP, 42-43
smart cards, 45-46
form factors, 201
Guest Services, 399-400. See also portals; WebAuth
guest-level access, Cisco ISE, 128
GUI, 150
Administration Dashboard, 161-162
Administration Home page, 162-164
Administration tab, 178-191
Operations tab, 165-172
Policy tab, 173-177
HA
MnT, 707-709
node groups, 710-712
PANs, 709-710
internal identity stores, 39-40
licensing packages, 178

https://t.me/learningnets

licensing in multinode cube, 706-707
logging events, 129
logging in, 155-156
network devices
NADs, 217-218
NDGs, 216
nodes
communication between, 138-139
configuring in distributed environment, 702-706
load balancing, 713-715
promoting to primary device, 702-703
onboarding process (BYOD), configuring, 538-569
patching, 716
personas, 129-131
Administration node, 129
IPN, 130-131
MnT, 130
Policy Service Node, 129-130
phased deployment approach, 681-694
Closed Mode, 692-694
Low-Impact Mode, 689-695
Monitor Mode, 685-689
preparing for, 683-685
transitioning to end state, 695
on wireless networks, 695
physical appliance specifications, 131
policies, 192-194
authentication, 233-232
portals, 384-389
posturing, 128, 648
assessment options, 652-654
authorization policy for compliance, modifying, 666-667
authorization policy for CPP, modifying, 663-665
building blocks, 658-659
CoA, 654-655
conditions, 659-660
functional components, 648
NAC agent types, 650-651
posture conditions, 652-654
remediation, 661
requirement function, 662-663
verifying, 667-674
profiling, 127, 445-447
probes, 447-459
RADIUS, 127
services, 138-139
syslog, 332-334

https://t.me/learningnets

trusted-level access, 128
user identity groups, 40
virtual appliance specifications, 132
web browser support, 150
X.509 certificates, 46
revocation, 48-49
validity dates, 47-48
Cisco ISE Profiler, 445-447. See also
and CoA, 478-482
global CoA, 479-480
per-profile CoA, 480-481
endpoint attribute filtering, 482
verifying profiling, 486-491
dashboard, 486-487
Endpoints Drill-down tool, 487-488
Global Search tool, 488
Cisco security certifications, comparing, 6
Client Policy view (AnyConnect Profile Editor), 76-78
client provisioning, 193
Client Provisioning subcomponent (Cisco ISE), 175-176
Closed Mode, 692-694
CoA (change of authorization), 31-32, 113
and Cisco ISE Profiler, 478-482
global CoA, 479-480
per-profile CoA, 480-481
posturing, 654-655
collection filters, 746-747
combining AND with OR operators, 281-286
command sets, 22
commands
debug commands, 336, 755
global 802.1X commands, 297
global configuration AAA commands, 293-294
global configuration RADIUS commands, 294-297
show aaa servers command, 329-330
show authentication session command, 331-332
show authentication sessions interface command, 668, 753-754
show device-sensor cache all command, 491
show monitor command, 460
test aaa command, 330-331
communication between Cisco ISE nodes, 138-139
communication flows, RADIUS, 29-30
comparing
authentication and authorization, 237, 265
EAP types, 62
RADIUS and TACACS+, 32

https://t.me/learningnets

security certifications, 6
virtual versus physical Cisco ISE deployments, 131-133
components of 802.1X, 56
computer authentication, 73
conditions, 241-243
combining AND with OR operators, 281-286
saving for reuse, 279-281
conditions (posture service), 117-118, 652-654, 659-660
configuration backups, 718-719
configuring
authentication
on wired switches, 293-306
on WLCs, 306-328
BYOD
Cisco ISE configuration, 538-569
onboarding, 532-538
certificate authentication, 506-516
allowed protocols, validating, 507-508
authorization policies, 511-512
CAP, verifying, 508-511
ensuring trust, 512-513
Cisco AnyConnect NAM supplicant, 75-88
Dogtag, 829-842
Guest Services, guest authorization, 400-415
ISS, 356-357
MDM onboarding, 584-589
Microsoft CA for BYOD, 796-819
requirements, 795-796
nodes in distributed environment, 702-706
ensuring accuracy of personas, 706
promoting node to primary device, 702-703
registering node to the deployment, 703-705
NTP, 826-827
portals
Friendly Names, 391-392
interfaces, 391
ports, 389-390
posture service, 655-674
conditions, 659-660
CPP, 657
remediation, 661
requirement function, 662-663
profiling, 459-464
device sensors, 462-463
DHCP helper, 459-460
SNMP settings, 481

https://t.me/learningnets

SPAN, 460
VACLs, 461-462
VMware, 463-464
sponsor portal
policies, 392-393
sponsor groups, 394-396
WebAuth
CWA, 350-359
device registration, 363-368
Windows native supplicant, 64-72
CONTINUE packets (TACACS+), 25
controlling access to networks
ingress ACLs, 603-604
VLAN assignment, 601-603
CPP (client provisioning policy), 657
authorization policy, modifying, 663-665
resources, downloading, 656-657
creating
individual accounts, 416
ISE cubes, 704-705
local ACLs, 297-298
random user accounts, 417
CRLs (certificate revocation lists), 49
cubes (ISE)
creating, 704-705
licensing, 706-707
logging targets, 729-730
customizing portals, 399-400
CWA (centralized web authentication), 104-106, 346-349
authorization policies, building, 360-362
captive portal bypass, 354-355
configuring, 350-359
verifying, 369-375

D

DACLs (downloadable access control lists),
dashboard, verifying profiling, 486-487
dashlets (Administration Dashboard), 162
databases
identity stores, 38-40
databases, identity stores
external identity stores, 41-50
internal identity stores, 39-40
debug commands, 336, 755
debug logs, 731-732
deploying Cisco ISE, 133. See also phased deployment approach

https://t.me/learningnets

four-node deployment, 136-137
fully distributed deployment, 137
single-node deployment, 133-135
two-node deployment, 135-136
device administration, 16, 21-22
command sets, 22
TACACS, 22
device sensors, configuring, 462-463
devices without supplicants, managing
MAB, 98-100
DHCP profiling, 116
MAC addresses, 115-116
WebAuth, 100-106
DHCP helper, 459-460
DHCP profiling, 116
probes, 449-452
DHCPSPAN probes, 449-452
diagnostic tools, 735-747
AnyConnect Diagnostics and Reporting tool, 748-750
collection filters, 746-747
endpoint diagnostics, 748
Evaluate Configuration Validator, 735-739
RADIUS Authentication Troubleshooting tool, 739-740
TCP Dump, 741-746
DIAMETER, 23
distributed environments
nodes, configuring, 702-706
ensuring accuracy of personas, 706
promoting node to primary device, 702-703
registering node to the deployment, 703-705
DNS probes, 454
Dogtag, 821-825
configuring, 829-842
installing, 829-830
installing packages with yum, 825
LDAP server, installing, 827-828
NTP service, configuring, 826-827
PHP services, installing, 828-829
prerequisites for use, 821-825
domains (AD), joining, 221-226
Dot1X. See 802.1X
Downlink MACSec, 634-635
downloading
CPP resources, 656-657
debug logs, 732
DRW (device registration WebAuth), 349

https://t.me/learningnets

dual-SSID onboarding, 530

E

EAP (Extensible Authentication Protocol)
- 802.1X, 56-58
- native types, 58-59
- tunneled types, 59-61
- types of, comparing, 62

EAP chaining, 89, 593-594
EAP-FAST, 60-61
EAP-GTC, 59
EAP-MD5, 58
EAP-MSCHAPv2, 59
EAPoL (EAP over LAN), 56-58
- components, 56

EAP-TEAP, 89
EAP-TLS, 58-59
- certificate authentication, 506

endpoint attribute filtering, 482
endpoint profile policies, 467-477
Endpoint Protection Service, 170
EndPointPolicy, 486
endpoints
- diagnostics, 748
- identities, 489
- local endpoint groups, 219
- managing, 590-593
- MDMs, 119
- posturing, 117-118, 649

Endpoints Drill-down tool, 487-488
enforcing traffic based on SGT, 628-632
enrolling devices for BYOD, 571
ensuring accuracy of personas, 706
ERROR message (TACACS+), 25
ERS Admin role (Cisco ISE), 156
ERS Guest role (Cisco ISE), 156
Evaluate Configuration Validator, 735-739
Evaluation license package (Cisco ISE), 178
exam. See SISAS 300-208 exam; practice exams
examples
- of authentication policies
  - alternative ID stores example, 253-254
  - remote access VPN example, 251-252
  - wireless SSID example, 248-251
- of authorization policies, 272-279

https://t.me/learningnets

extended logging, 751
external identity sources, 38
external identity stores, 41-50, 220-229
AD, 42, 221-226
domains, joining, 221-226
CAP, 226-227
ISS, 227-229
LDAP, 42-43
smart cards, 45-46

F

failure types, authentication, 401
features of MDMs, 119
Fedora
installing packages with yum, 825
proxy configuration, 825
Feed Service subcomponent (Cisco ISE), 191
FIPS (Federal Information Processing Standard), 77
Flex-Auth, configuring on wired switches, 299-302
flows (onboarding), verifying, 581-582
form factors for Cisco ISE, 201
format of SISAS 300-208 exam, 9-10
four-node deployment (Cisco ISE), 136-137
Friendly Names, configuring on web portals, 391-392
fully distributed deployment (Cisco ISE), 137
functional components of posture service, 648

G

global 802.1X commands, 297
global CoA, 479-480
global configuration AAA commands, 293-294
global configuration RADIUS commands, 294-297
Global Search tool, 488
goals
of authentication policies, 238-239
of authorization policies, 265-266
groups
local endpoint groups, 219
local user identity groups, 218
NDGs, 216
node groups, HA, 709-710
sponsor groups
configuring, 394-396
mapping, 396-397
guest accounts, provisioning, 416

https://t.me/learningnets

Guest Services. See also WebAuth
guest accounts, provisioning, 416
guest authorization, configuring, 400-415
guest user portal, screen elements, 386-389
importing user accounts, 418
individual accounts, creating, 416
portals, 384-389
customizing, 399-400
random user accounts, creating, 417
verifying guest access
on the switch, 428-438
on WLC, 419-427
web portals
configuring, 391-392
web portals, configuring, 389-390
guest user portal, 385
screen elements, 386-389
guest user types, 398
guest-level access, Cisco ISE, 128
GUI (Cisco ISE), 150
admin groups, 156
Administration Dashboard, 161-162
Administration Home page, 162-164
Administration tab, 178-191
Feed Service subcomponent, 191
Identity Management subcomponent, 183-186
Network Resources subcomponent, 186-189
System subcomponent, 178-182
Web Portal Management subcomponent, 189-190
logging in, 155-156
Operations tab, 165-172
Authentication subcomponent, 165-169
Reports subcomponent, 169
Troubleshoot subcomponent, 171-172
Operations tab (Cisco ISE), Endpoint Protection Service, 170
Policy tab, 173-177
Authentication subcomponent, 173
Authorization subcomponent, 173
Client Provisioning subcomponent, 175-176
Policy Elements subcomponent, 177
Posture subcomponent, 175
Profiling subcomponent, 174-175
Security Group Access subcomponent, 176
guidelines for load balancing, 713-714

H

https://t.me/learningnets

HA (high availability)
configuring on wired switches, 299-302
MnT, 707-709
node groups, 710-712
PANs, 709-710
RADIUS fallback, 309-310
hard tokens, 44
Help link (Administration Home page), 163-164
Helpdesk admin role (Cisco ISE), 156
host mode of switchport, setting, 302-303
hotfixes for Windows native supplicant, 752
HTTP probes, 457-459

I

identifying knowledge gaps in exam topics, 767-769
identities, 38
Identity Admin role (Cisco ISE), 156
identity groups, local user identity groups, 218
identity management, 35-34
endoint identities, 489
OTP services, 44-45
two-factor authentication, 43-44
X.509 certificates, 46
CRLs, 49
validity dates, 47-48
Identity Management subcomponent (Cisco ISE), 183-186
identity source sequence, 38
identity stores, 38-40, 247
external identity stores, 41-50, 220-229
AD, 221-226
CAP, 226-227
LDAP, 42-43
identity source sequence, 38
internal identity stores, 39-40
ISS, 227-229
local endpoint groups, 219
local users, 220
IEEE 802.1X, 23, 56-58
components, 56
computer authentication, 73
intended behavior, 680-681
NADs, 63
supplicants, 63-89
Cisco AnyConnect NAM supplicant, 75-88
devices without supplicants, 97-98
Windows native supplicant, 64-72

https://t.me/learningnets

user authentication, 72-73
implementing AnyConnect NAM profiles, 87-88
importing
public certificates, 513-515
user accounts, 418
individual accounts, creating, 416
ingress access control
ingress ACLs, 603-604
VLAN assignment, 601-603
installing
Dogtag, 829-830
LDAP server, 827-828
PHP services, 828-829
intended audience for this book, 6
intended behavior of 802.1X, 680-681
interfaces
configuring as switchports, 299
configuring on web portals, 391
dynamic interfaces for client VLAN, creating, 315
guest dynamic interface, creating, 317
internal endpoint database, 40
internal identity stores, 39-40
iOS devices, onboarding flow, 570-573
IOS load balancing, 715-716
IPN (Inline Posture Node), 130-131
ISE. See Cisco ISE (Identity Services Engine)
ISS (identity source sequences), 227-229
configuring, 356-357

J-K

Jobs, Steve, 522
joining AD domains, 221-226
knowledge gaps in exam topics, identifying, 767-769

L

LDAP (Lightweight Directory Access Protocol), 42-43
licensing packages
Cisco ISE, 178
licensing in multinode cube, 706-707
Live Authentication Log, 726-728
viewing, 336-337
Live Sessions Log, 728
viewing, 337
LLDP (Link Layer Discovery Protocol), 116
load balancing, 713-715

https://t.me/learningnets

IOS load balancing, 715-716
local ACLs, creating, 297-298
local user identity groups, 218
local users, 220
locating certificates, 204-205
logging
Authentication Live Log, 436-438
categories, 730
Cisco ISE, 129
debug logs, 731-732
extended logging, 751
Live Authentication Log, 726-728
viewing, 336-337
Live Sessions Log, 728
viewing, 337
log files, viewing from CLI, 733-734
supplicant provisioning logs, 753
support bundles, 734
syslog, 332-334
targets, 729-730
logging in to Cisco ISE, 155-156
logical profiles, 478
Low-Impact Mode, 689-695
LWA (local web authentication), 101-102, 346
with centralized portal, 102-104

M

MAB (MAC Authentication Bypass), 98-100, 113-117
authentication policies, 255-257
authorization policy, 403-406
DHCP profiling, 116
MAC addresses, 115-116
rules, 240
MAC addresses, 115-116
Mac OSX device onboarding flow, 577-580
MACSec, 632-641
Downlink MACSec, 634-635
Uplink MACSec, 638-640
maintenance
backup strategies, 718-719
patching Cisco ISE, 716
managing
devices without supplicants
MAB, 98-100, 113-117
WebAuth, 100-106
endpoints, 590-593

https://t.me/learningnets

mobile devices, MDMs, 119
mapping sponsor groups, 396-397
MDM (mobile device management), 118-119
features, 119
onboarding process, 583-589
integration points, 583
onboarding process (BYOD), configuring, 584-589
vendors, 119
messages
RADIUS
accounting messages, 30
authentication and authorization messages, 29-30
TACACS+
authentication messages, 25
authorization and accounting messages, 26-27
Metrics dashlet (Administration Dashboard), 162
Microsoft Active Directory, 42
Microsoft CA, configuring for BYOD, 796-819
requirements, 795-796
MnT (Monitoring and Troubleshooting node), 130
HA, 707-709
Mnt Admin role (Cisco ISE), 156
mobile devices
BYOD
Android device onboarding flow, 573-577
challenges, 528-529
configuring on Cisco ISE, 538-569
iOS onboarding flow, 570-573
Mac OSX device onboarding flow, 577-580
onboarding flows, verifying, 581-582
onboarding process, 529-530, 570-571
single-SSID onboarding, 531-530
Windows device onboarding flow, 577-580
MDMs (mobile device managers), 119
modifying authorization policy
for compliance, 666-667
for CPP, 663-665
Monitor Mode, 685-689
MS-CHAP (Microsoft CHAP), 23
multinode cubes, licensing, 706-707

N

NAC (network admission control) agent
agent types, 650-651
posture assessment, 649
supported remediation types, 651

https://t.me/learningnets

NADs (Network Access Devices), 23, 63, 217-218
configuring for onboarding, 532-538
verifying authentication, 329
NAM (Network Access Manager) profiles, implementing, 87-88
native EAP types, 58-59
native tagging, 621-628
NDGs (network device groups), 216
NetFlow probes, 457
network access, 16, 22-32
RADIUS, 22, 28-32
accounting messages, 30
authentication and authorization messages, 29-30
AV-pairs, 31
CoA, 31-32
communication flows, 29-30
comparing to TACACS+, 32
service types, 29
TACACS+, 23-27
authentication messages, 25
authorization and accounting messages, 26-27
comparing to RADIUS, 32
Network Device Admin role (Cisco ISE), 156
network devices
NADs, 217-218
NDGs, 216
Network Groups view (AnyConnect Profile Editor), 87
Network Resources subcomponent (Cisco ISE), 186-189
Networks view (AnyConnect Profile Editor), 79-86
NMAP probes, 453-454
node groups, HA, 709-710
nodes (Cisco ISE), 129-131
Administration node, 129
communication between, 138-139
configuring in distributed environment, 702-706
ensuring accuracy of personas, 706
promoting node to primary device, 702-703
registering node to the deployment, 703-705
four-node deployment, 136-137
IPN, 130-131
load balancers, 713-715
MnT, 130
HA, 707-709
multinode cubes, licensing, 706-707
PANs, HA, 709-710
Policy Service Node, 129-130
single-node deployment, 133-135

https://t.me/learningnets

two-node deployment, 135-136
nontunneled EAP types, 58-59
NTP (Network Time Protocol), 48
configuring, 826-827

O

OCSP (Online Certificate Status Protocol), 49-50
onboarding process (BYOD), 529-530. See also onboarding process (MDM)
Android onboarding flow, 573-577
configuring on Cisco ISE, 538-569
device enrollment, 571
device registration, 570
dual-SSID approach, 530
flows, verifying, 581-582
iOS onboarding flow, 570-573
Mac OSX onboarding flow, 577-580
NADs, configuring, 532-538
single-SSID approach, 531-530
Windows device onboarding flow, 577-580
onboarding process (MDM), 583-589
configuring, 584-589
integration points, 583
operational backups, 718-719
Operations tab (Cisco ISE), 165-172
Authentication subcomponent, 165-169
Endpoint Protection Service, 170
Reports subcomponent, 169
Troubleshoot subcomponent, 171-172
operators, combining AND with OR operators, 281-286
options
assessment options for posturing, 652-654
for authentication policy rules, 247
OR operator
combining with AND operator, 281-286
OTP (one-time password) services, 44-45
OUIs (organizationally unique identifiers), 115-116

P

packages, updating with yum, 826
PANs (policy administration nodes), HA, 709-710
PAP (Password Authentication Protocol), 23
passwords, OTP services, 44-45
patching Cisco ISE, 716
PEAP (Protected EAP), 60
Pearson VUE, registering for SISAS 300-208 exam, 5-6
permitting promiscuous traffic, 463-464

https://t.me/learningnets

per-profile CoA, 480-481
personas, 129-131
Administration node, 129
communication between, 138-139
four-node deployment, 136-137
IPN, 130-131
nodes, configuring in distributed environment, 702-706
ensuring accuracy of personas, 706
multinode cubes, licensing, 706-707
promoting node to primary device, 702-703
registering node to the deployment, 703-705
Policy Service Node, 129-130
single-node deployment, 133-135
two-node deployment, 135-136
phased deployment approach, 681-694
Closed Mode, 692-694
Low-Impact Mode, 689-695
Monitor Mode, 685-689
preparing Cisco ISE for, 683-685
transitioning to end state, 695
on wireless networks, 695
PHP services, installing, 828-829
physical appliance specifications (Cisco ISE), 131
PKI (Public Key Infrastructure), 180
Plus license package (Cisco ISE), 178
policies (Cisco ISE), 192-194
authentication, 233-232
authentication policies, 237-254
allowed protocols, 243-247
alternative ID stores example, 253-254
conditions, 241-243
goals of, 238-239
identity store, 247
MAB, 255-257
MAB rule flow chart, 240
options, 247
remote access VPN example, 251-252
restoring, 257
wireless SSID example, 248-251
authorization policies, 265-279
conditions, saving for reuse, 279-281
examples, 272-279
goals of, 265-266
rules, 266-279
CPP, 657
resources, downloading, 656-657

https://t.me/learningnets

guest authorization policies, configuring, 400-415
profiling policies, 464-478
endpoint profile policies, 467-477
logical profiles, 478
profiler feed service, 464-466
Policy Admin role (Cisco ISE), 156
Policy Elements subcomponent (Cisco ISE), 177
Policy Service Node (Cisco ISE), 129-130
Policy tab (Cisco ISE), 173-177
Authentication subcomponent, 173
Authorization subcomponent, 173
Client Provisioning subcomponent, 175-176
Policy Elements subcomponent, 177
Posture subcomponent, 175
Profiling subcomponent, 174-175
Security Group Access subcomponent, 176
Port Bounce CoA, 480
portals, 384-389
captive portal bypass, 354-355
customizing, 399-400
Friendly Names, configuring, 391-392
guest user portal, screen elements, 386-389
interfaces, configuring, 391
ports, configuring, 389-390
sponsor portal
guest accounts, provisioning, 416
guest user types, 398
policies, configuring, 392-393
sponsor groups, configuring, 394-396
sponsor groups, mapping, 396-397
types of sponsors, 393-396
ports
communication between Cisco ISE nodes, 138-139
configuring on web portals, 389-390
Posture Compliance dashlet (Administration Dashboard), 162
Posture subcomponent (Cisco ISE), 175
posturing, 117-118, 128, 648
assessment options, 652-654
building blocks, 658-659
CoA, 654-655
compliance, modifying authorization policy, 666-667
conditions, 659-660
configuring, 655-674
CPP, 657
authorization policy, modifying, 663-665
functional components, 648

https://t.me/learningnets

NAC
agent types, 650-651
supported remediation types, 651
posture conditions, 652-654
remediation, 661
requirement function, 662-663
verifying, 667-674
POTS (plain old telephone service), 22
PPP (Point-to-Point Protocol), 28
practice exams, 763-767
preparing
Cisco ISE for phased deployment, 683-685
for SISAS 300-208 exam, 759, 769-770
answering questions, 765-766
Cisco Certification Exam Tutorial, 759-760
exam-day advice, 762-763
features of this book, 13-14
knowledge gaps, identifying, 767-769
practice exams, 766-767
pre-exam suggestions, 762
time management, 760-762
topics covered, 4-5
probes, 447-459
DHCP probes, 449-452
DHCPSPAN probes, 449-452
DNS probes, 454
HTTP probes, 457-459
NetFlow probes, 457
NMAP probes, 453-454
RADIUS probes, 452-453
SNMP probes, 455-456
SNMP settings, configuring, 481
Profiler Activity dashlet (Administration Dashboard), 162
profiler feed service, 464-466
profiling, 127, 193, 445-447. See also profiling policies
authorization policies
endpoint identity groups, 483-485
EndPointPolicy, 486
endpoint attribute filtering, 482
infrastructure, configuring, 459-464
device sensors, 462-463
DHCP helper, 459-460
SPAN, 460
VACLs, 461-462
interfaces, VMware, 463-464
NetFlow probes, 457

https://t.me/learningnets

probes, 447-459
DHCP probes, 449-452
DHCPSPAN probes, 449-452
DNS probes, 454
HTTP probes, 457-459
NMAP probes, 453-454
RADIUS probes, 452-453
SNMP probes, 455-456
SNMP settings, configuring, 481
verifying, 486-491
dashboard, 486-487
Endpoints Drill-down tool, 487-488
Global Search tool, 488
profiling policies
endpoint profile policies, 467-477
logical profiles, 478
profiler feed service, 464-466
Profiling subcomponent (Cisco ISE), 174-175
promiscuous traffic, permitting, 463-464
promoting nodes to primary device, 702-703
proof of possession for certificates, verifying, 504-505
provisioning
client provisioning, 193
guest accounts from sponsor portal, 416
supplicant provisioning logs, 753
pseudo-browsers, 355
PSNs (policy service nodes), probes, 447-459
DHCP profiling, 449-452
DHCPSPAN probes, 449-452
DNS probes, 454
HTTP probes, 457-459
NetFlow probes, 457
NMAP probes, 453-454
RADIUS probes, 452-453
SNMP probes, 455-456

Q-R

RA (remote access), 106
RADIUS (Remote Authentication Dial-In User Service), 22, 28-32, 127
accounting messages, 30
authentication and authorization messages, 29-30
AV-pairs, 31
CoA, 31-32, 113
communication flows, 29-30
comparing to TACACS+, 32
IOS load balancing, 715-716

https://t.me/learningnets

probes, 452-453
service types, 29
RADIUS Authentication Troubleshooting tool, 739-740
random user accounts, creating, 417
RBAC Admin role (Cisco ISE), 156
Reauth CoA, 480
registering
devices for BYOD, 570
nodes to ISE cube, 703-705
for SISAS 300-208 exam, 5-6
WebAuth devices, 363-368
REJECT message (TACACS+), 25
remediation
NAC support for, 651
posture service, 661
remote access VPN example, authentication policies, 251-252
REPLY packets (TACACS+), 25
Reports subcomponent (Cisco ISE), 169
REQUEST message (TACACS+), 26-27
RESPONSE message (TACACS+), 26-27
responses to authentication failure, 402
restoring authentication policies, 257
reusing conditions, 279-281
revocation
OCSP, 49-50
verifying for certificates, 502-503
X.509 certificates, 48-49
role-specific authorization rules, 271
rules
802.1X authenticaton rule, 401
for authentication policies, 240
conditions, 241-243
options, 247
for authorization policies, 266-279
examples, 272-279
role-specific authorization rules, 271

S

sample switch configurations
Catalyst 2960/3560/3750 Series, 12.2(55)SE, 845-848
Catalyst 3560/3750 Series, 15.0(2)SE, 848-852
Catalyst 4500 Series, IOS-XE 3.3.0/15.1(1)SG, 852-856
Catalyst 6500 Series, 12.2(33)SXJ, 856-858
saving conditions for reuse, 279-281
screen elements, guest user portal, 386-389

https://t.me/learningnets

security certifications, comparing, 6
Security Group Access subcomponent (Cisco ISE), 176
selecting EAP type, 62
self-signed certificates, 206
Server Information pop-up (Administration Home page), 162
service types, 29
services (Cisco ISE), 138-139
posture service, 648
assessment options, 652-654
authorization policy for compliance, modifying, 666-667
authorization policy for CPP, modifying, 663-665
building blocks, 658-659
CoA, 654-655
conditions, 659-660
configuring, 655-674
CPP, 657
functional components, 648
NAC agent types, 650-651
posture conditions, 652-654
remediation, 661
requirement function, 662-663
verifying, 667-674
Setup Assistant link (Administration Home page), 163
SGA (security group access), 193-194. See also TrustSec
enforcement, 628-632
native tagging, 621-628
SGTs, 606-613
SXP, 613-621
SGACLs (Security Group ACLs), 629-632
SGTs (security group tags), 606-613
native tagging, 621-628
show aaa servers command, 329-330
show authentication session command, 331-332
show authentication session interface command, 753-754
show authentication sessions interface command, 668
show device-sensor cache all command, 491
show monitor command, 460
signing of certificates, verifying, 499-500
single-node deployment (Cisco ISE), 133-135
SISAS 300-208 exam
answering questions, 765-766
exam-day advice, 762-763
format of exam, 9-10
knowledge gaps, identifying, 767-769
practice exams, 763-767
pre-exam suggestions, 762

https://t.me/learningnets

preparing for, 759, 769-770
Cisco Certification Exam Tutorial, 759-760
features of this book, 13-14
registering for, 5-6
time management, 760-762
topics covered, 4-5
smart cards, 45-46
SNMP probes, 455-456
soft tokens, 44
sponsor groups
configuring, 394-396
mapping, 396-397
sponsor portal, 385
configuring, 392-393
guest accounts, provisioning, 416
guest user types, 398
sponsor groups
configuring, 394-396
mapping, 396-397
types of sponsors, 393-396
SponsorAllAccounts group, 394
SponsorGroupGrpAccounts group, 394
SponsorGroupOwnAccounts group, 394
SSL (Secure Sockets Layer), 42
standalone AnyConnect Profile Editor, configuring Cisco AnyConnect NAM supplicant, 75-88
Authentication Policy view, 78
Client Policy view, 76-78
Network Groups view, 87
Networks view, 79-86
START packets (TACACS+), 25
Super Admin role (Cisco ISE), 156
supplicants, 56, 63-89
Cisco AnyConnect NAM supplicant, 75-88
AnyConnect NAM profiles, implementing, 87-88
AnyConnect Profile Editor views, 76-87
EAP chaining, 89
devices without supplicants, managing, 97-98
MAB, 98-100, 113-117
MAC addresses, 115-116
WebAuth, 100-106
supplicant provisioning logs, 753
Windows native supplicant, 64-72
hotfixes, 752
support bundles, 734
switches
authentication, verifying, 329-334

https://t.me/learningnets

guest access, verifying, 428-438
wired switches, configuring authentication
ACL, applying, 305-306
creating local ACLs, 297-298
Flex-Auth, 299-302
global 802.1X commands, 297
global configuration AAA commands, 293-294
global configuration RADIUS commands, 294-297
HA, 299-302
host mode of switchport, setting, 302-303
settings, 303-305
switchports, 299
timers, 305
switchports
configuring on wired switches, 299
host mode, setting, 302-303
SXP (Security Group Exchange Protocol), 613-621
syslog, 332-334
System Admin role (Cisco ISE), 156
System subcomponent (Cisco ISE), 178-182
System Summary dashlet (Administration Dashboard), 162

T

TACACS (Terminal Access Controller Access Control System), 22
TACACS+ (Terminal Access Controller Access Control System Plus), 23-27
authentication messages, 25
authorization and accounting messages, 26-27
comparing to RADIUS, 32
TCP Dump, 741-746
test aaa command, 330-331
time management, SISAS 300-208 exam, 760-762
timers, configuring on wired switches, 305
topics covered in SISAS exam, 4-5, 10-13
transitioning from Monitor Mode to end state, 695
Triple-A, 21
Troubleshoot subcomponent (Cisco ISE), 171-172
troubleshooting tools
AnyConnect Diagnostics and Reporting tool, 748-750
diagnostic tools, 735-747
collection filters, 746-747
Evaluate Configuration Validator, 735-739
RADIUS Authentication Troubleshooting tool, 739-740
TCP Dump, 741-746
logging
categories, 730
debug logs, 731-732

https://t.me/learningnets

extended logging, 751
Live Authentication Log, 726-728
Live Sessions Log, 728
support bundles, 734
targets, 729-730
trusted-level access, Cisco ISE, 128
TrustSec, 605-632
enforcement, 628-632
native tagging, 621-628
SGTs, 606-613
SXP, 613-621
tunneled EAP types, 59-61
two-factor authentication, 43-44
two-node deployment (Cisco ISE), 135-136
types of sponsors, 393-396

U

Uplink MACSec, 638-640
user accounts
guest accounts, 416
importing, 418
individual accounts, creating, 416
local users, 220
random user accounts, creating, 417
user authentication, 72-73

V

VACLs (VLAN access control lists), 461-462
validity dates for certificates, 47-48
verifying, 501
vendors of MDMs, 119
verifying
authentication
on Cisco switches, 329-334
on WLCs, 334-336
BYOD onboarding flows, 581-582
certificate authentication, 516-519
CWA, 369-375
guest access
on the switch, 428-438
on WLC, 419-427
posturing, 667-674
profiling, 486-491
dashboard, 486-487
Endpoints Drill-down tool, 487-488
Global Search tool, 488

https://t.me/learningnets

proof of possession for certificates, 504-505
revocation of certificates, 502-503
signing of certificates, 499-500
validity dates of certificates, 501
viewing
Live Authentication Log, 336-337
Live Sessions Log, 337
log files from CLI, 733-734
WLC client details, 754
views (AnyConnect Profile Editor)
Authentication Policy view, 78
Client Policy view, 77
Network Groups view, 87
Networks view, 79-86
virtual appliance specifications (Cisco ISE), 132
VLAN assignment, controlling access to networks, 601-603
VMware, permitting promiscous traffic, 463-464
VPNs (virtual private networks), RA, 106

W

web authentication redirection ACLs, creating, 310-313
web browsers. See also GUI (Cisco ISE)
Cisco ISE support for, 150
pseudo-browsers, 355
Web Portal Management subcomponent (Cisco ISE), 189-190
web portals
customizing, 399-400
Friendly Names, configuring, 391-392
guest user portal, screen elements, 386-389
interfaces, configuring, 391
ports, configuring, 389-390
sponsor portal
configuring, 392-393
guest accounts, provisioning, 416
guest user types, 398
sponsor groups, configuring, 394-396
sponsor groups, mapping, 396-397
sponsor types, 393-396
WebAuth, 100-106, 340-341
CWA, 104-106, 346-349
authorization policies, building, 360-362
configuring, 350-359
verifying, 369-375
device registration, configuring, 363-368
DRW, 349
guest accounts, provisioning, 416

https://t.me/learningnets

guest authorization, configuring, 400-415
guest user portal, screen elements, 386-389
individual accounts, creating, 416
LWA, 101-102, 346
with centralized portal, 102-104
portals, 384-389
configuring, 389-390
customizing, 399-400
random user accounts, creating, 417
web portals
configuring, 391-392
web portals, configuring, 391
Windows device onboarding flow, 577-580
Windows native supplicant
configuring, 64-72
hotfixes, 752
Wired AutoConfig, 64
wired switches, configuring authentication
ACL, applying, 305-306
creating local ACLs, 297-298
Flex-Auth, 299-302
global 802.1X commands, 297
global configuration AAA commands, 293-294
global configuration RADIUS commands, 294-297
HA, 299-302
settings, 303-305
switchports, 299
timers, 305
Wireless license package (Cisco ISE), 178
wireless networks
phased deployment approach, 695
WLCs, configuring authentication, 306-328
ACLs, applying, 310
corporate SSID, creating, 324-328
dynamic interfaces for client VLAN, creating, 315
guest dynamic interface, creating, 317
guest WLAN, creating, 319-323
posture agent redirection ACL, creating, 313-314
RADIUS accounting servers, adding, 308-309
RADIUS authentication servers, adding, 306-308
RADIUS fallback, 309-310
web authentication redirection ACL, creating, 310-313
wireless SSID example, authentication policies, 248-251
Wireless Upgrade license package (Cisco ISE), 178
WLANs (wireless LANs), creating guest WLAN, 319-323
WLCs (Wireless LAN Controllers)

https://t.me/learningnets

authentication, configuring, 306-328
ACLs, applying, 310
corporate SSID, creating, 324-328
dynamic interfaces for client VLAN, creating, 315
guest dynamic interface, creating, 317
guest WLAN, creating, 319-323
posture agent redirection ACL, creating, 313-314
RADIUS accounting servers, adding, 308-309
RADIUS authentication servers, adding, 306-308
RADIUS fallback, 309-310
verifying authentication, 334-336
web authentication redirection ACL, creating, 310-313
client details, viewing, 754
debug commands, 336
guest access, verifying, 419-427
Woland, Aaron, 523

X-Y-Z

X.509 certificates, 46
revocation, 48-49
validity dates, 47-48
yum
installing Fedora packages, 825
updating system packages, 826

https://t.me/learningnets

https://t.me/learningnets

ciscopress.com: Your Cisco Certification and Networking Learning Resource



Subscribe to the monthly Cisco Press newsletter to be the first to learn about new releases and special promotions.

Visit **ciscopress.com/newsletters.**

While you are visiting, check out the offerings available at your finger tips.

–Free Podcasts from experts:
- OnNetworking
- OnCertification
- OnSecurity

Podcasts

View them at **ciscopress.com/podcasts.**

–Read the latest author **articles** and **sample chapters** at **ciscopress.com/articles.**

–Bookmark the Certification Reference Guide available through our partner site at **informit.com/certguide.**

Connect with Cisco Press authors and editors via Facebook and Twitter, visit **informit.com/socialconnect.**



https://t.me/learningnets

https://t.me/learningnets

PEARSON

ALWAYS LEARNING

PEARSON IT CERTIFICATION

Browse by Exams ▾ | Browse by Technology ▾ | Browse by Format | Explore ▾ | I'm New Here – Help!

Store | Forums | Safari Books Online

Pearson IT Certification

THE LEADER IN IT CERTIFICATION LEARNING TOOLS

Visit pearsonITcertification.com today to find:

- IT CERTIFICATION EXAM information and guidance for

  cisco | CompTIA | Microsoft | vmware

  Pearson is the official publisher of Cisco Press, IBM Press, VMware Press and is a Platinum CompTIA Publishing Partner—CompTIA's highest partnership accreditation

- EXAM TIPS AND TRICKS from Pearson IT Certification's expert authors and industry experts, such as
  - *Mark Edward Soper* – CompTIA
  - *David Prowse* – CompTIA
  - *Wendell Odom* – Cisco
  - *Kevin Wallace* – Cisco and CompTIA
  - *Shon Harris* – Security
  - *Thomas Erl* – SOACP



- SPECIAL OFFERS – pearsonITcertification.com/promotions
- REGISTER your Pearson IT Certification products to access additional online material and receive a coupon to be used on your next purchase

Articles & Chapters

Blogs

Books

Cert Flash Cards Online

eBooks

Mobile Apps

Newsletters

Podcasts

Question of the Day

Rough Cuts

Short Cuts

Software Downloads

Videos

CONNECT WITH PEARSON IT CERTIFICATION

Be sure to create an account on pearsonITcertification.com and receive members-only offers and benefits

https://t.me/learningnets

Code Snippets

```
IF conditions THEN ALLOW PROTCOLS IN LIST AllowedProtocolList
AND CHECK THE IDENTITY STORE IN LIST IdentityStore
```

https://t.me/learningnets

```
IF [conditions] THEN [AssignThesePermissions]
```

https://t.me/learningnets

```
C3560X(config)#aaa authentication dot1x default group radius
```

https://t.me/learningnets

```
C3560X(config)#aaa authorization network default group radius
```

https://t.me/learningnets

```
C3560X(config)# aaa accounting dot1x default start-stop group radius
```

https://t.me/learningnets

```
C3560X(config)#username radius-test password password
```

https://t.me/learningnets

```
C3560X(config)#radius-server host ise_ip_address auth-port 1812
acct-port 1813 test username radius-test key shared_secret
```

https://t.me/learningnets

```
C3560X(config)#username radius-test password password
```

https://t.me/learningnets

```
C3560X(config)#radius server [name]
```

https://t.me/learningnets

```
C3560X(config-radius-server)# address ipv4 [ip-address] auth-port 1812
acct-port 1813
```

https://t.me/learningnets

```
C3560X(config-radius-server)#key [shared-secret]
```

https://t.me/learningnets

```
C3560X(config-radius-server)# automate-tester username
[username-from-step-1]
```

https://t.me/learningnets

```
3850(config)# radius-server attribute 31 mac format ietf lower-case
```

https://t.me/learningnets

```
C3560X(config)#radius-server dead-criteria time 5 tries 3
```

https://t.me/learningnets

```
C3560X(config)#aaa server radius dynamic-author
C3560X(config-locsvr-da-radius)#client ise_ip_address server-key
shared_secret
```

https://t.me/learningnets

```
C3560X(config)#radius-server vsa send authentication
C3560X(config)#radius-server vsa send accounting
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
C3560X(config)#dot1x system-auth-control
```

https://t.me/learningnets

```
C3560X(config)#ip device tracking
```

https://t.me/learningnets

```
C3560X(config)#ip access-list extended ACL-ALLOW
C3560X(config-ext-nacl)#permit ip any any
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
C3560X(config)#interface range first_interface - last_interface
```

https://t.me/learningnets

```
C3560X(config-if-range)#switchport
```

https://t.me/learningnets

```
C3560X(config-if-range)#switchport host
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication priority dot1x mab
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication order dot1x mab
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication event fail action next-method
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication event server dead action
authorize
C3560X(config-if-range)#authentication event server dead action
authorize voice
C3560X(config-if-range)#authentication event server alive action
reinitialize
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication event server dead action
reinitialize vlan vlan-id
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication host-mode multi-auth
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication violation restrict
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication open
```

https://t.me/learningnets

```
C3560X(config-if-range)#dot1x pae authenticator
```

https://t.me/learningnets

```
C3560X(config-if-range)#dot1x timeout tx-period 10
```

https://t.me/learningnets

```
C3560X(config-if-range)#ip access-group ACL-ALLOW in
```

https://t.me/learningnets

```
C3560X(config-if-range)#authentication port-control auto
```

https://t.me/learningnets

```
3750-X#sho aaa servers

RADIUS: id 1, priority 1, host 10.1.100.232, 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
3750-X#
```

https://t.me/learningnets

```
3750-X# test aaa group radius employee1 Cisco123 legacy
Attempting authentication test to server-group radius using radius
User was successfully authenticated.

3750-X#
```

https://t.me/learningnets

```
3750-X# show authentication session int 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
      Session timeout: N/A
         Idle timeout: N/A
    Common Session ID: 0A013002000000110073D1F6
      Acct Session ID: 0x00000002
               Handle: 0xA9000012

Runnable methods list:
       Method   State
       mab      Authc Success
       dot1x    Not run
```

https://t.me/learningnets

```
C3560X(config)#logging monitor informational
C3560X(config)#logging origin-id ip
C3560X(config)#logging source-interface <interface_id>
C3560X(config)#logging host <ISE_MNT_PERSONA_IP_Address_x> transport
udp port 20514
```

https://t.me/learningnets

```
(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
```

https://t.me/learningnets

```
C3560X(config)# ip http server
```

https://t.me/learningnets

```
C3560X(config)# ip http secure-server
```

https://t.me/learningnets

```
C3560X(config)# ip http active-session-modules none
C3560X(config)# ip http secure-active-session-modules none
```

https://t.me/learningnets

```
C3560X# 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
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
> config network web-auth captive-bypass enable
```

https://t.me/learningnets

```
C3560X#  sho authen session int 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-PERMIT_ALL_TRAFFIC-51ef7db1
     URL Redirect ACL: ACL-WEBAUTH-REDIRECT
         URL Redirect: https://atw-cp-ise02.ise.local:8443/guestportal/gateway?sess
ionId=C0A8FE01000000270B2CAA43&action=cwa
      Session timeout: N/A
         Idle timeout: N/A
    Common Session ID: C0A8FE01000000270B2CAA43
      Acct Session ID: 0x0000002D
               Handle: 0x68000028

Runnable methods list:
       Method   State
       mab      Authc Success
       dot1x    Not run
```

https://t.me/learningnets

```
! Traffic to a DHCP server
permit udp any any eq bootps
! Traffic to a DNS server
permit udp any any eq domain
! Ping traffic and response for diagnostics purposes
permit icmp any any echo
permit icmp any any echo-reply
! Traffic to HTTP websites—must be permitted before being redirected
permit tcp any any eq www
! Traffic to HTTPS websites—must be permitted before being redirected
permit tcp any any eq 443
! Traffic to ISE portals, posturing, posture remediation, and provisioning
! functions
permit tcp any host 172.25.73.98 eq 8443
permit tcp any host 172.25.73.98 eq 8905
permit tcp any host 172.25.73.98 eq 8909
permit udp any host 172.25.73.98 eq 8905 8906
permit udp any host 172.25.73.98 eq 8909
```

https://t.me/learningnets

```
! Traffic from DHCP client
permit udp any any eq 68
! Traffic to a DHCP server
permit udp any any eq 67
! Deny all traffic to RFC1918 private/internal IP ranges
deny ip any 10.0.0.0 0.255.255.255
deny ip any 172.16.0.0 0.15.255.255
deny ip any 192.168.0.0 0.0.255.255
! Allow all other traffic—i.e. Internet-bound traffic
permit ip any any
```

https://t.me/learningnets

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B860000003B0231777D
      Acct Session ID:  0x00000049
               Handle:  0x2B000029
       Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:
         URL Redirect:  https://atw-cp-ise02.ise.local:8443/guestportal/
gateway?sessionId=0A742B860000003B0231777D&action=cwa
     URL Redirect ACL:  ACL-WEBAUTH-REDIRECT
              ACS ACL:  xACSACLx-IP-Pre-WebAuth-54bb3896
            SGT Value:  8

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

https://t.me/learningnets

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  ~68j1r6oz
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B860000003B0231777D
      Acct Session ID:  0x00000049
               Handle:  0x2B000029
       Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:
              ACS ACL:  xACSACLx-IP-Internet_Only-54cb00fc
            SGT Value:  8

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

https://t.me/learningnets

```
DC-4948(config)#monitor session [1-4] source [interface | vlan] [rx | tx]
DC-4948(config)#monitor session [1-4] destination interface [interface_name]
```

https://t.me/learningnets

```
C6K-DIST(config)#ip access-list extended HTTP_TRAFFIC
C6K-DIST(config-ext-nacl)#permit tcp any any eq www
```

https://t.me/learningnets

```
C6K-DIST(config)#ip access-list extended ALL_TRAFFIC
C6K-DIST(config-ext-nacl)#permit ip any any
```

https://t.me/learningnets

```
C6K-DIST(config)#vlan access-map HTTP_MAP 10
C6K-DIST(config-access-map)#match ip address HTTP_TRAFFIC
C6K-DIST(config-access-map)#action forward capture
```

https://t.me/learningnets

```
C6K-DIST(config)#vlan access-map HTTP_MAP 20
C6K-DIST(config-access-map)#match ip address ALL_TRAFFIC
C6K-DIST(config-access-map)#action forward
```

https://t.me/learningnets

```
C6K-DIST(config)#vlan filter HTTP_MAP vlan-list 41,42
```

https://t.me/learningnets

```
C6K-DIST(config)# interface gig1/0/1
C6K-DIST(config-if)#switchport capture allowed vlan 41
C6K-DIST(config-if)#switchport capture allowed vlan add 42
C6K-DIST(config-if)#switchport capture
```

https://t.me/learningnets

```
C3750X(config)#device-sensor filter-list dhcp list <dhcp_list_name>
C3750X(config-sensor-dhcplist)#option name host-name
C3750X(config-sensor-dhcplist)#option name class-identifier
C3750X(config-sensor-dhcplist)#option name client-identifier
```

https://t.me/learningnets

```
C3750X(config)#device-sensor filter-list cdp list <cdp_list_name>
C3750X(config-sensor-cdplist)#tlv name device-name
C3750X(config-sensor-cdplist)#tlv name platform-type
```

https://t.me/learningnets

```
C3750X(config)#device-sensor filter-list lldp list <lldp_list_name>
C3750X(config-sensor-lldplist)#tlv name port-id
C3750X(config-sensor-lldplist)#tlv name system-name
C3750X(config-sensor-lldplist)#tlv name system-description
```

https://t.me/learningnets

```
C3750X(config)#device-sensor filter-spec dhcp include list <dhcp_list_name>
C3750X(config)#device-sensor filter-spec lldp include list <cdp_list_name>
C3750X(config)#device-sensor filter-spec cdp include list <lldp_list_name>
```

https://t.me/learningnets

```
C3750X(config)#device-sensor accounting
C3750X(config)#device-sensor notify all-changes
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
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 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
```

https://t.me/learningnets

```
Network Access:AuthenticationMethod EQUALS MSCHAPV2
AND
Session:Device-OS EQUALS Android
```

https://t.me/learningnets

```
Network Access:AuthenticationMethod EQUALS MSCHAPV2
```

https://t.me/learningnets

```
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 ise.local/Users/Employees
```

https://t.me/learningnets

```
AND
x509_PKI
AND
SAN_is_Mac
AND
Endpoints:BYODRegistration EQUALS Yes
AND
MDM:DeviceRegisterStatus EQUALS Unregistered
```

https://t.me/learningnets

```
Employees
AND
x509_PKI
AND
SAN_is_Mac
AND
MDM:DeviceRegisterStatus EQUALS Registered
AND
MDM:DeviceComplianceStatus EQUAL Compliant
```

https://t.me/learningnets

```
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)
```

https://t.me/learningnets

```
NX7K-DIST(config)# int eth1/3
NX7K-DIST(config-if)# cts manual
NX7K-DIST(config-if-cts-manual)# policy static sgt 0x3
```

https://t.me/learningnets

```
C6K-DIST(config)#cts role-based sgt-map 192.168.26.0/24 sgt 4
```

https://t.me/learningnets

```
C6K-DIST(config)#cts role-based sgt-map vlan-list 40 sgt 4
```

https://t.me/learningnets

```
4503(config)#cts sxp enable
4503(config)#
*Aug  9 06:51:04.000: %CTS-5-SXP_STATE_CHANGE: CTS SXP enabled
4503(config)#cts sxp connection peer 10.1.40.1 password default mode peer listener
4503(config)#
*Aug 10 09:15:15.564: %CTS-6-SXP_TIMER_START: Connection <0.0.0.0, 0.0.0.0> retry
open timer started.
*Aug 10 09:15:15.565: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.1,
10.1.40.2>-1 state changed from Off to Pending_On.
4503(config)#
*Aug 10 09:15:15.566: %CTS-3-SXP_CONN_STATE_CHG_OFF: Connection <10.1.40.1,
10.1.40.2>-1 state changed from Pending_On to Off.
4503(config)#cts sxp default password TrustSec123
4503(config)#
*Aug 10 09:17:20.936: %CTS-5-SXP_DFT_PASSWORD_CHANGE: CTS SXP password changed.
```

https://t.me/learningnets

```
C6K-DIST(config)#cts sxp enable
C6K-DIST(config)#
Aug 10 16:16:25.719: %CTS-6-SXP_TIMER_START: Connection <0.0.0.0, 0.0.0.0> retry
open timer started.
C6K-DIST(config)#cts sxp default password TrustSec123
C6K-DIST(config)#cts sxp connection peer 10.1.40.2 password default mode peer
speaker
C6K-DIST(config)#
Aug 10 16:17:26.687: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.2, 10.1.40.1>-1
state changed from Off to Pending_On.
Aug 10 16:17:26.687: %CTS-6-SXP_CONN_STATE_CHG: Connection <10.1.40.2, 10.1.40.1>-1
state changed from Pending_On to On.
```

https://t.me/learningnets

```
C6K-DIST#sho cts sxp connections brief
 SXP              : Enabled
 Default Password : Set
 Default Source IP: Not Set
Connection retry open period: 120 secs
Reconcile period: 120 secs
Retry open timer is not running

-------------------------------------------------------------------------------
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
```

https://t.me/learningnets

```
C3750X(config)#cts role-based enforcement
C3750X(config)#interface Ten 1/1/1
C3750X(config-if)#cts manual
C3750X(config-if-cts-manual)#policy static sgt 2 trusted
```

https://t.me/learningnets

```
C3750X#sho cts interface Ten 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: ""
    Authorization Status: SUCCEEDED
        Peer SGT:         2
        Peer SGT assignment: Trusted
    SAP Status:           NOT APPLICABLE
        Configured pairwise ciphers:
            gcm-encrypt
            null

        Replay protection:   enabled
        Replay protection mode: STRICT

        Selected cipher:

    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:             0
        sap fail:                0
        authz success:           3
        authz fail:              0
        port auth fail:          0

    L3 IPM: disabled.
```

https://t.me/learningnets

```
C6K-DIST(config)#platform cts egress
C6K-DIST(config)#cts role-based enforcement
C6K-DIST(config)#interface Ten1/5
C6K-DIST(config-if)#cts manual
C6K-DIST(config-if-cts-manual)#policy static sgt 2 trusted
```

https://t.me/learningnets

```
C6K-DIST#show cts interface Ten1/5
Global Dot1x feature is Enabled
Interface TenGigabitEthernet1/5:
    CTS is enabled, mode: MANUAL
    IFC state:              OPEN
    Authentication Status: NOT APPLICABLE
        Peer identity:     "unknown"
        Peer's advertised capabilities: ""
    Authorization Status: SUCCEEDED
        Peer SGT:          2
        Peer SGT assignment: Trusted
    SAP Status:             NOT APPLICABLE
        Configured pairwise ciphers:
            gcm-encrypt
            null

        Replay protection:   enabled
        Replay protection mode: STRICT

        Selected cipher:

    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:                0
        sap fail:                   0
        authz success:              1
        authz fail:                 0
        port auth fail:             0

    L3 IPM: disabled.
```

https://t.me/learningnets

```
NX7K-CORE(config)# feature dot1x
NX7K-CORE(config)# cts enable
NX7K-CORE(config)# cts role-based enforcement
NX7K-CORE(config)# int eth1/26
NX7K-CORE(config-if)# cts manual
NX7K-CORE(config-if-cts-manual)# policy static sgt 0x2 trusted
```

https://t.me/learningnets

```
C3750X(config-if)#authentication linksec policy ?
  must-not-secure  Never secure sessions
  must-secure      Always secure sessions
  should-secure    OPTIONALLY secure sessions
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
C3750X# 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
```

https://t.me/learningnets

```
C3750X#conf t
Enter configuration commands, one per line. End with CNTL/Z.
C3750X(config)#int Ten1/1/1
C3750X(config-if)#cts manual
C3750X(config-if-cts-manual)#sap pmk 26 mode-list gcm-encrypt
C3750X(config-if-cts-manual)#end
C3750X#
```

https://t.me/learningnets

```
C3750X#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 00000000000000000000000000000000000000000000000000000000000000026
mode-list gcm-encrypt
 no macro auto processing
end
```

https://t.me/learningnets

```
C3750X#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.

C3750X#
```

https://t.me/learningnets

```
ip access-list extended CPP-REDIRECT-ACL
          ! Don't redirect DHCP traffic
          deny udp any any eq bootps
          ! Don't redirect DNS traffic
          deny udp any any eq domain
          ! Don't redirect traffic to ISE
          deny ip any host 10.1.100.232
          ! Redirect all other traffic
          permit ip any any
```

https://t.me/learningnets

```
KR-3750#show authentication sessions interface gigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B860000004E1047F6F4
      Acct Session ID:  Unknown
               Handle:  0xCB00002B
       Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:
         URL Redirect:  https://atw-cp-ise02.ise.local:8443/guestportal/
gateway?sessionId=0A742B860000004E1047F6F4&action=cpp
     URL Redirect ACL:  CPP-REDIRECT-ACL
              ACS ACL:  xACSACLx-IP-CPP-DACL-54d081dd

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

https://t.me/learningnets

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B860000005412E727CB
      Acct Session ID:  0x00000068
               Handle:  0x7A000030
       Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:
         URL Redirect:  https://atw-cp-ise02.ise.local:8443/guestportal/
gateway?sessionId=0A742B860000005412E727CB&action=cpp
     URL Redirect ACL:  CPP-REDIRECT-ACL
              ACS ACL:  xACSACLx-IP-CPP-DACL-54d09330

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

https://t.me/learningnets

```
KR-3750#show authentication sessions interface GigabitEthernet 1/0/3 details
            Interface:  GigabitEthernet1/0/3
          MAC Address:  0017.f2cb.5c34
         IPv6 Address:  Unknown
         IPv4 Address:  10.116.43.138
            User-Name:  00-17-F2-CB-5C-34
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
    Common Session ID:  0A742B86000000581446F323
      Acct Session ID:  0x0000006F
               Handle:  0x5C000033
       Current Policy:  POLICY_Gi1/0/3

Local Policies:
          Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
      Security Policy:  Should Secure
      Security Status:  Link Unsecure

Server Policies:

Method status list:
       Method           State
       dot1x            Stopped
       mab              Authc Success
```

https://t.me/learningnets

```
3750-X#show aaa server | i 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
```

https://t.me/learningnets

Example 21-4 Enabling IOS Load Balancing

```
3750-X(config)# radius-server load-balance method least-outstanding batch-size 5
```

https://t.me/learningnets

```
atw-cp-ise02/admin# show logging application
    5244288 Nov 03 2014 02:15:01  ad_agent.log.4
    5243450 Dec 05 2014 06:30:01  ad_agent.log.1
      22742 Dec 16 2014 16:33:31  mnt-report.log
```

https://t.me/learningnets

```
atw-cp-ise02/admin# show logging application ad_agent.log
2014-12-05T06:37:51.031851+00:00 atw-cp-ise02 adclient[7795]: WARN  <bg:machinep
w> daemon.main computer password change failed: Last error was unexpected discon
nect (GC); deferring reconnect.
2014-12-05T06:38:42.005040+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Start trusted domain discovery
2014-12-05T06:38:42.005074+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Delay /etc/krb5.conf update, Skipping trusted domains update beca
use the system is in disconnected mode
2014-12-05T06:48:42.004611+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Start trusted domain discovery
2014-12-05T06:48:42.004660+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Delay /etc/krb5.conf update, Skipping trusted domains update beca
use the system is in disconnected mode
2014-12-05T06:58:42.004231+00:00 atw-cp-ise02 adclient[7795]: INFO  <bg:krb5.con
f> daemon.main Start trusted domain discovery
```

https://t.me/learningnets

```
atw-cp-ise02/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
```

https://t.me/learningnets

```
3560-X#show run | include radius-server host
radius-server host 10.1.100.232 auth-port 1812 acct-port 1813 test username
radius-test key Cisco123
3560-X#
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE="eth0"
HWADDR="00:50:56:B8:BC:08"
ONBOOT="yes"
NM_CONTROLLED="yes"
BOOTPROTO=none
IPADDR=10.1.100.229
NETMASK=255.255.255.0
USERCTL=yes
TYPE=Ethernet
DNS1=10.1.100.103
GATEWAY=10.1.100.1
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# chkconfig network on
[root@atw-dogtag01 ~]# service network restart
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# vi /etc/yum.conf
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# cat /etc/yum.conf
[main]
cachedir=/var/cache/yum/$basearch/$releasever
keepcache=0
debuglevel=2
logfile=/var/log/yum.log
exactarch=1
obsoletes=1
gpgcheck=1
plugins=1
installonly_limit=3
proxy=http://proxy.woland.com:8080
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# yum install yum-plugin-fastestmirror
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# yum update
Loaded plugins: fastestmirror
Determining fastest mirrors

<<SNIP>>

Transaction Summary
================================================================================
Install       4 Package(s)
Upgrade     104 Package(s)

Total download size: 89 M
Is this ok [y/N]:
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# yum install ntp
[root@atw-dogtag01 ~]# chkconfig ntpd on
[root@atw-dogtag01 ~]# ntpdate [ip-address]
31 Jul 13:47:44 ntpdate[11361]: step time server [ip-address] offset 64.503042 sec
[root@atw-dogtag01 ~]# /etc/init.d/ntpd start
Starting ntpd (via systemctl):                             [  OK  ]
[root@atw-dogtag01 ~]#
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# yum install 389-ds
[root@atw-dogtag01 ~]# useradd ds389
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]#  /usr/sbin/setup-ds.pl
```

https://t.me/learningnets

```
================================================================================
The server must run as a specific user in a specific group.
It is strongly recommended that this user should have no privileges
on the computer (i.e. a non-root user). The setup procedure
will give this user/group some permissions in specific paths/files
to perform server-specific operations.

If you have not yet created a user and group for the server,
create this user and group using your native operating
system utilities.

System User [nobody]: ds389
System Group [nobody]: ds389
```

https://t.me/learningnets

```
Directory Manager DN [cn=Directory Manager]:
Password:
Password (confirm):
Your new DS instance 'atw-dogtag01' was successfully created.
Exiting . . .
Log file is '/tmp/setupo0Vx6g.log'
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# yum install php
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package php.i686 0:5.3.13-1.fc15 will be installed
--> Processing Dependency: php-common(x86-32) = 5.3.13-1.fc15 for package: php-
5.3.13-1.fc15.i686
--> Processing Dependency: php-cli(x86-32) = 5.3.13-1.fc15 for package: php-
5.3.13-1.fc15.i686
--> Running transaction check
---> Package php-cli.i686 0:5.3.13-1.fc15 will be installed
---> Package php-common.i686 0:5.3.13-1.fc15 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

================================================================================
 Package            Arch         Version                  Repository       Size
================================================================================
Installing:
 php                i686         5.3.13-1.fc15            updates         1.1 M
Installing for dependencies:
 php-cli            i686         5.3.13-1.fc15            updates         2.2 M
 php-common         i686         5.3.13-1.fc15            updates         547 k

Transaction Summary
================================================================================
Install       3 Package(s)
```

https://t.me/learningnets

```
Total download size: 3.9 M
Installed size: 13 M
Is this ok [y/N]: y
Downloading Packages:

Running Transaction
  Installing : php-common-5.3.13-1.fc15.i686                                1/3
  Installing : php-cli-5.3.13-1.fc15.i686                                   2/3
  Installing : php-5.3.13-1.fc15.i686                                       3/3

Installed:
  php.i686 0:5.3.13-1.fc15

Dependency Installed:
  php-cli.i686 0:5.3.13-1.fc15          php-common.i686 0:5.3.13-1.fc15

Complete!
[root@atw-dogtag01 ~]#
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# service httpd start
Starting httpd (via systemctl):                            [  OK  ]
[root@atw-dogtag01 ~]# service dirsrv start
Starting dirsrv:
    atw-dogtag01... already running                        [  OK  ]
[root@atw-dogtag01 ~]# chkconfig dirsrv on
[root@atw-dogtag01 ~]# chkconfig httpd on
[root@atw-dogtag01 ~]#
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# yum install pki-ca
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package pki-ca.noarch 0:9.0.20-1.fc15 will be installed
--> Processing Dependency: pki-selinux = 9.0.20-1.fc15 for package:
pki-ca-9.0.20-1.fc15.noarch
--> Processing Dependency: pki-common = 9.0.20-1.fc15 for package: pki-ca-9.0.20-1.
fc15.noarch
--> Processing Dependency: pki-ca-theme >= 9.0.0 for package: pki-ca-9.0.20-1.fc15.
noarch
--> Running transaction check
---> Package dogtag-pki-ca-theme.noarch 0:9.0.11-1.fc15 will be installed
--> Processing Dependency: dogtag-pki-common-theme = 9.0.11-1.fc15 for package:
dogtag-pki-ca-theme-9.0.11-1.fc15.noarch
---> Package pki-common.noarch 0:9.0.20-1.fc15 will be installed
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# service iptables stop
Stopping iptables (via systemctl):                         [  OK  ]
[root@atw-dogtag01 ~]# chkconfig iptables off
[root@atw-dogtag01 ~]#
```

https://t.me/learningnets

```
pkicreate    -pki_instance_root=/var/lib -pki_instance_name=ise-ca -subsystem_
type=ca -agent_secure_port=9443 -ee_secure_port=9444 -ee_secure_client_auth_
port=9446 -admin_secure_port=9447 -unsecure_port=9180 -tomcat_server_port=9701
-user=pkiuser -group=pkiuser -redirect conf=/etc/ise-ca -redirect logs=/var/log/
ise-ca -verbose
```

https://t.me/learningnets

```
pkiremove -pki_instance_root=/var/lib -pki_instance_name=ise-ca
```

https://t.me/learningnets

```
Installation information recorded in /var/log/ise-ca-install.log.
[debug] run_command(/sbin/service pki-cad restart ise-ca)
Before proceeding with the configuration, make sure
the firewall settings of this machine permit proper
access to this subsystem.

Please start the configuration by accessing:

https://atw-dogtag01.ise.local:9447/ca/admin/console/config/login?pin=UUVMDHRvTojQ
rdeod91e

After configuration, the server can be operated by the command:

    /sbin/service pki-cad restart ise-ca
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# /sbin/service pki-cad restart ise-ca
Stopping ise-ca:                                           [  OK  ]
Starting ise-ca:                                           [  OK  ]
[root@atw-dogtag01 ~]#
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# cp /etc/ise-ca/CS.cfg /etc/ise-ca/CS.cfg.bak
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# vi /etc/ise-ca/CS.cfg
```

https://t.me/learningnets

```
ca.scep.allowedEncryptionAlgorithms=DES3
ca.scep.allowedHashAlgorithms=SHA1,SHA256,SHA512
ca.scep.enable=true
ca.scep.encryptionAlgorithm=DES3
ca.scep.hashAlgorithm=SHA256
ca.scep.nonceSizeLimit=16
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# cp /var/lib/ise-ca/profiles/ca/caRouterCert.cfg /var/lib/
ise-ca/profiles/ca/caRouterCert.cfg.bak
```

https://t.me/learningnets

```
[root@atw-dogtag01 ~]# vi /var/lib/ise-ca/profiles/ca/caRouterCert.cfg
```

https://t.me/learningnets

```
[root@atw-dogtag01 ise-ca]# cat /var/lib/ise-ca/profiles/ca/caRouterCert.cfg
desc=This certificate profile is for enrolling router certificates.
visible=false
enable=true
enableBy=admin
auth.instance_id=
name=One Time Pin Router Certificate Enrollment
input.list=i1,i2
input.i1.class_id=certReqInputImpl
input.i2.class_id=submitterInfoInputImpl
output.list=o1
<<SNIP>>
```

https://t.me/learningnets

```
[root@atw-dogtag01 ise-ca]# service pki-cad restart
Stopping ise-ca:                                           [  OK  ]
Starting ise-ca:                                           [  OK  ]
```

https://t.me/learningnets

```
[root@atw-dogtag01 ise-ca]# mv /etc/httpd/conf.d/welcome.conf /etc/httpd/conf.d/
welcome.conf.bak
```

https://t.me/learningnets

```
[root@atw-dogtag01 ise-ca]# vi /var/www/html/scepproxy.php
```

https://t.me/learningnets

```
<?php
$ops = $_GET['operation'];
$msg= $_GET['message'];
$order   = array("\r\n", "\n", "\r", " ");
$msg = str_replace($order, "", $msg);
$msg = rawurldecode($msg);
$msg=urlencode($msg);
if ($ops == "GetCACaps")
{
echo "";
}
else
{
$url = "http://127.0.0.1/ca/cgi-bin/pkiclient.exe?operation=".$ops."&message=".$
msg;
$ch = curl_init();
curl_setopt($ch, CURLOPT_PORT, 9180);
curl_setopt($ch, CURLOPT_URL, $url);
curl_setopt($ch, CURLOPT_HEADER, 0);
curl_setopt($ch, CURLOPT_RETURNTRANSFER, 1);
//curl_setopt($ch, CURLOPT_POST, 1);
$body = curl_exec($ch);
curl_close($ch);
if ($ops=="PKIOperation")
{
header("Content-Type: application/x-pki-message");
}
else
{
header("Content-Type: application/x-x509-ca-cert");
}
echo $body;
}
?>
```

https://t.me/learningnets

```
[root@atw-dogtag01 ise-ca]# service httpd restart
Restarting httpd (via systemctl):                          [  OK  ]
[root@atw-dogtag01 ise-ca]#
```

https://t.me/learningnets

```
3560-X# show 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
!
!Dynamic Authorization is the official term for CoA, which allows a RADIUS server
to send commands to the NAD to change the authorization for the endpoint
aaa server radius dynamic-author
 client 10.1.103.232 server-key Cisco123
!
aaa session-id common
authentication mac-move permit
ip routing
!
```

https://t.me/learningnets

```
ip domain-name ise.local
ip name-server 10.1.100.100
! IP Device Tracking is required in order to insert the client IP address into
downloadable ACLs (dACLs).
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
```

https://t.me/learningnets

```
! VLAN 41 is the data VLAN for this particular switch.  This will vary in your
environment.
 switchport mode access
 switchport voice vlan 99
! VLAN 99 is the voice VLAN in this particular switch. This will vary in your
environment
 ip access-group ACL-ALLOW in
 authentication event fail action next-method
 authentication event server dead action authorize
 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
!
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
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
! The following logging commands should be used ONLY when troubleshooting,
or doing an initial roll-out.  Sending all the syslog to the MNT node can be
overwhelming and affect performance and storage.
See Chapter 22, "Troubleshooting Tools," for more information.
logging origin-id ip
logging source-interface Loopback0
logging host 10.1.103.4 transport udp port 20514
! Cisco ISE Monitoring uses a nontraditional port for Syslog.  It uses 20514
instead of 514.
!
! The following SNMP commands are used because this switch version does not have
Device Sensor capabilities.
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
!  The definitions of these attributes and their importance are covered in detail
in Chapter 12, "Implement Wired and Wireless Authentication."
radius-server dead-criteria time 5 tries 3
radius-server host 10.1.103.232 auth-port 1812 acct-port 1813 key Cisco123
radius-server vsa send accounting
radius-server vsa send authentication
!
end
```

https://t.me/learningnets

```
C3750X# show 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
```

https://t.me/learningnets

```
!
!
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 Sensors enables ISE to gather raw endpoint information and send that
information to ISE in a RADIUS accounting packet.
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
!
```

https://t.me/learningnets

```
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
!
```

https://t.me/learningnets

```
dot1x system-auth-control
!
interface Loopback0
 ip address 192.168.254.1 255.255.255.255
!
interface <ALL EDGE PORTS>
 switchport access vlan 10
 ! VlAN 10 is the data VLAN for this particular switch.  It may differ in your
environment.
 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
```

https://t.me/learningnets

```
 dot1x pae authenticator
 dot1x timeout tx-period 10
 spanning-tree portfast
 ip dhcp snooping information option allow-untrusted
!
! There are more VLANs on this switch than the 12.x version above, because it was
in a different part of the network with different requirements.  The author has
left the configuration this way to show that these are but sample configurations,
and you have many options.
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
!
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
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
!
!
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-02
 address ipv4 10.1.100.232 auth-port 1812 acct-port 1813
 automate-tester username radius-test
 key Cisco123
!
end
```

https://t.me/learningnets

```
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 10.1.103.232 server-key Cisco123
!
aaa session-id common
!
```

https://t.me/learningnets

```
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
!
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
authentication event server dead action authorize
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
!
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
 permit tcp any any eq www
 permit tcp any any eq 443
 remark all other traffic will be implicitly denied from the redirection
!
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 3radius-server vsa send accounting
radius-server vsa send authentication
!
!
radius server CP-02
 address ipv4 10.1.100.232 auth-port 1812 acct-port 1813
 automate-tester username radius-test
 key Cisco123
!


!
end
```

https://t.me/learningnets

```
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 10.1.103.232 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
```

https://t.me/learningnets

```
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
```

https://t.me/learningnets

```
 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
```

https://t.me/learningnets

```
 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 10.1.103.232 auth-port 1812 acct-port 1813 key Cisco123
radius-server vsa send accounting
radius-server vsa send authentication
!
end
```

https://t.me/learningnets