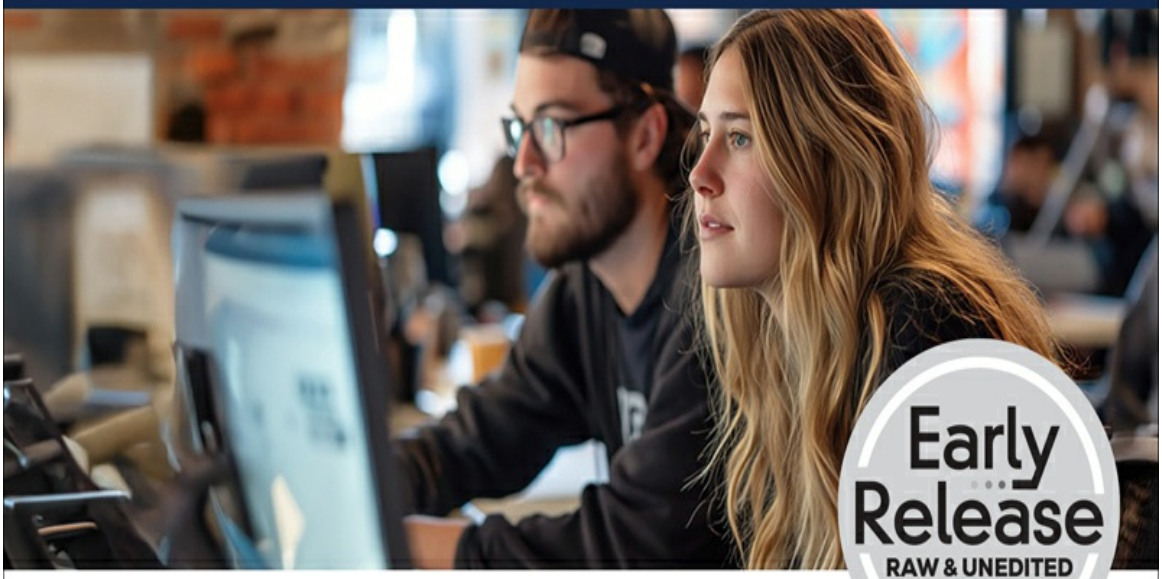
Cisco Identity Services Engine (ISE)

End-to-End Access Control



EUGENE KORNEYCHUK
SERHII KUCHERENKO
JONOTHAN EAVES

ciscopress.com

Cisco Identity Services Engine (ISE): End-to-End Access Control

Eugene Korneychuk, CCIE No. 43253 Serhii Kucherenko, CCIE No. 35182 Jonothan "Joff" Eaves, CCIE No. 4242

Cisco Press

A NOTE FOR EARLY RELEASE READERS

With Early Release eBooks, you get books in their earliest form-the author's raw and unedited content as they write-so you can take advantage of these technologies long before the official release of these titles.

If you have comments about how we might improve the content and/or examples in this book, or if you notice missing material within this title, please reach out to Pearson at PearsonITAcademics@pearson.com

Cisco Identity Services Engine (ISE): End-to-End Access Control

Eugene Korneychuk, Serhii Kucherenko, Jonothan Eaves

Copyright© 2027 Cisco Systems, Inc.

Published by:
Cisco Press
Hoboken, New Jersey

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

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

Please contact us with concerns about any potential bias at www.pearson.com/en-us/report-bias.html.

$PrintCode

Library of Congress Control Number: 2026952290

ISBN-13: 978-0-13-588123-1
ISBN-10: 0-13-588123-4

Warning and Disclaimer

This book is designed to provide information about Cisco Identity Services Engine. 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.

Trademark disclaimer: Android, Google Play, Chrome, Chromebook, and Google Workspace are trademarks of Google LLC. Apple and iOS are trademarks of Apple Inc. Windows is a trademark of Microsoft Corporation. Cisco Systems, Inc. and Cisco Identity Services Engine (ISE) are trademarks of Cisco Systems, Inc. All other trademarks are the

property of their respective owners.

Feedback Information

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

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

We greatly appreciate your assistance.

Please contact us with concerns about any potential bias at https://www.pearson.com/report-bias.html.

GM K12, Early Career and Professional Learning
Soo Kang

Alliances Manager, Cisco Press
Caroline Antonio

Executive Editor
James Manly

Managing Editor
Sandra Schroeder

Development Editor
Ellie C. Bru

Senior Project Editor
Mandie Frank

Copy Editor
Chuck Hutchinson

Technical Editors
Jesse Dubois, Adam Hollifield

Designer
Chuti Prasertsith

Composition
codeMantra

Indexer

Proofreader

About the Authors

Eugene Korneychuk, CCIE No. 43253, is a technical leader on the Security TAC Technical Leadership team at Cisco, where he specializes in Cisco Identity Services Engine (ISE) and broader network access security technologies. With more than 20 years of experience in networking and security, Eugene has worked extensively with enterprise customers to design, deploy, and troubleshoot ISE at scale, resolving some of the most complex escalations across on-premises, hybrid, and cloud environments.

A Cisco Live Hall of Fame speaker, Eugene has delivered numerous well-received sessions on ISE, covering topics such as troubleshooting, cloud deployments, automation, and integration with platforms like Microsoft Entra ID and Intune. Beyond his CCIE, he holds multiple certifications spanning AWS, Microsoft Azure, Terraform, and Cisco DevNet, reflecting his broad expertise across cloud platforms, infrastructure as code, and network automation.

Eugene brings a practitioner's perspective to this book, grounded in real-world TAC experience helping customers navigate the everyday challenges of deploying and maintaining identity-based access control at enterprise scale.

Serhii Kucherenko, CCIE No. 35182, is an Escalations technical leader within the Meraki Customer Support team at Cisco CX, where he specializes in resolving highly complex technical challenges involving

Meraki MX SD-WAN solutions. With a background deeply rooted in network security and access control, Serhii spent six years in the Cisco TAC AAA team, where he provided expert support for Cisco Identity Services Engine, TrustSec, and broad network access security architectures.

A subject matter expert in endpoint security, Serhii focuses extensively on the intricacies of endpoint posture assessment. As a long-time Cisco Live speaker, he has delivered numerous sessions on Cisco ISE and Meraki SD-WAN, consistently sharing his deep technical knowledge with the broader community.

Serhii brings a practitioner's perspective to his work, grounded in years of hands-on TAC experience helping customers navigate the everyday complexities of deploying and maintaining secure, scalable network environments.

Jonothan "Joff" Eaves, CCIE No. 4242, is a technical marketing engineer in the Cisco Campus Networking Group. His specialization is switching security with a focus toward group-based policies, automating network security functions and enabling software-defined segmentation.

Jonothan has presented network segmentation topics many times at Cisco Live, Webinars and Partner conferences, and works with Cisco Engineering teams to develop new features and functions as needed.

Prior to joining that group, Jonothan worked on solution validation covering network access security technologies (including security group mechanisms), voice, WAN, and routing projects and gained CCIE in 1998 (#4242). He has worked in the industry for 37 years, 29 of which have been with Cisco.

About the Technical Reviewers

Jesse Dubois has been supporting Cisco Systems' identity and NAC products for more than 20 years as part of the TAC organization. For the last 12 years, he has been in a technical leadership role with a passion for serviceability not only with Identity Services Engine but also the products ISE interacts with, like Cisco's switches and wireless controllers. Jesse has presented technical seminars, instructor-led labs, and most recently ISE performance, scale, and best practices at Cisco Live in both US and Europe events. He also regularly presents what is top of mind at both Security and Wireless TAC Time events and has been focusing on automated data analysis and problem detection in customer support data.

Adam Hollifield is an enterprise architect who has spent 15 years in the information technology industry performing security-focused network engineering, design, consulting, and advisory work. Prior to his current role, Adam worked for an IT hardware distributor and solutions aggregator as a network engineer helping to drive internal business and security objectives.

While specializing primarily in network segmentation and Zero Trust, Adam also has technical experience with general routing and switching, firewalls, and wireless networking.

Adam holds a bachelor of science in computer engineering from the University of Central Florida. He also holds a number of industry

certifications from multiple vendors. Adam is recognized as a Cisco Community Designated VIP, Cisco Champion, Aruba Airheads MVP Expert, and Fortinet Community Super User.

When not working, Adam is an avid runner and mountain biker. He also enjoys spending time with his wife and their three daughters.

Dedications

Eugene Korneychuk: This book is dedicated to my family, whose support, patience, and understanding have meant so much to me over the years. Thank you for always being there, for sharing in both the challenges and accomplishments, and for supporting me in everything I do.

Serhii Kucherenko: This book is dedicated to my beautiful and amazing wife, Uliana. It would not have been possible without your endless love, encouragement, and support. At the very beginning of my journey into networking, as a student in the Cisco Networking Academy, I often felt less capable and less knowledgeable than everyone around me. It was your belief in me[md]long before I believed in myself[md]that carried me forward and shaped the engineer I have become, and ultimately made this book possible. I am forever grateful for your understanding and patience through the many long nights and weekends I spent writing it.

Jonothan "Joff" Eaves: This book is dedicated to my wife and children, who have been my constant source of inspiration and support throughout my entire career. My professional journey has been incredibly rewarding, and I know that I could not have navigated the challenges or enjoyed the successes of this life without your patience, encouragement, and love. Thank you for being the foundation of everything I do.

Acknowledgments

Eugene Korneychuk: I would like to begin by thanking my co-authors and friends, Serhii and Jonothan. Your knowledge, dedication, and countless hours of collaboration made this book possible. It has been a privilege to work alongside both of you, and I am grateful for everything you brought to this project.

I would also like to express my sincere gratitude to my mentors, Aaron Woland and Gonzalo Salgueiro. Your guidance, encouragement, and willingness to share your knowledge have had a tremendous impact on my professional growth. Your mentorship has helped shape not only this book but also my career, and I am deeply grateful for your continued support over the years.

My sincere thanks go to the technical editors, Jesse Dubois and Adam Hollifield. Your careful reviews, technical expertise, and unique perspectives brought tremendous value to this book. Your thoughtful feedback challenged us to make the content more accurate, practical, and useful to readers. The experience you brought to the review process helped us look at the material from different perspectives and ultimately made this a better book.

A special thank you goes to my colleague, Max Abramsky, for his valuable contribution to one of the chapters. Your expertise and thoughtful input helped strengthen the content, and I truly appreciate the time and effort you dedicated to this project.

I would also like to thank James Manly, the executive editor, for his continued support, encouragement, and guidance throughout this project. I am grateful for your help in bringing this book from an idea to a finished work.

Thank you to Ellie Bru, the development editor for this book, for guiding me through multiple development and editing cycles with patience, professionalism, and care. Your thoughtful feedback, attention to detail, and willingness to answer my many questions made the process both enjoyable and rewarding.

Finally, I would like to thank the many customers, colleagues, and members of the Cisco community whom I have had the privilege of working with throughout my career. Every discussion, technical challenge, shared experience, and collaboration has contributed to my knowledge and professional growth. Much of what I have learned over the years has come from working alongside this remarkable community.

I hope this book gives something back and helps others in the same way that so many people in this community have helped me throughout my career.

Serhii Kucherenko: My path into networking started in a classroom, and so should these thanks. To my Cisco Networking Academy instructor, Dmytro Myronov, thank you for guiding my very first steps in this field. You gave me the foundation everything else was built on, and I still remember how much difference a patient, generous teacher can make.

To my first managers in the IT department at Pravex Bank, Anatoliy Pokosa and Dmytro Khalimonchuk, thank you for taking a chance on a CCNA student and turning him into a working engineer. The habits, discipline, and troubleshooting instincts I picked up under your

supervision have stayed with me for my entire career.

To Konstantin Lesnikov, my Technical Director at CNI, thank you for the trust and the leadership that let me grow into a technical lead, a network architect, and eventually a CCIE. You gave me room to take on problems that were bigger than I was, which turned out to be exactly what I needed.

To my coauthors and friends, Eugene and Jonothan, this book simply does not exist without you. Your expertise, persistence, and the countless hours we spent debating, drafting, and rewriting together made it what it is. Working beside you has been a genuine privilege.

To our technical editors, Jesse Dubois and Adam Hollifield, thank you for the rigor you brought to every chapter. Your reviews, deep technical knowledge, and different points of view pushed us toward greater accuracy and real-world usefulness, and the book is measurably better because you challenged us.

To James Manley, our executive editor, thank you for the steady support, guidance, and encouragement that carried this project from a rough idea to a finished book.

To Ellie Bru, our development editor, thank you for walking us through cycle after cycle of development and editing with patience, professionalism, and remarkable attention to detail. Your feedback and your willingness to answer an endless stream of questions made the process genuinely enjoyable.

Finally, thank you to the customers, colleagues, and members of the Cisco community I have had the privilege of working with over the years. Every escalation, every late-night troubleshooting session, and every shared lesson has shaped how I think about this technology. Most of what I know, I learned alongside this remarkable community.

I hope the years of distilled knowledge we translated into this book will help others on their path, so that the transfer of knowledge continues and benefits those who come after us.

Jonothan "Joff" Eaves: Writing a book on a subject as complex and evolving as Cisco Identity Services Engine is a significant undertaking that relies on the support and inspiration of many talented individuals. I am incredibly grateful to those who have shaped my professional journey and contributed to the insights presented in these pages.

My appreciation begins with Dave Brooks, who played a pivotal role early in my career. His influence taught me that even the most demanding technical challenges are best navigated when the work itself is enjoyable and fun, a philosophy I have carried with me ever since.

I am also deeply indebted to Nigel Townley, whose dedication to his reports and genuine interest in seeing others flourish have been a constant source of motivation. His leadership style set a standard for how I approach my own professional relationships and development.

Kevin Regan stands out as a true technology leader and innovator. His passion for teamwork and his ability to foster an environment where enjoyment and success go hand in hand have been instrumental in my growth. It is that same spirit of collaborative success that I have aimed to capture throughout this book.

Darrin Miller is without a doubt one of the most technical people I have ever had the pleasure of working with. Beyond his technical prowess, he has been a fantastic mentor and teacher, always willing to share his knowledge and guide me through the nuances of our field.

Working with Fay-Ann Lee has been a highlight of my career. She is an exceptional colleague, and our ability to bounce ideas off one another

has been vital to my thinking. We have always worked well as a team, and her contributions have undoubtedly strengthened the concepts discussed here.

Finally, I must thank Einar Nilsen-Nygaard, who is one of the most talented technical developers I have encountered. His unwavering dedication to doing what is right and pursuing the best possible outcome has served as a benchmark for excellence that I continue to strive for in my own work.

To all these individuals, thank you for your time, your wisdom, and your friendship. Your collective impact on my career has been profound, and this book is a reflection of the knowledge and values you have helped me cultivate.

Contents at a Glance

1. Introduction to Modern Access Control
2. ISE Deployment Principles
3. Implementing ISE in Hybrid Infrastructure
4. Policy Set Concepts
5. Authentication Options
6. Authorization Options
7. Introduction to TrustSec
8. Device Identification for Access Management
9. Managing Guest Users and Personal Devices
10. Implementing Endpoint Posture Services
11. Managing Network Access in the Meraki Environment
12. Meraki Adaptive Policy
13. Remote Access Management for On-Prem and SSE
14. Managing Network Access on Campus with Microsegmentation
15. Extending Microsegmentation to WAN and SD-WAN
16. Building Policy-Interconnect with Data Center
17. Classifying Cloud Workloads
18. ISE APIs and Advanced Automation Use Cases
19. Integrating Your Security Products to Enhance Context Using

pxGrid
20. Troubleshooting, Monitoring, Alerting, and System 360
21. ISE Maintenance

Reader Services

Register your copy at www.ciscopress.com/title/ISBN for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account*. Enter the product ISBN 9780135881231 and click **Submit**. When the process is complete, you will find any available bonus content under Registered Products.

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

Contents

Chapter 1. Introduction to Modern Access Control

- Network Evolution
- Network Access Control Response to the Changes

Chapter 2. ISE Deployment Principles

- Introduction
- ISE Node Personas
- ISE Deployment Models
- ISE Deployment Sizing
- ISE VM Disk Provisioning Guidance
- ISE Installation
- ISE Setup
- ISE Installation Verification
- Latency Requirements
- High Availability
- Building Distributed Deployment
- Licensing
- Summary

Chapter 3. Implementing ISE in Hybrid Infrastructure

- Introduction

ISE Cloud Deployments
ISE on AWS
ISE on Azure
Summary
Chapter 4. Policy Set Concepts
Introduction
Policy Sets Overview
Conditions Studio
RADIUS Server Sequence
Allowed Protocols
Policy Set Hit Counts
Eduroam
Summary
Chapter 5. Authentication Options
Introduction
Authentication Policies Concepts
Identity Sources
Summary
Chapter 6. Authorization Options
Introduction
Authorization Policies Concepts
Summary
Chapter 7. Introduction to TrustSec
Introduction
Classification

Propagation
Enforcement
Summary
Chapter 8. Device Identification for Access Management
Introduction
Context Visibility
Certainty Factor–Based Profiling
Multi-Factor Classification Profiling
Additional Profiler Settings
Summary
Chapter 9. Managing Guest Users and Personal Devices
Introduction
Guest Access and BYOD
Guest Access Components
BYOD Building Blocks
Summary
Chapter 10. Implementing Endpoint Posture Services
Introduction
ISE Posture Service Pillars
ISE Posture Flows
Stitching Attributes for Authorization
Configuring ISE Posture
MDM Flows
Configuring ISE MDM
Summary

Chapter 11. Managing Network Access in the Meraki Environment
Introduction
Cisco Meraki Dashboard and Cloud Networking Principles
Summary
Chapter 12. Meraki Adaptive Policy
Introduction
Microsegmentation Evolution in Meraki Networks
Implementing Adaptive Policy at the Organization Level
Implementing Adaptive Policy at the Device Level
Adaptive Policy Verification and Troubleshooting
Summary
Chapter 13. Remote Access Management for On-Prem and SSE
Introduction
Remote Access Evolution
Common NAC Use Cases with Cisco Secure Firewall
Secure Firewall Remote Access VPN: General Configuration Elements
Secure Firewall Remote Access VPN: Following the RA VPN Wizard
Secure Firewall Remote Access VPN: Configuring NAC Use Cases
Secure Firewall Remote Access VPN with ISE: Verification and Troubleshooting
Network Access Control for Cisco SSE with ISE
SSE: Putting Everything Together

SSE Remote Access VPN with ISE: Verification and Troubleshooting
Summary
Chapter 14. Managing Network Access on Campus with Microsegmentation
Introduction
Where to Enforce Traffic
Dynamic SGT Assignment
Static SGT Classification
Unknown SGT
Default Route SGT
Mitigating Malware Lateral Movement
Enforcement Using TCP Flags
PAC-less Operation
Use Cases
Reporting and Troubleshooting
Summary
Chapter 15. Extending Microsegmentation to WAN and SD-WAN
Introduction
SGT Propagation Across SD-WAN
Integrating SD-WAN Manager with ISE
SD-WAN Manager Subscribing to Security Group Tag (SGT)
Using ISE Context in SD-WAN Manager Unified Policy
SD-WAN Manager Subscribing to User/User Groups
SD-WAN Edge Devices Using SXP

Summary

Chapter 16. Building Policy-Interconnect with Data Center

Introduction

Prerequisites for Integrating ISE and APIC

ISE SGT Domains

Adding Workload Connection in ISE for APIC Integration

ISE Primary and Secondary SGTs via the Workload Classification Table

Learned Groups in ISE Available in Matrix and Authorization

APIC Learning Groups from ISE

ISE-ACI High Availability

Editing, Suspending, and Deleting Workload Connections

ACI Connectivity Report

Summary

Chapter 17. Classifying Cloud Workloads

Introduction

AWS EC2 Workloads and Access Rights

Workload Connections in ISE

ISE Workload Classification

ISE Inbound SGT Domain Rules

ISE Reporting on Workload Connections

Summary

Chapter 18. ISE APIs and Advanced Automation Use Cases

Introduction
API Gateway
OpenAPI
ERS API
MnT API
ISE Automation with Terraform
ISE Automation with Ansible
Summary

Chapter 19. Integrating Your Security Products to Enhance Context Using pxGrid
Introduction
Understanding pxGrid Architecture
ISE to Catalyst Center pxGrid Integration
ISE to Cisco Secure Firewall pxGrid Integration
pxGrid Direct
Summary

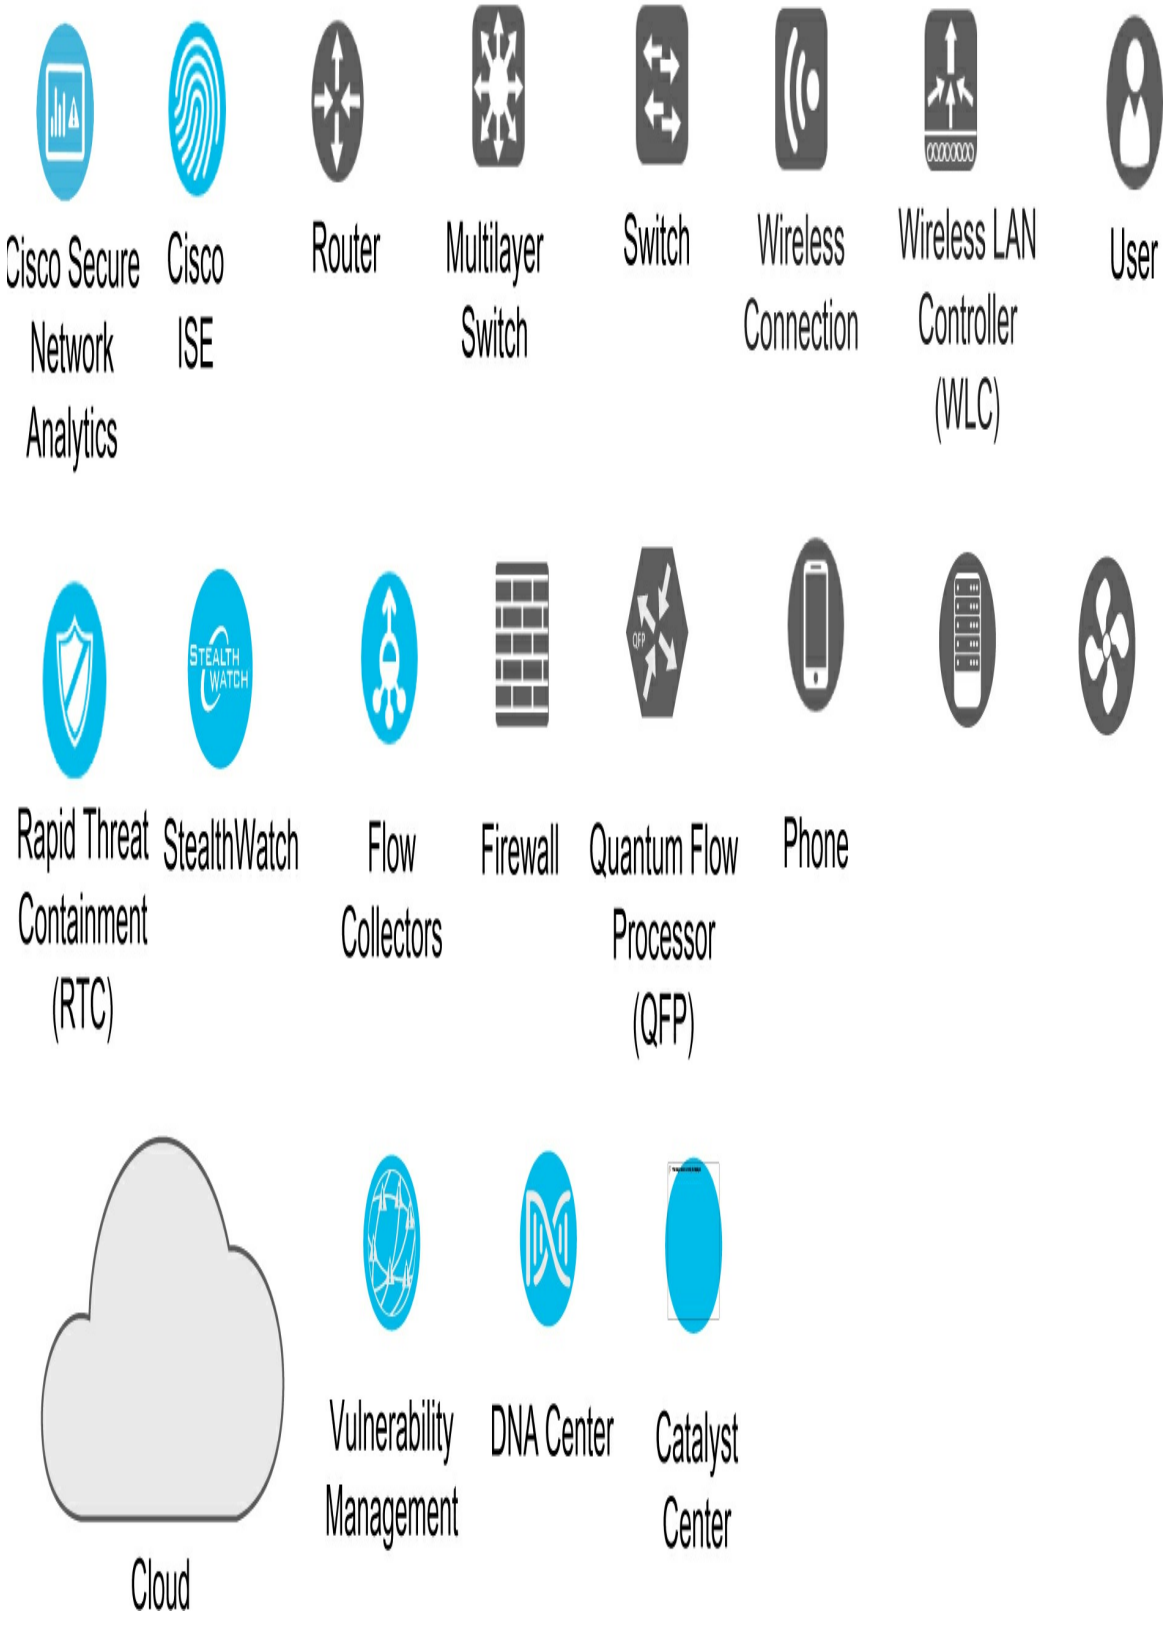
Chapter 20. Troubleshooting, Monitoring, Alerting, and System 360
Introduction
Troubleshooting Methodology
Alarms
Live Logs, Live Sessions, and Detailed Authentication Reports
Diagnostic Tools
Summary

Chapter 21. ISE Maintenance

Introduction
Repositories
Backup and Restore
ISE Upgrades
Patch Installation
Summary

Icons Used in This Book

1231_FMicons.png



Technet24

Command Syntax Conventions

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

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

Introduction

This book is the product of years of hands-on experience[md]designing, deploying, and troubleshooting Cisco Identity Services Engine (ISE) deployments across small campus environments and massive, multisite enterprise infrastructures. Rather than simply presenting configuration steps, this book focuses on the reasoning behind ISE design decisions: why certain deployment models scale better than others, why policy structure matters as much as policy content, and how the shift toward Zero Trust has fundamentally changed the role identity plays in network access control.

A core part of this book also focuses on how ISE operates across the full breadth of a modern hybrid environment[md]not just traditional on-premises campus networks, but hybrid and cloud deployments, Meraki-managed environments, remote access and SSE, data center integration, and multicloud workload classification. Identity-based access control is only as strong as its weakest integration point, and this book methodically walks through how to extend consistent, identity-driven policy across every one of these environments.

By the end of this book, you will have gained expert-level knowledge about the following topics:

- Core ISE deployment principles, sizing, high availability, and hybrid/cloud deployment models across AWS, Azure, and OCI

Technet24

- Policy set design, authentication and authorization options, including certificate-based, MFA, and external identity source integrations
- Cisco TrustSec and microsegmentation, from campus deployments through WAN/SD-WAN and data center policy interconnects
- Device identification and profiling, including multi-factor classification and Cloud MFC profiling
- Guest access and endpoint posture services
- Ways to extend ISE to Meraki environments, including Meraki Adaptive Policy
- Remote access management for on-premises and Security Service Edge (SSE) architectures
- Steps to classify and secure cloud workloads with dynamic, tag-based security group tag assignment
- ISE APIs, advanced automation use cases, and integration of third-party security products via pxGrid
- Troubleshooting methodology, monitoring, alerting, and System 360, along with day-to-day ISE maintenance and upgrades

Although this book is not written with the intent of helping you pass a specific certification exam, it does act as an excellent supplemental resource for studying toward Cisco Identity Services Engine[nd]focused certifications, including the CCNP Security and Cisco Certified Specialist[nd]Security Identity Management Implementation tracks.

Who Should Read This Book?

This book is designed for anyone tasked with securing the modern digital perimeter, where the traditional boundaries of the corporate network have effectively dissolved. As we move toward a Zero Trust architecture, identity has become the new firewall, and Cisco Identity Services Engine stands as the central nervous system for that strategy. Whether you are managing a small campus environment or orchestrating a massive global infrastructure that spans data centers, cloud workloads, and remote access points, the principles of identity-based security remain the same.

The primary audience for this book includes network engineers and security architects who are responsible for designing, deploying, and maintaining access control policies. If you find yourself bridging the gap between networking and security, this book will serve as your primary reference. We have structured the content to be accessible to those just beginning their journey with ISE, while providing the technical depth required by seasoned professionals who need to master advanced integrations like pxGrid, API-driven automation, and complex cloud-to-data-center segmentation.

For systems administrators and operations teams, this book provides a clear path through the complexities of daily maintenance and troubleshooting. We understand that in a fast-paced environment, you need more than just theory; you need actionable guidance on how to manage endpoint posture, handle guest access, and ensure that personal devices do not compromise the integrity of your corporate resources. By focusing on practical deployment principles and real-world policy requirements, we aim to move you from basic connectivity to a robust, automated security posture.

The level of expertise required to benefit from this book is intentionally broad. If you are a junior engineer looking to earn your certifications or simply trying to understand how to get your first ISE

Technet24

policy set up, you will find the introductory chapters and deployment principles to be a solid foundation. For those at the mid-level, the chapters on authentication, authorization, and TrustSec will help you refine your existing configurations and move toward more granular control. For the senior architects and technical leads who are tasked with integrating ISE into complex, hybrid environments, the later chapters on cloud workload classification, API automation, and data center integration will provide the advanced knowledge needed to push the boundaries of what your infrastructure can achieve.

Ultimately, we want this book to be a constant companion on your desk. We have avoided overly academic language in favor of a conversational, practical style that reflects the reality of working in IT today. Our goal is to help you understand not just how to configure a specific feature, but how to think about identity as a holistic security strategy. By the time you reach the final chapter on troubleshooting and maintenance, you should feel confident in your ability to design a resilient, scalable, and secure environment that supports the needs of your organization, regardless of the industry or the complexity of your network. Whether you are working in a traditional on-prem environment or navigating the complexities of a multicloud transition, the insights provided here will help you ensure that your users and devices are always verified, authorized, and secure.

How This Book Is Organized

Chapter 1, "Introduction to Modern Access Control": This chapter introduces the fundamental shift in enterprise networking, tracing the evolution from rigid, location-based "castle-and-moat" architectures to the modern, identity-centric models necessitated by today's hybrid, cloud-first environments. It explains why traditional, port-based access control is no longer sufficient and details how the transition to a Zero

Trust framework requires a more intelligent, context-aware approach to security. By positioning Cisco Identity Services Engine as a central orchestrator, the chapter sets the stage for the rest of the book, demonstrating how modern NAC integrates with broader technologies like SD-WAN and SASE to provide consistent, continuous verification for users and devices, regardless of their connection point.

Chapter 2, "ISE Deployment Principles": This chapter provides a comprehensive framework for planning and building a production-ready Cisco ISE deployment. It covers the four ISE node personas[md]Policy Administration, Policy Service, Monitoring and Troubleshooting, and pxGrid. It explains how to combine them into small, medium, or large deployment models based on the number of active endpoints. The chapter walks through installation methods across hardware appliances, virtual machines, and OpenStack, along with initial setup and post-installation verification. It also addresses critical design considerations such as latency requirements, disk provisioning, and high availability for PSN, MnT, and PAN nodes, as well as NIC bonding and ISE licensing, ensuring administrators can build a resilient, scalable cluster.

Chapter 3, "Implementing ISE in Hybrid Infrastructure": This chapter explores how to extend Cisco ISE deployments into the cloud to support hybrid infrastructure environments. It covers native ISE support on AWS, Azure, and OCI, detailing the available deployment methods for each platform, including Cloud Formation Templates on AWS and Azure Application deployment on Azure. The chapter walks through the installation process step by step for both AWS and Azure, using real console screenshots, and highlights the key limitations administrators should consider for each cloud provider. It concludes by emphasizing the flexibility ISE offers[md]whether organizations choose a fully cloud-based deployment or a hybrid model combining

Technet24

on-premises and cloud instances.

Chapter 4, "Policy Set Concepts": This chapter explains how Cisco ISE uses policy sets to logically group authentication and authorization policies, enabling administrators to manage complex access requirements in a scalable, organized way. It covers building policy sets with the Conditions Studio, configuring RADIUS server sequences to proxy authentication to upstream RADIUS servers, and restricting allowed protocols within a policy Set. The chapter also introduces policy set hit counts for monitoring and optimizing policy usage, along with a dedicated look at eduroam configuration for educational roaming use cases[md]giving you a complete picture of how ISE evaluates policies sequentially to grant access.

Chapter 5, "Authentication Options": This chapter explores the range of authentication options available in Cisco ISE, starting with built-in authentication policies and identity sources. It covers internal identity sources such as internal users and endpoints, along with certificate management using the System and Trusted Certificate Stores. The chapter then details external identity sources, including Active Directory, LDAP, and Microsoft Entra ID integrations, as well as native multifactor authentication with Duo. It closes with a look at network access devices.

Chapter 6, "Authorization Options": This chapter examines how Cisco ISE grants and controls network access through authorization policies and authorization profiles, following the principle of least privilege. It covers Change of Authorization (CoA), which lets ISE act on already-authenticated sessions, and explains how authorization policy conditions leverage a wider range of attributes than authentication policies. The chapter details authorization profiles and their results, including downloadable ACLs (DACLs), security group tags (SGTs), VLAN assignment, and URL redirection for use cases like Guest, BYOD,

and Posture. It closes with a look at network device profiles, which extend consistent authorization behavior to third-party network devices.

Chapter 7, "Introduction to TrustSec": This chapter introduces Cisco TrustSec as a software-defined segmentation framework that simplifies security policy management by decoupling access control from static network topologies. It explains the three fundamental pillars of the technology: classification, propagation, and enforcement. By leveraging centralized policy management through Cisco Identity Services Engine, the chapter demonstrates how organizations can achieve better visibility, implement least-privilege segmentation, and reduce the operational impact of security breaches.

Chapter 8, "Device Identification for access management**":** This chapter explains how Cisco ISE identifies and classifies devices to enable informed access decisions, starting with Context Visibility as the central tool for viewing endpoint data. It details certainty factor[nd]based profiling, covering the full range of probes[md]NetFlow, DHCP, HTTP, RADIUS, Device Sensor, SNMP, Active Directory, and pxGrid, along with profiling policies, logical Profiles, and how profiling enhances authorization policies. The chapter also introduces multi-factor classification (MFC) and Cloud MFC profiling, newer mechanisms that use simplified If-Then conditions and cloud telemetry to assign device labels. It closes by walking through a real profiling flow example and additional profiler settings, reinforcing that device visibility is the foundation for effective network security.

EK: I will reiterate – the correct chapter name is "Device Identification for access management"

Chapter 9, "Managing Guest Users and Personal Devices": This

Technet24

chapter focuses on the theory and practice of managing guest users and personal devices with Cisco Identity Services Engine. You are introduced to guest access and BYOD, and you learn how to implement guest access flows, including sponsored, employee, self-registration, and hotspot access. You explore approaches to web redirection and policy design, and examine the role of attribute stitching in authorization and policy enforcement. You review the configuration of guest access components, such as sponsor users, sponsor groups, guest portals, and guest types. You also study BYOD onboarding, client provisioning, and certificate management using both single-SSID and dual-SSID flows. Finally, you examine portal configuration for guest and device self-service management.

Chapter 10, "Implementing Endpoint Posture Services": This chapter focuses on the theory and practice of endpoint posture assessment with Cisco Identity Services Engine, one of the critical elements of the Zero Trust Network Access (ZTNA) approach. You learn why every endpoint must be verified before network access is granted, and how posture requirements[md]such as the presence and currency of anti-malware software or the installation of the latest patches[md]are derived from an organization's security policy. You examine the two approaches ISE offers for delivering posture services: ISE Posture, in which the Policy Service Node acts as the posture runtime and communicates directly with the posture agent on the endpoint, and ISE MDM services, in which compliance verification is delegated to an external mobile device management (MDM) or unified endpoint management (UEM) server and ISE consumes only the result. Finally, you compare the two approaches to understand when each is the better fit for a given deployment.

Chapter 11, "Managing Network Access in the Meraki Environment": This chapter focuses on network access control (NAC)

in Cisco Meraki cloud-managed networks, introducing the Meraki dashboard as the central platform for management, configuration, troubleshooting, and policy enforcement. You examine Meraki architecture and core operational elements and then learn how Meraki switches, access points, and security appliances integrate with Cisco Identity Services Engine through RADIUS to deliver wired and wireless NAC methods such as 802.1X, MAB, and iPSK. Finally, you review key implementation, monitoring, and troubleshooting tools[md]including policy controls, event visibility, CLI access on Catalyst/Smart switches, and Intelligent Capture[md]to validate authentication outcomes and resolve NAC issues effectively.

Chapter 12, "Meraki Adaptive Policy": This chapter focuses on Meraki Adaptive Policy and its role in delivering scalable, intent-based microsegmentation across cloud-managed networks. You trace the evolution from traditional Layer 3/4 ACLs and identity-based group policies to security group tag[nd]driven enforcement with security group ACLs (SGACLs), and you learn how this cloud-managed implementation of Cisco TrustSec is integrated and synchronized with Cisco Identity Services Engine through the ISE Meraki Connector. You also review platform requirements, policy design and enforcement workflows, SGT assignment behavior across network elements, and SD-WAN considerations for preserving segmentation across Auto VPN. Finally, you examine key monitoring and troubleshooting tools to validate policy operation and maintain consistent security outcomes across the Meraki infrastructure.

Chapter 13, "Remote Access Management for On-Prem and SSE": This chapter focuses on remote access evolution and its integration with network access control, from early VPN models to modern cloud-delivered architectures such as SASE and SSE. You examine how SASE combines SD-WAN and security services, while SSE delivers core

Technet24

security capabilities such as SWG, CASB, and ZTNA without the SD-WAN fabric. You then learn the key configuration elements for Cisco Secure Firewall remote access VPN, including certificate enrollment, NAT exemption handling, access policy adjustments, and authentication methods using certificates, RADIUS, and SAML, along with ISE handling of authorize-only flows. Finally, you review how Cisco SSE integrates with ISE through directory synchronization, IP pool design, and consistent policy application to extend NAC into cloud-managed environments.

Chapter 14, "Managing Network Access on Campus with Microsegmentation": This chapter explores the implementation of microsegmentation within a campus network using Cisco Identity Services Engine and security group tags to deliver granular, identity-based access control. It contrasts microsegmentation with macrosegmentation, advocating for destination-based policy enforcement to ensure scalability and operational efficiency. The text details methods for both dynamic and static SGT assignment, introduces specialized classifications, and discusses advanced security techniques. Furthermore, it outlines practical use cases for east-west, data center, and cloud traffic, while emphasizing modern operational improvements to maintain visibility and compliance across complex environments.

Chapter 15, "Extending MicroSegmentation to WAN and SD-WAN": This chapter explains how to extend microsegmentation and consistent security policies from local networks across wide area networks. It details the integration between SD-WAN Manager and Cisco ISE, showing how pxGrid and REST APIs are used to exchange security group tags and user information. By creating identity lists and unified security policies, organizations can enforce identity-based access control across their entire infrastructure. The chapter also

covers how the Overlay Management Protocol is used and how the Security Group Tag Exchange Protocol functions in this environment, ensuring robust and scalable security for modern multicloud and SASE environments.

Chapter 16, "Building Policy-interconnect with Data Center": This chapter details the integration of Cisco Identity Services Engine with Cisco Application Centric Infrastructure (ACI) to enable consistent, group-based policy management across campus and data center domains. It outlines the prerequisites and procedural steps for establishing workload connections between ISE and the Application Policy Infrastructure Controller (APIC), emphasizing the use of pxGrid for real-time, scalable context exchange. It addresses operational considerations, maintaining high availability and managing workload connections to ensure robust security policy enforcement across diverse network environments.

Chapter 17, "Classifying Cloud Workloads": This chapter introduces a dynamic methodology for securing hybrid cloud workloads in Cisco ISE by leveraging cloud-native tags to automate the classification of workloads using security group tags. It details the configuration of workload connectors for platforms such as AWS, Azure, GCP, and vCenter, explaining how administrators can map learned attributes to both primary and secondary SGTs to facilitate consistent, least-privilege access control. Furthermore, the chapter covers the use of inbound SGT domain rules to segregate binding propagation and describes how these automated bindings are shared with network devices via SXP and pxGrid, ultimately reducing the manual effort required to maintain security policies as workloads scale across cloud and on-prem environments.

Chapter 18, "ISE APIs and Advanced Automation Use Cases": This chapter explores how Cisco ISE APIs and automation tools simplify

Technet24

integration, configuration, and deployment at scale. It introduces the API Gateway and examines OpenAPI, External RESTful Services (ERS), and Monitoring and Troubleshooting (MnT) APIs, including documentation, authentication, filtering, pagination, session queries, and Change of Authorization operations. The chapter also demonstrates how Swagger UI, API clients, cURL, and Python can be used to test and execute API calls. Finally, it explains how Terraform automates ISE infrastructure deployment across cloud and virtual environments, while Ansible provides repeatable, idempotent configuration management through playbooks

Chapter 19, "Integrating Your Security Products to Enhance Context Using pxGrid": The chapter discusses the Cisco Platform Exchange Grid (pxGrid) as a vital framework for breaking down security silos by enabling seamless, real-time data exchange across diverse security infrastructure. It explains how pxGrid 2.0 uses REST and WebSocket technologies to provide the deep contextual awareness necessary for moving from reactive security to proactive orchestration. Beyond detailing the core architecture and its implementation within the Cisco ISE ecosystem, the chapter explores practical integrations with tools like Cisco Catalyst Center and Cisco Secure Firewall. Finally, it introduces pxGrid Direct, a powerful solution that extends these contextual capabilities to legacy or third-party products, ensuring that even non-native systems can contribute to a unified and intelligent security environment.

Chapter 20, "Troubleshooting, Monitoring, Alerting, and System 360": This chapter presents a structured methodology for troubleshooting ISE, walking through the process of defining the problem, understanding expected behavior, collecting data, and building a theory using a real Guest Authentication failure as a running example. It covers ISE's core monitoring and troubleshooting tools,

including live logs, live sessions, and detailed authentication reports, along with diagnostic tools such as TCP Dump, the Debug Wizard, and Support Bundle collection. The chapter also introduces alarms and reports for proactive monitoring, and concludes with System 360 Monitoring and Log Analytics, which provide dashboards for tracking node performance and analyzing RADIUS, TACACS, and profiler activity.

Chapter 21, "ISE Maintenance": This chapter covers the essential ongoing maintenance tasks required to keep an ISE deployment reliable, secure, and up to date. It details Repository management, covering their use in backups, upgrades, patches, and support bundles, followed by configuration and operational data backup and restore procedures. The chapter provides an in-depth walkthrough of ISE upgrades, including upgrade path considerations, preparation steps like the Upgrade Readiness Tool (URT), and the GUI, CLI, and backup-and-restore upgrade methods. It closes with Patch Installation guidance.

Technet24

Chapter 1. Introduction to Modern Access Control

Network Evolution

History tells us enterprise networking is a story of shifting boundaries.

For decades, the architecture of the corporate network was defined by physical boundaries and localized control. In this traditional model, the network was segmented into three primary domains: the data center, the campus, and the branch. Each of these segments functioned as a distinct entity, governed by static policies and anchored by a central hub. The data center served as the undisputed heart of the organization, housing critical applications and data stores within a climate-controlled fortress. The campus represented the high-density environment where employees gathered, connected through wired Ethernet or early wireless standards to access the resources residing in the data center. The branch acted as the peripheral extension, connecting remote offices back to the central hub, often through expensive, dedicated private circuits. This "castle-and-moat" approach relied on the assumption that the network perimeter was a fixed boundary and that security could be achieved by controlling the entry and exit points of that perimeter.

As the digital landscape matured, the rigidity of this model began to encounter significant friction. The proliferation of mobile devices, the

mandate for remote work, and the rapid adoption of Software-as-a-Service (SaaS) applications fundamentally altered the traffic patterns that the traditional architecture was designed to support. The focus shifted from connecting a user to a building to connecting a user to an application, regardless of where that application resided. As workloads migrated from on-prem servers to public and private clouds, the traffic that once flowed neatly from the branch to the data center began to experience hairpinning, creating inefficiencies and latency that hindered the user experience and strained the underlying infrastructure.

The transition to the modern hybrid model represents a departure from location-centric connectivity toward an identity-centric and application-centric paradigm. In this evolving environment, the network is no longer a static collection of segments but a dynamic fabric that extends to wherever the user and the workload happen to be. The traditional data center has been augmented by cloud infrastructure, creating a distributed environment where applications reside in multiple locations. Equally, the workforce has become increasingly dispersed, accessing corporate resources from home offices, coffee shops, serviced office spaces, and international locations. This shift has necessitated the move toward architectures like software-defined wide area networking (SD-WAN) and secure access service edge (SASE), and products such as Cisco Identity Services Engine (ISE), which decouple the network from the underlying hardware and provide consistent security and performance policies across the entire ecosystem.

In this hybrid world, the network must be capable of identifying the user, the device, and the application context to make real-time decisions about connectivity and security. The concept of a trusted internal network versus an untrusted external network has been largely

Technet24

replaced by the Zero Trust model, where every connection request is verified, regardless of its origin. Workloads are now orchestrated across hybrid cloud environments, requiring the network to provide seamless connectivity and consistent visibility. This evolution is not merely a change in technology but a fundamental reimagining of the network's role in the enterprise. It is no longer about maintaining the integrity of a physical perimeter; it is about ensuring the security and agility of the digital path between the user and the data, wherever they may reside.

The future of networking lies in this convergence of connectivity and security, where the network fabric intelligently adapts to the needs of the business. By moving away from the constraints of traditional segments, organizations can achieve the flexibility required to support a global, mobile workforce and a cloud-first application strategy. The evolution from well-defined segments to a fluid, hybrid model is the necessary response to a world where the office is no longer a place but a capability that follows the user. As we continue to integrate these advanced architectures, the focus remains on building a resilient, scalable, and secure foundation that empowers the enterprise to innovate at the speed of the cloud.

Network Access Control Response to the Changes

In the preceding section, we explored how the network has moved away from the rigid, perimeter-based "castle-and-moat" architecture toward a fluid, hybrid model. As the office ceases to be a fixed location and becomes a capability that follows the user, the way we handle network access must undergo a similar transformation. Traditional network access control, or NAC, was once primarily concerned with the physical port. It operated on the simple premise that if you were

plugged into a wall jack in the office, you were likely authorized to be there. This legacy approach, while effective for a time, simply cannot keep pace with a world where users connect from coffee shops, home offices, and various cloud environments.

The primary shift in NAC is moving from a gatekeeper of physical ports to an orchestrator of digital identity. In this modern landscape, the network needs to know far more than just the physical location of a device. It must understand the identity of the user, the health of the device they are using, the sensitivity of the application they are trying to reach, and the context of their current environment. This is where the evolution of tools like Cisco Identity Services Engine becomes critical. Instead of just asking if a device is allowed on a specific switch, modern NAC asks whether the user is who they say they are, whether their device is patched and secure, and whether that specific user should have access to that specific application at this particular moment in time.

This transition is the backbone of the Zero Trust model. Because the network perimeter has essentially dissolved, we can no longer rely on the assumption that anything inside the network is safe. Every connection request must be verified, regardless of where it originates. NAC has had to adopt this mindset by becoming a continuous verification engine rather than a one-time check at the point of entry. It now functions as the policy decision point that informs the rest of the network fabric, ensuring that security policies are consistent whether the user is sitting at their desk or working from a remote site.

Furthermore, NAC must now integrate seamlessly with the broader technology stack, including SD-WAN and SASE architectures. By decoupling access control from the underlying hardware, we gain the agility to apply consistent policies across a distributed environment. When a user moves from a local office to a remote location, the

Technet24

network should not require a new set of credentials or a different security posture. Instead, the NAC system should recognize the user's identity and maintain their access rights dynamically, ensuring that the digital path between the user and their data remains both secure and efficient.

Ultimately, the goal of modern NAC is to provide a seamless experience that does not sacrifice security for agility. By focusing on identity and context, we can move away from the constraints of traditional segments and embrace a model where the network intelligently adapts to the needs of the business. This evolution ensures that we are not just managing access but actively protecting the integrity of the digital path, allowing the enterprise to innovate with the speed and flexibility that the modern cloud-first world demands.

Chapter 2. ISE Deployment Principles

Introduction

Designing a Cisco Identity Services Engine deployment is one of the most crucial steps when working with ISE. You should ensure that by the time ISE is running and managing user and endpoint authentications, the ISE cluster is stable, designed according to best practices and Cisco recommendations, and set up for high availability. This chapter guides you through the most important considerations for ISE deployments and helps you design ISE for success.

ISE Node Personas

An ISE node is a single ISE instance, which can operate as a virtual machine (VM), hardware appliance, or cloud instance. ISE offers four node personas, as shown in Figure 2-1. You can assign a single persona to a node, or you can assign multiple personas to one node, depending on the deployment size you select.

Technet24

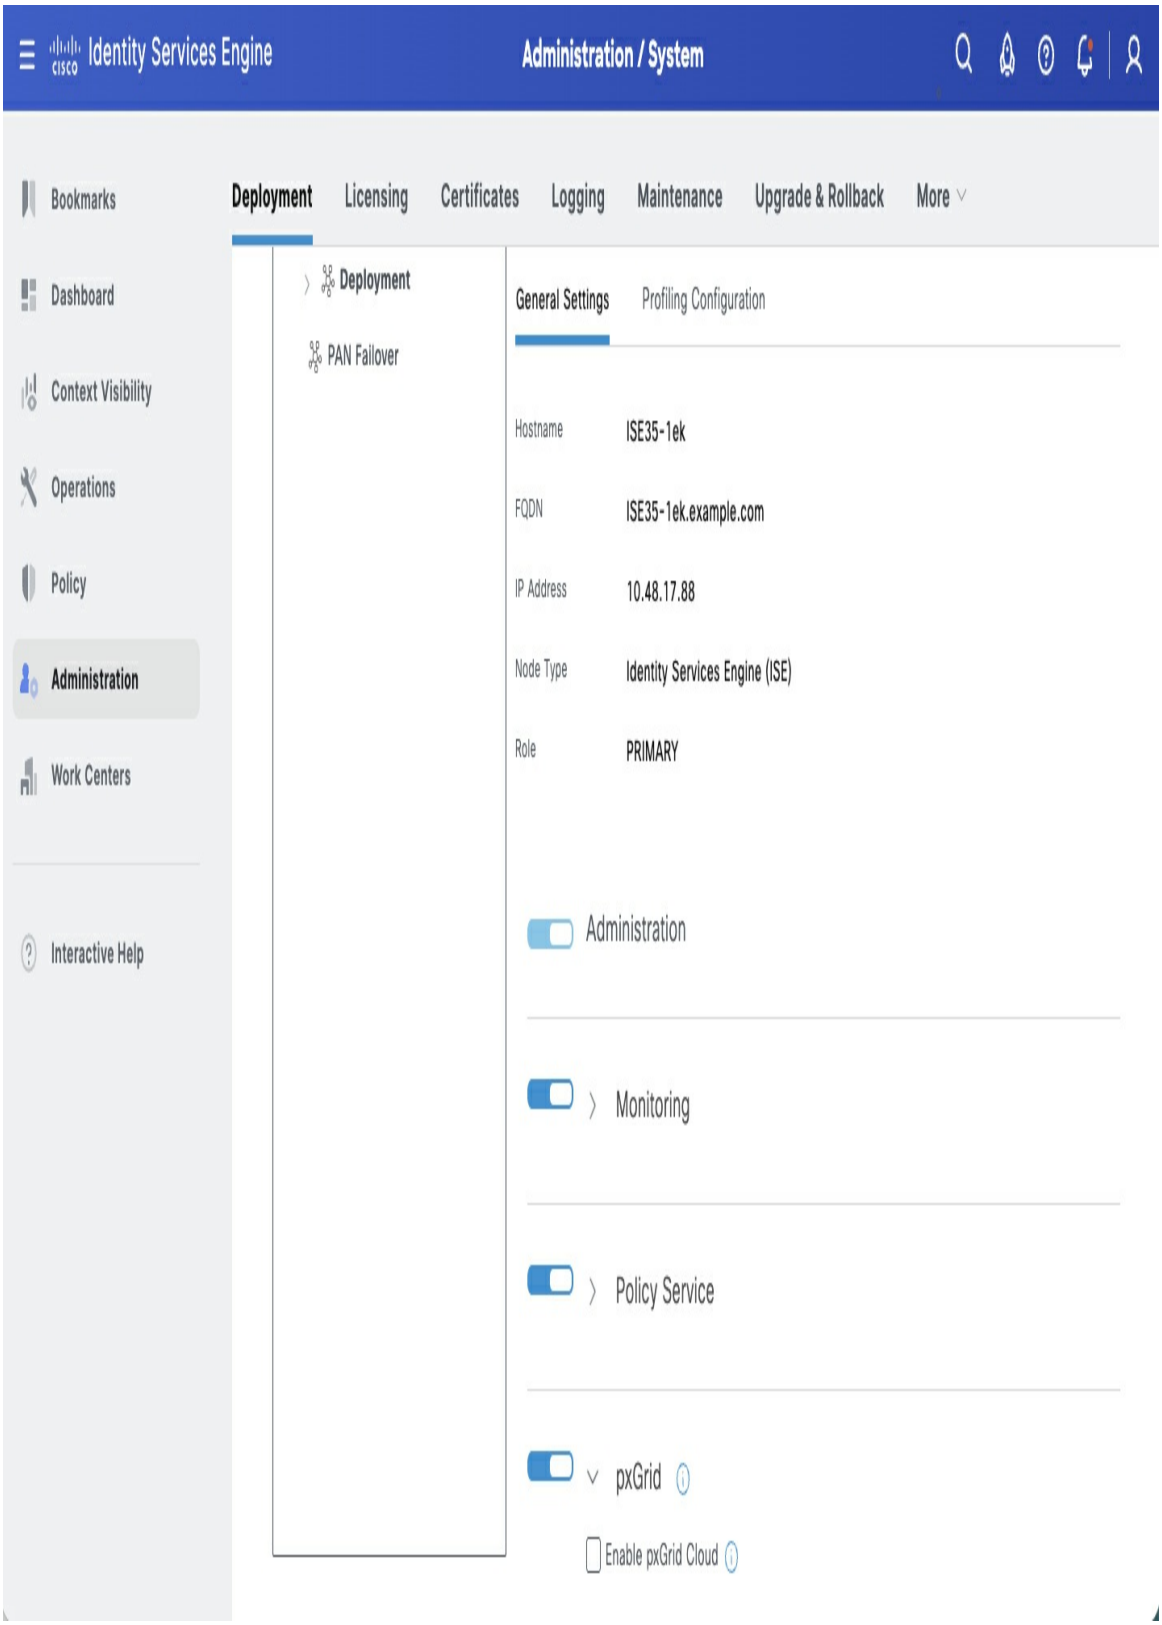


Figure 2-1 *ISE Personas in Distributed Deployment*

The Policy Administration Node, or PAN, serves as the interface for configuring ISE policies, settings, and certificates. You use it to add or remove endpoints, users, and network access devices. PAN provides a single pane of glass for ISE management. It maintains the ISE configuration database and replicates it to all nodes in the deployment. This design ensures that the ISE database remains consistent across all nodes at all times.

The Monitoring and Troubleshooting persona, or MnT, is responsible for collecting data for live logs, live sessions, and reporting. It serves as the destination for syslog messages from all ISE nodes. However, you still use the PAN to access this information and simplify management.

The Policy Service Node, or PSN, is the central node for making all policy decisions. You rely on it for authentication and authorization. It handles profiling, posture, guest authentication, TACACS+, and many other functions. The PSN provides most user and endpoint services. PSN nodes can serve multiple functions or just a single one. You can configure a PSN to handle the following:

- **RADIUS authentication with 802.1x and MAB:** PSNs act as Remote Authentication Dial-In User Service (RADIUS) servers to handle user and endpoint authentications.
- **Device administration with TACACS+:** PSNs process TACACS+ authentication, authorization, and accounting requests and make device-level access decisions.
- **Threat Centric Network Access Control (TC-NAC):** PSNs receive data from external TC-NAC solutions and make policy decisions based on the results.
- **Guest access:** PSNs host guest portals, authenticate guest users, and grant network access according to configured policies.

Technet24

- **TrustSec:** PSNs act as RADIUS servers for TrustSec-specific functions.
- **SXP:** PSNs peer with network devices over the SXP protocol to share IP to SGT mappings.
- **Passive ID:** PSNs handle passive authentication data from agents and other solutions to make policy decisions.

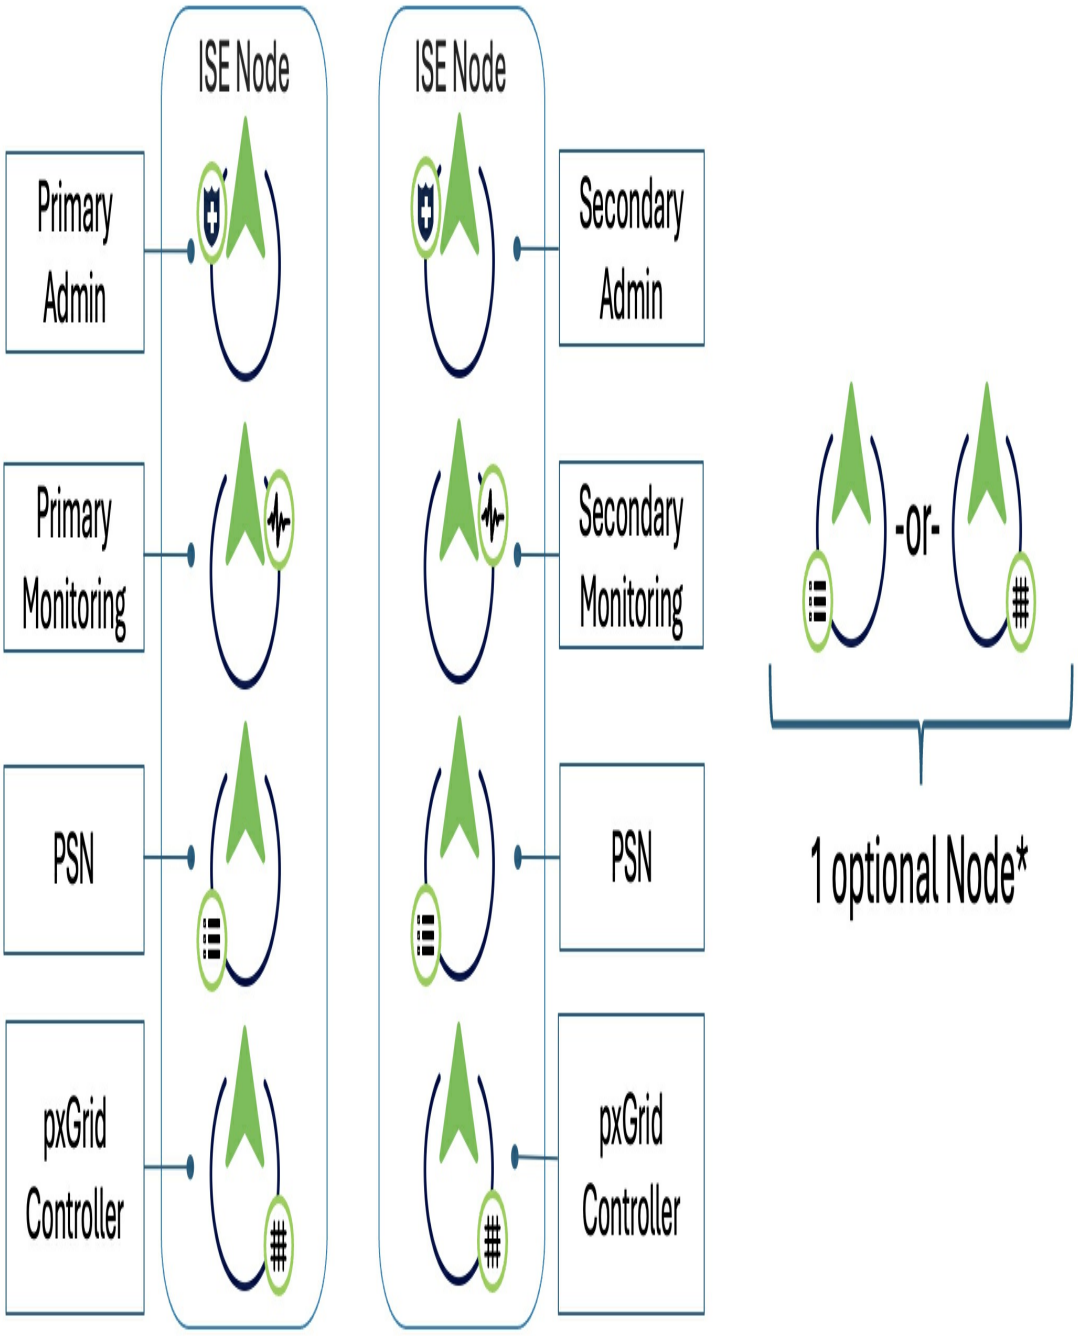
The PxGrid Controller persona facilitates context sharing between subscribers and publishers. For example, ISE can share authentication details with a firewall. This capability allows the firewall to make policy decisions based not only on the IP address of a packet but also on rich context received from ISE, such as Active Directory group, username, and endpoint profile.

ISE Deployment Models

You can choose from three deployment models when designing an ISE production cluster.

> **Note**
>
> The terms *cluster* and *deployment* are used interchangeably and refer to an ISE distributed deployment that consists of multiple nodes.

A small deployment, as shown in Figure 2-2, consists of two ISE nodes with all personas (PAN, MnT, PSN, and pxGrid) running on each node. One node acts as the primary node, and the other serves as the secondary node for redundancy. In a small deployment, you can optionally add one extra node. This additional node can function as a dedicated PSN, a pxGrid node, or a health check node.


ISE Node
ISE Node
Primary Admin
Secondary Admin
Primary Monitoring
Secondary Monitoring
PSN
PSN
pxGrid Controller
pxGrid Controller
-or-
1 optional Node*

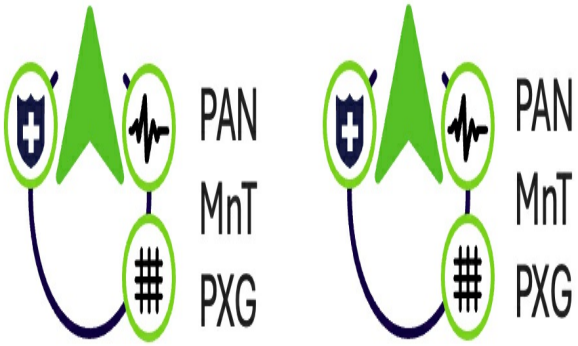
Technet24

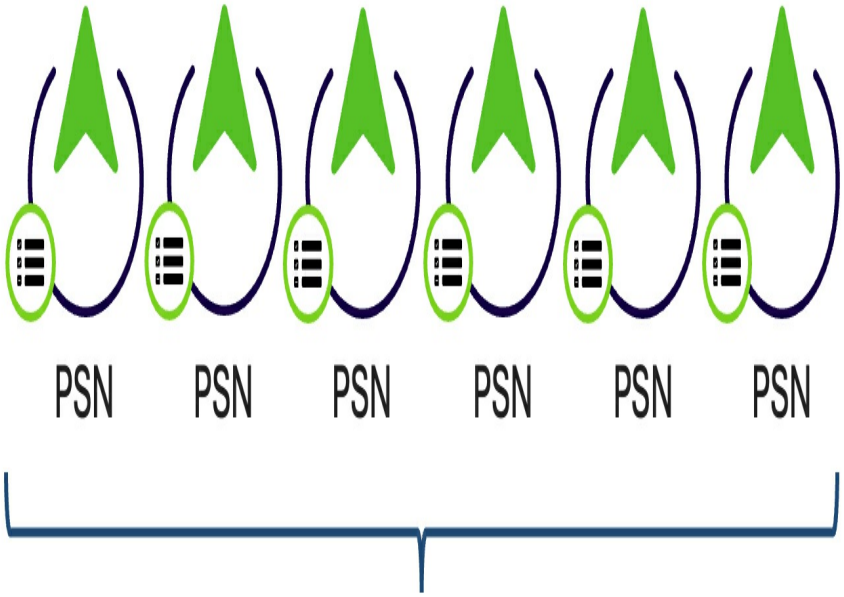
Figure 2-2 *ISE Small Deployment*

Note

Adding an additional node to a small deployment does not affect any scale numbers. You are recommended to include it for health check purposes if you choose to configure PAN high availability.

A medium deployment, as shown in Figure 2-3, consists of two nodes with PAN, MnT, and pxGrid personas assigned. One node serves as the primary, and the other acts as the secondary to provide redundancy. PSNs run on dedicated nodes, and you can have up to six PSNs in a medium deployment.





Max of 6 PSN's

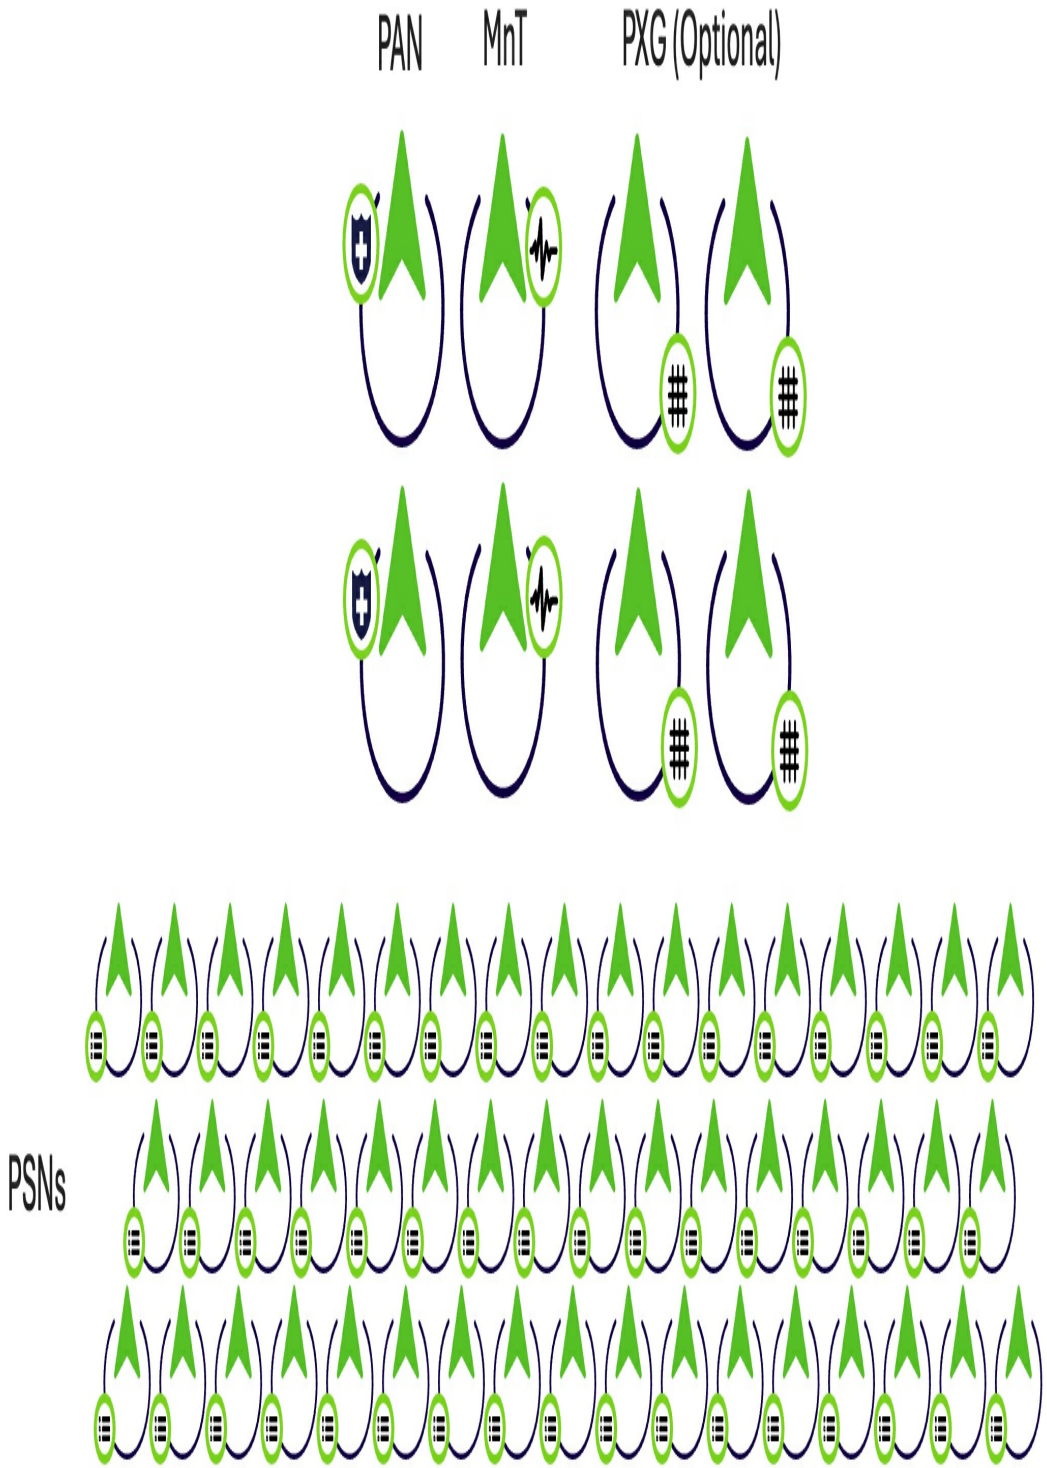
Technet24

Figure 2-3 *ISE Medium Deployment*

> **Note**
>
> You can run pxGrid on any PSN, or you can add up to two dedicated pxGrid nodes to the deployment.

The maximum supported configuration includes up to eight nodes.

A large deployment, as shown in Figure 2-4, consists of two PAN nodes (PPAN and SPAN), two MnT nodes (primary MnT and secondary MnT), up to four pxGrid nodes, and PSNs. All personas must run on dedicated nodes. This deployment supports up to 50 PSNs.



Technet24

Figure 2-4 *ISE Large Deployment*

The maximum supported configuration includes up to 58 nodes.

ISE Deployment Sizing

To choose the right deployment model, you need to consider the maximum number of active endpoints in your network. An active endpoint is an endpoint with a unique MAC address that is currently connected and authenticated on ISE. For example, if 10,000 endpoints connect at 8 a.m. and disconnect at 4 p.m., and then another 10,000 endpoints connect at 6 p.m. and disconnect at 12 a.m., the maximum number of active endpoints would still be 10,000 because these groups do not overlap.

> **Note**
>
> In practice, it is almost impossible to achieve complete separation between users working in different time zones.

After you determine the maximum number of concurrent active sessions, you can select the appropriate SNS appliance or virtual machine to serve as the PAN or MnT node.

> **Note**
>
> The scaling information presented in this book reflects the recommendations available at the time of writing, when Cisco Identity Services Engine (ISE) 3.5 patch 3 was the latest release. Because scaling recommendations may change with new releases, always refer to the latest *Performance and Scalability Guide for Cisco Identity Services Engine* to determine the appropriate deployment type and sizing for your environment.

Use Table 2-1 to make the correct choice.

Table 2-1 *Maximum Concurrent Active Sessions for Deployments with Different SNS Appliances Acting as PAN, MnT, or PAN/MnT*

Deployment	Cisco SNS 3595	Cisco SNS 3615	Cisco SNS 3715	Cisco SNS 3815	Cisco SNS 3655	Cisco SNS 3755, 3855	Cisco SNS 3695	Cisco SNS 3795, 3895
Large	500,000	Unsupported	Unsupported	Unsupported	500,000	750,000	2,000,000	2,000,000
Medium	20,000	12,500	75,000	75,000	25,000	150,000	50,000	150,000
Small	20,000	12,500	25,000	25,000	25,000	50,000	50,000	50,000

Suppose you expect up to 90,000 concurrent active sessions on ISE. You can immediately see that a small deployment will not meet this requirement, so you need to choose either a medium or large deployment. If you select a medium deployment, your PAN and MnT nodes can be any of the following: Cisco SNS 3755, 3855, 3795, or 3895.

Note

Some of the scaling values presented throughout this book are based on the Cisco SNS 36xx and 37xx appliances. For new deployments, however, Cisco recommends using the Cisco SNS 38xx appliance family because the Cisco SNS 36xx and 37xx appliance family has reached end of sale (EOS).

Note

Technet24

Always consider future expansion and potential increase in the number of concurrent sessions.

After selecting the PAN and MnT nodes for your deployment, you need to choose the appropriate PSNs. Use Table 2-2 for PSN selection.

Table 2-2 *Maximum Concurrent Active Sessions for Different ISE Appliances Acting as PSNs*

PSN Type	Cisco SNS 3595	Cisco SNS 3615	Cisco SNS 3715, 3815	Cisco SNS 3655	Cisco SNS 3755, 3855	Cisco SNS 3695	Cisco SNS 3795, 3895*
Dedicated PSN (Cisco ISE node has only PSN persona)	40,000	25,000	50,000	50,000	100,000	100,000	100,000
Shared PSN (Cisco ISE node has multiple personas)	20,000	12,500	25,000	25,000	50,000	50,000	50,000

The row with Shared PSN should be used with small deployments. The Dedicated PSN row should be used for medium and large deployments.

Because, in this example, you have chosen to proceed with a medium deployment and expect 90,000 concurrent active sessions, the following PSN models would meet these requirements: Cisco SNS 3755, 3855, 3695, 3795, and 3895. To achieve full high availability, you need at least two PSNs of the same model.

Note

Cisco SNS 3795 and Cisco SNS 3895 are equipped with more RAM and improved disk read/write performance. These models are best suited for dedicated PAN, dedicated MnT, or PAN/MnT personas and provide no added value when deployed as dedicated PSNs.

At this point, you may wonder why the tables include only hardware appliances and what options are available if you want to build your deployment with virtual machines or cloud instances. At the time of writing, Cisco ISE can be installed on VMware, KVM hypervisors, Hyper-V (Windows Server and Azure Stack HCI), Nutanix AHV, and OpenStack. It can also run as a virtual machine in AWS, Azure, and OCI.

After you have selected the hardware appliance for your personas, you can use Table 2-3 to choose the resources required for a virtual machine or cloud instance, if you decide to use those.

Table 2-3 *Specifications for Cisco ISE VM and Cloud Instances*

Technet24

Models	Cisco SNS 3615	Cisco SNS 3595	Cisco SNS 3655	Cisco SNS 3695	Cisco SNS 3715	Cisco SNS 3755	Cisco SNS 3795
VM Specification	16vCPU 32 GB	16vCPU 64 GB	24vCPU 96 GB	24vCPU 256 GB	24vCPU 32 GB	40vCPU 96 GB	40vCPU 256 GB
AWS	c5.4xlarge	m5.4xlarge	c5.9xlarge* m5.8xlarge	m5.16xlarge	c5.9xlarge* m5.8xlarge	—	m5.16xlarge
	c7i.4xlarge	—	m7i.8xlarge***	m7i.16xlarge***	m7i.8xlarge***	—	m7i.16xlarge***
Azure	Standard_F16s_v2*	Standard_D16s_v4	Standard_F32s_v2* Standard_D32s_v4	Standard_D64s_v4	Standard_F32s_v2* Standard_D32s_v4	—	Standard_D64s_v4
OCI	Optimized3.Flex* (8 OCPU and 32 GB)	Standard3.Flex (8 OCPU and 64 GB)	Optimized3.Flex (16 OCPU and 64 GB)* Standard3.Flex (16 OCPU and 128 GB)	Standard3.Flex (16 OCPU and 256 GB)	Optimized3.Flex (16 OCPU and 64 GB)* Standard3.Flex (16 OCPU and 128 GB)	—	Standard3.Flex (32 OCPU and 256 GB)

Authentication rates play a significant role as well. It is not enough just to choose the maximum number of concurrent sessions; transactions

per second (TPS) also should be part of the deployment choice decision. Authentication type and external database are two key factors that influence TPS on the PSNs. Table 2-4 lists authentication rates for RADIUS when PSN is a dedicated node in a deployment.

Table 2-4 *Authentication Rate for RADIUS When PSN Is a Dedicated Node in a Deployment*

Authentication Method	Cisco SNS 3615, 3715, 3815	Cisco SNS 3595	Cisco SNS 3655, 3755, 3695, 3795, 3855, 3895
PAP with internal user database	900	1100	1300
PAP with Active Directory/Open Database Connectivity (ODBC)	250	250	300
PAP with LDAP directory	300	300	350
PEAP (MSCHAPv2) with internal user database	150	150	200
PEAP (MSCHAPv2) with Active Directory/ODBC	150	150	175
PEAP (GTC) with internal user database	150	150	250
PEAP (GTC) with Active	100	150	175

Technet24

Directory/ODBC			
EAP-FAST (MSCHAPv2) with internal user database	350	400	500
EAP-FAST (MSCHAPv2) with Active Directory/ODBC	200	250	300
EAP-FAST (GTC) with internal user database	350	400	450
EAP-FAST (GTC) with Active Directory/ODBC	200	200	300
EAP-FAST (GTC) with LDAP directory	200	300	300
EAP-TLS with internal user database	150	150	200
EAP-TLS with Active Directory/ODBC	150	150	200
EAP-TLS with LDAP directory	150	200	200
EAP TEAP with internal user database	100	100	200
MAB with internal user database	500	900	1000
MAB with LDAP directory	400	500	600
EAP-TTLS PAP with Microsoft Entra ID	30	30	50
EAP-TLS with Microsoft Entra ID (Cisco ISE release 3.5 patch 2 and earlier)	30	30	50
EAP-TLS with Microsoft Entra ID (Cisco ISE release 3.5 patch 3 onwards)	40	30	100

You can build an ISE deployment using nodes of the same form factor, such as hardware appliances only, virtual machines only, or cloud instances only. Alternatively, you can combine nodes of different form factors, including hardware appliances, virtual machines, and cloud instances. You may mix and match all form factors as long as there is end-to-end connectivity between ISE nodes and they run the same major version and patch.

ISE VM Disk Provisioning Guidance

The minimum disk storage requirement for an ISE node is 300 GB for the PSN-only persona. Table 2-5 shows the minimum, maximum, and recommended disk space based on the ISE persona.

Table 2-5 *ISE Disk Space Allocation Recommendations*

Cisco ISE Persona	Minimum Disk Space	Recommended Disk Space	Maximum Disk Space
Standalone ISE	600 GB	600 GB to 2.4 TB	2.4 TB
PAN only	600 GB	600 GB	2.4 TB
MnT only	600 GB	600 GB to 2.4 TB	2.4 TB
PSN only	300 GB	300 GB	2.4 TB
PxGrid only	300 GB	300 GB	2.4 TB
PAN + MnT (+ PxGrid)	600 GB	600 GB to 2.4 TB	2.4 TB
PAN + MnT + PSN (+ PxGrid)	600 GB	600 GB to 2.4 TB	2.4 TB

Disk allocation may vary based on logging requirements. For extra log storage, you can increase the VM disk space. If you increase the disk space after installation, you must reimage the node to ensure proper

Technet24

disk allocation.

Hypervisor snapshots are not supported with ISE because there is no way to pause the database prior to creating a snapshot. You should use the ISE Backup and Restore functionality instead.

VMotion (hot migration) is supported. You do not need to power off ISE during hot migration, and ISE can be migrated without any interruption in availability.

Cold migration requires ISE to be in a shutdown state and powered off.

ISE Installation

Now that you have determined the number of nodes and the specifications required for your ISE deployment, you can focus on the ISE installation. There are multiple ways to install ISE.

Hardware Appliance

The installation process depends on the form factor you have chosen. For example, if you are installing ISE on a hardware appliance, the installation process would require you to follow these steps:

Step 1. Connect to the Cisco Integrated Management Controller (CIMC) for server management.

Step 2. Download the ISE software image of the desired version from https://software.cisco.com/download/home/283801620/typ

Step 3. From CIMC, launch the KVM console.

a. Select **Virtual Media > Activate Virtual Devices**.

b. Select **Virtual Media > Map CD/DVD**, select the ISO downloaded in Step 2, and click **Map Device**.

c. Select **Macros > Static Macros > Ctrl+Alt+Del** to boot the appliance with the ISO image you selected.

d. Press **F6** to open a boot menu. The screen in Figure 2-5 should appear.

e. Select **KVM-Mapped DVD** and start installation.

Technet24

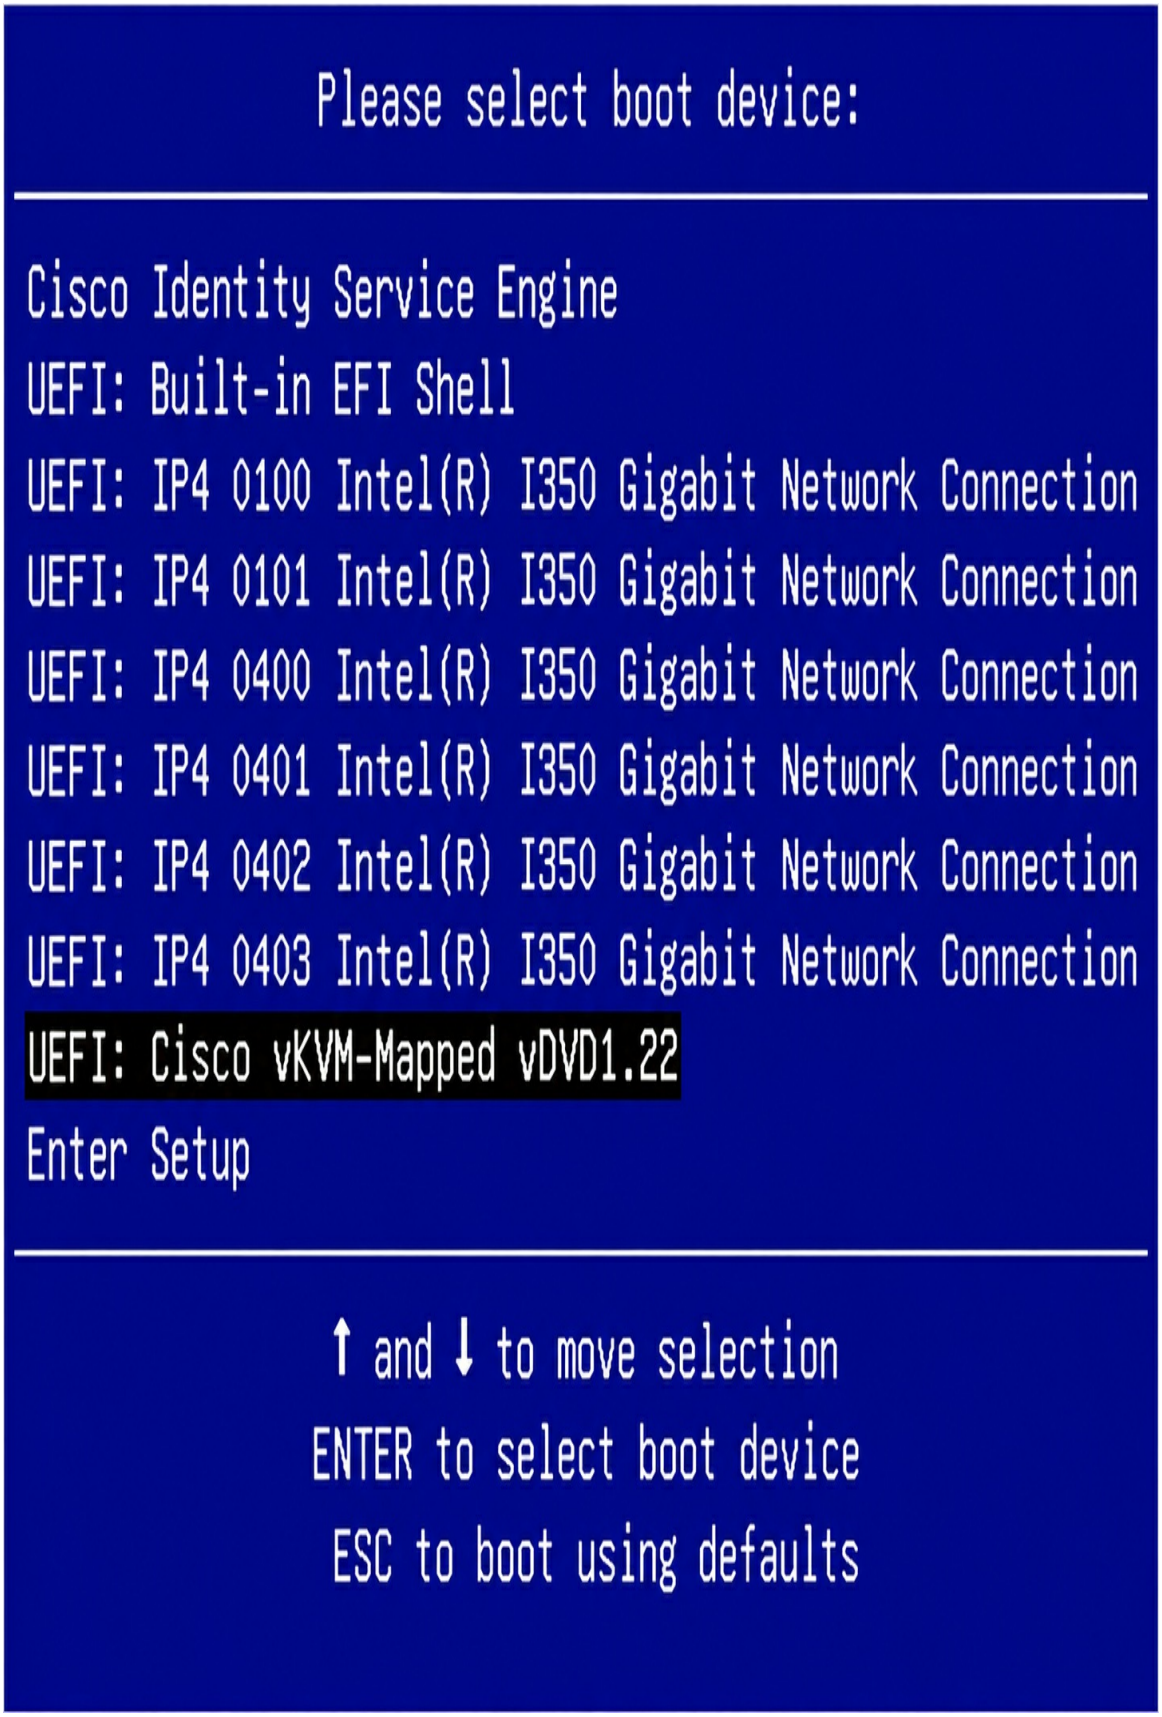

Please select boot device:
Cisco Identity Service Engine
UEFI: Built-in EFI Shell
UEFI: IP4 0100 Intel(R) I350 Gigabit Network Connection
UEFI: IP4 0101 Intel(R) I350 Gigabit Network Connection
UEFI: IP4 0400 Intel(R) I350 Gigabit Network Connection
UEFI: IP4 0401 Intel(R) I350 Gigabit Network Connection
UEFI: IP4 0402 Intel(R) I350 Gigabit Network Connection
UEFI: IP4 0403 Intel(R) I350 Gigabit Network Connection
UEFI: Cisco vKVM-Mapped vDVD1.22
Enter Setup
↑ and ↓ to move selection
ENTER to select boot device
ESC to boot using defaults

Figure 2-5 *ISE KVM Boot Device Menu*

Important:

Remote installation without physical access, using CIMC, may take several hours. It is recommended to copy the ISO file to a USB drive and use the USB drive during installation instead. Another option is to use the machine or jump host, which are logically close to the server. It is not recommended to install ISE via KVM over the Internet or WAN links.

Virtual Machine

For virtual machine installation, the process is similar. You need to boot the VM with the ISO image.

Best Practice:

Although the following example demonstrates Cisco ISE installation using the ISO image, Cisco recommends using OVA files whenever possible because they provide preconfigured settings and simplify deployment.

The screen in the Figure 2-6 will appear.

Technet24

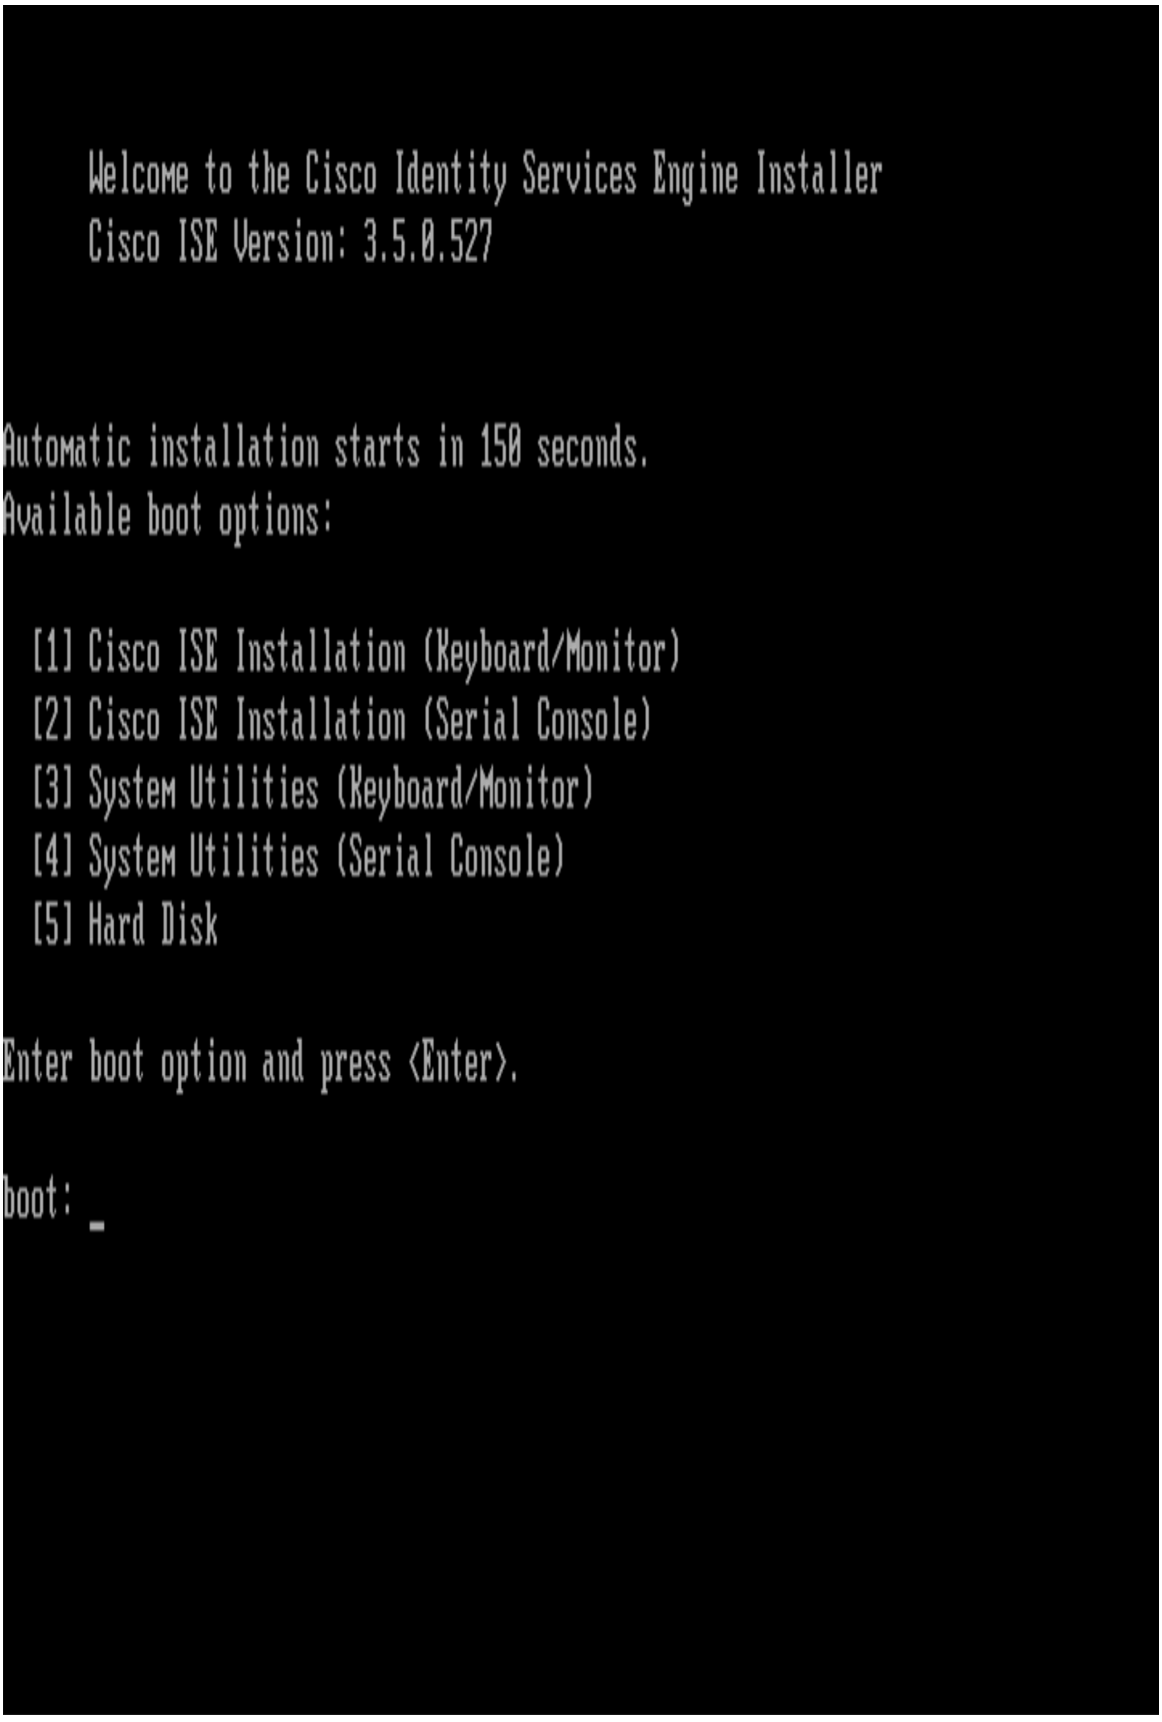

Welcome to the Cisco Identity Services Engine Installer
Cisco ISE Version: 3.5.0.527
Automatic installation starts in 150 seconds.
Available boot options:
[1] Cisco ISE Installation (Keyboard/Monitor)
[2] Cisco ISE Installation (Serial Console)
[3] System Utilities (Keyboard/Monitor)
[4] System Utilities (Serial Console)
[5] Hard Disk
Enter boot option and press <Enter>.
boot:

Figure 2-6 *ISE VM Boot Menu*

To proceed with ISE installation, select option 1: **Cisco ISE Installation (Keyboard/Monitor).**

Best Practice:

Always reserve CPU and memory resources to match the required allocation. Insufficient resource reservation can significantly affect ISE performance and stability.

OpenStack

Starting with ISE 3.4 patch 4, OpenStack ISE installation is also supported. The following steps use the OpenStack command-line interface (CLI) method for installation.

Step 1. Create a custom flavor to match the ISE appliance sizing you selected during the design phase. Here is a sample command for Cisco SNS 3815 node, with the name sns3815-openstack with 64 GB RAM, a 300 GB disk, and 32 virtual CPUs:

```
openstack flavor create sns3815-openstack --id auto --ram
32
```

Step 2. Create a glance image for ISE installation.

a. Create a blank Cinder volume for the VM's main hard drive:

```
openstack volume create --size
```

b. Create a temporary VM to copy the ISE filesystem onto the blank volume using this command:

```
openstack server create --image  --volume  --
```

Technet24

```
flavor  --network
```

Attach both the blank Cinder volume and installation ISO to the VM.

c. Install the OS through the VM console.

Access the VM console using this command:

```
openstack console url show
```

When the boot menu appears, select **[1] Cisco ISE Installation (Keyboard/Monitor).**

After installation completes, the console returns to the boot prompt. Now the volume contains bootable OS.

d. Configure the volume as bootable.

Delete the temporary VM to release the volume using this command:

```
openstack server delete
```

Verify that the volume status is "available" with this command:

```
watch openstack volume show
```

Mark the volume as bootable using this command:

```
openstack volume set --bootable
```

Follow these steps to create a QCOW2 image, install ISE using ISO, and upload the image to OpenStack.

a. Create QCOW2 image with this command:

```
qemu-img create -f qcow2 .qcow2
```

b. Install ISE ISO on the QCOW image. Run this command to boot the ISO and begin the installation on the disk

image:

```
/usr/libexec/qemu-kvm -enable-kvm -m  -smp  -cpu host \ -
drive file=.qcow2,format=qcow2 \ -cdrom  \ -boot d -net
nic,model=virtio -net user \ -nographic -serial mon:stdio
```

c. Perform the installation via a serial console. When the installation menu appears, select **2** to proceed with the installation using the serial console. Follow the on-screen prompts to complete the setup.

d. Upload the QCOW2 image to OpenStack. After the installation is complete and the image is prepared, use the OpenStack CLI to create the new image in your environment.

```
openstack image create --disk-format qcow2 --container-fo
 --private
```

Step 3. Create and launch the Cisco ISE VM with the prepared bootable volume:

```
openstack server create --volume  --flavor  --
network
```

Step 4. Configure network settings for the VM.

Access the VM console and enter the **setup** command.

ISE Setup

When the ISE installation is complete, the setup screen appears as shown in Figure 2-7.

Technet24

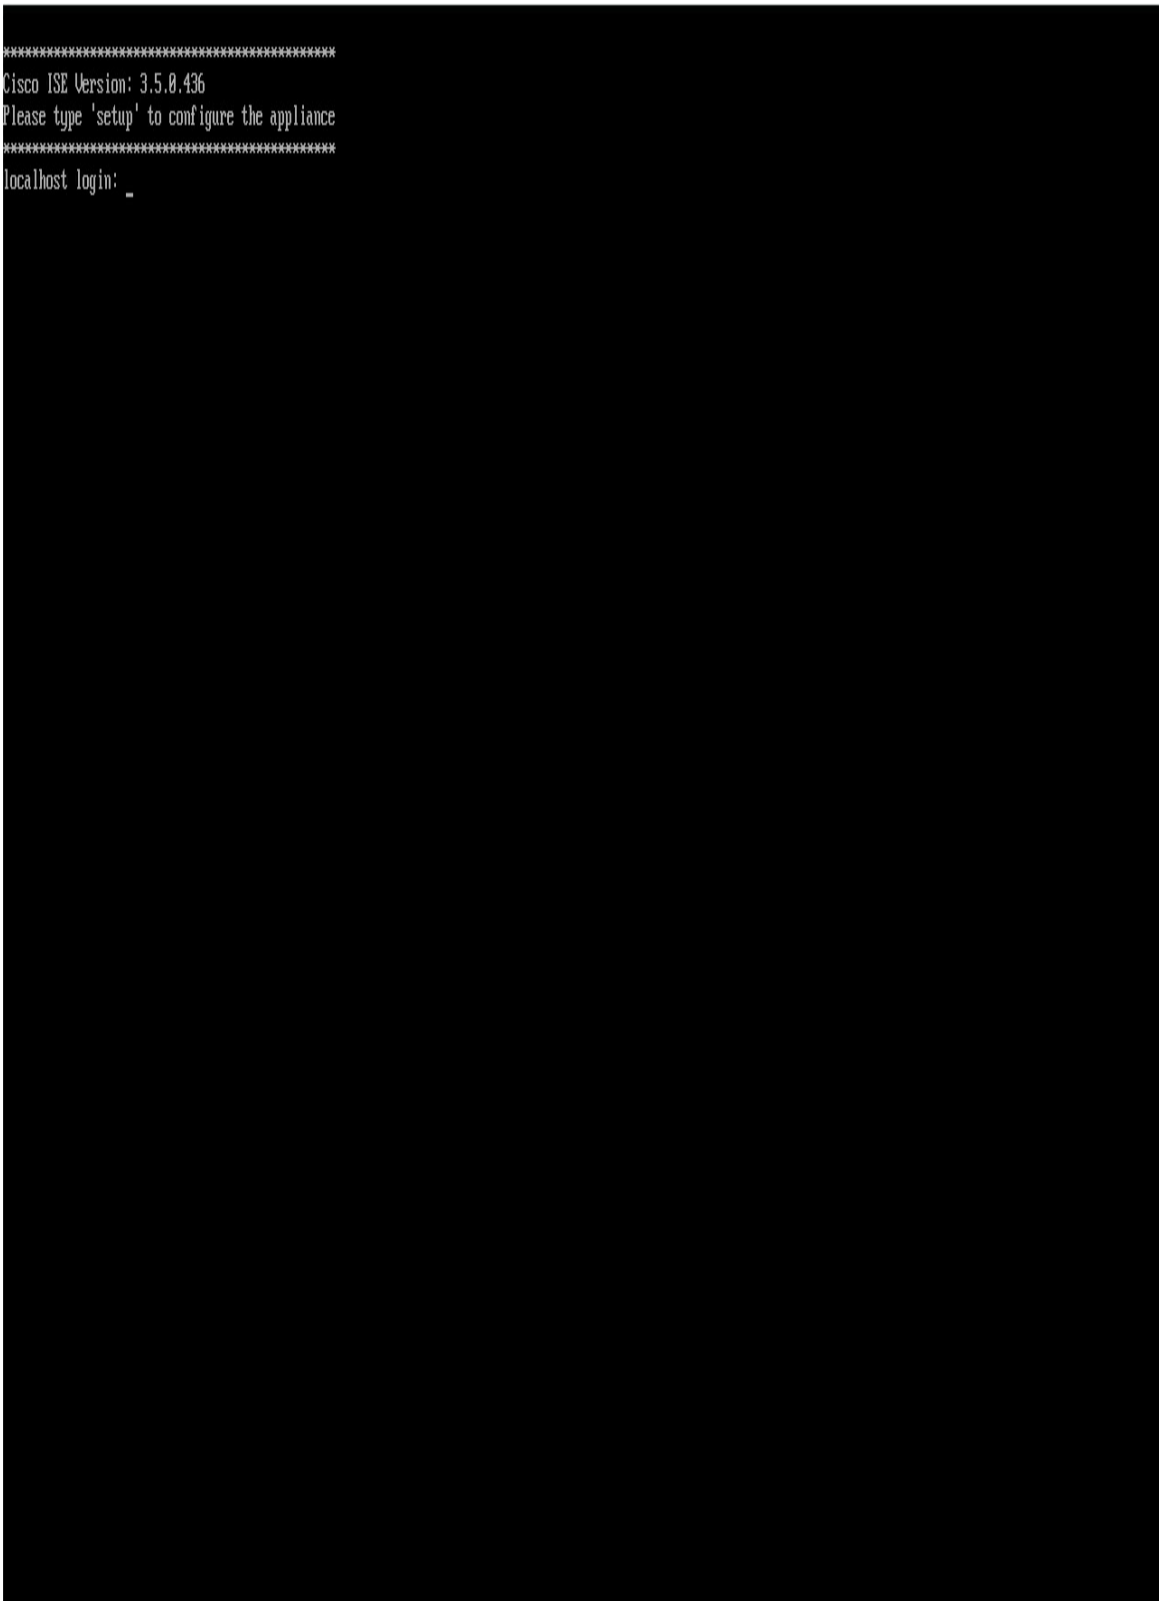


Figure 2-7 *ISE Setup Menu*

At this point, you need to type **setup** to configure your ISE node. ISE prompts you to configure several settings, as shown in Figure 2-8.

Technet24

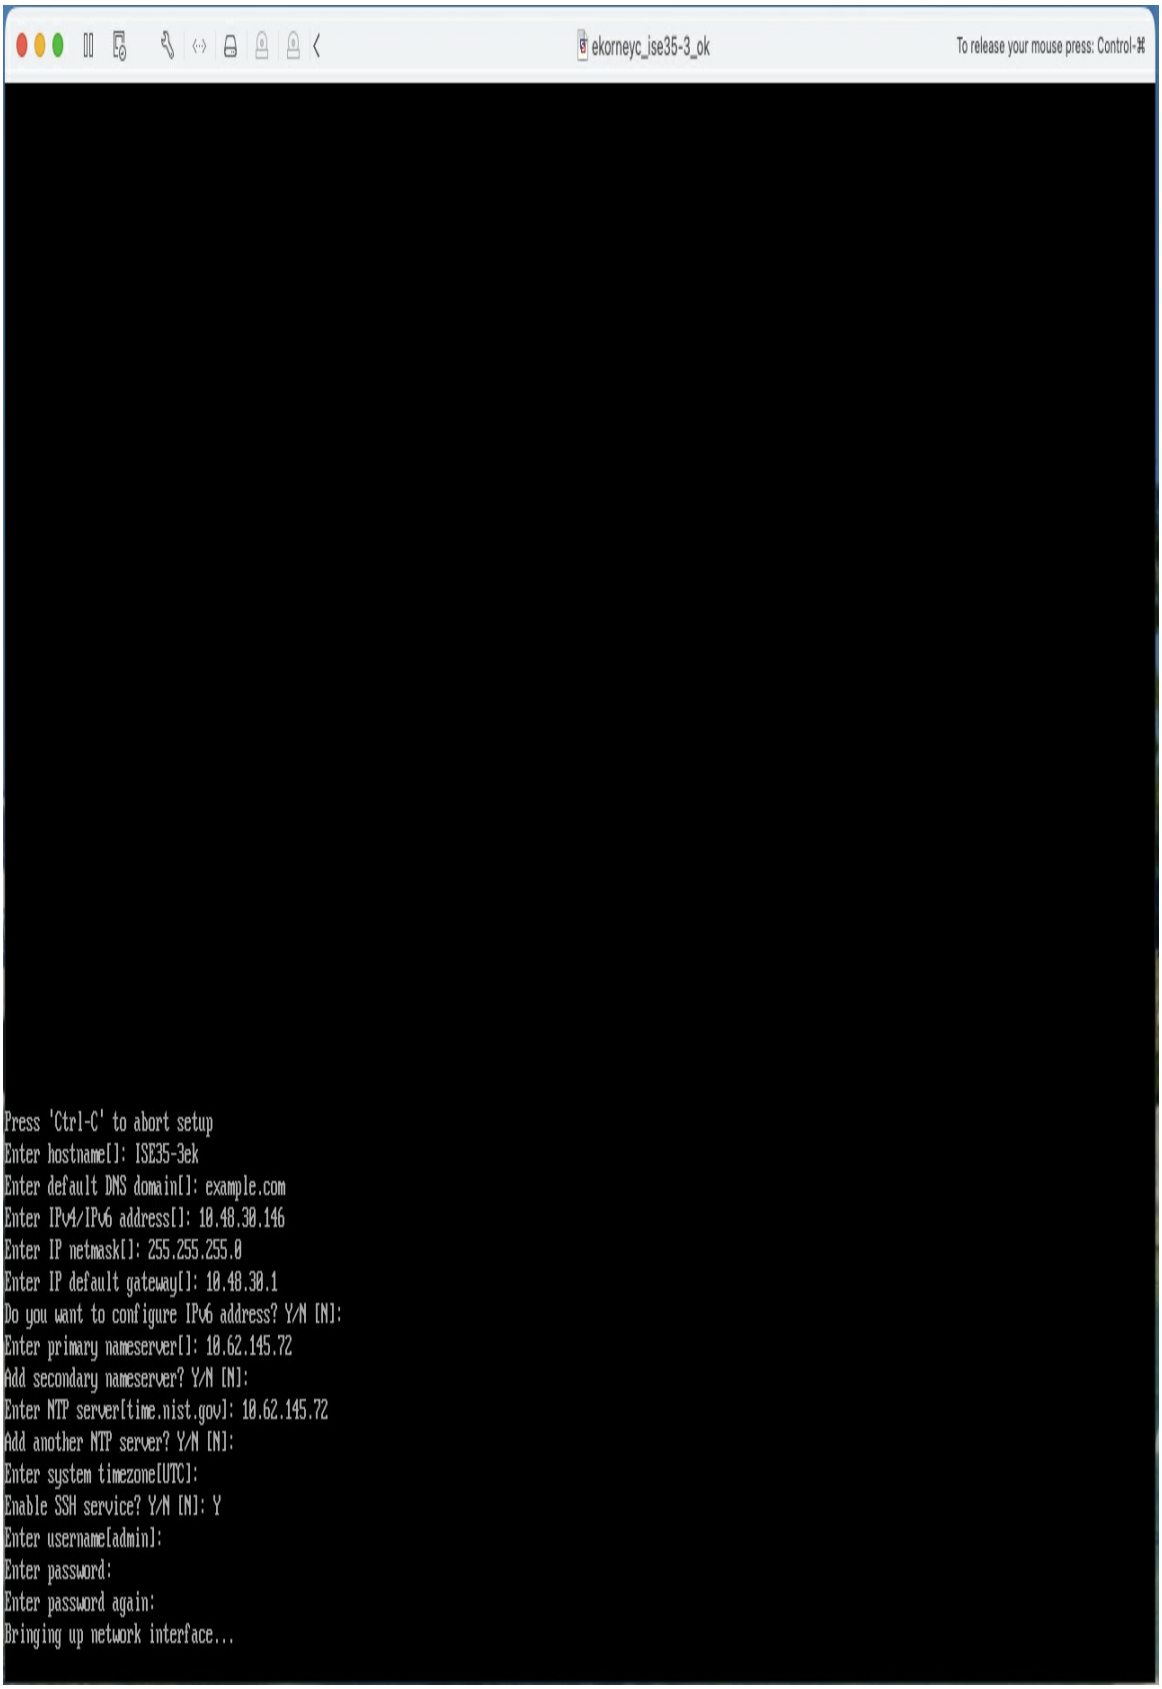

ekorneyc_ise35-3_ok
To release your mouse press: Control-⌘
Press 'Ctrl-C' to abort setup
Enter hostname[]: ISE35-3ek
Enter default DNS domain[]: example.com
Enter IPv4/IPv6 address[]: 10.48.30.146
Enter IP netmask[]: 255.255.255.0
Enter IP default gateway[]: 10.48.30.1
Do you want to configure IPv6 address? Y/N [N]:
Enter primary nameserver[]: 10.62.145.72
Add secondary nameserver? Y/N [N]:
Enter NTP server[time.nist.gov]: 10.62.145.72
Add another NTP server? Y/N [N]:
Enter system timezone[UTC]:
Enable SSH service? Y/N [N]: Y
Enter username[admin]:
Enter password:
Enter password again:
Bringing up network interface...

Figure 2-8 *ISE Initial Configuration*

Table 2-6 explains the purpose of each prompt.

Table 2-6 *ISE Configuration and Prompt Descriptions*

Technet24

Prompt	Description	Example
Hostname	Can be up to 19 characters, using only alphanumeric characters and hyphens. The first character must be a letter.	ISE35-3ek
Default DNS Domain	Must not be an IP address. Valid characters include ASCII characters, any numerals, the hyphen, and the period.	example.com
IPv4/IPv6 Address	Valid IPv4 or global IPv6 for the management interface. If IPv4 is entered, only IPv4 is accepted for the rest of the setup; similarly for IPv6 including the default gateway, name server, and NTP server. If you enter an IPv4 address and respond **yes** to the system prompt "Do you want to configure an IPv6 address?" the system will accept both IPv4 or IPv6 addresses for the rest of the configuration.	10.0.0.100
IP Netmask	Valid IPv4 or IPv6 netmask.	255.255.255.0
IP Default Gateway	Valid IPv4 or global IPv6 address for the default gateway.	10.0.0.1
Primary Name Server	Valid IPv4 or global IPv6 address for the primary name server.	10.0.0.250
NTP Server	Valid IPv4 or global IPv6 address for the primary name server.	10.0.0.250
System Timezone	Must be a valid time zone. For example, for Pacific Standard Time (PST), the System Timezone is PST8PDT, which is Coordinated Universal Time (UTC) minus 8 hours (UTC − 08:00 or 16:00).	UTC
SSH Service	Allows you to enable or disable SSH access to ISE. Select **Y** to allow ssh access and **N** to deny it.	Y
Username	Identifies the administrative username used for CLI access to the Cisco ISE system. If you choose not to use the default (admin), you must create a new username. The username must be three to eight characters in length and consist of valid alphanumeric characters (A—Z, a—z, or 0—9).	admin
Password	Identifies the administrative password that is used for CLI access to the Cisco ISE system. You must create this password to be able to continue because there is no default password. The password must be a minimum of six characters in length and include at least one lowercase letter (a—z), one uppercase letter (A—Z), and one numeral (0—9).	Cisc0123

Note

Starting with Cisco ISE release 3.5, Cisco SNS 37xx and Cisco SNS 38xx appliances support management connectivity through interfaces other than GigabitEthernet 0. During the installation process, the system prompts you to select the management interface.

ISE Installation Verification

To confirm that ISE is successfully installed, log in to the CLI via a console or an SSH connection and issue the command as shown in Example 2-1.

Example 2-1 *Confirming ISE Installation*

```
#show application status ise

ISE PROCESS NAME                       STATE            PROCESS I
------------------------------------------------------------------
Database Listener                      running          5059
Database Server                        running          138 PROCE
ISE Data Grid Service                  running          104080
Application Server                     running          13242
Profiler Database                      running          13772
ISE Elasticsearch                      running          52157
AD Connector                           running          58850
M&T Session Database                   running          9801
M&T Log Processor                      running          61142
Certificate Authority Service          running          66680
EST Service                            running          67706
SXP Engine Service                     disabled
TC-NAC Service                         disabled
PassiveID WMI Service                  disabled
PassiveID Syslog Service               disabled
PassiveID API Service                  disabled
PassiveID Agent Service                disabled
PassiveID Endpoint Service             disabled
```

Technet24

```
PassiveID SPAN Service                    disabled
DHCP Server (dhcpd)                       disabled
DNS Server (named)                        disabled
ISE Messaging Service                     running          14243
ISE API Gateway Database Service          running          14700
ISE API Gateway Service                   running          48638
ISE pxGrid Direct Service                 running          42923
ISE pxGrid Direct Pusher                  running          44557
Segmentation Policy Service               disabled
REST Auth Service                         running          53663
SSE Connector                             disabled
Hermes (pxGrid Cloud Agent)               disabled
MFA (Duo Sync Service)                    running          55291
McTrust (Meraki Sync Service)             disabled
aciconn (ACI Connection Service)          disabled
Workload Connector Service                disabled
ISE Prometheus Service                    disabled
ISE Prometheus Exporter                   disabled
ISE Grafana Service                       disabled
ISE MNT LogAnalytics Elasticsearch        running          57211
ISE Kibana Service                        disabled
ISE Native IPSec Service                  running          14112
Remote Support Authorization Service      disabled
MFC Profiler                              disabled
ISE Prometheus Alertmanager Service       disabled
Protocols Engine                          running          68667
```

Note

Some of the service outputs on your ISE node might not match exactly with Example 2-1. This difference is expected because the number of services depends on the personas selected and the configured services.

It is important to note that a "good" status for the services means either disabled or running. This status indicates that the services are turned off or working, respectively.

During ISE setup or restart, services may remain in a not running or initializing state for several minutes. This state is normal and reflects service initialization. However, when ISE is fully operational, no services should remain in a not running or initializing state. If this situation occurs, it indicates a problem with the affected services crashing or continuously restarting.

It is crucial to ensure that the ISE VM profile allocation has been configured correctly.

Profile allocation controls resource distribution within ISE. You can run the following command to verify it:

```
#show tech-support | begin "ISE Profile"
Displaying ISE Profile ...
**************************
Profile : sns3895
```

Latency Requirements

Latency requirements are especially important for communication between PAN and other nodes in the deployment. This channel handles database replication. The maximum round-trip latency is 300 ms.

The maximum does not mean that an ISE deployment with 310 ms latency between nodes will not work, but for optimal replication and a stable cluster, you should lower latency numbers as much as possible. Such latency requirements can influence design decisions, such as determining the optimal location for the primary PAN in a globally distributed deployment or deciding whether to send RADIUS traffic over a WAN link instead of replication traffic.

For example, in a deployment with nodes located in Asia, Europe, and the United States, placing the primary PAN in Europe may provide the

Technet24

best balance of latency between the US and Asia.

RADIUS latency is more tolerant of higher latency and is generally measured in seconds rather than milliseconds. Requirements may vary based on the number of endpoints, authentication rates, and protocols chosen. Optimal settings for the RADIUS server timeout should be set within 5 — 10 seconds.

There are no specific requirements for communication between other nodes in the cluster. However, it is recommended to avoid using slow links because all nodes communicate with each other within an ISE distributed deployment.

High Availability

High availability is crucial for every component, including network path and workload. ISE is no exception, and designing the ISE cluster for high availability is one of the most important tasks. If one component fails, another should be ready to take over to maintain business continuity.

PSN High Availability

PSN high availability is often misunderstood. All ISE PSNs are available to handle authentications at any time. The ISE cluster or an individual PSN does not control the switchover. Network access devices (NADs) or load balancers determine which PSN in the cluster to use. If one PSN goes down, the NAD or load balancer should switch authentications to another PSN, depending on timers and configuration settings.

For example, to configure a Cisco IOS XE switch to declare a RADIUS server as not reachable, you can use the following:

```
#radius-server dead-criteria time 10 tries 3
```

- **Time:** Define the period of time, in seconds, during which there is no response from the RADIUS server.
- **Tries:** Define the number of consecutive timeouts that must occur on the NAD before the RADIUS server is declared unreachable.

The following command controls the period of time, in minutes, during which the RADIUS server will not be used if it has been declared unresponsive:

```
#radius-server deadtime 15
```

If radius-server deadtime is not configured, it defaults to 0. After the RADIUS server is declared unresponsive, it becomes active again immediately.

Even if deadtime is configured, with the previous settings, there is no reliable way to confirm that the RADIUS server is available before it is brought back into service. The server becomes active based only on the timer, which can occur even if the RADIUS server is still unresponsive.

To address this issue, the RADIUS probe-on functionality was introduced:

```
radius server ISE1
  address ipv4 10.0.0.100 auth-port 1812 acct-port 1813
  automate-tester username test ignore-acct-port probe-on
  key
```

With this configuration, if the server was previously declared unresponsive, it is brought up only when it responds to the test probe, ensuring reliability.

It is best practice to configure the test user on both the switch and ISE

Technet24

and have a corresponding automated-tester policy set configured. Both Access-Accept and Access-Reject are considered valid responses from ISE to bring up the server.

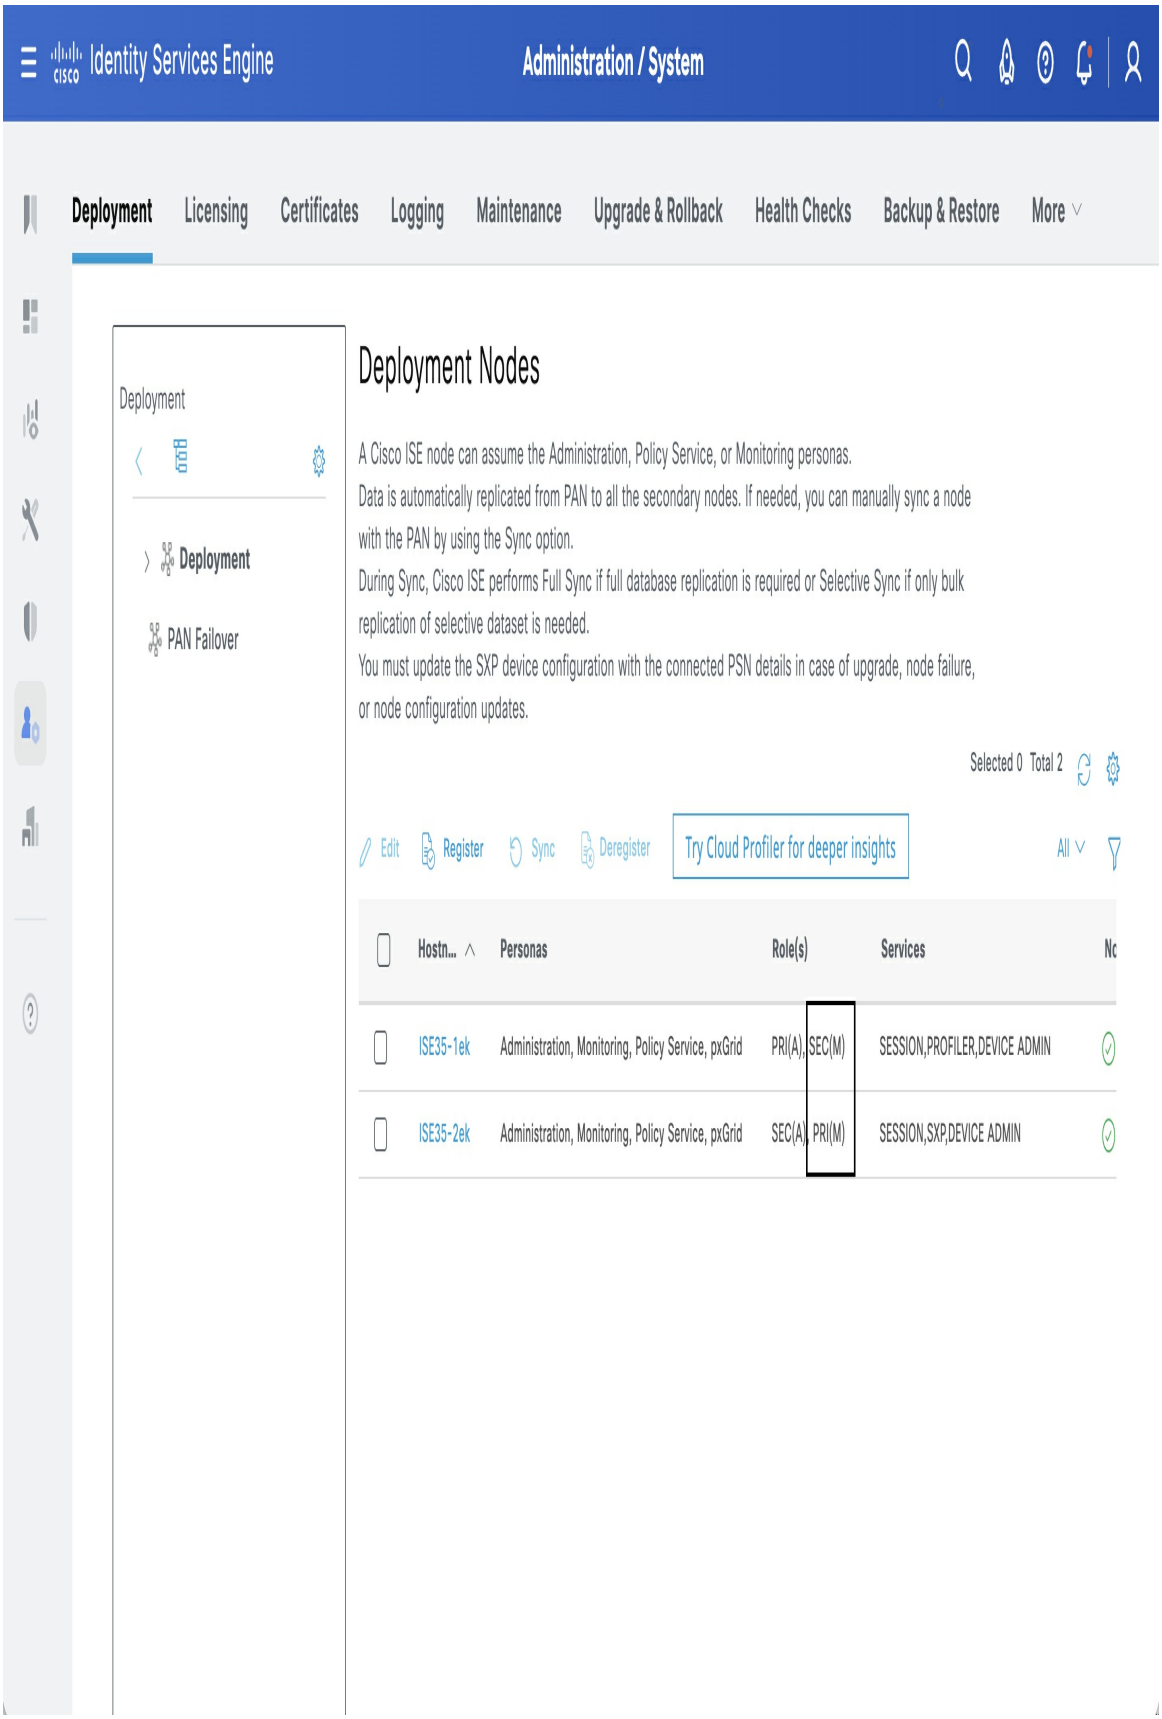
The probe-on functionality sends RADIUS authentication requests only when the server is declared unresponsive. The **ignore-acct-port** command instructs the switch not to validate RADIUS communication over accounting ports.

It is recommended to configure NADs with at least two RADIUS servers (PSNs) for high availability.

MnT High Availability

You can have a maximum of two MnT nodes in an ISE deployment. Both MnT nodes are active at all times and receive log messages in the form of syslogs from other nodes in the deployment. Syslog messages contain information about user authentication, device administration, node health, and more.

When configuring MnT nodes in the deployment, you can select which MnT node is primary (shown as **PRI(M)** in the **Role(s)** column on the deployment page) or secondary (shown as **SEC(M)** in the **Role(s)** column), as illustrated in Figure 2-9.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Deployment
Deployment
PAN Failover
Deployment Nodes
A Cisco ISE node can assume the Administration, Policy Service, or Monitoring personas.
Data is automatically replicated from PAN to all the secondary nodes. If needed, you can manually sync a node with the PAN by using the Sync option.
During Sync, Cisco ISE performs Full Sync if full database replication is required or Selective Sync if only bulk replication of selective dataset is needed.
You must update the SXP device configuration with the connected PSN details in case of upgrade, node failure, or node configuration updates.
Selected 0 Total 2
Edit
Register
Sync
Deregister
Try Cloud Profiler for deeper insights
All
Hostn...
Personas
Role(s)
Services
ISE35-1ek
Administration, Monitoring, Policy Service, pxGrid
PRI(A), SEC(M)
SESSION,PROFILER,DEVICE ADMIN
ISE35-2ek
Administration, Monitoring, Policy Service, pxGrid
SEC(A), PRI(M)
SESSION,SXP,DEVICE ADMIN

Technet24

Figure 2-9 *ISE Deployment MnT Roles*

While both MnT nodes are active and receive the same syslog messages from all other nodes in the deployment, only the primary MnT node is used when you log in to the PAN to retrieve monitoring and troubleshooting data, such as live logs, live sessions, or reports.

High availability for MnT nodes is built in. You only need to ensure that both primary and secondary MnT roles are allocated in your deployment. If the primary MnT node goes down, the PAN will automatically use the secondary node for log retrieval.

ISE uses ISE Messaging certificates to protect the transfer of syslog data. Syslog data is encapsulated within TLS tunnels protected by these certificates, and this feature is enabled by default. As an additional benefit, nodes buffer messages if an MnT node goes down and send them to the MnT node when it comes back online. This scenario would not occur with classic syslog messages, which would simply be lost, resulting in data inconsistency between MnT nodes. The MnT survivability limit is a 4 GB queue of messages.

For additional log retention and audit purposes, you can configure extra syslog targets. To add a syslog target, navigate to **Administration > System > Logging > Remote Logging Targets** and click **+Add**, as shown in Figure 2-10.

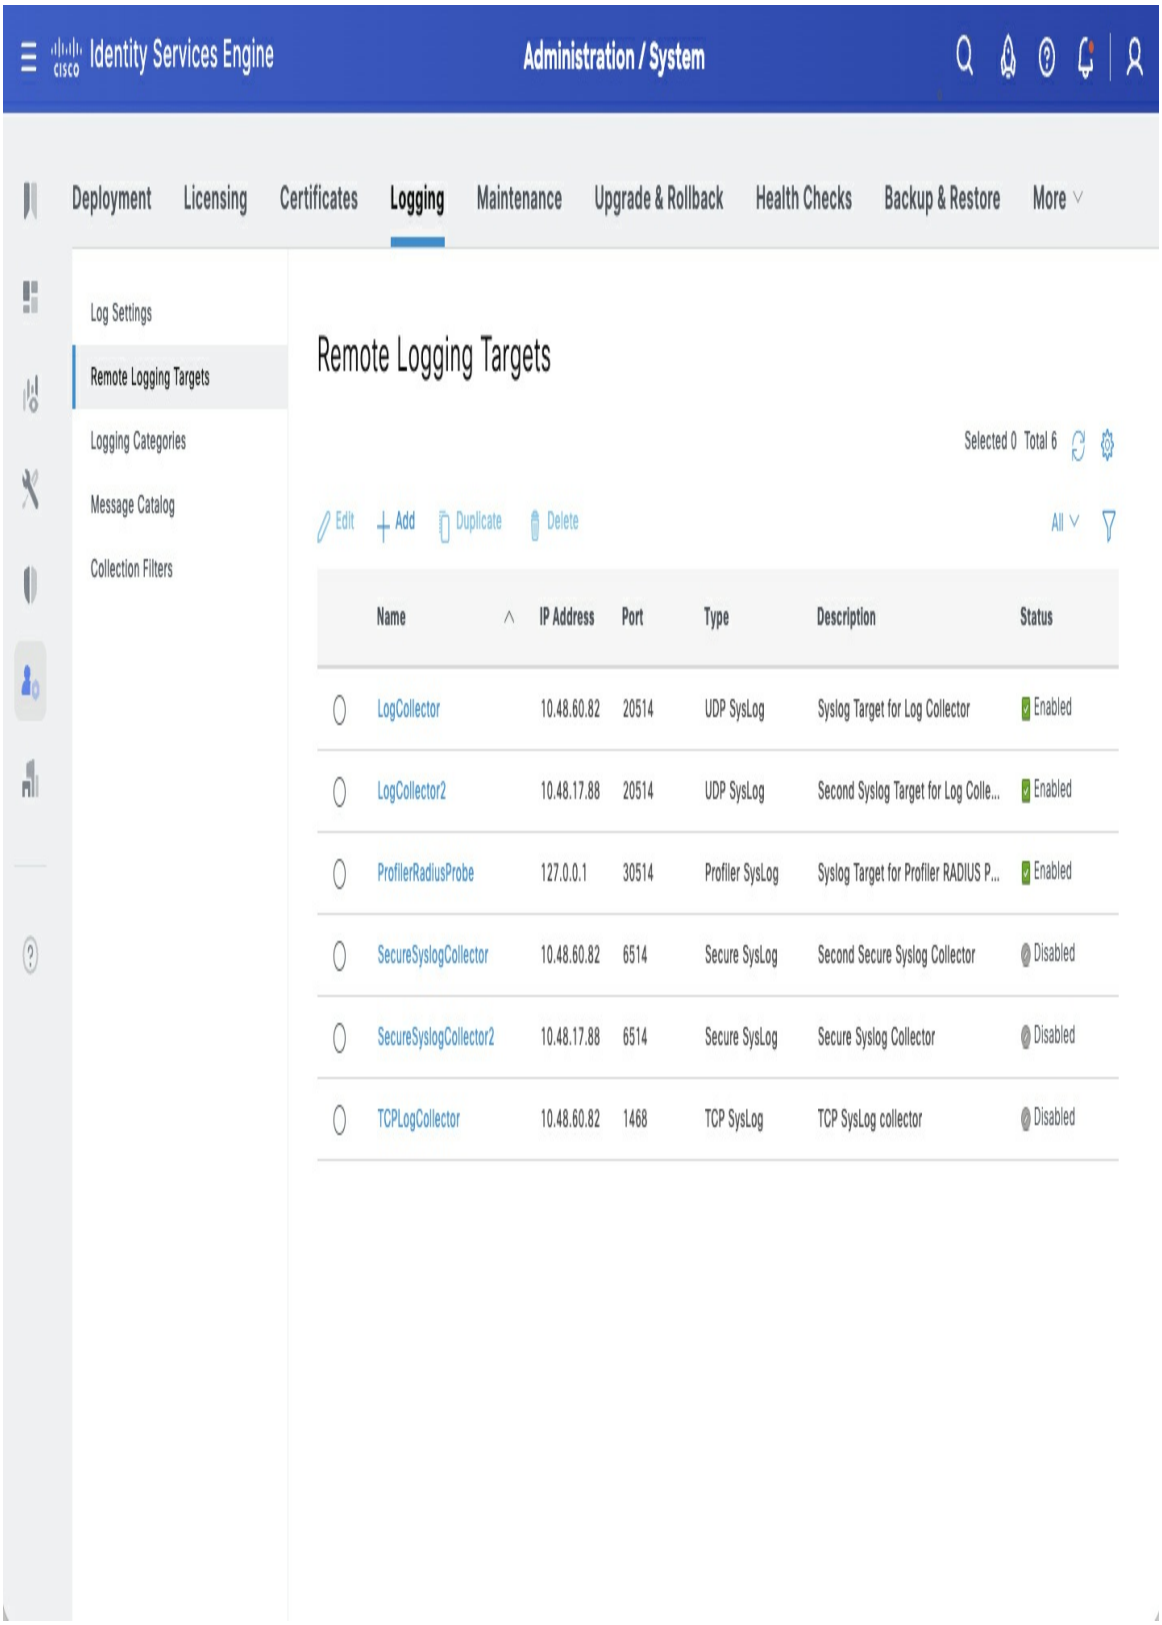


Figure 2-10 *Remote Logging Targets*

Technet24

To configure a Remote Logging Target, enter the name and the hostname or IP address of the syslog server, as shown in Figure 2-11. You can also specify a custom port and choose **Secure Syslog** or **TCP Syslog** as alternatives to UDP transport. When everything is configured, click **Submit**.

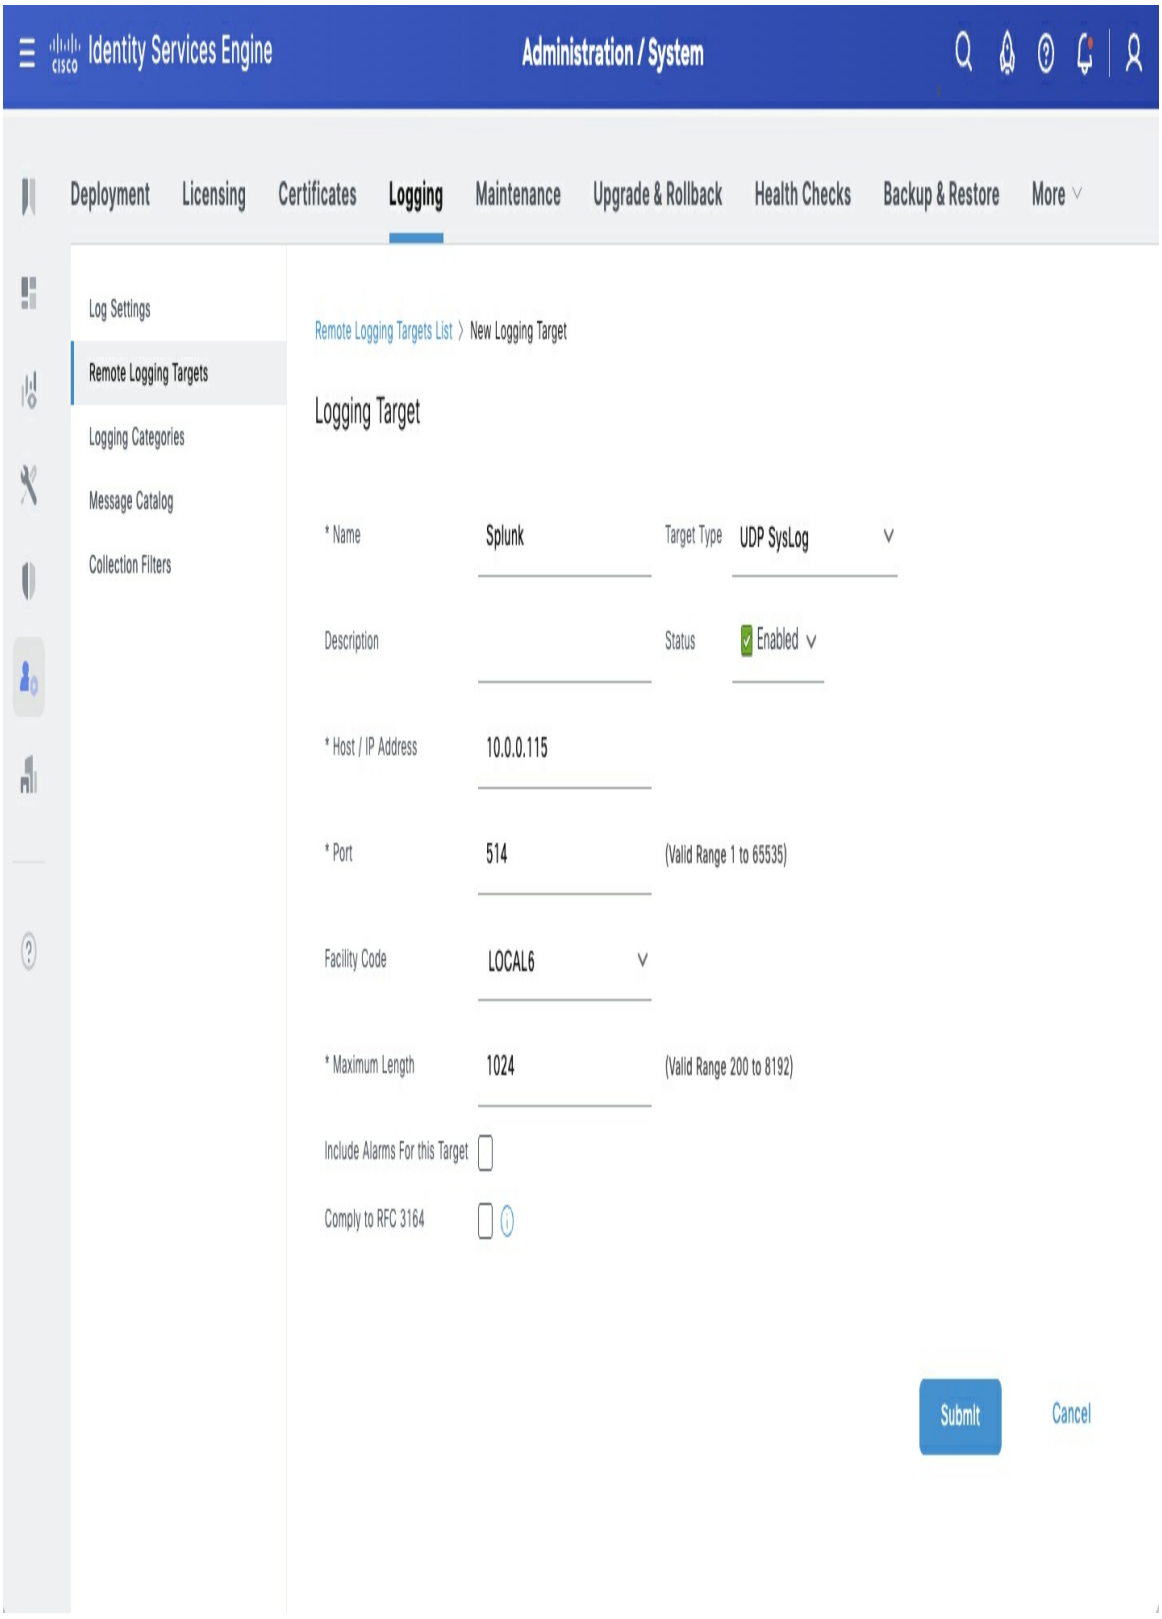
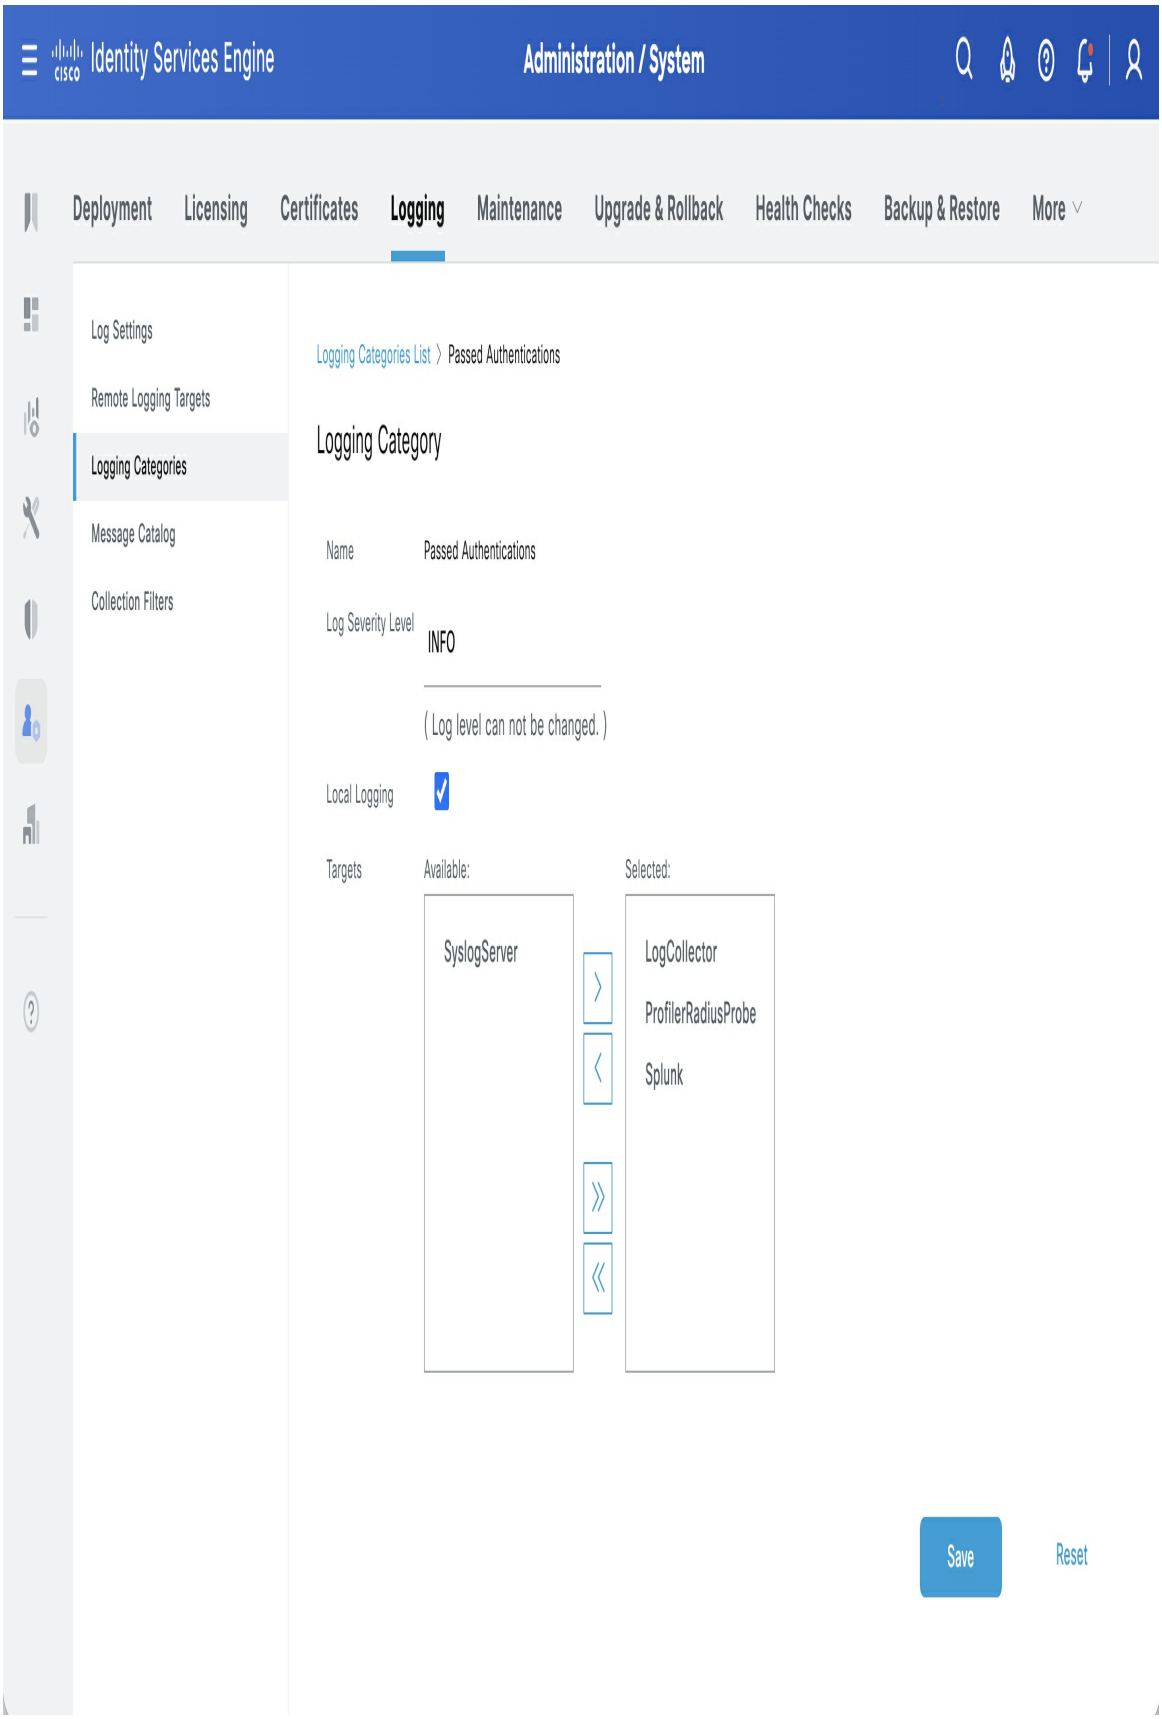


Figure 2-11 *ISE Logging Target Configuration*

Technet24

You can control which logging categories are sent to the newly configured Syslog Target under **Administration > System > Logging > Logging Categories**. Figure 2-12 shows the Splunk Syslog Target configured for Passed Authentications.

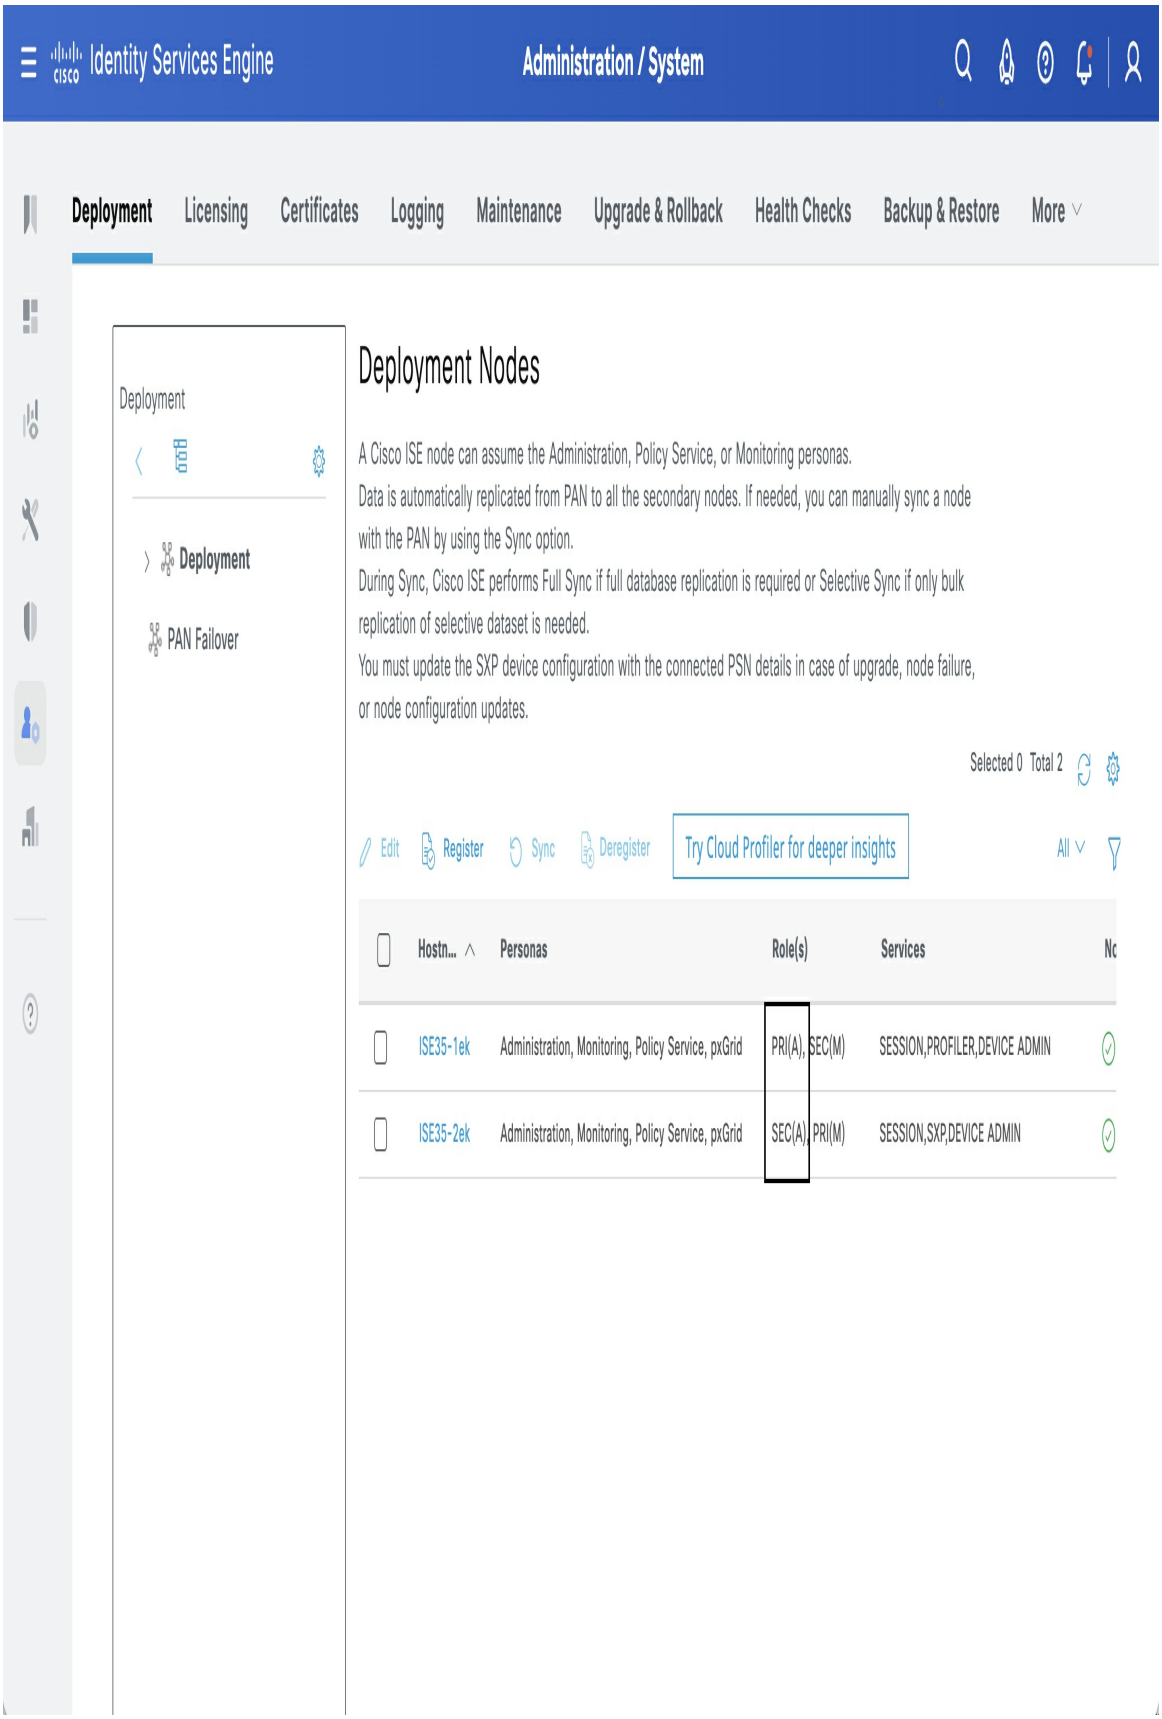

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Log Settings
Remote Logging Targets
Logging Categories
Message Catalog
Collection Filters
Logging Categories List
Passed Authentications
Logging Category
Name
Passed Authentications
Log Severity Level
INFO
(Log level can not be changed.)
Local Logging
Targets
Available:
Selected:
SyslogServer
LogCollector
ProfilerRadiusProbe
Splunk
Save
Reset

Technet24

Figure 2-12 *ISE Logging Category Settings*

PAN High Availability

Similar to MnT nodes, you can have up to two PAN nodes in the deployment. The primary policy administration node is referred to as PPAN (PRI(A)), and the secondary policy administration node is referred to as SPAN (SEC(A)), as shown in Figure 2-13.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Deployment
Deployment
PAN Failover
Deployment Nodes
A Cisco ISE node can assume the Administration, Policy Service, or Monitoring personas.
Data is automatically replicated from PAN to all the secondary nodes. If needed, you can manually sync a node with the PAN by using the Sync option.
During Sync, Cisco ISE performs Full Sync if full database replication is required or Selective Sync if only bulk replication of selective dataset is needed.
You must update the SXP device configuration with the connected PSN details in case of upgrade, node failure, or node configuration updates.
Selected 0 Total 2
Edit
Register
Sync
Deregister
Try Cloud Profiler for deeper insights
All
Hostn...
Personas
Role(s)
Services
ISE35-1ek
Administration, Monitoring, Policy Service, pxGrid
PRI(A), SEC(M)
SESSION,PROFILER,DEVICE ADMIN
ISE35-2ek
Administration, Monitoring, Policy Service, pxGrid
SEC(A), PRI(M)
SESSION,SXP,DEVICE ADMIN

Technet24

Figure 2-13 *ISE Deployment PAN Roles*

Unlike MnT nodes, PAN failover is not automatic by default. The GUI of all secondary nodes (nodes that are not PPAN) is limited, since most controls belong to the PPAN. You can perform PAN promotion from the GUI of the secondary node when the primary PAN is down.

To promote the SPAN to PPAN, navigate to **Administration > System > Deployment** and click **Promote to Primary**, as shown in Figure 2-14.

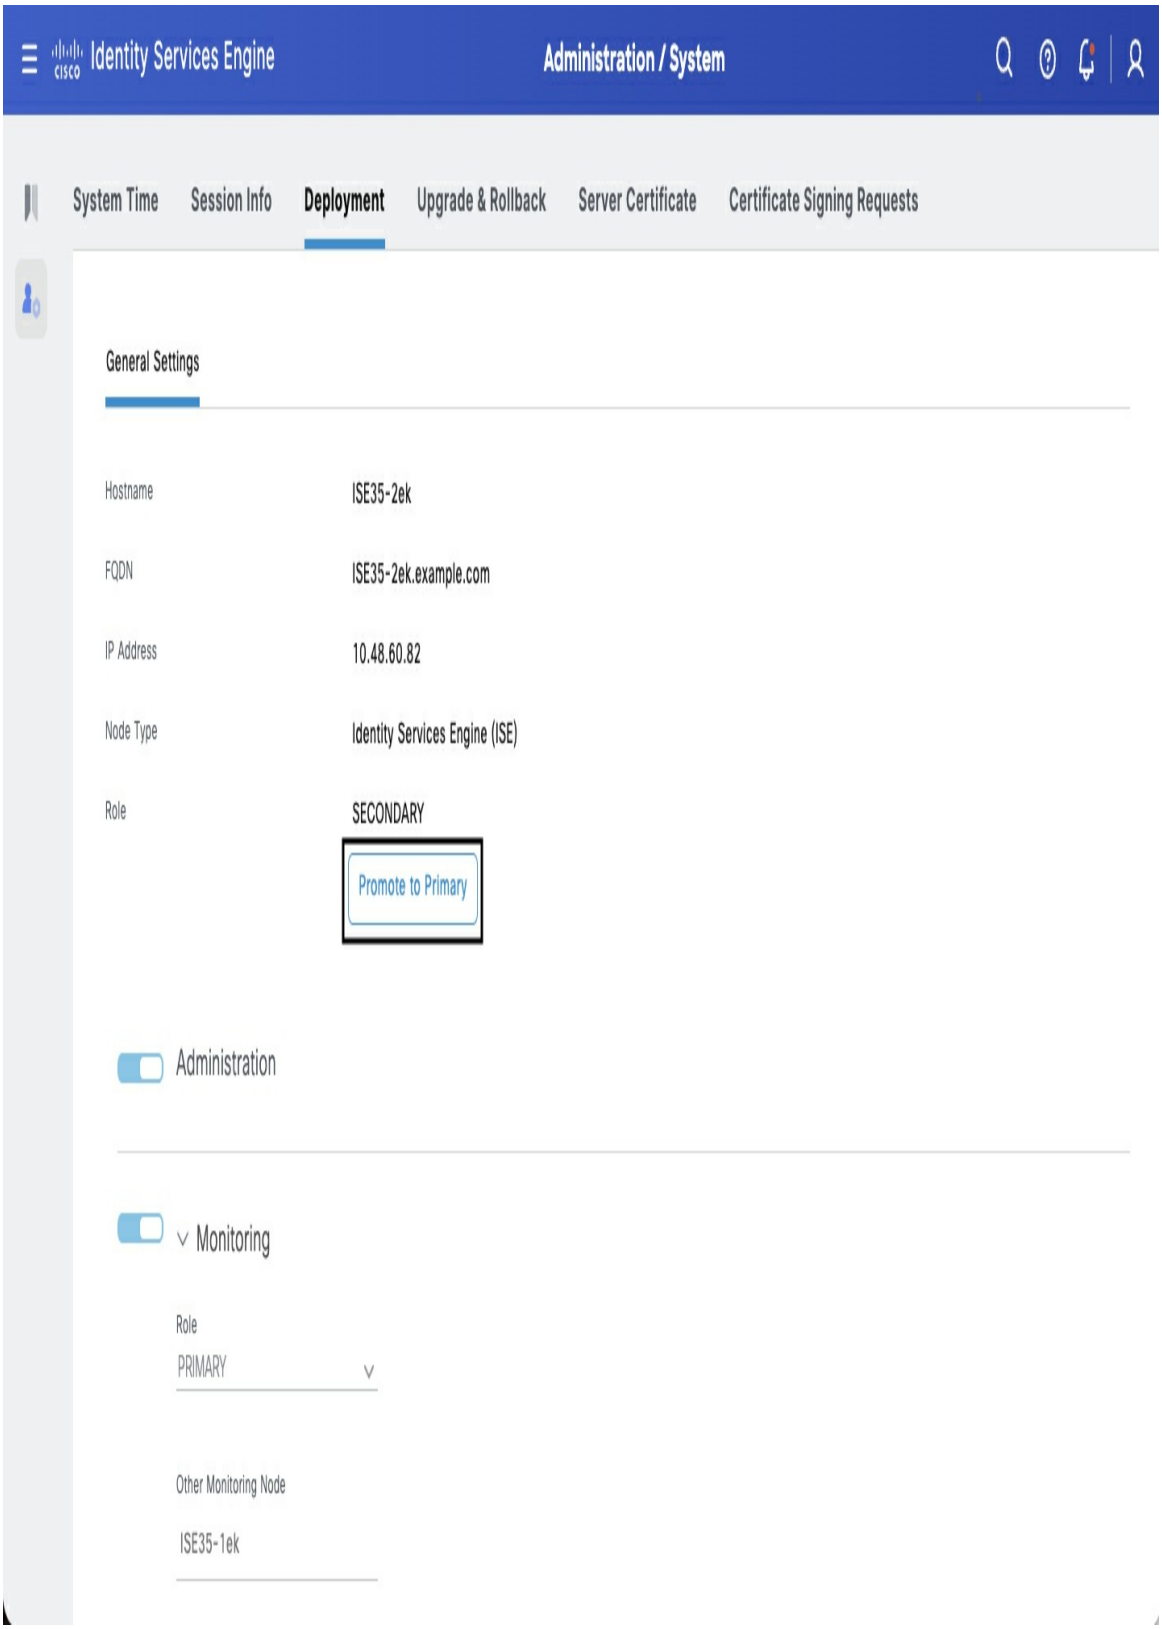
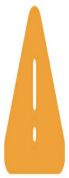


Figure 2-14 *Secondary PAN Promotion to Primary*

Technet24

When the promotion is requested, the warning shown in Figure 2-15 is displayed.



Warning

You have requested to promote this node to the primary administrative role. Upon completion of this request you will be logged out of the user interface. Services on both Administration nodes will be restarted. It will take up to 10 minutes for both nodes to attain new role. In case a manual Full Sync has been initiated for this node, you may want to wait till the completion of Sync. Are you sure you want to continue?

Cancel OK

Technet24

Figure 2-15 *Promotion to Primary PAN Warning*

Promotion to primary requires restarting the services on the node, making it unavailable during the restart.

PPAN is one of the most crucial components of an ISE deployment. Table 2-7 shows how features are affected if the PPAN goes down.

Table 2-7 *Features Availability When PAN Is Down*

Feature Name	Available When Primary PAN Is Down? (Yes/No)
Existing internal user RADIUS authentication	Yes
Existing or new AD user RADIUS authentication	Yes
Existing endpoint with no profile change	Yes
Existing endpoint with profile change	No
New endpoint learned through profiling.	No
Existing guest: Local Web Authentication (LWA)	Yes
Existing guest: Central Web Authentication (CWA)	Yes (apart from flows enabled for device registration, such as Hotspot, BYOD, and CWA with automatic device registration)
Guest change password	No
Guest: AUP	No
Guest: Max Failed Login Enforcement	No
New Guest (Sponsored or Self-registered)	No
Posture	Yes
BYOD with Internal CA	No
Existing Registered Devices	Yes
MDM on-boarding	No
pxGrid Service	No
System 360—Monitoring	No
System 360—Log Analytics	No (When the primary PAN also has a secondary MNT persona)
Log in to GUI of secondary nodes	Yes (The login process is delayed because a blocking call to the PAN is attempted to update the last login details. Login proceeds after this call times out.)

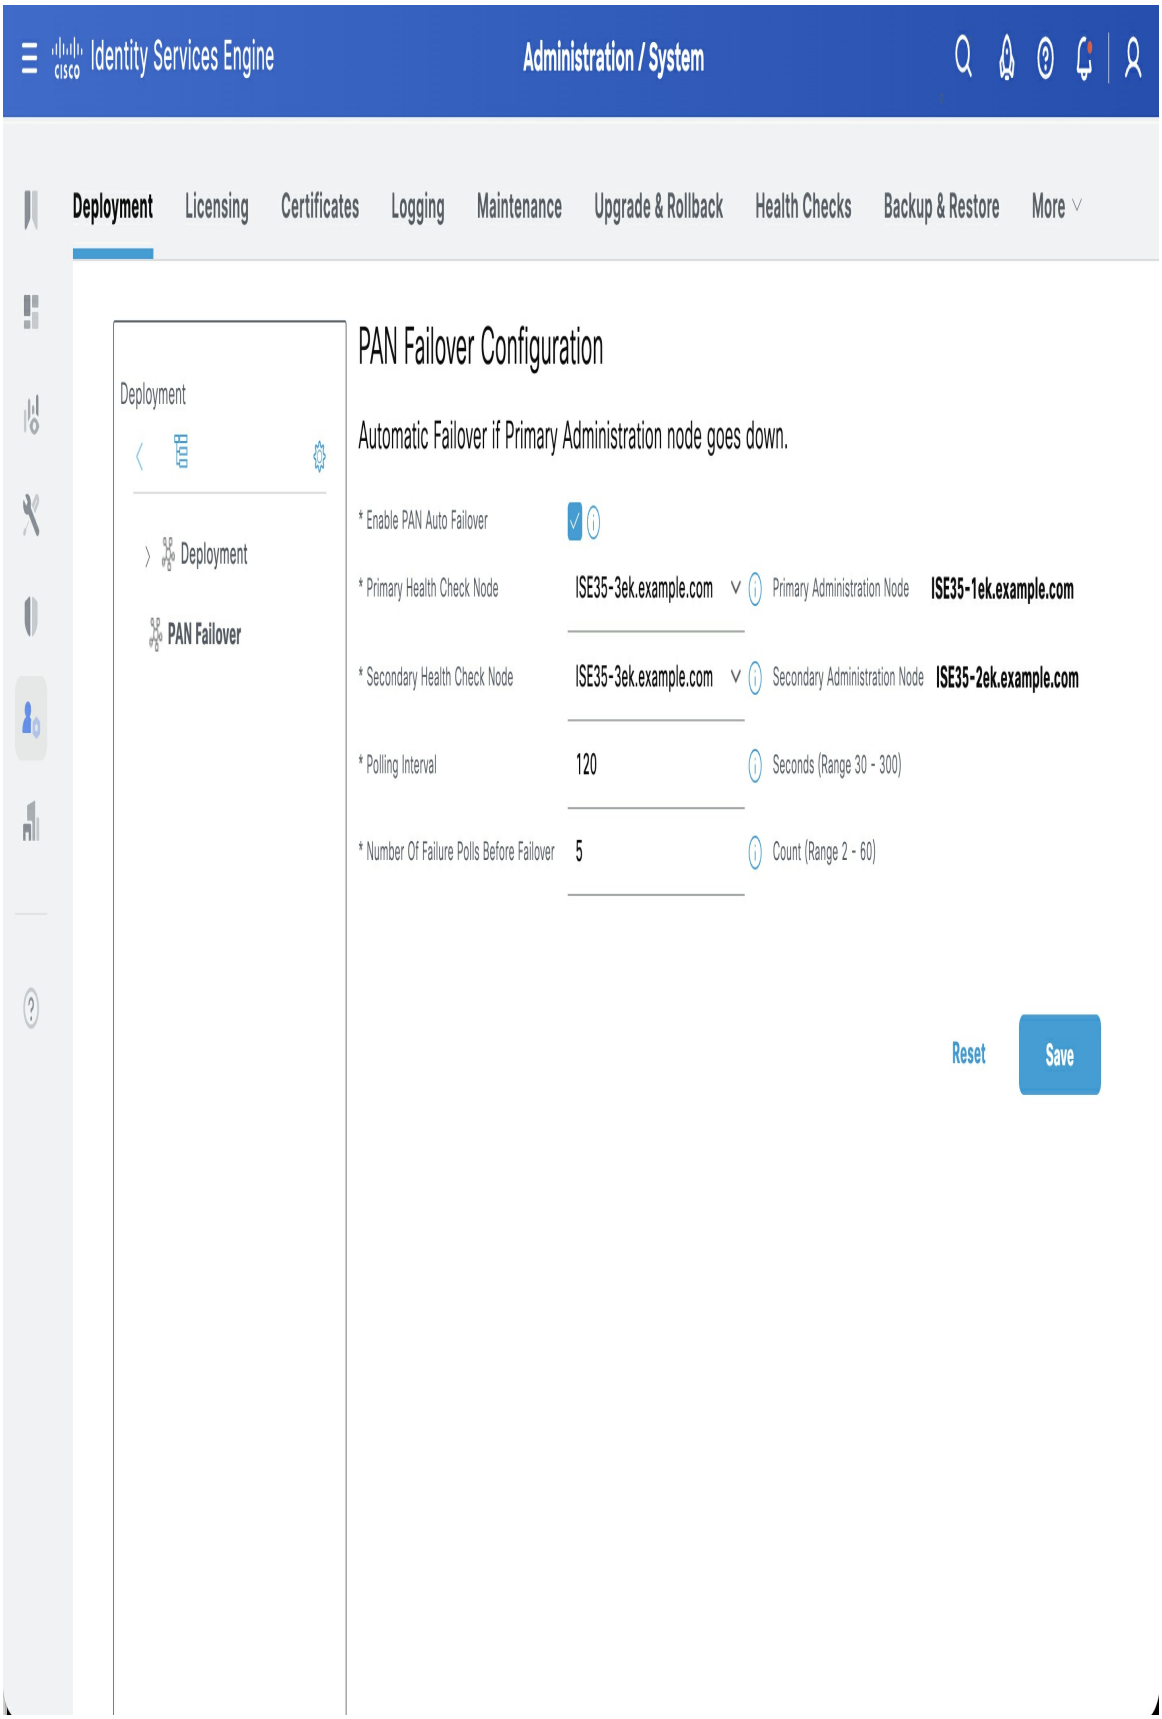
Technet24

An alternative to manual PAN failover is automatic PAN failover. Automatic PAN failover requires a third node to be present in the deployment. The purpose of this third node is to act as a health check node, determining when the PAN is down and when it comes back online.

The health check node can be any nonadministrative node in the deployment, such as an MnT, PSN, or pxGrid node, or any combination of these roles. You can have up to two health check nodes in the deployment. It is recommended to configure two health check nodes if the PPAN and SPAN are located in different data centers.

The active health check node monitors the availability of the PPAN. The health check node that monitors the SPAN is called the passive health check node; it remains inactive until the SPAN is promoted. When promotion occurs, the passive health check node becomes the active health check node and starts monitoring the new PPAN node in the deployment. At this point, the original active health check node becomes the passive health check node.

To configure the deployment for automatic PAN failover, navigate to **Administration > System > Deployment > PAN Failover**. Select the checkbox to enable **PAN Auto Failover**. Choose the health check nodes for both PPAN and SPAN, and configure the polling interval and number of retries, as shown in Figure 2-16. Then click **Save**.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Deployment
Deployment
PAN Failover
PAN Failover Configuration
Automatic Failover if Primary Administration node goes down.
* Enable PAN Auto Failover
* Primary Health Check Node
ISE35-3ek.example.com
Primary Administration Node
ISE35-1ek.example.com
* Secondary Health Check Node
ISE35-3ek.example.com
Secondary Administration Node
ISE35-2ek.example.com
* Polling Interval
120
Seconds (Range 30 - 300)
* Number Of Failure Polls Before Failover
5
Count (Range 2 - 60)
Reset
Save

Technet24

Figure 2-16 *PAN Auto Failover Configuration*

The polling interval controls how often a health check node verifies the health of the PAN node. The number of failures controls how many attempts are made before declaring that the PAN is down.

Several operations are affected when PAN failover is enabled. You cannot upgrade, install a patch, perform certain CLI operations, or restore a backup.

If you want to continue using automatic PAN failover but need to install a patch, you can temporarily disable automatic PAN failover and enable it again after the patch is installed.

NIC Bonding

ISE supports high availability at the interface level through network interface card (NIC) bonding or NIC teaming.

When you bond two interfaces, they appear as a single interface with one MAC address. NIC bonding allows you to protect ISE services from physical interface failures and loss of switchport connectivity due to shutdown or failure. However, it does not provide load balancing or link aggregation.

ISE uses fixed pairs of interfaces, where the first becomes the primary and the second acts as the backup. All traffic flows through the primary interface, and the backup is used only if a primary interface failure is detected. The IP and MAC addresses of the primary interface are assigned to the bond interface.

Table 2-8 shows the mapping between physical interfaces and bond interfaces.

Table 2-8 *Physical Interface Pairs Mapping to Bond Interfaces and*

Their Role in the Bond

Cisco ISE Physical NIC Name	Role in Bonded NIC	Bonded NIC Name
Gigabit Ethernet 0	Primary	Bond 0
Gigabit Ethernet 1	Backup	
Gigabit Ethernet 2	Primary	Bond 1
Gigabit Ethernet 3	Backup	
Gigabit Ethernet 4	Primary	Bond 2
Gigabit Ethernet 5	Backup	

NIC bonding is supported on ISE hardware appliances and virtual machines.

To configure NIC bonding, issue the following command under the primary interface configuration:

```
ISE35-1ek/admin(config-GigabitEthernet-0)#backup interface GigabitEt
```

As a result of this operation, services are restarted on the configured node.

To remove bonding from the interface, issue the same command with the no argument:

```
ISE35-1ek/admin(config-GigabitEthernet-0)#no backup interface Gigabi
```

Services are restarted during the removal of the backup interface as well.

Building Distributed Deployment

Technet24

When ISE is initially bootstrapped, it starts as a standalone node, as shown in Figure 2-17.

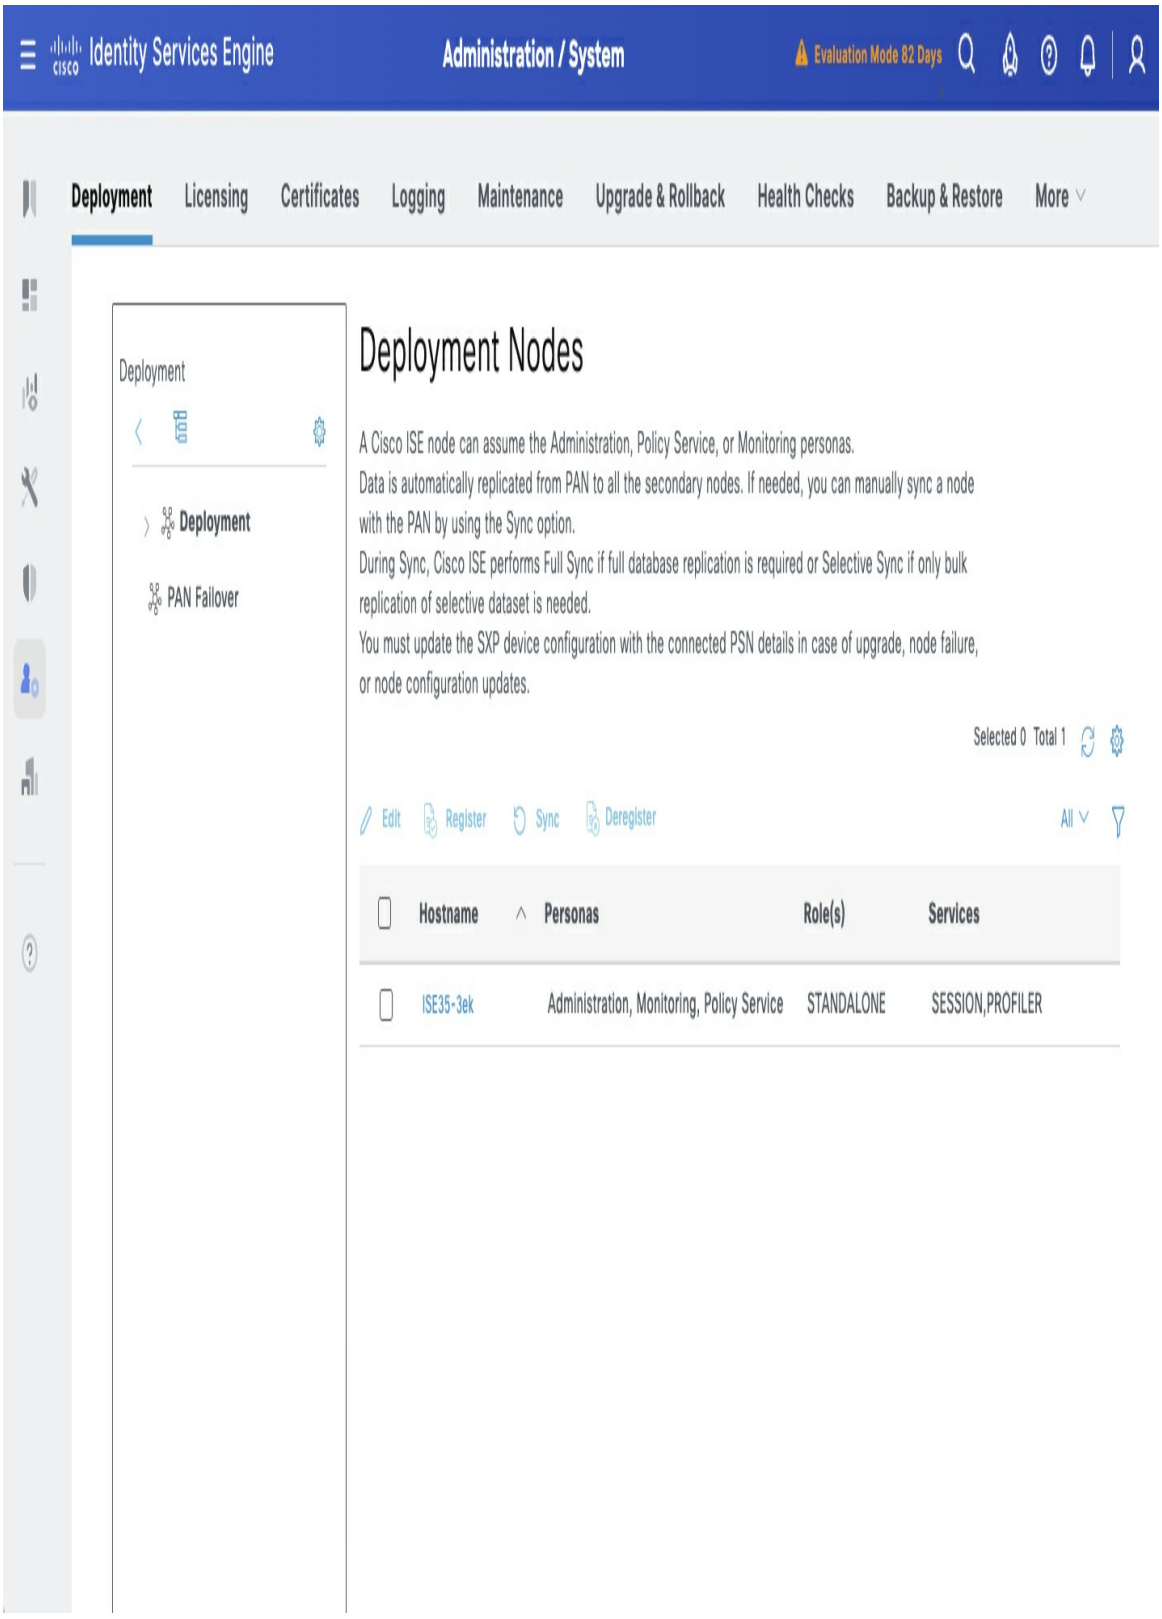


Figure 2-17 *ISE as a Standalone Node*

Technet24

If you want to designate the node as the PPAN of your deployment, the first step is to promote it to primary PAN. Navigate to **Administration > System > Deployment**, select the node, and click **Make Primary**, as shown in Figure 2-18.

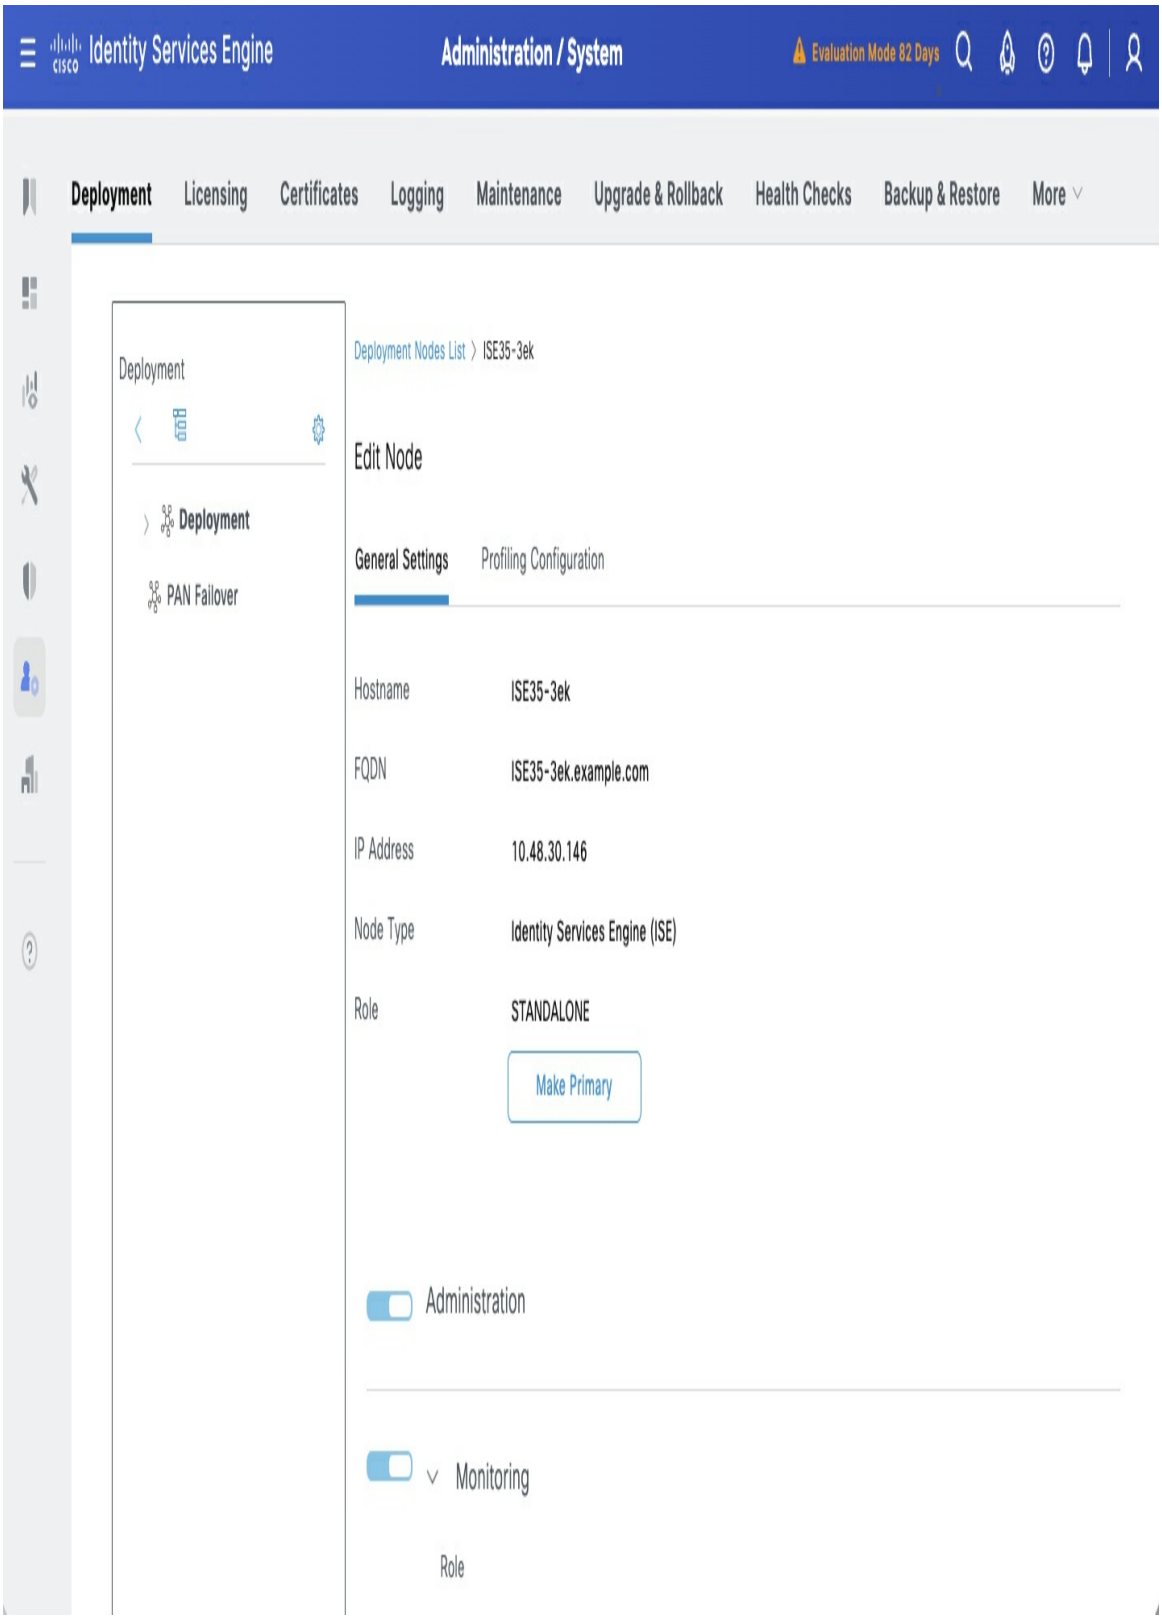


Figure 2-18 *ISE Standalone Promotion to Primary*

Technet24

> **Note**
>
> When you promote a standalone node to primary, services do not restart. However, if you change the role from primary to standalone, a service restart is required.

If you want to add a node to an existing deployment, the node role should be set to standalone.

Before you start building your distributed deployment, it is recommended to install CA-signed admin certificates, ideally generated by the same certificate authority. Relying on the self-signed certificates that ISE provides by default is not scalable.

ISE certificate installation is covered in the "Certificates" section in Chapter 5, "Authentication Options."

Although it is not recommended, ISE allows you to build a distributed deployment using self-signed certificates.

To add a secondary node to the deployment, navigate to **Administration > System > Deployment** and click the **Register** button. Figure 2-19 shows how a node with the FQDN ISE35-3ek.example.com is being added to an existing deployment consisting of two nodes — ISE35-1ek and ISE35-2ek.

You need to provide the FQDN of the node you are trying to add to deployment and admin credentials of that node.

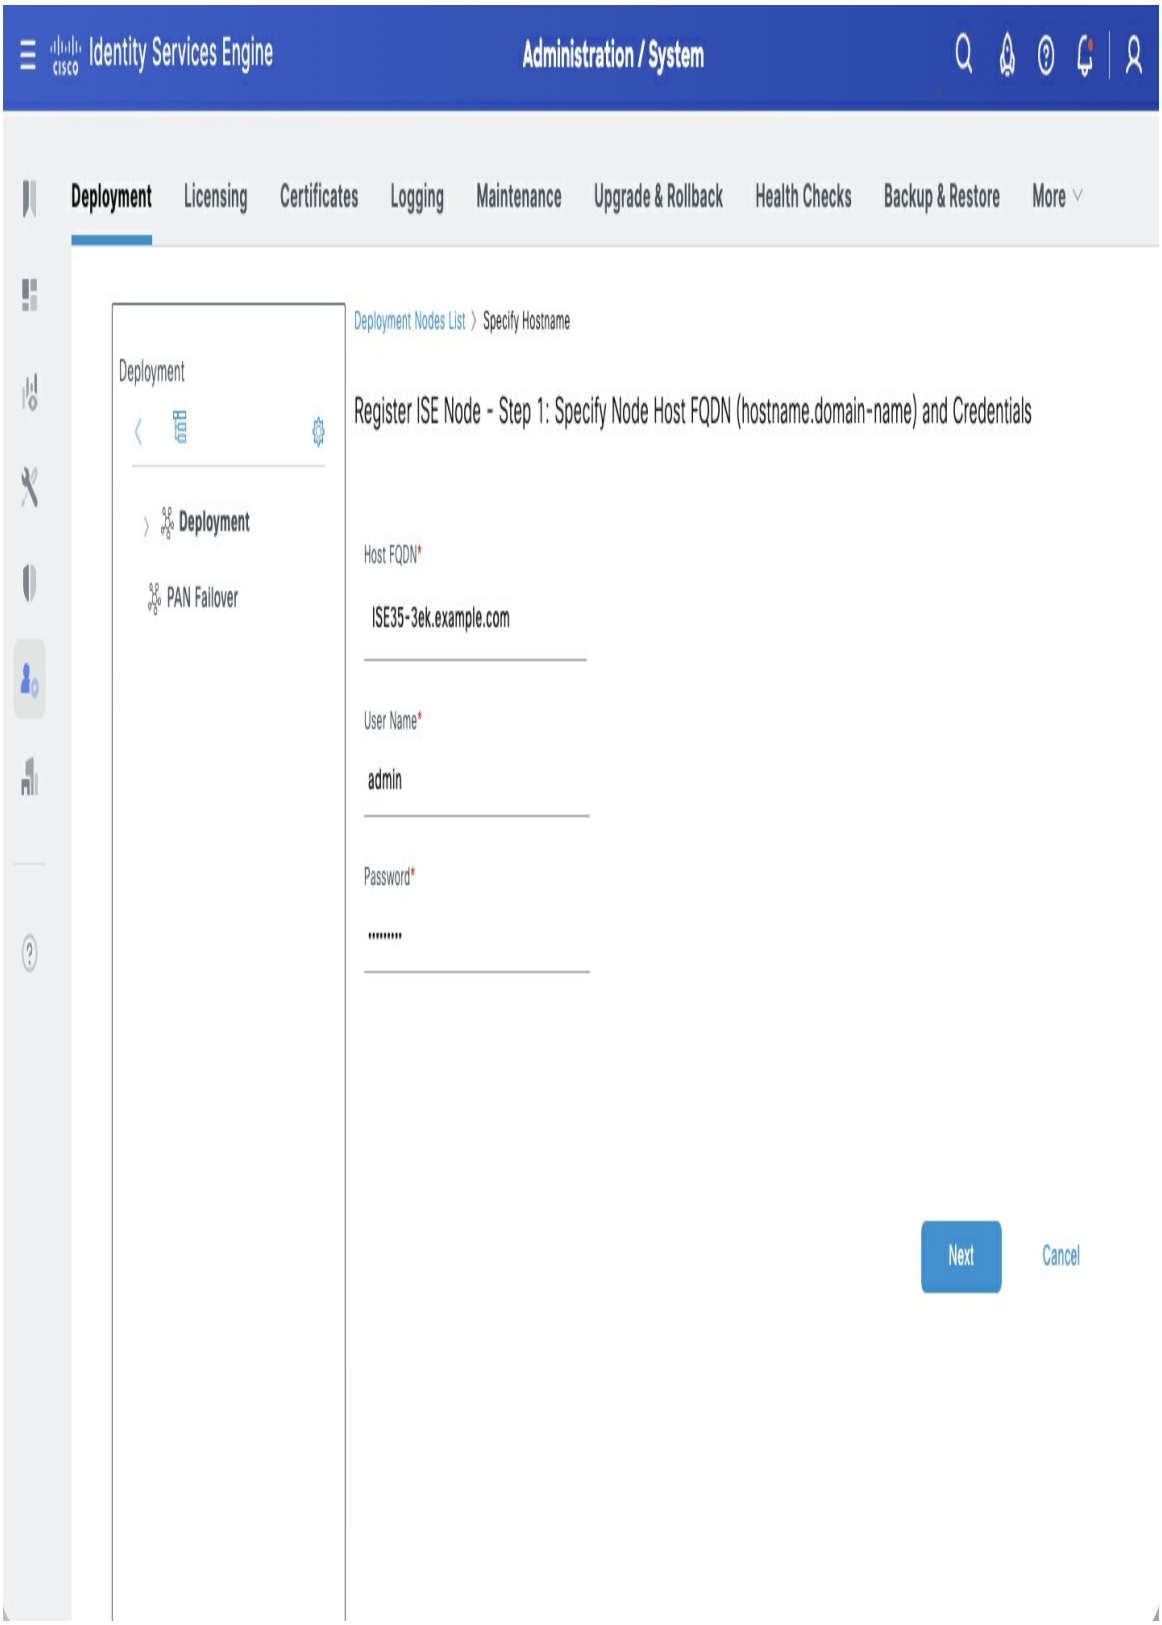
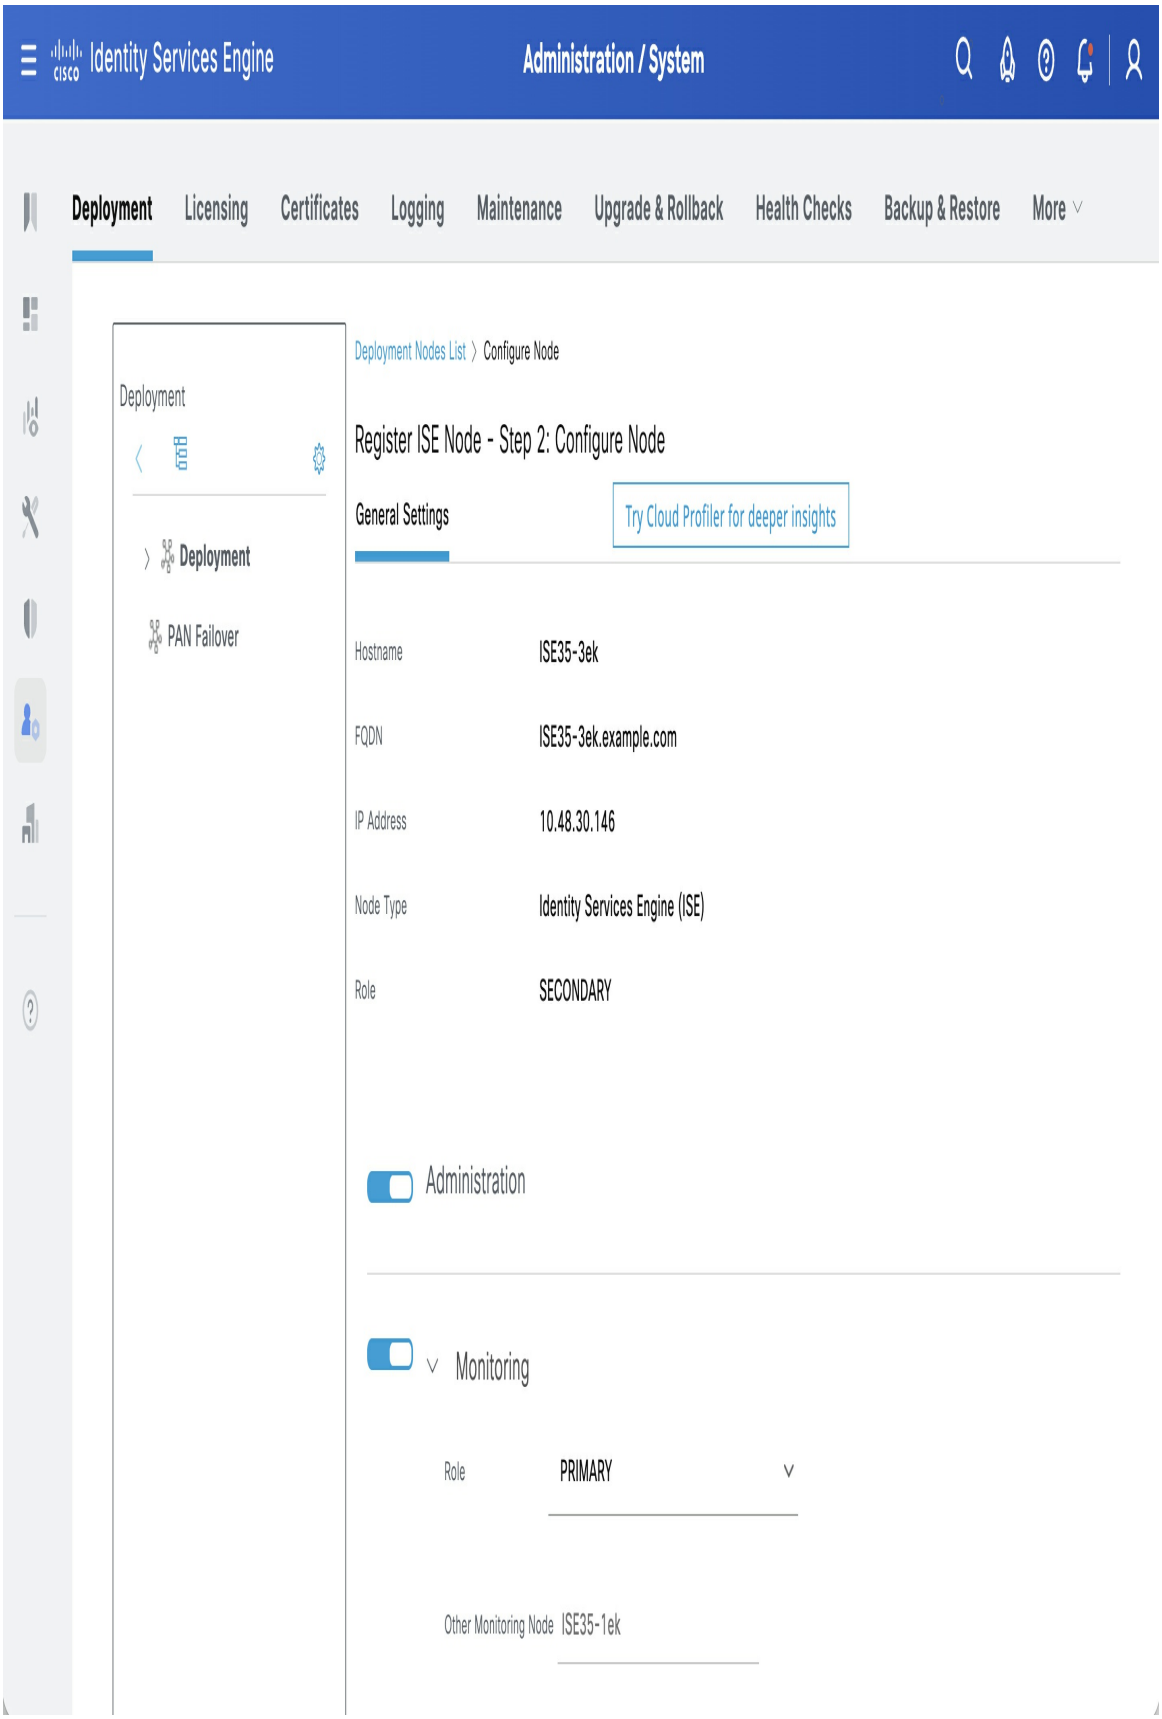


Figure 2-19 *Hostname and Credentials Menu During ISE Node*

Technet24

Registration

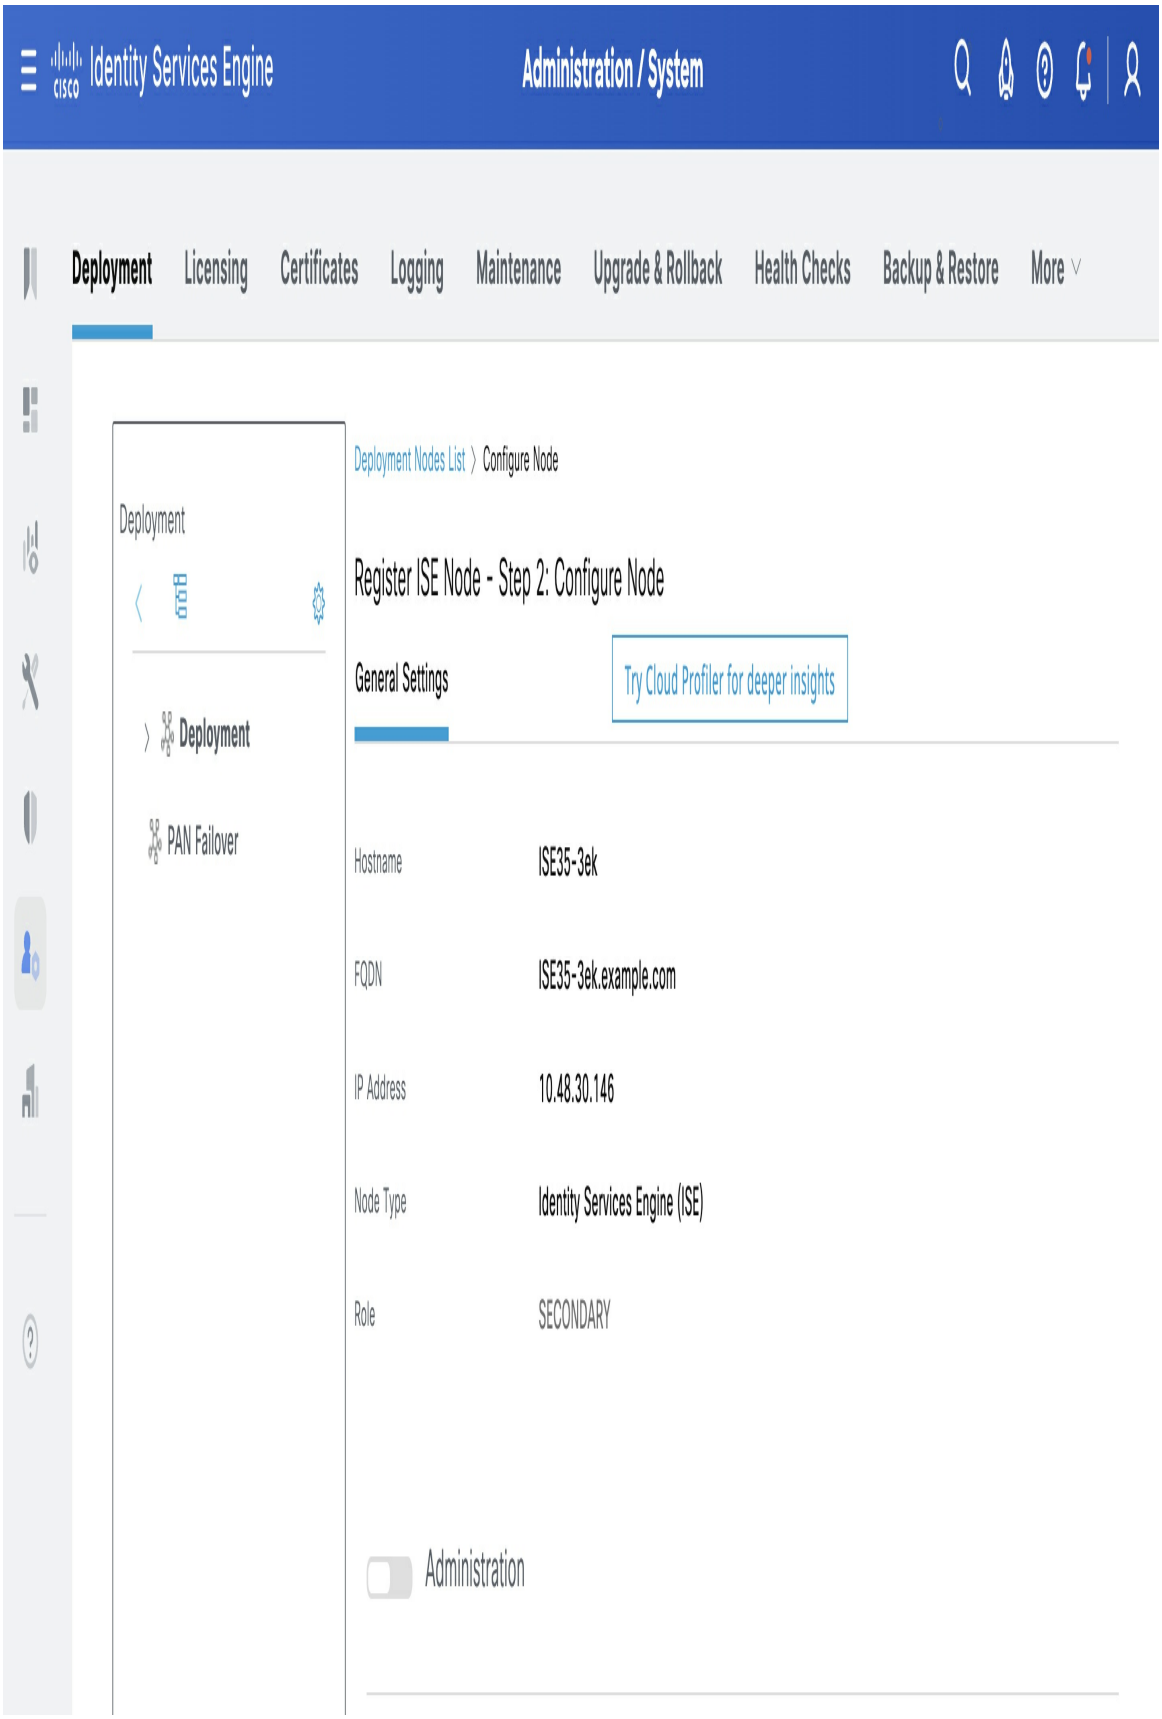
After you click **Next**, you are presented with options to select the personas for the new node you are adding to the ISE distributed deployment, as shown in Figure 2-20. You can select the roles immediately or adjust them later after the node becomes part of the deployment.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Deployment
Deployment
PAN Failover
Deployment Nodes List
Configure Node
Register ISE Node - Step 2: Configure Node
General Settings
Try Cloud Profiler for deeper insights
Hostname
ISE35-3ek
FQDN
ISE35-3ek.example.com
IP Address
10.48.30.146
Node Type
Identity Services Engine (ISE)
Role
SECONDARY
Administration
Monitoring
Role
PRIMARY
Other Monitoring Node
ISE35-1ek

Technet24

Figure 2-20 *Role Configuration Menu During ISE Node Registration*

If you are adding a secondary PAN node to the deployment, the Administration slider should be enabled. If PPAN and SPAN are already present in your deployment, disable the slider, as shown in Figure 2-21.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Deployment
Deployment
PAN Failover
Deployment Nodes List
Configure Node
Register ISE Node - Step 2: Configure Node
General Settings
Try Cloud Profiler for deeper insights
Hostname
ISE35-3ek
FQDN
ISE35-3ek.example.com
IP Address
10.48.30.146
Node Type
Identity Services Engine (ISE)
Role
SECONDARY
Administration

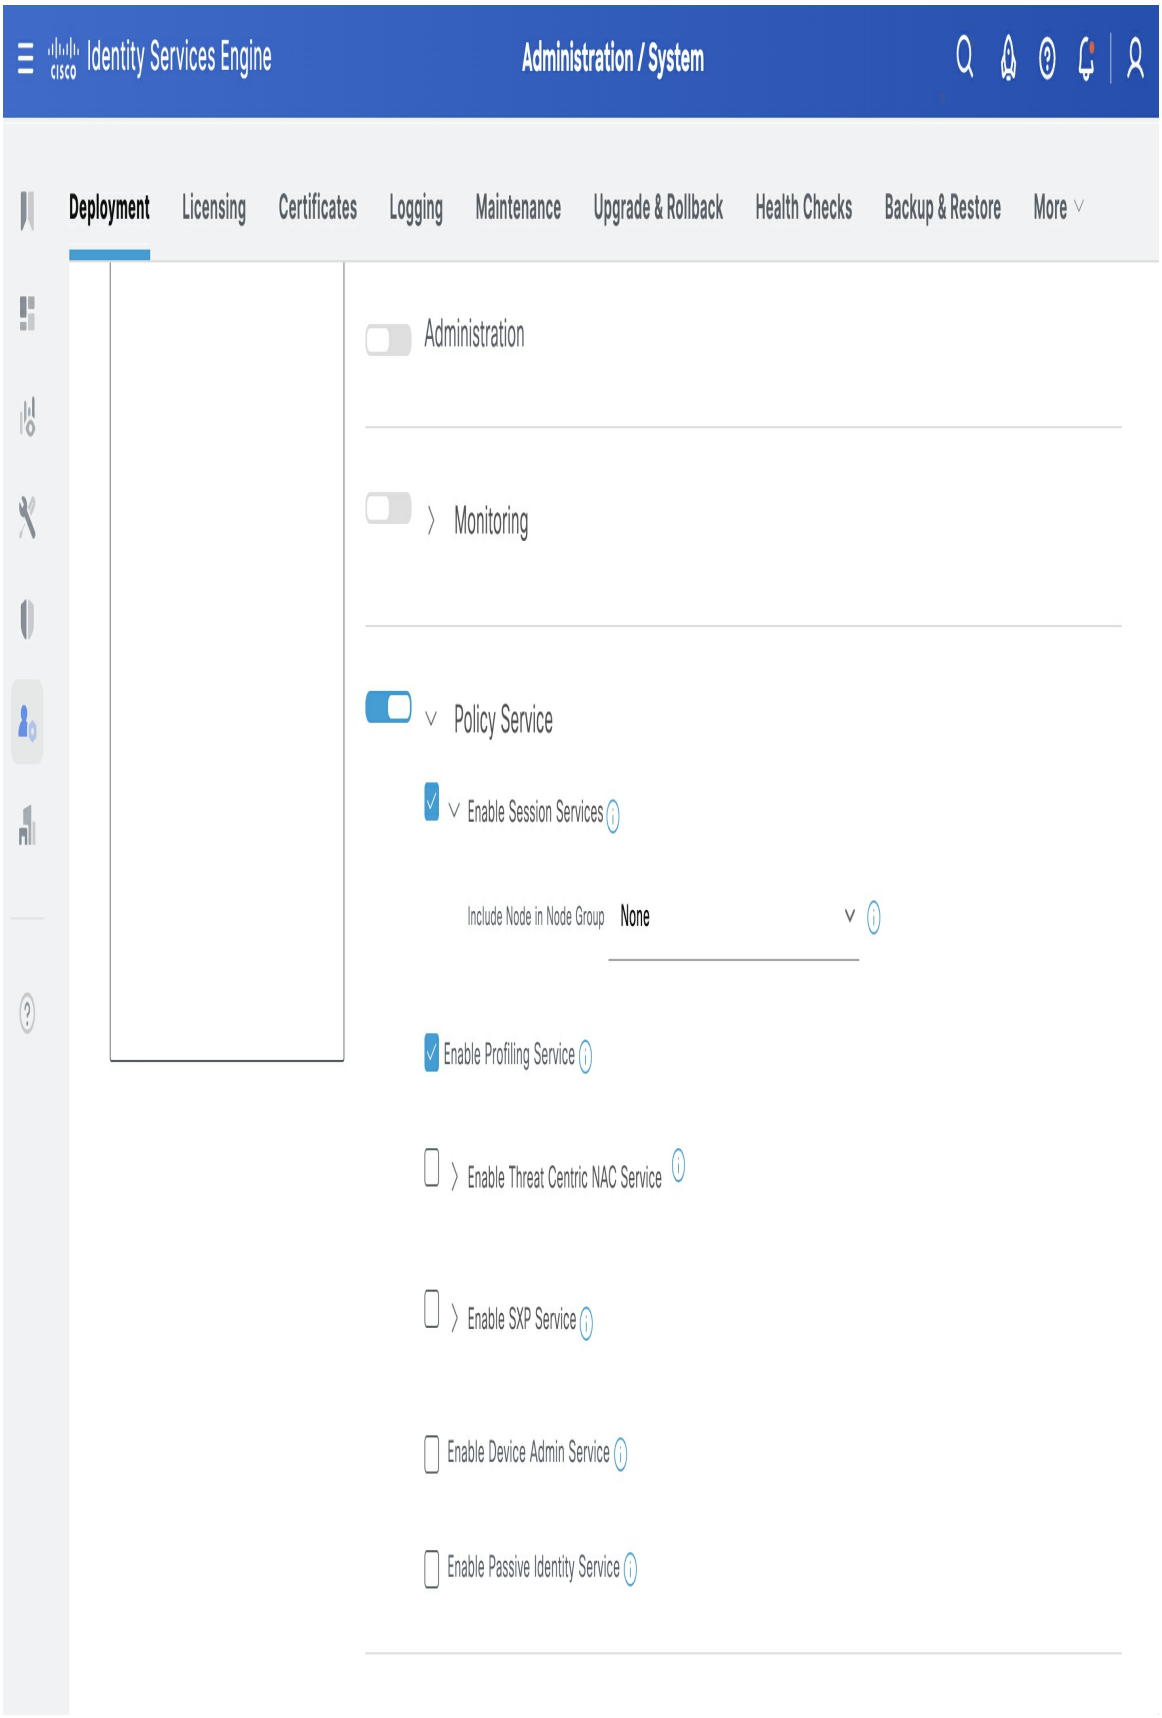
Technet24

Figure 2-21 *Administration Persona Configuration*

If you are adding a primary MnT node to the deployment, enable the slider and select **Primary** from the drop-down menu for the role.

If you are adding a secondary MnT node to the deployment, enable the slider and select **Secondary** from the drop-down menu for the role.

If you are adding a node that should not have an MnT role, disable the slider, as shown in Figure 2-22.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Administration
Monitoring
Policy Service
Enable Session Services
Include Node in Node Group
None
Enable Profiling Service
Enable Threat Centric NAC Service
Enable SXP Service
Enable Device Admin Service
Enable Passive Identity Service

Technet24

Figure 2-22 *MnT and PSN Persona Configuration*

Figure 2-22 also shows the options for configuring PSN functionality on the node. If you do not want the node to have the PSN persona, turn off the slider. If you want the node to be a PSN, turn on the slider and select the specific functions for the PSN. Figure 2-22 shows that the node being added is a PSN with only RADIUS and profiling services enabled.

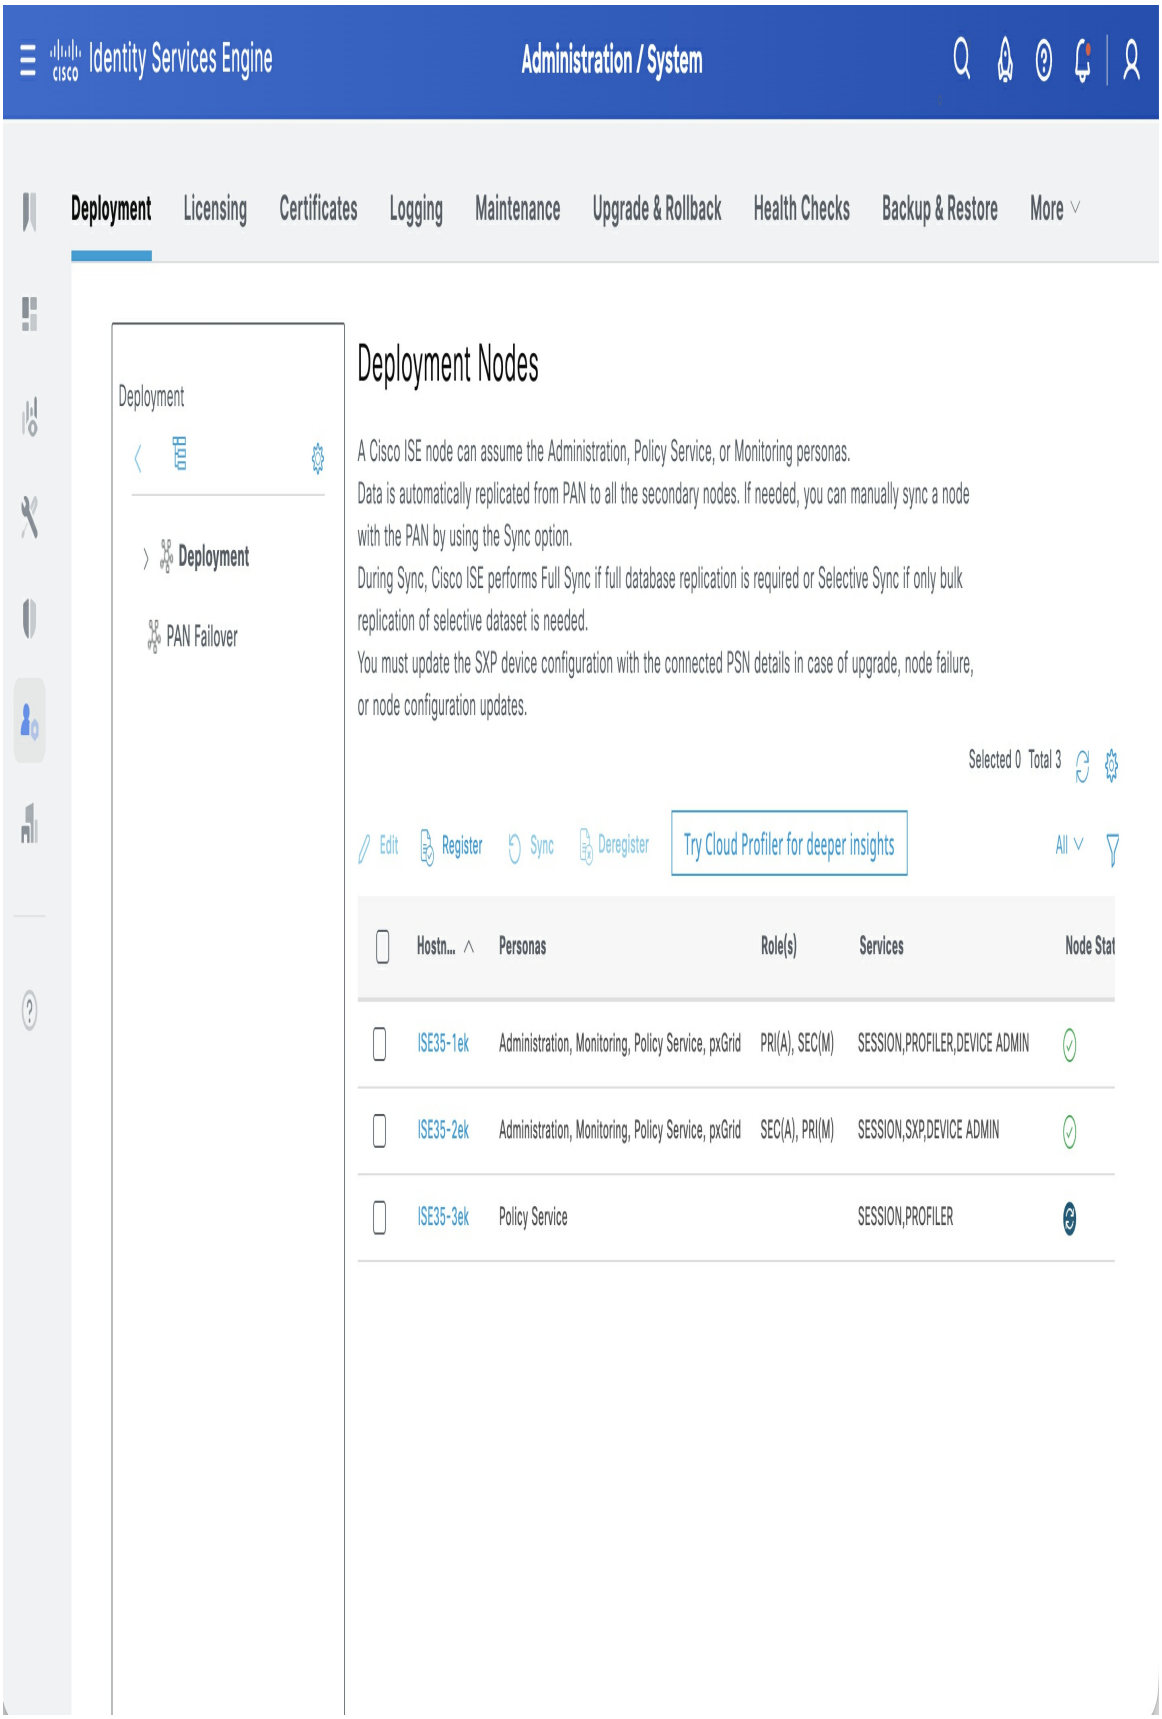
Best Practice:

Enable only the services that are required. If the PSN is not intended for device administration with TACACS+, do not enable it. If the PSN is not intended to be integrated with TC-NAC, do not enable it. You can always enable additional PSN services later if requirements change.

When the node is added, ISE restarts the services on the node you are adding. The PPAN or any other node in the deployment does not restart services during this operation.

After the node is added, ISE registration and synchronization begins. This multistep process includes database export, copying the database to the secondary node, database import on the secondary node, and replication of any missing database changes.

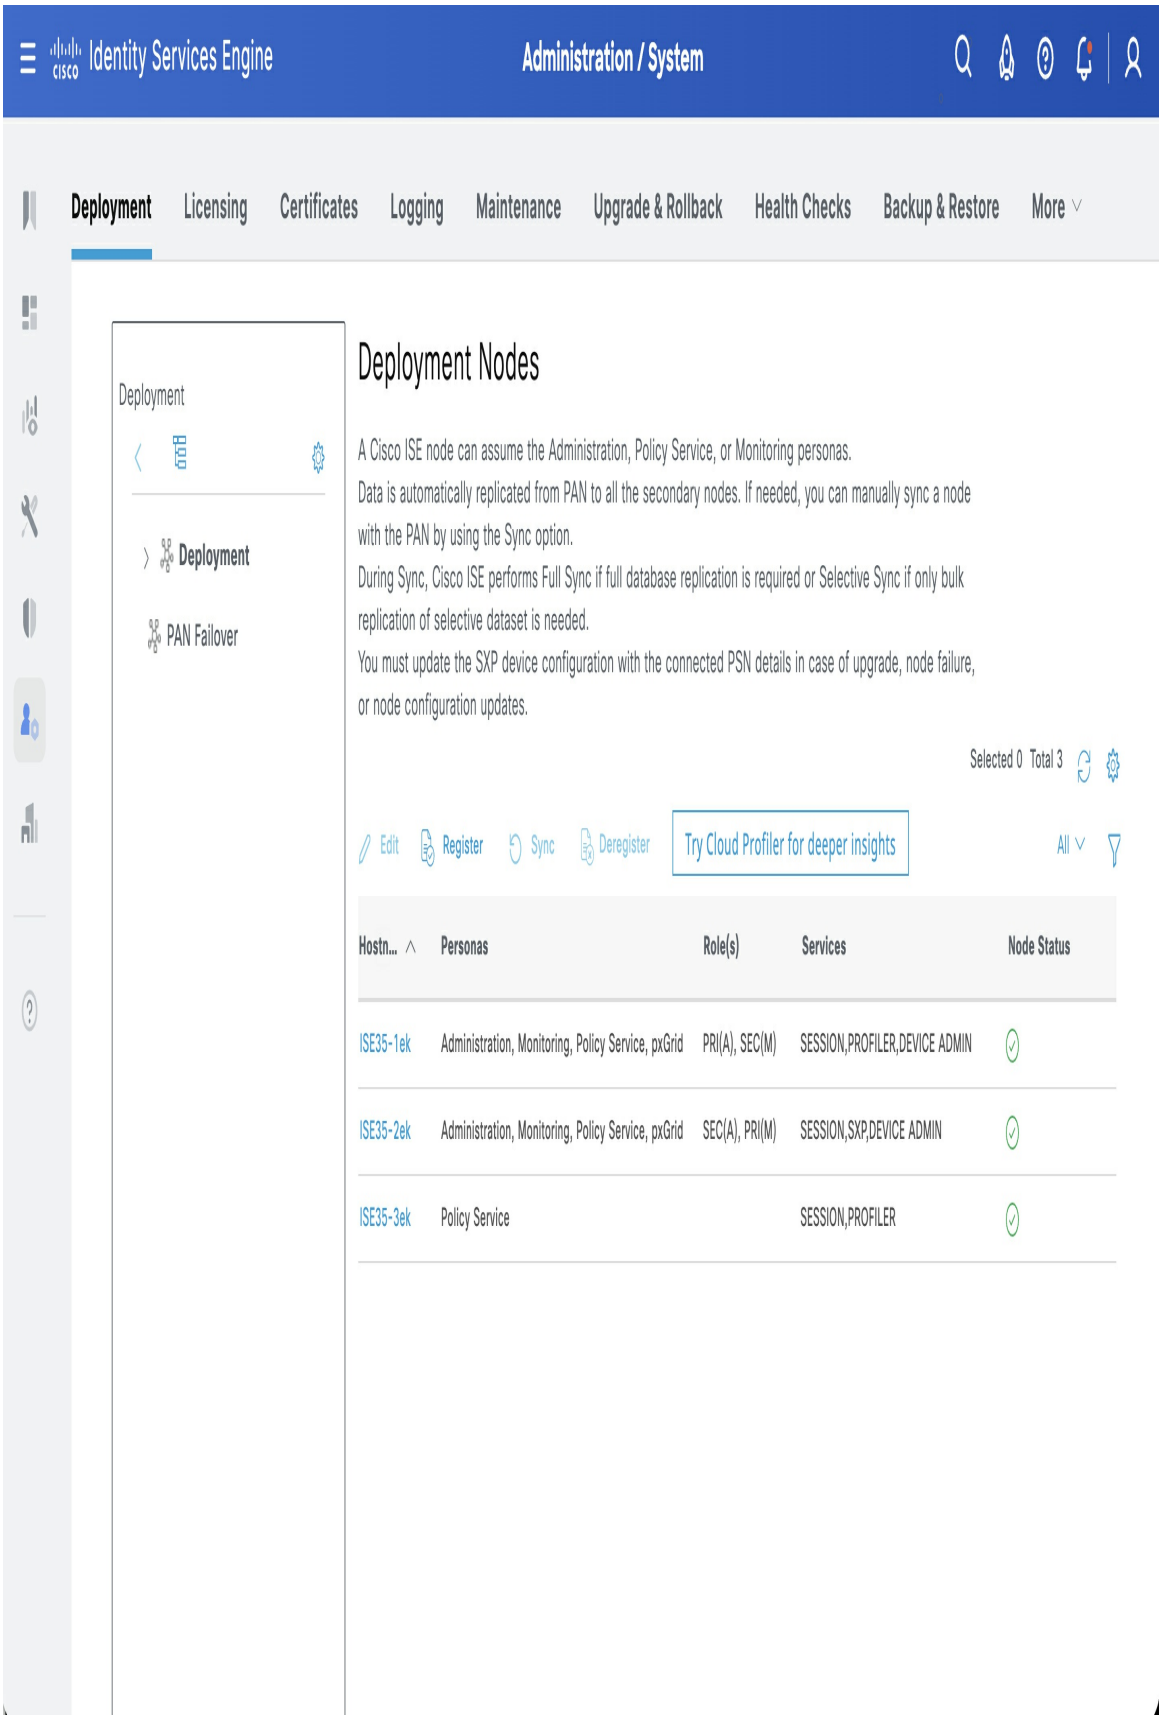
After successful registration, when you navigate to **Administration > System > Deployment** immediately after adding the node, you will see a **Registration or Sync in Progress** icon next to the secondary node, as shown in Figure 2-23.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Deployment
Deployment
PAN Failover
Deployment Nodes
A Cisco ISE node can assume the Administration, Policy Service, or Monitoring personas.
Data is automatically replicated from PAN to all the secondary nodes. If needed, you can manually sync a node with the PAN by using the Sync option.
During Sync, Cisco ISE performs Full Sync if full database replication is required or Selective Sync if only bulk replication of selective dataset is needed.
You must update the SXP device configuration with the connected PSN details in case of upgrade, node failure, or node configuration updates.
Selected 0 Total 3
Edit
Register
Sync
Deregister
Try Cloud Profiler for deeper insights
All
Hostn...
Personas
Role(s)
Services
Node Stat
ISE35-1ek
Administration, Monitoring, Policy Service, pxGrid
PRI(A), SEC(M)
SESSION,PROFILER,DEVICE ADMIN
ISE35-2ek
Administration, Monitoring, Policy Service, pxGrid
SEC(A), PRI(M)
SESSION,SXP,DEVICE ADMIN
ISE35-3ek
Policy Service
SESSION,PROFILER

Technet24

Figure 2-23 *ISE Registration or Sync in Progress*

When registration and replication are complete, the newly added node should transition to the connected state, as shown in Figure 2-24. The process may take around 20 minutes, but the exact duration depends on the size of the database, the node specifications, the personas, and the services enabled.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Deployment
Deployment
PAN Failover
Deployment Nodes
A Cisco ISE node can assume the Administration, Policy Service, or Monitoring personas.
Data is automatically replicated from PAN to all the secondary nodes. If needed, you can manually sync a node with the PAN by using the Sync option.
During Sync, Cisco ISE performs Full Sync if full database replication is required or Selective Sync if only bulk replication of selective dataset is needed.
You must update the SXP device configuration with the connected PSN details in case of upgrade, node failure, or node configuration updates.
Selected 0 Total 3
Edit
Register
Sync
Deregister
Try Cloud Profiler for deeper insights
All
Hostn...
Personas
Role(s)
Services
Node Status
ISE35-1ek
Administration, Monitoring, Policy Service, pxGrid
PRI(A), SEC(M)
SESSION,PROFILER,DEVICE ADMIN
ISE35-2ek
Administration, Monitoring, Policy Service, pxGrid
SEC(A), PRI(M)
SESSION,SXP,DEVICE ADMIN
ISE35-3ek
Policy Service
SESSION,PROFILER

Technet24

Figure 2-24 *ISE Registration or Sync Completed*

Licensing

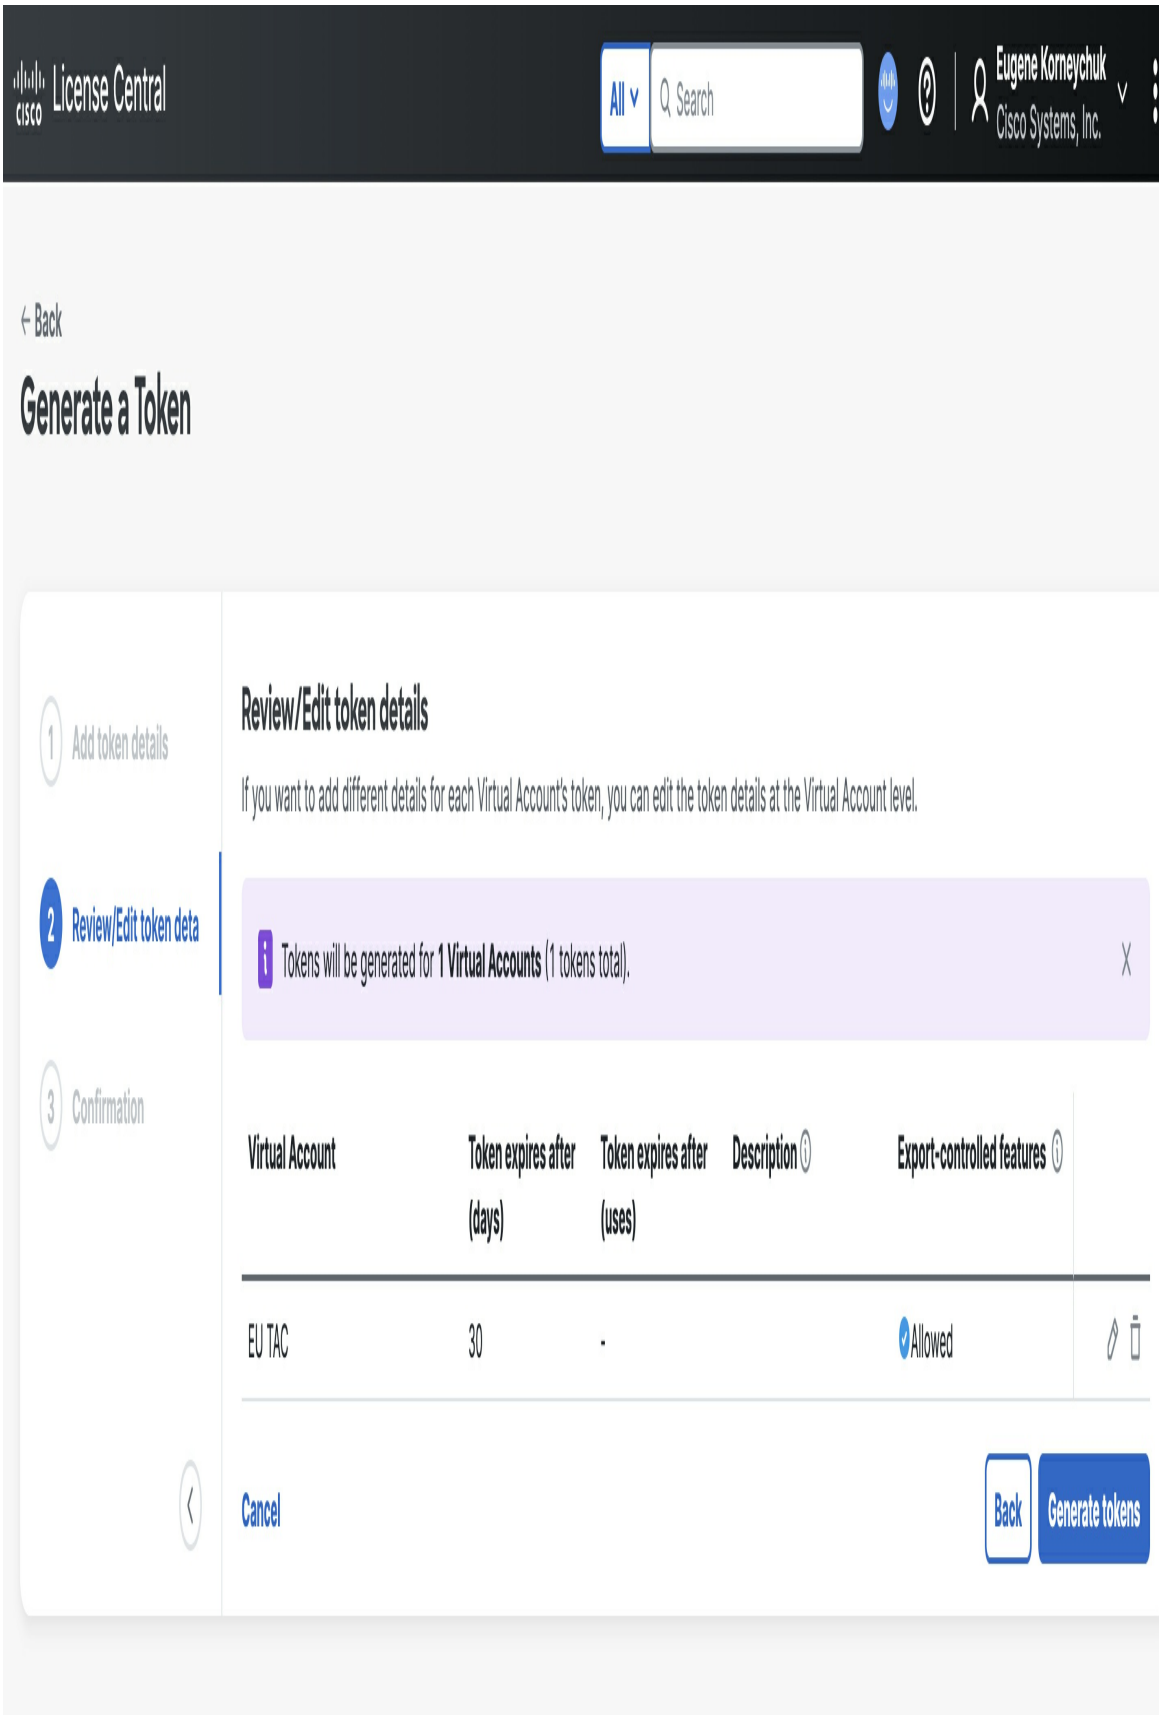
ISE licensing is one of the most important components to consider when planning and designing your ISE cluster. By default, ISE comes with a 90-day evaluation license, allowing you to test almost any feature on Cisco ISE before installing a permanent license.

There are three common ways to license ISE. If your ISE is connected to the Internet, you can use Smart Licensing Registration with Cisco Smart Software Manager (CSSM). ISE supports both direct connections and connections via proxy. This is the recommended method.

For air-gapped ISE deployments, you can use Smart Software Manager On-Premises (SSM On-Prem). This dedicated server is installed on premises, which eliminates the need for a direct connection from the ISE deployment to the Internet. It is suitable for air-gapped environments where direct Internet connectivity is not desired.

Specific License Reservation (SLR) is another method that allows customers to manually allocate ISE licenses. This method is recommended for highly secured, air-gapped networks.

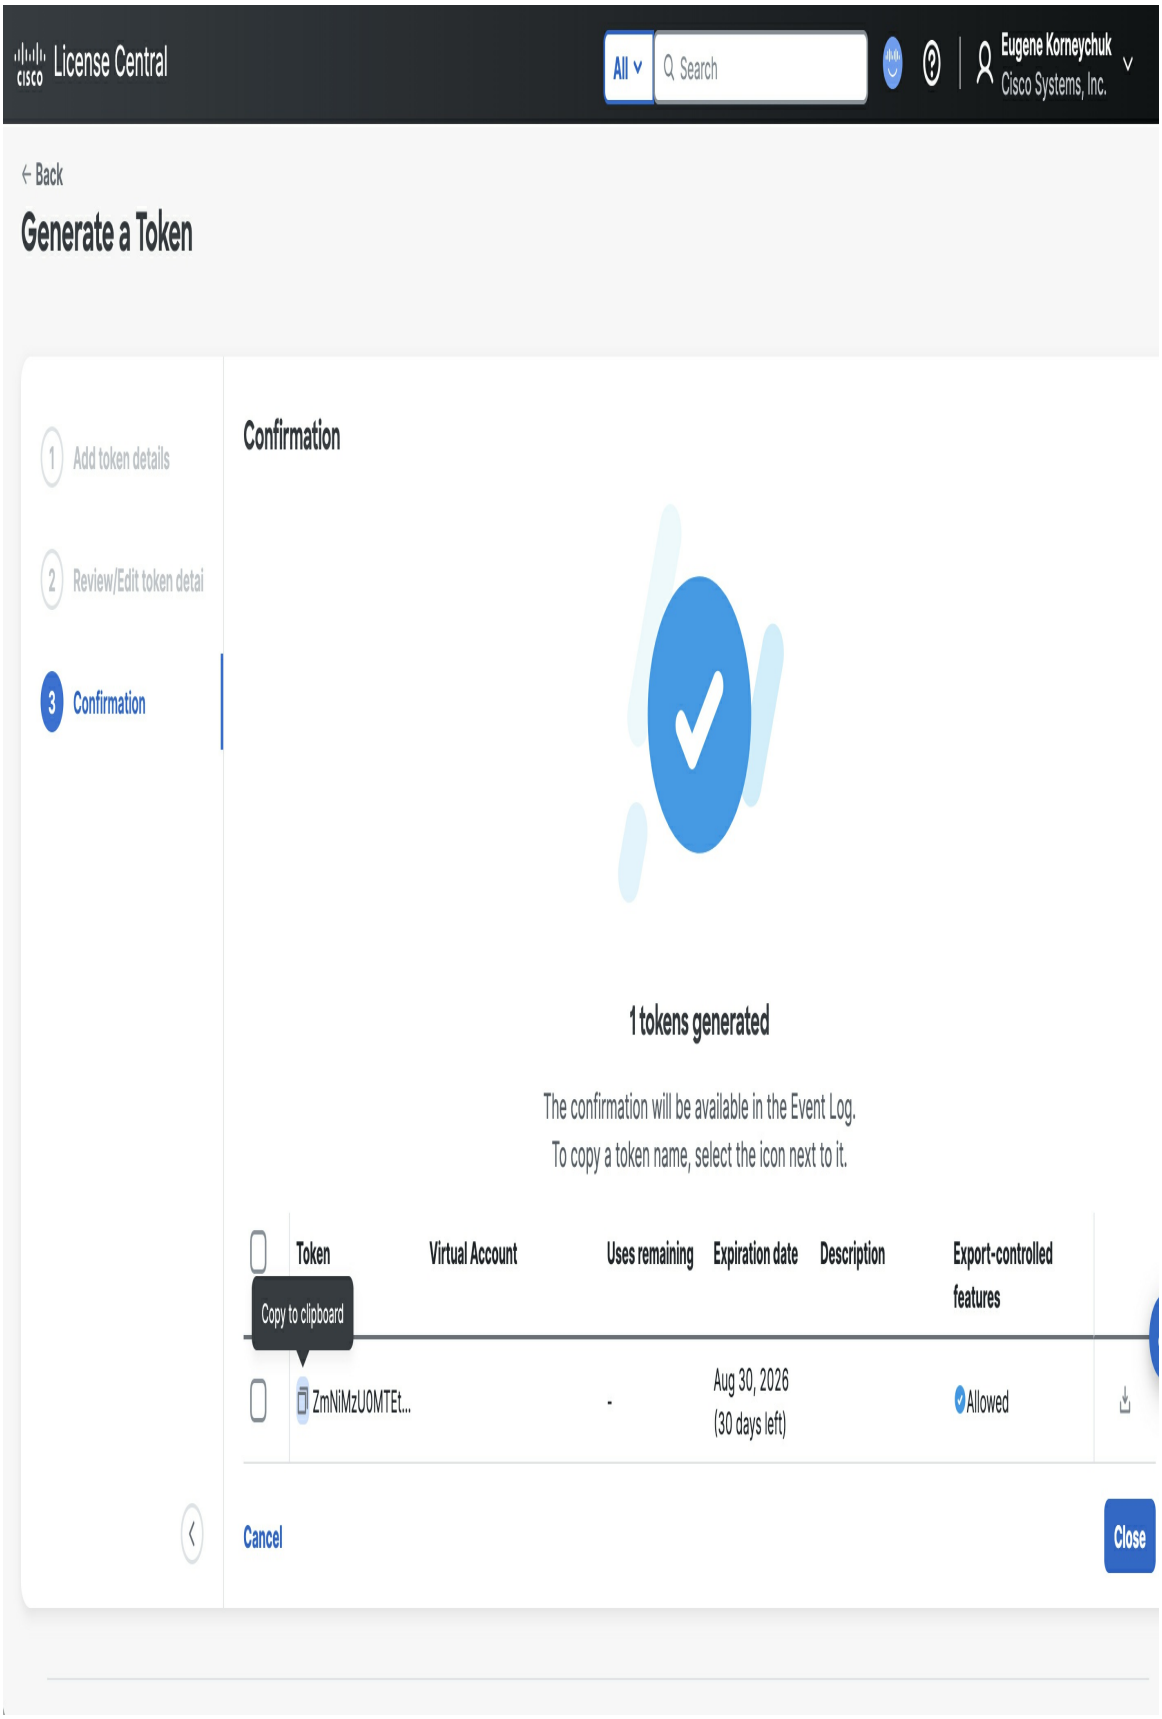
To register ISE with Smart Licensing using CSSM, navigate to **Licensing Central** at https://software.cisco.com/clc/overview. Under **Devices**, click **Generate a Token**. Select a virtual account and click **Generate Tokens** on the next page, as shown in Figure 2-25.


License Central
All
Search
Eugene Korneychuk
Cisco Systems, Inc.
Back
Generate a Token
1 Add token details
2 Review/Edit token deta
3 Confirmation
Review/Edit token details
If you want to add different details for each Virtual Account's token, you can edit the token details at the Virtual Account level.
Tokens will be generated for 1 Virtual Accounts (1 tokens total).
Virtual Account
Token expires after (days)
Token expires after (uses)
Description
Export-controlled features
EU TAC
30
-
Allowed
Cancel
Back
Generate tokens

Technet24

Figure 2-25 *Token Generation on Licensing Central*

Copy the value of the token, as shown in Figure 2-26.

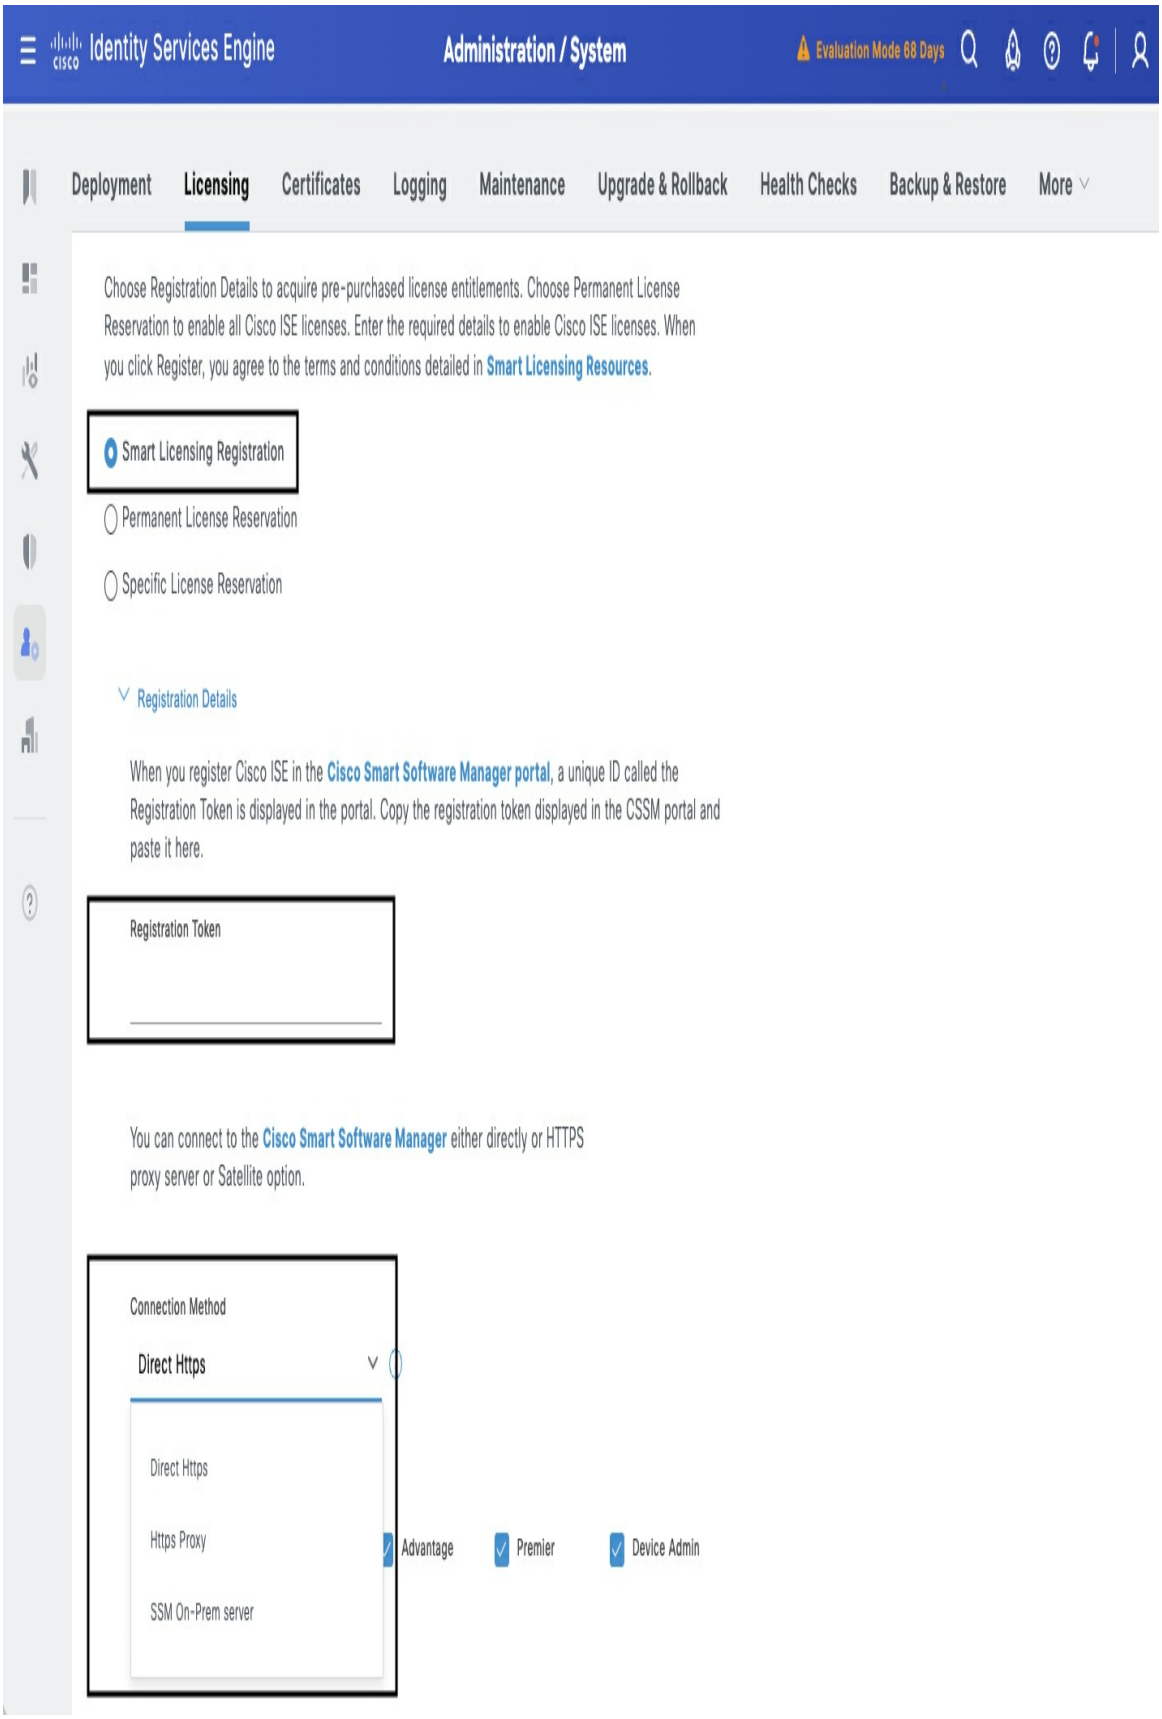

License Central
All
Search
Eugene Korneychuk
Cisco Systems, Inc.
Back
Generate a Token
1 Add token details
2 Review/Edit token detai
3 Confirmation
Confirmation
1 tokens generated
The confirmation will be available in the Event Log.
To copy a token name, select the icon next to it.
Token
Virtual Account
Uses remaining
Expiration date
Description
Export-controlled features
Copy to clipboard
ZmNiMzU0MTEt...
-
Aug 30, 2026
(30 days left)
Allowed
Cancel
Close

Technet24

Figure 2-26 *Token Copy from CSSM*

Navigate to **ISE > Administration > System > Licensing** and select the **Smart Licensing Registration** radio button, as shown in Figure 2-27. Enter the registration token value you copied in the previous step and select the appropriate connection method for your deployment. When you're done, click the **Register** button.

No service restart is associated with license registration.


Identity Services Engine
Administration / System
Evaluation Mode 68 Days
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Choose Registration Details to acquire pre-purchased license entitlements. Choose Permanent License Reservation to enable all Cisco ISE licenses. Enter the required details to enable Cisco ISE licenses. When you click Register, you agree to the terms and conditions detailed in Smart Licensing Resources.
Smart Licensing Registration
Permanent License Reservation
Specific License Reservation
Registration Details
When you register Cisco ISE in the Cisco Smart Software Manager portal, a unique ID called the Registration Token is displayed in the portal. Copy the registration token displayed in the CSSM portal and paste it here.
Registration Token
You can connect to the Cisco Smart Software Manager either directly or HTTPS proxy server or Satellite option.
Connection Method
Direct Https
Direct Https
Https Proxy
SSM On-Prem server
Advantage
Premier
Device Admin

Technet24

Figure 2-27 *ISE Smart Licensing Registration*

The example in Figure 2-28 shows successful license registration.

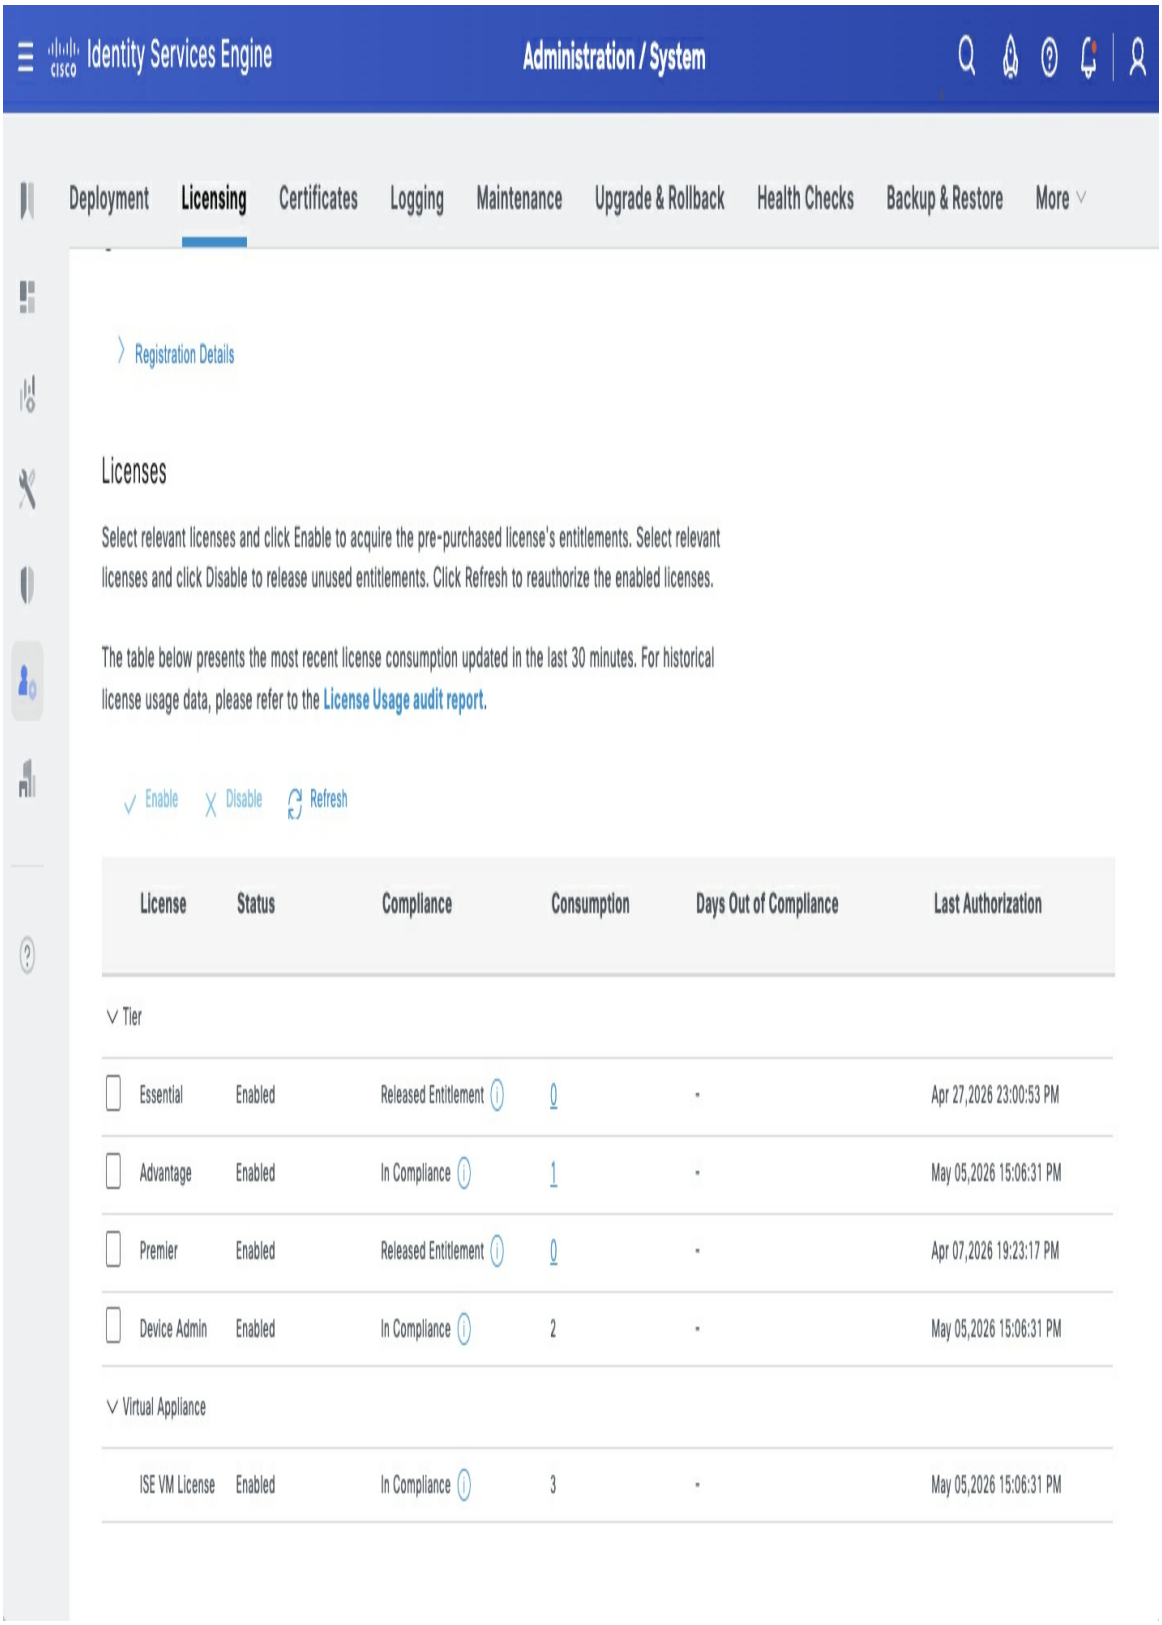


Figure 2-28 *Successful License Registration*

Technet24

There are three types of licenses with Cisco ISE:

- **Subscription license:** Includes Essential, Advantage, and Premium licenses.
- **Virtual machine license:** Required for virtual machines or cloud instance ISE servers.
- **Device admin license:** Required per PSN to perform device administration with TACACS+.

For CSSM registration, the number of virtual machine licenses should match the number of virtual machines and cloud instances in the deployment.

The number of device admin licenses should match the number of PSNs configured for device administration with TACACS+.

The number of subscription licenses should be based on the maximum number of concurrent active sessions you expect in the ISE deployment at any point in time. Subscription licenses use a nested model, meaning a Premium license includes all of the capabilities of the Advantage and Essential license tiers. Only one license is consumed by an active endpoint.

Table 2-9 shows which features are included within each license tier.

Table 2-9 *Feature to Subscription License Mapping*

Feature	Essentials	Advantage	Primer
AAA and 802.1X	✓	✓	✓
Guest (Hotspot, Self-Reg, Sponsored)	✓	✓	✓
Easy Connect (Passive ID)	✓	✓	✓
Profiling	X	✓	✓
BYOD	X	✓	✓
Context In or Out Using pxGrid, pxGrid Cloud, and pxGrid Direct	X	✓	✓
User-Defined Network	X	✓	✓
TrustSec (Group-Based Policy)	X	✓	✓
Endpoint Analytics Visibility and Enforcement	X	✓	✓
Rapid Threat Containment (Adaptive Network Control)	X	✓	✓
Posture	X	X	✓
Mobile Device Management (MDM) Compliance	X	X	✓
Threat-Centric Network Access Control (TC-NAC)	X	X	✓

Technet24

Best Practice:

Always refer to the Cisco ISE Licensing Guide for the most up-to-date information about license tiers; see https://cs.co/ise-licensing.

Subscription-based license consumption is measured through RADIUS sessions. When a RADIUS session ends, the license is released.

If more licenses than purchased are used, or when a license expires, ISE licenses go out of compliance. If the deployment is out of compliance for 30 days within a 60-day window, you lose administrative control of Cisco ISE until the license is purchased and activated. During this period, you will have limited visibility, management, and access to relevant features. Authentications and previously configured policies continue to operate as they did before expiration.

Summary

ISE deployment planning is the foundation for successful ISE operations. You need to ensure that proper sizing is selected, ISE and network devices are configured for high availability, and appropriate licensing is in place. Even if scaling is required later, ISE provides flexibility to add nodes to the deployment, which increases the number of supported active sessions and allows more users to connect to the network.

Chapter 3. Implementing ISE in Hybrid Infrastructure

Introduction

In today's world, cloud infrastructure has become commonplace. Over the years, organizations have migrated a significant number of workloads from on-premises environments to the cloud. Some companies now operate IT assets exclusively in the cloud. You can deploy Cisco Identity Services Engine as a native virtual machine in the cloud as well. This chapter guides you through supported ISE cloud deployments, discusses specific cloud providers, and explains the limitations and best practices for successful ISE implementations.

ISE Cloud Deployments

Cisco ISE is supported natively with

- Amazon Web Services (AWS) starting with release 3.1 and later
- Microsoft Azure Cloud Services (Azure) starting with release 3.2 and later
- Oracle Cloud Infrastructure (OCI) starting with release 3.2 and later

You can fully deploy ISE in the cloud, with all nodes hosted by a specific cloud provider. Alternatively, you can implement a hybrid

Technet24

deployment, placing some nodes on premises and others in the cloud.

You may combine all form factors within a single ISE deployment. ISE cloud deployments do not support SSH access with a username and password. Instead, you should access the cloud instance through SSH using a private/public key pair. Be sure to store the public key on ISE and ensure the private key remains secure on the machine you use to connect via SSH.

The default username for Cisco ISE instances is **iseadmin**. Even if you configure a different username through user data, ISE still creates the **iseadmin** account. You can create new admin accounts after successfully deploying ISE.

ISE deployments in the cloud rely on the bring your own license (BYOL) model. ISE nodes in the cloud require the same license types as ISE virtual machines. You must manage two primary licensing requirements:

- **Virtual machine licenses:** You require one per ISE node in the cloud.
- **Feature licenses:** You select these based on the implemented use case.

ISE on AWS

You should note that ISE, when deployed with cloud providers, does not function as a native service solution. The system operates as a virtual machine. In AWS, you can deploy ISE as an EC2 instance.

Overview of Cisco ISE on AWS

You can deploy ISE on AWS using two methods:

- ISE installation via CloudFormation Template (CFT)
- ISE installation via Amazon Machine Images (AMI)

Best Practice:

Cisco recommends installing ISE with a CFT for manual deployments. If you use Infrastructure as Code (IaC) tools such as Terraform or Ansible, choose AMI installation.

CloudFormation Template installation guides you through ISE setup steps using prepared prompts included in the template. This approach helps reduce errors. In contrast, AMI installation relies on a user data file and does not provide syntax validation during deployment, making the process more error prone.

While Chapter 2, "ISE Deployment Principles," covered the selection of ISE nodes based on concurrent active sessions, ISE instances in AWS also incorporate the concept of intended usage.

ISE supports only specific AWS instance types for resource specifications. Table 3-1 describes production instance types and their intended purpose.

Table 3-1 *AWS ISE Instances and Their Intended Usage*

Technet24

Instance Type	Intended Use	CPU Cores	RAM (in GB)
c5.4xlarge	PSN	16	32
c5.9xlarge		36	72
c7i.4xlarge*		16	32
m5.2xlarge	PAN or MnT or both	8	32
m5.8xlarge		32	128
m5.16xlarge		64	256
m7i.2xlarge*		8	32
m7i.8xlarge*		32	128
m7i.16xlarge*		64	256

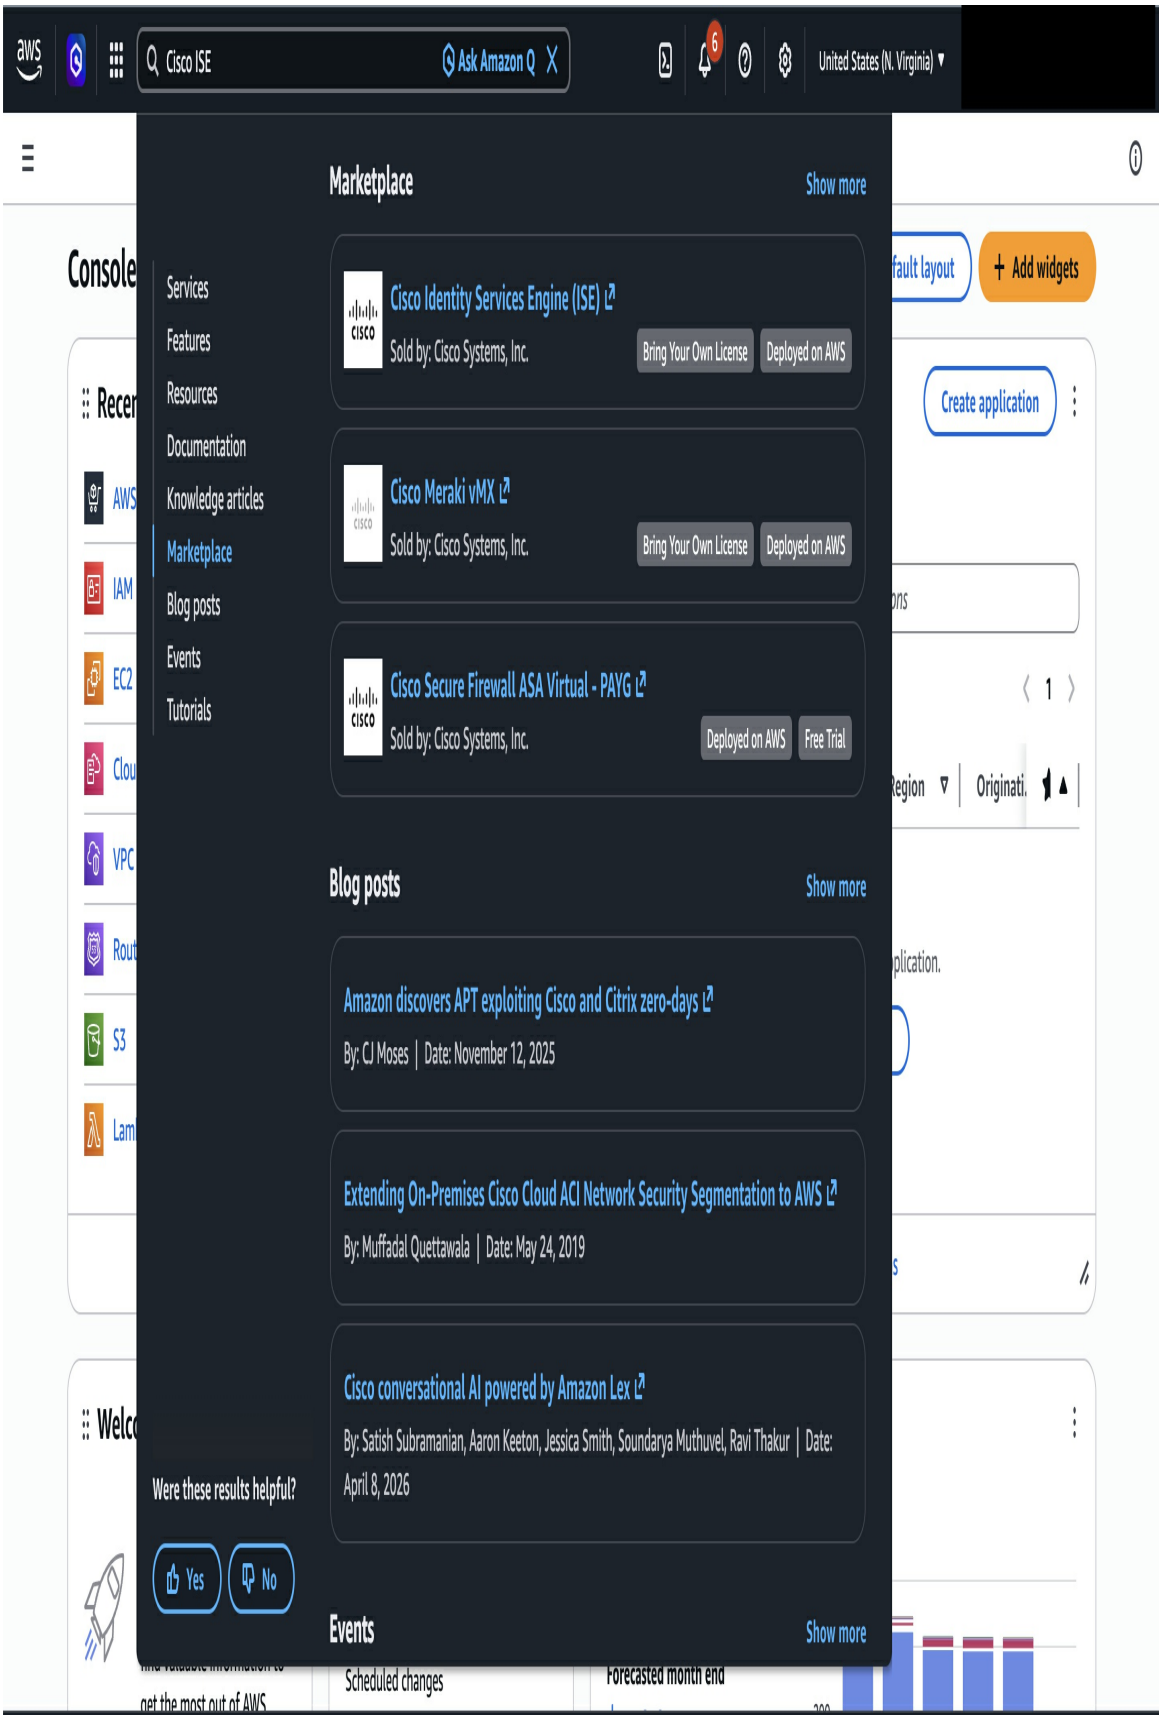
*Cisco ISE releases 3.4 patch 4, 3.5, and later support M7i and C7i instance types on AWS.

Note:

m5.2xlarge and m7i.2xlarge instances must be used as Extra Small PSN only.

ISE Installation via CFT

To install ISE from the Marketplace using a CloudFormation Template, search for Cisco ISE in the AWS Console Search window and select **Marketplace** from the left side of the screen. Click **Cisco Identity Services Engine (ISE)**, as shown in Figure 3-1.


Cisco ISE
Ask Amazon Q
United States (N. Virginia)
Services
Features
Resources
Documentation
Knowledge articles
Marketplace
Blog posts
Events
Tutorials
Were these results helpful?
Yes
No
Marketplace
Show more
Cisco Identity Services Engine (ISE)
Sold by: Cisco Systems, Inc.
Bring Your Own License
Deployed on AWS
Cisco Meraki vMX
Sold by: Cisco Systems, Inc.
Bring Your Own License
Deployed on AWS
Cisco Secure Firewall ASA Virtual - PAYG
Sold by: Cisco Systems, Inc.
Deployed on AWS
Free Trial
Blog posts
Show more
Amazon discovers APT exploiting Cisco and Citrix zero-days
By: CJ Moses | Date: November 12, 2025
Extending On-Premises Cisco Cloud ACI Network Security Segmentation to AWS
By: Muffadal Quettawala | Date: May 24, 2019
Cisco conversational AI powered by Amazon Lex
By: Satish Subramanian, Aaron Keeton, Jessica Smith, Soundarya Muthuvel, Ravi Thakur | Date: April 8, 2026
Events
Show more
Console
Create application
Add widgets
Scheduled changes
Forecasted month end

Technet24

Figure 3-1 *ISE on AWS Marketplace*

Click **View Purchase Options** (as shown in Figure 3-2) and then click **Subscribe** on the next screen.

Cisco Identity Services Engine (ISE) Info

Sold by: Cisco Systems, Inc.

View purchase options

Request private offer

Deployed on AWS

Cisco ISE on AWS provides secure network access control for IoT, BYOD, and corporate owned endpoints. Cisco ISE enables you to easily segment network access for employees, contractors, and guests across wired, wireles...

Show more

4.4 (97)

Overview

Try agent mode Create proposal Ask question

Cisco Identity Services Engine (ISE) on AWS enables Network Access Control (NAC) service workloads to be deployed and managed from the cloud while ensuring the flexibility required to meet each organizations unique cloud strategy. With Cisco ISE on AWS, you can unify the policy management of your organization for endpoint access control and network device administration. Cisco ISE is equipped with rich APIs to automate policy and lifecycle management, bringing ease of deployment and automation to the forefront of your NAC operations.

For GovCloud users launching directly from CloudFormation, please retrieve the Cisco Identity Services Engine (ISE) - GovCloud deployment template from the listing detail page linked below.
https://aws.amazon.com/marketplace/pp/prodview-uvsybra7r3iug

Highlights

- Gain visibility with context and control: Know who, what, where, and how endpoints and devices are connecting to your network to ensure compliance and limit risk, with or without the use of agents.
- Extend zero trust to contain threats: Software-Defined Network segmentation shrinks the attack surface, limits the spread of ransomware, and enables rapid threat containment.
- Accelerate the value of existing solutions: Integrate with other Cisco and third-party solutions to bring an active arm of protection into passive security solutions and increase your return on investment (ROI).

Technet24

Figure 3-2 *ISE on AWS Marketplace Subscription*

When you are subscribed, click **Launch Your Software**, as shown in Figure 3-3.

Subscribe to Cisco Identity Services Engine (ISE)

To create a subscription, review the pricing information and accept the terms for this software.

Existing subscription detected

You already have a subscription for this product.

View Subscription

Product details

Product name	Product ID	Offered by
Cisco Identity Services Engine (ISE)	bedef662-aba4-427e-b523-7c93cd50111c	Cisco Systems, Inc.

Deployed on AWS

Offer summary

You can download the offer, including its terms and conditions, before making a purchase. If configuration is required, configure the offer before downloading.

Offer ID	Offer type	Offer availability
basttrzv6xwc4yn2uup6bh730	Public offer	Purchased

Agreement ID	Date of purchase	License ID
agmt-25tx59wn47si4mdl69ln7ej9k	Aug 23, 2021	

You've already accepted this offer

Your AWS Marketplace agreement was created. You can launch your software or Manage subscriptions.

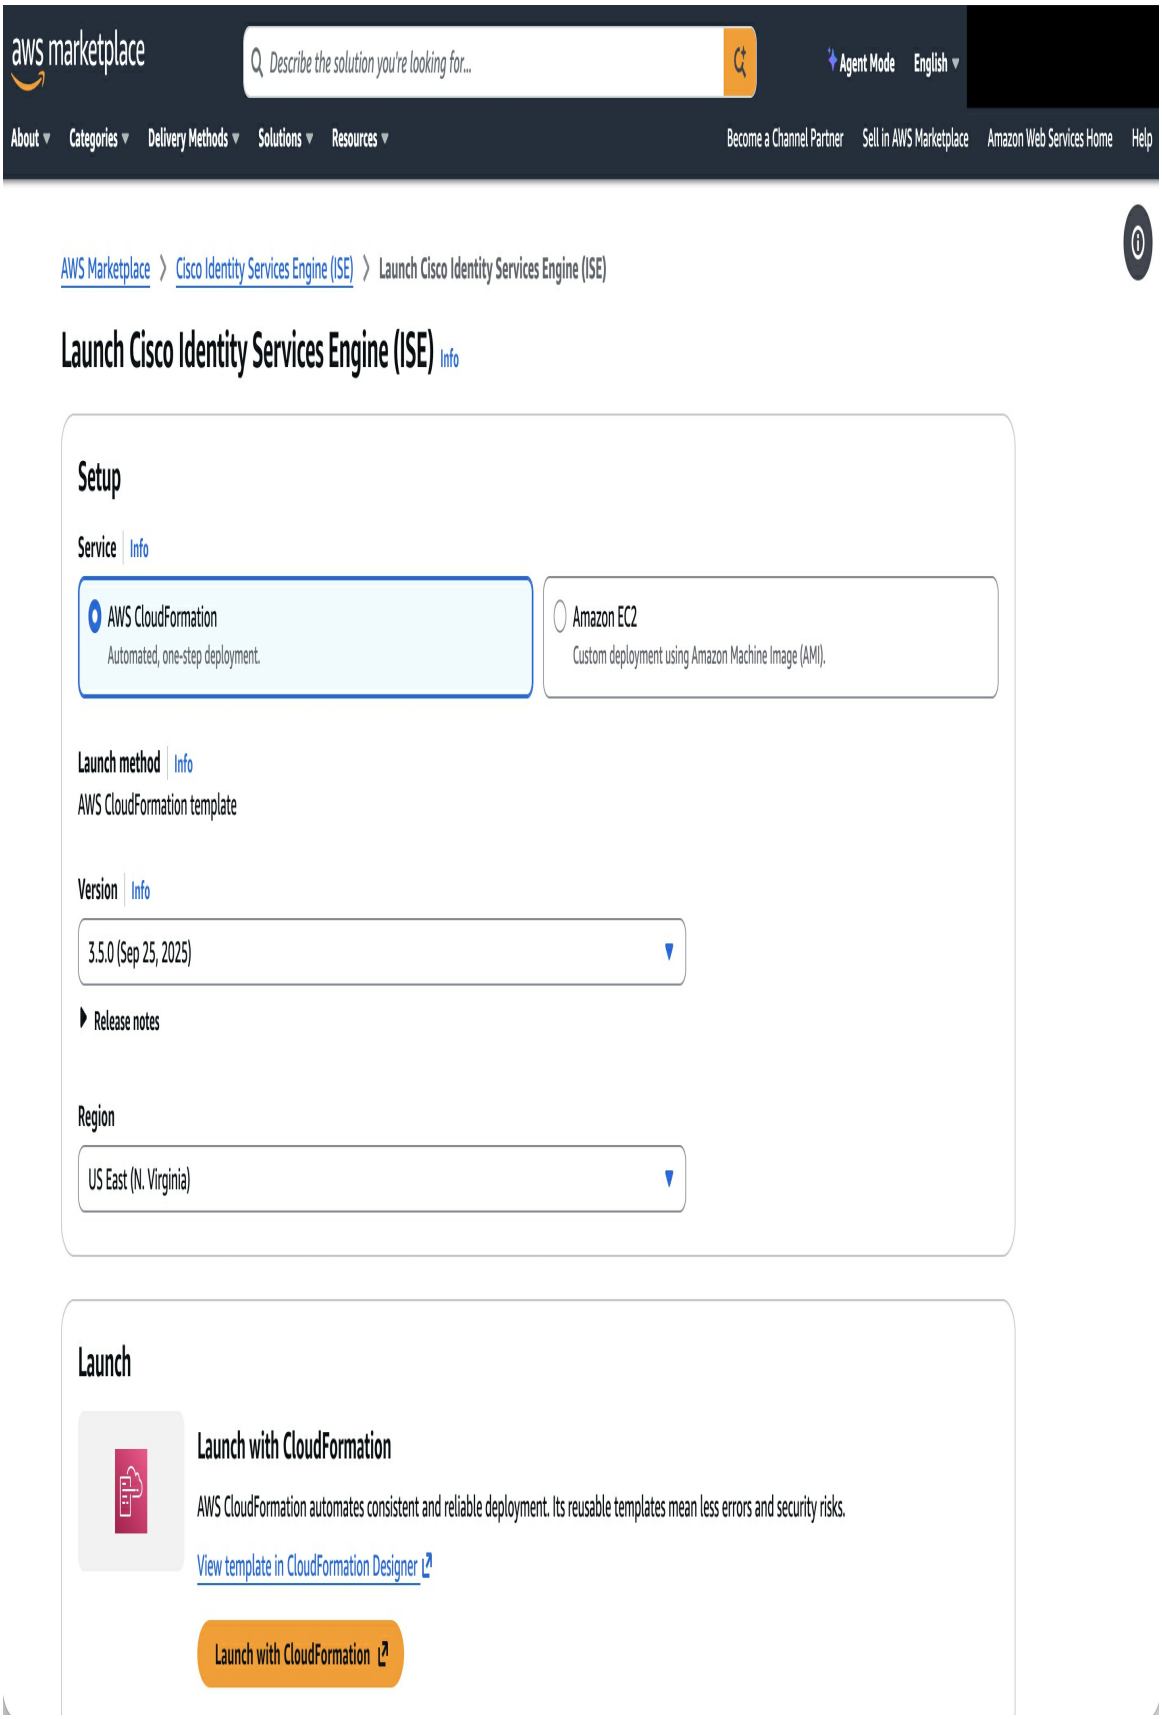
Launch your software

Pricing details

Technet24

Figure 3-3 *AWS Launch Your Software Menu*

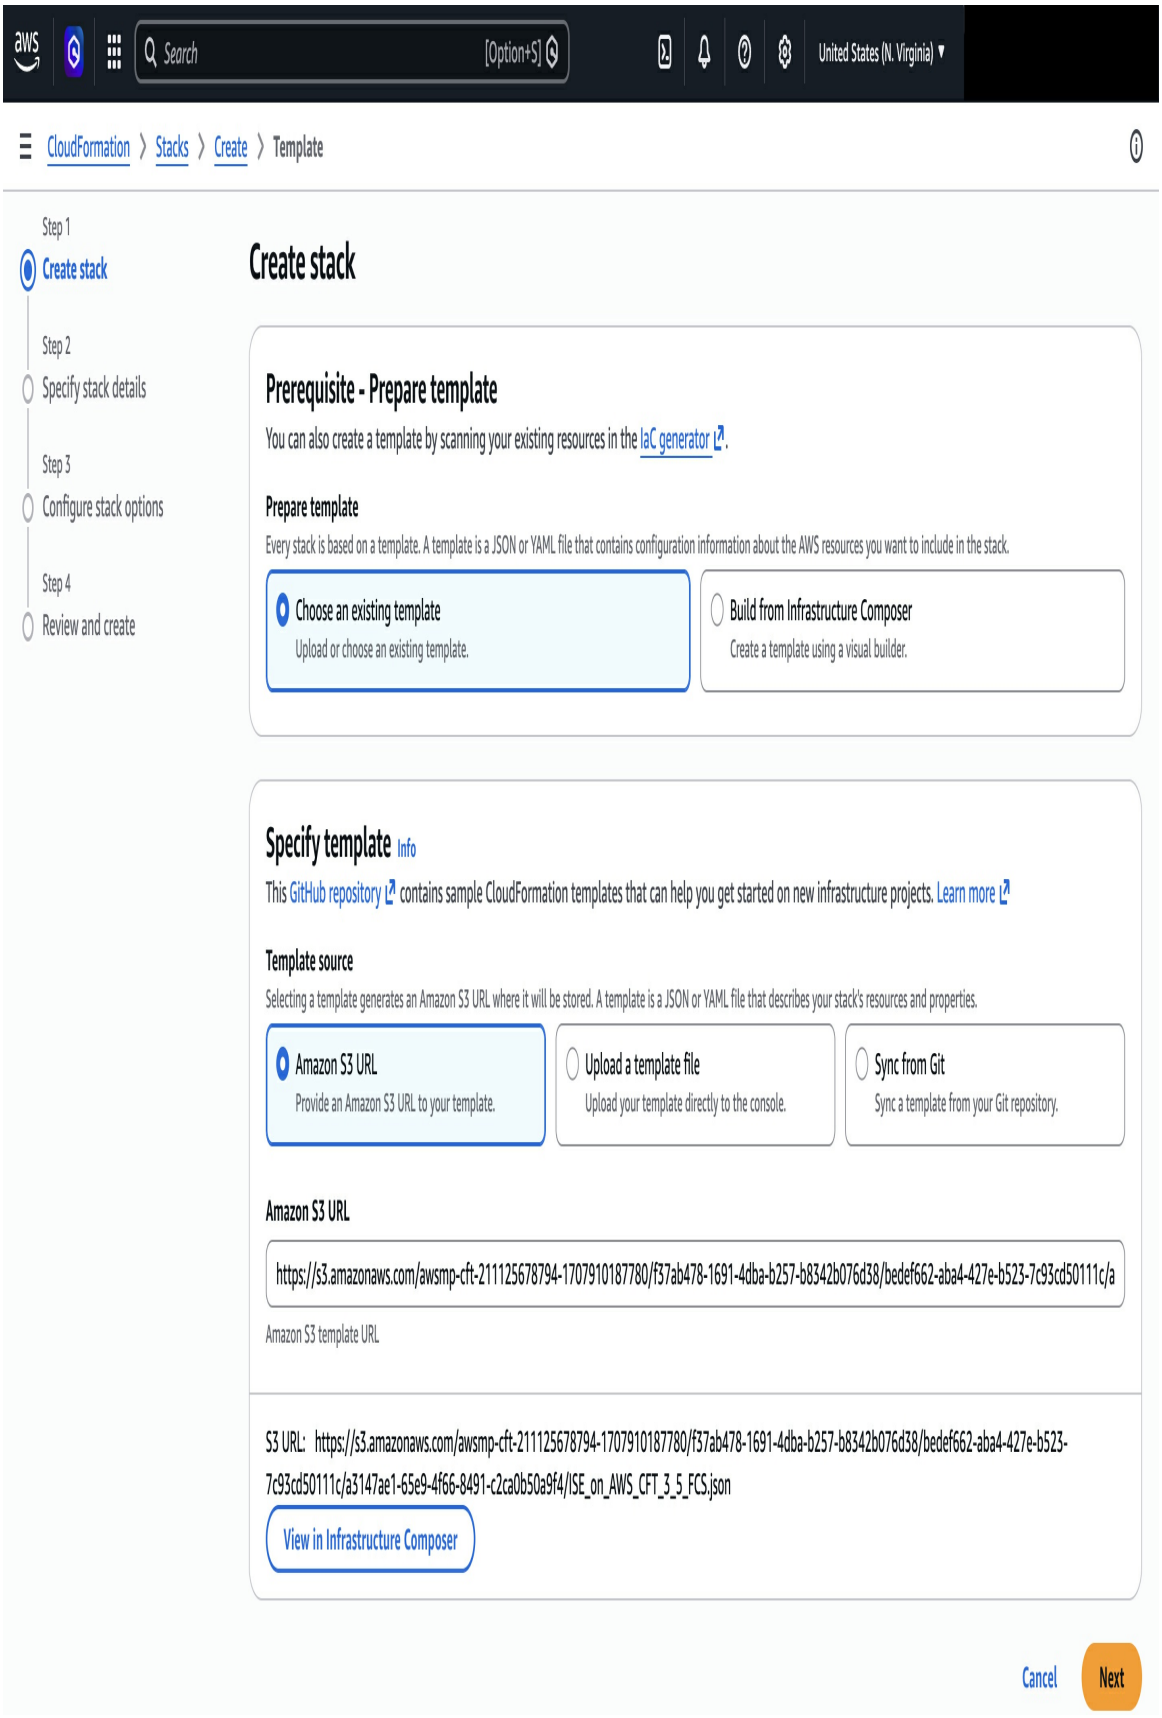
At this point, select **AWS CloudFormation**, choose the **Version** of ISE you want to install, and select the **Region** where the ISE node will be deployed. Finally, click **Launch with CloudFormation**, as shown in Figure 3-4.


aws marketplace
Describe the solution you're looking for...
Agent Mode
English
About
Categories
Delivery Methods
Solutions
Resources
Become a Channel Partner
Sell in AWS Marketplace
Amazon Web Services Home
Help
AWS Marketplace
Cisco Identity Services Engine (ISE)
Launch Cisco Identity Services Engine (ISE)
Launch Cisco Identity Services Engine (ISE) Info
Setup
Service Info
AWS CloudFormation
Automated, one-step deployment.
Amazon EC2
Custom deployment using Amazon Machine Image (AMI).
Launch method Info
AWS CloudFormation template
Version Info
3.5.0 (Sep 25, 2025)
Release notes
Region
US East (N. Virginia)
Launch
Launch with CloudFormation
AWS CloudFormation automates consistent and reliable deployment. Its reusable templates mean less errors and security risks.
View template in CloudFormation Designer
Launch with CloudFormation

Technet24

Figure 3-4 *AWS Launch with CloudFormation*

CloudFormation Template is a JSON file developed by Cisco that assists administrators with the initial ISE configuration. No changes are required on this screen. Click **Next**, as shown in Figure 3-5.


Search
[Option+S]
United States (N. Virginia)
CloudFormation
Stacks
Create
Template
Step 1
Create stack
Step 2
Specify stack details
Step 3
Configure stack options
Step 4
Review and create
Create stack
Prerequisite - Prepare template
You can also create a template by scanning your existing resources in the IaC generator.
Prepare template
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.
Choose an existing template
Upload or choose an existing template.
Build from Infrastructure Composer
Create a template using a visual builder.
Specify template Info
This GitHub repository contains sample CloudFormation templates that can help you get started on new infrastructure projects. Learn more
Template source
Selecting a template generates an Amazon S3 URL where it will be stored. A template is a JSON or YAML file that describes your stack's resources and properties.
Amazon S3 URL
Provide an Amazon S3 URL to your template.
Upload a template file
Upload your template directly to the console.
Sync from Git
Sync a template from your Git repository.
Amazon S3 URL
https://s3.amazonaws.com/awsmp-cft-211125678794-1707910187780/f37ab478-1691-4dba-b257-b8342b076d38/bedef662-aba4-427e-b523-7c93cd50111c/a
Amazon S3 template URL
S3 URL: https://s3.amazonaws.com/awsmp-cft-211125678794-1707910187780/f37ab478-1691-4dba-b257-b8342b076d38/bedef662-aba4-427e-b523-7c93cd50111c/a3147ae1-65e9-4f66-8491-c2ca0b50a9f4/ISE_on_AWS_CFT_3_5_FCS.json
View in Infrastructure Composer
Cancel
Next

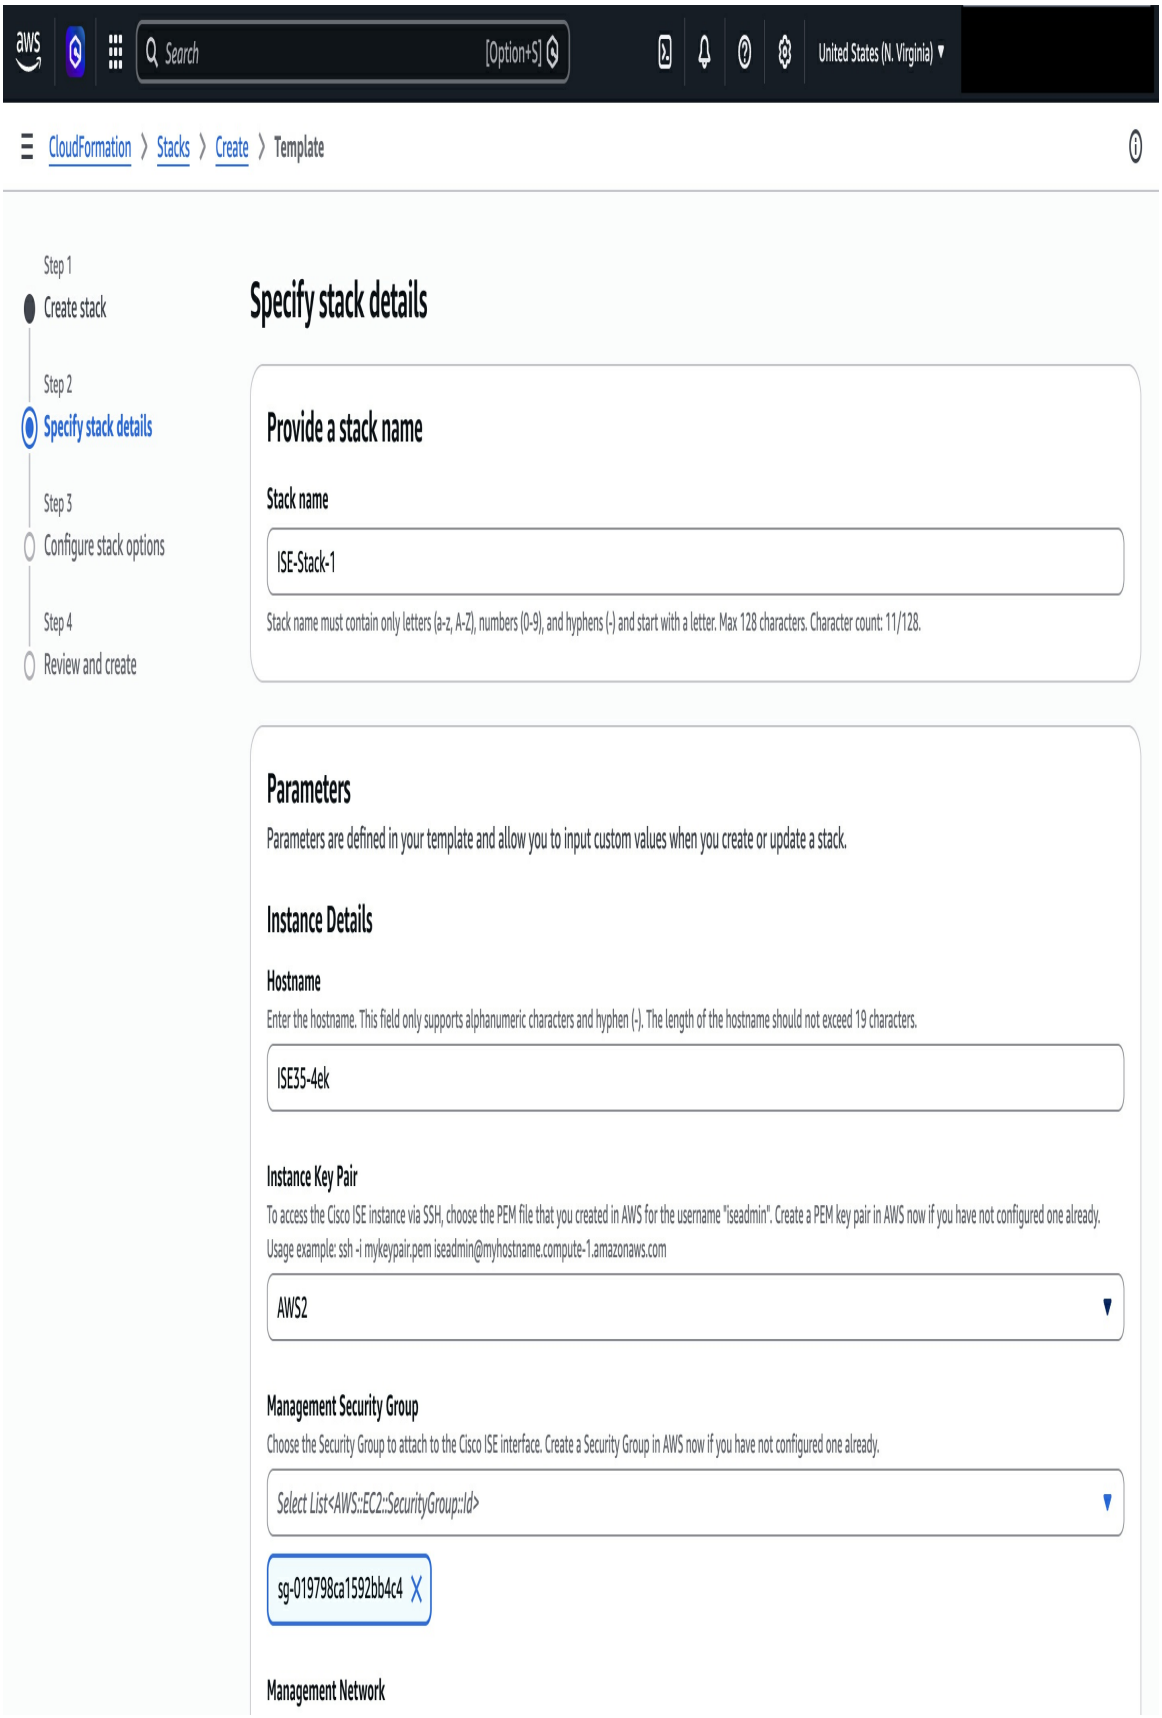
Technet24

Figure 3-5 *CloudFormation Stack Step 1*

> **Note:**
>
> Configure the Key Pair and Management Security Group before launching the CloudFormation Template.

The next steps include configuring specific AWS settings as well as ISE bootstrap settings. First, provide a stack name. This name serves as a unique identifier for the AWS resource collection and does not carry over to ISE.

Configure an ISE hostname, specify the existing **Key Pair** that will be used to access ISE via SSH, and select the **Management Security Group** that protects the EC2 instance, as shown in Figure 3-6.


Search
[Option+S]
United States (N. Virginia)
CloudFormation
Stacks
Create
Template
Step 1
Create stack
Step 2
Specify stack details
Step 3
Configure stack options
Step 4
Review and create
Specify stack details
Provide a stack name
Stack name
ISE-Stack-1
Stack name must contain only letters (a-z, A-Z), numbers (0-9), and hyphens (-) and start with a letter. Max 128 characters. Character count: 11/128.
Parameters
Parameters are defined in your template and allow you to input custom values when you create or update a stack.
Instance Details
Hostname
Enter the hostname. This field only supports alphanumeric characters and hyphen (-). The length of the hostname should not exceed 19 characters.
ISE35-4ek
Instance Key Pair
To access the Cisco ISE instance via SSH, choose the PEM file that you created in AWS for the username "iseadmin". Create a PEM key pair in AWS now if you have not configured one already.
Usage example: ssh -i mykeypair.pem iseadmin@myhostname.compute-1.amazonaws.com
AWS2
Management Security Group
Choose the Security Group to attach to the Cisco ISE interface. Create a Security Group in AWS now if you have not configured one already.
Select List<AWS::EC2::SecurityGroup::Id>
sg-019798ca1592bb4c4
Management Network

Technet24

Figure 3-6 *ISE Hostname, Instance Key Pair, and Security Group Settings*

Next, configure the management network where you will place ISE, and specify the **IP Address** from this network to assign to the G0 interface of ISE. Set the **Time Zone** and select the **Instance Type** from the supported list. Choose the appropriate **Volume Size** for the ISE node, as shown in Figure 3-7.





Management Network

Choose the subnet to be used for the Cisco ISE interface. To enable IPv6 addresses, you must associate an IPv6 CIDR block with your VPC and subnets. Create a Subnet in AWS now if you have not configured one already.

subnet-00c74ba0977cdcf9c

Management Private IPv4 Address

(Optional) Enter the IPv4 address from the subnet that you chose earlier. If this field is left blank, the AWS DHCP will assign an IP address.

172.31.200.150

Create IPv6-only Instance

Choose whether to create an IPv6-only instance. If yes, no IPv4 address will be assigned.

no

Management Private IPv6 Address

(Optional) Enter the IPv6 address from the subnet that you chose earlier. If this field is left blank, AWS will assign an IPv6 address automatically.

Enter String

Time Zone

Choose a system time zone.

Etc/UTC

Instance Type

Choose the required Cisco ISE instance type.

c5.9xlarge

Volume Size

Specify the storage in GB (Minimum 300GB and Maximum 2400GB). 600GB is recommended for production use, storage lesser than 600GB can be used for evaluation purpose only. On terminating the instance, volume will be deleted as well.

600

EBS Encryption

Technet24

Figure 3-7 *Management Network, IP Address, Time Zone, and Instance Settings*

Next, configure **EBS Encryption**, **DNS Domain**, and **Name Server**, and define the **NTP Server**, as shown in Figure 3-8.

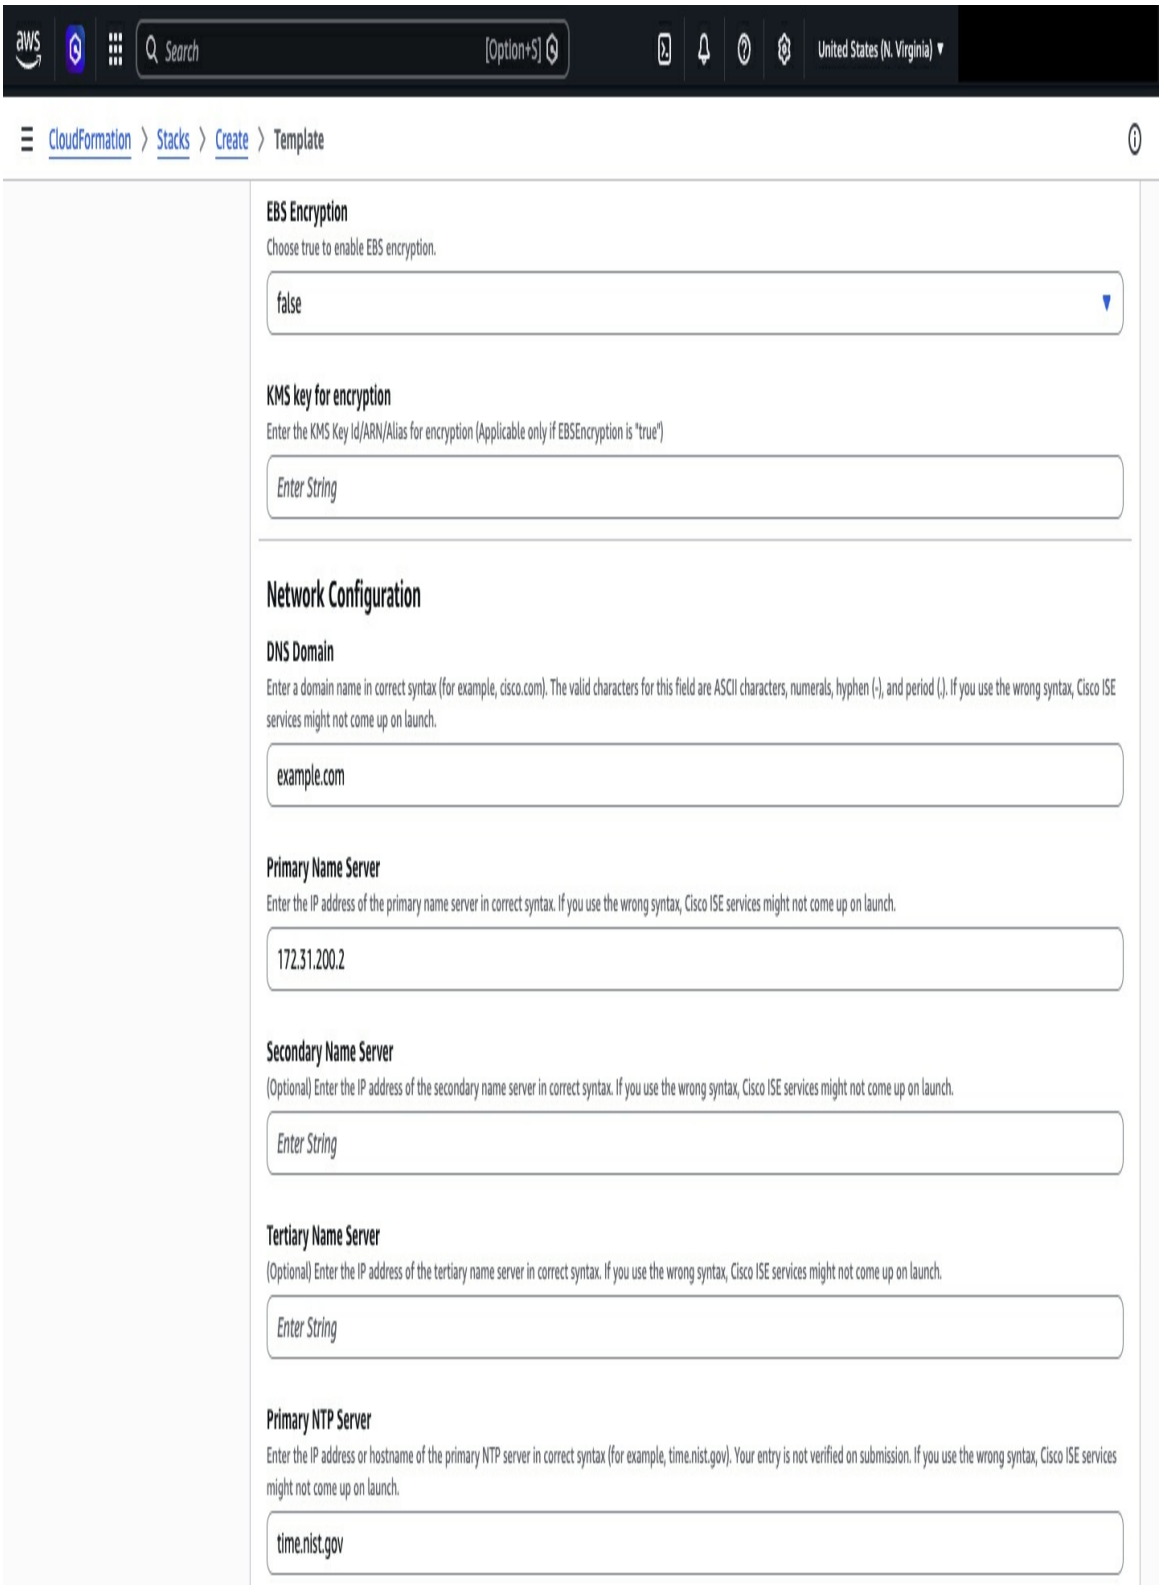
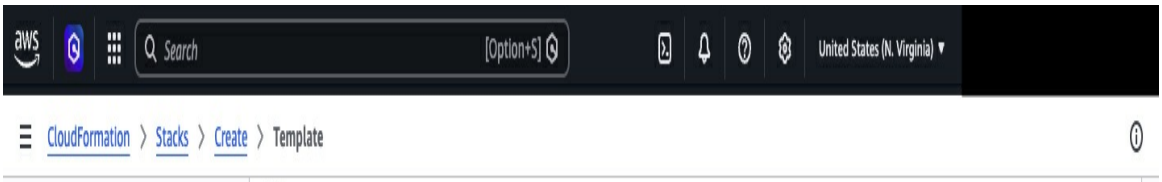


Figure 3-8 *ISE Domain, DNS Server, and NTP Server Settings*

Technet24

Finally, select whether you want **pxGrid** or **ERS** enabled by default, and configure the GUI password for the iseadmin user, as shown in Figure 3-9. Then select **Next**.



ERS

Do you wish to enable ERS?

yes

pxGrid

Do you wish to enable pxGrid?

no

User Details

Enter Password

Enter a password for the username "iseadmin". The password must be aligned with the Cisco ISE password policy. The configured password is used for Cisco ISE GUI access. Warning: The password is displayed in plaintext in the User Data section of the Instance settings window in the AWS Console.

••••••••

Confirm Password

Retype Password

••••••••

MPS3BucketName

staging-templates-prod-us-east-1-654654338570

MPS3BucketRegion

us-east-1

MPS3KeyPrefix

bedef662-aba4-427e-b523-7c93cd50111c/a3147ae1-65e9-4f66-8491-c2ca0b50a9f4/

MarketplaceImageId

This is the alias of the Marketplace AMI that will be deployed as part of this stack. Ensure this parameter is set to the following value: /aws/service/marketplace/prod-7wsm5sh6ugdle/3.5.0.

/aws/service/marketplace/prod-7wsm5sh6ugdle/3.5.0

Cancel Previous Next

Technet24

Figure 3-9 *ISE Services and iseadmin Password Settings*

Select **Next** to proceed to the Summary screen and click **Submit**.

When the ISE CloudFormation deployment is complete, you can access your ISE EC2 instance from the AWS Console, as shown in Figure 3-10.

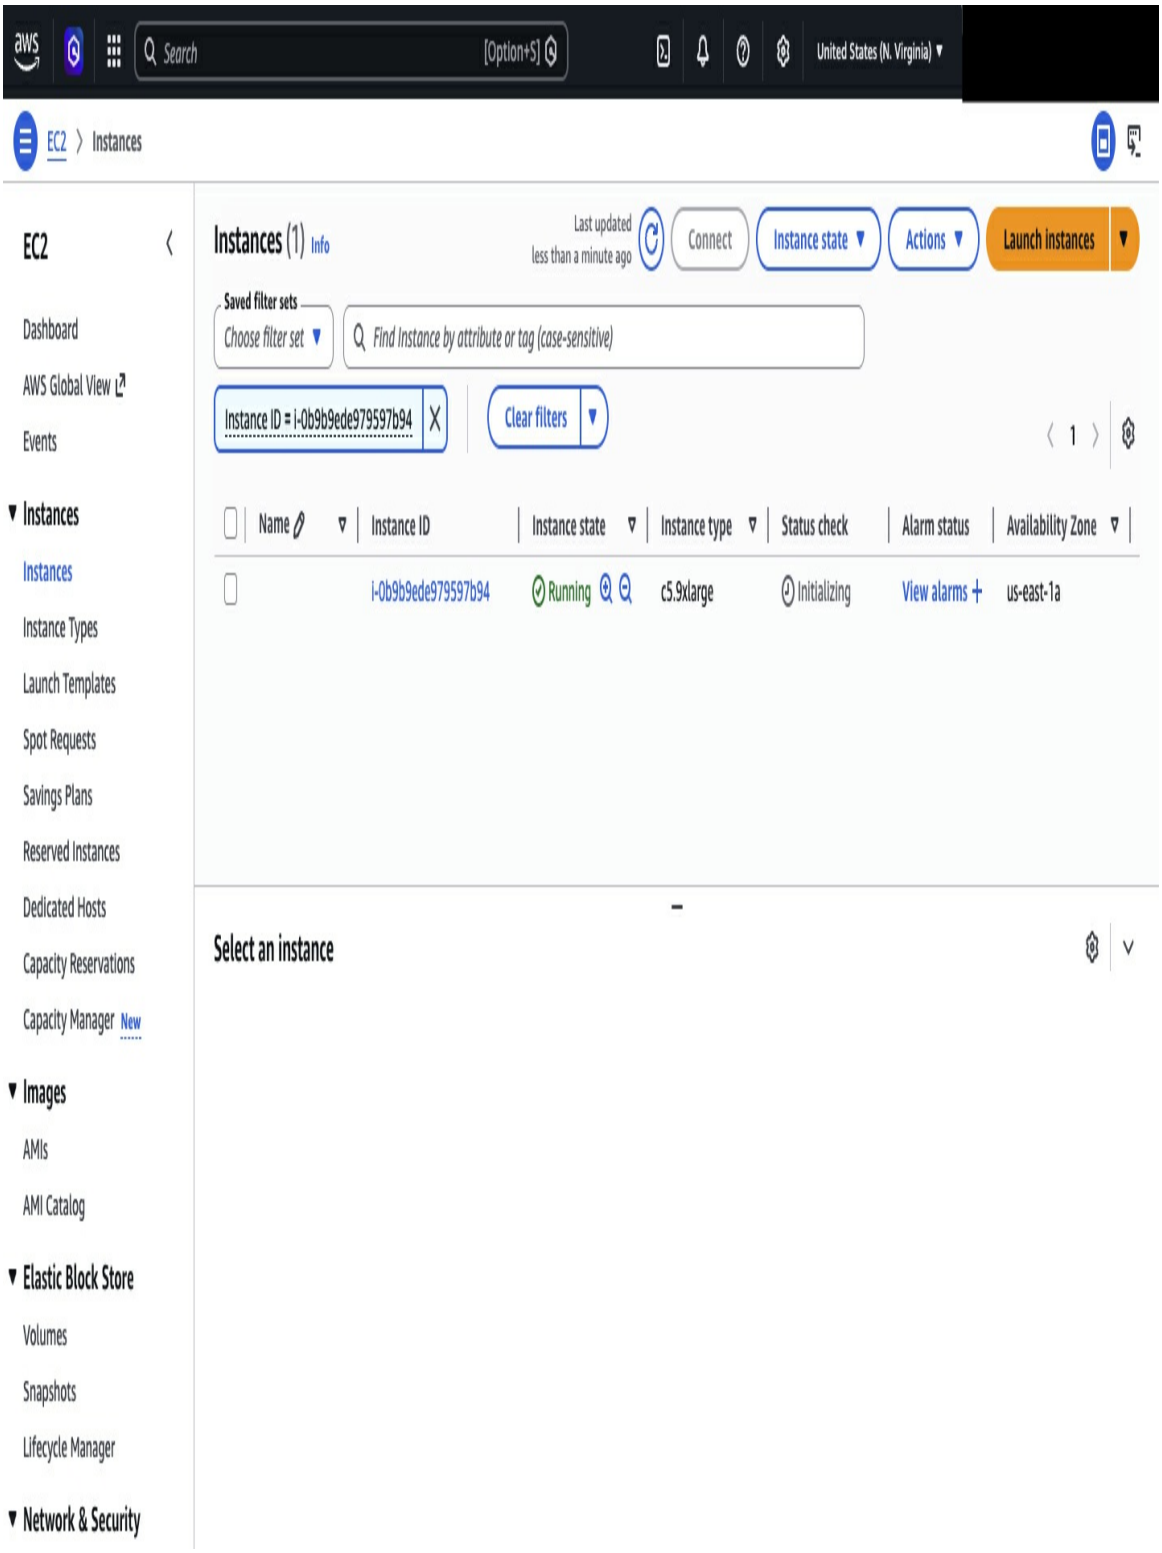


Figure 3-10 *ISE Instance Deployed via CloudFormation Template*

You can access the ISE CLI approximately 20–30 minutes after the ISE

Technet24

EC2 instance is built. To access the ISE instance via SSH, use the public key by entering the following command:

```
% ssh -i AWS2.pem iseadmin@
```

> **Note:**
>
> Change the key permissions before accessing the ISE node. For Linux or macOS, issue the following command:
>
> ```
> % sudo chmod 0400 mykeypair.pem
> ```

ISE services may require an additional 20–30 minutes to be fully operational, depending on the platform.

> **Note:**
>
> You can monitor the status of services from the AWS Serial Console. However, you should avoid interrupting the process and allow sufficient time for services to start.

Limitations of Cisco ISE on AWS

There are several limitations associated with ISE deployment on AWS. The following limitations are the most critical:

- NIC bonding is not supported in AWS.
- Amazon Elastic Block Store (EBS) snapshots are not supported.
- Changing the IP address is not supported after the instance is deployed. To use a new IP address, create a new instance.
- Resizing an EC2 instance is not supported.
- Upgrade workflow is not supported. To upgrade an ISE instance in the cloud, install the new version of ISE and perform a backup and restore of the ISE configuration database.

ISE on Azure

Cisco ISE is available on Azure Cloud in a manner similar to AWS. Although you can deploy instances using different methods, the ISE instance will ultimately operate as an Azure virtual machine.

Overview of Cisco ISE on Azure

You can deploy ISE in Azure using two methods:

- Azure Application
- Azure Virtual Machine

> **Best Practice:**
>
> Cisco recommends installing ISE using Azure Application for manual deployments. If you use IaC tools such as Terraform or Ansible, choose Azure Virtual Machine installation.

Azure Application deployment relies on Azure Resource Manager (ARM) templates. These JSON-based templates enable initial configuration directly from the Azure UI, rather than requiring manual input through a user data file.

Similar to AWS, ISE instances in Azure also incorporate the concept of intended usage. You should use this concept to select the instance types, in addition to considering concurrent sessions as discussed in Chapter 2.

ISE can be installed only on specific virtual machine types. Table 3-2 describes production instance types and their intended purpose as of ISE 3.5.

Table 3-2 *Azure ISE Instances and Their Intended Usage*

Technet24

Instance Type	Intended Use	CPU Cores	RAM (in GB)
Standard_D8s_v4	Extra Small PSN only	8	32
Standard_F16s_v2	PSN	16	32
Standard_F32s_v2		32	64
Standard_D32s_v4	PAN or MnT or both	32	128
Standard_D64s_v4		64	256

ISE Installation via Azure Application

To install ISE from Azure Cloud as an Azure Application, navigate to the **Marketplace**. Search for **Cisco ISE** in the Marketplace search window and select the **Cisco Identity Services Engine (ISE)** Azure Application, as shown in Figure 3-11.

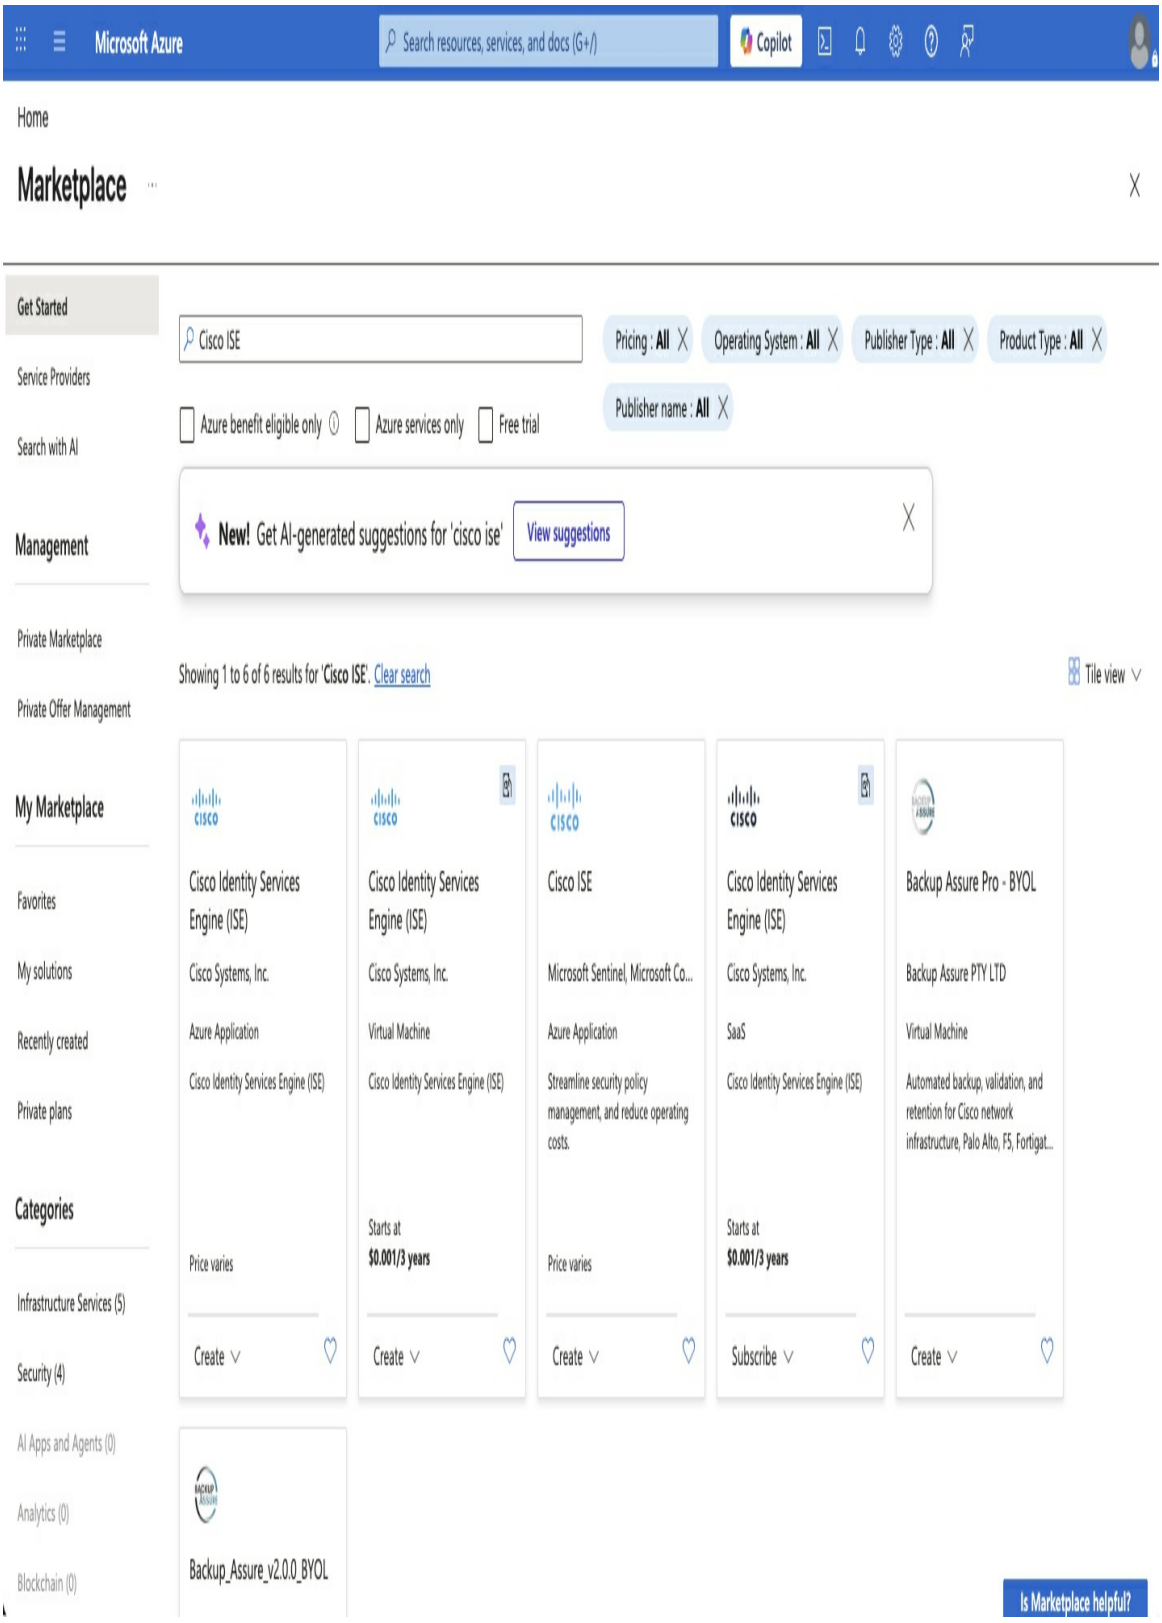


Figure 3-11 *ISE Azure Application on Azure Cloud Marketplace*

Technet24

Select the subscription where you want to deploy the ISE node and choose the ISE major release version from the **Plan** drop-down menu, as shown in Figure 3-12.

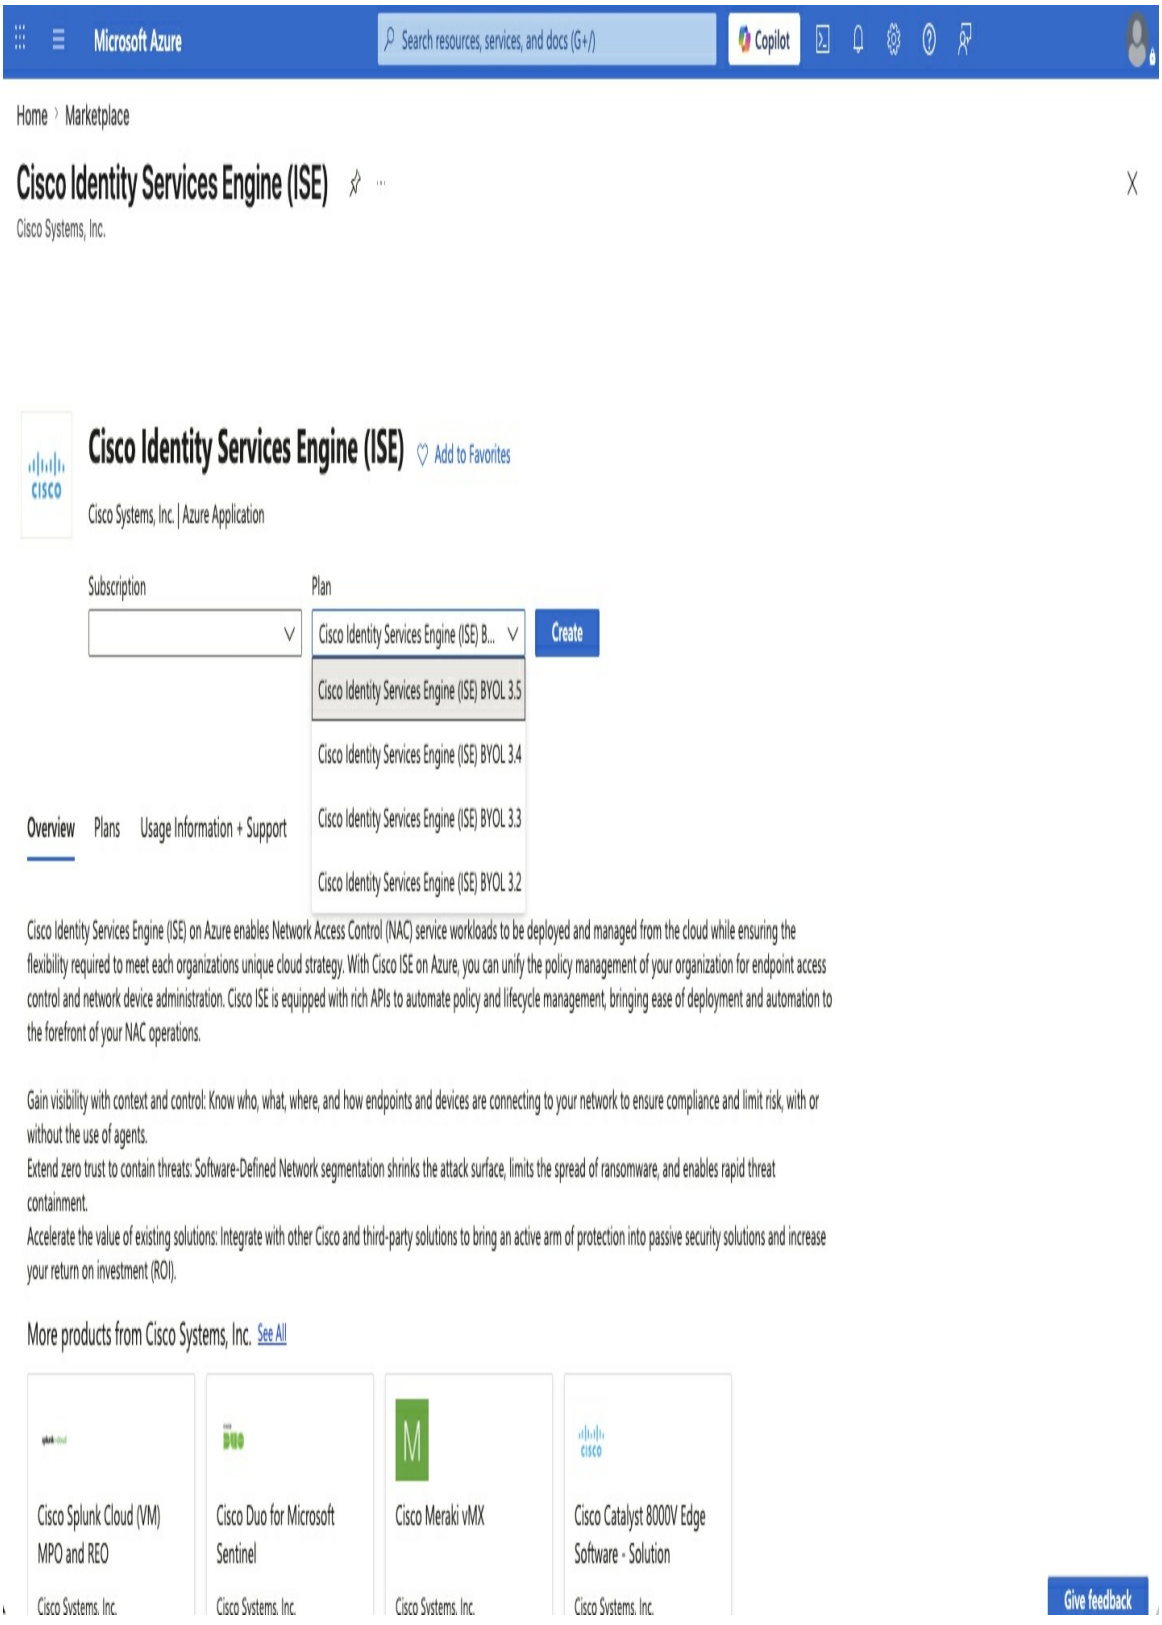


Figure 3-12 *ISE Subscription and Version Selection*

Technet24

Note:

Define the Key Pair, Security Group, Network, and Subnet before launching the ISE Application on Azure.

As shown in Figure 3-13, place resources in **Resource Group**. This is a convenient way to manage resources in Azure Cloud. Specify the **Region** for ISE deployment, assign a **Host Name** to the ISE node, select the **Time Zone**, and choose the **VM Size**. Allocate the **Volume Size**, which determines the disk size. Then click **Next**.

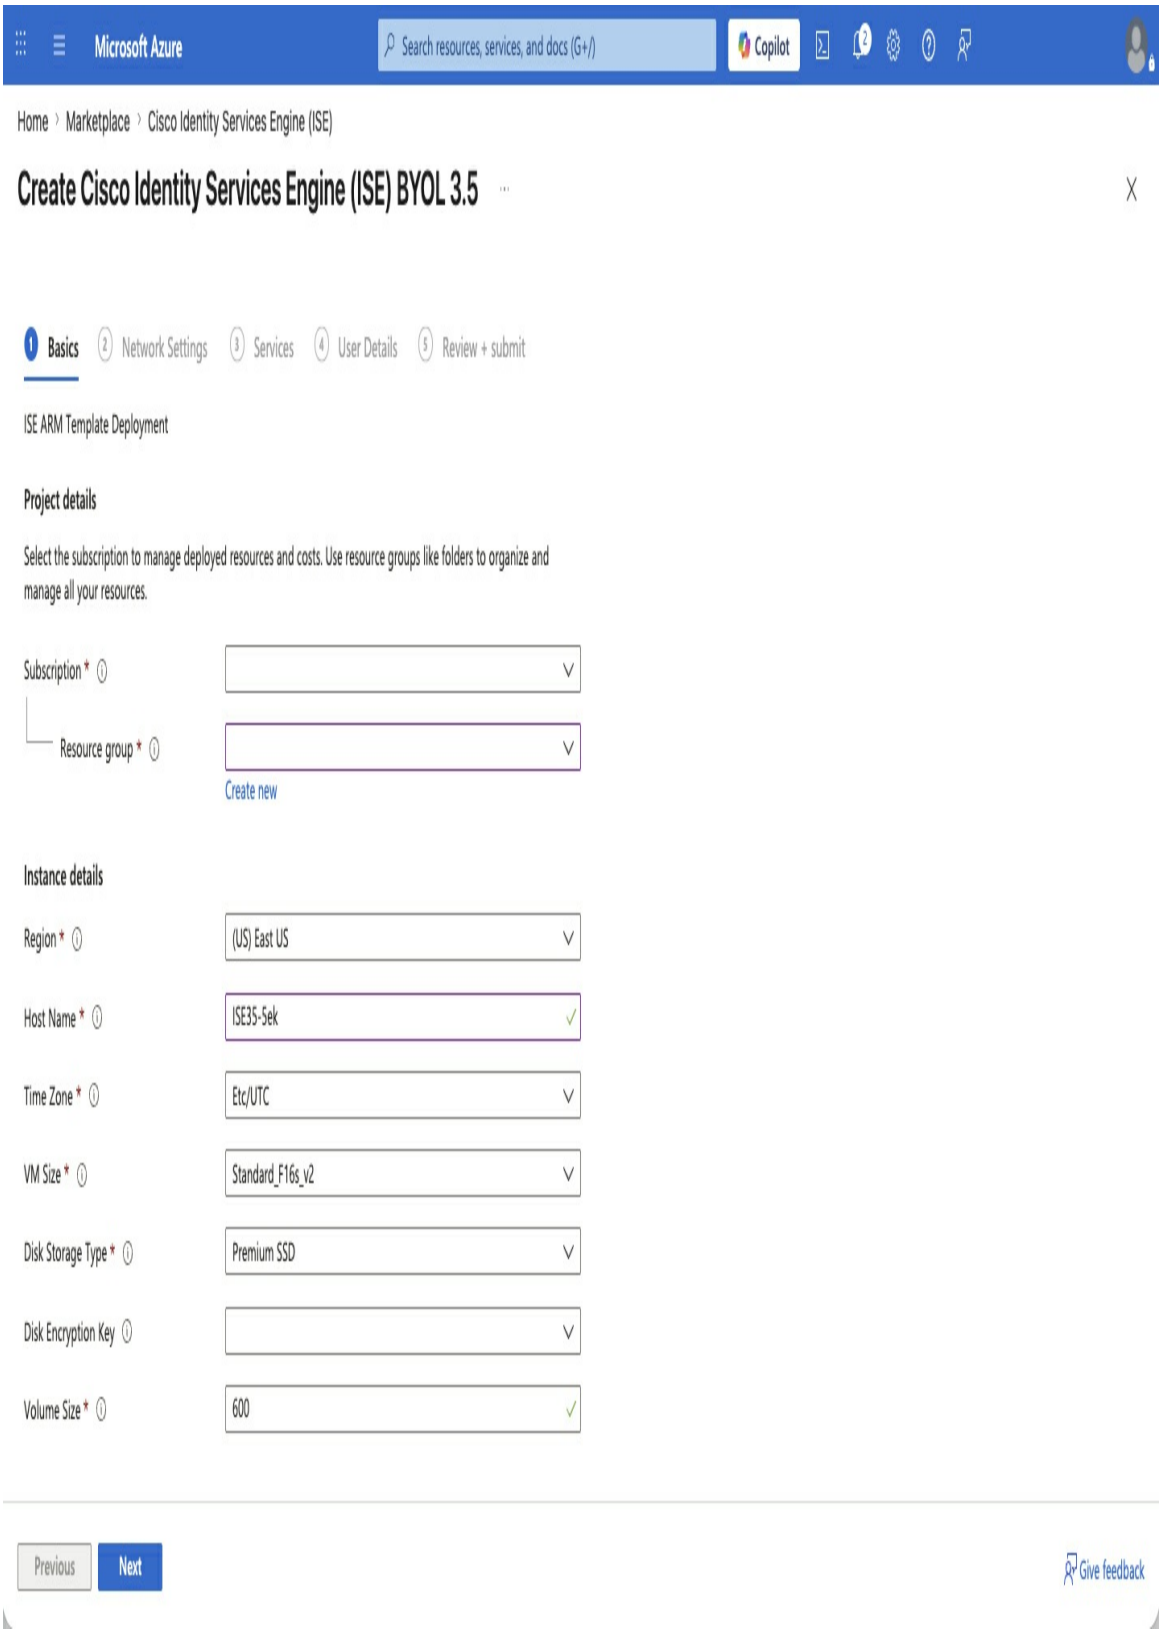


Figure 3-13 *Subscription, Resource Group, Region, Host Name,*

Technet24

Time Zone, VM Size, and Disk Settings

Next, configure the network settings, which include the **Virtual Network**, **Subnet**, and **Network Security Group** that controls which traffic is allowed to and from the ISE virtual machine. Define the key pair used to access ISE via SSH. Configure the **IP Address**, **DNS Domain Name**, **Name Server**, and **NTP Server**, as shown in Figure 3-14. Then click **Next**.

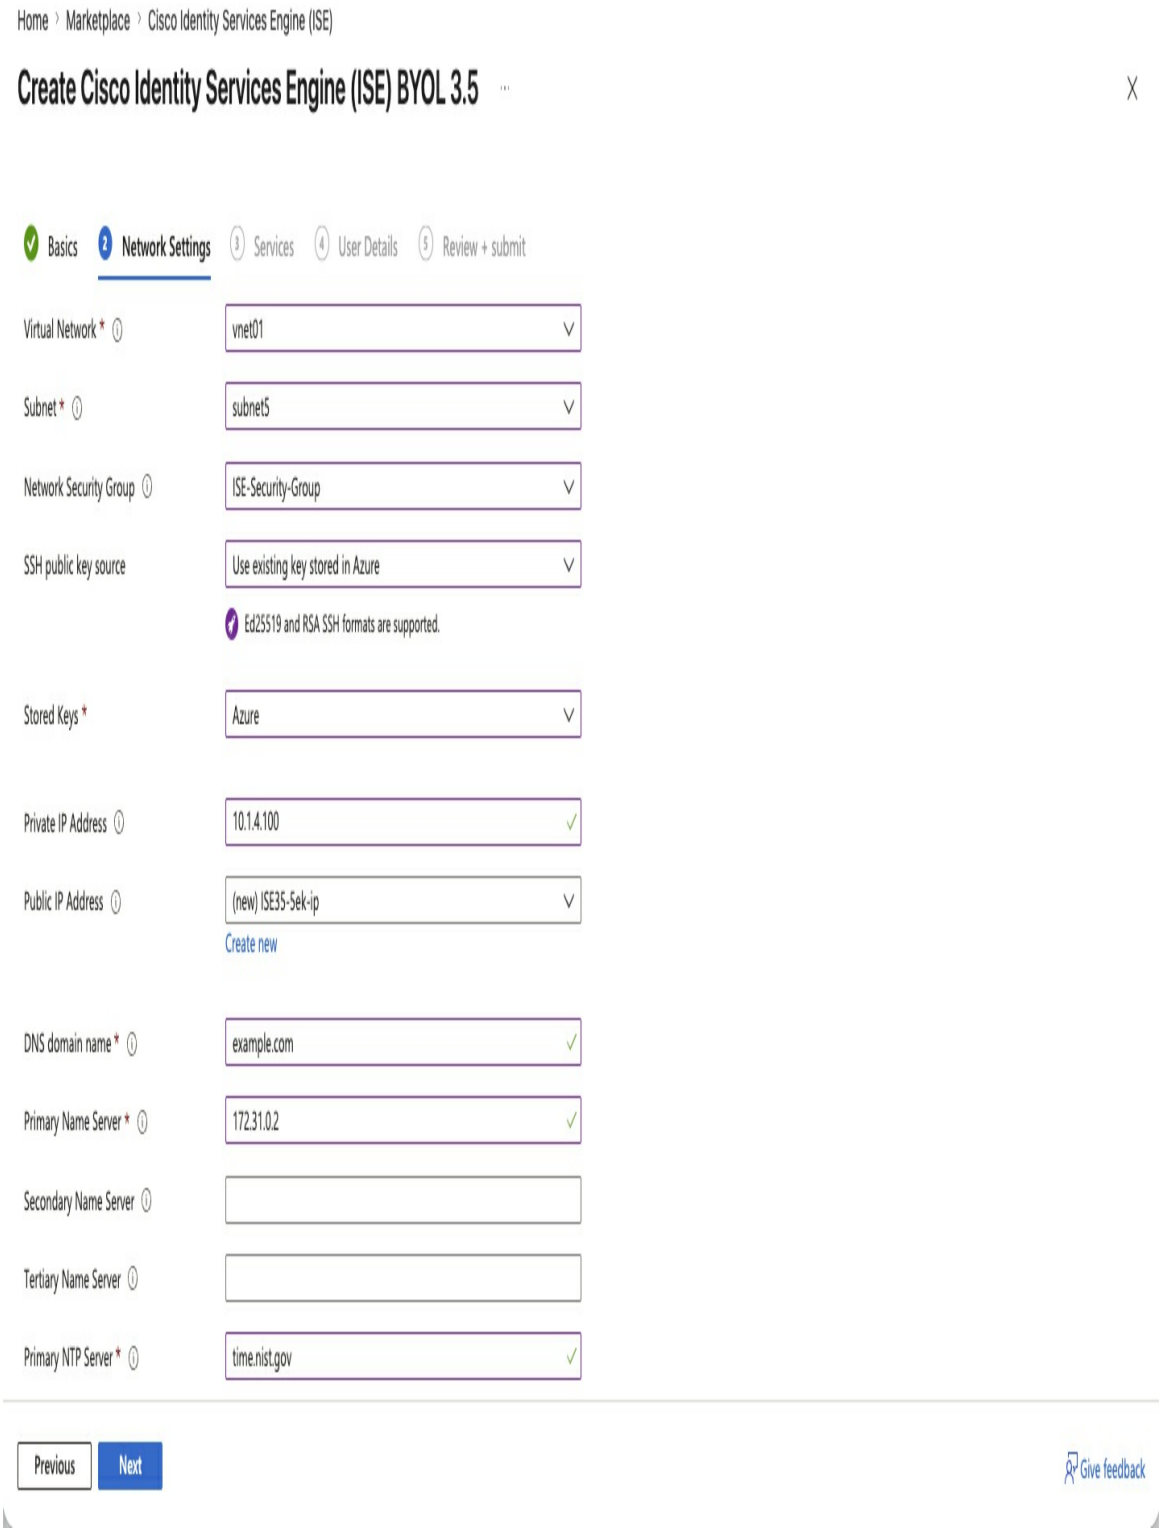


Figure 3-14 *ISE Network Settings*

On the next screen, select the additional **Services** you want to enable

Technet24

immediately after installation, as shown in Figure 3-15. Then click **Next**.



Figure 3-15 *ISE Services Settings*

Technet24

Next, define a password for the iseadmin user, as shown in Figure 3-16. Use this password to access the ISE GUI. Click **Next**.

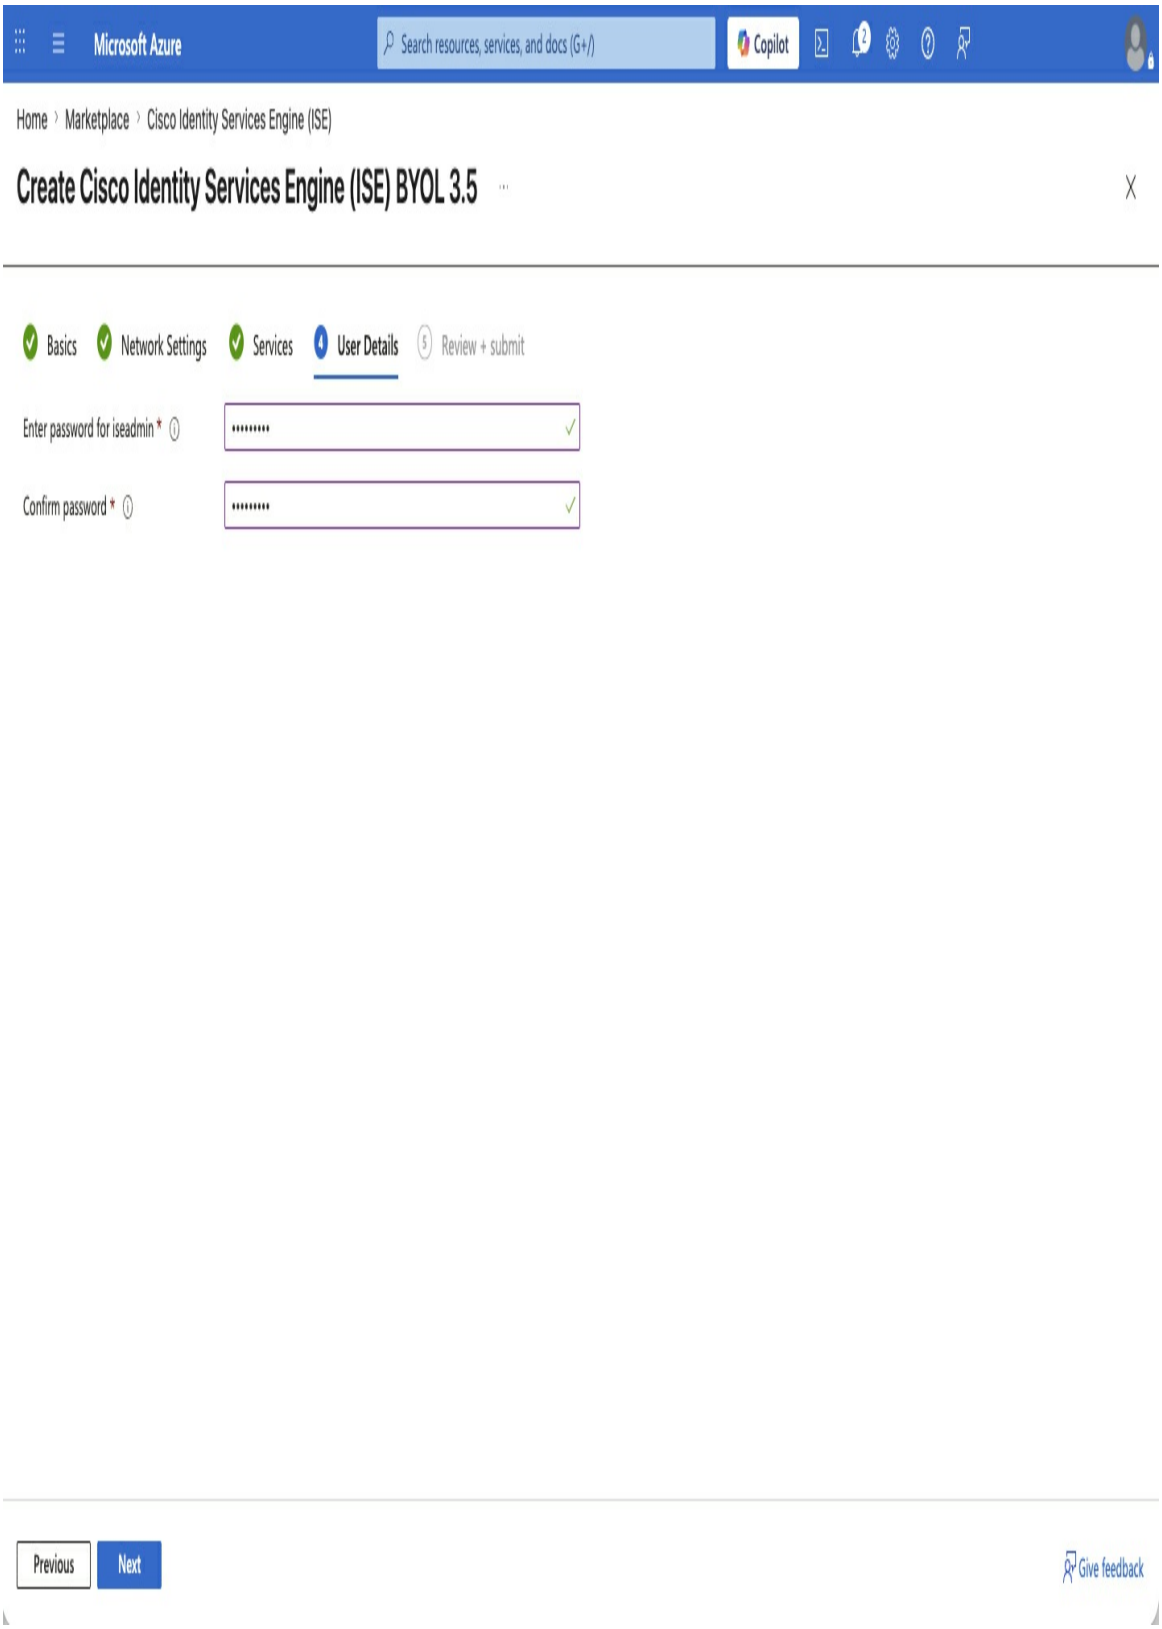


Figure 3-16 *ISE Password Settings*

Technet24

After you have configured all settings, review the configuration and click **Create** to deploy the ISE virtual machine in Azure Cloud, as shown in Figure 3-17.



Figure 3-17 *ISE Review and Submit Settings*

Technet24

To view the virtual machine created as part of the ISE Application deployment via Azure Marketplace, navigate to the **Azure Virtual Machines** menu, as shown in Figure 3-18.

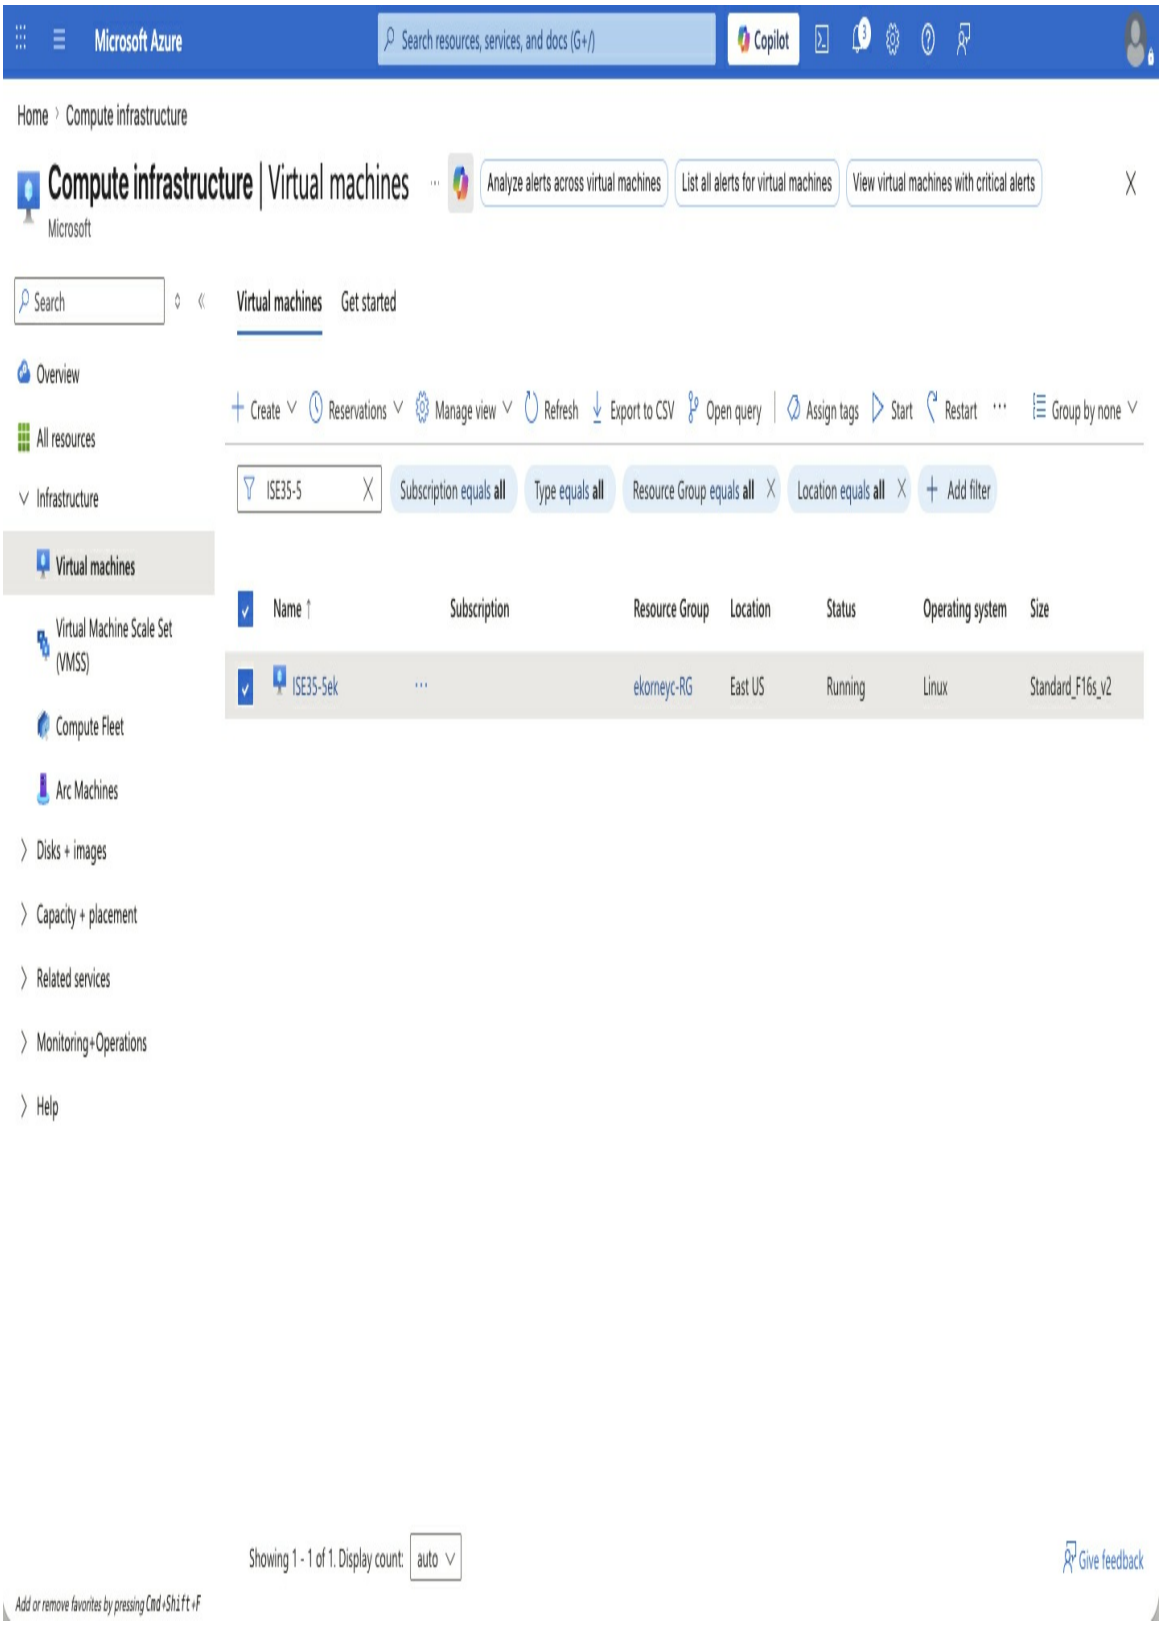


Figure 3-18 *Azure Virtual Machines Menu*

Technet24

You can access the ISE CLI approximately 20–30 minutes after the ISE virtual machine is deployed. To access the ISE instance via SSH, use the public key with the following command:

```
% ssh -i Azure.pem iseadmin@
```

Limitations of Cisco ISE on Azure

There are several limitations associated with ISE deployment on Azure. The following are the most critical:

- Accelerated networking is not supported in Azure Cloud. Enabling this feature can cause registration and deregistration issues.
- Upgrade workflow is not supported. To upgrade an ISE instance in the cloud, install the new version of ISE and perform a backup restore.
- Microsoft Azure, by default, drops out-of-order (OOO) IP fragments. This behavior can affect ISE deployment, especially if you use EAP-TLS authentication, which involves large packets containing certificates. You may encounter this issue with Azure VPN Gateway, ExpressRoute (with no FastPath enabled), or Azure Virtual WAN. To mitigate this issue:

  Deploy third-party VPN gateway (such as Catalyst 8000v or ASAv) as the Azure virtual machine. Ensure virtual reassembly is configured on the VPN gateway.

  If you continue to experience OOO drops, you can open a support ticket with Microsoft to allow OOO fragments.

Summary

Deploying ISE in the cloud provides significant flexibility in your ISE

cluster design. You can implement a hybrid ISE deployment with some instances in the cloud and others on premises, or you can choose a fully cloud-based deployment. ISE is supported in AWS, Azure, and OCI, offering multiple options for your deployment needs.

Technet24

Chapter 4. Policy Set Concepts

Introduction

One of the main functions of Cisco Identity Services Engine is to perform authentication and authorization of endpoints and users. To achieve better usability of the policies, ISE relies on policy sets. These logical rules help administrators group like combinations of authentication and authorization policies under a single policy set.

Policy Sets Overview

ISE processes policy sets in sequential order from top to bottom, and based on the first condition met, the matching policy set will be selected.

Two options can result in the selected policy set:

- RADIUS Server Sequence
- Allowed Protocols

ISE initially includes a single policy set called Default with no conditions, acting as a catch-all policy set, as shown in Figure 4-1. The default **Allowed Protocols/Server Sequence** is **Default Network Access**.

Best Practice:

In a production environment, you should maintain the default policy set as a catch-all rule, where the result is **Deny Access**. You should manage network access by defining and applying custom policy sets.

Technet24

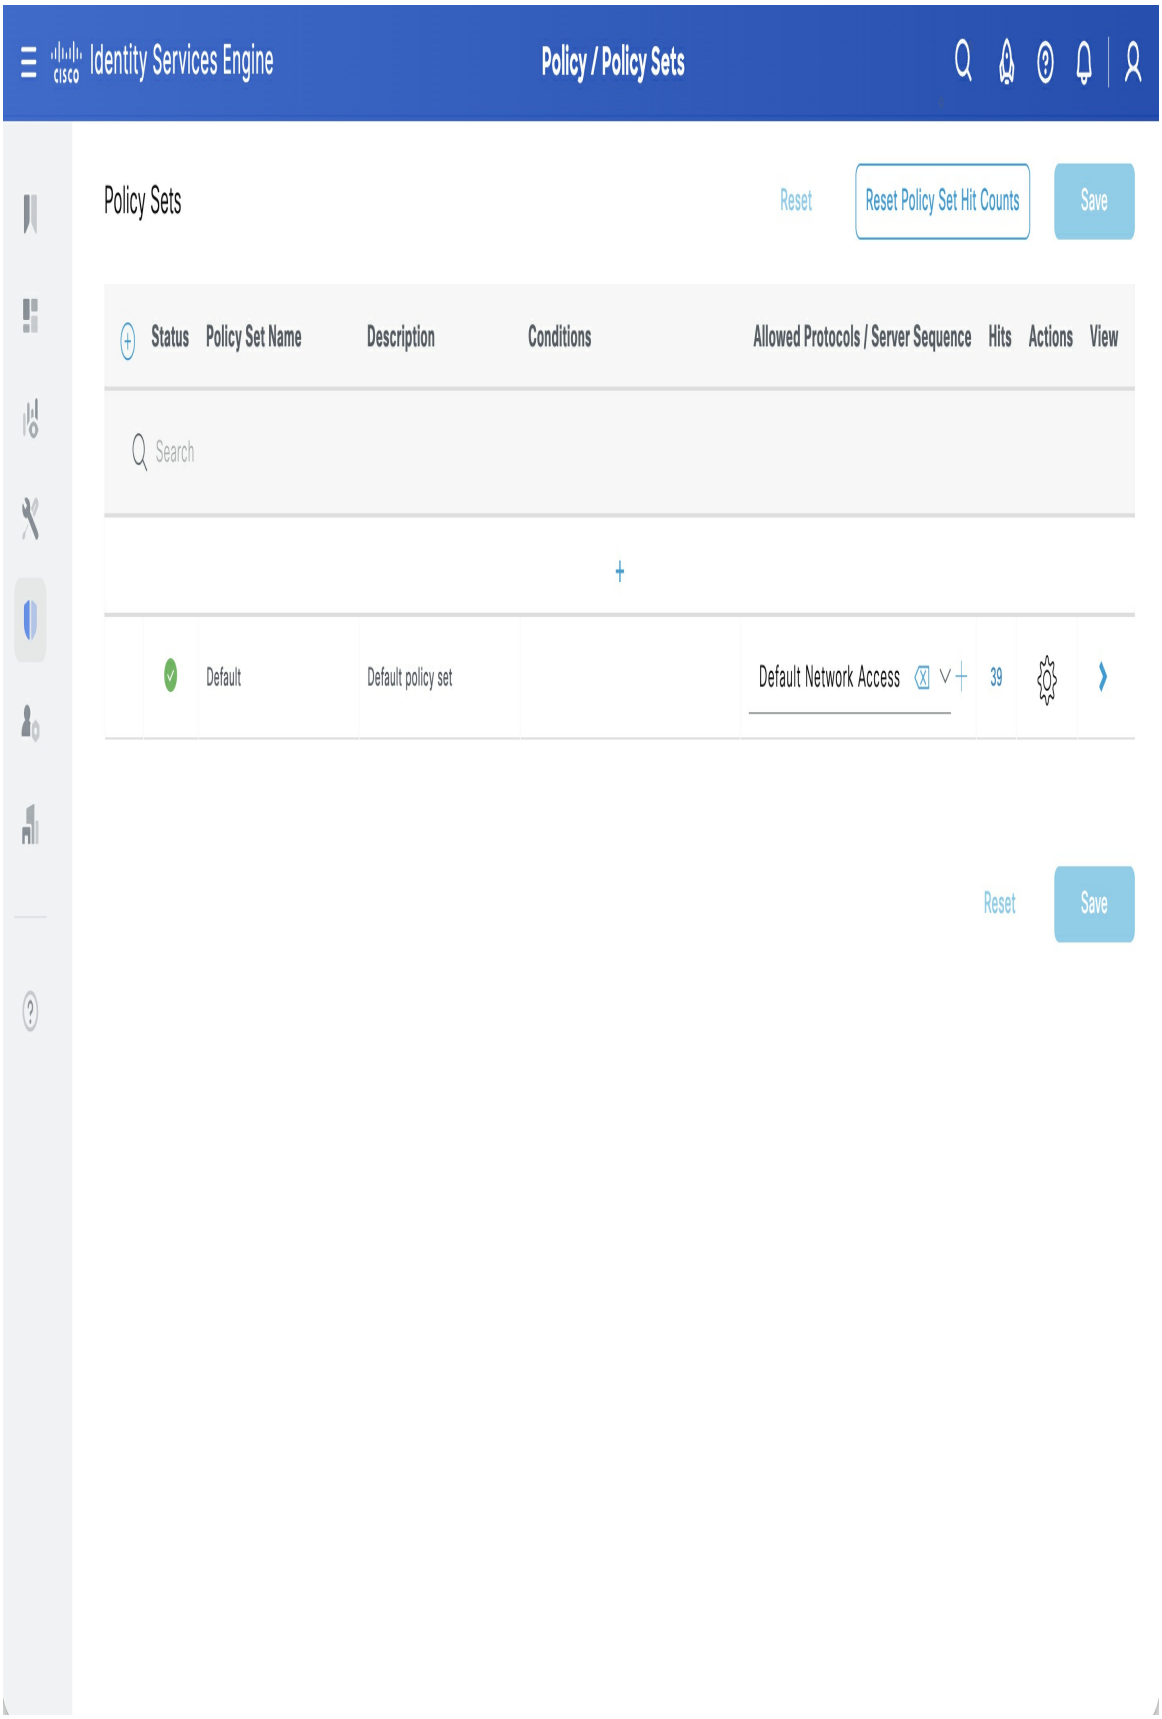

Identity Services Engine
Policy / Policy Sets
Policy Sets
Reset
Reset Policy Set Hit Counts
Save
Status
Policy Set Name
Description
Conditions
Allowed Protocols / Server Sequence
Hits
Actions
View
Search
Default
Default policy set
Default Network Access
39
Reset
Save

Figure 4-1 *Cisco ISE Policy Sets*

Granular policy sets allow you to achieve the following:

- Improve readability and reduce administrative overhead by shortening the lists of policies within each policy set.
- Enhance performance by minimizing the number of rules processed during authentication and authorization.
- Reduce the number of conditions in Authorization Policy rules by applying conditions at the policy set level rather than in multiple authorization policies.
- Review hit counts for each use case or location at a glance.

Most common use cases for policy sets are to distinguish authentications based on

- Location (One region versus another)
- Authentication method (wireless versus wired versus VPN)
- Authentication type (MAB versus 802.1X versus VPN)

Here are a few practical examples. Note that the actual Allowed Protocols would depend on specific use cases implemented with ISE. Figure 4-2 shows how to distinguish between authentication methods.

Technet24

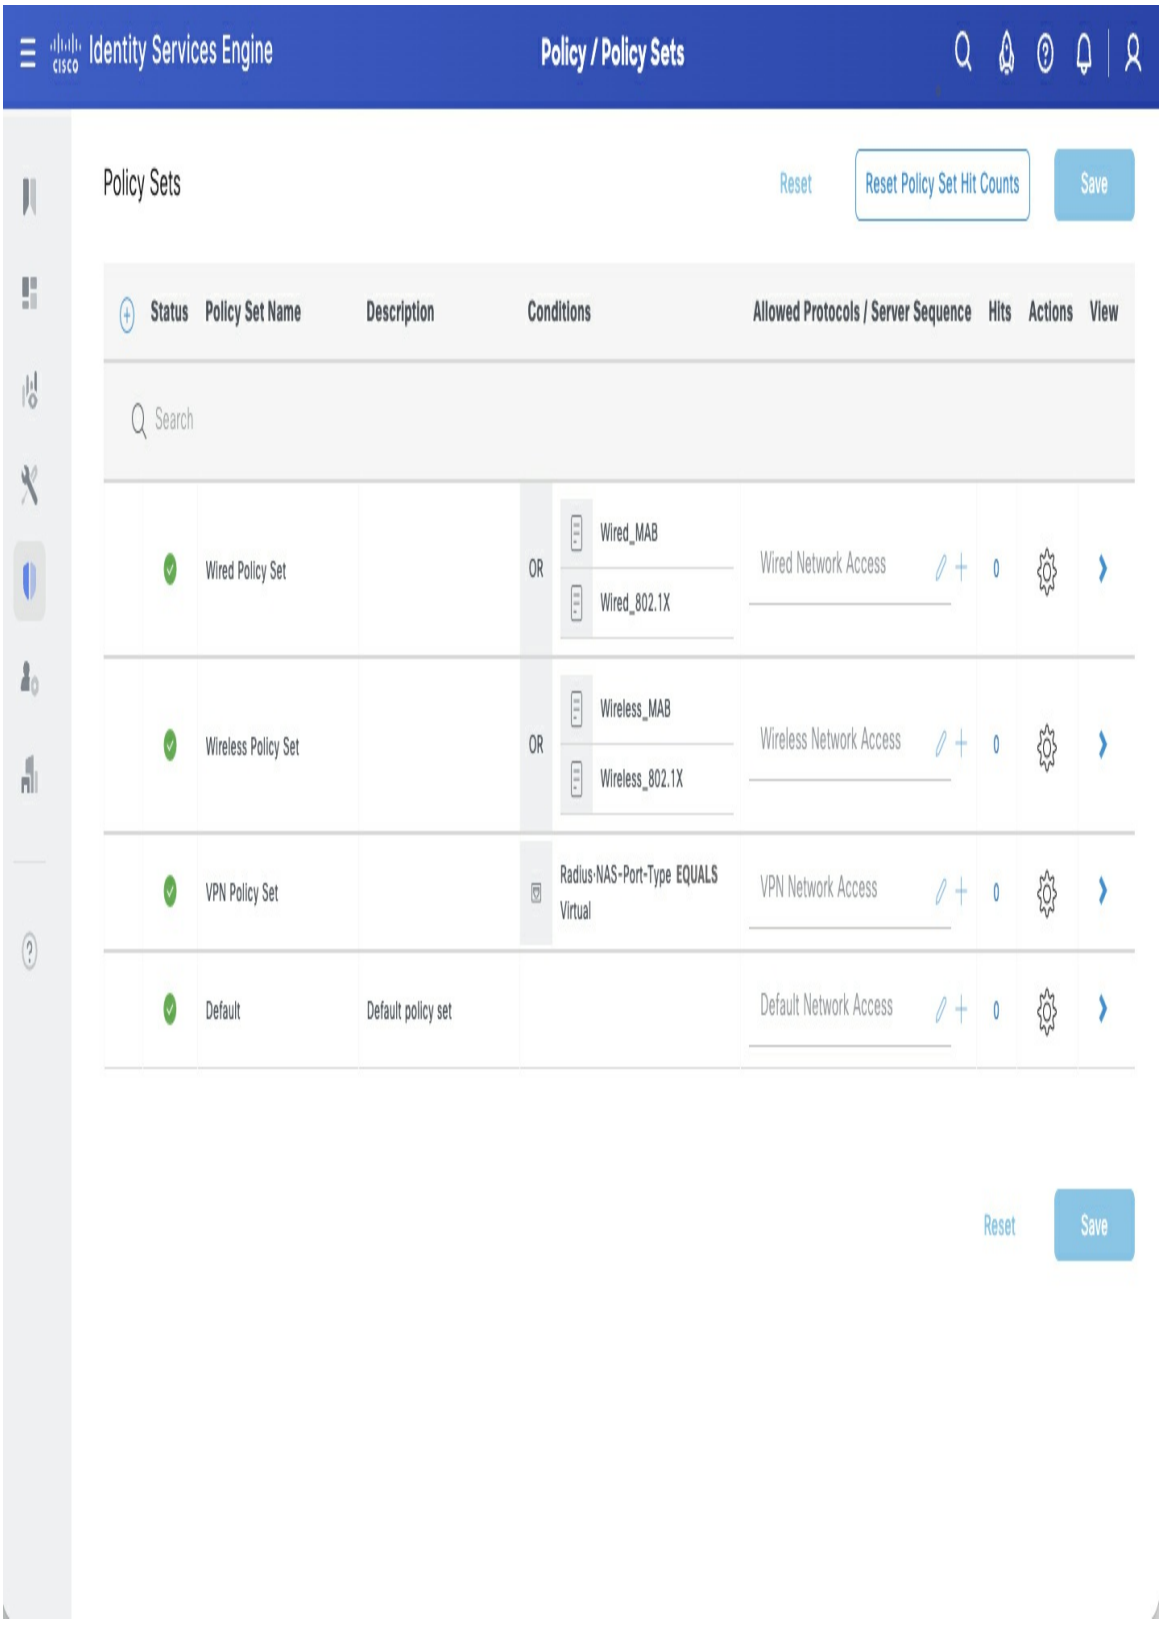


Figure 4-2 *Cisco ISE Policy Set Split Based on Authentication*

Method

Figure 4-3 shows how to distinguish between authentication types.

Technet24

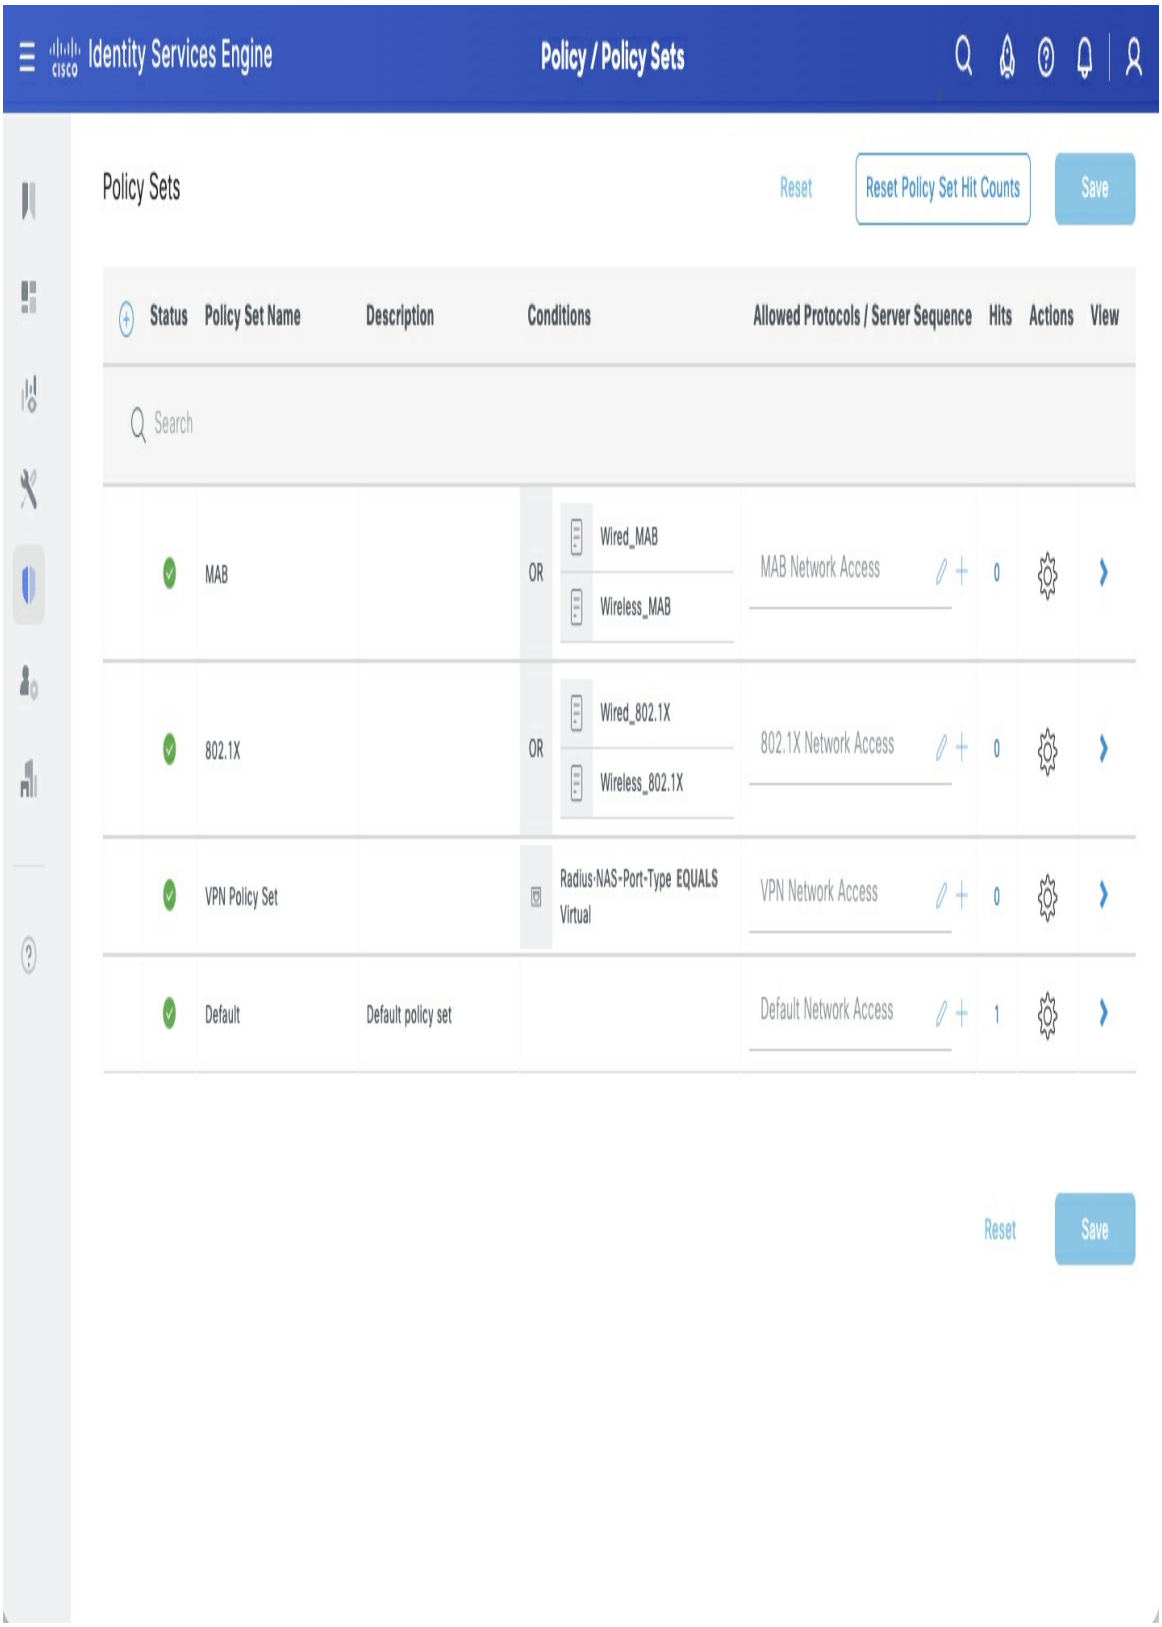


Figure 4-3 *Cisco ISE Policy Set Split Based on Authentication Type*

You can also combine the split based on authentication method and authentication type, as shown in Figure 4-4.

Technet24

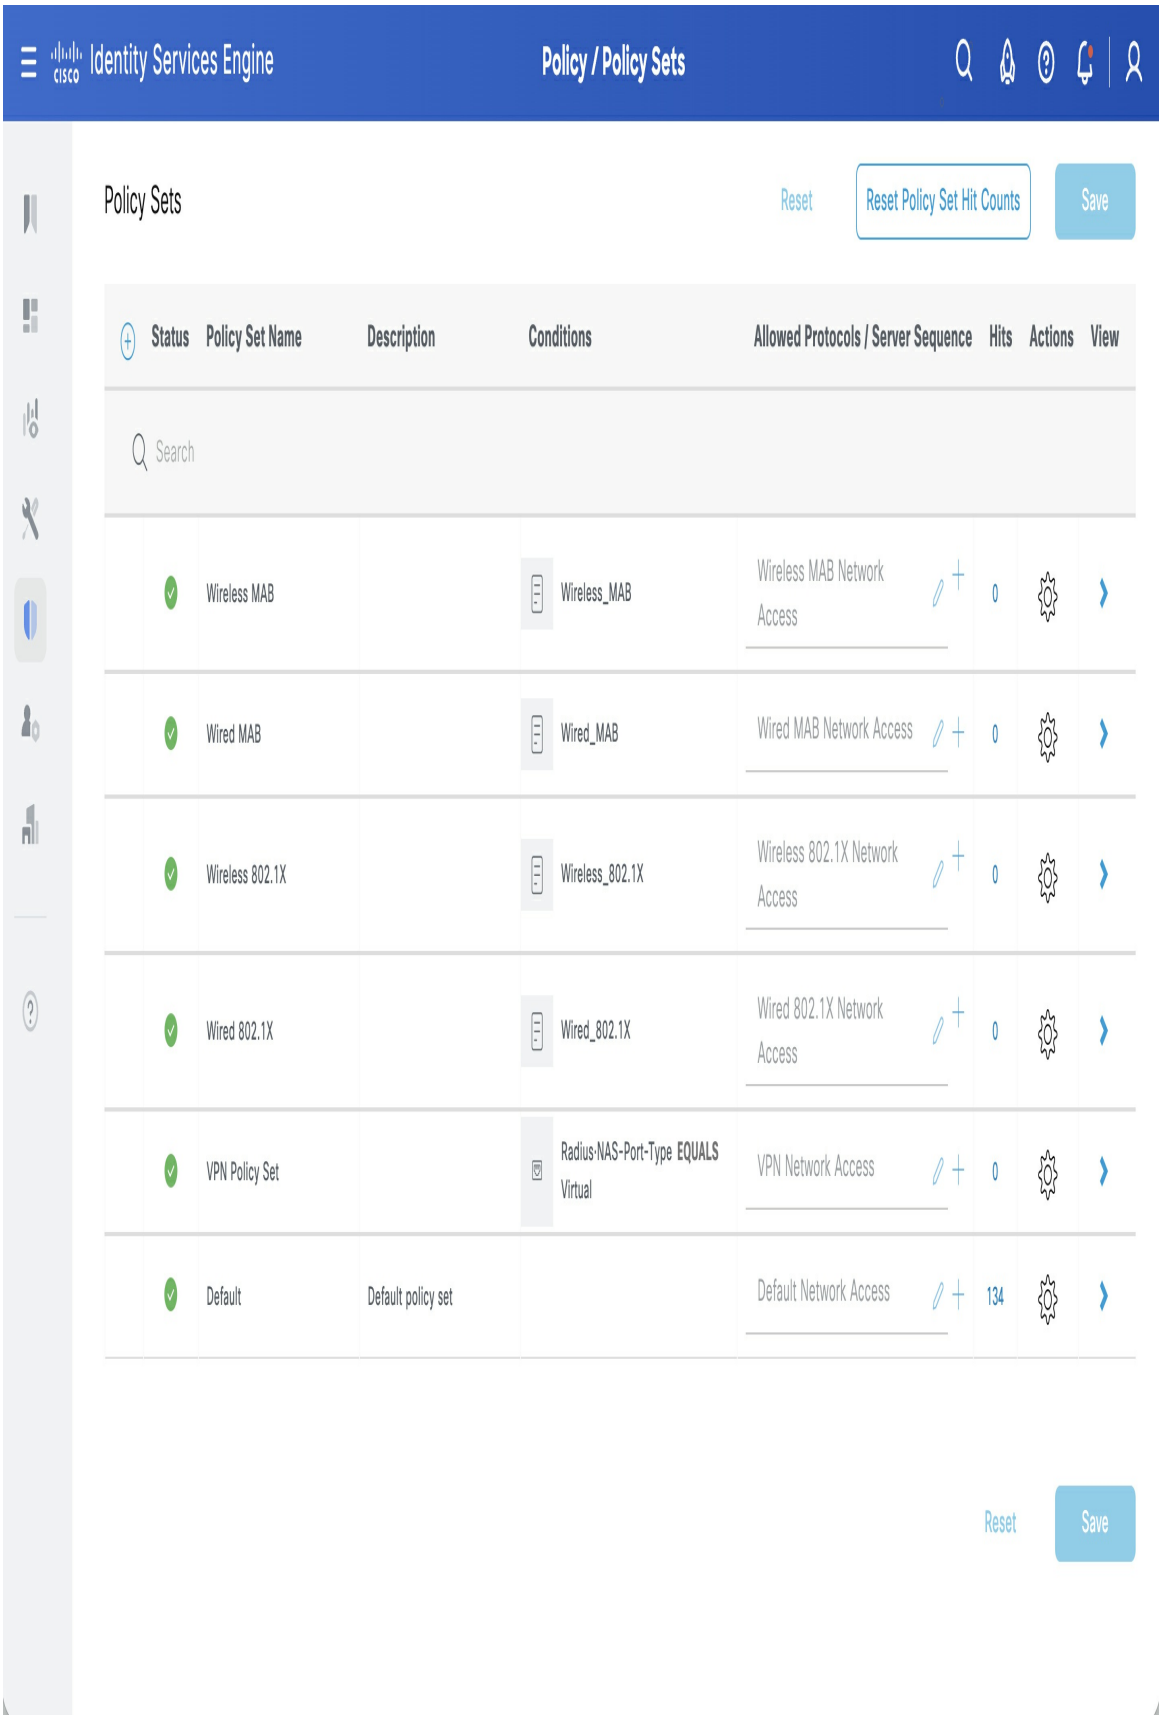

Identity Services Engine
Policy / Policy Sets
Policy Sets
Reset
Reset Policy Set Hit Counts
Save
Status
Policy Set Name
Description
Conditions
Allowed Protocols / Server Sequence
Hits
Actions
View
Search
Wireless MAB
Wireless_MAB
Wireless MAB Network Access
0
Wired MAB
Wired_MAB
Wired MAB Network Access
0
Wireless 802.1X
Wireless_802.1X
Wireless 802.1X Network Access
0
Wired 802.1X
Wired_802.1X
Wired 802.1X Network Access
0
VPN Policy Set
Radius·NAS-Port-Type EQUALS Virtual
VPN Network Access
0
Default
Default policy set
Default Network Access
134
Reset
Save

Figure 4-4 *Cisco ISE Policy Set Split Based on Authentication Method and Type*

Figure 4-5 shows how to distinguish between locations.

Technet24

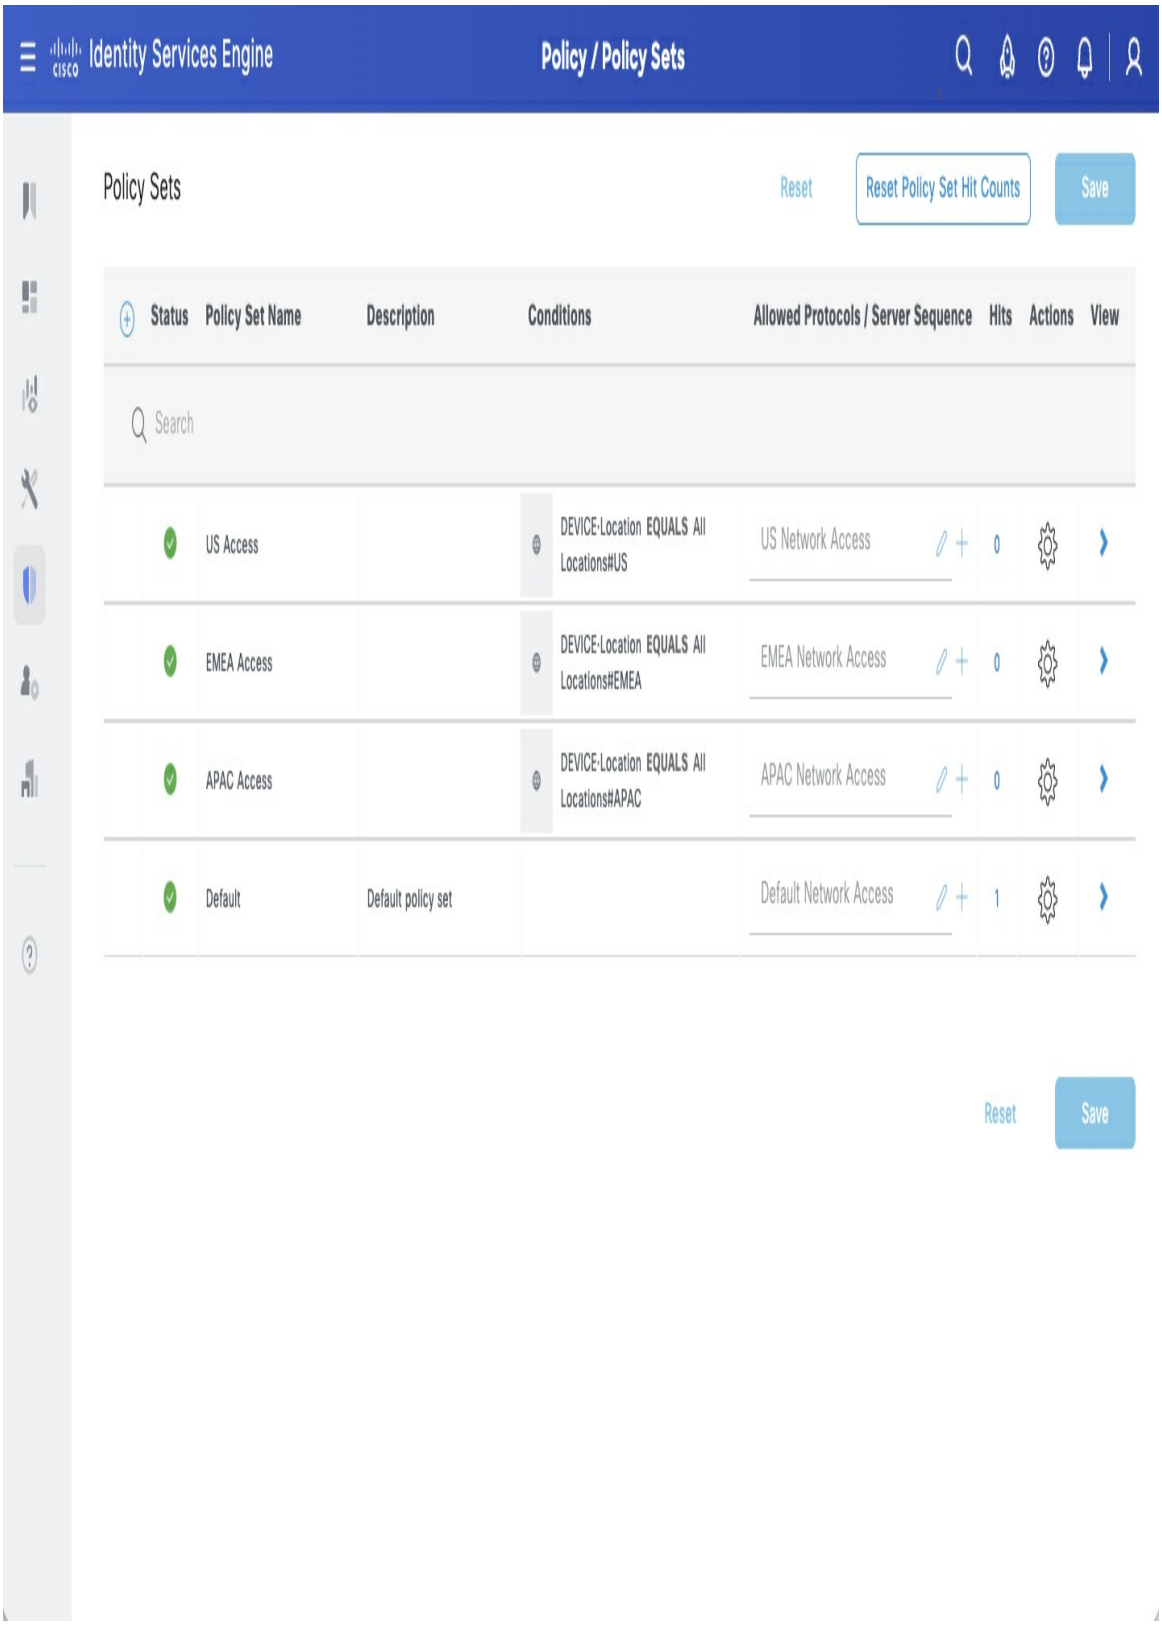


Figure 4-5 *Cisco ISE Policy Set Split Based on Location*

Best Practice:

Always use granular policy sets for easy management of the policies, along with optimal ISE performance.

Conditions Studio

To create a new policy set, you click the gear icon next to existing policy set and select **Insert New Rule Above**. If you want to save a new policy set, it must have a name, condition, and **Allowed Protocols/Server Sequence** configured.

To configure a condition, you, as administrator, click the plus icon within the new policy set definition, as shown in Figure 4-6.

Technet24

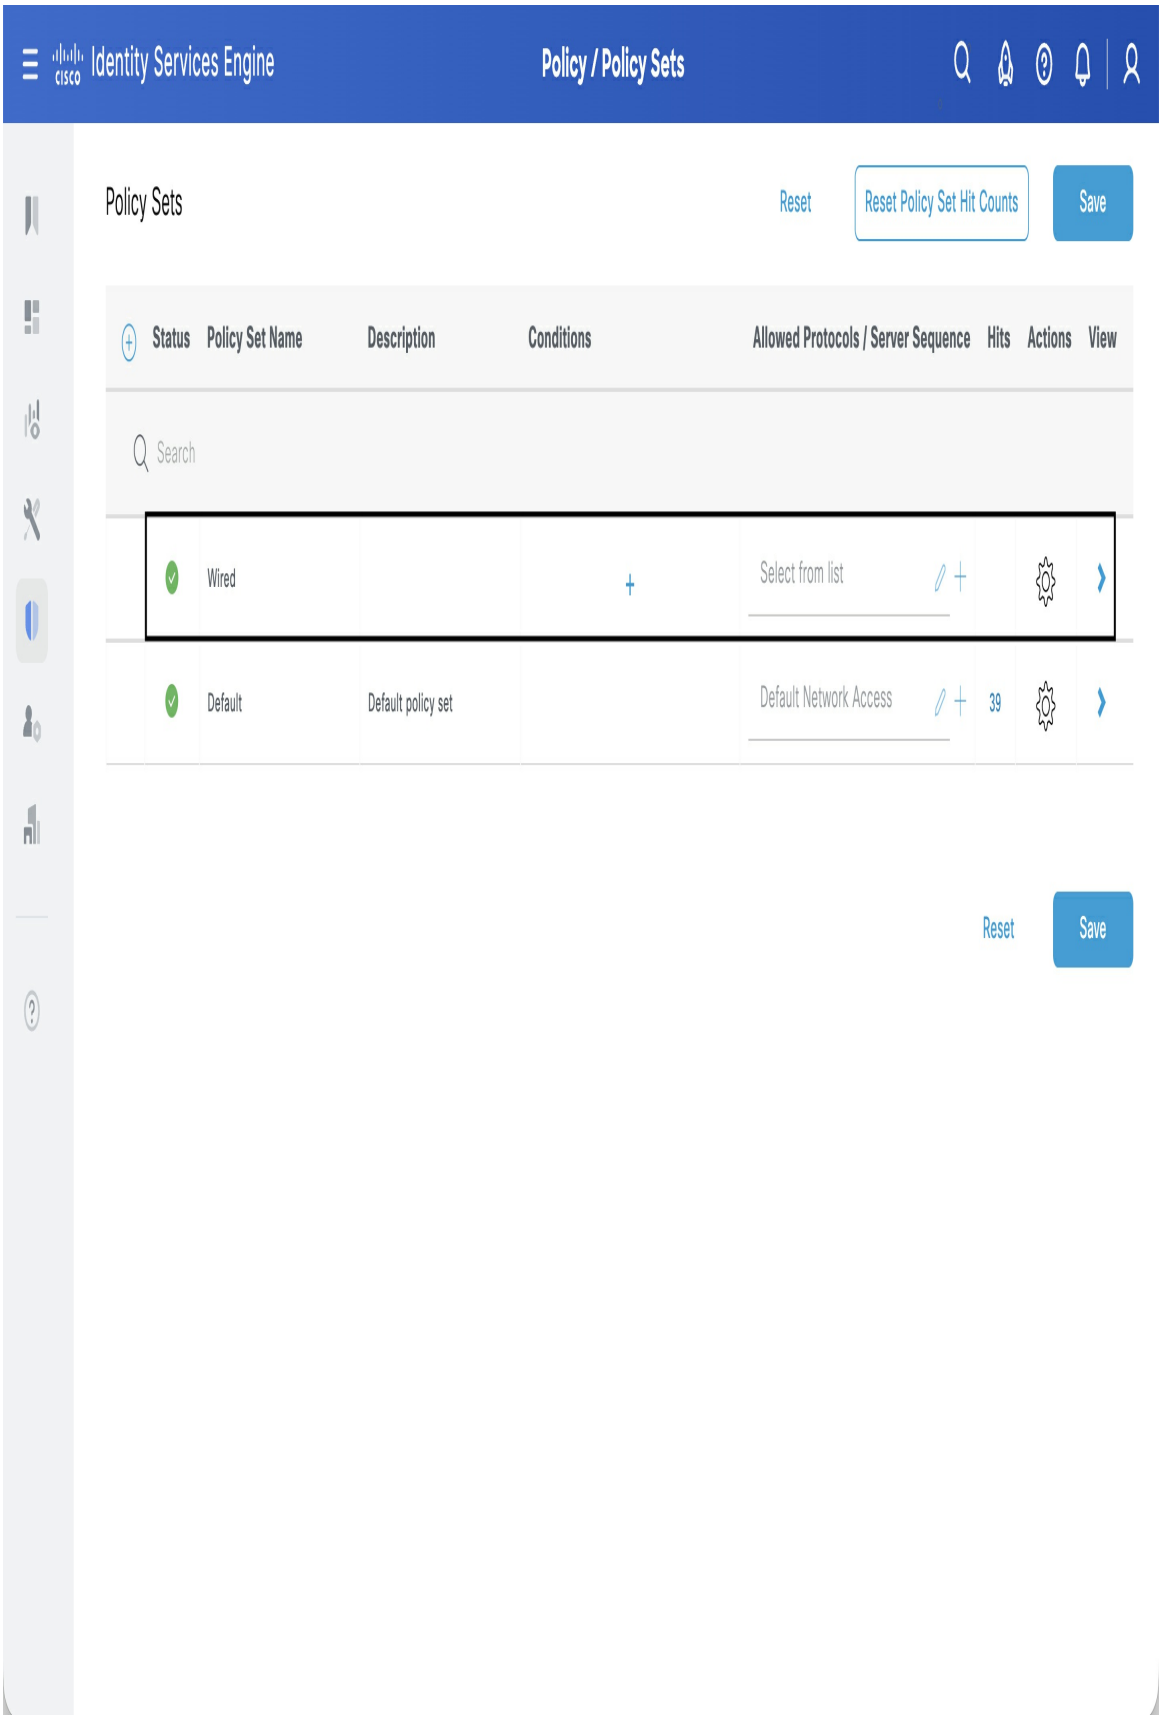

Identity Services Engine
Policy / Policy Sets
Policy Sets
Reset
Reset Policy Set Hit Counts
Save
Status
Policy Set Name
Description
Conditions
Allowed Protocols / Server Sequence
Hits
Actions
View
Search
Wired
Select from list
Default
Default policy set
Default Network Access
39
Reset
Save

Figure 4-6 *New Policy Set*

The Conditions Studio opens. It is one of the most critical tools used to build ISE policies. The Conditions Studio is used to configure

- Policy set conditions
- Authentication policy conditions
- MFA policy conditions
- Authorization policy conditions

The Conditions Studio consists of two parts — Library and Editor — as shown in Figure 4-7.

Technet24

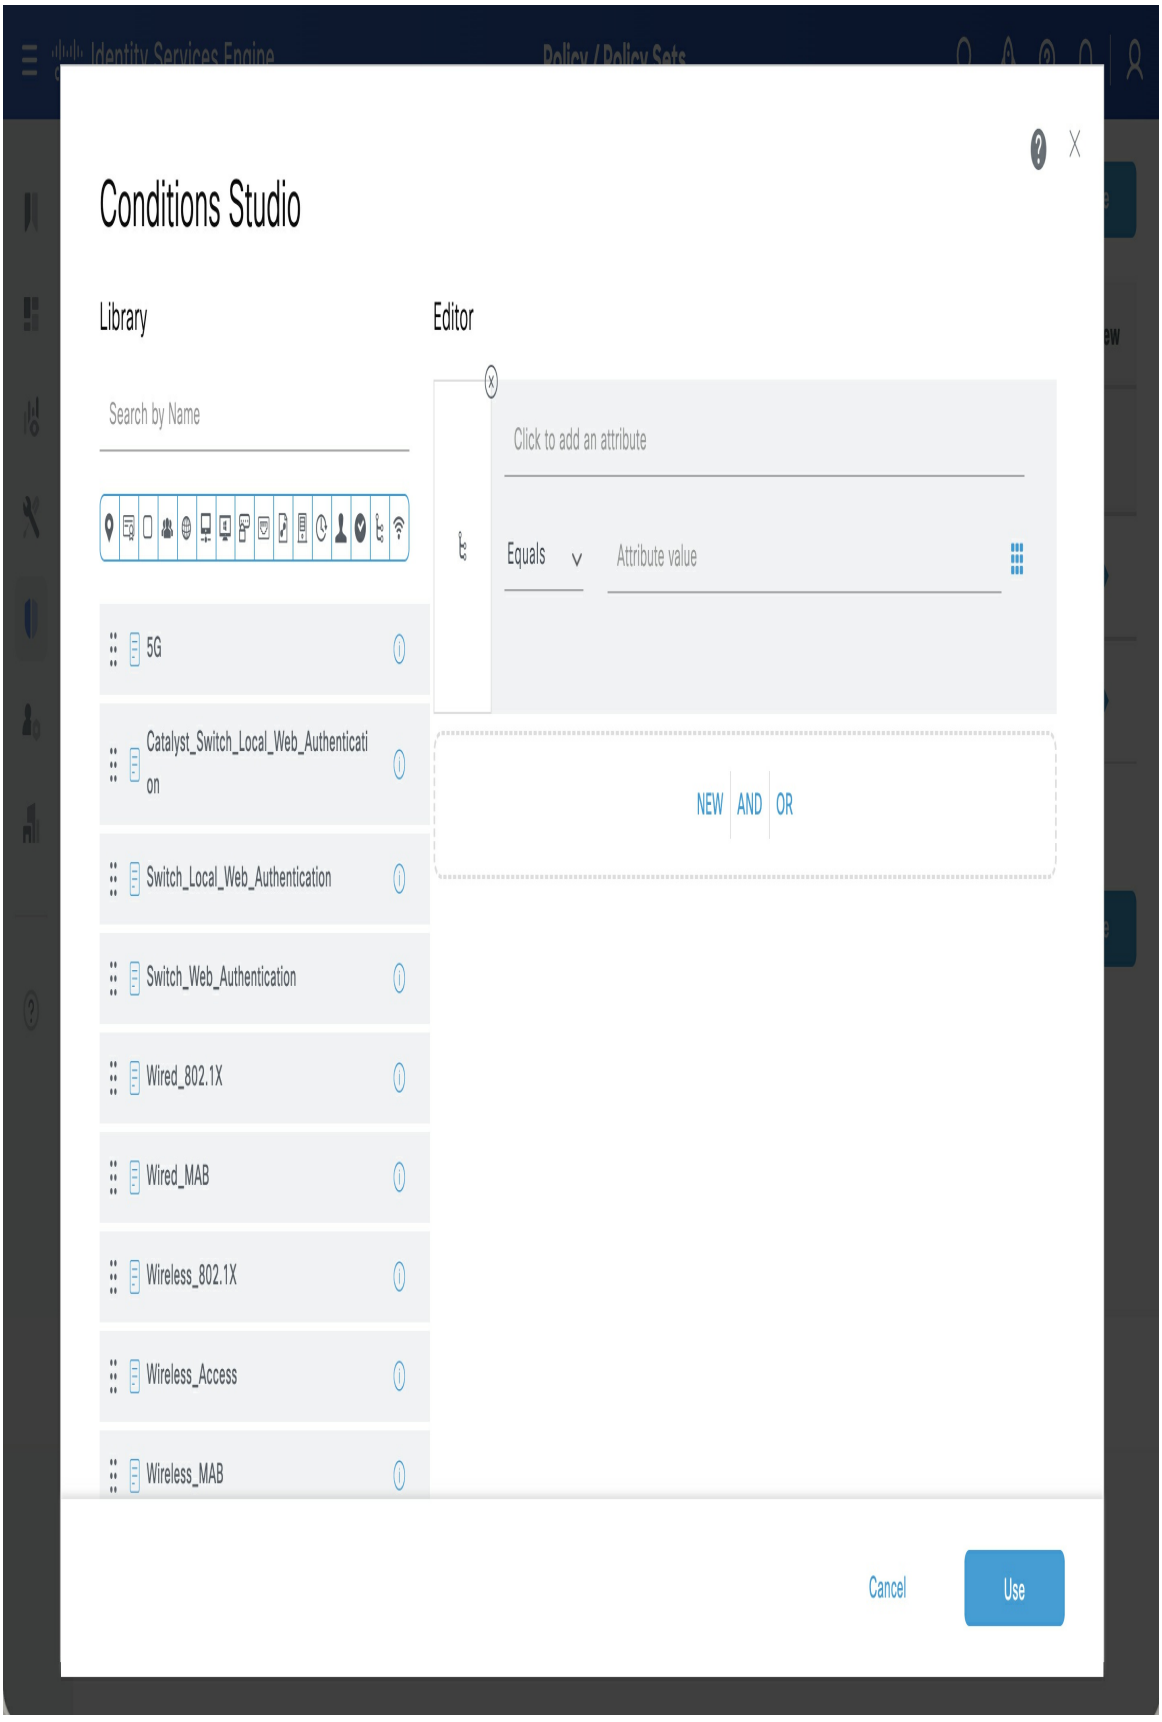

Identity Services Engine
Policy / Policy Sets
Conditions Studio
Library
Editor
Search by Name
5G
Catalyst_Switch_Local_Web_Authentication
Switch_Local_Web_Authentication
Switch_Web_Authentication
Wired_802.1X
Wired_MAB
Wireless_802.1X
Wireless_Access
Wireless_MAB
Click to add an attribute
Equals
Attribute value
NEW
AND
OR
Cancel
Use

Figure 4-7 *Cisco ISE Condition Studio*

The Library allows you, as administrator, to use preconfigured conditions. ISE provides a set of commonly used conditions by default. To use any of the Library conditions, simply drag the condition from the left side of the screen to the Editor and click **Use**.

Using Editor allows granular options to build a condition, using various combinations of logical OR and AND conditions.

To create a condition entry, select an attribute, choose an operator (**Equals** or **Not Equals)** and select an **Attribute Value**.

Attributes on ISE are selected based on dictionaries. ISE has a subset of dictionaries from which administrators choose the attributes, as shown in Figure 4-8. For example, the RADIUS dictionary contains attributes such as Calling Station ID and Service Type.

Best Practice:

ISE operates as a first-match policy engine. You should always use internal conditions before applying any external conditions to prevent unnecessary external lookups.

Technet24

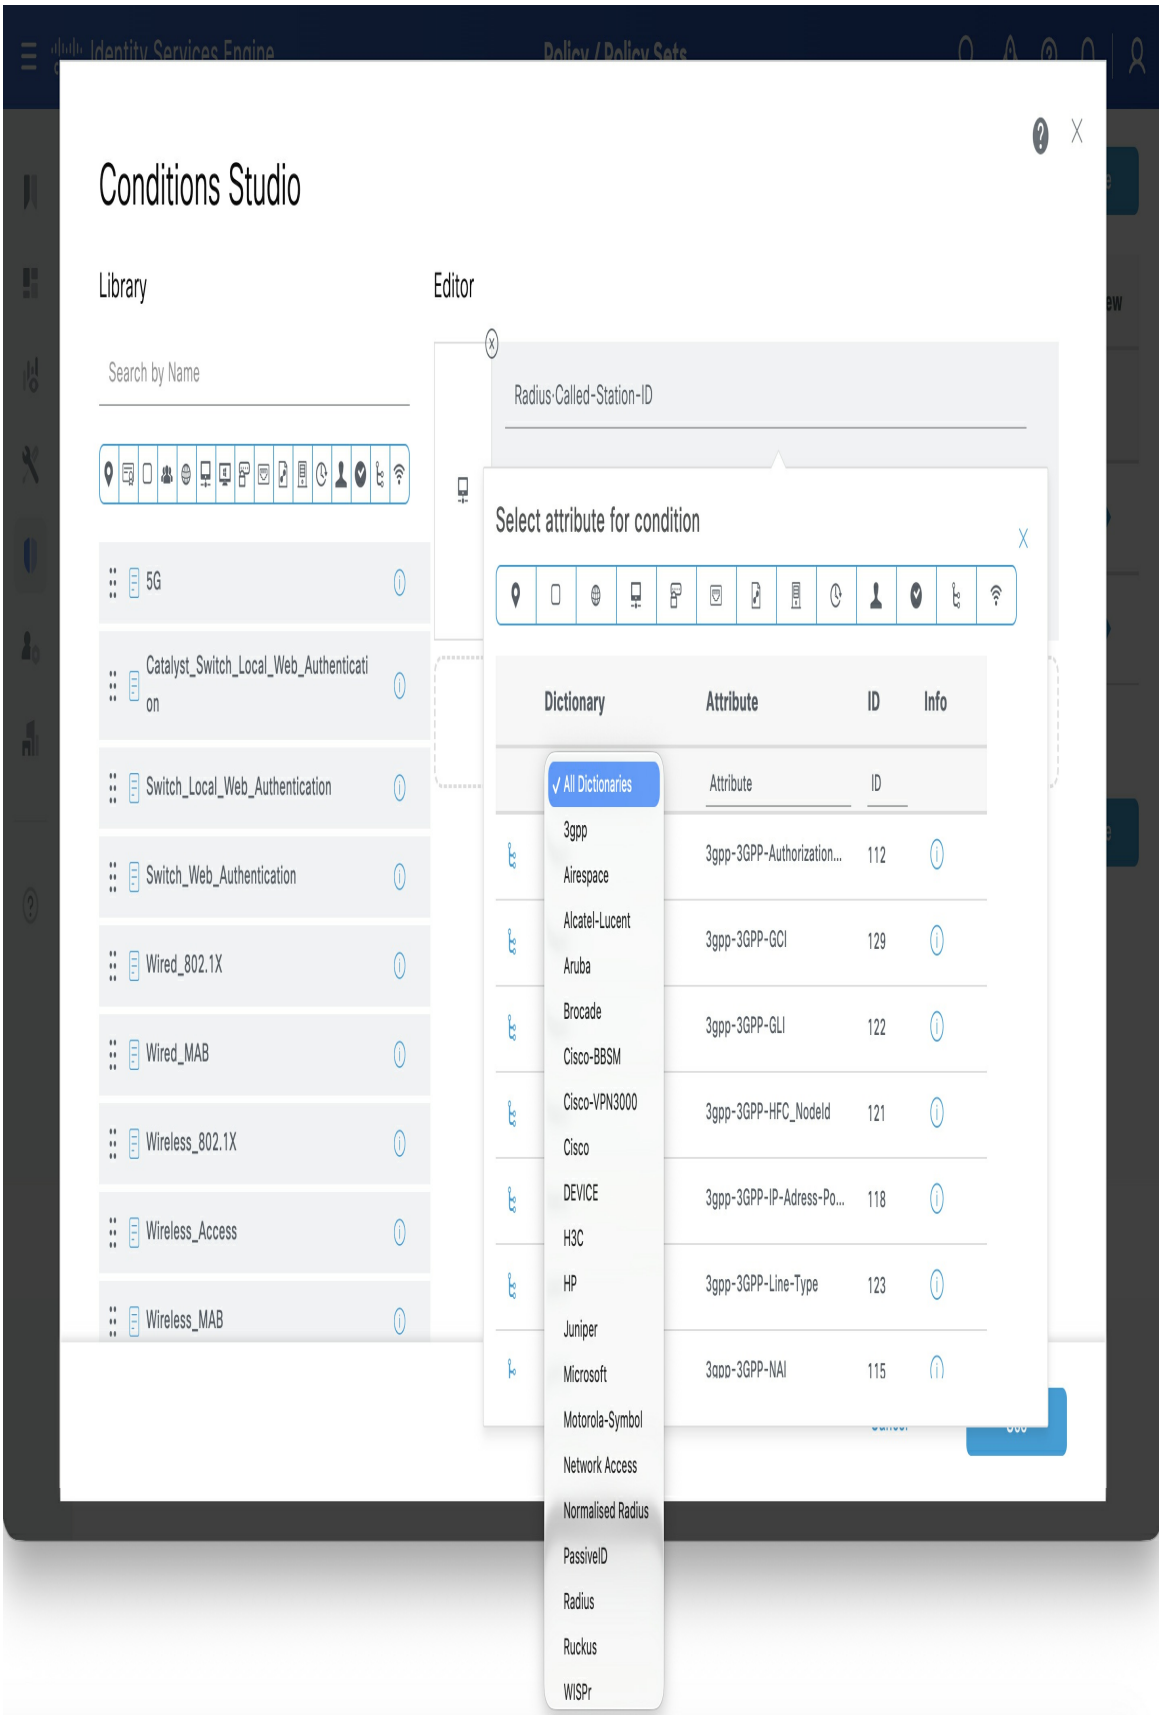

Conditions Studio
Library
Editor
Search by Name
5G
Catalyst_Switch_Local_Web_Authentication
Switch_Local_Web_Authentication
Switch_Web_Authentication
Wired_802.1X
Wired_MAB
Wireless_802.1X
Wireless_Access
Wireless_MAB
Radius·Called-Station-ID
Select attribute for condition
Dictionary
Attribute
ID
Info
All Dictionaries
3gpp
Airespace
Alcatel-Lucent
Aruba
Brocade
Cisco-BBSM
Cisco-VPN3000
Cisco
DEVICE
H3C
HP
Juniper
Microsoft
Motorola-Symbol
Network Access
Normalised Radius
PassiveID
Radius
Ruckus
WISPr
3gpp-3GPP-Authorization...
112
3gpp-3GPP-GCI
129
3gpp-3GPP-GLI
122
3gpp-3GPP-HFC_NodeId
121
3gpp-3GPP-IP-Adress-Po...
118
3gpp-3GPP-Line-Type
123
3gpp-3GPP-NAI
115

Figure 4-8 *Cisco ISE Conditions Studio: Attribute Selection*

For example, if you want to apply the policy set to the Network Device Group, you can select **Device Type Equals All Device Types**, as shown in Figure 4-9.

The Editor allows you to save newly created conditions to the Library by clicking the **Save** button. You can later reuse these conditions when they're needed.

Technet24

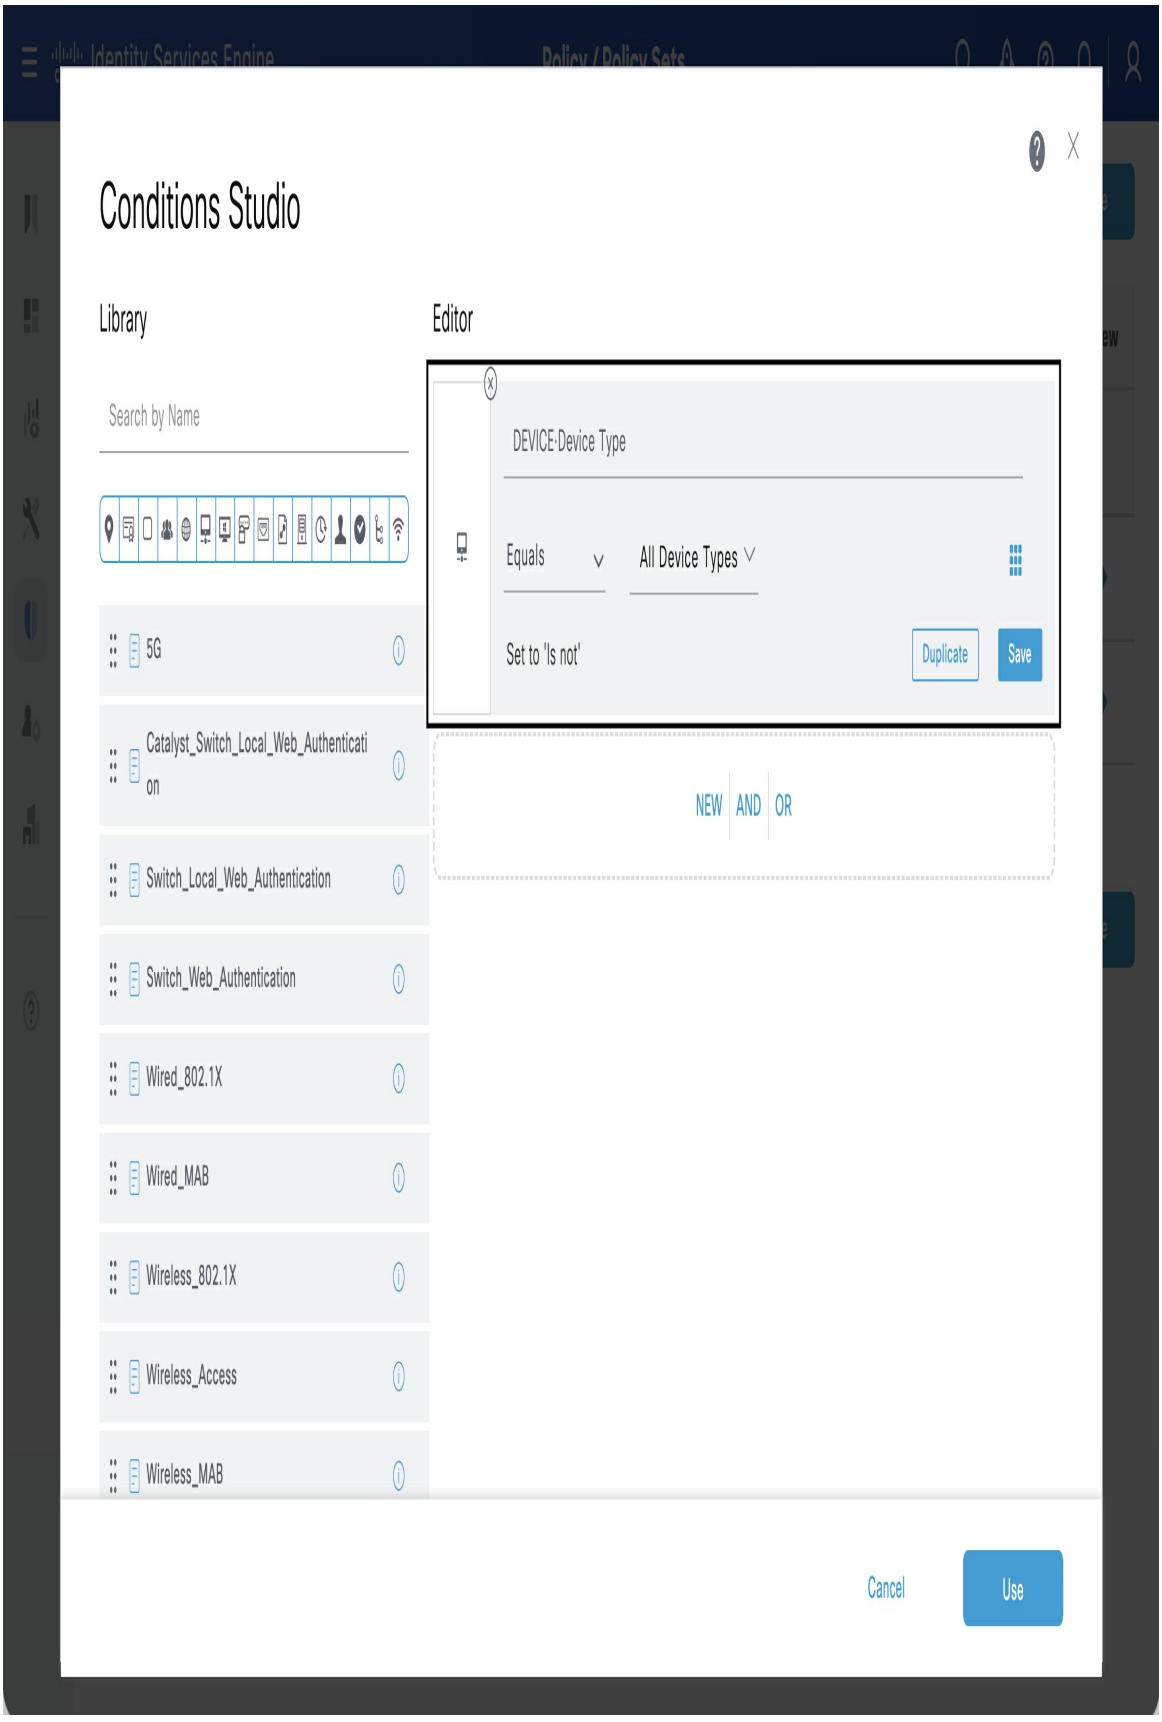

Identity Services Engine
Policy / Policy Sets
Conditions Studio
Library
Search by Name
5G
Catalyst_Switch_Local_Web_Authentication
Switch_Local_Web_Authentication
Switch_Web_Authentication
Wired_802.1X
Wired_MAB
Wireless_802.1X
Wireless_Access
Wireless_MAB
Editor
DEVICE·Device Type
Equals
All Device Types
Set to 'Is not'
Duplicate
Save
NEW
AND
OR
Cancel
Use

Figure 4-9 *Cisco ISE Condition Studio: Device Type Condition*

As administrator, you can also select the **Set to 'Is Not'** option, which would make this condition negative.

Administrators can combine AND and OR conditions within the same condition. Figure 4-10 illustrates a scenario where the policy is matched if user bob OR user alice is connected AND they are connected using Wired 802.1X authentication.

Technet24


Identity Services Engine
Policy / Policy Sets
Conditions Studio
Library
Search by Name
5G
Catalyst_Switch_Local_Web_Authentication
Switch_Local_Web_Authentication
Switch_Web_Authentication
Wired_802.1X
Wired_MAB
Wireless_802.1X
Wireless_Access
Wireless_MAB
Editor
Radius·User-Name
Equals
bob
Set to 'Is not'
Duplicate
Save
OR
Radius·User-Name
Equals
alice
Set to 'Is not'
Duplicate
Save
AND
NEW
AND
OR
Set to 'Is not'
Duplicate
Save
Wired_802.1X
Cancel
Use

Figure 4-10 *Cisco ISE Conditions Studio: Combining AND and OR Conditions*

Figure 4-11 shows how such a policy looks in the Policy Sets menu.

Technet24

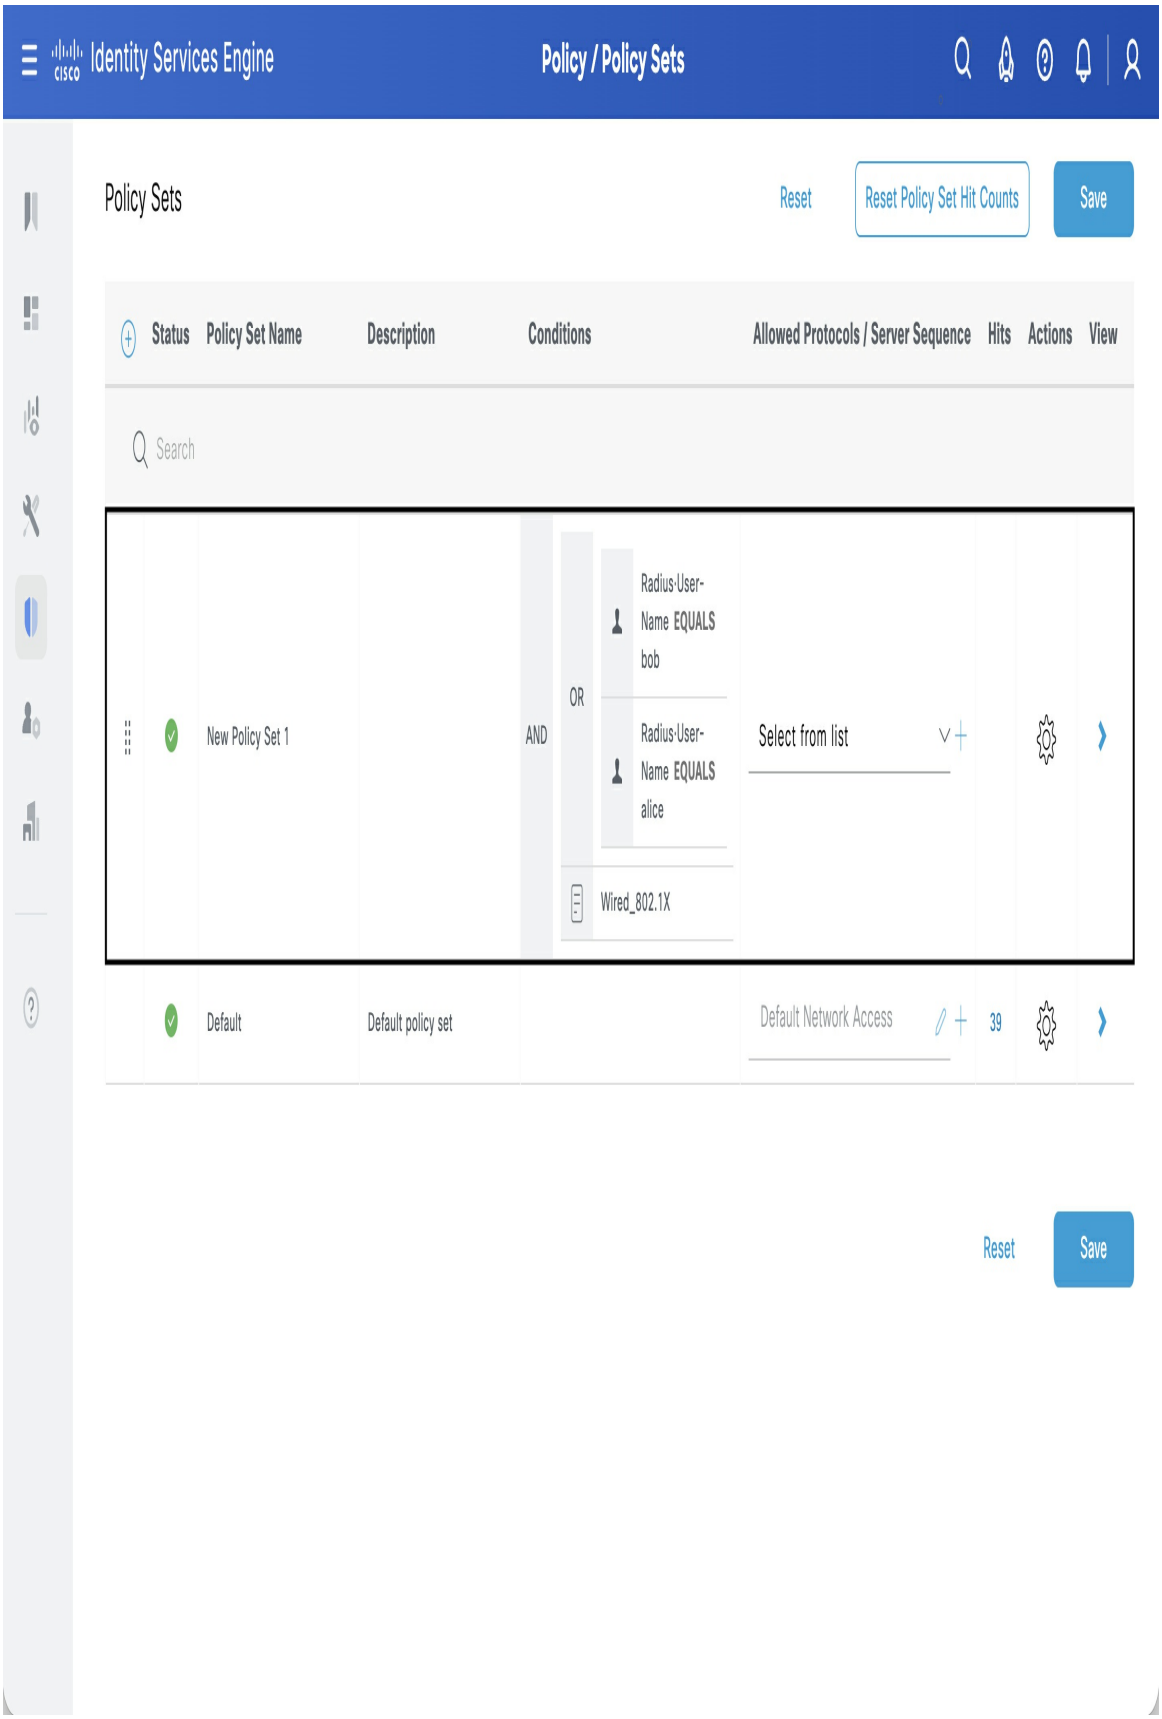

Identity Services Engine
Policy / Policy Sets
Policy Sets
Reset
Reset Policy Set Hit Counts
Save
Status
Policy Set Name
Description
Conditions
Allowed Protocols / Server Sequence
Hits
Actions
View
Search
New Policy Set 1
AND
OR
Radius·User-Name EQUALS bob
Radius·User-Name EQUALS alice
Wired_802.1X
Select from list
Default
Default policy set
Default Network Access
39
Reset
Save

Figure 4-11 *New Policy Set Created with Condition Studio*

RADIUS Server Sequence

RADIUS Server Sequence is used if you, as administrator, want proxy authentication to an upstream RADIUS server instead of ISE processing it. ISE forwards RADIUS packets received from network access devices (NADs) to the upstream RADIUS server, this server replies, and ISE forwards the replies to the NAD. Optionally, accounting can also be proxied to the upstream RADIUS server.

Configure RADIUS Server Sequence

To configure a RADIUS Server Sequence, you must first define one or more external RADIUS servers. Navigate to **Administration > Network Resources > External RADIUS Servers** (see Figure 4-12). Then click **Add**.

Technet24

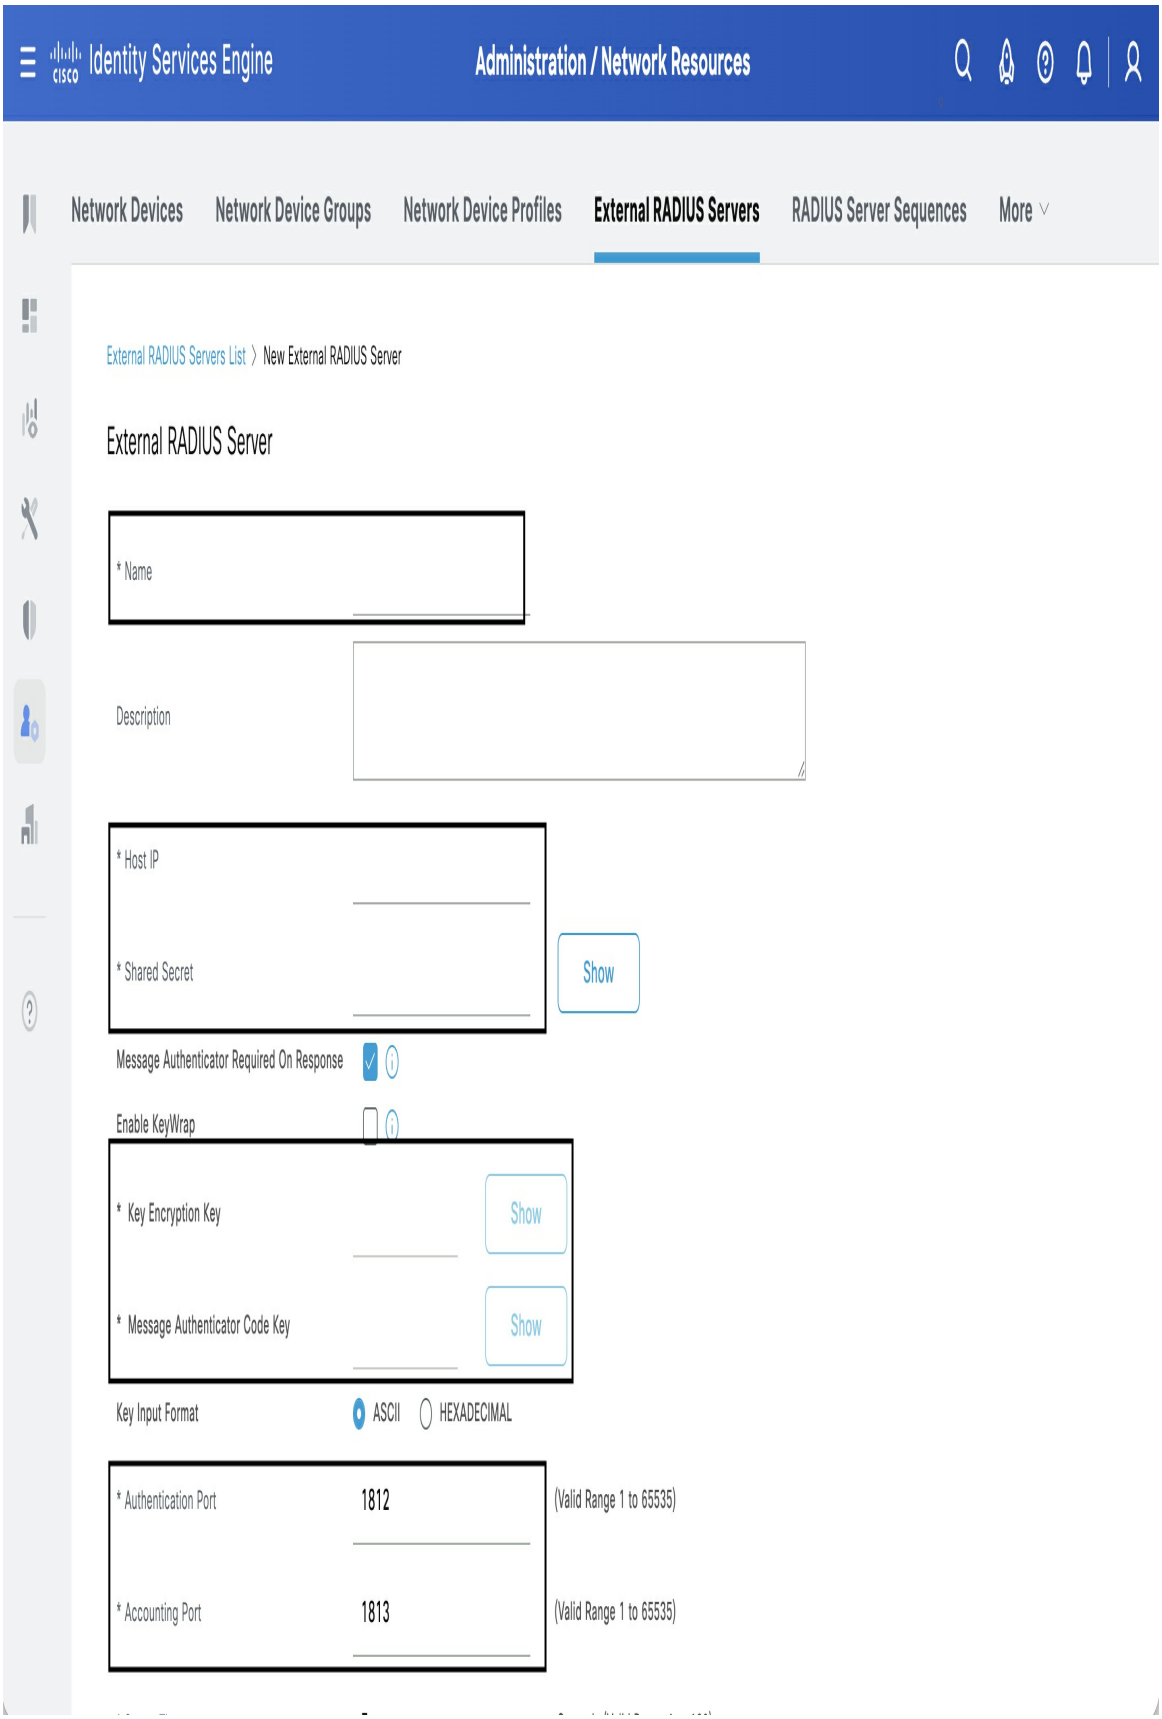

Identity Services Engine
Administration / Network Resources
Network Devices
Network Device Groups
Network Device Profiles
External RADIUS Servers
RADIUS Server Sequences
More
External RADIUS Servers List 〉 New External RADIUS Server
External RADIUS Server
* Name
Description
* Host IP
* Shared Secret
Show
Message Authenticator Required On Response
Enable KeyWrap
* Key Encryption Key
Show
* Message Authenticator Code Key
Show
Key Input Format
ASCII
HEXADECIMAL
* Authentication Port
1812
(Valid Range 1 to 65535)
* Accounting Port
1813
(Valid Range 1 to 65535)

Figure 4-12 *New External RADIUS Server*

Settings that are required to configure external RADIUS servers are explained in Table 4-1.

Table 4-1 *External RADIUS Server Settings*

Technet24

Field Name	Usage Guidelines
Name	(Mandatory) Enter a unique name for this external RADIUS server.
Description	(Optional) Enter a unique description for the external RADIUS server.
Host IP	(Mandatory) Set the IP address of the external RADIUS server.
Shared Secret	(Mandatory) Set the shared secret key for the external RADIUS server; it must match on the external server.
Message Authenticator	(Optional) If the checkbox is enabled, ISE requires you to include the Message-Authenticator attribute in the response from the external RADIUS server. If it is not included, the response will be ignored.
Enable KeyWrap	(Optional) If the checkbox is enabled, ISE enables the KeyWrap for RADIUS requests.
Key Encryption Key	(Mandatory if KeyWrap is enabled) Set to encrypt keying material with AES encryption.
Message Authenticator Code Key	(Mandatory if KeyWrap is enabled) Select to ensure integrity and authenticity of RADIUS messages.
Key Input Format	(Mandatory if KeyWrap is enabled) Select in which format the keys are entered—ASCII or HEXADECIMAL.
Authentication Port	(Mandatory) Set the RADIUS port used for authentication.
Accounting Port	(Mandatory) Set the RADIUS port used for accounting.
Server Timeout	(Mandatory) Set the timeout in seconds before the next RADIUS request is sent.
Connection Attempts	(Mandatory) Set the number of attempts before the server times out.
Radius ProxyFailover Expiration	(Mandatory) If the server is timed out based on a combination of Server Timeout and Connection Attempts, no new authentications will be sent to the server until the configured expiration timer expires. Configure the expiration timer in seconds.

External RADIUS servers cannot be used on their own in the policy sets; instead, they should be placed within a RADIUS server sequence, even if only one is in use.

Navigate to **Administration > Network Resources > RADIUS Server Sequence** (see Figure 4-13). Then click **Add**.

Technet24

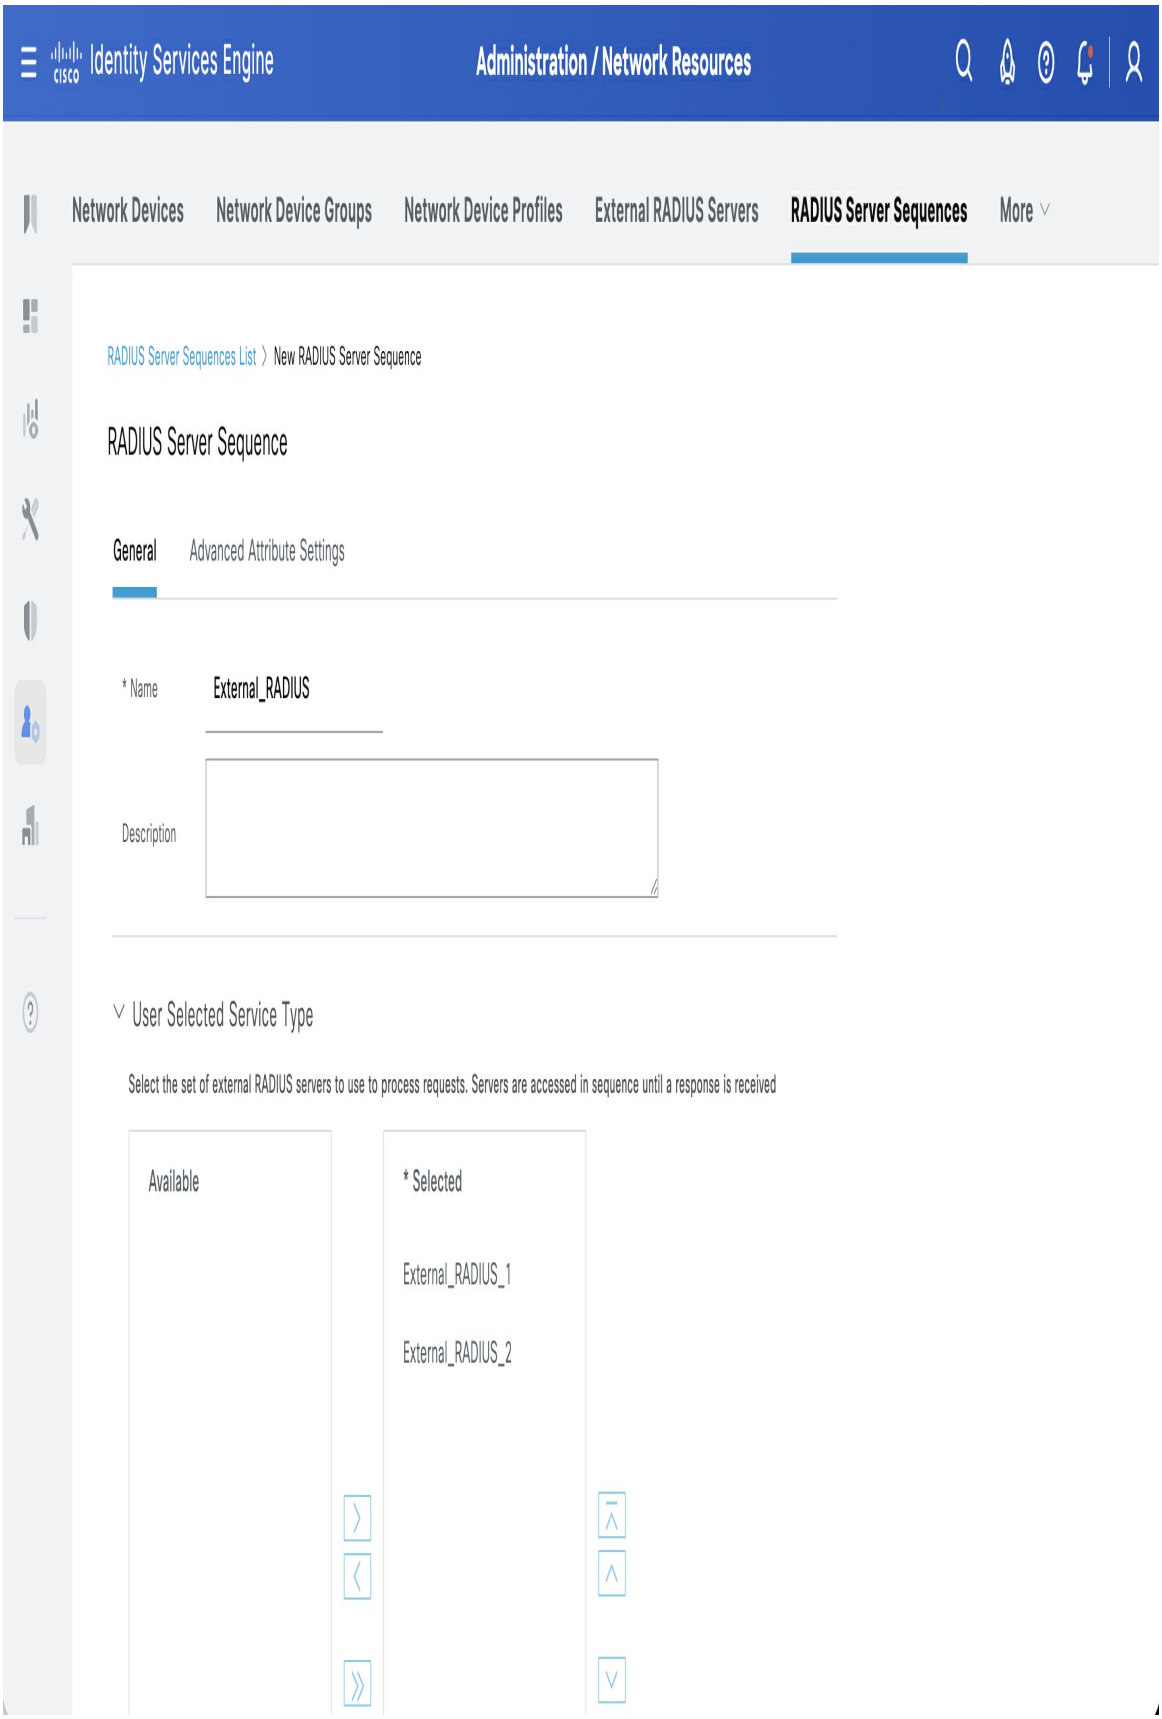

Identity Services Engine
Administration / Network Resources
Network Devices
Network Device Groups
Network Device Profiles
External RADIUS Servers
RADIUS Server Sequences
More
RADIUS Server Sequences List 〉 New RADIUS Server Sequence
RADIUS Server Sequence
General
Advanced Attribute Settings
* Name
External_RADIUS
Description
User Selected Service Type
Select the set of external RADIUS servers to use to process requests. Servers are accessed in sequence until a response is received
Available
* Selected
External_RADIUS_1
External_RADIUS_2

Figure 4-13 *New External RADIUS Server Sequence*

Give a unique name to the RADIUS server sequence and meaningful description (which is optional).

Under **User Selected Service Type**, move desired external RADIUS servers from the **Available** column to the **Selected** column.

Configure Advanced Attribute Settings

ISE allows certain modifications to requests sent to the external RADIUS servers.

As administrator, you can strip the start or the end of the subject name based on a configured separator — for example, removing the domain name from the subject and sending just the username to the external server.

You can also modify the attributes sent in the request to an upstream RADIUS server as well as in an Access-Accept back to the NAD. ISE allows you to add, remove, and modify attributes, as shown in Figure 4-14.

Technet24

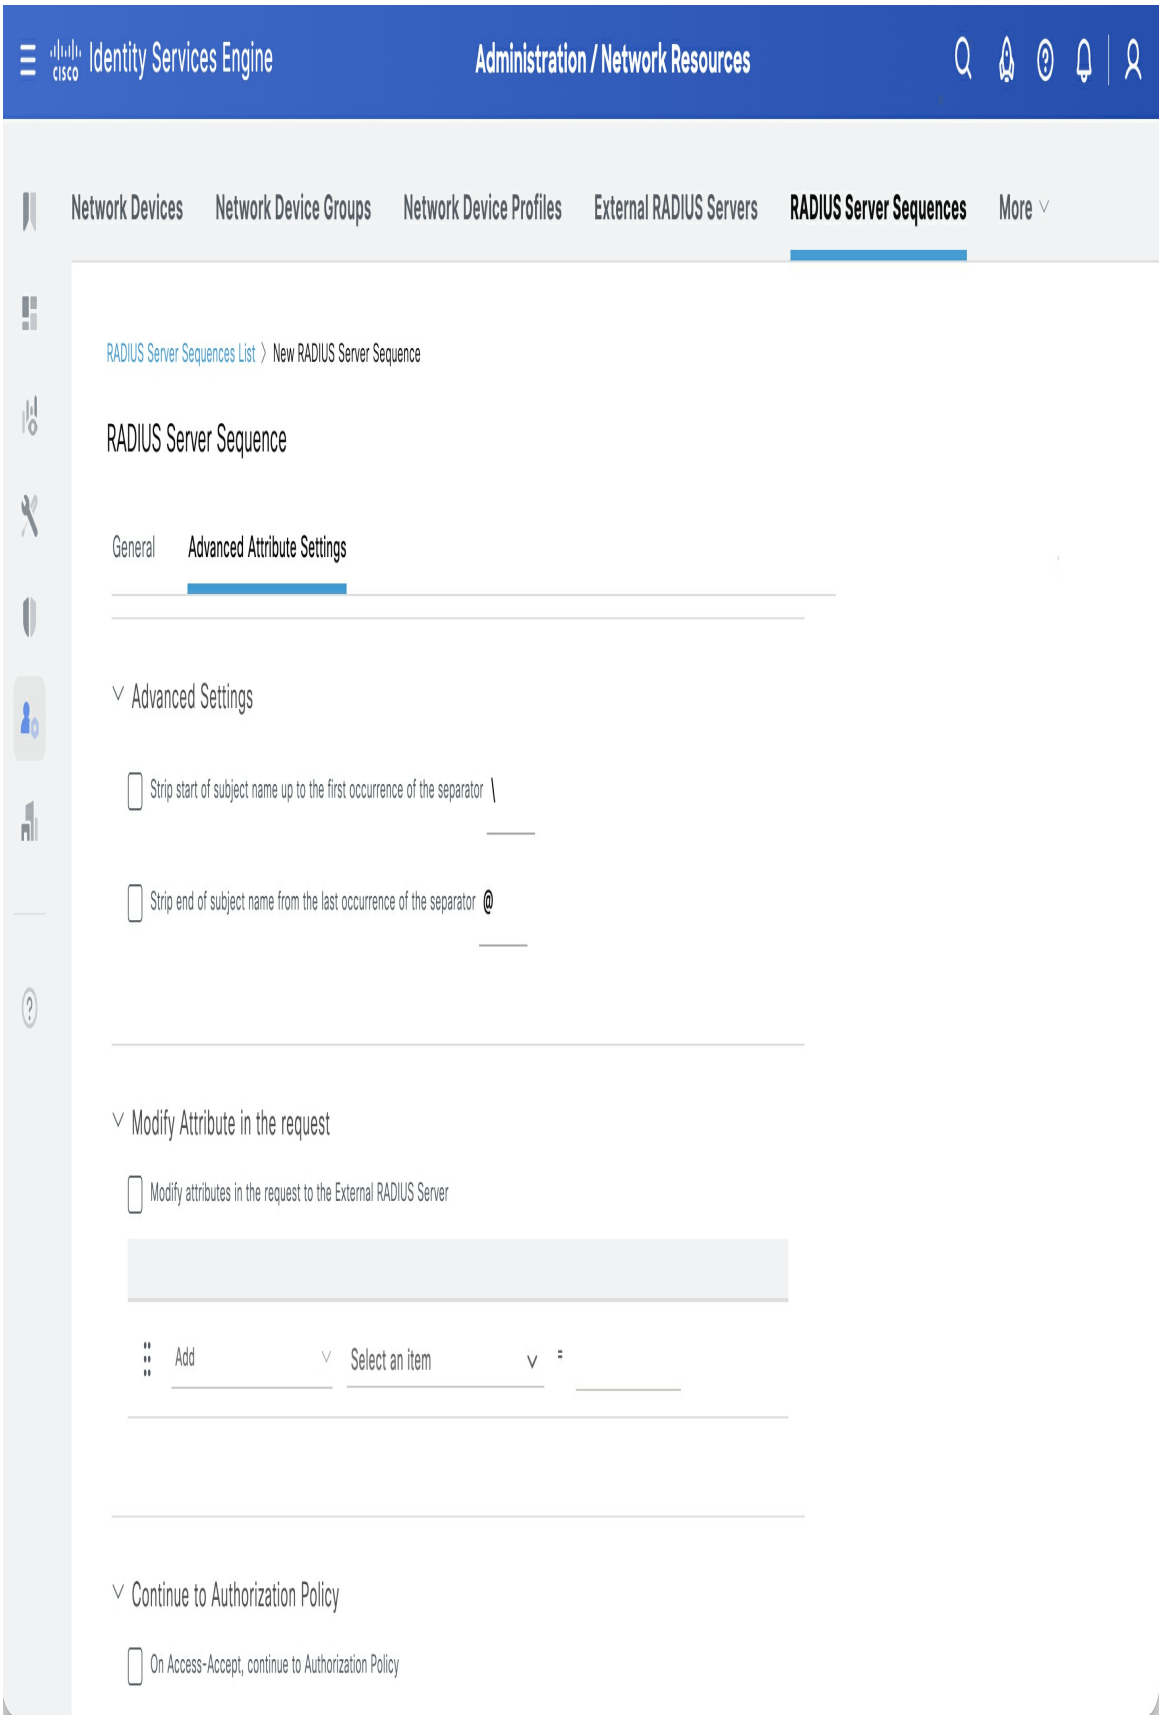

Identity Services Engine
Administration / Network Resources
Network Devices
Network Device Groups
Network Device Profiles
External RADIUS Servers
RADIUS Server Sequences
More
RADIUS Server Sequences List > New RADIUS Server Sequence
RADIUS Server Sequence
General
Advanced Attribute Settings
Advanced Settings
Strip start of subject name up to the first occurrence of the separator \
Strip end of subject name from the last occurrence of the separator @
Modify Attribute in the request
Modify attributes in the request to the External RADIUS Server
Add
Select an item
Continue to Authorization Policy
On Access-Accept, continue to Authorization Policy

Figure 4-14 *External RADIUS Server Sequence: Advanced Attribute Settings*

Finally, on top of attribute modification, you can also proceed to authorization to apply policies configured on ISE itself, as shown in Figure 4-15.

Technet24

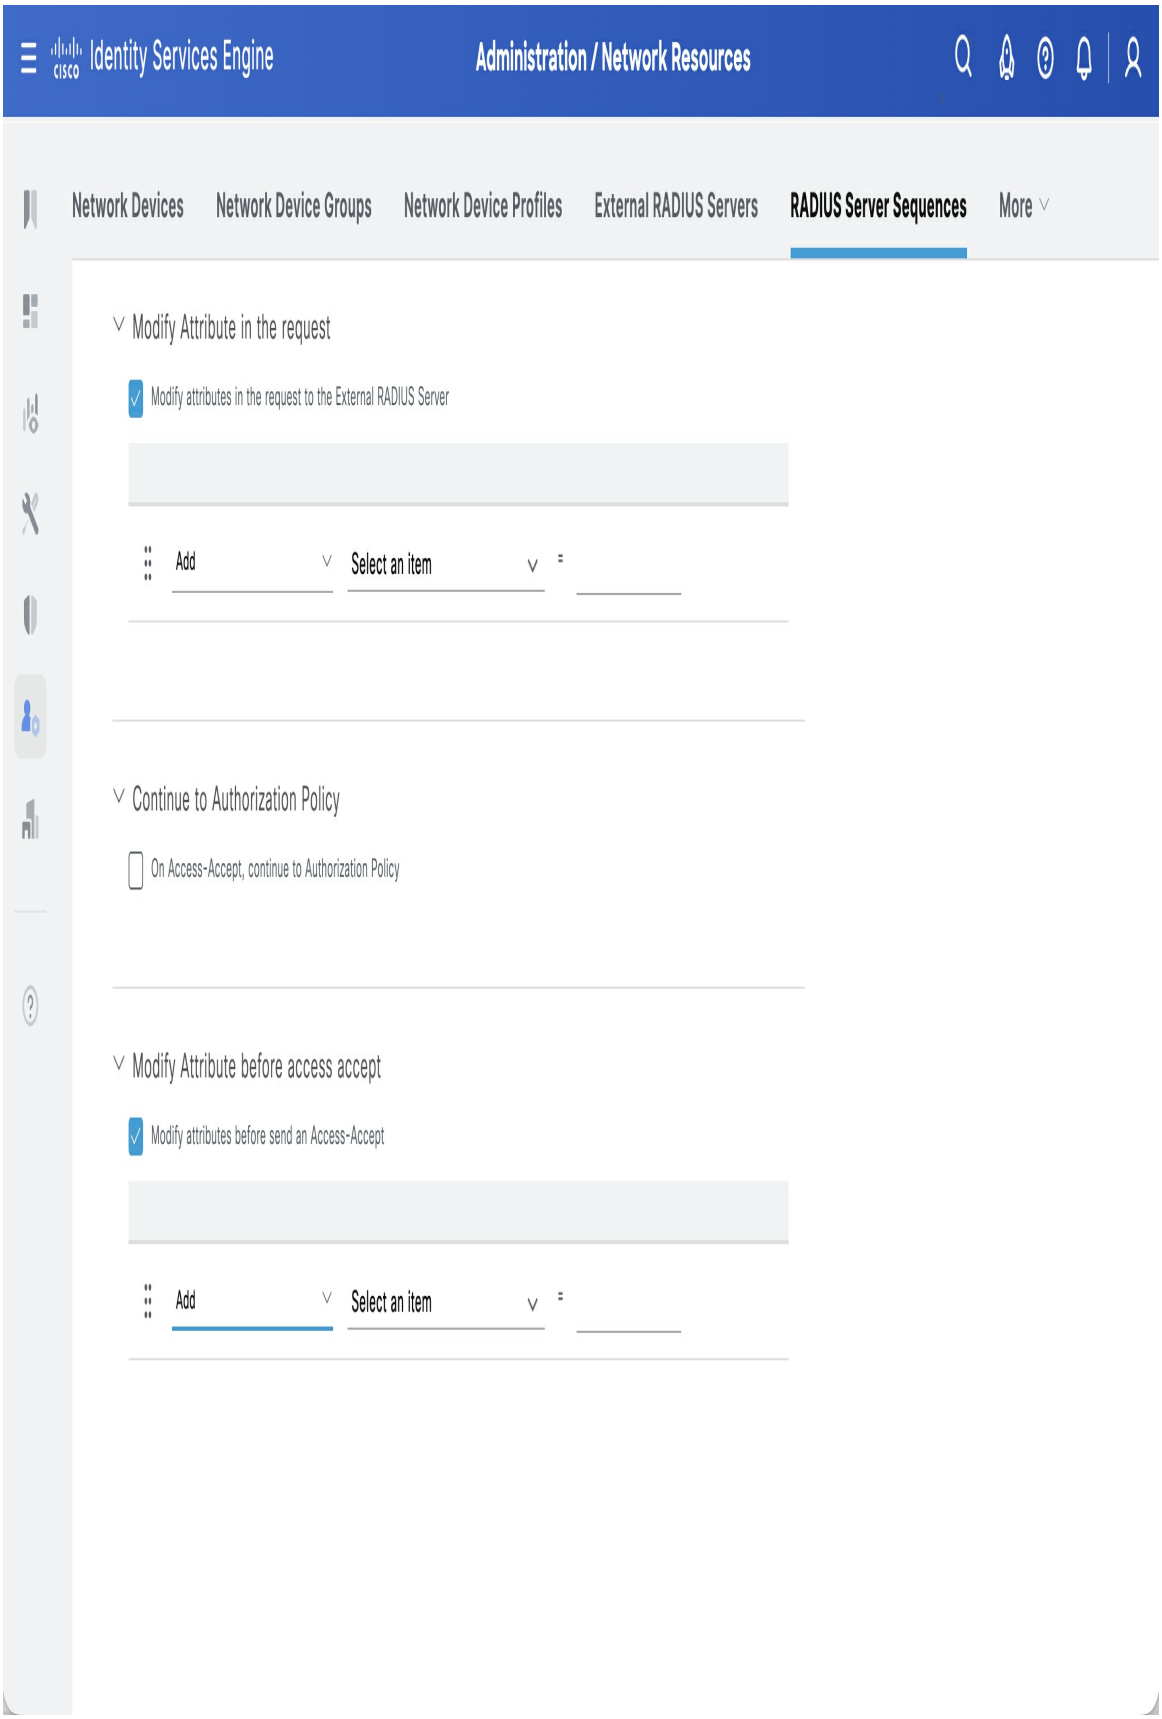

Identity Services Engine
Administration / Network Resources
Network Devices
Network Device Groups
Network Device Profiles
External RADIUS Servers
RADIUS Server Sequences
More
Modify Attribute in the request
Modify attributes in the request to the External RADIUS Server
Add
Select an item
Continue to Authorization Policy
On Access-Accept, continue to Authorization Policy
Modify Attribute before access accept
Modify attributes before send an Access-Accept
Add
Select an item

Figure 4-15 *External RADIUS Server Sequence: Advanced Attribute Settings Attributes Modification*

Allowed Protocols

The second option for the policy set result is Allowed Protocols. The purpose of Allowed Protocols is to limit the number of authentication protocols that can be used within a configured policy set.

Authentication protocols are crucial for verifying the identity of users and devices attempting to access a network or system. They vary significantly in their security mechanisms, complexity, and suitability for different environments.

Remember to enable only the protocols that are absolutely needed and expected for that policy set. ISE comes with one default allowed protocol, **Default Network Access**, which is shown in Figure 4-16. As a best practice, Default Network Access should never be directly used in production. To configure Allowed Protocols, navigate to **Policy > Policy Elements > Results > Authentication > Allowed Protocols**.

Technet24

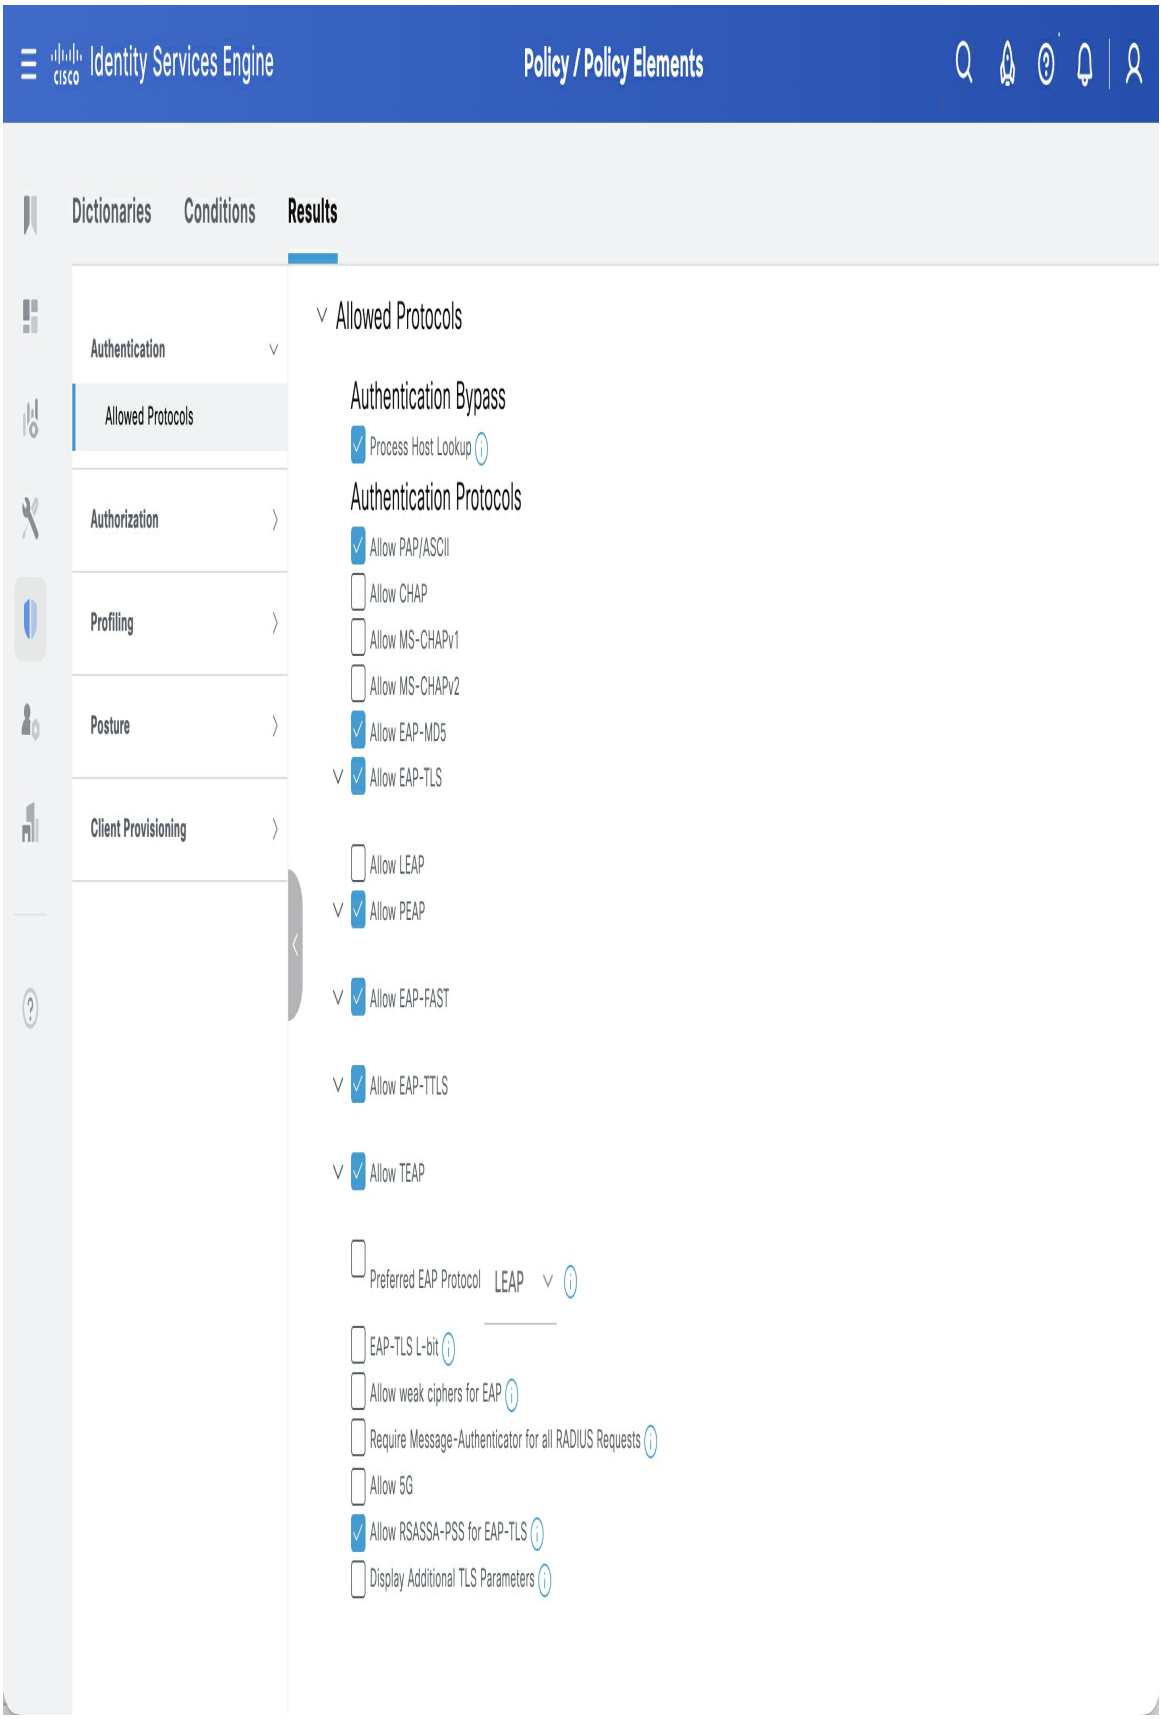

Identity Services Engine
Policy / Policy Elements
Dictionaries
Conditions
Results
Authentication
Allowed Protocols
Authorization
Profiling
Posture
Client Provisioning
Allowed Protocols
Authentication Bypass
Process Host Lookup
Authentication Protocols
Allow PAP/ASCII
Allow CHAP
Allow MS-CHAPv1
Allow MS-CHAPv2
Allow EAP-MD5
Allow EAP-TLS
Allow LEAP
Allow PEAP
Allow EAP-FAST
Allow EAP-TTLS
Allow TEAP
Preferred EAP Protocol
LEAP
EAP-TLS L-bit
Allow weak ciphers for EAP
Require Message-Authenticator for all RADIUS Requests
Allow 5G
Allow RSASSA-PSS for EAP-TLS
Display Additional TLS Parameters

Figure 4-16 *Cisco ISE Allowed Protocols*

The **Authentication Bypass** section has a single configuration option: **Process Host Lookup**. It controls MAC Authentication Bypass (MAB). If enabled, MAB authentications will be allowed within the policy set.

The **Authentication Protocols** section has control of which protocols are enabled for that particular Allowed Protocols set.

- **PAP/ASCII (Password Authentication Protocol):** PAP sends usernames and passwords in cleartext (unencrypted) over the network. It uses a two-way handshake for authentication.
- **CHAP (Challenge Handshake Authentication Protocol):** CHAP employs a three-way handshake and a challenge-response mechanism. The password itself is never sent over the network; instead, a hash (often MD5) of the password combined with a random challenge is transmitted.
- **MS-CHAPv1 (Microsoft Challenge Handshake Authentication Protocol version 1):** MS-CHAPv1 is an older version of Microsoft's implementation of CHAP. It uses a challenge-response mechanism. However, it is less secure than MS-CHAPv2 due to a relatively weak method for hashing the user's password, making it vulnerable to offline dictionary attacks.
- **MS-CHAPv2 (Microsoft Challenge Handshake Authentication Protocol version 2):** MS-CHAPv2 is an improved and more secure version of MS-CHAPv1. It employs a stronger challenge-response mechanism and supports mutual authentication, where both the client and the server authenticate each other. Passwords are not sent directly.
- **EAP-MD5 (Extensible Authentication Protocol–MD5):** EAP-MD5 is one of the earliest EAP methods defined. It uses an MD5

Technet24

hash function for authentication. It provides only one-way authentication (client to server) and does not support mutual authentication or key generation. It offers minimal security and is vulnerable to dictionary attacks and man-in-the-middle attacks due to the lack of server authentication and key generation.

- **EAP-TLS (Extensible Authentication Protocol–Transport Layer Security):** EAP-TLS is an open standard EAP method that utilizes the Transport Layer Security (TLS) protocol.

  It is certificate-based, requiring both the client and the authentication server to have digital certificates for mutual authentication. It is considered one of the most secure EAP standards available. It eliminates passwords, providing strong security against man-in-the-middle attacks and credential theft.

- **LEAP (Lightweight Extensible Authentication Protocol):** LEAP is a proprietary wireless authentication protocol developed by Cisco Systems. It uses dynamic WEP keys and provides mutual authentication between wireless clients and RADIUS servers. It authenticates users using usernames and passwords.

- **PEAP (Protected Extensible Authentication Protocol):** PEAP was developed jointly by Cisco Systems, Microsoft, and RSA Security. PEAP encapsulates EAP within a secure Transport Layer Security tunnel. It typically requires only a server-side certificate to establish this secure tunnel. A common inner method is EAP-MSCHAPv2.

- **EAP-FAST (Flexible Authentication via Secure Tunneling):** EAP-FAST is a protocol proposed by Cisco Systems as a replacement for LEAP. It was designed to address LEAP's weaknesses while maintaining a lightweight implementation. It uses a protected access credential (PAC) to establish a TLS tunnel. It was mostly superseded by TEAP.

- **EAP-TTLS (Extensible Authentication Protocol–Tunneled Transport Layer Security):** EAP-TTLS is similar in concept to PEAP; it establishes an encrypted TLS tunnel. It creates a secure, encrypted tunnel between the client and the authentication server, requiring only a server-side certificate. Within this tunnel, various less-secure authentication methods (such as PAP, CHAP, MS-CHAPv2) can be used.
- **TEAP (Tunnel Extensible Authentication Protocol):** TEAP is a modern tunnel-based EAP method defined in RFC 7170. It establishes a mutually authenticated TLS tunnel. It uses Type-Length-Value (TLV) objects to convey authentication data within the tunnel. A key feature is its support for EAP chaining, allowing multiple authentication methods (for example, machine and user authentication) within a single EAP conversation.

Best Practice:

Cisco recommends that you use EAP-TLS and TEAP whenever possible. These authentication protocols provide strong cryptographic security and resistance to phishing attacks

Within different authentication methods, you can also configure advanced settings. For example, within TEAP, you can configure which inner methods are enabled: EAP-MSCHAPv2 or EAP-TLS. You can also select **Enable EAP Chaining** under the protocol settings, as shown in Figure 4-17.

Technet24

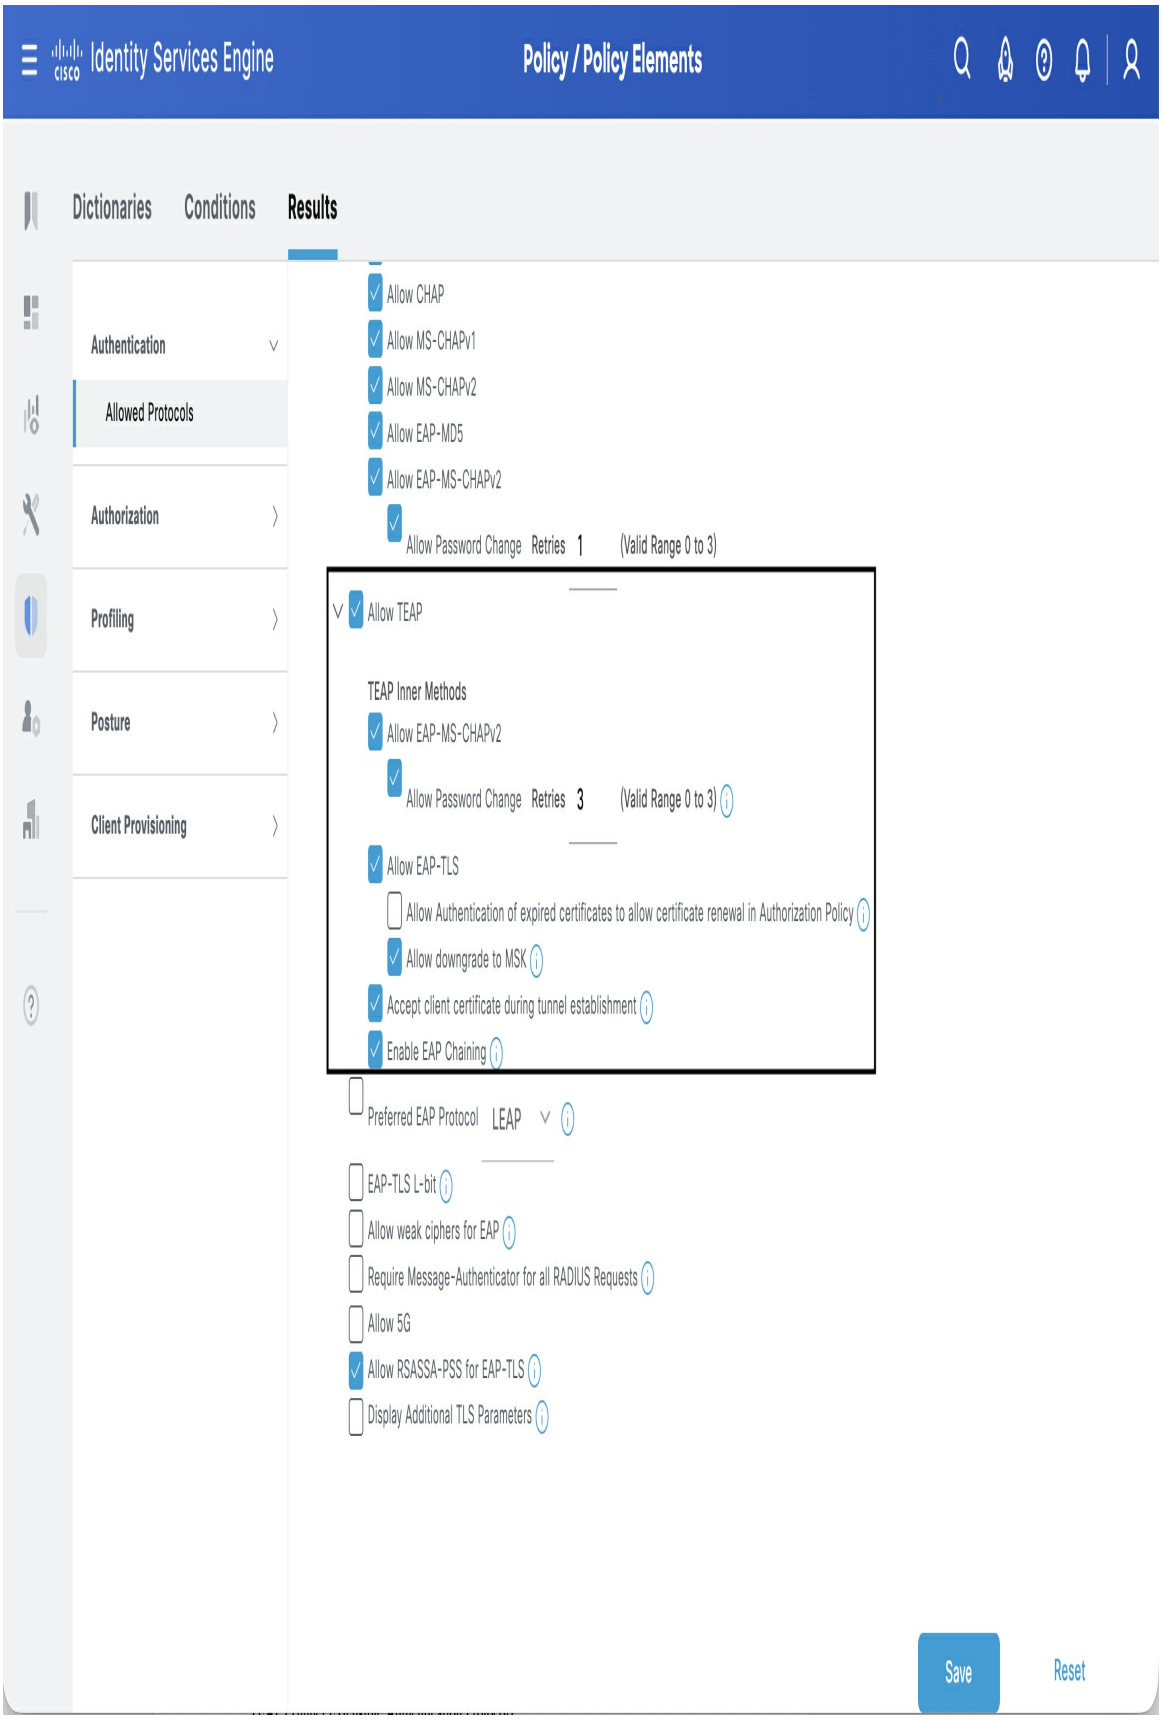

Identity Services Engine
Policy / Policy Elements
Dictionaries
Conditions
Results
Authentication
Allowed Protocols
Authorization
Profiling
Posture
Client Provisioning
Allow CHAP
Allow MS-CHAPv1
Allow MS-CHAPv2
Allow EAP-MD5
Allow EAP-MS-CHAPv2
Allow Password Change Retries 1 (Valid Range 0 to 3)
Allow TEAP
TEAP Inner Methods
Allow EAP-MS-CHAPv2
Allow Password Change Retries 3 (Valid Range 0 to 3)
Allow EAP-TLS
Allow Authentication of expired certificates to allow certificate renewal in Authorization Policy
Allow downgrade to MSK
Accept client certificate during tunnel establishment
Enable EAP Chaining
Preferred EAP Protocol LEAP
EAP-TLS L-bit
Allow weak ciphers for EAP
Require Message-Authenticator for all RADIUS Requests
Allow 5G
Allow RSASSA-PSS for EAP-TLS
Display Additional TLS Parameters
Save
Reset

Figure 4-17 *Allowed Protocols TEAP Settings*

EAP chaining is a powerful mechanism that allows both user and machine authentication to happen in a single session, rather than requiring two separate sessions. This is achieved by "chaining" multiple EAP authentications together and cryptographically binding them to ensure security and integrity.

Additional protocol settings are available right after the **Protocols** section, as shown in Figure 4-18.

Technet24

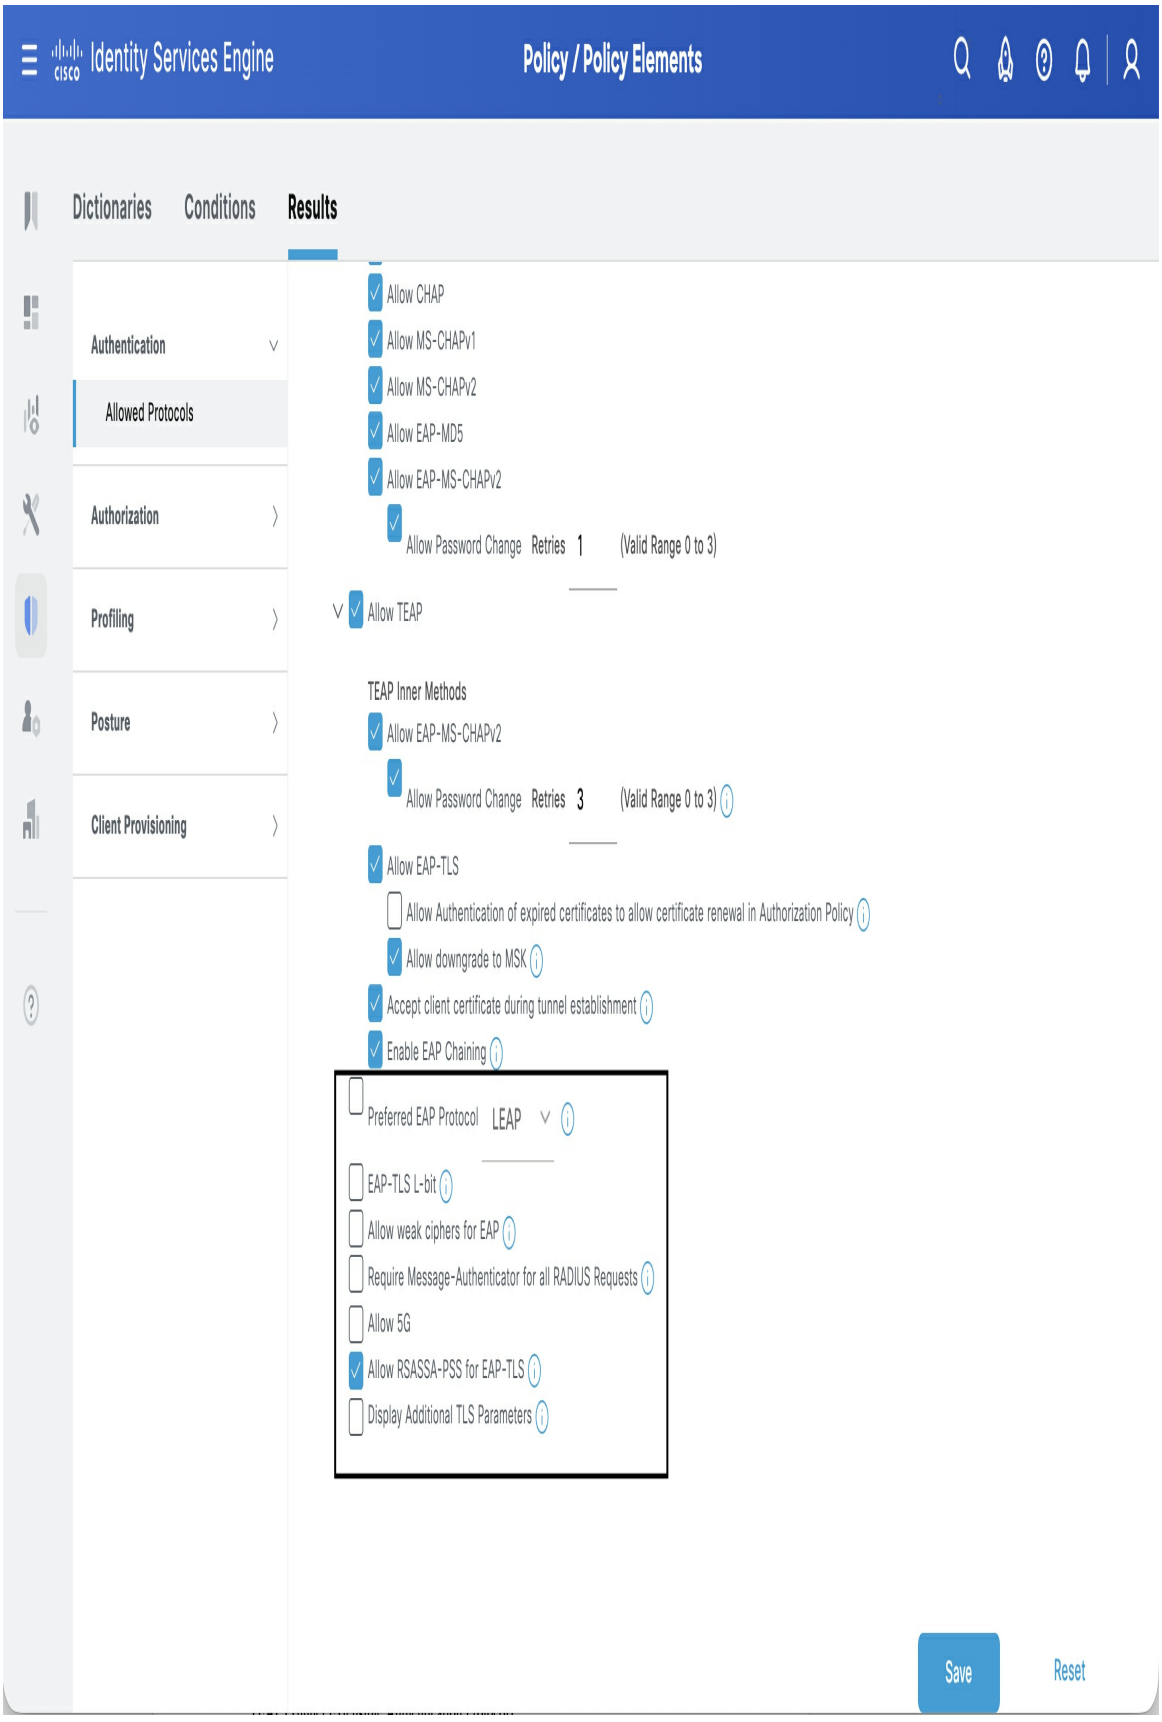

Identity Services Engine
Policy / Policy Elements
Dictionaries
Conditions
Results
Authentication
Allowed Protocols
Authorization
Profiling
Posture
Client Provisioning
Allow CHAP
Allow MS-CHAPv1
Allow MS-CHAPv2
Allow EAP-MD5
Allow EAP-MS-CHAPv2
Allow Password Change Retries 1 (Valid Range 0 to 3)
Allow TEAP
TEAP Inner Methods
Allow EAP-MS-CHAPv2
Allow Password Change Retries 3 (Valid Range 0 to 3)
Allow EAP-TLS
Allow Authentication of expired certificates to allow certificate renewal in Authorization Policy
Allow downgrade to MSK
Accept client certificate during tunnel establishment
Enable EAP Chaining
Preferred EAP Protocol LEAP
EAP-TLS L-bit
Allow weak ciphers for EAP
Require Message-Authenticator for all RADIUS Requests
Allow 5G
Allow RSASSA-PSS for EAP-TLS
Display Additional TLS Parameters
Save
Reset

Figure 4-18 *Allowed Protocols Additional Settings*

The following are additional protocol settings:

- **Allow Weak Cyphers on ISE:** This setting allows older ciphers to be used during EAP exchange. This capability might be useful for older devices such as old models of IP phones. If enabled, this setting does significantly compromise security by allowing weaker ciphers.

  ### Best Practice:

  Whenever possible, select more secure ciphers over enabling outdated security mechanisms.

- **Require Message-Authenticator for All RADIUS Requests:** This setting controls what ISE will do if it doesn't receive a Message-Authenticator in the RADIUS packet from a network access device. Message-Authenticator provides integrity protection for a RADIUS packet. This mechanism was introduced as a response to BlastRADIUS. For security purposes, it is a best practice to enable this setting. However, all devices matching the configured Allowed Protocols must send the Message-Authenticator in the Access-Request. You may need to upgrade or reconfigure NADs to support the Message-Authenticator before enabling this setting on ISE.

Policy Set Hit Counts

Policy set hit counts are a powerful mechanism that allows you, as administrator, to monitor the use of configured policy sets, as shown in Figure 4-19. Based on their usage, you can effectively optimize the configured policy sets or remove unused ones.

Technet24


Identity Services Engine
Policy / Policy Sets
Policy Sets
Reset
Reset Policy Set Hit Counts
Save
Status
Policy Set Name
Description
Conditions
Allowed Protocols / Server Sequence
Hits
Actions
View
Search
Default
Default policy set
Default Network Access
0
Refresh frequency Every 15 Minutes
Reset
Refresh

Figure 4-19 *Policy Sets Hit Counts*

Hits counts are automatically refreshed every 15 minutes, but you can also refresh them manually by hovering over the hit count value and clicking **Refresh**. A **Reset** option is available as well. You can reset them per policy set, or if you want to reset policy counts for all policy sets, you can select **Reset Policy Set Hit Counts**, which will reset all policy sets.

Best Practice:

When you configure policies according to best practices, you should arrange them from most hit to less hit based on the hit counts.

Eduroam

Eduroam, which stands for education roaming, is a Wi-Fi service for users at education and research institutions. It allows the same credentials to be used repeatedly across campuses of different institutions.

ISE allows integration with eduroam servers. Overall, there are three use cases:

1. Home user connecting at home institution
2. Home user connecting at remote institution
3. Remote user connecting at home institution

Figure 4-20 shows how home users connect at remote institutions.

Technet24

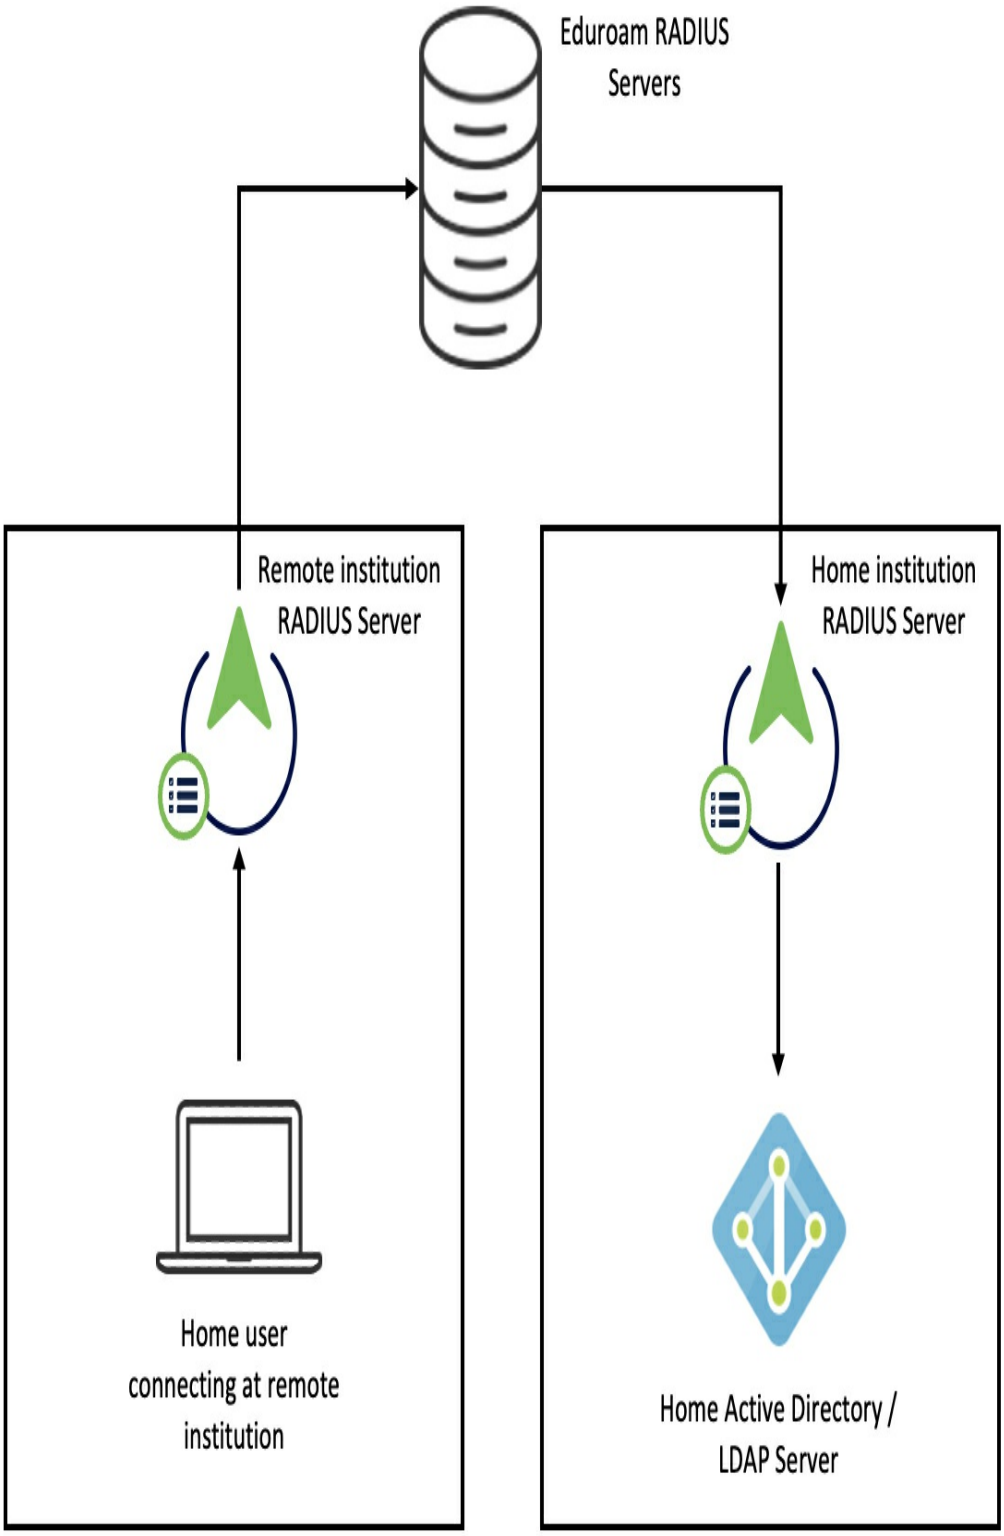

Eduroam RADIUS
Servers
Remote institution
RADIUS Server
Home institution
RADIUS Server
Home user
connecting at remote
institution
Home Active Directory /
LDAP Server

Figure 4-20 *Eduroam Authentication Flow for a Home User Connecting at Remote Institution*

Policy sets allow ISE to group the first and second use cases into one policy set and the third use case in another policy set.

Let's look closer into the third use case first. The third use case describes when a remote user connects at a home institution; hence, the home ISE server should proxy the authentication to eduroam servers so that it can be forwarded to the user's actual institution.

First, you should define external RADIUS servers and then configure the RADIUS server sequence using these servers, as shown in Figure 4-21.

Technet24

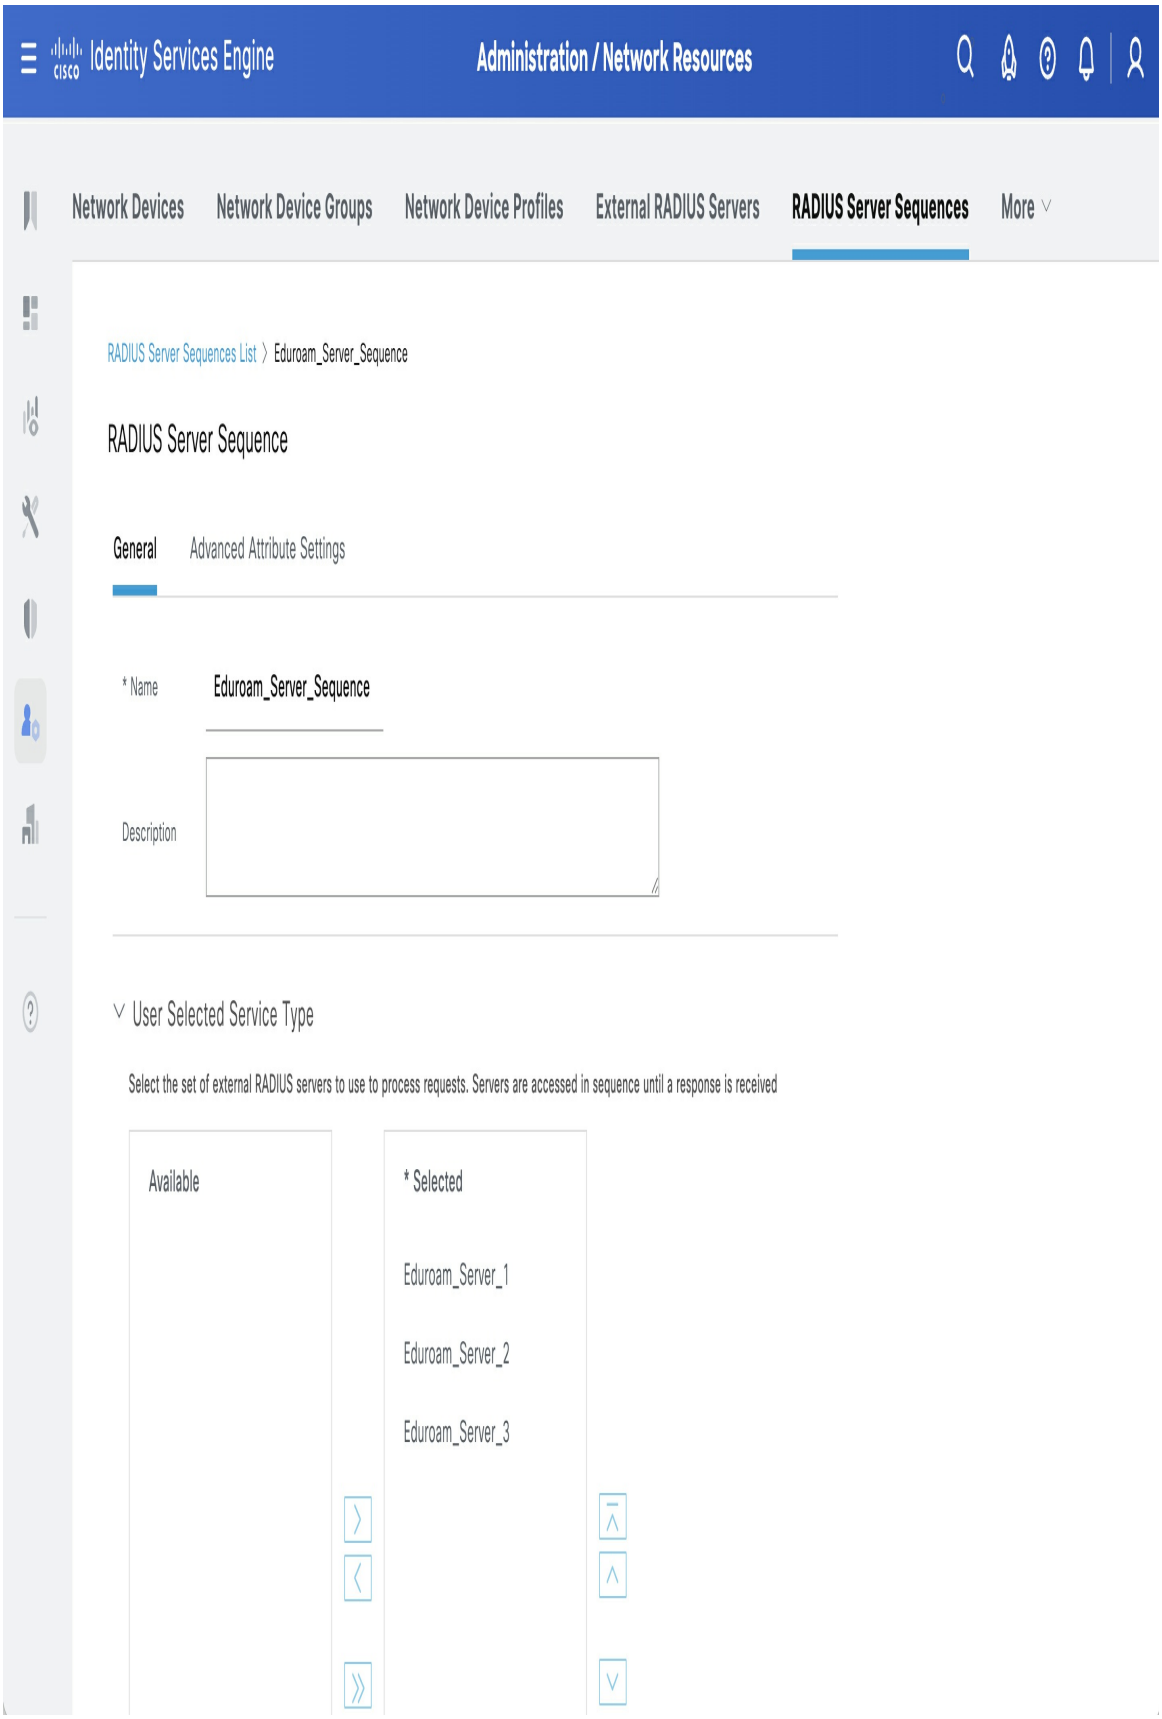

Identity Services Engine
Administration / Network Resources
Network Devices
Network Device Groups
Network Device Profiles
External RADIUS Servers
RADIUS Server Sequences
More
RADIUS Server Sequences List 〉 Eduroam_Server_Sequence
RADIUS Server Sequence
General
Advanced Attribute Settings
* Name
Eduroam_Server_Sequence
Description
User Selected Service Type
Select the set of external RADIUS servers to use to process requests. Servers are accessed in sequence until a response is received
Available
* Selected
Eduroam_Server_1
Eduroam_Server_2
Eduroam_Server_3

Figure 4-21 *RADIUS Server Sequence Configuration for Eduroam Use Case*

Now you must configure the policy set to forward RADIUS requests to those servers. To do so, you need first to create a compound condition. Navigate to **Policy > Policy Elements > Conditions > Library Conditions**. Then configure the condition according to Figure 4-22.

Technet24

Editor

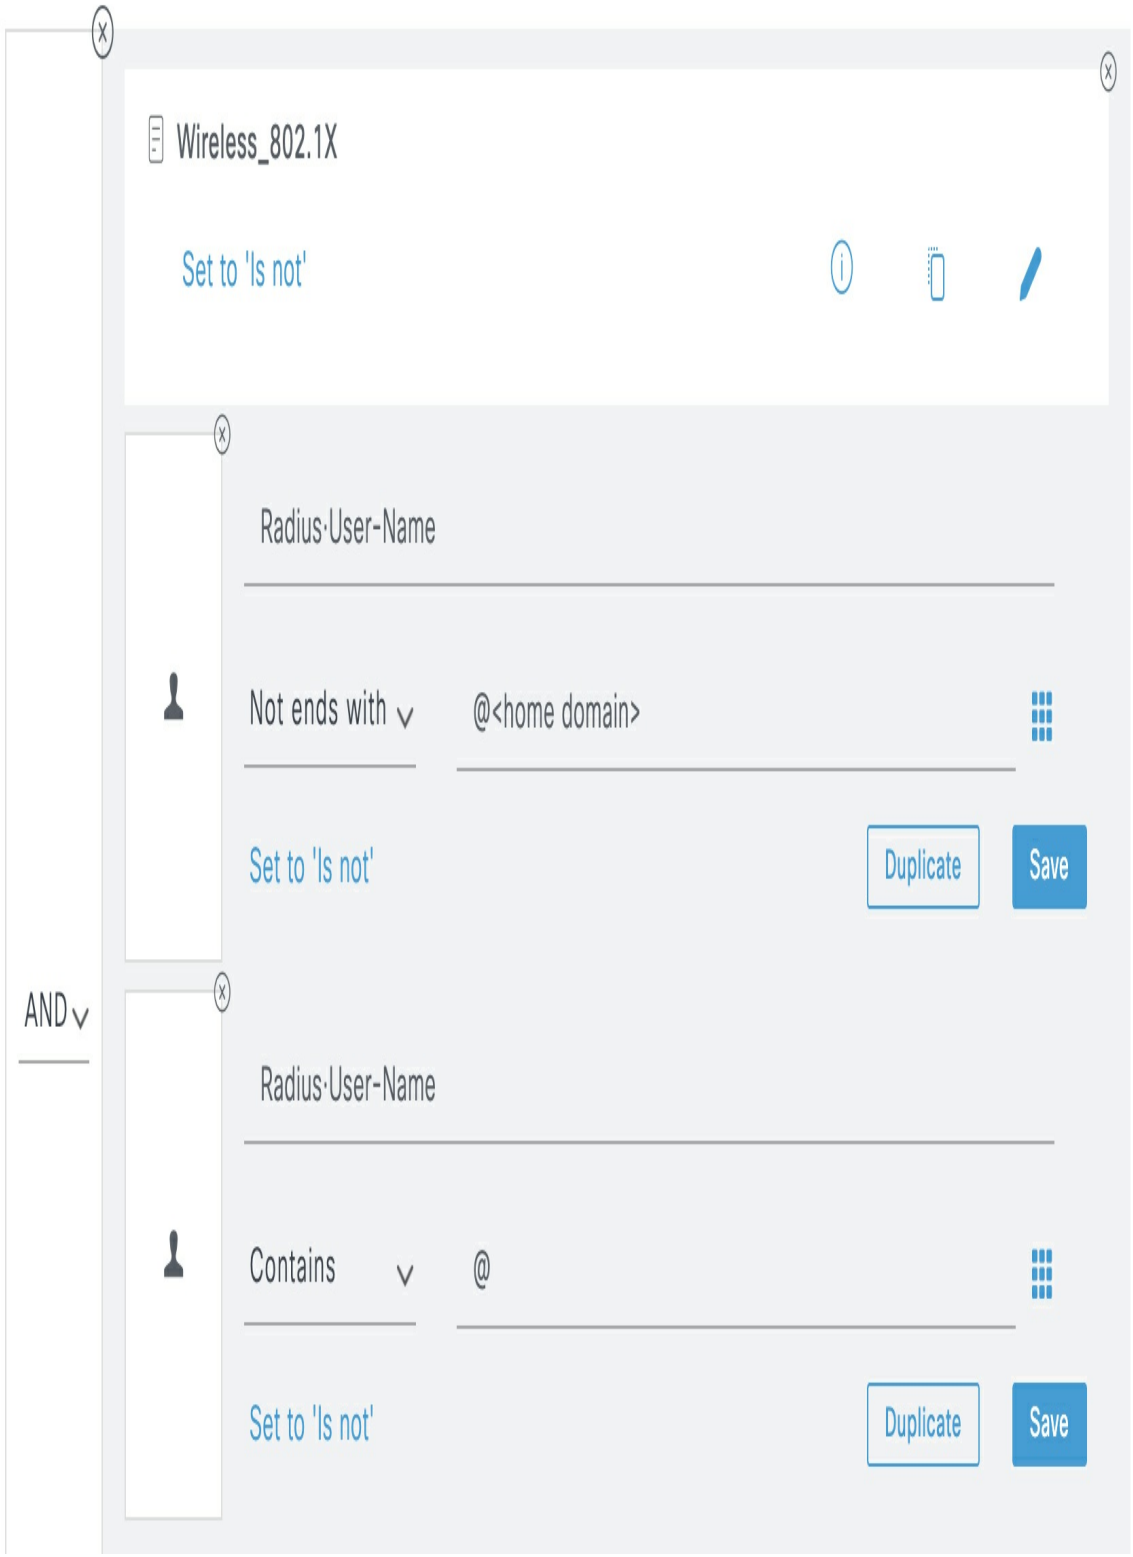

Wireless_802.1X
Set to 'Is not'
Radius·User-Name
Not ends with
@<home domain>
Set to 'Is not'
Duplicate
Save
AND
Radius·User-Name
Contains
@
Set to 'Is not'
Duplicate
Save

Figure 4-22 *Cisco ISE Library Condition Configuration for Eduroam Policy Set for Use Case Where a Home User Connects at a Remote Institution*

The following conditions are used:

- **Wireless_802.1X** ensures you match the 802.1X wireless authentications.
- **RADIUS:User-Name Not ends with @** ensures that this rule will not be matched for home users. Replace with your actual domain.
- **RADIUS:User-Name Contains @** ensures that the email address username will match this condition. It is expected that eduroam users will use their email address as their username.

Click **Save** to save this as new library condition and call it **Eduroam_User_External**.

Next, configure the policy set according to Figure 4-23. Note that WLAN id 6 is used as an example. You can configure it according to your WLC configuration or use an SSID name.

Technet24

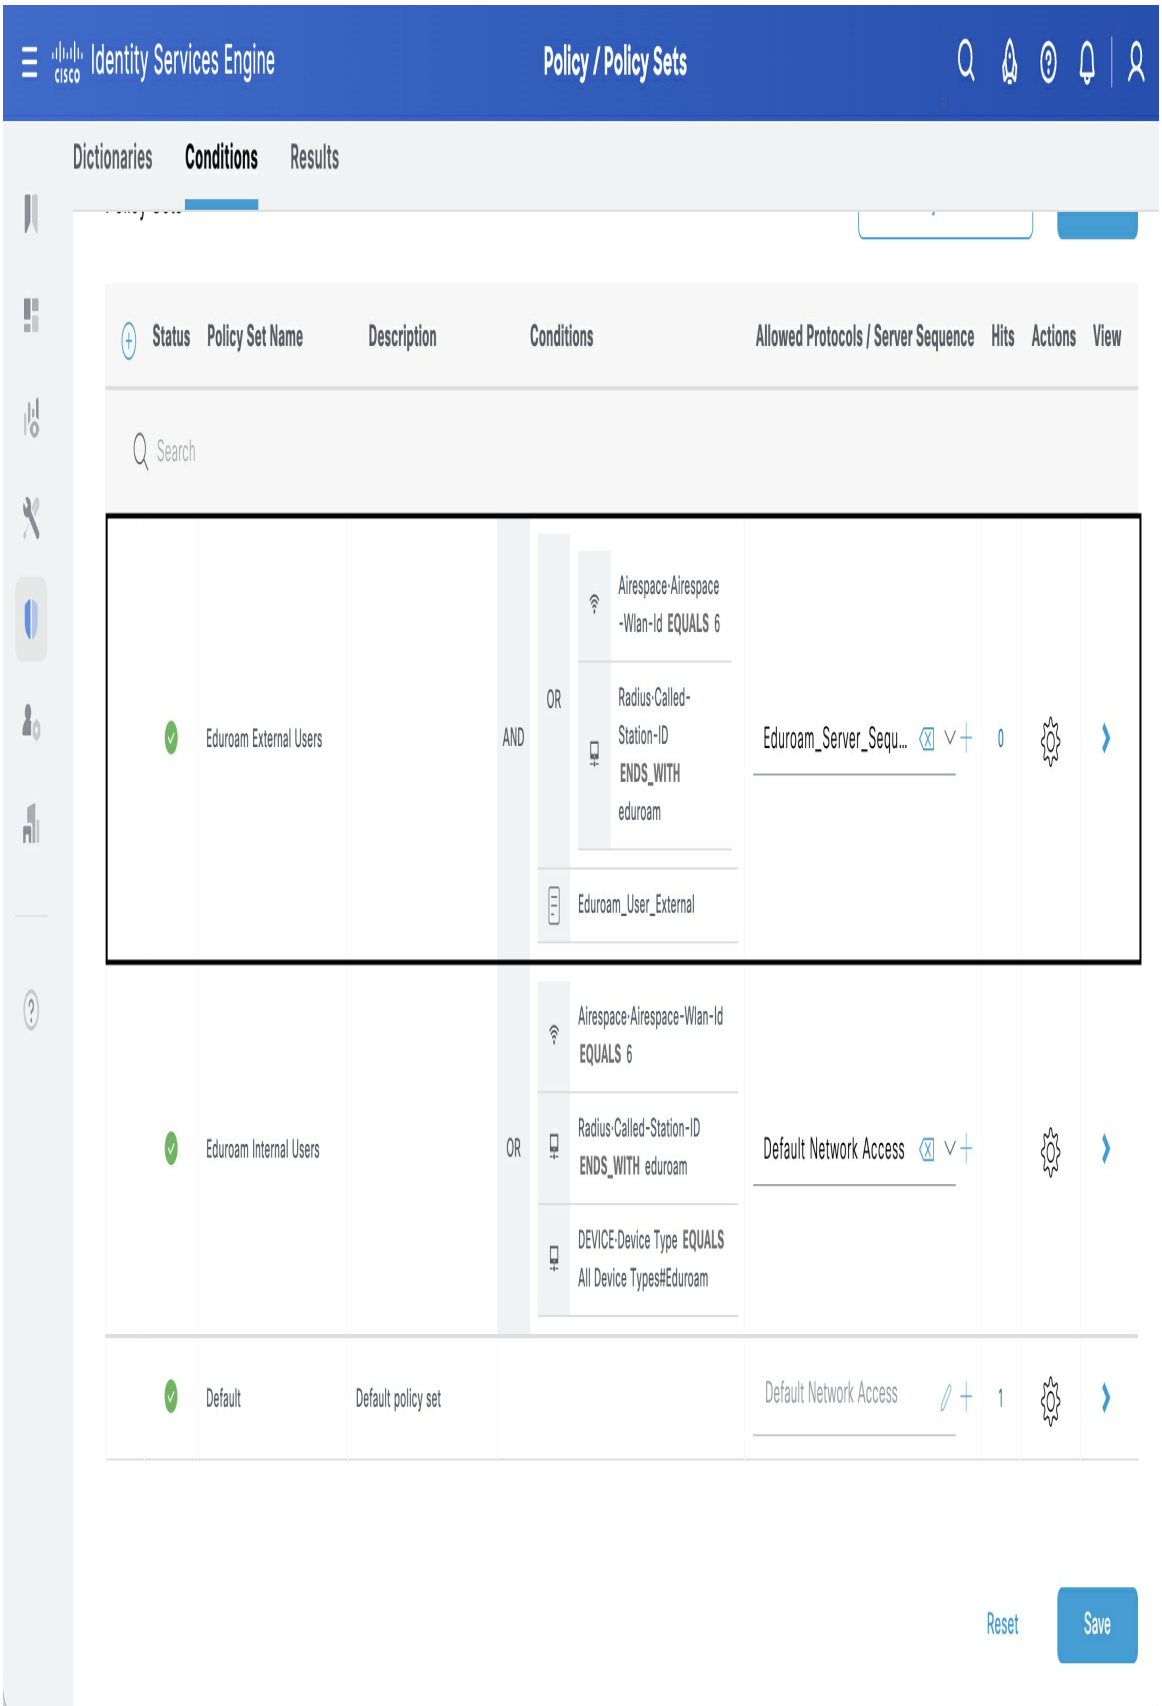

Identity Services Engine
Policy / Policy Sets
Dictionaries
Conditions
Results
Status
Policy Set Name
Description
Conditions
Allowed Protocols / Server Sequence
Hits
Actions
View
Search
Eduroam External Users
AND
OR
Airespace·Airespace -Wlan-Id EQUALS 6
Radius·Called-Station-ID ENDS_WITH eduroam
Eduroam_User_External
Eduroam_Server_Sequ...
0
Eduroam Internal Users
OR
Airespace·Airespace-Wlan-Id EQUALS 6
Radius·Called-Station-ID ENDS_WITH eduroam
DEVICE·Device Type EQUALS All Device Types#Eduroam
Default Network Access
Default
Default policy set
Default Network Access
1
Reset
Save

Figure 4-23 *Policy Set Configuration for External Users*

Now let's look into the first and second use cases. Because those examples contain home users, they should be authenticated locally with the institutions' own Active Directory/LDAP servers. Authentications should come from

- **First use case:** Local wireless LAN controllers, for home users connecting at their home institution
- **Second use case:** Eduroam servers, for home users connecting remotely

For the second use case, you need to configure eduroam servers as network access devices. To do this, place them into a dedicated network device group (Eduroam), as shown in Figure 4-24. Configure the network device group Eduroam prior to this step.

Technet24

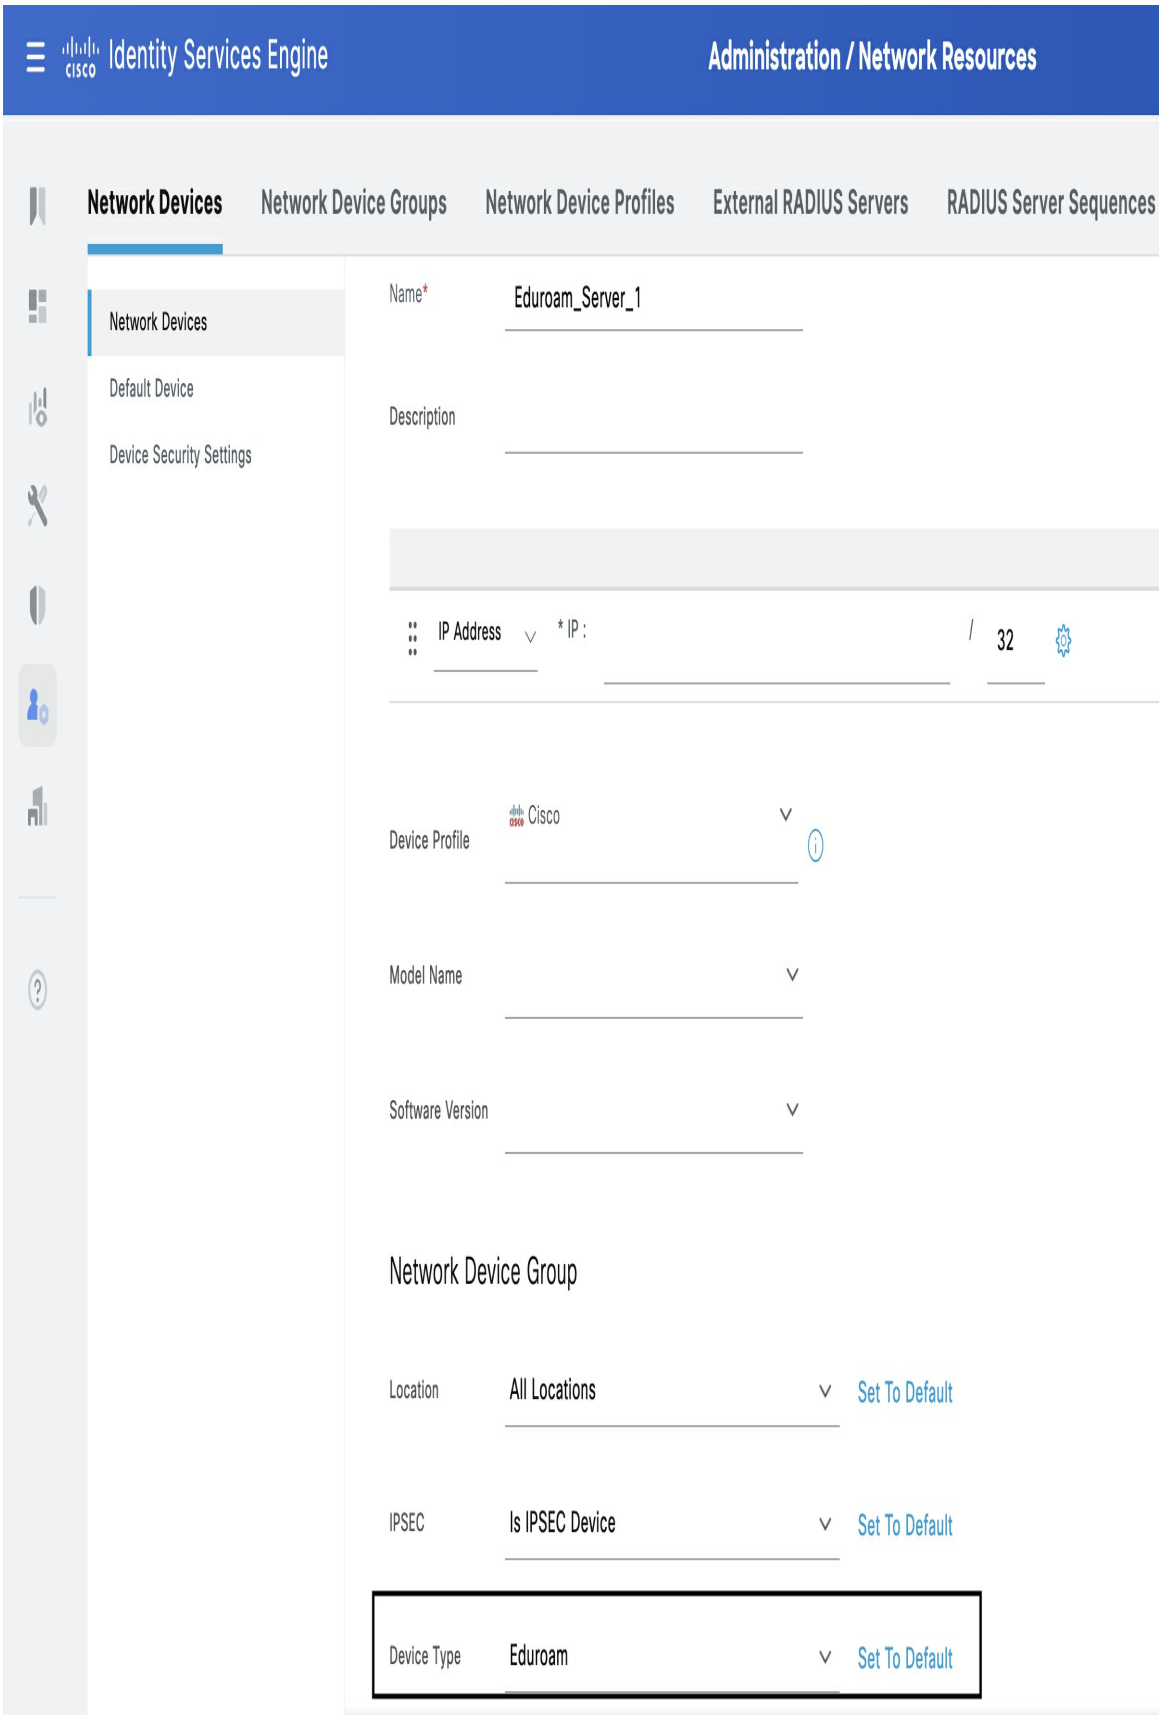

Identity Services Engine
Administration / Network Resources
Network Devices
Network Device Groups
Network Device Profiles
External RADIUS Servers
RADIUS Server Sequences
Network Devices
Default Device
Device Security Settings
Name*
Eduroam_Server_1
Description
IP Address
* IP :
/ 32
Device Profile
Cisco
Model Name
Software Version
Network Device Group
Location
All Locations
Set To Default
IPSEC
Is IPSEC Device
Set To Default
Device Type
Eduroam
Set To Default

Figure 4-24 *Network Device Configuration for Eduroam Servers*

Both use cases can be summed up under a single policy set condition, as shown in Figure 4-25.

Technet24

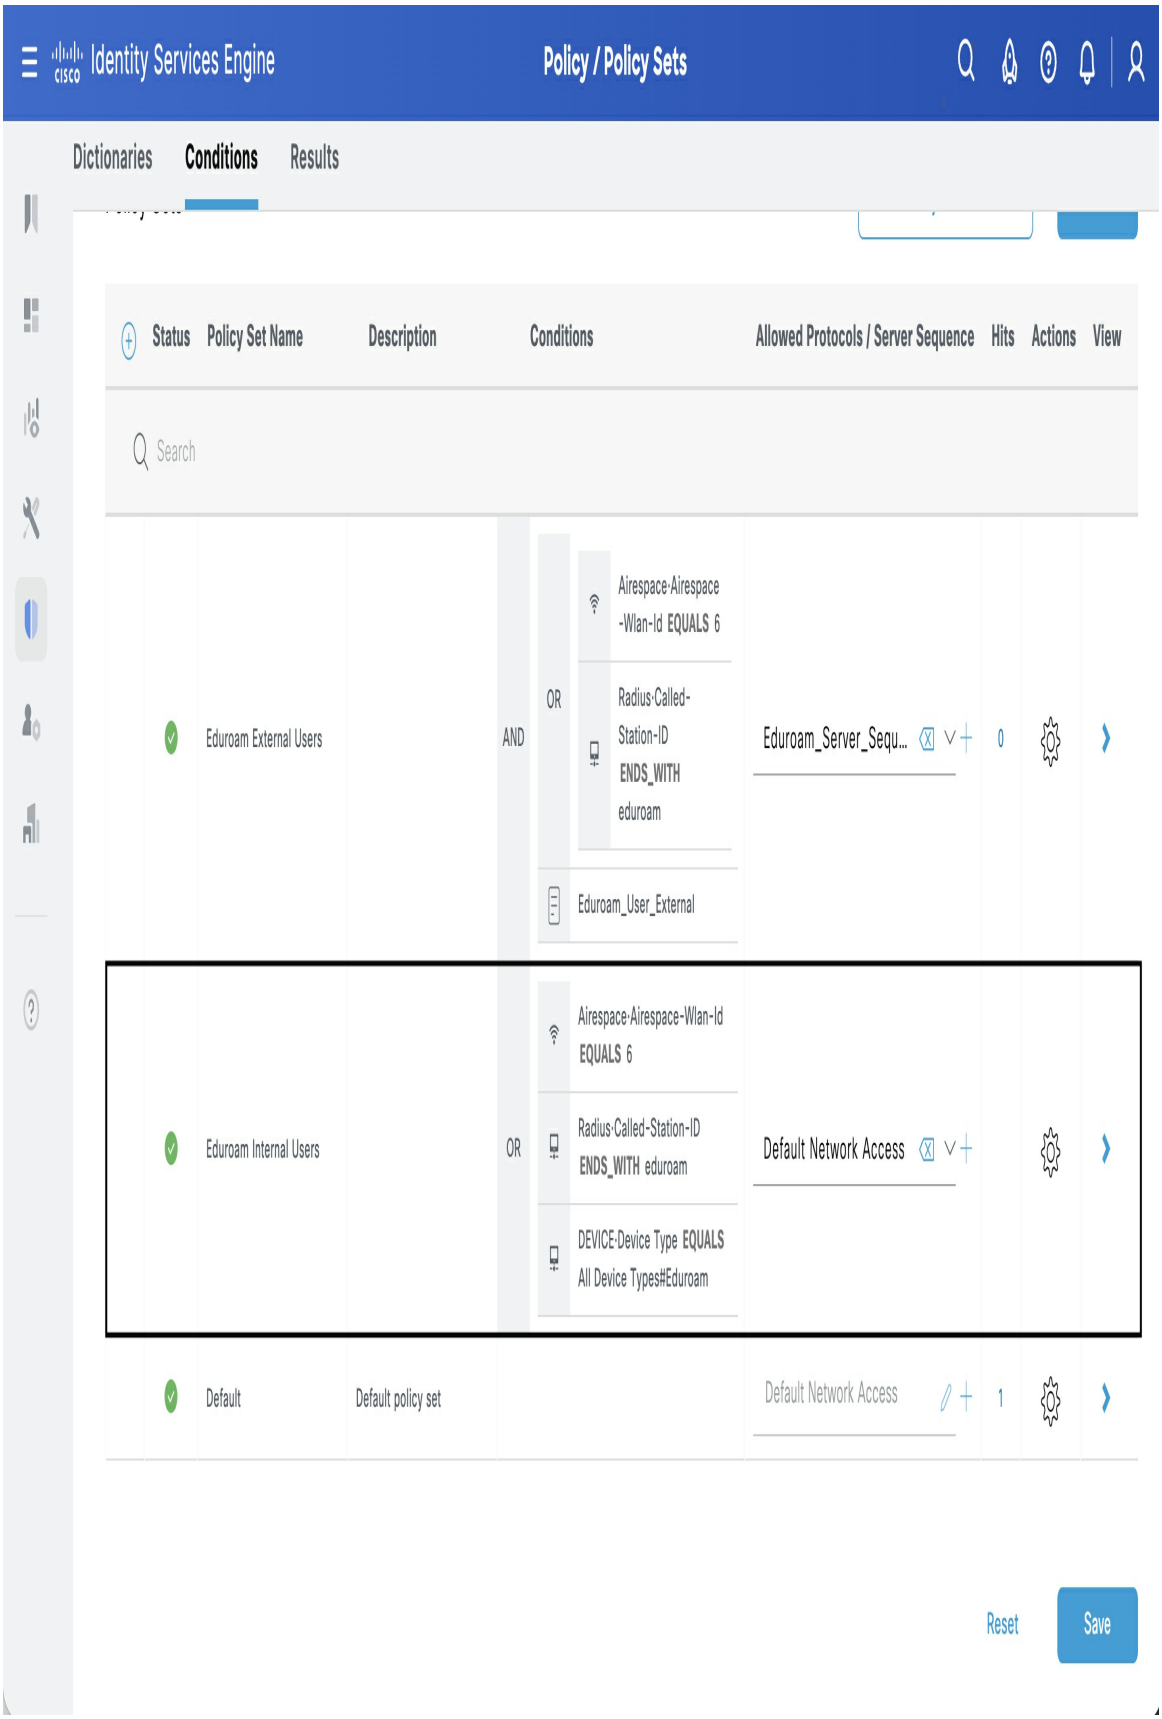

Identity Services Engine
Policy / Policy Sets
Dictionaries
Conditions
Results
Status
Policy Set Name
Description
Conditions
Allowed Protocols / Server Sequence
Hits
Actions
View
Search
Eduroam External Users
AND
OR
Airespace·Airespace-Wlan-Id EQUALS 6
Radius·Called-Station-ID ENDS_WITH eduroam
Eduroam_User_External
Eduroam_Server_Sequ...
0
Eduroam Internal Users
OR
Airespace·Airespace-Wlan-Id EQUALS 6
Radius·Called-Station-ID ENDS_WITH eduroam
DEVICE·Device Type EQUALS All Device Types#Eduroam
Default Network Access
Default
Default policy set
Default Network Access
1
Reset
Save

Figure 4-25 *Policy Set Configuration for Internal Users*

For accountability purposes within eduroam, an internal user's policy set can distinguish between the first and second use cases by configuring two different authentication policy conditions, as shown in Figure 4-26.

Technet24

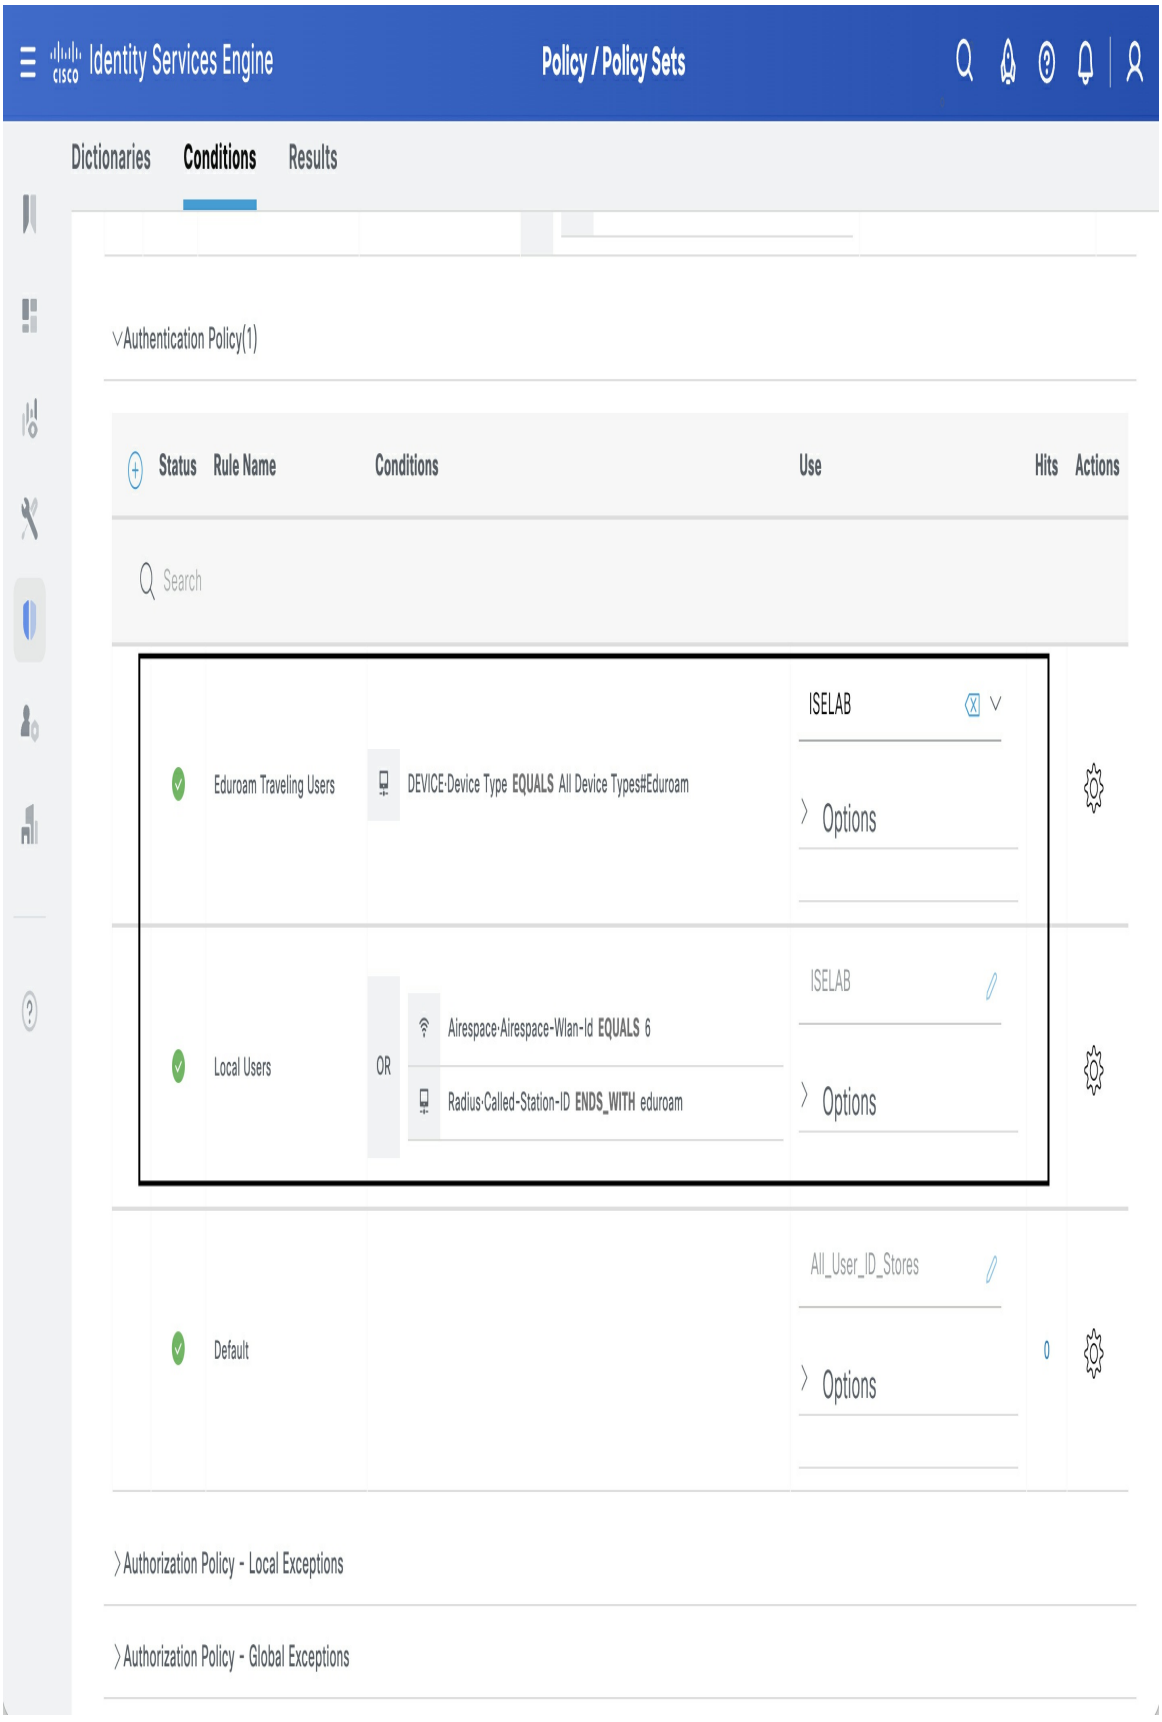

Identity Services Engine
Policy / Policy Sets
Dictionaries
Conditions
Results
Authentication Policy(1)
Status
Rule Name
Conditions
Use
Hits
Actions
Search
Eduroam Traveling Users
DEVICE·Device Type EQUALS All Device Types#Eduroam
ISELAB
Options
Local Users
OR
Airespace·Airespace-Wlan-Id EQUALS 6
Radius·Called-Station-ID ENDS_WITH eduroam
ISELAB
Options
Default
All_User_ID_Stores
0
Options
Authorization Policy - Local Exceptions
Authorization Policy - Global Exceptions

Figure 4-26 *Authentication Policy Conditions to Distinguish Traveling and Local Home Users*

Authorization policies for all three use cases should be configured based on company/institution security requirements.

Summary

Policy sets in ISE are logical concepts that combine authentication and authorization policies, enabling administrators to efficiently manage diverse access requirements in a scalable and organized manner. ISE policies are processed sequentially; ISE evaluates policy sets from top to bottom and applies the first one whose conditions are met.

As a result of policy set selection, either a specific RADIUS server sequence (for proxying authentication to external RADIUS servers) or a set of Allowed Protocols (to restrict which authentication methods are permitted) is used.

Administration can benefit from the Conditions Studio, which is a powerful tool within ISE that enables administrators to build complex and reusable matching criteria using logical AND/OR operations on attributes drawn from ISE dictionaries.

For optimal performance and easier administration, it is recommended to configure distinctions by location, authentication method, or type. After policies are configured, you can monitor their use with the help of hit counts.

ISE can seamlessly support eduroam deployments, distinguishing between local and visiting users and forwarding authentications as required by the scenario use case.

Technet24

Chapter 5. Authentication Options

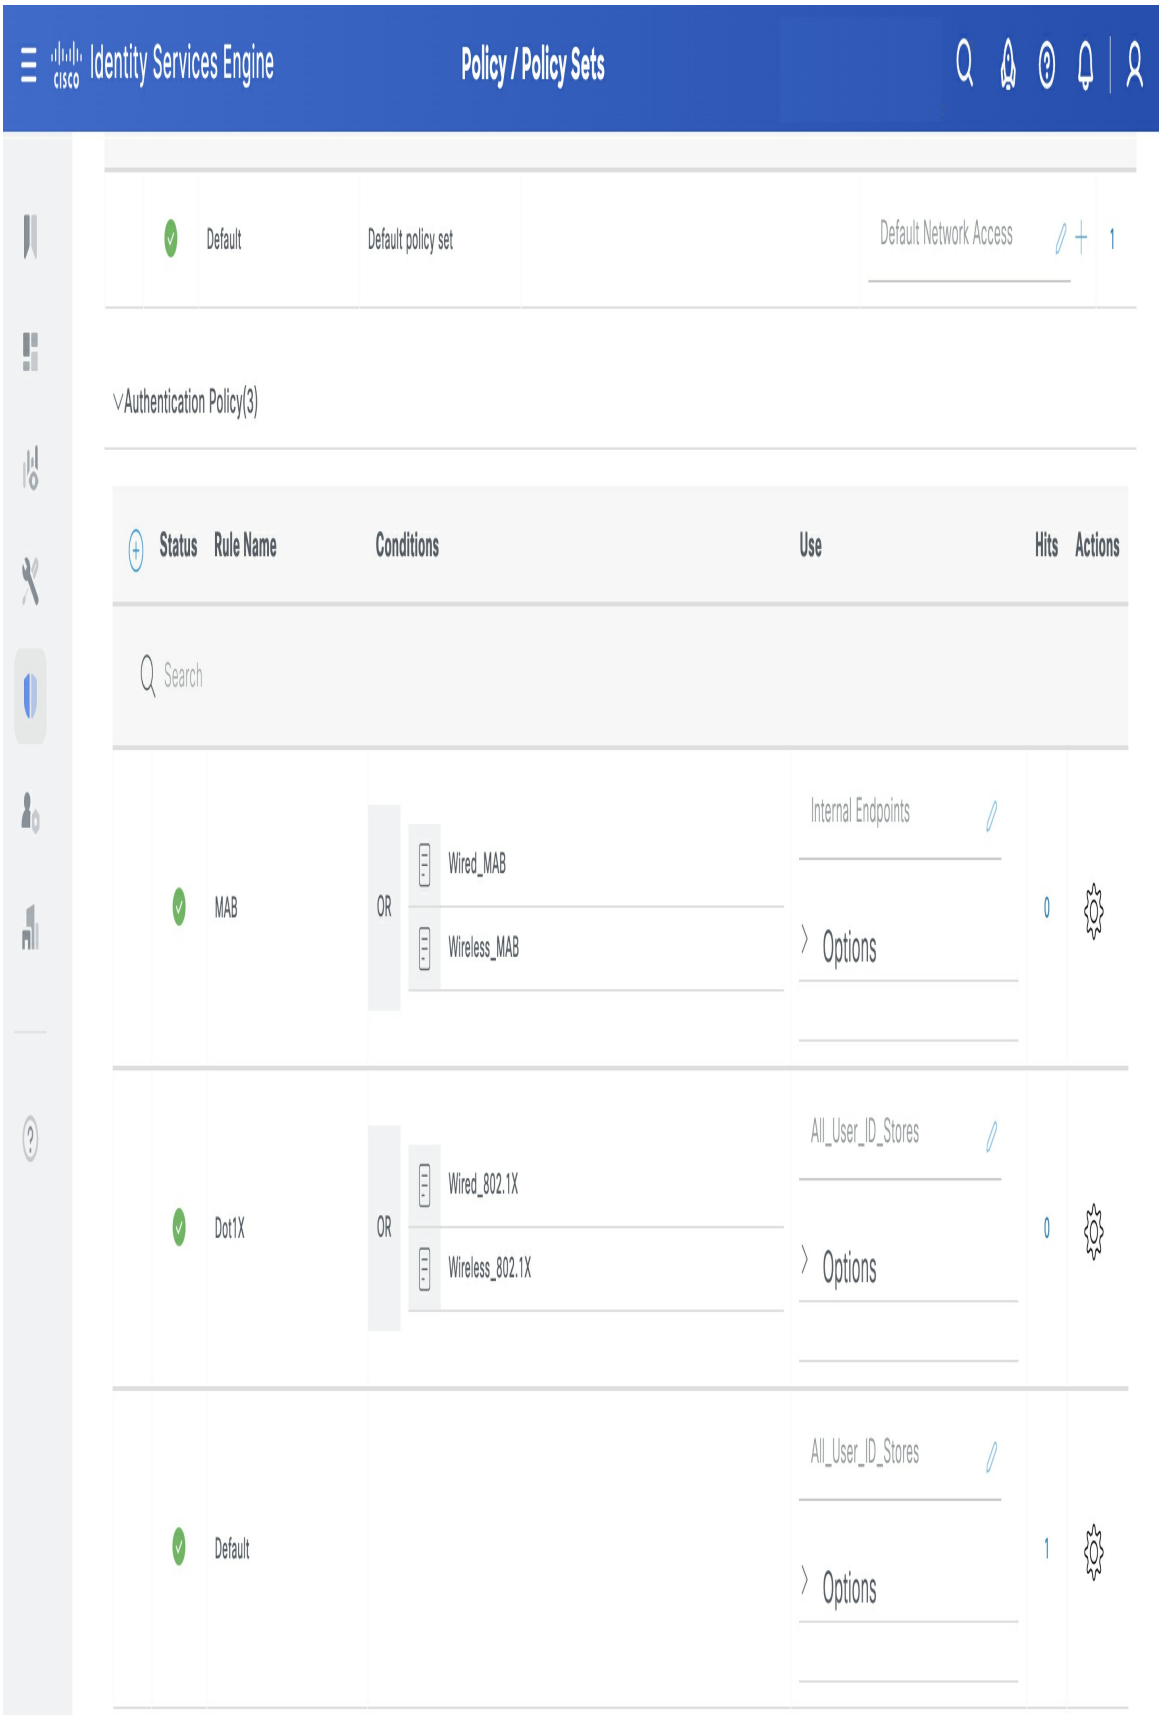
Introduction

Cisco Identity Services Engine is sometimes referred to as a AAA server. AAA stands for authentication, authorization, and accounting.

- Authentication helps confirm that you are who you claim to be. For example, at the airport, you present your passport or an ID document. Staff verify your identity by comparing your appearance with the photo in the document.
- Authorization determines which privileges you have. Continuing with the airport example, staff check the name on your passport to confirm which tickets you have purchased and grant access only to the privileges associated with your ticket. For instance, if you purchased an economy ticket, you are not allowed to sit in business class.
- Accounting records the authentication and authorization processes. In the airport example, this includes scanning your boarding pass and keeping a record that you boarded the plane.

Authentication Policies Concepts

Cisco ISE includes a built-in default policy set that contains its own built-in set of authentication policies, as illustrated in Figure 5-1.


Identity Services Engine
Policy / Policy Sets
Default
Default policy set
Default Network Access
1
Authentication Policy(3)
Status
Rule Name
Conditions
Use
Hits
Actions
Search
MAB
OR
Wired_MAB
Wireless_MAB
Internal Endpoints
Options
0
Dot1X
OR
Wired_802.1X
Wireless_802.1X
All_User_ID_Stores
Options
0
Default
All_User_ID_Stores
Options
1

Technet24

Figure 5-1 *Built-in Authentication Policies in Default Policy Set*

Built-in authentication policies help you understand why different policies are necessary. The main goal of differentiated authentication policies is to allow users and endpoints to authenticate using the correct identity source. You should avoid verifying users against the internal MAC address database of ISE. You should also prevent MAC Authentication Bypass (MAB) authentications from occurring against certificate authentication profiles, because all such attempts will fail.

The structure of an authentication policy is similar to that of policy sets. You have the rule name, condition, and result. The difference lies in the result: As a result of authentication policy selection, you select an identity source or an identity source sequence.

Additional options are associated with every authentication, as shown in Figure 5-2.

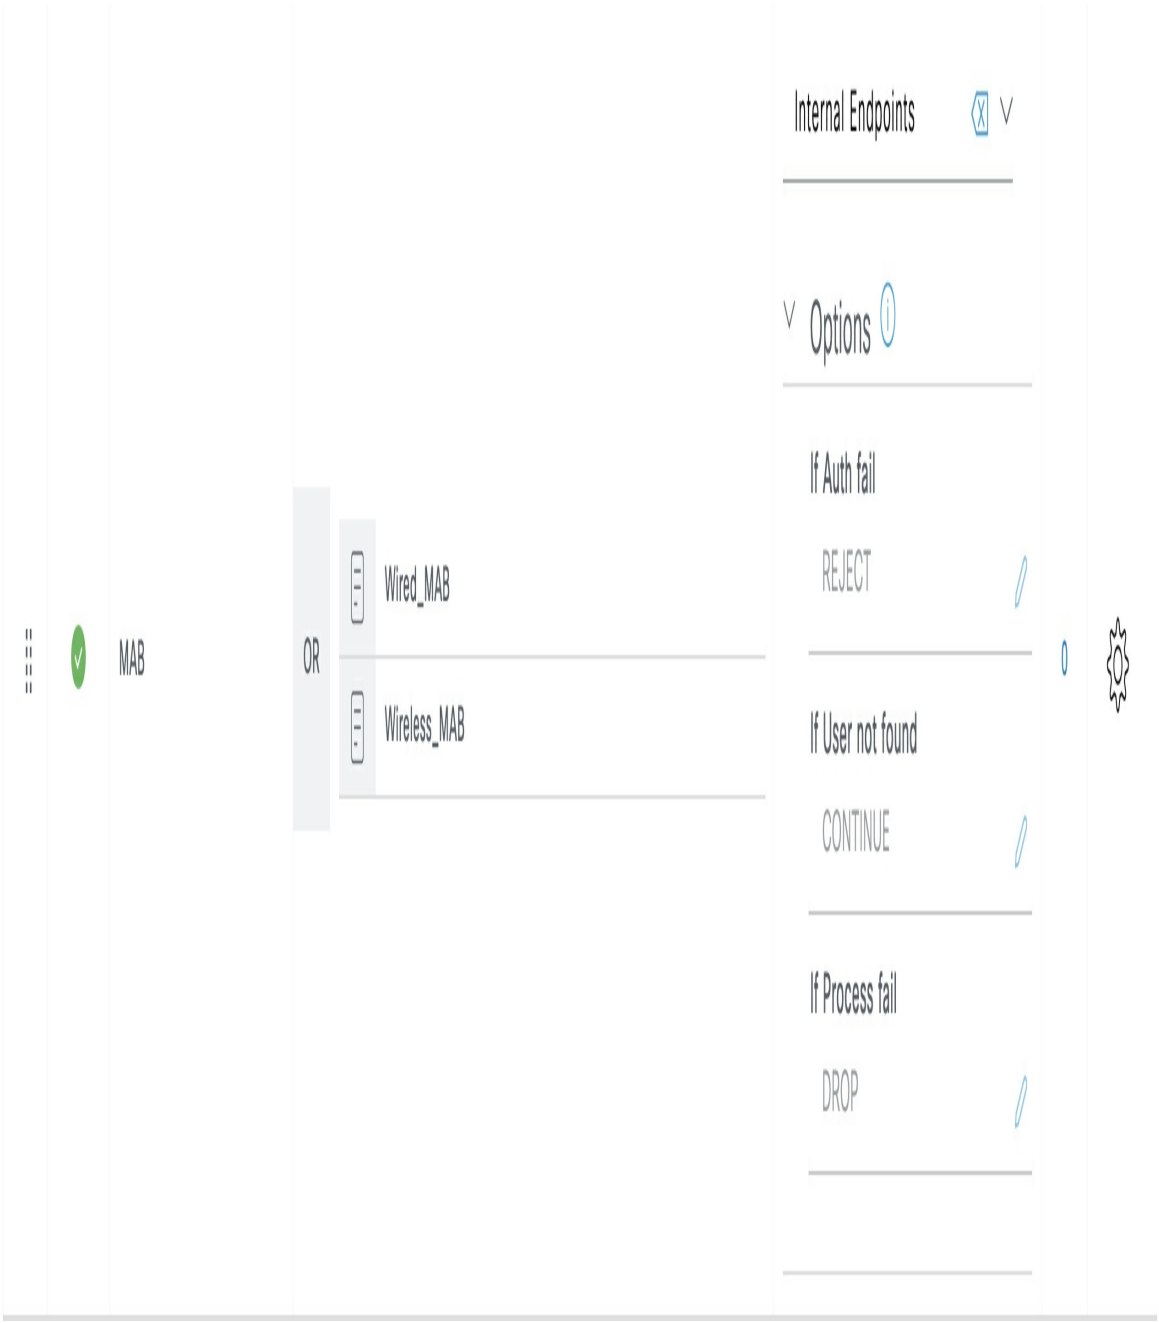


Figure 5-2 *Advanced Authentication Options*

These additional options provide flexibility, even if authentication fails. You can select to perform the following actions:

- **CONTINUE:** ISE continues authorization, even if authentication fails. You commonly use this option for guest flow scenarios, where an unknown endpoint attempts authentication. Despite the

Technet24

lack of knowledge about the endpoint, ISE continues authorization, allowing the guest user to be redirected to the guest portal.

- **DROP:** Authentication silently drops.
- **REJECT:** RADIUS Access-Reject is sent back to the network access device (NAD) if authentication fails.

There are also different types of failures. Authentication can end with one of the following:

- **Auth Failure:** For example, when the user enters a wrong password.
- **User not found:** For example, during MAB, the MAC address is not found in the internal database.
- **Process Failure:** For example, when ISE is unable to access the identity database or databases.

Best Practice:

Granular authentication policy configuration improves ISE performance and avoids unnecessary resource utilization. You should configure a unique identity source result per authentication policy whenever possible.

Identity Sources

There are multiple ways to verify user or endpoint identity. You can check whether the username and password combination is correct. You can also verify whether the presented user certificate is valid. Alternatively, you can check whether the MAC address of the endpoint is present in the Internal Endpoints Identity Store.

To ensure you verify identity in the correct place, Cisco ISE uses the

concept of an identity source. The authentication process checks credentials against the configured identity source. If you want to combine multiple sources, you can use an identity source sequence, which defines the order in which Cisco ISE searches these sources during authentication, as shown in Figure 5-3.

Technet24

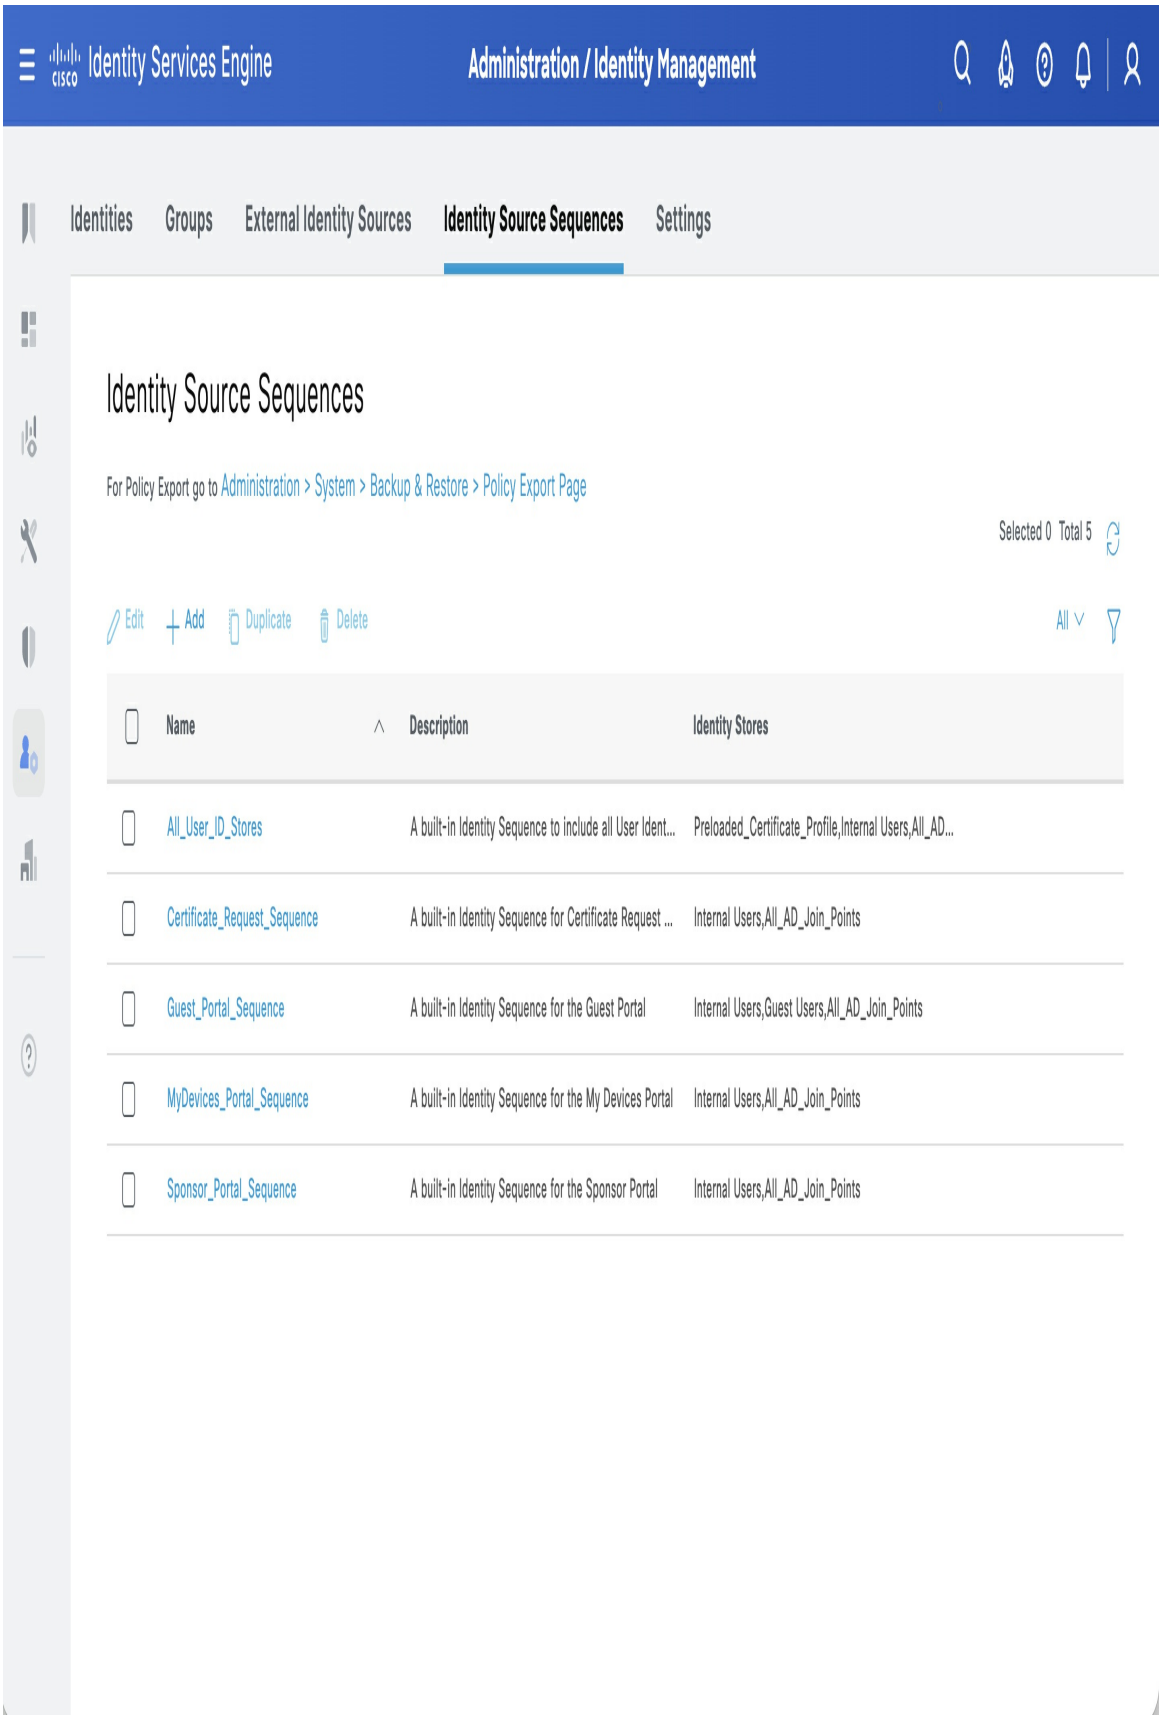

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
Identity Source Sequences
For Policy Export go to Administration > System > Backup & Restore > Policy Export Page
Selected 0 Total 5
Edit
Add
Duplicate
Delete
All
Name
Description
Identity Stores
All_User_ID_Stores
A built-in Identity Sequence to include all User Ident...
Preloaded_Certificate_Profile,Internal Users,All_AD...
Certificate_Request_Sequence
A built-in Identity Sequence for Certificate Request ...
Internal Users,All_AD_Join_Points
Guest_Portal_Sequence
A built-in Identity Sequence for the Guest Portal
Internal Users,Guest Users,All_AD_Join_Points
MyDevices_Portal_Sequence
A built-in Identity Sequence for the My Devices Portal
Internal Users,All_AD_Join_Points
Sponsor_Portal_Sequence
A built-in Identity Sequence for the Sponsor Portal
Internal Users,All_AD_Join_Points

Figure 5-3 *Built-in Identity Source Sequences*

You can either use existing sequences or create your own by combining different identity sources. For example, you can create a custom sequence in which you select only specific Active Directory join points.

ISE either stores user and endpoint data locally or relies on external sources. As a result, there are only two types of identity sources: internal identity sources and external identity sources.

Internal Identity Sources

Internal identity sources include internal users and internal endpoints.

To access the Internal Users Identity Store, navigate to **Administration > Identity Management > Identities > Users**, as shown in Figure 5-4.

Technet24

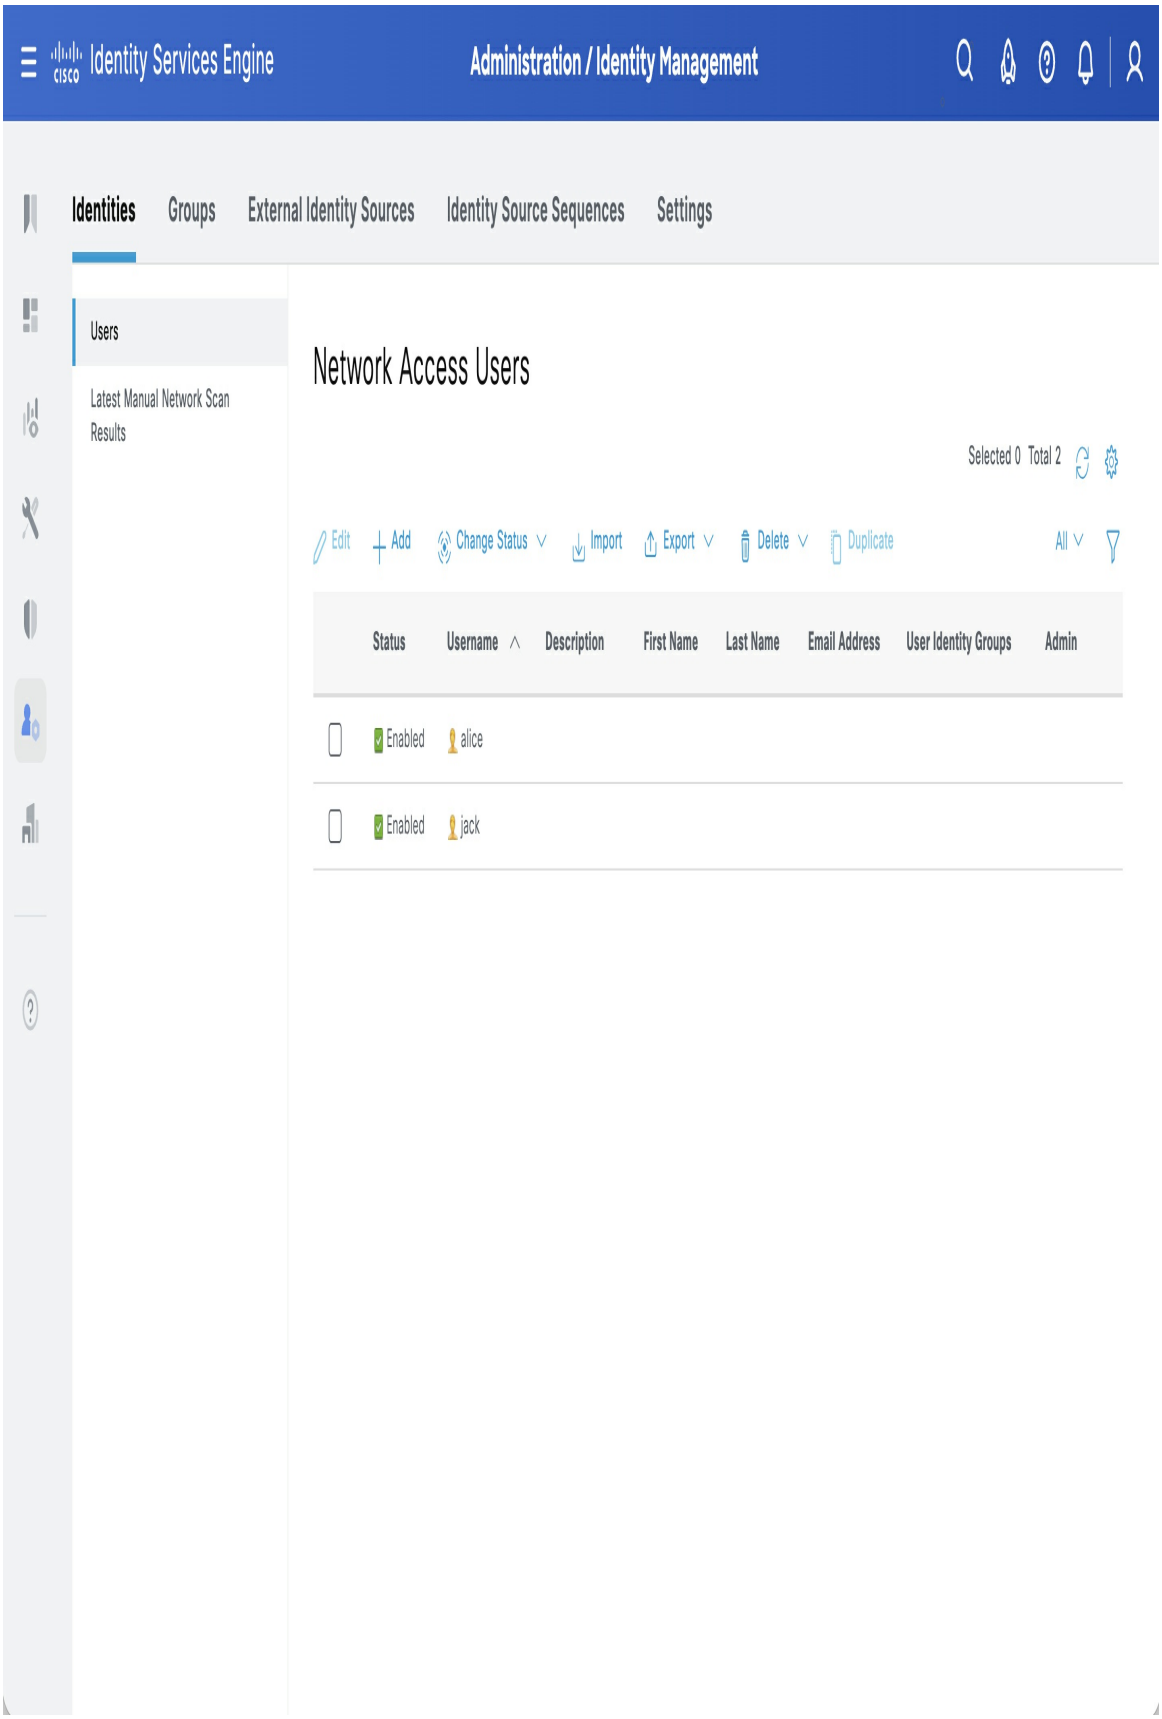

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
Users
Latest Manual Network Scan Results
Network Access Users
Selected 0 Total 2
Edit
Add
Change Status
Import
Export
Delete
Duplicate
All
Status
Username
Description
First Name
Last Name
Email Address
User Identity Groups
Admin
Enabled
alice
Enabled
jack

Figure 5-4 *Internal Users Identity Store*

You can populate the Internal Users Identity Store using different methods:

- Import users by using a CSV file.
- Create users by using the ERS API.
- Create users manually.

You can place users in local identity groups so that you can differentiate access based on group membership. You need to create the required groups first by navigating to **Administration > Identity Management > Identities > Groups > User Identity Groups**, as shown in Figure 5-5.

Technet24

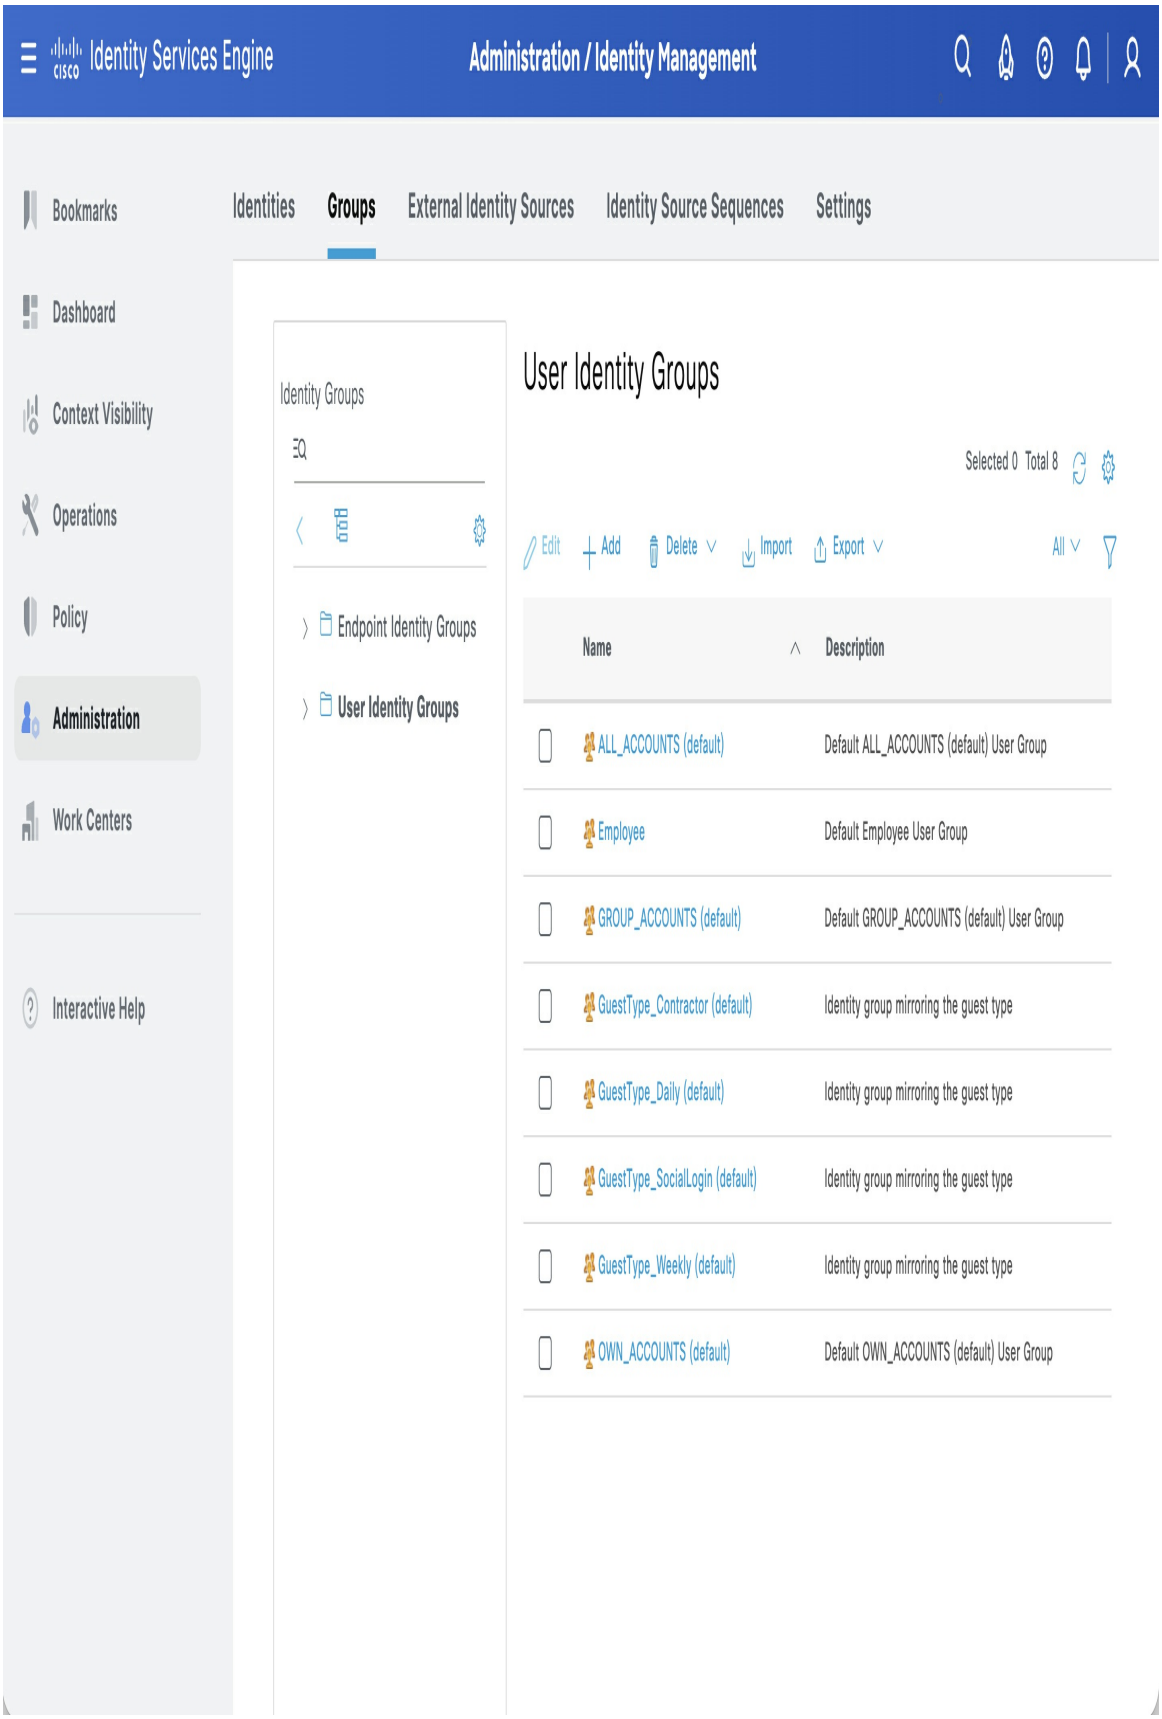

Identity Services Engine
Administration / Identity Management
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
Identity Groups
Endpoint Identity Groups
User Identity Groups
User Identity Groups
Selected 0 Total 8
Edit
Add
Delete
Import
Export
All
Name
Description
ALL_ACCOUNTS (default)
Default ALL_ACCOUNTS (default) User Group
Employee
Default Employee User Group
GROUP_ACCOUNTS (default)
Default GROUP_ACCOUNTS (default) User Group
GuestType_Contractor (default)
Identity group mirroring the guest type
GuestType_Daily (default)
Identity group mirroring the guest type
GuestType_SocialLogin (default)
Identity group mirroring the guest type
GuestType_Weekly (default)
Identity group mirroring the guest type
OWN_ACCOUNTS (default)
Default OWN_ACCOUNTS (default) User Group

Figure 5-5 *Internal Users Identity Groups*

The second Internal Identity Store is called Internal Endpoints. You use this store for MAB authentications, and you can also create groups to differentiate access between endpoints within **Administration > Identity Management > Identities > Groups > Endpoint Identity Groups**, as shown in Figure 5-6.

Technet24

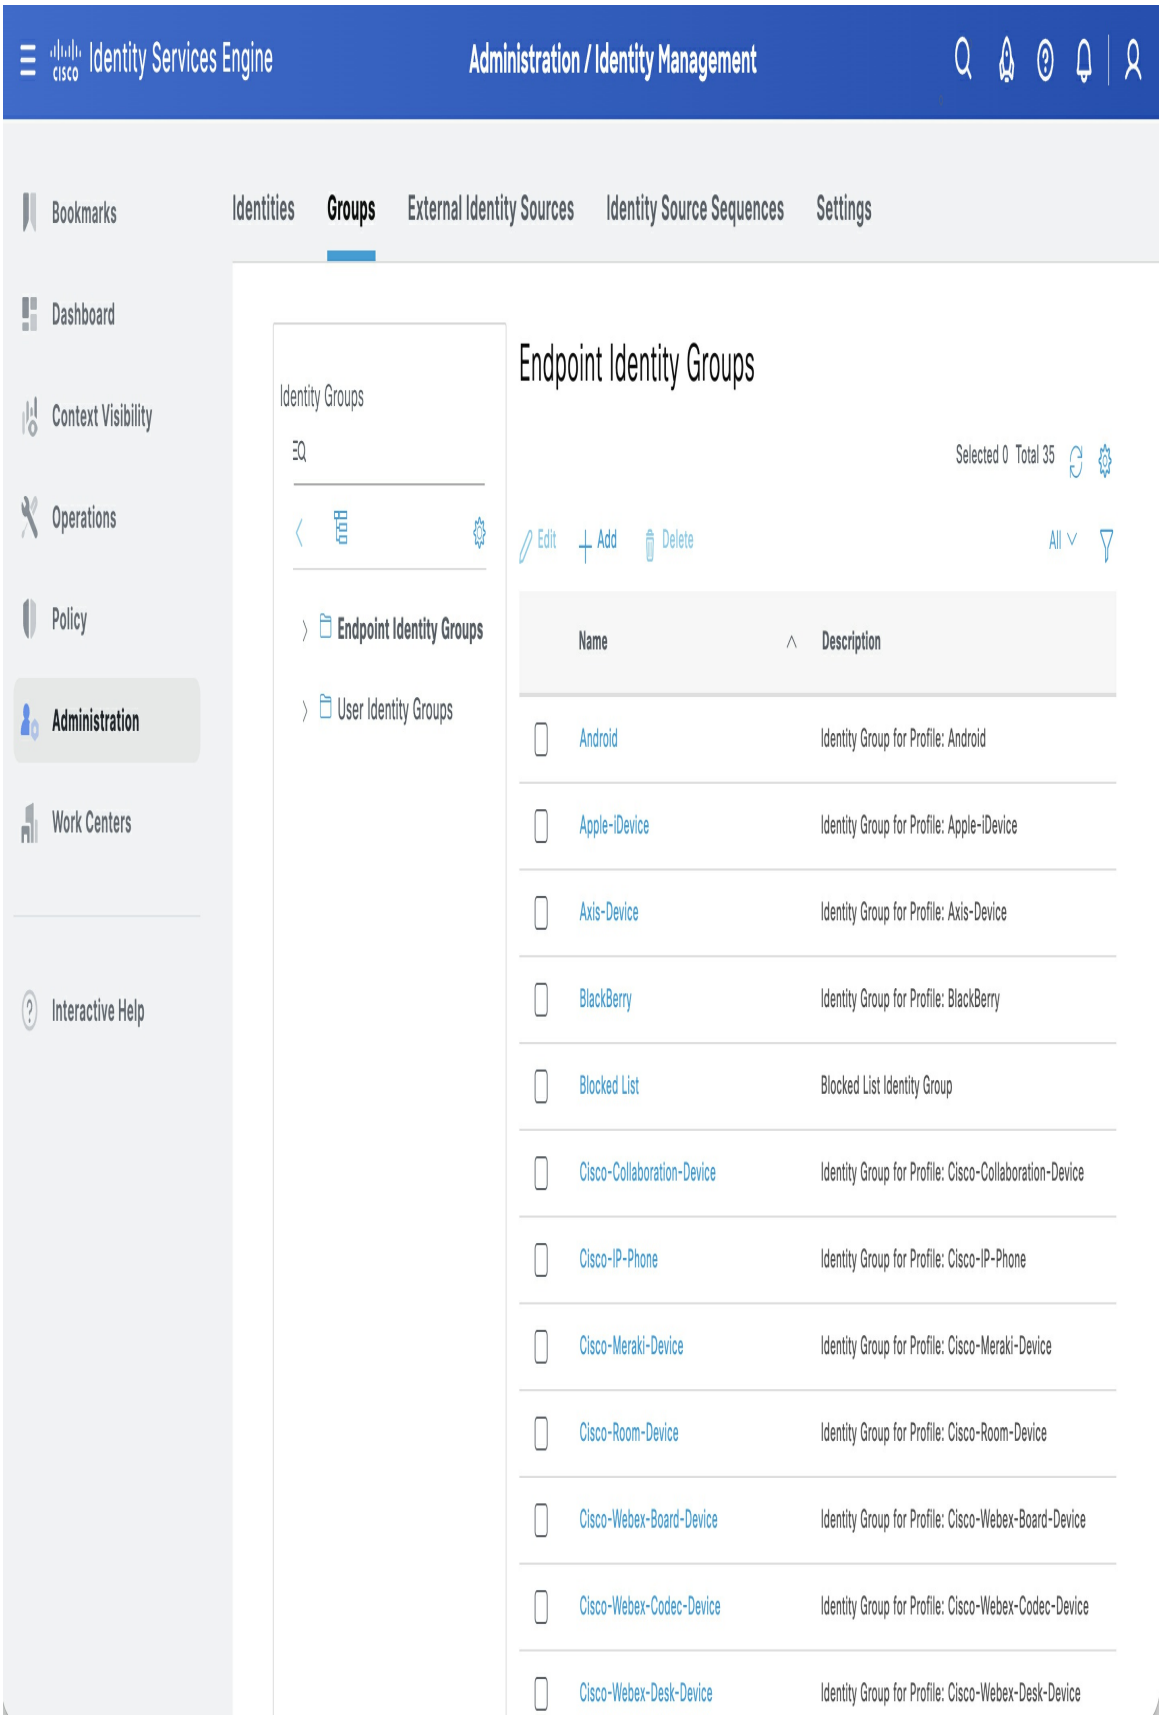

Identity Services Engine
Administration / Identity Management
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
Identity Groups
Endpoint Identity Groups
User Identity Groups
Endpoint Identity Groups
Selected 0 Total 35
Edit
Add
Delete
All
Name
Description
Android
Identity Group for Profile: Android
Apple-iDevice
Identity Group for Profile: Apple-iDevice
Axis-Device
Identity Group for Profile: Axis-Device
BlackBerry
Identity Group for Profile: BlackBerry
Blocked List
Blocked List Identity Group
Cisco-Collaboration-Device
Identity Group for Profile: Cisco-Collaboration-Device
Cisco-IP-Phone
Identity Group for Profile: Cisco-IP-Phone
Cisco-Meraki-Device
Identity Group for Profile: Cisco-Meraki-Device
Cisco-Room-Device
Identity Group for Profile: Cisco-Room-Device
Cisco-Webex-Board-Device
Identity Group for Profile: Cisco-Webex-Board-Device
Cisco-Webex-Codec-Device
Identity Group for Profile: Cisco-Webex-Codec-Device
Cisco-Webex-Desk-Device
Identity Group for Profile: Cisco-Webex-Desk-Device

Figure 5-6 *Internal Endpoint Identity Groups*

Context Visibility

The ISE Internal Endpoint Identity Store is powered by Context Visibility. To access it, navigate to **Context Visibility > Endpoints**, as shown in Figure 5-7.

Technet24

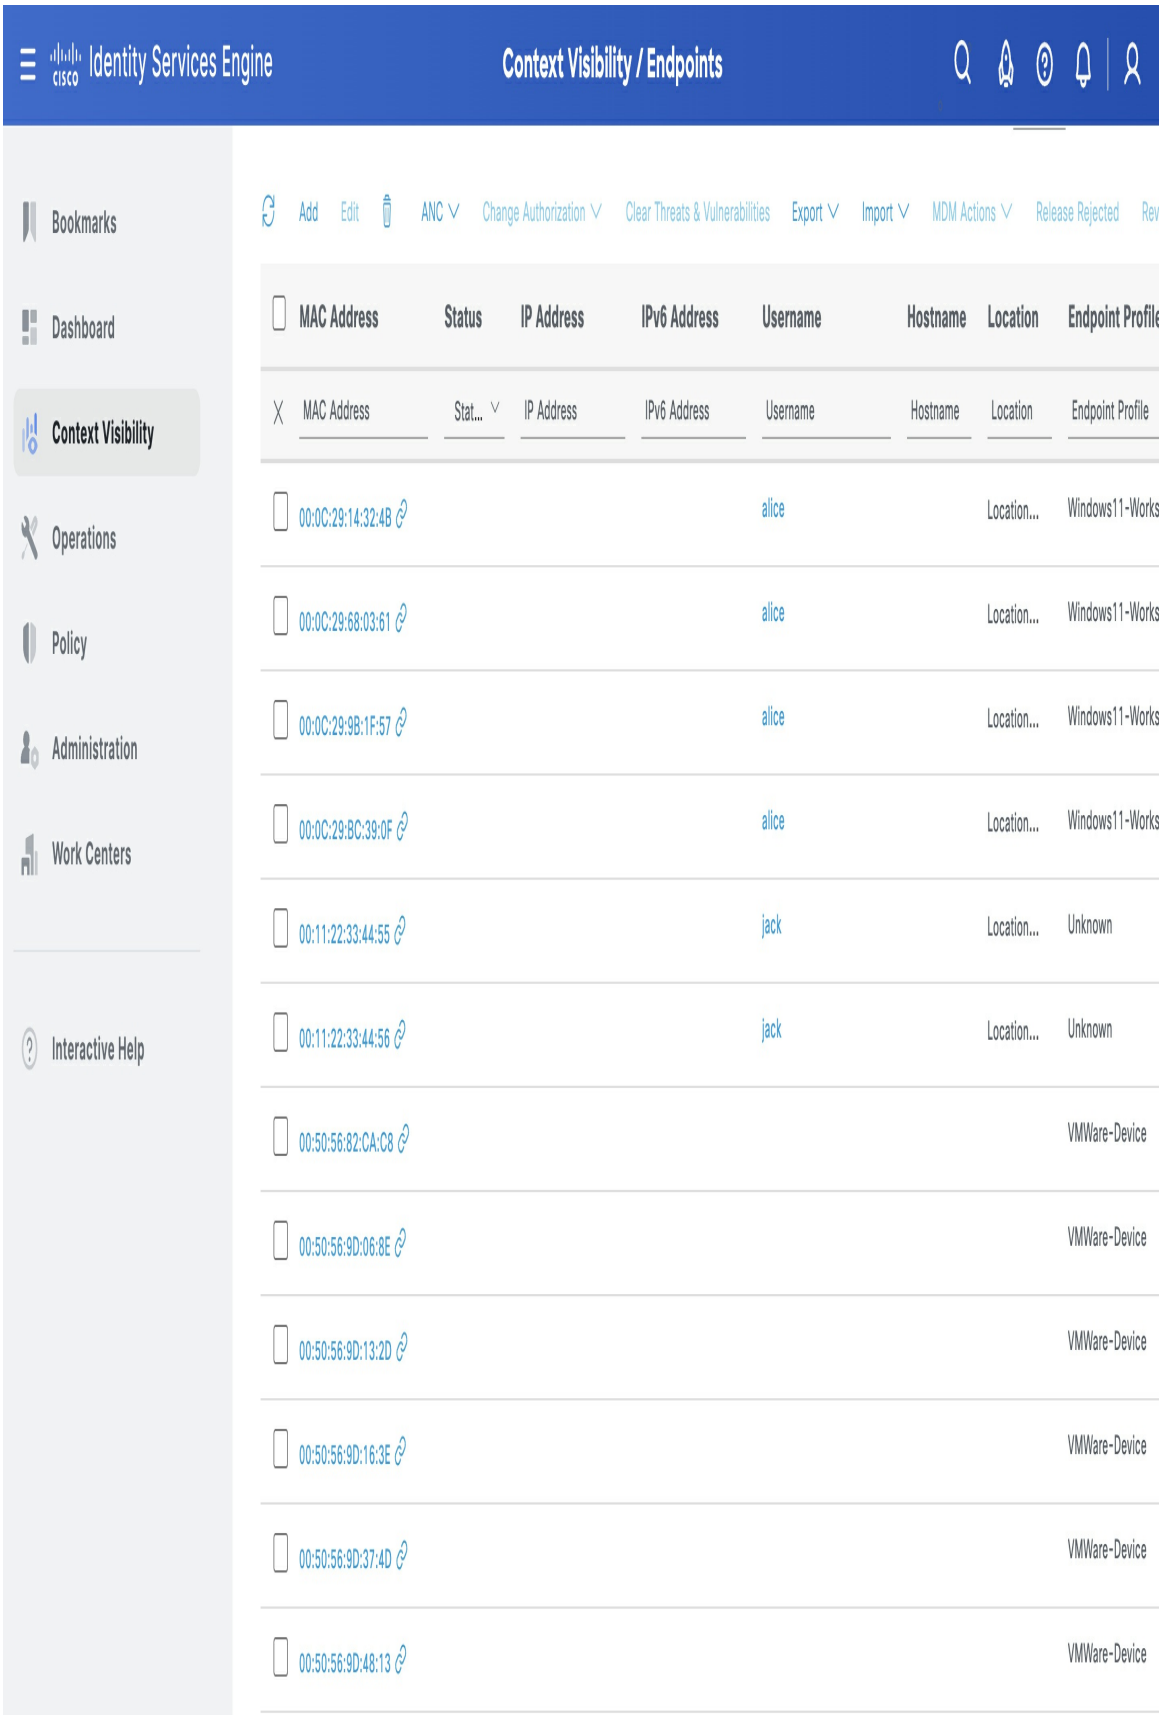

Identity Services Engine
Context Visibility / Endpoints
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Add
Edit
ANC
Change Authorization
Clear Threats & Vulnerabilities
Export
Import
MDM Actions
Release Rejected
Rev
MAC Address
Status
IP Address
IPv6 Address
Username
Hostname
Location
Endpoint Profile
MAC Address
Stat...
IP Address
IPv6 Address
Username
Hostname
Location
Endpoint Profile
00:0C:29:14:32:4B
alice
Location...
Windows11-Works
00:0C:29:68:03:61
alice
Location...
Windows11-Works
00:0C:29:9B:1F:57
alice
Location...
Windows11-Works
00:0C:29:BC:39:0F
alice
Location...
Windows11-Works
00:11:22:33:44:55
jack
Location...
Unknown
00:11:22:33:44:56
jack
Location...
Unknown
00:50:56:82:CA:C8
VMWare-Device
00:50:56:9D:06:8E
VMWare-Device
00:50:56:9D:13:2D
VMWare-Device
00:50:56:9D:16:3E
VMWare-Device
00:50:56:9D:37:4D
VMWare-Device
00:50:56:9D:48:13
VMWare-Device

Figure 5-7 *Context Visibility*

Context Visibility provides administrators with a detailed view of endpoints that ISE detects at any point in time. ISE creates an endpoint record whenever an endpoint attempts to authenticate, regardless of whether authentication succeeds. With Context Visibility, you can do the following:

- Identify connected and disconnected endpoints.
- Disconnect or reauthenticate an endpoint by using Change of Authorization (CoA).
- Clear threats and vulnerabilities.
- Release rejected endpoints.

Context Visibility populates automatically when an endpoint authenticates with ISE, and you can also populate it manually by doing the following:

- Importing endpoints via a CSV file.
- Creating a single endpoint at a time.
- Creating endpoints via the API.

Certificates

Before reviewing external identity sources and certificate-based authentication, you should first understand certificates on ISE. Certificates have historically been a challenging topic to understand and implement.

To navigate to the **Certificate Management** menu, go to **Administration > System > Certificates**, as shown in Figure 5-8.

Technet24

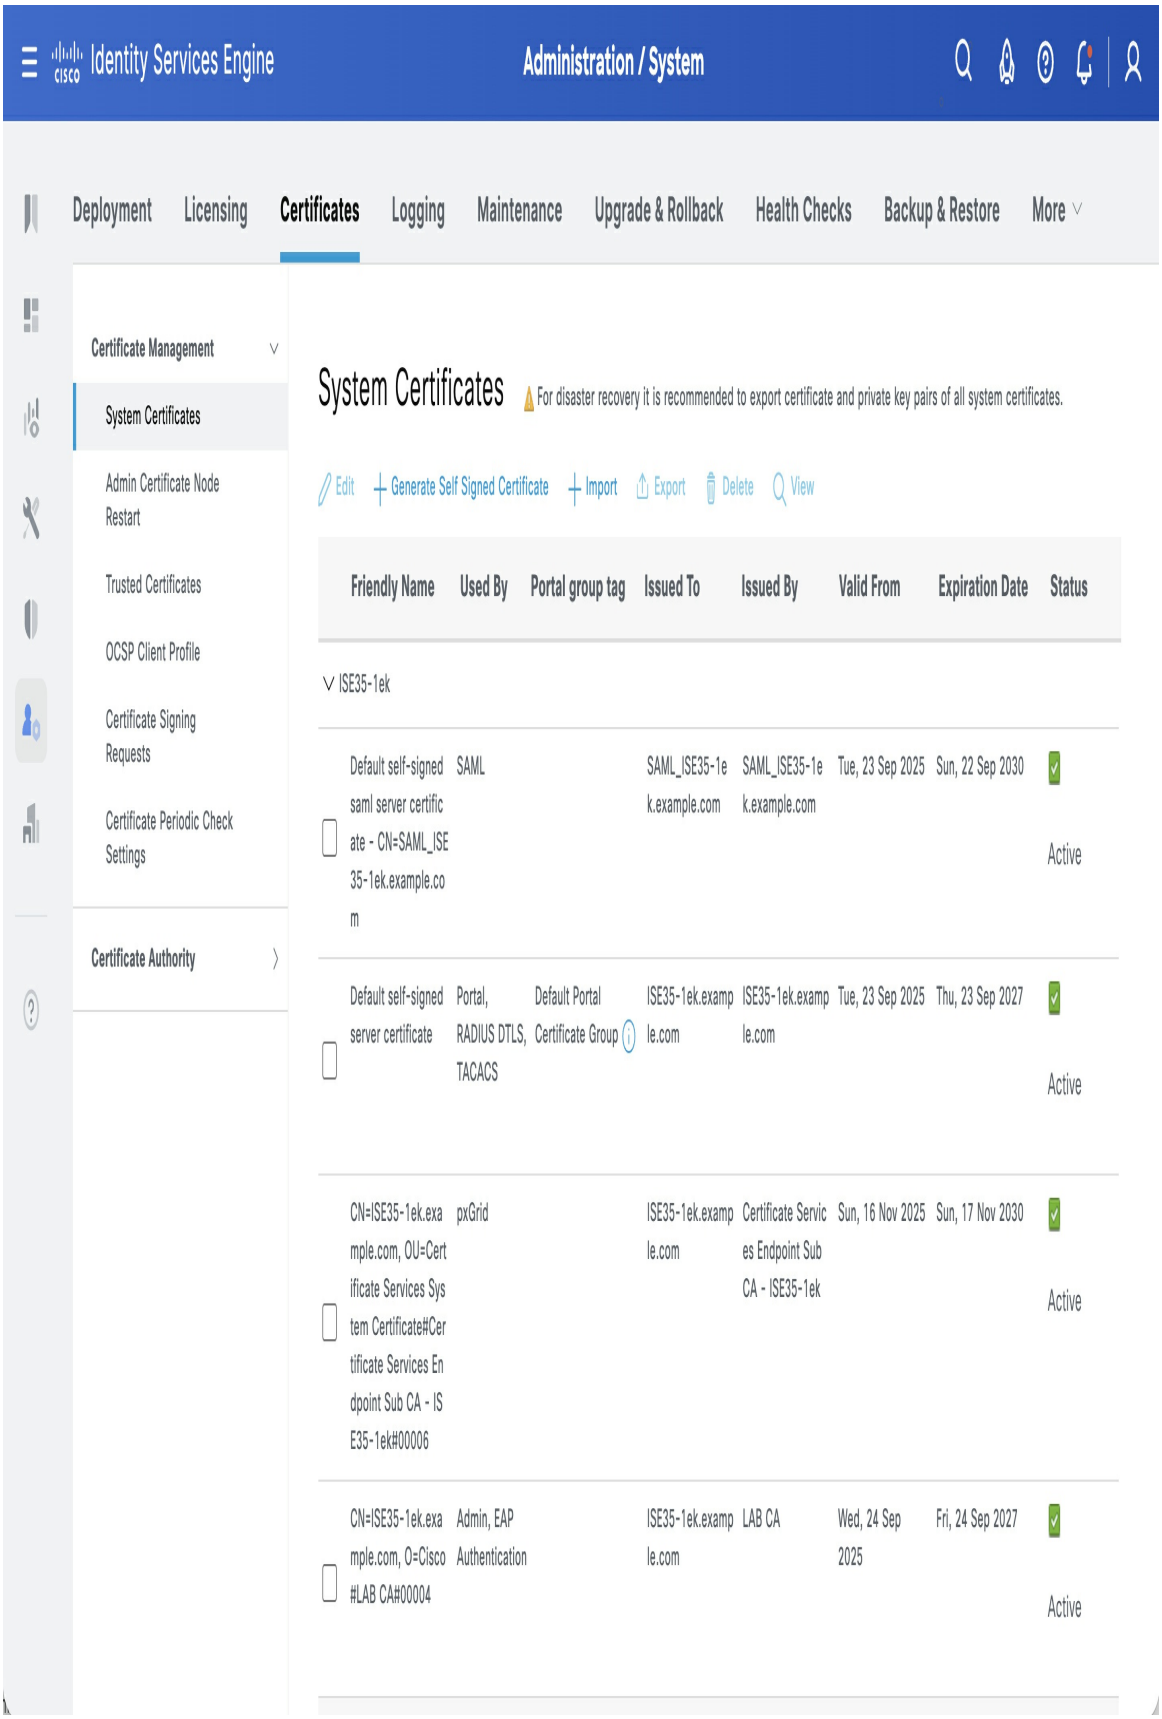

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
System Certificates
For disaster recovery it is recommended to export certificate and private key pairs of all system certificates.
Edit
Generate Self Signed Certificate
Import
Export
Delete
View
Friendly Name
Used By
Portal group tag
Issued To
Issued By
Valid From
Expiration Date
Status
ISE35-1ek
Default self-signed saml server certificate - CN=SAML_ISE35-1ek.example.com
SAML
SAML_ISE35-1ek.example.com
SAML_ISE35-1ek.example.com
Tue, 23 Sep 2025
Sun, 22 Sep 2030
Active
Default self-signed server certificate
Portal, RADIUS DTLS, TACACS
Default Portal Certificate Group
ISE35-1ek.example.com
ISE35-1ek.example.com
Tue, 23 Sep 2025
Thu, 23 Sep 2027
Active
CN=ISE35-1ek.example.com, OU=Certificate Services System Certificate#Certificate Services Endpoint Sub CA - ISE35-1ek#00006
pxGrid
ISE35-1ek.example.com
Certificate Services Endpoint Sub CA - ISE35-1ek
Sun, 16 Nov 2025
Sun, 17 Nov 2030
Active
CN=ISE35-1ek.example.com, O=Cisco#LAB CA#00004
Admin, EAP Authentication
ISE35-1ek.example.com
LAB CA
Wed, 24 Sep 2025
Fri, 24 Sep 2027
Active

Figure 5-8 *ISE Certificate Management*

ISE has two key certificate stores: the System Certificate Store, where ISE stores its own certificates, and the Trusted Certificate Store, where ISE stores certificates that ISE trusts during TLS handshakes.

System Certificate Store

ISE can assume different roles during TLS handshakes. It can act as a server — for example, when clients authenticate during EAP-TLS or present a captive portal web page during guest authentication. ISE can also act as a client — for example, when it retrieves Profiler Feed Service updates or connects to a mobile device management (MDM) server. Whenever ISE presents itself to a client or to a server, it uses certificates from the System Certificate Store.

To determine which certificate is used for a specific communication, ISE employs the concept of certificate usage. You can configure a single certificate for multiple functions or assign dedicated certificates to specific functions.

You can specify which certificates you want to use for the following functions:

- **Admin:** This certificate is used when you build a distributed deployment, when ISE connects to external services such as MDM or Posture Feed services, and when ISE presents the Admin portal.
- **EAP Authentication:** This certificate is used for EAP authentication protocols such as PEAP, EAP-TLS, and TEAP, where ISE presents its identity during SSL/TLS tunneling.
- **RADIUS DTLS:** This certificate is used to establish a secure connection between network access devices and ISE when the RADIUS DTLS protocol is used.

Technet24

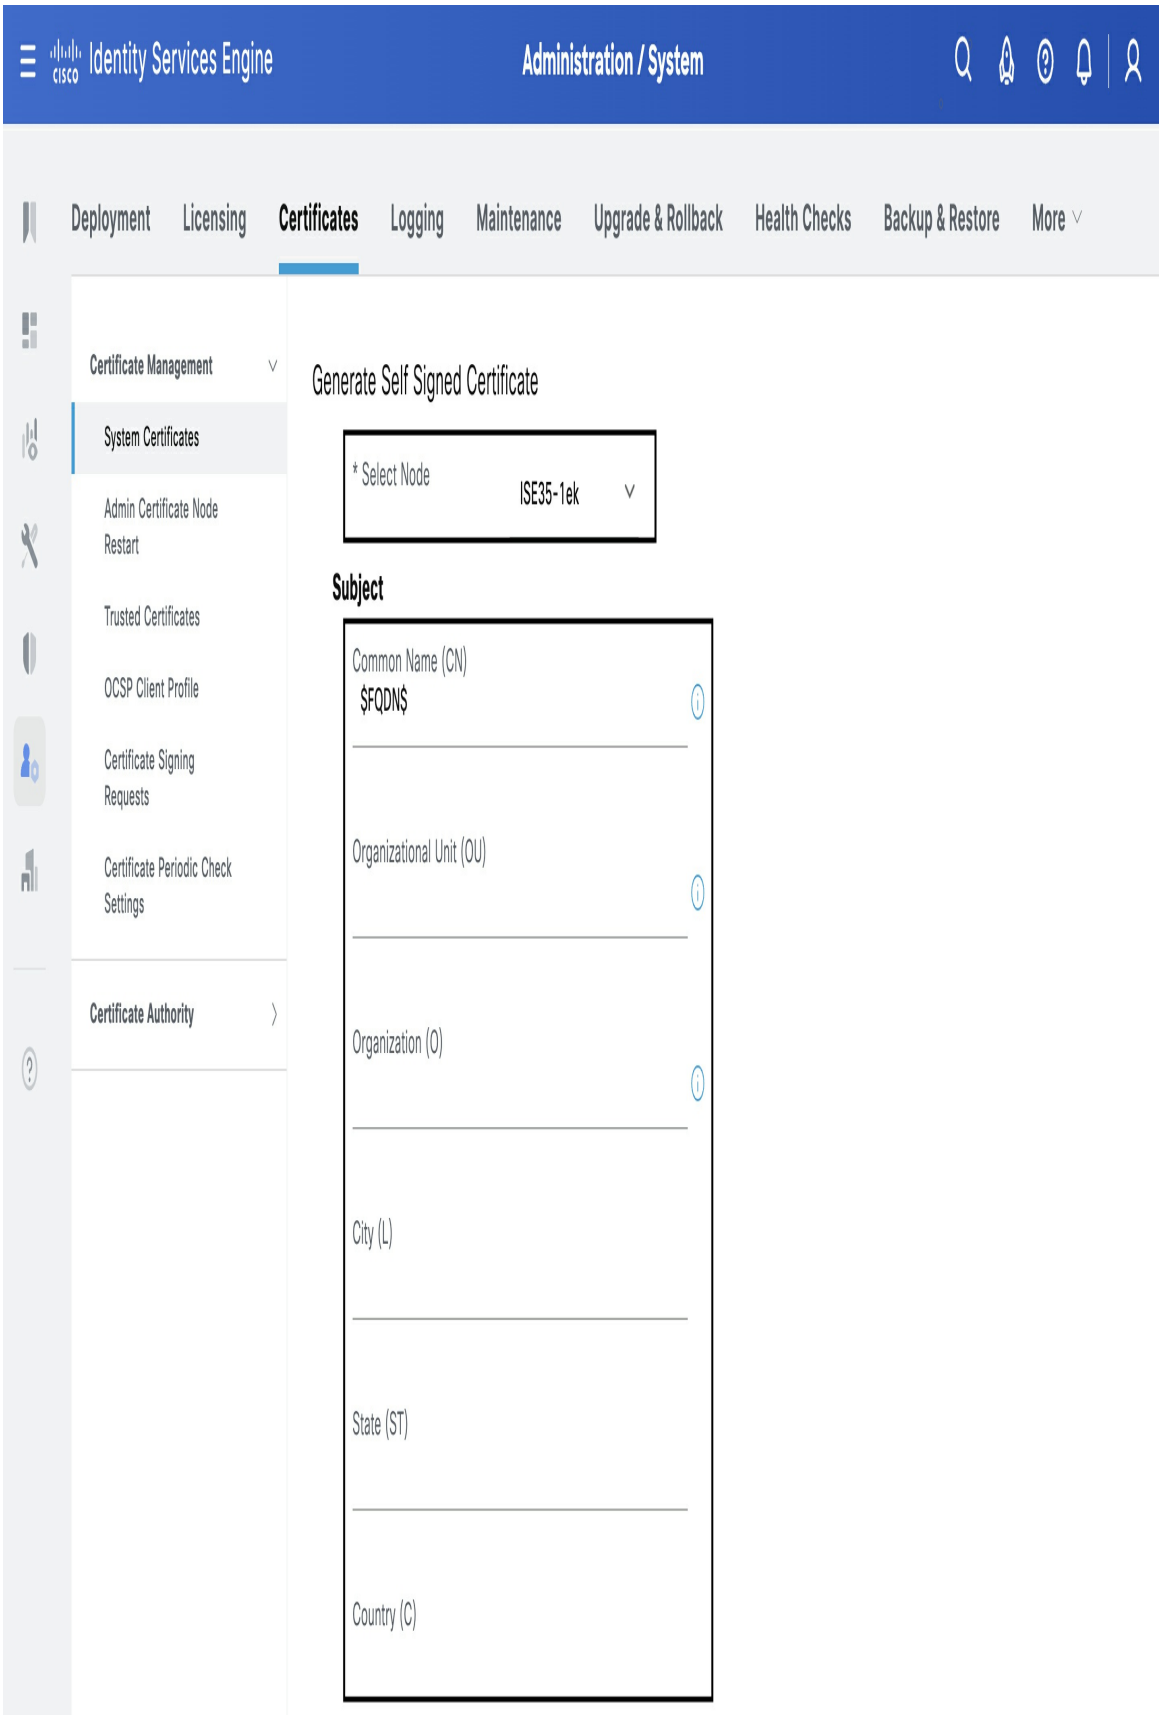
- **TACACS:** This certificate is used to establish a secure connection between network access devices and ISE when TACACS over TLS is used.
- **pxGrid:** This certificate is used during pxGrid integration for both client-side and server-side communication by ISE.
- **ISE Messaging Service:** This certificate is used to establish communication channels for internal services, such as transferring logs from PSN nodes to MNT nodes, in the deployment.
- **NativeIPsec:** This certificate is used to establish IPsec tunnels between network access devices and ISE when configured.
- **SAML:** This certificate is used for SAML signing.
- **Portal:** This certificate is used to present secure portal pages. Different certificates for different portals can be used by leveraging portal tags.

Out of the box, ISE includes most of these certificates. They consist of a mixture of self-signed certificates and internal CA-signed certificates.

Best Practice:

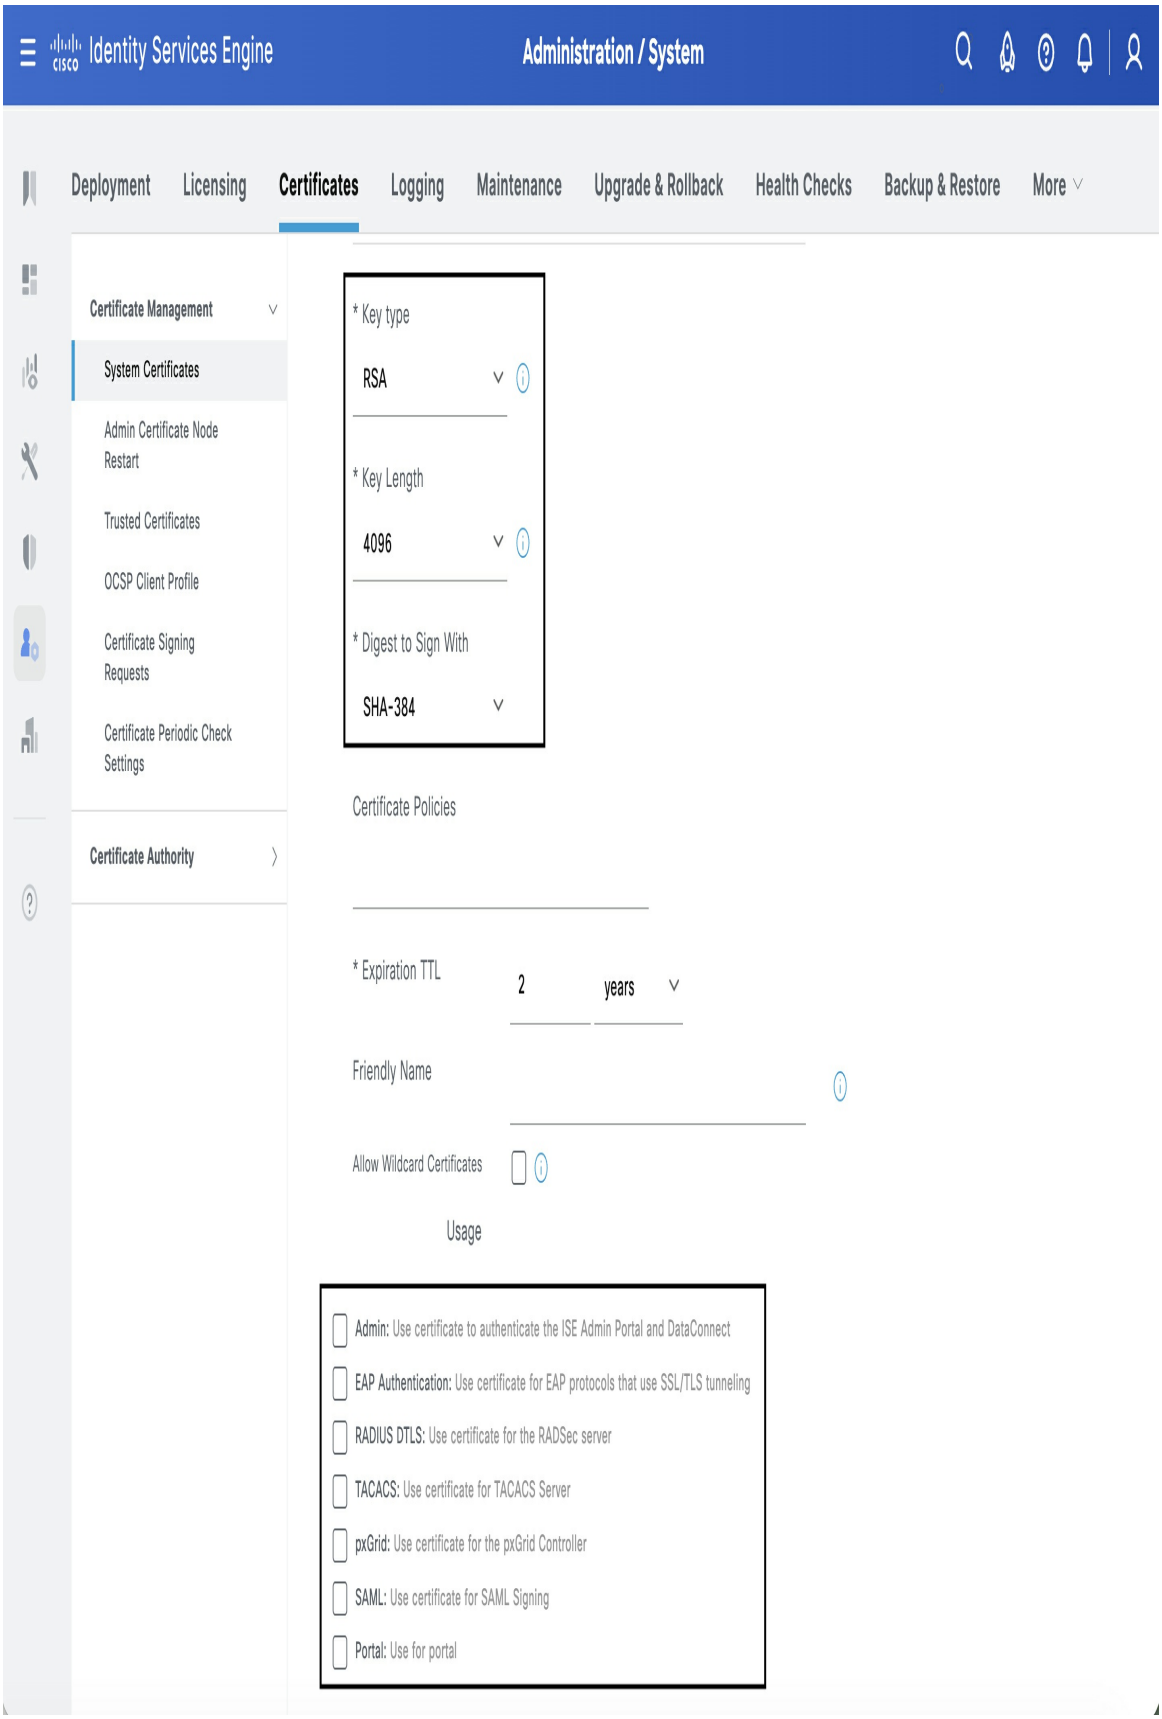
It is recommended to use CA-signed certificates, either internal or external, whenever possible, instead of self-signed certificates.

To generate a self-signed certificate, navigate to **Administration > System > Certificates > Certificate Management > System Certificates** and click **Generate Self Signed Certificate**. Populate the subject, select the node on which the certificate will be installed, provide key information, and select the appropriate usage, as shown in Figures 5-9 and 5-10.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
Generate Self Signed Certificate
* Select Node
ISE35-1ek
Subject
Common Name (CN)
$FQDN$
Organizational Unit (OU)
Organization (O)
City (L)
State (ST)
Country (C)

Technet24

Figure 5-9 *ISE Self-Signed Certificate Installation*

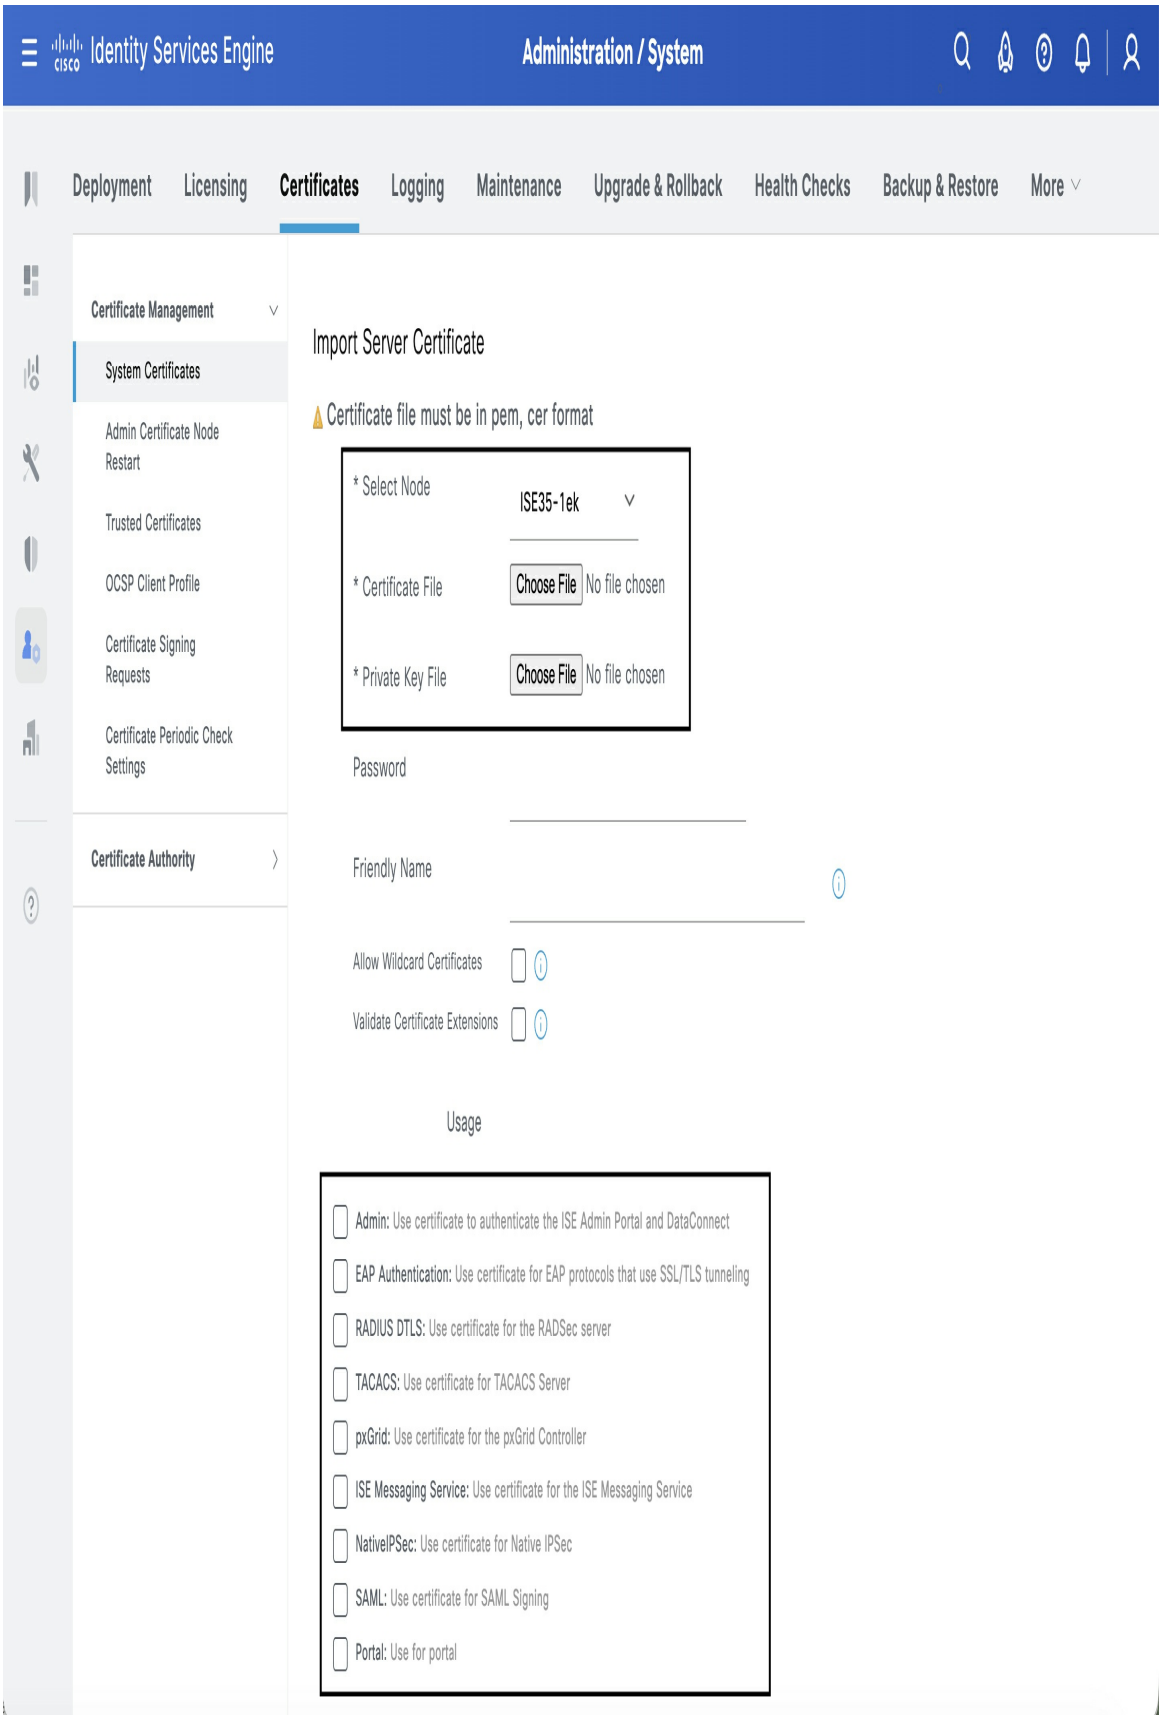

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
* Key type
RSA
* Key Length
4096
* Digest to Sign With
SHA-384
Certificate Policies
* Expiration TTL
2
years
Friendly Name
Allow Wildcard Certificates
Usage
Admin: Use certificate to authenticate the ISE Admin Portal and DataConnect
EAP Authentication: Use certificate for EAP protocols that use SSL/TLS tunneling
RADIUS DTLS: Use certificate for the RADSec server
TACACS: Use certificate for TACACS Server
pxGrid: Use certificate for the pxGrid Controller
SAML: Use certificate for SAML Signing
Portal: Use for portal

Technet24

Figure 5-10 *ISE Self-Signed Certificate Installation*

You can install a CA-signed system certificate on ISE in two ways. The first option is to navigate to **Administration > System > Certificates > Certificate Management > System Certificates** and click **Install**. You use this option when you do not generate the key pair locally on ISE and the certificate authority provides both the digital certificate and the private key.

You need to select the node, specify the certificate file and private key file, provide a password if required, and choose the usage, as shown in Figure 5-11.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
Import Server Certificate
Certificate file must be in pem, cer format
* Select Node
ISE35-1ek
* Certificate File
Choose File
No file chosen
* Private Key File
Choose File
No file chosen
Password
Friendly Name
Allow Wildcard Certificates
Validate Certificate Extensions
Usage
Admin: Use certificate to authenticate the ISE Admin Portal and DataConnect
EAP Authentication: Use certificate for EAP protocols that use SSL/TLS tunneling
RADIUS DTLS: Use certificate for the RADSec server
TACACS: Use certificate for TACACS Server
pxGrid: Use certificate for the pxGrid Controller
ISE Messaging Service: Use certificate for the ISE Messaging Service
NativeIPSec: Use certificate for Native IPSec
SAML: Use certificate for SAML Signing
Portal: Use for portal

Technet24

Figure 5-11 *ISE System Certificate Installation*

During the installation, you should select the **Allow Wildcard Certificates** checkbox if the certificate you are installing is a wildcard certificate. When you select this option, ISE replicates the certificate to all nodes in the deployment.

Wildcard certificates have a wildcard notation, consisting of an asterisk and a period, before the certificate domain name, which allows the certificate to be used by multiple hosts. Using wildcard certificates is considered generally less secure than assigning individual certificates for ISE nodes. An alternative to wildcard certificates is a certificate with multiple SAN fields, where you include the fully qualified domain name (FQDN) of each ISE node.

> ## Note
>
> When you install an admin certificate on the PAN node, you must restart all nodes. The certificate installation menu prompts you to either restart the individual nodes immediately or select a specific time and restart later. When you install an admin certificate on any other node in the deployment, only that node restarts.

For example, if your domain is iselab.net and you have multiple PSNs, you can use a single certificate with a CN or SAN field of *.iselab.net for all of them.

The second option is to generate the key pair locally on ISE and have the certificate authority sign the certificate. To do this, you first generate a certificate signing request (CSR). Navigate to **Administration > System > Certificates > Certificate Management > Certificate Signing Requests** and select **Generate Certificate Signing Requests (CSR)**.

As shown in Figure 5-12, CSR generation closely resembles self-signed certificate generation. You select the nodes and provide the key information and subject. In addition to the standard system certificate usages, the CSR menu introduces additional options. From this menu, you can do the following:

- Renew ISE internal certificates by selecting **ISE Root CA**.
- Renew ISE OCSP responder certificates by selecting **Renew ISE OCSP Responder Certificates**.
- Generate an **ISE Intermediate CA** if you want the ISE internal CA to be part of an enterprise CA.

Technet24

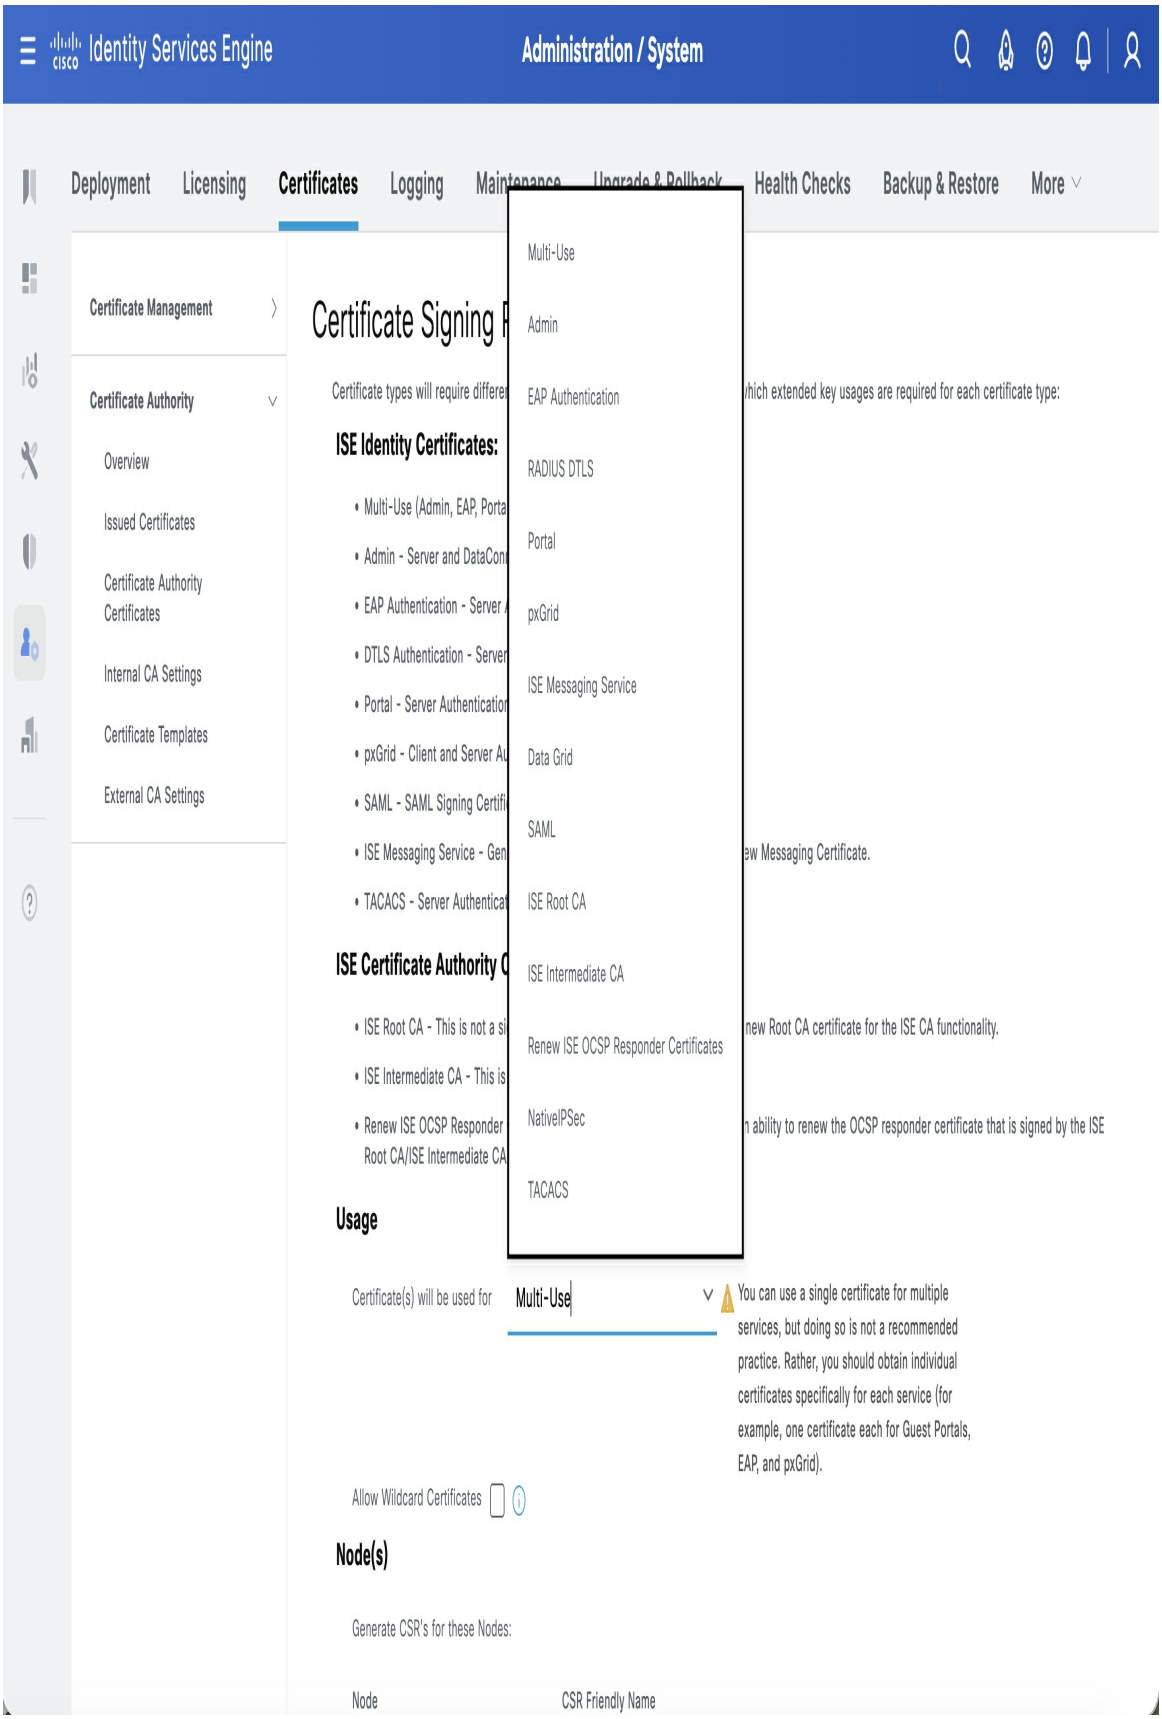

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
Certificate Authority
Overview
Issued Certificates
Certificate Authority Certificates
Internal CA Settings
Certificate Templates
External CA Settings
Certificate Signing
Certificate types will require differe
hich extended key usages are required for each certificate type:
ISE Identity Certificates:
• Multi-Use (Admin, EAP, Porta
• Admin - Server and DataCon
• EAP Authentication - Server
• DTLS Authentication - Server
• Portal - Server Authenticatio
• pxGrid - Client and Server Au
• SAML - SAML Signing Certifi
• ISE Messaging Service - Gen
ew Messaging Certificate.
• TACACS - Server Authentica
ISE Certificate Authority
• ISE Root CA - This is not a si
new Root CA certificate for the ISE CA functionality.
• ISE Intermediate CA - This is
• Renew ISE OCSP Responder
Root CA/ISE Intermediate CA
n ability to renew the OCSP responder certificate that is signed by the ISE
Usage
Certificate(s) will be used for
Multi-Use
You can use a single certificate for multiple services, but doing so is not a recommended practice. Rather, you should obtain individual certificates specifically for each service (for example, one certificate each for Guest Portals, EAP, and pxGrid).
Allow Wildcard Certificates
Node(s)
Generate CSR's for these Nodes:
Node
CSR Friendly Name
Multi-Use
Admin
EAP Authentication
RADIUS DTLS
Portal
pxGrid
ISE Messaging Service
Data Grid
SAML
ISE Root CA
ISE Intermediate CA
Renew ISE OCSP Responder Certificates
NativeIPSec
TACACS

Figure 5-12 *Generating a Certificate Signing Request*

After you generate the CSR, you must complete two additional steps. First, export the CSR and provide it to the certificate authority for signing. You can do this either by clicking **View** or by exporting it as a file, as shown in Figure 5-13.

Technet24

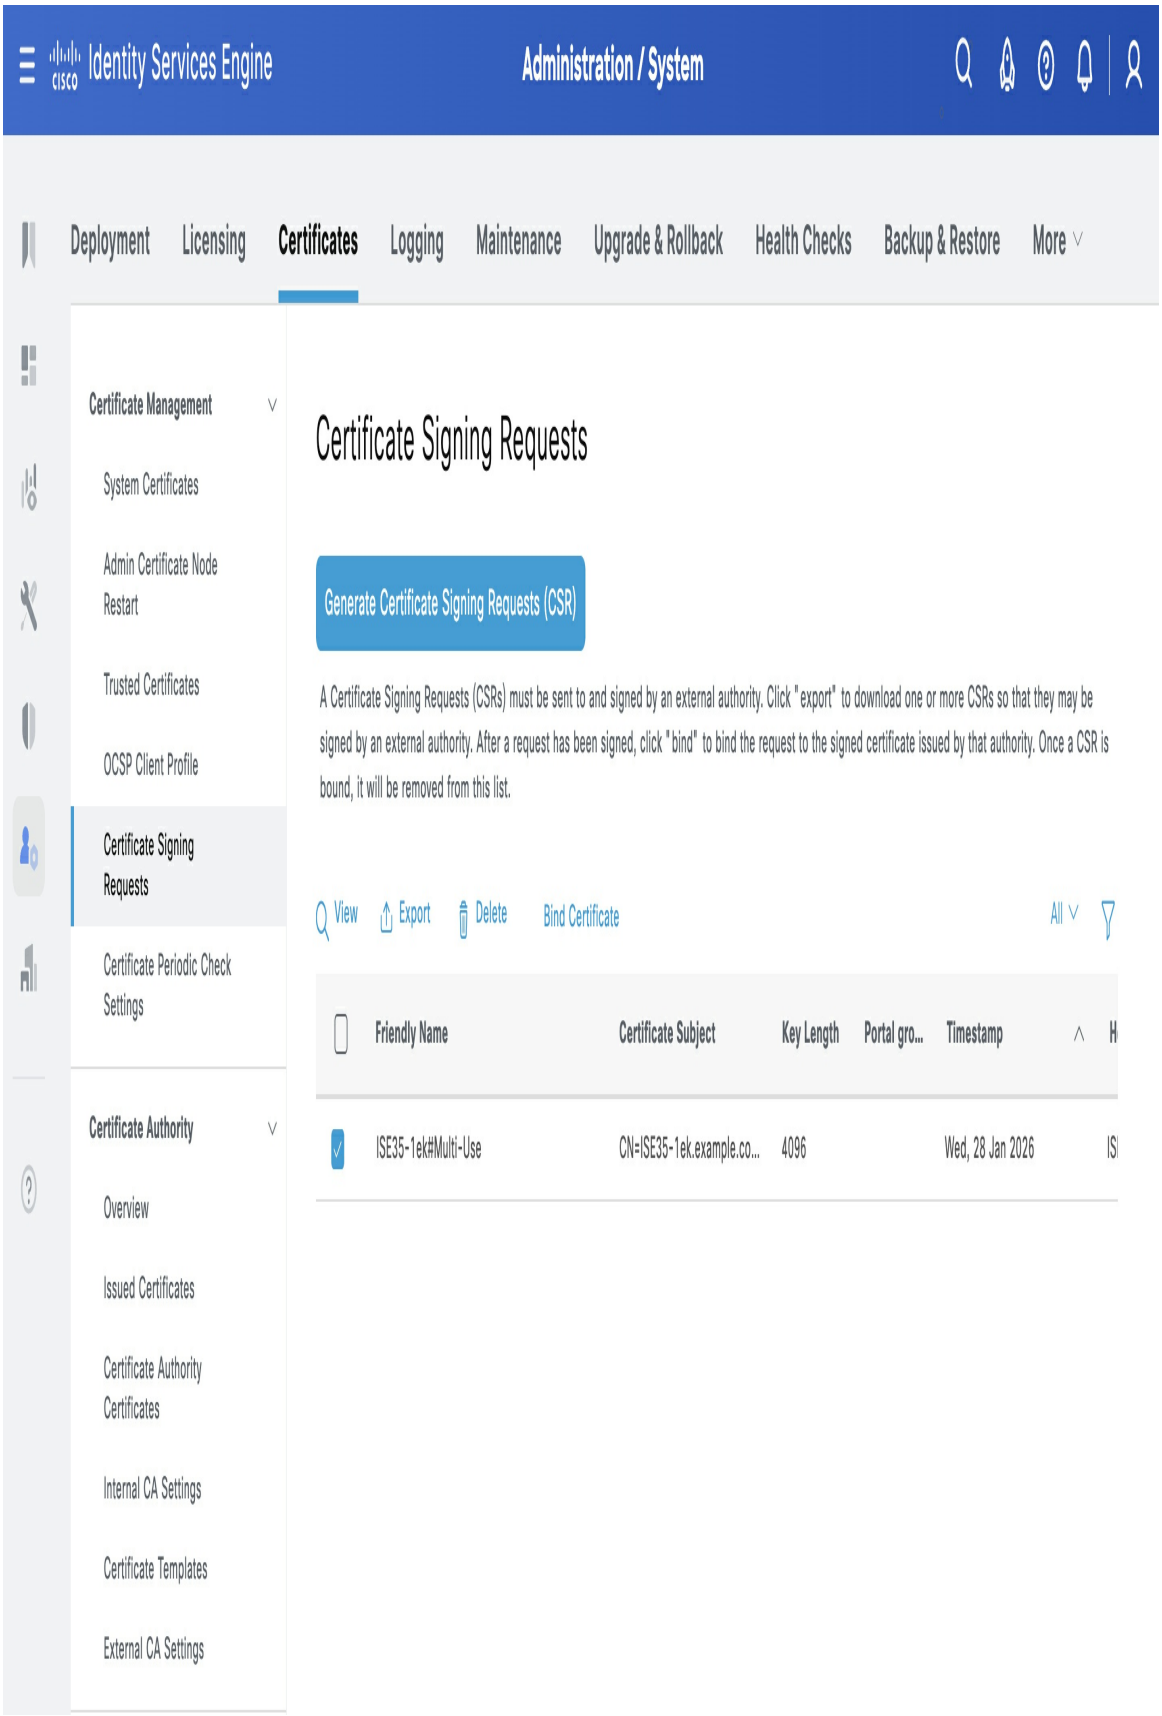

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
Overview
Issued Certificates
Certificate Authority Certificates
Internal CA Settings
Certificate Templates
External CA Settings
Certificate Signing Requests
Generate Certificate Signing Requests (CSR)
A Certificate Signing Requests (CSRs) must be sent to and signed by an external authority. Click "export" to download one or more CSRs so that they may be signed by an external authority. After a request has been signed, click "bind" to bind the request to the signed certificate issued by that authority. Once a CSR is bound, it will be removed from this list.
View
Export
Delete
Bind Certificate
All
Friendly Name
Certificate Subject
Key Length
Portal gro...
Timestamp
ISE35-1ek#Multi-Use
CN=ISE35-1ek.example.co...
4096
Wed, 28 Jan 2026

Figure 5-13 *CSR Export*

After the certificate authority signs the CSR, you bind the certificate to the private key by selecting **Bind Certificate**. As shown in Figure 5-14, this step requires minimal configuration. Because the certificate is already issued and the keys reside on ISE, you only need to select the certificate and assign a usage.

Technet24

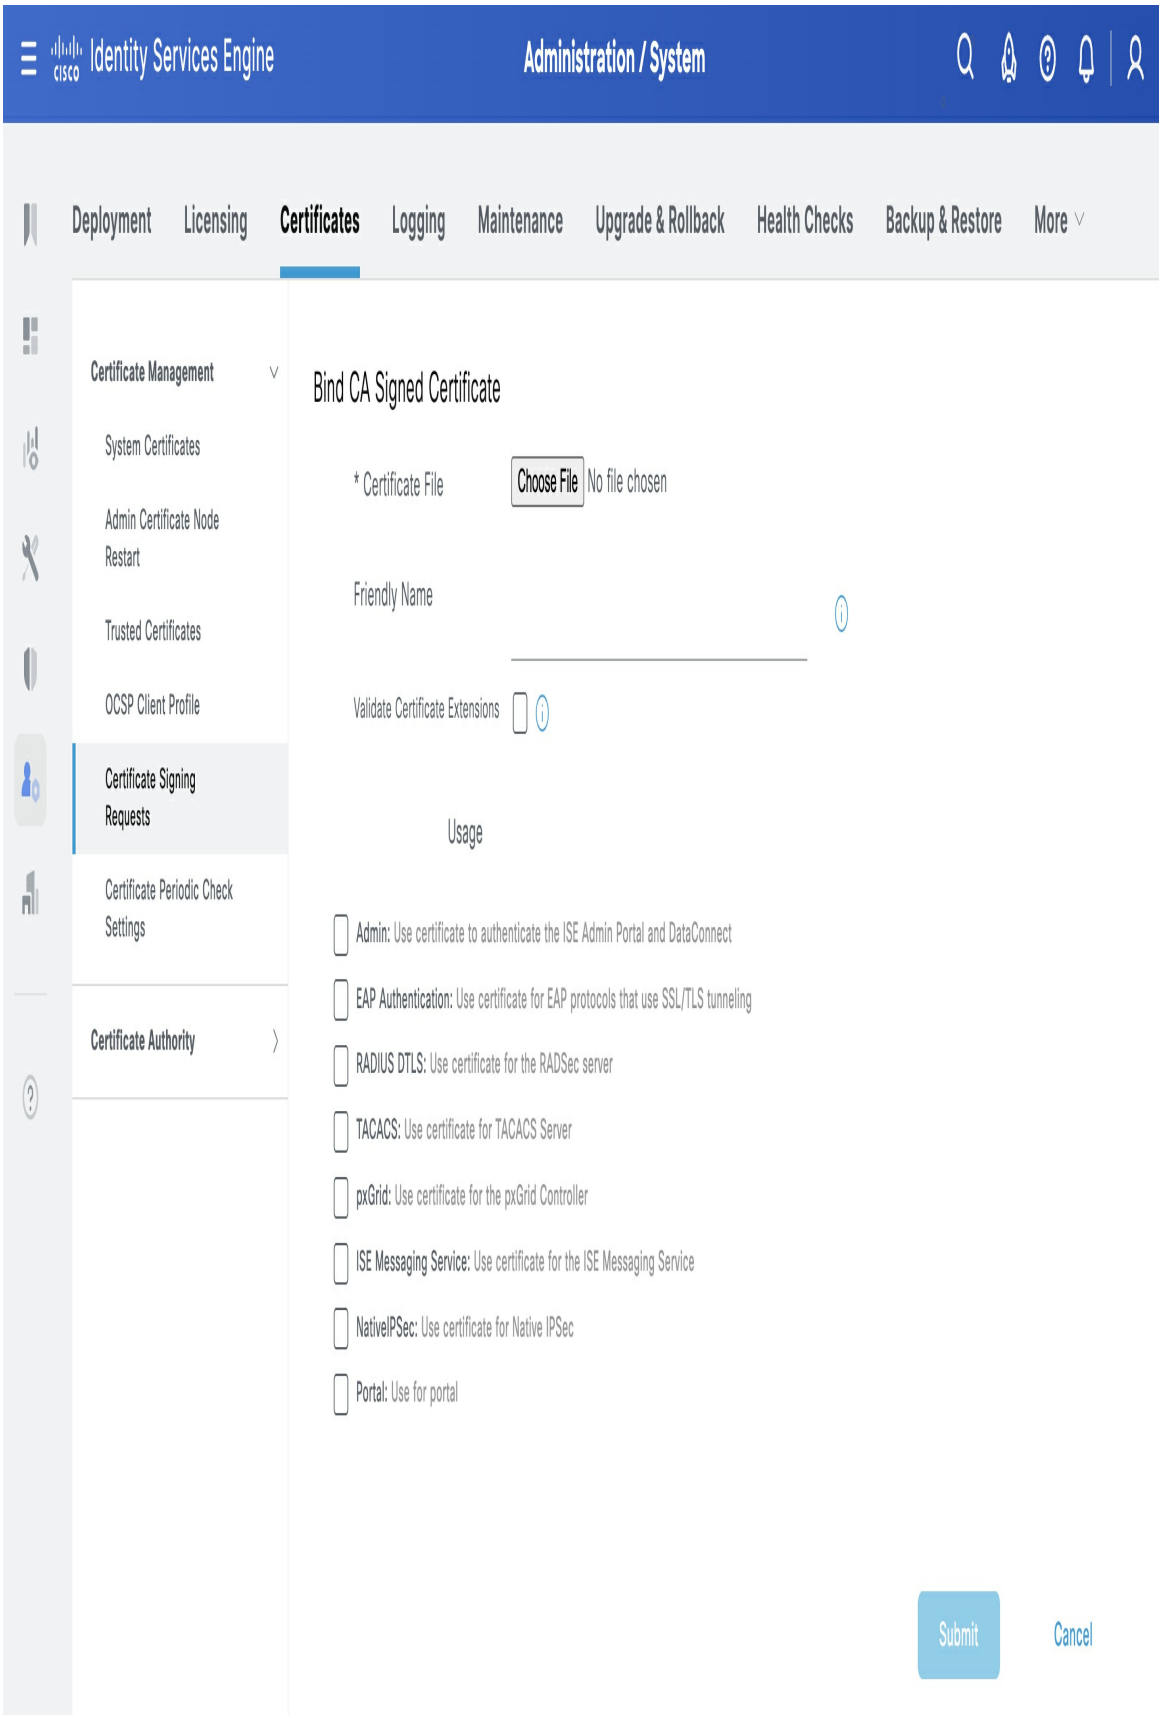

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
Bind CA Signed Certificate
* Certificate File
Choose File
No file chosen
Friendly Name
Validate Certificate Extensions
Usage
Admin: Use certificate to authenticate the ISE Admin Portal and DataConnect
EAP Authentication: Use certificate for EAP protocols that use SSL/TLS tunneling
RADIUS DTLS: Use certificate for the RADSec server
TACACS: Use certificate for TACACS Server
pxGrid: Use certificate for the pxGrid Controller
ISE Messaging Service: Use certificate for the ISE Messaging Service
NativeIPSec: Use certificate for Native IPSec
Portal: Use for portal
Submit
Cancel

Figure 5-14 *Certificate Bind*

Note

You must install all certificates that form the system certificate chain in the Trusted Certificate Store before you install the system certificate itself.

Trusted Certificate Store

The Trusted Certificate Store is used when ISE needs to validate certificates presented by a client or a server, depending on the direction of communication. For example, ISE must trust the issuer of an EAP-TLS client certificate, or it must trust the issuer of an SSL/TLS server certificate when it connects to an MDM server.

Similar to Browser Certificate Stores, ISE includes a preconfigured set of trusted certificates by default. To access the Trusted Certificate Store, navigate to **Administration > System > Certificates > Trusted Certificates**, as shown in Figure 5-15.

Technet24

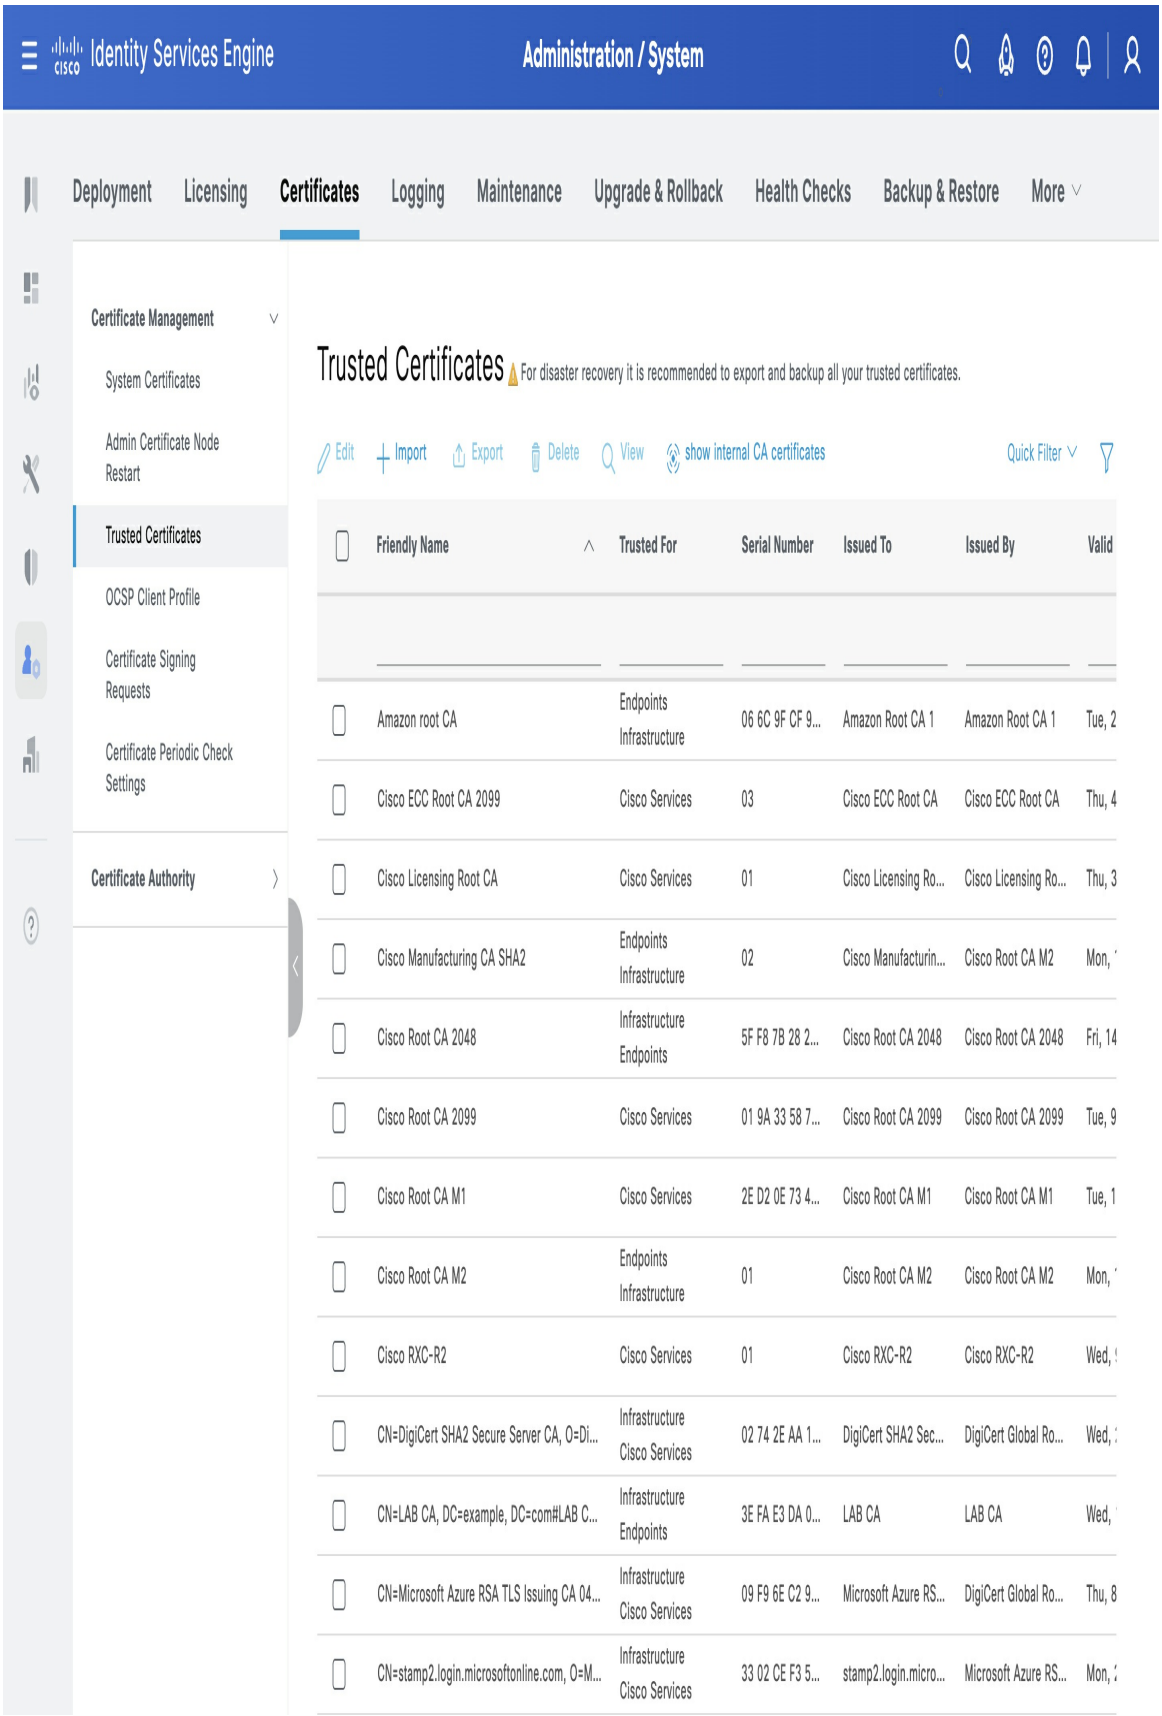

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
Trusted Certificates
For disaster recovery it is recommended to export and backup all your trusted certificates.
Edit
Import
Export
Delete
View
show internal CA certificates
Quick Filter
Friendly Name
Trusted For
Serial Number
Issued To
Issued By
Valid
Amazon root CA
Endpoints Infrastructure
06 6C 9F CF 9...
Amazon Root CA 1
Amazon Root CA 1
Tue, 2
Cisco ECC Root CA 2099
Cisco Services
03
Cisco ECC Root CA
Cisco ECC Root CA
Thu, 4
Cisco Licensing Root CA
Cisco Services
01
Cisco Licensing Ro...
Cisco Licensing Ro...
Thu, 3
Cisco Manufacturing CA SHA2
Endpoints Infrastructure
02
Cisco Manufacturin...
Cisco Root CA M2
Mon,
Cisco Root CA 2048
Infrastructure Endpoints
5F F8 7B 28 2...
Cisco Root CA 2048
Cisco Root CA 2048
Fri, 14
Cisco Root CA 2099
Cisco Services
01 9A 33 58 7...
Cisco Root CA 2099
Cisco Root CA 2099
Tue, 9
Cisco Root CA M1
Cisco Services
2E D2 0E 73 4...
Cisco Root CA M1
Cisco Root CA M1
Tue, 1
Cisco Root CA M2
Endpoints Infrastructure
01
Cisco Root CA M2
Cisco Root CA M2
Mon,
Cisco RXC-R2
Cisco Services
01
Cisco RXC-R2
Cisco RXC-R2
Wed,
CN=DigiCert SHA2 Secure Server CA, O=Di...
Infrastructure Cisco Services
02 74 2E AA 1...
DigiCert SHA2 Sec...
DigiCert Global Ro...
Wed,
CN=LAB CA, DC=example, DC=com#LAB C...
Infrastructure Endpoints
3E FA E3 DA 0...
LAB CA
LAB CA
Wed,
CN=Microsoft Azure RSA TLS Issuing CA 04...
Infrastructure Cisco Services
09 F9 6E C2 9...
Microsoft Azure RS...
DigiCert Global Ro...
Thu, 8
CN=stamp2.login.microsoftonline.com, O=M...
Infrastructure Cisco Services
33 02 CE F3 5...
stamp2.login.micro...
Microsoft Azure RS...
Mon,

Figure 5-15 *Trusted Certificate Store*

To populate the Certificate Store, select + **Import** on the same page. It is crucial that you select the required checkboxes during installation, as shown in Figure 5-16.

Technet24

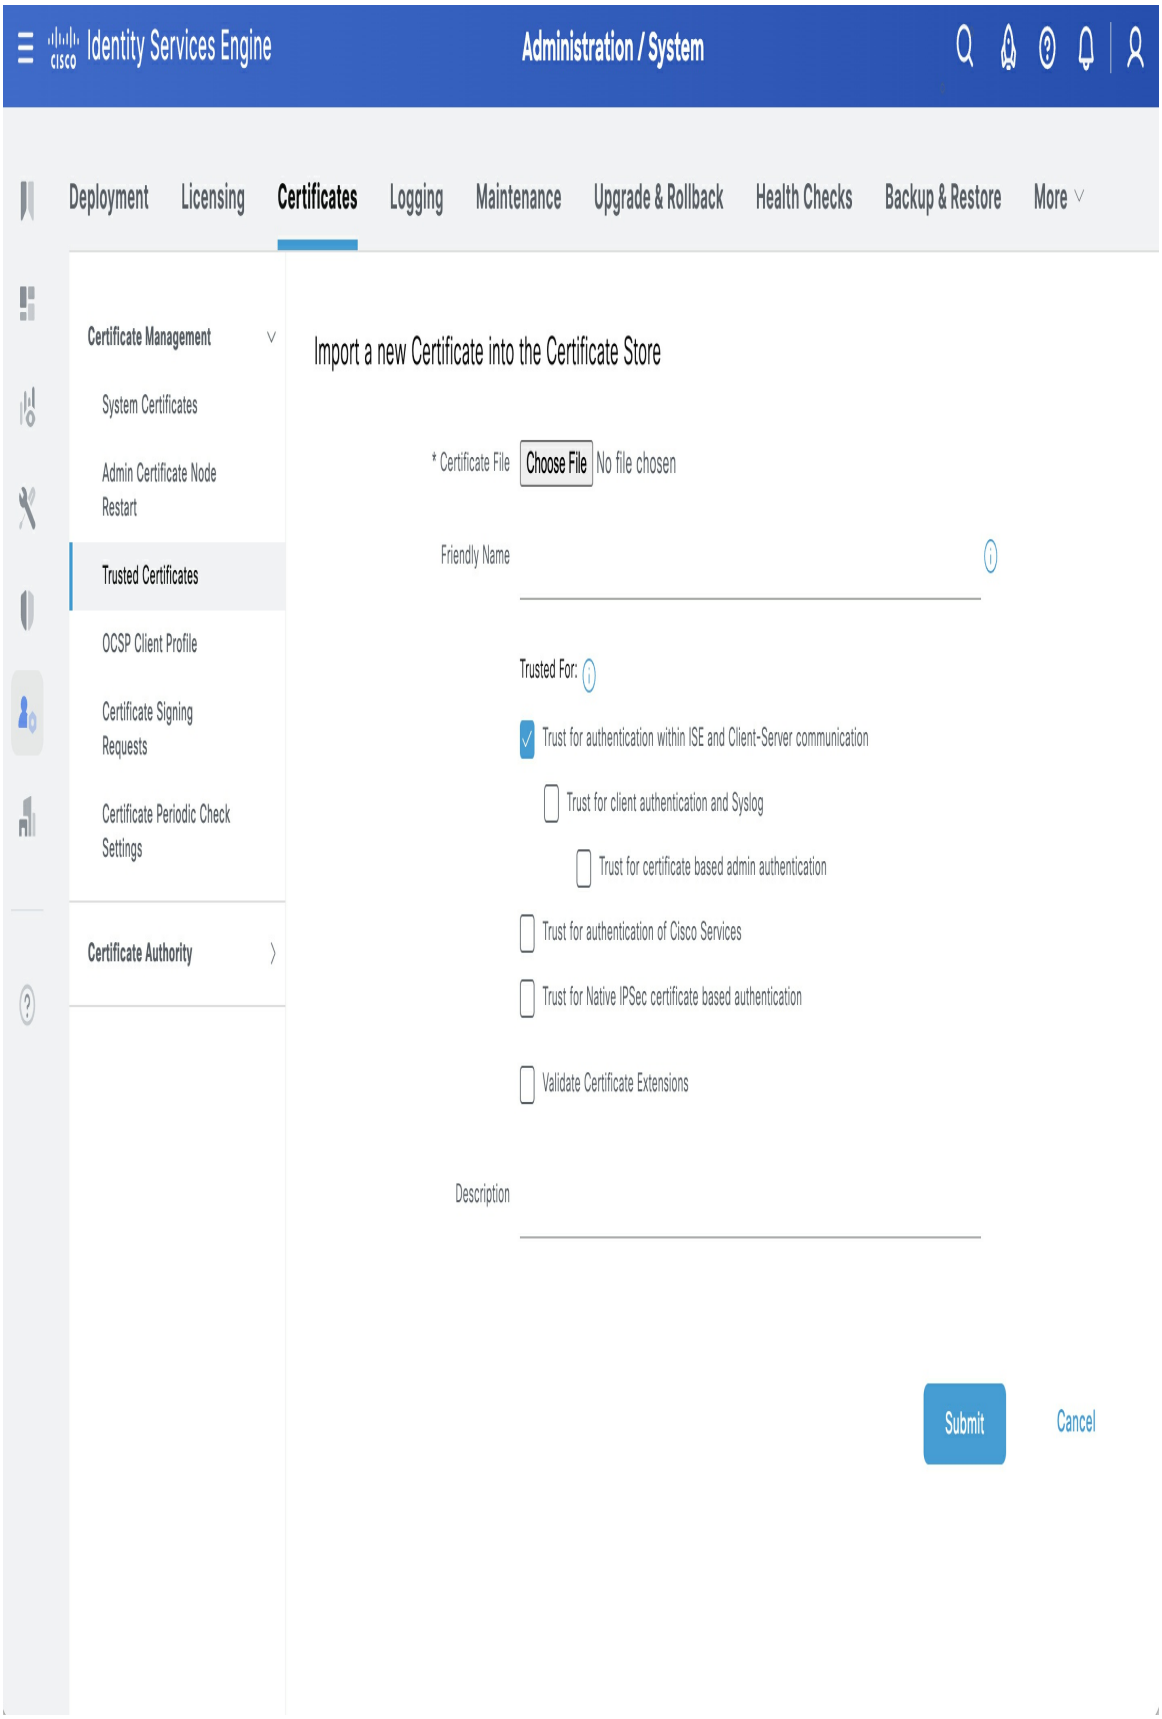

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
Import a new Certificate into the Certificate Store
* Certificate File
Choose File
No file chosen
Friendly Name
Trusted For:
Trust for authentication within ISE and Client-Server communication
Trust for client authentication and Syslog
Trust for certificate based admin authentication
Trust for authentication of Cisco Services
Trust for Native IPSec certificate based authentication
Validate Certificate Extensions
Description
Submit
Cancel

Figure 5-16 *Trusted Certificate Import*

Table 5-1 explains the purpose of each **Trusted For** usage.

Table 5-1 *Trusted Certificate Usage Guidelines*

Field Name	Purpose
Trust for Authentication Within ISE and Client-Server Communication	Select this checkbox to verify server certificates of other ISE nodes in the deployment or while communicating with LDAPs servers.
Trust for Client Authentication and Syslog	Select this checkbox to verify client certificates during EAP authentication or to verify the Syslog Server certificate.
Trust for Certificate-Based Admin Authentication	Select this checkbox to verify the client certificate during certificate-based authentication for ISE admin access.
Trust for Authentication of Cisco Services	Select this checkbox to verify server certificates of the external services, such as Feed Service.
Trust for Native IPsec Certificate-Based Authentication	Select this checkbox to verify the VPN client or server certificate during NAD <> ISE IPsec tunnel establishment.
Validate Certificate Extensions	Select this checkbox to verify that client certificate has keyUsage present and keyCertSign bit set and that the basicConstrains extension is present with the CA flag set to true.

Trusted certificates are stored on the PAN node and replicated across the deployment.

Best Practice:

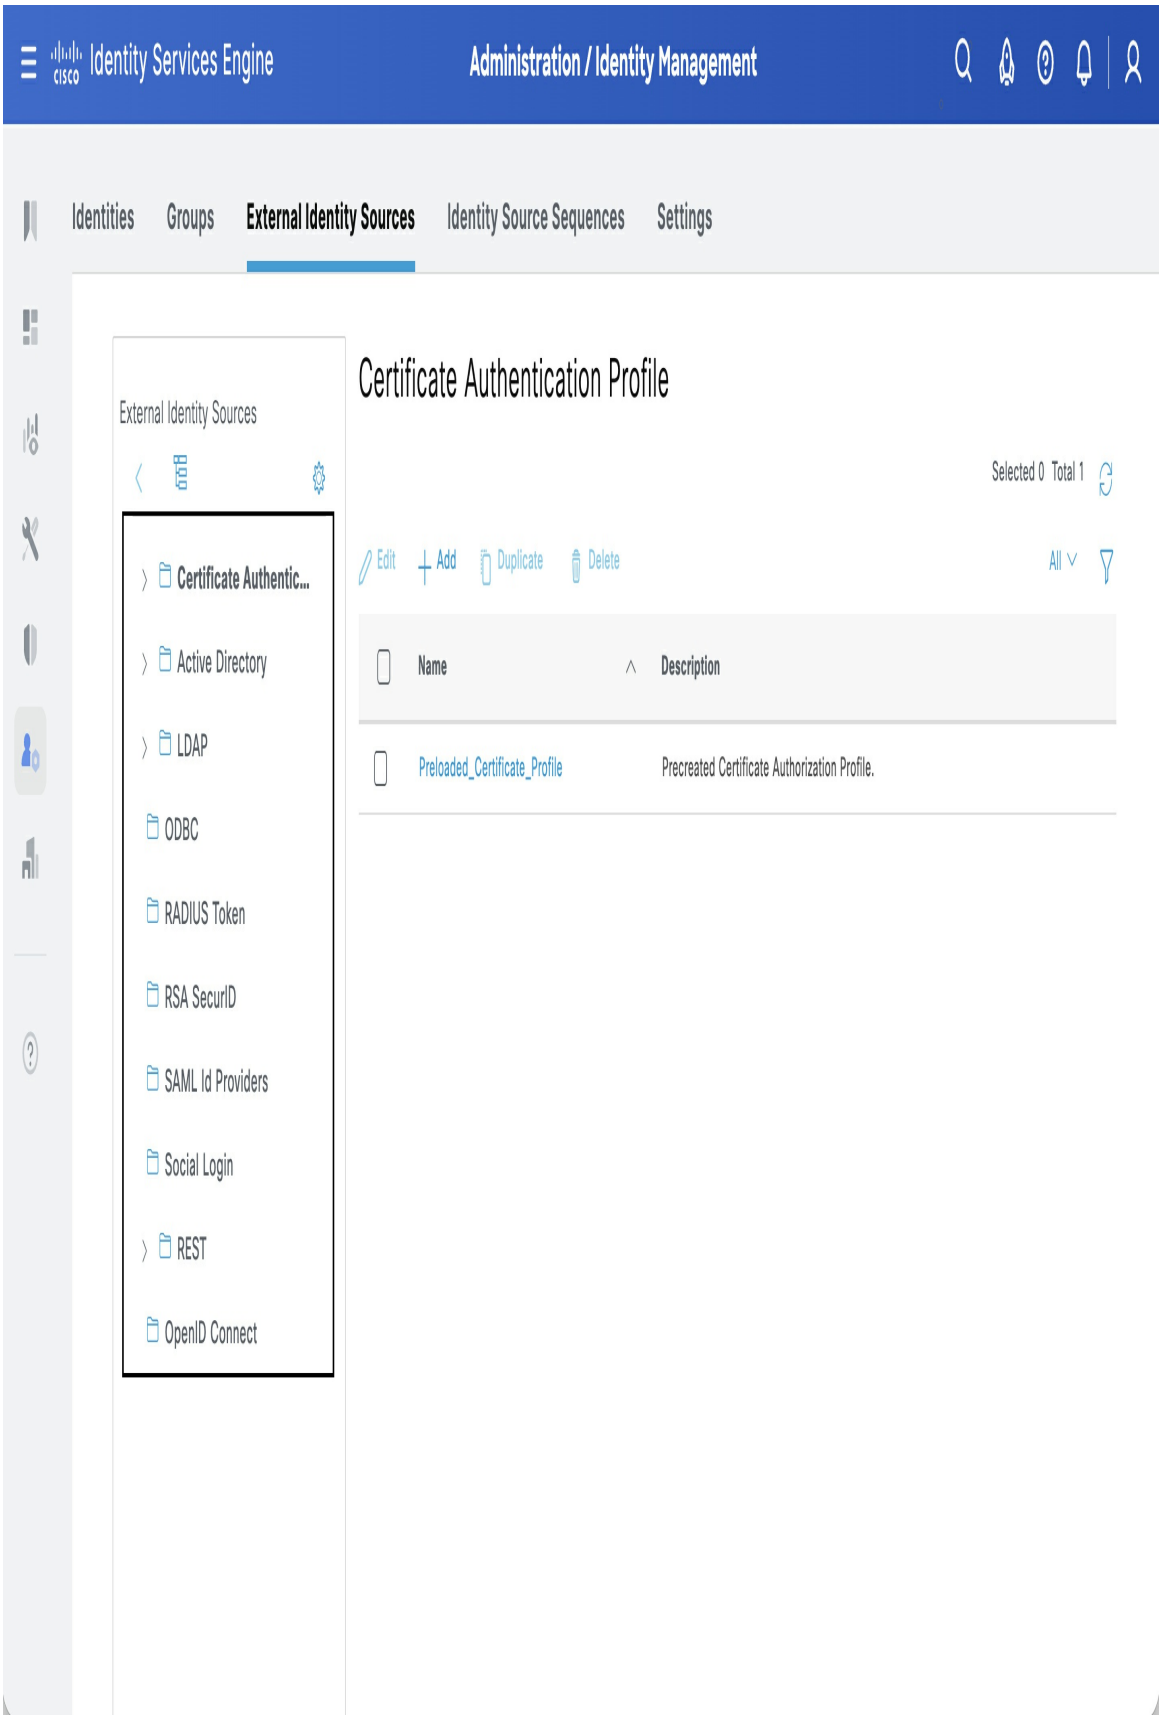
Always select the minimum required settings under Trusted

Technet24

For when you add a certificate to the Trusted Certificate Store.

External Identity Sources

Let's talk about each external identity source in more detail. To configure or view existing identity sources, navigate to **Administration > Identity Management > External Identity Sources**, as shown in Figure 5-17.

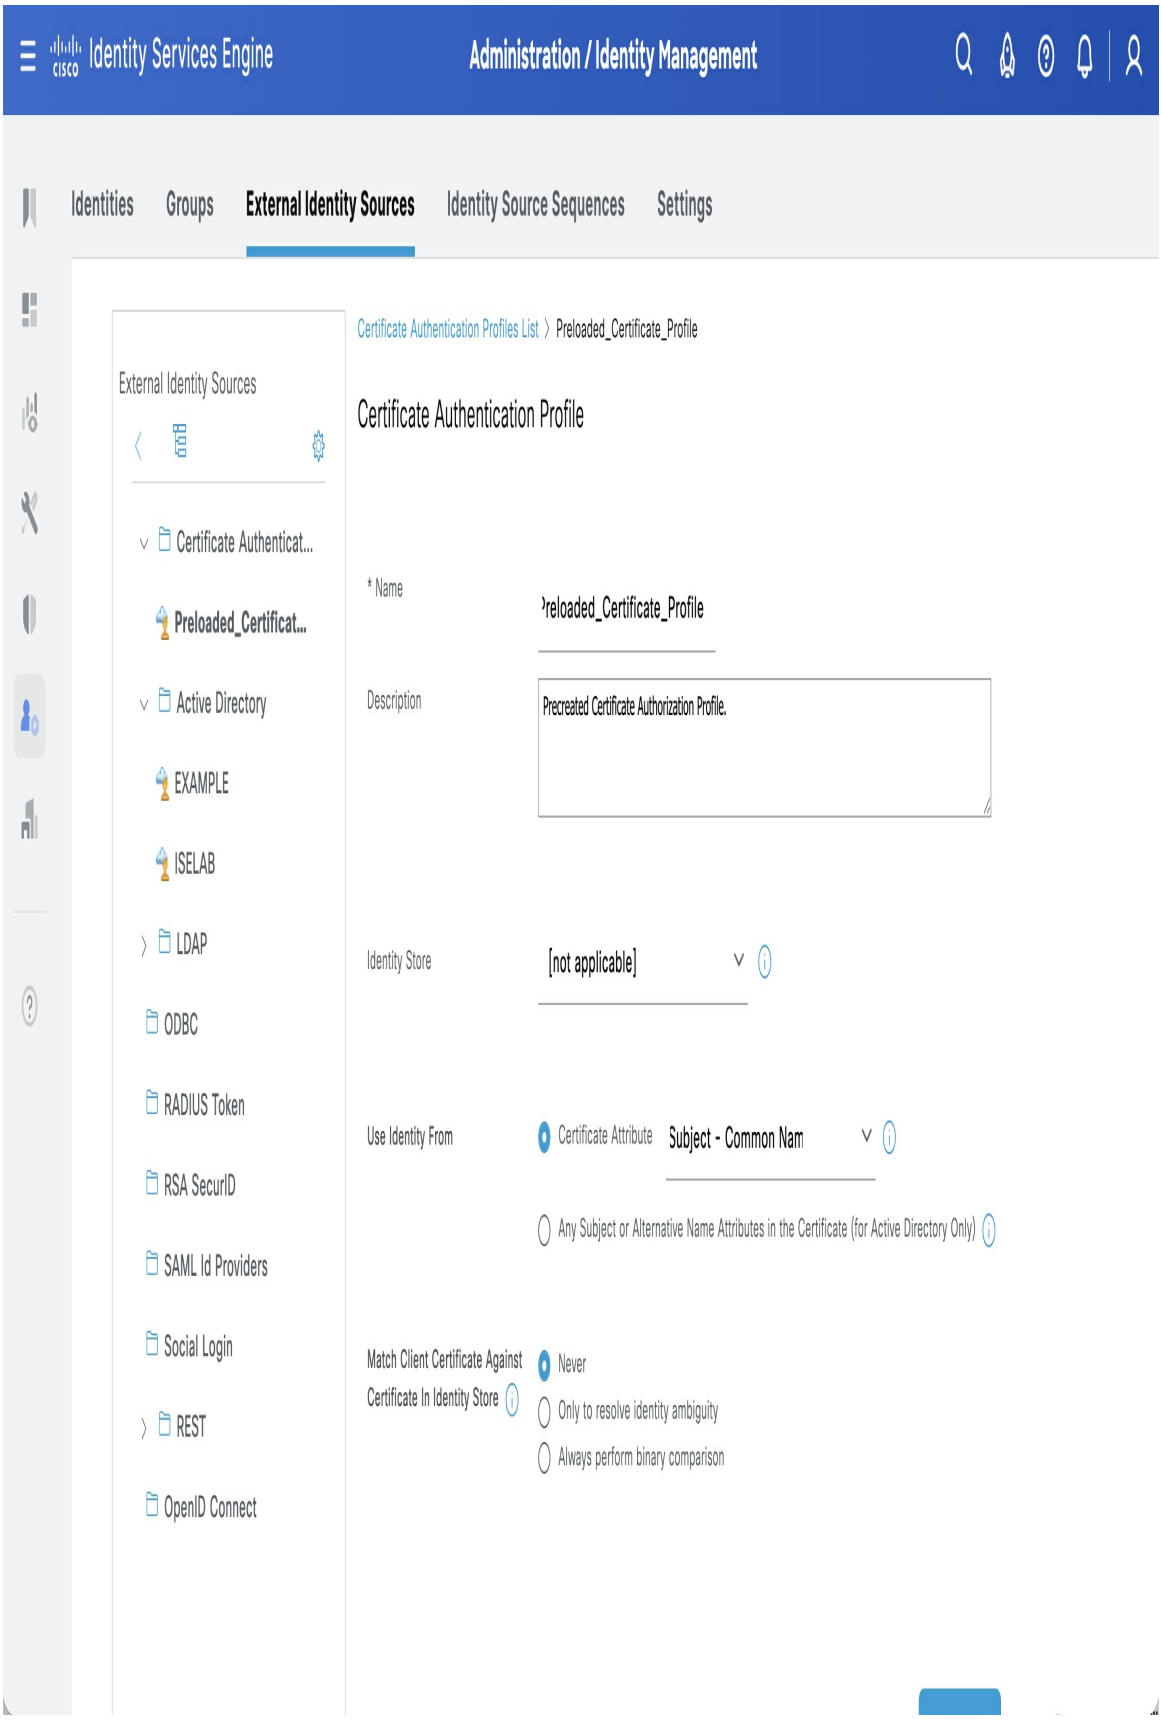

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authentic...
Active Directory
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Certificate Authentication Profile
Selected 0 Total 1
Edit
Add
Duplicate
Delete
All
Name
Description
Preloaded_Certificate_Profile
Precreated Certificate Authorization Profile.

Technet24

Figure 5-17 *Built-in External Identity Sources*

Certificate Authentication Profile

The Certificate Authentication Profile serves as one of the external identity sources, which may initially appear confusing. You do not store user or device information in this identity store. Instead, you use this store to perform certificate-based authentication. ISE includes a built-in Certificate Authentication Profile, as shown in Figure 5-18.

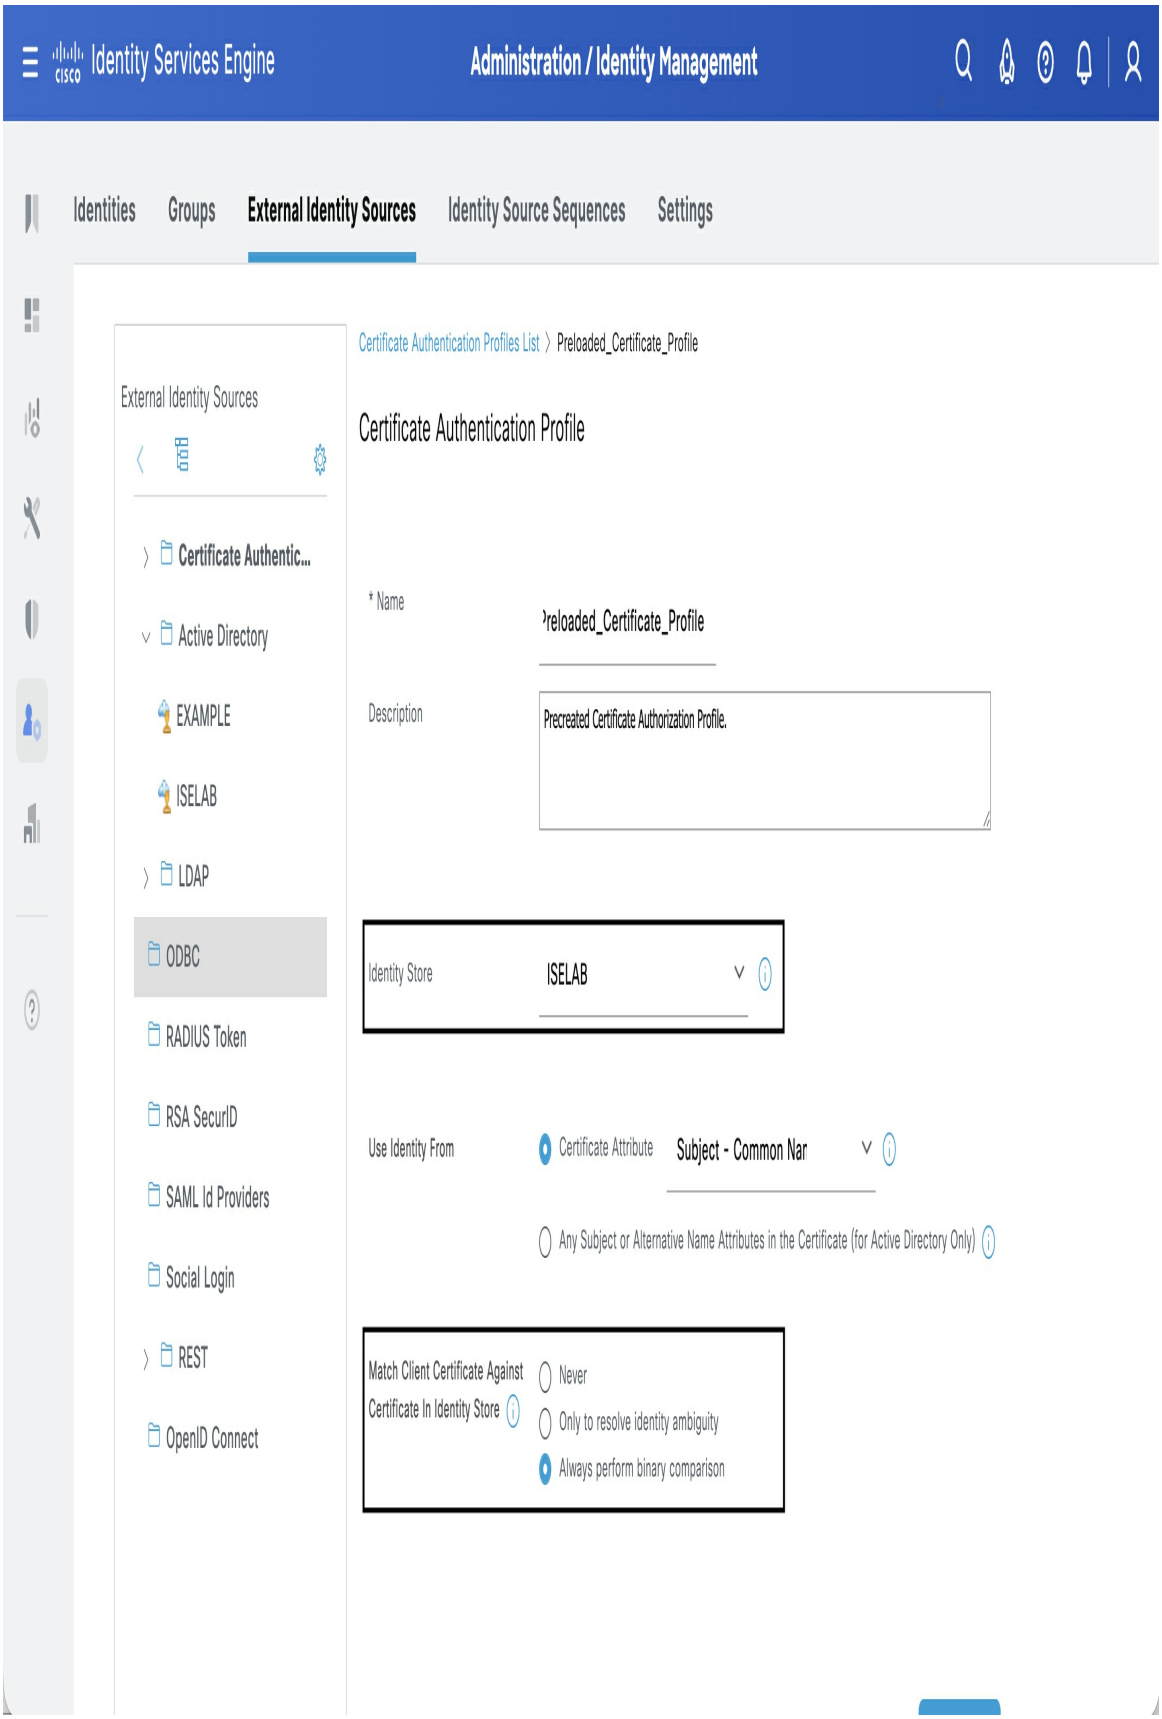

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Preloaded_Certificat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Certificate Authentication Profiles List
Preloaded_Certificate_Profile
Certificate Authentication Profile
* Name
Preloaded_Certificate_Profile
Description
Precreated Certificate Authorization Profile.
Identity Store
[not applicable]
Use Identity From
Certificate Attribute
Subject - Common Nam
Any Subject or Alternative Name Attributes in the Certificate (for Active Directory Only)
Match Client Certificate Against Certificate In Identity Store
Never
Only to resolve identity ambiguity
Always perform binary comparison

Technet24

Figure 5-18 *Built-in Certificate Authentication Profile*

Within the Certificate Authentication Profile, you can select which certificate field appears in the live logs and is used for lookups during authorization. Many users mistakenly assume that certificate-based authentication requires an external identity store. In practice, you do not need an external identity store to authenticate users with certificates. ISE can authenticate an endpoint by using EAP-TLS and verifying that the client certificate is valid and trusted. You can optionally enable binary certificate comparison if you want to compare the certificate presented by the client with a certificate stored in Active Directory or an LDAP server.

To perform binary certificate comparison, configure the Certificate Authentication Profile as shown in Figure 5-19.


Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authentic...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Certificate Authentication Profiles List
Preloaded_Certificate_Profile
Certificate Authentication Profile
* Name
Preloaded_Certificate_Profile
Description
Precreated Certificate Authorization Profile.
Identity Store
ISELAB
Use Identity From
Certificate Attribute
Subject - Common Nar
Any Subject or Alternative Name Attributes in the Certificate (for Active Directory Only)
Match Client Certificate Against Certificate In Identity Store
Never
Only to resolve identity ambiguity
Always perform binary comparison

Technet24

Figure 5-19 *ISE Binary Certificate Comparison Configuration*

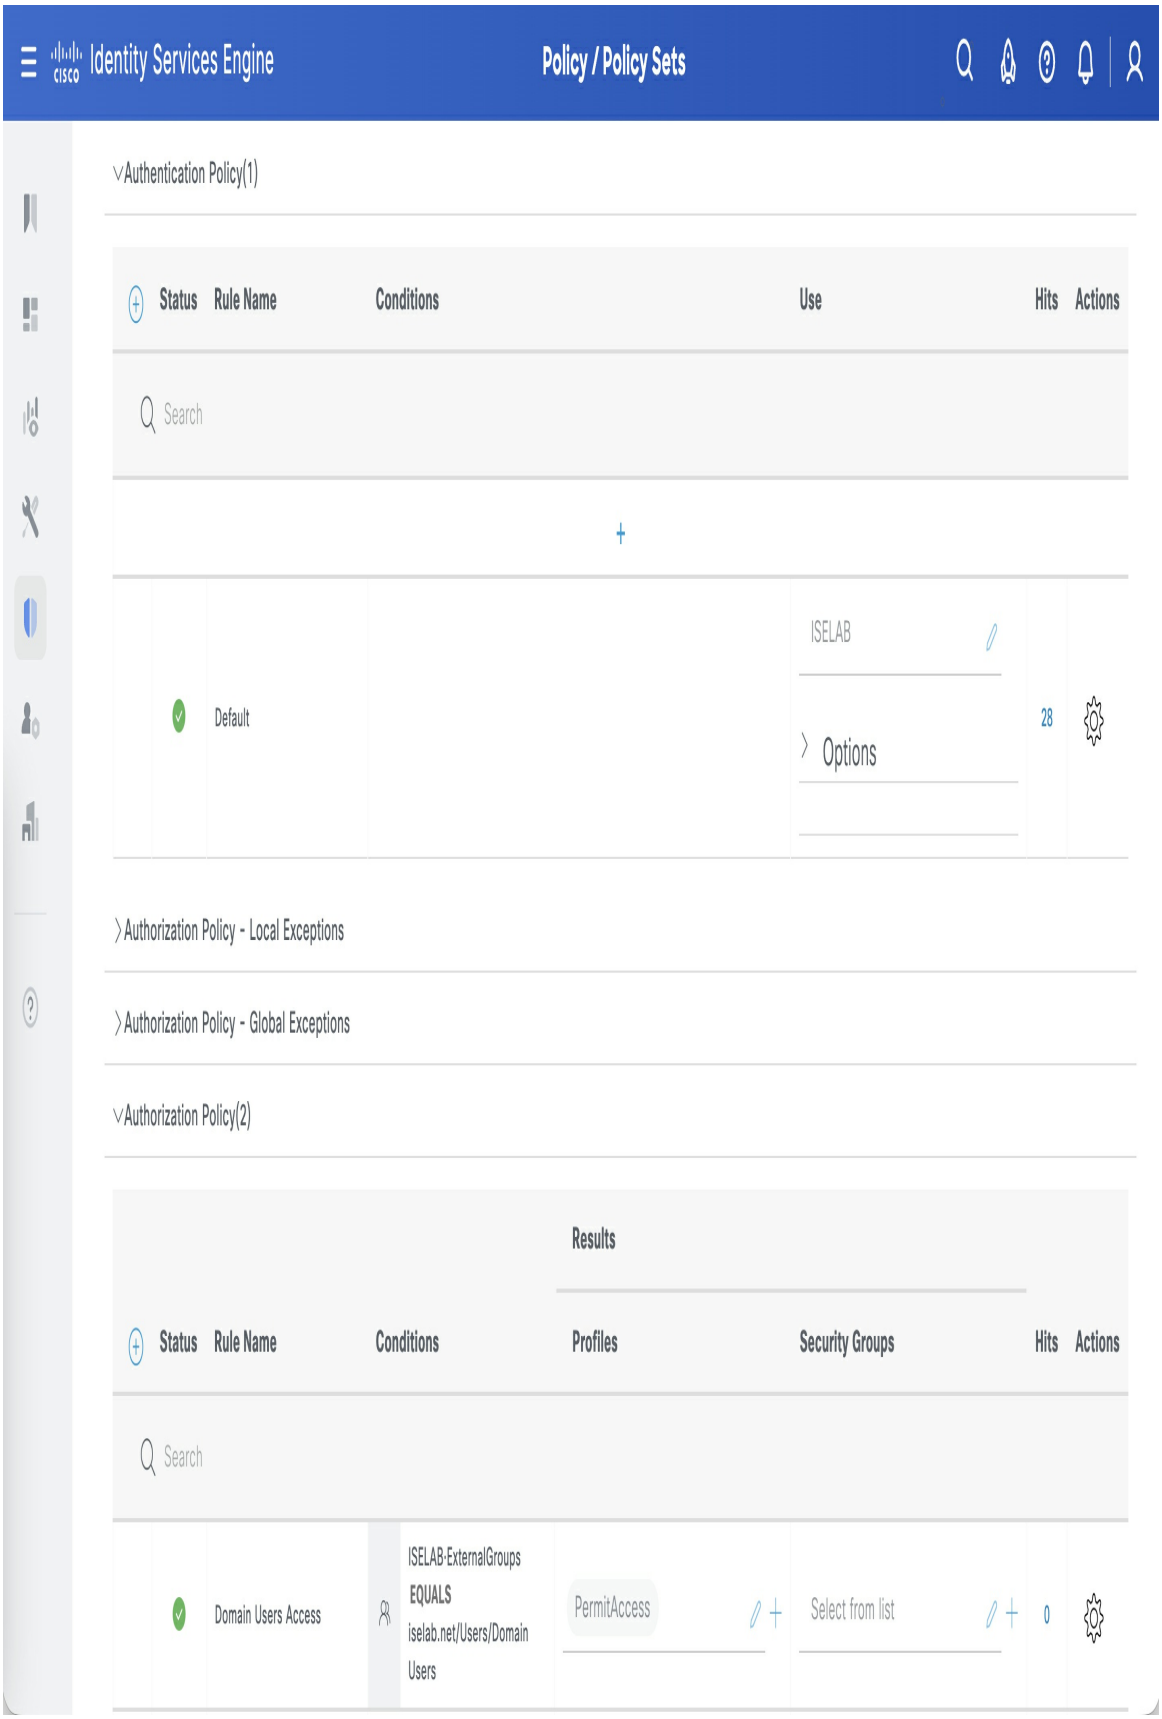
Note that Active Directory allows you to configure an alternative option to only do binary certificate comparison when you need to resolve identity ambiguity. With this option enabled, ISE performs fewer byte-by-byte comparisons.

Active Directory

Microsoft Active Directory integration is one of the most commonly implemented integrations with Cisco ISE.

You can use Active Directory groups and attributes in authorization policies to differentiate network access. Cisco ISE retrieves user and group information after successful authentication and can also retrieve this information when authentication occurs elsewhere.

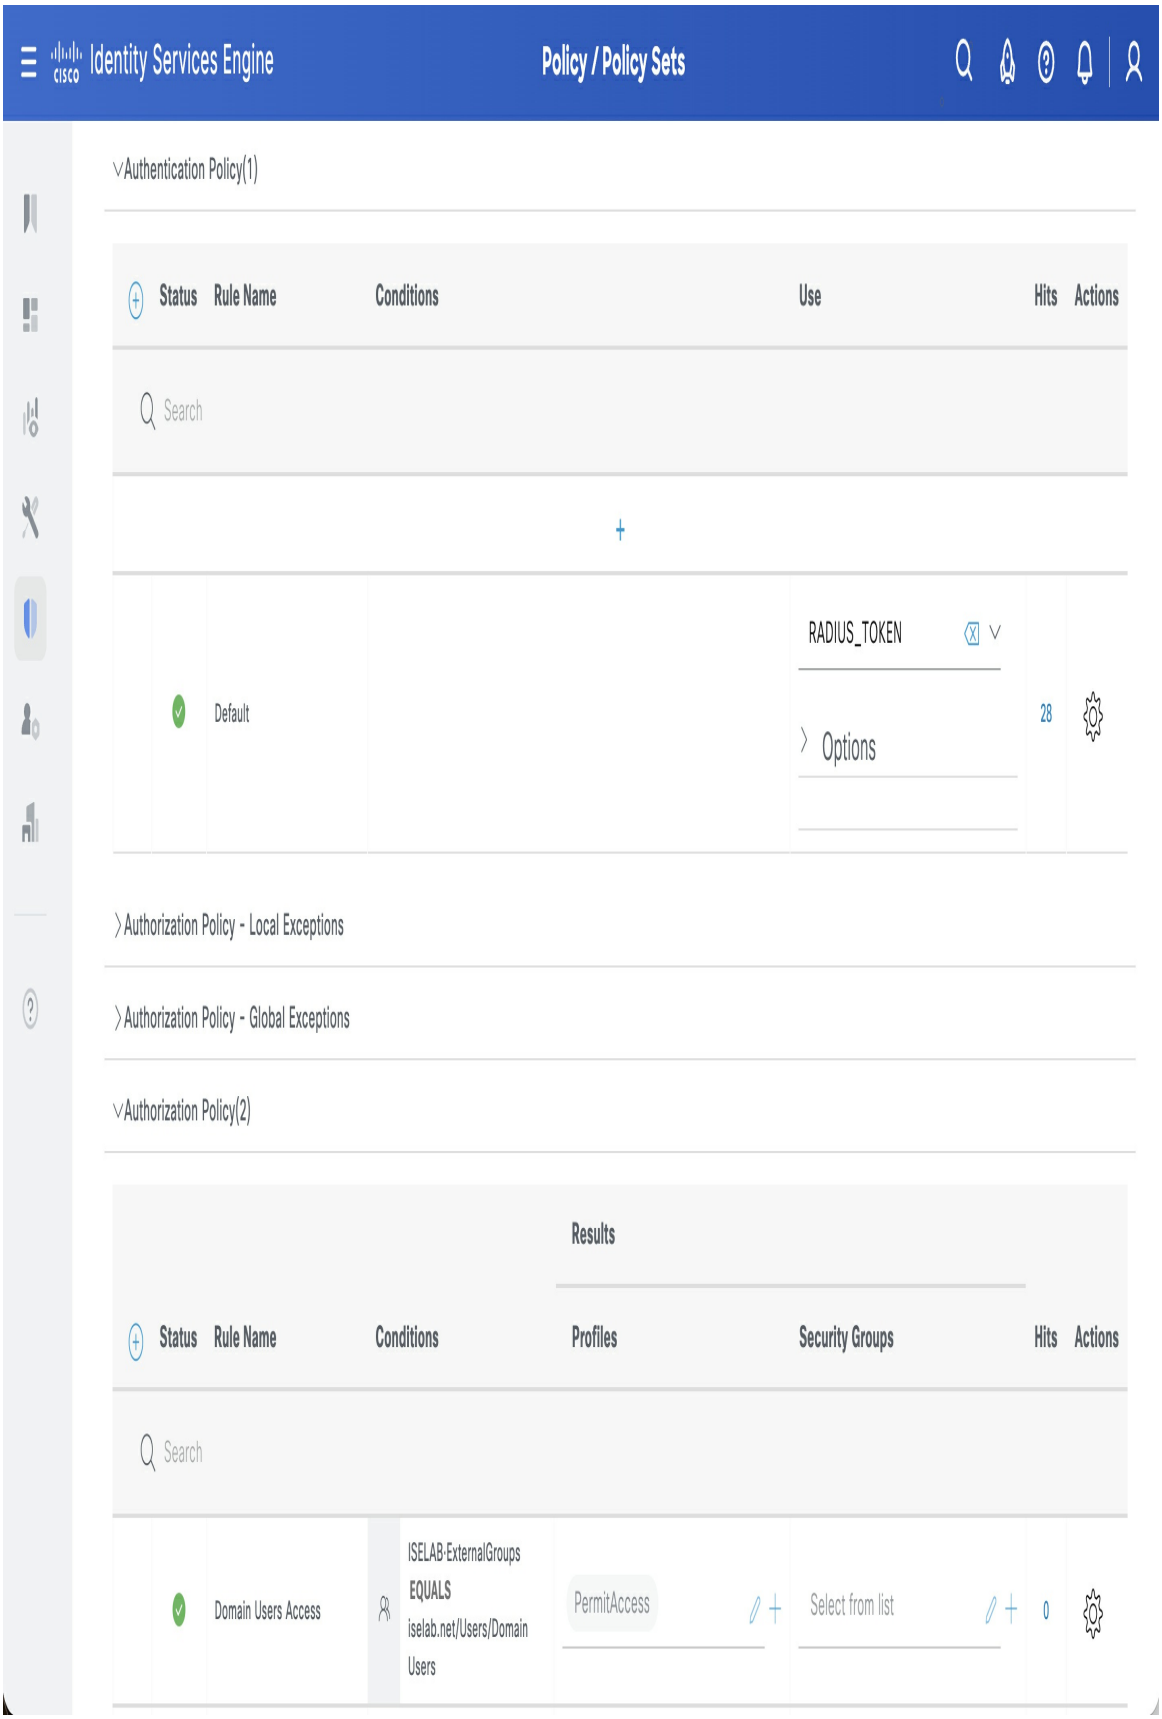
The following figures provide two examples of such scenarios. In Figure 5-20, authentication points to Active Directory, and the authorization condition also relies on Active Directory. This is a commonly used scenario, but it is not the only one that is supported.


Identity Services Engine
Policy / Policy Sets
Authentication Policy(1)
Status
Rule Name
Conditions
Use
Hits
Actions
Search
Default
ISELAB
Options
28
Authorization Policy - Local Exceptions
Authorization Policy - Global Exceptions
Authorization Policy(2)
Results
Status
Rule Name
Conditions
Profiles
Security Groups
Hits
Actions
Search
Domain Users Access
ISELAB-ExternalGroups
EQUALS
iselab.net/Users/Domain Users
PermitAccess
Select from list
0

Technet24

Figure 5-20 *Same Identity Source Is Used for Authentication and Authorization*

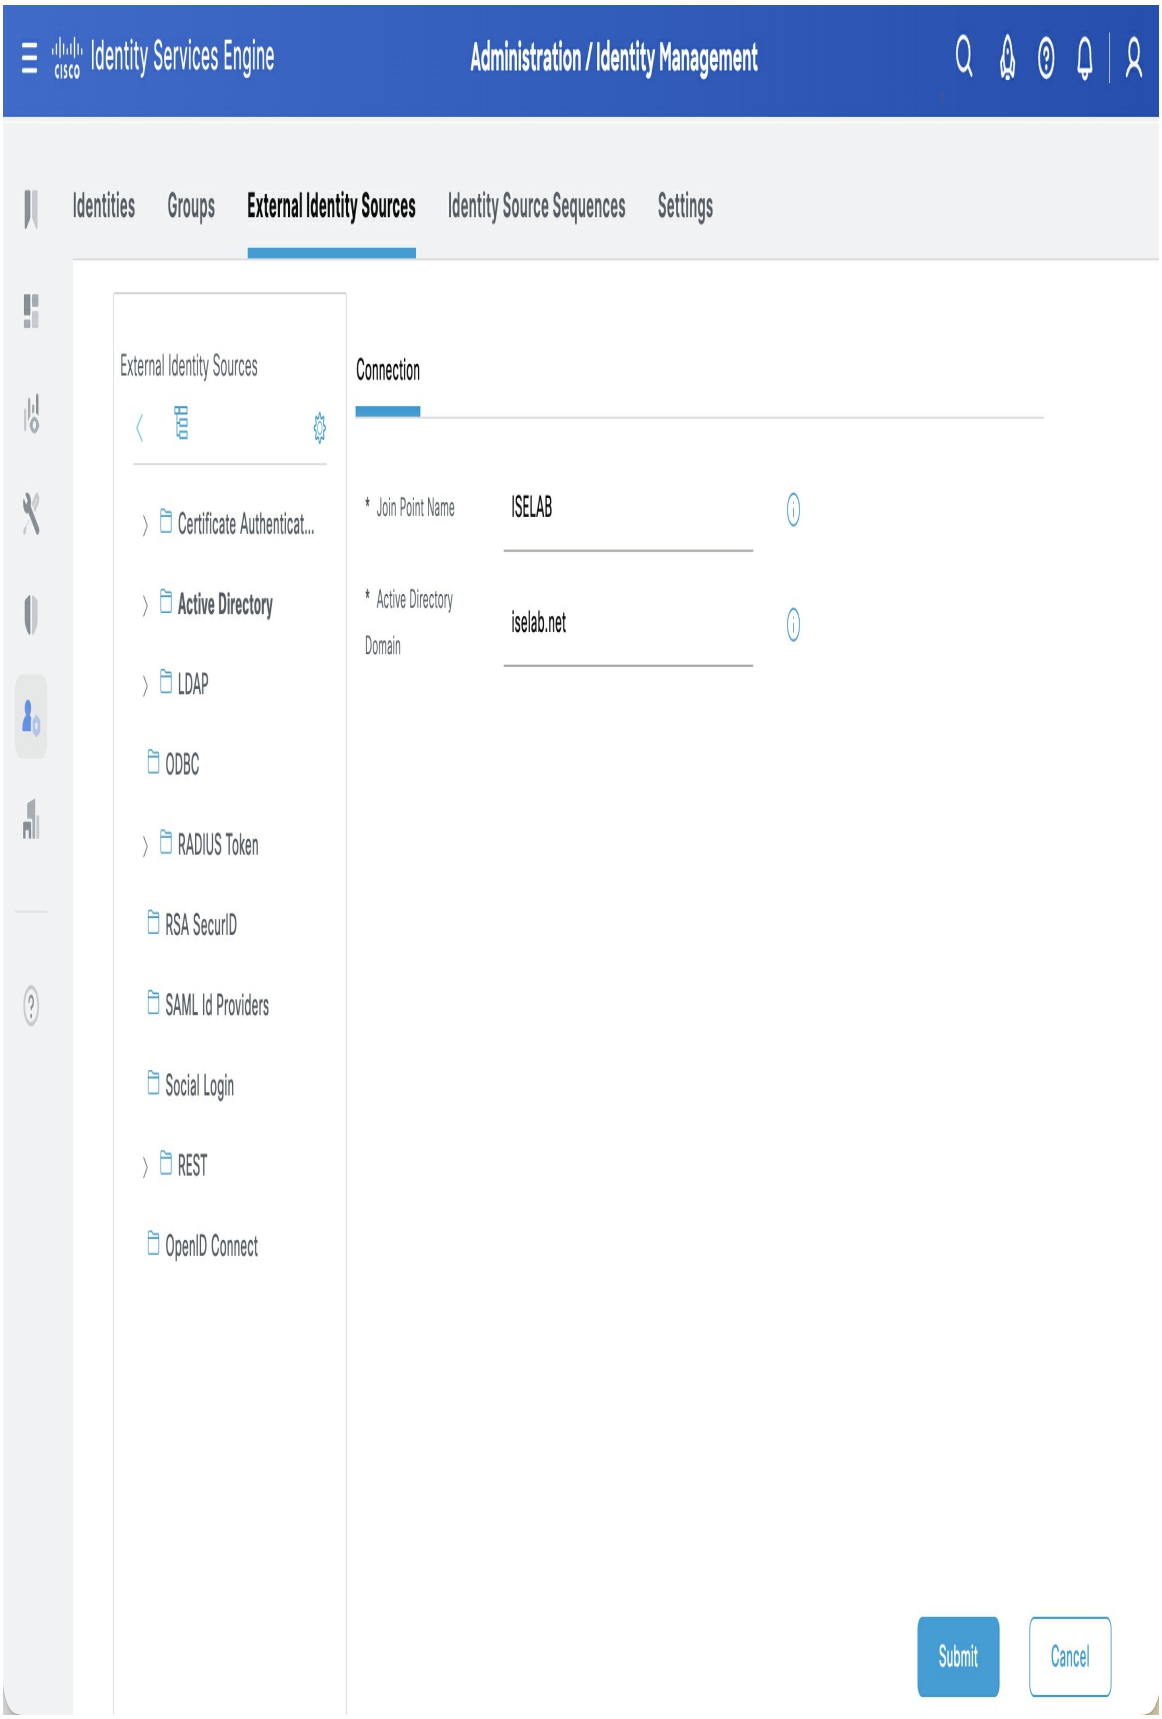
You can also configure ISE to authenticate users against completely different identity sources and still use authorization based on Active Directory groups and attributes. Figure 5-21 shows a RADIUS token server configured for authentication, while authorization relies on Active Directory, even though Active Directory does not appear in the authentication settings.


Identity Services Engine
Policy / Policy Sets
Authentication Policy(1)
Status
Rule Name
Conditions
Use
Hits
Actions
Search
Default
RADIUS_TOKEN
Options
28
Authorization Policy - Local Exceptions
Authorization Policy - Global Exceptions
Authorization Policy(2)
Results
Status
Rule Name
Conditions
Profiles
Security Groups
Hits
Actions
Search
Domain Users Access
ISELAB-ExternalGroups
EQUALS
iselab.net/Users/Domain Users
PermitAccess
Select from list
0

Technet24

Figure 5-21 *Different Identity Sources Are Used for Authentication and Authorization*

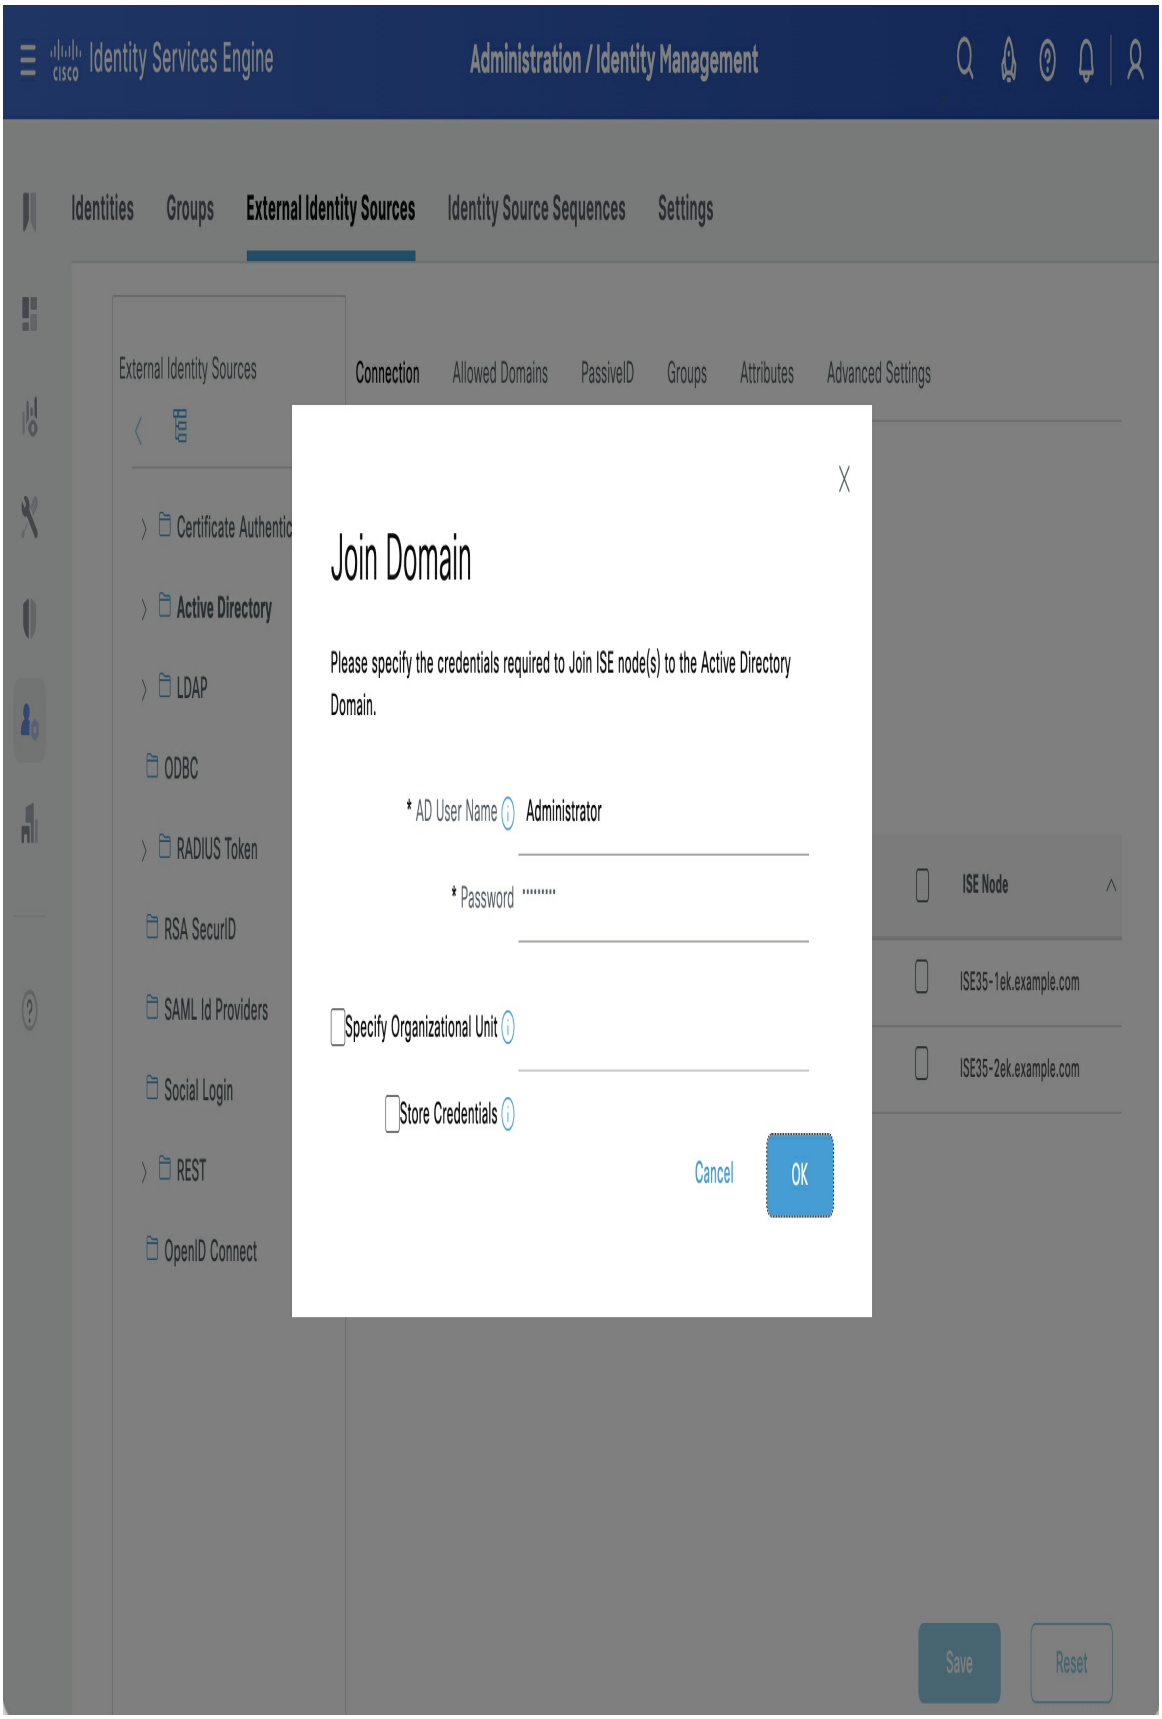
To integrate ISE with Active Directory, navigate to **Administration > Identity Management > External Identity Sources > Active Directory** and click **Add**. You need to enter the join point name, which defines how the Active Directory join point appears in ISE, and specify the Active Directory domain that you want to join, as shown in Figure 5-22.


Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Connection
* Join Point Name
ISELAB
* Active Directory Domain
iselab.net
Submit
Cancel

Technet24

Figure 5-22 *Active Directory Connection*

After you click **Submit**, ISE prompts you to join all nodes in the deployment to the domain. When you accept this prompt, ISE displays a menu where you enter the Active Directory credentials, as shown in Figure 5-23.


Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Connection
Allowed Domains
PassiveID
Groups
Attributes
Advanced Settings
Certificate Authentic
Active Directory
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
ISE Node
ISE35-1ek.example.com
ISE35-2ek.example.com
Save
Reset
Join Domain
Please specify the credentials required to Join ISE node(s) to the Active Directory Domain.
* AD User Name
Administrator
* Password
Specify Organizational Unit
Store Credentials
Cancel
OK

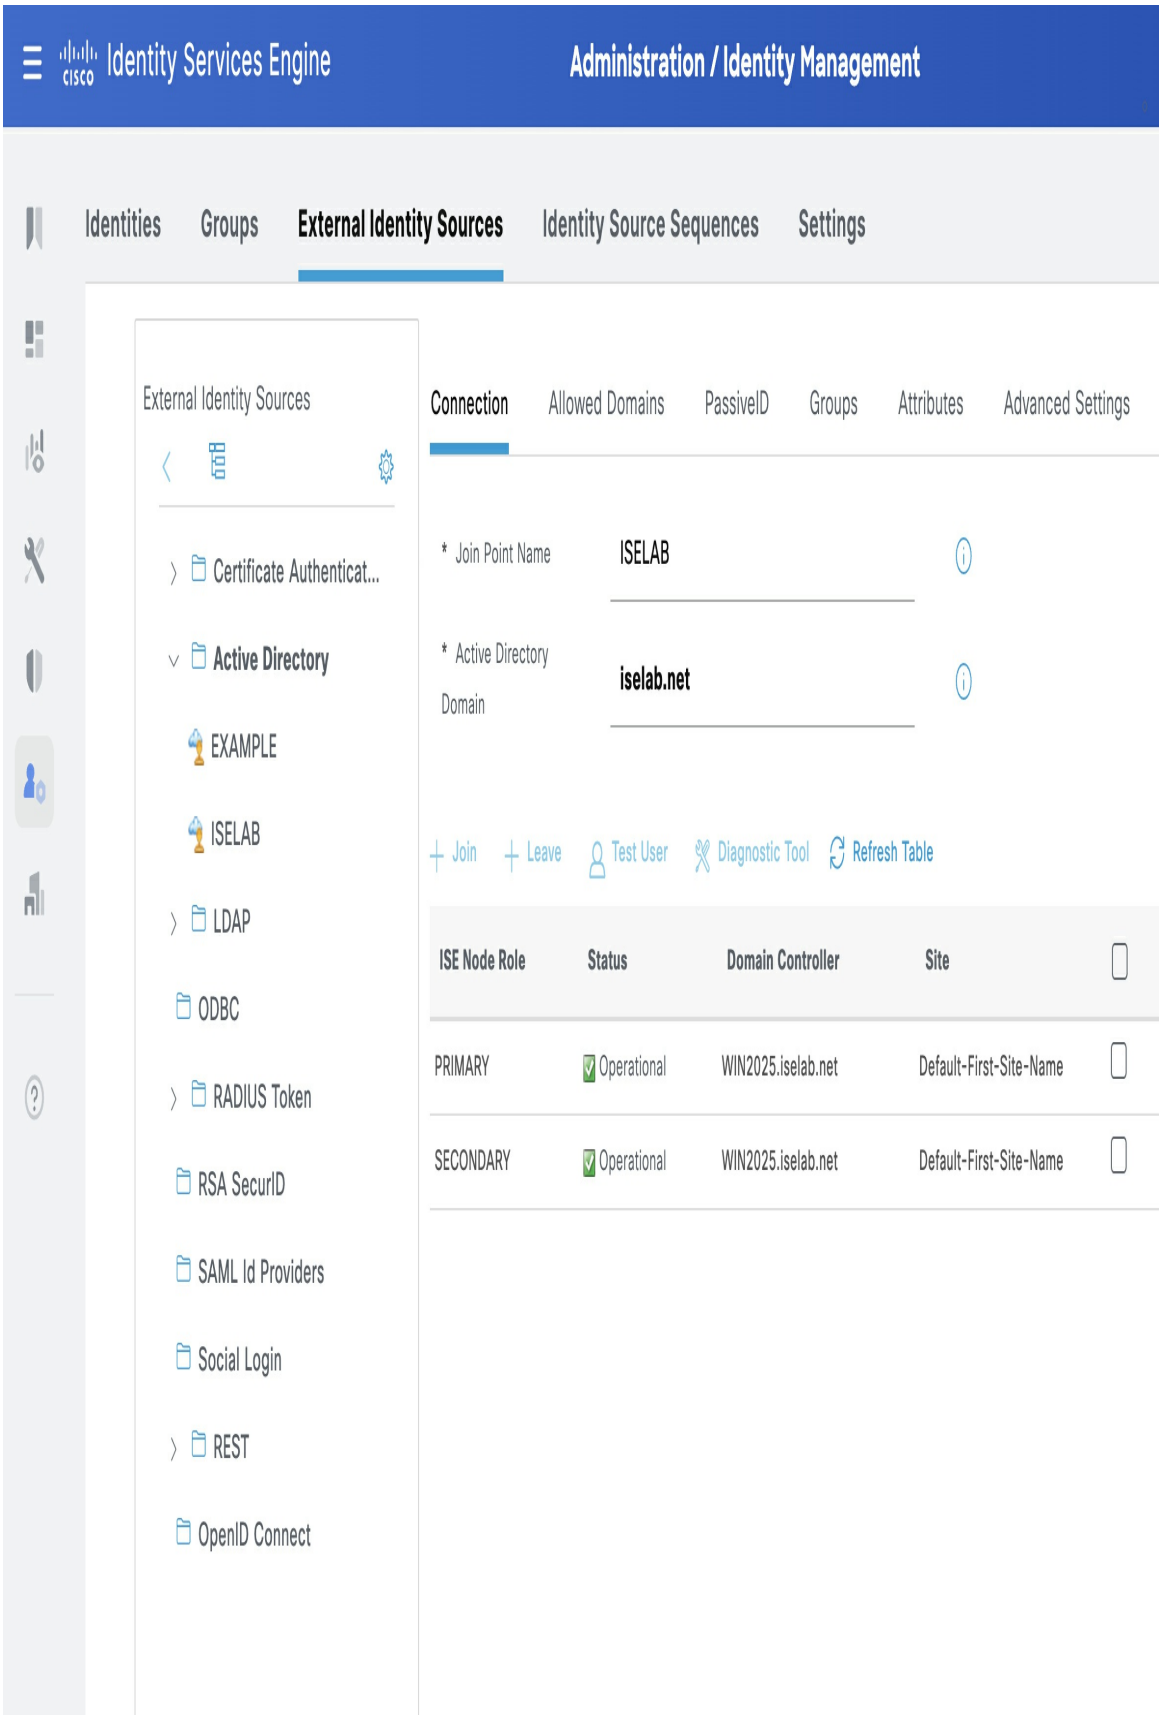
Technet24

Figure 5-23 *Active Directory Credentials Menu*

You must meet several prerequisites before you can successfully integrate ISE with Active Directory:

- Time must be synchronized between ISE and Active Directory. You should use an NTP server to achieve synchronization.
- The Active Directory account used for integration must have privileges to search, create, and remove machine accounts.

Figure 5-24 shows both ISE nodes successfully integrated with Active Directory. Note that each ISE node in the deployment is associated with a specific domain controller. This design allows ISE to connect to the domain controller that is closest to the PSN for efficient communication. ISE is site aware and will communicate with a domain controller in its own site unless all domain controllers in its own site are unavailable.


Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Connection
Allowed Domains
PassiveID
Groups
Attributes
Advanced Settings
* Join Point Name
ISELAB
* Active Directory Domain
iselab.net
Join
Leave
Test User
Diagnostic Tool
Refresh Table
ISE Node Role
Status
Domain Controller
Site
PRIMARY
Operational
WIN2025.iselab.net
Default-First-Site-Name
SECONDARY
Operational
WIN2025.iselab.net
Default-First-Site-Name

Technet24

Figure 5-24 *Successful Active Directory Integration.*

DC discovery on ISE is a complex process that involves multiple steps:

1. ISE sends a DNS query for all DC, GC, and KDC records. If the DNS response does not include any of these records, the integration fails with a DNS-related error.
2. ISE uses a CLDAP ping to discover all DCs and GCs by sending CLDAP requests to domain controllers based on their priority in the SRV records. ISE uses the first responding domain controller and then connects to that DC.

   One factor used to calculate domain controller priority is the response time to CLDAP pings. A faster response results in a higher priority.

Note

CLDAP is the mechanism that ISE uses to establish and maintain connectivity with domain controllers. It measures the response time until the first domain controller responds. The process fails if no response is received from a domain controller and generates a warning if the response time exceeds 2.5 seconds. CLDAP pings all domain controllers in the site, or all domain controllers in the domain if no site is defined. The CLDAP response includes both the domain controller site and the client site to which the ISE node is assigned.

3. ISE receives a ticket granting ticket (TGT) with join user credentials.
4. ISE creates the ISE machine account name by using MSRPC, including SAM and SPN.

5. ISE searches Active Directory by SPN to determine whether the ISE machine account already exists. If the account does not exist, ISE creates a new one.
6. ISE sets the ISE machine account password and verifies that the machine account is accessible.
7. ISE sets ISE machine account attributes, such as SPN, dnsHostname, and related attributes.
8. ISE obtains a TGT by using the ISE machine credentials through KRB5 and discovers all trusted domains.

When the domain controller to which ISE is connected becomes unavailable, ISE automatically fails over to the next domain controller in the priority list. Failover occurs when the domain controller is offline or otherwise unreachable.

Trusted domains discovered by ISE allow you to authenticate users who are not present in the domain that ISE joins. Figure 5-25 shows ISE integrated with the domain example.com, yet during integration, ISE also discovers other domains, such as child.example.com or ekorneyc.com.

Technet24

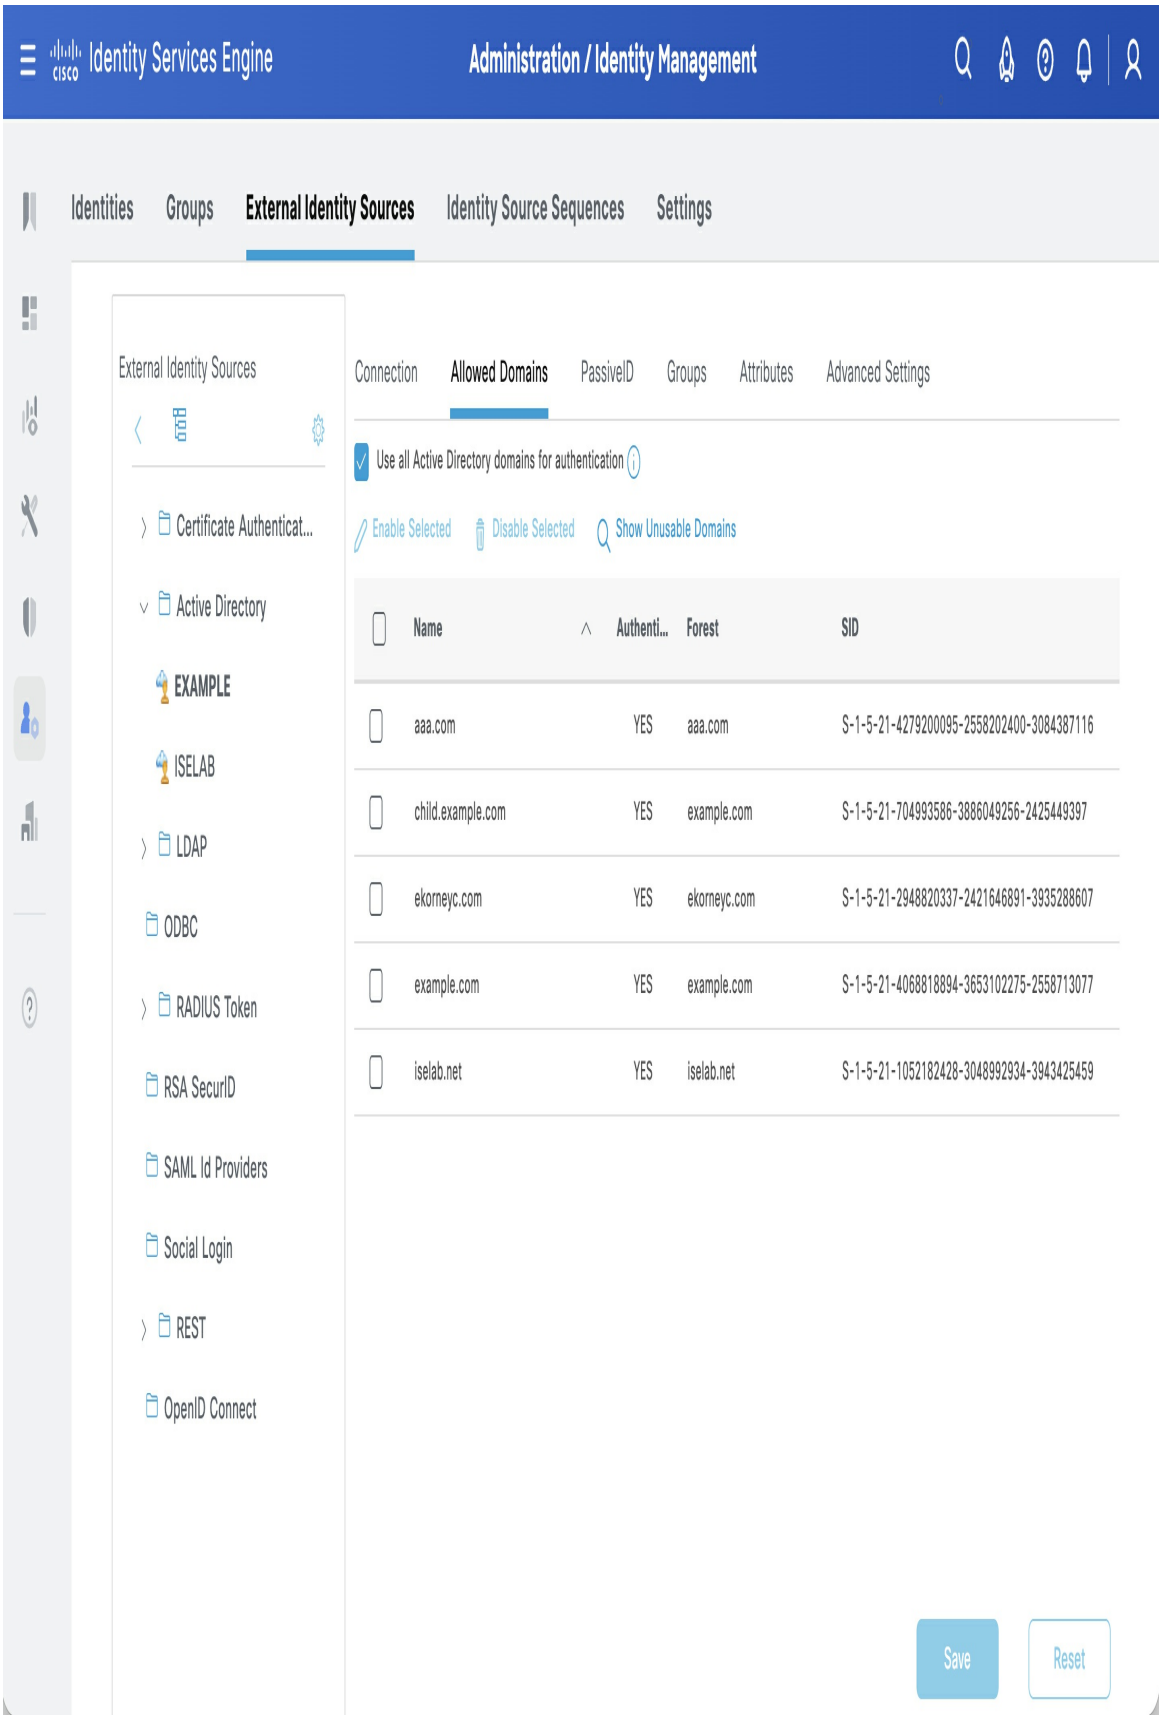

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Connection
Allowed Domains
PassiveID
Groups
Attributes
Advanced Settings
Use all Active Directory domains for authentication
Enable Selected
Disable Selected
Show Unusable Domains
Name
Authenti...
Forest
SID
aaa.com
YES
aaa.com
S-1-5-21-4279200095-2558202400-3084387116
child.example.com
YES
example.com
S-1-5-21-704993586-3886049256-2425449397
ekorneyc.com
YES
ekorneyc.com
S-1-5-21-2948820337-2421646891-3935288607
example.com
YES
example.com
S-1-5-21-4068818894-3653102275-2558713077
iselab.net
YES
iselab.net
S-1-5-21-1052182428-3048992934-3943425459
Save
Reset

Figure 5-25 *Active Directory Allowed Domains*

If you do not want ISE to authenticate users from a discovered domain, you can select the required domain and click **Disable Selected**.

To retrieve groups from Active Directory so that you can use them in policies, navigate to **Groups** and click **Add Groups**. Select the required groups and click **OK**. After you save the configuration, the groups become available for policy configuration and appear in the Active Directory settings, as shown in Figure 5-26.

Technet24

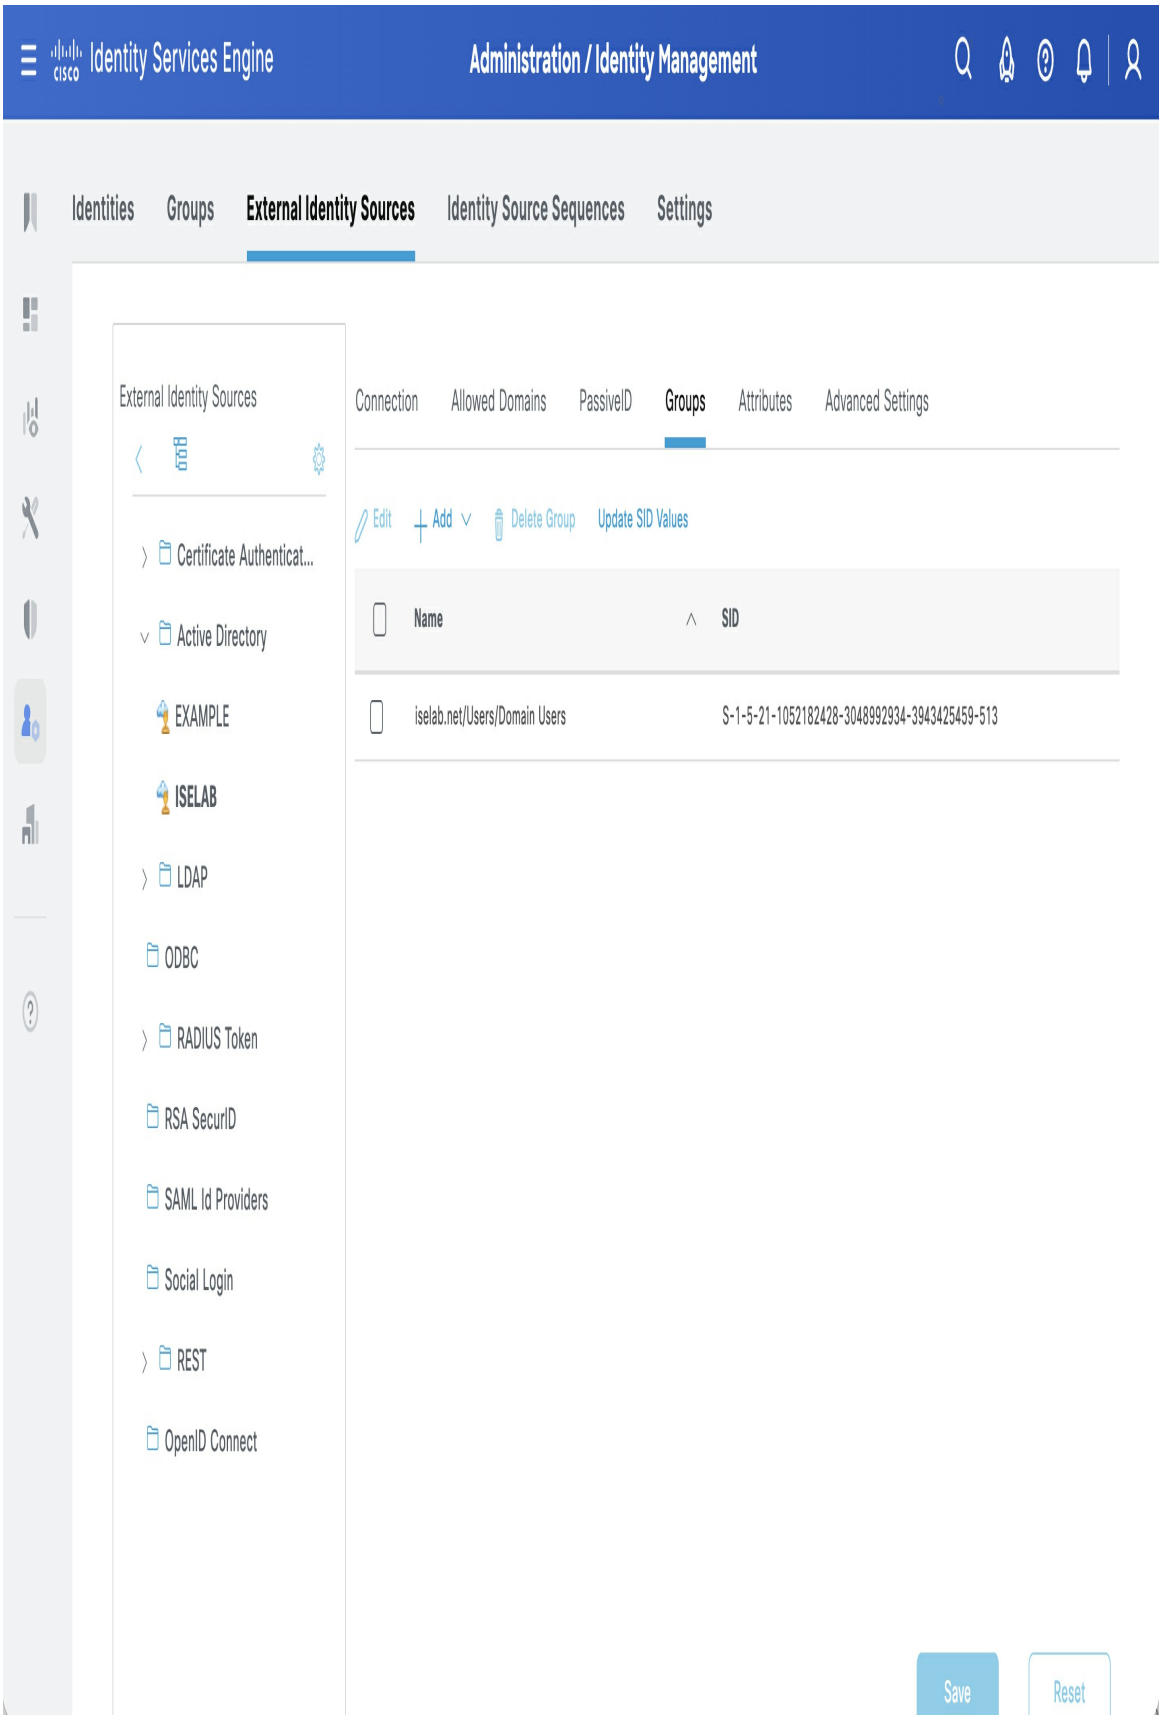

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Connection
Allowed Domains
PassiveID
Groups
Attributes
Advanced Settings
Edit
Add
Delete Group
Update SID Values
Name
SID
iselab.net/Users/Domain Users
S-1-5-21-1052182428-3048992934-3943425459-513
Save
Reset

Figure 5-26 *Active Directory Groups*

The same concept applies to attributes settings. You can select attributes from Active Directory by retrieving them from a sample user and saving them so that they appear on the Attributes tab. You can then use these attributes in policies.

If required, ISE also allows you to perform Identity Rewrite. You can configure this option from the advanced settings, as shown in Figure 5-27.

Technet24

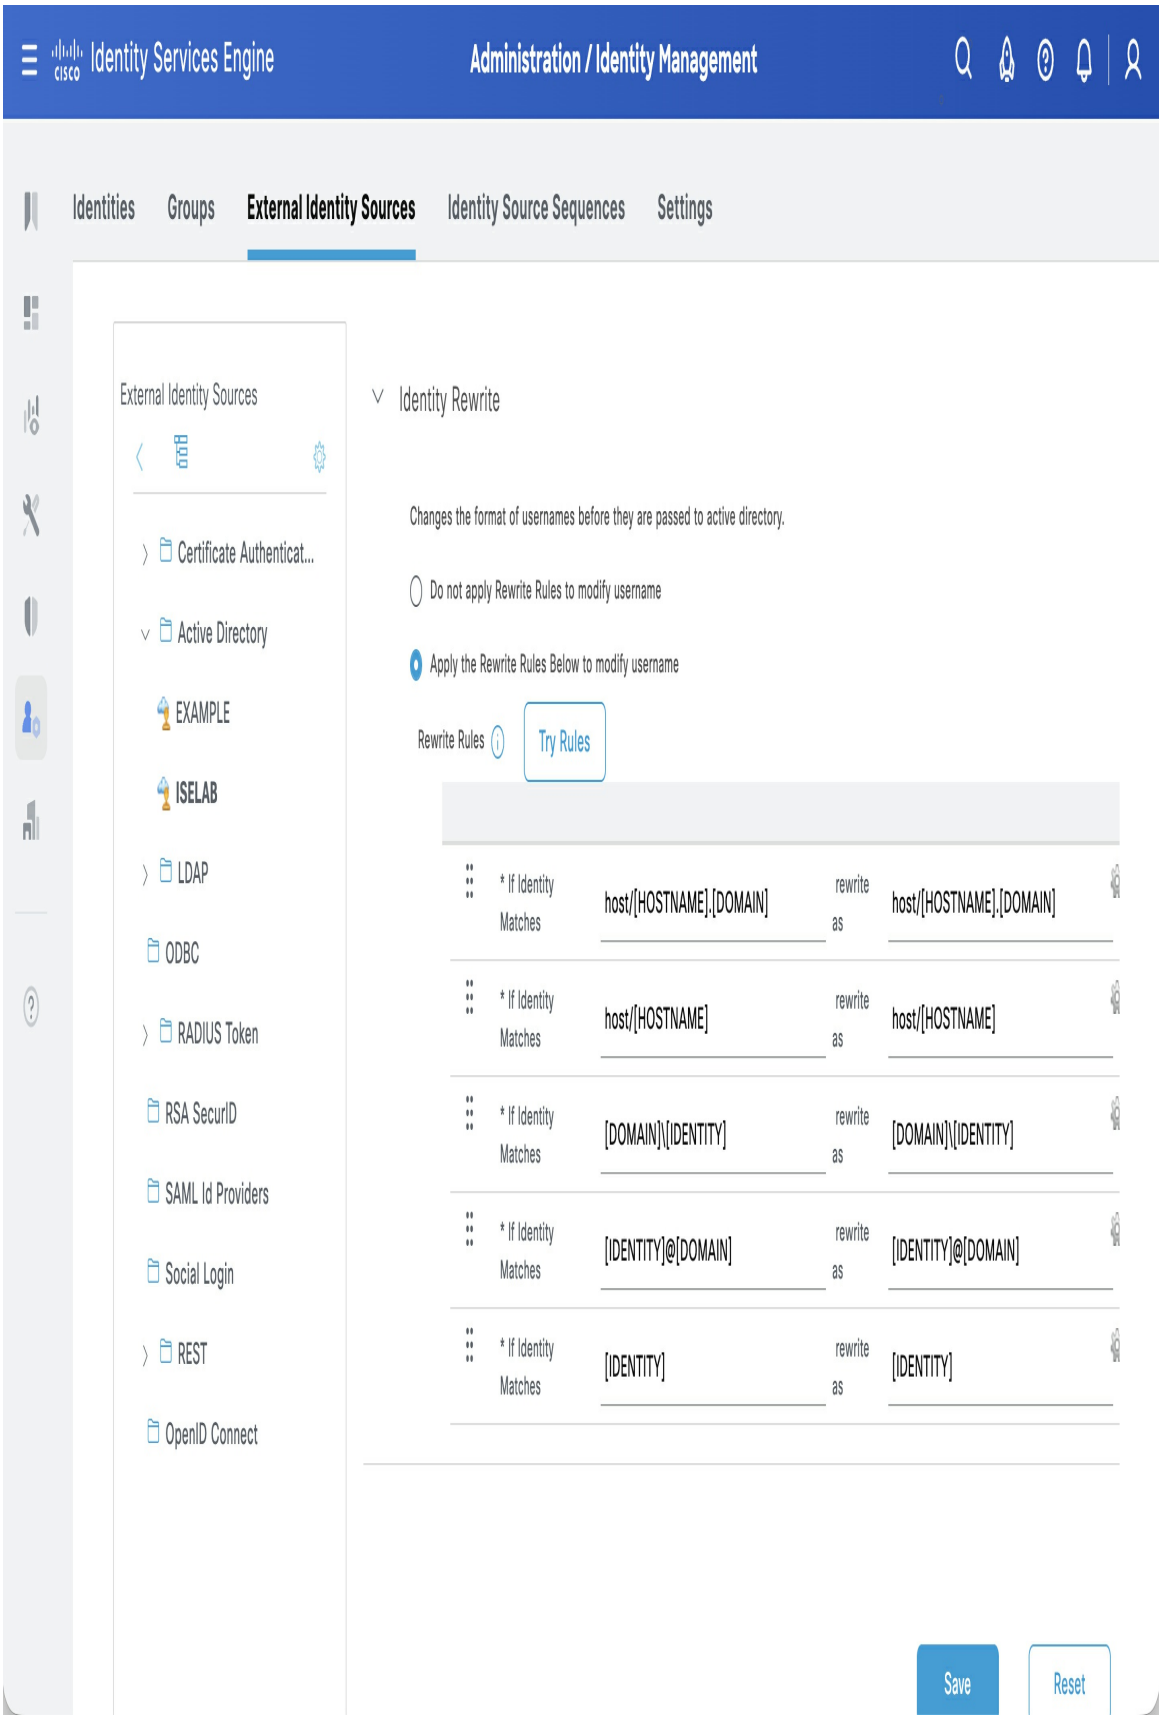

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
Identity Rewrite
Changes the format of usernames before they are passed to active directory.
Do not apply Rewrite Rules to modify username
Apply the Rewrite Rules Below to modify username
Rewrite Rules
Try Rules
* If Identity Matches
host/[HOSTNAME].[DOMAIN]
rewrite as
host/[HOSTNAME].[DOMAIN]
* If Identity Matches
host/[HOSTNAME]
rewrite as
host/[HOSTNAME]
* If Identity Matches
[DOMAIN]\[IDENTITY]
rewrite as
[DOMAIN]\[IDENTITY]
* If Identity Matches
[IDENTITY]@[DOMAIN]
rewrite as
[IDENTITY]@[DOMAIN]
* If Identity Matches
[IDENTITY]
rewrite as
[IDENTITY]
Save
Reset

Figure 5-27 *Active Directory Identity Rewrite*

Active Directory Identity Rewrite allows you to change the format of the username received in the request to another preferred format.

LDAP

In addition to Active Directory integration, ISE can integrate with Lightweight Directory Access Protocol (LDAP) for username/password and certificate-based or even MAB authentication. You can integrate ISE with multiple LDAP server instances. Additionally, within a single LDAP instance configuration, you can define primary and secondary LDAP servers. Failover occurs when the primary server stops responding to LDAP requests — for example, when the LDAP service is down or the server is unreachable.

> ### Note
>
> When you configure groups and attributes for LDAP, ISE always retrieves the configuration from the primary server. If the primary server is not reachable and you need to change the configuration, you must first restore connectivity to the primary server.

As with Active Directory, you can build authorization policies based on LDAP groups or LDAP attributes.

To configure LDAP integration, navigate to **Administration > Identity Management > External Identity Sources > LDAP** and click **Add**. You need to enter a name and select a schema, as shown in Figure 5-28.

Technet24

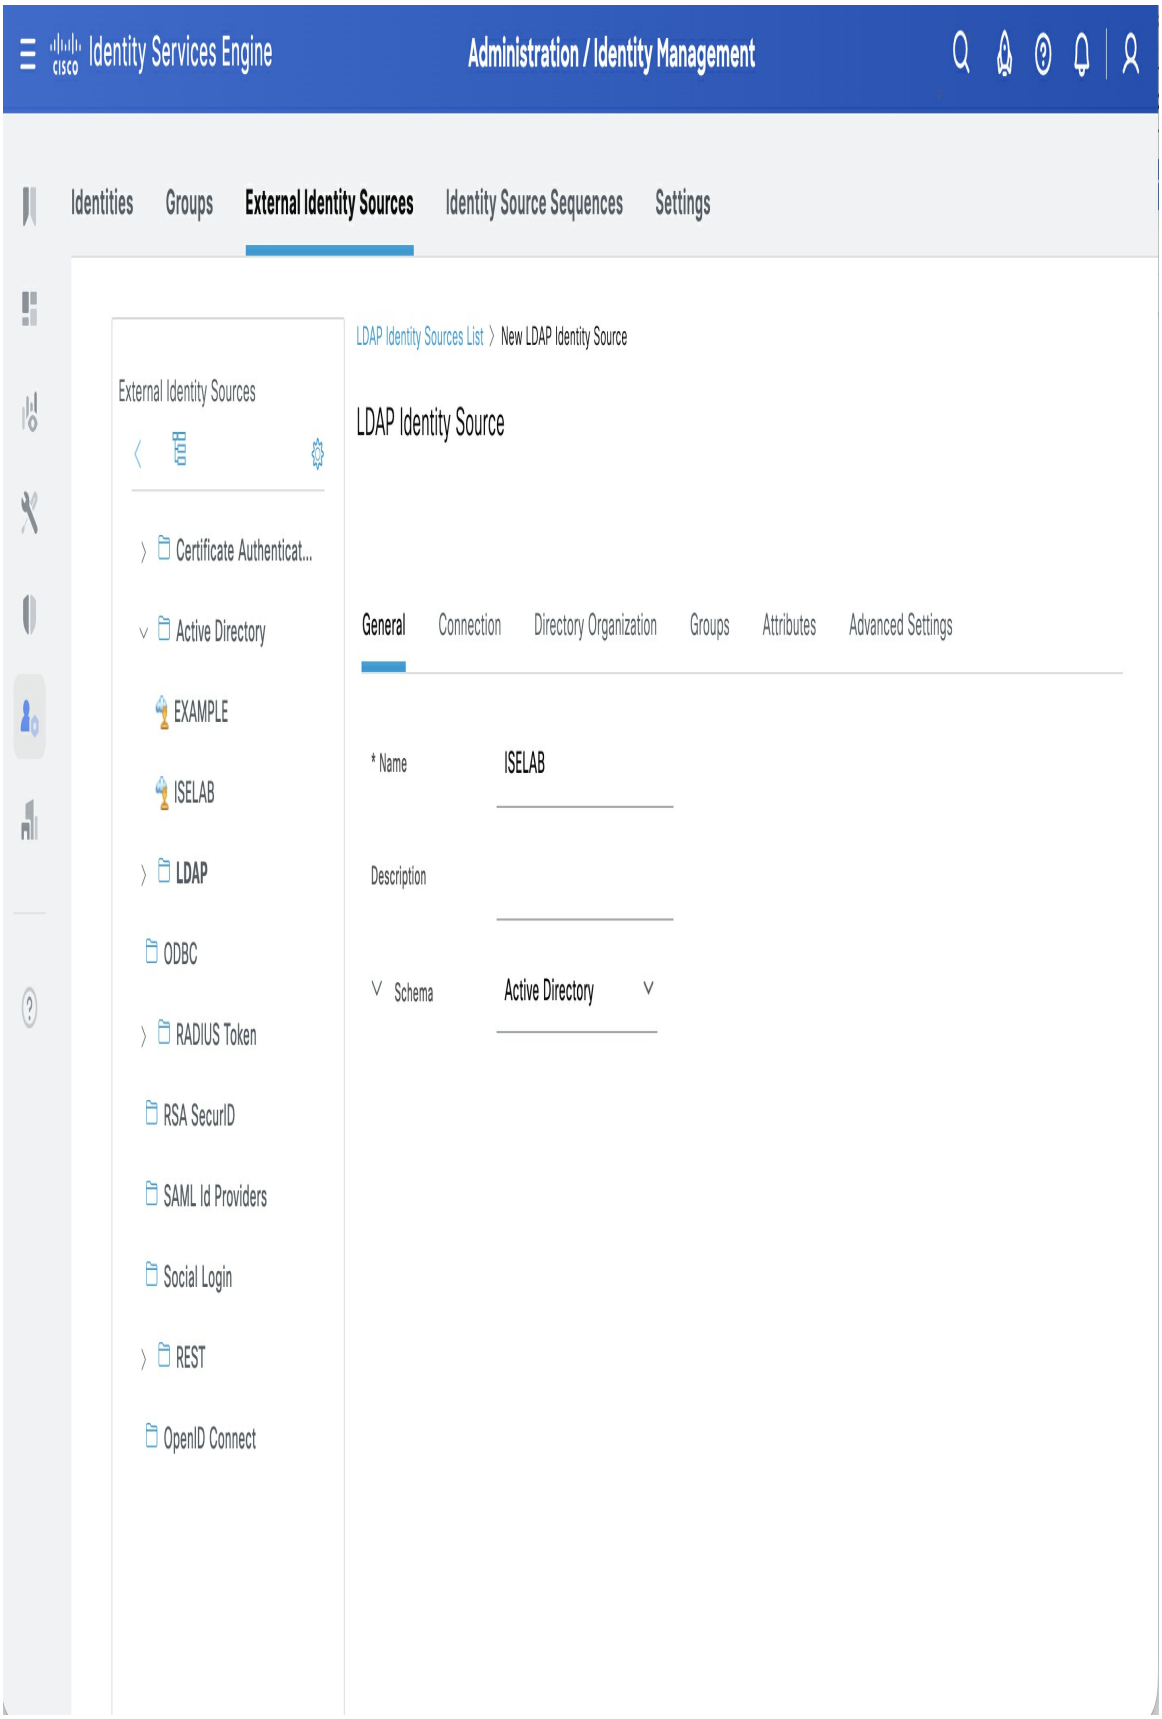

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
LDAP Identity Sources List
New LDAP Identity Source
LDAP Identity Source
General
Connection
Directory Organization
Groups
Attributes
Advanced Settings
* Name
ISELAB
Description
Schema
Active Directory

Figure 5-28 *LDAP Identity Source*

ISE includes several built-in schemas:

- Active Directory
- Sun Directory Server
- Novell eDirectory

You can also define a custom schema.

Next, under the **Connection** settings, you need to define the hostname or IP address of your primary and secondary LDAP servers, as shown in Figure 5-29.

Technet24

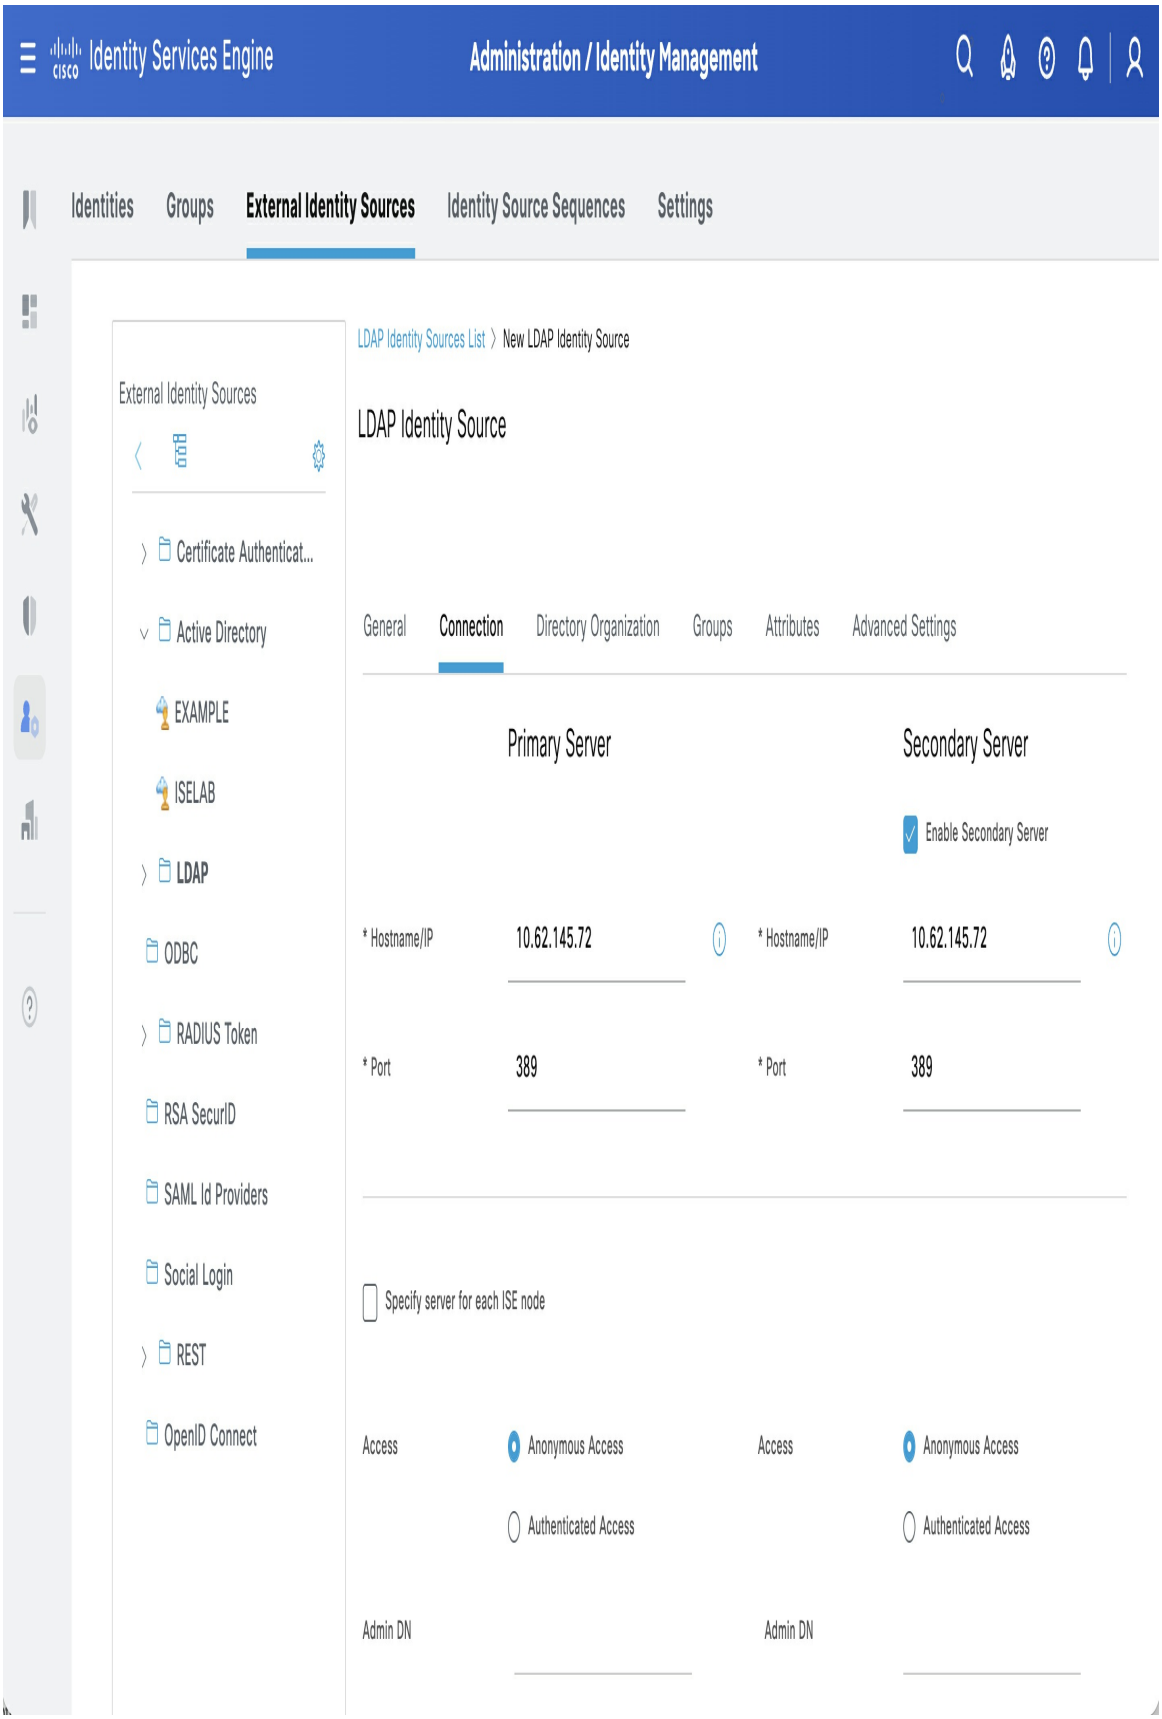

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
LDAP Identity Sources List
New LDAP Identity Source
LDAP Identity Source
General
Connection
Directory Organization
Groups
Attributes
Advanced Settings
Primary Server
Secondary Server
Enable Secondary Server
* Hostname/IP
10.62.145.72
* Hostname/IP
10.62.145.72
* Port
389
* Port
389
Specify server for each ISE node
Access
Anonymous Access
Authenticated Access
Access
Anonymous Access
Authenticated Access
Admin DN
Admin DN

Figure 5-29 *LDAP Identity Source Connection*

If you choose to configure **Authenticated Access**, you need to specify the **Admin DN** and the password for the Admin account. This configuration ensures that searches are performed by using administrative credentials, as shown in Figure 5-30.

Technet24

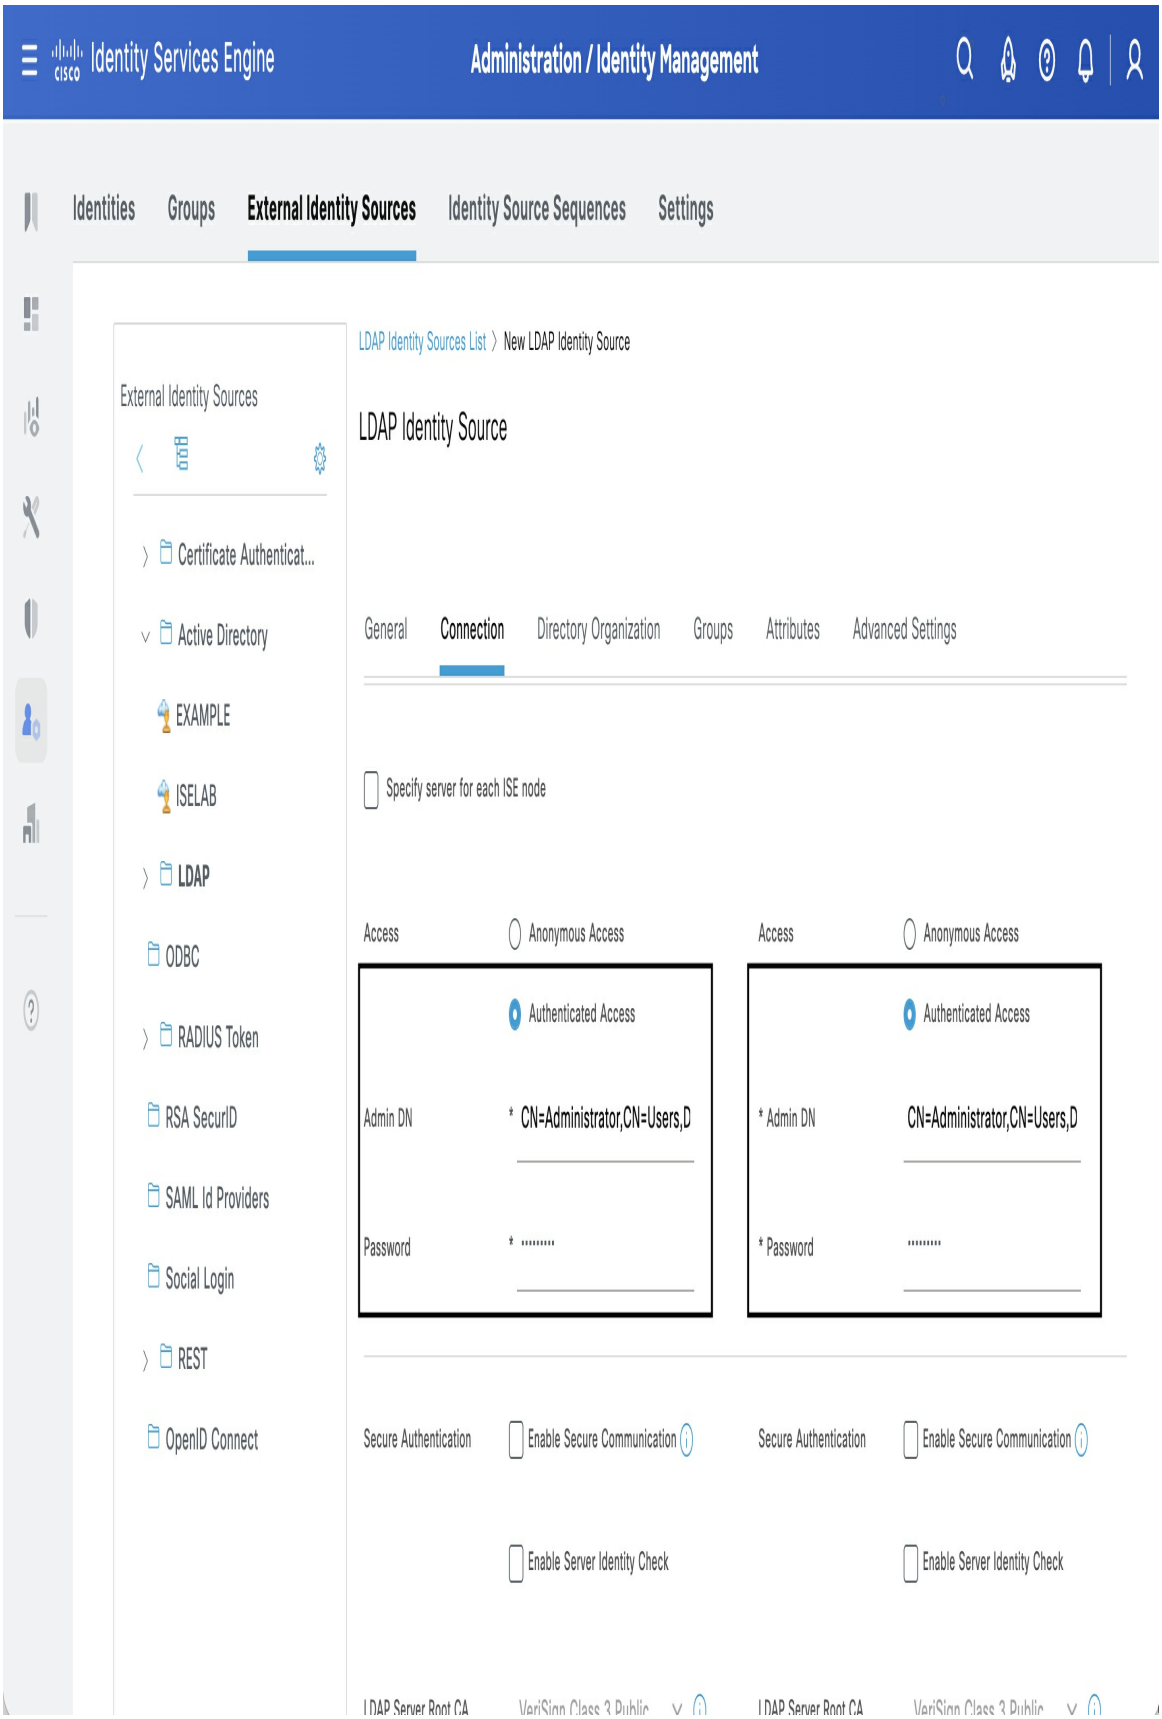

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
LDAP Identity Sources List
New LDAP Identity Source
LDAP Identity Source
General
Connection
Directory Organization
Groups
Attributes
Advanced Settings
Specify server for each ISE node
Access
Anonymous Access
Authenticated Access
Admin DN
* CN=Administrator,CN=Users,D
Password
Secure Authentication
Enable Secure Communication
Enable Server Identity Check
LDAP Server Root CA
VeriSign Class 3 Public

Figure 5-30 *LDAP Identity Source Authenticated Access Configuration*

Finally, you need to define the **Directory Organization**, as shown in Figure 5-31. In particular, you instruct ISE where to perform subject searches in the LDAP server and where to perform group searches.

Technet24

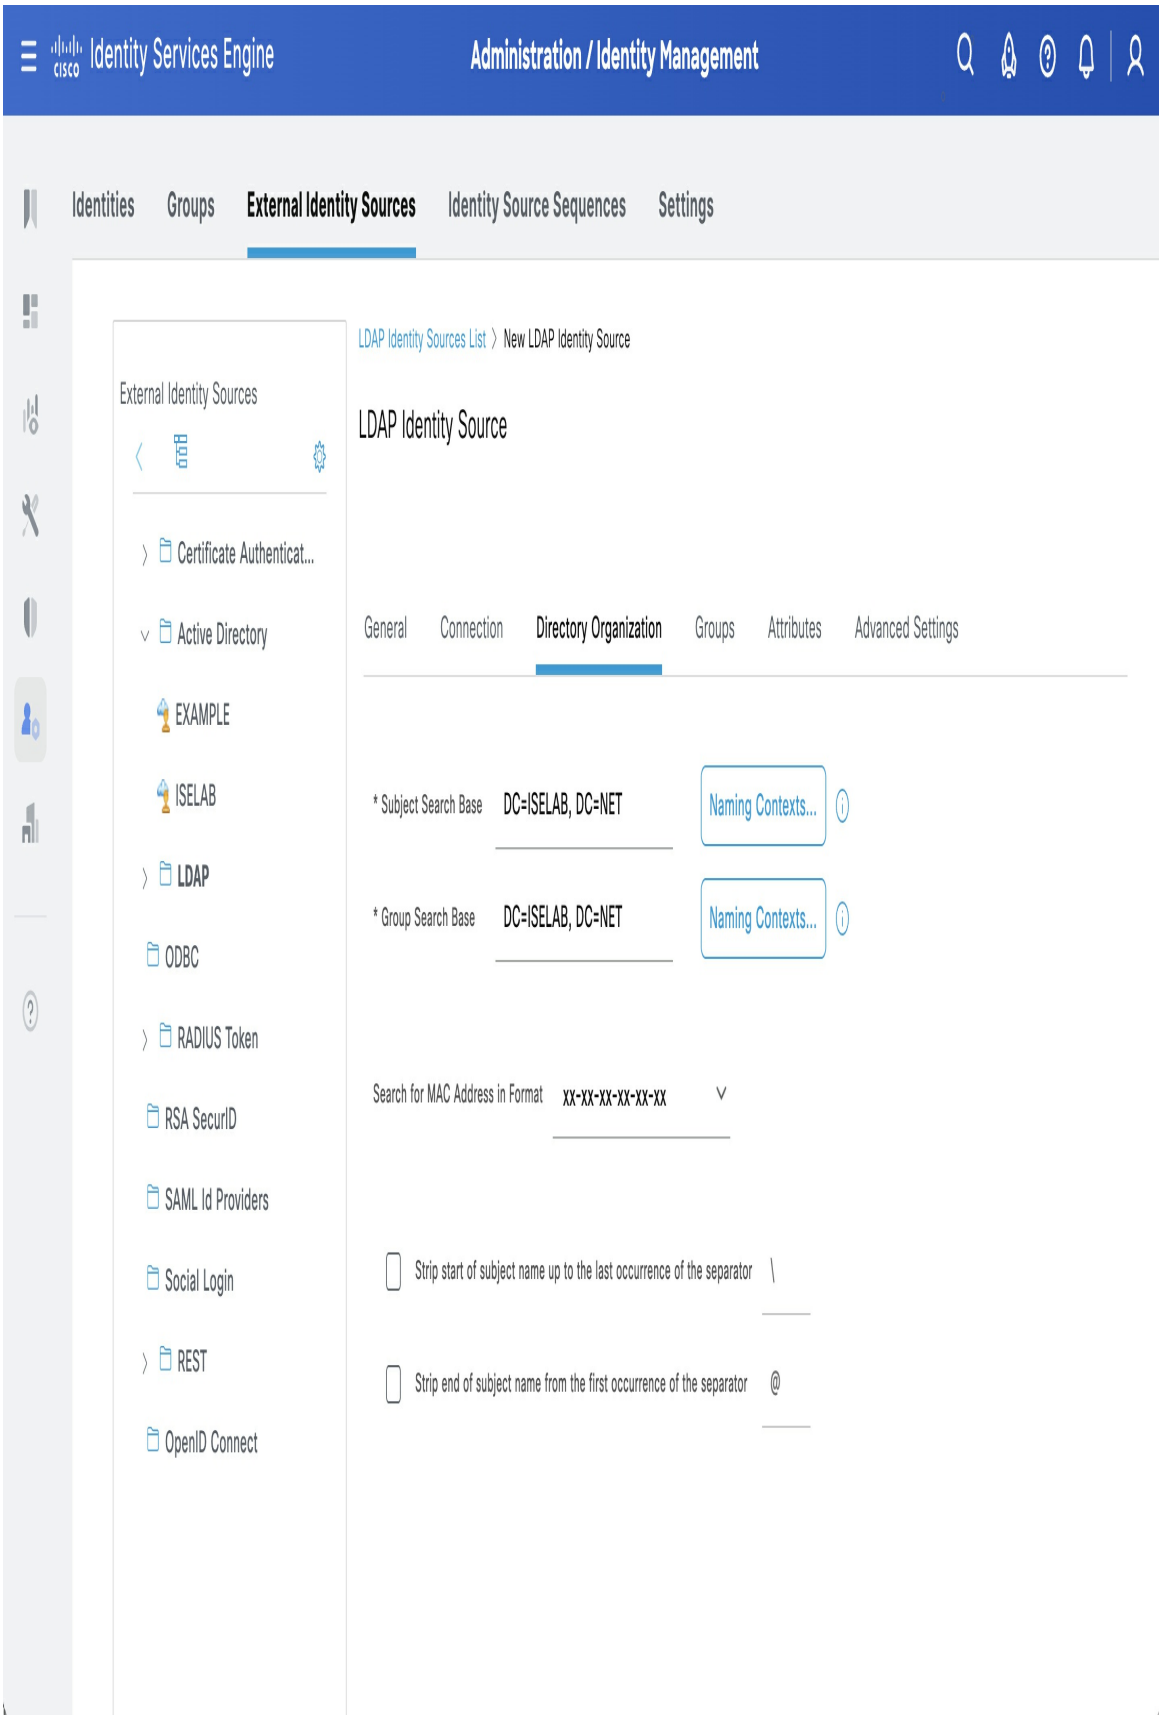

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
External Identity Sources
Certificate Authenticat...
Active Directory
EXAMPLE
ISELAB
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
LDAP Identity Sources List
New LDAP Identity Source
LDAP Identity Source
General
Connection
Directory Organization
Groups
Attributes
Advanced Settings
* Subject Search Base
DC=ISELAB, DC=NET
Naming Contexts...
* Group Search Base
DC=ISELAB, DC=NET
Naming Contexts...
Search for MAC Address in Format
xx-xx-xx-xx-xx-xx
Strip start of subject name up to the last occurrence of the separator
\
Strip end of subject name from the first occurrence of the separator
@

Figure 5-31 *LDAP Directory Organization Connection*

As with Active Directory, you can retrieve LDAP groups and attributes and use them in authorization policies.

You retrieve groups from the **Groups** tab. Click **Add** and choose **Select Groups from Directory**, as shown in Figure 5-32. After ISE retrieves the groups, you can select the ones you want to use and click **OK**. The same logic applies to LDAP attributes.

Technet24

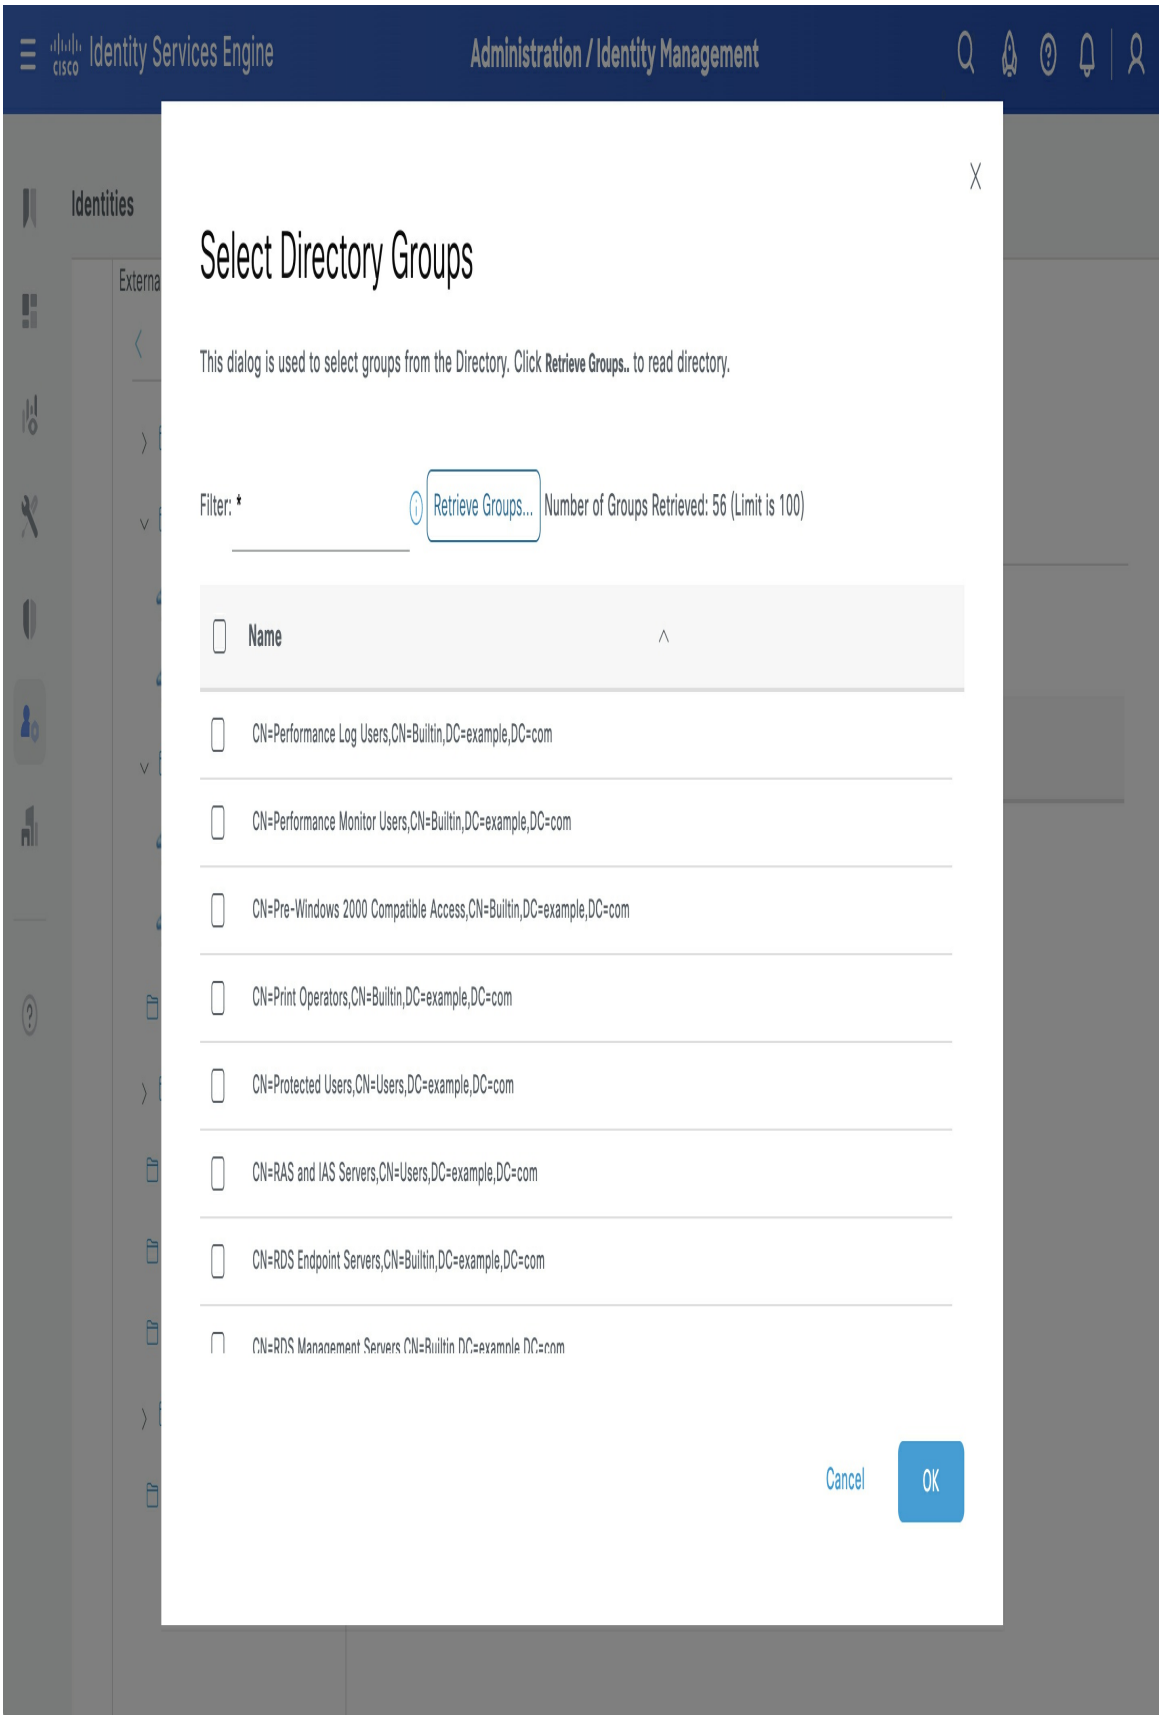

Identity Services Engine
Administration / Identity Management
Identities
Select Directory Groups
This dialog is used to select groups from the Directory. Click Retrieve Groups.. to read directory.
Filter: *
Retrieve Groups...
Number of Groups Retrieved: 56 (Limit is 100)
Name
CN=Performance Log Users,CN=Builtin,DC=example,DC=com
CN=Performance Monitor Users,CN=Builtin,DC=example,DC=com
CN=Pre-Windows 2000 Compatible Access,CN=Builtin,DC=example,DC=com
CN=Print Operators,CN=Builtin,DC=example,DC=com
CN=Protected Users,CN=Users,DC=example,DC=com
CN=RAS and IAS Servers,CN=Users,DC=example,DC=com
CN=RDS Endpoint Servers,CN=Builtin,DC=example,DC=com
Cancel
OK

Figure 5-32 *LDAP Directory Groups Selection*

REST

Many organizations are moving from traditional Active Directory to Microsoft Entra ID, formerly known as Azure Active Directory. A cloud-native identity system aligns better with modern security requirements, scales effectively, and is well suited for cloud application architectures.

Starting with ISE 3.1, ISE supports authentication and authorization with Microsoft Entra ID. Initially, only username and password-based authentication was supported, relying on the Resource Owner Password Credentials (ROPC) protocol. The flow diagram is shown in Figure 5-33.

Technet24

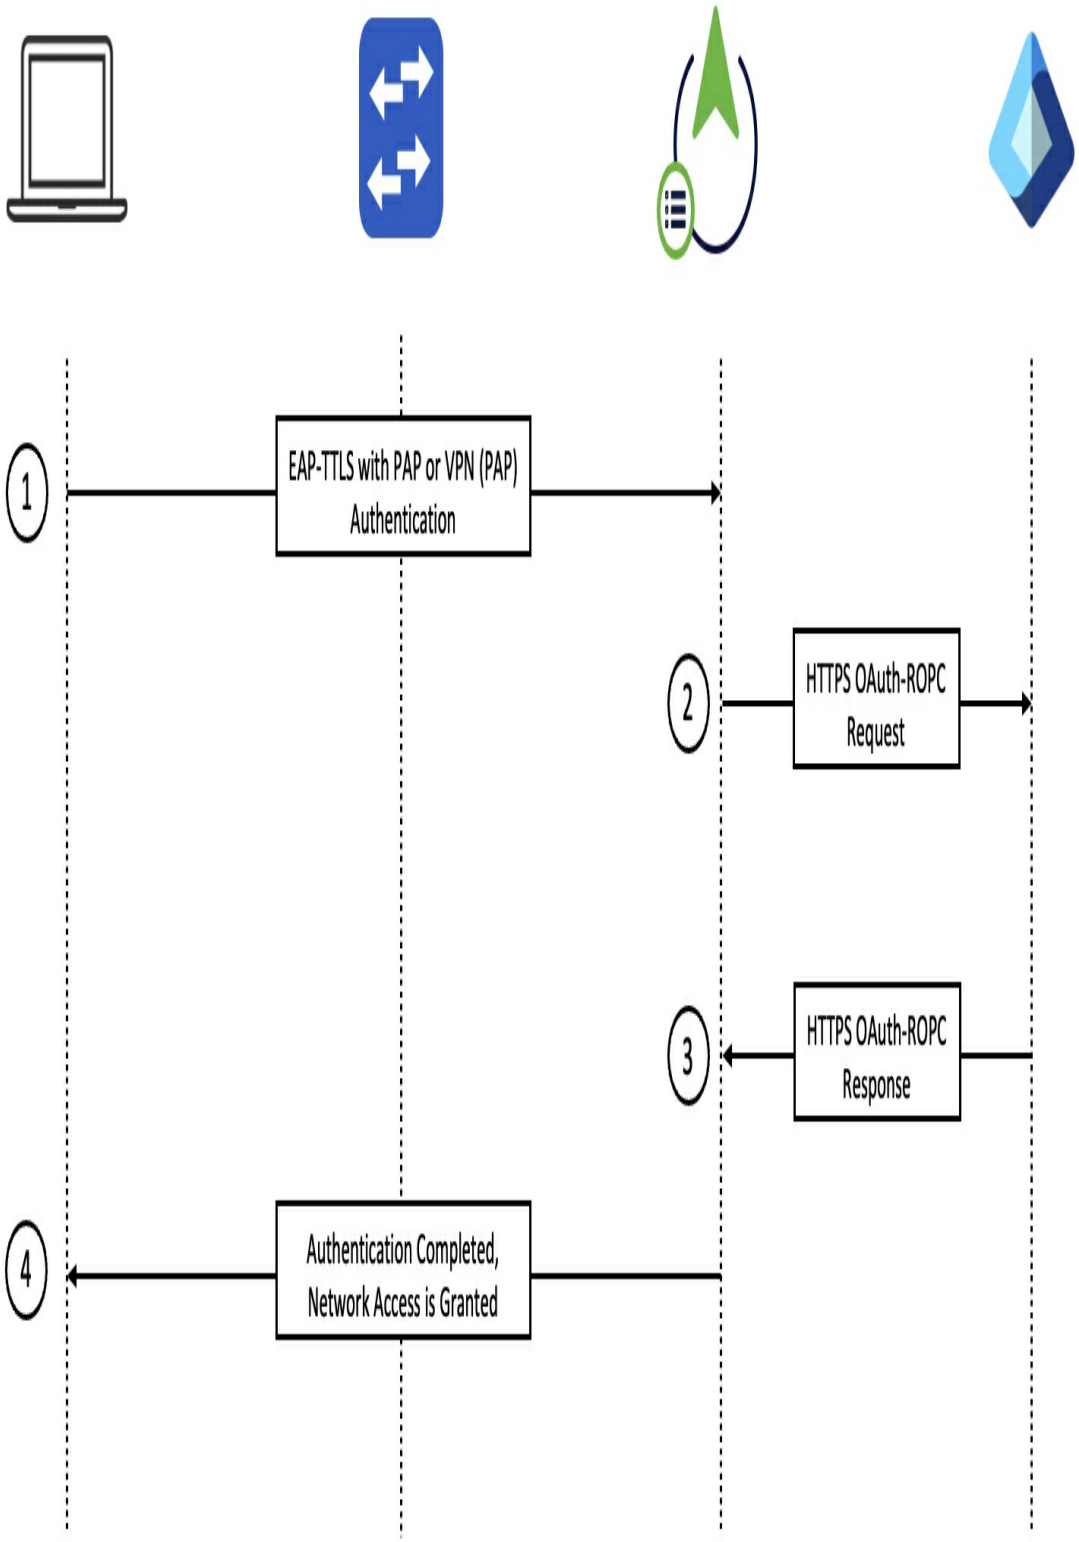

1
EAP-TTLS with PAP or VPN (PAP) Authentication
2
HTTPS OAuth-ROPC Request
3
HTTPS OAuth-ROPC Response
4
Authentication Completed, Network Access is Granted

Figure 5-33 *ISE ROPC Authentication Flow with Entra ID*

Starting with ISE 3.2, ISE also supports user certificate-based authentication. As a result, EAP-TLS and TEAP are available options. The flow diagram is shown in Figure 5-34.

Technet24

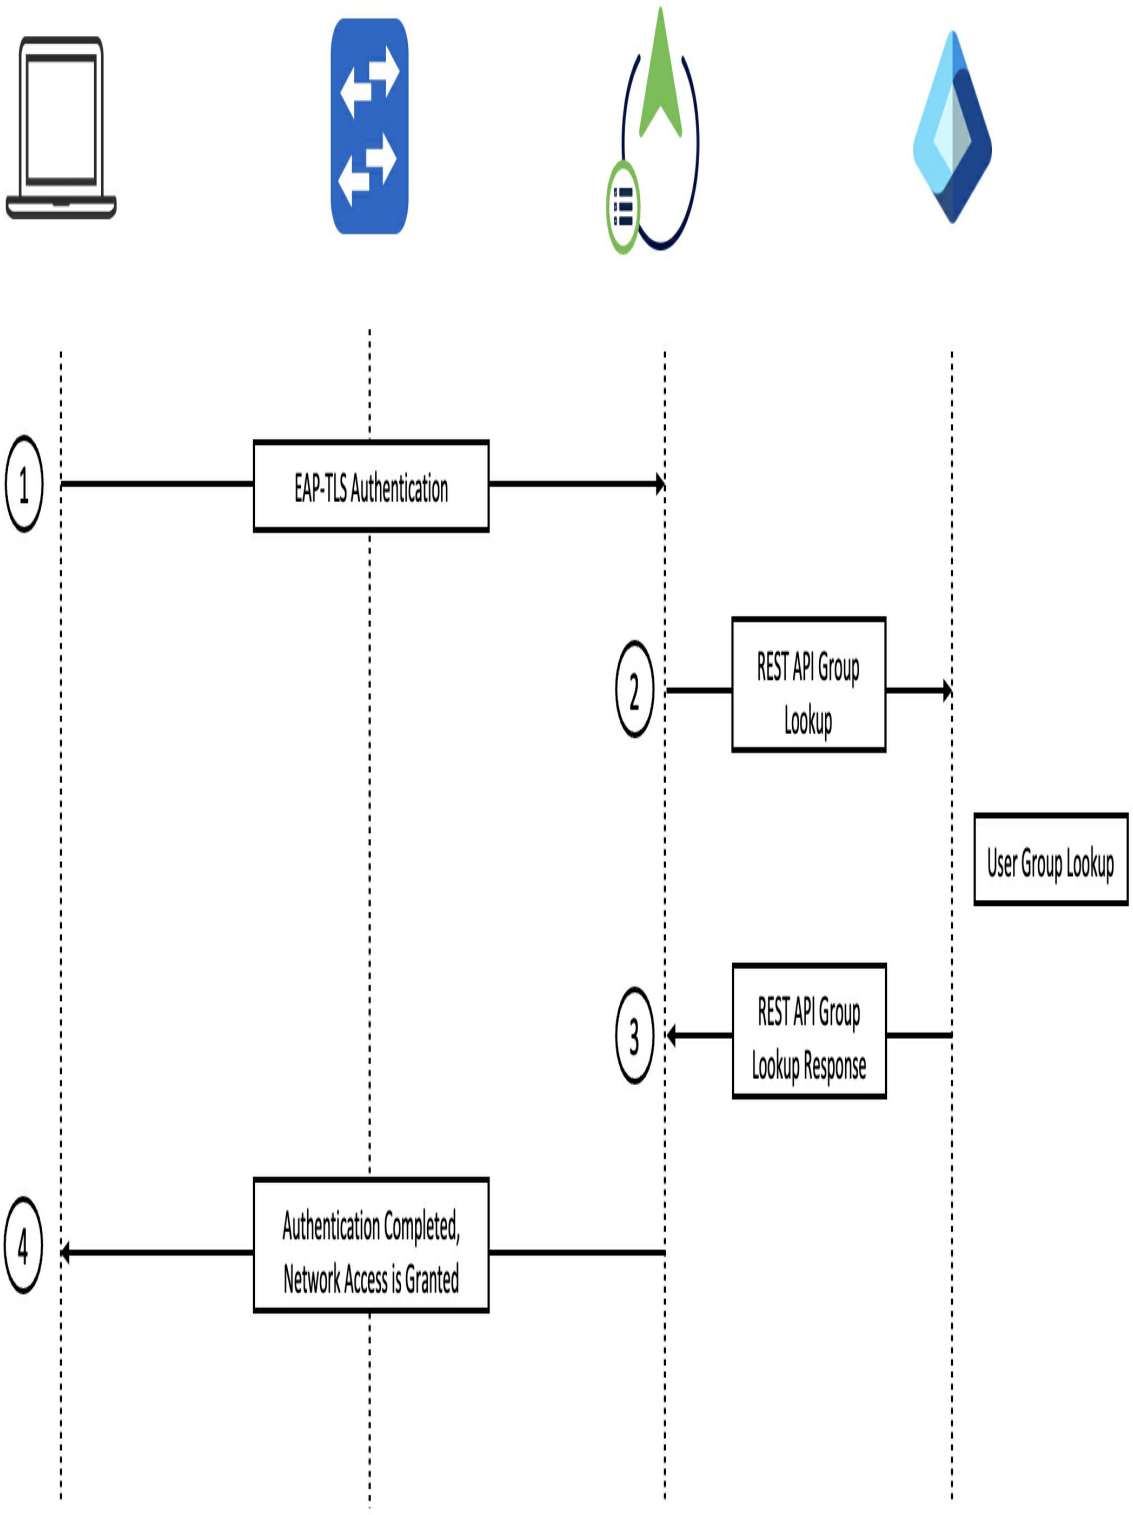

1
EAP-TLS Authentication
2
REST API Group Lookup
User Group Lookup
3
REST API Group Lookup Response
4
Authentication Completed, Network Access is Granted

Figure 5-34 *ISE EAP-TLS Authentication with Group Lookup in Entra ID*

Starting with ISE 3.4 patch 4 and ISE 3.5, ISE supports device authentication. This capability enables TEAP with EAP chaining as an option. With this configuration, you can perform both user and device group and attribute lookups. The flow diagram is shown in Figure 5-35.

Technet24

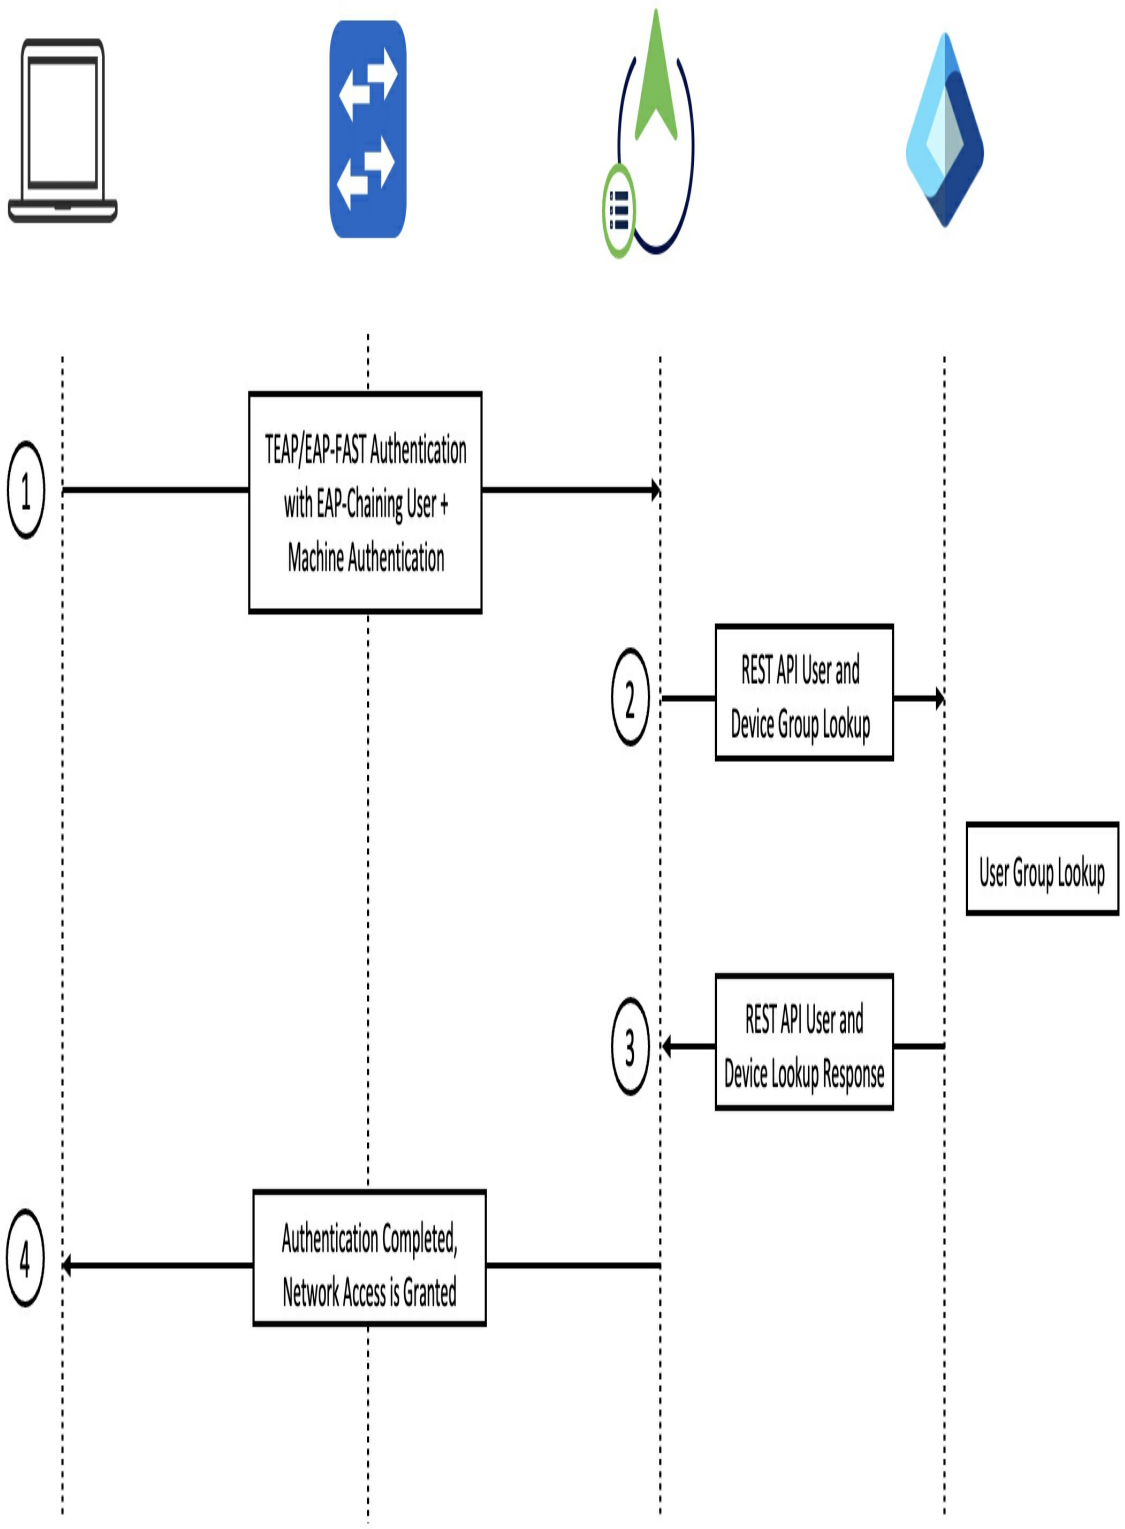

1
TEAP/EAP-FAST Authentication with EAP-Chaining User + Machine Authentication
2
REST API User and Device Group Lookup
User Group Lookup
3
REST API User and Device Lookup Response
4
Authentication Completed, Network Access is Granted

Figure 5-35 *ISE EAP-Chaining Authentication with User Group Lookup in Entra ID*

Starting with ISE 3.5, ISE also supports device administration by using TACACS+. The flow diagram is shown in Figure 5-36.

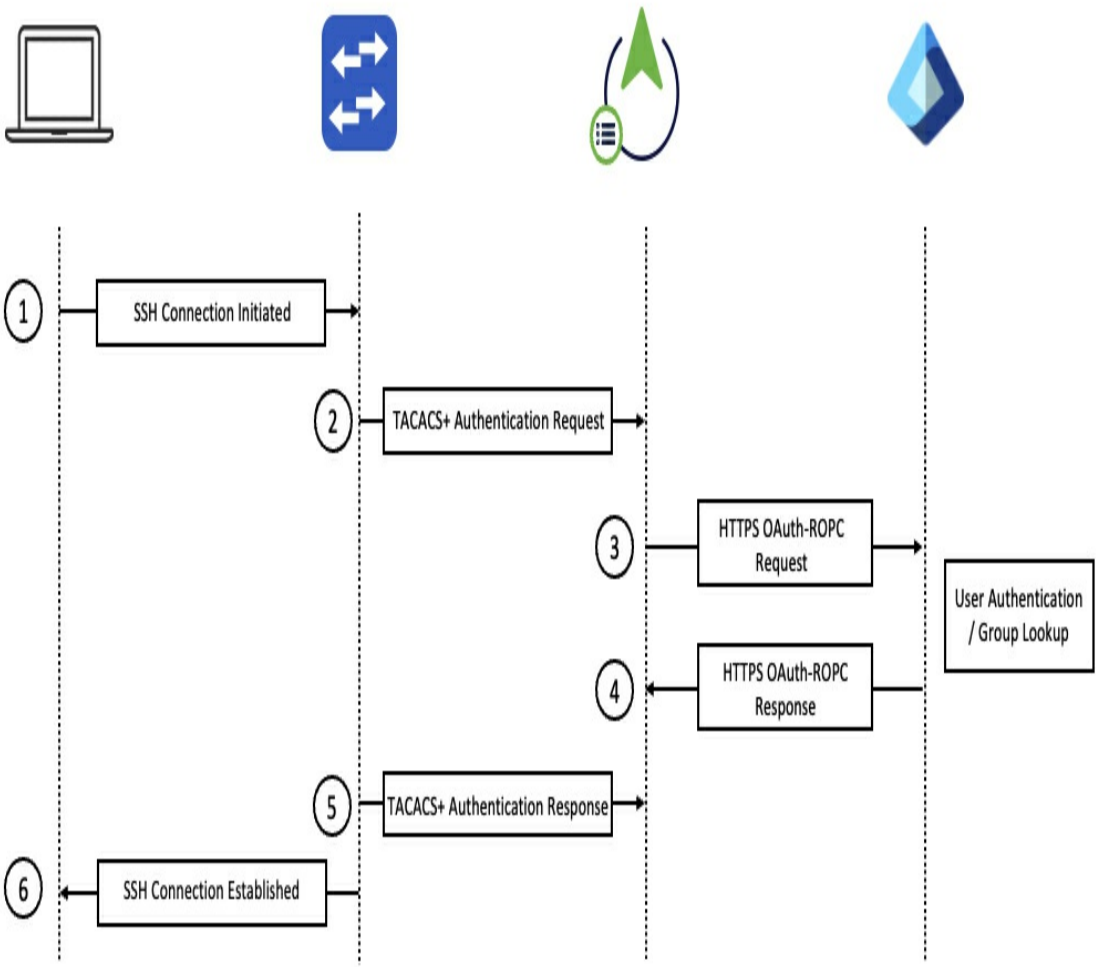


Figure 5-36 *Device Administration Using TACACS+ with User Group Lookup in Entra ID*

ISE integrates with the REST ID Store by using an application-to-

Technet24

application integration model. To configure the REST ID Store, navigate to **Administration > Identity Management > External Identity Sources > REST** and click **Add**.

For successful integration, you need to configure the following:

- Client ID
- Client Secret
- Tenant ID

You obtain all of these values from the Microsoft application that is responsible for the ISE and Entra ID integration. The ISE configuration is shown in Figure 5-37.

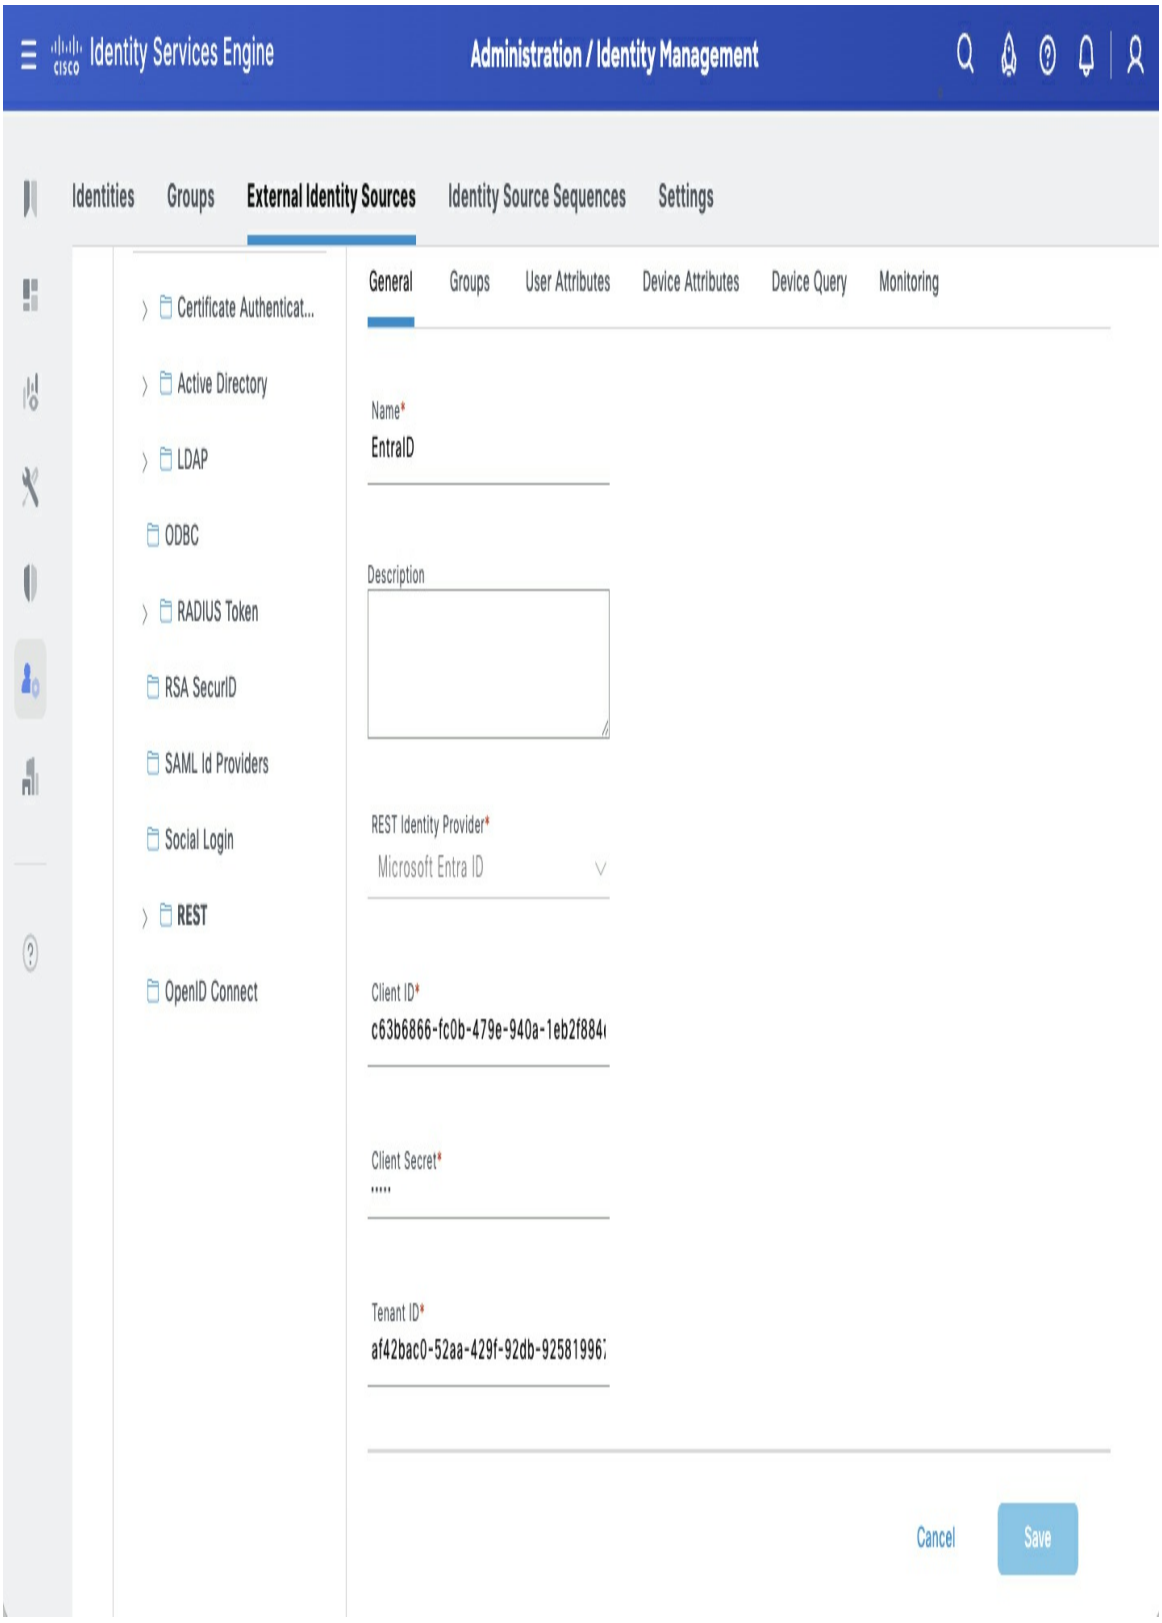
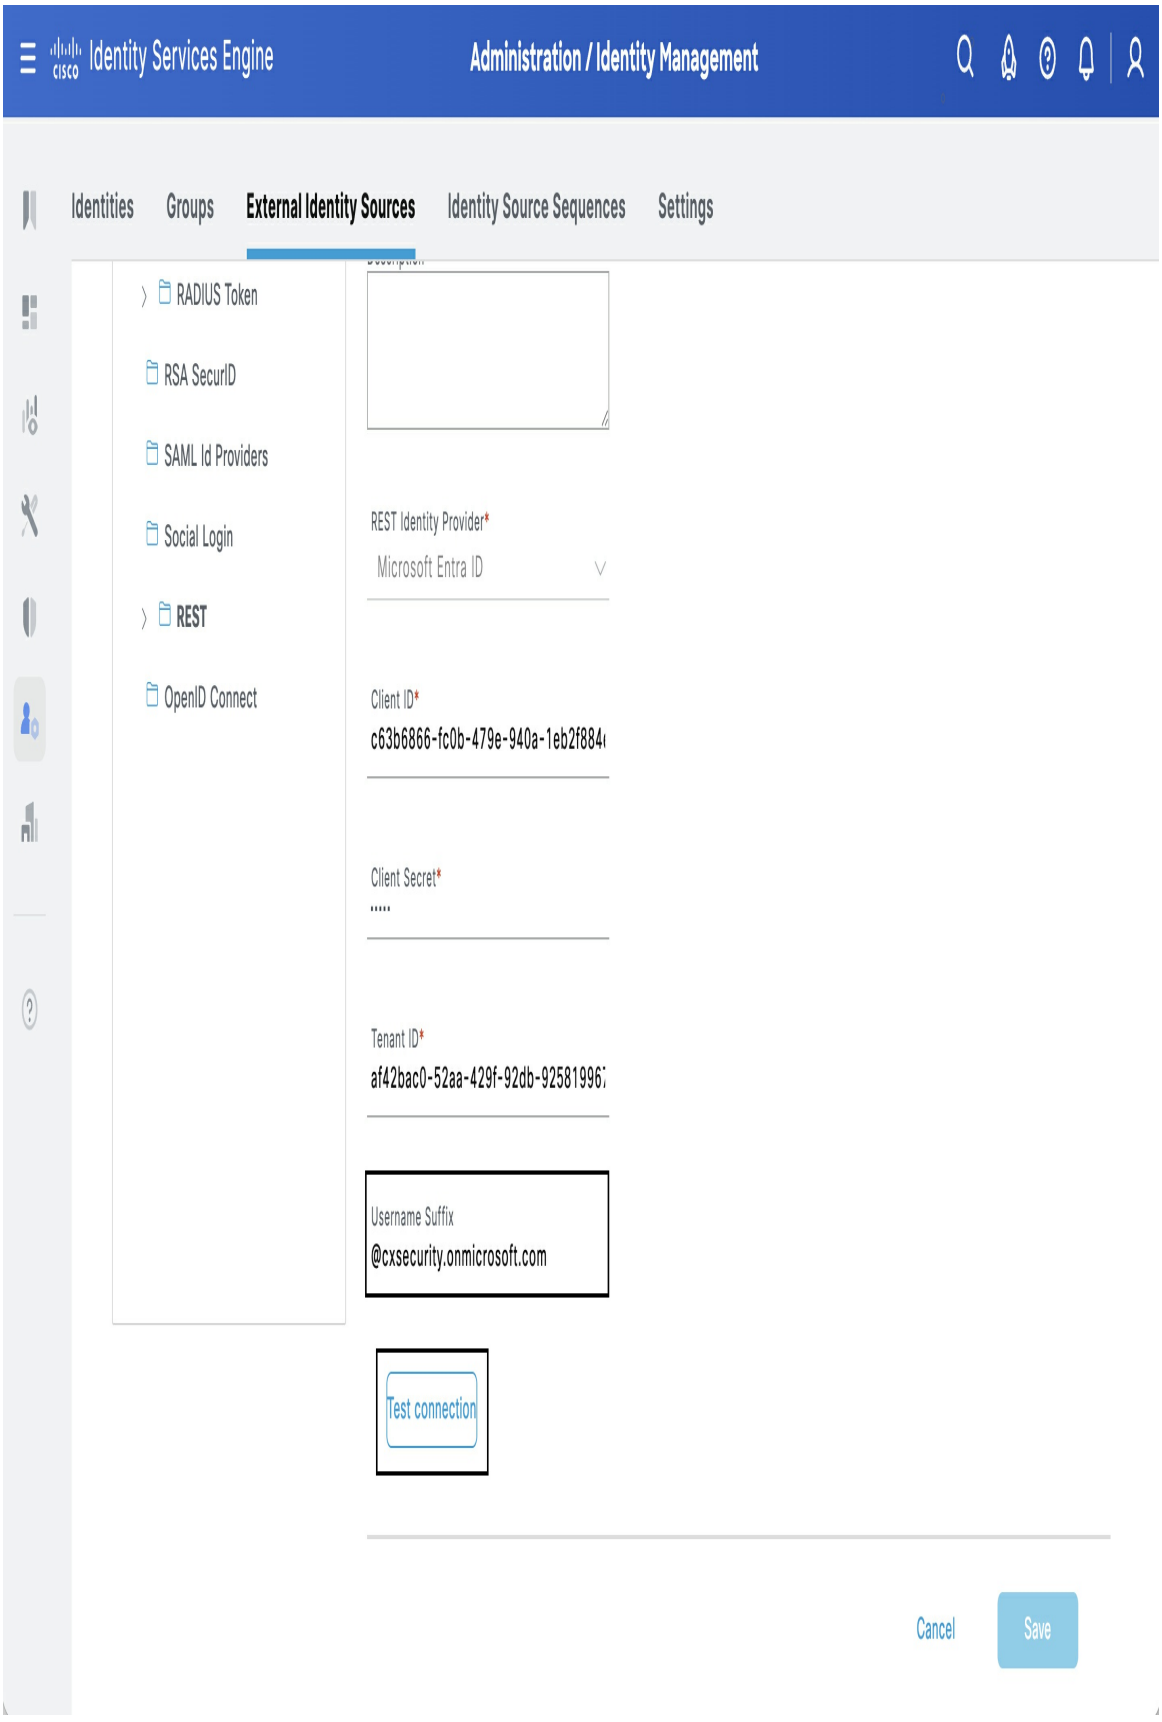


Figure 5-37 *Entra ID Configuration*

Technet24

Additionally, as shown in Figure 5-38, you can configure a username suffix. This configuration simplifies the user experience when you use ROPC for authentication. You do not need to enter the full user principal name. Instead, you can provide a short name, and ISE appends the configured suffix when it authenticates the user through Entra ID. Finally, you can select **Test Connection** to verify end-to-end connectivity between ISE and Entra ID and to confirm that the configuration data is valid.


Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
REST Identity Provider*
Microsoft Entra ID
Client ID*
c63b6866-fc0b-479e-940a-1eb2f884
Client Secret*
Tenant ID*
af42bac0-52aa-429f-92db-92581996
Username Suffix
@cxsecurity.onmicrosoft.com
Test connection
Cancel
Save

Technet24

Figure 5-38 *Entra ID Configuration*

To configure an application in Entra ID, navigate to portal.azure.com and select **App Registrations** from the search menu. Select **New Registration**, as shown in Figure 5-39.

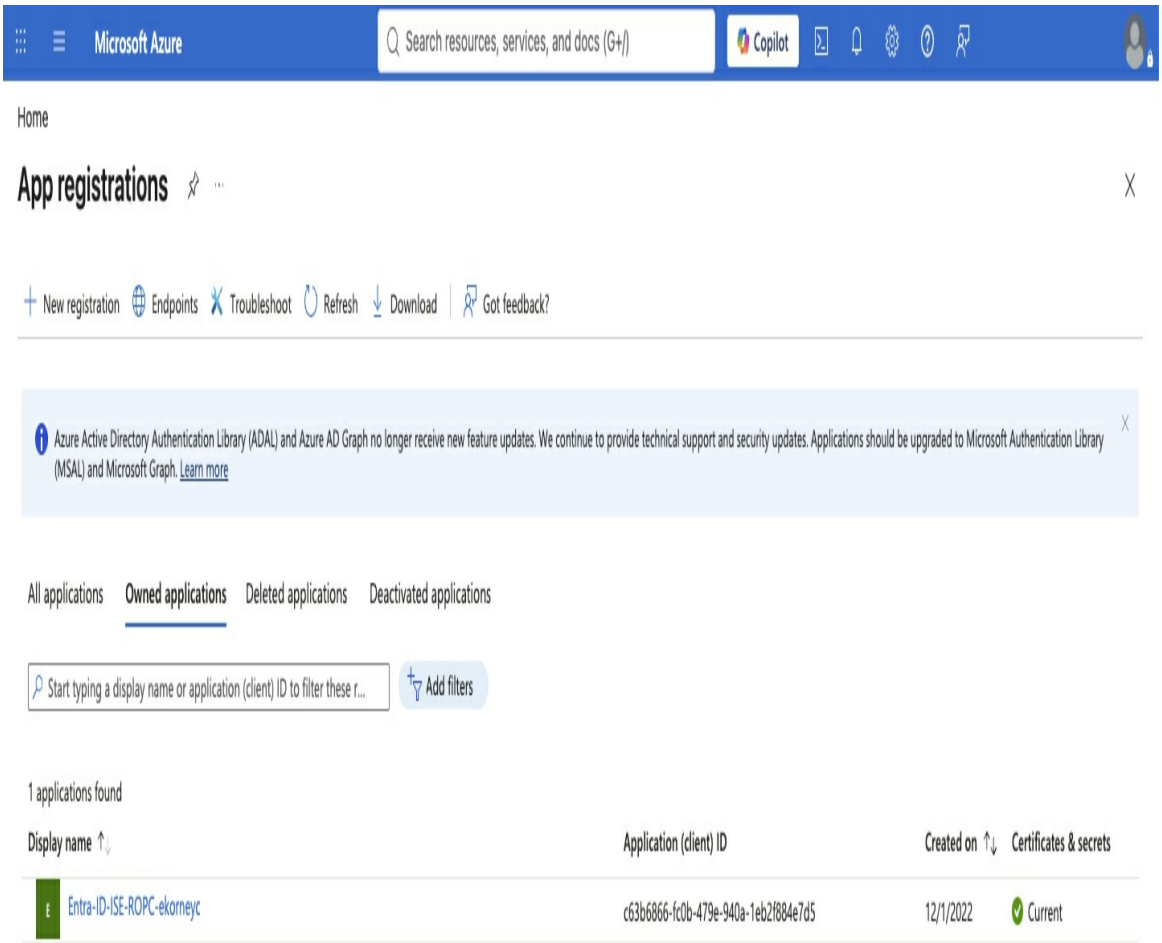


Figure 5-39 *Entra ID App Registration*

Technet24

Populate the name for the application and select **Register**, as shown in Figure 5-40.

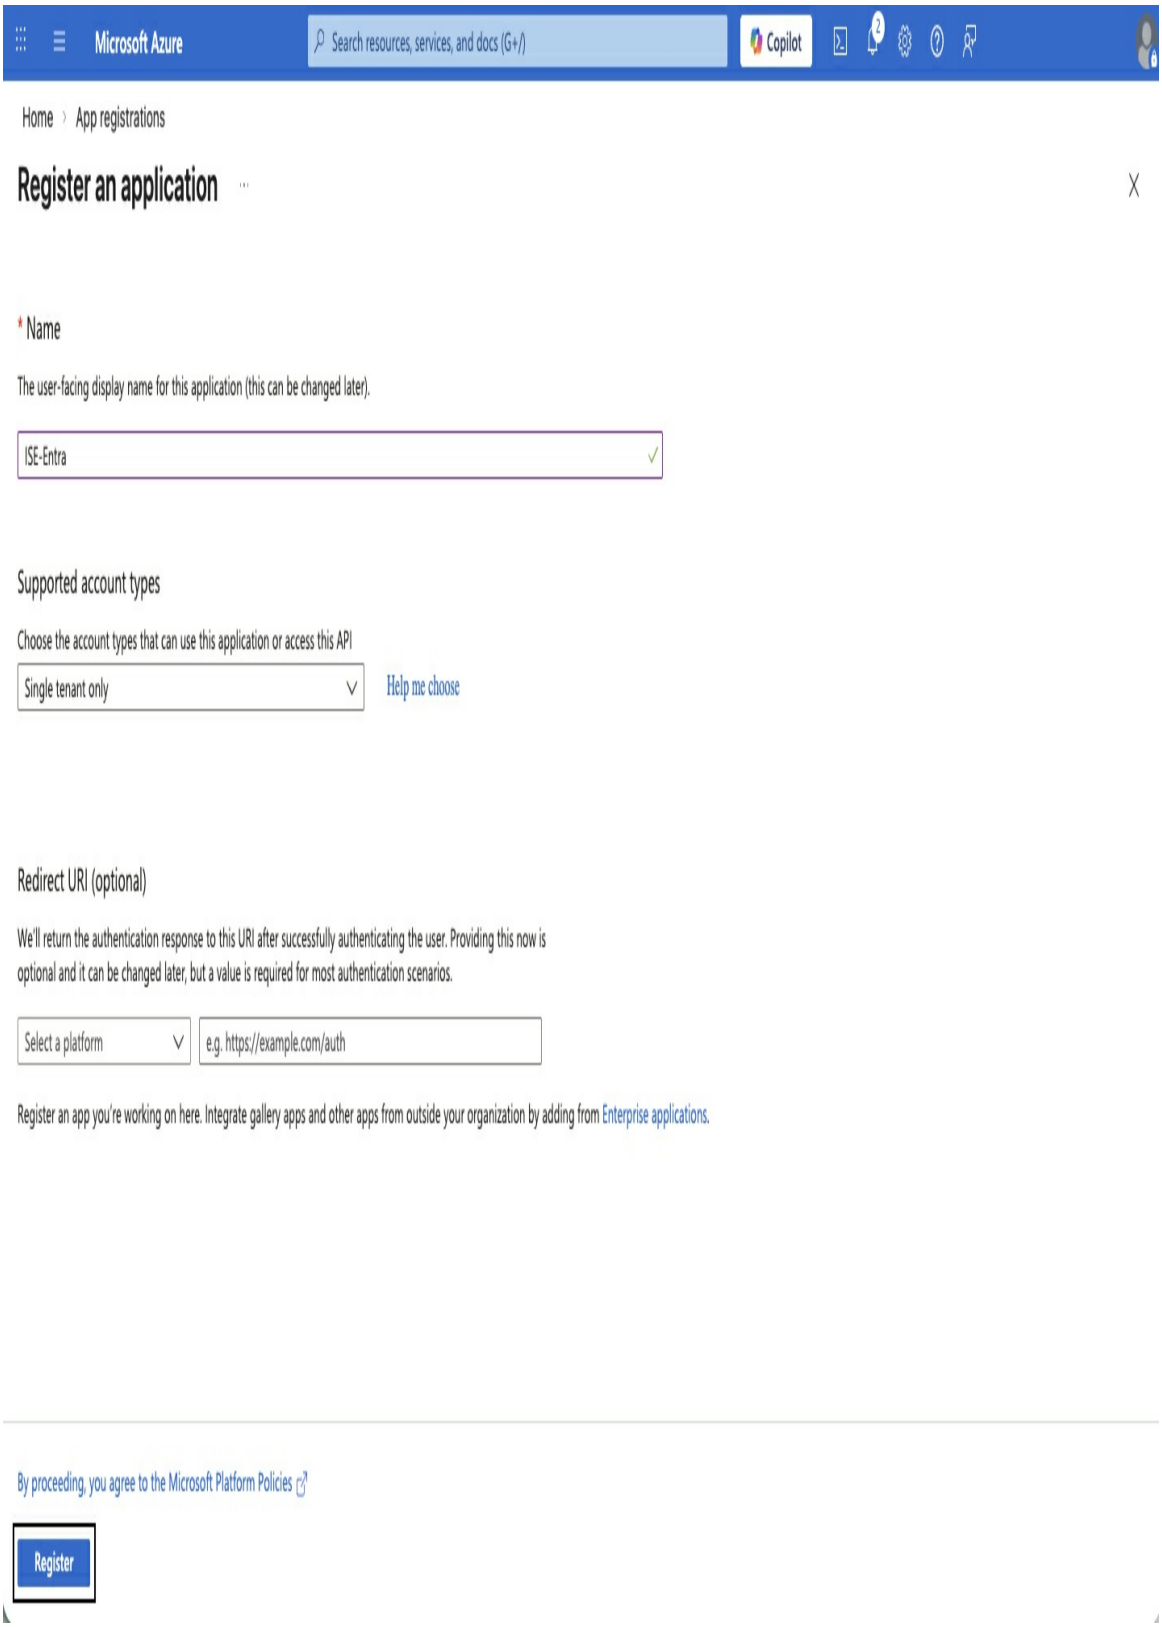


Figure 5-40 *Entra ID App Registration*

Technet24

After you register the application, note the **Application ID** and **Tenant ID**. You use these values to configure the ISE integration, as explained earlier. Their location is shown in Figure 5-41.

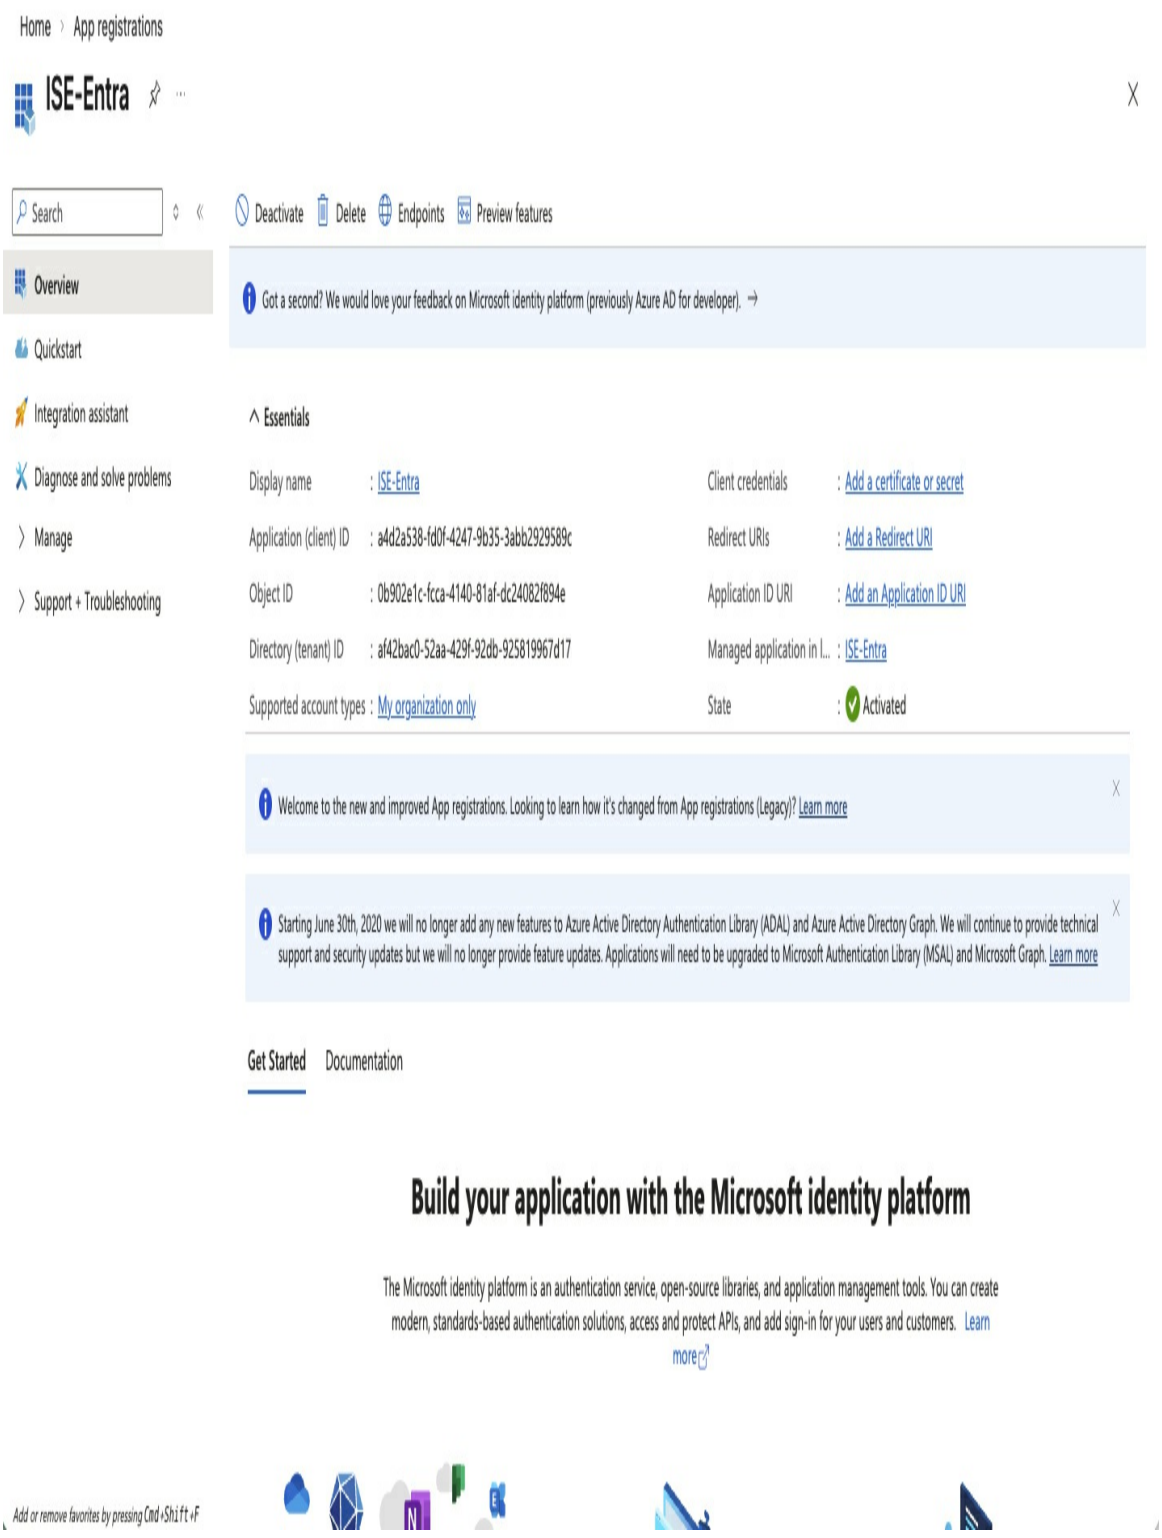


Figure 5-41 *Entra ID App Registration Overview*

Navigate to **Manage > Certificates & Secrets** and click **New Client**

Technet24

Secret. Provide a description and click **Add**, as shown in Figure 5-42.

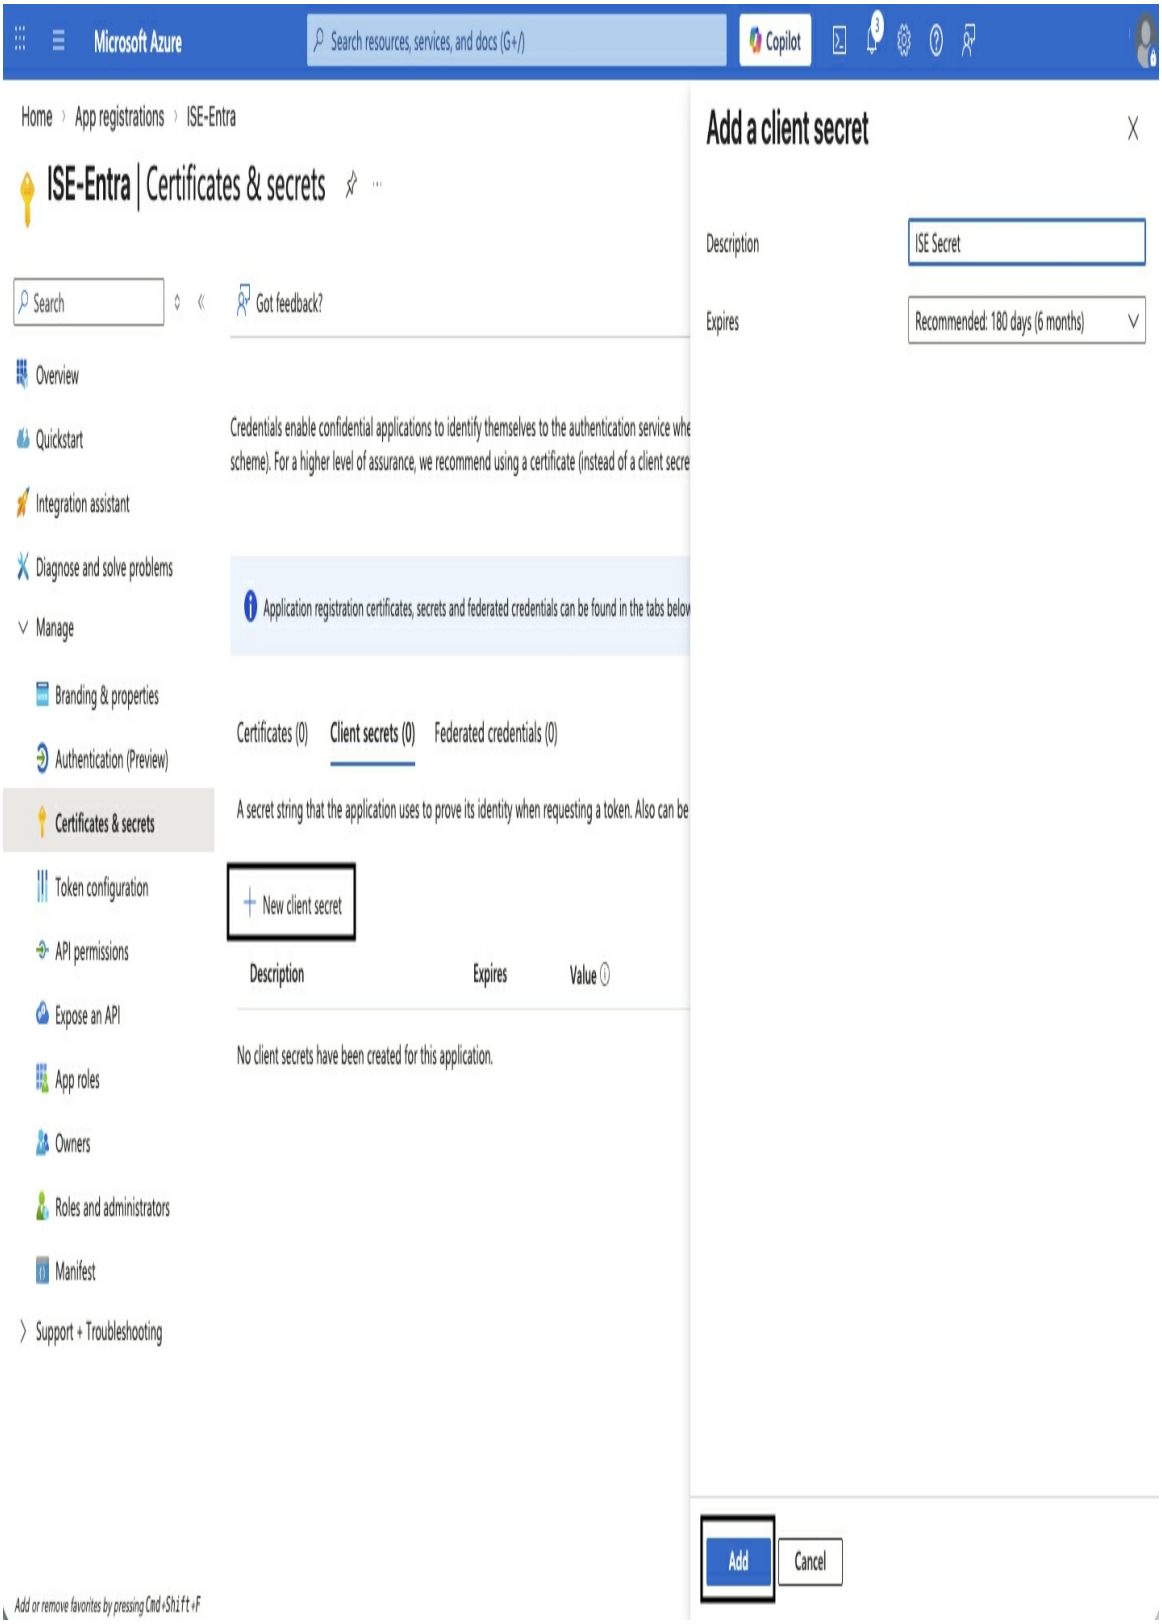


Figure 5-42 *Entra ID App Registration Client Secret Configuration*

It is crucial that you copy the value of the secret key, not the secret ID. You use this value to configure ISE, as shown in Figure 5-43.

Technet24

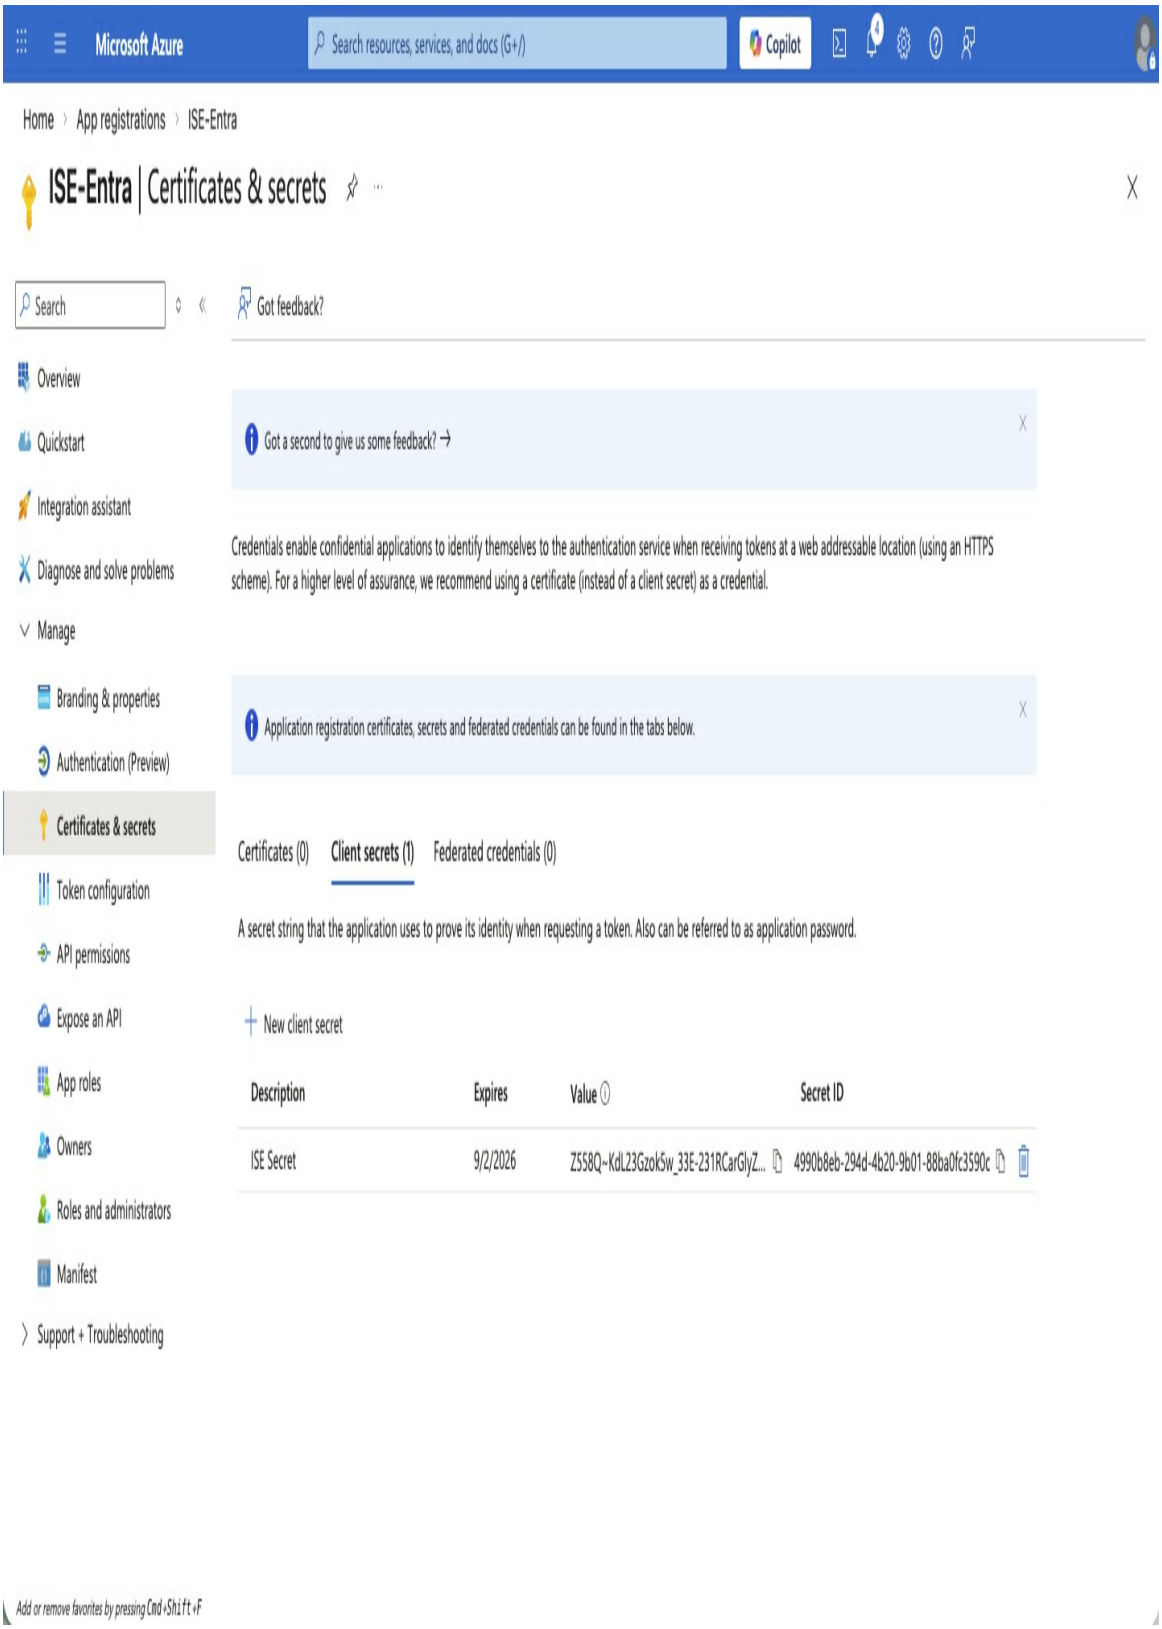


Figure 5-43 *Entra ID App Registration Client Secret Configuration*

Navigate to **Authentication > Settings**, enable **Allow Public Client Flows**, and then click **Save**, as shown in Figure 5-44. This setting allows the ROPC flow to run for this application.

Technet24

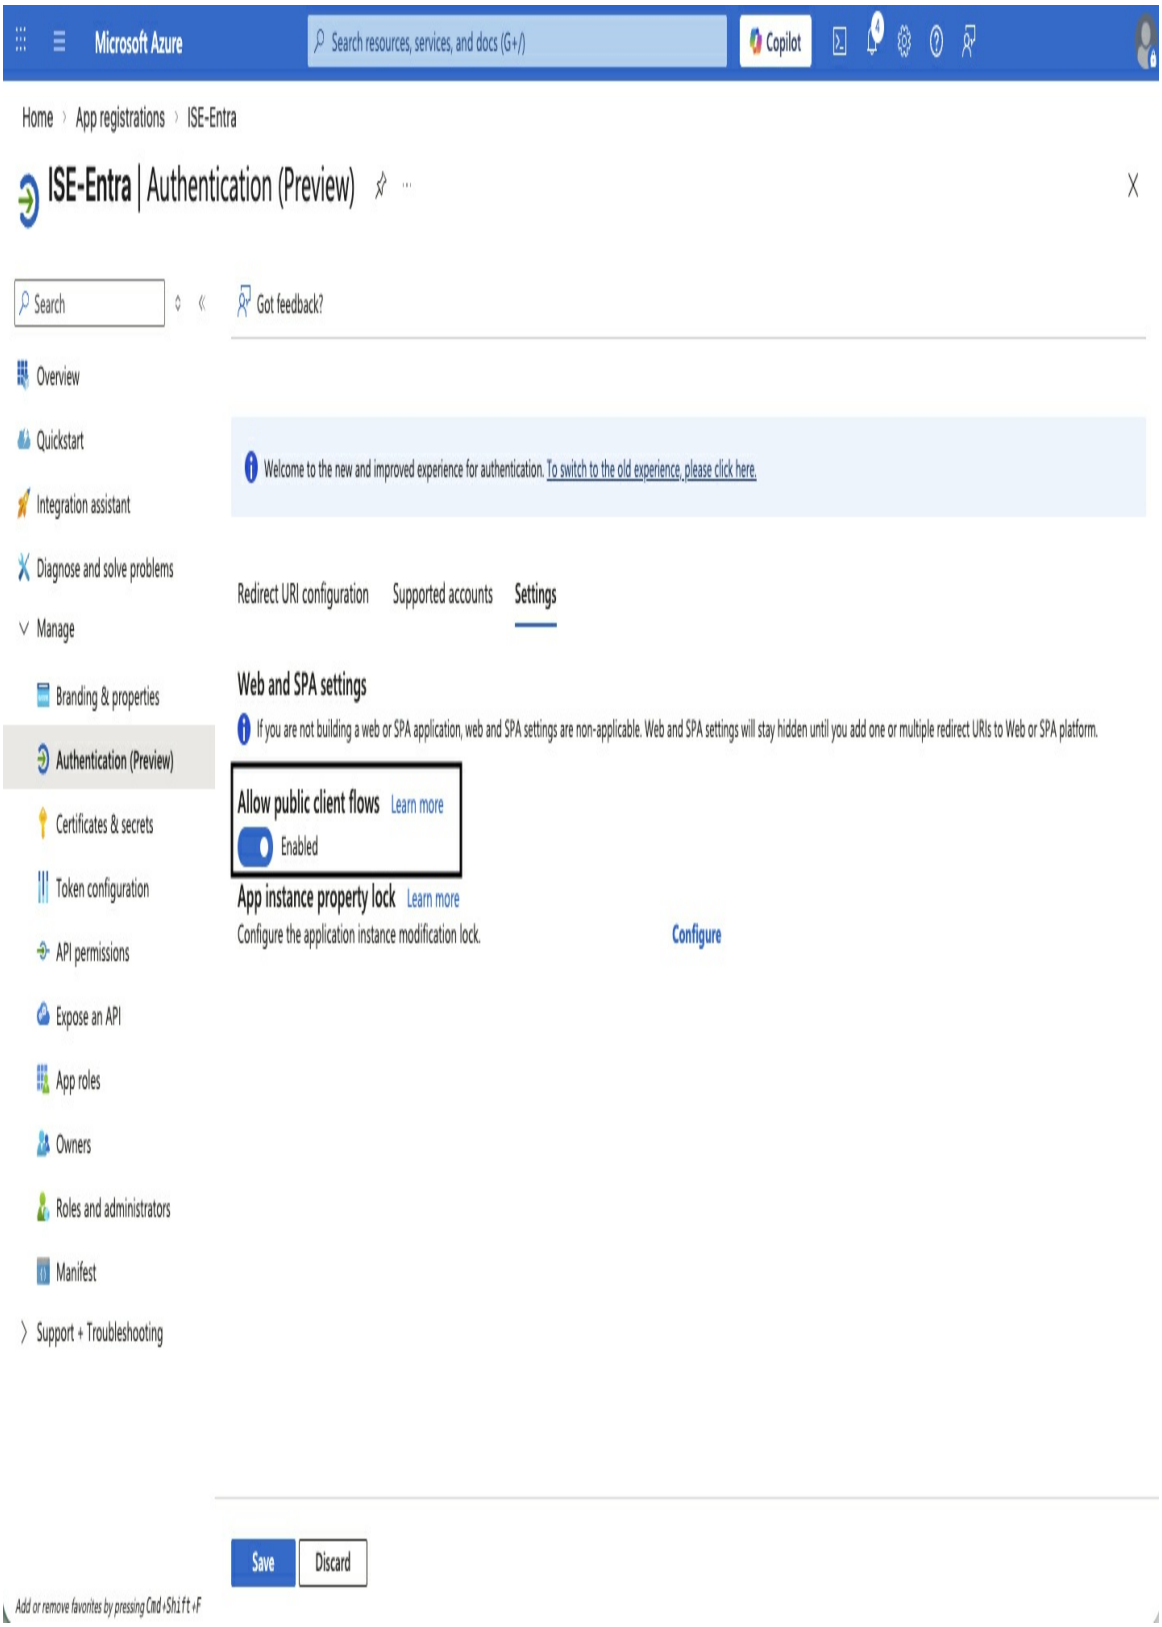


Figure 5-44 *Entra ID App Public Client Flow Configuration*

Navigate to **Token Configuration**, click **Add Groups Claim**, and select the groups that you want to use for this integration. Click **Add**, as shown in Figure 5-45.

Technet24

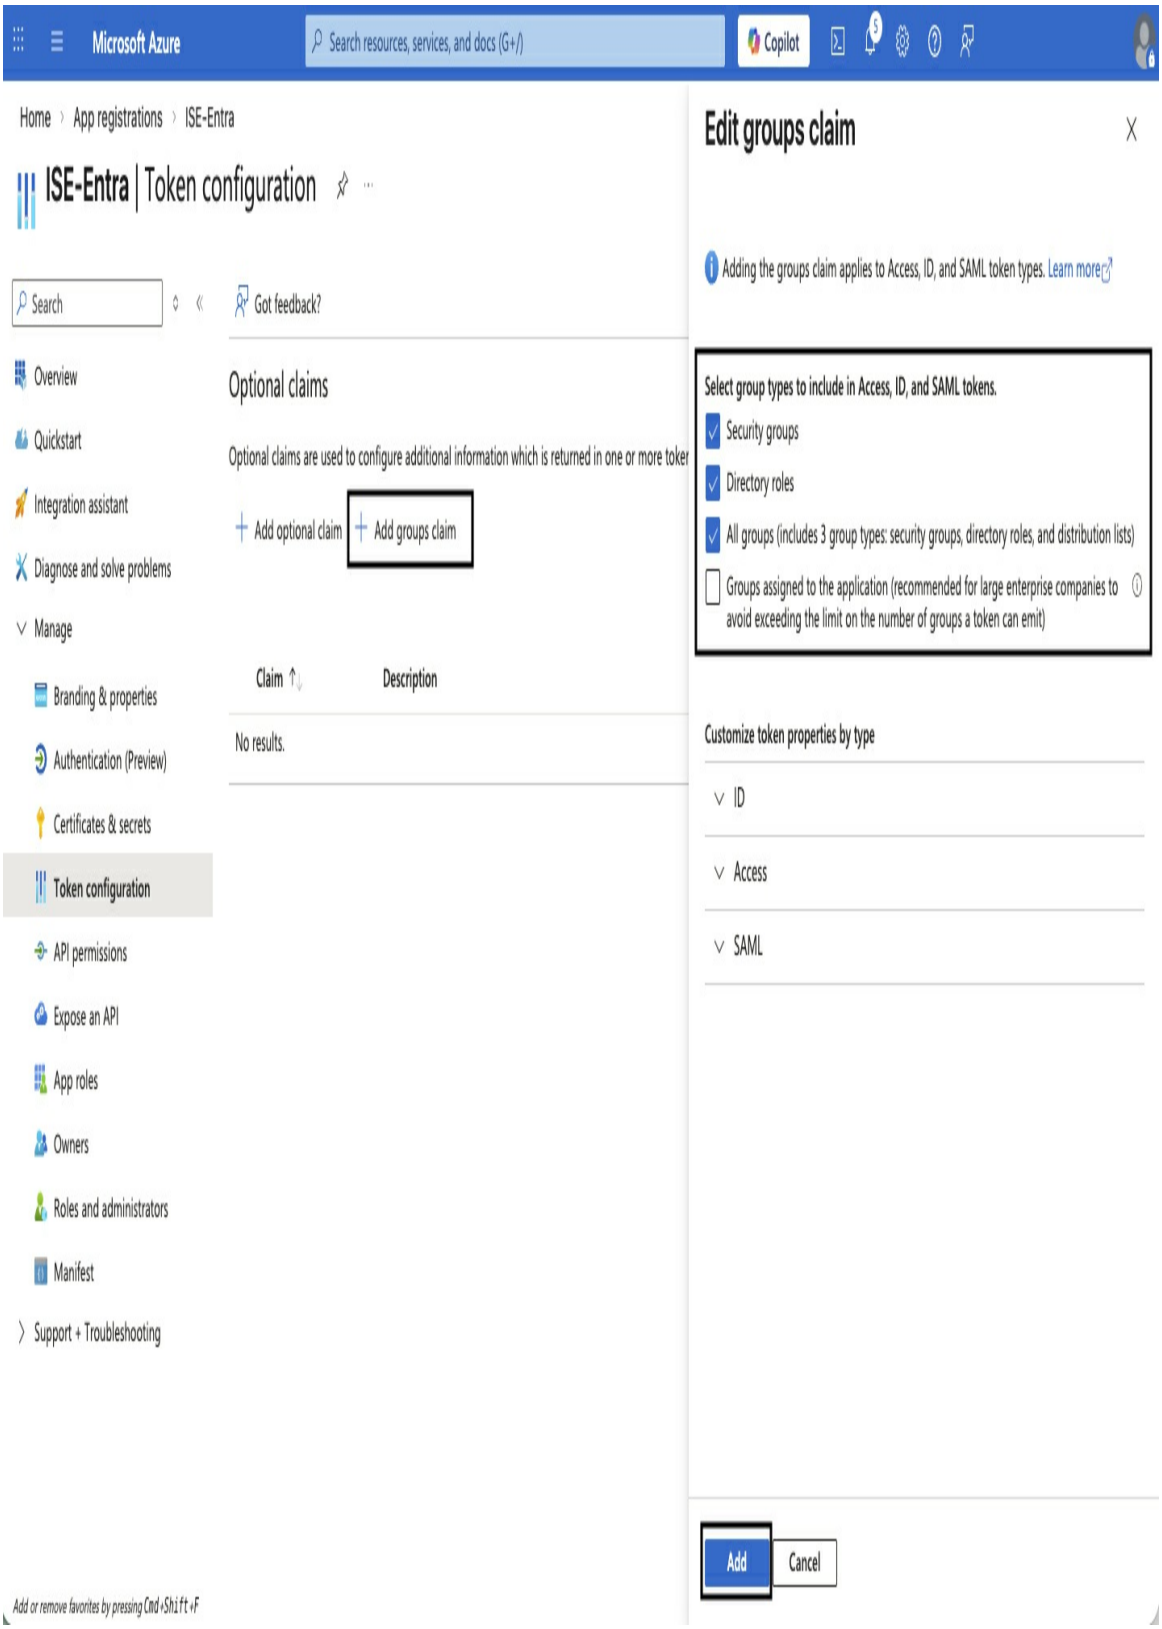


Figure 5-45 *Entra ID App Group Claim Configuration*

Navigate to **API Permissions**, select **+Add a Permission**, and choose **Microsoft Graph**, as shown in Figure 5-46.

Technet24

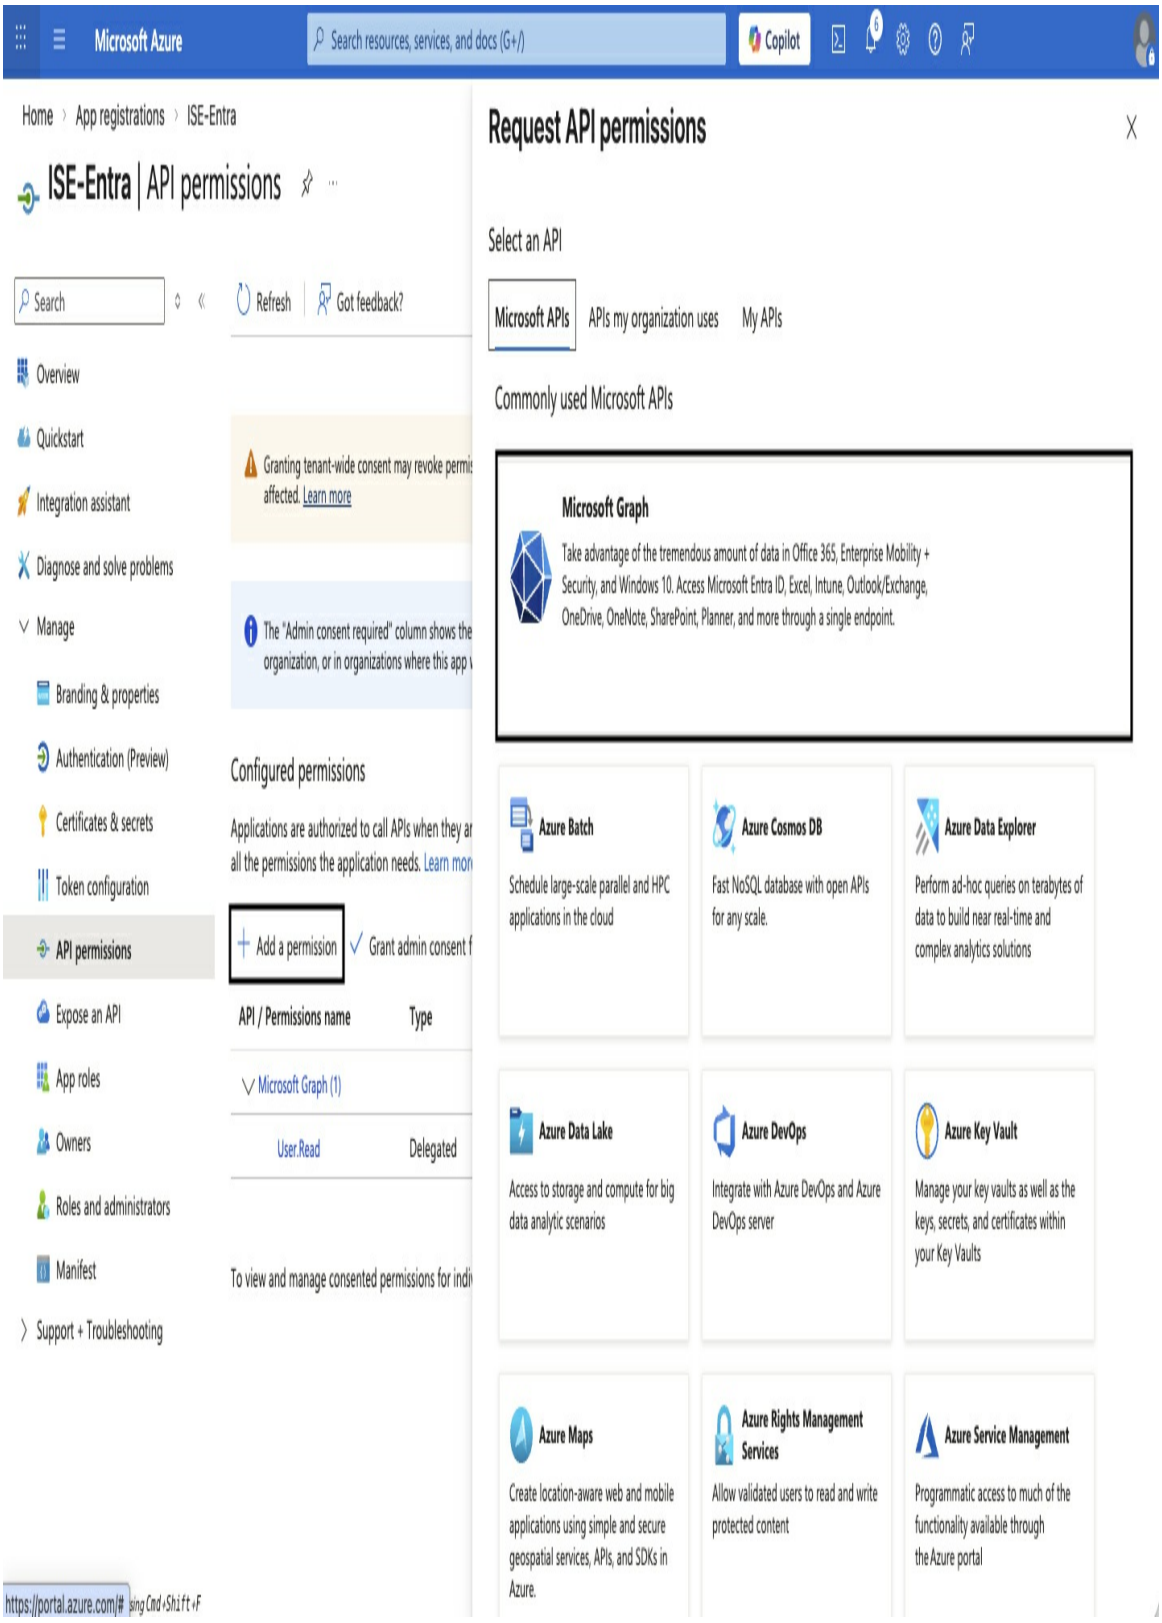
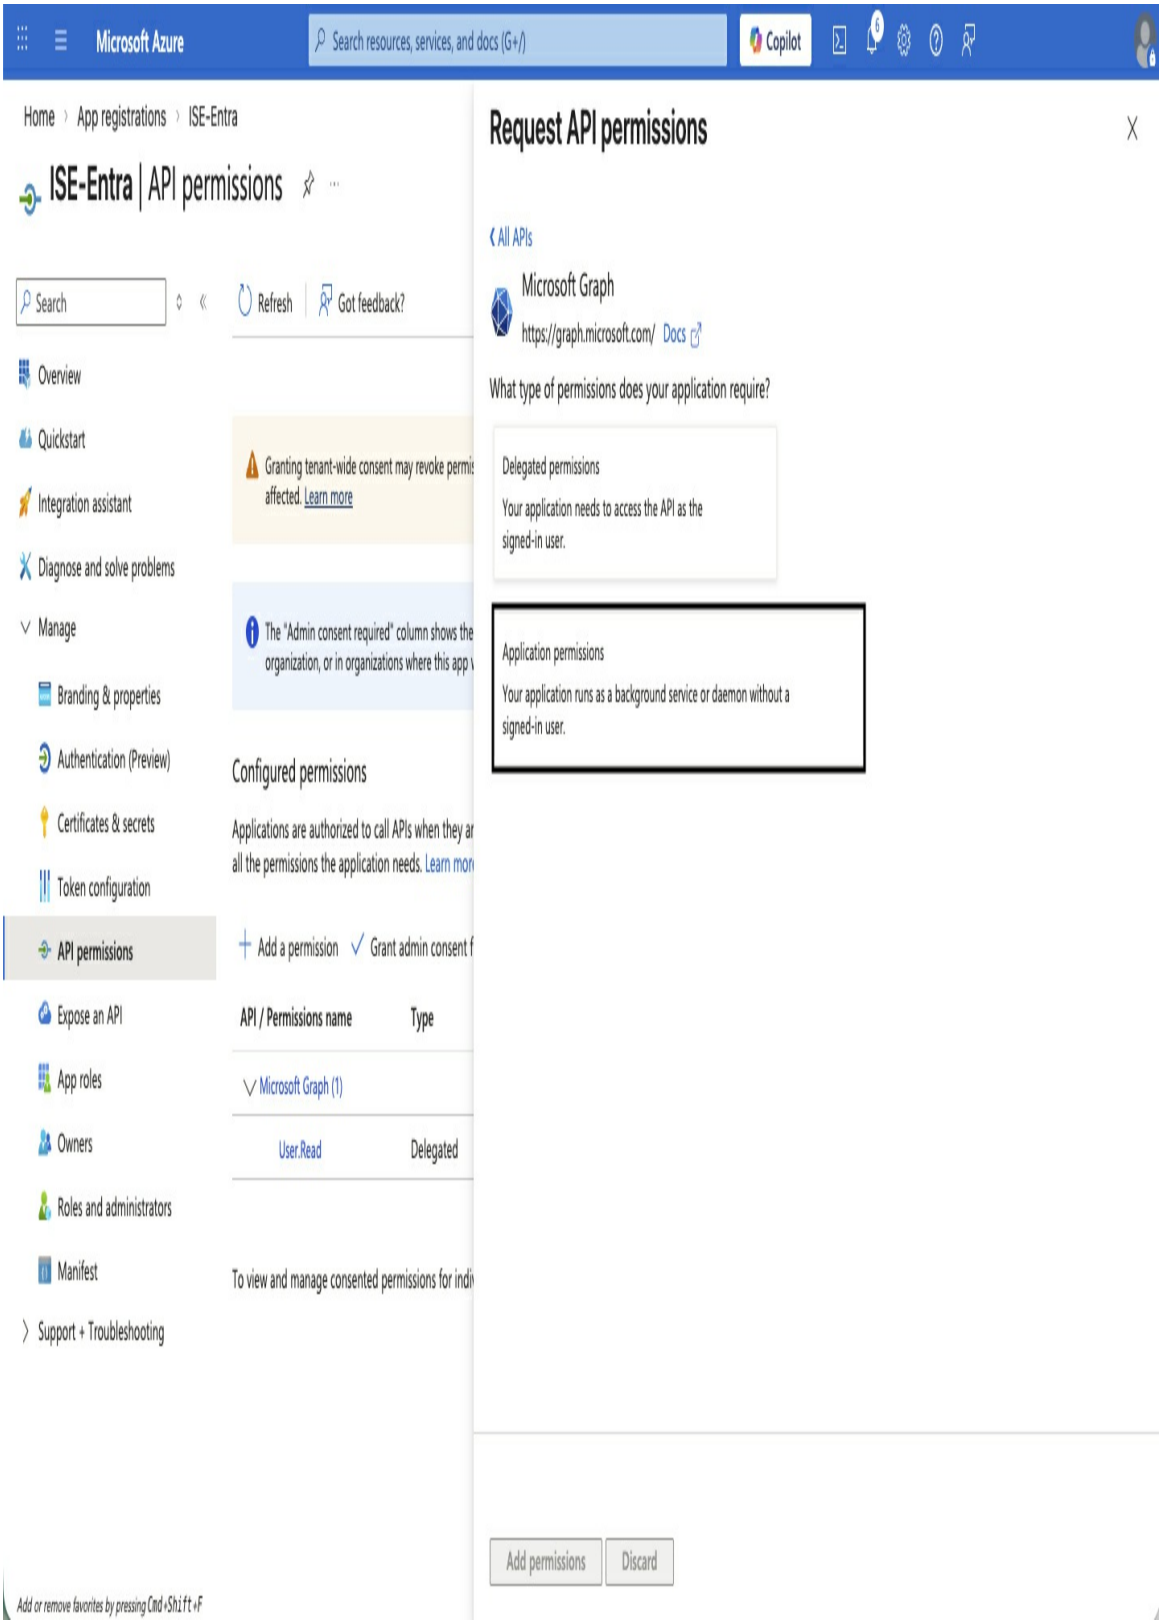


Figure 5-46 *Entra ID App API Permissions Configuration*

Select **Application Permissions**, as shown in Figure 5-47.



Technet24

Figure 5-47 *Entra ID Application Permissions Configuration*

Multiple permissions are required for different functionality. These permissions are explained in Table 5-2.

Table 5-2 *Required Application Permissions*

Workflow or Functionality	Required Permissions
Group retrieval (from group tab)	Group.Read.All: Needed to retrieve group information.
EAP-TLS—user authorization	User.Read.All: Needed only if user authorization with user attribute retrieval is required. If you are using only groups, this permission is not needed.
EAP-TLS—device authorization	Device.Read.All: Needed only if device authorization with device attribute retrieval is required. If you are using only groups, this permission is not needed.
TEAP (user and/or device authorization)	User.Read.All and/or Device.Read.All: Needed only if user and/or device authorization with user attribute and/or device retrieval is required. If you are using only groups, this permission is not needed.
TACACS+	User.Read: Needed to retrieve basic user information.

After you add the required permissions, click **Grant Admin Consent For**, as shown in Figure 5-48.

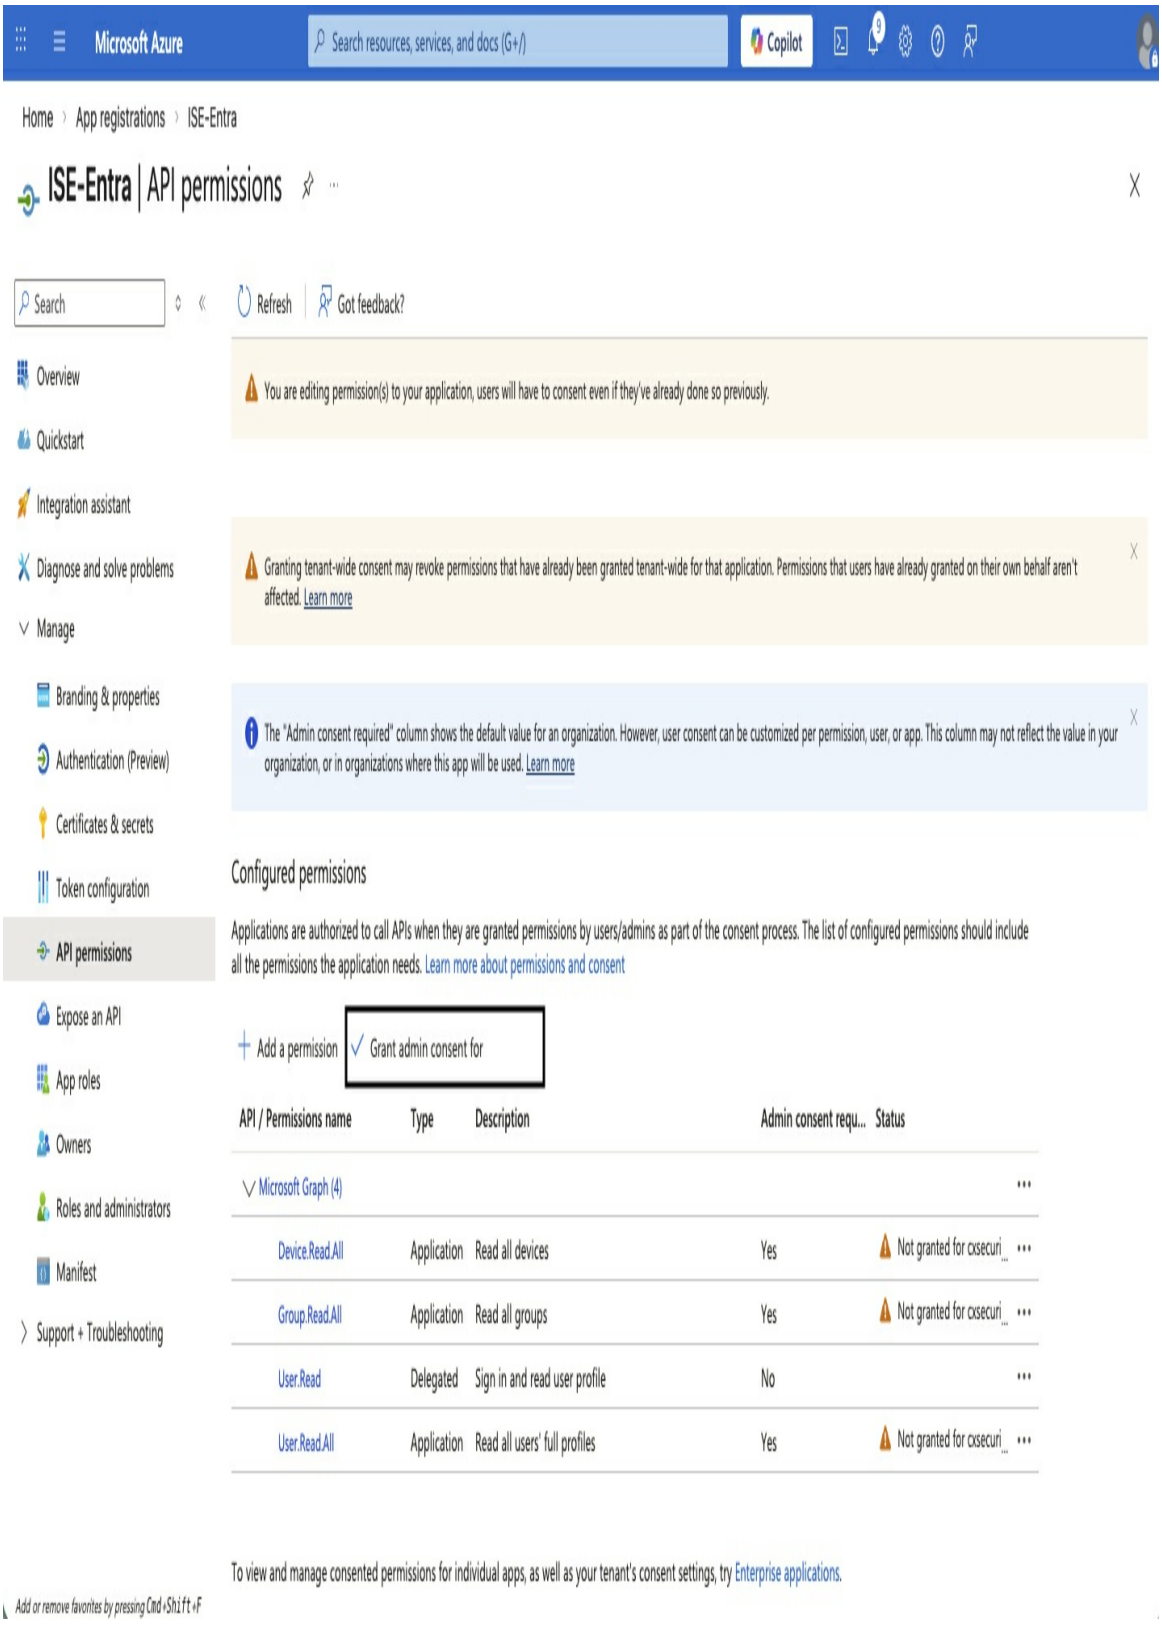


Figure 5-48 *Entra ID App Admin Consent Configuration*

Technet24

This step concludes the configuration of the Entra Application. After you complete this configuration, you can populate the ISE settings and use the **Test** button to verify connectivity.

You can then navigate to the **Groups** and **Attributes** tabs of the REST ID Store and download them so that you can use them in authorization policies. A few downloaded groups are shown in Figure 5-49.

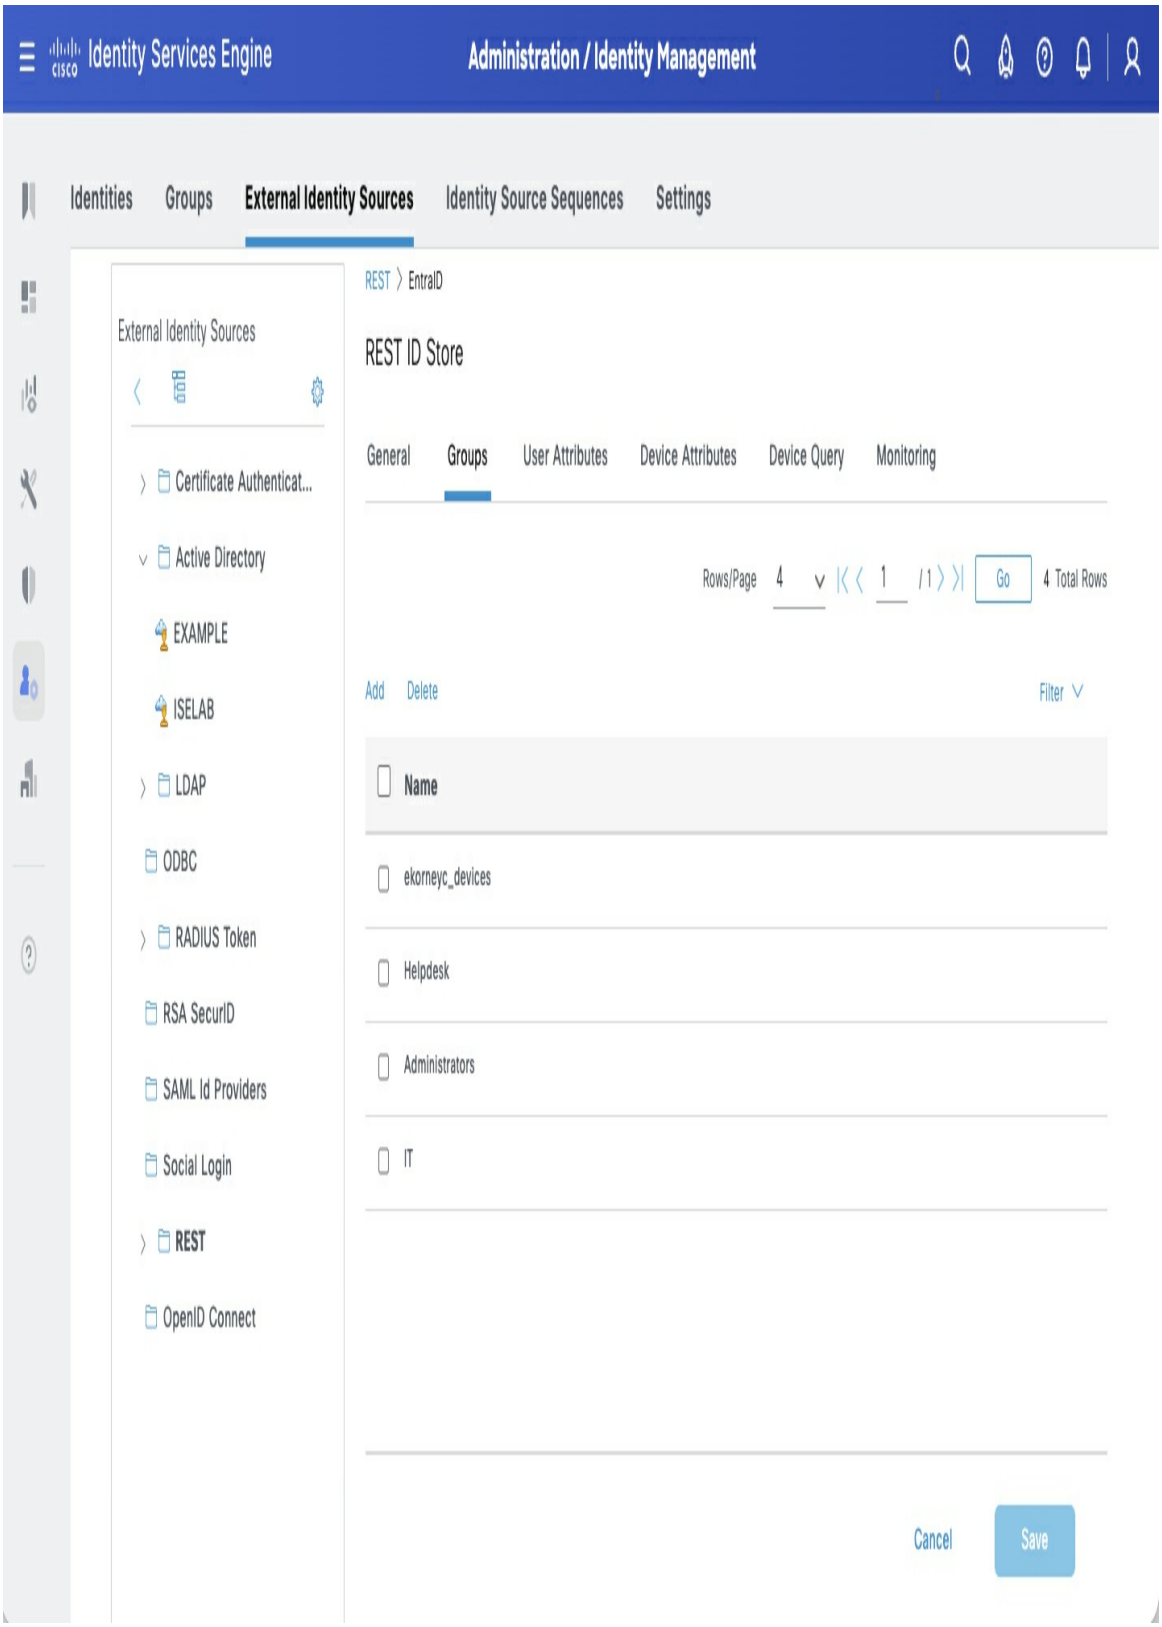


Figure 5-49 *REST ID Store Groups Configuration*

Technet24

To perform device authorization, navigate to the **Device Query** tab. You need to enable **Device Query** and configure the **Device Query String**. This setting defines how ISE builds queries to Entra ID and which attributes it uses to look up devices. You can select either the **Entra Device ID** or the **Entra Device Name**, as shown in Figure 5-50.

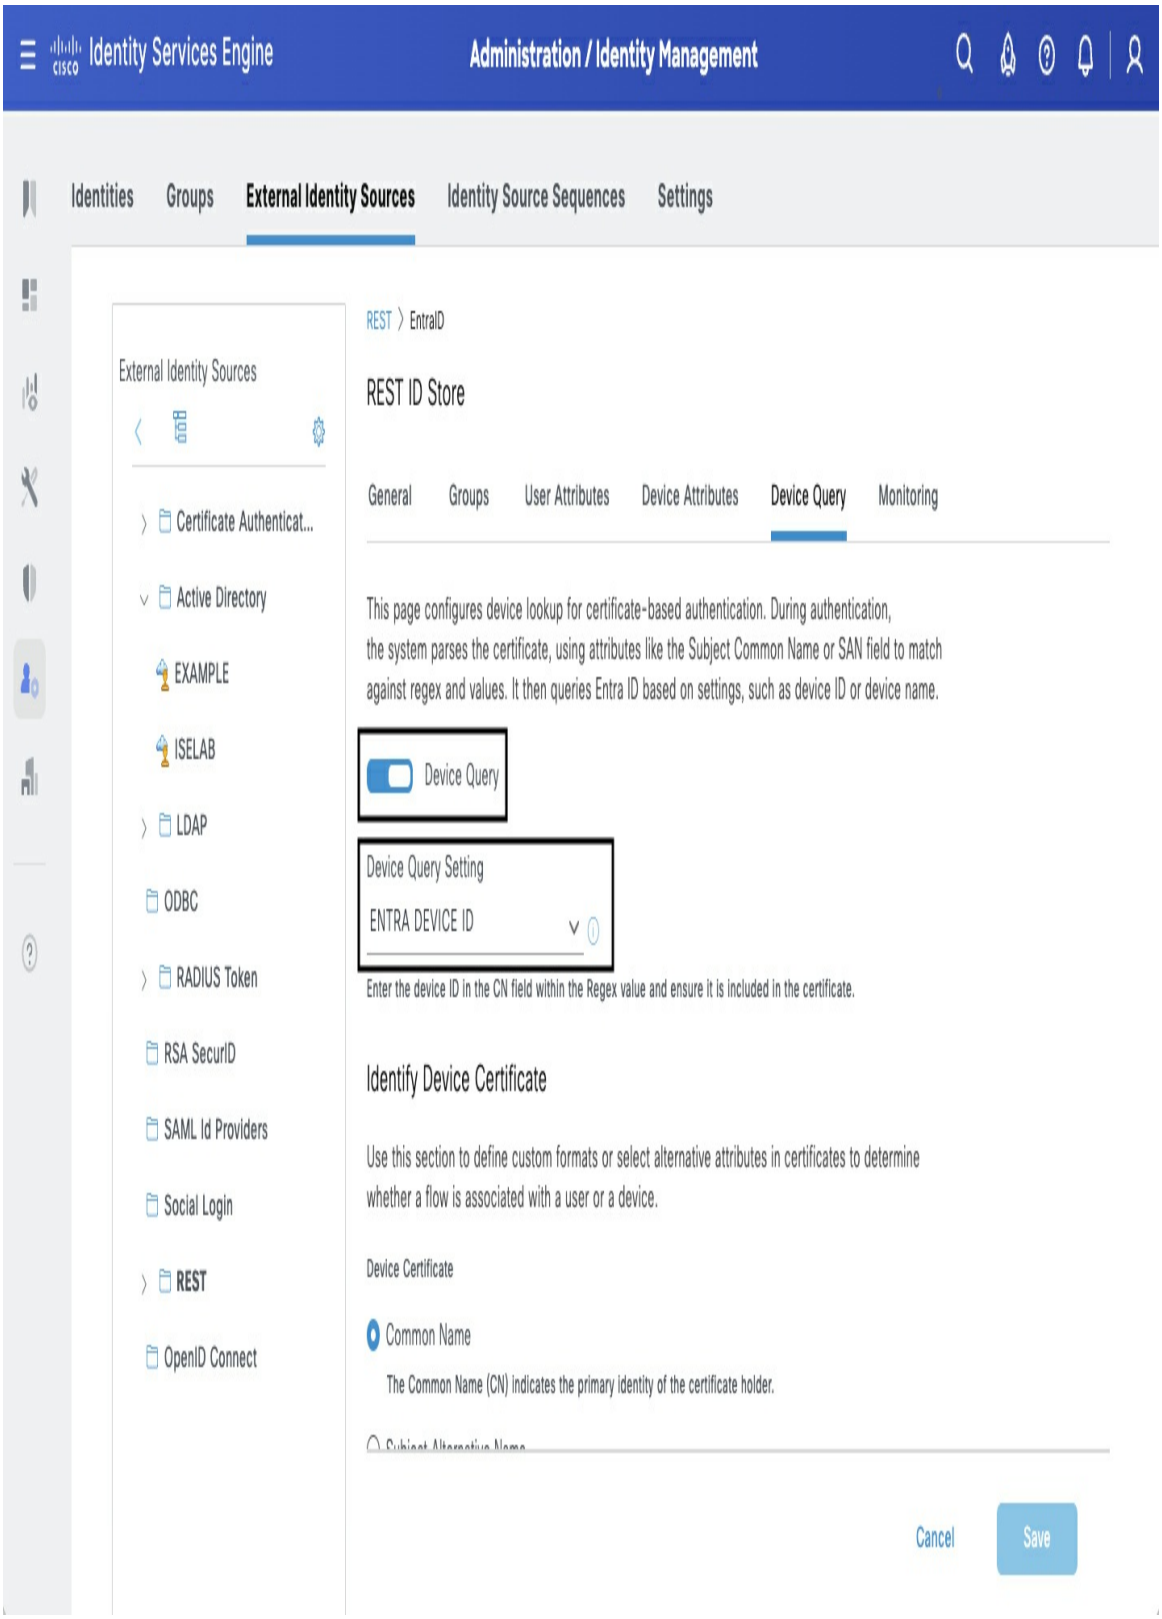


Figure 5-50 *REST ID Store Device Query Configuration*

Technet24

Next, you need to help ISE determine whether the certificate presented during authentication is a device certificate rather than a user certificate. First, select the attribute to inspect. You can use either the **Common Name** or the **Subject Alternative Name**, as shown in Figure 5-51. Finally, select the Regex that confirms the subject name format and use the **Validate** button to verify it.

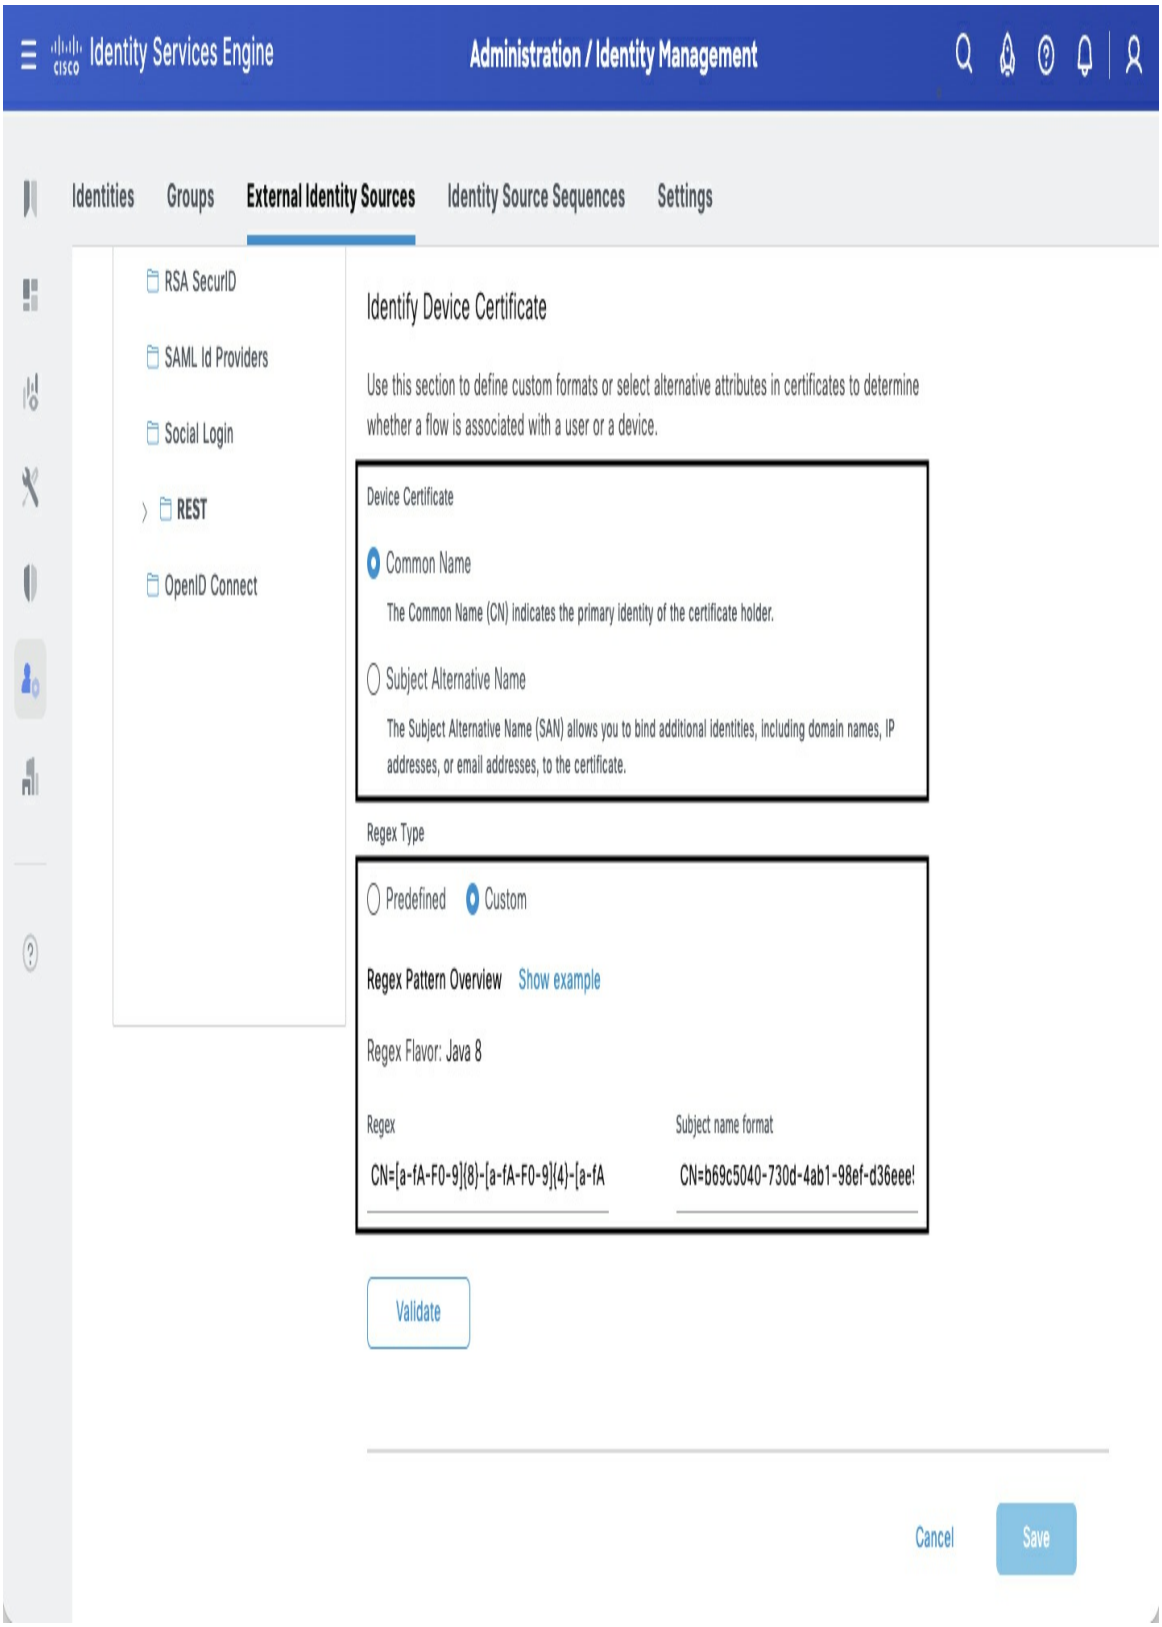


Figure 5-51 *REST ID Store Device Query Certificate Configuration*

Technet24

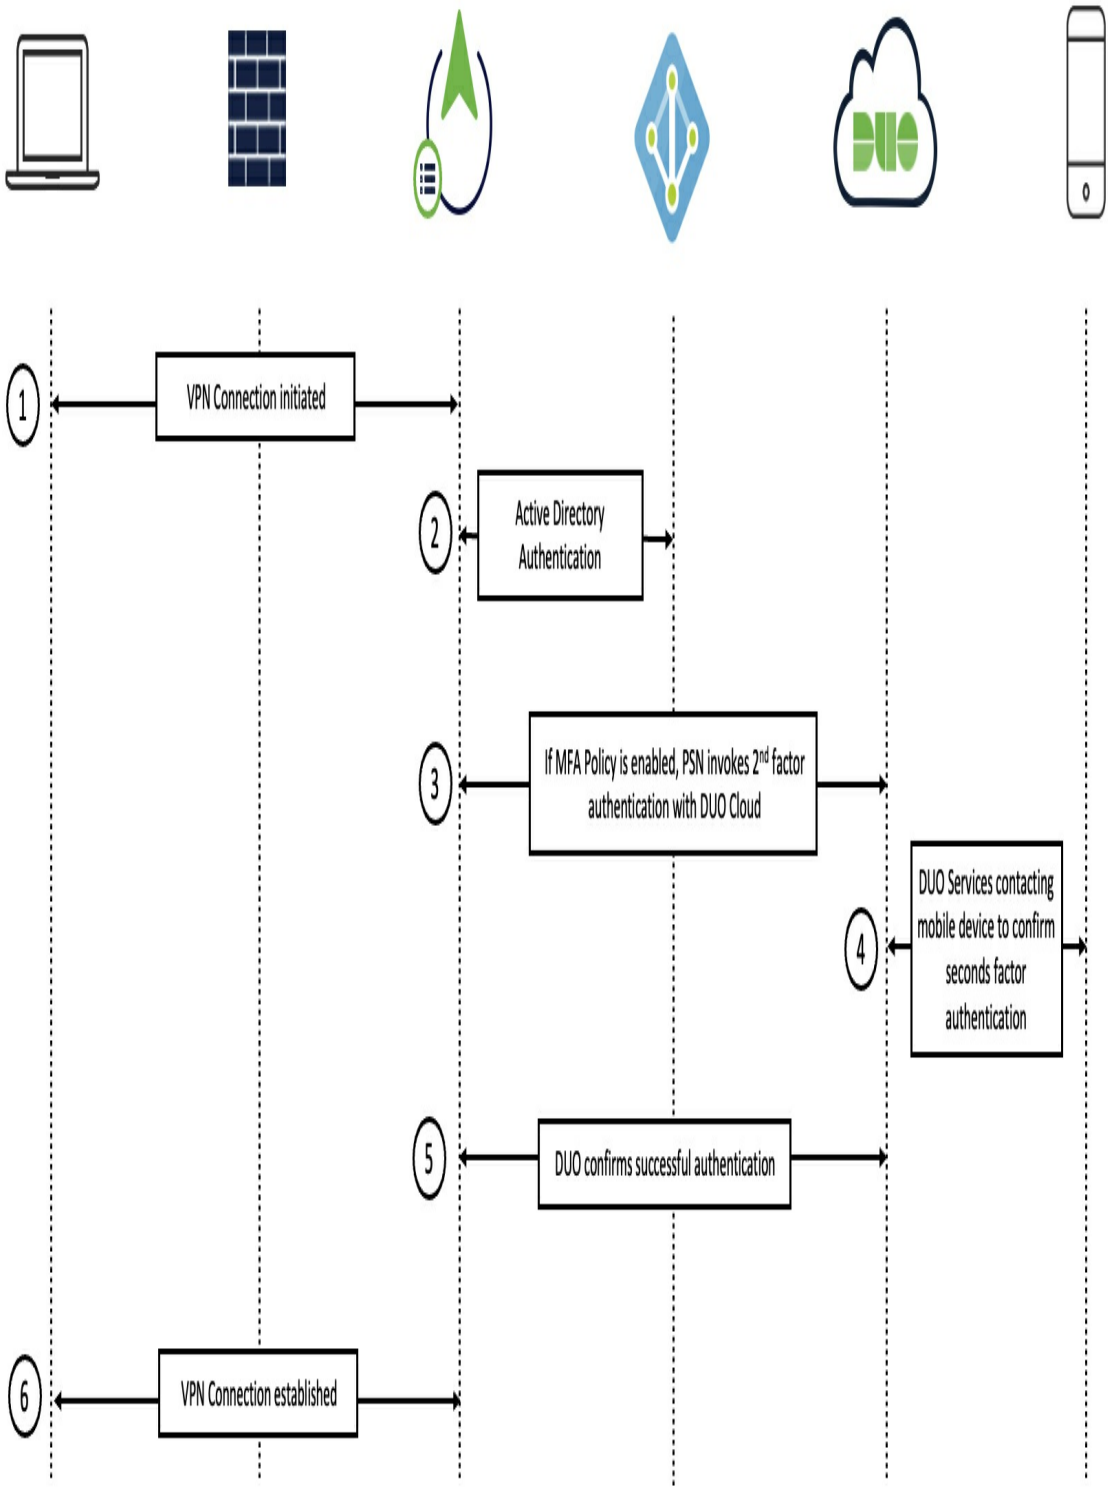
Duo

Starting with ISE 3.3 patch 1, you can integrate ISE natively with Duo for multifactor authentication (MFA). Native integration allows you to bypass the Duo Authentication Proxy service and send API calls for the second factor directly to the Duo cloud. This integration is supported for two use cases:

- VPN authentication
- Device administration

When you enable MFA on ISE, policy sets include an additional element called the MFA policy. The MFA policy sits between authentication and authorization and allows authorization to proceed only after MFA succeeds.

The complete flow is shown in Figure 5-52. In this example, Active Directory is used as the primary authentication method, and Duo is used as the second factor.


1
VPN Connection initiated
2
Active Directory Authentication
3
If MFA Policy is enabled, PSN invokes 2nd factor authentication with DUO Cloud
4
DUO Services contacting mobile device to confirm seconds factor authentication
5
DUO confirms successful authentication
6
VPN Connection established

Technet24

Figure 5-52 *VPN Authentication Flow with ISE and Duo as a Second Factor Authentication*

MFA is disabled on ISE by default. To configure the integration, you must first enable MFA by navigating to **Administration > Identity Management > Settings**, as shown in Figure 5-53.

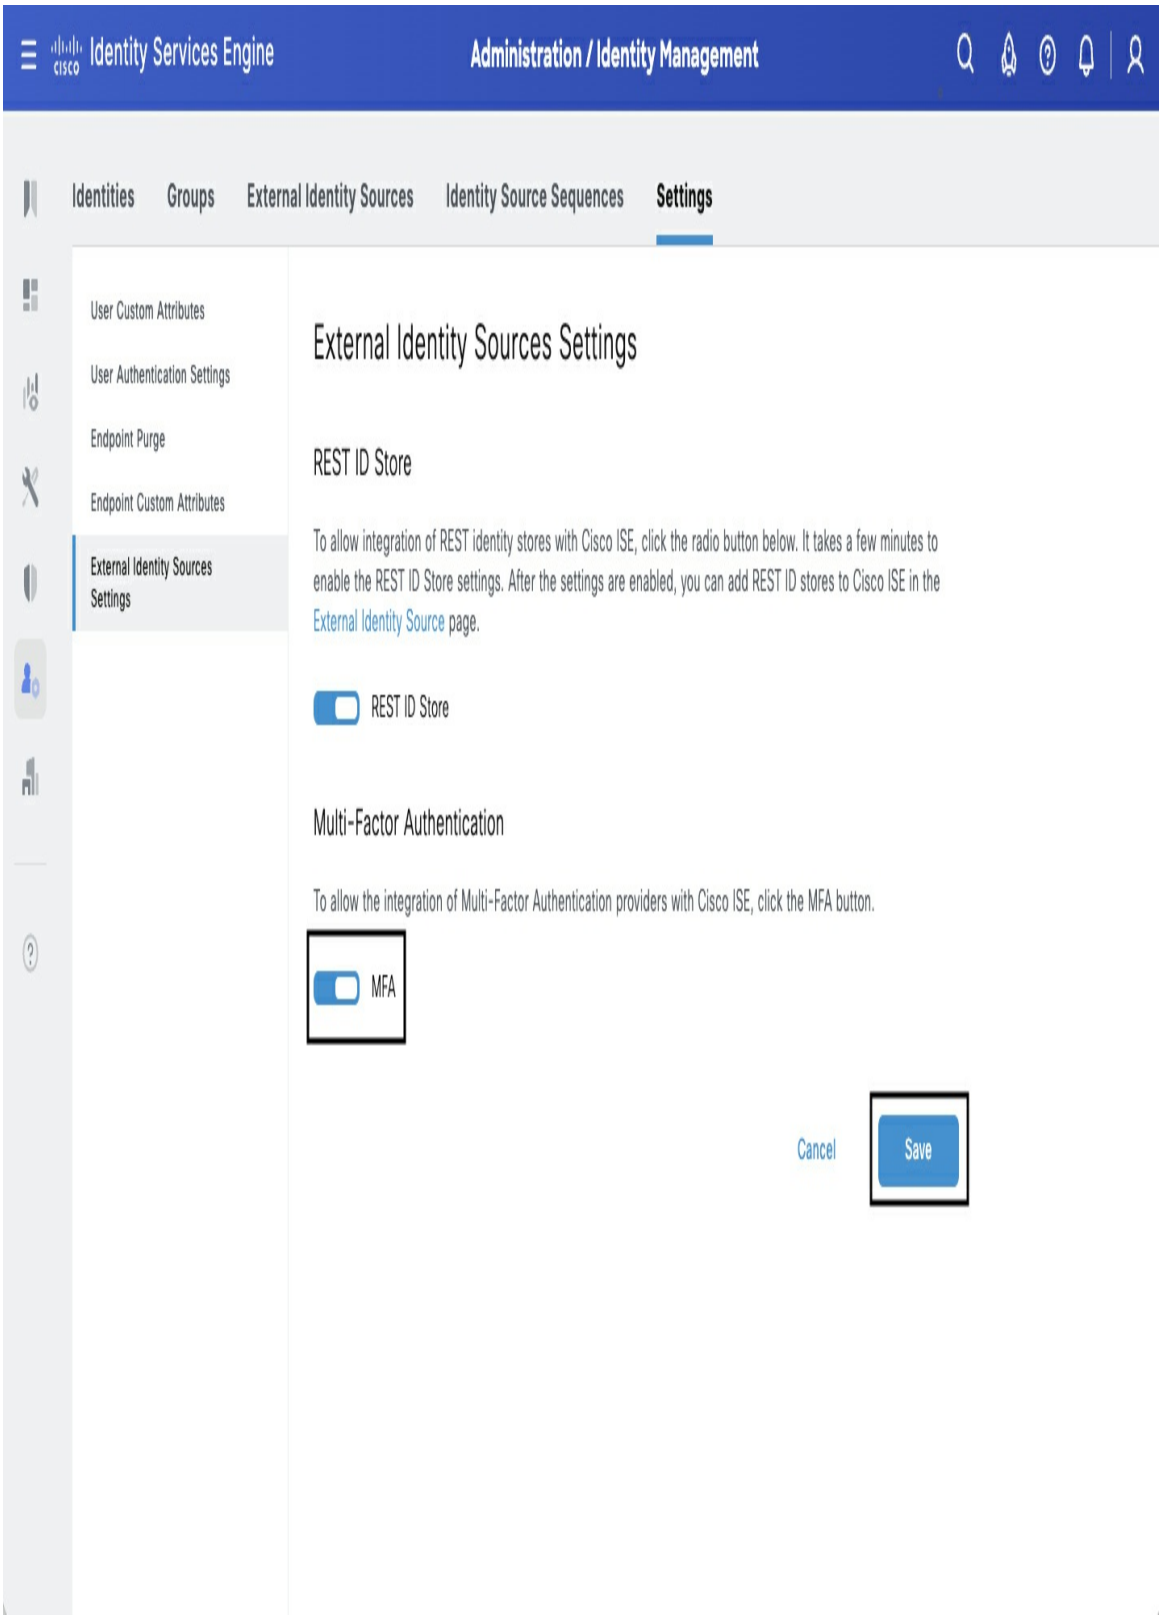


Figure 5-53 *ISE External Identity Sources Settings*

Technet24

After you save the setting, a new external identity store named MFA becomes available. To configure the integration, navigate to **Administration > External Identity Sources > MFA** and click **Add**. ISE presents a configuration wizard. After you click **Let's Do It**, provide a name for the identity store and click **Next**, as shown in Figure 5-54.

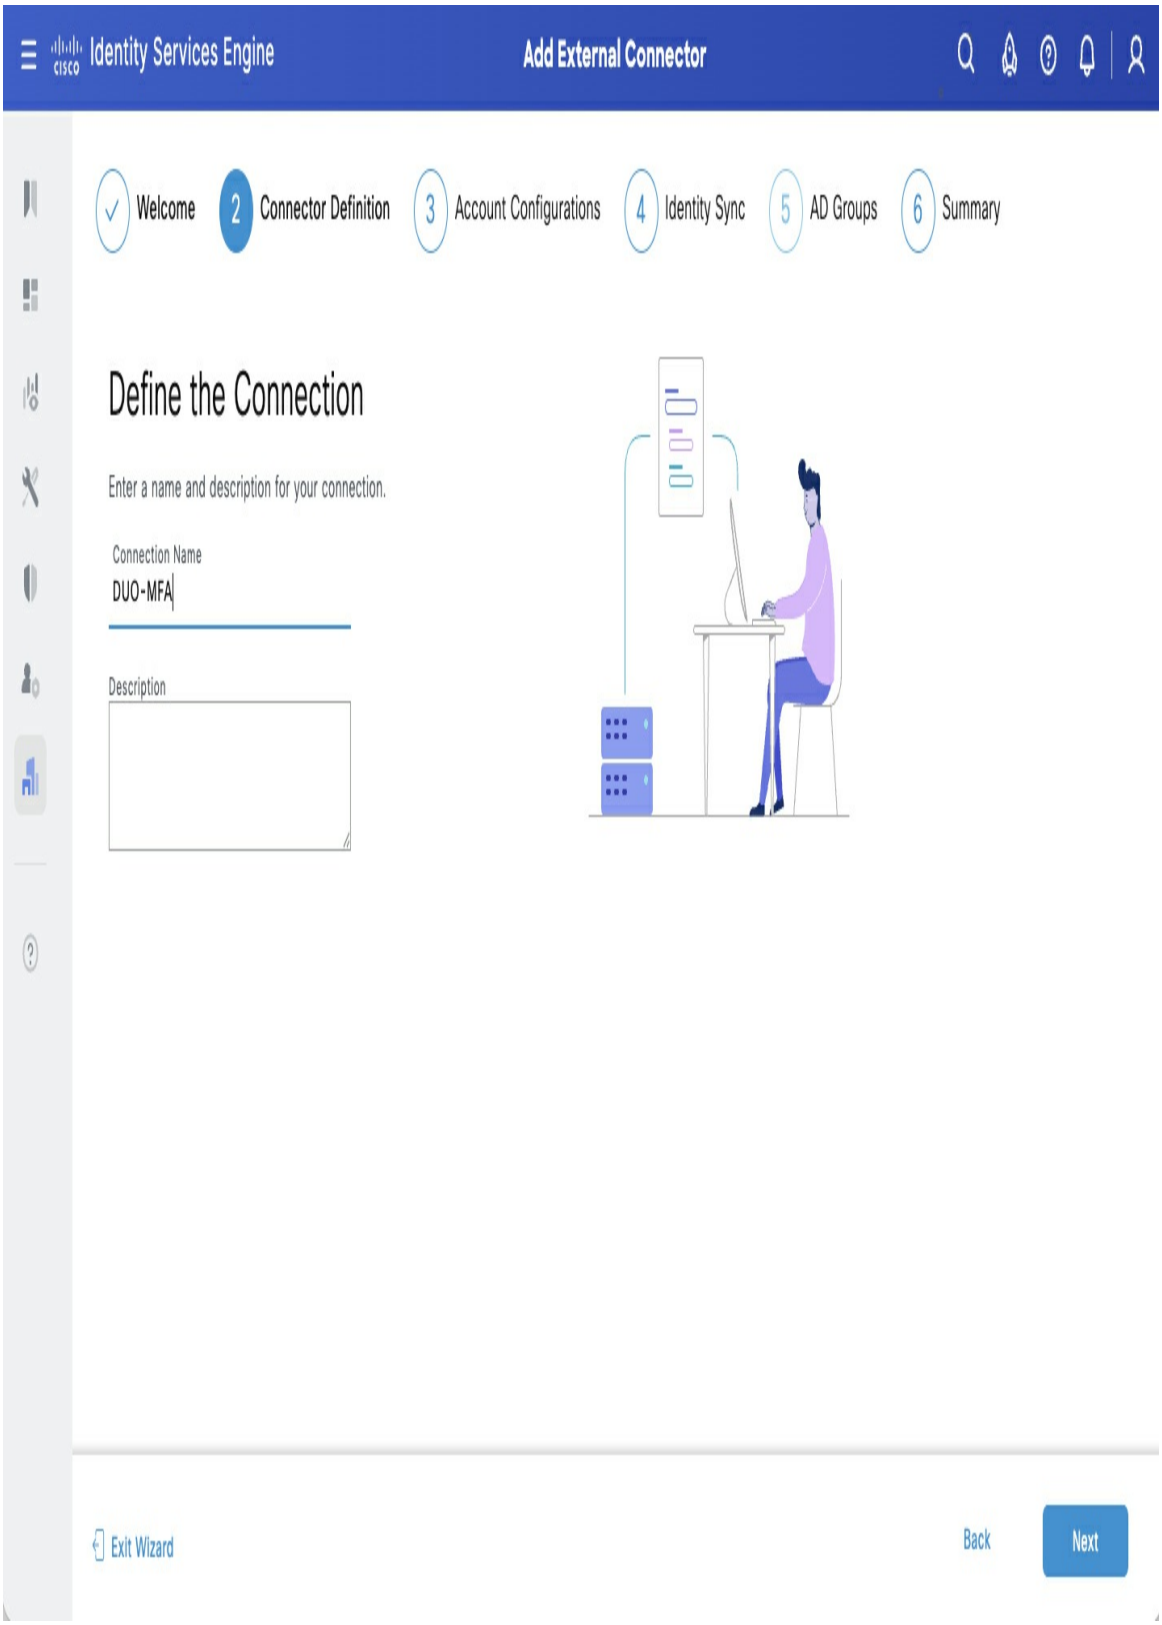


Figure 5-54 *ISE MFA Connector Definition*

Technet24

At this point, you need to populate the following parameters, as shown in Figure 5-55:

- **API Hostname**
- **ISE Admin API iKey**
- **ISE Admin API sKey**
- **ISE Authentication API iKey**
- **ISE Authentication API sKey**

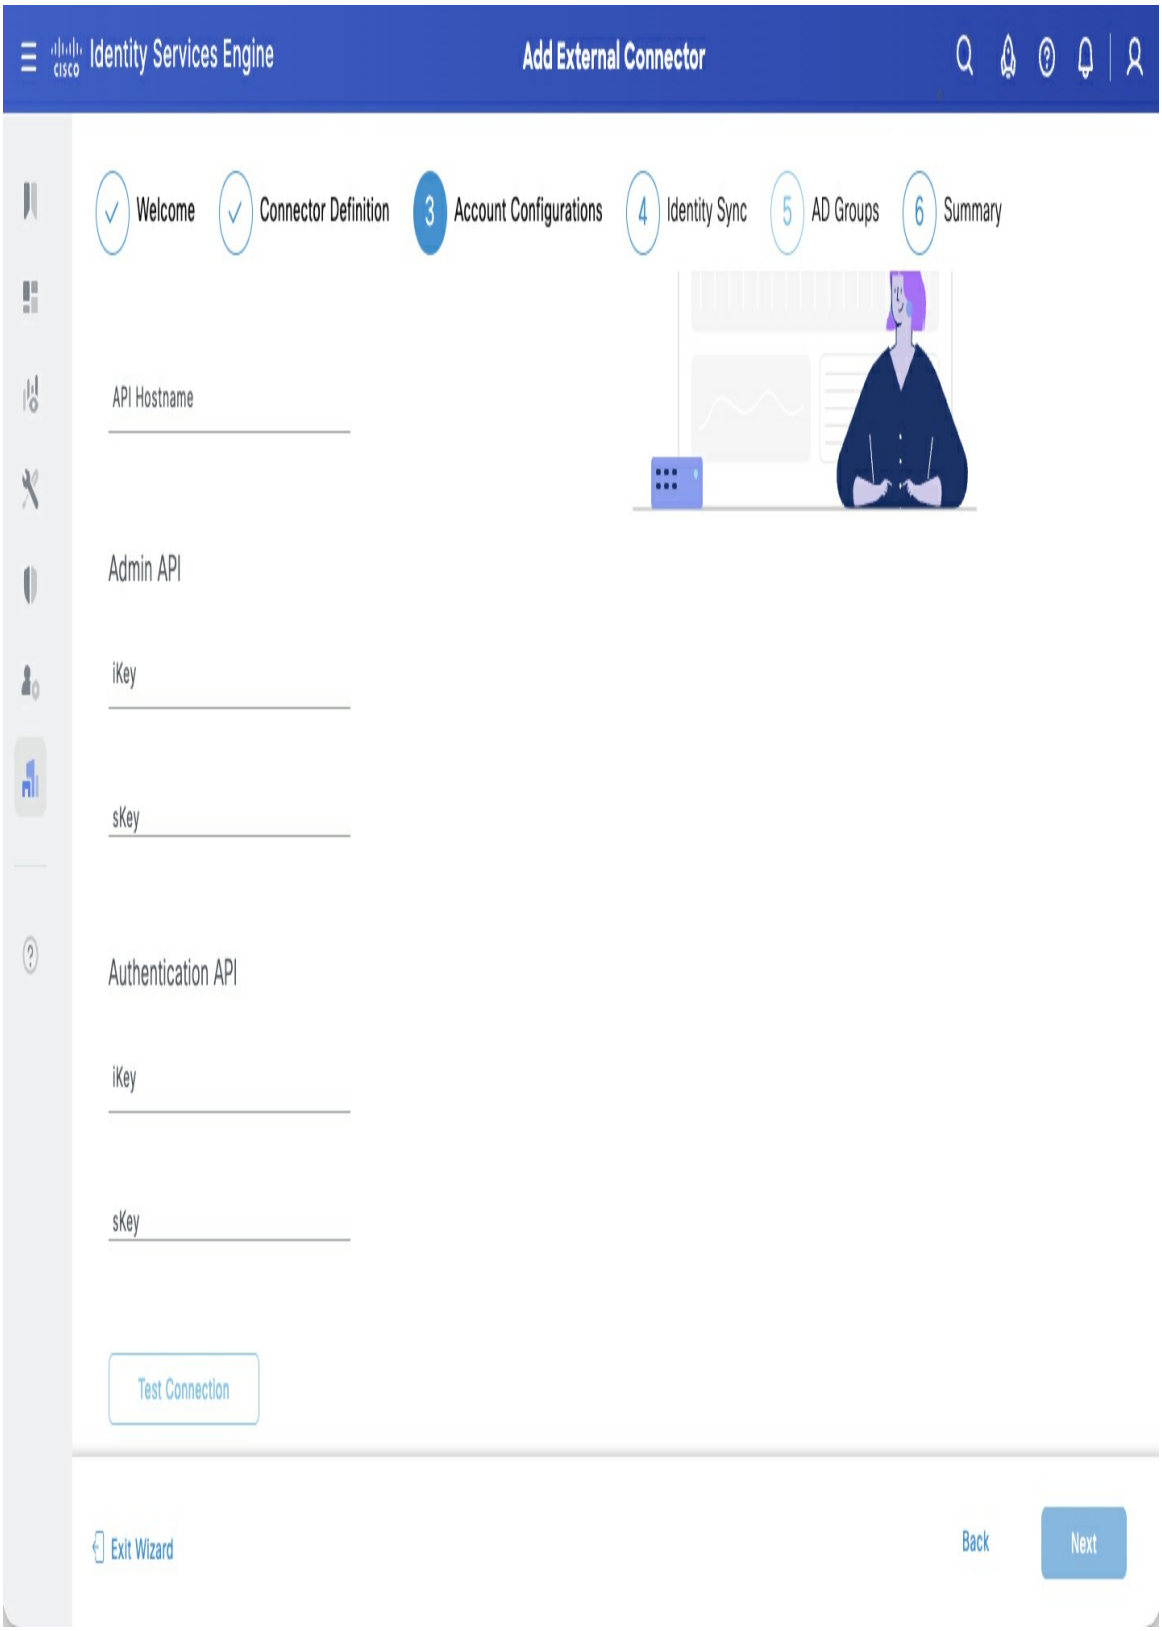


Figure 5-55 *ISE MFA Account Configurations*

Technet24

From the Duo Admin Portal, navigate to **Applications** and click **+Add Application**. Locate **ISE Admin API** and click **+Add**, as shown in Figure 5-56.

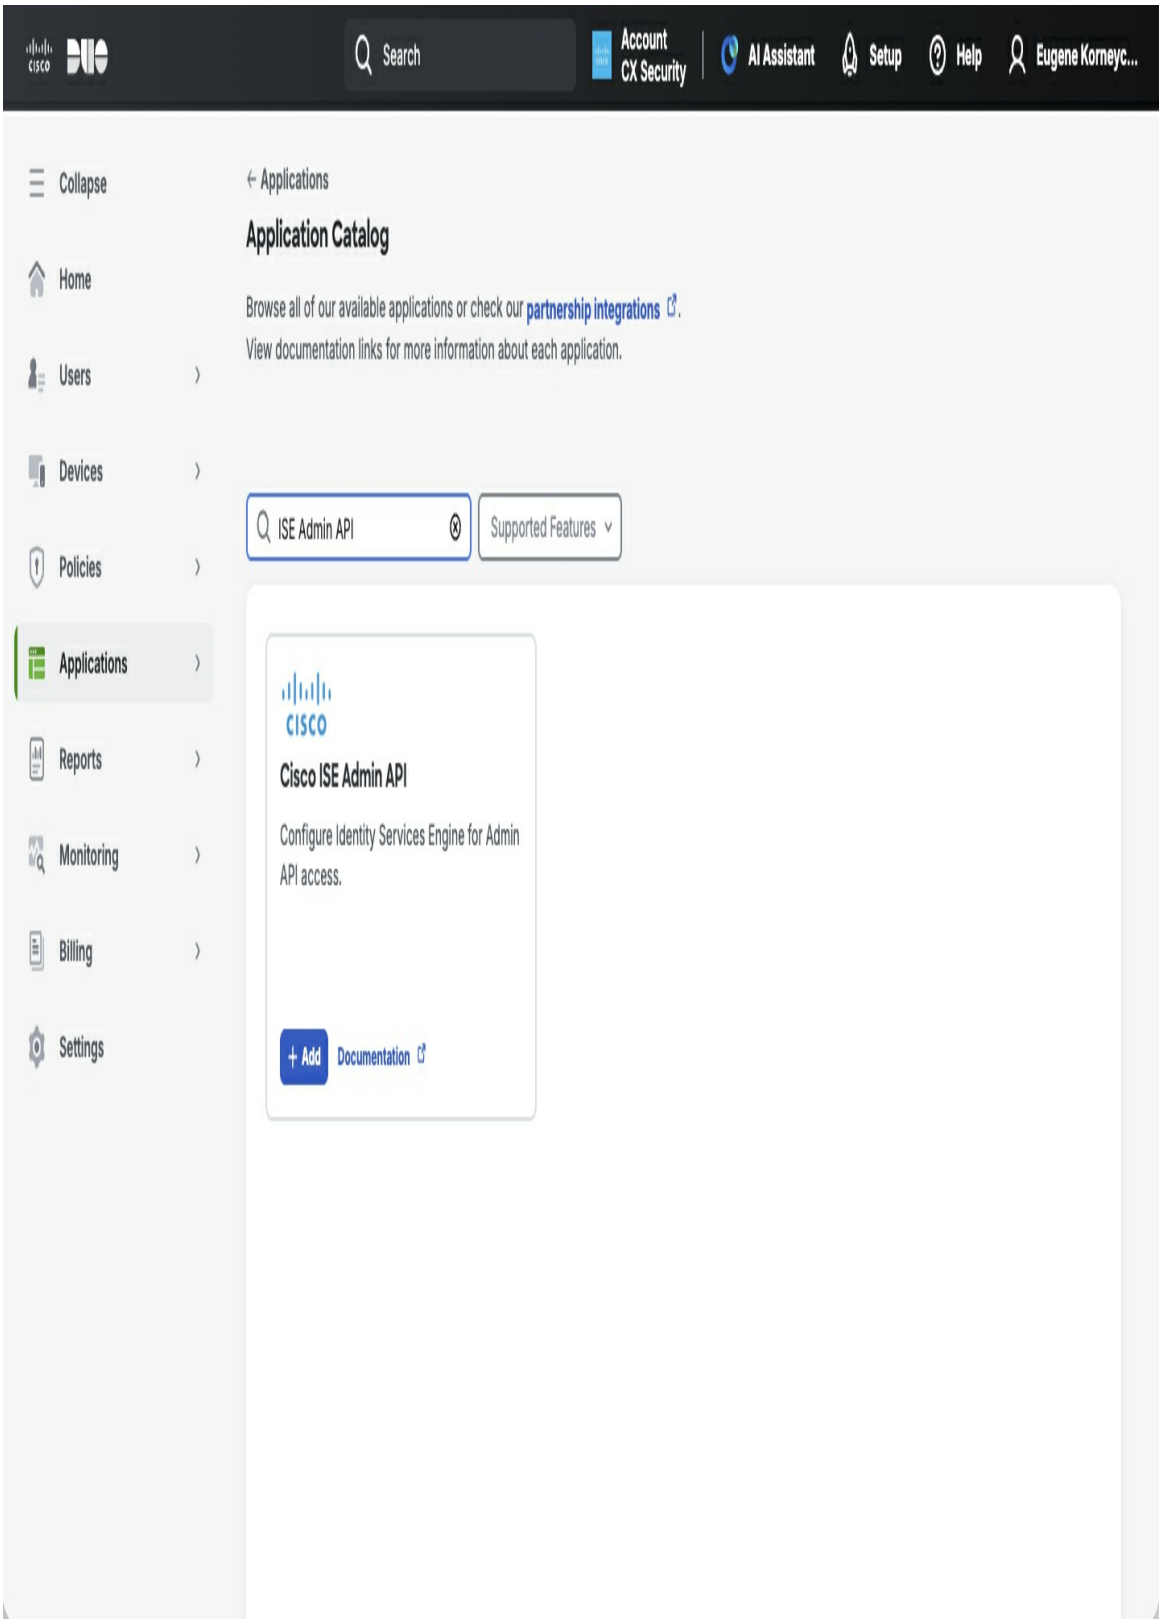


Figure 5-56 *Duo ISE Admin API Configuration*

Technet24

Make a note of the ISE Admin API **Integration Key**, **Secret Key**, and **API Hostname**, as shown in Figure 5-57. You use these parameters to configure ISE. Admin API is used to sync identities from Active Directory to Duo Cloud.

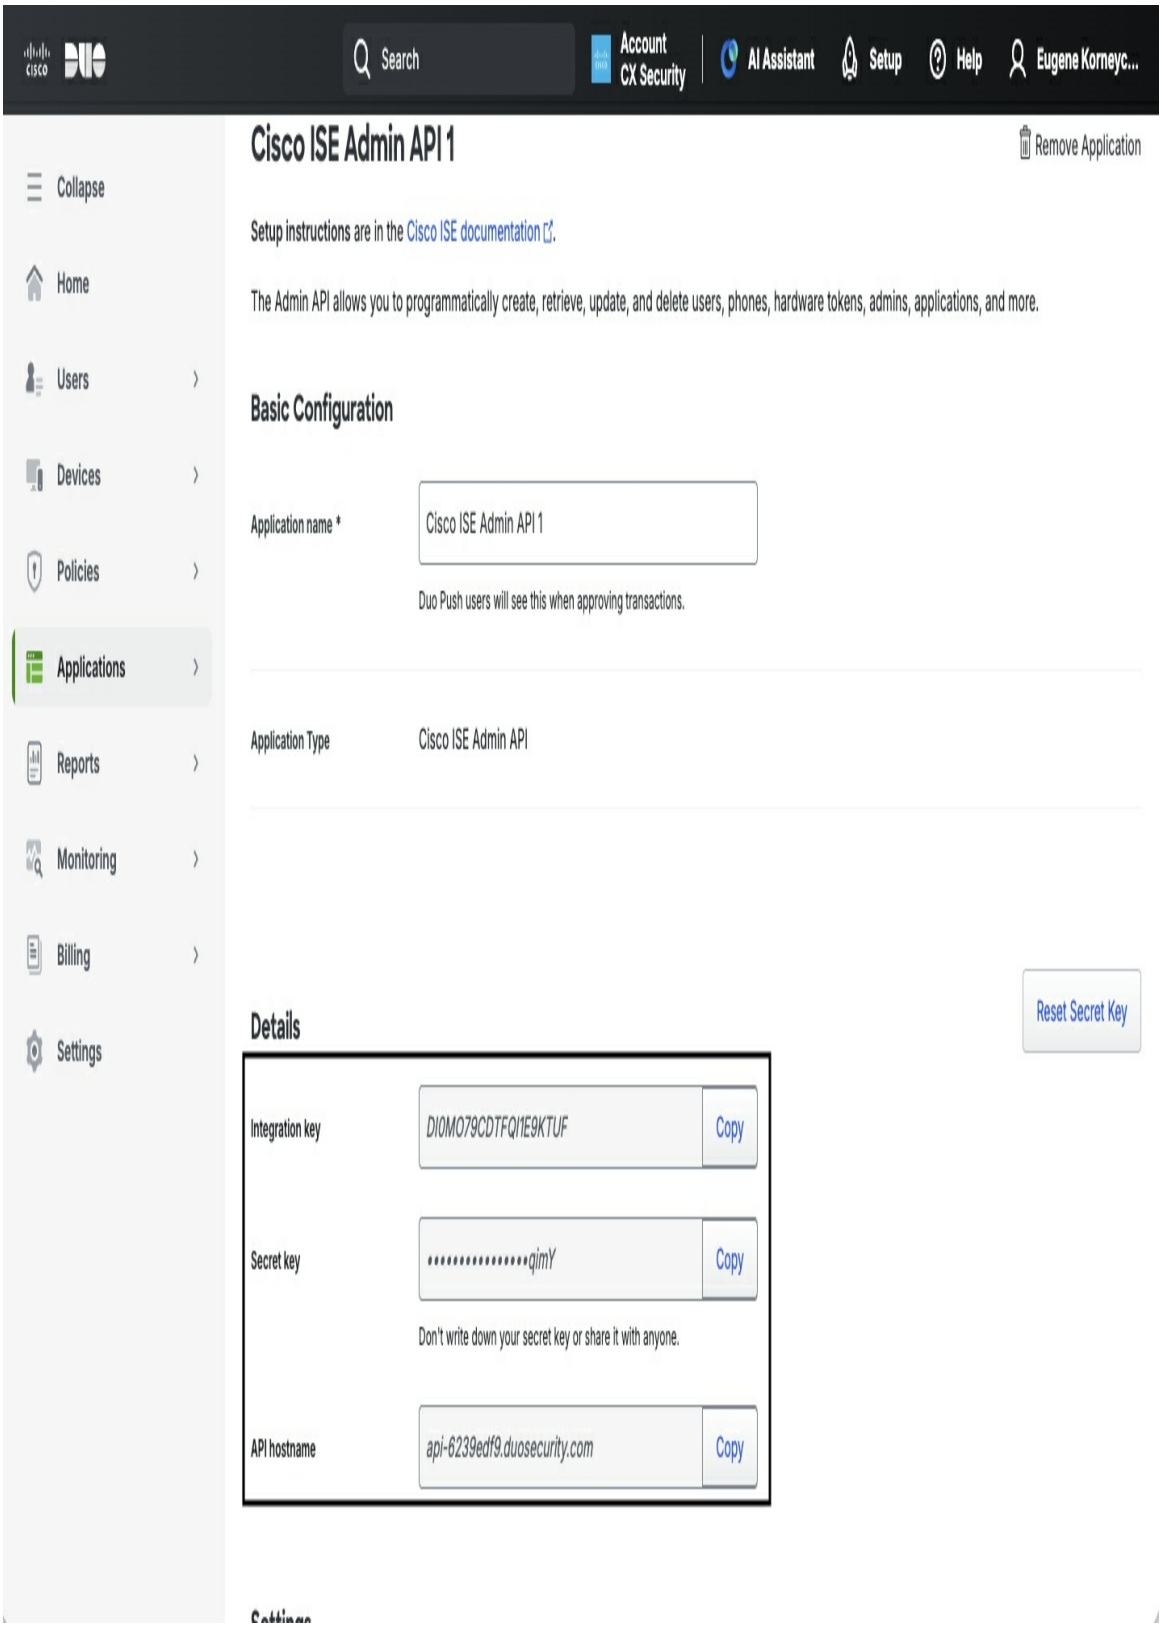


Figure 5-57 *Duo ISE Admin API Configuration Details*

Technet24

Scroll down to **Permissions** and grant resource access with **Read** and **Write** permissions, as shown in Figure 5-58. Save the changes.

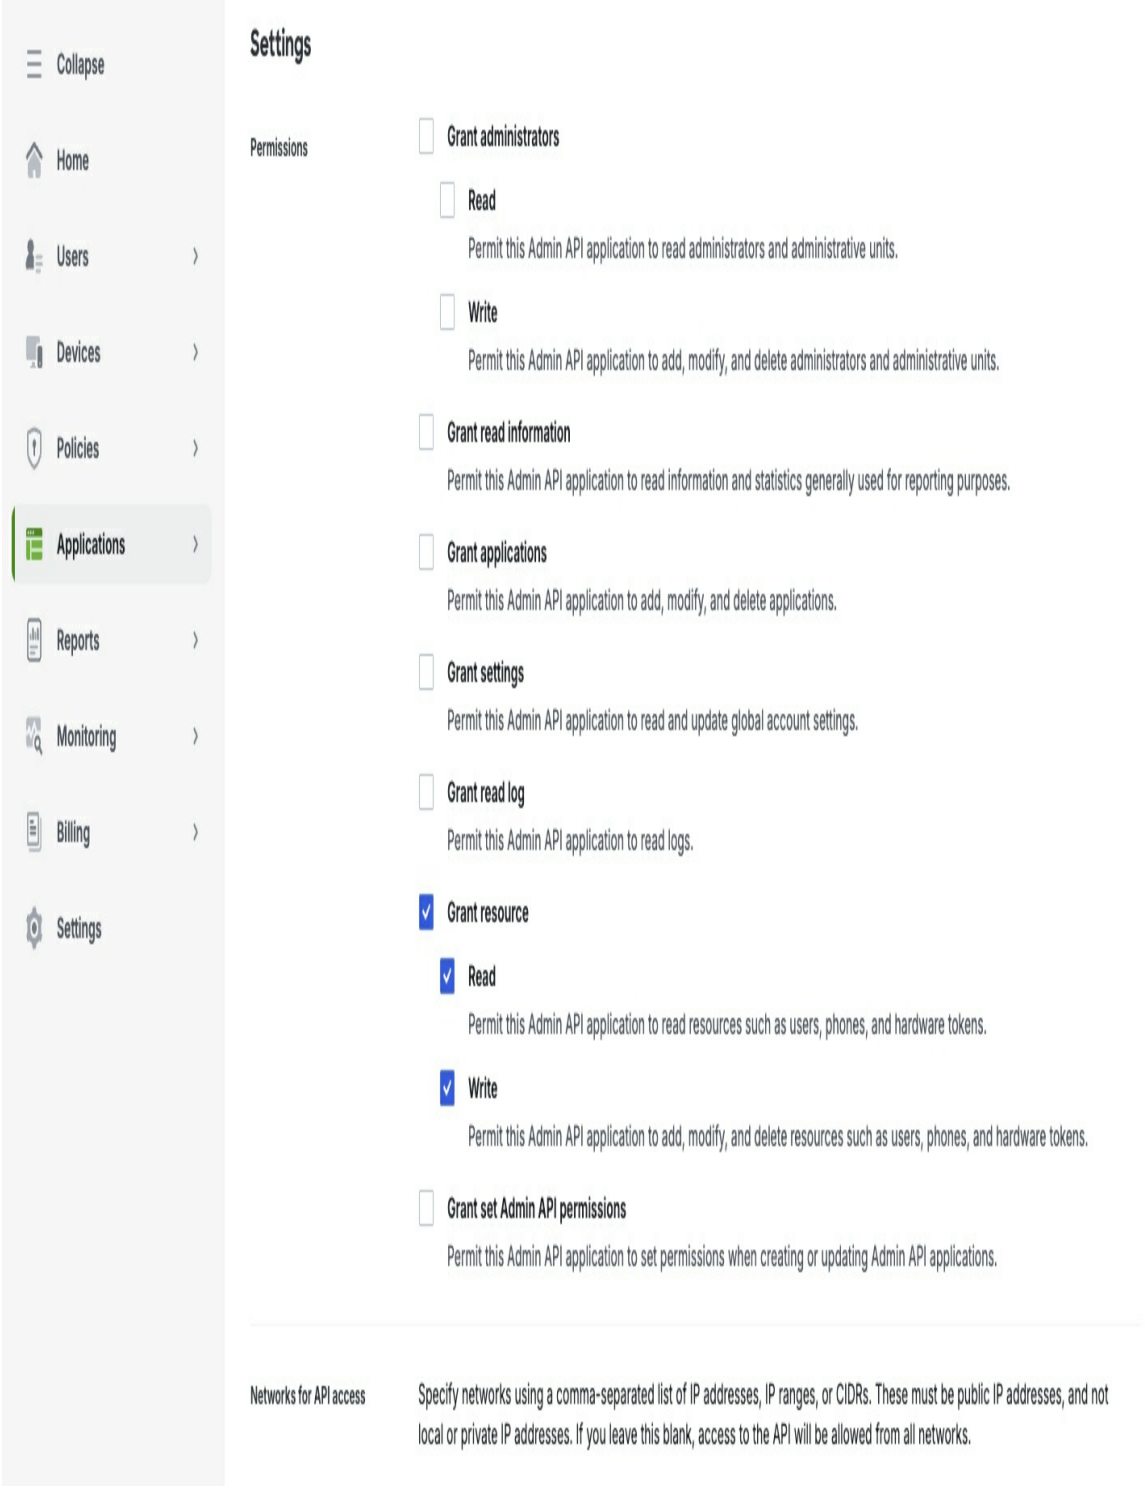


Figure 5-58 *Duo ISE Admin API Permissions*

Navigate back to the Duo Admin Portal, select **Applications**, and click **+Add Application**. Locate **ISE Auth API** and click **+Add**, as shown in Figure 5-59. ISE Auth API is used during authentication to verify second-factor authentication.

Technet24

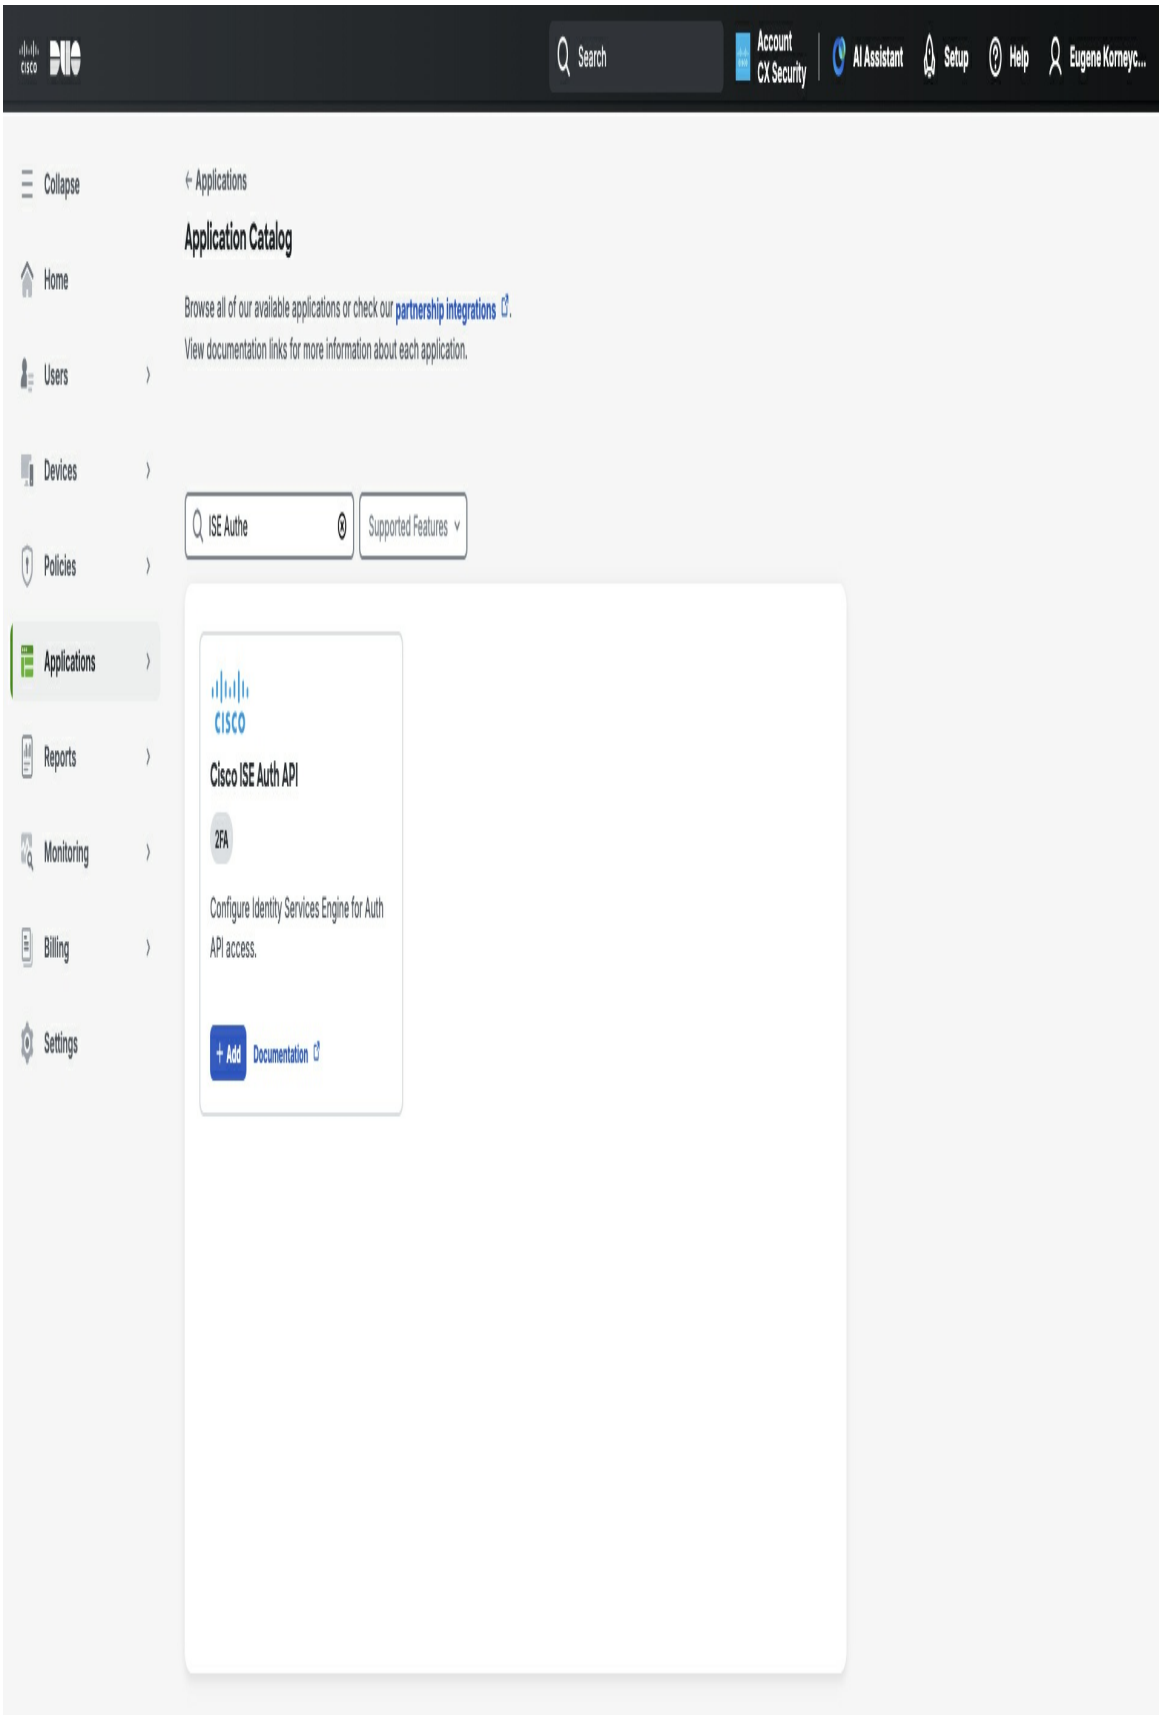

Search
Account
CX Security
AI Assistant
Setup
Help
Eugene Korneyc...
Collapse
Home
Users
Devices
Policies
Applications
Reports
Monitoring
Billing
Settings
Applications
Application Catalog
Browse all of our available applications or check our partnership integrations.
View documentation links for more information about each application.
ISE Authe
Supported Features
Cisco ISE Auth API
2FA
Configure Identity Services Engine for Auth API access.
+ Add
Documentation

Figure 5-59 *Duo ISE Auth API Configuration*

After you add the application, note the **Integration Key** and **Secret Key**, as shown in Figure 5-60.

Technet24

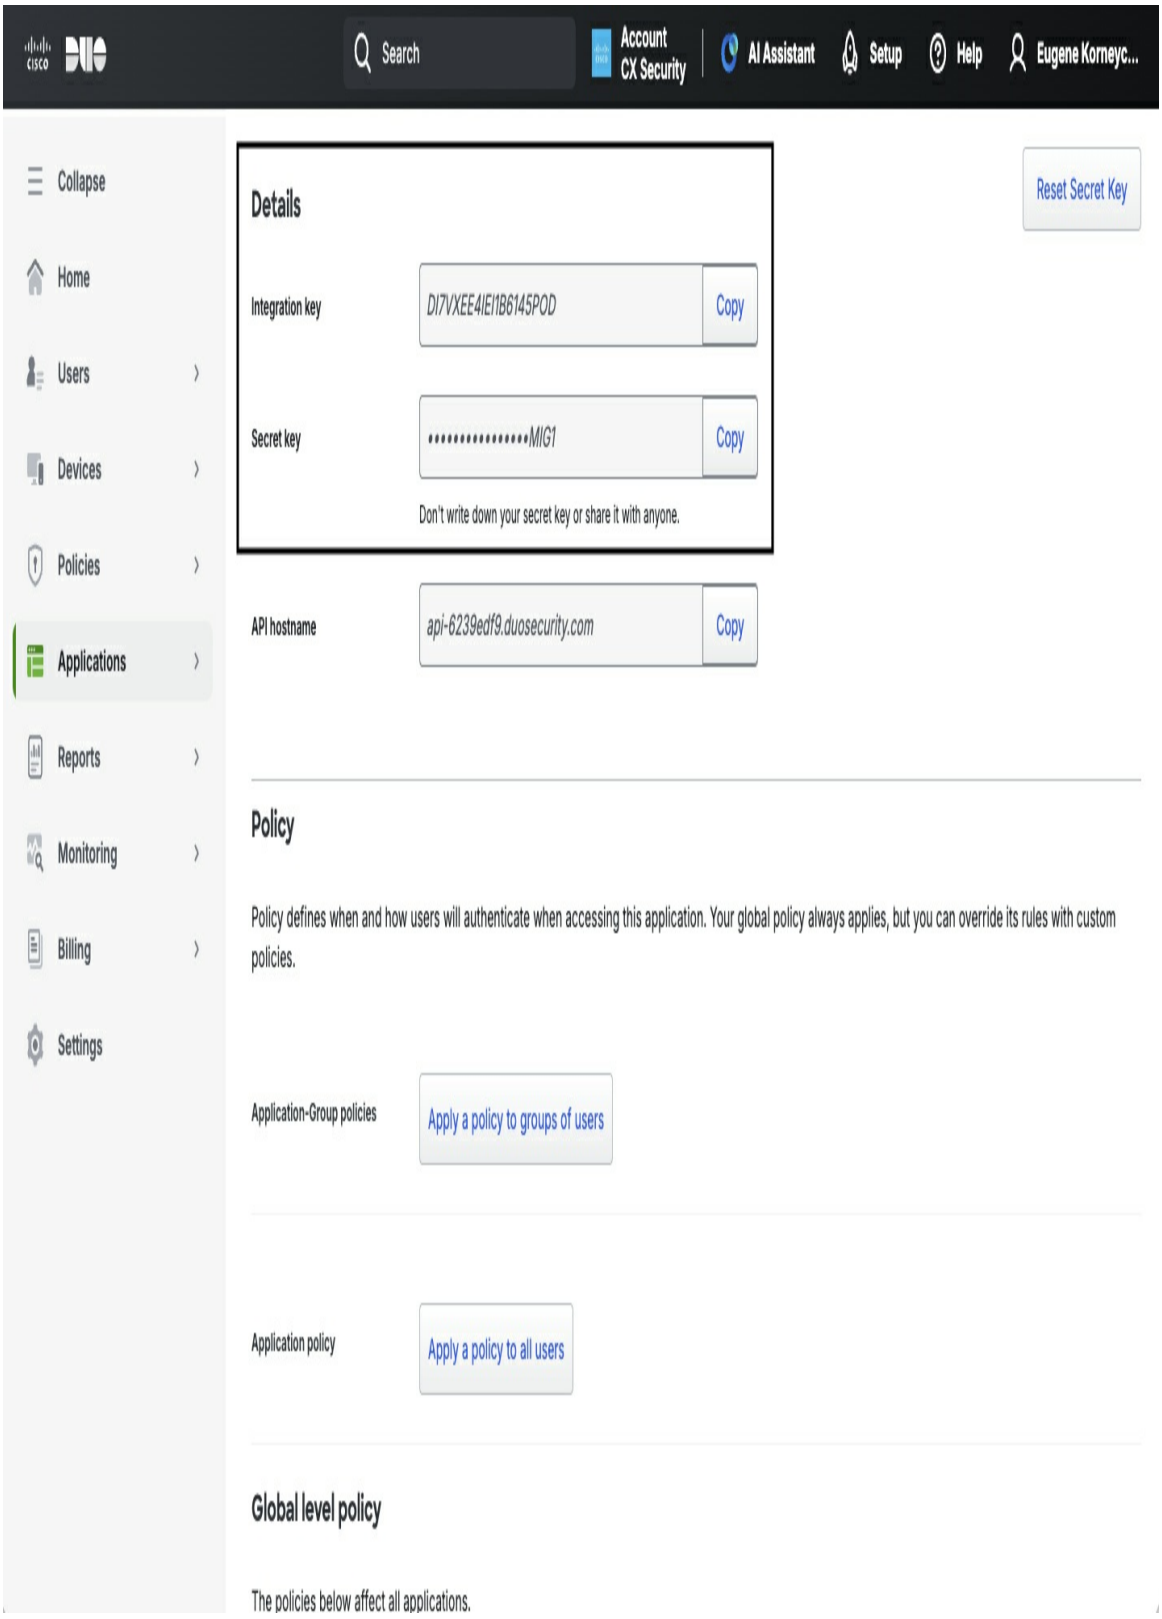


Figure 5-60 *Duo ISE Auth API Configuration Details*

After you populate the data on ISE, click **Test Connection** to verify that the MFA keys are entered correctly and that the required permissions are granted. After a successful test, as shown in Figure 5-61, click **Save**.

Technet24

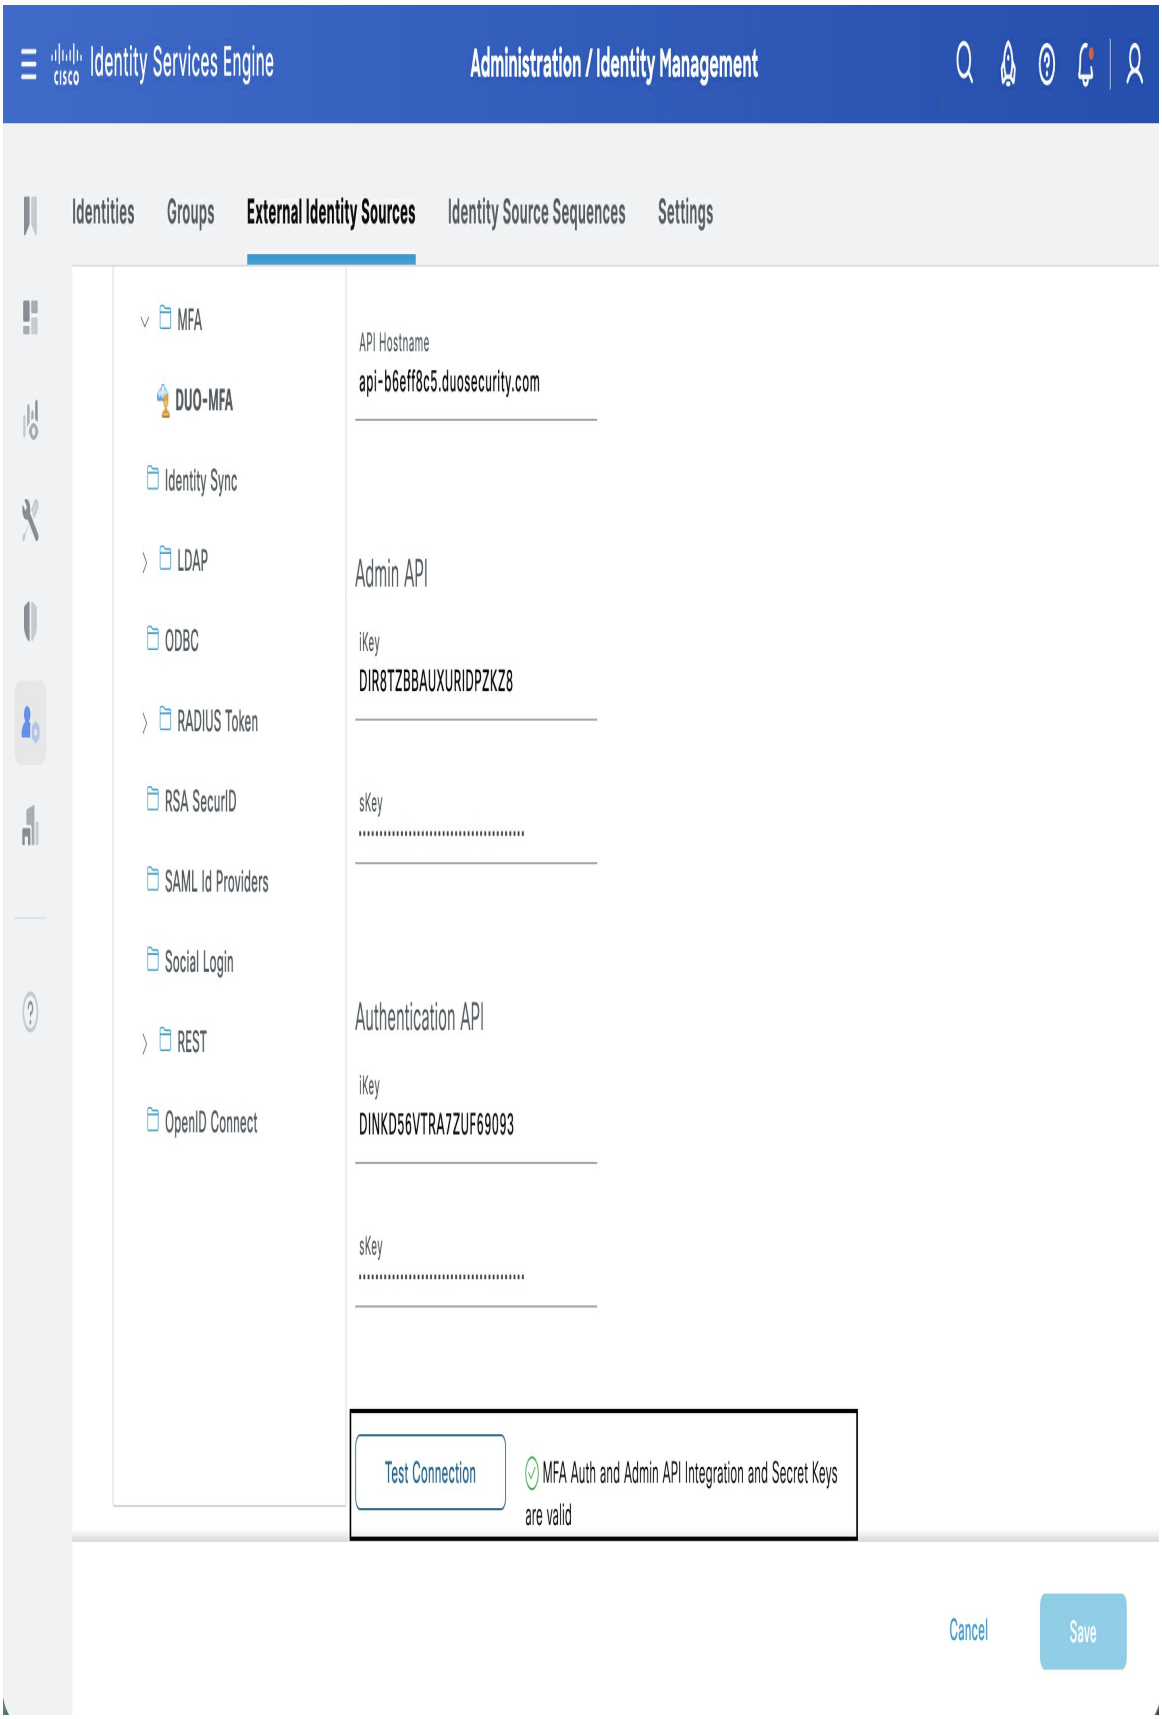

Identity Services Engine
Administration / Identity Management
Identities
Groups
External Identity Sources
Identity Source Sequences
Settings
MFA
DUO-MFA
Identity Sync
LDAP
ODBC
RADIUS Token
RSA SecurID
SAML Id Providers
Social Login
REST
OpenID Connect
API Hostname
api-b6eff8c5.duosecurity.com
Admin API
iKey
DIR8TZBBAUXURIDPZKZ8
sKey
Authentication API
iKey
DINKD56VTRA7ZUF69093
sKey
Test Connection
MFA Auth and Admin API Integration and Secret Keys are valid
Cancel
Save

Figure 5-61 *MFA Successful Test Connection*

Steps 4 and 5 of the MFA Connector Wizard allow Identity Sync, as shown in Figure 5-62. You can select the Active Directory join point, configure Active Directory groups, and synchronize users to Duo, or you can skip this step entirely.

Note

User enrollment is outside the scope of this book.

Technet24

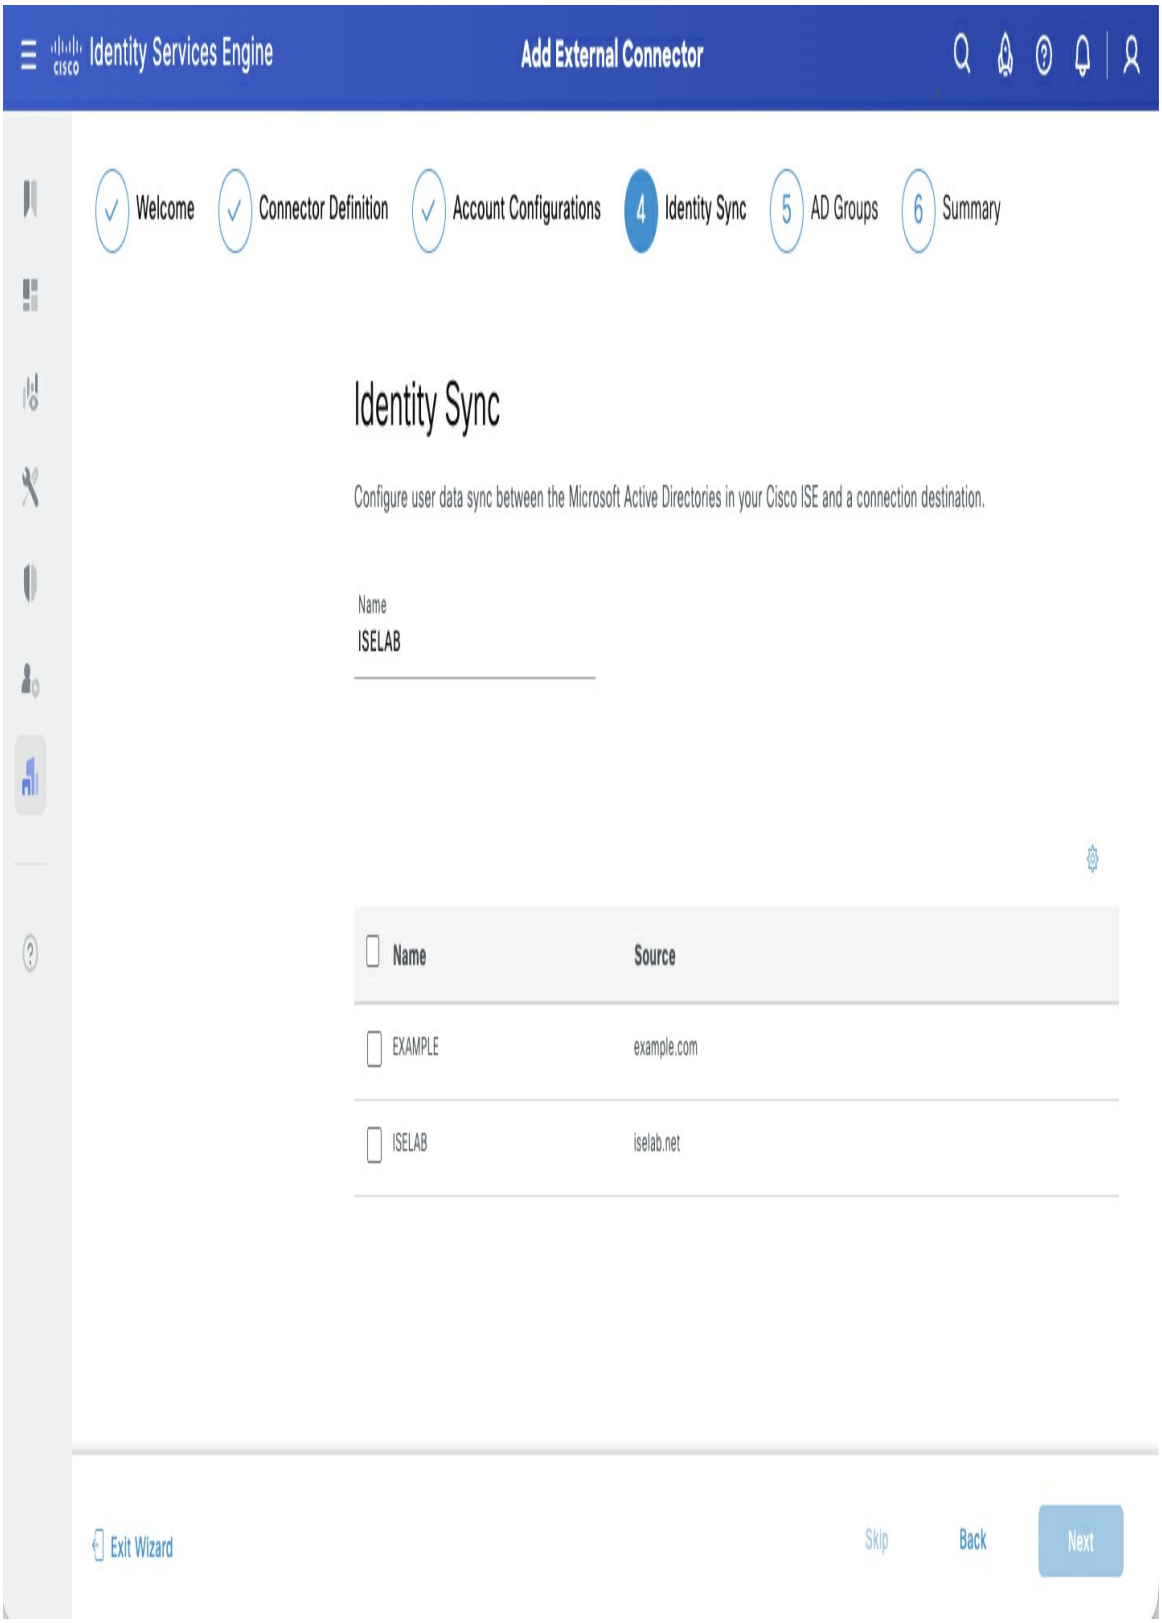


Figure 5-62 *ISE MFA Connector Identity Sync*

After you review the summary and select **Done**, the integration is complete, and you can use the new identity store in MFA policies.

Network Devices

A network access device on ISE is one of the most critical configuration elements. If you start authenticating users on ISE without configuring network access devices, ISE silently drops all Access-Request packets.

You can place network devices in different network device groups. This approach is useful when you want to differentiate network access based on device type or location. ISE includes several default parent groups, such as **All Device Types** and **All Locations**, as shown in Figure 5-63.

Technet24

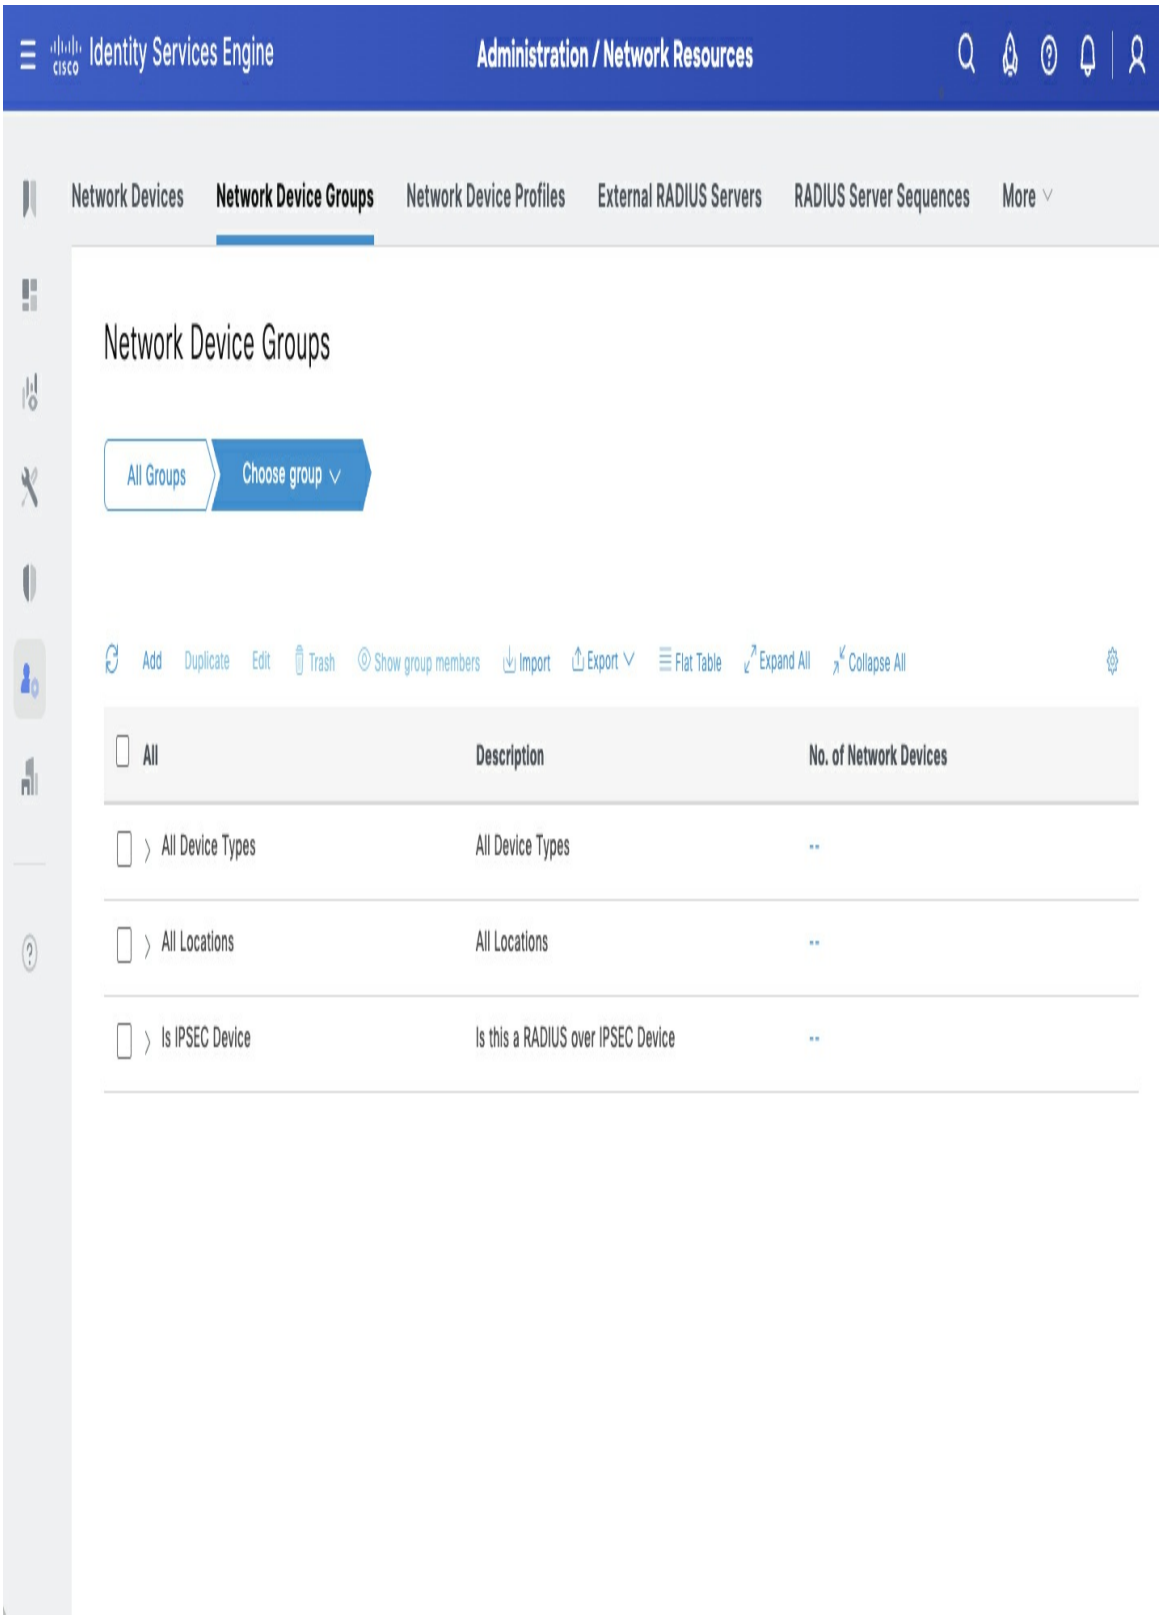


Figure 5-63 *ISE Network Device Groups Configuration*

You can create custom groups under existing groups or build a completely different group structure.

You can create network access devices by using different methods:

- Manual creation
- CSV import
- REST API

To configure a network device, navigate to **Administration > Network Resources > Network Devices** and click **+Add**. You need to assign a name to the network access device and configure its IP address, as shown in Figure 5-64.

Technet24

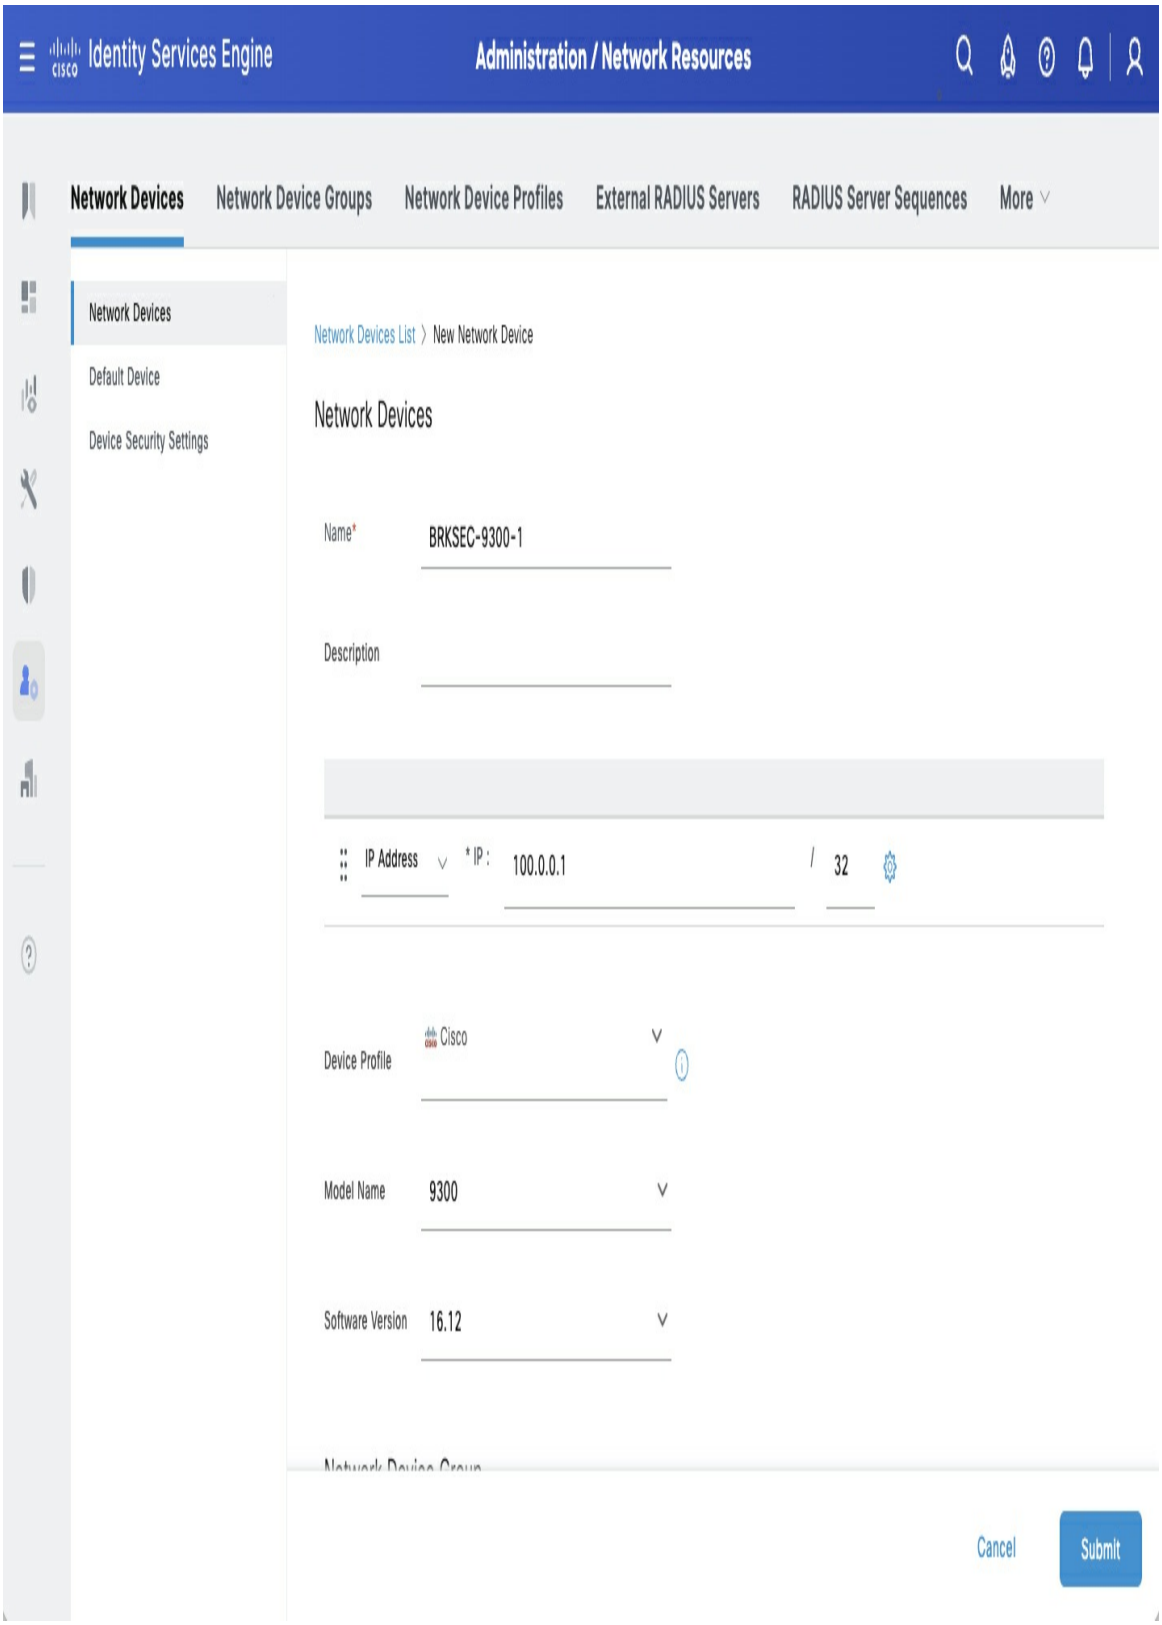


Figure 5-64 *ISE Network Device Configuration*

You can assign multiple IP addresses to a single network access device. You can also configure a broader network for a device, which allows ISE to treat multiple devices within that network as a single network access device. However, doing this will make it harder to create differentiated policies per network device in future.

> **Note**
>
> If there is an overlapping network access device space, the most specific entry takes precedence. For example, when you configure Switch1 with 10.0.0.1 and Switch2 with 10.0.0.0/24, Switch1 will authenticate users if the packet originates from 10.0.0.1.

Next, you configure the network device groups associated with the network access device, as shown in Figure 5-65.

> **Note**
>
> The IPSEC network device group is a legacy configuration element from earlier ISE designs that used an ESR router to terminate IPsec tunnels. In current ISE releases, this element has no functional use. You should not assign any values to this group.

Technet24

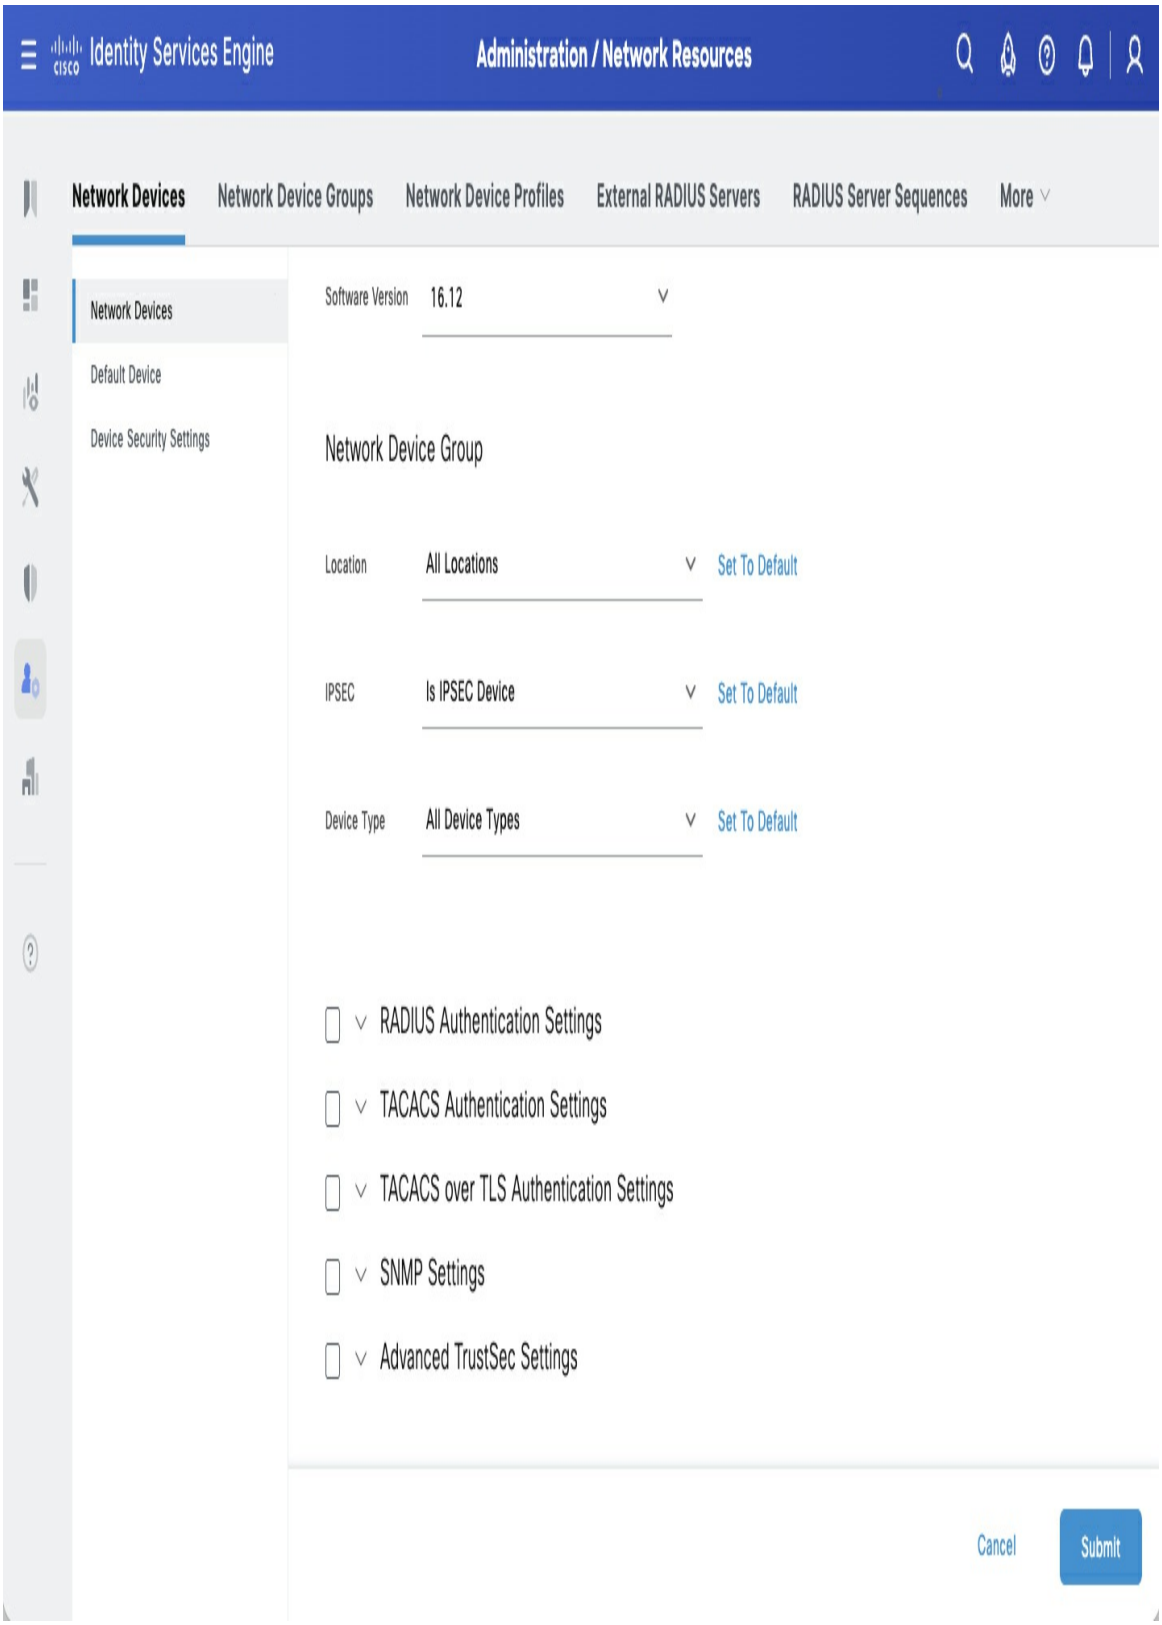


Figure 5-65 *ISE Network Device Configuration*

If the device uses RADIUS, you need to define a RADIUS shared secret that matches the value configured on the network device, as shown in Figure 5-66. From the same menu, you can also configure DTLS settings if you want to establish a TLS tunnel before sending RADIUS authentication requests. This configuration requires certificates to be provisioned on the network devices and trusted by ISE.

Technet24

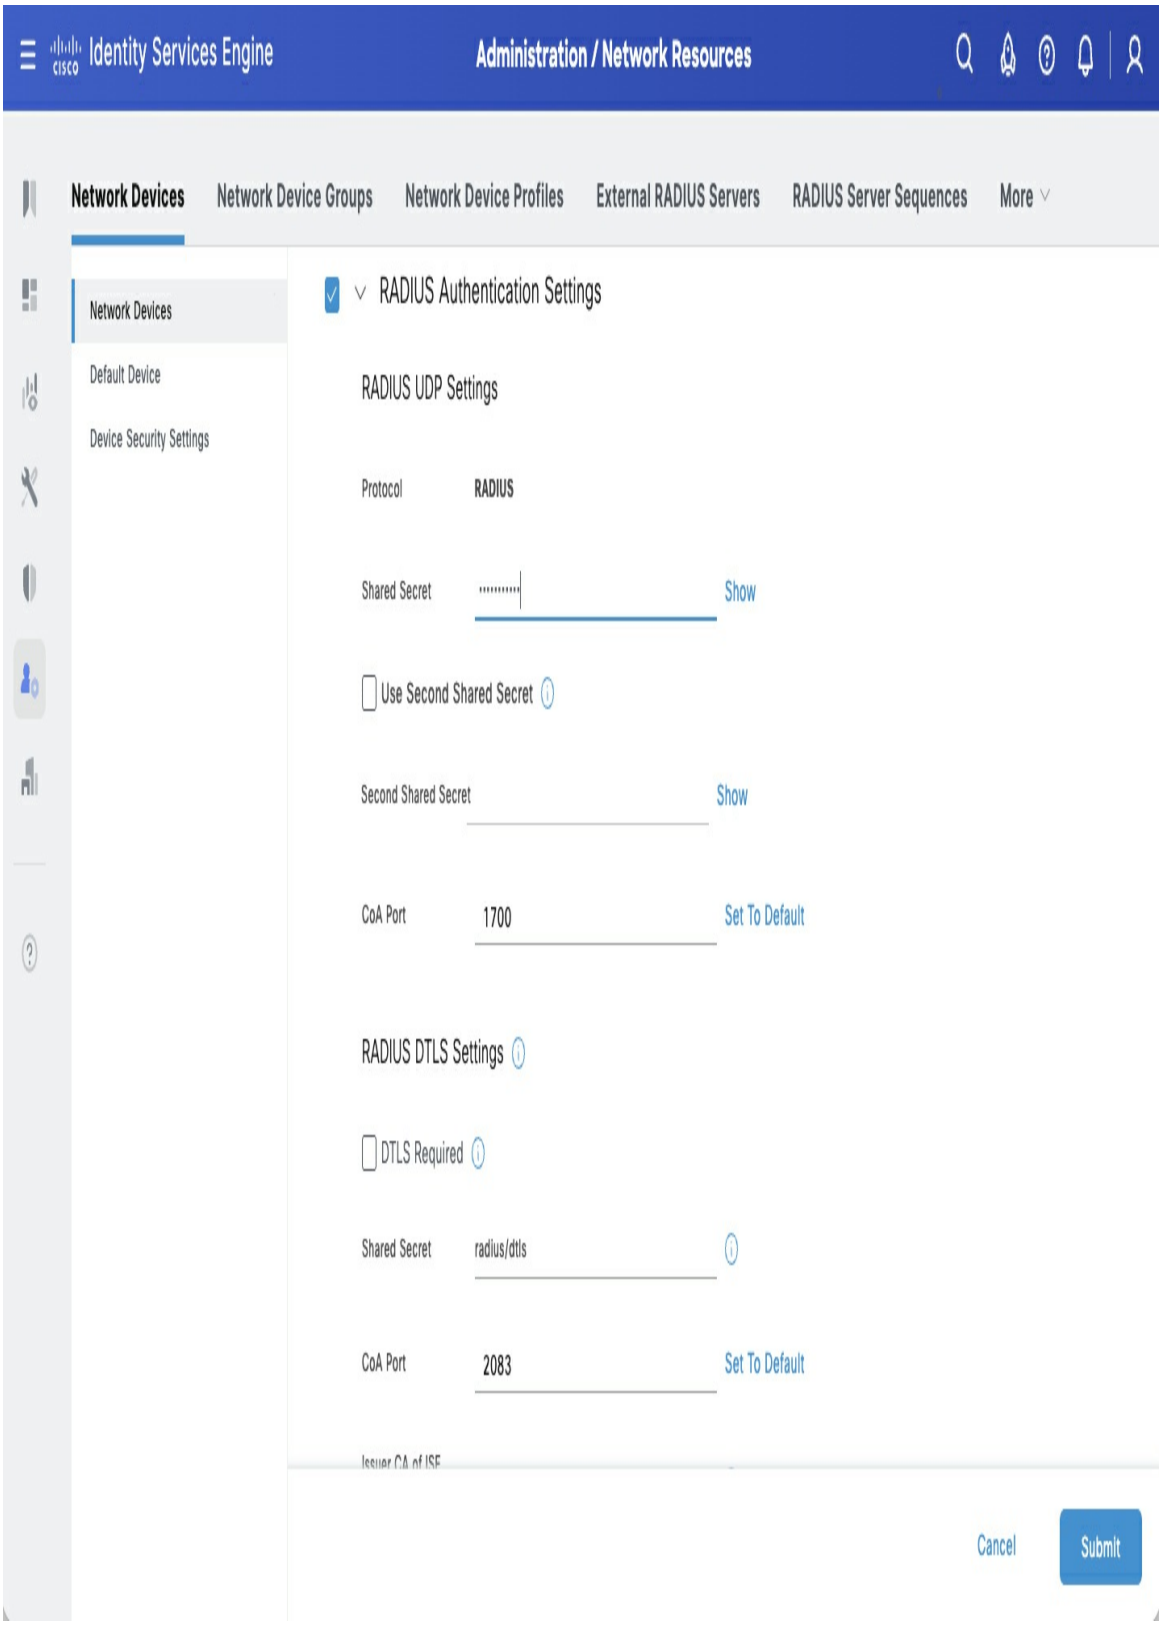


Figure 5-66 *ISE Network Device RADIUS Configuration*

Similar to RADIUS, when you configure device administration with TACACS+, you need to define a TACACS+ shared secret, as shown in Figure 5-67.

Technet24

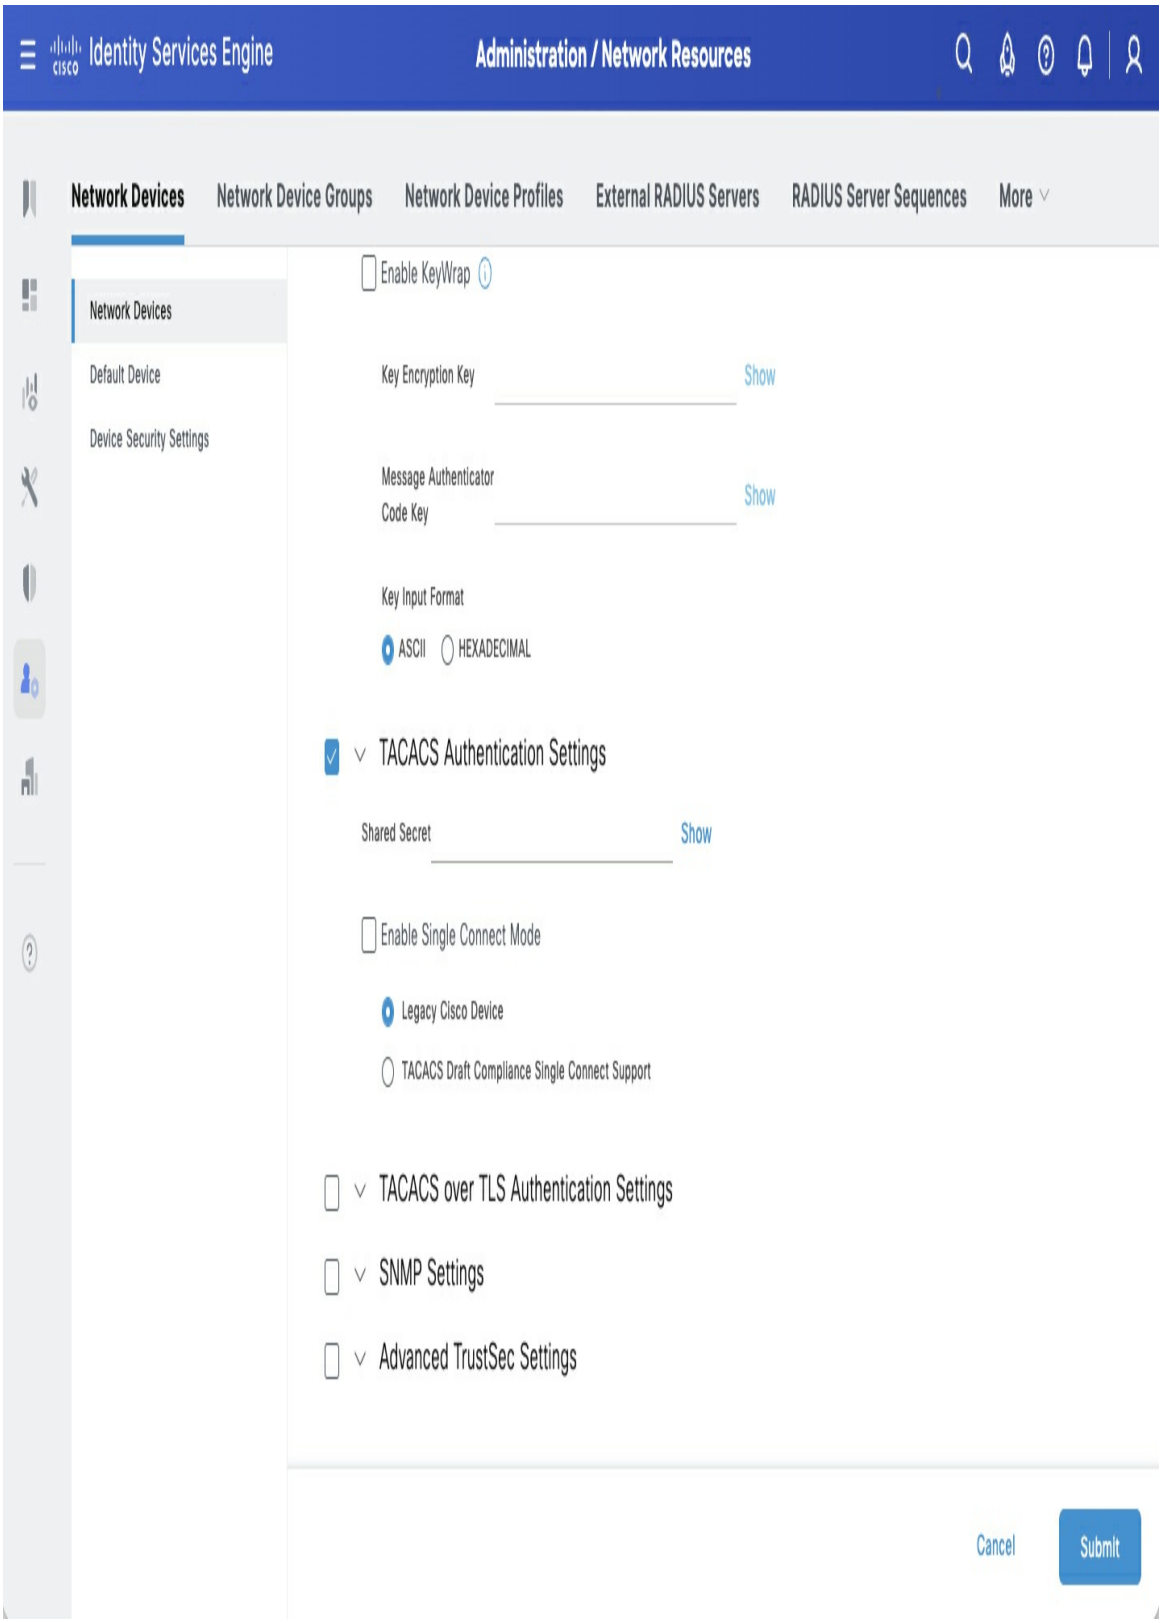


Figure 5-67 *ISE Network Device TACACS Configuration*

Summary

Authentication is one of the most critical elements of any network access control system, including ISE. To authenticate users and devices, you must decide which identity store ISE uses for authentication. ISE provides a broad range of authentication options, including Active Directory, LDAP, and Entra ID for username and password-based authentication. You can also use the Certificate Authentication Profile to authenticate endpoints based on certificates. ISE supports native MFA integration with Duo to add an additional layer of security. Finally, you must define network access devices for authentication to succeed.

Technet24

Chapter 6. Authorization Options

Introduction

Authorization is a key element of network access control. It differentiates the privileges that you, as administrator, assign to users and devices. To improve security, you should grant network access only to the required resources, or use the principle of least privilege. On Cisco Identity Services Engine, authorization policies and authorization profiles serve this purpose.

Authorization Policies Concepts

You build authorization policies in the same way you build policy sets and authentication policies. You define a condition that a policy must meet to match. ISE then evaluates policies from top to bottom, and only a single policy matches on a first-match basis.

By default, ISE includes some built-in policies. Figure 6-1 shows these policies.

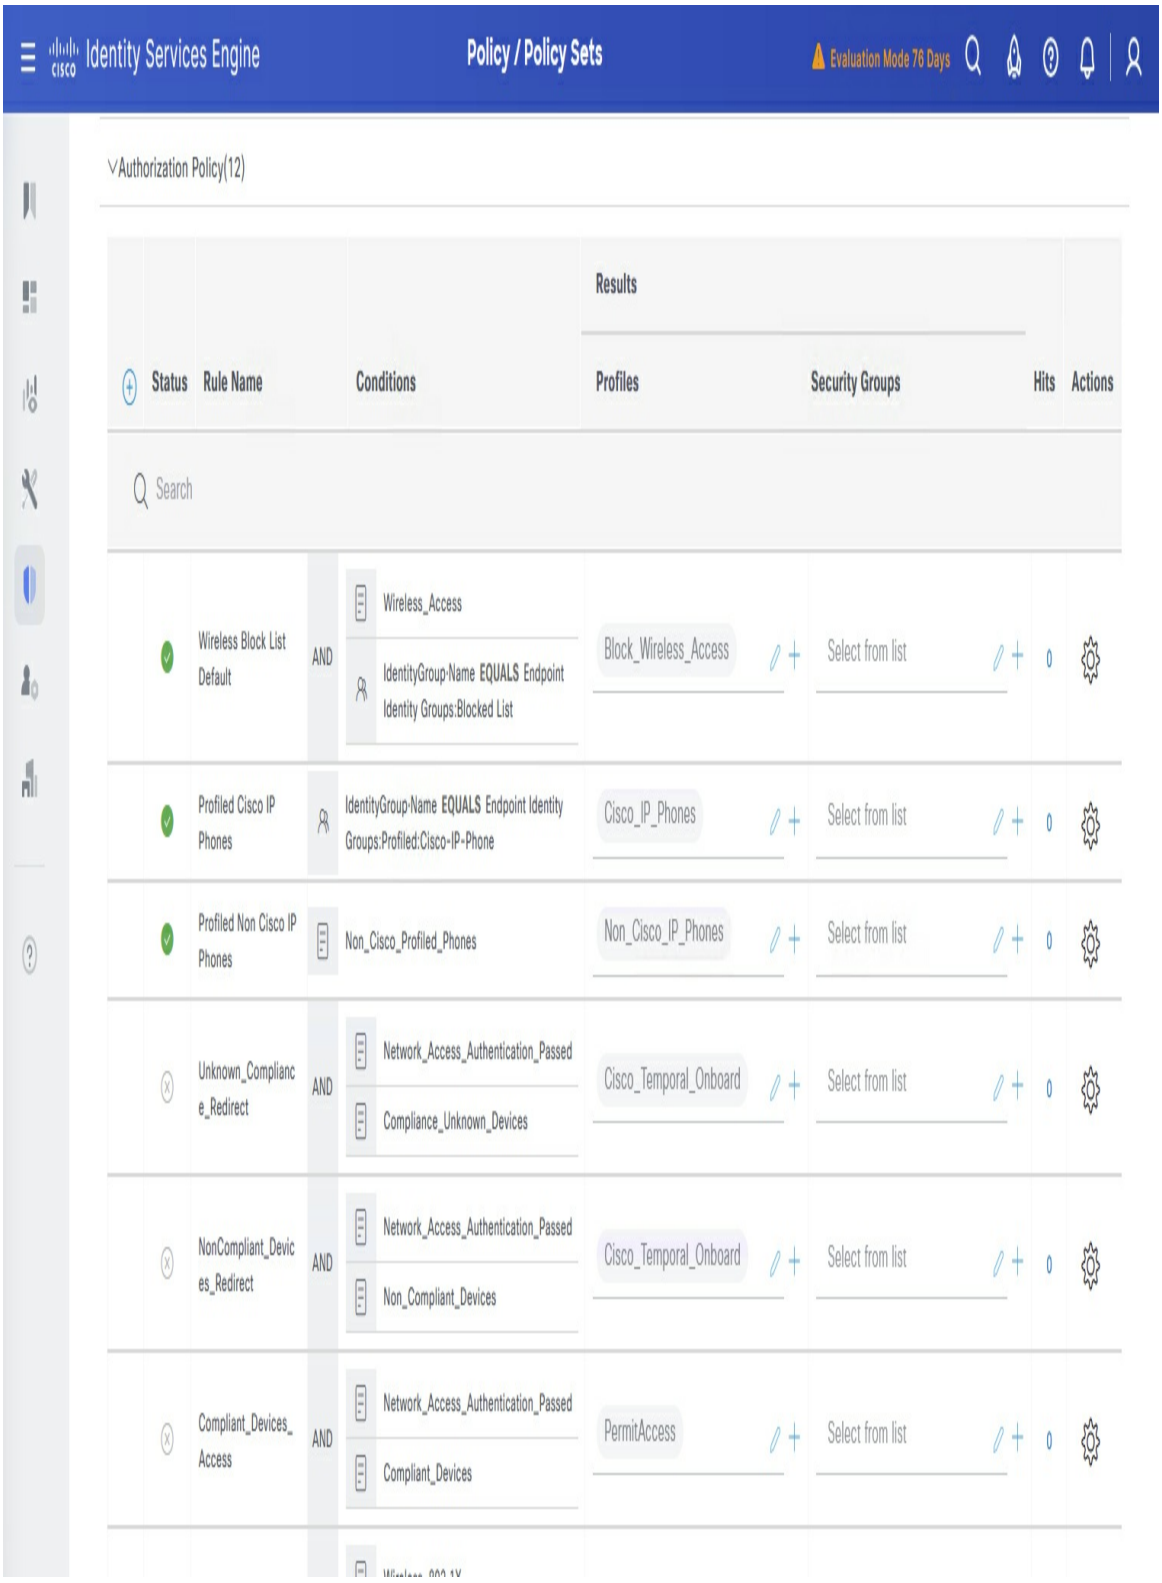


Figure 6-1 *Built-in Authorization Policies on ISE*

Technet24

ISE enables some of these policies by default and disables others by default. Preconfigured policies simplify the administrator experience. Instead of configuring use cases from scratch, you can adjust the policies that ISE already preconfigured. You should consider use cases in authorization policies carefully. Because authorization determines who receives which privileges, most features rely on authorization policies. Profiling, Posture, BYOD, MDM, and Guest all require part of their configuration in authorization.

Best Practice:

Cisco recommends that you keep the authorization conditions of the default policy set empty, with a single default policy with DenyAccess result. This approach ensures that no one receives network access unless authenticated according to the configured policies. This recommendation holds true provided you have created dedicated policy sets with appropriate authentication and authorization policies within them.

Let's examine one of the policy examples to understand the logic behind policies for use cases. Figure 6-2 shows two built-in authorization policies for the Guest use case.

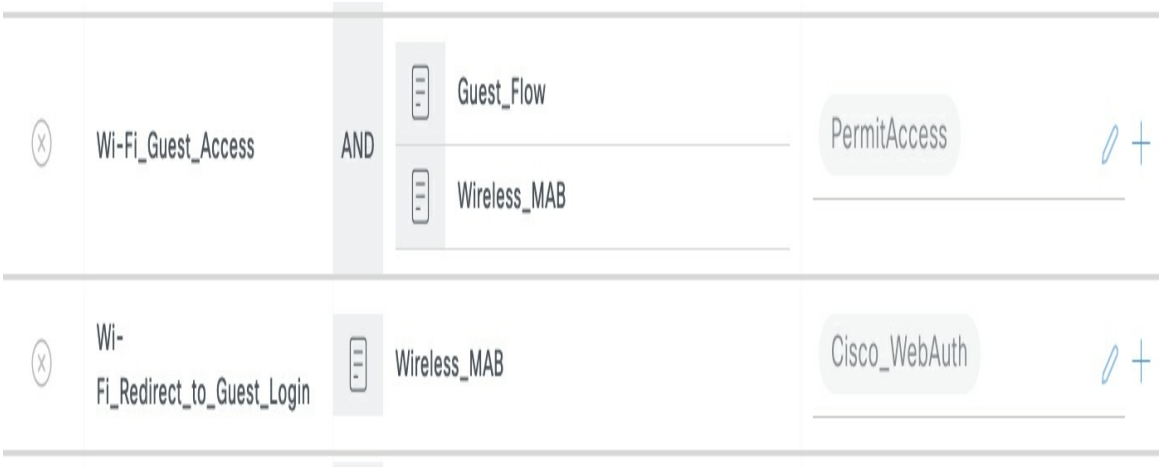


Figure 6-2 *Built-in Authorization Policies for the Guest Use Case*

The Wi-Fi_Redirect_to_Guest_Login policy serves as the catch-all policy. If a user matches this policy, ISE redirects the user to the guest portal. The user then completes guest authentication, and ISE issues a Change of Authorization (CoA). After that, the user authenticates again, matches Wi-Fi_Guest_Access, and receives new permissions that grant full guest access.

This concept is crucial to understanding ISE because most features rely on it. The process follows this sequence: initial authentication, CoA, and final authentication.

Change of Authorization (CoA)

Historically, the client always initiated RADIUS authentication to the server. Authentication started with an Access-Request and finished with either Access-Accept or Access-Reject. After the user was authenticated, the server had limited options and relied on reauthentication mechanisms to change authentication or authorization results when necessary.

Change of Authorization lets you act on already-authenticated sessions. The server sends CoA to the network access device (NAD) and can either disconnect the endpoint or force reauthentication. This mechanism fits most ISE features. Limited access is granted first. In the Guest use case, for example, the user can access only the guest portal initially. After the user passes authentication, full Internet access is granted.

You generally control CoA settings and behavior within the feature settings. For example, you control which CoA type the Profiler component sends from **Work Centers > Profiler > Settings >**

Technet24

Profiler Settings, as shown in Figure 6-3.

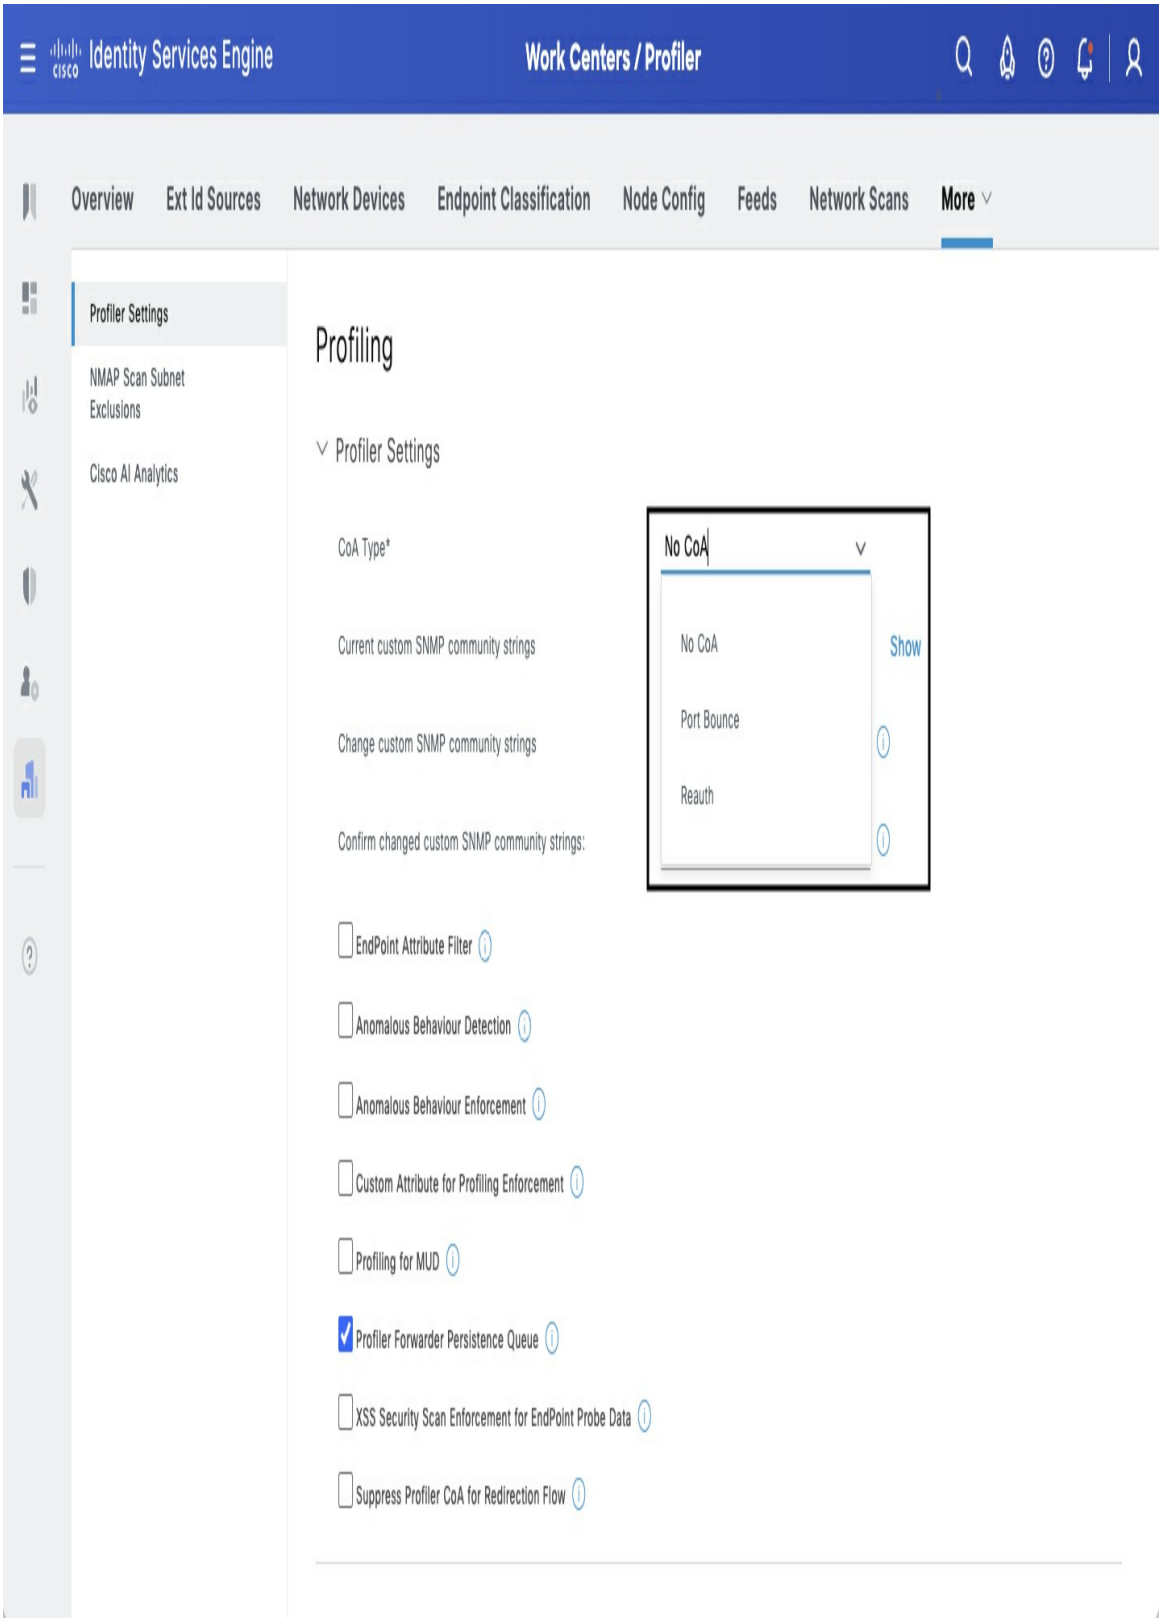


Figure 6-3 *Profiler CoA Settings*

However, for Guest CoA, you can choose the type from the Portal Settings in the case of the Hotspot portal. The type of CoA would depend on what the network access device supports for the specific feature. You can access this option under **Work Centers > Guest Access > Settings > Guest Portals > Hotspot Portal**, as shown in Figure 6-4.

Technet24

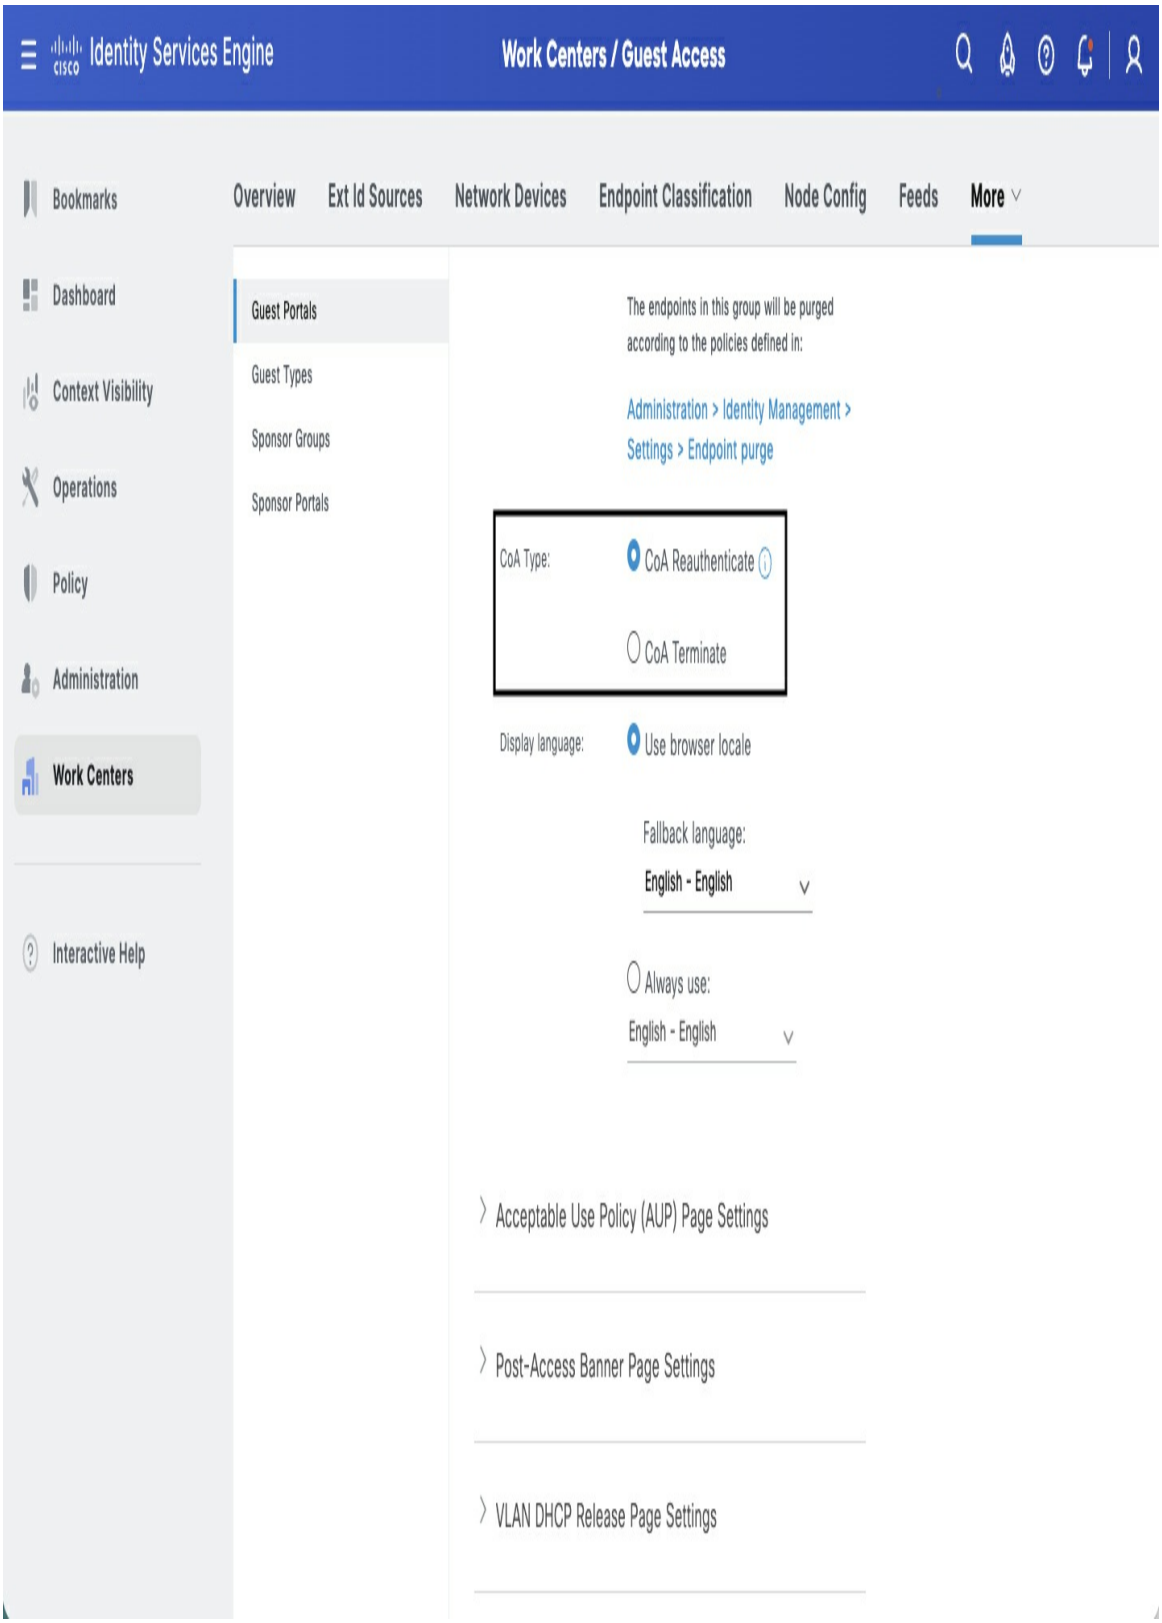


Figure 6-4 *Guest CoA Settings*

You can also issue a CoA manually for any connected endpoint. You can do this from **Operations > RADIUS > Live Sessions**, as shown in Figure 6-5.

Technet24

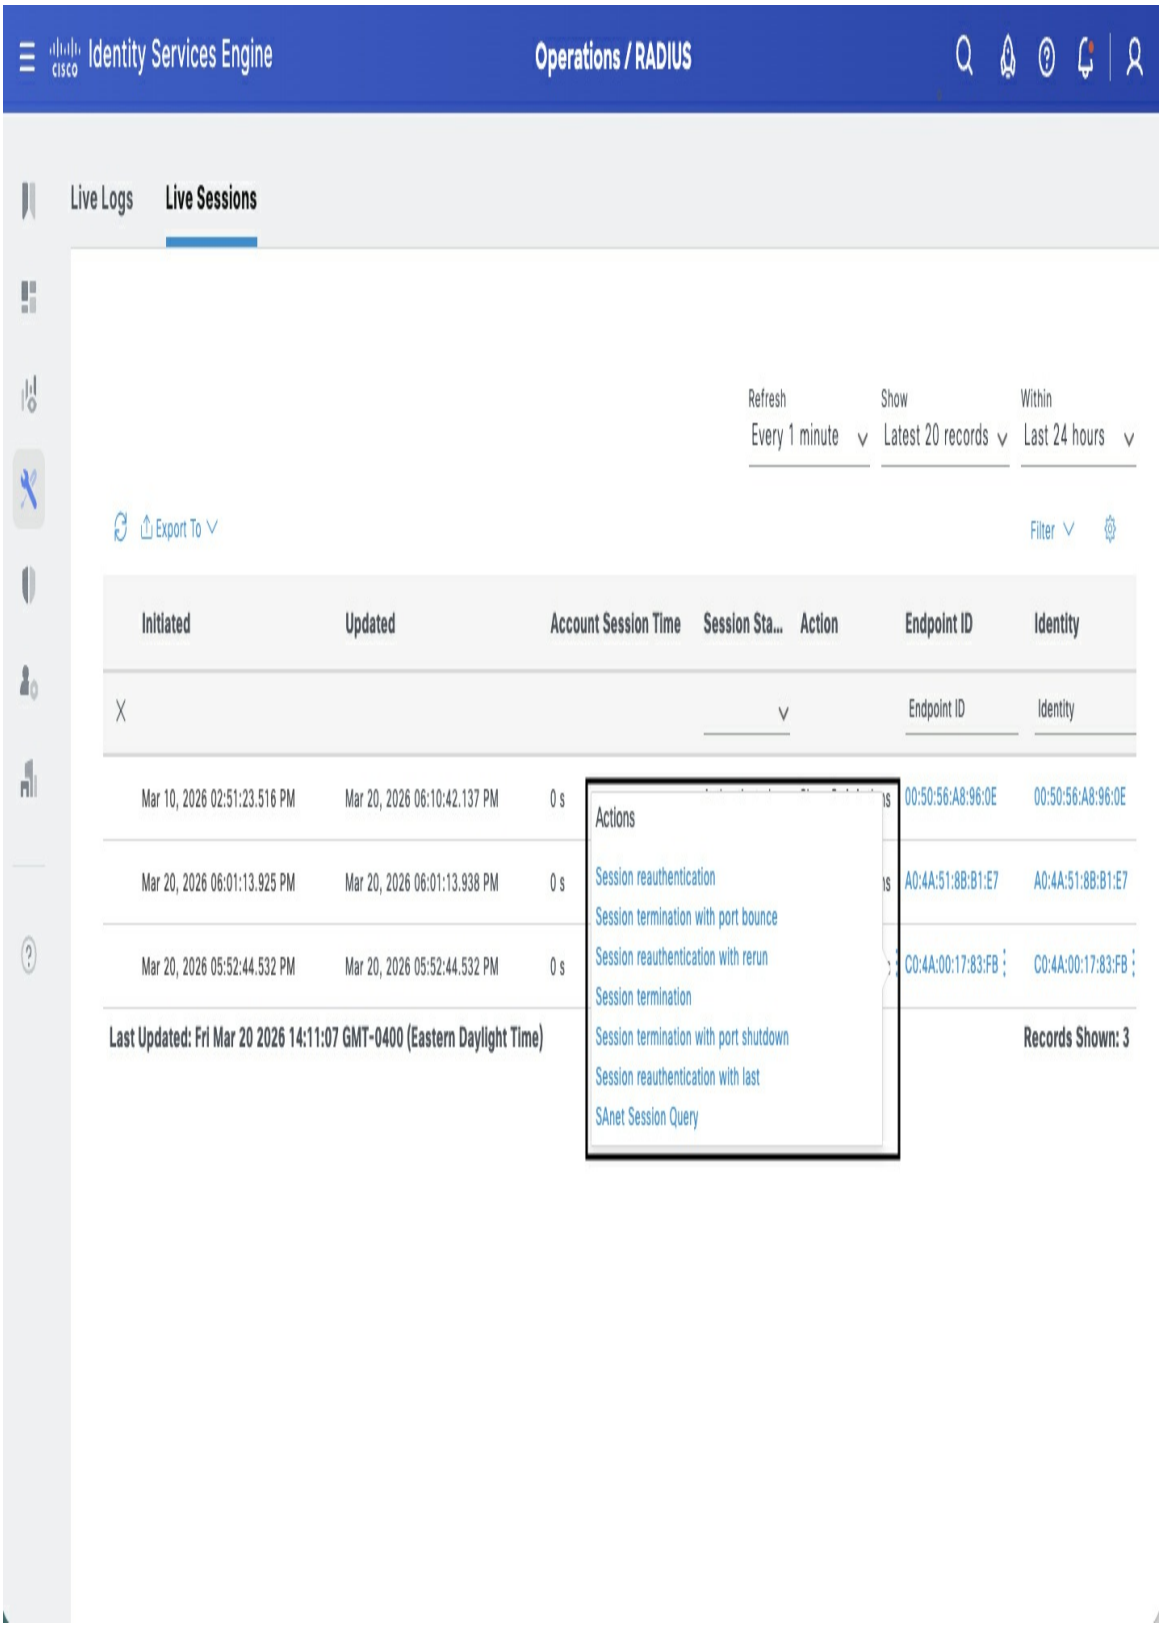


Figure 6-5 *Live Sessions CoA Options*

In the example shown in Figure 9-6, a CoA Reauthenticate is issued for the endpoint with MAC address A0:4A:51:8B:B1:E7. Live Logs show a successful CoA attempt. CoA entries look very different compared to Authentication Live Logs. When you issue a CoA, the Live Log entry contains empty fields for **Identity**, **Endpoint Profile**, **Authentication Policy**, and **Authorization Policy**. This makes it easier to spot CoA in the Live Logs, as shown in Figure 6-6.

Technet24

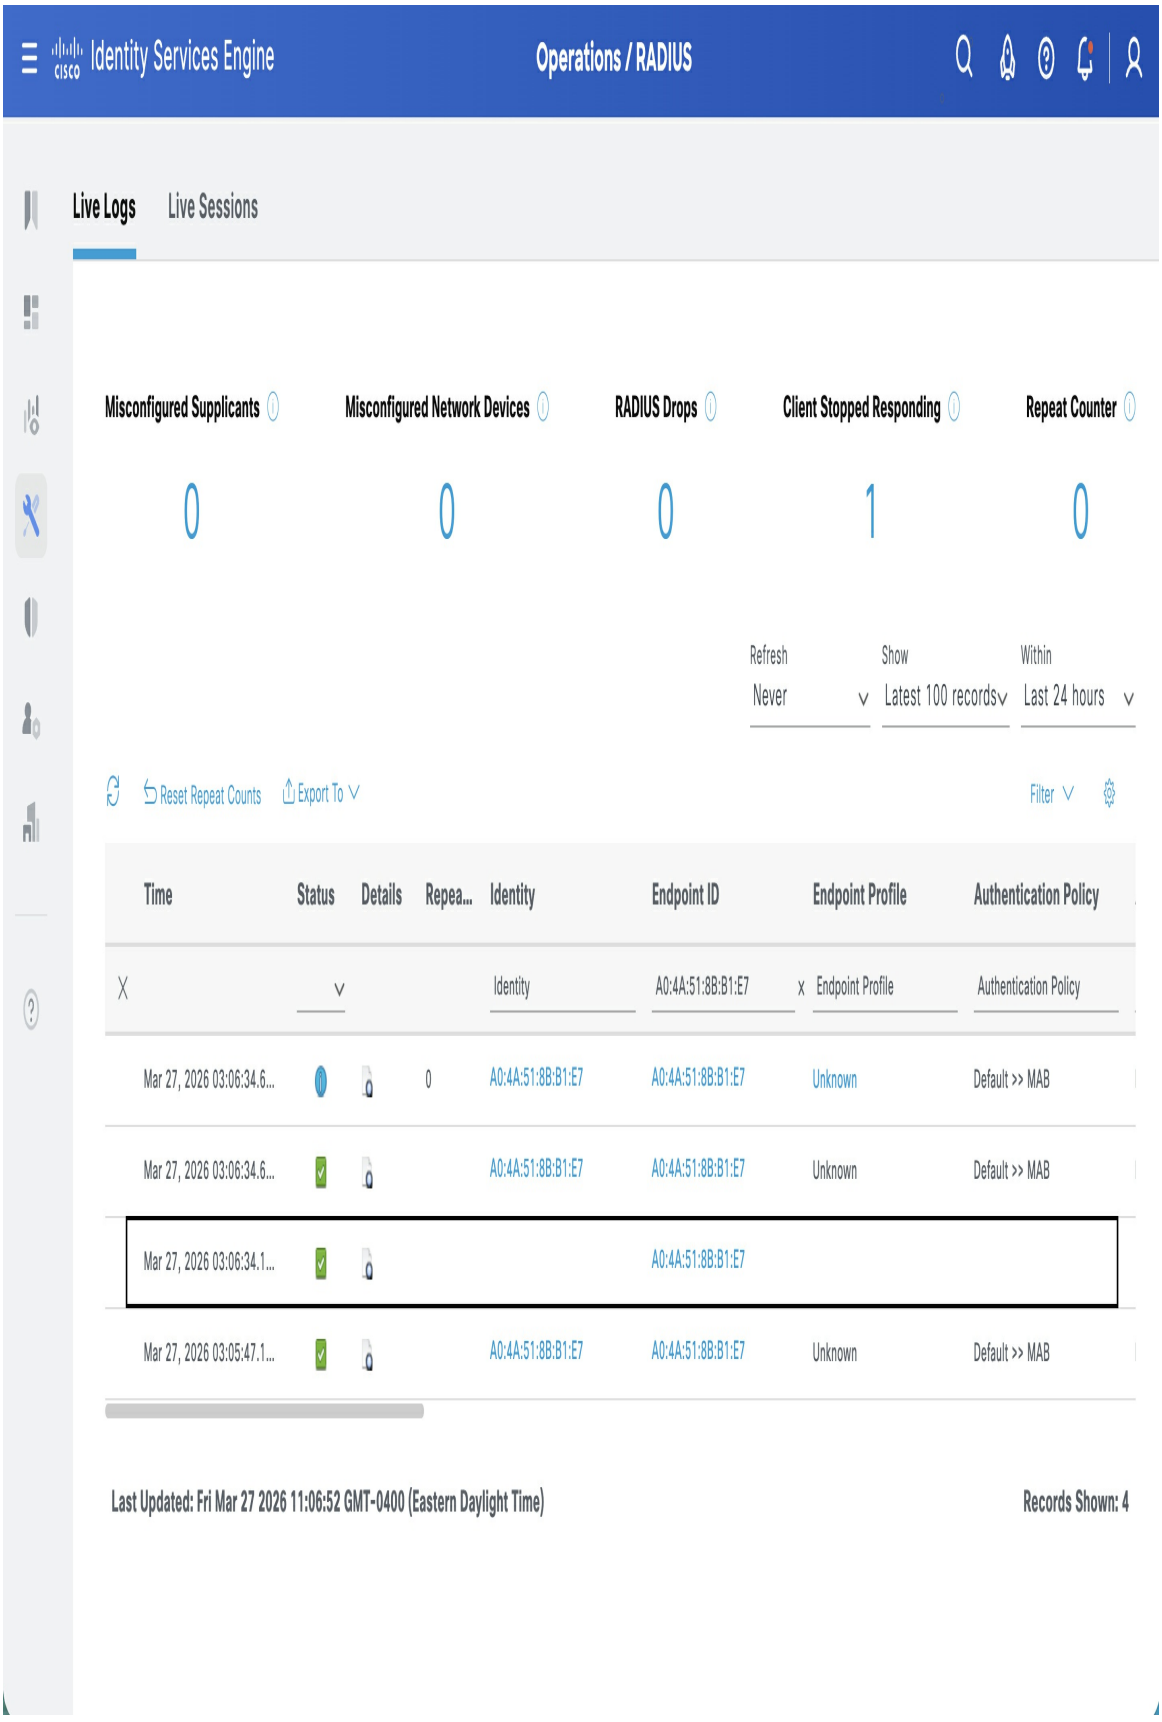

Identity Services Engine
Operations / RADIUS
Live Logs
Live Sessions
Misconfigured Supplicants
Misconfigured Network Devices
RADIUS Drops
Client Stopped Responding
Repeat Counter
0
0
0
1
0
Refresh
Never
Show
Latest 100 records
Within
Last 24 hours
Reset Repeat Counts
Export To
Filter
Time
Status
Details
Repea...
Identity
Endpoint ID
Endpoint Profile
Authentication Policy
Identity
A0:4A:51:8B:B1:E7
Endpoint Profile
Authentication Policy
Mar 27, 2026 03:06:34.6...
0
A0:4A:51:8B:B1:E7
A0:4A:51:8B:B1:E7
Unknown
Default >> MAB
Mar 27, 2026 03:06:34.6...
A0:4A:51:8B:B1:E7
A0:4A:51:8B:B1:E7
Unknown
Default >> MAB
Mar 27, 2026 03:06:34.1...
A0:4A:51:8B:B1:E7
Mar 27, 2026 03:05:47.1...
A0:4A:51:8B:B1:E7
A0:4A:51:8B:B1:E7
Unknown
Default >> MAB
Last Updated: Fri Mar 27 2026 11:06:52 GMT-0400 (Eastern Daylight Time)
Records Shown: 4

Figure 6-6 *Successful CoA Live Logs*

It is crucial not to miss CoA entries when you troubleshoot specific endpoint flows. When you filter the logs, use the **Endpoint ID** column and populate it with the MAC address, not **Identity**, since the **Identity** field is empty for CoA.

When you click the magnifying glass next to the CoA, you see several crucial elements, as shown in Figure 6-7.

Technet24

Other Attributes

IPSEC	IPSEC#Is IPSEC Device#No
Device CoA type	Cisco CoA
TotalAuthenLatency	802500902
CiscoAVPair	subscriber:command=reauthenticate
Device Type	Device Type#All Device Types
IsThirdPartyDeviceFlow	false
Network Device Profile	Cisco
CoASourceComponent	Live Session UI or ERS
Device IP Address	10.48.26.61
Event-Timestamp	1774623993
Device CoA port	1700
ClientLatency	802500730
ConfigVersionId	317
CoAReason	COA initiated by admin
AcsSessionID	cc2ef066-e025-4f2a-8976-9958428644fe
NetworkDeviceProfileId	b0699505-3150-4215-a80e-6753d45bf56c

Figure 6-7 *CoA Detailed Authentication Report*

CiscoAVPair subscriber-command=reauthenticate indicates that you sent a CoA Reauthenticate.

CoASourceComponent Live Sessions UI or ERS confirms that you issued a CoA directly from Live Sessions. If the source is Profiling, Guest, or Posture, ISE shows the respective source component.

CoAReason is CoA initiated by admin, which confirms that you issued CoA manually.

Network access devices must be configured to accept a CoA packets from ISE Policy Service Nodes (PSNs).

Authorization Policy Conditions

Authorization policies rely on the same Conditions Studio as policy sets and authentication policies, but authorization policy conditions can be constructed on a wider variety of attributes. For instance, authorization conditions can use attributes and groups from Active Directory, LDAP, Entra, and RADIUS token servers. They can also rely on internal user and endpoint groups and attributes. This approach allows greater flexibility during authorization and enables proper differentiated access.

ISE supports two types of conditions: internal and external. Internal conditions are available locally on ISE and do not require external calls. Examples include Endpoint Identity Group, Network Device Group, and Endpoint Logical Profile. When ISE evaluates authorization policies, it processes them from top to bottom. Within each authorization policy, conditions are also evaluated from top to bottom.

Suppose you are configuring a policy for IP cameras within a policy set

Technet24

dedicated to a certain region. The policy set already contains the Domain Users Access Rule, which grants network access based on an external identity group. Figure 6-8 shows a poor example. As a result of this policy implementation, each authorization for an IP camera triggers an "expensive" external lookup. Since policies are evaluated from top to bottom, ISE must ensure the camera does not belong to an Active Directory user group.


Authorization Policy(2)
Results
Status
Rule Name
Conditions
Profiles
Security Groups
Hits
Actions
Search
Domain Users Access
ISELAB-ExternalGroups EQUALS iselab.net/Users/Domain Users
PermitAccess
Select from list
0
IP Cameras Access
EndPoints-LogicalProfile EQUALS Cameras
PermitAccess
Select from list
Default
DenyAccess
Select from list
0

Technet24

Figure 6-8 *Poor Example of Authorization Policy Configuration*

Instead, you should place the logical profile condition at the top. This condition is "cheap" because it does not consume resources for external calls and allows ISE to check the logical profile for the endpoint quickly and internally. Figure 6-9 shows the correct policy placement.

Authorization Policy(2)

Status	Rule Name	Conditions	Results: Profiles	Results: Security Groups	Hits	Actions
Search						
✓	IP Cameras Access	EndPoints·LogicalProfile EQUALS Cameras	PermitAccess	Select from list		
✓	Domain Users Access	ISELAB-ExternalGroups EQUALS iselab.net/Users/Domain Users	PermitAccess	Select from list	0	
✓	Default		DenyAccess	Select from list	0	

Technet24

Figure 6-9 *Correct Example of Authorization Policy Configuration*

It is equally important to build conditions properly within a single authorization policy. Figure 6-10 shows a poor example of an MDM compliant access authorization policy. When ISE evaluates this policy, it first queries the external MDM server to check whether the endpoint is compliant. Only then does it check whether the endpoint is a Windows 11 workstation, which ISE performs locally.



Figure 6-10 *Poor Example of Authorization Policy Condition Configuration*

Technet24

Instead, you should reverse the policy, as shown in Figure 6-11. ISE should first perform a quick local verification of the endpoint profile. If the endpoint is not a Windows 11 workstation, you should not trigger an MDM query.



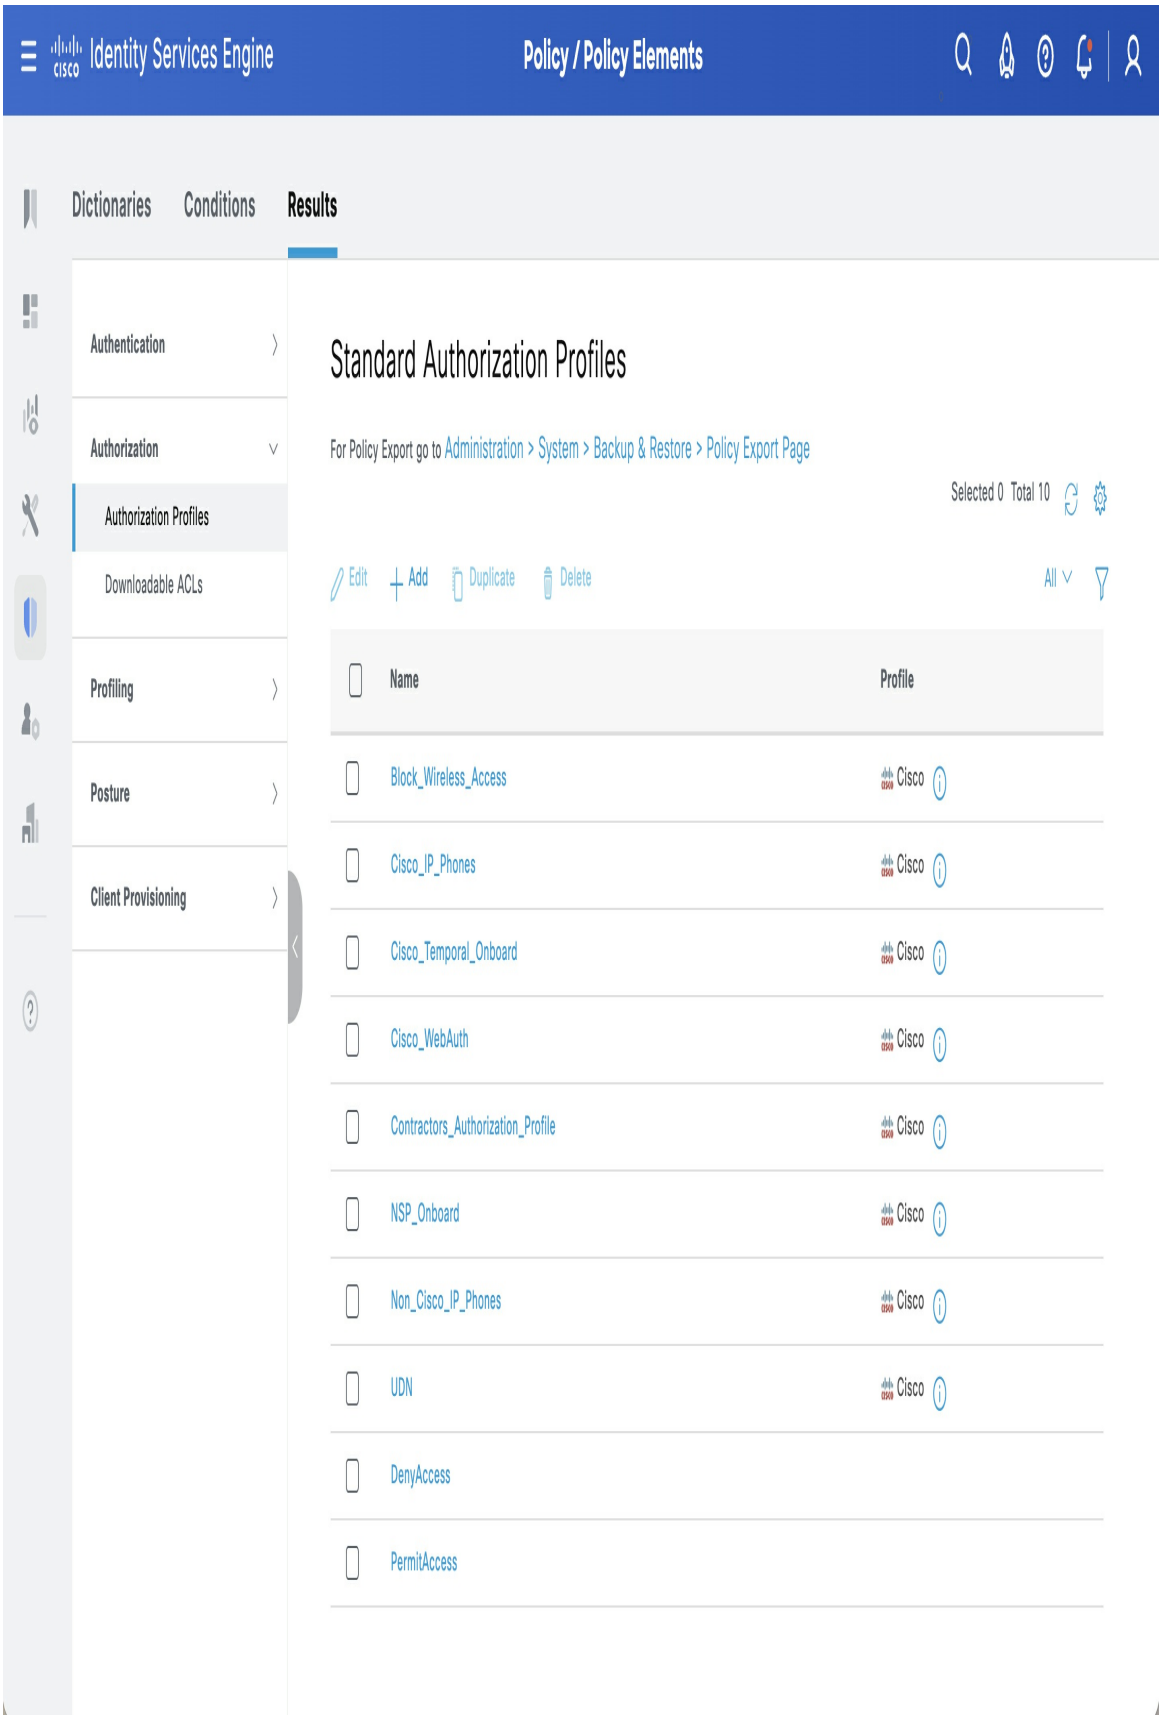
Figure 6-11 *Correct Example of Authorization Policy Condition Configuration*

Technet24

Authorization Profiles

Authorization profiles control the permissions you assign to the user or the endpoint. They allow granular access through downloadable access control lists (DACLs), enable microsegmentation by assigning security group tags (SGTs), and support advanced use cases with redirect ACLs and redirect URLs.

ISE includes some preconfigured authorization profiles, as shown in Figure 6-12.


Identity Services Engine
Policy / Policy Elements
Dictionaries
Conditions
Results
Authentication
Authorization
Authorization Profiles
Downloadable ACLs
Profiling
Posture
Client Provisioning
Standard Authorization Profiles
For Policy Export go to Administration > System > Backup & Restore > Policy Export Page
Selected 0 Total 10
Edit
Add
Duplicate
Delete
All
Name
Profile
Block_Wireless_Access
Cisco_IP_Phones
Cisco_Temporal_Onboard
Cisco_WebAuth
Contractors_Authorization_Profile
NSP_Onboard
Non_Cisco_IP_Phones
UDN
DenyAccess
PermitAccess
Cisco

Technet24

Figure 6-12 *Built-in Authorization Profiles*

Some of these built-in profiles are used in default authorization policies. For example, the Cisco_IP_Phones authorization profile assigns a DACL called PERMIT_ALL_IPV4_TRAFFIC and places the device in the voice domain, as shown in Figure 6-13.



Figure 6-13 *Cisco_IP_Phones Authorization Profile Attribute Details*

The Cisco_WebAuth authorization profile redirects the user to the built-in self-registered guest portal, as shown in Figure 6-14.



Figure 6-14 *Cisco_WebAuth Authorization Profile Attribute Details*

DenyAccess and PermitAccess authorization profiles send RADIUS Access-Accept and Access-Reject, respectively, with no extra attributes.

To configure an authorization profile, navigate to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**. Click **+Add** to add a profile, as shown in Figure 6-15.

Technet24

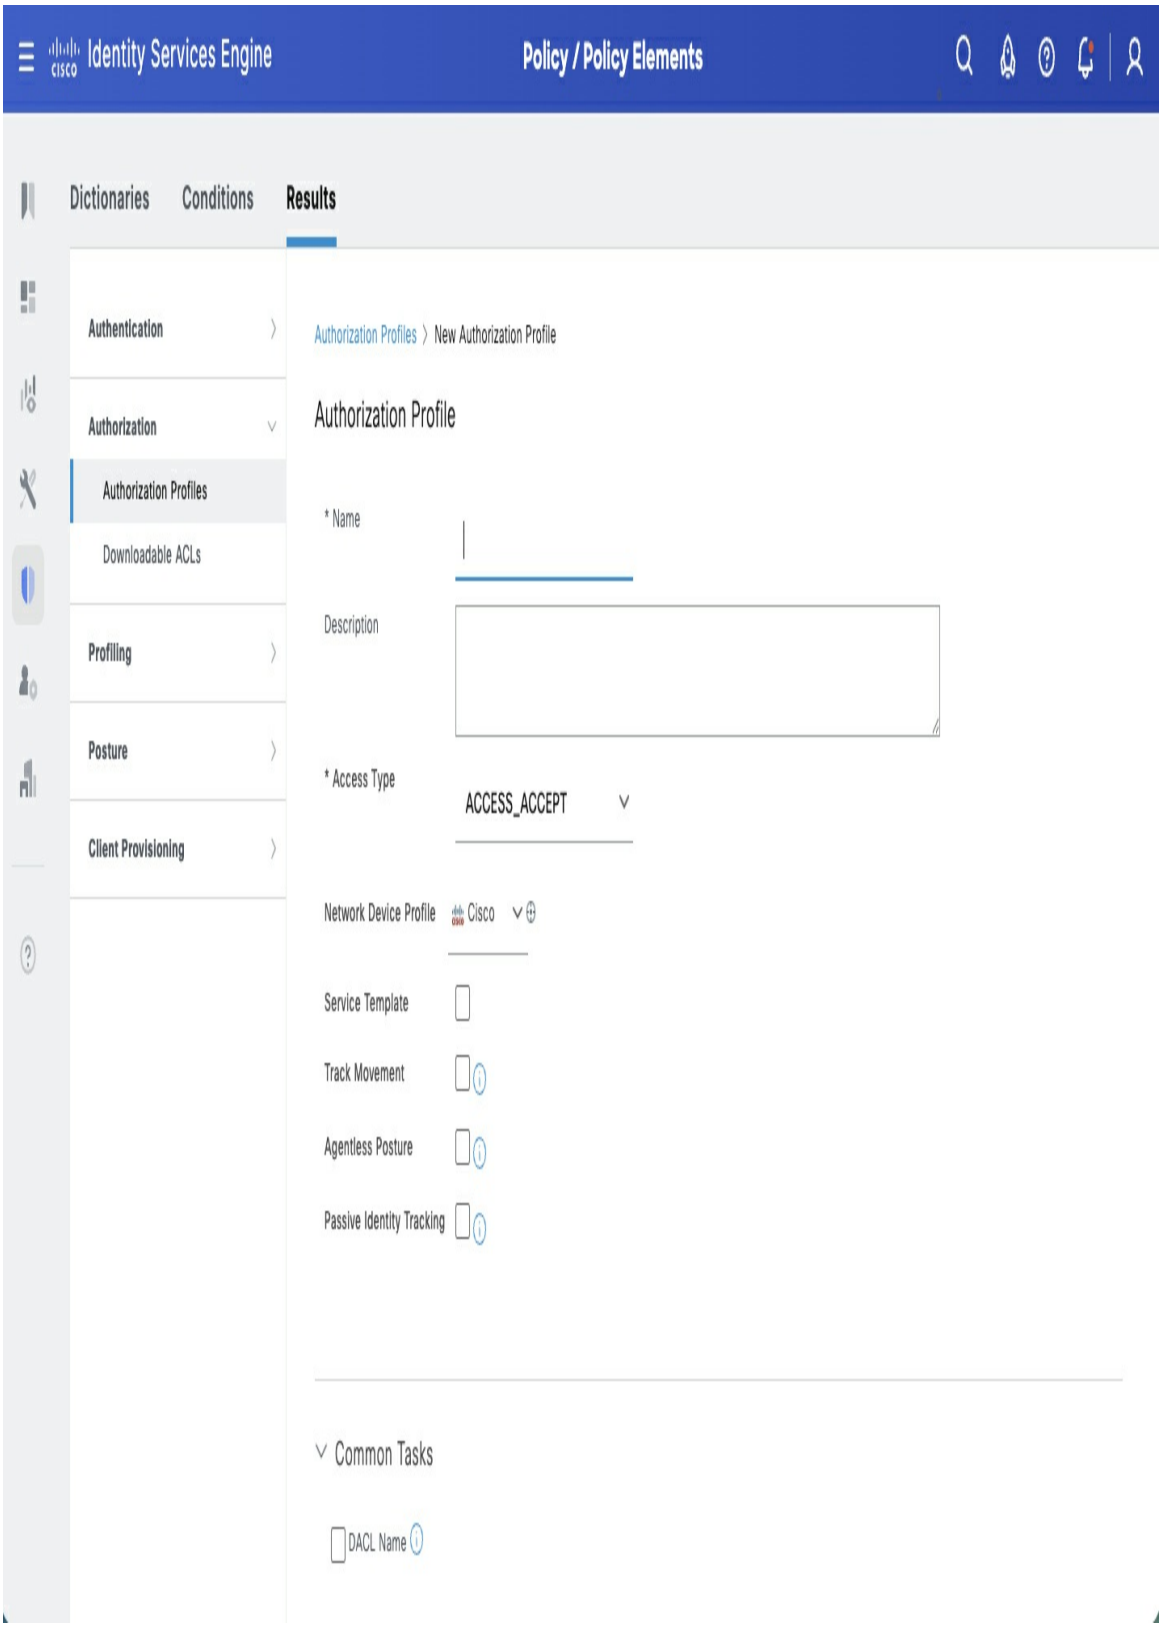


Figure 6-15 *Authorization Profile Configuration*

You need to provide a name for the profile and choose the **Access Type**. In most cases, you should select **ACCESS_ACCEPT** and add additional attributes.

The **Attributes** section contains two parts: **Common Tasks** and **Advanced Attributes Settings**.

Common Tasks include the attributes most widely used, as shown in Figure 6-16. Let's explore them in more detail.

Technet24

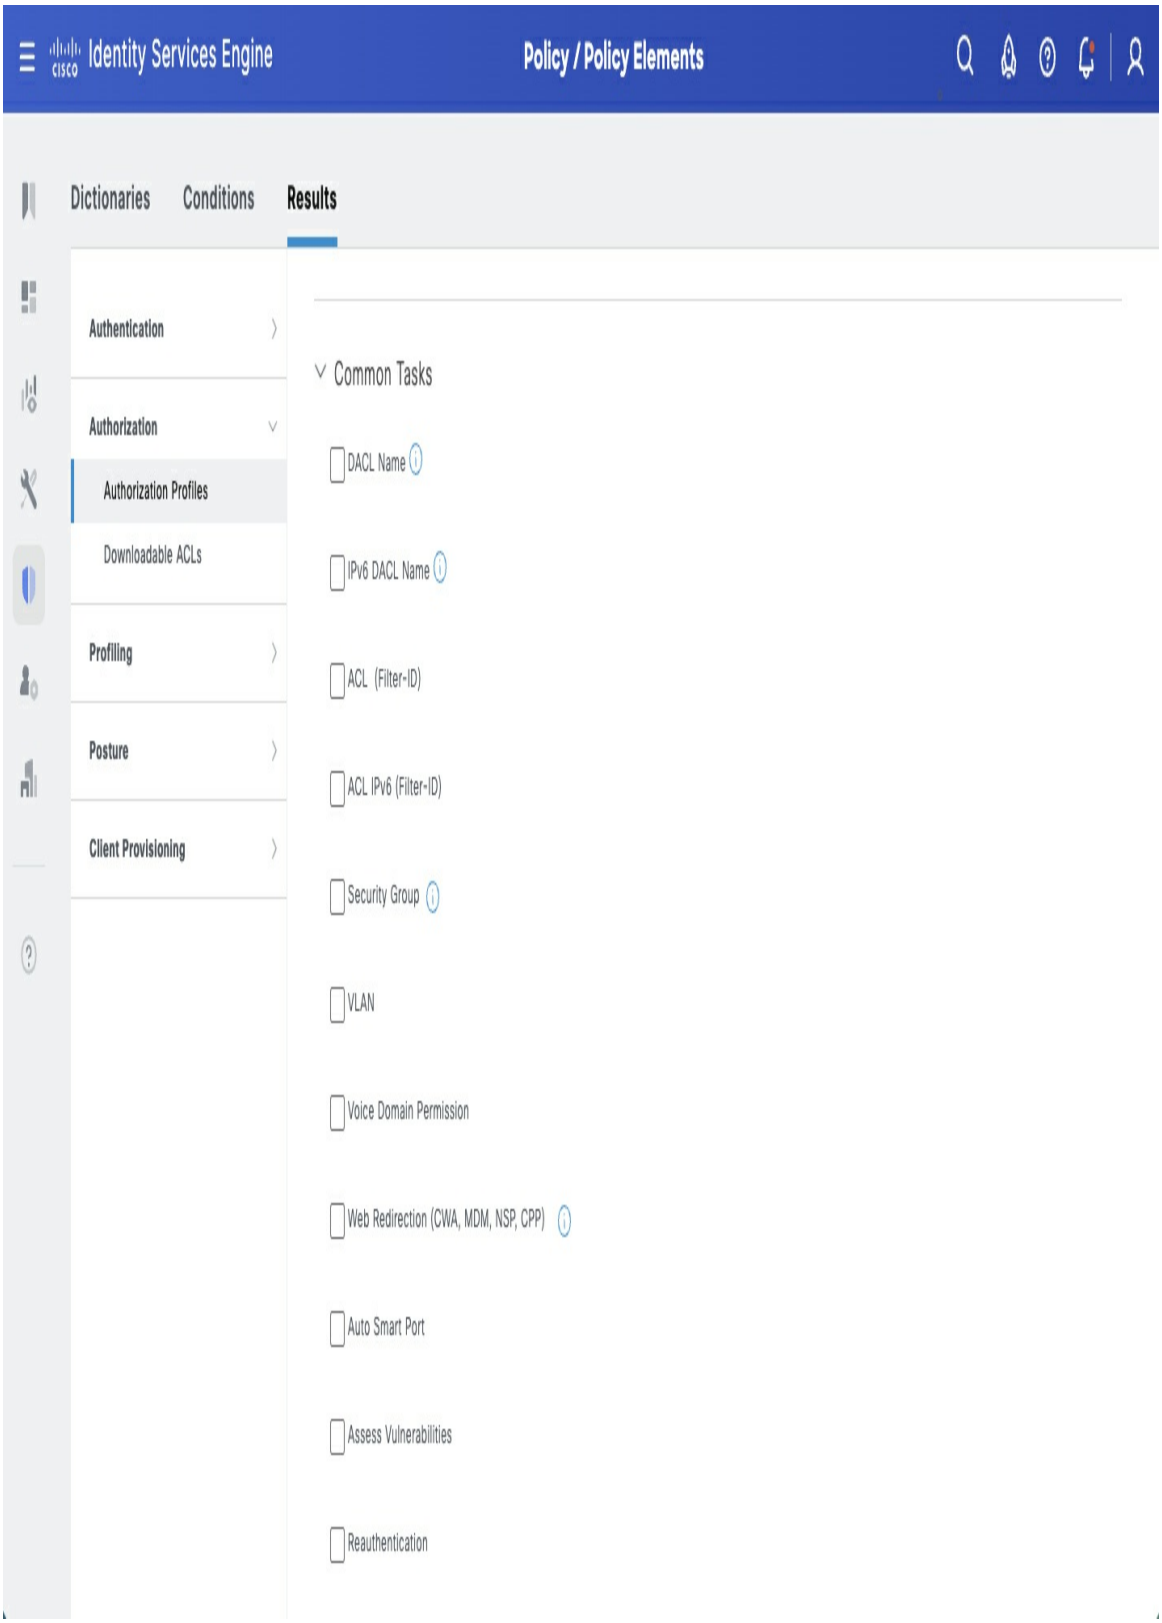


Figure 6-16 *Authorization Profile Common Tasks*

You can configure a dedicated DACL on ISE and assign it as part of the authorization. After you select this option, a drop-down menu appears, allowing you to choose the DACL, as shown in Figure 6-17. After the network access device receives it, the device sends an additional Access-Request to ISE to download the content of the DACL.

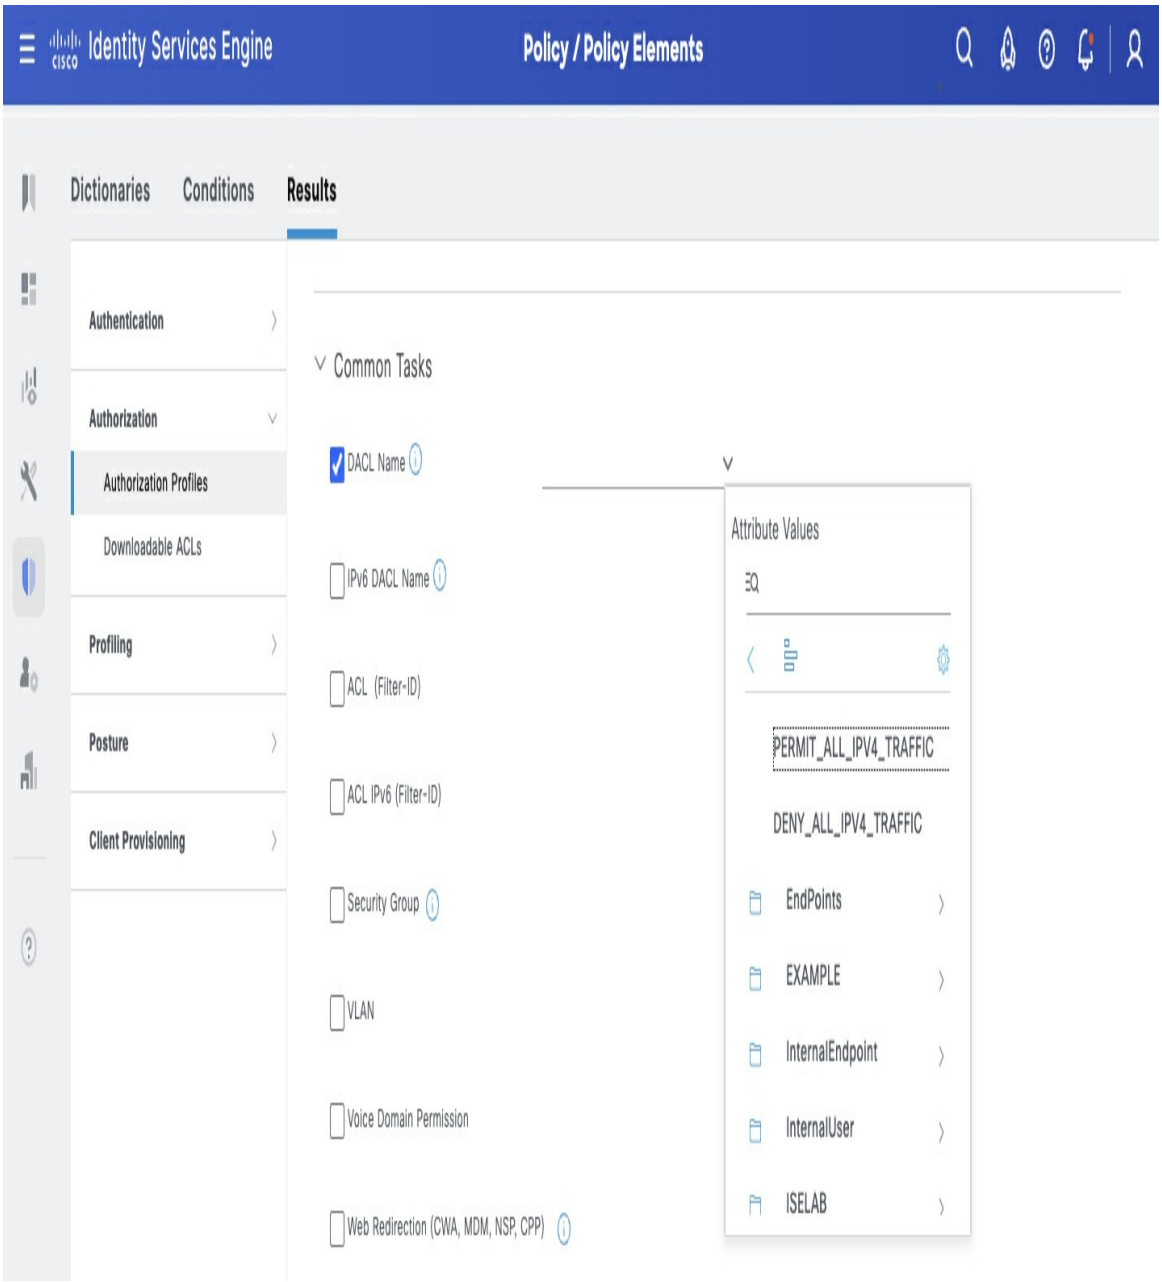


Figure 6-17 *Authorization Profile DACL Configuration*

Technet24

ACL (Filter-ID) allows you to configure an access control list that is already defined on the device and associate it with an authenticated user or endpoint. You can use an arbitrary name or an attribute from Active Directory or LDAP server, as shown in Figure 6-18.

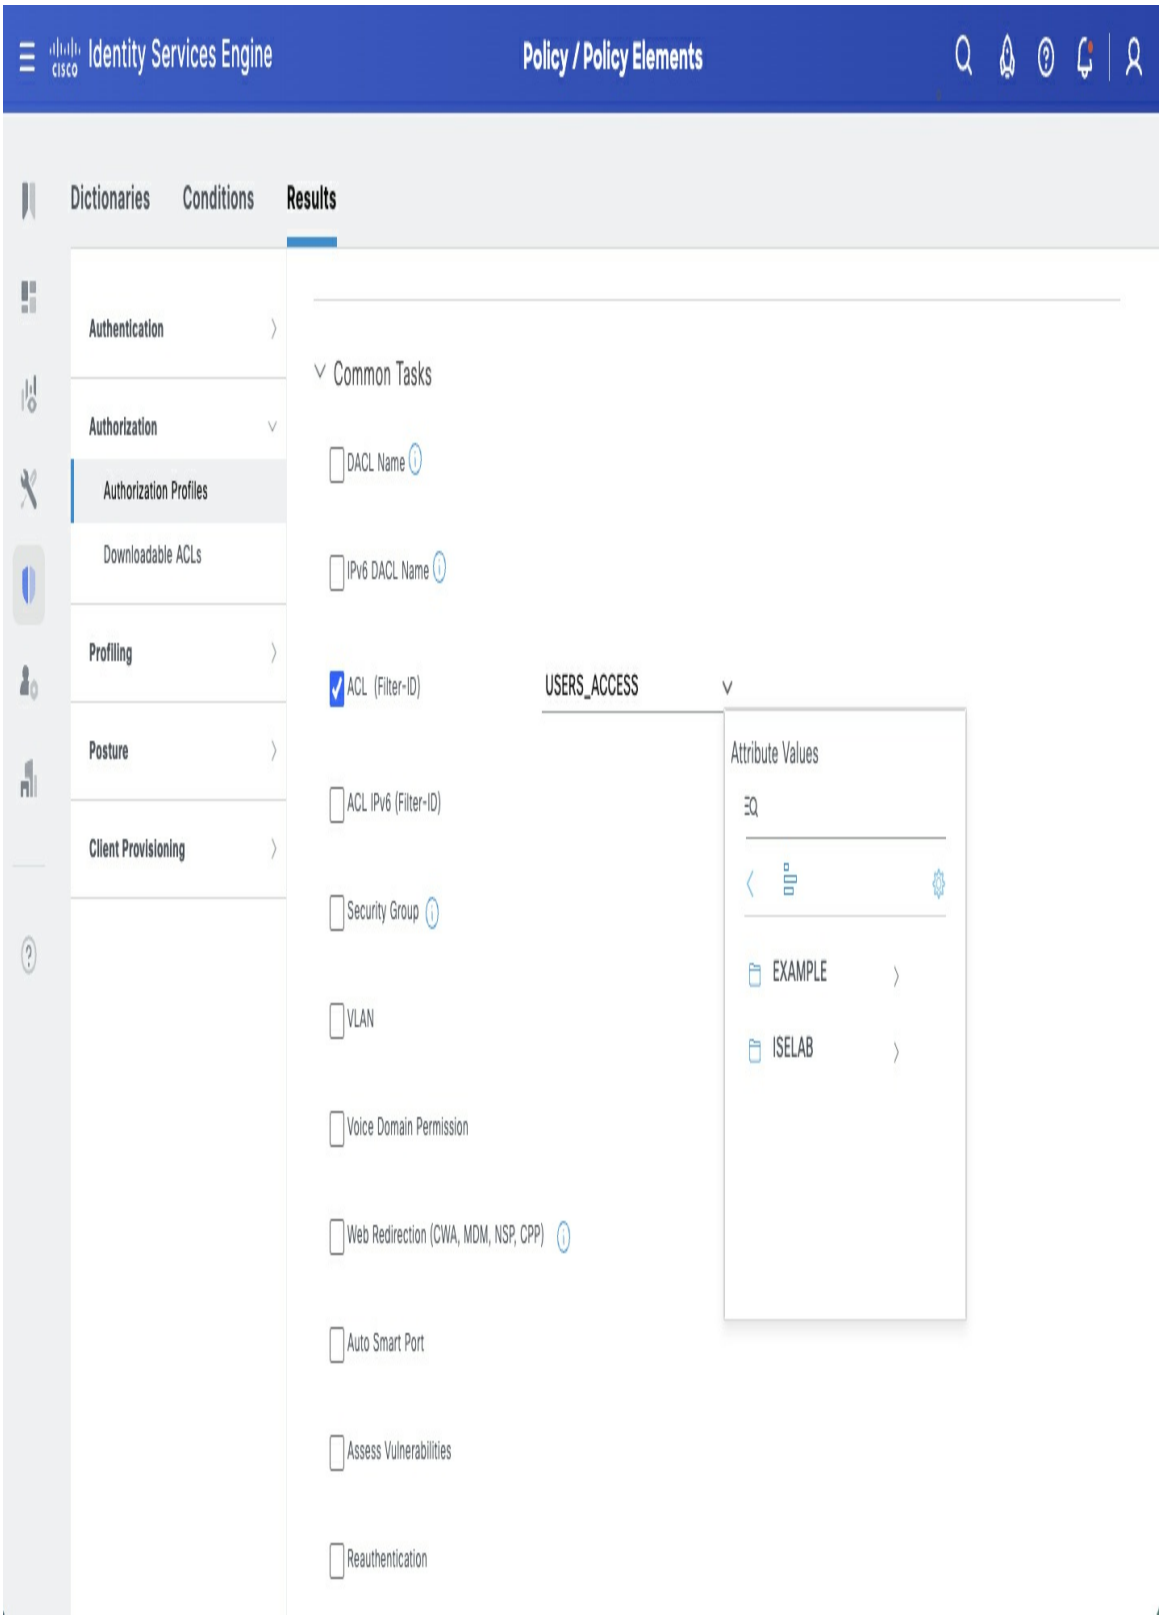


Figure 6-18 *Authorization Profile ACL (Filter-ID) Configuration*

Technet24

You can configure a security group directly in authorization policies, as shown in Figure 6-19.



Figure 6-19 *SGT Assignment via Authorization Policy Results*

However, you gain more flexibility when you configure this security group through authorization profiles. Once you check **Security Group**, you can select the SGT from the drop-down menu and configure the **VLAN** and **Virtual Network** where the user or endpoint should be placed, as shown in Figure 6-20.

Note

If you select the **Security Group** checkbox, you must configure the **VLAN** through **Security Group**. The **VLAN** option outside of the SGT becomes grayed out. This avoids configuration conflicts and prevents you from defining **VLAN** in two places, which could lead to potential conflicts. SGTs are covered in more detail in later chapters.

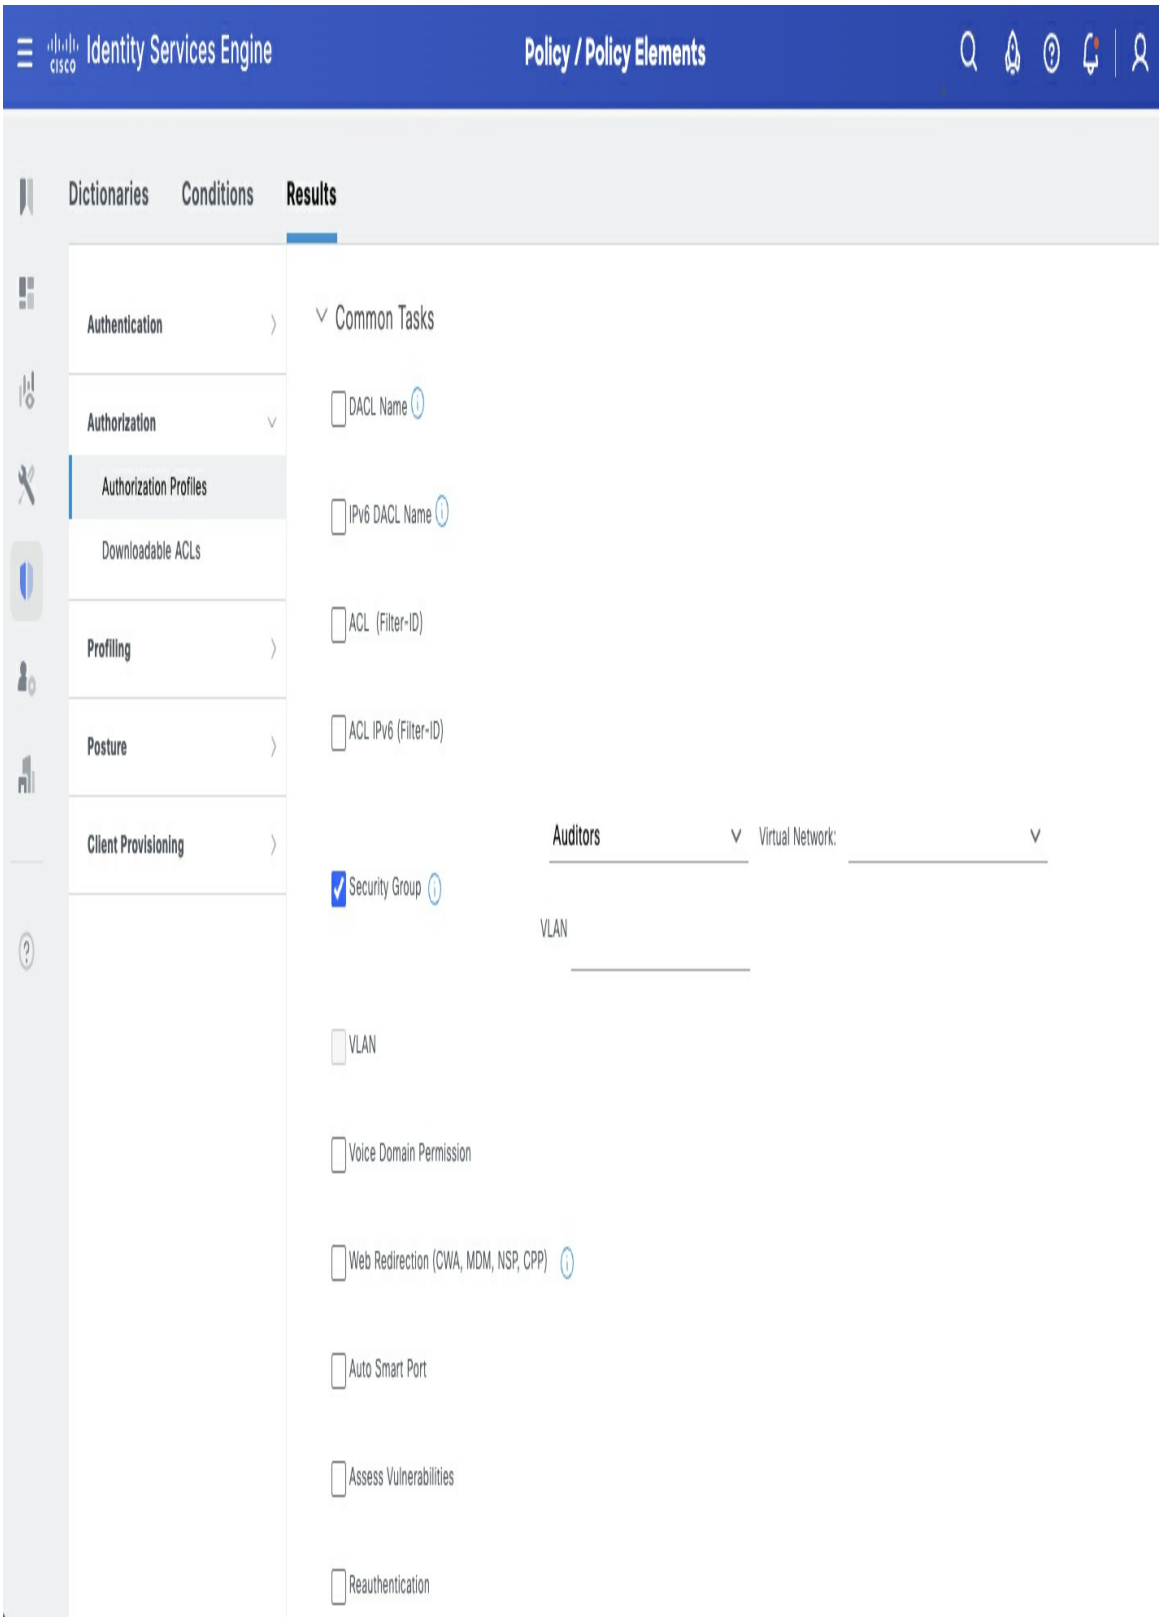


Figure 6-20 *Authorization Profile Security Group Configuration*

Technet24

VLAN selection allows you to configure dynamic VLAN assignment and place users or endpoints in different VLANs based on specific conditions. On ISE, select the **VLAN** checkbox and provide the **VLAN** name or **ID**, as shown in Figure 6-21.

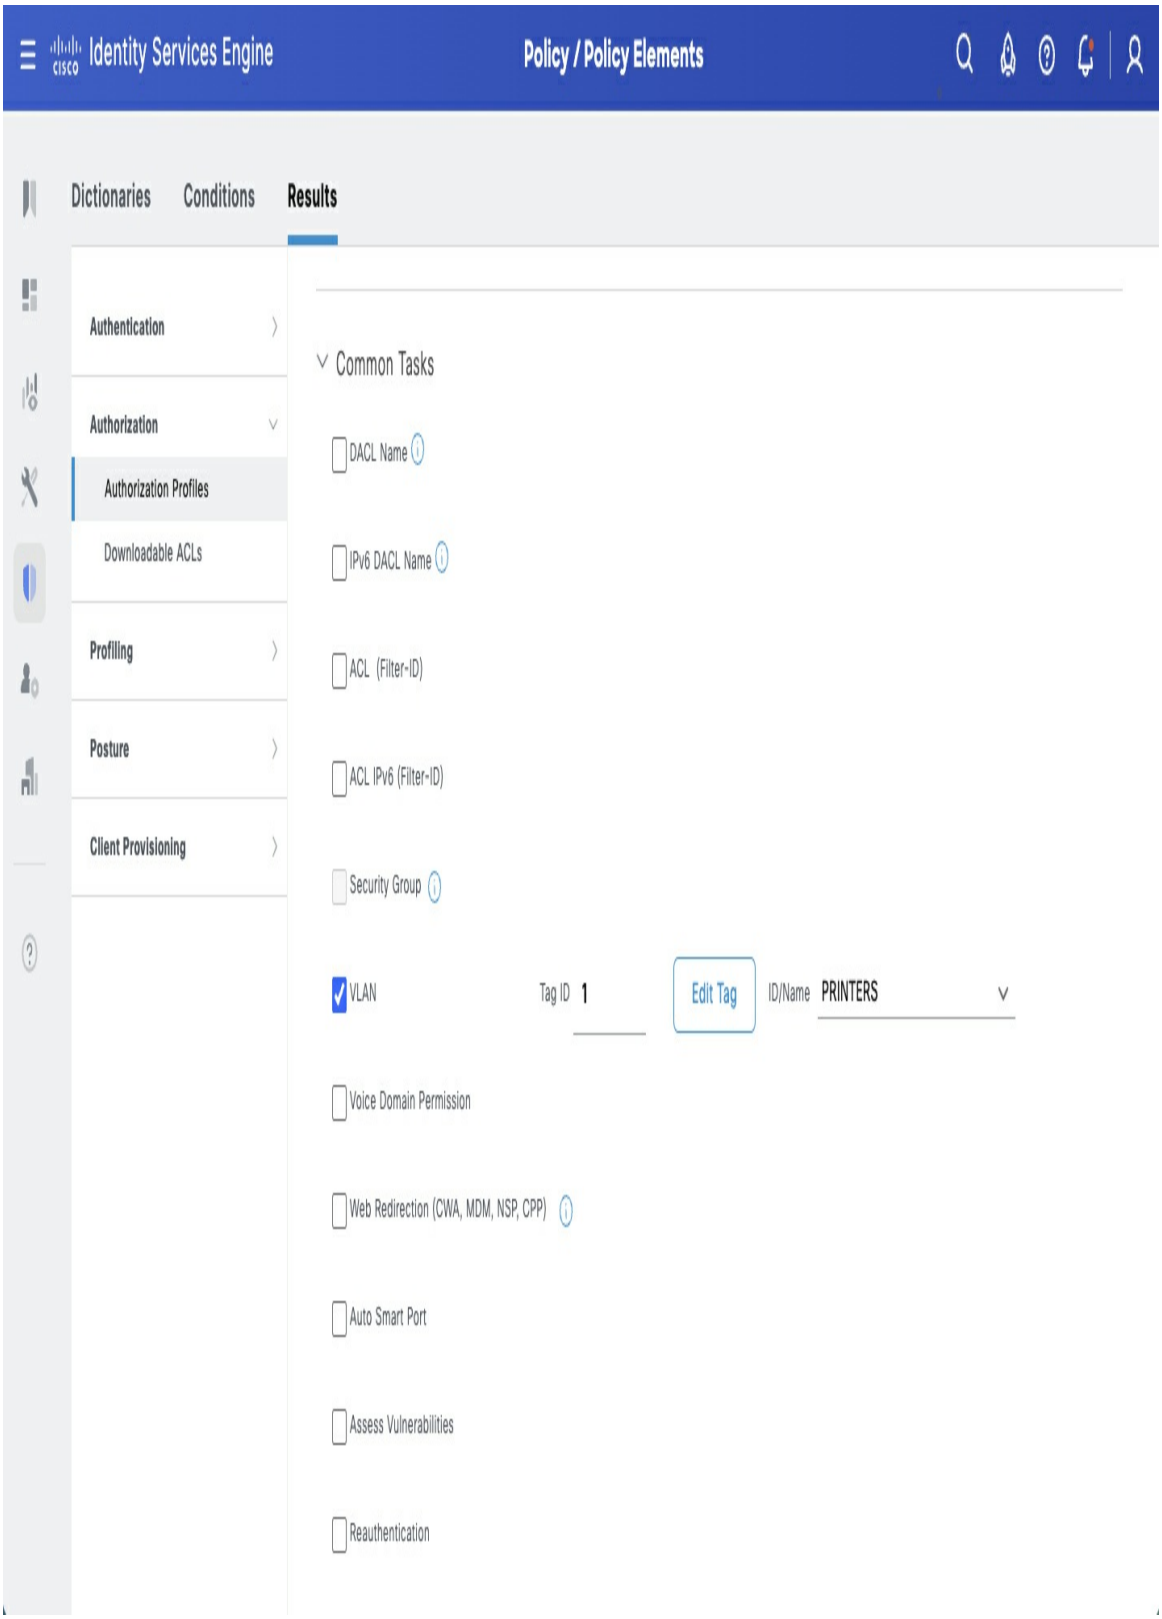


Figure 6-21 *Authorization Profile VLAN Configuration*

Technet24

It is crucial to periodically reauthenticate endpoints, and the most scalable way to do this is from ISE. You can configure the reauthentication timer and choose how to maintain connectivity during reauthentication, as shown in Figure 6-22.

Best Practice:

Configure endpoint reauthentication to 24 hours (86400 seconds). You can set a lower interval, but it should not be less than 2 hours (7200 seconds). You should always check the values allowed at the network access device level.

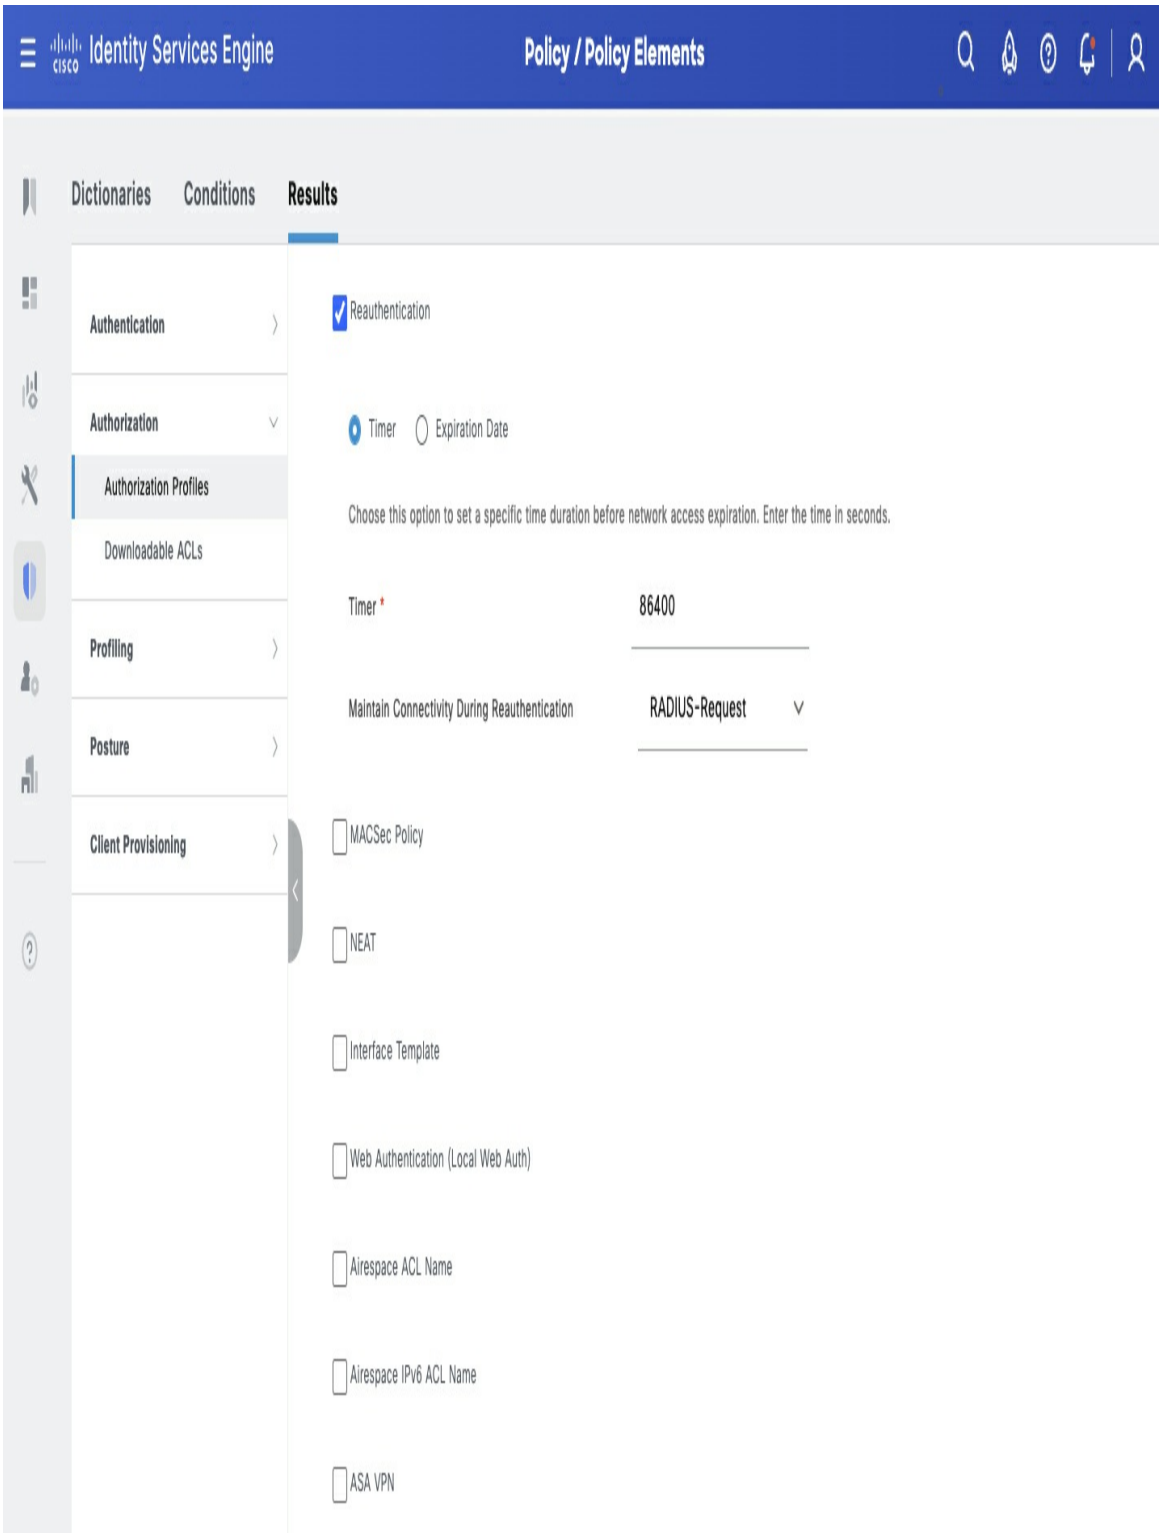


Figure 6-22 *Authorization Profile Reauthentication Configuration*

If you want to configure settings outside of the Common Tasks, use

Technet24

the **Advanced Attributes Settings**. You can use this section to send specific Cisco AV pairs or to apply third-party vendor dictionaries, as shown in Figure 6-23.

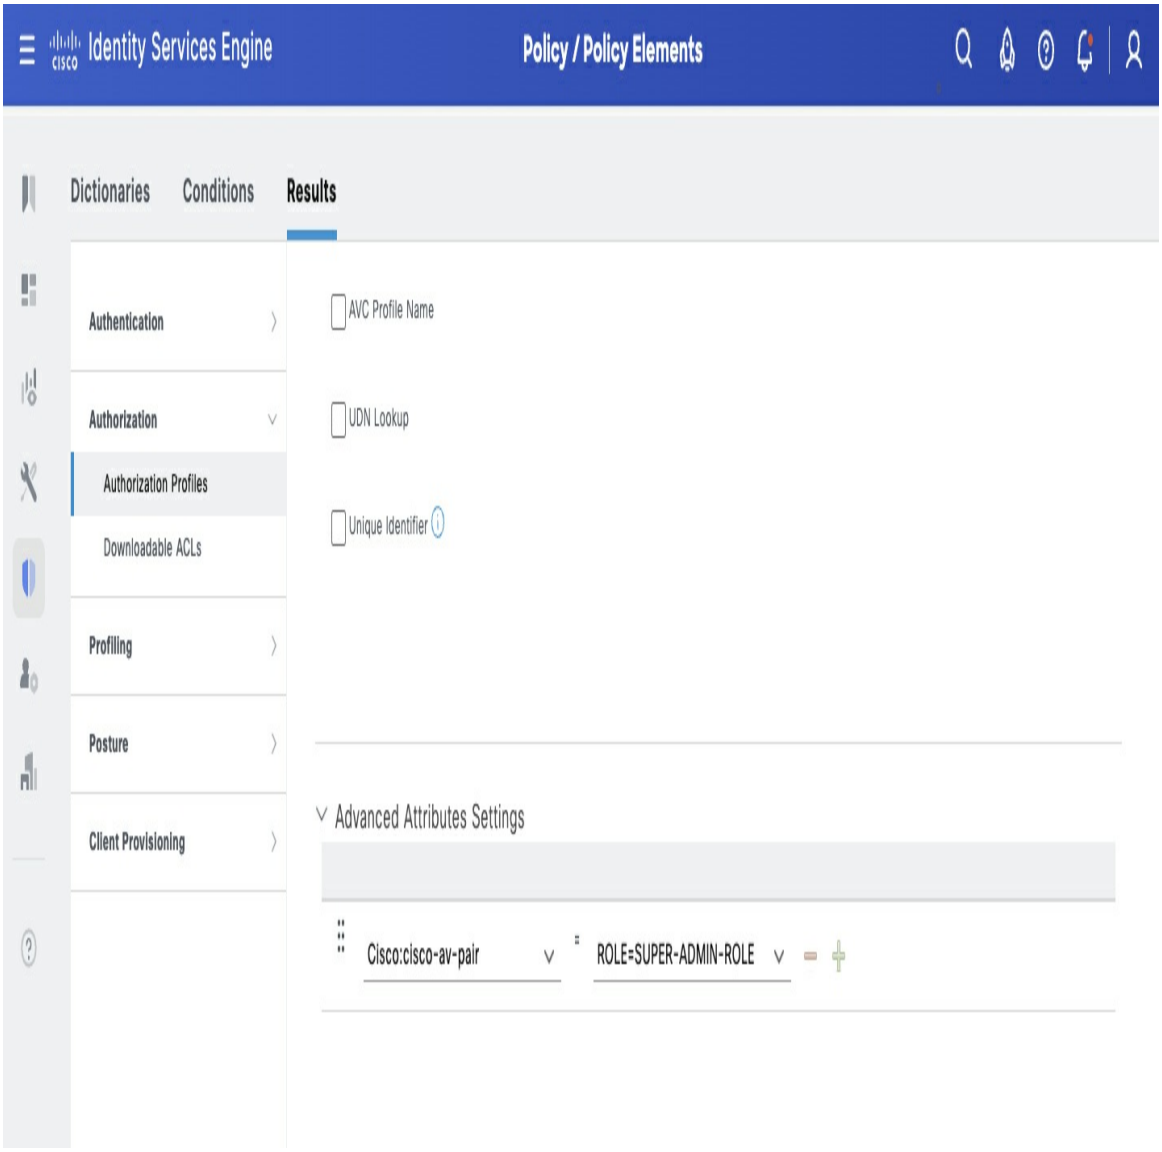


Figure 6-23 *Authorization Profile Advanced Attributes Settings*

Finally, the **Attributes Details** section contains the raw output of what ISE will send back in the RADIUS packet, as shown in Figure 6-24.



Figure 6-24 *Authorization Profile Attribute Details*

URL Redirection

One crucial concept to understand with ISE is URL redirection. You use URL redirection in multiple ISE features, including Guest, BYOD, and Posture. The goal of URL redirection with ISE is to redirect the user to one of the ISE portals, depending on your configuration.

As a result of initial authentication, ISE returns two Cisco AV pairs:

- url-redirect
- url-redirect-acl

Url-redirect controls where you should redirect the user. Url-redirect-acl specifies which ACL on the network access device you should use to control which traffic will be redirected. You can configure **Web Redirection** in the **Common Tasks** section, as shown in Figure 6-25.

Technet24



Figure 6-25 *Authorization Profile Web Redirection Configuration*

As a result, the Attributes Details are shown in Figure 6-26. One confusing aspect is that ISE displays ip:port in the Redirect URL. In practice, ISE substitutes the IP with the fully qualified domain name (FQDN) of the PSN that handled the initial authentication. If PSN1 terminated the RADIUS authentication and triggered the redirection, ISE returns the FQDN of PSN1. If PSN2 terminated the RADIUS authentication and triggered the redirection, ISE returns the FQDN of PSN2.



Figure 6-26 *Web Redirection Authorization Profile Attribute*

Details

You can also configure a static IP or FQDN instead of exposing the PSN's FQDNs, as shown in Figure 6-27.



Figure 6-27 *Authorization Profile Web Redirection Static IP/Hostname/FQDN Configuration*

However, you need to be extra careful with this configuration and ensure you have

- Correct certificates that support the custom FQDN
- Correct load balancing set up to allow the same PSN that handled RADIUS authentication to handle guest redirection

Figure 6-28 shows how redirection works on the network device itself.

Technet24

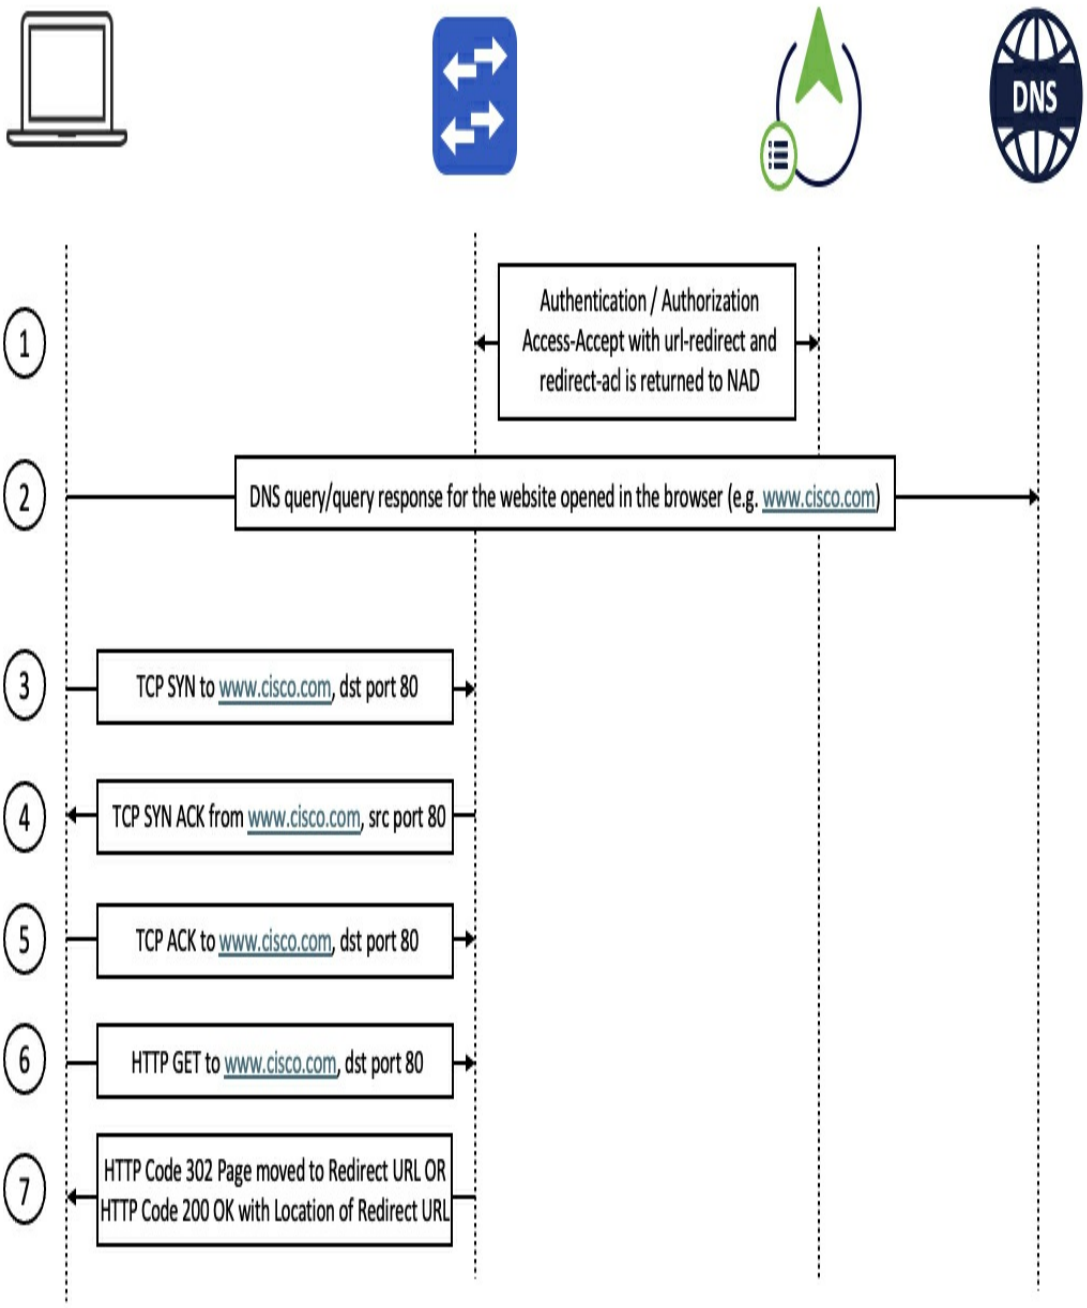


Figure 6-28 *URL Redirection Flow on Cisco NAD*

The NAD intercepts the three-way TCP handshake and responds on behalf of the server. Once the handshake is complete, the NAD redirects the client browser to the URL specified in the redirect, using

HTTP status code 200 or 302, depending on the implementation.

It is crucial to configure Redirect ACLs properly on network access devices. If you configure a Catalyst Switch or a 9800 WLC, you need to deny access from and to the PSNs and deny DNS traffic as well. This way, you ensure that DNS packets and packets to ISE nodes are allowed without redirection. The following output shows the ACL configured on a 9800 WLC, which denies traffic to ISE PSNs in the 10.48.26.0/24 network segment, denies DNS, and permits HTTP:

```
WLC-9800#show running-config | section access-list
ip access-list extended ACL_WEBAUTH_REDIRECT
10 deny ip any 10.48.26.0 0.0.0.255
20 deny ip 10.48.26.0 0.0.0.255 any
30 deny udp any any eq domain
40 deny udp any eq domain any
50 permit tcp any any eq www
WLC-9800#
```

A similar output shows the 9300 Catalyst Switch configuration, with two PSNs denied at 10.48.26.60 and 10.48.26.160:

```
SWITCH-9300#show running-config | section access-list
ip access-list extended ACL_WEBAUTH_REDIRECT
10 deny ip any host 10.48.26.60
20 deny ip any host 10.48.26.160
30 permit tcp any any eq www
SWITCH-9300#
```

DACL

Another powerful authorization option is the use of downloadable access control lists. ISE includes four DACLs that either allow or deny IPv4 and IPv6 traffic, as shown in Figure 6-29.

Technet24

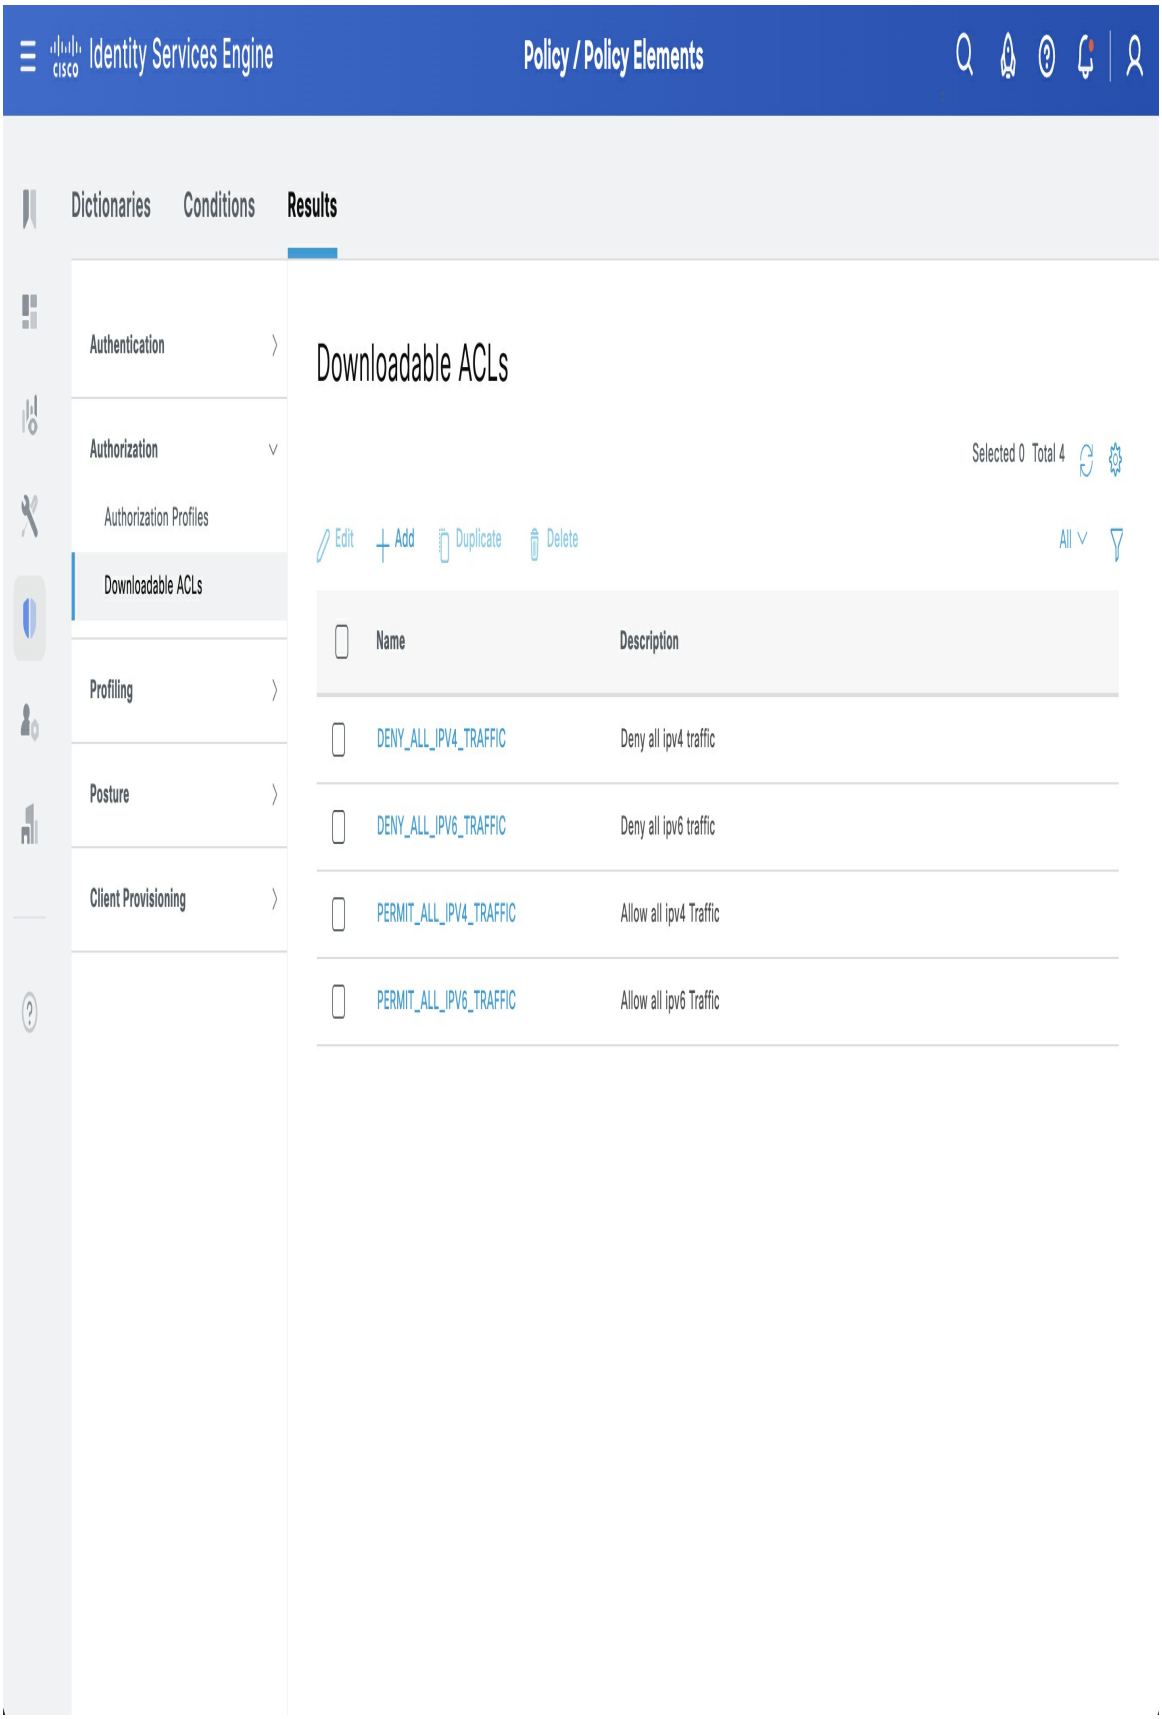

Identity Services Engine
Policy / Policy Elements
Dictionaries
Conditions
Results
Authentication
Authorization
Authorization Profiles
Downloadable ACLs
Profiling
Posture
Client Provisioning
Downloadable ACLs
Selected 0 Total 4
Edit
Add
Duplicate
Delete
All
Name
Description
DENY_ALL_IPV4_TRAFFIC
Deny all ipv4 traffic
DENY_ALL_IPV6_TRAFFIC
Deny all ipv6 traffic
PERMIT_ALL_IPV4_TRAFFIC
Allow all ipv4 Traffic
PERMIT_ALL_IPV6_TRAFFIC
Allow all ipv6 Traffic

Figure 6-29 *Built-in Downloadable ACLs*

To add a DACL, navigate to **Policy > Policy Elements > Authorization > Downloadable ACLs** and click **+Add**, as shown in Figure 6-30.

Technet24

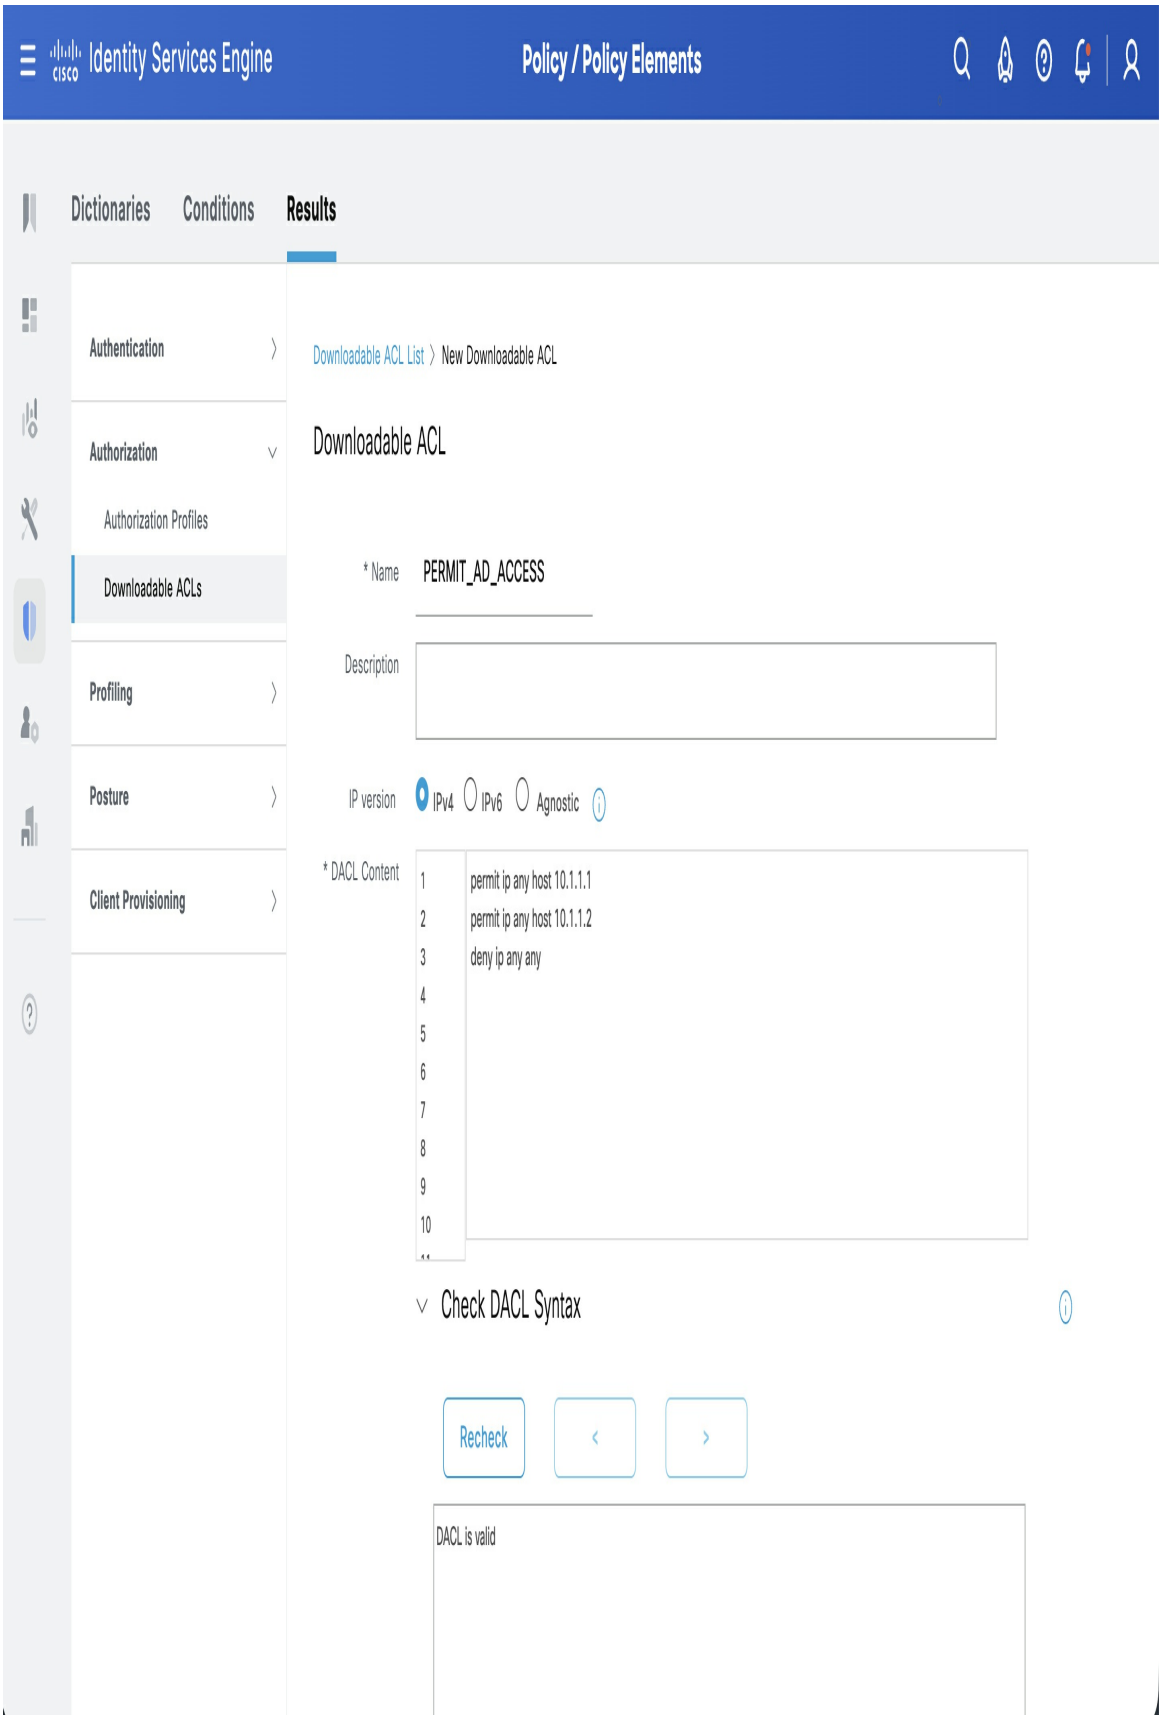

Identity Services Engine
Policy / Policy Elements
Dictionaries
Conditions
Results
Authentication
Authorization
Authorization Profiles
Downloadable ACLs
Profiling
Posture
Client Provisioning
Downloadable ACL List > New Downloadable ACL
Downloadable ACL
* Name
PERMIT_AD_ACCESS
Description
IP version
IPv4
IPv6
Agnostic
* DACL Content
1 permit ip any host 10.1.1.1
2 permit ip any host 10.1.1.2
3 deny ip any any
4
5
6
7
8
9
10
Check DACL Syntax
Recheck
DACL is valid

Figure 6-30 *Downloadable ACL Configuration*

After it is configured, you must add the DACL to the authorization profile and assign the authorization profile to the Results section of the authorization policy. When ISE returns the DACL during RADIUS authentication and authorization, only the name is included in Access-Accept. Subsequently, the NAD sends a follow-up request for the content of the DACL, as shown in Figure 6-31.

Technet24

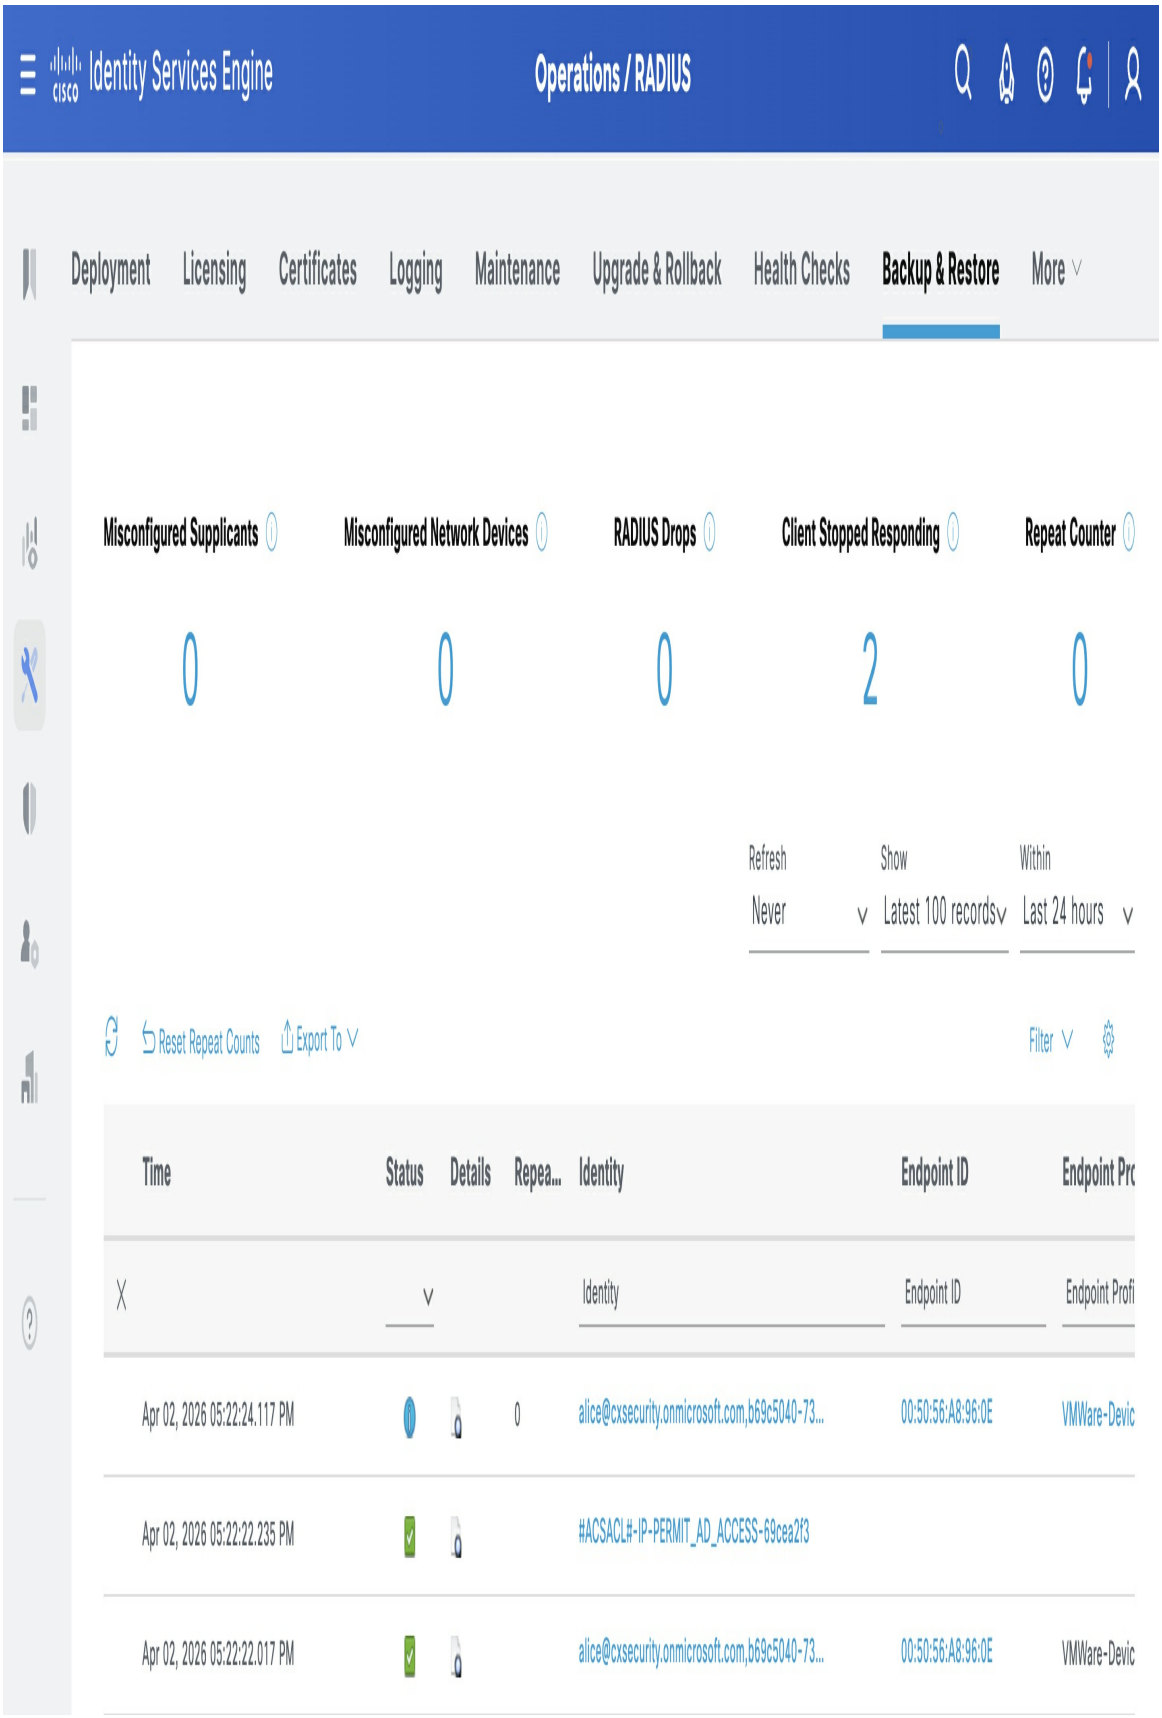

Identity Services Engine
Operations / RADIUS
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Misconfigured Supplicants
Misconfigured Network Devices
RADIUS Drops
Client Stopped Responding
Repeat Counter
0
0
0
2
0
Refresh
Never
Show
Latest 100 records
Within
Last 24 hours
Reset Repeat Counts
Export To
Filter
Time
Status
Details
Repea...
Identity
Endpoint ID
Endpoint Pro
Identity
Endpoint ID
Endpoint Profi
Apr 02, 2026 05:22:24.117 PM
0
alice@cxsecurity.onmicrosoft.com,b69c5040-73...
00:50:56:A8:96:0E
VMWare-Devic
Apr 02, 2026 05:22:22.235 PM
#ACSACL#-IP-PERMIT_AD_ACCESS-69cea2f3
Apr 02, 2026 05:22:22.017 PM
alice@cxsecurity.onmicrosoft.com,b69c5040-73...
00:50:56:A8:96:0E
VMWare-Devic

Figure 6-31 *Live Logs DACL Flow*

The Detailed Authentication Report shows the successful DACL download event, as shown in Figure 6-32.

Technet24

Cisco ISE

Overview

Event	5232 DACL Download Succeeded
Username	#ACSACL#-IP-PERMIT_AD_ACCESS-69cea2f3
Endpoint Id	
Endpoint Profile	
Authorization Result	

Steps

Step ID	Description	Latency (ms)
11001	Received RADIUS Access-Request	
11017	RADIUS created a new session	0
11117	Generated a new session ID	1
11002	Returned RADIUS Access-Accept	0

Authentication Details

Source Timestamp	Apr 02, 2026 05:22:22.000 PM
Received Timestamp	Apr 02, 2026 05:22:22.000 PM
Policy Server	ISE35-1ek
Event	5232 DACL Download Succeeded
Failure Reason	
Username	#ACSACL#-IP-PERMIT_AD_ACCESS-69cea2f3

Figure 6-32 *DACL Detailed Authentication Report Overview*

The Result portion of this report contains the content of the DACL, as shown in Figure 6-33.

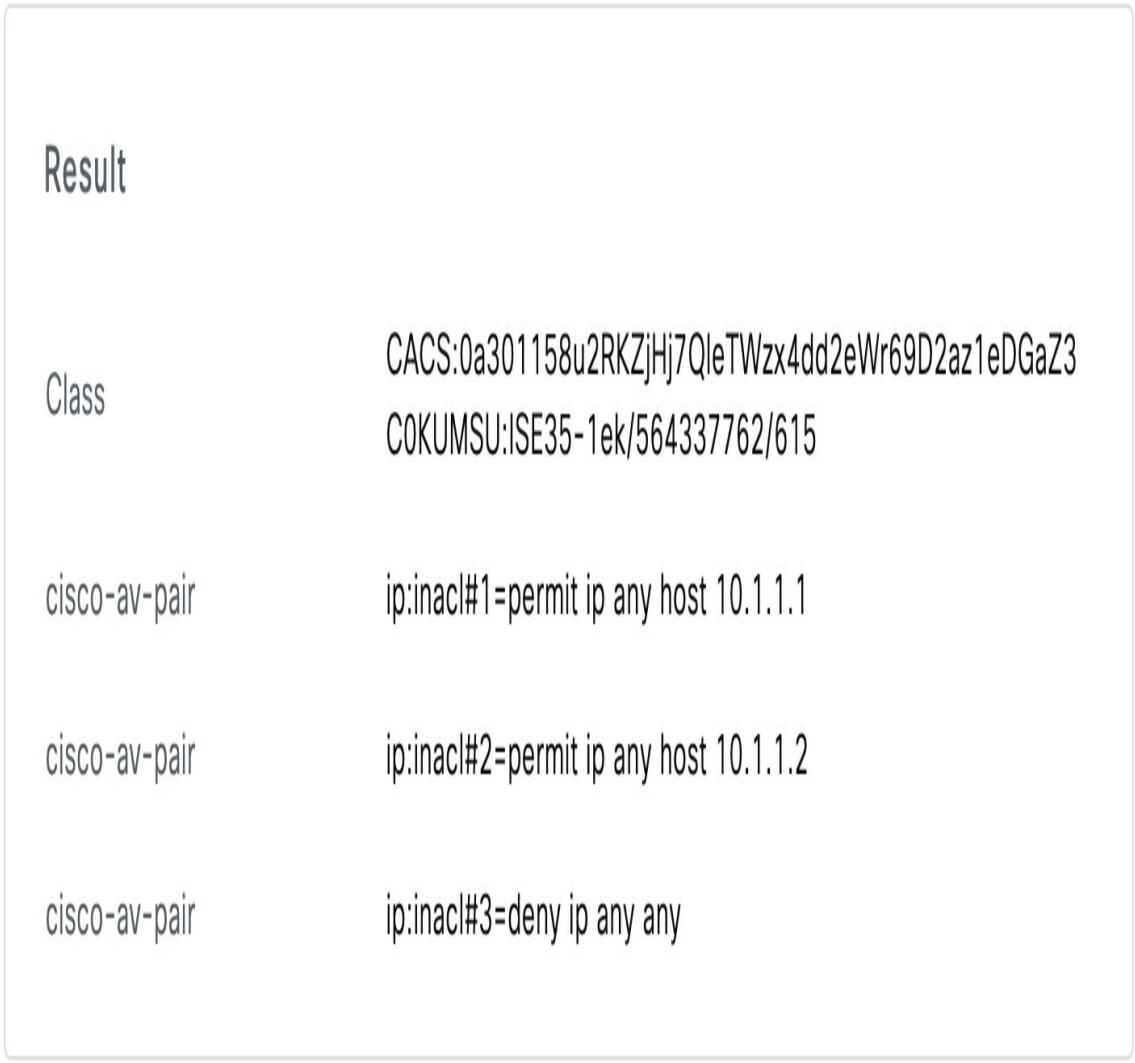


Figure 6-33 *DACL Detailed Authentication Report Results*

Network Device Profiles

ISE natively supports third-party network devices through the Network Device Profiles functionality.

Technet24

You can create your own profiles, import or export them, or use the Cisco Communities page to access profiles developed through community efforts. To navigate to Network Device Profiles, go to **Administration > Network Resources > Network Device Profiles**, as shown in Figure 6-34.

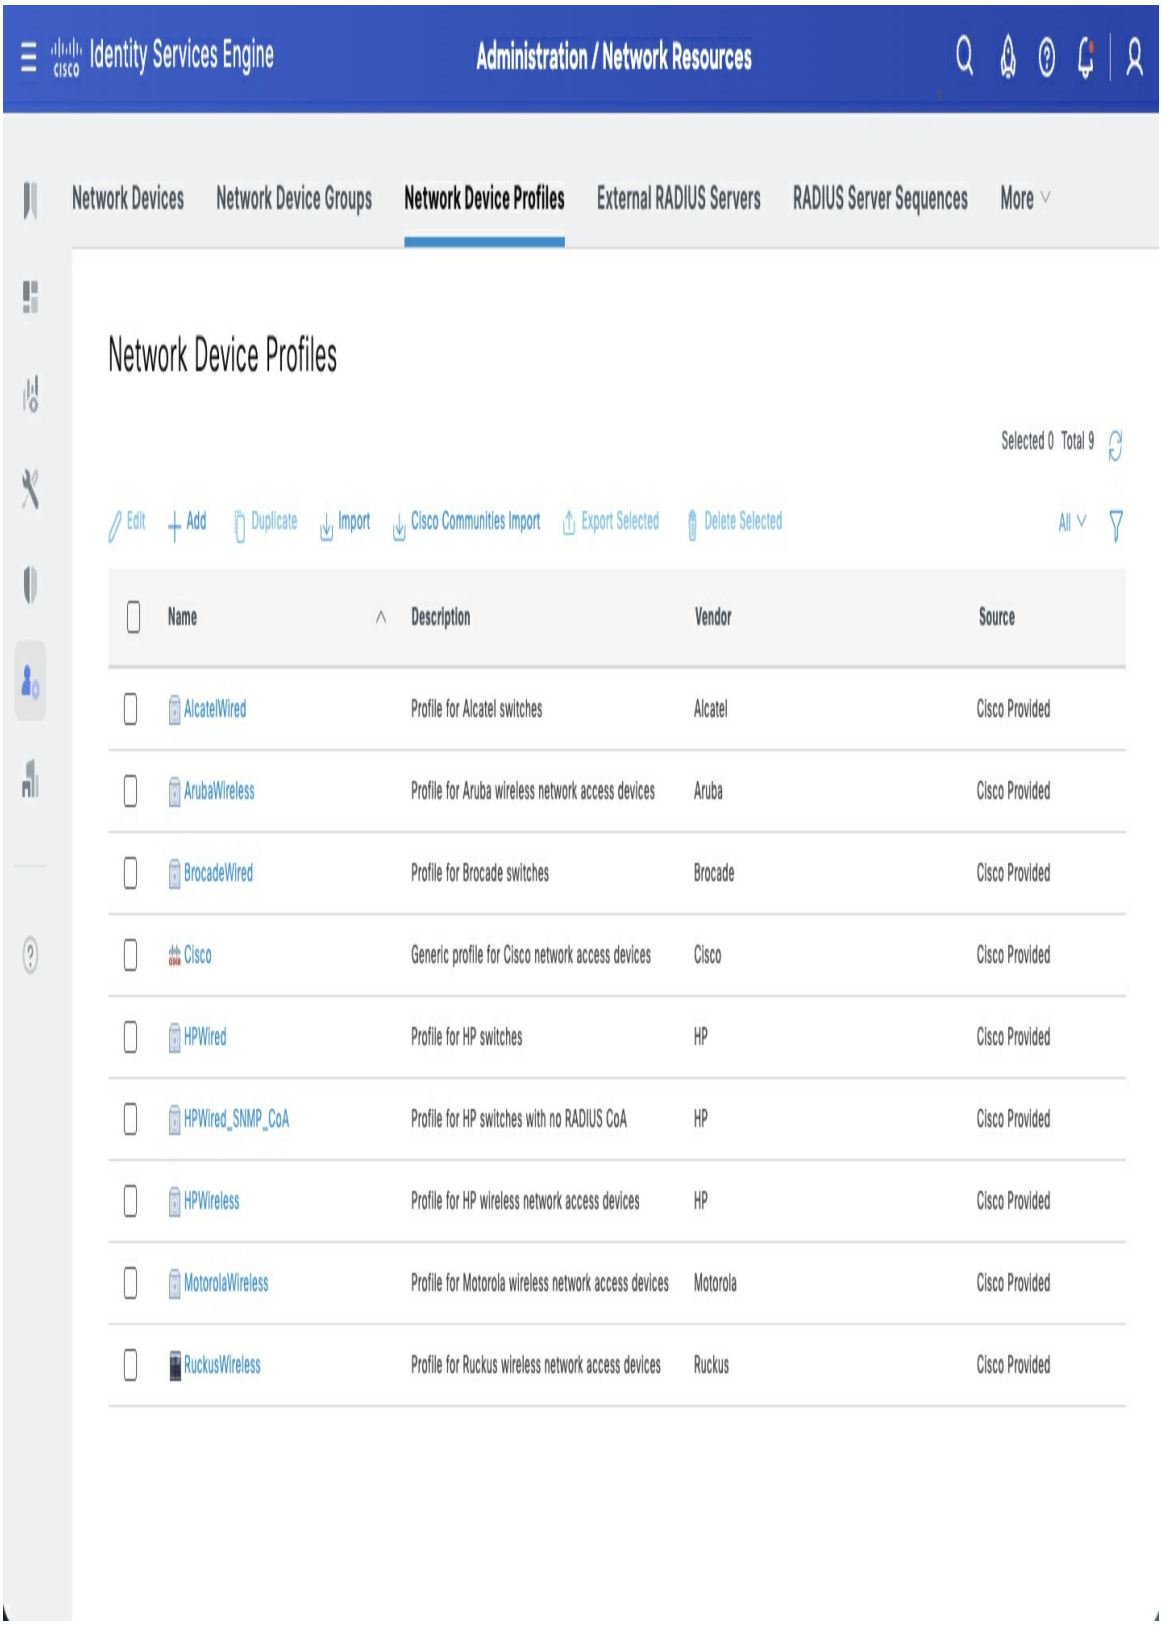


Figure 6-34 *Built-in Network Device Profiles*

Technet24

Network device profiles help with the following:

- Define the protocols that the network device supports, along with RADIUS dictionary associations, as shown in Figure 6-35.

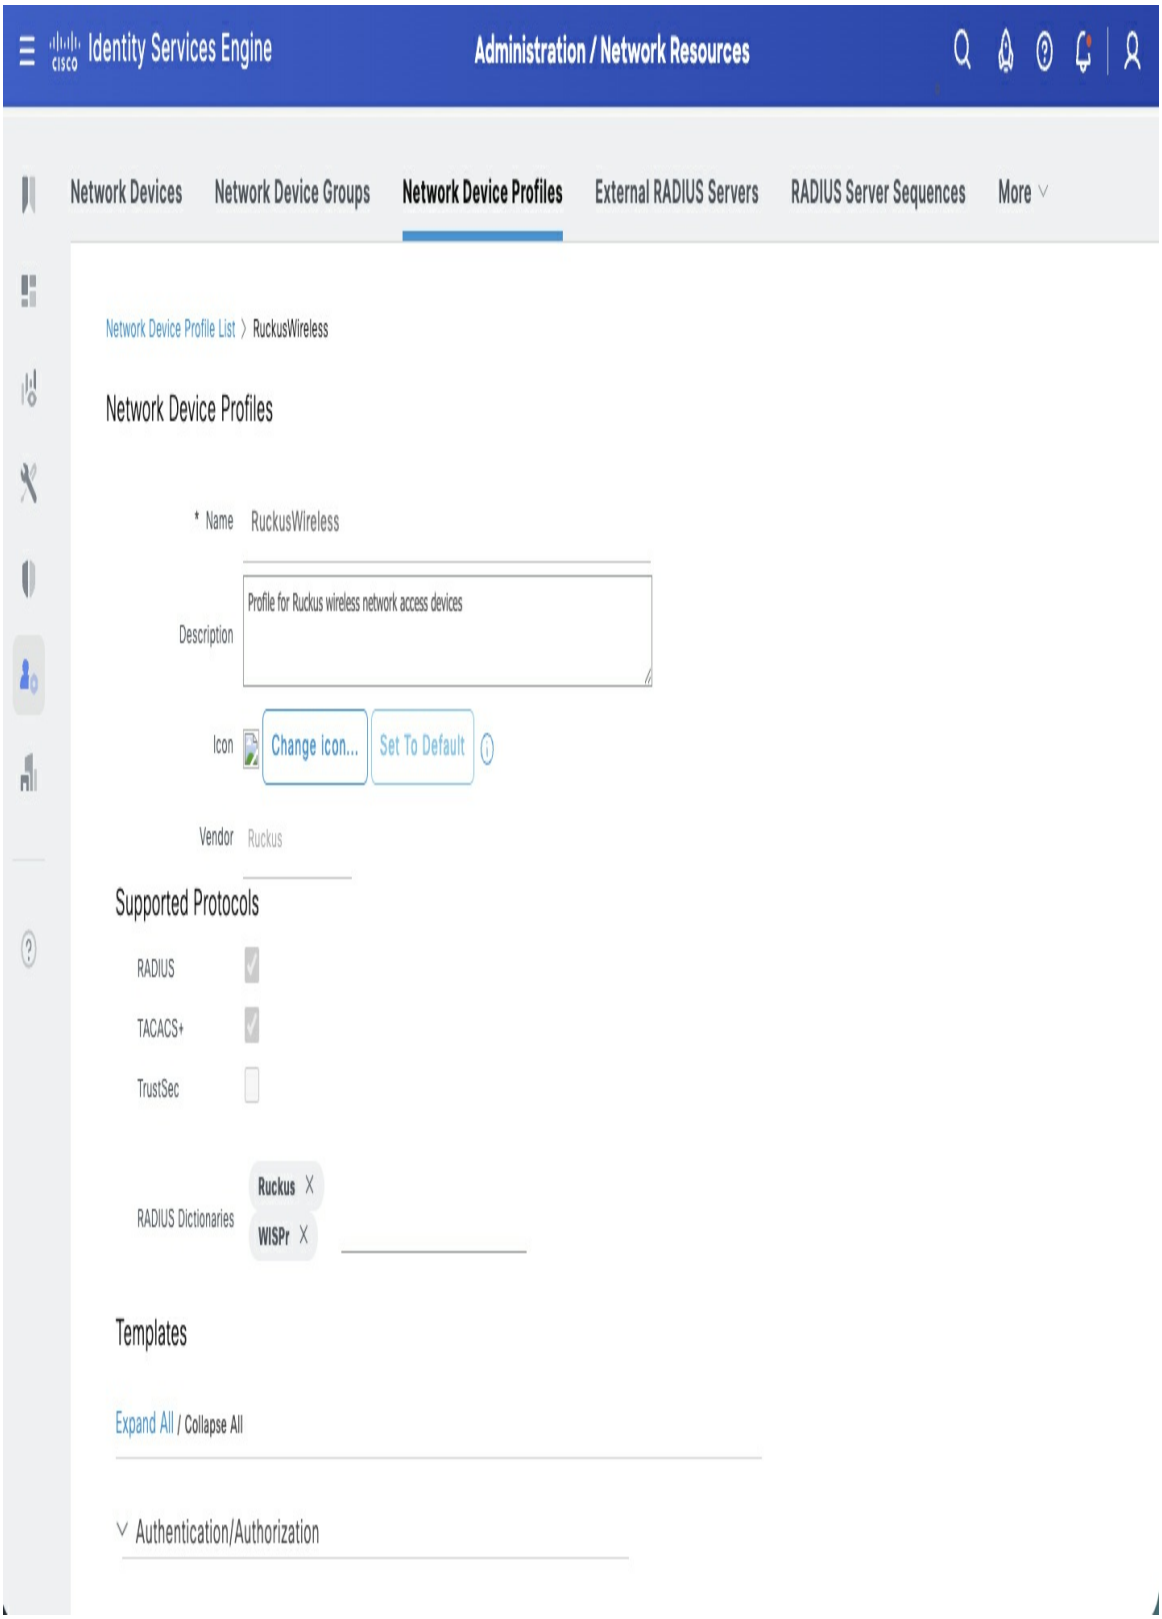


Figure 6-35 *Network Device Profile Configuration*

Technet24

- Identify dynamically the type of authentication flow the third-party device uses. For example, with the same condition for wireless MAB, ISE can identify different vendors performing MAB authentication, even if they use different attributes to identify a MAB flow. This is accomplished through the **Flow Type Conditions** section, as shown in Figure 6-36.

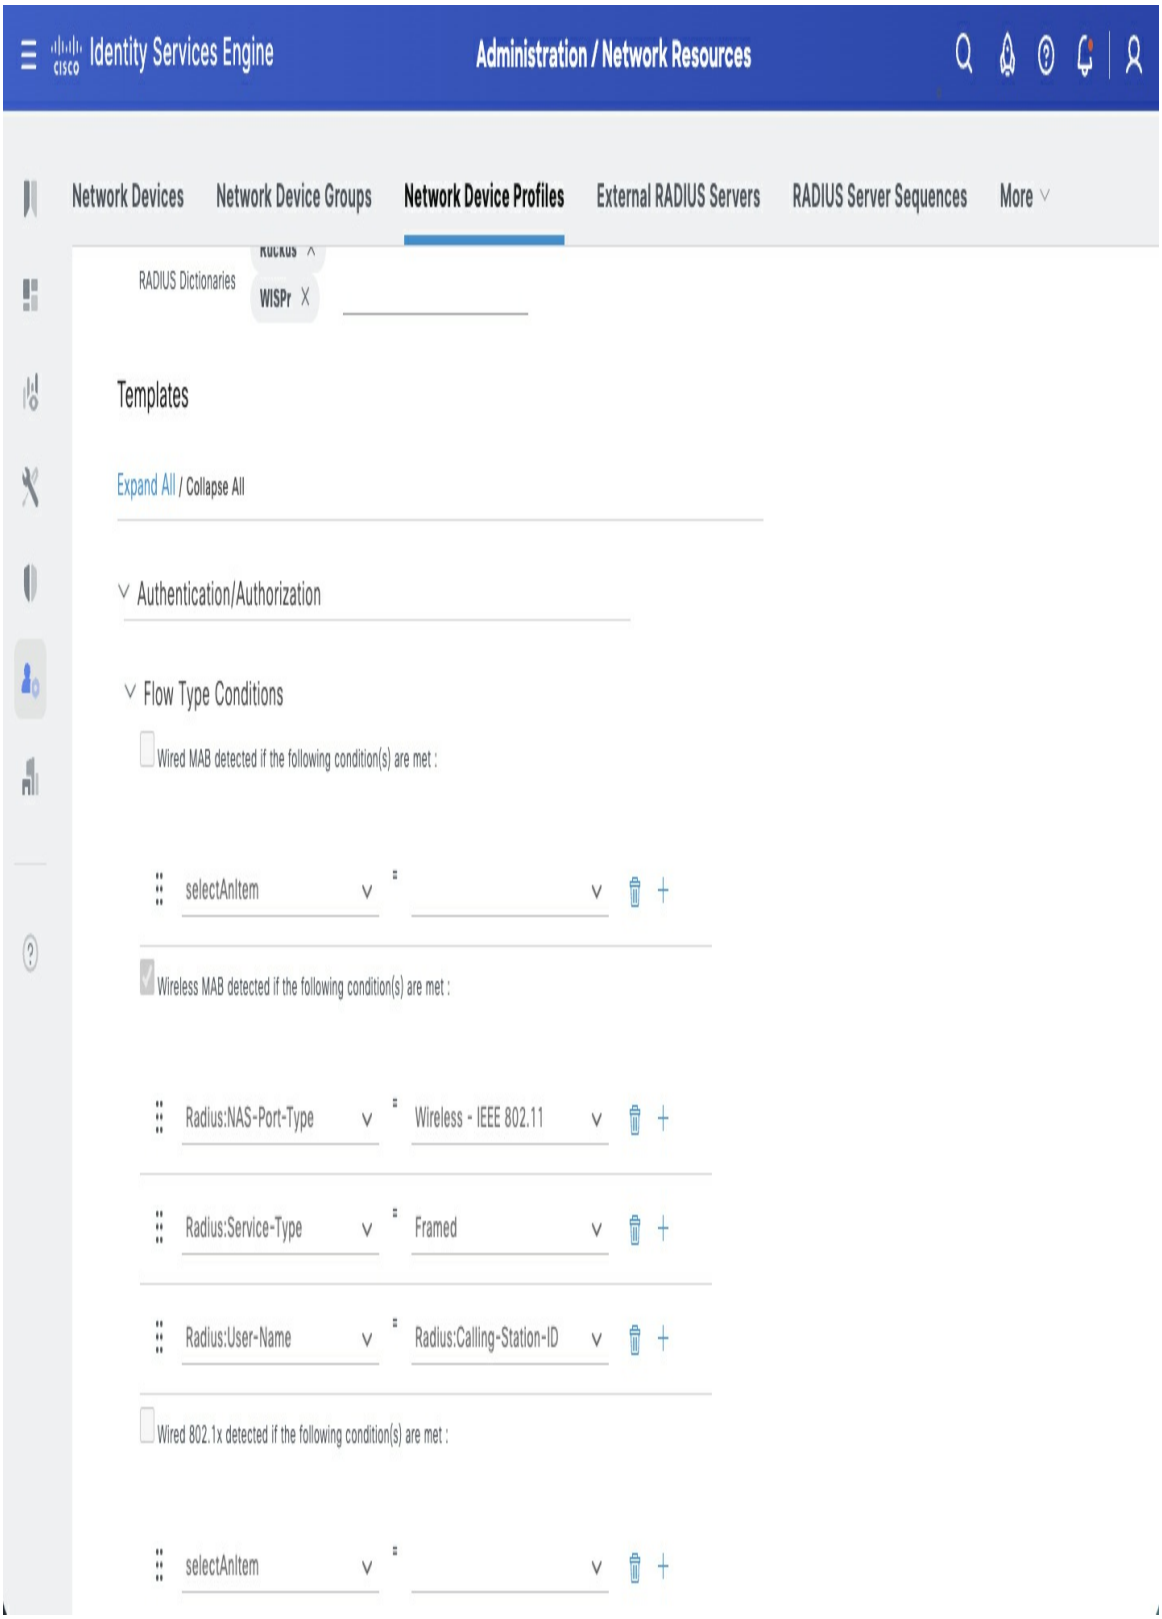


Figure 6-36 *Network Device Profile Flow Type Conditions*

Technet24

- Identify which CoA packets are sent for specific CoA actions. Different vendors might use different ports for CoA and may include additional attributes within the packets. The network device profile allows you to dynamically select the correct CoA type based on the network device type communicating with ISE. CoA options are shown in Figure 6-37.

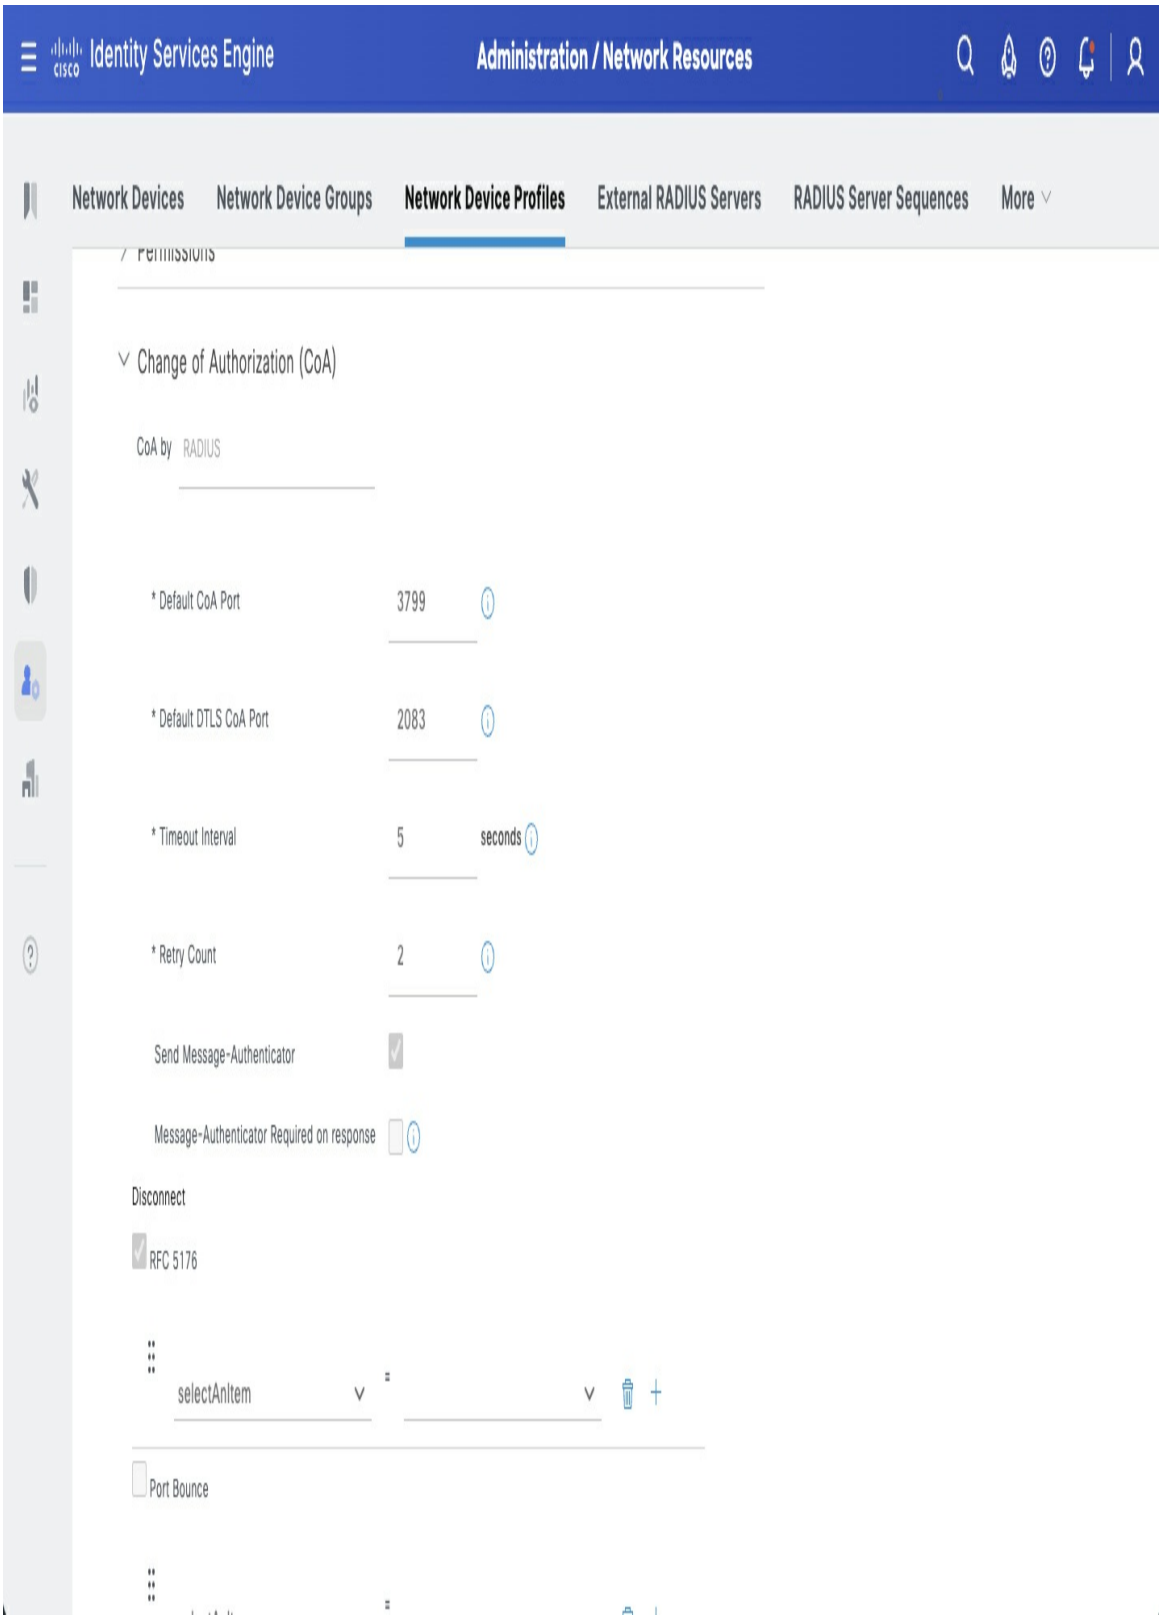


Figure 6-37 *Network Device Profile Change of Authorization*

Technet24

Configuration

Additionally, the network device profile helps you set the correct VLAN and ACL permissions used by different vendors and apply the appropriate URL redirection technique based on vendor requirements.

All network access devices must have a network device profile selected; the default profile is Cisco. To select the matching profile from Network Device Settings, navigate to **Administration > Network Resources > Network Devices**, as shown in Figure 6-38.

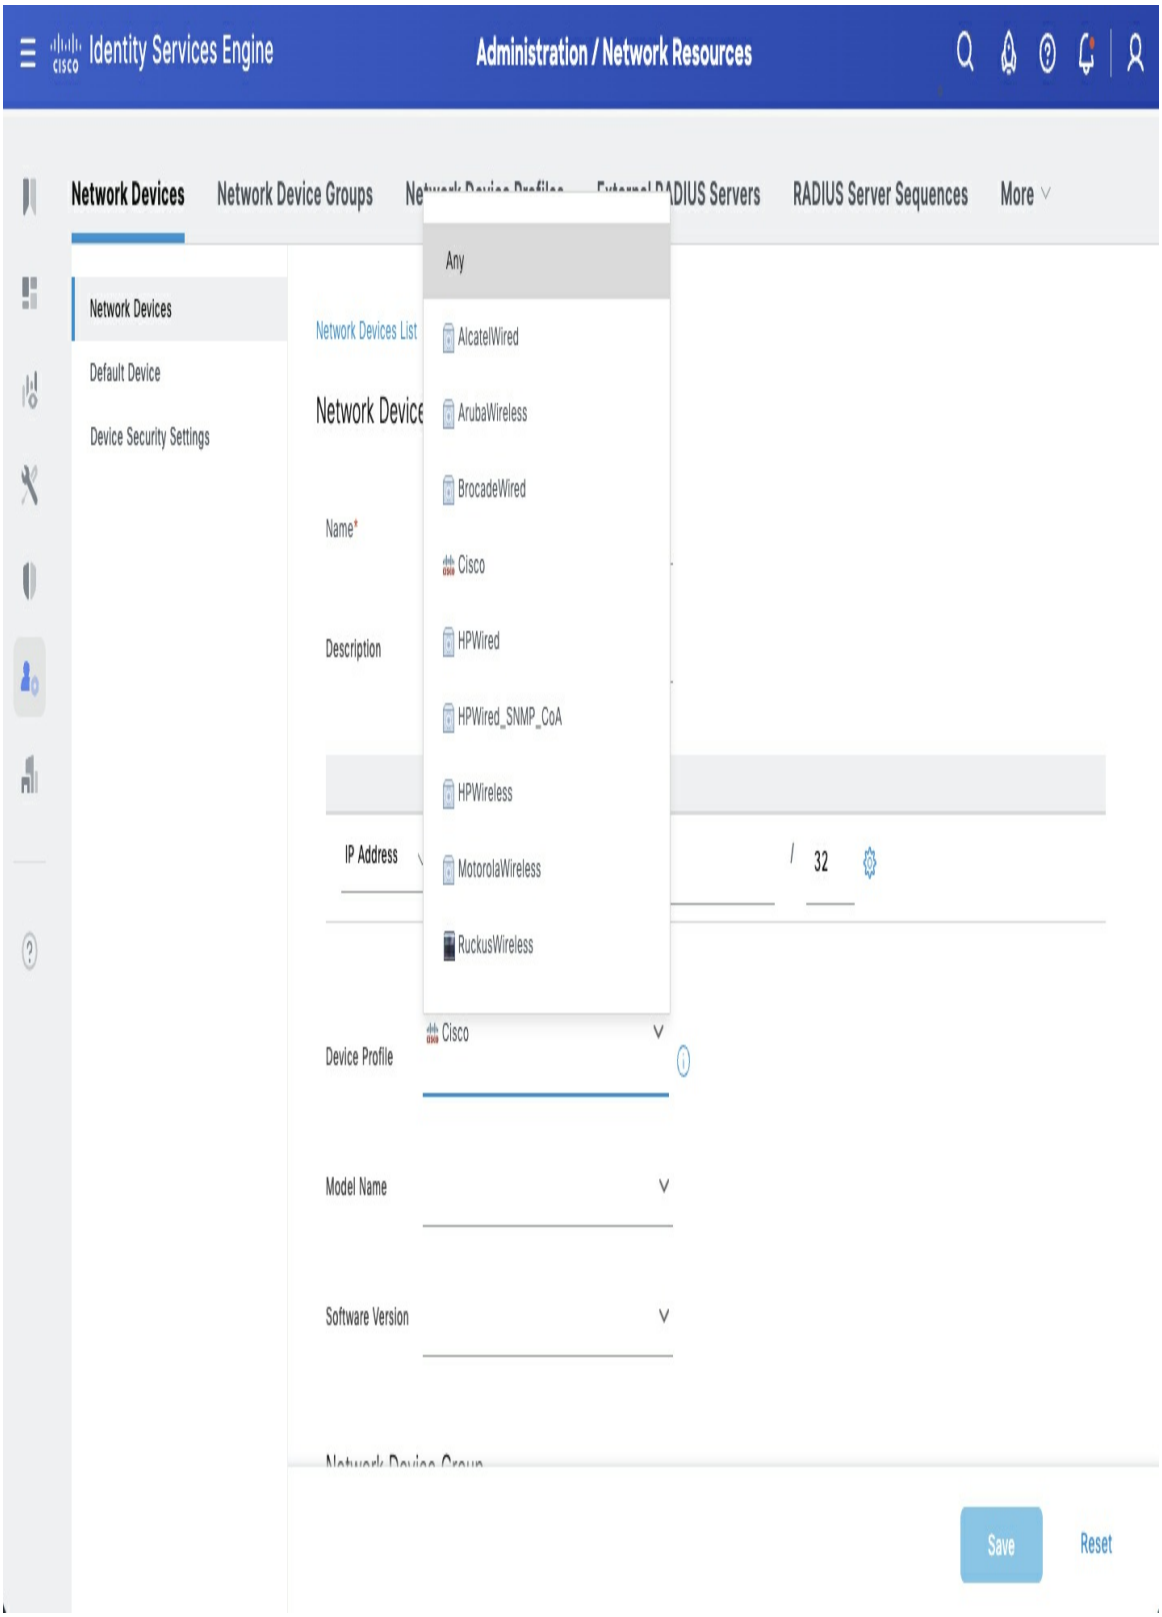


Figure 6-38 *Network Device Profile Assignment*

Technet24

Summary

Authorization (along with authorization profiles) is one of the most crucial parts of ISE configuration. It allows you to configure granular access and assign only the permissions required for users and endpoints. Authorization profiles offer great flexibility and can be built based on various internal and external conditions. The results of authorization profiles may include dynamic VLANs, DACLs, Redirect ACLs, Redirect URLs, and many other options.

Chapter 7. Introduction to TrustSec

Introduction

Cisco coined the term *TrustSec* in the first decade of the 21st century. This term covers the simplification and management of security policy using groups to classify network traffic and enforce permissions. Traffic classification is based on endpoint identity and context, not IP address, enabling policy change without network redesign. A centralized policy management platform gathers advanced contextual data about who and what is accessing the network; uses security group tags (SGTs) to define roles and access rights; and then pushes the associated policy to network devices such as switches, routers, and security platforms. This approach provides better visibility through richer contextual information and allows an organization to be better able to provide least-privilege access and segmentation, reducing the impact and costs associated with a potential breach.

TrustSec can be seen as a software-defined segmentation technology that is embedded within network switches, routers, wireless LAN infrastructure, and firewalls. The technology consists of three primary concepts: classification, propagation, and enforcement. When users connect to the network, they are authenticated using various network access methods, and then authorization follows, which entails classifying the user's traffic at network ingress, leveraging rich contextual information such as identity, group membership, location,

Technet24

or access type, for example. After the user's traffic is classified, network devices propagate the classification information toward a network device assigned to be an enforcement point, typically at network egress, where dynamically downloaded policies dictate whether the traffic should be permitted or denied.

Security groups are well supported across the Cisco enterprise portfolio of network devices. The name of the technology may vary between products — for example, the Meraki portfolio uses Adaptive Policy as a descriptor — but the mechanisms are the same.

Legacy network devices also support security groups, but the embedded application-specific integrated circuit (ASIC) dictates the level at which support is possible. Products such as the 3750x, 3560x, 3560CX, and earlier IE platforms use an ASIC that was built around features being bound and delivered by MAC/Port/VLAN, which has implications on classification and enforcement. From a classification point of view, the main limitation is that the MAC address of the endpoint needs to be learned as well as the IP address (the IP Device Tracking entry is required). These platforms use what is commonly referred to as a port-based architecture. This means that endpoints needing classification must be Layer 2 adjacent to the switch. There are also enforcement limitations on these platforms (such as bindings cannot be used by an SXP listener for enforcement, and only one SGT is supported on a port and VLAN combination).

Later platforms like the 3650, 3850, Catalyst 4000, Catalyst 6000, Catalyst 9000 (9200, 9300, 9400, 9500, and 9600), and C-series platforms like the C9350, C9550, and C9610 (utilizing Silicon One ASICs) use an IP-based architecture and do not exhibit the same limits as the products requiring Layer 2 adjacency. When discussing network device operation in this chapter, we will describe only the current hardware and software operation.

Classification

TrustSec classification is essentially the process of assigning a security group tag to traffic as it enters the network, which allows security policies to be enforced based on the specific identity of the user or device rather than relying on traditional, less flexible IP addresses.

Classification Introduction

You learned about TrustSec classification at a high level in the Introduction. The idea is to assign tags to groups of users, endpoints, servers, and services and use that to abstract the security policy from the static network topology. This approach offers flexibility, simplicity, and efficiency, and provides intuitive networking. It is intuitive because policies are created using role names that are meaningful to administrators; for example, it is much more meaningful building and managing policies using group names like Employees, Guests, Confidential_Servers, and Building_Control_Systems rather than a bunch of IP addresses, ports, and protocols.

Classification is the process of assigning security group tags to IP addresses, and network elements require visibility into this classification to produce an end-to-end security solution. IP addresses can be assigned to endpoints statically or dynamically (via DHCP) and can also be assigned into security groups statically or dynamically.

Classification and IP Address Learning

If the IP address is assigned dynamically on an IOS-XE device (via DHCP), then the network device where the actual classification takes place must track which IP is assigned to the endpoint. This tracking is accomplished using the IP Device Tracking (IPDT) feature. In many platforms this feature is now called Switch Integrated Security Feature

Technet24

(SISF) to have a common way of tracking IPv4 and IPv6 addresses in the system. Traffic received by the network device is snooped to extract the assigned IP address and related MAC address and is stored in a binding table for use by various features, including TrustSec. If the IP address is assigned statically, then snooping ARP requests determine the IP.

Sample SISF binding table command-line interface output is shown in Example 7-1.

Example 7-1 *IOS-XE SISF Binding Table*

```
Switch#show device-tracking database
Binding Table has 3 entries, 2 dynamic (limit 100000)
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Addr
DH4 - IPv4 DHCP, DH6 - IPv6 DHCP, PKT - Other Packet, API - APId
Preflevel flags (prlvl):
0001:MAC and LLA match     0002:Orig trunk            0004:Orig a
0008:Orig trusted trunk    0010:Orig trusted access   0020:DHCP a
0040:Cga authenticated     0080:Cert authenticated    0100:Static

    Network Layer Address   Link Layer Address Interface   vlan
Time left
L   10.6.5.254           0000.0c9f.fa49         Vl1021     1021 0
REACHABLE
DH4 10.6.5.100        0050.568f.784f        Gi1/0/1     1021 0025
REACHABLE  281 s(1881 s)
ND  FE80::CC34:AC53       0050.568f.784f        Gi1/0/1     1021
REACHABLE  4 s try 0
```

IP TO SGT Bindings/Mappings

The network device must keep track of the authenticated endpoints, associated IP addresses, and assigned SGTs. An SGT assigned to an endpoint is bound to the associated IP address; hence, the IP TO SGT association for an endpoint is called a binding or a mapping.

As well as the SGT being updated or changed, the binding can also be removed. A binding will be removed if pertinent configuration is deleted, if there is a DHCP release request from an endpoint, if there is a DHCP lease expiration, or when ARP probes detect that a host is disconnected (or becomes nonresponsive for any reason). Whether just updated or deleted, a AAA accounting message is generated and sent to ISE. This message allows ISE to be aware of the session status and be able to immediately update internal and external systems (via SXP or pxGrid).

IP TO SGT bindings reside in a virtual routing and forwarding (VRF) table on network devices. That could be a user-defined VRF or the default VRF if there are no user-defined VRFs. If an endpoint is connected to a VRF and has an IP address assigned, then any SGT associated with that endpoint will create an IP TO SGT binding in the same VRF. There are well-defined scale numbers advertised for the number of bindings supported per platform; these numbers are platform-wide across all VRFs.

IP TO SGT Dynamic Classification

Dynamic classification means that the SGT is assigned by a server, typically by Cisco Identity Services Engine through a protocol called Remote Authentication Dial-In User Service (RADIUS). When users and endpoints authenticate with ISE, the SGT is assigned as part of the authorization phase.

If an SGT is assigned dynamically from ISE as part of authorization, then the access device receives the SGT in the RADIUS Access-Accept and tracks its use with the endpoint's session. Session information within access devices can be displayed with **show access-session** or **show authentication** commands, as shown in Example 7-2. Near the bottom, SGT value 18 has been assigned from ISE for the endpoint

Technet24

with the MAC address ending in 4f.

Example 7-2 ***show access-session*** *CLI Output*

```
Switch#show access-session
Interface   MAC Address    Method  Domain  Status Fg   Session ID
-----------------------------------------------------------------
Gi1/0/1     0050.568f.784f dot1x   DATA    Auth        0219040A00
Gi1/0/1     000c.295e.4932 dot1x   DATA    Auth        0101010400
Gi1/0/37    0050.568f.4d8c dot1x   DATA    Auth        0101010400
Gi1/0/2     0050.5694.d054 dot1x   DATA    Auth        0101010400
Gi1/0/37    0050.56a0.1a50 dot1x   DATA    Auth        0101010400
Gi1/0/37    0050.56a0.c867 mab     DATA    Unauth      0101010400

Session count = 1

Key to Session Events Blocked Status Flags:

  A - Applying Policy (multi-line status for details)
  D - Awaiting Deletion
  F - Final Removal in progress
  I - Awaiting IIF ID allocation
  P - Pushed Session
  R - Removing User Profile (multi-line status for details)
  U - Applying User Profile (multi-line status for details)
  X - Unknown Blocker

Switch#show access-session mac 0050.568f.784f details
            Interface:  GigabitEthernet1/0/1
               IIF-ID:  0x1F7BCB65
          MAC Address:  0050.568f.784f
         IPv6 Address:  fe80::cc34:7361:cd10:ac53
         IPv4 Address:  10.6.5.100
            User-Name:  hvac2
               Status:  Authorized
               Domain:  DATA
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
  Acct update timeout:  172800s (local), Remaining: 103271s
```

```
    Common Session ID:  0219040A0000001191D49891
      Acct Session ID:  0x00000009
               Handle:  0x33000007
       Current Policy:  PMAP_DefaultWiredDot1xOpenAuth_1X_MAB

Local Policies:

Server Policies:
             VN Value:  BuildingMgmt
          Vlan Group:  Vlan: 1021
           SGT Value:  18

Method status list:
       Method           State
        dot1x           Authc Success
```

Note:

The CLI output shown in Example 7-2 also shows a **VN Value** and a **Vlan Group** assigned via the server policies (from ISE). The VN is a virtual network within SD-Access networks and manifests itself as a VRF in network devices.

Once sessions have an IP address and an SGT assigned, then the IP TO SGT bindings are tracked as per Example 7-3. The dynamic session in Example 7-2 with IP address 10.6.5.100 and SGT 18 is tracked as a source type of LOCAL.

Example 7-3 *Displaying IP and SGT Assignment Within IOS-XE*

```
Switch#show cts role-based sgt-map all
Active IPv4-SGT Bindings Information

IP Address              SGT     Source
============================================
1.1.1.8                 2       INTERNAL
10.6.5.100              18      LOCAL
```

Technet24

```
50.50.50.1              20      CLI
51.51.0.0/16            21      CLI
52.52.52.10             22      VLAN
68.100.19.8             62      SXP
172.16.13.0/24          20      L3IF

IP-SGT Active Bindings Summary
============================================
Total number of CLI      bindings = 2
Total number of INTERNAL bindings = 1
Total number of active   bindings = 6
```

LOCAL entries as well as other source types are described here:

- LOCAL represents IP TO SGT bindings formed as the result of a dynamic authentication with a RADIUS server such as ISE. LOCAL is also used to show bindings for IP addresses learned (via IP device tracking or ARP snooping) from the presence of Port to SGT configuration commands.
- INTERNAL signifies an SGT assigned to IP interfaces on the network device itself.
- CLI entries represent static IP TO SGT or Subnet to SGT bindings configured on the network device.
- VLAN represents an SGT assigned to a learned IP address (via IP device tracking or ARP snooping) from the presence of a VLAN TO SGT configuration command.
- SXP signifies the IP TO SGT binding is learned from Security Group Tag Exchange Protocol (SXP).
- L3IF entries are SGTs assigned to IP prefixes learned from L3IF configuration commands on a Layer 3 interface.

Note:

SXP is discussed later in this chapter.

Note:

Several different classification methods could be present on a network device for the same IP address. For example, an IP TO SGT mapping could be learned via IP device tracking for VLAN TO SGT configuration; plus an IP TO SGT configuration command could also be configured. A built-in order of precedence is used to determine which mapping takes precedence, and that order of precedence is discussed later.

Once network devices have an IP TO SGT binding dynamically assigned, then Accounting messages are sent to ISE to keep ISE informed of the assignment. This approach allows ISE to keep track of the global bindings assigned and share with the ecosystem using pxGrid and SXP. When network devices have any IP TO SGT bindings removed, then ISE is updated accordingly to allow the whole ecosystem to be synchronized.

The network device needs to continually monitor the status of each endpoint, including any change or removal of the IP TO SGT bindings. As well as through a configuration change, an assigned SGT can change as part of a reauthorization; this may occur through the ISE Profiling function, for example. After initially assigning an SGT, ISE learns of further endpoint attributes over time, and using RADIUS Change of Authorization (CoA), the classification can be updated on the network device appropriately.

Another example of an assigned SGT changing is using what is known as rapid threat containment (RTC). RTC is a Cisco term explaining the detection of threats by ISE ecosystems and using automation for threat mitigation. When an ISE ecosystem platform detects some sort of anomaly in the endpoint posture or behavior, then that platform

Technet24

can send a message to ISE using the ISE Adaptive Network Control (ANC) function to quarantine the endpoint. This mechanism works by changing the SGT assigned to the endpoint, often by assigning an SGT called quarantine. Figure 7-1 shows RTC being used by the Cisco Secure Network Analytics platform.

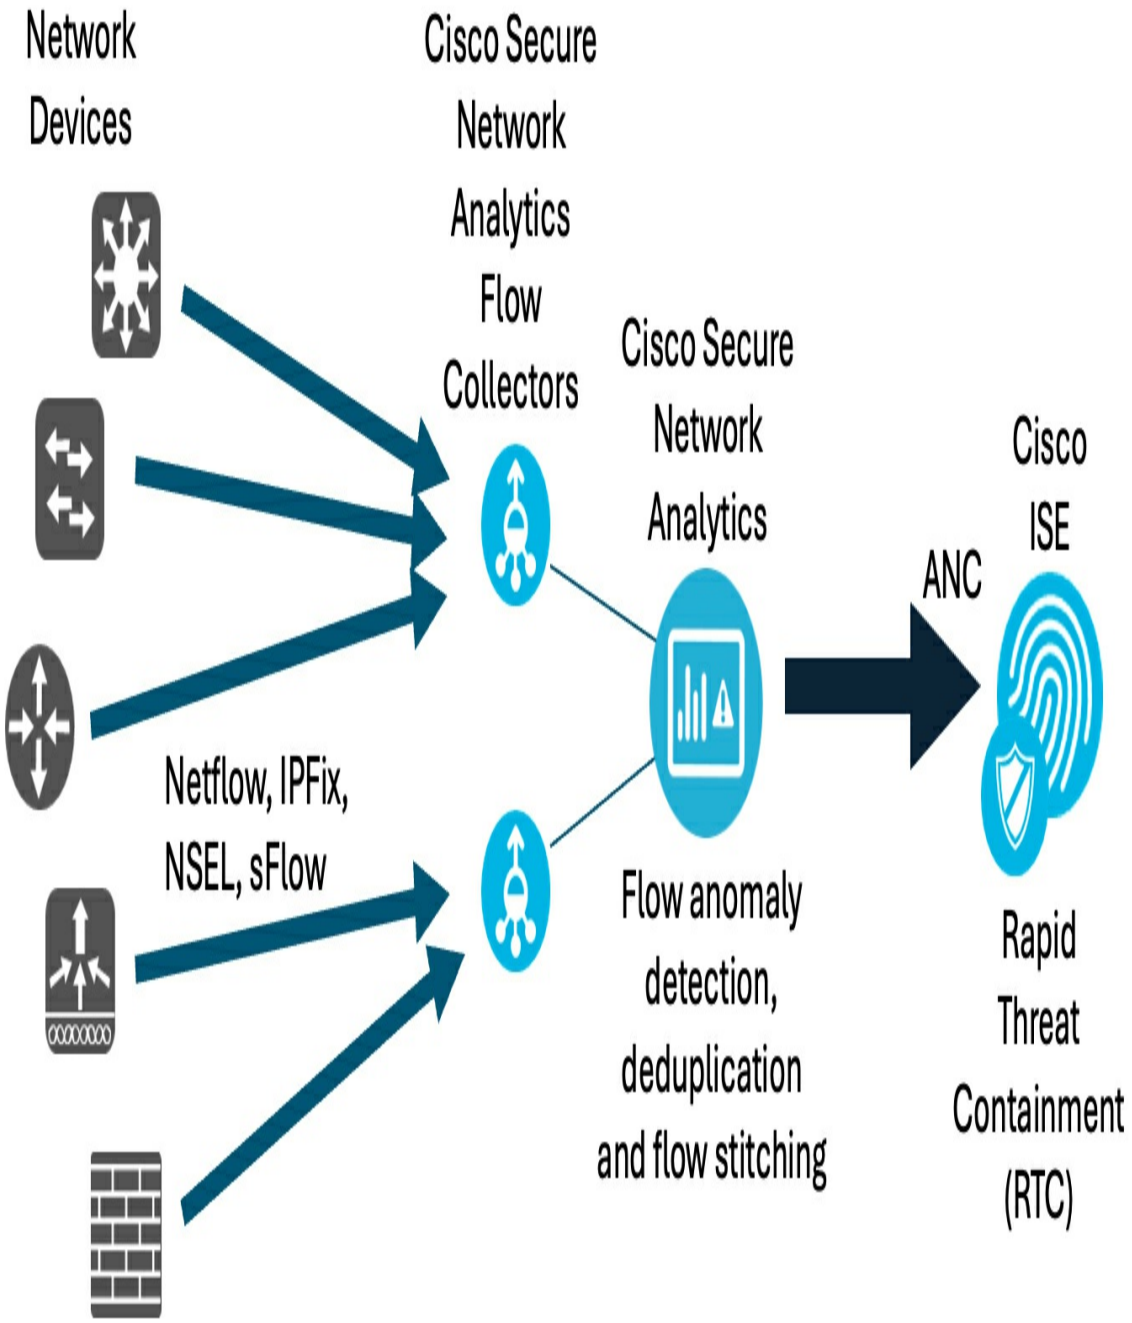


Figure 7-1 *Cisco Secure Network Analytics Utilizing ISE ANC*

IP TO SGT Static Classification

Static classification, in contrast, is used when ISE is not used for authentication or authorization, providing support for servers, subnets, VLANs, and Layer 2 and Layer 3 interfaces. The following configuration commands are available for the static methods on network devices:

For IP TO SGT binding:

```
cts role-based sgt-map  sgt
cts role-based sgt-map vrf   sgt
```

For Subnet to SGT binding:

```
cts role-based sgt-map / sgt
cts role-based sgt-map vrf  / sgt
```

For VLAN TO SGT binding:

```
cts role-based sgt-map vlan-list <1-4094 or all> sgt
```

For Port to SGT binding:

```
interface
cts manual
  policy static sgt
  no propagate sgt
```

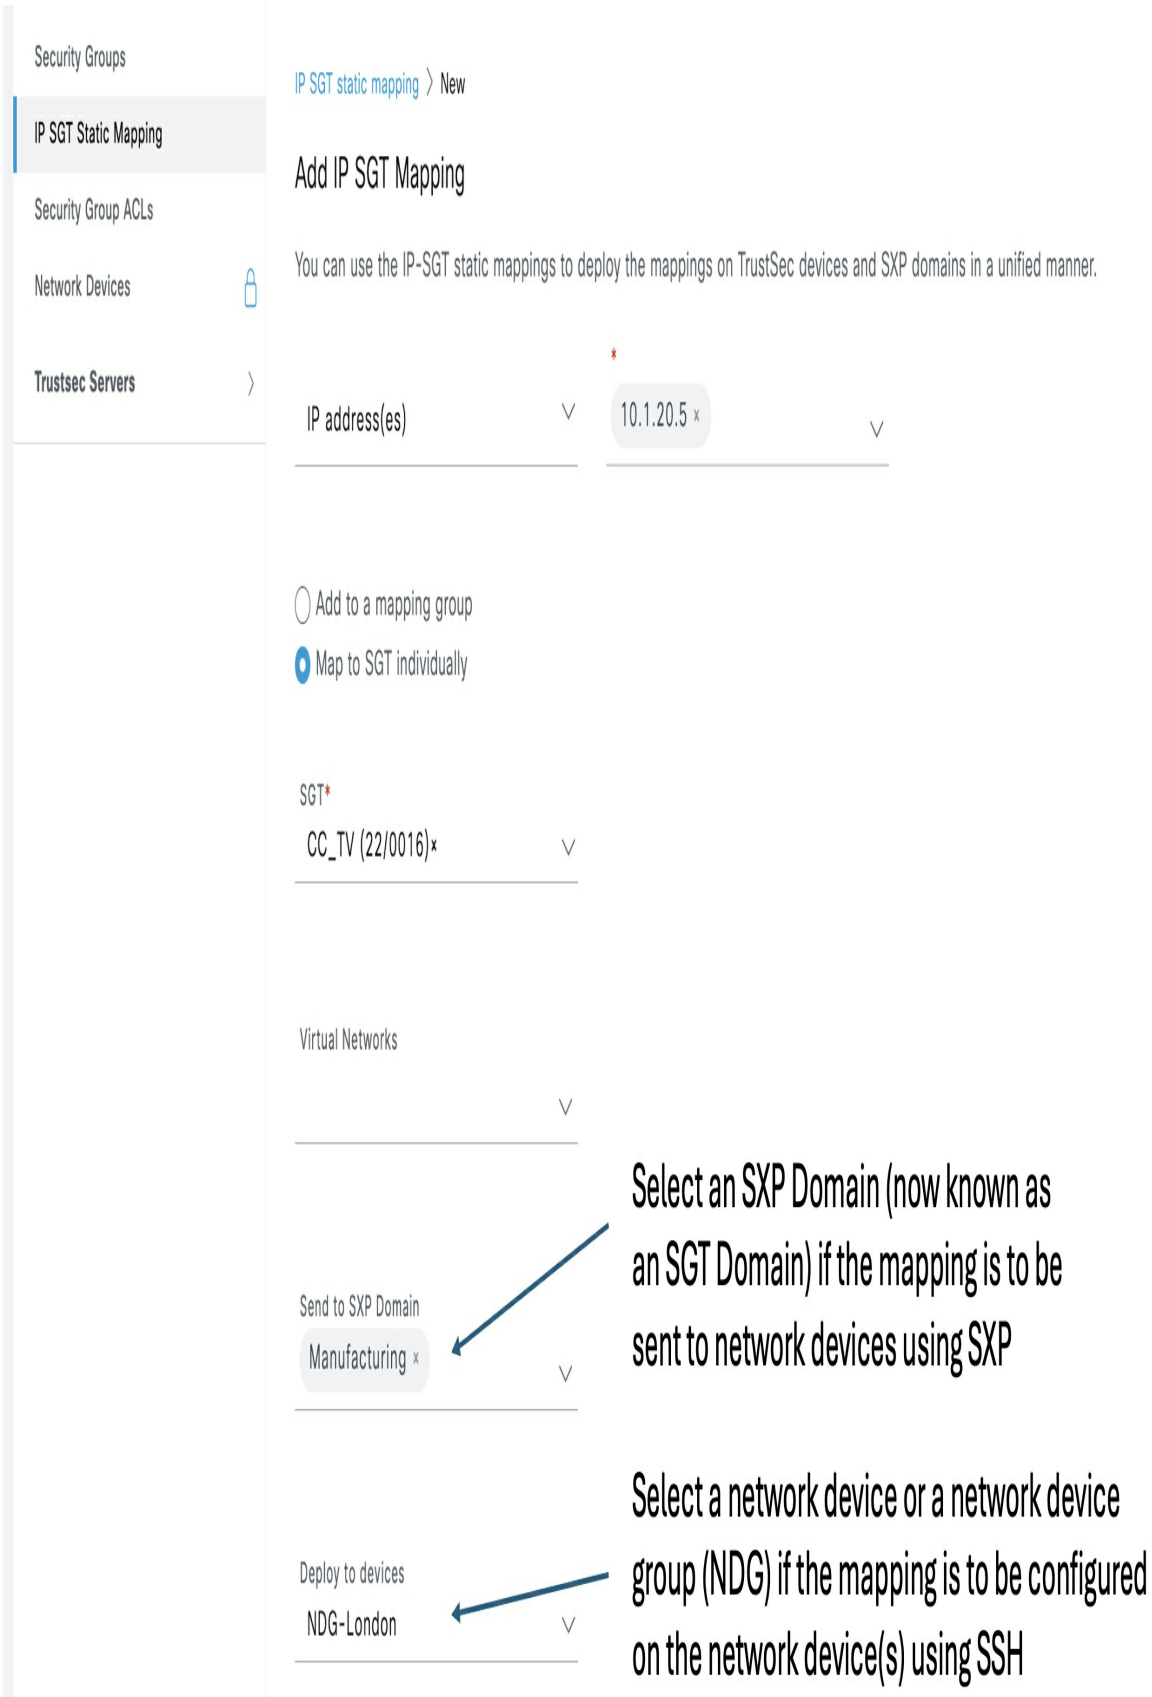
For L3IF:SGT binding:

```
interface
  no switchport
  cts role-based sgt-map sgt
  or cts role-based sgt-map security-group
```

Static IP and Subnet to SGT bindings, or mappings, can be managed by ISE. You can manage them by navigating to **Work Centers >**

Technet24

TrustSec > Components > IP SGT Static Mapping (see Figure 7-2). ISE can send the added mappings to network devices using SXP, explained later in the "Propagation" section, and/or by using SSH to log in to the devices to configure via the CLI. It is typical to add mappings to the SXP table or to configure directly using SSH, not both, but the option is available to send by both should it be required.


Security Groups
IP SGT Static Mapping
Security Group ACLs
Network Devices
Trustsec Servers
IP SGT static mapping 〉 New
Add IP SGT Mapping
You can use the IP-SGT static mappings to deploy the mappings on TrustSec devices and SXP domains in a unified manner.
IP address(es)
10.1.20.5
Add to a mapping group
Map to SGT individually
SGT*
CC_TV (22/0016)
Virtual Networks
Send to SXP Domain
Manufacturing
Deploy to devices
NDG-London
Select an SXP Domain (now known as an SGT Domain) if the mapping is to be sent to network devices using SXP
Select a network device or a network device group (NDG) if the mapping is to be configured on the network device(s) using SSH

Technet24

Figure 7-2 *Adding IP TO SGT Static Mapping*

If the mappings are set to be configured using SSH via **Deploy to Devices**, then the configuration is not carried out by saving the mapping. Several mappings can be added and saved this way with the information staying within ISE. When you are ready to configure onto the network devices, select the **Deploy** option to deploy the mappings at the top of the **IP SGT Static Mapping** UI (see Figure 7-3). This **Deploy** function is required only for the SSH function; it is not required when only adding mappings into the SXP database.



Figure 7-3 *Deploying Static Mappings Utilizing SSH*

When **Deploy** is selected, ISE displays the progress of the communication with the network device(s) and informs whether the deployment was successful or not. (In Figure 7-4, Kernow-9350 is a network device within the network device group called NDG-London.)



Figure 7-4 *Reporting Progress of Deploying Static Mappings via SSH*

Technet24

The newly created mapping is displayed within the **IP SGT Static Mapping** table shown in Figure 7-5.

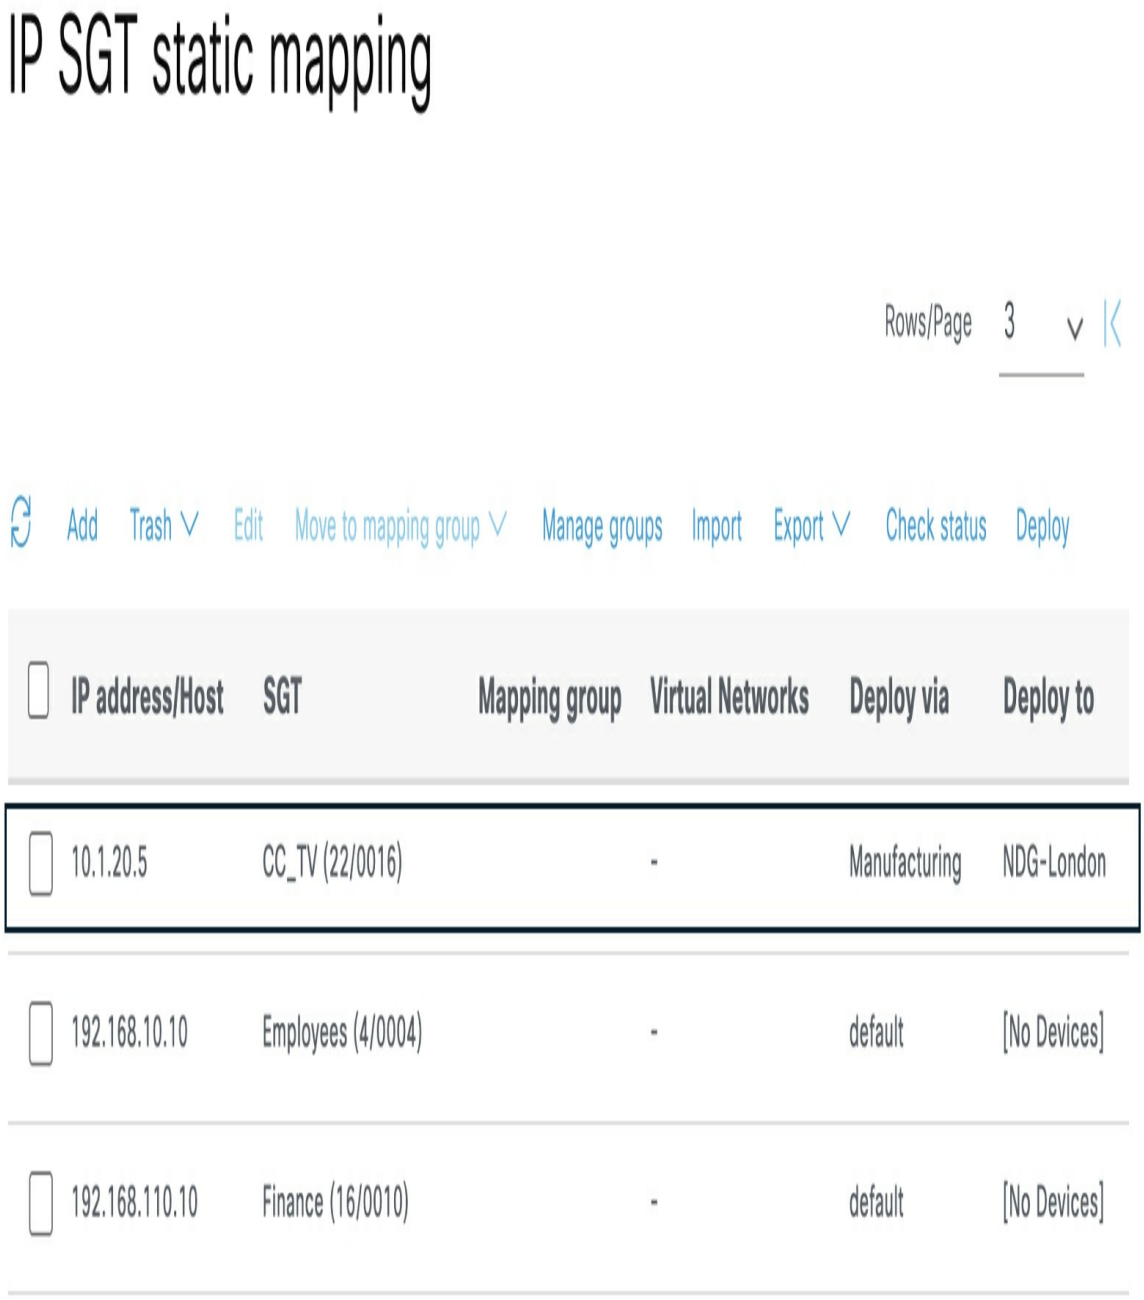


Figure 7-5 *Mappings Shown in ISE Static Mapping Table*

Because the 10.1.20.5-to-SGT CC_TV mapping was also sent to the SXP database, the binding can also be seen in the **SGT Bindings** UI (see

Figure 7-6). To locate it, navigate to **Work Centers > TrustSec > SXP > SGT Bindings**.

Technet24

SGT Bindings

This page displays all the configured SGT bindings, giving you an easy view of assigned primary SGTs, the source of the learned data, and any workload classifications applied. You can review classification details and the specifics of defined inbound & outbound SGT domain rules by clicking the corresponding column values.

Rows/Page

Add SGT Domain filter Manage SGT Domain filters

IP Address	Primary SGT	VN	Learned From	Learned ...	SGT Domain
10.1.1.126/32	Demo_WebEPG_ACI ...	DEFAULT_VN	10.1.101.37,7.5.9.136	Session	default
10.1.20.5/32	CC_TV (22/0016)		10.1.101.37	Local	Manufacturing
10.2.2.127/32	Demo_ClientEPG_AC...	DEFAULT_VN	10.1.101.37,7.5.9.136	Session	default
10.6.5.185/32	RandD (115/0073)	DEFAULT_VN	10.1.101.37,1.1.1.8	Session	default
192.168.10.10/32	Employees (4/0004)		10.1.101.37	Local	default
192.168.110.10/32	Finance (16/0010)		10.1.101.37	Local	default

Figure 7-6 *Mappings Shown in ISE SGT Binding Table*

Bindings are created and reside in a VRF table. If no VRF is specified when adding a binding via CLI, then the binding is placed within the global/default VRF. As can be seen from the ISE UI when adding an **IP SGT Static Mapping**, there is no field to enter VRF information. The **Virtual Networks** field in the UI is for use internally by ISE. Because there is no location to specify the VRF for provisioning mappings via SSH, then ISE can only use that SSH function to configure mappings in the global/default VRF on network devices.

Currently, an SGT is only ever assigned to an IP address.

Unclassified Traffic/Unknown SGT

A best practice when deploying group-based policies is to start small with just a few SGTs on a limited number of network devices. So, if just a small number of SGTs are used to classify a subset of users and endpoints, then what happens to the remaining users and endpoints on the network? Because they are not classified, they are known as unclassified, and from a policy point of view, an SGT of Unknown/SGT 0 is implied.

A policy using Unknown can be configured in the ISE matrix, as shown in Figure 7-7.

Technet24

Production Matrix

Production

Populated cells:19

Create Policy | Monitor All Disabled | Deploy | Verify Deploy

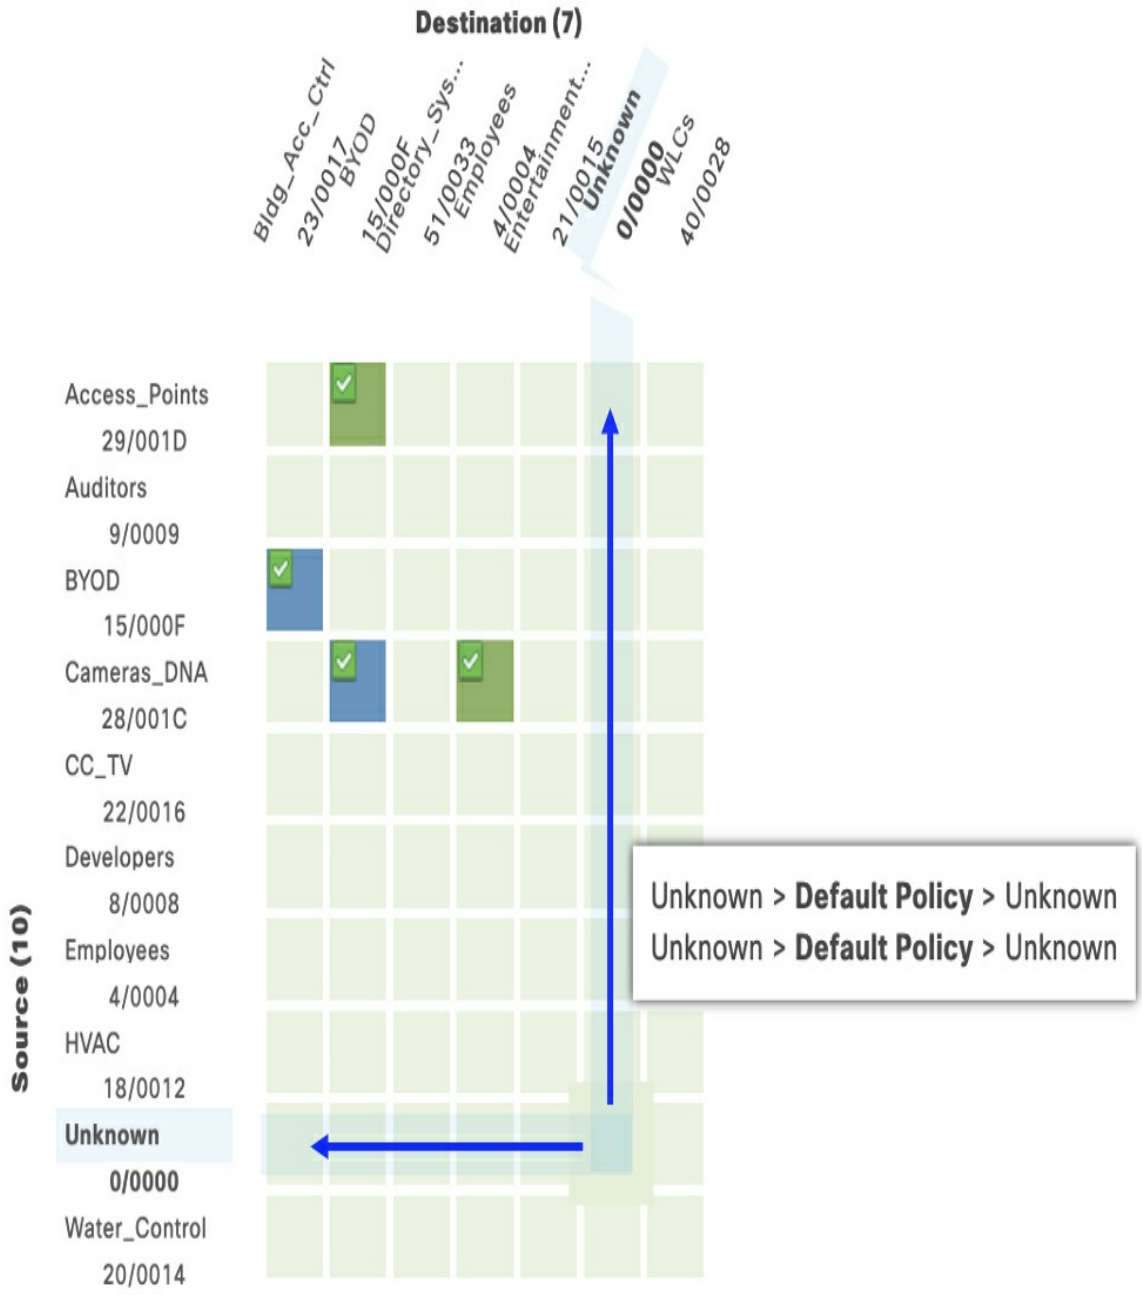


Figure 7-7 *Unknown/SGT 0 Entry Available in TrustSec Policy Matrix*

In this way, policies can be deployed from any SGT to Unknown, or from Unknown to any SGT, or indeed from Unknown to Unknown.

Classification Order of Precedence

Network devices rely on multiple classification methods, so establishing a clear order of precedence is vital. Consider what happens when a device encounters conflicting data: It might have a static IP TO SGT binding, a learned binding from SXP, and then receive a packet with a different SGT in the Cisco Meta Data (CMD) field. When these sources disagree, which mapping does the device actually trust?

The order of precedence used by Cisco IOS-XE devices, from highest priority to lowest, is as follows:

1. Inline tagging (SGT in CMD field or protocol header), when trusted
2. (INTERNAL) Local network device interfaces and network device SGT
3. SGT caching
4. (LOCAL) Local bindings learned via authentication, device tracking, and ARP snooping or Port to SGT static binding
5. IP ARP: Tagged ARP packets received inline
6. Bindings received by SXP
7. Layer 3 interface static binding (L3IF)
8. (CLI) IP TO SGT and Subnet to SGT static binding (longest match

Technet24

takes precedence)

9. VLAN TO SGT static binding

Classification Summary

Classification is where network traffic is assigned into a security group, typically at network ingress. It is always the IP address of a user or endpoint that is assigned into a group.

The assignment can be made statically via CLI on network devices, or from ISE, where ISE has the capability to use SXP or SSH to accomplish the static designation.

The assignment can also be made dynamically using the RADIUS protocol and a AAA server such as ISE. When users or endpoints connect to the network, the RADIUS protocol provides authentication and authorization whereby an SGT is assigned via an ISE authorization rule. The accounting component keeps ISE updated with the active bindings.

Propagation

Propagation is simply the process of sending IP TO SGT bindings from one location to another. Bindings can be propagated in the data plane or control plane, depending on the use case and hardware available.

Propagation Introduction

We have discussed classifying IP addresses into security groups using static and dynamic methods and by utilizing mechanisms such as DHCP and ARP snooping. Dynamic assignment happens at the point where the endpoint enters the network — that is, at the access switch or VPN gateway (with ISE assigning the SGT). Static assignment

happens wherever that static binding is configured; examples are access, distribution, core, on switches, routers, or firewalls. Classification information, or bindings, can also be propagated from one location to another.

So, why propagate the SGT? The technology dictates that any enforcement point needs to be able to understand both the source and destination classification (IP address and SGT) when making a policy decision. If an enforcement point is at the ingress of the network, then it would need to know all the destinations and classifications that are deeper in the network. This would mean that if the network had 500,000 active endpoints, the enforcement point would, in theory, need a table in memory that included all 500,000 endpoint classifications plus all the enforcement rules needed to protect those endpoints. There are enforcement points that can accommodate a large number of entries, but they are typically very expensive and very large-scale routers/switches/firewalls. This makes them unsuited for enforcement in access layers. If you were somehow able to forward the source classification, then you could enforce on the egress of the network. Then the only aspect the enforcement point needs to understand is the directly connected users/endpoints that are being defended by the enforcement point, because the source classification information would be received with every packet. Using this technique, hundreds or thousands of classifications can be catered for providing an economical solution. So, the answer to the question "Why propagate the SGT?" is for scaling reasons.

If the enforcement point is on a different platform than the device carrying out the classification, then that source classification information must be passed from classification device to enforcement device. If two users are connected to the same access switch and the use case is to deny traffic between those users, then there would be

Technet24

no need for propagation in that instance. Both source and destination SGTs would be on the same switch, and enforcement would occur on that switch.

Inline Tagging Propagation

Inserting SGTs directly in the data plane is a highly efficient way to propagate classification information across your network. While Cisco ISE manages the policy environment, it does not participate directly in the data-plane functional layer, but understanding these mechanisms is vital for a complete view of your solution.

The primary method for this is the use of Cisco Meta Data, or CMD. When you use CMD, the SGT is inserted directly into the frame header of Layer 2 traffic. Because most endpoints are not aware of this tagging, network devices handle the heavy lifting, inserting the tag and passing it along to the next hop. As frames travel, receiving devices extract the SGT and the source IP, effectively learning the binding and allowing for policy enforcement or further propagation.

Beyond CMD, you have several other ways to carry this classification information. You can propagate tags within VXLAN headers or encapsulate them using secure technologies like GRE, IPsec, DMVPN, and GETVPN. Regardless of the specific transport method, the goal remains the same: ensuring the network egress device receives the SGT, allowing it to apply the correct security policy based on the known destination binding. By embedding the SGT directly into the data path, you significantly improve the scalability of your group-based access control across the entire organization.

Propagation via pxGrid

Inline tagging is not supported by all network devices or vendors. If

inline tagging is not suitable, then other propagation methods are available. ISE pxGrid is one method that is widely used by Cisco platforms as well as other vendors to receive binding information from ISE. Examples are Cisco Secure Network Analytics, Cisco Secure Firewall, Cisco Secure Web Appliance, Check Point, Palo Alto, and Fortinet. In this way, these ISE ecosystems can gain visibility of the assigned IP TO SGT bindings and act accordingly — for example, to enforce traffic using group-based policies.

Cisco ISE needs to be configured to propagate bindings over pxGrid. You can find the setting under **Work Centers > TrustSec > Settings > SXP Settings** (see Figure 7-8).



Figure 7-8 *Setting ISE to Propagate SXP Bindings via pxGrid*

Propagation via Security Group Tag Exchange Protocol (SXP)

Another way to propagate IP TO SGT bindings is via the Security Group Tag Exchange Protocol, or SXP. Unlike inline tagging, which relies on the data plane to carry the SGT, SXP handles this process through the control plane and is supported by both network devices and Cisco ISE. You establish a peer-to-peer TCP overlay connection between devices over port 64999. When setting up this connection, you define the roles for each end. One device can act as a speaker to

Technet24

send mappings to a listener, or you can configure both ends to act as both a speaker and a listener, which enables bidirectional propagation. Because this protocol uses TCP as its transport, you do not need to worry about the vendor of any intermediate devices because they simply pass the traffic along without needing to understand the SXP protocol itself.

SXP Versions

Over time a number of different SXP versions have been developed. Table 7-1 lists the versions and the supported functions.

Table 7-1 *SXP Versions and Functions*

SXP Version	Function
SXP Version 1	Initial SXP version supporting IPv4 binding propagation.
SXP Version 2	Includes support for IPv6 binding propagation and version negotiation.
SXP Version 3	Adds support for Subnet to SGT binding propagation. If speaking to a lower version, then the subnet will be expanded to individual IP-SGT entries.
SXP Version 4	Provides loop detection and prevention, capability exchange, and a built-in keepalive mechanism (most prevalent today and is critical if it's possible for bindings to be sent in a loop in the network.).
SXP Version 5	Carries VRF and VLAN information, reducing the number of SXP connections required when multiple VRFs are deployed

Note:

IP TO SGT bindings sit in a VRF (global/default or specified) on network devices and in an SGT domain in ISE. SXP connections terminated in a network device VRF or ISE SGT domain only send and receive bindings within that VRF/SGT domain. Only SXP Version 5 can send or receive bindings across VRFs (which ISE does not support at the time of writing).

SXPv4 Loop Detection and Mitigation, SXP dB, and RBM

Technet24

Figure 7-9 illustrates SXPv4 loop detection and high availability.

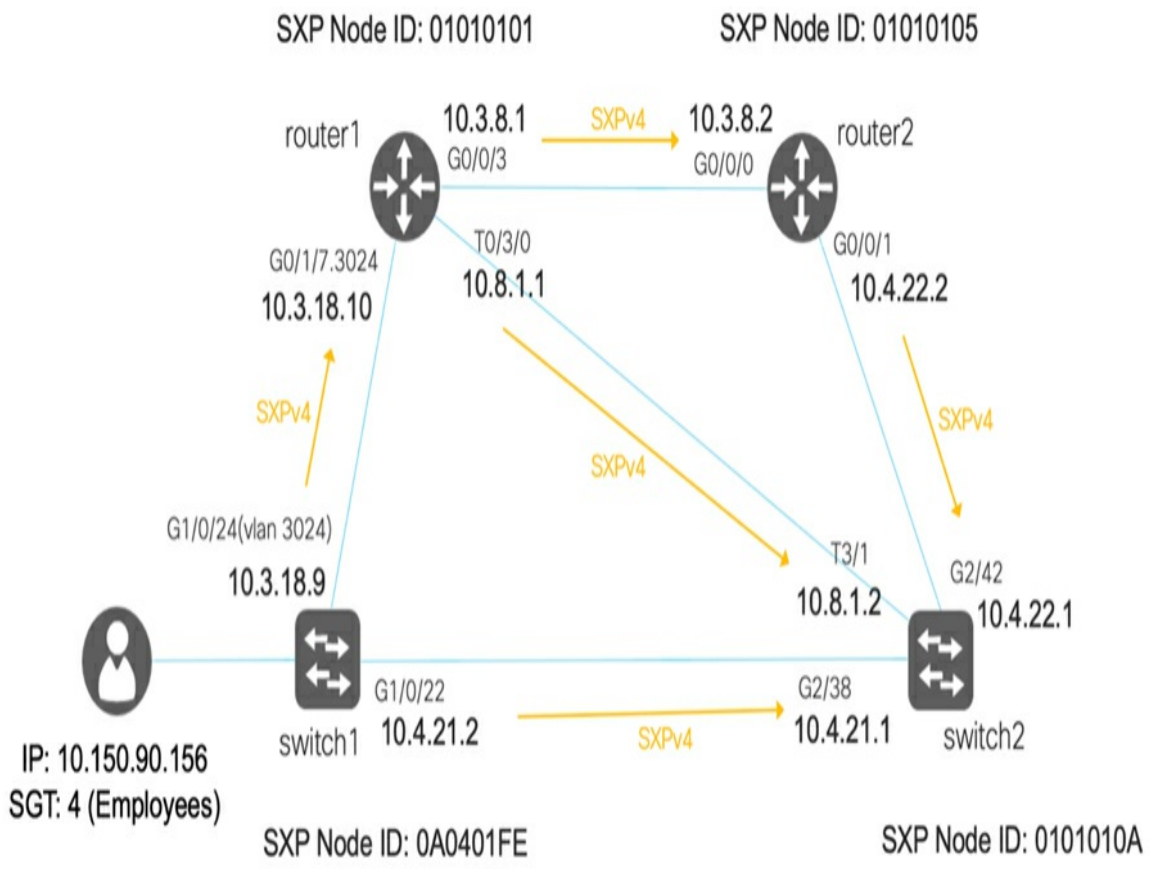


Figure 7-9 *Multiple SXPv4 Paths*

In this figure, an endpoint is connected to switch1 and is assigned SGT 4 (Employees). Five SXP peer-to-peer connections are configured as depicted by the dotted arrows. The associated IP TO SGT binding is sent from switch1 to both router1 and switch2 via SXPv4. Router1 then also sends the binding to both router2 and switch2. Finally, router2 sends the binding to switch2. Ultimately, switch2 receives the IP TO SGT binding three times — one from each possible path from switch1 to switch2.

Note that Cisco ISE supports SXPv4 and therefore would behave similarly to the propagating devices depicted in Figure 7-9.

Let's investigate the state of the SXP connections using switch2, as shown in Example 7-4.

Example 7-4 *SXP Information Shown by switch2*

```
switch2#show cts sxp connections brief
 SXP              : Enabled
 Highest Version Supported: 4
 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
------------------------------------------------------------------
10.4.21.2        10.4.21.1     On             0:00:13:14 (dd:hr:m
10.4.22.2        10.4.22.1     On             0:00:11:11 (dd:hr:m
10.8.1.1         10.8.1.2      On             0:00:09:16 (dd:hr:m

Total num of SXP Connections = 3

switch2#show cts sxp connections
 SXP              : Enabled
 Highest Version Supported: 4
 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          : 10.4.21.2
Source IP        : 10.4.21.1
Conn status      : On
Conn version     : 4
Conn capability  : IPv4-IPv6-Subnet
Conn hold time   : 120 seconds
Local mode       : SXP Listener
Connection inst# : 1
```

Technet24

```
TCP conn fd      : 1
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 0:00:14:19 (dd:hr:mm:sec)

---------------------------------------------
Peer IP          : 10.4.22.2
Source IP        : 10.4.22.1
Conn status      : On
Conn version     : 4
Conn capability  : IPv4-IPv6-Subnet
Conn hold time   : 120 seconds
Local mode       : SXP Listener
Connection inst# : 1
TCP conn fd      : 2
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 0:00:12:16 (dd:hr:mm:sec)

---------------------------------------------
Peer IP          : 10.8.1.1
Source IP        : 10.8.1.2
Conn status      : On
Conn version     : 4
Conn capability  : IPv4-IPv6-Subnet
Conn hold time   : 120 seconds
Local mode       : SXP Listener
Connection inst# : 1
TCP conn fd      : 3
TCP conn password: default SXP password
Hold timer is running
Duration since last state change: 0:00:10:21 (dd:hr:mm:sec)


Total num of SXP Connections = 3
```

Note:

The commands shown in Example 7-4 do not specify a VRF

and therefore relate to the global VRF.

There is no SXP loop in Figure 7-9; all the SXP connections are unidirectional and do not form a loop. If, however, the SXP connection from switch1 to switch2 were reversed (that is, switch2 set as a speaker and switch1 as a listener), then there would be a loop around switch1, router1, and switch2. If undetected and unsuppressed, this approach would cause the SXP packets to loop continually, causing platform resource issues.

To prevent continual looping, SXPv4 has introduced an SXP Node ID. It can be set manually or assigned automatically when SXP is enabled per network device. When automatically assigning a Node ID, a network device (including ISE) will use the highest configured loopback address, or if no loopback is configured, the highest configured address on any other interface. The Node IDs in Figure 7-9 have been assigned automatically.

Example 7-5 shows an SXP CLI command on switch2 that displays the Node ID (automatically generated) for the switch itself and the received bindings (otherwise known as mappings).

Example 7-5 *Switch2 Node-ID and Received Bindings/Mappings*

```
switch2#show cts sxp sgt-map
SXP Node ID(generated):0x0101010A(1.1.1.10)
IP-SGT Mappings as follows:
IPv4,SGT: <10.150.90.156 , 4:Employees>
source  : SXP;
Peer IP : 10.4.21.2;
Ins Num : 1;
Status  : Active;
Seq Num : 13
Peer Seq: 0A0401FE,
IPv4,SGT: <10.150.90.156 , 4:Employees>
source  : SXP;
```

Technet24

```
Peer IP : 10.4.22.2;
Ins Num : 1;
Seq Num : 97
Peer Seq: 01010105,01010101,0A0401FE,
IPv4,SGT: <10.150.90.156 , 4:Employees>
source  : SXP;
Peer IP : 10.8.1.1;
Ins Num : 1;
Seq Num : 55
Peer Seq: 01010101,0A0401FE,
Total number of IP-SGT Mappings: 3
```

In Example 7-5, three bindings are received for IP 10.150.90.156 and SGT 4 (Employees). The most relevant piece of context within the binding information is the Peer Seq (peer sequence), which shows the Node IDs of the network devices traversed from the source. The first binding in the output has a Peer Sequence of 0A0401FE, which in Figure 7-9 is switch1. This binding correlates with the Peer IP shown as well. The last binding in the output has a Peer Sequence of 01010101,0A0401FE, so has traversed router1 and switch1. Similarly, the middle binding has a Peer Sequence of 01010105,01010101,0A0401FE so has been propagated via router2, router1, and switch1.

These Node IDs are used for loop detection and prevention. If a node receives an SXP packet with its own Node ID within the Peer Sequence, then the packet will be discarded. As stated previously, ISE itself can be an SXP listener or speaker and supports SXPv4; therefore, it generates an SXP Node ID for itself. If ISE were using GigabitEthernet 0 with assigned IP 10.1.101.37, then the self-generated SXPv4 Node ID would be 0A016525.

The SXP CLI used in Example 7-5 shows SXP bindings from the SXP database. Commands commencing with **show cts sxp** retrieve information from this SXP database. Example 7-5 shows three entries

for a particular IP TO SGT binding. The entry with the least number of Node IDs traversed (and, if a tie, the last entry learned) will be actively used as the source of the binding. This entry is marked with **Status: Active** in Example 7-5 (it has only traversed one Node ID, which is less than the other two bindings). This entry is used by the network device Role-Based Manager (RBM) function to track active bindings and use for operations. CLI commands to investigate RBM output start with **show cts role-based**, as shown in Example 7-6.

Example 7-6 *Switch2 Active Bindings from the Role-Based Manager*

```
switch2#show cts role-based sgt-map all
Active IPv4-SGT Bindings Information

IP Address              SGT     Source
============================================
10.150.90.156           4       SXP
```

In this example, there is just one entry in this table for this particular binding, and it is actively used by this network device.

If any of the traversed paths disappear for any reason (for example, the link between switch1 and switch2 goes down), then switch2 can immediately revert to using the next preferred entry in the SXP database (the path traversing router1).

Because SXP uses a binding table, there are limits to the number of bindings each platform can support. The scale numbers supported per platform are documented on the Cisco website, and those scale numbers refer to the RBM scale, not the SXP database.

SXPv5 Propagating VRF and VLAN Information

As was discussed previously, SXP bindings reside in a VRF (default/global or user-defined). SXP connections propagating these

Technet24

bindings from network device to network device are therefore terminated in those VRFs.

Figure 7-10 shows SXPv4 propagating bindings from VRF2 in one network device to the global VRF in another network device.

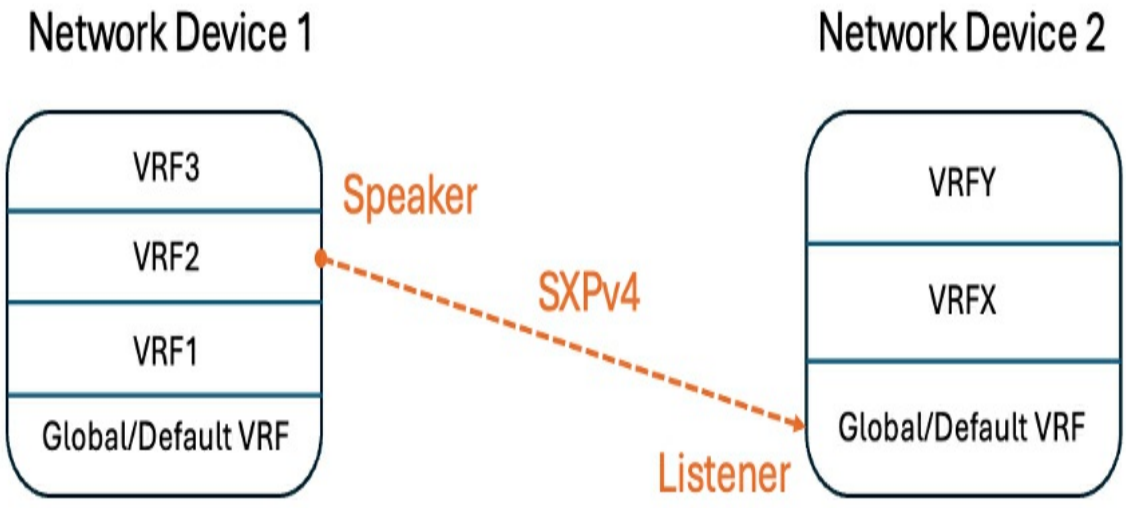


Figure 7-10 *SXPv4 Propagating Bindings Between Devices*

There is no intelligence built into SXPv4 for propagating specific bindings across VRFs; it just propagates or receives bindings within the VRF where the terminated SXP IP addresses reside. This bloats the configuration required when multiple VRFs require bindings from SXP. As an example, consider an SD-Access Border device, where five overlay VRFs are deployed requiring SXP. Figure 7-11 shows five SXPv4 connections are required.

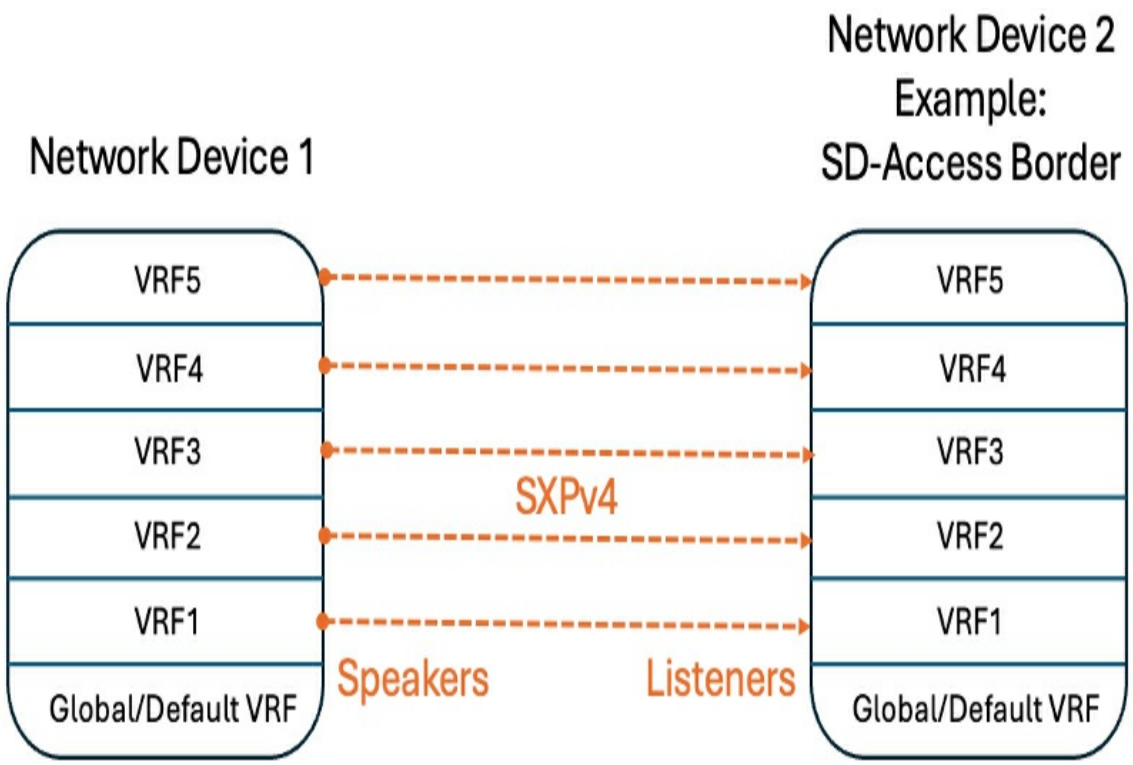


Figure 7-11 *SXPv4 Connection per VRF*

Note that network device 1 here could be Cisco ISE. Rather than VRFs, Cisco ISE uses SGT domains where IP TO SGT bindings reside.

Cisco IOS-XE release 17.9.1 delivered SXPv5. At the time of writing, Cisco ISE does not yet support SXPv5. We have included SXPv5 here because it is hoped ISE will support it in the future. With support in Cisco IOS-XE, it is supported on Cisco switches, Cisco routers, and Cisco 9800 wireless LAN controllers. SXPv5 not only carries the IP and SGT bindings, but also the VRF where the binding resides and the VLAN is assigned.

Through configuration, bindings from specific VRFs can be transmitted from network device 1 and received into the same VRFs on network device 2, all using just one SXP connection. In the configuration displayed in Figure 7-12, bindings from all VRFs on network device 1 are exported over the one SXPv5 connection, no matter which VRF

Technet24

that SXP connection is terminated on that network device. The associated VRF is transmitted, along with the IP and SGT of each binding. Network device 2, on the receiving side, shows the configuration to use the VRF received with each binding, to place each binding into the correct VRF on network device 2.

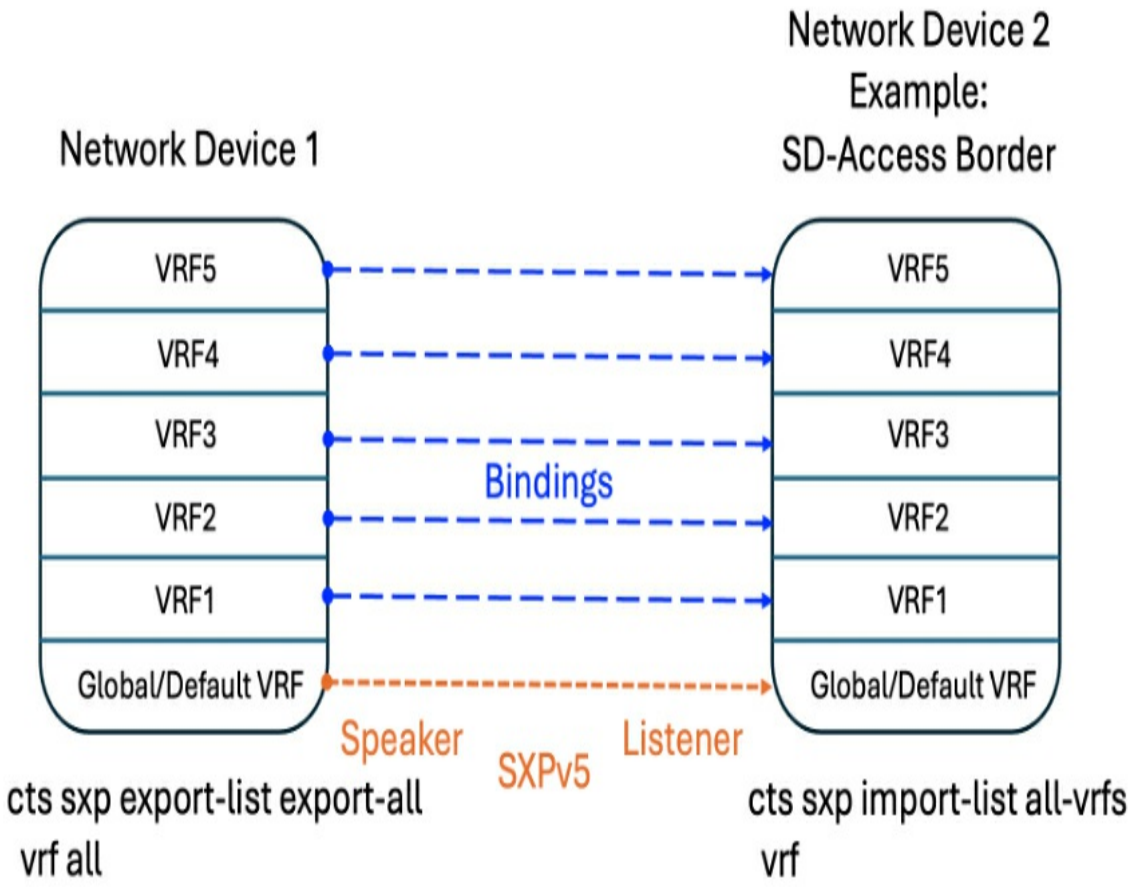


Figure 7-12 *One SXPv5 Connection Required to Propagate Bindings in Multiple VRFs*

SXPv5 allows the propagation of bindings in multiple VRFs across a single SXP connection. Different export and import list configurations are available to suit varying use cases. The advantage is that SXP connection configuration is reduced when deploying an architecture with multiple VRFs.

SXP Configuration and Timers

SXP is supported by Cisco ISE, routers, switches, and Cisco Adaptive Security Appliances (ASA). The configuration in this section applies to all these products.

SXP needs to be enabled as a service on Cisco ISE; to do so, navigate to **Administration > Deployment**, within the **Policy Service** section. On network devices, you enable it using the following command:

```
Switch(config)#cts sxp enable
```

When you're building SXP connections from one network device to another, a password is optional, along with setting the source IP address to be used for the connection. Both of these can be specified within the connection configuration itself, or a global configuration can be used. If you use global settings, then the command structure is as follows:

```
Switch(config)#cts sxp default password
Switch(config)#cts sxp default source-ip
```

Within ISE, you can set the global password under **Work Centers > TrustSec > Settings > SXP Settings**.

When a peer-to-peer SXP connection is configured, one end of the connection is set as a speaker and the other end as a listener. IP TO SGT bindings will be propagated from the network device with an SXP connection configured as a speaker to the network device with an SXP connection configured as a listener. If required, a connection can be configured to send or receive bindings in both directions.

A sample command configuring an SXP speaker on a network device is as follows:

```
Switch(config)#cts sxp connection peer  source
```

Technet24

```
password  mode local speaker hold-time  vrf
```

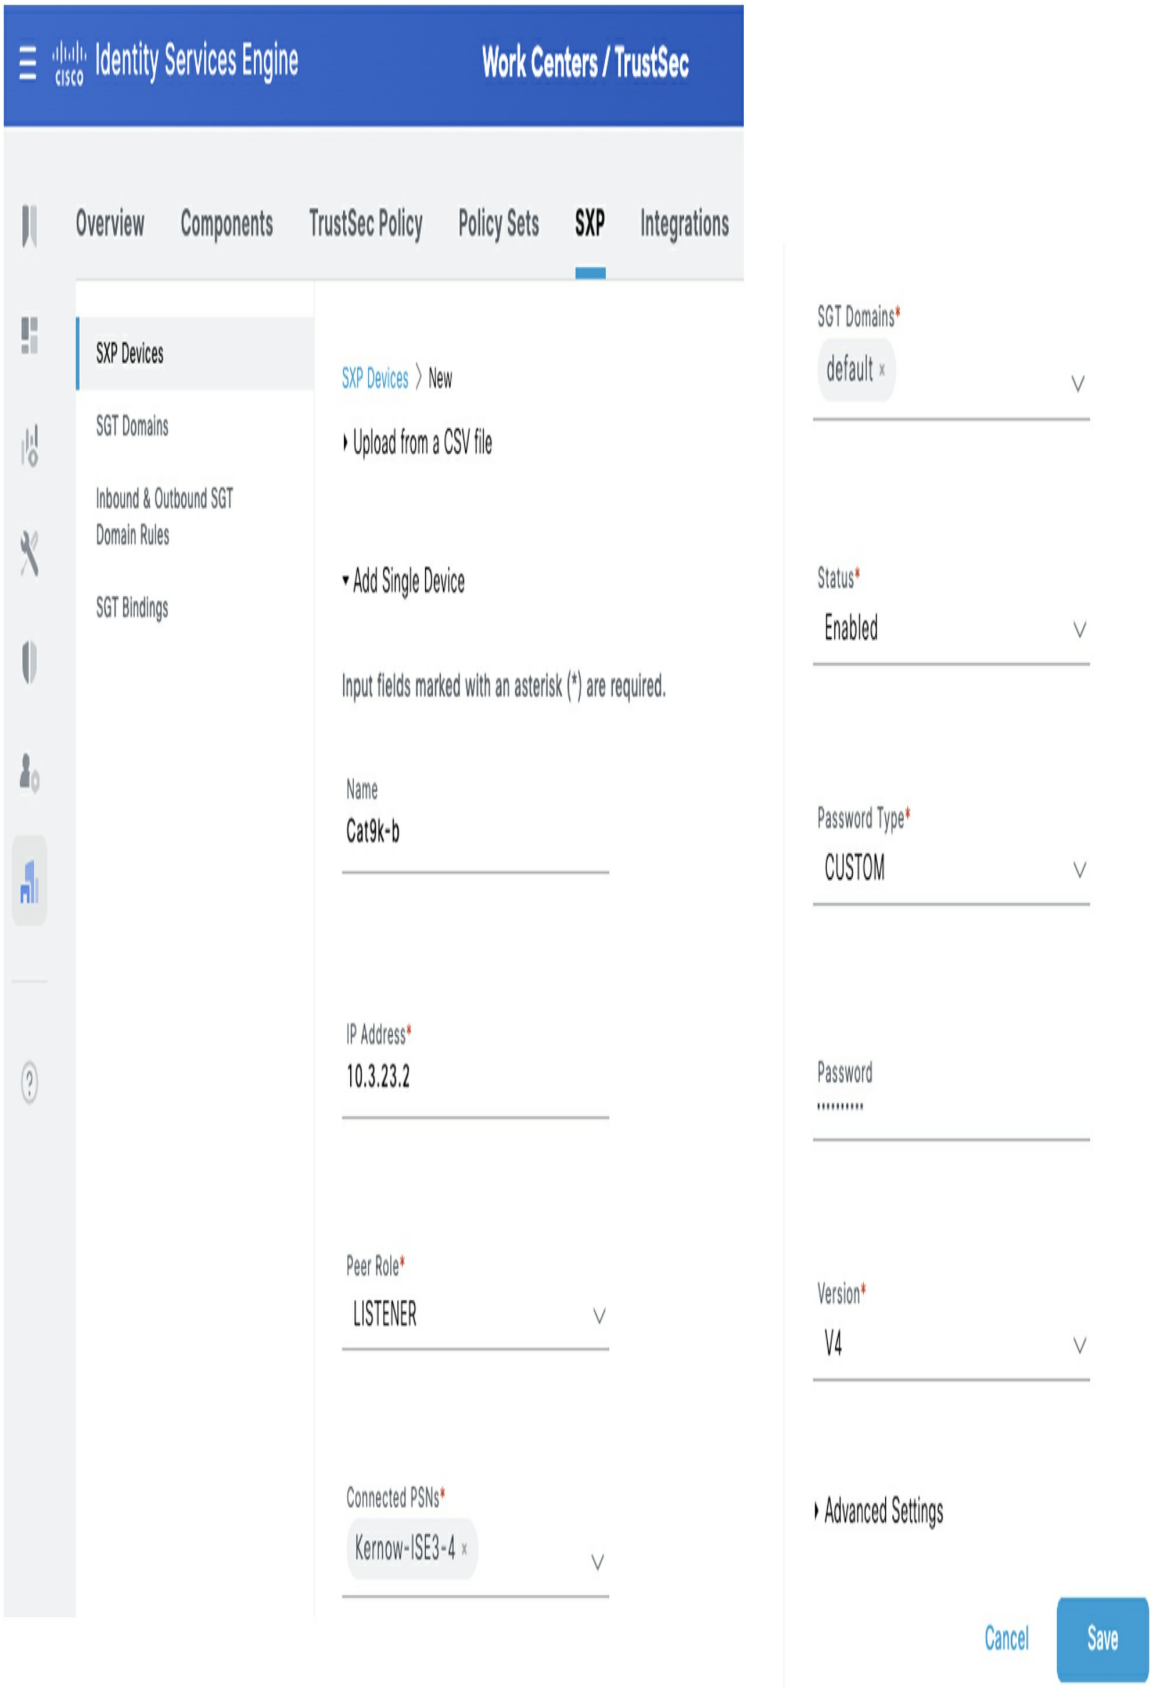
A sample command configuring an SXP listener on a network device follows:

```
Router(config)#cts sxp connection peer  source
password  mode local listener hold-time   vrf
```

As stated previously, bindings sit in a VRF on network devices, so setting the VRF is part of adding an SXP connection. If the VRF is not configured when adding a connection, then the default/global VRF is used.

The mode local setting within the commands specifies that the speaker or listener setting is set on the local device. As an alternative to using *local*, the term *peer* can be used to specify that the speaker or listener setting will be set for the peer device. As well as setting speaker or listener, the term *both* can be used to indicate bindings will be sent in both directions.

Figure 7-13 shows the options when utilizing the ISE UI. In ISE, an **SXP Device** is added, but it can be thought of as adding an SXP connection between ISE and the device. The **Peer Role** specifies whether the peer device is listening for bindings, speaking (or sending) bindings, or set to both speaker and listener. **SGT Domains** specifies which SGT domain bindings are either sent over this connection or the SGT domains to store the bindings when received from the peer; multiple SGT domains can be specified.


Identity Services Engine
Work Centers / TrustSec
Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
SXP Devices
SGT Domains
Inbound & Outbound SGT Domain Rules
SGT Bindings
SXP Devices 〉 New
Upload from a CSV file
Add Single Device
Input fields marked with an asterisk (*) are required.
Name
Cat9k-b
IP Address*
10.3.23.2
Peer Role*
LISTENER
Connected PSNs*
Kernow-ISE3-4
SGT Domains*
default
Status*
Enabled
Password Type*
CUSTOM
Password
Version*
V4
Advanced Settings
Cancel
Save

Technet24

Figure 7-13 *ISE SXP Device Configuration*

As can be seen in the commands listed earlier, some timers govern how SXP behaves. Different timers are available depending on whether the connection is a speaker or a listener. Similarly, Figure 7-14 shows the different timers available in ISE under the SXP Device **Advanced Settings** when the device is set for speaker, listener, or both.

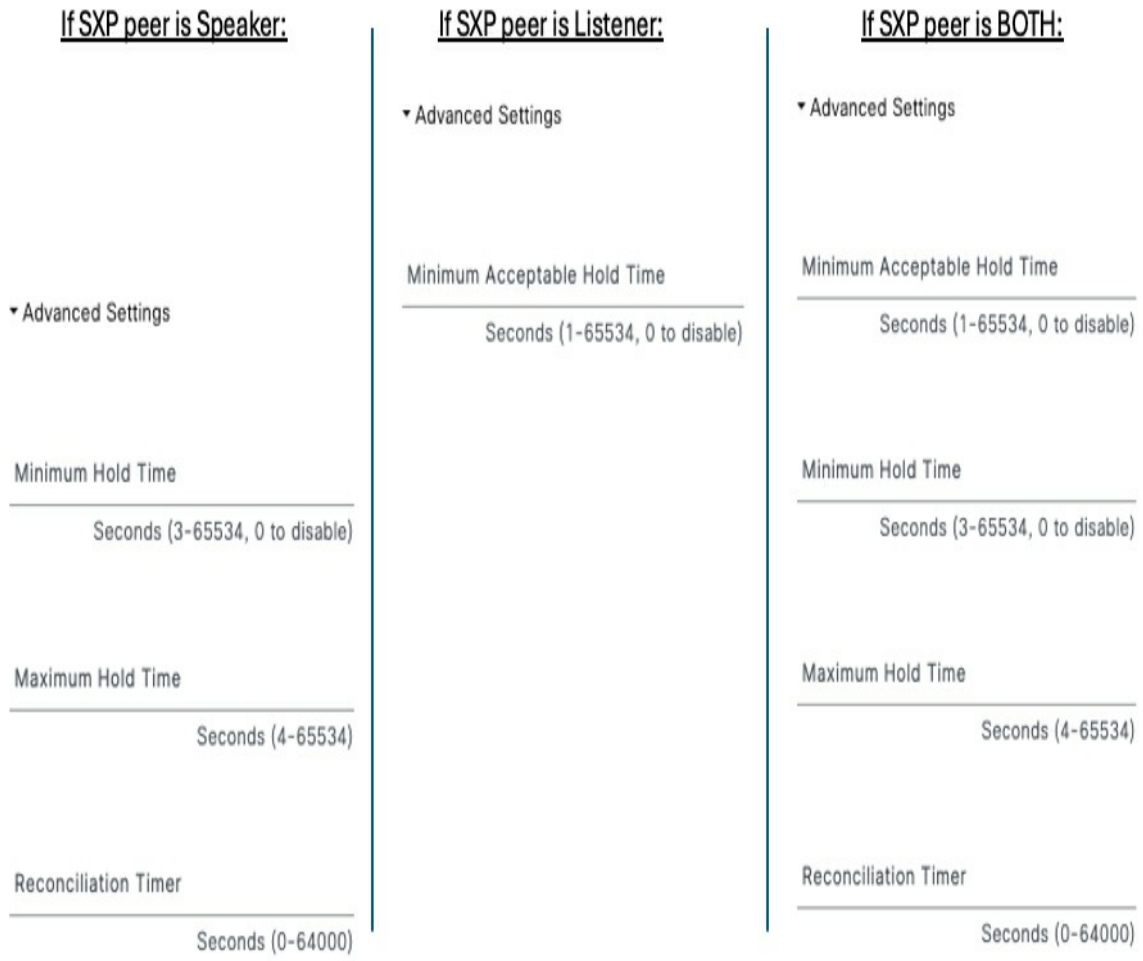
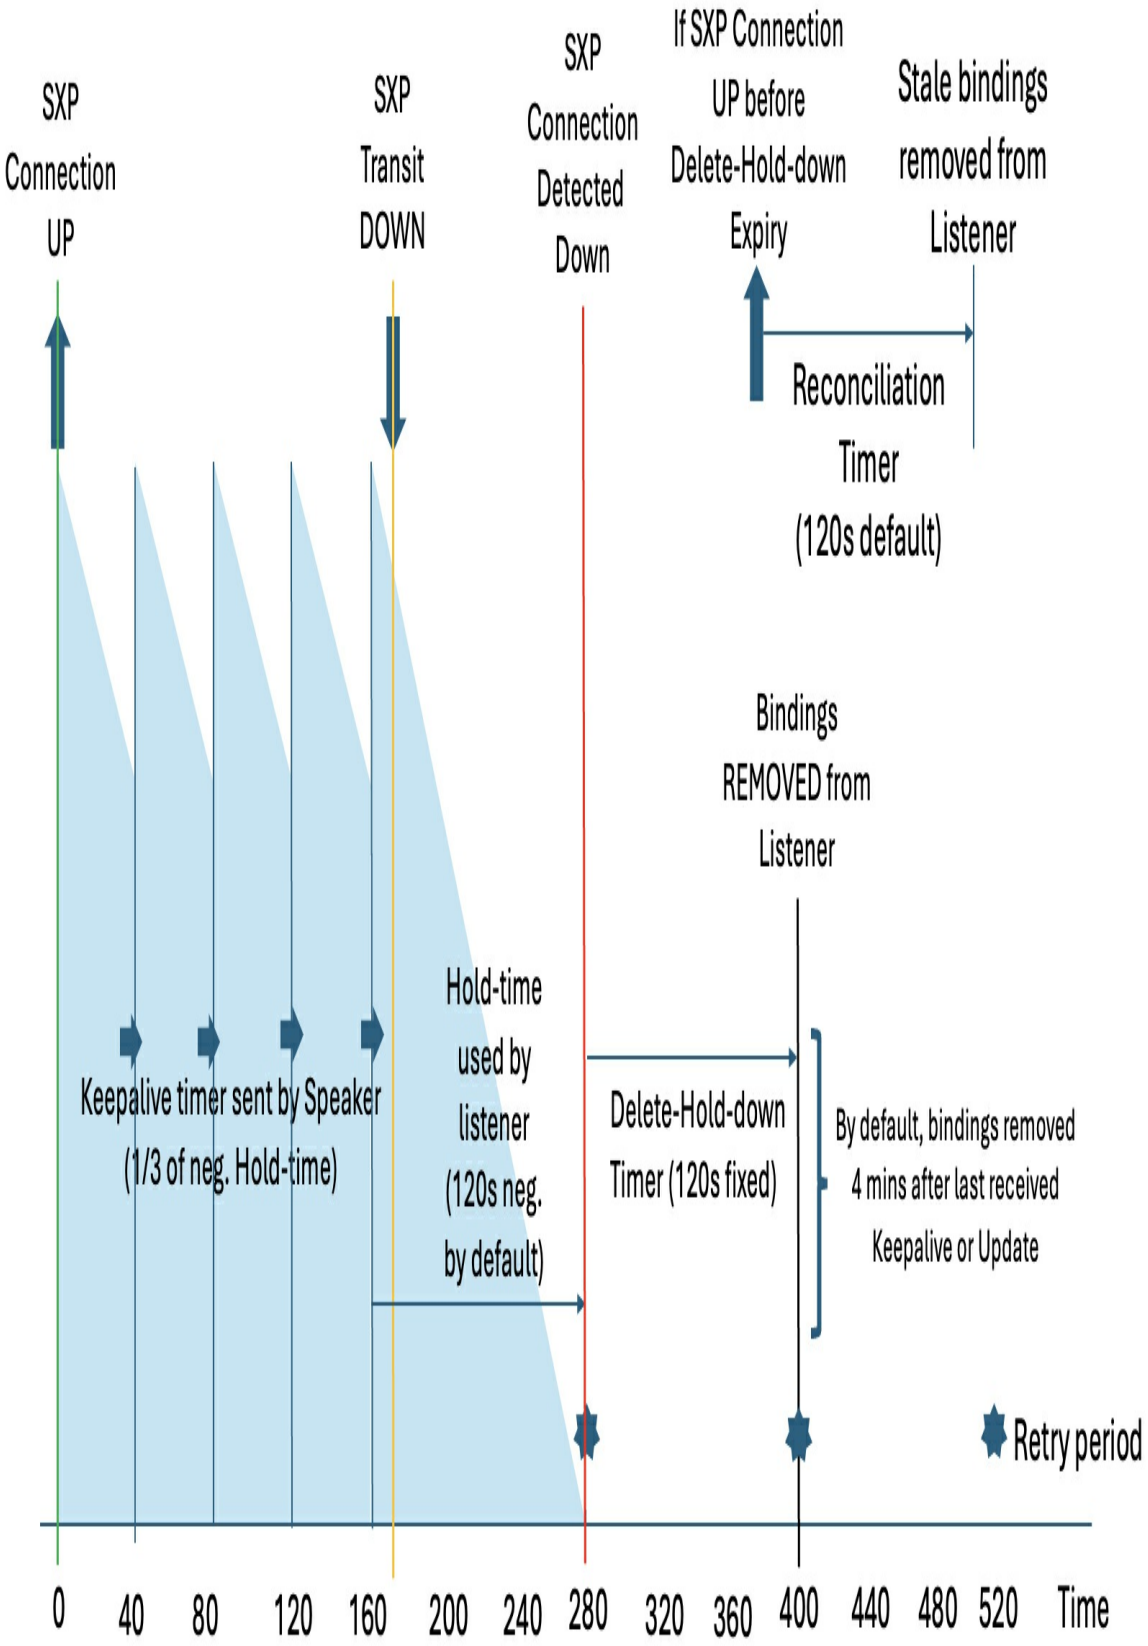


Figure 7-14 *ISE SXP Device Advanced Settings for Timer Details*

Figure 7-15 shows these timers and their interaction.


SXP Connection UP
SXP Transit DOWN
SXP Connection Detected Down
If SXP Connection UP before Delete-Hold-down Expiry
Stale bindings removed from Listener
Reconciliation Timer (120s default)
Bindings REMOVED from Listener
Keepalive timer sent by Speaker (1/3 of neg. Hold-time)
Hold-time used by listener (120s neg. by default)
Delete-Hold-down Timer (120s fixed)
By default, bindings removed 4 mins after last received Keepalive or Update
Retry period
0
40
80
120
160
200
240
280
320
360
400
440
480
520
Time

Technet24

Figure 7-15 *SXP Timers*

Once an SXP connection comes up, the peer configured as a speaker starts to send periodic keepalives to the listener. The period between keepalive messages is calculated as being one-third of the negotiated hold time.

The hold time is negotiated based on timer configuration settings. A global platform-wide hold time can be specified on the devices terminating the SXP connection, or the hold time can be set for each connection. Negotiation employs a minimum hold time setting on the speaker side and a minimum and maximum setting on the listener side. The configuration commands for both global and per connection hold time are as follows:

Global hold time:

- Speaker side: **cts sxp speaker hold-time**
- Listener side: **cts sxp listener hold-time**

Hold time per SXP connection:

- Speaker side: **cts sxp connection .. hold-time <0=min global / min 1-65534 secs>**
- Listener side: **cts sxp connection .. hold-time <0=min global / min 1-65534 secs> <0=max global / max 1-65534 secs>**

By default, the hold time negotiates to 120 seconds. This is shown in the shaded area in the lower left in Figure 7-15 along with the corresponding keepalive timer of 40 seconds.

If a transitory network device or WAN goes down carrying the SXP overlay, then further keepalives will not be received by the SXP listener. A transit outage is depicted in the figure at just over 160 seconds. Only when the hold time reduces to 0 will the network

devices detect the outage, and the Retry timer and Delete-Hold-Down timers commence. The Retry period is configurable (0–64,000 seconds) with a default of 120 seconds and is just the time between SXP connection reestablishment attempts. The Delete-Hold-Down timer is more critical though, because upon expiry, all IP TO SGT bindings received through this connection (received by the listener) will be deleted. This timer is fixed at 120 seconds.

We can deduce, then, that if the SXP connection remains down, received IP TO SGT bindings will be deleted after hold time expiry plus Delete-Hold-Down-time expiry, which in total, by default, equates to 240 seconds, or 4 minutes. Because the Delete-Hold-Down-time is not configurable, the only way to increase this total time is to increase the hold time.

The only other timer shown in Figure 7-15 is the Reconciliation time. This timer controls what happens to bindings if an SXP connection comes back up before the Delete-Hold-down expiry time. A number of minutes will have elapsed without binding propagation, and therefore, the bindings within the device on the speaker side may have changed without the receiver side being aware of those changes. The Reconciliation time specifies when any stale bindings would be removed from the listener.

ISE SXP Propagation (Including SGT Domains and Filtering)

SXP is not just supported on network devices like switches, routers, and wireless LAN controllers, but ISE itself can be an SXP speaker and/or listener. SXP is enabled on the ISE Policy Service Node (PSN), and the recommendation is to reserve dedicated PSNs for this feature. Up to four SXP PSN pairs (eight PSNs in total) can be enabled to

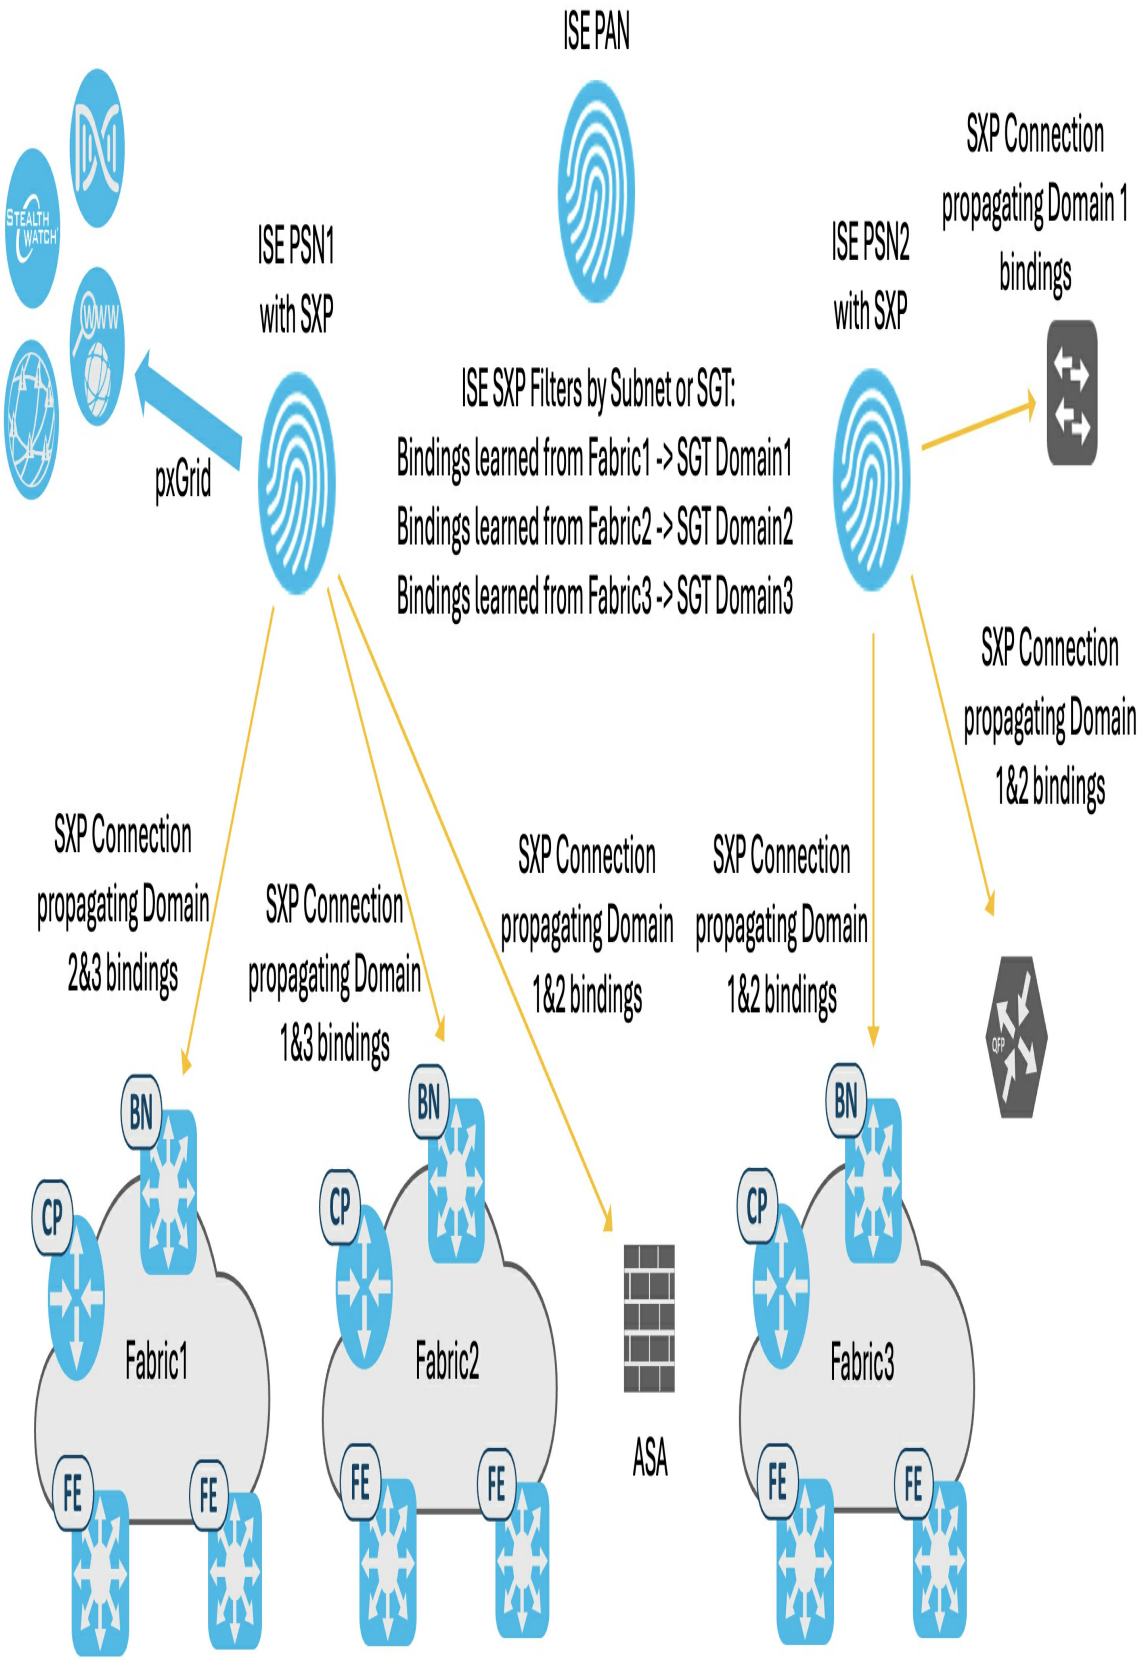
Technet24

support high availability and scale.

While SXP is highly effective at propagating group tag mappings across the network, the sheer volume of these mappings in a large-scale enterprise can create significant overhead. This is where SXP filtering becomes a critical design consideration. Without filtering, every SXP listener in the network potentially receives the entire database of SGT mappings, which consumes valuable memory and processing resources on access-layer devices, for example, that have no functional need for that specific intelligence. By implementing SXP filters, you can restrict the propagation of SGT mappings to only those that are relevant to a specific segment or set of devices, effectively reducing the footprint of the SXP table and minimizing the risk of unnecessary information exposure.

ISE uses the concept of domains to filter SGT bindings. SXP connections or devices on ISE are assigned to a domain, or a number of domains, when added to the system. Static IP TO SGT bindings can be configured on ISE (under **Work Centers > TrustSec > Components > IP SGT Static Mapping**), and when added, the domain or domains can be specified. When dynamic mappings or bindings are formed (through authorization), those bindings are, by default, automatically sent to the default SGT domain.

The settings on ISE shown in Figure 7-8 determine whether SGT bindings are sent to pxGrid and whether RADIUS mappings are sent to the SGT table. Filters are available on ISE to detect these dynamic bindings (by subnet or SGT) and move them to a different domain if desired. Bindings within an SGT domain are sent to network devices where that domain is specified in the connection/devices configuration. Figure 7-16 shows this concept within a multi-fabric SD-Access use case utilizing an IP transit between fabrics.


ISE PAN
ISE PSN1
with SXP
ISE PSN2
with SXP
SXP Connection
propagating Domain 1
bindings
pxGrid
ISE SXP Filters by Subnet or SGT:
Bindings learned from Fabric1 -> SGT Domain1
Bindings learned from Fabric2 -> SGT Domain2
Bindings learned from Fabric3 -> SGT Domain3
SXP Connection
propagating Domain
1&2 bindings
SXP Connection
propagating Domain
2&3 bindings
SXP Connection
propagating Domain
1&3 bindings
SXP Connection
propagating Domain
1&2 bindings
SXP Connection
propagating Domain
1&2 bindings
BN
CP
FE
Fabric1
Fabric2
Fabric3
ASA

Technet24

Figure 7-16 *ISE Propagation Methods Including SGT Domains*

Dynamic bindings created from endpoints within Fabric1 are automatically placed in the ISE default SGT domain, but a configured SXP filter (by subnet or SGT) moves them into an SGT domain called Domain1. The same happens for endpoints within Fabric2 and Fabric3 and moved to SGT domains called Domain2 and Domain3, respectively.

To reclassify inbound traffic to each fabric, each fabric border only needs to learn of bindings from the remote sites. To achieve this, the SXP connection from ISE to the border only needs to propagate the bindings within the SGT domains for the remote sites. In Figure 7-16, bindings in Domain2 and Domain3 are sent to the Fabric1 border, bindings in Domain1 and Domain3 are sent to the Fabric2 border, and bindings in Domain1 and Domain2 are sent to the Fabric3 border.

Note that the bindings within SGT Domain1 and Domain2 are sent to the ASA firewall. Additionally, specific domain mappings are sent to the network devices, as shown on the right side of Figure 7-16.

The ISE PAN manages all the SGT bindings. In the figure, PSN2 has SXP connections or devices only for Domain1 and Domain2. For this reason, the PAN sends bindings only within those two domains to PSN2.

The ISE GUI displays bindings only for the domains that are specified in enabled SXP connections or devices. This leads to the SGT binding table showing no entries if no connections or devices are added. If bindings are required to be sent to pxGrid and there are no SXP devices to send mappings to, to populate the SGT Mappings table, you can add a dummy SXP connection or device for the required SGT domains. You should use a dummy/unused IP address as the device IP address (that connection does not have to be ON/UP to populate the

mapping table and allow propagation over pxGrid).

SXP Filtering on Network Devices

SXP filtering is also available on network devices. These devices can be applied on a peer-to-peer basis or globally, which caters for all SXP connections on a device. Filtering uses an IP prefix or an SGT and can act on either a speaker or listener connection. Configuration utilizes an SXP filter list, containing the rules, and an SXP filter group tying the filter list to the peer. Following is a configuration example:

```
cts sxp filter-enable
cts sxp filter-list filter-list-example
  deny sgt 4
  deny ipv4 10.150.90.0/24 deny sgt 5
cts sxp filter-group speaker filter-group-example
  filter filter-list-example
  peer ipv4 10.3.18.10
```

You can use the following command for filter visibility:

```
Router#show cts sxp filter-group speaker detailed
Global Speaker Filter: Not configured
Filter-group: filter-group-example
    Filter-name: filter-list-example
    Filter-rules:
        10 deny sgt 4 (1)
        20 deny ipv4 10.150.90.0/24 deny sgt 5 (0)
    Total Matches: 1
    Default Deny Count: 0
    peer 10.3.18.10
```

Propagation Summary

Figure 7-17 shows the different propagation methods. No recommendation is portrayed here; this figure is purely to indicate the method possibilities.

Technet24

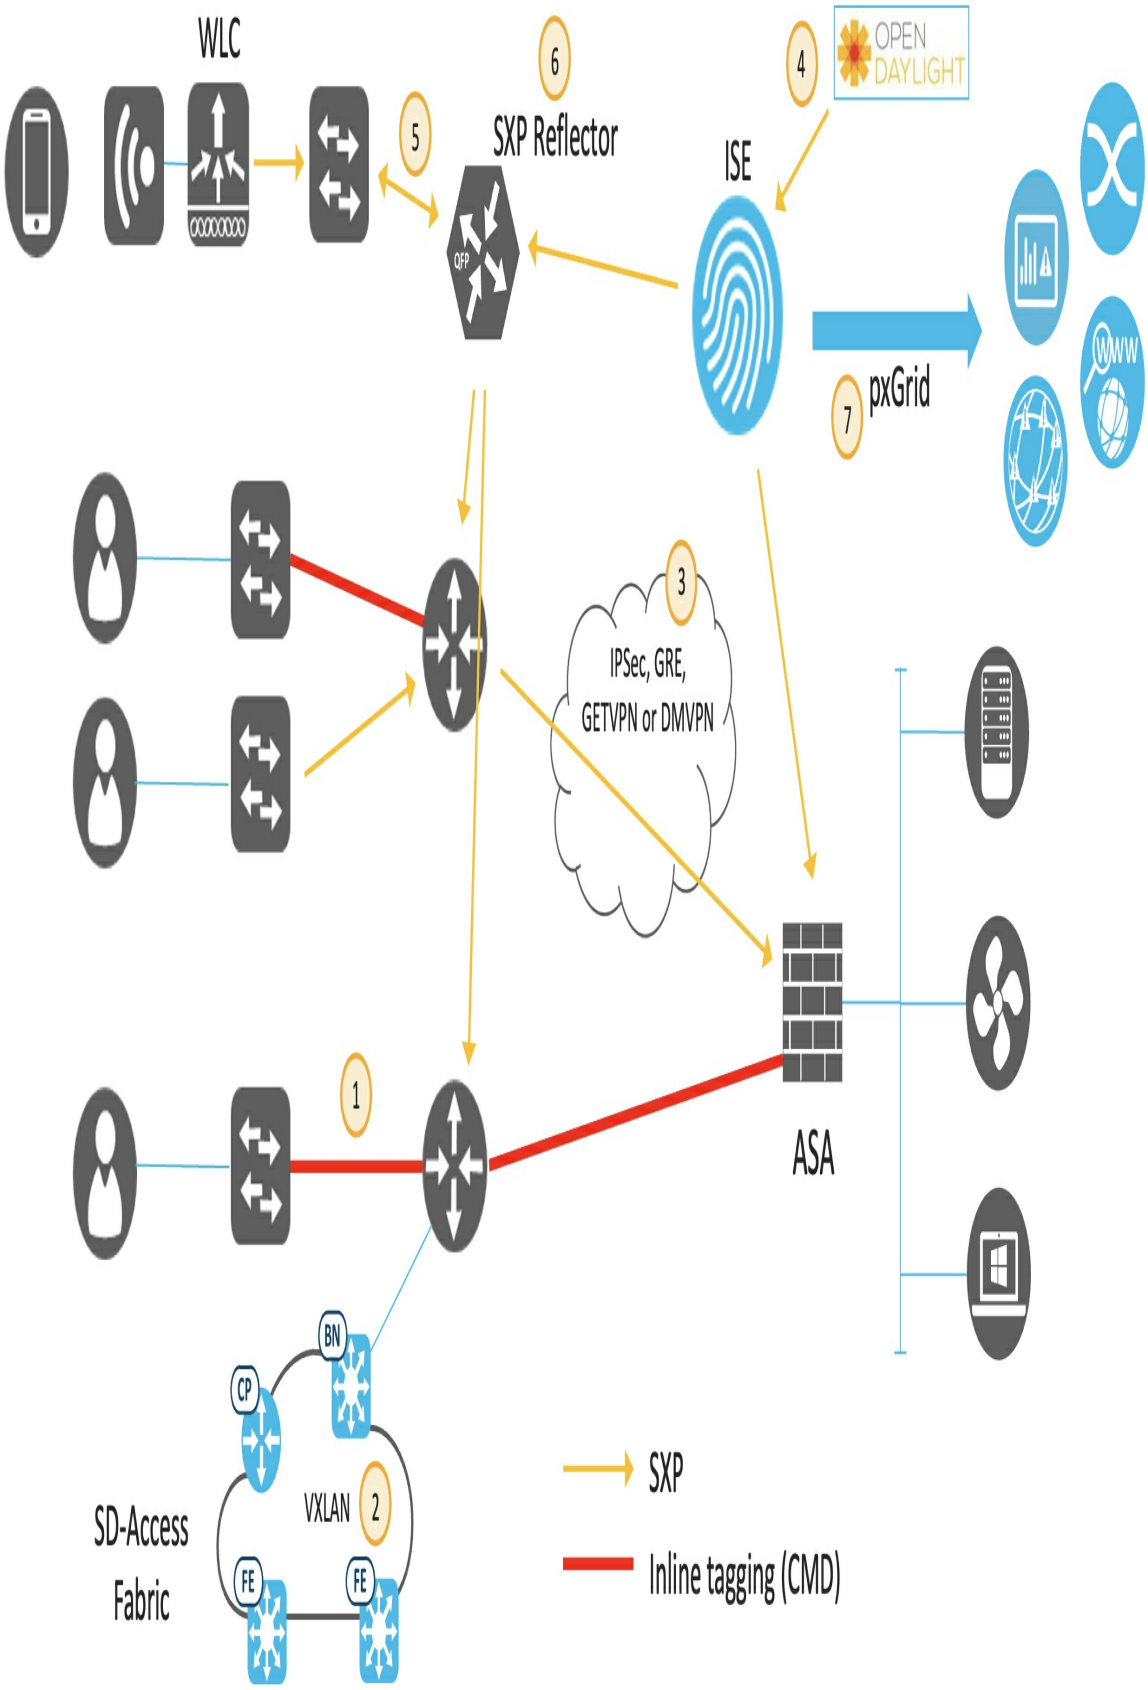

WLC
6
SXP Reflector
5
4
OPEN DAYLIGHT
ISE
QFP
pxGrid
7
3
IPSec, GRE,
GETVPN or DMVPN
ASA
1
BN
CP
VXLAN
2
FE
FE
SD-Access
Fabric
SXP
Inline tagging (CMD)

Figure 7-17 *Propagation Methods*

The functions in Figure 7-17 are indicated by the numbers:

1. The red lines illustrate the use of inline tagging (source SGT carried in the data plane within the CMD field of the Layer 2 frame). Inline tagging should be used where possible as propagation is carried out in hardware and at line rate.
2. The source SGT is also carried in the data plane across an SD-Access fabric within VXLAN.
3. If the SGT needs to be propagated across a WAN, then other data plane propagation techniques can be used, like inserting the source SGT into headers of IPsec, GRE, GETVPN, or DMVPN.
4. Both inline tagging and SXP specifications have been made available to vendors via the Internet Engineering Task Force (IETF) and Request for Comment (RFC) documentation. OpenDaylight is one platform that has made use of these publications, and the SXP protocol has been implemented. Note that Figure 7-17 shows SXP bindings being received by ISE from OpenDaylight and other SXP connections being sourced from ISE. That is purely to make it clear that ISE can send SXP bindings (SXP configured as a speaker), receive SXP bindings (SXP configured as a listener), or send and receive bindings (SXP configured as both ways).
5. This SXP connection shows that bindings can be sent in both directions if required.
6. SXP bindings can be sent from hop to hop or from peer to peer. If there are many SXP connections from one region or location to another, then sometimes it's best to send bindings via one SXP connection to the destination region and fan out from there.

Technet24

The platform carrying out this fan-out function is often called an SXP reflector.

7. Identity, session, and SGT information is propagated via pxGrid. This allows systems such as Cisco Secure Network Analytics, Cisco Catalyst Center, Cisco Secure Web Appliance, Cisco Secure Firewall Management Center, and many ISE ecosystem platforms to gain insight into all this context.

Enforcement

Enforcement in a Cisco TrustSec network refers to the process where network devices apply access control policies based on SGTs assigned to source and destination devices. This approach enables scalable, identity-based policy enforcement that is topology-independent.

Enforcement Introduction

Group-based enforcement occurs on network devices where enforcement is enabled. If not enabled, then group-based policy is not downloaded, and no enforcement takes place on that device. Enforcement can be enabled globally on L3 interfaces, individually on VLANs, and can be disabled per interface if required.

An enforcement policy consists of a source SGT, a destination SGT, plus a security group access control list (SGACL) that holds the details of which TCP and UDP ports to permit or deny. In some systems such as Cisco Catalyst Center, an SGACL is known as a contract. An SGACL is made up of one or multiple security group access control entries (SGACEs). As well as SGACLs being used between source and destination SGTs, an SGACL also is used as a default enforcement, the default being either a permit or deny.

When you're operating with a default permit policy, enforcement

occurs on a platform only when the following four aspects are true:

1. Enforcement is enabled.
2. The source IP TO SGT binding is known/learned.
3. The destination IP TO SGT binding is known/learned.
4. There is a policy present to enforce.

> **Note:**
> When operating with a default deny policy, then a platform will also enforce traffic that is unclassified, implicitly assigned the Unknown classification with SGT 0.

Enforcement Onboarding

When enforcement is enabled, CTS onboarding takes place, preparing the network device for enforcement. Onboarding includes downloading a protected access credential (PAC) if used, and the CTS environment data from ISE.

The PAC is a unique shared credential used for mutual EAP-FAST authentication between a network device and an authentication server (like Cisco ISE) without requiring a full public key infrastructure (PKI). It creates a tunnel to download environment data and security policies. However, using a PAC does not provide adequate security because MD5 is utilized; plus, customers have experienced PAC issues related to failed downloads and refresh times. As a result, from ISE release 3.4 and IOS-XE 17.15.1, a PAC-less mode is supported, eliminating the use of EAP-FAST and PAC and solely relies on a shared secret for network device to ISE authentication. There is no configuration on ISE to support the PAC-less methodology; the network devices are just configured to use a PAC or not under the **radius server** configuration.

Technet24

ISE 3.4 accepts either request, validates the shared secret, and acts appropriately as requested by the network device.

When authenticated and a PAC downloaded, if used, the onboarding procedure moves on to downloading the environment data. This includes the following important components:

1. CTS server list
2. Local network device SGT
3. List of all security group tags (SGTs) in ISE
4. Timers

When downloaded, the environment data can be displayed using the command shown in Example 7-7.

Example 7-7 *Network Device Downloaded Environment Data*

```
switch#show cts environment-data
CTS Environment Data
====================
Current state = COMPLETE
Last status = Successful
Service Info Table:
Local Device SGT:
  SGT tag = 2-01:TrustSec_Devices
Server List Info:
Installed list: CTSServerList1-0001, 1 server(s):
 *Server: 10.1.101.37, port 1812, A-ID 61FFFBA027C503F5E0F00126B7
          Status = ALIVE
          auto-test = TRUE, keywrap-enable = FALSE, idle-time = 6
secs
Security Group Name Table:
    0-01:Unknown
    2-01:TrustSec_Devices
    3-01:Network_Services
    4-05:Employees
    ….
```

```
    14-01:PCI_Servers
    15-04:BYOD

Environment Data Lifetime = 86400 secs
Last update time = 09:45:36 UTC Tue Feb 17 2026
Env-data expires in   0:21:24:16 (dd:hr:mm:sec)
Env-data refreshes in 0:21:24:16 (dd:hr:mm:sec)
Cache data applied           = NONE
State Machine is running
Retry_timer (60 secs) is not running
switch#
```

You can define the CTS server list within ISE under **Work Centers > TrustSec > Components > TrustSec Servers > TrustSec AAA Servers**, as shown in Figure 7-18.

Technet24

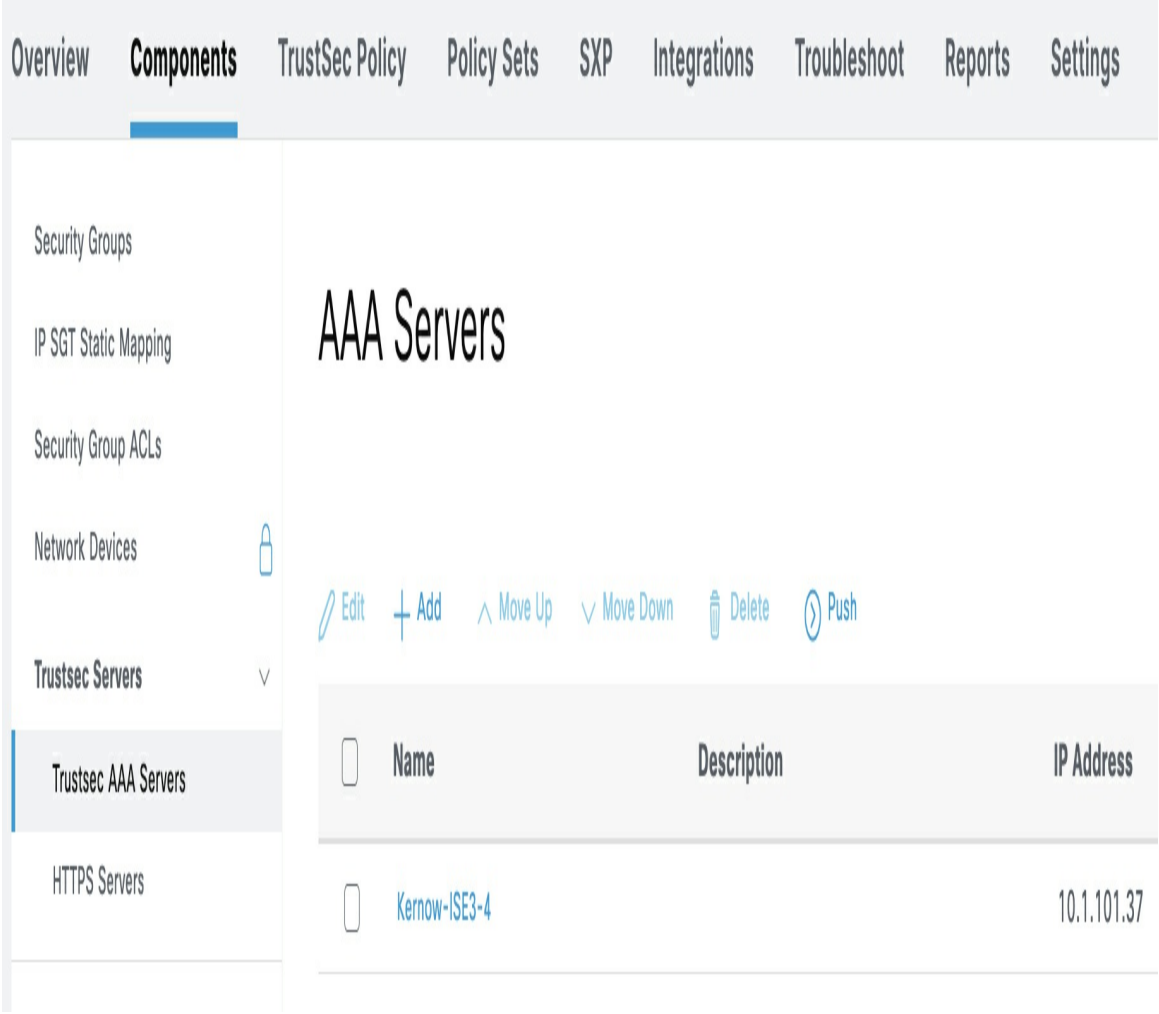


Figure 7-18 *TrustSec AAA Servers*

Typically, there would be multiple entries in the AAA Servers table in ISE, one entry for each PSN from which network devices would request policy. The network device uses the first entry in the list to request policy. If that PSN is unavailable for any reason, then the second entry in the list is tried, and so on.

Rather than using AAA to download policy, ISE also supports a more secure and reliable HTTPS method of downloading policy.

The Local Network Device SGT allows the network device to assign an SGT to traffic generated by the network device itself. This is assigned in ISE under **Work Centers > TrustSec > TrustSec Policy > Network**

Device Authorization (see Figure 7-19). By default, ISE allocates the SGT named TrustSec_Devices with an SGT of 2 in this table, but as administrator, you can choose any SGT you wish and even assign different SGTs to different network devices based on rule conditions.



Figure 7-19 *ISE Table to Assign Network Device SGT*

The list of all the security group tags (SGTs) in ISE is downloaded by the network devices in the environment data. The names and numbers are useful if you're building policies via the CLI and for show command visibility.

Lastly, the environment data shows relevant timers along with expiry and refresh times.

After the PAC (if used) and the environment data have been downloaded, then the relevant policy is also downloaded. The policy download mechanism is explained later in Figure 7-32, but no matter whether SGTs are assigned to users and endpoints or not, the default policy is downloaded as part of onboarding. The default policy is defined in ISE at the top of the TrustSec policy matrix, as shown to the right of Figure 7-20.

Technet24



Figure 7-20 *Default Policy Shown with TrustSec Matrix*

The default policy can be a permit, whereby all traffic is permitted unless explicitly denied with defined rules, or can be a deny, whereby all traffic is denied unless explicitly permitted with defined rules. When ISE is first installed, the default policy is set as a permit. If you change this to a deny, then you must make sure that explicit permit policies are in place to ensure networking protocols and management are not mistakenly blocked.

ISE initially has one policy matrix. All policy requests that come into ISE are serviced by this one matrix. Under **Work Centers > TrustSec > Settings > Work Process Settings**, you can enable the use of multiple policy matrices, as shown in Figure 7-21. (If ISE is integrated with Catalyst Center, then multiple matrices are not supported at the time of writing.)

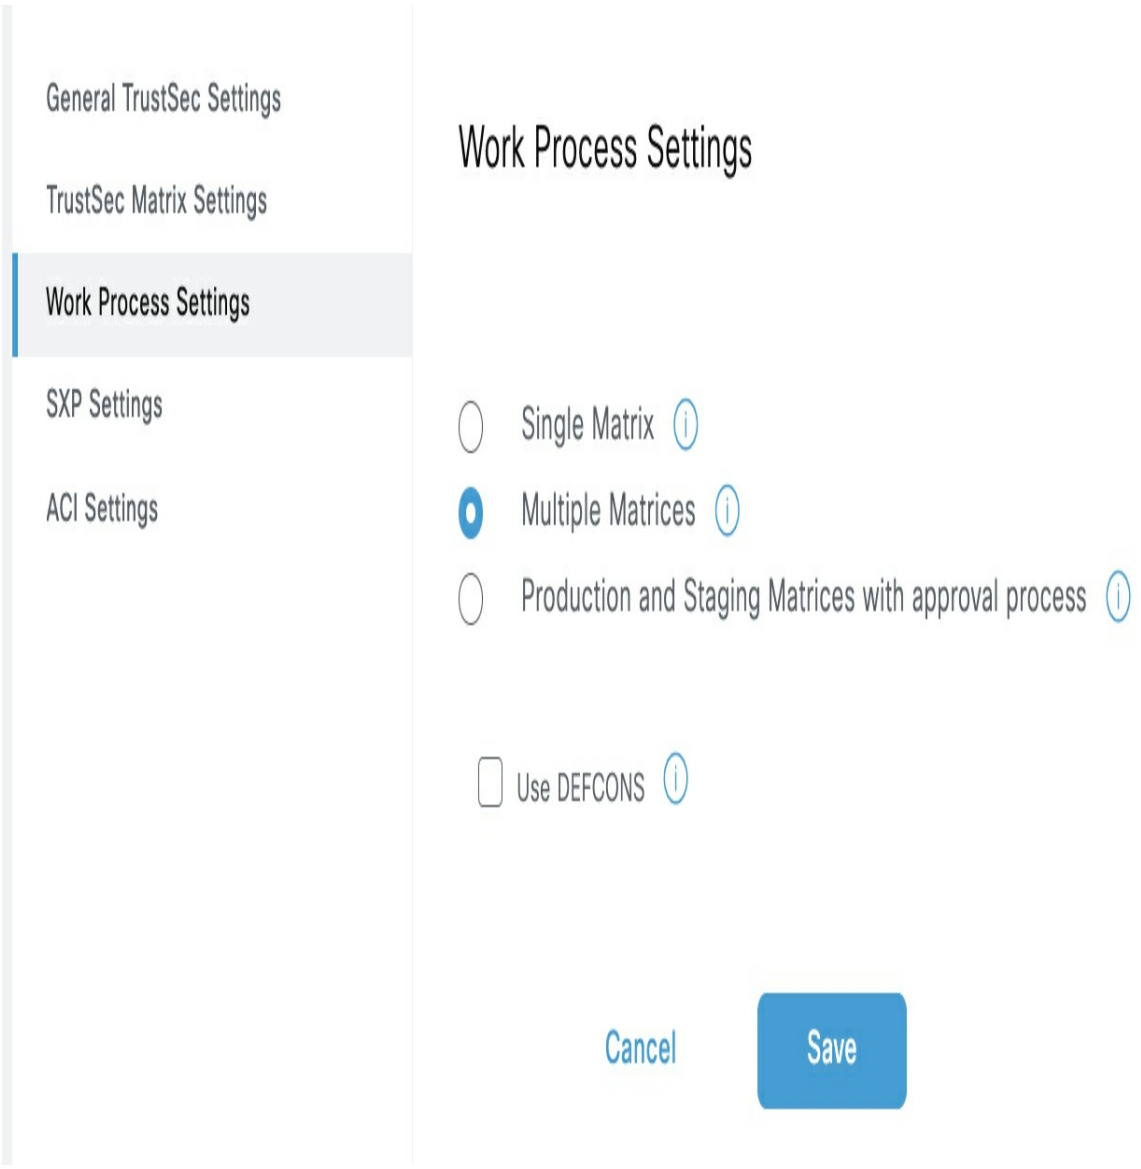


Figure 7-21 *ISE Setting for Enabling Multiple TrustSec Matrices*

When enabled, the default matrix under **Work Centers > TrustSec > TrustSec Policy > Egress Policy > Matrices List** becomes known as the Production matrix.

Using that UI, you can add further matrices by providing a matrix name and copying policies from existing matrices.

When saved, the new matrix is visible, but no policies are downloaded from that new matrix until further configuration is made.

Technet24

As can be deduced from Figure 7-22, the policy downloaded is related to the NAD making the download request and the policy matrix that NAD is assigned to.

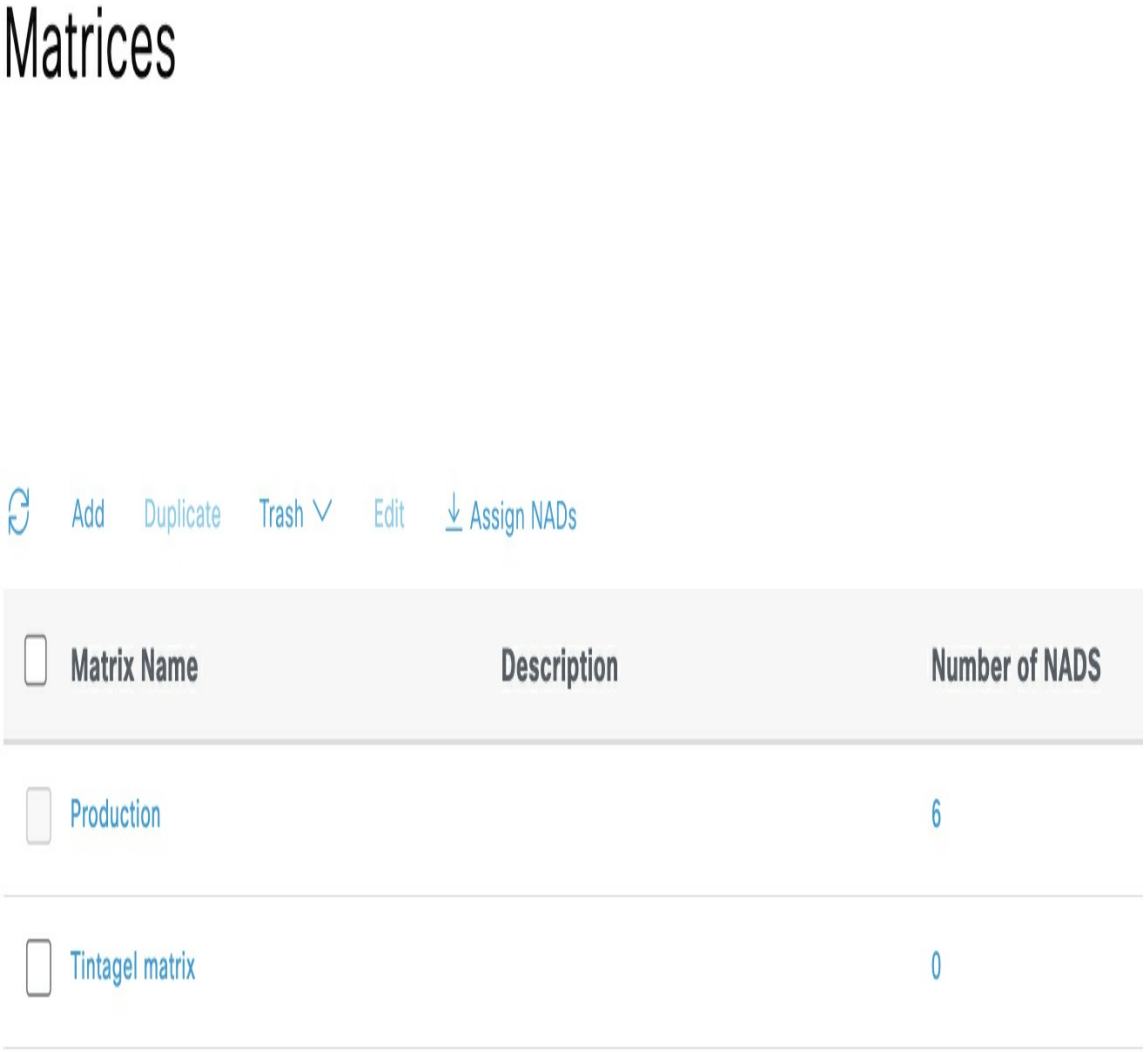


Figure 7-22 *Matrix List Showing Newly Added Matrix*

Using the **Assign NADs** function in Figure 7-24, you can assign NADs to different matrices. In Figure 7-23, you can select the NADs at the top of the UI and assign the matrix at the bottom.

Assign Network Devices

1 Select network devices. (Filters may be used)

Rows/Page 6 | 1 /1 | Go | 6 Total Rows

Filter

	Name	IP	Location	Type	Matrix
X	Name	IP	Location	Type	
☐	Kernow-3850-orig.kernow.com	10.1.1.9/32	Location#All Locations	Device Type#All Device Types	Production
☑	Cat9k-NT.Kernow.com	10.1.1.14/32	Location#All Locations	Device Type#All Device Types	Production
☐	Cat9k-NB.Kernow.com	10.1.1.15/32	Location#All Locations	Device Type#All Device Types	Production
☑	Kernow-3850.kernow.com	10.1.1.4/32	Location#All Locations	Device Type#All Device Types	Production
☐	Kernow-Cat9300-b.kernow.com	10.1.1.8/32	Location#All Locations#NDG-C...	Device Type#All Device Types	Production
☐	Kernow-Cat9500	10.1.1.11/32	Location#All Locations	Device Type#All Device Types	Production

2 Assign these to a matrix Tintagel matrix×

Close Assign

Technet24

Figure 7-23 *Assigning NADs to TrustSec Matrices*

Once assigned, the NAD assignments are visible in the matrix list, as shown in Figure 7-24.

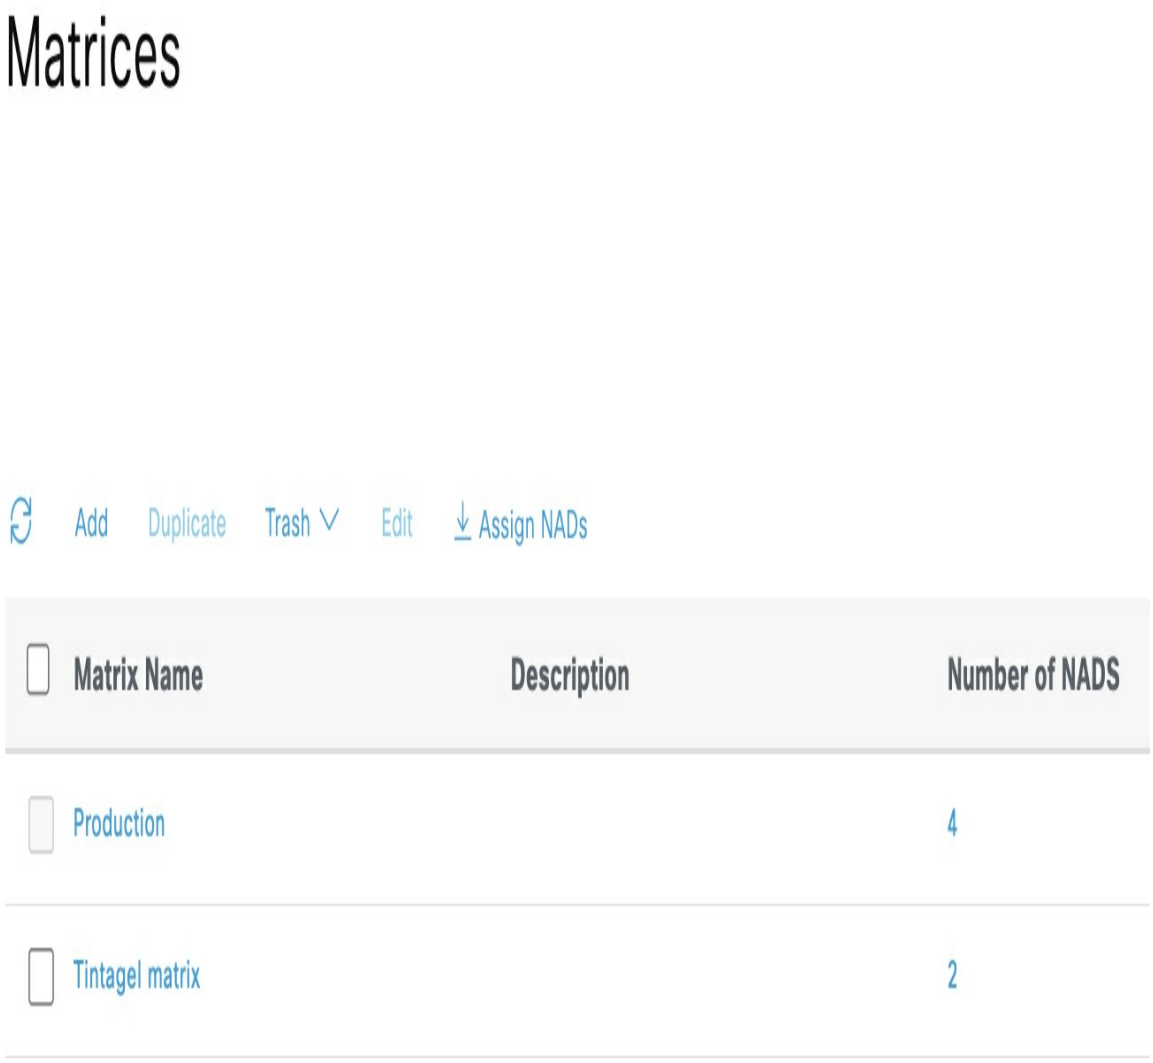


Figure 7-24 *Matrix List Showing NADs Assigned*

You can display and manage the different matrices by navigating to **Work Centers > TrustSec > TrustSec Policy > Egress Policy > Matrix**. The drop-down box shown in Figure 7-25 enables you to display the required matrix.

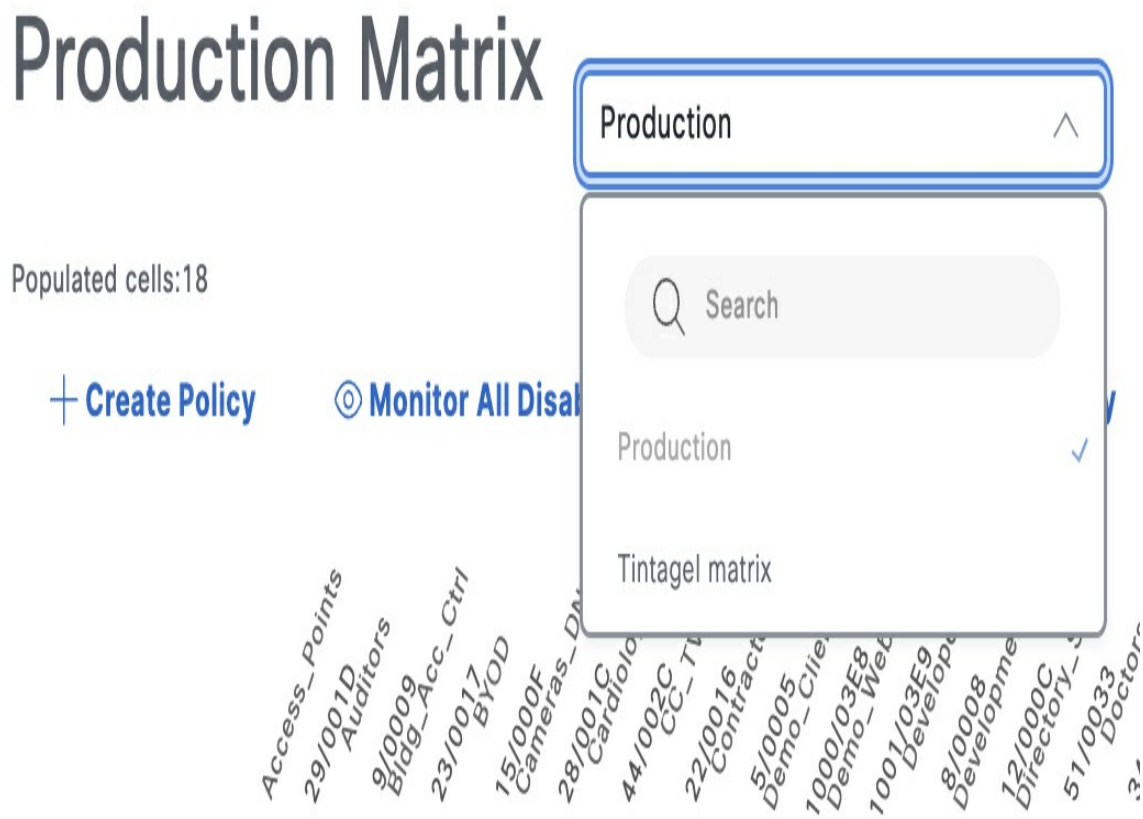


Figure 7-25 *Selecting the TrustSec Matrix to Be Displayed*

When a matrix is displayed, you can add policies by clicking a matrix cell or by using the **+Create Policy** option at the top of the matrix. When you create a policy, you select a source SGT plus a destination SGT, and then choose a security group access control list (SGACL) and/or a final catch-all rule. Both an SGACL and a final catch-all rule can be selected only if multiple SGACLs are allowed under **Work Centers > TrustSec > Settings > TrustSec Matrix Settings**, as shown in Figure 7-26. (If ISE is integrated with Catalyst Center, then multiple SGACLs are not supported at the time of writing.)

Technet24

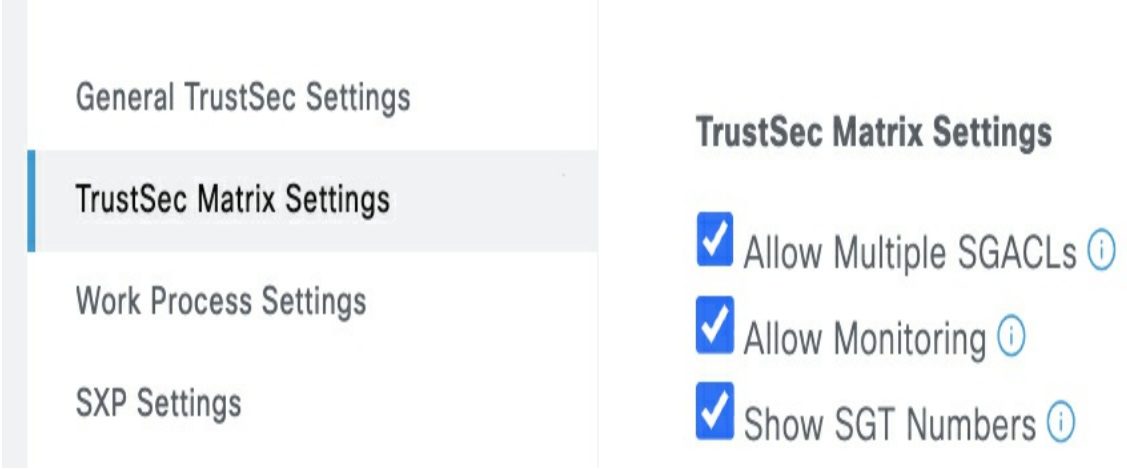


Figure 7-26 *ISE Setting to Allow Multiple SGACLs per Policy*

By default, two SGACLs exist in ISE from installation — namely, Permit IP and Deny IP. To add further SGACLs permitting or denying specific UDP or TCP ports, navigate to **Work Centers > TrustSec > Components > Security Group ACLs**.

When new entries are added, the ACL content field is free-form and is not currently validated by ISE. The content must be in the correct format to be accepted by the network devices. A sample format is shown here:

```
permit udp src eq domain
deny tcp dst eq 21
```

At the time of writing, an ISE feature has been developed to validate the SGACL syntax, and it will be made generally available in an upcoming release.

Figure 7-27 shows a sample pop-up UI when a policy is added via the matrix. The source and destination security group are selected along with an SGACL.

Create Security Group ACL Mapping...

Source Security Group:

Auditors

Destination Security Group:

Contractors

Status Enabled

Description

Assigned Security Group ACLs

Anti_Malware

Final Catch All Rule None

Technet24

Figure 7-27 *Creating a TrustSec Security Policy*

When adding a policy, you can specify its status. In addition to enabled and disabled, you can configure a status of **Monitor**, as shown in Figure 7-28. This status allows the downloaded policy to be in monitor mode; if, for example, the monitor mode is a deny statement, then the deny action is visible, but traffic is not blocked.

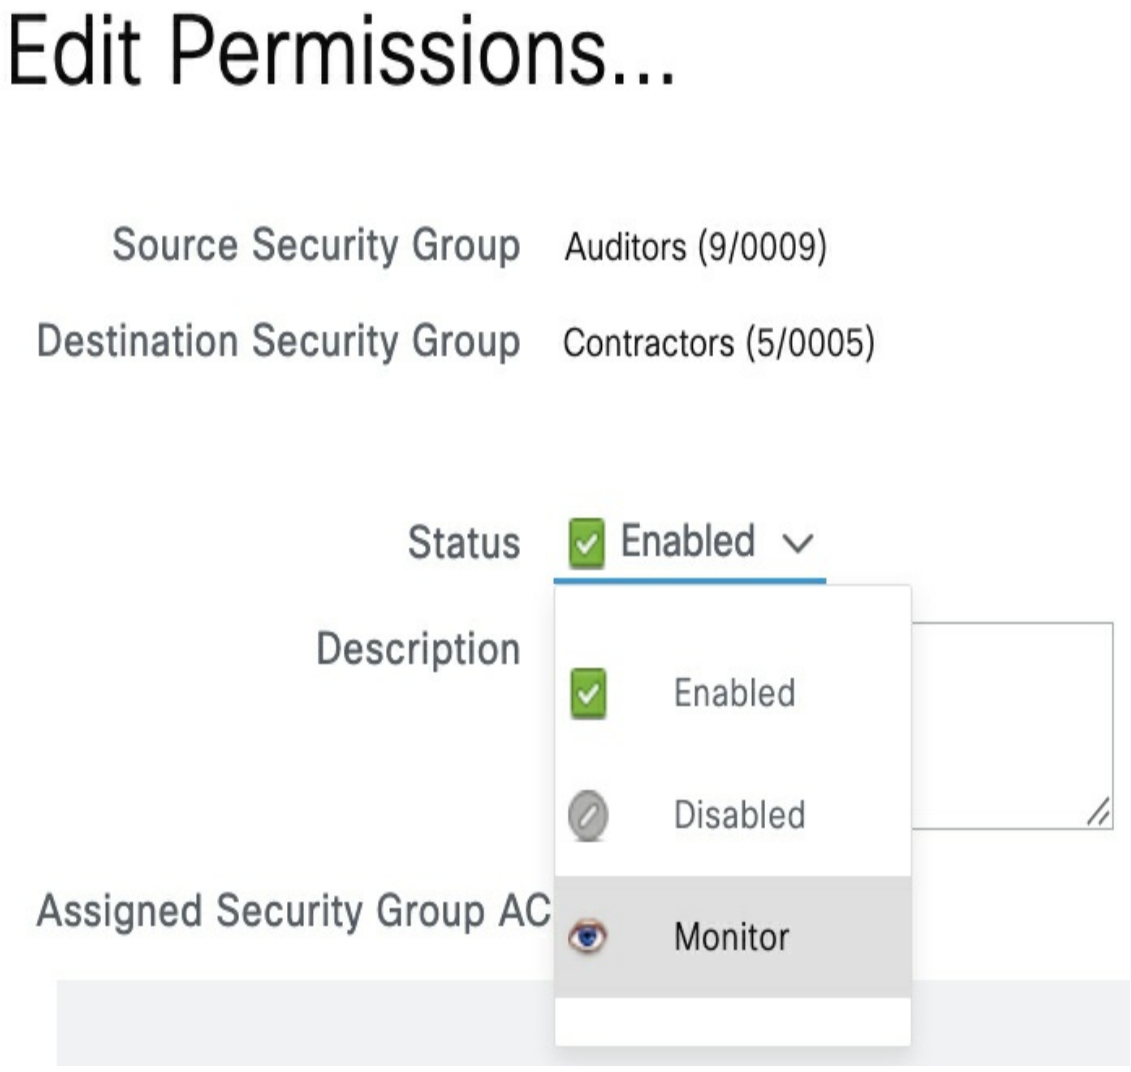


Figure 7-28 *A TrustSec Policy Can Be Set in Monitor Mode*

Once added, the policy is shown in the matrix, as shown in Figure 7-29.



Figure 7-29 *Newly Added Policy Displayed in the TrustSec Matrix*

Technet24

When policies are added and saved to the matrix, network devices are not immediately aware of the change. By default, timers on the NADs refresh their policies every 24 hours. To make network access devices immediately aware of recent changes, you can use the **Deploy** function at the top of the matrix. When **Deploy** is selected, that selection instructs ISE to send a RADIUS Change of Authorization to the network devices to inform them of the changes and to allow those changes to be downloaded. Best practice is to make just a few policy changes between instigating Deploy functions; in other words, use Deploy frequently. This reduces the number of CoA messages being sent per Deploy and ensures continued end-user experience, ensuring that connectivity remains stable and seamless while the security policies evolve.

Policy Download Mechanism

At this point, we've discussed classification, how it takes place, and how a network device learns of the IP TO SGT bindings. What happens immediately after this occurs? If a network device is configured to be an SGT enforcement point, then as soon as it learns of an IP TO SGT binding classification that has been configured, learned dynamically, or learned via SXP, that network device will request the policies from ISE to protect that SGT as a destination. Figure 7-32 illustrates this concept. Remember that the TrustSec policy matrix is built with source SGTs down the left side of the matrix and destination SGTs across the top. When an employee connects to the network device and authenticates with ISE, the employee's SGT is assigned and sent to that network device via a RADIUS Access-Accept message. Immediately after a network device learns of this SGT to protect, it makes a request to download the policies destined toward this SGT (equating to the matrix column where that SGT is the destination). So, the switch that the employee is connected to downloads the employee's SGT column

(that is, employees-to-employees permit and SCADA-to-employees deny). Other than the default policy, these are the only policies downloaded to that access switch in relation to the visible matrix in the figure. Similarly, when the camera connects to a network device and authenticates with ISE, the camera SGT is assigned and sent to that network device via RADIUS authorization. Immediately after that network device learns of this SGT to protect, it makes a request to download the policies destined toward this SGT (equating to the matrix column where the camera SGT is the destination). So, the switch that the camera is connected to downloads the camera SGT column (that is, camera-to-camera custom SGACL and employees-to-camera deny). Again, other than the default policy, these are the only policies downloaded to that access switch in relation to the visible matrix in Figure 7-30.

Technet24

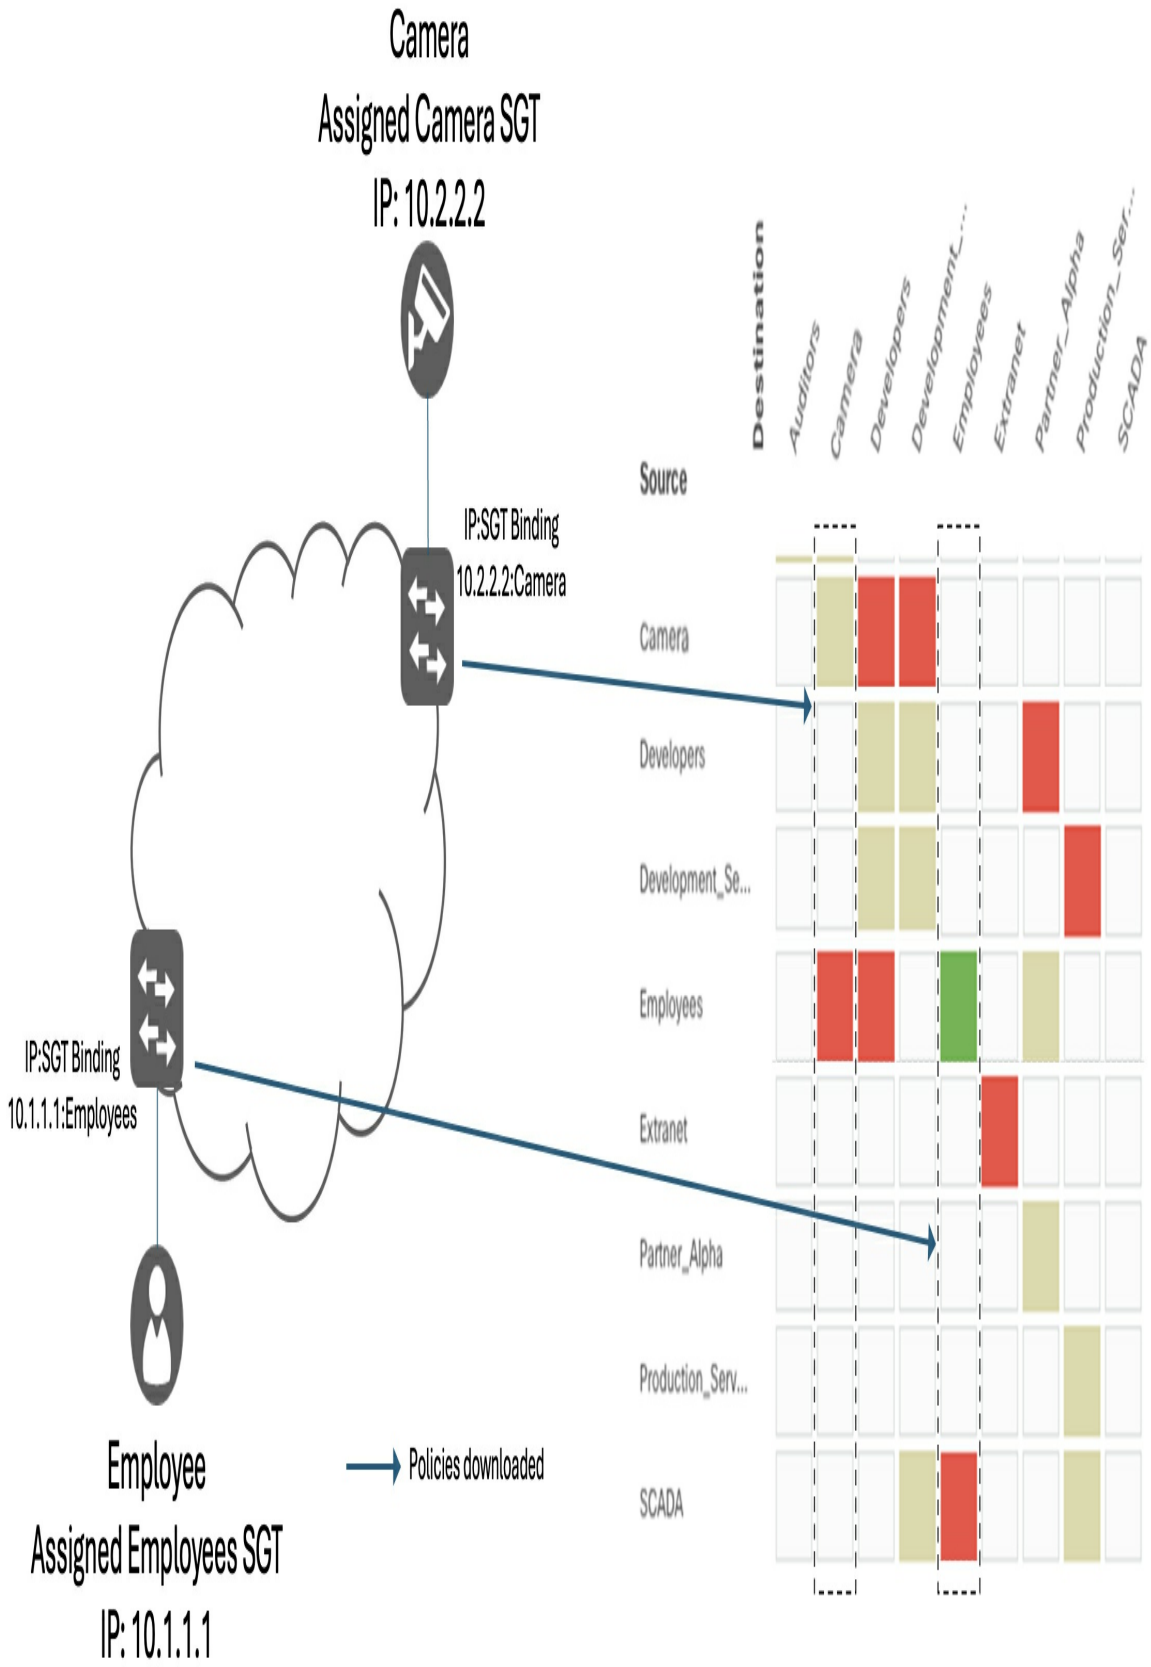

Camera
Assigned Camera SGT
IP: 10.2.2.2
IP:SGT Binding
10.2.2.2:Camera
IP:SGT Binding
10.1.1.1:Employees
Employee
Assigned Employees SGT
IP: 10.1.1.1
Policies downloaded
Destination
Auditors
Camera
Developers
Development_...
Employees
Extranet
Partner_Alpha
Production_Ser...
SCADA
Source
Camera
Developers
Development_Se...
Employees
Extranet
Partner_Alpha
Production_Serv...
SCADA

Figure 7-30 *Policy Download Concept*

This concept helps us to understand that network devices only need to download policies for devices directly attached or for IP TO SGT bindings they are aware of. This is one reason why this technology scales. Even if the TrustSec matrix is very large with many SGTs, only a very small subset of policies is downloaded to each network-enforcing device (depending on the bindings the device has learned).

Now we can understand why the technology is said to enforce at network egress. The closer enforcement gets to the destination, the less policies need to be downloaded and installed by the enforcement device. Additionally, that destination policy is in place as soon as endpoints connect to the network. Thus, this mechanism is dynamic and scalable.

Enforcement Logging

Visibility is key when enforcing traffic. Statistics, syslog, and NetFlow are all useful for monitoring operations.

From a statistics point of view, enforcing devices display hit counts for enforcement. A sample command is shown in Example 7-8.

Example 7-8 *Policy Hit Counts on Network Devices*

```
switch#show cts role-based counters
Role-based IPv4 counters
From    To      SW-Denied  HW-Denied  SW-Permitt HW-Permitt SW-Mo
*       *       0          0          0          11367      0
28      4       0          218        0          0          0
6       16      0          0          0          2367       0
10      16      0          0          0          7853       0
14      16      0          10922      0          0          0
22      1001    0          0          0          62         0
```

Technet24

This capture shows actual permits and denies between security groups but also includes columns for monitor mode to show whether enforcement would be applied if switched out of monitor mode.

Syslog is another useful tool. Logging messages are generated when the **log** keyword is appended to a security group access control entry (SGACE).

For example,

a) if the following SGACL exists:

```
Kernow_Permit_IP
permit udp src eq 53 log
```

b) the logging level is set to informational level 6 (the default level):

```
logging level cts 6
```

c) the syslog messages of that level are enabled on the console or syslog server:

```
logging console informational
logging trap informational
```

Then a syslog message will be generated by the network device, which could look like the following if a UDP packet using source port 53 is detected from source group 1002 to destination group 17000:

```
Feb 5 08:15:45.309 GMT: %RBM-6-SGACLHIT: ingress_interface='Gigab
sgacl_name='Kernow_Permit_IP-10' action='Permit' protocol='udp' s
ip='185.136.228.13' src-port='53' dest-ip='10.72.1.2' dest-port='
dgt='17000' logging_interval_hits='5'
```

Additionally, NetFlow can export flows, including the source and destination SGT. You can monitor them using the CLI command **show flow monitor cache**.

Syslog and NetFlow messages can be sent to any security information and event management (SIEM) system, but a recommendation is to use Cisco Secure Network Analytics or Splunk to provide a centralized view of security events, allowing security teams to investigate and identify anomalies.

From IOS-XE release 17.13.1, the export of FirewallEvent (233) is also supported on Doppler ASIC platforms (Catalyst 9200, 9300, 9400, 9500, and 9600), where Event 1 donates a permit action and Event 3 donates a deny action. Cisco Secure Network Analytics release 7.4.2 and later support the receiving of these events and is very useful for security policy compliance. Figure 7-31 shows an event in Secure Network Analytics displaying a denied flow action.

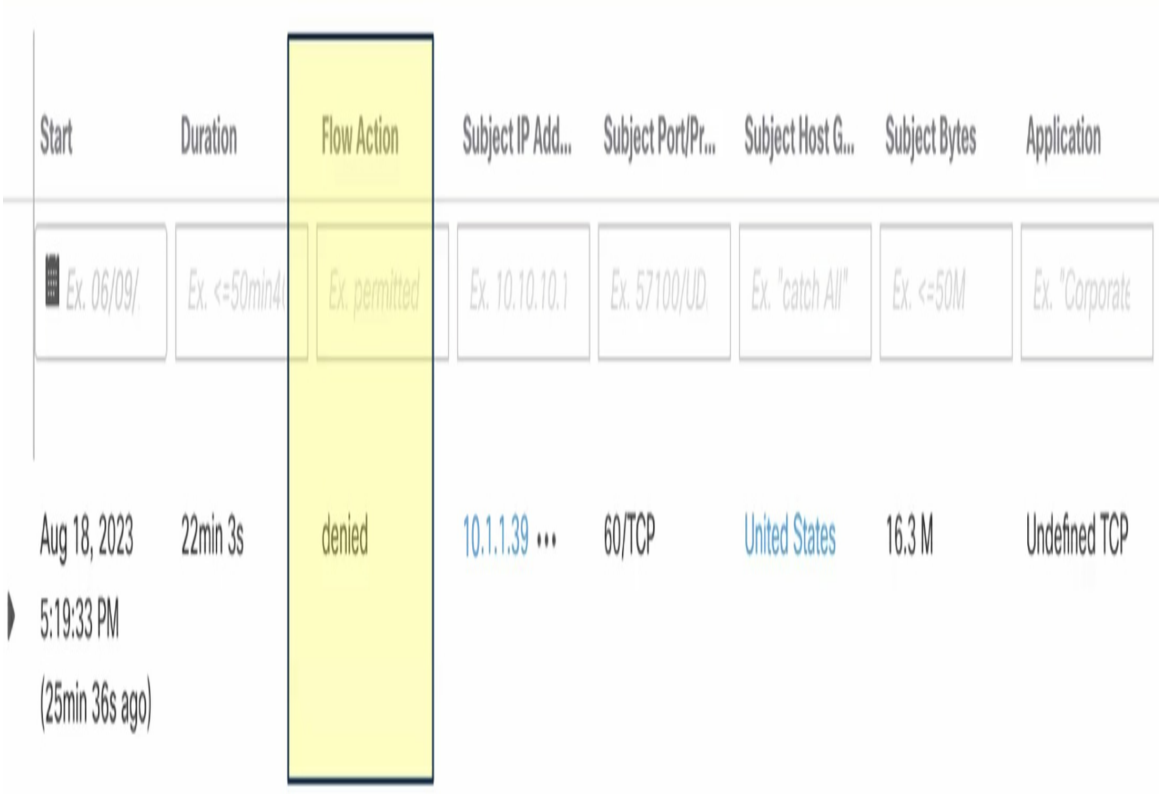


Figure 7-31 *Secure Network Analytics Displaying Flow Action of Denied from NetFlow*

Technet24

Enforcement Summary

Enforcement on a platform is contingent upon three specific conditions: Enforcement must be explicitly enabled, the source and destination IP TO SGT bindings must be known or learned, and a valid policy must be present. To maintain scalability, this enforcement is typically executed at the network egress point, because this approach ensures optimal performance and efficient resource utilization while maintaining security at the perimeter of the destination.

The onboarding process for TrustSec has evolved significantly, particularly with the introduction of PAC-less mode in ISE release 3.4, which is now considered the best practice over the traditional protected access credential used with EAP-FAST. During this onboarding phase, the network device downloads essential environment data from ISE, including the server list, the local network device SGT, the complete list of SGTs, and various operational timers. This download ensures the device has the necessary context to participate in the TrustSec solution effectively.

Policy management within ISE offers flexibility through the use of multiple TrustSec matrices and multiple SGACLs per policy, though it is important to note that these features are not supported when ISE is integrated with Catalyst Center. To assist with implementation, an SGACL monitor mode is available, providing valuable visibility into traffic patterns before you commit to any restrictive deny actions. Furthermore, policy downloads are handled dynamically, meaning they are pushed to the device only when necessary to protect known SGTs and are removed when they are no longer required, which optimizes resource utilization.

To maintain a healthy and secure environment, robust logging mechanisms are essential. You can leverage counters, syslog, and

NetFlow to monitor the TrustSec environment, all of which serve as critical components for ongoing network operations and troubleshooting. These tools provide the necessary insight to verify that policies are functioning as intended and allow for proactive adjustments to the security posture.

Summary

Cisco TrustSec represents a fundamental shift in network security by moving away from traditional IP-based access control toward a software-defined segmentation model. By decoupling security policy from the underlying network topology, TrustSec allows organizations to manage access based on endpoint identity and context rather than static IP addresses. This architecture simplifies policy management, reduces the complexity of network redesigns, and limits the impact of potential breaches through a "least-privilege" approach. The technology is built on three functional stages: classification, propagation, and enforcement.

Classification is the foundational process of assigning a security group tag (SGT) to a user, endpoint, or service as it enters the network. This process abstracts the user or device into a logical group, such as Employees, Guests, or Confidential Servers. Dynamic classification is the most common method, where Cisco ISE assigns an SGT during the RADIUS authorization phase. The network device then tracks the resulting IP-to-SGT binding and updates ISE using accounting messages. For devices that do not support dynamic authentication, such as servers, static classification methods allow administrators to map SGTs to specific IP addresses, subnets, VLANs, or physical ports. Because multiple classification methods may exist on a single device, Cisco IOS-XE follows a strict order of precedence to determine which tag to apply, with inline tagging typically taking the highest priority.

Technet24

Traffic that remains unclassified is assigned a default Unknown tag (SGT 0), allowing for a phased deployment where security policies can be introduced gradually.

Once traffic is classified, the source SGT must be communicated across the network to the point where the policy will be applied. This process, known as propagation, is essential for scalability. By carrying the source classification with every packet, enforcement points only need to maintain the destination bindings for their locally attached devices rather than a global table of every endpoint on the network. The most efficient method of propagation is inline tagging, which inserts the SGT directly into the Cisco Meta Data field of the Layer 2 frame. This allows the tag to be carried at line rate in hardware. In environments where inline tagging is not supported, the Security Group Tag Exchange Protocol (SXP) provides a control-plane alternative. SXP uses a TCP-based peer-to-peer connection to share IP-to-SGT bindings between devices. Modern iterations of SXP, such as version 4, include critical loop detection and mitigation features, while version 5 allows for the propagation of VRF and VLAN information across a single connection. Additionally, ISE can propagate identity context to other security platforms, such as firewalls and analytics engines, via the pxGrid framework.

Enforcement is the final stage where the network applies security policies to permit or deny traffic. This typically occurs at the network egress — the device closest to the destination — to ensure the system scales effectively. Before a device can enforce policy, it must undergo an onboarding process to download environment data from ISE, which includes the global list of SGT names and numbers, server lists, and relevant timers. This communication traditionally uses a protected access credential (PAC) via EAP-FAST, though newer versions of ISE and IOS-XE support a PAC-less mode for a simplified onboarding

process. The actual security rules are managed through a centralized TrustSec policy matrix in ISE, where administrators define permissions between source and destination SGTs using security group ACLs (SGACLs). These SGACLs specify which protocols and ports are allowed. To ensure a smooth deployment, TrustSec supports a monitor mode that allows administrators to view hit counts and potential denials without actually blocking traffic. Comprehensive logging through Syslog and NetFlow provides deep visibility into these enforcement actions, allowing security teams to identify anomalies and ensure compliance across the entire infrastructure.

Technet24

Chapter 8. Device Identification for Access Management

Introduction

You cannot protect what you cannot see. You must understand what connects to your network and where these connections occur. You may recognize most corporate devices through asset management databases, but manually entered data can be unreliable or compromised. Your organization needs a mechanism that enables you, as administrator, to confirm, with certainty, that a Windows 11 workstation connects behind the switch on the second floor in the third rack or that a Cisco 8851 IP phone occupies the same port. Knowing only the MAC address or IP address of the endpoint does not provide sufficient information.

Cisco Identity Services Engine profiling allows you to dynamically assign specific profiles to endpoints based on received attributes. ISE collects endpoint data directly from endpoints, network access devices (NADs), and external databases to classify endpoints.

In addition to network visibility, profiling enhances security. You do not rely solely on endpoint authentication and authorization; you also ensure that only specific types of devices gain access to the network.

Context Visibility

Context Visibility powers ISE profiling. You can access it from the **Context Visibility > Endpoints** menu, as shown in Figure 8-1.

Technet24

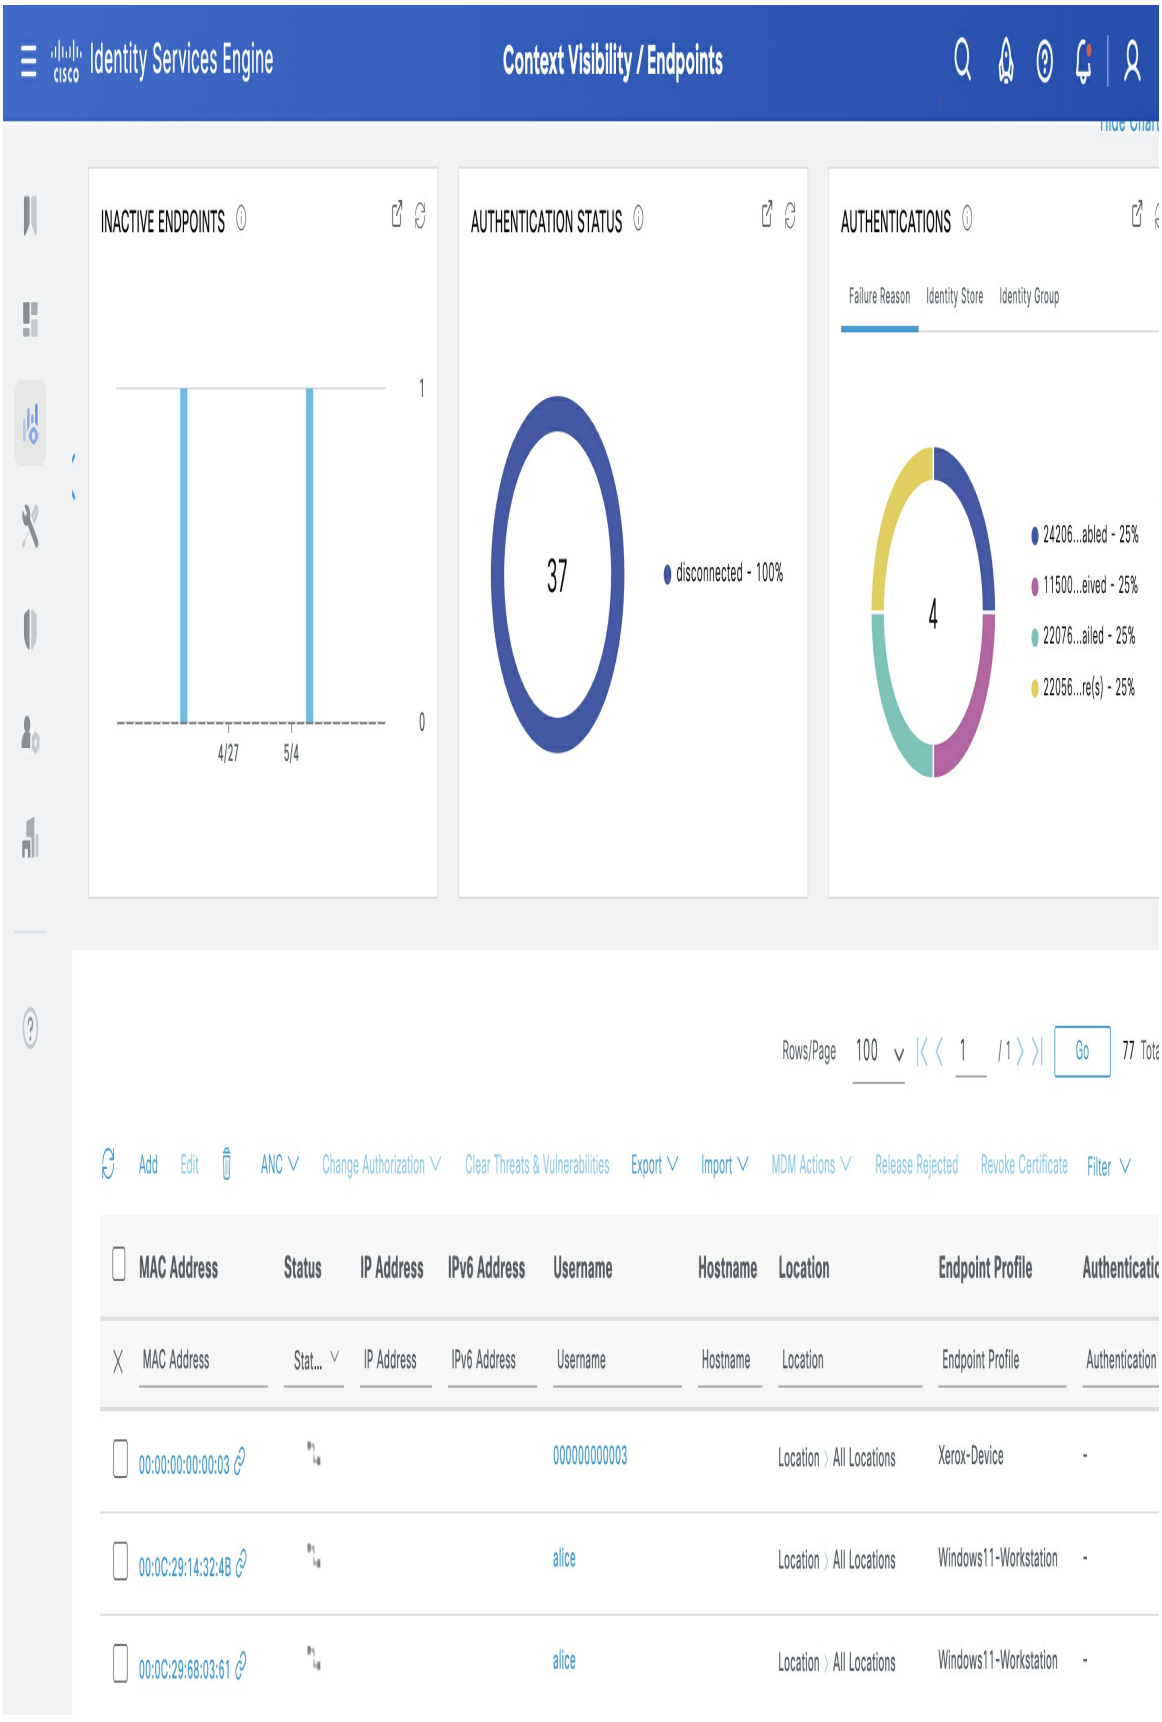

Identity Services Engine
Context Visibility / Endpoints
INACTIVE ENDPOINTS
1
0
4/27
5/4
AUTHENTICATION STATUS
37
disconnected - 100%
AUTHENTICATIONS
Failure Reason
Identity Store
Identity Group
4
24206...abled - 25%
11500...aived - 25%
22076...ailed - 25%
22056...re(s) - 25%
Rows/Page 100
1 /1
Go
77 Tota
Add
Edit
ANC
Change Authorization
Clear Threats & Vulnerabilities
Export
Import
MDM Actions
Release Rejected
Revoke Certificate
Filter
MAC Address
Status
IP Address
IPv6 Address
Username
Hostname
Location
Endpoint Profile
Authenticatio
00:00:00:00:00:03
000000000003
Location > All Locations
Xerox-Device
00:0C:29:14:32:4B
alice
Location > All Locations
Windows11-Workstation
00:0C:29:68:03:61
alice
Location > All Locations
Windows11-Workstation

Figure 8-1 *Context Visibility*

Context Visibility provides access to all endpoints ever seen by ISE, regardless of whether authentication was successful or unsuccessful. Context Visibility serves as the primary tool for ISE administrators to examine the specifics of a particular endpoint on ISE. You can easily filter endpoints based on MAC address, IP address, username, or many other parameters.

You can also view which endpoints are currently connected, which profiles are assigned to specific endpoints, and the associations between IP addresses and MAC addresses.

By clicking an endpoint's MAC address, you gain quick access to the simplified **Endpoint Attributes** screen, as shown in Figure 8-2. From this screen, you can access **General Attributes**, **MFC Attributes**, **Endpoint Profile** information, and **IP Addresses**.

Technet24

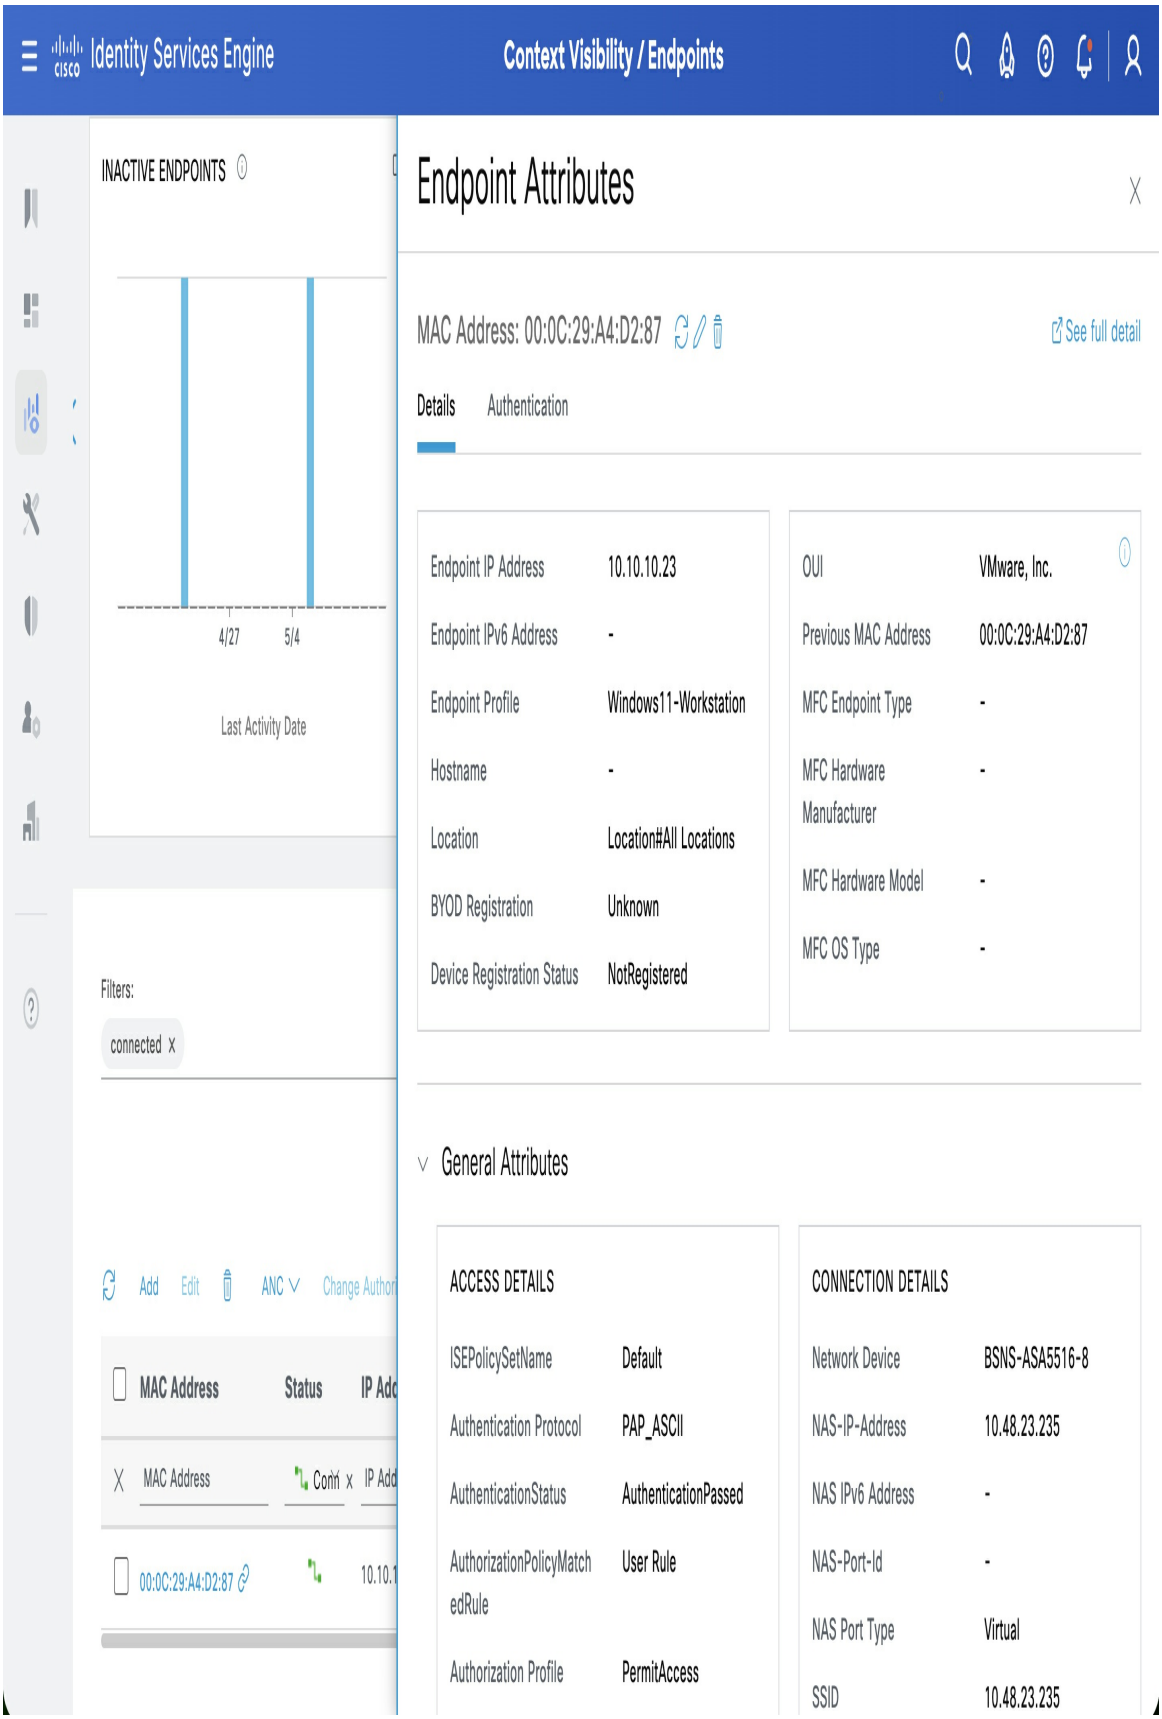

Identity Services Engine
Context Visibility / Endpoints
INACTIVE ENDPOINTS
4/27
5/4
Last Activity Date
Filters:
connected x
Add
Edit
ANC
MAC Address
Status
MAC Address
00:0C:29:A4:D2:87
Endpoint Attributes
MAC Address: 00:0C:29:A4:D2:87
See full detail
Details
Authentication
Endpoint IP Address 10.10.10.23
Endpoint IPv6 Address -
Endpoint Profile Windows11-Workstation
Hostname -
Location Location#All Locations
BYOD Registration Unknown
Device Registration Status NotRegistered
OUI VMware, Inc.
Previous MAC Address 00:0C:29:A4:D2:87
MFC Endpoint Type -
MFC Hardware Manufacturer -
MFC Hardware Model -
MFC OS Type -
General Attributes
ACCESS DETAILS
ISEPolicySetName Default
Authentication Protocol PAP_ASCII
AuthenticationStatus AuthenticationPassed
AuthorizationPolicyMatchedRule User Rule
Authorization Profile PermitAccess
CONNECTION DETAILS
Network Device BSNS-ASA5516-8
NAS-IP-Address 10.48.23.235
NAS IPv6 Address -
NAS-Port-Id -
NAS Port Type Virtual
SSID 10.48.23.235

Figure 8-2 *Context Visibility Endpoint Attributes*

You can quickly change the **Policy Assignment** or configure **Static Group Assignments** for the endpoint by clicking the **Edit** button, as shown in Figure 8-3.

Technet24

Edit Endpoint

General Attributes

Mac Address*

00:0C:29:A4:D2:87

Description

Static Assignment

Static Group Assignment

Policy Assignment

Windows11-Workstation

Identity Group Assignment

Windows11-Workstation

Custom Attributes

Other Attributes

Cancel

Refresh

Save

Figure 8-3 *Context Visibility Edit Endpoint Menu*

You can obtain all information known to ISE about a specific endpoint by clicking **See Full Detail**, as shown in Figure 8-4.

Technet24

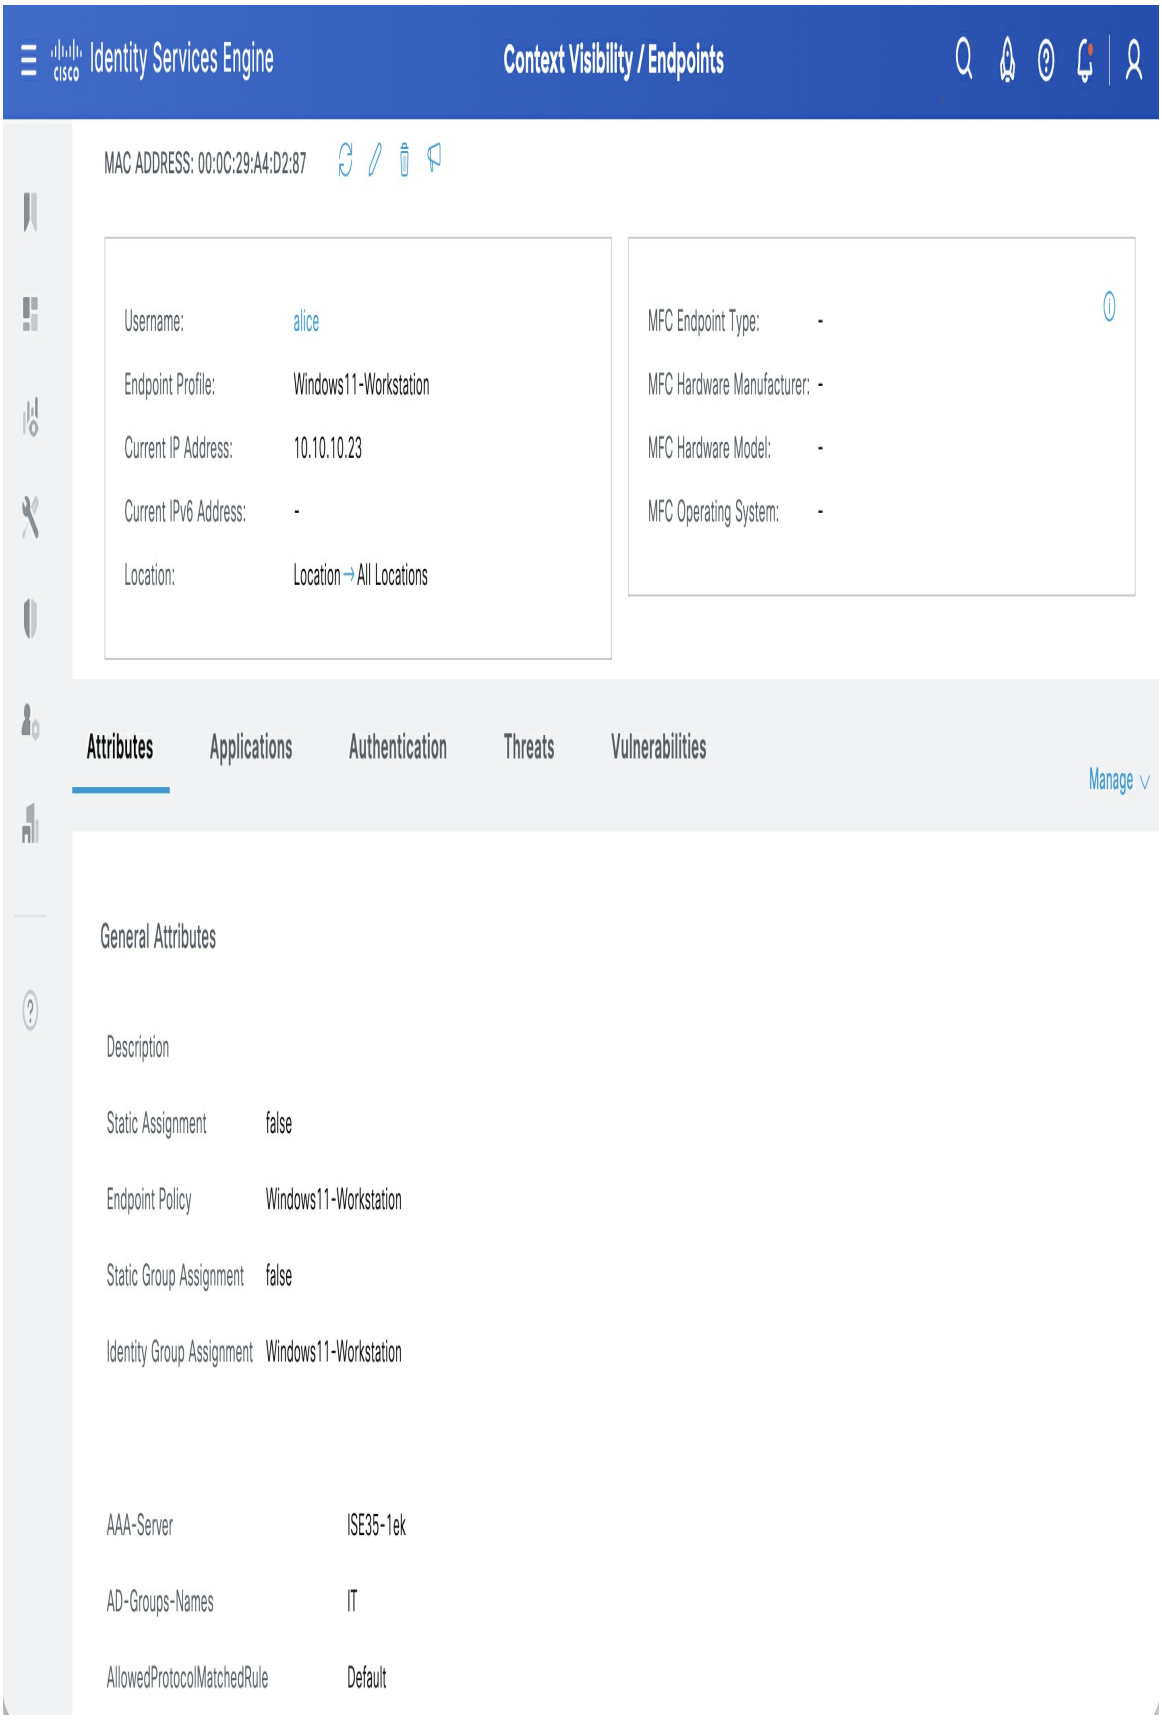

Identity Services Engine
Context Visibility / Endpoints
MAC ADDRESS: 00:0C:29:A4:D2:87
Username: alice
Endpoint Profile: Windows11-Workstation
Current IP Address: 10.10.10.23
Current IPv6 Address: -
Location: Location → All Locations
MFC Endpoint Type: -
MFC Hardware Manufacturer: -
MFC Hardware Model: -
MFC Operating System: -
Attributes
Applications
Authentication
Threats
Vulnerabilities
Manage
General Attributes
Description
Static Assignment false
Endpoint Policy Windows11-Workstation
Static Group Assignment false
Identity Group Assignment Windows11-Workstation
AAA-Server ISE35-1ek
AD-Groups-Names IT
AllowedProtocolMatchedRule Default

Figure 8-4 *Context Visibility Full Detail Endpoint Menu*

You can navigate to five different menu options to obtain information about their attributes. Figure 8-5 shows a small portion of the endpoint attributes learned by ISE during authentication of the endpoint.

Technet24

AAA-Server	ISE35-1ek
AD-Groups-Names	IT
AllowedProtocolMatchedRule	Default
AuthenticationIdentityStore	EntraID
AuthenticationMethod	PAP_ASCII
AuthenticationStatus	AuthenticationPassed
AuthorizationPolicyMatchedRule	User Rule
BYODRegistration	Unknown
CVPN3000/ASA/PIX7x-Client-Type	2
CVPN3000/ASA/PIX7x-Session-Subtype	3
CVPN3000/ASA/PIX7x-Session-Type	1
CVPN3000/ASA/PIX7x-Tunnel-Group-Name	RA
Called-Station-ID	10.48.23.235
Calling-Station-ID	10.26.161.171
ClientLatency	0
DTLSSupport	Unknown
DestinationIPAddress	10.48.17.88
DestinationPort	1645
Device IP Address	10.48.23.235
Device Identifier	63B68BA3FC66F0A3607DD1B930D55C0DC008D371

Figure 8-5 *Context Visibility Edit Endpoint Menu*

Additionally, from Context Visibility, you can view which applications are installed and running on the endpoint, discovered through ISE Posture. You can also review the details of the last authentication attempt and see the threats and vulnerabilities detected on the endpoint via TC-NAC integrations.

You can populate the endpoint database on ISE in three ways:

- Dynamically, by allowing ISE to learn endpoints through profiling. This method is the most scalable because ISE performs it automatically.
- Manually, by adding the endpoint via Context Visibility or by importing endpoints from a CSV file or an LDAP server.
- Programmatically, by using OpenAPI or ERS API.

Certainty Factor–Based Profiling

You can profile your endpoints in two different ways. The first method uses probes and relies on the certainty factor, which is the original approach to profiling endpoints. The second method relies on cloud services and manual policies to profile endpoints and is known as cloud profiling.

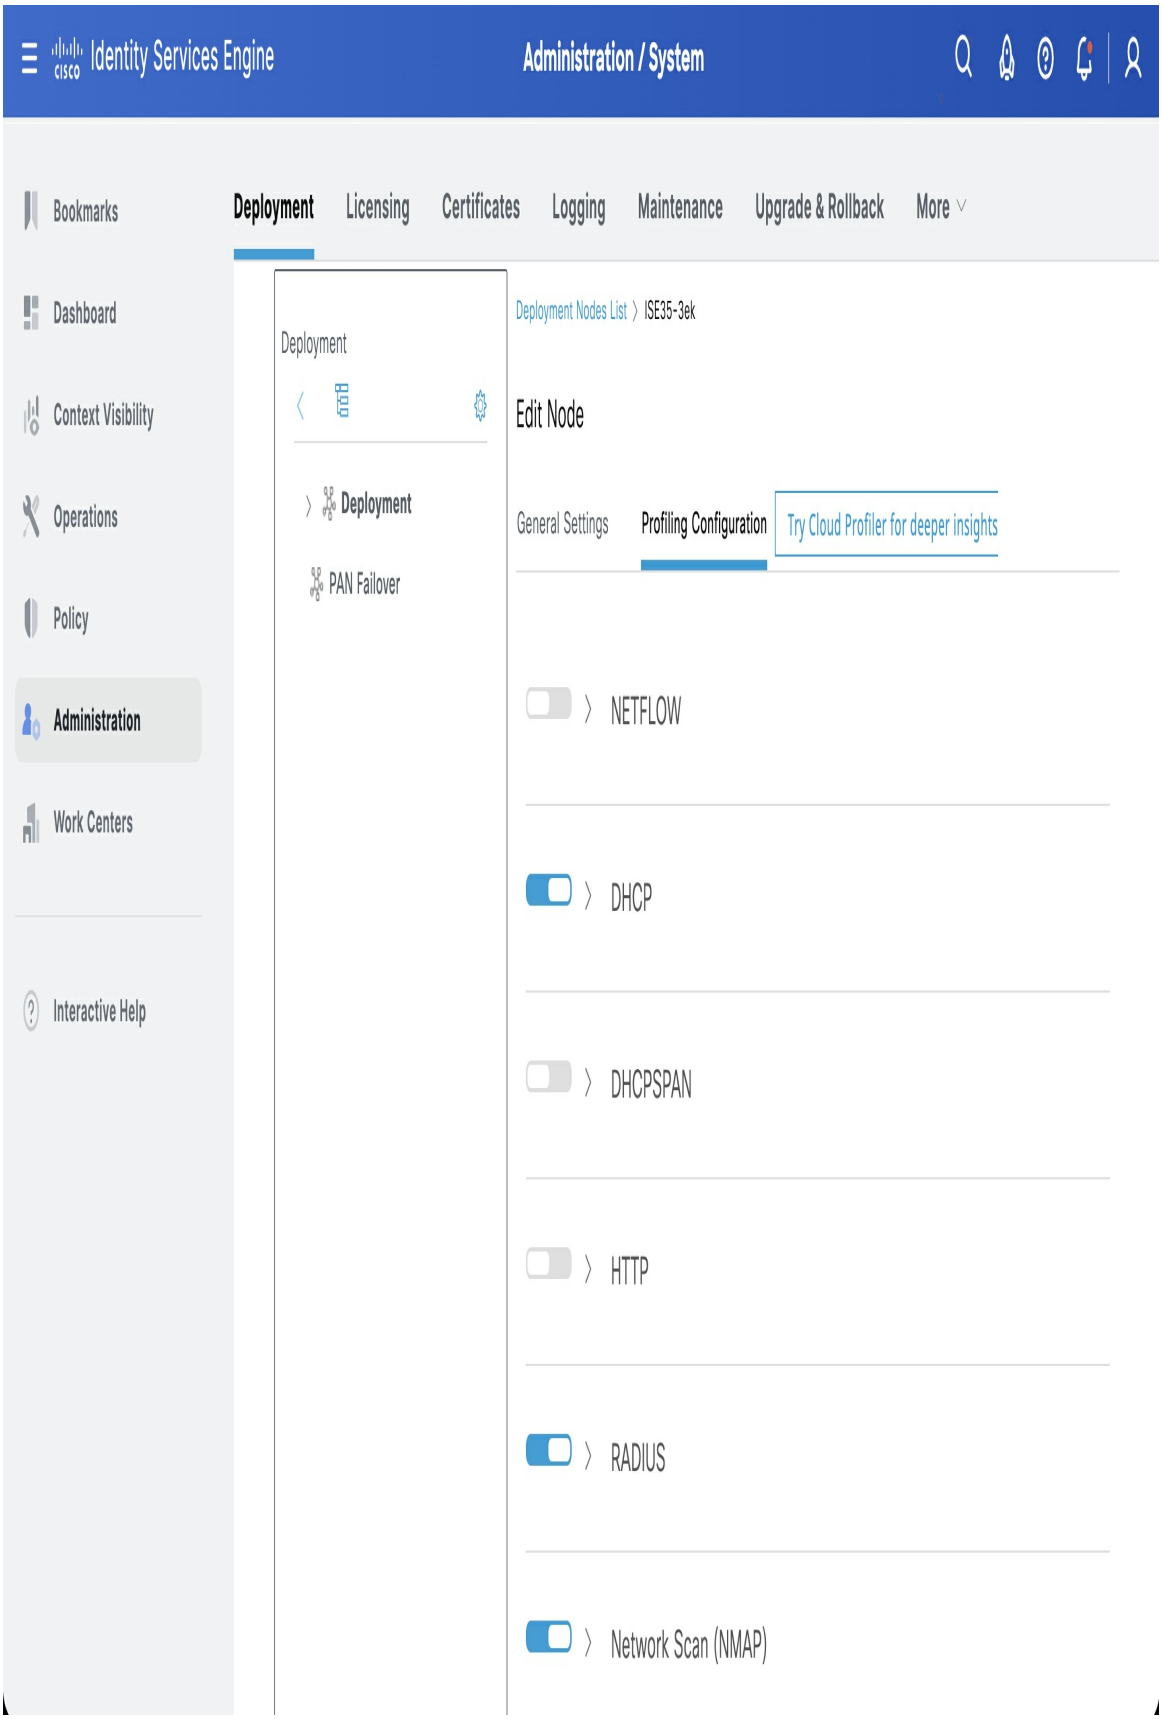
First, let's examine certainty factor–based profiling to understand how it works.

Probe Configuration

Certainty factor–based profiling relies on probe information. Probes consist of protocols and services that ISE uses to learn more about the endpoint. You configure probes under **Administration > System >**

Technet24

Deployment > (Selected PSN) > Profiling Configuration, as shown in Figure 8-6.


Identity Services Engine
Administration / System
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
More
Deployment
Deployment
PAN Failover
Deployment Nodes List
ISE35-3ek
Edit Node
General Settings
Profiling Configuration
Try Cloud Profiler for deeper insights
NETFLOW
DHCP
DHCPSPAN
HTTP
RADIUS
Network Scan (NMAP)

Technet24

Figure 8-6 *ISE Probe Configuration*

There are eleven probes in total:

- NetFlow
- DHCP
- DHCPSPAN
- HTTP
- RADIUS
- NMAP
- DNS
- SNMPQUERY
- SNMPTRAP
- Active Directory
- pxGrid

You configure probes for each Policy Service Node (PSN). Now let's examine each probe in detail.

NetFlow Probe

NetFlow is a protocol developed by Cisco to collect data about IP traffic traversing network devices. Devices cache metadata and send it to NetFlow collectors. ISE can function as a collector. Examples of metadata include

- Source IP address
- Destination IP address
- Source port

- Destination port
- Protocol

You can access the full list of attributes supported by ISE via the NetFlow dictionary from **Policy > Policy Elements > Dictionaries > System > NetFlow > Dictionary Attributes**, as shown in Figure 8-7.

Technet24

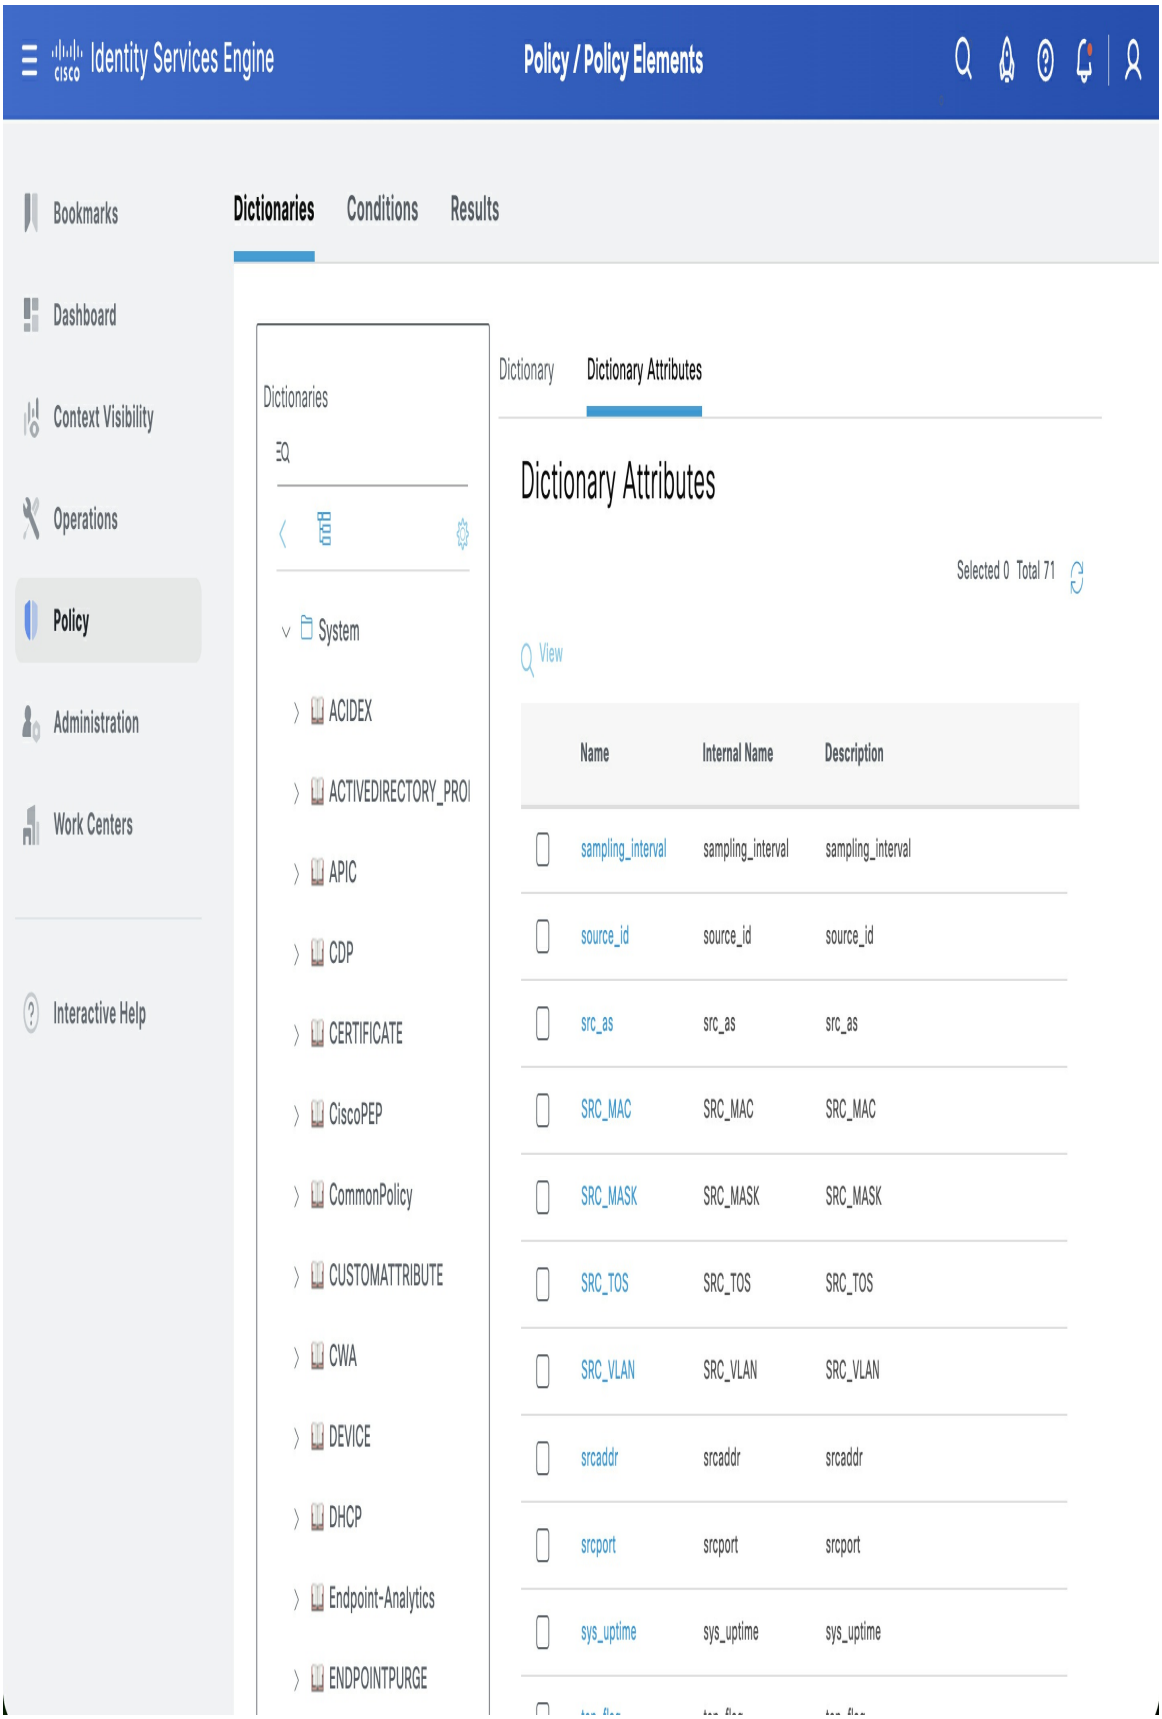

Identity Services Engine
Policy / Policy Elements
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Dictionaries
Conditions
Results
Dictionaries
System
ACIDEX
ACTIVEDIRECTORY_PROI
APIC
CDP
CERTIFICATE
CiscoPEP
CommonPolicy
CUSTOMATTRIBUTE
CWA
DEVICE
DHCP
Endpoint-Analytics
ENDPOINTPURGE
Dictionary
Dictionary Attributes
Dictionary Attributes
Selected 0 Total 71
View
Name
Internal Name
Description
sampling_interval
sampling_interval
sampling_interval
source_id
source_id
source_id
src_as
src_as
src_as
SRC_MAC
SRC_MAC
SRC_MAC
SRC_MASK
SRC_MASK
SRC_MASK
SRC_TOS
SRC_TOS
SRC_TOS
SRC_VLAN
SRC_VLAN
SRC_VLAN
srcaddr
srcaddr
srcaddr
srcport
srcport
srcport
sys_uptime
sys_uptime
sys_uptime

Figure 8-7 *NetFlow Dictionary Attributes*

ISE can use NetFlow to profile endpoints based on traffic patterns. This capability is well suited for endpoints that generate unique flows and use distinctive protocols, such as Internet of Things (IoT) devices.

To configure the NetFlow probe, enable the checkbox for NetFlow in the ISE Probe Configuration under the PSNs intended to receive NetFlow data. Ensure that network devices along the path of endpoint traffic flows are configured to export NetFlow data to ISE. ISE supports NetFlow versions 5 through 9, with version 9 recommended.

Best Practice:

Because NetFlow generates a significant amount of traffic, it is advisable to use alternative probes to uniquely identify endpoints. You may still use NetFlow when other probes cannot obtain the necessary data.

DHCP and DHCP SPAN Probes

When the DHCP or DHCP SPAN probe is enabled, ISE collects attributes from DHCP packets.

DHCP probes require network access devices to send copies of DHCP broadcast packets to a unicast address, which would be an ISE PSN. ISE does not assign an IP address to the endpoint; instead, it collects attributes and profiles the endpoint based on these attributes.

Typically, you need to configure the **ip helper-address** with the IP address of the PSN on the switch to forward DHCP packets to ISE.

In addition to IP and MAC addresses, DHCP probes provide valuable information that helps identify the operating system of endpoints. The most useful attributes are as follows:

Technet24

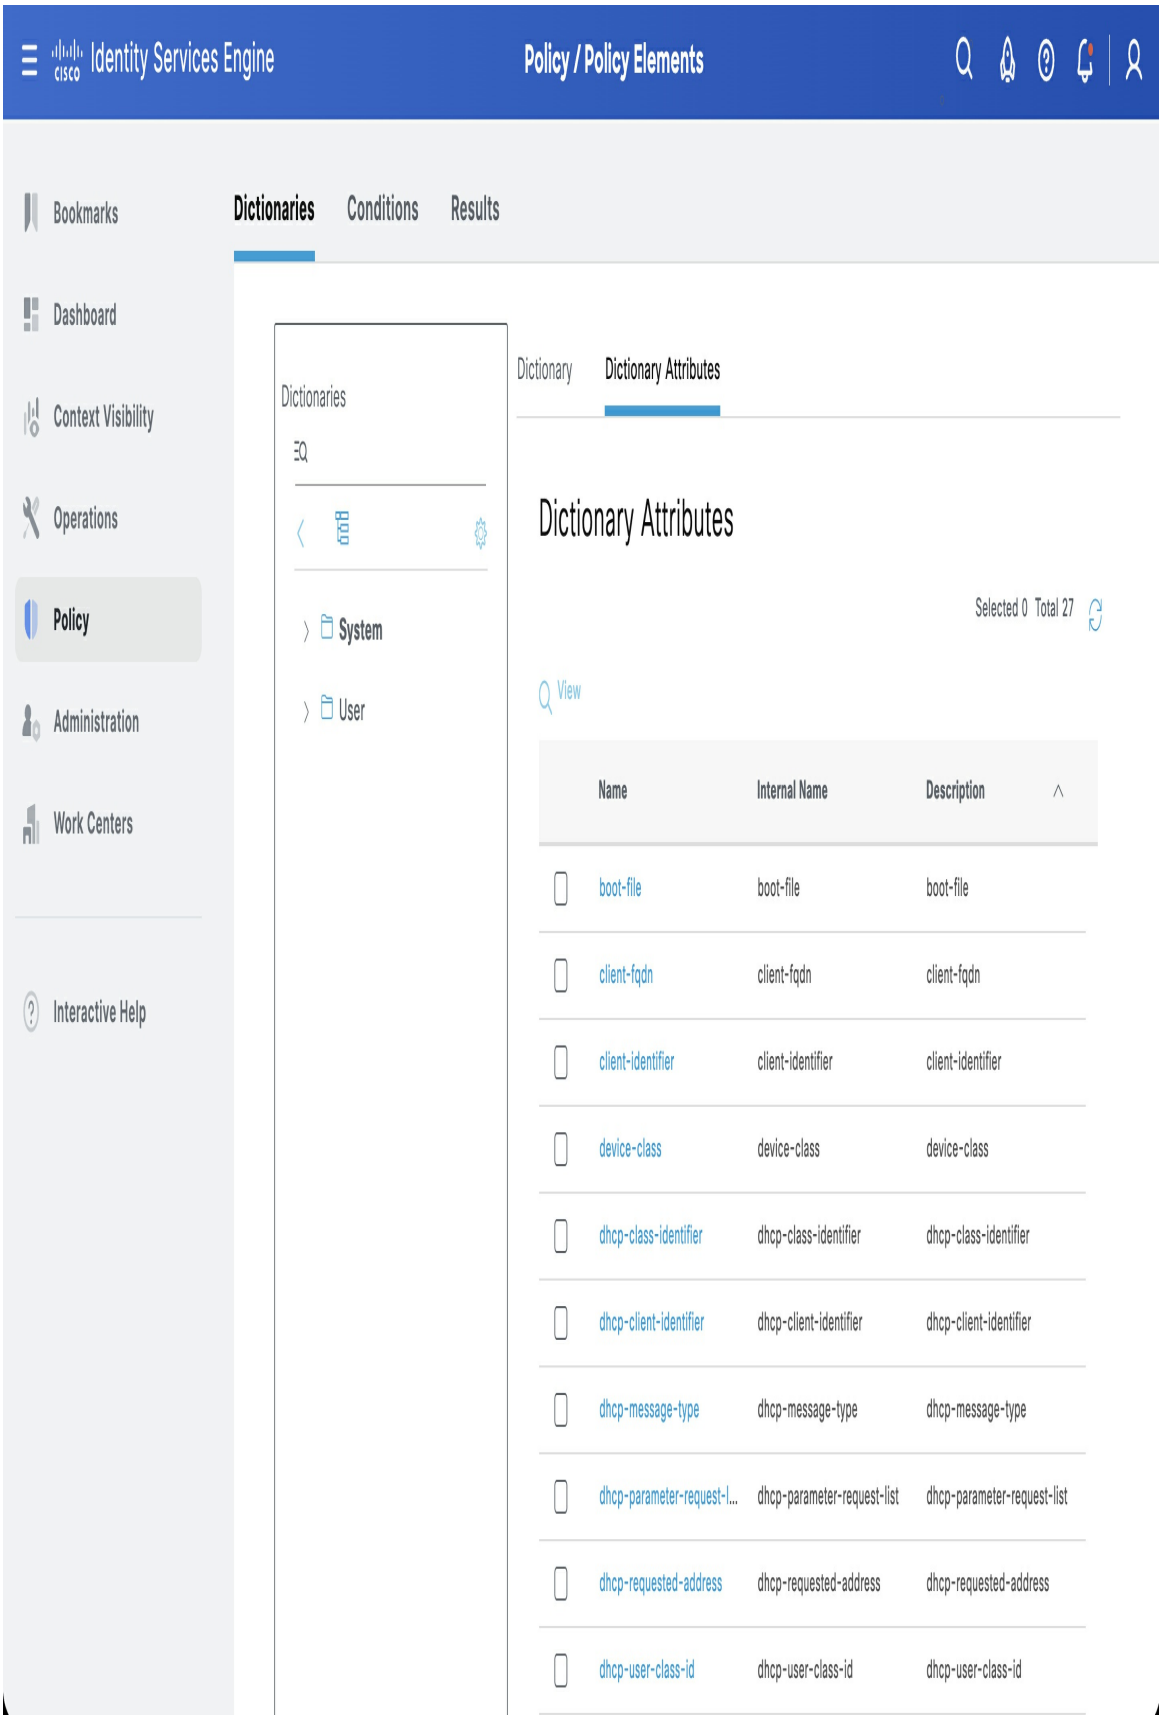
- Vendor class identifier
- Client identifier
- Hostname

Example 8-1 is a snippet from Wireshark showing attributes from a DHCP packet sent by a Windows virtual machine.

Example 8-1 *Attributes from a DHCP Packet Sent by a Windows Virtual Machine*

```
Option: (60) Vendor class identifier
    Length: 8
    Vendor class identifier: MSFT 5.0
Option: (12) Host Name
    Length: 14
    Host Name: ekorneyc-Win10
Option: (61) Client identifier
    Length: 7
    Hardware type: Ethernet (0x01)
    Client MAC address: VMware_a8:b1:b2 (00:50:56:a8:b1:b2)
```

You can access the full list of attributes supported by ISE via the DHCP dictionary from **Policy > Policy Elements > Dictionaries > System > DHCP > Dictionary Attributes**, as shown in Figure 8-8.


Identity Services Engine
Policy / Policy Elements
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Dictionaries
Conditions
Results
Dictionaries
System
User
Dictionary
Dictionary Attributes
Dictionary Attributes
Selected 0 Total 27
View
Name
Internal Name
Description
boot-file
boot-file
boot-file
client-fqdn
client-fqdn
client-fqdn
client-identifier
client-identifier
client-identifier
device-class
device-class
device-class
dhcp-class-identifier
dhcp-class-identifier
dhcp-class-identifier
dhcp-client-identifier
dhcp-client-identifier
dhcp-client-identifier
dhcp-message-type
dhcp-message-type
dhcp-message-type
dhcp-parameter-request-l...
dhcp-parameter-request-list
dhcp-parameter-request-list
dhcp-requested-address
dhcp-requested-address
dhcp-requested-address
dhcp-user-class-id
dhcp-user-class-id
dhcp-user-class-id

Technet24

Figure 8-8 *DHCP Dictionary Attributes*

DHCP SPAN relies on copies of traffic sent to ISE using SPAN functionality.

Best Practice:

Similar to NetFlow, DHCP SPAN can generate a significant amount of traffic, so it is recommended not to use it unless a specific use case exists. DHCP is a reliable probe that does not generate substantial traffic and is especially useful for identifying the underlying operating system of the endpoint. Device Sensor is a feature on network access devices that is superior to DHCP Probe alone and should be used instead, if possible, because it includes data from DHCP.

HTTP Probe

The HTTP probe allows ISE to capture data from the HTTP User-Agent field. This method is helpful with identifying the operating system of the endpoint.

The following snippet is from a Wireshark capture of an HTTP packet and the User-Agent field:

```
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15; rv:99.0) G
Firefox/99.0
```

The HTTP probe is used when HTTP traffic is sent directly to ISE, such as in a redirected flow where users access the ISE guest portal hosted on the PSNs through guest redirection.

Best Practice:

While HTTP data is useful for profiling, Device Sensor

supersedes this functionality. You should configure Device Sensor when possible and include HTTP data within it, rather than configuring direct HTTP profiling.

RADIUS Probe

The RADIUS probe is the simplest method for profiling endpoints. User and endpoint authentications include the Calling-Station-Id as a RADIUS attribute, which in most cases contains the MAC address of the endpoint. The MAC address consists of two parts: organizational unique identifier (OUI) and universally administered address (UAA). The OUI enables ISE-to-profile endpoints based on their organization, such as grouping endpoints into categories like Cisco Devices or VMware Devices.

The following snippet is from a Wireshark capture where the RADIUS packet contains the Calling-Station-Id.

```
AVP: t=Calling-Station-Id (31)
Type: 31
Length: 19
Calling-Station-Id: 24-7E-12-66-C1-98
```

Device Sensor

Although Device Sensor is not a dedicated probe, it is a crucial feature for profiling endpoints. Device Sensor enables switches and wireless LAN controllers to collect DHCP, LLDP, CDP, and HTTP data from endpoints. The devices cache this information and send it to ISE through RADIUS Accounting. It is important to note that RADIUS Accounting Start is sent only after successful authentication, which means that endpoints must first pass authentication to be profiled based on Device Sensor data. Initial authentication can assign very limited permissions. As long as RADIUS Access-Accept is returned,

Technet24

RADIUS Accounting Start will carry cached attributes from the network device to ISE.

The most powerful aspect of Device Sensor is its use of RADIUS Accounting as a native method of profiling, eliminating the need to introduce additional traffic to ISE. The combination of multiple attributes makes Device Sensor the preferred method for profiling endpoints.

> **Note**
>
> With Identity-Based Networking Services (IBNS) 2.0, you can send Device Sensor data even without authentication.

Device Sensor requires NAD configuration. Example 8-2 shows the configuration commands needed for Device Sensor to send cached CDP, LLDP, and DHCP data.

Example 8-2 *Device Sensor Configuration for Sending Cached CDP, LLDP, and DHCP Data*

```
cdp run
lldp run
ip dhcp snooping
!
device-sensor filter-list dhcp list dhcp-list
option name host-name
option name domain-name
option number 50
option name parameter-request-list
option name class-identifier
option name client-identifier
!
device-sensor filter-list lldp list lldp-list
tlv name system-description
!
device-sensor filter-list cdp list cdp-list
tlv name device-name
```

```
tlv name platform-type
device-sensor filter-spec dhcp include list dhcp-list
device-sensor filter-spec lldp include list lldp-list
device-sensor filter-spec cdp include list cdp-list
device-sensor accounting
device-sensor notify all-changes
```

> **Note**
>
> RADIUS Authentication and Accounting must be configured on the device.

SNMP Query and SNMP Trap Probes

The primary purpose of an SNMP trap is to trigger an SNMP query. With an SNMP query, you can obtain information such as

- cdpCacheEntry
- lldpLocalSystemData
- lldpRemoteSystemsData

There are two types of SNMP queries:

- Polled SNMP query
- Triggered SNMP query

A polled SNMP query occurs periodically based on configured intervals, with the default interval set to 28,800 seconds (8 hours). A triggered SNMP query is sent based on Accounting Start or an SNMP Trap.

To configure SNMP query or trap settings, navigate to **Administration > System > Deployment > (Selected PSN) > Profiling Configuration**, as shown in Figure 8-9.

Technet24

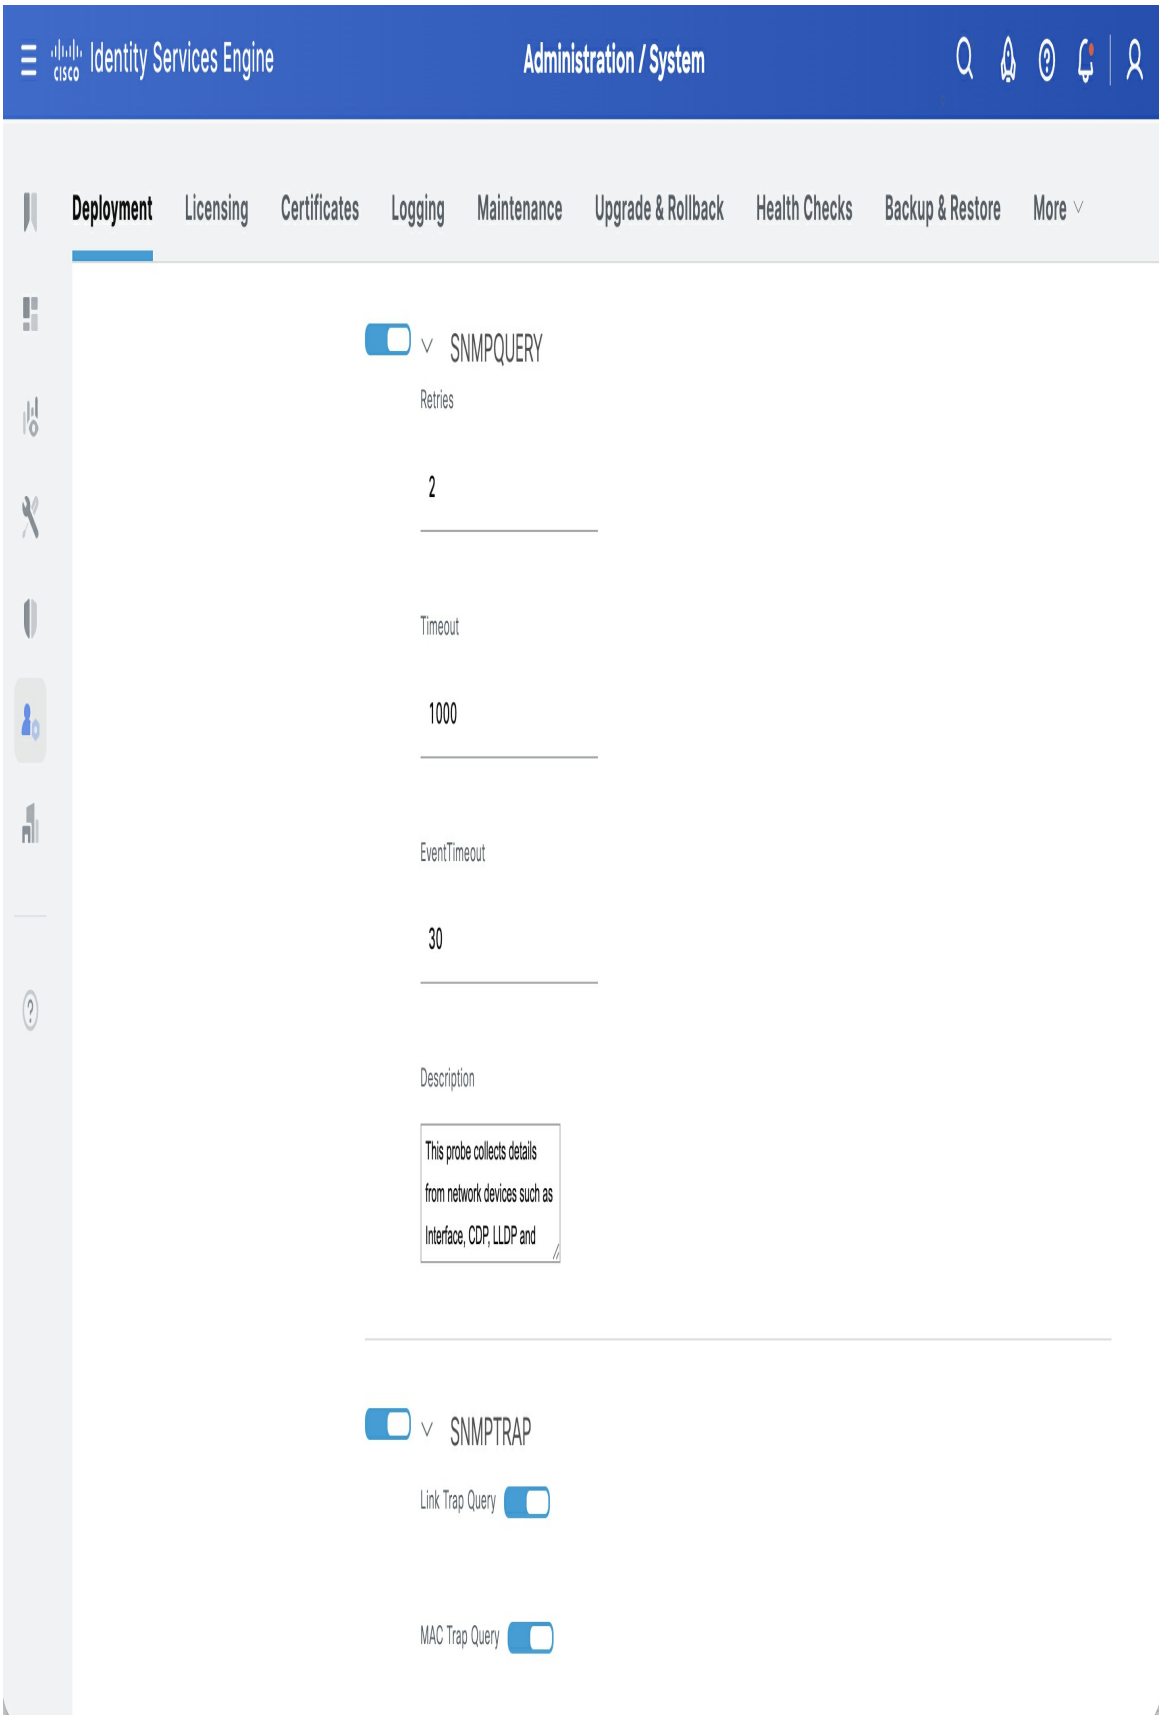

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
SNMPQUERY
Retries
2
Timeout
1000
EventTimeout
30
Description
This probe collects details from network devices such as Interface, CDP, LLDP and
SNMPTRAP
Link Trap Query
MAC Trap Query

Figure 8-9 *SNMP Query and SNMP Trap Settings*

Additionally, you need to configure SNMP-specific settings for each network device under **Administration > Network Resources > Network Devices > (Network Device) > SNMP Settings**, as shown in Figure 8-10.

Technet24

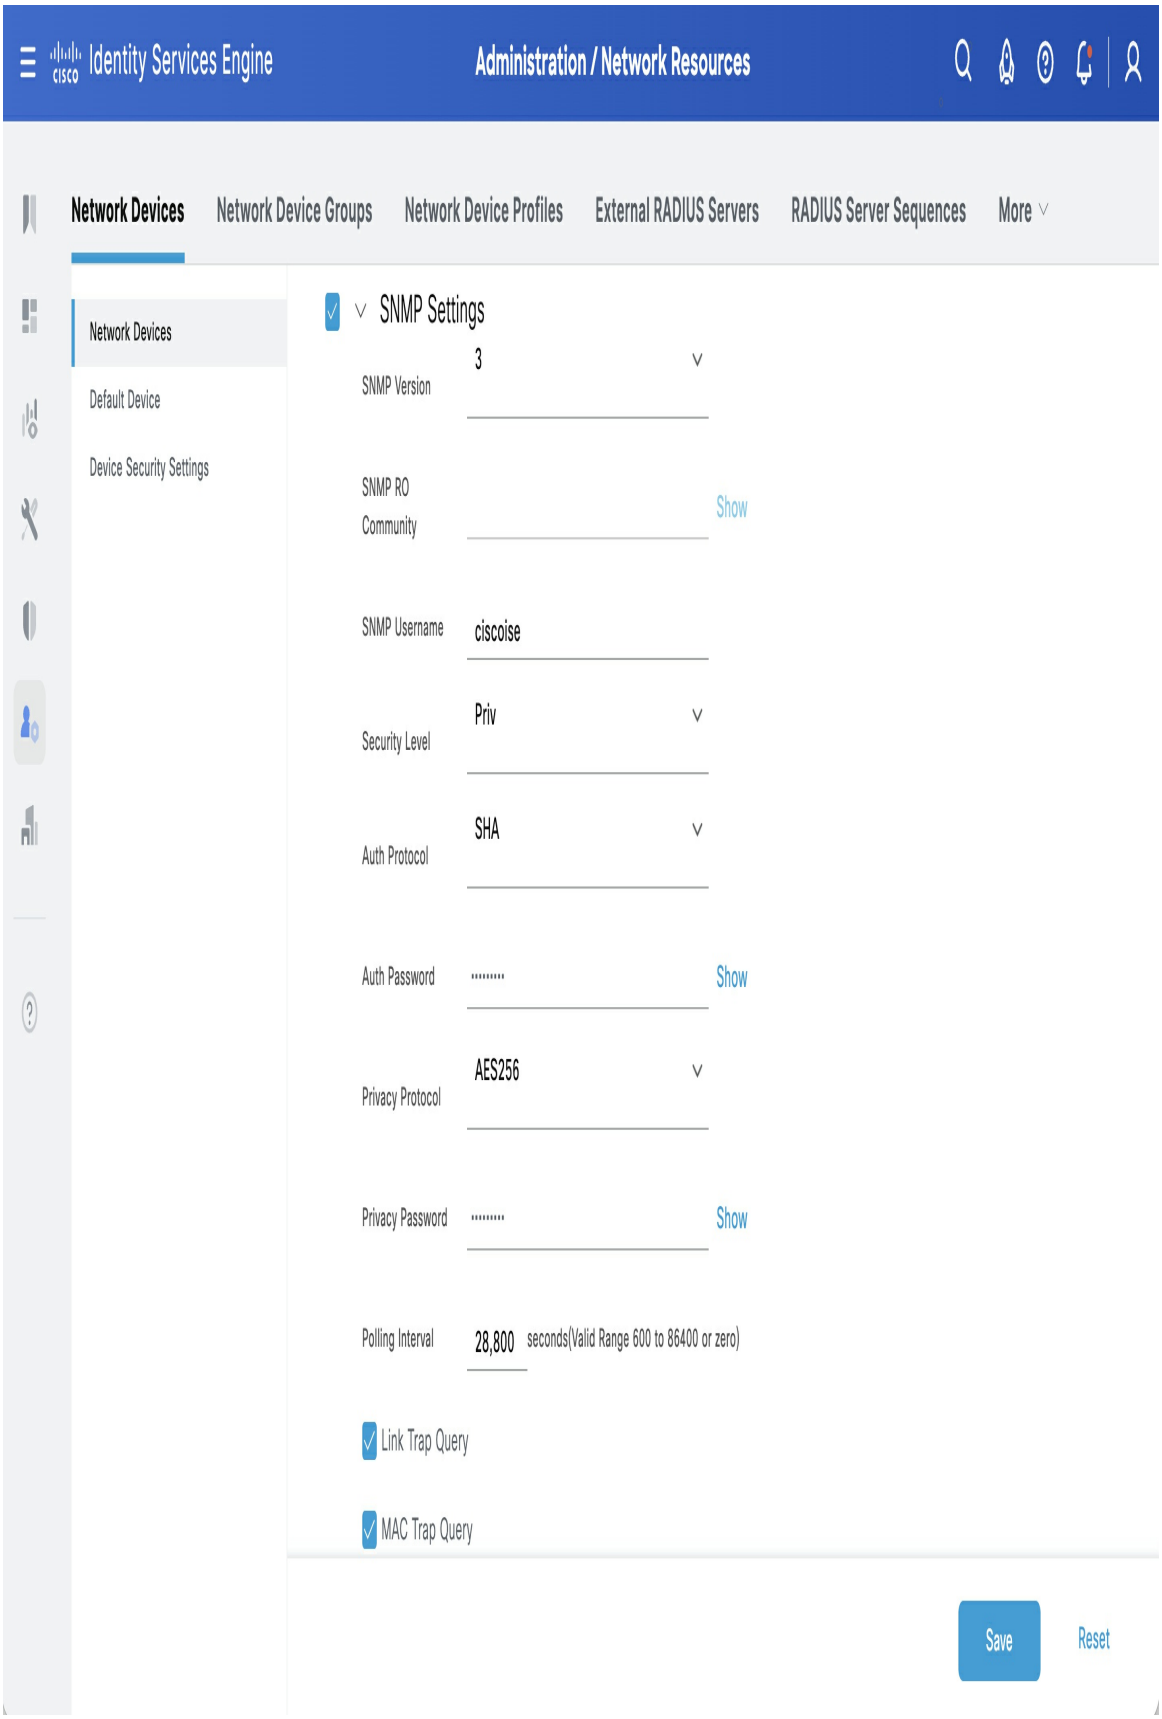

Identity Services Engine
Administration / Network Resources
Network Devices
Network Device Groups
Network Device Profiles
External RADIUS Servers
RADIUS Server Sequences
More
Network Devices
Default Device
Device Security Settings
SNMP Settings
SNMP Version
3
SNMP RO Community
Show
SNMP Username
ciscoise
Security Level
Priv
Auth Protocol
SHA
Auth Password
Show
Privacy Protocol
AES256
Privacy Password
Show
Polling Interval
28,800
seconds(Valid Range 600 to 86400 or zero)
Link Trap Query
MAC Trap Query
Save
Reset

Figure 8-10 *Network Device SNMP Settings*

Active Directory Probe

An Active Directory probe is especially useful when you want to use Active Directory data to make policy decisions. The following attributes can be fetched using the AD probe:

- AD-Host-Exists
- AD-Join-Point
- AD-Operating-System
- AD-OS-Version
- AD-Service-Pack

It is common to use the AD probe to distinguish access between corporate and noncorporate devices. You can accomplish this by configuring a profiling policy that combines the AD-Host-Exists attribute, which confirms the machine is present in Active Directory, with the AD-Join-Point attribute, which confirms the domain the machine is joined to.

The host lookup is performed based on the computer name of the endpoint, so it is essential for ISE to have this information. The hostname can be retrieved using the following probes:

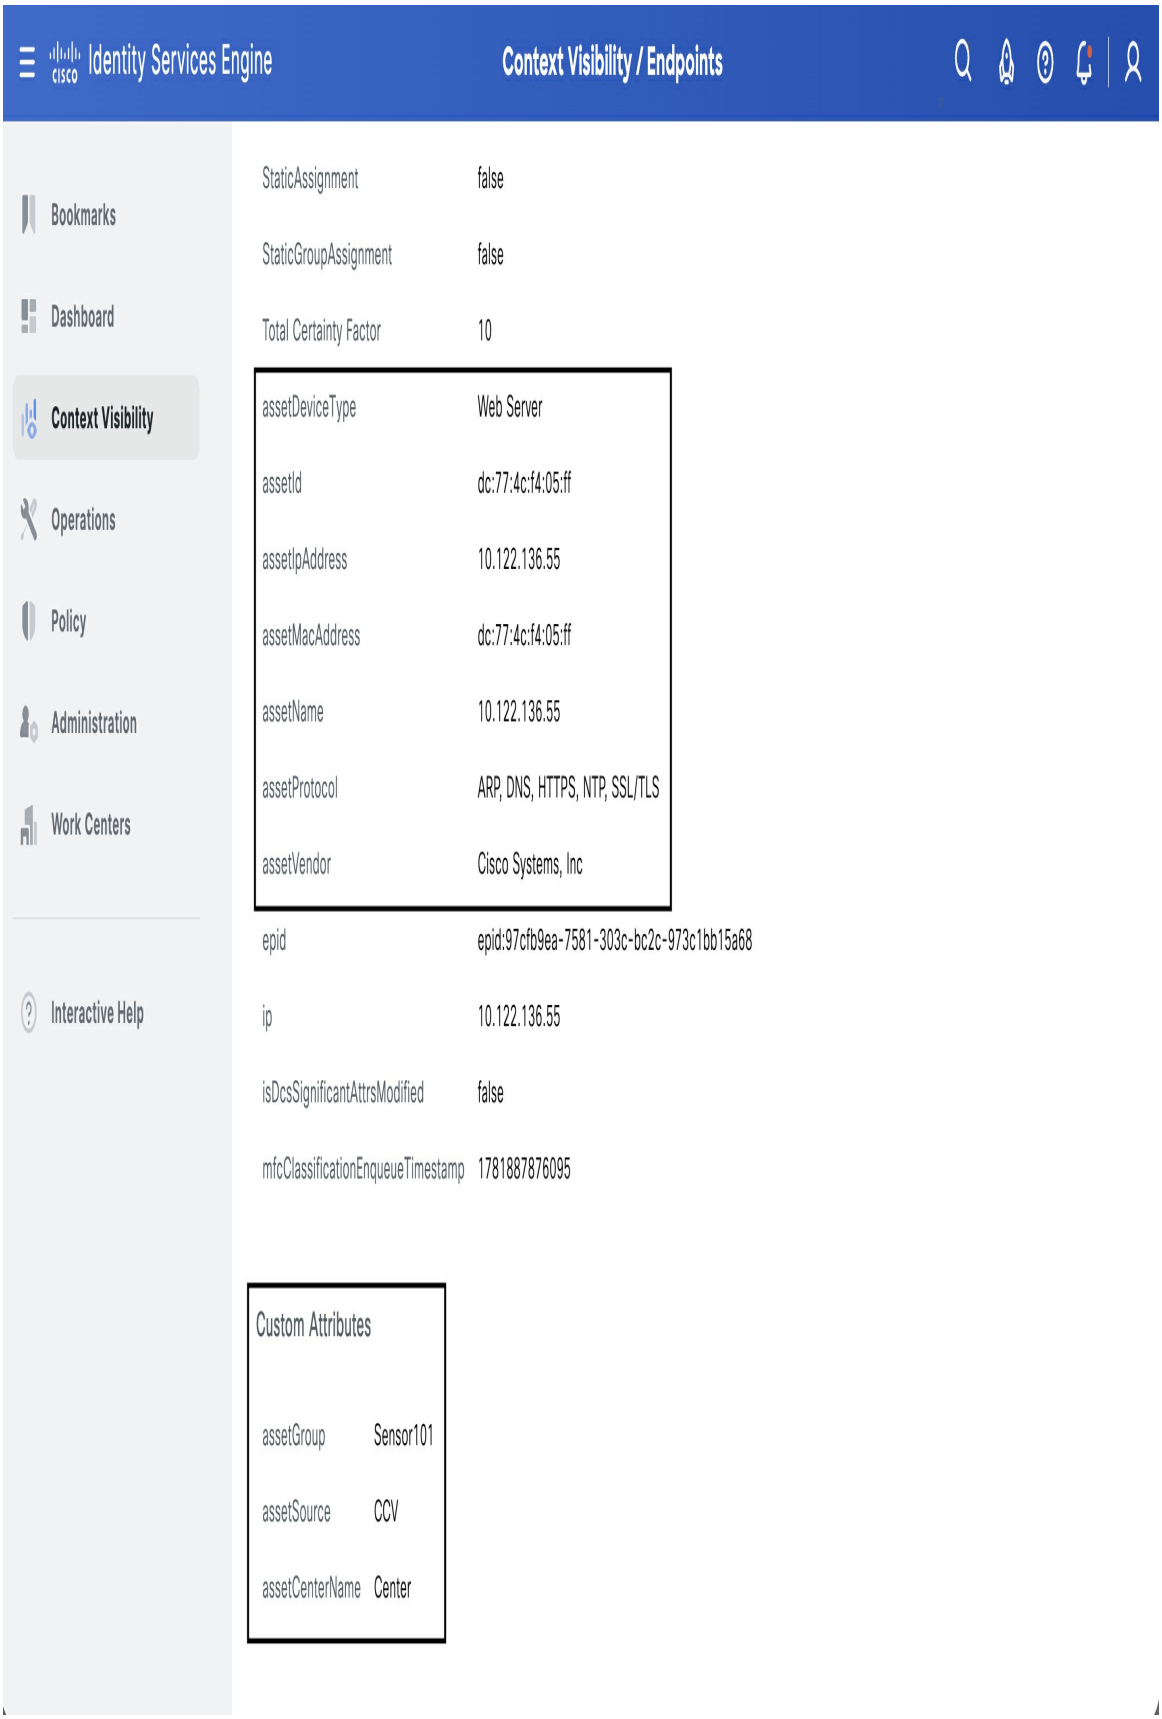
- DHCP Probe
- DNS Probe

Once you enable the AD probe and integrate ISE nodes with Active Directory, the AD probe will query Active Directory for attributes as soon as the hostname of the endpoint is received.

Technet24

Profiling Using pxGrid Probe

Another capability of ISE is the utilization of the pxGrid probe. The pxGrid probe allows ISE to create endpoints based on data received via the pxGrid protocol. A notable example of this integration is Cisco Cyber Vision. During pxGrid registration, Cisco Cyber Vision acts as a publisher and sends a pxGrid STOMP message to ISE for detected OT/IT endpoints, including attributes it has learned about the devices, so the corresponding endpoint is created. Figure 8-11 shows the endpoint attributes from Context Visibility that are received from Cyber Vision. You can see some built-in asset attributes, such as assetDeviceType and assetProtocol, as well as custom attributes like assetGroup and assetSource.

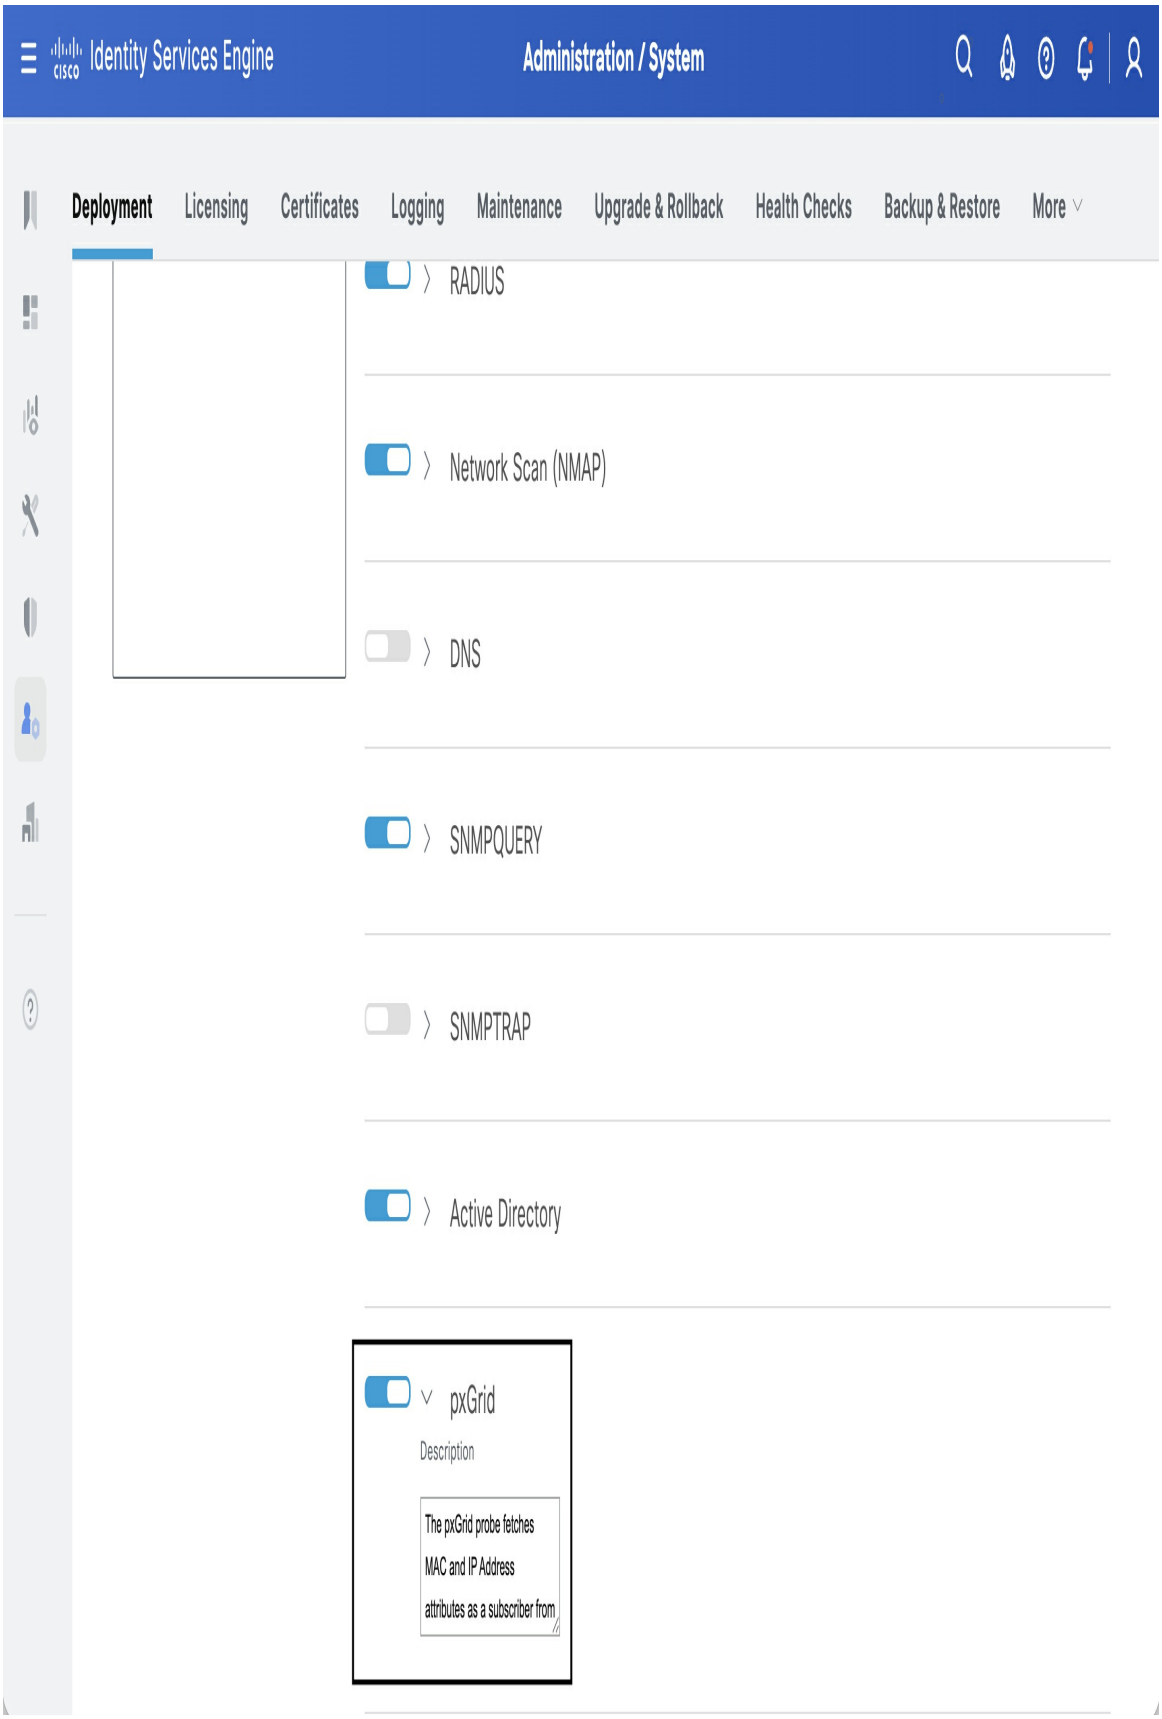

Identity Services Engine
Context Visibility / Endpoints
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
StaticAssignment false
StaticGroupAssignment false
Total Certainty Factor 10
assetDeviceType Web Server
assetId dc:77:4c:f4:05:ff
assetIpAddress 10.122.136.55
assetMacAddress dc:77:4c:f4:05:ff
assetName 10.122.136.55
assetProtocol ARP, DNS, HTTPS, NTP, SSL/TLS
assetVendor Cisco Systems, Inc
epid epid:97cfb9ea-7581-303c-bc2c-973c1bb15a68
ip 10.122.136.55
isDcsSignificantAttrsModified false
mfcClassificationEnqueueTimestamp 1781887876095
Custom Attributes
assetGroup Sensor101
assetSource CCV
assetCenterName Center

Technet24

Figure 8-11 *Endpoint Asset Attributes in Context Visibility*

You can also use these attributes in profiling policies to create dedicated profiles. A sample use case is SGT assignment based on assetGroups. The strength of this approach is in its automation: When you change the group in Cyber Vision, ISE is automatically updated with the new information. The profile changes if corresponding profiling policies are configured, and a Change of Authorization (CoA) is sent to update the endpoint context, assigning a new SGT based on the new assetGroup if corresponding authorization policies are in place.

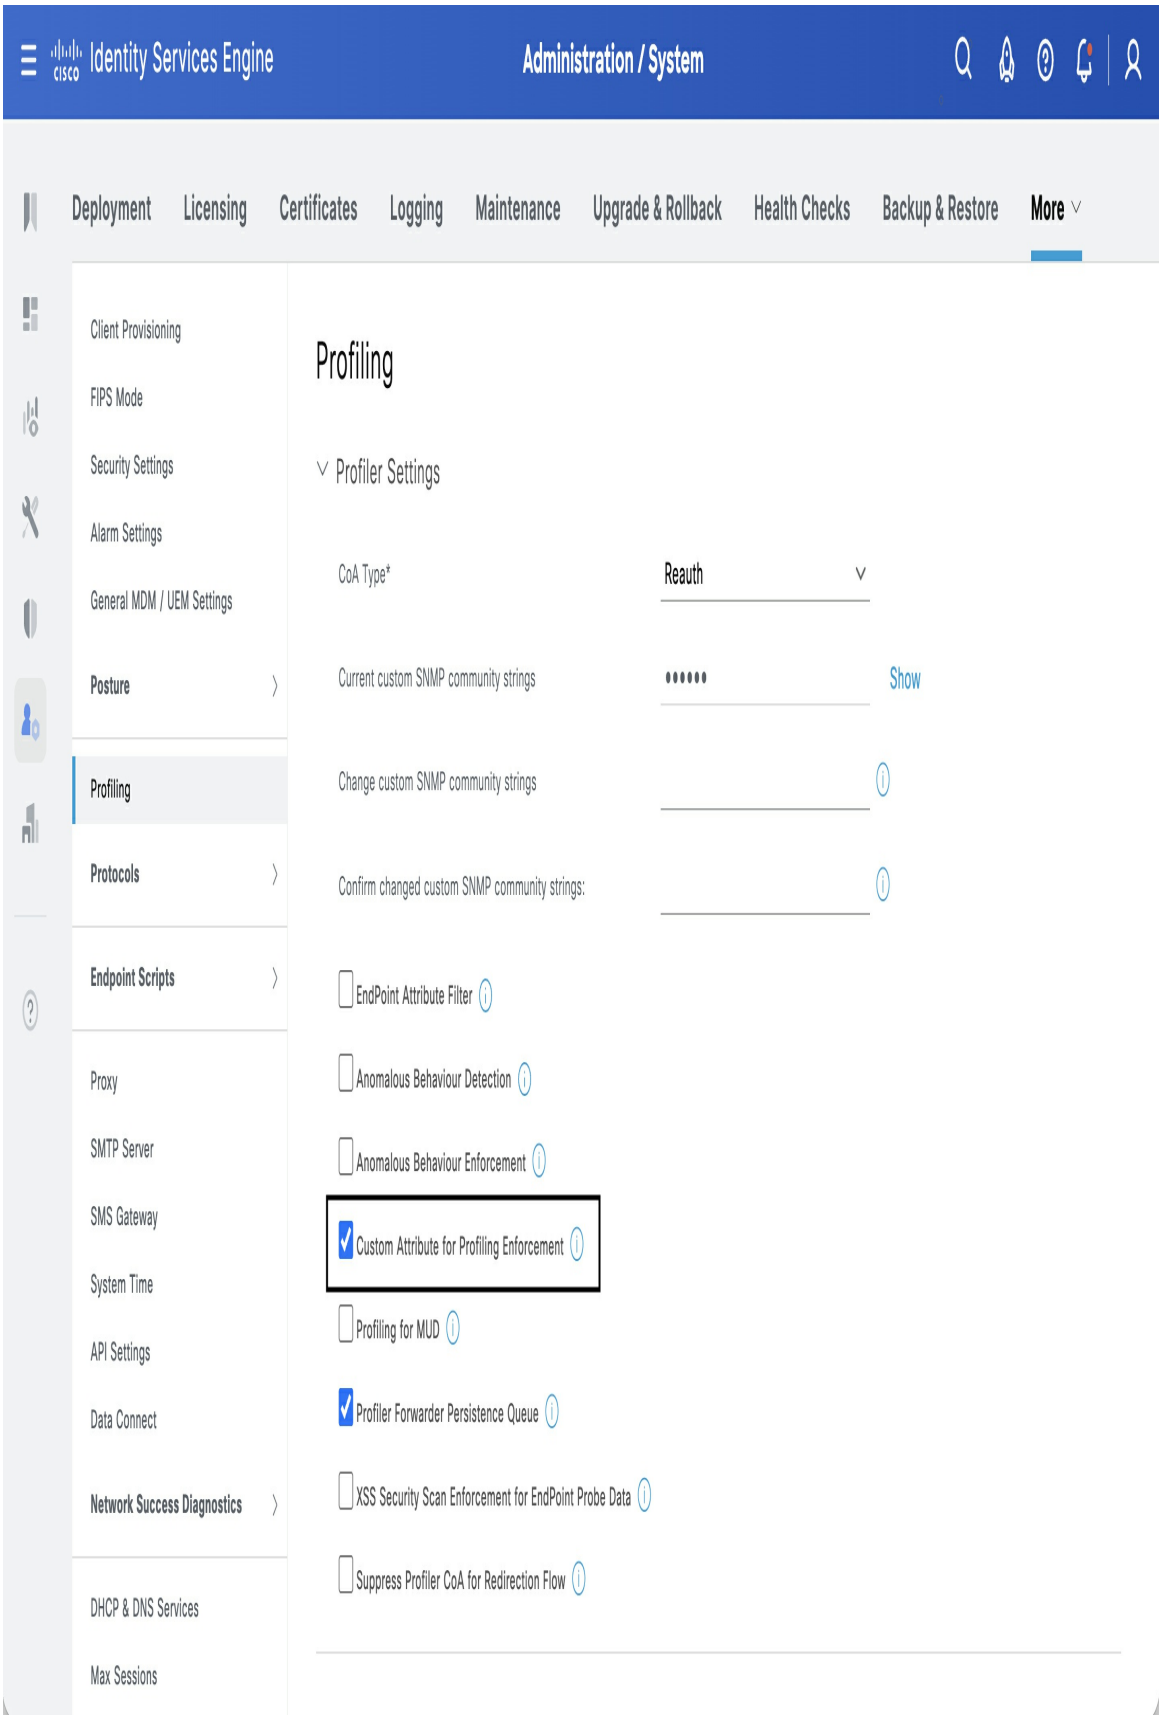
For ISE to create endpoints when integrated with Cyber Vision, you must enable the pxGrid probe under **Administration > Deployment > (Selected PSN) > Profiling Configuration**, as shown in Figure 8-12.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
RADIUS
Network Scan (NMAP)
DNS
SNMPQUERY
SNMPTRAP
Active Directory
pxGrid
Description
The pxGrid probe fetches
MAC and IP Address
attributes as a subscriber from

Technet24

Figure 8-12 *pxGrid Probe Configuration*

To populate custom attributes and use them in profile creation, enable the **Custom Attributes for Profiling Enforcement** setting under **Administration > System > Profiling**, as shown in Figure 8-13.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Client Provisioning
FIPS Mode
Security Settings
Alarm Settings
General MDM / UEM Settings
Posture
Profiling
Protocols
Endpoint Scripts
Proxy
SMTP Server
SMS Gateway
System Time
API Settings
Data Connect
Network Success Diagnostics
DHCP & DNS Services
Max Sessions
Profiling
Profiler Settings
CoA Type*
Reauth
Current custom SNMP community strings
Show
Change custom SNMP community strings
Confirm changed custom SNMP community strings:
EndPoint Attribute Filter
Anomalous Behaviour Detection
Anomalous Behaviour Enforcement
Custom Attribute for Profiling Enforcement
Profiling for MUD
Profiler Forwarder Persistence Queue
XSS Security Scan Enforcement for EndPoint Probe Data
Suppress Profiler CoA for Redirection Flow

Technet24

Figure 8-13 *Custom Attributes for Profiling Enforcement Setting*

Note

Complete pxGrid integration between ISE and Cisco Cyber Vision is not covered in this book. For more information on pxGrid integrations, refer to Chapter 19, "Integrating Your Security Products to Enhance Context Using pxGrid."

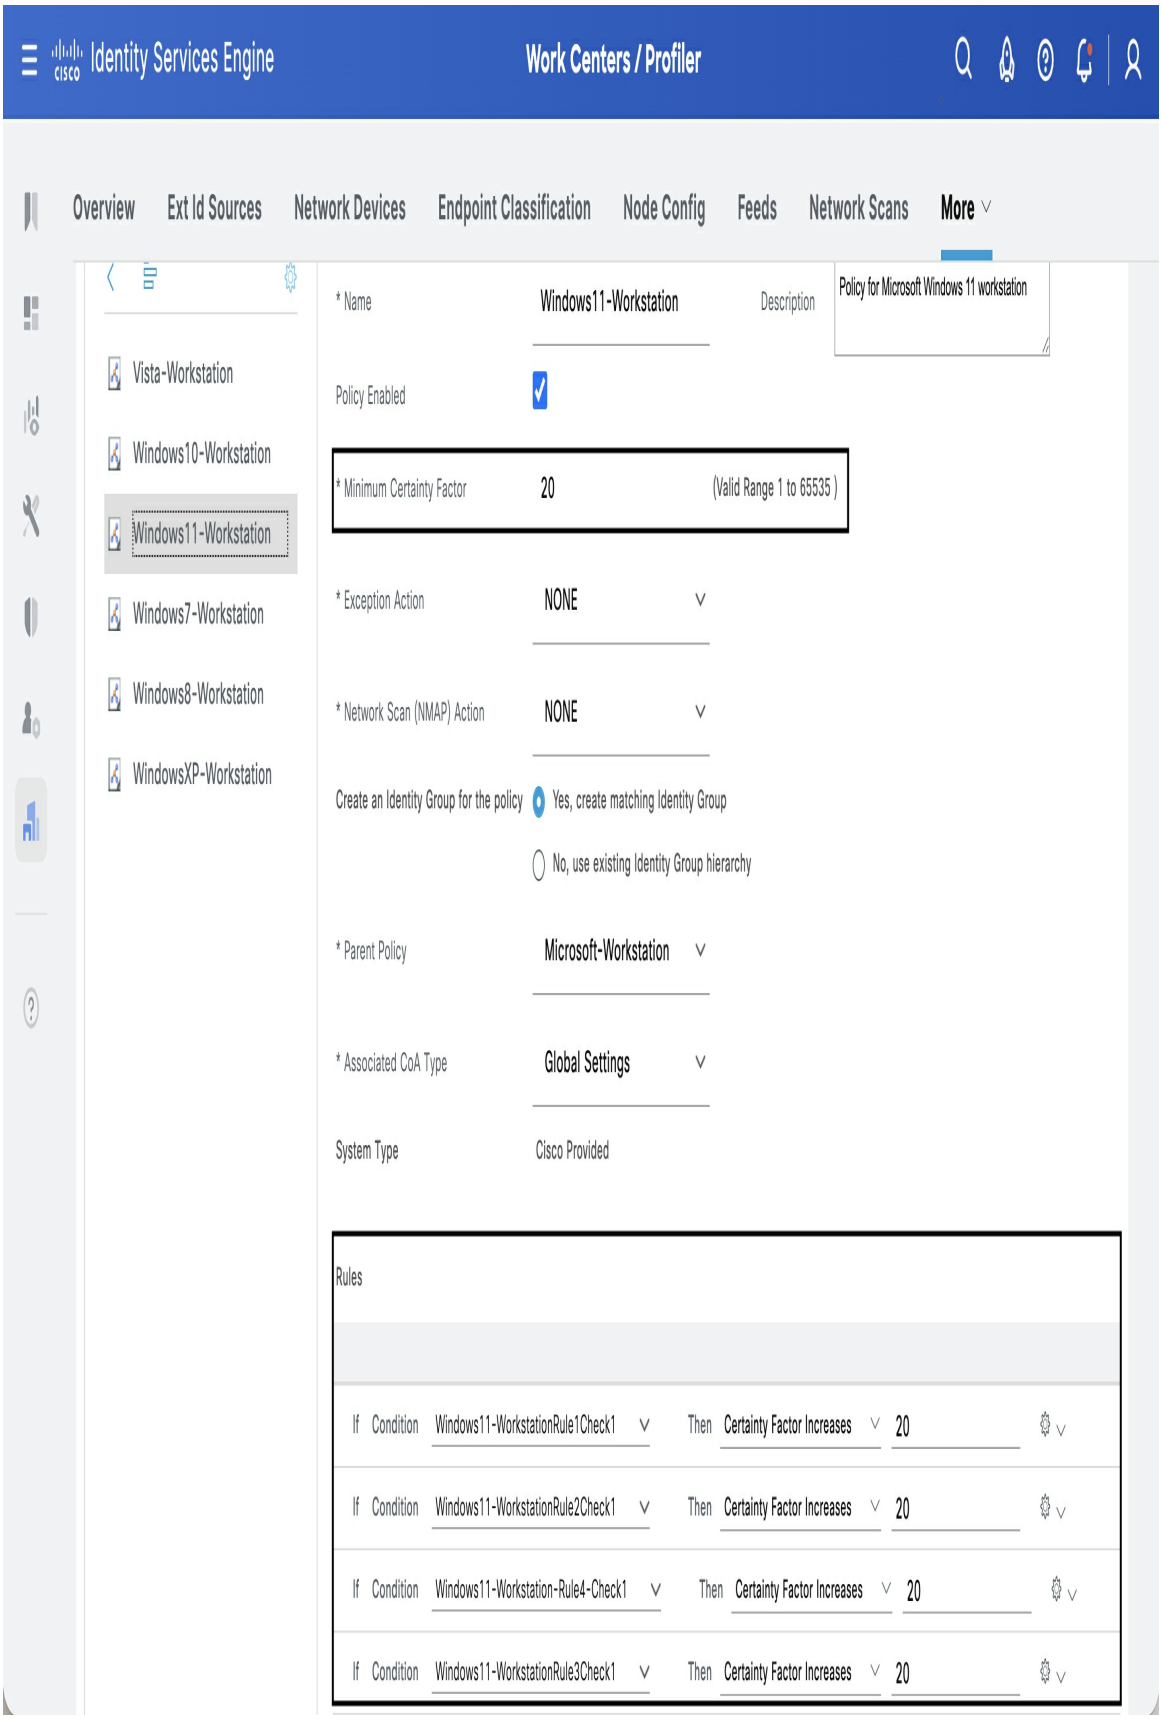
Profiling Policies

There are two types of profiling policies:

- Cisco provided (via Feed Service)
- Administrator created (manually created)

Profiling policies operate based on certainty factors. In simple terms, this is the number of points an endpoint must earn to be profiled. For example, the Windows11-Workstation profile shown in Figure 8-14 has a Minimum Certainty Factor of 20, meaning the endpoint must obtain 20 points to be profiled as a Windows11-Workstation.

How does an endpoint earn points? Points are awarded based on matching rules. As shown in Figure 8-14, the Windows11-Workstation profile has four rules, and each individual match grants the endpoint 20 points, which means matching a single rule is sufficient for this profile to be assigned.


Identity Services Engine
Work Centers / Profiler
Overview
Ext Id Sources
Network Devices
Endpoint Classification
Node Config
Feeds
Network Scans
More
Vista-Workstation
Windows10-Workstation
Windows11-Workstation
Windows7-Workstation
Windows8-Workstation
WindowsXP-Workstation
* Name
Windows11-Workstation
Description
Policy for Microsoft Windows 11 workstation
Policy Enabled
* Minimum Certainty Factor
20
(Valid Range 1 to 65535)
* Exception Action
NONE
* Network Scan (NMAP) Action
NONE
Create an Identity Group for the policy
Yes, create matching Identity Group
No, use existing Identity Group hierarchy
* Parent Policy
Microsoft-Workstation
* Associated CoA Type
Global Settings
System Type
Cisco Provided
Rules
If Condition Windows11-WorkstationRule1Check1 Then Certainty Factor Increases 20
If Condition Windows11-WorkstationRule2Check1 Then Certainty Factor Increases 20
If Condition Windows11-Workstation-Rule4-Check1 Then Certainty Factor Increases 20
If Condition Windows11-WorkstationRule3Check1 Then Certainty Factor Increases 20

Technet24

Figure 8-14 *Windows11-Workstation Profiling Policy*

Before examining the rules, let's review other settings available in the profiling policy.

To classify an endpoint, you must ensure that the profile is enabled. You can create a dedicated identity group, where endpoints are placed once profiled. The **Parent Policy** is part of endpoint hierarchy; for example, you can have multiple Windows endpoints with different operating systems, and they will all belong to the parent policy — Microsoft Workstation.

> ## Note
>
> If the endpoint is part of a hierarchy, it must first meet the minimum certainty factor for the parent policy.

The **Associated CoA Type** controls which type of CoA packet is sent when an endpoint receives a new profile assignment.

If an endpoint matches two different policies at the same time, there is no logic to determine which profile is assigned. In such cases, it is recommended to adjust the certainty factors of one of the policies. In most cases, however, the default policies should be left with their default values.

Unlike authentication or authorization policies, you can match multiple profiling rules simultaneously. For example, the Windows11-Workstation profile has four rules. If each condition is met, the endpoint can theoretically earn 80 points.

Now let's examine the conditions to understand which probes can contribute to Windows11-Workstation profile in Table 8-1.

Table 8-1 *Windows 11 Workstation Condition Details*

Name	Expression
Windows11-WorkstationRule1Check1	ACTIVEDIRECTORY_PROBE:AD-Operating-System CONTAINS Windows 11
Windows11-WorkstationRule2Check1	ACIDEX:device-platform-version STARTSWITH 10.0.2
Windows11-Workstation-Rule4-Check1	IP:User-Agent CONTAINS WINDOWS_11
Windows11-WorkstationRule3Check1	Wifi_Device_Analytics:DEVICE_INFO_MODEL_NUM CONTAINS Windows 11

These conditions indicate that the Windows 11 profile can be assigned to the endpoint if

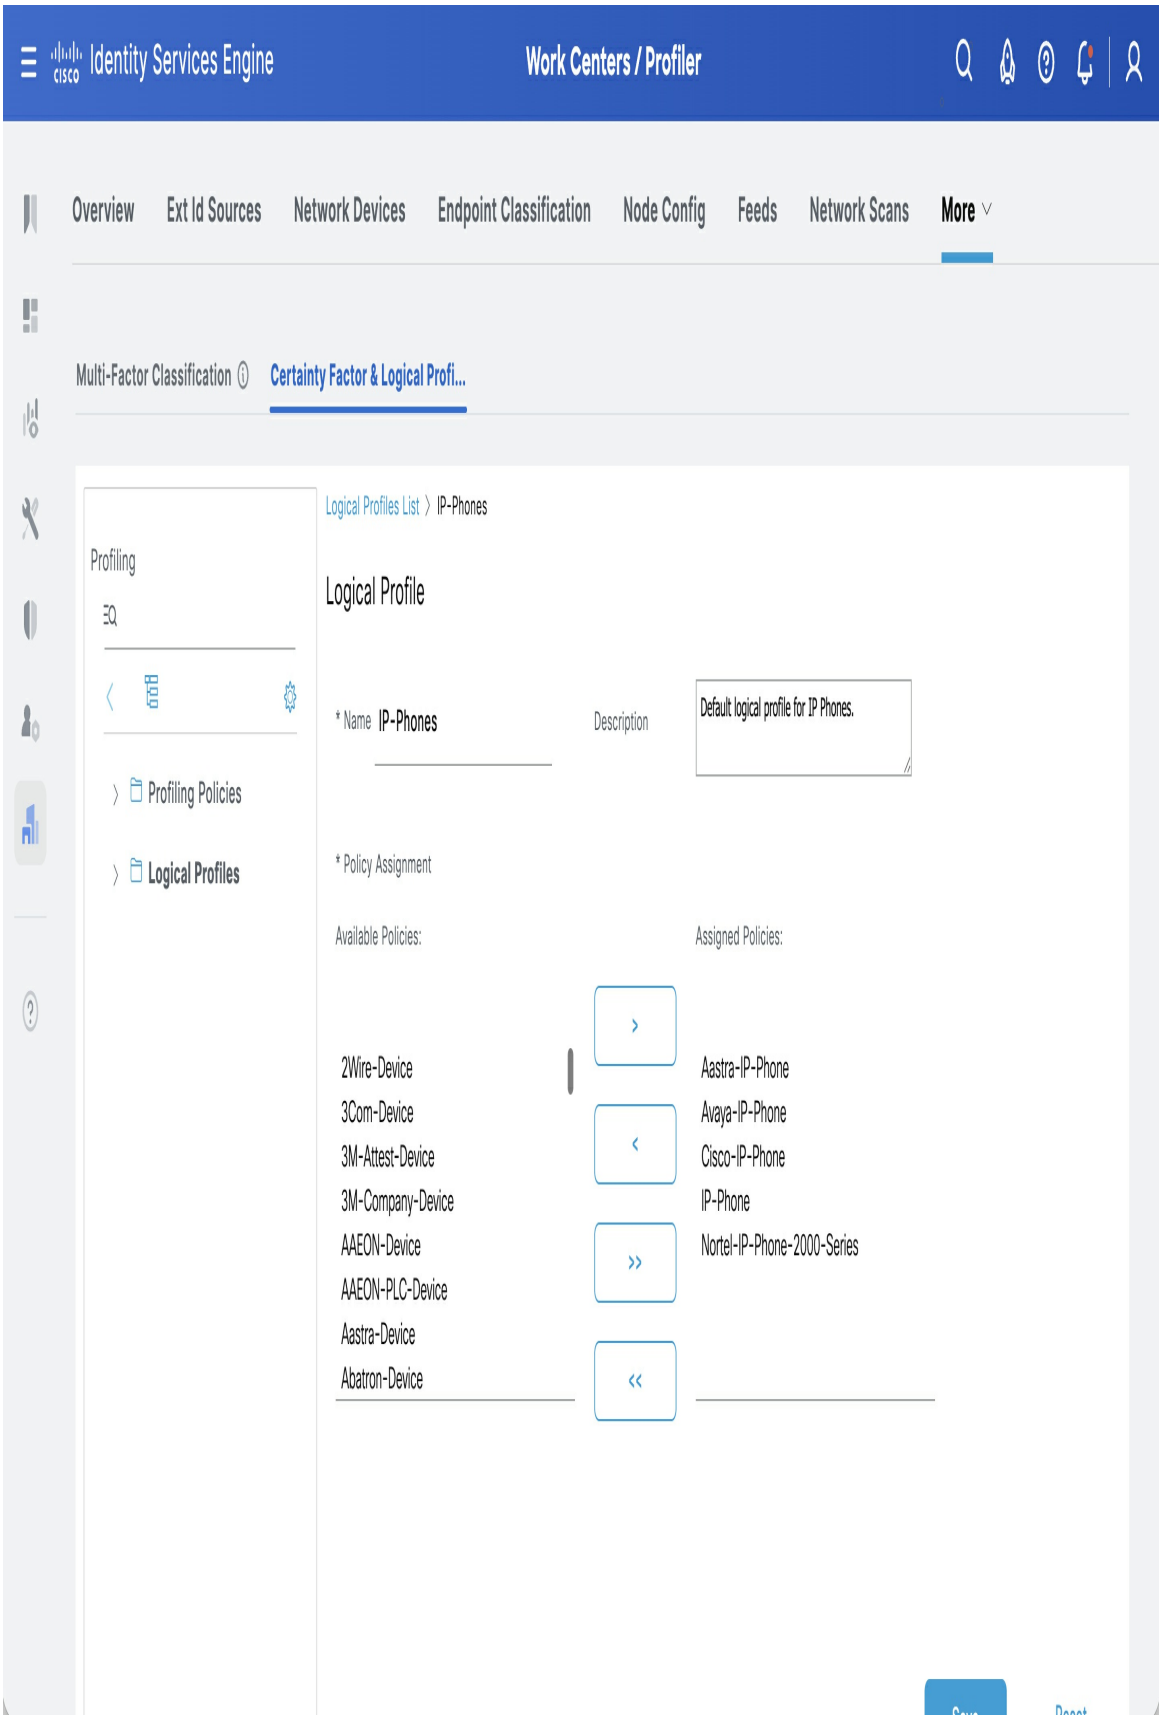
- The Active Directory Probe is configured, and the AD-Operating-System attribute is present in Active Directory.
- The endpoint connects via VPN, and ACIDex attributes are present in RADIUS Accounting (in this case, the RADIUS Probe is used). ACIDex attributes are Cisco Secure Client Identity Extensions, which provide a powerful method for profiling. Each time an endpoint connects to ISE via VPN, Secure Client sends a list of learned attributes. These attributes are sent in mdm-tlv attributes.
- The HTTP Probe is used, and User-Agent information is populated.
- The endpoint connects via wireless, and Wi-Fi Device Analytics is configured on the wireless LAN controller (in this case, the RADIUS Probe is used).

Logical Profiles

Technet24

In a multivendor environment, it is important to maintain security without sacrificing visibility. For example, you may have multiple vendors of phones, yet all devices should belong to the same voice VLAN. You may also have different IP camera manufacturers, but all should have the necessary permissions to access network video recorders.

Logical profiles support such use cases. As shown in Figure 8-15, ISE provides multiple built-in logical profiles. The IP-Phones profile combines several parent profiles from different vendors.

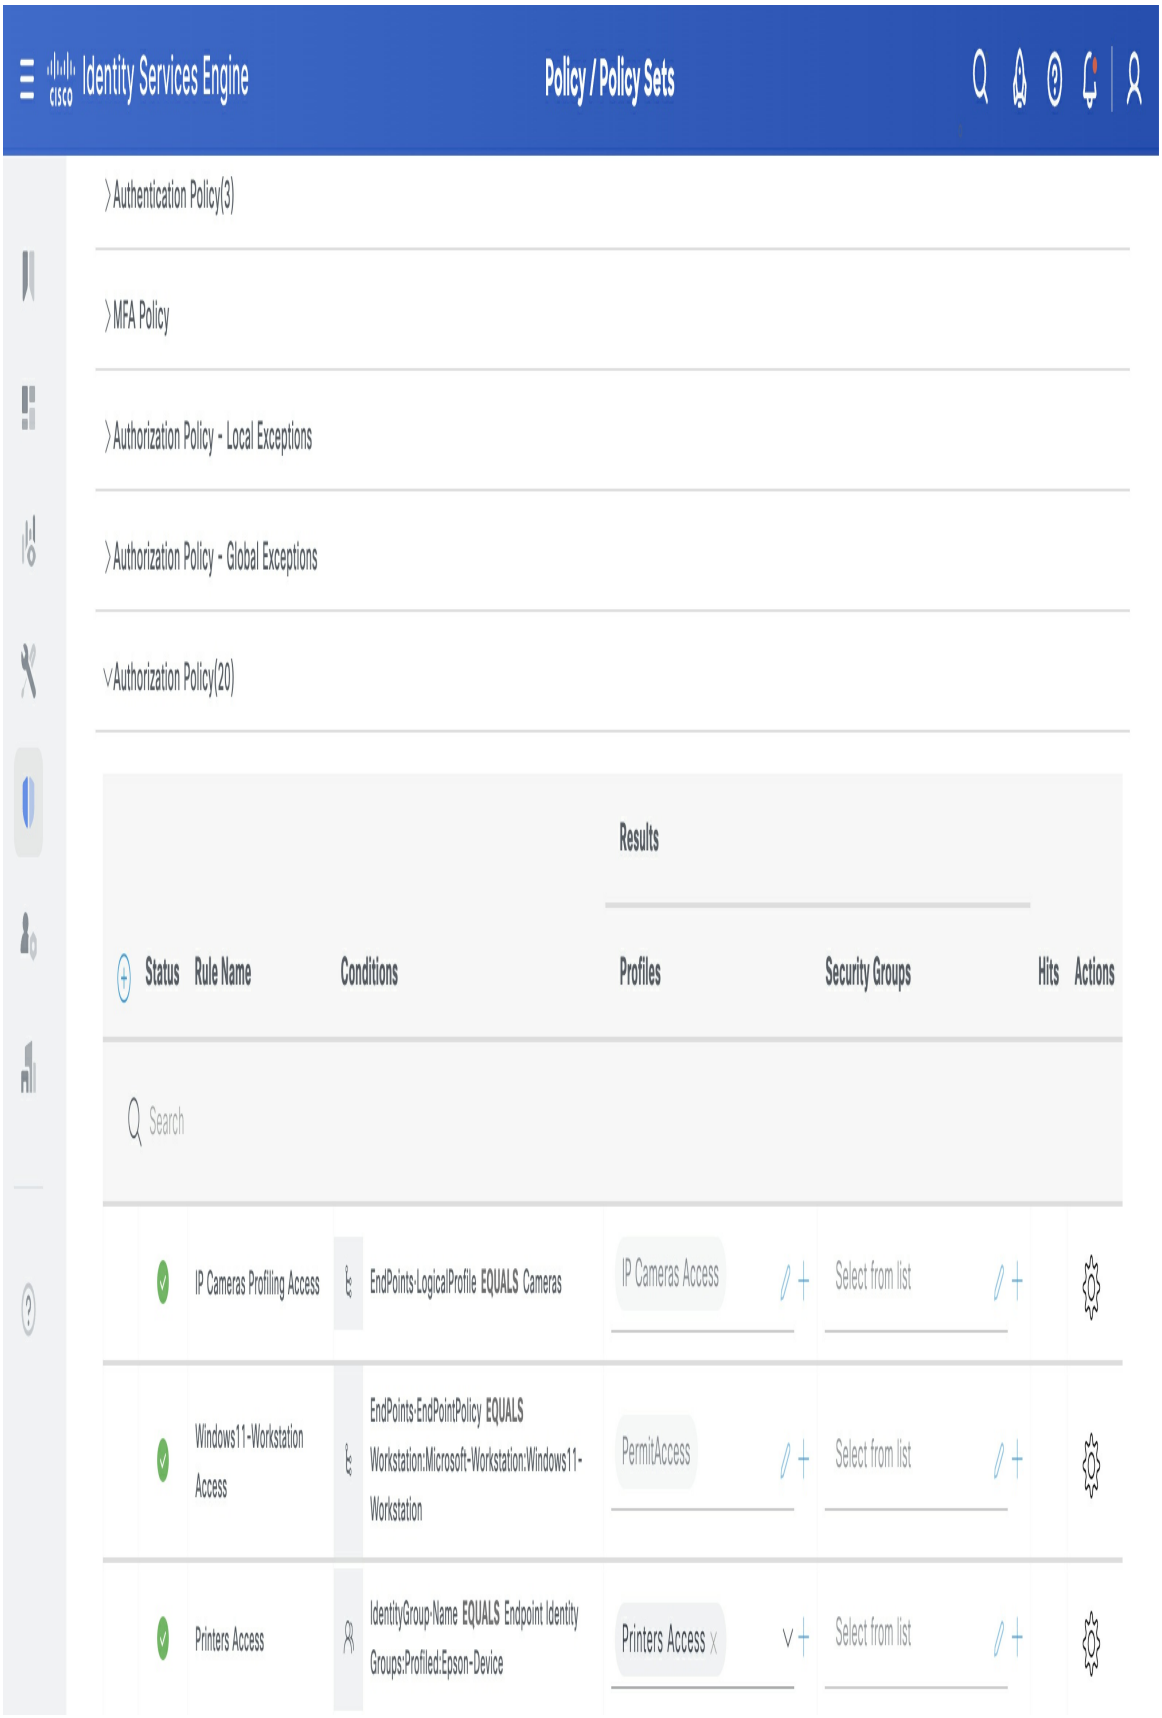

Identity Services Engine
Work Centers / Profiler
Overview
Ext Id Sources
Network Devices
Endpoint Classification
Node Config
Feeds
Network Scans
More
Multi-Factor Classification
Certainty Factor & Logical Profi...
Profiling
Profiling Policies
Logical Profiles
Logical Profiles List > IP-Phones
Logical Profile
* Name IP-Phones
Description
Default logical profile for IP Phones.
* Policy Assignment
Available Policies:
2Wire-Device
3Com-Device
3M-Attest-Device
3M-Company-Device
AAEON-Device
AAEON-PLC-Device
Aastra-Device
Abatron-Device
Assigned Policies:
Aastra-IP-Phone
Avaya-IP-Phone
Cisco-IP-Phone
IP-Phone
Nortel-IP-Phone-2000-Series

Technet24

Figure 8-15 *ISE Logical Profiles*

Authorization Using Profiling

With the knowledge gained about endpoint profiles, let's examine how they can be used to enhance authorization policies. Figure 8-16 illustrates several scenarios where profiling information improves the authorization policy. The IP Cameras Profiling Access policy relies on a logical profile condition. The Windows11-Workstation Access policy relies on the endpoint profile, and the Printers Access policy relies on the identity group.


Identity Services Engine
Policy / Policy Sets
Authentication Policy(3)
MFA Policy
Authorization Policy - Local Exceptions
Authorization Policy - Global Exceptions
Authorization Policy(20)
Results
Status
Rule Name
Conditions
Profiles
Security Groups
Hits
Actions
Search
IP Cameras Profiling Access
EndPoints·LogicalProfile EQUALS Cameras
IP Cameras Access
Select from list
Windows11-Workstation Access
EndPoints·EndPointPolicy EQUALS Workstation:Microsoft-Workstation:Windows11-Workstation
PermitAccess
Select from list
Printers Access
IdentityGroup·Name EQUALS Endpoint Identity Groups:Profiled:Epson-Device
Printers Access
Select from list

Technet24

Figure 8-16 *ISE Authorization Policy Configuration Based on Profiling Information*

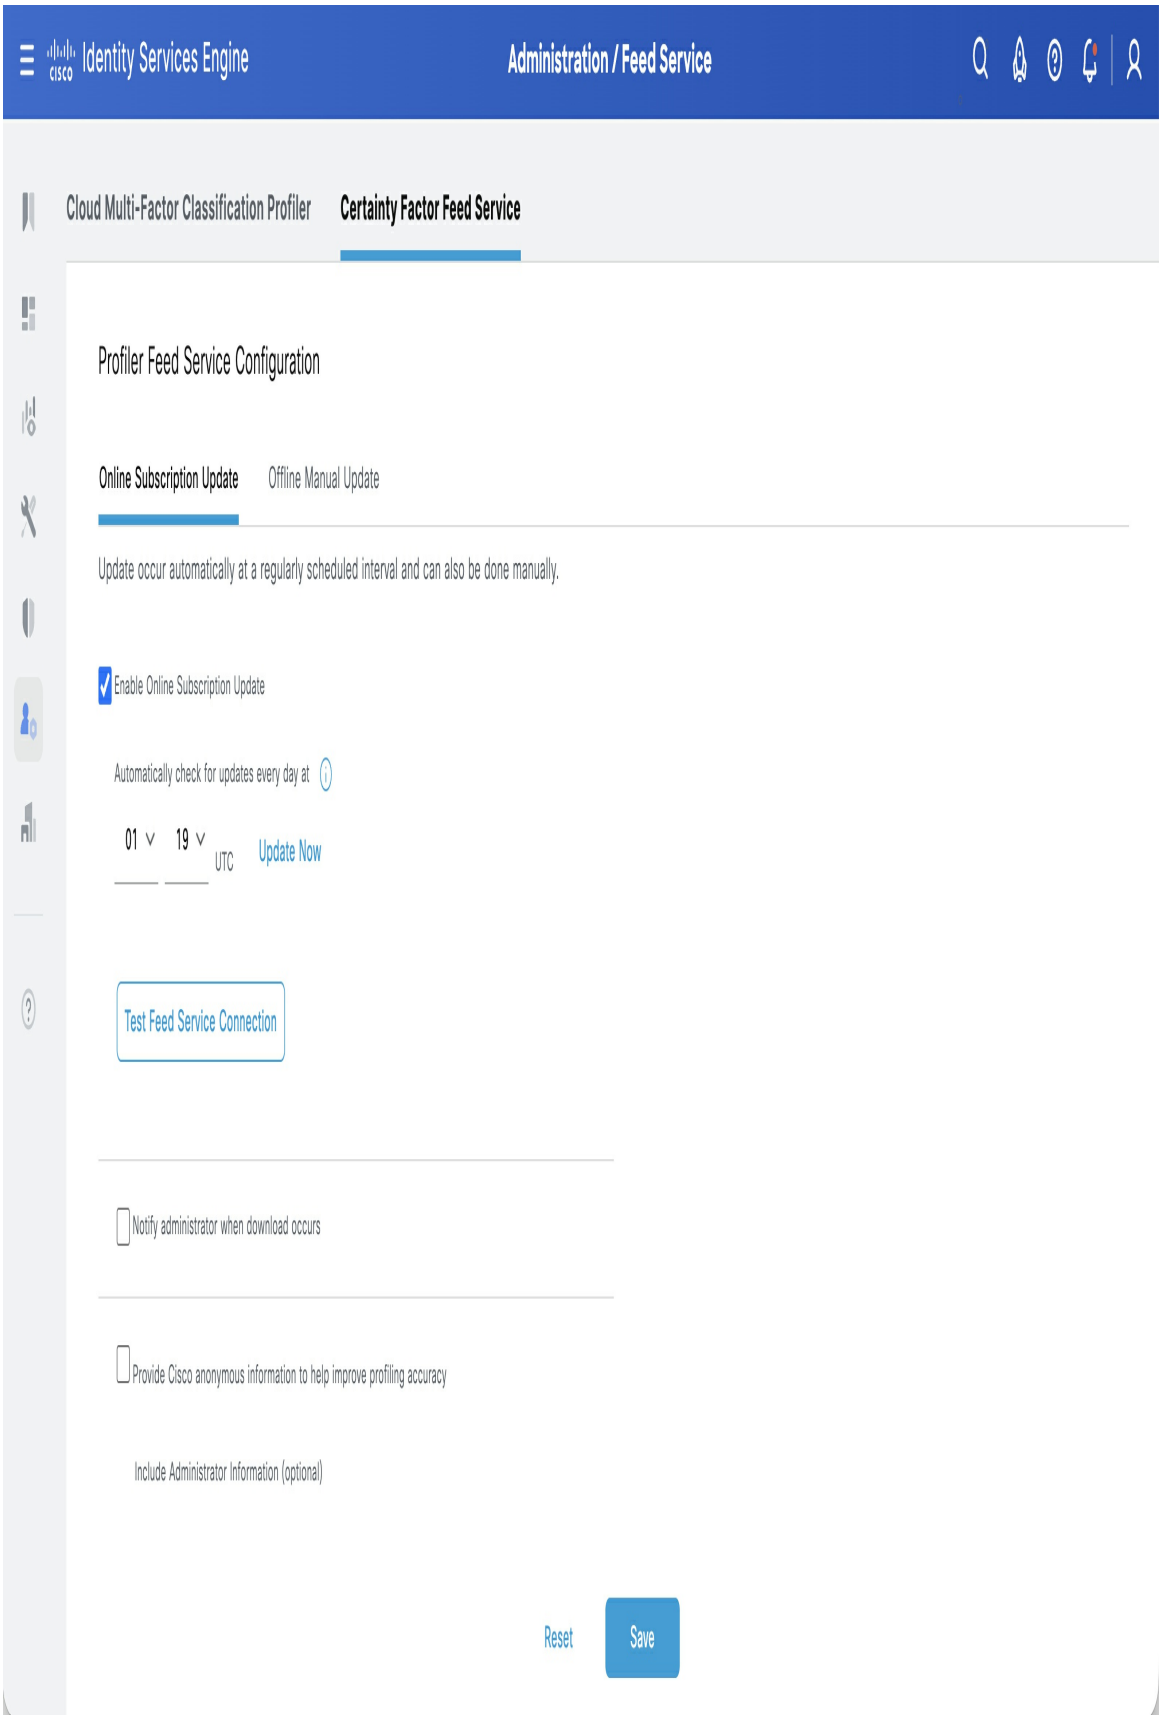
Important:

Profiling policy is not a replacement for authentication or authorization. The purpose of profiling conditions in the policy is to enhance existing policies, such as those relying on Active Directory groups for user access or internal endpoint identity groups for endpoint access.

Feed Service

Feed Service on ISE enables automatic downloading of new profiles approved by Cisco and updates OUI information.

You configure Feed Service under **Administration > Feed Service > Certainty Factor Feed Service**, as shown in Figure 8-17.


Identity Services Engine
Administration / Feed Service
Cloud Multi-Factor Classification Profiler
Certainty Factor Feed Service
Profiler Feed Service Configuration
Online Subscription Update
Offline Manual Update
Update occur automatically at a regularly scheduled interval and can also be done manually.
Enable Online Subscription Update
Automatically check for updates every day at
01
19
UTC
Update Now
Test Feed Service Connection
Notify administrator when download occurs
Provide Cisco anonymous information to help improve profiling accuracy
Include Administrator Information (optional)
Reset
Save

Technet24

Figure 8-17 *Certainty Factor Feed Service*

You can download updates automatically from the Internet using the **Online Subscription Update** tab. From the same screen, you can test the Feed Service connection to ensure end-to-end connectivity.

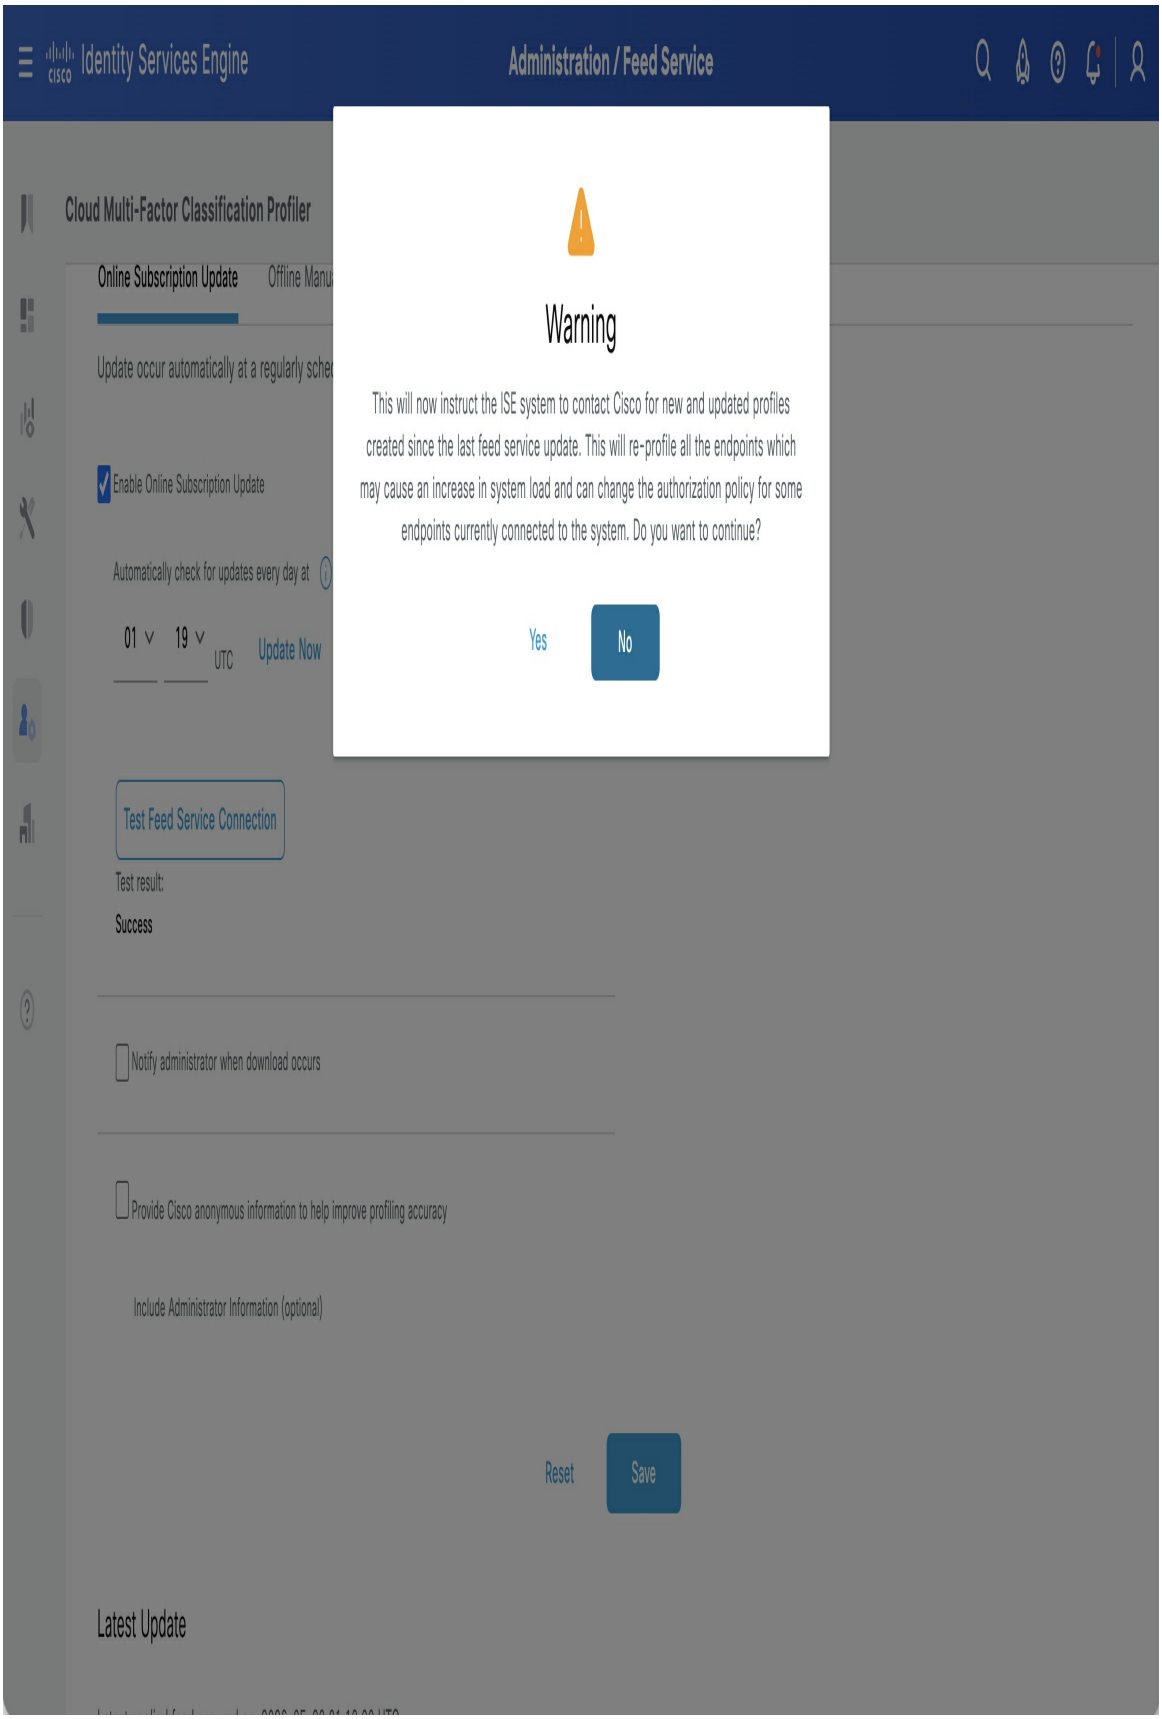
Alternatively, if ISE does not have an Internet connection, you can use the **Offline Manual Update** option. Offline Feed updates can be downloaded from https://ise.cisco.com.

Best Practice:

You should enable Online Subscription Updates to maintain an up-to-date OUI database and profiles. As vendors introduce new products and acquisitions occur, updated profiles ensure better visibility into your network.

Note

When a Feed Update is completed or when you click **Update Now**, as shown in Figure 8-18, reprofiling of all endpoints occurs. It is recommended to perform updates during offline hours.


Identity Services Engine
Administration / Feed Service
Cloud Multi-Factor Classification Profiler
Online Subscription Update
Offline Manu
Update occur automatically at a regularly sche
Enable Online Subscription Update
Automatically check for updates every day at
01
19
UTC
Update Now
Test Feed Service Connection
Test result:
Success
Notify administrator when download occurs
Provide Cisco anonymous information to help improve profiling accuracy
Include Administrator Information (optional)
Reset
Save
Latest Update
Warning
This will now instruct the ISE system to contact Cisco for new and updated profiles created since the last feed service update. This will re-profile all the endpoints which may cause an increase in system load and can change the authorization policy for some endpoints currently connected to the system. Do you want to continue?
Yes
No

Technet24

Figure 8-18 *Certainty Factor Feed Service Update Now Execution*

Best Practice:

You should create any admin profile with a minimum certainty factor of 1000 or higher to prevent Cisco-provided feed profiles from overriding admin-created profiles, if this outcome is desired.

Profiler Flow Overview

Next, let's examine the profiling flow and what to expect from endpoints being profiled on ISE. Before starting the analysis, let's review the configured policies. Figure 8-19 shows three built-in policies of interest:

- Default Authorization Policy with Deny Access result
- Basic_Authenticated_Access Authorization Policy with PermitAccess result
- Profiled Cisco IP Phones Authorization Policy with Cisco_IP_Phones result

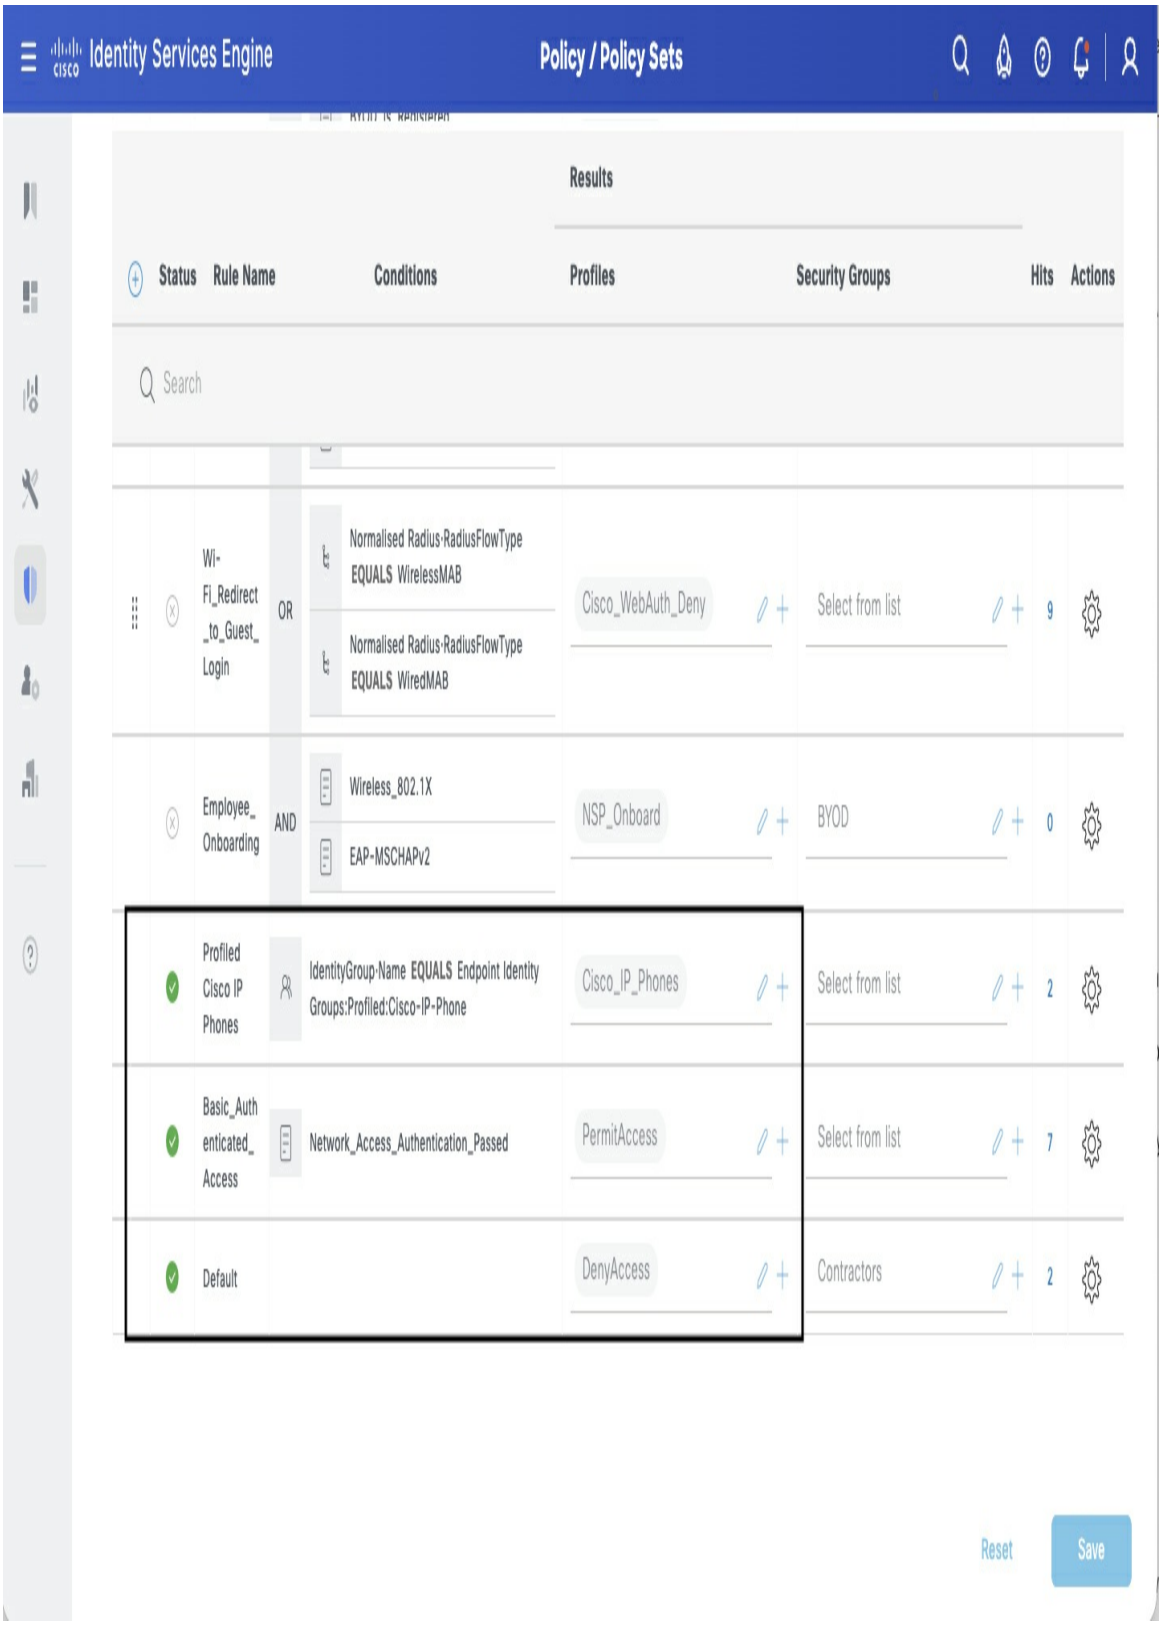


Figure 8-19 *Built-in Authorization Policies Used for Profiling Flow*

Technet24

Overview

Figure 8-20 shows the end-to-end profiling flow for the Cisco IP Phone 8865.



Jun 12, 2026 04:42:18.778 PM		AC:7E:8A:B6:AB:32	AC:7E:8A:B6:AB:32	Cisco-IP-Phone-8865	Default >> MAB	Default >> Profiled Cisco IP Phones
Jun 12, 2026 04:41:26.634 PM		#ACSACL#-IP-PERMIT_...				
Jun 12, 2026 04:41:26.583 PM		AC:7E:8A:B6:AB:32	AC:7E:8A:B6:AB:32	Cisco-IP-Phone-8865	Default >> MAB	Default >> Profiled Cisco IP Phones
Jun 12, 2026 04:41:26.537 PM			AC:7E:8A:B6:AB:32			
Jun 12, 2026 04:41:24.950 PM		AC:7E:8A:B6:AB:32	AC:7E:8A:B6:AB:32	Cisco-Device	Default >> MAB	Default >> Basic_Authenticated_Access
Jun 12, 2026 04:41:18.891 PM			AC:7E:8A:B6:AB:32			
Jun 12, 2026 04:41:07.685 PM		AC:7E:8A:B6:AB:32	AC:7E:8A:B6:AB:32		Default >> MAB	Default >> Default

Figure 8-20 *Live Logs for Profiling Flow Overview*

The flow illustrates the IP phone onboarding use case. This phone has never been seen by ISE and is being connected to the network for the first time. Review the process step by step, starting from the initial authentication, to understand what occurs during profiling.

- 04:41:07.685: The endpoint with MAC address AC:7E:8A:B6:AB:32 connected for the first time. Since ISE did not recognize it, authentication failed and Access-Reject was sent. Although Access-Reject was sent, the endpoint was added to the ISE Internal Endpoints database. No access was granted.
- 04:41:18.891: Ten seconds later, the first CoA packet was sent to the network access device because the endpoint was profiled for the first time. The 10-second interval is a hard-coded period for endpoints new to ISE. At this point, ISE identified the device as a

Cisco device based on OUI. CoA triggered reauthentication.

- 04:41:24.950: The second authentication attempt occurred. This time, authentication succeeded, and the Basic_Authenticated_Access Authorization Policy matched, which checks whether the endpoint is authenticated. Access-Accept was sent. Since the network access device received Access-Accept, it triggered an Accounting Start request containing Device Sensor information.

Note

Even though Access-Accept was sent in this case, the production recommendation is to grant limited permissions for basic authentication, such as Access-Accept with a DACL denying traffic. The goal is to limit access while ensuring Accounting data is sent for improved profiling evaluation.

- 04:41:26.537: A CoA packet was sent because the endpoint matched a more specific profiling policy, and it is used in the authorization policy. The endpoint was profiled as Cisco-IP-Phone-8865 and placed in the Cisco-IP-Phone Endpoint Identity Group.
- 04:41:26.583: The third authentication attempt occurred. Since the endpoint was part of the Cisco-IP-Phone Endpoint Identity Group, the "final" policy Profiled Cisco IP Phones matched, and Cisco_IP_Phones permissions were granted.
- 04:41:26.634: The built-in Cisco_IP_Phones Authorization Profile contains a DACL. The network access device requested the DACL content and assigned it to the endpoint.

As a result, the endpoint was automatically profiled as Cisco-IP-Phone-8865 out of the box and received appropriate permissions.

Technet24

Important:

The purpose of the profiling flow overview is to explain how profiling functions. Profiling should not replace authentication; you should use it only to enhance authentication. Whenever possible, use secure authentication protocols such as EAP-TLS or TEAP.

Note

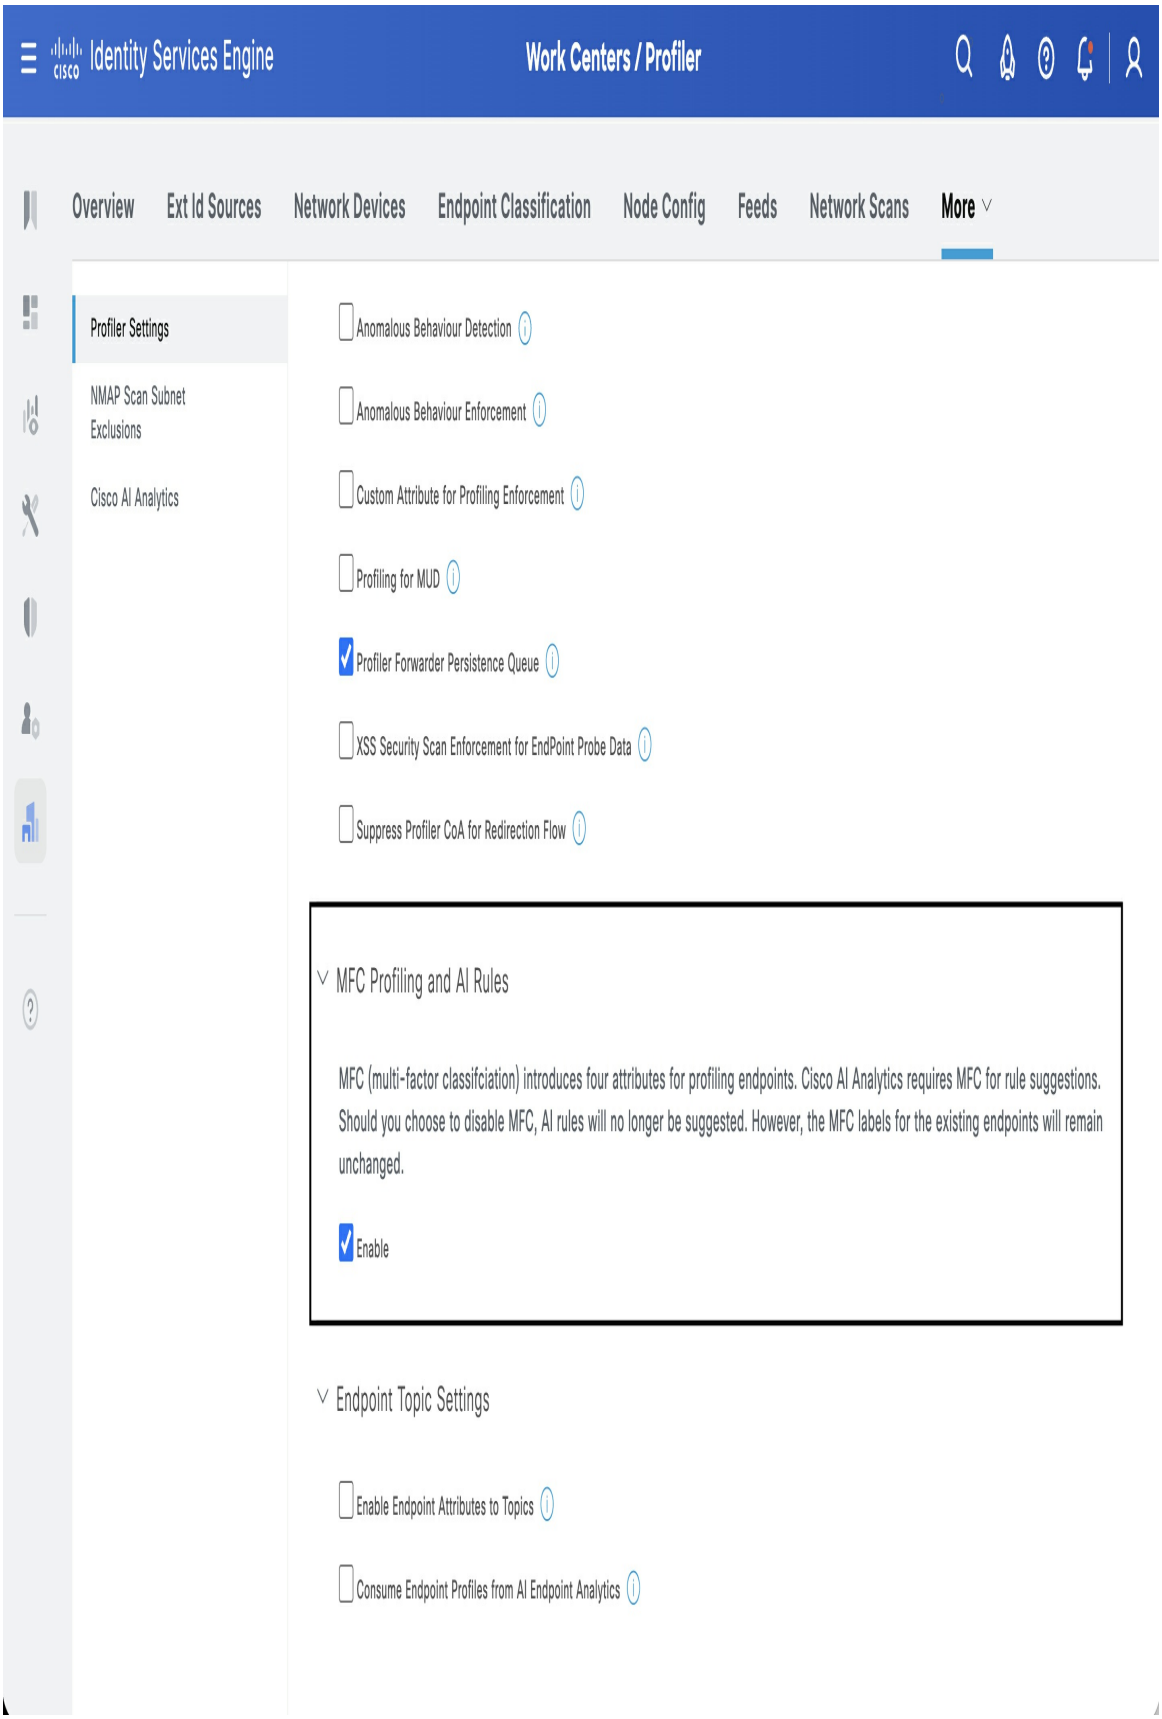
This flow occurs only when an endpoint is profiled on ISE for the first time. If you disconnect the endpoint and reconnect it after completing the process described earlier, it will immediately match the final policy.

Multi-Factor Classification Profiling

Starting with ISE 3.3, multi-factor classification (MFC) was introduced. MFC relies on probe data and allows you to group endpoints based on new endpoint attributes (labels):

- MFC endpoint type
- MFC hardware manufacture
- MFC hardware model
- MFC operating system

To enable MFC profiling, navigate to **Work Centers > Profiler > Settings > Profiler Settings**. Enable the checkbox under **MFC Profiling and AI Rules**, as shown in Figure 8-21.


Identity Services Engine
Work Centers / Profiler
Overview
Ext Id Sources
Network Devices
Endpoint Classification
Node Config
Feeds
Network Scans
More
Profiler Settings
NMAP Scan Subnet Exclusions
Cisco AI Analytics
Anomalous Behaviour Detection
Anomalous Behaviour Enforcement
Custom Attribute for Profiling Enforcement
Profiling for MUD
Profiler Forwarder Persistence Queue
XSS Security Scan Enforcement for EndPoint Probe Data
Suppress Profiler CoA for Redirection Flow
MFC Profiling and AI Rules
MFC (multi-factor classifciation) introduces four attributes for profiling endpoints. Cisco AI Analytics requires MFC for rule suggestions. Should you choose to disable MFC, AI rules will no longer be suggested. However, the MFC labels for the existing endpoints will remain unchanged.
Enable
Endpoint Topic Settings
Enable Endpoint Attributes to Topics
Consume Endpoint Profiles from AI Endpoint Analytics

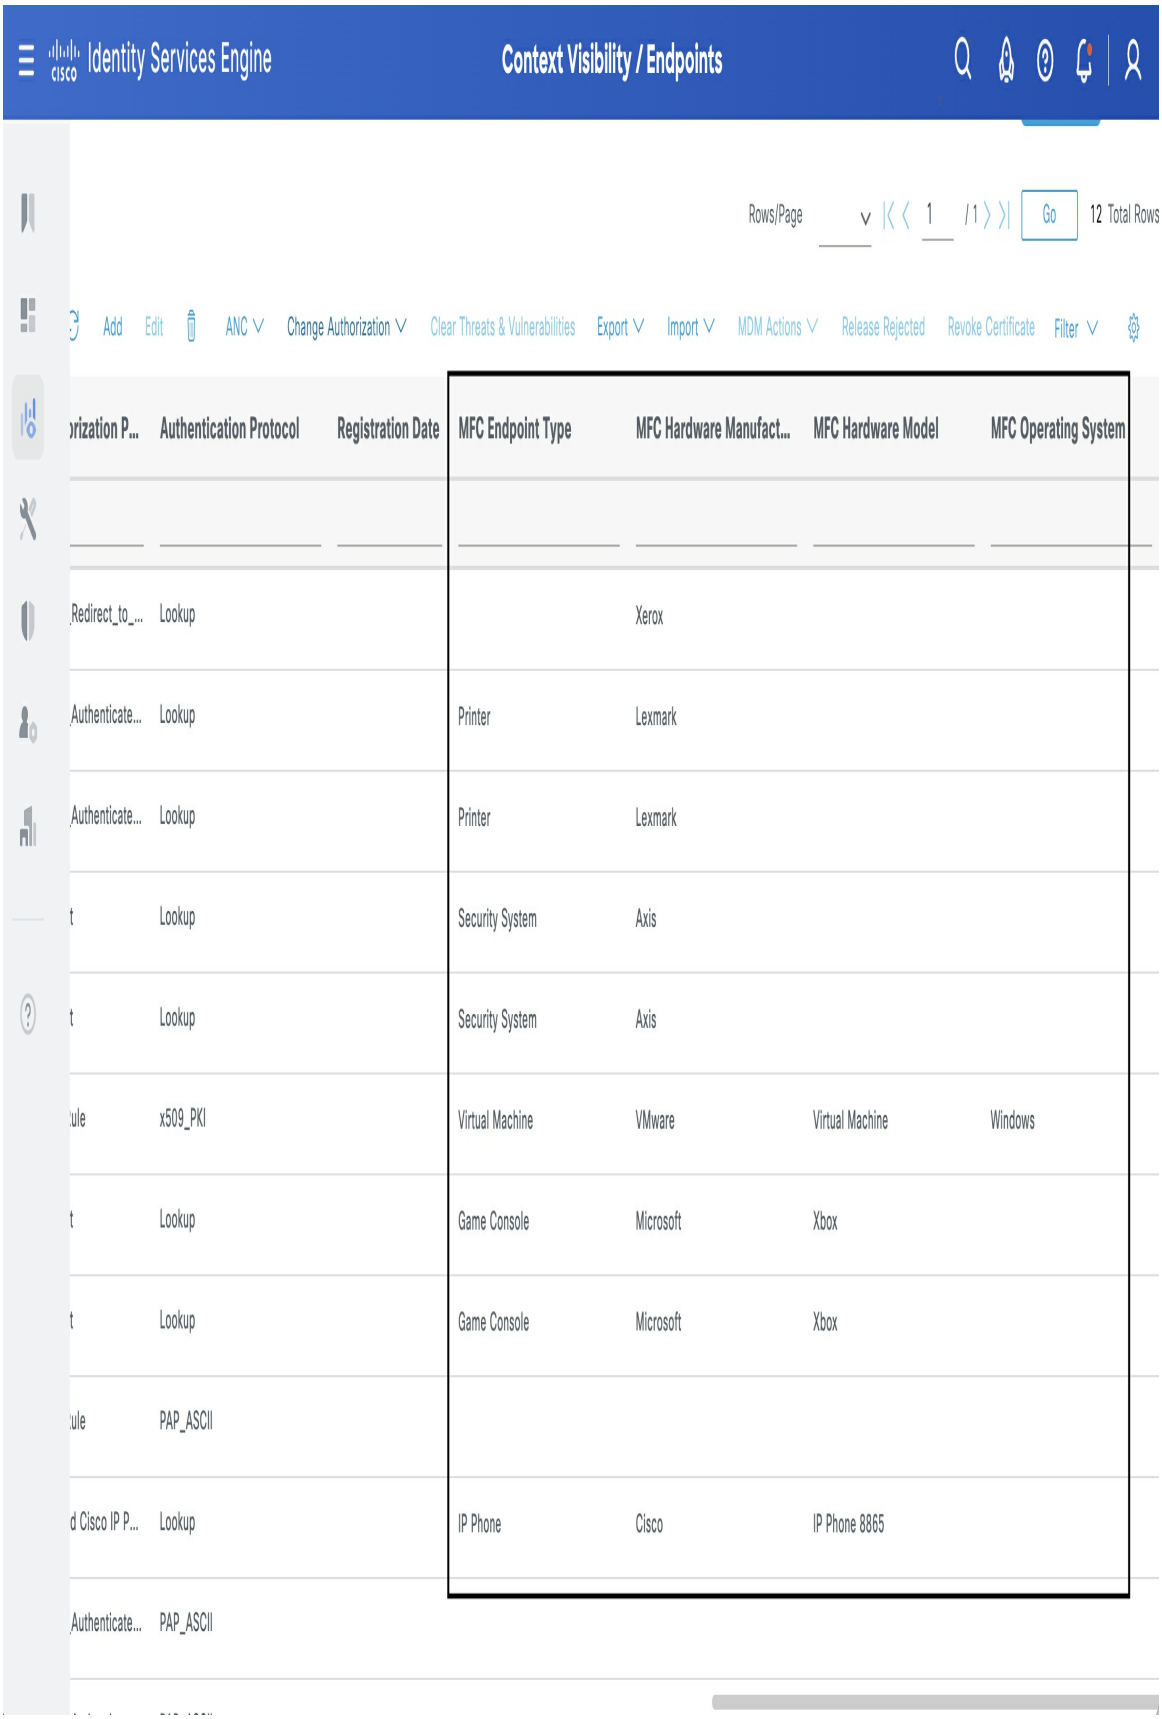
Technet24

Figure 8-21 *MFC Profiling Configuration*

Once MFC Profiling is enabled, ISE begins assigning MFC labels to endpoints based on profiling data. You can view the assigned MFC attributes from **Context Visibility > Endpoints**, as shown in Figure 8-22.

Note

The terms *attributes* and *labels* are used interchangeably and share the same definition.


Identity Services Engine
Context Visibility / Endpoints
Rows/Page
1
/1
Go
12 Total Rows
Add
Edit
ANC
Change Authorization
Clear Threats & Vulnerabilities
Export
Import
MDM Actions
Release Rejected
Revoke Certificate
Filter
orization P...
Authentication Protocol
Registration Date
MFC Endpoint Type
MFC Hardware Manufact...
MFC Hardware Model
MFC Operating System
Redirect_to_...
Lookup
Xerox
Authenticate...
Lookup
Printer
Lexmark
Authenticate...
Lookup
Printer
Lexmark
t
Lookup
Security System
Axis
t
Lookup
Security System
Axis
ule
x509_PKI
Virtual Machine
VMware
Virtual Machine
Windows
t
Lookup
Game Console
Microsoft
Xbox
t
Lookup
Game Console
Microsoft
Xbox
ule
PAP_ASCII
d Cisco IP P...
Lookup
IP Phone
Cisco
IP Phone 8865
Authenticate...
PAP_ASCII

Technet24

Figure 8-22 *MFC Attributes in Context Visibility*

When building authorization policies, you can rely not only on a single endpoint profile but also on an MFC attribute or a set of attributes. You can design your policy as a combination of certainty factor–based profiling policies and MFC attributes, as shown in Figure 8-23.


Authorization Policy(20)
Results
Status
Rule Name
Conditions
Profiles
Security Groups
Hits
Actions
Search
Access -Point Access
AND
EndPoints:EndPointPolicy EQUALS Cisco-Device:Cisco-Wireless-Access-Point:Cisco-Wireless-AP-CW9179F
AND
EndPoints:MFCInfoHardwareManufacturer CONTAINS Cisco
EndPoints:MFCInfoEndpointType CONTAINS Access-Point
PermitAccess
Select from list

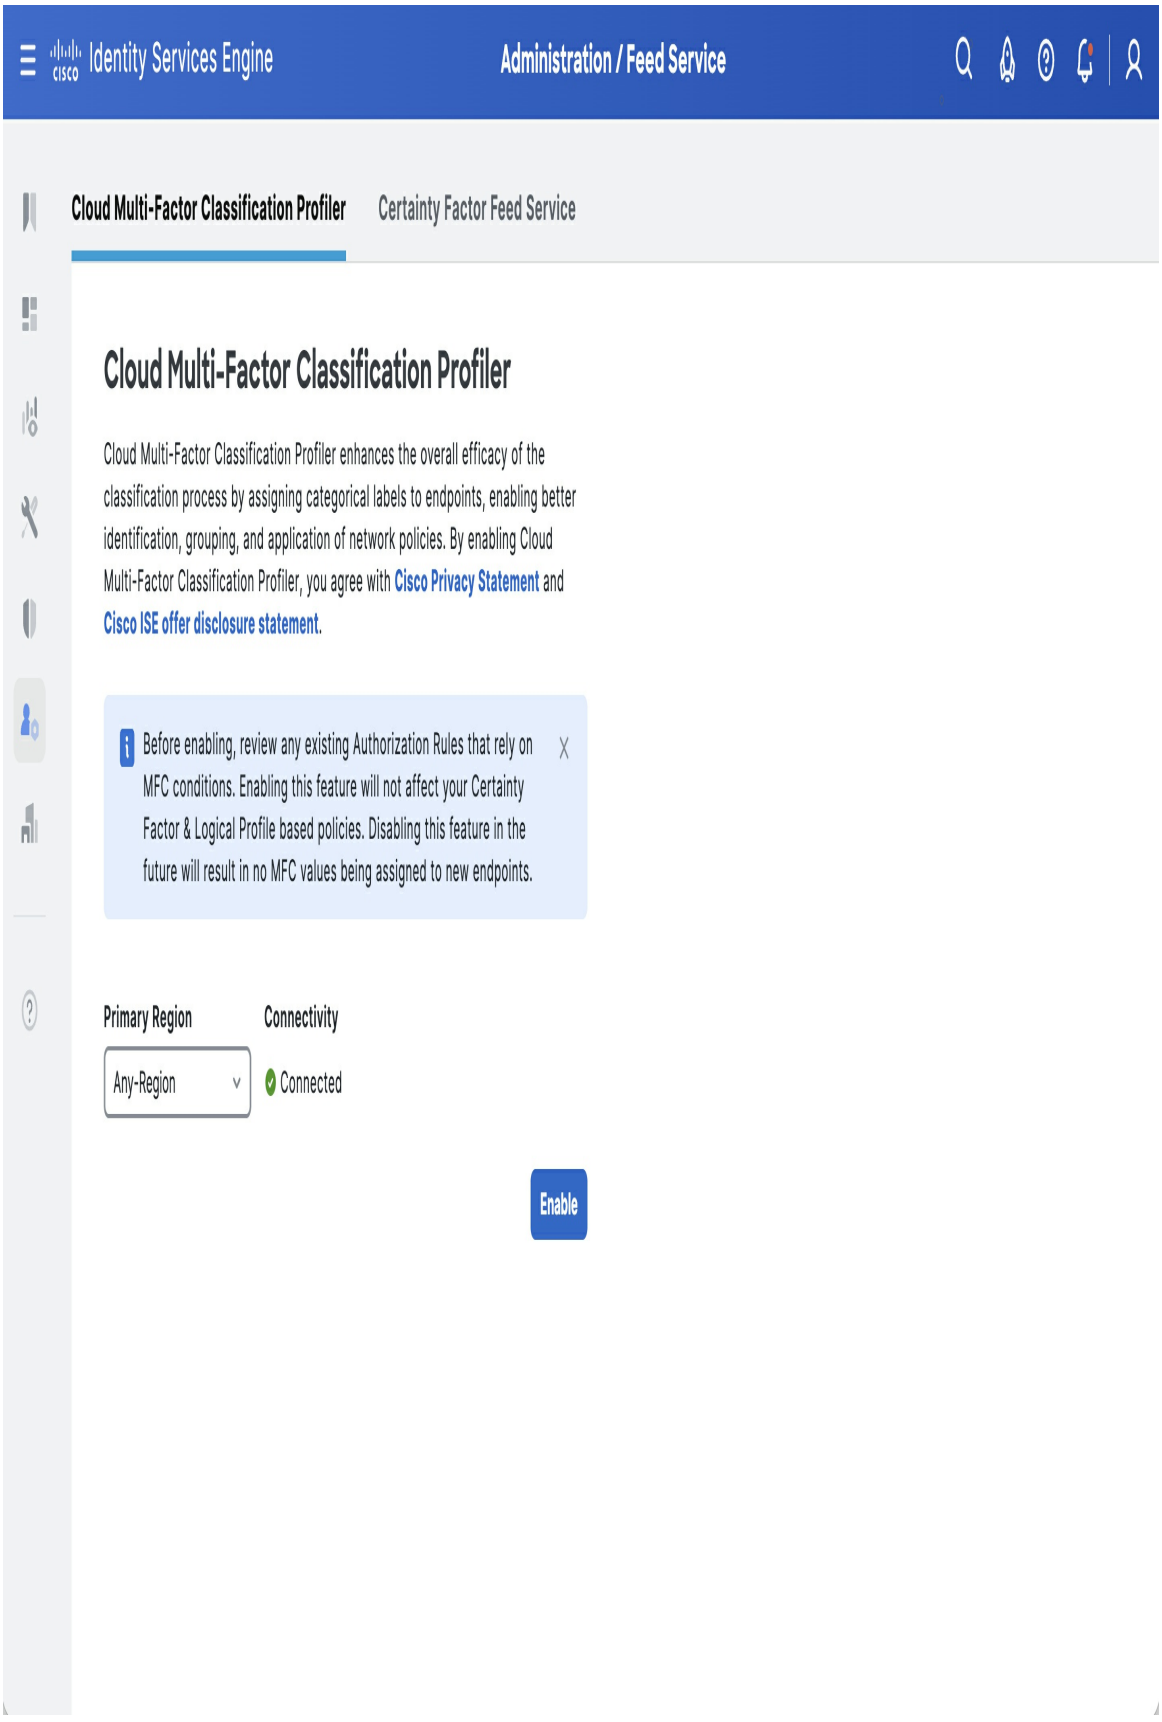
Technet24

Figure 8-23 *MFC Attributes Used in Authorization Policy*

Combining certainty factor–based profiling policies with MFC attributes or using multiple MFC attributes within a single policy provides greater flexibility for rule creation.

> **Note**
>
> Administrators have no influence over system rules in ISE that are responsible for assigning MFC attributes.

When multi-factor classification is enabled, ISE sends profiling data (telemetry) to the cloud and receives MFC profiling attributes assigned to endpoints. To enable the Cloud MFC Profiler, navigate to **Administration > Feed Service > Cloud Multi-Factor Classification Profiler**, as shown in Figure 8-24. Select your preferred region and click **Enable**.


Identity Services Engine
Administration / Feed Service
Cloud Multi-Factor Classification Profiler
Certainty Factor Feed Service
Cloud Multi-Factor Classification Profiler
Cloud Multi-Factor Classification Profiler enhances the overall efficacy of the classification process by assigning categorical labels to endpoints, enabling better identification, grouping, and application of network policies. By enabling Cloud Multi-Factor Classification Profiler, you agree with Cisco Privacy Statement and Cisco ISE offer disclosure statement.
Before enabling, review any existing Authorization Rules that rely on MFC conditions. Enabling this feature will not affect your Certainty Factor & Logical Profile based policies. Disabling this feature in the future will result in no MFC values being assigned to new endpoints.
Primary Region
Any-Region
Connectivity
Connected
Enable

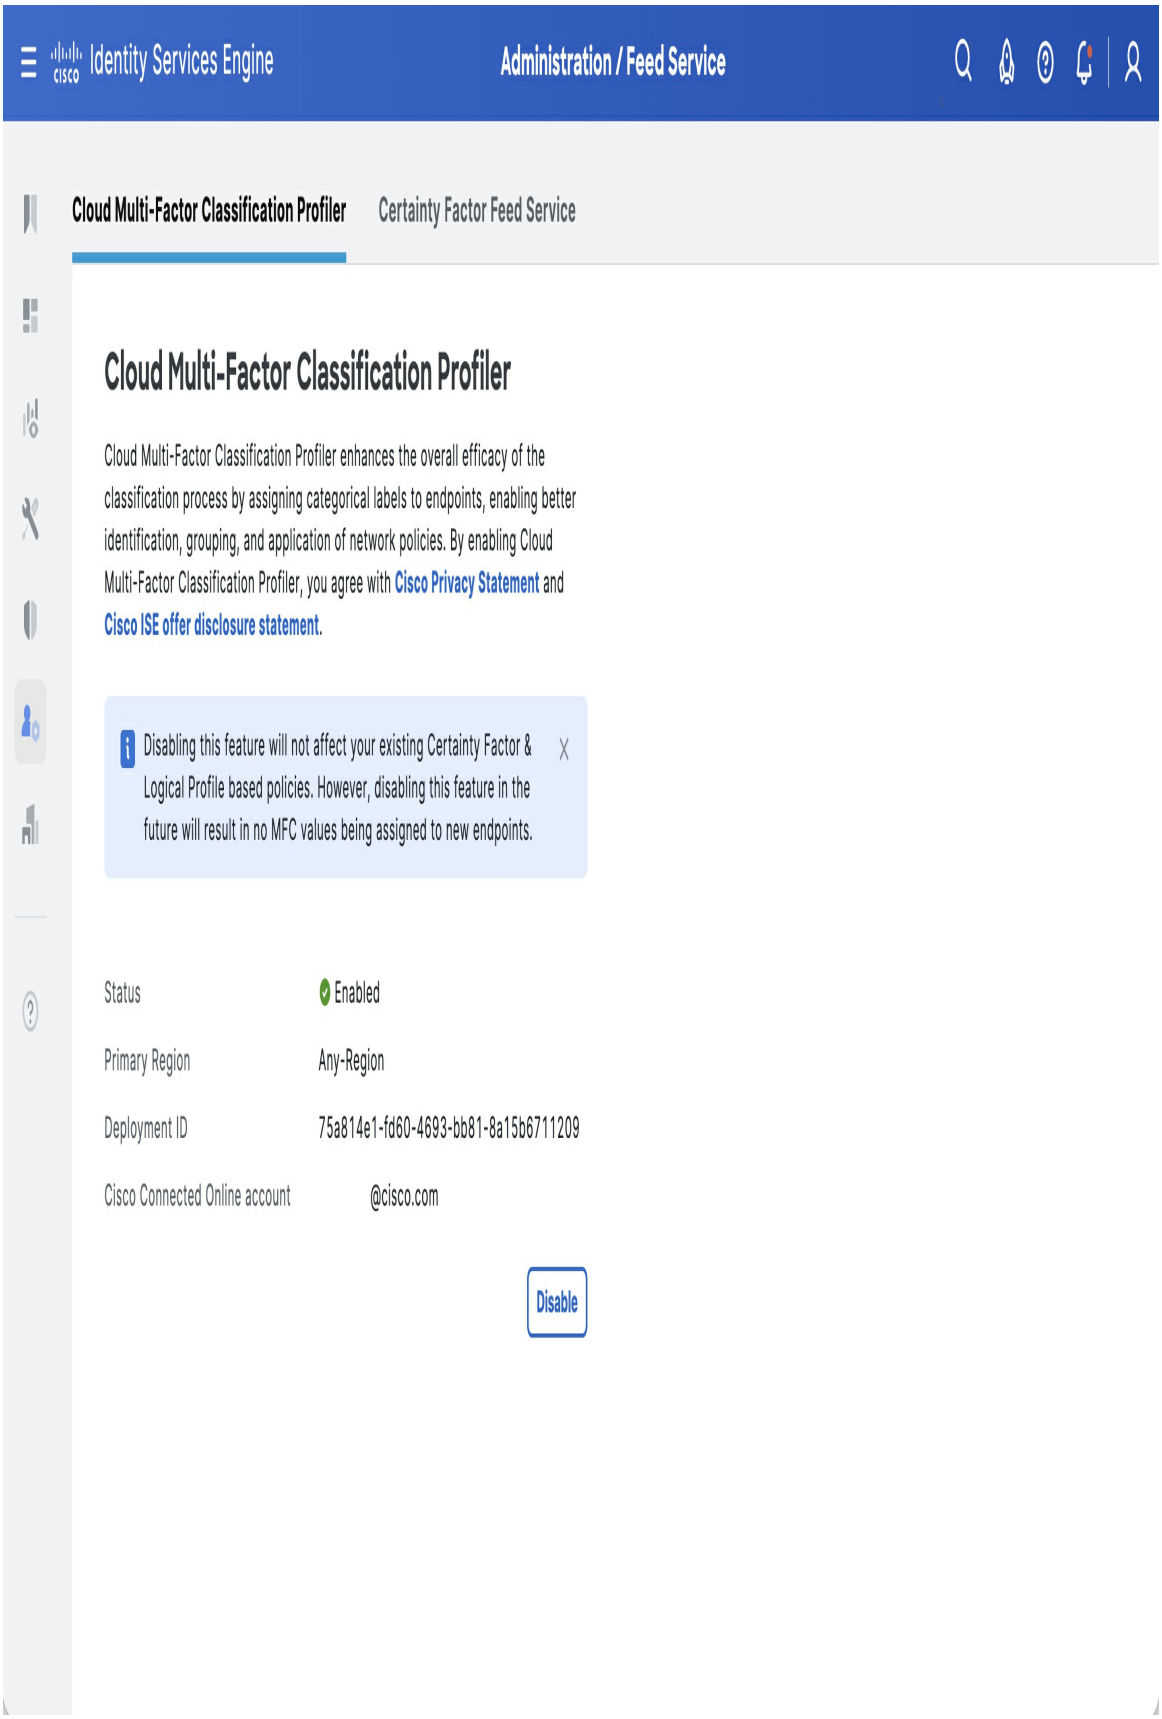
Technet24

Figure 8-24 *Cloud MFC Profiler Configuration*

Note

When you enable the Cloud MFC Profiler, ISE will reclassify all existing active endpoints.

Enabling the feature will guide you through authorization with your Cisco Connected Online (CCO) account. After you provide your credentials, the successful registration screen should appear as shown in Figure 8-25.


Identity Services Engine
Administration / Feed Service
Cloud Multi-Factor Classification Profiler
Certainty Factor Feed Service
Cloud Multi-Factor Classification Profiler
Cloud Multi-Factor Classification Profiler enhances the overall efficacy of the classification process by assigning categorical labels to endpoints, enabling better identification, grouping, and application of network policies. By enabling Cloud Multi-Factor Classification Profiler, you agree with Cisco Privacy Statement and Cisco ISE offer disclosure statement.
Disabling this feature will not affect your existing Certainty Factor & Logical Profile based policies. However, disabling this feature in the future will result in no MFC values being assigned to new endpoints.
Status
Enabled
Primary Region
Any-Region
Deployment ID
75a814e1-fd60-4693-bb81-8a15b6711209
Cisco Connected Online account
@cisco.com
Disable

Technet24

Figure 8-25 *Enabled Cloud MFC Profiler Configuration*

> **Note**
>
> If you disable the Cloud MFC Profiler, no new labels will be assigned, even by system rules.

Cloud profiling introduced the concept of multi-factor classification profiling policies. Certainty factor–based profiling can be challenging because determining the appropriate certainty factor for custom profiles often requires trial and error to ensure it fits within the overall framework.

The new MFC profiling policy relies on simple If-Then conditions, making it easier to configure and troubleshoot.

You can configure two types of policies:

- Custom policies
- Direct mapping policies

Direct mapping policies are basic in their configuration. They allow you to map different attributes directly to an MFC label. To configure them, navigate to **Policy > Profiling > Multi-Factor Classification** and select **+ Add Profiling Policy**.

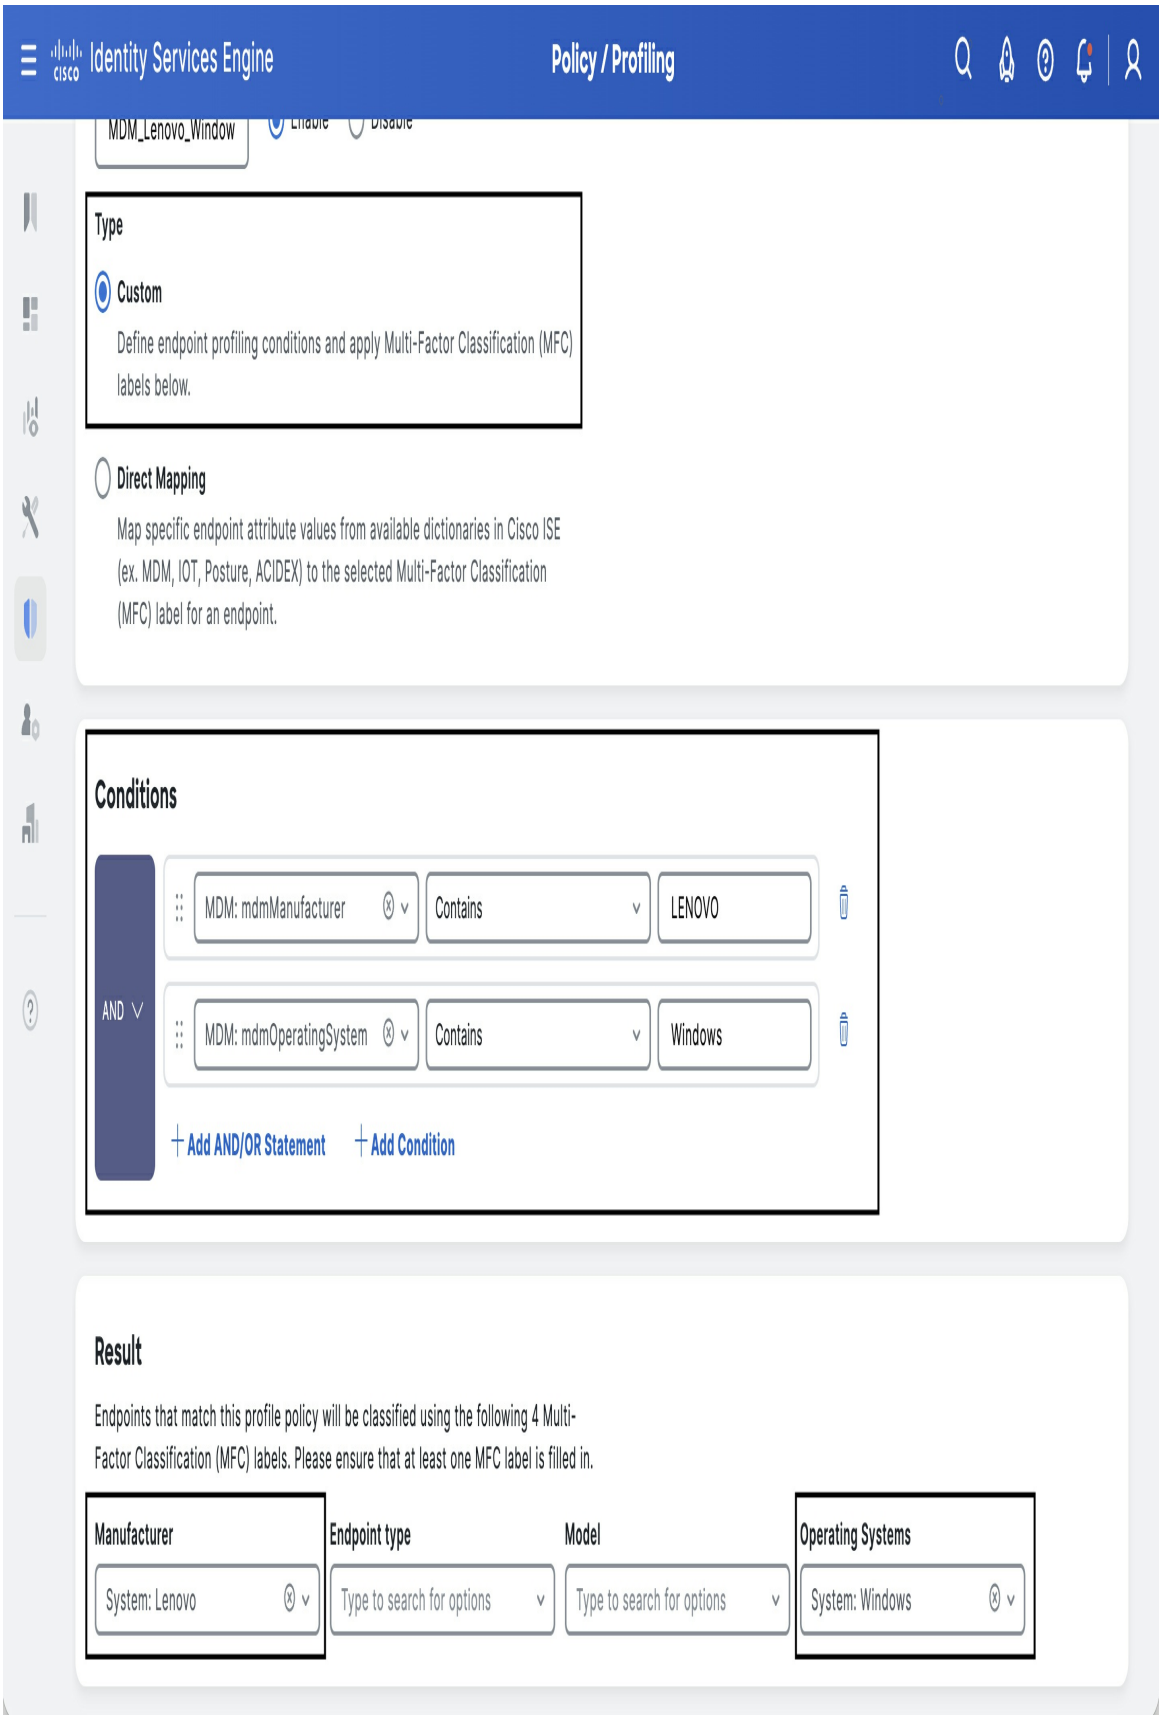
Figure 8-26 shows the configuration of a direct mapping rule within a Cloud MFC profiling policy, where the IOTASSET attribute iotAssetDeviceType is mapped directly to the endpoint type MFC label.


Identity Services Engine
Policy / Profiling
Add Profiling Policy
Details
Policy name
iotAssetDeviceType_D
Status
Enable
Disable
Type
Custom
Define endpoint profiling conditions and apply Multi-Factor Classification (MFC) labels below.
Direct Mapping
Map specific endpoint attribute values from available dictionaries in Cisco ISE (ex. MDM, IOT, Posture, ACIDEX) to the selected Multi-Factor Classification (MFC) label for an endpoint.
Mapping
Choose the authoritative sources and attributes you want to directly use as MFC (Multi-Factor Classification). The endpoints matching this profile policy will be profiled directly with the attributes from the source of your selection. Please fill at least 1 mapping.
Select attribute
IOTASSET: iotAssetDeviceTy...
map to MFC label
Endpoint type
+ Add mapping
Cancel
Continue

Technet24

Figure 8-26 *MFC Direct Mapping Policy*

Custom mapping policies allow you to assign specific values to MFC attributes based on a set of conditions. The example in Figure 8-27 shows how you can assign values for manufacturer and operating system based on MDM attributes from the endpoint.


Identity Services Engine
Policy / Profiling
MDM_Lenovo_Window
Enable
Disable
Type
Custom
Define endpoint profiling conditions and apply Multi-Factor Classification (MFC) labels below.
Direct Mapping
Map specific endpoint attribute values from available dictionaries in Cisco ISE (ex. MDM, IOT, Posture, ACIDEX) to the selected Multi-Factor Classification (MFC) label for an endpoint.
Conditions
AND
MDM: mdmManufacturer
Contains
LENOVO
MDM: mdmOperatingSystem
Contains
Windows
Add AND/OR Statement
Add Condition
Result
Endpoints that match this profile policy will be classified using the following 4 Multi-Factor Classification (MFC) labels. Please ensure that at least one MFC label is filled in.
Manufacturer
System: Lenovo
Endpoint type
Type to search for options
Model
Type to search for options
Operating Systems
System: Windows

Technet24

Figure 8-27 *MFC Custom Policy*

You are not limited to MDM attributes; you can use any combination of the following attributes in the conditions.

- ACIDEX
- ACTIVDIRECTORY_PROBE
- CDP
- DHCP
- IOTASSET
- IP
- LLDP
- MAC
- MDM
- MUD
- Multimedia
- NMAP
- Posture
- RADIUS
- SNMP
- Wifi_Device_Analytics

To access configured MFC policies, navigate to **Policy > Profiling > Multi-Factor Classification**, as shown in Figure 8-28.

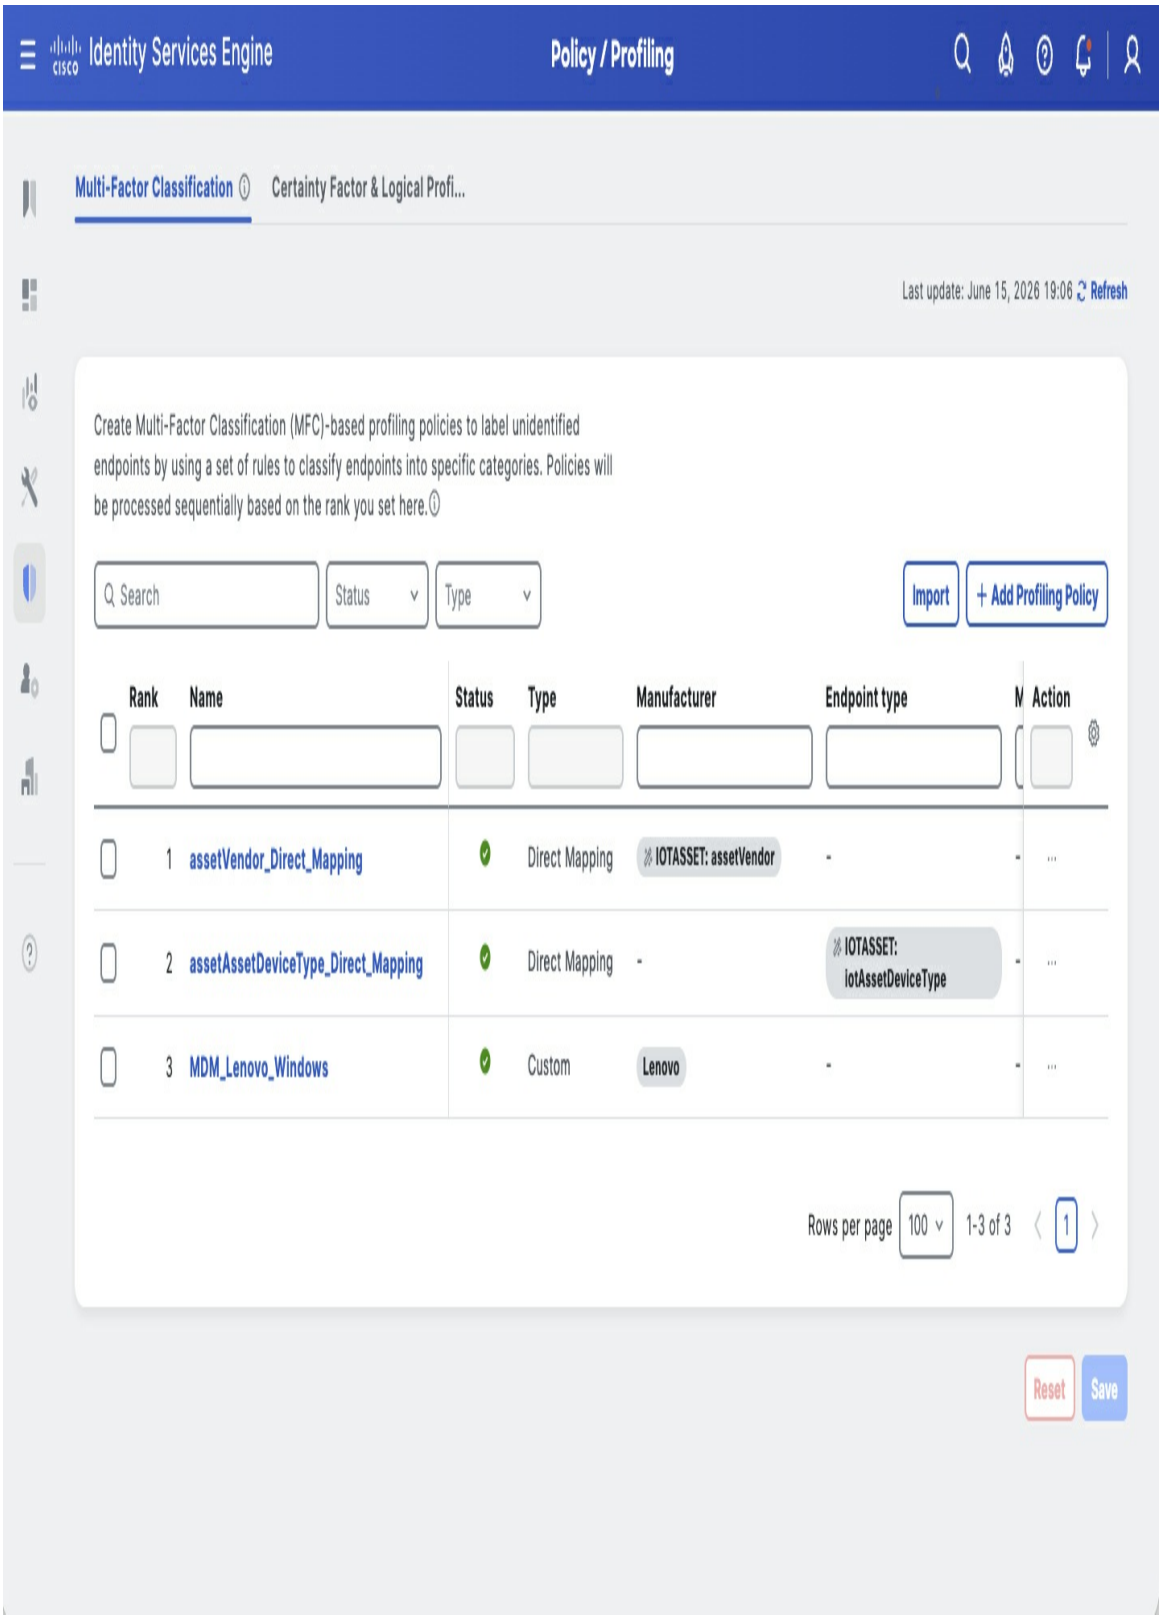


Figure 8-28 *MFC Profiling Policies*

Technet24

It is crucial to understand the order of operations and the impact of enabling the Cloud MFC Profiler.

If the Cloud MFC Profiler was never enabled in the ISE deployment, custom and direct mapping policies are responsible for MFC assignment. The priority of MFC attributes in the result is determined by the rule order; the higher the rule ranks, the more priority it has in case of conflict. System rules populate MFC labels that remain empty, provided there is enough profiling data to do so.

If the Cloud MFC Profiler was ever enabled, three types of rules are executed in parallel: custom rules, direct mapping rules, and cloud rules (responses from the cloud). Results are merged from all three sources. In case of conflict, custom and direct mapping rules are evaluated first based on their rank, and then cloud rules populate the remaining attributes.

Additional Profiler Settings

One of the additional settings is the Endpoint Attribute Filter. You can configure it under **Administration > Work Centers > Profiler > Settings > Profiler Settings**, as shown in Figure 8-29.

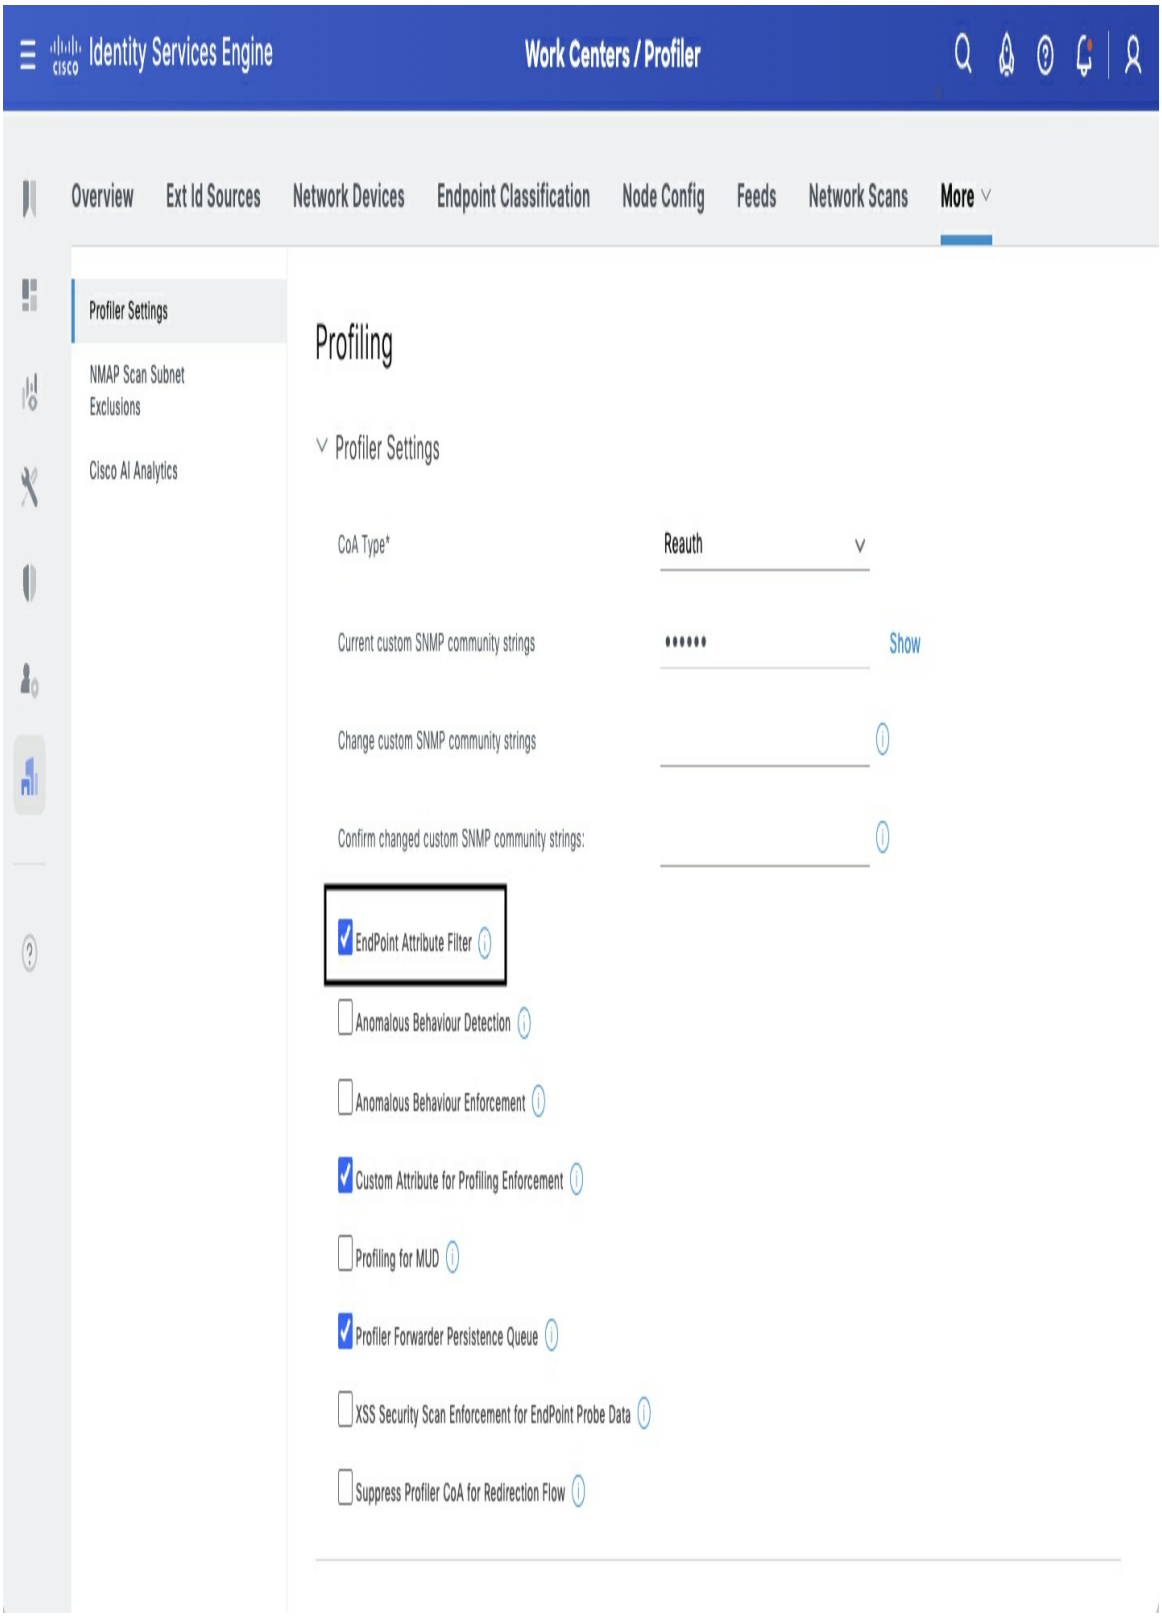


Figure 8-29 *Endpoint Attribute Filter*

Technet24

The Endpoint Attribute Filter reduces endpoint data collection and replication to a specific subset of attributes. When enabled, it uses only the attributes required to support default (Cisco-provided) profiling policies, critical RADIUS operations, BYOD, and WebAuth. All other attributes are dropped at the time of collection, protecting the profiler component from unnecessary data and enhancing system security.

Best Practice:

Always enable the Endpoint Attribute Filter unless a specific use case requires additional profiling.

You can configure the Profiler Forwarder Persistence Queue under **Administration > Work Centers > Profiler > Settings > Profiler Settings**, as shown in Figure 8-30.

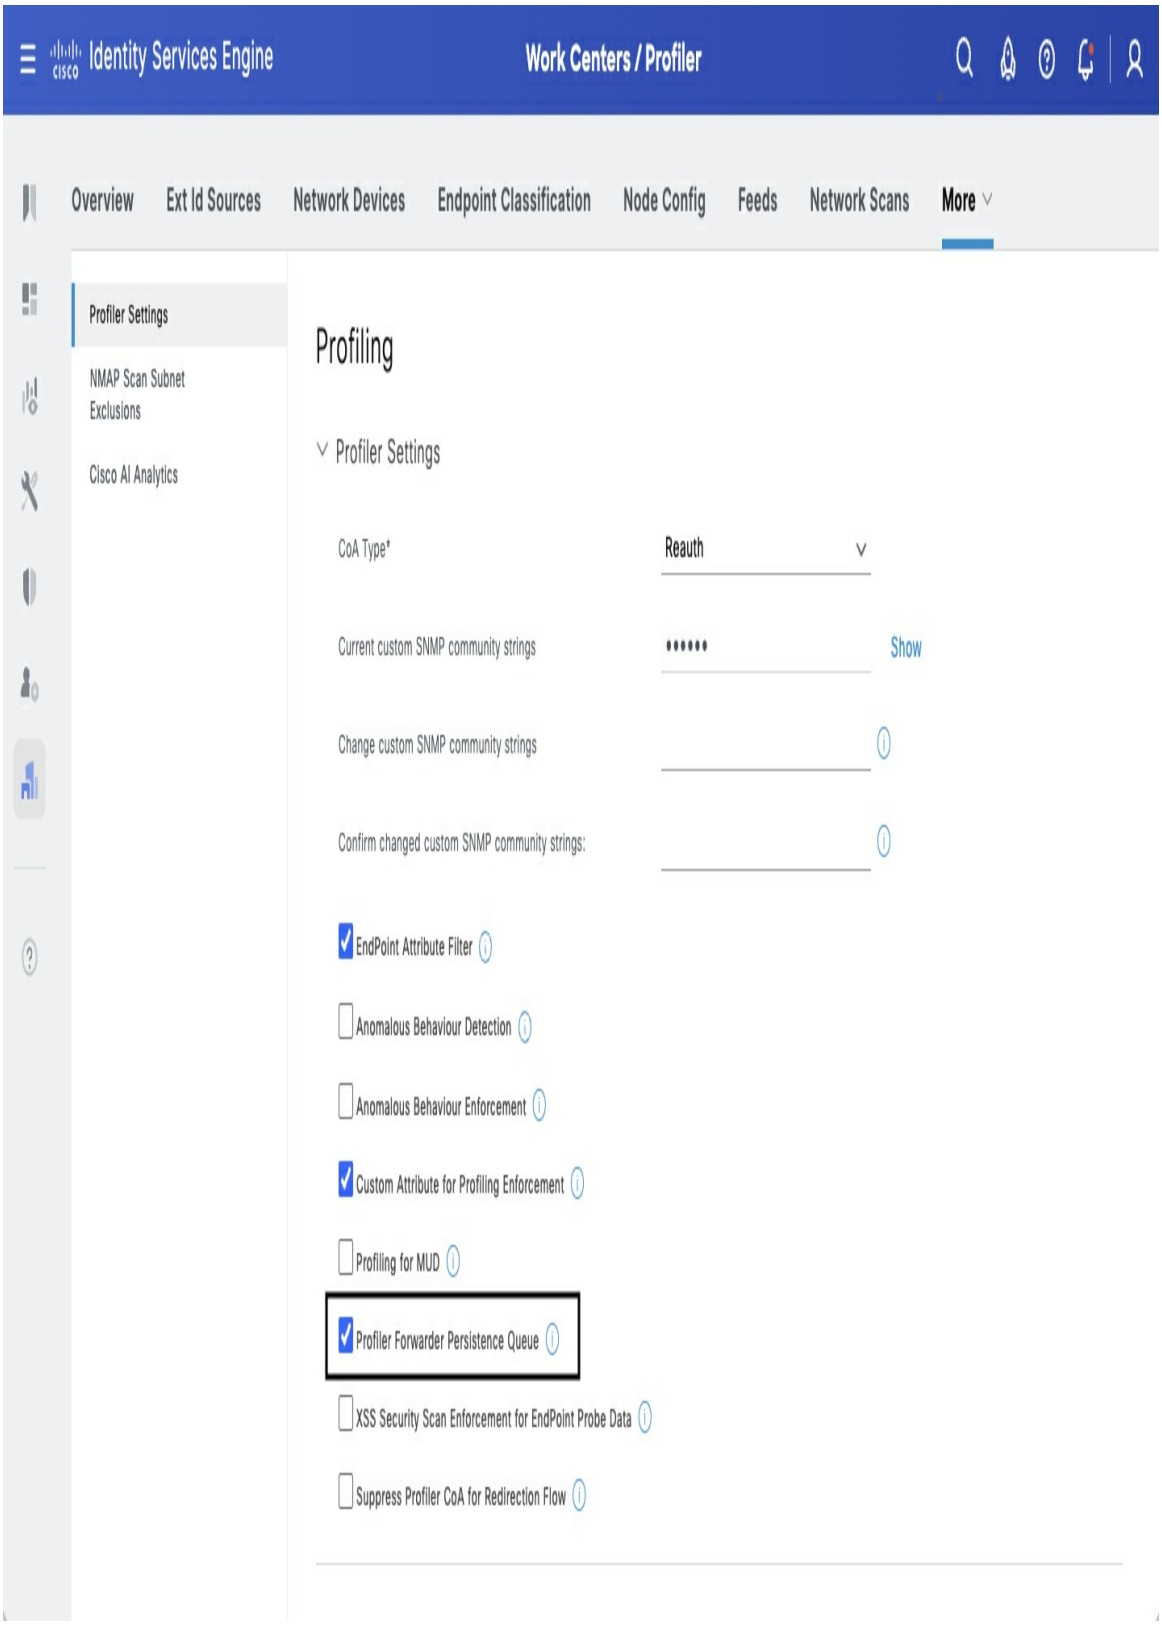


Figure 8-30 *Profiler Forwarder Persistence Queue*

Technet24

Profiler Forwarder Persistence enables an internal queueing mechanism within ISE for profiling data processing. This feature helps manage spikes in profiling events and prevents potential loss of probe data.

Profiler CoA settings can be configured globally, as shown in Figure 8-31. The available options are as follows:

- Reauth
- Port Bounce
- No CoA

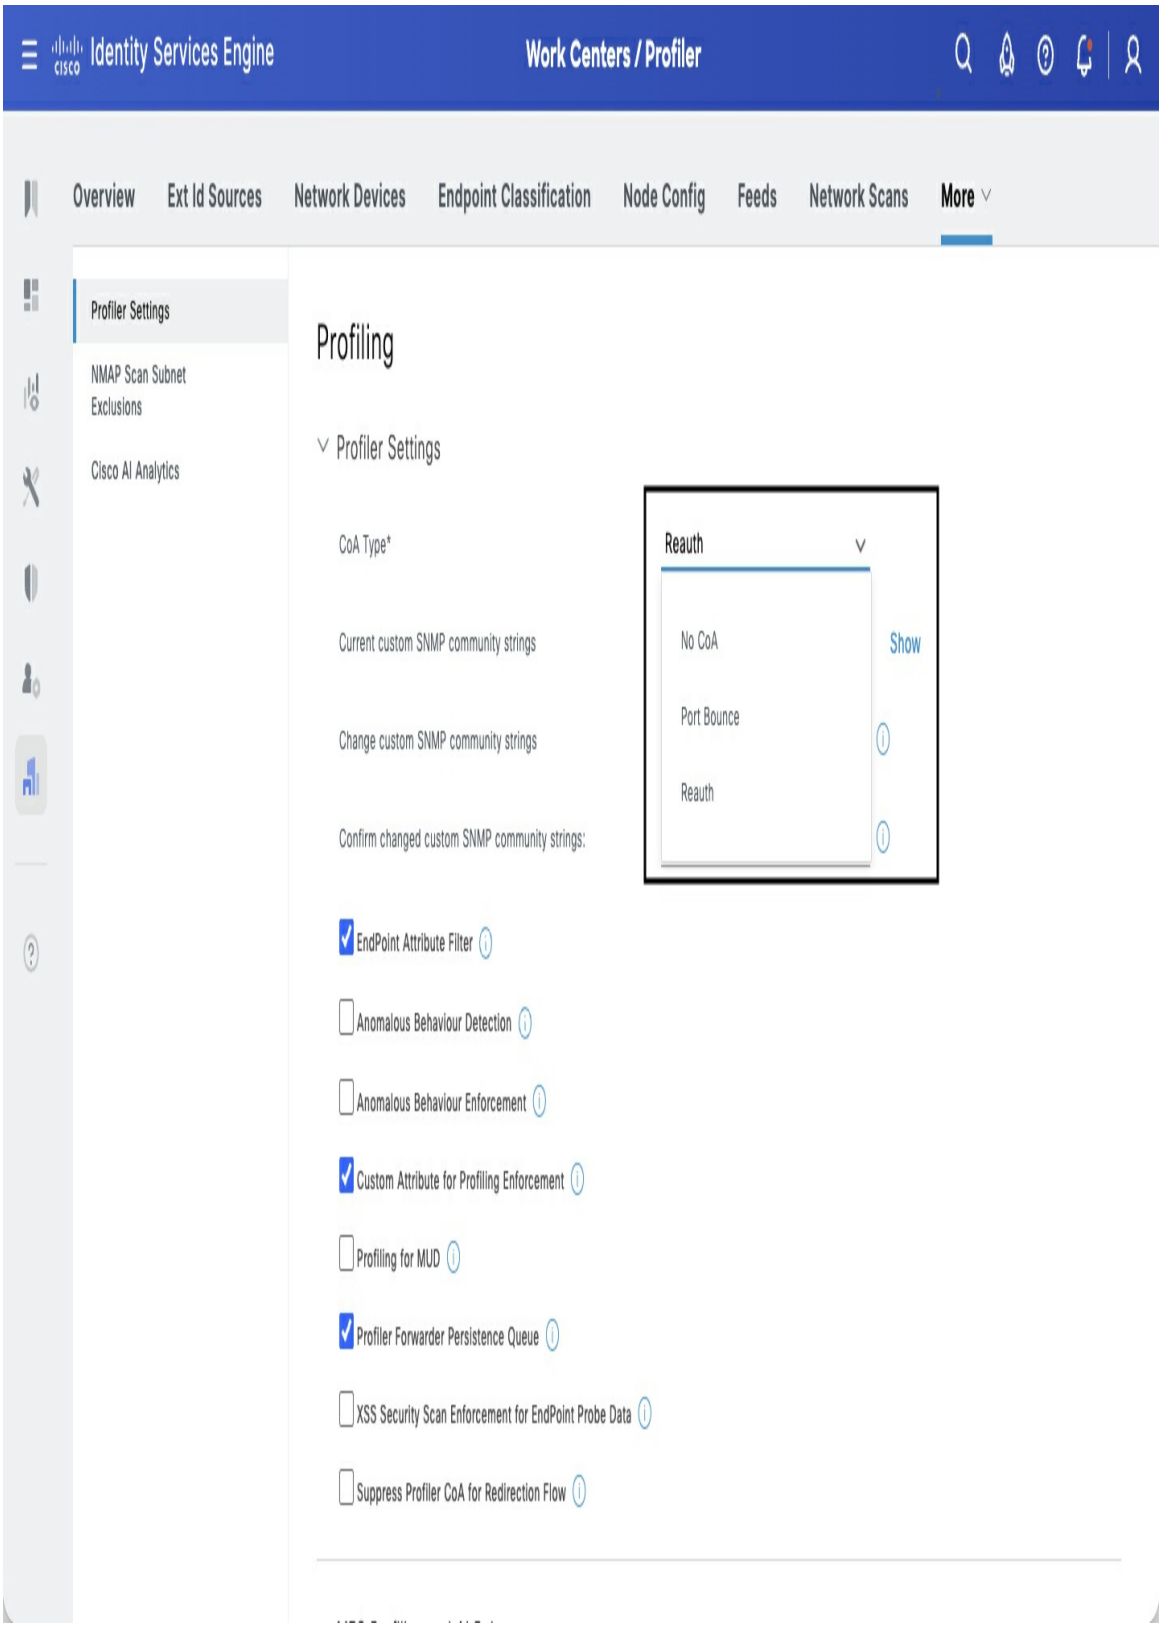


Figure 8-31 *Global Profiler CoA Settings*

Technet24

Once an endpoint is profiled, the CoA decision is made based on the specific profile settings, as shown in Figure 8-32. If the profile is configured to use global settings, it will inherit the options defined globally, as previously described.

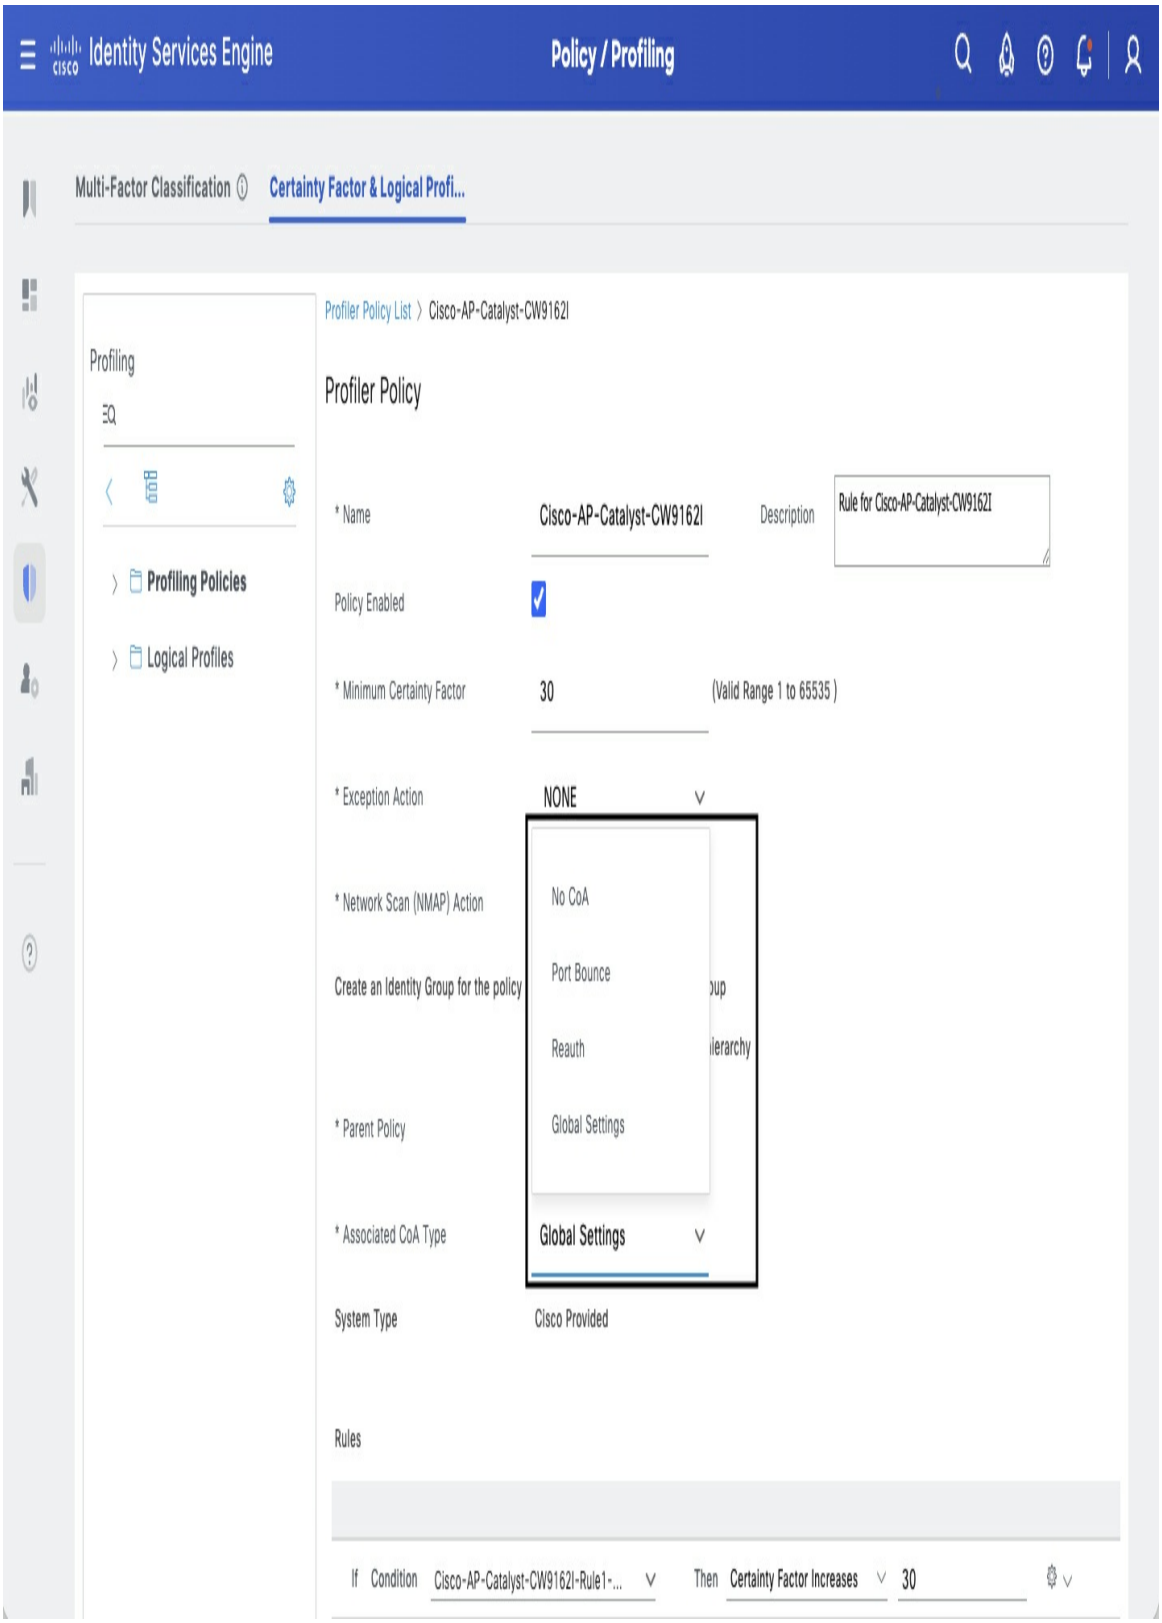


Figure 8-32 *Endpoint Profile CoA Settings*

Technet24

Important:

If Global Profiler Settings is set to No CoA, it overrides all configuration settings in Endpoint Profiler Policies.

Summary

Profiling is a powerful mechanism that adds visibility to your network. You cannot protect what you do not know about, and profiling addresses this problem. In addition to visibility, profiling enables you to enhance your security policies. Certainty factor–based profiling and Cloud MFC profiling provide flexibility in profile and label assignment to suit the requirements of nearly any scenario.

Chapter 9. Managing Guest Users and Personal Devices

Introduction

This chapter focuses on the theory and practice of managing guest users and personal devices with Cisco Identity Services Engine. You are introduced to guest access and BYOD, and learn how to implement guest access flows, including sponsored, employee, self-registration, and hotspot guest access. The text explores approaches to web redirection and policy design, and explains the role of attribute stitching for authorization and policy enforcement. It reviews the configuration of guest access components, such as sponsor users, sponsor groups, guest portals, and guest types. The chapter also details BYOD onboarding, client provisioning, and certificate management using both single and dual SSID flows. Finally, it examines portal configuration for guest and device self-service management.

Guest Access and BYOD

There is probably no organization without visitors. Sooner or later, almost every IT department finds itself in a situation where temporary access must be provided to users who are not company employees. They may include guests, auditors, vendor contractors, and similar roles. ISE is a centralized network access control (NAC) engine that

Technet24

handles this by using Guest Access features that enable administrators to deliver the full lifecycle of guest access management.

The following is an example of the guest lifecycle stages:

- **Create:** A sponsor or guest initiates an account within administrator-defined rules.
- **Deliver:** Credentials are sent via short message service (SMS), email, or printout.
- **Authenticate:** A guest logs in through the portal to gain network access.
- **Expire:** The account automatically deactivates after its validity period.
- **Audit:** Every action, from creation to usage, is logged for compliance.

Another common challenge that IT departments encounter is the need to securely manage employees' personal devices. If the IT team completely ignores personal devices, the eventual result is "shadow IT" appearing on the company's premises. Cisco ISE addresses this concern by providing comprehensive network access control for employees' or company-owned mobile devices using the built-in bring your own device (BYOD) mechanism.

The BYOD implementation on Cisco ISE has the following characteristics:

- **Multiple device support:** Cisco ISE enables users to register and manage several personal devices, such as phones, tablets, and laptops. Each device follows the same policy logic, ensuring a consistent and secure experience across platforms like iOS, macOS, Android, Windows, and Chromebook.

- **Multiple network topologies:** Cisco ISE enforces consistent policies across wired, wireless, and VPN connections. Regardless of access method, authorization decisions are based on the same identity and device trust framework.
- **Self-registration and self-service model:** Users can register their devices without IT assistance through a guided onboarding portal and later manage them through the My Devices portal.
- **Certificate provisioning:** Cisco ISE automatically handles certificate enrollment during onboarding, issuing a digital certificate that uniquely identifies the device. This allows secure, password-free reauthentication and maintains persistent trust across future sessions. A Cisco ISE internal certificate authority (CA) can be used to provision certificates to BYOD devices.
- **Blocking and reinstating devices:** Users can suspend a lost or stolen device and later reinstate it through the self-service My Devices portal. This capability prevents unauthorized use and reduces IT involvement in routine device recovery scenarios.

Guest Access Components

Now let's examine the guest service building blocks in ISE. You can consider guest services as an advanced ISE feature alongside capabilities such as posture and TrustSec. Guest access configuration requires working with multiple components and paying close attention to a variety of parameters critical for successful deployment. The most important guest access elements are as follows.

- **The guest flow types:** The type determines how guest access is processed and how the guest lifecycle is managed. The flows include Sponsored Guest Flow, Employees as Guests, Self-Registration, and Hotspot Access.

Technet24

- **Guest redirection:** Redirection is a critical aspect of the guest service. You can perform redirection in several ways, such as through Classic Cisco Central Web Authentication (CWA), DHCP-based redirection, and Static Portal URL. When you're planning guest service implementation, it is essential to understand how each redirection method operates.
- **Attribute stitching:** Any guest flow begins when the user encounters the guest portal. At this stage, ISE typically has limited information about the endpoint or the user. Once the user passes through the guest portal, ISE gathers additional details about both the user and the endpoint. Attribute stitching enables ISE to provide granular network access to guest users and endpoints after authentication on the guest portal.
- **Sponsor users and portals:** When designing guest services, you need to determine who can create guest users, what types of guest users can be created, and where those users can be generated.
- **Guest users and portals:** Ultimately, the created guest users of a specific type gain network access after being redirected to the appropriate guest portal.

You can configure the majority of the settings required to deploy guest services in the corresponding work center, which you access by navigating to **Work Centers > Guest Access**.

Types of Guest Flows

With Cisco ISE, you can deliver guest services using different approaches, often referred to as flows. These flows provide you with additional flexibility in implementing guest access. The following subsections explain each flow in detail.

Sponsored Guest Flow

In Sponsored guest flow, a sponsor user creates guest user accounts. After the accounts are established, the sponsor can share guest access credentials with the end user using various methods, such as print, email, or SMS. Figure 9-1 shows an example of the sponsored guest flow.

Technet24

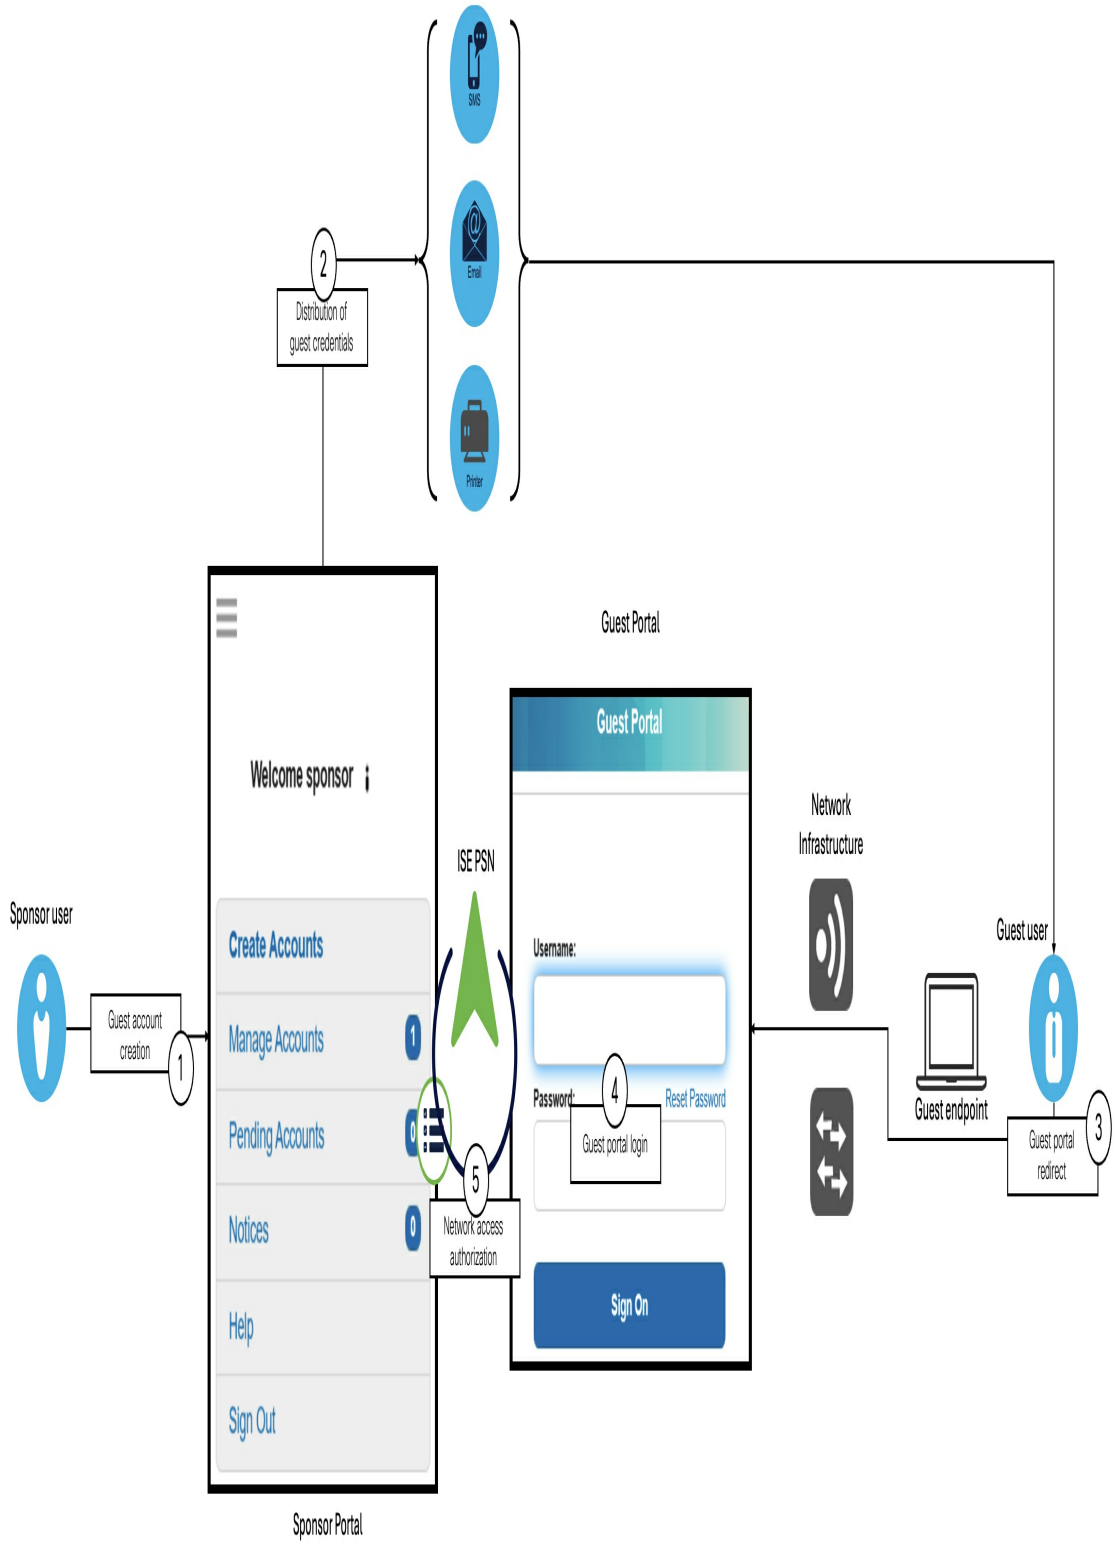

SMS
Email
Printer
2
Distribution of guest credentials
Guest Portal
Welcome sponsor
Create Accounts
Manage Accounts
1
Pending Accounts
0
Notices
0
Help
Sign Out
Sponsor Portal
Sponsor user
Guest account creation
1
ISE PSN
5
Network access authorization
Guest Portal
Username:
Password:
Reset Password
4
Guest portal login
Sign On
Network Infrastructure
Guest endpoint
Guest user
Guest portal redirect
3

Figure 9-1 *Sponsored Guest Flow*

The main steps of the sponsored guest flow are as follows:

Step 1. Sponsor users access the sponsor portal and create the guest account.

Step 2. The guest account credentials are communicated to the guest user. You can deliver credentials in several ways, such as a simple printed copy or more advanced methods like sending credentials over email or SMS. The options for credential delivery are configurable in the Sponsor group settings.

Step 3. The guest user is redirected to the guest portal as part of the initial network authentication. The different redirect options will be discussed later in this chapter.

Step 4. The guest user enters the credentials provided by the sponsor.

Step 5. ISE performs authorization policy evaluation to determine the next policy and triggers a Change of Authorization (CoA) back to the network access device, which results in a new authentication. At this stage, ISE grants the network access defined for this specific guest use case based on session and endpoint attributes. The endpoint and session attributes, as well as session stitching, will be described later in the chapter.

For guest user notifications, besides completing the necessary settings at the sponsor group level (which will be explained later in the chapter), you need to enable the notification connections themselves.

To set up both email and SMS notification settings, navigate to **Work Centers > Guest Access > Administration**. Choose **SMTP Server**

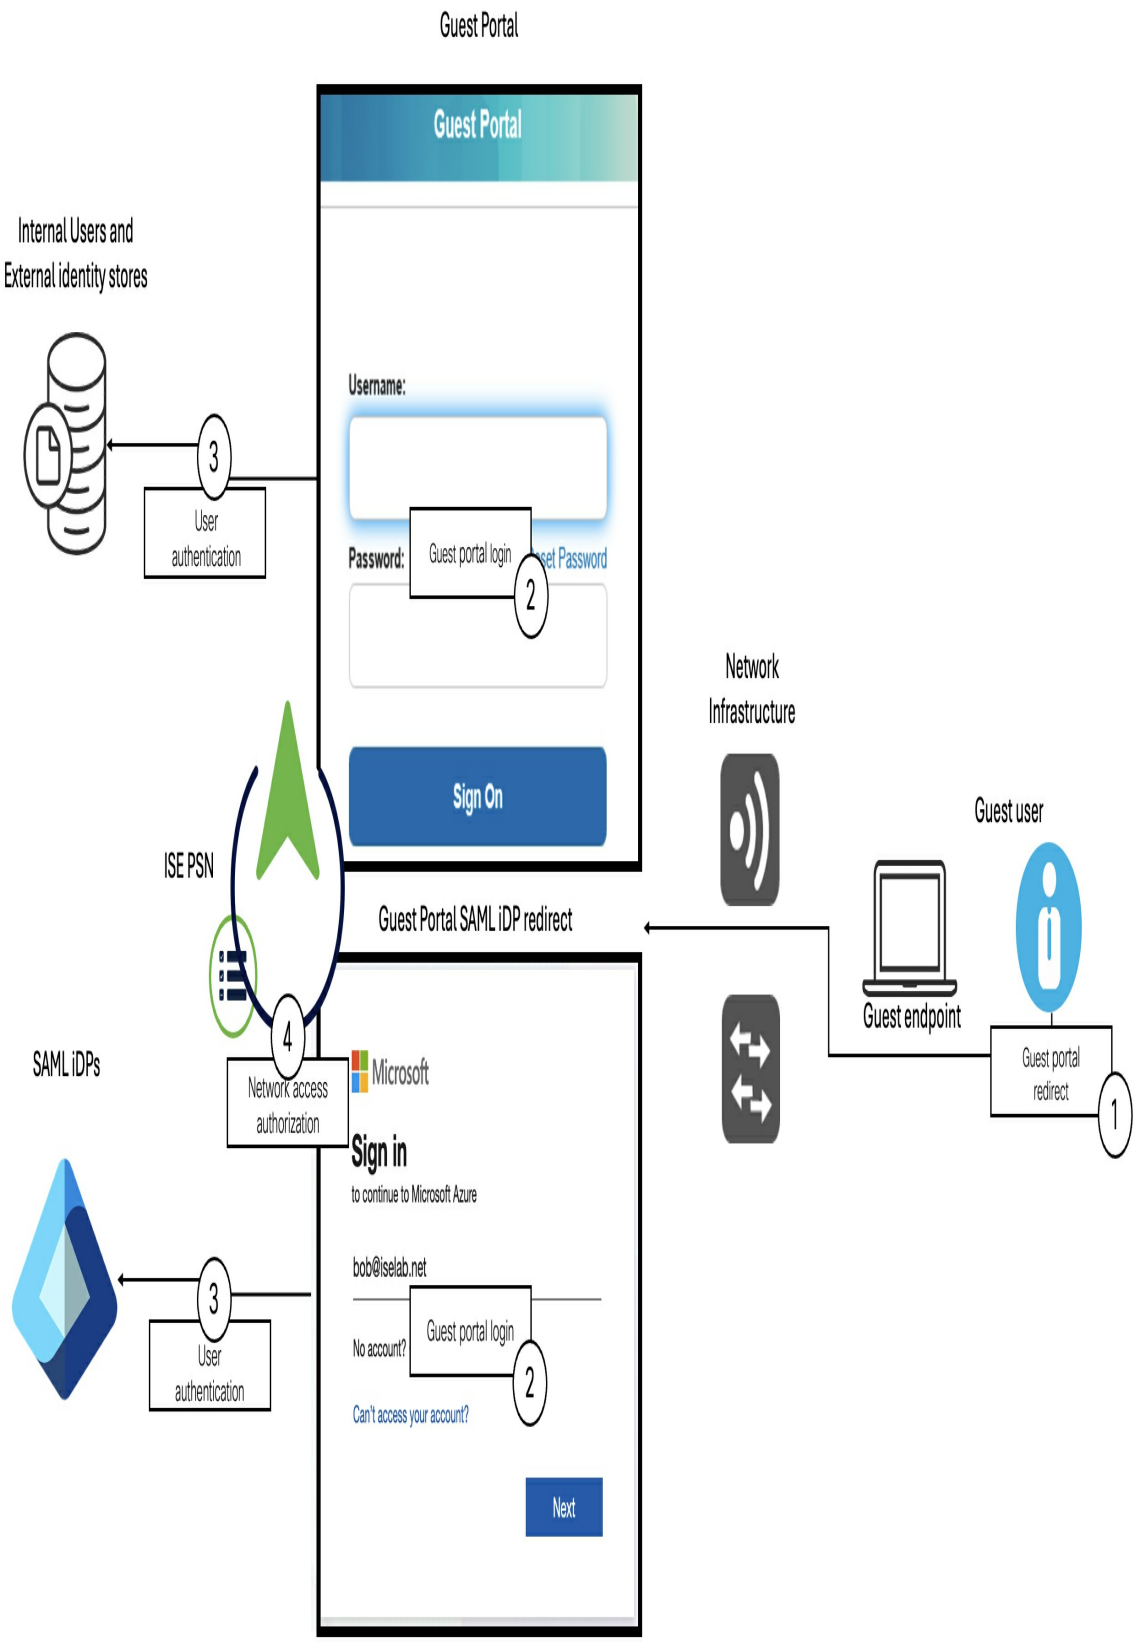
Technet24

Settings to configure email notifications or go to SMS Gateway Provider settings to set up SMS notifications.

On the **SMS Gateway Providers** page, you can review the settings of the predefined gateways or define your own SMS gateway.

Employee Guest Flow

The Employee guest flow is typically not described as a separate guest flow type in the documentation, but it is important to remember that employees can use their credentials to authenticate on the guest portals by default. Figure 9-2 outlines the Employee guest flow.


Guest Portal
Guest Portal
Internal Users and
External identity stores
Username:
Password:
Guest portal login
Sign On
3
User
authentication
2
Network
Infrastructure
Guest user
ISE PSN
Guest Portal SAML iDP redirect
Guest endpoint
Guest portal
redirect
1
4
Network access
authorization
SAML iDPs
Microsoft
Sign in
to continue to Microsoft Azure
bob@iselab.net
No account?
Guest portal login
Can't access your account?
Next
3
User
authentication
2

Technet24

Figure 9-2 *Employee Guest Flow*

The following steps outline the Employee guest flow:

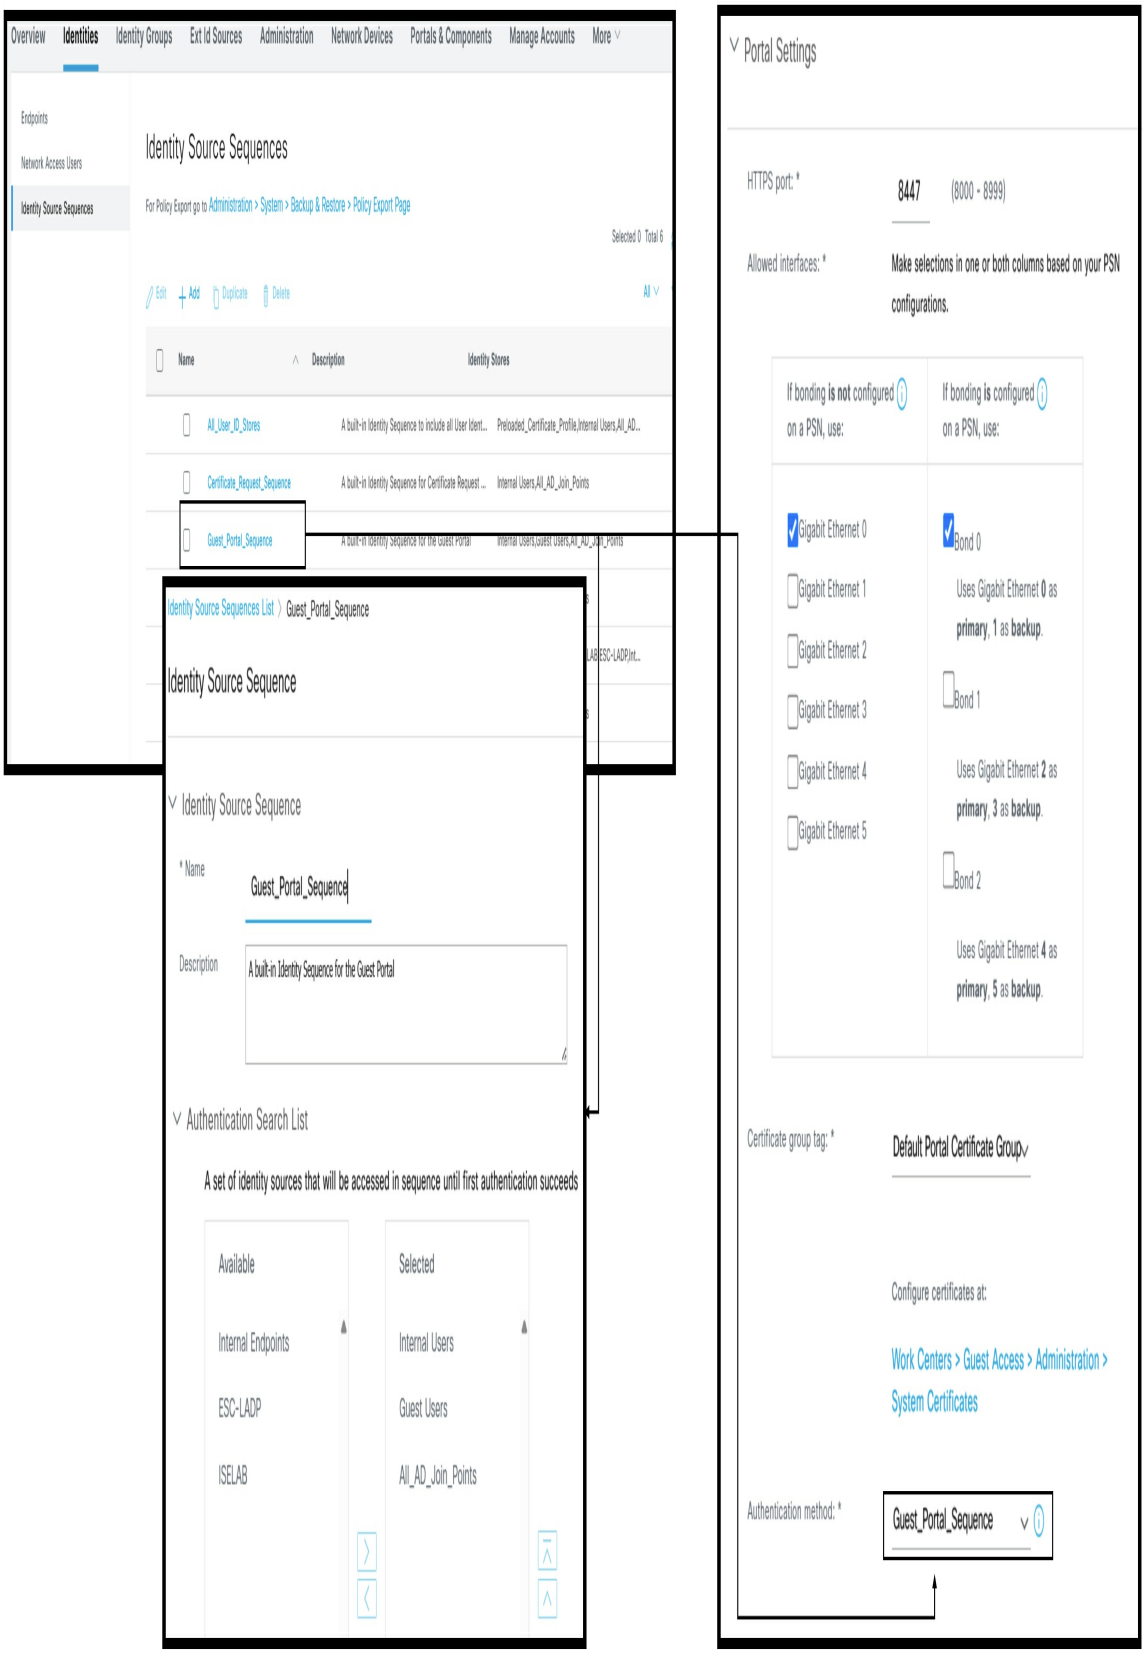
Step 1. The user is redirected to the guest portal as part of the initial network authentication.

Step 2. The user logs in to the guest portal. The login may use standard ISE identity stores such as Internal Users, Active Directory, LDAP, and similar sources. Alternatively, the portal can be configured to use an external SAML identity provider (IdP) for authentication.

Step 3. Authentication occurs in the selected identity store.

Step 4. ISE performs authorization policy evaluation to determine the next policy and triggers a CoA back to the network access device, resulting in a new authentication. At this stage, ISE grants the network access defined for this specific guest use case based on session and endpoint attributes.

For traditional user authentication methods, the identity stores available on the guest portal are determined by the identity store sequence selected in the **Portal Settings** section of the portal configuration. By default, a preconfigured Guest_Portal_Sequence is used, which includes Guest users, Internal users, and All_AD_Join_Points. Figure 9-3 provides an example of the Guest_Portal_Sequence settings and the authentication settings of the guest portal.


Overview
Identities
Identity Groups
Ext Id Sources
Administration
Network Devices
Portals & Components
Manage Accounts
More
Endpoints
Network Access Users
Identity Source Sequences
Identity Source Sequences
For Policy Export go to Administration > System > Backup & Restore > Policy Export Page
Selected 0 Total 6
Edit
Add
Duplicate
Delete
All
Name
Description
Identity Stores
All_User_ID_Stores
A built-in Identity Sequence to include all User Ident...
Preloaded_Certificate_Profile,Internal Users,All_AD...
Certificate_Request_Sequence
A built-in Identity Sequence for Certificate Request ...
Internal Users,All_AD_Join_Points
Guest_Portal_Sequence
Identity Source Sequences List > Guest_Portal_Sequence
Identity Source Sequence
Identity Source Sequence
* Name
Guest_Portal_Sequence
Description
A built-in Identity Sequence for the Guest Portal
Authentication Search List
A set of identity sources that will be accessed in sequence until first authentication succeeds
Available
Internal Endpoints
ESC-LADP
ISELAB
Selected
Internal Users
Guest Users
All_AD_Join_Points
Portal Settings
HTTPS port: *
8447
(8000 - 8999)
Allowed interfaces: *
Make selections in one or both columns based on your PSN configurations.
If bonding is not configured on a PSN, use:
If bonding is configured on a PSN, use:
Gigabit Ethernet 0
Gigabit Ethernet 1
Gigabit Ethernet 2
Gigabit Ethernet 3
Gigabit Ethernet 4
Gigabit Ethernet 5
Bond 0
Uses Gigabit Ethernet 0 as primary, 1 as backup.
Bond 1
Uses Gigabit Ethernet 2 as primary, 3 as backup.
Bond 2
Uses Gigabit Ethernet 4 as primary, 5 as backup.
Certificate group tag: *
Default Portal Certificate Group
Configure certificates at:
Work Centers > Guest Access > Administration > System Certificates
Authentication method: *
Guest_Portal_Sequence

Technet24

Figure 9-3 *Guest Authentication Methods*

You can create a new identity store sequence or update existing ones by navigating to **Work Centers > Guest Access > Identities > Identity Source Sequences**. You can update the authentication settings of the guest portals by going to **Work Centers > Guest Access > Portals & Components > Guest Portals**.

In cases where you need to use a SAML IdP to authenticate guest users, use the following steps to prepare the authentication settings for the SAML flow:

Step 1. Ensure that the SAML IdP certificate issuer is trusted by ISE with the **Trust for Authentication with ISE** setting enabled. You can import certificates by navigating to **Administration > System > Certificates > Trusted Certificates**.

Step 2. Configure new SAML identity providers by navigating to **Work Centers > Guest Access > Ext Id Sources**. At the beginning, there is no need to configure the IdP in ISE; you simply create it with an assigned name.

Step 3. Attach the IdP created in Step 1 to the guest portal where SAML authentication needs to be enabled. Select the IdP in the **Authentication Method** drop-down within the **Portal Settings** section of the portal.

Step 4. After attaching the IdP to the guest portal, navigate back to the IdP settings and export the IdP metadata by going to the **Service Provider Info** tab and clicking the **Export** button.

Step 5. Use the details from the exported metadata file to set up a SAML service provider (SP) application on the IdP side,

such as Cisco Duo.

Step 6. Most modern IdPs allow you to export the metadata file for use on the SP side after the application has been configured. Use this file to upload it back to the IdP settings in ISE. In the IdP settings, go to the **Identity Provider Config** tab and import the IdP metadata.

Another question you need to address during Employee guest flow planning is which guest type settings are assigned to employees and their endpoints. Each guest portal provides an option to configure the guest type, which ensures that all settings from the selected guest type apply to employees logging in through that portal. Figure 9-4 shows an example of the **Employees Using This Portal as Guests Inherit Login Options and Dynamic Endpoint Group Assignment From** setting.

Technet24

Portals Settings and Customization

Portal Name: *

Self-Registered Guest Portal (default)

Description:

Guests may create their own accounts

Language File

Portal test URL

Portal Behavior and Flow Settings | Portal Page Customization

Portal & Page Settings

Portal Settings

HTTPS port: * 8443 (8000 - 8999)

Allowed interfaces: * Make selections in one or both columns based on your PSN configurations.

Employees using this portal as guests inherit login options and dynamic endpoint group assignment from: *

Contractor (default)

Contractor (default)
Daily (default)
Weekly (default)

Display language:

Fallback language:

English - English

Always use: English - English

Figure 9-4 *Guest Type for Employees*

Self-Registration Guest Flow

As part of the Self-Registration guest flow, guest users can provision guest accounts for themselves. The guest account credentials can be displayed at the end of the self-registration process and shared with the user via SMS or email. The Self-Registration flow may optionally include a sponsor approval step, during which a sponsor user must approve the creation of a self-registration account. Figure 9-5 shows an example of the Self-Registration guest flow.

Technet24

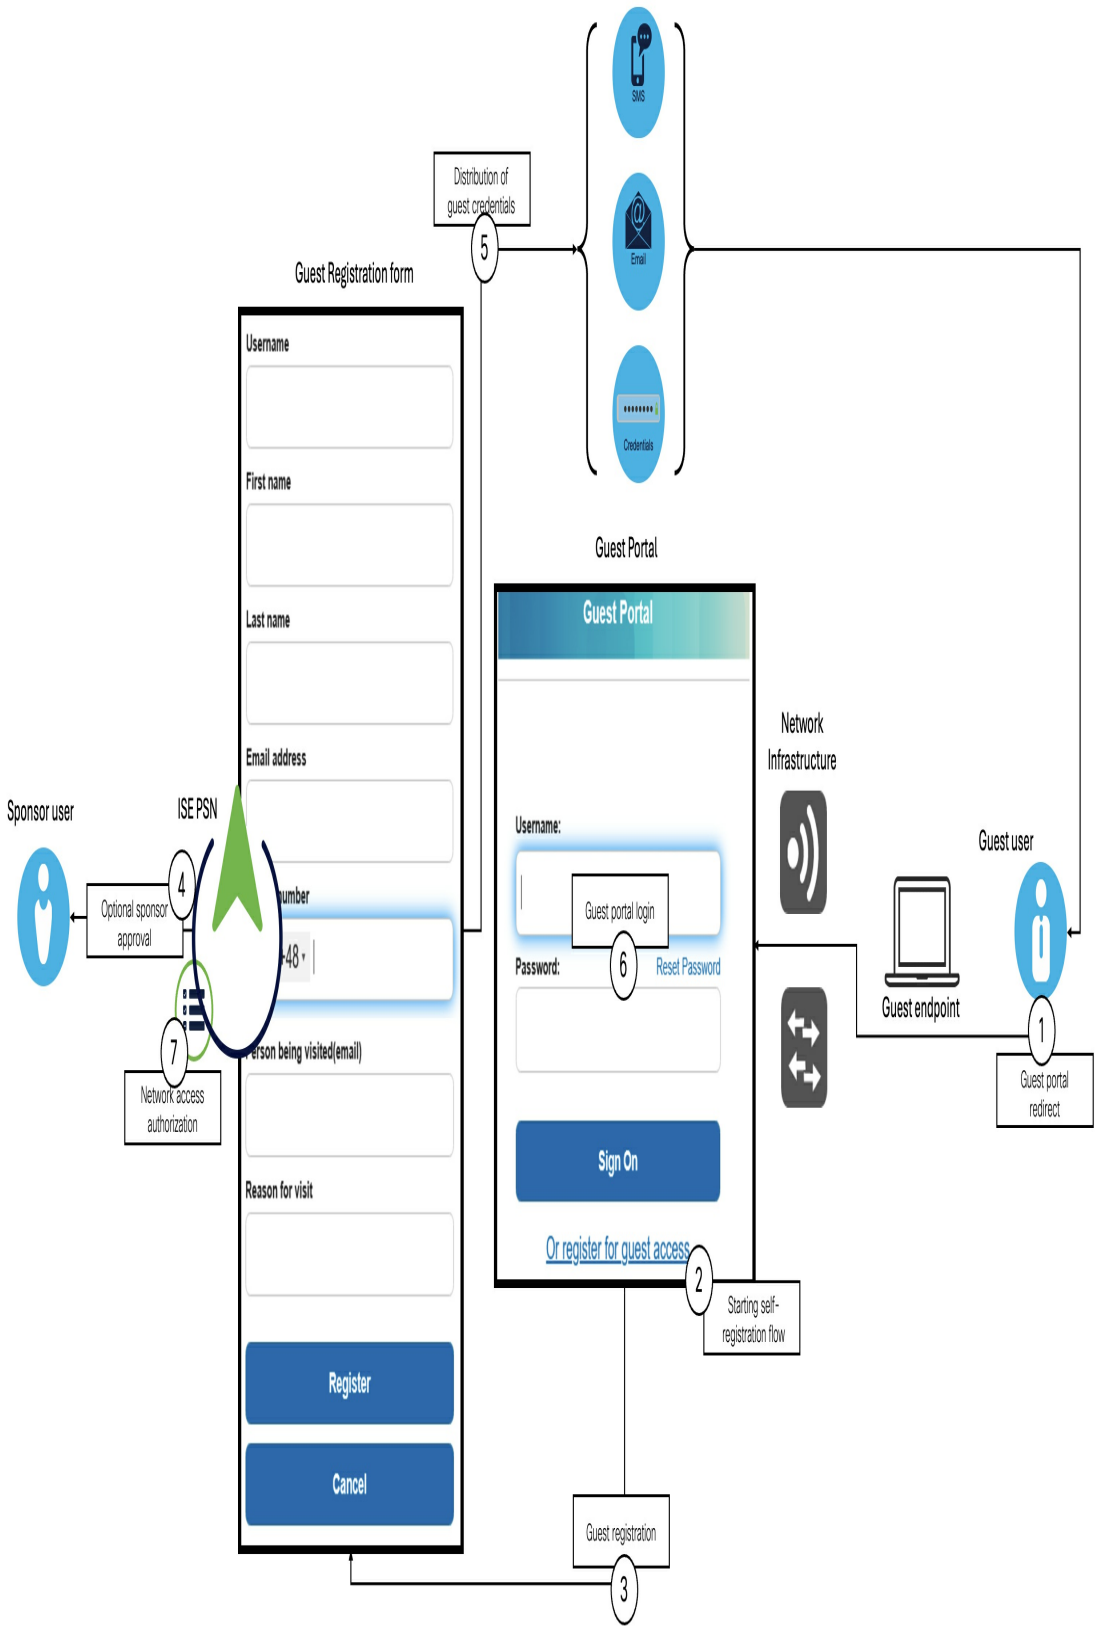

SMS
Email
Credentials
Distribution of guest credentials
5
Guest Registration form
Username
First name
Last name
Email address
Person being visited(email)
Reason for visit
Register
Cancel
Guest Portal
Guest Portal
Username:
Guest portal login
Password:
6
Reset Password
Sign On
Or register for guest access
2
Starting self-registration flow
Guest registration
3
Network Infrastructure
Guest endpoint
Guest user
1
Guest portal redirect
Sponsor user
ISE PSN
4
Optional sponsor approval
7
Network access authorization

Figure 9-5 *Self-Registration Guest Flow*

The following steps explain the Self-Registration guest flow:

Step 1. The user is redirected to the guest portal as part of the initial network authentication.

Step 2. The user selects the self-registration option on the guest portal instead of logging in. You enable the self-registration flow by setting **Allow Guests to Create Their Own Accounts** in the **Login Page Settings** of the portal configuration.

Step 3. The user fills out the registration form.

Step 4. (Optional sponsor approval for the guest account.) You enable sponsor approval by checking **Require Guests to Be Approved** in the **Registration Form Settings** section of the portal settings.

Step 5. Guest credentials are delivered to the end user.

Step 6. The user uses the credentials provided in Step 5 to log in to the guest portal.

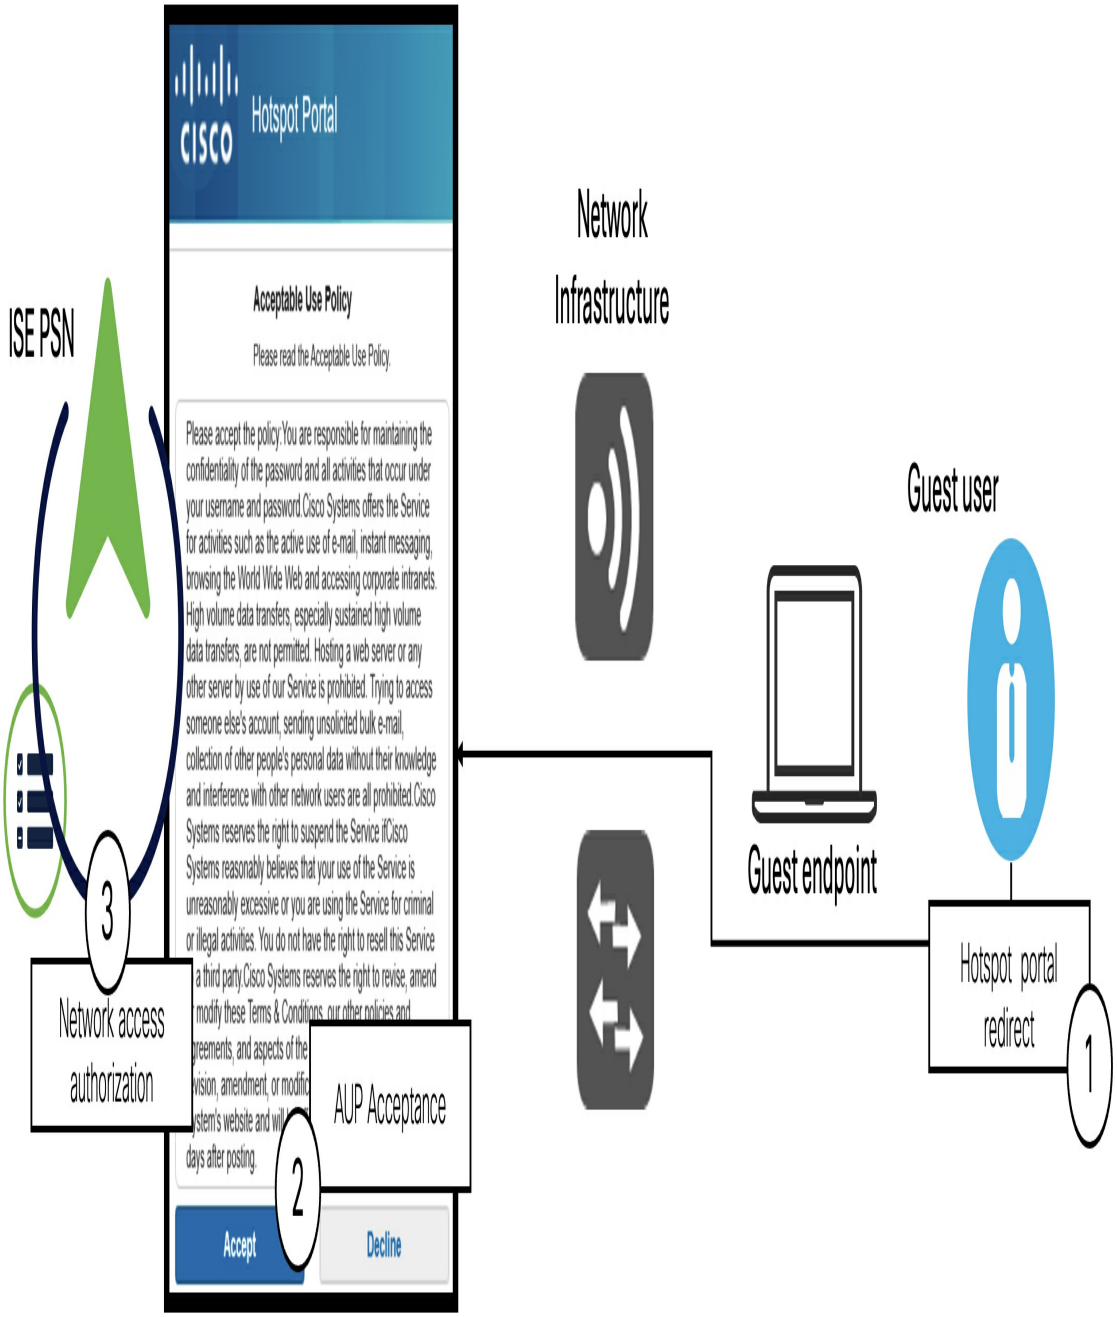
Step 7. ISE performs authorization policy evaluation to determine the next policy and triggers a CoA back to the network access device, resulting in a new authentication. At this stage, ISE grants the network access defined for this specific guest use case based on session and endpoint attributes.

Hotspot Guest Flow

With the Hotspot guest flow, no user authentication is required. The typical use case involves environments where you need to provide quick Internet access with minimal inconvenience to the end user.

Technet24

Common examples include restaurants, coffee shops, airports, and similar venues. Administrators can enable access code-based authentication, requiring any user redirected to the portal to provide the same access code. You can accomplish this by checking the **Require an Access Code** box in the **Acceptable Use Policy (AUP) Page Settings** section of the hotspot portal. Figure 9-6 demonstrates an example of the hotspot flow.


Hotspot Portal
Acceptable Use Policy
Please read the Acceptable Use Policy.
Please accept the policy.You are responsible for maintaining the confidentiality of the password and all activities that occur under your username and password.Cisco Systems offers the Service for activities such as the active use of e-mail, instant messaging, browsing the World Wide Web and accessing corporate intranets. High volume data transfers, especially sustained high volume data transfers, are not permitted. Hosting a web server or any other server by use of our Service is prohibited. Trying to access someone else's account, sending unsolicited bulk e-mail, collection of other people's personal data without their knowledge and interference with other network users are all prohibited.Cisco Systems reserves the right to suspend the Service if Cisco Systems reasonably believes that your use of the Service is unreasonably excessive or you are using the Service for criminal or illegal activities. You do not have the right to resell this Service a third party.Cisco Systems reserves the right to revise, amend modify these Terms & Conditions, our other policies and
greements, and aspects of the
vision, amendment, or modific
stem's website and will
days after posting.
Accept
Decline
ISE PSN
Network access authorization
3
AUP Acceptance
2
Network Infrastructure
Guest endpoint
Guest user
Hotspot portal redirect
1

Technet24

Figure 9-6 *Hotspot Guest Flow*

The following steps explain the Hotspot guest flow:

Step 1. The user is redirected to the guest portal as part of the initial network authentication.

Step 2. The user confirms the Acceptable Use Policy (AUP) on the portal.

Step 3. ISE performs authorization policy evaluation to determine the next policy and triggers a CoA back to the network access device, resulting in a new authentication. In the hotspot use case, since ISE does not have much information about the actual user, the typical next policy uses the endpoint group assigned to the device as part of the hotspot flow. This approach may allow network access without additional redirection for a significant period, depending on the endpoint purge settings located in **Administration > Identity Management > Settings > Endpoint Purge**. If you need users to be redirected to the hotspot portal regularly to reconfirm the AUP, you can use the LastAUPAcceptanceHours endpoint attribute in combination with the endpoint identity group.

Different Approaches to Web Redirection

Web redirection, also known as Centralized Web Authentication, functions as a central component in solutions such as guest access, BYOD, and Endpoint posture assessment. In these scenarios, you typically encounter an initial authentication with limited access, during which either the user or a software agent on the endpoint — such as Cisco Secure Client ISE posture module or Network Setup Assistant — needs to communicate with the web portal on the ISE Policy Service

Node (PSN).

This communication with the portal generally results in additional attributes becoming available for the session or endpoint — for example, use cases being set to guest flow, posture status changing to compliant, or the endpoint identity group being defined. These new attributes trigger authorization policy reevaluation, a CoA, and subsequent reauthentication, with a new access level assigned. This process operates effectively because both the user and the endpoint agent can be directed to the appropriate PSN using CWA.

Web redirection can be performed using multiple approaches:

- **Cisco CWA:** This is the recommended method for all Cisco NADs. All Cisco devices participating in network access control, whether as part of branch or campus networks or as VPN headends, have built-in support for CWA.
- **Third-party NADs with dynamic redirection support:** Some NAD vendors allow the redirect URL to be provisioned from ISE as part of the authentication flow. Before considering this solution, you should consult the documentation of the specific NAD vendor.
- **DNS sinkholing:** ISE PSN can act as a DHCP and DNS server for endpoints when they are in "limited" access. With this approach, redirection occurs by replacing the actual web server IP that the end user tries to access in the DNS reply with the IP of the ISE PSN. You can use this method with third-party NADs that do not support web redirection but have CoA capability.

Cisco CWA

Cisco CWA implementation was explained in Chapter 6, "Authorization Options." Figure 9-7 shows a visualization of the CWA process.

Technet24

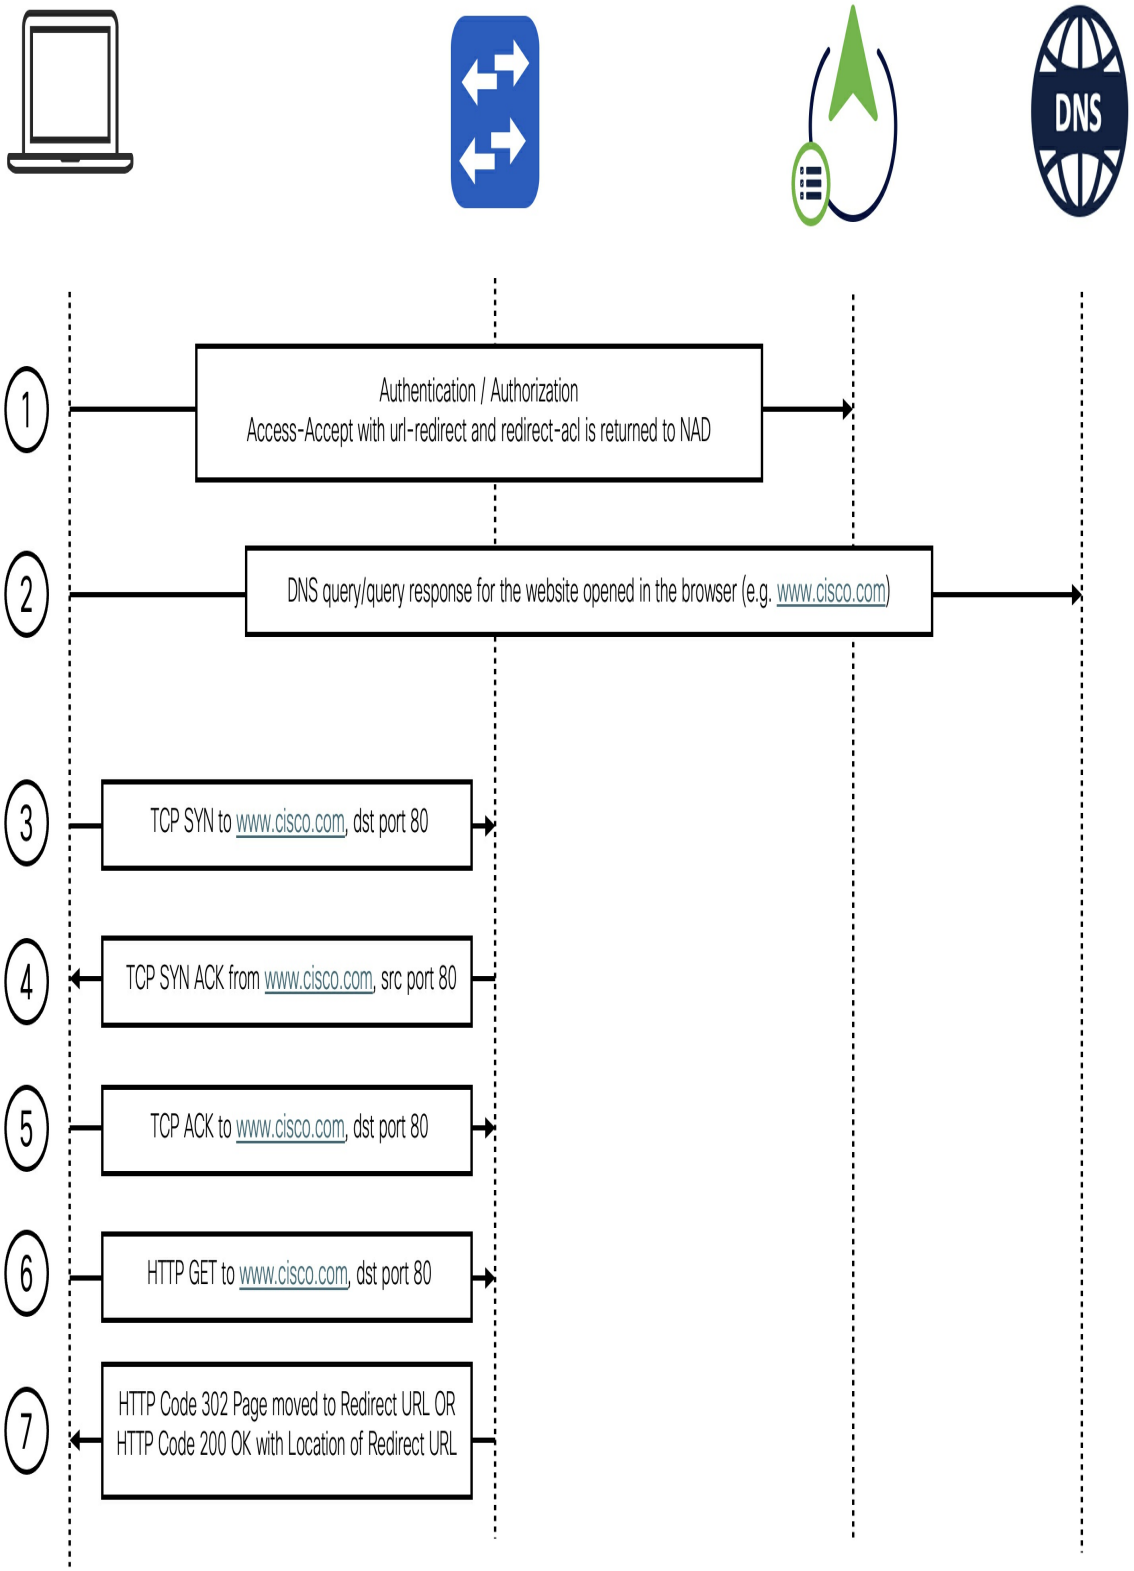

DNS
1
Authentication / Authorization
Access-Accept with url-redirect and redirect-acl is returned to NAD
2
DNS query/query response for the website opened in the browser (e.g. www.cisco.com)
3
TCP SYN to www.cisco.com, dst port 80
4
TCP SYN ACK from www.cisco.com, src port 80
5
TCP ACK to www.cisco.com, dst port 80
6
HTTP GET to www.cisco.com, dst port 80
7
HTTP Code 302 Page moved to Redirect URL OR
HTTP Code 200 OK with Location of Redirect URL

Figure 9-7 *Cisco CWA Flow*

In any web redirection flow, it is important to ensure that the redirection lands on the same PSN where authentication occurred. Normally, the redirect URL resolves this issue because it contains the fully qualified domain name (FQDN) of the ISE PSN where the user or endpoint was authenticated. One exception to this rule arises when administrators manually override the CWA setting with a static FQDN, as demonstrated in Figure 9-8.

Technet24

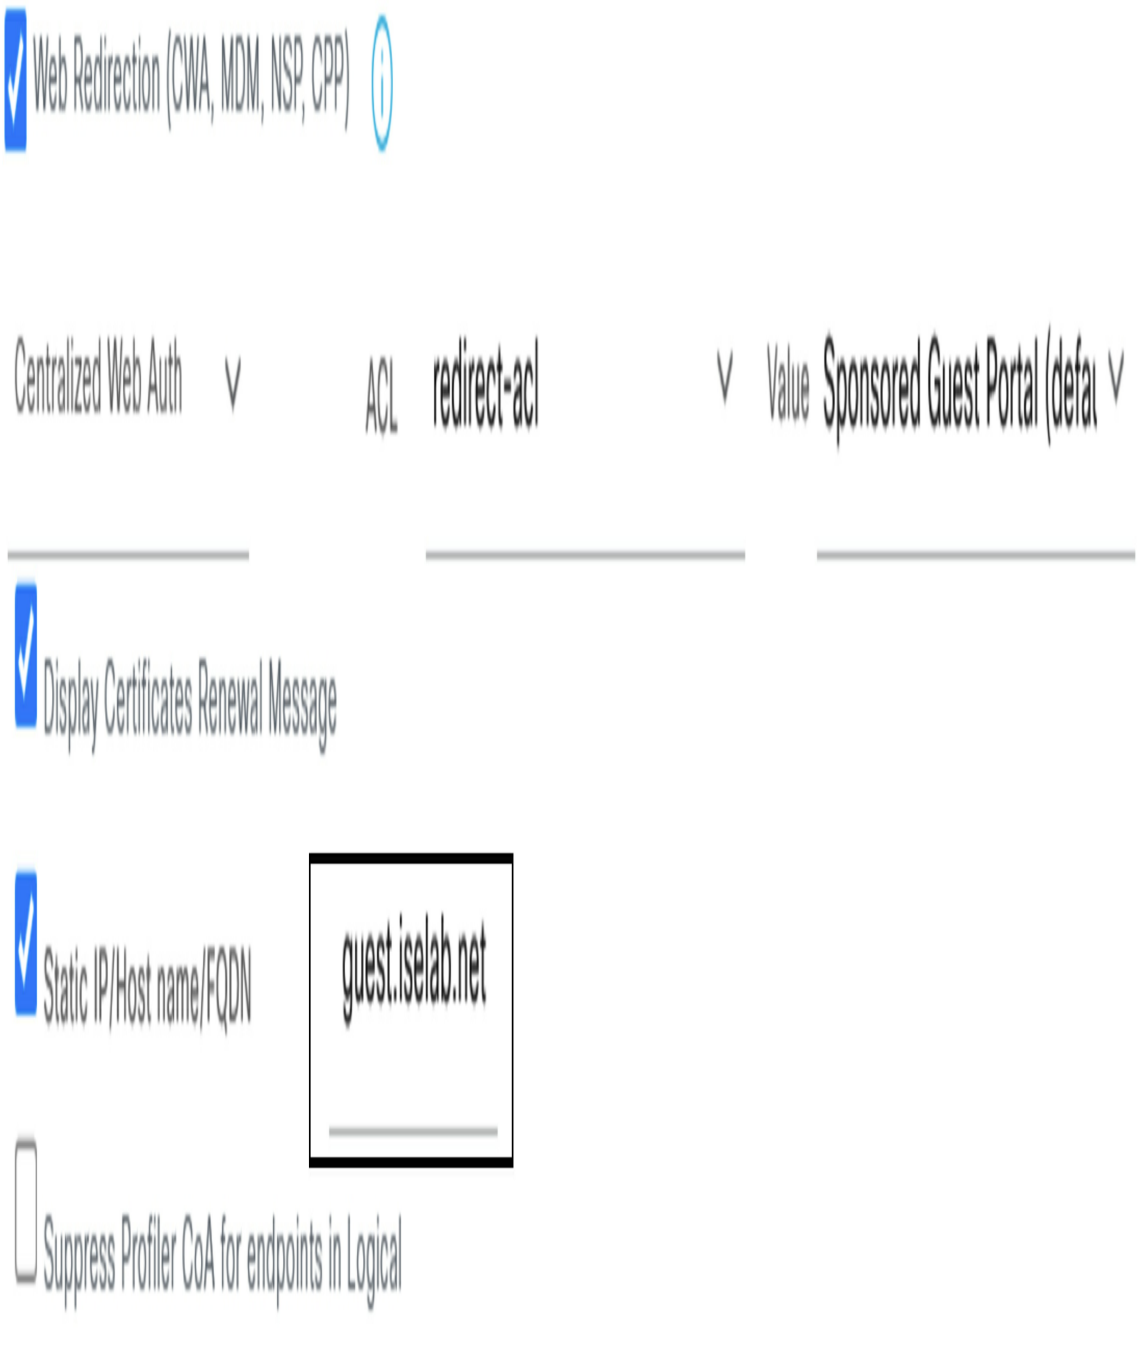

Web Redirection (CWA, MDM, NSP, CPP)
Centralized Web Auth
ACL
redirect-acl
Value
Sponsored Guest Portal (defau
Display Certificates Renewal Message
Static IP/Host name/FQDN
guest.iselab.net
Suppress Profiler CoA for endpoints in Logical

Figure 9-8 *Redirection to a Static FQDN*

This approach often proves problematic because the FQDN is typically configured to resolve to any ISE PSN node in a round-robin fashion on the DNS server side. As a result, when authentication and DNS resolution land on the same PSN, the flow works as expected. However, if the PSN IP returned by DNS is different from the one where authentication occurred, the user sees Error 400 instead of the guest portal. It is recommended to avoid this configuration and rely on a dynamic redirect URL.

DNS Sinkholing Redirection

You should not use redirection with Cisco NADs. Before implementation, you need to ensure that the NAD you plan to use supports CoA. While redirection can be performed directly by the PSN using DNS sinkholing, CoA remains necessary to move the user from the redirect phase to the next access level. Figure 9-9 shows a visualization of DNS sinkholing redirection.

Technet24

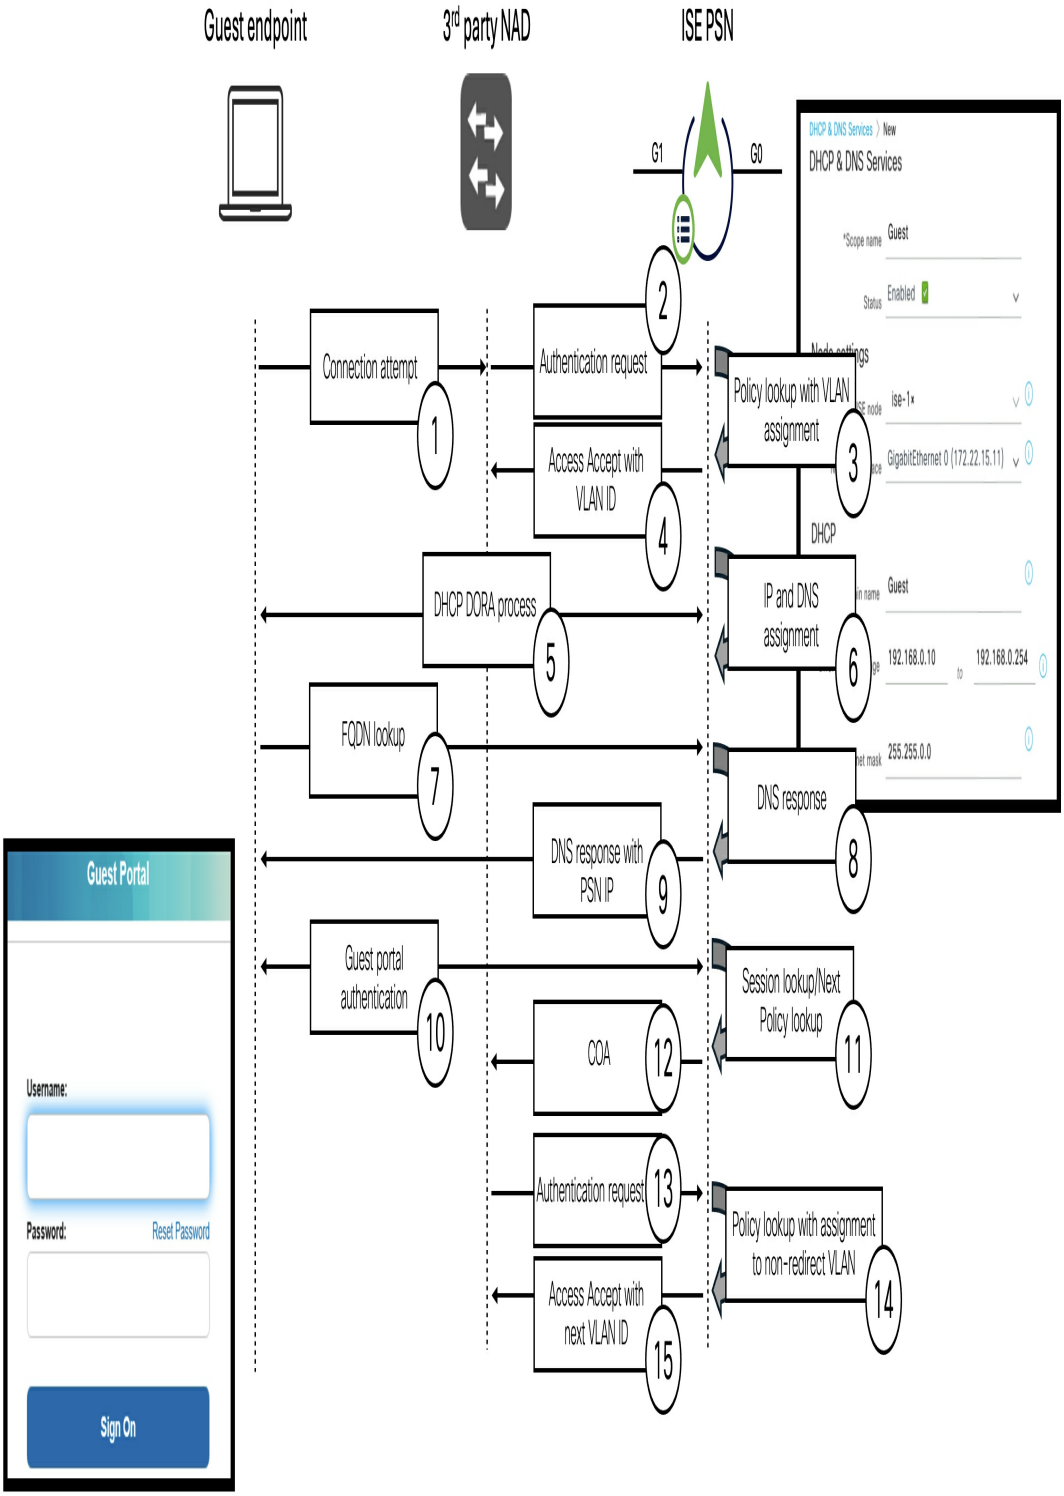

Guest endpoint
3rd party NAD
ISE PSN
G1
G0
DHCP & DNS Services > New
DHCP & DNS Services
*Scope name
Guest
Status
Enabled
Node settings
ISE node
ise-1
Interface
GigabitEthernet 0 (172.22.15.11)
DHCP
Scope name
Guest
IP range
192.168.0.10
to
192.168.0.254
Subnet mask
255.255.0.0
Connection attempt
1
Authentication request
2
Policy lookup with VLAN assignment
3
Access Accept with VLAN ID
4
DHCP DORA process
5
IP and DNS assignment
6
FQDN lookup
7
DNS response
8
DNS response with PSN IP
9
Guest portal authentication
10
Session lookup/Next Policy lookup
11
COA
12
Authentication request
13
Policy lookup with assignment to non-redirect VLAN
14
Access Accept with next VLAN ID
15
Guest Portal
Username:
Password:
Reset Password
Sign On

Figure 9-9 *DNS Sinkholing Redirection*

The following steps explain DNS sinkholing redirection:

Step 1. The endpoint connects to a third-party NAD.

Step 2. The NAD initiates simple MAB authentication.

Step 3. The PSN performs policy lookup and assigns a "limited access" authorization profile with the VLAN name or ID. The same profile should include CWA.

Step 4. The switch places the endpoint in the VLAN provided by PSN. At this point, one of the ISE interfaces needs to be either directly present in the VLAN (typically a non-G0 interface), or if there is an L3 border between NAD and PSN, you must configure the NAD to relay DHCP requests to the PSN IP.

Step 5. The endpoint requests network configuration using DHCP.

Step 6. The PSN assigns an IP address and DNS server, which is the PSN IP, to the endpoint from the preconfigured DHCP pool.

Step 7. The endpoint sends DNS requests for the web resource that the end user tries to open.

Step 8. The PSN processes the DNS request. For all requests except the domains present in the External domains list, The PSN always returns its own IP address.

Step 9. The DNS response includes the PSN IP instead of the IP of the requested resource.

Step 10. The web request is initiated to the PSN IP. At this point, the PSN needs to decide if this particular endpoint is in the "redirect" state. To confirm this, the PSN performs a session

Technet24

lookup using the source IP to locate the MAC address in the DHCP bindings. Then, the MAC address is used to locate the authentication session in the PSN session cache and check the session attributes. If The PSN confirms that the endpoint is in the redirect state, the guest portal is displayed where the user can perform authentication.

Step 11. ISE performs authorization policy evaluation to determine the next policy and triggers a CoA back to the network access device, resulting in a new authentication.

Step 12. The CoA is sent to the NAD.

Step 13. The NAD submits a new authentication request.

Step 14. The PSN performs policy lookup and assigns a "full access" authorization profile using the attributes added by the PSN into the session cache after portal authentication. The authorization profile in this scenario should contain a VLAN name or ID where the endpoint can obtain new network settings from the regular DHCP server. The ISE DHCP server is used only for redirection.

Step 15. Authorization attributes are returned to the NAD.

DNS sinkholing redirection can be used for NAD profiles that have redirection disabled and CoA support enabled. You can review NAD profiles by navigating to **Administration > Network Resources > Network Device Profiles**. Preconfigured profiles cannot be modified, but you can copy an existing profile and edit the settings in the copy.

After you identify or configure the applicable NAD profile, you need to create the DHCP pool. To do this, navigate to **Administration > System > Settings > DHCP & DNS Services** and click **Add**. You cannot enable the DHCP service for the ISE interface where DHCP profiling is running. It is recommended to use an interface different

from G0 for the DHCP service or disable DHCP-based profiling probes.

All DNS requests from clients that receive their IP addresses from ISE and use ISE PSN as their DNS server will be redirected to the ISE IP, except for the domains specified in the **External Domain** section. For those domains, proper name resolution will occur using the DNS servers defined in the **External DNS Servers** section.

The final two steps are to create the authorization profiles and assign them to the authorization policy. You need to ensure that authentication requests from the specific NAD always land in the authorization policy that includes an authorization profile created for that NAD. To achieve this, it is recommended to create a dedicated policy set for each third-party NAD type. As a policy set match statement, you can use the DEVICE:Network Device Profile condition, pointing to the specific Network Device Profile. Figure 9-10 shows an example of a policy set created for the HPWired NAD profile.

Technet24

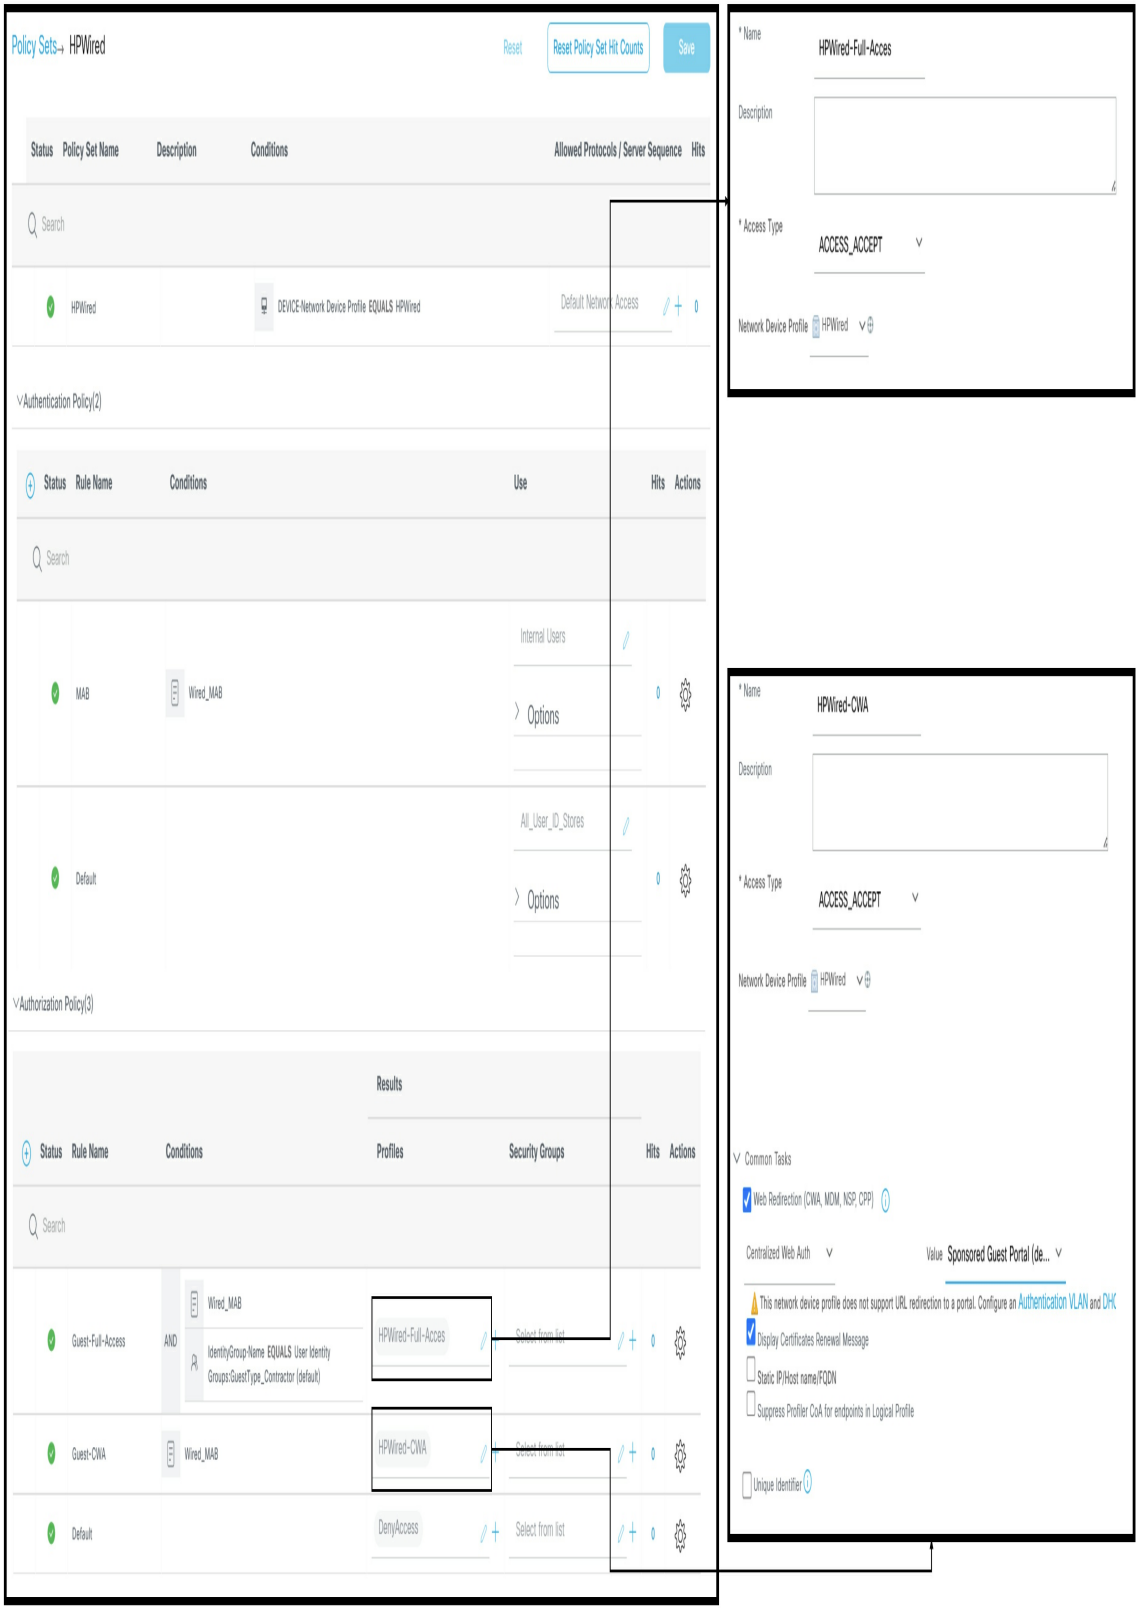

Policy Sets→ HPWired
Reset
Reset Policy Set Hit Counts
Save
Status Policy Set Name Description Conditions Allowed Protocols / Server Sequence Hits
Search
HPWired
DEVICE-Network Device Profile EQUALS HPWired
Default Network Access
Authentication Policy(2)
Status Rule Name Conditions Use Hits Actions
Search
MAB
Wired_MAB
Internal Users
Options
Default
All_User_ID_Stores
Options
Authorization Policy(3)
Results
Status Rule Name Conditions Profiles Security Groups Hits Actions
Search
Guest-Full-Access
AND
Wired_MAB
IdentityGroup-Name EQUALS User Identity Groups:GuestType_Contractor (default)
HPWired-Full-Acces
Select from list
Guest-CWA
Wired_MAB
HPWired-CWA
Select from list
Default
DenyAccess
Select from list
* Name
HPWired-Full-Acces
Description
* Access Type
ACCESS_ACCEPT
Network Device Profile
HPWired
* Name
HPWired-CWA
Description
* Access Type
ACCESS_ACCEPT
Network Device Profile
HPWired
Common Tasks
Web Redirection (CWA, MDM, NSP, CPP)
Centralized Web Auth
Value
Sponsored Guest Portal (de...
This network device profile does not support URL redirection to a portal. Configure an Authentication VLAN and DHC
Display Certificates Renewal Message
Static IP/Host name/FQDN
Suppress Profiler CoA for endpoints in Logical Profile
Unique Identifier

Figure 9-10 *ISE Authorization Profiles and Polices for DNS Sinkholing Redirection*

In the policy set, MAB authentication policy is configured to "Continue" if the user is not found, enabling initial authentication for the guest. Two authorization policies are present: One redirects to the guest portal, and the other provides full access once the user group is identified as Contractor, which is the default guest type used for employee login.

Redirection and HTTPS Traffic

While devices such as wireless LAN controllers (WLCs) and switches theoretically support web redirection for HTTPS traffic, you should generally avoid enabling it whenever possible. HTTPS redirection presents several serious drawbacks:

- **User confusion due to certificate warnings:** When a user attempts to access an HTTPS resource, the WLC or switch performs a TCP intercept, which then leads to a TLS handshake for HTTPS traffic. The issue is that switches and WLCs can only present a local certificate, which does not match the requested resource's name, resulting in a certificate error in the browser.
- **Guest and other portals being blocked by HTTP Strict Transport Security (HSTS) during redirection:** HSTS forces browsers to use HTTPS for sites previously accessed over HTTPS, regardless of what the user enters in the address bar. During redirection to such resources, the browser blocks redirection due to the HSTS max-age feature, which strictly prohibits access in case of certificate violations.
- **Increased load on NADs:** Regular HTTP redirection requires only TCP intercept from the NAD, but HTTPS redirection also

Technet24

necessitates building a TLS tunnel, which is a resource-intensive operation.

How does redirection work in the real world without HTTPS, considering that most web resources today use HTTPS? Modern browsers use HTTP probes to predefined external resources to detect when an endpoint is in a redirect state. Because of these probes, there is no genuine need for HTTPS redirection.

To ensure that HTTPS redirection is not used:

- For switches, you can disable the HTTPS server with the configuration mode command **no ip http secure-server**.
- For Catalyst WLCs, you can keep the HTTPS server enabled but specifically disable HTTPS web authentication, as shown in Example 9-1.

Example 9-1 *Cisco WLC 9800 Series Controller* ***parameter-map*** *Example*

```
parameter-map type webauth global
   webauth-http-enable
```

Stitching Attributes for Authorization

In the context of guest access, as explained previously, you typically encounter three distinct events:

- **Initial authentication:** Also called "limited access" or "redirect state." At this point, a simple MAB authentication is usually performed with minimal context available to ISE. Often, even the endpoint MAC address is unknown to ISE. The information typically known at this stage includes the type of NAD, NAD location, SSID name, and other generic attributes. This stage

generally results in CWA.

- **Guest portal authentication:** At this stage, the user enters credentials on the guest portal or at least accepts the AUP of the hotspot portal. As a result, the guest endpoint is typically added to the endpoint identity group, and additional user context is collected, such as user group membership. This extra information can be used as conditions in the next authorization policy, which provides less restrictive network access. When new details are added to the endpoint or session context, authorization policies are reevaluated, and if the next policy can be selected, a CoA is issued to the NAD.
- **Subsequent authentication:** Also called the "full access" state. As a result of the CoA, the session is authenticated again, and the next authorization policy is applied.

When you're planning your authorization policies, it is important to understand which attributes belong to the session and which belong to the endpoint. Session attributes are typically available only while the session is active and have not been terminated with an accounting-stop message. In contrast, endpoint attributes remain accessible as long as the endpoint is known to ISE. The endpoint may become "unknown" to ISE if the MAC address changes due to MAC randomization or if the endpoint is removed from the ISE database through manual deletion or endpoint purge policy action.

Examples of session attributes may include the following:

- **Called Station ID:** This typically contains the SSID name in Wi-Fi use cases.
- **Internal or external user groups:** For example, these are groups of ISE internal users such as guest types, or groups retrieved from external identity sources like Active Directory or LDAP.

Technet24

- **Network Access UseCase set to Guest Flow:** ISE sets the UseCase to Guest Flow after successful authentication on the guest portal.

There are many more session attributes than those listed here, but this short selection provides a general understanding of what session attributes include. Figure 9-11 shows an example of authorization policies used for guest access that are based on session attributes. The lower policy named Guest-CWA matches for wireless MAB authentication when the RADIUS Called-Station-ID includes the word *Guest*. Once the user passes authentication on the guest portal, the next authorization policy — Guest-Access — can be applied, which uses the same checks but additionally includes the user identity group GuestType_Contractor (default).

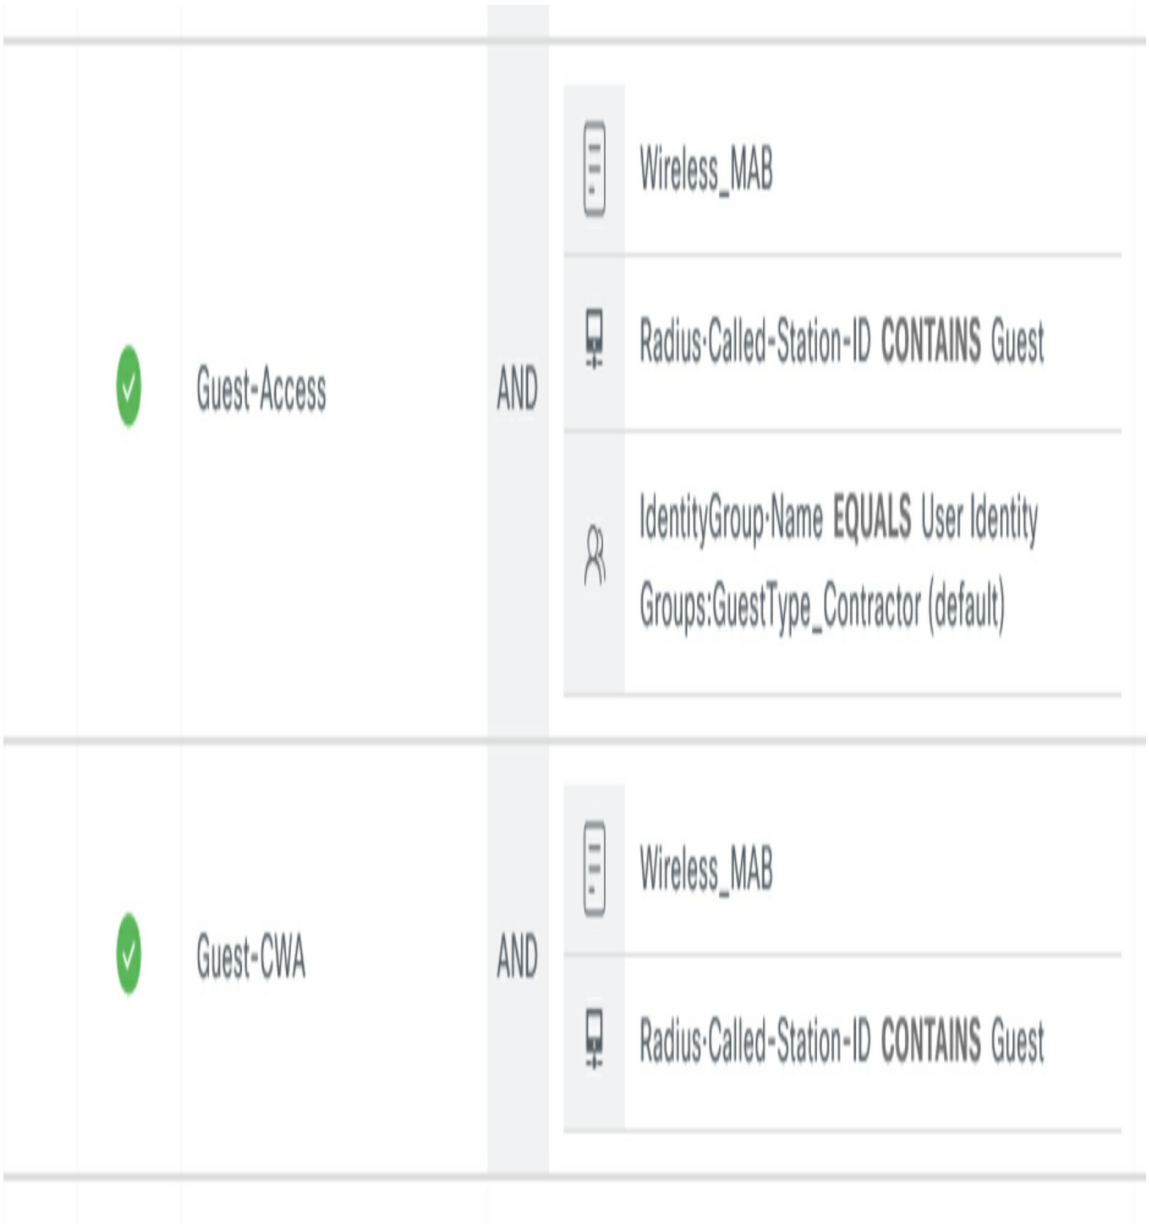


Figure 9-11 *Guest Access Authorization Policies Based on Session Attributes*

Endpoint attributes may include the following:

- **Endpoint identity group:** This value for the guest endpoint is typically set according to the Guest Type settings. Within each guest type, you can configure the group name that ISE should use to place the endpoint after the user successfully passes

Technet24

authentication on the guest portal. You can configure guest types by navigating to **Work Centers > Guest Access > Portals & Components > Guest Types**. Inside each guest type, you specify the device identity group in the section called **Login Options**, as shown in Figure 9-12.

Guest Type

Guest type name: *

Daily (default)

Description:

Default settings allow network access for just 1-5 days.

Login Options

☑ Maximum simultaneous logins 3 (1-999)

When guest exceeds limit:

- ◉ Disconnect the oldest connection
- ○ Disconnect the newest connection

☐ Redirect user to a portal page showing an error message ⓘ

This requires the creation of an authorization policy rule

Maximum devices guests can register: 5 (1-999)

Assign endpoint identity groups for guest device registration workflows:

◉ Standard ○ Dynamic ⓘ

Identity Group

GuestEndpoints ⓘ

Technet24

Figure 9-12 *Endpoint Group Assignment in the Guest Type Settings*

- **LastAUPAcceptanceHours:** When a guest endpoint is saved in ISE, the AUP acceptance timestamp is added as an endpoint attribute. Later, you can use the LastAUPAcceptanceHours check in authorization policy to redirect endpoints to the ISE web portal after a certain period, regardless of how many disconnects or reconnects (new sessions) occur within that range.
- **DeviceRegistrationStatus:** After guest devices pass through the guest portal for the first time, the DeviceRegistrationStatus attribute is added. You can use this in authorization policy to match all guest endpoints, regardless of their identity group.

As with session attributes, there are many more endpoint attributes than those explained here.

Figure 9-13 contains an example of authorization policies built using endpoint attributes. The Hotspot-CWA policy performs initial user redirection to the hotspot portal. When the endpoint's LastAUPAcceptanceHours is less than 8, the endpoint belongs to the Guest Endpoints identity group, the connection is authenticated over Wireless MAB, and the Called Station ID contains Public-WiFi, the policy named Hotspot-Full-Access grants full access. If, on a subsequent connection, ISE detects that LastAUPAcceptanceHours is more than 8, the policy Hotspot-CWA-Old-AUP redirects the user to the hotspot portal again.

✓	Hotspot-Full-Access	AND	Wireless_MAB Radius·Called-Station-ID **CONTAINS** Public-WiFi IdentityGroup·Name **EQUALS** Endpoint Identity Groups:GuestEndpoints EndPoints·LastAUPAcceptanceHours **LESS** 8	PermitAccess
✓	Hotspot-CWA-Old-AUP	AND	Wireless_MAB Radius·Called-Station-ID **CONTAINS** Public-WiFi IdentityGroup·Name **EQUALS** Endpoint Identity Groups:GuestEndpoints EndPoints·LastAUPAcceptanceHours **GREATER** 8	Hotspot-CWA
✓	Hotspot-CWA	AND	Wireless_MAB Radius·Called-Station-ID **CONTAINS** Public-WiFi	Hotspot-CWA

Technet24

Figure 9-13 *Guest Access Authorization Policies Based on Endpoint Attributes*

Sponsor Groups, Users and Portals

For the Sponsored guest flow in ISE, the following high-level components are used:

- **Sponsor portal(s):** This portal allows sponsor users to create and manage guest accounts. Sponsor portals are accessible on all ISE nodes with the PSN role enabled. End users can access sponsor portals using the exact portal URL or FQDN defined for the portal. If you use an FQDN, the DNS server should resolve the sponsor portal FQDN to one or more ISE PSN IPs. You can also use a load balancer virtual IP (VIP) to distribute sponsor access among different ISE nodes. Sponsor users must authenticate when accessing the sponsor portal.
- **Sponsor groups:** Sponsor groups serve as authorization for sponsor users. If a sponsor user can authenticate but is not assigned to any group, access to the portal will fail. Sponsor groups also determine how sponsors can create and manage guest users, such as which guest users a sponsor can view and which guest types a sponsor can use to create guest accounts.

Both sponsor portals and sponsor groups rely on other configuration elements, such as the requirement for a certificate on the sponsor portal and the need for guest locations in sponsor groups. However, it is important to remember that sponsor portals and sponsor groups are the key elements to which other components are connected. Figure 9-14 explains the interconnection between different configuration elements needed for the sponsored guest flow to function.


CN=ise-1.iselab.net, OU=ISE LAB#iselab-LAB-DC-1-CA#00005
Admin, EAP Authentication, RADIUS DTLS, Portal, TACACS
Default Portal Certificate Group
ise-1.iselab.net
1
When portal loads the TLS connection is protected by the certificate assigned to the sponsor portal through the Portal Tag
Sponsor Portal
Sign On
Use the Sponsor portal to manage guest accounts. Sign on with your username and password after verifying that you agree to the terms and conditions, if required.
Username:
Password:
Sign On
Certificate group tag:*
Default Portal Certificate Group
Configure certificates at:
Work Centers > Guest Access > Administration > System Certificates
Fully qualified domain names (FQDN) and host names:
Identity source sequence: *
Sponsor_Portal_Sequence
Sponsor users are authenticated using identity source sequence assigned to the portal
2
Identity Source Sequence
Identity Source Sequence
* Name
Sponsor_Portal_Sequence
Authentication Search List
A set of identity sources that will be accessed in sequence until first authentication succeeds
Available
Internal Endpoints
Guest Users
Selected
Internal Users
All_AD_Join_Points
SponsorGroup
Disable Sponsor Group
Sponsor group name:*
GROUP_ACCOUNTS (default)
Description:
Sponsors assigned to this group can manage just the guest accounts created by sponsors from the same sponsor group. By default, users in the GROUP_ACCOUNTS user identity group are members of this sponsor group
Match Criteria
Member Groups - Sponsor must belong to at least one of the selected groups.
Members...
GROUP_ACCOUNTS (default)
ISELAB:iselab.net/Users/Domain Users
As a result of the authentication the sponsor user needs to be assigned to one of the sponsor groups
3

Technet24

Figure 9-14 *Logic Behind Sponsor Portals and Groups*

The following steps explain the interaction between sponsor users and sponsor portals:

Step 1. When a sponsor user enters the sponsor portal URL, a TLS connection is established with one of the PSNs. The exact PSN is determined by the specific sponsor portal URL, FQDN name resolution, or a third-party load balancer. When the TLS connection is established, the PSN returns the certificate assigned to the sponsor portal to the user by using a certificate group tag. The certificate group tag functions by first assigning the certificate to an existing or new group tag at the time of certificate installation at **Administration > System > Certificates > System Certificates**, and then using the same group tag in the portal settings.

Step 2. Sponsor authentication can be done using standard authentication methods. For this, the default identity store sequence called Sponsor_Portal_Sequence is used, which includes Internal User and All_AD_Join_Point. Alternatively, similar to guest access, you can assign a SAML IdP external identity source to the sponsor portal. In the case of SAML IdP, the configuration is similar to what is used for guest users, with one exception: The SAML assertion must return the group for the sponsor user.

Step 3. The authenticated sponsor user is now authorized to select the sponsor group to which the user belongs. For standard authentication methods, internal and external groups are directly assigned to the sponsor group, as demonstrated in Figure 9-10. When SAML is used, sponsor group assignment is a two-step process. First, you specify in the

SAML IdP settings on ISE which attribute in the SAML assertion represents group membership and add the mappings between group IDs and group names. Second, you add the group names from the previous step to the sponsor groups, as demonstrated in Figure 9-15.

Technet24

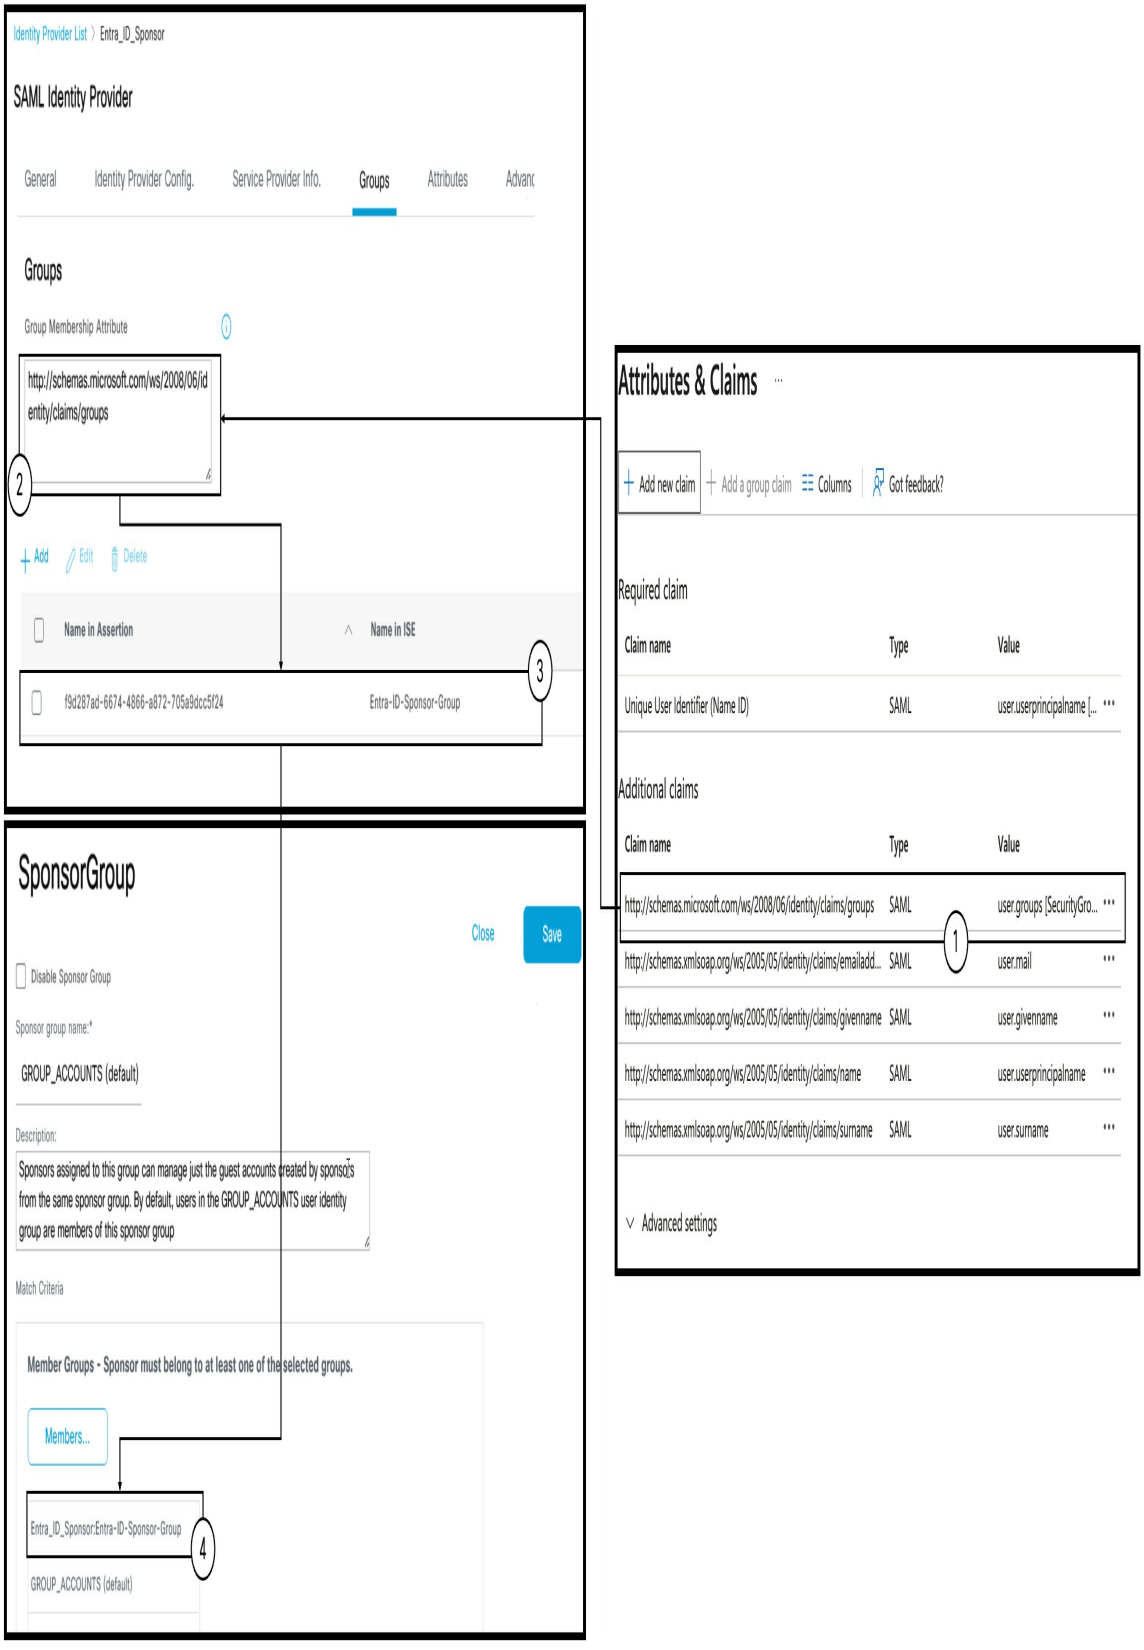

Identity Provider List > Entra_ID_Sponsor
SAML Identity Provider
General
Identity Provider Config.
Service Provider Info.
Groups
Attributes
Advanc
Groups
Group Membership Attribute
http://schemas.microsoft.com/ws/2008/06/identity/claims/groups
2
Add
Edit
Delete
Name in Assertion
Name in ISE
f9d287ad-6674-4866-a872-705a9dcc5f24
Entra-ID-Sponsor-Group
3
Attributes & Claims
Add new claim
Add a group claim
Columns
Got feedback?
Required claim
Claim name
Type
Value
Unique User Identifier (Name ID)
SAML
user.userprincipalname [...
Additional claims
Claim name
Type
Value
http://schemas.microsoft.com/ws/2008/06/identity/claims/groups
SAML
user.groups [SecurityGro...
1
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailadd...
SAML
user.mail
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/givenname
SAML
user.givenname
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name
SAML
user.userprincipalname
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/surname
SAML
user.surname
Advanced settings
SponsorGroup
Close
Save
Disable Sponsor Group
Sponsor group name:*
GROUP_ACCOUNTS (default)
Description:
Sponsors assigned to this group can manage just the guest accounts created by sponsors from the same sponsor group. By default, users in the GROUP_ACCOUNTS user identity group are members of this sponsor group
Match Criteria
Member Groups - Sponsor must belong to at least one of the selected groups.
Members...
Entra_ID_Sponsor:Entra-ID-Sponsor-Group
4
GROUP_ACCOUNTS (default)

Figure 9-15 *Sponsor SAML Authorization*

In Figure 9-15, Entra ID is used as an example of a SAML IdP, and the following steps explain the required process to authorize a SAML user into the ISE sponsor group:

Step 1. The SAML settings of the SP application must include the group attribute. For Entra ID, you can configure this attribute by navigating to **Enterprise Application > SAML-Based Sign-On > Attributes & Claims**, where you can add a group claim as shown in the example.

Step 2. On the ISE side, in the SAML IdP settings, you can add the attribute used in the assertion to represent groups by navigating to the **Groups** tab in the **SAML IdP** settings.

Step 3. Once ISE recognizes the attribute used to carry group information, you can add mappings between the group IDs in the assertion and friendly group names to be used on the ISE side.

Step 4. The friendly group name is assigned to the sponsor group.

ISE provides three default sponsor groups, each with access to a different scope of guest accounts:

- **ALL_ACCOUNTS:** Sponsor users authorized into this group can manage all guest accounts present on ISE. You can consider this group as a Super sponsor group.
- **GROUP_ACCOUNTS:** Users in this group can manage all guest accounts created by any other sponsor belonging to the same group. The most common use case for this user type is reception workers who need to create guest accounts for visitors and manage accounts created by their colleagues.

Technet24

- **OWN_ACCOUNTS:** Sponsor users in this group can manage only the guest accounts they have created. They do not have access to guest accounts created by other sponsor users. This group is mostly used for regular employees so they can create guest accounts for their own visitors.

Sponsor groups contain a wide variety of settings that determine what types of accounts the group can create and manage, as well as how sponsors can perform these actions. Some of the most important settings include

- **This Sponsor Group Can Create Accounts Using These Guest Types:** Here you define which specific guest types that members of the sponsor group can create.
- **Select the Locations That Guests Will Be Visiting:** In this setting, you specify the locations for which sponsor users can create guest accounts. By default, only the "San Jose" location is added, because it is the only default location present on ISE. You can add additional locations by navigating to **Work Centers > Guest Access > Configure > Guest Types**. The location is important for guest account duration. When guest accounts are created with **Account Duration Starts: From Sponsor-Specified Date**, the actual start time is based on the time zone in which the guest account was created, and the time zone must be configured inside the location.

Examples of both settings are shown in Figure 9-16.

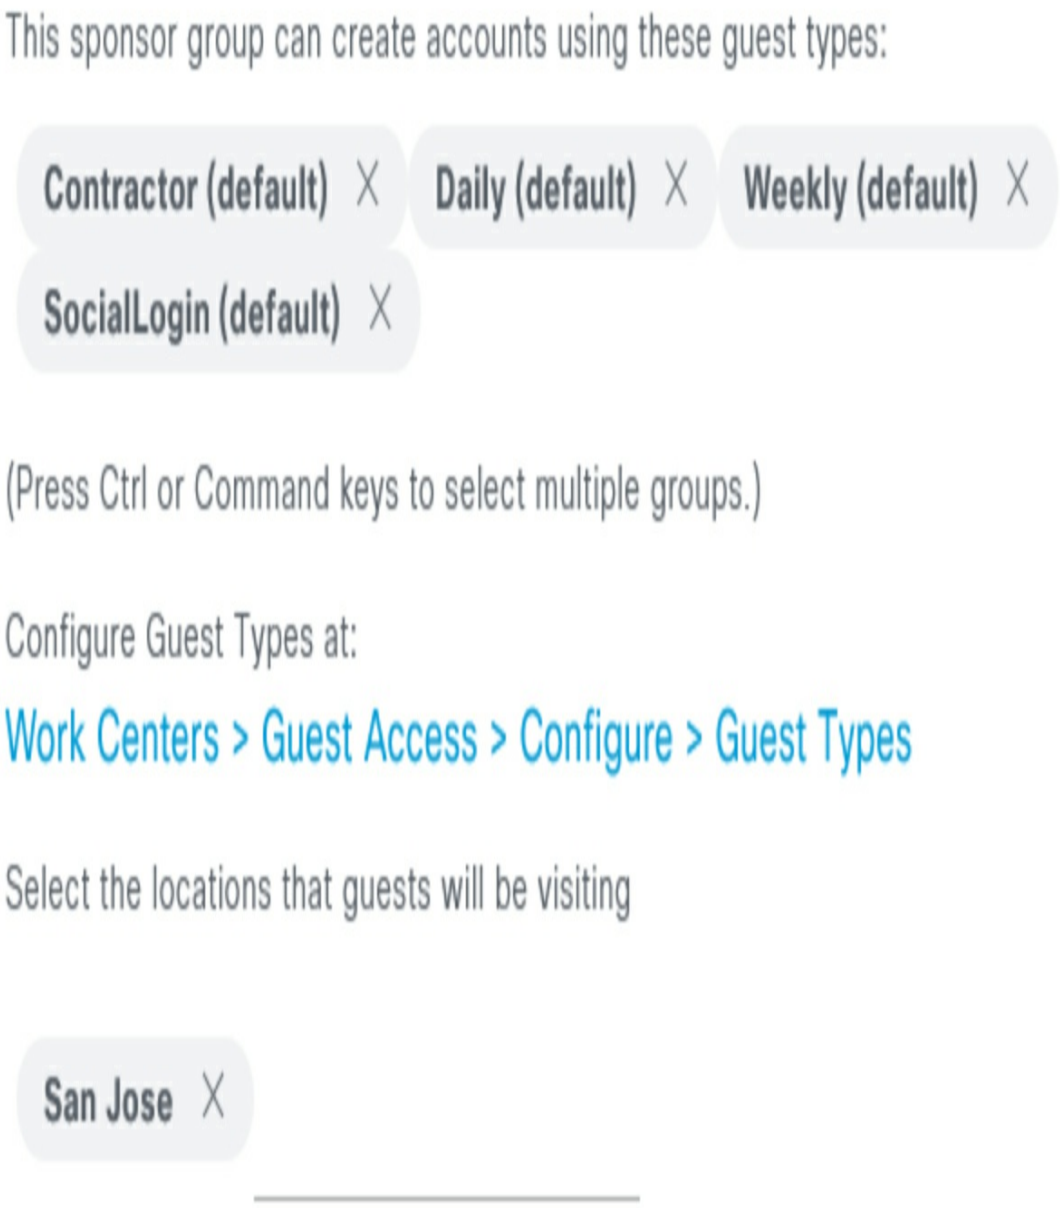


Figure 9-16 *Sponsor Group Settings*

Guest Users and Portals

From the perspective of guest users, two main configuration components exist, each depending on additional elements such as portal certificates, external identity stores, identity store sequences,

Technet24

and more. These components are guest portals and guest types.

We will first discuss guest portals. As previously explained, users are redirected to the guest portal as part of the initial network authentication. TLS is used to establish the connection with the portal, and the connection is protected by the certificate assigned to the guest portal using certificate group tags, as described earlier in the section "Sponsor Users and Portals."

Three types of guest portals are available on ISE:

- **Sponsored Guest Portal:** This portal type, by default, allows authentication for guest users and employees (internal ISE users or users from external identity sources). Guest users are assigned their guest type by sponsors. For employees, as explained earlier, the guest type is defined in the portal settings.
- **Self-Registered Guest Portal:** This is an extended version of the Sponsored Guest Portal, because it enables the option **Allow Guests to Create Their Own Accounts**, allowing guest users to register their own guest accounts with or without sponsor approval. You can configure the guest type for self-registered guests in the portal settings in the Registration Form Settings drop-down **Assign to Guest Type**, as demonstrated in Figure 9-17.

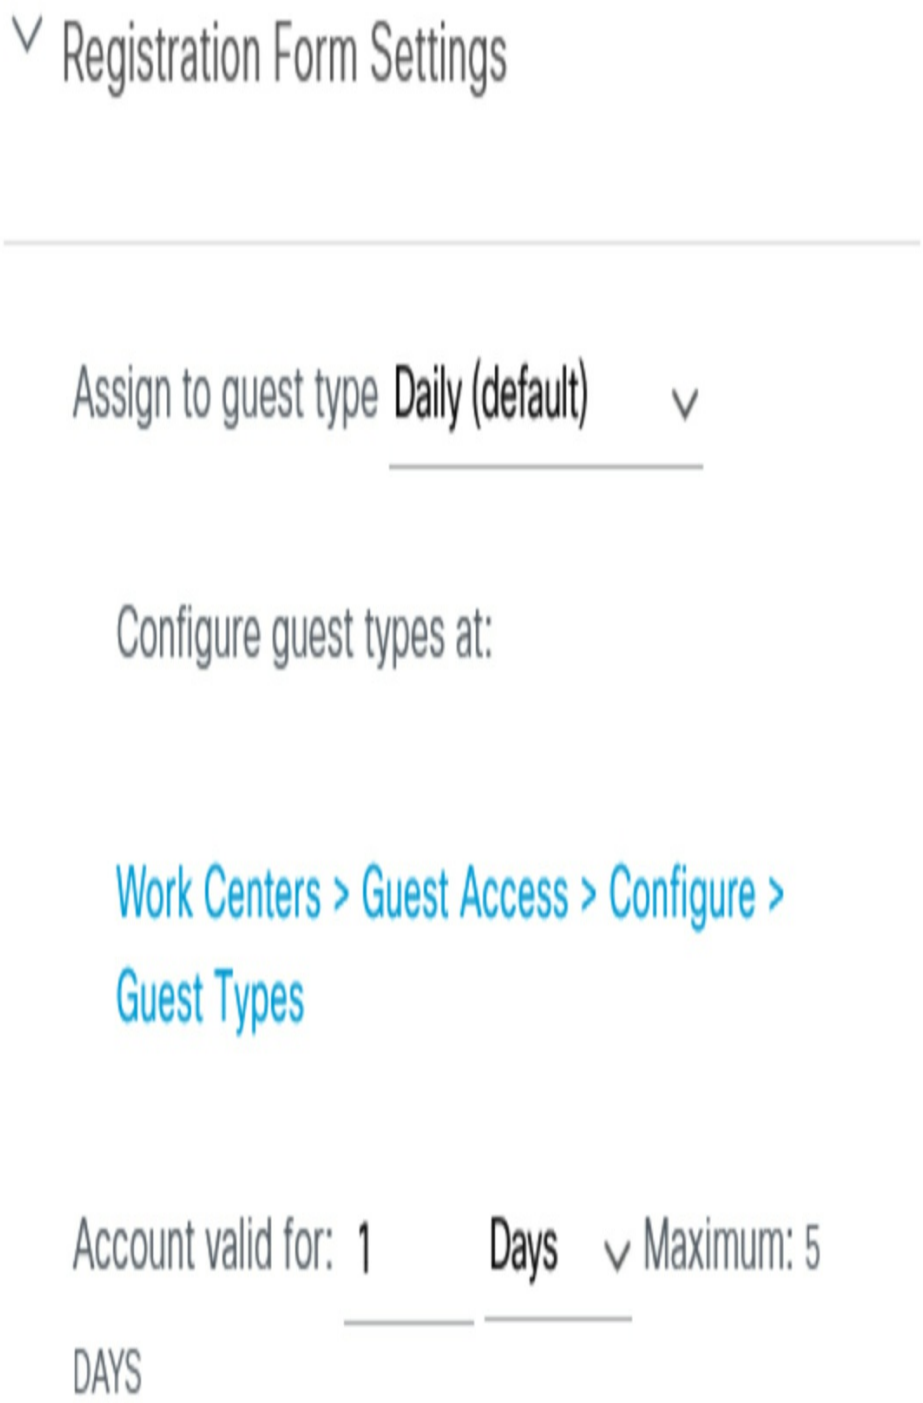


Figure 9-17 *Defining Guest Type for Self-Registered Guests*

- **Hotspot Guest Portal:** This portal type differs from the previous

Technet24

two because it does not offer any form of individual user authentication. You can use an access code, which remains the same for all users, as a simple authentication mechanism. When working with the hotspot portal, remember that this flow supports both CoA Reauthenticate and CoA Terminate, with CoA Reauthenticate being the default option. When deciding which CoA type is most suitable for your needs, review the earlier section "Stitching Attributes for Authorization." With CoA Reauthenticate, both session and endpoint attributes can be used in the final authorization policy, while with CoA Terminate, you are limited to endpoint attributes only, since CoA Terminate removes the original session. Figure 9-18 shows an example of the CoA settings on the Hotspot Guest Portal.

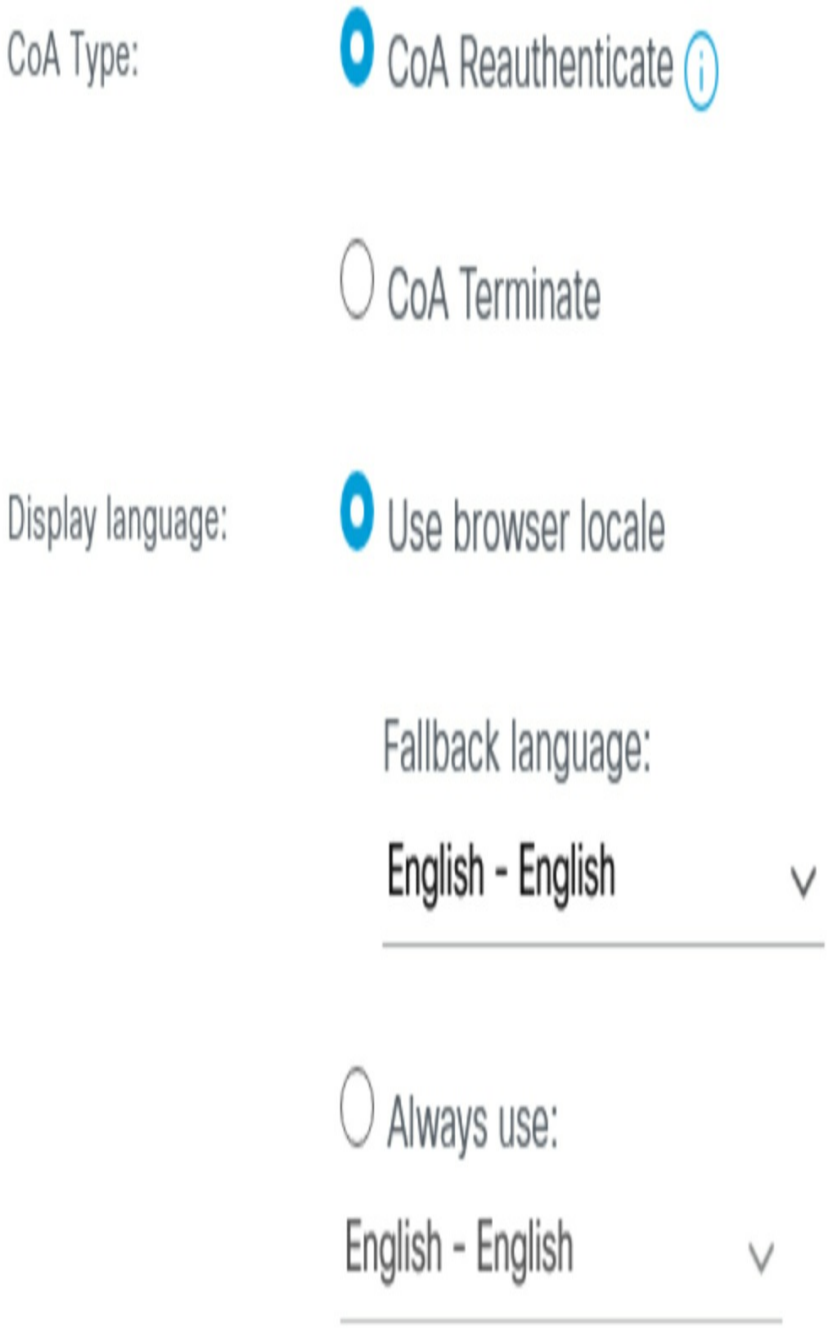


Figure 9-18 *Hotspot Portal CoA Settings*

When it comes to guest types, ISE provides four default guest types with settings predefined according to best practices. The following

Technet24

steps outline preconfigured guest types:

- **Contractor:** Guest user account with up to one year lifetime
- **Daily:** Short-term guest accounts valid for one to five days
- **Weekly:** Accounts designed for longer durations, lasting a couple of weeks
- **SocialLogin:** Guest type for users with the Facebook social login feature, with a lifetime of one to five days

In the guest type settings, the following important options are configurable:

- **The start of the account duration:** This defines how the account becomes active. An example is provided in Figure 9-19.



Figure 9-19 *Account Duration Start Configuration*

- **Account duration and time-of-day access:** In this area, you can configure the maximum duration for the given guest type account and specify time-of-day restrictions for guest user access. An example is provided in Figure 9-20.



Figure 9-20 *Guest Account Duration and Time-of-Day Access*

- **Maximum login and maximum devices:** Here you can specify how many simultaneous logins are allowed for the same guest user and how many endpoints can belong to the same guest account. The example is shown in Figure 9-21.

Technet24

Maximum simultaneous logins 3 (1-999)

When guest exceeds limit:

- Disconnect the oldest connection
- Disconnect the newest connection

Redirect user to a portal page showing an error message

This requires the creation of an authorization policy rule

Maximum devices guests can register: 5 (1-999)

Assign endpoint identity groups for guest device registration workflows:

Standard Dynamic

Identity Group

GuestEndpoints

Figure 9-21 *Maximum Sessions and Maximum Endpoints*

BYOD Building Blocks

ISE BYOD services allow end users to onboard personal devices for secure connectivity using single or dual SSID provisioning flows. Onboarded devices can be provisioned with network settings and a certificate for secure EAP-TLS authentication. Users also gain the ability to manage their devices through the My Devices portal. The main BYOD components on ISE include the following:

- **Endpoints:** The consumer devices that undergo BYOD onboarding on ISE. Depending on the device type, ISE provides different sets of BYOD capabilities. The device type also dictates how the BYOD flow occurs. Supported devices include Windows, macOS, iOS, Android, and Chromebooks.
- **Client provisioning:** This component detects the correct device type using the "Cisco supported OS version" package and assigns the appropriate client provisioning policy based on device OS and other parameters such as user identity group. Each client provisioning policy determines what configuration should be provisioned to the endpoint and how.
- **Certificate management:** A major benefit of BYOD is the ability to provision a device with an identity certificate with minimal complexity for the end user. ISE supports certificate provisioning from both the internal CA and external CAs using SCEP or EST.
- **Portals:** Similar to Guest Access, BYOD services rely on CWA to redirect users or applications to ISE BYOD portals for required configuration provisioning. BYOD also includes a special portal called My Devices portal, allowing users to manage onboarded devices.

Technet24

- **Policy and policy components:** Like other features on ISE, BYOD services depend on authentication and authorization, which require proper setup of ISE policies and their components.

Types of BYOD Flows

When ISE BYOD flows are discussed, you often encounter terms like *single SSID BYOD* flow and *dual SSID BYOD*, which may prompt the question: Is BYOD only for wireless connections? The answer is no. BYOD services can be used for wired, wireless, and VPN connections. However, the most common BYOD use case involves wireless connectivity, which is why the terminology often references *SSIDs*. Before exploring each flow type in detail, let's consider a high-level overview of them in the context of any connection method, not just wireless.

- **Single SSID BYOD flow:** When you hear about a single SSID BYOD flow, it usually means the user initially connects to the network using a strong authentication method, such as password-based EAP (for example, EAP-PEAP MSCHAPv2) for wired or wireless access, or username/password authentication for VPN. In this scenario, ISE already recognizes the user and has details such as group membership, allowing ISE to perform a redirect directly to the Native Supplicant Provisioning (NSP) portal for BYOD-assisted registration and configuration provisioning. After BYOD registration, for wired and wireless connections, the endpoint simply reconnects using a stronger authentication method, such as EAP-TLS. For VPN connections, this flow can be used to perform BYOD device registration and assist with certificate provisioning for future VPN access.
- **Dual SSID BYOD flow:** In this case, the process starts with no authentication, because the device is typically connected to guest

Wi-Fi or authenticated over MAB in the wired network. This approach is generally not used for VPN connections. The device connected to the guest network is redirected to the guest portal, where you, as an administrator, can enable BYOD flow for employees. As explained earlier in the chapter, employees from the perspective of guest access are any non-guest users, such as users in ISE internal databases or those in external identity sources. After authentication on the guest portal, BYOD-assisted registration and configuration provisioning occurs. In a Wi-Fi environment, the endpoint reconnects to the more secure SSID where 802.1X authentication is performed, typically using EAP-TLS. For wired access, the endpoint starts dot1x authentication using the EAP method configured as part of the BYOD process.

Each BYOD flow can be divided into three phases:

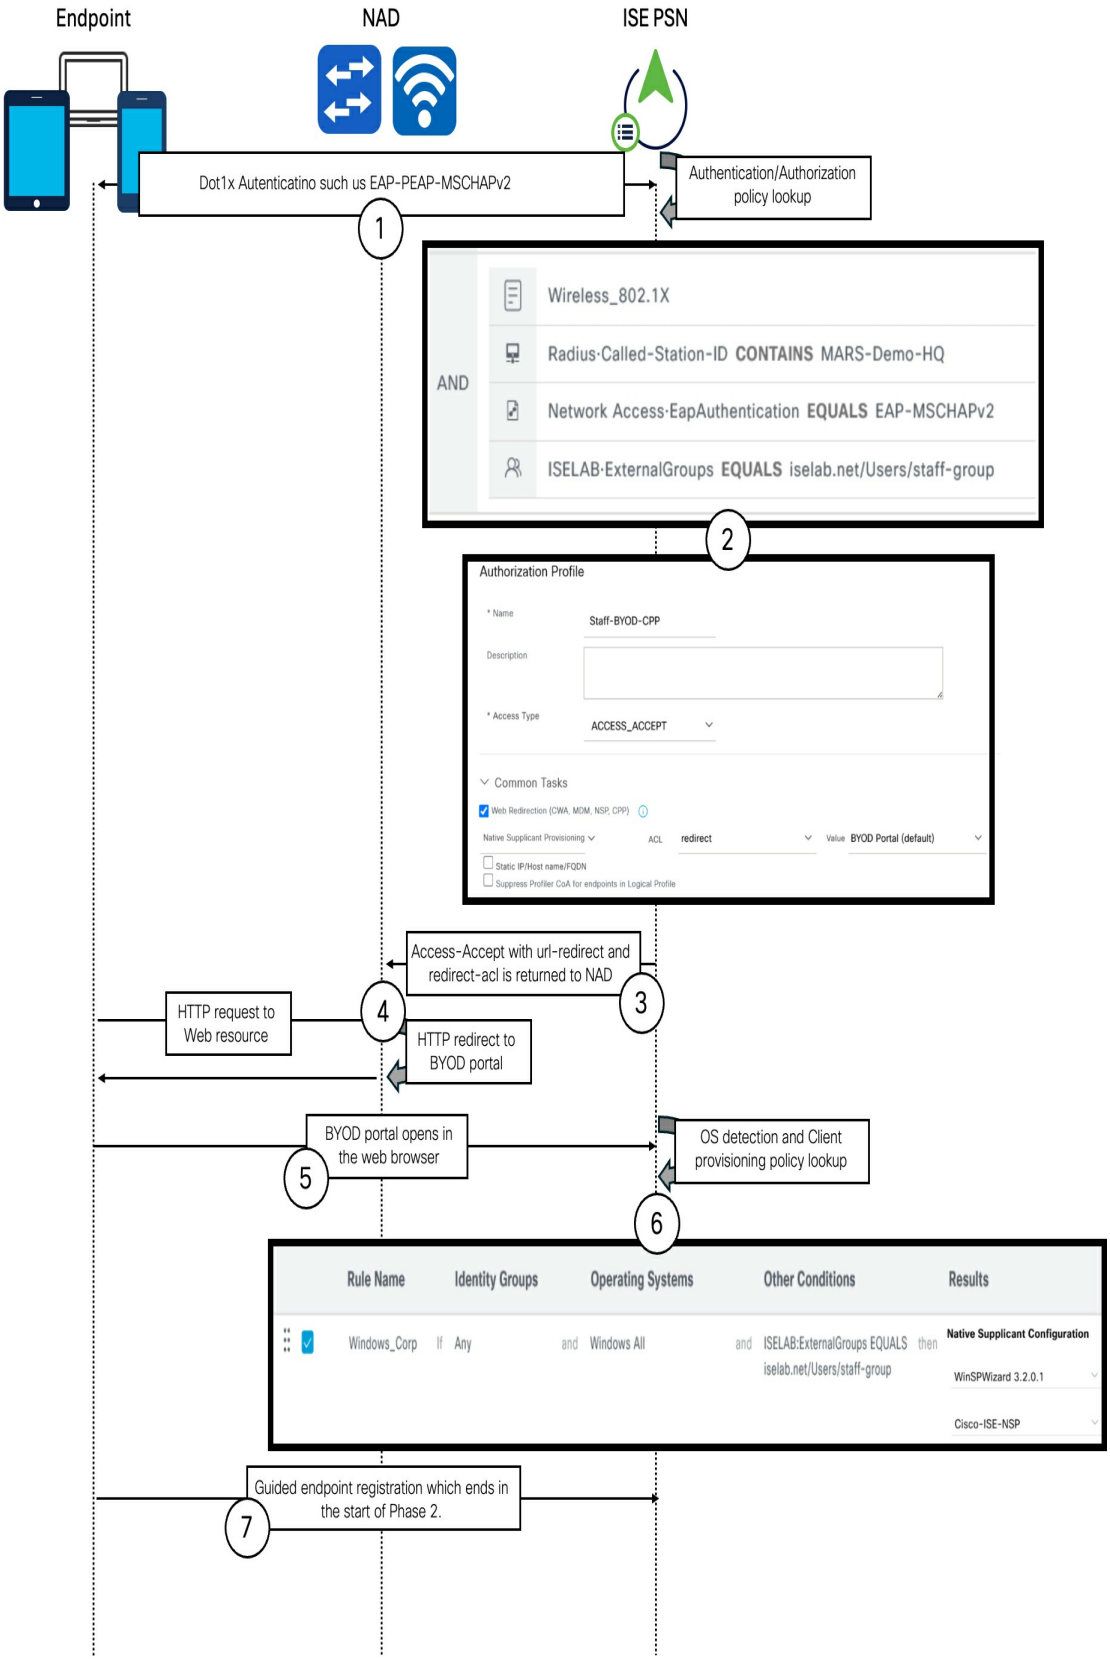
- **Phase 1:** Initial authentication and client provisioning. During this phase, as a result of the initial authentication, the endpoint is redirected to the BYOD or guest portal depending on the flow type, where client provisioning and endpoint registration take place.
- **Phase 2:** Native supplicant setup. This phase varies for different supported operating systems and device types. The high-level summary is that native supplicant configuration is performed using either a software package from ISE, built-in OS software, or a package from the relevant application store. This software package handles the supplicant configuration for Phase 3, including setting up the network profile and provisioning the identity certificate from ISE, which is later used for EAP-TLS authentication.
- **Phase 3:** Final BYOD authentication. At this stage, the endpoint authenticates on the network using the new settings and receives

Technet24

a less restrictive authorization policy assigned by ISE.

Single SSID BYOD Flow

Now let's examine the different phases of the single SSID BYOD flow. Phase 1 is represented in Figure 9-22.


Endpoint
NAD
ISE PSN
Dot1x Autenticatino such us EAP-PEAP-MSCHAPv2
1
Authentication/Authorization policy lookup
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
ISELAB·ExternalGroups EQUALS iselab.net/Users/staff-group
2
Authorization Profile
* Name
Staff-BYOD-CPP
Description
* Access Type
ACCESS_ACCEPT
Common Tasks
Web Redirection (CWA, MDM, NSP, CPP)
Native Supplicant Provisioning
ACL
redirect
Value
BYOD Portal (default)
Static IP/Host name/FQDN
Suppress Profiler CoA for endpoints in Logical Profile
Access-Accept with url-redirect and redirect-acl is returned to NAD
3
HTTP request to Web resource
4
HTTP redirect to BYOD portal
BYOD portal opens in the web browser
5
OS detection and Client provisioning policy lookup
6
Rule Name
Identity Groups
Operating Systems
Other Conditions
Results
Windows_Corp
If
Any
and
Windows All
and
ISELAB:ExternalGroups EQUALS iselab.net/Users/staff-group
then
Native Supplicant Configuration
WinSPWizard 3.2.0.1
Cisco-ISE-NSP
Guided endpoint registration which ends in the start of Phase 2.
7

Technet24

Figure 9-22 *Phase 1 of a Single SSID BYOD Flow*

Phase 1 of the single SSID BYOD flow consists of the following steps:

Step 1. Initial authentication occurs between ISE and the supplicant. In wired or wireless use cases, this is typically a password-based 802.1X authentication such as EAP-PEAP-MSCHAPv2.

Step 2. Authorization takes place on ISE. The authorization policy in the example is based on a wireless use case. For wired use cases, you need to change the default compound condition to Wired_802.1X and remove the SSID name check (Radius Called Station ID). The authorization profile in this example includes a redirection to the default BYOD portal.

Step 3. ISE returns Access-Accept to the NAD with redirect attributes.

Step 4. The NAD redirects the endpoint to the ISE BYOD portal.

Step 5. The ISE BYOD portal loads in the web browser.

Step 6. As part of Step 5, ISE receives the User-Agent from the browser and uses it to detect the OS of the endpoint. Once the OS is identified, this information is included in the client provisioning policy lookup process along with session cache attributes such as username, user group membership, and connection type. All these attributes can be used to select the appropriate client provisioning policy. The policy in the example checks for user group membership and, as a result, assigns the Native Supplicant provisioning app and network configuration profile.

Step 7. The user registers the endpoint on the BYOD portal and

proceeds to Phase 2.

Phase 2 behavior depends on the endpoint and operating system used. Figure 9-23 provides a summary of the flows for the supported endpoints.

Technet24

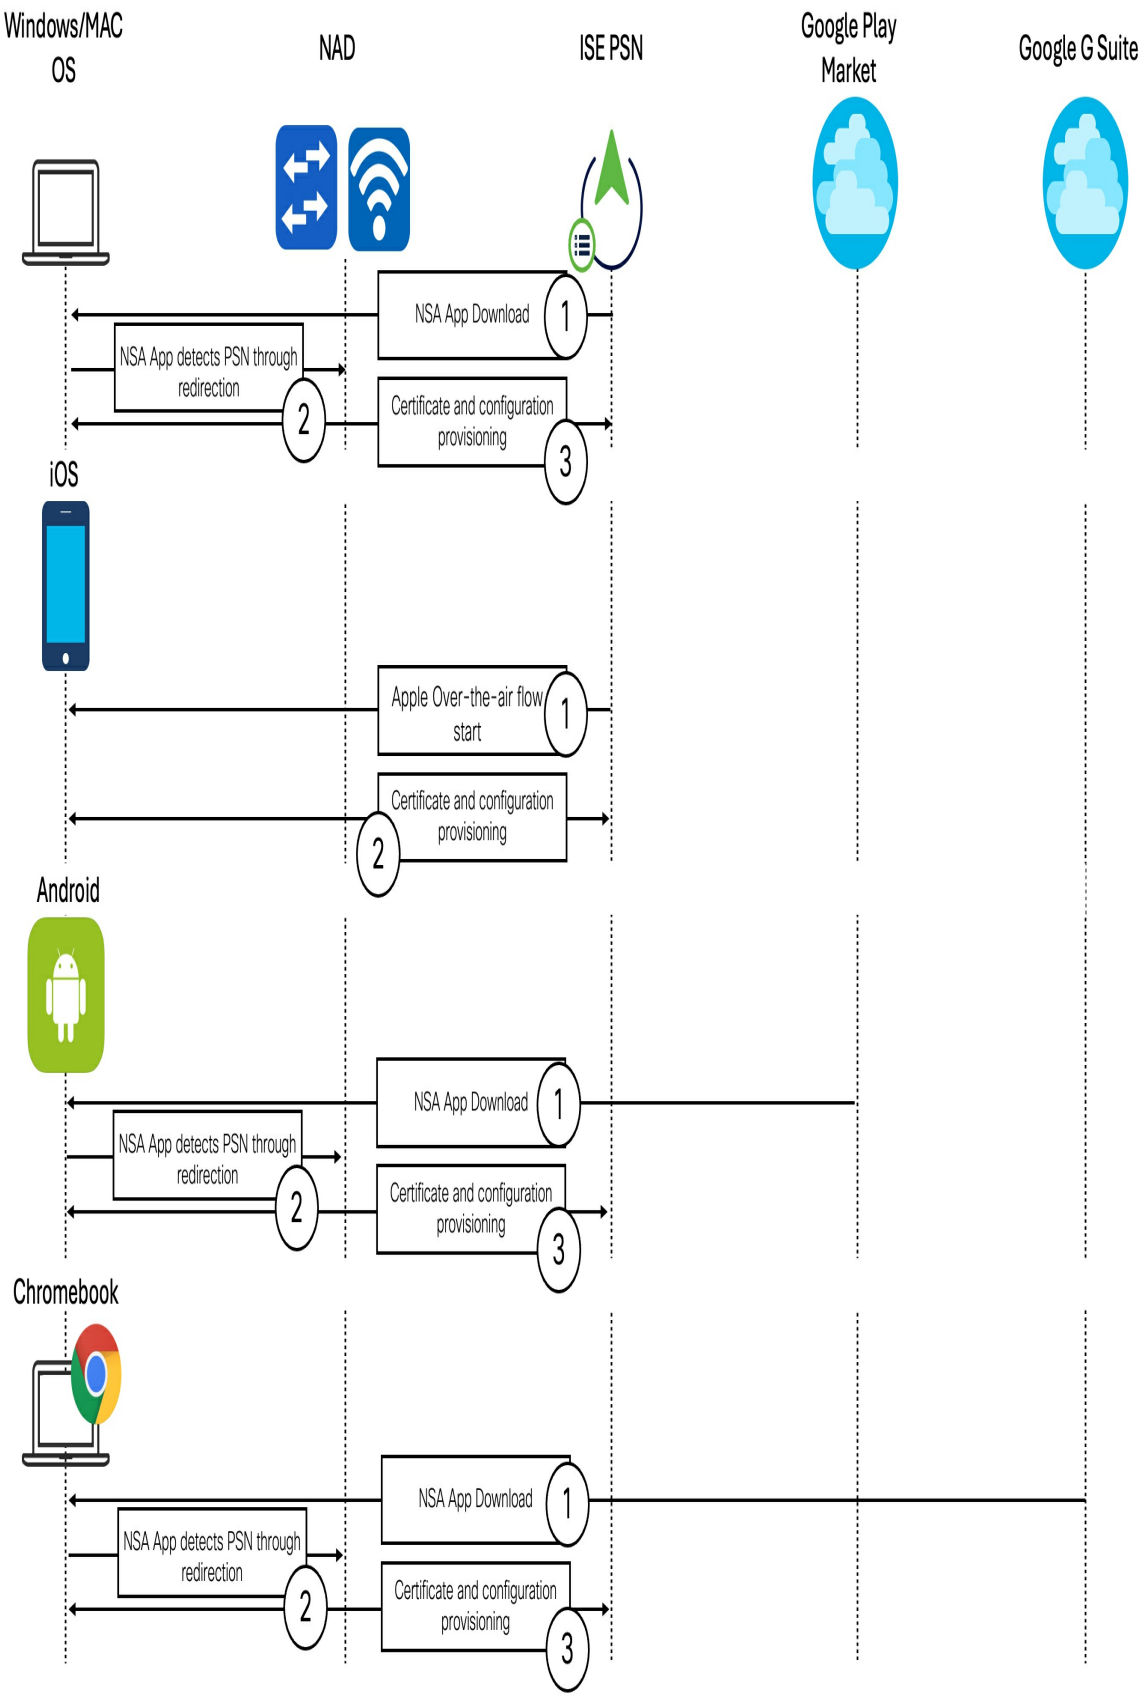

Windows/MAC OS
NAD
ISE PSN
Google Play Market
Google G Suite
NSA App Download
1
NSA App detects PSN through redirection
2
Certificate and configuration provisioning
3
iOS
Apple Over-the-air flow start
1
Certificate and configuration provisioning
2
Android
NSA App Download
1
NSA App detects PSN through redirection
2
Certificate and configuration provisioning
3
Chromebook
NSA App Download
1
NSA App detects PSN through redirection
2
Certificate and configuration provisioning
3

Figure 9-23 *Phase 2 of a Single SSID BYOD Flow*

The following steps outline Phase 2:

Steps for All OS endpoints except iOS:

Step 1. Network Setup Assistant (NSA) app download. In this step, the endpoint downloads the NSA software, which assists the user with the onboarding process.

For Windows and macOS, the NSA app is downloaded directly from ISE.

For Android devices, the app must be installed from the Google Play Store. ISE provides a link to the app on the BYOD portal. The challenge with Android devices is the requirement to provide endpoint-restricted Internet access during the redirect flow. This is typically achieved by specifying exceptions in the Redirect ACL using FQDNs instead of IP addresses.

For Chromebooks, the device should already be registered in G Suite before the BYOD process begins. G Suite must also be configured to provision the NSA app to the Chromebook.

Step 2. Redirection. NSA uses its own probes, which are HTTP requests to the default gateway IP and HTTP requests to enroll.cisco.com, to obtain the redirect URL from the NAD.

Step 3. **(All OS endpoints except iOS)** NSA communicates with the PSN using the URL obtained in Step 2 to retrieve the network configuration and send the CSR for certificate signing by ISE.

Steps for iOS devices:

Technet24

Step 1. **(iOS)** For iOS devices, ISE uses the native iOS feature called Apple Over-the-Air (OTA), which allows certificate provisioning and network configuration. ISE triggers OTA directly from the BYOD portal by pushing the OTA profile to the endpoint.

Step 2. **(iOS)** OTA communication with ISE occurs, during which a CSR is generated and signed, and network settings are provisioned to the endpoint.

Step 3. **(iOS)** Not applicable.

Phase 2 occurs only when the client provisioning lookup results in a policy with NSP settings. If the policy match does not happen, the user receives an error on the portal, preventing the flow from continuing. In some situations, you may need to register devices as BYOD assets even when configuration and certificate provisioning are impossible or not preferred. In such cases, you can allow network access when the client provisioning policy is not selected. To do this, navigate to **Work Centers > BYOD > Settings > Client Provisioning** and select **Allow Network Access** in the *** Native Supplicant Provisioning Policy Unavailable:** drop-down. Figure 9-24 provides an example of this page.

Client Provisioning

* Enable Provisioning: Enable

* Enable Automatic Download: Disable

* Update Feed URL: https://www.cisco.com/web/secure/spa/p

* Native Supplicant Provisioning Policy Unavailable: Apply Defined Authorization Policy

Allow Network Access

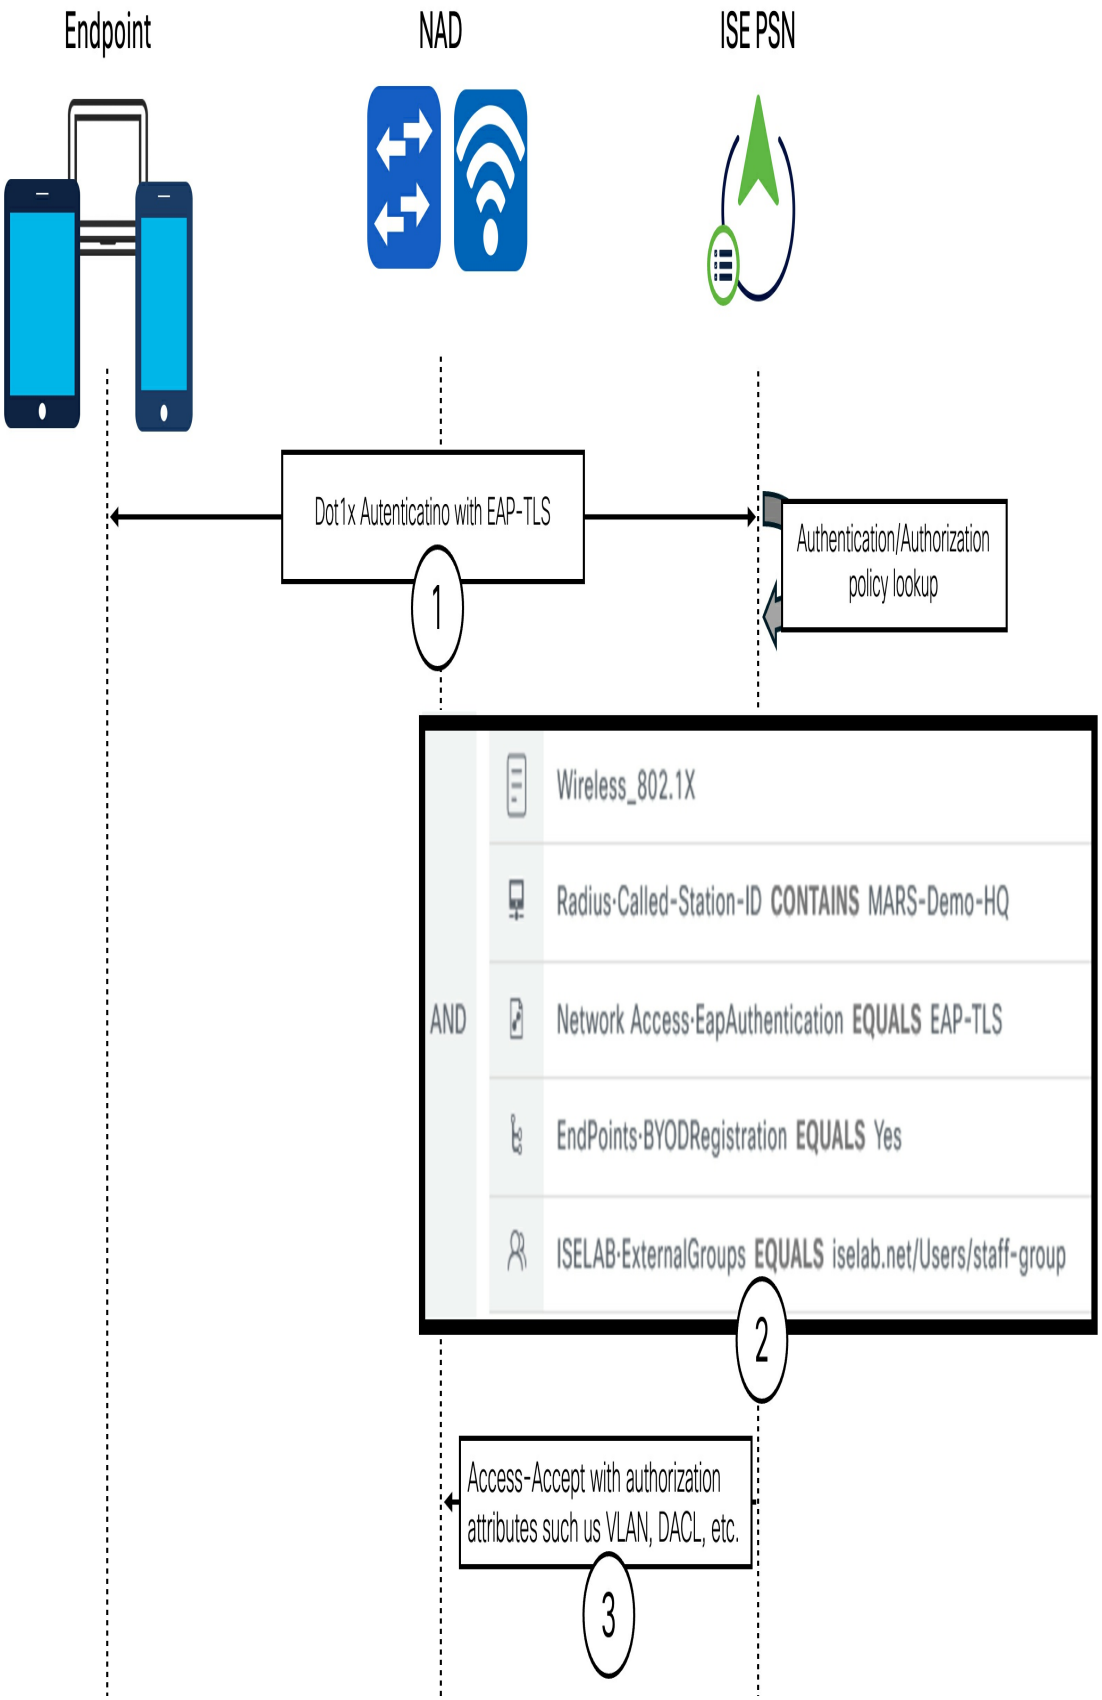
Apply Defined Authorization Policy

Technet24

Figure 9-24 *Allow Network Access When Native Supplicant Provisioning Policy Unavailable*

When Allow Network Access is enabled and the endpoint does not match a client provisioning policy, you can use a BYOD registration flag or the default endpoint group Registered Devices in your subsequent authorization policy. Users will also have access to their devices through the My Devices portal.

Finally, the flow concludes with Phase 3. The steps of Phase 3 are outlined in Figure 9-25.


Endpoint
NAD
ISE PSN
Dot1x Autenticatino with EAP-TLS
1
Authentication/Authorization policy lookup
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
Network Access·EapAuthentication EQUALS EAP-TLS
EndPoints·BYODRegistration EQUALS Yes
ISELAB·ExternalGroups EQUALS iselab.net/Users/staff-group
2
Access-Accept with authorization attributes such us VLAN, DACL, etc.
3

Technet24

Figure 9-25 *Phase 2 of a Single SSID BYOD Flow*

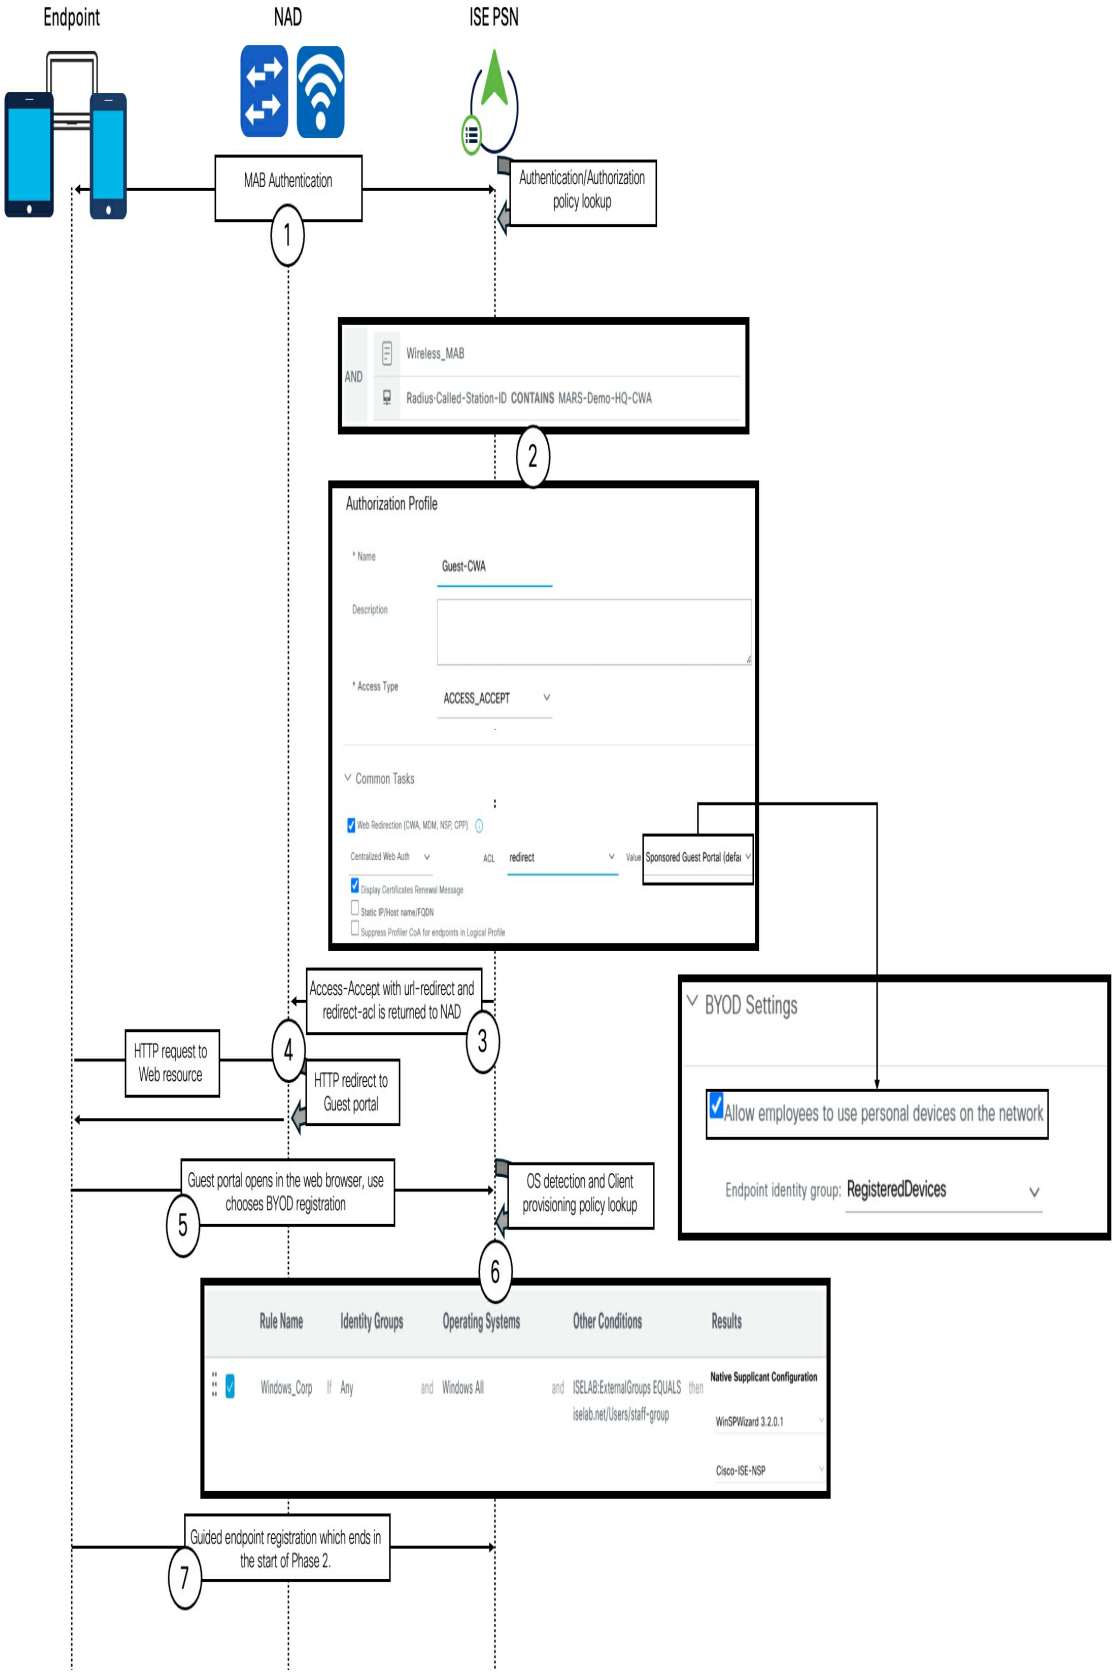
Step 1. The new 802.1X authentication is performed using EAP-TLS.

Step 2. ISE conducts an authorization policy lookup. The example contains the policy for the wireless use case. For wired use cases, you need to change the default compound condition to Wired_802.1X and remove the SSID name check (Radius Called Station ID).

Step 3. ISE returns authorization attributes to the NAD as part of the Access-Accept. Typically, endpoints may be placed in a different VLAN and receive a new DACL assignment.

Dual SSID BYOD Flow

Phase 1 of the dual SSID BYOD flow is represented in Figure 9-26.


Endpoint
NAD
ISE PSN
MAB Authentication
1
Authentication/Authorization policy lookup
AND
Wireless_MAB
Radius-Called-Station-ID CONTAINS MARS-Demo-HQ-CWA
2
Authorization Profile
* Name
Guest-CWA
Description
* Access Type
ACCESS_ACCEPT
Common Tasks
Web Redirection (CWA, MDM, NSP, CPP)
Centralized Web Auth
ACL
redirect
Value
Sponsored Guest Portal (defa
Display Certificates Renewal Message
Static IP/Host name/FQDN
Suppress Profiler CoA for endpoints in Logical Profile
Access-Accept with url-redirect and redirect-acl is returned to NAD
3
4
HTTP request to Web resource
HTTP redirect to Guest portal
BYOD Settings
Allow employees to use personal devices on the network
Endpoint identity group: RegisteredDevices
Guest portal opens in the web browser, use chooses BYOD registration
5
OS detection and Client provisioning policy lookup
6
Rule Name
Identity Groups
Operating Systems
Other Conditions
Results
Windows_Corp
If
Any
and
Windows All
and
ISELAB:ExternalGroups EQUALS iselab.net/Users/staff-group
then
Native Supplicant Configuration
WinSPWizard 3.2.0.1
Cisco-ISE-NSP
Guided endpoint registration which ends in the start of Phase 2.
7

Technet24

Figure 9-26 *Phase 1 of a Dual SSID BYOD Flow*

Phase 1 of the Dual SSID BYOD flow consists of the following steps:

Step 1. Initial authentication. This is typically performed using MAB and is similar to the initial authentication in the guest flow.

Step 2. Authorization occurs on ISE. For the dual SSID BYOD flow, ISE does not have much detail about the endpoint or user during initial authentication. As in the guest flow, it is important to ensure that the ISE authentication policy uses a continue action for the "user not found" scenario, so authorization can proceed for unknown MAC addresses. The authorization policy will typically include a compound MAB condition for wired or wireless connections, depending on the connection type. Additionally, for wireless connections, you may include the Radius-Called-Station-ID condition to match the SSID name. The authorization profile should contain Central Web Authentication with a redirect to one of the guest portals. The corresponding guest portal needs to have the **Allow Employees to Use Personal Devices on the Network** checkbox enabled in the **BYOD** section.

Step 3. ISE returns Access-Accept to the NAD with redirect attributes.

Step 4. The NAD redirects the endpoint to the ISE guest portal.

Step 5. The ISE guest portal loads in the web browser.

Step 6. As part of Step 5, ISE receives the User-Agent from the browser and uses it to detect the OS of the endpoint. After the end user authenticates on the guest portal, ISE has enough information to determine if this is a guest user or

an employee. The employee can then start the BYOD flow. Once the BYOD flow begins, the endpoint OS information combines with other session attributes such as username, user group membership, and connection type. All these attributes can be used to select the appropriate client provisioning policy. The policy in the example checks for user group membership and, as a result, assigns the Native Supplicant provisioning app and network configuration profile.

Step 7. The user registers the endpoint on the guest portal and proceeds to Phase 2.

Phase 2 of the dual SSID BYOD flow matches the process used in the single SSID scenario.

Phase 3 of the dual SSID BYOD flow is mostly the same as in the single SSID flow, with the distinction that after BYOD configuration provisioning, the endpoint must either connect to the new SSID using WPA2-Enterprise with EAP-TLS authentication for wireless connections, or the NIC is reconfigured with 802.1X and EAP-TLS for wired connections. Figure 9-27 provides an example of Phase 3 for the dual SSID BYOD flow.

Technet24

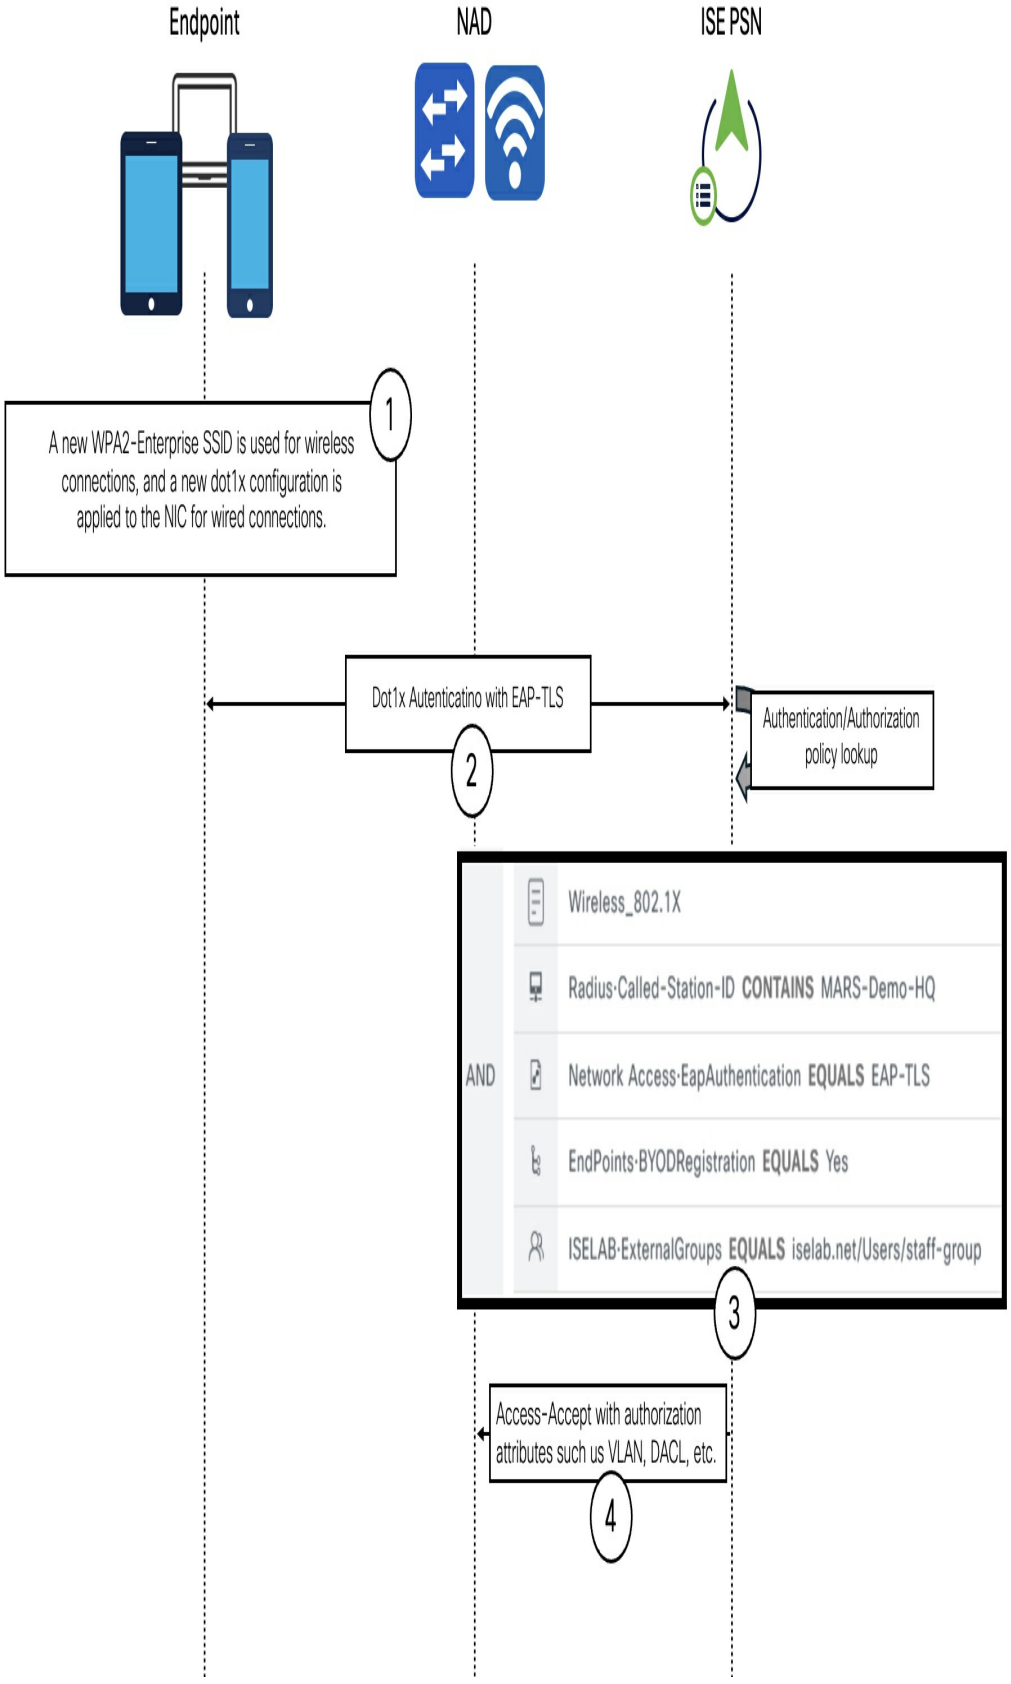

Endpoint
NAD
ISE PSN
A new WPA2-Enterprise SSID is used for wireless connections, and a new dot1x configuration is applied to the NIC for wired connections.
1
Dot1x Autenticatino with EAP-TLS
2
Authentication/Authorization policy lookup
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
Network Access·EapAuthentication EQUALS EAP-TLS
EndPoints·BYODRegistration EQUALS Yes
ISELAB·ExternalGroups EQUALS iselab.net/Users/staff-group
3
Access-Accept with authorization attributes such us VLAN, DACL, etc.
4

Figure 9-27 *Phase 3 of a Dual SSID BYOD Flow*

ISE BYOD Components

In this section, we will review the following ISE BYOD components in more detail:

- Client provisioning
- Certificate management
- BYOD portals

Endpoints and policies are not covered here, because they were explained in detail in the previous section.

Client Provisioning

The Client Provisioning component of the BYOD service is responsible for detecting the endpoint and operating system type, and providing the required software and configuration to the endpoint.

As previously explained, endpoint type and OS detection occur when the Guest or BYOD portal is opened by the user as a result of web redirection. ISE uses a specialized OS fingerprint database, which you should update regularly to ensure support for all new endpoints and operating systems. You can check the current version of the OS fingerprint database by navigating to **Administration > System > Settings > Posture > Updates**. The Cisco supported OS version indicates the version of the OS fingerprint database. On the same page, you can enable automatic periodic updates to ensure that OS fingerprints remain current.

BYOD client provisioning resources can be divided into two categories:

- **NSA software:** As explained earlier, this category is approached

Technet24

differently depending on the endpoint and OS type.

- **Native Supplicant Profile:** This configuration file contains the settings that should be applied to the device as part of BYOD onboarding. It includes certificate provisioning details and the actual network settings to be used after successful onboarding.

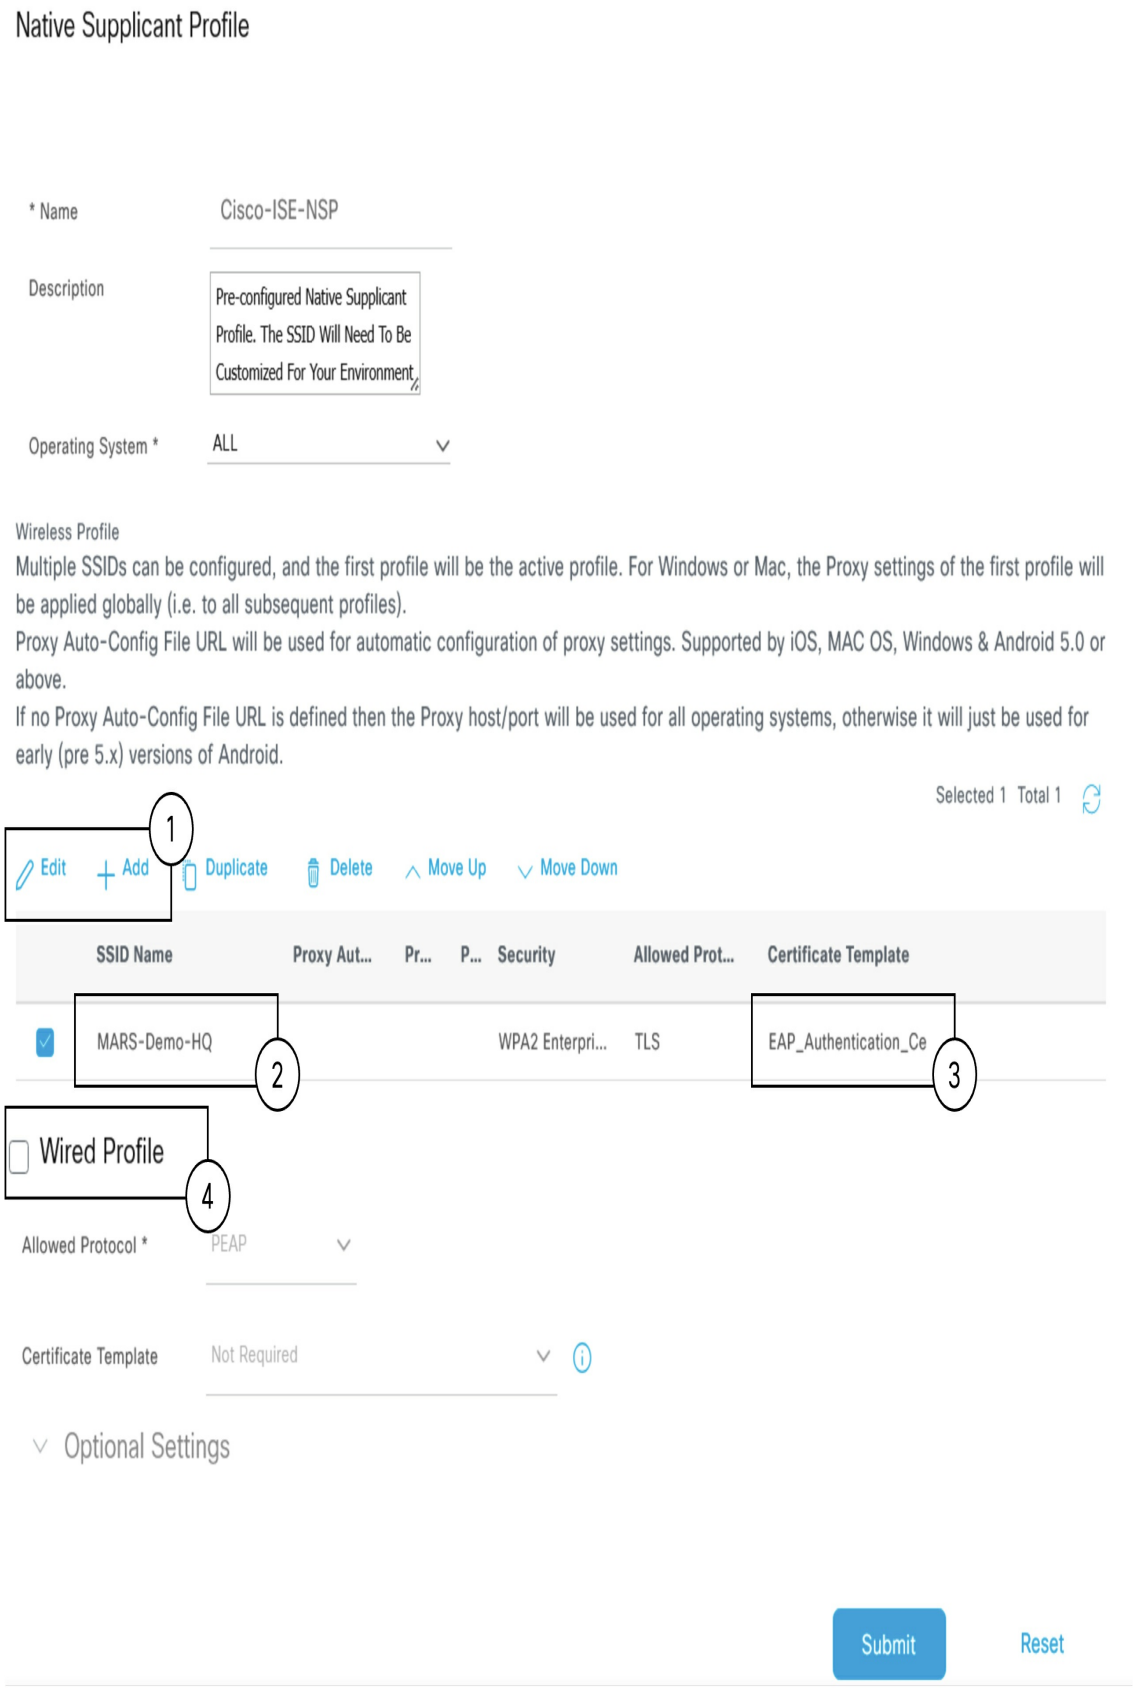
To access and manage BYOD client provisioning, navigate to **Work Centers > BYOD > Client Provisioning > Resources**. By clicking **Add**, as demonstrated in Figure 9-28, you can add agent resources from the Cisco site to download the latest version of Supplicant Provisioning Wizard (SPW) software (NSA package for Windows and macOS) or add a Native Supplicant Profile.


Overview
Identities
Identity Groups
Network Devices
Ext Id Sources
Client Provisioning
Portals & Components
More
Client Provisioning Policy
Resources
Resources
Selected 0 Total 18
Edit
Add
Duplicate
Delete
Export
Quick Filter
Agent resources from Cisco site
Agent resources from local disk
Native Supplicant Profile
Agent Configuration
Agent Posture Profile
Type
MacOsXSPWizard
WinSPWizard
CiscoSecureClientDesktopWindows

Technet24

Figure 9-28 *BYOD Client Provisioning Resources*

As an administrator, you can edit and reuse the existing Native Supplicant Profile called Cisco-ISE-NSP or create a new one. Figure 9-29 demonstrates an example of the NSP profile.


Native Supplicant Profile
* Name
Cisco-ISE-NSP
Description
Pre-configured Native Supplicant Profile. The SSID Will Need To Be Customized For Your Environment.
Operating System *
ALL
Wireless Profile
Multiple SSIDs can be configured, and the first profile will be the active profile. For Windows or Mac, the Proxy settings of the first profile will be applied globally (i.e. to all subsequent profiles).
Proxy Auto-Config File URL will be used for automatic configuration of proxy settings. Supported by iOS, MAC OS, Windows & Android 5.0 or above.
If no Proxy Auto-Config File URL is defined then the Proxy host/port will be used for all operating systems, otherwise it will just be used for early (pre 5.x) versions of Android.
Selected 1 Total 1
1
Edit
Add
Duplicate
Delete
Move Up
Move Down
SSID Name
Proxy Aut...
Pr...
P...
Security
Allowed Prot...
Certificate Template
MARS-Demo-HQ
2
WPA2 Enterpri...
TLS
EAP_Authentication_Ce
3
Wired Profile
4
Allowed Protocol *
PEAP
Certificate Template
Not Required
Optional Settings
Submit
Reset

Technet24

Figure 9-29 *Native Supplicant Profile Configuration*

The following outlines the main steps in the Native Supplicant Profile configuration:

Step 1. You can edit an existing wireless profile or create a new one.

Step 2. Each wireless profile requires the SSID name to be specified.

Step 3. For EAP-TLS connections, you need to select the certificate template. By default, ISE provides a preconfigured certificate template called EAP_Authentication_Certificate_Template, which is set to obtain the client certificate from the ISE internal CA.

Step 4. If necessary, you can configure the wired profile as well.

Certificate Management

When a client certificate needs to be provisioned as part of the BYOD process, ISE supports two approaches for this:

- **ISE internal CA:** The certificate can be issued by the ISE internal CA, which is considered the preferred approach because it simplifies the setup.
- **External CA:** ISE can use the SCEP registration authority (RA) functionality of an existing external CA to obtain a user certificate. In this scenario, ISE acts as a proxy: It receives the certificate request from the client, forwards it to the preconfigured CA, and then returns the certificate to the client.

You can create the external CA profile by navigating to **Work Centers > BYOD > Portals & Components > Certificates > External CA**

Templates and clicking **Add**. Figure 9-30 shows an example of the external CA profile configuration. Before configuring the profile, ensure that the CA is properly set up for the SCEP RA role.

Technet24

Add SCEP RA Profile

* Name: Win-CA-SCEP

Description:

* URL:

https://172.22.15.11/certsr

Test Connection

Submit

Figure 9-30 *External CA Profile Configuration*

You can configure certificate settings by editing an existing template or creating a new one. Certificate templates are found by navigating to **Work Centers > BYOD > Portals & Components > Certificates > Certificate Templates**. The templates include a wide range of certificate settings. When you configure a template for an external CA, ensure that the settings added to the template are supported on the external CA side; otherwise, certificate provisioning will fail.

Portals

The BYOD service on ISE uses three types of portals to facilitate device onboarding and assist users with self-management of onboarded endpoints:

- **BYOD portal:** This portal is used during the single SSID BYOD flow. The portal does not require many additional settings except for assigning the proper certificate using certificate group tags and adjusting the portal's appearance to match company requirements. Users do not need to authenticate on this portal, because user identification occurs during the initial 802.1X authentication in the single SSID BYOD flow.
- **Guest portal:** Sponsored and self-registered guest portals can be configured to enable BYOD service for employees. Users must authenticate on these portals so ISE can determine whether the user is a guest or an employee. Only employee users can register and onboard their devices using the BYOD service.
- **My Devices portal:** Employees can use this portal after device onboarding to manage their registered devices. The portal requires authentication and, similar to the sponsor portal, typically has an FQDN assigned for easier access.

Technet24

To edit an existing BYOD portal or create a new one, navigate to **Work Centers > BYOD > Portals & Components > BYOD Portals**. By default, one BYOD portal is available on ISE. Figure

To adjust portal settings for the dual SSID BYOD flow, navigate to **Work Centers > Guest Access > Portals & Components > Guest Portals** and edit the portal you intend to use. Employee BYOD settings are located in the **BYOD** section. In this area, you can configure the following options:

- **Allow Employees to Use Personal Devices on the Network:** This setting enables the BYOD flow for employees.
- **Endpoint Identity Group:** The drop-down specifies which endpoint identity group should be used to place BYOD registered devices.
- **Allow Employees to Choose Guest Access Only:** Employees can opt out of the BYOD flow and continue with regular guest Internet access.
- **Display Device ID Field During Registration:** This option shows device details during registration.

Figure 9-31 demonstrates BYOD settings for the guest portal.

BYOD Settings

☑ Allow employees to use personal devices on the network

Endpoint identity group:

RegisteredDevices

Configure endpoint identity groups at

Administration > Identity Management > Groups > Endpoint Identity Groups

The endpoints in this group will be purged according to the policies defined in:

Administration > Identity Management > Settings > Endpoint purge

☑ Allow employees to choose to guest access only

☐ Display Device ID field during registration

Technet24

Figure 9-31 *Guest Portal BYOD Settings*

To configure the My Devices portal, navigate to **Work Centers > BYOD > Portals & Components > My Devices Portal**. The most important settings for My Devices portal are as follows:

- **Certificate group tag:** As with guest and sponsor portals, certificates are assigned to the portals using portal tags. The tag must be assigned to the ISE identity certificate, and the same tag should be assigned to the portal for certificate changes to take effect. Remember that all portals on a single TCP port always use the same certificate.
- **Fully qualified domain names and hostnames:** As with the sponsor portal, you can define a user-friendly FQDN for users to access the portal. My Devices portals run on all ISE PSN nodes.
- **Authentication method:** Select the identity store sequence that will be used to authenticate users when they attempt to access the My Devices portal.

Figure 9-32 demonstrates the My Devices portal settings.

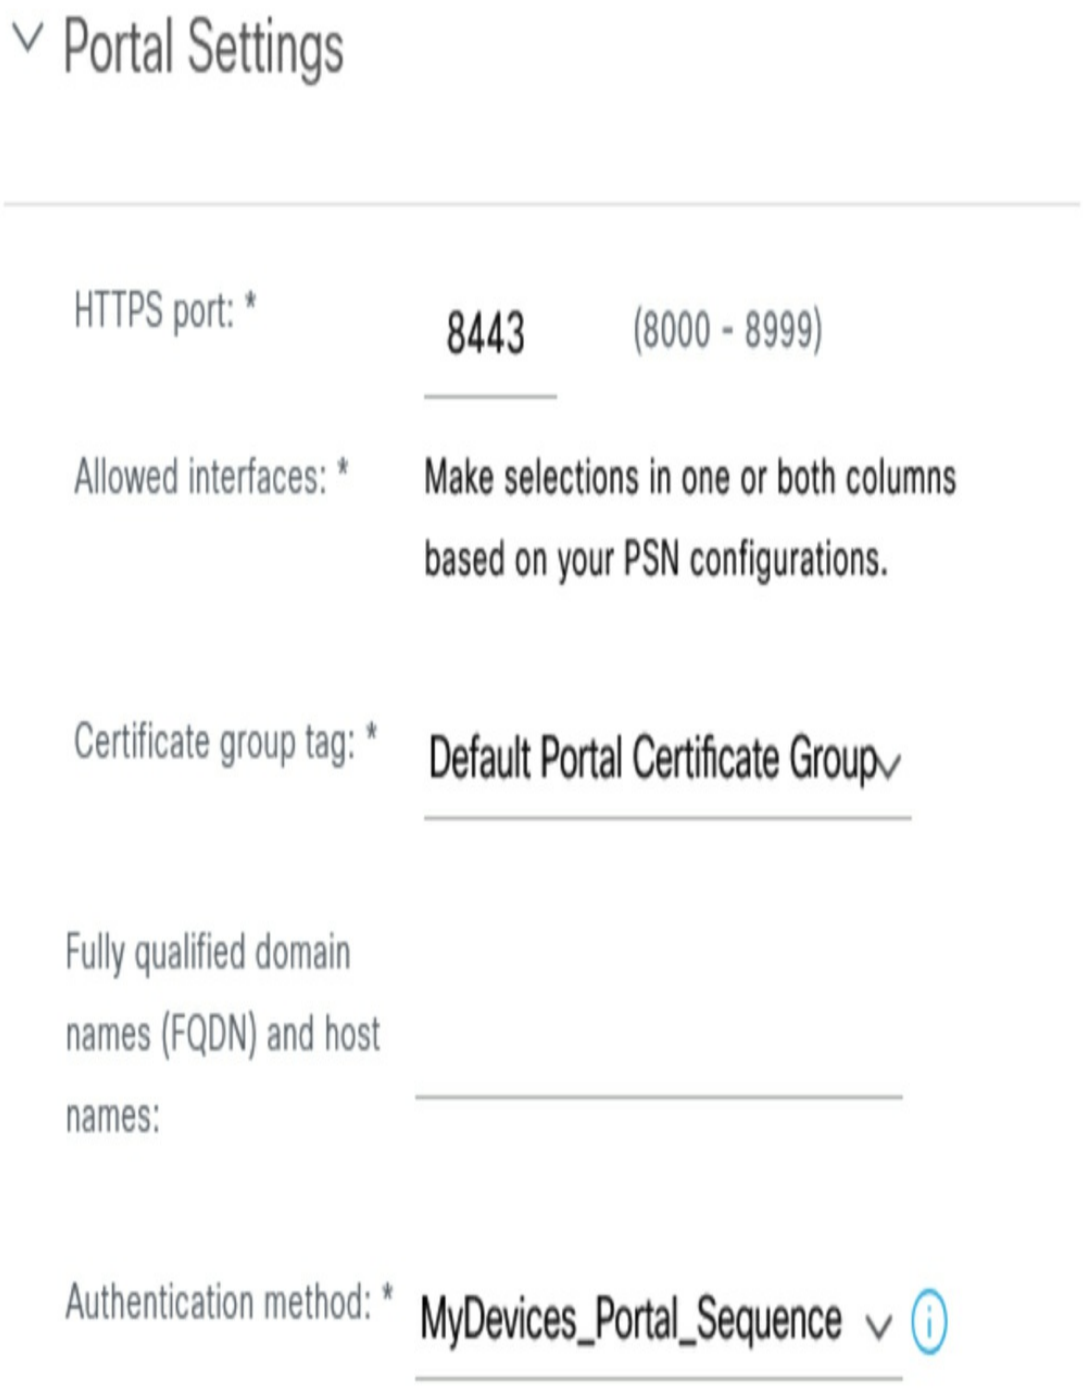


Figure 9-32 *My Devices Portal Settings*

Technet24

Summary

This chapter examines the management of guest users and personal devices in modern enterprise environments through Cisco ISE. The text outlines the essential processes for providing secure, temporary network access to visitors, auditors, and contractors by leveraging the full guest access lifecycle — account creation, credential delivery, authentication, expiration, and audit logging. Key guest access flows are covered, including sponsored access, self-registration, employee guest access, and hotspot methods, each with a breakdown of technical steps, notification options, and the role of attribute stitching for granular and compliant authorization.

The chapter explores core building blocks of guest services in ISE, such as sponsor users, sponsor portals, guest types, and guest portals. The content addresses guest notification setup, authentication source selection, and the integration of SAML identity providers for flexible access management. Secure guest access is demonstrated using various web redirection approaches, including Cisco CWA and DNS sinkholing, with guidance on best practices for HTTPS handling and authorization policy design.

Onboarding, provisioning, and management of personal devices are detailed using ISE's workflows. Distinctions between single SSID and dual SSID BYOD flows are explained, as well as the multi-phase onboarding process — from initial authentication and client provisioning to native supplicant setup and certificate-based network access. The discussion includes client provisioning policies, certificate management strategies, and the use of device self-service portals for secure and efficient personal device integration.

Throughout the chapter, step-by-step guidance is provided for configuring policy sets, authorization profiles, and client provisioning

resources to enforce consistent security across wired, wireless, and VPN connections. Applying these practices helps organizations deliver secure, scalable, and user-friendly guest and BYOD experiences.

Technet24

Chapter 10. Implementing Endpoint Posture Services

Introduction

Endpoint posture assessment is one of the critical elements of the entire Zero Trust Network Access (ZTNA) approach. Every endpoint must be verified before network access is granted to ensure its security posture aligns with corporate requirements and security policy.

Endpoint posture verification may include checks for the presence of anti-malware software, whether the anti-malware software is up to date, whether the endpoint has the latest patches installed, and similar requirements. The specific requirements are determined by the details of the organization's security policies.

ISE posture services are key components of network access control. From the Cisco Identity Services Engine perspective, posture services can be delivered in two different approaches:

- **ISE posture:** This approach means the ISE Policy Service Node (PSN) acts as a posture runtime and communicates with the endpoint posture agent installed on the device to verify compliance.
- **ISE MDM services:** With this approach, ISE does not perform

verifications on the endpoint directly. Instead, it queries an external system — the mobile device management (MDM) or unified endpoint management (UEM) server. The endpoint compliance verification is performed by the MDM/UEM, and ISE simply receives the result of such verification, which can be used to assign the appropriate level of access to the endpoint.

Table 10-1 provides a comparison between ISE posture and MDM services.

Table 10-1 *ISE Posture and MDM Services Comparison*

Aspect	ISE Posture	ISE MDM Services
Compliance Management	ISE	UEM/MDM Vendor
ISE License	Premier	Premier
Agent	Cisco Secure Client	MDM Agent/Profiles
Remediation	Yes	No
Other Features	Agent, Agentless, Stealth, Temporal, Posture Lease, PRA and Grace Period	Visibility and Profiling
Policies	Granular	Granular
Network Access	Real-Time Evaluation	MDM Status
OS Coverage	Windows, macOS, and Linux	Windows, macOS, Linux, iOS, Android, etc.
Operations	Deploying and Maintaining Compliance Checks, Agents, OS Updates	Deploying Certificates and Maintenance of Cert Attributes
Deploying Apps	Yes (Remediation)	Yes

Technet24

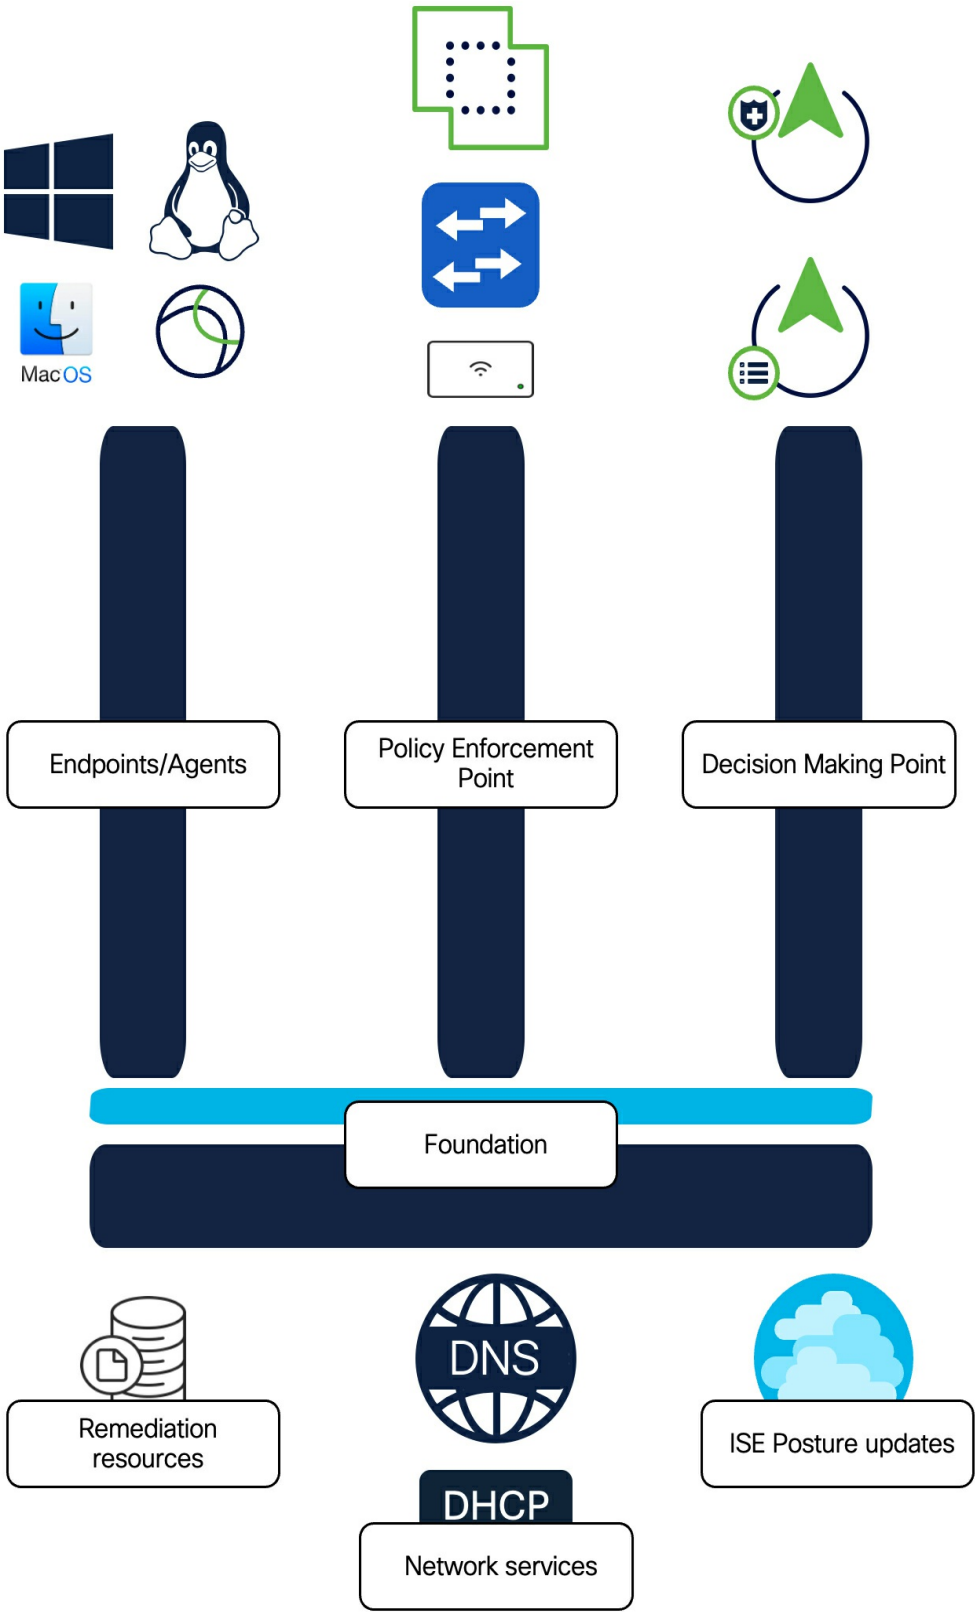
Note

The exact support for different features on the MDM server side depends on the MDM server vendor.

ISE Posture Service Pillars

The ISE posture service is considered one of the most complex use cases to implement, not only because of configuration complexity — although some complexity does exist — but due to the pervasive nature of posture implementation. When you deploy posture, it is no longer sufficient to focus solely on general endpoint settings such as certificates or 802.1X supplicant configuration. You need to understand the entire endpoint security ecosystem, including which security software is used on your endpoints, how this software is updated, and what constitutes a "compliant" endpoint state.

The best way to understand how to properly deploy and configure the ISE posture service is to start with the broader perspective and then narrow the focus to specific components and features. The entire ISE posture service can be represented by three main pillars and the foundational services that support them. Figure 10-1 demonstrates the ISE posture service pillars and the foundation services.


MacOS
Endpoints/Agents
Policy Enforcement Point
Decision Making Point
Foundation
Remediation resources
DNS
DHCP
Network services
ISE Posture updates

Technet24

Figure 10-1 *ISE Posture Services Pillars*

Now let's discuss the individual pillars and their roles in the ISE posture service:

- **Endpoint and Agents:** This pillar includes all the endpoints supported for the ISE posture service, such as Windows, macOS, and Linux, as well as all available posture agent types like the Secure Client ISE posture module and temporal agent. This category also covers third-party security applications and services that ISE posture can validate and control.
- **Policy Enforcement Points:** For the ISE posture service, Policy Enforcement Points represent all types of network access devices (NADs) supported in the posture flow. Later in the chapter, we will discuss all available variations of ISE posture flows and their requirements for NADs.
- **Decision-Making Points:** This pillar includes the various nodes in the ISE deployment. The ISE Administration node is used to perform posture configuration, which is propagated to all other nodes in the deployment. The Administration node is also responsible for acquiring all updates required for the posture service and distributing those updates across the deployment. The ISE PSNs act as the posture runtime elements where all endpoint verifications are performed and authorization decisions are made.
- **The Foundation:** All three pillars require additional components for the posture service to function. This includes update and remediation resources used by endpoints during the process; network services such as network transport, DNS, and DHCP to ensure successful communication among system elements; and ISE posture updates, which are stored in the cloud and regularly provided by Cisco.

Posture Flow: Key Milestones

Now that we've identified the key components, it is important to focus on the elements of the ISE posture flow. The posture flow generally describes the steps required for the endpoint posture assessment process to be successfully completed. Although the steps may vary depending on the flow type, understanding each one clearly simplifies the comprehension of individual flows later in the chapter. Figure 10-2 demonstrates the maximum set of steps within the ISE posture flow.

Technet24



Optional: Agent pre-install



Initial Authentication/Authorization



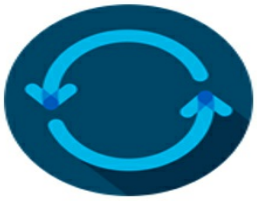
Client Provisioning



Posture assessment process



Optional: Remediation process



Change of Authorization (CoA)



Final Authentication/Authorization

Figure 10-2 *ISE Posture Flow Steps*

The key milestones of the ISE posture flow are as follows:

- **Agent pre-installation (Optional):** While the agent can be provisioned directly from ISE as part of initial authentication and redirection, it is also possible to have the agent pre-installed. The agent that supports pre-installation is Cisco Secure Client with the ISE posture module. Administrators can install Secure Client manually or use software management tools such as Microsoft SCCM. The agent can be pre-installed with all required profiles. Pre-installation is the preferred approach when users do not have administrative rights on their computers, making software installation directly from ISE challenging. Secure Client pre-installation is also commonly used for the redirection-less posture flow.
- **Initial authentication/authorization:** At this step, limited network access is provided. Depending on the posture flow type, access may be restricted by web redirection for redirect-based flows or by other methods such as VLAN or ACL for redirection-less flows. When redirection is used, users may be redirected to either the client provisioning portal or the guest portal for authentication (when posture is combined with the guest flow) and client provisioning.
- **Client provisioning:** This step can be executed in two ways. The first involves end-user participation, which occurs when the endpoint does not have a posture agent installed or a temporary agent needs to be used. In this scenario, the user is redirected to the portal where the client provisioning policy is identified using session details such as user group membership and connection type. The user is then guided through the process of obtaining and running the posture agent. The second approach happens

Technet24

when the agent is already present on the endpoint. In this case, the agent contacts the ISE PSN and receives all necessary updates without end-user involvement.

- **Posture assessment process:** At this stage, the agent contacts an ISE PSN and requests the list of posture policies and conditions that require verification. On the PSN side, the posture policy lookup is performed. Session information, such as user group membership and connection type, is combined with details from the posture request submitted by the agent, including endpoint OS and agent type, to locate the appropriate posture policy or policies. Multiple policies can be selected if matches exist. If no match is found, the endpoint receives a notification that its status is "compliant," meaning posture operates in a fail-open mode and network access is granted as if the endpoint is in a compliant state. When posture policies are identified, the agent receives the list of requirements and performs their verification locally. The agent uses compliance modules to identify third-party security tools, programs, and services, and to remediate them if necessary. After local assessment, the report is returned to ISE, where the final compliance decision is made based on policy settings, such as whether the requirement is mandatory, optional, or set to audit mode.
- **Remediation process (Optional):** If an endpoint does not meet mandatory requirements, the ISE system typically provides remediation instructions. For example, when anti-malware software is outdated, automatic remediation can trigger a software update. Before the remediation period ends, the agent will revalidate the status of the endpoint to confirm that all mandatory requirements are fulfilled. If the device remains noncompliant, the session state on ISE changes to noncompliant, and a Change of Authorization will be issued. As an administrator,

you may configure a dedicated authorization policy in ISE that enforces restricted access for noncompliant endpoints.

- **Change of Authorization (CoA):** When the posture status of a session changes, PSN sends a CoA request to the NAD to reauthenticate the connection so ISE can apply the new authorization policy. For wired and wireless access, the CoA triggers reauthentication, resulting in a subsequent authentication attempt. In remote access VPN scenarios, reauthentication typically disconnects the end user. To prevent disruption in VPN environments, the CoA Push is used. With CoA Push, the next authentication is not required because the necessary authorization attributes are included in the CoA itself.
- **Final authentication/authorization:** At this stage, the ISE PSN grants the access level intended for a compliant endpoint.

Note

While posture assessment is technically possible for both corporate users and guests (both approaches will be covered later in the chapter), posture for guest users is not generally recommended except for contractor use cases.

Endpoints and Agents

ISE posture services support a variety of posture agents, each with unique characteristics to address a broad spectrum of customer requirements. The types of ISE posture agents are as follows:

- **Secure Client ISE posture module:** This agent is often referred to as the "persistent" posture agent. The ISE posture module is a recommended method for delivering posture assessment because it offers the broadest coverage of conditions and supports both

Technet24

manual and automatic remediation. From an administrator's perspective, this approach delivers the most flexible way to control endpoints.

Cisco Secure Client ISE posture module is available for Windows, macOS, and Linux.

- **Secure Client ISE posture module Stealth:** This configuration allows you to run Cisco Secure Client with the ISE posture module in the background as a service without a user interface. In this mode, the ISE posture module can deliver notifications to the end user and provide a limited set of remediation capabilities.

  Cisco Secure Client ISE posture module Stealth is available for Windows and macOS.

- **ISE posture module Temporal Agent:** The temporal agent is delivered to the endpoint from the guest or client provisioning portal during each authentication. This agent operates as a portable application, which means installation is not necessary. After downloading, the user simply runs the app. Once the posture status is determined, the application can no longer be used. The temporal agent offers a limited set of conditions for verification and a restricted range of remediation actions.

  ISE posture module Temporal Agent is available for Windows and macOS.

- **ISE Agentless posture:** With this approach, ISE automatically pushes a lightweight agent to the endpoint without any interaction from the user. The posture status is evaluated, and then the software is removed from the device. This method provides limited support for posture conditions and remediation actions.

  ISE Agentless posture is available for Windows and macOS.

Table 10-2 presents a high-level comparison of agents from the perspective of deployment complexity and visibility to end users.

Table 10-2 *High-Level Comparison of Agents*

Feature	Persistent Agent	Stealth Agent	Temporal Agent	Agentless Posture
Install required	Once	Once	No, Run every session	No
End-user interaction	Minimal, after first install	None	Each session	None
Deployment effort	Highest upfront (install/profile), lowest ongoing	Highest upfront (install/profile), lowest ongoing	Lowest	Low, but NAD accounting + admin-credential access required

A comparison of functionality available with different agent types for various endpoint operating systems is provided in Table 10-3.

Table 10-3 *Comparison of Agent Functionality*

Technet24

Capability	Cisco Secure Client (SC)			SC Stealth		Temporal Agent		Agentless Posture	
	Windows	macOS	Linux	Windows	macOS	Windows	macOS	Windows	macOS
Anti-Malware Checks	+	+	+	+	+	+	+	+	+
Firewall Installation Checks	+	+	—	+	+	+	+	+	+
Application Inventory	+	+	—	+	+	+	+	+	+
Hardware Inventory	+	+	—	+	+	+	+	+	+
Process Checks	+	+	+	+	+	+	+	+	+
Dictionary Conditions	+	+	+	+	+	+	+	+	+
Application Checks	+	+	—	+	+	+	+	+	+
File Checks	+	+	Limited	+	+	+	+	+	+
Service Checks	+	+	—	+	+	+	+	+	+
Disk Encryption	+	+	—	+	+	+	+	+	+
Patch Management	+	+	Limited	+	+	+	+	+	+
Registry Checks	+	N/A	N/A	+	N/A	+	N/A	Limited	N/A
USB Checks	+	—	—	+	—	+	—	+	—
WSUS Remediation (legacy)	+	N/A	N/A	+	N/A	—	N/A	—	N/A
Remediation	Auto / Manual	Partial	Partial	Partial Auto	Partial	Text	Text	—	—
Reassessment	+	+	+	+	+	—	—	—	—

The following examples provide some reasoning behind selecting a specific agent type:

- **Secure Client ISE posture module:** You can select this option when the endpoint is corporate-managed or requires ongoing and continuous compliance monitoring. The reasoning behind this choice: It is the only agent type that supports the full condition set, including periodic and continuous reassessment. This module serves as the default for anything you need monitored, not just checked once.
- **Secure Client ISE posture module Stealth:** Choose this agent when posture assessment must run as a service with minimal or no interaction with the end user. The reasoning behind this selection: It supports ongoing and continuous compliance monitoring.
- **ISE posture moduleTemporal Agent:** Use this type when posture is required for guest, contractor, or 802.1X portal flows where you do not want to install anything but still need a legitimate posture check. The common reasoning behind this decision: It allows you to run an assessment without leaving a footprint and trades reassessment capability for zero installation.
- **ISE Agentless posture:** This agent is the best choice for managed endpoints you can access via SSH or WinRM with administrative credentials, where installation and user interaction are not possible. The reasoning behind this selection: It enables you to deliver point-in-time compliance audits.
- **ISE posture compliance module:** Every agent type uses an ISE posture compliance module. The compliance module is a library that determines whether the endpoint complies with defined posture conditions. For Secure Client implementations, both regular and stealth, you, as an administrator, must specify which

Technet24

version of the compliance module to use in the Secure Client configuration. For temporal agent and agentless posture, the compliance module is built into the package itself.

The compliance module exists in two main branches: 3.x and 4.x. The 4.x branch is more modern and should be used in all new implementations. You specify compliance module branch versions in the posture policy. Agents can only retrieve posture policies for the same branch. For example, if an agent's compliance module version is 4.100, only policies with the 4.x branch can be assigned by ISE during the posture assessment process.

ISE Posture Flows

Understanding the different ISE posture flows enables you to manage ISE posture services more effectively during planning, implementation, ongoing operation, and troubleshooting. You can deliver the ISE posture service using three primary flows. The following provides a high-level overview of these flows:

- **Redirect-based posture:** This approach is recommended for all Cisco network access devices because it ensures that users or posture agents can reach the correct Policy Service Node during client provisioning and posture assessment stages. This is achieved with Central Web Authentication mechanisms available in Cisco NADs.
- **Redirection-less posture:** This flow should be used only with non-Cisco network access devices that do not support Central Web Authentication. While the posture process can operate without redirection, the lack of this feature increases the chances of failure when agents are unable to locate the correct Policy Service Node or determine the proper posture status for the

session.

- **Agentless posture:** This flow does not depend on redirection. Instead of the traditional method, where the agent locates the PSN, the agentless posture approach requires the PSN to locate the endpoint and push the agentless posture module to determine the actual posture status.

In all three flows, the essential requirement for NADs is to support Change of Authorization. Similar to the guest flow, there is an initial authentication with limited access, which then transitions to the required access level after posture assessment and reauthentication triggered by CoA.

Redirect-Based Posture

The redirect-based posture flow includes two variations. The first occurs when the endpoint does not have a posture agent pre-installed or requires the use of the temporal agent. The second variation applies when the agent is already installed and no user interaction with the client provisioning portal is necessary.

Both variations begin identically; however, the sequence of steps diverges based on whether a pre-installed agent exists. If the agent is not present, the flow may split again depending on the client provisioning policy selection, since the steps for the temporal agent differ from those needed to install Cisco Secure Client with the ISE posture module.

The following description explains all possible variants of the redirect-based flow. Figure 10-3 shows the beginning of the redirect posture flow.

Technet24

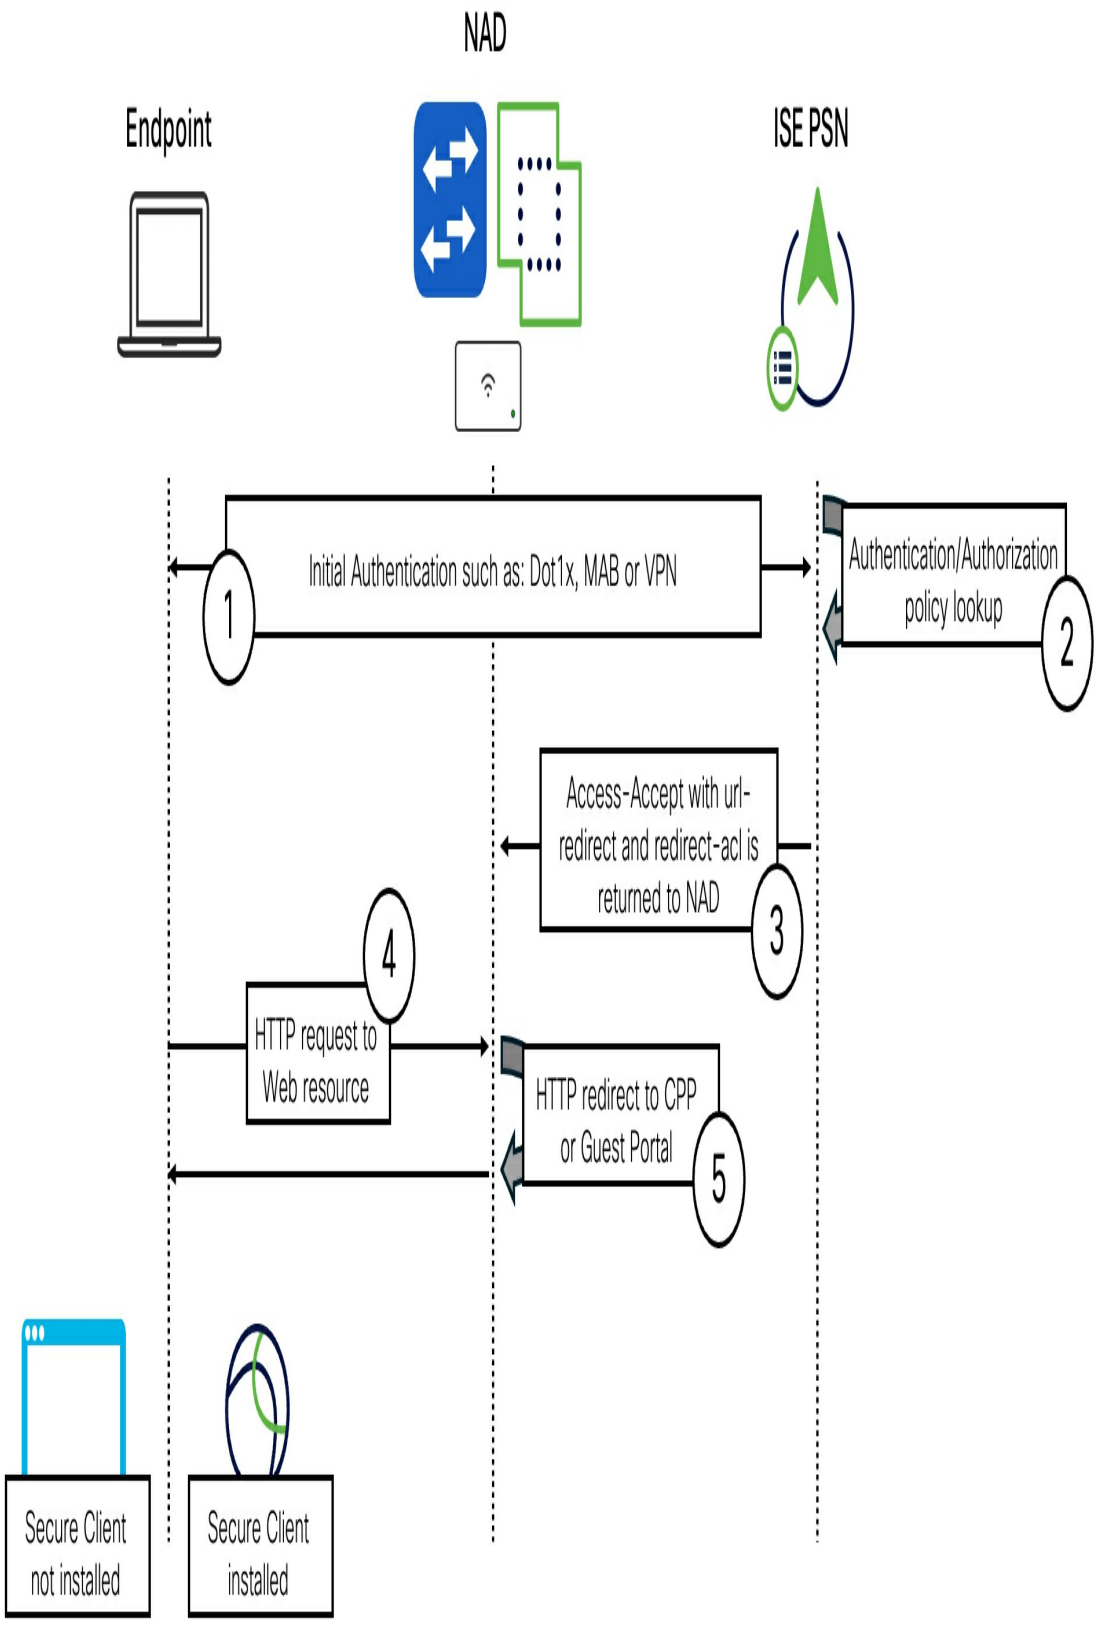

NAD
Endpoint
ISE PSN
Initial Authentication such as: Dot1x, MAB or VPN
1
Authentication/Authorization policy lookup
2
Access-Accept with url-redirect and redirect-acl is returned to NAD
3
4
HTTP request to Web resource
HTTP redirect to CPP or Guest Portal
5
Secure Client not installed
Secure Client installed

Figure 10-3 *Initial Authentication and Redirection in the Redirect Posture Flow*

Step 1. Initial Dot1x, MAB, or VPN authentication.

Step 2. Authentication and authorization policy selection. When you deploy posture as part of the guest flow, the policy set should include an authentication policy that allows authorization to continue when the user is not found. For the guest posture use case, the initial authorization policy typically specifies the SSID name as a condition along with a posture session status check. For initial authentication in all use cases and flows, it is recommended to use the condition from the session dictionary — PostureStatus not equal to compliant, as shown in Figure 10-4.

Technet24



Figure 10-4 *Example of Initial Authorization Policy*

Step 3. As part of the Access-Accept response, redirection attributes are returned to the network access device. These

attributes include the name of the redirect ACL, which must be configured on the NAD, and the redirect URL. For the guest posture flow, you should redirect the user to the guest portal. In other cases, redirection occurs to the client provisioning portal. Figure 10-5 shows a comparison of the redirection settings in the authorization profiles.

Technet24

Initial Authentication authorization profile with Client Provisioning Posture redirect



Initial Authentication authorization profile with Guest portal redirect



Figure 10-5 *Redirection Options for Initial Authentication*

Step 4. HTTP activity from the user or from the browser, such as captive portal detection probes, reaches the network access device where redirection has been applied as part of the previous step.

Step 5. The NAD redirects the connection to the URL obtained from ISE in the Access-Accept response.

The subsequent steps depend on whether the Cisco Secure Client ISE posture module is present. First, let's review what occurs when the module is not installed.

Figure 10-6 shows the sequence of steps performed when the ISE posture module is not installed. This includes user redirection to the portal and client provisioning steps. This series concludes with a split in the flow, depending on the outcome of the client provisioning policy, because the subsequent steps differ for the temporal agent and Secure Client installation.

Technet24

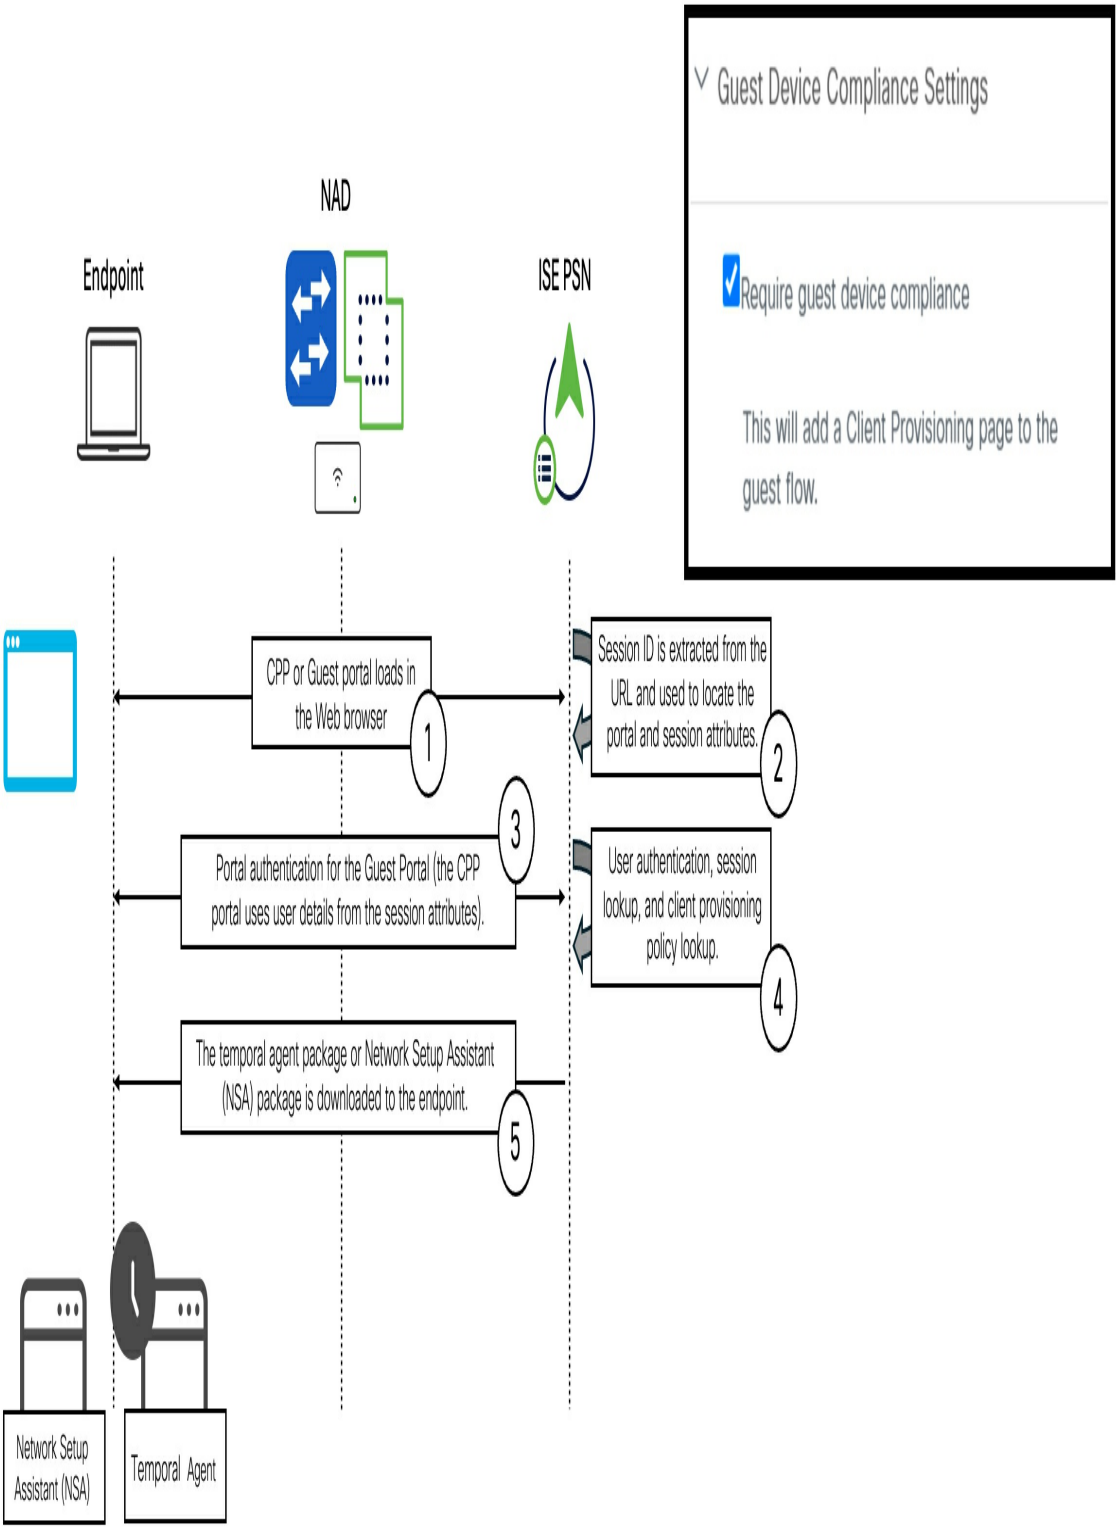

Guest Device Compliance Settings
Require guest device compliance
This will add a Client Provisioning page to the guest flow.
NAD
Endpoint
ISE PSN
CPP or Guest portal loads in the Web browser
1
Session ID is extracted from the URL and used to locate the portal and session attributes.
2
3
Portal authentication for the Guest Portal (the CPP portal uses user details from the session attributes).
User authentication, session lookup, and client provisioning policy lookup.
4
The temporal agent package or Network Setup Assistant (NSA) package is downloaded to the endpoint.
5
Network Setup Assistant (NSA)
Temporal Agent

Figure 10-6 *Redirect-Based Posture Portal Interaction and Client Provisioning*

Step 1. The client provisioning portal or guest portal loads in the user's browser.

Step 2. On the ISE side, the URL that users open as a result of the redirect is used to locate session details such as user group membership, connection type, SSID name, and other relevant information. ISE performs this lookup using the audit session ID included in the redirect URL. In parallel with session lookup, ISE also uses the user-agent provided by the web browser to identify the endpoint's operating system type. You will use this information during the client provisioning policy lookup.

Step 3. If posture is performed as part of the guest flow, guest authentication takes place on the portal for the user. Authentication details are added as session attributes, which you can also use for client provisioning policy selection. For the client provisioning portal, authentication does not occur, so you typically use this portal in scenarios where the user is authenticated during the initial authentication step, such as Dot1x or VPN.

Note

In the guest flow, for client provisioning policy selection and subsequent posture agent download, you need to enable the **Require Guest Device Compliance** checkbox in the **Guest Device Compliance Settings** of the relevant guest portal, as shown in Figure 10-6.

Step 4. Client provisioning policy lookup occurs on the ISE side to

Technet24

determine which agent should be used for posture assessment. ISE performs client provisioning policy selection in a top-down manner until the first match is found. If no match is found, an error will be displayed to the user. Figure 10-7 shows an example of client provisioning policies for guest and Dot1x use cases, where Secure Client is used for Dot1x and the temporal agent is applied for guest users. You can find client provisioning policies by navigating to **Work Centers > Posture > Client Provisioning**.

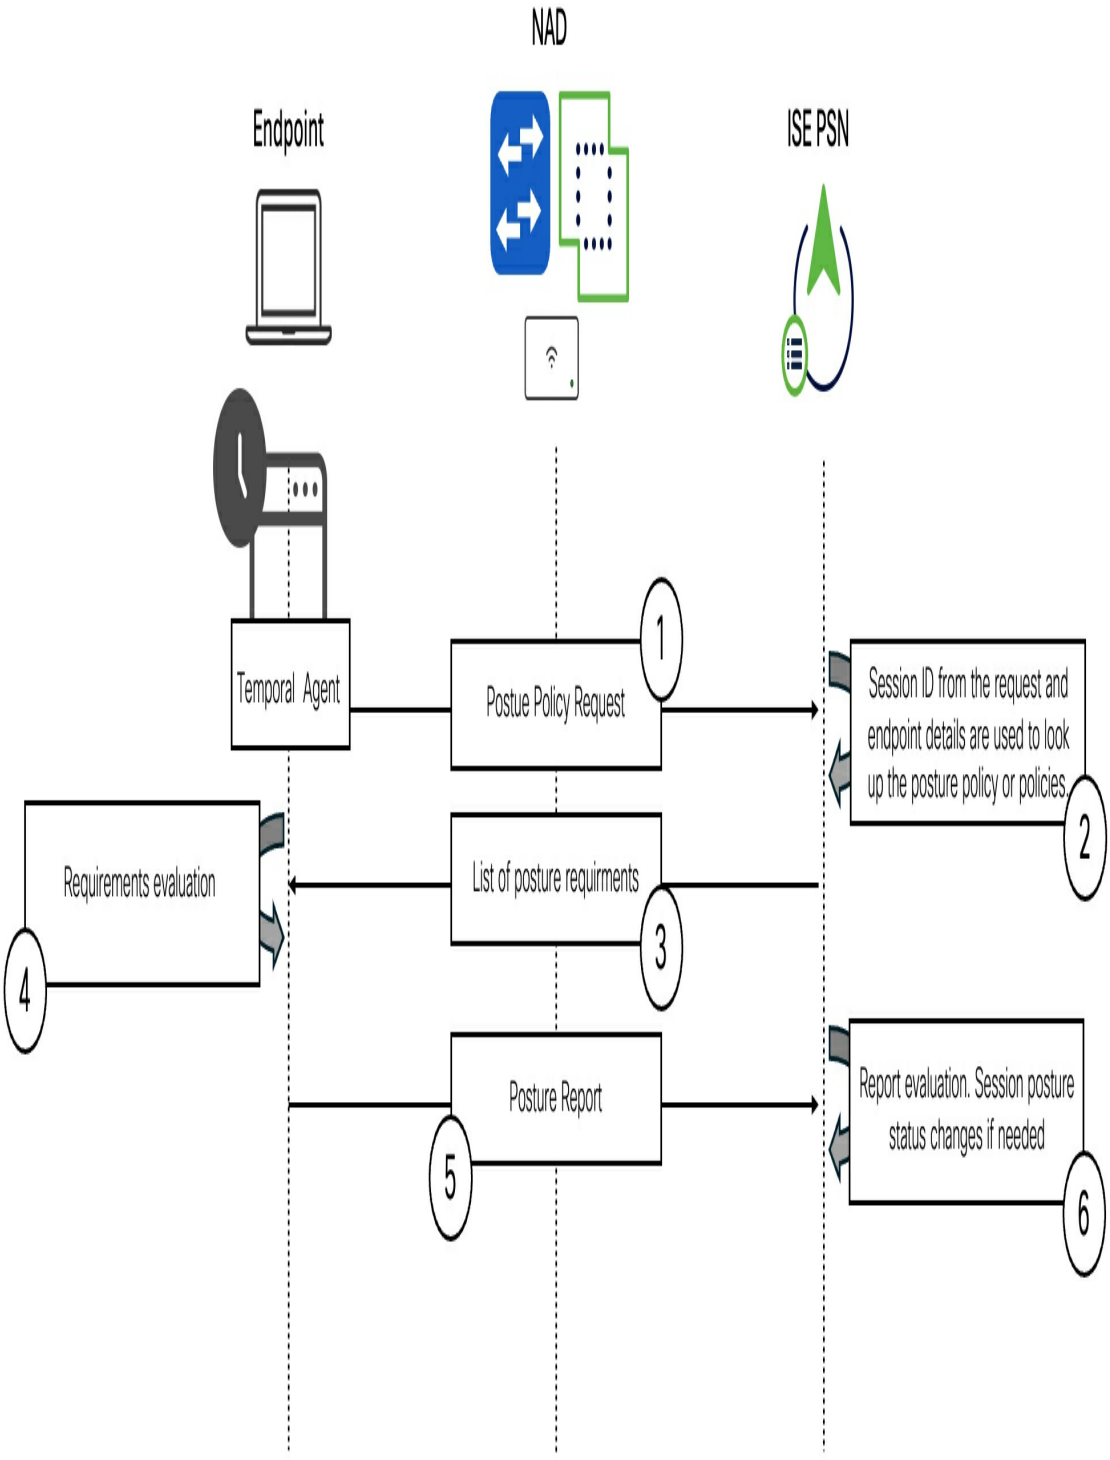

Rule Name
Identity Groups
Operating Systems
Other Conditions
Results
Win-SC-dot1x
If Any
and Windows All
and ISELAB:ExternalGroups EQUALS iselab.net/Users/Domain Users AND Radius:Called-Station-ID CONTAINS MARS-Demo-HQ
then SC-5-1-17-3394
Policy for users authenticated over Dot1x.
AD user group membership and SSID name are used as conditions
Result is Secure Client config
Win-Temporal-Agent
If GuestType_Daily (default)
and Windows All
and Radius:Called-Station-ID CONTAINS MARS-Demo-Guest
then CiscoTemporalAgentWindows 5.1.10.41
Policy for Guest users
Guest users' identity group and SSID name are used as conditions
Result is Temporal Agent

Technet24

Figure 10-7 *Client Provisioning Policies Example*

Step 5. Depending on the client provisioning policy selection, either the temporal agent or Network Setup Assistant (NSA) is downloaded by the user. At this point, the posture flow splits again.

Figure 10-8 demonstrates the variation of the redirect-based flow with the temporal agent.


NAD
Endpoint
ISE PSN
Temporal Agent
Postue Policy Request
1
Session ID from the request and endpoint details are used to look up the posture policy or policies.
2
Requirements evaluation
List of posture requirments
3
4
Posture Report
5
Report evaluation. Session posture status changes if needed
6

Technet24

Figure 10-8 *The Redirect-Based Flow with Temporal Agent*

Step 1. The agent submits a request to obtain the list of posture policies and requirements for evaluation. This request includes audit session ID, operating system details, and the agent type/version. The temporal agent does not perform PSN discovery to send the request, because the necessary details are supplied with the agent file itself.

Step 2. On the PSN side, details from the request are used to perform posture policy selection. By default, ISE provides multiple preconfigured posture policies, which you can use as is or modify for your specific use case. You can find posture policies at **Work Centers > Posture > Posture Policy**. Figure 10-9 demonstrates preconfigured policies for the temporal agent.

Posture Policy

Guide Me

Define the Posture Policy by configuring rules based on operating system and/or other conditions.

Temporal

Status	Policy Options	Rule Name	Identity Groups	Operating Systems	Compliance Module	Posture Type	
	Policy Options	Default_AntiMalware_Policy_Mac_temporal	If Any	and Mac OSX	and 4.x or later	and Temporal Agent	and
	Policy Options	Default_AntiMalware_Policy_Win_temporal	If Any	and Windows All	and 4.x or later	and Temporal Agent	and
	Policy Options	Default_AppVis_Policy_Mac_temporal	If Any	and Mac OSX	and 4.x or later	and Temporal Agent	and
	Policy Options	Default_AppVis_Policy_Win_temporal	If Any	and Windows All	and 4.x or later	and Temporal Agent	and
	Policy Options	Default_Firewall_Policy_Mac_temporal	If Any	and Mac OSX	and 4.x or later	and Temporal Agent	and
	Policy Options	Default_Firewall_Policy_Win_temporal	If Any	and Windows All	and 4.x or later	and Temporal Agent	and

Technet24

Figure 10-9 *Default Posture Policies for Temporal Agent*

Note

You can use the **Guide Me** button to follow structured instructions on how to configure posture services in ISE. The **Guide Me** button is provided by the WalkMe integration.

Step 3. The list of posture requirements and their corresponding remediation actions is returned to the agent.

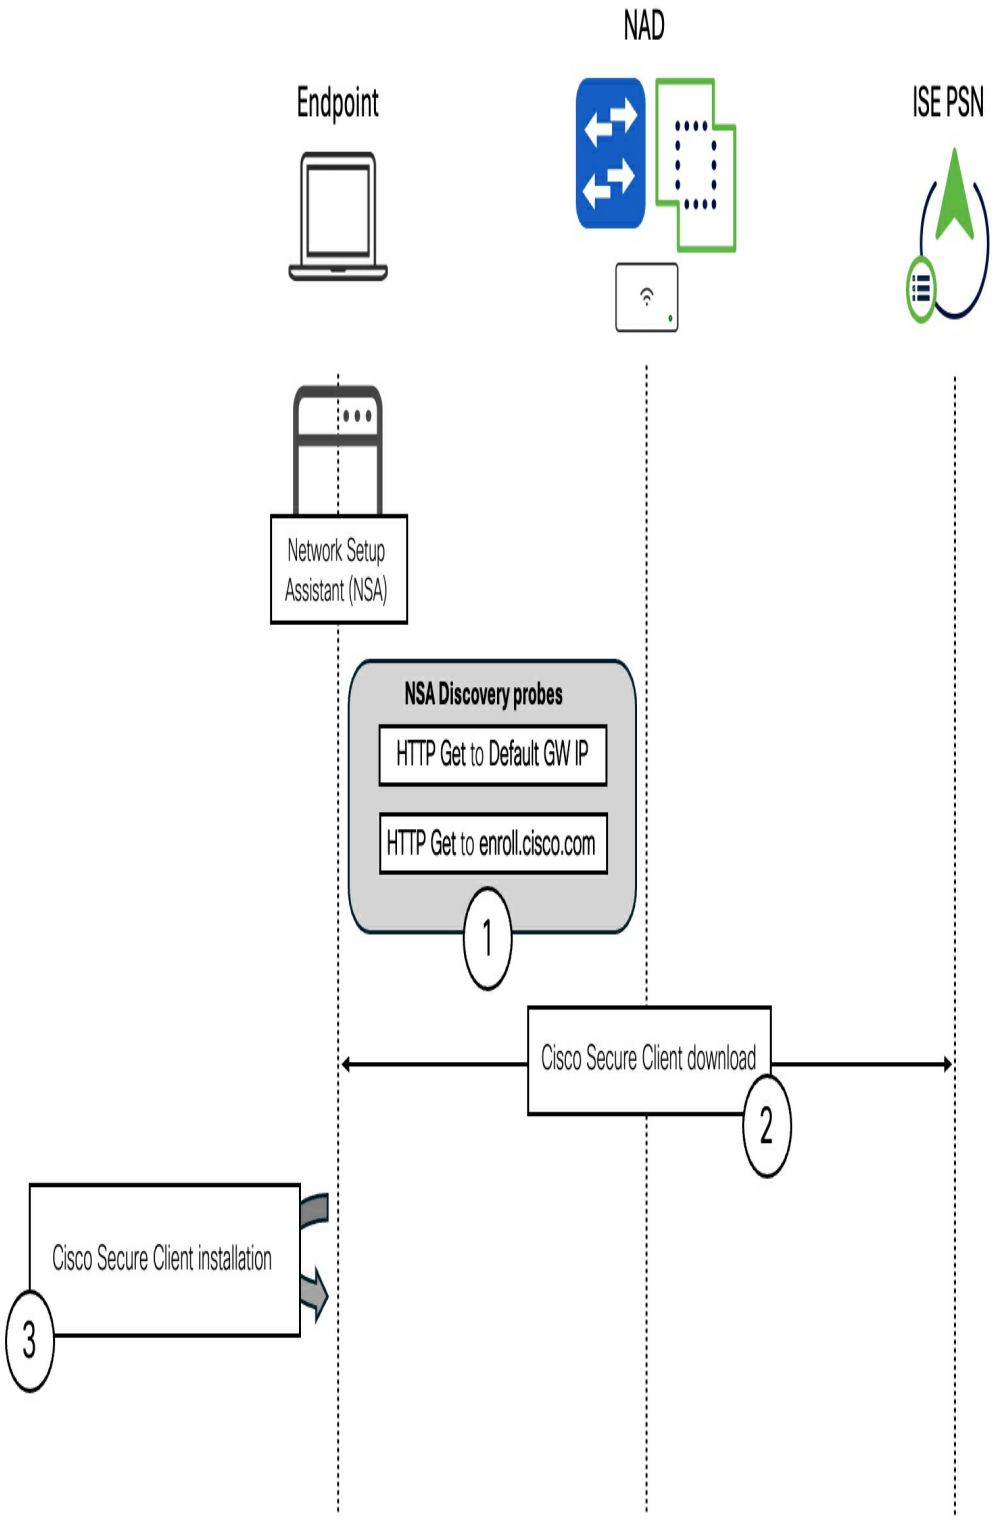
Step 4. The agent verifies the requirements locally using the compliance module and prepares the posture assessment report.

Step 5. The posture report is submitted to ISE.

Step 6. ISE evaluates the report according to the posture policy settings. If any mandatory conditions defined in the posture policy fail, the agent is notified about noncompliance, and the remediation process starts on the agent side. If all mandatory conditions are met, the session status changes to compliant.

After Step 6, both subtypes of the redirect-based flow — with Secure Client and with the temporal agent — continue with Change of Authorization and final authentication and authorization, which will be demonstrated later in the chapter.

Now, return to the variation of the flow where Secure Client installation is required. Figure 10-10 shows the process of installing Secure Client with the help of Network Setup Assistant. The NSA is a lightweight Cisco-provided software responsible for obtaining Secure Client modules from Cisco ISE. The NSA is available for both Windows and macOS.


NAD
Endpoint
ISE PSN
Network Setup Assistant (NSA)
NSA Discovery probes
HTTP Get to Default GW IP
HTTP Get to enroll.cisco.com
1
Cisco Secure Client download
2
Cisco Secure Client installation
3

Technet24

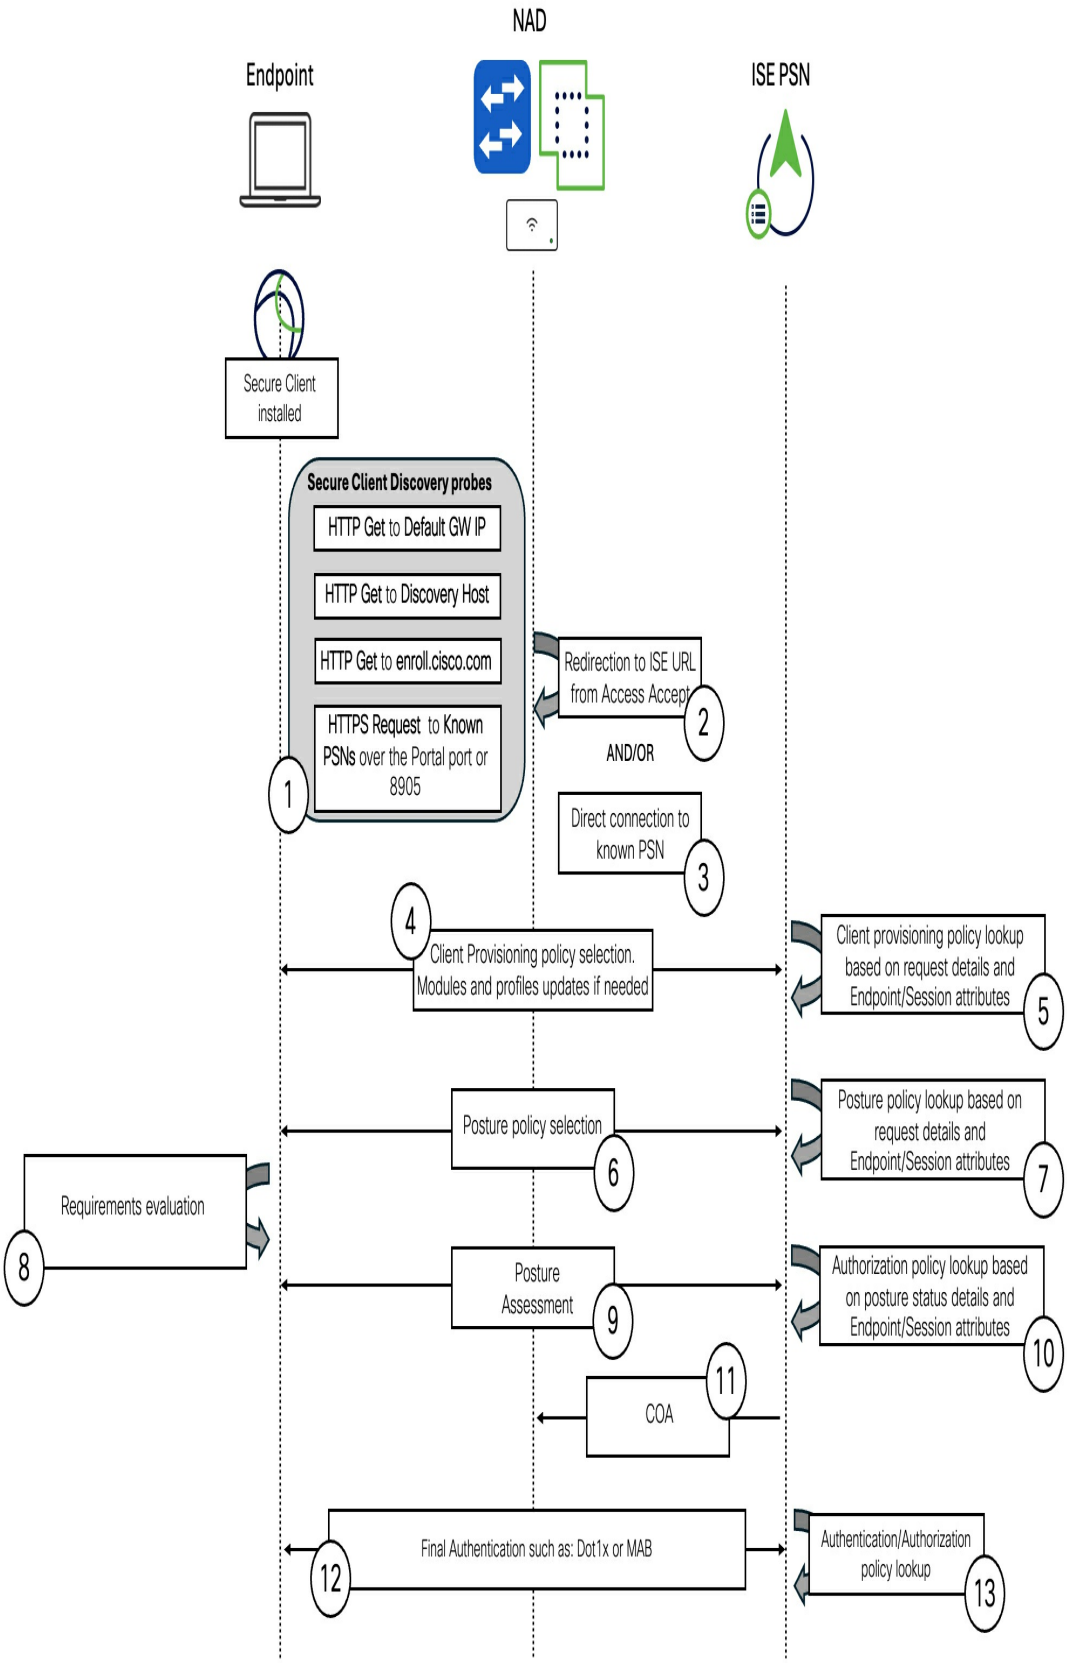
Figure 10-10 *Secure Client Installation with NSA*

Step 1. Network Setup Assistant uses HTTP probes to the IP address of the default gateway and the FQDN enroll.cisco.com to obtain the redirect URL from the network access device.

Step 2. Network Setup Assistant sends a request to the Policy Service Node and downloads Secure Client modules and profiles.

Step 3. Secure Client installs on the endpoint.

Once Secure Client is installed, the process continues with the version of the flow that uses the pre-installed posture agent. Figure 10-11 shows the redirect-based posture flow with Secure Client installed.


NAD
Endpoint
ISE PSN
Secure Client installed
Secure Client Discovery probes
HTTP Get to Default GW IP
HTTP Get to Discovery Host
HTTP Get to enroll.cisco.com
HTTPS Request to Known PSNs over the Portal port or 8905
1
Redirection to ISE URL from Access Accept
2
AND/OR
Direct connection to known PSN
3
4
Client Provisioning policy selection. Modules and profiles updates if needed
Client provisioning policy lookup based on request details and Endpoint/Session attributes
5
Posture policy selection
6
Posture policy lookup based on request details and Endpoint/Session attributes
7
Requirements evaluation
8
Posture Assessment
9
Authorization policy lookup based on posture status details and Endpoint/Session attributes
10
11
COA
Final Authentication such as: Dot1x or MAB
12
Authentication/Authorization policy lookup
13

Technet24

Figure 10-11 *Redirect-Based Posture Flow with Secure Client Installed*

Step 1. The ISE posture module in Secure Client starts the discovery process. The triggers for this process are

- Initial login
- Layer 3 (L3) change or network interface card (NIC) change (new IP address, NIC state change).

Note

Reauthentication does not trigger the PSN discovery process. The shortcomings of this behavior and methods to address them will be explained later in the chapter.

The agent executes discovery probes simultaneously. The result from the first successful probe is used. Results from redirect-based probes take precedence over results from probes to known PSNs.

Step 2. Redirect-based probes reach the network access device, and as a result, the redirect URL is returned to the agent.

Step 3. The probe to the previously connected Policy Service Node reaches the PSN, where a local session lookup is performed using the endpoint's IP address and MAC address to determine if the node owns the session for that endpoint. When the ISE posture module is freshly installed, the agent has no knowledge of previously connected PSNs. After the first successful posture attempt, a special file called ConnectionData.xml is created on the system. This file is updated after every successful posture attempt.

Step 4. After the PSN is identified, the posture module performs

update checks. During this step, client provisioning policy selection occurs on the PSN side. If any profiles or agent components require an update, the agent receives newer versions from the PSN.

Step 5. The PSN performs a client provisioning policy lookup to determine what needs to be updated. PSNs use the audit session ID from the agent request, along with agent and endpoint details, during client provisioning policy selection.

Step 6. The agent submits a request to obtain the list of posture policies and requirements for evaluation. This request includes audit session ID, operating system details, and the agent type/version.

Step 7. On the PSN side, details from the request are used to perform posture policy selection. By default, ISE provides multiple preconfigured posture policies, which you can use as is or modify for your specific use case. Once policies are selected, the list of posture requirements and their corresponding remediation actions is returned to the agent.

Step 8. The agent verifies the requirements locally using the compliance module and prepares the posture assessment report.

Step 9. The posture report is submitted to ISE.

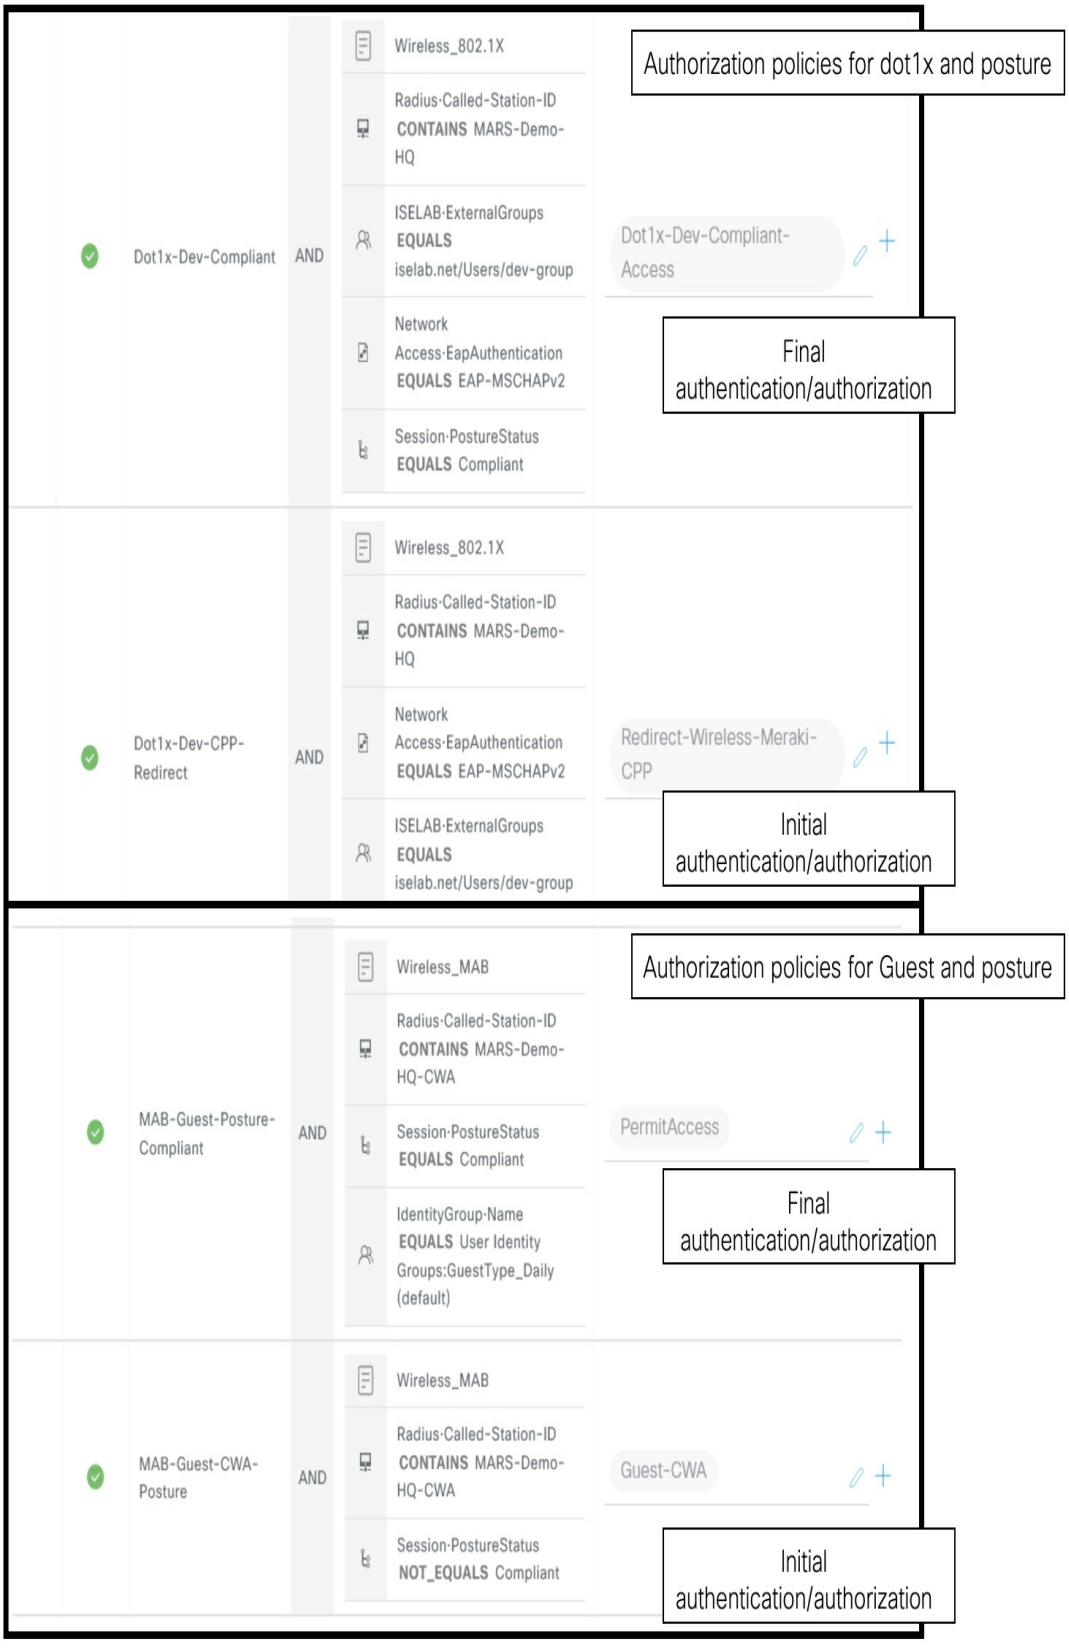
Step 10. ISE evaluates the report according to the posture policy settings. If any mandatory conditions defined in the posture policy fail, the agent is notified about noncompliance and the remediation process starts on the agent side. If all mandatory conditions are met, the session status changes to compliant.

Technet24

Step 11. ISE sends a CoA to the network access device to trigger reauthentication, resulting in a subsequent authentication attempt. For posture over VPN, CoA Push is used instead of CoA Reauth to keep the user connected.

Step 12. The final authentication happens using Dot1x or MAB.

Step 13. On the PSN side, authentication and authorization policy lookup is performed. At this stage, the authorization policy will typically have the PostureStatus from the session dictionary set to compliant, along with other attributes used in the initial authorization policy. For the guest use case, the authorization policy will usually include the guest user identity group. Figure 10-12 demonstrates examples of authorization policies for Dot1x and guest posture.


Authorization policies for dot1x and posture
Dot1x-Dev-Compliant
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
ISELAB·ExternalGroups EQUALS iselab.net/Users/dev-group
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
Session·PostureStatus EQUALS Compliant
Dot1x-Dev-Compliant-Access
Final authentication/authorization
Dot1x-Dev-CPP-Redirect
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
ISELAB·ExternalGroups EQUALS iselab.net/Users/dev-group
Redirect-Wireless-Meraki-CPP
Initial authentication/authorization
Authorization policies for Guest and posture
MAB-Guest-Posture-Compliant
AND
Wireless_MAB
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ-CWA
Session·PostureStatus EQUALS Compliant
IdentityGroup·Name EQUALS User Identity Groups:GuestType_Daily (default)
PermitAccess
Final authentication/authorization
MAB-Guest-CWA-Posture
AND
Wireless_MAB
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ-CWA
Session·PostureStatus NOT_EQUALS Compliant
Guest-CWA
Initial authentication/authorization

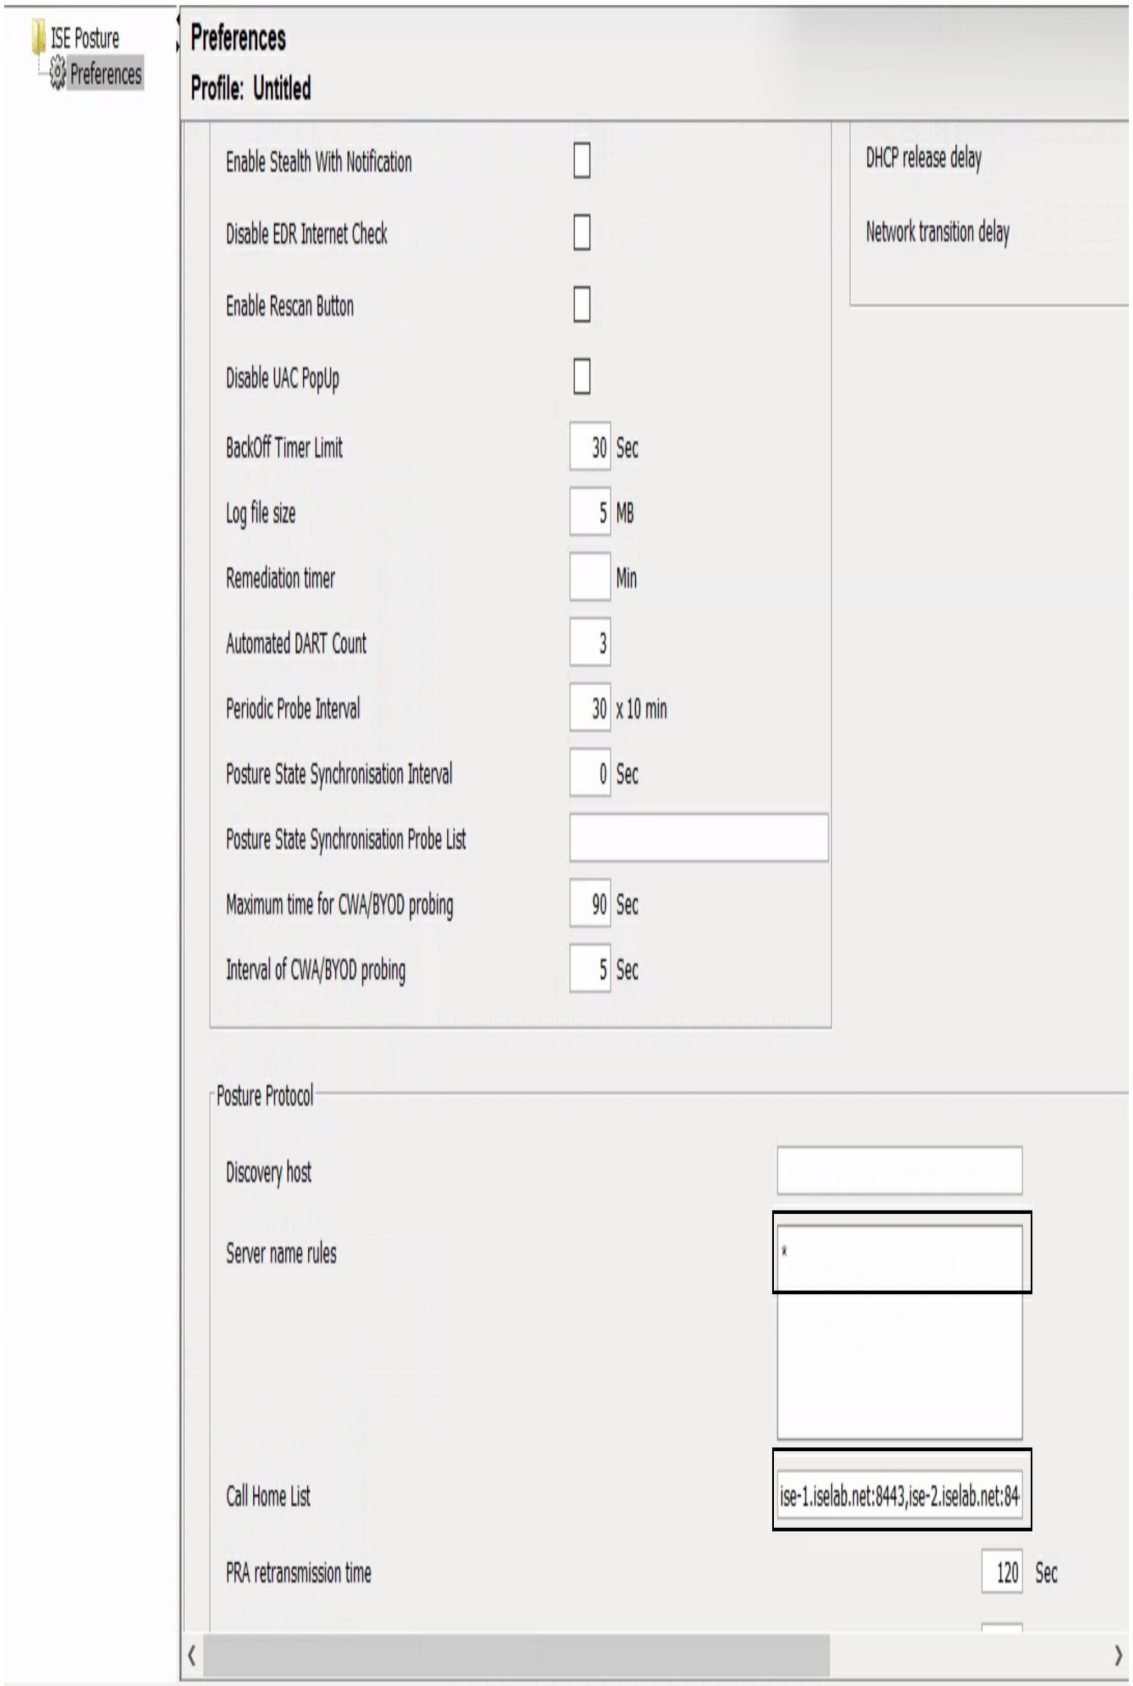
Technet24

Figure 10-12 *Examples of Initial and Final Authorization Policies*

Redirection-less Posture

The redirection-less flow is only possible with the Secure Client ISE posture module. This flow requires you to pre-install the agent and provision the ISE posture profile together with the agent. You can create the posture profile using the Secure Client profile editor. When preparing the profile for the redirection-less flow, add the FQDNs of the ISE Policy Service Nodes to the profile, along with the port number for the client provisioning portal. If you do not specify a port number, the agent will use TCP port 8905 for probes. Connections over this port are protected by the ISE admin certificate.

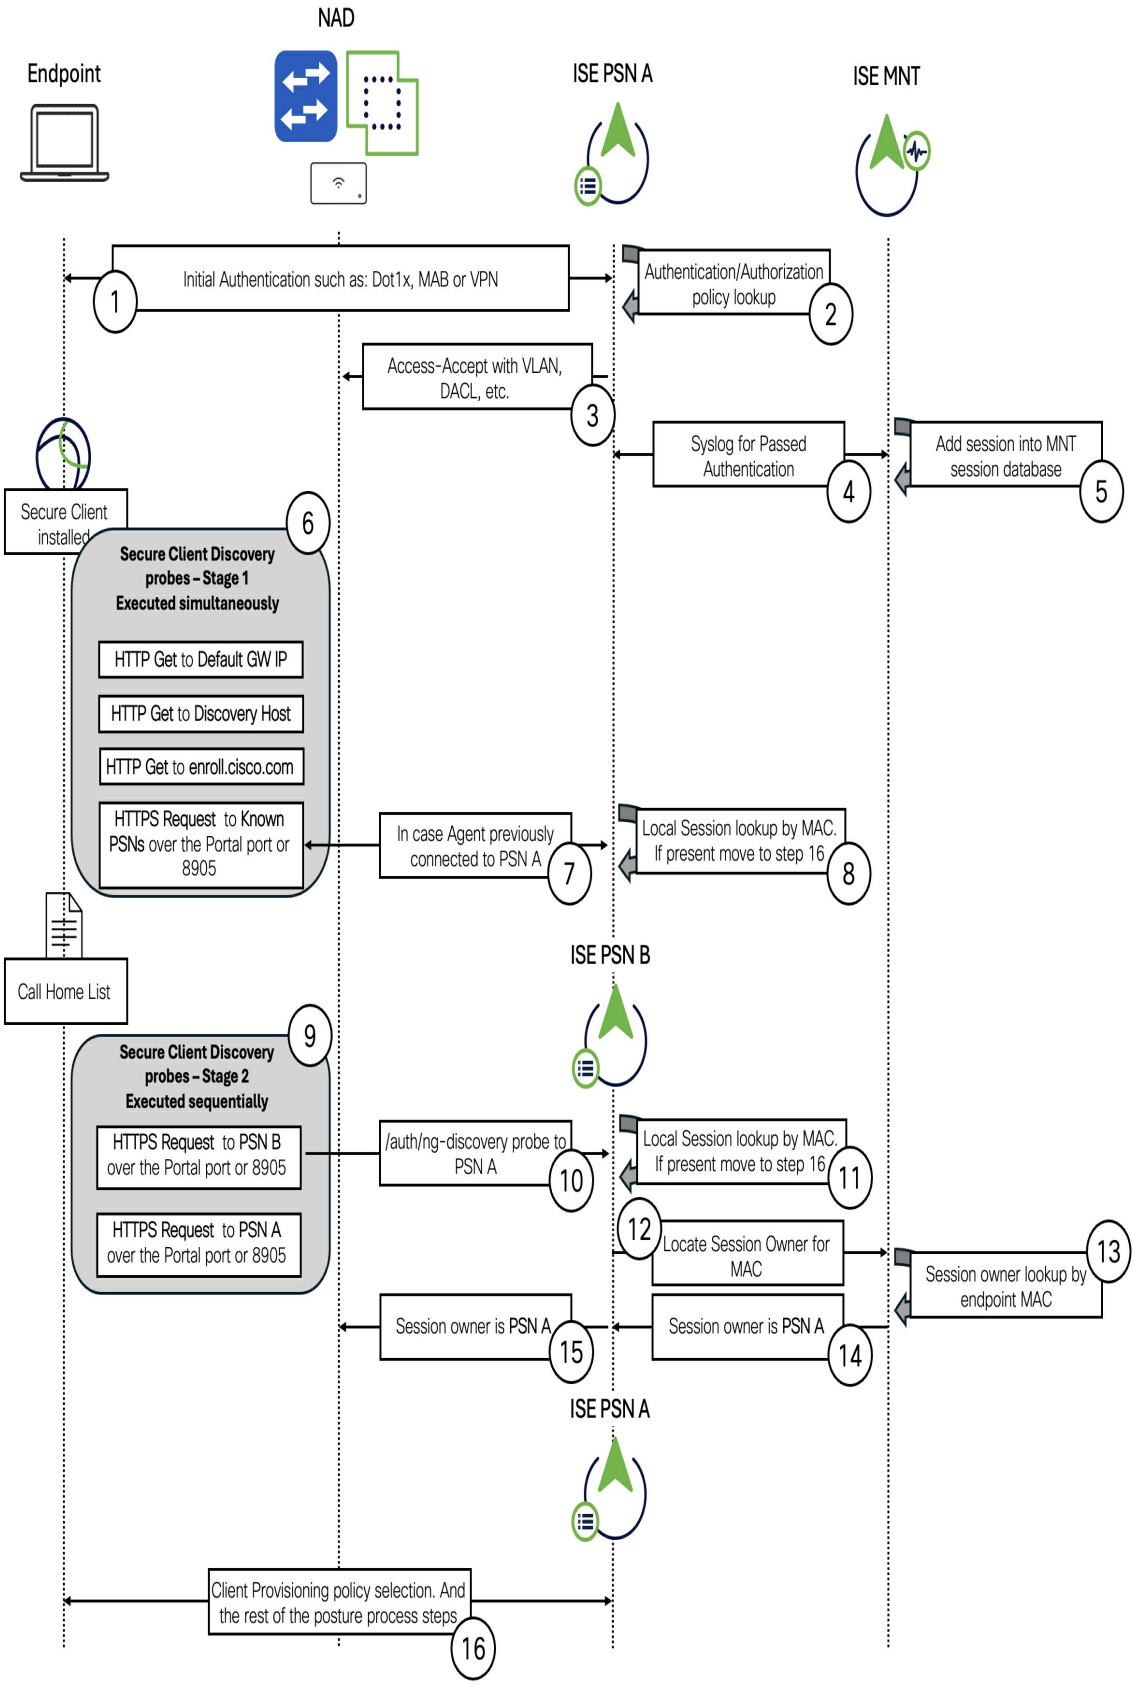
Figure 10-13 provides an example of ISE posture profile configuration. In the **Server Name Rules** section, specify which CN or SAN values in the ISE certificates should be acceptable. To accept any CN or SAN, specify *****. In the Call Home List, you can specify your PSNs' FQDNs with the client provisioning portal port after the colon (:). Use commas to specify multiple PSNs.


ISE Posture
Preferences
Preferences
Profile: Untitled
Enable Stealth With Notification
Disable EDR Internet Check
Enable Rescan Button
Disable UAC PopUp
BackOff Timer Limit
30
Sec
Log file size
5
MB
Remediation timer
Min
Automated DART Count
3
Periodic Probe Interval
30
x 10 min
Posture State Synchronisation Interval
0
Sec
Posture State Synchronisation Probe List
Maximum time for CWA/BYOD probing
90
Sec
Interval of CWA/BYOD probing
5
Sec
DHCP release delay
Network transition delay
Posture Protocol
Discovery host
Server name rules
*
Call Home List
ise-1.iselab.net:8443,ise-2.iselab.net:84
PRA retransmission time
120
Sec

Technet24

Figure 10-13 *Secure Client ISE Posture Profile Example*

The redirection-less posture flow is outlined in Figure 10-14.


NAD
Endpoint
ISE PSN A
ISE MNT
Initial Authentication such as: Dot1x, MAB or VPN
1
Authentication/Authorization policy lookup
2
Access-Accept with VLAN, DACL, etc.
3
Syslog for Passed Authentication
4
Add session into MNT session database
5
Secure Client installed
6
Secure Client Discovery probes – Stage 1
Executed simultaneously
HTTP Get to Default GW IP
HTTP Get to Discovery Host
HTTP Get to enroll.cisco.com
HTTPS Request to Known PSNs over the Portal port or 8905
In case Agent previously connected to PSN A
7
Local Session lookup by MAC. If present move to step 16
8
Call Home List
ISE PSN B
9
Secure Client Discovery probes – Stage 2
Executed sequentially
HTTPS Request to PSN B over the Portal port or 8905
HTTPS Request to PSN A over the Portal port or 8905
/auth/ng-discovery probe to PSN A
10
Local Session lookup by MAC. If present move to step 16
11
12
Locate Session Owner for MAC
13
Session owner lookup by endpoint MAC
Session owner is PSN A
14
Session owner is PSN A
15
ISE PSN A
Client Provisioning policy selection. And the rest of the posture process steps
16

Technet24

Figure 10-14 *Redirection-Less Posture Flow*

Step 1. Initial Dot1x, MAB, or VPN authentication.

Step 2. Initial authentication and authorization policy lookup. The authorization policy should typically include the condition from the session dictionary — PostureStatus not equal to compliant, along with any other required conditions such as user group membership, connection type, and similar attributes.

Step 3. The authorization attributes are returned to the network access device. Since no redirection occurs, you should restrict user access to network resources using the ACL supported by the specific NAD type. Users may also be placed into a designated restricted VLAN.

Step 4. The Syslog message reporting successful authentication is submitted to the primary and secondary Monitoring and Troubleshooting (MNT) nodes.

Step 5. The primary MNT node adds the session to the sessions database. The MNT session database differs from the PSN session cache. The distinction will be explained later in the chapter.

Note

Steps 4 and 5 are present for any authentication and authorization, whether it is regular Dot1x, guest, or redirect-based posture. Policy Service Nodes always report authentication and authorization outcomes for every connection attempt to both Monitoring and Troubleshooting nodes. Previous flows did not include these steps because they were not critical to the execution of the flow.

Step 6. The agent initiates the first set of discovery probes. This set includes all probes explained previously and is referred to as Stage 1 discovery. The probes are executed simultaneously. In environments without redirection, all redirect-based probes will fail.

Step 7. If the agent performed posture assessment successfully in the past, it can use the list of previously connected Policy Service Nodes to send direct discovery probes to them.

Step 8. Each Policy Service Node that receives the probe performs a lookup in the local session cache to determine if it owns the session for the given endpoint. For this lookup, the PSN uses the endpoint's MAC and IP addresses provided in the probe.

Step 9. In environments without redirection, and when the endpoint has not performed a successful posture assessment in the past, the agent moves to the next set of probes. This time, the agent sequentially uses values specified in the Call Home List. These probes use the URL /auth/ng-discovery, enabling Policy Service Nodes to distinguish between regular posture probes and those specifically designed for the non-redirect flow. As a fallback, if posture was successful previously, the agent sends probes to the PSNs listed in ConnectionData.xml using the /auth/ng-discovery URL. If all probes fail, the agent sends a request to https://enroll.cisco.com:8905 as a last resort. For this probe to work, your DNS system must resolve enroll.cisco.com to one of the PSN IP addresses. Additionally, to avoid certificate warnings, the PSN admin certificate must include enroll.cisco.com as a Subject Alternative Name. This set of probes is called Stage 2

Technet24

discovery.

Step 10. One of the /auth/ng-discovery probes reaches the Policy Service Node.

Step 11. PSN performs a lookup in the local session cache to determine if it owns the session for the given endpoint. For this lookup, the PSN uses the endpoint's MAC and IP addresses provided in the probe.

Step 12. If the Policy Service Node does not have a session for the endpoint, the nature of the /auth/ng-discovery probe instructs the PSN to forward the request to the primary Monitoring and Troubleshooting node. The request to the Monitoring and Troubleshooting node contains only the endpoint's MAC address(es) and does not include the IP address(es).

Step 13. The Monitoring and Troubleshooting node performs a session lookup to identify the Policy Service Node that owns the session for the endpoint.

Note

Since only the MAC address can be used for session lookup on the Monitoring and Troubleshooting node, redirection-less posture typically cannot be used with VPN on third-party network access devices. In most implementations, the RADIUS server receives only the endpoint IP information during VPN authentication.

Step 14. Information about the session owner is returned to the Policy Service Node that submitted the request.

Step 15. The Policy Service Node returns information about the session owner to the agent.

Step 16. The agent sends a request that triggers client provisioning policy lookup and subsequent update of required modules and profiles if necessary. From this step onward, the posture process continues in the same manner as described for the redirect-based posture flow.

Agentless Posture

Agentless posture can be implemented for Windows and macOS devices and requires several important prerequisites on the endpoint and NAD sides. The prerequisites are as follows:

- Network access devices must be configured to include the client IP address in RADIUS accounting updates.
- For Windows clients, port 5985 must be open to allow PowerShell access on the client, and PowerShell must be version 5.1 or above.
- For macOS clients, port 22 must be open to allow SSH access to the client, and curl must be version 7.34 or above.
- Client credentials for shell login must have local administrator privileges.
- The following entry should be updated in the sudoers file to prevent certificate installation failure on the endpoints: **ALL = (ALL) NOPASSWD: /usr/bin/security, /usr/bin/osascript**.
- For macOS, the user account configured must be an administrator account.

The agentless posture flow is shown in Figure 10-15.

Technet24

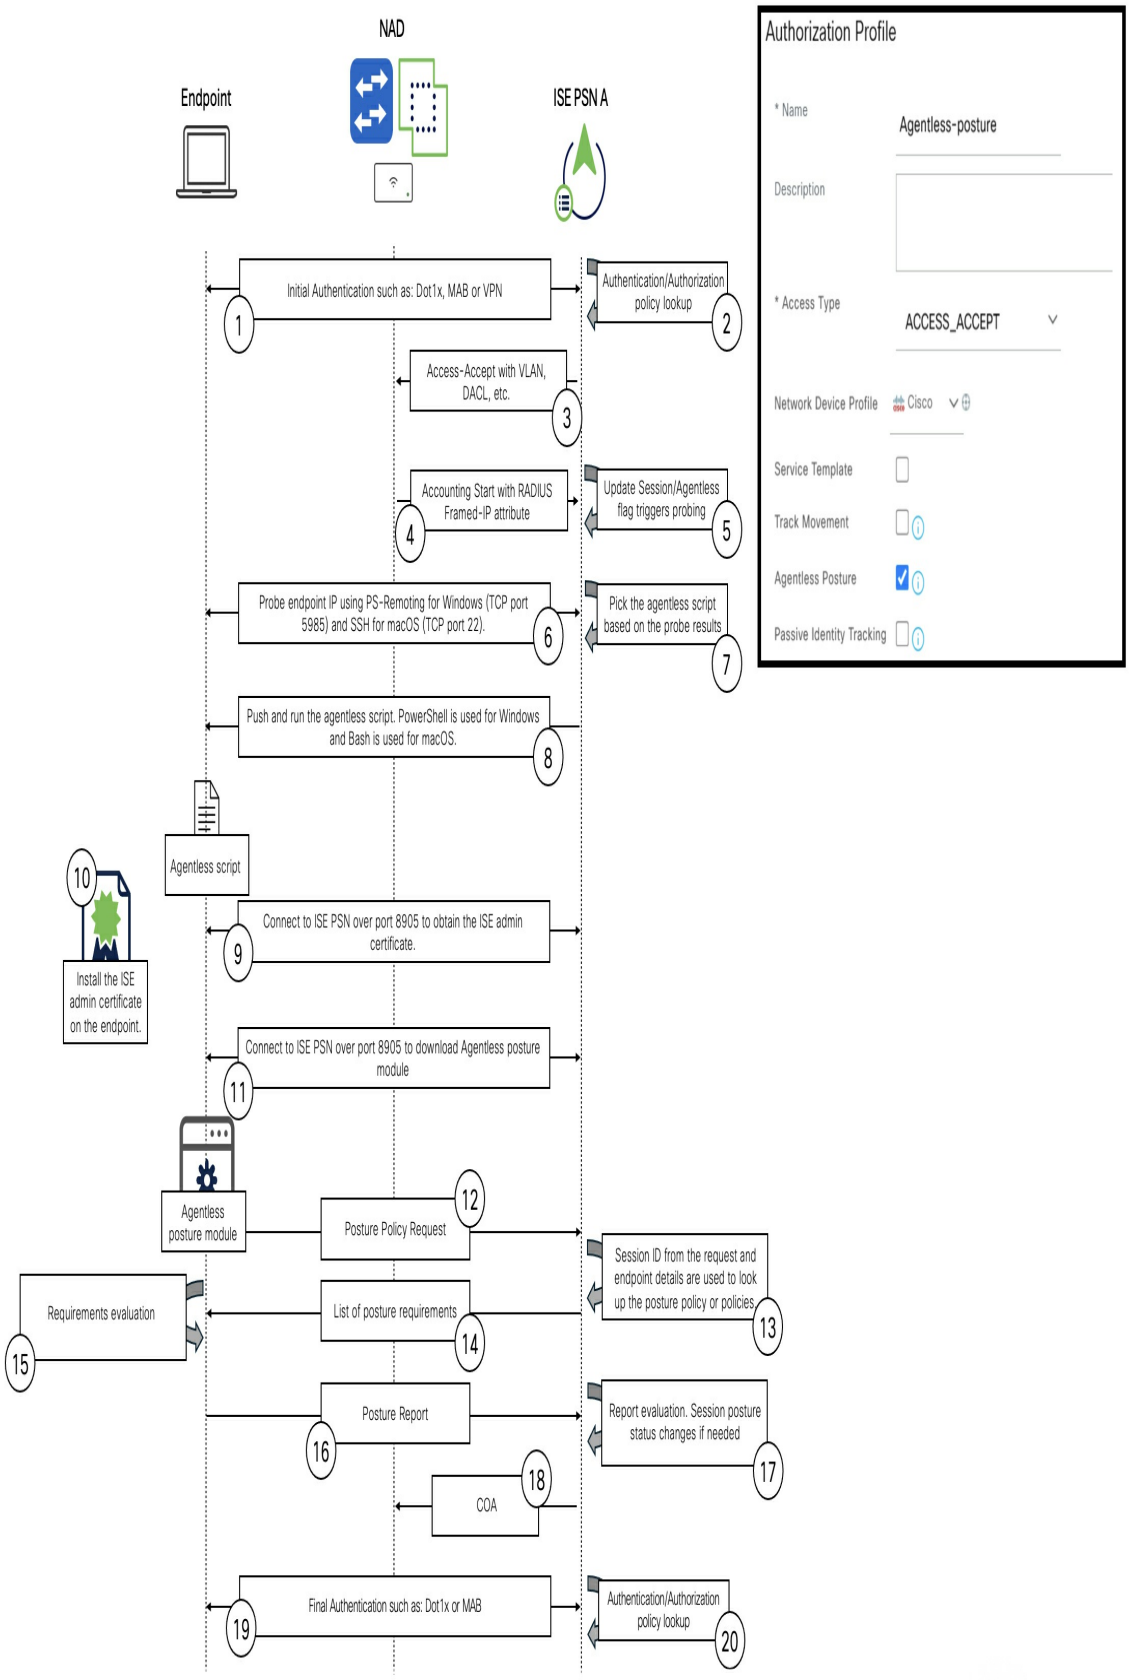

NAD
Endpoint
ISE PSN A
Initial Authentication such as: Dot1x, MAB or VPN
1
Authentication/Authorization policy lookup
2
Access-Accept with VLAN, DACL, etc.
3
Accounting Start with RADIUS Framed-IP attribute
4
Update Session/Agentless flag triggers probing
5
Probe endpoint IP using PS-Remoting for Windows (TCP port 5985) and SSH for macOS (TCP port 22).
6
Pick the agentless script based on the probe results
7
Push and run the agentless script. PowerShell is used for Windows and Bash is used for macOS.
8
Agentless script
10
Install the ISE admin certificate on the endpoint.
Connect to ISE PSN over port 8905 to obtain the ISE admin certificate.
9
Connect to ISE PSN over port 8905 to download Agentless posture module
11
Agentless posture module
Posture Policy Request
12
Session ID from the request and endpoint details are used to look up the posture policy or policies
13
Requirements evaluation
15
List of posture requirements
14
Posture Report
16
Report evaluation. Session posture status changes if needed
17
COA
18
Final Authentication such as: Dot1x or MAB
19
Authentication/Authorization policy lookup
20
Authorization Profile
* Name
Agentless-posture
Description
* Access Type
ACCESS_ACCEPT
Network Device Profile
Cisco
Service Template
Track Movement
Agentless Posture
Passive Identity Tracking

Figure 10-15 *Agentless Posture Flow*

Step 1. Initial Dot1x, MAB, or VPN authentication.

Step 2. Initial authentication and authorization policy lookup. For the agentless posture flow to function, the selected authorization profile must have the Agentless Posture checkbox enabled, as shown in Figure 10-15. This setting instructs the PSN to tag a session in a specific manner and monitor it for incoming accounting messages with the endpoint IP address present in the RADIUS Framed-IP-Address field.

Step 3. The authorization attributes are returned to the network access device. Since no redirection occurs, you should restrict user access to network resources using the ACL supported by the specific NAD type. Users may also be placed into a designated restricted VLAN.

Step 4. Accounting start or interim accounting update from the network access device delivers the endpoint IP address in the RADIUS Framed-IP-Address field to the Policy Service Node.

Step 5. When the endpoint IP address is added to a session marked as part of the agentless flow, this instructs the Policy Service Node to probe the endpoint using both SSH and PowerShell Remoting.

Step 6. SSH and PowerShell Remoting reach the endpoint. The PSN records which connection was successful to identify the endpoint type.

Step 7. Based on the probe result, the Policy Service Node selects the type of script that should be used to continue the flow.

Technet24

Step 8. The agentless script is pushed to the endpoint and executed there.

Step 9. The agentless script sends an HTTPS request to the Policy Service Node over port 8905 to retrieve the PSN admin certificate.

Step 10. The Policy Service Node certificate is added as trusted on the endpoint side.

Step 11. The agentless script connects to the Policy Service Node using port 8905 to download the agentless posture module.

Step 12. The agent submits a request to obtain the list of posture policies and requirements for evaluation. This request includes audit session ID, operating system details, and the agent type/version.

Step 13. On the PSN side, details from the request are used to perform posture policy selection. By default, ISE provides multiple preconfigured posture policies, which you can use as is or modify for your specific use case.

Step 14. The list of posture requirements and their corresponding remediation actions is returned to the agent.

Step 15. The agent verifies the requirements locally using the compliance module and prepares the posture assessment report.

Step 16. The posture report is submitted to ISE.

Step 17. ISE evaluates the report according to the posture policy settings. If any mandatory conditions defined in the posture policy fail, the agent is notified about noncompliance, and the remediation process starts on the

agent side. If all mandatory conditions are met, the session status changes to compliant.

Step 18. ISE sends a CoA to the network access device to trigger reauthentication, resulting in a subsequent authentication attempt. For posture over VPN, CoA Push is used instead of CoA Reauth to keep the user connected.

Step 19. The final authentication happens using Dot1x or MAB.

Step 20. On the PSN side, authentication and authorization policy lookup is performed. At this stage, the authorization policy will typically have the PostureStatus from the session dictionary set to compliant, along with other attributes used in the initial authorization policy.

Stitching Attributes for Authorization

While guest access involves stitching a wide variety of attributes, in the posture use case, attribute stitching is limited to one attribute called PostureStatus, which is present in the session dictionary. The dictionary name immediately suggests that this attribute relates to the "session" and is likely associated with what was described as "session attributes" in Chapter 9, "Managing Guest Users and Personal Devices." This description is both accurate and misleading. PostureStatus is one of those elements where the answer is "It's complicated" or "It depends."

In essence, for every new session, the PostureStatus attribute should be "Pending/Unknown," since the Policy Service Node runtime has not yet communicated with the posture agent to determine the actual posture status. After the Change of Authorization, the PostureStatus is typically either "Compliant" or "Not compliant." However, the reality is more complex than this simple rule. The PostureStatus value depends

Technet24

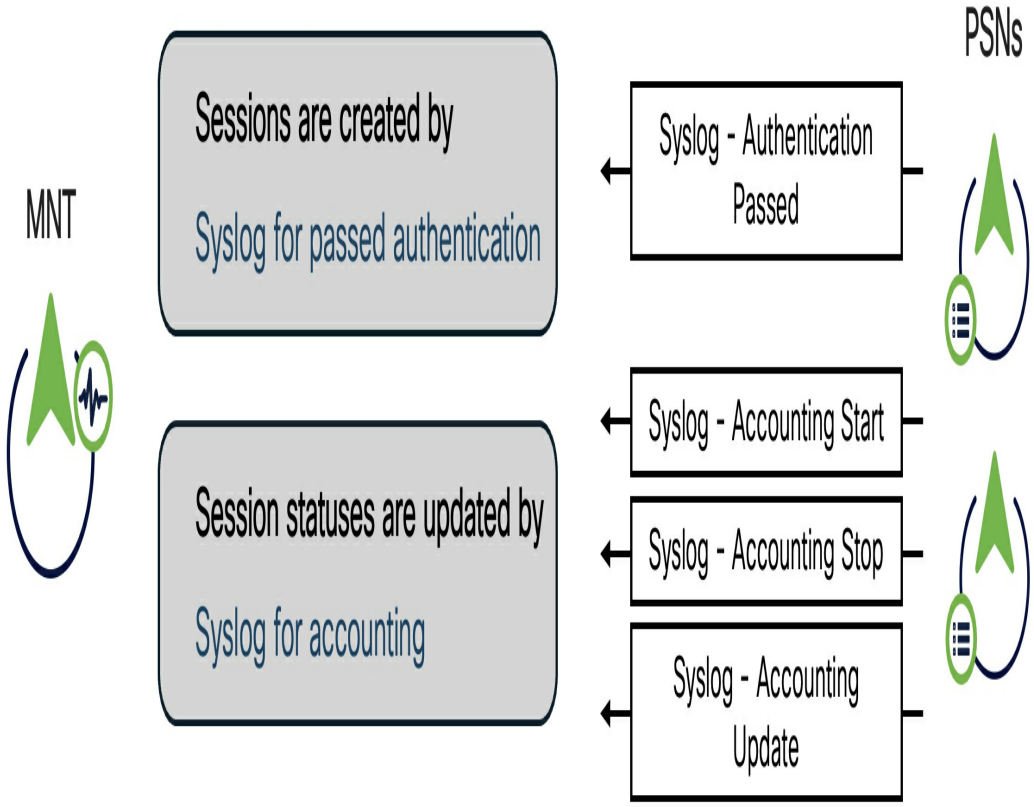
on ISE session management logic as well as specific endpoint attributes. Understanding all relevant components is essential for a complete view of how the ISE posture service operates. The following section explores the details of how the PostureStatus value is determined.

ISE Session Management and Posture

Before we examine how the PostureStatus attribute is assigned, it is important to understand that PostureStatus is a session attribute. When a session is terminated and no longer exists, the PostureStatus attribute is also removed. In an ISE deployment, two different ISE personas manage the session lifecycle: Monitoring and Troubleshooting (MNT) and Policy Service Node (PSN).

MNT manages the session lifecycle globally across the deployment, while each PSN is responsible only for sessions terminated on it by network access devices. The posture runtime relies on the session and its attributes stored on the PSN side, but a complete understanding of session management across all responsible ISE personas is essential.

Figure 10-16 demonstrates the session lifecycle management logic on the MNT side.



Rules for sessions removal

a. Sessions without accounting start (Authenticated) removed after 60 minutes,

b. Sessions with accounting stop (Terminated) removed after 15 minutes

c. Sessions in 'Started' state (MNT got accounting start) removed after 120 hours without Interim update.

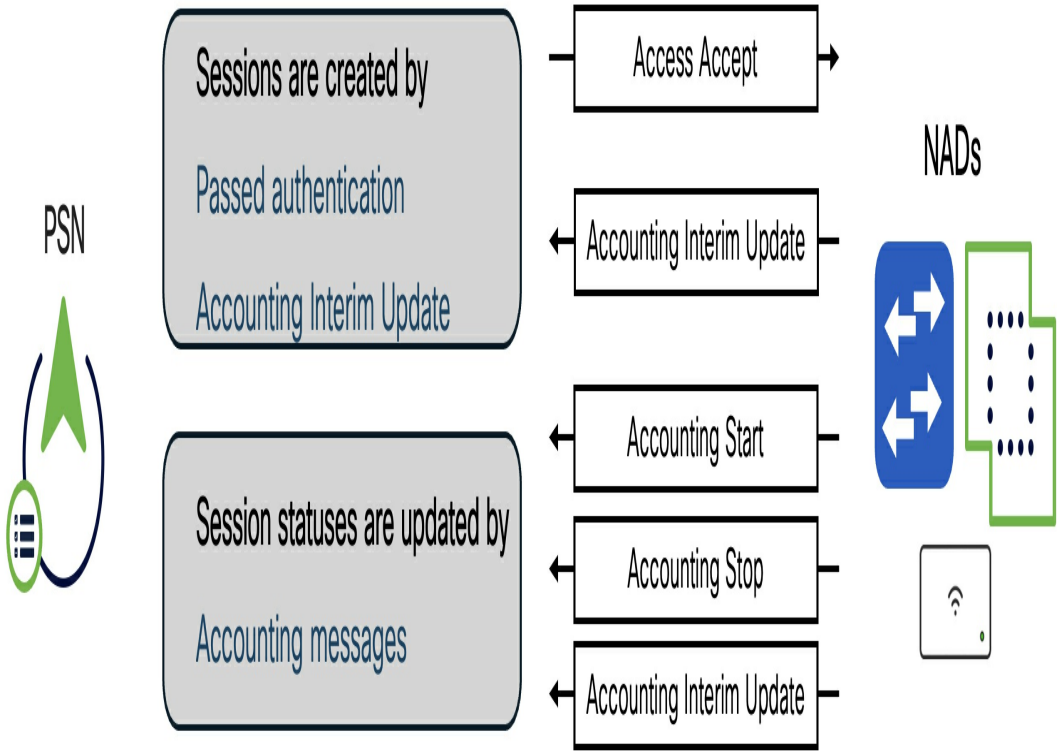
Technet24

Figure 10-16 *Session Management Lifecycle on MNT*

On the MNT node, the following logic applies to the session management process:

- Sessions are created by Passed Authentication Syslog messages from Policy Service Nodes.
- Sessions are removed using the following logic:
  - Authenticated sessions (those where an accounting start syslog event has not been received) are removed after 60 minutes.
  - Terminated sessions (those for which an accounting stop syslog event was processed by MNT) are removed after 15 minutes.
  - Started sessions (those for which MNT received an accounting start syslog event) are removed after 120 hours without any new accounting interim update syslog event.

The logic of session lifecycle management performed by the PSN is displayed in Figure 10-17.



Rules for sessions removal

a. Sessions are removed upon processing Accounting stop

b. Least recently used sessions are removed after reaching platform limit

c. Session cache is cleared upon PSN reload or Application Server restart

Technet24

Figure 10-17 *Session Management Lifecycle on MNT*

On the Policy Service Node, the following logic applies to the session management process:

- Sessions are created by Passed Authentication events resulting from processing authentication requests from network access devices.
- Accounting Interim Update messages can create an entry in the session cache if no session exists for the given session ID or endpoint MAC address.
- Sessions are removed using the following logic:
  - Accounting Stop events remove the specific session from the cache.
  - If the session cache is full and a new session needs to be added, the PSN uses the Least Recently Used (LRU) algorithm to remove the session that has not received updates for the longest period. Session cache capacity depends on the ISE hardware or VM specifications. This information is available in Table 2-2, "Maximum Concurrent Active Sessions for Different ISE Appliances Acting as PSNs," in Chapter 2, "ISE Deployment Principles."
  - The entire content of the session cache is cleared by a PSN or application service restart.

The session management process on the Policy Service Node side may result in two scenarios that can directly affect the posture process. Although these scenarios are rare, you can further reduce their likelihood by following best practices on network access devices and load balancers. However, they cannot be completely avoided. These

scenarios include the following:

- **Stale session:** This situation occurs when a session remains in the PSN session cache because an accounting stop message was not received. It typically happens due to the accounting stop message being forwarded to the wrong PSN, often as a result of load balancer stickiness failure or a brief network interruption between the NAD and PSN. A stale session can persist in the PSN session cache for a significant amount of time until it is removed by the LRU algorithm or by a PSN or application service restart. When posture services are enabled, a stale session retains the PostureStatus established for the endpoint previously. Later probes that reach the PSN directly may be influenced by outdated information, which, from an administrator's perspective, appears as follows: ISE live logs show that initial posture authorization is assigned while the agent reports to the end user that it is in the "Compliant" state. The best protection against stale sessions is properly configured stickiness on the load balancer, ensuring that updates for the same session always reach the same PSN, and a network path with sufficient redundancy between NAD and PSN.
- **Phantom session:** This type of session is another anomaly in the session cache, created when an accounting interim update is misrouted to a different PSN. Phantom sessions are typically generated for the same reasons as stale sessions and present similar risks during posture assessment. You can mitigate these risks to some extent by applying the same best practices.

As explained, neither phenomenon can be completely avoided, because it is impossible to guarantee that each accounting stop and interim update will always reach the correct PSN. Therefore, it is recommended to always use redirect-based posture, since the agent prefers results from the redirect-based probe, which always points to

Technet24

the correct PSN and session.

For non-redirect posture, you should build authorization policies and profiles so that probes can only reach the PSN where the current session was authenticated. You can achieve this by combining ACLs supported by the given NAD type with authorization policies applied when authentication reaches a specific PSN. In this approach, the number of initial authorization policies equals the number of PSNs in your deployment. Figure 10-18 shows an example of authorization policies, profiles, and ACLs recommended for redirection-less posture. The example is based on DACLs, a Cisco NAD feature, so you can adjust the configuration according to the same logic to fit the requirements of your specific NAD type.


Authorization Profile
* Name
NG-Posture-ISE-2
Description
* Access Type
ACCESS_ACCEPT
Common Tasks
DACL Name
NG-Posture-Unknown-ISE-2
* Name
NG-Posture-Unknown-ISE-2
Description
IP version
IPv4
IPv6
Agnostic
* DACL Content
1 remark allow access to ISE-2
2 permit ip any host 172.22.15.21
3 remark allow DNS
4 permit udp any any eq 53
5 remark permit dhcp
6 permit udp any any eq bootpc any eq bootps
7 remark allow remediation servers
8 permit ip any host 172.22.15.10
NG-Posture-Compliant
AND
Wireless_802.1X
Session-PostureStatus EQUALS Compliant
PermitAccess
NG-Posture-Unknown-ISE-2
AND
Wireless_802.1X
Network Access-ISE Host Name EQUALS ise-2.iselab.net
Session-PostureStatus NOT_EQUALS Compliant
NG-Posture-ISE-2
NG-Posture-Unknown-ISE-1
AND
Wireless_802.1X
Network Access-ISE Host Name EQUALS ise-1.iselab.net
Session-PostureStatus NOT_EQUALS Compliant
NG-Posture-ISE-1
Authorization Profile
* Name
NG-Posture-ISE-1
Description
* Access Type
ACCESS_ACCEPT
Common Tasks
DACL Name
NG-Posture-Unknown-ISE-1
* Name
NG-Posture-Unknown-ISE-1
Description
IP version
IPv4
IPv6
Agnostic
* DACL Content
1 remark allow access to ISE-1
2 permit ip any host 172.22.15.11
3 remark allow DNS
4 permit udp any any eq 53
5 remark permit dhcp
6 permit udp any any eq bootpc any eq bootps
7 remark allow remediation servers
8 permit ip any host 172.22.15.10

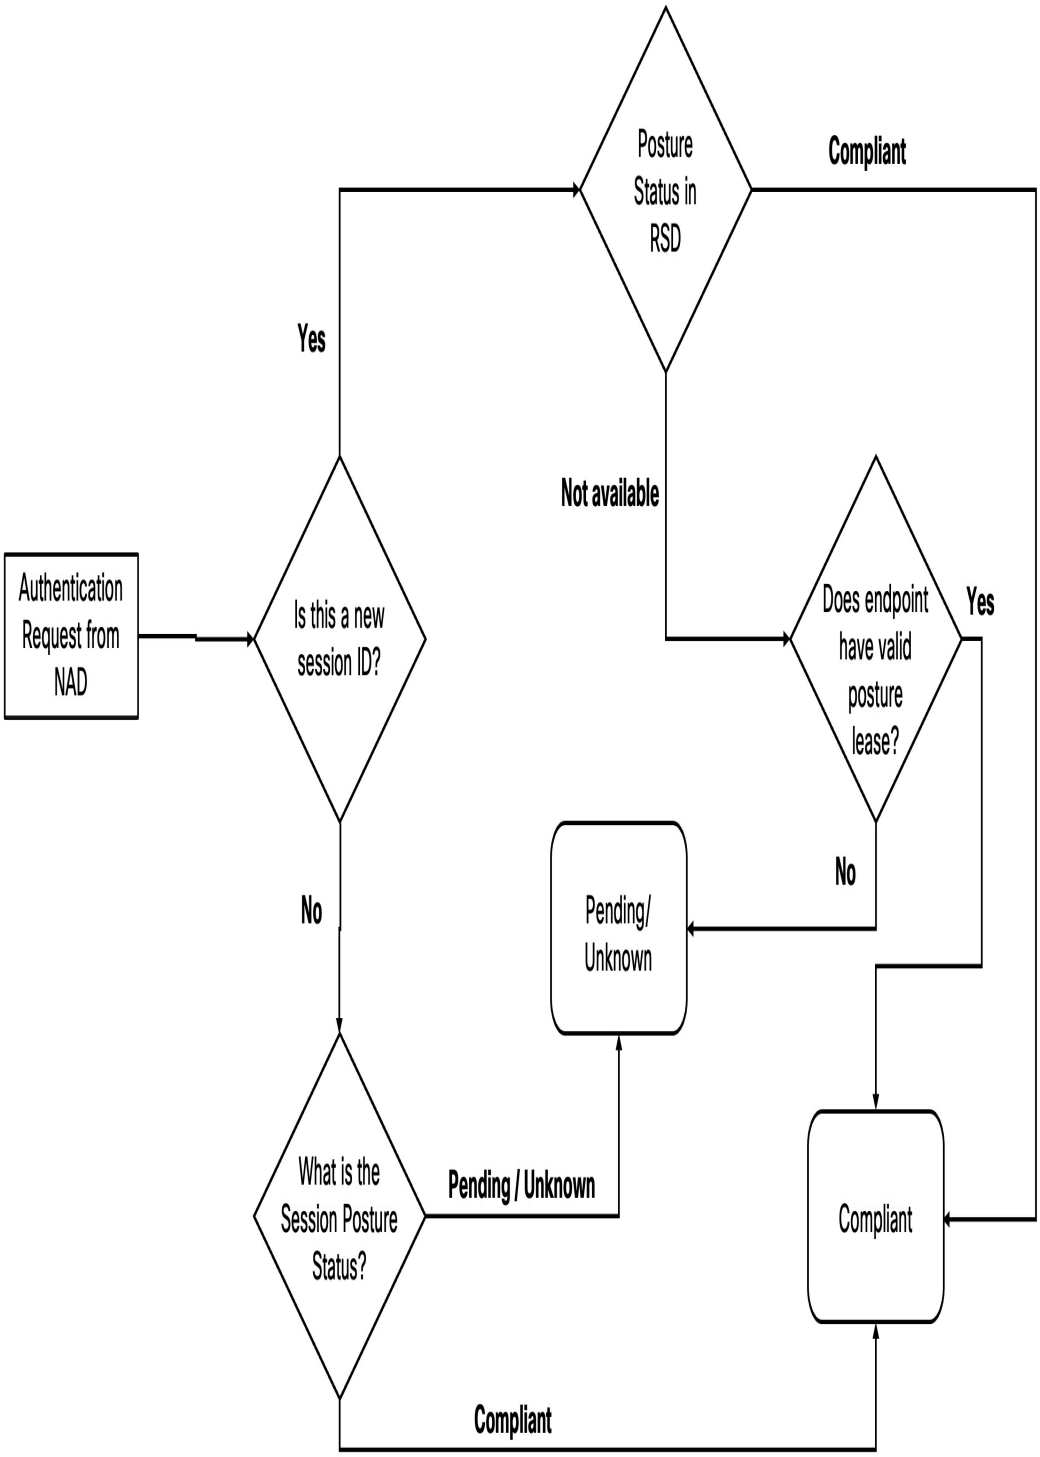
Technet24

Figure 10-18 *Authorization Policies Recommended for Redirection-Less Posture*

In these examples, initial authorization policies use the attribute ISE hostname from the Network Access dictionary. This ensures that authentication on each individual PSN receives its own authorization profile. In the authorization profiles, ACLs are configured so that when authentication occurs on PSN ise-1.iselab.net, probes from the posture agent can reach only ISE-1; when authentication occurs on ISE-2, probes are limited to ISE-2. This approach protects the agent from receiving the wrong posture status from a PSN with a stale or phantom session.

Logic Behind Posture Status

When an authentication request is processed and authorization policies include the PostureStatus attribute, the Policy Service Node runtime must decide and assign the PostureStatus value. This assignment occurs before the agent communicates with the PSN. The PostureStatus obtained from the actual posture assessment process is then used to update the PostureStatus value, which results in issuing a Change of Authorization and leads to the next authentication and authorization. The logic of PostureStatus value selection at the time of authentication and authorization is illustrated in Figure 10-19.


Authentication Request from NAD
Is this a new session ID?
Yes
No
Posture Status in RSD
Compliant
Not available
Does endpoint have valid posture lease?
Yes
No
Pending/ Unknown
What is the Session Posture Status?
Pending / Unknown
Compliant
Compliant

Technet24

Figure 10-19 *PostureStatus Selection Logic*

When authentication and authorization are processed, the Policy Service Node first checks if the request is for an existing session. If the request arrives for a session already present in the session cache, the existing PostureStatus value is used. If the request arrives for a new session ID, the PSN checks whether this session ID is present in the RADIUS Session Directory (RSD).

The RSD relies on Light Data Distribution (LDD) to replicate a limited set of session attributes, including PostureStatus, between all PSNs in the deployment. It is important to note that the PSN session cache remains local to each PSN, while the RSD only replicates specific attributes. If PostureStatus for the session is set as compliant on another PSN, the local PSN will use this value, which prevents issues related to stale and phantom sessions. If the session ID is not present in the RSD, the algorithm continues to the next step.

RSD and LDD features are enabled by default but can be disabled by navigating to **Administration > Settings > Light Data Distribution**.

As a final step, the PSN verifies whether the endpoint has a valid posture lease value. When the posture lease is valid, the session automatically receives a compliant posture status. If the posture lease is not present, the posture status is set as pending/unknown.

ISE Posture Lease

By default, ISE is configured to perform posture assessment for each new session ID. In dynamic environments where authentication may occur frequently, such as during Wi-Fi roaming, this behavior may not be preferred.

As an alternative, the posture lease feature can be used. When a

posture lease is enabled, the ISE Policy Service Nodes save the posture lease expiration timestamp into the endpoint attributes. During the period between the initial posture assessment and lease expiration, PostureStatus is assigned as compliant for every authentication attempt for the same endpoint. When the lease expires, the session for the next authentication after lease expiration will receive PostureStatus equal to pending/unknown, which triggers assignment of the restricted authorization profile, typically with web redirection attributes. This, in turn, results in a full posture assessment performed on the endpoint.

You can enable the posture lease by navigating to **Work Centers > Posture > Settings > Posture General Settings** and changing the setting in the **Posture Lease** section from **Perform Posture Assessment Every Time a User Connects to the Network** to **Perform Posture Assessment Every** with the preferred interval, as shown in Figure 10-20.

Technet24

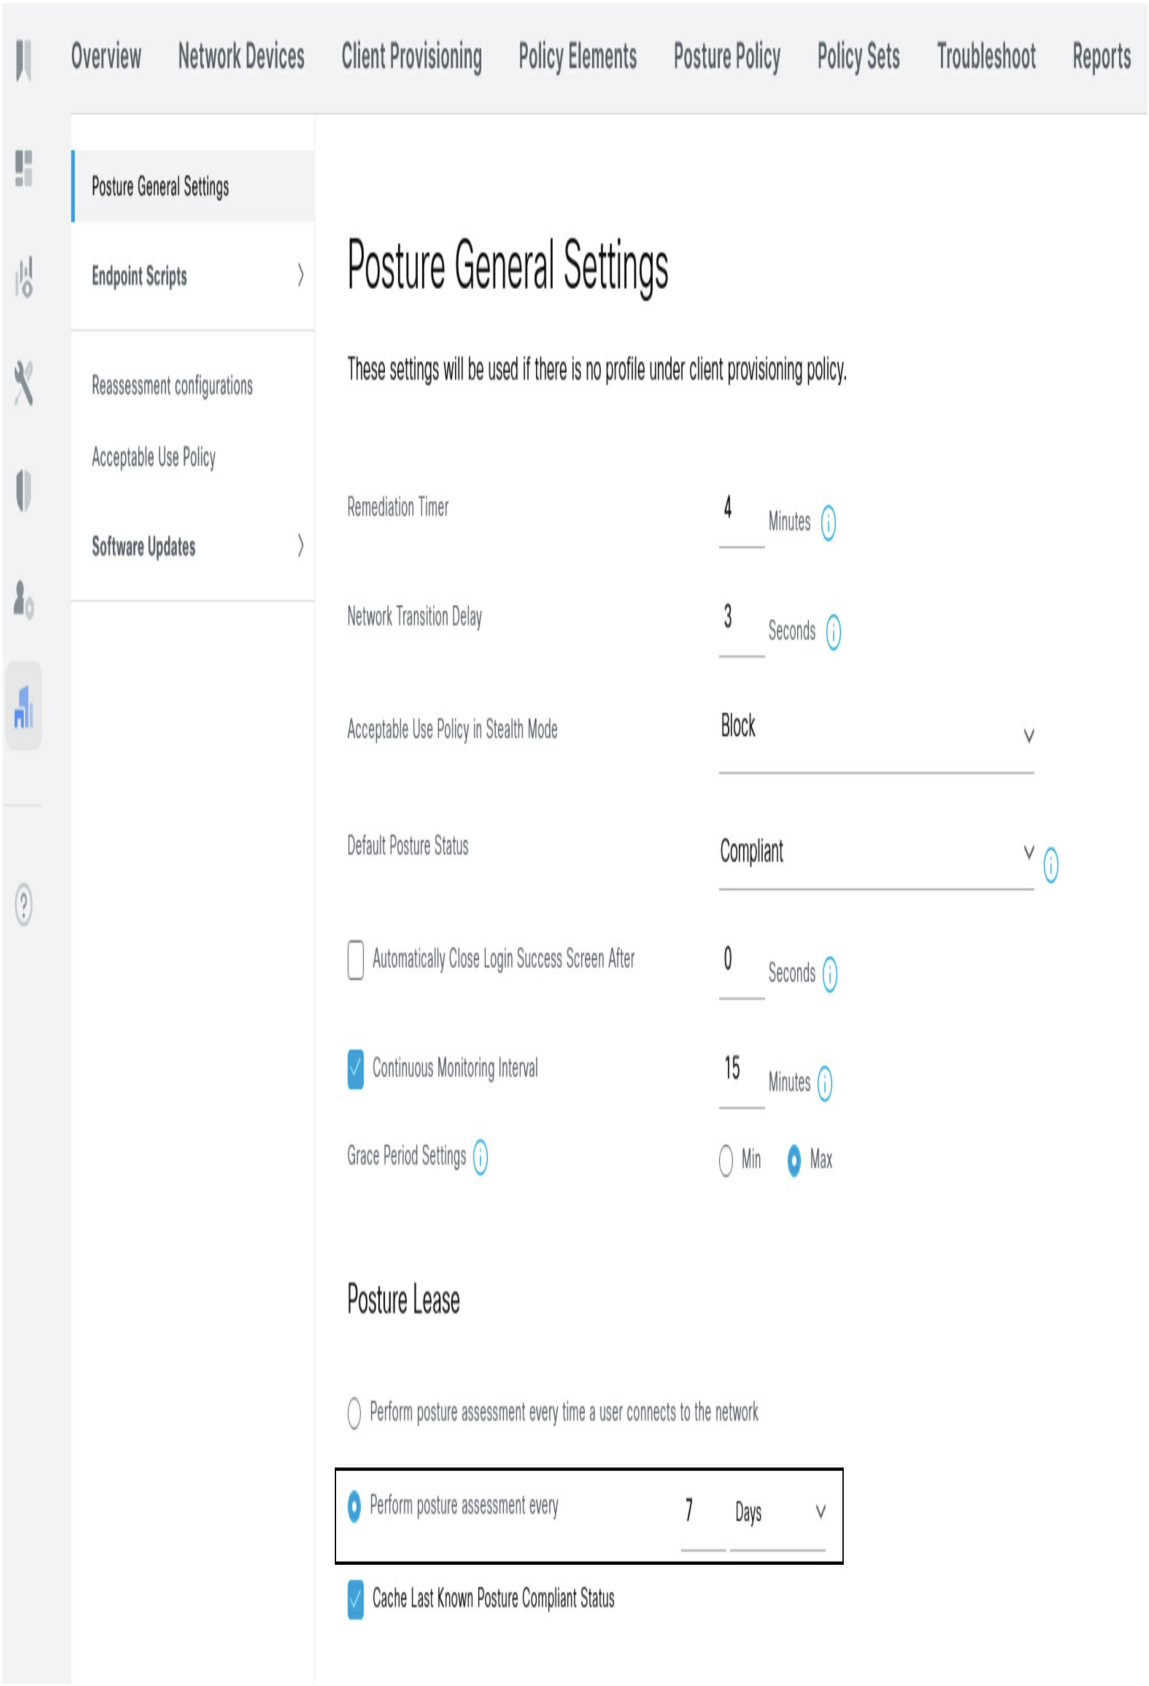

Overview
Network Devices
Client Provisioning
Policy Elements
Posture Policy
Policy Sets
Troubleshoot
Reports
Posture General Settings
Endpoint Scripts
Reassessment configurations
Acceptable Use Policy
Software Updates
Posture General Settings
These settings will be used if there is no profile under client provisioning policy.
Remediation Timer
4
Minutes
Network Transition Delay
3
Seconds
Acceptable Use Policy in Stealth Mode
Block
Default Posture Status
Compliant
Automatically Close Login Success Screen After
0
Seconds
Continuous Monitoring Interval
15
Minutes
Grace Period Settings
Min
Max
Posture Lease
Perform posture assessment every time a user connects to the network
Perform posture assessment every
7
Days
Cache Last Known Posture Compliant Status

Figure 10-20 *Posture Lease Configuration*

The logic behind the posture lease feature is illustrated in Figure 10-21.

Technet24

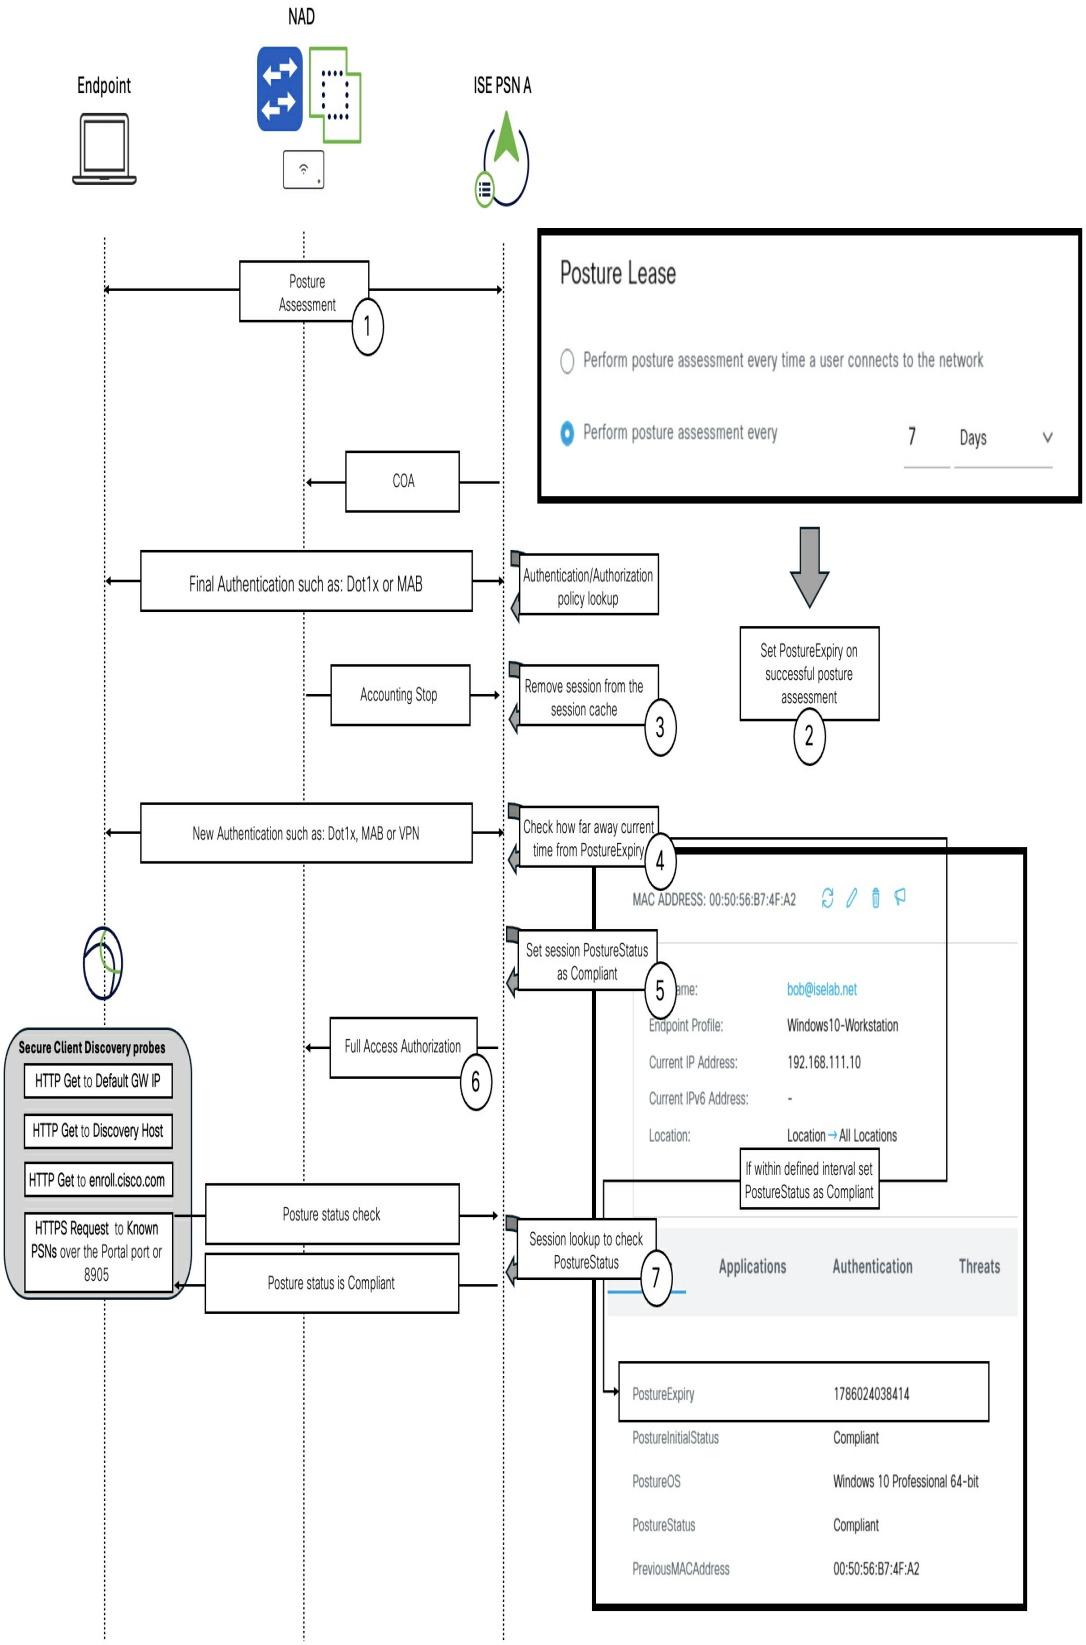

NAD
Endpoint
ISE PSN A
Posture Assessment
1
Posture Lease
Perform posture assessment every time a user connects to the network
Perform posture assessment every
7
Days
COA
Final Authentication such as: Dot1x or MAB
Authentication/Authorization policy lookup
Set PostureExpiry on successful posture assessment
2
Accounting Stop
Remove session from the session cache
3
New Authentication such as: Dot1x, MAB or VPN
Check how far away current time from PostureExpiry
4
MAC ADDRESS: 00:50:56:B7:4F:A2
Set session PostureStatus as Compliant
5
bob@iselab.net
Endpoint Profile:
Windows10-Workstation
Current IP Address:
192.168.111.10
Current IPv6 Address:
-
Location:
Location→All Locations
Full Access Authorization
6
If within defined interval set PostureStatus as Compliant
Secure Client Discovery probes
HTTP Get to Default GW IP
HTTP Get to Discovery Host
HTTP Get to enroll.cisco.com
HTTPS Request to Known PSNs over the Portal port or 8905
Posture status check
Posture status is Compliant
Session lookup to check PostureStatus
7
Applications
Authentication
Threats
PostureExpiry
1786024038414
PostureInitialStatus
Compliant
PostureOS
Windows 10 Professional 64-bit
PostureStatus
Compliant
PreviousMACAddress
00:50:56:B7:4F:A2

Figure 10-21 *ISE Posture Lease Logic*

The following steps explain the posture lease logic. Standard posture process steps are omitted; only aspects relevant to posture lease are covered.

Step 1. Posture assessment occurs between the agent and the Policy Service Node.

Step 2. In the case of a successful posture assessment with posture lease enabled, the PostureExpiry attribute is added to the endpoint. This attribute contains the expiration date and time in the Milliseconds Epoch format. This timestamp marks when the posture lease is considered expired and a full posture process must be performed again. To view the timestamp, navigate to **Context Visibility > Endpoints** and open the endpoint where PostureExpiry needs to be verified. Any standard Epoch timestamp converter can be used to obtain the actual posture lease expiration date and time.

Step 3. When the endpoint is disconnected from the network, the session is removed from the PSN session cache along with all associated attributes, including PostureStatus.

Step 4. As part of the next authentication, the PostureExpiry value is checked to determine the PostureStatus value.

Step 5. If the current timestamp is within the Posture Lease window, the session receives PostureStatus assigned as compliant.

Step 6. Since PostureStatus is set to compliant, the authorization policy with full network access is selected.

Step 7. During the PSN discovery process, the agent sends a

Technet24

request to previously connected PSNs, and the current PSN responds that the endpoint is already in a compliant state. This cancels the posture process on the endpoint side, and posture status is displayed as compliant in the agent's UI.

PRA, CRA, and Grace Period

To better support long-lasting network connections on the endpoint side and handle scenarios where a policy becomes noncompliant but network access must still be provided during a "grace period," ISE supports three features:

- **Posture reassessment (PRA):** This feature allows the Policy Service Node to instruct the agent, after successful posture assessment, to periodically contact the PSN and request a new posture assessment. PRA is typically used in environments where posture lease is enabled, ensuring that endpoint noncompliance can be detected without waiting for the full posture lease interval to expire.
- **Continuous reassessment (CRA):** This feature is similar to PRA but operates much faster. On the agent side, certain events are monitored continuously, and changes are reported to the Policy Service Node in near real-time. Other events and changes are monitored regularly and reported to the PSN at short polling intervals, typically every 10 minutes. CRA is significantly faster, because the smallest interval supported for PRA is 60 minutes. CRA is supported starting with ISE 3.5 patch 3 and requires Secure Client 5.1.17.3382 or later.
- **Grace period:** This feature allows configuration of a grace period for specific posture policies. The grace period enables the endpoint to maintain "compliant" network access if it was

compliant previously, when requirements specified in the posture policy are not met during a designated time interval. This feature is useful for lengthy but less critical posture checks, such as patch management verification. The endpoint can maintain network access if updates require significant time to install.

Now, we will discuss each feature in more detail.

Posture Reassessment (PRA)

The logic of PRA operations is shown in Figure 10-22.

Technet24

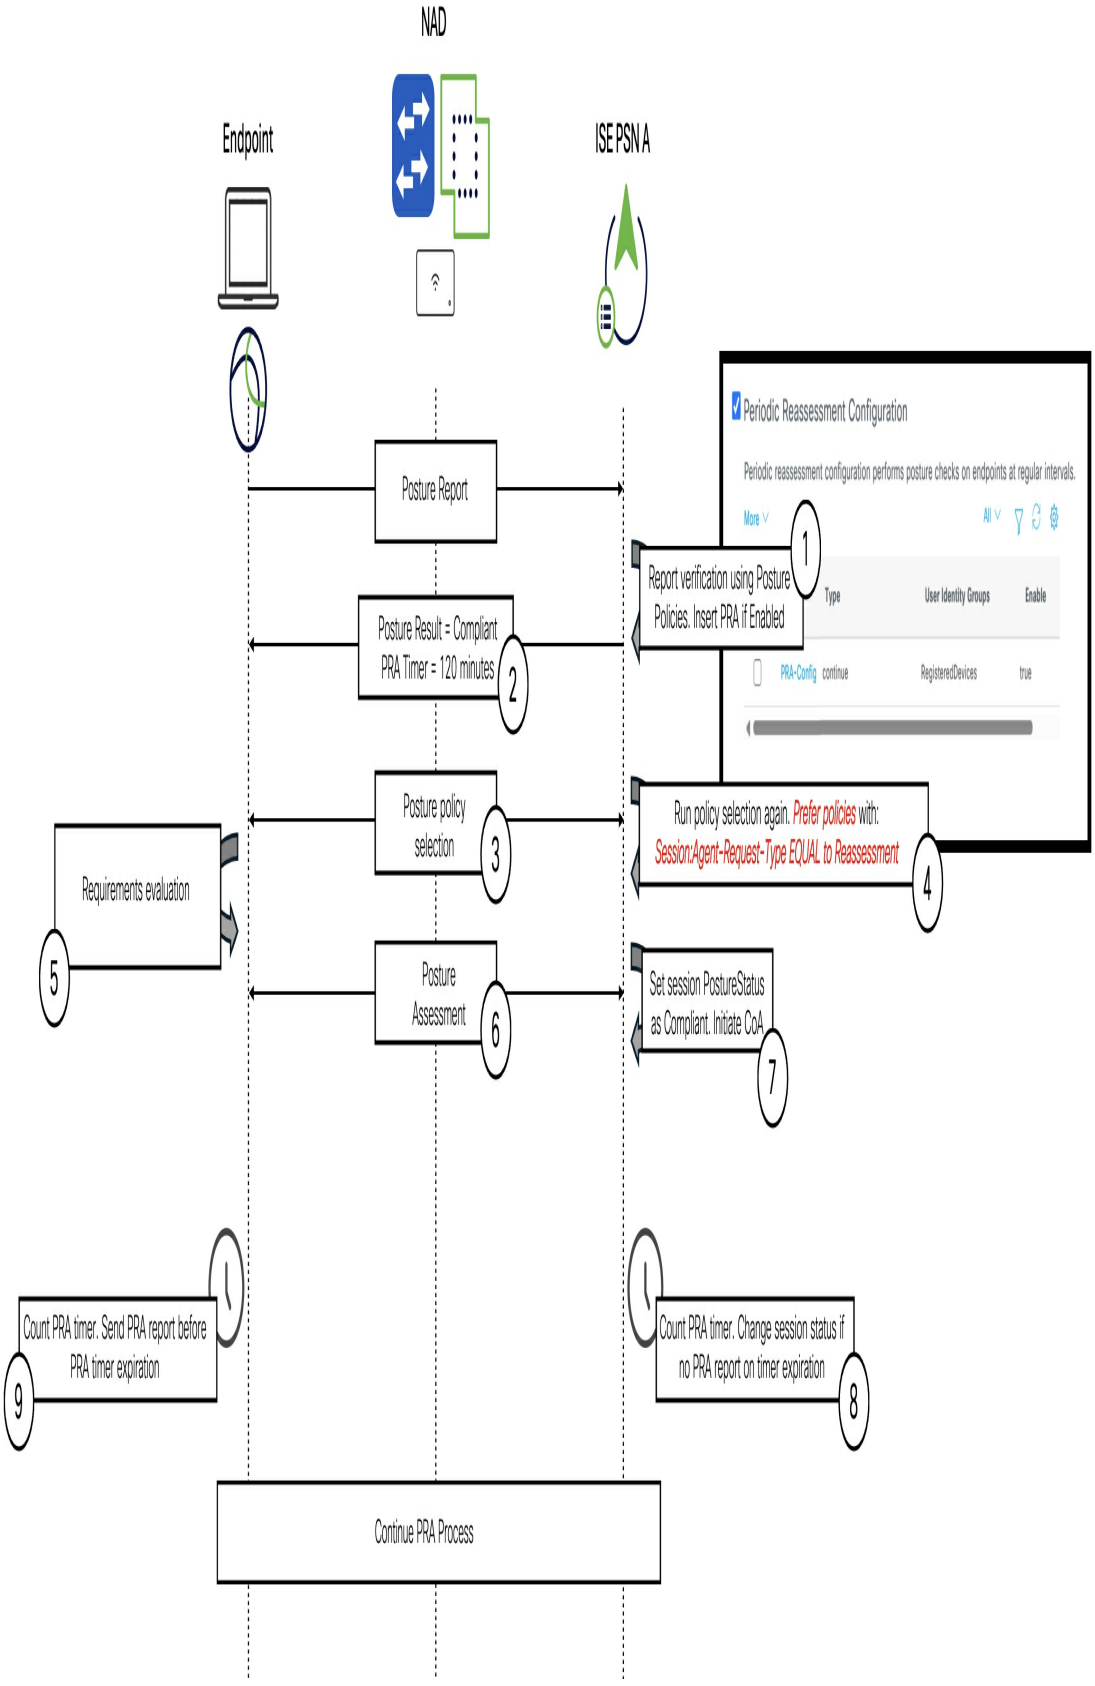

NAD
Endpoint
ISE PSN A
Periodic Reassessment Configuration
Periodic reassessment configuration performs posture checks on endpoints at regular intervals.
More
All
Type
User Identity Groups
Enable
PRA-Config continue
RegisteredDevices
true
Posture Report
Report verification using Posture Policies. Insert PRA if Enabled
1
Posture Result = Compliant
PRA Timer = 120 minutes
2
Posture policy selection
3
Run policy selection again. Prefer policies with:
Session:Agent-Request-Type EQUAL to Reassessment
4
Requirements evaluation
5
Posture Assessment
6
Set session PostureStatus as Compliant. Initiate CoA
7
Count PRA timer. Change session status if no PRA report on timer expiration
8
Count PRA timer. Send PRA report before PRA timer expiration
9
Continue PRA Process

Figure 10-22 *Posture Reassessment Operations*

The following steps are involved in the posture reassessment (PRA) process:

Step 1. The Policy Service Node verifies the agent's posture report using posture policies. If the result is compliant, the PSN checks whether PRA is enabled for this particular endpoint or user.

To configure PRA, navigate to **Work Centers > Posture > Settings > Reassessment Configurations**.

Step 2. When PRA is enabled, the PRA interval is inserted into the posture response that the PSN sends back to the agent.

Step 3. The presence of the PRA timer in the posture response instructs the agent to send a new posture request. This request is marked as "reassessment." The first reassessment request is always sent as part of the initial posture assessment before the PSN issues the Change of Authorization.

Step 4. The PSN performs a Posture Policy lookup again. At this stage, you can use the condition **Session:Agent-Request-Type EQUAL to Reassessment** in your posture policies to narrow down the selection to only those checks required for reassessment. By default, when Agent-Request-Type is not specified in the policies, the same list of policies is selected. If needed, you can specify **Session:Agent-Request-Type EQUAL to Reassessment** and **Session:Agent-Request-Type EQUAL to Initial** in your policies to distinguish between posture policies used for the initial request and for reassessment.

Technet24

Step 5. Posture requirements are evaluated on the agent side.

Step 6. The first PRA posture report is sent to the PSN.

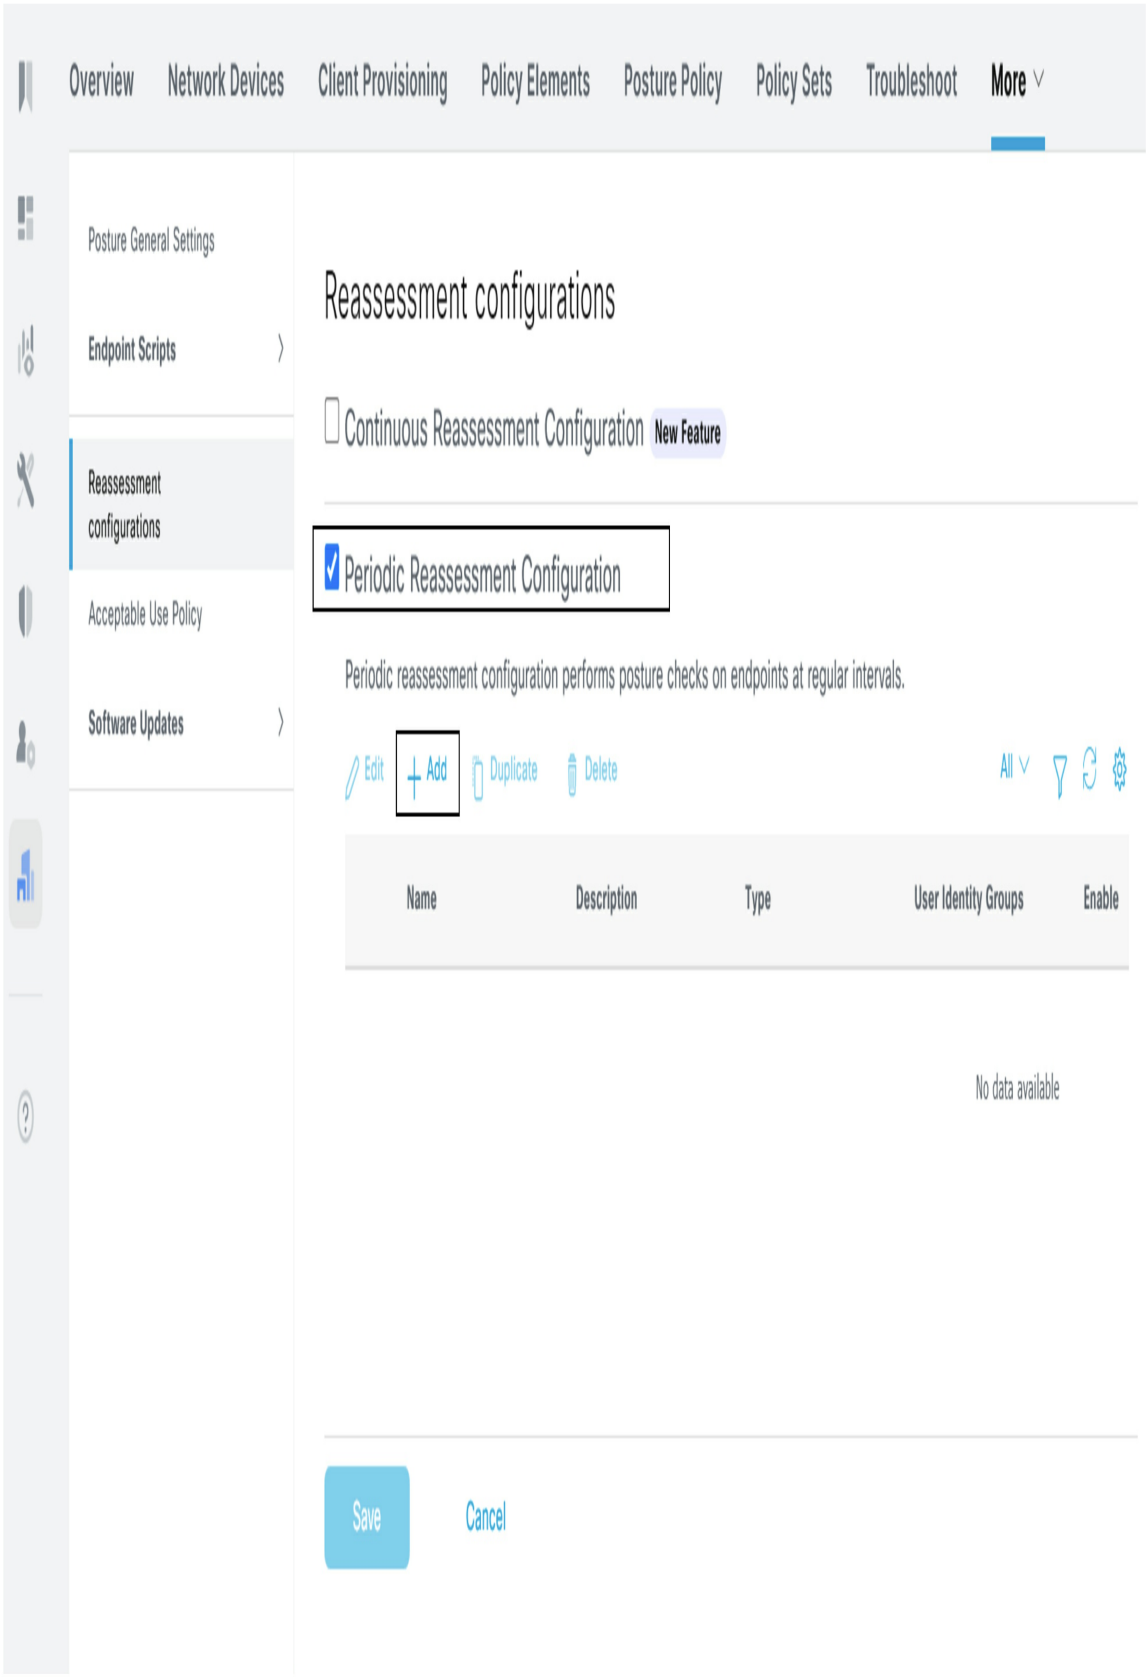
Step 7. The posture process concludes on the ISE side by marking the session as compliant, reevaluating authorization policies, and issuing a Change of Authorization.

Step 8. The ISE PSN starts the PRA timer and assigns it to the session. If the session is terminated with an accounting stop, PRA is removed. If, upon PRA timer expiration, the PSN does not receive a new report from the agent, the session is marked as noncompliant and the Change of Authorization is issued to reauthenticate the endpoint and reduce its network access level.

Step 9. On the agent side, the PRA timer is monitored. The PRA report must be sent before the PRA timer expires. The agent removes the PRA timer if a new posture discovery process starts, for example, due to user logout/login or network disconnect/reconnect.

The PRA configuration is a two-step process: First, you need to enable PRA itself and then add one or more PRA configurations.

To enable PRA, navigate to **Work Centers > Posture > Settings > Reassessment Configurations**, select the checkbox next to **Periodic Reassessment Configuration**, and then click **+ Add** to create a new PRA configuration, as demonstrated in Figure 10-23.


Overview
Network Devices
Client Provisioning
Policy Elements
Posture Policy
Policy Sets
Troubleshoot
More
Posture General Settings
Endpoint Scripts
Reassessment configurations
Acceptable Use Policy
Software Updates
Reassessment configurations
Continuous Reassessment Configuration
New Feature
Periodic Reassessment Configuration
Periodic reassessment configuration performs posture checks on endpoints at regular intervals.
Edit
Add
Duplicate
Delete
All
Name
Description
Type
User Identity Groups
Enable
No data available
Save
Cancel

Technet24

Figure 10-23 *Adding a New PRA Configuration*

In the new window, specify the configuration name and select at least one endpoint or user identity group. The remaining settings are optional. For users, only internal identity groups are supported, so you cannot apply a PRA configuration to Active Directory or other external groups. If PRA needs to be enabled for all endpoints and users, use the identity group Any, as shown in the example. PRA policies should not overlap, meaning the same group cannot be added to two different policies. An example of PRA configuration is demonstrated in Figure 10-24.

Reassessment Configurations List 〉 New Reassessment Configuration

Reassessment Configuration

* Configuration Name	PRA-Any-Endpoint
Configuration Description	
Use Reassessment Enforcement ?	☑
Enforcement Type	continue
Interval	240 minutes.
Grace Time	5 minutes.
Group Selection Rules	1. Each configuration must have a unique group or a unique combination of groups. 2. No two configurations may have any group in common. 3. If a config already exists with a group of 'Any', then no other configs can be created unless - i. the existing config with a group of 'Any' is updated to reflect a group (or groups) other than 'Any', or ii. the existing config with a group of 'Any' is deleted 4. If a config with a group of 'Any' must be created,delete all other configs first.
* Select User Identity Groups	Any

PRA configurations

Configurations list

Existing Reassessment Configurations	User Identity Groups

Submit Cancel

Technet24

Figure 10-24 *PRA Configuration Settings*

Other important PRA settings include the following:

- **Use Reassessment Enforcement checkbox:** This setting enables ISE to react when the endpoint enters a noncompliant state during PRA.
- **Enforcement Type:** This setting instructs ISE on what action to take if the endpoint becomes noncompliant. The options are
  - **Continue:** The endpoint continues to have network access.
  - **Logoff:** The endpoint is considered noncompliant and loses network access.
  - **Remediate:** The appropriate remediation action is taken on the endpoint.
- **Interval:** This setting specifies the posture reassessment interval.
- **Grace Time:** This setting instructs the Policy Service Node how long to wait after the posture reassessment interval expires before issuing the Posture Change of Authorization.

PRA is supported only with Cisco Secure Client and is not available with the temporal agent or agentless posture.

Continuous Reassessment (CRA)

The main difference between CRA and PRA is that CRA enables the agent to use both event-based and polling-based approaches to monitor changes on the endpoint, allowing the Policy Service Node to react to posture status changes in a manner that is as close to real time as possible. Table 10-4 demonstrates which type of monitoring is supported for each type of posture condition.

Table 10-4 *Posture Conditions and CRA Monitoring Approaches*

	Windows	macOS	Linux
Firewall	Polling	Polling	Not Supported
USB	Event based	Not Supported	Not Supported
File	Event based	Event based	Event based
Script	Polling	Polling	Polling
Registry	Event based	Not Supported	Not Supported
Service	Event based	Polling	Not Supported
Disk Encryption	Polling	Polling	Not Supported
Patch Management	Polling	Polling	Polling
Anti-Malware	Polling	Polling	Polling
Application	Polling	Polling	Not Supported
Process	Event based	Polling	Polling
Osquery	Not Supported	Not Supported	Polling

The logic of continuous reassessment is demonstrated in Figure 10-25.

Technet24

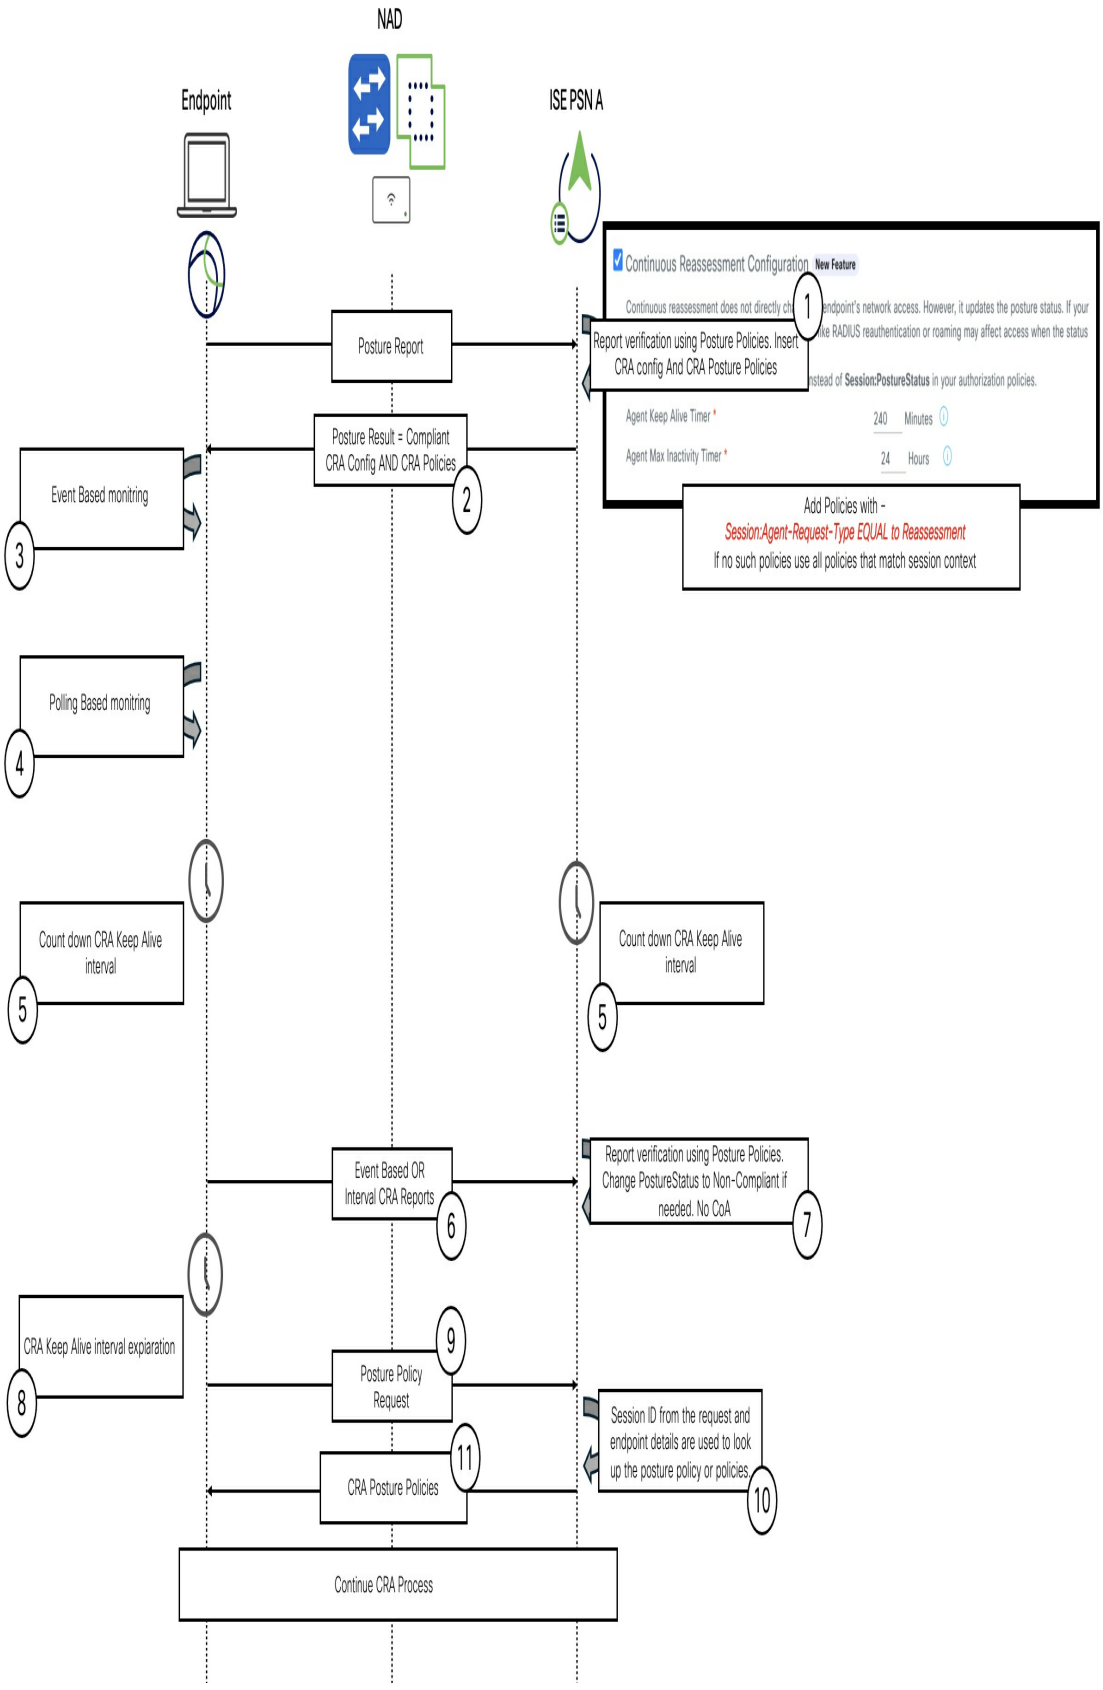

NAD
Endpoint
ISE PSN A
Continuous Reassessment Configuration
New Feature
Continuous reassessment does not directly ch
endpoint's network access. However, it updates the posture status. If your
ke RADIUS reauthentication or roaming may affect access when the status
nstead of Session:PostureStatus in your authorization policies.
Agent Keep Alive Timer *
240
Minutes
Agent Max Inactivity Timer *
24
Hours
1
Posture Report
Report verification using Posture Policies. Insert CRA config And CRA Posture Policies
Posture Result = Compliant CRA Config AND CRA Policies
2
Add Policies with -
Session:Agent-Request-Type EQUAL to Reassessment
If no such policies use all policies that match session context
Event Based monitring
3
Polling Based monitring
4
Count down CRA Keep Alive interval
5
Count down CRA Keep Alive interval
5
Event Based OR Interval CRA Reports
6
Report verification using Posture Policies. Change PostureStatus to Non-Compliant if needed. No CoA
7
CRA Keep Alive interval expiaration
8
9
Posture Policy Request
Session ID from the request and endpoint details are used to look up the posture policy or policies.
10
11
CRA Posture Policies
Continue CRA Process

Figure 10-25 *Continuous Reassessment Steps*

The following steps outline the process involved in the continuous reassessment:

Step 1. Upon receiving a posture report from the agent, the Policy Service Node performs verification to determine the posture status using posture policies. If the posture status is compliant, CRA configuration is verified.

Step 2. If CRA is enabled, the agent receives CRA policies and the CRA keepalive interval in the posture response. CRA policies are posture policies applicable to this specific session/endpoint. You can narrow policy selection, as with PRA, by using the **Session:Agent-Request-Type EQUAL to Reassessment** condition in your posture policies.

Step 3. The agent starts event-based monitoring for the supported posture conditions present in CRA policies.

Step 4. The agent begins polling-based monitoring for the supported conditions present in CRA policies.

Step 5. Both the agent and PSN count the CRA keepalive interval.

Step 6. When changes are detected on the agent side, a report is sent to ISE.

Step 7. ISE reviews the report, and if the endpoint moves out of the compliant state, the session PostureStatus changes to noncompliant. However, CRA-detected changes in posture status do not trigger a CoA, meaning the endpoint will maintain the same access level as before. If reauthentication is triggered by the network access device, the restricted authorization policy for the noncompliant state will be assigned. To avoid this scenario, you can use

Technet24

the condition **PostureInitialStatus** from the session dictionary in your authorization policies instead of **PostureStatus**. **PostureInitialStatus** is not affected by changes discovered during CRA.

Step 8. The keepalive interval expires on the agent side.

Step 9. Upon expiration of the keepalive interval, the agent sends a new posture request to obtain an updated list of posture policies and requirements. This ensures configuration changes in posture policies are picked up by endpoints connected to the network for an extended period.

Step 10. Information from the request is used to look up posture policies.

Step 11. Posture policies and requirements are sent to the agent.

While the session remains active, this process continues, and on each keepalive interval expiration, the list of CRA policies is requested again.

CRA is configurable in the same location as PRA. Navigate to **Work Centers > Posture > Settings > Reassessment Configurations** and select the checkbox next to **Continuous Reassessment Configuration**. An example of CRA settings is demonstrated in Figure 10-26.



Figure 10-26 *CRA Settings*

Grace Period

The grace period allows you to maintain network access for endpoints that were compliant in the past, even when one or more posture policies become noncompliant.

The grace period operates based on the following logic:

Technet24

Step 1. When **Grace Period** is enabled, during the first posture assessment request that results in a compliant outcome, the Policy Service Node assigns an attribute called PostureLastCompliantExpiry to the endpoint. This attribute contains a future timestamp representing when the last known compliant posture status expires.

Step 2. On each subsequent assessment that ends in a compliant result, the timestamp in PostureLastCompliantExpiry is updated.

Step 3. If the endpoint becomes noncompliant, the PSN checks whether the current timestamp is before PostureLastCompliantExpiry. If so, and the posture policy resulting in noncompliance has a grace period assigned, ISE populates LAST_GRACE_EXPIRY for this endpoint and policy.

Step 4. While LAST_GRACE_EXPIRY remains within the grace period, the PSN assigns the compliant status to the endpoint.

Step 5. After the noncompliant condition is remediated, LAST_GRACE_EXPIRY is removed and PostureLastCompliantExpiry is updated with a new timestamp.

The grace period configuration is a two-step process. First, you need to enable the Cache Last Known Posture Compliant Status feature. To do this, navigate to **Work Centers > Posture > Settings > Posture General Settings**, enable the **Cache Last Known Posture Compliant Status** checkbox, and specify the duration of the **Last Known Posture Compliant Status** period. An example of the Last Known Posture Compliant Status configuration is demonstrated in Figure 10-27.

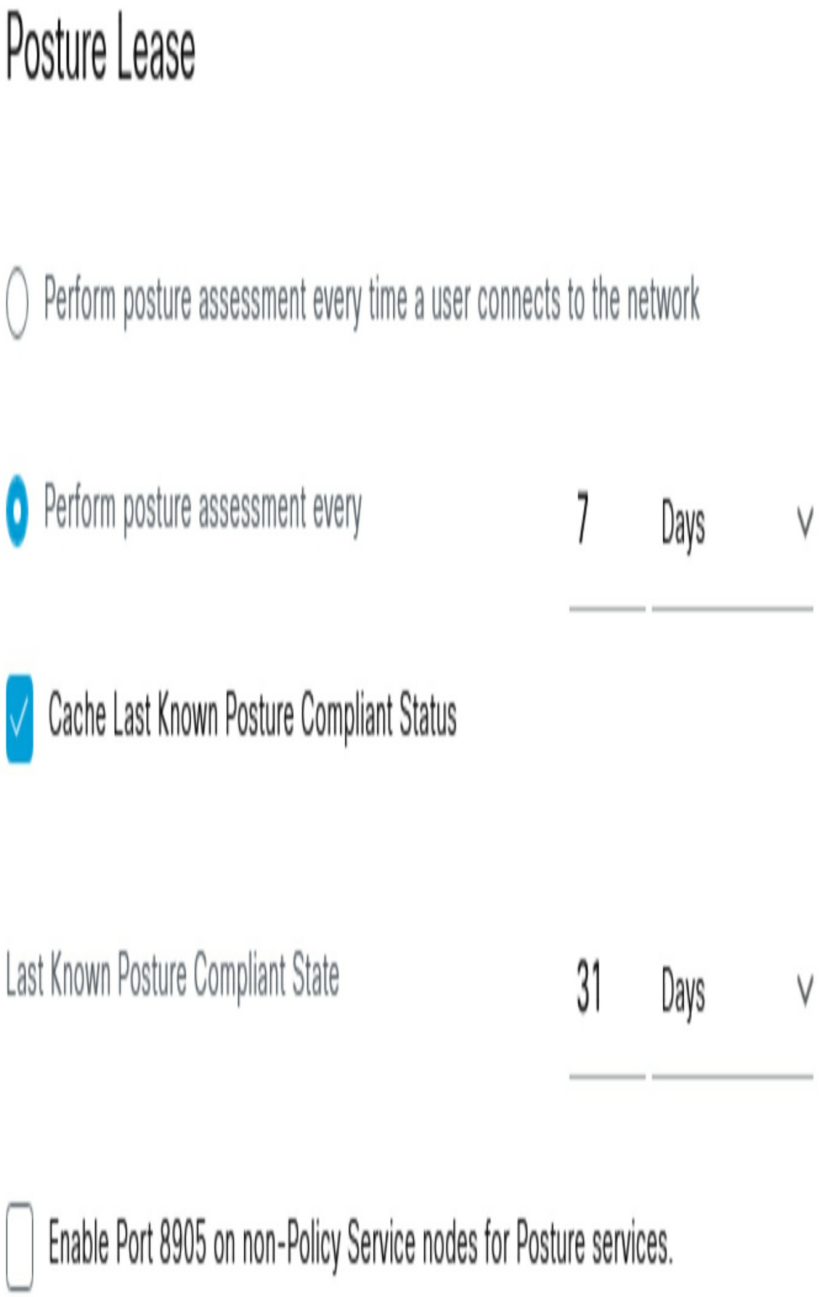


Figure 10-27 *Last Known Posture Compliant Status Configuration*

After the Cache Last Known Posture Compliant Status feature is

Technet24

enabled, you can enable the grace period in individual posture policies. To do this, navigate to **Work Centers > Posture > Posture Policy** and edit the policies where the grace period needs to be enabled. An example of the Posture Policy for Grace Period is demonstrated in Figure 10-28.

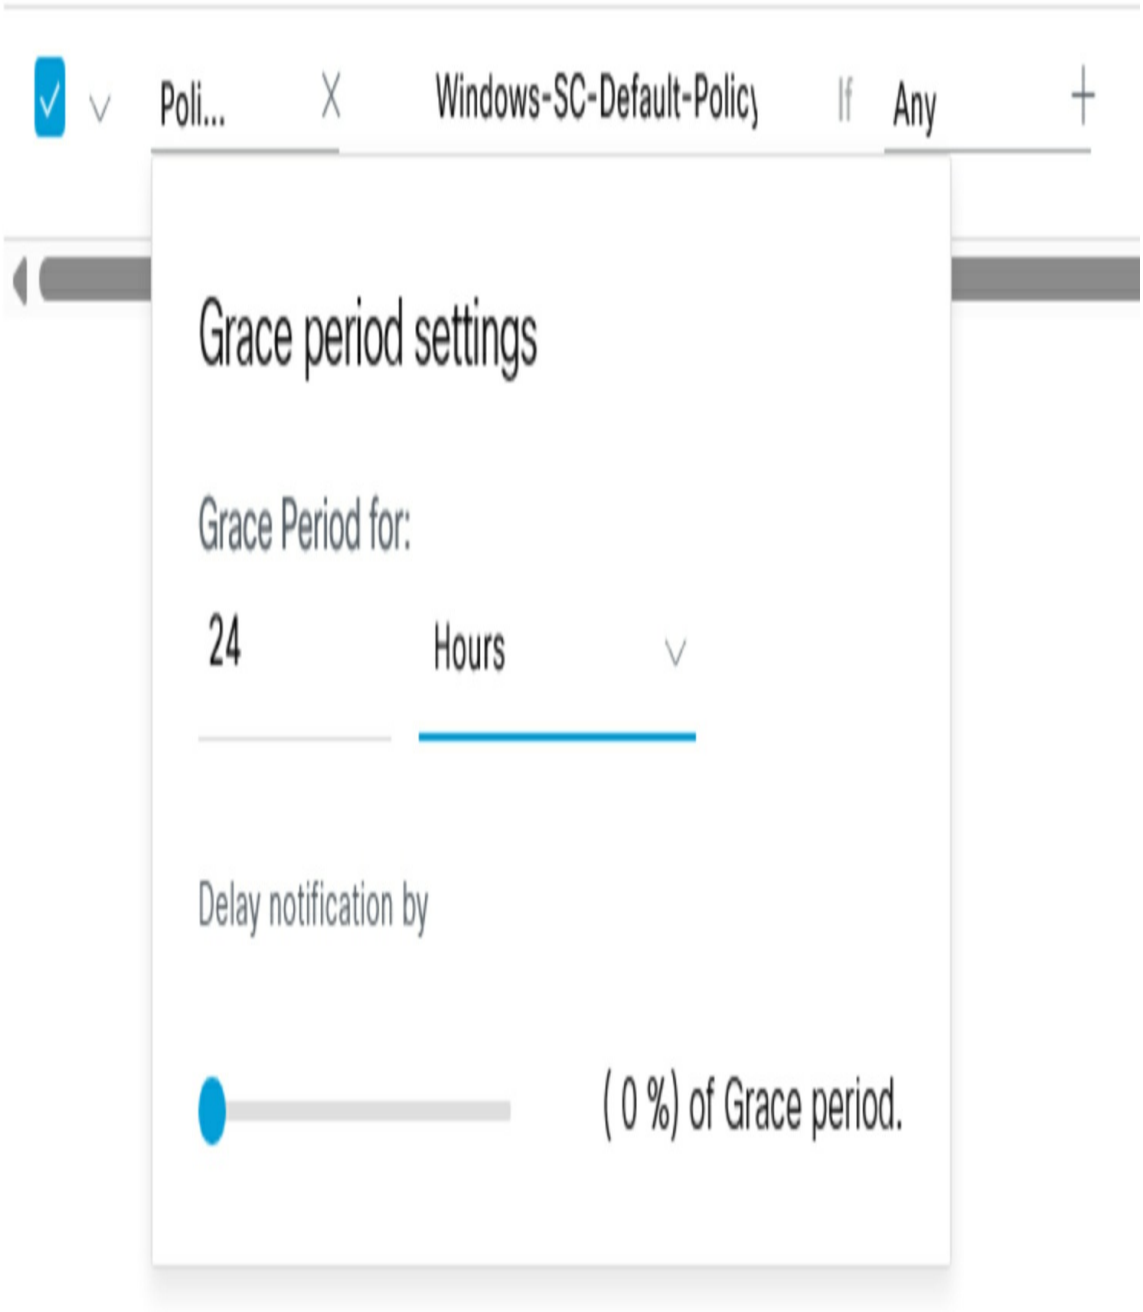


Figure 10-28 *Grace Period Configuration*

Configuring ISE Posture

ISE posture services configuration is a multistep process, with the complexity of each step depending significantly on the specific posture requirements to be implemented. This section covers the most important steps at a high level to provide a general overview of how a complete posture service configuration appears.

All individual configuration steps can be grouped into four main phases:

- **Preparation:** During these steps, ISE and network access devices are prepared for posture service deployment. On the ISE side, required posture updates are installed, and NADs are prepared to handle Central Web Authentication and process Change of Authorization requests from the Policy Service Node.
- **Client Provisioning:** This category depends on the agent type to be deployed. Here, client provisioning resources are configured or uploaded to ISE, and client provisioning policies are created.
- **Posture Configuration:** This group of steps includes all posture settings needed to create actual posture policies. As an administrator, you define all posture requirements and connect them with the corresponding remediation actions so they can be added to posture policies.
- **Authorization Policies Configuration:** In these steps, you create authorization profiles for different posture states, such as a CWA profile for initial authentication, a permit access profile for compliant endpoints, and restricted access for noncompliant endpoints. These profiles are then used to define the required authorization policies.

Technet24

Preparation Phase

On the ISE side, as part of the preparation phase, it is important to ensure that posture conditions and OS fingerprints are updated. To check when the last update occurred, navigate to **Work Centers > Posture > Settings > Software Updates > Posture Updates**. On this page, you can perform updates manually by clicking **Update Now**, or you can configure your own update schedule, as demonstrated in Figure 10-29.



Update Information

Last successful update on	**2026/08/24 20:29:56**
Last update status since ISE was started	**Last update attempt at 2026/08/24 20:29:56 was successful**
Cisco conditions version	**310440.0.0.0**
Cisco AV/AS support chart version for windows	**316.0.0.0**
Cisco AV/AS support chart version for Mac OSX	**233.0.0.0**
Cisco AV/AS support chart version for Linux	**86.0.0.0**
Cisco supported OS version	**112.6.3.0**

Technet24

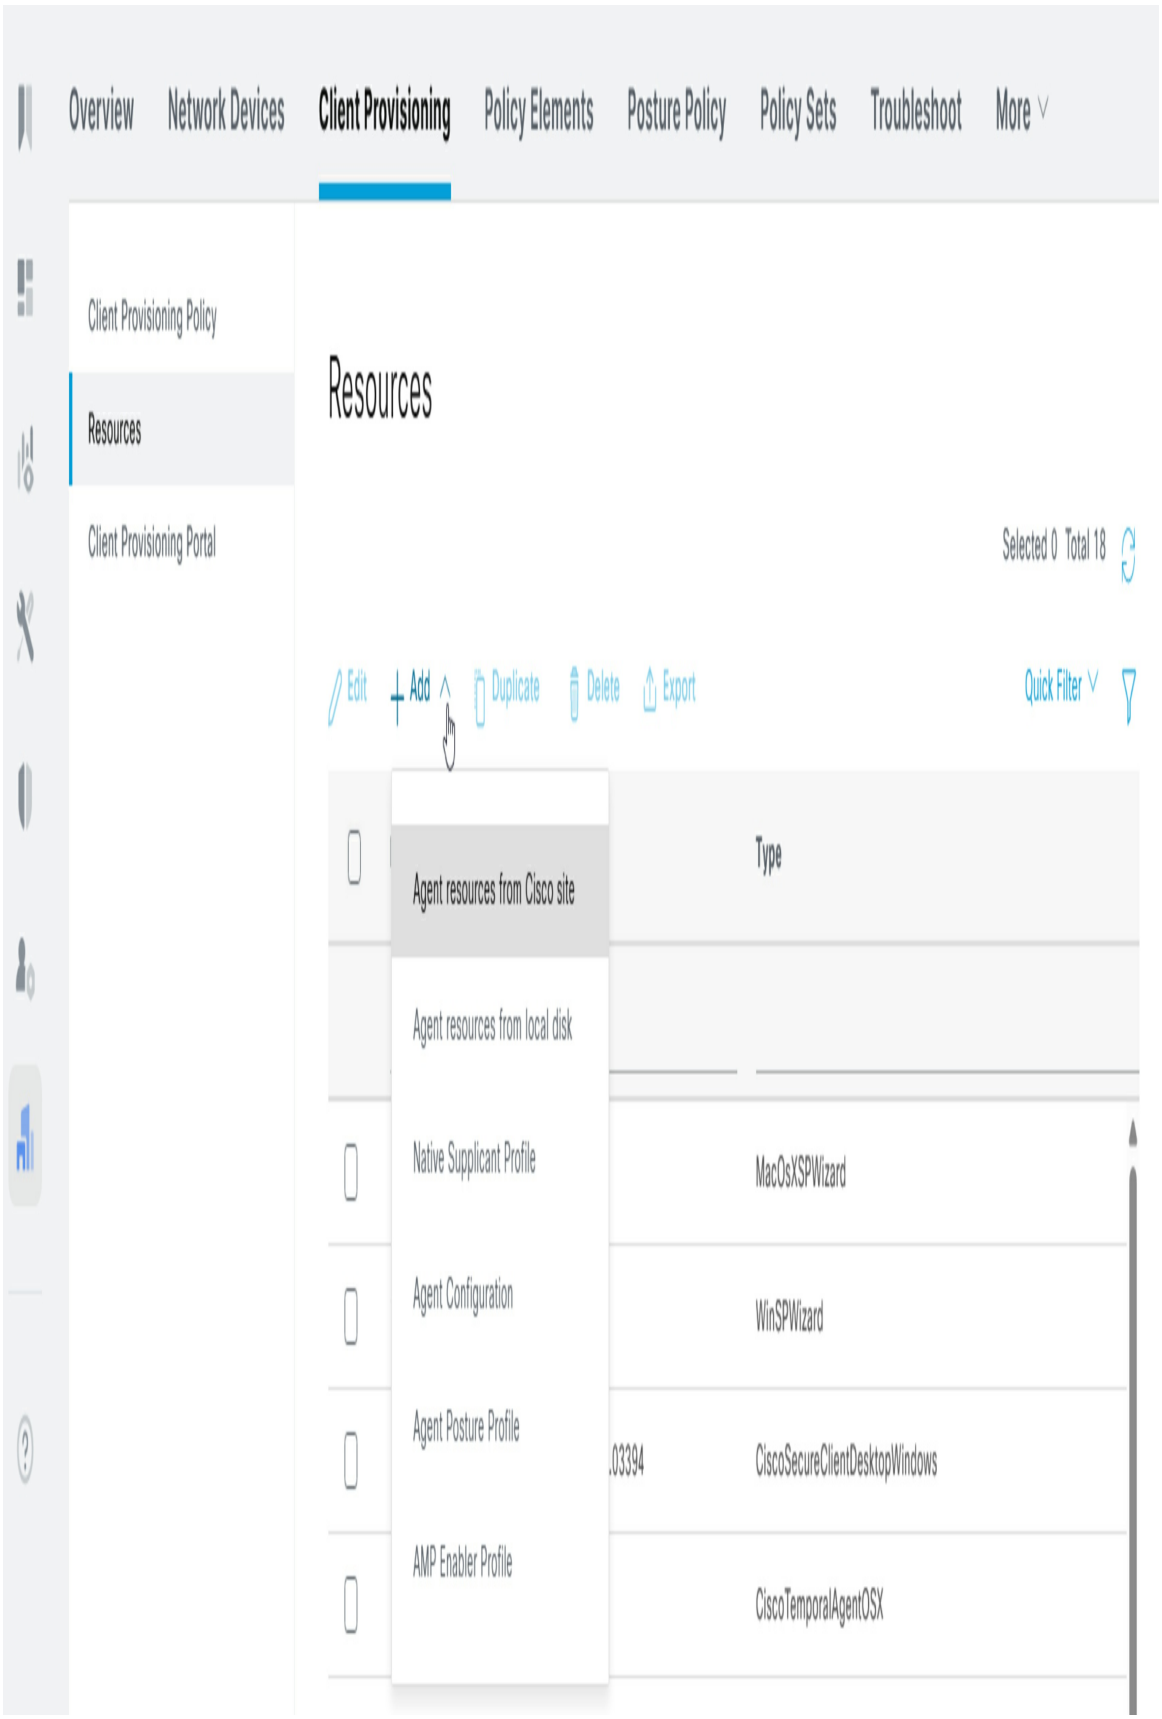
Figure 10-29 *Posture Update Settings*

On the network access device side, the steps required for posture services to operate depend on the NAD type. However, from a high-level perspective, the following configuration needs to be in place for redirect-based posture. This excludes any general settings required to enable Dot1x or MAB:

- The web server must be enabled.
- The web redirect ACL should be defined on the NAD.
- ISE Policy Service Nodes need to be defined as CoA clients.

Client Provisioning Phase

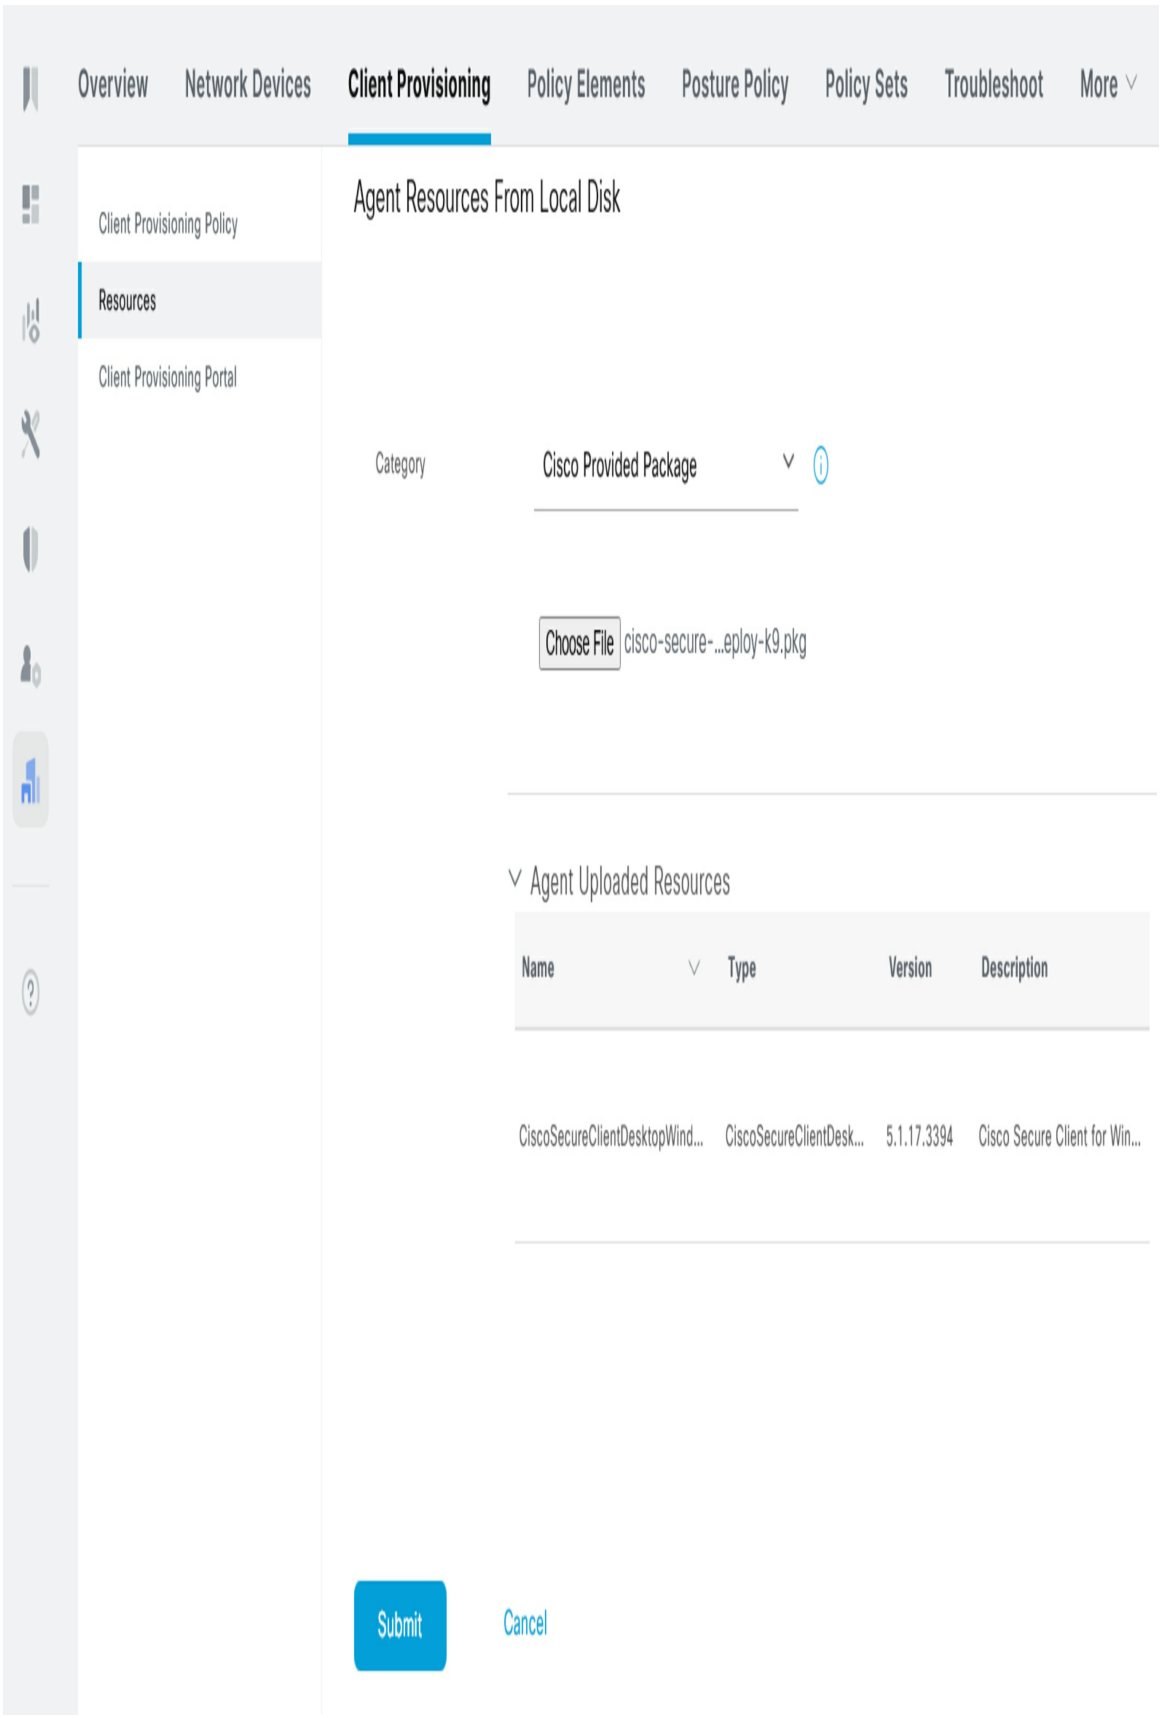
If the temporal agent or agentless module needs to be used, you can download them to ISE directly from cisco.com. To do this, navigate to **Work Centers > Posture > Client Provisioning > Resources**, click **Add**, and choose **Agent Resources from the Cisco Site**, as demonstrated in Figure 10-30. After the list of resources loads, select the latest version of the temporal agent and agentless module and download them to ISE.


Overview
Network Devices
Client Provisioning
Policy Elements
Posture Policy
Policy Sets
Troubleshoot
More
Client Provisioning Policy
Resources
Client Provisioning Portal
Resources
Selected 0 Total 18
Edit
Add
Duplicate
Delete
Export
Quick Filter
Agent resources from Cisco site
Agent resources from local disk
Native Supplicant Profile
Agent Configuration
Agent Posture Profile
AMP Enabler Profile
Type
MacOsXSPWizard
WinSPWizard
CiscoSecureClientDesktopWindows
CiscoTemporalAgentOSX

Technet24

Figure 10-30 *Adding Agent Resources from the Cisco Site*

For Secure Client, the configuration is more complex. First, you need to download the latest version of the Secure Client package from cisco.com. Then select the package for your OS type that contains **Headend Deployment Package** in the name. When the package is available on the same computer from which the ISE UI is accessible, navigate to **Work Centers > Posture > Client Provisioning > Resources**. Click **Add** and choose **Agent Resources from Local Disk**. On the new page, select **Cisco Provided Package** in the **Category** and then upload the Secure Client PKG file. An example is demonstrated in Figure 10-31.


Overview
Network Devices
Client Provisioning
Policy Elements
Posture Policy
Policy Sets
Troubleshoot
More
Client Provisioning Policy
Resources
Client Provisioning Portal
Agent Resources From Local Disk
Category
Cisco Provided Package
Choose File
cisco-secure-...eploy-k9.pkg
Agent Uploaded Resources
Name
Type
Version
Description
CiscoSecureClientDesktopWind...
CiscoSecureClientDesk...
5.1.17.3394
Cisco Secure Client for Win...
Submit
Cancel

Technet24

Figure 10-31 *Adding a Secure Client Package to ISE*

The next step required to create Secure Client configuration is to add the compliance module. You can download the compliance module file in the same way as the temporal agent or agentless module. After the list of resources is loaded, select the latest compliance module file for the appropriate OS type.

The final step in preparing to create the Secure Client configuration is to define the ISE posture profile. On the **Client Provisioning Resources** page, click **Add** and choose **Agent Posture Profile.** The profile configuration example is demonstrated in Figure 10-32. Only two mandatory settings exist for each profile: First, specify the profile name; second, define the value in the **Server Name Rules** section.

Name *

SC-5-1-13-Posture-Profile

Description:

Agent Behavior

Parameter	Value	Description
Enable debug log	No	Enables the debug log on the agent
Preferred reassessment type	Continuo	Preferred reassessment configuration will be applied when both reassessment configurations are enabled.
Operate on non-802.1X wireless	No	Enables the agent to operate on non-802.1X wireless networks.

Posture Protocol

Parameter	Value	Description
Reassessment retransmission time	120 secs	This is the agent retry period if there is a reassessment communication failure
Retransmission Delay	60 secs	Time (in seconds) to wait before retrying.
Retransmission Limit	4	Number of retries allowed for a message.
Discovery host		Enter any IP address or FQDN that is routed through a NAD. The NAD detects and redirects that HTTP traffic to the Client Provisioning portal.
Discovery Backup Server List	Choose	By default, AnyConnect sends discovery probes to all the Cisco ISE PSNs sequentially if the PSN is unreachable. Choose specific PSNs as the backup list and restrict the nodes to which AnyConnect sends discovery probes.
Server name rules *	*	A list of wildcarded, comma-separated names that defines the servers that the agent can connect to. E.g. "*.cisco.com"

Technet24

Figure 10-32 *Agent Posture Profile*

The Server Name Rules determine which server names are allowed to communicate with the agent. The agent compares the value from the **Server Name Rules** section with the CN or SAN in the PSN certificate. If any name should be allowed, use *****. Alternatively, use a wildcard that covers your domain — for example, *.iselabnet.net. It is important to note that * covers only one level in the domain hierarchy. For example, *.iselabnet.net will work for ise1.iselabnet.net and ise2.iselabnet.net but will not work for ise1.us.iselabnet.net and ise2.us.iselabnet.net.

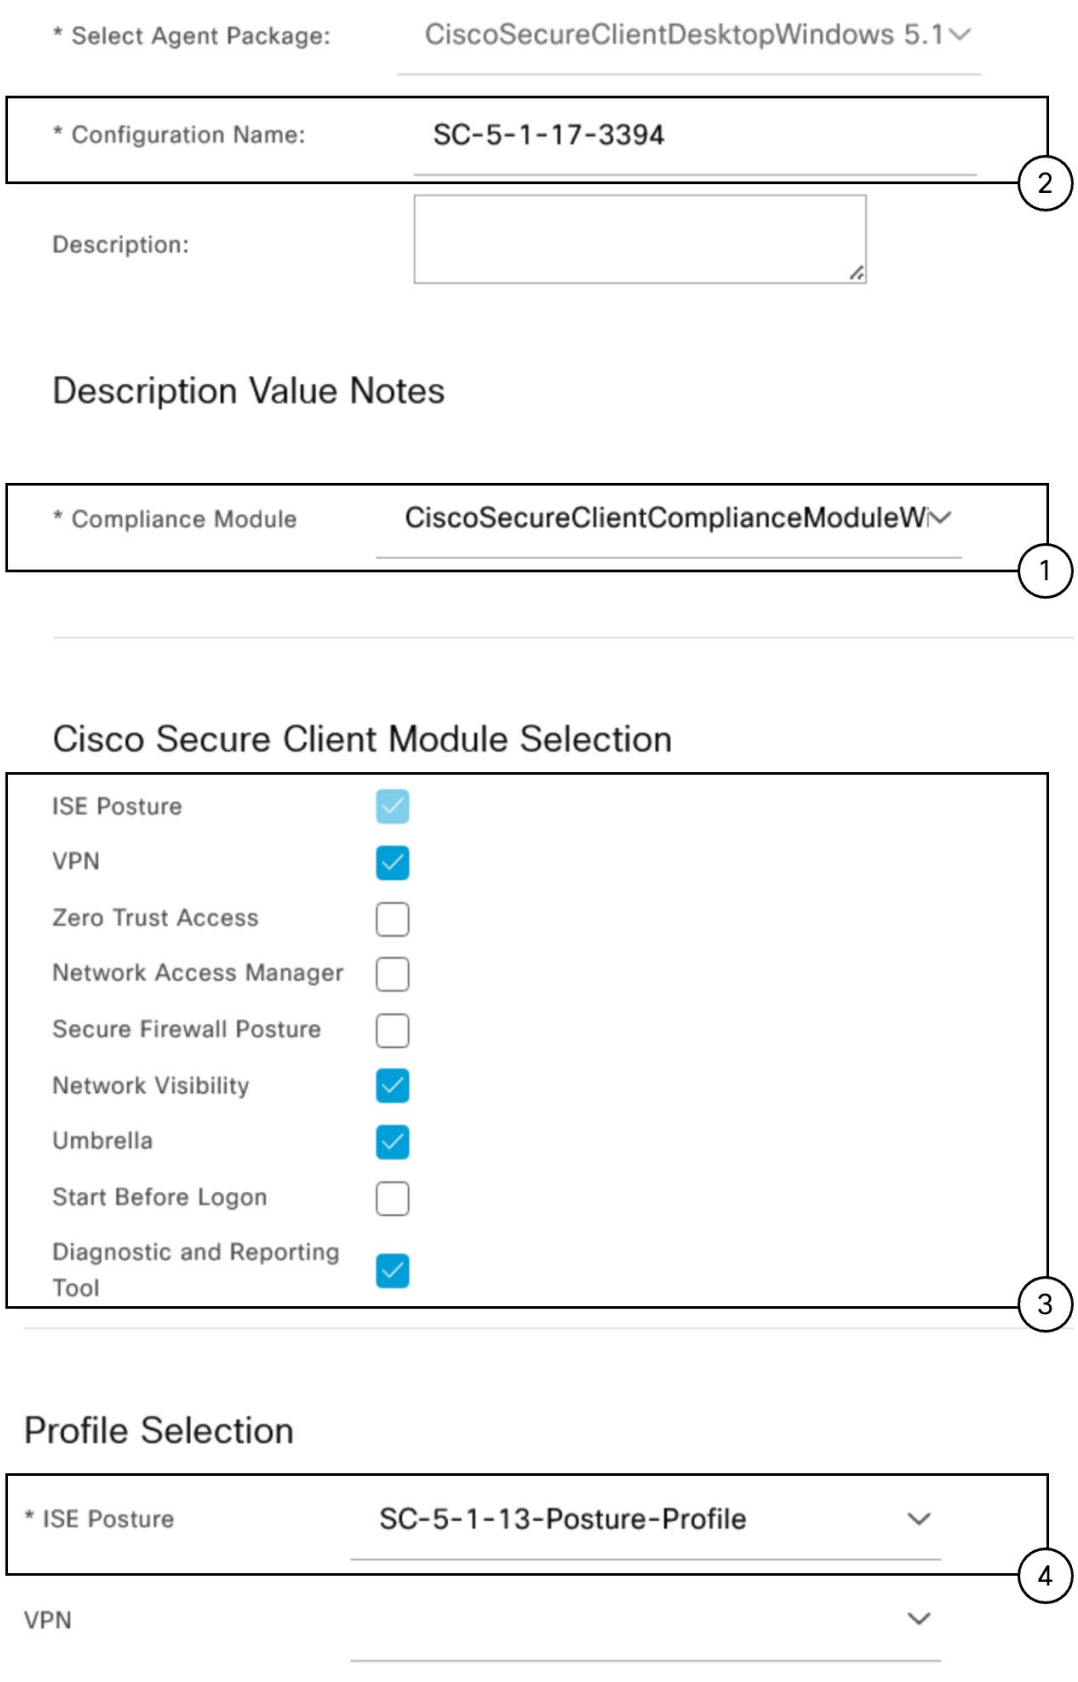
When all elements are ready, combine them in the agent configuration, which is later used in client provisioning policies. To create the agent configuration, go to the **Client Provisioning Resources** page, click **Add**, and choose **Agent Configuration**. On the new page, select your Secure Client package file from the list named **Select Agent Package**. After the package is selected, the full configuration page will be displayed. On this page, as demonstrated in Figure 10-33, do the following:

Step 1. Select your compliance module version.

Step 2. Specify the configuration name.

Step 3. Choose the modules to provision in the **Cisco Secure Client Module Selection** section.

Step 4. Select the ISE posture profile in the **Profile Selection** section.


* Select Agent Package:
CiscoSecureClientDesktopWindows 5.1
* Configuration Name:
SC-5-1-17-3394
2
Description:
Description Value Notes
* Compliance Module
CiscoSecureClientComplianceModuleW
1
Cisco Secure Client Module Selection
ISE Posture
VPN
Zero Trust Access
Network Access Manager
Secure Firewall Posture
Network Visibility
Umbrella
Start Before Logon
Diagnostic and Reporting Tool
3
Profile Selection
* ISE Posture
SC-5-1-13-Posture-Profile
4
VPN

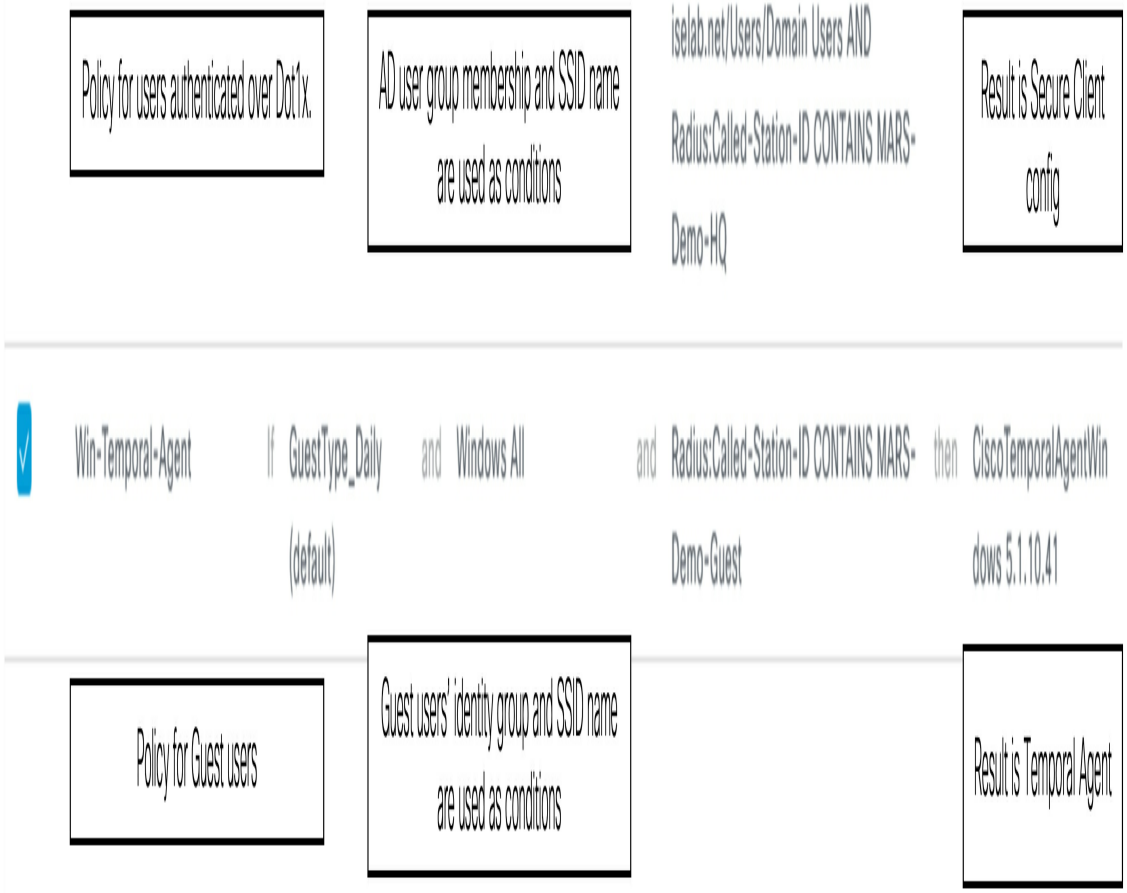
Technet24

Figure 10-33 *Agent Configuration*

The rest of the settings on this page are optional.

The final step of the client provisioning phase is to configure the client provisioning policies. In the **Other Conditions** section, use parameters specific to the given user or endpoint type to ensure the correct client will be provisioned. An example of client provisioning policies is shown in Figure 10-34.

Rule Name	Identity Groups	Operating Systems	Other Conditions	Results
Win-SC-dot1x	If Any	and Windows All	and ISELAB:ExternalGroups EQUALS iselab.net/Users/Domain Users AND Radius:Called-Station-ID CONTAINS MARS-Demo-HQ	then SG-S-1-17-3394
Win-Temporal-Agent	If GuestType_Daily (default)	and Windows All	and Radius:Called-Station-ID CONTAINS MARS-Demo-Guest	then CiscoTemporalAgentWindows 5.1.10.41



Technet24

Figure 10-34 *Client Provisioning Policies Example*

Posture Configuration Phase

When you are configuring posture policies, many common use cases — such as verifying if anti-malware software is up to date, checking if a USB device is attached, or ensuring the patch management client is up to date — are already covered by preconfigured posture conditions, remediations, and corresponding policies.

However, if you need to configure something from scratch, it is important to understand the order of operations:

- **Posture conditions:** When a custom posture check is required, begin by defining a posture condition of the specific type. Posture conditions are grouped by their respective types, such as Anti-Malware, USB, and Patch Management. To view existing conditions or create a new one, navigate to **Work Centers > Posture > Policy Elements > Conditions.**
- **Posture remediations:** Each posture condition requires a corresponding remediation action. The remediation action dictates what the agent should do if the posture condition is not met. Remediations are also grouped by specific categories. Some posture condition categories have direct remediation categories; for example, the Anti-Malware posture condition category has an Anti-Malware remediation category, where you can create an automatic remediation to force the endpoint's anti-malware software to update itself. For other posture condition categories without direct remediation categories, use more generic remediation categories, such as Script remediation, which can be paired with the Osquery condition category to execute a specific action based on the Osquery result. To configure remediation

actions, navigate to **Work Centers > Posture > Policy Elements > Remediations**. The remediation action could be as simple as displaying text on the screen or providing a link that the user needs to open to proceed with the next remediation steps.

- **Posture requirements:** Before configuring the posture policy, connect conditions with their corresponding remediations through posture requirements. You can configure posture requirements by navigating to **Work Centers > Posture > Policy Elements > Requirements**.

Note

Posture conditions and requirements require you to select the compliance module version. When a condition is created for the compliance module branch 4.x, it can only be selected in requirements created for the same branch. If you attempt to select a 4.x condition in a 3.x requirement, your condition will not be displayed.

After the preconfigurations are completed, posture policies can be defined. To do this, navigate to **Work Centers > Posture > Posture Policy**. On the **Posture Policy** page, you will not see an Add button, so to add a new policy, click the down arrow next to the **Edit** button in one of the existing policies and choose **Insert a New Policy**, as shown in Figure 10-35.

Technet24

Posture Policy

Guide Me

Define the Posture Policy by configuring rules based on operating system and/or other conditions.

and	Mac OSX	and	4.x or later	and	Temporal Agent	and	then	Default_Hardware_Attributes_Requirement_Mac_temporal	Edit
and	Windows All	and	4.x or later	and	Agent	and	then	Default_Hardware_Attributes_Requirement_Win	Edit
and	Windows All	and	4.x or later	and	Temporal Agent	and	then	Default_Hardware_Attributes_Requirement_Win_temporal	Edit
and	Windows All	and	4.x or later	and	Agent	and	then	USB_Block	Edit
and	Windows All	and	4.x or later	and	Temporal Agent	and	then	USB_Block_temporal	Edit
and	Windows All	and	4.x or later	and	Agent	and	then	WinDefender-Req	Edit

Duplicate

Insert new policy

Delete

Reset

Figure 10-35 *Adding a New Posture Policy*

For each policy, you can define the following parameters:

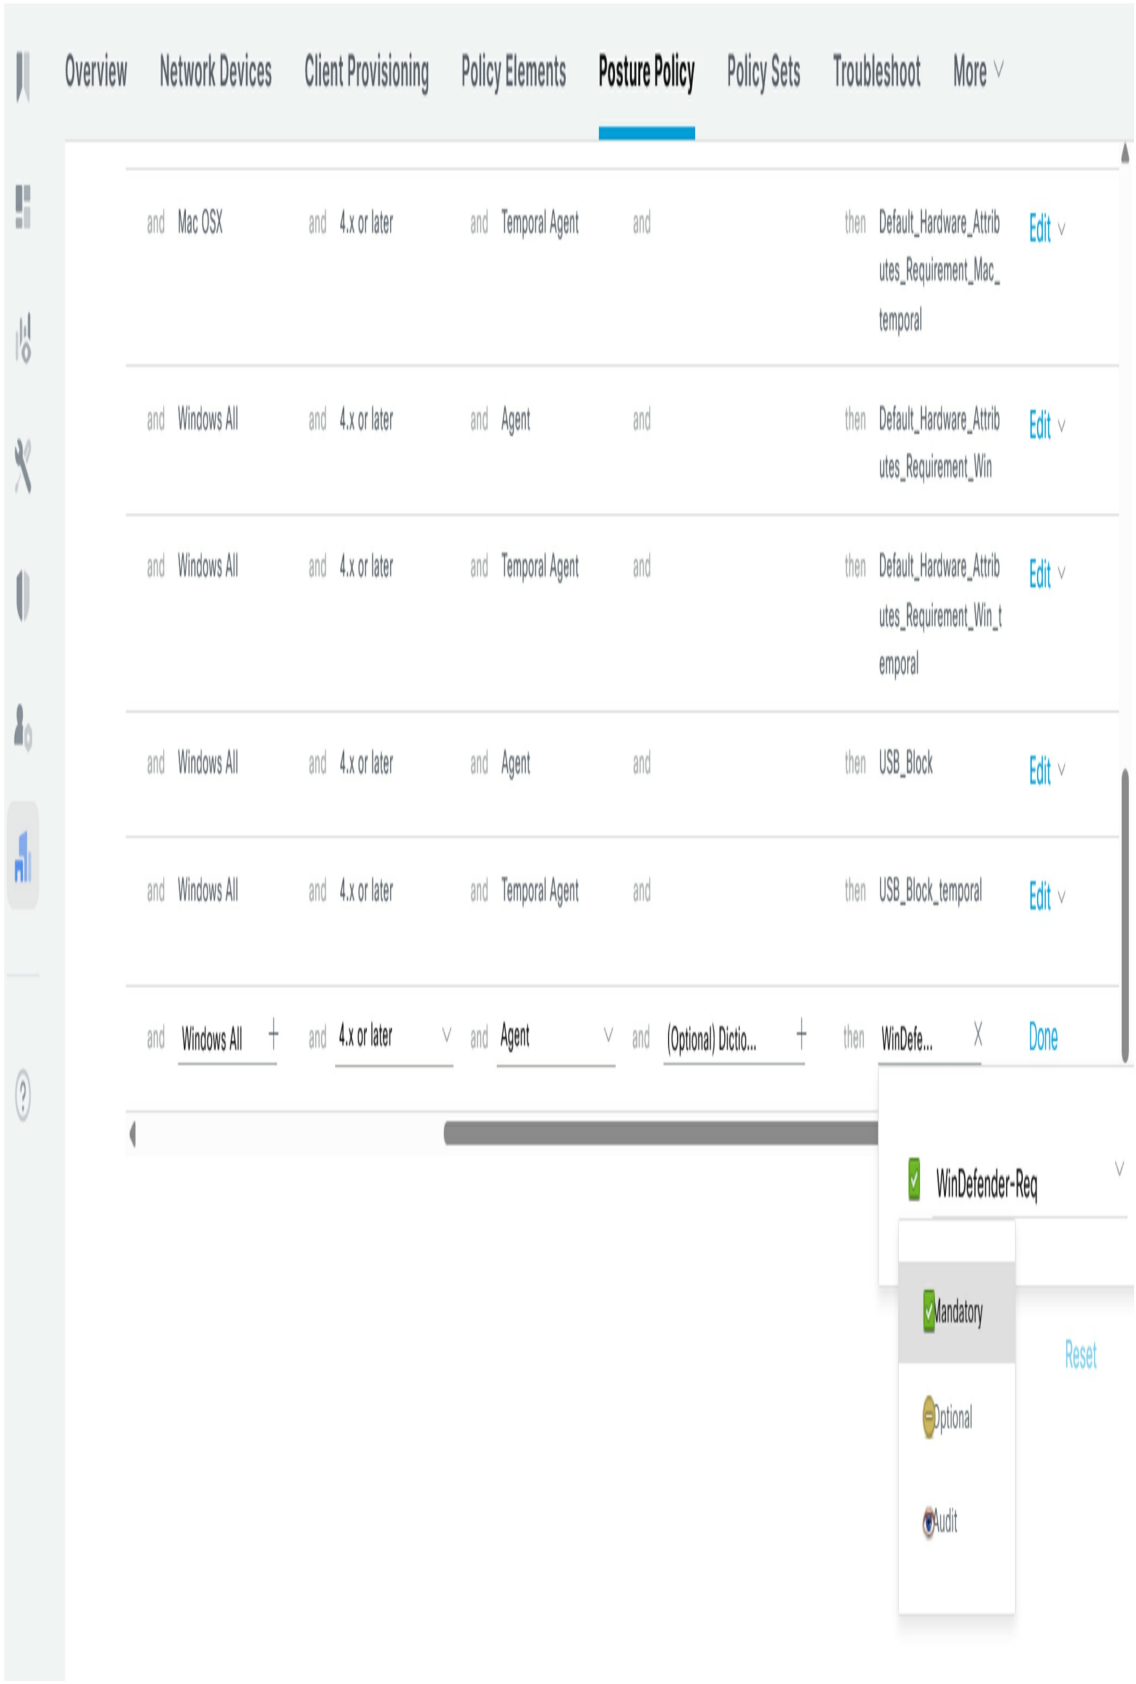
- **Status:** This parameter determines whether the policy is enabled or disabled. Only enabled policies can be used. By default, all preconfigured policies are disabled.
- **Policy Options:** This parameter allows you to specify Grace Period settings, which are disabled by default.
- **Rule Name:** Specify the name of your posture policy.
- **Identity Groups:** Select internal endpoint or user identity groups to which this policy should apply. By default, **Any** is selected in all policies.
- **Operating Systems:** Select the OS type for which this policy can be applied. OS type functions as a filter. Later, in the **Requirements** column, you can only select requirements created for this OS type.
- **Compliance Module:** Select the compliance module branch for which this policy can be applied. Compliance Module acts as a filter. Later, in the **Requirements** column, you can only select requirements created for this Compliance Module branch.
- **Posture Type:** Choose the type of agent to which the policy should apply.
- **Other Conditions:** Select items such as user Active Directory group, SSID name, and posture assessment type.
- **Requirements:** Add one or more requirements to the policy. For each requirement, specify whether it is Mandatory, Optional, or Audit.
    - **Mandatory:** The requirement must be passed for the session to

Technet24

become compliant.

- **Optional:** The user is notified about the failure and has the option to bypass this condition.
- **Audit:** The result is logged, but the user is not notified.

An example of selecting the enforcement type for a specific requirement is demonstrated in Figure 10-36.


Overview
Network Devices
Client Provisioning
Policy Elements
Posture Policy
Policy Sets
Troubleshoot
More
and Mac OSX and 4.x or later and Temporal Agent and then Default_Hardware_Attributes_Requirement_Mac_temporal Edit
and Windows All and 4.x or later and Agent and then Default_Hardware_Attributes_Requirement_Win Edit
and Windows All and 4.x or later and Temporal Agent and then Default_Hardware_Attributes_Requirement_Win_temporal Edit
and Windows All and 4.x or later and Agent and then USB_Block Edit
and Windows All and 4.x or later and Temporal Agent and then USB_Block_temporal Edit
and Windows All and 4.x or later and Agent and (Optional) Dictio... then WinDefe... Done
WinDefender-Req
Mandatory
Optional
Audit
Reset

Technet24

Figure 10-36 *Enforcement Type for Posture Requirement*

Authorization Policies Configuration Phase

To configure authorization policies, you must first create an authorization profile. At minimum, two authorization profiles are required: one with CWA redirection to the client provisioning or guest portal, depending on how the posture flow will be performed, and one for full access, which is assigned after successful posture assessment. The CWA profile can also be used to cover a noncompliant scenario. Examples of authorization profiles for CWA are demonstrated in Figure 10-37. An authorization profile with CWA is required only for the redirect-based posture.

Initial Authentication authorization profile with Client Provisioning Posture redirect



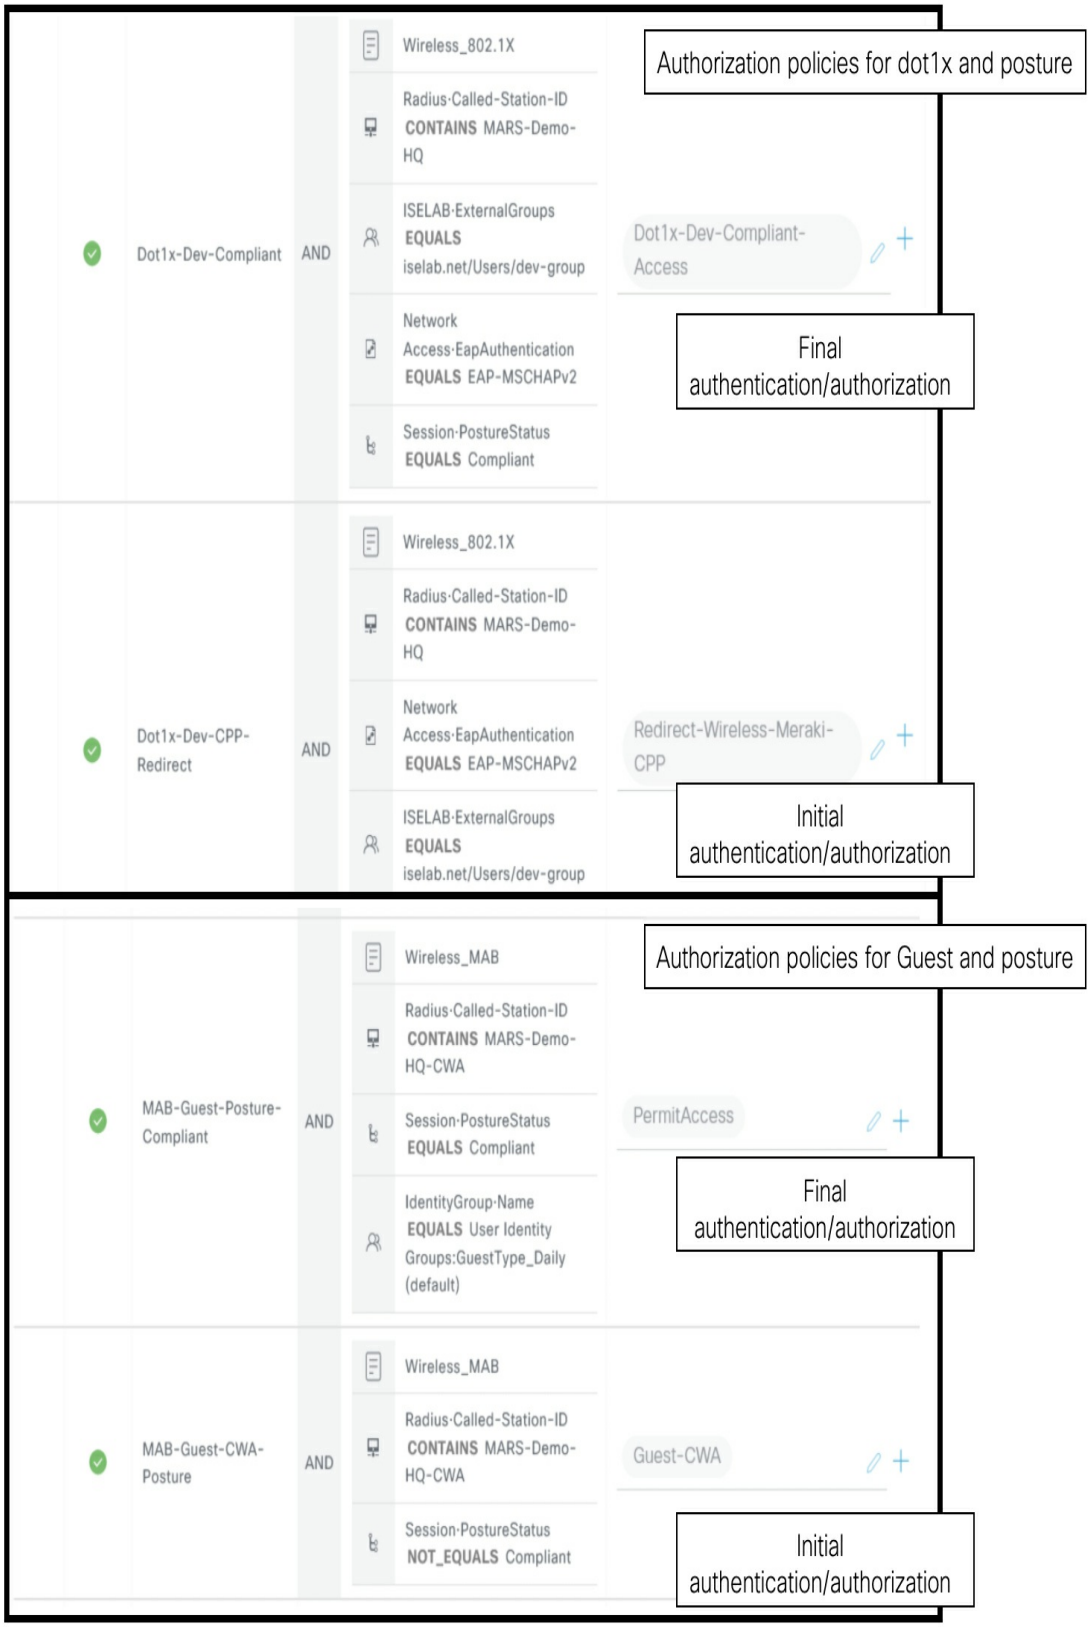
Initial Authentication authorization profile with Guest portal redirect



Technet24

Figure 10-37 *Examples of Authorization Profiles*

Examples of the authorization policies that can be used for posture are shown in Figure 10-38.


Authorization policies for dot1x and posture
Dot1x-Dev-Compliant
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
ISELAB·ExternalGroups EQUALS iselab.net/Users/dev-group
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
Session·PostureStatus EQUALS Compliant
Dot1x-Dev-Compliant-Access
Final authentication/authorization
Dot1x-Dev-CPP-Redirect
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
ISELAB·ExternalGroups EQUALS iselab.net/Users/dev-group
Redirect-Wireless-Meraki-CPP
Initial authentication/authorization
Authorization policies for Guest and posture
MAB-Guest-Posture-Compliant
AND
Wireless_MAB
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ-CWA
Session·PostureStatus EQUALS Compliant
IdentityGroup·Name EQUALS User Identity Groups:GuestType_Daily (default)
PermitAccess
Final authentication/authorization
MAB-Guest-CWA-Posture
AND
Wireless_MAB
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ-CWA
Session·PostureStatus NOT_EQUALS Compliant
Guest-CWA
Initial authentication/authorization

Technet24

Figure 10-38 *Examples of Authorization Policies*

MDM Flows

Cisco ISE supports integration with a wide variety of mobile device management (MDM) vendors. It is recommended to check the current list of supported vendors online. Some MDM solutions have evolved into unified endpoint management (UEM), which means they support more than just mobile devices. Microsoft Intune and JAMF are good examples, because they also support desktop operating systems.

As explained earlier, the results of MDM and posture flows are similar, but the methods used to achieve these results differ significantly.

MDM and UEM integrations are categorized as follows:

- **Mobile device managers:** MDM/UEM solutions such as Microsoft Intune, JAMF, AirWatch, and others. This group can be further divided by
  - **Authentication type:** Some MDM/UEM solutions use basic authentication with a username and password, which may represent an actual user or API credentials on the MDM side. Others require advanced authentication methods such as OAuth. The primary example here is Microsoft Intune.
  - **Supported API version:** ISE currently supports MDM API versions 2 and 3. The main difference between these APIs is how ISE and MDM identify the endpoint. With version 2, ISE uses the RADIUS Calling-Station-ID value from the access-request as an endpoint identifier. For wired and wireless access, the RADIUS Calling-Station-ID contains the endpoint MAC address, which is then used to locate the endpoint on the MDM server by adding it as an endpoint ID in the API request. With API version 3,

selected MDM servers use a unique endpoint ID (UID), which resolves issues such as MAC randomization and WiFi/Ethernet USB dongles shared between multiple endpoints. This API version 3 approach is recommended. Importantly, ISE does not have access to the endpoint ID generated by the MDM server. For this flow to work, the endpoint ID is included as one of the parameters in the endpoint identity certificate, which must be used during Dot1x authentication. This allows ISE to obtain the endpoint ID value and query the MDM server using this value.

- **Microsoft Configuration Manager:** This is a distinct type of MDM/UEM solution consisting solely of Microsoft System Center Configuration Manager (SCCM). SCCM integration does not rely on ISE MDM API mechanisms. Instead, it uses a WMI/MS-RPC connection between ISE Policy Service Nodes and the SCCM server to obtain endpoint details, specifically the statuses of individual configuration baselines (CBs) assigned to the endpoint. The compliance decision with SCCM works differently than with classic MDM. Classic MDM servers provide ISE with posture status over an API, which can be compliant or noncompliant. With SCCM, ISE receives the compliance statuses of individual CBs, and the endpoint is considered compliant by the Policy Service Node only if all CBs have a compliant status.

The MDM APIv2 flow is shown in Figure 10-39.

Technet24

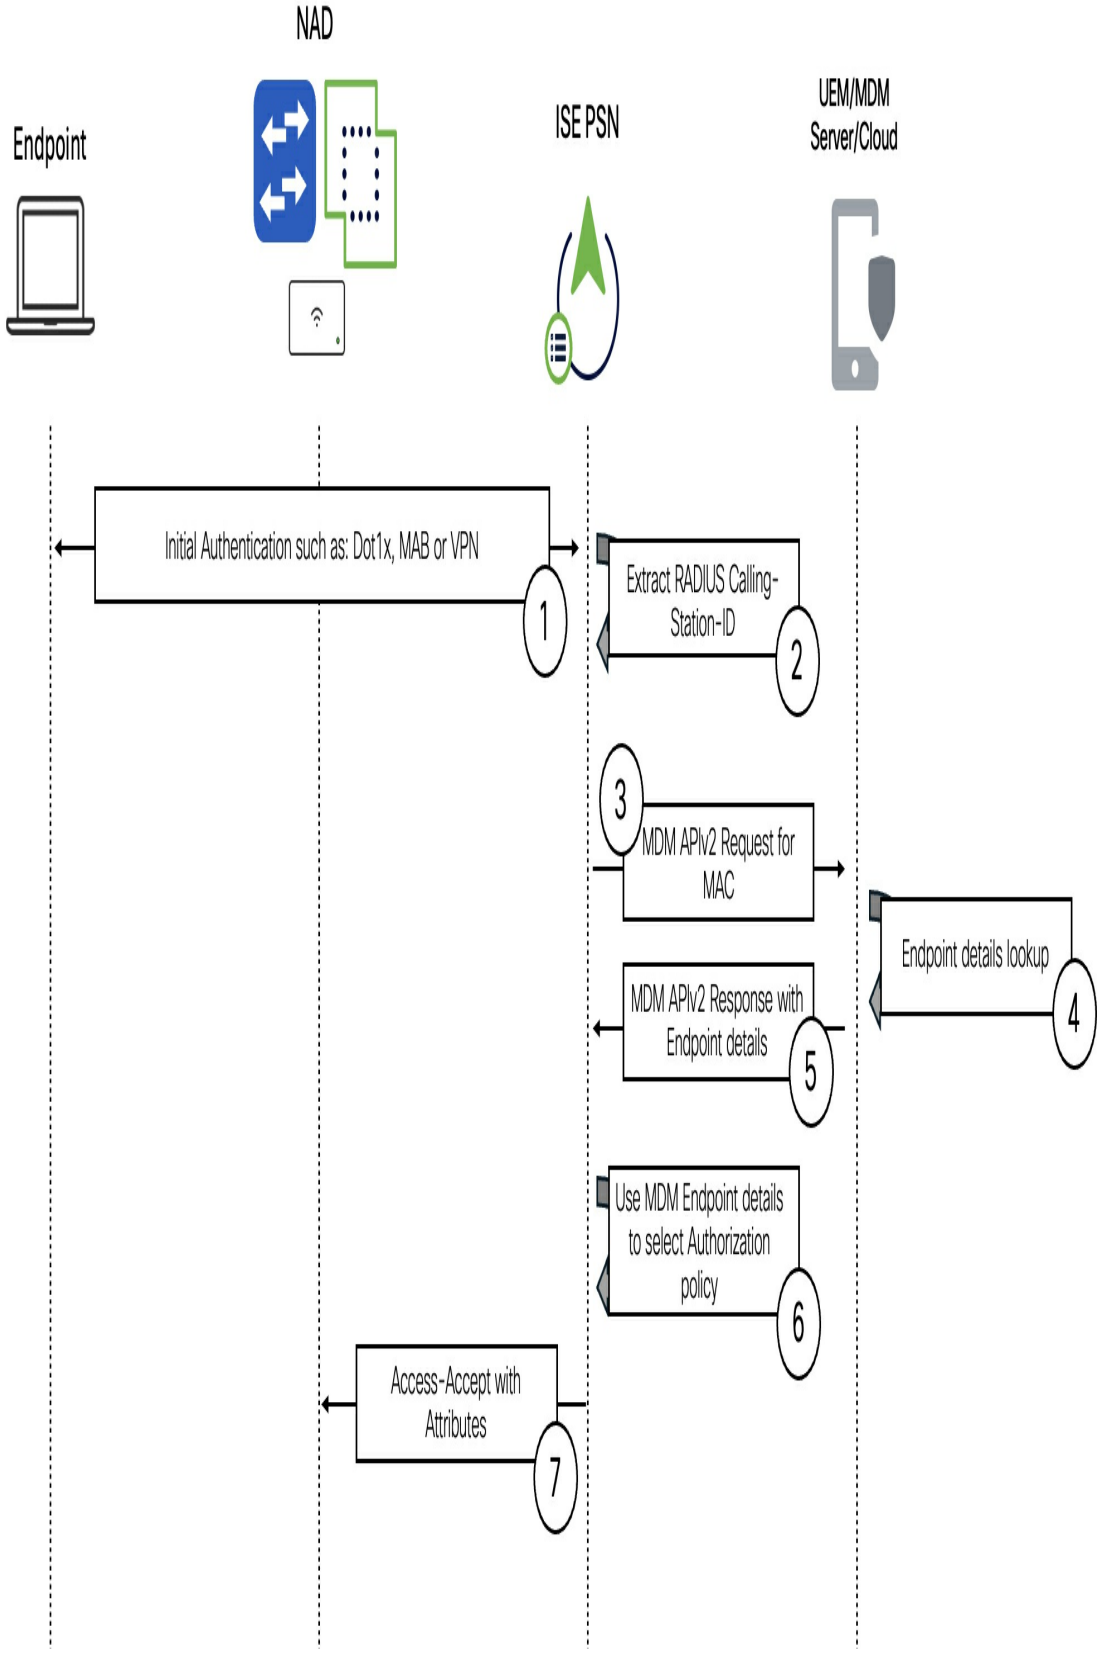

NAD
Endpoint
ISE PSN
UEM/MDM
Server/Cloud
Initial Authentication such as: Dot1x, MAB or VPN
1
Extract RADIUS Calling-Station-ID
2
3
MDM APIv2 Request for MAC
Endpoint details lookup
4
MDM APIv2 Response with Endpoint details
5
Use MDM Endpoint details to select Authorization policy
6
Access-Accept with Attributes
7

Figure 10-39 *Steps of the MDM APIv2 Flow*

The steps involved in the ISE MDMv2 flow are as follows:

Step 1. Authentication is performed. As part of this process, the network access device uses the endpoint MAC address as the RADIUS Calling-Station-ID value in all Access-Requests for wired and wireless Dot1x. In the VPN use case, the endpoint MAC address is delivered to ISE within the AnyConnect Identity Extensions (ACIDEX) attribute called mdm-tlv=device-mac. This process works only for Cisco VPN NADs, such as Secure Firewall and Secure Access Remote Access VPN, and requires Cisco Secure Client.

Step 2. The Policy Service Node extracts the MAC address from the RADIUS Calling-Station-ID or ACIDEX attribute.

Step 3. The v2 API request is sent to the MDM server.

Step 4. The MDM server performs device details lookup using the MAC address.

Step 5. The response is returned to the ISE Policy Service Node. The following attributes represent the maximum amount of detail the MDM server can return for the endpoint. Consult your MDM server documentation to determine which attributes are supported by your specific solution.

- DaysSinceLastCheckin
- Manufacturer
- OsVersion
- DeviceComplianceStatus
- MDMFailureReason
- PhoneNumber

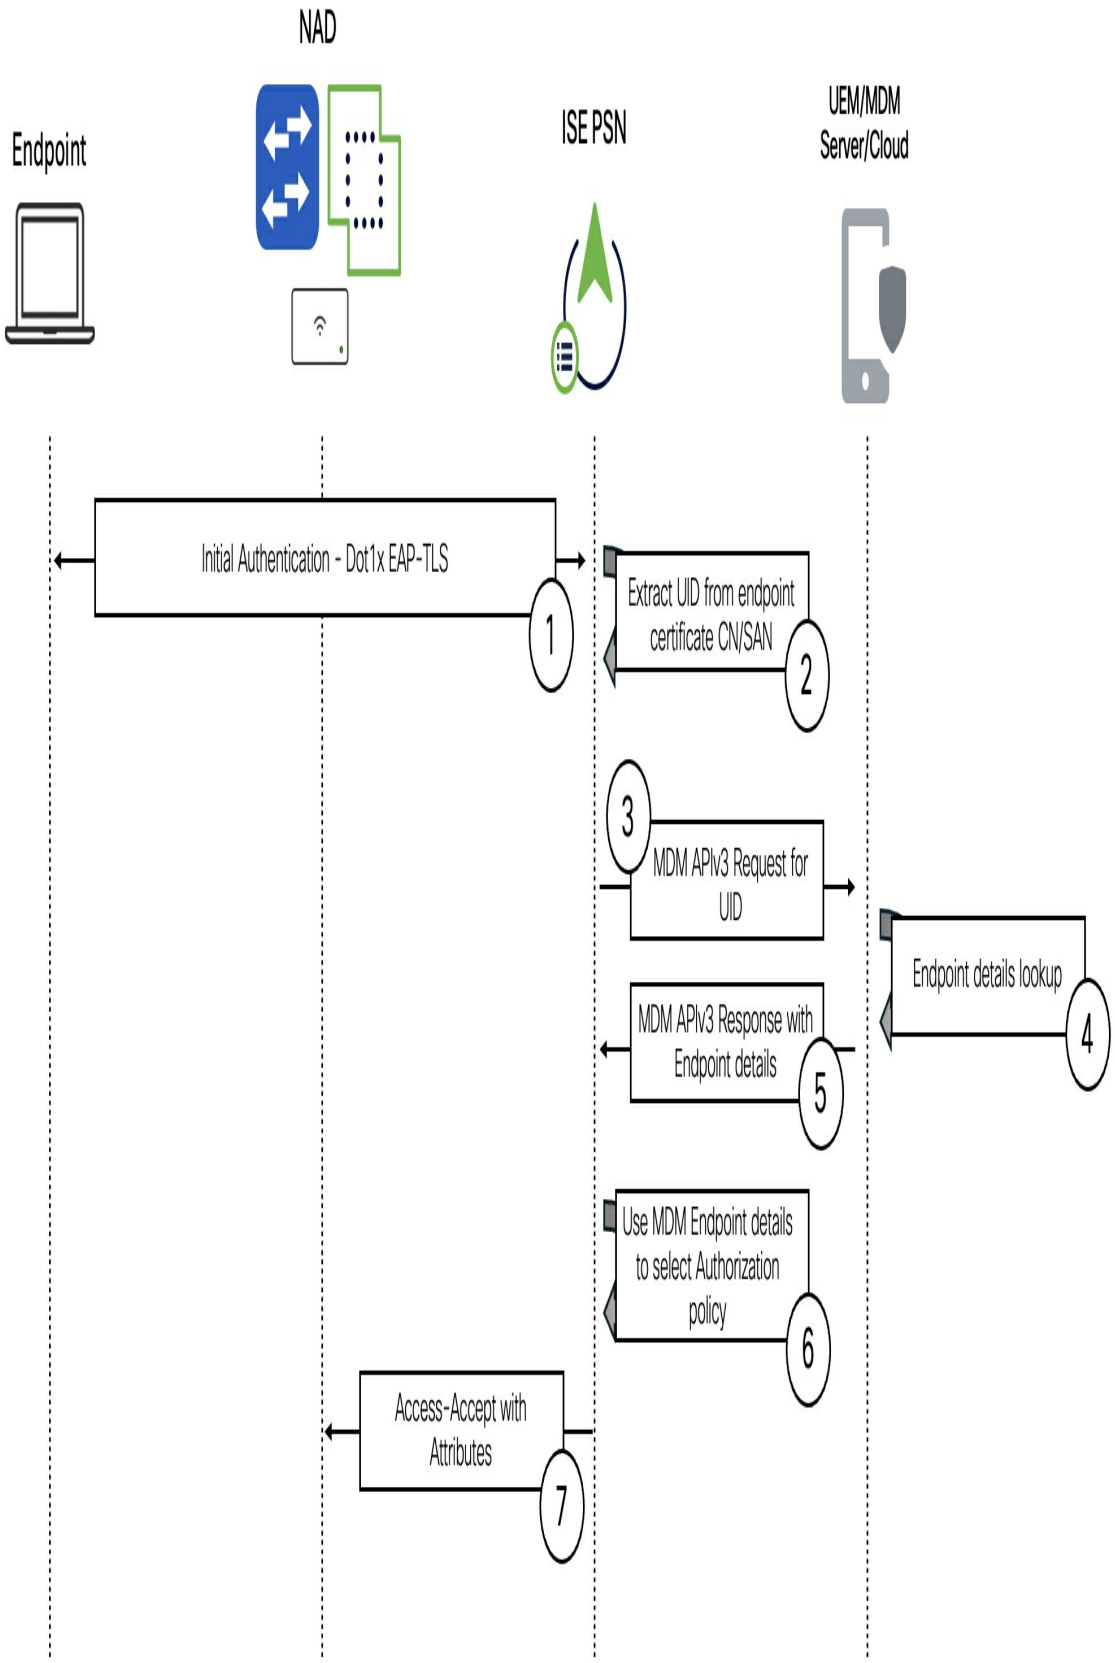
Technet24

- DeviceRegistrationStatus
- MDMServerName
- PinLockStatus
- DiskEncryptionStatus
- MDMServerReachable
- SerialNumber
- IMEI
- MEID
- ServerType
- JailBroken
- Model
- UDID
- UserNotified

Step 6. ISE adds MDM attributes into the session or endpoint context and uses these contexts to look up the authorization policy.

Step 7. The Access-Accept response with the appropriate authorization attributes is returned to the network access device.

The MDM APIv3 flow is shown in Figure 10-40.


NAD
Endpoint
ISE PSN
UEM/MDM
Server/Cloud
Initial Authentication - Dot1x EAP-TLS
1
Extract UID from endpoint certificate CN/SAN
2
3
MDM APIv3 Request for UID
Endpoint details lookup
4
MDM APIv3 Response with Endpoint details
5
Use MDM Endpoint details to select Authorization policy
6
Access-Accept with Attributes
7

Technet24

Figure 10-40 *Steps of the MDM APIv3 Flow*

The following steps are involved in the ISE MDMv3 flow:

Step 1. Authentication is performed. For MDMv3 API, authentication must be completed using the endpoint identity certificate so the UID can be extracted from the certificate. Certificate provisioning needs to occur before authentication. The best approach is to use MDM as a certificate authority to provision such certificates. The UID value is generated by the MDM for each endpoint.

The following MDM/UEM servers currently support Cisco ISE MDM API version 3:

- Ivanti (previously MobileIron UEM) core and cloud UEM services
- Microsoft Endpoint Manager Intune
- JAMF Casper Suite
- Omnissa (previously AirWatch)

Step 2. The Policy Service Node extracts the UID from the endpoint certificate CN and SAN fields.

Step 3. The v3 API request is sent to the MDM server.

Step 4. The MDM server performs device details lookup using the UID value.

Step 5. The response is returned to the ISE Policy Service Node.

Step 6. ISE adds MDM attributes into the session or endpoint context and uses these contexts to look up the authorization policy.

Step 7. The Access-Accept response with the appropriate

authorization attributes is returned to the network access device.

The MDM SCCM flow is shown in Figure 10-41.

Technet24

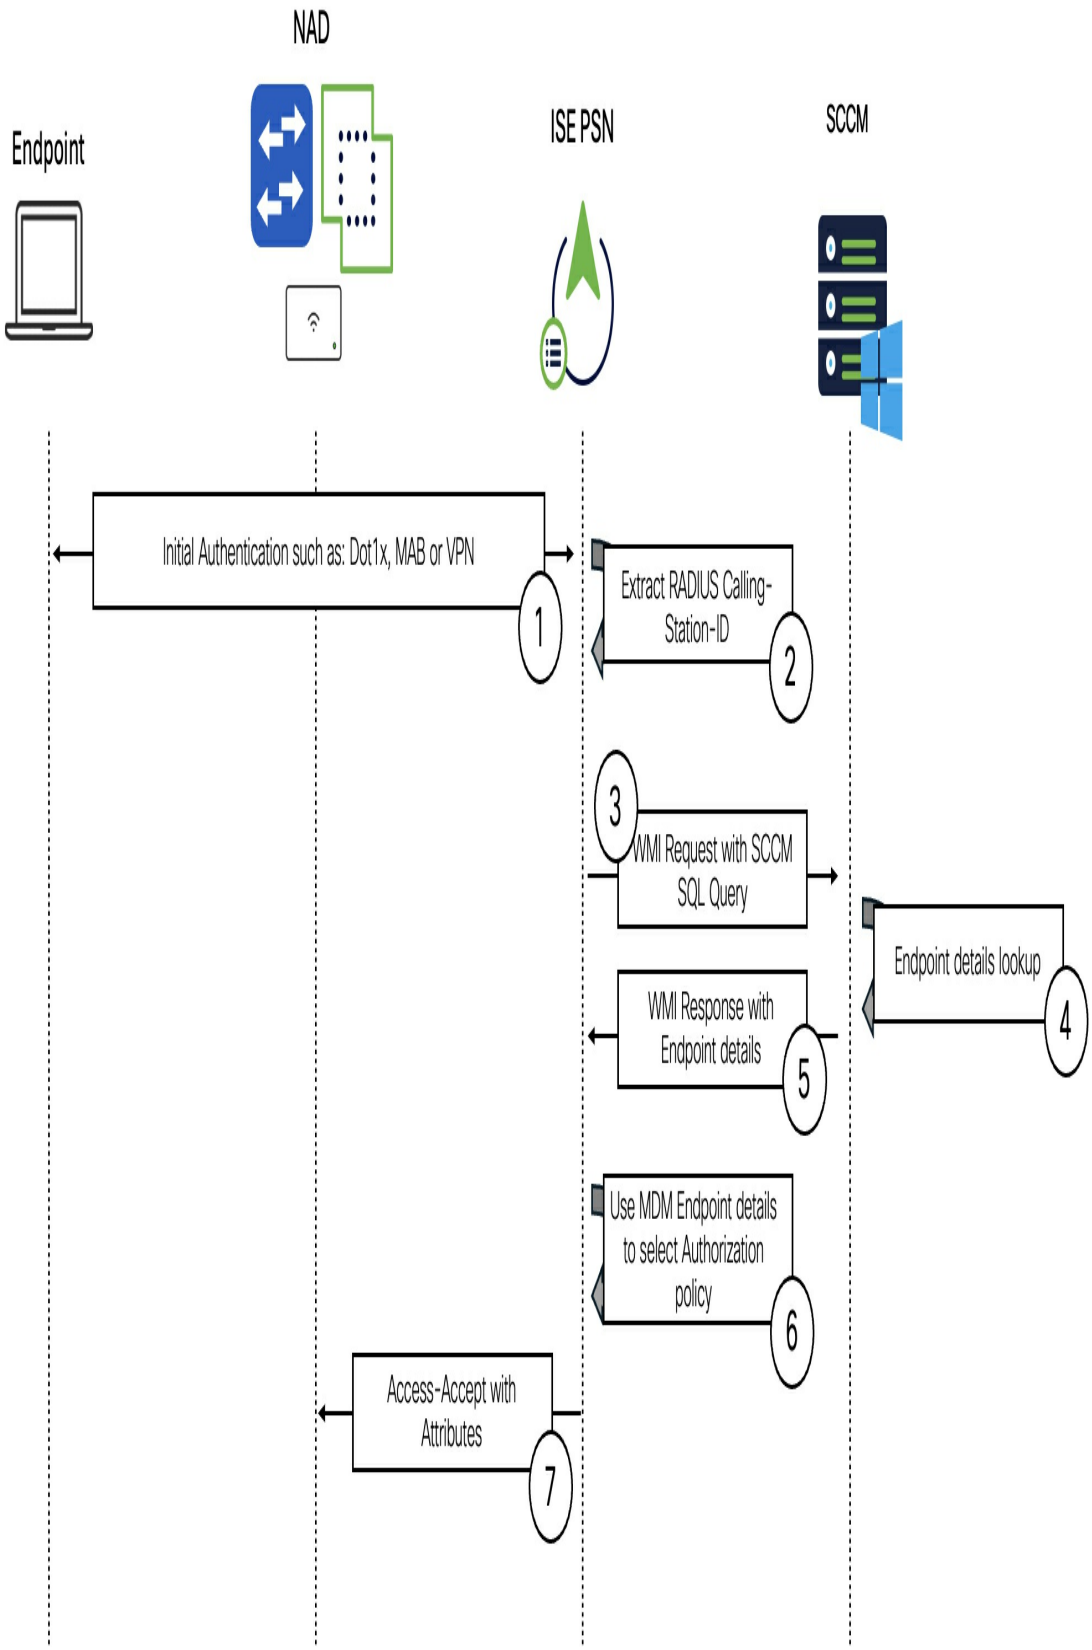

NAD
Endpoint
ISE PSN
SCCM
Initial Authentication such as: Dot1x, MAB or VPN
1
Extract RADIUS Calling-Station-ID
2
3
WMI Request with SCCM SQL Query
Endpoint details lookup
4
WMI Response with Endpoint details
5
Use MDM Endpoint details to select Authorization policy
6
Access-Accept with Attributes
7

Figure 10-41 *Steps of the MDM SCCM Flow*

The following steps outline the ISE MDM SCCM flow:

Step 1. Authentication is performed.

Step 2. The Policy Service Node extracts the MAC address from the RADIUS Calling-Station-ID or ACIDEX attribute.

Step 3. The PSN sends an SCCM SQL query using WMI/MS-RPC mechanisms to obtain the compliance status of all configuration baselines assigned to the given endpoint MAC address.

Step 4. SCCM executes the query and prepares a response.

Step 5. The response is returned to the ISE Policy Service Node. The PSN evaluates the compliance status of all CBs and marks the endpoint as compliant only if all CBs have a compliant status. There is no actual DeviceRegistrationStatus with SCCM, so the PSN assumes that the endpoint is registered if SCCM is aware of it and can return at least some results.

Step 6. The DeviceComplianceStatus and DeviceRegistrationStatus attributes are used by the PSN to match the authorization policy. In the case of SCCM, the PSN determines the DeviceRegistrationStatus based on the result returned from the SCCM SQL query. When the result is not empty, DeviceRegistrationStatus is assumed as Registered.

Step 7. The Access-Accept response with the appropriate authorization attributes is returned to the network access device.

Configuring ISE MDM

Technet24

The MDM integration steps can be divided into three phases:

- **Preparation:** On the ISE side, this phase includes verifying connectivity to the MDM server and confirming that the MDM server certificate is trusted by ISE. The remaining settings must be performed on the MDM server side and will vary depending on the server.
- **MDM server configuration:** Before you can populate authorization policies with MDM conditions, the actual integration must be configured. You need to add the MDM server to ISE and confirm that the connection is successful.
- **Authorization policies configuration:** Once the integration is ready, you can adjust the authorization policy settings to include MDM attributes.

Preparation

When you're considering connectivity requirements between ISE and the MDM server, it is important to address the following questions:

Is the MDM server located inside or outside the corporate network?
For on-premises services, establish the internal network path and identify any firewall or other network security appliances along the way. For servers located on the Internet, a proxy server may need to be configured on ISE if required in the given network.

What are the connectivity requirements for the specific server type?
For most mobile device managers using APIv2 or APIv3, connectivity relies on HTTPS. Typically, connections are established over TCP port 443, but ISE also supports custom destination port

numbers.

For Microsoft Configuration Manager, connectivity requirements are broader. You need to allow TCP 135 and TCP 445 from ISE to the SCCM server.

Is the MDM server certificate trusted by ISE?

Consult with the MDM server administrator to clarify which certificate is used to secure the MDM ISE API connection. If the issuer of this certificate is unknown to ISE, address this by installing the issuer and any other certificate authorities in the chain into the ISE trusted store.

To verify proxy settings on ISE or specify the proxy server, navigate to **Administration > System > Settings > Proxy**.

To install the issuer of the MDM server certificate, navigate to **Administration > System > Certificates > Trusted Certificates** and click **+ Import** to add the certificate. Ensure that the **Trust for Authentication of Cisco Services** checkbox is checked during the certificate import. An example of the certificate import is demonstrated in Figure 10-42.

Technet24

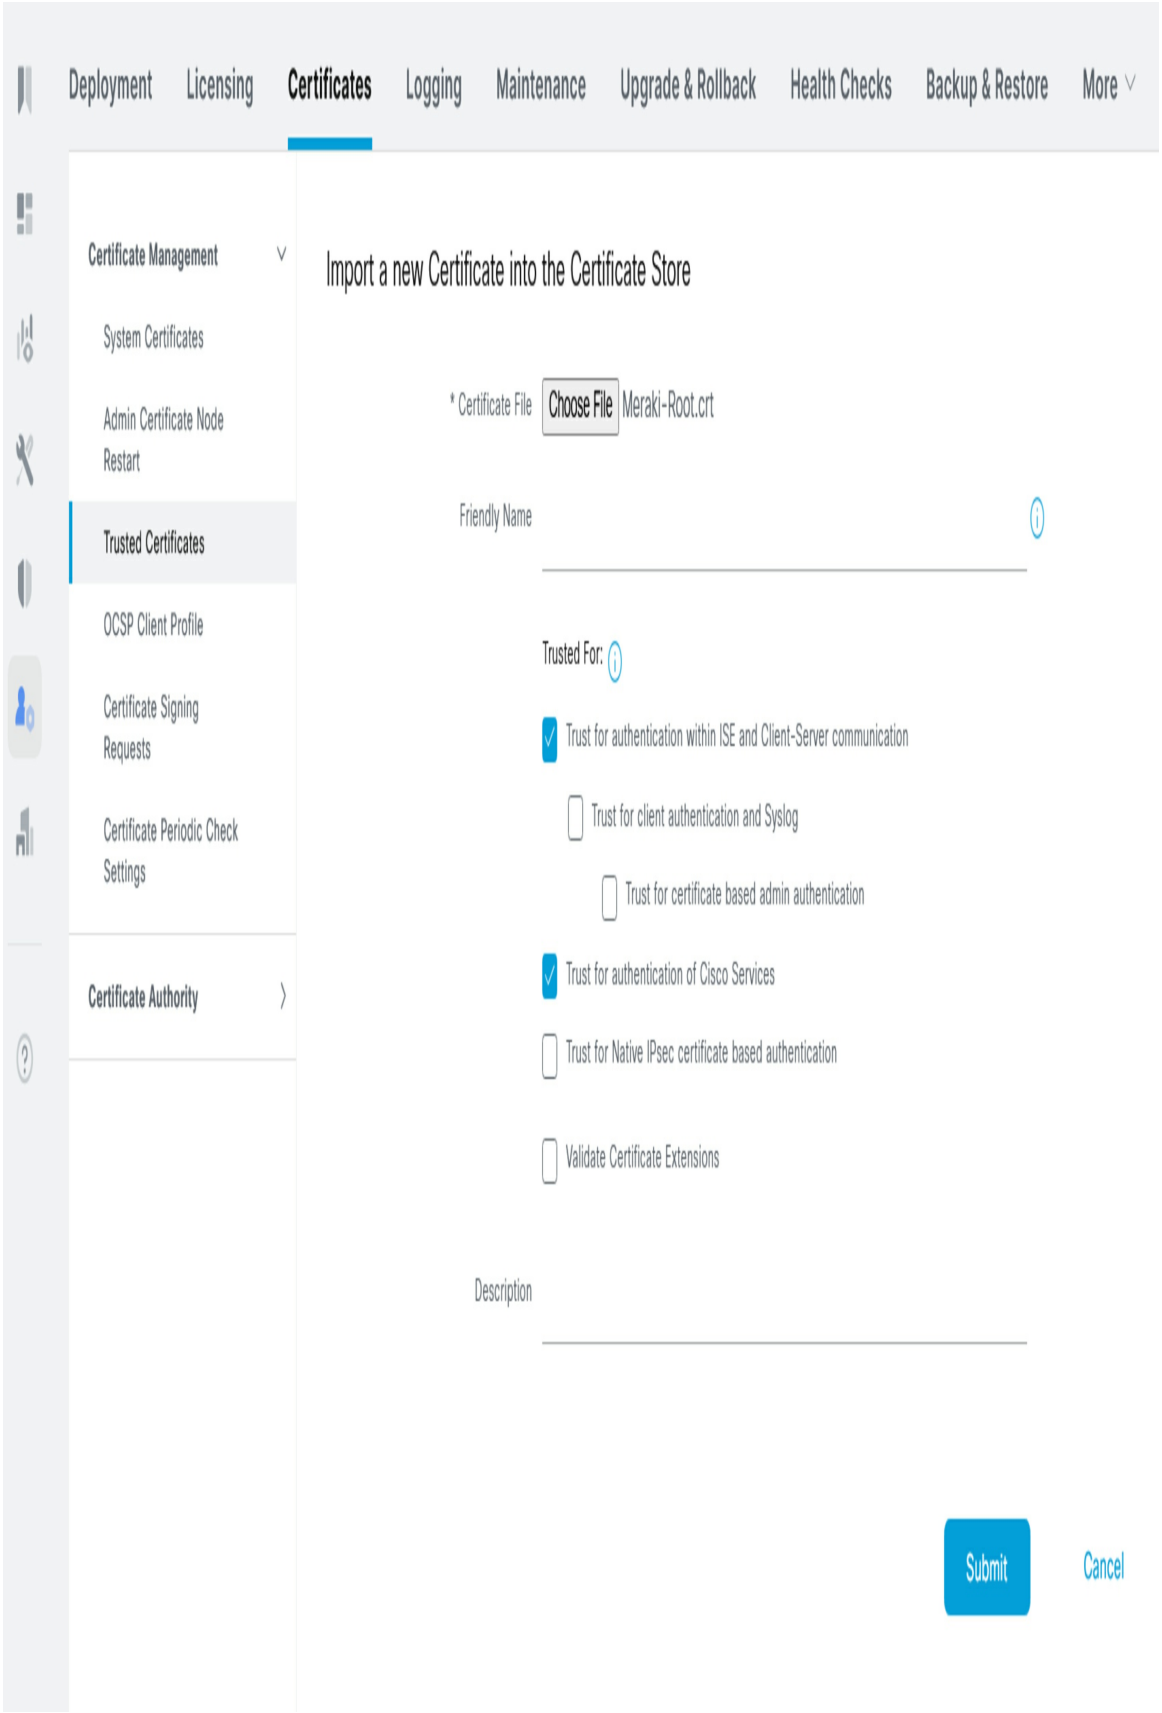

Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Certificate Management
System Certificates
Admin Certificate Node Restart
Trusted Certificates
OCSP Client Profile
Certificate Signing Requests
Certificate Periodic Check Settings
Certificate Authority
Import a new Certificate into the Certificate Store
* Certificate File
Choose File
Meraki-Root.crt
Friendly Name
Trusted For:
Trust for authentication within ISE and Client-Server communication
Trust for client authentication and Syslog
Trust for certificate based admin authentication
Trust for authentication of Cisco Services
Trust for Native IPsec certificate based authentication
Validate Certificate Extensions
Description
Submit
Cancel

Figure 10-42 *MDM Issuer Certificate Installation*

MDM Server Configuration

To configure the MDM server on ISE, navigate to **Administration > Network Resources > External MDM** and click **Add** to configure a new MDM server. The server configuration is outlined in Figure 10-43.

Technet24

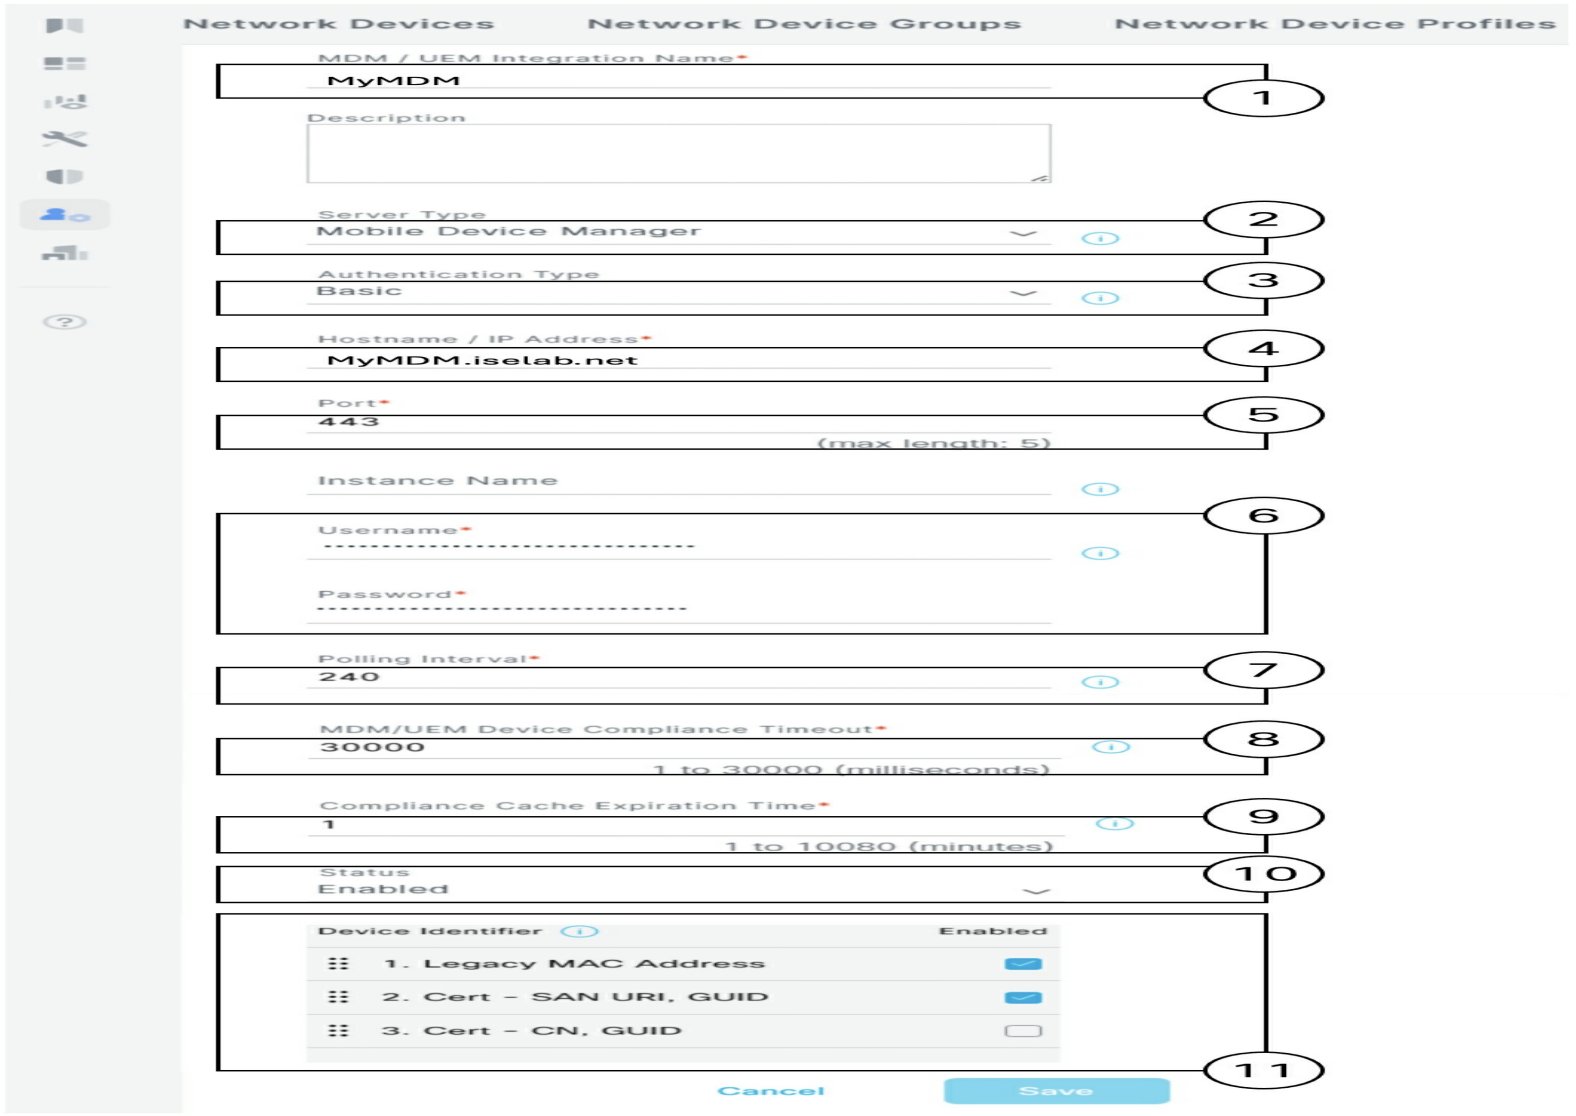

Network Devices
Network Device Groups
Network Device Profiles
MDM / UEM Integration Name*
MyMDM
1
Description
Server Type
Mobile Device Manager
2
Authentication Type
Basic
3
Hostname / IP Address*
MyMDM.iselab.net
4
Port*
443
(max length: 5)
5
Instance Name
Username*
Password*
6
Polling Interval*
240
7
MDM/UEM Device Compliance Timeout*
30000
1 to 30000 (milliseconds)
8
Compliance Cache Expiration Time*
1
1 to 10080 (minutes)
9
Status
Enabled
10
Device Identifier
Enabled
1. Legacy MAC Address
2. Cert – SAN URI, GUID
3. Cert – CN, GUID
11
Cancel
Save

Figure 10-43 *MDM Server Configuration Steps*

The following steps are involved in MDM server configuration:

Step 1. Configure the MDM server name. This name identifies the server in the UI and is also used as the MDMServerName endpoint attribute after the endpoint status has been obtained from the server. This simplifies MDM server selection for ISE if more than one MDM server is configured.

Step 2. Choose the server type. For this example, the Mobile Device Manager type is used.

Step 3. Select the authentication method.

Step 4. Specify the IP address or FQDN of the MDM server.

Step 5. Specify the MDM server port.

Step 6. Enter MDM server credentials.

Step 7. Define a polling interval, which by default has a value of 240 seconds. Upon polling interval expiration, ISE contacts the MDM server to obtain the list of all noncompliant endpoints so they can be disconnected from the network.

Step 8. Specify device compliance timeout. This is the time ISE waits for a reply from the MDM server. If the multiple MDM server feature is used, the recommended value for the timeout is 500 milliseconds.

Step 9. The Compliance Cache Expiration Time feature is similar to the posture lease but is used with MDM. When a sufficiently long timer is set, ISE will not query the MDM server if it already received a result for the endpoint in the past and the Compliance Cache Expiration Time is still

Technet24

valid.

Step 10. Define the MDM server status. The default value is Disabled. For the server to be used, it must be in the Enabled state.

Step 11. Configure the attribute that ISE should use to query information on a specific endpoint from the MDM. When both MAC and Certificate fields are used, ISE can choose either the v2 or v3 version of the API. The attributes need to be set in priority order to instruct ISE which attribute should be tried first.

Authorization Policies Configuration

As usual, authorization policies configuration begins with defining authorization profiles. You may choose from the following approaches:

- **Authorization profiles and policies that include inline enrollment:** In this case, ISE can redirect the user to the MDM server as part of the CWA flow, simplifying the MDM enrollment process for the user.
- **Authorization profiles and policies without inline enrollment:** In this use case, profiles and policies typically handle compliant and noncompliant scenarios. In the compliant scenario, the user receives full network access, whereas in the noncompliant scenario, some form of restricted access is provided, for example, using DACL or VLAN.

Note

Inline enrollment is only supported with MDM servers of the type Mobile Device Manager.

For the inline enrollment process, the ISE MDM portal is used. The portal simply assists users by redirecting them to the appropriate MDM server and, in most cases, does not require any additional configuration beyond assigning the correct certificate using the Portal Certificate Group Tag. You can edit or check settings of an existing portal or create a new one if needed by navigating to **Administration > Device Portal Management > Mobile Device Management**.

Later, you can specify the MDM portal within the authorization profile by selecting both the portal and the MDM server, as shown in Figure 10-44.

Technet24

Authorization Profiles 〉 New Authorization Profile

Authorization Profile

* Name: MyMDM-Enroll

Description:

* Access Type: ACCESS_ACCEPT

Network Device Profile: Cisco

Common Tasks

Web Redirection (CWA, MDM, NSP, CPP)

MDM Redirect | ACL: mdm_redirect | Value: MDM Portal (default) | MDM Server: MyMDM

Static IP/Host name/FQDN

Suppress Profiler CoA for endpoints in Logical Profile

Figure 10-44 *Inline MDM Enrollment Authorization Profile*

An example of the authorization policies that can be used for the inline enrollment use case is shown in Figure 10-45.

Technet24

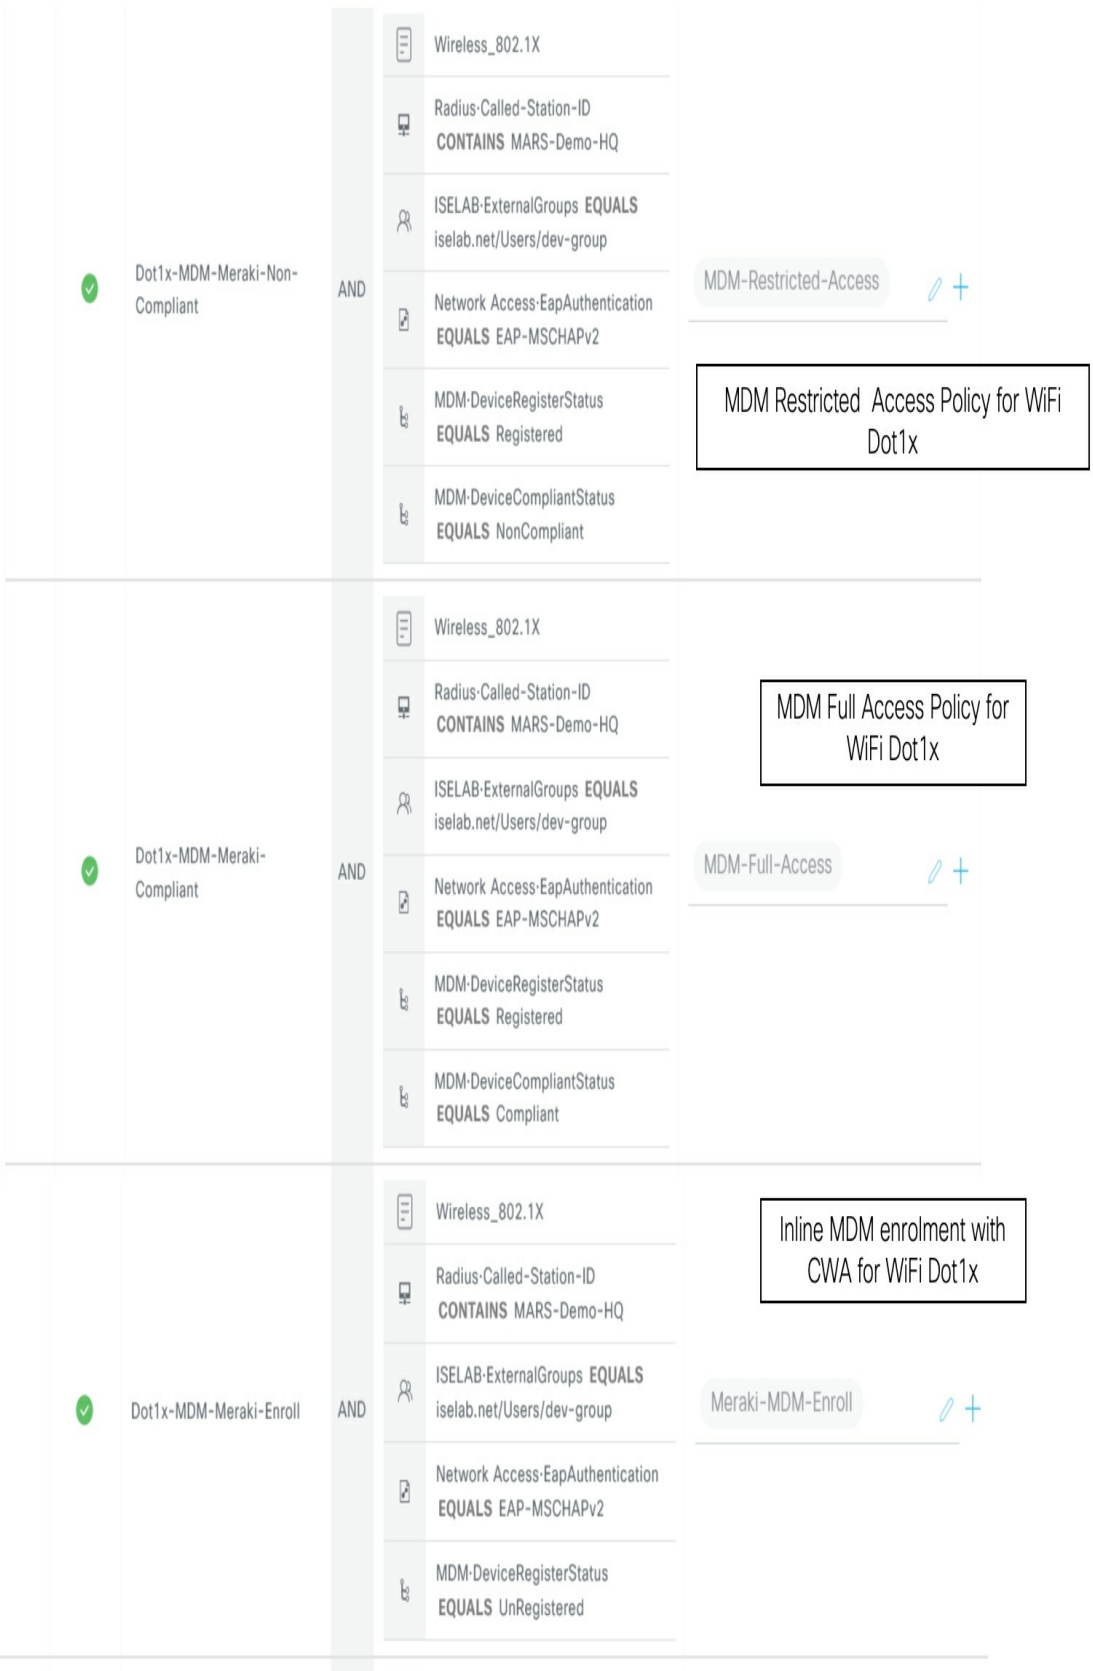

Dot1x-MDM-Meraki-Non-Compliant
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
ISELAB·ExternalGroups EQUALS iselab.net/Users/dev-group
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
MDM·DeviceRegisterStatus EQUALS Registered
MDM·DeviceCompliantStatus EQUALS NonCompliant
MDM-Restricted-Access
MDM Restricted Access Policy for WiFi Dot1x
Dot1x-MDM-Meraki-Compliant
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
ISELAB·ExternalGroups EQUALS iselab.net/Users/dev-group
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
MDM·DeviceRegisterStatus EQUALS Registered
MDM·DeviceCompliantStatus EQUALS Compliant
MDM-Full-Access
MDM Full Access Policy for WiFi Dot1x
Dot1x-MDM-Meraki-Enroll
AND
Wireless_802.1X
Radius·Called-Station-ID CONTAINS MARS-Demo-HQ
ISELAB·ExternalGroups EQUALS iselab.net/Users/dev-group
Network Access·EapAuthentication EQUALS EAP-MSCHAPv2
MDM·DeviceRegisterStatus EQUALS UnRegistered
Meraki-MDM-Enroll
Inline MDM enrolment with CWA for WiFi Dot1x

Figure 10-45 *MDM Authorization Policies*

To adjust your scenario for MDM without inline enrollment, you can use the same policies, simply omitting the policy responsible for inline enrollment.

Summary

This chapter explains how to implement endpoint posture services using Cisco ISE to support Zero Trust security principles in your organization. It explores the foundational components of posture services — endpoints, network devices, policy decision points, and ISE infrastructure — and how each element contributes to effective posture assessment and enforcement.

The text details the key milestones in the posture flow, including authentication, provisioning, posture assessment, remediation, and change of authorization. It describes the different posture agents available — Persistent, Stealth, Temporal, and Agentless — with clear guidance on their use cases, deployment considerations, and session lifecycle implications.

Comprehensive analysis is provided on posture flows, including redirect-based, redirection-less, and agentless methods. You also gain insight into how Cisco Secure Client discovers the Policy Service Node and the technical reasons behind probe precedence and endpoint lookup.

The chapter clarifies the logic that governs posture status, session management, and data replication. It examines the differences between MNT and PSN session management, the causes of stale and phantom sessions, and the mechanisms behind RADIUS Session Directory and Light Data Distribution for posture attribute

Technet24

synchronization.

In addition, the chapter addresses advanced posture features, including periodic reassessment (PRA), continuous reassessment (CRA), and the grace period. It reviews how these features use key attributes — such as PostureExpiry and PostureLastCompliantExpiry — to manage compliance windows and authorization behavior during reassessment or loss of connectivity.

The content guides you through the multiphase configuration of ISE posture services, from preparation and agent provisioning to posture policy creation and authorization logic. Integration with external device management platforms is covered in depth, highlighting the technical differences between MDM API v2 and v3, the use of unique endpoint IDs, and the specific process for integrating Microsoft SCCM through WMI/MS-RPC.

By following the chapter's step-by-step explanations and configuration recommendations, you can design and maintain robust endpoint posture services that support secure access for managed and unmanaged devices across wired, wireless, and VPN environments.

Chapter 11. Managing Network Access in the Meraki Environment

Introduction

This chapter explains the principles of network access control (NAC) in Cisco Meraki cloud-managed networks. It describes the Cisco Meraki dashboard and the core principles of cloud networking. The text covers the history and current state of network access control in Cisco Meraki environments. You will examine NAC implementations for switches, wireless networks, and security appliances.

Cisco Meraki Dashboard and Cloud Networking Principles

Cloud-managed networking always includes both the device component and the controller component. This section explains how to manage devices using the Cisco Meraki dashboard. You will discover how the dashboard enables centralized administration, offering streamlined configuration and monitoring for all connected devices.

Cisco Meraki Dashboard Structure

The Cisco Meraki dashboard, often referred to simply as the dashboard, serves as the central management platform for Meraki

Technet24

classic devices such as access points (MRs), switches (MSs), and security appliances (MXs). You also use the dashboard to manage and monitor selected Catalyst switches and routers. This unified interface streamlines oversight and configuration across diverse network hardware. The dashboard offers a single pane of glass for administration and acts as the main contact point for all cloud-managed devices. It delivers configuration services and collects events and telemetry data. At a high level, the dashboard provides the following services for both administrators and managed devices:

- **Zero-Touch provisioning:** You can connect all cloud-managed devices to the dashboard with minimal or no initial configuration. This approach significantly simplifies the deployment of cloud-managed networks compared to traditional network management methods.
- **Centralized configuration management:** You use the dashboard as the sole interface for configuring your network. You do not need to access individual devices to make configuration changes or updates. The dashboard stores all device configurations centrally and automatically distributes any changes you make to the relevant devices in real time.
- **Centralized software management:** You can manage all software upgrades and downgrades from the dashboard, eliminating the need to update devices individually.
- **Centralized maintenance and troubleshooting:** The dashboard collects telemetry and events from each device, making this information easily accessible to you. The unified interface streamlines daily operations and troubleshooting.

Figure 11-1 illustrates the logical structure of the Meraki dashboard from the administrator's perspective.

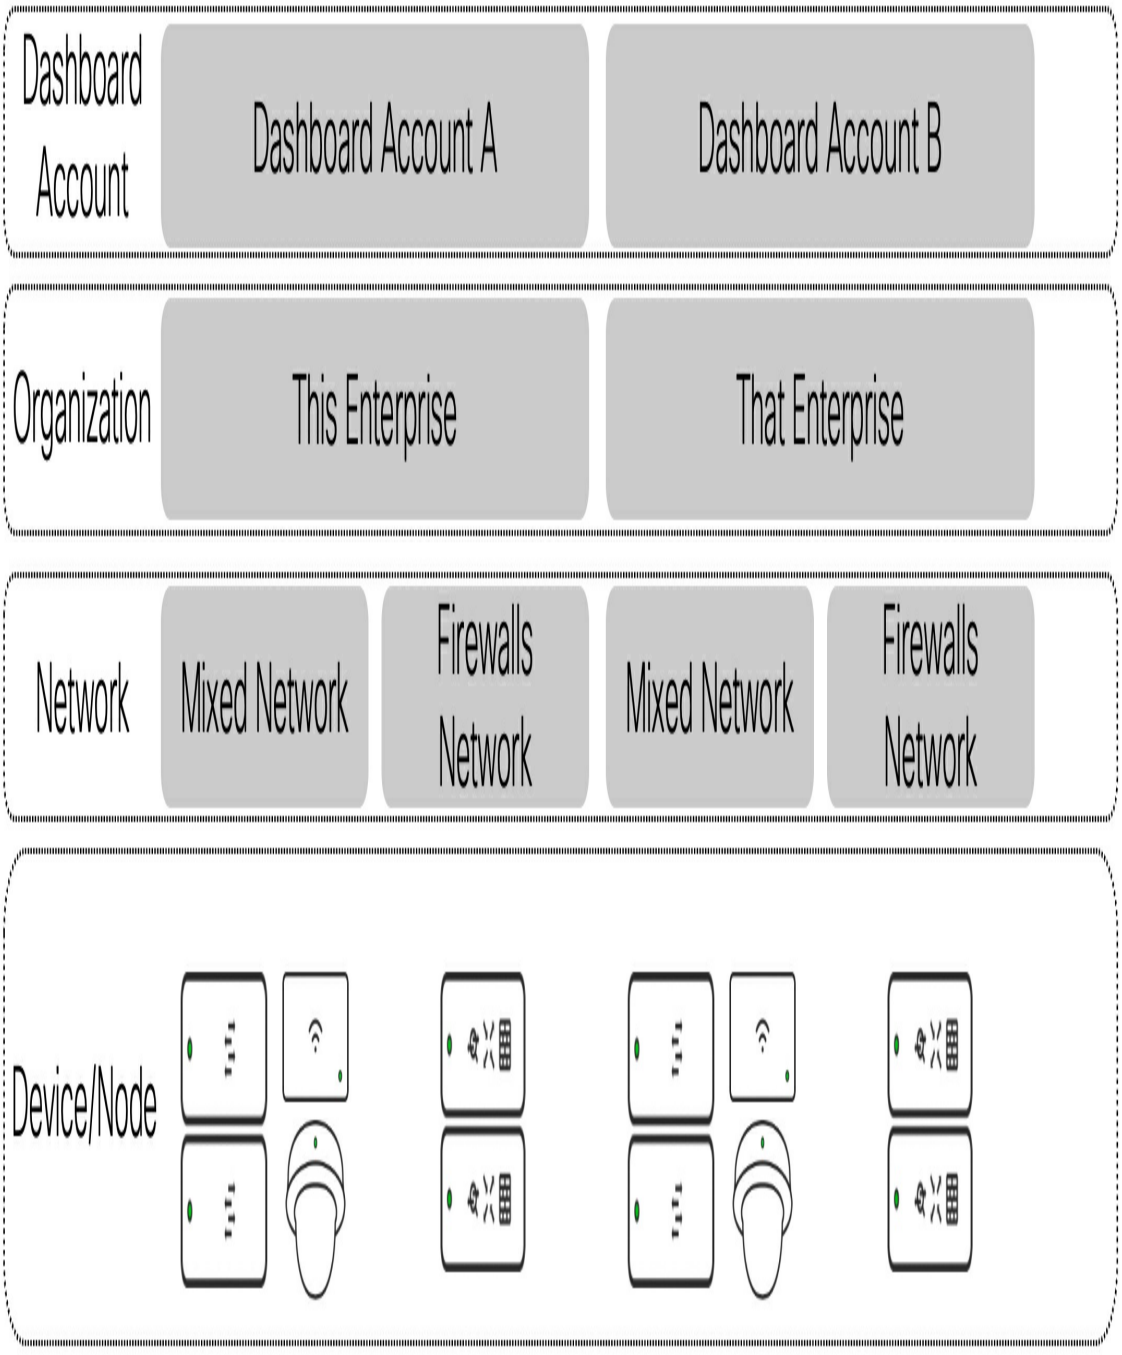


Figure 11-1 *Meraki Dashboard Logical Structure*

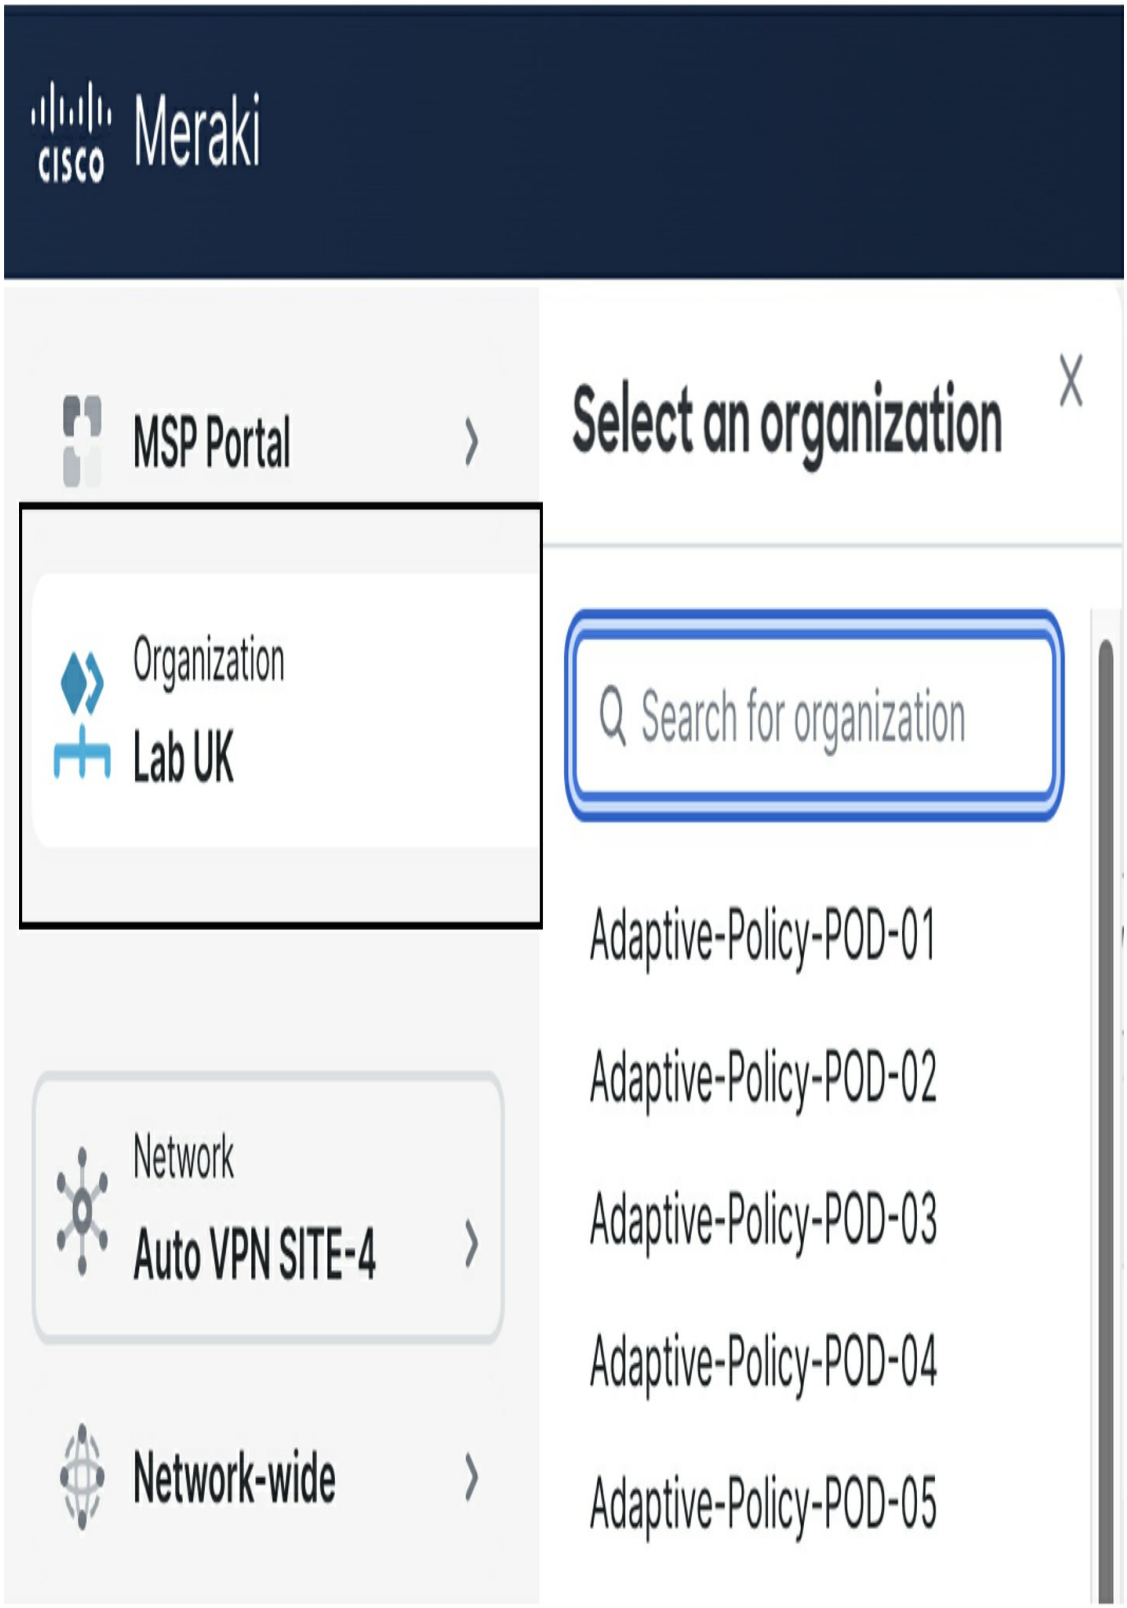
Now let's review different dashboard logical components one by one.

Technet24

The administrator account is the main element of the dashboard. With an administrator account, you can manage one or multiple organizations:

- **Single-organization administrator account:** This scenario is the most common and applies when your real-life organizational structure fits within a single dashboard organization. For example, a typical business, such as a bank or hotel chain with centralized management, uses a single organization from the dashboard perspective.
- **Multi-organization administrator account:** This scenario applies to business groups where each business entity requires a separate organization. Another common example is an administrator in a managed service provider (MSP) environment, where you need to manage multiple organizations.

Figure 11-2 shows the Cisco Meraki dashboard multi-organization view, where you can select the appropriate organization by using the **Organization** drop-down list.


Meraki
MSP Portal
Organization
Lab UK
Network
Auto VPN SITE-4
Network-wide
Select an organization
Search for organization
Adaptive-Policy-POD-01
Adaptive-Policy-POD-02
Adaptive-Policy-POD-03
Adaptive-Policy-POD-04
Adaptive-Policy-POD-05

Technet24

Figure 11-2 *Cisco Meraki Dashboard Multi-Organization View*

The next element in the logical dashboard structure is the *organization*. In the Cisco Meraki dashboard, an organization consists of a collection of networks that belong to a single entity, such as a company or school district. The main components of an organization include, besides networks themselves, the following:

- **Inventory:** Lists all devices that you have added to the organization.
- **Licenses:** Displays the licenses required to support the available devices and their respective functions and feature set levels.

Beneath the organization level, you will find *networks*. A network serves as the primary building block within each organization. Networks contain Cisco Meraki devices, their configurations, statistics, and information about client devices.

Meraki networks can be of two types:

- **Combined network:** This network type allows you to group different device types together, typically based on physical location.
- **Device-specific networks:** These networks focus on a single device type, such as switching or wireless. This approach gives you greater flexibility in grouping devices for effective management, maintenance, and troubleshooting.

The last element — and the core operational component — of the logical dashboard structure is the *device* or *node*. Devices include individual components such as switches, access points, and cameras.

Now that we've reviewed the logical structure of the dashboard, let's examine its physical aspects, including its components and operational

principles. This section describes what the dashboard consists of and explains how it functions.

As the saying goes, "The cloud is just someone else's computer." This statement holds true for the Cisco Meraki dashboard. On a physical level, the dashboard consists of geographically distributed data centers that serve different regions. Within each region, data centers are organized into highly available groups to ensure that an outage at a single regional data center does not impact operations within that region. Figure 11-3 provides a high-level overview of the dashboard physical structure.

Technet24

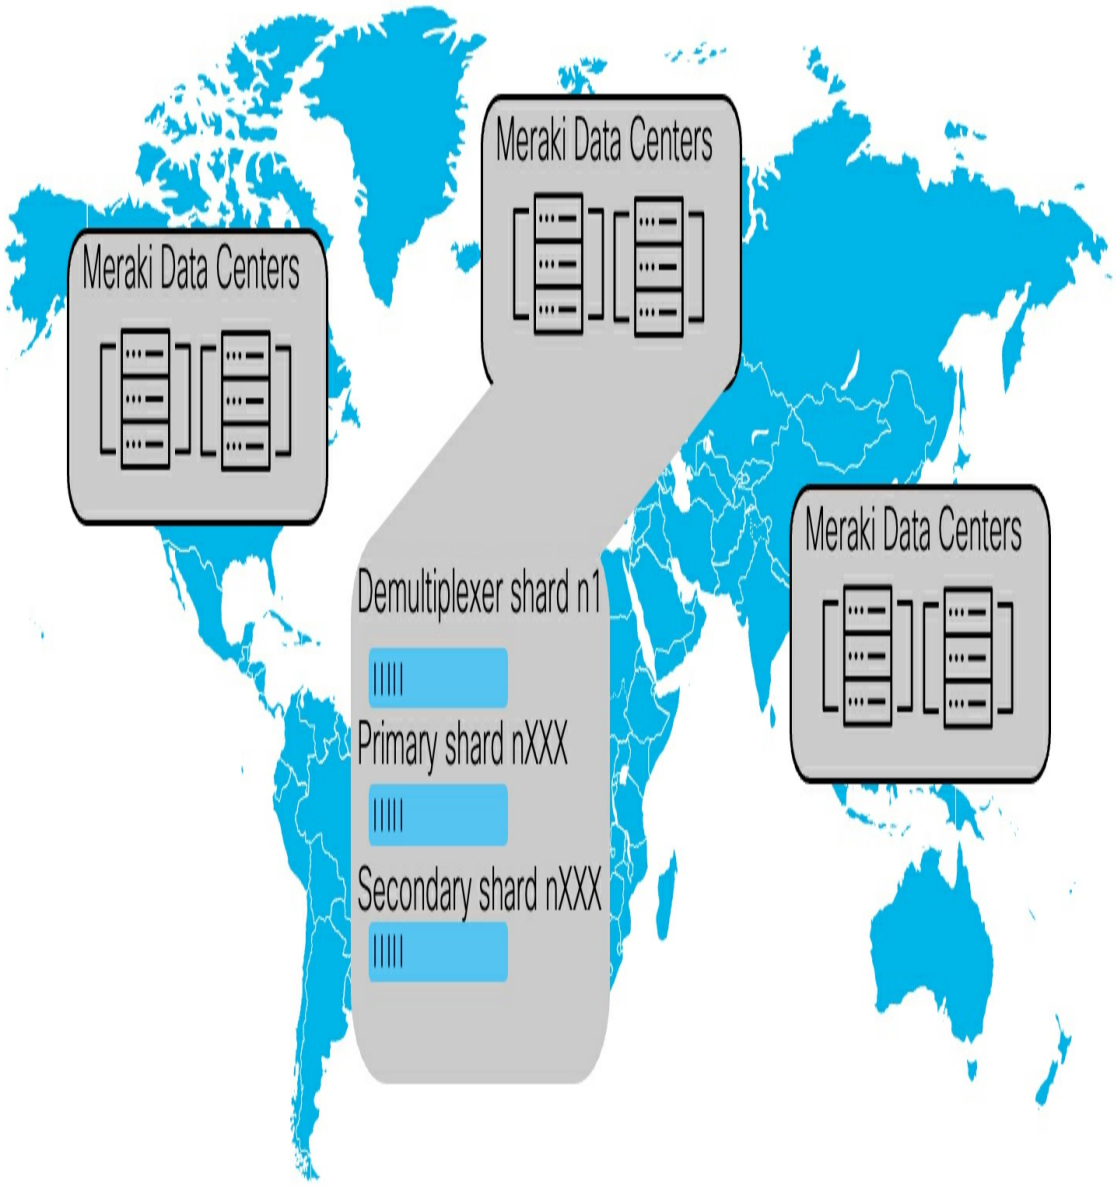


Figure 11-3 *Cisco Meraki Dashboard Physical Structure*

If you examine the individual data centers more closely, you will notice an icon depicting servers labeled with the word *shard*. A shard is a dashboard computing element, like a large and powerful server. It

consists of two main components: the front end, which provides the dashboard user interface for administrators, and the back end, which translates administrators' commands into configurations for cloud-managed devices and propagates those configurations to the devices. Each shard manages a subset of customer organizations. When you, as administrator, perform operations on the dashboard or when a network device registers to the dashboard, communication always occurs with the specific shard assigned to your organization at the time of its creation.

If you examine the browser address bar while the Meraki dashboard is open, you may notice a URL in the format n*XXX*.dashboard.meraki.com, where n*XXX* represents the ID of your shard.

Let's review the types of Meraki dashboard shard:

- **Demultiplexer shard n1:** Each region contains a pair of n1 demultiplexer shards, consisting of a primary and secondary shard. These shards do not host any customer organizations. Instead, they serve as a sole source of truth when an administrator or device needs to be directed to the appropriate shard. The n1 shard maintains a copy of the entire Meraki dashboard's logical structure and global inventory, enabling it to redirect requests from devices and users to the correct shards.
- **Primary shards:** Administrators and devices interact with their respective primary shards. These shards provide the user interface for administrators and deliver back-end connectivity for managed devices.
- **Secondary shard:** Each primary shard has a corresponding secondary shard that serves as a backup in case any issues arise with the primary. Secondary shards reside in a different data

Technet24

center within the same region to ensure that a data center outage does not interrupt Meraki service.

Meraki Use Cases with ISE for Network Access Control

When you implement network access control in Meraki networks, you typically work with three types of Meraki devices, each with its own set of use cases.

Meraki MS and CS Switches

Meraki MS and CS switches enable you to implement wired network access control within a Meraki environment. Meraki MS switches offer a traditional Meraki switching solution, whereas Meraki CS switches introduce a new approach by allowing you to manage Cisco Catalyst switches through the Meraki dashboard.

Both switch families provide essential NAC features, such as these:

- Support for open and closed authentication modes
- 802.1X authentication
- MAC Authentication Bypass (MAB)
- RADIUS accounting with START, STOP, and INTERIM updates
- RADIUS server monitoring and dead server detection
- Caching of RADIUS authorization results to maintain consistent access for endpoints and users when RADIUS servers are unreachable
- Support for RADIUS Change of Authorization (CoA)
- VLAN assignment for voice and data clients

- Voice and Data domains assignment
- Support for various local VLAN assignment use cases, including Failed Auth VLAN, Guest VLAN, and Critical Auth VLANs for voice and data
- Authentication bypass for IP phones
- Meraki group policy assignment
- Device sensor capabilities for reporting endpoint details to Cisco Identity Services Engine over RADIUS, enabling more granular policy assignment based on device profiles
- URL redirection to enable guest access, posture assessment, and BYOD onboarding
- Adaptive Policy support on selected models

> **Note**
>
> Adaptive Policy is a Meraki implementation of Cisco TrustSec.

Although both MS and CS switches offer these features, you may encounter differences in their implementation across the two families. Feature support can also vary within a single family of switches. For example, not all MS models support device sensor capabilities. You will find more details on supported features and specific models later in the chapter.

Meraki MR Access Points

Meraki MR access points provide network access control for cloud-managed wireless networks. At the time of writing, the Meraki MR portfolio consists of two main groups: classic Meraki wireless devices (models that include "MR") and Cisco wireless access points (models that include "CW"). Support for specific features and use cases, as well

Technet24

as their implementation details, may vary between access point models and across different device families.

The supported wireless NAC use cases are as follows:

- 802.1X authentication for WPA Enterprise networks
- MAC-based authentication for open networks
- MAC-based authentication for WPA Personal networks with Identity Pre-Shared Key (iPSK) lookup
- RADIUS accounting with START, STOP, and INTERIM updates
- RADIUS server monitoring and dead server detection
- Support for RADIUS CoA
- VLAN assignment for clients
- Meraki group policy assignment
- Device sensor capabilities for reporting endpoint details to Cisco ISE over RADIUS, enabling more granular policy assignment based on device profiles
- URL redirection to enable use cases such as guest access, posture assessment, and BYOD onboarding
- Adaptive Policy support on selected models

MX Security Appliances

On the security appliances side, support for NAC use cases varies depending on the MX model type. MX appliances fall into three groups: low-end appliances, low-end appliances with integrated wireless, and high-end appliances.

As of this writing, the low-end appliances include the following models: MX67/68(C/W/CW), MX75, Z3(C), and Z4(C). Low-end

appliances with integrated wireless include models that contain a "W" in the model name — for example, MX67W. High-end appliances comprise models that do not fall into the previous two categories.

Low-end MX appliances:

- Port-based 802.1X authentication.
- MAC Authentication Bypass.
- Hybrid authentication. In this mode, the MX appliance first attempts 802.1X authentication. If the client does not respond, the appliance fails over to MAB.
- RADIUS-based authentication for remote access VPN.
- Meraki group policy assignment for remote access VPN.
- Adaptive Policy SGT inter-VLAN transport.
- Adaptive Policy SGT Auto VPN transport.

Low-end MX appliances with integrated wireless:

- Port-based 802.1X authentication.
- 802.1X authentication for WPA Enterprise networks.
- MAC Authentication Bypass.
- Hybrid authentication. In this mode, the MX appliance first attempts 802.1X authentication. If the client does not respond, the appliance fails over to MAB.
- RADIUS-based authentication for remote access VPN.
- Meraki group policy assignment for remote access VPN.
- Adaptive Policy SGT inter-VLAN transport.
- Adaptive Policy SGT Auto VPN transport.

Technet24

Note

At the time of writing, port-based authentication on the MX does not support any type of authorization attribute assignment from the RADIUS server. RADIUS accounting is not supported for port-based authentication or remote access VPN use cases.

High-end MX appliances:

- RADIUS-based authentication for remote access VPN.
- Meraki group policy assignment for remote access VPN.
- Adaptive Policy SGT inter-VLAN transport.
- Adaptive Policy SGT Auto VPN transport.

RADIUS Flow with Different Meraki Platforms

It is important to understand how RADIUS communications are handled by different Meraki platforms because this can impact planning and design decisions. This section explains the RADIUS communication flow for various Meraki devices. You'll learn how each platform processes authentication requests and manages network access through RADIUS interactions.

High-Level Flow Explanation

When you know the different device families and their NAC use cases, it is important to understand the flow of control plane traffic required for the NAC system to operate. In this context, control plane traffic refers specifically to RADIUS communication.

Although Meraki devices establish a management tunnel with the cloud, they always communicate with RADIUS servers over the local

network. This traffic never traverses the cloud. The only exception occurs when you use cloud RADIUS servers integrated into the Meraki dashboard, such as Access Manager or Meraki Authentication. When designing and operating the network, you should always consider how RADIUS traffic flows between the network access device (NAD), which, in the Meraki environment, may be a switch (MS or CS), a wireless access point (MR or CW), or a security appliance (MX). Figure 11-4 provides a summarized representation of RADIUS flows in different scenarios.

Technet24

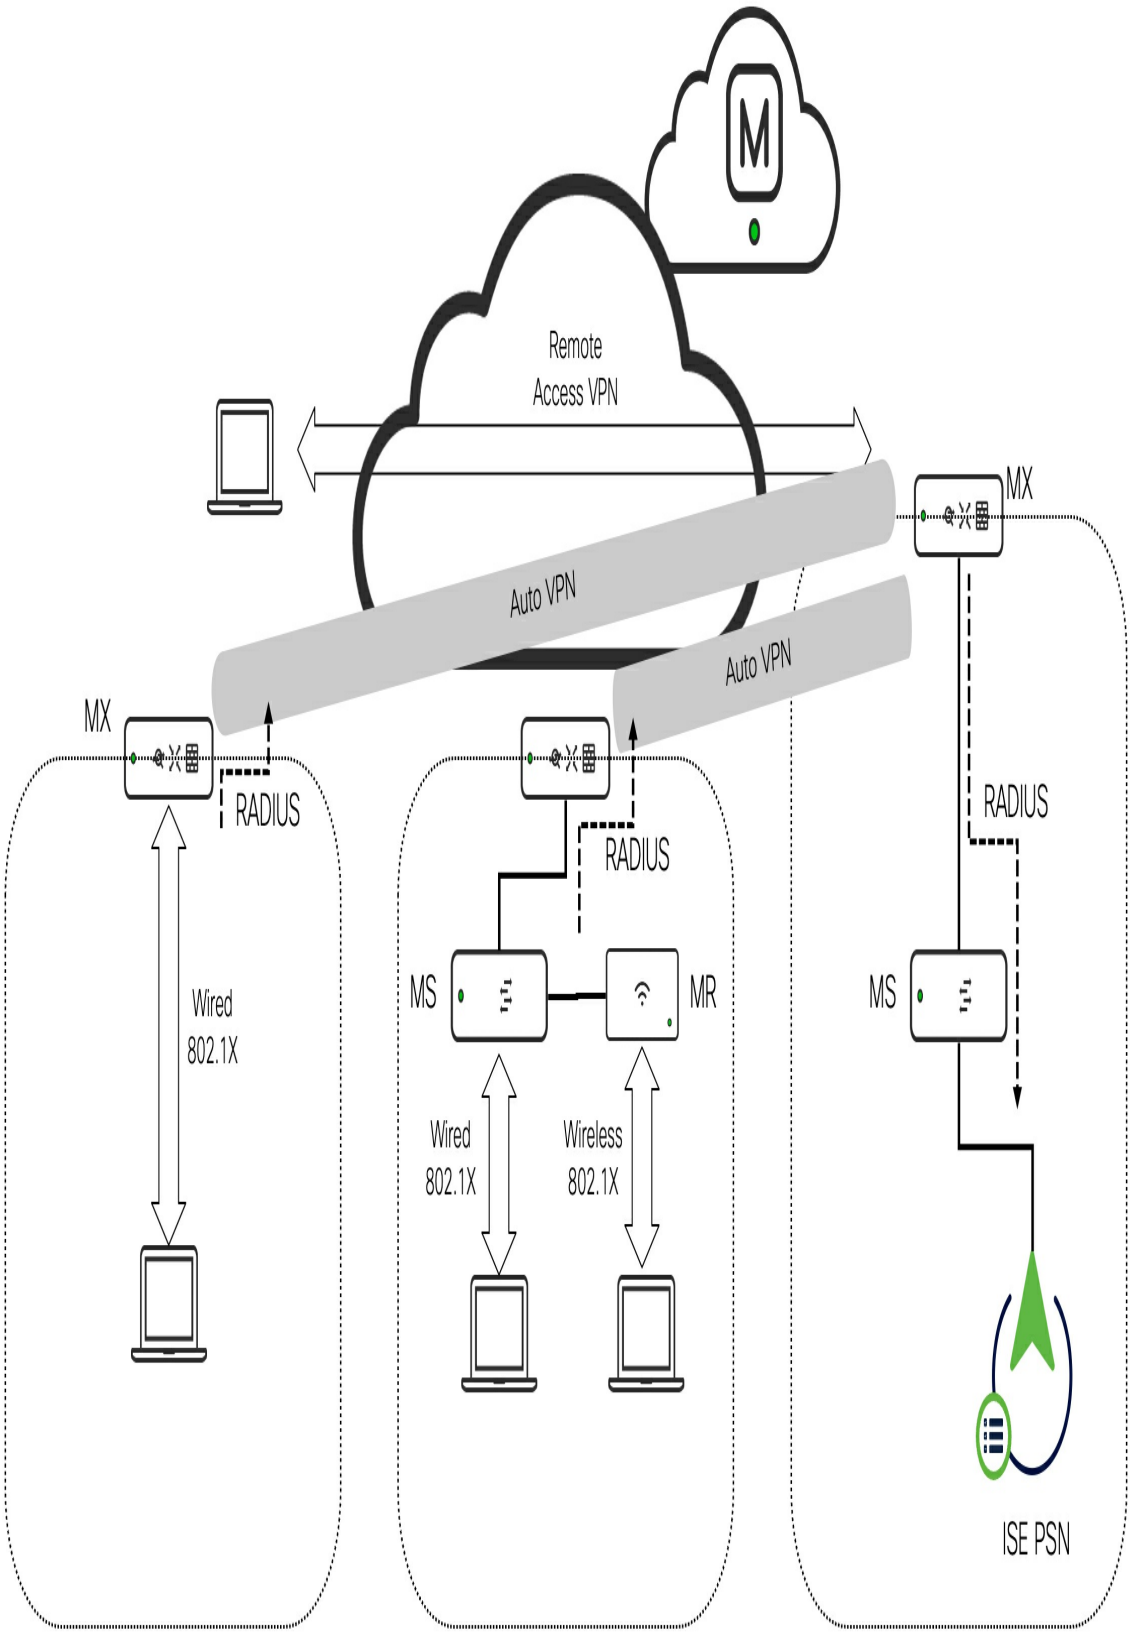

M
Remote
Access VPN
Auto VPN
Auto VPN
MX
MX
RADIUS
RADIUS
RADIUS
Wired
802.1X
MS
MR
MS
Wired
802.1X
Wireless
802.1X
ISE PSN

Figure 11-4 *RADIUS Path in Different Scenarios*

Now that we have established the general pattern of RADIUS traffic flow in a Meraki environment, it is important to learn how the source IP address is selected for RADIUS packets. Understanding this process is crucial because Cisco ISE uses the source IP address in each RADIUS packet to identify the network access device. Figure 11-5 illustrates what an ISE administrator may observe when the source IP address of RADIUS traffic is not properly configured.

Technet24

Live Logs | Live Sessions

Never

Reset Repeat Counts | Export To

Time	Status	Details	Repea...	Failure Reason
X				11007
Nov 10, 2025 09:53:52.224 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:51:17.861 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:51:12.859 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:51:07.859 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:51:02.855 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:50:57.855 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:50:52.851 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:50:47.850 PM				11007 Could not locate Network Device or AAA Clien
Nov 10, 2025 09:50:42.847 PM				11007 Could not locate Network Device or AAA Clien

Figure 11-5 *List of Failed Authentication Attempts*

You can investigate a detailed authentication report to identify the source IP address used by the device. To do this, click the magnifying glass in the **Details** column of the authentication attempt you want to verify. In the report, scroll down to the **Device IP Address** field, as shown in Figure 11-6. This field displays the actual source IP address from the RADIUS packet.

Technet24

NAS IPv4 Address	10.110.255.1
NAS IPv6 Address	
NAS Port Id	
NAS Port Type	Ethernet
Authorization Profile	
Posture Status	
Security Group	

Other Attributes

DestinationPort	1812
NAS-Port	3
Called-Station-ID	A8-46-9D-79-F5-33:
Framed-MTU	1400
Device IP Address	10.110.28.1

Figure 11-6 *Detailed Authentication Report*

Figure 11-6 also demonstrates that the RADIUS attribute NAS IPv4 Address does not always match the actual IP address from which the packet arrives, and ISE always uses the source IP address to locate the network access device.

How is the source IP address selected for RADIUS packets on different platforms? The common rule is that RADIUS uses the IP address of the device's management interface. However, this process is not always straightforward and requires a detailed review of each platform to understand the specific caveats.

Meraki MS and CS Switches

Classic MS switches and CS switches prior to the CS 17.15 release use a dedicated management interface, which also serves as the source IP address for RADIUS communication. You can view management interface IP addresses by navigating to the switches table of the relevant Meraki network in the dashboard. Go to **Switching > Switches**. The management IP addresses appear in the **Local IP** column, as demonstrated in Figure 11-7.

Technet24


Meraki
Search
MSP Portal
Switches
Last day
Add a switch
Organization
Ikarem MARS Demo
0 Offline
0 Alerting
1 Online
0 Dormant
Search
Filters 1 result
Download
Network
Branch-Greenfield-POL
Status
Name
MAC address
Connectivity (UTC+1)
Cloud ID
Local IP
Firmware vers
Branch-Access
bc:b1:d3:f7:18:2b
Q4CC-EBRM-FKHE
10.110.255.5
MS 18.1.2
Secure Connect
Network-wide
Assurance

Figure 11-7 *Switches Table*

The management interface of the Classic MS switch and CS switch prior to the CS 17.15 release supports both static IP configuration and IP address assignment via DHCP. You can find the management interface settings on the left side of the Device Summary tab. Figure 11-8 provides an example.

Technet24

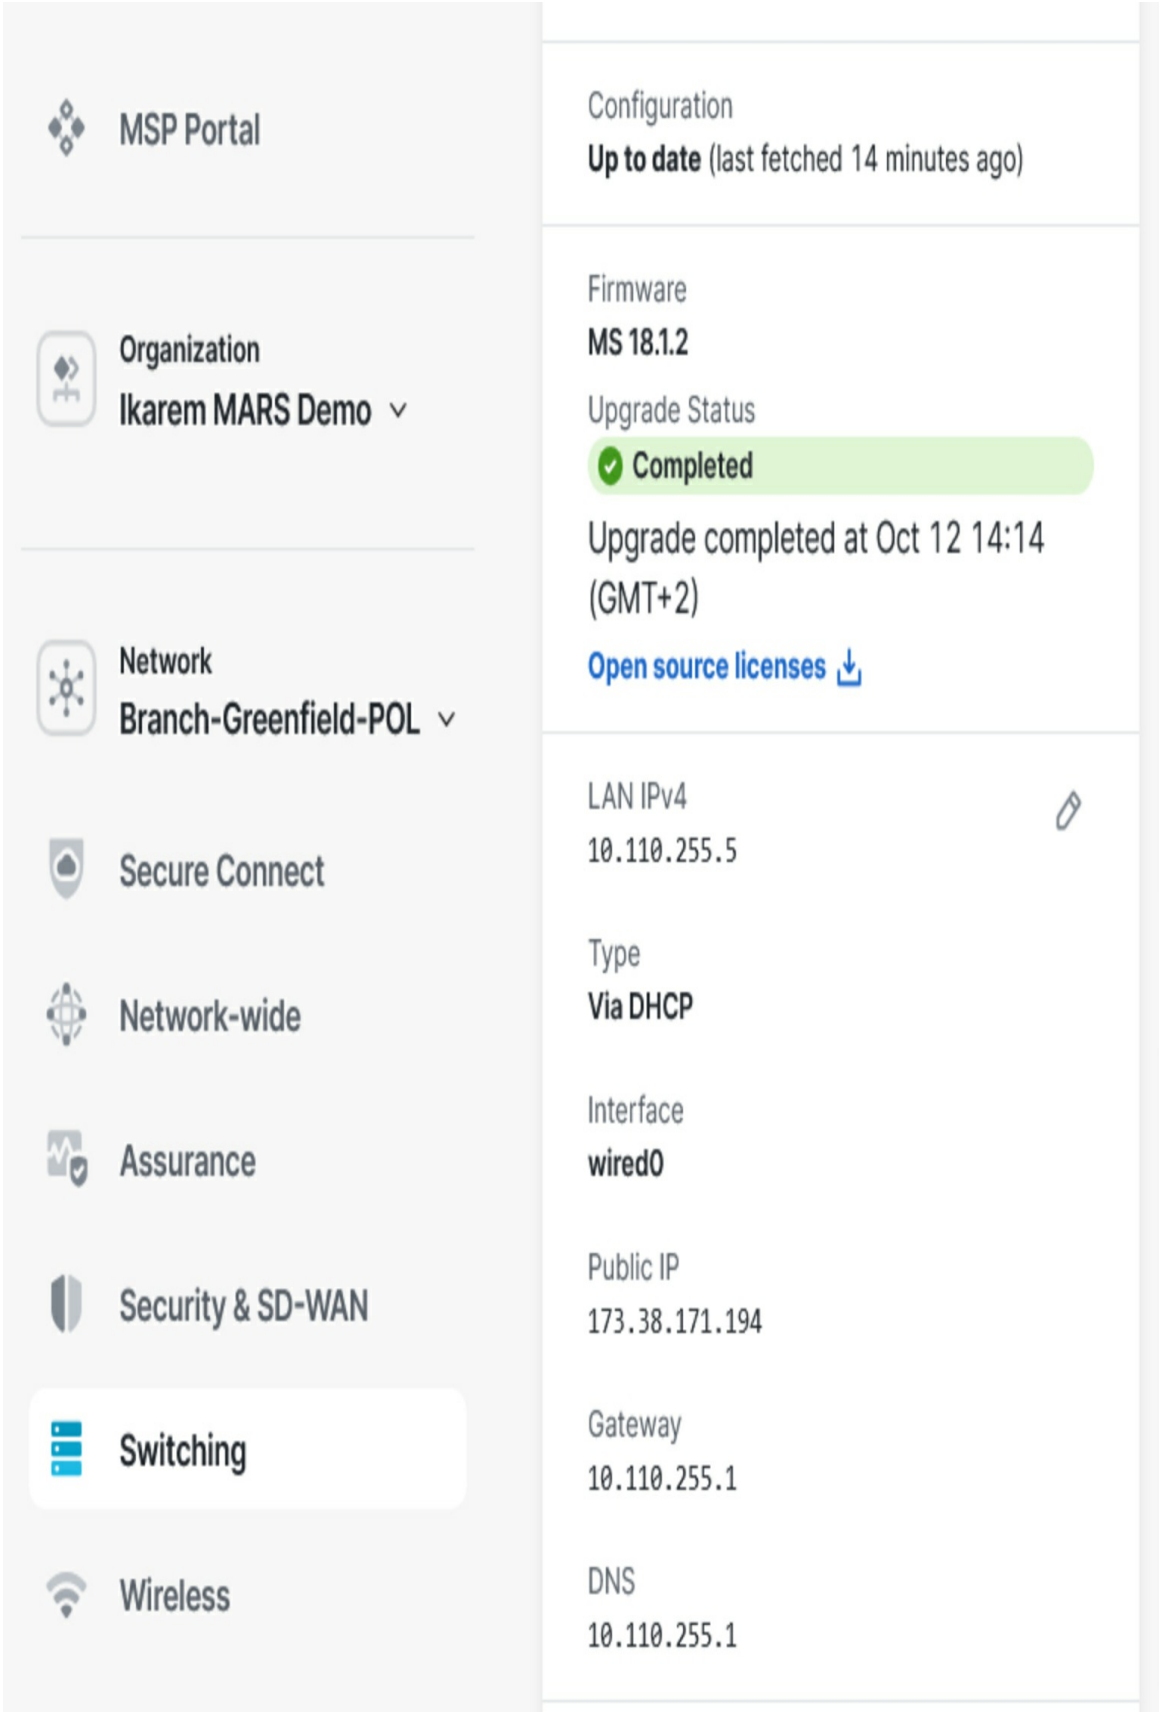

MSP Portal
Organization
Ikarem MARS Demo
Network
Branch-Greenfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Configuration
Up to date (last fetched 14 minutes ago)
Firmware
MS 18.1.2
Upgrade Status
Completed
Upgrade completed at Oct 12 14:14 (GMT+2)
Open source licenses
LAN IPv4
10.110.255.5
Type
Via DHCP
Interface
wired0
Public IP
173.38.171.194
Gateway
10.110.255.1
DNS
10.110.255.1

Figure 11-8 *Management Interface Settings*

The configuration section for the management interface is labeled **LAN IPv4**. You can change the management interface settings by clicking the pencil icon.

For CS switches running firmware 17.15 and later, you can no longer configure the management interface from the Summary tab. This change is a result of the introduction of cloud-native IOS XE, beginning with IOS XE release 17.15. In this context, *cloud-native* means that the underlying IOS XE operating system used by Catalyst switches managed through the Meraki dashboard natively supports connectivity with the Meraki cloud. This update introduces a new approach to management interfaces. With cloud-native IOS XE, the device does not use a dedicated management interface. Instead, the system selects one of the switch virtual interfaces (SVIs) as the uplink to the Meraki cloud.

The new approach that cloud-native CS switches use to determine which SVI serves as the path to the Meraki cloud relies on two features: Uplink Auto Config (UAC) and UAC uplink selection. You can still assign static IPv4 and IPv6 addresses and designate an SVI as a "preferred" uplink to ensure predictable flow for management plane traffic such as RADIUS.

To assign a static IP address to the SVI, you need to navigate to **Switching > Routing & DHCP**. There, you can create a new SVI by clicking **Add**, or you can edit an existing SVI by selecting it from the SVI list. When you reach the configuration page for the individual SVI, you can define it as the primary uplink by selecting the **Preferred Uplink** checkbox, as demonstrated in Figure 11-9.

Technet24

Interface Editor

Switch or switch stack HQ-Access-1

Name Uplink Vlan 51

VLAN 51

IP Toggle IPv4 Only | IPv6 Only | Both

Choose the Both tab to configure both IPv4 and IPv6

IPv4 management connectivity

Preferred uplink

Figure 11-9 *Preferred Uplink Checkbox*

At the same time both UAC and UAC uplink selection are enabled by default, and you cannot deactivate them. Therefore, you should understand the potential impact of these features.

UAC identifies active VLANs in the network by passively monitoring data plane traffic and measuring the number of packets observed in each VLAN. The switch detects the presence of individual VLANs by examining the 802.1Q tags of packets passing through. When any VLAN exhibits activity of at least 60 packets within a two-minute interval, UAC considers it active. UAC then attempts to obtain an IP address in this VLAN using DHCP unless you have defined a static IP address. If DHCP succeeds and the default gateway information is available, the switch performs a set of tests — an ARP probe of the gateway, a ping to the gateway, and a ping to catalyst.meraki.com. When these tests succeed, the switch designates this uplink as the primary.

UAC uplink selection continuously monitors the primary uplink using the same probes. If connectivity issues occur, the switch can promote another SVI to primary uplink status. This process plays a crucial role in the context of the RADIUS flow. By default, when the uplink that you configured as primary recovers, the switch does not automatically revert to it.

You, as the administrator, have two options for failing over to the manually configured primary uplink:

1. Enable the **Automatically Fallback to Preferred Uplink When Available** option by navigating to **Switching > Switch Settings**, as shown in Figure 11-10.

Technet24

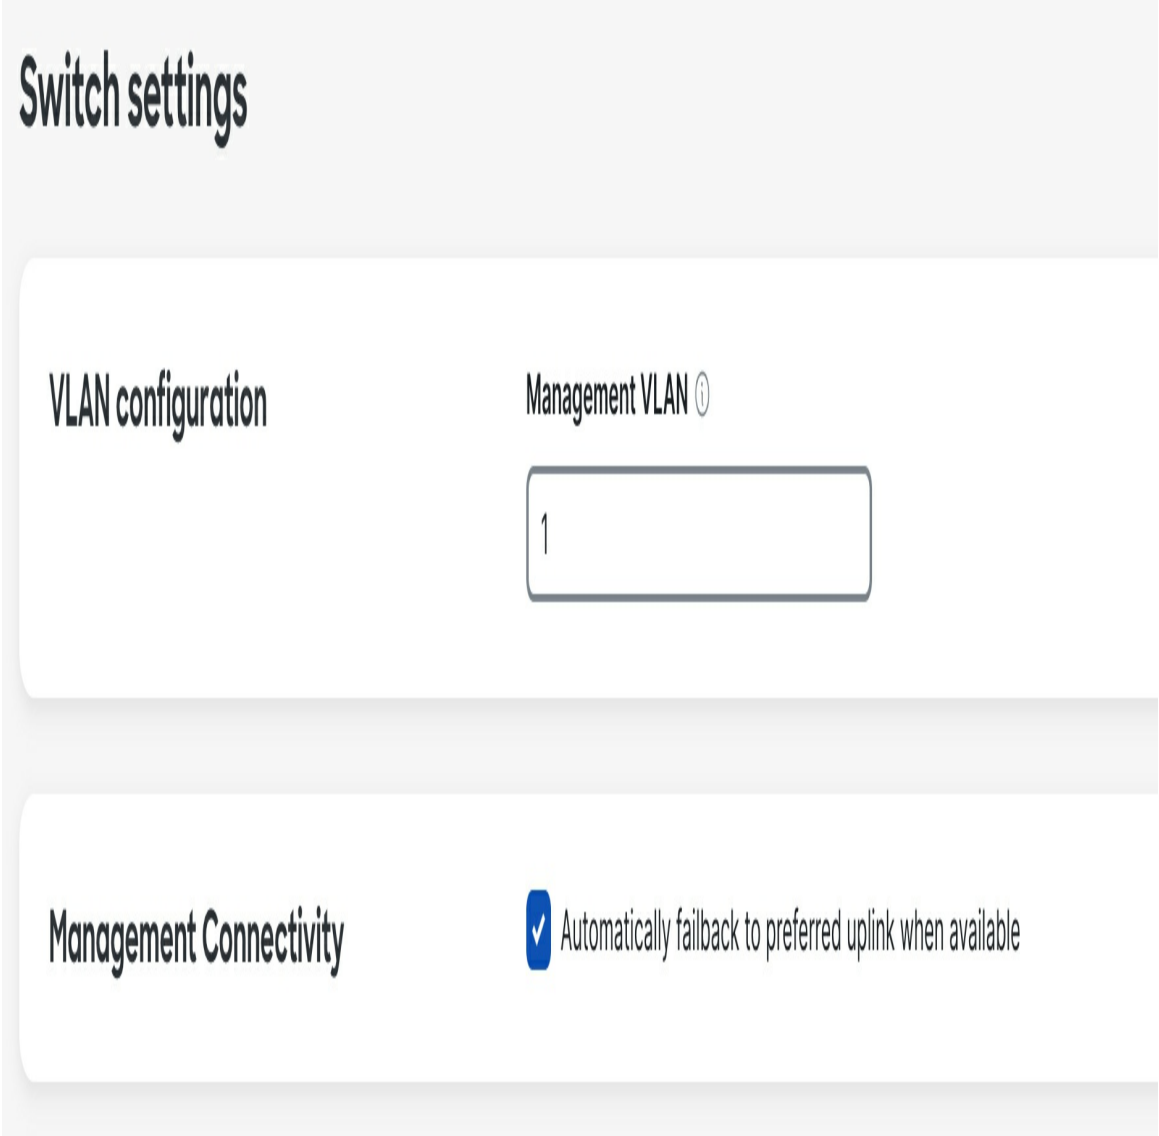
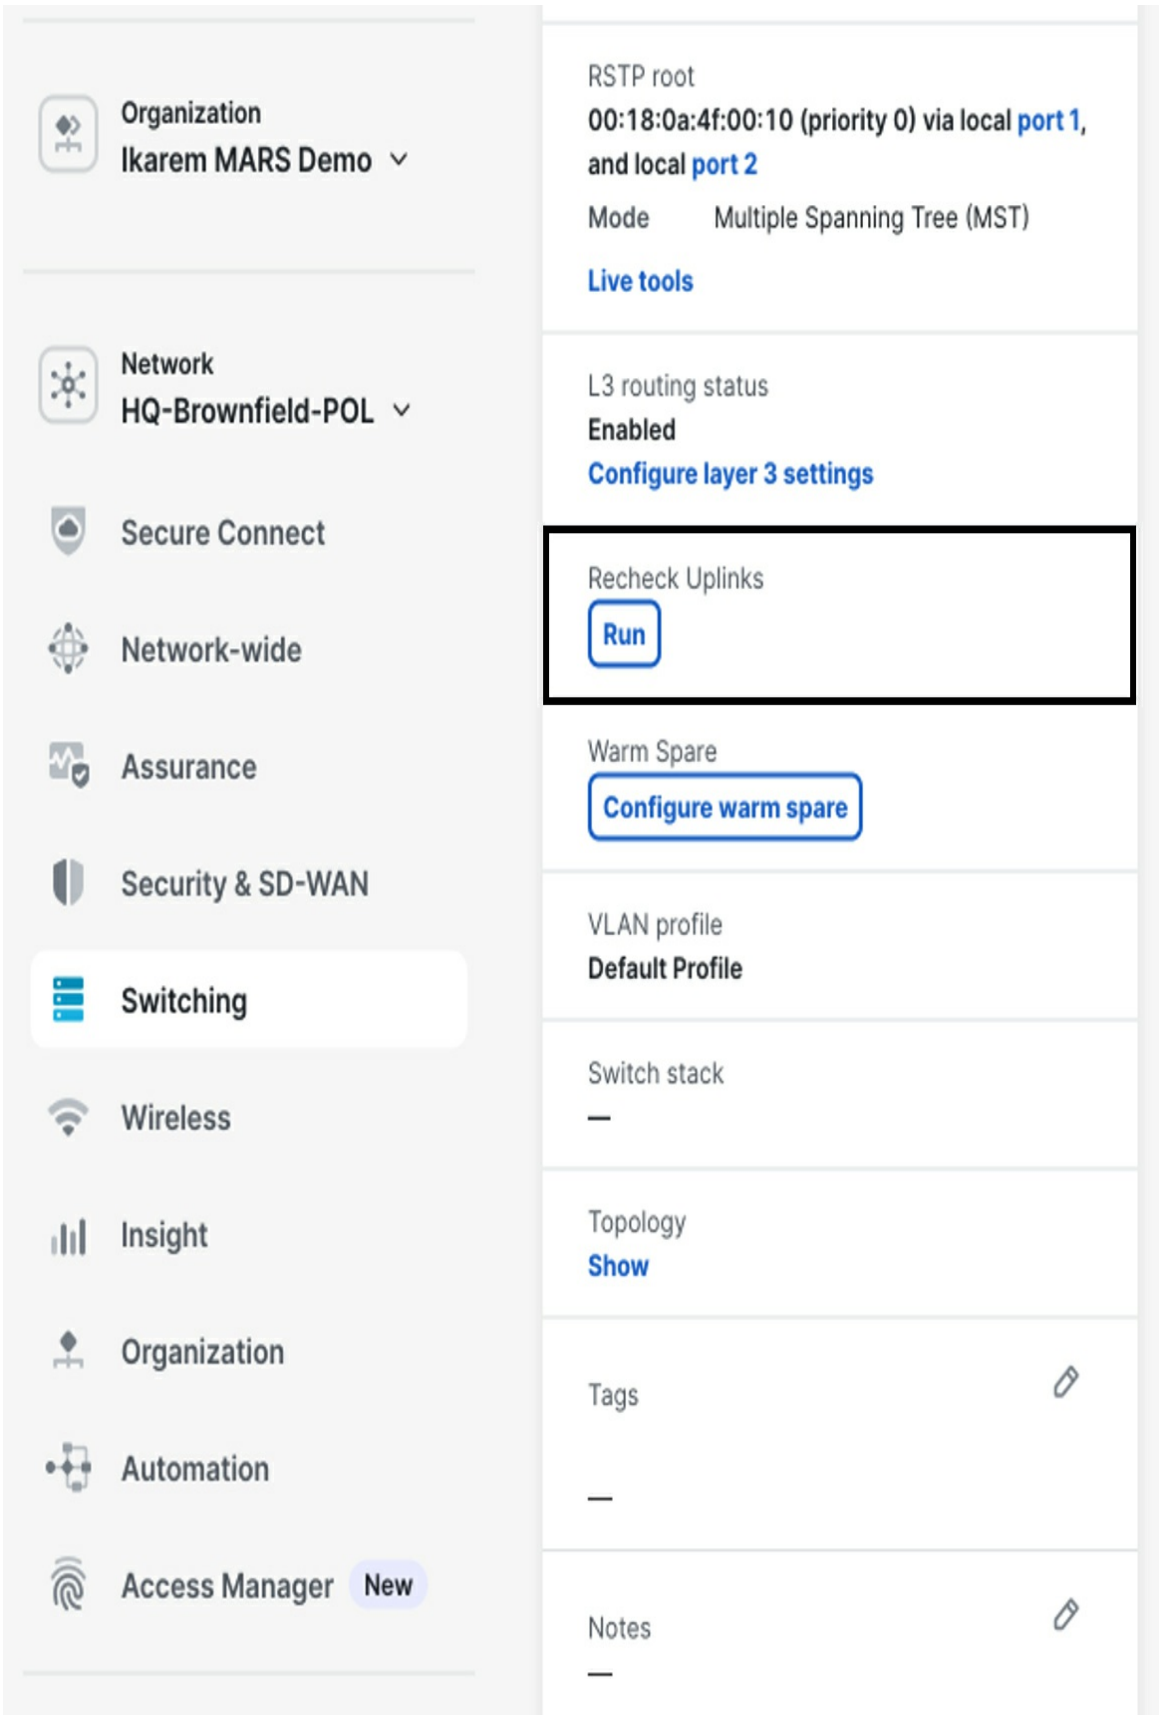


Figure 11-10 *Automatic Fallback to Preferred Unlink Setting*

2. If you do not want to enable automatic fallback, you can use the **Recheck Uplinks** feature as an alternative. Use this feature after connectivity has been restored on the preferred primary uplink. You can find the button for this feature at the bottom of the **Switch Summary** page, as demonstrated in Figure 11-11.


Organization
Ikarem MARS Demo
Network
HQ-Brownfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Automation
Access Manager
New
RSTP root
00:18:0a:4f:00:10 (priority 0) via local port 1, and local port 2
Mode Multiple Spanning Tree (MST)
Live tools
L3 routing status
Enabled
Configure layer 3 settings
Recheck Uplinks
Run
Warm Spare
Configure warm spare
VLAN profile
Default Profile
Switch stack
—
Topology
Show
Tags
—
Notes
—

Technet24

Figure 11-11 *Recheck Uplinks Feature*

Note

The **Automatically Fallback to Preferred Uplink When Available** and **Recheck Uplinks** features first became available in CS 17.15.1.

You can also verify the status and statistics for UAC and UAC uplink selection through a terminal session on the switch. To access the terminal, navigate to **Switching > Switches**, select the specific switch, and open the **Tools** tab. There, you can launch the terminal and execute the command **show uac uplink db**, as shown in Figure 11-12.

Terminal Opens a troubleshooting console to this device, similar to an IOS-XE CLI window.

Close Terminal | Detach | Detach (full)

This terminal allows limited access to read-only IOS-XE commands, such as **show**.

```
HQ-Access-1#show uac uplink db
Uplink Autoconfig: Enable
IPV4 Uplink: Vlan 51
      Ping Pass Count: 1247
      GW ARP Pass Count:2
IPV6 Uplink: None
      Ping Pass Count: 0
      GW ARP Pass Count:0
IfName    Score  State  IPAddress       Subnet/Prefix  ARPFail  PingFail   GWProbe     Ping ReScore
Vlan10        4      7  10.100.10.21    255.255.255.0      2        0         0        0     0
Vlan11        4      7  10.100.11.21    255.255.255.0      2        0         0        0     0
Vlan12        4      7  10.100.12.21    255.255.255.0      2        0         0        0     0
Vlan51        8     15  10.100.254.21   255.255.255.0      0        0      3227A    8E92373     0
Vlan99        4      9  10.100.101.21   255.255.255.0      0        0         0        0     0
IfName    Score  State  IPv6Address           Prefix  ARPFail  PingFail   GWProbe     Ping ReScore
HQ-Access-1#
```

Technet24

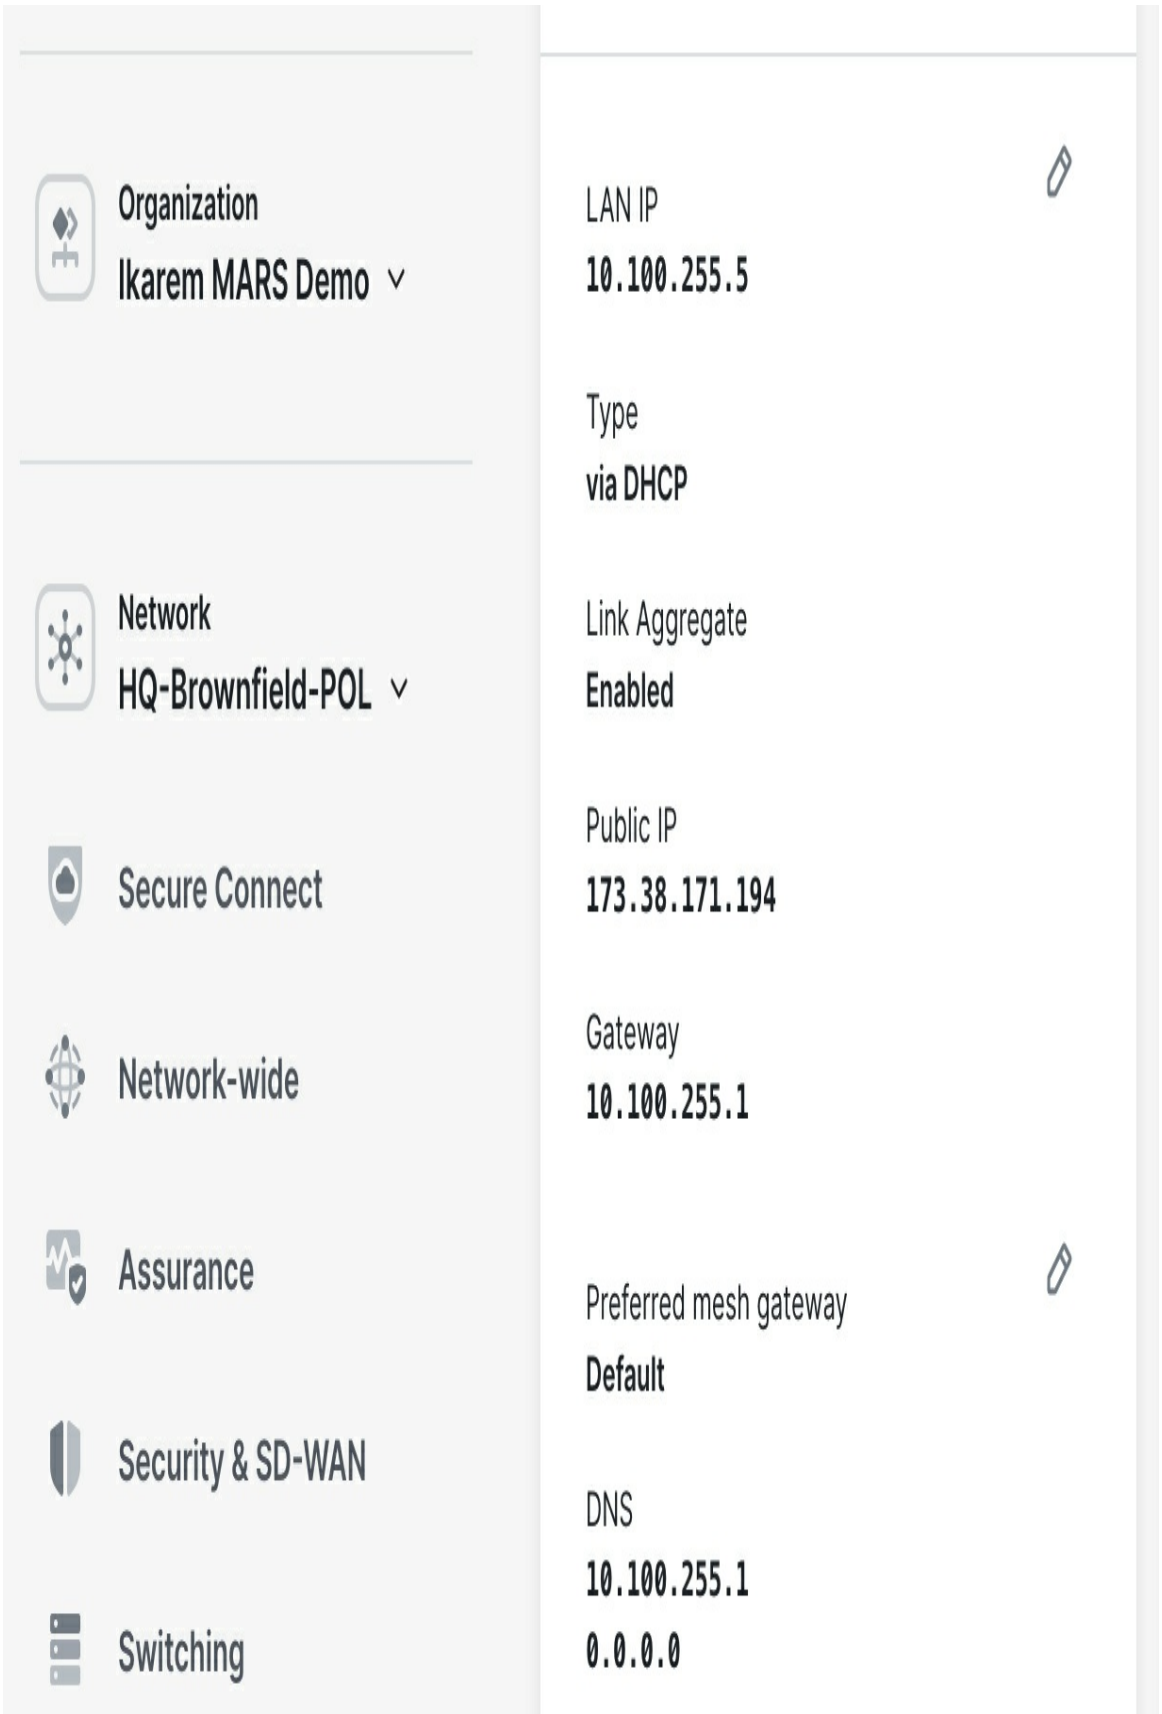
Figure 11-12 *Output of* ***show uac uplink db*** *Command*

Note

It is recommended to add all potential IP addresses to the network device definition on Cisco ISE.

Meraki MR and CW Access Points

In the case of access points, the situation resembles that of classic MS switches. Each access point includes a dedicated management interface, which you can configure on the device's **Summary** page, as shown in Figure 11-13.


Organization
Ikarem MARS Demo
Network
HQ-Brownfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
LAN IP
10.100.255.5
Type
via DHCP
Link Aggregate
Enabled
Public IP
173.38.171.194
Gateway
10.100.255.1
Preferred mesh gateway
Default
DNS
10.100.255.1
0.0.0.0

Technet24

Figure 11-13 *Access Point Management Interface Settings*

A common pattern with Meraki access points is that administrators often place all access points in each location into the same management VLAN, which is then added as a NAD IP. For example, Figure 11-14 illustrates how the entire management subnet 10.100.255.0/24 is configured for the MyAPs NAD.

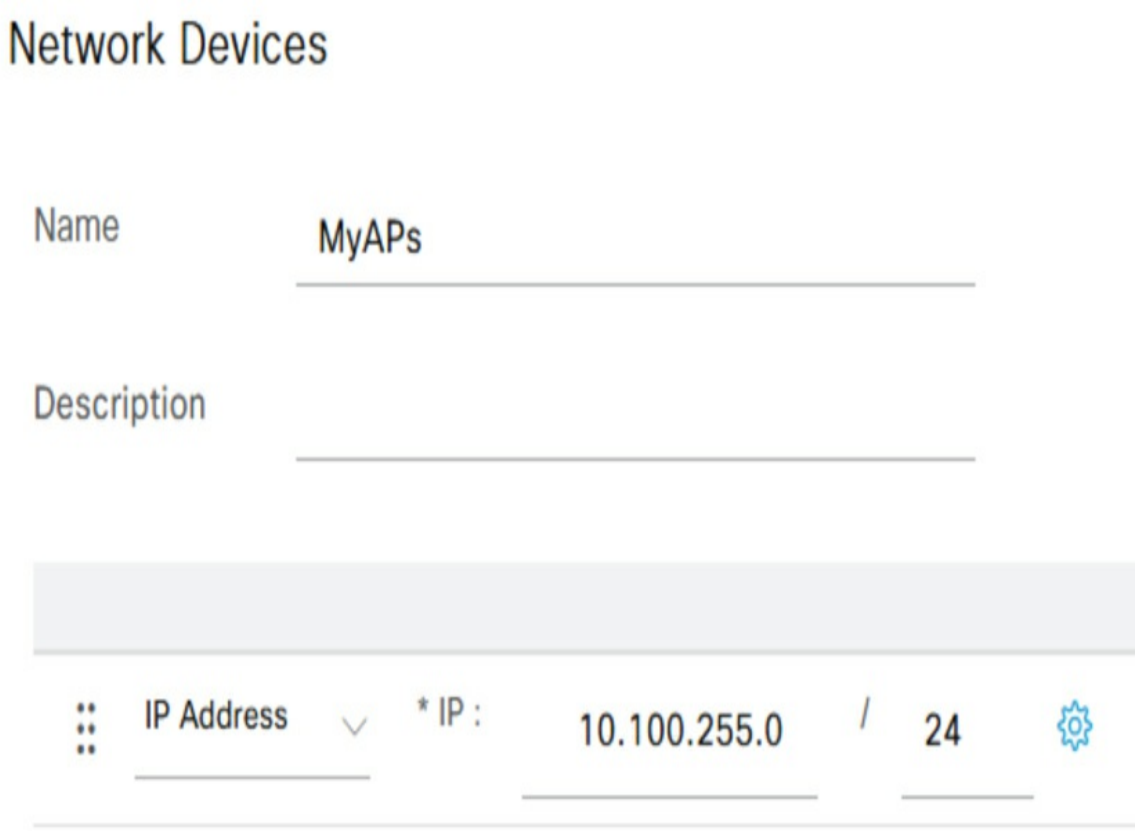


Figure 11-14 *NAD with Subnet Defined as Device IP*

Although this configuration simplifies initial setup, it can create challenges during troubleshooting. Specifically, it may become difficult to quickly identify the correct access point by using ISE live logs, as demonstrated in Figure 11-15.

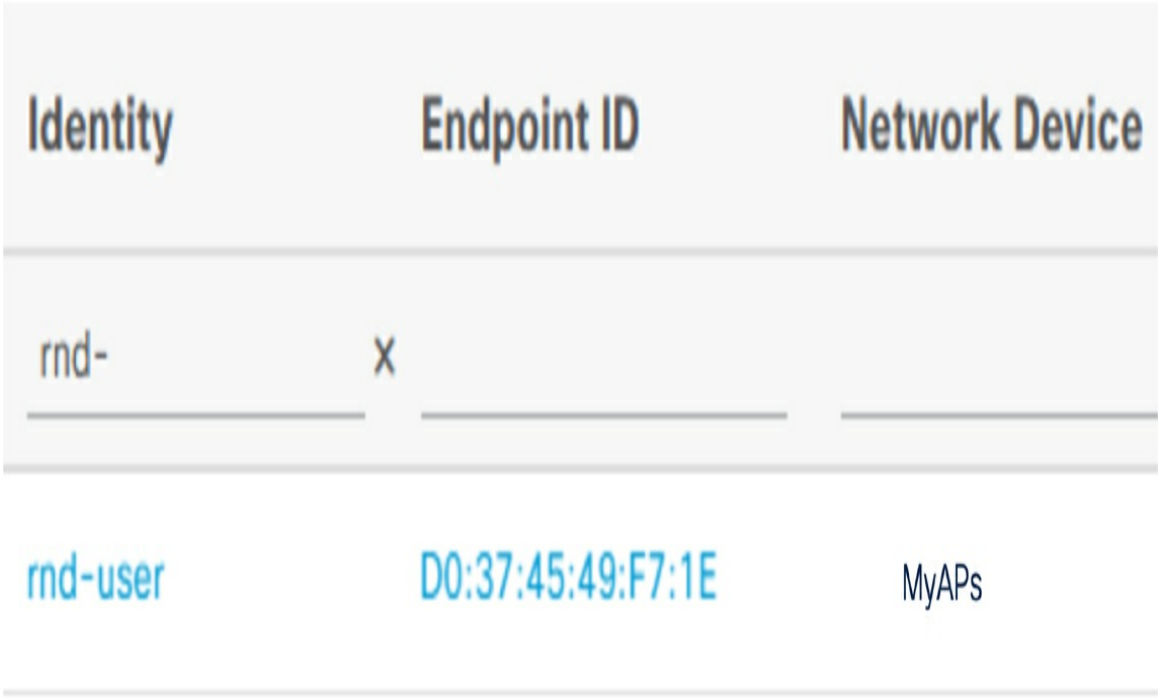


Figure 11-15 *NAD Live Logs Example*

It is recommended that you add each access point as an individual NAD in ISE for this reason. When devices obtain their management IP addresses through DHCP, you can configure DHCP servers to always assign the same IP address to each MAC address. This setup ensures predictability of RADIUS traffic.

MX Appliances and Teleworker Gateways

MX appliances and teleworker gateways (Z-series appliances) running MX 19.1.6 or later use the IP address of the interface with the highest VLAN ID added to the VPN as the source IP for RADIUS packets. You can review the IP settings of the MX by navigating to **Security & SD-WAN > Addressing & VLANs**, as shown in Figure 11-16. In this example, the MX uses the IP address of VLAN 100 as the source for RADIUS packets because VLAN 100 has the highest VLAN ID and is also included in the Meraki VPN.

Technet24

Routing

LAN setting

VLANs | Single LAN

Subnets

Search by VLAN name,

Delete | Add VLAN

ID ▲	VLAN name	Version	Config	VLAN interface IP	Uplink	Group policy	VPN mode	Adaptive policy
1	Default	4	Manual	10.110.128.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
10	Mng Greenfield	4	Manual	10.110.255.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
23	Clients	4	Manual	10.110.23.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
24	Servers	4	Manual	10.110.24.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
25	Guest	4	Manual	10.110.25.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
26	Auth-Failed	4	Manual	10.110.26.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
27	Critical-Data	4	Manual	10.110.27.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
28	Critical-Voice	4	Manual	10.110.28.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			
100	Core-Interconnect	4	Manual	10.110.100.1/24	Any	None	Enabled	None
		6	Disabled	--	Any			

Figure 11-16 *MX VLANs List*

When MX appliances or teleworker gateways do not have Meraki VPN enabled, or when no interfaces are added to the VPN, the source IP address in RADIUS packets falls within the 6.x.x.x range. The appliance uses this internal IP address for specific management plane and control plane operations between the device and the dashboard. Although you can theoretically add this IP address to the NAD list, practical implementation is more complex. To ensure that RADIUS functions correctly, you must verify that packets with a 6.x.x.x source IP can be routed through all transit devices, which is not always feasible. Figure 11-17 shows a failed authentication attempt received from an MX appliance with no networks added to the Meraki VPN.

Technet24

Other Attributes

DestinationPort	1812
NAS-Port	3
Called-Station-ID	A8-46-9D-79-F5-33:
Framed-MTU	1400
Device IP Address	6.13.184.161
ConfigVersionId	151
Protocol	Radius
Device Port	33722
DTLSSupport	Unknown
AcsSessionID	ise-1/551983081/7934
Acct-Session-Id	620B8BFFDC477160

Figure 11-17 *MX RADIUS Request Sourced from 6.x.x.x IP*

Methods of Access Control for Wired Connections

This section explains various NAC features based on the type of network connection. We've used this approach because, although some NAC features may appear identical from the ISE perspective, their implementation can differ significantly in the Meraki dashboard depending on the connection type.

To better understand this concept, consider how different connection types vary in their implementation of NAC:

- **Wireless authentication:** With wireless authentication, the outcome for network access is always binary. When authentication succeeds, the endpoint receives network access with the appropriate authorization/segmentation settings. If authentication fails, the endpoint receives no network access. The binary nature of authentication for wireless connections also applies to the authentication type. With wireless networks, you must define in advance whether a given SSID uses MAC Authentication Bypass or 802.1X. In contrast, a wired port can support both authentication methods concurrently.
- **Wired authentication:** Wired authentication is significantly more complex than wireless or VPN authentication. Multiple authentication modes, methods, and authorization mechanisms exist, as well as various ways to apply authorization results.

Due to the differences between wireless and wired authentication, it is logical to begin with the methods of access control used for wired clients. This approach allows for coverage of concepts that are later applied to wireless and VPN connections.

Technet24

Main Dashboard Tools for Wired NAC

In the Meraki dashboard, the main elements required for wired NAC configuration are distributed across different locations. The following list outlines these locations and the purpose of the features configurable there. Throughout this chapter, we will revisit each location to review the available features in greater detail.

- **Organization-wide RADIUS servers (Organization > Settings):** This configuration element allows you, as administrator, to define RADIUS servers for the entire organization and later reuse them at the network level for wired NAC. At the time of writing, organization-wide RADIUS servers are only supported for wired NAC.
- **Network-level group policies (Network-wide > Group Policies):** Group policy is a unique Meraki feature that enables you, as administrator, to assign specific network-level access restrictions to endpoints and users. You can apply group policies directly by attaching them to specific endpoints or VLANs, or indirectly by assigning a group policy based on the value of the Filter-ID attribute returned by the RADIUS server during the authentication and authorization process.

  ### Note

  RADIUS Filter-ID is the default attribute used to carry the group policy name. However, this can be modified in certain cases.

- **Network-level VLAN profiles (Network-wide > VLAN Profiles):** This feature allows the use of named VLANs assigned by the RADIUS server. By default, Meraki access points and switches assign numbered VLANs based on information from the

RADIUS server, but in some scenarios, it may be beneficial to use named VLANs instead.

- **Network-level access policies (Switching > Access Policies):** An access policy serves as the primary configuration container that consolidates most of the settings required at the switch port level for NAC functionality. The access policy includes parameters such as authentication type (open or closed), RADIUS settings, and network-level permissions for various scenarios. The following sections will focus primarily on access policy features and their configuration.
- **Organization-level Adaptive Policy settings (Organization > Adaptive Policy):** This section serves as the primary location for Adaptive Policy configuration. Meraki Adaptive Policy provides a specialized implementation of Cisco TrustSec for Meraki network environments. While some minor settings must be configured at the individual device or network level, the majority of the configuration occurs at the organization level.

Organization-wide RADIUS Servers

To define organization-wide RADIUS servers, navigate to **Organization > Settings** and scroll to the **RADIUS Servers** section. Figure 11-18 shows this section with a preconfigured RADIUS server. As an administrator, you can add new RADIUS servers to the list by clicking **Add a RADIUS Server** (1) or edit/delete existing servers (2) by clicking the three dots in the **Actions** section.

Technet24

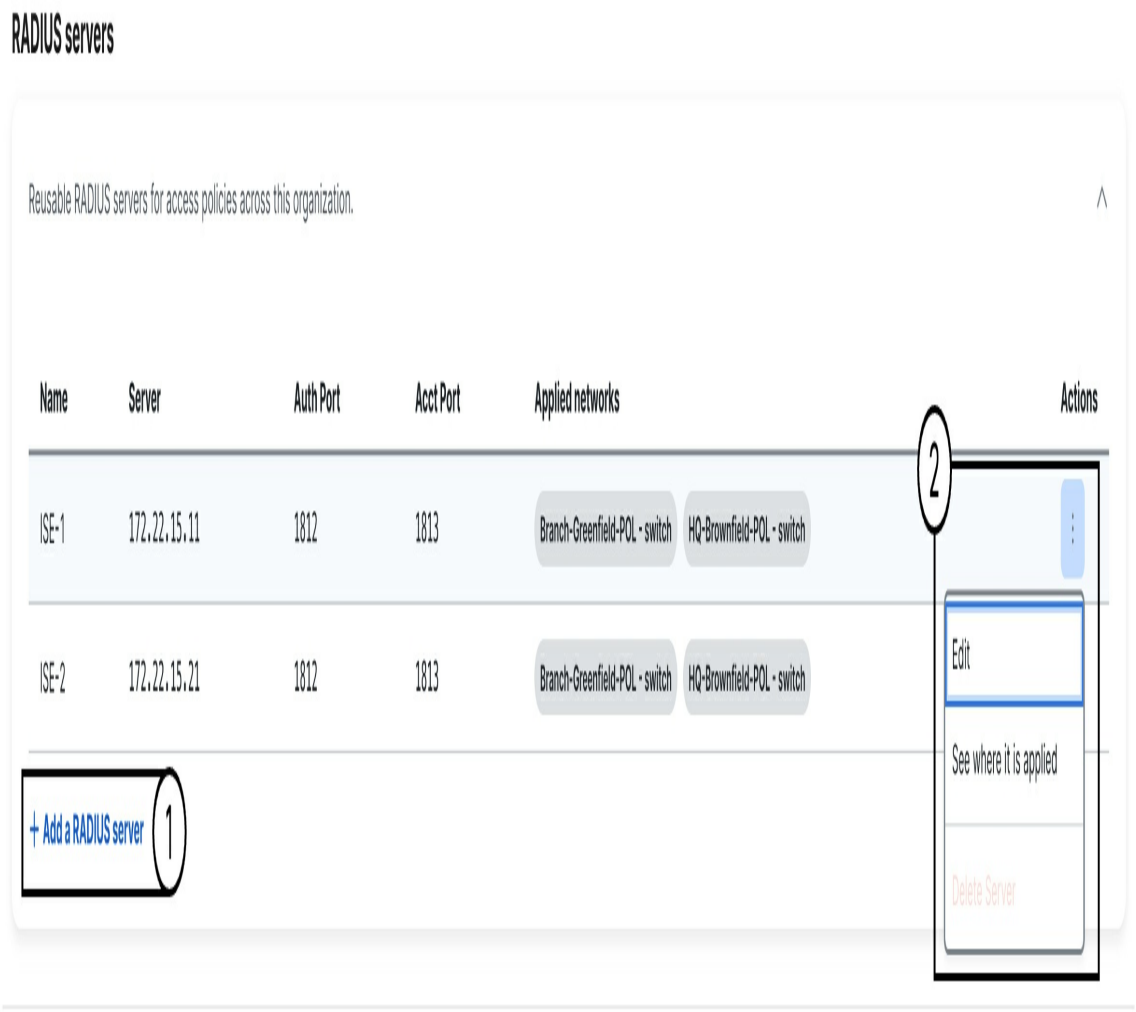


Figure 11-18 *Organization-wide RADIUS Servers*

Note

The delete option is not available if the server is already in use at the network or SSID level. To delete such a server, you must first remove it from the access policy or SSID. The **See Where It Is Used** option can help you locate the server in question.

To configure a new server, define the settings as outlined here. Figure

11-19 provides an example of these settings:

- **Address:** Specify the IPv4, IPv6 address, or FQDN.
- **Name:** Enter the RADIUS server name to be used in the dashboard.
- **Server Type:** Enable authentication, accounting service, or both. When Cisco ISE acts as a RADIUS server, enable both authentication and accounting.
- **Authentication Port/Accounting Port:** Specify the UDP port for the service. The default authentication port is 1812, and the default accounting port is 1813.
- **Secret:** Specify the secret to use for the connection with the RADIUS server.

Technet24

Add RADIUS server

Address*

172.22.15.31

Name

ISE-3

Server type

☑ Authentication ☑ Accounting

Authentication port*

1812

Accounting port*

1813

Secret*

•••••••••• Show

Cancel Save

Figure 11-19 *Organization-wide RADIUS Servers*

Note

RADIUS server timeout and retry count are not configurable for Meraki MS and CS switches in the dashboard. Both switch families use a five-second RADIUS timeout and three retries, which should be suitable for most environments. If you require different values, contact Meraki customer support.

From the ISE perspective, you configure a Meraki network access device as you would any other device. For Meraki appliances, use **Cisco** as the device type. It is recommended to specify both the location and device type to provide greater flexibility during policy and policy set configuration.

Network-Level Group Policies

Group policies represent a Meraki network security concept that enables you to apply a wide variety of restrictions to endpoints and users connected to the network. A group policy serves as a configuration container with various security settings, such as firewall rules and traffic shaping. The settings that can be applied depend on the device type to which the user or endpoint is connected.

In the switching domain, group policy firewall rules function as group policy access control lists (GPACLs). Switches do not use other security settings from group policies. When creating policies to serve as GPACLs, you do not need to define settings beyond the required firewall rules. If any other settings are present in the group policy used as a GPACL, the switch ignores them. However, these settings can still be applied on Meraki devices that support them.

First, let's review the minimum hardware and software requirements

Technet24

for using GPACLs. Table 11-1 outlines these requirements.

Note

When not explicitly specified, the term *9K-M* refers to all Catalyst models managed by the Meraki dashboard.

Table 11-1 *GPACLs Support on Meraki Switches*

MS Switch Family	MS Switch Model	Minimum Firmware Required
MS100 Series	MS130	MS 18
	MS150	MS 18
MS200 series	MS210	MS 14.5
	MS225	MS 14.5
	MS250	MS 14.5
MS300 series	MS350	MS 14.5
	MS355	MS 14.5
	MS390	MS 15.8
MS400 series	MS410	MS 14.5
	MS425	MS 14.5
	MS450	MS 14.5
Catalyst 9K-M	C9300/L/X -M	CS 16+
Catalyst 9K-M	Cloud Native IOS XE	ISO XE 17.15.2+

Another important consideration is the platform limits related to GPACLs. You must carefully evaluate these limits during the planning phase and continuously review them during normal network operations. Improper planning or a lack of ongoing monitoring may

result in TCAM overruns, which can lead to silent failures and cause connectivity issues for any clients attached to an active group policy on the switch.

Table 11-2 summarizes the platform limits at the time of writing.

Table 11-2 *GPACLs Platform Limits*

Platforms	Limit Type	Limit Value
MS130 MS150 MS210 MS225 MS250 MS350 MS355	Active GPACLs per switch	20
	Total number of active rules with Layer-4 port ranges per switch	32
	Maximum number of active ACL entries	200
MS390/9K-M	Maximum number of ACL entries per ACL	1000
	Active GPACLs per switch	50
	Maximum number of active ACL entries	4900
C9200L-M	Active GPACLs per switch	50
	Maximum number of active ACL entries	1000

When planning a GPACL deployment, consult the current version of the Meraki MS group policy access control lists documentation to understand the latest limitations and recommendations for each platform.

The following are other important GPACL considerations:

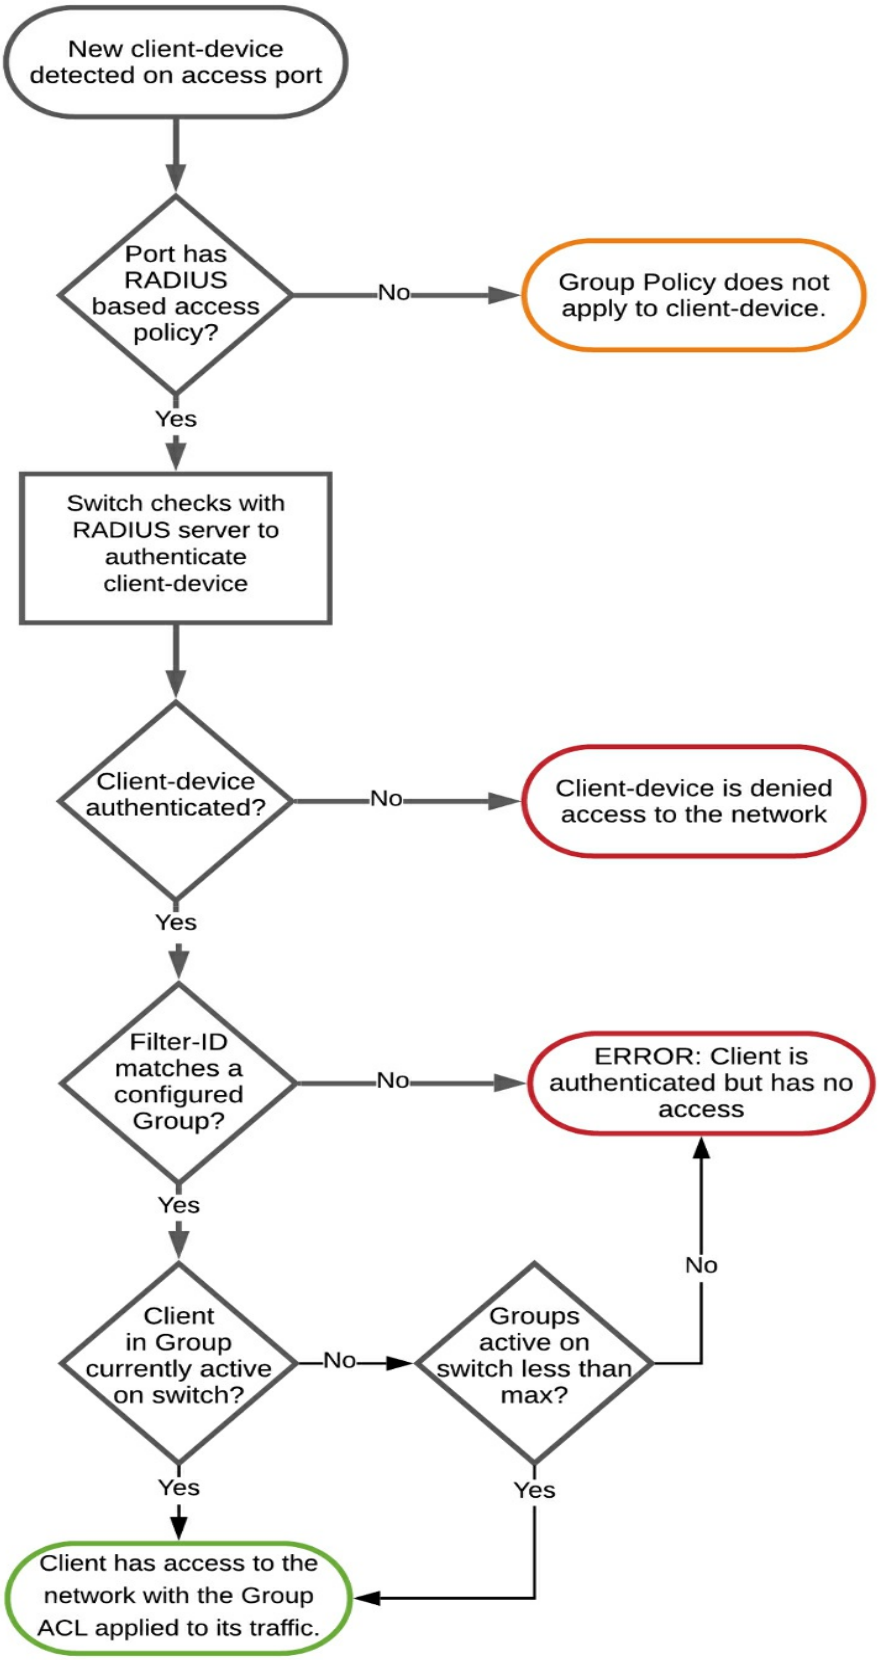
- Only IP or CIDR-based rules are supported. MS switches do not

Technet24

support groups containing rules that use FQDNs.

- MS switches apply GPACLs through client authentication access policies and require a RADIUS server. Static assignment of a group to a client for GPACL application is not possible on MS switches. **Note** Some exceptions to this rule exist with Open authentication mode. These exceptions will be explained in the following sections.
- Supported access policy host modes for GPACLs include single-host, multi-auth, and multi-domain. GPACLs cannot be applied to clients authenticated by an access policy using multi-host mode.
- MS switches cannot apply GPACLs to clients connecting on trunk ports.
- MS switches implement GPACLs as stateless access control entries.
- GPACL rules take precedence over switch ACL rules (configured in the **Switch > ACL** section), but only on the switch where the client has been authenticated.
- GPACL names on MS switches must begin with an alphanumeric character and may only include alphanumeric characters, underscores, or hyphens.

The flow chart in Figure 11-20 illustrates the GPACL assignment process.


New client-device detected on access port
Port has RADIUS based access policy?
No
Group Policy does not apply to client-device.
Yes
Switch checks with RADIUS server to authenticate client-device
Client-device authenticated?
No
Client-device is denied access to the network
Yes
Filter-ID matches a configured Group?
No
ERROR: Client is authenticated but has no access
Yes
Client in Group currently active on switch?
No
Groups active on switch less than max?
No
Yes
Yes
Client has access to the network with the Group ACL applied to its traffic.

Technet24

Figure 11-20 *GPACL Assignment Process*

You can configure group policies at **Network-wide > Group Policies**. Figure 11-21 shows the page with two group policies already configured. Use the **Add a Group** link to create a new policy. In the **Actions** column, you have the option to clone or delete an existing policy.



Figure 11-21 *Network-wide Group Policies Page*

To configure a new group policy, click the **Add a Group** link. Figure 11-22 provides a step-by-step example of group policy configuration.

Technet24

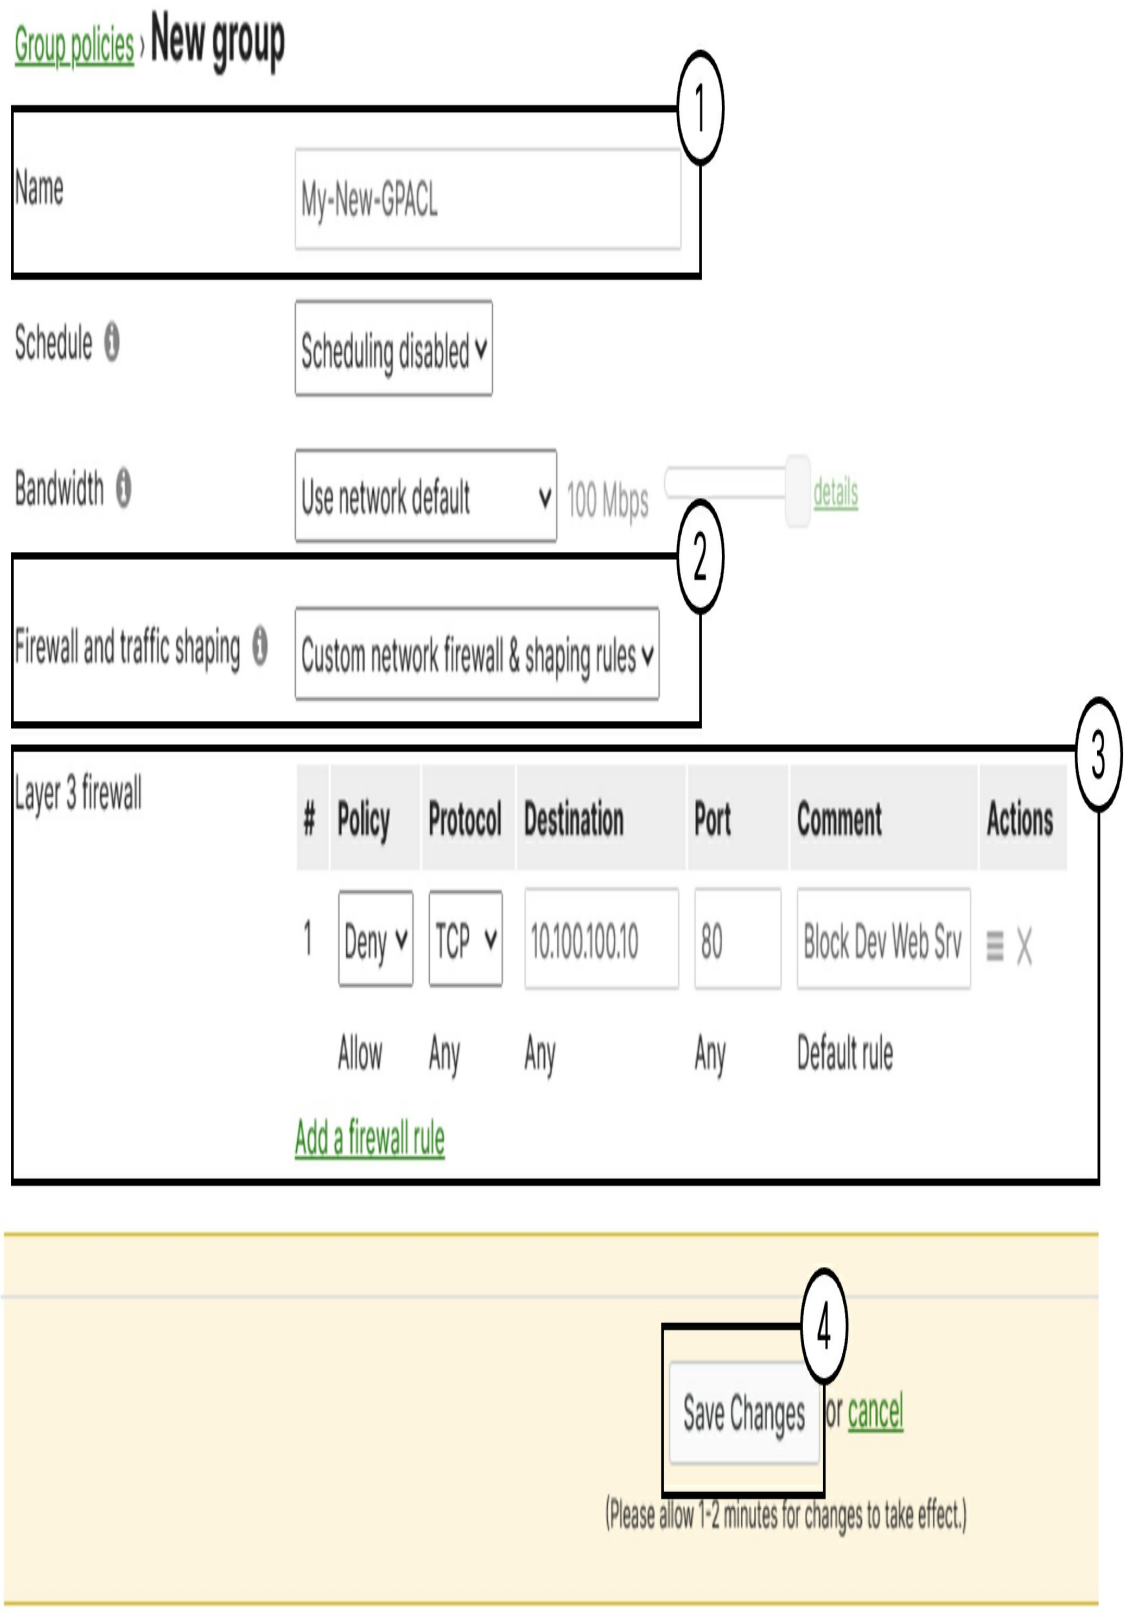

Group policies › New group
1
Name
My-New-GPACL
Schedule
Scheduling disabled
Bandwidth
Use network default
100 Mbps
details
2
Firewall and traffic shaping
Custom network firewall & shaping rules
3
Layer 3 firewall
Policy Protocol Destination Port Comment Actions
1 Deny TCP 10.100.100.10 80 Block Dev Web Srv
Allow Any Any Any Default rule
Add a firewall rule
4
Save Changes
or cancel
(Please allow 1-2 minutes for changes to take effect.)

Figure 11-22 *New Group Policy Configuration*

Follow these steps to configure a new group policy:

Step 1. Specify the group policy name.

Step 2. Select **Custom Network Firewall & Shaping Rules** from the **Firewall and Traffic Shaping** drop-down list.

Step 3. Add the required rules in the **Layer 3 Firewall** section using the **Add a Firewall Rule** link.

Step 4. Scroll down to the bottom of the page and save your changes.

Note

GPACL includes a default explicit **permit IP any** as a last rule.

VLAN Profiles

In the context of integration with ISE or other RADIUS solutions, VLAN profiles are useful because they allow assignment of VLAN names, rather than VLAN numbers, in the Access-Accept response for both MS and MR devices.

Using a name instead of a number is convenient because logical function names such as "Guests" typically remain consistent across the organization, whereas VLAN numbers for this role may vary from switch to switch.

Table 11-3 summarizes the hardware and software requirements for VLAN profile support.

Table 11-3 *Hardware and Software Requirements for VLAN Profiles*

Technet24

MS Switch Family	MS Switch Model	Minimum Firmware Required
MS100 Series	MS120	MS 15+
	MS125	MS 15+
	MS130	MS 15.22+
MS200 series	MS210	MS 15+
	MS225	MS 15+
	MS250	MS 15+
MS300 series	MS350 MS355	MS 15+ MS 15+
	MS390	CS 16+
MR	20/30H/33/42/42E/52/53E/70/74/84	MR 30+
CW WIFI6/7	all models	MR 30+
Catalyst C9K-M	C9K-M	CS 16+/Cloud Native IOS XE

Before you begin configuring VLAN profiles, consider the following points:

- Although you can create and use non-default VLAN profiles, you should manage (create or delete) VLANs only in the default profile.
- All VLANs added to the default profile automatically propagate to all other VLAN profiles configured within the network.
- Non-default profiles are designed to allow VLAN mappings to be overridden from the default but are not intended to provide a completely separate list of names and groups.

You can configure VLAN profiles by navigating to **Network-wide > VLAN Profiles**. The **VLAN Profiles** page consists of three tabs:

- **VLAN Profile:** Displays the list of all VLAN profiles defined for the network. By default, only the profile named **Default Profile** is available. You can add a new profile by clicking the **Add VLAN Profile** button. You can edit existing profiles by clicking the three dots next to the profile name.
- **Profile Assignment:** Allows you to bind custom profiles to devices within the network.
- **Settings:** Provides additional configuration options. Select the **Enable the Named VLAN** checkbox in this tab to activate RADIUS assignment for named VLANs.

Figure 11-23 shows the **VLAN Profiles** page with the **VLAN Profiles** tab in focus.

Technet24

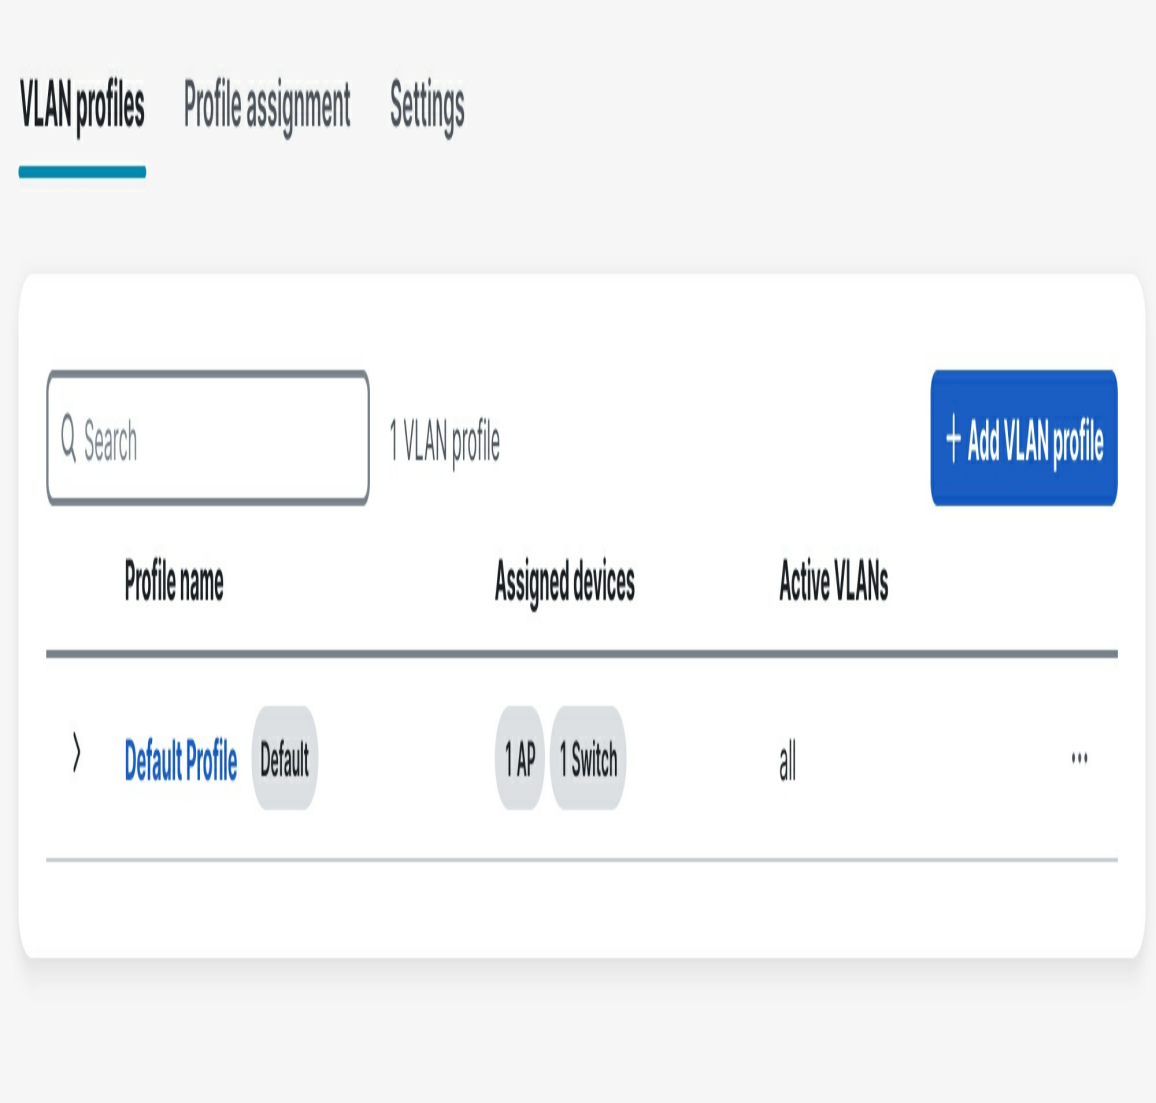


Figure 11-23 *VLAN Profiles Page*

To support named VLAN assignment, follow these minimum steps:

- Configure VLANs with their respective names in the **Default Profile** by clicking the three dots and editing the profile.
- Enable **Named VLAN** in the **Settings** tab.

Figure 11-24 displays the default VLAN profile. To add VLAN mappings, click **Add Named VLAN**, and after adding all VLANs, save the profile by clicking **Save Profile Changes** at the bottom of the page.

Edit profile

Profile name

Default Profile

Iname

Default

Active VLANs

all

VLAN name

+ Add Named VLAN

#	VLAN name	VLAN ID	Adaptive p	Actions
1	default	1		
2	Mng-Greenfield	10		
3	Clients	23		
4	Servers	24		
5	Core-Interconnect	100		

Group name

+ Add VLAN Group

#	Group Name	VLAN list	Actions

No record

Technet24

Figure 11-24 *Default VLAN Profile Page*

Within the VLAN profile, you can also map VLANs to specific Adaptive Policy tags. On the right side, you have the option to create VLAN groups. In this case, the name returned by the RADIUS server refers to the group instead of a specific VLAN.

As a final step to enable a named VLAN assignment, navigate to the **Settings** tab and enable the **Named VLAN** checkbox. Figure 11-25 provides an example of this step.

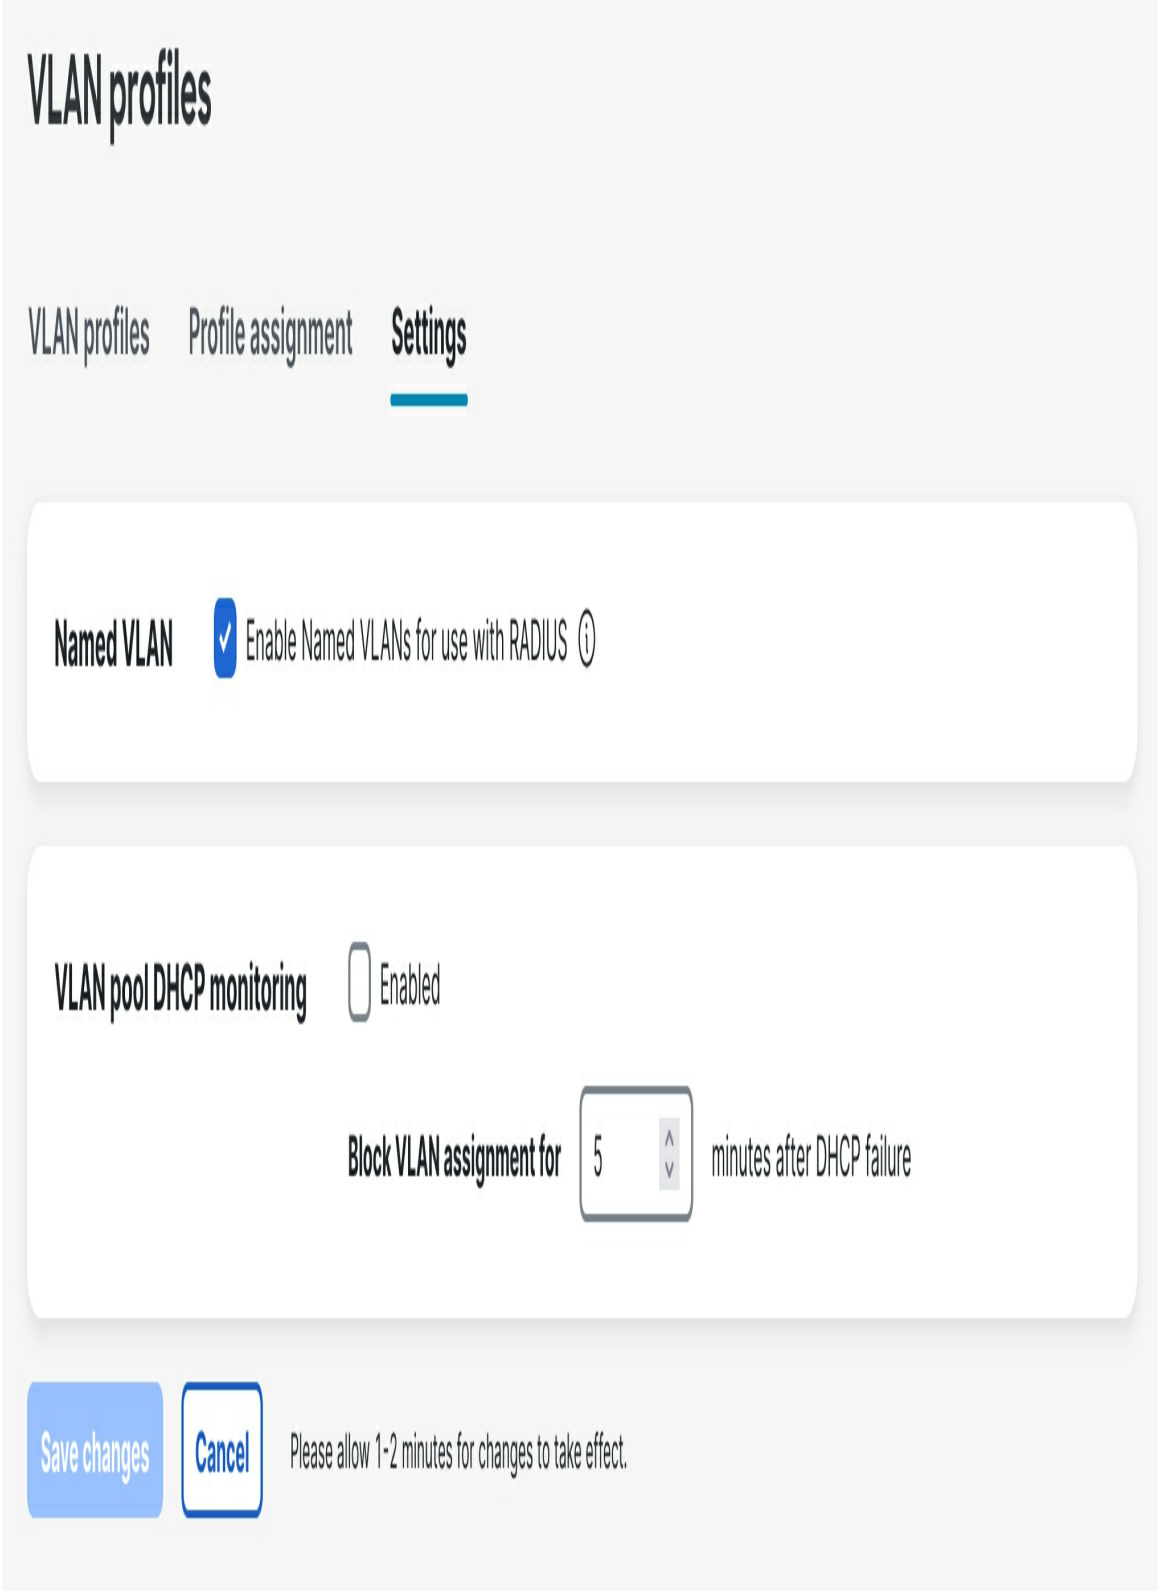


Figure 11-25 *Settings Tab of VLAN Profiles Page*

Technet24

Access Policies Structure

Before we review the individual features configurable in access policies, it is beneficial to understand the structure of an access policy. This understanding makes it easier to locate each feature within the access policy container as we discuss individual features in the following sections.

Switching access policies are located at **Switching > Access Policies**. Figure 11-26 shows the **Access Policies** page with two preconfigured access policies.


View old version
Access policies
+ Add policy
Search
2 policies
Policy name
Affected ports
Host mode
Actions
Branch-Closed
2
Single-Host
Branch-Open
1
Multi-Auth
Edit
Copy
Delete

Technet24

Figure 11-26 *The Access Policies Page*

You can create a new policy by clicking **Add Policy**. You also can edit, copy, or delete existing policies by clicking the three dots in the **Actions** column.

> ## Note
>
> Remove the access policy from any interface where it is assigned before deleting it.

The **Access Policy Settings** page can be divided into four logical sections:

- **Main settings:** This section allows you to define the access policy name, specify the authentication method, select or configure RADIUS servers for policies that use the RADIUS server authentication method, and configure general authentication settings such as authentication mode and RADIUS server testing.
- **Port settings:** In this section, you can configure port-level authentication settings such as host mode, 802.1X control direction, and related options.
- **Authorization parameters:** This section allows you to set authorization parameters, such as VLANs assigned in different scenarios for Closed mode. You can also assign GPACLs and adaptive policy tags based on various authentication and authorization outcomes in Open mode.
- **Additional settings:** This section includes additional settings such as URL redirection parameters and RADIUS result caching options.

> ## Note

The settings displayed in the access policy depend on the selected authentication method and authentication mode. For example, the options available with the RADIUS Server authentication method differ from those shown when you select **Meraki Authentication**. The same distinction applies between Open and Closed authentication modes.

Figure 11-27 shows the access policy and its configuration sections. The following sections discuss each configuration segment in greater detail.

Technet24

Name: MyAccessPolicy

Main settings

Authentication method: RADIUS server

Authentication Mode: Closed

RADIUS servers

- [ ] RADIUS Server testing
- [ ] RADIUS CoA support
- [ ] Enable RADIUS accounting servers

#	Name	Host	Secret	Auth	Port	Accounting	Port	Actions

No Servers Added

Select RADIUS

+ Add a server

RADIUS attribute specifying group policy name: None

Connection

Policy Type: None

Host mode: Single-Host

802.1X control direction: Both

Re-authentication interval:

Options

- [ ] Voice auth
- [ ] Disable port bounce

Port settings

Critical Auth VLAN

Data: None

Voice: None

Guest VLAN: None

Failed Auth VLAN: None

Authorization parameters

URL redirect walled garden

RADIUS caching

Caching timeout: 24

Additional settings

Figure 11-27 *Access Policy Logical Configuration Sections*

Access Policy: Main Settings Section

In the main section, you can begin by defining the name of a new access policy. Next, in the **Authentication Method** drop-down list, select **RADIUS Server**. While other options are available, this book focuses on integration with ISE and therefore always uses **RADIUS Server** as the authentication method. At the next step, select the **Authentication Mode**, as demonstrated in Figure 11-28.

Technet24

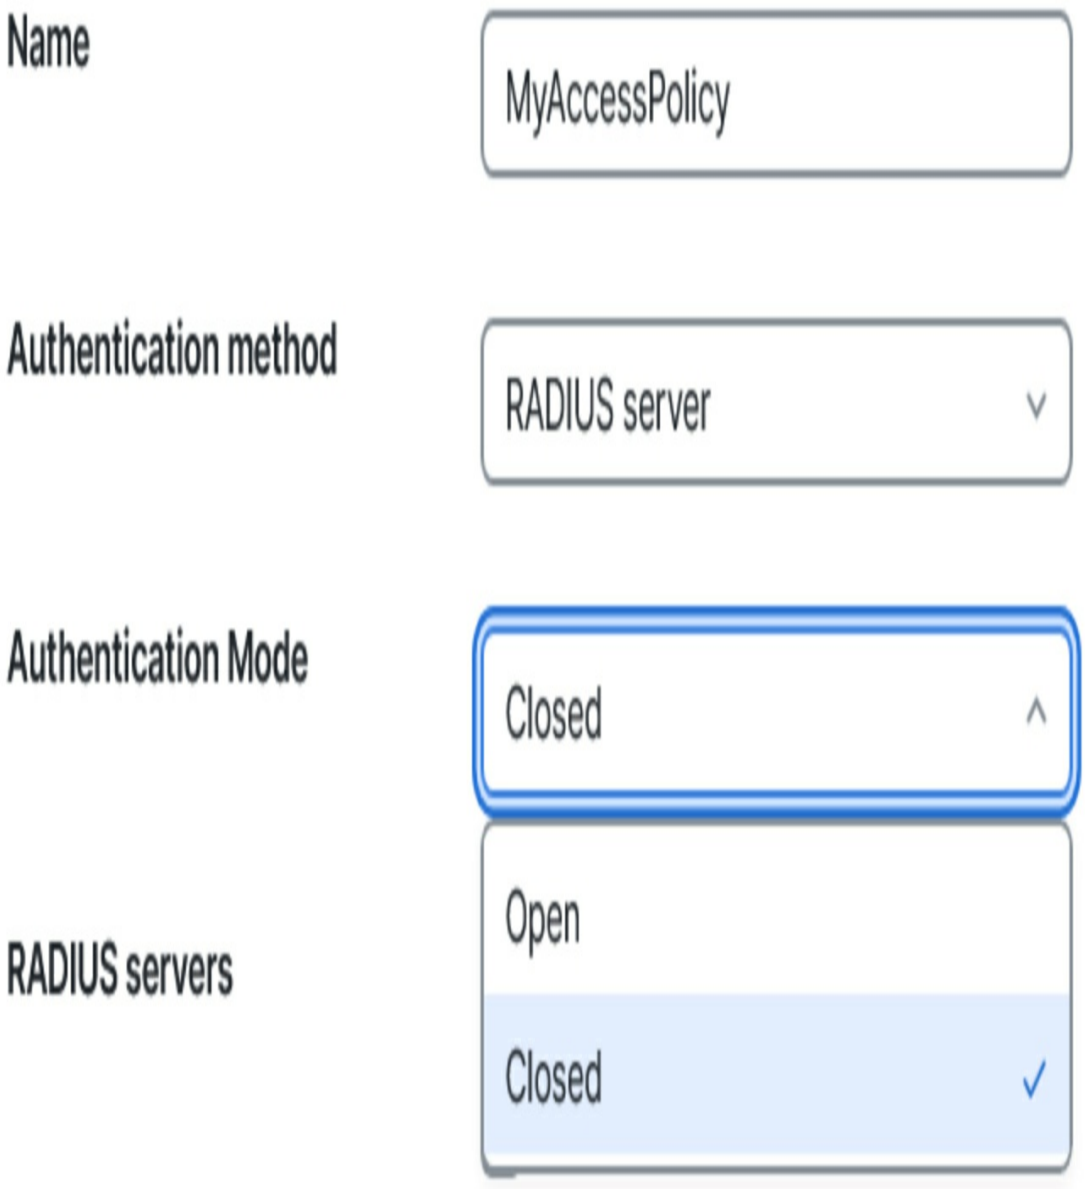


Figure 11-28 *Authentication Mode Selection*

Closed authentication mode on MS and CS switches functions in the same way as on classic Cisco Catalyst switches. Before authentication, the port allows only EAP frames; all other traffic is blocked. Depending on the specifics of the authentication and authorization result, the switch may assign multiple settings to the endpoint. These settings will

be discussed in more detail later.

At the time of writing, Open authentication mode on Meraki switches includes the following important features and limitations:

- **Supported switch models:** MS130 and MS150.
- **Minimum software version:** MS 18.
- **Supported RADIUS authorization attributes:** Filter-ID. The switch ignores other attributes in the Access-Accept message.
- **Local VLAN support:** The switch can assign endpoints only to the Data and Voice VLANs configured on the port. Features such as Guest VLAN, Failed VLAN, and Critical Auth VLAN for Data and Voice are not supported. Instead, in Open mode, you can assign Guest GPACL, Failed Auth GPACL, and similar policies, along with SGTs.

After you select the authentication mode, the remaining settings on the page adjust to reflect the specific mode. As a result, some of the following configuration settings will be explained twice: once for closed mode and again for open mode. Figure 11-29 displays the default state of the remaining settings available in the **Main Settings** section.

Technet24

Name

MyAccessPolicy

Authentication method

RADIUS server

Authentication Mode

Closed

RADIUS servers

RADIUS Server testing

RADIUS CoA support

Enable RADIUS accounting servers

#	Name	Host	Secret	Auth	Port	Accounting	Port	Actions



No Servers Added

Select RADIUS

+ Add a server

RADIUS attribute specifying group policy name

None

Figure 11-29 *Default State in the Main Settings Section for Closed Authentication Mode*

Figure 11-30 shows the default settings and user interface for the **Main Settings** section when you select **Open Authentication Mode**.

Technet24

Name

MyAccessPolicy

Authentication method

RADIUS server

Authentication Mode

Open

Pre-Authentication GPACL

None

RADIUS servers

RADIUS Server testing

RADIUS CoA support

Enable RADIUS accounting servers

#	Name	Host	Secret	Auth	Port	Accounting	Port	Actions



No Servers Added

Select RADIUS

+ Add a server

RADIUS attribute specifying group policy name

Filter-Id

Figure 11-30 *Default State in the Main Settings Section for Open Authentication Mode*

Pre-Authentication GPACL

The **Pre-Authentication GPACL** drop-down list is available only in Open authentication mode. Here you can select the name of the group policy that will act as the pre-authentication GPACL.

The pre-authentication GPACL determines which traffic is permitted on the port before authentication takes place. By default, all traffic is allowed in Open authentication mode.

Depending on the pre-authentication GPACL configuration, Open mode can operate in two ways:

- When a permit-all pre-authentication ACL is applied, Open mode functions like Monitor mode on classic Catalyst switches, where all traffic is allowed.
- When a granular pre-authentication ACL allowing only specific traffic is applied, Open mode behaves like Low Impact mode, providing limited access to essential services (as defined by the ACL) until authentication is complete.

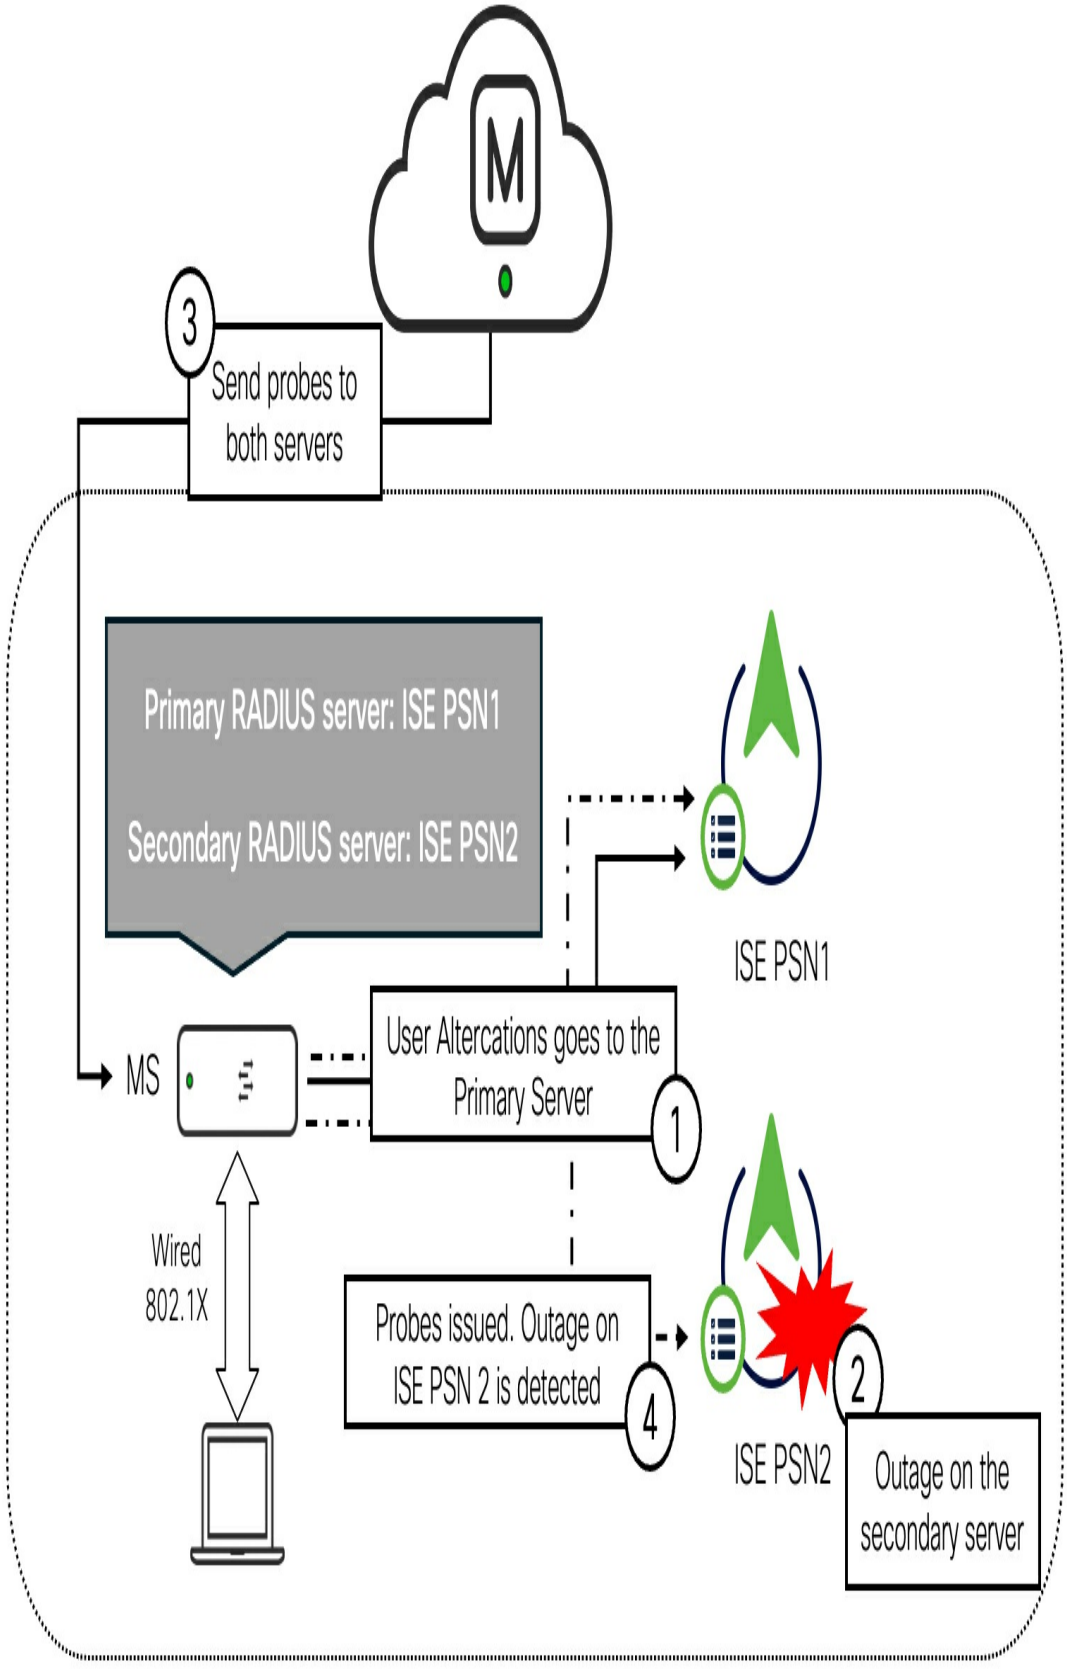
RADIUS Server Testing

As part of recent updates to the access policy, the RADIUS testing and RADIUS monitoring features were combined into a single checkbox called **RADIUS Server Testing**. In Closed mode, the default state for this feature is off. We will explain in detail how this feature operates after the integration of RADIUS server testing and monitoring:

- RADIUS testing is responsible for periodically sending synthetic authentication requests to all RADIUS servers configured in the

Technet24

access policy, using the test username meraki_8021x_test. The dashboard initiates these tests by instructing the switch to perform the authentication. Tests occur at long intervals, typically once every 24 hours. The interval is not configurable in the dashboard, and the dashboard generally requests the switch to perform the test once a day for each RADIUS server. However, the actual time between tests may vary depending on the dashboard's background task queue. **Note** Both Access-Accept and Access-Reject responses are considered valid results. Failures occur only in the event of RADIUS timeouts. This feature helps detect unresponsive RADIUS servers that are not actively handling authentication. For a better understanding, see Figure 11-31.


M
3
Send probes to both servers
Primary RADIUS server: ISE PSN1
Secondary RADIUS server: ISE PSN2
ISE PSN1
MS
User Altercations goes to the Primary Server
1
Wired 802.1X
Probes issued. Outage on ISE PSN 2 is detected
4
2
ISE PSN2
Outage on the secondary server

Technet24

Figure 11-31 *RADIUS Server Testing*

The RADIUS testing flow operates as follows:

- The switch uses the first server listed in the access policy for all authentication, as long as it remains available. Typically, the switch detects an outage of this server through RADIUS timeouts that occur during regular authentication requests.
  1. In the event of an outage on any secondary server, the switch cannot detect its unavailability through regular authentication, since secondary servers are not actively used for endpoint or user authentication.
  2. The dashboard periodically instructs the switch to send synthetic authentication requests to all servers configured in the access policy.
  3. These requests can detect outages on secondary servers and mark them as dead.
- The RADIUS monitoring process only starts when RADIUS servers become unreachable. When this situation occurs, the switch begins sending test authentication requests for the user meraki_8021x_test. At the time of writing, MS switches initiate the RADIUS monitoring process only when all RADIUS servers in the access policy are marked as unreachable and then send requests to all configured servers every minute. For CS switches, monitoring starts as soon as the first server becomes unreachable. CS switches send test authentication requests every five minutes.

Note

RADIUS server testing is disabled by default with Closed authentication mode. Enabling it is recommended. RADIUS server testing is always enabled and cannot be disabled with

Open authentication mode.

Normally, switches fail over a specific authentication to the next available RADIUS server when a timeout occurs, and subsequent authentication attempts experience the same delay. RADIUS testing and RADIUS monitoring help minimize user impact by redirecting authentication requests away from unresponsive RADIUS servers and by tracking unreachable servers to ensure they can be reactivated with minimal disruption to ongoing authentications.

RADIUS CoA Support

The **RADIUS CoA Support** checkbox enables Change of Authorization support for the RADIUS servers specified in the access policy. CoA allows policy changes to be applied during an existing authenticated session in response to updates on the ISE side, such as posture status changes, profile updates, or guest authentication.

CS and MS switches support CoA reauthentication, CoA disconnect, and CoA port-bounce. **Note** In the case of Open authentication, CoA is enabled by default.

Enable RADIUS Accounting Servers

By default, servers configured in the access policy are used only for authentication. To enable RADIUS accounting, you must check the **Enable RADIUS Accounting Servers** option. If you add servers from the organization-wide RADIUS servers list and they have accounting enabled, you must enable this checkbox for the servers to appear in the drop-down list next to **Add a Server**.

Enabling RADIUS accounting also enables Device Sensor functionality. Meraki switches send interim-accounting updates every 20 minutes by default, and these updates include device sensor attributes discovered

Technet24

by the switch.

As of MS 10.19, device sensor functionality for enhanced device profiling is available by including CDP and LLDP information in the RADIUS Accounting message on models MS120, MS125, MS130, MS130R, MS150, MS220, MS225, MS250, MS320, MS350, MS355, MS410, MS425, and MS450. As of version 14.19 and later, the MS390 and C9300-M also support device sensors with enhanced attributes across LLDP, CDP, mDNS, and DHCP for profiling.

RADIUS Attribute Specifying Group Policy Name

The **RADIUS Attribute Specifying Group Policy Name** drop-down list allows you to select an attribute in the RADIUS Access-Accept message whose value will be used to identify the group policy. At the time of writing, only Filter-ID can be selected. In the case of **Open Authentication**, **Filter-ID** is selected by default and cannot be removed. Figure 11-32 displays the **Main Settings** section with a completed configuration for Closed mode.


Name
Branch-Closed
Authentication method
RADIUS server
Authentication Mode
Closed
RADIUS servers
RADIUS Server testing
RADIUS CoA support
Enable RADIUS accounting servers
Name Host Secret Auth Port Accounting Port Actions
1 ISE-1 172.22.15.1 1812 1813
2 ISE-2 172.22.15.2 1812 1813
2 Select RADIUS
+ Add a server
RADIUS attribute specifying group policy name
Filter-Id

Technet24

Figure 11-32 *Main Settings Section with a Completed Configuration for Closed Mode*

Figure 11-33 displays the **Main Settings** section with a completed configuration for Open mode.


Name
Branch-Open
Authentication method
RADIUS server
Authentication Mode
Open
Pre-Authentication GPACL
Branch-Open-Auth - 100
RADIUS servers
RADIUS Server testing
RADIUS CoA support
Enable RADIUS accounting servers
Name Host Secret Auth Port Accounting Port Actions
1 ISE-1 172.22.15.1 1812 1813
2 ISE-2 172.22.15.2 1812 1813
2 Select RADIUS
+ Add a server
RADIUS attribute specifying group policy name
Filter-Id

Technet24

Figure 11-33 *Main Settings Section with a Completed Configuration for Open Mode*

Access Policy: Port Settings

The settings available for configuration in the **Port Settings** section depend on the selected authentication mode. Figure 11-34 demonstrates the settings in their default state when you select **Closed Authentication Mode**.

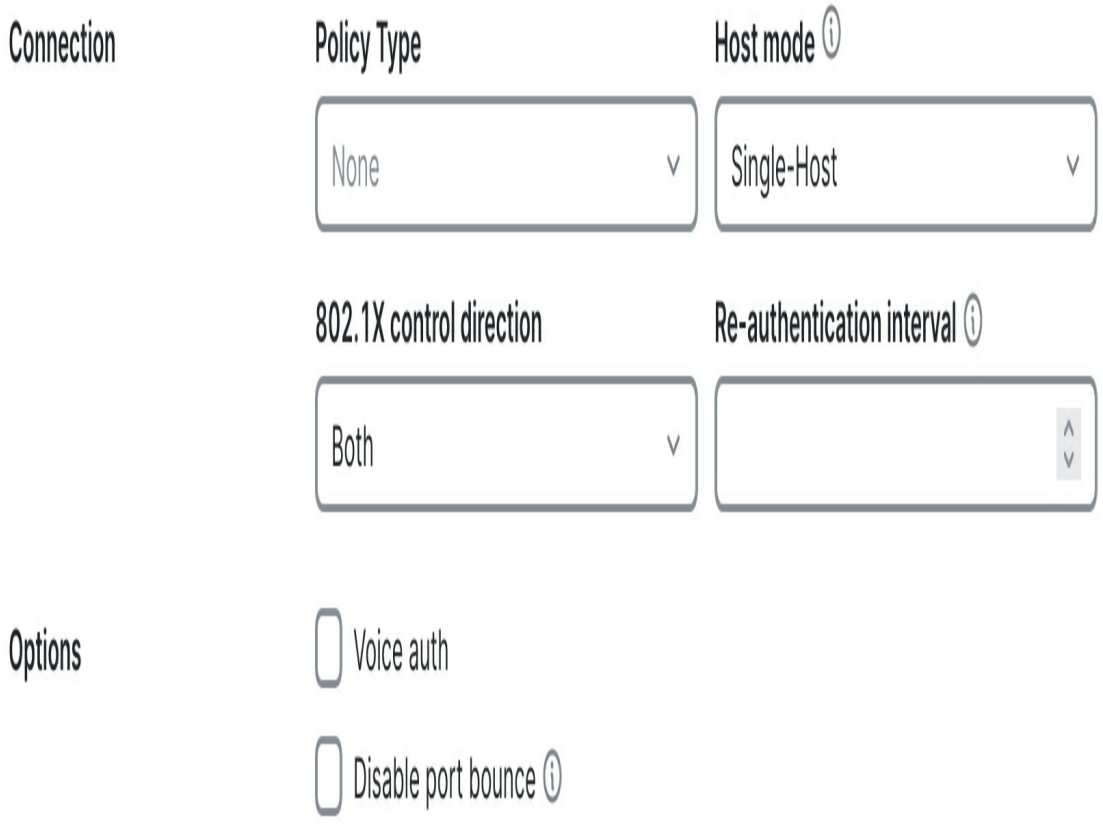


Figure 11-34 *Defaults in the Port Settings Section for Closed Mode*

To demonstrate the difference, Figure 11-35 shows the settings in their default state when you select **Open Authentication Mode**.



Figure 11-35 *Defaults in the Port Settings Section for Open Mode*

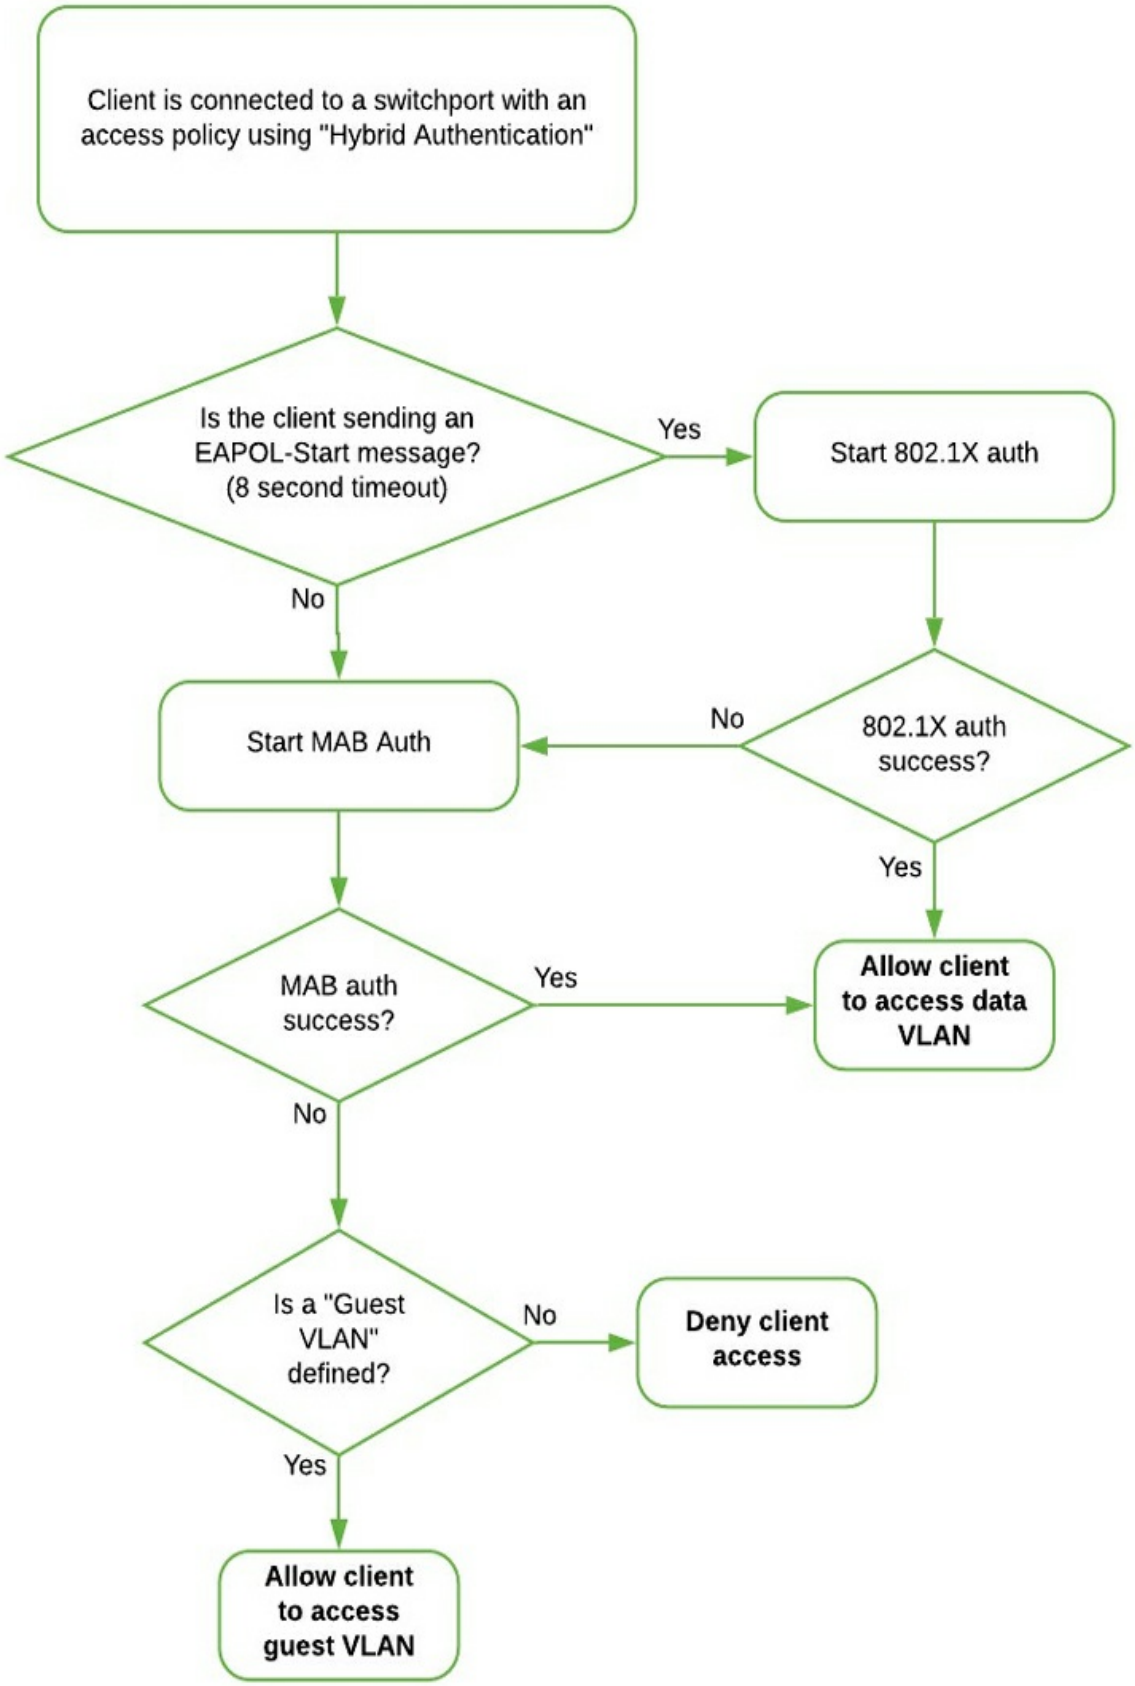
Now let's review the settings and their options one by one:

- **Policy Type:** This setting determines how different authentication types, such as MAB or 802.1X, are performed on the port. The available approaches are similar to those supported on classic Catalyst switches, so only a summary is provided here.
  - **802.1X:** When you enable an 802.1X access policy on a switch port, the switch prompts the connected client with EAP Identity Request messages to initiate 802.1X authentication using one of the EAP methods supported by both the client and the server. If authentication succeeds, the switch assigns session attributes from the RADIUS Access-Accept message, such as VLAN, Filter-ID, and web redirect parameters. If the client does not respond

Technet24

or authentication fails, the authorization result depends on factors such as the authentication mode and the configured VLANs or GPACLs for guest access, or failed authentication access. These options and their combinations will be discussed later in greater detail.

- **MAC Authentication Bypass:** With MAC authentication, the switch listens to frames on the port and uses the source MAC address as the username in RADIUS packets. Although this method is not considered secure, it is often used in combination with technologies such as profiling to provide authenticated access for endpoints without 802.1X support.
- **Hybrid Authentication:** With hybrid authentication, the switch combines 802.1X and MAB methods, giving priority to 802.1X. Figure 11-36 outlines the general flow that Meraki switches use on interfaces configured for hybrid authentication. The specifics of network access control, including which VLANs and other authorization attributes apply in different scenarios, will be discussed in detail later in this chapter.


Client is connected to a switchport with an access policy using "Hybrid Authentication"
Is the client sending an EAPOL-Start message? (8 second timeout)
Yes
Start 802.1X auth
No
802.1X auth success?
No
Start MAB Auth
Yes
MAB auth success?
Yes
Allow client to access data VLAN
No
Is a "Guest VLAN" defined?
No
Deny client access
Yes
Allow client to access guest VLAN

Technet24

Figure 11-36 *Hybrid Authentication Flow*

- **Concurrent Authentication:** In concurrent authentication, the switch processes both MAB and 802.1X simultaneously. Cisco ISE does not currently support this authentication method. **Note** The **Concurrent Authentication** checkbox appears only in Closed mode when you select **Hybrid Authentication** as the policy type.

- **Host Mode:** This setting determines how the switch handles MAC address learning on the port before and after authentication. The options are the same as those available on classic Catalyst switches. The modes are as follows:

  - **Single-Host:** With single-host authentication, a connected device attempts authentication, and if it fails, the client is denied access. This mode is recommended for switch ports that support only one client. If multiple devices connect to the same switch port, such as through a hub or daisy-chain off a VoIP phone, only one client will receive network access upon successful authentication. The switch ignores all subsequent authentication requests from other clients, and they will not receive access.

  - **Multi-Host:** With multi-host mode, a single successful authentication puts the port into a forwarding state. The switch ignores all subsequent authentication attempts, and all next endpoints get the same access level as the first authenticated device. This mode is recommended in deployments where the authenticated device acts as a point of access to the network, such as hubs and access points.

  - **Multi-Domain:** With multi-domain authentication, the switch can authenticate one device on the data VLAN and one device on the voice VLAN. If a second device is detected on either

VLAN, the switch does not grant access to that device. This mode uses hybrid authentication and requires voice VLAN authentication. Multi-domain authentication is recommended for switch ports connected to a phone with a device located behind the phone. Authentication operates independently on each VLAN and does not affect the forwarding state of the other VLAN. To ensure proper authentication, you must specify a voice VLAN on the switch port. Failure to specify a voice VLAN will prevent client devices from passing traffic correctly after authentication.

- **Multi-Auth:** With Multi-Auth mode, the switch requires each connected device to authenticate. Multiple devices may connect to each port. After the switch assigns a VLAN to an endpoint on the port, subsequent endpoints must use matching VLAN information, or the switch will deny them access to the port. Only one client is supported on the voice VLAN. Guest VLANs are not supported in this mode.

Note

The default interface mode is a single-host for the Closed mode and Multi-Auth for the Open mode.

- **802.1X Control Direction:** The control direction drop-down list appears only in Closed authentication mode. This setting allows you to define how the switch handles the direction of packets that are allowed or denied before authentication.
  - **Both:** This is the default option. With control direction set to **Both**, the switch permits only EAP frames in the inbound direction prior to successful authentication and blocks all outbound frames before authentication is complete.

Technet24

- **Inbound-Only:** With inbound-only control direction, the switch port blocks all ingress traffic but allows limited egress traffic from the network to the connected device. This setting is often used to permit Wake-on-LAN magic packets to reach a sleeping host. After receiving the packet, the host can attempt normal 802.1X or MAB authentication to authorize the switch port for full ingress and egress traffic.

- **Voice Auth:** By default, Meraki switches can identify IP phones using information from CDP and LLDP frames. When this option is disabled (the default state), the switch does not perform authentication for devices identified as IP phones. To enable authentication for these devices over RADIUS, you must check the **Voice Auth** checkbox.

- **Re-authentication Interval:** This setting enables periodic reauthentication for clients. The reauthentication interval is defined in seconds. You should carefully consider the reauthentication frequency to avoid overwhelming RADIUS servers with frequent authentication requests.

Note

It is not recommended to set reauthentication intervals lower than 7200 seconds (2 hours).

- **Disable Port Bounce:** In the default state, when this setting is not enabled, the switch bounces the ports of all endpoints placed into Critical VLANs after all RADIUS servers become unreachable, after at least one of the RADIUS servers comes back online. This action retriggers authentication when RADIUS monitoring detects that server availability has been restored.

  If you enable this setting, the endpoint remains in the Critical VLAN even after the RADIUS server becomes reachable again. In

this case, the endpoint must reauthenticate naturally, such as through user logout/login or device power cycle.

Figure 11-37 displays the **Port Settings** section with a completed configuration for Closed mode.

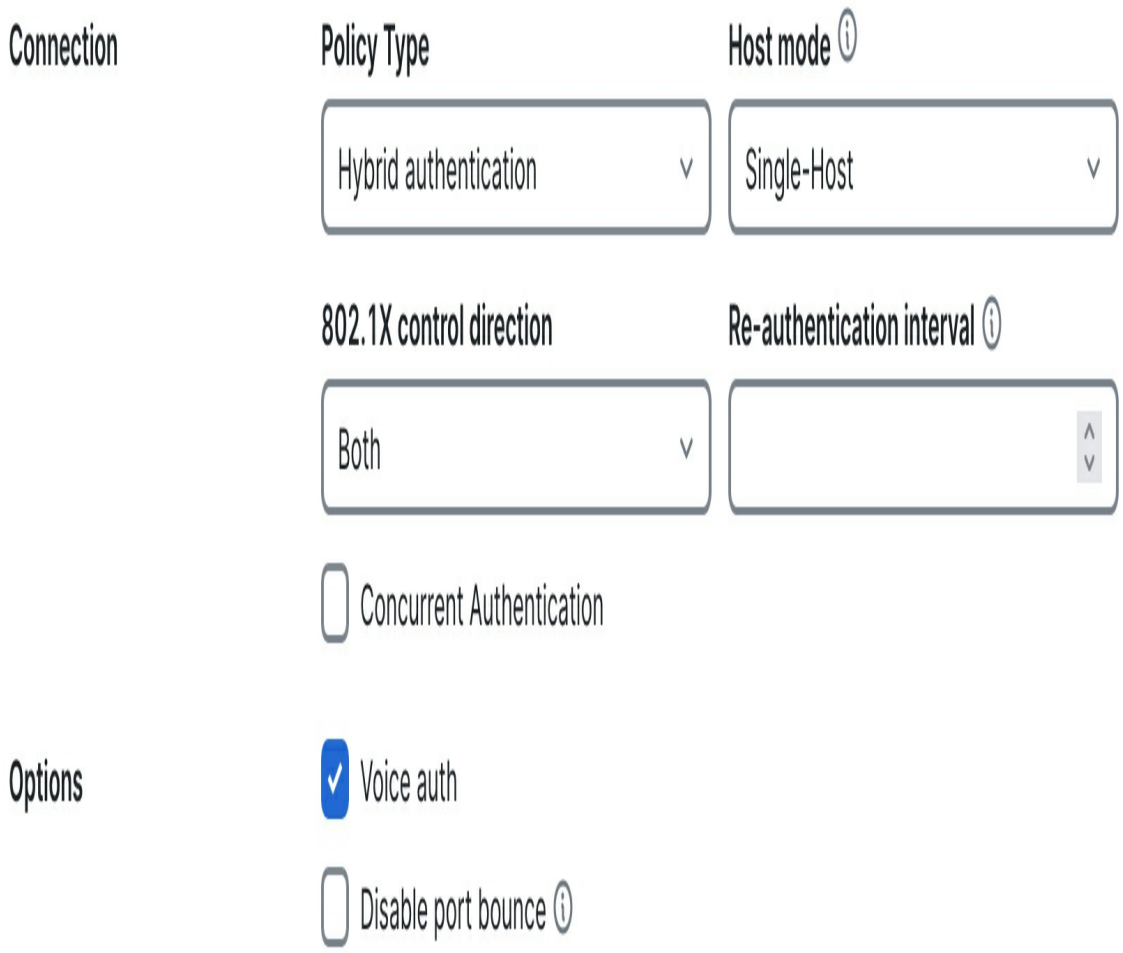


Figure 11-37 *Port Settings Section with a Completed Configuration for Closed Mode*

Figure 11-38 displays an example of the **Port Settings** section with a completed configuration for Open mode.

Technet24



Figure 11-38 *Port Settings Section with a Completed Configuration for Closed Mode*

Access Policy: Authorization Parameters Section

The settings available for configuration in the **Authorization Parameters** section depend heavily on the selected authentication mode.

In simple terms, in Closed authentication mode, most authorization settings are based on VLANs. In Open authentication mode, VLAN support is limited at the time of writing, so most additional authorization settings rely on GPACLs and Adaptive Policy tags. For this reason, it is logical for us to outline and explain the settings available in both modes separately.

Figure 11-39 shows the default state of the settings available under the **Authorization Parameters** section when Closed mode is selected.

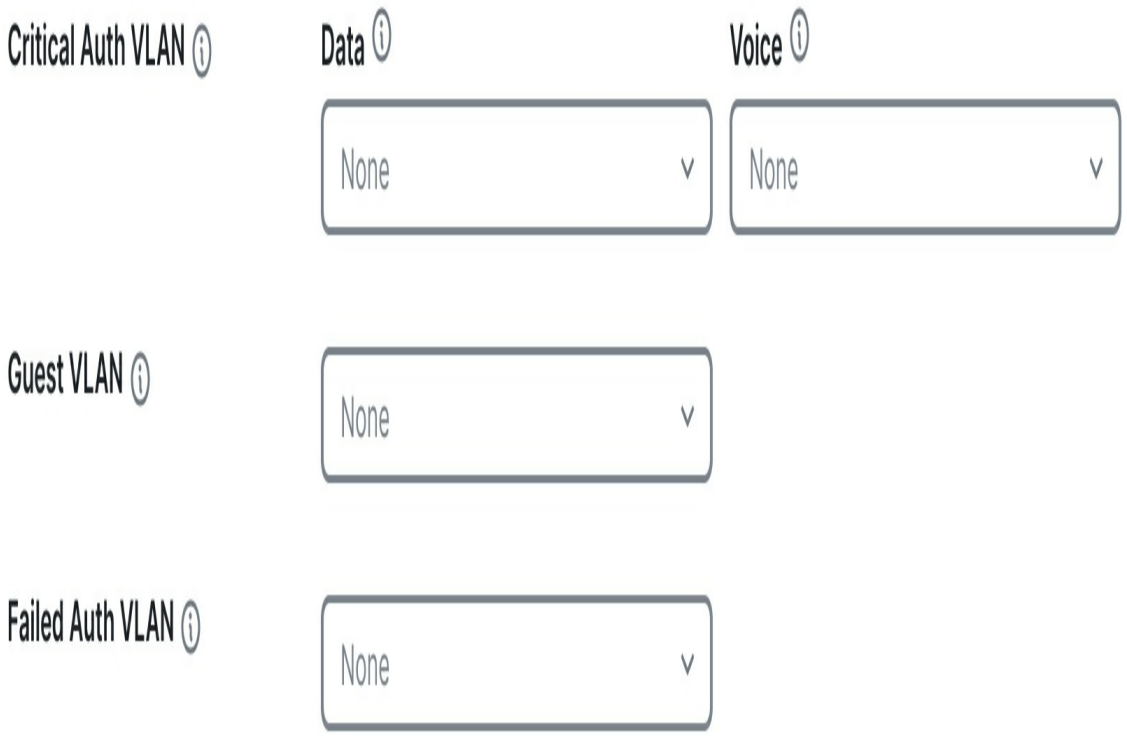


Figure 11-39 *Authorization Parameters Section Defaults in the Closed Mode*

- **Critical Auth VLAN:** Here you can assign VLANs for data and voice domains that will be used when all RADIUS servers are unreachable and a new authentication request must be processed. When named VLANs are used, you can select VLANs from the list; when named VLANs are not enabled, you can define VLANs manually.

  Keep in mind that, by default, 802.1X clients do not permit network access when authentication fails. When such clients are placed into Critical Auth VLANs, the switch must respond with an EAP failure to the EAP start from the clients. To ensure 802.1X

Technet24

clients can access the network when RADIUS servers are unavailable, configure supplicants to fall back to unauthorized network access.

- **Guest VLAN:** This VLAN is assigned when the endpoint does not respond to the authentication request from the switch; it is primarily useful for ports configured for 802.1X without hybrid authentication.
- **Failed Auth VLAN:** The purpose of this VLAN is to provide limited access when the RADIUS server returns an Access-Reject for the endpoint; this ensures some level of connectivity even if authentication is rejected. For 802.1X endpoints, fallback to unauthorized network access must be enabled on the supplicant to allow connectivity after authentication is rejected.

Table 11-4 shows various combinations of VLANs and authentication outcomes, depending on which VLANs you configure in this section.

Table 11-4 *Authentication Results Combination for Different VLANs Combinations*

Configured Options	Authentication Result		
	EAP Timeout (For 802.1X Policies Only)	RADIUS Timeout (Server Unreachable)	Authentication Fail (Access-Reject)
Guest (existing behavior)	Guest VLAN	Guest VLAN	Access denied*
Failed	Access denied	Access denied	Failed Auth VLAN
Critical	Access denied	Critical Auth VLAN	Access denied
Guest and Failed	Guest VLAN	Guest VLAN	Failed Auth VLAN
Guest and Critical	Guest VLAN	Critical Auth VLAN	Access denied*
Critical and Failed	Access denied	Critical Auth VLAN	Failed Auth VLAN
Guest, Failed and Critical	Guest VLAN	Critical Auth VLAN	Failed Auth VLAN

*When you use hybrid authentication without concurrent authentication enabled, a client that fails both 802.1X and MAB authentication will be placed in the Guest VLAN.

Figure 11-40 shows the **Authorization Parameters** section configured in Closed mode.

Technet24

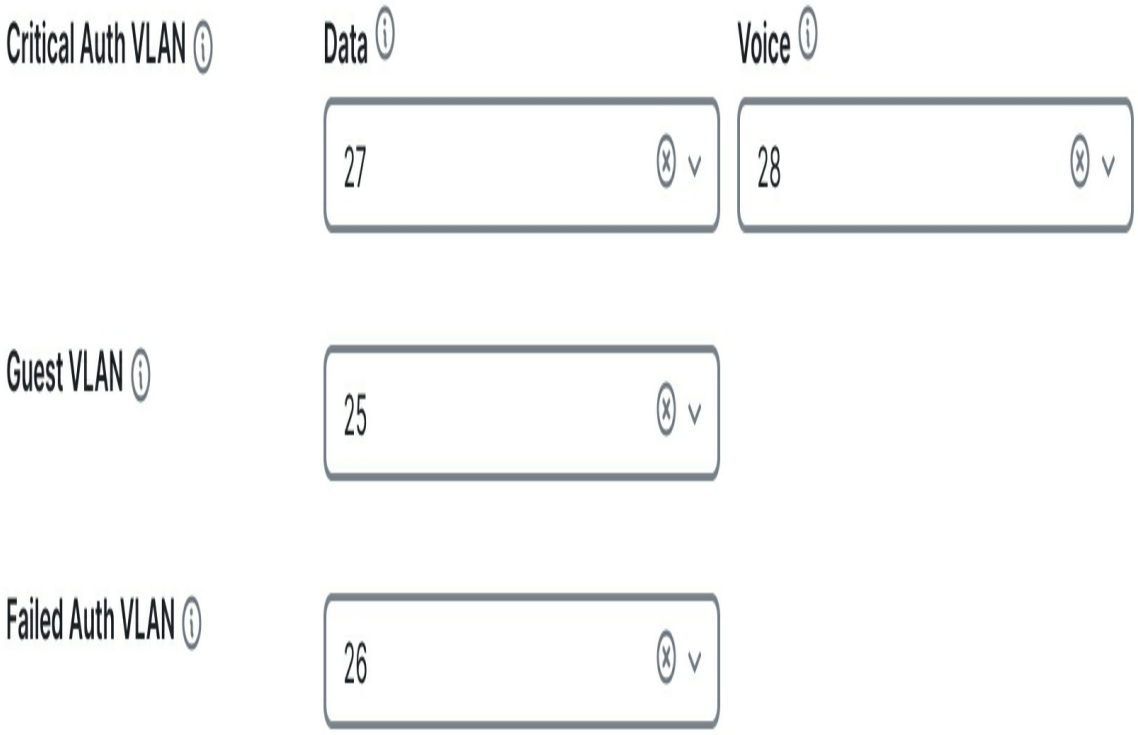


Figure 11-40 *The Authorization Parameters Section in Closed Mode*

Figure 11-41 shows the default state of the settings available under the **Authorization Parameters** section when Open mode is selected.

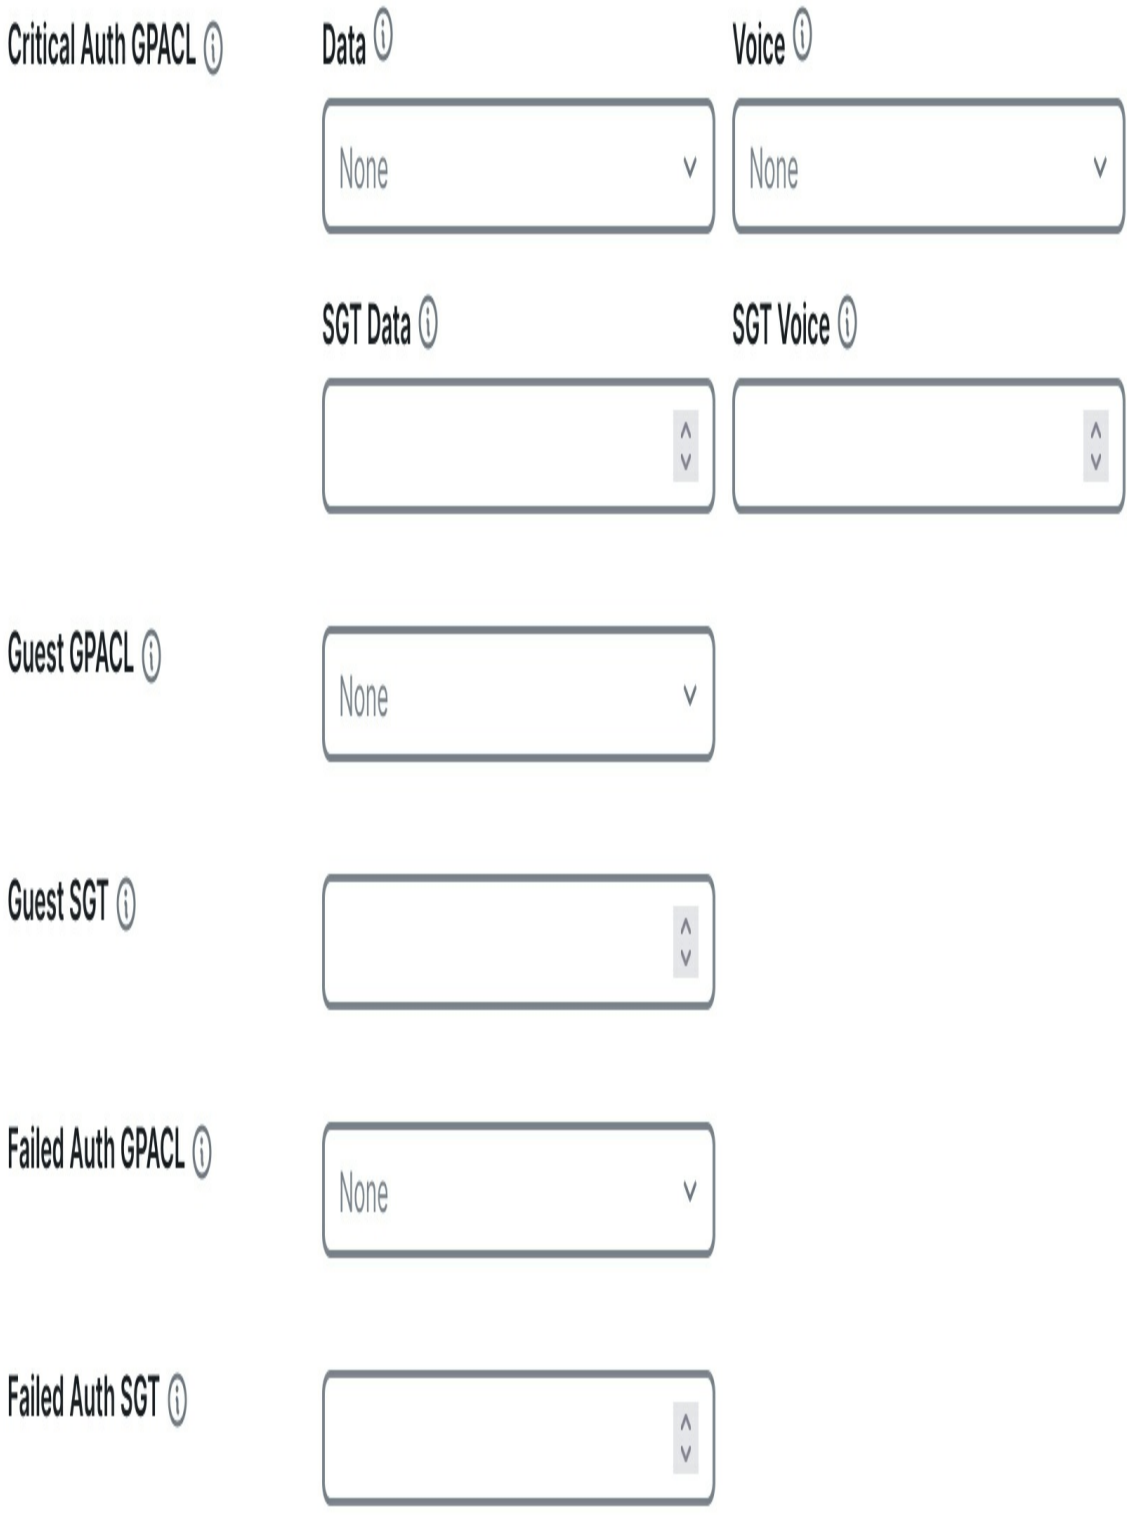


Figure 11-41 *Authorization Parameters Section Defaults in the Open Mode*

Technet24

In Open mode, you can specify GPACLs for different scenarios such as **Critical Data**, **Critical Voice**, **Failed Auth**, and **Guest Auth** instead of VLANs. For each of these scenarios, you can also define SGTs to assign to the endpoint packets if the session results in one of these outcomes.

Access Policy: Additional Settings

The last section of Access Policy settings includes two configuration options: URL redirect settings and RADIUS caching configuration. URL redirect is not supported with Open authentication mode. RADIUS caching is supported only on CS switches and MS390 switches that meet the minimum software version requirements — CS 17+ or IOS-XE. The default settings in this section, when Closed authentication mode is used, are shown in Figure 11-42.

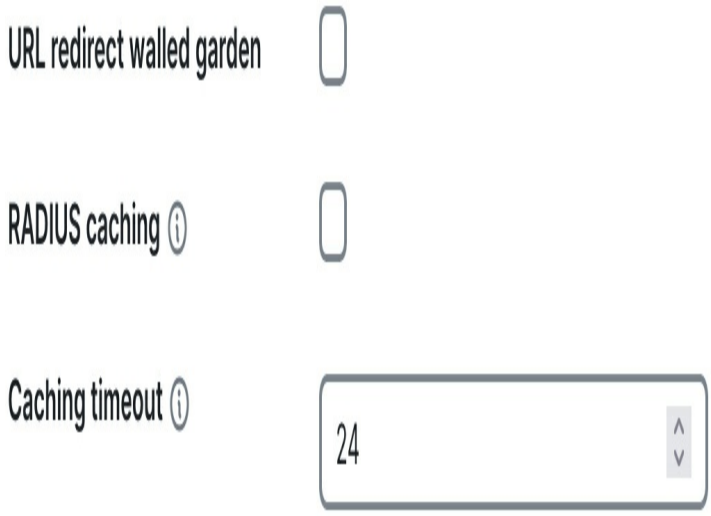


Figure 11-42 *Defaults in Additional Settings Section*

- **URL Redirect Walled Garden:** This setting enables web redirection support for interfaces assigned to the specified access policy; this is used for advanced scenarios such as Guest, BYOD,

and Posture.

Web redirection is supported on the following switch models: MS210, MS225, MS250, MS350, MS355, MS390, MS410, MS420, MS425, and 9K-M, provided that CoA is enabled in the access policy.

For URL redirection on MS390 and 9K-M switches, the Access-Accept message from the RADIUS server must include a walled-garden ACL name. The first access policy in the list requires an ACL name of Walled-Garden-00, with each subsequent policy incrementing by one (for example, Access Policy 2 requires Walled-Garden-01). The Walled Garden ACL on MS390 and C9300-M switches permits DHCP, DNS, and IP access to the specified walled garden IP addresses.

Figure 11-43 shows access policies with their corresponding redirect ACL names that must be configured on the ISE side, along with the authorization profile created for guest web-redirect on ISE for the policy Branch-Closed.

Technet24

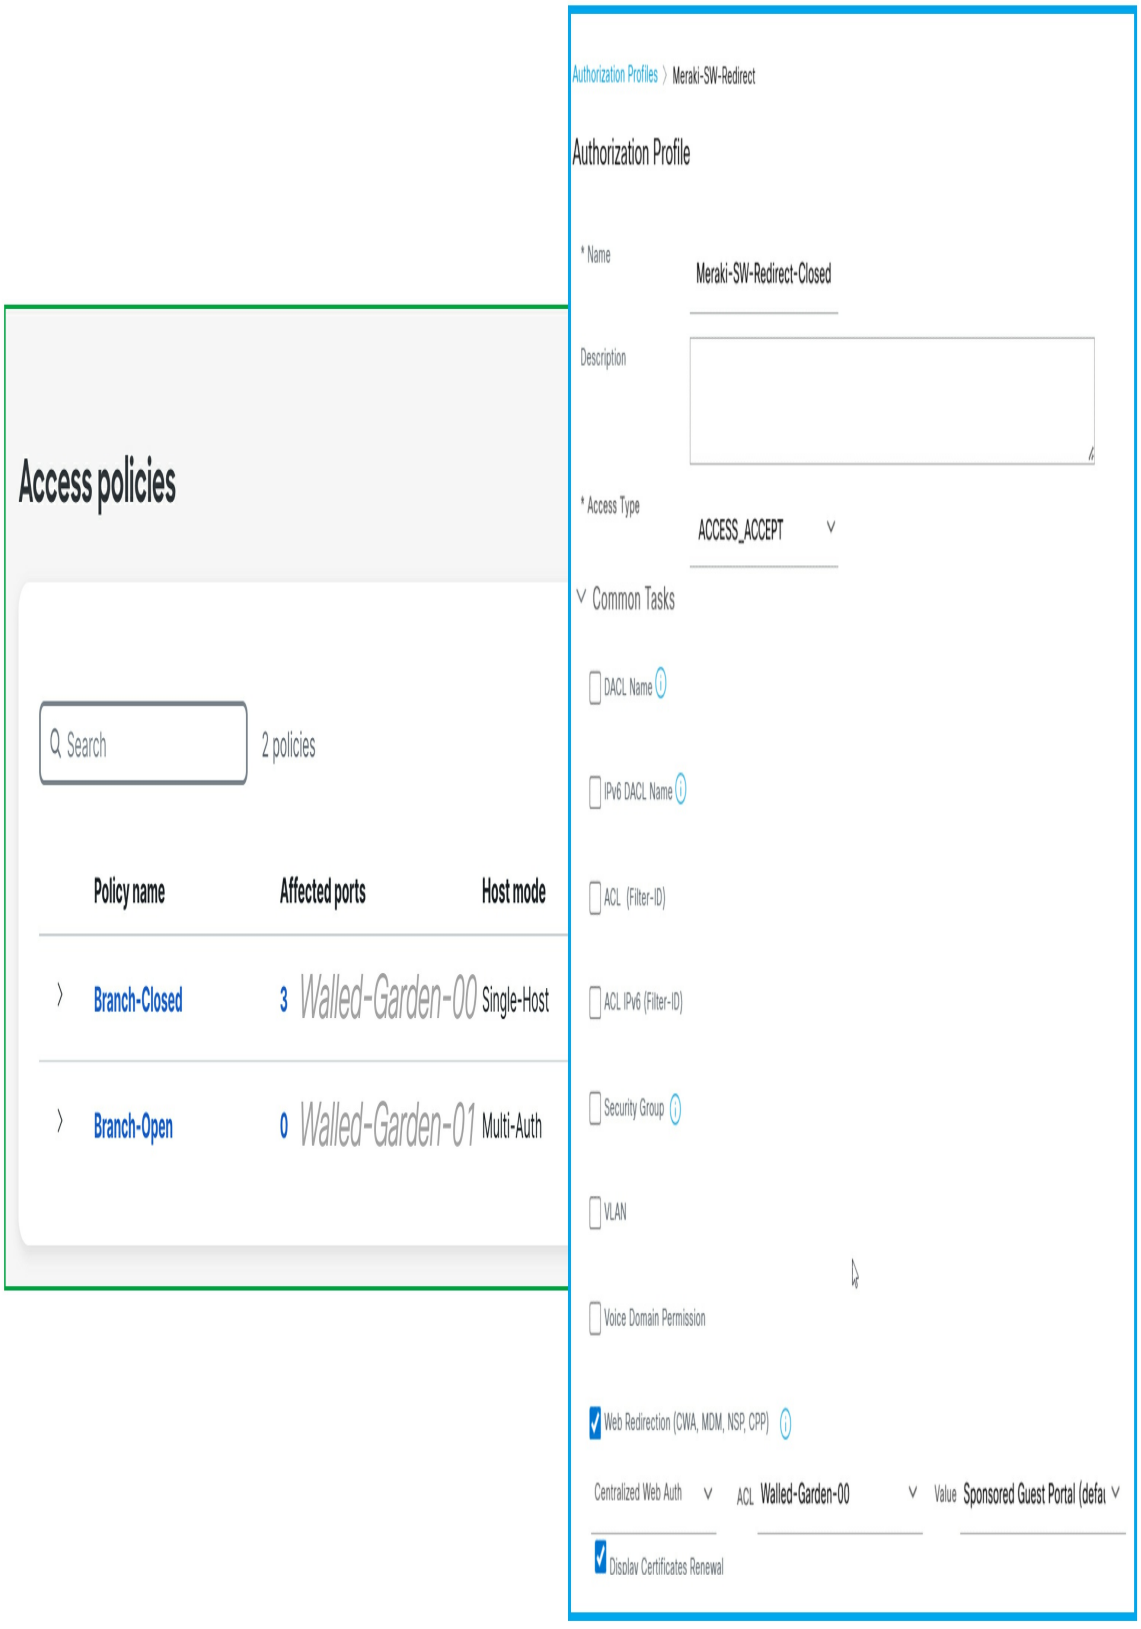

Access policies
Search
2 policies
Policy name
Affected ports
Host mode
Branch-Closed
3
Walled-Garden-00
Single-Host
Branch-Open
0
Walled-Garden-01
Multi-Auth
Authorization Profiles
Meraki-SW-Redirect
Authorization Profile
* Name
Meraki-SW-Redirect-Closed
Description
* Access Type
ACCESS_ACCEPT
Common Tasks
DACL Name
IPv6 DACL Name
ACL (Filter-ID)
ACL IPv6 (Filter-ID)
Security Group
VLAN
Voice Domain Permission
Web Redirection (CWA, MDM, NSP, CPP)
Centralized Web Auth
ACL
Walled-Garden-00
Value
Sponsored Guest Portal (defa
Display Certificates Renewal

Figure 11-43 *Access Policy and ISE Authorization Profile for Web Redirection on MS390 and C9300-M Switches*

- **URL Redirect Walled Garden Ranges:** This setting becomes available when you enable **URL Redirect Walled Garden**. You can specify a list of IP addresses and subnets; traffic to these destinations will be excluded from web redirection.
- **RADIUS Caching:** Selecting this checkbox instructs switches that support RADIUS caching to temporarily save the RADIUS authorization results for connected endpoints. When you enable caching and RADIUS servers become unreachable, the switch can apply the previous authorization attributes to the endpoints. This helps users and devices maintain their previous level of access during outages. Compared to Critical VLANs, this feature is less disruptive because the endpoint's access does not change when reauthentication is needed and RADIUS servers are unavailable. This functionality is based on the IOS-XE feature called AAA Authorization and Authentication Cache.
- **Caching Timeout:** This setting instructs the switch how long to retain the last authorization result in the cache for connected endpoints.

Methods of Access Control for Wireless Connections

This section focuses on methods for implementing network access control in the wireless Meraki environment, including configuration steps and best practices.

Main Dashboard Tools for Wireless NAC

In the case of wireless NAC, most dashboard tools required for

Technet24

configuration remain the same as those used for wired NAC. The following are the available tools and those not yet covered:

- Network-level group policies (**Network-wide > Group Policies**)
- Network-level VLAN profiles (**Network-wide > VLAN Profiles**)
- Wireless access control settings (**Wireless > Access Control**). Access control serves as the primary configuration container that consolidates most settings required at the SSID level including those needed for NAC functionality. In the case of NAC, access control includes parameters such as security type (open, password-based, enterprise, etc.), RADIUS settings, and other essential SSID parameters used for access control. The following sections focus primarily on access control features and their configuration.
- Organization-level Adaptive Policy settings (**Organization > Adaptive Policy**)

Network-level group policies and network-level VLAN profiles apply to wireless NAC in the same way as they do to wired NAC. Because Adaptive Policy is a comprehensive topic on its own, it will be discussed separately in greater detail. The following sections focus primarily on the wireless access control settings relevant to Meraki wireless integration with Cisco ISE.

Network-Level Group Policies

The general concept of Meraki group policy was discussed in the "Wired NAC" section.

MR devices also support group policies, but the range of supported features is broader compared to wired NAC setups, where policies can only function as GPACLs.

The following group policy features are supported by MR/CW:

- **Scheduling:** Group policies can be scheduled using the Schedule option; this allows the policy to remain active or inactive only during the specified times.
- **Per-client bandwidth limit:** This setting ensures that users do not consume more bandwidth than permitted.
- **VLAN tagging:** Group policy can include a VLAN ID; this allows assignment of all endpoints with the specified group policy to a particular VLAN.
- **Splash page authorization:** You can use a group policy to bypass the Splash page enabled in **Access Control**.
- **Layer 3 firewall rules:** Layer 3 firewall rules provide you, as administrator, with granular access control over outbound client traffic. For the MR/CW series, outbound traffic refers to client traffic originating from the wireless network and destined for the wired LAN or the Internet.
- **Layer 7 firewall rules:** Cisco Meraki access points provide the ability to create Layer 7 firewall rules that deny specific traffic types. While most firewall rules inspect headers at Layer 3 (IP address) and Layer 4 (Port), Layer 7 rules inspect the packet payload to match against known traffic types.
- **Traffic shaping rules:** You can create shaping policies to apply per-user controls on a per-application basis. This allows throttling of recreational applications such as peer-to-peer file-sharing programs and prioritization of enterprise applications such as Salesforce.com, ensuring that business-critical application performance is not compromised.

For MR devices, the group policy name can be returned in multiple

Technet24

ways, not just through the Filter-ID attribute as with MS and CS switches. The following RADIUS attributes are supported for group policy assignment on MR access points:

- Filter-ID
- Reply Message
- Airspace-ACL-name
- Aruba-User-Role

Wireless Access Control Structure

Before we review the individual features configurable in the **Access Control** tab, it is beneficial to understand its structure. This knowledge makes it easier to locate each feature within the Access Control container.

To access the Access Control settings, navigate to **Wireless > Access Control**. This page gives you access to every configurable SSID in the current network:

- **SSID drop-down list:** These options allow you to switch between SSIDs present in the network to configure them. Figure 11-44 shows the SSID drop-down list.

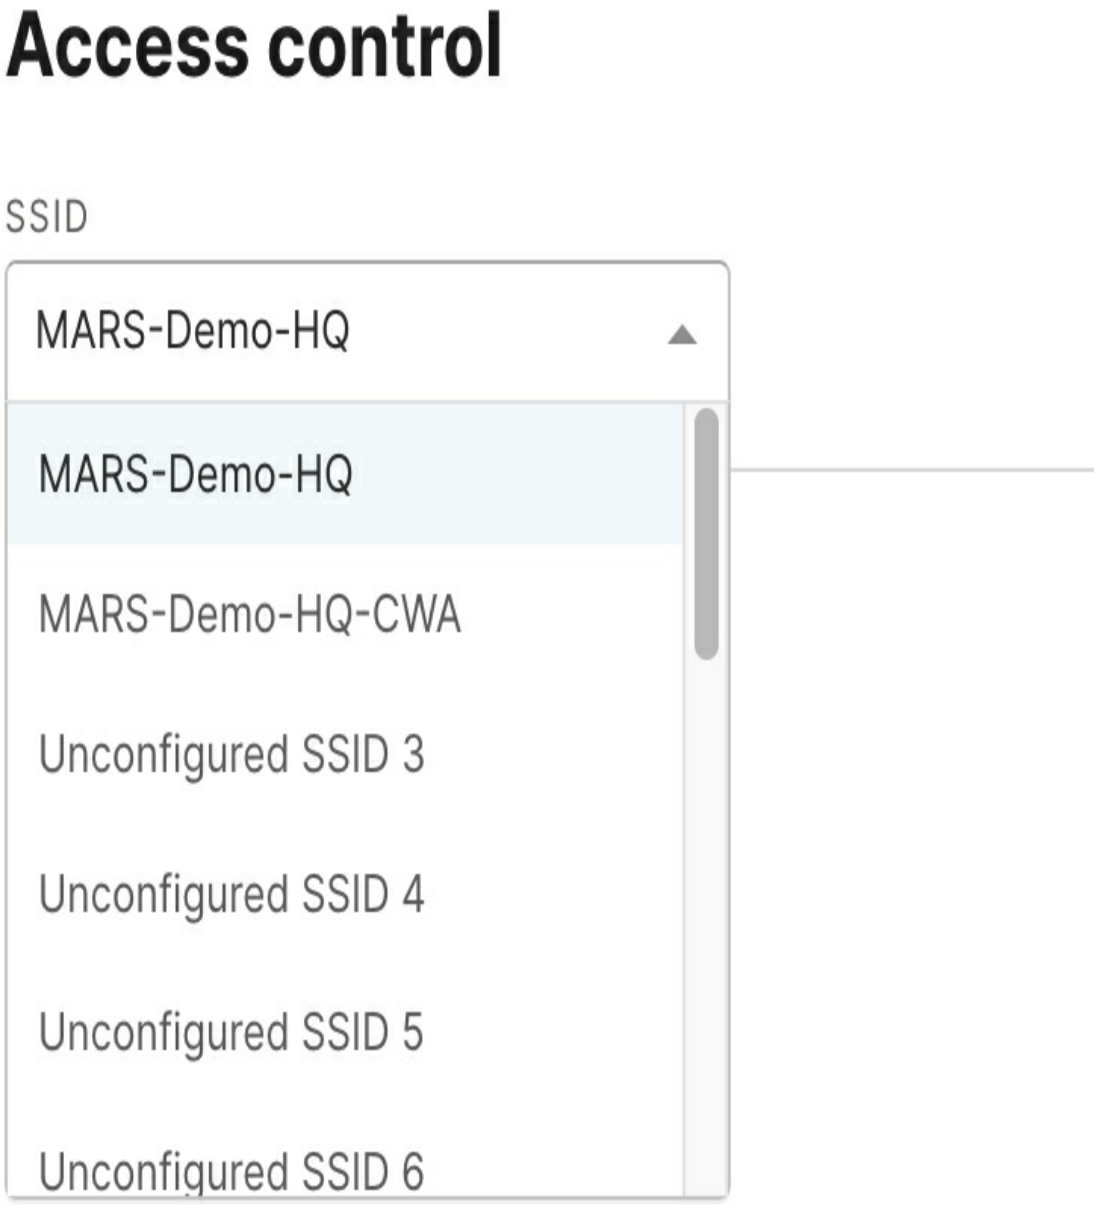


Figure 11-44 *SSID Drop-Down List*

- **General SSID settings:** This section allows you to configure basic SSID parameters such as SSID name, status, and visibility. Figure 11-45 shows the General SSID settings section.

Technet24



Figure 11-45 *General SSID Settings*

- **SSID security settings:** Here you select the type of protection to enable for the SSID. Multiple options are available, but for ISE integration, the primary options of interest are **MAC-Based Access Control, Enterprise With**, and **Identity PSK with**

RADIUS. Figure 11-46 shows the **Security** section.

Technet24

Security *WPA2 Enterprise with 2 RADIUS servers and 2 accounting servers*

Not all security methods are compatible with Cisco ISE splash page

This SSID will not broadcast on the 6 GHz band. Use WPA3 to enable this band.

- Open (no encryption)
  Any user can associate
- Opportunistic Wireless Encryption (OWE)
  Any user can associate with data encryption
- Password
  Users must enter a passphrase to associate
- MAC-based access control (no encryption)
  my RADIUS server
  RADIUS server is queried at association time
- **Enterprise with**
  my RADIUS server
  User credentials are validated with 802.1X at association time
- Identity PSK with RADIUS
  MAC-based Authentication
  RADIUS server is queried at association time to obtain a passphrase for a device based on its MAC address
- Identity PSK without RADIUS
  Devices are assigned a group policy based on its passphrase

Figure 11-46 *SSID Security Settings*

- **SSID Additional settings:** In this section (shown in Figure 11-47), you can specify additional SSID settings such as WPA encryption and mandatory DHCP address assignment for clients. From an ISE perspective, the most important setting is 802.11r, which enables fast roaming for clients that support it. Later in the chapter, we will discuss fast roaming and its impact on integration with ISE in greater detail.



Figure 11-47 *SSID Additional Settings*

- **Splash Page settings:** Meraki wireless supports a wide variety of Splash page options. You can select **None** to completely disable redirection or choose **Cisco Identity Services Engine (ISE) Authentication** when you need to integrate the wireless environment with ISE for Central Web Authentication (CWA). When you enable the ISE option, the MR access point includes the Cisco audit-session ID attribute in RADIUS packets sent to ISE.

  By default, if RADIUS authentication is performed and the Splash page section is not configured to use ISE, the RADIUS packets do not include the audit-session ID attribute. It is important to

Technet24

highlight that for CWA to function correctly with Cisco device types — including MR/CW — the session ID must be included because ISE uses this attribute for session stitching before and after CoA. Also, enabling the ISE option in the Splash settings automatically enables CoA. The **Splash Page** settings section is illustrated in Figure 11-48.

Splash page *Cisco ISE authentication*

Not all splash authentication methods are compatible with WPA2-Enterprise authentication

- None (direct access)
  Users can access the network as soon as they associate
- Click-through
  Users must view and acknowledge your splash page before being allowed on the network
- Sponsored guest login
  Guests must enter a valid sponsor and own email address before being allowed on the network
- Sign-on with
  Meraki Cloud Authentication
  Users must enter a username and password before being allowed on the network
- Sign-on with SMS Authentication
  Users enter a mobile phone number and receive an authorization code via SMS.
  After a trial period of 25 texts, you will need to connect with your Twilio account on the Network-wide settings page.
- **Cisco Identity Services Engine (ISE) Authentication**
  Users are redirected to the Cisco ISE web portal for device posturing and guest access
- Billing (paid access)
  Users choose from various pay-for-access options, or an optional free tier. Only one enabled SSID may be configured to 'Billing'

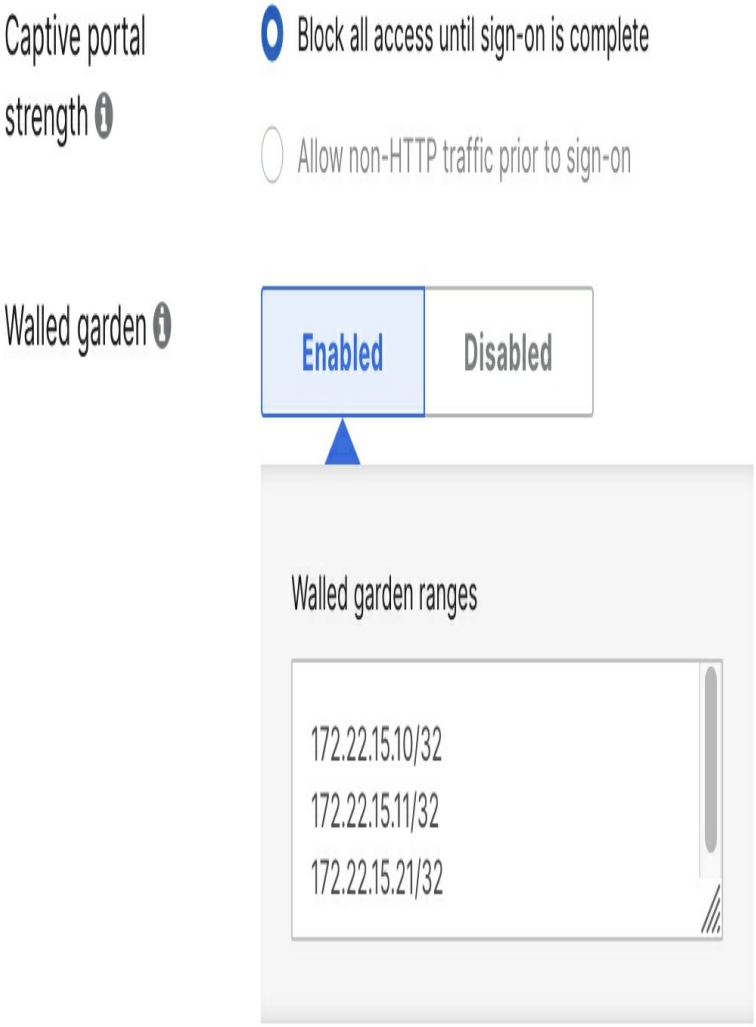
Technet24

Figure 11-48 *Splash Page Settings*

- **Advanced splash settings:** When you enable the Splash page with ISE in this section, you can define walled garden rules to exclude certain destinations from redirection. Figure 11-49 shows advanced Splash settings with walled garden rules that allow traffic to two ISE nodes and an DNS server.

Advanced splash settings

(Captive portal strength, walled garden)



Customize the look and feel of your splash page from Splash page settings.

Technet24

Figure 11-49 *Advanced Splash Settings*

- **SSID RADIUS settings:** RADIUS settings on the SSID are divided into two groups: Basic RADIUS settings and Advanced RADIUS settings.
  - In the Basic settings, you define RADIUS servers and specify main parameters such as Device Sensor, RADIUS testing, and CoA support.
  - Advanced settings allow you to manage options such as EAP and RADIUS timers, Called-Station-ID, and NAS-ID.

  Figure 11-50 shows basic RADIUS settings for an SSID.

RADIUS *2 RADIUS servers, 2 accounting servers - Testing enabled, CoA supported*

RADIUS servers

#	Host IP or FQDN	Auth port	Secret	RadSec	Test	Actions
1	172.22.15.21	1812	••••••••••••	☐	Test	•••
2	172.22.15.11	1812	••••••••••••	☐	Test	•••

Add server 3 max.

RADIUS accounting servers

#	Host IP or FQDN	Acct port	Secret	RadSec	Actions
1	172.22.15.21	1813	••••••••••••	☐	•••
2	172.22.15.11	1813	••••••••••••	☐	•••

Add server 3 max.

Accounting interim interval: 10 minutes

Accounting start delay: 0 seconds

☑ RADIUS Accounting Device Profiling support

☑ RADIUS testing

☑ RADIUS CoA support

RADIUS attribute specifying group policy name: Filter-Id

Technet24

Figure 11-50 *SSID Basic RADIUS Settings*

You can see the **Advanced RADIUS Settings** section in Figure 11-51.

Advanced RADIUS settings

(NAS ID, Called-station-ID, DAS clients, RADIUS timeout, retry count, fallback, EAP timers)

Called-station-ID

#	Category
1	AP MAC address
2	SSID name

Add identifier 4 max.

NAS ID

When 802.11r is enabled the NAS-ID can be a maximum of 48 characters, NAS-ID beyond the maximum will result in network stability issues. We have limited the options allowed to create the NAS-ID to ensure the network remains stable but still allow for flexibility

#	Category
1	AP MAC address
2	SSID number

Add identifier 3 max.

Server timeout 5 second(s)

Retry count 3

RADIUS fallback **Active** Off

EAP timers

EAP timeout 5 second(s)

EAP max retries 5 time(s)

EAP identity timeout 5 second(s)

EAP identity retries 5 time(s)

EAPOL key timeout 500 milliseconds

EAPOL key retries 4 time(s)

Technet24

Figure 11-51 *Advanced RADIUS Settings*

- **IP address assignment:** In this section, you specify how endpoints are bridged into the wired part of the network. From a NAC perspective, the most relevant setting is **RADIUS Override**. When enabled, this option allows MR/CW to assign VLANs retrieved from the RADIUS server in the Access-Accept message. Figure 11-52 shows the IP address assignment section.

Client IP and VLAN *Bridge mode*

Meraki AP assigned (NAT mode)

Clients receive IP addresses in an isolated 10.0.0.0/8 network. Clients cannot communicate with each other, but they may communicate with devices on the wired LAN if the SSID firewall settings permit.

External DHCP server assigned

Meraki devices operate transparently (do not perform NAT or DHCP). Wireless clients will receive DHCP leases from a server on the LAN or use static IPs. Use this for wireless clients requiring seamless roaming, shared printers, and wireless cameras.

Bridged | Tunneled

Layer 3 roaming

Add DHCP option 82: Enabled | Disabled

RADIUS override: Override VLAN tag | Ignore VLAN attribute

RADIUS guest VLAN: Disabled

Bonjour forwarding
Bridge mode only: Enabled | Disabled

VLAN tagging: VLAN ID

#	Access point tags	VLAN ID
	Default	10

+ Add VLAN ID

Assign group policies by device type: Enabled | Disabled

Adaptive policy group
Bridge mode, NAT mode, MCG mode only: 0: Unknown

Technet24

Figure 11-52 *IP Address Assignment Section*

Wireless Access Control: Enterprise with RADIUS

The process of integrating Meraki wireless access control with ISE for enterprise authentication includes the following steps. At this point, it is assumed that all other preconfiguration steps discussed in this book have already been completed. However, to provide a comprehensive checklist, the full list of steps includes the following:

Step 1. Integrate ISE with external identity sources.

Step 2. Add MR/CW devices as NADs with attributes such as device type and location to ensure the most granular policy definitions.

Step 3. Adjust the allowed protocols to include only required authentication methods.

Step 4. Ensure that you complete all required preconfiguration on the Meraki side outside of the SSID access control setup. For example, define group policies and VLAN profiles if you intend to use them as part of the authorization flow. Define a policy set for Meraki MR/CW. This step is optional but highly recommended.

Step 5. If you use the same policy set for both 802.1X and MAB, ensure that an authentication policy for MAB is added. We will discuss this step in more detail in the section dedicated to MAC-based access control.

Step 6. Adjust authentication policies if needed. By default, each policy set includes the default authentication policy, which points to all identity sources.

Step 7. Configure authorization policies.

Step 8. Set up wireless access control.

In the following sections, we will discuss only the necessary steps or those where best practices and recommendations should be implemented.

Creating Authorization Profiles for Your Users

For MR devices, you primarily configure the Filter-ID to return the group policy name and the VLAN ID or name to assign the user to a specific network. You can also enable redirection for features such as Posture, BYOD, and Guest; these features will be explained separately. Figure 11-53 shows an ISE authorization profile that assigns a group policy named Staff-Group-Policy and a VLAN named Staff-VLAN.

Technet24

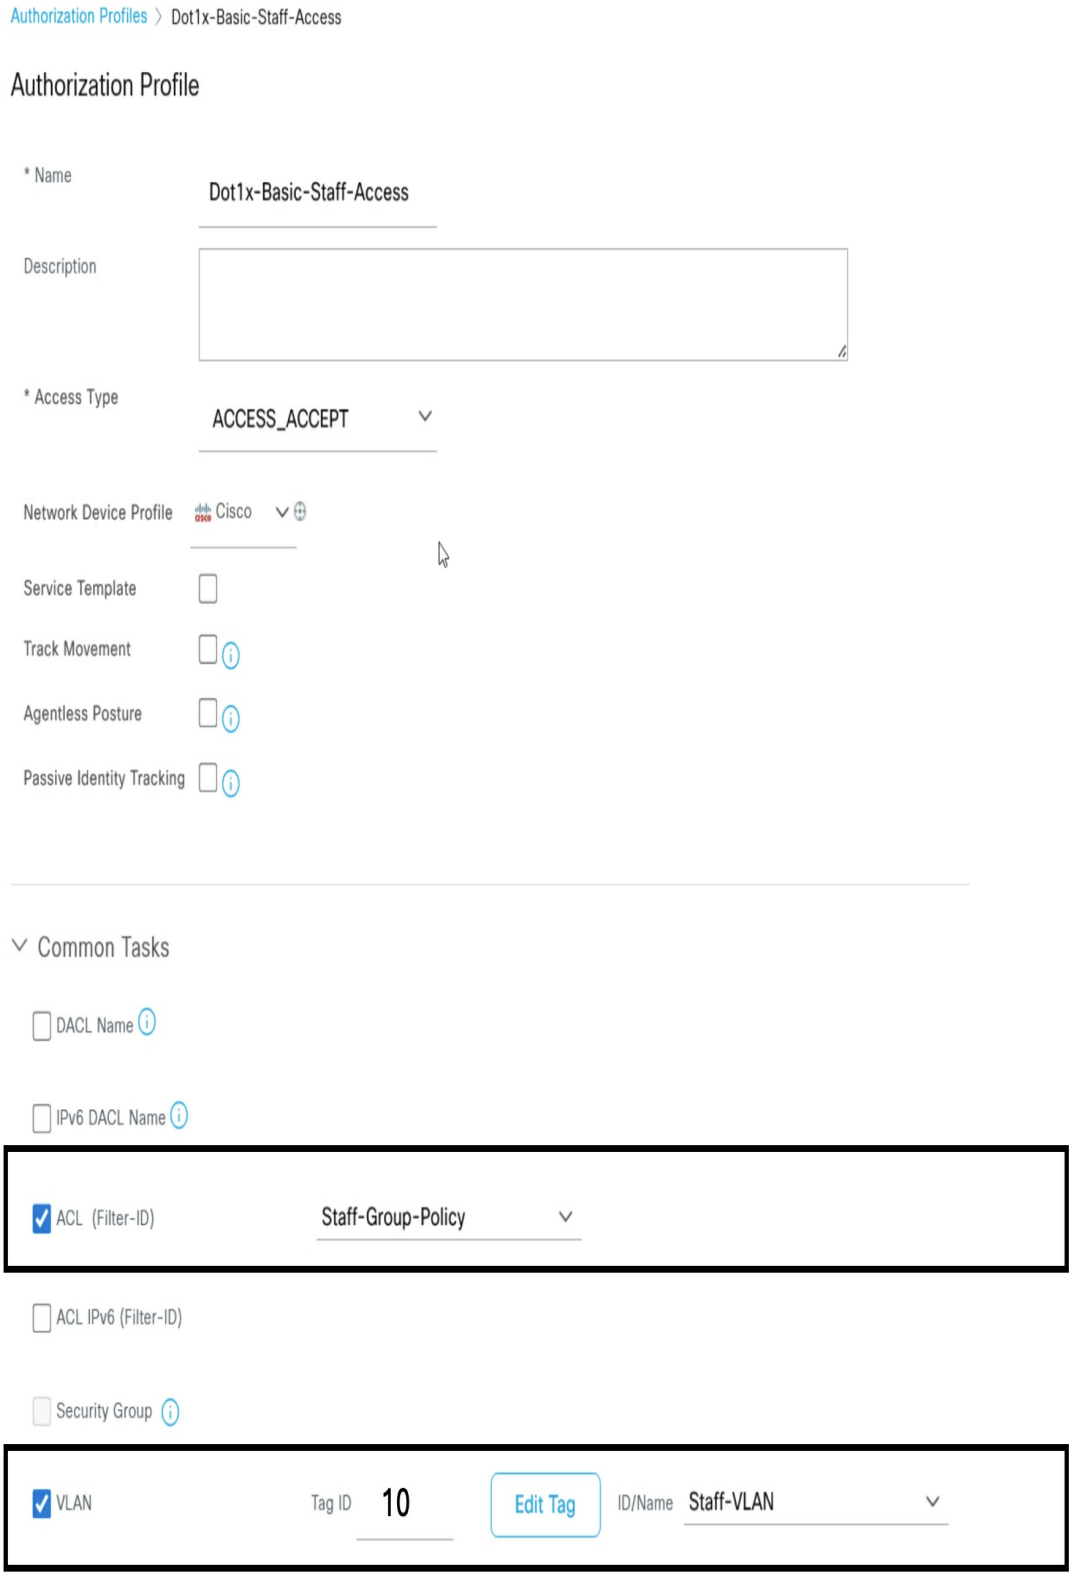

Authorization Profiles
Dot1x-Basic-Staff-Access
Authorization Profile
* Name
Dot1x-Basic-Staff-Access
Description
* Access Type
ACCESS_ACCEPT
Network Device Profile
Cisco
Service Template
Track Movement
Agentless Posture
Passive Identity Tracking
Common Tasks
DACL Name
IPv6 DACL Name
ACL (Filter-ID)
Staff-Group-Policy
ACL IPv6 (Filter-ID)
Security Group
VLAN
Tag ID
10
Edit Tag
ID/Name
Staff-VLAN

Figure 11-53 *ISE Authorization Profile for Wireless 802.1X*

Note

While this example is provided in the "Wireless Access Control" section, you can use the same profile for wired NAC as well.

Figure 11-54 shows objects used in the ISE authorization profile on the Meraki side.

Technet24

VLAN profiles

Network-wide > VLAN profiles > VLAN profiles > Default Profile

Edit profile

Profile name

Default Profile

Iname

Default

Active VLANs

all

VLAN name

+ Add Named VLAN

#	VLAN name	VLAN ID	Adaptive policy group	Actions
1	default	1		
2	Staff-VLAN	10		

Group name

+ Add VLAN Group

#	Group Name	VLAN list	Actions

No record

Network-wide > Group Policies

Group policies

Name	Affecting	Bandwidth	VLAN	Splash	Bonjour	Traffic	Hostname visibility	AMP	Content	Actions
Staff-Group-Policy	0 clients	Default	Default	Default	Default	2 rules applied	Default	Default	Default	Clone ✖

Add a group

Figure 11-54 *VLAN Profile and Group Policy in Network-Wide Settings*

Defining a Policy Set and Policies for Meraki MR

For a mixed environment that consists of Catalyst and legacy Meraki devices, the policy sets shown in Figure 11-55 are used. MR/CW devices have their own policy set, which should be matched based on the device type defined in the NAD settings at the time of NAD creation.

Technet24

Policy Sets

Reset | Reset Policy Set Hit Counts | Save

Status	Policy Set Name	Description	Conditions	Allowed Protocols / Server Sequence	Hits	Actions	View
Search							
✓	Cisco-Wireless		DEVICE·Device Type EQUALS All Device Types#Cisco-Wireless	Default Network Access	0	⚙	›
✓	Cisco-Switches		DEVICE·Device Type EQUALS All Device Types#Cisco-Switches	Default Network Access	0	⚙	›
✓	Cisco-VPN		DEVICE·Device Type EQUALS All Device Types#Cisco-VPN	Default Network Access	0	⚙	›
✓	Meraki-MR-CW		DEVICE·Device Type EQUALS All Device Types#Meraki-MR	Default Network Access	66	⚙	›
✓	Meraki-MS-CS		DEVICE·Device Type EQUALS All Device Types#Meraki-MS-CS	Default Network Access	1209	⚙	›
✓	Meraki-MX		DEVICE·Device Type EQUALS All Device Types#Meraki-MX	Default Network Access	493	⚙	›
✓	Default	Default policy set		Default Network Access	43	⚙	›

Figure 11-55 *Policy Sets*

Figure 11-56 shows a policy suitable for generic 802.1X user access. You can also apply this policy to wired access with slight modifications. To do so, remove the Called-Station ID value and change the authentication type compound condition to Wired_802.1X.

Technet24

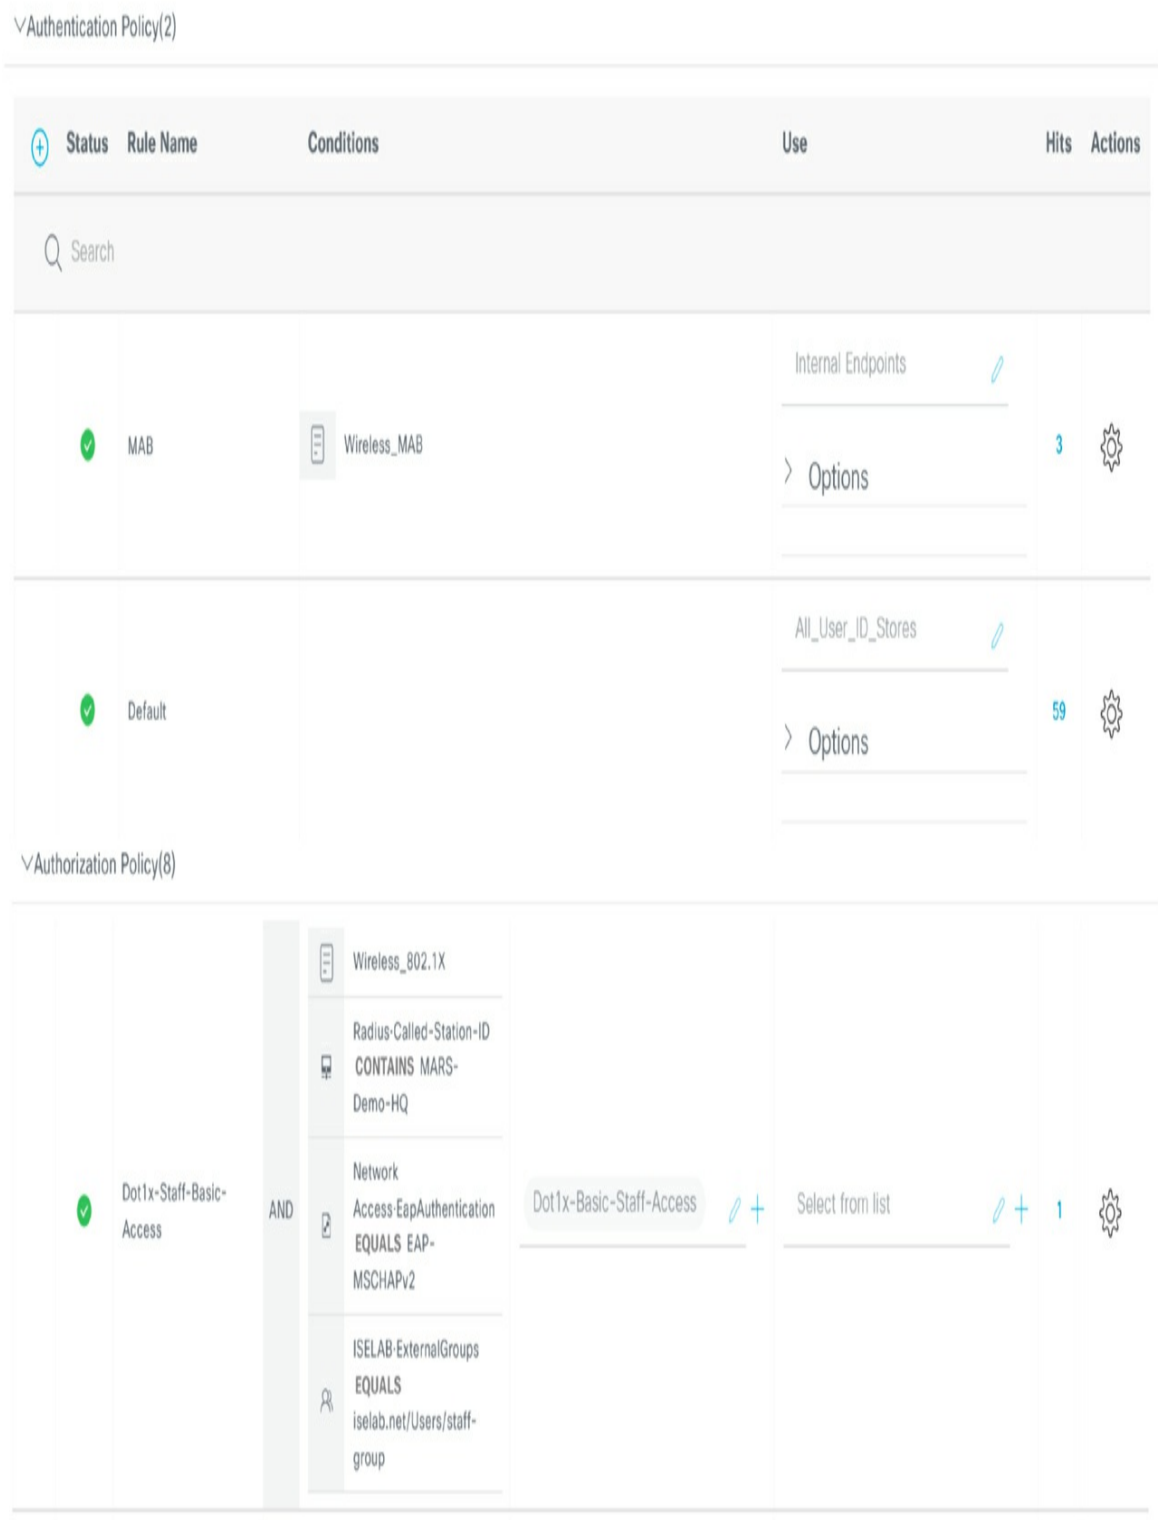


Figure 11-56 *ISE Authorization Policy for MR 802.1X*

The policy in this example includes multiple authorization conditions,

arranged from less specific to more specific, which is considered best practice. For example, the policy begins with conditions that can be identified early in the authentication flow, such as authentication type and SSID name, and proceeds to conditions that require authentication to fully complete before matching, such as Active Directory group membership.

Examine the conditions used in the policy in greater detail. The name on the left before the colon represents the dictionary name for the condition, where applicable:

- **Authentication type compound condition** — The first condition to check verifies whether the request uses wireless 802.1X.
- **Radius: Called-Station-ID** — This condition matches the SSID name by using the **contains** pattern match method. By default, MR access points use the following format for the Called-Station-ID attribute — AP MAC:SSID name.
- **Network Access: EapAuthentication** — If you need to assign different permissions based on EAP type, use this condition. For example, endpoints authenticated with EAP-TLS can receive less restricted access because they are considered corporate devices.
- **ISELAB: ExternalGroups** — This condition matches based on Active Directory group membership. The dictionary name in this context refers to the name of the AD join point.

Setting Up Wireless Access Control

For an 802.1X SSID, you should select **Enterprise with My RADIUS Server** in the SSID security settings, as shown in Figure 11-57.

Technet24



Figure 11-57 *SSID with Enterprise Security*

Enable 802.11r (fast roaming) if required. Fast roaming and its impact on ISE will be discussed later in the chapter. Figure 11-58 shows the SSID Additional Settings section with fast roaming enabled.

Wi-Fi Personal Network (WPN) ⓘ

Enabled | Disabled

WPA encryption ⓘ

WPA3 only ▾

802.11r ⓘ

- ○ Enabled
- ○ Adaptive
- ◉ Disabled

802.11w ⓘ

- ○ Enabled (allow unsupported clients)
- ◉ Required (reject unsupported clients)
- ○ Disabled (never use)

⚠ Traffic from static IP address clients will be blocked on this SSID.

Mandatory DHCP

Enabled | Disabled

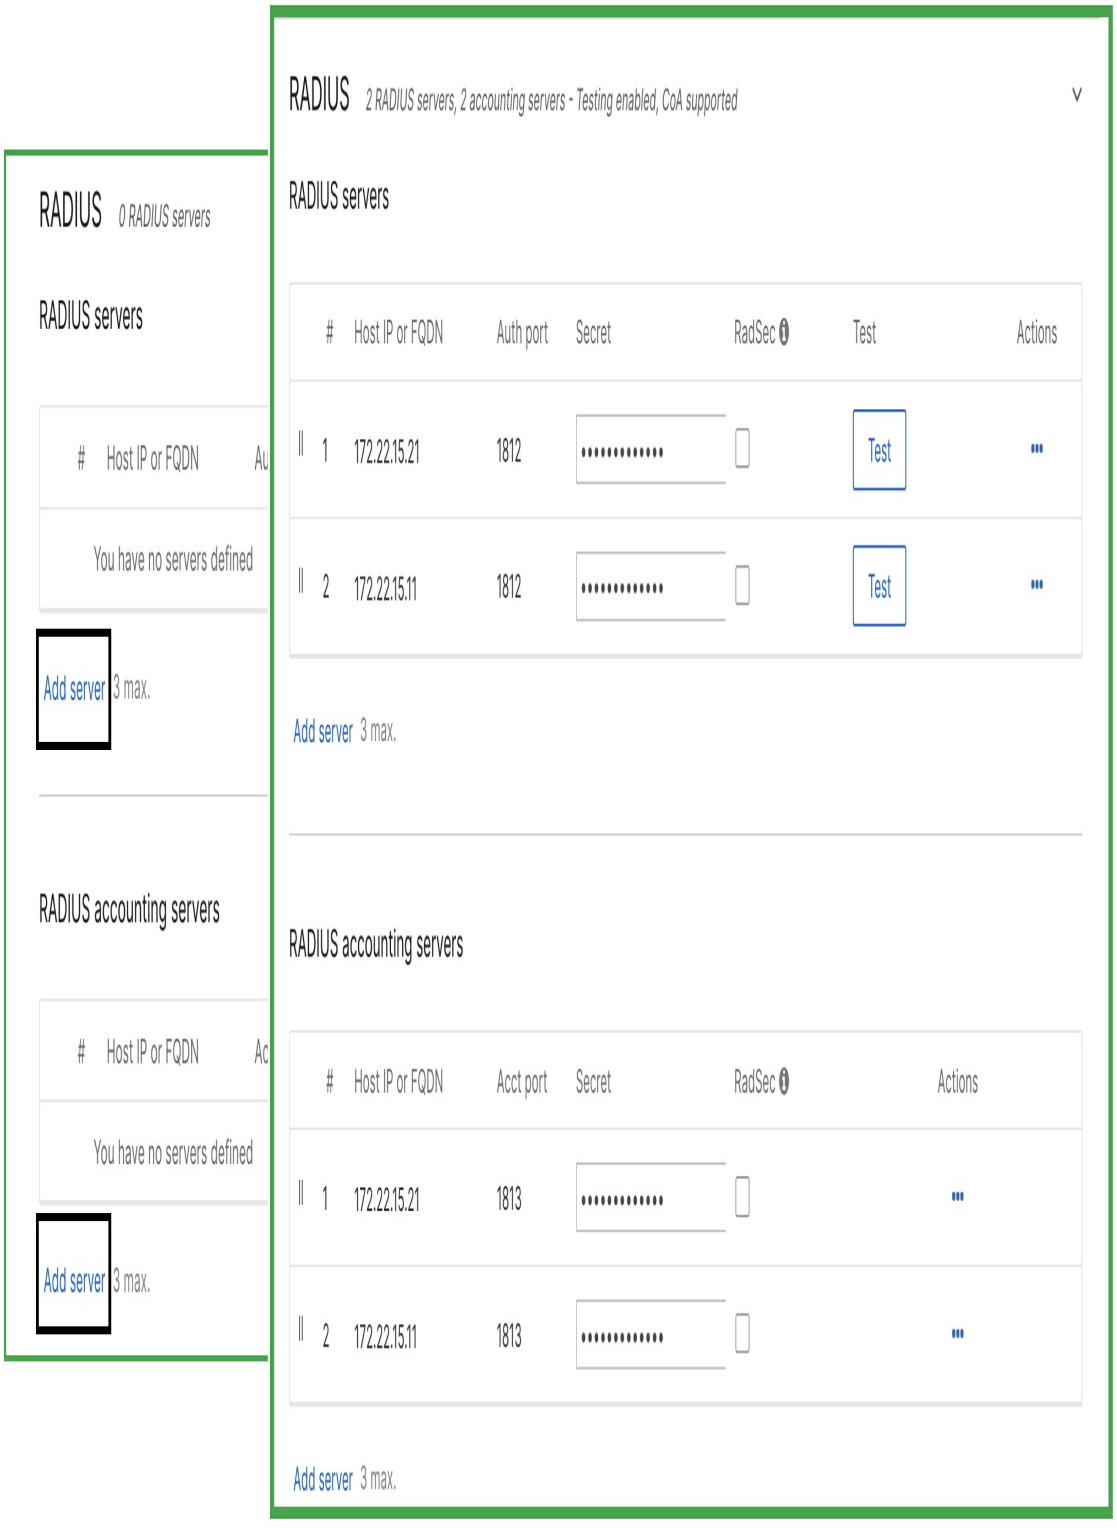
Local Auth Fallback ⓘ

Enabled | Disabled

Technet24

Figure 11-58 *Fast Roaming Enabled*

Now you can add your RADIUS authentication and accounting servers by clicking **Add Server** in the **SSID RADIUS Settings** section, as demonstrated in Figure 11-59.


RADIUS 0 RADIUS servers
RADIUS servers
Host IP or FQDN Au
You have no servers defined
Add server 3 max.
RADIUS accounting servers
Host IP or FQDN Ac
You have no servers defined
Add server 3 max.
RADIUS 2 RADIUS servers, 2 accounting servers - Testing enabled, CoA supported
RADIUS servers
Host IP or FQDN Auth port Secret RadSec Test Actions
1 172.22.15.21 1812 Test
2 172.22.15.11 1812 Test
Add server 3 max.
RADIUS accounting servers
Host IP or FQDN Acct port Secret RadSec Actions
1 172.22.15.21 1813
2 172.22.15.11 1813
Add server 3 max.

Technet24

Figure 11-59 *Configuring RADIUS Servers for SSID*

You can complete the RADIUS general settings configuration by defining the following parameters:

- **Accounting Interim Interval:** Set the interval for sending interim update accounting messages. The default interval is 10 minutes, but consider increasing it to at least 60 minutes to prevent ISE from becoming overloaded with accounting messages.
- **Accounting Start Delay:** This setting introduces a configurable delay of 0 to 60 seconds before the access point sends RADIUS accounting messages to the RADIUS server. This delay ensures that the Framed-IP-Address is included, supporting accurate IP-based access control policies.
- **RADIUS Accounting Device Profiling Support:** You can enable device profiling data collection and transmission to ISE using RADIUS accounting messages. MR/CW supports the following protocols for data collection: HTTP, DHCP, mDNS, DHCPv6, CDP, LLDP, and Device Ecosystem, also known as WiFi Device Analytics on Catalyst 9800 WLCs.
- **RADIUS Testing:** This checkbox enables periodic testing of all defined RADIUS servers using the username meraki_8021x_test. This feature operates identically to RADIUS testing in wired environments.
- **RADIUS CoA Support:** This setting enables CoA support for the SSID. After it is enabled, all authentication requests include the Cisco Audit Session-ID attribute. When the Splash page with ISE is enabled for the SSID, CoA is enabled by default and cannot be disabled.
- **RADIUS Attribute Specifying Group Policy Name:** This drop-

down list allows you to select an attribute in the RADIUS Access-Accept message whose value will be used to identify the group policy.

Figure 11-60 shows these RADIUS settings.

Technet24

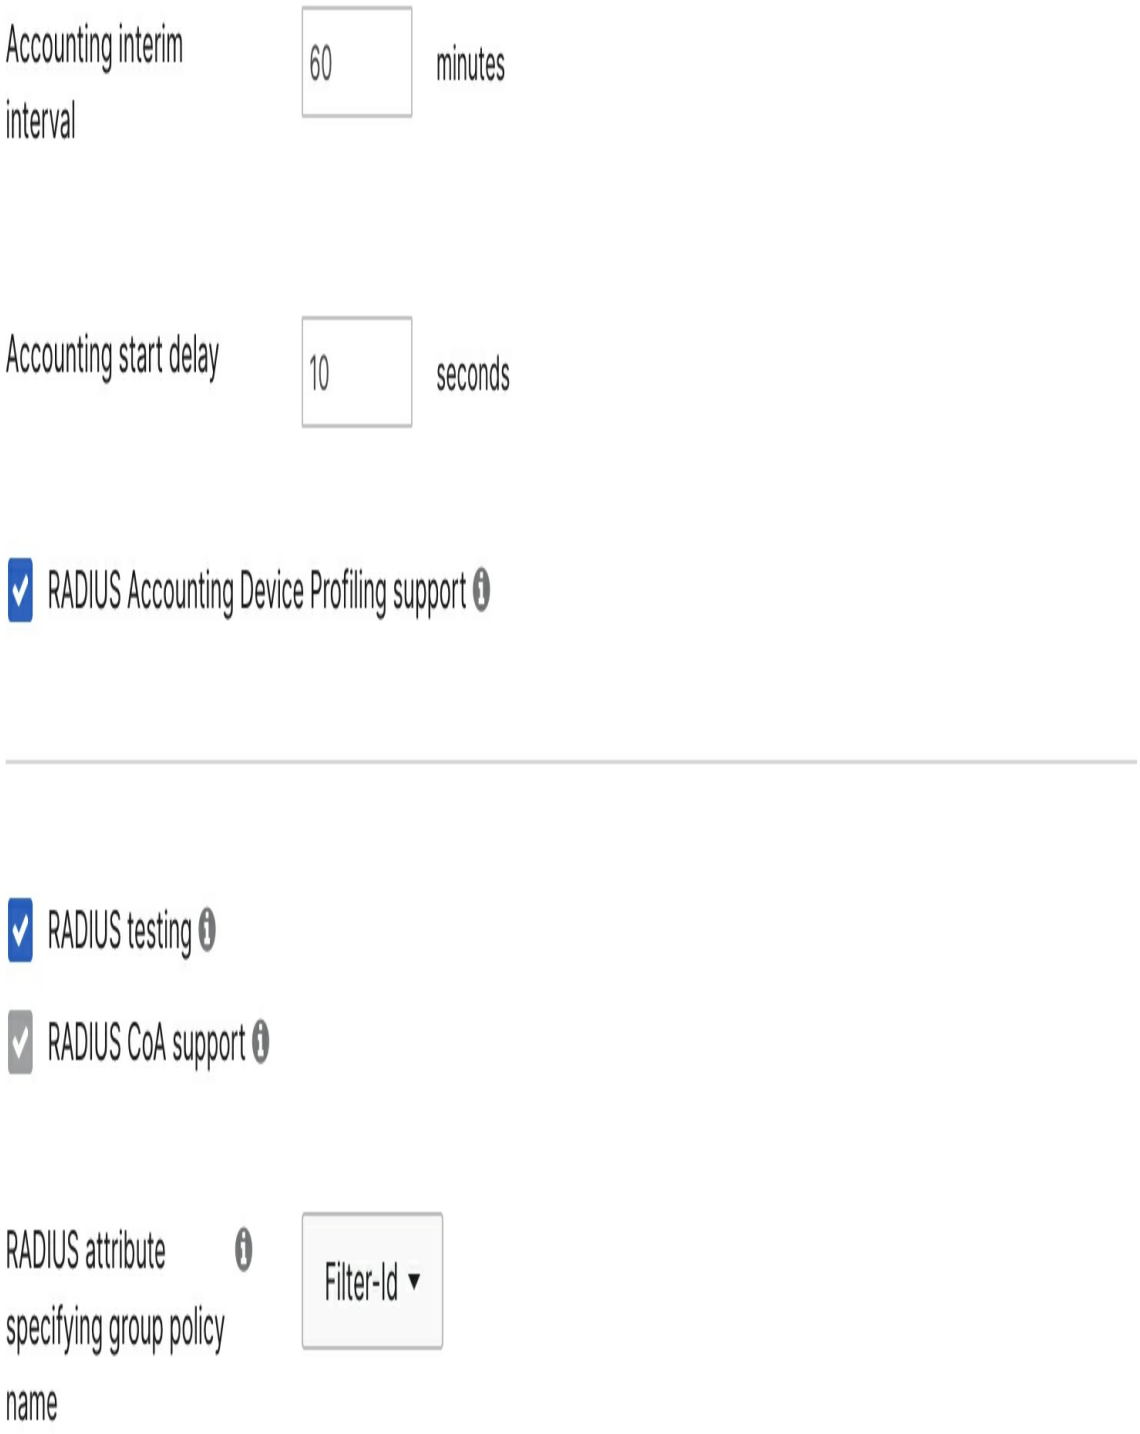


Figure 11-60 *Additional RADIUS Settings*

Now you're ready to configure advanced RADIUS settings. This step is

optional and should be performed only with a clear understanding of how changes in this section may affect the authentication flow. Typically, you should make changes in this section only based on recommendations from Cisco Meraki Support. One setting in this section that deserves attention is RADIUS server fallback, which is off by default. When you set RADIUS server fallback to **Active**, the MR access points send RADIUS messages of type Status-Server to the server with higher priority when it is marked as dead. This process enables fallback to the preferred server when it becomes active again. RADIUS server priority is determined by the order of servers in the access control settings; the top server has priority one, the second has priority two, and the third has priority three.

> **Note**
>
> At the time of writing, ISE does not support RADIUS Status-Server messages defined in RFC 5997. This means there is no need to configure Active fallback when ISE is used as RADIUS server.

The final step in configuring SSID access control is to enable RADIUS VLAN assignment if required. An example of this setting appears in Figure 11-61.

Technet24

External DHCP server assigned

Meraki devices operate transparently (do not perform NAT or DHCP). Wireless clients will receive DHCP leases from a server on the LAN or use static IPs. Use this for wireless clients requiring seamless roaming, shared printers, and wireless cameras.

Bridged | Tunneled

Layer 3 roaming

Add DHCP option 82: Enabled | Disabled

RADIUS override: Override VLAN tag | Ignore VLAN attribute

RADIUS guest VLAN: Disabled

Bonjour forwarding
Bridge mode only: Enabled | Disabled

Figure 11-61 *VLAN RADIUS Override*

Wireless Access Control: MAC-Based Access Control

For MAB-based wireless authentication, you follow the same checklist as before; however, only the settings specific to the MAB authentication method are described in the following list. The main change to the checklist is the absence of a requirement to configure integration with external identity stores on ISE. Instead, as administrator, you typically need to define your endpoints to allow you to connect.

- **Create authorization profiles for your users:** Authorization profiles can be created for different device types in the network that need to establish wireless connections using MAB authentication. Typical profiles include attributes such as Filter-ID to assign the appropriate group policies and VLAN IDs or names to place devices in the correct network segment.
- **Define a policy set and policies for MR/CW:** One important consideration when performing MAB is that, by default, a newly created policy set does not include an authentication policy that checks the endpoint identity store. Figure 11-62 highlights such a policy in the policy set created for MR/CW devices and provides an authorization policy that can be used for generic MAB access.

Technet24

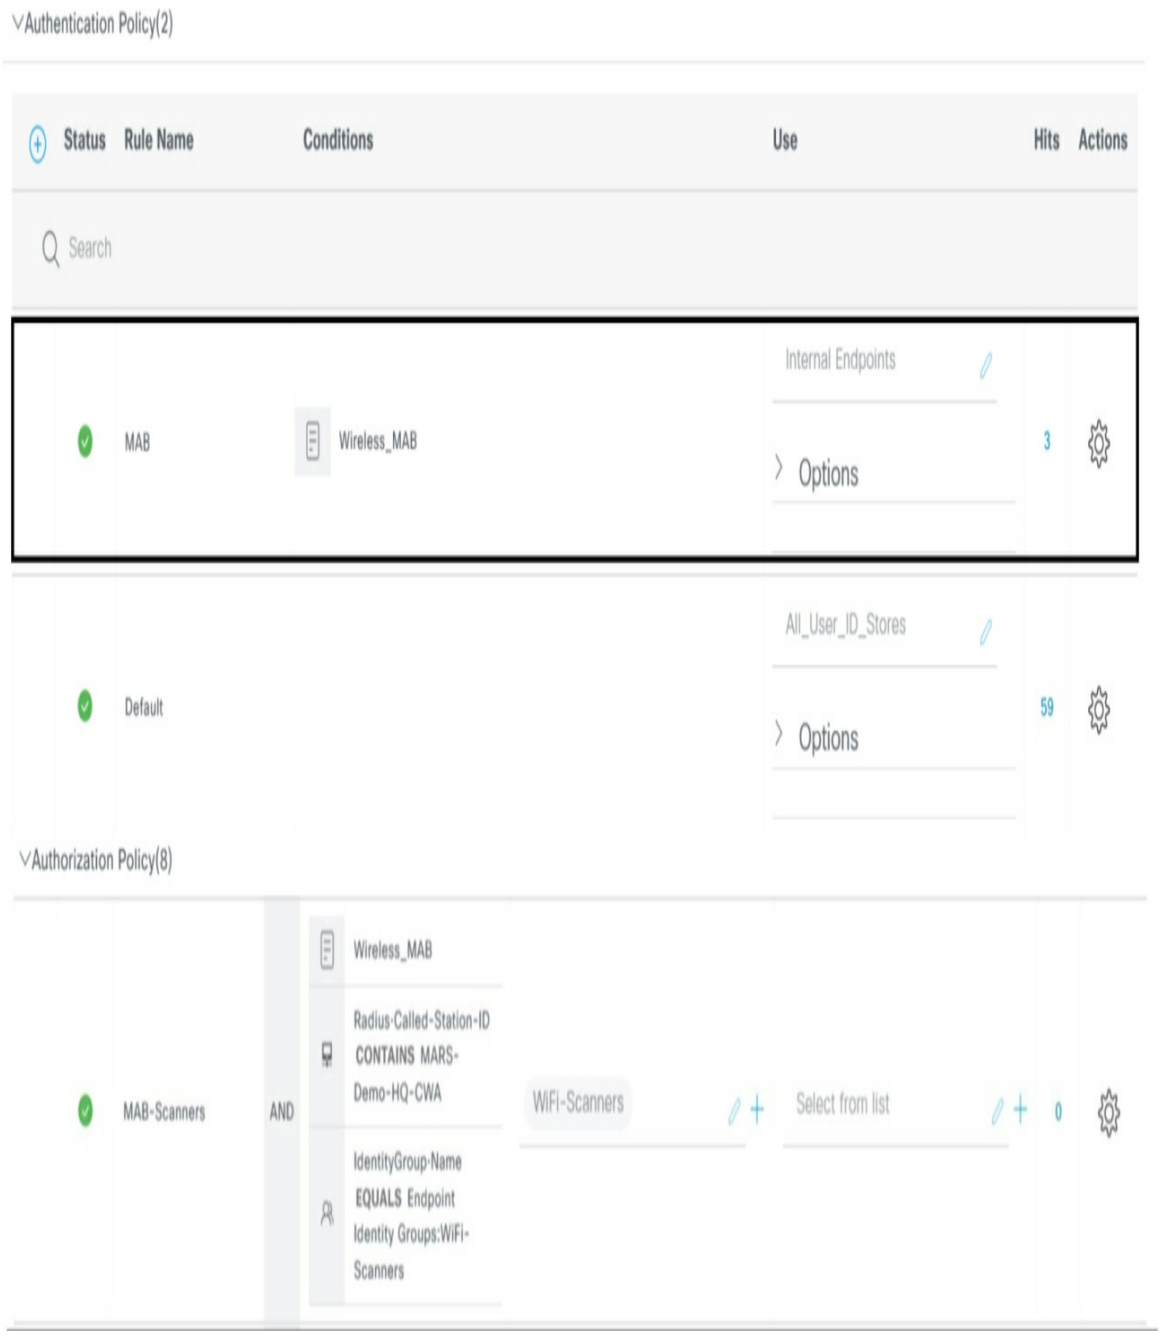


Figure 11-62 *Authentication and Authorization Policy Example for Wireless MAB*

The authorization policy here follows the same approach discussed in the section on 802.1X. The endpoint identity group WiFi-Scanners distinguishes this authentication from other authentication attempts.

Note

Generic MAB is considered insecure, so you should use more secure methods whenever possible. For userless endpoints, consider migrating to iPSK, which offers greater security.

- **Set up wireless access control:** Setting up SSID access control for MAB follows the same approach as for 802.1X. The primary difference is the requirement to enable MAC-based access control (no encryption) in the Security settings, as shown in Figure 11-63.



Figure 11-63 *MAC-Based Access Control (No Encryption)*

Wireless Access Control: Identity PSK with RADIUS

The configuration flow for iPSK closely follows the process described in the section on MAB. The main difference is that the authorization profile for the specific endpoint group or device type must include the attribute Radius:Tunnel-Password:

- **Create authorization profiles for your users:** ISE can return the pre-shared secret key in the Access-Accept message when its value is defined as Radius:Tunnel-Password in the authorization profile. You can add this attribute in the **Advanced Attribute Settings** section of the ISE authorization profile, as shown in

Technet24

Figure 11-64.

Authorization Profiles 〉 Meraki-iPSK

Authorization Profile

* Name: Meraki-iPSK

Description:

* Access Type: ACCESS_ACCEPT

Network Device Profile: Cisco

Service Template

Track Movement

Agentless Posture

Passive Identity Tracking

Advanced Attributes Settings

Radius:Tunnel-Password = Meraki Tag ID 1 Edit Tag

Technet24

Figure 11-64 *Radius Tunnel-Password in ISE Authorization Profile*

- **Define a policy set and policies for MR/CW:** In the iPSK scenario, you can use a policy similar to the one demonstrated in the MAB example. The only differences are the authorization profile and identity group that need to be applied.
- **Set up wireless access control:** Settings required in access control for an iPSK SSID are similar to those explained in the "Enterprise with RADIUS" section. The primary difference is the security method, which must be configured as **Identity PSK with RADIUS**, as shown in Figure 11-65.

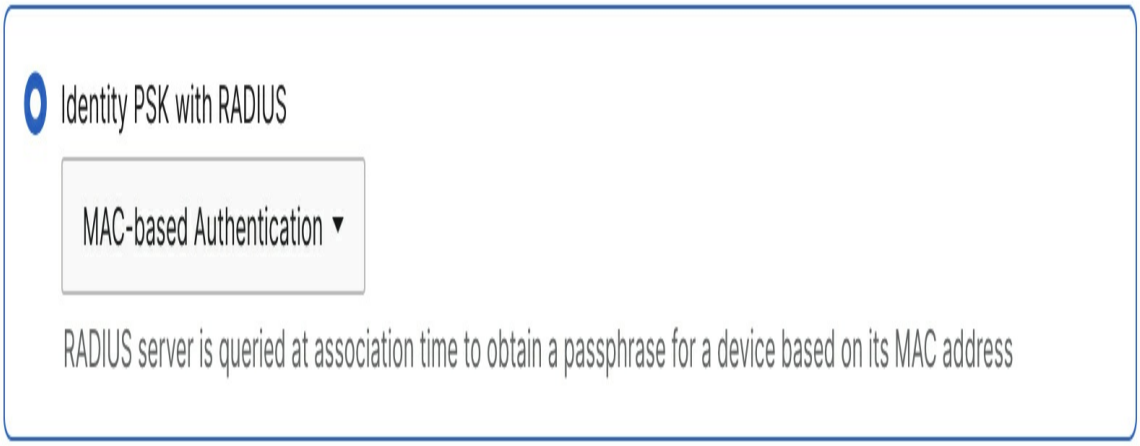


Figure 11-65 *Identity PSK with RADIUS Access Control*

Wireless Access Control: Enabling Splash Page with ISE

Implementing a Splash page with ISE is typically required for one of the following three scenarios:

- **Client Provisioning Portal (CPP):** This scenario is commonly used for endpoint posture assessment. In most cases, it requires 802.1X authentication, which involves two stages. The first stage is an initial redirection to the CPP, during which the endpoint is

provisioned with the posture agent and a posture assessment is performed. The second stage occurs after the endpoint receives full network access following a CoA. In this case, the first authentication results in Splash page redirection, where the Splash page is an ISE CPP.

- **Guest Central Web Authentication (CWA):** This scenario also involves two stages. The first stage typically uses MAB authentication, resulting in a redirect to one of the guest portals hosted on ISE PSN nodes. The second stage provides the necessary level of guest access to the network, which is normally restricted using a group policy.
- **Endpoint Registration (BYOD):** This flow also consists of two authentication stages. The first stage can use either 802.1X or MAB, with a redirect to the BYOD portal for endpoint registration and client provisioning, including certificate provisioning. The second phase occurs after CoA, when the endpoint authenticates as a registered device, most often using EAP-TLS.

This section provides examples for the first and second scenarios. The ISE settings required for CPP, Guest, and BYOD flows have already been covered in detail earlier in this book. This chapter focuses exclusively on Meraki-specific settings and recommendations.

Creating Authorization Profiles for Your Users

When you implement redirection with MR, the access control settings on the SSID manage redirection exceptions. On the ISE side, a redirect-ACL is not required. However, because you must create an authorization profile for the Cisco device type, you need to define a redirect-ACL value; in this example, use Null as the redirect-ACL name. Figure 11-66 shows an ISE authorization profile for CPP redirection.

Technet24

Authorization Profiles 〉 Redirect-Wireless-Meraki-CPP

Authorization Profile

* Name: Redirect-Wireless-Meraki-CPI

Description:

* Access Type: ACCESS_ACCEPT

Network Device Profile: Cisco

Service Template

Track Movement

Agentless Posture

Passive Identity Tracking

Common Tasks

Web Redirection (CWA, MDM, NSP, CPP)

Client Provisioning (Posture) ACL Null Value Client Provisioning Portal (defi

Static IP/Host name/FQDN

Suppress Profiler CoA for endpoints in Logical Profile

Figure 11-66 *ISE Authorization Profile for CPP Redirection*

The CWA authorization profile is similar to the CPP profile, except for the Web Redirection type, which must be configured for CWA. Figure 11-67 shows an authorization profile for guest redirection.

Technet24

Authorization Profiles 〉 Redirect-Wireless-Meraki-Guest

Authorization Profile

* Name: Redirect-Wireless-Meraki-Gue

Description:

* Access Type: ACCESS_ACCEPT

Network Device Profile: Cisco

Service Template

Track Movement

Agentless Posture

Passive Identity Tracking

Web Redirection (CWA, MDM, NSP, CPP)

Centralized Web Auth | ACL: Null | Value: Sponsored Guest Portal (defa

Display Certificates Renewal Message

Static IP/Host name/FQDN

Suppress Profiler CoA for endpoints in Logical Profile

Figure 11-67 *ISE Authorization Profile for CWA Redirection*

Authorization profiles for post-CoA authentication can include the Filter-ID attribute to assign the group policy on the MR side.

In case you decide to use VLANs as part of the authorization, it is important to remember that MAB endpoints do not handle VLAN changes well, because unlike 802.1X supplicants, they lack mechanisms to detect network changes. If you use Splash redirection with ISE for iPSK connections, ensure that both pre- and post-CoA authorization profiles define the same key.

Defining a Policy Set and Policies for MR/CW

Policies for the 802.1X CPP flow follow the same pattern as generic 802.1X access policies. It is important to note that two policies are required: The lower policy handles redirection and uses the additional attribute Session:PostureStatus not equal to compliant, allowing the policy to match all authentication attempts except those with compliant status. The upper policy is identical but uses Session:PostureStatus equal to compliant. Figure 11-68 provides these policies.

Technet24

✓	Dot1x-Dev-Compliant	AND	Wireless_802.1X Radius-Called-Station-ID **CONTAINS** MARS-Demo-HQ ISELAB-ExternalGroups **EQUALS** iselab.net/Users/dev-group Network Access-EapAuthentication **EQUALS** EAP-MSCHAPv2 Session-PostureStatus **EQUALS** Compliant	Dot1x-Dev-Compliant-Access	Select from list	2	
✓	Dot1x-Dev-CPP-Redirect	AND	Wireless_802.1X Radius-Called-Station-ID **CONTAINS** MARS-Demo-HQ Network Access-EapAuthentication **EQUALS** EAP-MSCHAPv2 ISELAB-ExternalGroups **EQUALS** iselab.net/Users/dev-group Session-PostureStatus **NOT_EQUALS** Compliant	Redirect-Wireless-Meraki-CPP	Select from list	7	

Figure 11-68 *CPP Authorization Policies*

For the guest use case, ensure that the MAB authentication policy has the option **If User Not Found** set to **Continue**. This setting allows MAC addresses unknown to ISE to successfully authenticate and match the initial redirect policy. This setting is highlighted in Figure 11-69.

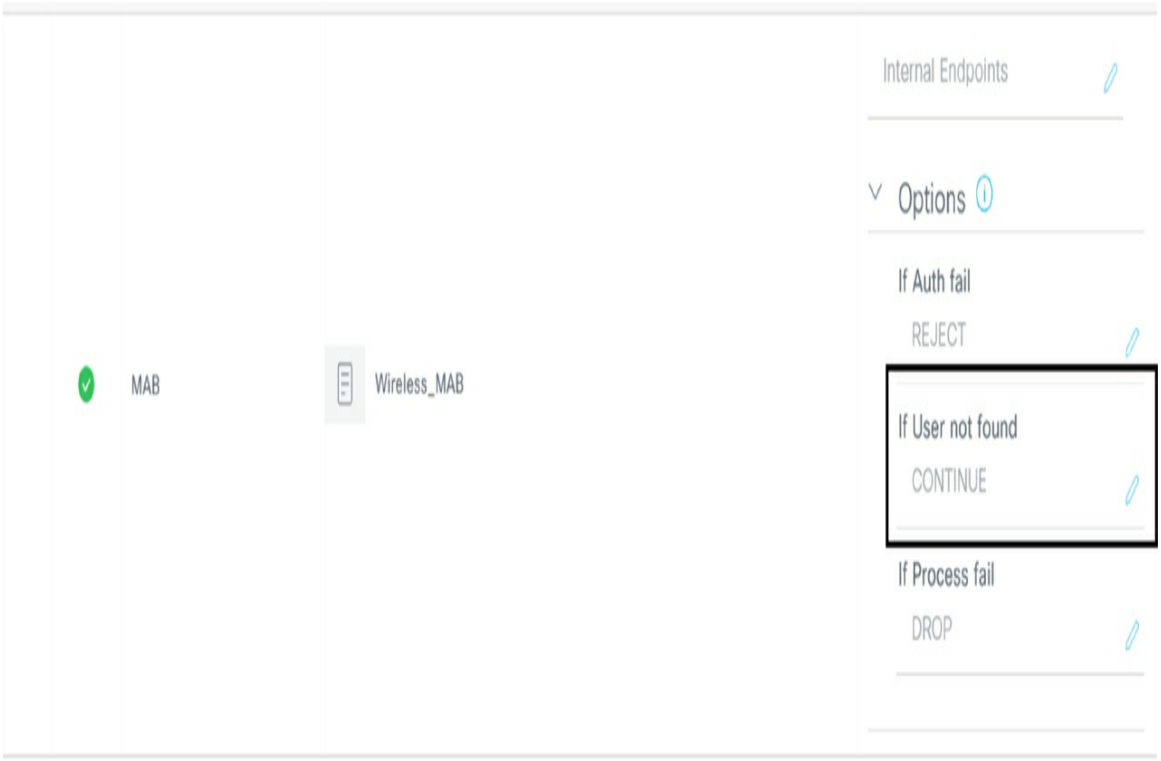


Figure 11-69 *If User Not Found Authentication Policy Option*

As with the CPP scenario, you also need two policies for the guest use case. The lower policy is a catch-all policy that enforces redirection. Do not use this policy for userless devices that require profiling because they will not be able to exit the redirect flow. The upper policy is assigned to the session after CoA, after the guest user has completed the required actions on the guest portal. An example of Guest authorization policies appears in Figure 11-70.

Technet24

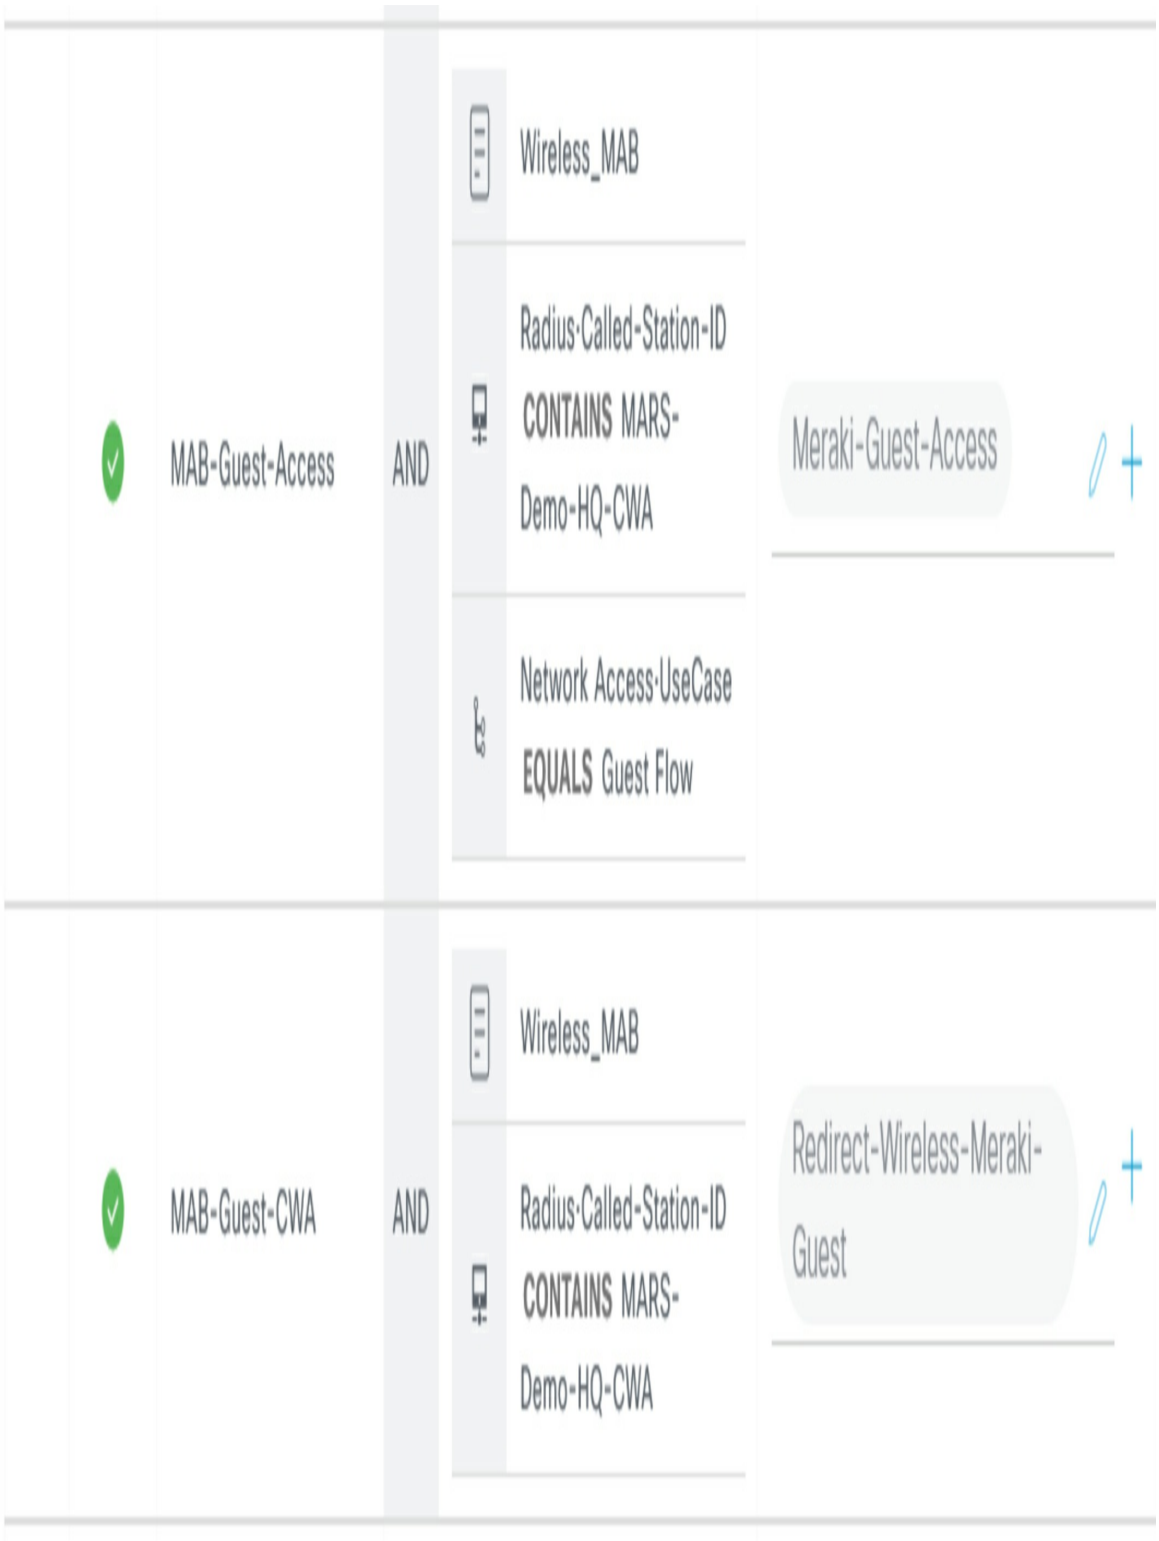


Figure 11-70 *CWA Authorization Policies*

Wireless Access Control: Fast Roaming

Fast roaming is often recommended as a feature to reduce the load on RADIUS servers, especially in highly mobile environments such as university campuses and retail stores. To better understand the impact of fast roaming, it is important to first examine how authentication typically occurs during regular roaming. Figure 11-71 shows the authentication flow when an endpoint roams from AP1 to AP2.

Technet24

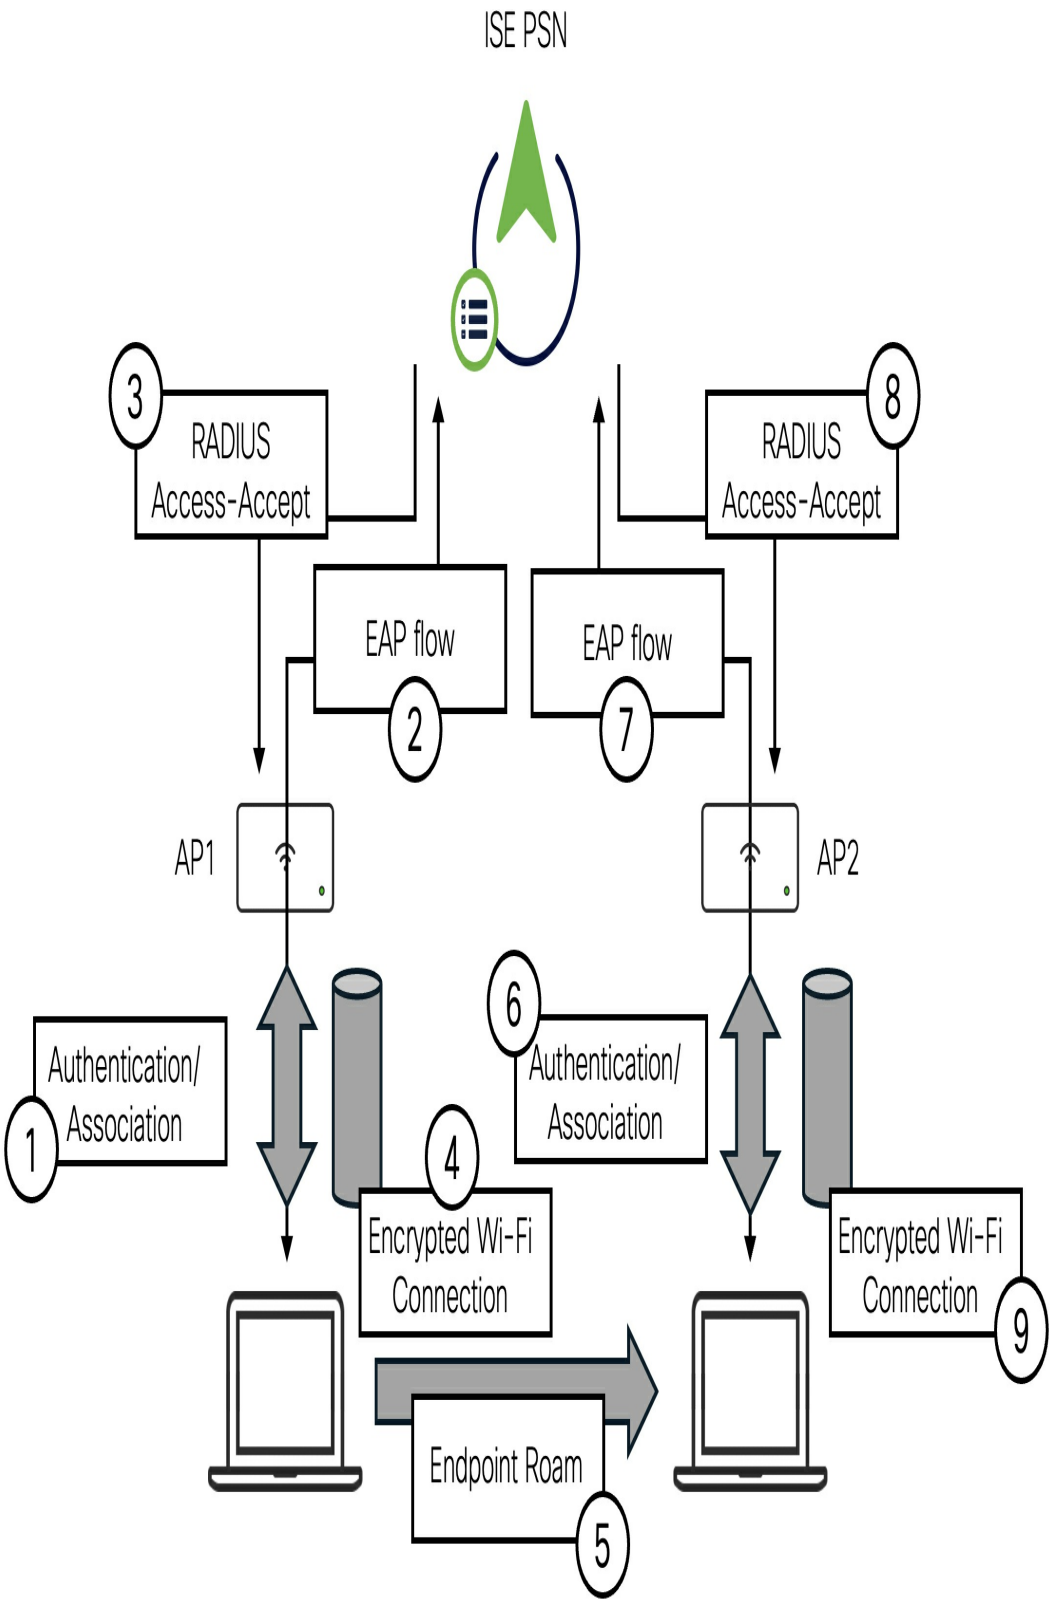

ISE PSN
3
RADIUS
Access-Accept
8
RADIUS
Access-Accept
EAP flow
2
EAP flow
7
AP1
AP2
1
Authentication/
Association
6
Authentication/
Association
4
Encrypted Wi-Fi
Connection
Encrypted Wi-Fi
Connection
9
Endpoint Roam
5

Figure 11-71 *Roaming Example*

Review the steps involved in a standard client roam between two access points:

Step 1. The client performs authentication and associates with AP1.

Step 2. An EAP flow is executed between the supplicant and the authentication server. As a result of a successful exchange, the supplicant receives keying material to protect the Wi-Fi connection.

Step 3. As part of the RADIUS Access-Accept, AP1 also receives the keying material.

Step 4. The Wi-Fi connection between the client and AP1 is now protected.

Step 5. The client roams from AP1 to AP2.

Step 6. The client performs authentication and associates with AP2.

Step 7. Another EAP flow is executed between the supplicant and the authentication server. Upon success, the supplicant again receives keying material to protect the new Wi-Fi connection.

Step 8. As part of the RADIUS Access-Accept, AP2 receives the keying material.

Step 9. The Wi-Fi connection between the client and AP2 is now protected.

When you enable 802.11r, the Meraki dashboard, through its connection to every access point in the network, ensures that keys required to protect Wi-Fi connections are securely distributed among all APs. This results in the fast-roaming flow illustrated in Figure 11-72.

Technet24

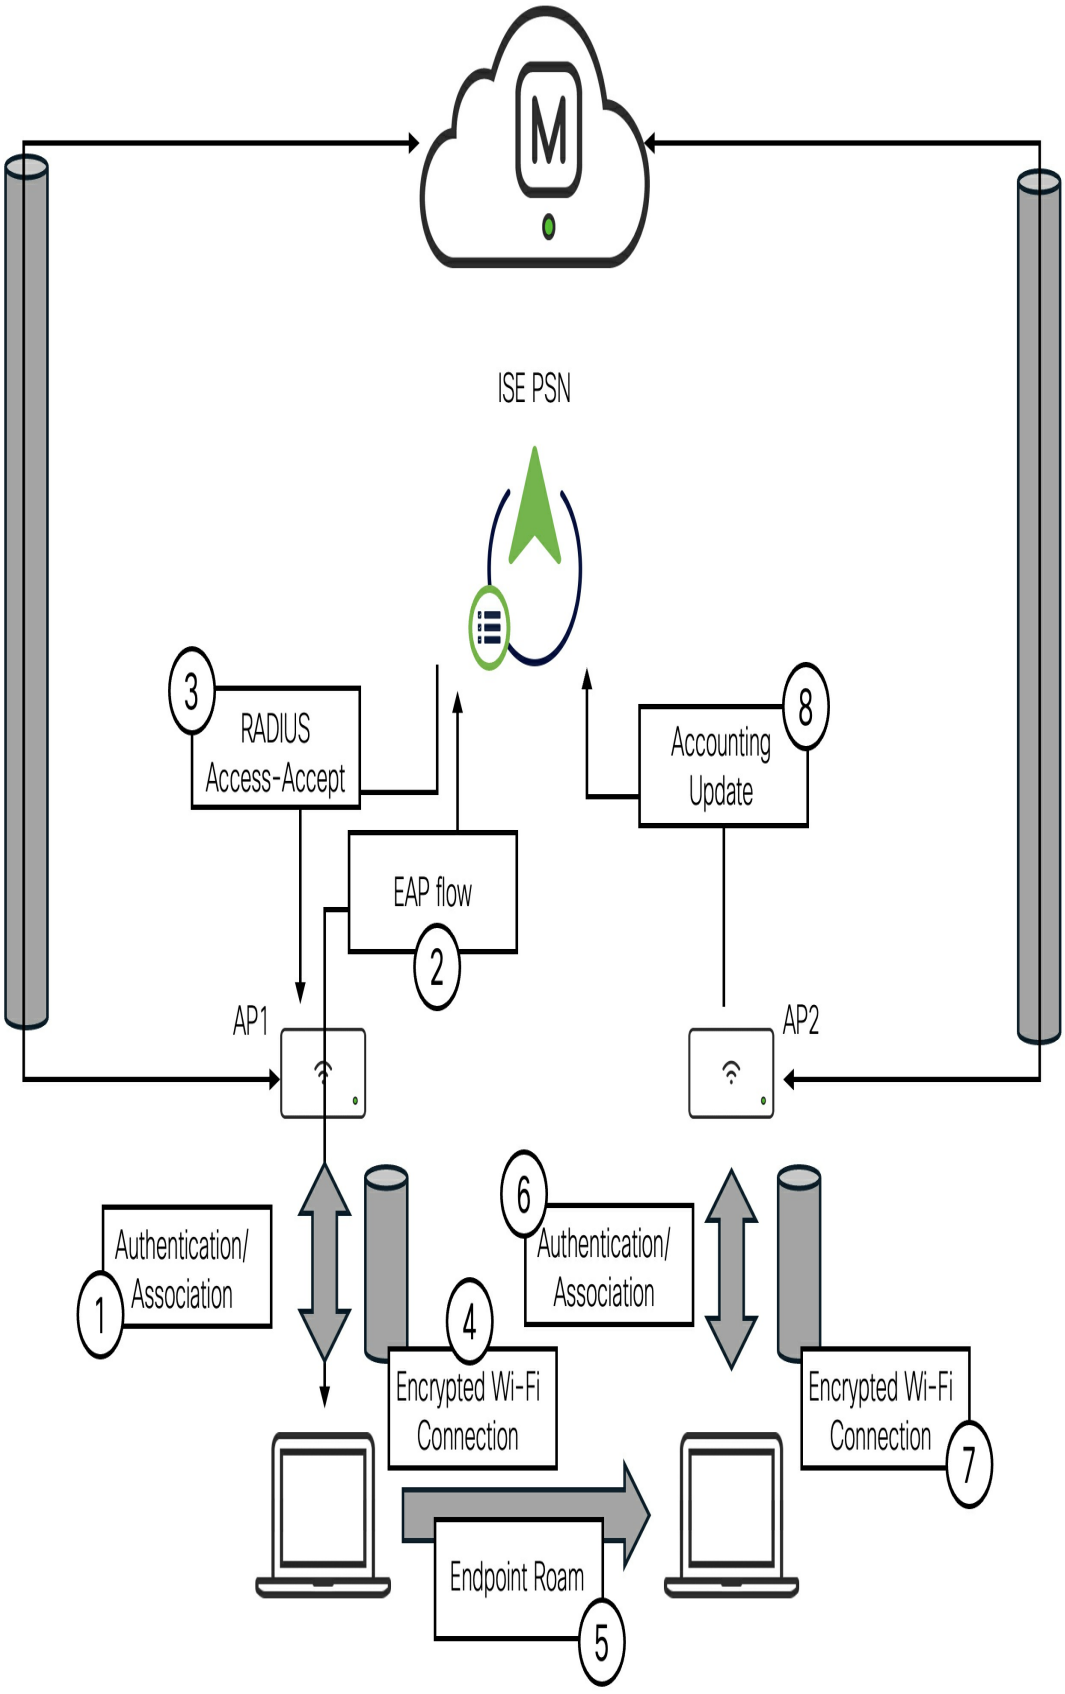

M
ISE PSN
3
RADIUS
Access-Accept
8
Accounting
Update
EAP flow
2
AP1
AP2
Authentication/
Association
1
6
Authentication/
Association
4
Encrypted Wi-Fi
Connection
Encrypted Wi-Fi
Connection
7
Endpoint Roam
5

Figure 11-72 *Fast Roaming flow*

Review the steps involved in client fast roaming between two access points:

Step 1. The client performs authentication and associates with AP1.

Step 2. An EAP flow is executed between the supplicant and the authentication server. As a result of a successful exchange, the supplicant receives keying material to secure the Wi-Fi connection.

Step 3. As part of the RADIUS Access-Accept, AP1 also receives the keying material.

Step 4. The Wi-Fi connection between the client and AP1 is now protected.

Step 5. The client roams from AP1 to AP2.

Step 6. The client performs authentication and associates with AP2.

Step 7. Due to the 802.11r mechanism, the key for the session is already present on AP2, and the connection is immediately protected.

Step 8. AP2 updates ISE to ensure that the session state reflects the new NAD so that any required CoA can reach the correct access point.

Note

Meraki wireless supports fast roaming with the ISE Splash page enabled, starting from the MR32. ISE support for this feature was introduced in ISE 3.3 patch 5, ISE 3.4 patch 2, and ISE 3.5.

Technet24

Meraki Dashboard NAC Verification and Troubleshooting

The Meraki dashboard provides a wide variety of tools for monitoring and troubleshooting. The goal of this section is to focus on monitoring and troubleshooting for NAC scenarios. You will learn about the following tools here:

- **Clients table** (**Network-wide > Clients**)**:** This tool allows you to view all clients discovered on the given network in a single location. The table allows you to modify the view, perform filtering in a wide variety of ways, and investigate different client details.
- **Event log** (**Network-wide > Event log** or **Assurance > Event log**)**:** Here you can investigate event logs for wired and wireless access attempts by using filters such as client MAC address, time of day, network device, and event log type.
- **Wireless Overview** (**Wireless > Overview**)**:** This page provides you with a comprehensive summary of the overall health of your wireless network. From a NAC perspective, this page gives you a summary of failed connection attempts and a general overview of network services health such as RADIUS. The **Overview** page also provides tabs to help you pivot into individual authentication attempts or gain a better understanding of overall authentication and connectivity patterns. Those tabs are **Connection Log** and **Timeline**.
- **CS Cloud native switches Terminal** (**Switching > Switches > (Select your switch) > Tools > Launch Terminal**)**:** The terminal provides a powerful tool for cloud-native switches, allowing you to execute standard show commands to investigate authentication status in great detail.

- **Packet captures** (**Network-wide > Intelligent Capture** or **Assurance > Intelligent Capture**)**:** Dashboard capture tools allow you to troubleshoot authentication issues at the packet level.

Clients Table

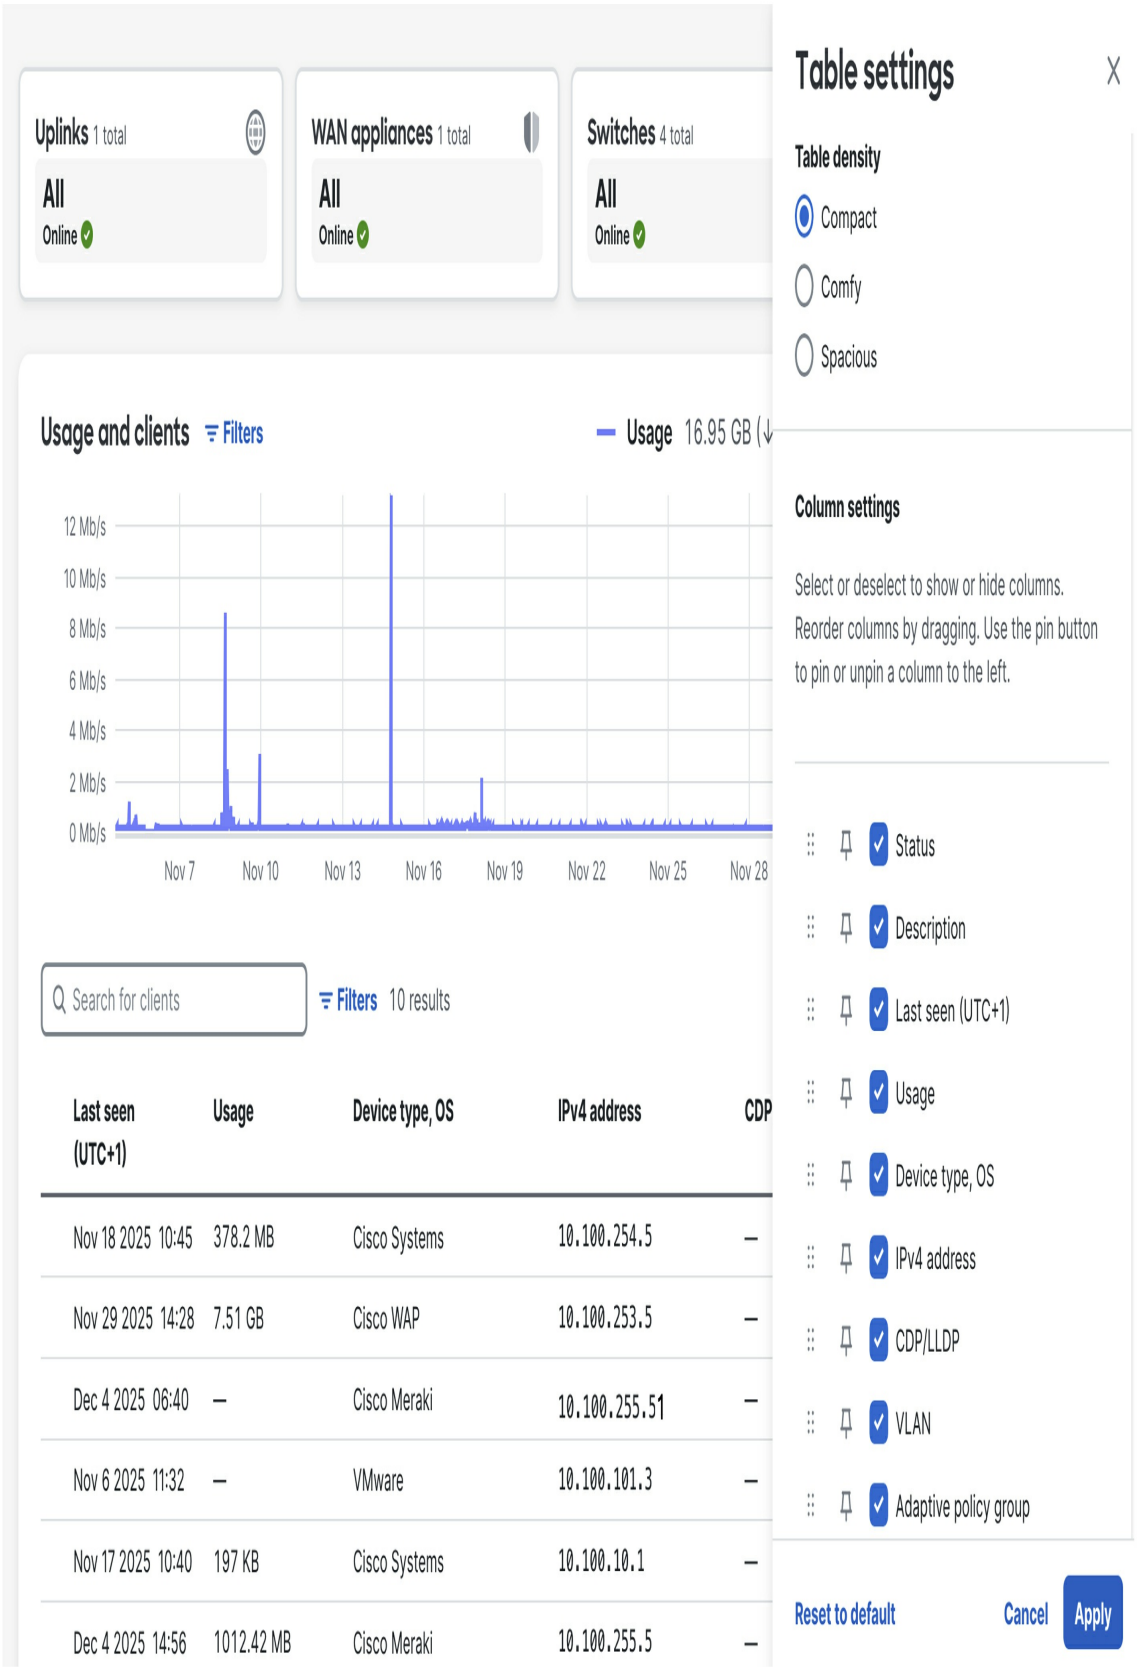
To see all the clients discovered in the given network, navigate to **Network-wide > Clients**. On the right side of the table, you can click the gear icon to add a column. The following columns are recommended to effectively investigate and monitor NAC:

- **Status:** Indicates the type of connection, such as wired or wireless.
- **Description:** Displays the endpoint name or MAC address, which also links to the **Endpoint Details** page.
- **Last Seen:** Shows when the endpoint was last seen on the network.
- **Usage:** Displays data usage for the endpoint.
- **Device Type, OS:** Shows the results of Meraki profiling for the endpoint.
- **IPv4 Address:** Displays the address of the endpoint.
- **CDP/LLDP:** Provides information collected about the endpoint over CDP or LLDP.
- **VLAN:** Indicates the VLAN assigned to the endpoint.
- **Adaptive Policy:** Displays the SGT tag assigned to the endpoint.
- **Named VLAN:** Indicates the named VLAN assigned to the endpoint, if any.
- **User:** Displays the username seen during authentication for this endpoint.

Technet24

- **802.1X Policy:** Shows the group policy assigned as a result of the 802.1X process.
- **Connected To:** Indicates where the endpoint is currently connected.
- **Policy:** Displays the group policy assigned to the endpoint.

Figure 11-73 shows how to adjust columns in the client table.


Uplinks 1 total
All
Online
WAN appliances 1 total
All
Online
Switches 4 total
All
Online
Usage and clients
Filters
Usage 16.95 GB
12 Mb/s
10 Mb/s
8 Mb/s
6 Mb/s
4 Mb/s
2 Mb/s
0 Mb/s
Nov 7
Nov 10
Nov 13
Nov 16
Nov 19
Nov 22
Nov 25
Nov 28
Search for clients
Filters 10 results
Last seen (UTC+1)
Usage
Device type, OS
IPv4 address
CDP
Nov 18 2025 10:45 378.2 MB Cisco Systems 10.100.254.5 –
Nov 29 2025 14:28 7.51 GB Cisco WAP 10.100.253.5 –
Dec 4 2025 06:40 – Cisco Meraki 10.100.255.51 –
Nov 6 2025 11:32 – VMware 10.100.101.3 –
Nov 17 2025 10:40 197 KB Cisco Systems 10.100.10.1 –
Dec 4 2025 14:56 1012.42 MB Cisco Meraki 10.100.255.5 –
Table settings
Table density
Compact
Comfy
Spacious
Column settings
Select or deselect to show or hide columns. Reorder columns by dragging. Use the pin button to pin or unpin a column to the left.
Status
Description
Last seen (UTC+1)
Usage
Device type, OS
IPv4 address
CDP/LLDP
VLAN
Adaptive policy group
Reset to default
Cancel
Apply

Technet24

Figure 11-73 *Adjusting Columns in the Client Table*

A detailed client view can provide more information about the specific endpoint and the user connected behind it. To open this view, click the value in the **Description** column of the client table. Figure 11-74 shows the detailed client view, with key information important for NAC monitoring and troubleshooting highlighted.



Current | Offline | Wireless | Dec 04, 2025, 17:27 (UTC +1)

Client connections



Client details

Status	Currently connected since Dec 4 17:26 View location
SSID	MARS-Demo-HQ
Access point	HQ-AP-1 Topology
User	bob@iselab.net (802.1X login)
Signal	67db Channel 11 2.4GHz
Device type, OS	TP-LINK Technologies
Device capable Wi-Fi standards	802.11n - 2.4 GHz, 802.11r Capability details
Tools	History Packet capture Disconnect client Ping

IPv4 address	10.100.10.10 Edit
IPv6 address (link-local)	fe80:0:0:0:a79c:2841:2989:748e
MAC address	50:3e:aa:da:da:69
VLAN ID	10
Port forwarding	None Edit
1:1 NAT IPs	None
Notes	None Edit

Device policy	Normal Edit
802.1X policy	Staff-Group-Policy
Bandwidth	Unlimited
Layer 3 firewall	0 rules
Layer 7 firewall	0 rules

Technet24

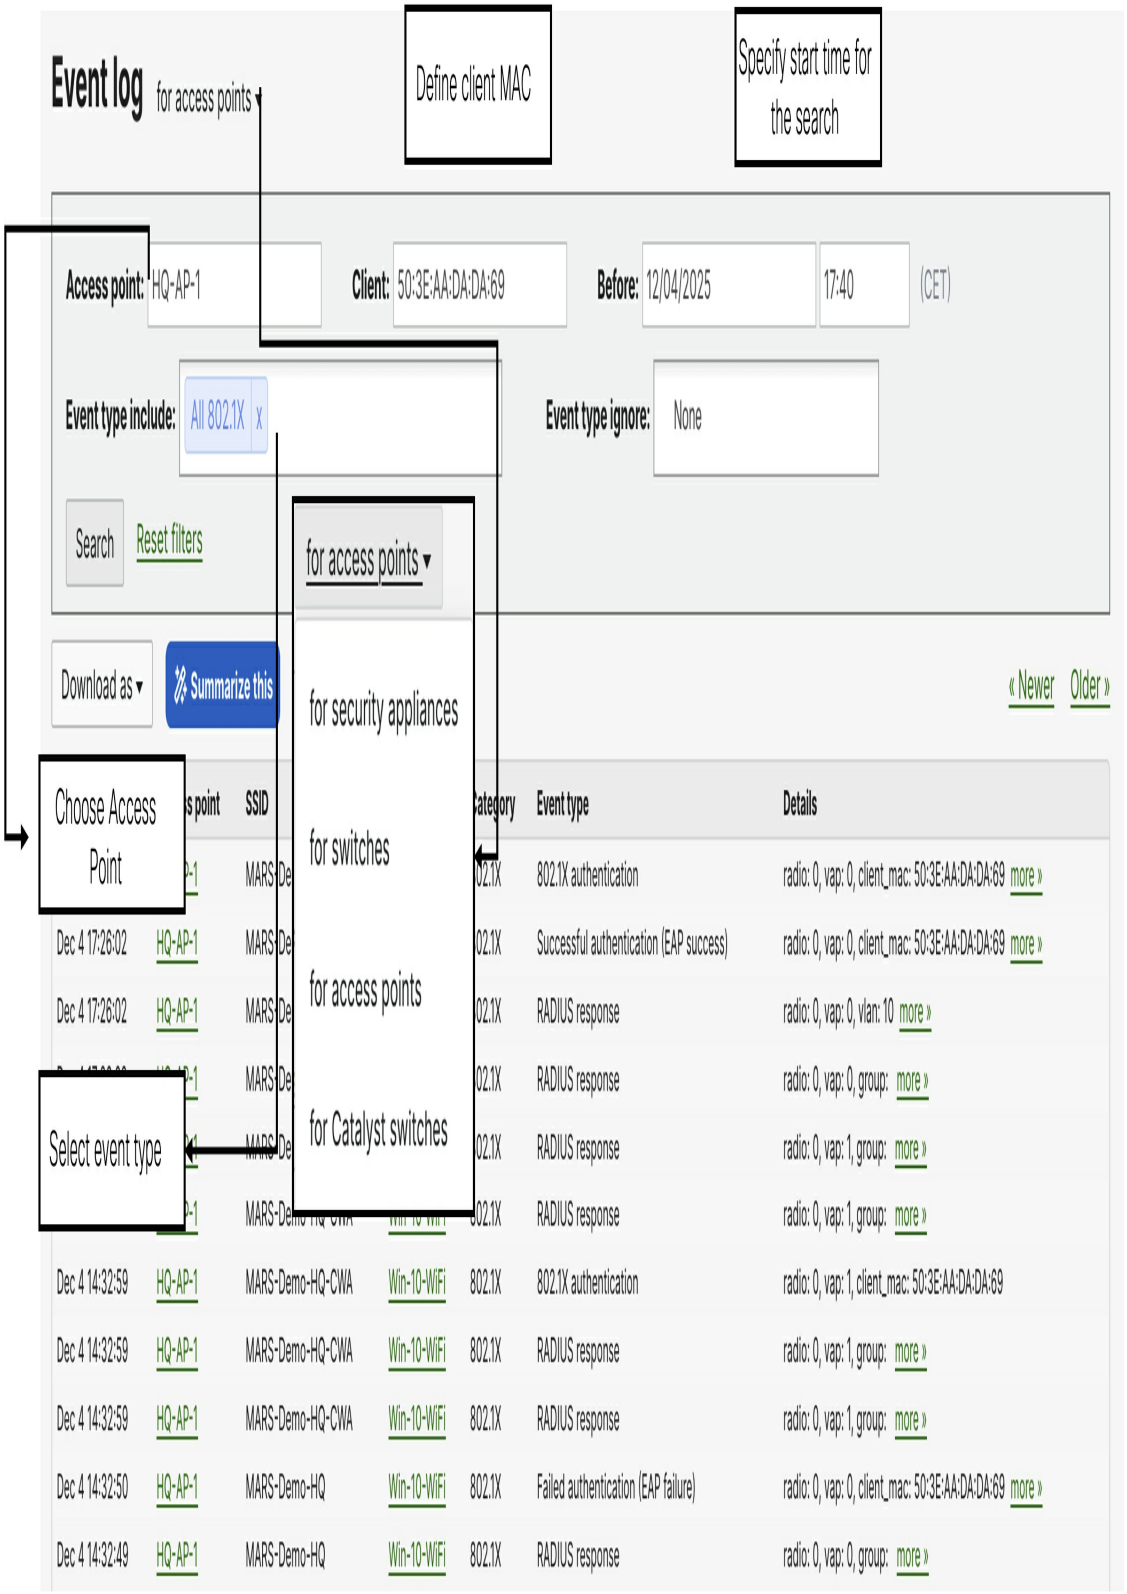
Figure 11-74 *A Client Detailed View*

Event Log

To check detailed event logs in a specific network, navigate to **Network-wide > Event Log** or **Assurance > Event Log**. At the top of the page, you need to select the **Event Log Source**, such as **For Access Points**, **For Switches**, or **For Catalyst Switches**. After selecting the **Device Type**, you can narrow your investigation by selecting an individual switch or access point. To refine your search further, specify the endpoint MAC address in the **Client** section. You can also filter information based on timestamp and event type. The most common event type for investigating authentication problems is 802.1X All. Figure 11-75 shows the **Event Log** page with explanations for its different parts.

Note

The **Event log** page UI may differ for various network devices.


Event log
for access points
Define client MAC
Specify start time for the search
Access point: HQ-AP-1
Client: 50:3E:AA:DA:DA:69
Before: 12/04/2025
17:40
(CET)
Event type include: All 802.1X
Event type ignore: None
Search
Reset filters
for access points
for security appliances
for switches
for access points
for Catalyst switches
Download as
Summarize this
« Newer
Older »
Choose Access Point
Select event type
SSID
Event type
Details
Dec 4 17:26:02 HQ-AP-1 MARS-De 802.1X Successful authentication (EAP success) radio: 0, vap: 0, client_mac: 50:3E:AA:DA:DA:69 more »
Dec 4 17:26:02 HQ-AP-1 MARS-De 802.1X RADIUS response radio: 0, vap: 0, vlan: 10 more »
802.1X RADIUS response radio: 0, vap: 0, group: more »
802.1X RADIUS response radio: 0, vap: 1, group: more »
802.1X RADIUS response radio: 0, vap: 1, group: more »
Dec 4 14:32:59 HQ-AP-1 MARS-Demo-HQ-CWA Win-10-WiFi 802.1X 802.1X authentication radio: 0, vap: 1, client_mac: 50:3E:AA:DA:DA:69
Dec 4 14:32:59 HQ-AP-1 MARS-Demo-HQ-CWA Win-10-WiFi 802.1X RADIUS response radio: 0, vap: 1, group: more »
Dec 4 14:32:59 HQ-AP-1 MARS-Demo-HQ-CWA Win-10-WiFi 802.1X RADIUS response radio: 0, vap: 1, group: more »
Dec 4 14:32:50 HQ-AP-1 MARS-Demo-HQ Win-10-WiFi 802.1X Failed authentication (EAP failure) radio: 0, vap: 0, client_mac: 50:3E:AA:DA:DA:69 more »
Dec 4 14:32:49 HQ-AP-1 MARS-Demo-HQ Win-10-WiFi 802.1X RADIUS response radio: 0, vap: 0, group: more »

Technet24

Figure 11-75 *Event Log Page*

Wireless Overview

Wireless Overview provides a powerful tool to assess the health of your wireless network at a glance. The page contains multiple tabs, with the most important for NAC use cases described here:

- **Overview:** This tab displays a summary of your network health status.
- **Connection Log:** This tab presents a summary table of all connection attempts in the network, regardless of SSID or access point. You can filter the view to focus on specific Wi-Fi networks, access points, or clients.
- **Timeline:** This tool shows how connectivity changes over time for endpoints present in the network.

The last two tools are also available at the access point level.

An example of the **Wireless Overview** tab appears in Figure 11-76. Every hyperlink on the page acts as a filter and can help you narrow your investigation further.

Wireless overview

Last day | All SSIDs | New

Connection health



Failed clients
0/1 No change

Time to connect
0.5 s 0.14 s
Expected < 5 s

Roaming
N/A

Performance health



Latency
94.17 ms 55.51 ms
Expected < 60 ms

Packet loss
1% No change
Expected < 10%

Signal quality (SNR)
50 dB No change
Expected > 27 dB

Network service health



RADIUS success
40% -60 %

DHCP success
86% -14 %

DNS success
100% No change

Technet24

Figure 11-76 *Wireless Overview Tab*

For example, by selecting **RADIUS Success** in the **Network Service Health** section, you access a filtered view that gives you a better understanding of RADIUS statistics. From this view, you can narrow your search even further by using one of the available hyperlinks. An example of the **RADIUS Success** view appears in Figure 11-77.

Network service health



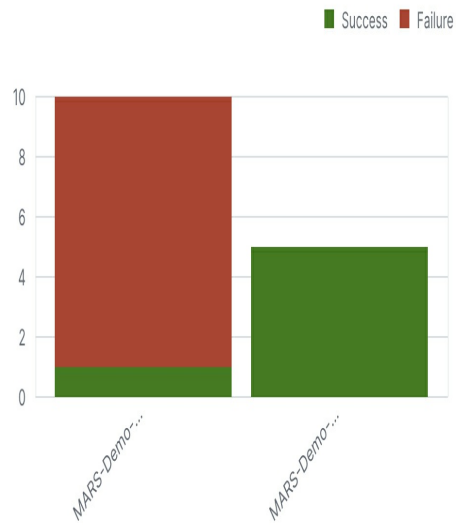
RADIUS success
40% ↘ -60 %

DHCP success
86% ↘ -14 %

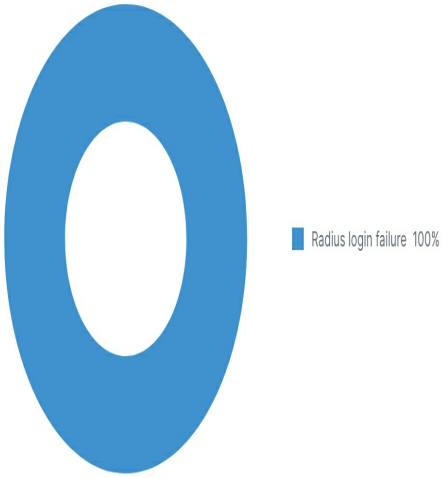
DNS success
100% No change

RADIUS Server IP All IPs

RADIUS client connection attempts by SSID



RADIUS failure reason codes



Impact

1/1 Access points with most failures

Access Point	Failures
HQ-AP-1	9/15

Impact

1/1 Clients with most failures

Client	Failures
Win-10-WiFi	9/15

Technet24

Figure 11-77 *RADIUS Success View*

The **Connection Log** tab, shown in Figure 11-78, provides you with access to all connection attempts recorded in the network, along with advanced filtering capability. You can start by viewing all authentications and narrow your investigation further using the available filters, such as filtering authentications for a specific SSID.

Access points

Overview | List | Health | Map | **Connection log** | Timeline

Connection Log

05:21 to 17:21

SSID	BAND	ACCESS POINT TAG	ACCESS POINT	CLIENT	FAILURE STEP
All SSIDs	All bands	All tags	All Access Points	All clients	All steps

Showing connection failures from 12 hours ago

Connection start time	Client device	Access point	SSID	Failure stage	Failure reason
Thu Dec 4, 2025, 14:34:26	Win-10-WiFi	HQ-AP-1	MARS-Demo-HQ	Authentication	Client failed 802.1X authentication to the RADIUS server. auth_mode='wpa3-802.1x' vlan_id='10' radius_proto='ipv4' radius_ip='172.22.15.21' reason='radius_login_failure' reassoc='1' radio='0' vap='0' channel='11' rssi='67'
Thu Dec 4, 2025, 14:34:20	Win-10-WiFi	HQ-AP-1	MARS-Demo-HQ	Authentication	Client failed 802.1X authentication to the RADIUS server. auth_mode='wpa3-802.1x' vlan_id='10' radius_proto='ipv4' radius_ip='172.22.15.21' reason='radius_login_failure' reassoc='1' radio='0' vap='0' channel='11' rssi='67'
Thu Dec 4, 2025, 14:33:38	Win-10-WiFi	HQ-AP-1	MARS-Demo-HQ	Authentication	Client failed 802.1X authentication to the RADIUS server. auth_mode='wpa3-802.1x' vlan_id='10' radius_proto='ipv4' radius_ip='172.22.15.21' reason='radius_login_failure' radio='0' vap='0' channel='11' rssi='66'
Thu Dec 4, 2025, 14:33:22	Win-10-WiFi	HQ-AP-1	MARS-Demo-HQ	Authentication	Client failed 802.1X authentication to the RADIUS server. auth_mode='wpa3-802.1x' vlan_id='10' radius_proto='ipv4' radius_ip='172.22.15.21' reason='radius_login_failure' radio='0' vap='0' channel='11' rssi='67'
Thu Dec 4, 2025, 14:33:01	Win-10-WiFi	HQ-AP-1	MARS-Demo-HQ-CWA	DHCP	Client made a request to the DHCP server, but the DHCP server rejected the client's request. server_ip='10.100.11.21' vlan_id='11' details='dhcp_nack' radio='0' vap='1' channel='11' rssi='66'
Thu Dec 4, 2025, 14:32:50	Win-10-WiFi	HQ-AP-1	MARS-Demo-HQ	Authentication	Client failed 802.1X authentication to the RADIUS server. auth_mode='wpa3-802.1x' vlan_id='10' radius_proto='ipv4' radius_ip='172.22.15.21' reason='radius_login_failure' radio='0' vap='0' channel='11' rssi='67'

Technet24

Figure 11-78 *Connection Log Tab*

The last tool that proves extremely helpful for NAC monitoring and troubleshooting on the **Wireless Overview** page is **Timeline**, shown in Figure 11-79. The timeline gives you visibility into the historical state of all endpoints present in the network from the perspective of their wireless connectivity.

Access points

Overview | List | Health | Map | Connection log | Timeline

Network Timeline

Last 12 hours

SSID	BAND	CONNECTION STEP	STATUS	CLIENT	AP
All SSIDs	All bands	All connections	All statuses	All clients	All APs

Showing the most recent 22 timeline events.

Dec 4 14:34:26

Client Win-10-WiFi had a failed connection to SSID MARS-Demo-HQ on access point HQ-AP-1 during authentication because the auth server rejected the auth request.

CHANNEL	BAND	SNR	VLAN ID	RADIUS SERVER	802.11 REASON (CODE 23)
11	2.4 GHz	67 dB	10	172.22.15.21	802.1X RADIUS failed

Dec 4 14:34:26

Client Win-10-WiFi had a successful disassociation to SSID MARS-Demo-HQ on access point HQ-AP-1.

CHANNEL	BAND	SNR
11	2.4 GHz	67 dB

Dec 4 14:34:20

Client Win-10-WiFi had a failed connection to SSID MARS-Demo-HQ on access point HQ-AP-1 during authentication because the auth server rejected the auth request.

CHANNEL	BAND	SNR	VLAN ID	RADIUS SERVER	802.11 REASON (CODE 23)
11	2.4 GHz	67 dB	10	172.22.15.21	802.1X RADIUS failed

Dec 4 14:34:20

Client Win-10-WiFi had a successful disassociation to SSID MARS-Demo-HQ on access point HQ-AP-1.

CHANNEL	BAND	SNR
11	2.4 GHz	67 dB

Technet24

Figure 11-79 *Timeline Tab*

CS Cloud Native Switches Terminal

Terminal is a powerful tool for monitoring and troubleshooting tasks that you need to perform on cloud-native Cisco switches managed by the Meraki dashboard. You can access it by navigating to **Switching > Switches > (Select your switch) > Tools > Launch Terminal**.

In the terminal, you can use standard IOS XE show commands typically used to monitor and troubleshoot NAC on regular Catalyst/Smart switches. Some examples of common commands useful for NAC monitoring and troubleshooting on Catalyst/Smart switches are explained in Table 11-5.

Table 11-5 *Common Commands Useful for NAC Monitoring and Troubleshooting on Catalyst Switches*

Command	Action
show aaa accounting	Displays accounting information for AAA (authentication, authorization, and accounting) services.
show aaa authentication	Shows the status and configuration of AAA authentication.
show aaa authorization	Displays AAA authorization information.
show access-lists	Lists all configured access control lists (ACLs) and their statistics.
show authentication sessions interface [*interface*] **details**	Displays detailed information about authentication sessions on a specific interface, including status, methods, and session IDs.
show ip access-list interface [*interface*]	Displays ACLs currently assigned to the interface.
show dot1x	Displays IEEE 802.1X authentication status.
show dot1x interface ethernet [*interface*]	Shows 802.1X status on a specific Ethernet interface.
show ip dhcp snooping	Displays DHCP snooping status and configuration.
show ip dhcp snooping binding	Lists DHCP snooping bindings.
show ip arp inspection	Shows Dynamic ARP Inspection (DAI) status.
show ip arp inspection statistics	Displays statistics for ARP inspection.
show radius	Displays RADIUS server status and statistics.
show radius-server	Shows detailed RADIUS server configuration and status

Figure 11-80 provides an example of the **show authentication session interface Fi1/0/45 details** command in use.

Technet24

Terminal Opens a troubleshooting console to this device, similar to an IOS-XE CLI window.

Close Terminal | Detach | Detach (full)

This terminal allows limited access to read-only IOS-XE commands, such as **show**.

```
HQ-Access-1#show authentication session interface Fi1/0/45 details
Interface:  FiveGigabitEthernet1/0/45
              IIF-ID:  0x18EE47CE
         MAC Address:  2c31.246a.1d97
        IPv6 Address:  Unknown
        IPv4 Address:  Unknown
           User-Name:  2C-31-24-6A-1D-97
         Device-type:  Cisco-IP-Phone-8841
         Device-name:  Cisco IP Phone 8841
              Status:  Authorized
              Domain:  VOICE
      Oper host mode:  multi-auth
    Oper control dir:  both
     Session timeout:  N/A
  Acct update timeout:  172800s (local), Remaining: 88724s
   Common Session ID:  15FE640A00000029C6B7AA76
     Acct Session ID:  0x00000008
              Handle:  0xff000011
      Current Policy:  POLMAP_AP_0


Server Policies:


Method status list:
       Method          State
        dot1x          Authc Failed
          mab          Authc Success

HQ-Access-1#
```

Figure 11-80 *Details of the* ***show authentication sessions interface [interface]*** *Command*

Packet Captures

Modern access points and MS/CS switches provide support for Intelligent Capture. This feature allows you to store and view your captures in the Meraki cloud. Intelligent Capture also introduces the following enhancements that were not available with classic captures:

- **Scheduled capture:** You can to create a schedule to execute your capture. This approach is supported for access points and switches.
- **Proactive PCAP:** This feature is supported on access points and allows you to start packet captures based on specific events, such as authentication failure.

You can configure your capture by navigating to **Network-wide > Intelligent Capture** or **Assurance > Intelligent Capture**. At the top of the page, you need to select the device type to access the Intelligent Capture user interface specific to the given device type. Figure 11-81 shows the **Intelligent Capture** page with the drop-down list to select device type.

Technet24

Intelligent Capture

For access points

- For security appliances
- For switches
- For access points ✓

Enhancing Intelligent Capt **P.**

Explore how Proactive PCA pture, automatically diagnosing issues and delivering

New capture | Stored captur ptures | Proactive PCAP Enablement

Access Points

Access Points Select

Interface

Wireless

Output

Save to cloud

Duration (secs)

60

Capture filters

View example filters

File name

HQ-Brownfield-POL - wireless-wireless

Notes

We recommend including a description of the reason for your capture, this will help us improve the analysis capabilities of our capture tool.

Start capture | Schedule capture

Figure 11-81 *Intelligent Capture*

If you need to capture RADIUS traffic on the access point, select **Wired Interface** and apply a filter such as **port 1812** or **port 1813**. Figure 11-82 shows a capture for an access point configured to filter RADIUS traffic with the output type **download**, which means the capture will be downloaded directly to your computer once completed.

Technet24

New capture | Stored captures | Scheduled captures | Proactive PCAP Enablement

Access Points

All Access Points X HQ-AP-1 X

Interface

Wired

Output

Download

Duration (secs)

60

Capture filters

port 1812 or port 1813

View example filters

File name

HQ-Brownfield-POL-Multiple-APs-wired

Notes

We recommend including a description of the reason for your capture, this will help us improve the analysis capabilities of our capture tool.

Start capture

Figure 11-82 *Access Point Intelligent Capture Example*

In the case of the switch, you can use the same filter on the uplink interface.

Summary

This chapter presents a focused overview of network access control in Cisco Meraki cloud-managed networks. You learn how the Meraki dashboard functions as the central platform for unified management, configuration, and troubleshooting. The text explains Meraki's logical and physical architecture, including the roles of administrators, organizations, networks, devices, and the shard infrastructure that supports secure, reliable operations.

The chapter describes how Meraki devices, such as switches, access points, and security appliances, integrate with Cisco Identity Services Engine to enforce network access control using RADIUS authentication. It reviews how proper source IP selection and device configuration are essential for successful integration, whether with older or newer devices.

Wired and wireless NAC methods are covered in detail, including the use of 802.1X, MAC Authentication Bypass (MAB), Identity PSK (iPSK), and integration with ISE for advanced access scenarios. The text also outlines key dashboard tools and settings for applying group policies, VLAN profiles, and authentication options, along with advanced features like URL redirects and RADIUS caching.

Finally, the chapter explores the main tools for NAC monitoring and troubleshooting in the Meraki dashboard. These tools include the Clients Table, Event Log, Wireless Overview, and the terminal on Catalyst/Smart switches for IOS XE commands. The chapter concludes

Technet24

with the use of Intelligent Capture for packet analysis, ensuring you have the resources to verify authentication status and resolve NAC issues effectively.

Chapter 12. Meraki Adaptive Policy

Introduction

This chapter focuses on the theory and practice of implementing Adaptive Policy microsegmentation mechanisms in cloud-managed networks. You will explore the evolution of microsegmentation in Meraki networks, learn how to implement Adaptive Policy at the organization and device levels, and review methods for verifying and troubleshooting Adaptive Policy.

Microsegmentation Evolution in Meraki Networks

Over time, methods of microsegmentation in cloud-managed Meraki networks have evolved from simple network access control procedures, using access control lists and firewall rules, to advanced Adaptive Policy microsegmentation. You can identify the following stages in the evolution of microsegmentation in Meraki cloud-managed networks:

- **Access control lists and firewall rules:** These are the fundamental means of implementing microsegmentation at the network level in the dashboard. Access control rules apply to switches, whereas firewall rules apply to access points. This approach is easy to deploy but lacks flexibility and does not

Technet24

provide identity-based microsegmentation.

- **Group policies:** The next stage involves using group policies, which serve as common configuration containers for a wide variety of security settings across all network devices, such as switches, access points, and MX security appliances. For both wired and wireless microsegmentation, firewall rules within group policies enable granular access control for specific users and endpoints. Since you can assign group policies using RADIUS servers, this method delivers identity-based microsegmentation. However, unlike downloadable access control lists (DACLs) on ISE, group policies lack centralized management, requiring you to create them individually at each network level.
- **Adaptive Policy:** This approach introduces a centralized, organization-level microsegmentation model. Well-known Cisco TrustSec technologies, such as security group tags (SGTs) and enforcement through security group access control lists (SGACLs), are used to classify, propagate, and enforce traffic using intent-based policies.
- **Common Policy:** Although Adaptive Policy addresses microsegmentation effectively in cloud-managed networks, you may encounter challenges due to the duality of the control plane when Adaptive Policy and TrustSec-related settings exist on both the ISE and Meraki dashboard. This situation creates an administrative burden, because you must ensure that all elements remain synchronized. The introduction of the ISE Meraki Connector resolves this issue by allowing you to manage policies and environmental data components, such as SGTs and SGACLs, on ISE, ensuring that this information stays synchronized with the Meraki dashboard.

Segmentation with ACLs and Group Policies

As in traditional on-premises Cisco networks, microsegmentation in Meraki cloud-managed networks historically begins with implementing ACLs and firewall rules:

- You use ACLs for wired access control with Meraki switches.
- You deploy firewall rules for wireless access control with Meraki access points.

You can configure ACL rules for switches at the network level, which means that rules you create will apply to all switches within the specified network. Only one set of rules is supported per network, so you must integrate all requirements for the network into a single rule set. To configure ACLs for switches, navigate to **Switching > ACL**. By default, only the Allow rule appears in the **User-Defined Rules** section. All rules you create are added before the default Allow rule. Figure 12-1 shows an ACL rule set for switches.

Technet24

Access control list

Dashboard service rules

In order to help maintain connectivity with dashboard, dashboard service rules are added to the access control list. These rules consist of an explicit allow for all IPv4 traffic to and from the listed dashboard IP addresses.

Dashboard IP address
158.115.128.0/19
209.206.48.0/20
216.157.128.0/20

User-defined rules

#	Policy	IP Version	Protocol	Source	Src port	Destination	Dst port	Vlan	Established
1	Allow	IPv4	TCP	10.100.10.0/24	Any	10.110.1.15	8080		
2	Deny	IPv4	TCP	10.100.10.0/24	Any	10.110.1.17	8081		
3	Deny	IPv4	TCP	10.100.10.0/24	Any	10.110.1.17	443		
4	Deny	IPv4	TCP	10.100.10.0/24	Any	10.110.11.0	445		
5	Deny	IPv4	UDP	10.100.10.0/24	Any	10.110.1.19	53		
	Allow	Any	Any	Any	Any	Any	Any	Any	

Add a rule

Figure 12-1 *Switch ACL*

Following are some important characteristics of switch ACL rules:

- The rule set ends with a default Allow Any rule, so if you require a strict security policy, you need to prepend a default Deny Any rule after all other rules in your rule set. If you plan to implement a default Deny Any rule, carefully consider the potential impact.
- The rules apply in the inbound direction on all ports of all switches in the network. When building your rule set, make sure the rules permit traffic in both directions.
- ACL rules are stateless, but the Established option is supported with IOS-XE Meraki (Catalyst switches managed by Meraki dashboard) switches, allowing only packets that are part of an already-established connection.
- You can associate rules with specific VLANs. Note that this functionality is not supported with Catalyst switches and MS-390.

For MR/CW access points, the most basic way to implement microsegmentation is by defining Layer 3 firewall rules per network and SSID. You can configure rules for each dashboard network and then select the specific SSID to which the rules will apply. This approach enables you to create a unique rule set for each SSID in every network. To configure Layer 3 firewall rules for MR/CW, navigate to **Wireless > Firewall & Traffic Shaping**, and select the target SSID from the drop-down list at the top of the page. Layer 3 firewall rules are located in the **Outbound Rules** section. You can add rules by clicking **Add New**. Figure 12-2 demonstrates a Firewall & Traffic Shaping configuration.

Technet24

Firewall & traffic shaping

SSID: MARS-Demo-HQ

MARS-Demo-HQ

Block IPs and ports

Layer 2 LAN isolation: Disabled (bridge mode or tunneled to CG only)

Allow Bonjour forwarding exception

DHCP guard: Disabled

RA guard: Enabled

RA allowed routers: One IPv6 address per line

Outbound rules

Search...

Add new

#	Policy	IP Version	Protocol	Destination	Dst port	Rule description	Actions
No custom rules defined							
	Allow	IPv4	Any	Local LAN	Any	Wireless clients accessing LAN	
	Allow	IPv4	Any	Any	Any	Default rule	

Figure 12-2 *MR L3 Firewall Rules*

The following are some important characteristics of MR Layer 3 firewall rules:

- All rules are stateless and do not provide any session state verification.
- All MR/CW models can support a maximum of 1800 Layer 3 firewall rules.
- Layer 3 firewall rules are designed to block traffic exiting the access point. To prevent communication between clients on the same SSID, you should enable Layer 2 isolation.

While both approaches described can be used to implement microsegmentation, they have some significant disadvantages:

- All rules must be defined at the network level, which results in a lack of organization-level visibility into microsegmentation policies.
- An increase in the number of required policies may lead to rules becoming extremely difficult to manage.
- Both approaches lack identity-based network access control because they do not consider user and endpoint details.
- Rules for wired and wireless microsegmentation do not provide intent-based filtering, which is necessary in modern networks.

As the next stage in the evolution of microsegmentation, group policies — and specifically the Layer 3 firewall rules defined within them — have become the standard for implementing identity-based, granular network-level access control. Chapter 11, "Managing Network Access in the Meraki Environment," describes group policies, so in this section, we will focus on their impact on microsegmentation and the

Technet24

challenges with group policies that typically lead to the adoption of Adaptive Policy.

Group policies are configurable at the network level, and as previously described, they serve as universal configuration containers for all access control security settings that you can apply to different device types. To view group policies defined in a specific network, navigate to **Network-wide > Group Policies**. This action brings you to the **Group Policies** configuration page, as shown in Figure 12-3. As an administrator, you can choose the **Add** (Add a group), **Delete** (×), and **Clone Group Policies** actions.



Figure 12-3 *Group Policies Configuration Page*

Figure 12-4 shows a group policy created to implement granular network access control for a specific user group.

Technet24

Name: IT-Group-policy

Schedule: Scheduling disabled

Bandwidth: Use network default — 100 Mbps — details

Firewall: Custom network firewall & shaping rules

#	Policy	Protocol	Destination	Port	Comment	Actions
1	Allow	ICMP	Any	Any		≡ X
2	Allow	TCP	10.1.98.10/32	8083		≡ X
3	Allow	TCP	10.1.101.10/32	8083		≡ X
4	Allow	TCP	10.1.100.10/32	8083		≡ X
5	Allow	UDP	10.1.100.10/32	53		≡ X
6	Allow	TCP	10.1.98.10/32	8081		≡ X
7	Allow	TCP	10.1.101.10/32	8081		≡ X
8	Allow	TCP	10.1.100.10/32	8081		≡ X
9	Allow	UDP	Any	67		≡ X
10	Deny	Any	Any	Any		≡ X
	Allow	Any	Any	Any	Default rule	

Add a firewall rule

Delete group

Figure 12-4 *Group Policy L3 Firewall Rules*

Identity-based assignment of the group for specific users and endpoints occurs as part of the authentication flow. When authentication succeeds, you can configure the RADIUS server, including Cisco ISE, to supply a Filter-ID attribute in the Access-Accept message. The switch or access point verifies the value from the Filter-ID attribute and compares it against the names of the group policies available in the given network. Figure 12-5 illustrates this process in greater detail.

Technet24

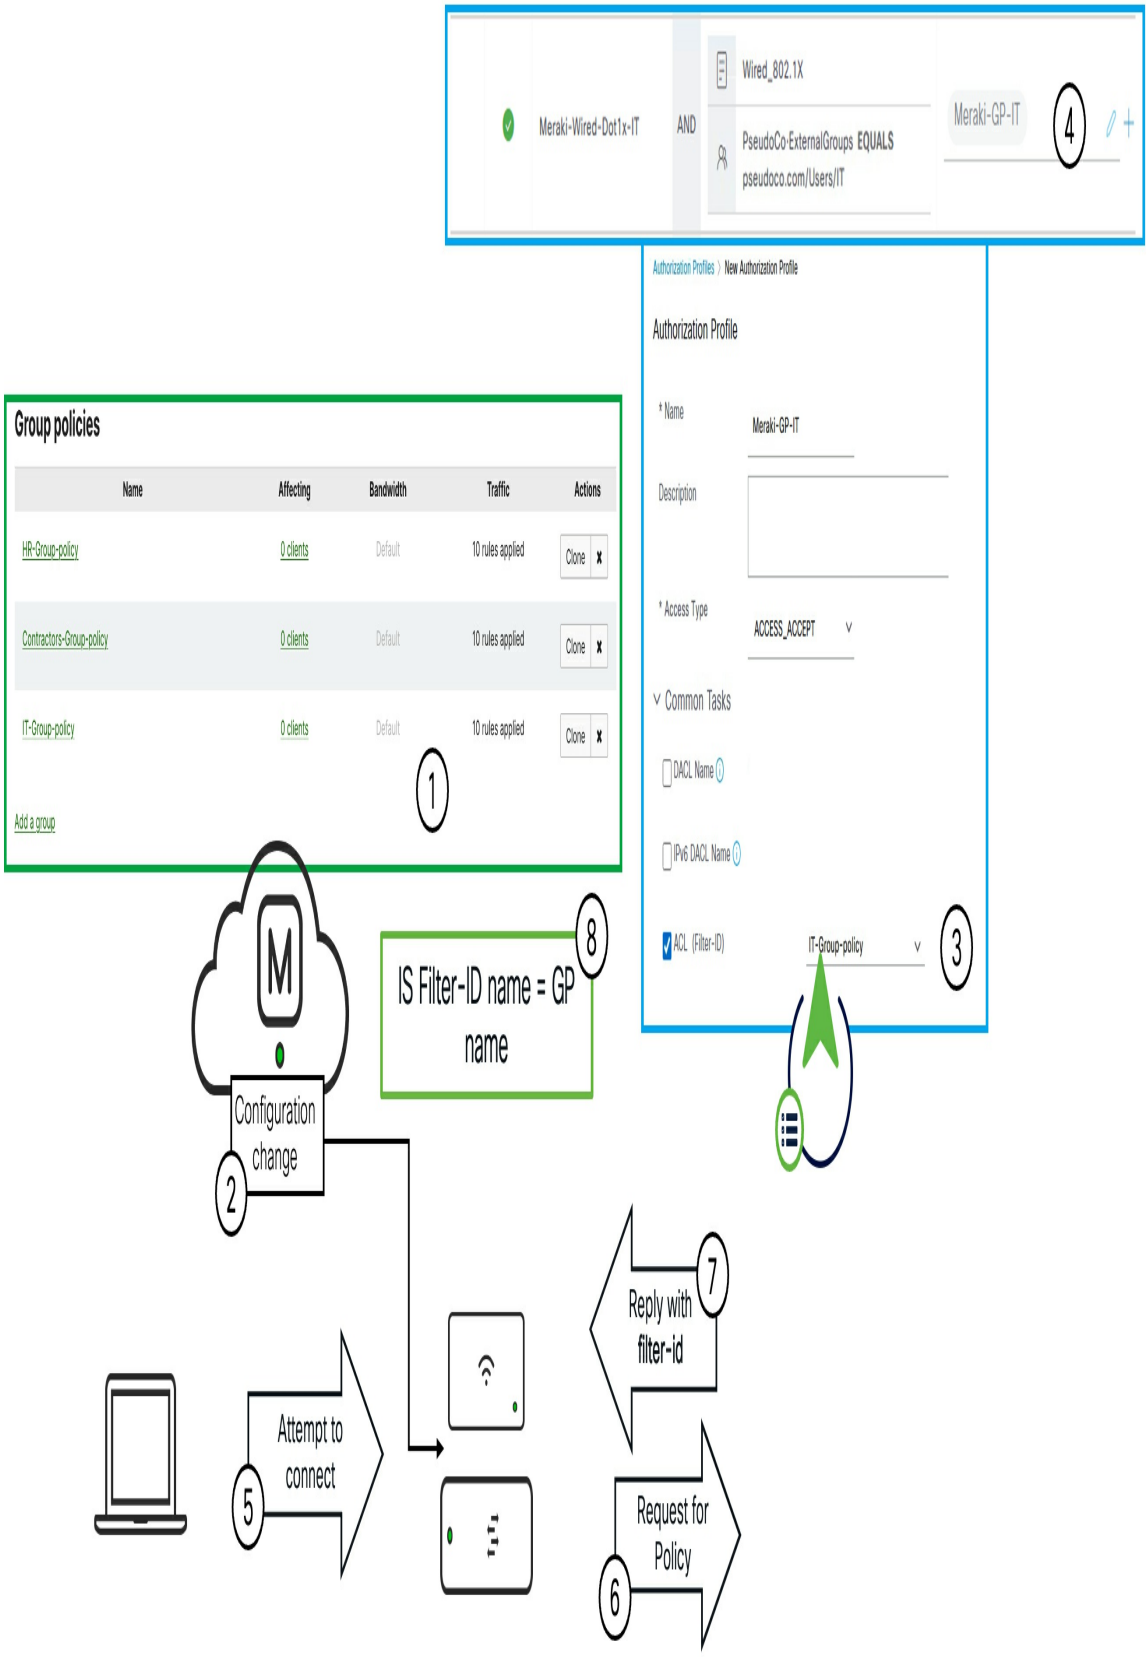

Meraki-Wired-Dot1x-IT
AND
Wired_802.1X
PseudoCo-ExternalGroups EQUALS pseudoco.com/Users/IT
Meraki-GP-IT
4
Authorization Profiles > New Authorization Profile
Authorization Profile
* Name
Meraki-GP-IT
Description
* Access Type
ACCESS_ACCEPT
Common Tasks
DACL Name
IPv6 DACL Name
ACL (Filter-ID)
IT-Group-policy
3
Group policies
Name
Affecting
Bandwidth
Traffic
Actions
HR-Group-policy
0 clients
Default
10 rules applied
Clone
Contractors-Group-policy
0 clients
Default
10 rules applied
Clone
IT-Group-policy
0 clients
Default
10 rules applied
Clone
Add a group
1
8
IS Filter-ID name = GP name
Configuration change
2
Attempt to connect
5
Reply with filter-id
7
Request for Policy
6

Figure 12-5 *Group Policy Configuration and Assignment Flow*

The following high-level steps are involved in group policy configuration and assignment flow:

Step 1. The administrator creates a group policy with Layer 3 firewall rules in the Meraki dashboard.

Step 2. The group policy propagates to network devices present in the network where the administrator defined the group policy.

Step 3. On the ISE side, the administrator creates an authorization profile with an Access Type of Access-Accept and assigns the name of the group policy to the Filter-ID defined in the Common Tasks section of the authorization profile.

Step 4. The administrator configures the authorization policy on ISE, referencing the authorization profile defined in the previous step.

Step 5. User or endpoint authentication starts between the supplicant and the access point or switch.

Step 6. The access point or switch performs the authenticator role, which results in a policy lookup if authentication succeeds.

Step 7. ISE returns Access-Accept with the Filter-ID attribute set to the name of the group policy.

Step 8. The access point or switch performs a policy name lookup in the local configuration and assigns the resulting set of rules to the endpoint or user.

While group policies address some of the challenges encountered with ACLs and network-level firewall rules, they still present limitations, especially as network size and complexity increase. Some of these

Technet24

challenges include the following:

- Group policies lack centralized management across the organization, requiring administrators to configure policies individually for each network.
- The number of group policies can become difficult to manage as the network expands, leading to increased administrative overhead.
- Group policies offer limited visibility into organization-wide access control, making it challenging to maintain consistency and compliance.
- Identity-based access control is not as robust as intent-based solutions, which can hinder granular segmentation in complex environments.

Microsegmentation with Adaptive Policy

Adaptive Policy was introduced as a next-generation solution to deliver intent-based microsegmentation in cloud-managed networks, addressing the limitations of ACLs, firewall rules, and group policies. To enable a comprehensive understanding of this solution, we will explore the following topics:

- **Introduction to adaptive policies:** This section provides a clear understanding of the solution and its underlying components.
- **Key differences between Cisco TrustSec and Meraki Adaptive Policy:** This review prepares you for designing and implementing microsegmentation in hybrid environments.
- **Meraki platforms and Adaptive Policy support:** This overview helps you understand what is supported and where.

Introduction into Adaptive Polices

Adaptive Policy is a Cisco TrustSec implementation for cloud-managed network devices operated through the Meraki dashboard. The dashboard serves as a single pane of glass where most Adaptive Policy elements are managed at the organization level. Some aspects of enabling Adaptive Policies and specific configurations must be performed at the device level.

Some important Adaptive Policy characteristics include the following:

- Adaptive Policy delivers organization-wide intent-based policy. All policies are defined at the organization level with an intent-first approach. Administrators manage policies based on user and resource groups rather than network parameters such as IP addresses and subnets, which clarifies intent not only for network administrators but also for security administrators and auditors.
- Adaptive Policy utilizes inline security group tags. Cisco SGTs tag frames at the network level to provide classification, propagation, and enforcement at any network point where Adaptive Policy is supported.
- Adaptive Policy delivers context shared over the data plane. SGTs allow connection context, specifically user or endpoint roles, to be included in each packet during inline tagging and propagation.
- Adaptive Policy provides IP and topology agnostic security. Adaptive Policy supports both IPv4 and IPv6, as well as various types of network transports such as wired, wireless, and SD-WAN, enabling interoperability and maintaining connection context across different environments.

In a discussion of Adaptive Policy, it is important to understand the naming convention for different configuration elements and how this

Technet24

convention aligns with the names used in classic TrustSec terminology. Table 12-1 explains the connection between these terms.

Table 12-1 *Comparison of Element Naming in Adaptive Policy and TrustSec*

Dashboard Adaptive Policy Element Naming	ISE TrustSec Element Naming	Explanation
Custom ACLs	Security group ACLs	Security group access control lists (SGACLs) are policy elements that define permissions between designated source and destination groups.
Groups	Security groups	Groups define security group tag (SGT) values, which are assigned to endpoints and their frames during classification, propagation, and enforcement. Groups typically have meaningful names that clearly convey identity and intent. In the background, these names are mapped to actual SGT values used on network devices. SGTs are represented as decimal or hexadecimal values within the range of 0 to 65,535 and 0000 to FFFF, respectively.
Policies	TrustSec policy	Policies define the traffic direction by specifying from which group to which group traffic flows and assign a set of permissions to this direction using SGACLs.
Adaptive Policy Objects	IP SGT static mapping	Adaptive Policy Objects and IP-SGT static mapping are typically used for endpoints where dynamic classification is not possible. A common example is a server that requires IP to SGT mapping but cannot receive an assignment through 802.1X authentication.

Note

Throughout this chapter, both regular TrustSec and Adaptive Policy terminology are used interchangeably. It is important to keep in mind the mappings between different definitions.

Adaptive Policy vs. TrustSec: Key Differences

TrustSec implementation with Cisco ISE and traditional on-premises network equipment relies on ISE as the central point for control-plane functionality. ISE authenticates network access devices when they join TrustSec domains using the EAP-FAST protocol and provides protected access credentials (PACs) to these devices. The PACs are used for authentication and authorization on the ISE side. Once provisioned, network devices can download environmental data from ISE, including all SGTs defined on the platform. At this stage, network devices remain unaware of the actual TrustSec policies and SGACLs configured on ISE. When an endpoint receives an IP to SGT mapping through any supported method, the SGT from the mapping enables the device to locate and download all policies from ISE where this group (SGT) acts as the destination.

This process requires additional effort from the administrator:

- All TrustSec network access devices must be preconfigured with Cisco TrustSec credentials.
- Administrators need to define the same credentials on both the ISE and device sides to ensure EAP-FAST communication.
- When changes are made on ISE, such as adding a new SGT or creating, deleting, or reconfiguring an SGACL, administrators must push these changes to all devices. Typically, this is done by clicking a button in the ISE UI to initiate propagation. If this step is not performed, devices will obtain updated settings during the next environmental data refresh. The default interval for

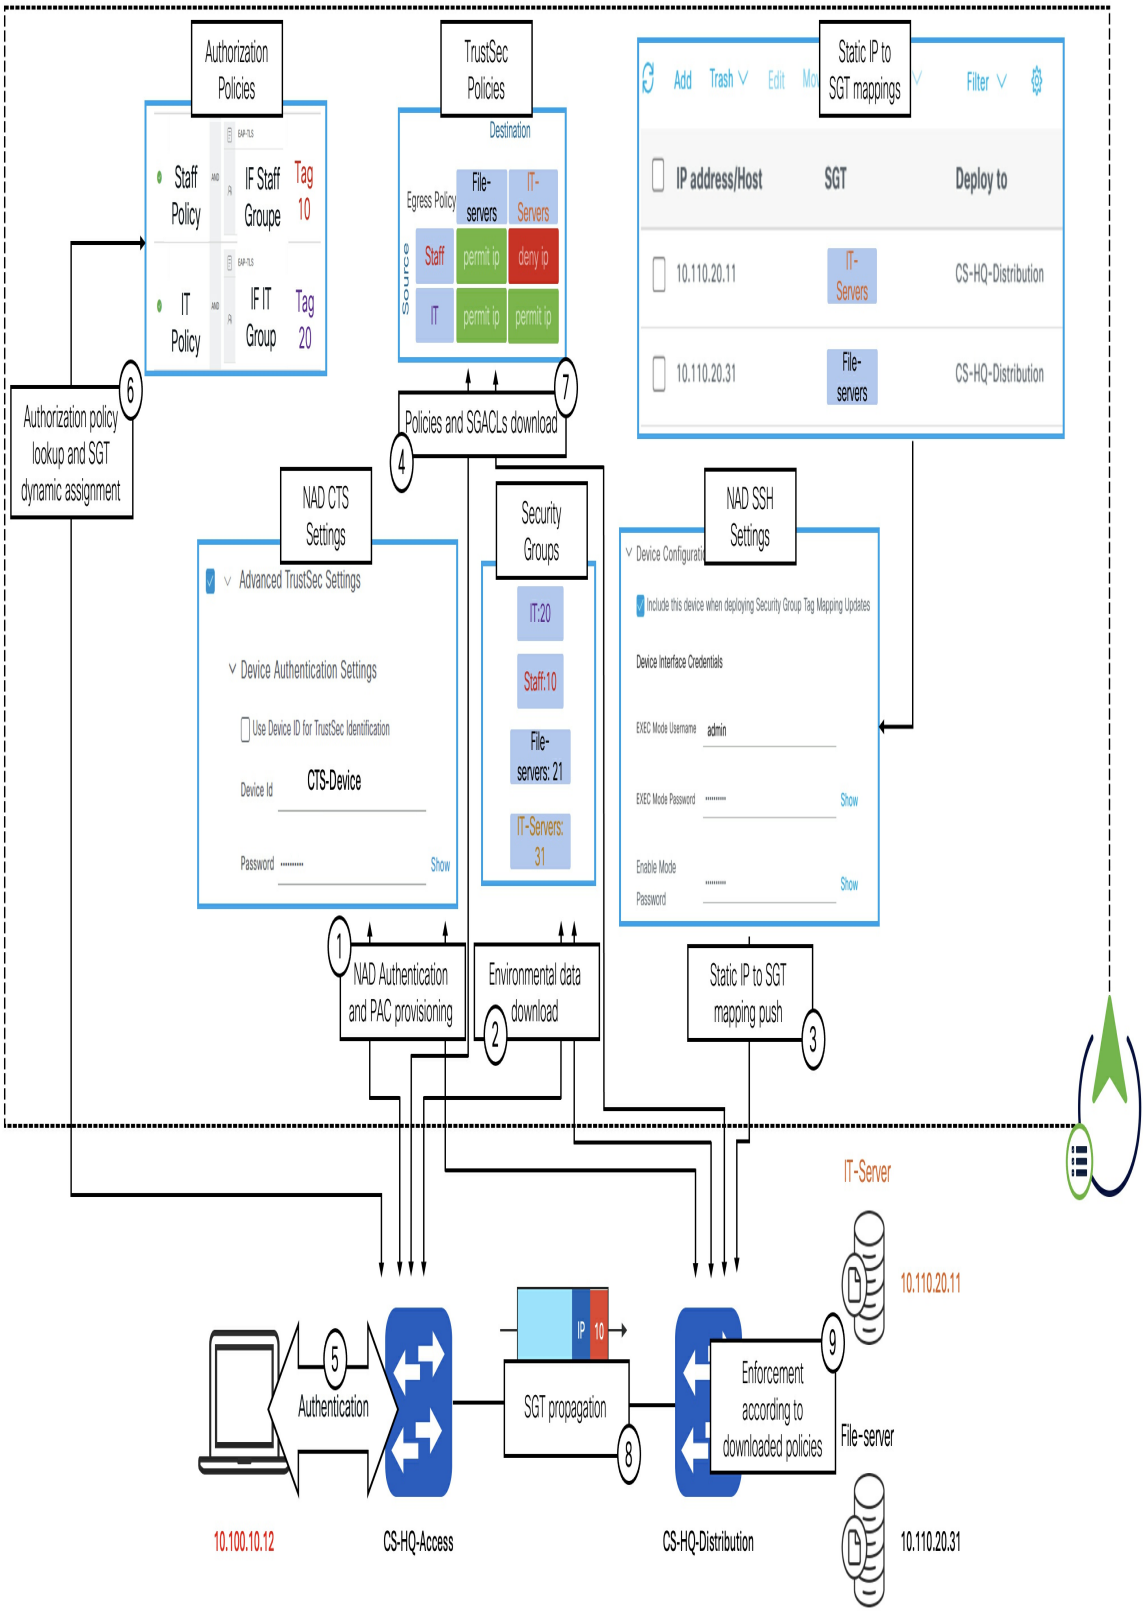
Technet24

environmental data refresh is two hours.

As an alternative to EAP-FAST, you can use PAC-less provisioning of policies and configuration, which is supported starting from ISE version 3.4. With this method, the network access device uses credentials to authenticate itself and then establishes a secure tunnel with Cisco ISE to provision policies and environmental data. This approach is recommended.

> **Note**
>
> It is recommended to perform changes to TrustSec policies and environmental data outside of business hours.

Figure 12-6 provides a high-level overview of the flow between the network access device and ISE needed to provision all required TrustSec components.


Authorization Policies
Staff Policy
IF Staff Groupe
Tag 10
IT Policy
IF IT Group
Tag 20
TrustSec Policies
Destination
Egress Policy
File-servers
IT-Servers
Source
Staff
permit ip
deny ip
IT
permit ip
permit ip
Static IP to SGT mappings
Add
Trash
Edit
Filter
IP address/Host
SGT
Deploy to
10.110.20.11
IT-Servers
CS-HQ-Distribution
10.110.20.31
File-servers
CS-HQ-Distribution
6
Authorization policy lookup and SGT dynamic assignment
7
Policies and SGACLs download
4
NAD CTS Settings
Advanced TrustSec Settings
Device Authentication Settings
Use Device ID for TrustSec Identification
Device Id
CTS-Device
Password
Show
Security Groups
IT:20
Staff:10
File-servers: 21
IT-Servers: 31
NAD SSH Settings
Include this device when deploying Security Group Tag Mapping Updates
Device Interface Credentials
EXEC Mode Username
admin
EXEC Mode Password
Show
Enable Mode Password
Show
1
NAD Authentication and PAC provisioning
2
Environmental data download
Static IP to SGT mapping push
3
IT-Server
10.110.20.11
IP
10
5
Authentication
SGT propagation
8
9
Enforcement according to downloaded policies
File-server
10.110.20.31
10.100.10.12
CS-HQ-Access
CS-HQ-Distribution

Technet24

Figure 12-6 *TrustSec Component Provisioning Flow Between ISE and Cisco Network Access Devices*

The following steps outline the process of establishing a TrustSec domain with subsequent classification, propagation, and enforcement:

Step 1. NADs perform EAP-FAST authentication on ISE to authenticate and authorize themselves as TrustSec NADs and to obtain the EAP-FAST PAC, which is later used to authenticate and authorize requests for TrustSec configuration components such as environmental data (SGTs), TrustSec policies, and TrustSec ACLs (SGACLs).

Step 2. Environmental data download.

Step 3. Provisioning of static IP to SGT mappings. This step can occur at various points in the flow, but it is a prerequisite for the TrustSec policies and SGACLs download performed by CS-HQ-Distribution (in this example). If you plan to use static IP to SGT mapping provisioned from ISE, NADs must be configured as SXP listeners, or SSH credentials should be defined in NAD settings. If SSH is used, you need to ensure that ISE nodes are permitted to establish SSH connections with NADs.

Step 4. In this example, CS-HQ-Distribution performs the TrustSec policies and SGACLs download triggered by the provisioning of static IP to SGT mappings in Step 3. Each NAD is responsible for contacting ISE and requesting TrustSec policies and SGACLs when a new IP to SGT mapping is added (static or dynamic) if the NAD does not have TrustSec policies for the SGT in the new mapping. The NAD downloads all policies and their associated SGACLs where the SGT from the new mapping is used as the

destination.

Step 5. A new endpoint is authenticated on CS-HQ-Access.

Step 6. Authentication triggers policy lookup on the ISE side with SGT assignment on the NAD. Once IP device tracking functionality learns the IP of the endpoint, the IP to SGT mapping becomes available on the switch.

Step 7. Due to the new IP to SGT mapping being added, CS-HQ-Access requests TrustSec policies and associated SGACLs from ISE. The logic of this step is identical to what was described in Step 4.

Step 8. All packets belonging to the endpoint now leave CS-HQ-Access tagged with the appropriate SGT when forwarded over interfaces enabled for TrustSec.

Step 9. Enforcement is performed on the destination end, which is CS-HQ-Distribution in this example.

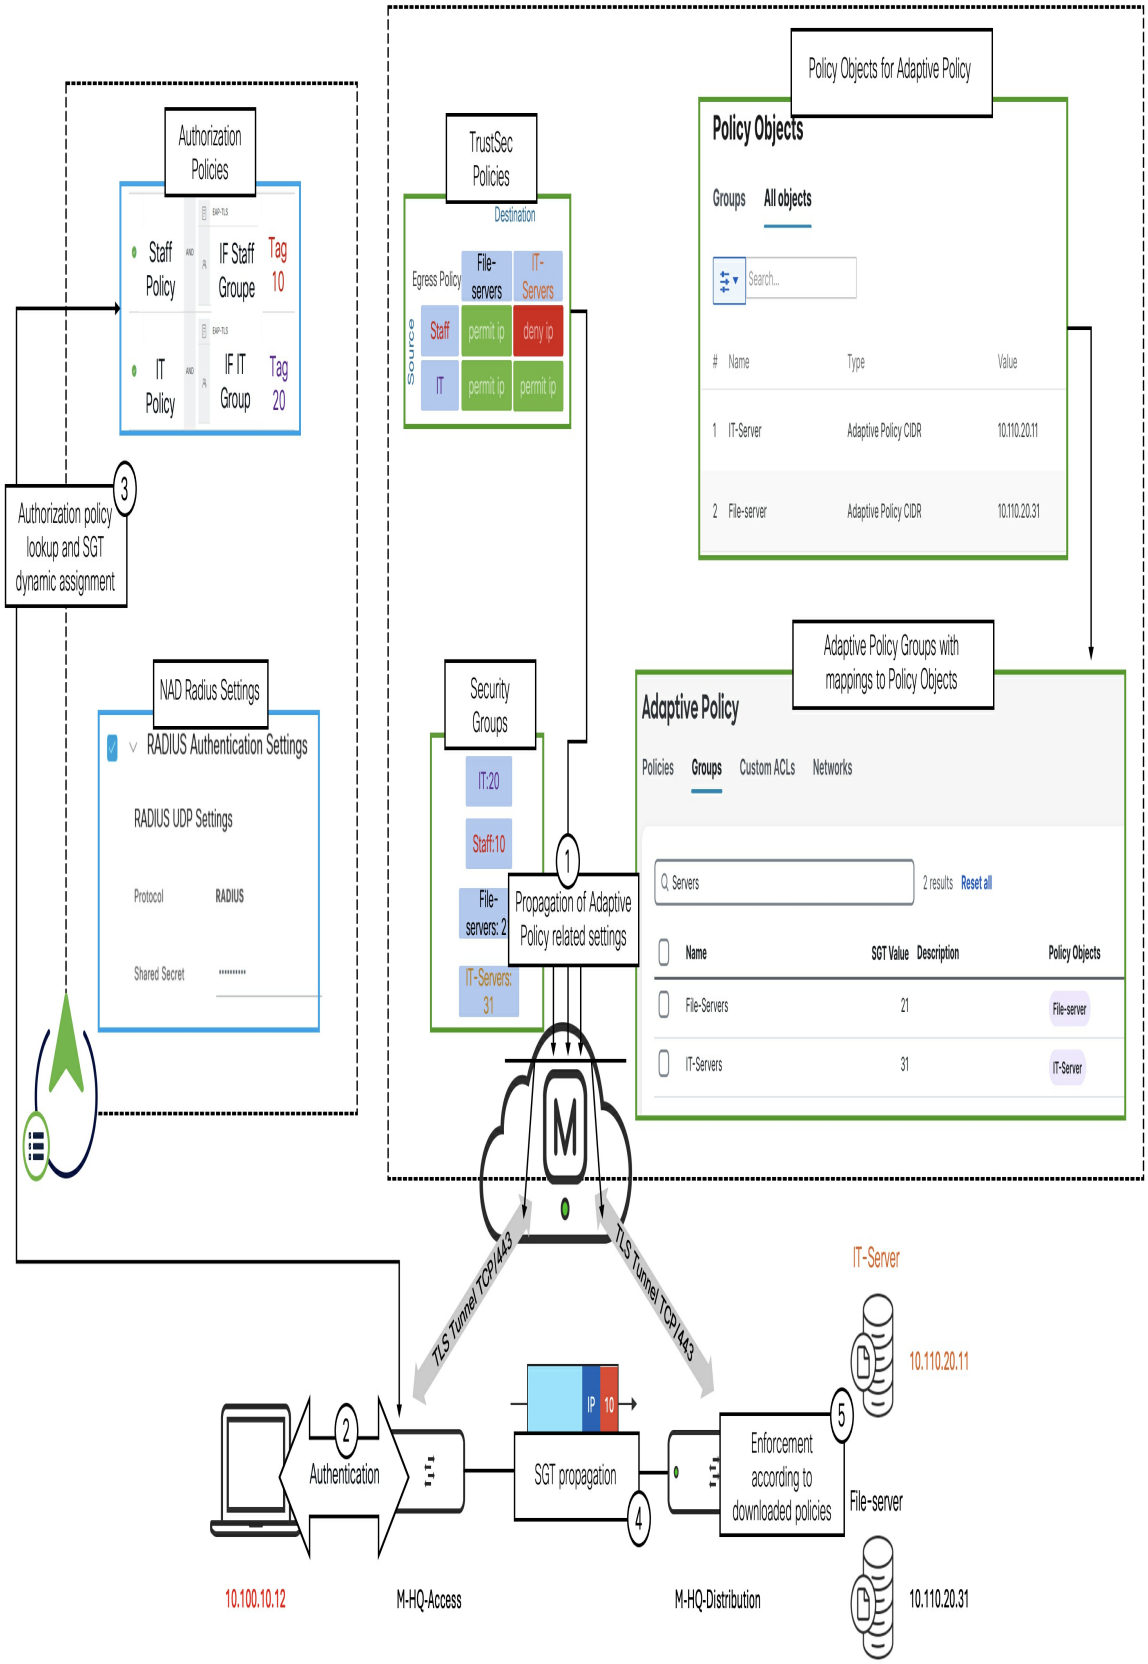
This example demonstrates how ISE uses Cisco CTS and EAP-FAST to establish the TrustSec domain and propagate required settings back to network access devices. Beginning with ISE 2.7 and Cisco IOS XE 17.1.1, a REST API connection between devices and ISE can be used as an alternative to the classic RADIUS-based CTS and EAP-FAST approach. However, this option does not completely eliminate the administrative burden of joining devices to the TrustSec domain.

With the Meraki dashboard and cloud-managed network access devices, most of the administrative overhead associated with the TrustSec flow has been eliminated. The Meraki dashboard already maintains a secure and authenticated connection to every managed device, so administrators do not need to perform additional steps to join devices to the TrustSec domain. The dashboard also ensures that configuration between each device and the dashboard remains

Technet24

synchronized, which removes the responsibility from the administrator to manually push new changes to TrustSec policies and elements to network access devices.

Figure 12-7 provides a high-level overview of the flow between the network access device and Meraki dashboard needed to provision all required TrustSec components.


Authorization Policies
Staff Policy
IF Staff Groupe
Tag 10
IT Policy
IF IT Group
Tag 20
Authorization policy lookup and SGT dynamic assignment
3
NAD Radius Settings
RADIUS Authentication Settings
RADIUS UDP Settings
Protocol
RADIUS
Shared Secret
TrustSec Policies
Destination
Egress Policy
File-servers
IT-Servers
Source
Staff
permit ip
deny ip
IT
permit ip
permit ip
Policy Objects for Adaptive Policy
Policy Objects
Groups
All objects
Search...
Name Type Value
1 IT-Server Adaptive Policy CIDR 10.110.20.11
2 File-server Adaptive Policy CIDR 10.110.20.31
Adaptive Policy Groups with mappings to Policy Objects
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Servers
2 results
Reset all
Name
SGT Value
Description
Policy Objects
File-Servers
21
File-server
IT-Servers
31
IT-Server
Security Groups
IT:20
Staff:10
File-servers: 2
IT-Servers: 31
1
Propagation of Adaptive Policy related settings
TLS Tunnel TCP/1443
TLS Tunnel TCP/1443
IT-Server
10.110.20.11
IP
10
2
Authentication
SGT propagation
4
5
Enforcement according to downloaded policies
File-server
10.110.20.31
10.100.10.12
M-HQ-Access
M-HQ-Distribution

Technet24

Figure 12-7 *TrustSec Component Provisioning Flow Between Meraki Dashboard and Cloud-Managed Network Access Device*

The following steps outline the process of Adaptive Policy propagation with subsequent classification, propagation, and enforcement:

Step 1. The dashboard automatically propagates all Adaptive Policy elements to networks where Adaptive Policy is enabled. After elements are created, edited, or deleted, administrators do not need to take additional action to ensure that changes reach the corresponding network access devices. The dashboard propagates all changes through secure connections established between the dashboard and cloud-managed devices.

Step 2. A new endpoint is authenticated on MS-HQ-Access.

Step 3. Authentication triggers a policy lookup on the ISE side with an SGT assigned on the NAD. Once IP device tracking functionality learns the IP of the endpoint, the IP to SGT mapping becomes available on the switch.

Step 4. All packets belonging to the endpoint now leave MS-HQ-Access tagged with the appropriate SGT when forwarded over interfaces enabled for TrustSec.

Step 5. Enforcement is performed on the destination end, which is MS-HQ-Distribution in this example.

As demonstrated by the two types of flows, Adaptive Policy offers a simpler process and requires significantly less work from the administrator to ensure that microsegmentation is in place. The key difference to highlight is that when you use TrustSec with Cisco NADs and ISE, all related settings are provisioned on ISE. In contrast, with Adaptive Policy, all settings must be configured in the Meraki

dashboard. However, because dynamic SGT assignment occurs from ISE, the administrator must ensure that SGT values remain synchronized between ISE and the Meraki dashboard. By default, this manual process requires careful review of group names and their assigned tag values on both platforms whenever changes are made. Failing to maintain alignment can result in ISE assigning a tag value different from what the Meraki dashboard expects, which can cause enforcement issues with microsegmentation.

Requirements for Platforms Supporting Adaptive Policy

In the context of cloud-managed Meraki networks, Adaptive Policy is supported on three main platforms:

- MS and CS cloud-managed switches
- MR and CW cloud-managed access points
- MX security appliances

Each platform has unique characteristics regarding Adaptive Policy, and different sets of Adaptive Policy features are supported by each. At a minimum, you need to understand the set of hardware, software, and licensing requirements. All platform-specific details presented in this book are valid as of the time of writing, but in the future, the list of hardware, software, and licensing requirements may change, as well as the features supported by each individual platform.

Table 12-2 explains the current status of hardware, software, and licensing requirements for each platform.

Table 12-2 *Adaptive Policy Hardware, Software, and Licensing Requirements*

Technet24

Platform	Model	Minimum License	Minimum Firmware Version	Recommended Version
Switches	MS390	Advanced	MS14.33.1	Latest Cloud Management with IOS XE release
	C9300-M	Advanced	CS15.21.1	
	C9300\L\X-M	Advanced	CS16.7	
	C9200L	Advanced	Cloud Management with IOS XE 17.15+	
	C9200/CX	Advanced	Cloud Management with IOS XE 17.18+	
	C9500H	Advanced	Cloud Management with IOS XE 17.18+	
	MS130X/R & MS150	Advanced	MS17/MS18	Latest MS release
Access points	Meraki MR and CW Access Points*	Advanced	MR27	Latest MR release
MX security appliances	All models that can upgrade to firmware 18.1 except MX84	Advanced for SGT Transport	MX18.1	Latest MX Release
	All models that can upgrade to firmware 18.1 except MX84	Secure SD-WAN Plus for SGT Assignment	MX18.2	Latest MX Release

* Wi-Fi 5 wave 2, Wi-Fi 6, Wi-Fi 6E (MR57 and CW916x), and Wi-Fi 7 (CW917x) access points.

Migration to Common Policy

As previously mentioned, one of the significant challenges in managing microsegmentation with Adaptive Policies is maintaining alignment between Adaptive Policies and TrustSec settings across ISE and the Meraki dashboard. Figure 12-8 demonstrates a scenario that may result in improper access being granted to a user or endpoint.

Technet24

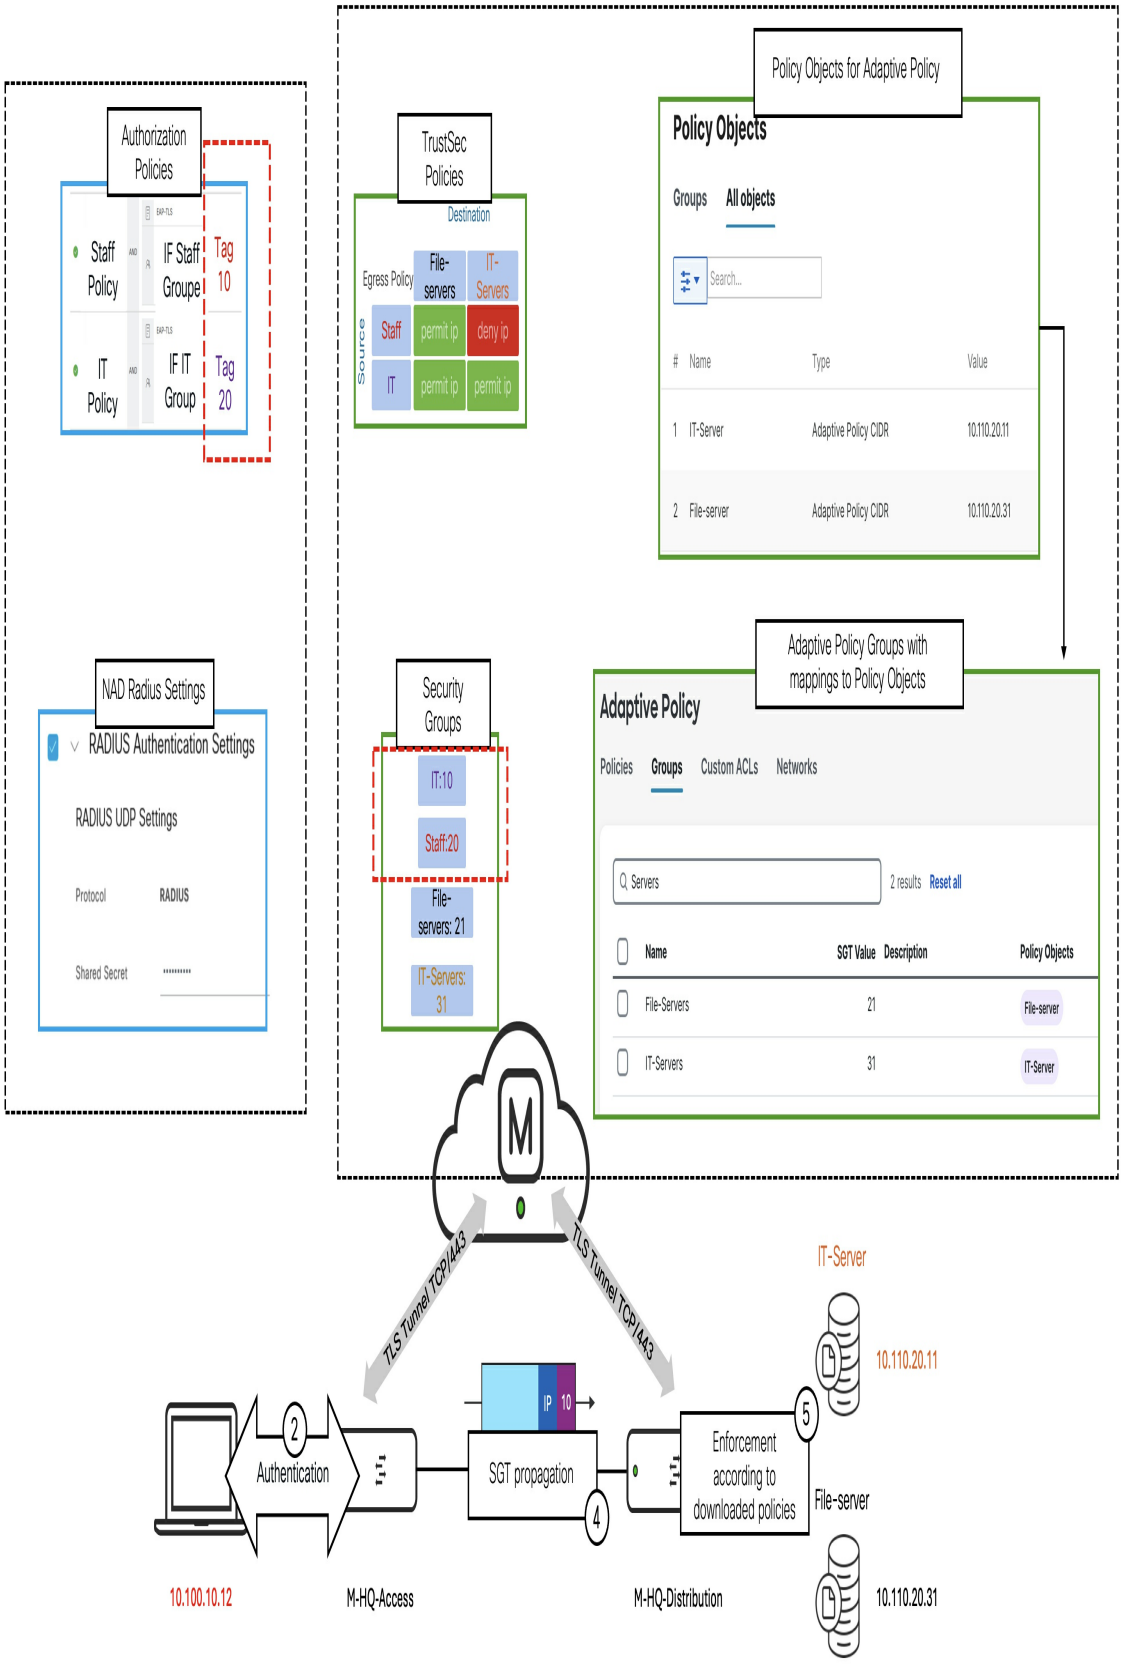

Authorization Policies
Staff Policy
IF Staff Groupe
Tag 10
IT Policy
IF IT Group
Tag 20
TrustSec Policies
Destination
Egress Policy
File-servers
IT-Servers
Source
Staff
permit ip
deny ip
IT
permit ip
permit ip
Policy Objects for Adaptive Policy
Policy Objects
Groups
All objects
Search...
Name
Type
Value
1 IT-Server
Adaptive Policy CIDR
10.110.20.11
2 File-server
Adaptive Policy CIDR
10.110.20.31
NAD Radius Settings
RADIUS Authentication Settings
RADIUS UDP Settings
Protocol
RADIUS
Shared Secret
Security Groups
IT:10
Staff:20
File-servers: 21
IT-Servers: 31
Adaptive Policy Groups with mappings to Policy Objects
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Servers
2 results
Reset all
Name
SGT Value
Description
Policy Objects
File-Servers
21
File-server
IT-Servers
31
IT-Server
TLS Tunnel TCP/443
TLS Tunnel TCP/443
IT-Server
10.110.20.11
IP
10
2
Authentication
SGT propagation
4
5
Enforcement according to downloaded policies
File-server
10.110.20.31
10.100.10.12
M-HQ-Access
M-HQ-Distribution

Figure 12-8 *Conflicting SGT Usage Between ISE and Meraki Dashboard*

As demonstrated in the figure, ISE assigns tag 10 to Staff users and tag 20 to IT users. On the Meraki side, the situation is reversed: Adaptive Policy group 10 is used for IT, and group 20 is used for Staff users. This results in inappropriate access being assigned to each user group. It is worth noting that this issue can be difficult to identify during troubleshooting. The situation becomes even more challenging because, by default, ISE automatically generates SGT numbers, and if the administrator does not pay close attention, it is easy to end up with similar misalignment.

To address these issues and unify policy management between the Meraki and ISE dashboards, Cisco introduced the Meraki Connector in ISE 3.2 patch 1. The Meraki Connector enables you to synchronize TrustSec configurations from ISE into the Meraki dashboard. This integration allows administrators to use the same policies across the entire hybrid environment. Even in a fully cloud-managed environment, it is beneficial to rely on the Meraki Connector to prevent operational errors.

At a high level, ISE establishes a connection with the Meraki dashboard using standard API calls and performs initial synchronization of policy and configuration objects. Ongoing synchronization occurs at predefined intervals to ensure that ISE and the dashboard remain aligned for TrustSec and Adaptive Policy configurations.

The main benefits of the Meraki Connector include the following:

- It helps with hybrid deployments by maintaining microsegmentation configuration integrity.

Technet24

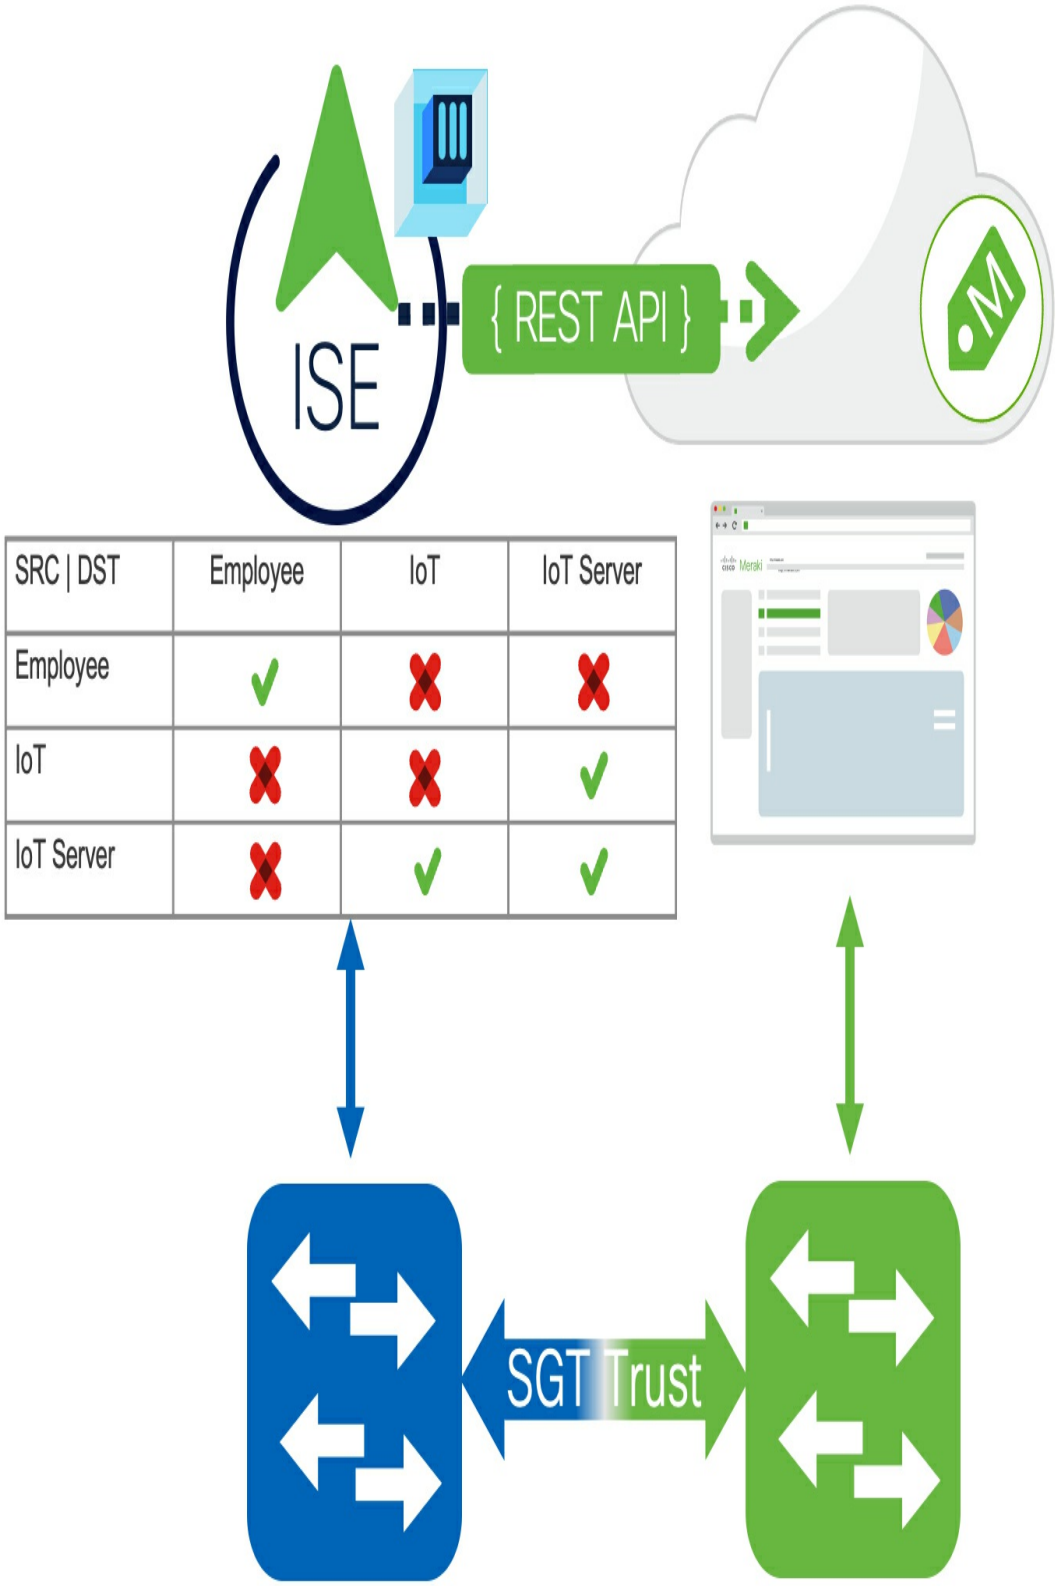
- It removes operational redundancy by eliminating the need to replicate policies across multiple platforms.
- It reduces the possibility of policy errors that could affect user experience.
- It simplifies administration by allowing ISE to serve as the single source of truth and interface for policy management.

The ISE Meraki Connector has the following scaling characteristics:

- ISE can connect to up to 20 Meraki organizations.
- Forty SGTs can be synchronized from ISE to each Meraki organization.
- The connector can synchronize a matrix of up to 25-by-25 SGTs, resulting in 625 policies with up to seven ACLs per policy.
- The connector can synchronize up to 95 SGACLs with a maximum of 16 ACEs (rules) per SGACL.

When planning your deployment, always consider the lowest limit. For example, while the ISE connector theoretically supports 625 policies with up to seven SGACLs per policy, the actual number of unique SGACLs is 95.

Figure 12-9 demonstrates a hybrid network where both TrustSec and Adaptive Policy must be implemented, and the ISE Meraki Connector is used to simplify the administration of policies and their elements to deliver consistent policy across all network segments.


ISE
{ REST API }
SRC | DST
Employee
IoT
IoT Server
Employee
IoT
IoT Server
SGT Trust

Technet24

Figure 12-9 *Meraki ISE Connector High-Level Overview*

Important ISE Meraki connector considerations are as follows:

- ISE TrustSec Matrix maps to organizations, not individual networks.
- Synchronization occurs from one ISE deployment to up to 20 Meraki organizations.
- All Meraki organizations receive the same default ISE TrustSec Matrix.
- Objects can only be synchronized from ISE to Meraki dashboard.
- ISE overrides any existing objects in Meraki.
- ISE does not delete Meraki Adaptive Policies; it only creates policies.
- Only the default ISE TrustSec Matrix is supported, not multiple matrices.
- Meraki supports only limited SGACL syntax.
- No indication exists in the Meraki dashboard regarding which policies originated from ISE.

When you are planning and operating with SGACL syntax and content, it is important to remember that not all syntax supported by ISE can be propagated to the Meraki dashboard. The main structure of each individual entry in the SGACLs selected for synchronization should appear as follows:

```
{ Deny | Allow } { TCP | UDP | ICMP | Any } src-port dst-port
```

Keep the following caveats in mind:

- For protocols ICMP and Any, both the src-port and dst-port must

be any. You cannot specify a port, port list, or port range in these cases.

- Port numbers can be in the range of 1–65535. Named ports are not allowed.
- A comma-separated list of port numbers can be used — for example, 80, 443, 1521.
- A range of port numbers can be used — for example, 33–44.

Recommendations and Best Practices for Hybrid Environments

Hybrid Environment Example

When you're operating hybrid environments with cloud-managed and classic Cisco switches or other TrustSec-capable devices, it is important to consider the following:

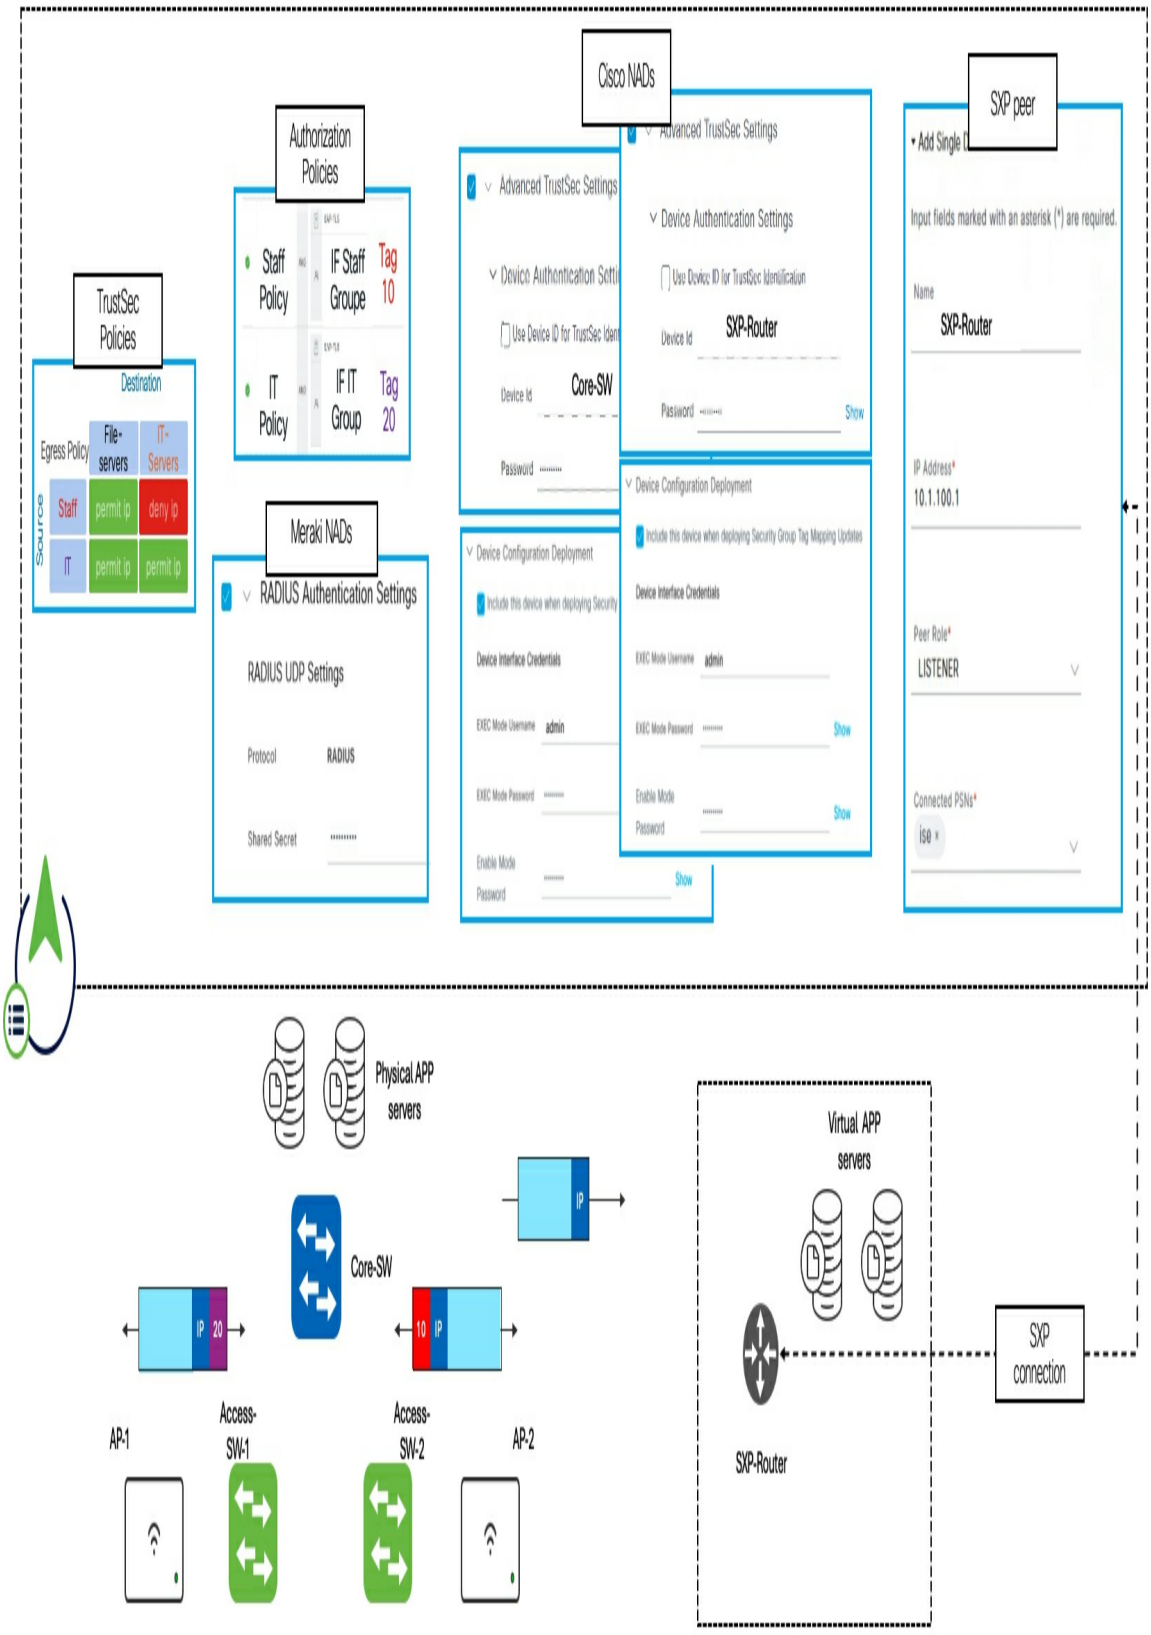
- ISE serves as the central control plane element and single point of policy and policy element management for any on-premises Cisco device. This means that network access devices must be properly configured for CTS on both ends — ISE and the network device itself — for settings and policies to be provisioned.
- For cloud-managed devices, by default, only authorization policies are used on ISE for SGT assignment.
- On-premises Cisco switches are commonly used as hardware for the core or distribution layer in hybrid networks. When you're connecting cloud-managed and on-premises switches, it is important to pay attention to inline tagging settings on both ends. You need to ensure the same SGT is used at the interface level and enable tagging only when both ends are prepared for

Technet24

TrustSec or Adaptive Policy.

- On-premises Cisco switches and other TrustSec devices can be integrated with ISE using SXP to deliver enforcement for endpoints located behind cloud-managed switches when inline tagging cannot be used.

Figure 12-10 shows an example of a hybrid network environment where inline tagging is implemented between Meraki access points, cloud-managed switches, and a classic on-premises core switch. The core switch in this case enforces policies for traffic destined to physical application servers connected to it. Mappings for application servers can be deployed statically or via SXP. All traffic that needs to reach virtual application servers leaves the network segment untagged. The SXP router delivers microsegmentation enforcement once packets arrive, due to its SXP connection with ISE. By default, ISE propagates all dynamic IP to SGT mappings for authenticated clients and endpoints to each SXP listener. This allows the SXP router to use the source IP in the packet from the remote endpoint to locate the local IP to SGT mapping.


Cisco NADs
SXP peer
Authorization Policies
TrustSec Policies
Meraki NADs
Staff Policy
IF Staff Groupe
Tag 10
IT Policy
IF IT Group
Tag 20
Destination
Egress Policy
File-servers
IT-Servers
Source
Staff
IT
permit ip
deny ip
permit ip
permit ip
Advanced TrustSec Settings
Device Authentication Settings
Use Device ID for TrustSec Identification
Device Id
SXP-Router
Password
Show
Core-SW
Device Configuration Deployment
Include this device when deploying Security Group Tag Mapping Updates
Device Interface Credentials
EXEC Mode Username
admin
EXEC Mode Password
Enable Mode Password
RADIUS Authentication Settings
RADIUS UDP Settings
Protocol
RADIUS
Shared Secret
Input fields marked with an asterisk (*) are required.
Name
IP Address*
10.1.100.1
Peer Role*
LISTENER
Connected PSNs*
ise
Physical APP servers
Virtual APP servers
Core-SW
IP
20
10
AP-1
Access-SW-1
Access-SW-2
AP-2
SXP connection
SXP-Router

Technet24

Figure 12-10 *Hybrid Network Environment Example*

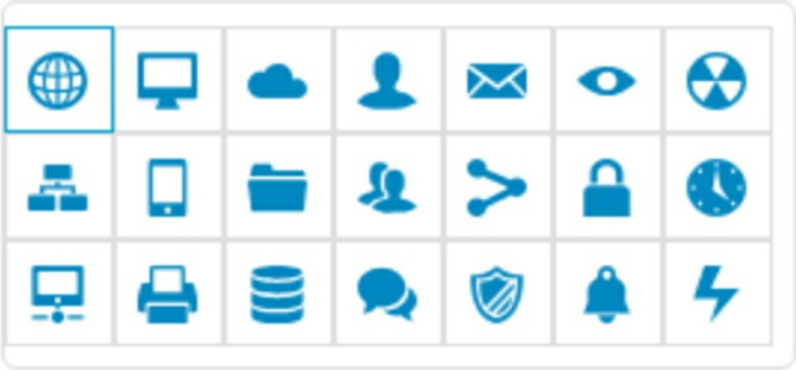
Manual SGT Management

By default, when SGTs are created on ISE, the tag value is assigned dynamically. ISE uses the next available tag number for the newly created SGT. You can verify this by navigating to **Work Centers > TrustSec > Components > Security Groups** and clicking **Add**. Figure 12-11 shows a new group being added. In the highlighted area, you can see that ISE dynamically selected the SGT value as (Dec / Hex): 16 / 0010.

* Name

IT

* Icon



Description

Security Group Tag (Dec / Hex): 16/0010

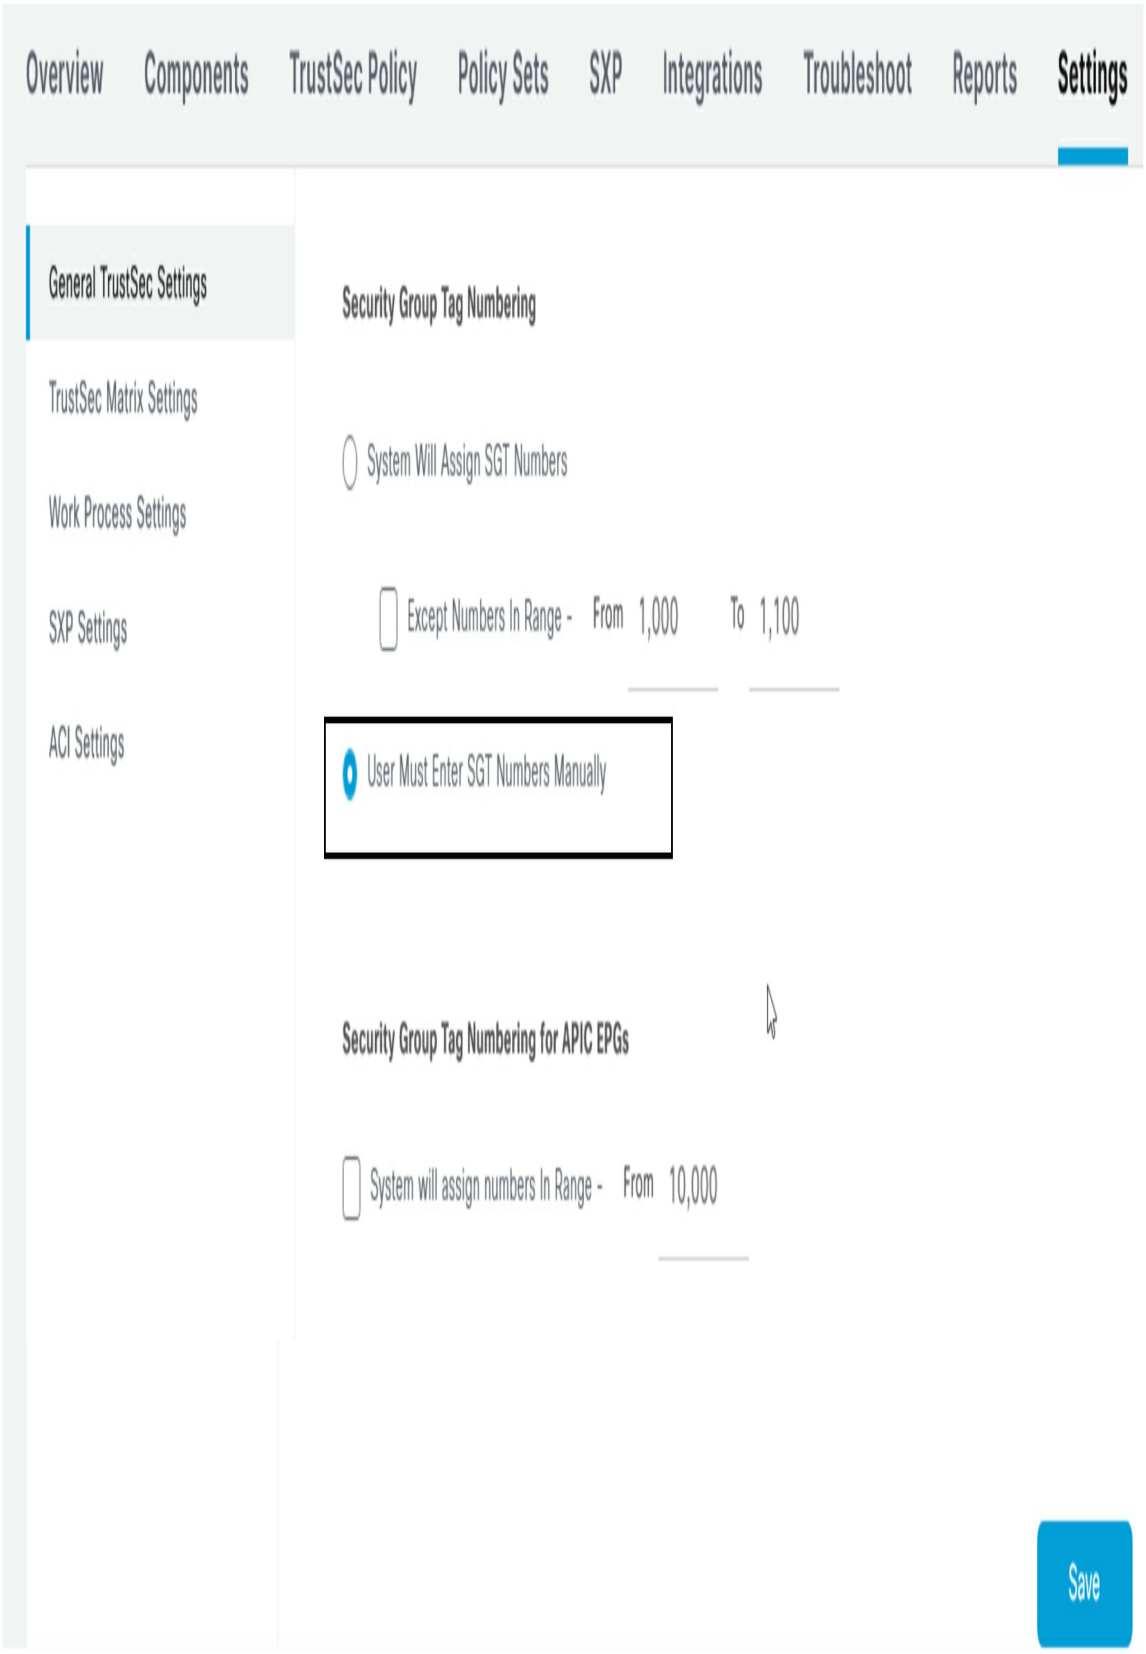
Generation Id: 0

Technet24

Figure 12-11 *Dynamic SGT Tag Example*

With this default behavior in hybrid environments, it is recommended to first create the security group on ISE and then replicate the same group name and value in the Meraki dashboard. This workflow helps protect you from conflicting configurations if the ISE Meraki Connector is not in use.

To gain greater flexibility, you can change the default ISE behavior and assign tag values manually. To achieve this, navigate to **Work Centers > TrustSec > Settings** and select **User Must Enter SGT Numbers Manually** in the **Security Group Tag Numbering** section. Click **Save** at the bottom of the page for the setting to take effect. Figure 12-12 shows how to switch to manual tag assignment.


Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
General TrustSec Settings
TrustSec Matrix Settings
Work Process Settings
SXP Settings
ACI Settings
Security Group Tag Numbering
System Will Assign SGT Numbers
Except Numbers In Range - From 1,000 To 1,100
User Must Enter SGT Numbers Manually
Security Group Tag Numbering for APIC EPGs
System will assign numbers In Range - From 10,000
Save

Technet24

Figure 12-12 *Switching to Static SGT Tags*

Configuring ISE Meraki Connector

Before you start configuring the ISE Meraki Connector, there are two important prerequisites:

- You must obtain a Meraki dashboard write API key.
- You must confirm that ISE trusts the issuer of the Meraki API endpoint.

Note

The ISE Meraki Connector establishes a TLS-protected connection to the Meraki dashboard and also uses TLS connections between its internal components within the Primary Policy Administrative Node (PPAN). For this internal connection to function properly, the ISE PPAN admin certificate must include a Subject Alternative Name (SAN) of type DNS with either the node FQDN or a wildcard FQDN. The ISE Meraki Connector does not support PPAN admin certificates that have only a Subject Name defined.

To obtain an API key with write permission, make sure you are logged in to the dashboard with a full-access admin account. After logging in, navigate to **Organization > API & Webhooks > API Keys and Access** and then click **Generate API Key**, as shown in Figure 12-13.

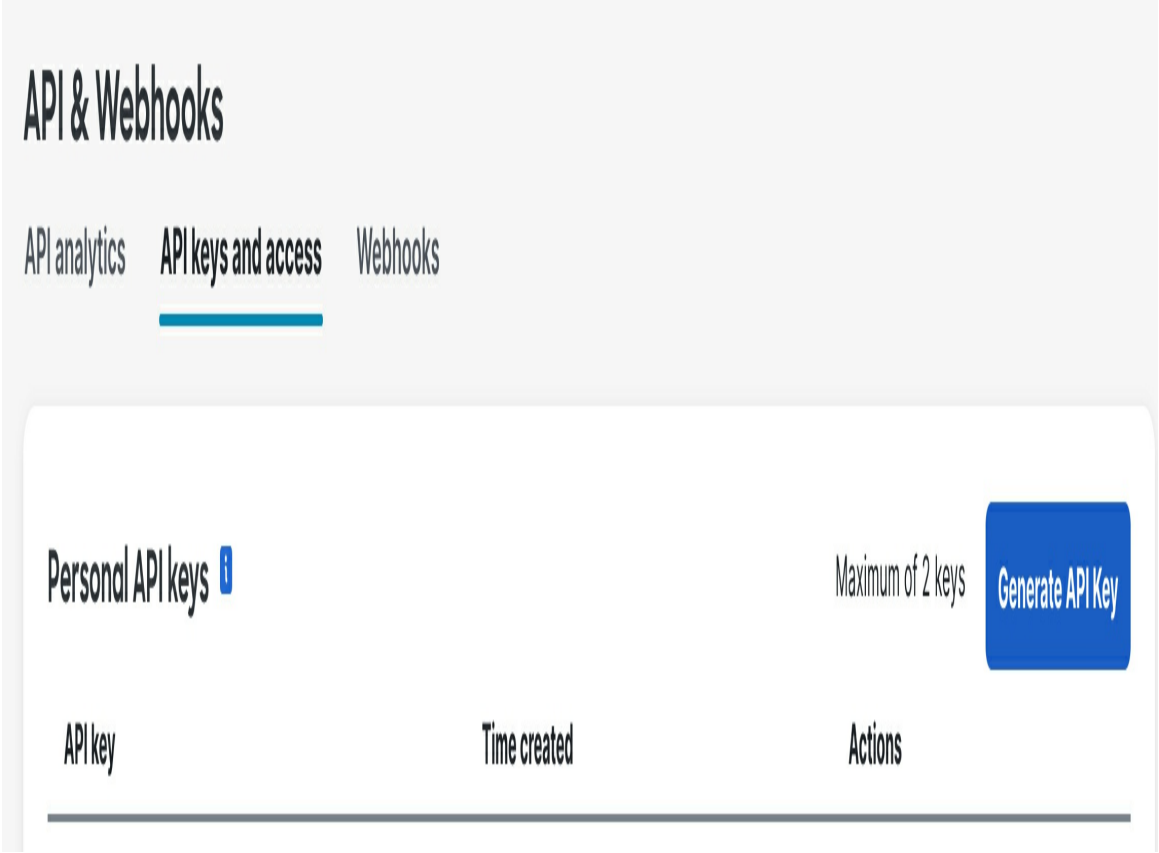


Figure 12-13 *Meraki API Key Generation*

At the time of writing, the Meraki API endpoint used for the ISE Meraki Connector is available in two dashboard regions. The Global: api.meraki.com with the server certificate chain is outlined in Example 12-1.

Example 12-1 *Global: api.meraki.com with Server Certificate Chain*

```
Certificate chain
 0 s:/CN=api.meraki.com
   i:/C=US/O=Google Trust Services/CN=WR1

 1 s:/C=US/O=Google Trust Services/CN=WR1
   i:/C=US/O=Google Trust Services LLC/CN=GTS Root R1

 2 s:/C=US/O=Google Trust Services LLC/CN=GTS Root R1
   i:/C=BE/O=GlobalSign nv-sa/OU=Root CA/CN=GlobalSign Root CA
```

Technet24

The China: api.meraki.ch with the server certificate chain is outlined in Example 12-2.

Example 12-2 *China: api.meraki.ch with the Server Certificate Chain*

```
Certificate chain
 0 s:/C=US/ST=California/L=San Jose/O=Cisco Systems Inc./CN=ios.m
   i:/C=US/O=IdenTrust/OU=HydrantID Trusted Certificate Service/C

 1 s:/C=US/O=IdenTrust/OU=HydrantID Trusted Certificate Service/C
   i:/C=US/O=IdenTrust/CN=IdenTrust Commercial Root CA 1

 2 s:/C=US/O=IdenTrust/CN=IdenTrust Commercial Root CA 1
   i:/C=US/O=IdenTrust/CN=IdenTrust Commercial Root CA 1
```

The ISE trusted store must have all certificates from the chain installed, except for the identity certificate of the server, with **Trust for Authentication Within ISE and Client-Server Communication** and **Trust for Authentication of Cisco Services** enabled. You can check whether certificates are already present in the ISE trusted store by navigating to **Administration > System > Certificates > Trusted Certificates.**

To verify the current status of the API endpoint certificate chain, use the following **openssl** command:

```
openssl s_client -connect api.meraki.com:443 -servername api.meraki.
showcerts
```

The **openssl s_client** tool with the **showcerts** option enabled displays all certificates returned by the web server in the chain in privacy-enhanced mail (PEM) format. You can save the certificate directly from the terminal and import it to ISE by clicking the **Import** button on the **Trusted Certificates** page.

After you have verified the prerequisites, you can proceed with configuring the ISE Meraki Connector. Navigate to **Work Centers > TrustSec > Integrations > Meraki > Overview** and click **Connect Meraki**, as demonstrated in Figure 12-14.

Technet24

Overview Components TrustSec Policy Policy Sets SXP Integrations Troubleshoot Reports Settings

Workload Connectors

Overview

Workload Connections

Attributes Dictionary

Workload Classification

Meraki

Overview

Sync Status

Connections

Sync Selections

Overview

Warning: Known issue with ISE-Meraki integration

You have carefully segmented your Cisco network by defining Cisco TrustSec policies and propagating them across the network using ISE. Now, you can easily extend these policies into the Meraki domain by using this Integration. Specify a set of policies that you want to replicate to one or more Meraki organizations, and the Integration will sync them automatically to Meraki Dashboard. From there, the policies will be propagated to your Meraki network devices.



1. Prepare

Identify the TrustSec policies currently configured in ISE that you wish to sync to Meraki

Identify the Meraki Organizations that you want to sync to. You must have access to each of these Organizations from your Meraki Dashboard account.

2. Define

Perform the setup workflow by pressing the Connect Meraki button.

Connect Meraki

In the workflow you will:

- Create a Meraki Dashboard connection.

3. Go Live & Monitor

When setup workflow is completed, each selected policy will be synced to each of the selected Organizations.

Thereafter, any policy changes made in ISE will be detected at the next sync interval, and the corresponding policies in your Meraki Organizations will be updated

Figure 12-14 *Starting Meraki Integration Wizard*

The connector wizard includes the following steps:

Step 1. **Welcome.** At this step, you receive basic information about the wizard.

Step 2. **Add Connections.** At this step, you define the API key and select the Meraki organization with which policies and other TrustSec elements will be synchronized.

Step 3. **Sync Interval.** Define the periodic sync interval.

Step 4. **Select Egress Policies.** Choose policies that require synchronization. This step determines what elements will be automatically selected in Steps 5 and 6.

Step 5. **Select ACLs.** ISE selects SGACLs according to the policies chosen in Step 4.

Step 6. **Select SGTs.** ISE selects SGTs according to the policies chosen in Step 4.

Step 7. **Summary.** A summary of the configuration that will be propagated to the Meraki organization is displayed.

The following example demonstrates the complete integration process:

Step 1. **Welcome.** At this step, shown in Figure 12-15, you receive basic information about the wizard.

Technet24

Overview Components TrustSec Policy Policy Sets SXP Integrations Troubleshoot Reports Settings

1 Welcome 2 Add Connections 3 Sync Interval 4 Select Egress Policies 5 Select ACLs 6 Select SGTs 7 Summary

Welcome

You can now sync SGT, ACL and Egress policies from ISE to Meraki! Once the setup is complete, you can make changes via the TrustSec page.



Exit All changes saved

Let's do it!

Figure 12-15 *The Welcome Step in ISE Meraki Connector Wizard*

Step 2. **Add Connections.** At this step, shown in Figure 12-16, you can modify the connection name (1), choose the API endpoint (2), enter your Meraki dashboard API key (3), click **Connect** and select the organization with which ISE should be integrated from the drop-down list (4). After you complete all required steps, you can click the **Next** button (5).

Technet24

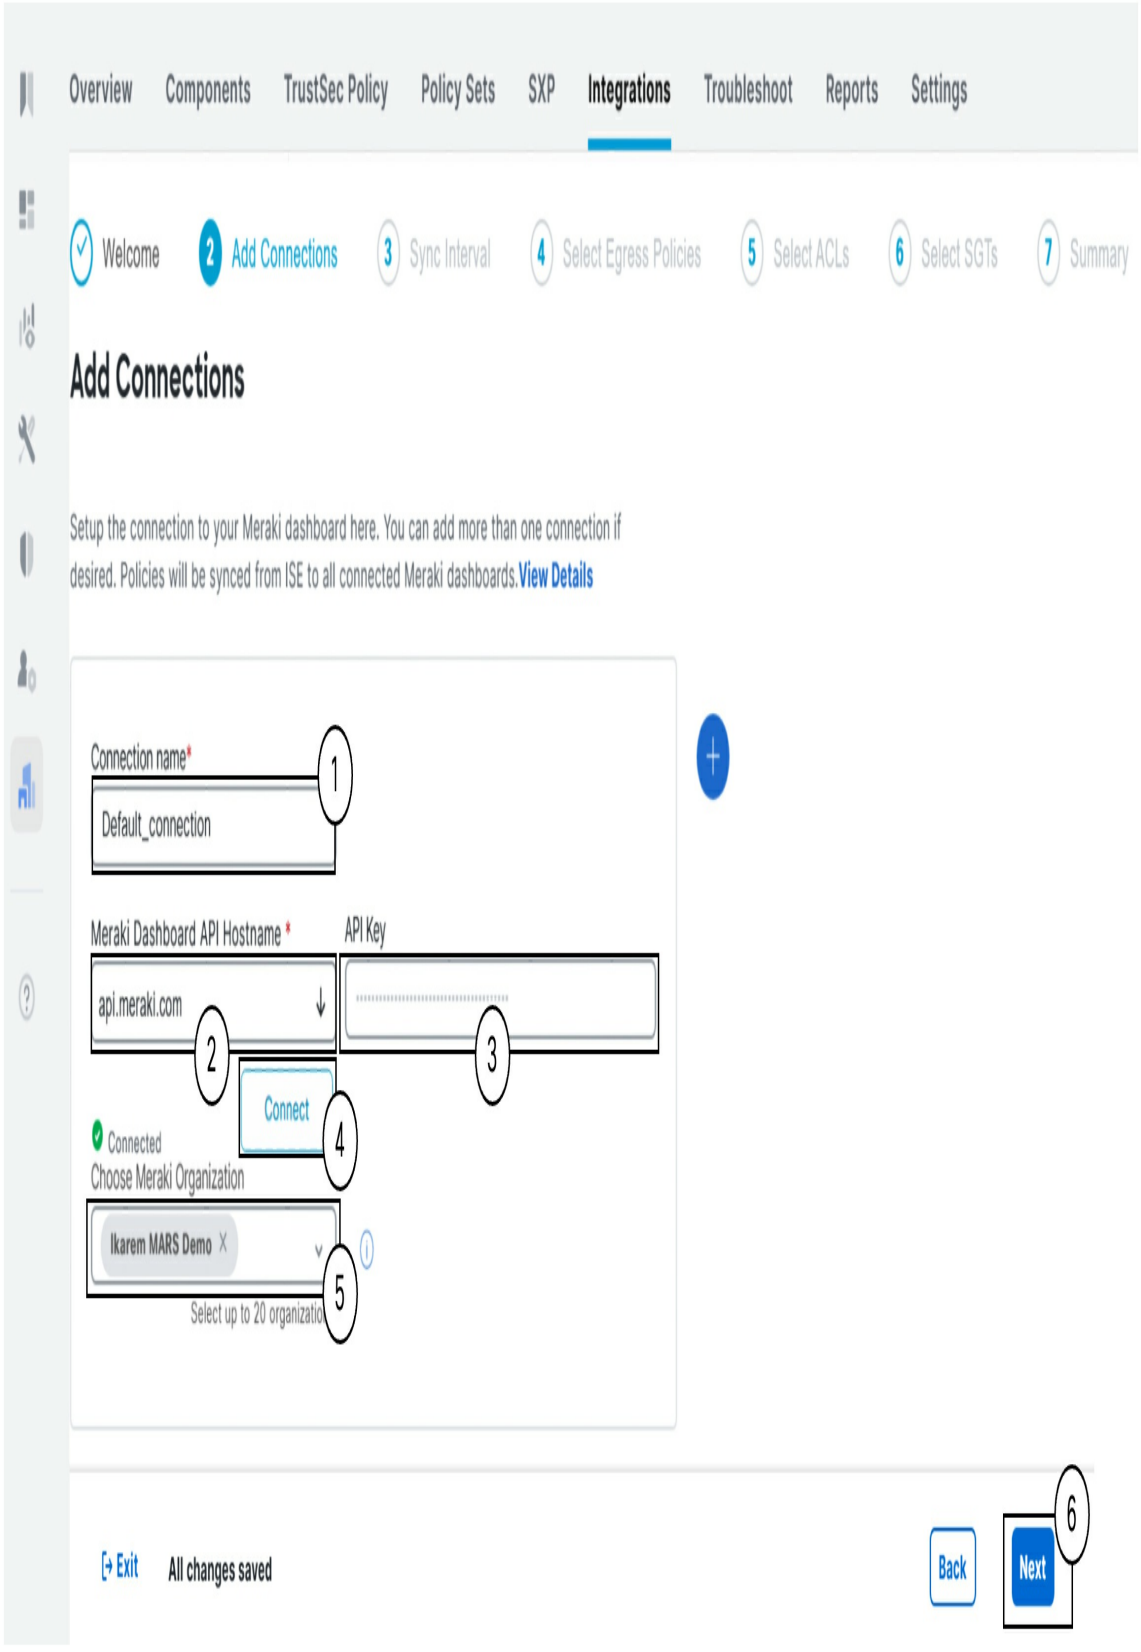

Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Welcome
2 Add Connections
3 Sync Interval
4 Select Egress Policies
5 Select ACLs
6 Select SGTs
7 Summary
Add Connections
Setup the connection to your Meraki dashboard here. You can add more than one connection if desired. Policies will be synced from ISE to all connected Meraki dashboards.View Details
Connection name*
Default_connection
1
Meraki Dashboard API Hostname *
API Key
api.meraki.com
2
3
Connect
4
Connected
Choose Meraki Organization
Ikarem MARS Demo X
5
Select up to 20 organizatio
Exit
All changes saved
Back
Next
6

Figure 12-16 *The Add Connections Step in ISE Meraki Connector Wizard*

Step 3. **Sync Interval.** In this step, shown in Figure 12-17, you can specify a periodic interval (1) after which ISE will perform synchronization verification and propagate any new changes to the Meraki dashboard.

Technet24


Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Welcome
Add Connections
3 Sync Interval
4 Select Egress Policies
5 Select ACLs
6 Select SGTs
7 Summary
Set up Sync Interval
Set up the time for Interval Syncing. You will be able to change interval and sync manually after.
Sync Interval*
12
1
5 - 30 (Minutes)
Exit
All changes saved
Back
Next

Figure 12-17 *The Sync Interval Step in ISE Meraki Connector Wizard*

Step 4. **Select Egress Policies.** In this step, shown in Figure 12-18, you are shown the list of egress policies (source SGT/destination SGT pairs) for which the ISE administrator has specified SGACLs. Select the egress policies you wish to propagate to the Meraki dashboard (1) and click **Next** (2).

Technet24

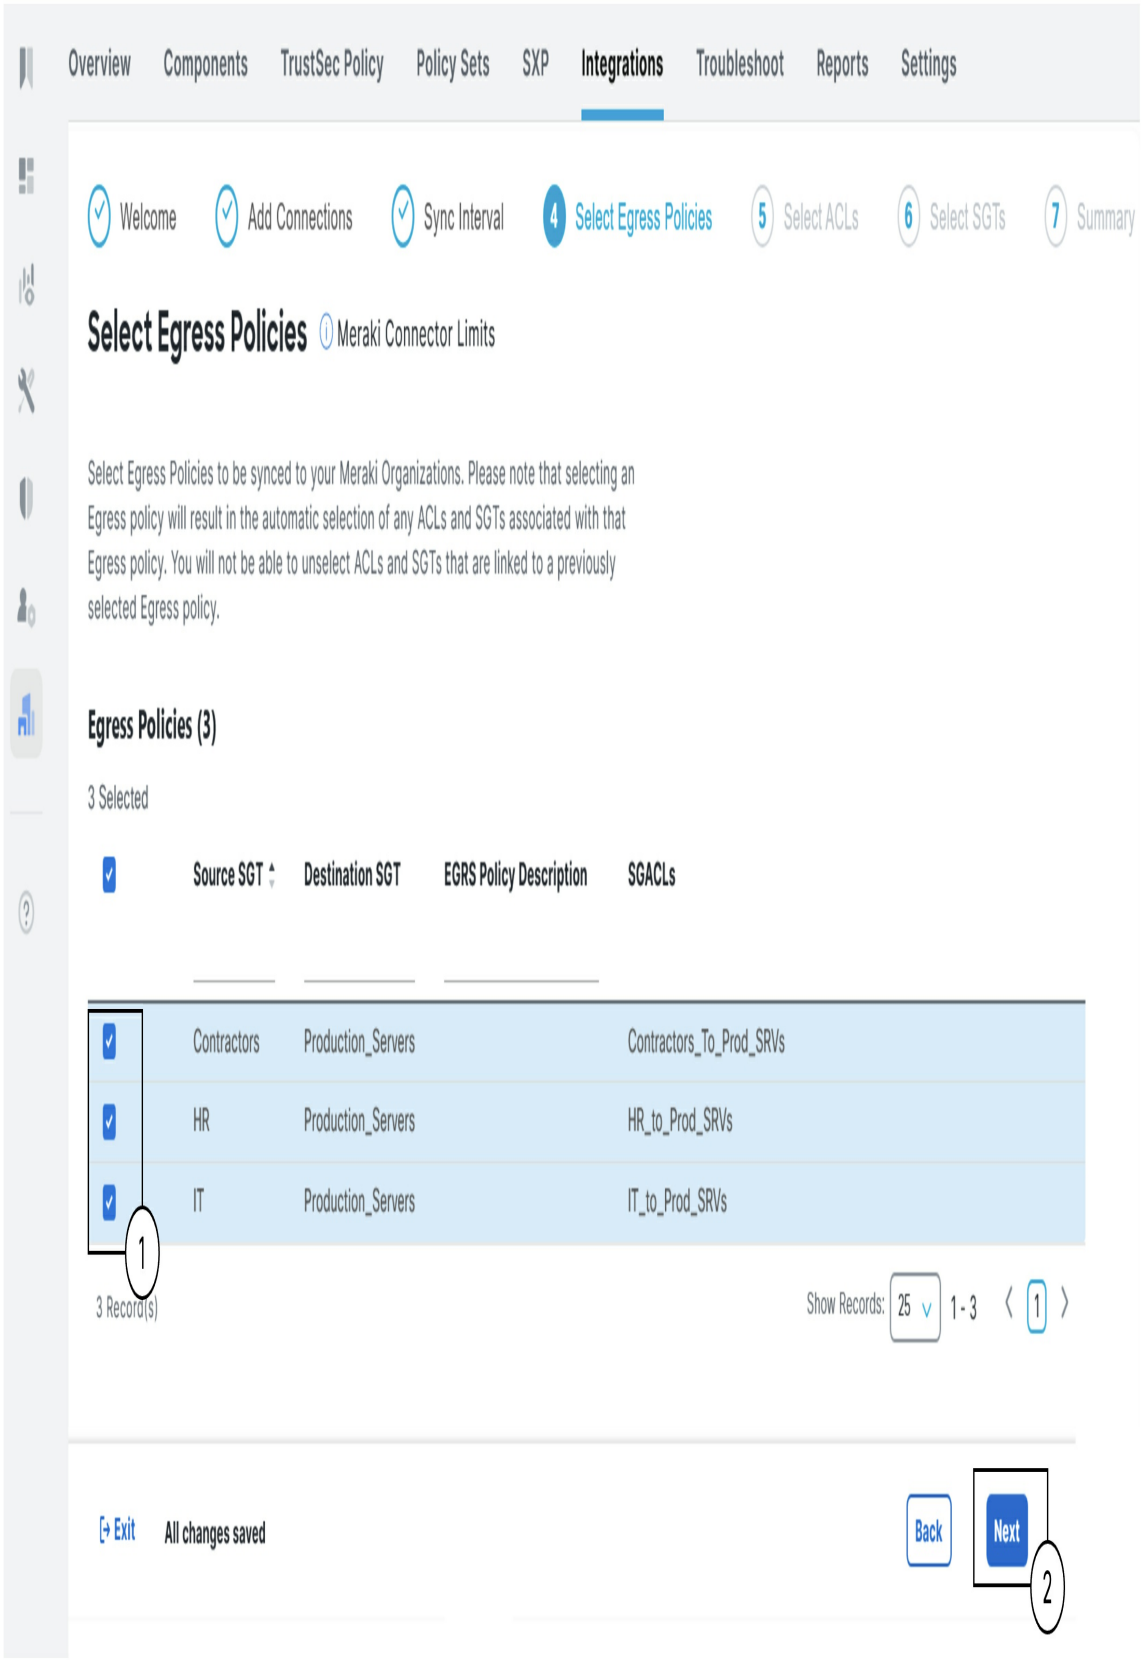

Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Welcome
Add Connections
Sync Interval
4 Select Egress Policies
5 Select ACLs
6 Select SGTs
7 Summary
Select Egress Policies
Meraki Connector Limits
Select Egress Policies to be synced to your Meraki Organizations. Please note that selecting an Egress policy will result in the automatic selection of any ACLs and SGTs associated with that Egress policy. You will not be able to unselect ACLs and SGTs that are linked to a previously selected Egress policy.
Egress Policies (3)
3 Selected
Source SGT
Destination SGT
EGRS Policy Description
SGACLs
Contractors
Production_Servers
Contractors_To_Prod_SRVs
HR
Production_Servers
HR_to_Prod_SRVs
IT
Production_Servers
IT_to_Prod_SRVs
1
3 Record(s)
Show Records: 25
1 - 3
1
Exit
All changes saved
Back
Next
2

Figure 12-18 *The Select Egress Policies Step in ISE Meraki Connector Wizard*

Step 5. **Select ACLs.** In this step, shown in Figure 12-19, you are presented with all SGACLs created by the ISE administrator. The SGACLs associated with the policies selected in the previous step are automatically selected and cannot be unchecked. You may mark any additional SGACLs for synchronization if needed. If only the ACLs related to the previously selected policies are required on the Meraki side, click **Next** (1).

Technet24

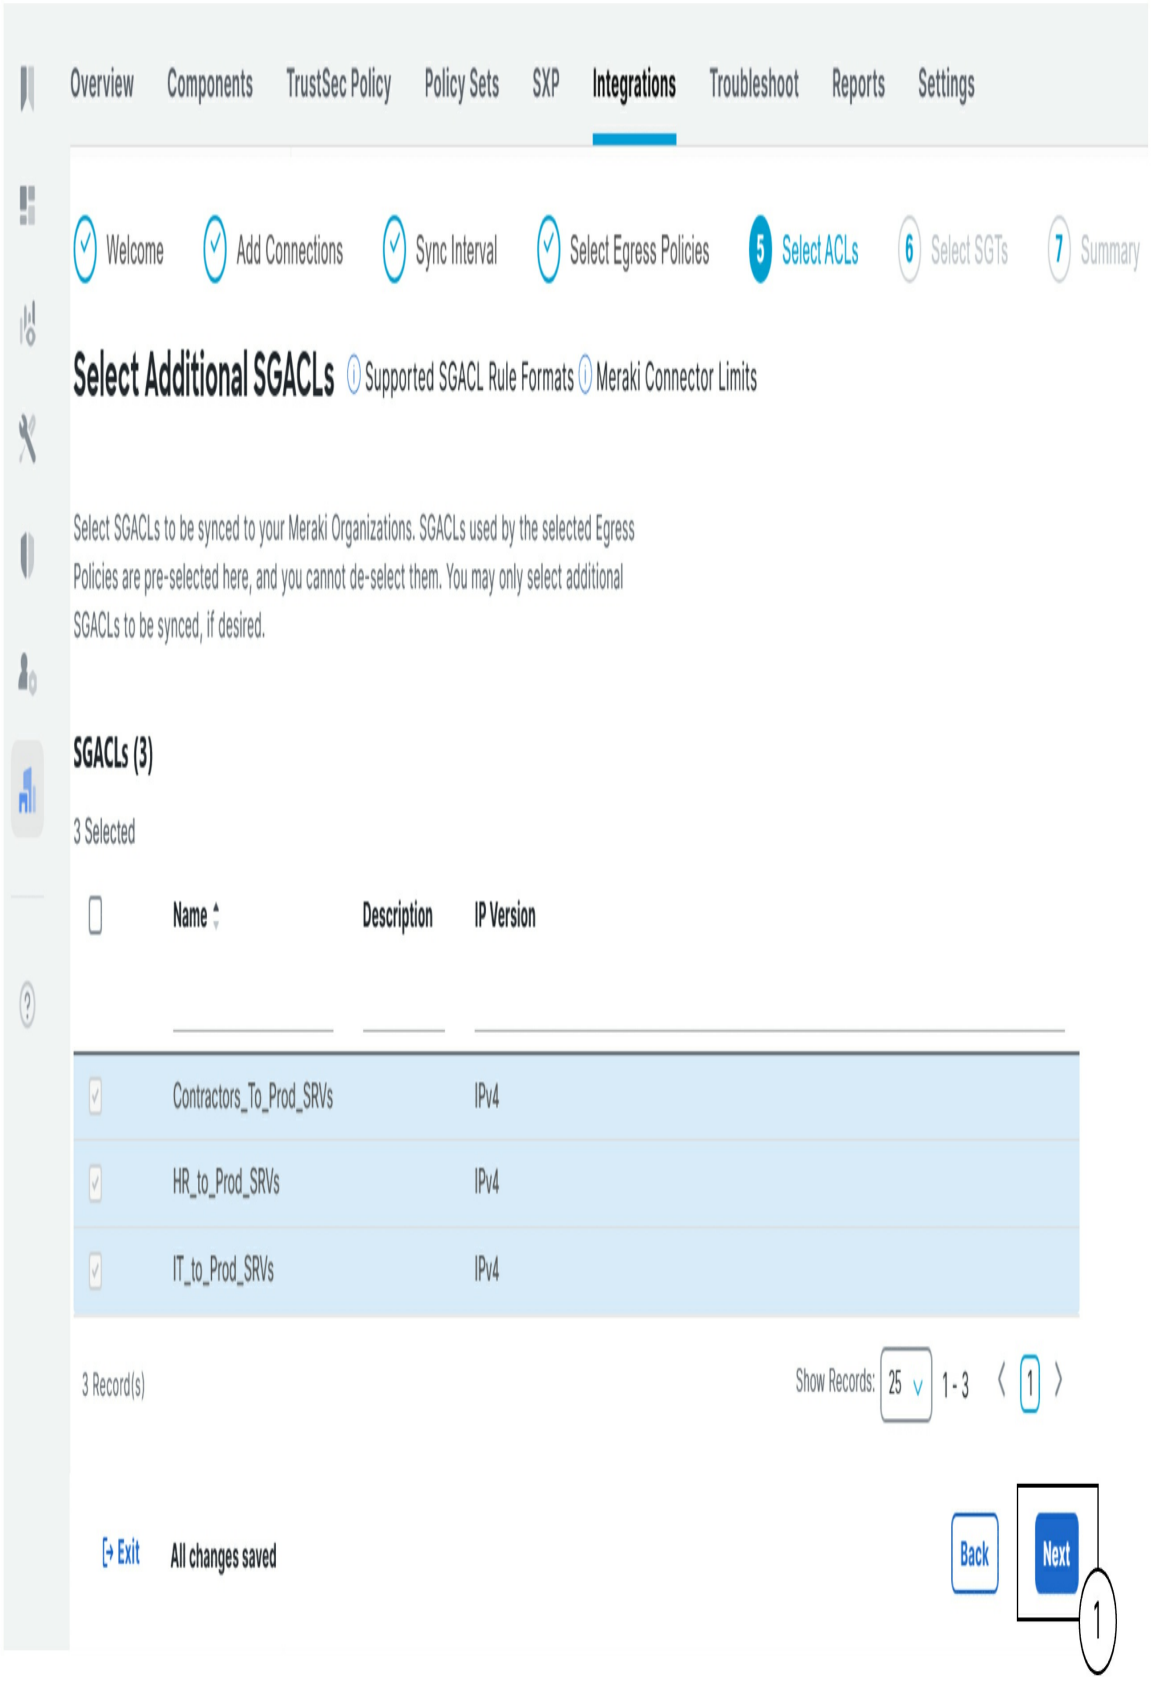

Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Welcome
Add Connections
Sync Interval
Select Egress Policies
5 Select ACLs
6 Select SGTs
7 Summary
Select Additional SGACLs
Supported SGACL Rule Formats
Meraki Connector Limits
Select SGACLs to be synced to your Meraki Organizations. SGACLs used by the selected Egress Policies are pre-selected here, and you cannot de-select them. You may only select additional SGACLs to be synced, if desired.
SGACLs (3)
3 Selected
Name
Description
IP Version
Contractors_To_Prod_SRVs
IPv4
HR_to_Prod_SRVs
IPv4
IT_to_Prod_SRVs
IPv4
3 Record(s)
Show Records: 25
1 - 3
Exit
All changes saved
Back
Next
1

Figure 12-19 *The Select ACLs Step in ISE Meraki Connector Wizard*

Step 6. **Select SGTs.** In this step, shown in Figure 12-20, you see a list of all SGTs created by the ISE administrator. SGTs tied to the policies chosen in the previous step are automatically selected and cannot be removed from selection. If needed, you can select additional SGTs for synchronization. If you only require SGTs associated with the previously selected policies on the Meraki side, click **Next** (1).

Technet24

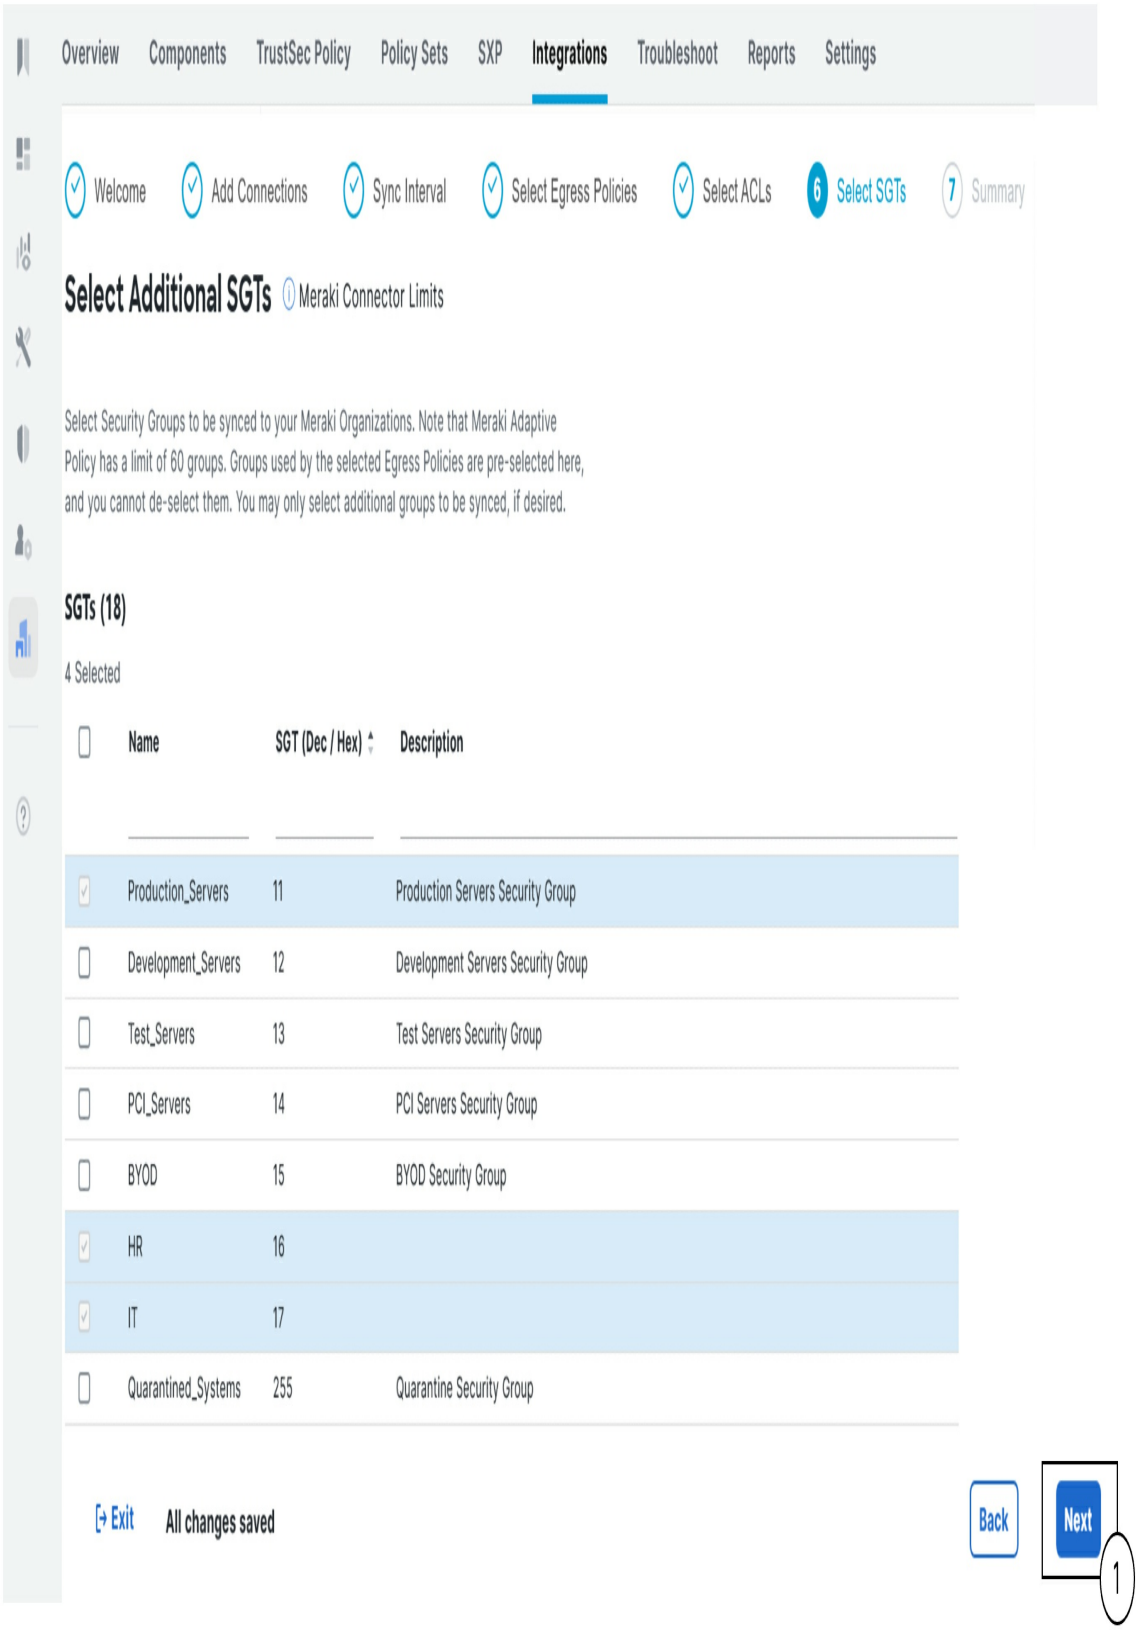

Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Welcome
Add Connections
Sync Interval
Select Egress Policies
Select ACLs
6 Select SGTs
7 Summary
Select Additional SGTs
Meraki Connector Limits
Select Security Groups to be synced to your Meraki Organizations. Note that Meraki Adaptive Policy has a limit of 60 groups. Groups used by the selected Egress Policies are pre-selected here, and you cannot de-select them. You may only select additional groups to be synced, if desired.
SGTs (18)
4 Selected
Name
SGT (Dec / Hex)
Description
Production_Servers 11 Production Servers Security Group
Development_Servers 12 Development Servers Security Group
Test_Servers 13 Test Servers Security Group
PCI_Servers 14 PCI Servers Security Group
BYOD 15 BYOD Security Group
HR 16
IT 17
Quarantined_Systems 255 Quarantine Security Group
Exit
All changes saved
Back
Next
1

Figure 12-20 *Select SGTs Step in ISE Meraki Connector Wizard*

Step 7. **Summary.** At this step, shown in Figure 12-21, a summary of the wizard is displayed. Click **Next** (1) to initiate configuration synchronization.

Technet24

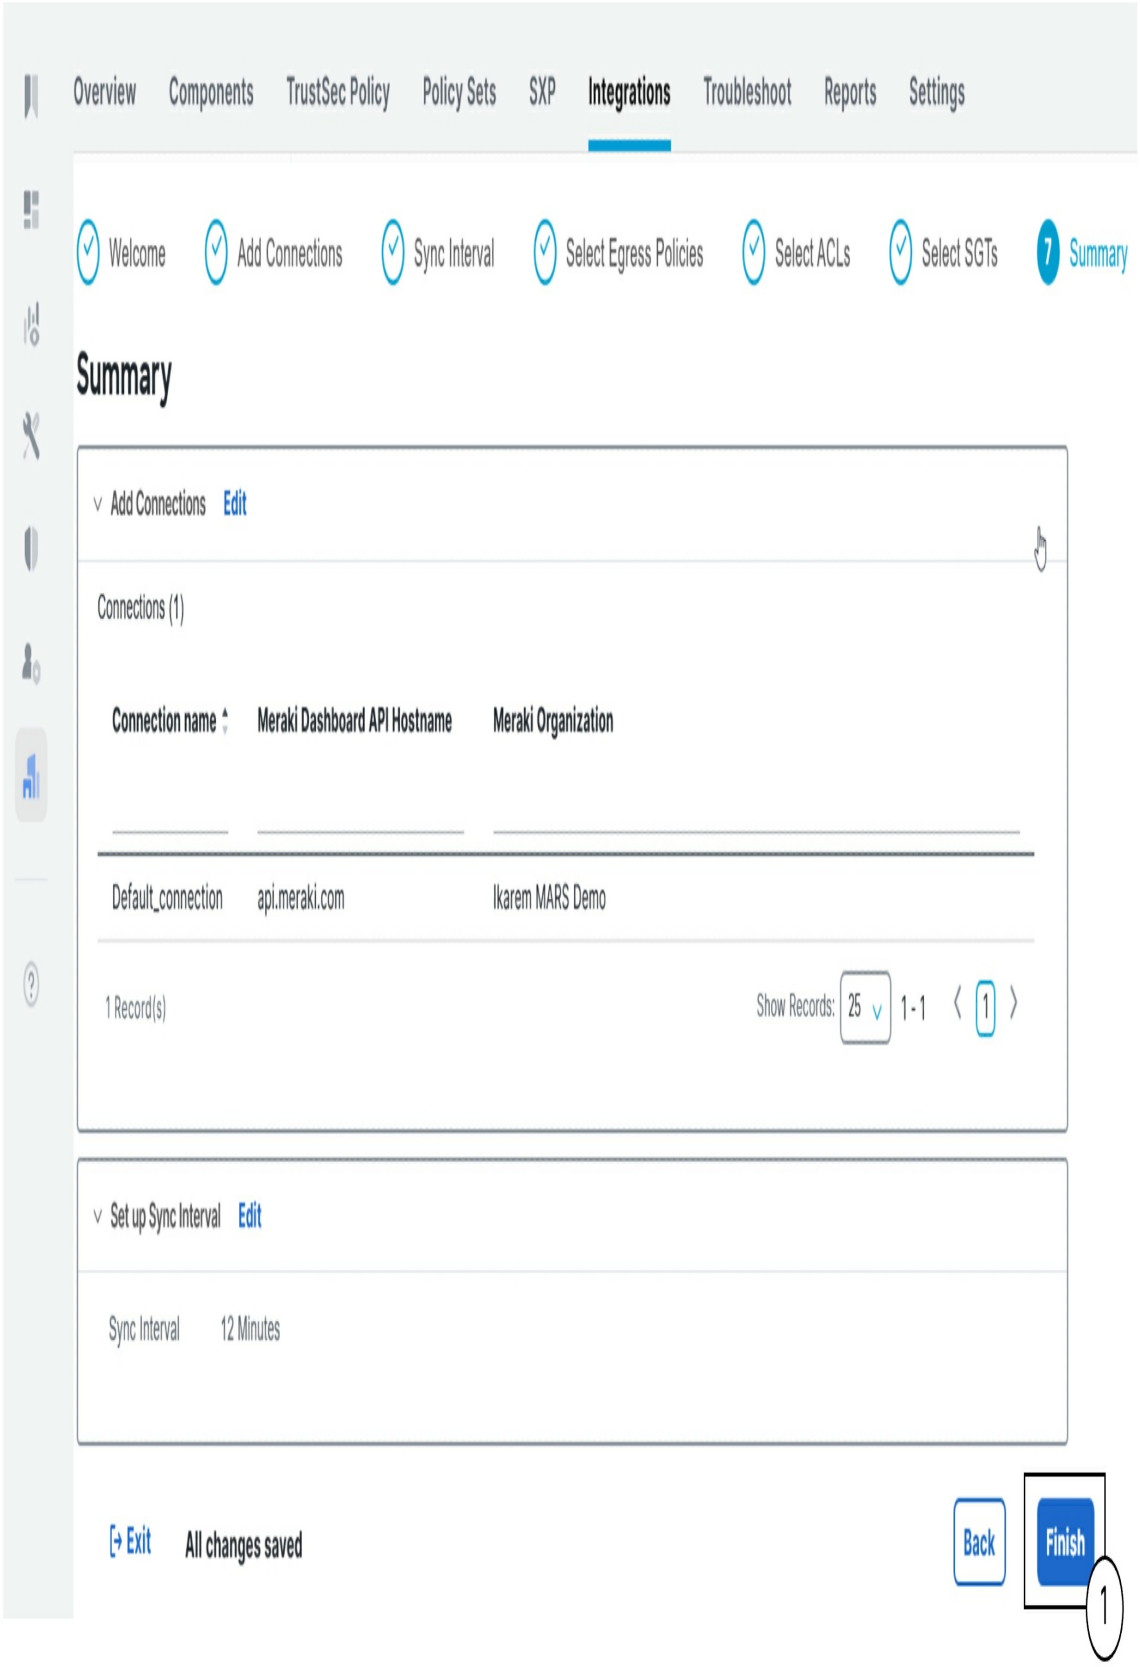

Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Welcome
Add Connections
Sync Interval
Select Egress Policies
Select ACLs
Select SGTs
7
Summary
Summary
Add Connections
Edit
Connections (1)
Connection name
Meraki Dashboard API Hostname
Meraki Organization
Default_connection
api.meraki.com
Ikarem MARS Demo
1 Record(s)
Show Records:
25
1 - 1
1
Set up Sync Interval
Edit
Sync Interval
12 Minutes
Exit
All changes saved
Back
Finish
1

Figure 12-21 *The Summary Step in ISE Meraki Connector Wizard*

After completing the Meraki Connector configuration wizard, you can verify the synchronization status by navigating to **Work Centers > TrustSec > Integrations > Meraki > Sync Status**. Figure 12-22 shows a successful sync.

Technet24


Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Workload Connectors
Meraki
Overview
Sync Status
Connections
Sync Selections
Sync Status
Sync Interval: Every 12 Minutes
Data Sync in: 4:01
Sync Now
Pause Sync
View sync status summary. Check remediations of failed for sync Egress Policies, ACLs and SGTs. Check each Cloud Connection status on Connections.
ISE TO MERAKI SYNC
3/3
Egress Policies
3/3
ACLs
4/4
SGTs
ORGANIZATIONS
Last Synchronized Time: Mar 22, 2026 2:17 PM
1
Fully Synced
0
Partially Synced
0
Failed to Sync

Figure 12-22 *ISE Meraki Connector Sync Status*

Implementing Adaptive Policy at the Organization Level

Most Adaptive Policy settings must be configured at the organization level, except for port-level settings on specific devices to enable Adaptive Policy and certain group mapping tasks, such as VLAN-to-group mapping and SSID-to-group mapping. These settings are performed either on individual devices or at the network level. From an organizational perspective, you, as the administrator, need to address the following main areas of configuration:

- Enable Adaptive Policy for specific networks.
- Configure Adaptive Policy groups.
- Define Adaptive Policy objects for static IP to SGT mappings.
- Create and manage Adaptive Policy ACLs.
- Define Adaptive Policies.

Enabling Adaptive Policy for Specific Networks

To enable Adaptive Policy for the networks requiring microsegmentation, navigate to **Organization > Adaptive Policy > Networks**. Select the networks where Adaptive Policy should be enabled and click **Enable**. At this step, the licenses are verified. If you enable Adaptive Policy for a combined network, licenses for MX, MS, and MR must be present. If you do not plan to extend the network with other device types, it is recommended to create a network for each specific device type. Figure 12-23 shows how to enable Adaptive Policy for specific networks.

Technet24


MSP Portal
View old version
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Organization
Ikarem MARS Demo
Network
alexo-sc-spoke-template
Secure Connect
Network-wide
Cellular Gateway
Security & SD-WAN
Switching
Wireless
Cameras
Sensors
Insight
Organization
Search
Status
9 results
3 items selected
Select all 9 items
Cancel
Enable
Disable
Networks
Enablement
HQ-Brownfield-POL
Enabled
DC-EMEA-1
Disabled
Branch-Greenfield-POL
Enabled
Fabric-1-POL
Disabled
Switch-Upgrade
Disabled
DC-EMEA-2
Disabled
Secure Connect
Disabled
Azure vMX 1
Disabled
ESX-vMX-1-POL
Disabled
Rows per page
30
1-9 of 9
1

Figure 12-23 *Adaptive Policy Networks Tab*

Configuring Adaptive Policy Groups

To create a new Adaptive Policy group, navigate to **Organization > Adaptive Policy > Groups** and click **+ Add Group**. Figure 12-24 shows the Adaptive Policy Groups tab.

Technet24

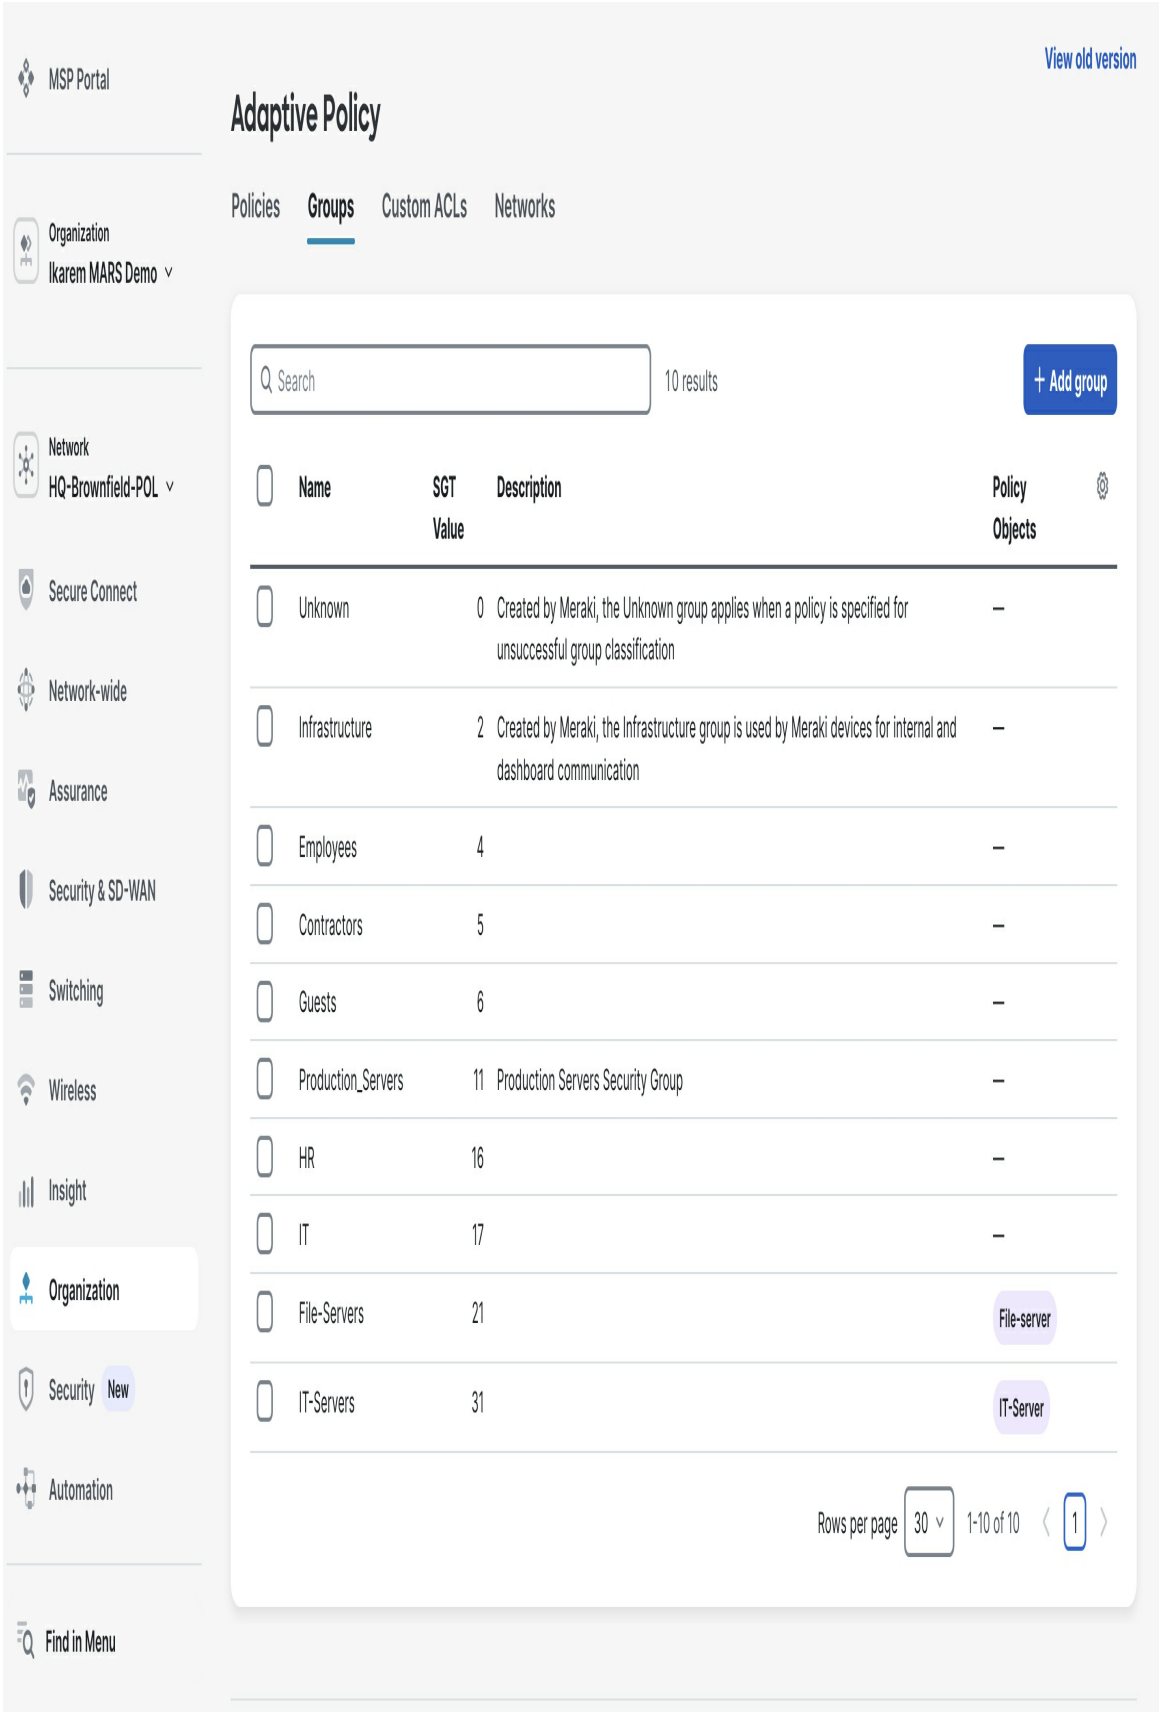

MSP Portal
Organization
Ikarem MARS Demo
Network
HQ-Brownfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security
New
Automation
Find in Menu
View old version
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Search
10 results
+ Add group
Name
SGT Value
Description
Policy Objects
Unknown
0
Created by Meraki, the Unknown group applies when a policy is specified for unsuccessful group classification
–
Infrastructure
2
Created by Meraki, the Infrastructure group is used by Meraki devices for internal and dashboard communication
–
Employees
4
–
Contractors
5
–
Guests
6
–
Production_Servers
11
Production Servers Security Group
–
HR
16
–
IT
17
–
File-Servers
21
File-server
IT-Servers
31
IT-Server
Rows per page
30
1-10 of 10
1

Figure 12-24 *Adaptive Policy Groups Tab*

In the new window, enter the Adaptive Policy group name (1), SGT value (2), and description (3) if needed; then click **Create** (4). Figure 12-25 shows a new Adaptive Policy group configuration.

Technet24

Add Group

Name

RND-Users

1

SGT Value

101

2

Description

3

Policy Object Binding

Select

Cancel

Create

3

Figure 12-25 *Creating an Adaptive Policy Group*

Note

When the Meraki dashboard and ISE are integrated, you should avoid creating any Adaptive Policy configuration in the Meraki dashboard except for static IP to SGT mapping. If you use ISE and Meraki separately, make sure to use the same SGT numbers for the same group names on both platforms.

To manage existing Adaptive Policy groups, such as performing edit or delete operations, select the group in the list and click the required action in the top-right corner of the **Adaptive Policy Groups** tab, as demonstrated in Figure 12-26.

Technet24

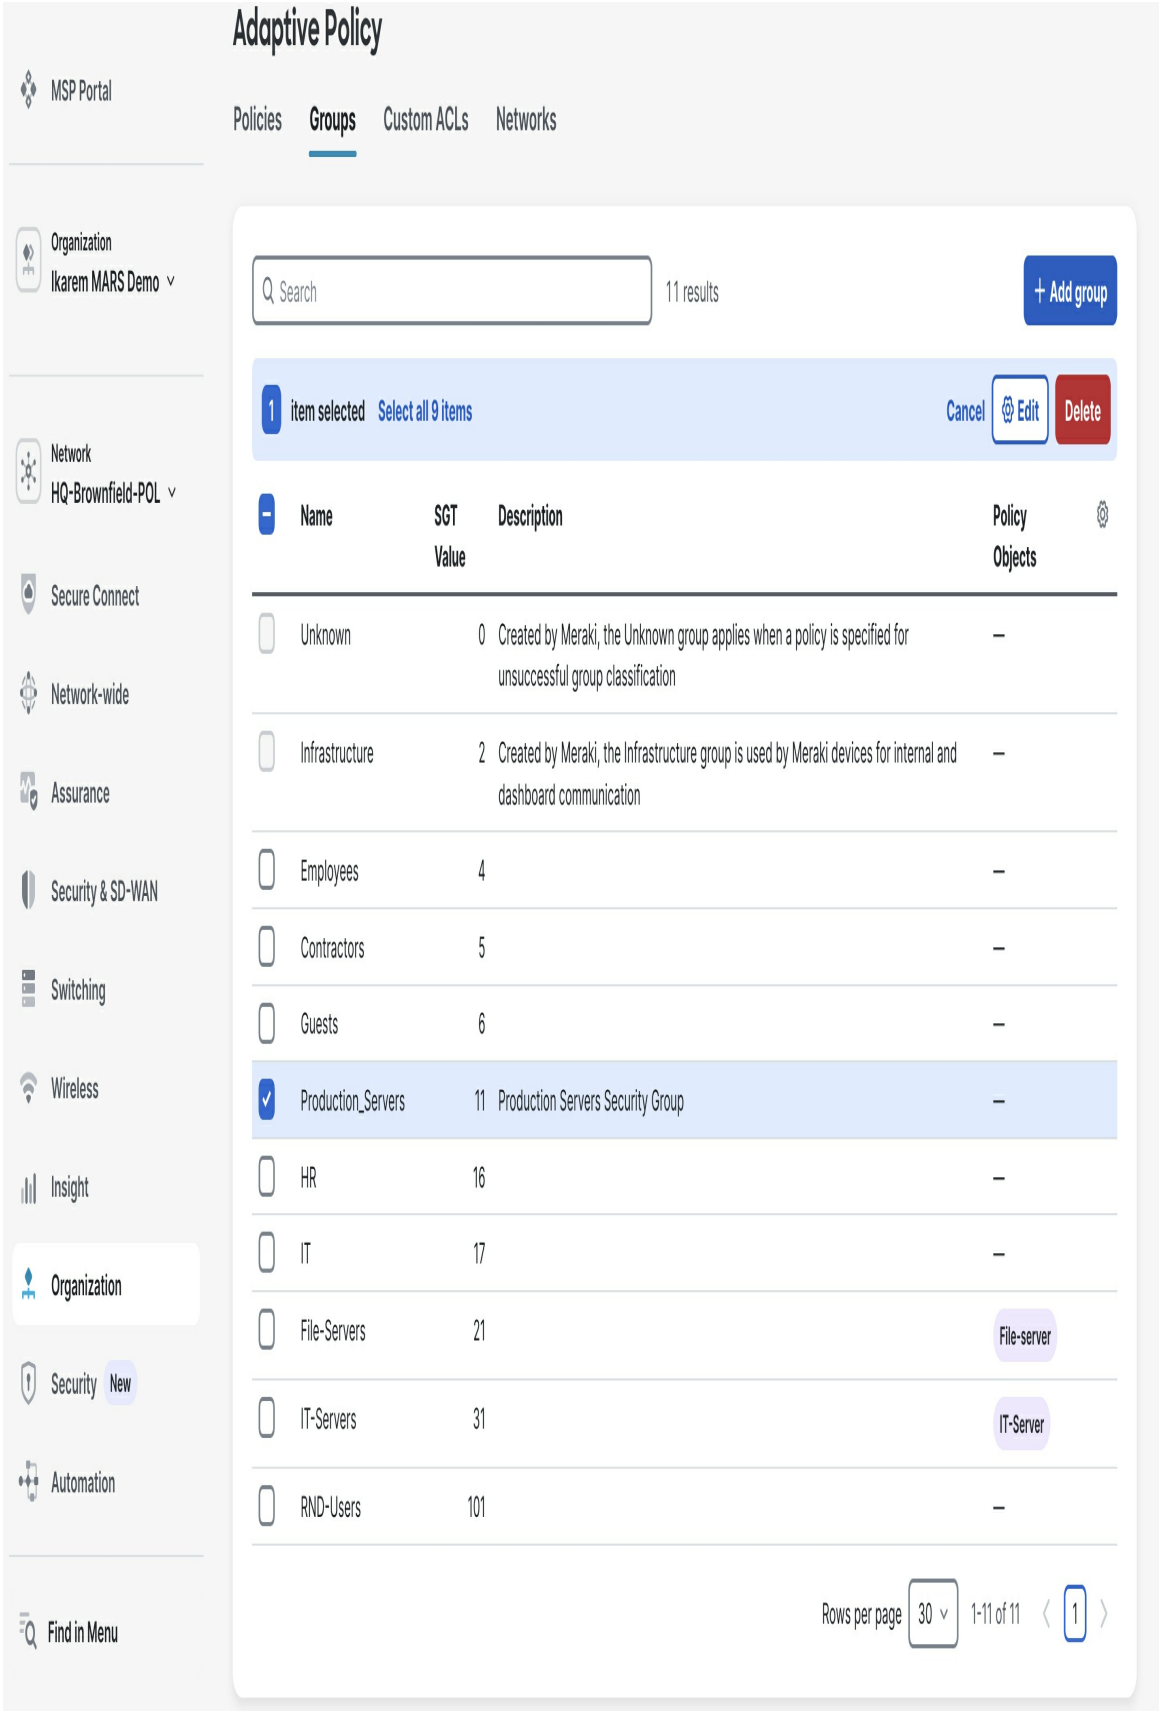

MSP Portal
Organization
Ikarem MARS Demo
Network
HQ-Brownfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security
New
Automation
Find in Menu
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Search
11 results
+ Add group
1 item selected
Select all 9 items
Cancel
Edit
Delete
Name
SGT Value
Description
Policy Objects
Unknown 0 Created by Meraki, the Unknown group applies when a policy is specified for unsuccessful group classification –
Infrastructure 2 Created by Meraki, the Infrastructure group is used by Meraki devices for internal and dashboard communication –
Employees 4 –
Contractors 5 –
Guests 6 –
Production_Servers 11 Production Servers Security Group –
HR 16 –
IT 17 –
File-Servers 21 File-server
IT-Servers 31 IT-Server
RND-Users 101 –
Rows per page 30
1-11 of 11
1

Figure 12-26 *Managing Adaptive Policy Groups*

Defining Adaptive Policy Objects for Static IP to SGT Mappings

Adaptive Policy objects are used to map IP addresses in Classless Inter-Domain Routing (CIDR) format, represented by IP address and prefix length, such as IP address/prefix length, to Adaptive Policy groups. This approach creates a static mapping between the host IP and SGT when the prefix length equals 32, and between subnet IP and SGT for any other prefix length.

Mapping configurations include two steps:

Step 1. Create a policy object of the Adaptive Policy type.

Step 2. Associate the policy object with an Adaptive Policy group.

The steps are outlined here in more detail:

Step 1. Create a policy object of the Adaptive Policy type. To create a new object, navigate to **Organization > Policy Objects** and click the **Add New** button on the **All Objects** tab (see Figure 12-27).

Technet24


MSP Portal
Organization
Ikarem MARS Demo
Network
HQ-Brownfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Policy Objects
Groups
All objects
Search...
Add new
Name Type Value Used in Last updated
1 IT-Server Adaptive Policy CIDR 10.110.20.11 Mar 12, 2026
2 File-server Adaptive Policy CIDR 10.110.20.31 Mar 12, 2026
2 results

Figure 12-27 *Policy Objects All Objects Tab*

In the new window, enter the object name (1), select **Adaptive Policy** (2) in the **Category** section, and specify the IP address in CIDR format (3). When all required fields are populated, click **Create Object** (4). Figure 12-28 shows an Adaptive Policy object configuration.

Technet24

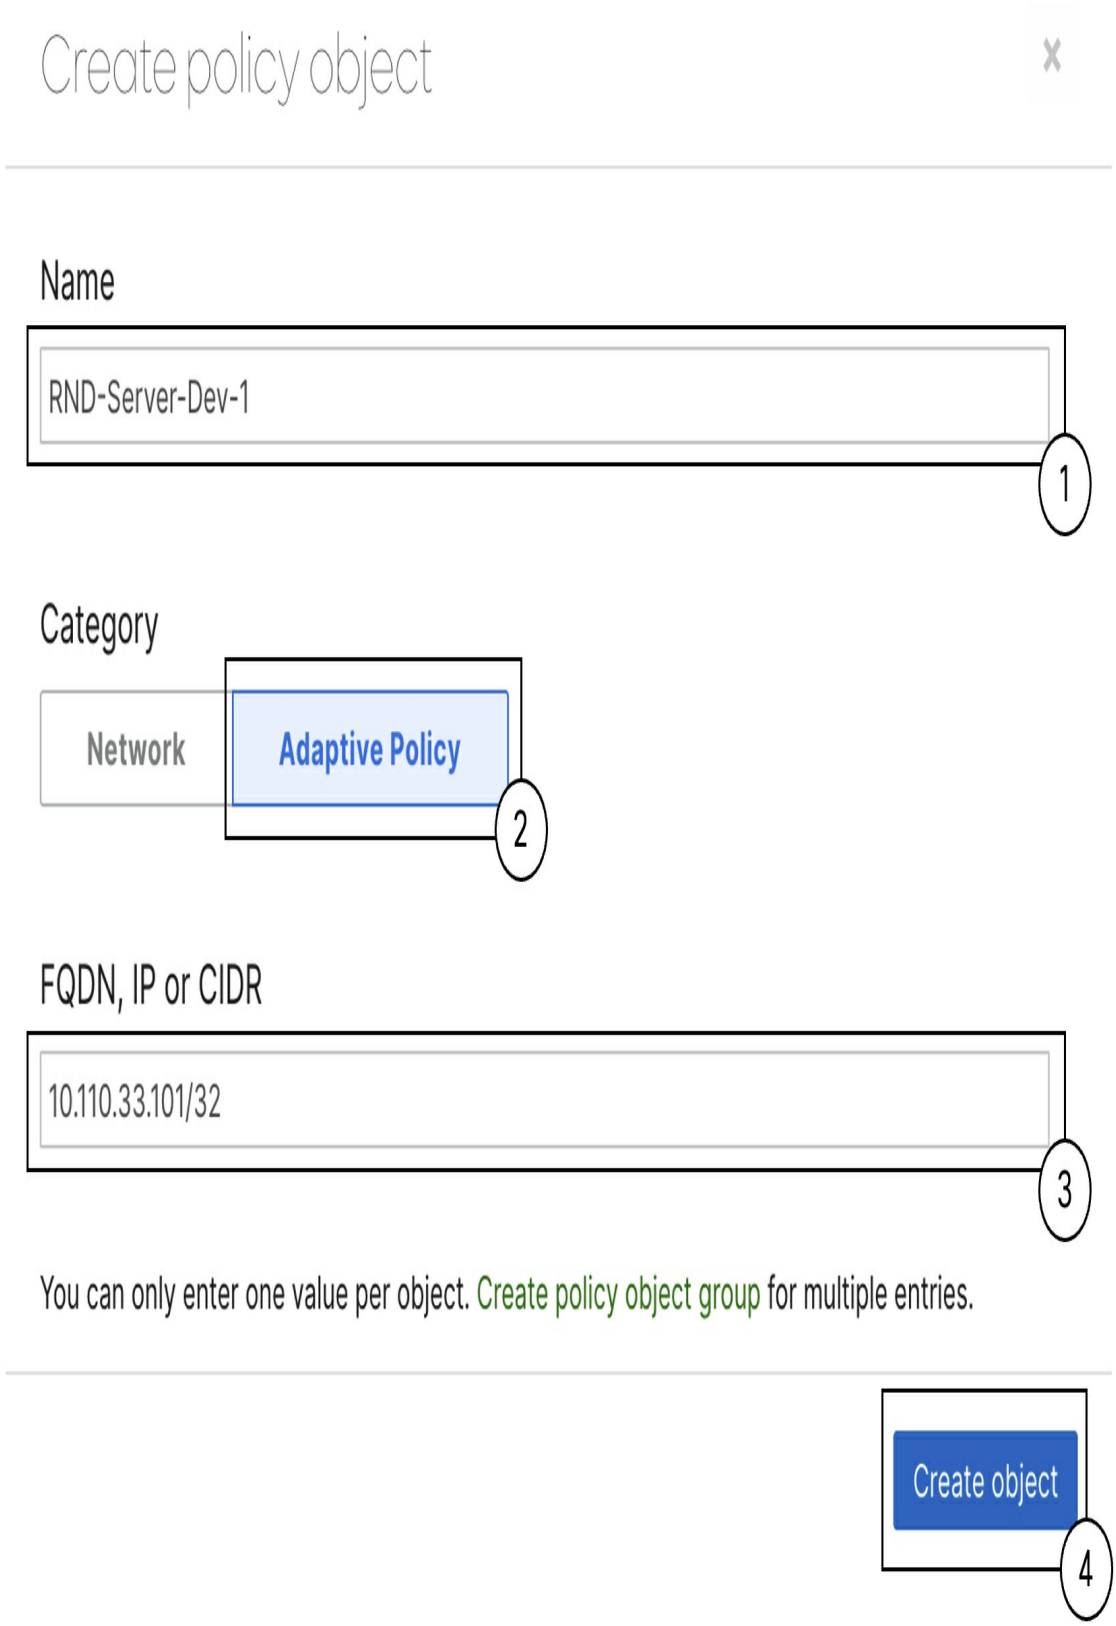

Create policy object
Name
RND-Server-Dev-1
1
Category
Network
Adaptive Policy
2
FQDN, IP or CIDR
10.110.33.101/32
3
You can only enter one value per object. Create policy object group for multiple entries.
Create object
4

Figure 12-28 *Create Policy Object Example*

Step 2. Associate the policy object with an Adaptive Policy group. To associate the policy object with an Adaptive Policy group, navigate to **Organization > Adaptive Policy > Groups**. Select the group where the new object needs to be added and click **Edit.** In the new window, you can select the object from the drop-down list called **Policy Object Binding**, as demonstrated in Figure 12-29.

Technet24

Edit Group

Name

Production_Servers

SGT Value ⓘ

11

Description

Production Servers Security Group

Policy Object Binding ⓘ

1 Select

- [ ] File-server
- [ ] IT-Server
- [x] RND-Server-Dev-1

Figure 12-29 *Associating the Policy Object with an Adaptive Policy Group*

Creating and Managing Adaptive Policy ACLs

SGACLs in the Meraki dashboard are managed from the **Custom ACLs** tab located in **Organization > Adaptive Policy**. By navigating to this tab, you can select existing ACLs in the list for editing or deletion, as demonstrated in Figure 12-30.

Technet24

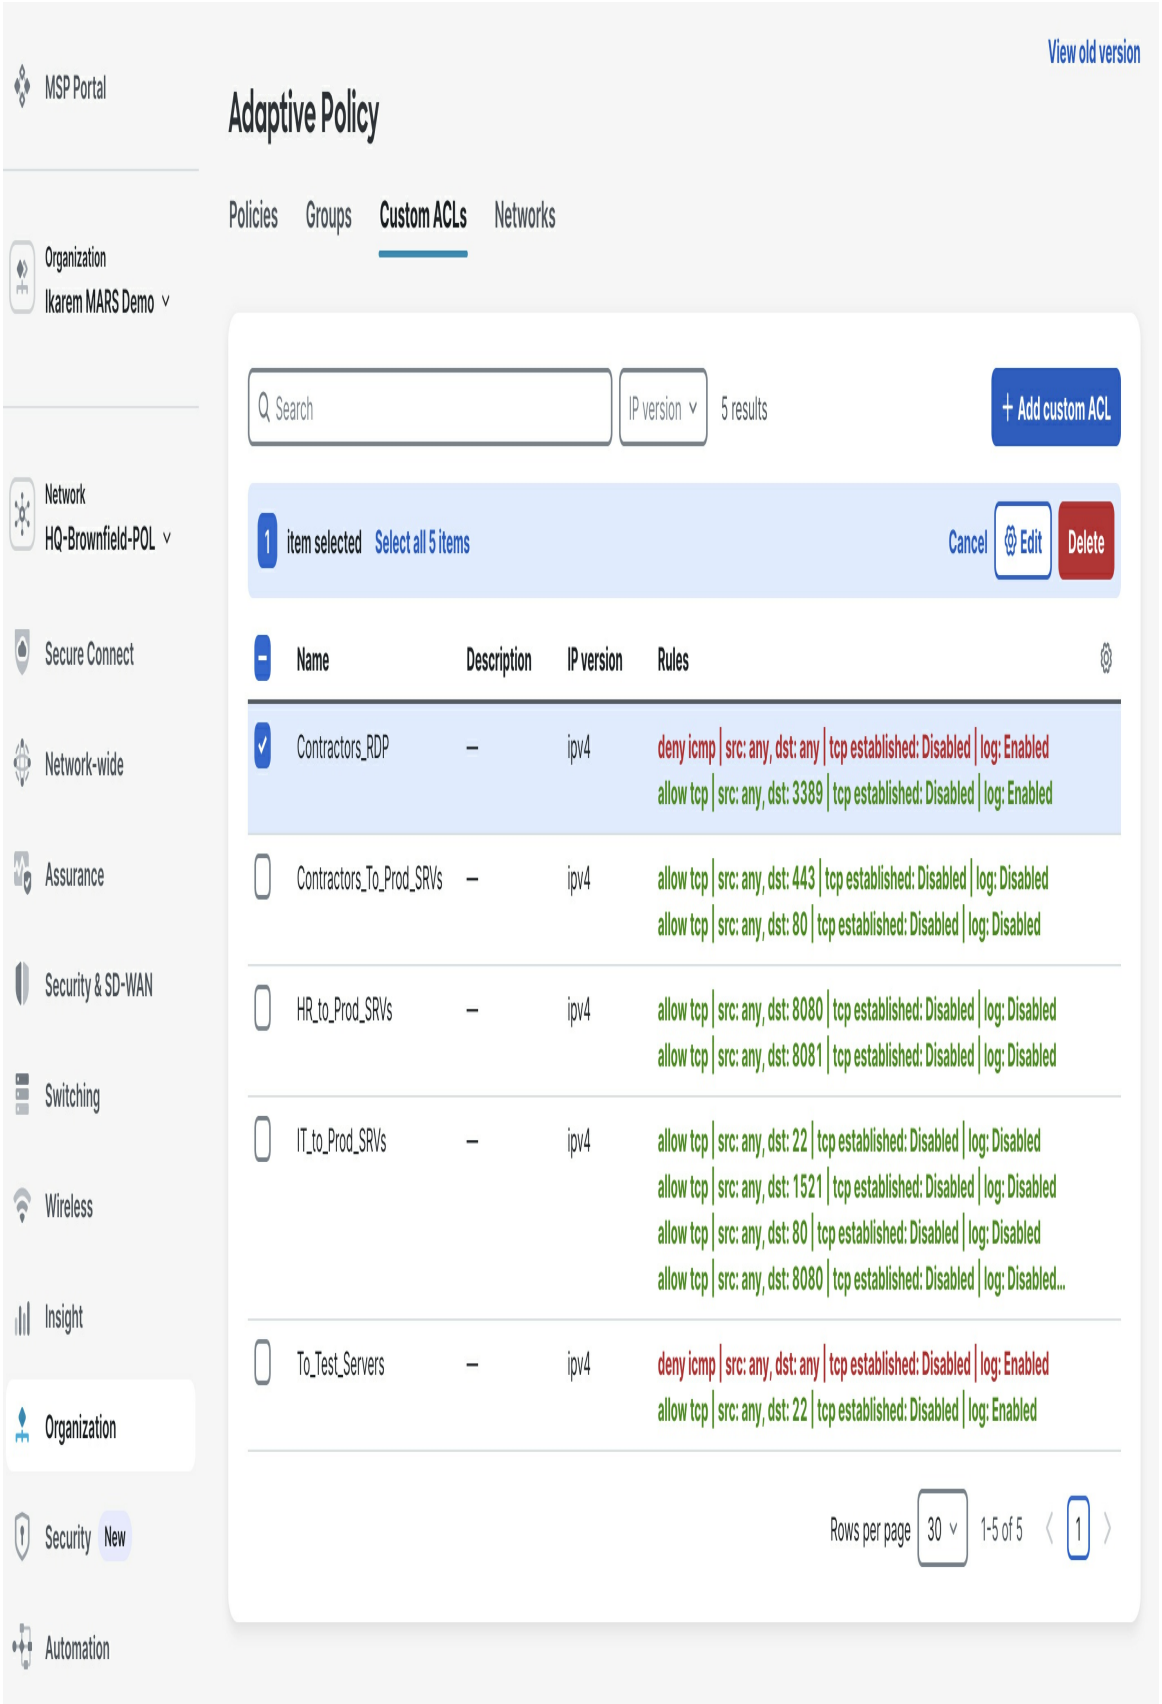

View old version
MSP Portal
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Organization
Ikarem MARS Demo
Search
IP version
5 results
+ Add custom ACL
Network
HQ-Brownfield-POL
1 item selected
Select all 5 items
Cancel
Edit
Delete
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security
New
Automation
Name
Description
IP version
Rules
Contractors_RDP
–
ipv4
deny icmp | src: any, dst: any | tcp established: Disabled | log: Enabled
allow tcp | src: any, dst: 3389 | tcp established: Disabled | log: Enabled
Contractors_To_Prod_SRVs
–
ipv4
allow tcp | src: any, dst: 443 | tcp established: Disabled | log: Disabled
allow tcp | src: any, dst: 80 | tcp established: Disabled | log: Disabled
HR_to_Prod_SRVs
–
ipv4
allow tcp | src: any, dst: 8080 | tcp established: Disabled | log: Disabled
allow tcp | src: any, dst: 8081 | tcp established: Disabled | log: Disabled
IT_to_Prod_SRVs
–
ipv4
allow tcp | src: any, dst: 22 | tcp established: Disabled | log: Disabled
allow tcp | src: any, dst: 1521 | tcp established: Disabled | log: Disabled
allow tcp | src: any, dst: 80 | tcp established: Disabled | log: Disabled
allow tcp | src: any, dst: 8080 | tcp established: Disabled | log: Disabled...
To_Test_Servers
–
ipv4
deny icmp | src: any, dst: any | tcp established: Disabled | log: Enabled
allow tcp | src: any, dst: 22 | tcp established: Disabled | log: Enabled
Rows per page
30
1-5 of 5
1

Figure 12-30 *Custom ACLs Management*

To create a new custom ACL, click **+ Add Custom ACL**. This action will open the ACL configuration page, as shown in Figure 12-31. Now, complete this initial steps of the Adaptive Policy ACL configuration:

Step 1. Specify the ACL name.

Step 2. Define the ACL IP version.

Step 3. After these parameters are defined, you can begin adding individual ACL rules.

Step 4. After adding all required individual rules, you need to click **Create** to save the new custom ACL.

Technet24

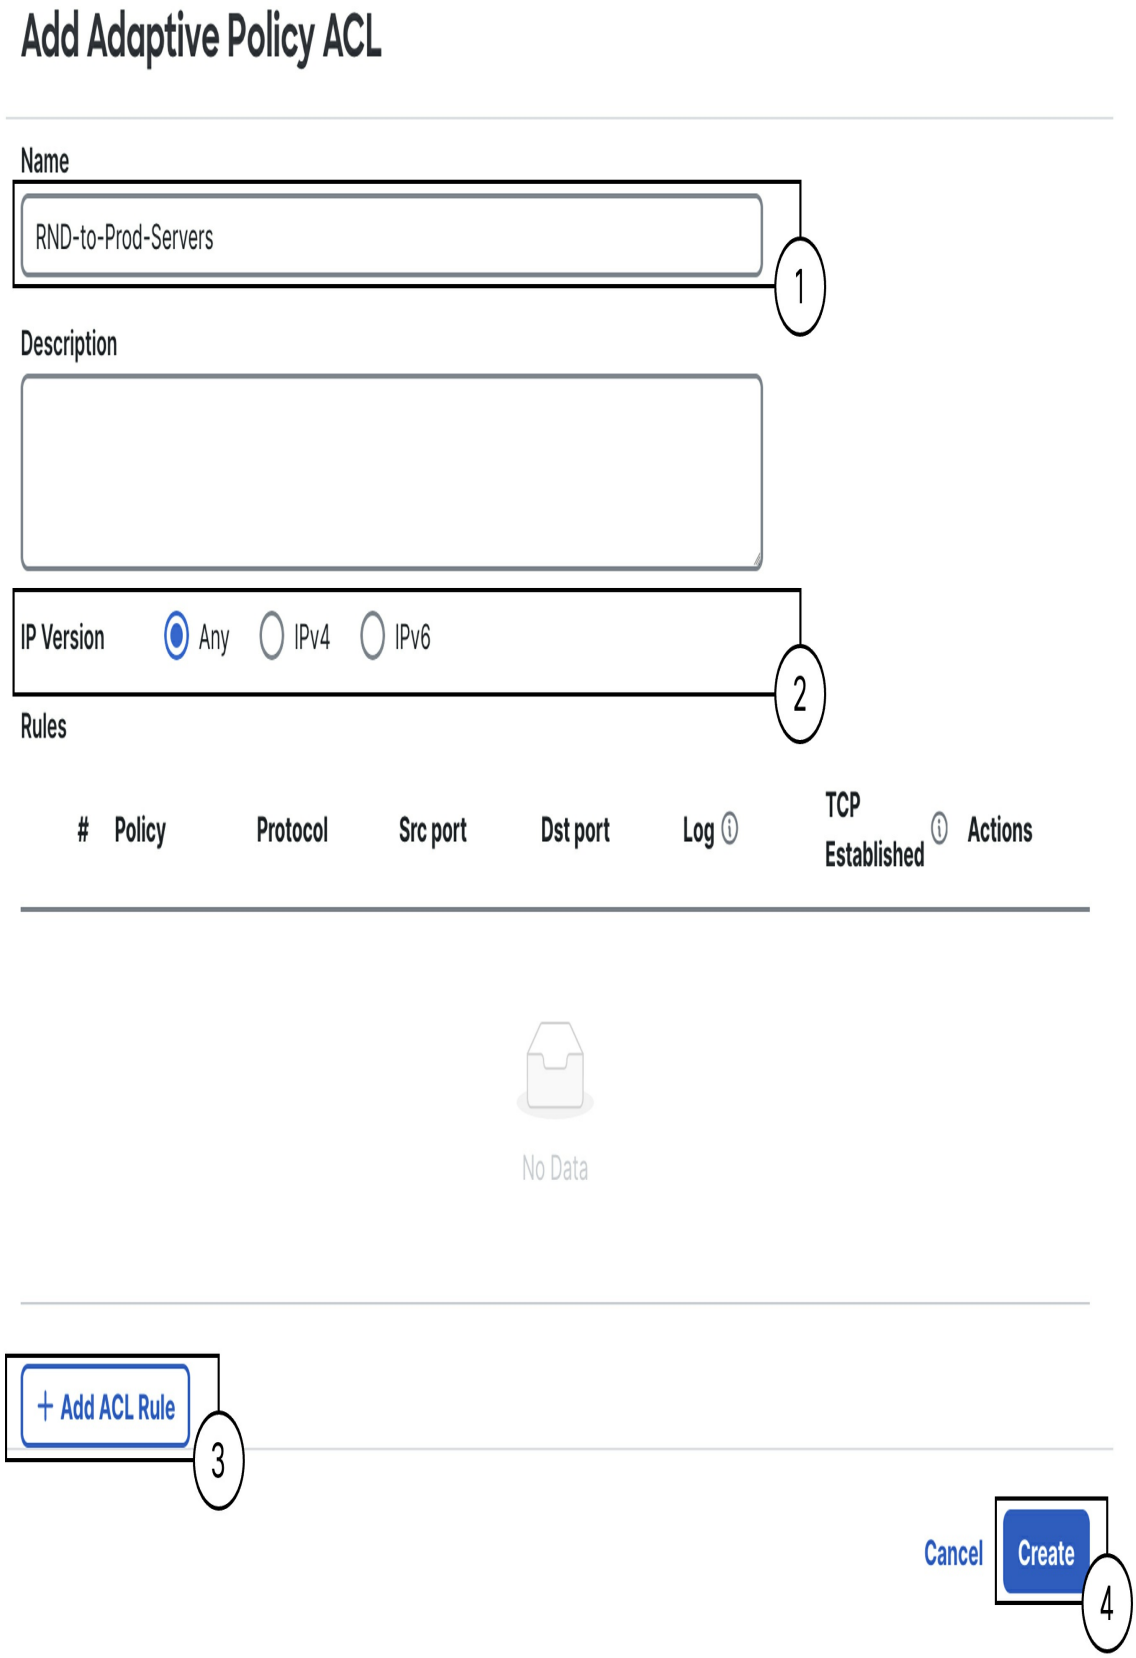

Add Adaptive Policy ACL
Name
RND-to-Prod-Servers
1
Description
IP Version
Any
IPv4
IPv6
2
Rules
Policy Protocol Src port Dst port Log TCP Established Actions
No Data
+ Add ACL Rule
3
Cancel
Create
4

Figure 12-31 *Custom ACLs Settings*

Configuration of an individual rule involves multiple steps, which are displayed in Figure 12-32.

Technet24

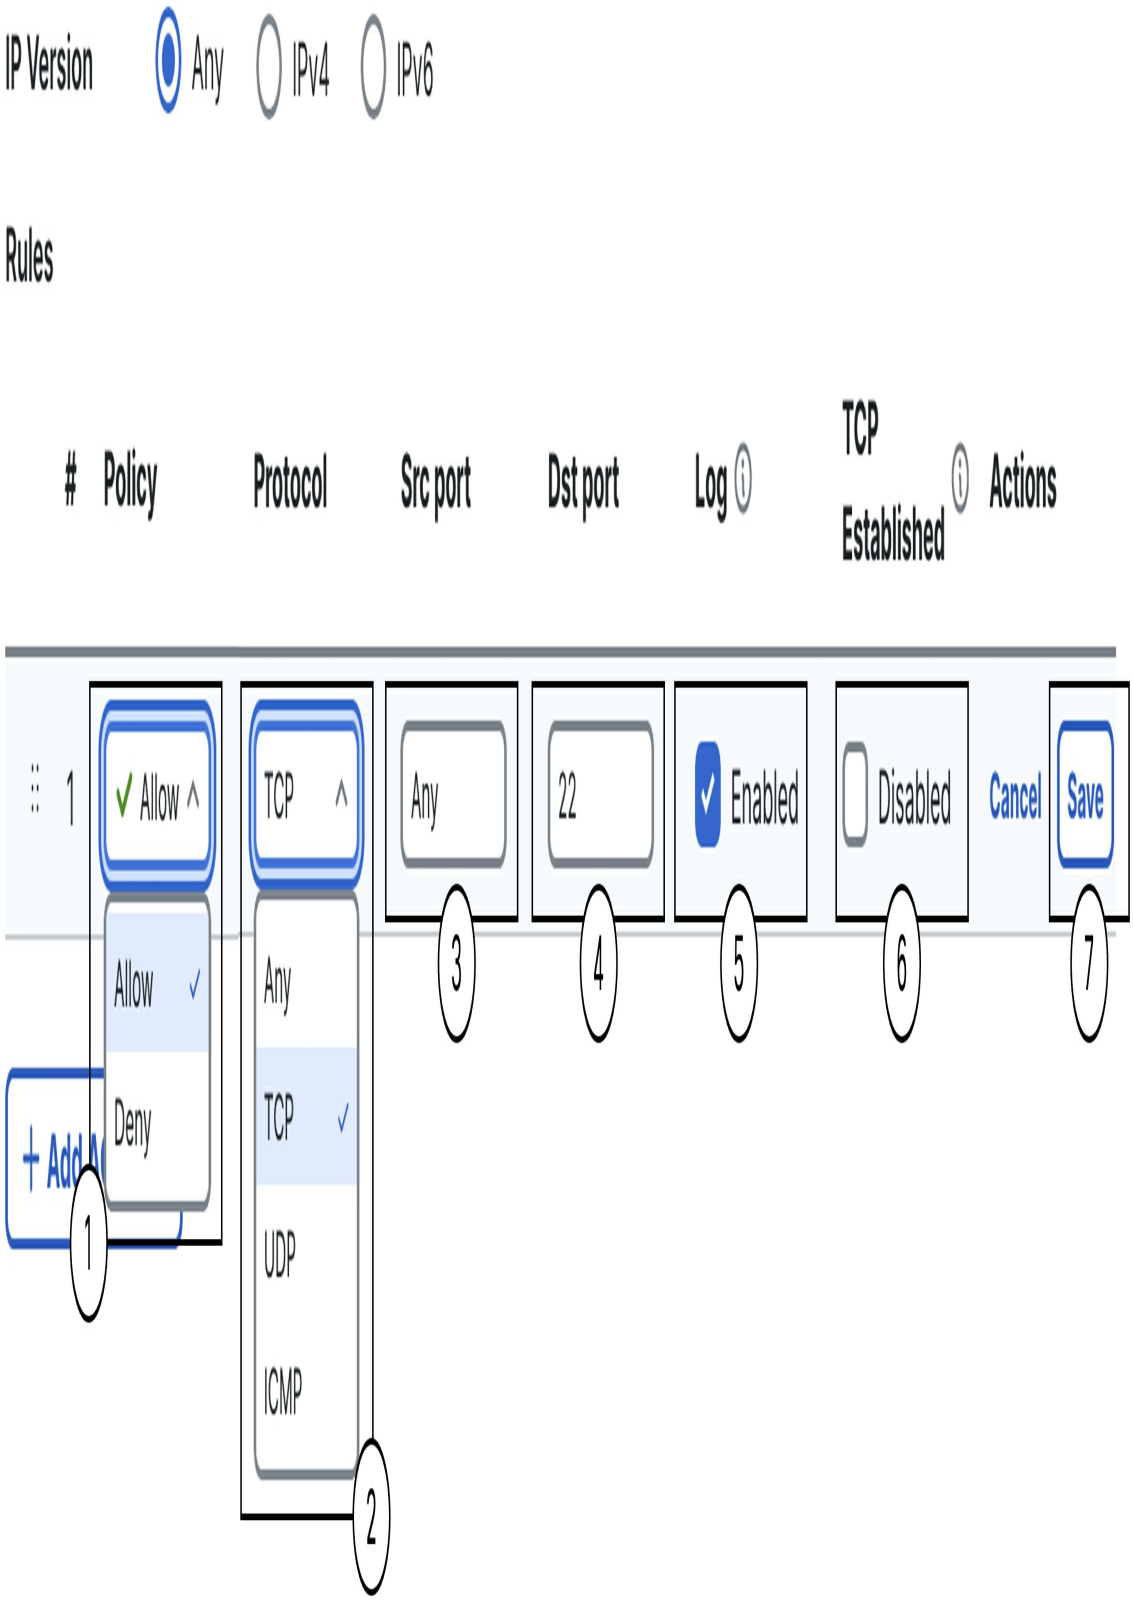

IP Version
Any
IPv4
IPv6
Rules
Policy
Protocol
Src port
Dst port
Log
TCP Established
Actions
1
Allow
TCP
Any
22
Enabled
Disabled
Cancel
Save
Allow
Deny
Any
TCP
UDP
ICMP
1
2
3
4
5
6
7

Figure 12-32 *Custom ACL Rule Settings*

The following steps need to be taken:

Step 1. Choose the policy type: **Allow** or **Deny**.

Step 2. Select the protocol from the drop-down list.

Step 3. Select the source port for **TCP** or **UDP**-based applications or keep it as **Any**.

Step 4. Select the destination port for TCP or UDP-based applications or keep it as **Any**.

Step 5. Enable logging for the rule.

Step 6. Turn on TCP established verification. This setting will only allow traffic for already-established connections.

Step 7. Click **Save** to save the individual rule.

When you're planning or configuring individual rules, it is important to consider the following:

- Although logging can be enabled in any custom ACL, at the time of writing, it is supported only by MS390 and C9X00-M platforms. For all other platforms, the dashboard will not propagate the logging setting in a custom ACL rule.
- Although a TCP-established check can be enabled in any custom ACL, at the time of writing, it is supported only by MS390 and C9X00-M platforms. For all other platforms, the dashboard will not propagate the TCP established setting in a custom ACL rule.
- Custom ACLs do not include a default implicit deny rule at the end, unlike regular ACLs. When you use Adaptive Policies, you need to either configure an explicit deny at the end of each custom ACL or use the Deny default action in the policy itself. The

Technet24

second approach will be described later in the chapter.

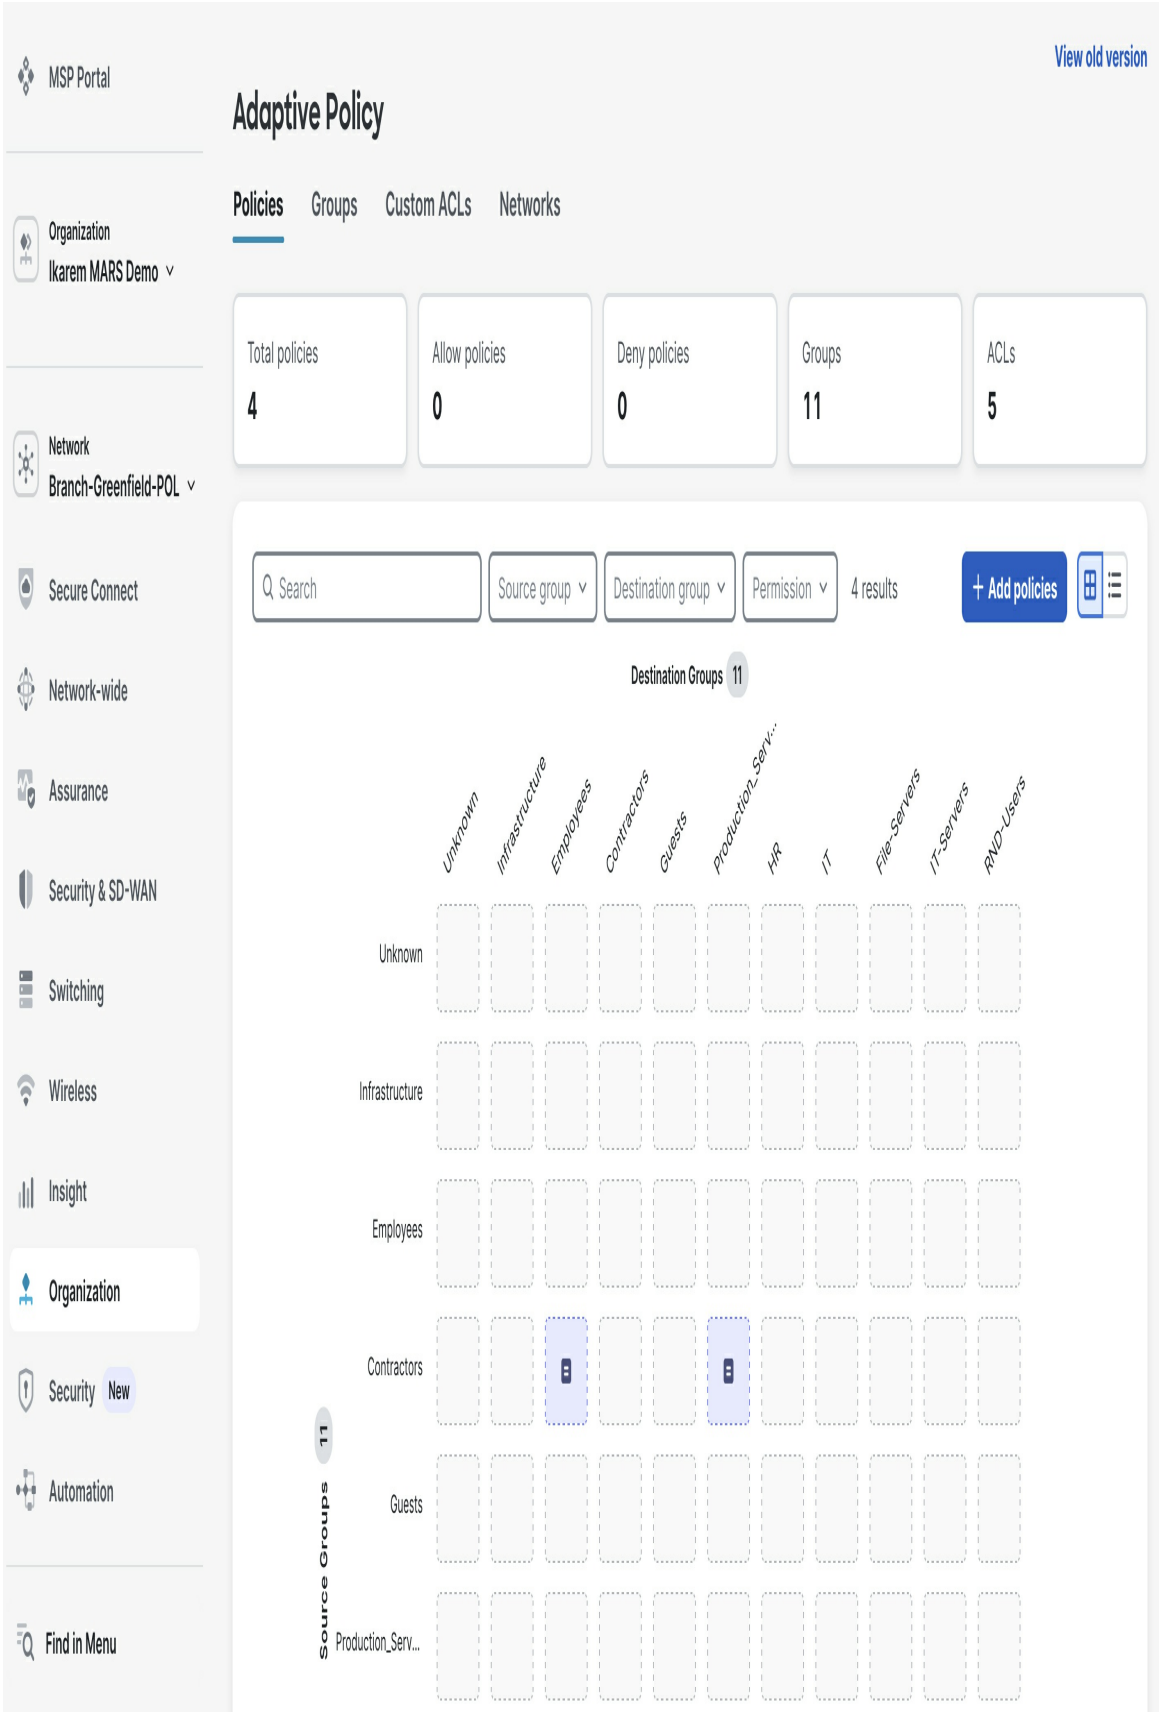
Note

When you're considering the implementation of TCP-established verification, it is important to remember that this setting needs to be enabled for return traffic. If you create a rule describing access from the client to the server, do not configure the rule to verify the TCP-established state because the client must initiate the connection to the server first and the access rule must allow the initial SYN request.

Defining Adaptive Policies

Adaptive policies are managed on the **Policies** tab located in **Organization > Adaptive Policy**. Here, you have a centralized view of your Adaptive Policy matrix with the ability to define new policies as well as edit or delete existing ones.

To create a new policy, click **+ Add Policies**, as demonstrated in Figure 12-33.


MSP Portal
Organization
Ikarem MARS Demo
Network
Branch-Greenfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security New
Automation
Find in Menu
View old version
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Total policies
4
Allow policies
0
Deny policies
0
Groups
11
ACLs
5
Search
Source group
Destination group
Permission
4 results
+ Add policies
Destination Groups 11
Unknown
Infrastructure
Employees
Contractors
Guests
Production_Serv...
HR
IT
File-Servers
IT-Servers
RND-Users
Source Groups 11
Unknown
Infrastructure
Employees
Contractors
Guests
Production_Serv...

Technet24

Figure 12-33 *Adaptive Policy Policies Tab*

After clicking the **+ Add Policies** button, you need to configure the actual policy in the new window. Figure 12-34 shows an Adaptive Policy configuration.


View old version
MSP Portal
Adaptive Policy
Add Policies
Organization
Ikarem MARS Demo
Network
Branch-Greenfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security
New
Automation
Find in Menu
Source groups
Search
Name
SGT Value
Description
Unknown (0) 0 Created by Meraki, the Unknown group applies when a policy is specified f...
Infrastructure (0) 2 Created by Meraki, the Infrastructure group is used by Meraki devices for...
Employees (0) 4 –
Contractors (2) 5 –
Guests (0) 6 –
Production_Servers (0) 11 Production Servers Security Group
HR (1) 16 –
IT (1) 17 –
File-Servers (0) 21 –
IT-Servers (0) 31 –
RND-Users (0) 101 –
1
Destination groups
Search
3
1 item selected
Select all 11 items
Allow
Deny
Custom
Delete
Name
Permission
SGT Value
Description
Unknown (0) 0 Created by Meraki, the Unknown group applies when a...
Infrastructure (0) 2 Created by Meraki, the Infrastructure group is used ...
Employees (1) 4 –
Contractors (0) 5 –
Guests (0) 6 –
Production_Servers (3) 11 Production Servers Security Group
2
HR (0) 16 –
IT (0) 17 –

Technet24

Figure 12-34 *Adaptive Policy Configuration Options*

The following steps are needed to configure an Adaptive Policy:

Step 1. Select a source group.

Step 2. Select a destination group.

Step 3. Choose the action for the traffic between source and destination groups. The options are as follows:

- **Deny:** This option denies any traffic from the source group to the destination group.
- **Allow:** This option allows all traffic from the source group to the destination group.
- **Custom:** Select this option to assign your custom ACL to traffic from the source group to the destination group.

If you select **Allow** or **Deny**, a new pop-up window will appear asking you to confirm the selection. In this pop-up window, you can select the **Would You Like to Automatically Create the Inverse Policy** checkbox, which instructs the dashboard to apply the same policy in the opposite direction. This option is useful for straightforward Allow and Deny policies, because it saves time, but it should not be used with custom ACLs. If you select **Would You Like to Automatically Create the Inverse Policy** for custom ACLs, the dashboard will apply the same ACLs in the opposite direction of traffic, which may not meet your requirements. Figure 12-35 shows the confirmation pop-up window.

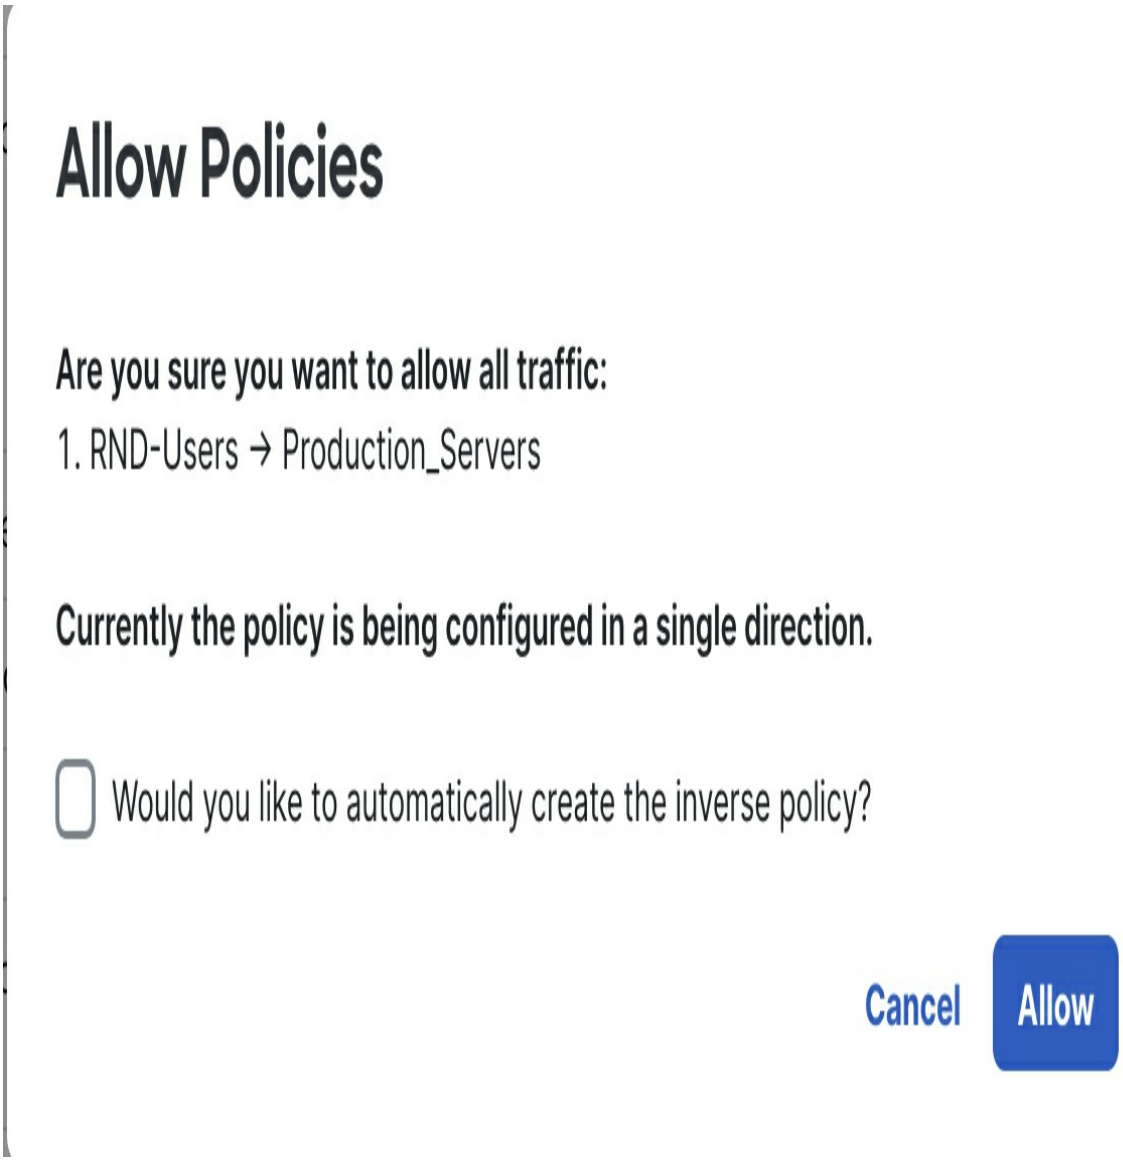


Figure 12-35 *Adaptive Policy Confirmation Pop-Up Window*

If you select **Custom** as the policy action, the pop-up window will appear differently. In this window, you must assign existing custom ACLs to the policy and select the default action. Figure 12-36 shows the required configuration steps.

Technet24

Edit Policy

Source group	Destination group	Custom ACLs	Last-entry rule
RND-Users	Production_Servers		Allow

Configure ACLs

Configure custom ACLs on selected policies. Note if multiple policies are selected, change made here will overwrite existing ACLs

#	ACL name	Description	Rules	Actions
1	RND_users_to_Prod_SRVs			Cancel Save

+ Add an entry

1

2

3

Last-entry rule will apply as the last ACL entry.

Allow

4

Description

Cancel Review change

5

Figure 12-36 *Custom Adaptive Policy Configuration*

The following steps outline how to configure a custom Adaptive Policy:

Step 1. Click **+ Add an Entry** to add a custom ACL. You can define multiple ACLs within the same Adaptive Policy.

Step 2. Select the custom ACL from the drop-down list.

Step 3. Save the policy line.

Step 4. Configure the default action. If not configured, the default action is set to **Allow**. If this is the case, explicit deny rules are required in the custom ACLs to enforce least-privileged access.

Step 5. Click **Review Change** to save the configuration.

Figure 12-37 demonstrates how to edit existing policies in the matrix; the following steps outline how to do this:

Step 1. Click the cell that requires a change. This action opens an edit menu on the right side of the screen, where you can modify your custom policy.

Step 2. Click the pencil icon to modify your custom policy.

Technet24

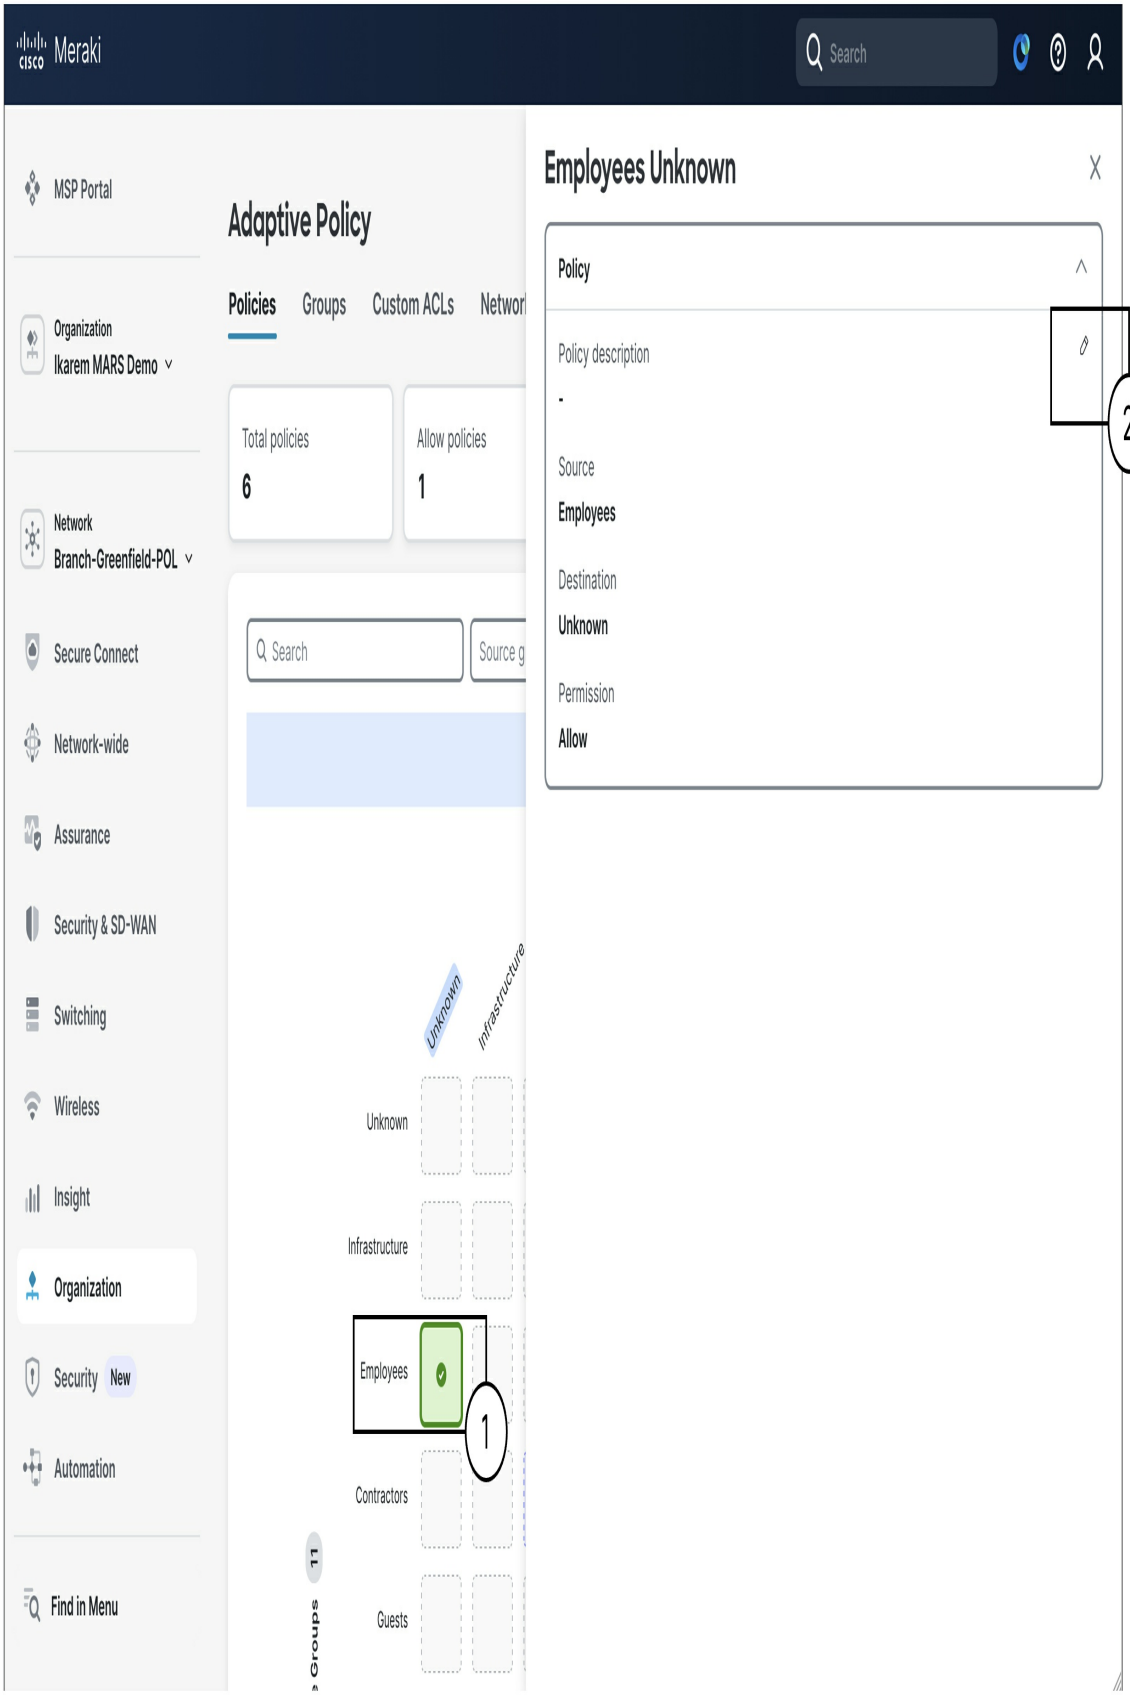

Meraki
Search
MSP Portal
Organization
Ikarem MARS Demo
Network
Branch-Greenfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security
New
Automation
Find in Menu
Adaptive Policy
Policies
Groups
Custom ACLs
Total policies
6
Allow policies
1
Search
Unknown
Infrastructure
Employees
Contractors
Guests
1
Employees Unknown
Policy
Policy description
-
Source
Employees
Destination
Unknown
Permission
Allow
2

Figure 12-37 *Editing Custom Adaptive Policies*

After closing the edit menu on the right side of the screen, you gain access to Adaptive Policy management elements such as Edit and Delete, as demonstrated in Figure 12-38.

Technet24

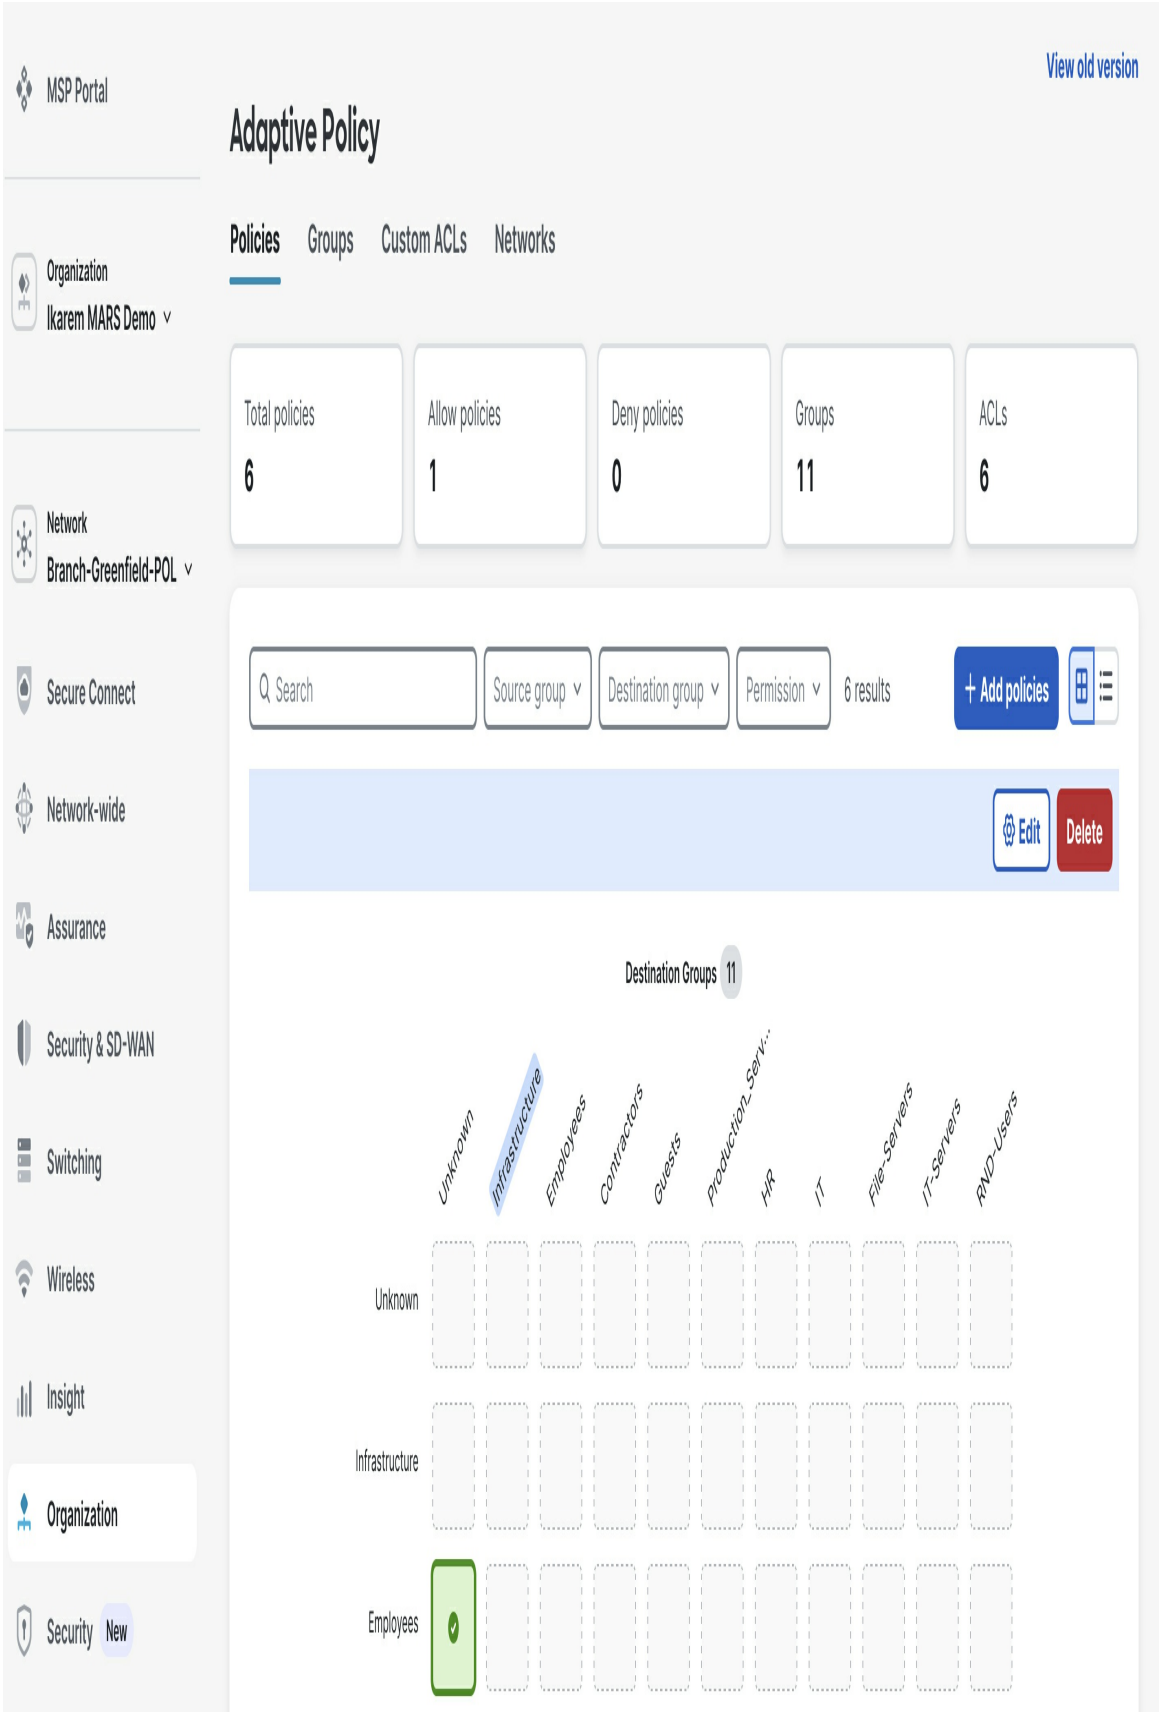

MSP Portal
Organization
Ikarem MARS Demo
Network
Branch-Greenfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security
New
View old version
Adaptive Policy
Policies
Groups
Custom ACLs
Networks
Total policies
6
Allow policies
1
Deny policies
0
Groups
11
ACLs
6
Search
Source group
Destination group
Permission
6 results
+ Add policies
Edit
Delete
Destination Groups 11
Unknown
Infrastructure
Employees
Contractors
Guests
production_Serv...
HR
IT
File-Servers
IT-Servers
RND-Users
Unknown
Infrastructure
Employees

Figure 12-38 *Adaptive Policies Management*

Implementing Adaptive Policy at the Device Level

Adaptive Policy implementation at the device level varies based on device type. The subsequent sections explain how to implement Adaptive Policy for specific devices.

Implementing Adaptive Policy for Wired Network Access Control

Access Policy Features for Wired Connections

MS and CS switches support the following list of Adaptive Policy features:

- Port-based Adaptive Policy group assignment
- VLAN-based Adaptive Policy group assignment
- RADIUS-based Adaptive Policy group assignment
- Static IP to SGT and Subnet:SGT mappings
- Adaptive Policy inline tagging
- Adaptive Policy enforcement

Now, we will review each feature separately to help you gain a better understanding of how they can be used to deliver microsegmentation.

Port-Based Adaptive Policy Group Assignment

Port-based Adaptive Policy group assignment is the simplest mechanism for assigning an Adaptive Policy group. To assign the

Technet24

group, you must define it in **Organization > Adaptive Policy > Group**. Once you statically assign the group to the port, all traffic from any device connected to this port is considered as belonging to this specific group. Each IP address behind the port is added to the IP to SGT table next to the group selected for the port. All frames from any MAC address located behind the interface are tagged with the group selected on the port when they leave the switch through an interface enabled for systems using adaptive policy. Such an interface is typically connected to another Adaptive Policy-compatible switch, access point, or security appliance.

Figure 12-39 shows the Adaptive Policy Groups drop-down list on the switch interface.

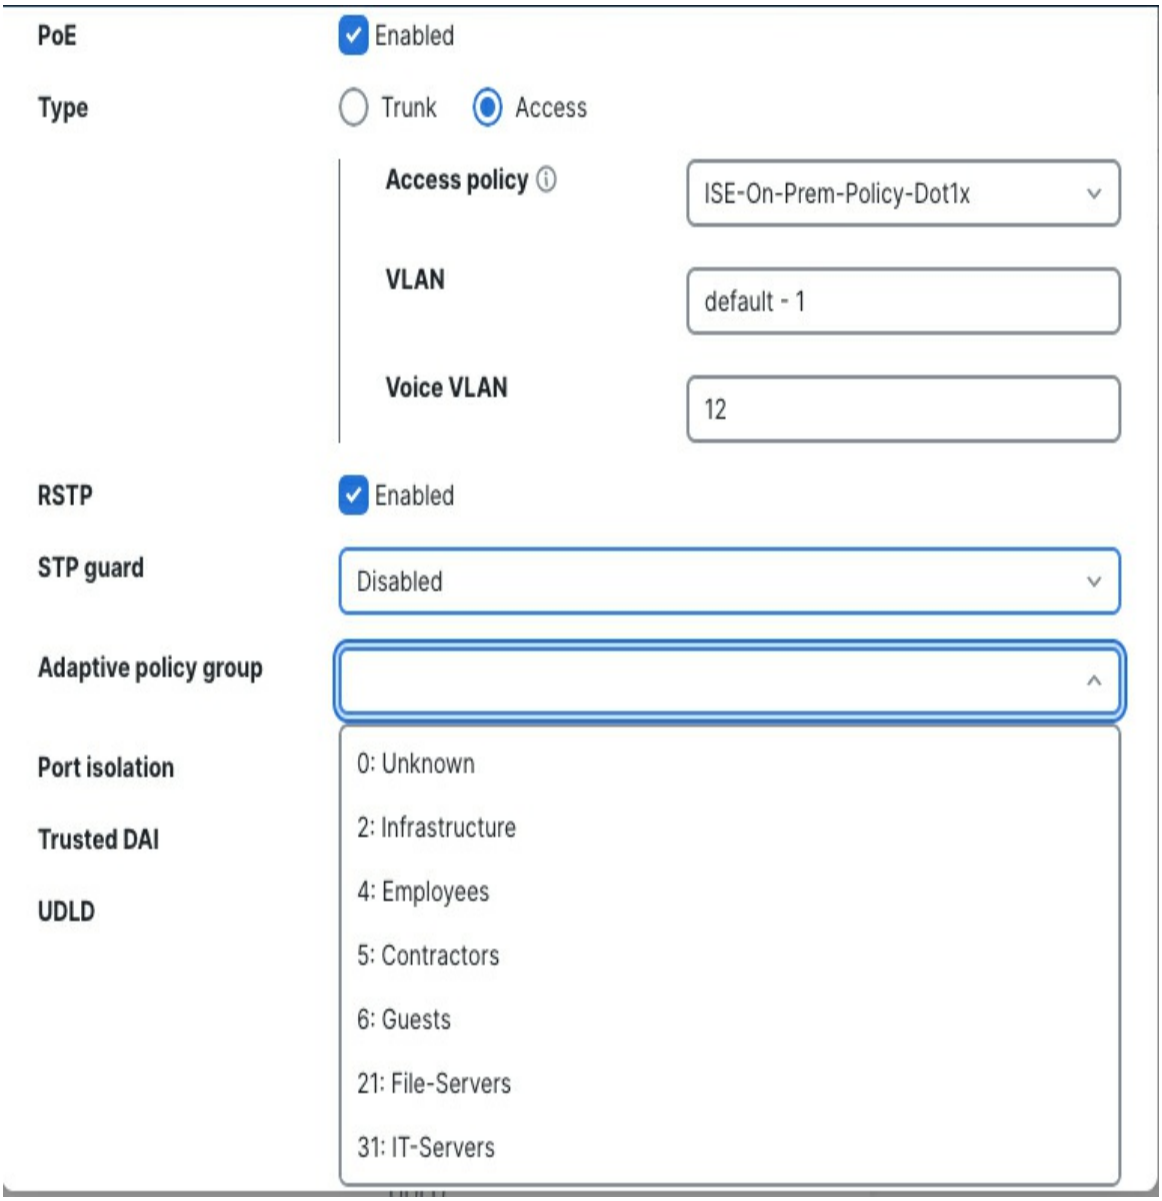


Figure 12-39 *Static Adaptive Policy Group Assignment on the Interface Level*

VLAN-Based Adaptive Policy Group Assignment

VLAN-based Adaptive Policy group assignment was mentioned in the previous chapter. As the administrator, you can map an existing security group to a VLAN listed in the VLAN profile of the given Meraki network. VLAN profiles are located at **Network-wide > VLAN Profiles**. Figure 12-40 shows default VLAN profiles with two VLANs

Technet24

assigned to Adaptive Policy groups.

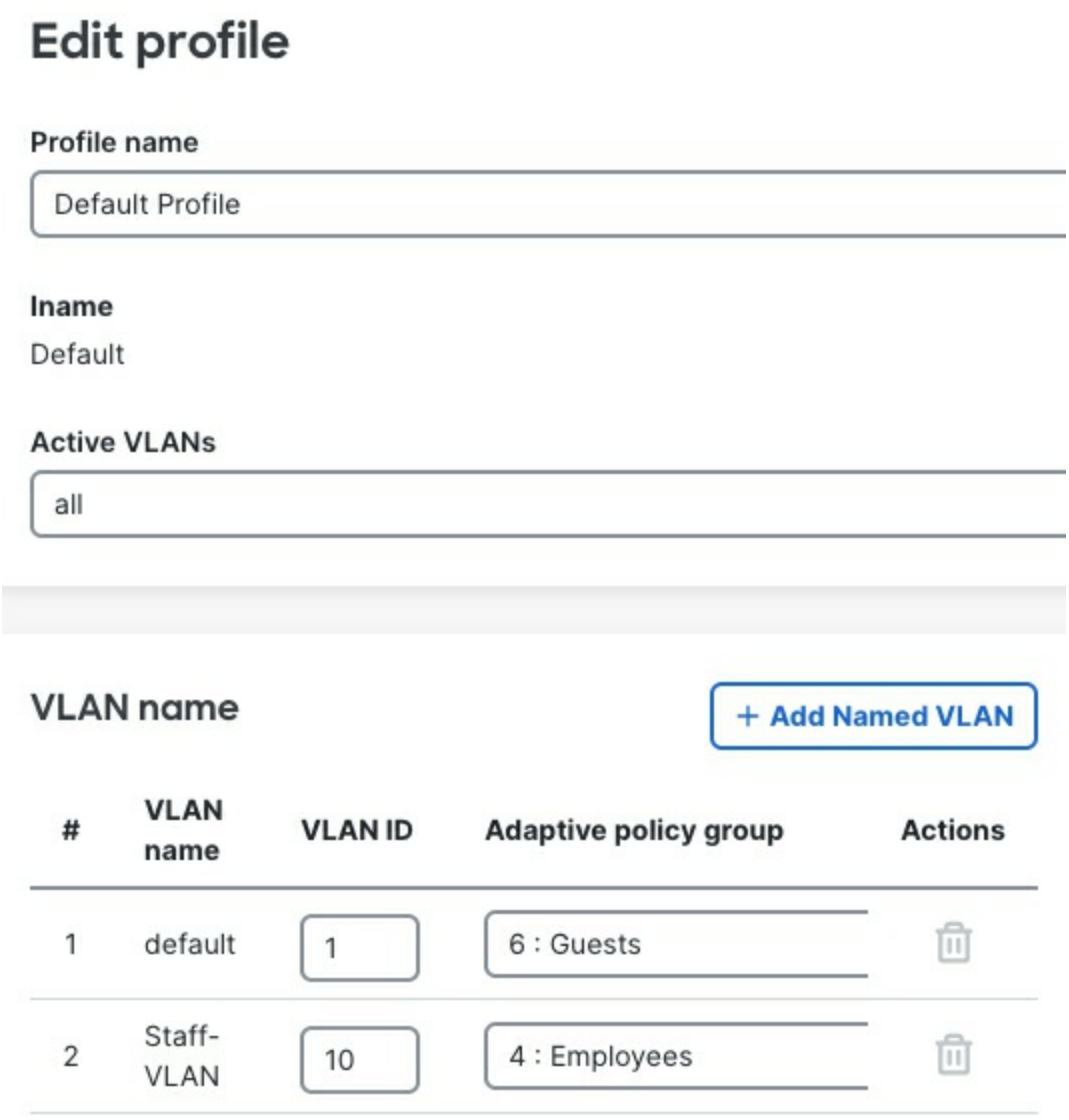


Figure 12-40 *Static Adaptive Policy Group Assignment on the VLAN Level*

RADIUS-Based Adaptive Policy Group Assignment

In the case of RADIUS Adaptive Policy group assignment, you do not need to perform specific configuration on the switch. Most settings are handled on ISE, where SGTs and corresponding authorization policies should be defined. On the Meraki switch, the port must be

configured with an access policy to enable dot1x or MAB authentication for the endpoint, and at the organization level, the network where the switch is located must be enabled for Adaptive Policy.

Static IP to SGT and Subnet:SGT Mappings

IP to SGT and Subnet:SGT binding allows you to create and propagate mappings for devices that cannot be dynamically assigned SGTs, such as servers and other resources that cannot undergo MAB or 802.1X authentication. The process of assigning and propagating bindings consists of two steps:

Step 1. Create the policy object of the category Adaptive Policy. You perform this step on the **Organization > Policy Objects > All Objects** page. Figure 12-41 shows such an object.

Technet24

Edit policy object

Name

IT-Server

Category

Network | **Adaptive Policy**

FQDN, IP or CIDR

10.110.20.11/32

You can only enter one value per object. Attach multiple objects to an adaptive policy group for multiple entries.

Edit object

Figure 12-41 *The Policy Object of the Category Adaptive Policy*

Step 2. Assign the previously created policy object to an Adaptive Policy group by navigating to **Organization > Adaptive Policy > Groups**. Figure 12-42 shows Adaptive Policy groups with assigned objects.

Technet24

Adaptive Policy

Policies | **Groups** | Custom ACLs | Networks

Search | 7 results | + Add group

Name	SGT Value	Description	Policy Objects
Unknown	0	Created by Meraki, the Unknown group applies when a policy is specified for unsuccessful group classification	–
Infrastructure	2	Created by Meraki, the Infrastructure group is used by Meraki devices for internal and dashboard communication	–
Employees	4		–
Contractors	5		–
Guests	6		–
File-Servers	21		File-server
IT-Servers	31		IT-Server

Figure 12-42 *Adaptive Policy Group with Assigned Policy Objects*

Adaptive Policy Inline Tagging

You can enable inline tagging on trunk switch ports by checking the **Enable** box next to **Peer SGT Capable**. Along with this setting, you must select an Adaptive Policy group to save the changes. Both ends of the connection should have inline tagging enabled and the same Adaptive Policy group (SGT) configured. Inline tagging is typically enabled on uplink ports connected to Adaptive Policy-capable distribution switches, on ports where access points are connected, and on ports linked to other Cisco equipment that supports inline tagging, such as firewalls and wireless LAN controllers. Figure 12-43 provides an example of port configuration for inline tagging.

Technet24

Update 1 Ports

Port status	☑ Enabled
Port profile	☐ Enabled
Link negotiation	Auto negotiate
EEE ⓘ	☐ Enabled
Port schedule	Unscheduled
PoE	☑ Enabled
Type	◉ Trunk ○ Access
Native VLAN	default - 1
Allowed VLANs	1-1000
RSTP	☑ Enabled
STP guard	Disabled
Peer SGT capable	☑ Enabled
Adaptive policy group	2: Infrastructure

Cancel Update

Figure 12-43 *Port Settings for the Inline Tagging*

Adaptive Policy Enforcement

No specific configuration is required at the device level for enforcement to occur. Enforcement depends on all the settings outlined earlier in the chapter, as well as the configuration implemented at the organization level.

The main prerequisites for enforcement to occur are as follows:

- The packet should either arrive tagged with an SGT, or if packet switching or routing occurs within the same switch, the IP to SGT mapping for the source IP must be available.
- The destination IP must have an IP to SGT mapping, which informs the switch of the destination SGT value.
- A policy must be assigned between the destination SGT and the source SGT.

Adaptive Policy Guidelines for Wired Network Segmentation

Adaptive Policy Group Assignment: Order of Operations

When implementing adaptive policies, you must consider the priority with which the adaptive policy group value is assigned to a frame when IP to SGT mapping is available from multiple sources, such as dynamic mapping via RADIUS and static mapping via policy objects. The following order defines the precedence of different adaptive policy group sources:

- Port:SGT or RADIUS session SGT assignment (802.1X/MAB). This has the highest priority for group assignment. The group mapping from the RADIUS server always takes precedence over

Technet24

static mapping on the port.

- IP/Subnet:SGT map. Group mapping using policy objects is considered the second priority.
- VLAN:SGT map. An adaptive policy group assigned to the VLAN is the third priority.
- Last resort assignment of 0 (unknown).

To clarify the priority of operations, consider the following example: The switch port is configured with SGT 20, the RADIUS server assigns SGT value 10, there is an IP to SGT mapping for the endpoint with SGT 5, and the VLAN is assigned to SGT 100. In this case, the SGT value from the RADIUS server (SGT 10) will be used. The priority order for SGTs is SGT 10, SGT 20, SGT 5, and SGT 100.

Adaptive Policy Enforcement Point

When an endpoint is assigned to an Adaptive Policy group, there are two possibilities for enforcement:

- Both source and destination are connected to the same switch. In this case, enforcement occurs within the switch itself.
- Source and destination are connected to different network devices. In this scenario, the enforcement point is typically the device closest to the destination, provided inline tagging is enabled along the path.

The second scenario deserves closer attention because the actual enforcement point may vary depending on the network setup and where inline tagging is enabled. Understanding the enforcement point location for specific types of traffic is critical for design, operation, and troubleshooting. Consider the following two subscenarios for this option.

- **Scenario 1: Inline tagging is present along the entire path between source and destination.**

Figure 12-44 outlines the first scenario. In this case, while Access-SW-1 has all the information needed to perform enforcement, the presence of inline tagging enabled on the port where the destination is attached means Access-SW-1 does not enforce locally. Instead, it tags the frames with SGTs and forwards them. The actual enforcement occurs on Access-SW-2.

Technet24

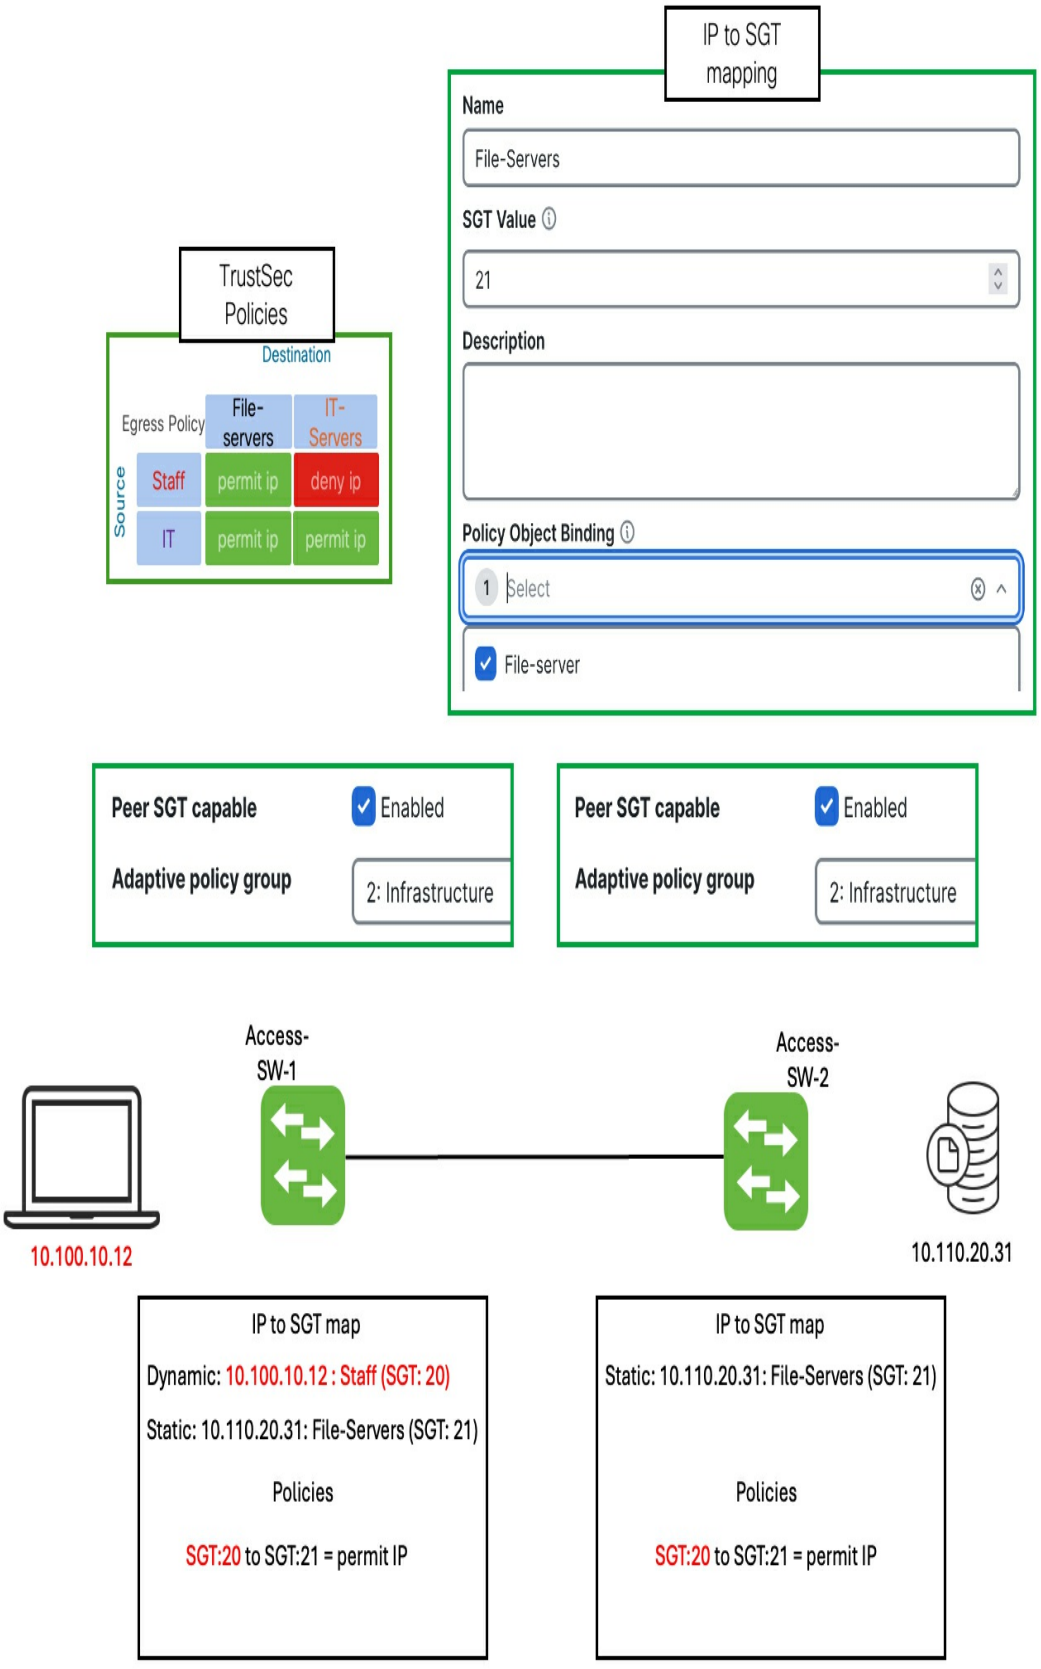

IP to SGT mapping
Name
File-Servers
SGT Value
21
Description
Policy Object Binding
1 Select
File-server
TrustSec Policies
Destination
Egress Policy
File-servers
IT-Servers
Source
Staff
permit ip
deny ip
IT
permit ip
permit ip
Peer SGT capable
Enabled
Adaptive policy group
2: Infrastructure
Peer SGT capable
Enabled
Adaptive policy group
2: Infrastructure
Access-SW-1
Access-SW-2
10.100.10.12
10.110.20.31
IP to SGT map
Dynamic: 10.100.10.12 : Staff (SGT: 20)
Static: 10.110.20.31: File-Servers (SGT: 21)
Policies
SGT:20 to SGT:21 = permit IP
IP to SGT map
Static: 10.110.20.31: File-Servers (SGT: 21)
Policies
SGT:20 to SGT:21 = permit IP

Figure 12-44 *Enforcement with Inline Tagging Present on the Path*

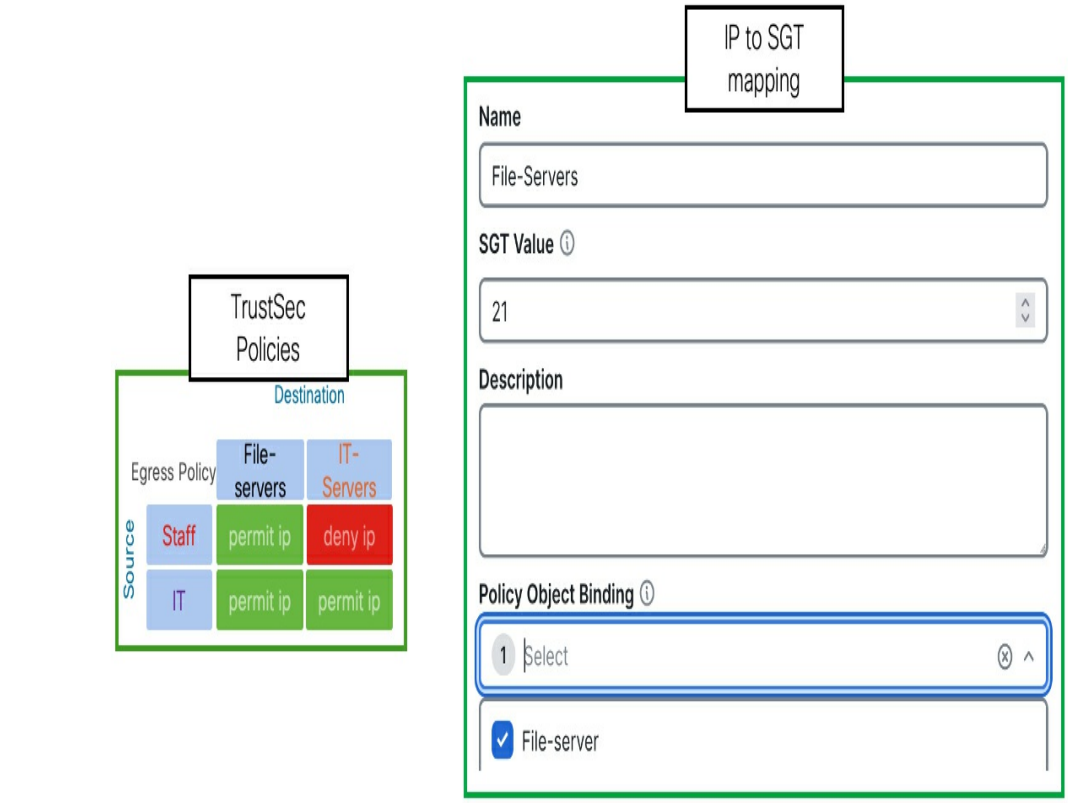
- **Scenario 2: Inline tagging is not present either along the entire path or at a specific point within the path.**

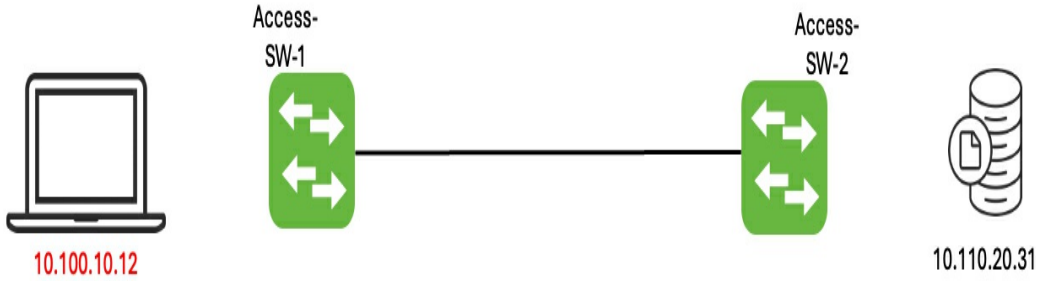
Figure 12-45 outlines the second scenario. Since Access-SW-1 does not have inline tagging enabled on the path to the destination but has all the information required for enforcement, enforcement will occur locally. Another possible case in this scenario arises when the switch where the source is connected does not have information about the destination IP to SGT mapping. In this situation, frames are propagated without tagging. This subscenario typically represents a connection across non-TrustSec–capable network infrastructure. On the destination side, when the destination is connected to a Cisco switch, you can use an SXP connection to ISE to receive source IP to SGT mappings for authenticated users, since Meraki switches do not support SXP directly.

Note

To learn mappings for authenticated users, ISE must receive accounting updates from the devices where the user is authenticated, which include the Framed-IP attribute.

Technet24





IP to SGT map

Dynamic: 10.100.10.12 : Staff (SGT: 20)

Static: 10.110.20.31: File-Servers (SGT: 21)

Policies

SGT:20 to SGT:21 = permit IP

IP to SGT map

Static: 10.110.20.31: File-Servers (SGT: 21)

Policies

SGT:20 to SGT:21 = permit IP

Figure 12-45 *No Inline Tagging Present on the Path*

Implementing Adaptive Policy for Wireless Network Access Control

Implementing Adaptive Policies for wireless devices requires an understanding of the supported features and adherence to best practices to ensure reliable deployment. The following sections focus on these aspects of Adaptive Policies for access points.

Access Policy Features for Wireless Connections

Both MR and CW access points support the following Adaptive Policy features:

- SGT support detection on Uplink
- SSID-based Adaptive Policy group assignment
- Group policy–based Adaptive Policy group assignment
- Static IP to SGT and Subnet:SGT mappings
- RADIUS-based Adaptive Policy group assignment
- Adaptive Policy inline tagging
- Adaptive Policy enforcement

Now, we will review each feature separately to help you gain a better understanding of how they can be used to deliver microsegmentation.

SGT Support Detection on Uplink

SGT support detection on Uplink allows the access point to confirm whether the port on the upstream switch, where the AP is connected, is configured for inline tagging. An inline tagging configuration mismatch may result in traffic being dropped, and for cloud-managed

Technet24

equipment, this may cause an inability to establish a connection with the dashboard.

After the AP boots, it begins checking the frames received on its uplink for 30 seconds. If the AP detects SGT-tagged frames, it assumes that inline tagging is enabled upstream and starts propagating tagged packets. Conversely, if the AP does not detect any tagged frames, it assumes that inline tagging is disabled upstream and will not propagate any tagged frames, regardless of available IP to SGT mappings or their source (dynamic or static).

To verify the AP's decision for uplink SGT support detection, navigate to **Network-wide > Event Log**. Use the following steps to filter the logs:

Step 1. Ensure that **For Access Points** is selected as the event log source.

Step 2. Filter in **Event Type Include** to add the Adaptive Policy state change.

Step 3. Additionally, you can filter logs for a specific AP by specifying the AP MAC address or name in the **Access Point** field.

Step 4. All event logs can also be filtered for a specific time range.

Figure 12-46 shows an Adaptive Policy state change event during which the AP confirmed that the upstream switch is configured for inline tagging.

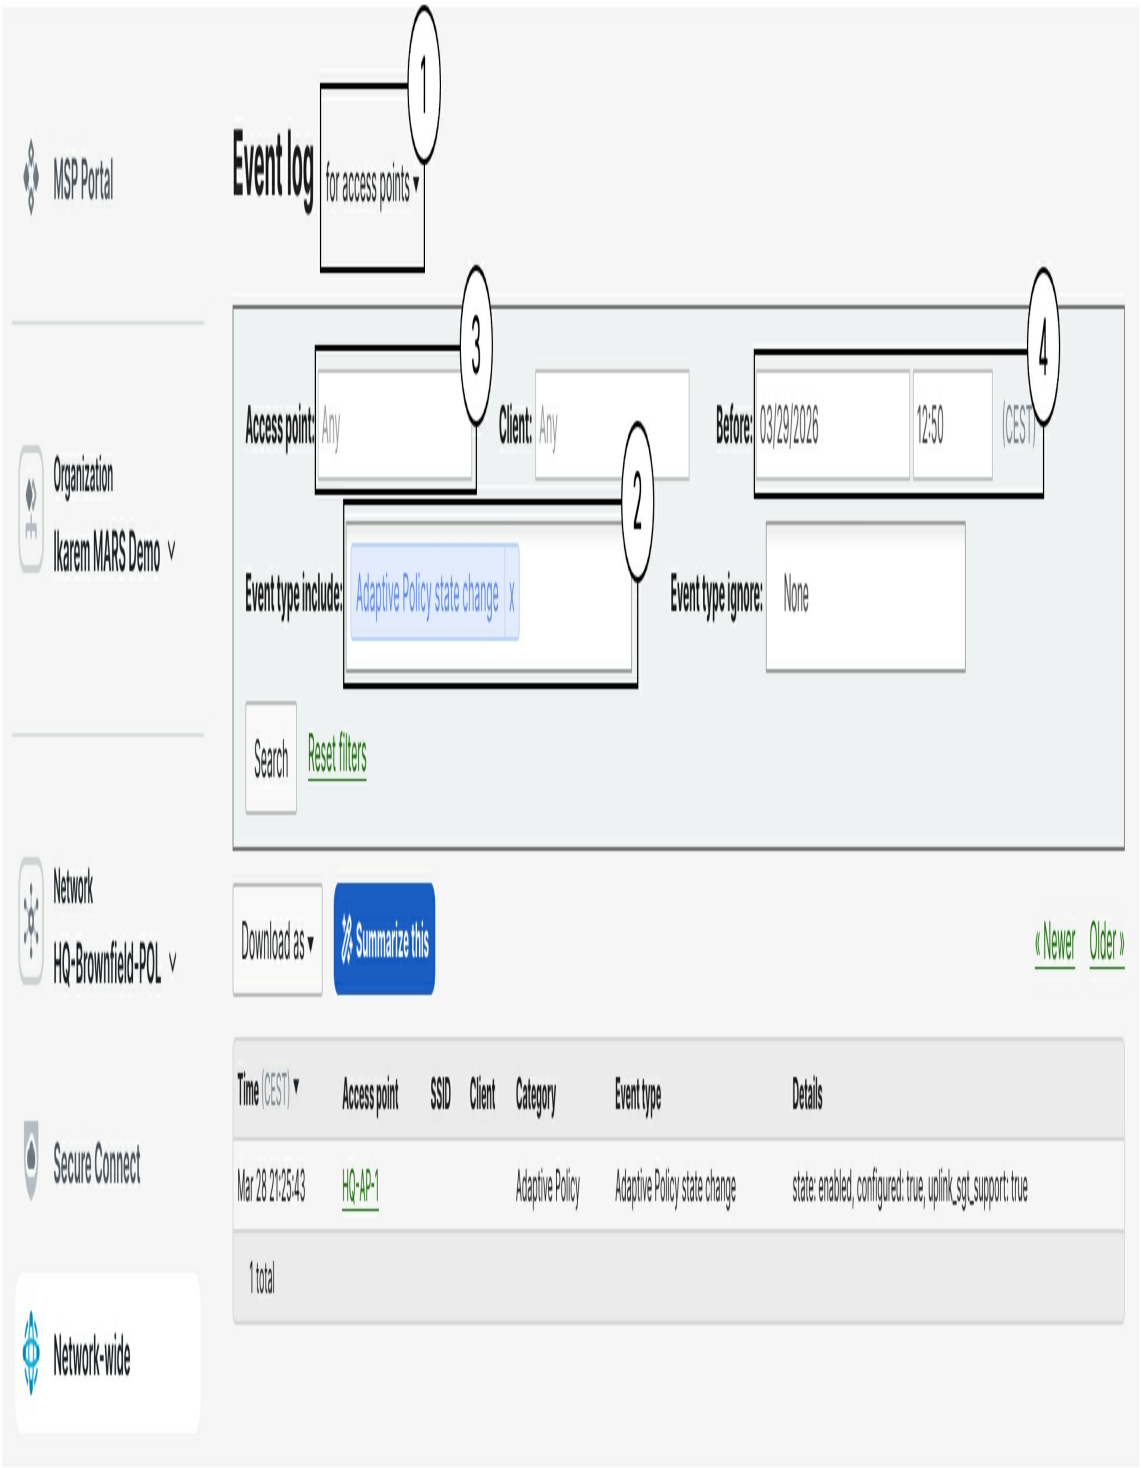


Figure 12-46 *Port Settings for the Inline Tagging*

From the perspective of the switch to which the AP is connected, the

Technet24

port is typically configured as a trunk with a specific native VLAN if required and enabled for Adaptive Policy inline tagging, as demonstrated in the previous section.

SSID-Based Adaptive Policy Group Assignment

Traffic from all clients connected to a specific SSID can be assigned to an Adaptive Policy group at the SSID level. This configuration instructs all APs with the given SSID to mark any frame received on that SSID with the SSID-level Adaptive Policy group. Common examples where this approach is used include guest SSIDs or SSIDs created for specific IoT devices.

To define an Adaptive Policy group at the SSID level, navigate to **Wireless > Access Control**. Then use the following steps to define an Adaptive Policy group at the SSID level:

Step 1. Select the SSID name where the Adaptive Policy group should be assigned from the drop-down list.

Step 2. Scroll to the bottom of the page and choose the group from the **Adaptive Policy Group** selection drop-down.

Figure 12-47 shows Adaptive Policy group configuration for an SSID.


Access control
SSID
1
MARS-Demo-HQ
Basic info
SSID (name)
MARS-Demo-HQ
SSID status
Enabled
Disabled
Hide SSID
Scroll down
Adaptive policy group
Bridge mode, NAT mode, MCG mode only
6: Guests
2

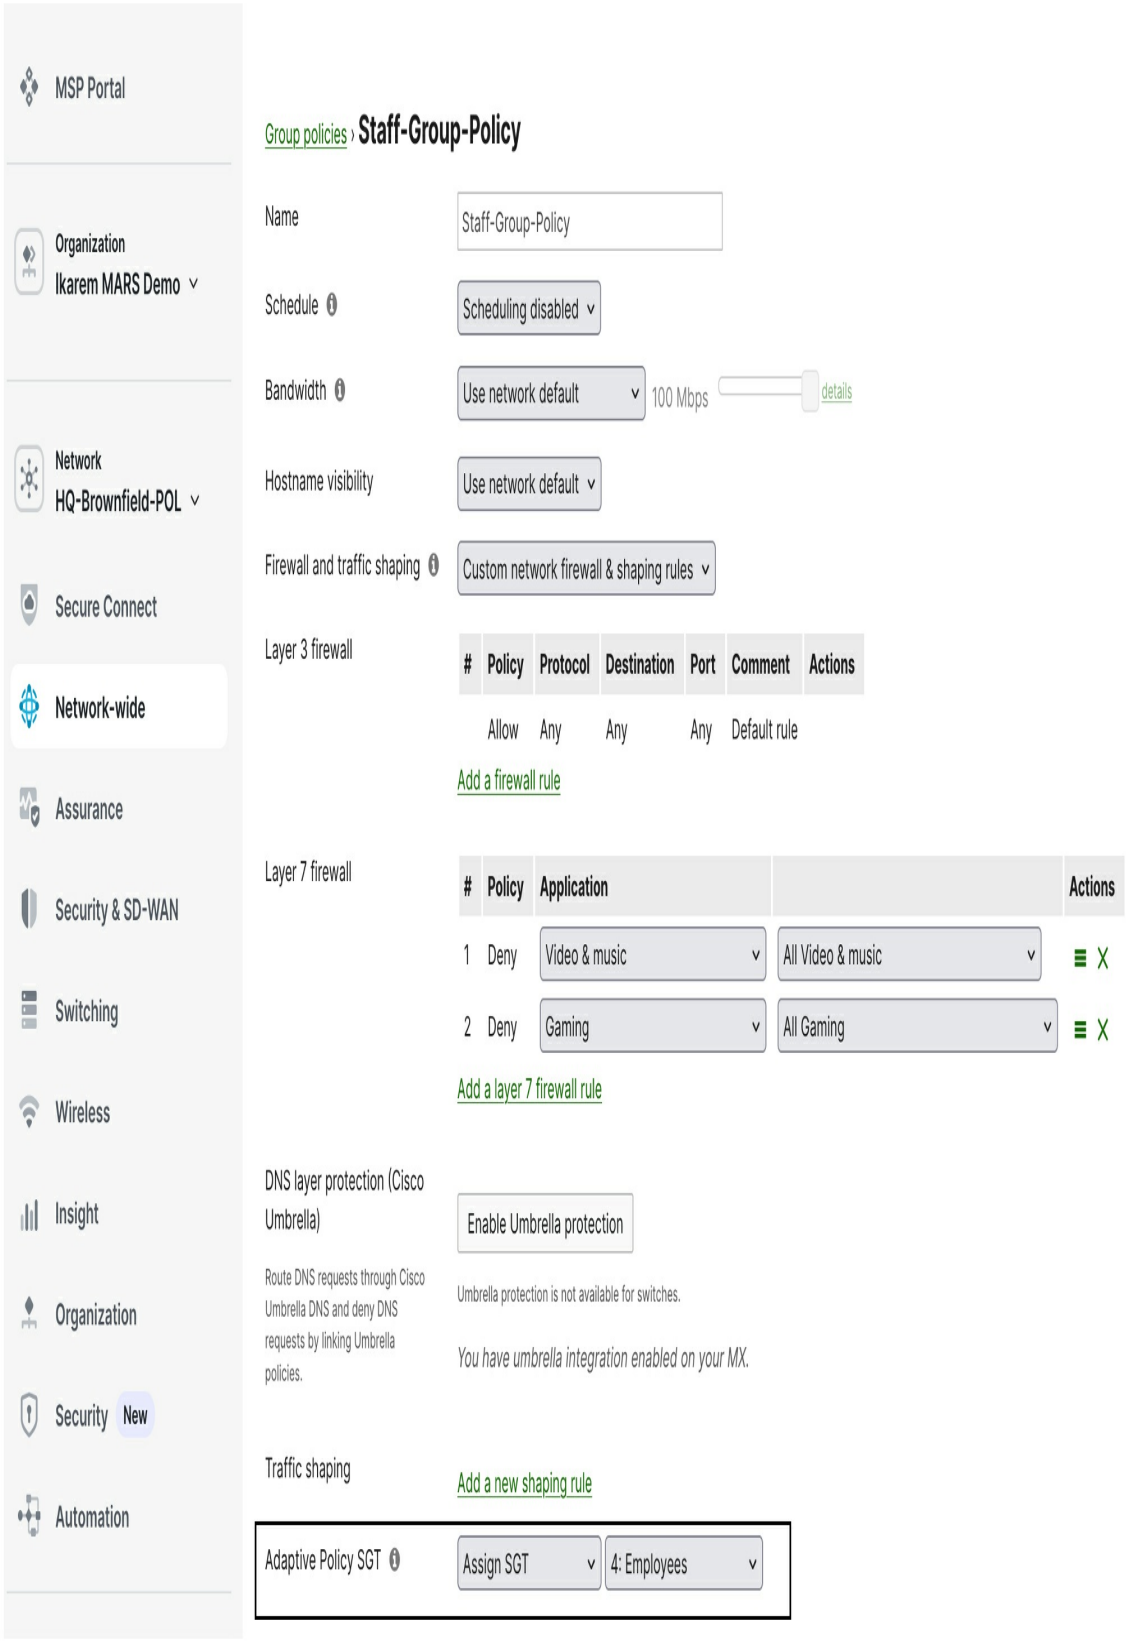
Technet24

Figure 12-47 *SSID Adaptive Policy Group*

Group Policy–Based Adaptive Policy Group Assignment

Each group policy defined at the network level can be used to assign an Adaptive Policy group to wireless users or endpoints. If the RADIUS server assigns the group policy name in the Filter-ID attribute, this is considered dynamic group assignment. Alternatively, as the administrator, you may manually assign specific group policies, such as a quarantine group policy, which tags all traffic originating from the endpoint and significantly restricts endpoint access.

To configure group policy–based SGT assignment, navigate to **Network-wide > Group Policies**. When creating a new group policy or editing an existing one, you can define the SGT for any packet belonging to an endpoint with this specific group policy assigned. Figure 12-48 illustrates group policy–level SGT assignment.


MSP Portal
Organization
Ikarem MARS Demo
Network
HQ-Brownfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security New
Automation
Group policies › Staff-Group-Policy
Name
Staff-Group-Policy
Schedule
Scheduling disabled
Bandwidth
Use network default
100 Mbps
details
Hostname visibility
Use network default
Firewall and traffic shaping
Custom network firewall & shaping rules
Layer 3 firewall
Policy Protocol Destination Port Comment Actions
Allow Any Any Any Default rule
Add a firewall rule
Layer 7 firewall
Policy Application Actions
1 Deny Video & music All Video & music
2 Deny Gaming All Gaming
Add a layer 7 firewall rule
DNS layer protection (Cisco Umbrella)
Enable Umbrella protection
Route DNS requests through Cisco Umbrella DNS and deny DNS requests by linking Umbrella policies.
Umbrella protection is not available for switches.
You have umbrella integration enabled on your MX.
Traffic shaping
Add a new shaping rule
Adaptive Policy SGT
Assign SGT
4: Employees

Technet24

Figure 12-48 *Group Policy SGT Assignment*

Static IP to SGT and Subnet:SGT Mappings

As explained in the "Access Policy Features for Wired Connections" section, IP to SGT mappings can be propagated by assigning policy objects created for specific IPs or subnets to policy groups. The configuration steps for wireless connections are the same. However, as previously mentioned, this method is primarily used for resources such as servers and workloads, and due to the specifics of wireless connections, it is not widely used for wireless devices.

RADIUS-Based Adaptive Policy Group Assignment

For RADIUS Adaptive Policy group assignment, you do not need to perform specific configuration on the access point. Most settings are handled on ISE, where SGTs and corresponding authorization policies should be defined. On the network level, the SSID must be configured for authentication with the RADIUS server, and Adaptive Policy must be enabled for the network at the organization level.

Adaptive Policy Inline Tagging

No specific settings are required at the access point level for inline tagging to function. If SGT support detection on the uplink results in confirmation of SGT support on the upstream switch, the access point will tag every user frame according to the existing mappings.

Adaptive Policy Enforcement

No specific configuration is required at the device level for enforcement to occur. Enforcement depends on all the settings outlined earlier in the chapter, as well as the configuration implemented at the organization level.

The main prerequisites for enforcement to occur are as follows:

- The packet should either arrive tagged with an SGT, or if the packet is switched locally on the access point, the IP to SGT mapping for the source IP must be available.
- The destination IP must have an IP to SGT mapping, which informs the access point of the destination SGT value.
- A policy must be assigned between the destination SGT and the source SGT.

Adaptive Policy Guidelines for Wireless Network Segmentation

Adaptive Policy Group Assignment: Order of Operations

Meraki access points use a slightly different priority order than switches to assign SGT when mapping is available from multiple sources. The priority order is as follows:

- SGT assignment based on the group policy. This method has the highest priority to ensure that administrators can always override priority if necessary.
- RADIUS session SGT assignment. When the endpoint does not have a group policy assigned, the SGT from the RADIUS session holds the next highest priority.
- IP/Subnet:SGT map. Group mapping using policy objects is considered the third priority.
- SGT mapping based on SSID. This is the final option to select.
- Last resort assignment of 0 (unknown).

To clarify the priority of operations, consider the following example: The Group Policy is configured with SGT 20, the RADIUS server assigns

Technet24

SGT value 10, there is an IP to SGT mapping for the endpoint with SGT 5, and the SSID is assigned to SGT 100. In this case, the SGT value from the Group Policy (SGT 20) will be used. The priority order for SGTs is SGT 20, SGT 10, SGT 5, and SGT 100.

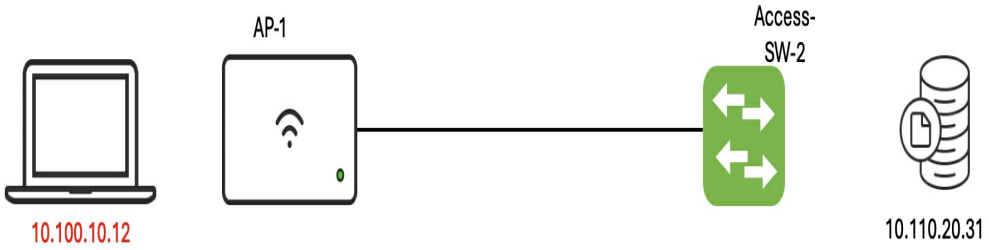
Adaptive Policy Enforcement Point

In the case of wireless network access, the same options available for wired access apply:

- Both the source and destination are connected to the same access point. In this scenario, enforcement occurs directly on the access point.
- One side of the communication is connected to the access point, while the other side is connected elsewhere in the network.

Figure 12-49 demonstrates a scenario where the access point is connected to a switch with inline tagging support. In this situation, packets destined for a remote destination leave the access point tagged with the appropriate source SGT, and enforcement occurs at the location where the destination is connected.





IP to SGT map

Dynamic: 10.100.10.12 : Staff (SGT: 20)

Static: 10.110.20.31: File-Servers (SGT: 21)

Policies

SGT:20 to SGT:21 = permit IP

IP to SGT map

Static: 10.110.20.31: File-Servers (SGT: 21)

Policies

SGT:20 to SGT:21 = permit IP

Technet24

Figure 12-49 *Enforcement for Wireless Client with Inline Tagging*

When inline tagging is not supported somewhere along the path, this may occur in two ways:

- The access point is connected to a switch that does not support TrustSec. If RADIUS authentication is performed on ISE, the Cisco network device (such as a switch, router, or firewall) where the destination is connected can be integrated with ISE over SXP to obtain dynamic IP to SGT mappings from ISE for authenticated users.
- The access point is connected to a TrustSec-capable switch, but inline tagging is not supported elsewhere along the path. In this case, you can use the same approach by connecting the Cisco device where the destination is connected to ISE over SXP.

> **Note**
>
> To learn mappings for authenticated users, ISE must receive accounting updates from the devices where the user is authenticated, which include the Framed-IP attribute.

Implementing Adaptive Policy for MX SD-WAN

Implementing Adaptive Policies for MX devices requires an understanding of the supported features and adherence to best practices to ensure reliable deployment. The following sections focus on these aspects of Adaptive Policies for MX Security Appliances.

Access Policy Features for MX SD-WAN

At the time of writing, Meraki MX Security Appliances differ from switches and access points in terms of Adaptive Policy support because they do not support dynamic SGT assignment from the

RADIUS server. The reason is that the primary role of MX appliances is not to authenticate end users but to ensure that the Adaptive Policy domain can be extended across SD-WAN, connecting different locations and guaranteeing consistent control and microsegmentation throughout the network.

MX security appliances support the following Adaptive Policy features:

- Port-level SGT assignment and inline tagging preservation (Routed mode MX)
- VLAN-level SGT assignment (Routed mode MX)
- Group policy–level SGT assignment (Routed mode MX)
- Inline tagging in Auto VPN (Routed mode and VPN concentrator mode MXs)

Port-Level SGT Assignment and Inline Tagging Preservation

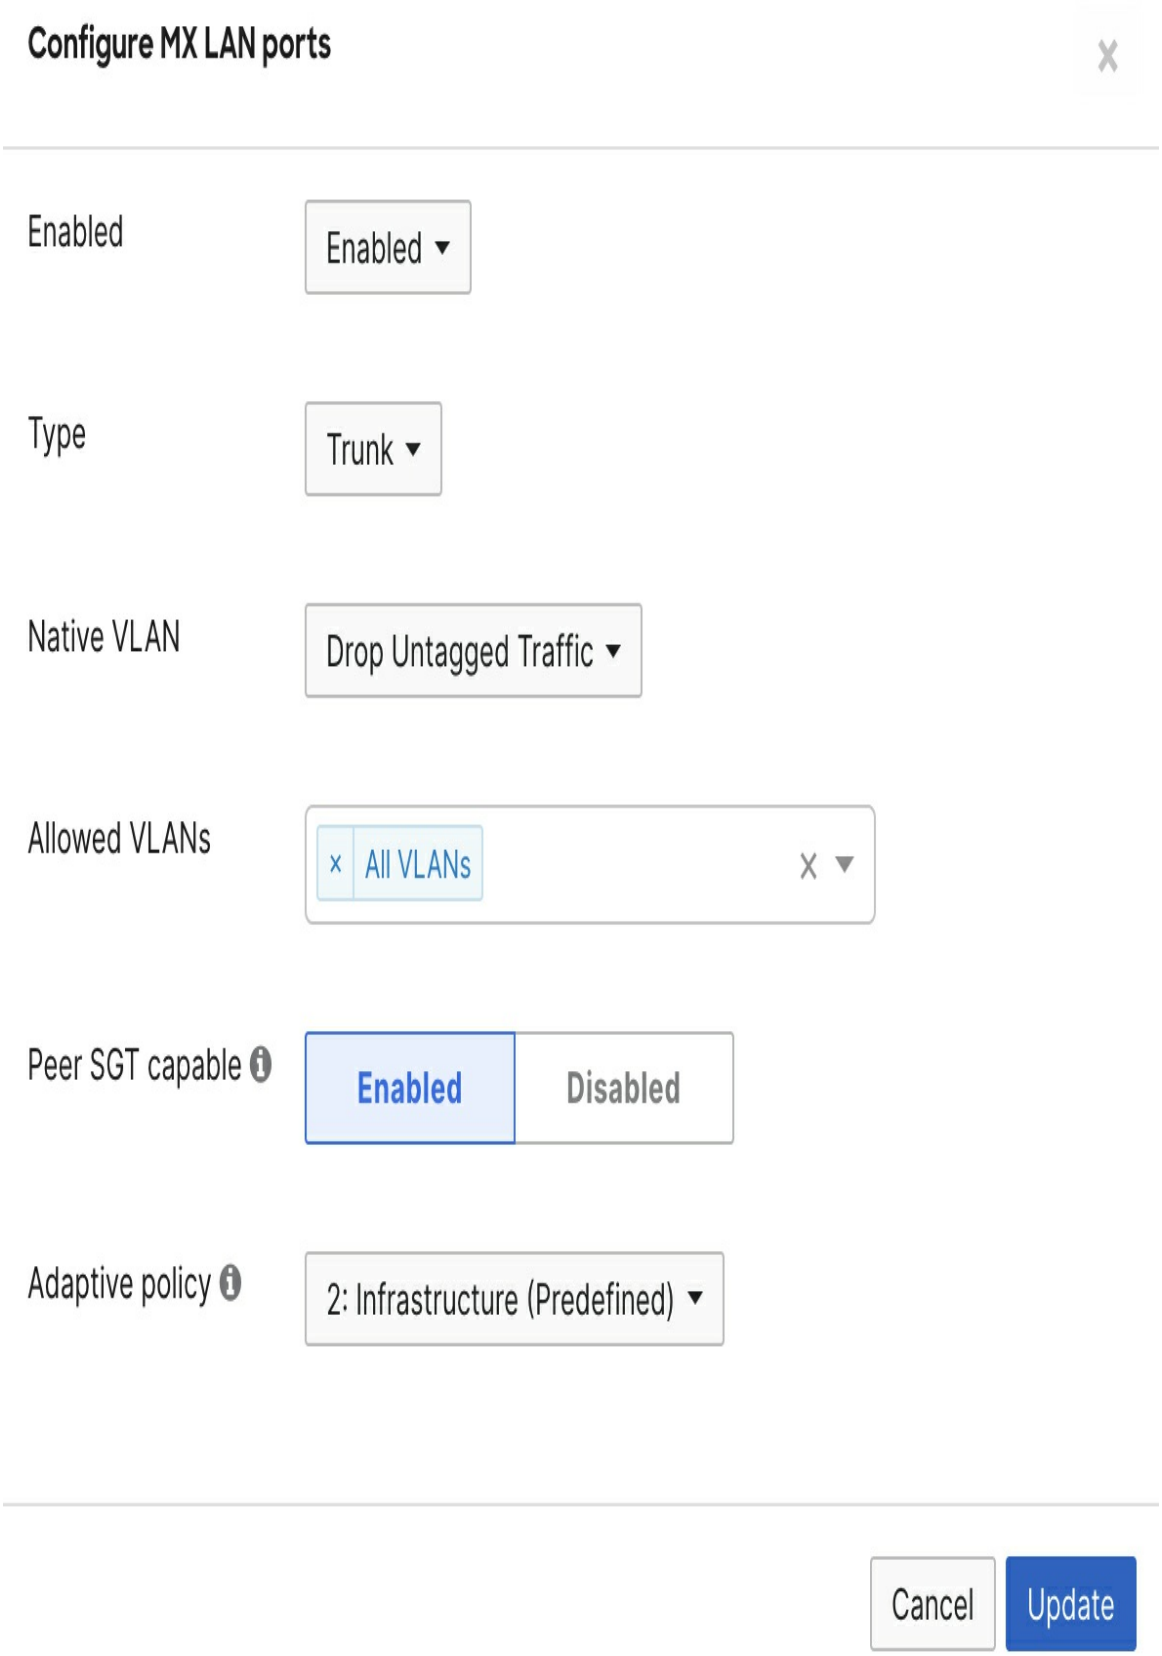
MX appliances support SGT assignment at the physical port level, ensuring that all untagged traffic received on the interface can be assigned an SGT. This setting can be combined with the **Peer SGT Capable** option.

When **Peer SGT Capable** is enabled and the MX is connected to an SGT-capable peer, the MX will receive and send frames with SGT tags. The SGT assigned to the port will only be used for outgoing frames if they do not have SGT tags.

When **Peer SGT Capable** is disabled and the port is configured with an SGT, the MX is instructed to tag incoming frames when the destination port has **Peer SGT Capable** enabled or packets are destined for an SGT-capable Auto VPN peer. Outgoing frames in this case remain untagged.

Technet24

Figure 12-50 shows an MX physical interface configured with a static SGT and **Peer SGT Capable** set to **Enabled**. To reconfigure an MX physical interface, navigate to **Security & SD-WAN > Addressing & VLANs**, scroll down to the **Per-port VLAN Settings** table, and click the port you wish to edit.


Configure MX LAN ports
Enabled
Enabled
Type
Trunk
Native VLAN
Drop Untagged Traffic
Allowed VLANs
All VLANs
Peer SGT capable
Enabled
Disabled
Adaptive policy
2: Infrastructure (Predefined)
Cancel
Update

Technet24

Figure 12-50 *MX Port-Level Adaptive Policy Settings*

VLAN-Level SGT Assignment

As an administrator, you can configure the MX to assign all traffic in a specific VLAN to the SGT of your choice. This functions similarly to SSID-level SGT assignment in wireless networks. To change VLAN settings, navigate to **Security & SD-WAN > Addressing & VLANs** and click the VLAN that needs to be reconfigured in the **Subnets** table. Figure 12-51 demonstrates how a VLAN can be configured with a specific SGT.

Modify VLAN

VLAN name

Default

VLAN ID

1

Group policy

None

Adaptive policy group

4: Employees

VPN mode

Enabled | Disabled

Next

Technet24

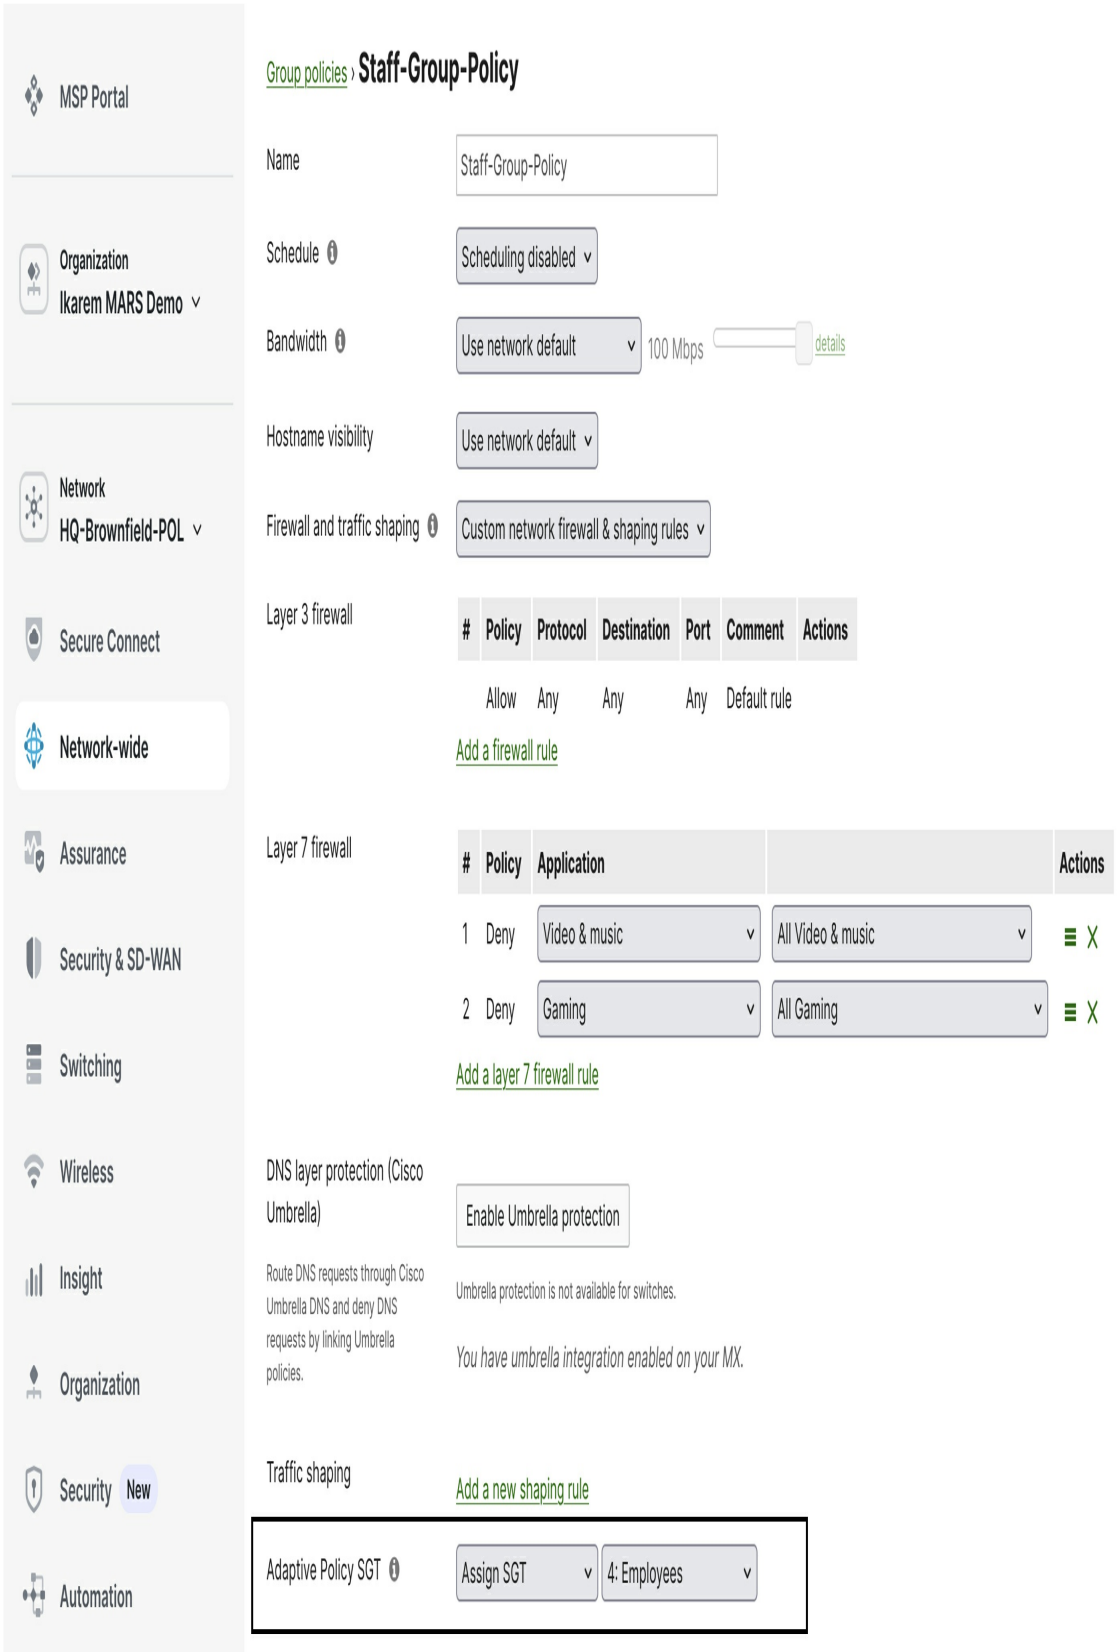
Figure 12-51 *MX VLAN-Level Adaptive Policy Settings*

Group Policy–Level SGT Assignment (Routed Mode MX)

Starting with MX 26.1.2, SGTs in the group policies can be used for the following purposes:

- **Classification:** In this case, traffic from the endpoint to which a group policy with an SGT was assigned will be marked with this SGT when packets leave the MX over a LAN port enabled with **Peer SGT Capable**, or if packets are destined for an Auto VPN peer enabled for inline tagging.
- **Enforcement:** It is important to understand that no Adaptive Policy is actually enforced on the endpoint itself. However, if traffic arrives at the MX on a LAN port with **Peer SGT Capable** enabled, or from an Auto VPN peer configured for inline tagging with the tag specified in the group policy, all settings from that group policy will be assigned to the traffic.

To configure SGT in the group policy, navigate to **Network-wide > Group Policies**, click the policy you wish to edit, and in the **Adaptive Policy SGT** section, select **Assign SGT** and choose the group name to be used. Figure 12-52 shows SGT assignment at the group policy level.


MSP Portal
Organization
Ikarem MARS Demo
Network
HQ-Brownfield-POL
Secure Connect
Network-wide
Assurance
Security & SD-WAN
Switching
Wireless
Insight
Organization
Security New
Automation
Group policies › Staff-Group-Policy
Name
Staff-Group-Policy
Schedule
Scheduling disabled
Bandwidth
Use network default
100 Mbps
details
Hostname visibility
Use network default
Firewall and traffic shaping
Custom network firewall & shaping rules
Layer 3 firewall
Policy Protocol Destination Port Comment Actions
Allow Any Any Any Default rule
Add a firewall rule
Layer 7 firewall
Policy Application Actions
1 Deny Video & music All Video & music
2 Deny Gaming All Gaming
Add a layer 7 firewall rule
DNS layer protection (Cisco Umbrella)
Enable Umbrella protection
Route DNS requests through Cisco Umbrella DNS and deny DNS requests by linking Umbrella policies.
Umbrella protection is not available for switches.
You have umbrella integration enabled on your MX.
Traffic shaping
Add a new shaping rule
Adaptive Policy SGT
Assign SGT
4: Employees

Technet24

Figure 12-52 *MX Group Policy SGT Assignment*

Inline Tagging in Auto VPN

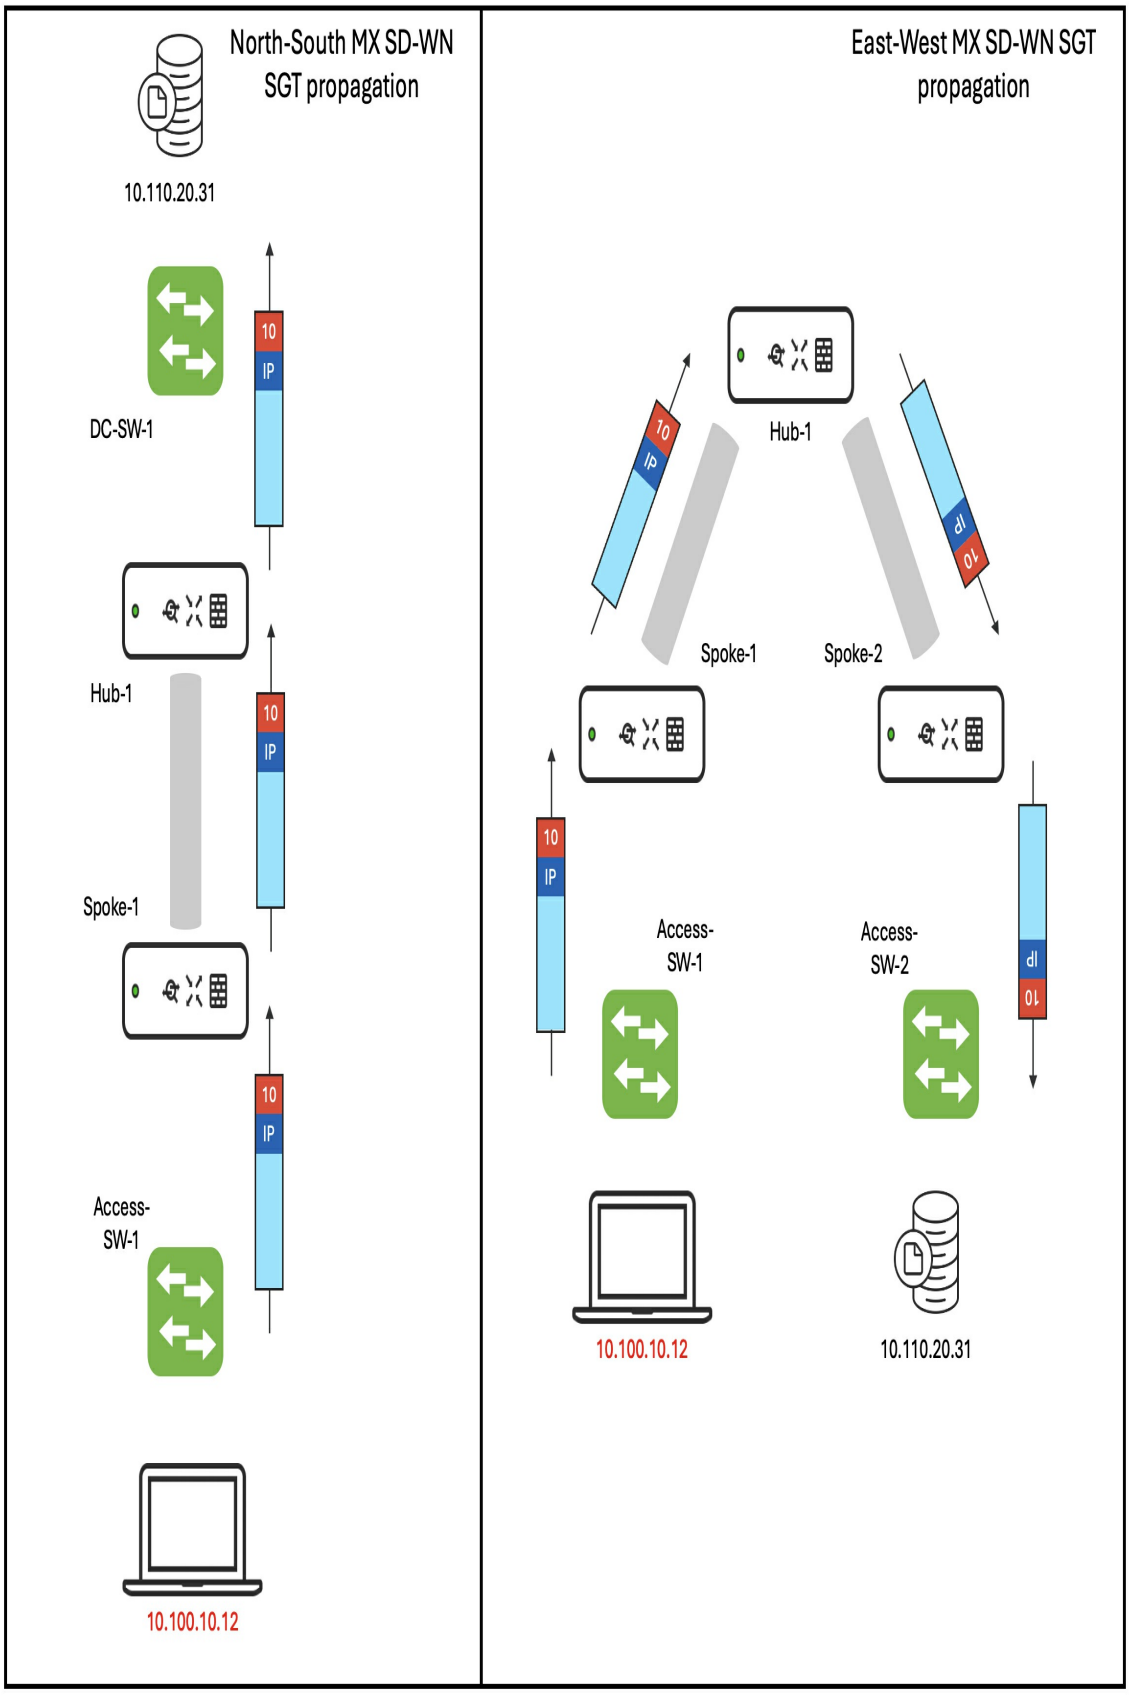
MX appliances can preserve SGT tagging when packets cross Auto VPN. Auto VPN is the Meraki implementation of SD-WAN fabric, interconnecting different sites with MX appliances installed across a wide variety of connectivity methods such as Internet, MPLS, and 4G, in a secure manner through Auto VPN tunnels.

For proper SGT preservation, you need to ensure that inline tagging is enabled at both locations on the LAN sides of the MX appliances connected over Auto VPN and that Auto VPN itself is configured for SGT propagation.

Since Auto VPN is most often implemented using a hub-and-spoke approach, end-to-end SGT propagation may include two scenarios:

- **Spoke-to-hub traffic:** This is typical north-south traffic between clients and resources either outside your organization or in on-premises data centers. If this is the only case that needs to be covered, it is sufficient to enable SGT support for Auto VPN on both the hub and the spoke.
- **Spoke-to-spoke traffic:** This is a standard example of east-west traffic when two devices connected behind different spokes exchange information. In this scenario, SGT support for Auto VPN must be enabled on both spokes as well as on the hub.

Figure 12-53 demonstrates both scenarios of SGT propagation across Auto VPN.


North-South MX SD-WN SGT propagation
10.110.20.31
DC-SW-1
10
IP
Hub-1
Spoke-1
Access-SW-1
10.100.10.12
East-West MX SD-WN SGT propagation
Hub-1
Spoke-1
Spoke-2
Access-SW-1
Access-SW-2
10.100.10.12
10.110.20.31

Technet24

Figure 12-53 *Scenarios of SGT Propagation over Auto VPN*

To enable SGT propagation over Auto VPN, navigate to **Security & SD-WAN > Site-to-Site VPN** and set **Peer SGT Capable** to **Enabled**. This applies to all Auto VPN peers. When this setting is enabled on the spoke, it applies to connections to all hubs defined in the spoke Auto VPN settings. When this configuration is enabled on the hub, it applies to all spokes connected to the hub as well as to all hub-to-hub tunnels. Figure 12-54 shows an example of **Peer SGT Capable** set to **Enabled** on the spoke.

Site-to-site VPN

Type

- ○ Off
  Do not participate in site-to-site VPN.
- ○ Hub (Mesh)
  Establish VPN tunnels with all hubs and dependent spokes.
- ◉ Spoke
  Establish VPN tunnels with selected hubs.

IPv6 default routing over VPN has not been implemented at this time. If you want to use an IPv4 default route over VPN, please either disable VPN access for your IPv6-enabled VLANs or disable IPv6 on VLANs where a default route over VPN is in use.

Hubs

#	Name	IPv4 default route	Actions
1	HQ-Brownfield-POL - appliance	☐	≡ X

Add a hub

VPN settings

Peer SGT capable — **Enabled** | Disabled

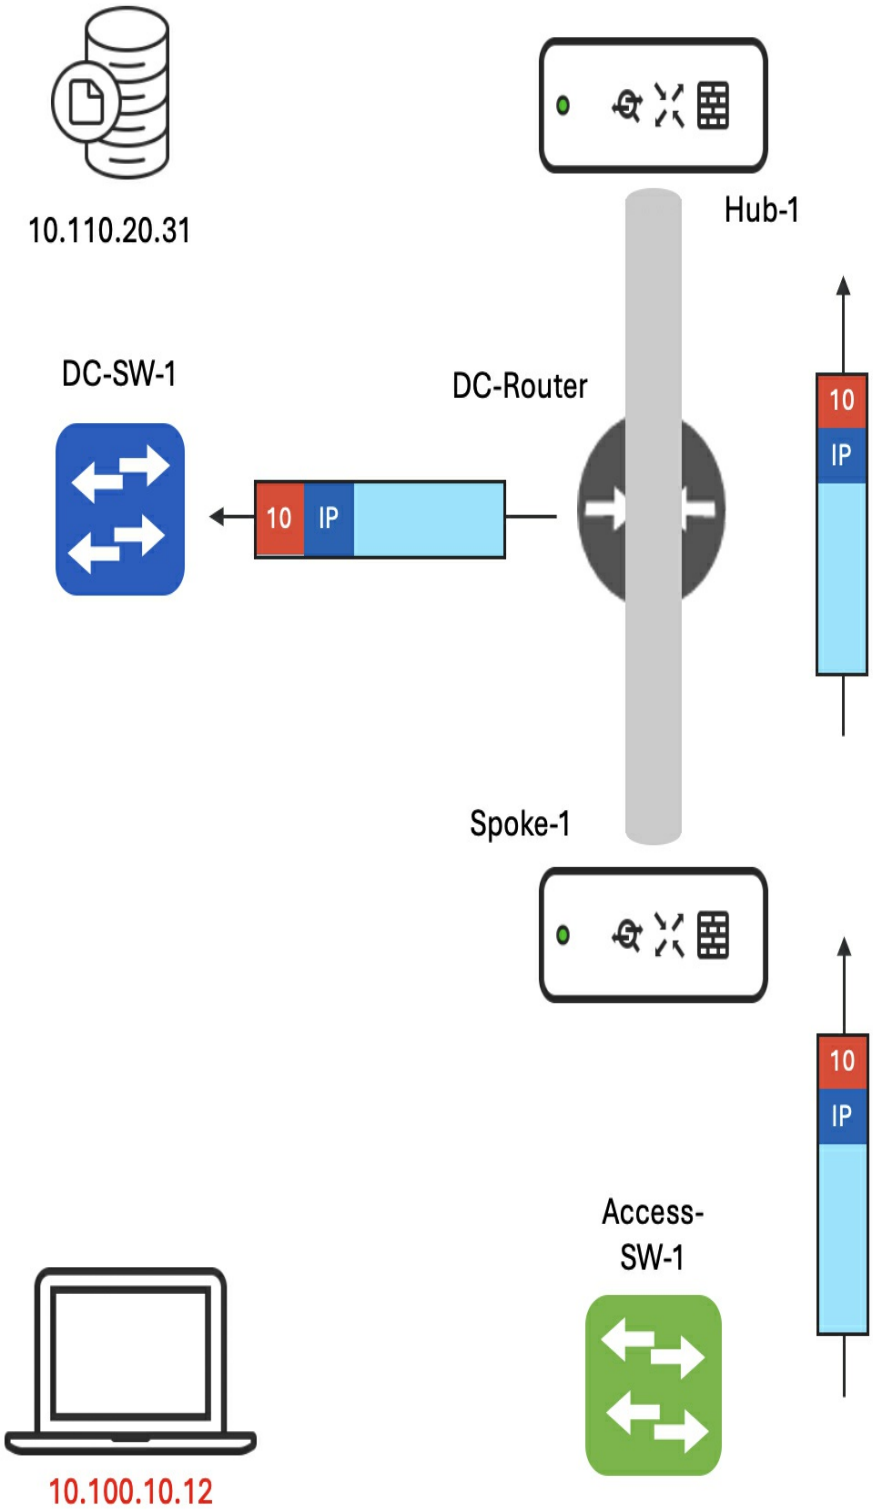
IPv4 VPN subnet translation — Enabled | **Disabled**

Technet24

Figure 12-54 *Auto VPN SGT Propagation Settings*

If the Hub MX operates in VPN concentrator mode, you need to enable SGT propagation on the WAN interface, since in concentrator mode the MX has only one active interface used for both internal and external connectivity. When connected within the data center, the upstream switch should also be enabled for SGT propagation. This configuration allows for classification, propagation, and enforcement of traffic received inside the Auto VPN tunnels and destined for resources in the data center.

Figure 12-55 explains SGT propagation with MX operating in VPN concentrator mode. The packet from the client is tagged on Access-SW1 and propagated to Spoke-1 MX. Since Spoke-1 has **Peer SGT Capable** enabled on the LAN port where Access-SW1 is connected, the tagged packet is accepted and propagated with the tag across the Auto VPN tunnel to Hub-1 MX because **Peer SGT Capable** is also enabled for Auto VPN. On the Hub-1 side, the packet is decrypted and forwarded through the same WAN interface with the SGT tag because the WAN interface is enabled for SGT propagation.


10.110.20.31
Hub-1
DC-SW-1
DC-Router
10
IP
10
IP
Spoke-1
10
IP
Access-
SW-1
10.100.10.12

Technet24

Figure 12-55 *Auto VPN SGT Propagation with Hub in VPN Concentrator Mode*

To configure SGT propagation on an MX operating in VPN concentrator mode, navigate to **Security & SD-WAN > Appliance Stats > Uplink** and click the pencil icon next to **Adaptive Policy**. This will open a new dialog where you can set **Peer SGT Capable** to **Enabled**, as demonstrated in Figure 12-56.


Summary
Uplink
IPv6 Prefixes
Location
Tools
Configuration
General
PUBLIC IP
86.48.219.122
hostname unknown
WAN 1
TYPE
IPv4
IPv6
CONFIGURED AS
Dynamic
Auto (Stateless)
STATUS
Not connected
Not connected
WAN 2
TYPE
IPv4
IPv6
CONFIGURED AS
Dynamic
Auto (Stateless)
STATUS
Not connected
Not connected
Revert port 2 from WAN to LAN
Adaptive Policy
WAN adaptive policy applies to b
Configure WAN
WAN adaptive policy applies to both IPv4 and IPv6
Peer SGT Capable
Enabled
Cancel
Update

Technet24

Figure 12-56 *Auto VPN SGT Propagation on VPN Concentrator Mode MX*

Adaptive Policy Guidelines for MX SD-WAN Network Segmentation

Adaptive Policy Group Assignment: Order of Operations

On MX security appliances, SGTs are assigned according to the following priority, with the top being the highest priority:

- **Group Policy Assignment:** SGTs assigned through group policy have the highest priority among all sources. On MX, group policy can be assigned to an endpoint manually, added to the VLAN, or assigned to a VPN user either in the VPN settings (for all VPN users) or from the RADIUS server.
- **RADIUS Assignment:** At the time of writing, MX does not support SGT assignment directly over RADIUS, but for future support, RADIUS assignment is designated as the second-highest priority.
- **Static Port Assignment:** SGT assigned directly to the port holds third priority.
- **IP to SGT Assignment:** Currently, the dashboard does not propagate static IP to SGT mappings to MX appliances, but for future support, this method is assigned fourth priority.
- **VLAN Assignment:** SGT assignment based on VLAN is last in the priority list.

Adaptive Policy Group Propagation Considerations

For traffic between different MX LAN interfaces, the east-west traffic handling rules apply. The rules for Adaptive Policy group propagation

in the east-west use case are explained in Table 12-3.

Table 12-3 *Adaptive Policy Group Propagation Rules for East-West Traffic*

Technet24

Ingress Interface Peer SGT Capable?	Ingress Packet Type	Ingress Interface Behavior	Egress Interface Peer SGT Capable?	Egress Packet Type	Egress Interface Behavior
Enable	SGT packet	Preserve and forward the SGT packet	Enable	SGT packet	Forward the SGT packet
Disable	SGT packet	Strip the SGT tag and forward the non-SGT packet	Enable	Add interface-level SGT and forward the SGT packet	Forward the SGT packet
Enable	Non-SGT packet	Add interface-level SGT and forward the SGT packet	Enable	SGT packet	Forward the SGT packet
Disable	Non-SGT packet	Forward the non-SGT packet	Enable	Add interface-level SGT and forward the SGT packet	Forward the SGT packet
Enable	SGT packet	Preserve and forward the SGT packet	Disable	Strip the SGT tag and forward the non-SGT packet	Forward the non-SGT packet
Disable	SGT packet	Strip the SGT tag and forward the non-SGT packet	Disable	Non-SGT packet	Forward the non-SGT packet
Enable	Non-SGT packet	Add interface-level SGT and forward the SGT packet	Disable	Strip the SGT tag and forward the non-SGT packet	Forward the non-SGT packet
Disable	Non-SGT packet	Forward the non-SGT packet	Disable	Non-SGT packet	Forward the non-SGT packet

For Auto VPN traffic between different MX appliances, the north-south traffic handling rules apply. The rules for Adaptive Policy group propagation in the north-south use case are explained in Table 12-4.

Table 12-4 *Adaptive Policy Group Propagation Rules for North-South Traffic*

Technet24

Local MX: VPN Peer SGT Capable	Ingress Packet Type	Ingress Behavior for AutoVPN Transport	Remote MX: VPN Peer SGT Capable	Egress Packet Type	Egress Interface Behavior
Enable	SGT packet	Preserve and forward the SGT packet	Enable	SGT packet	Forward the SGT packet
Disable	SGT packet	Strip the SGT tag and forward the non-SGT packet	Enable	Non-SGT packet	Forward the non-SGT packet
Enable	Non-SGT packet	Add SGT tag=0 and forward the SGT packet	Enable	SGT packet	Forward the SGT packet
Disable	Non-SGT packet	Forward the non-SGT packet	Enable	Non-SGT packet	Forward the non-SGT packet
Enable	SGT packet	Strip the SGT tag and forward the non-SGT packet	Disable	Non-SGT packet	Forward the non-SGT packet
Disable	SGT packet	Strip the SGT tag and forward the non-SGT packet	Disable	Non-SGT packet	Forward the non-SGT packet
Enable	Non-SGT packet	Forward the non-SGT packet	Disable	Non-SGT packet	Forward the non-SGT packet
Disable	Non-SGT packet	Forward the non-SGT packet	Disable	Non-SGT packet	Forward the non-SGT packet

Adaptive Policy Verification and Troubleshooting

From the perspective of Adaptive Policy monitoring and troubleshooting, the following features are available in the Meraki dashboard at the time of writing:

- **Visibility into assigned Adaptive Policy groups in client details:** This functionality is available for both wired and wireless connections. Administrators can verify which Adaptive Policy group is currently assigned to the authenticated client.
- **SGACL logs:** When administrators enable logging in Adaptive Policy custom ACLs, it is possible to monitor ACL hits in the event logs. At the time of writing, this feature is only supported for MS390 and C9X00-M platforms. However, a separate logging category is already available for the rest of the Adaptive Policy–capable switches to ensure future support.
- **Adaptive Policy Hit Counters Live Tool:** This dashboard tool displays packet counters for packets processed between different Adaptive Policy groups. The counters include permitted and denied packets for both IPv4 and IPv6. Currently, this feature is only available for MS390 and C9X00-M platforms.

Visibility into Assigned Adaptive Policy Groups in Client Details

To verify Adaptive Policy groups for authenticated clients, navigate to **Network-wide > Clients** in the network where the client is connected and ensure that the **Adaptive Policy Group** column is included in the clients table. We explained how to add additional columns in the section "Meraki Dashboard NAC Verification and Troubleshooting" in

Technet24

Chapter 11. Figure 12-57 shows the clients table with the highlighted client named Win-10-WiFi assigned to the Adaptive Policy group Employees with SGT tag 4.

MSP Portal

Organization
Ikarem MARS Demo

Network
Branch-Greenfield-POL

Secure Connect

Network-wide

Assurance

Security & SD-WAN

Switching

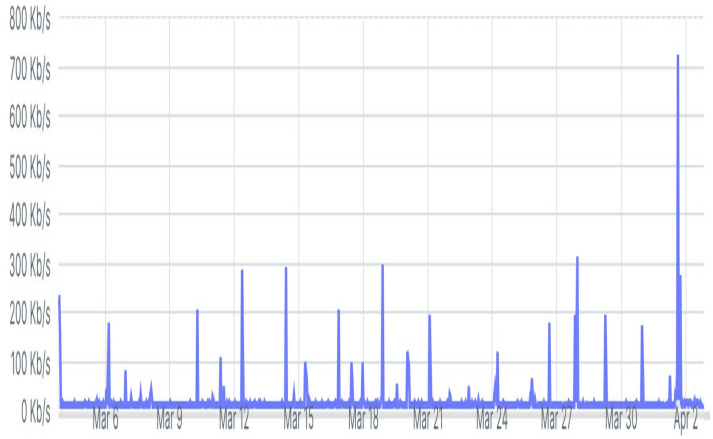
Wireless

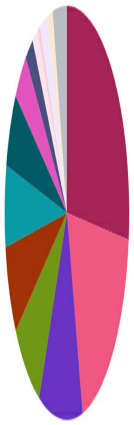
Insight

Usage and clients

Filters

Usage 3.23 GB (↓2.56 GB, ↑683.47 MB)





Applications details

Search for clients

Filters 5 results

Download

Status	Description	Usage	Device type, OS	IPv4 address	Adaptive policy group	CDP/LLDP
Wired	Win-10-WiFi	47.71 MB	VMware	172.22.15.161	–	–
Wired	SEP006440B4102E	57.88 MB	Cisco IP Phone CP-7965G	10.110.28.2	–	SEP00644
Wired	Aim into the Cup	2.59 GB	Cisco Systems	10.110.23.2	–	–
Wired	lnx1-krk	210.36 MB	VMware	10.100.11.3	–	–
Wireless	Win-10-WiFi	344.29 MB	Windows 10	10.110.23.4	4: Employees	–

Technet24

Figure 12-57 *Adaptive Policy Groups in the Client Table*

SGACL Logs

To view the syslog events for the SGACL line configured with logging enabled, navigate to **Network-wide > Event Log** or **Assurance > Event Log** and change the logging source to **For Catalyst Switches**. SGACLs generate two types of events that can be verified and investigated in the event logs: IPACCESSLOGP and SGACLHIT. Figure 12-58 shows the IPACCESSLOGP event type.



Figure 12-58 *IPACCESSLOGP SGACL Events*

Figure 12-59 shows the SGACLHIT events.

Technet24



Figure 12-59 *SGACLHIT SGACL Events*

For classic Meraki MS switches, use **For Switches** as the event source and filter with the event type **Adaptive Policy ACL Log**.

> **Note**
>
> At the time of writing, the Adaptive Policy ACL log is not implemented on classic MS switches.

Figure 12-60 shows the filter that can be used in the future with classic MS switches.



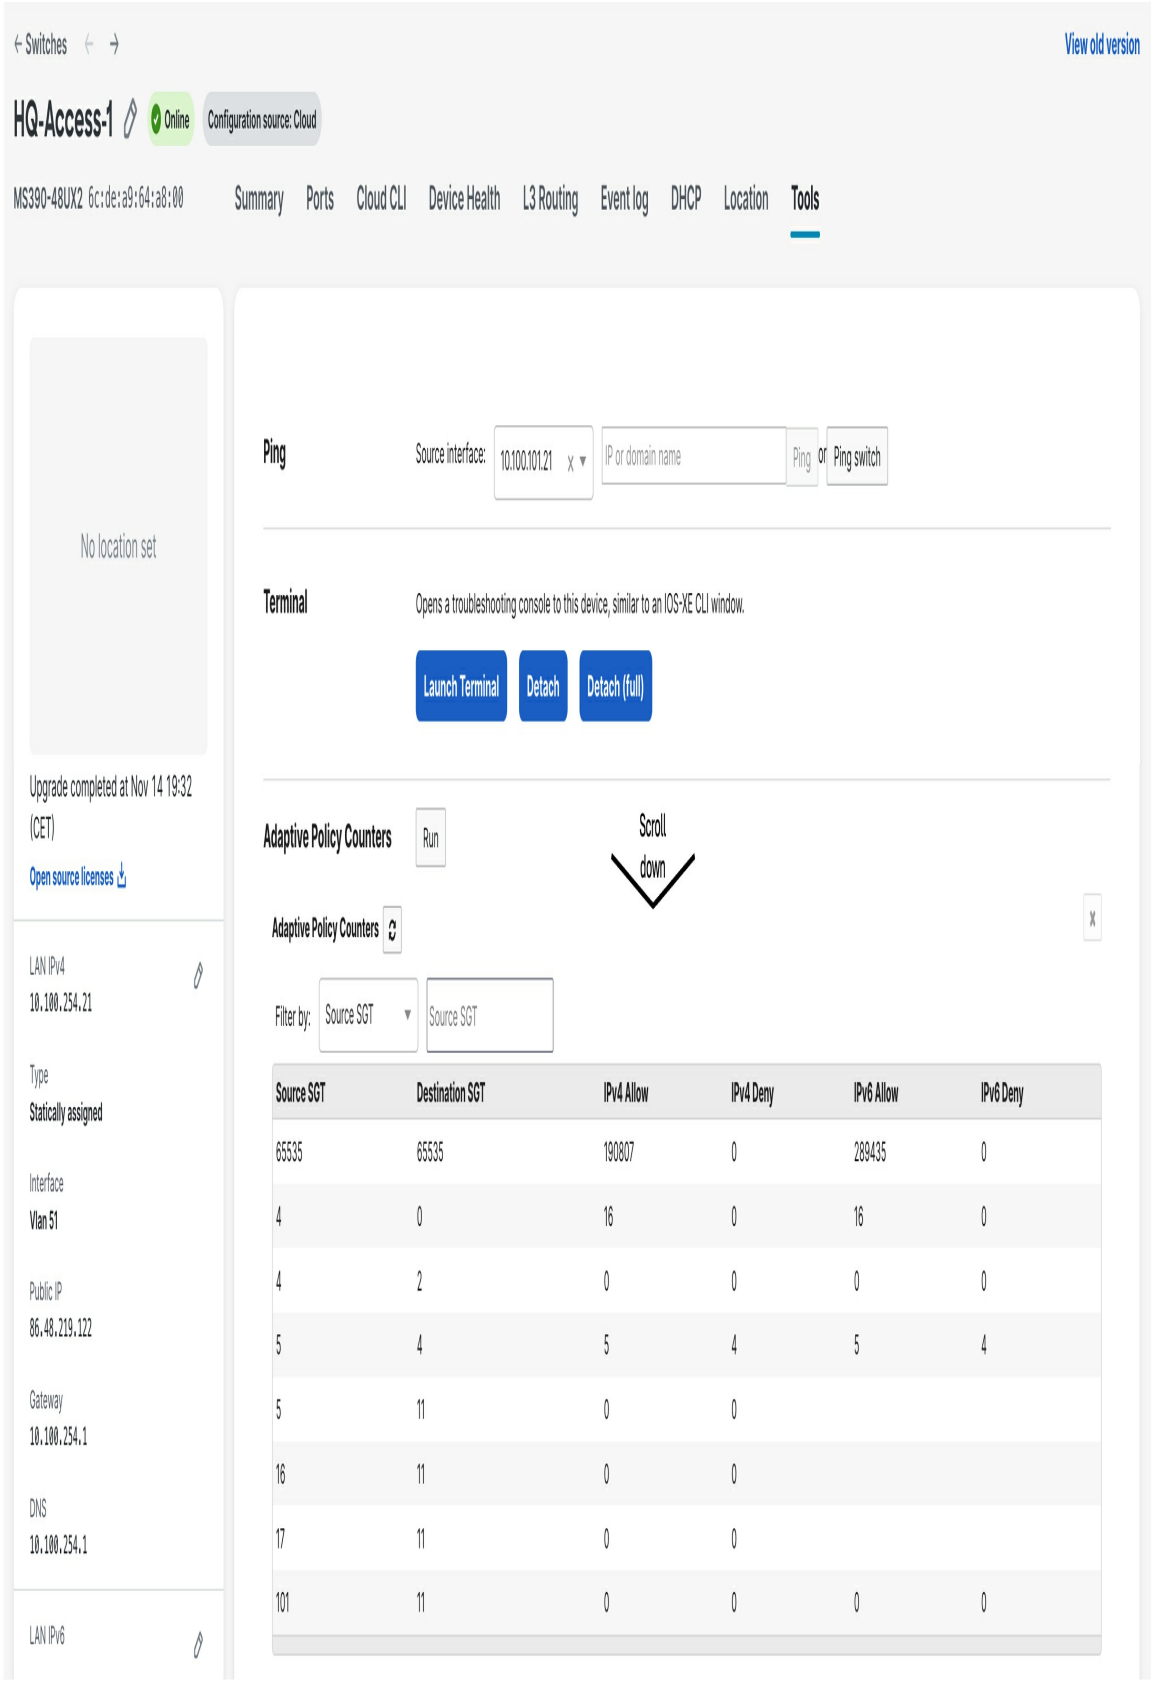
Figure 12-60 *Adaptive Policy ACL Log Filter*

Adaptive Policy Hit Counters Live Tool

The Hit Counters Live Tool is available on the **Tools** page of supported switches. To access it, navigate to **Switching > Switches** and click the switch where counters need to be investigated. When the **Switch** page

Technet24

opens, go to the **Tools** tab and scroll down to the **Adaptive Policy Counters** section. Figure 12-61 shows Adaptive Policy Hit Counters Live Tool statistics.


← Switches
View old version
HQ-Access-1
Online
Configuration source: Cloud
MS390-48UX2 6c:de:a9:64:a8:00
Summary
Ports
Cloud CLI
Device Health
L3 Routing
Event log
DHCP
Location
Tools
No location set
Upgrade completed at Nov 14 19:32 (CET)
Open source licenses
LAN IPv4
10.100.254.21
Type
Statically assigned
Interface
Vlan 51
Public IP
86.48.219.122
Gateway
10.100.254.1
DNS
10.100.254.1
LAN IPv6
Ping
Source interface:
10.100.101.21
IP or domain name
Ping
or
Ping switch
Terminal
Opens a troubleshooting console to this device, similar to an IOS-XE CLI window.
Launch Terminal
Detach
Detach (full)
Adaptive Policy Counters
Run
Scroll down
Adaptive Policy Counters
Filter by:
Source SGT
Source SGT
Source SGT
Destination SGT
IPv4 Allow
IPv4 Deny
IPv6 Allow
IPv6 Deny
65535 65535 190807 0 289435 0
4 0 16 0 16 0
4 2 0 0 0 0
5 4 5 4 5 4
5 11 0 0
16 11 0 0
17 11 0 0
101 11 0 0 0 0

Technet24

Figure 12-61 *Adaptive Policy Hit Counters Live Tool*

Summary

The chapter provides an in-depth overview of Meraki Adaptive Policy, charting the progression from basic access control methods to advanced, intent-based microsegmentation. It examines how Meraki's approach evolved from Layer 3/4 ACLs to identity-based group policies and ultimately to Adaptive Policy, which enables dynamic security segmentation with SGTs and SGACLs. The chapter outlines the advantages of Adaptive Policy as a cloud-managed implementation of Cisco TrustSec, allowing you to streamline policy creation and enforcement through the Meraki dashboard while maintaining policy consistency with Cisco Identity Services Engine via the ISE Meraki Connector.

The text explores how to synchronize SGTs and SGACLs between ISE and Meraki using the Connector's guided setup, and reviews the hardware, software, and licensing requirements for enabling Adaptive Policy across switches, access points, and security appliances. The implementation section guides you through the process of configuring policy objects, assigning tags, and building enforcement matrices to ensure that security policies operate effectively at every point in your network.

The chapter also discusses the specifics of SGT assignment priority on different platforms, including switches, access points, and MX security appliances, so you understand how enforcement decisions are made. In addition, it reviews how Adaptive Policy integrates with SD-WAN, preserving SGTs across Auto VPN tunnels to maintain segmentation and security in distributed environments. Monitoring and troubleshooting tools such as event logs, client details, and SGACL counters are covered to help you maintain visibility and operational

assurance. By mastering these concepts, you are equipped to implement scalable, granular security controls throughout your Meraki infrastructure.

Technet24

Chapter 13. Remote Access Management for On-Prem and SSE

Introduction

This chapter focuses on the theory and practice of remote access virtual private networking and explains how you can protect remote access VPNs with Cisco ISE network access control (NAC). You will study the progression of remote access technologies, from early VPN protocols to advanced Secure Access Service Edge and Security Service Edge architectures. The chapter guides you through deploying secure remote connectivity using Cisco Secure Firewall and Firewall Management Center. You will learn how to set up network access control and authentication processes with Cisco Identity Services Engine and how to integrate Cisco SSE cloud solutions with ISE for directory synchronization and uniform policy enforcement. Finally, you will review troubleshooting and validation methods for remote access and network access control configurations in both on-premises and cloud-managed environments.

Remote Access Evolution

Remote access VPN technologies started to reach end users in the mid to late 1990s with the development and release of Point-to-Point Tunneling Protocol (PPTP) by Microsoft, Ascend Communications, and 3Com. With the era of low Internet prices and the wide adoption of

technologies such as Wi-Fi, remote access VPN became an essential component for any enterprise.

The introduction of Cisco Borderless Network architecture in 2009 acknowledged the growing importance of mobile users and technologies and delivered enterprise security architecture to securely connect anyone, anywhere, on any device, to any resource, seamlessly and reliably.

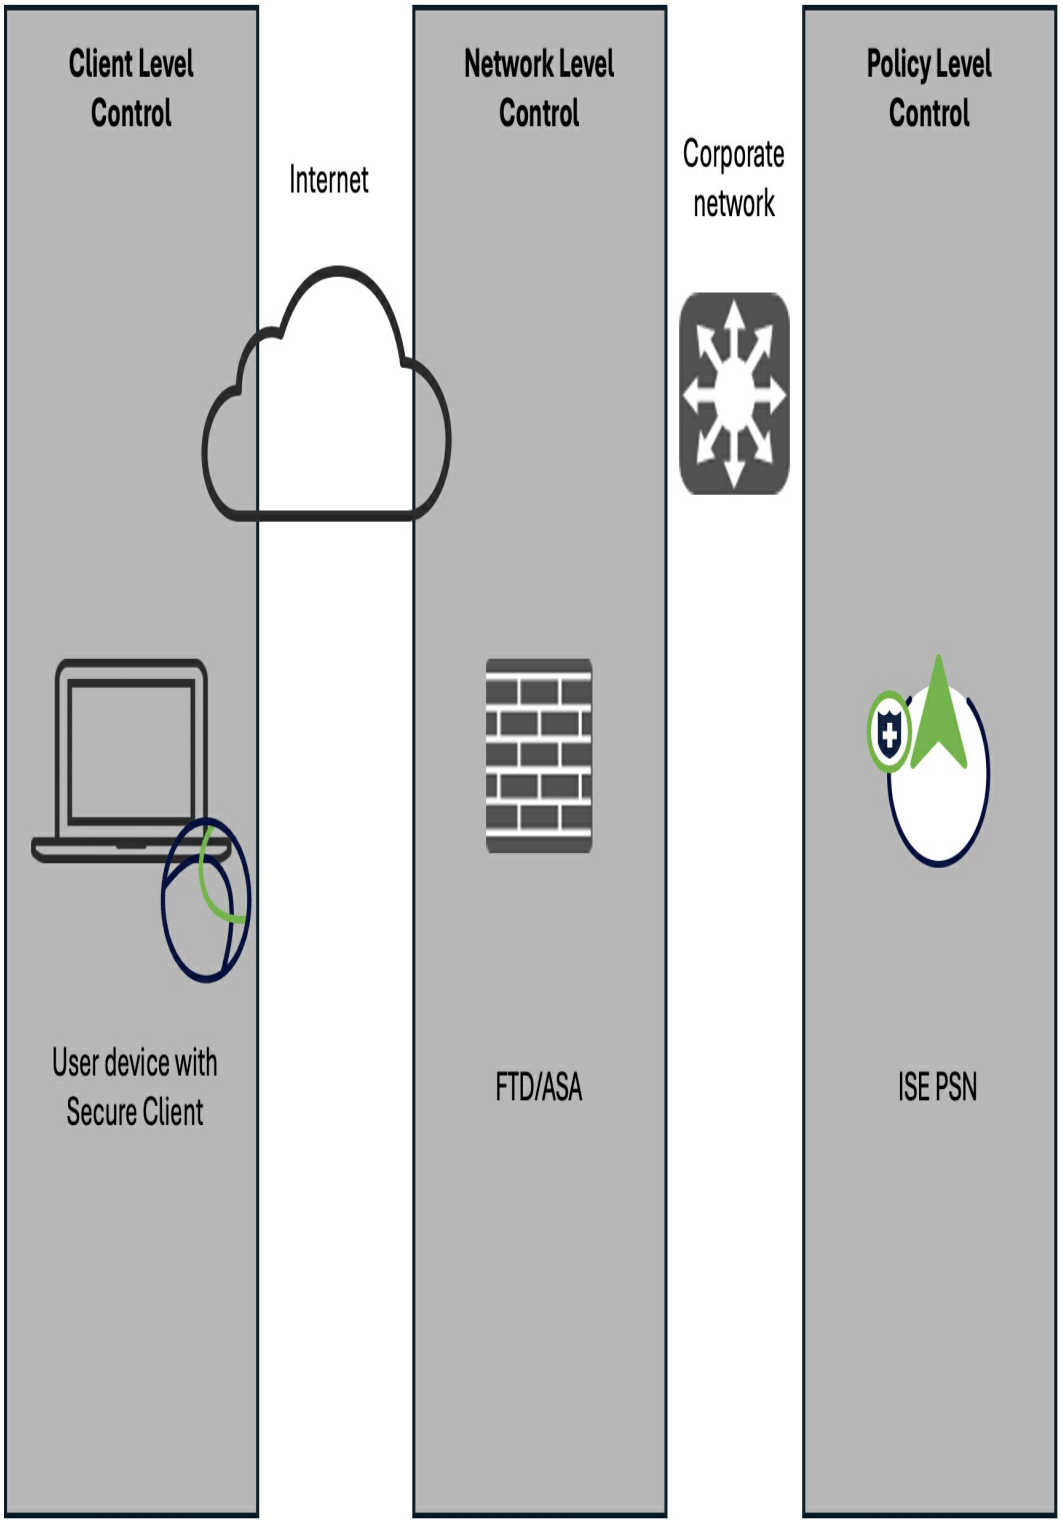
The Cisco Borderless Network architecture itself was broader than just a remote access use case, but the remote access VPN was one of the key areas of focus due to rising demand for users to work from anywhere. From the perspective of components, the remote access VPN was represented by

- **Client-level control:** Cisco AnyConnect Secure Client, later renamed Cisco Secure Client, represents the client-level control. This modular software delivers a variety of security services required for remote users. The key component of Cisco Secure Client is the VPN module, which establishes and maintains connections to remote access VPN headends. Cisco Secure Client also contains other modules designed to enhance the security of remote access VPN, such as the ISE posture module, which verifies endpoint posture when you establish a VPN connection.
- **Network-level control:** Cisco Adaptive Security Appliance (ASA) initially represented the network-level control, which was later complemented and ultimately replaced by Cisco Secure Firewall. The network control level delivers a secure entry point into the corporate network, where remote users and endpoints authenticate and receive network access based on the full connection context, such as location, posture state, and time of day when the connection is attempted.

Technet24

- **Policy-level control:** Cisco Access Control Server (ACS) 5.X initially delivered the policy-level control, which was later complemented and ultimately replaced by Cisco Identity Services Engine. The purpose of the policy level is to ensure that any remote access attempt is authenticated and authorized based on rich connection context.

The control levels of the traditional remote access VPN system are shown in Figure 13-1.


Client Level Control
Internet
Network Level Control
Corporate network
Policy Level Control
User device with Secure Client
FTD/ASA
ISE PSN

Technet24

Figure 13-1 *Remote Access VPN Control Levels*

While this system can deliver network access control, you should consider several elements that are not present in the figure but are essential for the modern organization's cybersecurity posture, such as network-level intrusion prevention systems (IPSs), SSL and deep packet inspection, data loss prevention (DLP), and, with the rise of cloud-based applications, modern security solutions such as Cloud Access Security Broker (CASB). Implementing all these solutions at each remote access VPN entry point requires significant investment both operationally, because someone needs to control and operate all these distributed systems, and financially, because the growth of remote access VPN entry points means that more security systems must be purchased and deployed.

These factors, combined with the growing realization that previous security approaches no longer provide adequate protection due to flat network access after authentication and authorization, have triggered the development and adoption of new technologies such as Zero Trust Network Access (ZTNA), Security Service Edge (SSE), and Secure Access Service Edge (SASE).

All three approaches are typically delivered as cloud-based security solutions, and it is often unclear what each of them does specifically, where they overlap, and where they differ from each other.

First, you need to understand that SSE and SASE describe the entire solution, while ZTNA is a service within those solutions, similar to other security services such as remote access VPN or Cloud-Delivered Firewall.

Both SSE and SASE deliver next-generation cybersecurity solutions in the cloud with the Zero Trust approach in mind. SSE and SASE enable organizations to reduce the number of security systems required on-

premises, since cloud providers now deliver these services in a centralized manner.

Here are the common security services offered by both SSE and SASE:

- **ZTNA:** Allows granular access control when users access specific applications. It works both in the browser and with the help of a client — for example, the ZTNA module inside Cisco Secure Client.
- **Remote Access VPN:** Delivers traditional remote access VPN with the headend located in the cloud. When you establish a connection to the cloud-based VPN headend, you can apply all cloud security features and controls to both authentication and authorization, as well as to the traffic.
- **Secure Web Gateway (SWG):** Delivers granular access control to web resources, typically at both DNS and HTTP/HTTPS levels.
- **Cloud Access Security Broker (CASB):** Enforces security policies between cloud service users and cloud applications.
- **Cloud-Delivered Firewall (CDF):** Provides granular access control at the network level. It is typically combined with cloud-based IPS solutions.
- **Data Loss Prevention (DLP):** Controls user interactions with critical company data to prevent this data from being stolen or exfiltrated.
- **Network connectivity:** Typically represented by the ability of SSE and SASE to establish IPsec tunnels from the cloud back to on-premises sites, connecting users who access the network with ZTNA or remote access VPN to existing applications and services.

Where exactly do SSE and SASE differ? The difference lies in the level of network connectivity support. While both solutions support IPsec connectivity to corporate networks, SASE delivers a new approach by

Technet24

allowing organizations to terminate SD-WAN connections in the same security cloud. This makes the SASE cloud a central point that connects everything and everyone within the organization.

Figure 13-2 shows how the combination of SD-WAN and SSE features creates the SASE solution.

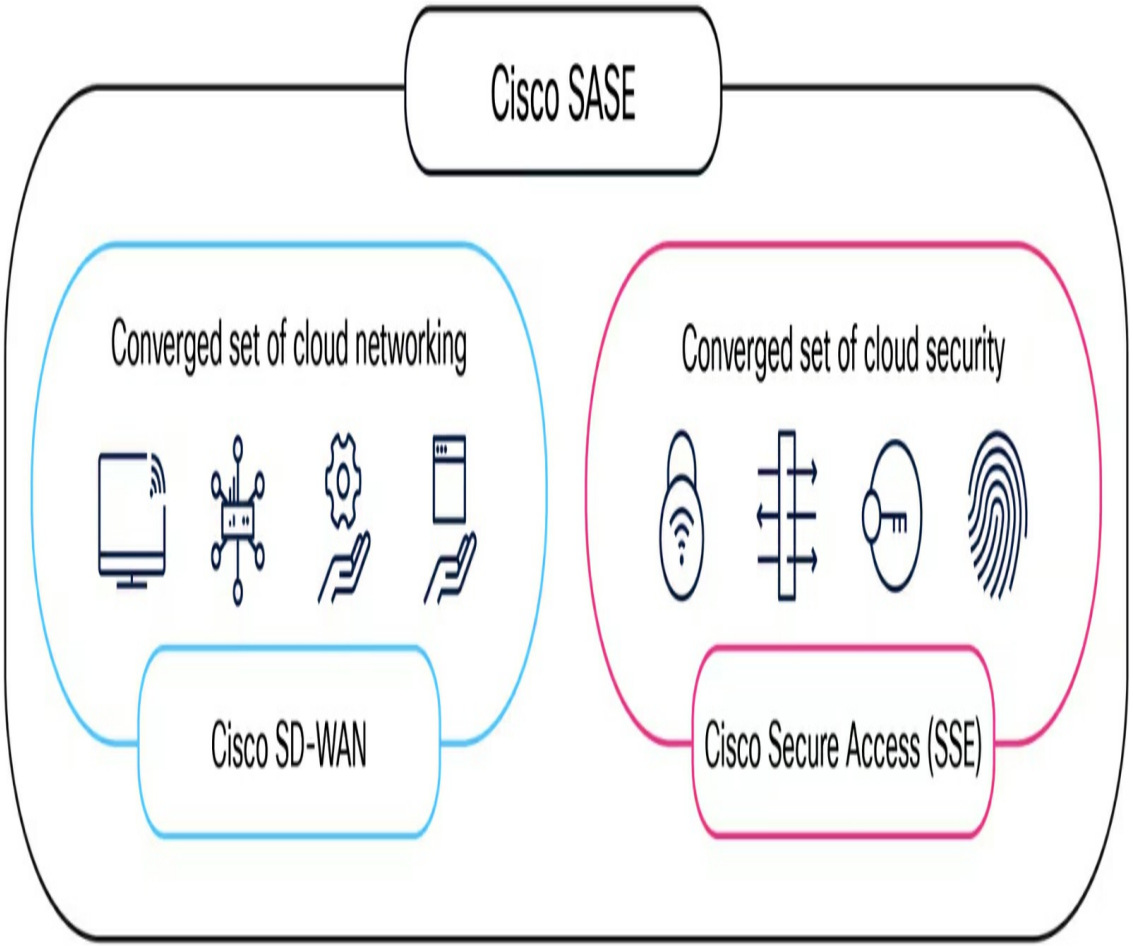


Figure 13-2 *SASE Solution Components*

Common NAC Use Cases with Cisco Secure Firewall

When you use Cisco ISE as a network access control solution for

remote access VPN with Cisco Secure Firewall, the available use cases fall into two categories:

- **Basic network access control:** You can manage and control all possible combinations of authentication and authorization using ISE as the primary or secondary authentication and authorization source. Basic network access control is also provided through the assignment of downloadable ACLs or VPN group policies.
- **Advanced network access control:** This category includes use cases that deliver the next level of endpoint control by providing posture services, performing endpoint state validation through unified endpoint management (UEM) or mobile device management (MDM) solutions, or by combining both approaches.

Basic Network Access Control with Secure Firewall

The basic network access control category includes three main use cases, described next.

Basic Authentication – overview

Basic authentication refers to standard username- and password-based authentication with ISE as the primary or secondary authentication server. With this approach, you can use ISE to assign the authorization policy based on user or device details, connection type, and other factors. As authorization attributes, you can primarily use Framed-IP-Address, downloadable ACL (DACL), and VPN group policy name, which ISE supplies back to Secure Firewall inside the RADIUS Class attribute.

Figure 13-3 shows the basic authentication flow.

Technet24

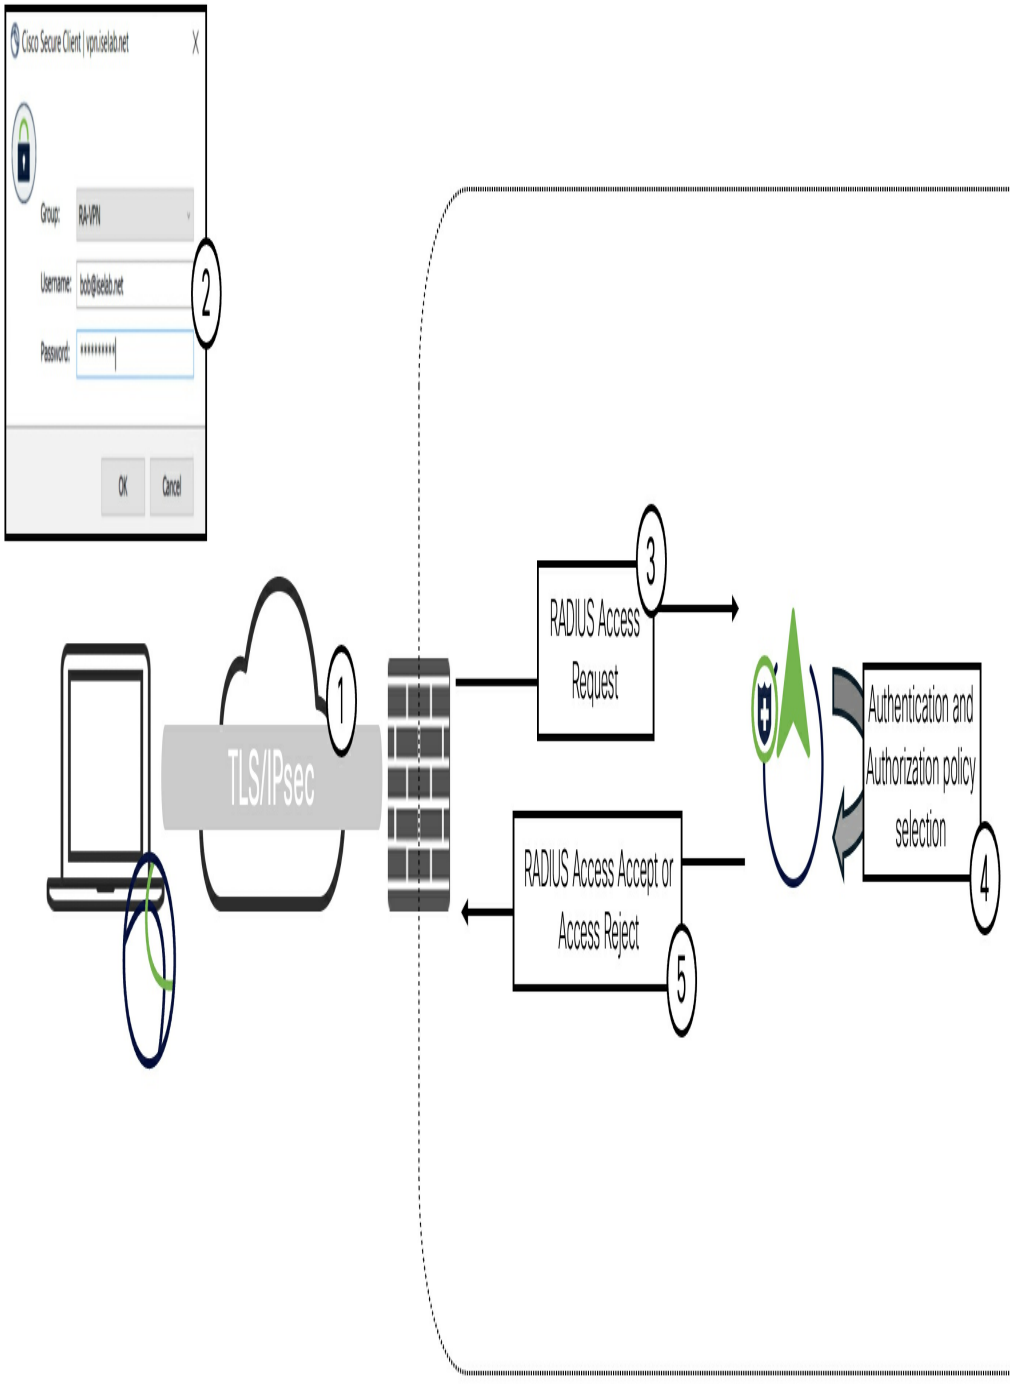

Cisco Secure Client | vpn.iselab.net
Group:
RA-VPN
Username:
bob@iselab.net
Password:
OK
Cancel
2
1
TLS/IPsec
3
RADIUS Access Request
RADIUS Access Accept or Access Reject
5
Authentication and Authorization policy selection
4

Figure 13-3 *Basic Authentication Flow*

As part of the basic authentication flow, the following main steps occur:

Step 1. Cisco Secure Client attempts to establish a VPN connection.

Step 2. The user is prompted for credentials.

Step 3. Secure Firewall sends user credentials in the RADIUS Access-Request to ISE.

Step 4. ISE processes authentication and authorization policies to select the identity store for user authentication and picks an authorization policy to assign the proper access level.

Step 5. If authentication and authorization succeed, ISE returns the attributes that need to be assigned on the Secure Firewall side, such as DACL, Framed-IP address, or group policy, in the Access-Accept message. If authentication or authorization fails, ISE returns Access-Reject.

Note

Basic authentication is considered insecure and should be avoided, when possible, in favor of more secure methods, such as certificate-based authentication and Security Assertion Markup Language (SAML).

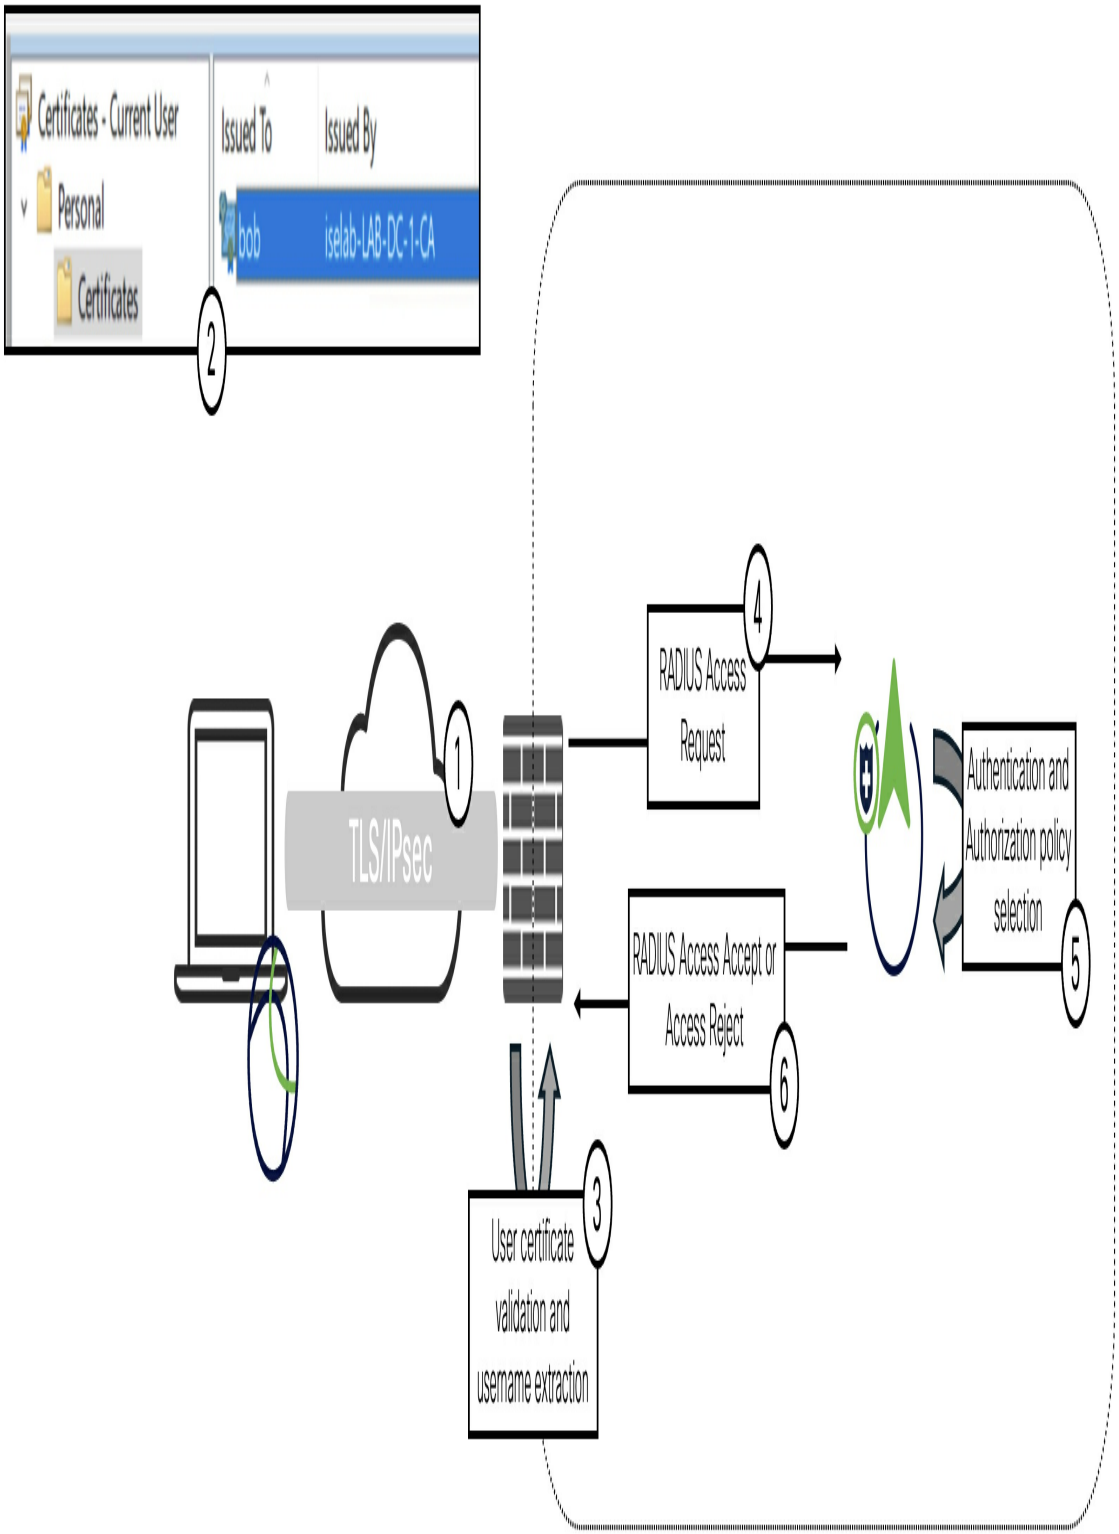
Certificate-Based Authentication with Authorization on ISE – Overview

In the use case employing certificate-based authentication with authorization on ISE, you authenticate the end user based on the identity certificate installed in the trusted store of the endpoint. Secure

Technet24

Firewall validates the certificate itself. Secure Firewall then sends an Authorize-Only RADIUS request to ISE for the username extracted from the certificate. ISE can make an authorization decision based on the user lookup. For an Authorize-Only request, authentication always fails on ISE, and this requires specific policy configuration that will be covered later in the "Certificate-Based Authentication with Authorization on ISE" section.

Figure 13-4 shows the flow for this use case.


Certificates - Current User
Personal
Certificates
Issued To
Issued By
bob
iselab-LAB-DC-1-CA
2
TLS/IPsec
1
RADIUS Access Request
4
RADIUS Access Accept or Access Reject
6
Authentication and Authorization policy selection
5
User certificate validation and username extraction
3

Technet24

Figure 13-4 *Certificate-Based Authentication Flow*

During certificate-based authentication with authorization on ISE, the following main steps occur:

Step 1. Secure Client attempts to establish a VPN connection.

Step 2. Secure Firewall requests certificate-based authentication. This triggers a certificate lookup on the endpoint. Once the user certificate is detected, the endpoint submits it to Secure Firewall for verification.

Step 3. Secure Firewall authenticates the user's certificate and, if authentication succeeds, extracts a predefined field from the certificate and uses it as the username to send a RADIUS Access-Request to ISE.

On the Secure Firewall side, user certificate validation includes the following steps:

a. Certificate chain validation: Confirms that all certificates in the user's certificate chain are correct.

b. Trust validation: Verifies whether Secure Firewall trusts the issuer of the user's certificate by checking trustpoints on Secure Firewall.

c. Certificate validity verification: Confirms that the user's certificate is already valid or has not expired.

d. Revocation verification: Verifies the certificate revocation status when the corresponding trustpoint on Secure Firewall includes CRL or OCSP settings.

e. Private key presence on the endpoint: As part of the TLS handshake, Secure Firewall verifies that the client can sign with the private key associated with the user's

certificate.

f. **Extended key usage (EKU) verification:** Requires the user's certificate to have the Client Authentication EKU.

Step 4. Secure Firewall sends the RADIUS Access-Request to ISE with the username extracted from the certificate and the RADIUS Service-Type set to Authorize Only.

Step 5. ISE processes authentication and authorization policies. Authentication for Authorize-Only requests always fails, so you must configure an authentication policy that continues on authentication failure. This approach allows ISE to proceed with authorization policy selection, using user lookup based on the username from the access request.

Step 6. If ISE selects an authorization policy with the Access-Accept authorization profile, it returns the attributes from the profile that need to be assigned on the Secure Firewall side, such as DACL, Framed-IP address, or group policy, in the Access-Accept message. If authorization fails, ISE returns Access-Reject.

SAML-Based Authentication with Authorization on ISE – Overview

This SAML-based authentication use case closely resembles the previous one, with one key difference. Instead of certificates, users authenticate with the help of a SAML identity provider (IdP) preconfigured on Secure Firewall. In this scenario, Secure Firewall acts as the SAML service provider (SP). The user identity for the Authorize-Only request to ISE is extracted from the SAML assertion.

Figure 13-5 shows the flow for this use case.

Technet24

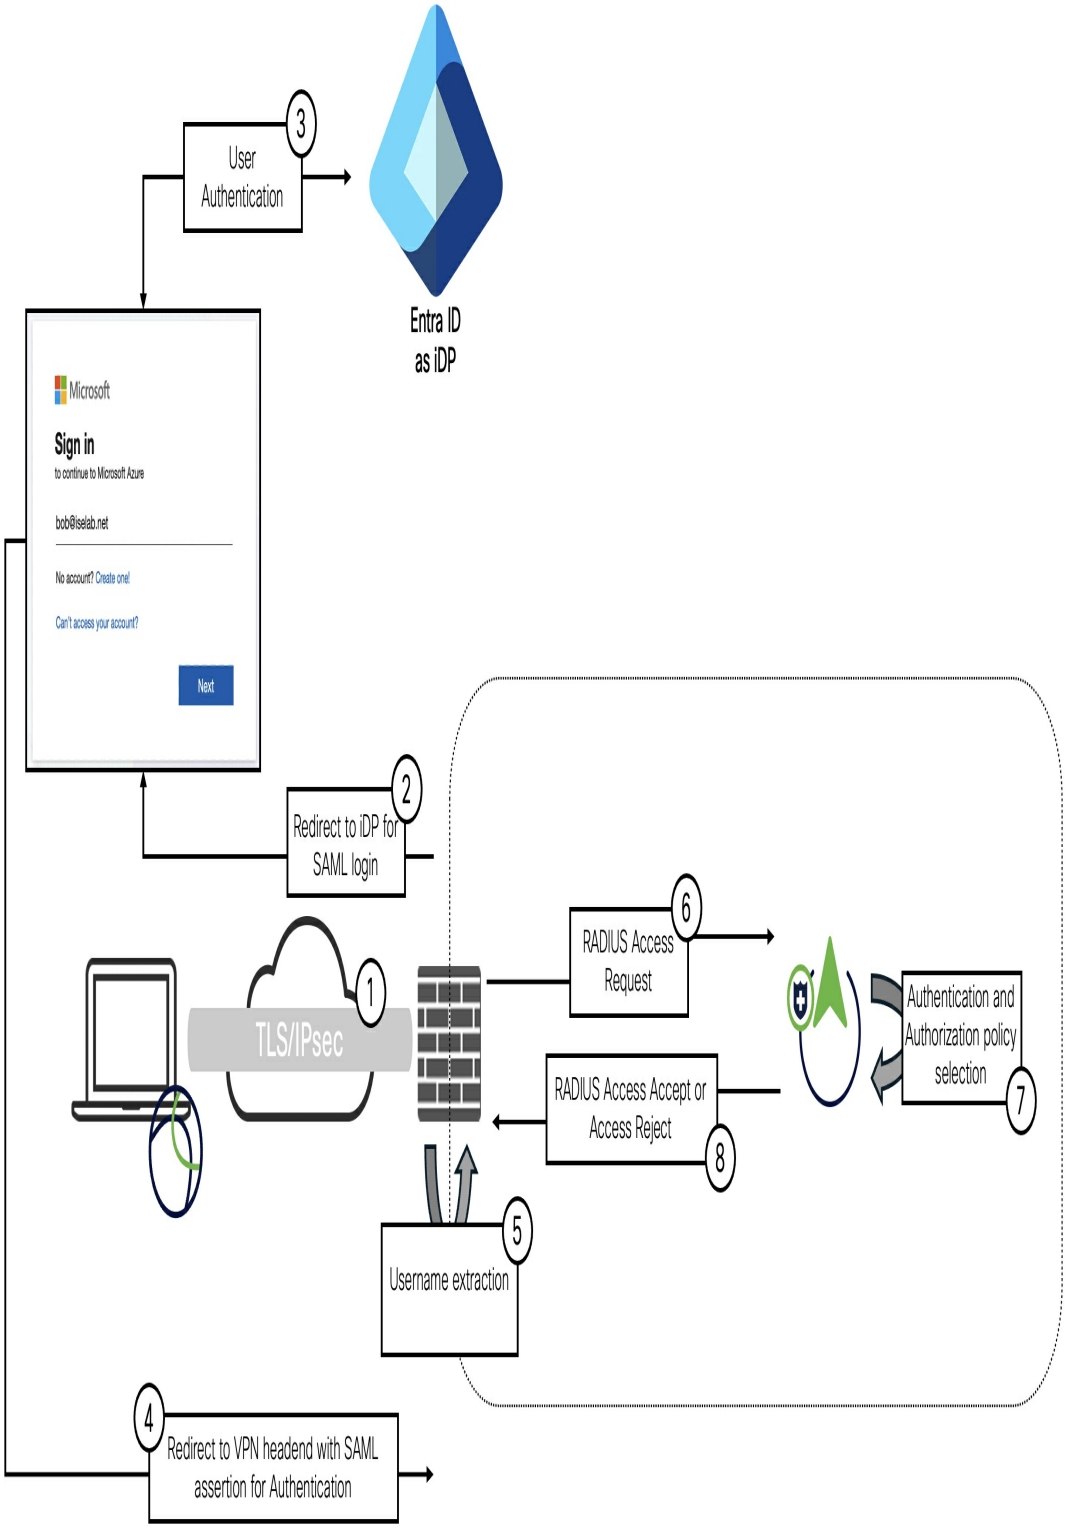

3
User Authentication
Entra ID as iDP
Microsoft
Sign in
to continue to Microsoft Azure
bob@iselab.net
No account? Create one!
Can't access your account?
Next
2
Redirect to iDP for SAML login
1
TLS/IPsec
6
RADIUS Access Request
Authentication and Authorization policy selection
7
RADIUS Access Accept or Access Reject
8
5
Username extraction
4
Redirect to VPN headend with SAML assertion for Authentication

Figure 13-5 *SAML Authentication Flow*

SAML-based authentication flow with authorization on ISE consists of the following main steps:

Step 1. Secure Client attempts to establish a VPN connection.

Step 2. Secure Firewall redirects Secure Client to the SSO URL of the SAML IdP configured in the tunnel group.

Step 3. The user authenticates with the IdP.

Step 4. If authentication succeeds, the IdP redirects Secure Client to the Secure Firewall SAML Assertion Consumer Service (ACS) URL with the assertion document.

Step 5. Secure Firewall extracts the username from the SAML assertion.

Step 6. Secure Firewall sends the RADIUS Access-Request to ISE with the username extracted from the SAML assertion and the RADIUS Service-Type set to Authorize Only.

Step 7. ISE processes authentication and authorization policies for Authorize Only requests in the same way as explained in the certificate-based authentication use case.

Step 8. If ISE selects an authorization policy with the Access-Accept authorization profile, it returns the attributes from the profile that need to be assigned on the Secure Firewall side, such as DACL, Framed-IP address, or group policy, in the Access-Accept message. If authorization fails, ISE returns Access-Reject.

Advanced Network Access Control with Secure Firewall

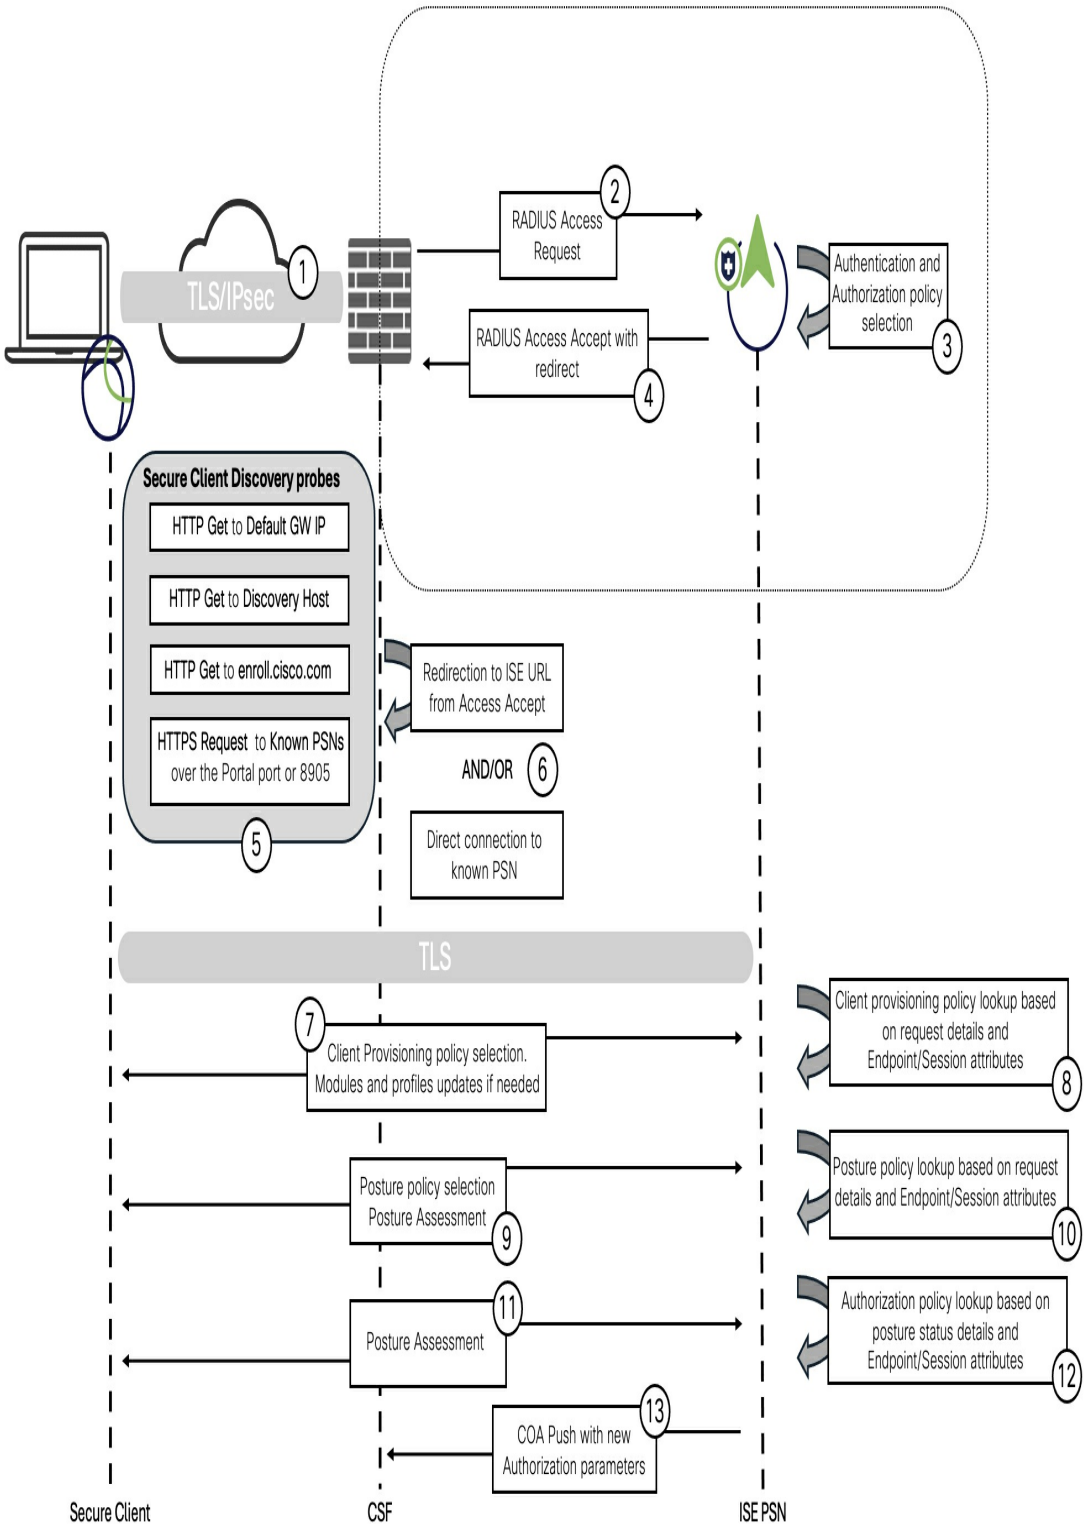
Technet24

This advanced network access control category includes two use cases.

Posture for the Remote Access VPN Users - Overview

The purpose of this posture use case is to ensure that endpoints connecting to the corporate network through Secure Firewall as a remote access VPN headend are properly assessed against the company's endpoint posture policy.

Figure 13-6 outlines the main steps of the posture flow for the remote access VPN connection.


TLS/IPsec
1
RADIUS Access Request
2
Authentication and Authorization policy selection
3
RADIUS Access Accept with redirect
4
Secure Client Discovery probes
HTTP Get to Default GW IP
HTTP Get to Discovery Host
HTTP Get to enroll.cisco.com
HTTPS Request to Known PSNs over the Portal port or 8905
5
Redirection to ISE URL from Access Accept
AND/OR
6
Direct connection to known PSN
TLS
7
Client Provisioning policy selection. Modules and profiles updates if needed
Client provisioning policy lookup based on request details and Endpoint/Session attributes
8
Posture policy selection Posture Assessment
9
Posture policy lookup based on request details and Endpoint/Session attributes
10
11
Posture Assessment
Authorization policy lookup based on posture status details and Endpoint/Session attributes
12
13
COA Push with new Authorization parameters
Secure Client
CSF
ISE PSN

Technet24

Figure 13-6 *Posture Flow for the Remote Access VPN*

While posture flows were generally outlined and discussed in Chapter 10, "Implementing Endpoint Posture Services," remote access VPN posture has its own specifics. The main steps of this flow are as follows:

Step 1. Secure Client attempts to establish a VPN connection.

Step 2. An authentication request is submitted to ISE.

Step 3. Authentication and authorization policy lookup occurs on ISE. For the posture process to begin, the authorization policy with the posture-status attribute set to Unknown or not equal to Compliant must be selected. This policy typically references the authorization profile with Web Redirection enabled and set to Client Provisioning (Posture) along with a redirection ACL.

Note

Although there are methods to perform posture assessment without redirection in the VPN use case, they are not recommended with Cisco network access devices. Redirect-based posture guarantees that posture requests always reach the appropriate PSN. If redirection is considered user-unfriendly, you should focus on minimizing its impact on the user rather than disabling it entirely. More details on this topic are provided in Chapter 10.

Step 4. Access-Accept or Access-Reject is returned to Secure Firewall. The following steps are executed only after successful authentication and authorization. In redirect-based posture, the Access-Accept typically contains the redirect URL and redirect ACL parameters.

Step 5. The ISE posture module in Secure Client initiates PSN discovery once the VPN connection is established.

Step 6. As a result of redirection or a direct connection to a previously known PSN, the posture process begins and the TLS tunnel is established directly between Secure Client and the ISE posture module.

Step 7. The ISE posture module requests a client provisioning policy to verify if any updates to profiles or modules are needed.

Step 8. The PSN runtime processes the request from the ISE posture module and uses information in the request, such as session ID, to locate the session in the session cache for this specific authentication. Data from the session, including user group membership and connection details, as well as information from the ISE posture module, are used to select the appropriate client provisioning policy. The policy details are returned to the agent, where the final decision is made regarding any necessary updates.

Note

In the VPN scenario, the Core VPN module of Secure Client cannot be updated from ISE and must be updated from the Secure Firewall side.

Step 9. Similar to Step 7, the agent sends a request to locate the posture policies that need to be verified.

Step 10. The PSN runtime performs posture policy selection in the same way as Step 8.

Step 11. The ISE posture module sends the posture report back to PSN after completing the local posture assessment.

Technet24

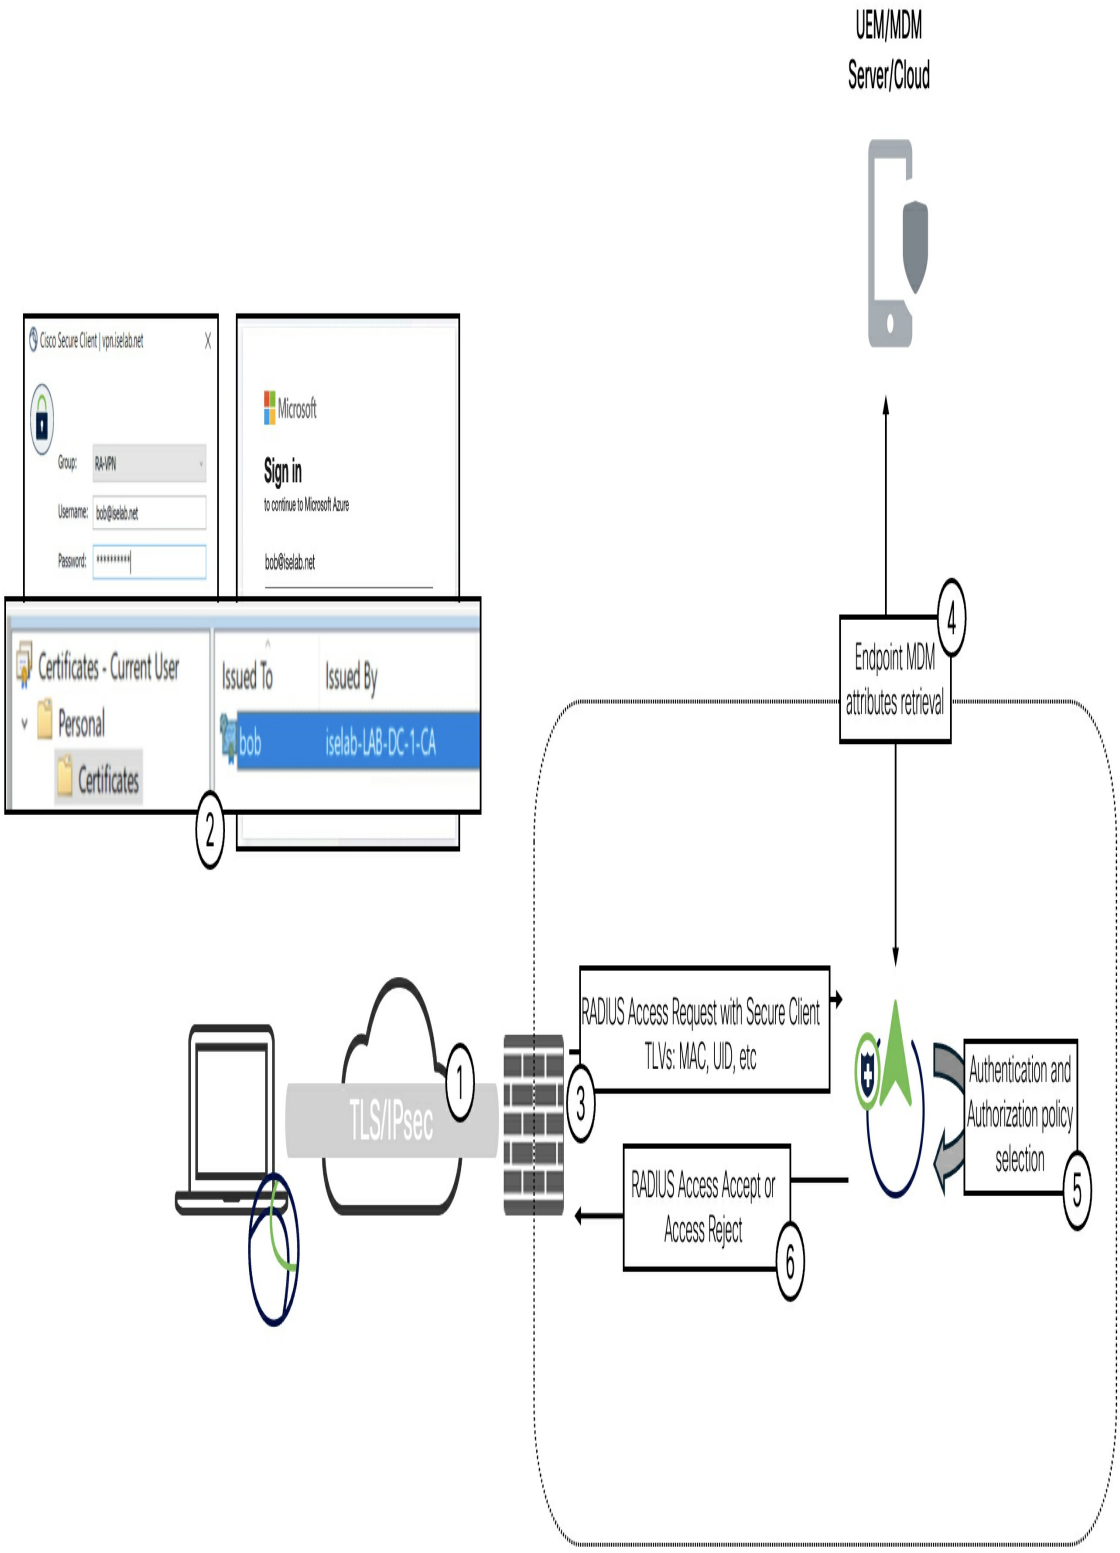
Step 12. Based on the posture module report, the PSN runtime selects the next authorization policy and extracts authorization attributes from the profile specified in the policy. Typically, after successful posture assessment, the authorization policy with posture status equal to Compliant is selected.

Step 13. PSN issues a COA-push that includes the attributes to be assigned to the session, typically a less-restrictive DACL. The COA-push enables assignment of a new access level without requiring reauthentication, since reauthentication in the VPN use case would disconnect the user.

MDM for Remote Access VPN Users

The next advanced use case involves verifying MDM attributes for the endpoint on the ISE PSN side before allowing the VPN connection. In this scenario, authorization policies must be configured to instruct PSN to request endpoint MDM attributes from the MDM server, such as MDM compliance status, MDM registration status, and similar properties.

Figure 13-7 shows the MDM flow for remote access VPN.


UEM/MDM
Server/Cloud
Cisco Secure Client | vpn.iselab.net
Group:
RA-IPN
Username:
bob@iselab.net
Password:
Microsoft
Sign in
to continue to Microsoft Azure
bob@iselab.net
Certificates - Current User
Personal
Certificates
Issued To
Issued By
bob
iselab-LAB-DC-1-CA
2
4
Endpoint MDM attributes retrieval
TLS/IPsec
1
3
RADIUS Access Request with Secure Client TLVs: MAC, UID, etc
Authentication and Authorization policy selection
5
RADIUS Access Accept or Access Reject
6

Technet24

Figure 13-7 *Remote Access VPN with MDM Verification*

Remote access VPN connection with MDM attribute retrieval and verification on the ISE PSN side involves the following main steps:

Step 1. Secure Client attempts to establish a VPN connection.

Step 2. Authentication of the user takes place using one of the supported methods.

Step 3. An authentication request is submitted to ISE. During processing, ISE PSN conducts initial authorization policy evaluation to determine whether it needs to query the MDM server and, if necessary, which MDM server to use. Once the server is identified, PSN submits the request to the server.

> **Note**
>
> In Figure 13-7, authentication and authorization are combined in Step 5 to simplify the flow. In the actual process, authentication always takes place before PSN selects the MDM server, and authorization is performed after PSN receives a response from the MDM server or after the request to the MDM server times out.

Step 4. The request is sent to the selected MDM server. PSN selects the MDM server based on endpoint attributes, authorization policy settings (MDM server name attribute), or by using the Query Multiple UEM/MDM feature when enabled. (Refer to Chapter 10 for more information on the MDM flow.)

Step 5. Previously collected attributes and attributes obtained from the MDM server are used to select the appropriate

authorization policy.

Step 6. If authentication and authorization succeed, ISE returns the attributes that need to be assigned on the Secure Firewall side, such as DACL, Framed-IP address, or group policy, in the Access-Accept message. If authentication or authorization fails, ISE returns Access-Reject.

Secure Firewall Remote Access VPN: General Configuration Elements

Before you begin implementing various NAC use cases for remote access VPN on Secure Firewall, you must complete some general configuration steps. This section covers the configuration elements required for different NAC use cases to operate. The following are preconfiguration tasks:

- **Secure Firewall identity certificate provisioning:** To handle TLS or IKEv2 remote access VPN requests, Secure Firewall must authenticate itself. This task is accomplished with a Secure Firewall identity certificate. You must complete certificate provisioning before performing any other configuration related to remote access VPN.
- **NAT exemptions and Access Control permissions:** For remote VPN clients to communicate directly with ISE, which is necessary in VPN posture use cases, you must exempt traffic between the clients and ISE from Network Address Translation (NAT) rules. Also, you must ensure that this communication is permitted in the Secure Firewall Access Control Policy (ACP).
- **Routing toward ISE servers:** As an administrator, you must verify that Secure Firewall has a valid data plane route to ISE servers so the RADIUS exchange can take place.

Technet24

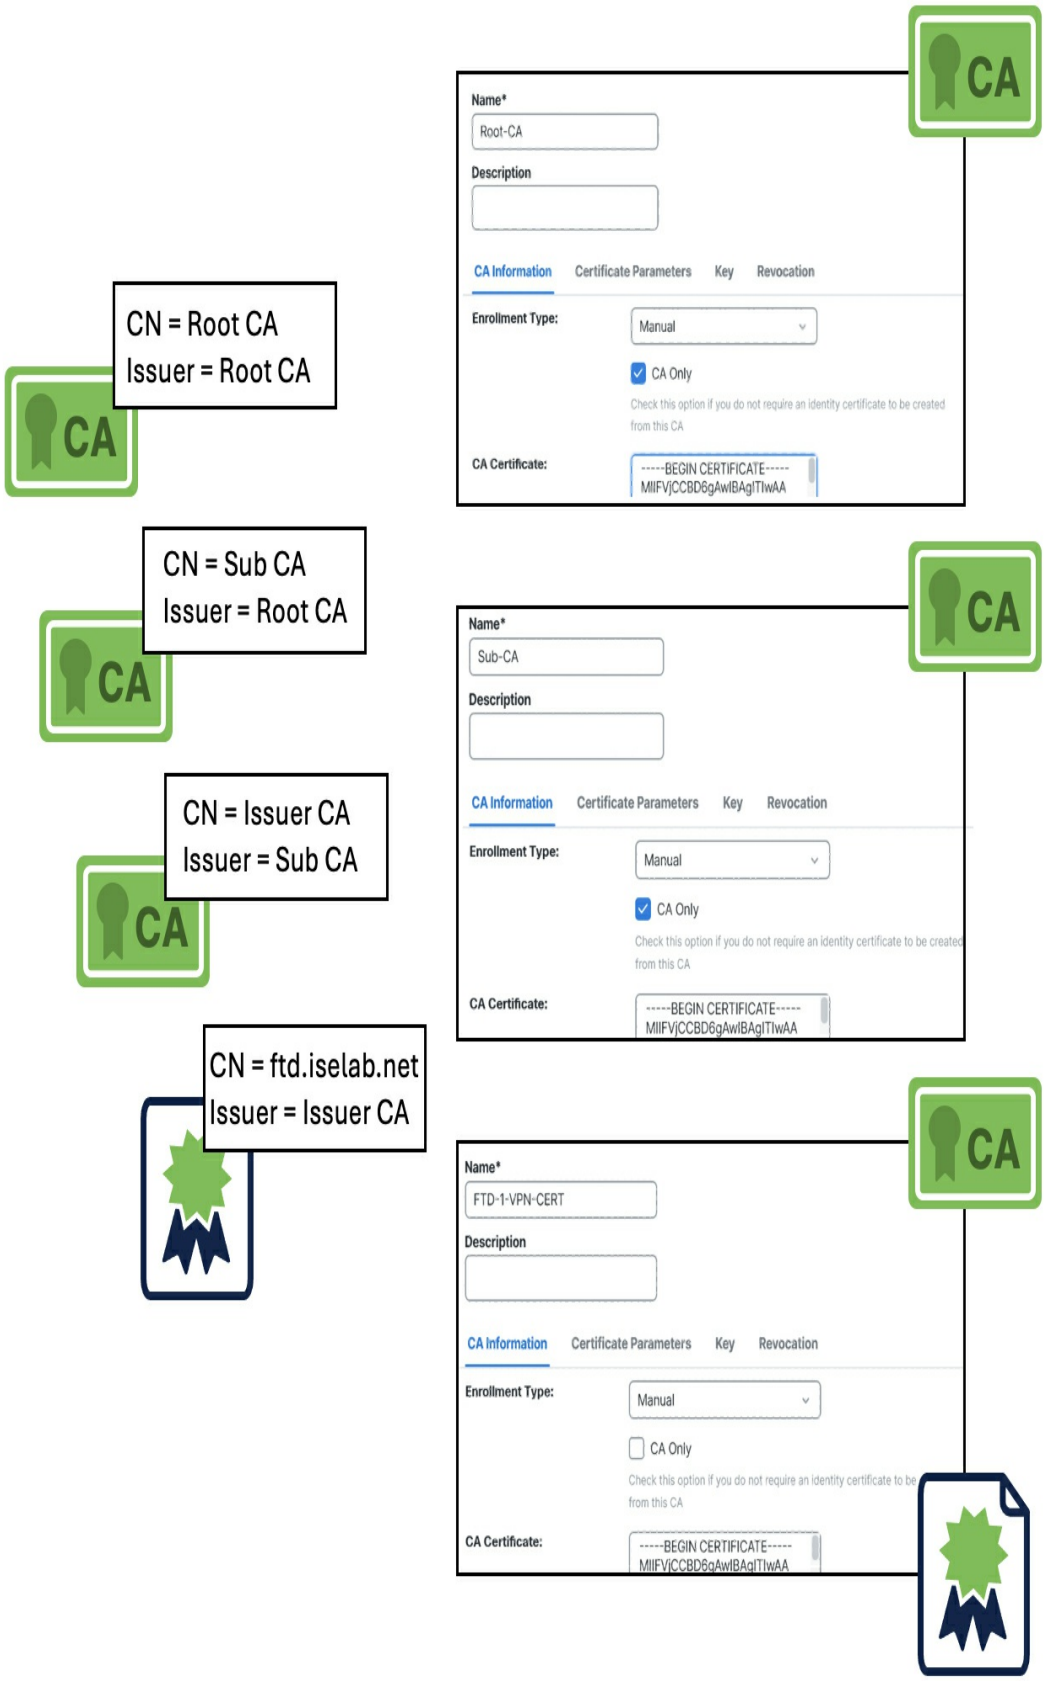
- **Secure Client image and profiles:** Before configuring remote access VPN on Secure Firewall, you must ensure the Secure Client image is uploaded to the Firewall Management Center. If Secure Client requires custom profiles, you can upload them to FMC at this step as well.

Secure Firewall Identity Certificate Provisioning

Several methods exist for provisioning the Secure Firewall identity certificate, but all of them require two main steps:

Step 1. Create one or more certificate enrollment objects. At this step, define objects that will be used as trust points on Secure Firewall. The number of objects depends on the depth of the certificate chain. When a Secure Firewall certificate is issued directly by the root certificate authority (CA), only one object is needed that contains both the identity and CA certificate. If the certificate chain is longer and includes sub-CAs, create an enrollment object for each CA in the chain. The lowest object in the CA hierarchy should include the Secure Firewall certificate and the CA certificate of its issuer. The remaining objects should be created as CA only.

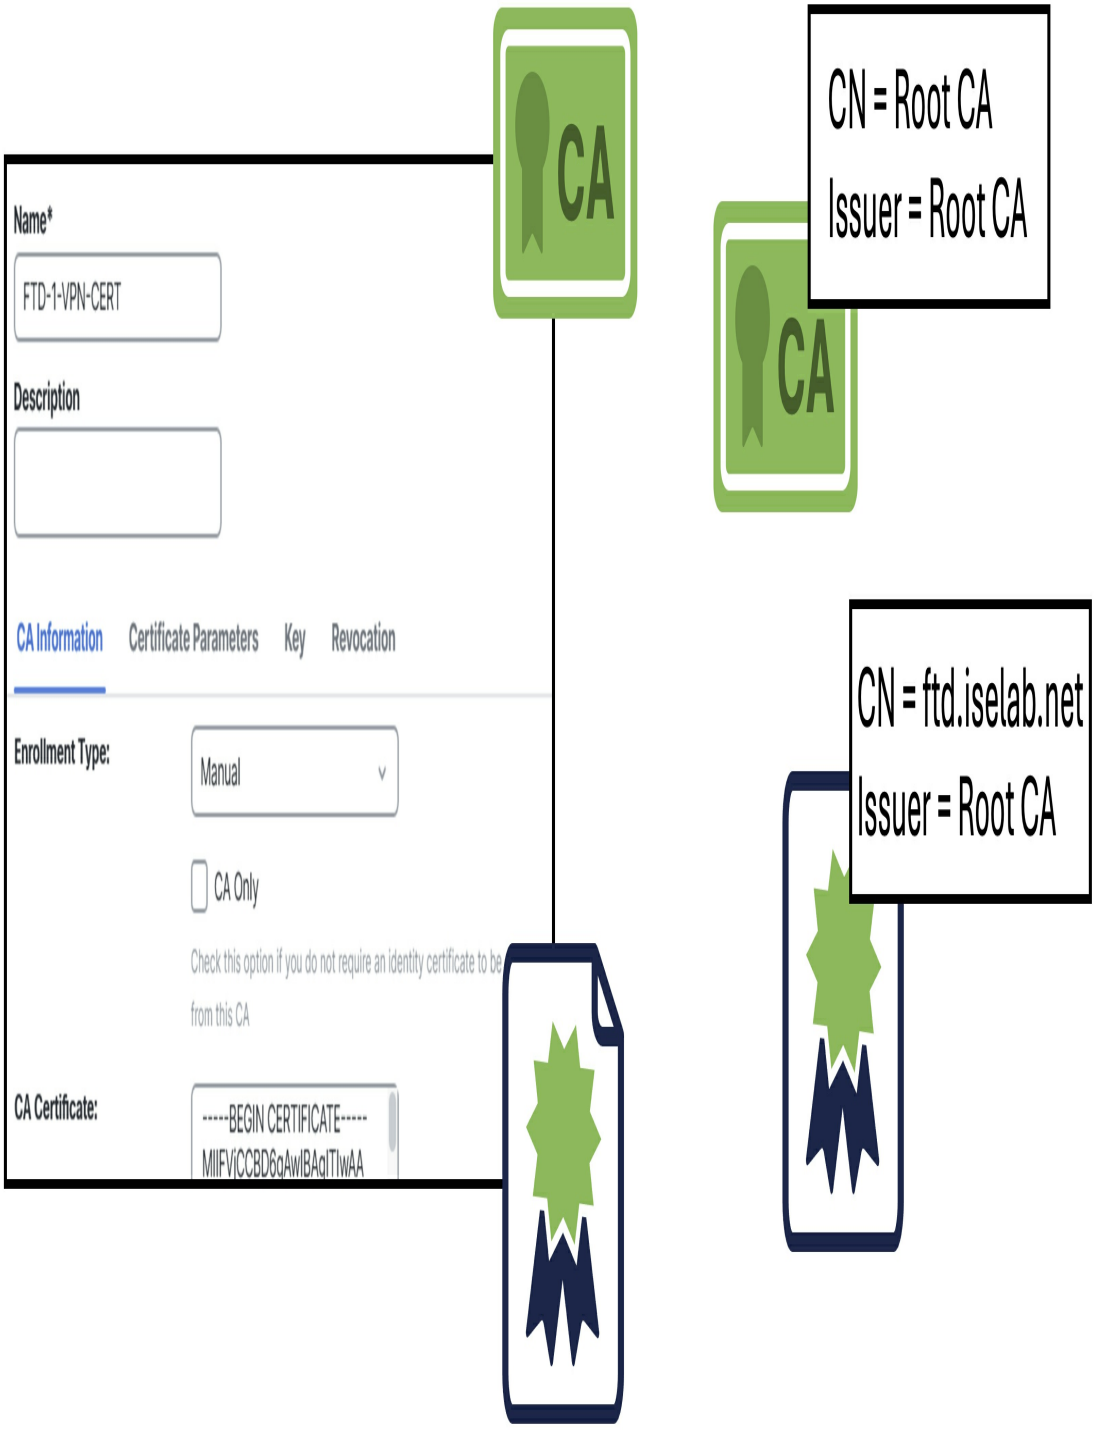
Figure 13-8 shows a certificate chain consisting of four levels with associated certificate enrollment objects.


CN = Root CA
Issuer = Root CA
Name*
Root-CA
Description
CA Information
Certificate Parameters
Key
Revocation
Enrollment Type:
Manual
CA Only
Check this option if you do not require an identity certificate to be created from this CA
CA Certificate:
-----BEGIN CERTIFICATE-----
MIIFVjCCBD6gAwIBAgITIwAA
CN = Sub CA
Issuer = Root CA
Name*
Sub-CA
Description
CA Information
Certificate Parameters
Key
Revocation
Enrollment Type:
Manual
CA Only
Check this option if you do not require an identity certificate to be created from this CA
CA Certificate:
-----BEGIN CERTIFICATE-----
MIIFVjCCBD6gAwIBAgITIwAA
CN = Issuer CA
Issuer = Sub CA
CN = ftd.iselab.net
Issuer = Issuer CA
Name*
FTD-1-VPN-CERT
Description
CA Information
Certificate Parameters
Key
Revocation
Enrollment Type:
Manual
CA Only
Check this option if you do not require an identity certificate to be created from this CA
CA Certificate:
-----BEGIN CERTIFICATE-----
MIIFVjCCBD6gAwIBAgITIwAA

Technet24

Figure 13-8 *Certificate Chain Consisting of Four Levels*

The example provided earlier is typical for publicly signed certificates, where the depth of the certificate chain is significant. If the Secure Firewall certificate is signed directly by the root CA, the certificate chain and associated enrollment object will appear as shown in Figure 13-9.


CA
CN = Root CA
Issuer = Root CA
CA
Name*
FTD-1-VPN-CERT
Description
CA Information
Certificate Parameters
Key
Revocation
Enrollment Type:
Manual
CA Only
Check this option if you do not require an identity certificate to be
from this CA
CA Certificate:
-----BEGIN CERTIFICATE-----
MIIFVjCCBD6gAwIBAgITIwAA
CN = ftd.iselab.net
Issuer = Root CA

Technet24

Figure 13-9 *Device Certificate Signed by Root CA*

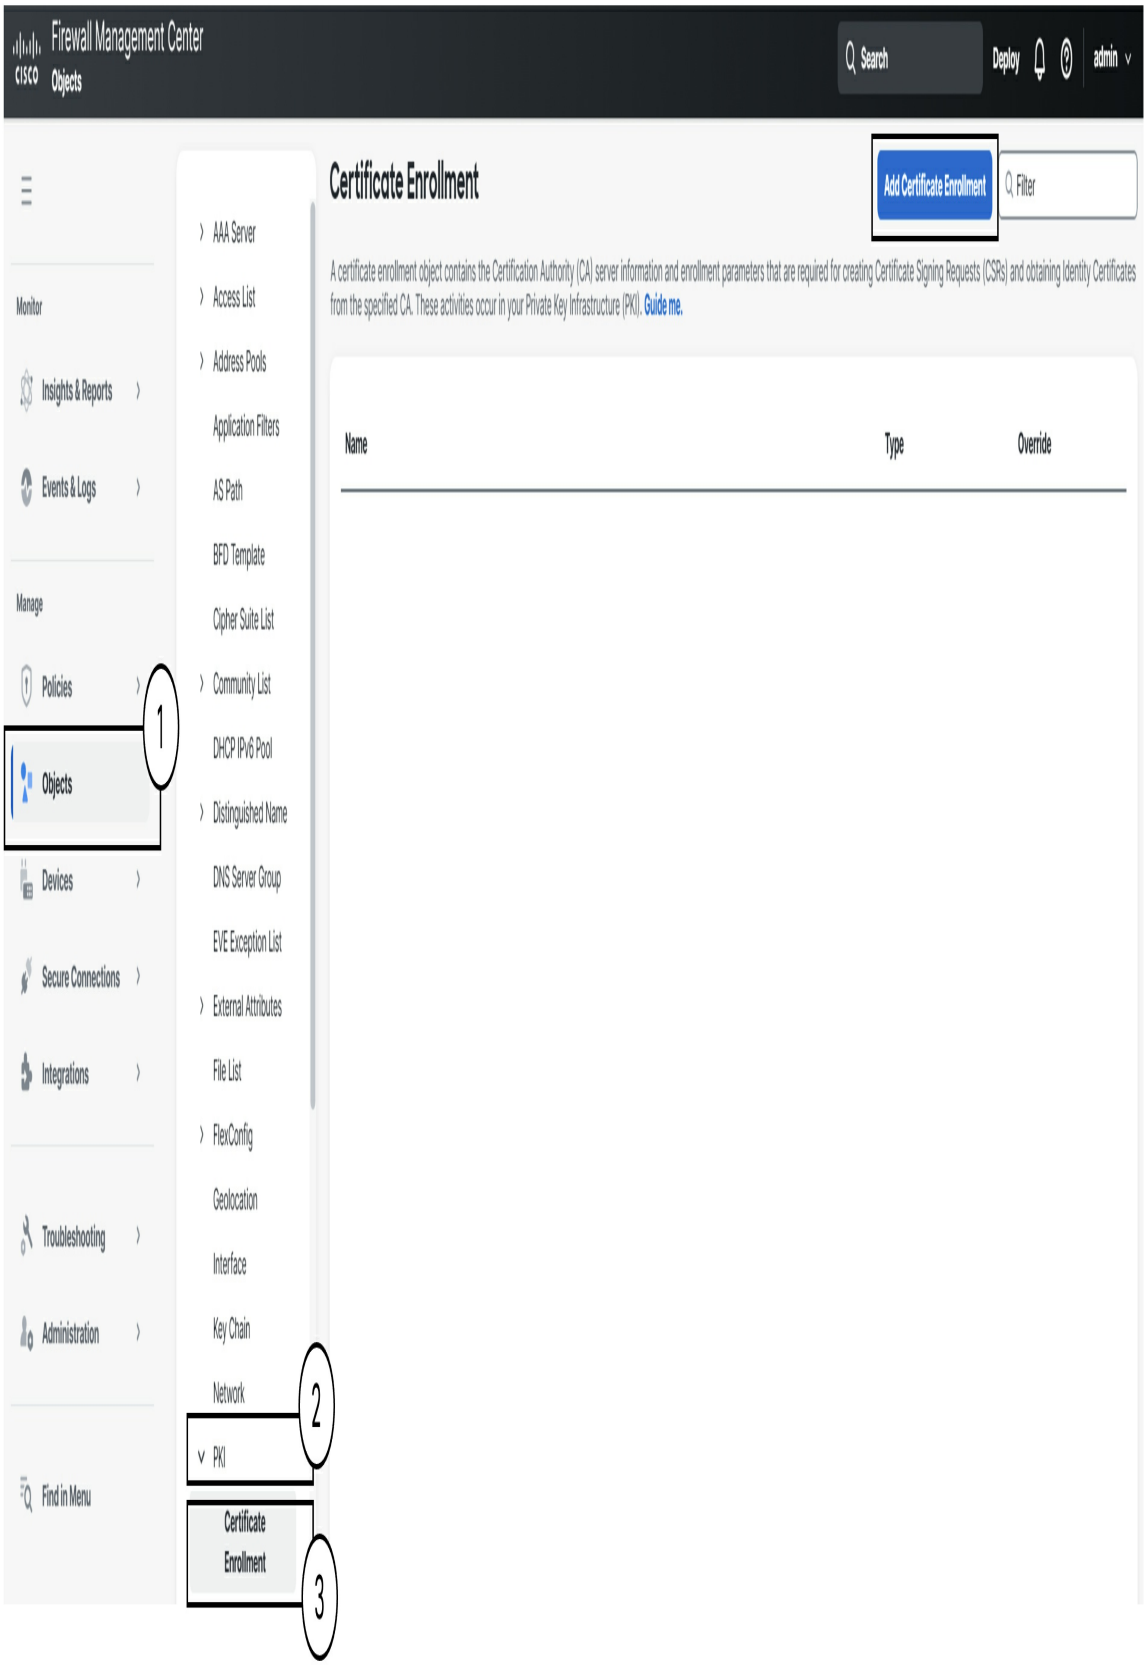
The earlier examples use the manual certificate enrollment method, but the same logic applies to any other approach.

To create certificate enrollment objects, refer to Figure 13-10 and use the following steps:

Step 1. Navigate to **Objects**.

Step 2. Click the **PKI** drop-down.

Step 3. Choose **Certificate Enrollment** in the FMC interface and click **Add Certificate Enrollment**.


Firewall Management Center
Objects
Search
Deploy
admin
Monitor
Insights & Reports
Events & Logs
Manage
Policies
Objects
Devices
Secure Connections
Integrations
Troubleshooting
Administration
Find in Menu
AAA Server
Access List
Address Pools
Application Filters
AS Path
BFD Template
Cipher Suite List
Community List
DHCP IPv6 Pool
Distinguished Name
DNS Server Group
EVE Exception List
External Attributes
File List
FlexConfig
Geolocation
Interface
Key Chain
Network
PKI
Certificate Enrollment
Certificate Enrollment
Add Certificate Enrollment
Filter
A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. These activities occur in your Private Key Infrastructure (PKI). Guide me.
Name
Type
Override
1
2
3

Technet24

Figure 13-10 *Certificate Enrollment Object*

Configuration of individual enrollment objects depends on the object type (CA + identity certificate or CA-only) and the enrollment method. CA-only enrollment objects require only a name and an imported CA certificate in PEM Base64 format. Example 13-1 provides a snippet of a PEM Base64 certificate.

Example 13-1 *PEM Base64 Certificate Snippet*

```
-----BEGIN CERTIFICATE-----
MIIDXTCCAkWgAwIBAgIJAN...[Base64 Encoded Data]...
-----END CERTIFICATE-----
```

When the object is for the identity certificate, you need to configure the following minimum settings:

- Certificate of the identity certificate issuer imported in PEM Base64 format.
- Certificate parameters such as email, company, and organizational unit, so the target Secure Firewall Appliance can construct the Common Name for the certificate.

These minimum settings apply to the manual enrollment method. Other methods may require different settings. For example, with PKCS12 (Public-Key Cryptography Standards) enrollment, the device certificate, private key, and CA certificate are supplied directly by the CA inside a password-protected PKCS12 file, so you only need to provide the certificate enrollment object name.

After you create all required certificate enrollment objects, assign them to the specific Secure Firewall appliance. After assignment, the subsequent behavior depends on the enrollment type:

- **Manual:** In this method, generate the CSR, sign it, and later bind

it to the pending certificate request on the Secure Firewall appliance.

- **Simple Certificate Enrollment Protocol (SCEP):** The Secure Firewall appliance uses SCEP settings in the enrollment object to request the certificate from the SCEP server.
- **Enrollment over Secure Transport (EST):** The Secure Firewall appliance uses EST settings in the enrollment object to request the certificate from the EST server.
- **Automated Certificate Management Environment (ACME):** The Secure Firewall appliance uses ACME settings in the enrollment object to request the certificate from the ACME server.
- **PKCS12:** The device identity certificate, private key, and CA certificate are installed on the Secure Firewall as is and no additional steps are needed.

To assign certificate enrollment objects in the FMC interface, navigate to **Devices > Certificates** and click **Add**, as shown in Figure 13-11.

Technet24

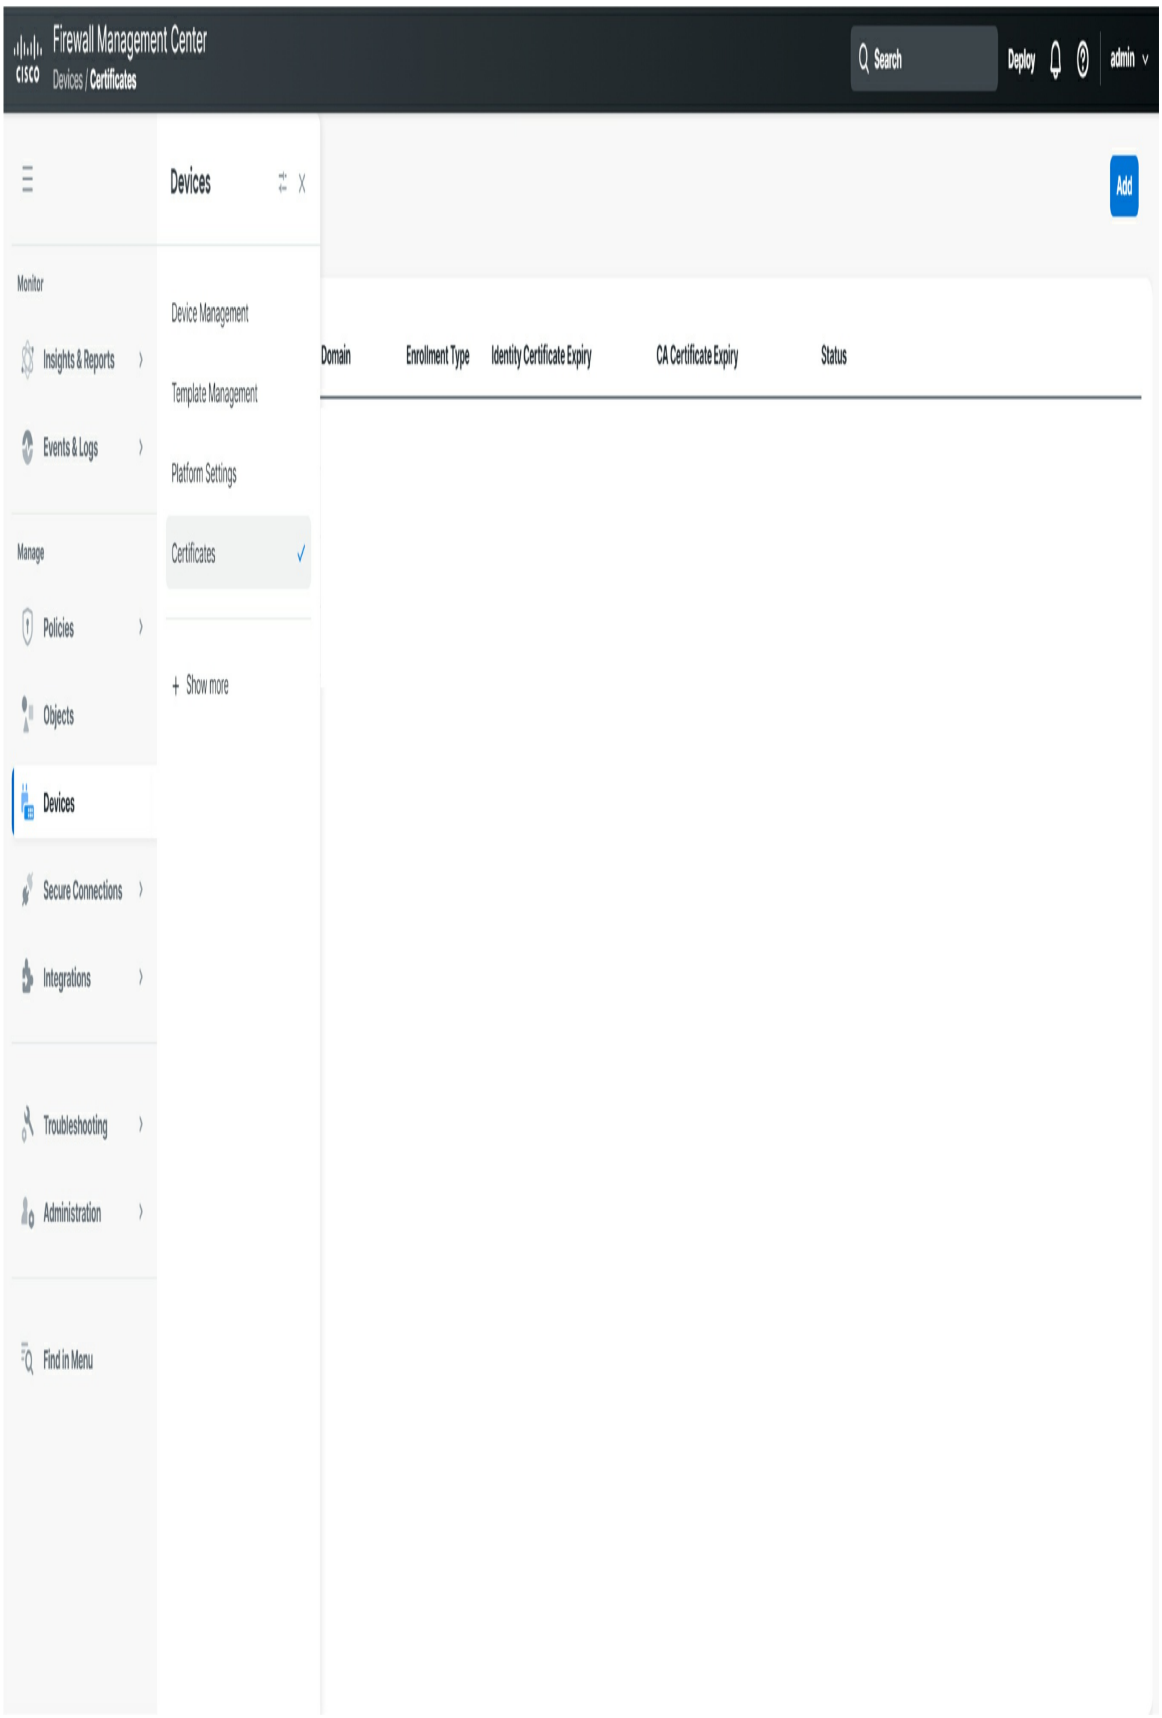

Firewall Management Center
Devices / Certificates
Search
Deploy
admin
Devices
Add
Monitor
Insights & Reports
Events & Logs
Manage
Policies
Objects
Devices
Secure Connections
Integrations
Troubleshooting
Administration
Find in Menu
Device Management
Template Management
Platform Settings
Certificates
+ Show more
Domain
Enrollment Type
Identity Certificate Expiry
CA Certificate Expiry
Status

Figure 13-11 *Adding a Certificate to the Device*

In the pop-up window that appears, select the device that requires assignment of the certificate enrollment object, choose the object from the list, and click **Add**, as shown in Figure 13-12.

Technet24

Add New Certificate

Add a new certificate to the device using cert enrollment object which is used to generate CA and identify certificate.

Device*:

KRK-FTD-1

Certificate Enrollment*:

VPN-CERT

Certificate Enrollment Details:

Name:	VPN-CERT
Enrollment Type:	Manual (CA & ID)
Enrollment URL:	N/A

Cancel Add

Figure 13-12 *Assigning an Enrollment Object*

Secure Firewall NAT Exception and Access Control

In most NAC use cases covered in this chapter, direct communication between the client and ISE is not required. However, one scenario necessitates connectivity: posture for remote access VPN. In this use case, Secure Client must establish a connection with ISE PSNs to perform client provisioning, posture policy lookup, and actual posture assessment. It is common for remote users to access internal company resources when their endpoints connect over remote access VPN. In the posture over VPN scenario, the ISE PSN operates as one of these resources.

When such connectivity must be permitted, Secure Firewall administrators typically need to ensure the following:

- The firewall policies allow connectivity from remote access VPN IP pools.
- If Secure Firewall performs NAT for remote access VPN IP pools, traffic between VPN users and internal resources is exempted from NAT.

Allowing Traffic from Remote Access VPN Users

As an administrator, you can choose between allowing traffic from remote access VPN IPs in the Access Control Policy (ACP), which offers greater flexibility and improved security but requires additional configuration, or bypassing ACP for client VPN traffic. Both options are explained in the subsequent sections.

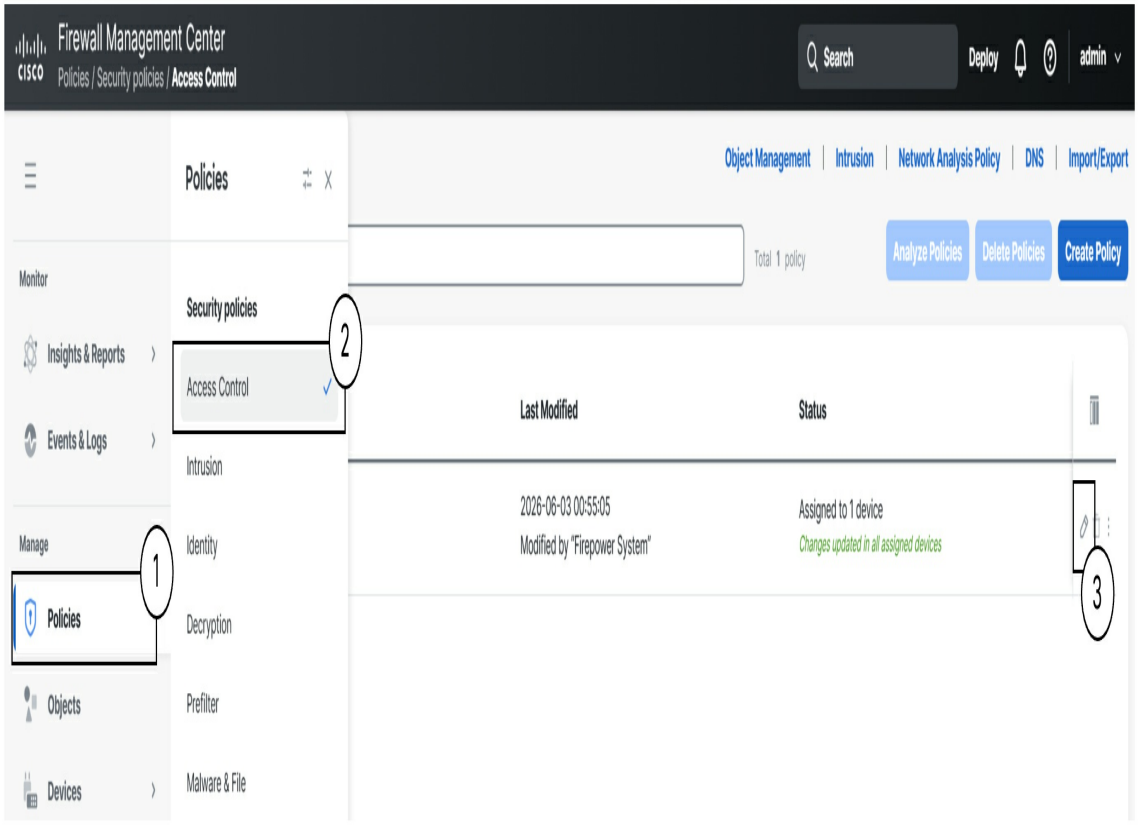
Allowing Traffic from Remote Access VPN IPs in the

Technet24

Access Control Policy

As an administrator, you permit connections from the remote access VPN IP pool as a source from the outside interface of Secure Firewall to any inside interfaces connected to required resources. This is the default approach. Using ACP provides the ability to build granular access policies, since ACP enables not only L3/L4 firewalling but also L7 access rules, including user identity, AppID, and TLS decryption. However, this approach requires explicit permission for traffic from remote access VPN IP pools on the outside interface of the firewall. Since this approach may introduce additional risks due to source IP address spoofing, it is recommended to confirm with your ISP that filtering for RFC 1918 ranges is in place. If that is not the case, Cisco recommends enabling anti-spoofing (Unicast Reverse Path Forwarding, or uRPF).

The process of adding a rule for remote access VPN users to the ACP, along with an example of a simple rule that can be used, is shown in Figure 13-13.


Firewall Management Center
Policies / Security policies / Access Control
Search
Deploy
admin
Policies
Object Management
Intrusion
Network Analysis Policy
DNS
Import/Export
Total 1 policy
Analyze Policies
Delete Policies
Create Policy
Monitor
Insights & Reports
Events & Logs
Manage
Policies
Objects
Devices
Security policies
Access Control
Intrusion
Identity
Decryption
Prefilter
Malware & File
Last Modified
Status
2026-06-03 00:55:05
Modified by "Firepower System"
Assigned to 1 device
Changes updated in all assigned devices
1
2
3


Packets
Prefilter Rules
Decryption
Security Intelligence
Identity
Access Control
More
Assigned to 1 device
Type to search
Total 2 rules
Add Category
Add Rule
4
Name
Action
Source
Destination
Applications
Zones
Networks
Zones
Networks
Ports
Mandatory 2 rules (1 - 2)
1 New-Rule-#1-ALLOW-I...
Allow
Any
any
Any
any
Any
Any
2 New-Rule-#2-ALLOW
Allow
Outside
ISE-LAB-RAVPN...
Any
ISE-LAB-Inte...
Any
Any
Default (No rules)
There are no rules in this section. Add Rule or Add Category

Technet24

Figure 13-13 *Access Policy for Remote Access VPN*

You can add the new policy for remote access VPN users by following these steps:

Step 1. Navigate to **Policies**.

Step 2. Click **Access Control**.

Step 3. Locate the policy for the Secure Firewall where the ACP needs to be added and click the pencil icon to edit the existing ACP.

Step 4. On the new page, click **Add Rule**.

Figure 13-13 shows a simple policy that allows traffic from the remote access VPN IP pool to all internal subnets on all interfaces.

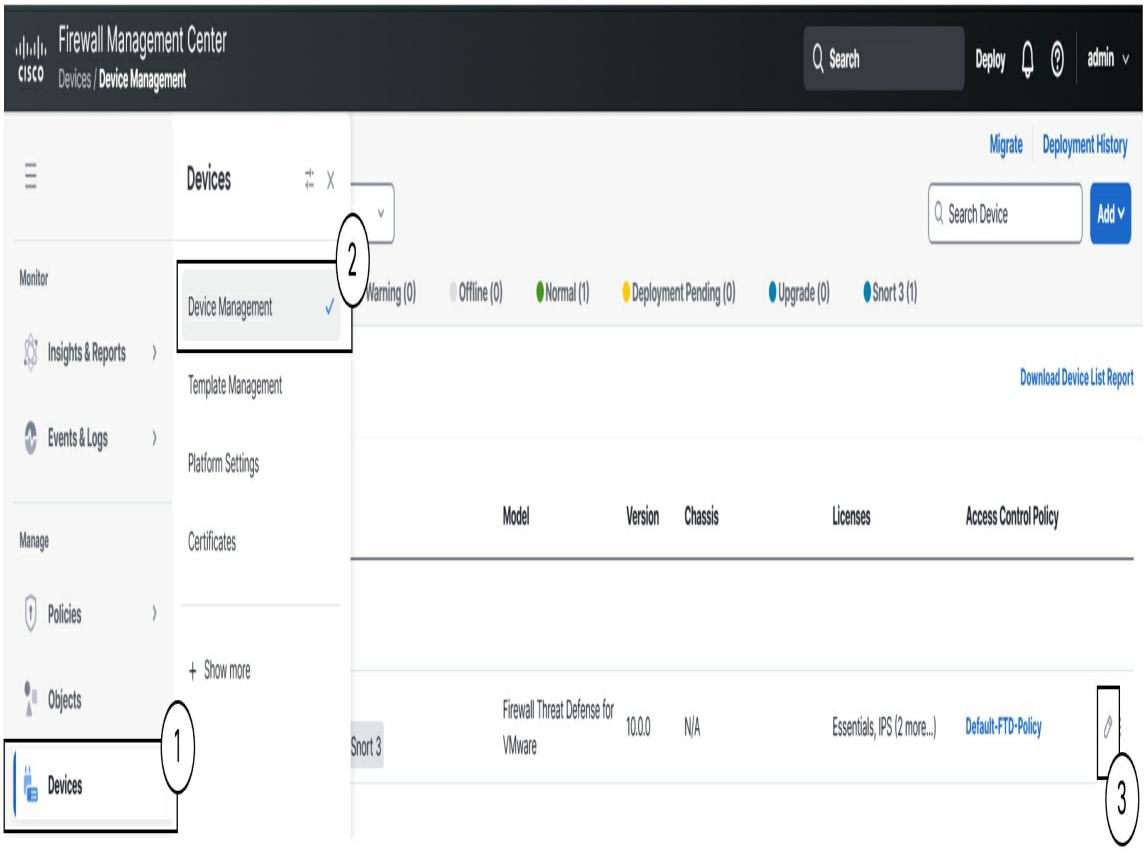
To additionally enable Anti-Spoofing protection, using Figure 13-14 as a reference, do the following:

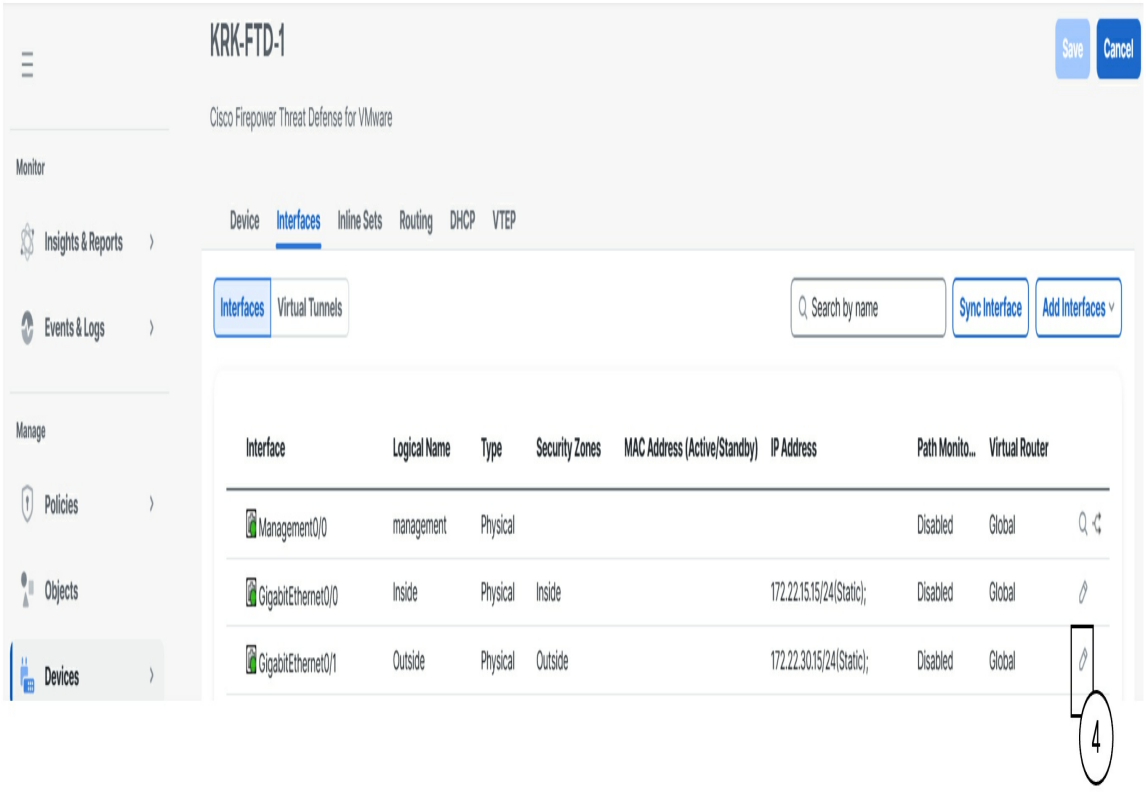
Step 1. Navigate to **Devices**.

Step 2. Click **Device Management**.

Step 3. Click the pencil icon for the Secure Firewall appliance that requires settings modification.

Step 4. On the new page, click the pencil icon next to the interface where remote access VPN is terminated.


Firewall Management Center
Devices / Device Management
Search
Deploy
admin
Devices
Migrate
Deployment History
Search Device
Add
Monitor
Insights & Reports
Events & Logs
Device Management
Template Management
Platform Settings
Certificates
+ Show more
Warning (0)
Offline (0)
Normal (1)
Deployment Pending (0)
Upgrade (0)
Snort 3 (1)
Download Device List Report
Model
Version
Chassis
Licenses
Access Control Policy
Manage
Policies
Objects
Devices
Snort 3
Firewall Threat Defense for VMware
10.0.0
N/A
Essentials, IPS (2 more...)
Default-FTD-Policy
1
2
3


KRK-FTD-1
Save
Cancel
Cisco Firepower Threat Defense for VMware
Monitor
Insights & Reports
Events & Logs
Manage
Policies
Objects
Devices
Device
Interfaces
Inline Sets
Routing
DHCP
VTEP
Interfaces
Virtual Tunnels
Search by name
Sync Interface
Add Interfaces
Interface
Logical Name
Type
Security Zones
MAC Address (Active/Standby)
IP Address
Path Monito...
Virtual Router
Management0/0
management
Physical
Disabled
Global
GigabitEthernet0/0
Inside
Physical
Inside
172.22.15.15/24(Static);
Disabled
Global
GigabitEthernet0/1
Outside
Physical
Outside
172.22.30.15/24(Static);
Disabled
Global
4

Technet24

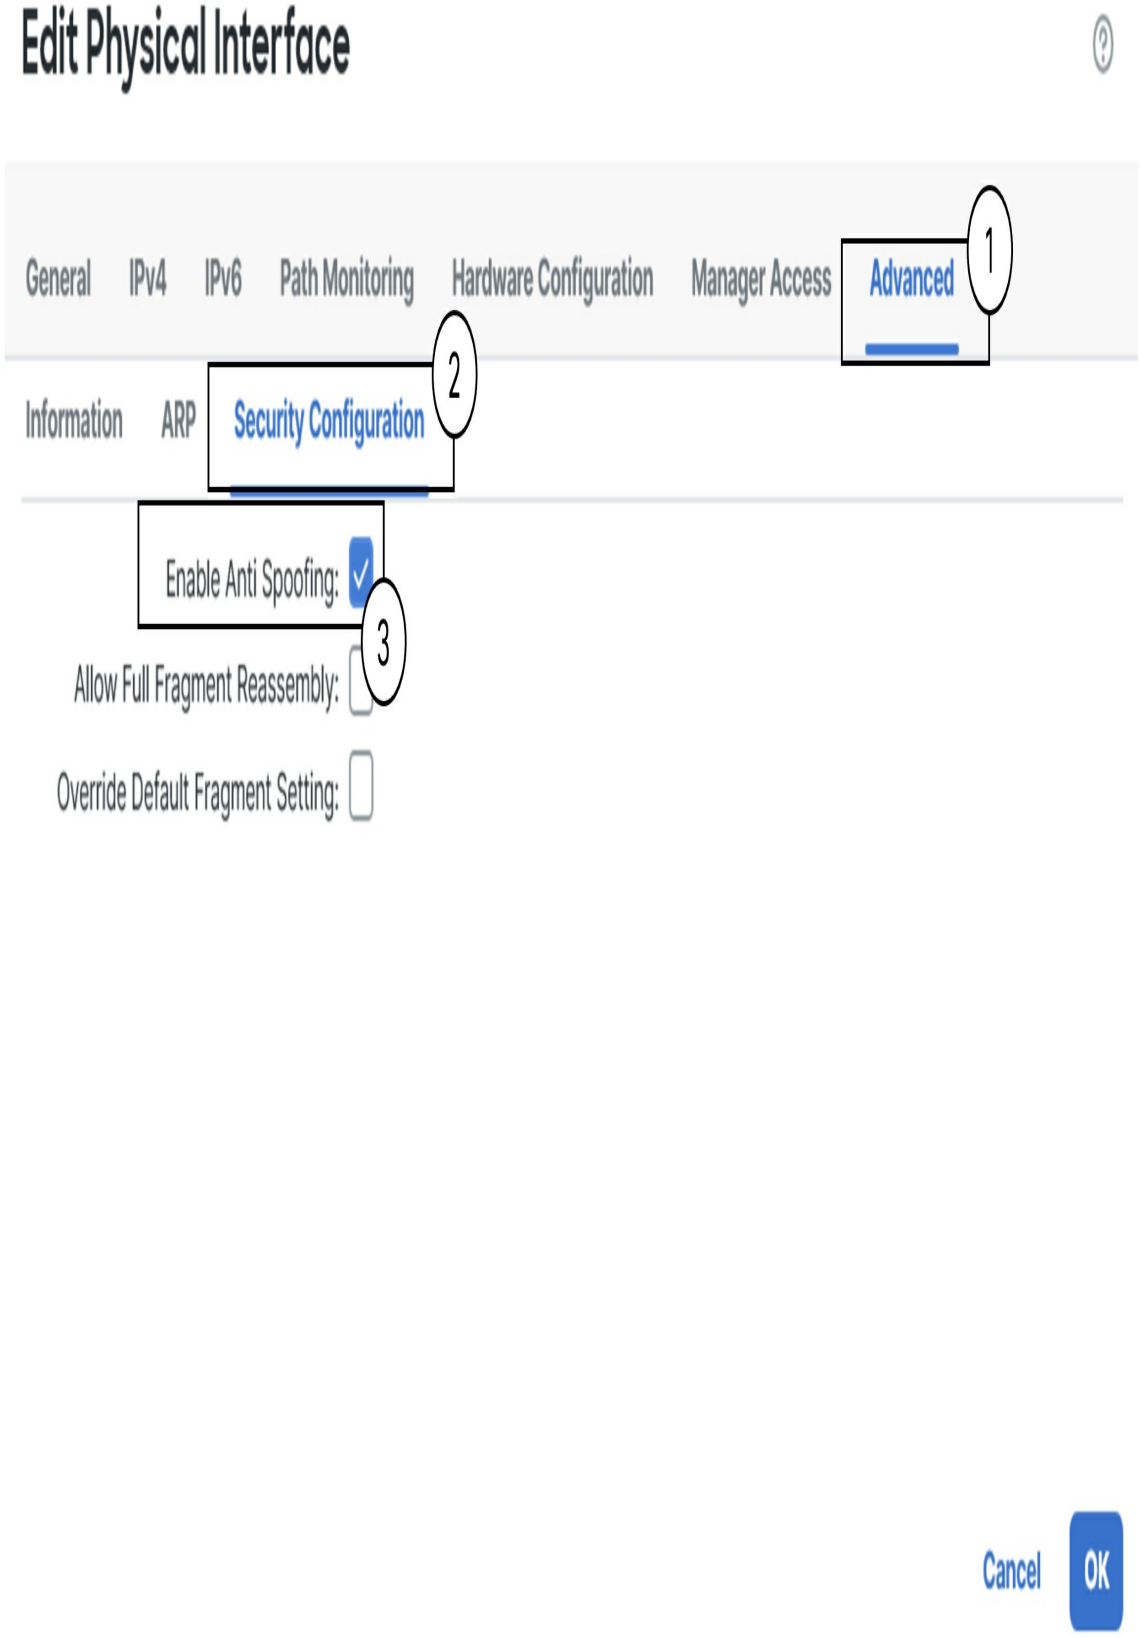
Figure 13-14 *Accessing Interface Settings*

An example of enabling anti-spoofing is shown in Figure 13-15. Follow these steps to complete the configuration:

Step 1. In the interface settings, go to the **Advanced** tab.

Step 2. Navigate to the **Security Configuration** tab when it opens.

Step 3. Click the **Enable Anti-Spoofing** checkbox.


Edit Physical Interface
General
IPv4
IPv6
Path Monitoring
Hardware Configuration
Manager Access
Advanced
1
Information
ARP
Security Configuration
2
Enable Anti Spoofing:
3
Allow Full Fragment Reassembly:
Override Default Fragment Setting:
Cancel
OK

Technet24

Figure 13-15 *Interface-Level Anti-Spoofing*

Bypassing ACP

You can bypass ACP by selecting **sysopt connection permit-vpn** at the end of the **Remote Access VPN** wizard. This action results in bypassing the ACP entirely for remote access VPN traffic. If this method is chosen, VPN group policies can filter traffic for remote access VPN users by specifying ACLs within the policies. With the ACL approach, filtering is limited to L3/L4 information. Figure 13-16 shows the Bypass ACP checkbox.



Figure 13-16 *ACP Bypass for RA VPN*

Excluding Remote Access VPN Traffic from NAT

Connectivity to internal resources also requires traffic from remote access VPN networks to be excluded from NAT when communication

Technet24

occurs with internal resources, including ISE PSNs.

The first question to consider is whether NAT exclusion is truly necessary. The answer depends on your network design. If Secure Firewall serves as the RA VPN headend and Internet access for users is handled elsewhere, such as on another Secure Firewall appliance, you might find that the RA VPN firewall does not have any NAT policies, so there is no need to exclude anything from NAT. However, a more common setup involves the RA VPN headend and Internet access control roles combined in the same Secure Firewall appliance. If this is the case, Secure Firewall is likely already configured with NAT policies, and it is important to ensure that traffic to internal resources from RA VPN users is excluded from NAT translation.

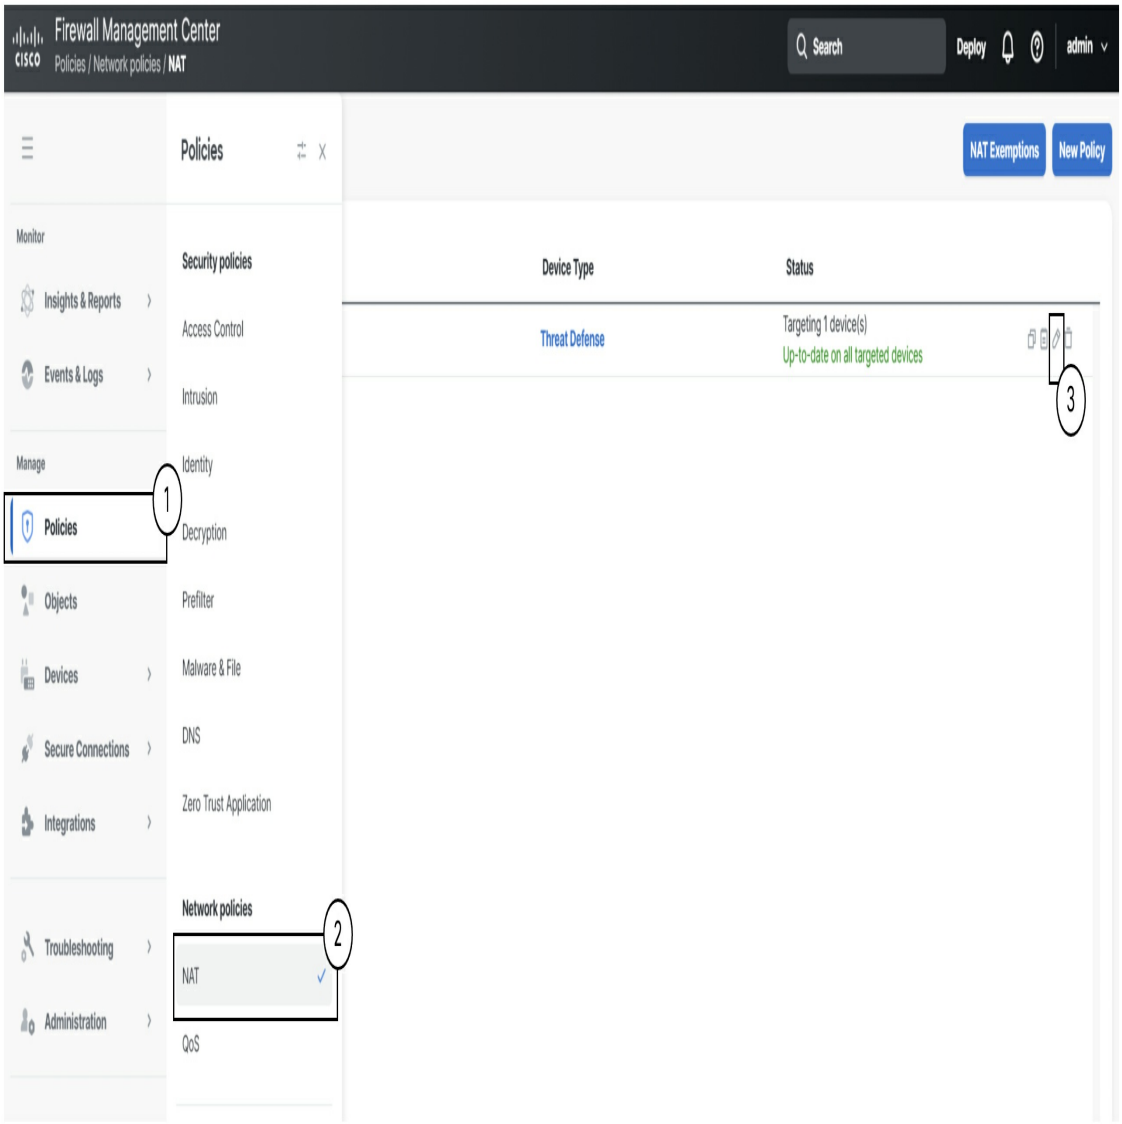
To add the NAT exemption, refer to Figure 13-17 and use the following steps:

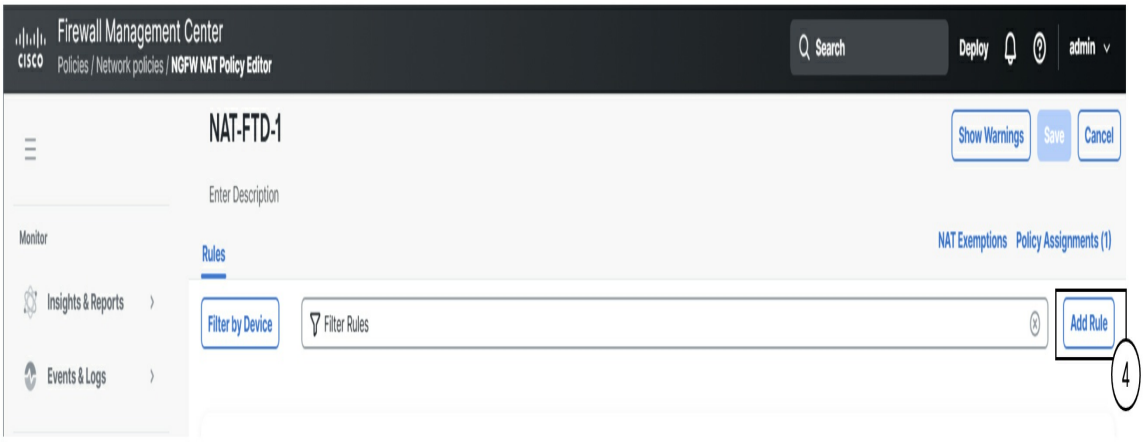
Step 1. Navigate to **Policies**.

Step 2. Click **NAT**.

Step 3. Locate the policy for the Secure Firewall where the exemption needs to be added and click the pencil icon.

Step 4. On the new page, click **Add Rule** to begin configuring the NAT exemption rule.

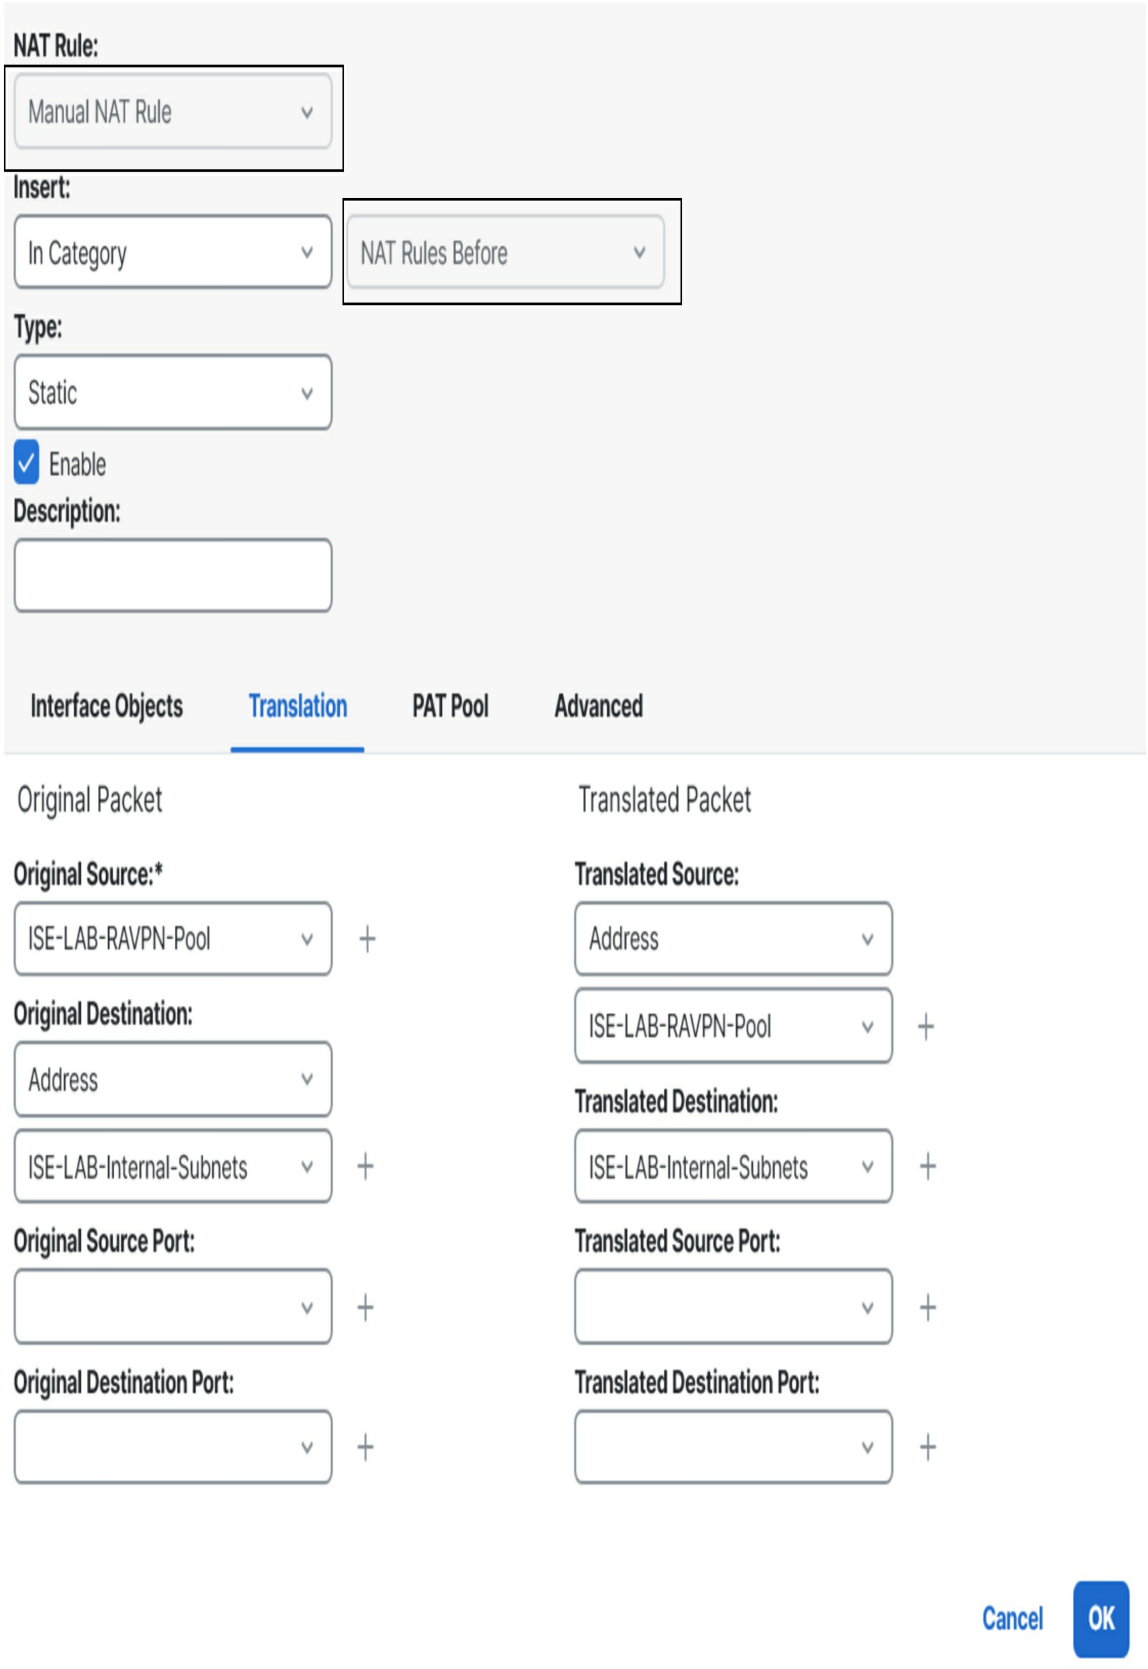

Firewall Management Center
Policies / Network policies / NAT
Search
Deploy
admin
Policies
NAT Exemptions
New Policy
Monitor
Insights & Reports
Events & Logs
Manage
Policies
Objects
Devices
Secure Connections
Integrations
Troubleshooting
Administration
Security policies
Access Control
Intrusion
Identity
Decryption
Prefilter
Malware & File
DNS
Zero Trust Application
Network policies
NAT
QoS
Device Type
Status
Threat Defense
Targeting 1 device(s)
Up-to-date on all targeted devices
1
2
3


Firewall Management Center
Policies / Network policies / NGFW NAT Policy Editor
Search
Deploy
admin
NAT-FTD-1
Show Warnings
Save
Cancel
Enter Description
Monitor
Insights & Reports
Events & Logs
Rules
NAT Exemptions
Policy Assignments (1)
Filter by Device
Filter Rules
Add Rule
4

Technet24

Figure 13-17 *Editing the NAT Policy*

A sample rule is demonstrated in Figure 13-18. Make sure the rule has a type of manual NAT, the used category is NAT Rules Before, and the translation section specifies the remote access VPN pool as the source and internal networks as the destination. You should configure the translation to translate the source and destination to themselves, which effectively excludes packets from NAT translation for the given source and destination IP pairs.


NAT Rule:
Manual NAT Rule
Insert:
In Category
NAT Rules Before
Type:
Static
Enable
Description:
Interface Objects
Translation
PAT Pool
Advanced
Original Packet
Original Source:*
ISE-LAB-RAVPN-Pool
Original Destination:
Address
ISE-LAB-Internal-Subnets
Original Source Port:
Original Destination Port:
Translated Packet
Translated Source:
Address
ISE-LAB-RAVPN-Pool
Translated Destination:
ISE-LAB-Internal-Subnets
Translated Source Port:
Translated Destination Port:
Cancel
OK

Technet24

Figure 13-18 *NAT Exemption Rule*

Secure Firewall Routing

Another important requirement for client communication with ISE PSNs is properly implemented routing on the Secure Firewall. Typically, routes to the ISE servers are added alongside routes to other internal resources that should be accessible by remote VPN clients.

Secure Firewall supports different types of routing technologies, each requiring varying levels of knowledge and experience to implement. As a starting point, you need to verify whether routing toward ISE servers is already present by executing the command **show route** in the diagnostic CLI, as shown in Figure 13-19.

```
krk-esc-ftd-1# show route 172.22.15.11


Routing entry for 172.22.15.0 255.255.255.0
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via Inside
      Route metric is 0, traffic share count is 1
```

Figure 13-19 *Secure Firewall* ***show route*** *Command*

To access the diagnostic CLI, log in to the Secure Firewall appliance over SSH and execute the command **system support diagnostic-cli**.

If no routes exist toward ISE servers, implement the routing technology of your choice by following these steps:

Step 1. Navigate to **Devices**.

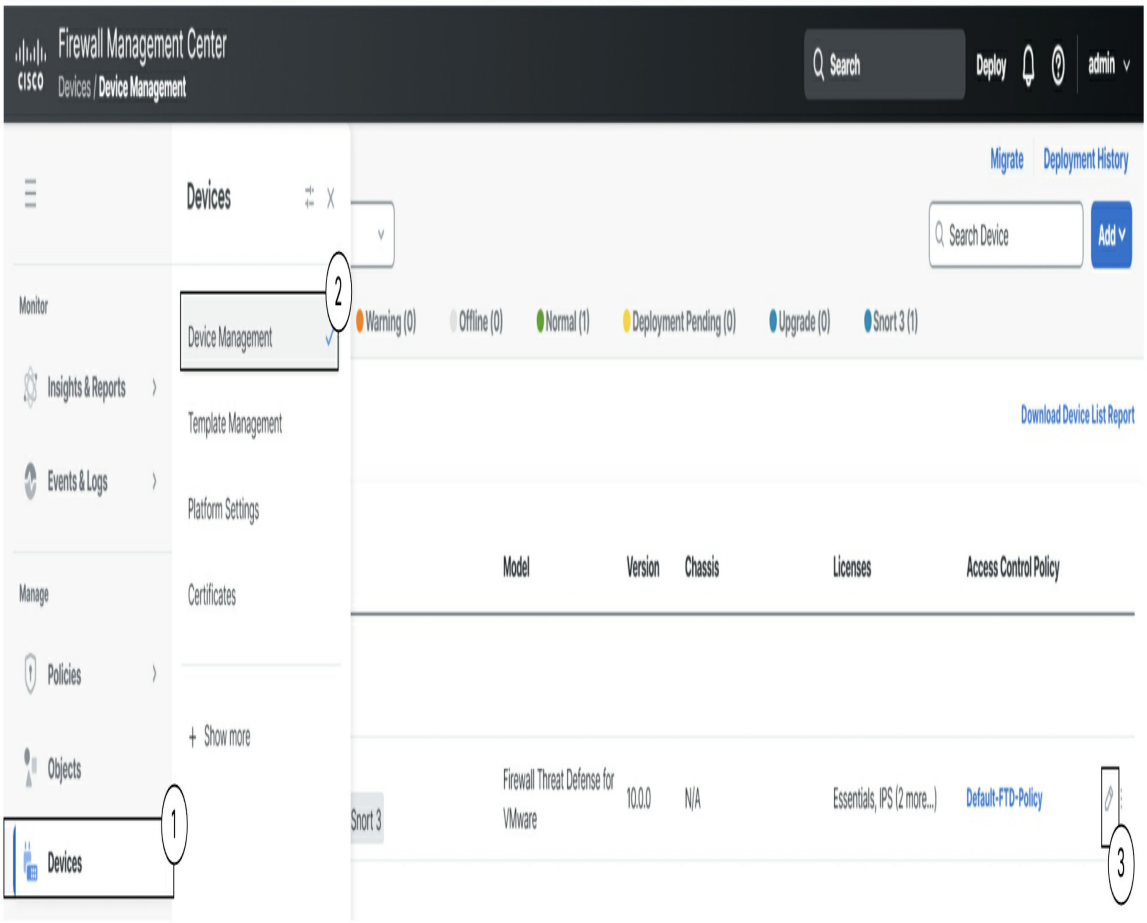
Step 2. Click **Device Management**.

Step 3. Click the pencil icon for the Secure Firewall appliance that requires settings modification.

Step 4. On the new page, go to the **Routing** tab and configure the appropriate routing.

Figure 13-20 shows an example.

Technet24


Firewall Management Center
Devices / Device Management
Search
Deploy
admin
Devices
Migrate
Deployment History
Search Device
Add
Monitor
Insights & Reports
Events & Logs
Device Management
Template Management
Platform Settings
Certificates
+ Show more
Warning (0)
Offline (0)
Normal (1)
Deployment Pending (0)
Upgrade (0)
Snort 3 (1)
Download Device List Report
Model
Version
Chassis
Licenses
Access Control Policy
Manage
Policies
Objects
Devices
Snort 3
Firewall Threat Defense for VMware
10.0.0
N/A
Essentials, IPS (2 more...)
Default-FTD-Policy
1
2
3

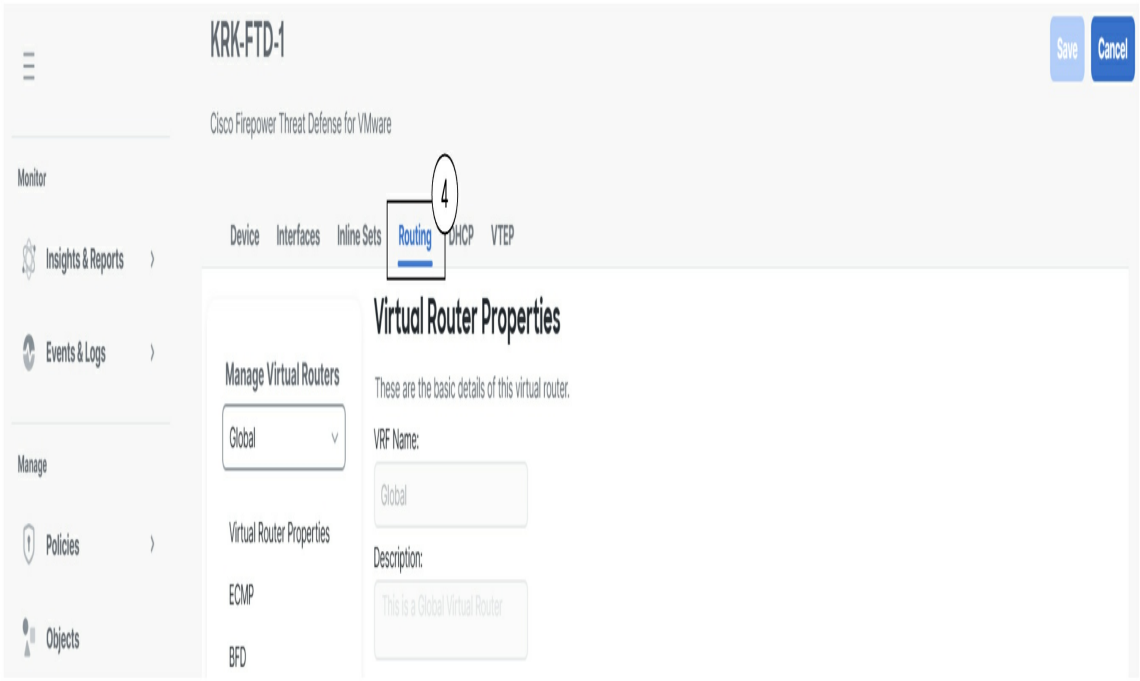

KRK-FTD-1
Save
Cancel
Cisco Firepower Threat Defense for VMware
Monitor
Insights & Reports
Events & Logs
Manage
Policies
Objects
Device
Interfaces
Inline Sets
Routing
DHCP
VTEP
4
Virtual Router Properties
Manage Virtual Routers
These are the basic details of this virtual router.
Global
VRF Name:
Global
Description:
This is a Global Virtual Router
Virtual Router Properties
ECMP
BFD

Figure 13-20 *Secure Firewall Routing Settings*

Secure Client Image and Profiles

For the remote access VPN to function, you need to download the secure client image from Cisco Software Central (software.cisco.com) and upload it to the FMC.

For Secure Firewall, the Headend Deployment package is required. This package is available for various operating systems, as shown in Figure 13-21.

Technet24

Cisco Secure Client Headend Deployment Package (Linux x86_64) cisco-secure-client-linux64-5.1.17.3394-webdeploy-k9.pkg Advisories	11-May-2026	175.38 MB	
Cisco Secure Client Pre-Deployment Package (Mac OS) - Administrator rights or managed device required for install or upgrade. See Administrator Guide and Release Notes for details. cisco-secure-client-macos-5.1.17.3394-predeploy-k9.dmg Advisories	11-May-2026	189.92 MB	
Application Programming Interface [API] (Mac OS) cisco-secure-client-macos-5.1.17.3394-vpnapi.tgz Advisories	11-May-2026	17.32 MB	
Cisco Secure Client Headend Deployment Package (Mac OS) - Administrator rights or managed device required for install or upgrade. See Administrator Guide and Release Notes for details. cisco-secure-client-macos-5.1.17.3394-webdeploy-k9.pkg Advisories	11-May-2026	187.43 MB	
Cisco Secure Client Pre-Deployment Package (Windows) - includes individual MSI files cisco-secure-client-win-5.1.17.3394-predeploy-k9.zip Advisories	11-May-2026	171.56 MB	
Application Programming Interface [API] (Windows) cisco-secure-client-win-5.1.17.3394-vpnapi.zip Advisories	11-May-2026	128.08 MB	
Cisco Secure Client Headend Deployment Package (Windows) cisco-secure-client-win-5.1.17.3394-webdeploy-k9.pkg Advisories	11-May-2026	152.85 MB	
Cisco Secure Client Pre-Deployment Package (Windows ARM64) - includes individual MSI files cisco-secure-client-win-arm64-5.1.17.3394-predeploy-k9.zip Advisories	11-May-2026	125.74 MB	
Cisco Secure Client Headend Deployment Package (Windows ARM64) cisco-secure-client-win-arm64-5.1.17.3394-webdeploy-k9.pkg Advisories	11-May-2026	143.40 MB	

Figure 13-21 *Secure Client Software*

After you obtain the client image, you can upload it to the FMC by referencing Figure 13-22 and using the following steps:

Step 1. Navigate to **Objects**.

Step 2. Click **VPN**.

Step 3. Click **Secure Client File**.

Step 4. Click **Add Secure Client File**.

Technet24

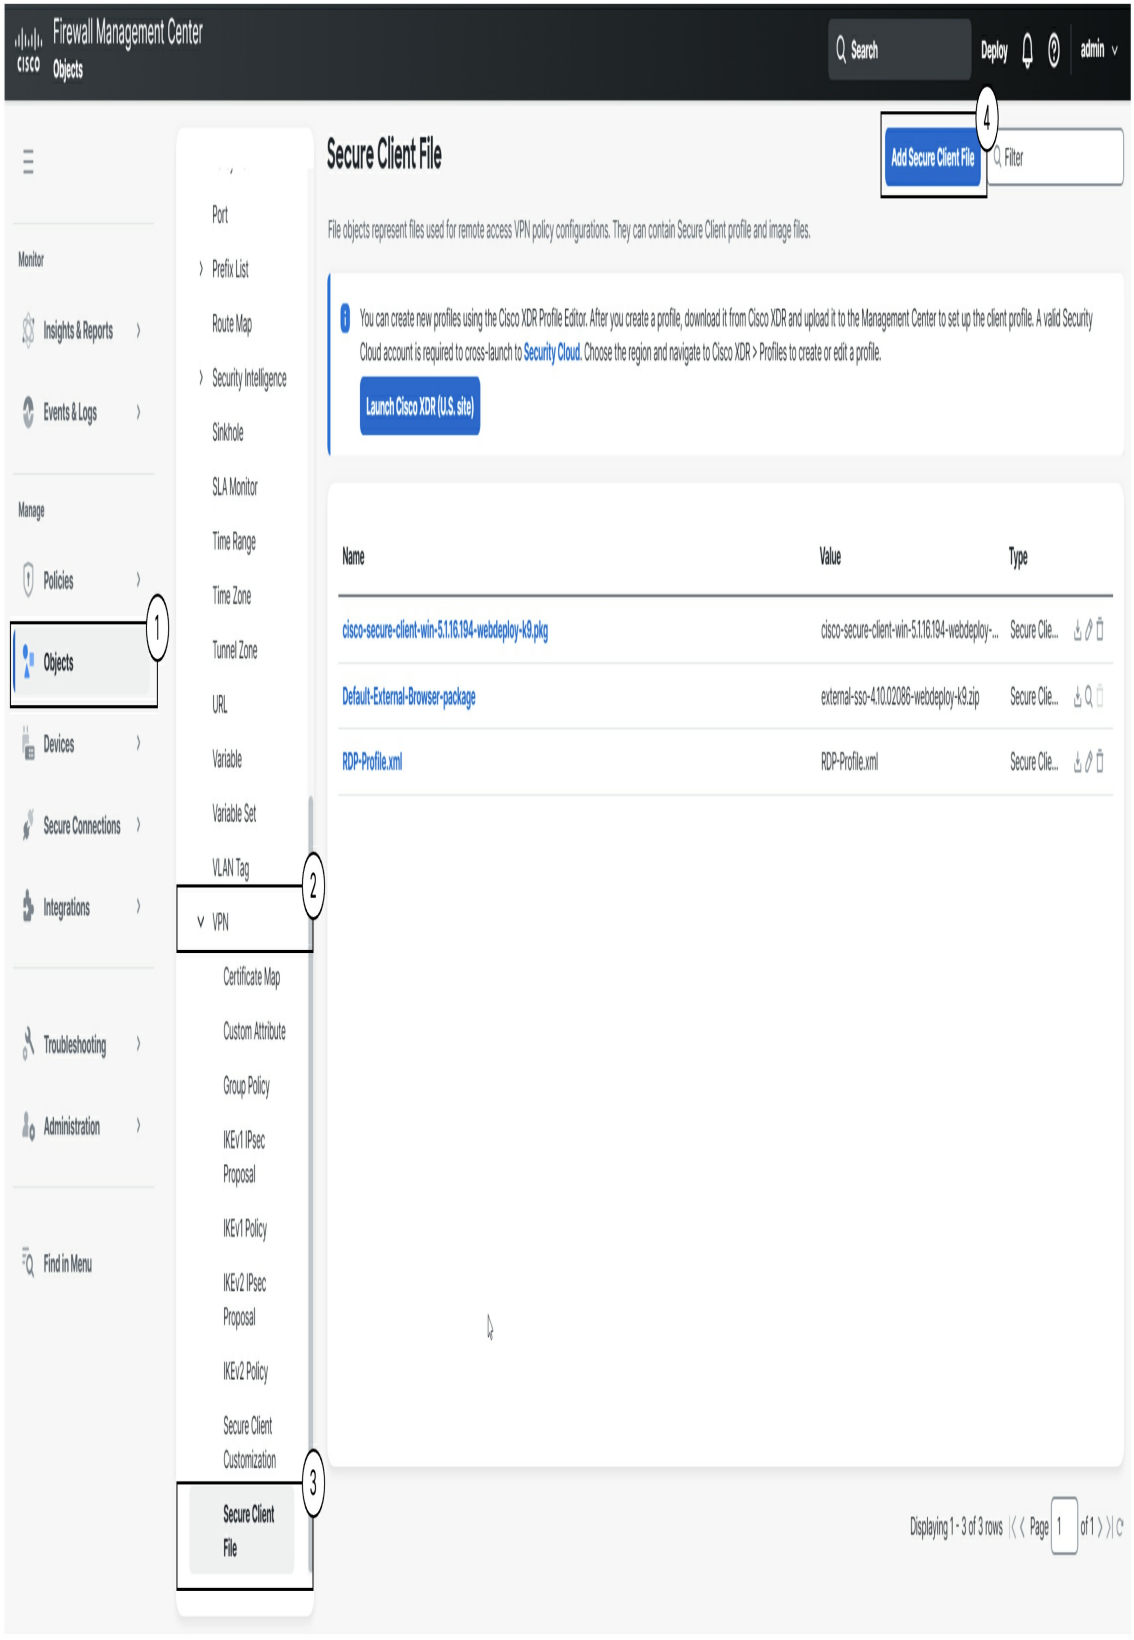

Firewall Management Center
Objects
Search
Deploy
admin
Monitor
Insights & Reports
Events & Logs
Manage
Policies
Objects
Devices
Secure Connections
Integrations
Troubleshooting
Administration
Find in Menu
1
Port
Prefix List
Route Map
Security Intelligence
Sinkhole
SLA Monitor
Time Range
Time Zone
Tunnel Zone
URL
Variable
Variable Set
VLAN Tag
VPN
2
Certificate Map
Custom Attribute
Group Policy
IKEv1 IPsec Proposal
IKEv1 Policy
IKEv2 IPsec Proposal
IKEv2 Policy
Secure Client Customization
Secure Client File
3
Secure Client File
4
Add Secure Client File
Filter
File objects represent files used for remote access VPN policy configurations. They can contain Secure Client profile and image files.
You can create new profiles using the Cisco XDR Profile Editor. After you create a profile, download it from Cisco XDR and upload it to the Management Center to set up the client profile. A valid Security Cloud account is required to cross-launch to Security Cloud. Choose the region and navigate to Cisco XDR > Profiles to create or edit a profile.
Launch Cisco XDR (U.S. site)
Name
Value
Type
cisco-secure-client-win-5.1.16.194-webdeploy-k9.pkg
cisco-secure-client-win-5.1.16.194-webdeploy-...
Secure Clie...
Default-External-Browser-package
external-sso-4.10.02086-webdeploy-k9.zip
Secure Clie...
RDP-Profile.xml
RDP-Profile.xml
Secure Clie...
Displaying 1 - 3 of 3 rows
Page 1 of 1

Figure 13-22 *Secure Client File Page*

Figure 13-23 shows how to upload the Secure Client image from the local machine.

Technet24

Add Secure Client File

Name:*

cisco-secure-client-win-5.1.16.194-we...

File Type:*

Secure Client Image

File Name:*

cisco-secure-client-win-5.1.16.194-we...

Browse..

Description:

Cancel

Save

Figure 13-23 *Secure Client Image Upload*

Optionally, you may choose to upload secure client profiles for various modules. You can perform this task on the same page, with the only distinction that, during the upload, you need to select the appropriate profile type (for example, Secure Client VPN profile) instead of Secure Client Image.

You can create profiles in two ways:

- Locally, on a Windows machine by using the Secure Client profile editor software.
- In Cisco XDR, by leveraging the Client Management section. Profiles created in this environment can be downloaded to your local machine and uploaded to FMC.

Secure Firewall Remote Access VPN: Following the RA VPN Wizard

The first step for remote access NAC configuration in Secure Firewall is the remote access VPN wizard available in FMC. The wizard requires you to enable an authentication method for remote access VPN. In the later example within the "RA VPN Wizard Step-by-Step" section, basic RADIUS authentication is used to complete the wizard. Most settings required for the wizard are specified during its flow. Important prerequisite settings not included in the wizard were discussed previously.

To access the wizard, refer to Figure 13-24 and use the following steps:

Step 1. In FMC, navigate to **Secure Connections**.

Step 2. Click **Remote Access VPN**.

Step 3. Click **Add**.

Technet24


Firewall Management Center
Secure Connections / Remote Access VPN
Search
Deploy
admin
Secure Connections
Add
Monitor
Site-to-Site VPN & SD-WAN
Status
Last Modified
Insights & Reports
Remote Access VPN
Events & Logs
Dynamic Access Policy
Manage
Policies
Objects
Devices
Secure Connections
Integrations
1
2
3

Figure 13-24 *Initiating the Remote Access VPN Wizard*

RA VPN Wizard Step-by-Step

At the first wizard page, you can define the basic settings, such as configuration name, the type of remote access VPN protocol, and other required parameters. Figure 13-25 shows this initial step.

Technet24

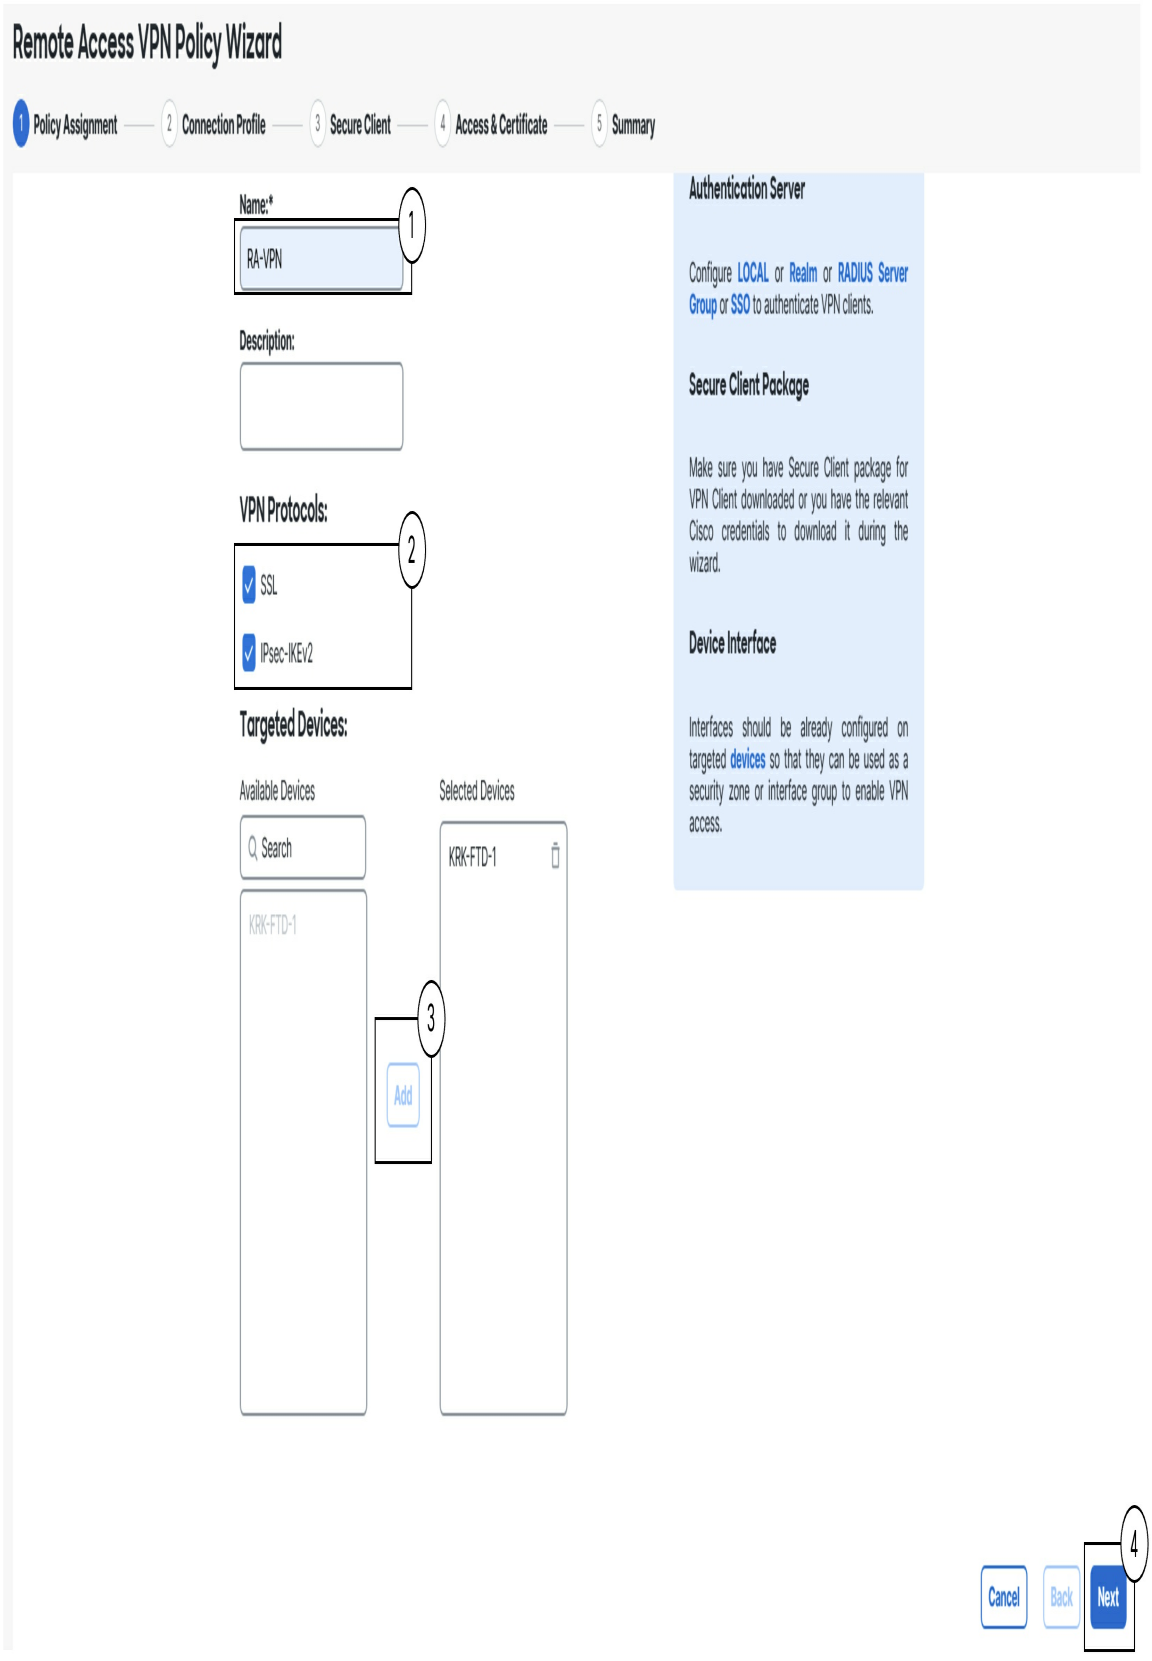

Remote Access VPN Policy Wizard
1 Policy Assignment
2 Connection Profile
3 Secure Client
4 Access & Certificate
5 Summary
Name:*
RA-VPN
1
Description:
VPN Protocols:
SSL
IPsec-IKEv2
2
Targeted Devices:
Available Devices
Search
KRK-FTD-1
Add
3
Selected Devices
KRK-FTD-1
Authentication Server
Configure LOCAL or Realm or RADIUS Server Group or SSO to authenticate VPN clients.
Secure Client Package
Make sure you have Secure Client package for VPN Client downloaded or you have the relevant Cisco credentials to download it during the wizard.
Device Interface
Interfaces should be already configured on targeted devices so that they can be used as a security zone or interface group to enable VPN access.
Cancel
Back
Next
4

Figure 13-25 *RA VPN Wizard Policy Assignment Step*

As part of this wizard page, complete the following steps:

Step 1. Define the remote access VPN configuration name.

Step 2. Select the protocol that will be used for remote access VPN.

Step 3. Specify which Secure Firewall appliance this configuration should apply to.

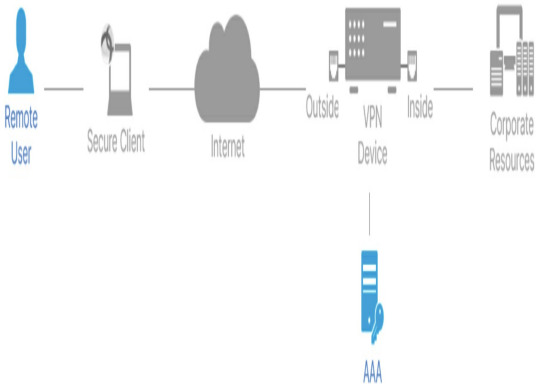
Step 4. Click **Next** to proceed to the next wizard page.

The second page of the wizard focuses on connection profile settings, including AAA settings, IP address assignment options, and group policy selection. Figure 13-26 shows the top of this wizard page.

Technet24

Remote Access VPN Policy Wizard

1 Policy Assignment — 2 Connection Profile — 3 Secure Client — 4 Access & Certificate — 5 Summary



Connection Profile:

Connection Profiles specify the tunnel group policies for a VPN connection. These policies pertain to creating the tunnel itself, how AAA is accomplished and how addresses are assigned. They also include user attributes, which are defined in group policies.

Connection Profile Name:*

RA-VPN (1)

This name is configured as a connection alias, it can be used to connect to the VPN gateway

Authentication, Authorization & Accounting (AAA):

Specify the method of authentication (AAA, certificates or both), and the AAA servers that will be used for VPN connections.

Authentication Method:

AAA Only (2)

Authentication Server:*

(3)

(LOCAL or Realm or RADIUS)

Figure 13-26 *RA VPN Connection Profile Settings*

Complete the following steps on this page:

Step 1. Change the remote access VPN connection profile name if needed. By default, the name matches the VPN configuration.

Step 2. Select how AAA should be performed for the remote access VPN. Figure 13-27 shows the list of available options.

Technet24

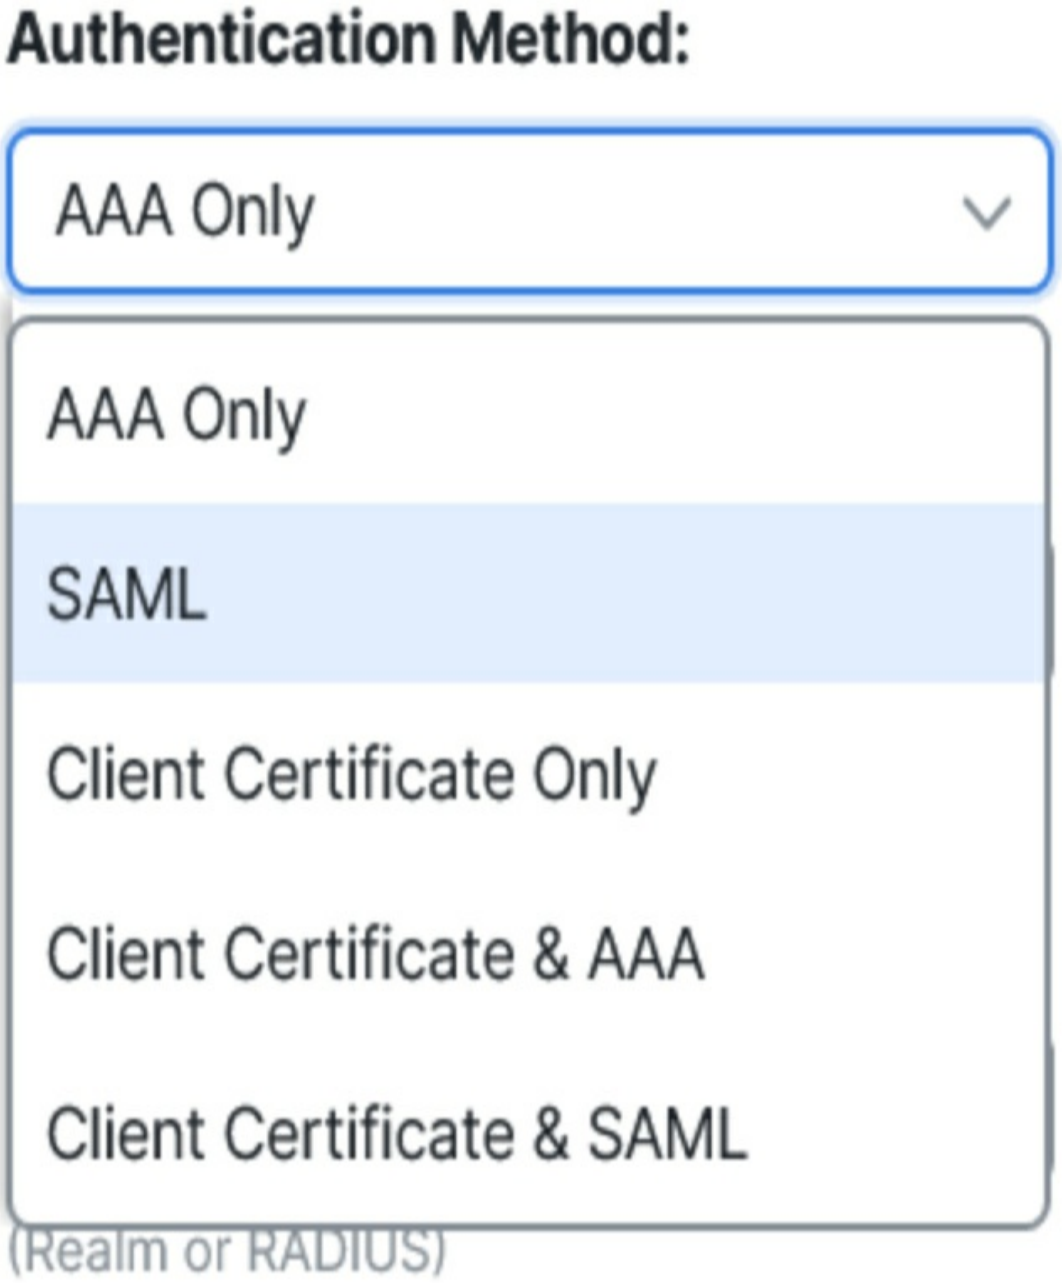


Figure 13-27 *RA VPN AAA Options*

Step 3. Select the authentication server group. Add a new group by clicking the **+** button.

To define a new RADIUS server group, click the **+** button next to the

Authentication Server drop-down and select **RADIUS Server Group**, as shown in Figure 13-28.

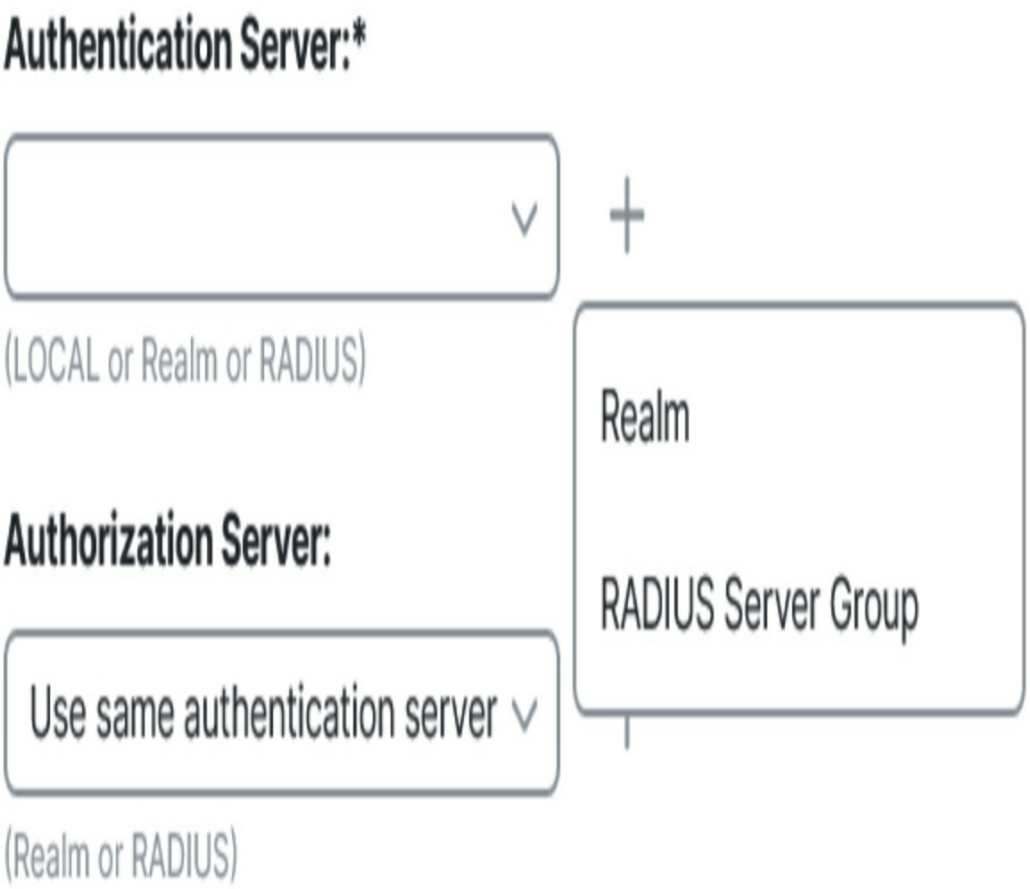


Figure 13-28 *Adding a RADIUS Server Group to the VPN Connection Profile*

An example of RADIUS server group configuration is provided in Figure 13-29. Server group-level settings are as follows:

- **Group Name:** This name is selected within the VPN connection profile.
- **Group Accounting Mode:** Choose **Single** or **Multiple** to determine how accounting information is propagated. With ISE, it

Technet24

is recommended to use Single to ensure accounting messages reach a single PSN at a time.

- **Retry Interval:** This value defines after how many seconds the RADIUS request should be retried. The group-level value overrides values set for individual servers.
- **Enable Authorize Only:** Enable this setting for NAC use cases where only authorization requests are sent to ISE.
- **Enable Interim Accounting Updates:** Enable interim accounting for connected clients and specify the interim accounting update interval.
- **Enable Dynamic Authorization:** Use this setting to enable CoA.
- **RADIUS Servers:** Add individual RADIUS servers by clicking **+**.

Edit RADIUS Server Group

Name:*

ISE-LAB-PSNs

Description:

Group Accounting Mode:

Single

Retry Interval:* (1-10) Seconds

7

Realms:

- [x] Enable authorize only
- [x] Enable interim account update

Interval:* (1-120) hours

24

- [x] Enable dynamic authorization

Port:* (1024-65535)

1700

- [ ] Merge Downloadable ACL with Cisco AV Pair ACL
  - ( ) After Cisco AV Pair ACL
  - ( ) Before Cisco AV Pair ACL

RADIUS Servers (Maximum 16 servers)

IP Address/Hostname

Cancel Save

Technet24

Figure 13-29 *RADIUS Server Group Settings*

An example of individual RADIUS server configuration is shown in Figure 13-30. By default, Secure Firewall uses the routing table to identify the interface where the RADIUS server is connected. You can override this setting by selecting **Specific Interface**, as shown in Figure 13-30.

Edit RADIUS Server

IP Address/Hostname:*

172.22.15.11

Configure DNS at Threat Defense Platform Settings to resolve hostname

☑ RADIUS Server-Enabled Message Authenticator

The RADIUS server must be configured correctly for the message authenticator to work.

Authentication Port:* (1-65535)

1812

Key:*

••••••••

Confirm Key:*

••••••••

Accounting Port: (1-65535)

1813

Timeout: (1-300) Seconds

10

Connect using:

○ Routing ◉ Specific Interface

Inside +

Redirect ACL:

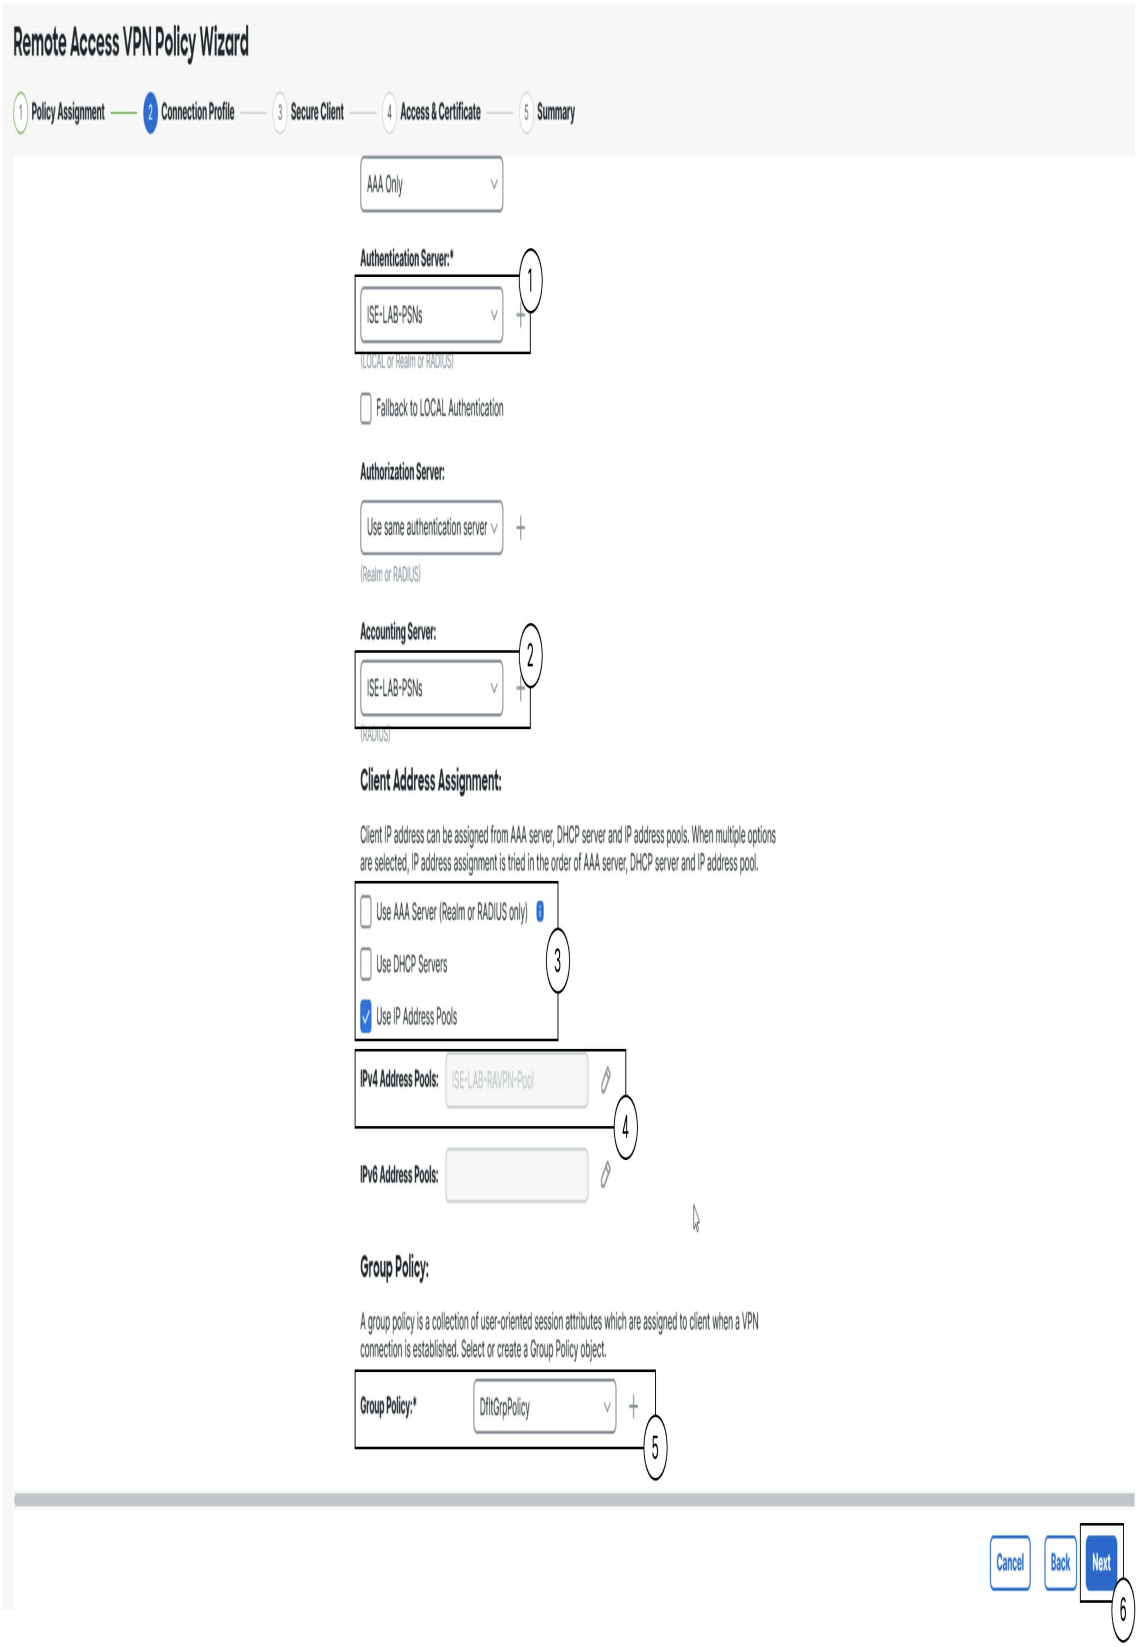
+

Cancel Save

Technet24

Figure 13-30 *RADIUS Server Settings*

When the RADIUS group configuration is completed, you return to the **Connection Profile** page of the wizard. Figure 13-31 displays the next steps in the Connection Profile configuration.


Remote Access VPN Policy Wizard
1 Policy Assignment
2 Connection Profile
3 Secure Client
4 Access & Certificate
5 Summary
AAA Only
Authentication Server:*
ISE-LAB-PSNs
(LOCAL or Realm or RADIUS)
Fallback to LOCAL Authentication
Authorization Server:
Use same authentication server
(Realm or RADIUS)
Accounting Server:
ISE-LAB-PSNs
(RADIUS)
Client Address Assignment:
Client IP address can be assigned from AAA server, DHCP server and IP address pools. When multiple options are selected, IP address assignment is tried in the order of AAA server, DHCP server and IP address pool.
Use AAA Server (Realm or RADIUS only)
Use DHCP Servers
Use IP Address Pools
IPv4 Address Pools:
ISE-LAB-RAVPN-Pool
IPv6 Address Pools:
Group Policy:
A group policy is a collection of user-oriented session attributes which are assigned to client when a VPN connection is established. Select or create a Group Policy object.
Group Policy:*
DfltGrpPolicy
Cancel
Back
Next
1
2
3
4
5
6

Technet24

Figure 13-31 *Connection Profile Configuration Continuation*

Step 1. Select the previously configured RADIUS server group for authentication.

Step 2. Choose the same group for accounting.

Step 3. Configure your IP address assignment method. Multiple methods can be enabled simultaneously. Address assignment is attempted in the following priority order: AAA server, followed by DHCP assignment, and finally by the IP local pool. Local pools can be defined at both the Connection Profile and Group Policy levels, with group policy pools always taking higher priority. The method that successfully assigns the address is used, and others are not attempted.

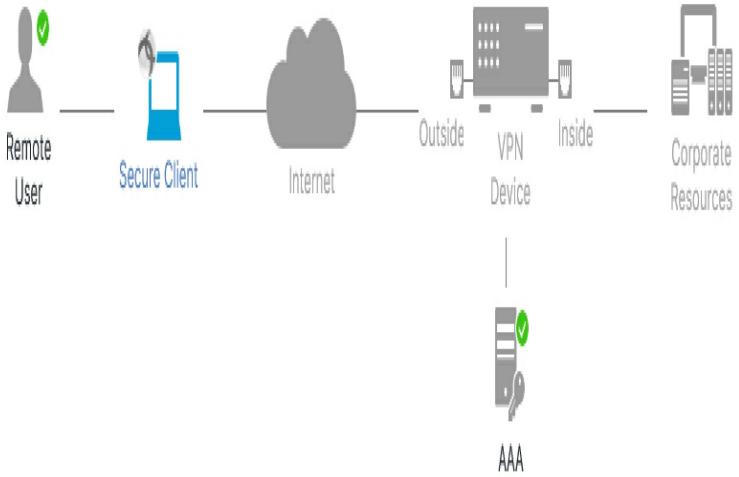
Step 4. Select a local address pool if needed.

Step 5. Modify an existing group policy or create a new one. DefaultGroupPolicy is always selected by default. Use the **+** button to create a new named group policy. After it is created, assign it to the Connection Profile. If you plan to assign group policies from ISE, use the **+** button to create those group policies without assigning them to the Connection Profile.

Step 6. Click **Next** to proceed to the next step.

The next wizard page is called **Secure Client**, as shown in Figure 13-32. On this page, select the Secure Client image previously uploaded to FMC.

1 Policy Assignment — 2 Connection Profile — 3 Secure Client — 4 Access & Certificate — 5 Summary



Secure Client Image

The VPN gateway can automatically download the latest Secure Client package to the client device when the VPN connection is initiated. Minimize connection setup time by choosing the appropriate OS for the selected package.

Download Secure Client packages from **Cisco Software Download Center**.

Show Re-order buttons

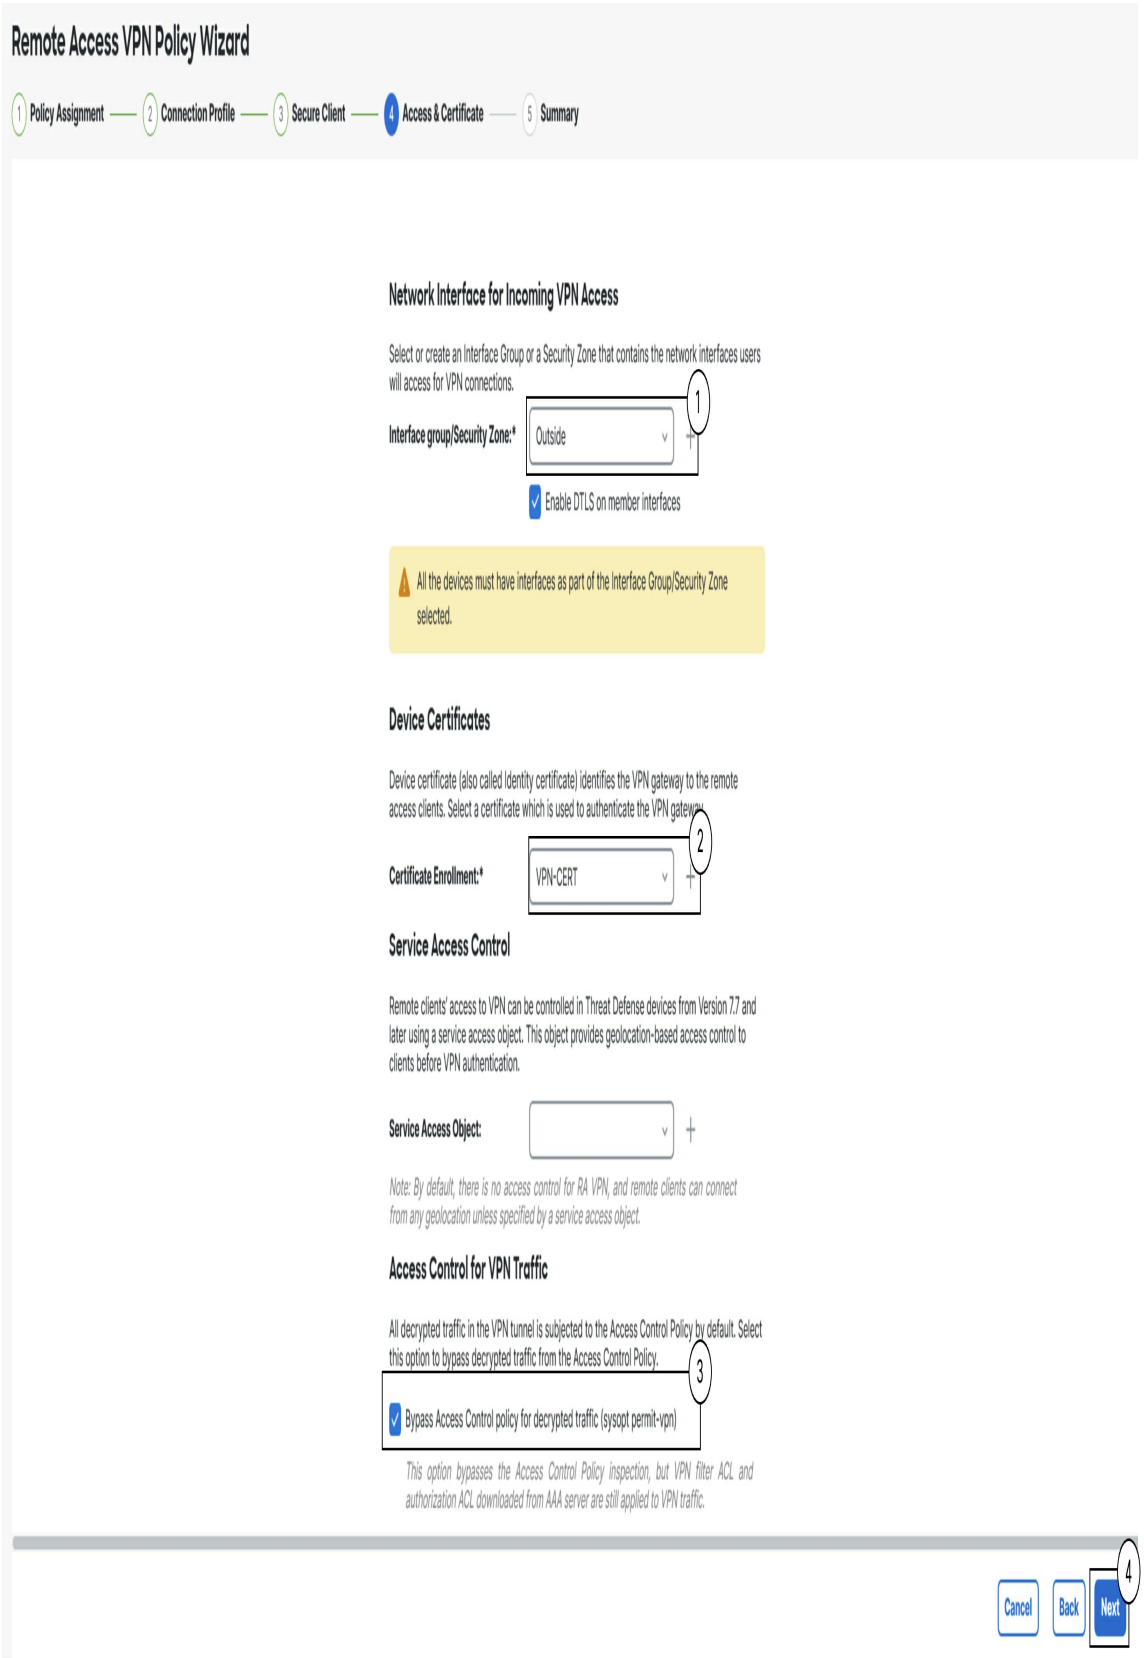
Secure Client File Object Name	Secure Client Package Name	Operating System
cisco-secure-client-win-5.1.16.19...	cisco-secure-client-win-5.1.16.194-webdep...	Windows

Cancel Back Next

Technet24

Figure 13-32 *The Secure Client Page of RA VPN Wizard*

The following wizard steps represent the final set of actions where configuration is required. After you complete the settings on the **Access & Certificate** wizard page, as shown in Figure 13-33, proceed to the **Summary** page, which displays all configuration details for the remote access VPN policy. After reviewing the configuration, save your changes and conclude the wizard.


Remote Access VPN Policy Wizard
1 Policy Assignment
2 Connection Profile
3 Secure Client
4 Access & Certificate
5 Summary
Network Interface for Incoming VPN Access
Select or create an Interface Group or a Security Zone that contains the network interfaces users will access for VPN connections.
Interface group/Security Zone:*
Outside
1
Enable DTLS on member interfaces
All the devices must have interfaces as part of the Interface Group/Security Zone selected.
Device Certificates
Device certificate (also called Identity certificate) identifies the VPN gateway to the remote access clients. Select a certificate which is used to authenticate the VPN gateway.
Certificate Enrollment:*
VPN-CERT
2
Service Access Control
Remote clients' access to VPN can be controlled in Threat Defense devices from Version 7.7 and later using a service access object. This object provides geolocation-based access control to clients before VPN authentication.
Service Access Object:
Note: By default, there is no access control for RA VPN, and remote clients can connect from any geolocation unless specified by a service access object.
Access Control for VPN Traffic
All decrypted traffic in the VPN tunnel is subjected to the Access Control Policy by default. Select this option to bypass decrypted traffic from the Access Control Policy.
3
Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)
This option bypasses the Access Control Policy inspection, but VPN filter ACL and authorization ACL downloaded from AAA server are still applied to VPN traffic.
Cancel
Back
Next
4

Technet24

Figure 13-33 *Access & Certificate Page of RA VPN Wizard*

Complete the following required steps on the **Access & Certificate** wizard page:

Step 1. Select the zone or interface where remote access VPN needs to be enabled. If the policy affects multiple Secure Firewall appliances, ensure that the same zone or interface is present on all devices.

Step 2. Select the Certificate Enrollment profile available on each Secure Firewall appliance affected by the policy.

Step 3. Enable ACP bypass for remote access VPN traffic if required.

Step 4. Click **Next** to proceed to the next wizard page.

On the next page, review the configuration and save your changes.

Basic ISE Setup for Remote Access VPN with Secure Firewall

From the ISE perspective, basic integration with Secure Firewall requires the following configuration elements:

- **NAD entry for Secure Firewall appliances:** Add Secure Firewall appliances that need to authenticate remote access VPN users as NADs in ISE. By using the device type NAD attribute, you can group them to later create a dedicated policy set.
- **DACLs:** If you plan to assign DACLs as part of remote access VPN authentication and authorization, create them so they can later be assigned by ISE.
- **Authorization profiles:** Create profiles to assign IP addresses,

DACLs, group policies, and other attributes to remote access VPN users according to organizational security policy requirements.

- **Policy set:** Create a dedicated policy set for remote access VPN.
- **Authentication policies:** If identity stores are already preconfigured, the default authentication policy can usually be used. If more specific requirements exist, add additional authentication policies.
- **Authorization policies:** Set policies that will provision granular access for users based on organizational security policy requirements.

A NAD entry created for the Secure Firewall appliance is shown in Figure 13-34. This NAD is assigned to the Device Type Cisco-VPN, which allows later use in policy set selection.

Technet24

Name* HQ-FTD-1

Description

IP Address * IP : 172.22.15.15 / 32

Device Profile Cisco

Model Name

Software Version

Network Device Group

Location All Locations Set To Default

IPSEC No Set To Default

Device Type Cisco-VPN Set To Default

Figure 13-34 *NAD Settings for Secure Firewall Appliance*

An authorization profile created for a remote access VPN client connecting to Secure Firewall can include DACL, Framed-IP-Address, and ASA VPN, which translates into the RADIUS Class attribute and carries the group policy name. Figure 13-35 shows an authorization profile used to assign DACL and ASA VPN parameters. If you need to assign Framed-IP-Address, select it in the **Advanced Attributes Settings** section of the authorization profile. Assigning a specific IP directly in the authorization profile is uncommon because it would apply to multiple users. Administrators typically prefer to map Framed-IP-Address to another attribute obtained by ISE during authentication and authorization that contains an IP for the specific user. The most common approach is to assign IPs in Active Directory using one of its attributes and later map this attribute to Framed-IP-Address in the authorization profile.

Technet24

Name* HQ-FTD-1

Description

IP Address * IP : 172.22.15.15 / 32

Device Profile Cisco

Model Name

Software Version

Network Device Group

Location All Locations Set To Default

IPSEC No Set To Default

Device Type Cisco-VPN Set To Default

Figure 13-35 *NAD Settings for Secure Firewall Appliance*

Figure 13-36 shows a policy set created for requests arriving from NADs assigned to the Cisco-VPN device type, with the default authentication policy and a simple authorization rule that assigns a DACL and a group policy to domain users. It is important to configure the group policy with the same name as used in the authorization profile on the Secure Firewall side.

Technet24

Status	Policy Set Name	Description	Conditions	Allowed Protocols / Server Sequence	Hits
Search					
	Cisco-VPN		DEVICE-Device Type EQUALS All Device Types#Cisco-VPN	Default Network Access	18

∨Authentication Policy(2)

Status	Rule Name	Conditions	Use	Hits	Actions
Search					
	Default		All_User_ID_Stores > Options	1	

〉Authorization Policy - Local Exceptions

〉Authorization Policy - Global Exceptions

∨Authorization Policy(2)

			Results			
Status	Rule Name	Conditions	Profiles	Security Groups	Hits	Actions
Search						
	Authorize-Only-Catch-All	ISELAB-ExternalGroups EQUALS iselab.net/Users/Domain Users	Secure-Firewall-Basic-...	Select from list	8	
	Default		DenyAccess	Select from list	11	

Figure 13-36 *Simple RA VPN Authentication and Authorization Policies on ISE*

Secure Firewall Remote Access VPN: Configuring NAC Use Cases

This section examines the configuration for more complex NAC use cases applicable to remote access VPN. They include the following:

- Certificate-based authentication with authorization on ISE
- SAML-based authentication with authorization on ISE
- Posture for remote access VPN users

The MDM use case is excluded because it does not require any specific settings on the Secure Firewall side. From an ISE perspective, you use the same settings described in Chapter 10.

Certificate-Based Authentication with Authorization on ISE

For certificate-based authentication of remote access VPN users with authorization on ISE, consider the following key points:

- Secure Firewall handles certificate authentication. This process requires the user certificate chain to be trusted by Secure Firewall, accomplished by installing the CAs from the user's chain as enrollment objects. Certificate verification also involves chain reconstruction to confirm that the certificate chain provided by Secure Client contains the appropriate certificates, as well as certificate validation, revocation checking if configured, and username extraction.
- ISE performs authorization for the username extracted from the

Technet24

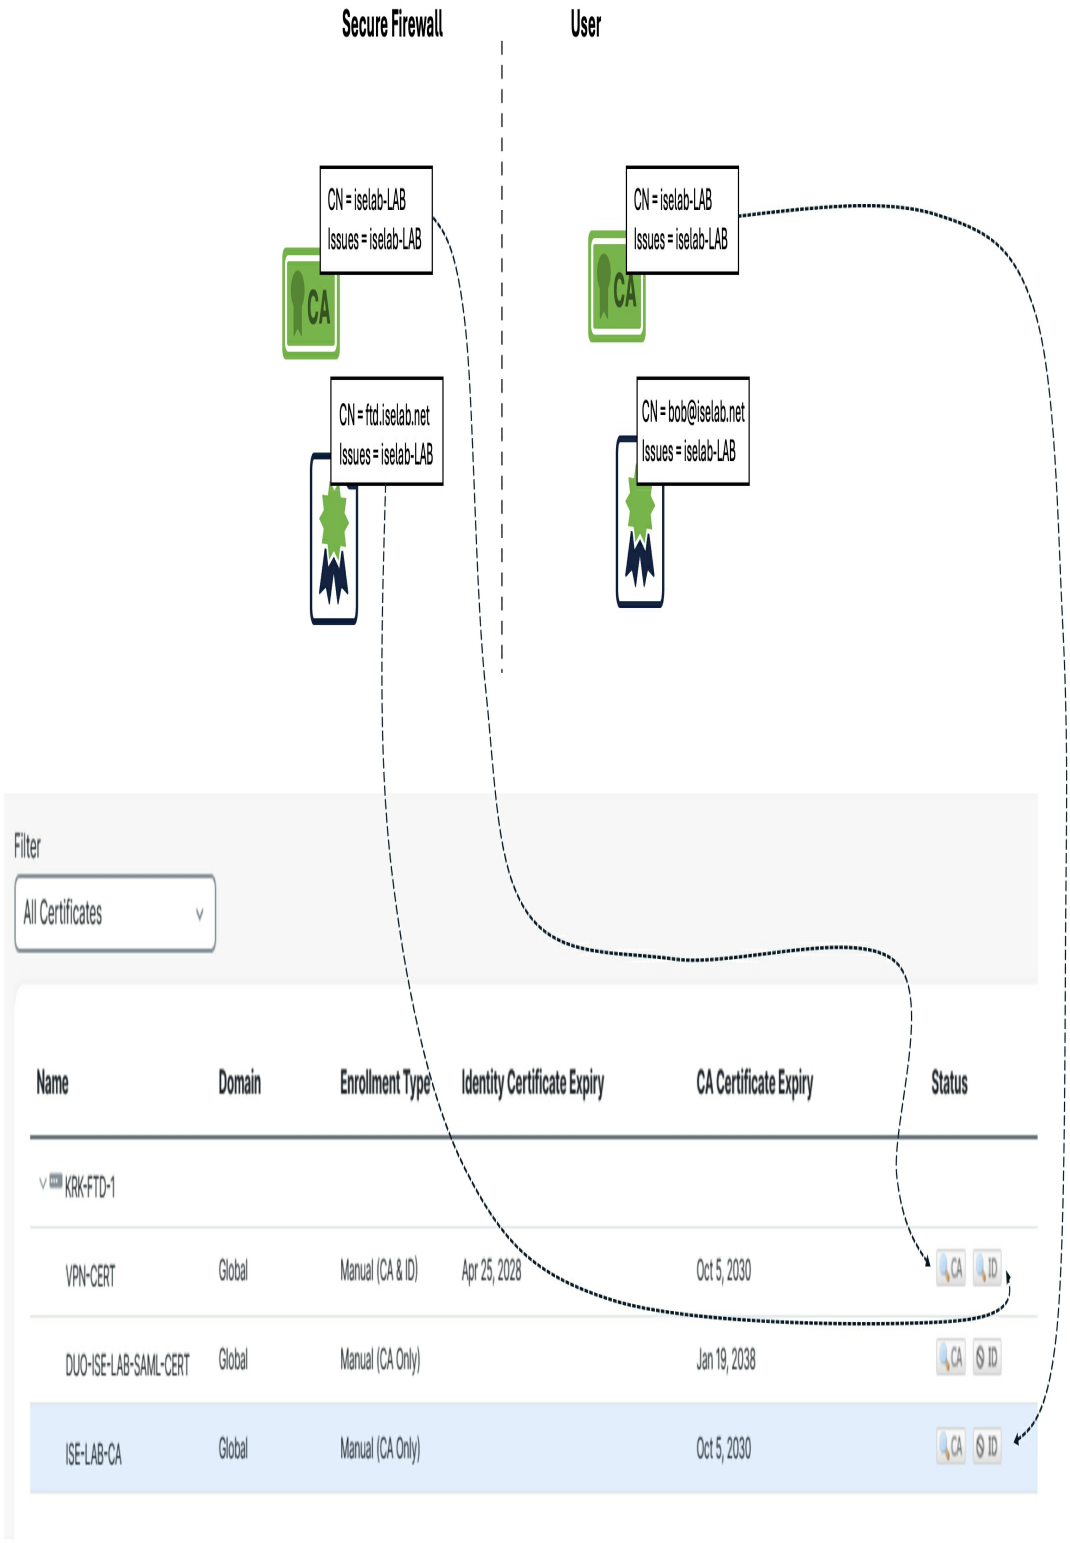
certificate. This process functions as a RADIUS access request with the Authorize-Only service type. Authentication always fails on ISE for this request, but you can configure policies to proceed with authorization.

- Authorization in ISE primarily relies on user lookup in the configured identity stores. You must ensure that Secure Firewall provides the correct username in the request for accurate authorization.

Note

The described flow requires the RADIUS server group to have the **Enable Authorize Only** feature enabled.

When certificate authentication is performed on Secure Firewall, remember that if both the firewall certificate and user certificate are issued by the same CA, the CA certificate must be installed twice — once as part of the enrollment object used for the firewall certificate and a second time as a CA-only enrollment object. The enrollment object containing the identity certificate cannot be used to validate end-user certificates. Figure 13-37 illustrates this requirement.


Secure Firewall
User
CN = iselab-LAB
Issues = iselab-LAB
CA
CN = iselab-LAB
Issues = iselab-LAB
CA
CN = ftd.iselab.net
Issues = iselab-LAB
CN = bob@iselab.net
Issues = iselab-LAB
Filter
All Certificates
Name
Domain
Enrollment Type
Identity Certificate Expiry
CA Certificate Expiry
Status
KRK-FTD-1
VPN-CERT
Global
Manual (CA & ID)
Apr 25, 2028
Oct 5, 2030
CA
ID
DUO-ISE-LAB-SAML-CERT
Global
Manual (CA Only)
Jan 19, 2038
CA
ID
ISE-LAB-CA
Global
Manual (CA Only)
Oct 5, 2030
CA
ID

Technet24

Figure 13-37 *Certificate Enrollment Objects Mapping on Secure Firewall*

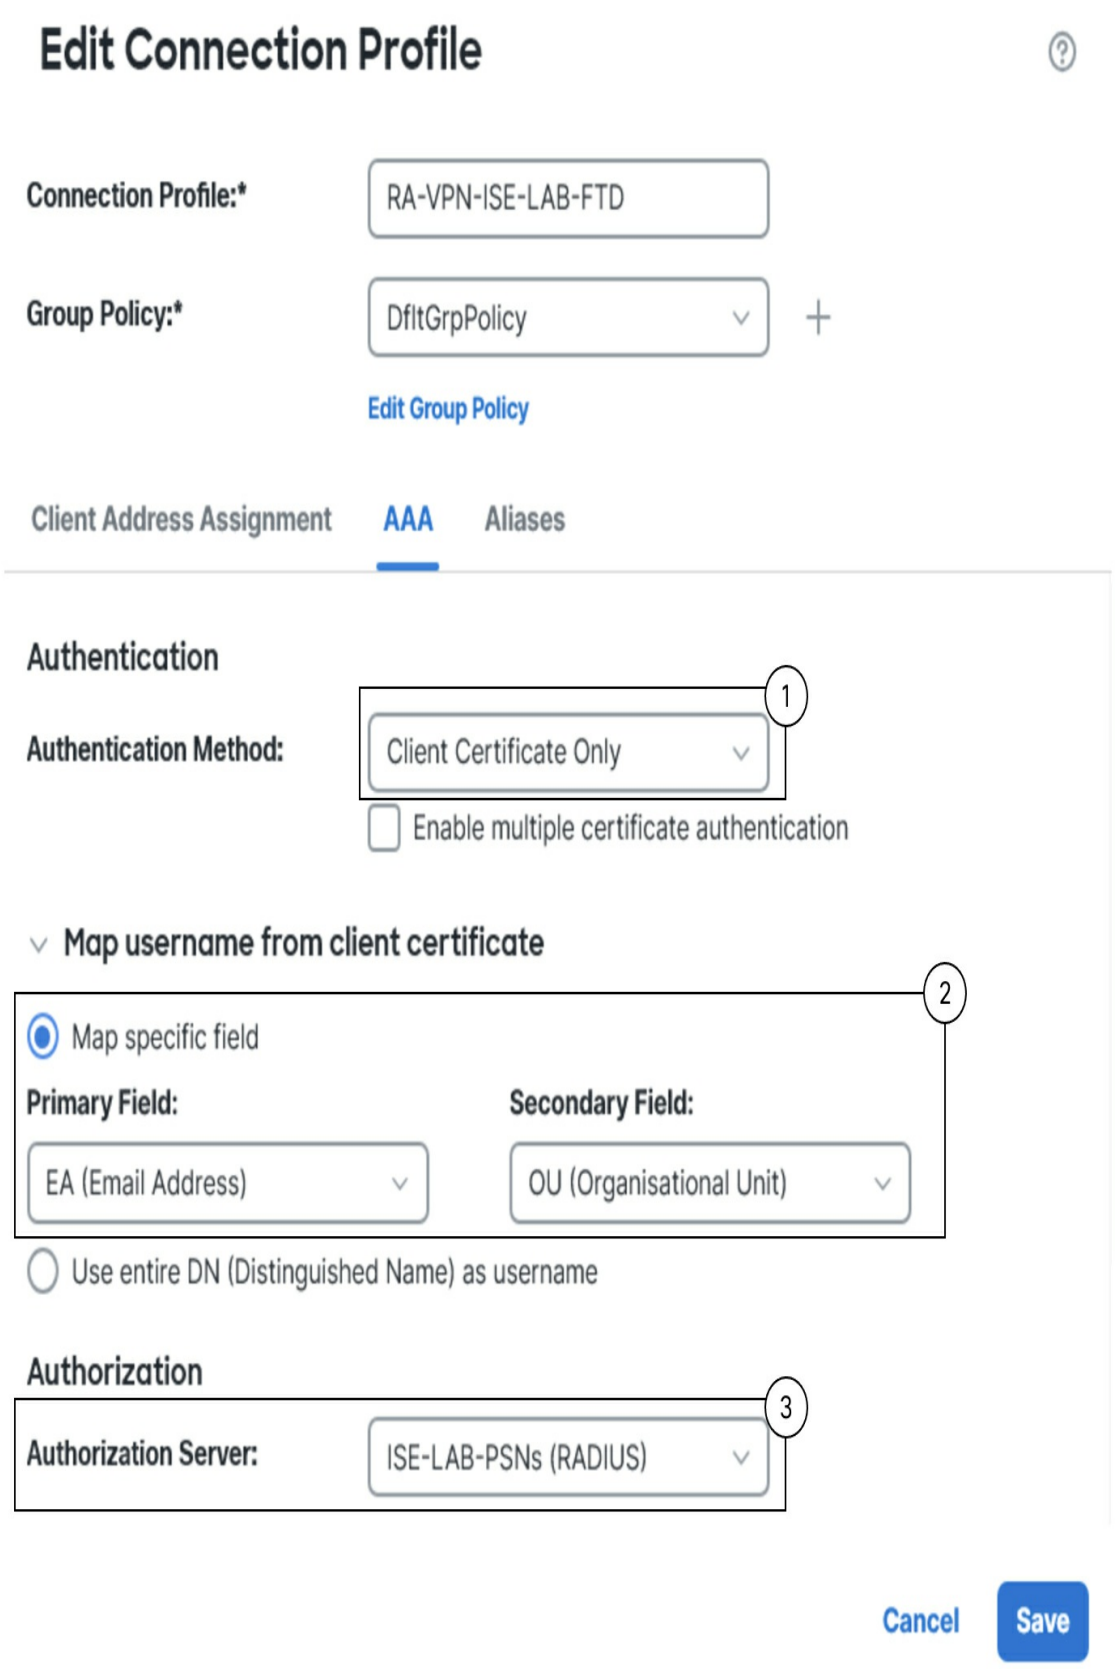
To change the authentication settings for an existing connection profile, navigate to **Secure Connections > Remote Access VPN** and click the pencil icon next to the policy that requires editing. Within the policy, choose the connection profile where authentication settings need to be modified by clicking the pencil icon next to the profile name.

Figure 13-38 provides an example of a connection profile configured for Client Certificate Only authentication, with authorization pointing to Cisco ISE and the email address from the CN used as the username. To transition your authentication to Client Certificate, follow these steps:

Step 1. Change the **Authentication Method** to **Client Certificate Only**.

Step 2. Use the **Map Specific Field** drop-down menu to map the appropriate attribute (**Primary Field** and, if necessary, **Secondary Field**) from the client certificate into the username.

Step 3. In the **Authorization Server** drop-down, select your RADIUS server group.


Edit Connection Profile
Connection Profile:*
RA-VPN-ISE-LAB-FTD
Group Policy:*
DfltGrpPolicy
Edit Group Policy
Client Address Assignment
AAA
Aliases
Authentication
Authentication Method:
Client Certificate Only
1
Enable multiple certificate authentication
Map username from client certificate
2
Map specific field
Primary Field:
EA (Email Address)
Secondary Field:
OU (Organisational Unit)
Use entire DN (Distinguished Name) as username
Authorization
Authorization Server:
ISE-LAB-PSNs (RADIUS)
3
Cancel
Save

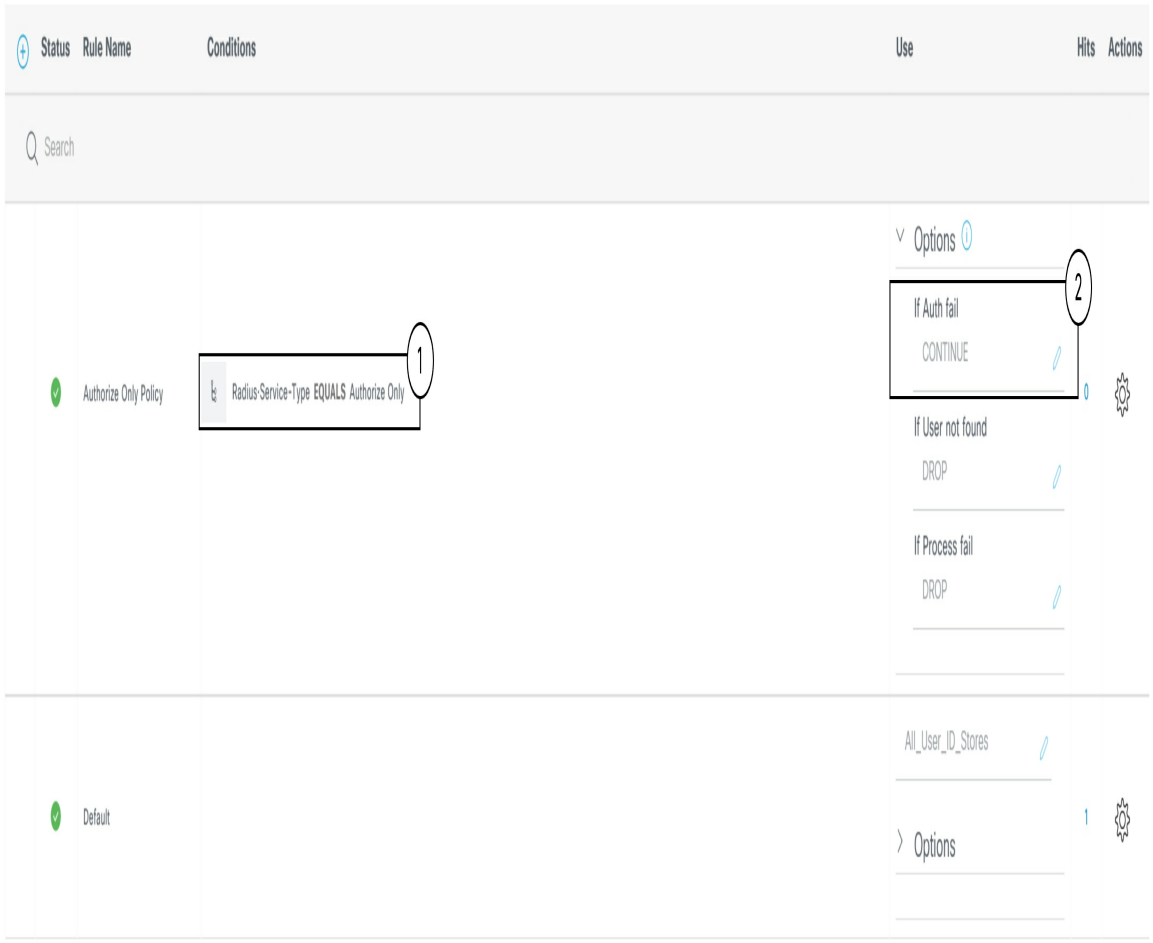
Technet24

Figure 13-38 *A Connection Profile Configured for Certificate-Based Authentication*

Note

It is possible to combine certificate authentication with username and password authentication. For this case, select **Client Certificate & AAA** as the authentication method.

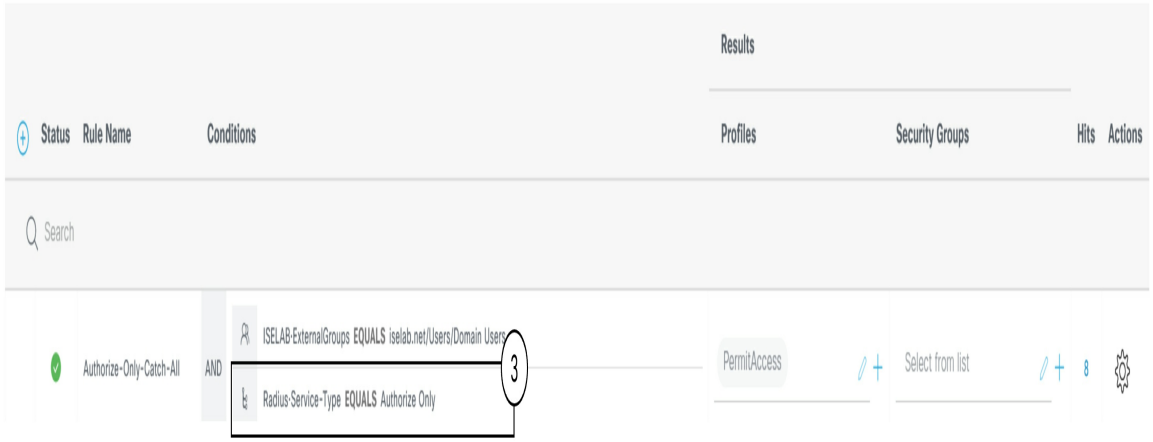
On the ISE side, add an additional authentication policy to the VPN policy set. This policy can match the Authorize-Only RADIUS Service-Type, along with any other required attributes. In the Policy Option, set **If Auth Fail** to **CONTINUE**, which enables ISE to proceed to authorization policy selection. Additionally, Authorize-Only can be used within the authorization policies. Figure 13-39 shows the ISE configuration.



Authorization Policy - Local Exceptions

Authorization Policy - Global Exceptions

Authorization Policy(2)



Technet24

Figure 13-39 *ISE Policies Configuration for Authorize-Only Use Case*

SAML-Based Authentication with Authorization on ISE

When SAML authentication is performed by Secure Firewall and authorization for the same access request occurs on ISE, you need to take the following configuration considerations into account:

- A SAML application for Secure Firewall must be added on the IdP side.
- The CA(s) of the SAML IdP certificate must be added to Secure Firewall through a certificate enrollment object.
- Create a new single sign-on AAA server for SAML authentication and assign it to the connection profile.
- Extract SAML details from Secure Firewall and provision them on the IdP side.
- Update ISE authentication and authorization policies if required.

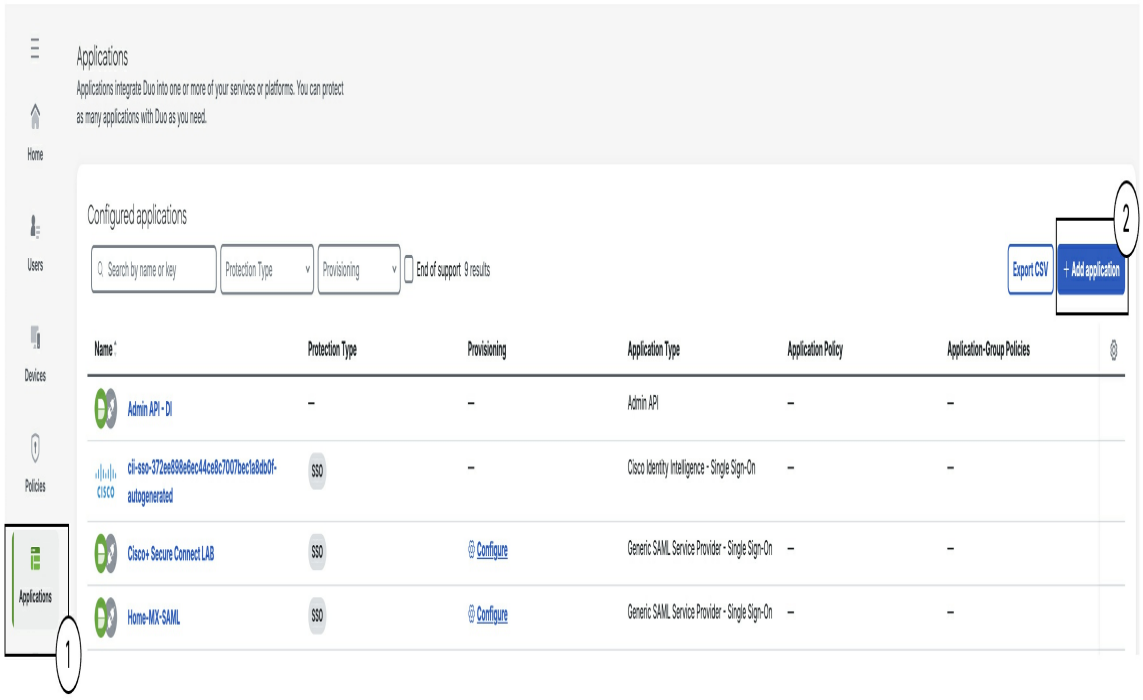
In this book, Cisco Duo is used as an example of a SAML IdP. When Secure Firewall needs to integrate with Duo for SAML authentication, you need to use a Generic SAML service provider application on the Duo side. To create this application, refer to Figure 13-40 and use the following steps:

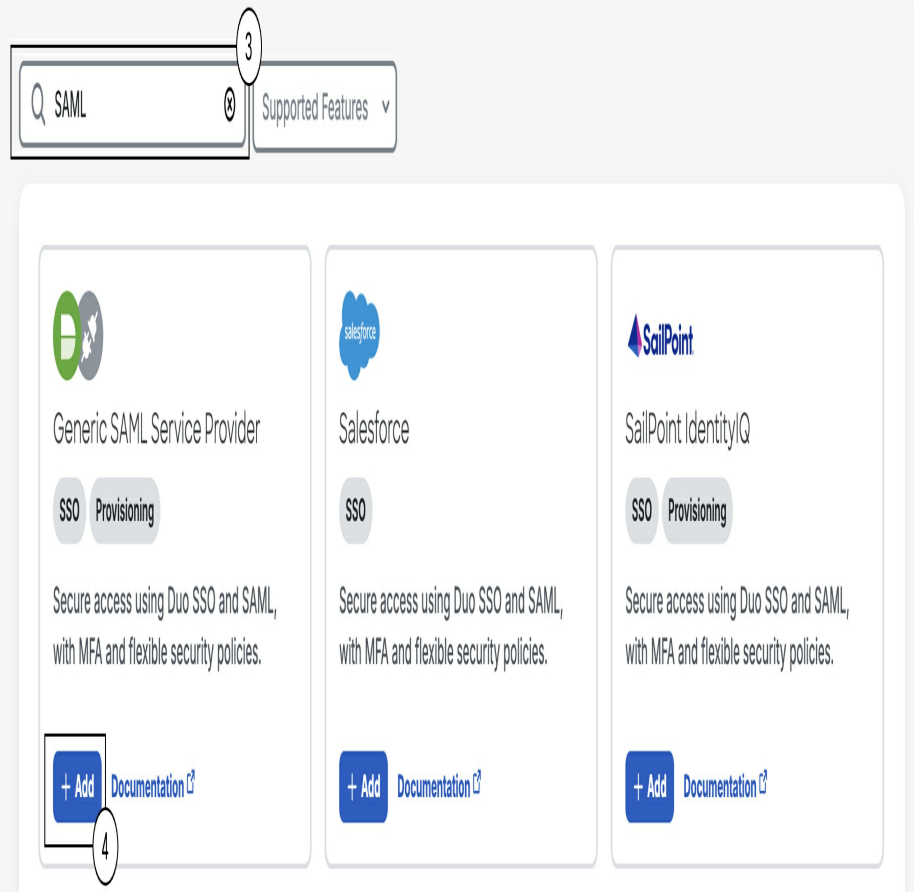
Step 1. Navigate to **Applications** in the **Duo** dashboard.

Step 2. Click **+ Add Application**.

Step 3. On the new page, search for **SAML**.

Step 4. Click **+ Add** for the Generic SAML Service Provider app.


Applications
Applications integrate Duo into one or more of your services or platforms. You can protect as many applications with Duo as you need.
Home
Users
Devices
Policies
Applications
Configured applications
Search by name or key
Protection Type
Provisioning
End of support
9 results
Export CSV
+ Add application
Name
Protection Type
Provisioning
Application Type
Application Policy
Application Group Policies
Admin API - DI
Admin API
cli-sso-372ee898e6ec44ce8c7007bec1a8db0f-autogenerated
Cisco Identity Intelligence - Single Sign-On
Cisco+ Secure Connect LAB
Configure
Generic SAML Service Provider - Single Sign-On
Home-MX-SAML
1
2


SAML
Supported Features
3
Generic SAML Service Provider
SSO
Provisioning
Secure access using Duo SSO and SAML, with MFA and flexible security policies.
+ Add
Documentation
4
Salesforce
SSO
Secure access using Duo SSO and SAML, with MFA and flexible security policies.
+ Add
Documentation
SailPoint IdentityIQ
SSO
Provisioning
Secure access using Duo SSO and SAML, with MFA and flexible security policies.
+ Add
Documentation

Technet24

Figure 13-40 *Create SAML App in DUO Dashboard*

The certificate that must be installed on Secure Firewall can be downloaded from the app, as shown in Figure 13-41. In the case of Duo, the certificate is self-signed and should be installed in FMC as a manual enrollment object, CA-only. After installing the certificate enrollment object in FMC, verify that it has been provisioned to the Secure Firewall where the remote access VPN configuration will be applied before proceeding to the next steps. More information on certificate installation is available in the section "Secure Firewall Identity Certificate Provisioning" earlier in this chapter. If your IdP uses a certificate signed by a CA, install all CA certificates in the chain on Secure Firewall as CA-only enrollment objects. Additionally, for the upcoming step, save the following:

1. Entity ID
2. Single Sign-On URL
3. Single Log-Out URL

Application name *

ISA-LAB-VPN

The unique identifier of the service provider

Application Type

Generic SAML Service Provider - Single Sign-On

User access

- Disable for all users
  Some users may still have access. See exceptions below.
- Enable only for permitted groups
  Some users may still have access. See exceptions below.
- **Enable for all users**

Metadata

1

Entity ID

https://sso-05bd4b19.sso.duosecurity.com/saml2/sp/DI9FUN6P47L0SI79CU25/meta Copy

2

Single Sign-On URL

https://sso-05bd4b19.sso.duosecurity.com/saml2/sp/DI9FUN6P47L0SI79CU25/sso Copy

3

Single Log-Out URL

https://sso-05bd4b19.sso.duosecurity.com/saml2/sp/DI9FUN6P47L0SI79CU25/slo Copy

Metadata URL

https://sso-05bd4b19.sso.duosecurity.com/saml2/sp/DI9FUN6P47L0SI79CU25/meta Copy

Certificate Fingerprints

SHA-1 Fingerprint

D0:38:0B:4B:E5:AD:C1:5E:3F:FC:51:DC:FF:7F:4D:1C:35:28:C5:E0 Copy

SHA-256 Fingerprint

21:53:77:21:32:55:B0:A1:22:11:54:C7:D7:6A:FB:F9:CD:CD:CE:E2:3D:DC:09:FC:0F:0F:2 Copy

Downloads

Certificate

Download certificate | Copy certificate | Expires: 01-19-2038

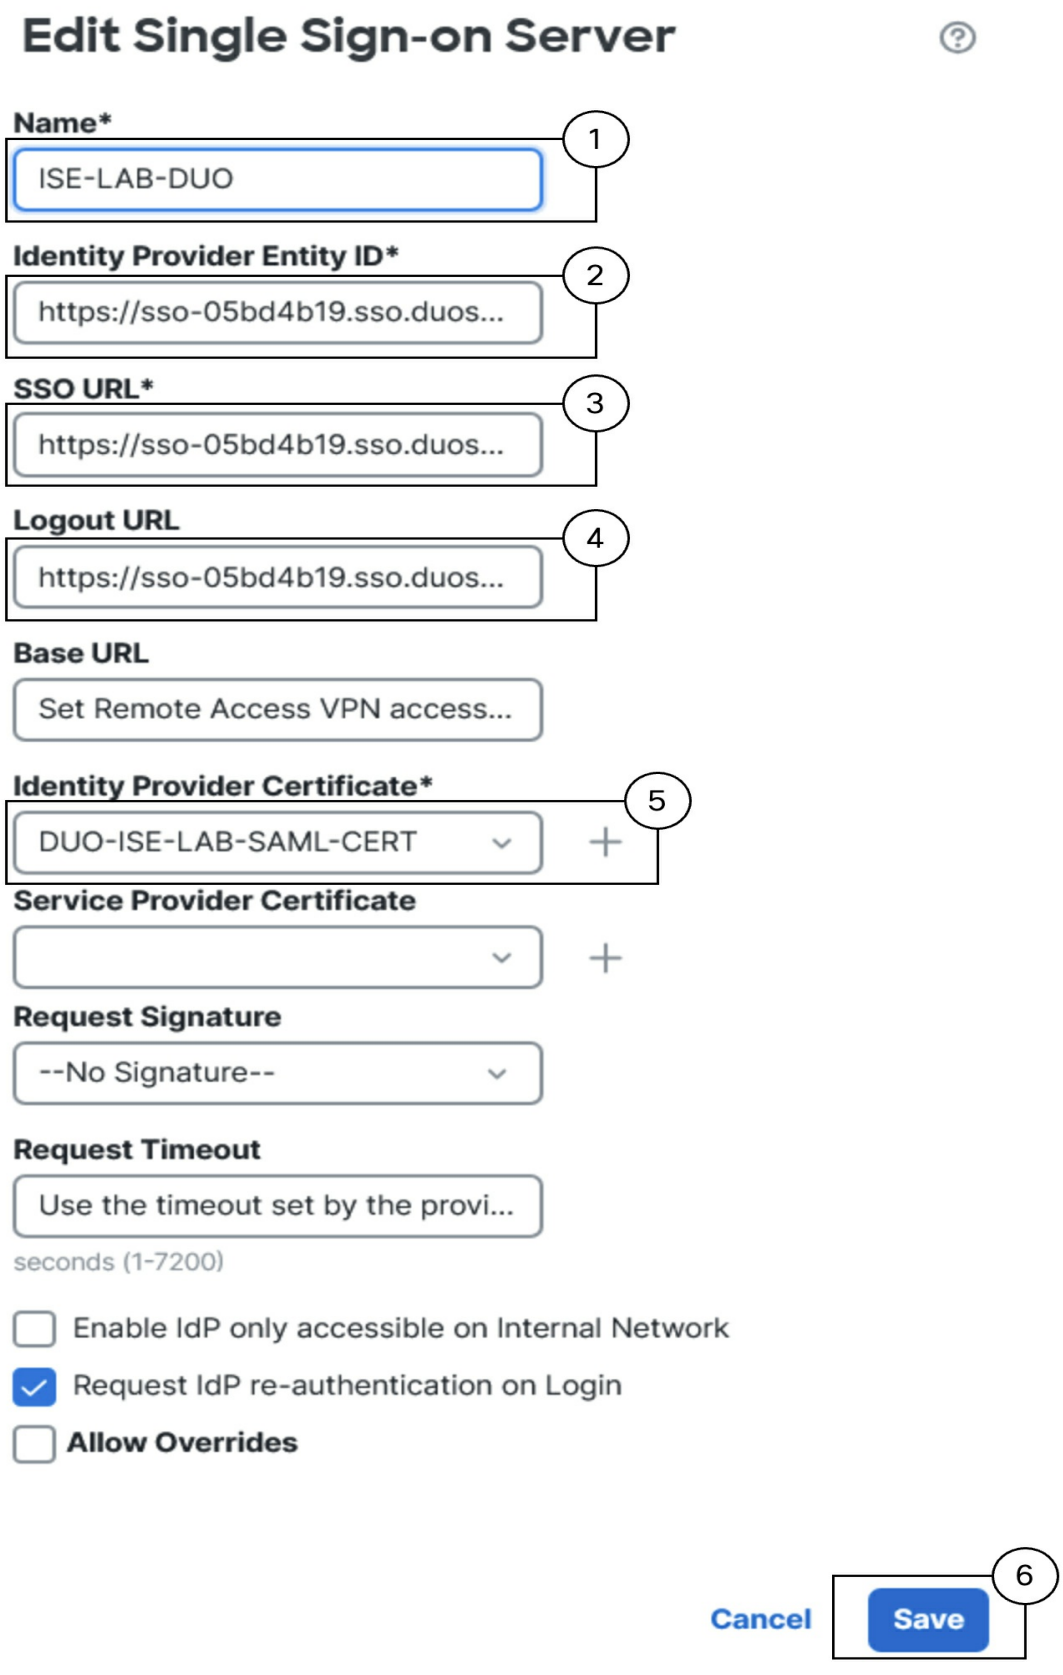
SAML Metadata

Download XML | Copy XML

Technet24

Figure 13-41 *Certificate Download from DUO SAML App.*

As the next step, create an SSO AAA server in the FMC interface. Navigate to **Objects > AAA > Single Sign-on Server** and click **Add Single Sign-on Server**. Figure 13-42 shows a single sign-on server configuration.


Edit Single Sign-on Server
Name*
ISE-LAB-DUO
1
Identity Provider Entity ID*
https://sso-05bd4b19.sso.duos...
2
SSO URL*
https://sso-05bd4b19.sso.duos...
3
Logout URL
https://sso-05bd4b19.sso.duos...
4
Base URL
Set Remote Access VPN access...
Identity Provider Certificate*
DUO-ISE-LAB-SAML-CERT
5
Service Provider Certificate
Request Signature
--No Signature--
Request Timeout
Use the timeout set by the provi...
seconds (1-7200)
Enable IdP only accessible on Internal Network
Request IdP re-authentication on Login
Allow Overrides
Cancel
Save
6

Technet24

Figure 13-42 *Single Sign-on AAA Server Configuration*

As a minimum, configure the server settings by following these steps:

Step 1. Define the SSO server **Name**.

Step 2. Set the **Identity Provider Entity ID** using details from the IdP app page.

Step 3. Specify the IdP **SSO URL** using information from the IdP app page.

Step 4. Set the IdP **Logout URL** using details from the IdP app page.

Step 5. Provide the **Identity Provider Certificate**.

Step 6. Click **Save** to save the configuration.

> ## Note
>
> For additional security, you can enable request signing. This feature must also be available and enabled on the IdP side. From the SSO server configuration perspective, when this feature is enabled, you are required to select a certificate enrollment profile for the Service Provider certificate. This profile must contain both the identity certificate and CA certificate.

Lastly, reconfigure the connection profile in the VPN policy for SAML authentication with authorization on ISE. Figure 13-43 shows the connection profile AAA configuration for this use case.

Edit Connection Profile

Connection Profile:* RA-VPN-ISE-LAB-FTD

Group Policy:* DfltGrpPolicy

Edit Group Policy

Client Address Assignment | AAA | Aliases

Authentication

Authentication Method: SAML

Authentication Server: ISE-LAB-DUO (SSO)

Override Identity Provider Certificate

SAML Login Experience:

- (•) VPN client embedded browser
- ( ) Default OS Browser

Authorization

Authorization Server: ISE-LAB-PSNs (RADIUS)

[x] Allow connection only if user exists in authorization database

Accounting

Cancel | Save

Technet24

Figure 13-43 *Connection Profile SAML AAA Configuration*

After assigning the SAML AAA server to the connection profile, extract its details over the CLI to complete the IdP app configuration for Secure Firewall integration. For this task, log in to the Secure Firewall appliance diagnostic CLI and execute the command **show saml metadata** , as shown in Figure 13-44. Copy the values of the entityID and AssertionConsumerService attributes from the output, as shown in the figure.

krk-esc-ftd-1# show saml metadata RA-VPN-ISE-LAB-FTD

SP Metadata

<?xml version="1.0" encoding="UTF-8" standalone="yes"?>

<EntityDescriptor entityID="https://172.22.30.15/saml/sp/metadata/RA-VPN-ISE-LAB-FTD" xm

lns="urn:oasis:names:tc:SAML:2.0:metadata">

<SPSSODescriptor AuthnRequestsSigned="false" WantAssertionsSigned="true" protocolSupport

Enumeration="urn:oasis:names:tc:SAML:2.0:protocol">

<KeyDescriptor use="signing">

<ds:KeyInfo xmlns:ds="http://www.w3.org/2000/09/xmldsig#">

<ds:X509Data>

<ds:X509Certificate></ds:X509Certificate>

</ds:X509Data>

</ds:KeyInfo>

</KeyDescriptor>

<AssertionConsumerService index="0" isDefault="true" Binding="urn:oasis:names:tc:SAML:2. 0:bindings:HTTP-POST"

Location="https://172.22.30.15/+CSCOE+/saml/sp/acs?tgname=RA-VPN-I SE-LAB-FTD" />

<SingleLogoutService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect" Locati

on="https://172.22.30.15/+CSCOE+/saml/sp/logout"/><SingleLogoutService Binding="urn:oasi s:names:tc:SAML:2.0:bindings:HTTP-

POST" Location="https://172.22.30.15/+CSCOE+/saml/sp/l ogout"/></SPSSODescriptor>

</EntityDescriptor>

Technet24

Figure 13-44 *SAML AAA Server Metadata*

Configure these values on the IdP side. In the case of Cisco Duo, insert the entityID value into the **Entity ID** field and the AssertionConsumerService value into the **Assertion Consumer Service (ACS) URL** field. Figure 13-45 shows the application on the Duo side for Secure Firewall integration.

Service Provider

Metadata Discovery

None (manual input)

Entity ID *

https://172.22.30.15/saml/sp/metadata/RA-VPN-ISE-LAB-F

The unique identifier of the service provider.

Assertion Consumer Service (ACS) URL *

https://172.22.30.15/+CSCOE+/saml/sp/acs?tgname=RA-VI

+ Add an ACS URL

The service provider endpoint that receives and processes SAML assertions.

Single Logout URL

Single Logout URL

Optional: The service provider endpoint that receives and processes SAML logout requests.

Technet24

Figure 13-45 *Finalizing iDP Side Configuration*

From an ISE perspective, you can use the same policies created for certificate-based authentication with authorization on ISE. There are two caveats for the SAML-based flow:

- Identity stores configured on ISE must resolve the identity provided by the user during SAML login on the IdP. This is typically a username in the format of an email address, such as bob@iselab.net.
- The AAA RADIUS server group on Secure Firewall must have the **Enable Authorize Only** feature enabled.

Posture for the Remote Access VPN Users

The posture configuration for remote access VPN closely resembles the posture settings required for wired or wireless access. Authorization policies on ISE can follow the same structure discussed in Chapter 10.

The key difference in Secure Firewall scenarios is the need to properly configure the Redirect ACL, which can later be used as part of the posture assessment in the authorization policy applied when posture status is unknown or not equal to compliant.

On the Secure Firewall, you need to complete the following stages to define a redirect ACL and apply it to the RADIUS server:

- Configure an extended ACL that bypasses redirection for traffic destined to ISE and other critical services such as DNS.
- Assign the configured redirect ACL to the RADIUS servers.
- Confirm that the VPN client can resolve the fully qualified domain names (FQDNs) required for the posture process to function.

To configure the redirect ACL, refer to Figure 13-46 and use the following steps:

Step 1. Navigate to **Objects**.

Step 2. Click **Access List**.

Step 3. Click **Extended**.

Step 4. Click **Add Extended Access List**.

Technet24



Figure 13-46 *FMC Extended ACL Page*

When configuring redirect ACLs for ISE posture, place deny statements at the top for resources that must be excluded from redirection attempts. Typically, these resources include the following:

- ISE PSN addresses where posture assessment can be performed. Exclude all ISE PSNs.
- IP addresses of DNS servers.
- IP addresses of any critical systems that must be accessible before the posture process completes. Keep the list minimal and include only what is absolutely necessary.
- IP addresses of remediation servers. If internal resources are needed for the posture process to complete, such as servers for anti-malware updates, exclude them from redirection.
- Exclude resources from redirection either per service by specifying individual ports or by using permit any statements for specific IP addresses. The latter option is less secure.

The redirect ACL typically ends with a permit any statement, indicating that Secure Firewall will attempt to redirect traffic to any destination when it is an HTTP request over port 80. Figure 13-47 shows objects that cover resources excluded from redirection and the redirect ACL itself. At the top of the figure, the object group named ISE-LAB-Exclude-from-Redirect appears; it is a collection of host objects, each representing a specific destination.

Technet24

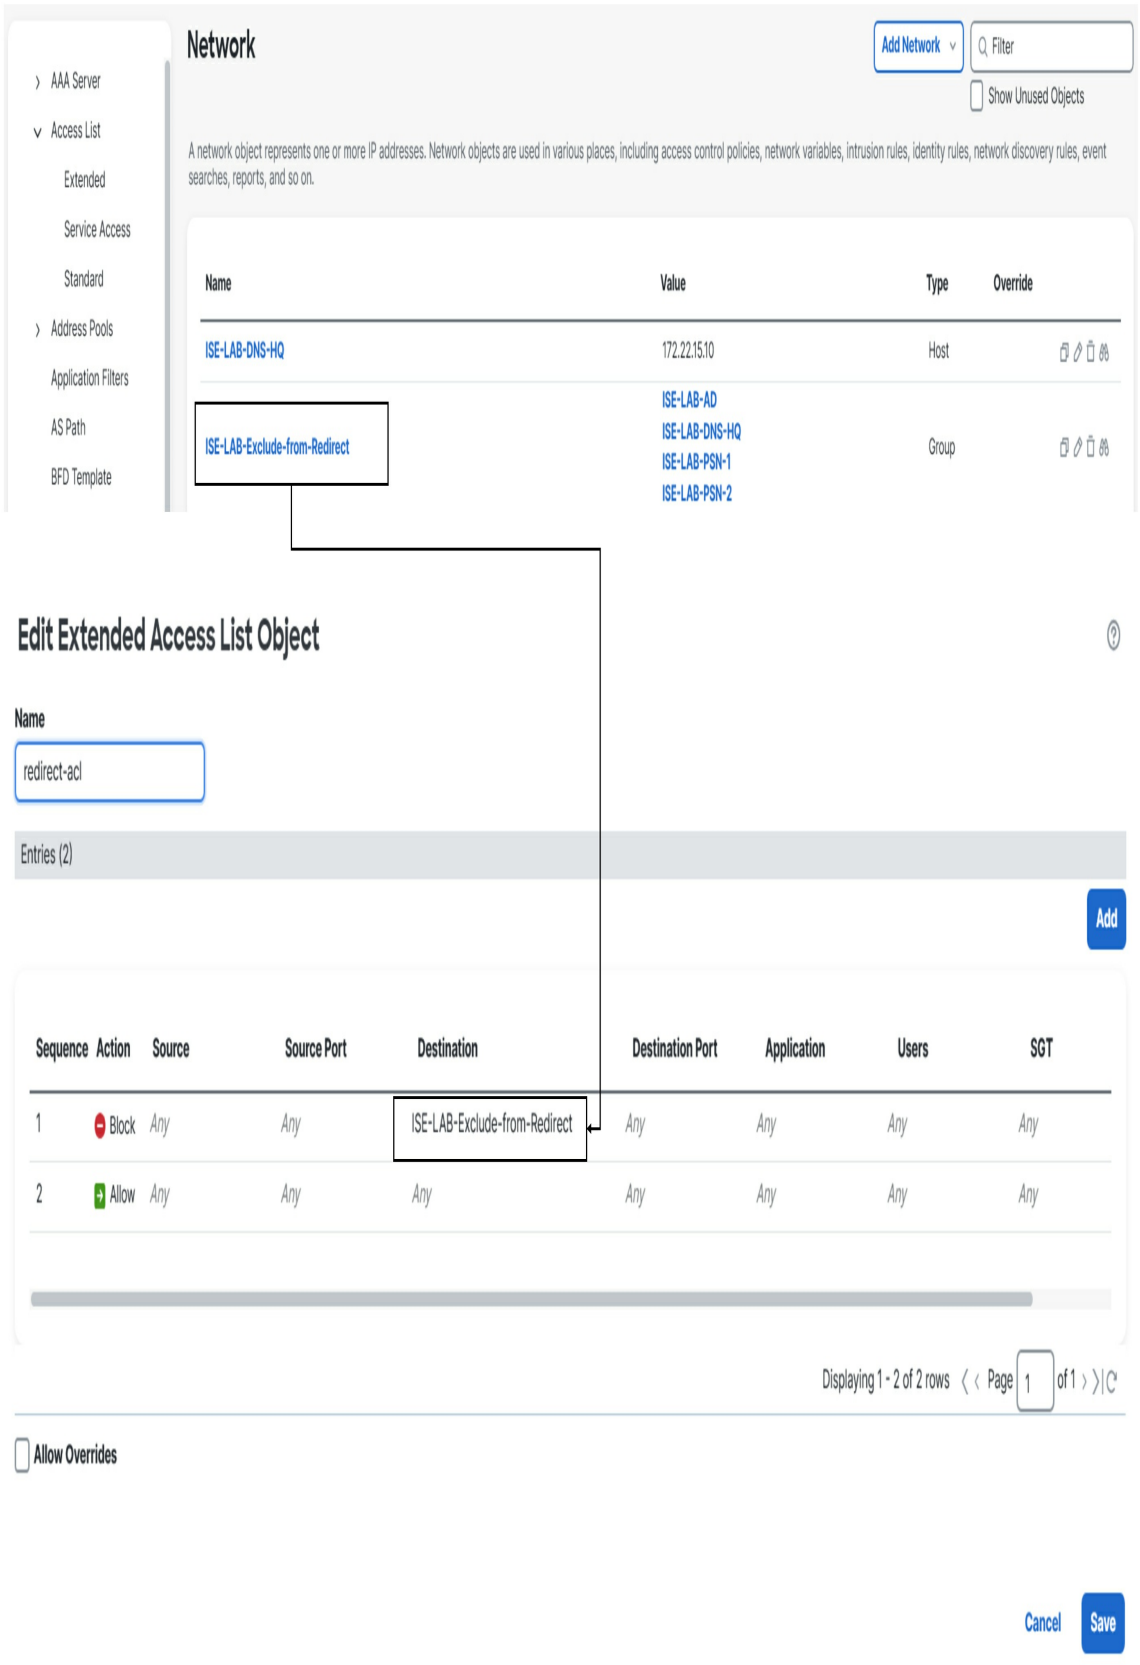

Network
Add Network
Filter
Show Unused Objects
A network object represents one or more IP addresses. Network objects are used in various places, including access control policies, network variables, intrusion rules, identity rules, network discovery rules, event searches, reports, and so on.
AAA Server
Access List
Extended
Service Access
Standard
Address Pools
Application Filters
AS Path
BFD Template
Name
Value
Type
Override
ISE-LAB-DNS-HQ
172.22.15.10
Host
ISE-LAB-Exclude-from-Redirect
ISE-LAB-AD
ISE-LAB-DNS-HQ
ISE-LAB-PSN-1
ISE-LAB-PSN-2
Group
Edit Extended Access List Object
Name
redirect-acl
Entries (2)
Add
Sequence
Action
Source
Source Port
Destination
Destination Port
Application
Users
SGT
1
Block
Any
Any
ISE-LAB-Exclude-from-Redirect
Any
Any
Any
Any
2
Allow
Any
Any
Any
Any
Any
Any
Any
Displaying 1 - 2 of 2 rows
Page 1 of 1
Allow Overrides
Cancel
Save

Figure 13-47 *FMC Redirect ACL Configuration*

After you create the ACL, it is not propagated to the Secure Firewall appliance automatically. To trigger propagation, assign the redirect ACL to the RADIUS servers. Navigate to **Objects > AAA > RADIUS Server Group** and click the pencil icon next to the group used in the remote access VPN policy. When the group is open, use the pencil icon next to each assigned RADIUS server to edit its settings and select the previously created extended ACL in the **Redirect ACL** drop-down, as shown in Figure 13-48.

Technet24

Edit RADIUS Server

IP Address/Hostname:*

172.22.15.11

Configure DNS at Threat Defense Platform Settings to resolve hostname

RADIUS Server-Enabled Message Authenticator

The RADIUS server must be configured correctly for the message authenticator to work.

Authentication Port:* (1-65535)

1812

Key:*

••••••••

Confirm Key:*

••••••••

Accounting Port: (1-65535)

1813

Timeout: (1-300) Seconds

10

Connect using:

Routing Specific Interface

Inside

Redirect ACL:

redirect-acl

Cancel Save

Figure 13-48 *Redirecting an ACL Assignment to the RADIUS Server*

Another important configuration step required for posture over VPN to function is setting up name resolution on the VPN client side so that ISE PSNs FQDNs and other FQDNs required for the posture process can be properly resolved. Client DNS settings are configurable in the group policy assigned to the connection profile. To adjust these settings, navigate to **Secure Connections > Remote Access VPN** and click the pencil icon next to the policy that requires editing. Within the policy, choose the connection profile where the group policy needs to be updated by clicking the pencil icon next to the profile name. Inside the profile, click **Edit Group Policy** under the group policy name, as shown in Figure 13-49.

Technet24


Edit Connection Profile
Connection Profile:*
RA-VPN-ISE-LAB-FTD
Group Policy:*
DfltGrpPolicy
Edit Group Policy

Figure 13-49 *Remote Access VPN Connection Profile, Access Group Policy Settings*

In the group policy window, switch to the **DNS/WINS** section in the **General** tab and ensure that the correct internal DNS servers are defined, along with the appropriate internal default domain suffix. Figure 13-50 shows the DNS/WINS configuration.

Technet24

Edit Group Policy

Name:*

DfltGrpPolicy

Description:

General | Secure Client | Advanced

- VPN Protocols
- IP Address Pools
- Banner
- DNS/WINS
- DHCP
- Split Tunneling

Primary DNS Server:

ISE-LAB-AD

Secondary DNS Server:

Primary WINS Server:

Secondary WINS Server:

Default Domain:

iselab.net

Cancel | Save

Figure 13-50 *Group Policy DNS/WINS Settings*

From the perspective of ISE posture module installation and version control for Secure Client software, consider the following important points:

- When split tunneling is enabled for remote access VPN users, ensure that posture probes are routed over the VPN tunnel. For Windows machines, where the default gateway is assigned by the OS to the VPN adapter, the probe to the gateway typically triggers the redirection. For macOS machines, where VPN adapters are considered point-to-point and do not have a default gateway value, make sure that enroll.cisco.com is routed inside the VPN tunnel.
- If a user does not have the ISE posture module preinstalled and split tunneling is enabled, typical browsing behavior will not trigger redirection. As well, in this scenario, standard redirection detection mechanisms in web browsers will not activate the redirect process. Consider sending the traffic of browser probes used to detect redirection over the VPN tunnel. More information about browser methods for redirection detection is provided in Chapter 10.
- From the perspective of agent version management, Secure Firewall and ISE typically use the same Secure Client version during initial deployment. If the ISE posture module is not installed from Secure Firewall, the agent version on ISE must match the agent version on Secure Firewall; otherwise, the ISE posture module will not be installed during the first connection attempt due to a version mismatch between headends.
- Once the ISE Posture agent is provisioned on the endpoint as part of the Secure Client installation, the Secure Client version on ISE

Technet24

can be lower than the version used by Secure Firewall. When an agent upgrade is required, perform the upgrade from the Secure Firewall side, because the core VPN module cannot be updated from ISE when a VPN connection is established.

Secure Firewall Remote Access VPN with ISE: Verification and Troubleshooting

From the Secure Firewall side, the main verification command to confirm details of the end user connection is **show vpn-sessiondb detail anyconnect**, which returns information for all AnyConnect VPN sessions. To verify session details for a specific user, use **show vpn-sessiondb detail anyconnect filter name** .

Execute these commands in the Secure Firewall appliance diagnostic CLI. Figure 13-51 shows a session for the user bob@iselab.net.

```
krk-esc-ftd-1# show vpn-sessiondb detail anyconnect filter name bob@iselab.net

Session Type: AnyConnect Detailed

Username    : bob@iselab.net      Index     : 28
Assigned IP : 192.168.111.10      Public IP   : 172.22.31.39
Protocol    : AnyConnect-Parent SSL-Tunnel DTLS-Tunnel
License     : AnyConnect Premium
Encryption  : AnyConnect-Parent: (1)none  SSL-Tunnel: (1)AES-GCM-128  DTLS-Tunnel: (1)AES-GCM-256
Hashing     : AnyConnect-Parent: (1)none  SSL-Tunnel: (1)SHA256  DTLS-Tunnel: (1)SHA384
Bytes Tx    : 166434753         Bytes Rx    : 57179507
Pkts Tx     : 424764            Pkts Rx     : 452094
Pkts Tx Drop : 0                Pkts Rx Drop : 0
Group Policy : DfltGrpPolicy      Tunnel Group : RA-VPN-ISE-LAB-FTD
Login Time  : 13:47:06 UTC Thu Jul 30 2026
Duration    : 24d 16h:42m:41s
Inactivity  : 0h:00m:00s
VLAN Mapping : N/A              VLAN      : none
Audt Sess ID : ac160f0f0001c0006a6b55da
Security Grp : none             Tunnel Zone  : 0
```

Technet24

Figure 13-51 *Anyconnect Session Details*

In the output, pay attention to the following attributes:

- **Username:** The provided username. This is especially important for use cases such as certificate-based authentication and SAML-based authentication.
- **Assigned IP:** The IP address assigned to the user.
- **Tunnel Group:** The tunnel group (connection profile) where the connection landed.
- **Group Policy:** The group policy assigned to the connection.
- **Client Ver:** The Secure Client version on the connecting device.
- **Filter Name:** The name of the downloadable ACL.
- **ISE Posture:** Values for Redirect URL and Redirect ACL are visible when the session is in a redirect state.

From the ISE perspective, the same verification procedures described in previous chapters can be performed. Figure 13-52 shows the authentication and authorization flow in the RADIUS live logs during posture assessment.



Figure 13-52 *ISE RADIUS Live Logs for Posture over VPN Flow*

The typical authentication and authorization flow for the posture over VPN use case includes the following steps:

Step 1. Initial authentication, which results in redirection. The authorization profile name indicates that Client

Technet24

Provisioning Posture (CPP_ results were assigned.

Step 2. COA push after successful posture assessment. This is identified as a COA push because the next authorization policy and profile are present in the Authorization Policy and Authorization Profile columns, respectively.

Step 3. An internal ISE session event representing the current state of the session — full network access.

Network Access Control for Cisco SSE with ISE

Network access control for Cisco SSE with ISE extends identity-based policy enforcement beyond traditional on-premises environments into cloud-delivered secure access. This approach combines Cisco ISE's contextual identity and posture capabilities with Cisco SSE's policy enforcement points to provide consistent access decisions for users and devices.

Cisco SSE and ISE: Design Considerations

When planning an integration between ISE and SSE for delivering NAC services, consider the following points:

- **Connectivity between ISE nodes and SSE:** Ensure proper planning for connectivity, since SSE is a cloud-native solution, whereas ISE servers are located within the corporate network, even when deployed in public clouds such as Azure or AWS.
- **NAD identification:** When SSE sends RADIUS requests to ISE PSNs, the requests originate from IP addresses within the SSE fabric. Properly define these addresses during SSE configuration and make sure that NADs on the ISE side are provisioned with them.

When considering connectivity between SSE and ISE, evaluate the following transport options and select the most appropriate based on your requirements and the current state of your network:

- **The IPsec tunnel:** SSE can establish an IPsec tunnel to any on-premises device that supports IPsec. RADIUS traffic can then be routed between SSE and ISE PSNs through the IPsec tunnel. SSE supports IPsec tunnel redundancy. The tunnel can be either policy-based or route-based. For route-based tunnels, SSE supports routing using BGP. To configure a site-to-site IPsec tunnel on SSE, navigate to **Connect > Network Connections**, go to the **Network Tunnel Groups** tab, and click **+Add** to configure a new IPsec tunnel.
- **Integration with Catalyst SD-WAN:** This integration, which also relies on individual IPsec tunnels, enables automated configuration. SSE API keys are installed in the Catalyst SD-WAN manager, making the entire integration easier to implement, manage, and monitor.
- **Integration with Meraki SD-WAN:** The Meraki dashboard can be integrated with SSE by adding virtual MX hubs into the SSE fabric. This allows the Meraki Auto VPN SD-WAN fabric to be extended into SSE. This integration also relies on API. SSE API keys must be provisioned in the Meraki dashboard.

For NAD identification on SSE, specify the System IP pool during each regional IP pool configuration. This IP range is used to source SSE-originated control plane traffic such as RADIUS. Later, on the ISE side, ensure you define a NAD for each SSE region and use the corresponding System IP pool as the NAD IP address in the NAD settings.

Technet24

SSE: Preparing Configuration Elements

The main prerequisites on the SSE side for NAC to work successfully with ISE are as follows:

- Users who need authentication or authorization through ISE must be present on the SSE side; otherwise, authorization of remote access VPN connections will fail. To address this issue, integrate SSE with one of the supported directories.
- RADIUS server groups must be created and assigned to the regions.

SSE Directory Options

As of this writing, SSE supports the following options for directory integration:

- **Identity Provider (IdP):** This option allows you to provision user details to SSE, such as usernames and group membership, from SAML identity providers like Cisco Duo, Okta, and others.
- **Manual Provisioning:** As an administrator, you can import user and group information using a CSV file.
- **Active Directory:** This option provides integration with on-premises Active Directory using the AD agent.
- **Google Workspace:** This option allows integration of SSE with Google Workspace as a directory to obtain users.

For this book, Active Directory is used as the directory example.

To add a directory connection in SSE, refer to Figure 13-53 and use the following steps:

Step 1. Navigate to **Connect**.

Step 2. Click **Users, Groups, and Endpoint Devices**.

Step 3. Click **Configuration Management**.

Technet24

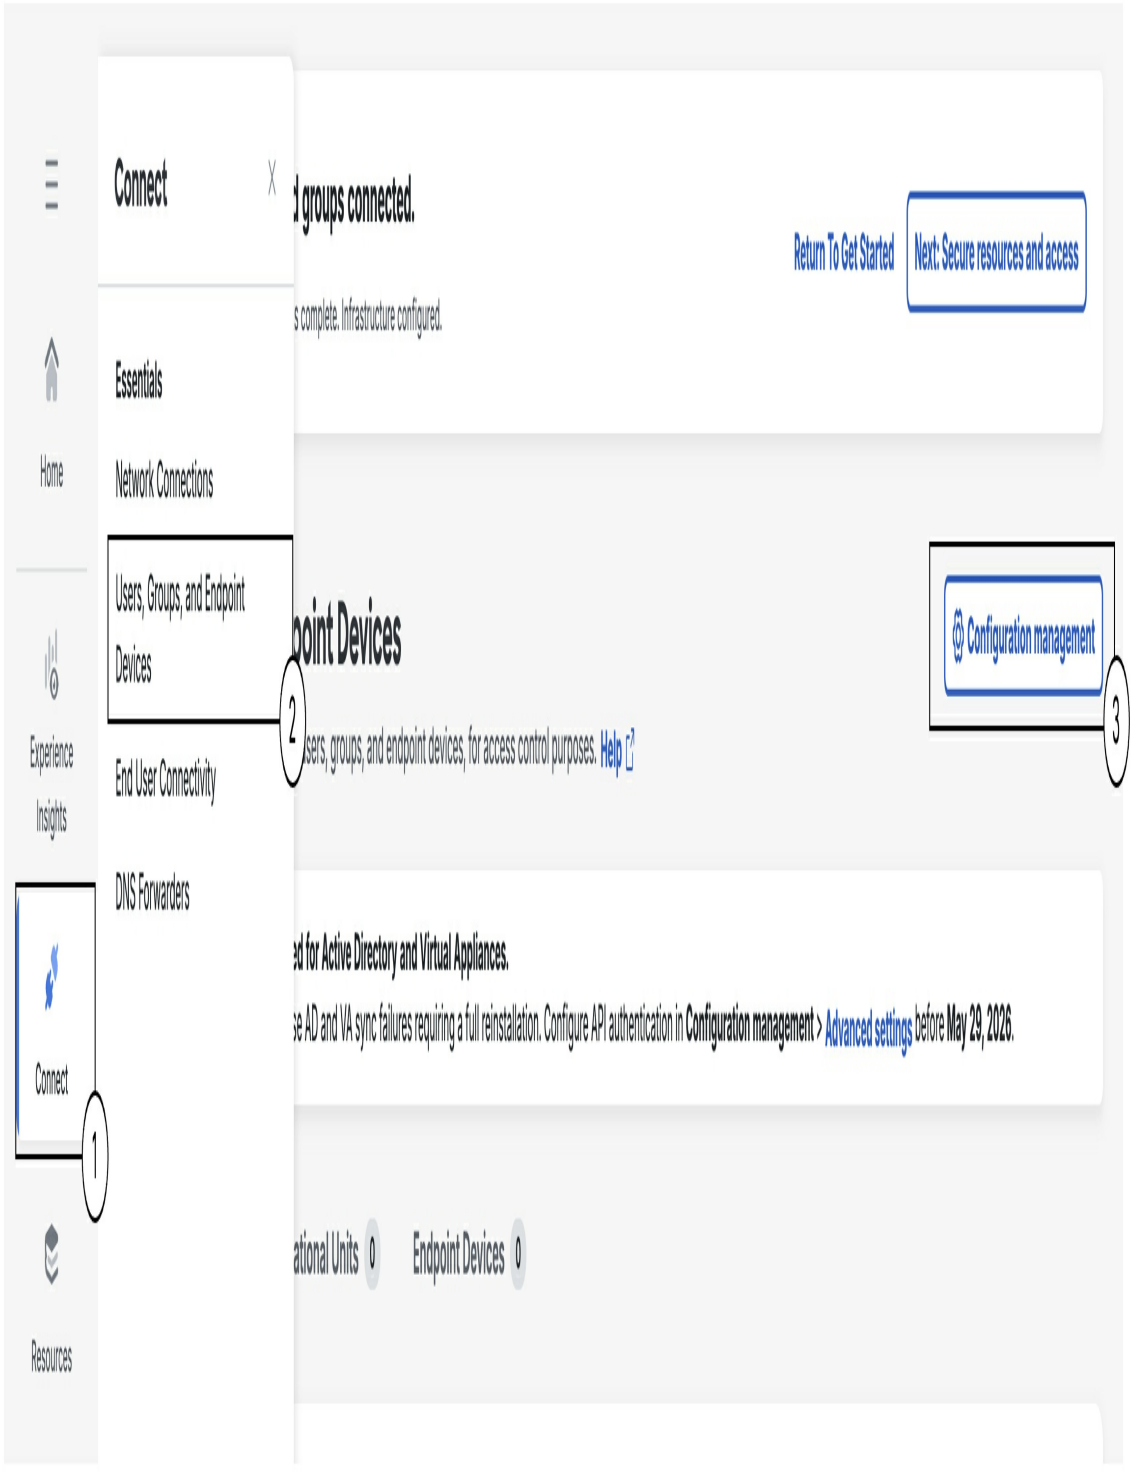
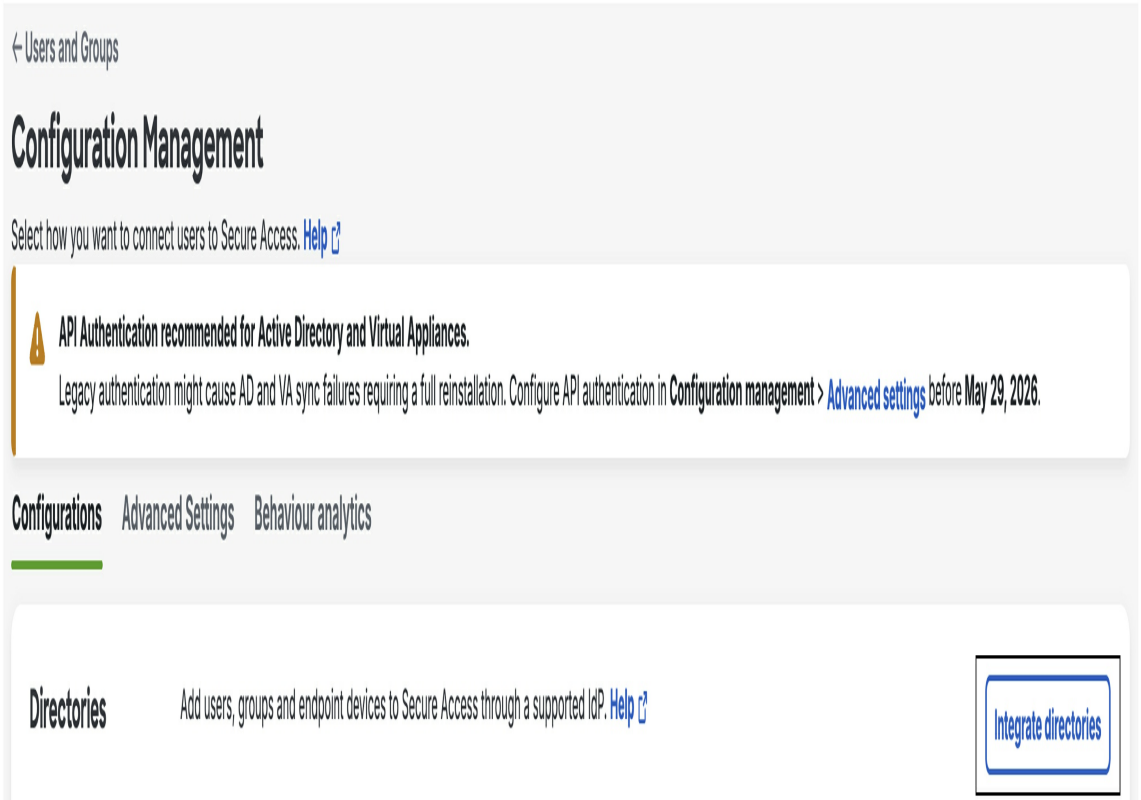

Connect
groups connected.
Return To Get Started
Next: Secure resources and access
complete. Infrastructure configured.
Home
Essentials
Network Connections
Users, Groups, and Endpoint Devices
2
Experience Insights
End User Connectivity
DNS Forwarders
point Devices
Configuration management
3
users, groups, and endpoint devices, for access control purposes. Help
Connect
1
ed for Active Directory and Virtual Appliances.
e AD and VA sync failures requiring a full reinstallation. Configure API authentication in Configuration management > Advanced settings before May 29, 2026.
ational Units 0
Endpoint Devices 0
Resources

Figure 13-53 *Access SSE Directory Configuration*

On the **Configuration Management** page, click the **Integrate Directories** button, and on the next page, proceed with the directory you would like to use. Figure 13-54 provides an example.

Technet24


← Users and Groups
Configuration Management
Select how you want to connect users to Secure Access. Help
API Authentication recommended for Active Directory and Virtual Appliances.
Legacy authentication might cause AD and VA sync failures requiring a full reinstallation. Configure API authentication in Configuration management > Advanced settings before May 29, 2026.
Configurations
Advanced Settings
Behaviour analytics
Directories
Add users, groups and endpoint devices to Secure Access through a supported IdP. Help
Integrate directories


Provisioning Method
Select a provisioning method to add users to Secure Access.
Methods
Identity provider(IdP)
Provision Users and Groups though a supported identity provider (IdP) service.
Manual Provisioning
Manually provision users and groups through the Secure Access CSV file and endpoint devices through API.
Active Directory
Connect Secure Access to your Active Directory and import your users and groups.
Google Workspace
Sync Google Workspace identities with Secure Access.

Figure 13-54 *Access SSE Directory Options*

Each directory type requires a unique set of configuration steps.

To begin configuring Active Directory integration, first add an Admin API SSE key and secret to the Advanced Settings in the **Users and Groups > Configuration Management** section. Refer to Figure 13-55, which demonstrates this process. To create the key, follow these steps:

Step 1. Click **Admin**.

Step 2. Select **API Keys**.

Step 3. Select **KeyAdmin Keys**.

Step 4. Click **Add**.

Technet24

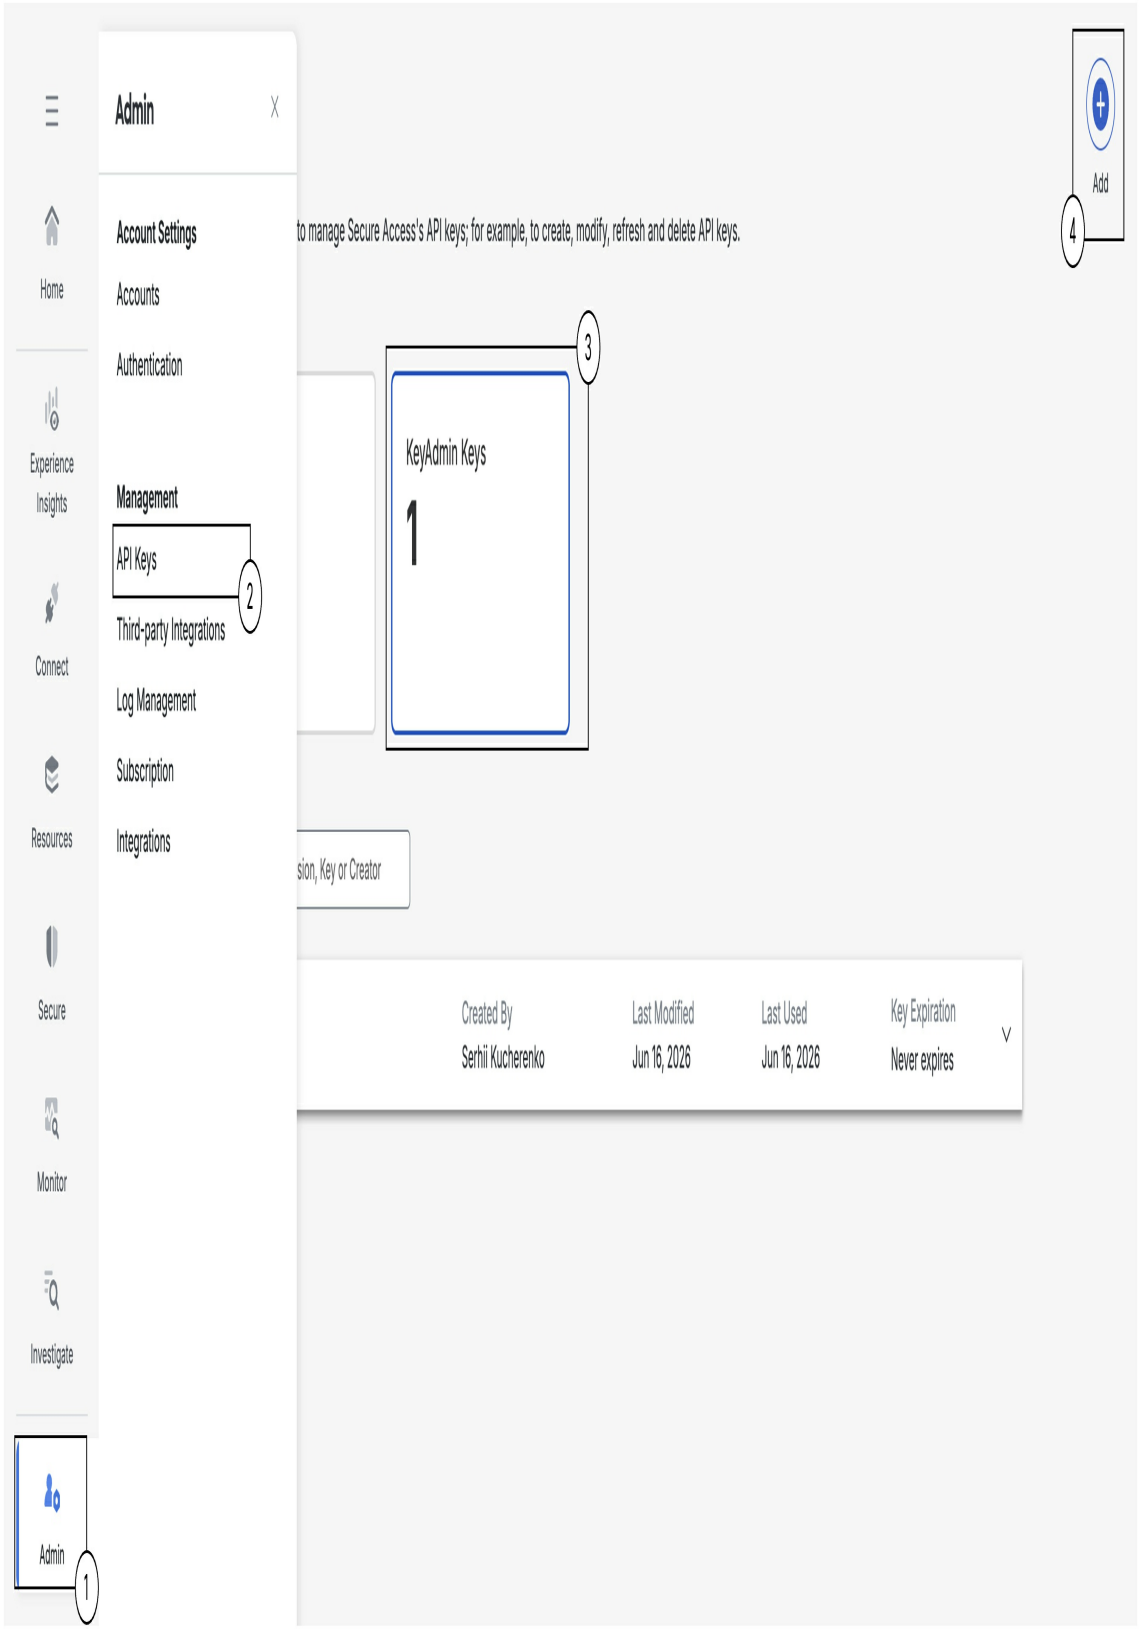

Admin
Account Settings
Accounts
Authentication
Management
API Keys
Third-party Integrations
Log Management
Subscription
Integrations
Home
Experience Insights
Connect
Resources
Secure
Monitor
Investigate
Admin
1
2
3
4
Add
to manage Secure Access's API keys; for example, to create, modify, refresh and delete API keys.
KeyAdmin Keys
1
sion, Key or Creator
Created By
Serhii Kucherenko
Last Modified
Jun 16, 2026
Last Used
Jun 16, 2026
Key Expiration
Never expires

Figure 13-55 *SSE API Keys Configuration Page*

As the next step, add permissions to the API key. For Active Directory integration, enable all available permissions, as shown in Figure 13-56. After you click **Create Key**, the key and secret will be displayed; be sure to save them in a secure location.

Technet24

KeyAdmin Key Name

AD-Key

Description *(Optional)*

Permissions

Select the areas of Secure Access API Key functionality that this key is allowed to manage.

☑ List API Keys	Read-Only
☑ Create API Keys	Write
☑ Update API Keys	Write
☑ Refresh API Keys	Write
☑ Delete API Keys	Write

Expiry Date

◉ Never expire

○ Expire on Sep 14 2026

Network Restrictions *(Optional)*

Optionally, add up to 10 networks from which this key can perform authentications. Add networks using a comma separated list of public IP addresses or CIDRs.

IP Addresses

For example: 100.10.10.0/24, 1.1.1.1

ADD

CANCEL

CREATE KEY

Figure 13-56 *SSE Admin API Key Configuration*

When the key and secret are saved, navigate back to **Connect > Users, Groups, and Endpoint Devices** and click **Configuration Management**. On the new page, go to the **Advanced Settings** tab, as shown in Figure 13-57. Insert your Admin API key and secret; then click **Generate Client API Key Pair**.

Technet24

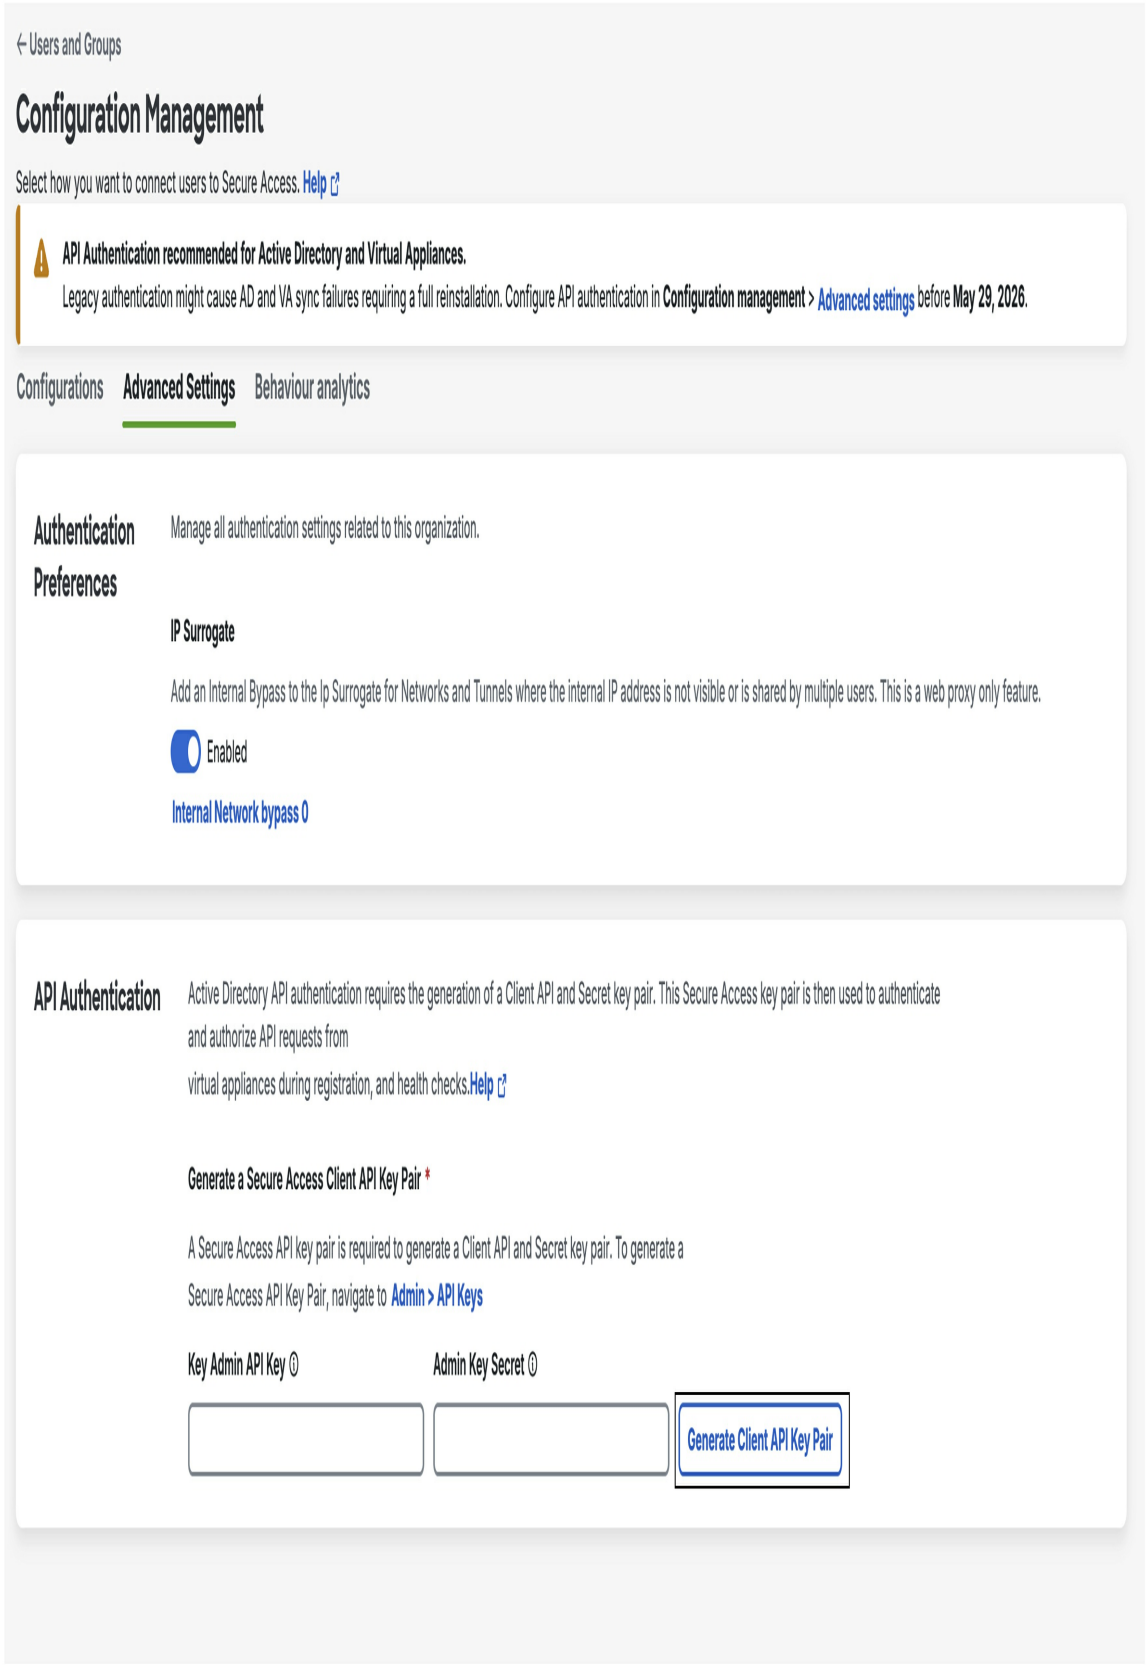

← Users and Groups
Configuration Management
Select how you want to connect users to Secure Access. Help
API Authentication recommended for Active Directory and Virtual Appliances.
Legacy authentication might cause AD and VA sync failures requiring a full reinstallation. Configure API authentication in Configuration management > Advanced settings before May 29, 2026.
Configurations
Advanced Settings
Behaviour analytics
Authentication Preferences
Manage all authentication settings related to this organization.
IP Surrogate
Add an Internal Bypass to the Ip Surrogate for Networks and Tunnels where the internal IP address is not visible or is shared by multiple users. This is a web proxy only feature.
Enabled
Internal Network bypass 0
API Authentication
Active Directory API authentication requires the generation of a Client API and Secret key pair. This Secure Access key pair is then used to authenticate and authorize API requests from
virtual appliances during registration, and health checks.Help
Generate a Secure Access Client API Key Pair *
A Secure Access API key pair is required to generate a Client API and Secret key pair. To generate a Secure Access API Key Pair, navigate to Admin > API Keys
Key Admin API Key
Admin Key Secret
Generate Client API Key Pair

Figure 13-57 *Adding an Admin API Key to Users and Groups Settings*

After adding the API keys, return to the Active Directory configuration. For this integration to function, install agent software on one of the domain controllers. The link to the executable and detailed installation instructions are available as part of the Active Directory integration process. Figure 13-58 provides an example of Active Directory configuration.

Technet24

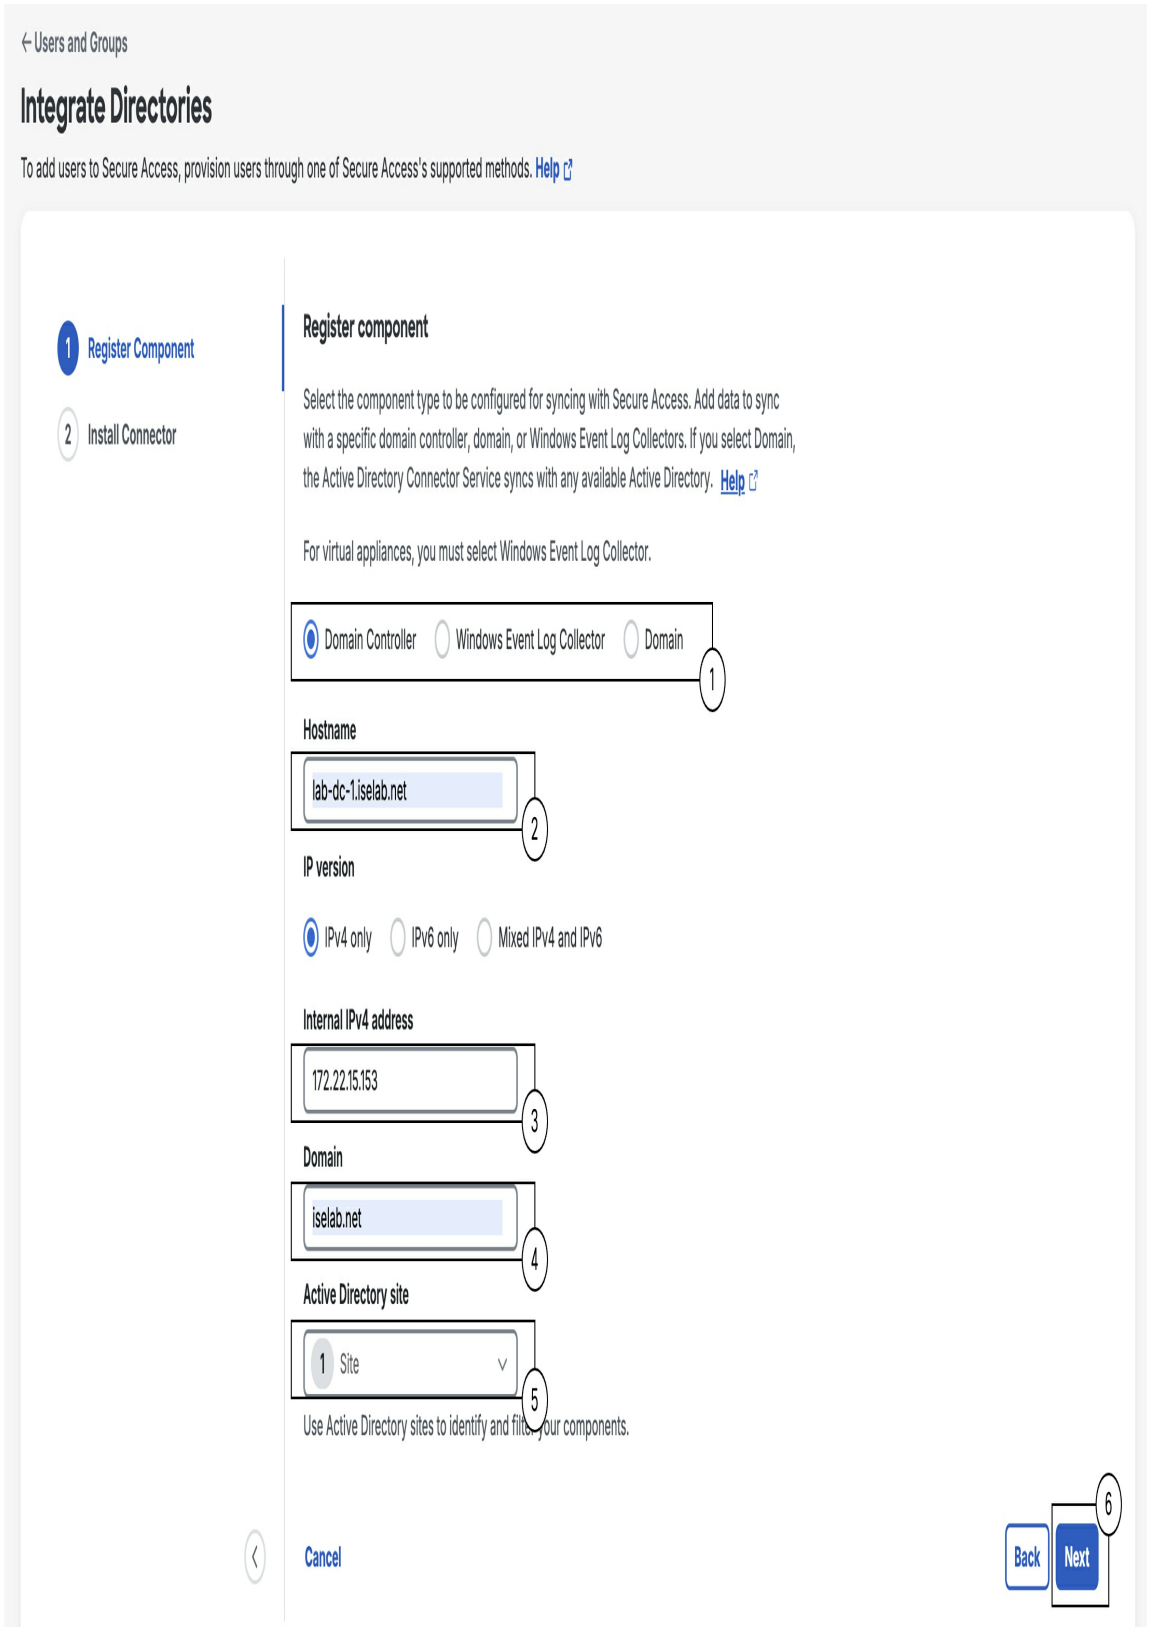

← Users and Groups
Integrate Directories
To add users to Secure Access, provision users through one of Secure Access's supported methods. Help
1 Register Component
2 Install Connector
Register component
Select the component type to be configured for syncing with Secure Access. Add data to sync with a specific domain controller, domain, or Windows Event Log Collectors. If you select Domain, the Active Directory Connector Service syncs with any available Active Directory. Help
For virtual appliances, you must select Windows Event Log Collector.
Domain Controller
Windows Event Log Collector
Domain
1
Hostname
lab-dc-1.iselab.net
2
IP version
IPv4 only
IPv6 only
Mixed IPv4 and IPv6
Internal IPv4 address
172.22.15.153
3
Domain
iselab.net
4
Active Directory site
1 Site
5
Use Active Directory sites to identify and filter your components.
Cancel
Back
Next
6

Figure 13-58 *Active Directory Configuration in SSE*

To complete Active Directory integration, use the following the steps:

Step 1. Define the component. For the Active Directory agent, you can choose either **Domain Controller (DC) —** in this case, provide the details of the DC — or use **Domain**, which allows the connector to select any available Active Directory.

Step 2. For the Domain Controller scenario, specify the DC FQDN.

Step 3. For the Domain Controller scenario, specify the DC IP address.

Step 4. Specify the domain for the integration.

Step 5. Choose the Active Directory site. By default, Microsoft Default-First-Site is used.

Step 6. Click **Next**.

On the next page, if the Active Directory connector has been installed properly, you should see information about the established connection. After returning to the **Users, Groups, and Endpoint Devices** page, in the **Users** tab, you should see the users who were synced with Active Directory. Figure 13-59 provides an example of the **Users** page.

Technet24


Users, Groups, and Endpoint Devices
Provision and manage the authentication of users, groups, and endpoint devices, for access control purposes. Help
Home
Users 10
Groups and Organizational Units 21
Endpoint Devices 3
Experience Insights
Connect
Resources
Secure
Monitor
Investigate
Users
Manage your organization's users and their devices connections and enrollments. To add users, go to Configuration management > Integrate directories
At any time, you can disconnect or unenroll a user's device. Help
Search
Source
Directory
Risk Alert
Trust Level
10 + results
Name
Risk Alerts
User Principal Name (UPN)
Auth Property
Source
Directory
Trust Level
Connected(VPN)
Enrolled(ZTNA)
Administrator - Administrator@iselab.net Administrator@iselab.net OnPrem Active Directory Profile - 0 0 -
alice - alice@iselab.net alice@iselab.net OnPrem Active Directory Profile - 0 0 -
bob - bob@iselab.net bob@iselab.net OnPrem Active Directory Profile - 0 0 -
bruce - bruce@iselab.net bruce@iselab.net OnPrem Active Directory - 0 0 -

Figure 13-59 *Users in SSE*

Configuring SSE Regions

RADIUS server groups required for remote access VPN authentication and authorization can be created separately in **Resources > Settings > AAA Servers** or defined directly during region configuration.

To define region settings, refer to Figure 13-60 and follow the steps:

Step 1. Click **Connect**.

Step 2. Click **End User Connectivity**.

Step 3. Open the **Virtual Private Network** tab.

Step 4. Click **Manage** in the **Regions and IP Pools** section.

Technet24

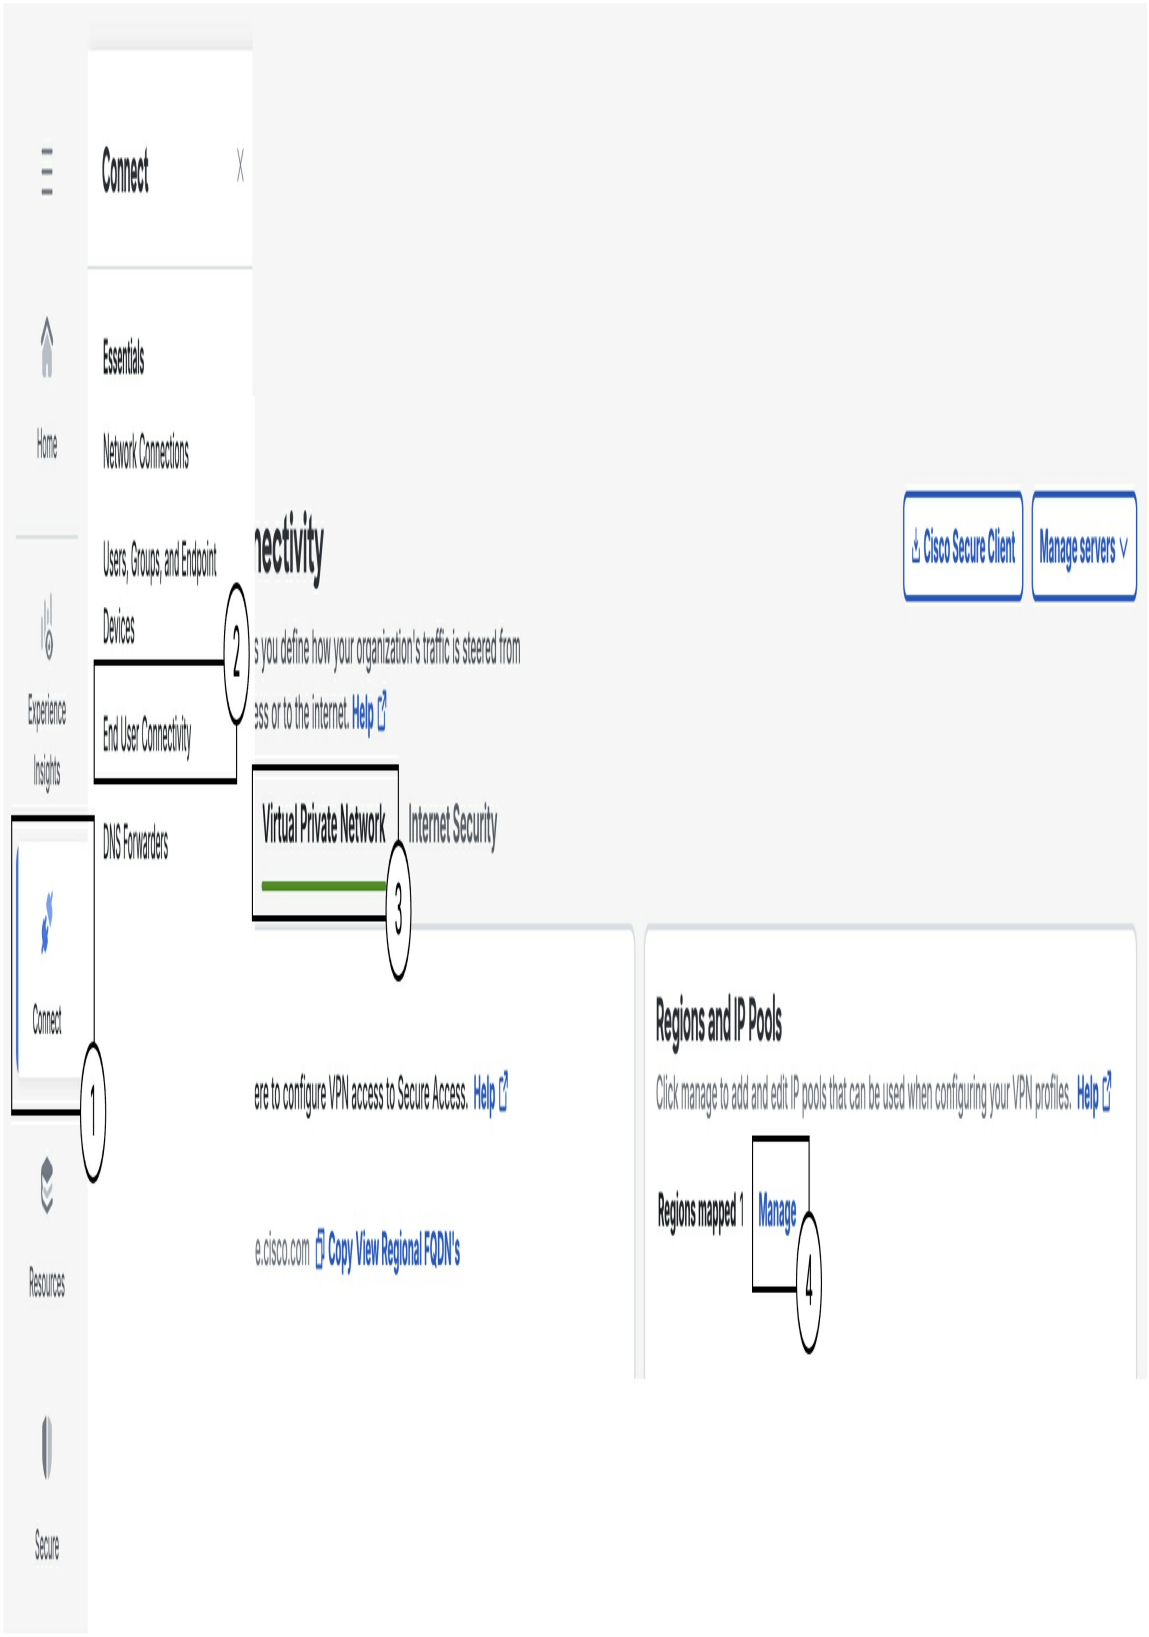

Connect
Essentials
Network Connections
Users, Groups, and Endpoint Devices
End User Connectivity
DNS Forwarders
Home
Experience Insights
Connect
Resources
Secure
nectivity
s you define how your organization's traffic is steered from
ess or to the internet. Help
Virtual Private Network
Internet Security
ere to configure VPN access to Secure Access. Help
e.cisco.com
Copy View Regional FQDN's
Cisco Secure Client
Manage servers
Regions and IP Pools
Click manage to add and edit IP pools that can be used when configuring your VPN profiles. Help
Regions mapped 1
Manage
1
2
3
4

Figure 13-60 *Regions and IP Pools*

Referencing Figure 13-61, on the next page, do the following:

Step 1. Click **Add IP Pool**.

Technet24

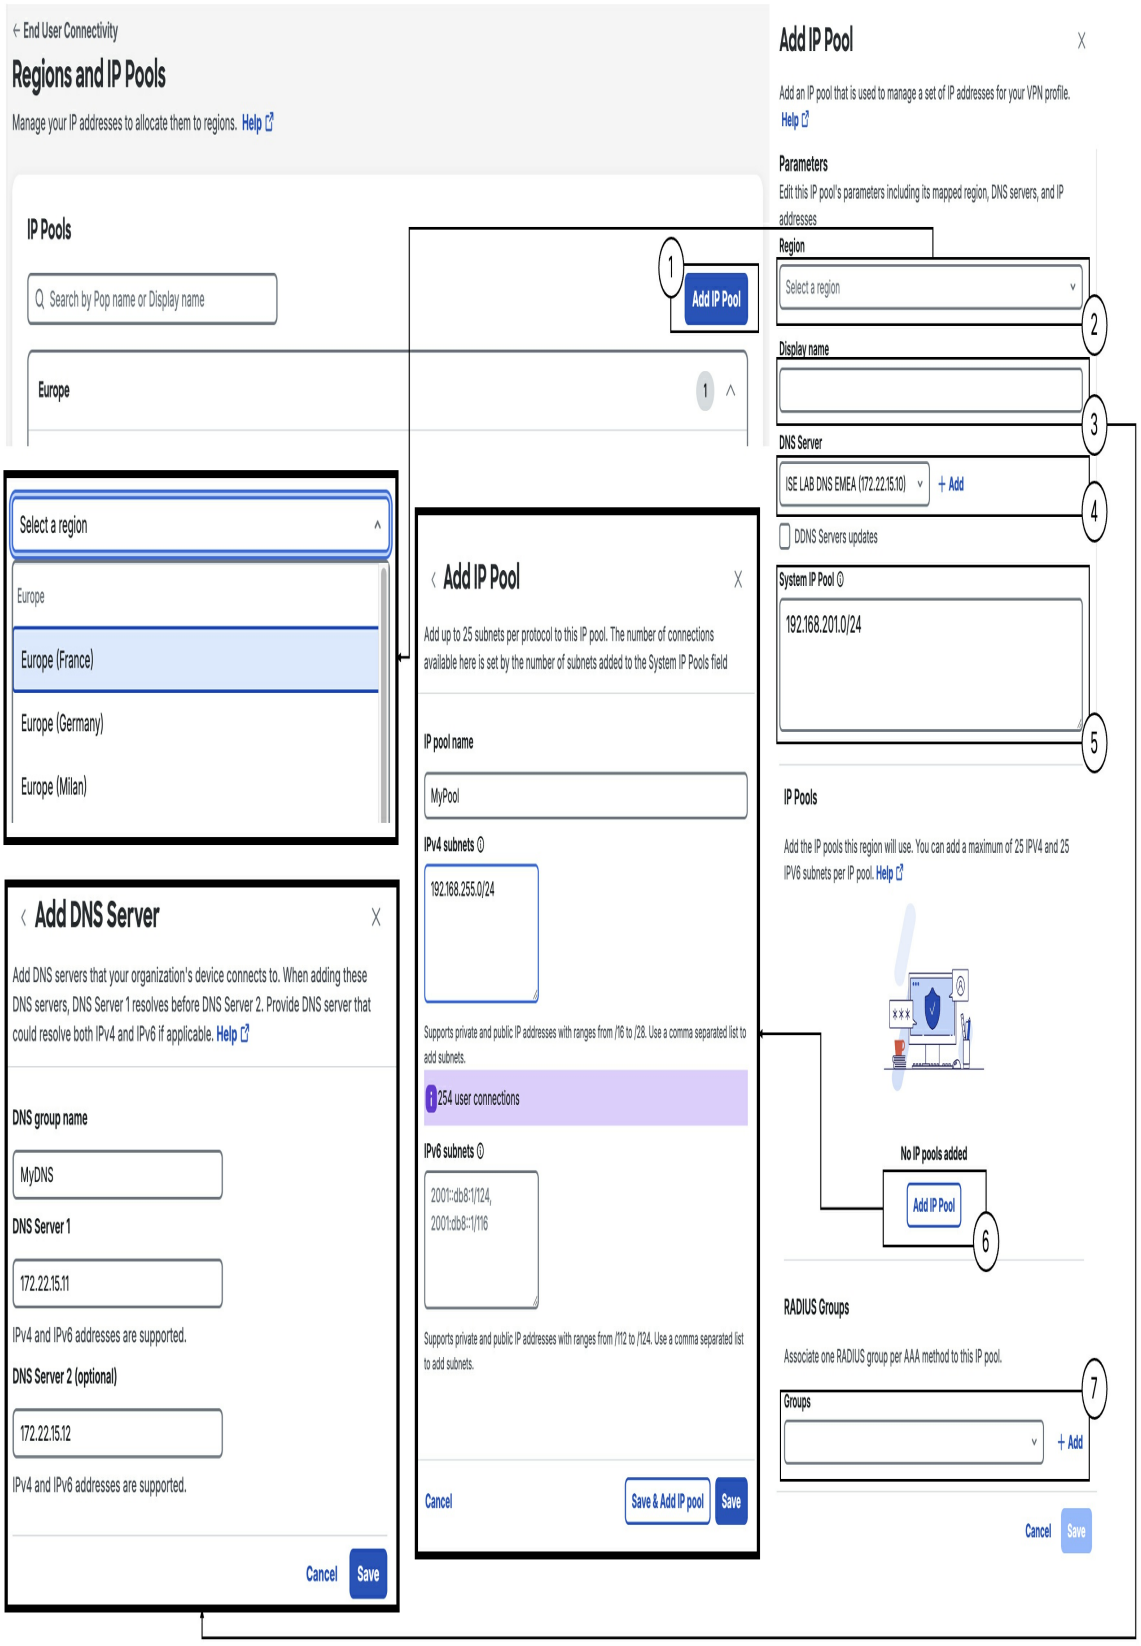

← End User Connectivity
Regions and IP Pools
Manage your IP addresses to allocate them to regions. Help
IP Pools
Search by Pop name or Display name
Europe
1
Add IP Pool
Add IP Pool
Add an IP pool that is used to manage a set of IP addresses for your VPN profile.
Help
Parameters
Edit this IP pool's parameters including its mapped region, DNS servers, and IP addresses
Region
Select a region
Display name
DNS Server
ISE LAB DNS EMEA (172.22.15.10)
+ Add
DDNS Servers updates
System IP Pool
192.168.201.0/24
IP Pools
Add the IP pools this region will use. You can add a maximum of 25 IPV4 and 25 IPV6 subnets per IP pool. Help
No IP pools added
Add IP Pool
RADIUS Groups
Associate one RADIUS group per AAA method to this IP pool.
Groups
+ Add
Cancel
Save
1
2
3
4
5
6
7
Select a region
Europe
Europe (France)
Europe (Germany)
Europe (Milan)
Add IP Pool
Add up to 25 subnets per protocol to this IP pool. The number of connections available here is set by the number of subnets added to the System IP Pools field
IP pool name
MyPool
IPv4 subnets
192.168.255.0/24
Supports private and public IP addresses with ranges from /16 to /28. Use a comma separated list to add subnets.
254 user connections
IPv6 subnets
2001:db8:1/124,
2001:db8::1/116
Supports private and public IP addresses with ranges from /112 to /124. Use a comma separated list to add subnets.
Cancel
Save & Add IP pool
Save
Add DNS Server
Add DNS servers that your organization's device connects to. When adding these DNS servers, DNS Server 1 resolves before DNS Server 2. Provide DNS server that could resolve both IPv4 and IPv6 if applicable. Help
DNS group name
MyDNS
DNS Server 1
172.22.15.11
IPv4 and IPv6 addresses are supported.
DNS Server 2 (optional)
172.22.15.12
IPv4 and IPv6 addresses are supported.
Cancel
Save

Figure 13-61 *Region IP Pool Configuration*

Step 2. Select the SSE region from the drop-down list.

Step 3. Specify a friendly name for the IP pool.

Step 4. Select a DNS configuration from the drop-down list or click **+Add** to create a new configuration. In the DNS configuration, specify the configuration name and provide IP addresses of the primary and secondary DNS servers.

Step 5. Specify the network/prefix for the System IP pool. The System IP pool is used as the source IP addresses for RADIUS traffic and other types of control plane communications.

Step 6. Configure the IP pool. This pool is used to assign IP addresses to endpoints. Specify the IP pool name and define IPv4 and IPv6 subnets for IP assignment.

Step 7. Select a RADIUS server group from the drop-down list or add a new one by clicking **+Add**.

The RADIUS server group configuration includes multiple steps, outlined in the following figures.

Following along with Figure 13-62, complete these steps:

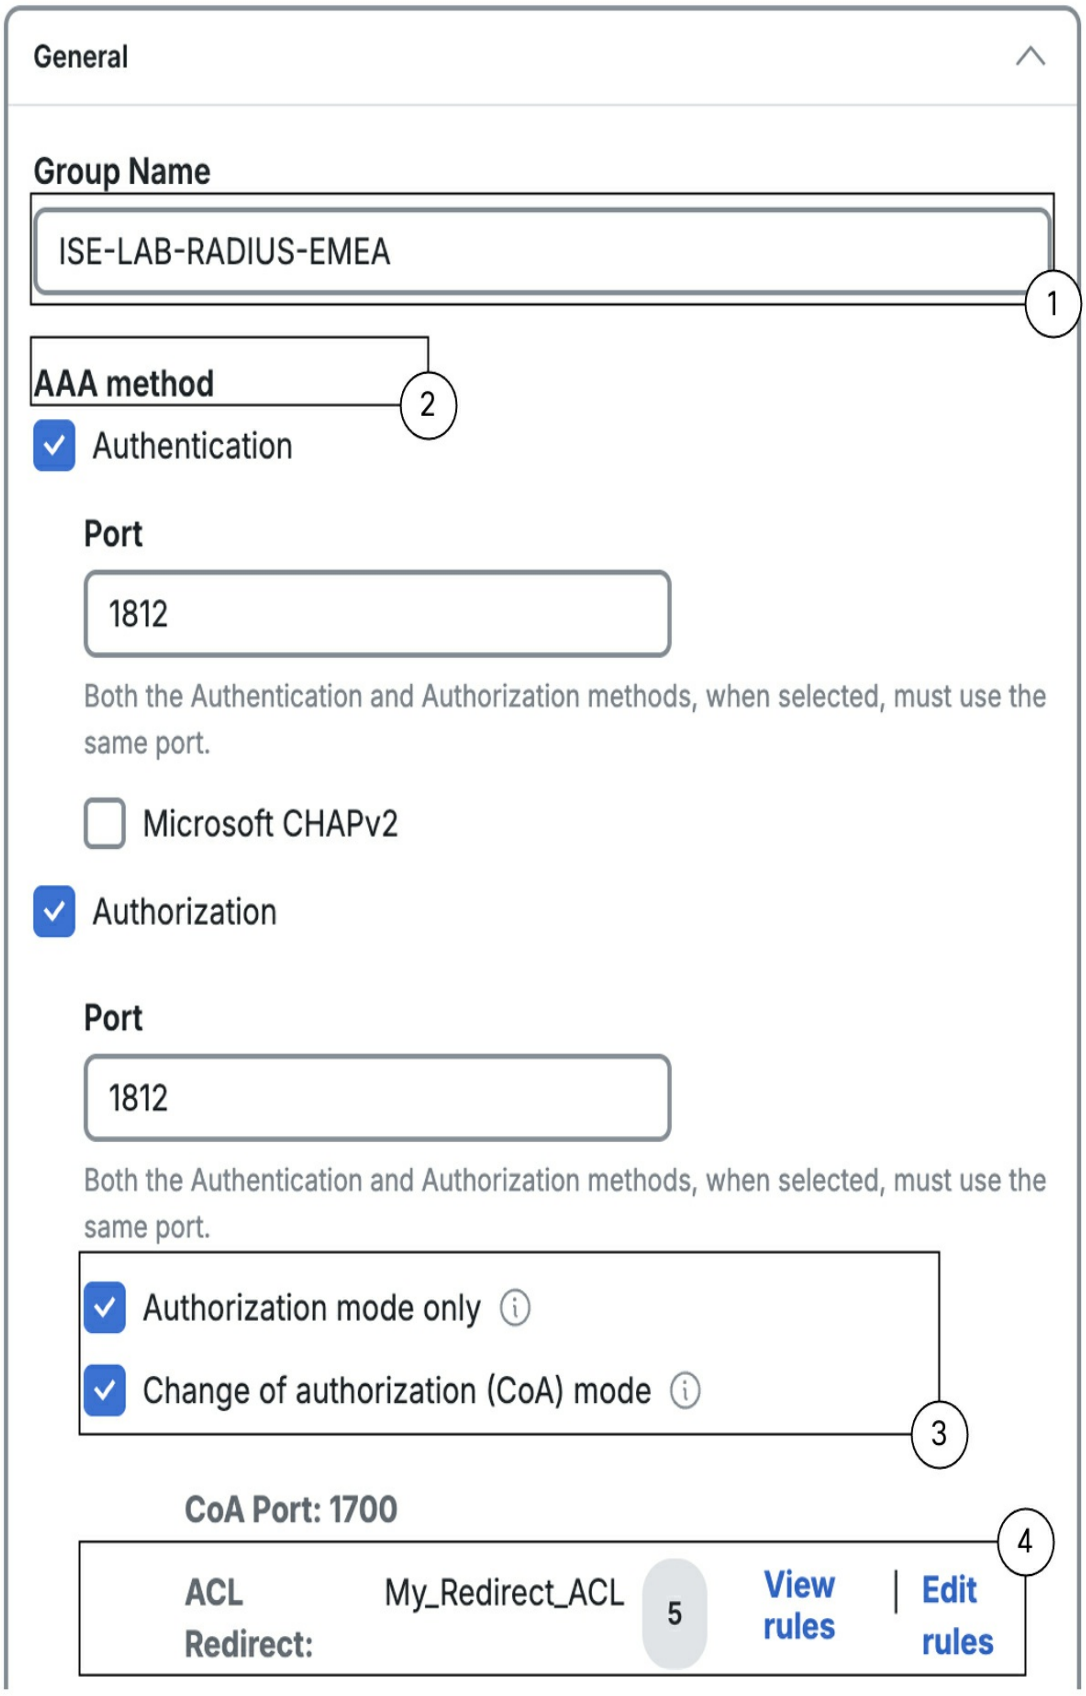
Step 1. Provide the RADIUS server group name.

Step 2. Enable authentication, authorization, and accounting in the **AAA Method** section.

Step 3. Enable **Authorization Mode Only** in the **Authorization** section, which is required for certificate-based and SAML-based authentication, and enable **Change of Authorization (CoA) Mode** to support posture over VPN.

Technet24

Step 4. Additionally, if needed, you can view or edit the redirect ACL used with the RADIUS server group. This ACL is automatically populated with Deny statements for the RADIUS servers and DNS servers defined in the Region IP pool, ensuring that ISE PSNs used in the region are excluded from redirection, as well as DNS traffic.


General
Group Name
ISE-LAB-RADIUS-EMEA
1
AAA method
2
Authentication
Port
1812
Both the Authentication and Authorization methods, when selected, must use the same port.
Microsoft CHAPv2
Authorization
Port
1812
Both the Authentication and Authorization methods, when selected, must use the same port.
Authorization mode only
Change of authorization (CoA) mode
3
CoA Port: 1700
4
ACL Redirect:
My_Redirect_ACL
5
View rules
Edit rules

Technet24

Figure 13-62 *RADIUS Group Settings Part 1*

The default state of the Redirect ACL is shown in Figure 13-63. If you wish to edit the default redirect ACL, clicking the **Edit** button brings you to the **Resources > AAA Servers** page, where you can start editing the default ACL by clicking **Manage CoA ACL Redirect Rules**.

ACL redirect rules

Rules defined for ACL redirect. To add new rules or edit the existing ones, go to ACL redirect rules page.

#	Action	Service	Source	Destination	Port	Comment
01	Deny	UDP	Any	Any	domain	-
02	Deny	ICMP	Any	Any	-	-
03	Deny	IP	Any	RADIUS Servers	-	-
04	Allow	TCP	Any	Any	www	-
05	Allow	TCP	Any	Any	https	-

Close

AAA Servers

Manage CoA ACL redirect rules

Authentication, authorization and accounting (AAA) servers are used to control access to resources, private apps and the internet. Once added to Secure Access, AAA servers can be added to a RADIUS group and used to authenticate VPN profiles. Help

RADIUS Groups | RADIUS Servers

Search by group name

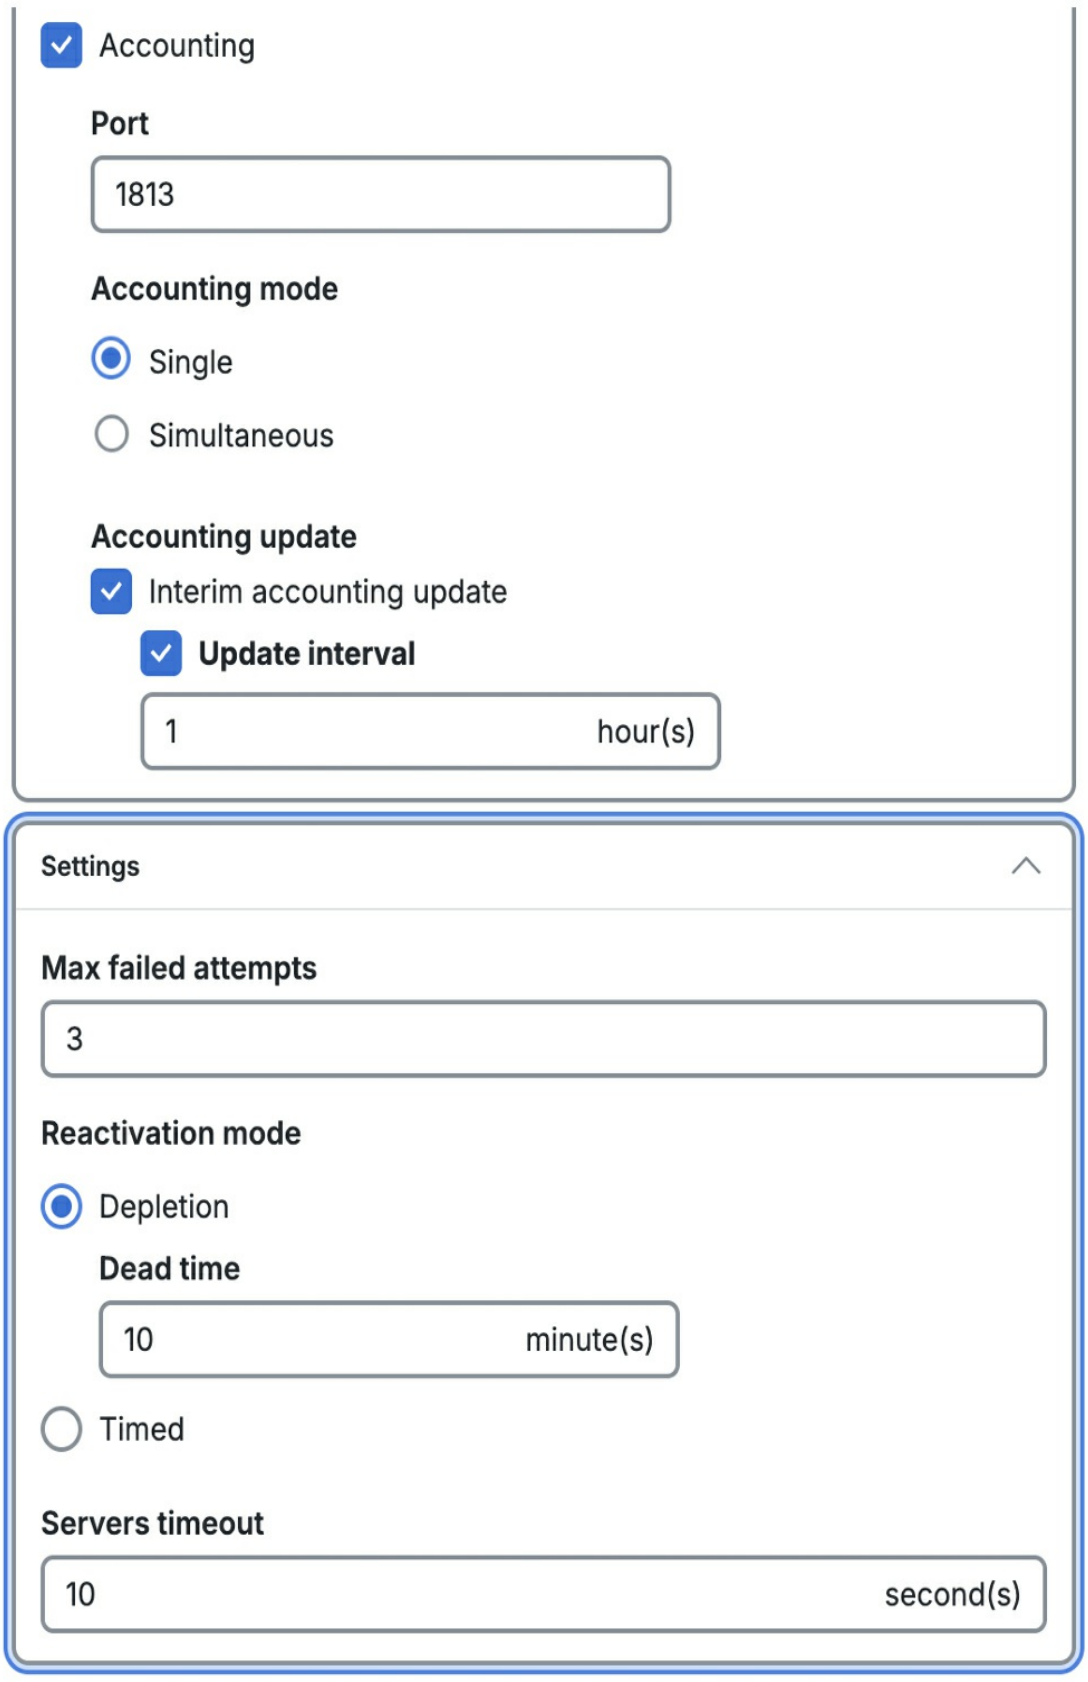
Add group

Group Name	AAA Methods	Servers in Group	Regions	VPN Profiles
ISE-LAB-RADIUS-EMEA	Authentication, Authorization, Accounting	2	ISE-LAB-EMEA-UK	ISE-LAB-UK-RA-VPN

Technet24

Figure 13-63 *Redirect ACL View and Edit Options*

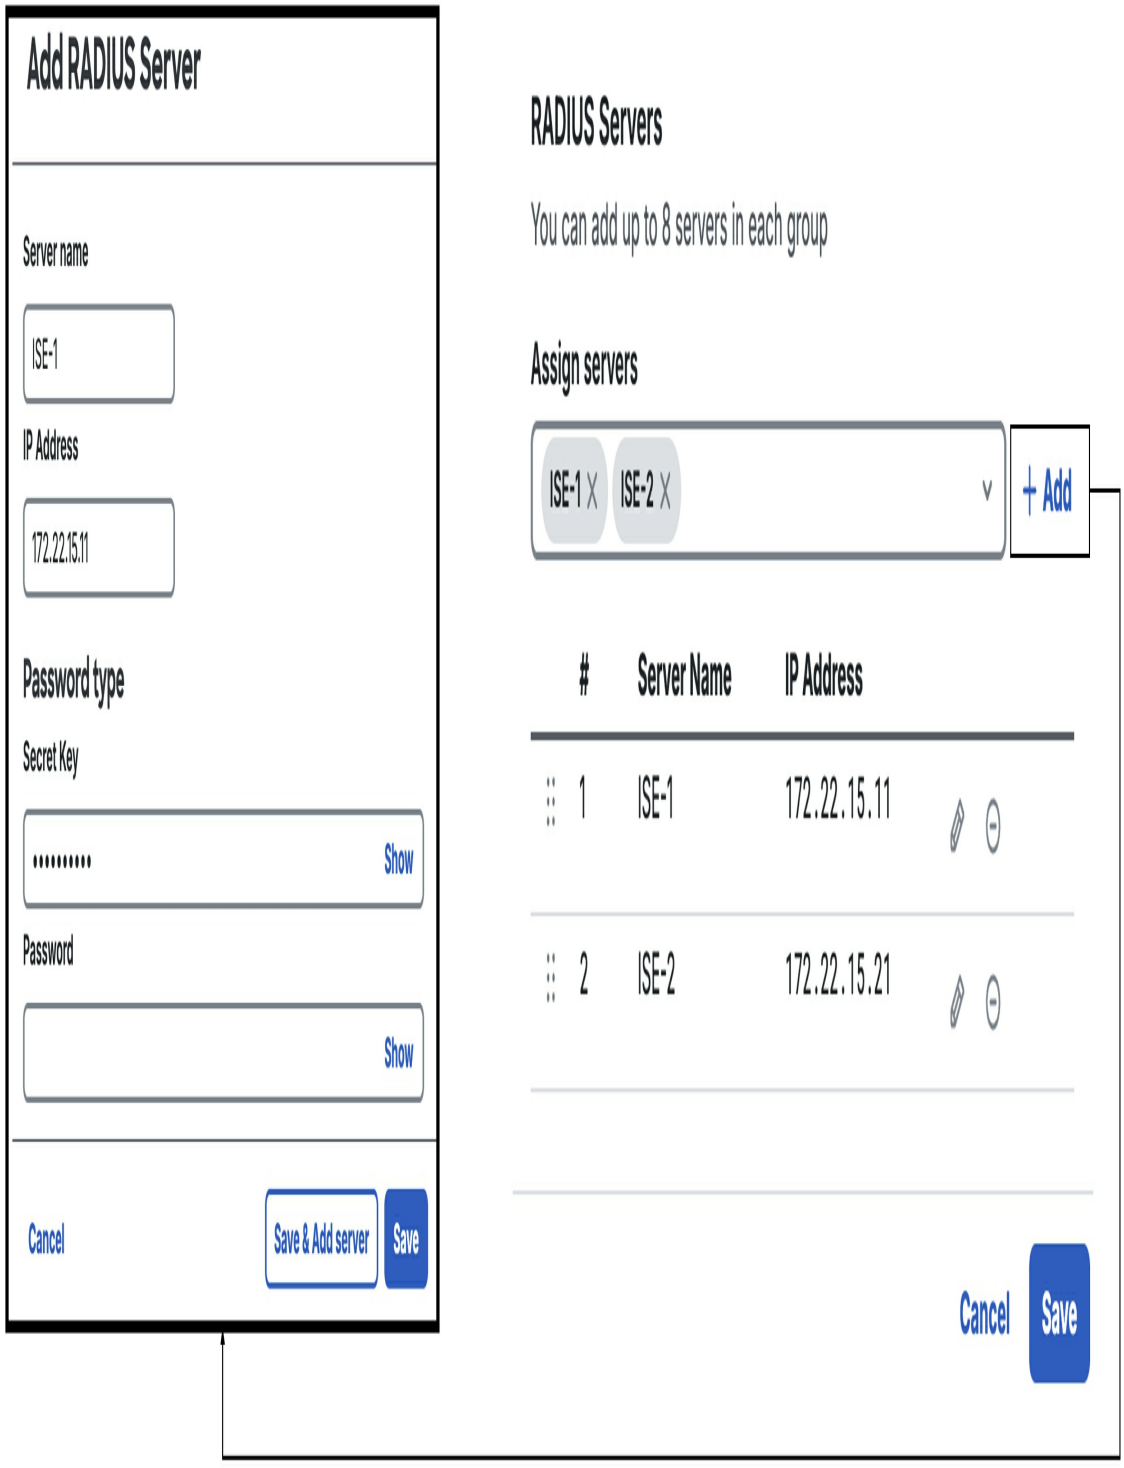
By scrolling down in the RADIUS server group settings, you can define accounting parameters such as enabling interim accounting updates and setting the update interval, as well as configure RADIUS server dead detection parameters. These settings are shown in Figure 13-64.


Accounting
Port
1813
Accounting mode
Single
Simultaneous
Accounting update
Interim accounting update
Update interval
1
hour(s)
Settings
Max failed attempts
3
Reactivation mode
Depletion
Dead time
10
minute(s)
Timed
Servers timeout
10
second(s)

Technet24

Figure 13-64 *RADIUS Group Settings Part 2*

As the final step in Region IP pool configuration, add RADIUS servers to the RADIUS server group, as shown in Figure 13-65.


Add RADIUS Server
Server name
ISE-1
IP Address
172.22.15.11
Password type
Secret Key
Show
Password
Show
Cancel
Save & Add server
Save
RADIUS Servers
You can add up to 8 servers in each group
Assign servers
ISE-1 X
ISE-2 X
Add
Server Name IP Address
1 ISE-1 172.22.15.11
2 ISE-2 172.22.15.21
Cancel
Save

Technet24

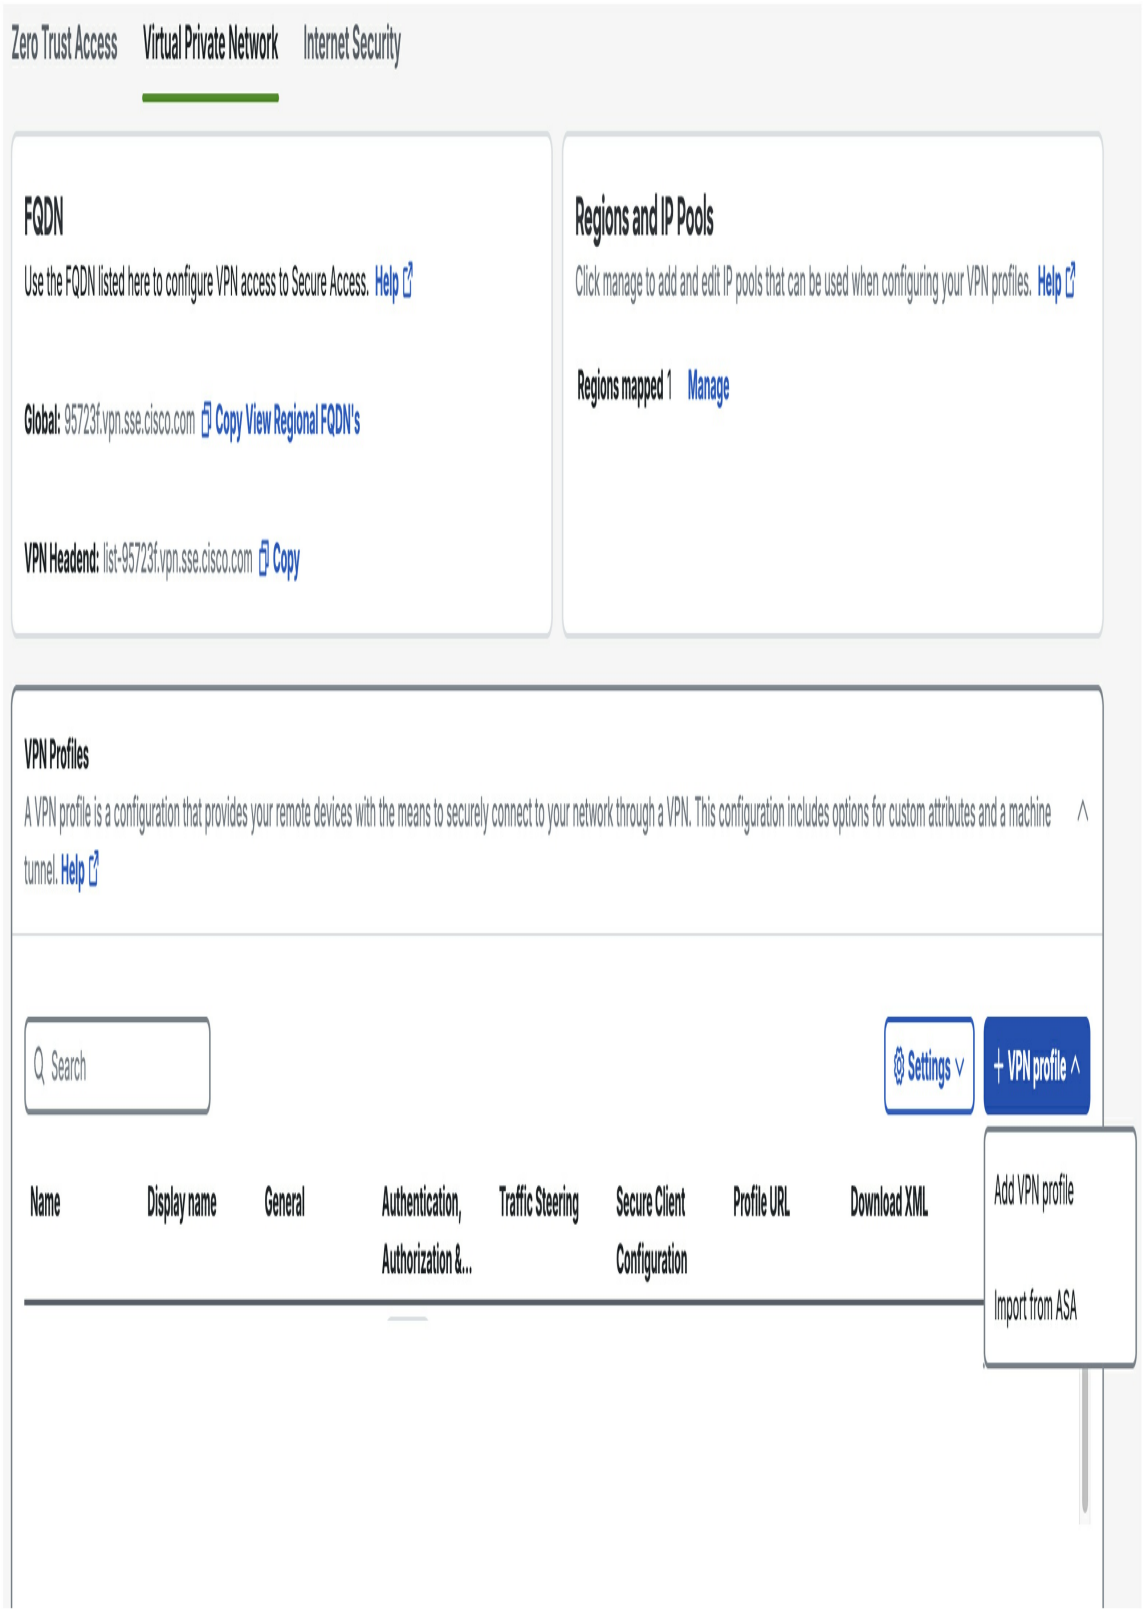
Figure 13-65 *RADIUS Group Settings Part 3*

SSE: Putting Everything Together

This section describes how to connect previously created components. On the SSE side, all necessary elements are linked together as part of the VPN profile configuration, and on the ISE side, the required configuration steps for SSE are explained.

VPN Profile Configuration

On the SSE side, the next step after synchronizing users with one of the supported directories and defining a regional IP pool that includes a RADIUS group is to proceed with VPN profile configuration. Navigate to **Connect > End User Connectivity**, open the **Virtual Private Network** tab, and click **+ Add VPN Profile**. Select **Add VPN Profile** from the drop-down menu, as shown in Figure 13-66.


Zero Trust Access
Virtual Private Network
Internet Security
FQDN
Use the FQDN listed here to configure VPN access to Secure Access. Help
Global: 95723f.vpn.sse.cisco.com Copy View Regional FQDN's
VPN Headend: list-95723f.vpn.sse.cisco.com Copy
Regions and IP Pools
Click manage to add and edit IP pools that can be used when configuring your VPN profiles. Help
Regions mapped 1 Manage
VPN Profiles
A VPN profile is a configuration that provides your remote devices with the means to securely connect to your network through a VPN. This configuration includes options for custom attributes and a machine tunnel. Help
Search
Settings
+ VPN profile
Name
Display name
General
Authentication, Authorization &...
Traffic Steering
Secure Client Configuration
Profile URL
Download XML
Add VPN profile
Import from ASA

Technet24

Figure 13-66 *Adding a VPN Profile*

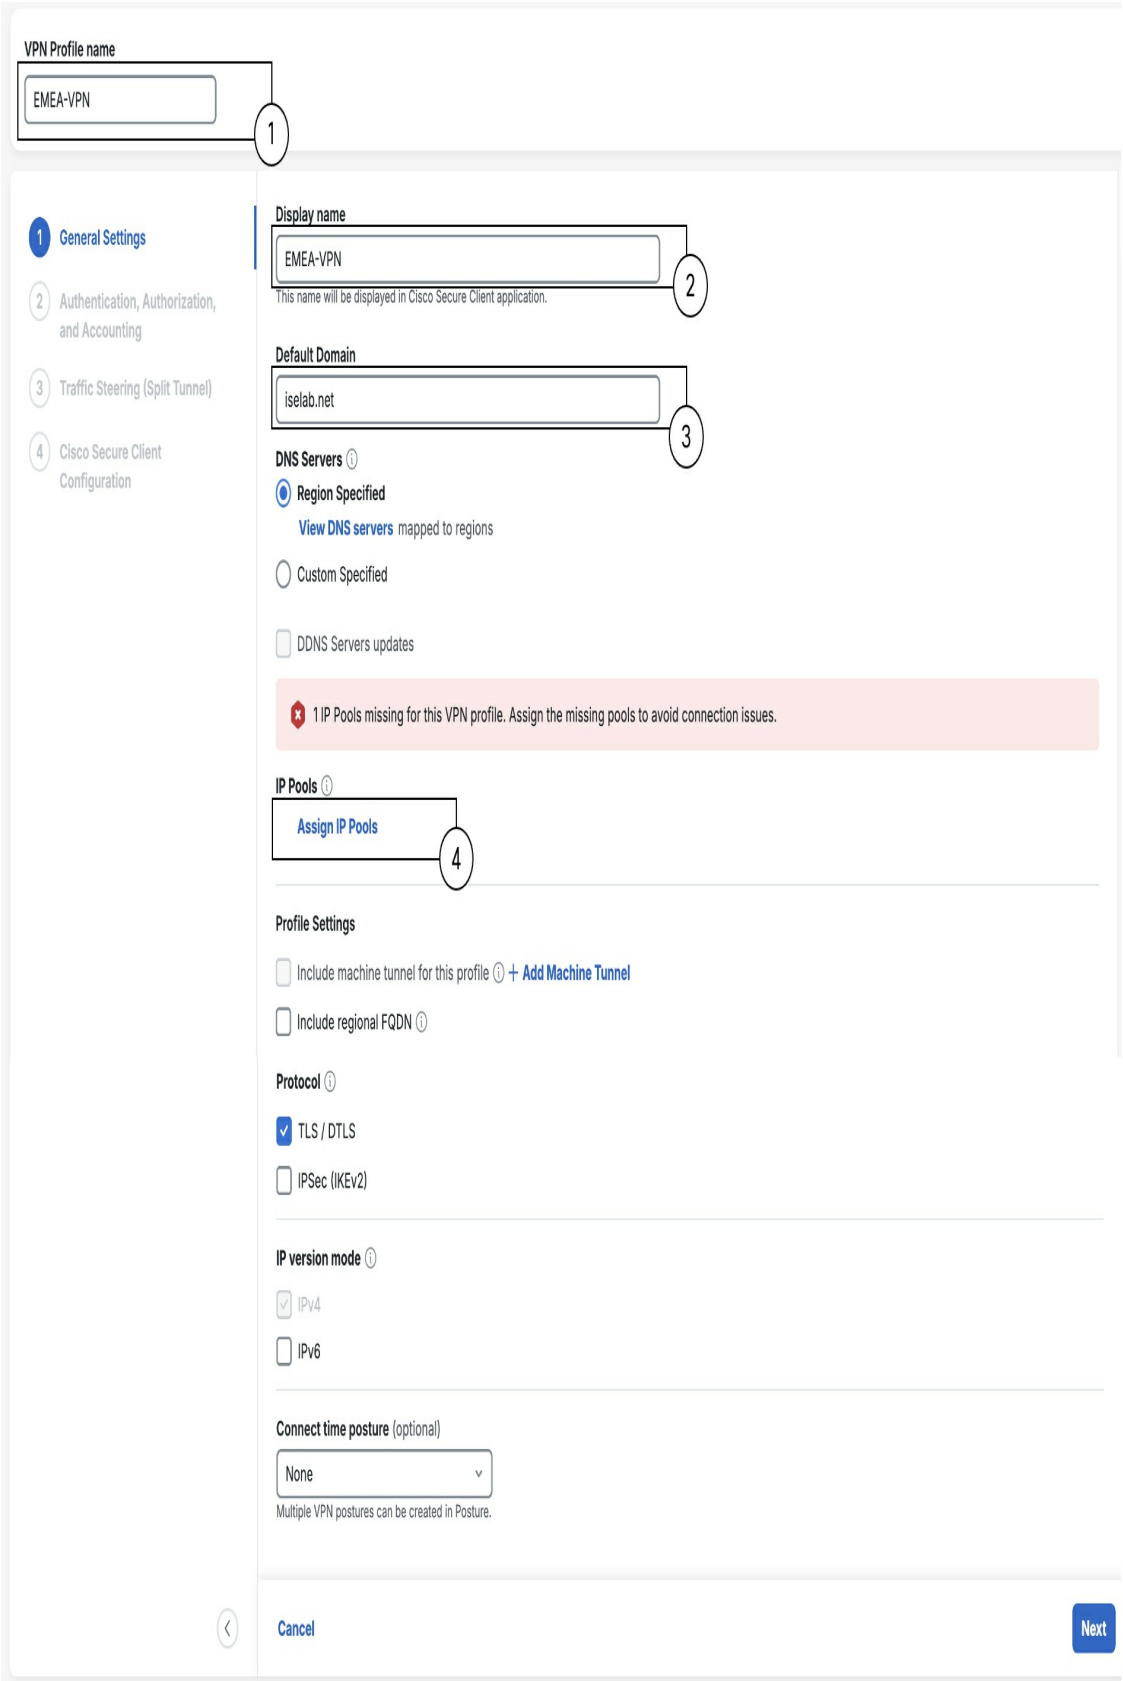
In the **General Settings** section of the profile configuration (refer to Figure 13-67), follow these steps:

Step 1. Define the profile name in **VPN Profile Name**.

Step 2. Define the profile name that will be displayed in Secure Client as the **Display Name**.

Step 3. Set the **Default Domain** name for the connection.

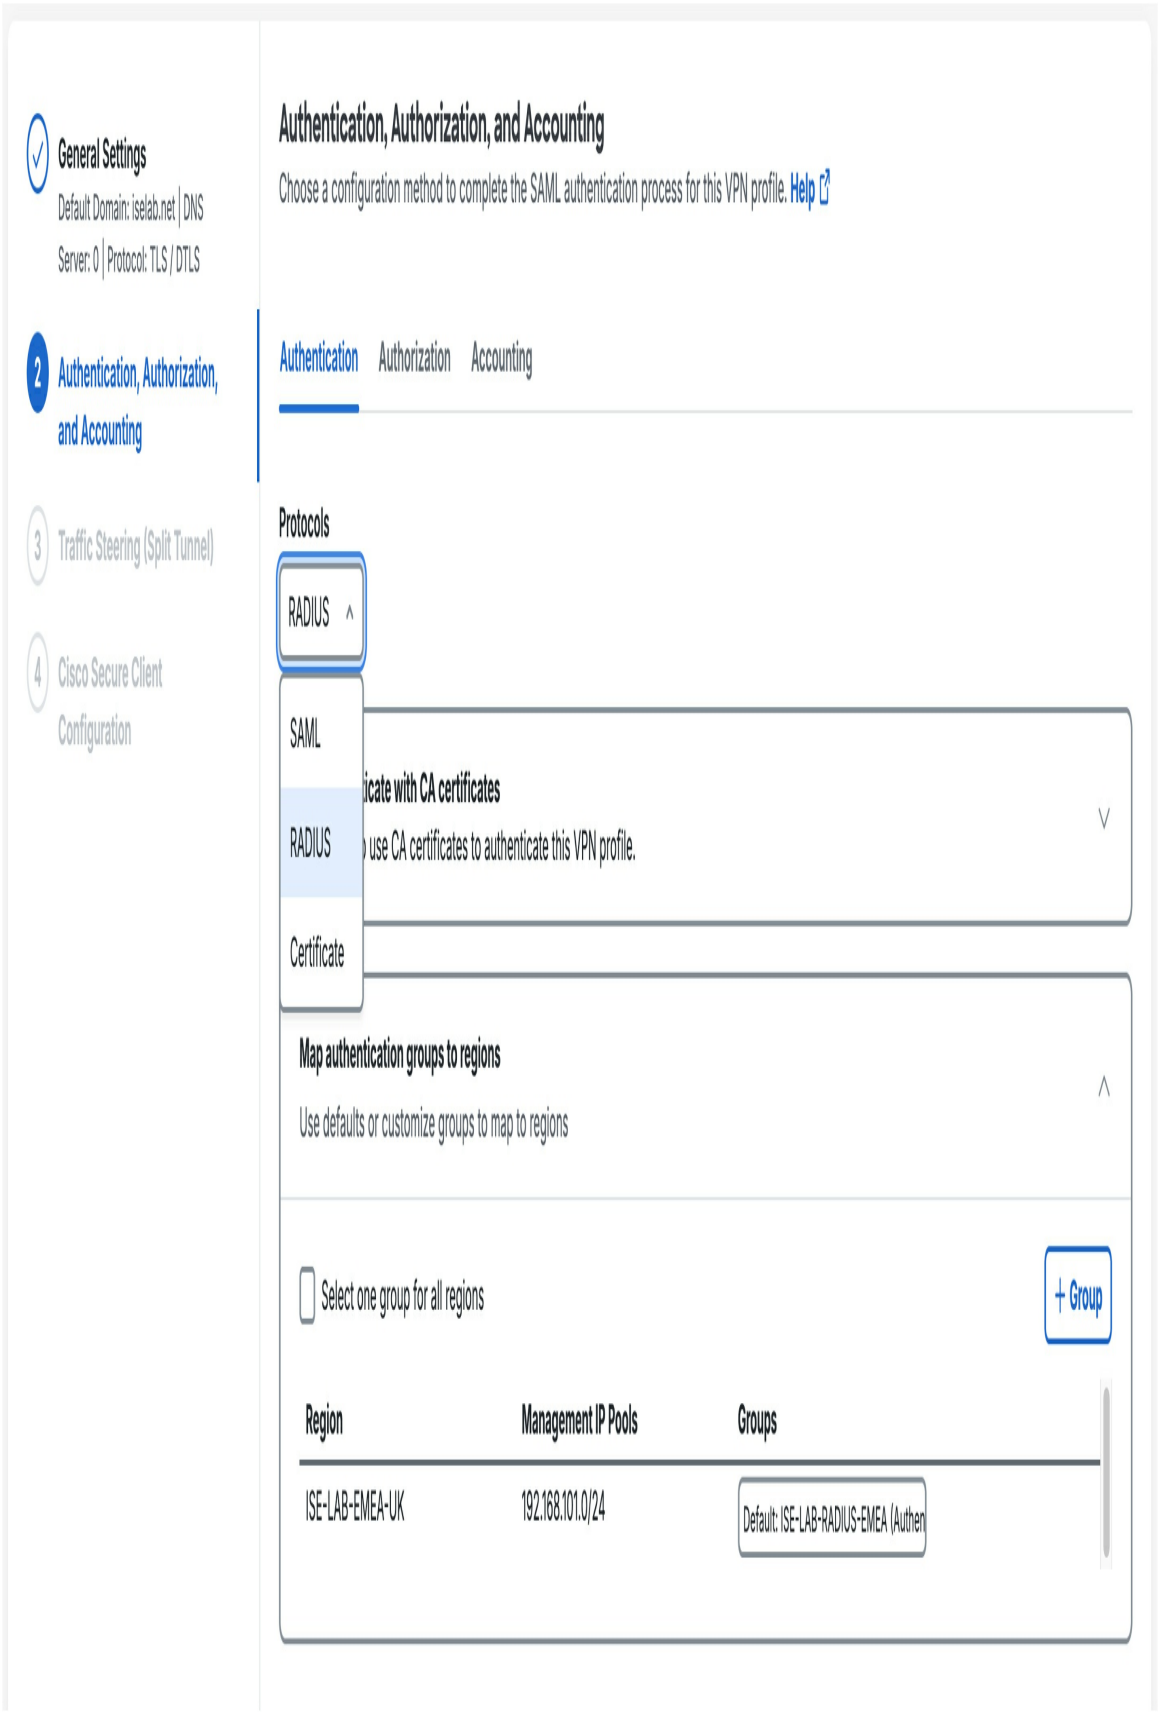
Step 4. Select the previously configured pool in **IP Pools**.


VPN Profile name
EMEA-VPN
1
1 General Settings
2 Authentication, Authorization, and Accounting
3 Traffic Steering (Split Tunnel)
4 Cisco Secure Client Configuration
Display name
EMEA-VPN
This name will be displayed in Cisco Secure Client application.
2
Default Domain
iselab.net
3
DNS Servers
Region Specified
View DNS servers mapped to regions
Custom Specified
DDNS Servers updates
1 IP Pools missing for this VPN profile. Assign the missing pools to avoid connection issues.
IP Pools
Assign IP Pools
4
Profile Settings
Include machine tunnel for this profile
+ Add Machine Tunnel
Include regional FQDN
Protocol
TLS / DTLS
IPSec (IKEv2)
IP version mode
IPv4
IPv6
Connect time posture (optional)
None
Multiple VPN postures can be created in Posture.
Cancel
Next

Technet24

Figure 13-67 *VPN Profile General Settings*

Next, configure authentication, authorization, and accounting settings. The combination of authentication/authorization options is similar to those available for Secure Firewall. Figure 13-68 shows RADIUS authentication for the VPN connection.


General Settings
Default Domain: iselab.net | DNS Server: 0 | Protocol: TLS / DTLS
2
Authentication, Authorization, and Accounting
3
Traffic Steering (Split Tunnel)
4
Cisco Secure Client Configuration
Authentication, Authorization, and Accounting
Choose a configuration method to complete the SAML authentication process for this VPN profile. Help
Authentication
Authorization
Accounting
Protocols
RADIUS
SAML
RADIUS
Certificate
icate with CA certificates
use CA certificates to authenticate this VPN profile.
Map authentication groups to regions
Use defaults or customize groups to map to regions
Select one group for all regions
+ Group
Region
Management IP Pools
Groups
ISE-LAB-EMEA-UK
192.168.101.0/24
Default: ISE-LAB-RADIUS-EMEA (Authen

Technet24

Figure 13-68 *VPN Profile Authentication Settings*

Now, finish profile configuration by setting up Traffic Steering and Cisco Secure Client configuration.

ISE Configuration for SSE

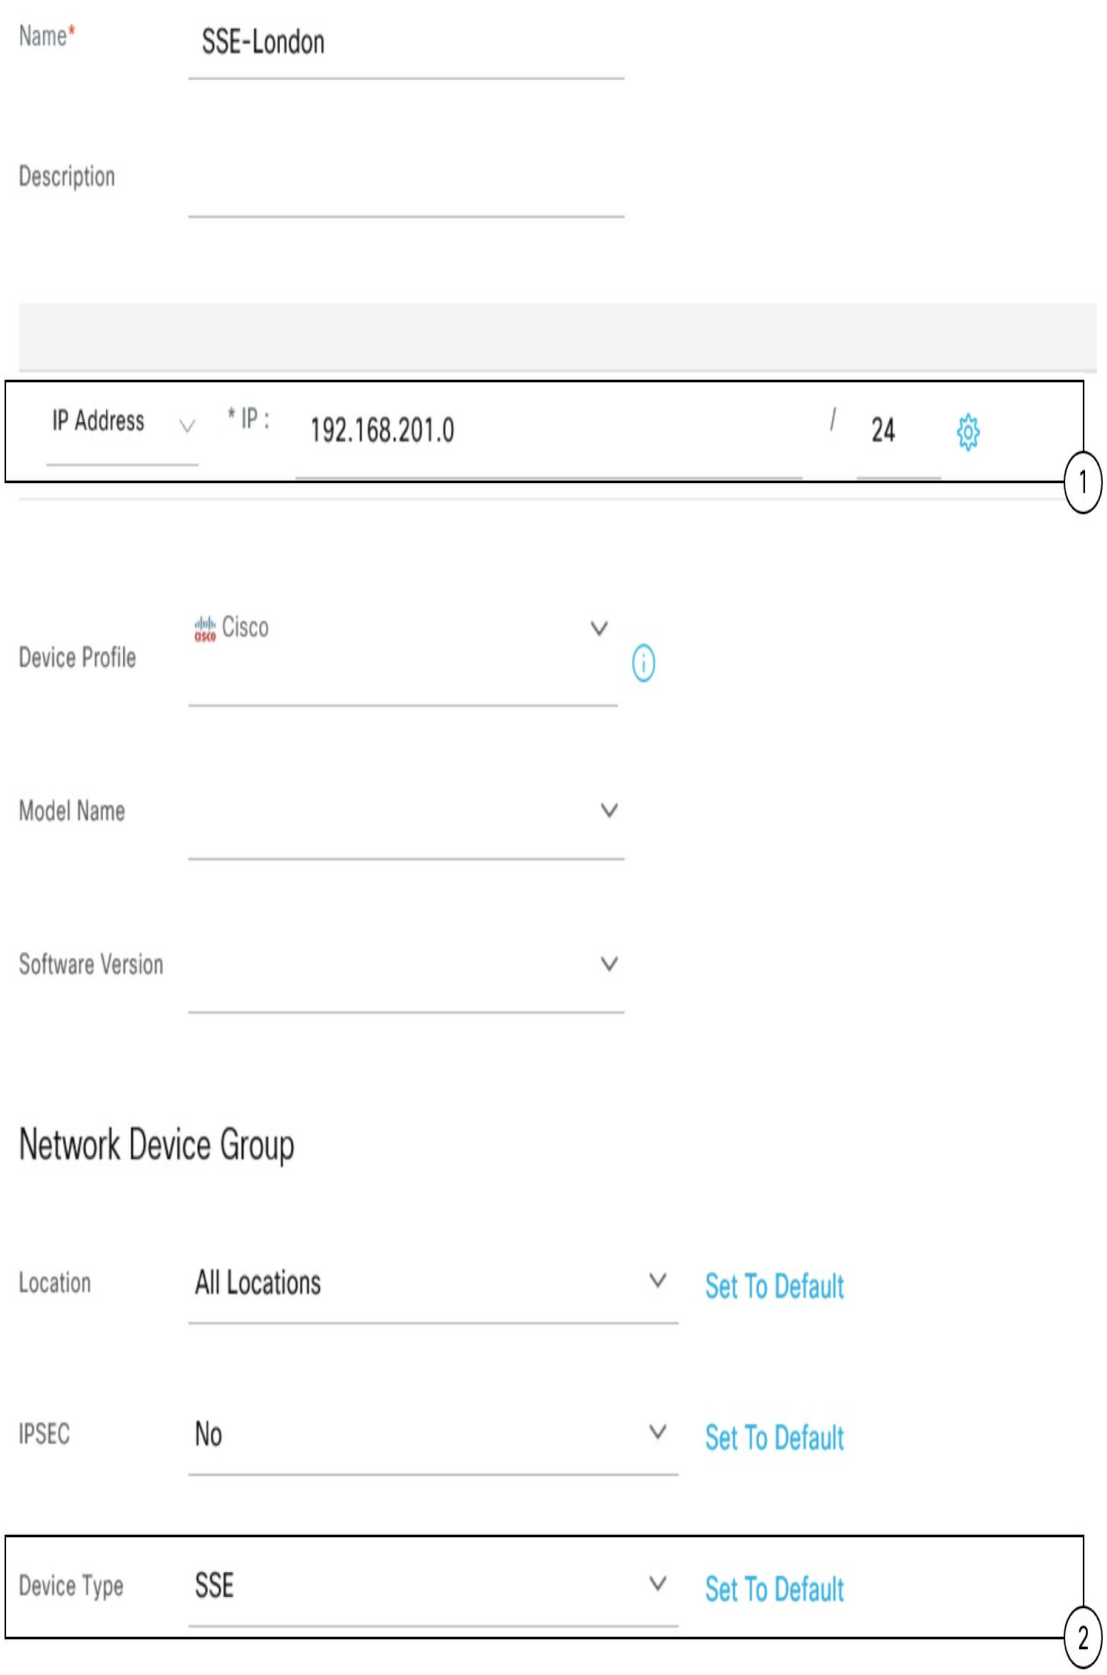
When configuring ISE for integration with SSE, pay special attention to two aspects of the ISE configuration:

- **Network Access Device (NAD) for SSE:** The source IP address of the NAD will be within the range configured at the Region IP pool level as the System IP pool. Ensure that the NAD IP for SSE covers the entire range.
- **Authentication and Authorization Policies:** Configure these policies similarly to those used for Secure Firewall. If Authorize-Only use cases are required, ensure this requirement is reflected in the authentication and authorization policy settings.

An example of SSE NAD configuration is shown in Figure 13-69. During NAD configuration, follow these steps:

Step 1. Define the System IP pool range as an IP address.

Step 2. You may also consider creating and assigning a specific device type to all SSE NADs to later use it for steering all requests coming from SSE NADs into a single policy set.


Name*
SSE-London
Description
IP Address
* IP :
192.168.201.0
/
24
1
Device Profile
Cisco
Model Name
Software Version
Network Device Group
Location
All Locations
Set To Default
IPSEC
No
Set To Default
Device Type
SSE
Set To Default
2

Technet24

Figure 13-69 *SSE NAD Configuration on ISE*

Next, set up the authentication and authorization policy within the policy set according to how authentication and authorization are configured on the SSE side.

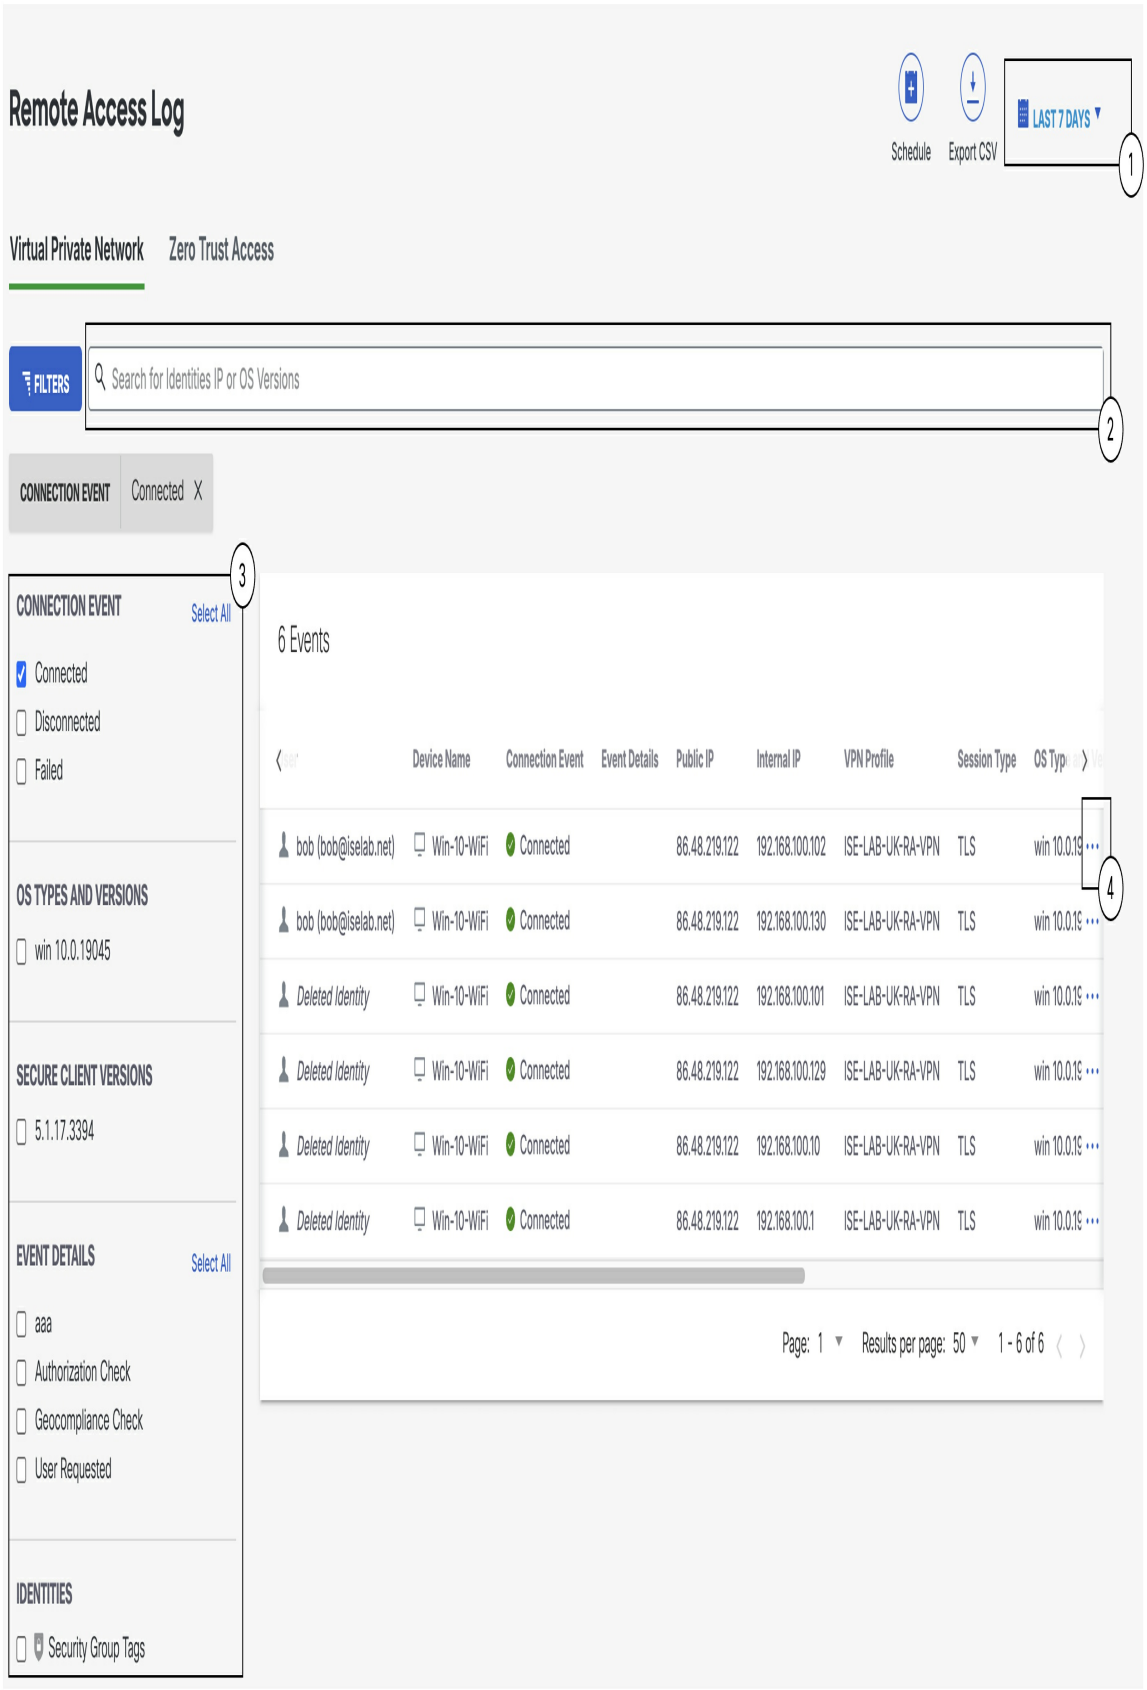
When SAML or certificate-based authentication is used with authorization on ISE, ensure that the authentication policy is configured to process Authorize-Only requests. The settings are the same as for Secure Firewall, which are covered in the sections "Certificate-Based Authentication with Authorization on ISE" and "SAML-Based Authentication with Authorization on ISE."

For both authentication scenarios, consider how the username is obtained:

- For certificate-based authentication, SSE uses the configuration in the VPN profile to determine which certificate attribute should be used as the username.
- For SAML authentication, the user enters the username directly on the SAML IdP login page, and it is returned to SSE inside the SAML assertion. SSE then submits this username in the Authorize-Only request to ISE.

SSE Remote Access VPN with ISE: Verification and Troubleshooting

From the SSE side, the main source of information for monitoring and troubleshooting NAC integration with ISE is the Remote Access Log report located in **Monitor > Remote Access Log**. Figure 13-70 provides an example of this report.


Remote Access Log
Schedule
Export CSV
LAST 7 DAYS
1
Virtual Private Network
Zero Trust Access
FILTERS
Search for Identities IP or OS Versions
2
CONNECTION EVENT
Connected X
3
CONNECTION EVENT
Select All
Connected
Disconnected
Failed
OS TYPES AND VERSIONS
win 10.0.19045
SECURE CLIENT VERSIONS
5.1.17.3394
EVENT DETAILS
Select All
aaa
Authorization Check
Geocompliance Check
User Requested
IDENTITIES
Security Group Tags
6 Events
Device Name
Connection Event
Event Details
Public IP
Internal IP
VPN Profile
Session Type
OS Typ
bob (bob@iselab.net)
Win-10-WiFi
Connected
86.48.219.122
192.168.100.102
ISE-LAB-UK-RA-VPN
TLS
win 10.0.19
bob (bob@iselab.net)
Win-10-WiFi
Connected
86.48.219.122
192.168.100.130
ISE-LAB-UK-RA-VPN
TLS
win 10.0.19
Deleted Identity
Win-10-WiFi
Connected
86.48.219.122
192.168.100.101
ISE-LAB-UK-RA-VPN
TLS
win 10.0.19
Deleted Identity
Win-10-WiFi
Connected
86.48.219.122
192.168.100.129
ISE-LAB-UK-RA-VPN
TLS
win 10.0.19
Deleted Identity
Win-10-WiFi
Connected
86.48.219.122
192.168.100.10
ISE-LAB-UK-RA-VPN
TLS
win 10.0.19
Deleted Identity
Win-10-WiFi
Connected
86.48.219.122
192.168.100.1
ISE-LAB-UK-RA-VPN
TLS
win 10.0.19
4
Page: 1
Results per page: 50
1 - 6 of 6

Technet24

Figure 13-70 *SSE Remote Access Log Report*

Some of the most important report parameters include the numbered fields in Figure 13-70 and further described here:

1. Timeframe selection, which by default is set to the last 24 hours
2. Filter bar to search events for specific users or IPs
3. Quick filters to display events of specific types, for particular devices, and similar criteria
4. Report details

From ISE, the most helpful tools for monitoring and troubleshooting integration with SSE are as follows:

- **Operations > RADIUS Live Logs:** Verify connection attempts in real time and review events from the last 24 hours.
- **Operations > Reports > Endpoints and Users > RADIUS Authentications:** Review historical authentication events.
- **Operations > Reports > Endpoints and Users > Posture Assessment by Endpoint:** Review details of posture assessment when it is performed for remote access SSE users.

All reports on ISE provide flexible filtering capabilities by username, endpoint MAC address, NAD, and other criteria.

Summary

This chapter guides you through the evolution of remote access technologies and their integrations with NAC solutions, starting with early VPN protocols and progressing toward modern architectures like SASE and SSE. It explains how SASE combines SD-WAN and security services in the cloud, while SSE focuses on delivering security functions

such as SWG, CASB, and ZTNA without the SD-WAN fabric.

You learn how to configure secure remote connectivity using Cisco Secure Firewall and Firewall Management Center. You set up certificate enrollment, manage NAT exemptions, adjust access control policies, and follow step-by-step procedures for enabling certificate-based, RADIUS, and SAML authentication. The chapter also addresses Authorize-Only requests in ISE to ensure seamless user authentication and authorization.

Finally, you explore how to integrate cloud-managed services with Cisco ISE. You configure directory synchronization, define system IP pools for network device identification, and apply consistent access policies across cloud-managed endpoints. These steps help you extend robust network access control to modern, cloud-based environments.

Technet24

Chapter 14. Managing Network Access on Campus with Microsegmentation

Introduction

Cisco Identity Services Engine is a key component for granting network access for users and endpoints. Whether the use case is to provide access to data center and cloud resources, allow east-west collaboration, or purely accommodate Internet access, Cisco ISE equips administrators with the technology to enforce granular, identity-based access policies, gain comprehensive visibility into network activity, and automate security operations across complex, multidomain environments.

Connectivity may be provided using a wired connection to an access port, using wireless from access points and wireless LAN controllers, or via secure remote access to a concentrator. Access methods include using 802.1X, MAC Authentication Bypass (MAB), and WebAuth. No matter the connectivity or access method used, ISE can be used to authenticate the user or endpoint and authorize least-privilege access using constructs like access control lists (ACLs) and security group tags (SGTs).

The topic for this chapter is microsegmentation, which is accomplished using SGTs. This is in contrast with macrosegmentation, which is

provided by utilizing virtual routing and forwarding (VRF) instances, VLANs, or firewalls/ACLs. Macrosegmentation may be suitable for some separation needs but is static in nature, is architecture dependent, and soon becomes an administrative burden to manage at scale. Microsegmentation allows granular separation within a routing table, within a VLAN, and even between multiple sessions within the same physical access port or interface.

Where to Enforce Traffic

A question that administrators are often asked is about best practice for where to enforce traffic within a network. In a modern network architecture, traffic flows from a source network device to multiple destination network devices. Within this framework, each network device operates with localized knowledge; the source device is aware only of the IP addresses and group assignments of the endpoints directly connected to its own ports, whereas destination devices possess similar knowledge regarding their own locally connected endpoints. To maintain optimal performance and scalability, the industry best practice is to enforce group-based security policies at the destination network device. When source classification metadata is embedded directly into the data plane with every packet, the destination device gains the necessary context regarding the source binding alongside its local knowledge of the destination binding, enabling precise, context-aware policy enforcement. This approach facilitates distributed enforcement across multiple network devices, which inherently scales as the network grows. Conversely, attempting to enforce group-based policies at the network ingress is not recommended; such a strategy would require the ingress device to maintain comprehensive knowledge of all possible destination bindings, while simultaneously forcing the source device to download exhaustive policies for every potential destination, creating a

Technet24

significant operational bottleneck that fails to scale in dynamic, large-scale environments.

Dynamic SGT Assignment

When users or endpoints connect to a network access port or SSID, the network device — be it a switch, access point, or wireless LAN controller — will send a RADIUS access-request to ISE (acting as the authenticator). ISE will use collected context to authenticate the user or endpoint and subsequently move to the authorization phase to authorize or deny network access. ISE uses authorization policies and authorization profiles to determine the level of network access to assign. One method of controlling the level of access is to assign an appropriate SGT. Once authorization is determined, ISE sends a RADIUS access-accept (or access-reject) back to the network device to implement the desired level of access — one example being to classify incoming traffic from the user or endpoint with the assigned SGT. This assigned SGT is tied to the user or endpoint session on the network device rather than the physical port or interface. This approach allows for multiple users or endpoints to be connected and authorized on the same access port. Take an IP phone, for example; it generally creates two sessions on an interconnected access port, one for voice and one for data, so a data device (a PC, for example) can be connected to the back of the phone, providing data services to the user of the phone. Another example is connecting an unintelligent switch to a Cisco access port. Several IoT endpoints can be connected to this unintelligent switch, but the Cisco network device can assign different access sessions to these detected IoT endpoints, providing segmented access between them.

Static SGT Classification

Not all endpoints will authenticate with ISE. Servers, for example, will most likely be connected directly to the network without utilizing access methods such as 802.1X or MAB. These can still be classified and play a part in the TrustSec design because their IP addresses can be statically assigned into groups rather than using RADIUS dynamic methods. Static classification methods include using CLI to add IP to SGT mappings (or bindings), Subnet to SGT mappings, VLAN to SGT mappings, or Port to SGT classifications. An alternative to managing static mappings via CLI, they can be managed centrally by ISE; however, ISE only supports IP TO SGT and Subnet to SGT mappings. To support ISE static mappings, navigate to **Work Centers > TrustSec > Components > IP SGT Static Mapping**.

As explained in Chapter 7, "Introduction to TrustSec," when adding static mappings into the ISE UI, you enter the IP address or subnet along with the SGT to be assigned. Additionally, you choose the method by which the mapping(s) is to be sent to the network devices. There are two choices: one is by using SXP, and the other is by using SSH for ISE to log in to specified network devices and provision the mappings directly into the running-config.

If you're adding static mappings into the SXP database (separated by domains), then those mappings are added immediately after you select **Save**, as shown in Figure 14-1. If you are deploying by SSH, then selecting **Save** does not immediately initiate the SSH function; the Deploy facility also needs to be initiated. See Chapter 7 for details.

Technet24

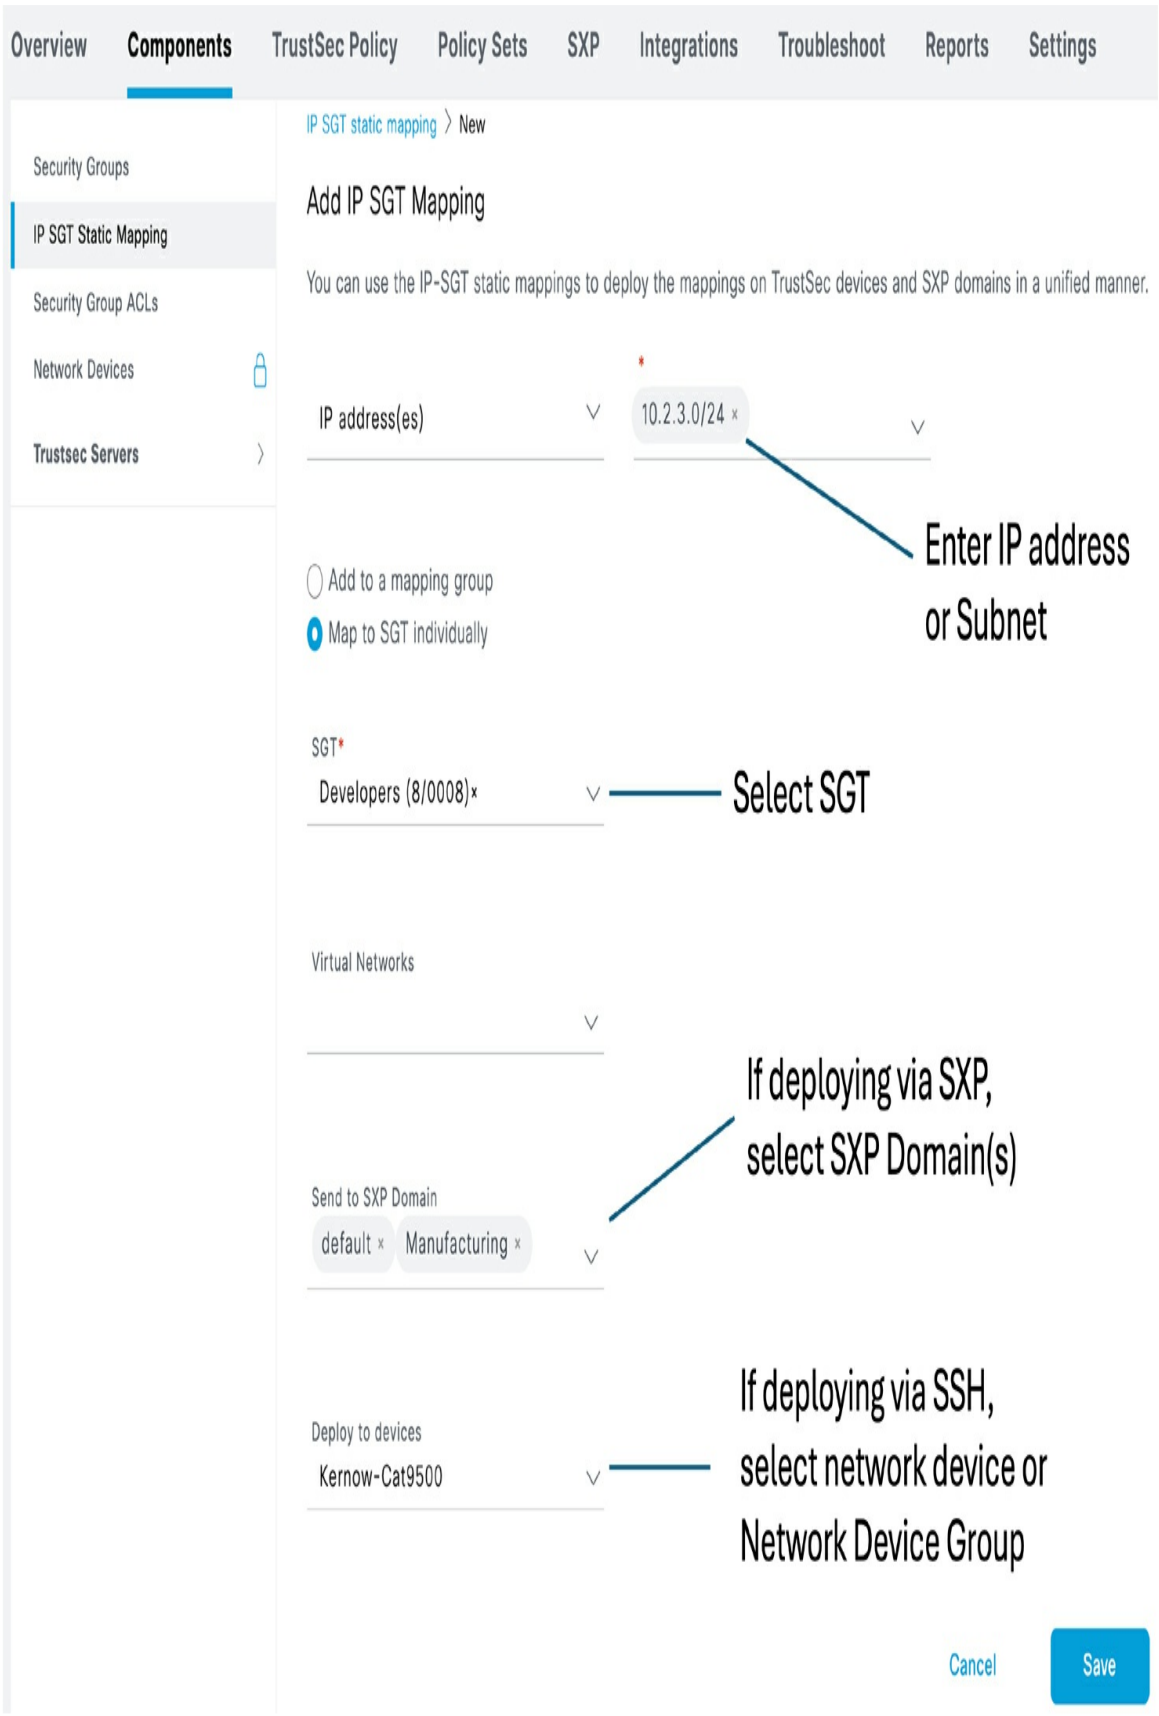

Overview
Components
TrustSec Policy
Policy Sets
SXP
Integrations
Troubleshoot
Reports
Settings
Security Groups
IP SGT Static Mapping
Security Group ACLs
Network Devices
Trustsec Servers
IP SGT static mapping > New
Add IP SGT Mapping
You can use the IP-SGT static mappings to deploy the mappings on TrustSec devices and SXP domains in a unified manner.
IP address(es)
10.2.3.0/24
Enter IP address or Subnet
Add to a mapping group
Map to SGT individually
SGT*
Developers (8/0008)
Select SGT
Virtual Networks
If deploying via SXP, select SXP Domain(s)
Send to SXP Domain
default
Manufacturing
If deploying via SSH, select network device or Network Device Group
Deploy to devices
Kernow-Cat9500
Cancel
Save

Figure 14-1 *Adding Static SGT Mappings into ISE*

Unknown SGT

When you are deploying SGTs, not all clients need to be classified. This allows for a phased approach, starting with a small number of groups, a limited number of endpoints or users, and a confined number of participating network devices. The clients without an SGT assignment have a classification implicitly understood to be the Unknown SGT. The number associated with the Unknown SGT is 0. The Unknown SGT cannot be explicitly specified as a classification for any client, but it is available to be used in security policy, either as a source group classification or a destination group classification.

An example using the Unknown SGT is in one method of classifying traffic from the Internet. If all clients on a network have explicit classifications, then the fact that traffic from the Internet is not classified allows policies to be added, such as source Unknown SGT to destination explicit SGT and equally from source explicit SGT to destination Unknown SGT. Adding a policy from Unknown SGT to Unknown SGT is also possible. Figure 14-2 shows examples of TrustSec policies in ISE using the Unknown SGT as either a source or destination criteria, or both.

Technet24

Production Matrix

Populated cells:22

Create Policy Monitor All Disabled Deploy



Figure 14-2 *ISE TrustSec Matrix Highlighting Unknown SGT*

Default Route SGT

As discussed, traffic from the Internet can be implied to be classified with Unknown SGT, but there is a more definitive method of classifying Internet traffic. It relates to a network device having a default route to reach the Internet, because the classification is for traffic directed via the default route. The network device command used to carry this out is as follows:

For IPv4: **cts role-based sgt-map 0.0.0.0/0 sgt x**
For IPv6: **cts role-based sgt-map 0:0::/0 sgt x**

This command is not automated by Catalyst Center, Meraki Dashboard, or pushed by ISE; it is implemented via the use of the CLI only. If Internet traffic is directed by the default route, then traffic coming from the Internet will be classified by the specified SGT, and traffic destined toward the Internet will be subjected to any TrustSec enforcement policies with the specified SGT as a destination criterion.

Mitigating Malware Lateral Movement

When TrustSec enforcement is discussed, it is often assumed enforcement is between two different groups. However, enforcement is also possible between members of the same group. As an example, a deployment may include classifications for employees with users assigned an SGT with the same name — Employees. SGT enforcement can be utilized to permit traffic between these users or deny traffic between them. Alternatively, an SGACL can be deployed to permit some ports and protocols and deny others. This is a method by which malware lateral movement can be mitigated between users. Even if some ports need to be open to allow collaboration between users,

Technet24

ports typically used by threat actors to propagate malware can be denied and therefore contain any compromise. As a best practice, it is recommended to use an SGACL to permit the ports that are required between users and to use an explicit default deny to block all other traffic.

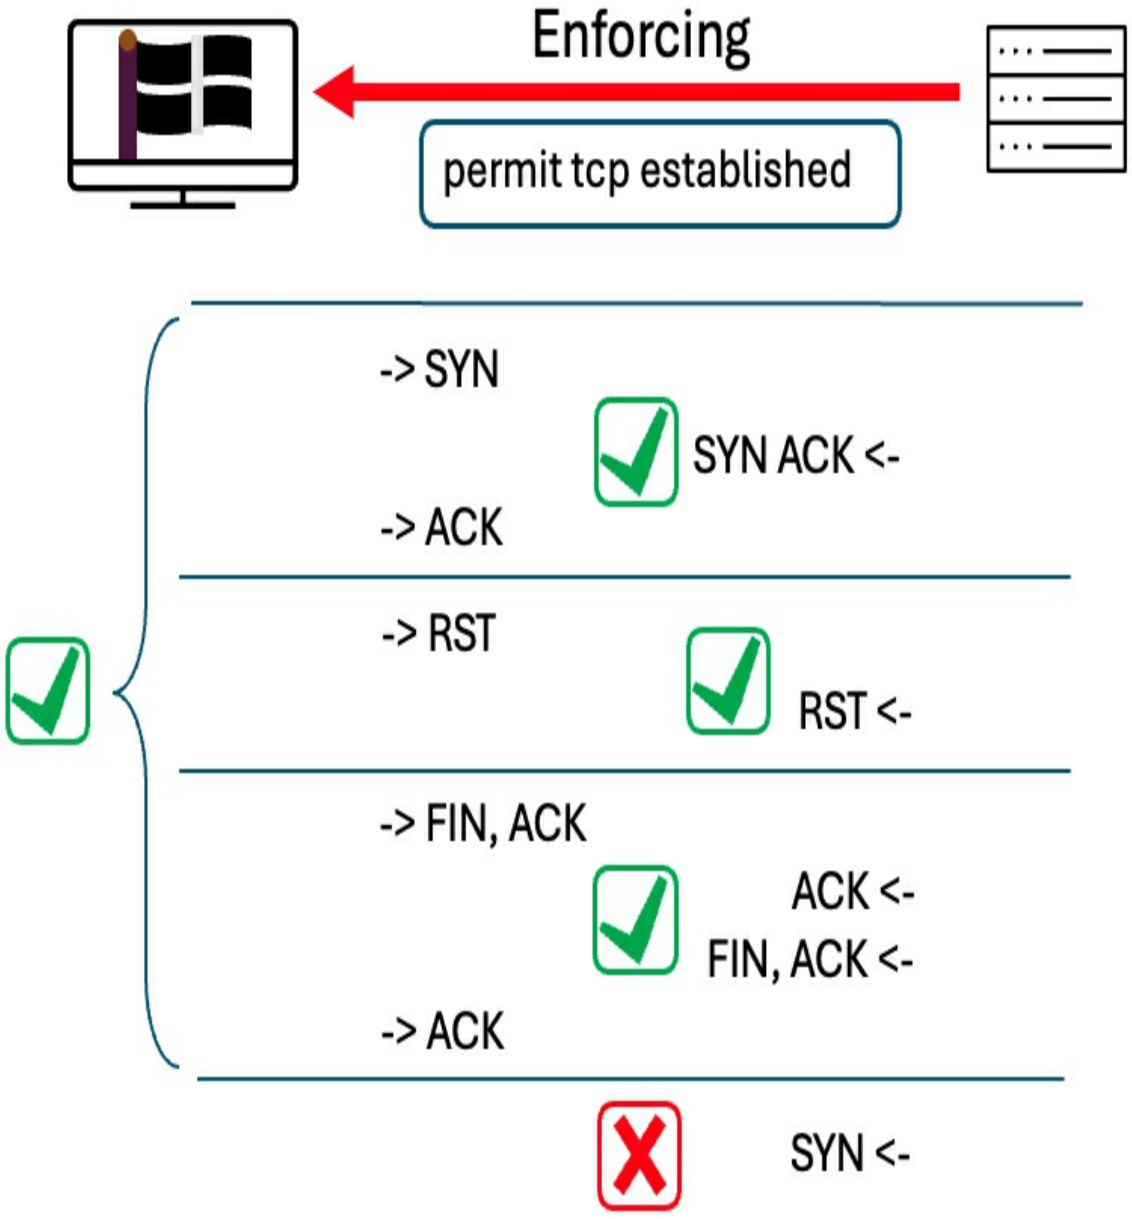
Enforcement Using TCP Flags

Layer 4 TCP and UDP ports are typically used within an SGACL to permit or deny traffic. However, TCP flags can also be monitored and the state of the flags used for enforcement purposes. The available flags can be seen via the CLI using the question mark option, as shown in Example 14-1.

Example 14-1 *Network Device TCP Options Within a Role-Based ACL*

```
Switch(config-rb-acl)#deny tcp ?
  ack               Match on the ACK bit
  dscp              Match packets with given dscp value
  dst               Destination port conditions
  established       Match established connections
  fin               Match on the FIN bit
  fragments         Check non-initial fragments
  log               Log matches against this entry
  log-input         Log matches against this entry, including inpu
  match-all         Match if all specified flags are present
  match-any         Match if any specified flag is present
  option            Match packets with given IP Options value
  precedence        Match packets with given precedence value
  psh               Match on the PSH bit
  rst               Match on the RST bit
  src               Source port conditions
  syn               Match on the SYN bit
  tos               Match packets with given TOS value
  urg               Match on the URG bit
```

One of the interesting options shown in Example 14-1 is **established**. It uses a combination of ACK and RST flags to detect stateful connectivity. SGACLs are stateless in operation, but the use of TCP flags allows stateful operation to be detected. Figure 14-3 shows an SGACL containing **permit tcp established** between a server and a client, which only allows packets that include the ACK or RST TCP flags. As such, it allows traffic established from the client, but not when attempted to be established from the server.



Technet24

Figure 14-3 *TCP Flag Operation in SGACLs*

PAC-less Operation

TrustSec traditionally uses a Protected Access Credential (PAC), which is a unique shared credential used to mutually authenticate the client and the server. It establishes a secure TLS tunnel between the client and the server during EAP-FAST authentication. The PAC key is used to create this unique TLS tunnel, enabling mutual authentication without requiring digital certificates. Chapter 7 goes into more detail and explains the PAC-less feature. It is recommended to make use of the PAC-less feature from ISE release 3.4 and IOS-XE 17.15.1 because it does not add significant security benefits and potential PAC key operational issues can be avoided. Better still, you should download policy via the HTTPS option to provide a secure and reliable transport.

Use Cases

Enforcement on the campus can be provided at the access layer, at the distribution or core, or at gateways leading to the Internet, cloud, or data center locations. It is available on switches, routers, wireless LAN controllers, access points, and firewalls and enforces wired, wireless, and remote access traffic. ISE is key in all use cases because it dynamically assigns endpoints and users into secure groups and provides the segmentation policy to accommodate microsegmentation.

East-West Use Case

One traffic flow use case, from endpoint to endpoint, is typically known as an east-west flow. Traffic flows from one access device across the campus to another access device. In this case, for non-

fabric, it is best to utilize inline tagging from a source access switch to a destination access switch where the source SGT classification is embedded in the Layer 2 Cisco Meta Data (CMD) field. For a campus using a fabric, the source classification is embedded in a Virtual eXtensible Local Area Network (VXLAN). This way, the destination switch learns of the source classification via the received traffic, knows the destination classification due to the destination endpoint's authenticated session, and can enforce using the downloaded policy.

User-to-Data Center Use Case

Another use case may be user-to-data center access control, providing least privilege, and protecting confidential data. If the data center is operating with a compatible group-based mechanism, such as Cisco Application Centric Infrastructure (ACI), then group information can be exchanged. However, if incompatible, then enforcement can be provided on campus at a gateway function between the campus and data center. This gateway function could be provided by a switch, a router, or a firewall. Similarly to the east-west flow, inline tagging could be utilized to propagate the source SGT from the campus access layer to the gateway in a nonfabric deployment. The destination SGT mapping is not known by the gateway natively though; it must be informed of the destination information by out-of-band methods. This could be SXP for switches, routers, and Cisco Secure ASA firewall, or pxGrid for the Cisco Secure Firewall Threat Defense or third-party firewalls. Cisco ISE can provide this destination control-plane context to the gateway via SXP or pxGrid, but first ISE needs to learn this destination mapping information. The data center resource's IP TO SGT mapping information can be entered into the ISE UI, but the programmatic methodology preferred today is automation via API. When data center resources are spun up, you can design the automation system to automatically add the assigned IP and SGT into

Technet24

ISE via API to create the static mapping. This way, ISE can immediately propagate the mapping to the gateway via SXP or pxGrid, triggering the gateway to download the required security policy and enforce traffic appropriately.

User-to-Workload Use Case

A user-to-workload use case, where those workloads are in the cloud, is like the user-to-data center use case, whereby a gateway can act as a classifier for traffic from the cloud and an enforcer for traffic toward the cloud. Again, an automation system could add the required IP TO SGT mappings into ISE when cloud workloads are spun up, which will allow ISE to propagate those mappings to the gateway via SXP or pxGrid as appropriate.

However, there is a mechanism by which ISE can learn of workloads directly — by using the Workload Connections feature. This feature is discussed in detail in Chapter 17, "Classifying Cloud Workloads." In this case, ISE uses a workload connector to access the cloud provider or VMware vCenter to learn of workloads, their IP addresses, and the workload attributes. These attributes are then used in an ISE classification table to assign SGTs, allowing generated IP TO SGT mappings to be propagated to network devices, such as a gateway, via SXP, or pxGrid. Ultimately, traffic from the workloads in the cloud or VMware vCenter is classified by these mappings in the gateway, and this classification is propagated toward the campus edge, where enforcement takes place. In the other direction, traffic flowing from users in the campus propagates user classification to the gateway where the cloud or vCenter workload mapping resides and enforcement takes place.

Reporting and Troubleshooting

Chapter 7 details some mechanisms used for reporting and troubleshooting. They include using syslog to inform of SGACL hits. Syslog messages are generated when entries in an SGACL are suffixed with the log keyword. It is best practice to use the log keyword sparingly when deploying new groups, services, or policies and to remove it when any investigation has completed. This approach saves on the CPU resources required. Using NetFlow to notify of group-based operations is an alternative to syslog and makes better use of hardware resources. NetFlow can export flows utilizing the use of SGTs with the result being consumed by the likes of Cisco Secure Network Analytics (SNA) and Splunk. The match entries in the flow record are configured as shown in Example 14-2.

Example 14-2 *NetFlow Configuration with* ***match*** *Commands for SGTs*

```
flow record
  match ipv4 version
  match ipv4 protocol
  match ipv4 source address
  match ipv4 destination address
  match transport source-port
  match transport destination-port
  match interface
  match flow cts source group-tag
  match flow cts destination group-tag
  collect counter bytes long
  collect counter packets long
  collect timestamp absolute first
  collect timestamp absolute last
  collect transport tcp flags
```

SGACL drops can also be reported via NetFlow from IOS-XE release 17.13. Cisco Secure Network Analytics can make use of this from release 7.4.2, helping with policy compliance visibility. The collect entry required in the flow record is shown in Example 14-3, and this must be

Technet24

configured when matching the output interface because SGACL drops are carried out at egress.

Example 14-3 *NetFlow Commands Showing a Collect Entry for SGACL Drops*

```
flow record
  match ipv4 version
  match ipv4 protocol
  match ipv4 source address
  match ipv4 destination address
  match transport source-port
  match transport destination-port
  match interface output
  match flow cts source group-tag
  match flow cts destination group-tag
  collect counter bytes long
  collect counter packets long
  collect timestamp absolute first
  collect timestamp absolute last
  collect policy firewall event
```

Summary

This chapter explores the implementation of microsegmentation within a campus using Cisco Identity Services Engine and security group tags (SGTs) to provide granular, identity-based network access. Unlike macrosegmentation, which relies on static virtual routing and forwarding (VRF) instances, VLANs, or firewalls/ACLs, microsegmentation allows for flexible, scalable security within VLANs and routing tables. The industry best practice for enforcement is to perform policy checks at the destination switch, where source classification metadata is combined with local destination knowledge to avoid the operational bottlenecks associated with ingress-based enforcement.

SGTs can be assigned dynamically through ISE via RADIUS authentication, which supports multiple sessions per physical port, or statically for nonauthenticating devices like servers. Special classifications include the Unknown SGT for unclassified traffic and the Default route SGT for managing traffic from and to the Internet.

Security capabilities are further enhanced by monitoring TCP flags, such as the established flag, to detect stateful connectivity within stateless SGACLs. The chapter also recommends transitioning to PAC-less operation and HTTPS-based policy downloads in modern ISE and IOS-XE releases to simplify credential management. Various use cases illustrate how these concepts apply to east-west traffic, data center access, and cloud workloads, with the latter often utilizing pxGrid Direct for automated discovery and classification. Finally, the chapter covers troubleshooting and reporting, highlighting the use of syslog and NetFlow for deep visibility into group-based operations and policy drops, which can be integrated with platforms like Cisco Secure Network Analytics to ensure comprehensive policy compliance.

Technet24

Chapter 15. Extending Microsegmentation to WAN and SD-WAN

Introduction

Microsegmentation is covered elsewhere in this book for local area networks (LANs), whether that be for fabric or non-fabric sites. This chapter covers the case where multiple sites exist, and the requirement is to propagate policy from site to site over a wide area network (WAN).

When fabrics are involved, the WAN transport is catered for within the technologies. With a LISP-based fabric like Cisco software-defined access (SDA), an SDA-Transit is deployed across the WAN that uses LISP as the control plane and VXLAN as the data plane encapsulation. Transit control plane nodes handle inter-site prefix registration and signaling. With a BGP EVPN–based fabric, the transit is part of the fabric, where BGP EVPN is the control plane protocol providing a standards-based overlay for multisite connectivity and, similarly to LISP, uses VXLAN encapsulation for the data plane. In both cases, VXLAN allows the propagation of the assigned security group.

Where fabric technologies are not available for use across the WAN, other methods are required to propagate the security policy. One method is to use inline tagging where the Cisco Meta Data (CMD) field

in the Layer 2 frame is utilized to insert the assigned source security group tag (SGT) for the receiving site to use in enforcing policy decisions. Inline tagging can be used to propagate policy between fabric sites, between non-fabric sites, and between fabric and non-fabric sites. Not all network devices used across the WAN support inline tagging though, so an investigation would be required to determine if this is a possible solution in the deployment in question.

A great solution for deploying, managing, and supporting a WAN installation is Cisco Catalyst software-defined wide area networking (SD-WAN). This transformative technology simplifies complex network operations by decoupling the control plane from the forwarding plane, enabling granular visibility, and propagating the SGTs to ensure consistent connectivity and policy across diverse enterprise environments. Cisco Catalyst SD-WAN includes several key components that work together to provide a secure and manageable WAN solution:

- The Cisco Catalyst SD-WAN Manager, formerly known as vManage, acts as a centralized network management system with a graphical dashboard. It simplifies device configuration, network management, and monitoring; automates deployment with preconfigured and user-defined templates; and offers application visibility, troubleshooting, and operational insights.
- The Cisco Catalyst SD-WAN Controller, previously called the vSmart Controller, serves as the centralized brain of the SD-WAN solution by controlling data traffic flow. It establishes and maintains control plane connections with edge routers using Datagram Transport Layer Security (DTLS) tunnels; runs the Overlay Management Protocol (OMP) to manage routing, policy, and key distribution; and enforces policies created on the Manager.

Technet24

- Another important component is the Cisco Catalyst SD-WAN Validator, formerly the vBond Orchestrator, which handles the initial authentication and onboarding of WAN edge devices. It orchestrates connectivity between edge routers, controllers, and the manager, and facilitates NAT traversal for devices behind NAT.
- The WAN edge routers, which include Cisco IOS XE Catalyst SD-WAN devices, are physical or virtual routers located at branch offices, data centers, or cloud sites. They provide secure data plane connectivity between sites over multiple WAN transports and handle traffic forwarding, encryption, quality of service (QoS), and routing protocols such as BGP and OSPF.

Together, these components form the orchestration, management, control, and data planes of the Cisco Catalyst SD-WAN solution, enabling secure, scalable, and manageable WAN connectivity optimized for multicloud and SASE architectures.

SGT Propagation Across SD-WAN

SD-WAN edge devices need to be configured to propagate the source SGT across the WAN. For instance, a fabric or non-fabric LAN site can use inline tagging to insert the source classification (SGT) of a user or endpoint into the L2 frames up to an SD-WAN ingress edge router. That ingress router can be configured to receive the SGT from the LAN connection and transport it across a tunnel to a receiving egress SD-WAN router. This egress router can take the SGT from the tunnel and place it back into L2 frames for traffic sent to the destination LAN devices where enforcement can take place. Alternatively, the SD-WAN devices could enforce the traffic themselves based on the transported SGT and group-based policy implemented.

To enable SGT inline tagging propagation through the SD-WAN IPsec

tunnels, the latest and strategic approach is to use configuration groups rather than feature templates. Configuration groups are also used to configure the LAN side of the connectivity, where the SGT will be received from the site on the ingress router and where the SGT will be sent by the egress router to the destination site.

To enable SGT propagation on the IPsec tunnel, navigate in SD-WAN Manager to **Configuration > Configuration Groups > Transport & Management Profiles**. When configuring an interface within a profile, click the **Tunnel** tab. Figure 15-1 shows the tunnel interface and SGT propagation enabled.

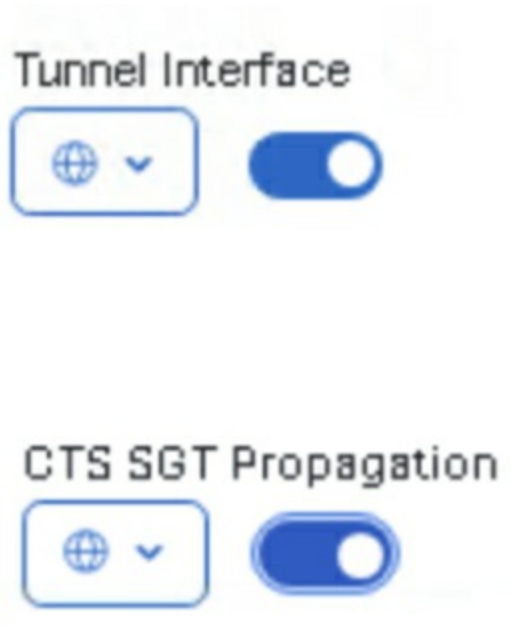


Figure 15-1 *Tunnel Interface and SGT Propagation Enabled*

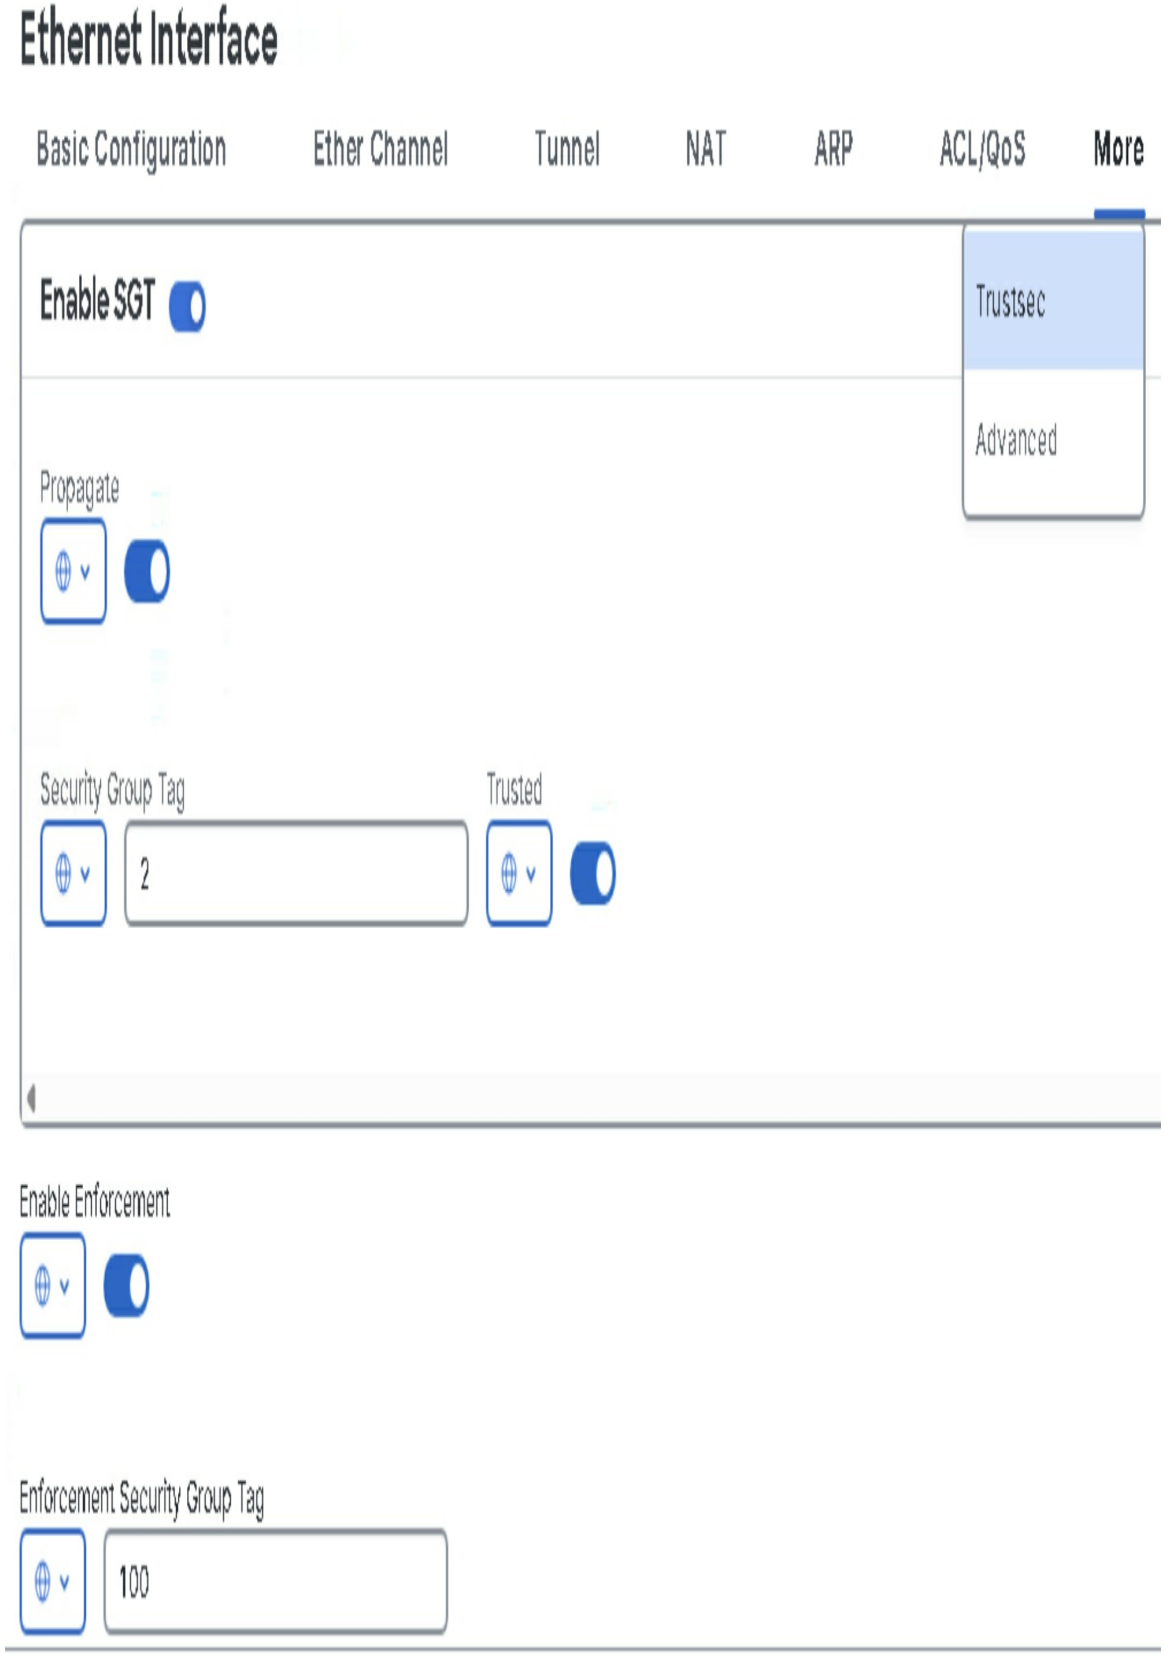
Deploying this update leads to configuration like that shown next. Note the tunnel is not shut down and cts manual is configured to enable SGT propagation:

```
interface Tunnel1
  no shutdown
  ip unnumbered GigabitEthernet1
```

Technet24

```
  no ip redirects
  ipv6 unnumbered GigabitEthernet1
  no ipv6 redirects
  cts manual
  exit
  tunnel source GigabitEthernet1
  tunnel mode sdwan
exit
```

To enable SGT support on the LAN side of the connectivity, again use **Configuration > Configuration Groups > Transport & Management Profiles**. When configuring an interface within a profile, click the **TrustSec** tab. Figure 15-2 shows all the options enabled and fields entered.


Ethernet Interface
Basic Configuration
Ether Channel
Tunnel
NAT
ARP
ACL/QoS
More
Enable SGT
Trustsec
Advanced
Propagate
Security Group Tag
2
Trusted
Enable Enforcement
Enforcement Security Group Tag
100

Technet24

Figure 15-2 *Enabling SGT Propagation on the SD-WAN LAN Interfaces*

If the settings in Figure 15-2 were deployed, the following indicates how the configuration may appear on the SD-WAN edge device:

```
interface GigabitEthernet1
  description LAN
  no shutdown
  arp timeout 1200
  ip address 172.16.1.5 255.255.255.0
  no ip redirects
  ip mtu 1500
  ip nat outside
  load-interval 30
  mtu 1500
  cts manual
    policy static sgt 2 trusted
  exit
  cts role-based enforcement
  cts role-based sgt-map sgt 100
  negotiation auto
exit
```

To explain, the cts manual enables the SGT propagation, policy static sgt 2 is from the Security Group Tag UI setting, and trusted is enabled. The SGT 2 here classifies inbound traffic if no SGT is received in the L2 frame or if the SGT in the L2 frame is not trusted. If an SGT is received in the L2 frame, and it is configured to be trusted, then the SGT within the frame will take precedence. Enabling enforcement in the UI manifests as cts role-based enforcement in the CLI, and Enforcement Security Group Tag in the UI as cts role-based sgt-map sgt x in the CLI. This is an inbound classifier for a Layer 3 interface classifying the learned IP prefixes and available for enforcement at egress. Figure 15-3 highlights the IP prefix-to-SGT binding created on the SD-WAN edge from this deployed configuration; the source is shown from a

Layer 3 interface classification (L3IF).

```
SanJose-BR1Edge1#show cts role-based sgt-map all
Active IPv4-SGT Bindings Information

IP Address              SGT     Source
========================================
172.16.1.0/24           100     L3IF

IP-SGT Active Bindings Summary
========================================
Total number of L3IF     bindings = 1
Total number of active   bindings = 1

Active IPv6-SGT Bindings Information

IP Address                                SGT     Source
============================================================
SanJose-BR1Edge1#
```

Figure 15-3 *IP Prefix (L3IF) Binding Created on SD-WAN Edge*

Integrating SD-WAN Manager with ISE

The Cisco SD-WAN Manager integrates with Cisco Identity Services Engine to learn security groups and group memberships and allow

Technet24

traffic enforcement within the WAN platforms. The integration with ISE occurs via Representational State Transfer (REST) handled by the ISE External Restful Services (ERS) implementation and pxGrid. Initially, you should ensure ISE has ERS enabled by navigating to **Administration > System > Settings > API Settings > API Service Settings**.

This ISE ERS setting is crucial for SD-WAN Manager to programmatically create the necessary client accounts on ISE, retrieve security groups, and exchange configuration and identity data.

pxGrid must also be enabled in ISE. To do so, navigate to **Administration > System > Deployment > [Select the ISE node in question] > General Settings**, as shown in Figure 15-4.



Figure 15-4 *Enabling ISE pxGrid*

When SD-WAN Manager integrates with ISE, it is recommended that certificates are used for a secure coupling. In the following example, a password-based methodology is used to highlight the simplicity. To enable a password-based integration as well as an automatic approval of the integration request, in ISE navigate to **Administration > pxGrid Services > Settings** and enable both options.

By default, a subscription to the ISE pxGrid SXP Topic will only retrieve

static IP:SGT bindings. To also retrieve dynamic bindings (created when using 802.1x or MAB, for example), you need to make a selection in the ISE settings. Navigate to **Work Centers > TrustSec > Settings > SXP Settings** and enable **Publish SGT Bindings on pxGrid**, as shown in Figure 15-5.

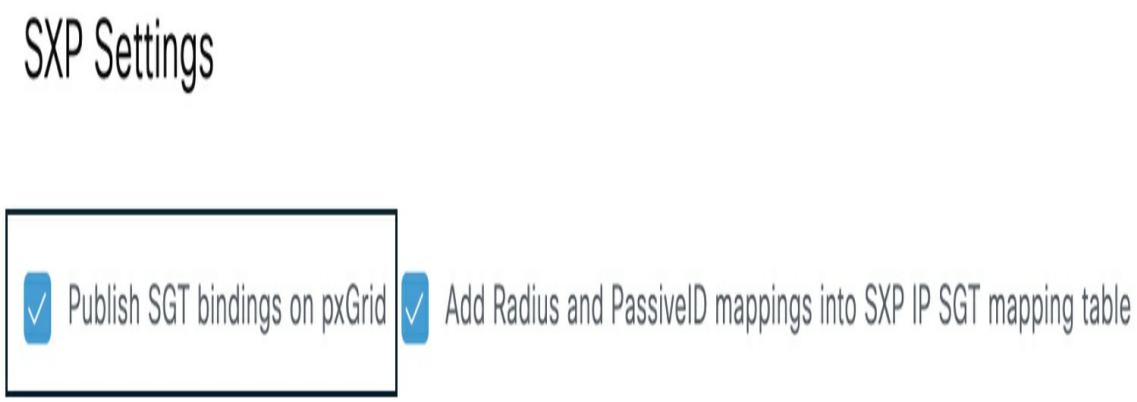


Figure 15-5 *Enabling the Publication of Dynamic IP:SGT Bindings in the SXP Topic*

The following examples use SD-WAN release 20.18.1. The details may differ slightly if you are using a different release. Within SD-WAN Manager, navigate to **Administration > Integration Management** (see Figure 15-6).

Technet24

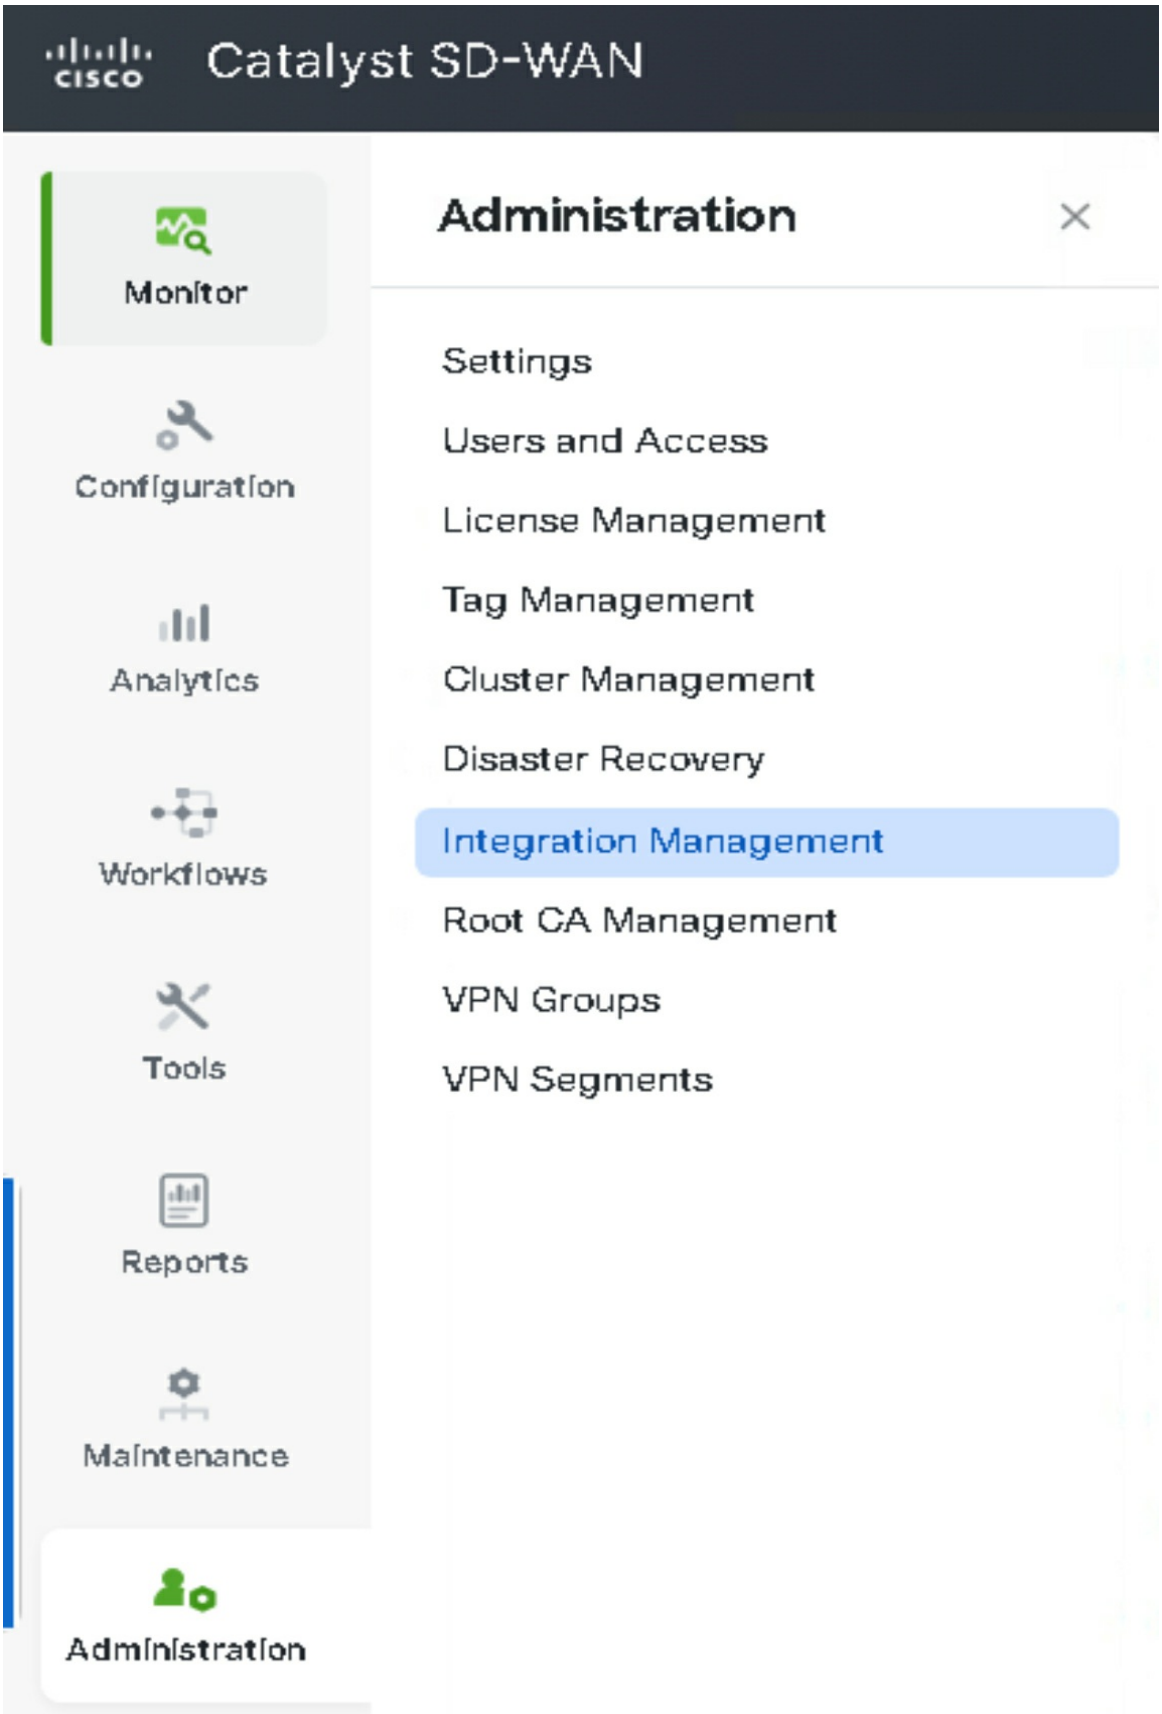

Catalyst SD-WAN
Monitor
Configuration
Analytics
Workflows
Tools
Reports
Maintenance
Administration
Administration
Settings
Users and Access
License Management
Tag Management
Cluster Management
Disaster Recovery
Integration Management
Root CA Management
VPN Groups
VPN Segments

Figure 15-6 *SD-WAN Manager Integration Management Menu*

When the **Integration Management** page opens, select the **Identity Services Engine** submenu, as shown in Figure 15-7.

Technet24

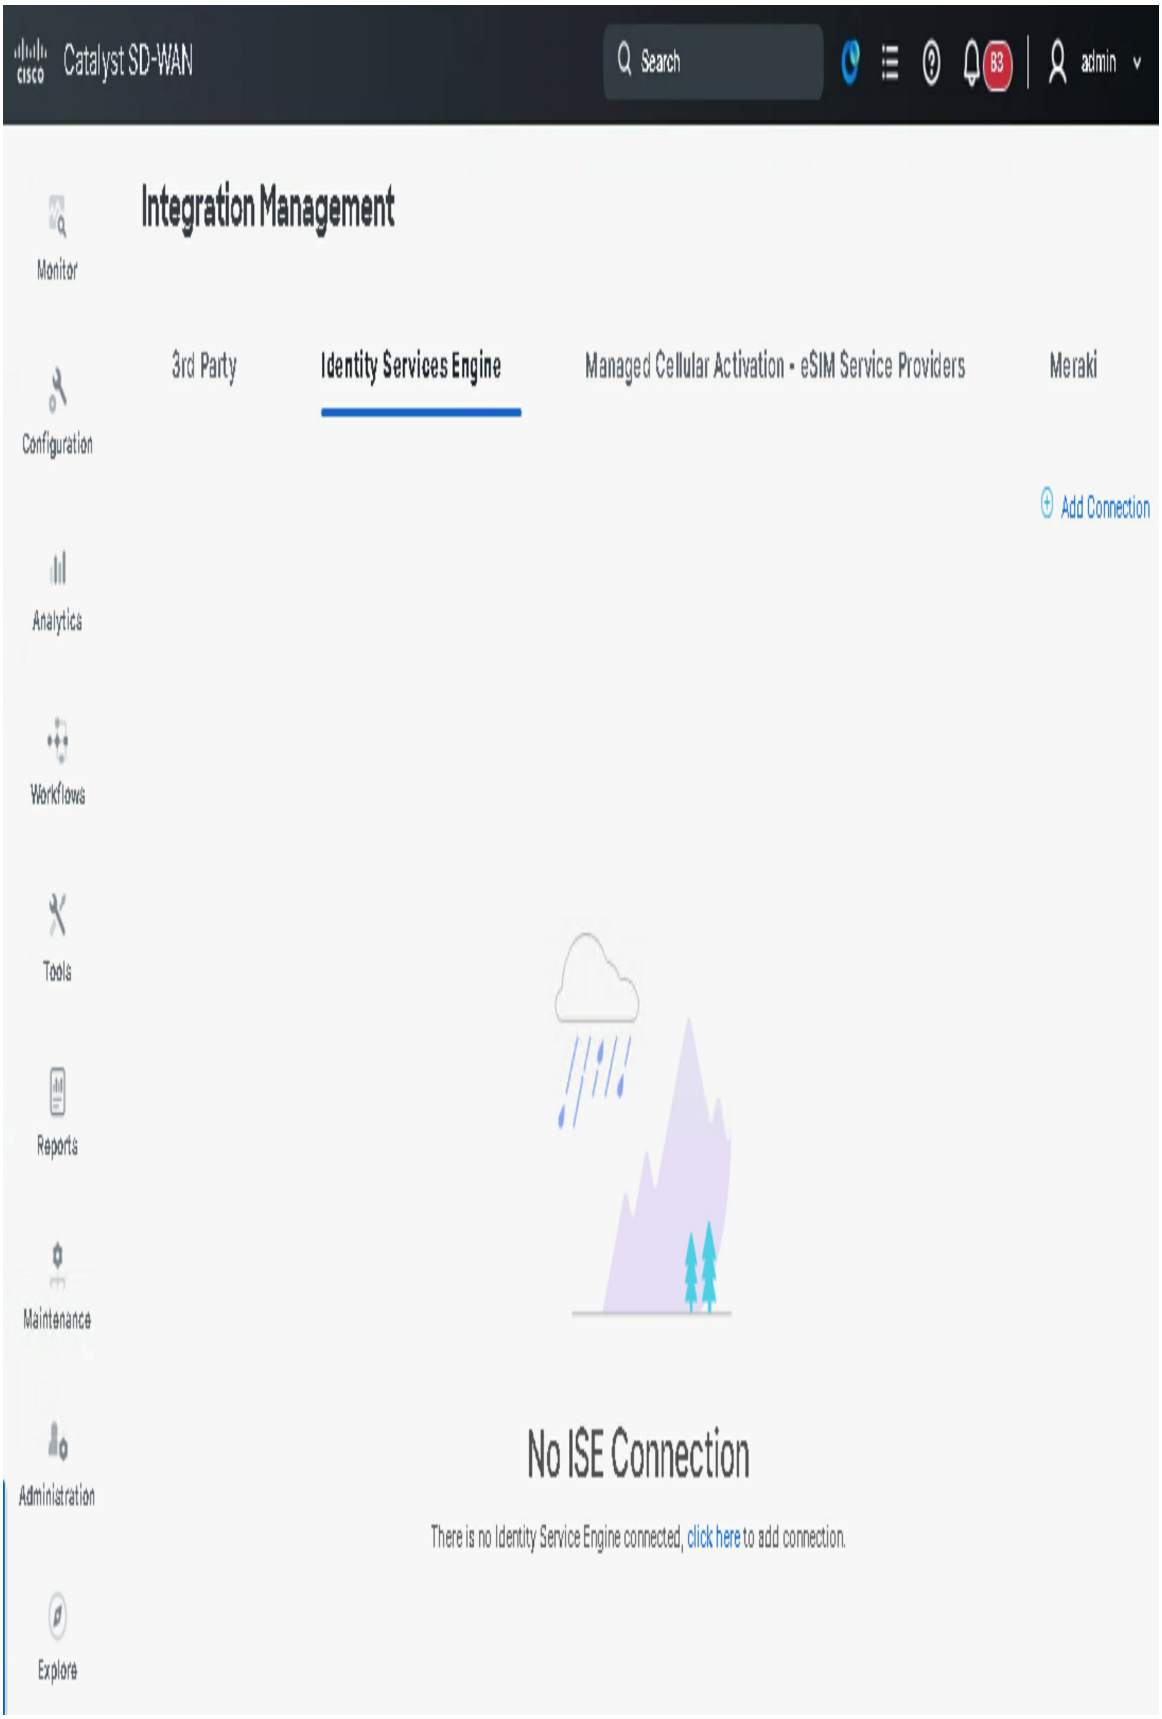

Catalyst SD-WAN
Search
83
admin
Integration Management
Monitor
Configuration
Analytics
Workflows
Tools
Reports
Maintenance
Administration
Explore
3rd Party
Identity Services Engine
Managed Cellular Activation - eSIM Service Providers
Meraki
Add Connection
No ISE Connection
There is no Identity Service Engine connected, click here to add connection.

Figure 15-7 *SD-WAN Manager Integration with ISE*

To add an integration with ISE, either click **Add Connection** near the top right of the screen or click the link provided in the message near the bottom. The window in Figure 15-8 appears.

Technet24

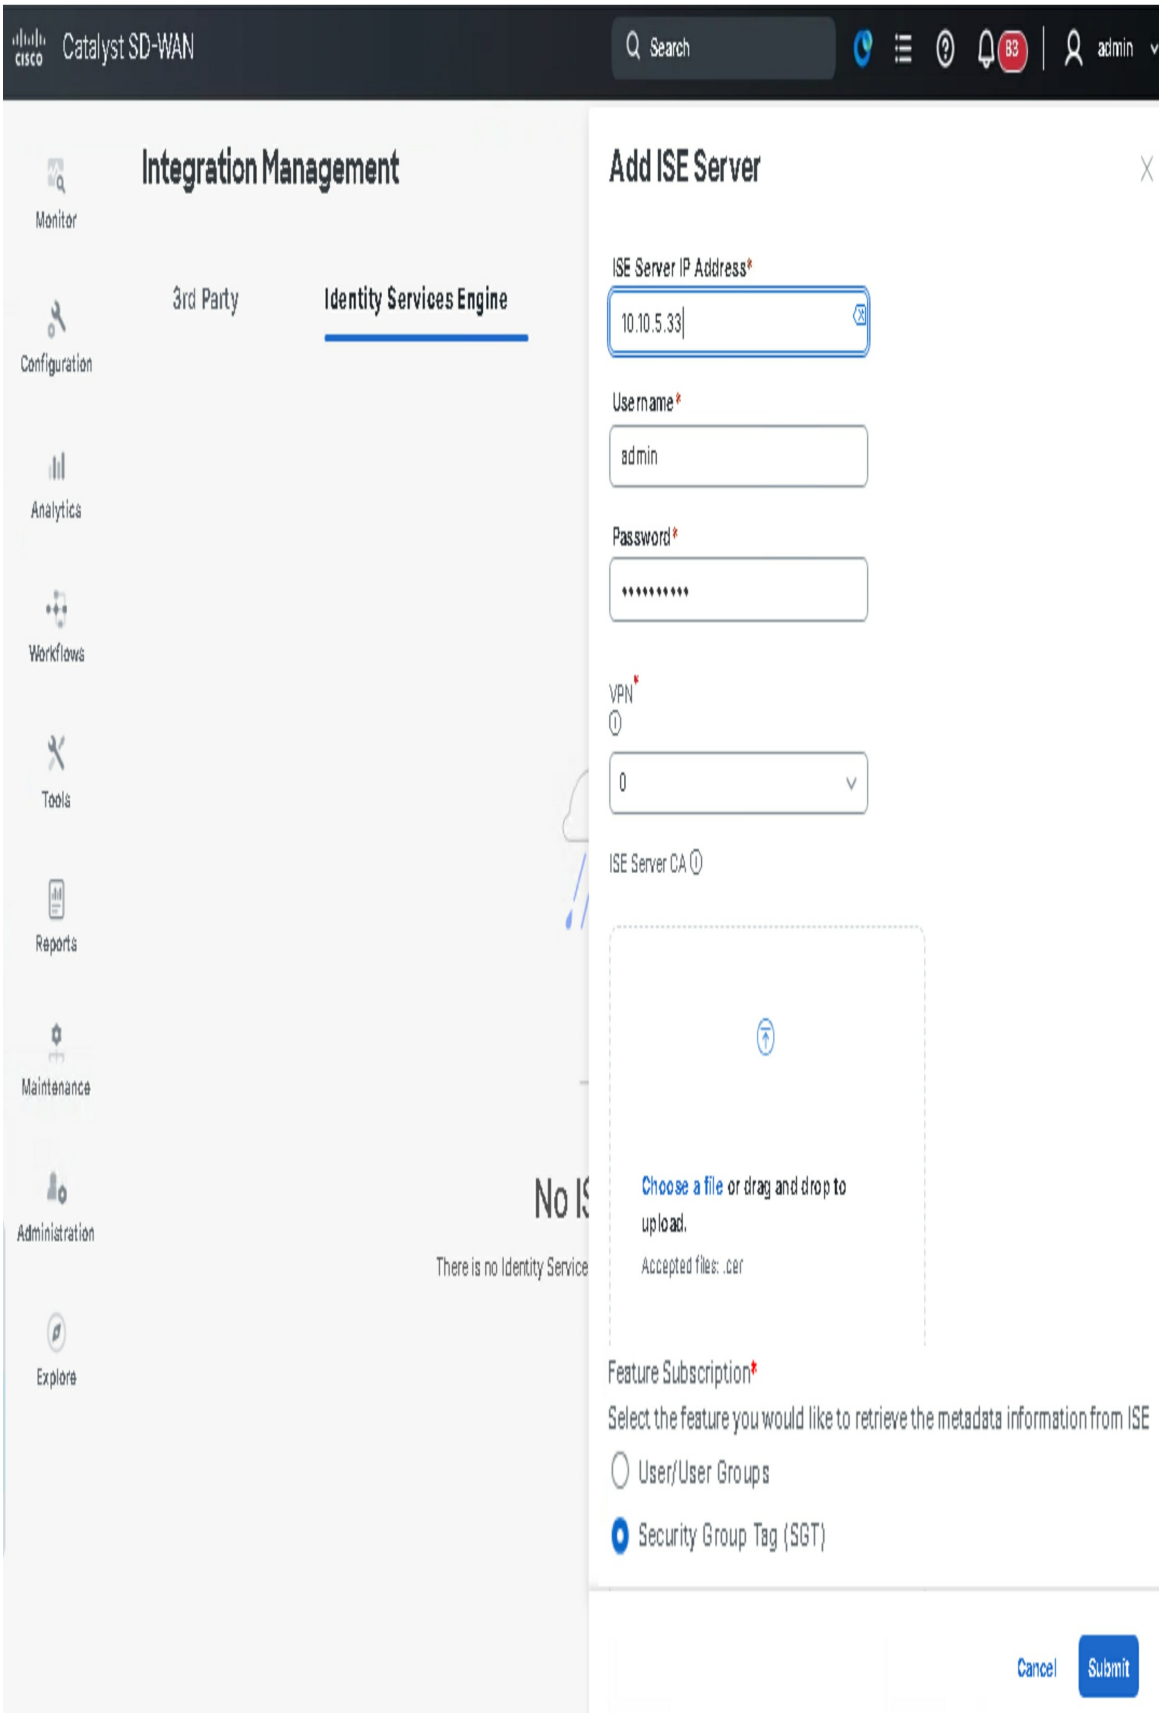

Catalyst SD-WAN
Search
admin
Integration Management
Monitor
Configuration
Analytics
Workflows
Tools
Reports
Maintenance
Administration
Explore
3rd Party
Identity Services Engine
No IS
There is no Identity Service
Add ISE Server
ISE Server IP Address*
10.10.5.33
Username*
admin
Password*
VPN*
0
ISE Server CA
Choose a file or drag and drop to upload.
Accepted files: .cer
Feature Subscription*
Select the feature you would like to retrieve the metadata information from ISE
User/User Groups
Security Group Tag (SGT)
Cancel
Submit

Figure 15-8 *Adding ISE Details into SD-WAN Manager*

Here, enter the ISE server FQDN or IP address and the username and password of ISE. Because this example does not use certificates, the ISE and pxGrid server certificate authority do not need to be uploaded, but they are recommended in a live deployment. Additional required information is the VPN over which connectivity to the ISE server must be established, plus a choice between subscribing to the **User/User Groups** or the **Security Group Tag (SGT)**. If selecting **User/User Groups**, you, as administrator, will be asked to input the ISE Active Directory identity store join point and the Active Directory domain.

SD-WAN Manager Subscribing to Security Group Tag (SGT)

In this first example, you will select the **Security Group Tag (SGT)** option when adding an ISE server so that a subscription to ISE pxGrid will learn the IP:SGT bindings.

After you submit the input, the SD-WAN Manager establishes a pxGrid connection to ISE and shows the status as **Connected**.

Under the **Action** column, you can select the three dots to display options. Figure 15-9 shows these options.

Technet24

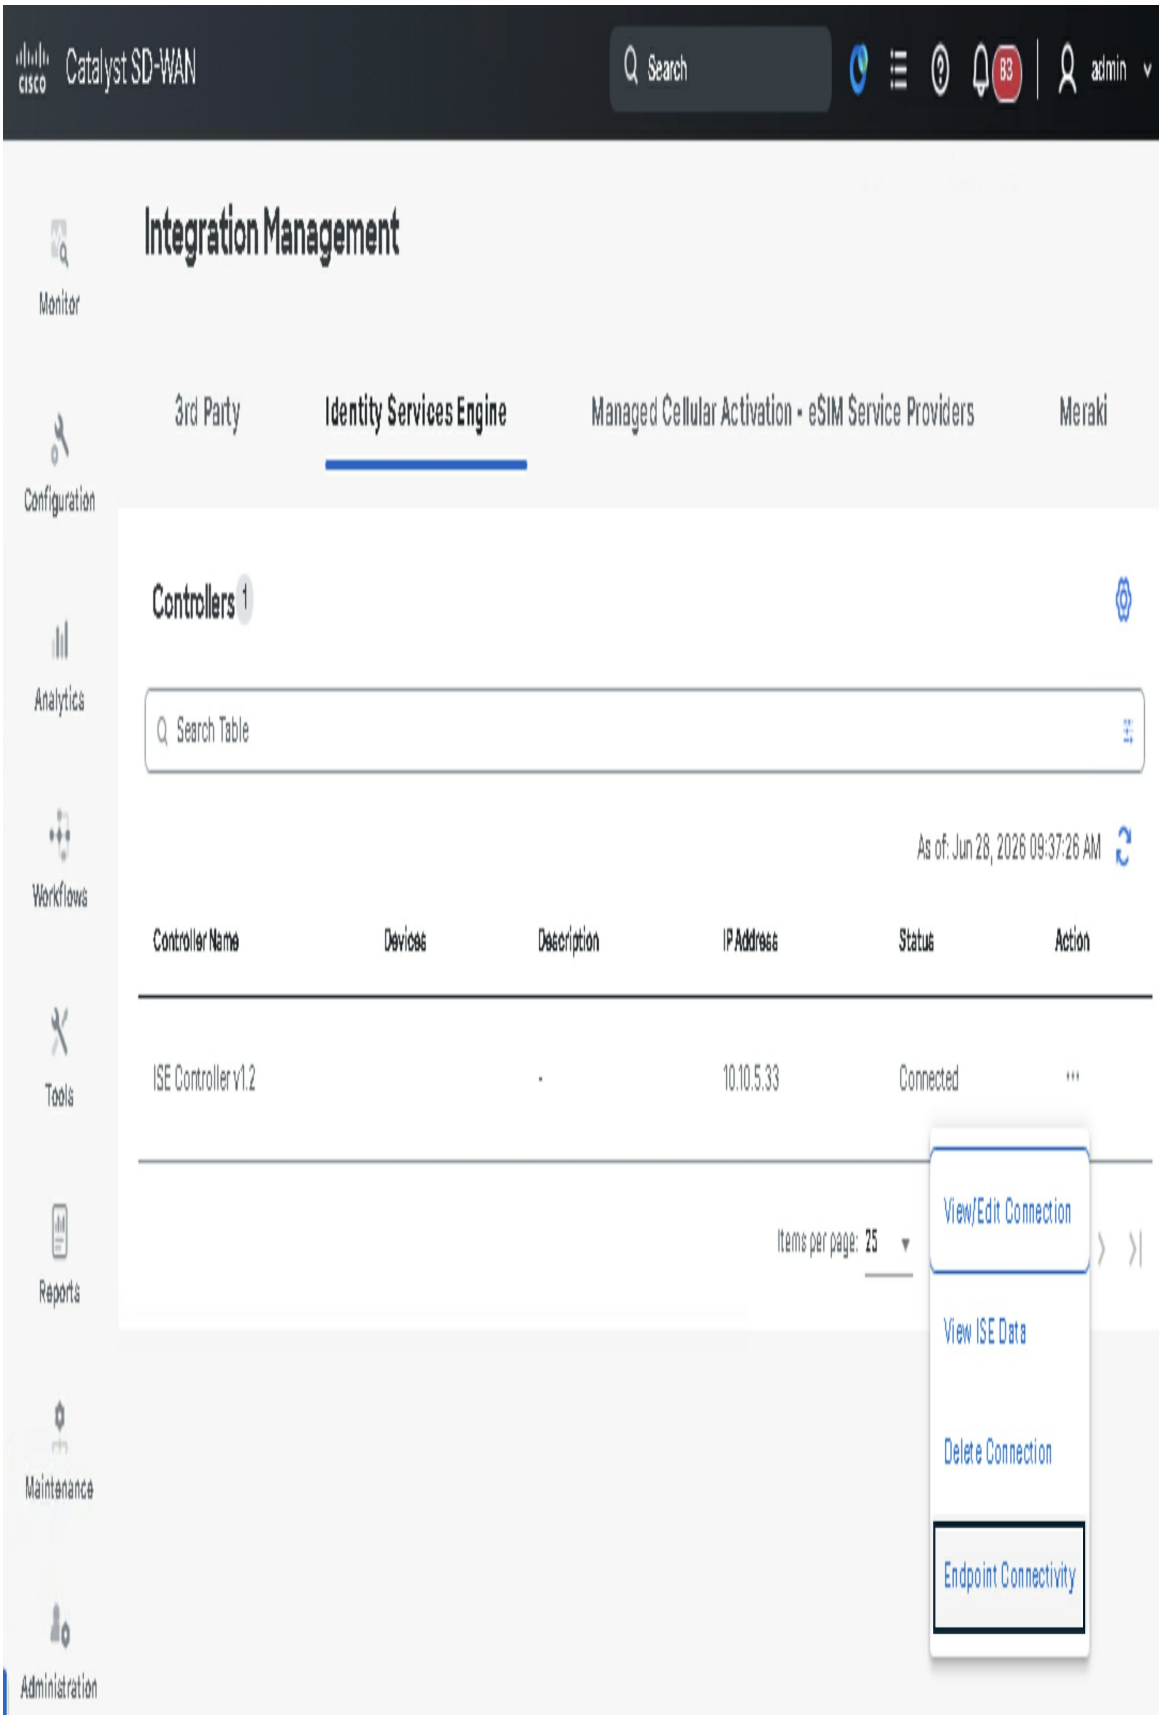

Catalyst SD-WAN
Search
admin
Integration Management
Monitor
Configuration
Analytics
Workflows
Tools
Reports
Maintenance
Administration
3rd Party
Identity Services Engine
Managed Cellular Activation - eSIM Service Providers
Meraki
Controllers 1
Search Table
As of: Jun 28, 2026 09:37:26 AM
Controller Name
Devices
Description
IP Address
Status
Action
ISE Controller v1.2
10.10.5.33
Connected
Items per page: 25
View/Edit Connection
View ISE Data
Delete Connection
Endpoint Connectivity

Figure 15-9 *Viewing Endpoint Connectivity of ISE Integration*

Selecting **Endpoint Connectivity** instructs SD-WAN Manager to display all the WebSocket connections that are active. Figure 15-10 shows that SD-WAN Manager has a WebSocket connection to ISE through which context is learned. SD-WAN Manager also has a connection to the SD-WAN Controller so that IP:SGT bindings are distributed from ISE to the Manager, then to the Controller, and on to the SD-WAN edge platforms.

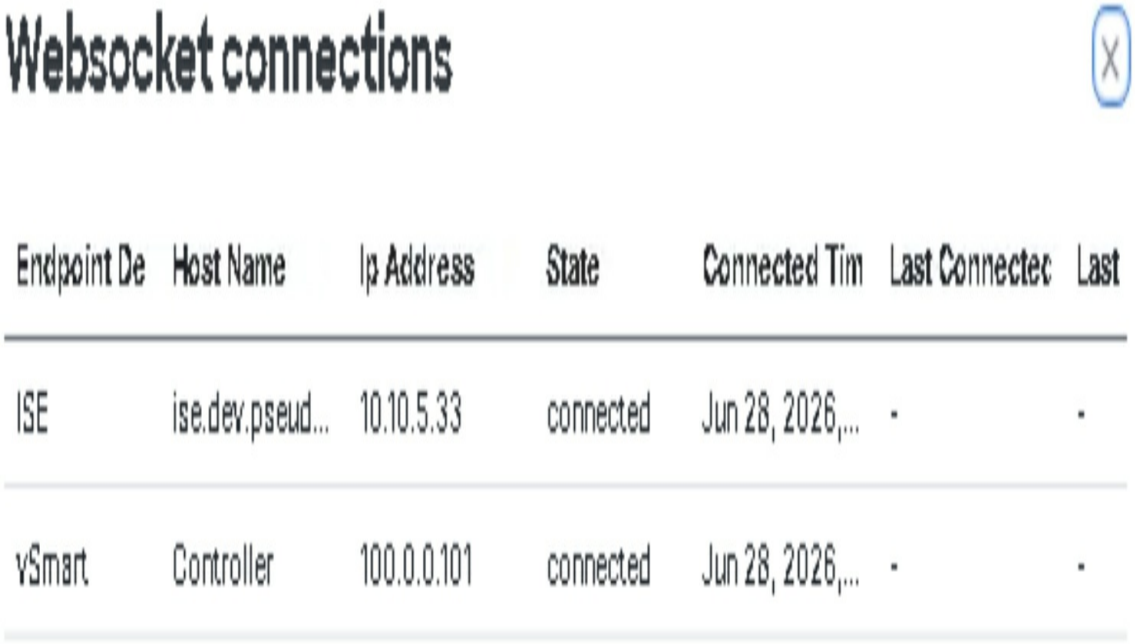


Figure 15-10 *SD-WAN Manager WebSocket Connections*

Figure 15-11 shows the ISE side of the WebSocket connectivity from the SD-WAN Manager. The last entry is the relevant one here; the pxGrid subscription is for learning bindings from /topic/com.cisco.ise.sxp.binding. This is how SD-WAN Manager, SD-WAN Controller, and ultimately the SD-WAN edge devices learn about the IP:SGT bindings from ISE.

Technet24

WebSocket

Connections Topics

Connections

Rows/Page 5 | 1 /1

Connection Name	Connect To	Session Id	Certificate	Subscriptions	Publications	IP Address
-ise-fanout-ise	ise	ise:0	CN=ise.dev.pseu...	/topic/wildcard		127.0.0.1
-ise-admin-ise	ise	ise:1	CN=ise.dev.pseu...	/topic/com.cisco.ise.pxgri...	/topic/com.cisco.ise...	10.10.5.33
-ise-mnt-ise	ise	ise:2	CN=ise.dev.pseu...	/topic/com.cisco.ise.sessi...	/topic/com.cisco.ise...	10.10.5.33
-ise-fanout-ise	ise	ise:3	CN=ise.dev.pseu...	/topic/distributed	/topic/distributed	10.10.5.33
dCloud-Lab-Modality_vmana...	ise	ise:4	No Certificate	/topic/com.cisco.ise.sxp.b...		172.16.0.1

/topic/com.cisco.ise.sxp
.binding

Figure 15-11 *ISE pxGrid Subscription from SD-WAN Manager*

Back in SD-WAN Manager, another option from clicking the three dots in the **Action** column of the ISE integration entry is **View ISE Data**.

Selecting the **View ISE Data** option displays the security groups, or SGTs, learned from ISE, as shown in Figure 15-12. You can use security group information from here to create security policies in SD-WAN Manager.

Technet24

View SGTs

Security Group Tag 22

Search Table

As of: Jun 28, 2026 09:40:37 AM

Name	Tag
BYOD	15
PCI_Servers	14
ANY	65535
Auditors	9
Contractors	5
Developers	8
Development_Servers	12
Employees	4
Guests	6
Network_Services	3
Point_of_Sale_Systems	10

Figure 15-12 *Security Group Tags in SD-WAN Manager Learned from ISE*

After the SD-WAN Controller learns the IP:SGT bindings, they are sent to the SD-WAN edge devices. They are sent over OMP, which handles the Cisco SD-WAN routing and policy control plane. An edge router shows the bindings learned from the Controller via source OMP, as seen in Figure 15-13.

```
SanJose-BR1Edge1#sh cts role sgt all
Active IPv4-SGT Bindings Information

IP Address              SGT     Source
============================================
10.10.102.12            5       OMP
10.10.102.34            19      OMP
10.10.102.100           4       OMP
10.10.102.106           20      OMP
10.10.102.109           17      OMP
10.10.102.170           17      OMP
192.168.0.10            16      OMP

IP-SGT Active Bindings Summary
============================================
Total number of OMP      bindings = 7
Total number of active   bindings = 7

Active IPv6-SGT Bindings Information

IP Address                                  SGT     Source
============================================================
SanJose-BR1Edge1#
```

Figure 15-13 *Edge Device Showing SGT Bindings Learned via OMP*

Using ISE Context in SD-WAN Manager Unified Policy

To use these learned security groups in SD-WAN policy, you first need

Technet24

to create identity lists in SD-WAN Manager. Navigate to **Configuration** > **Security**, select **Custom Options**, and click **Lists**. Then choose **Identity** from the list types.

Then click **New Identity List**, provide a name, and click the **Security Group Tag** drop-down box to select one or more SGTs, as shown in Figure 15-14.

New Identity List

Identity List Name*

New-People-Identity-List

Description

Subscription Type User/ User Group Security Group Tag (SGT)

Security Group Tag

Employees - 4 x Developers - 8 x Contractors - 5 x Guests - 8 x

Add Cancel

Name	Entries	Reference Count	Updated By	Last Updated	Action

No data available

Technet24

Figure 15-14 *SD-WAN Manager Adding a New Identity List*

Next, click **Add** to add this identity list.

Now you can create a second identity list, as shown in Figure 15-15. You can use one list as source criteria in the security policy and the other as destination criteria.

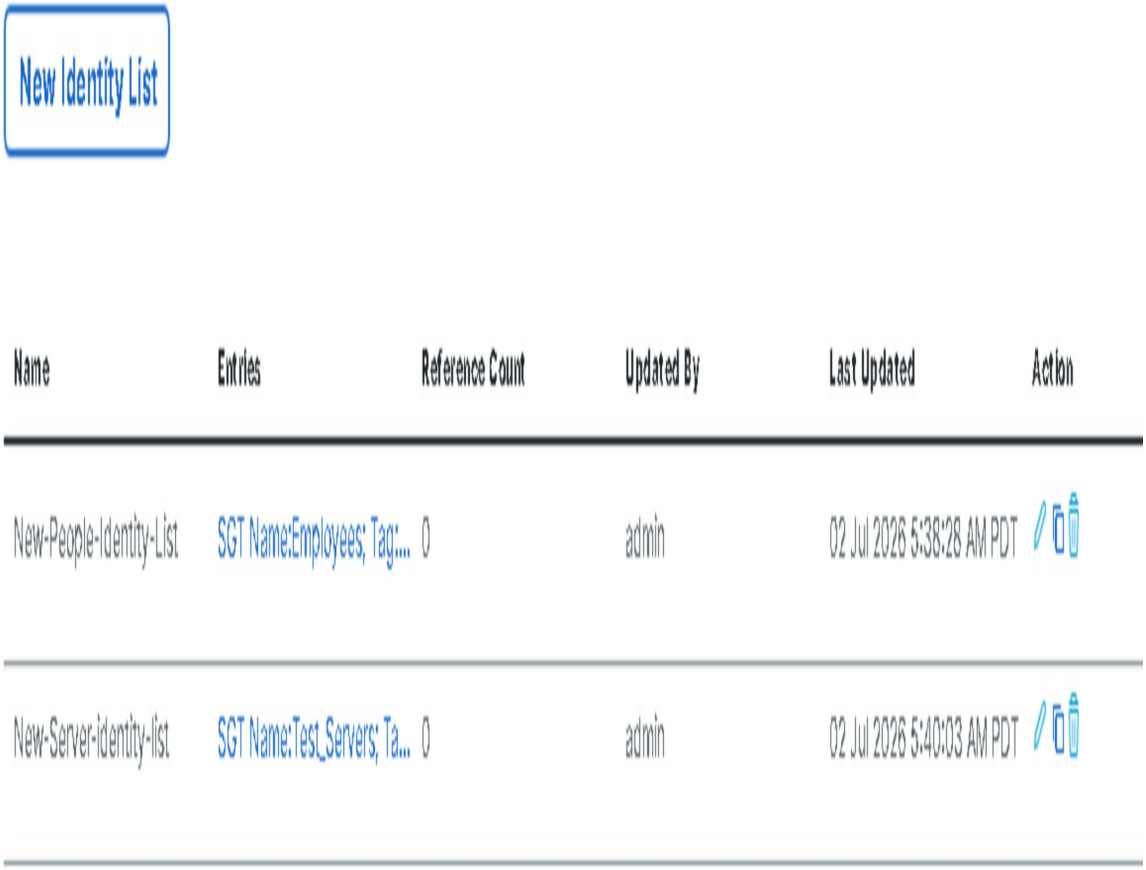


Figure 15-15 *Two Identity Lists Added*

Navigate to **Configuration > Security**. Then choose **Add Unified Policy**.

In the following UI, select **Add NG Firewall Policy** and select **Create New**.

Within the new Unified Security Policy, you can add rules.

Add rules by clicking **Add Rule/Rule Set Rule** to add a new rule. This is where you can enter the Source, Destination, Protocol and Application list details, as shown in Figure 15-16.

Technet24


New Firewall Rule
X
Order 1
Name Rule 1
Type IPv4
Action Drop
Log
Advanced Inspection Profile Select an Advanced Inspection Profile
Source / Destination
Source
Destination
Protocol
Application List
Any
Any
Any
Any

Figure 15-16 *Details for Adding a New Firewall Rule*

Click the **+** sign to the left of **Source** to add the source criteria. For **Type**, choose **Identity** to add an identity list containing your selected SGTs.

The identity lists previously created are offered for selection. As an added touch, SD-WAN Manager displays the SGTs assigned to each identity list, as shown in Figure 15-17.

Technet24

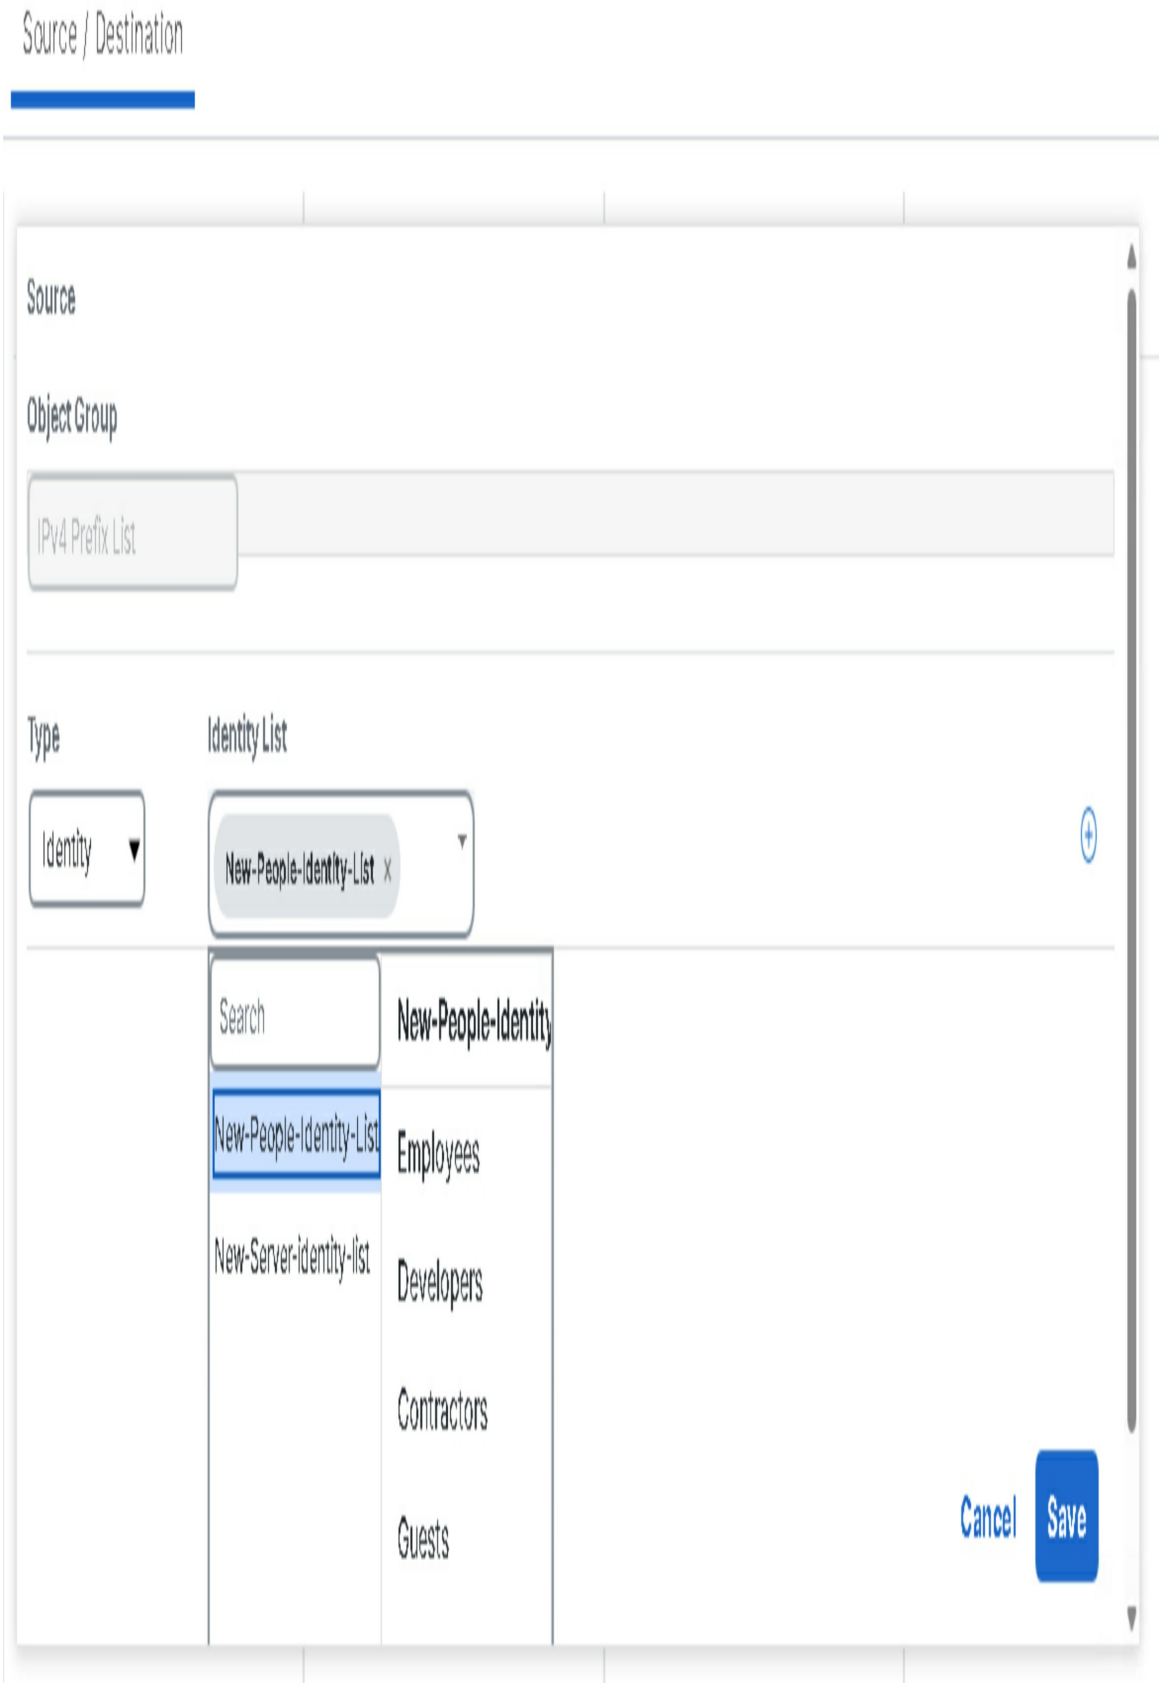

Source / Destination
Source
Object Group
IPv4 Prefix List
Type
Identity List
Identity
New-People-Identity-List x
Search
New-People-Identity
New-People-Identity-List
Employees
New-Server-identity-list
Developers
Contractors
Guests
Cancel
Save

Figure 15-17 *Selecting a Defined Identity List in the Source Criteria*

In the same way, you can select the destination criteria.

After you save the rule, SD-WAN Manager displays the details with the source and destination identity lists selected, as shown in Figure 15-18.

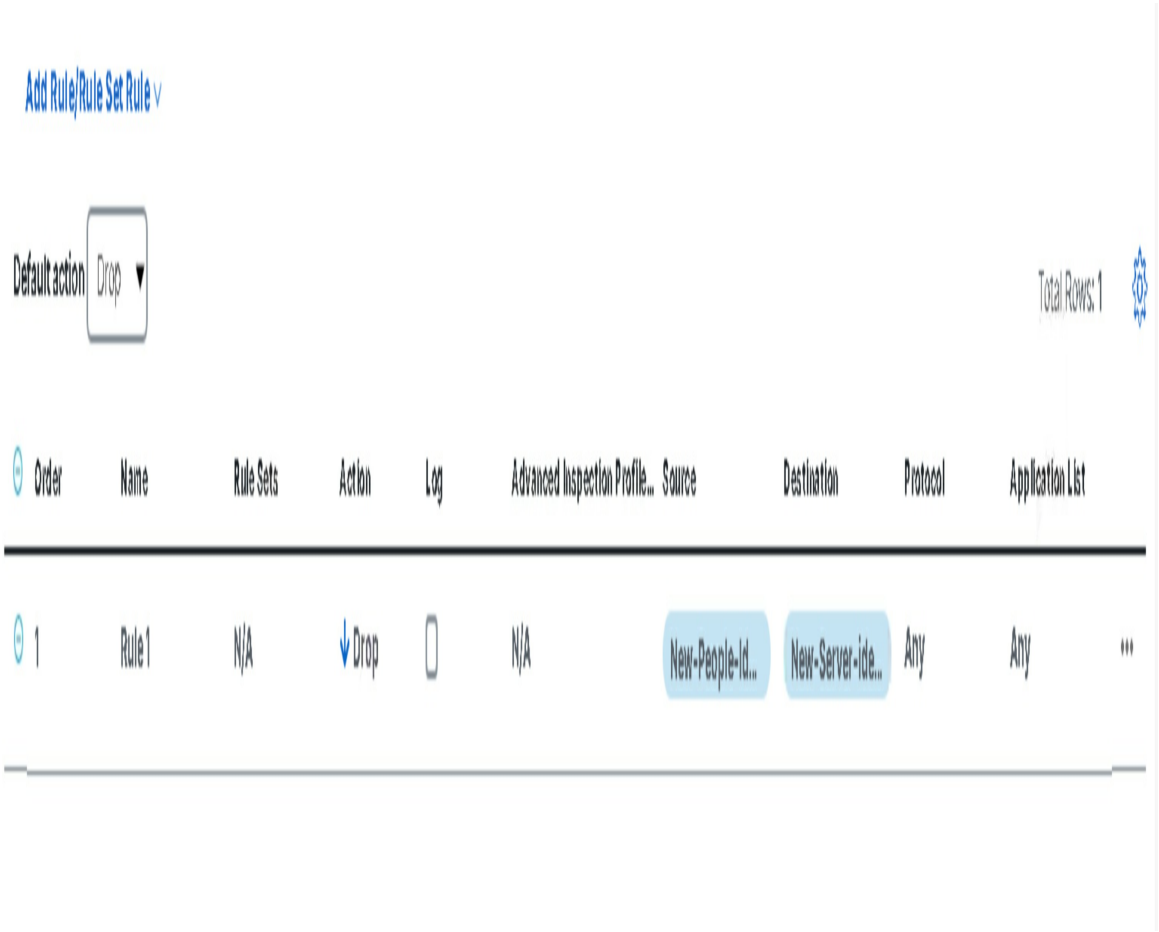


Figure 15-18 *Firewall Rule Within the Unified Security Policy*

SD-WAN Manager Subscribing to User/User Groups

When you're adding an integration with ISE, you have the option to subscribe to **User/User Groups** (see Figure 15-19) rather than **Security Group Tag (SGT)**. This is still an ISE pxGrid subscription but

Technet24

targeting learning user and user group information from an ISE integrated Active Directory system.

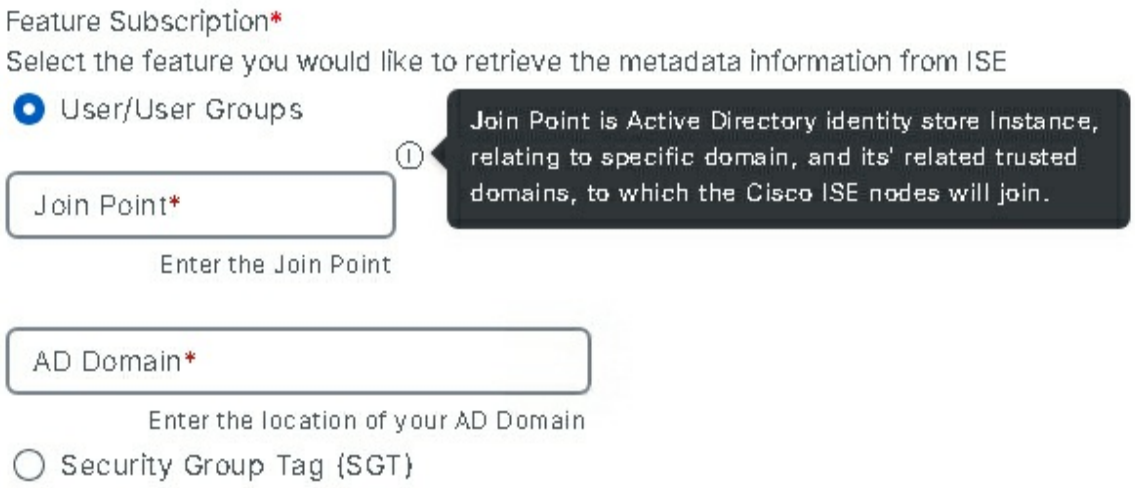


Figure 15-19 *Details to Enter When Selecting to Subscribe to User/User Groups*

When connected, SD-WAN Manager displays the integration status, as shown in Figure 15-20.



Figure 15-20 *SD-WAN Manager to ISE Integration Status*

Technet24

In the **Action** column, selecting the three dots and then **View ISE Data** shows users and user groups learned via pxGrid from ISE but ultimately from the Active Directory connected to ISE. Figure 15-21 shows examples of users downloaded.

View Users / User Groups

Users | User Groups

Users 46

Search Table

As of: Jul 03, 2026 10:34:40 AM

Name	Source
administrator	dev.pseudoco.com
guest	dev.pseudoco.com
krbtgt	dev.pseudoco.com
amar.gupta	dev.pseudoco.com
paul.cole	dev.pseudoco.com
anita.copper	dev.pseudoco.com
hayley.rob	dev.pseudoco.com

Technet24

Figure 15-21 *SD-WAN Retrieved User List from ISE/Active Directory*

Click the **User Groups** submenu to see the user groups downloaded (see Figure 15-22).

View Users / User Groups

Users | **User Groups**

User Groups 67

Search Table

As of: Jul 03, 2026 10:37:02 AM

Name	Source
Users/DnsUpdateProxy	dev.pseudoco.com
Builtin/Certificate Service DCOM Access	dev.pseudoco.com
Users/Domain Users	dev.pseudoco.com
Builtin/Incoming Forest Trust Builders	dev.pseudoco.com
Users/Denied RODC Password Replication Group	dev.pseudoco.com
Builtin/Account Operators	dev.pseudoco.com
DEV/Groups/HR-Staff	dev.pseudoco.com

Technet24

Figure 15-22 *SD-WAN Retrieved User Groups List from ISE/Active Directory*

When integrated, SD-WAN Manager has similar WebSocket connections created as when the **Security Group Tag (SGT)** option is selected. It shows one connection to ISE and one to the Controller, as shown in Figure 15-23.

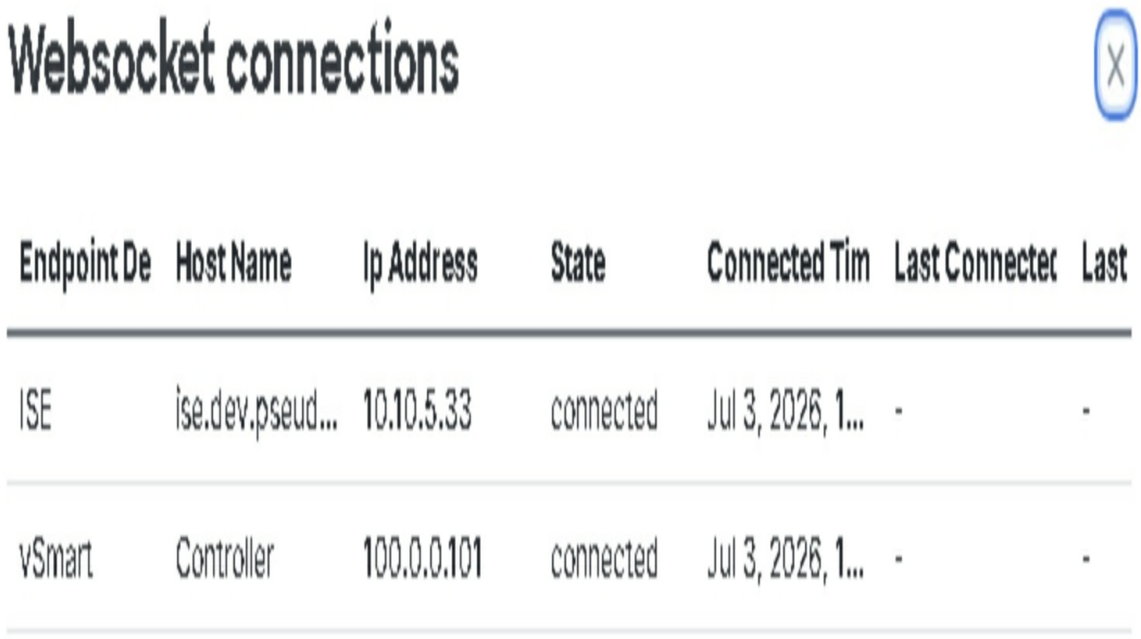


Figure 15-23 *SD-WAN Manager WebSocket Connections*

ISE shows the SD-WAN Manager connectivity under **pxGrid Clients**.

With the **User/User Groups** option, SD-WAN Manager subscribes to the following ISE pxGrid topics:

- /topic/com.cisco.ise.session
- /topic/com.cisco.ise.session.group

This allows the SD-WAN system to learn user session information rather than SGT assignments. The ISE pxGrid subscription for the SD-

WAN Manager is shown in Figure 15-24.

Technet24

Connections Topics

Connections

Rows/Page 5 | 1 /1 | Go | 5 Te

Filter

Connection Name	Connect To	Session Id	Certificate	Subscriptions	Publications	IP Address
-ise-fanout-ise	ise	ise:0	CN=ise.dev.pseudoco.c...	/topic/wildcard		127.0.0.1
-ise-mnt-ise	ise	ise:1	CN=ise.dev.pseudoco.c...	/topic/com.cisco.ise.sessi...	/topic/com.cisco.ise...	10.10.5.33
-ise-fanout-ise	ise	ise:2	CN=ise.dev.pseudoco.c...	/topic/com.cisco.ise.sessi...	/topic/distributed	10.10.5.33
-ise-admin-ise	ise	ise:4	CN=ise.dev.pseudoco.c...	/topic/com.cisco.ise.sessi...	/topic/com.cisco.ise...	10.10.5.33
dCloud-Lab-Modality_vmanage178309961810 5	ise	ise:5	No Certificate	/topic/com.cisco.ise.sessi...		172.16.0.1

/topic/com.cisco.ise.sessi...
/topic/com.cisco.ise.sessi...
/topic/com.cisco.ise.sessi...

Figure 15-24 *ISE Showing pxGrid Topics That SD-WAN Manager Subscribes To*

Ultimately, the retrieved users or user groups can be assigned when you're creating a new identity list and therefore available when you're creating a unified policy in SD-WAN Manager. Figure 15-25 shows how to use users and user groups in an identity list.

Technet24

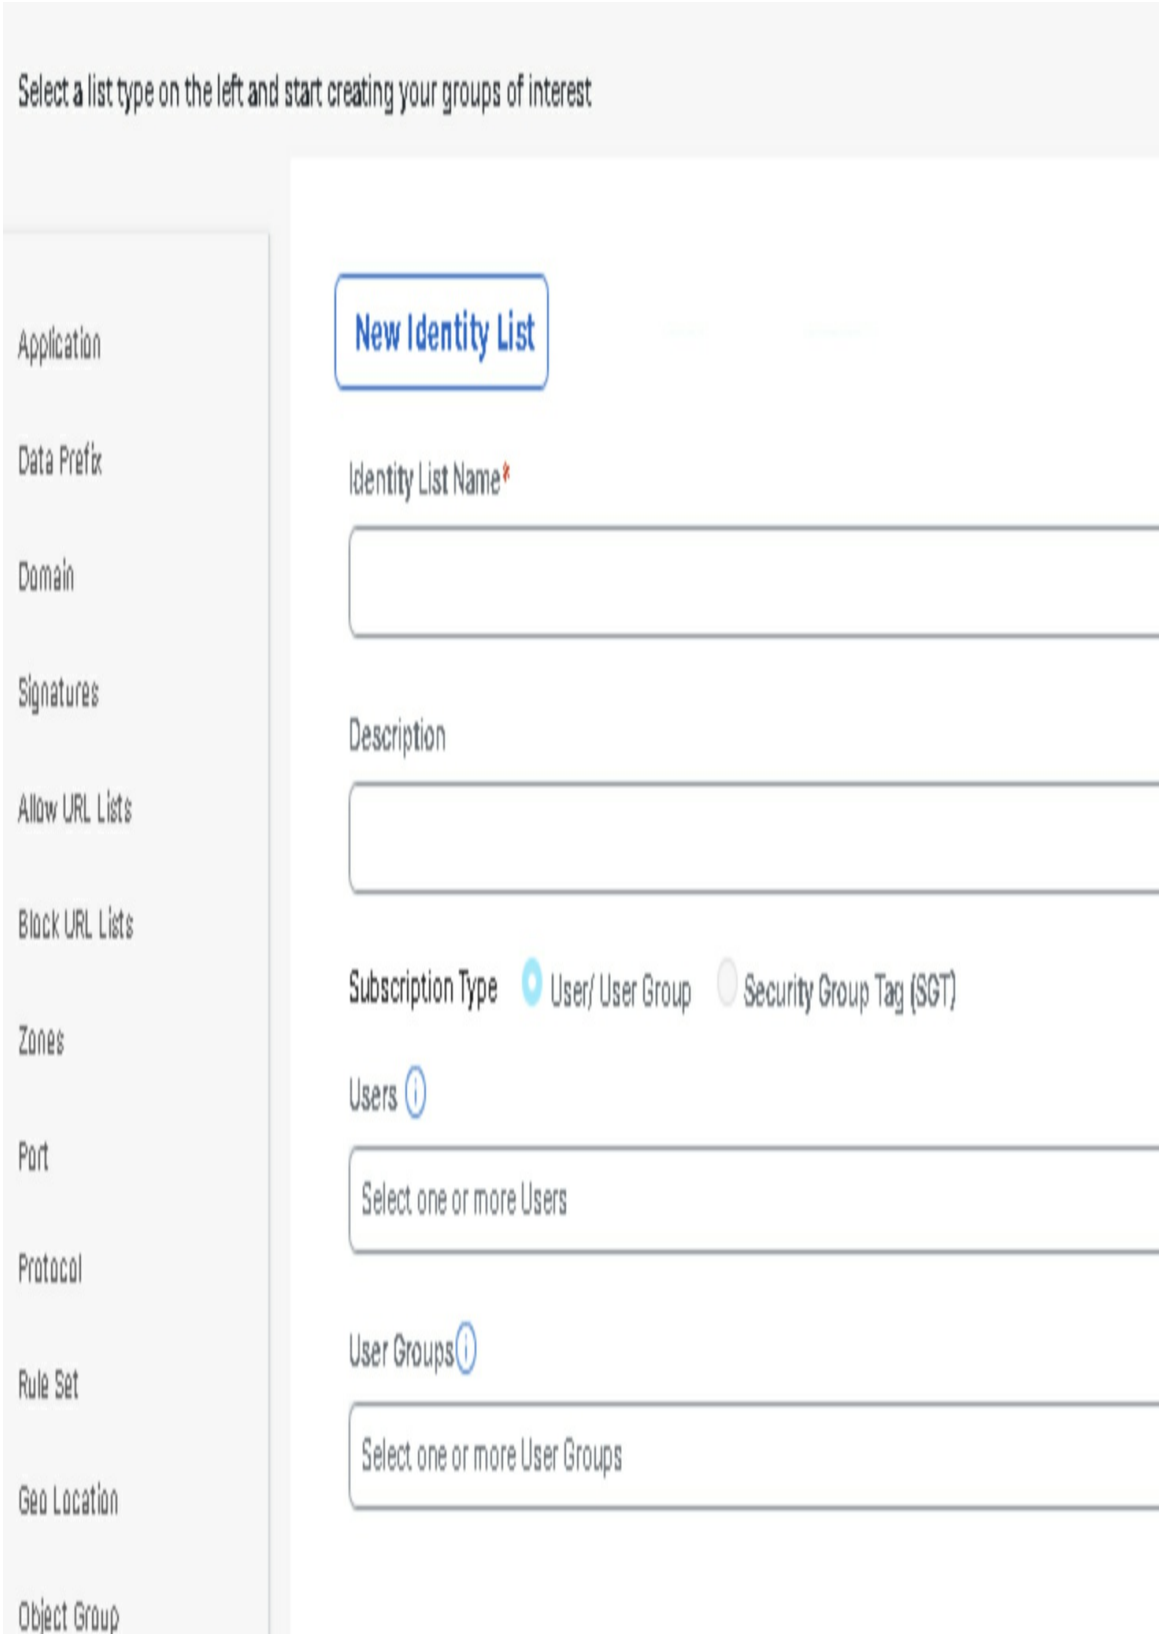


Figure 15-25 *SD-WAN Identity List Showing User/User Group*

Selections

SD-WAN Edge Devices Using SXP

Previously, you saw SD-WAN edge devices learning IP:SGT bindings from the SD-WAN Controller over OMP. You can also use the IOS-XE functionality of the devices to utilize Security Group Tag Exchange Protocol (SXP). You can create an SXP device in ISE to connect to a TCP SXP connection on the edge device, and any bindings known by ISE are propagated to the edge device over the SXP connection. Figure 15-26 shows a command run on an SD-WAN edge device to display the SGT bindings; the bindings are learned from an SXP source.

```
SanJose-BR1Edge1#sh cts role sgt vrf 10 all
Active IPv4-SGT Bindings Information

IP Address              SGT     Source
============================================
1.1.1.1                 13      SXP
10.10.102.12            5       SXP
10.10.102.34            19      SXP
10.10.102.106           20      SXP
10.10.102.109           17      SXP
10.10.102.170           17      SXP

IP-SGT Active Bindings Summary
============================================
Total number of SXP      bindings = 6
Total number of active   bindings = 6
```

Figure 15-26 *SD-WAN Edge Device Showing Bindings Learned from ISE SXP*

Whether the bindings are learned via OMP or SXP, if enforcement is enabled, a security policy is requested from the ISE TrustSec matrix for policies in the SGT matrix column (related to the destination SGT). The network devices making the policy request need to be added into ISE

Technet24

as network devices.

For example, Figure 15-26 and Figure 15-10 include an SXP binding for SGT 13 and an OMP binding for SGT 16. The edge device downloads all policies with SGT 13 and 16 as destinations. In this case, just one policy is present with SGT 13 as a destination and just one with SGT 16 as a destination. Figure 15-27 shows these two policies downloaded to an SD-WAN edge device.

```
SanJose-BR1Edge1#sh cts role per
IPv4 Role-based permissions default:
        Permit IP-00
IPv4 Role-based permissions from group 15:BYOD to group 13:Test_Servers:
        Deny IP-00
IPv4 Role-based permissions from group 15:BYOD to group 16:Administrators:
        Permit IP-00
RBACL Monitor All for Dynamic Policies : FALSE
RBACL Monitor All for Configured Policies : FALSE
```

Figure 15-27 *SD-WAN Edge Device Showing Policies Downloaded from ISE*

Summary

This chapter explores the extension of microsegmentation principles from local area networks to the wide area network, including Cisco Catalyst SD-WAN. It highlights the necessity of maintaining consistent security policies across geographically dispersed sites, particularly where traditional fabric technologies might not be available. The Cisco Catalyst SD-WAN solution facilitates this through a robust architecture comprising the SD-WAN Manager, Controller, Validator, and edge routers, which collectively decouple the control and data planes to ensure granular visibility and policy enforcement.

To ensure identity information is synchronized, Cisco Catalyst SD-WAN ensures that security group tags are propagated across the entire

enterprise, effectively extending the policy domain from a single site to the global network. This propagation allows security policies to remain consistent regardless of where a user or device connects. Furthermore, the architecture provides flexibility by allowing administrators to reclassify incoming traffic based on specific requirements and enforce granular security policies at egress, ensuring that microsegmentation remains effective even as traffic traverses the wide area network.

The integration between the SD-WAN Manager and Cisco Identity Services Engine provides additional benefits. By utilizing Representational State Transfer (REST) and pxGrid, the SD-WAN Manager can programmatically retrieve security group tags or user and user group information from ISE. This process begins with enabling the necessary API and pxGrid services within ISE, followed by configuring the connection within the SD-WAN Manager. Once established, the system uses secure WebSocket connections to exchange identity data, allowing the SD-WAN system to learn IP:SGT bindings or user session details.

To implement these security policies, administrators create identity lists within the SD-WAN Manager, which categorize the learned SGTs or user groups. These lists are then integrated into unified security policies, enabling the creation of zone-based firewall rules based on source and destination identity rather than just traditional IP addresses. IP:SGT bindings are distributed from the SD-WAN Manager to the SD-WAN Controller and subsequently to the edge routers via the OMP, ensuring that security policies are enforced consistently across the WAN.

For hybrid environments, group-based policy can also be dynamically downloaded from the ISE TrustSec matrix when OMP or SXP bindings are learned by the SD-WAN edge devices. In addition to SD-WAN

Technet24

OMP, Security Group Tag Exchange Protocol (SXP) can be configured between ISE and the edge devices, so bindings are propagated directly to the routers. Whether the bindings are learned through OMP or SXP, the edge devices effectively download the relevant policies from the ISE TrustSec matrix. This comprehensive approach ensures that security remains scalable, manageable, and highly effective for modern multicloud and SASE architectures.

Chapter 16. Building Policy-Interconnect with Data Center

Introduction

Administrators often struggle with maintaining consistent security policies when traffic moves between campus environments and data centers. This chapter explores how to bridge those domains using Cisco Identity Services Engine (ISE) and Cisco Application Centric Infrastructure (ACI). By leveraging the Common Policy initiative and pxGrid, you will discover how to synchronize group context, such as security group tags and endpoint groups, to ensure that security policies remain effective regardless of where an endpoint connects.

The following pages walk through the practical steps of setting up and managing these integrations. You will learn how to configure SGT domains for segregating group bindings, establish workload connections between ISE and APIC, and utilize primary and secondary SGTs to scale your security architecture. Beyond the initial setup, the discussion covers how to maintain high availability across your infrastructure and how to monitor the health of your connections using built-in reporting tools. By the end of this chapter, you will have a clear understanding of how to create a unified, scalable security framework that extends from your branch offices into your data center fabric.

Technet24

While users and endpoints connect to the branch or campus domain, resources and applications tend to reside in the cloud or an on-prem data center. How security context is managed across these domains is discussed within this book. This chapter concentrates on when the data center domain is deployed using Cisco Application Centric Infrastructure. Even if Cisco Catalyst Center is used for network orchestration in the campus, the policy integration between domains is established between Cisco ISE and APIC.

The ACI policy constructs consist of groups of resources called endpoint groups (EPGs) and endpoint security groups (ESGs). The equivalent on the ISE side for policy management of users and endpoints is called security groups — applied in network devices by assigning security group tags (SGTs).

Without integration, it is difficult for administrators in one domain to build security policies for traffic destined to or sourced from another. Consider a data center administrator, for example. How can security policies be built for traffic flowing from doctors or sales teams when they are accessing secure resources? The mobile nature of users dictates where their traffic will be sourced from plus what source IP address will be used, so building security policies in the data center for these dynamic allocations is traditionally very challenging.

Cisco has solved this challenge by synchronizing group context between domains. Cisco ISE shares selected security groups (SGTs) along with IP address membership to Cisco APIC, and Cisco APIC shares selected data center endpoint groups (EPGs and ESGs) along with IP address membership to Cisco ISE. Today, within the Cisco Common Policy initiative, Cisco pxGrid is utilized for this integration, providing a real-time, scalable, and robust context-sharing mechanism.

The original ISE-to-APIC integration was delivered circa 2019. That integration has now been superseded due to limits imposed with the original design, such as support for a single APIC, single VRF instance, and single L3Out interface. ISE also only accessed APIC with REST APIs to push and pull data, so sharing was not accomplished using a dynamic method such as pxGrid. At the time of writing, the original integration is still supported for migration purposes, but for new ISE-to-APIC integrations, the Common Policy methodology must be used. An overview of the integration is shown in Figure 16-1.

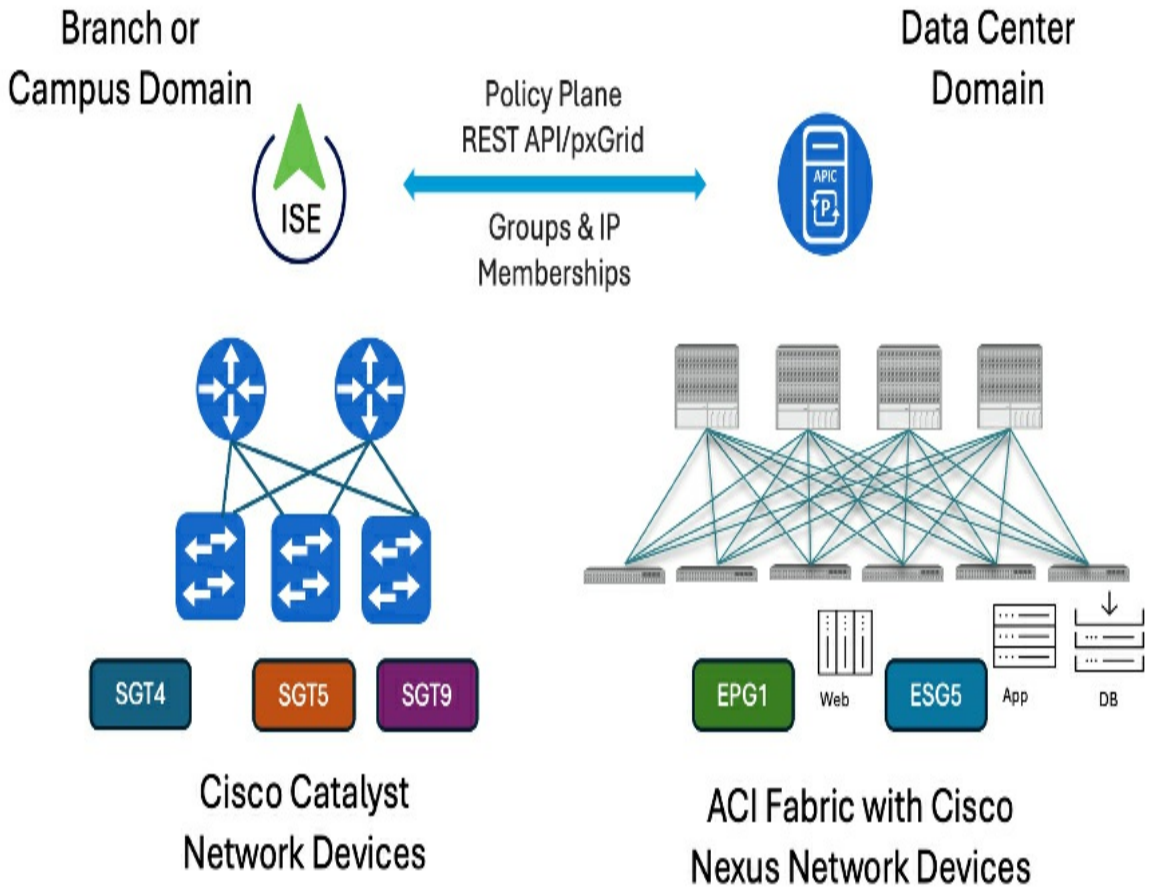


Figure 16-1 *Context Exchange Between Domains*

This new ISE-to-ACI integration offers customers many use cases, leading to desirable outcomes. First, the integration supports multiple

Technet24

ACI tenants, multiple VRFs, and multiple L3Out interfaces. Filters (discussed later) allow only the applicable groups to be shared. Additionally, multiple connections can be added to multiple APIC clusters, allowing multiple ACI fabrics in multiple data centers and sites to be integrated. Enforcement is made flexible by allowing permit or deny rules in both the campus and/or the data center orchestrated by the applicable controller. Furthermore, enforcement using this learned context is extended to the likes of firewalls, Cisco SD-WAN controllers, and Cisco's Security Service Edge (SSE) solution Cisco Secure Access, by utilizing the pxGrid propagation mechanism of ISE.

Prerequisites for Integrating ISE and APIC

Before the two domains are integrated, some prerequisites must be met:

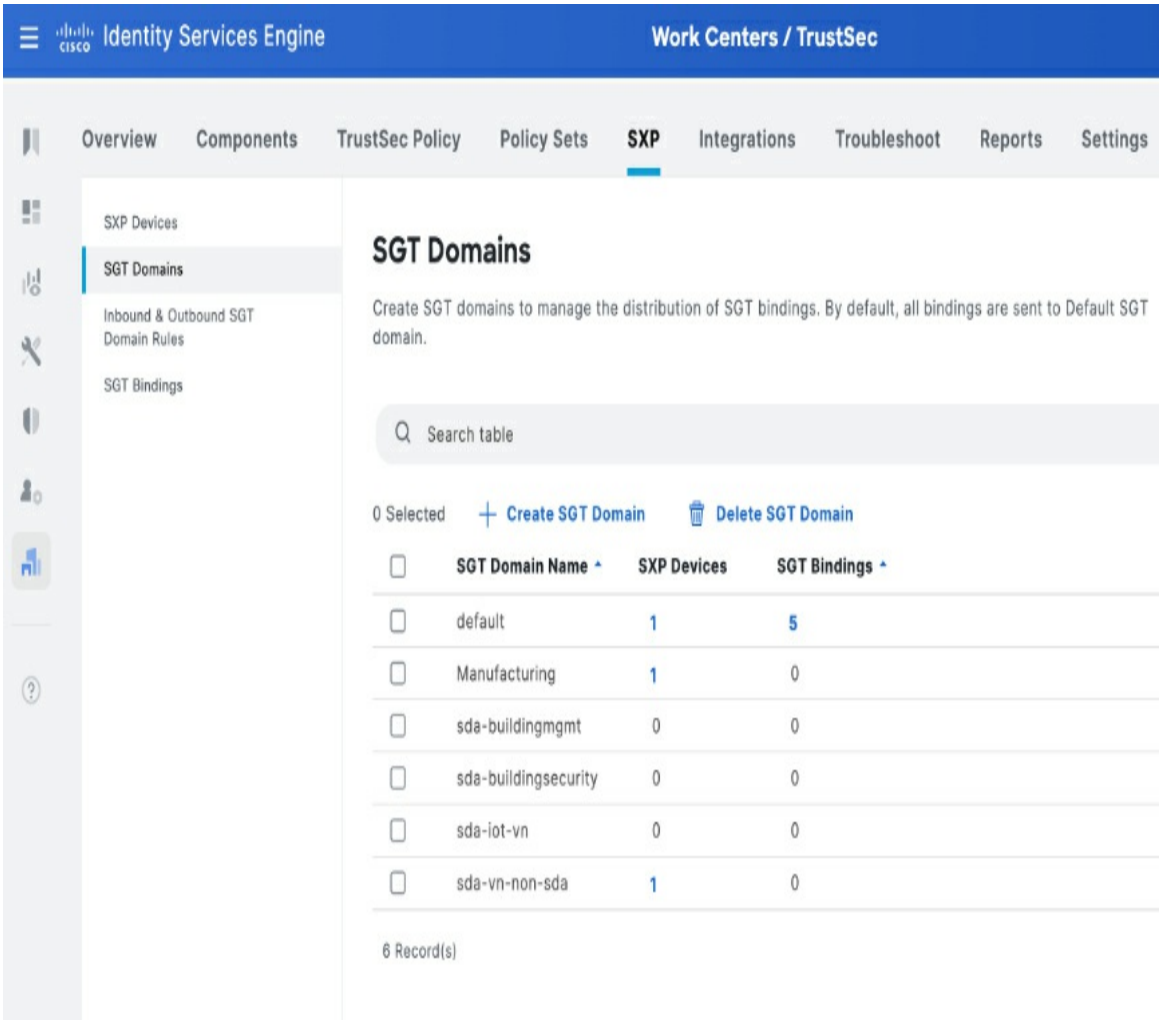
- The minimum software version must be running on both ISE and APIC cluster. For ISE, that is release 3.4 patch 1, and for APIC, it is release 6.1(2).
- Security Group Tag Exchange Protocol (SXP) services and pxGrid must be enabled on ISE.
- DNS must contain entries for both ISE and APIC, especially to ensure APIC is able to resolve the FQDN name of the ISE node where the pxGrid service is running.

Configuring nondefault inbound SGT domain rules is an optional task. By default, Cisco ISE includes an inbound SGT domain rule that maps ACI group memberships to the default SGT domain. Should your deployment require the mapping of bindings across multiple SGT domains, you can achieve this by defining custom inbound SGT domain rules.

ISE SGT Domains

Think of an SGT domain as a space in ISE to keep or store group memberships, otherwise known as IP-to-group bindings. Multiple SGT domains can be created on ISE, used to segregate bindings, and ultimately provide the capability of sending specific bindings to different destinations.

You can manage SGT domains in ISE by navigating to **Work Centers > TrustSec > SXP > SGT Domains** (see Figure 16-2). Before ISE 3.4, these domains were called Security Group Tag Exchange Protocol (SXP) domains.



Technet24

Figure 16-2 *ISE SGT Domains*

You also can create an SGT domain under **Work Centers > TrustSec > SXP > SXP Devices > Assign SGT Domain**, as shown in Figure 16-3.



Figure 16-3 *Assigning an SGT Domain from SXP Devices*

SXP devices are remote network devices that ISE has an SXP peering with. SXP is a peer-to-peer out-of-band IP:SGT binding propagation protocol running over TCP port 64999. Peer examples are Cisco switches, Cisco routers, Cisco ASAs, or other ISE deployments. Each of these SXP devices specified in ISE has one or more SGT domains associated with it. If the SXP device is specified as being an SXP speaker in ISE, then the IP:SGT bindings received from that network device will be placed in the assigned SGT domain(s). If the SXP device is specified as being an SXP listener in ISE, then the IP to SGTs within those specified ISE SGT domains will be sent to the network device. An SXP device in ISE can be an SXP speaker and a listener at the same time, the peer role being set to **Both** to achieve this result. We recommend enabling binding propagation only in the direction required for your specific use case, because this approach optimizes the solution for better scalability. You should enable propagation both ways *only* if that is indeed the required outcome.

If multiple SGT domains are to be used in the design, they can be defined before adding an ACI workload connection and set as the destination in inbound SGT domain rules. To add inbound SGT domain rules, in ISE navigate to **Work Centers > TrustSec > SXP > Inbound & Outbound SGT Domain Rules**; then select the **Inbound SGT Domain Rules** submenu.

To add an inbound SGT domain rule, select the option **Add Inbound Rule**. Provide a rule name and select the SGT domain destination that learned IP:SGT bindings are to be placed in ISE. Last, create rule configuration, or criteria, used to select the group information that is to be learned. In Figure 16-4, the group information will be learned from the ACI connection called ISE_APIC1, from the Demo Tenant and VRF called vrf-1, and the IP:SGT bindings will be placed in the Manufacturing SGT domain.

Technet24

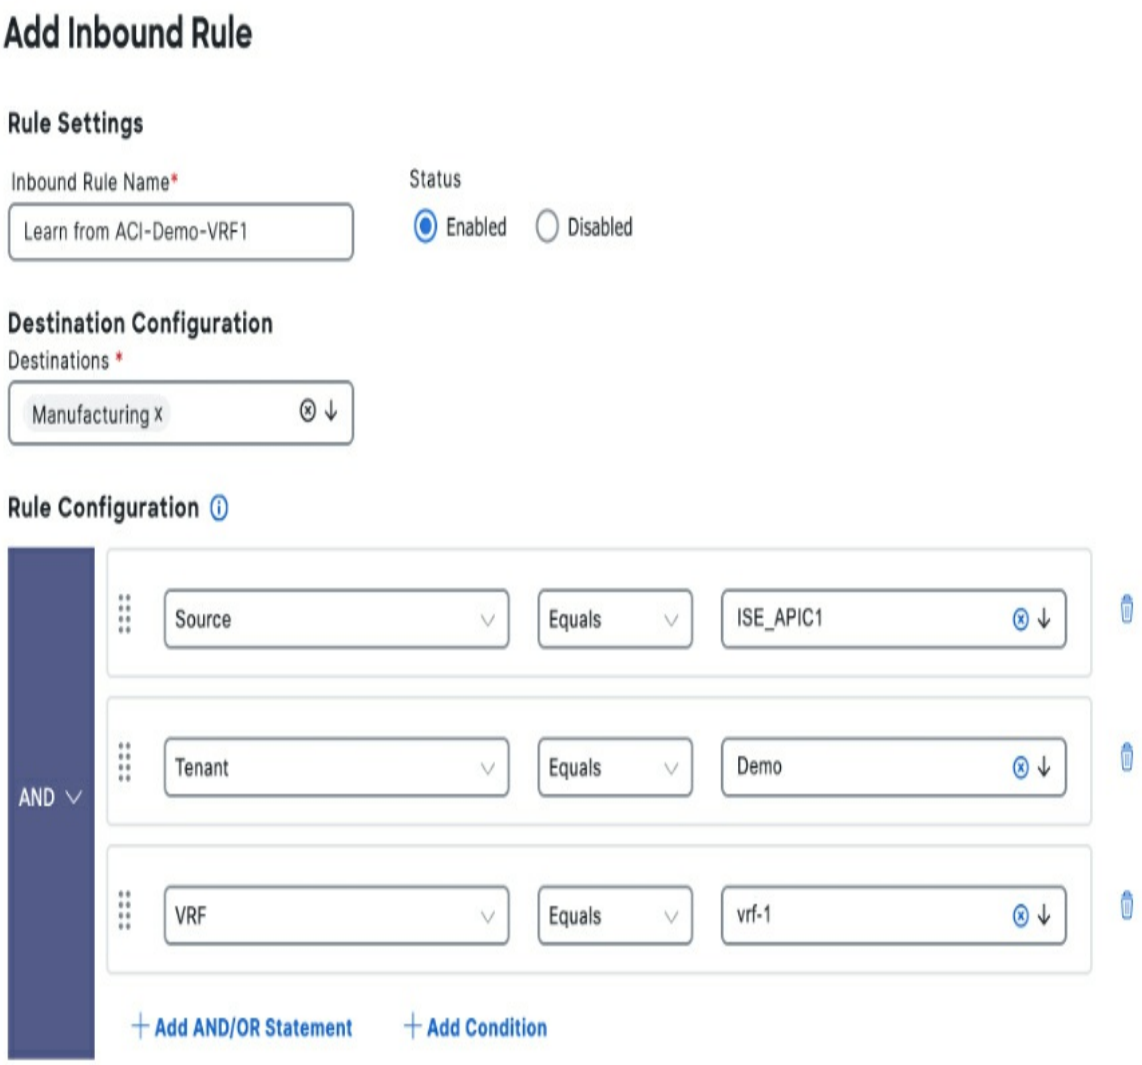


Figure 16-4 *ISE Inbound SGT Domain Rule*

The inbound SGT domain rule table works by using multimatch. If any of the rules match (including the default rule), then they will be acted upon. This means that all learned bindings will always be placed in the default SGT domain as instructed by the default inbound SGT domain rule, as well as any other SGT domain as dictated by the user-defined rules.

If multiple SGT domain rules are not required, the default inbound SGT domain rule will be used, sending received ACI bindings into the default ISE SGT domain. If ISE SGT domains are created on Day-N, the

ISE inbound SGT domain rules can be edited or added, and all bindings, previously received or to be received in the future, will be placed in the new SGT domain(s).

Adding Workload Connection in ISE for APIC Integration

To integrate the two domains, you, as administrator on ISE, create a workload connector to an APIC cluster. Multiple connectors can be added to multiple APIC clusters. The APICs can be standalone, part of a multipod fabric, or part of a multisite fabric where subnets are not stretched between sites. When you're adding a workload connection, the following input is requested as part of an ISE workflow:

1. Provide a name for the connection.
2. Include the FQDN or IP address for APIC. Enter the address for one APIC in the cluster. This address can be the APIC out-of-band or in-band IP address. The IP addresses of the other APICs in the cluster will be retrieved automatically by ISE and used for high availability purposes.
3. Enter an APIC username and password (should be admin or a user with admin privileges).
4. Validate the ACI certificate. Validation requires the APIC certificate to be added to the ISE certificate store.
5. Use the naming convention for SGT names for EPGs and ESGs published by APIC. By default, the SGT name will be the EPG/ESG name followed by the tenant name and then the type (EPG or ESG). The name can be created from other APIC objects like the application profile and VRF or user-defined prefix or suffix. The maximum character limit of the SGT name was restricted to 32 if

Technet24

a Catalyst Center release prior to 2.3.7.6 was integrated with ISE. Catalyst Center release 2.3.7.6 introduced a character limit of 64.

6. From a list of all EPGs and ESGs, select the entries ISE is to subscribe to.
7. Allocate an SGT numbering range in ISE for learned EPGs and ESGs.

As can be seen from point 6, the groups that ISE learns from ACI are selected when the workload connection is added; you can edit the connection if changes are to be made to that group selection. No groups are shared from ISE to ACI at this stage; that is accomplished via ISE outbound SGT domain rules, described later.

After the workload connection is established, the status is shown in the ISE Workload Connections UI, shown in Figure 16-5.

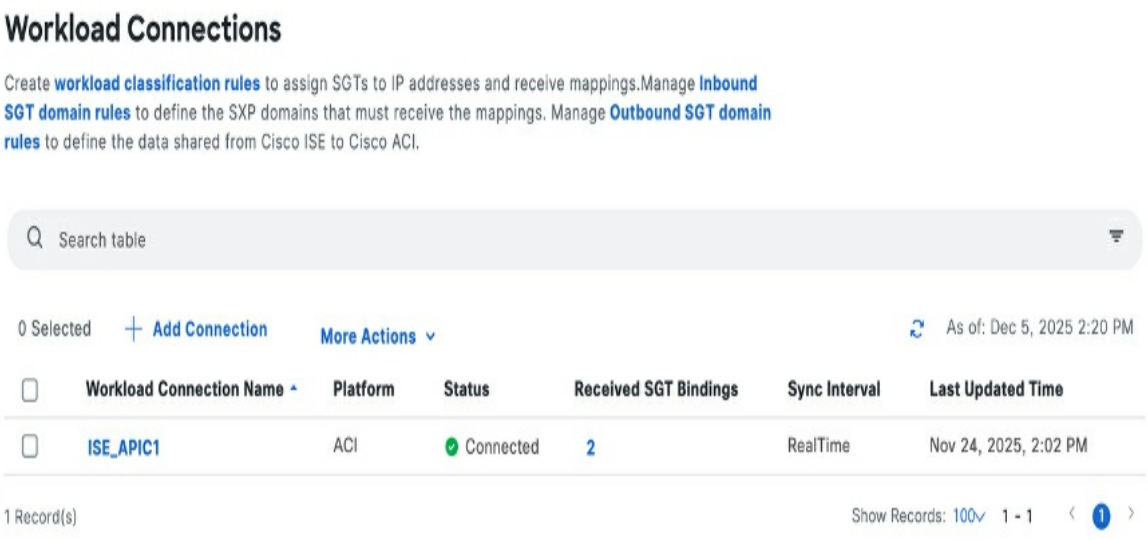


Figure 16-5 *ISE Workload Connection Status*

Status is also shown in the APIC UI, with a particular mention of the new **Integrations** menu, as shown in Figure 16-6.



Figure 16-6 *APIC Integrations Menu*

The **Connection Details** submenu will show a pxGrid **Publisher** role as groups are shared, or published, to ISE.

A pxGrid subscriber role will also be displayed after ISE outbound SGT domain rules have been created and groups shared to APIC.

When ISE learns groups from ACI, they are placed in the Security Groups table. The name of the group and assigned number are defined as configured in the workload connection settings, as shown in Figure 16-7.

Technet24



Figure 16-7 *ISE Security Groups*

As well as groups being learned, the group memberships are learned (that is, the IP-to-EPG and IP-to-ESG bindings). ISE will receive these memberships, normalize them into IP:SGT bindings, and store them in the ISE SGT Bindings table, as shown in Figure 16-8.

Technet24

SGT Bindings

This page displays all the configured SGT bindings, giving you an easy view of assigned primary SGTs, the source of the learned data, and any workload classifications applied. You can review classification details and the specifics of defined inbound & outbound SGT domain rules by clicking the corresponding column values.

Rows/Page 5 | 1 /1 | Go | 5 Total Rows

Add SGT Domain filter | Manage SGT Domain filters | Filter

IP Address	Primary SGT	VN	Learned From	Learned ...	SGT Domain	PSNs Involved
10.1.1.126/32	Demo_WebEPG_ACI ...	DEFAULT_VN	10.1.101.37,7.5.9.136	Session	default	Kernow-ISE3-4
10.2.2.127/32	Demo_ClientEPG_AC...	DEFAULT_VN	10.1.101.37,7.5.9.136	Session	default	Kernow-ISE3-4
10.6.5.185/32	RandD (115/0073)	DEFAULT_VN	10.1.101.37,1.1.1.8	Session	default	Kernow-ISE3-4
192.168.10.10/32	Employees (4/0004)		10.1.101.37	Local	default	Kernow-ISE3-4
192.168.110.10/32	Finance (16/0010)		10.1.101.37	Local	default	Kernow-ISE3-4

Figure 16-8 *ISE SGT Bindings*

Once the bindings are in the table, ISE will share subsets to destinations defined in the outbound SGT domain rules, SXP Devices table, and to suitably subscribed pxGrid clients.

ISE Primary and Secondary SGTs via the Workload Classification Table

From release 3.4, ISE utilizes the concept of a primary SGT and secondary SGT(s). When ISE learns of a group from ACI, ISE stores this as a primary SGT. This primary SGT is published out via SXP and out the pxGrid session topic as Security Group. Secondary SGTs (there can be more than one) are not sent out via SXP currently because SXP does not yet have that capability. The network device data plane does not have the capability to propagate the secondary SGT either, so the Cisco Meta Data (CMD) field in the frame only carries the source primary SGT today. However, ISE can send the secondary SGT(s) out via the pxGrid session topic as an ordered array named Secondary Security Groups. At the time of writing, there are no known consumers of these secondary SGTs yet, but in the future, firewall products could make use of this context via pxGrid and ultimately use access rules such as matching SGT A AND SGT B.

One great use case for this is when multiple sites are defined using SGTs and there are varying security requirements per site. For example, a customer has 20 sites with 100 user and endpoint roles per site. At first glance, it may seem as if 20 × 100 = 2,000 SGTs would be required to satisfy that requirement. However, you could assign a secondary SGT to be the site number and the primary SGT to be the role within the site, dramatically reducing the number of SGTs required down to about 120, which is a 94 percent saving in group numbers.

Technet24

This approach would help with scale but also with operational management of the deployment.

When ISE learns of groups from ACI, ISE always learns the EPGs and ESGs as primary SGTs. ISE treats them as SGTs have been treated for the past few years in terms of storage and propagation. If you want to assign secondary SGTs, you can use the ISE Workload Classification table. Figure 16-9 illustrates how the Workload Classification feature operates in ISE and displays the one default disabled rule.

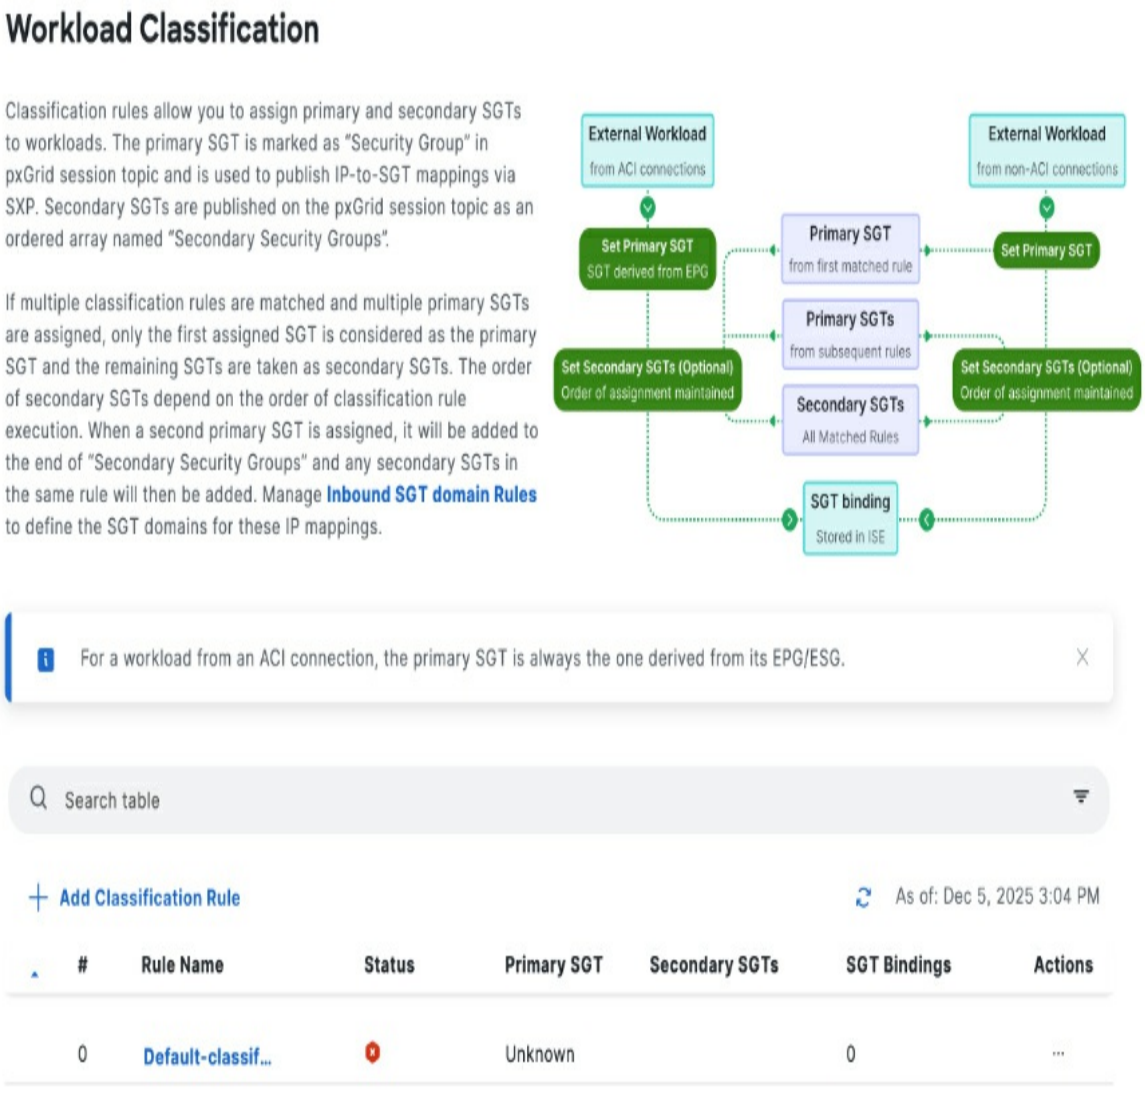


Figure 16-9 *ISE Workload Classification Table*

Rules and conditions defined in that table allow specific EPGs and ESGs to be selected, and any SGTs assigned in the table become secondary SGTs, as shown in Figure 16-10.

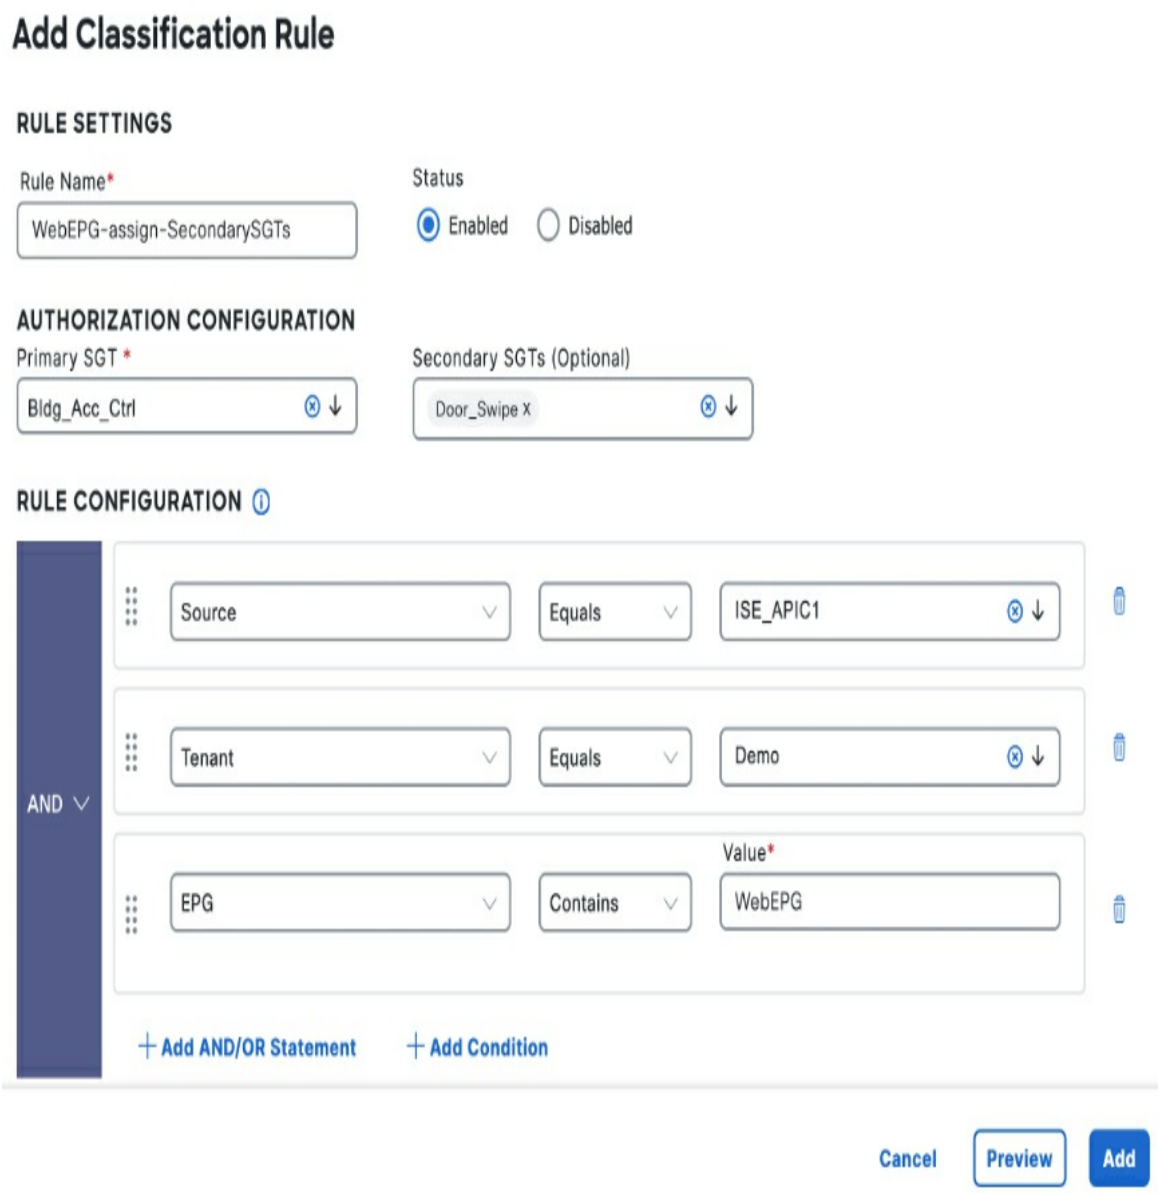


Figure 16-10 *ISE Classification Rule*

In Figure 16-10, the first assigned SGT in the Authorization Configuration section is labeled as the primary SGT. This actually becomes the first secondary SGT because the primary SGT is allocated

Technet24

from the learned EPG or ESG itself. The second assigned SGT in the Authorization Configuration section, labeled Secondary SGTs, becomes the second secondary SGT.

As stated previously, at the time of writing, there are no consumers of this secondary SGT yet, so the use of the Workload Classification table in ISE is of little practical use for this use case currently; it is more useful when used with classifying workloads such as from AWS, Azure, GCP, or vCenter.

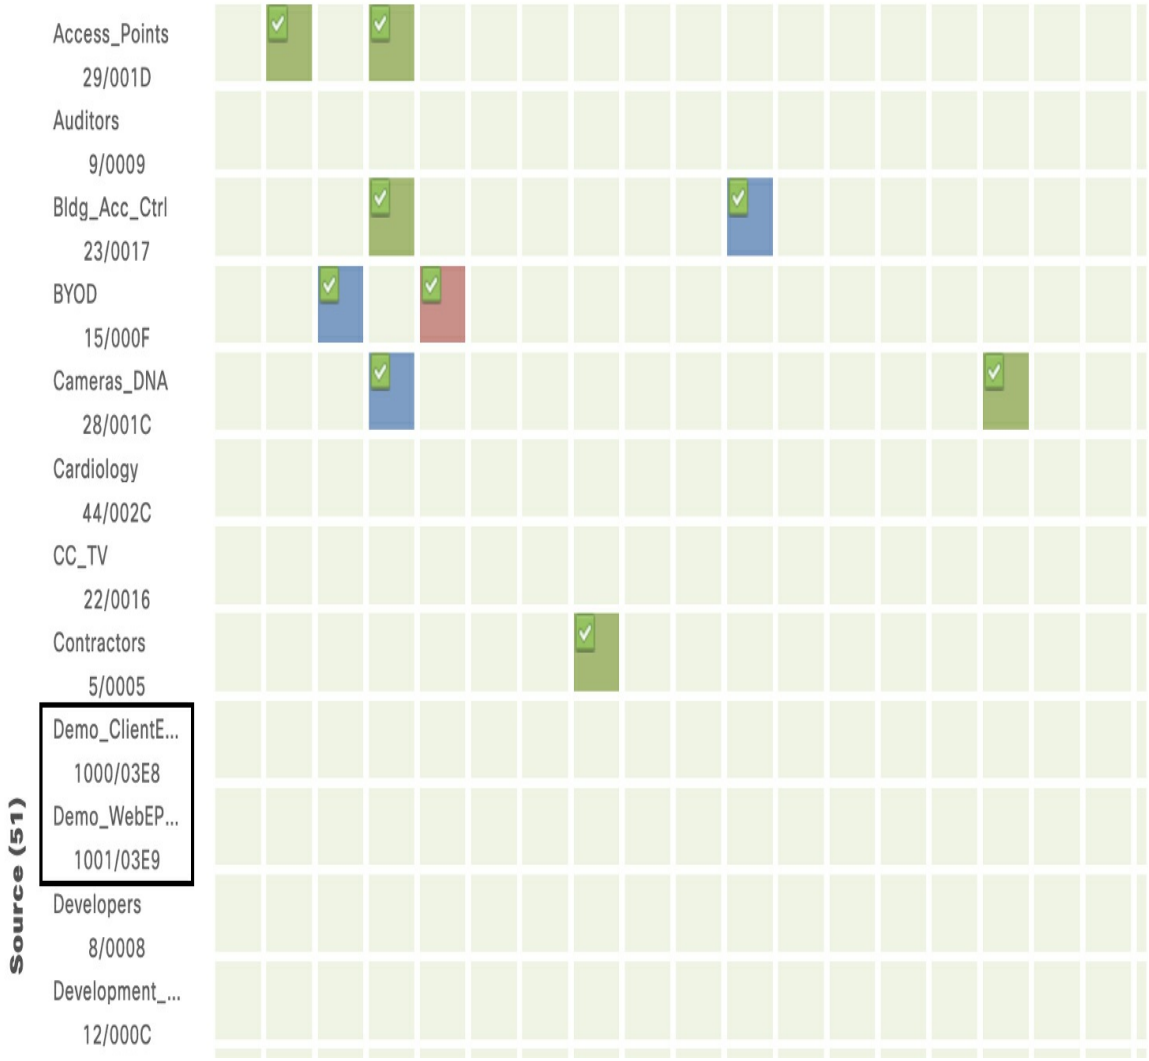
Learned Groups in ISE Available in Matrix and Authorization

After ISE has learned groups from ACI, the groups are added to the Security Groups table, and they become available in the TrustSec policy matrix, as shown in Figure 16-11

Production Matrix

Populated cells:16

Create Policy | Monitor All Disabled | Deploy | Verify Deploy | Default Policy: Permit IP



Technet24

Figure 16-11 *Learned Groups in ISE Available in the TrustSec Matrix*

In this way, group-based policies can be built using the learned groups. For example, Figure 16-12 shows how a policy could be built to permit or deny certain TCP or UDP ports sourced from a group assigned by ISE to a group learned from ACI and could possibly be enforced on a campus gateway router between the campus and data center domains.

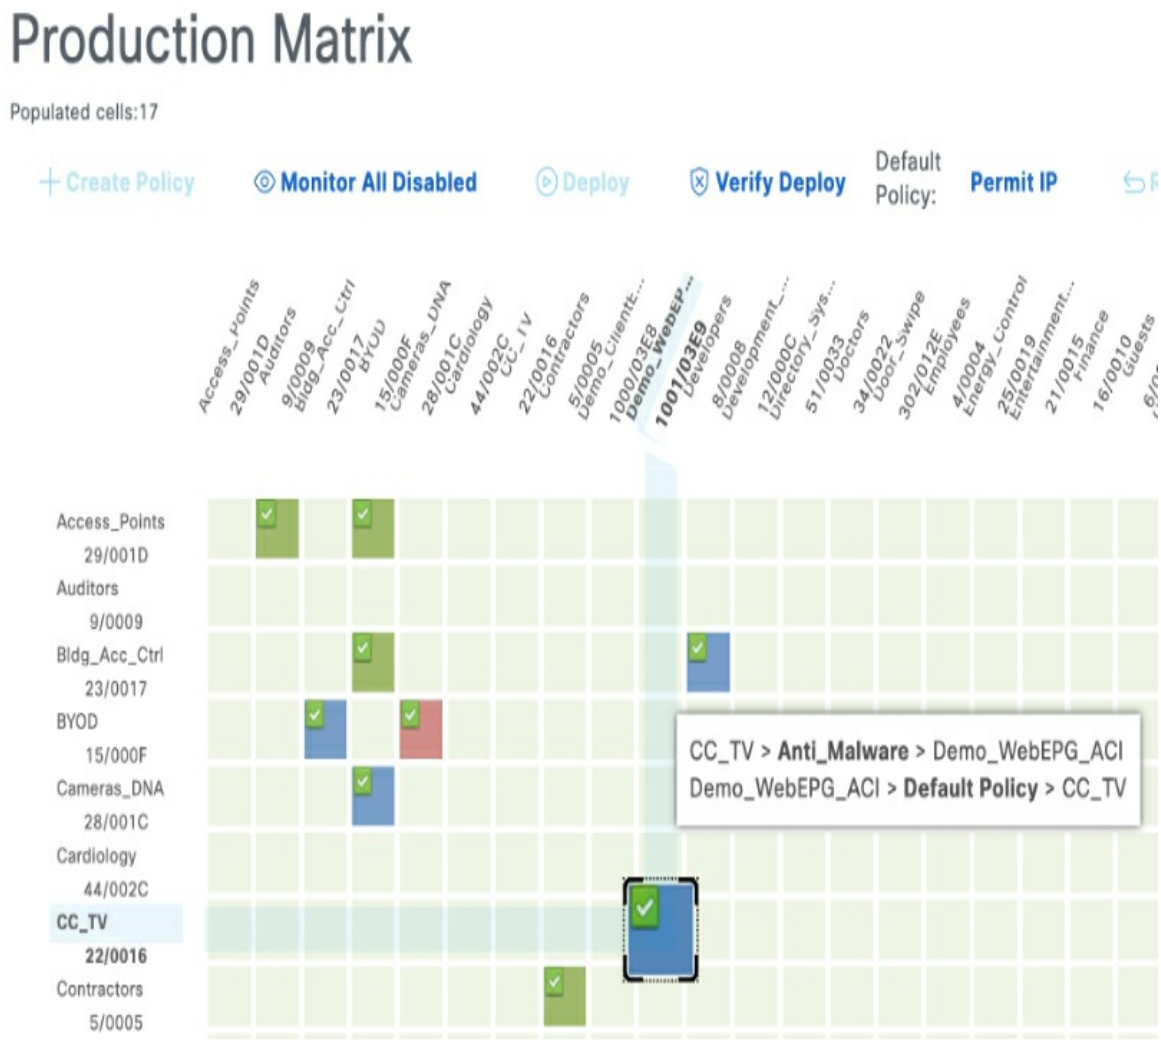


Figure 16-12 *Policy in ISE Using a Group Learned from APIC as a Destination*

Another example, shown in Figure 16-13, is a policy sourced from a group learned from ACI to a group assigned by ISE; this could be enforced at the campus edge.

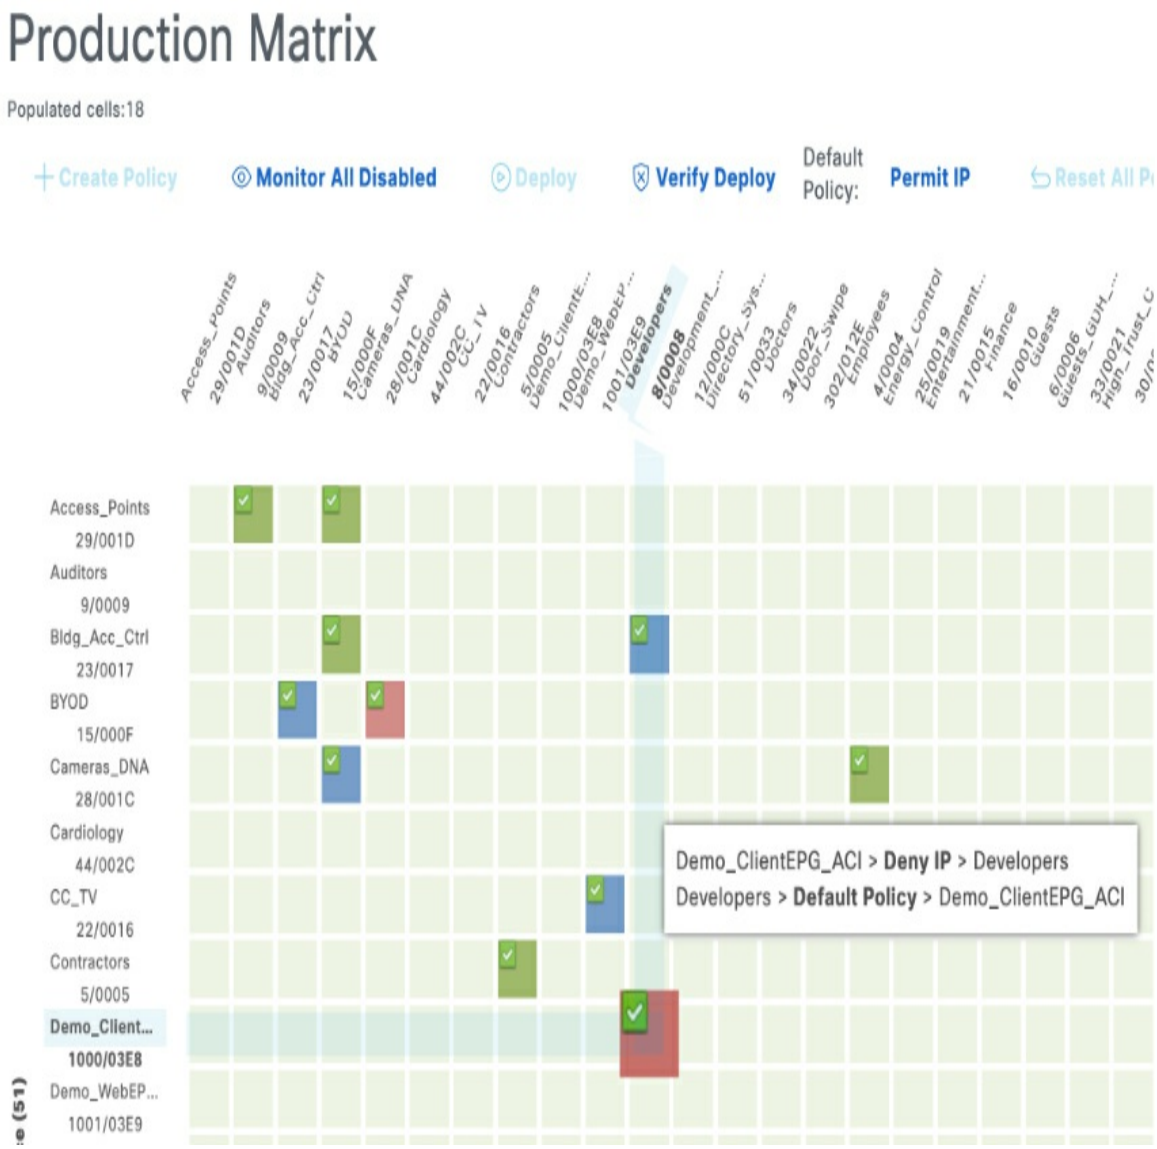
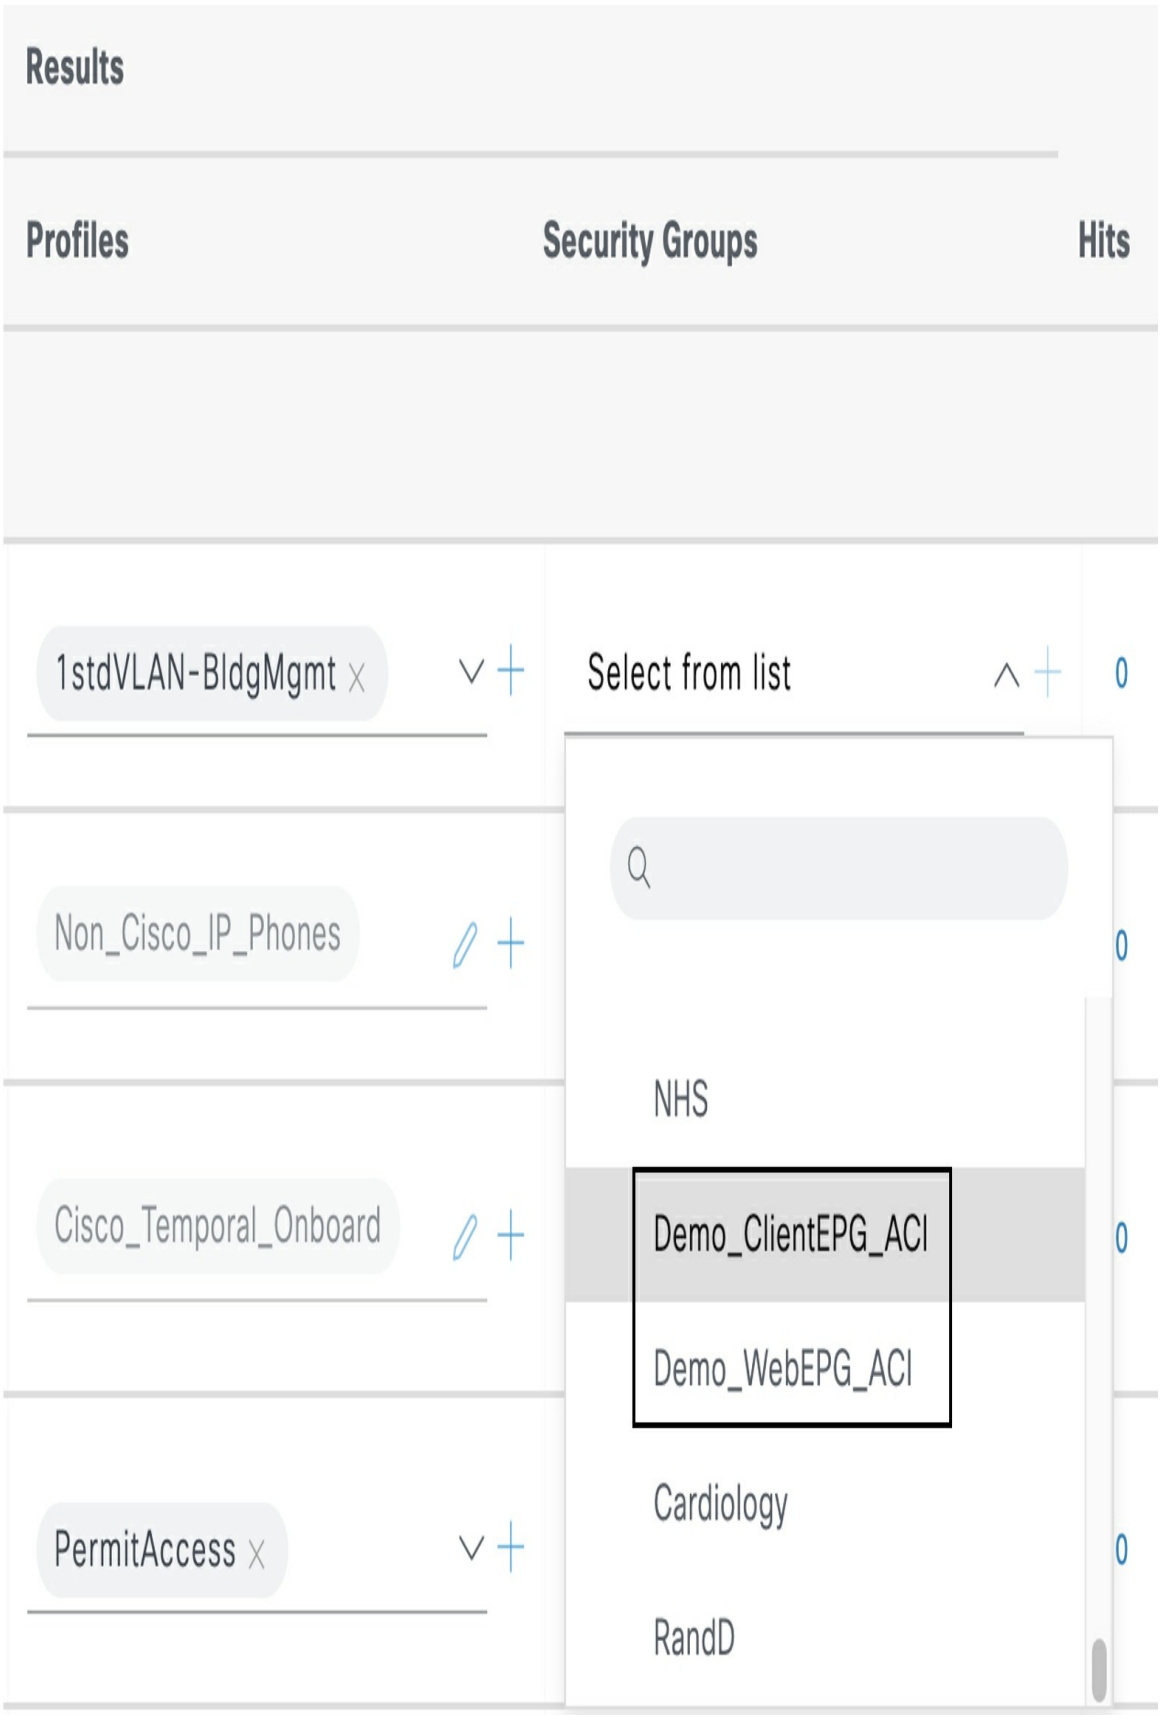


Figure 16-13 *Policy in ISE Using a Group Learned from APIC as a Source*

The learned groups from ACI have been seen to be available in the TrustSec policy matrix, but they are also available to be selected in the ISE authorization profiles and authorization rules. See the example in

Technet24

Figure 16-14. A use case for this is when endpoints or services within the different domains require the same classification. Consider PCI servers, for example. They may be deployed within the data center as well as in the campus. Being able to assign the same group to these servers in the campus via ISE authorization promotes consistency and provides a better integrated solution.


Results
Profiles
Security Groups
Hits
1stdVLAN-BldgMgmt
Select from list
0
Non_Cisco_IP_Phones
0
NHS
Demo_ClientEPG_ACI
Demo_WebEPG_ACI
Cisco_Temporal_Onboard
0
Cardiology
PermitAccess
0
RandD

Technet24

Figure 16-14 *Learned ACI Groups Available in ISE Authorization Rule*

The security group result can be selected within the authorization rule itself, as shown in Figure 16-14. Other possible results of authorization can be selected within the assigned authorization profile. The 1stdVLAN-BldgMgmt authorization profile is assigned in the top rule shown in Figure 16-14. You can find authorization profiles in ISE by navigating to **Policy > Policy Elements > Results > Authorization > Authorization Profiles**.

The capability to select the **Security Group** within the **Common Tasks** section of the **Authorization Profile** (see Figure 16-15) was introduced when ISE added integration support for Cisco Catalyst Center. This made it easier to provision the security group, virtual network (VN), and VLAN as part of the endpoint onboarding integration workflow. Whether ISE is integrated with Cisco Catalyst Center or not, it is a choice to assign the security group within the authorization rule or in the authorization profile. Even though it is possible to assign a security group in both the authorization rule plus in the assigned authorization profile for that rule, doing so is not supported. This approach instructs ISE to send an authorization reply with two security groups for the same endpoint, and it is not supported on the network access devices currently. The network access device may accept this authorization reply, but it is undetermined which security group will actually classify the endpoint traffic.


Security Groups
Common Tasks
DACL Name
IPv6 DACL Name
ACL (Filter-ID)
ACL IPv6 (Filter-ID)
Contractors
Demo_ClientEPG_ACI
Demo_WebEPG_ACI
Developers
Development_Servers
Directory_Systems
Demo_ClientEPG_ACI
Security Group
Virtual Network:
VLAN

Technet24

Figure 16-15 *Learned ACI Groups Available in ISE Authorization Profiles*

Now we are faced with the question of whether there is a recommendation for configuring the security group in the authorization rule or in the authorization profile. A recommendation has not been published so can be thought of as an administrator's choice. However, if multiple SGTs are deployed per VLAN, configuring the SGT assignment in the authorization rule allows you to minimize the number of authorization profiles, because the same profile (purely assigning the VLAN) can then be used for multiple rules.

APIC Learning Groups from ISE

So far, the discussion has covered ISE learning groups from ACI. We have already stated that no ISE SGTs are sent to ACI by default. The SGT groups sent to ACI are controlled via the ISE outbound SGT domain rules. The underlying mechanism behind the table uses server-side filtering. The server in this case is the pxGrid node, not the APIC client. The server/pxGrid node actually carries out the filtering, but the ISE table settings inform APIC which filter criteria to use when requesting the group information from ISE. So, the ISE outbound SGT domain rule table includes the filter criteria (rule configuration), a rule name, and the APIC L3Out interface where APIC is to receive the group information.

When APIC receives the SGTs, they are normalized into external EPGs for use in policy.

There is another aspect to the ISE outbound SGT domain rule that is apparent only when the rule has the name, conditions, and destination L3Out defined. The UI expands and allows ACI contract selection(s) to be made, where ISE retrieves all relevant providers and consumers for

the chosen L3Out, allowing one of each to be selected per shared group to establish the necessary communication policy. See the selections being made in Figure 16-16.

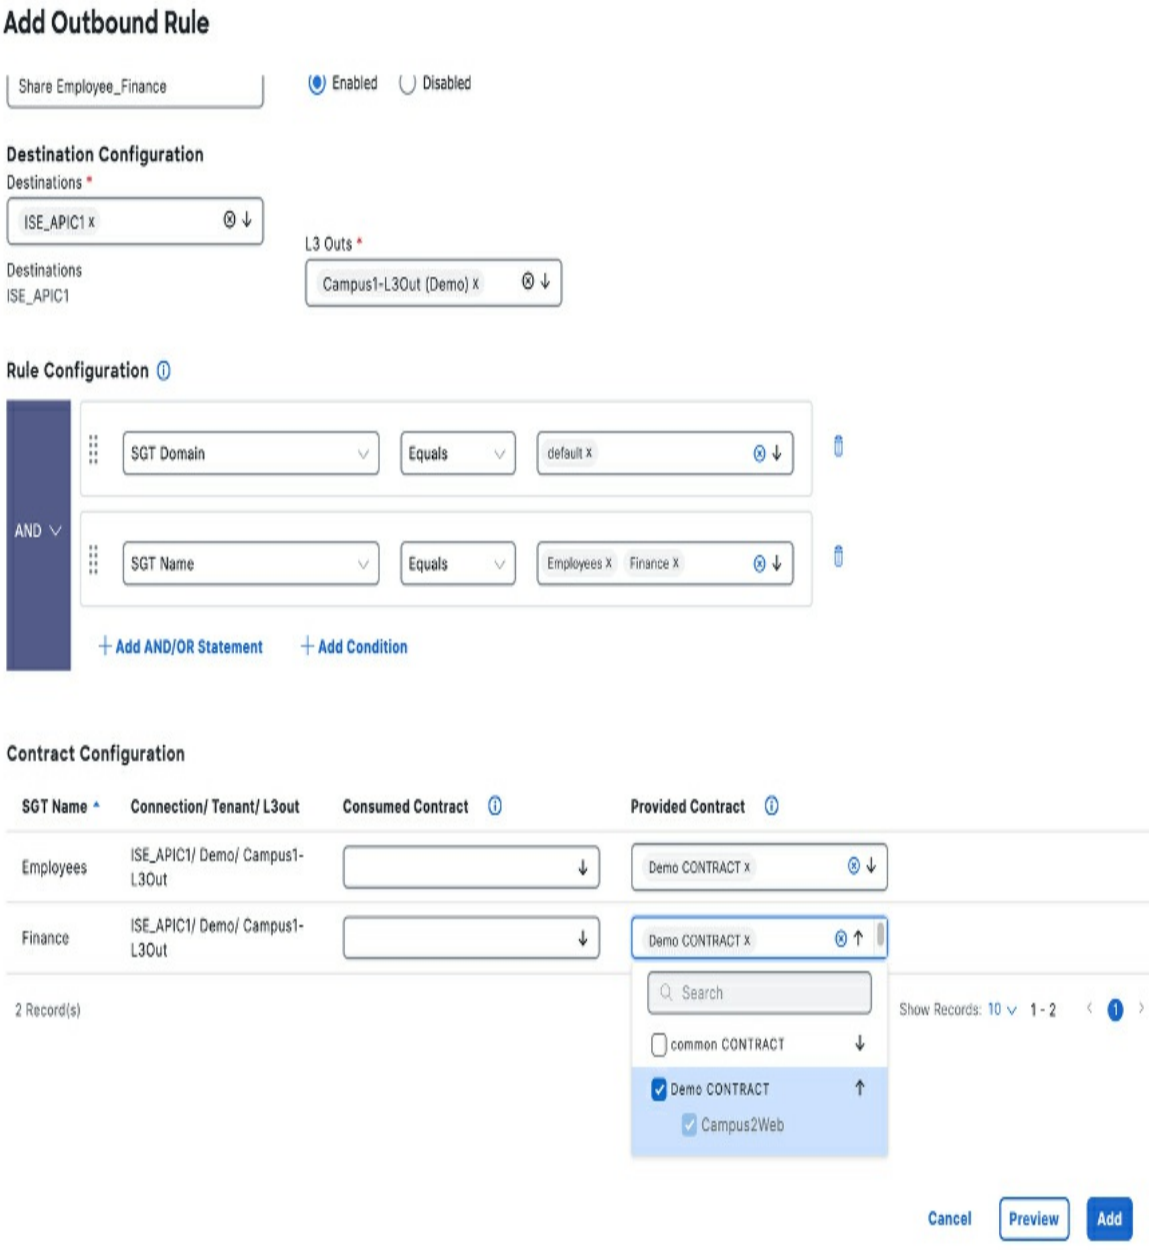
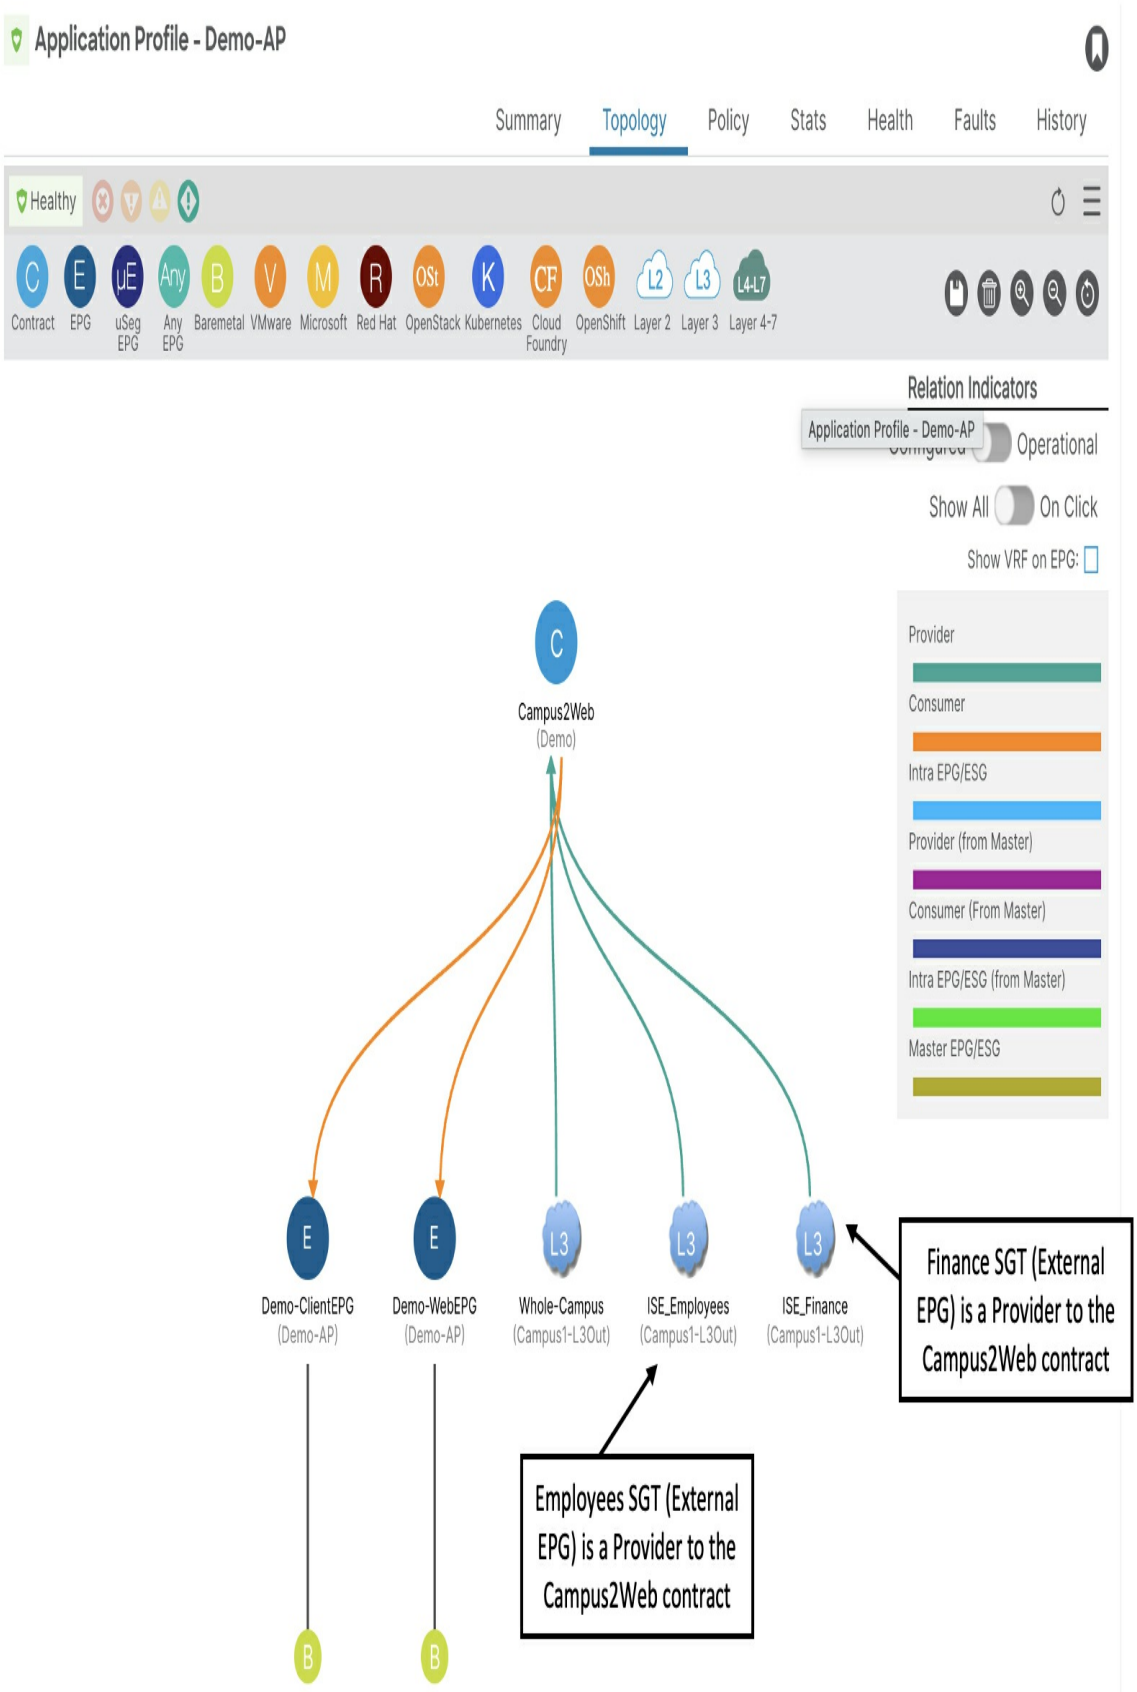


Figure 16-16 *Selecting ACI Contracts in ISE Outbound SGT Domain Rule*

After the configuration is saved and the security group shared with

Technet24

APIC, the selected **Consumed Contract** and **Provided Contract** are pre-allocated to that security group in APIC to alleviate the need to jump over to the APIC UI to allocate the contracts outside of any workflow. See how the selections made in Figure 16-16 are reflected in the topology map in Figure 16-17.


Application Profile - Demo-AP
Summary
Topology
Policy
Stats
Health
Faults
History
Healthy
Contract
EPG
uSeg EPG
Any EPG
Baremetal
VMware
Microsoft
Red Hat
OpenStack
Kubernetes
Cloud Foundry
OpenShift
Layer 2
Layer 3
Layer 4-7
Relation Indicators
Application Profile - Demo-AP
Operational
Show All
On Click
Show VRF on EPG:
Provider
Consumer
Intra EPG/ESG
Provider (from Master)
Consumer (From Master)
Intra EPG/ESG (from Master)
Master EPG/ESG
Campus2Web
(Demo)
Demo-ClientEPG
(Demo-AP)
Demo-WebEPG
(Demo-AP)
Whole-Campus
(Campus1-L3Out)
ISE_Employees
(Campus1-L3Out)
ISE_Finance
(Campus1-L3Out)
Finance SGT (External EPG) is a Provider to the Campus2Web contract
Employees SGT (External EPG) is a Provider to the Campus2Web contract

Technet24

Figure 16-17 *APIC Showing Contracts Assigned After Integration*

You can see these shared SGTs (external EPGs) in APIC by navigating to the **Integrations** menu and **Configuration > Subscribed SGTs** submenu, which displays the external EPGs and the Tenant and L3Out through which the groups were learned (see Figure 16-18).

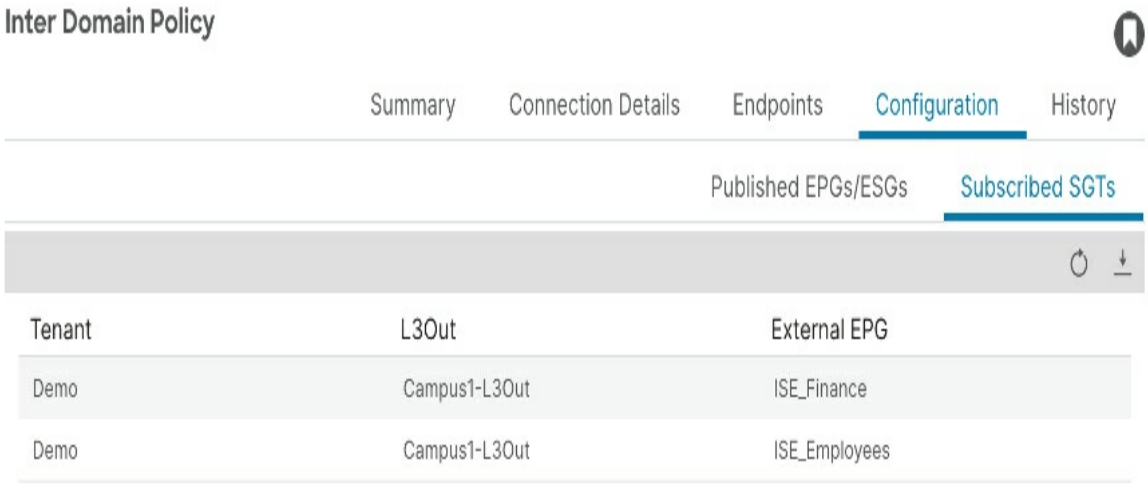


Figure 16-18 *APIC Displaying Learned SGTs*

The **APIC Endpoints > SGT Endpoints** submenu shows the SGT numbers learned, along with the IP bindings or memberships of those groups (as shown in Figure 16-19).

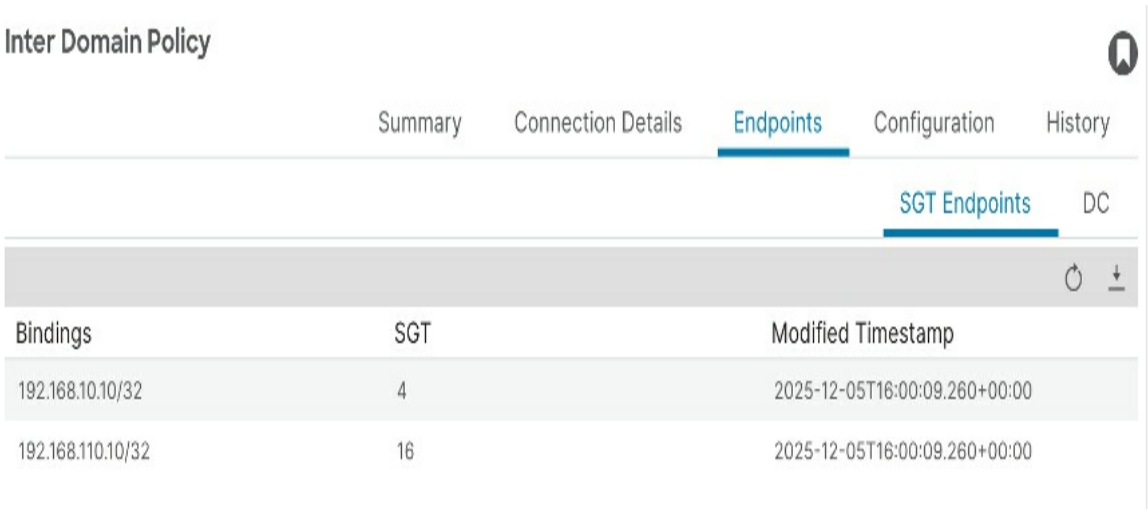


Figure 16-19 *APIC Displaying the Learned IP:SGT Bindings*

Figure 16-20 shows the **APIC Tenants** menu hierarchy and the learned groups under the **Networking > L3Outs > *contract* > External EPGs** sublevel. By highlighting one of these learned groups and selecting the **Operational > SGT Endpoints** submenu, you can display the SGT name and number along with any IP to group bindings learned.

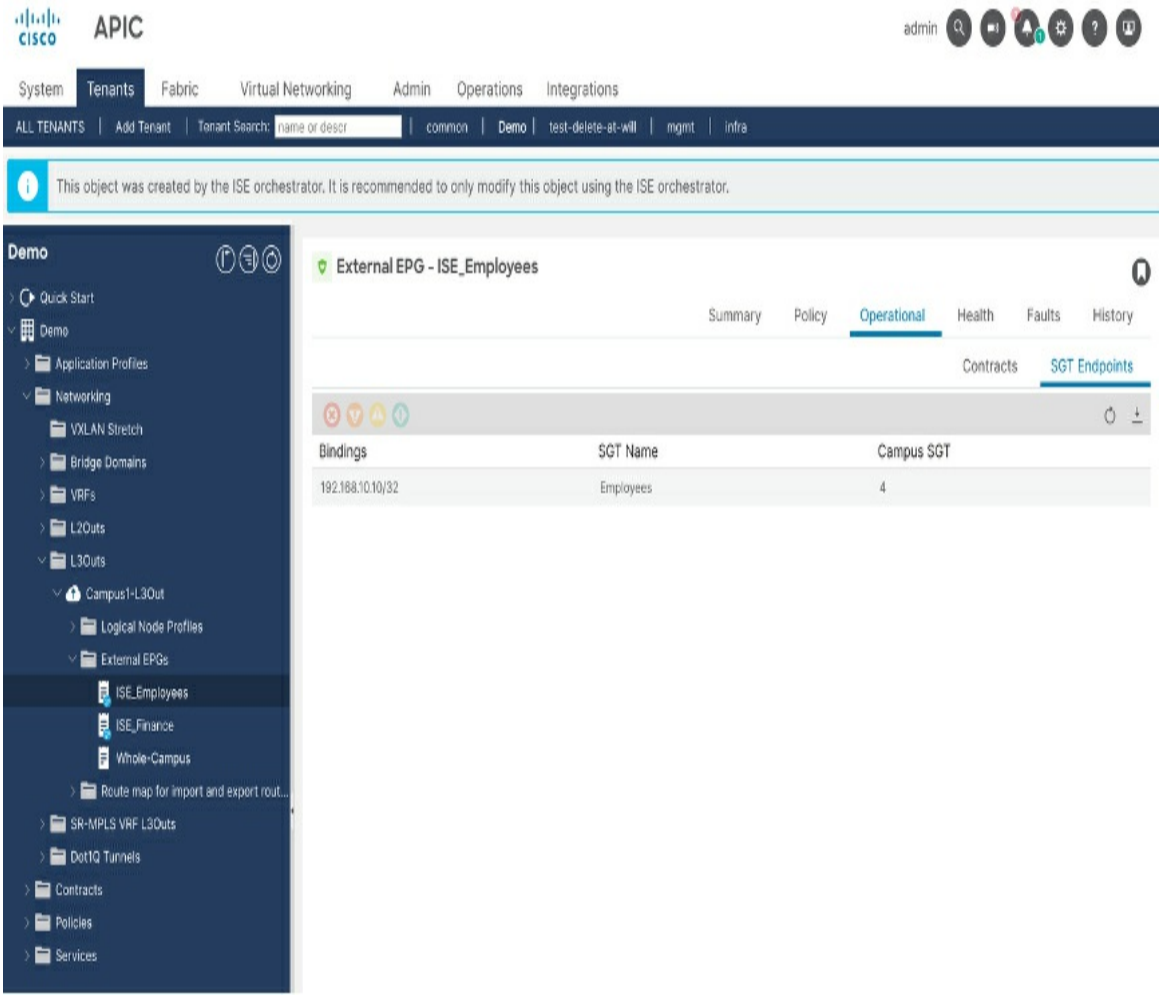


Figure 16-20 *APIC Tenants Hierarchy Showing the Learned IP:SGT Bindings*

ISE-ACI High Availability

Technet24

Cisco ISE supports a highly available architecture for all service components. The components that interact with ACI are the Policy Administration Nodes (PANs) and the pxGrid nodes. The PAN nodes are themselves highly available with a primary instance and a secondary instance. Up to four pxGrid nodes can be deployed within an ISE deployment, and each one is active. If the primary PAN node fails, then the secondary PAN node is promoted, and the integration with ACI continues to function using the secondary PAN. If the pxGrid node used for ACI integration fails, then APIC subscribes and publishes to one of the other active pxGrid nodes.

On the data center ACI side, the controller is deployed in a cluster of three APICs. On integration, an ISE administrator enters the IP address of one of the APIC controllers, but ISE learns the IP addresses of all three of the controllers. If the integrated APIC fails, then ISE switches integration over to one of the other APICs in the cluster. This makes the ISE-to-ACI integration highly available and allows for any possible single failure to occur and operation to continue.

Editing, Suspending, and Deleting Workload Connections

After ACI connections have been added on Cisco ISE, they can be edited, suspended or deleted.

To edit the connection, click the workload connection name.

Once it is clicked, a workflow similar to adding a connection is shown via a slide-in window from the right-hand side, as shown in Figure 16-21.


TrustSec Policy
Policy Sets
SX
Workload Connections
Create workload classification rules to
SGT domain rules to define the SXP do
rules to define the data shared from Ci
Search table
0 Selected
Add Connection
Workload Connection Name
ISE_APIC1
1 Record(s)
ACI Connections Details
Configuration
Name Conversion
Synced EPG/ESGs
SGT Numbering Range
ACI Connection Name*
ISE_APIC1
FQDN or IP Address*
10.5.9.136
Non-editable after creation
ACI Username*
admin
ACI Password
Login Domain
No Domain
View Details
Validate ACI certificate
LEARNED FQDN OR IP ADDRESSES
10.5.9.136
In Service
10.5.9.135
In Service
10.5.9.134
In Service

Technet24

Figure 16-21 *Editing an ACI Workload Connection*

When the slide-in window opens from the right, the **Configuration** tab is selected, showing the name and credential information. This information can be edited if required. In turn, you can select the other tabs, namely **Name Conversion**, **Synced EPG/ESGs,** and **SGT Numbering Range**, and you can also modify the details within those tabs.

Under the **Name Conversion** tab, there is a description to how any modifications are handled. See this description in Figure 16-22.

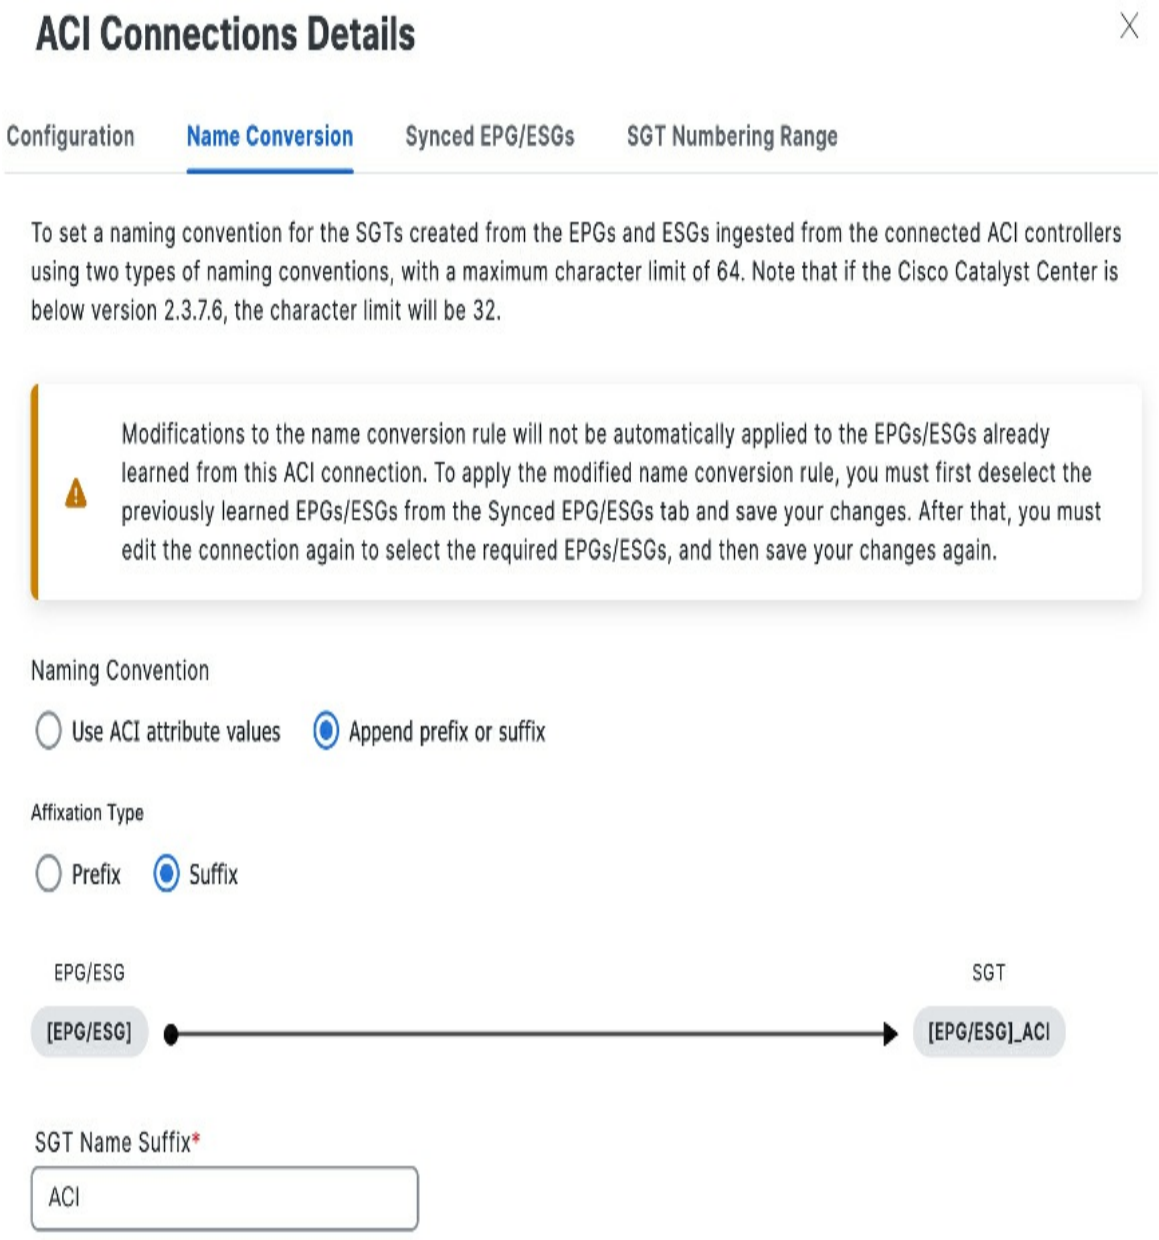


Figure 16-22 *Editing the Name Conversion Details*

Rather than deleting a workload connection on ISE, you can choose an option to suspend it instead. This option removes bindings and session data related to that connection, but the learned groups remain in place. APIC will change to "listen" only so that existing and new groups and bindings in ISE will be received by APIC.

The suspend state allows all the configuration information to remain in

Technet24

place for a quick re-establishment when required. You can find the **Suspend** option by selecting the checkbox to the left of the workload connection entry and dropping down the **More Actions** options (see Figure 16-23). Once the connection is suspended, its status will show as **Suspended**. To re-establish the connection, select the checkbox to the left of the workload connection entry, drop down the **More Actions** options, and select **Connect**.

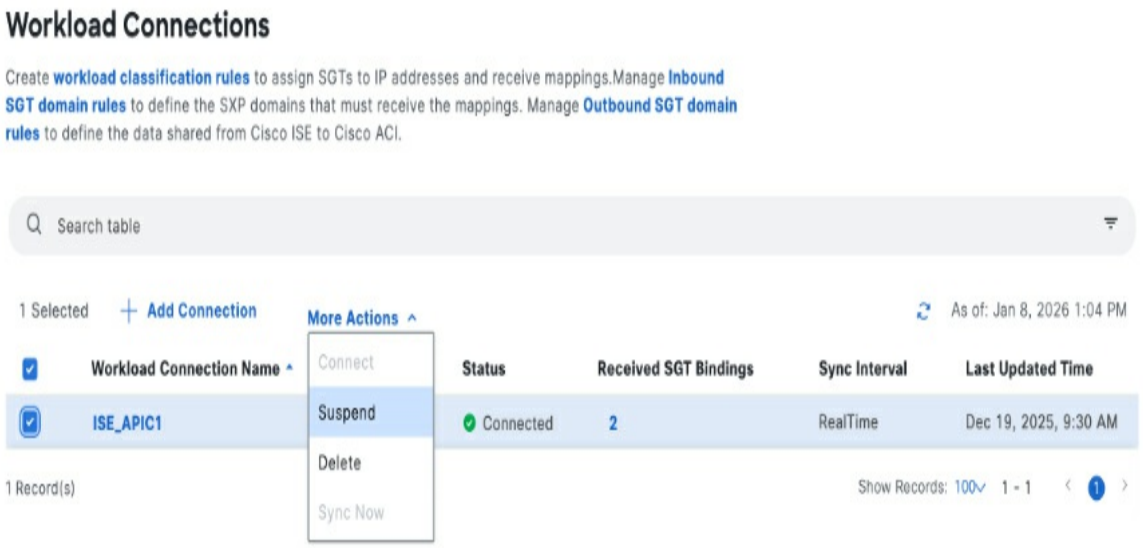


Figure 16-23 *Suspending a Workload Connection*

When a workload connection is no longer required, you can delete it. Please note that this action also removes all associated shared groups and bindings. You can locate the delete option by following the steps described and shown in Figure 16-23.

ACI Connectivity Report

To show an ACI connectivity report in ISE, navigate to **Operations > Reports > Reports > Workloads > ACI** (or **All Workloads**), as shown in Figure 16-24.

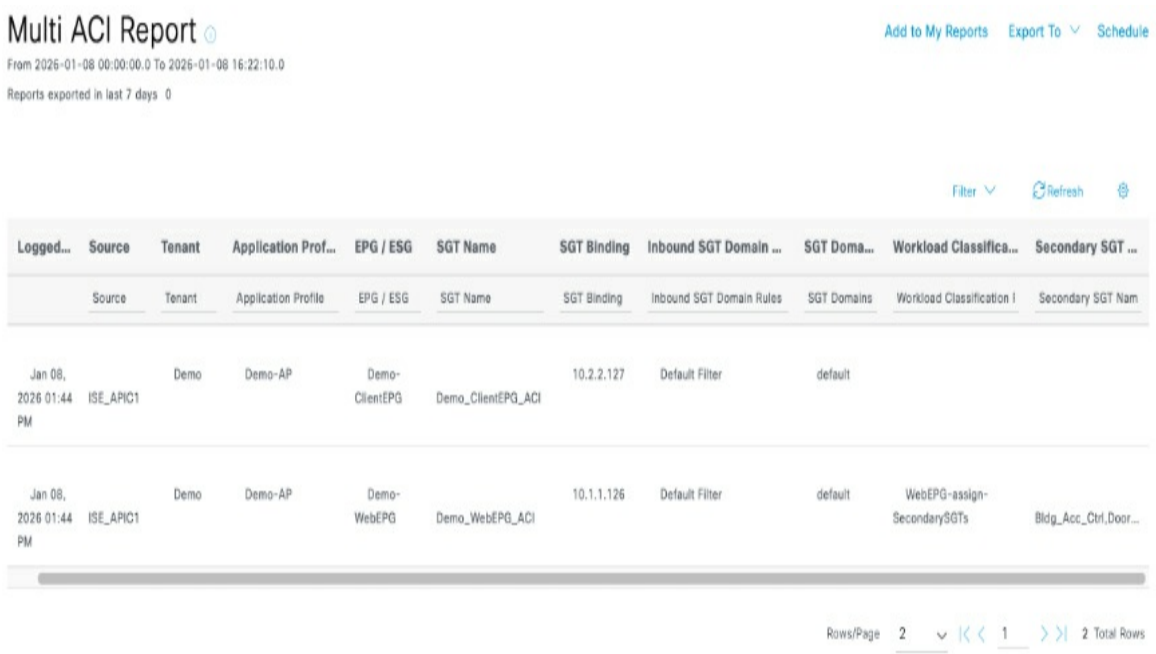


Figure 16-24 *ACI Connectivity Report*

Summary

The integration between Cisco Identity Services Engine and Cisco Application Centric Infrastructure enables consistent security policy management across campus and data center domains. By utilizing the Common Policy methodology and pxGrid, administrators can synchronize group context, such as security group tags and endpoint groups or endpoint security groups, in real time. This approach addresses the challenges of managing security for mobile users and dynamic workloads by allowing security policies to follow the endpoint regardless of its location or IP allocation.

The integration process requires a minimum software release for compatibility, and configuration of SXP and pxGrid services. ISE administrators manage group membership using SGT domains, which allow for the segregation of IP-to-group bindings. When workload connections are defined to APIC clusters, ISE can subscribe to EPGs and ESGs, which are then treated as primary SGTs. These learned

Technet24

groups become available within the TrustSec policy matrix and authorization rules, facilitating granular policy enforcement. Furthermore, ISE can share its own SGTs with APIC as external EPGs, where they can be mapped to specific ACI contracts to streamline policy deployment.

Advanced features such as primary and secondary SGTs provide a mechanism for scaling security policies, especially in multisite environments where secondary SGTs can represent site-specific requirements. High availability is maintained through redundant ISE policy administration and pxGrid nodes, alongside the inherent clustering of APIC controllers, ensuring that the integration remains resilient to individual component failures. Ongoing management is supported through the ability to edit, suspend, or delete workload connections, and connectivity status can be monitored through dedicated reports within the ISE operations dashboard.

Chapter 17. Classifying Cloud Workloads

Introduction

Deploying workloads in the cloud is very common; several providers now offer the required services. Group-based policies can be used to protect those workloads and provide least privilege for service access.

One method of protection is to enforce policies at network egress before traffic is sent to the cloud. This is accomplished by sending source IP-to-SGT bindings to a campus egress network device utilizing inline or control plane methods. Destination IP-to-SGT workload bindings are typically sent to that egress network device from Cisco Identity Services Engine using SXP or pxGrid, depending on which protocol the particular network device supports. These destination bindings can be added as part of an automation ecosystem where workload IP addresses are learned and bindings provisioned into ISE accordingly. Alternatively, an automation system could program static IP-to-SGT bindings directly into the network device if appropriate.

Since ISE version 3.4, a more dynamic method of classifying workloads has been introduced. Workloads in the cloud can be managed by cloud computing services such as Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform (GCP), but also managed by VMware vCenter.

Technet24

In ISE, a workload connection can be added to connect to these providers to retrieve attributes about the present workloads. These attributes can then be used in an ISE classification table to map the learned IP addresses of the workloads to assigned security group tags. Once allocated by ISE, these mappings, or bindings, can be shared with network devices utilizing SXP or pxGrid for enforcement purposes.

AWS EC2 Workloads and Access Rights

Let's use AWS as an example. In this case, EC2 instances could be listed as shown in Figure 17-1.

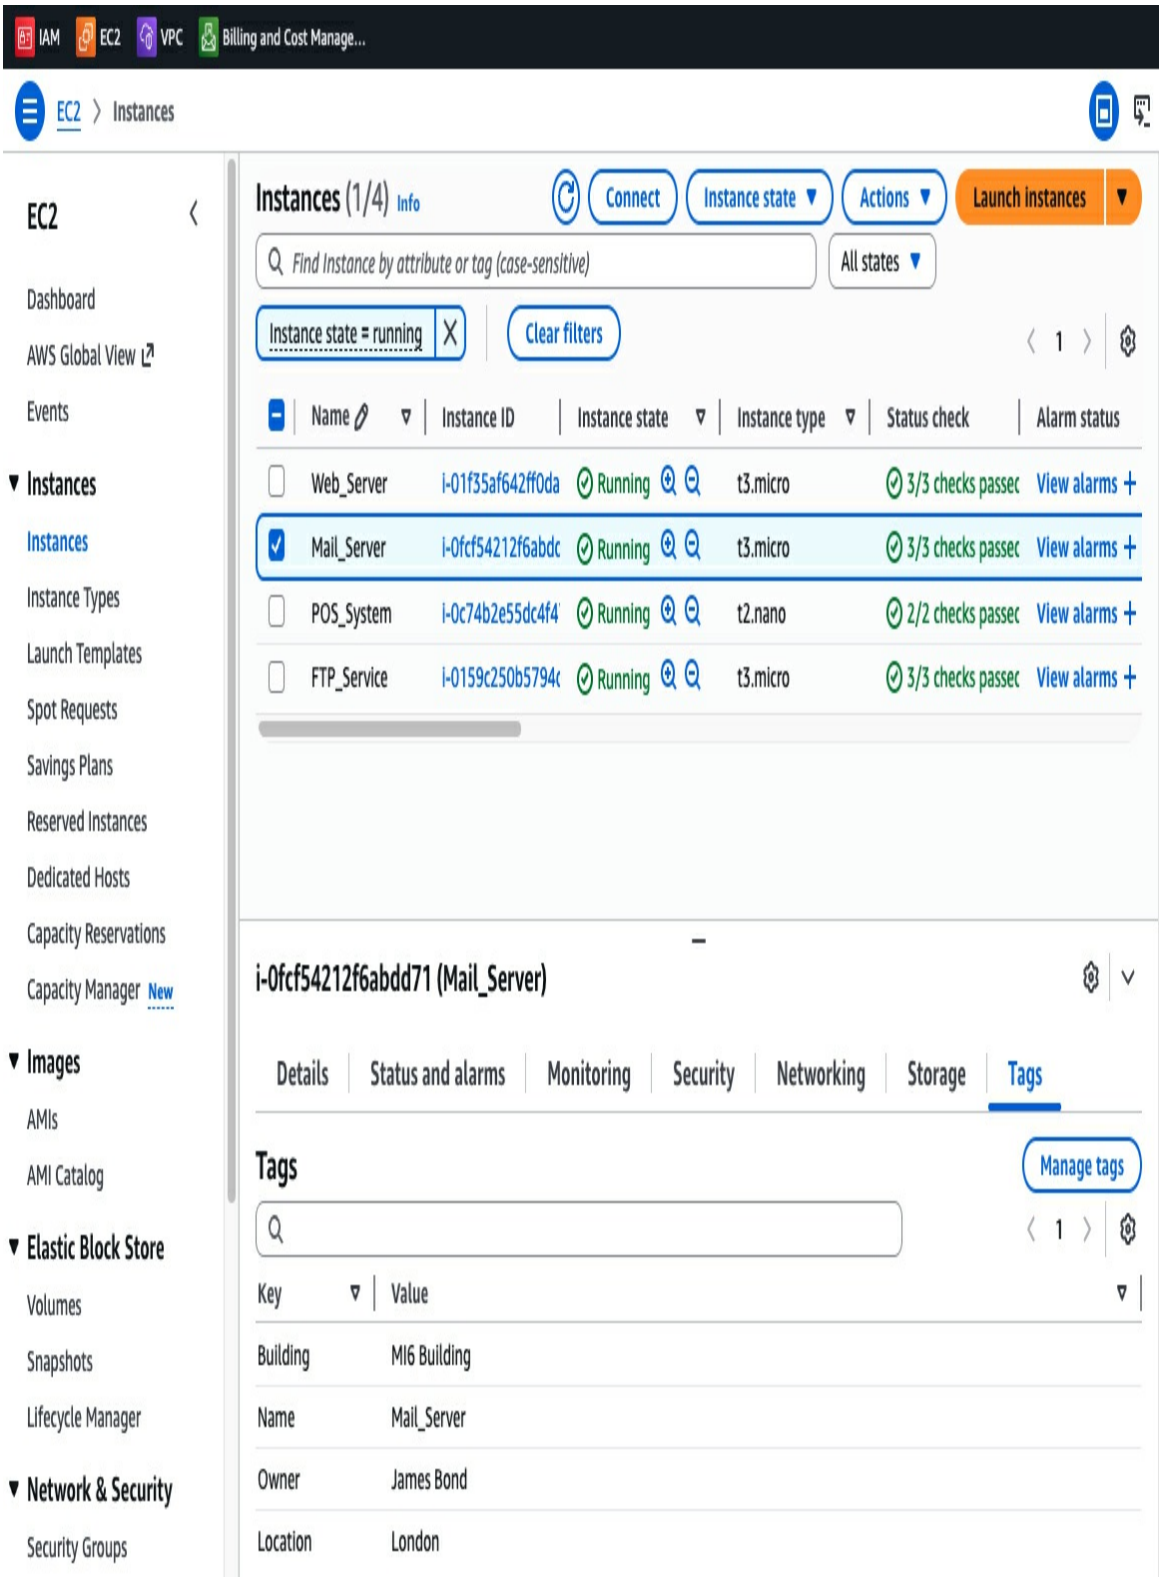


Figure 17-1 *AWS EC2 Instances with Assigned Tags*

Technet24

In Figure 17-1, the Mail_Server instance is highlighted, and the bottom part of the UI has the **Tags** submenu selected. Tags such as Name, Owner, Location, and Building have been added in relation to the server. Each of the four EC2 servers, or workloads, has similar or different tags associated with it. These tags, along with the assigned IP addresses, are what ISE uses to classify the workloads, which in turn is used for traffic enforcement.

For ISE to access the workload data, you can add an identity and access management (IAM) user in AWS with an access key for programmatic access. The same access key is entered into ISE in Figure 17-5. To navigate AWS, go to **Identity and Access Management > IAM > Access Management > IAM Users**. Under **Permissions Policies**, you can add a policy that permits ec2:DescribeTags, ec2:DescribeVpcs, and ec2:DescribeInstances as a minimum for ISE to be able to import the required attributes.

Workload Connections in ISE

To set up a workload connection in ISE, navigate to **Work Centers > TrustSec > Integrations > Workload Connections**.

Select **Add Connection**, which begins a workflow. Then select **Let's Do It!** on the subsequent screen.

The second part of the workflow allows you to select the workload platform. You can choose between ACI, AWS, Azure, vCenter, and GCP. For this exercise, select **AWS** and then **Next** to move on to the next part of the workload.

You need to enter a name for the connection, along with the AWS region and access details of an account in the AWS service.

After you select **Next**, ISE lists the attributes learned from all the

workloads in the signed-in account. These attributes are the AWS tags, as shown in Figure 17-1. As administrator, you can select any attributes, and they will be placed into an ISE dictionary to subsequently be used in ISE rules. For instance, if you need classifications for workloads in different locations, then selecting the location attribute would be appropriate; if you need classifications based on owner, then selecting the owner attribute would be appropriate, and so on. As you can see an example for Figure 17-2, you can select **Name**, **Building**, **Location**, and **Owner** as attributes that can subsequently be used in ISE rules. After you do so, select **Next**.

Technet24

Welcome | Workload Platform | Create Workload Connection | 4 Manage Attributes | 5 Summary

Manage Attributes

All the attributes learned from the workload connection is listed here. For optimum use of search lists and memory, include in the Cisco ISE dictionary only those attributes that you want to use to create classification policies, inbound and outbound SGT domain rules, and in authorization policies. At least one attribute must be added to dictionary to proceed to next step. If you have an existing connection with the same platform, any attributes already added to the Cisco ISE dictionary are automatically enabled in this list as well. You can also edit the attribute name to be displayed in the Cisco ISE dictionary.

Include All in Dictionary

Search Table

Included in Dictionary	Workload Attribute	Attribute Name in Dictionary	References
○	Access Permission	Access Permission	--
○	Last Updated	Last Updated	--
●	Name	Name	--
●	Building	Building	--
●	Location	Location	--
●	Owner	Owner	--

Exit Wizard

Back Next

Figure 17-2 *Selecting Workload Attributes*

When these steps are completed, a summary of the configuration is displayed. If it is correct, select **Create**.

The newly created workload connector is then shown in the workload connections table. See the second entry in the table in Figure 17-3.

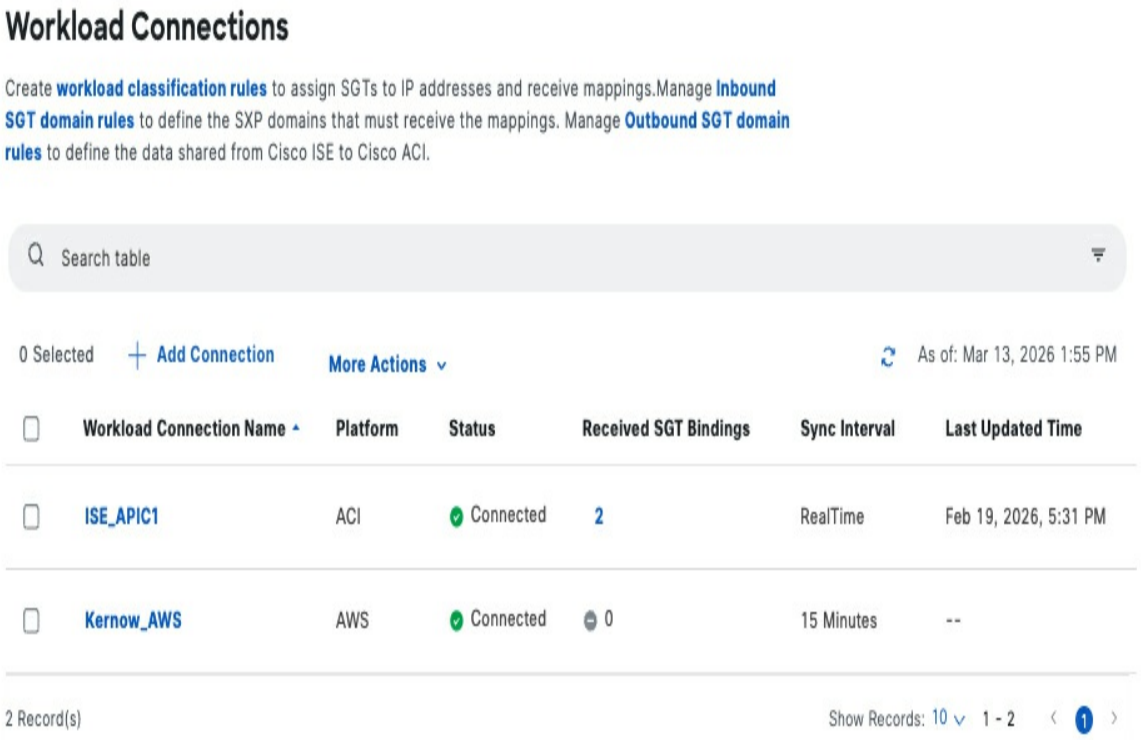


Figure 17-3 *Workload Connections*

No received SGT bindings are displayed at this stage because no classification rules have been configured.

After setting up the connection, you might need to edit the attributes to be used in classification rules. You may need to remove some attributes selected in the workflow, and you may need to add other attributes. You can accomplish this by navigating to **Work Centers > TrustSec > Integrations > Workload Connectors > Attributes Dictionary**, as shown in Figure 17-4.

Technet24

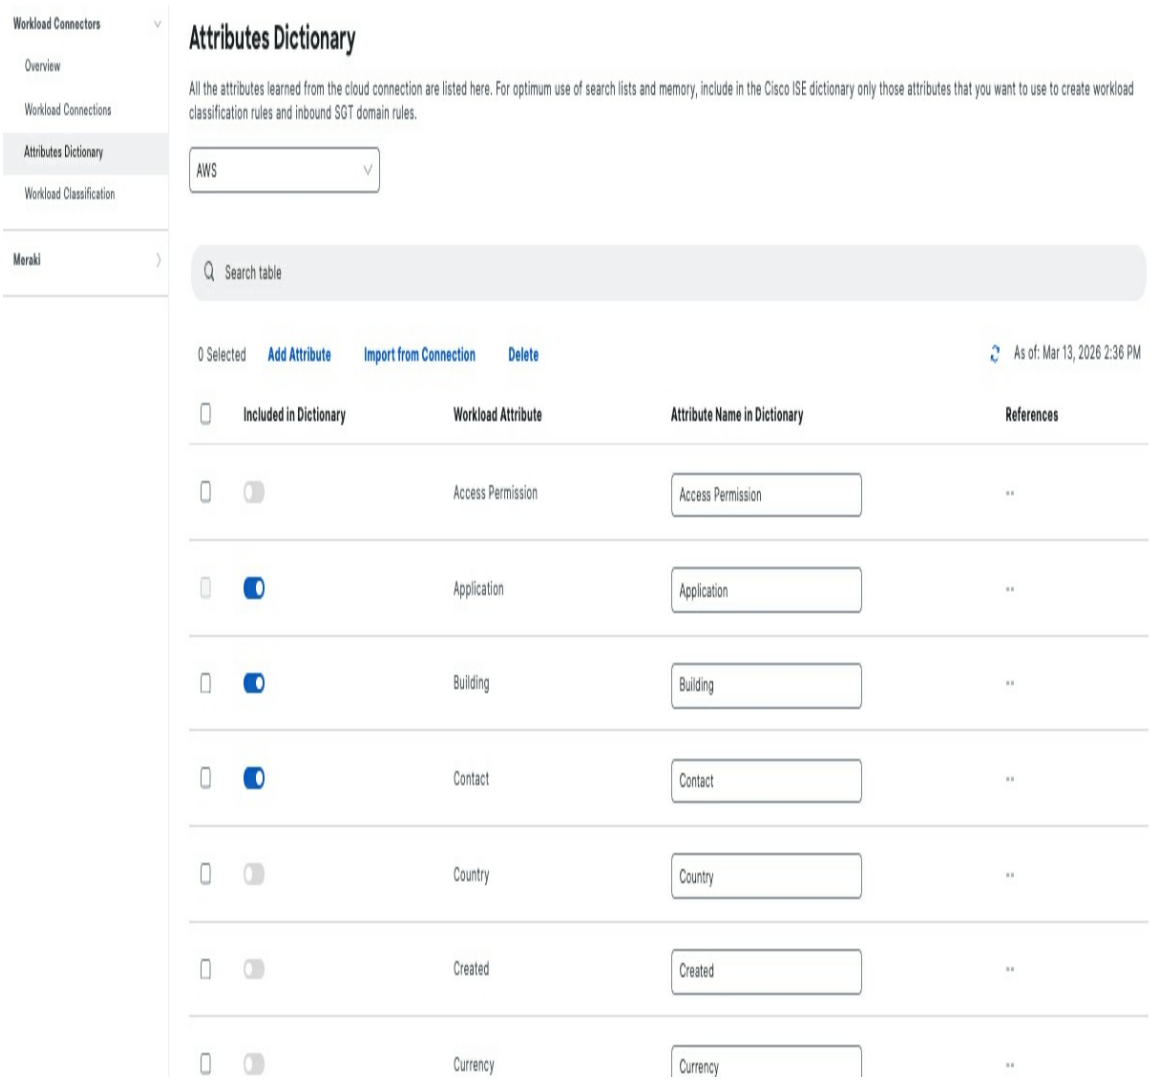


Figure 17-4 *Reselecting Workload Attributes After a Connection Is Established*

ISE Workload Classification

After the workload attributes have been selected and the connection successfully established, you can add ISE classification rules to create the required IP-to-SGT bindings. In ISE, navigate to **Work Centers > TrustSec > Integrations > Workload Connectors > Workload Classification**. Here, as administrator, you can select specific attributes, enter the matching criteria for each attribute, and select the

SGT to assign. ISE will then learn the IP addresses of the matching workloads, assign the entered SGT, and place the created IP-to-SGT bindings into the SGT binding table. These bindings are then propagated via SXP and pxGrid, depending on which method is enabled and configured. This is how workload destination IP-to-SGT bindings are sent to network devices to enforce traffic destined toward the cloud.

An existing default classification rule is disabled by default. You can utilize and edit it and add new rules. Figure 17-5 shows a new classification rule being added.

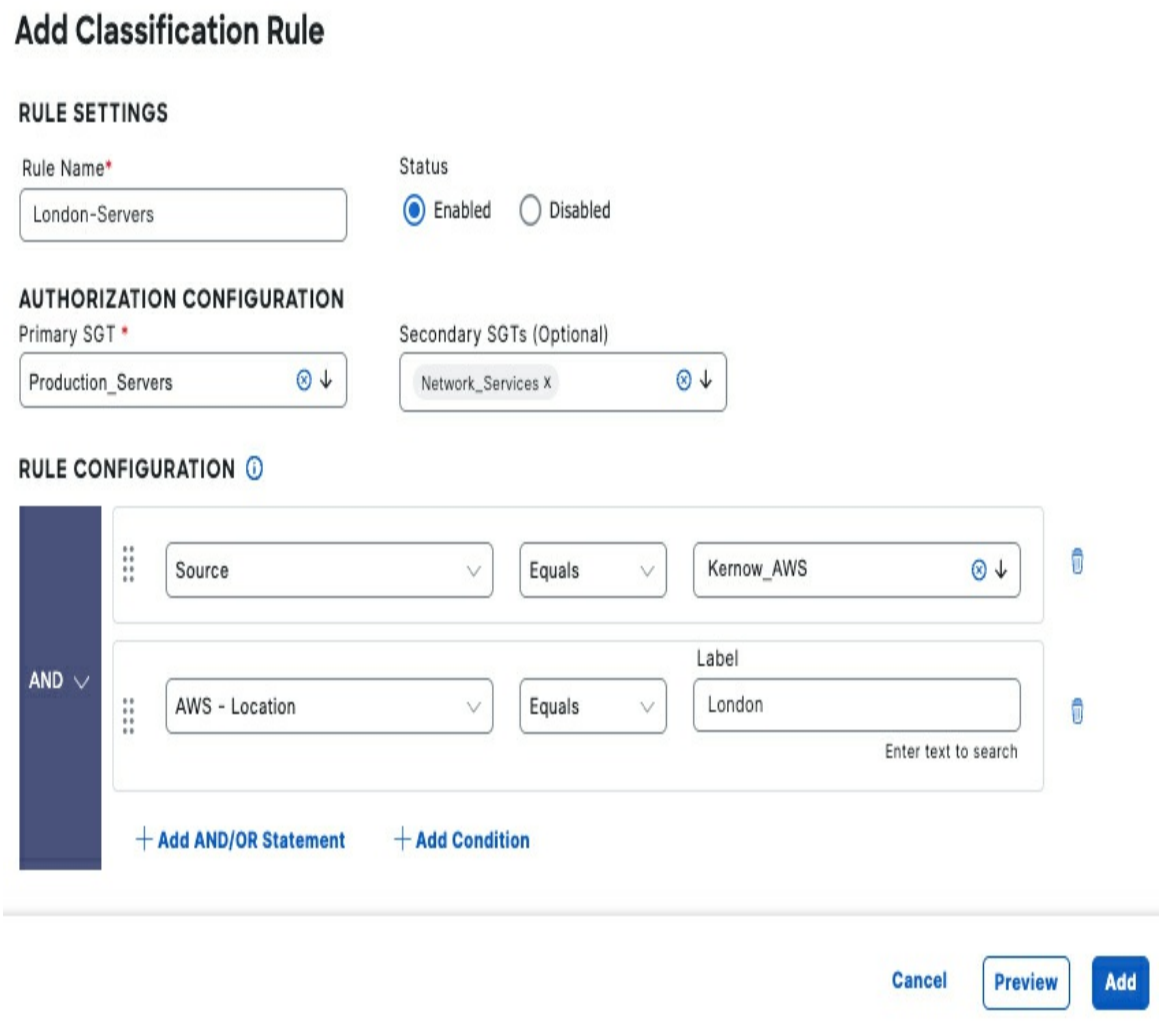


Figure 17-5 *Adding a Classification Rule*

Technet24

When adding a classification rule, you provide a name, an enabled or disabled state, followed by the SGTs that will be assigned to the IP addresses of the workloads defined by this rule. The first SGT box is labeled **Primary SGT**, and you can think of it as the SGT that has been used since the technology's inception; that is, it is the SGT that has always been assigned through static means, through ISE authorization, and always been propagated via SXP and inline tagging.

Defining the term *primary SGT* has allowed ISE to also assign and use a *secondary SGT*. This means that since version 3.4, ISE can assign an IP address of a workload to more than one SGT at the same time. At the time of writing, this secondary SGT cannot be used in ISE authorization, SXP, or inline tagging, but it is sent out via pxGrid in an ordered array named Secondary Security Groups. Future support is envisioned for Cisco firewalls and other ISE ecosystem partners to make use of this ordered array in the near future.

After the SGTs have been defined, you can enter the matching conditions. Figure 17-7 shows the **Source** attribute selected, along with an **equals** operator and matching condition of **Kernow_AWS**. The second entry has the **AWS-Location** attribute selected; this represents the Location Tag attribute learned from AWS. It is selectable here because **Location** was previously selected as an attribute to be placed in an ISE dictionary and therefore available within ISE rules (refer to Figure 17-4). Along with the **AWS-Location**, the **equals** operator is selected and the matching condition **London**. Together, these conditions match workloads learned from AWS with a location of London.

The configuration of the classification rule is now complete. Select **Add** and then **Save** to ensure the configuration is saved.

After the configuration is saved, ISE acts on the conditions within the

workload classification rules. It finds the workloads in AWS with a location of London and assigns the primary SGT of **Production_Servers** and secondary SGT of **Network_Services** to the learned IP addresses of those workloads, as shown in Figure 17-6. The classification rule entry is updated with the number of bindings created using that rule.

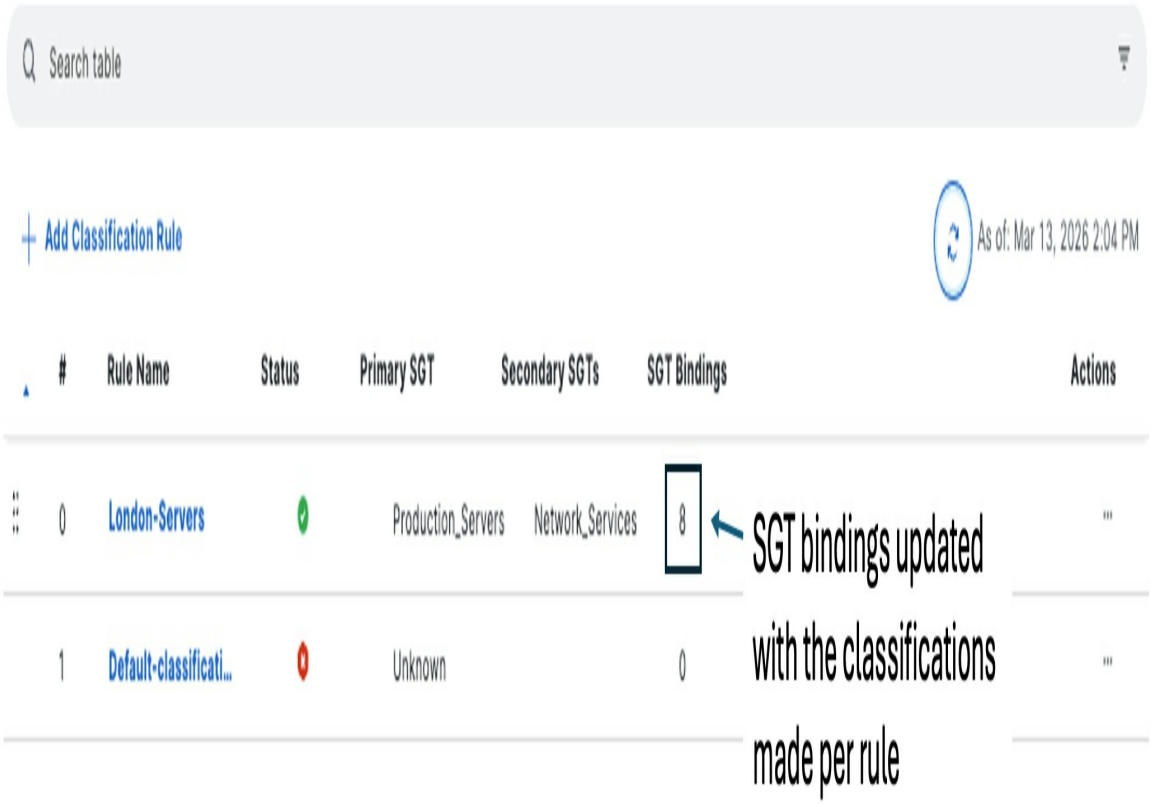


Figure 17-6 *Classification Rule Showing the SGT Bindings Created*

You can view the bindings in the **SGT Bindings** table found at **Work Centers > TrustSec > SXP** (see Figure 17-7).

Technet24

SGT Bindings

This page displays all the configured SGT bindings, giving you an easy view of assigned primary SGTs, the source of the learned data, and any workload classifications applied. You can review classification details and the specifics of defined inbound & outbound SGT domain rules by clicking the corresponding column values.

Rows/Page 25 | 1 /1 | Go | 17 Total Rows

Add SGT Domain filter | Manage SGT Domain filters | Filter

IP Address	Primary SGT	VN	Learned From	Learned ...	SGT Dom...	PSNs Involv...	Workload Cl...	Inbound SGT Do...	Source T...	Secondary Se...	Outbound SGT Do...
10.1.1.126/32	Demo_WebEPG_ACI (1001/0...	DEFAULT_VN	10.1.101.37,7.5.9.136	Session	default	Kernow-ISE3-4		Default-inbound-rule	ACI		
10.1.20.5/32	CC_TV (22/0016)		10.1.101.37	Local	Manufacturing	Kernow-ISE3-4					
10.2.2.127/32	Demo_ClientEPG_ACI (1000/...	DEFAULT_VN	10.1.101.37,7.5.9.136	Session	default	Kernow-ISE3-4		Default-inbound-rule	ACI		
10.6.5.180/32	Employees (4/0004)	DEFAULT_VN	10.1.101.37,1.1.1.8	Session	default	Kernow-ISE3-4			Radius		Share Employee_Finance
10.6.5.189/32	RandD (115/0073)	DEFAULT_VN	10.1.101.37,1.1.1.8	Session	default	Kernow-ISE3-4			Radius		
10.6.5.192/32	IP_Phones (17/0011)	DEFAULT_VN	10.1.101.37,1.1.1.8	Session	default	Kernow-ISE3-4			Radius		
10.6.5.200/32	Test_Servers (13/000D)	DEFAULT_VN	10.1.101.37,1.1.1.8	Session	default	Kernow-ISE3-4			Radius		
18.215.156.227/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
54.242.123.37/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
54.242.156.87/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
100.54.213.158/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
172.31.19.20/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
172.31.19.43/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
172.31.25.33/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
172.31.28.89/32	Production_Servers (11/000B)	DEFAULT_VN	10.1.101.37,10.1.101.1	Session	default	Kernow-ISE3-4	London-Servers	Default-inbound-rule	AWS	Network_Services	
192.168.10.10/32	Employees (4/0004)		10.1.101.37	Local	default	Kernow-ISE3-4					Share Employee_Finance
192.168.110.10/32	Finance (16/0010)		10.1.101.37	Local	default	Kernow-ISE3-4					Share Employee_Finance

Figure 17-7 *Bindings Created for AWS Workloads in the SGT Bindings Table*

These bindings are propagated by ISE using SXP and/or pxGrid to network devices where traffic can be enforced from/to AWS cloud workloads.

ISE Inbound SGT Domain Rules

By default, the assigned IP-to-SGT bindings classified by cloud workloads are placed into the default ISE SGT domain. SGT domains are explained in Chapter 7, "Introduction to TrustSec," but basically are spaces in ISE to separate bindings and help in propagating the correct subset of bindings to appropriate destinations. Even though classifying workloads places the bindings into the default SGT domain by default, you can add rules to the Inbound SGT Domain rules, which select the appropriate SGT domain when workloads are classified.

Note that the default Inbound SGT Domain rule can be edited. However, there are no condition rules, which means it applies to all SGT bindings; plus, it cannot be disabled. The only attribute that can be changed in the default rule is the SGT domain; it can be changed from default to any user-defined domains, and multiples can be selected.

When adding an Inbound SGT Domain rule, you need to enter a name for the rule plus the status of the rule when it is created, whether enabled or disabled. You then enter the destination SGT domain for the workload classifications assigned. If the desired SGT domain is not pre-created, then you can create one from the **Destinations** drop-down box. Finally, you can add the conditions or rule configurations. The example in Figure 17-8 uses the source of data from the **Kernow_AWS** workload connection and the workload attribute **AWS-**

Technet24

Name containing the term **POS**.

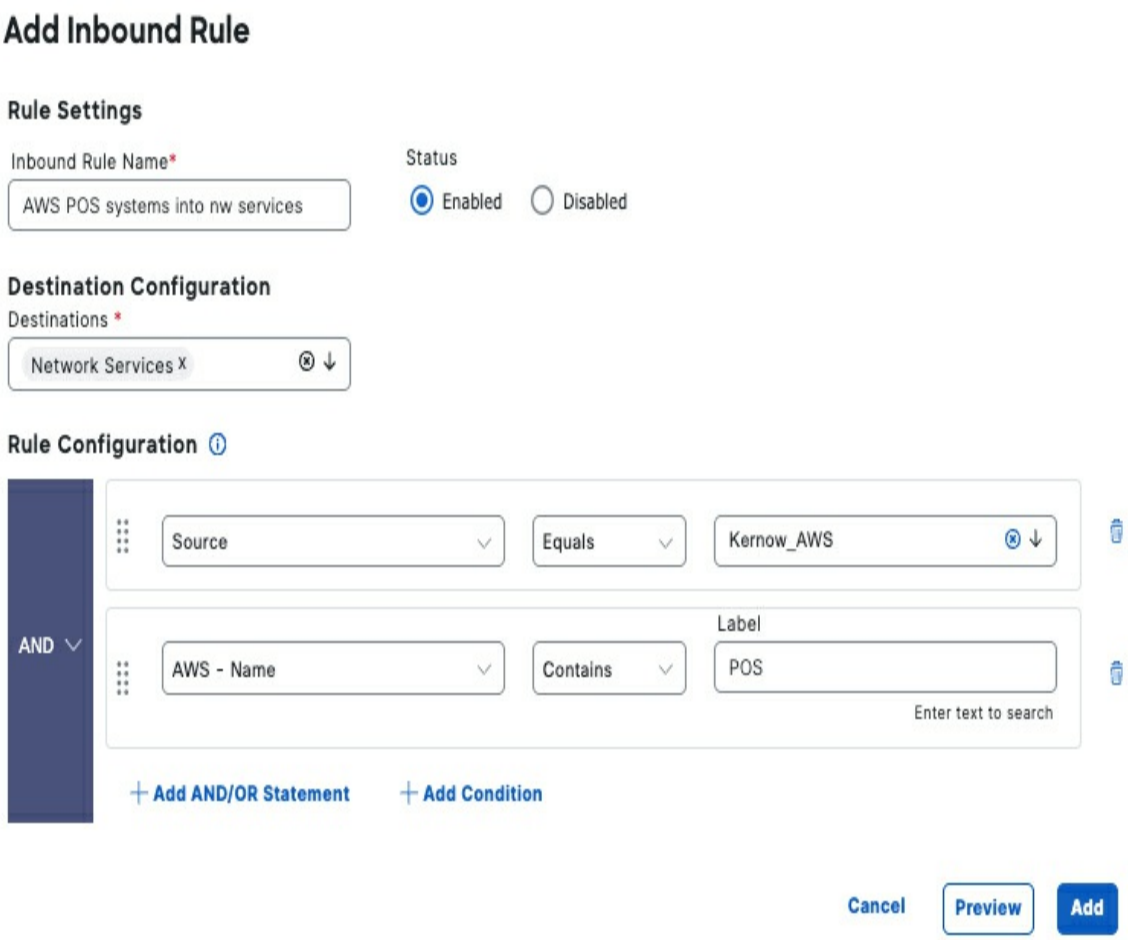


Figure 17-8 *Adding an ISE Inbound SGT Domain Rule*

After the new rule is added, it takes a few seconds for the number of matching SGT bindings to be displayed on the **Inbound SGT Domain Rules** page (see Figure 17-9). Bear in mind that the default rule is also matched because it cannot be disabled; it is a multi-match operation.

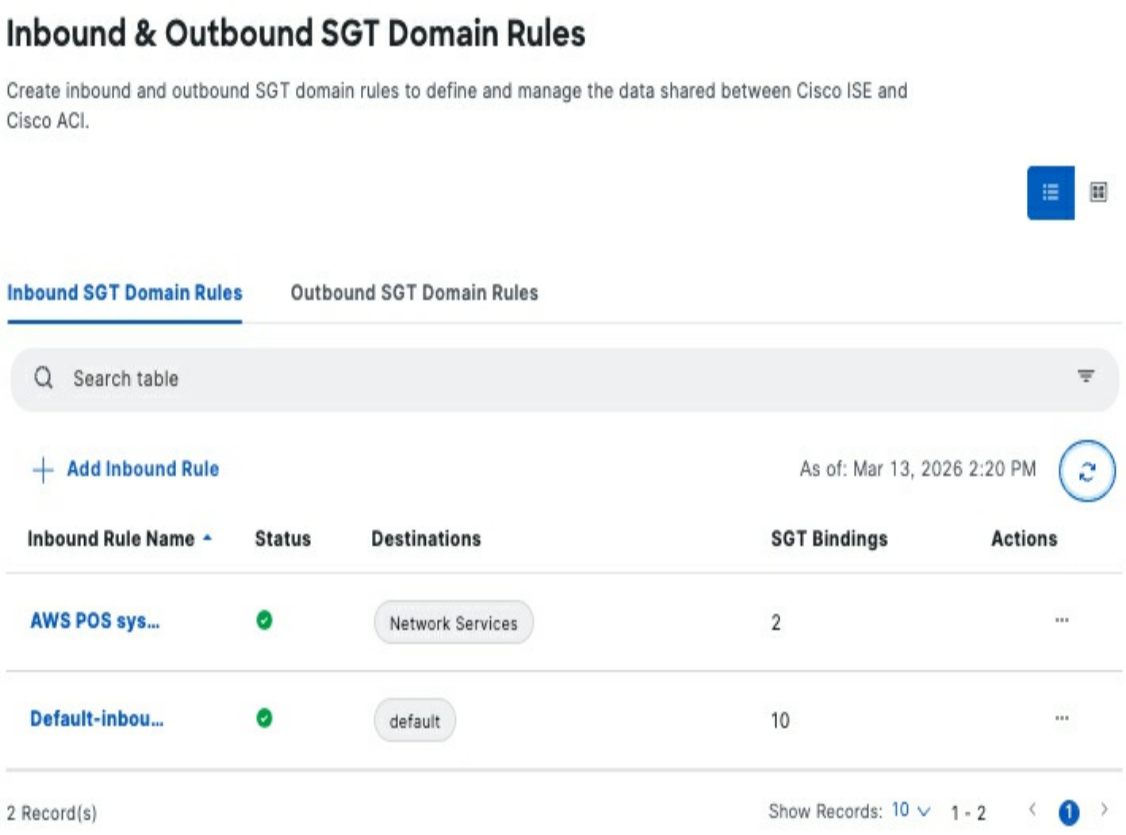


Figure 17-9 *Number of SGT Bindings Matched by Inbound SGT Domain Rules*

You can view those two bindings assigned by the newly added rule in the **SGT Bindings** screen by navigating to **Work Centers > TrustSec > SXP > SGT Bindings**, as shown in Figure 17-10.

Technet24

SGT Bindings

This page displays all the configured SGT bindings, giving you an easy view of assigned primary SGTs, the source of the learned data, and any workload classifications applied. You can review classification details and the specifics of defined inbound & outbound SGT domain rules by clicking the corresponding column values.

Rows/Page 2 | 1 /1 | Go | 2 Total Rows

Add SGT Domain filter | Manage SGT Domain filters | Filter

IP Address	Primary SGT	Learned From	SGT Doma...	Workload Cl...	Inbound SGT Do...	Source T...	Secondary Sec...
X			Network Ser. x				
100.54.213.158/32	Production_Servers ...	10.1.101.37,10.1.101.1	Network Service	London-Servers	Default-inbound-rule,AWS POS systems into nw services	AWS	Network_Services
172.31.28.89/32	Production_Servers ...	10.1.101.37,10.1.101.1	Network Service	London-Servers	Default-inbound-rule,AWS POS systems into nw services	AWS	Network_Services

Figure 17-10 *SGT Bindings Filtered by SGT Domain*

Filtering the SGT bindings by the SGT domain name **Network Services** finds the two bindings imported into that domain by the added ISE Inbound SGT rule created in Figure 17-10. When in that SGT domain, ISE propagates those bindings via pxGrid and any SXP connection associated with that domain. When the receiving network devices or appliances receive those bindings, they can be used for visibility or enforcement purposes.

ISE Reporting on Workload Connections

To report on the actions of the Workload Connections, navigate to **Operations > Reports > Reports > Workloads**.

Summary

This chapter outlines the methodology for securing cloud-based workloads using Cisco Identity Services Engine. As organizations increasingly deploy services across hybrid cloud environments (AWS, Azure, GCP, and vCenter), maintaining consistent security policies is critical. From version 3.4 onwards, Cisco ISE provides a dynamic mechanism to classify these workloads and enforce group-based policies using security group tags.

The integration process begins with establishing cloud service connections through which ISE retrieves workload metadata or tags, which are then mapped to SGTs using classification rules. Within the classification framework, administrators can define primary and secondary SGTs. Primary SGTs are used as SGTs have always been used, whereas secondary SGTs allow for multiple SGT assignments per IP. These secondary tags are currently sent via pxGrid as an ordered array, intended for future ecosystem integrations.

Technet24

To maintain binding separation, administrators can utilize Inbound SGT Domain rules to categorize these bindings into specific domains. This allows binding filtering to occur, so network devices are sent only the bindings they require.

By leveraging cloud-native tags and integrating them into the ISE TrustSec framework, organizations can achieve a "least-privilege" access model for cloud workloads. This automated approach reduces the manual overhead of managing static IP-to-SGT bindings and ensures that security policies remain consistent as workloads scale or migrate across cloud and on-premises environments.

Chapter 18. ISE APIs and Advanced Automation Use Cases

Introduction

An application programming interface (API) defines a specific interface or connection point within an application and exposes selected functionality or data to another application in a consistently structured manner. In the context of Cisco Identity Services Engine and other security platforms, APIs enable external tools and applications to interact directly with the security platform without requiring user interaction in the graphical user interface or command-line interface.

A highly desirable feature of APIs in modern systems is their ability to establish a layer of abstraction between the defined outcome or response of the interface and the back-end implementation in the application code. When you update or modify back-end systems or functions, the API response remains consistent and unaffected. This architecture allows you to make improvements in code with the assurance that, as long as the API response remains unchanged, you do not risk breaking external dependencies within the system.

The large-scale implementation of APIs enables system interaction without direct human involvement. Before APIs became commonplace, external interactions with systems and applications occurred exclusively through human interface methods, such as the CLI or GUI.

Technet24

APIs now allow you to automate system interactions using scripting and integrate functionalities across systems to combine separate entities into a single logical solution capable of addressing modern challenges.

ISE provides a variety of APIs across several categories, including OpenAPI, ERS, and MnT APIs, which are discussed in detail in this chapter. All ISE APIs use HTTPS for transport and support multiple execution methods. Common approaches include utilizing the built-in Swagger UI for OpenAPI, making direct browser calls, employing API clients such as Bruno or Postman, using the Requests library in Python, or issuing cURL commands from the terminal.

API Gateway

Before ISE 3.1, you needed to manually route API calls to the appropriate API service through the client used for the request. For example, you directed Monitoring and Troubleshooting (MnT) API calls to an MnT node in the deployment on TCP port 9443. You sent External Restful Services (ERS) calls for write operations to the Primary Admin Node on port 9060, while ERS read operations could be sent to any node on port 9060. These different locations and ports often caused confusion and increased overhead when making API calls. To address this situation, ISE 3.1 and later releases introduced the API gateway solution.

The API gateway operates on TCP port 443 and serves as a single entry and exit point for service API calls, including OpenAPI, ERS, and MnT APIs. The gateway uses HTTPS as the transport method, which encrypts communications between the client and ISE. For these requests, ISE presents the system certificate designated for "Admin."

In a standalone ISE deployment, the API gateway remains enabled by

default and cannot be disabled. In a distributed deployment, the API gateway is enabled by default on the Primary Admin Node unless another node in the deployment has the feature already enabled, in which case it is not enabled by default on the Primary Admin Node. You can toggle the API gateway feature on nodes within the deployment by using the ISE GUI under **Administration > System > Settings > API Settings > API Gateway Settings,** as shown in Figure 18-1. For high availability of API services, you should enable the API gateway service on at least two nodes.

Technet24

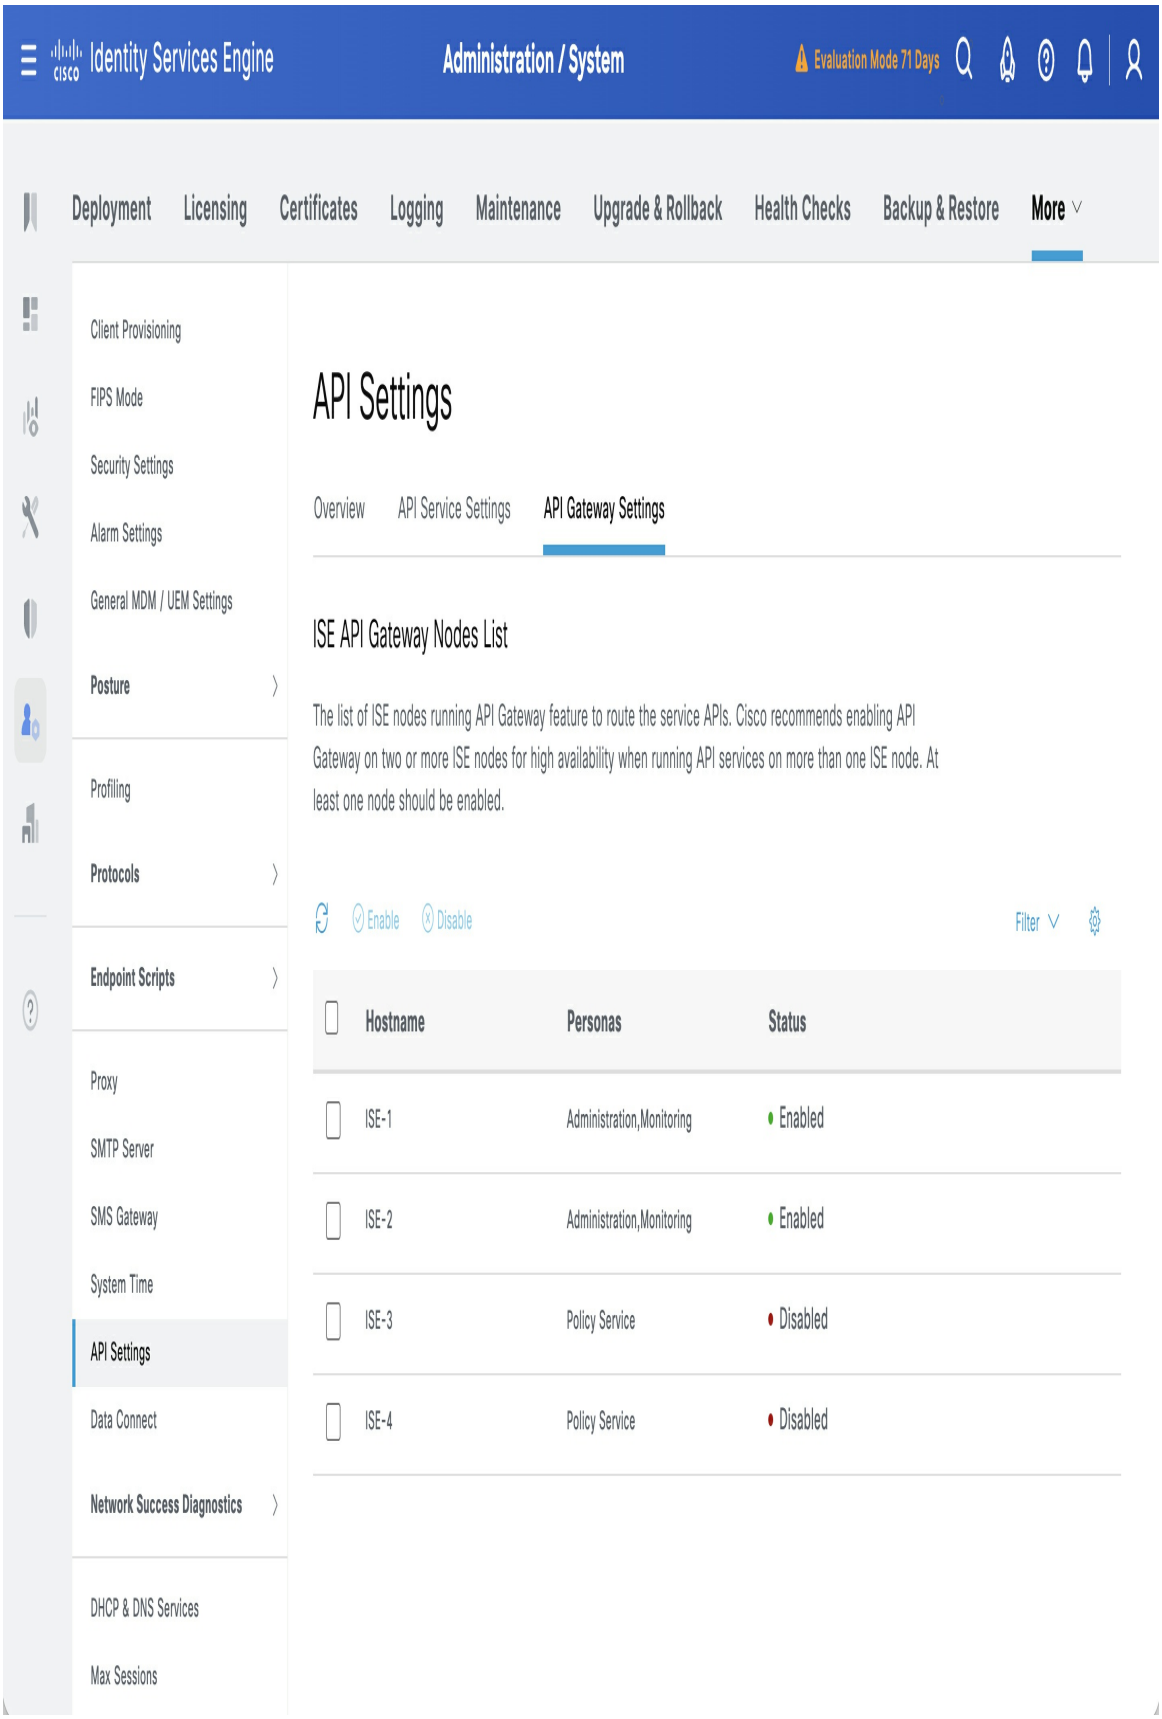

Identity Services Engine
Administration / System
Evaluation Mode 71 Days
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Client Provisioning
FIPS Mode
Security Settings
Alarm Settings
General MDM / UEM Settings
Posture
Profiling
Protocols
Endpoint Scripts
Proxy
SMTP Server
SMS Gateway
System Time
API Settings
Data Connect
Network Success Diagnostics
DHCP & DNS Services
Max Sessions
API Settings
Overview
API Service Settings
API Gateway Settings
ISE API Gateway Nodes List
The list of ISE nodes running API Gateway feature to route the service APIs. Cisco recommends enabling API Gateway on two or more ISE nodes for high availability when running API services on more than one ISE node. At least one node should be enabled.
Enable
Disable
Filter
Hostname
Personas
Status
ISE-1
Administration,Monitoring
Enabled
ISE-2
Administration,Monitoring
Enabled
ISE-3
Policy Service
Disabled
ISE-4
Policy Service
Disabled

Figure 18-1 *API Gateway Settings*

The API gateway appears as an ISE process in the output of the **show application status ise** command in the ISE CLI. The API gateway process should remain in the running state on all ISE nodes, regardless of whether the service is enabled from the ISE GUI. The application server process depends on the API gateway process being active, so if the API gateway service fails to start, you may not be able to access the GUI and may experiences a loss of functionality on the node.

OpenAPI

OpenAPI is an industry standard specification for defining and documenting REST APIs. This specification is commonly used with the Swagger open-source toolkit, which provides tools designed to aid the OpenAPI development lifecycle of the design, building, documentation, and consumption of the APIs. This allows for more efficient development and testing of APIs, but perhaps the key benefit is the automatically generated, interactive documentation available with the Swagger UI.

Before ISE 3.4, you needed to enable OpenAPI manually from **Administration > System > Settings > API Settings > API Service Settings**. On this page, you could enable OpenAPI for the Primary Administration Node as well as for all other nodes. In ISE 3.4 and later releases, OpenAPI remains enabled by default, and the option to disable it was removed from this page.

The Swagger UI is available in ISE at https:///api/swagger-ui/index.html and requires an administrator account for access. You can also reach the Swagger UI through a hyperlink under **Administration > System > Settings > API Settings > Overview**, as shown in Figure 18-2.

Technet24

Identity Services Engine | Administration / System

Deployment | Licensing | Certificates | Logging | Maintenance | Upgrade & Rollback | Health Checks | Backup & Restore | More

- Client Provisioning
- FIPS Mode
- Security Settings
- Alarm Settings
- General MDM / UEM Settings
- Posture
- Profiling
- Protocols
- Endpoint Scripts
- Proxy
- SMTP Server
- SMS Gateway
- System Time
- API Settings
- Data Connect
- Network Success Diagnostics
- DHCP & DNS Services
- Max Sessions
- Light Data Distribution

API Settings

Overview | API Service Settings | API Gateway Settings

API Services Overview

You can manage Cisco ISE nodes through two sets of API formats–External Restful Services (ERS) and OpenAPI.
Starting Cisco ISE Release 3.1, new APIs are available in the OpenAPI format.
The ERS and OpenAPI services are HTTPS-only REST APIs that operate over port 443.
Currently, ERS APIs also operate over port 9060. However, port 9060 will not be supported for ERS APIs in later Cisco ISE releases. We recommend that you only use port 443 for ERS APIs.
OpenAPI service is enabled by default and cannot be disabled, ERS API service is disabled by default. Enable ERS API service by clicking the toggle button in the API Service Settings tab.
To use either API service, you must have the ERS-Admin or ERS-Operator user group assignment.

For more information on ISE ERS API, please visit:
https://ise-1.iselab.net/ers/sdk

For openapi documention for ERS, click below:
ERS_V1

For more information on ISE Open API, please visit:
https://ise-1.iselab.net/api/swagger-ui/index.html

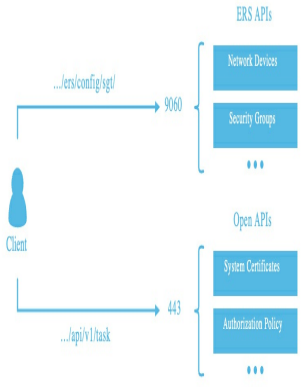


Figure 18-2 *API Swagger Documentation URL*

Within the Swagger UI, you will find numerous definitions that contain API endpoints related to various features and resources in ISE. You can select these definitions using the drop-down menu located in the top-right corner of the page, as shown in Figure 18-3. Within each definition, various API endpoints are listed, allowing you to perform different functions on the selected definition.

Technet24

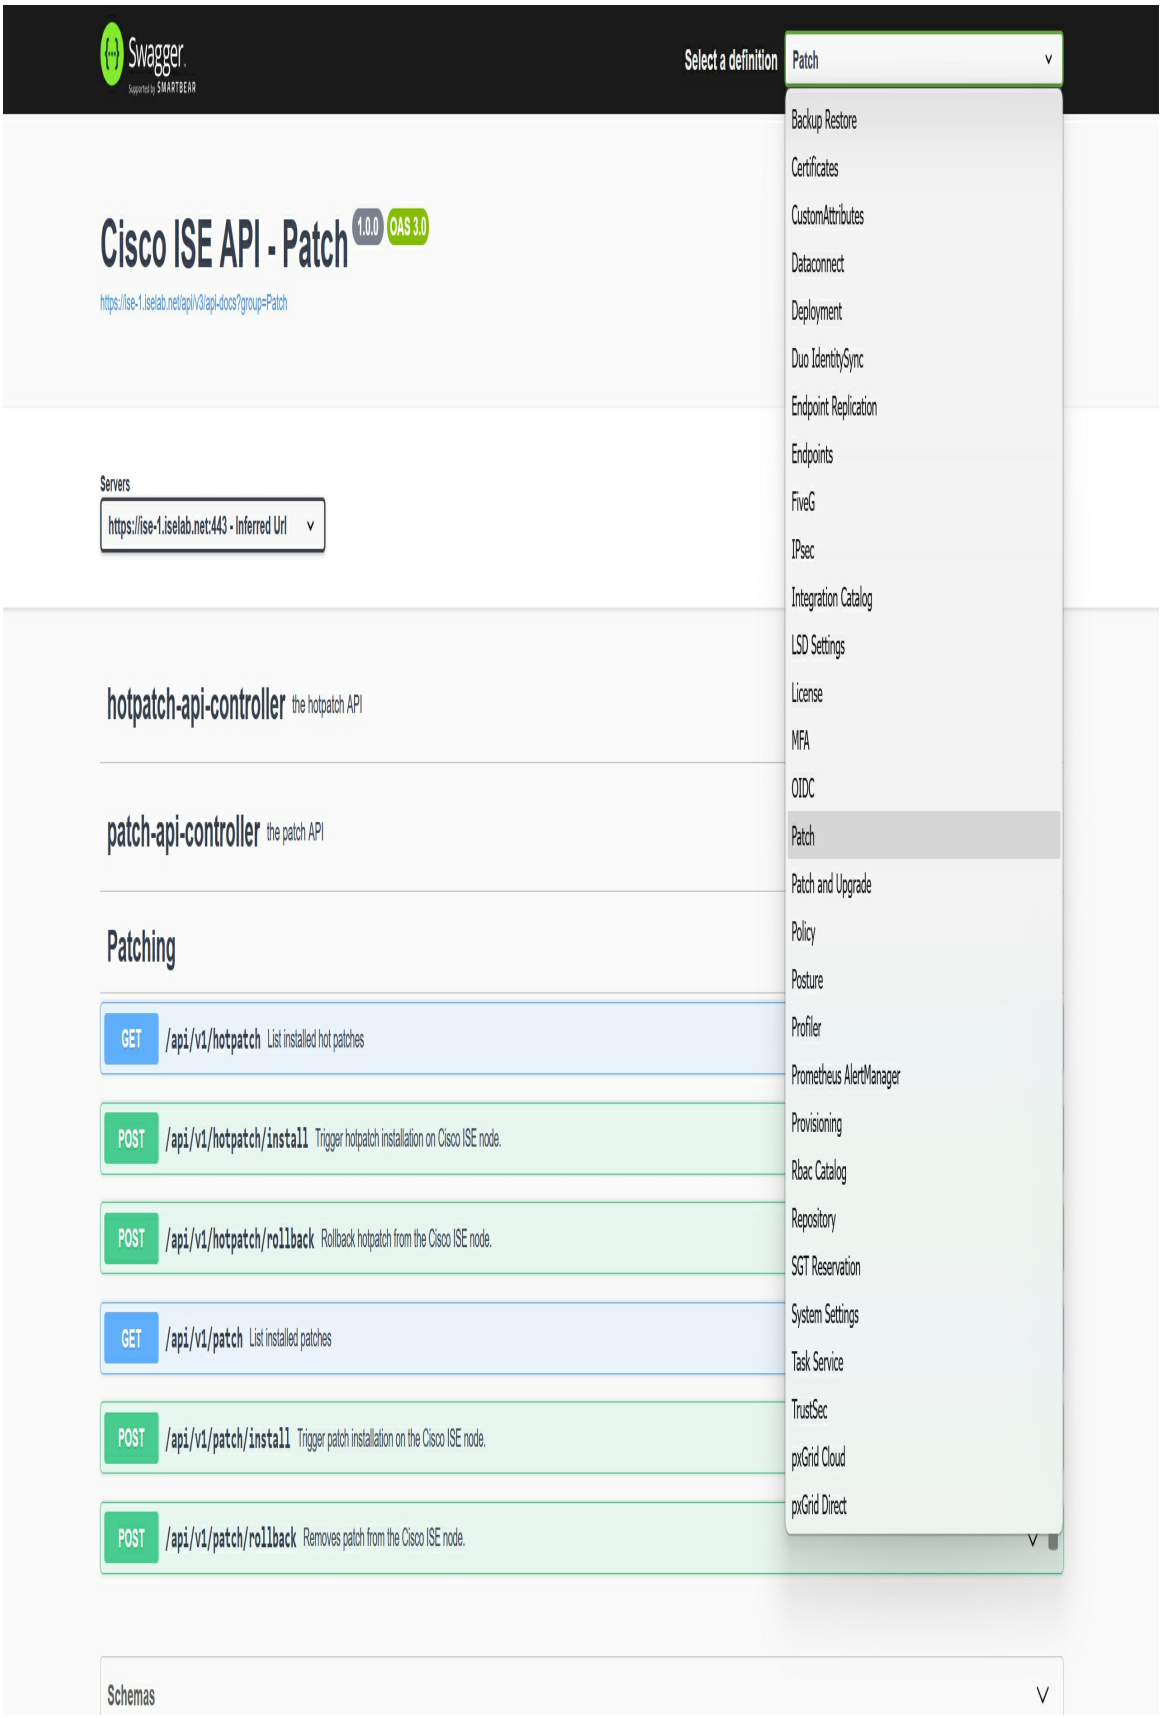

Swagger
Supported by SMARTBEAR
Select a definition
Patch
Backup Restore
Certificates
CustomAttributes
Dataconnect
Deployment
Duo IdentitySync
Endpoint Replication
Endpoints
FiveG
IPsec
Integration Catalog
LSD Settings
License
MFA
OIDC
Patch
Patch and Upgrade
Policy
Posture
Profiler
Prometheus AlertManager
Provisioning
Rbac Catalog
Repository
SGT Reservation
System Settings
Task Service
TrustSec
pxGrid Cloud
pxGrid Direct
Cisco ISE API - Patch
1.0.0
OAS 3.0
https://ise-1.iselab.net/api/v3/api-docs?group=Patch
Servers
https://ise-1.iselab.net:443 - Inferred Url
hotpatch-api-controller the hotpatch API
patch-api-controller the patch API
Patching
GET /api/v1/hotpatch List installed hot patches
POST /api/v1/hotpatch/install Trigger hotpatch installation on Cisco ISE node.
POST /api/v1/hotpatch/rollback Rollback hotpatch from the Cisco ISE node.
GET /api/v1/patch List installed patches
POST /api/v1/patch/install Trigger patch installation on the Cisco ISE node.
POST /api/v1/patch/rollback Removes patch from the Cisco ISE node.
Schemas

Figure 18-3 *Cisco ISE Swagger UI*

Each API endpoint within the definition can be expanded to display a brief description of the endpoint's functionality and any required inputs for the API call. The Swagger UI provides a useful feature that allows you to issue calls using interactive buttons. Within the API endpoint, the **Try It Out** button becomes available; when you click this button, the **Execute** button appears, enabling you to trigger the API call, as shown in Figure 18-4.

Technet24

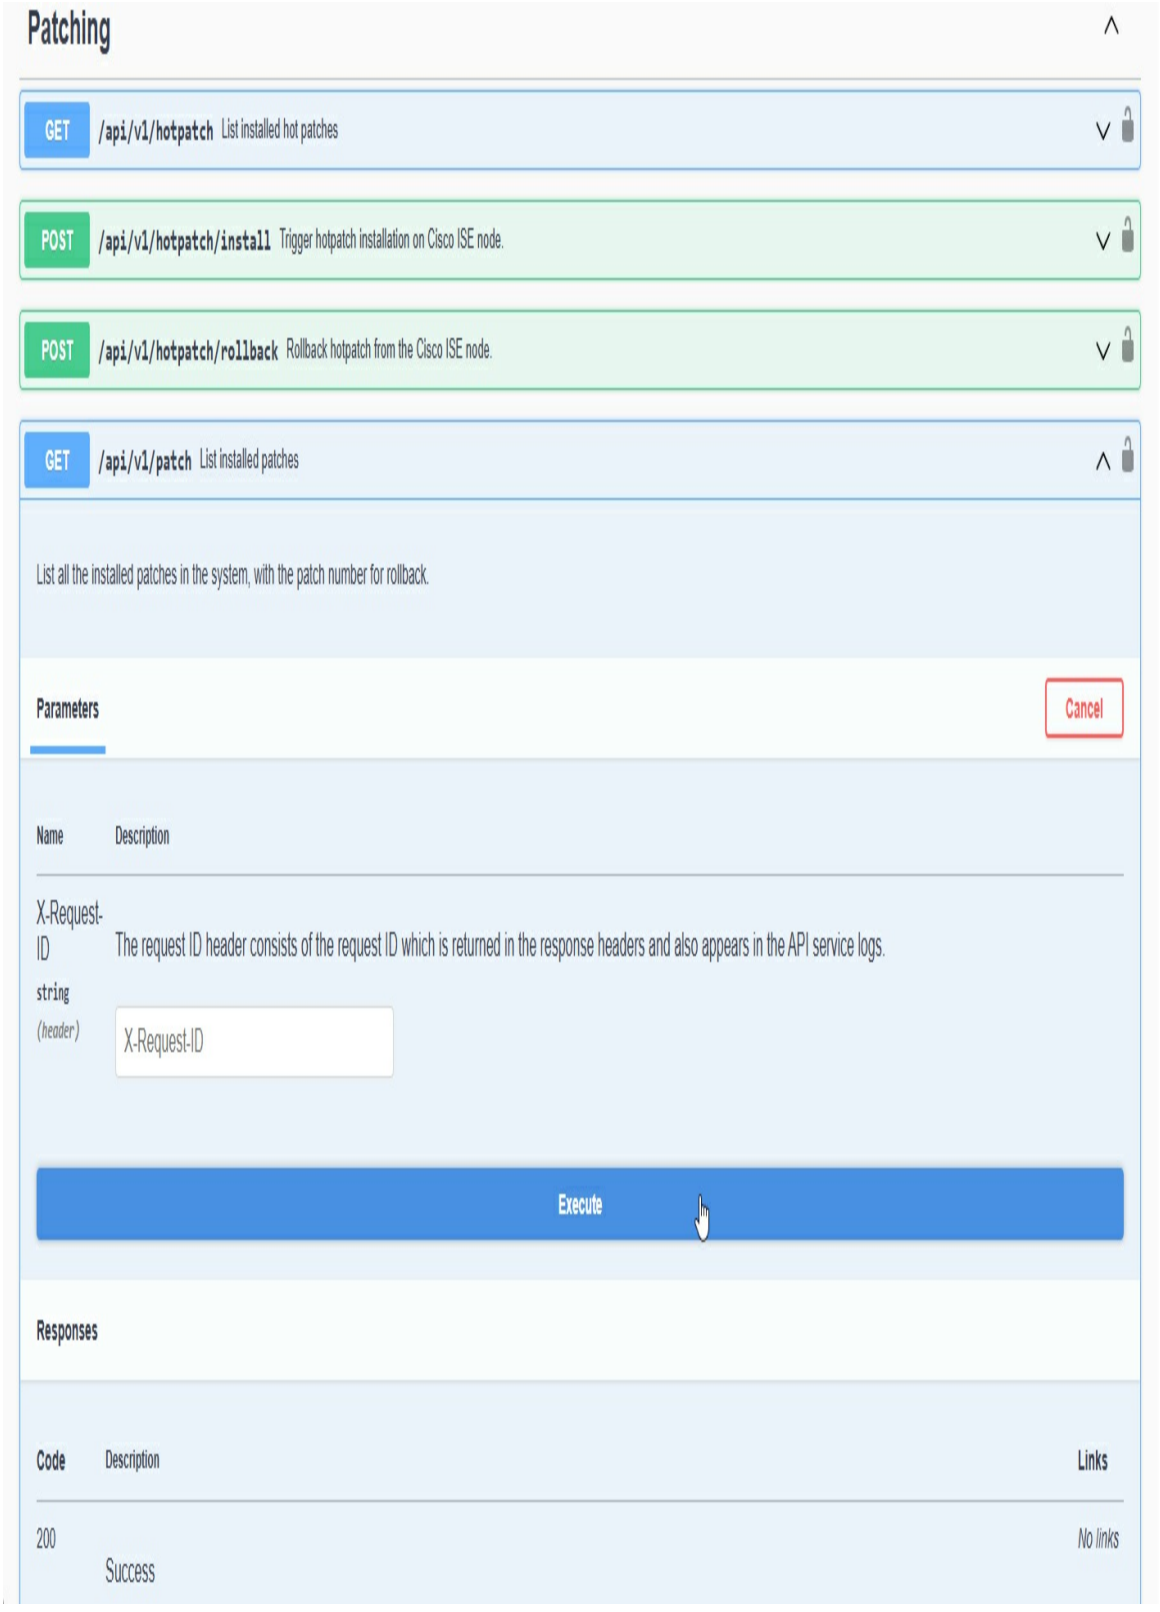


Figure 18-4 *Using the ISE Swagger UI Try It Out Option*

After you execute a call, the server displays the response code and body. In this example, the executed call lists all installed patches, and the response body indicates that patch 3 is installed on ISE 3.5, as shown in Figure 18-5.

Technet24

Curl

```
curl -X 'GET' \
  'https://ise-1.iselab.net:443/api/v1/patch' \
  -H 'accept: application/json'
```

Request URL

```
https://ise-1.iselab.net:443/api/v1/patch
```

Server response

Code	Details

200

Response body


```
{
  "iseVersion": "3.5.0.527",
  "patchVersion": [
    {
      "patchNumber": 3,
      "installDate": "Fri Jul 24 17:36:54 2026"
    }
  ]
}
```


Download

Response headers

```
cache-control: no-cache,no-store,must-revalidate
content-encoding: gzip
content-security-policy: default-src 'self'; script-src 'self' 'unsafe-inline' 'unsafe-eval'; style-src 'self' 'unsafe-inline'; img-src 'self' data:;
content-type: application/json
date: Wed,19 Aug 2026 16:18:51 GMT
expires: Thu,01 Jan 1970 00:00:00 GMT
pragma: no-cache
server:
strict-transport-security: max-age=31536000; includeSubDomains
vary: accept-encoding
x-content-security-policy: default-src 'self'; script-src 'self' 'unsafe-inline' 'unsafe-eval'; style-src 'self' 'unsafe-inline'; img-src 'self' data:;
x-content-type-options: nosniff
x-frame-options: SAMEORIGIN
x-request-id: aa8f9da0-9be9-11f1-8554-1e9c98a9a05d
x-webkit-csp: default-src 'self'; script-src 'self' 'unsafe-inline' 'unsafe-eval'; style-src 'self' 'unsafe-inline'; img-src 'self' data:;
x-xss-protection: 1; mode=block
```

Responses

Code	Description	Links
200	Success	*No links*

Figure 18-5 *The ISE Swagger UI Try It Out Response*

In addition to the result, another useful section after execution is the cURL output, which you can copy and run from Linux-based terminals. You need to add authorization to the generated cURL command because it is not included automatically. To add Basic authorization, use the **–u** flag with either **username:password** for no interaction after the call, or just **username** if you prefer to be prompted interactively for the password. The interactive prompt ensures that the password does not appear in the terminal history. Example 18-1 shows a complete **cURL OpenAPI** command to retrieve the installed patches from ISE, authorized by the admin user.

Example 18-1 *OpenAPI Call to Get ISE Patch Information*

```
shell# curl -u 'admin' -X 'GET' \
>   'https://:443/api/v1/patch' \
>   -H 'accept: application/json'
Enter host password for user 'admin':
{
  "iseVersion" : "3.5.0.527",
  "patchVersion" : [ {
    "patchNumber" : 3,
    "installDate" : "Mon Jul 13 09:41:43 2026"
  } ]
}
```

ERS API

External Restful Services (ERS) provide a set of APIs that enable external clients to perform Create, Read, Update, Delete (CRUD) operations on ISE resources. These calls use HTTPS as the transport protocol, similar to the OpenAPI format. ERS APIs operate on two ports within ISE: TCP 443 and 9060. ERS requests sent to port 443 are routed through the API gateway to the appropriate API component,

Technet24

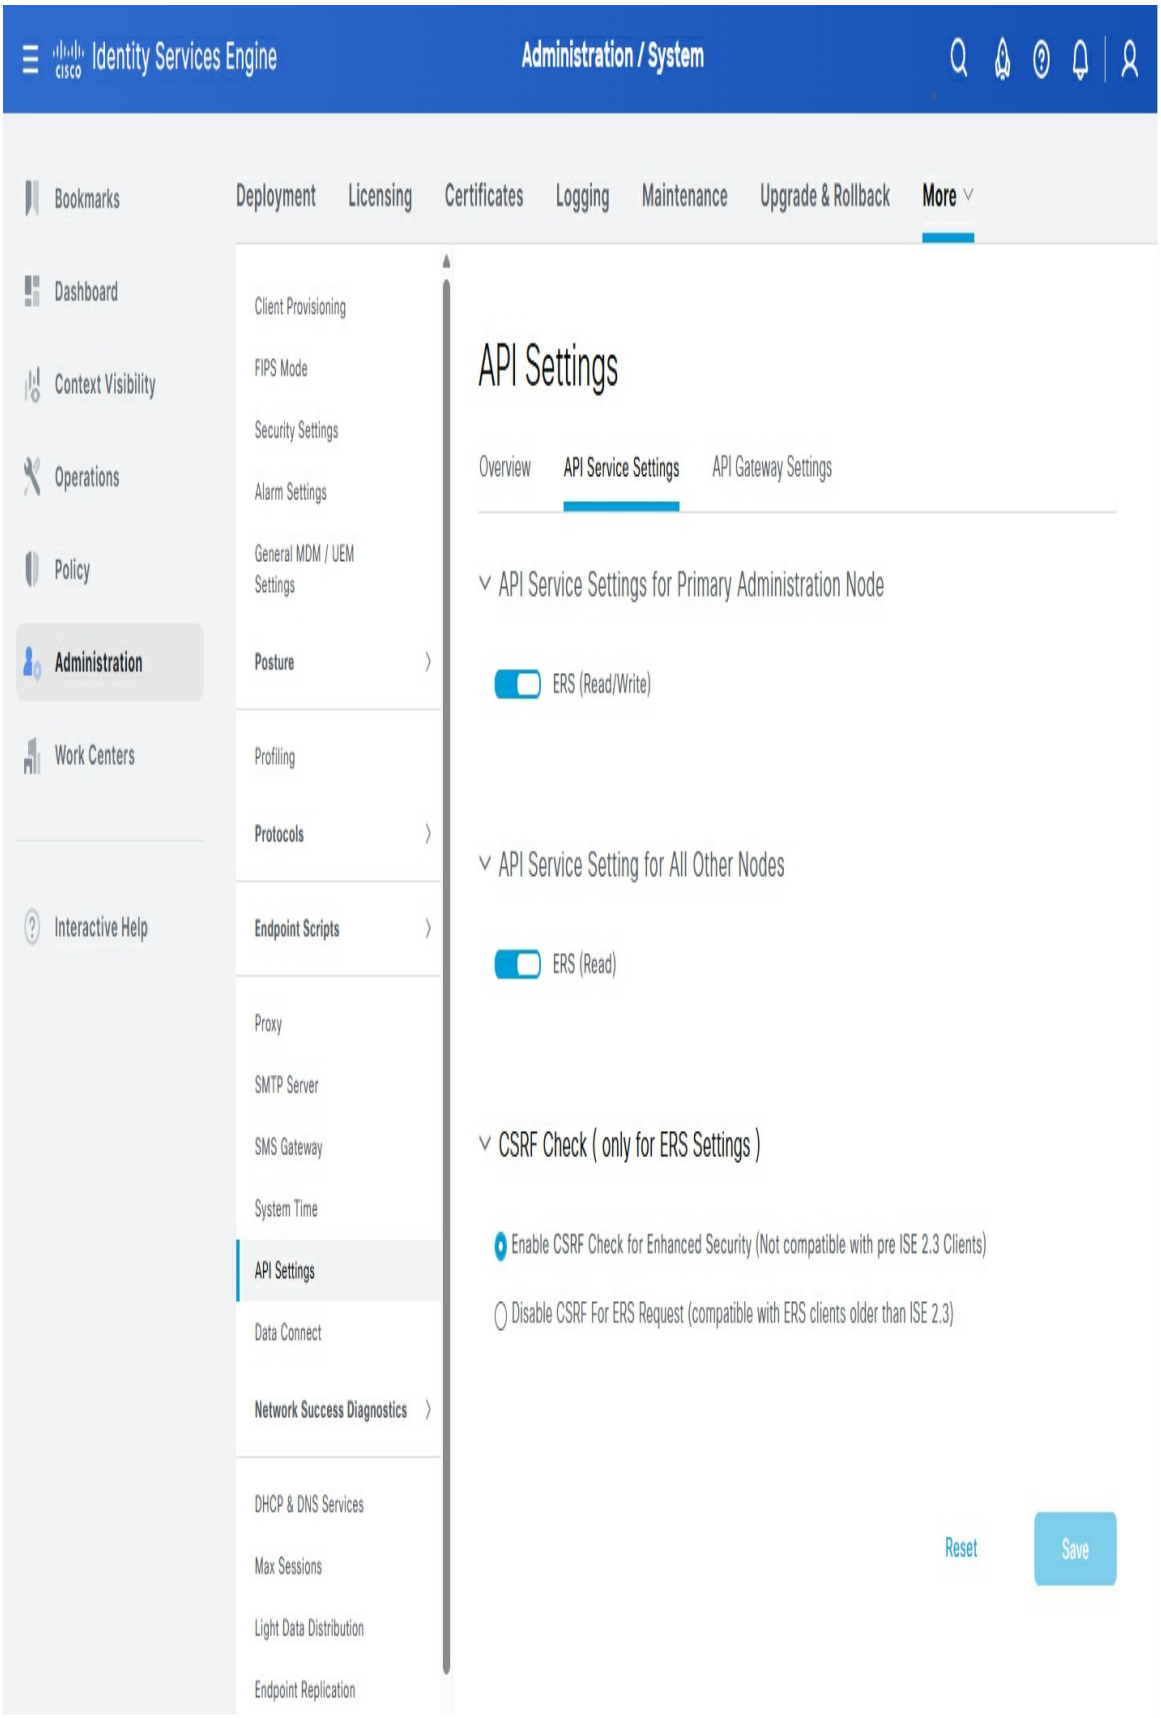
whereas the legacy port 9060 supports only ERS requests.

Best Practice:

You should use port 443 for ERS calls because port 9060 may not be supported for ERS APIs in future Cisco ISE releases.

A common use case for ERS APIs is the integration of external tools with ISE. For example, Cisco Catalyst Center uses ERS APIs to automate the addition and management of network devices, as well as the management of Cisco TrustSec components and policies. Numerous third-party guest management applications also integrate with ISE through ERS APIs to automate guest account creation and guest access management.

ERS is not enabled by default and must be enabled via **Administration > System > Settings > API Settings > API Service Settings**. On this page, you can enable ERS for read and write access on the Primary Administration Node and read-only access for all other nodes, as shown in Figure 18-6. You also can enable a Cross-Site Request Forgery (CSRF) check, in which case the External Restful Services client must send a GET request to fetch the CSRF token from Cisco ISE and include the CSRF token in the requests that are sent to Cisco ISE. Cisco ISE will then validate a CSRF token when a request is received from the External Restful Services client. Cisco ISE processes the request only if the token is valid.


Identity Services Engine
Administration / System
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
More
Client Provisioning
FIPS Mode
Security Settings
Alarm Settings
General MDM / UEM Settings
Posture
Profiling
Protocols
Endpoint Scripts
Proxy
SMTP Server
SMS Gateway
System Time
API Settings
Data Connect
Network Success Diagnostics
DHCP & DNS Services
Max Sessions
Light Data Distribution
Endpoint Replication
API Settings
Overview
API Service Settings
API Gateway Settings
API Service Settings for Primary Administration Node
ERS (Read/Write)
API Service Setting for All Other Nodes
ERS (Read)
CSRF Check (only for ERS Settings)
Enable CSRF Check for Enhanced Security (Not compatible with pre ISE 2.3 Clients)
Disable CSRF For ERS Request (compatible with ERS clients older than ISE 2.3)
Reset
Save

Technet24

Figure 18-6 *ISE ERS API Settings in ISE 3.5*

After you enable ERS, two useful links appear on the **Administration > System > Settings > API Settings > Overview** page, as shown in Figure 18-7. The first link directs you to the ERS Online Software Development Kit (SDK) at https:///ers/sdk. The second link provides a JSON documentation file that contains all available ERS API definitions.

Identity Services Engine | Administration / System

Deployment | Licensing | Certificates | Logging | Maintenance | Upgrade & Rollback | Health Checks | Backup & Restore | More

- Client Provisioning
- FIPS Mode
- Security Settings
- Alarm Settings
- General MDM / UEM Settings
- Posture
- Profiling
- Protocols
- Endpoint Scripts
- Proxy
- SMTP Server
- SMS Gateway
- System Time
- API Settings
- Data Connect
- Network Success Diagnostics
- DHCP & DNS Services
- Max Sessions
- Light Data Distribution

API Settings

Overview | API Service Settings | API Gateway Settings

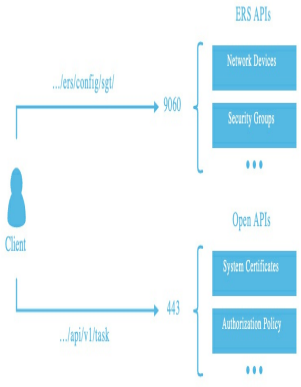
API Services Overview

You can manage Cisco ISE nodes through two sets of API formats–External Restful Services (ERS) and OpenAPI.
Starting Cisco ISE Release 3.1, new APIs are available in the OpenAPI format.
The ERS and OpenAPI services are HTTPS-only REST APIs that operate over port 443.
Currently, ERS APIs also operate over port 9060. However, port 9060 will not be supported for ERS APIs in later Cisco ISE releases. We recommend that you only use port 443 for ERS APIs.
OpenAPI service is enabled by default and cannot be disabled, ERS API service is disabled by default. Enable ERS API service by clicking the toggle button in the API Service Settings tab.
To use either API service, you must have the ERS-Admin or ERS-Operator user group assignment.

For more information on ISE ERS API, please visit:
https://ise-1.iselab.net/ers/sdk

For openapi documention for ERS, click below:
ERS_V1

For more information on ISE Open API, please visit:
https://ise-1.iselab.net/api/swagger-ui/index.html



Technet24

Figure 18-7 *ISE ERS Documentation*

The ERS Online SDK serves as the authoritative source for ERS documentation and definitions. The SDK provides ERS setup resources, release notes detailing ERS changes in the current ISE version and previous releases, definitions for all available ERS APIs, and sample use cases in cURL and Python, as shown in Figure 18-8. You must authenticate with valid admin credentials to access this page.

External RESTful Services (ERS) Online SDK

Quick Reference

- Setting Up
- Successful Flows
- Failure Flows
- Request Headers
- Response Headers
- Error Codes
- Searching a Resource
- Creating a Resource
- Updating a Resource
- Deleting a Resource
- Other Operations
- Bulk Operations

API Documentation

Developer Resources

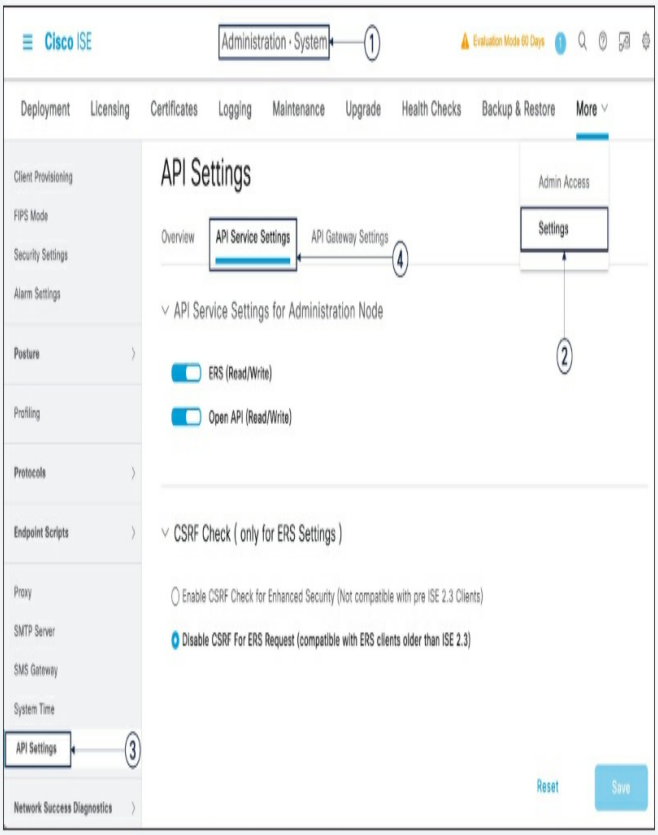
Setting Up

This Online SDK is an application programming reference that explains the operations, parameters, responses of the ISE ERS API. ERS is designed to allow external clients to perform CRUD (Create, Read, Update, Delete) operations on ISE resources. ERS is based on the HTTP protocol and REST methodology. In this section you will learn how to easily enable ERS and to start building your own applications.

Note: The ERS is limited to Active Directory operations only for ISE-PIC nodes.

Enable ERS

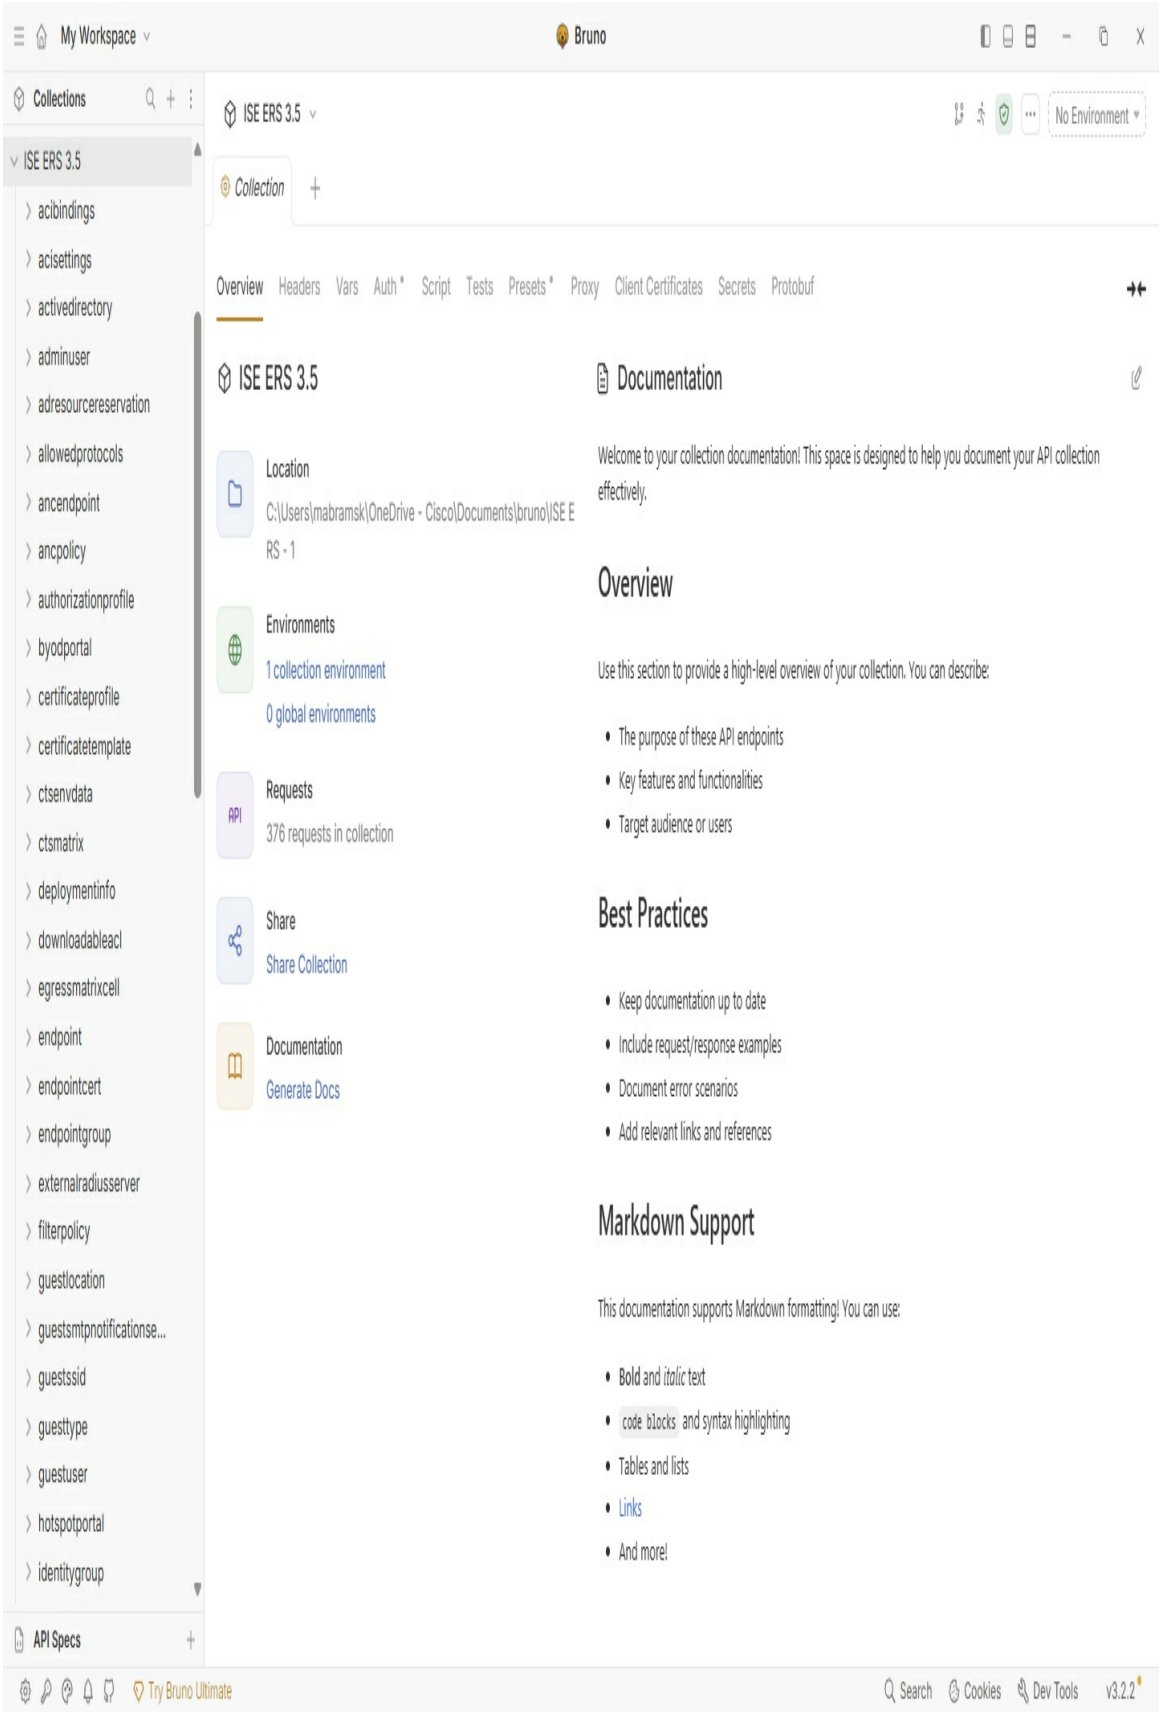
ERS APIs are HTTPS-only REST APIs that operate over port 443 and port 9060.
However, port 9060 might not be supported for ERS APIs in later Cisco ISE releases.
The server responds with a timeout error if you try to access port 443 without first enabling ERS services.
You must first enable ERS services from the Cisco ISE GUI (From the main menu, choose Administration > Settings > API Settings > API Service Settings tab. Click the ERS (Read/Write) toggle button.)



Technet24

Figure 18-8 *ISE ERS Online SDK*

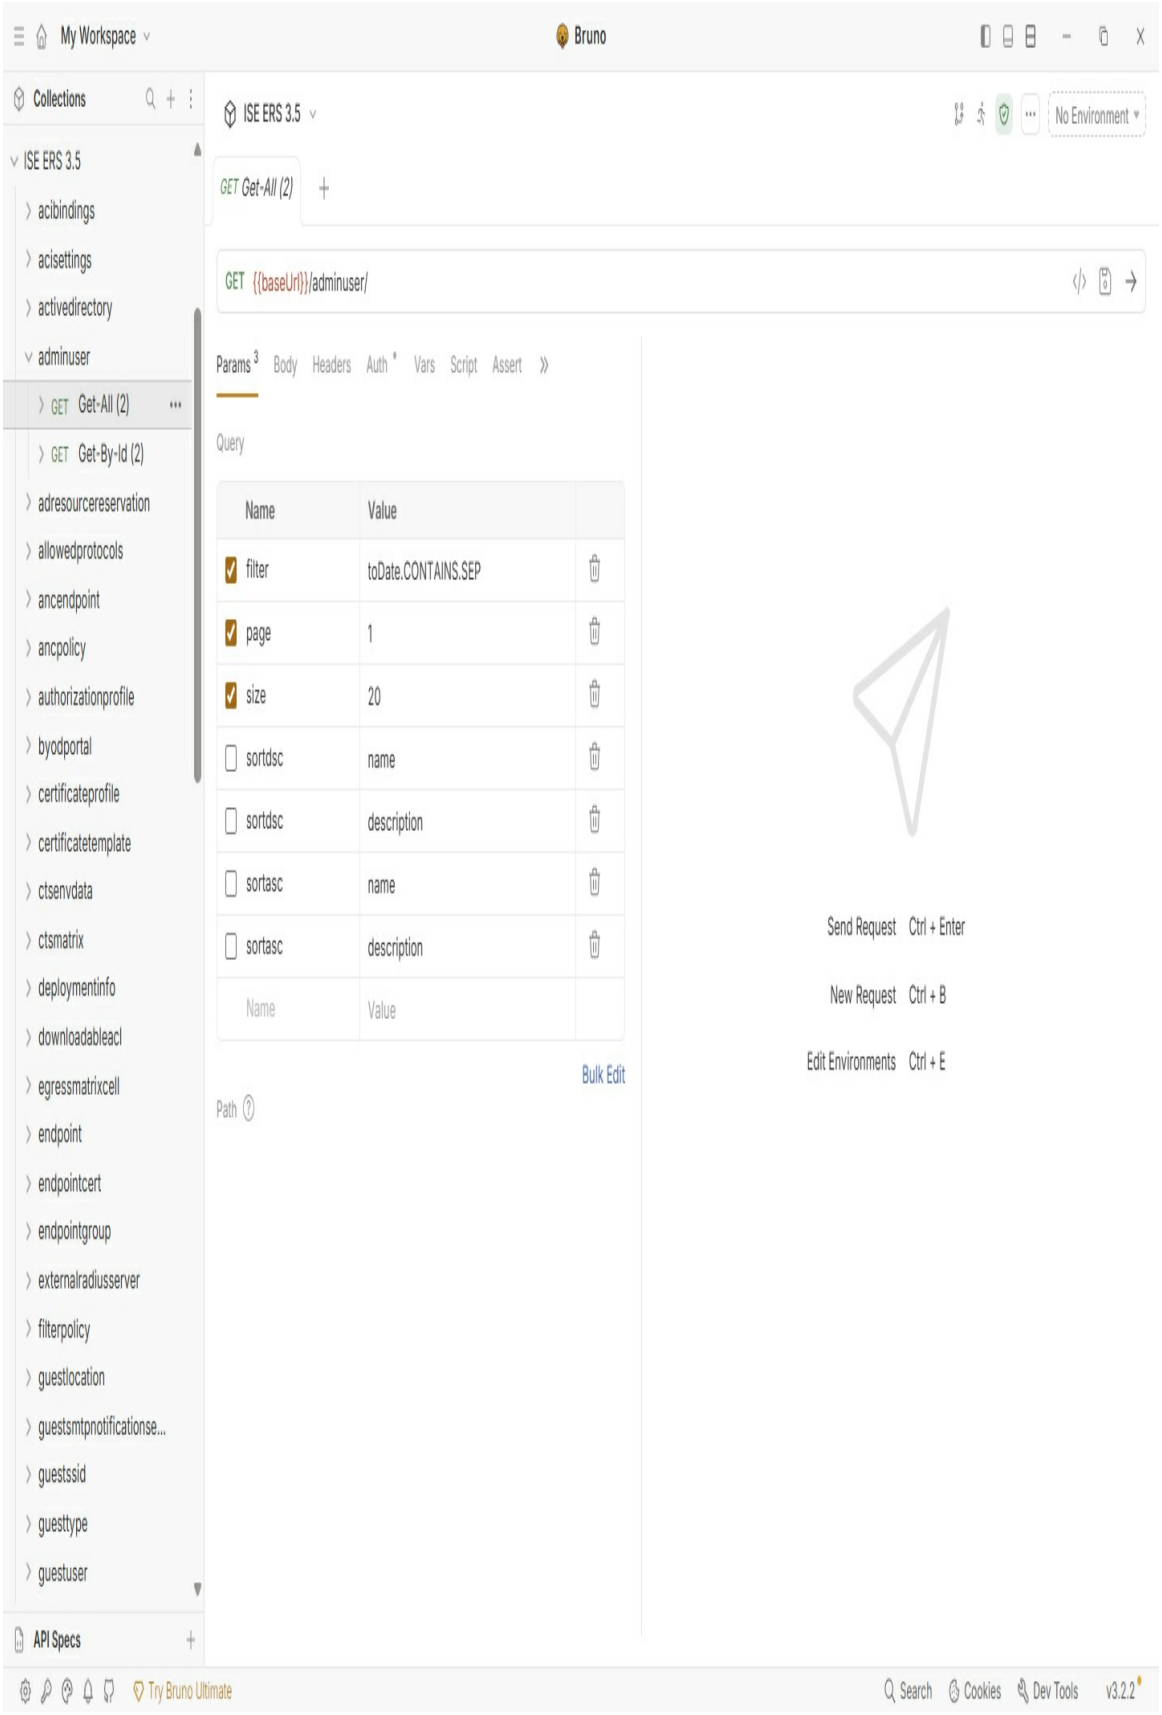
The ERS_V1 JSON file, available from the **Administration > System > Settings > API Settings > Overview** page, can be downloaded and imported as a collection into API client tools such as Postman and Bruno, as shown in Figure 18-9. This file provides the ERS API definitions, enables you to make API calls directly, and allows you to convert these definitions into runnable code in various programming languages within a single interface. The definitions in the JSON collection are organized into logical categories that group the API definitions according to their effects on the ISE application.


My Workspace
Bruno
Collections
ISE ERS 3.5
No Environment
Collection
acibindings
acisettings
activedirectory
adminuser
adresourcereservation
allowedprotocols
ancendpoint
ancpolicy
authorizationprofile
byodportal
certificateprofile
certificatetemplate
ctsenvdata
ctsmatrix
deploymentinfo
downloadableacl
egressmatrixcell
endpoint
endpointcert
endpointgroup
externalradiusserver
filterpolicy
guestlocation
guestsmtpnotificationse...
guestssid
guesttype
guestuser
hotspotportal
identitygroup
API Specs
Overview
Headers
Vars
Auth
Script
Tests
Presets
Proxy
Client Certificates
Secrets
Protobuf
ISE ERS 3.5
Location
C:\Users\mabramsk\OneDrive - Cisco\Documents\bruno\ISE ERS - 1
Environments
1 collection environment
0 global environments
Requests
376 requests in collection
Share
Share Collection
Documentation
Generate Docs
Documentation
Welcome to your collection documentation! This space is designed to help you document your API collection effectively.
Overview
Use this section to provide a high-level overview of your collection. You can describe:
The purpose of these API endpoints
Key features and functionalities
Target audience or users
Best Practices
Keep documentation up to date
Include request/response examples
Document error scenarios
Add relevant links and references
Markdown Support
This documentation supports Markdown formatting! You can use:
Bold and italic text
code blocks and syntax highlighting
Tables and lists
Links
And more!
Try Bruno Ultimate
Search
Cookies
Dev Tools
v3.2.2

Technet24

Figure 18-9 *ISE ERS JSON Imported in Bruno*

The ERS API definitions from the collection include all required parameters, optional parameters, and sample bodies for API definitions where applicable. You can easily edit these elements within the API client tool to test various calls and scenarios for use in scripts and automation, as shown in Figure 18-10.


My Workspace
Bruno
Collections
ISE ERS 3.5
No Environment
GET Get-All (2)
GET {{baseUrl}}/adminuser/
Params 3
Body
Headers
Auth
Vars
Script
Assert
Query
Name
Value
filter
toDate.CONTAINS.SEP
page
1
size
20
sortdsc
name
sortdsc
description
sortasc
name
sortasc
description
Bulk Edit
Path
Send Request Ctrl + Enter
New Request Ctrl + B
Edit Environments Ctrl + E
acibindings
acisettings
activedirectory
adminuser
GET Get-All (2)
GET Get-By-Id (2)
adresourcereservation
allowedprotocols
ancendpoint
ancpolicy
authorizationprofile
byodportal
certificateprofile
certificatetemplate
ctsenvdata
ctsmatrix
deploymentinfo
downloadableacl
egressmatrixcell
endpoint
endpointcert
endpointgroup
externalradiusserver
filterpolicy
guestlocation
guestsmtpnotificationse...
guestssid
guesttype
guestuser
API Specs
Try Bruno Ultimate
Search
Cookies
Dev Tools
v3.2.2

Technet24

Figure 18-10 *ISE ERS Call Execution with Bruno*

Many Get-All ERS API calls provide optional parameters that allow you to filter and sort the results, as well as adjust the pagination of the output. All filter parameters for ERS APIs follow the same structure: a triplet consisting of the field, operator, and value, separated by dots, as shown in Figure 18-11.

External RESTful Services (ERS) Online SDK

Quick Reference

API Documentation

- Admin User
- Advanced Customization Global Set
- Allowed Protocols
- Authorization Profile
- BYOD Portal
- Certificate Template
- CertificateProfile
- Clear Threats and vulnerabilities
- Downloadable ACL
- Egress Matrix Cell
- End Point
- End Point Certificates
- EndPoints Identity Group
- External Radius Server
- Filter Policy
- Guest Location
- Guest Smtp Notification Configurati
- Guest Ssid
- Guest Type
- Guest User
- Hotspot Portal
- IP To SGT Mapping
- IP To SGT Mapping Group
- ISE Service Information
- Identity Group
- Identity Sequence
- Internal User
- Ldap
- My Device Portal
- Native Supplicant Profile
- Network Device
- Network Device Group
- Node Details
- PSN Node Details with Radius Servi
- Portal
- Portal Theme

Developer Resources

Identity Group

```
    "nextPage" : {
      "rel" : "next",
      "href" : "link-to-next-page",
      "type" : "application/xml"
    },
    "previousPage" : {
      "rel" : "previous",
      "href" : "link-to-previous-page",
      "type" : "application/xml"
    }
  }
}
```

Supported Filter and Sorting Fields:

```
Filter: [name, description]

                                To search guest users by using toDate column, please follow the below format:
                                DD-MON-YY (Ex: 13-SEP-18)

                                * Day or Year: GET /ers/config/guestuser/?filter=toDate.CONTAINS.13
                                * Month: GET /ers/config/guestuser/?filter=toDate.CONTAINS.SEP
                                * Date: GET /ers/config/guestuser/?filter=toDate.CONTAINS.13-SEP-18


Sorting: [name, description]
```

Back to top

Delete

Request:

Method:	DELETE
URI:	https://ise-1.iselab.net/ers/config/identitygroup/{id}
HTTP 'Content-Type' Header:	application/xml | application/json
HTTP 'Accept' Header:	application/xml | application/json

Technet24

Figure 18-11 *ISE ERS Calls Filtering*

The field portion of the filter triplet specifies the resource field you want to filter. For example, filtering on the "name" of the resource is a common practice. You can find the available filterable fields for API calls in the API Documentation section of the ERS Online SDK.

The operator within the triplet defines which operation is performed on the value within the selected field. Table 18-1 lists all supported operations and their descriptions. The most suitable operator for a given value depends on the type of data stored in the selected field. For instance, you may find the GT or LT operators more useful in fields containing numeric values than in fields containing string values.

Table 18-1 *API Operators and Description*

Operator	Description
EQ	Equals
NEQ	Not Equals
GT	Greater Than
LT	Less Than
STARTSW	Starts With
NSTARTSW	Not Starts With
ENDSW	Ends With
NENDSW	Not Ends With
CONTAINS	Contains
NCONTAINS	Not Contains

The last member of the triplet, the value, specifies what the operator evaluates within the selected field. You enter exactly what you want to search for in this field, and the input depends on the type of data present in the field being filtered. For example, if you want to filter all Profiler profiles where the name contains "Windows," use **name** as the field, **CONTAINS** as the operator, and **Windows** as the value. The complete filter for the ERS call becomes **filter=name.CONTAINS.Windows**, as shown in Figure 18-12. If you want to filter specifically for the "Windows11-Workstation" profile, set the operator to **EQ** and the value to **Windows11-Workstation**. The complete filter in this case is **filter=name.EQ.Windows11-Workstation**.

Technet24

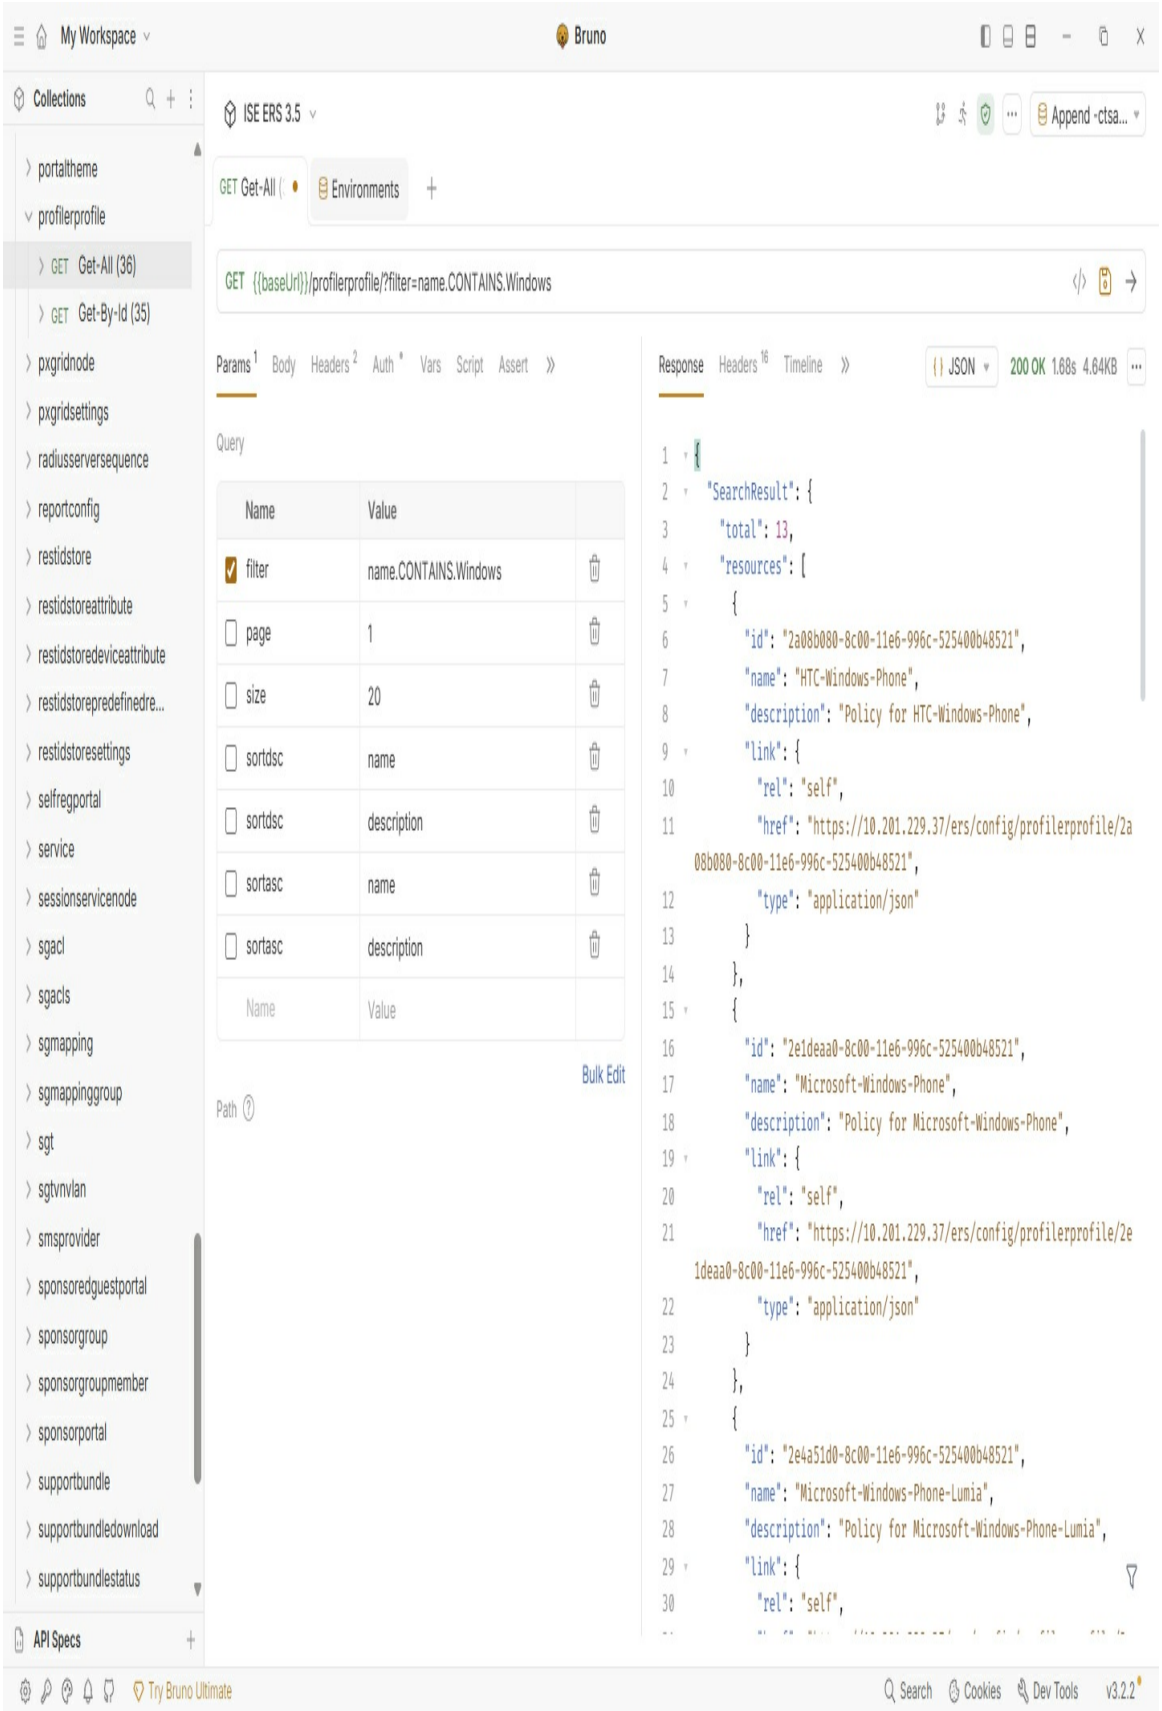

My Workspace
Bruno
Collections
ISE ERS 3.5
Append -ctsa...
GET Get-All
Environments
portaltheme
profilerprofile
GET Get-All (36)
GET Get-By-Id (35)
pxgridnode
pxgridsettings
radiusserversequence
reportconfig
restidstore
restidstoreattribute
restidstoredeviceattribute
restidstorepredefinedre...
restidstoresettings
selfregportal
service
sessionservicenode
sgacl
sgacls
sgmapping
sgmappinggroup
sgt
sgtvnvlan
smsprovider
sponsoredguestportal
sponsorgroup
sponsorgroupmember
sponsorportal
supportbundle
supportbundledownload
supportbundlestatus
API Specs
GET {{baseUrl}}/profilerprofile/?filter=name.CONTAINS.Windows
Params 1
Body
Headers 2
Auth
Vars
Script
Assert
Query
Name
Value
filter
name.CONTAINS.Windows
page
1
size
20
sortdsc
name
sortdsc
description
sortasc
name
sortasc
description
Bulk Edit
Path
Response
Headers 16
Timeline
JSON
200 OK 1.68s 4.64KB
{
"SearchResult": {
"total": 13,
"resources": [
{
"id": "2a08b080-8c00-11e6-996c-525400b48521",
"name": "HTC-Windows-Phone",
"description": "Policy for HTC-Windows-Phone",
"link": {
"rel": "self",
"href": "https://10.201.229.37/ers/config/profilerprofile/2a08b080-8c00-11e6-996c-525400b48521",
"type": "application/json"
}
},
{
"id": "2e1deaa0-8c00-11e6-996c-525400b48521",
"name": "Microsoft-Windows-Phone",
"description": "Policy for Microsoft-Windows-Phone",
"link": {
"rel": "self",
"href": "https://10.201.229.37/ers/config/profilerprofile/2e1deaa0-8c00-11e6-996c-525400b48521",
"type": "application/json"
}
},
{
"id": "2e4a51d0-8c00-11e6-996c-525400b48521",
"name": "Microsoft-Windows-Phone-Lumia",
"description": "Policy for Microsoft-Windows-Phone-Lumia",
"link": {
"rel": "self",
Try Bruno Ultimate
Search
Cookies
Dev Tools
v3.2.2

Figure 18-12 *ISE ERS Call Execution Using Bruno with a Filter*

Another useful optional operation for Get-All ERS calls is sorting. You can use two parameters to sort results: **sortasc** and **sortdsc**. To sort, add the relevant parameter and set it to the field you want to sort. For example, to sort Profiler profiles in descending alphabetical order by profile name, add the parameter "**sortdsc=name**". Example 18-2 demonstrates this query using the Requests Python library, which is commonly employed for scripting API calls.

Example 18-2 *Using the Python Requests Library*

```
import requests

url = "https:///ers/config/profilerprofile/"

querystring = {"sortdsc":"name"}

headers = {
    "accept": "application/json",
}

response = requests.get(url, headers=headers, params=querystring,

print(response.json())
```

The last optional parameters for bulk GET requests are the paging parameters, known as pagination. When an API response includes numerous resources, the response returns them in pages instead of all at once. The page and size parameters control the index of the page and the number of resources returned per response, respectively. By default, the size is set to 20 and the page is set to 1, so the first 20 resources are returned if no parameters are specified. To iterate through pages of 20 resources, you increase the page parameter to 2, then 3, and so forth. You can modify the size parameter to any positive

Technet24

value between 1 and 100 to change the number of resources returned on each page.

MnT API

ISE Monitoring and Troubleshooting (MnT) APIs offer an interface for retrieving and interacting with active and historical RADIUS sessions within a deployment. These APIs are accessible on nodes with the Monitoring persona enabled using TCP port 9443, but you can more easily access them via TCP port 443 on any node with the API gateway enabled. The base URL for MnT APIs is https:///admin/API/mnt/. MnT APIs are supported only for ISE internal admins, and the admin must belong to the Super Admin, System Admin, or MnT Admin group within ISE.

Unlike ERS and OpenAPI, ISE does not provide internal documentation for MnT APIs. However, you can access a Swagger UI documenting MnT APIs at $$$a href="https://developer.cisco.com/docs/identity-services-engine/latest/monitoring-open-api-monitoring-open-api/">https://developer.cisco.com/docs/identity-services-engine/latest/monitoring-open-api-monitoring-open-api/$$$/a>. Another notable difference is that MnT API responses are formatted in XML, whereas ERS and OpenAPI responses typically use JSON. XML formatting displays well in the browser, as shown in Figure 18-13, making browser-based requests a practical choice for quick checks or testing GET MnT APIs.

This XML file does not appear to have any style information associated with it. The document tree is shown below.


```
▼<sessionCount>
   <count>10</count>
 </sessionCount>
```


Technet24

Figure 18-13 *ISE MNT API Call in the Browser*

There are three major categories of MnT API calls: Session Management, Troubleshooting, and Change of Authorization. Session Management API calls allow you to retrieve session counts, session lists, and detailed session attributes for active sessions tracked on the MnT node. The detailed session attribute calls support searches using multiple session identifiers, including the username, NAS IP address, endpoint IP address, session ID, and endpoint MAC address. Figure 18-14 provides an example of retrieving detailed attributes for a session by searching with the endpoint MAC address.


ise-1.iselab.net/admin/API/mnt
Not secure
https://ise-1.iselab.net/admin/API/mnt/Session/MACAddress/8E:A7:6B:EA:78:46
This XML file does not appear to have any style information associated with it. The document tree is shown below.
<sessionParameters>
<passed xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="xs:string">1</passed>
<failed xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="xs:string">0</failed>
<user_name>8E:A7:6B:EA:78:46</user_name>
<nas_ip_address>10.201.229.167</nas_ip_address>
<calling_station_id>8E:A7:6B:EA:78:46</calling_station_id>
<orig_calling_station_id>8E:A7:6B:EA:78:46</orig_calling_station_id>
<cpmsession_id>20010db8000000010000000000000003Sx59/DtnzXDGiXgpbZ2nLodOH8QGiWVG9xl1esQ9OY</cpmsession_id>
<destination_ip_address>10.201.229.37</destination_ip_address>
<device_ip_address>10.201.229.167</device_ip_address>
<network_device_name>SPRT-RCDN</network_device_name>
<acs_server>ise-1</acs_server>
<authentication_method>mab</authentication_method>
<authentication_protocol>Lookup</authentication_protocol>
<framed_ip_address>192.168.87.14</framed_ip_address>
<auth_acs_timestamp>2026-08-19T11:47:27.729-05:00</auth_acs_timestamp>
<execution_steps>11001,11017,11117,11027,15049,15008,15048,15048,15048,15048,15041,22072,15013,24210,24216,15013,24432,24325,24313,24318,24313,24366,24322,24352,24412,15013,24631,24633,15013,2443
<response>{User-Name=8E-A7-6B-EA-78-46; Class=CACS:20010db8000000010000000000000003Sx59/DtnzXDGiXgpbZ2nLodOH8QGiWVG9xl1esQ9OY:ise-3-5-909-37/575232512/84; LicenseTypes=1; }</response>
<posture_status/>
<selected_azn_profiles>PermitAccess</selected_azn_profiles>
<service_type>Call Check</service_type>
<message_code>5200</message_code>
<auth_acsview_timestamp>2026-08-19T11:47:27.729-05:00</auth_acsview_timestamp>
<auth_id>1785869666272478</auth_id>
<location>All Locations</location>
<device_type>All Device Types#SPRT</device_type>
<response_time>41</response_time>
<other_attr_string>StepLatency=1=0;2=0;3=0;4=1;5=0;6=4;7=1;8=0;9=0;10=6;11=3;12=1;13=0;14=1;15=1;16=0;17=4;18=0;19=0;20=0;21=0;22=0;23=0;24=0;25=1;26=0;27=2;28=0;29=0;30=1;31=0;32=0;33=0;34=0;35=
Users:!:SelectedAuthenticationIdentityStores=All_AD_Join_Points:!:SelectedAuthenticationIdentityStores=Guest
Users:!:SelectedAuthenticationIdentityStores=mabramskAD:!:CPMSessionID=20010db8000000010000000000000003Sx59/DtnzXDGiXgpbZ2nLodOH8QGiWVG9xl1esQ9OY:!:Device Type=Device Type#All Device
Types#SPRT:!:AuthenticationStatus=UnknownUser:!:IdentitySelectionMatchedRule=Default:!:OriginalUserName=8E:A7:6B:EA:78:46:!:StepData="6= Network Access.NetworkDeviceName","7= DEVICE.Device
Type","8= Radius.Calling-Station-ID","9= Radius.NAS-IP-Address","11=All_User_ID_Stores","12=Internal Users","15=All_AD_Join_Points","16=All_AD_Join_Points","17=8E-A7-6B-EA-78-
46","18=mabramsk.lab","19=mabramsk.lab","20=test.join.point","21=test.join.point","23=ERROR_NO_SUCH_USER","24=All_AD_Join_Points","25=Guest Users","28=mabramskAD","29=mabramskAD","30=8E-A7-6B-
EA-78-46","31=mabramsk.lab","32=mabramsk.lab","34=ERROR_NO_SUCH_USER","35=mabramskAD","43=mabramskAD","44=8E-A7-6B-EA-78-
46","45=mabramsk.lab","46=mabramsk.lab","48=ERROR_NO_SUCH_USER","49=mabramskAD","50= mabramskAD.ExternalGroups"=StepData:!:Protocol=Radius:!:ISEPolicySetName=SPRT
Tests:!:TrustSec=TrustSec#TrustSec:!:UseCase=Host Lookup:!:DTLSSupport=Unknown:!:NetworkDeviceProfileId=b0699505-3150-4215-a80e-6753d45bf56c:!:Acct-Session-
Id=4657/8E:A7:6B:EA:78:46/E50F:!:IPSEC=IPSEC#Is IPSEC Device#No:!:IdentityPolicyMatchedRule=Default:!:IsMachineIdentity=false:!:RADIUS
Username=8E:A7:6B:EA:78:46:!:CiscoAVPair=:!:IsThirdPartyDeviceFlow=false:!:Framed-MTU=1300:!:Network Device Profile=Cisco:!:EndPointMACAddress=8E-A7-6B-EA-78-46:!:AcsSessionID=ise-3-5-909-
37/575232512/84:!:AuthorizationPolicyMatchedRule=Default:!:Location=Location#All Locations:!:</other_attr_string>
<acct_id>1785869666272479</acct_id>
<acct_acs_timestamp>2026-08-19T11:47:27.838-05:00</acct_acs_timestamp>
<acct_acsview_timestamp>2026-08-19T11:47:27.838-05:00</acct_acsview_timestamp>
<acct_session_id>4657/8E:A7:6B:EA:78:46/E50F</acct_session_id>
<acct_status_type>Start</acct_status_type>
<acct_input_octets>0</acct_input_octets>
<acct_output_octets>0</acct_output_octets>
<acct_input_packets>0</acct_input_packets>

Technet24

Figure 18-14 *ISE MNT API Session Call*

Troubleshooting API calls enable you to quickly retrieve basic information from ISE. These calls include obtaining the installed ISE version and retrieving descriptions and codes for all defined failure reasons in ISE. The most general utility is provided by calls that query available authentication and accounting details for endpoints based on MAC addresses. You can use these calls to automate queries and obtain information similar to what appears on the **RADIUS Live Logs** page within ISE.

The last category of MnT APIs includes calls that can not only retrieve information but also directly affect live sessions. This category features the CoA APIs. There are two CoA-related calls: reauthentication and disconnect. A reauthentication CoA causes the session to be immediately reauthenticated and reauthorized, which matches the CoA performed in profiling, posture, and guest flows within ISE. The disconnect CoA forces the session to disconnect immediately. With the disconnect API, you have two options corresponding to port actions. The Port Bounce option performs a "shutdown" followed by a "no shutdown" on the connected port, causing the endpoint to temporarily disconnect and then reconnect. The Port Shutdown option performs a "shutdown" on the connected port and leaves it in that state.

ISE Automation with Terraform

Previous chapters explored various ISE installation methods. You can install ISE on dedicated hardware appliances, virtual machines using supported hypervisors, and in the cloud. However, this approach requires separate installation for each ISE node in the deployment. Although setting up a single ISE node is straightforward, building a

deployment with 20 nodes significantly increases the setup time. Infrastructure as Code (IaC) tools address this challenge by allowing you to declare the desired setup in text files and deploy the entire infrastructure with a single command. You can also remove the entire infrastructure just as easily. IaC tools help you save time on infrastructure maintenance; for example, by adding an extra code block, you can introduce an additional ISE node. Terraform is not limited to ISE deployments. When deploying ISE nodes, you rely not on ISE APIs, but on hypervisor or cloud APIs, depending on where ISE is deployed.

Consider a basic example of deploying an ISE node in AWS. Example 18-3 demonstrates the results from macOS, which serves as the platform for initiating the Terraform deployment.

Example 18-3 *Terraform Apply Execution*

```
% terraform apply

Terraform used the selected providers to generate the following ex
are indicated with the following symbols:
  + create

Terraform will perform the following actions:

  < Output omitted. Terraform is presenting what changes will be m

Plan: 1 to add, 0 to change, 0 to destroy.

Do you want to perform these actions?
  Terraform will perform the actions described above.
  Only 'yes' will be accepted to approve.

  Enter a value: yes

aws_instance.ise1: Creating...
aws_instance.ise1: Still creating... [10s elapsed]
```

Technet24

```
aws_instance.ise1: Creation complete after 13s [id=i-0b0e754157c6

Apply complete! Resources: 1 added, 0 changed, 0 destroyed.
%
```

In this example, Terraform deployed a single ISE instance injust 13 seconds. The process scales efficiently; when you deploy four instances instead of one, the deployment also completes in 13 seconds, as shown in Example 18-4.

Example 18-4 *Terraform Apply Completion*

```
aws_instance.ise3: Creating...
aws_instance.ise1: Creating...
aws_instance.ise2: Creating...
aws_instance.ise4: Creating...
aws_instance.ise4: Still creating... [10s elapsed]
aws_instance.ise1: Still creating... [10s elapsed]
aws_instance.ise3: Still creating... [10s elapsed]
aws_instance.ise2: Still creating... [10s elapsed]
aws_instance.ise4: Creation complete after 13s [id=i-0e6fcc51b618
aws_instance.ise2: Creation complete after 13s [id=i-07975e10ee24
aws_instance.ise3: Creation complete after 13s [id=i-05a4369196f0
aws_instance.ise1: Creation complete after 13s [id=i-045eb4f18d2d

Apply complete! Resources: 4 added, 0 changed, 0 destroyed.
```

Terraform code for deployment is straightforward. In Example 18-5, the first block defines the providers and provider settings used with Terraform. The second block specifies the resources to be created, in this case, the ISE virtual machine.

Example 18-5 *Terraform Configuration File*

```
terraform {
  required_providers {
```

```
    aws = {
      source  = "hashicorp/aws"
      version = "~> 5.0"
    }
  }
  required_version = ">= 1.10.0"
}

provider "aws" {
  profile = "default"
  region  = "us-east-1"
}

# Create AWS ISE VM Instance

resource "aws_instance" "ise1" {
  ami             = "ami-0430b7c7e99705989"
  instance_type   = "m5.8xlarge"
  user_data       = file("user_data_ise1")
  key_name        = "AWS2"
  subnet_id       = "subnet-00c74ba0977cdcf9a"
  private_ip      = "172.31.200.200"
  security_groups = ["sg-019798ca1592bb4c5"]
  tags = {
    Name = "ISE-AWS"
  }
}
```

To deploy the ISE instance as a virtual machine in AWS, you must define several parameters:

- AMI image ID of the required version, which you can obtain from the Marketplace
- instance_type that matches the resources allocated for the ISE VM
- key_name used for SSH access to the ISE VM
- User data, which is the most critical parameter and contains attribute-value pairs for ISE bootstrap settings

Technet24

- AWS network settings

Best Practice:

Always consult the latest ISE documentation for the syntax of the user data based on the ISE version you are deploying. Most issues with Terraform deployments arise from mistakes in the syntax of the ISE user data file.

Example 18-5 demonstrates ISE deployment in AWS with Terraform, but you are not limited to AWS. You also can deploy ISE using IaC in other cloud providers and hypervisors.

ISE Automation with Ansible

Ansible is another IaC tool. In practice, you typically use Terraform for infrastructure management and Ansible for configuration management. You can use Ansible with ISE as well. Ansible offers several benefits: It automates configuration management, operates without agents, uses YAML for easy-to-read configuration, and is idempotent, meaning it applies only necessary changes even if you run the playbook multiple times.

The playbook serves as a core component of Ansible, where you define the configuration to deploy. Example 18-6 shows the execution of the playbook.

Example 18-6 *Ansible Playbook Execution*

```
(A) ekorneyc@EKORNEYC-PC Ansible % ansible-playbook -i hosts ise-
PLAY [ise_servers]
**********************************************************************
**************************************************
TASK [Create User Group VPN_Users]
**********************************************************************
************************************ changed: [ise_server]
```

```
TASK [Create User Alice]
*********************************************************************
********************************************** changed: [ise_s
PLAY RECAP
*********************************************************************
********************************************************************* i
unreachable=0 failed=0 skipped=0 rescued=0 ignored=0
(A) ekorneyc@EKORNEYC-PC Ansible %
```

The playbook shown in Example 18-6 configures a user group and a user, then adds the user to the group. The content of the playbook is provided in Example 18-7.

Example 18-7 *Ansible Playbook*

```
- hosts: ise_servers
  gather_facts: no
  tasks:

  - name: Create User Group VPN_Users
    cisco.ise.identity_group:
      ise_hostname: "{{ise_hostname}}"
      ise_username: "{{ise_username}}"
      ise_password: "{{ise_password}}"
      ise_verify: "{{ise_verify}}"
      state: present
      name: VPN_Users
    register: group

  - name: Create User Alice
    cisco.ise.internal_user:
      ise_hostname: "{{ise_hostname}}"
      ise_username: "{{ise_username}}"
      ise_password: "{{ise_password}}"
      ise_verify: "{{ise_verify}}"
      state: present
      identityGroups: "{{ group.ise_response.id }}"
      name: Alice
      password: "{{vpn_password}}"
```

Technet24

```
    register: result
```

Variables enclosed in curly brackets are stored separately in the group_vars directory.

Example 18-7 creates a basic configuration, but you also can develop advanced configurations. Always consult the latest Ansible Galaxy library at https://galaxy.ansible.com/ui/repo/published/cisco/ise/docs/ for the list of supported modules.

Summary

ISE exposes multiple types of APIs and enables you to develop custom applications for automating critical tasks. In addition, ISE integrates with Infrastructure as Code (IaC) tools, allowing you to scale infrastructure deployments, minimize human errors, and significantly reduce the time required for deploying infrastructure and configurations. The ISE Github page contains multiple examples of Ansible and Terraform automations as well as SDK examples; it is located at https://github.com/CiscoISE.

Chapter 19. Integrating Your Security Products to Enhance Context Using pxGrid

Introduction

In the modern enterprise, security infrastructure is rarely composed of products from a single vendor. While the "best-of-breed" approach offers specialized protection, it often creates security silos. These silos prevent disparate systems from sharing critical information, leading to fragmented visibility and delayed incident response.

This chapter explores how the Cisco Platform Exchange Grid (pxGrid) acts as the connective tissue for your security architecture, enabling seamless data exchange and deep contextual awareness. We will also examine how pxGrid Direct extends these capabilities to environments where native integration might be limited.

Traditional security models rely on static rules. However, modern threats are dynamic. To defend against them, security products need "context" — the "who, what, where, when, and how" of a network connection.

By sharing contextual data, such as user identity, device posture, and threat intelligence, across your stack, you can move from reactive security to proactive orchestration. For example, when a firewall

Technet24

detects a suspicious flow, it can query the identity engine to determine the user and device, then trigger a policy change on the network access control (NAC) system to isolate the endpoint.

Understanding pxGrid Architecture

Cisco pxGrid is an open, scalable, and secure framework that allows security products to exchange information in real time. It follows a hub-and-spoke model consisting of three primary roles. The pxGrid Controller acts as the centralized broker that manages connections, subscriptions, and security certificates. The Publisher is a product that provides information, such as Cisco ISE publishing user identity data. The Subscriber is a product that consumes information, such as a SIEM or firewall subscribing to identity updates.

This framework offers several key benefits. Standardization is achieved by using a common language for cross-vendor communication. Security is maintained through mutual TLS (mTLS), which ensures that only authenticated and authorized nodes can exchange data. Finally, scalability is improved because the model decouples publishers and subscribers, allowing the network to grow without the need for complex point-to-point configurations.

From Cisco Identity Services Engine release 3.1, all pxGrid connections must use what is known as pxGrid 2.0. Prior to that, pxGrid 1.0 used the Extensible Messaging and Presence Protocol (XMPP), which had disadvantages such as requiring heavy XML processing (which is not ideal for application-to-application communication) and needing to integrate with unsupported or undocumented client software development kits (SDKs). Only pxGrid 2.0 will be covered in this book.

pxGrid 2.0 utilizes REST and WebSocket. These use industry standards, they are simple and readily supported, and message bodies are

formatted in the popular JSON format.

REST is used for one-shot queries providing quick extensible querying mechanisms.

WebSocket provides a quick, scalable, and persistent bidirectional data transfer; can transport any number of protocols; and can provide message delivery without polling. The messages over WebSocket conform to the Simple Text Oriented Messaging Protocol (STOMP), where the following commands are supported on pxGrid:

- **CONNECT**
- **DISCONECT**
- **SUBSCRIBE**
- **UNSUBSCRIBE**
- **SEND**
- **MESSAGE**
- **ERROR**

pxGrid consumers implement the **SUBSCRIBE** command, whereas providers implement **SEND**.

pxGrid 2.0 services are published on Github. Figure 19-1 shows the pxGrid wiki home page (https://github.com/cisco-pxgrid/pxgrid-rest-ws/wiki), along with the available ISE services on the right side.

Technet24

Home

alei121 edited this page on Aug 13, 2019 · 7 revisions

pxGrid

pxGrid is a protocol framework that defines the control mechanisms to facilitate machine-to-machine communications.

There are 2 types of pxGrid clients:

1. pxGrid service consumer
2. pxGrid service provider

Cisco ISE services comprise of HTTP APIs and WebSocket pubsub. It provides access to session context and other features in Cisco ISE.

Pages 18

Find a page or section...

- Home
  - pxGrid
- ANC configuration
- Cisco ISE services
- Endpoint
- Endpoint Asset
- MDM
- Profiler configuration
- Pubsub
- pxGrid Consumer
- pxGrid Filtering
- pxGrid Overview
- pxGrid Provider
- Radius Failure
- Session Directory
- System Health
- TrustSec
- TrustSec configuration
- TrustSec SXP

Figure 19-1 *ISE pxGrid 2.0 Home Page and Services Within Github*

To consume context via these pxGrid services, three core components are used:

- **restBaseURL:** This is the primary entry point for the pxGrid REST API. It acts as the "discovery" gateway. When a client first connects to the pxGrid node, it uses the restBaseURL to authenticate and query the available services. It provides the metadata necessary for the client to understand what services are offered and how to reach them. The client initiates the connection by contacting the restBaseURL. Through a series of GET requests, the client retrieves the service registry. This registry contains the locations of various services, including the specific wsPubsubService URL.
- **wsPubsubService:** This represents the WebSocket-based Publish/Subscribe (Pub/Sub) service. While the REST API is used for initial configuration and discovery, the actual real-time exchange of data occurs over a persistent WebSocket connection. The wsPubsubService is the specific endpoint URL that the client must connect to for receiving or sending real-time notifications. Once the client has discovered the wsPubsubService URL, it upgrades its connection to a WebSocket. This persistent channel is the "pipe" through which all real-time events will flow. Without the wsPubsubService endpoint, the client would be restricted to manual polling via REST, which is inefficient for high-volume, time-sensitive data.
- **Topic:** A topic is the logical category of data being exchanged. Examples include session directory updates, endpoint profiling data, or Adaptive Network Control (ANC) status. Topics provide the granularity required to ensure that a client receives only the specific data it is authorized to consume or is interested in

Technet24

processing. With the WebSocket connection established, the client must specify what it wants to hear. The client sends a subscription request to the wsPubsubService identifying the specific topic (such as com.cisco.ise.session). The wsPubsubService then filters the stream of data, ensuring that only messages associated with that topic are pushed to the client.

In essence, the restBaseURL allows the client to find the wsPubsubService, and the wsPubsubService allows the client to subscribe to a specific topic to receive actionable, real-time intelligence from the ISE ecosystem. This hierarchy allows robust integrations by decoupling the discovery process from the high-throughput data delivery mechanism, ensuring seamless and scalable interoperability between Cisco ISE and other Cisco and third-party security solutions.

The following are some sample pxGrid topics:

- /topic/com.cisco.ise.session
- /topic/com.cisco.ise.config.profiler
- /topic/com.cisco.ise.config.trustsec.security.group
- /topic/com.cisco.ise.sxp.binding
- /topic/com.cisco.ise.endpoint
- /topic/com.cisco.ise.config.anc.status
- /topic/com.cisco.ise.radius.failure
- /topic/com.cisco.ise.config.trustsec.security.group.acl
  /topic/com.cisco.ise.config.trustsec.egress.policy

Each topic includes several objects that define the data exchanged. For example, the objects within the /topic/com.cisco.ise.sxp.binding topic may look like the following:

```
{
    "bindings": [{
        "ipPrefix": "10.4.1.10/32",
        "tag": 33,
        "source": "10.5.1.21",
        "peerSequence": "10.5.1.21"
    }]
}
```

With each topic containing a range of varying objects, clients can subscribe to the topics relevant to the context needed to be exchanged. You can navigate to a full list of objects on Github at https://github.com/cisco-pxgrid/pxgrid-rest-ws/wiki

ISE supports up to four pxGrid nodes providing high availability. A client can connect to any of the four nodes and be able to connect to the others if an outage occurs. pxGrid capability is enabled per ISE node, as shown in Figure 19-2.

Technet24

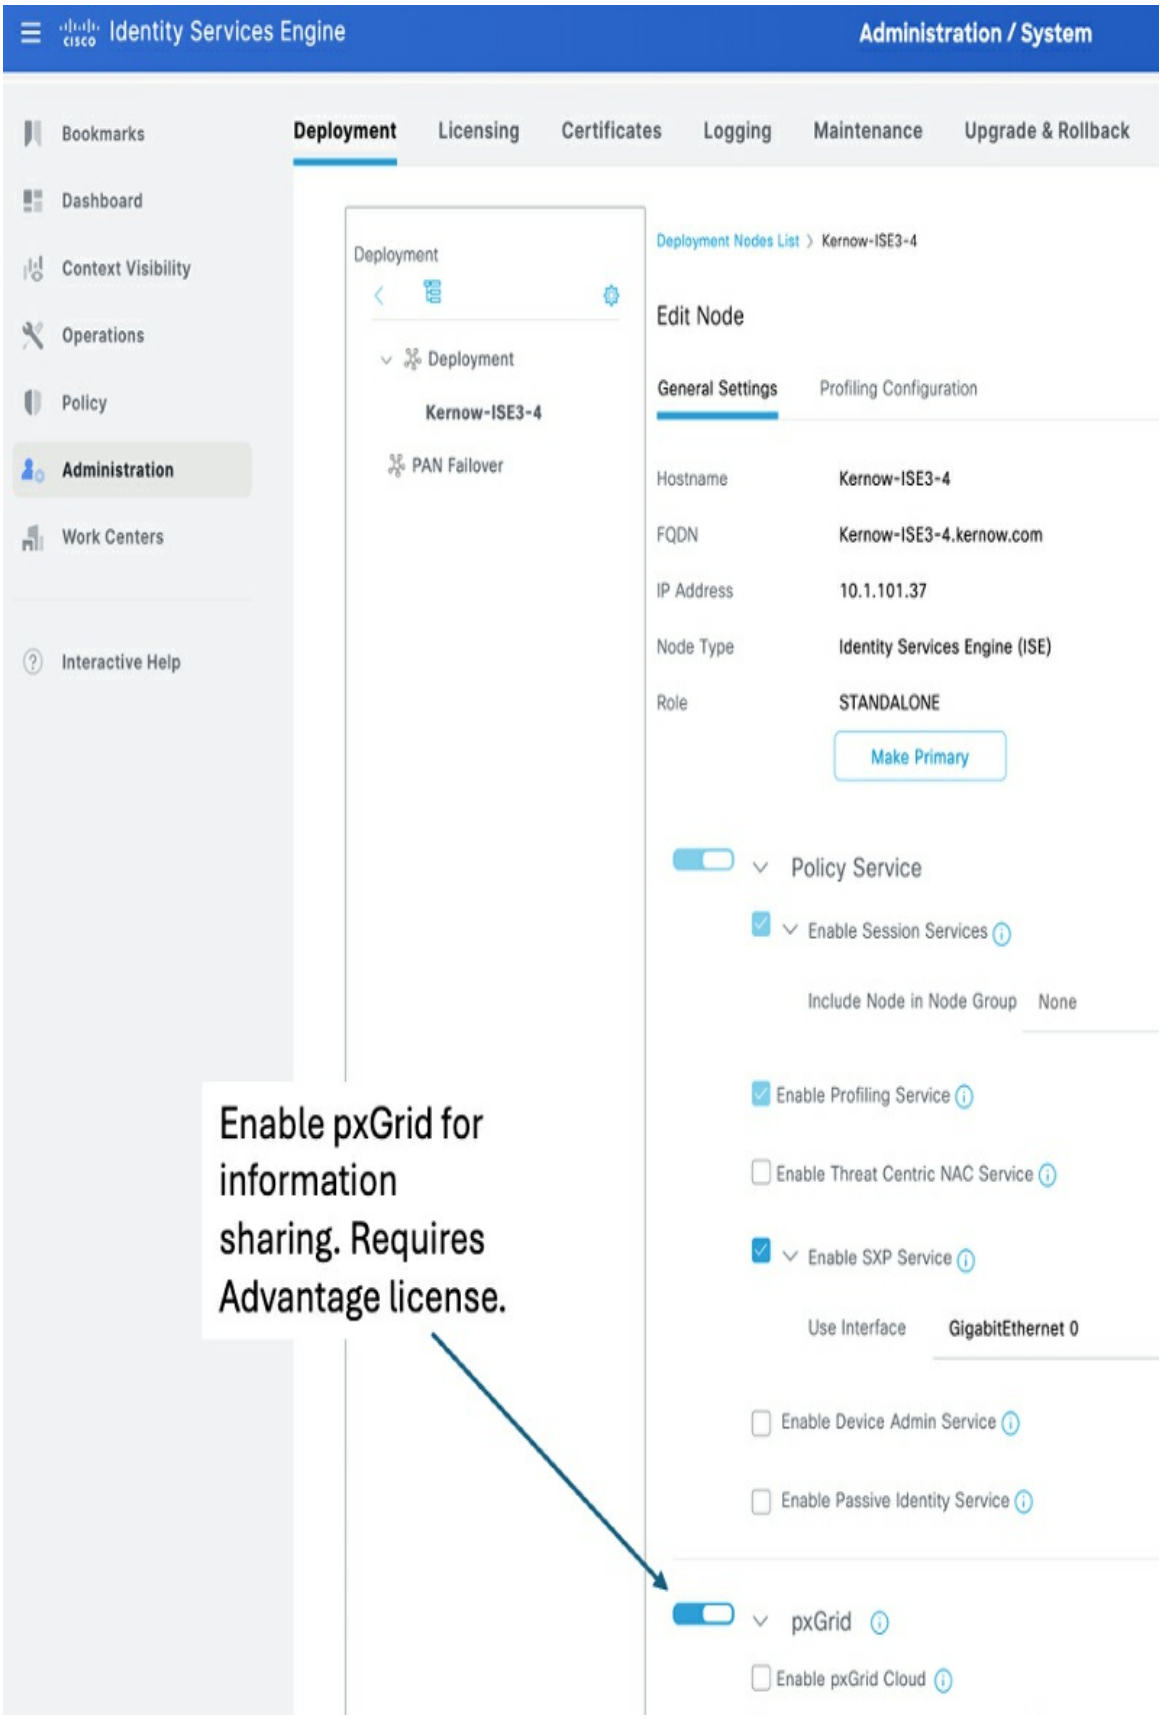

Identity Services Engine
Administration / System
Bookmarks
Dashboard
Context Visibility
Operations
Policy
Administration
Work Centers
Interactive Help
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Deployment
Deployment
Kernow-ISE3-4
PAN Failover
Deployment Nodes List > Kernow-ISE3-4
Edit Node
General Settings
Profiling Configuration
Hostname
Kernow-ISE3-4
FQDN
Kernow-ISE3-4.kernow.com
IP Address
10.1.101.37
Node Type
Identity Services Engine (ISE)
Role
STANDALONE
Make Primary
Policy Service
Enable Session Services
Include Node in Node Group
None
Enable Profiling Service
Enable Threat Centric NAC Service
Enable SXP Service
Use Interface
GigabitEthernet 0
Enable Device Admin Service
Enable Passive Identity Service
pxGrid
Enable pxGrid Cloud
Enable pxGrid for information sharing. Requires Advantage license.

Figure 19-2 *Enabling pxGrid on the ISE Node*

When clients first connect to the ISE pxGrid node, the connectivity must be approved within ISE. The user interface will inform you, as administrator, when an approval is required. It may be beneficial to allow ISE to automatically approve these requests if that makes sense in your environment. To automatically approve the requests, navigate to **Administration > pxGrid Services > Settings** and select and save the option **Automatically Approve New Certificate-Based Accounts**, as shown in Figure 19-3.

Technet24

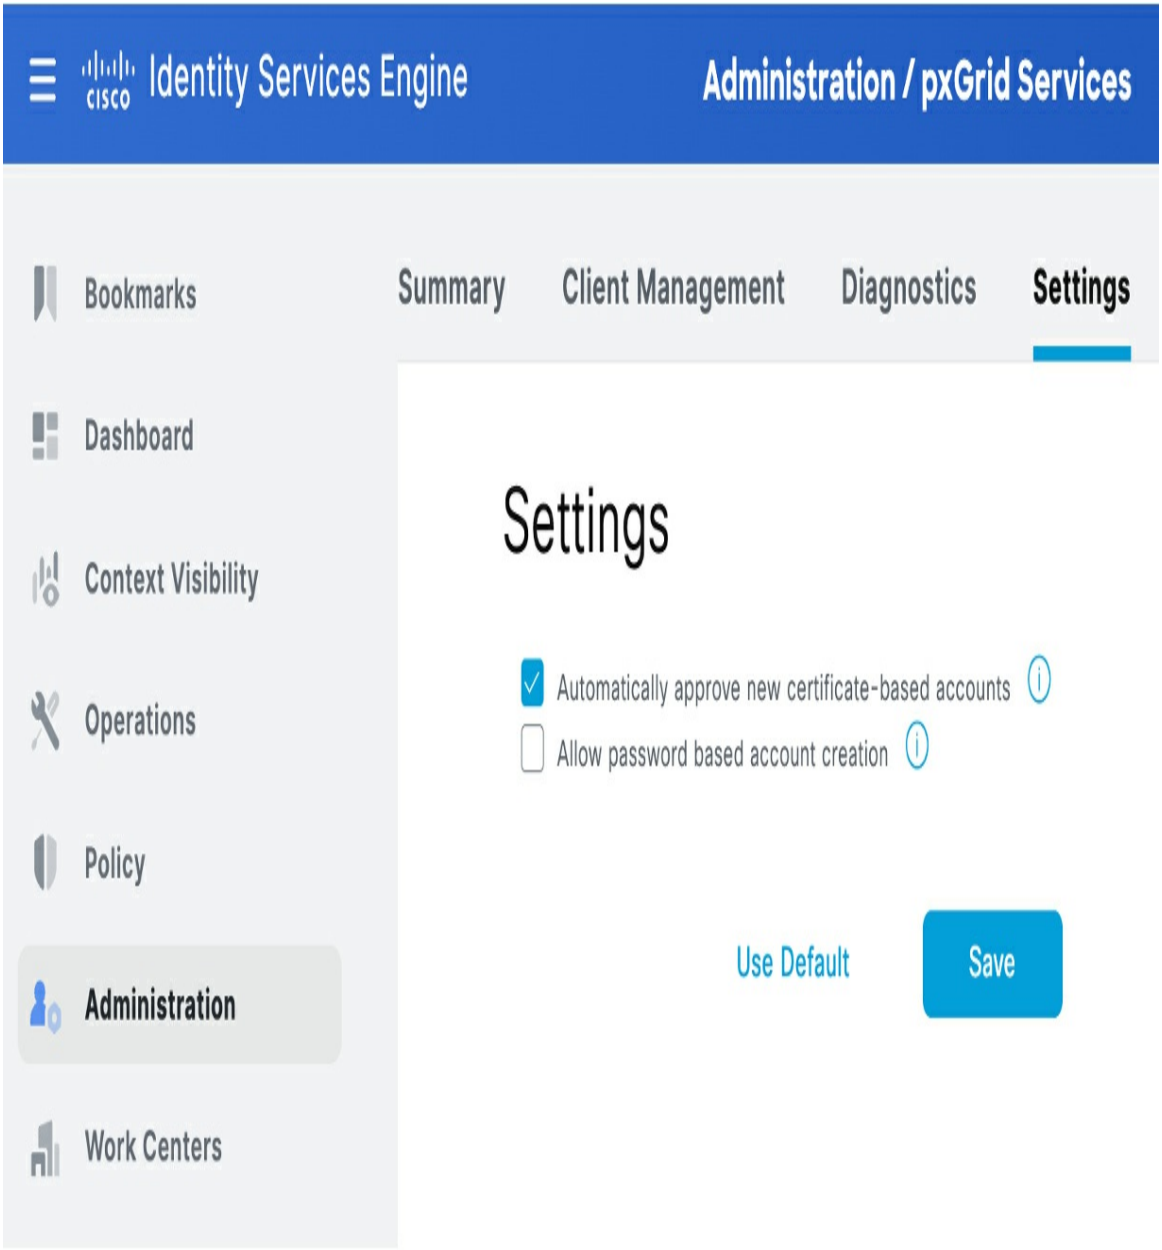
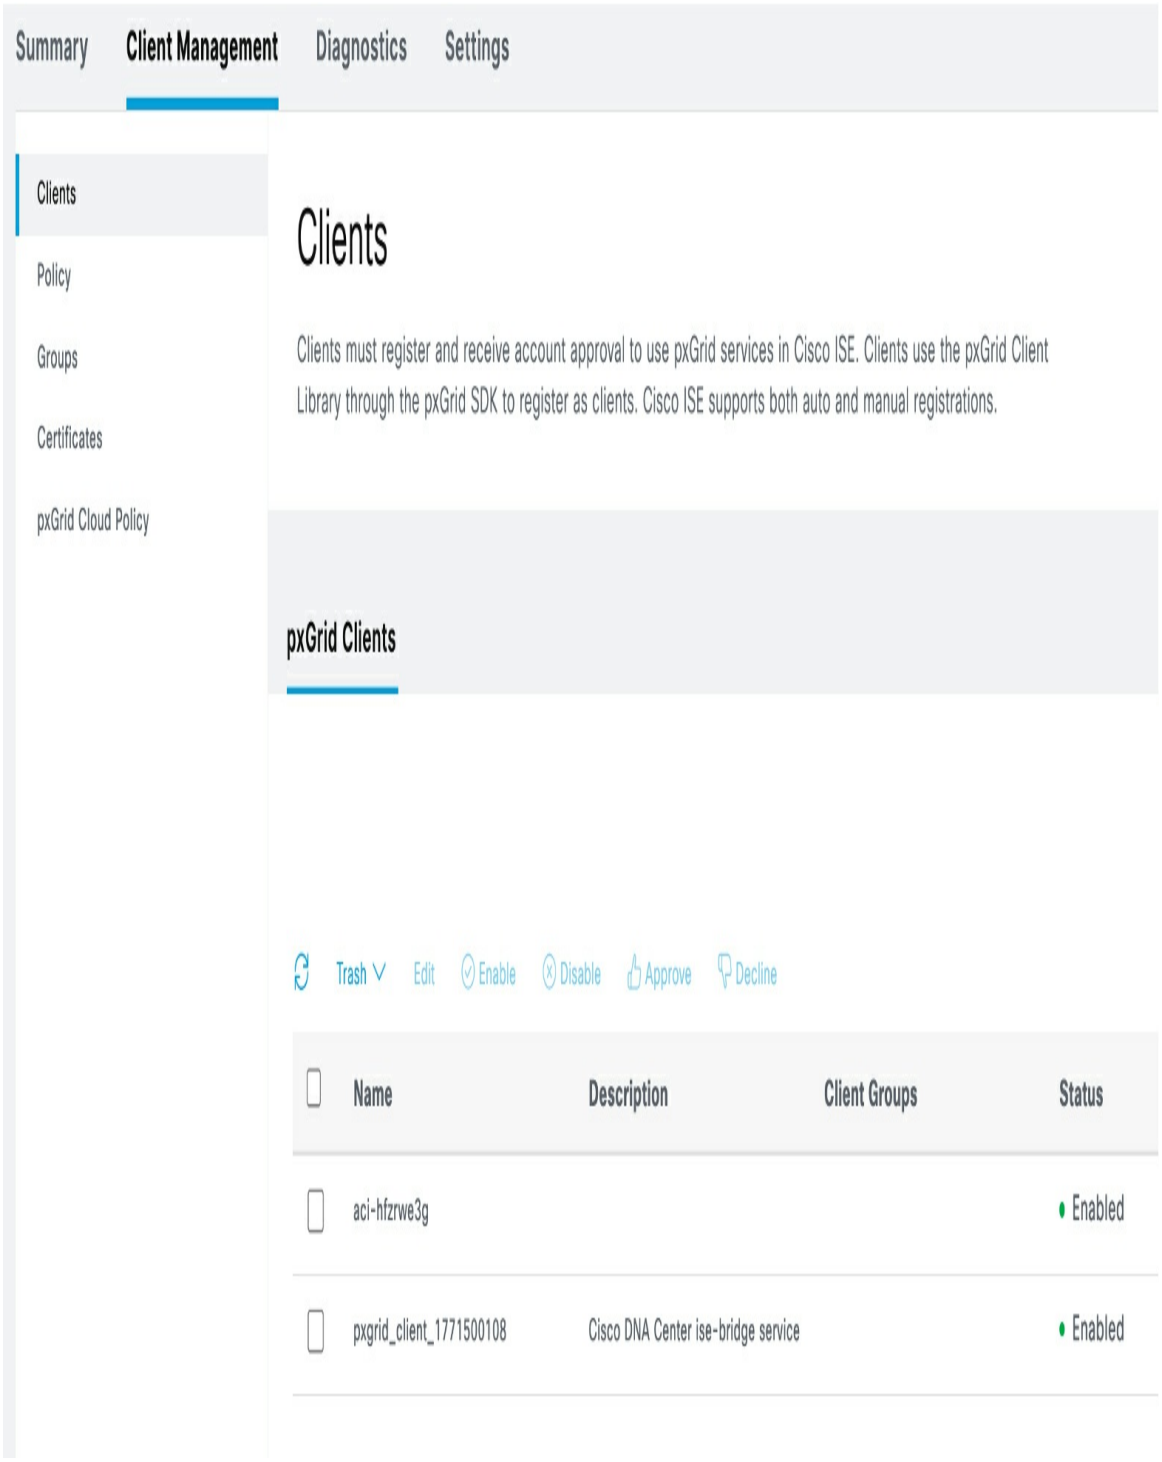


Figure 19-3 *pxGrid Setting to Automatically Approve Connection Requests*

If necessary, also enable **Allow Password Based Account Creation**.

After pxGrid is enabled, the summary of any client connectivity is available. You can find it by navigating to **Administration > pxGrid Services > Summary**. You can see connected pxGrid clients by navigating to **Administration > pxGrid Services > Client**

Management > Clients. Figure 19-4 shows approved connectivity to an Application Centric Infrastructure (ACI) controller and a Catalyst Center appliance.

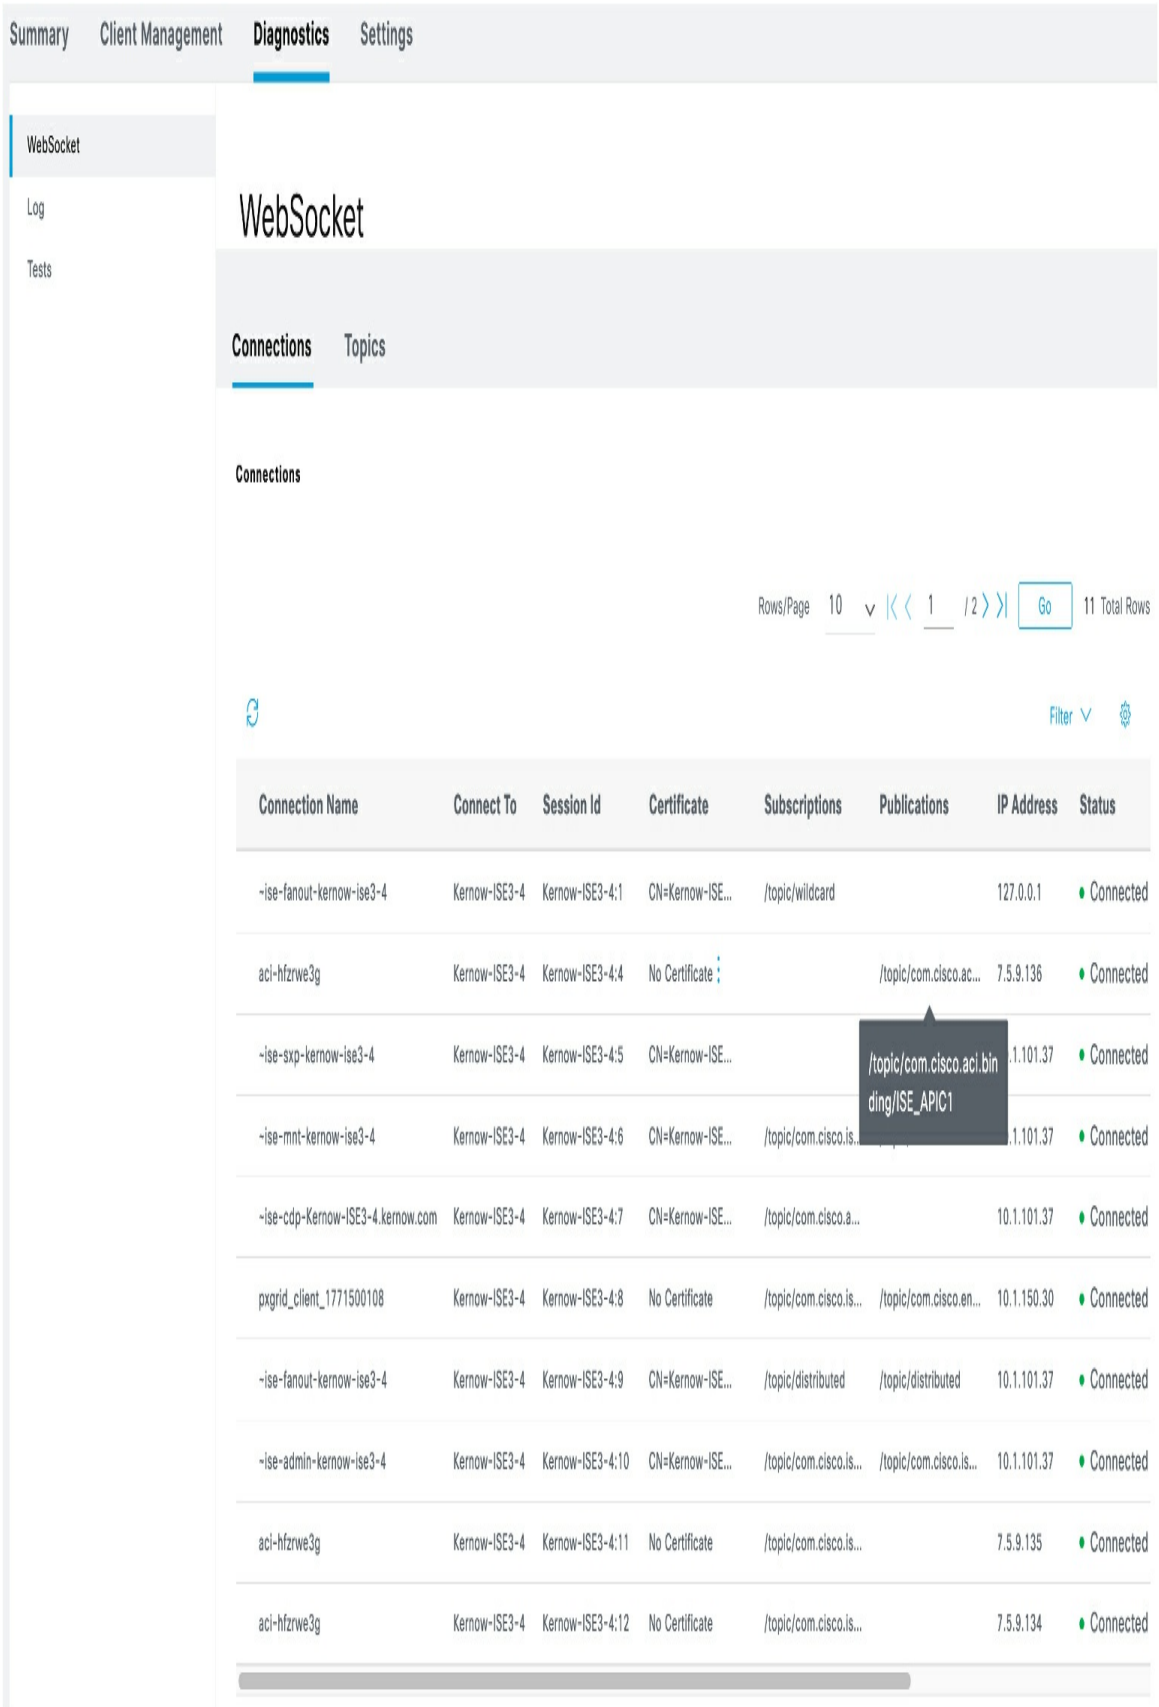


Technet24

Figure 19-4 *pxGrid Client Example*

When it comes to troubleshooting pxGrid connectivity, the **Diagnostics** menu is helpful. To use it, navigate to **Administration > pxGrid Services > Diagnostics**. The **WebSocket** submenu is useful for displaying the topic subscriptions and publications, and the **Log and Tests** submenus serve as useful supplementary tools.

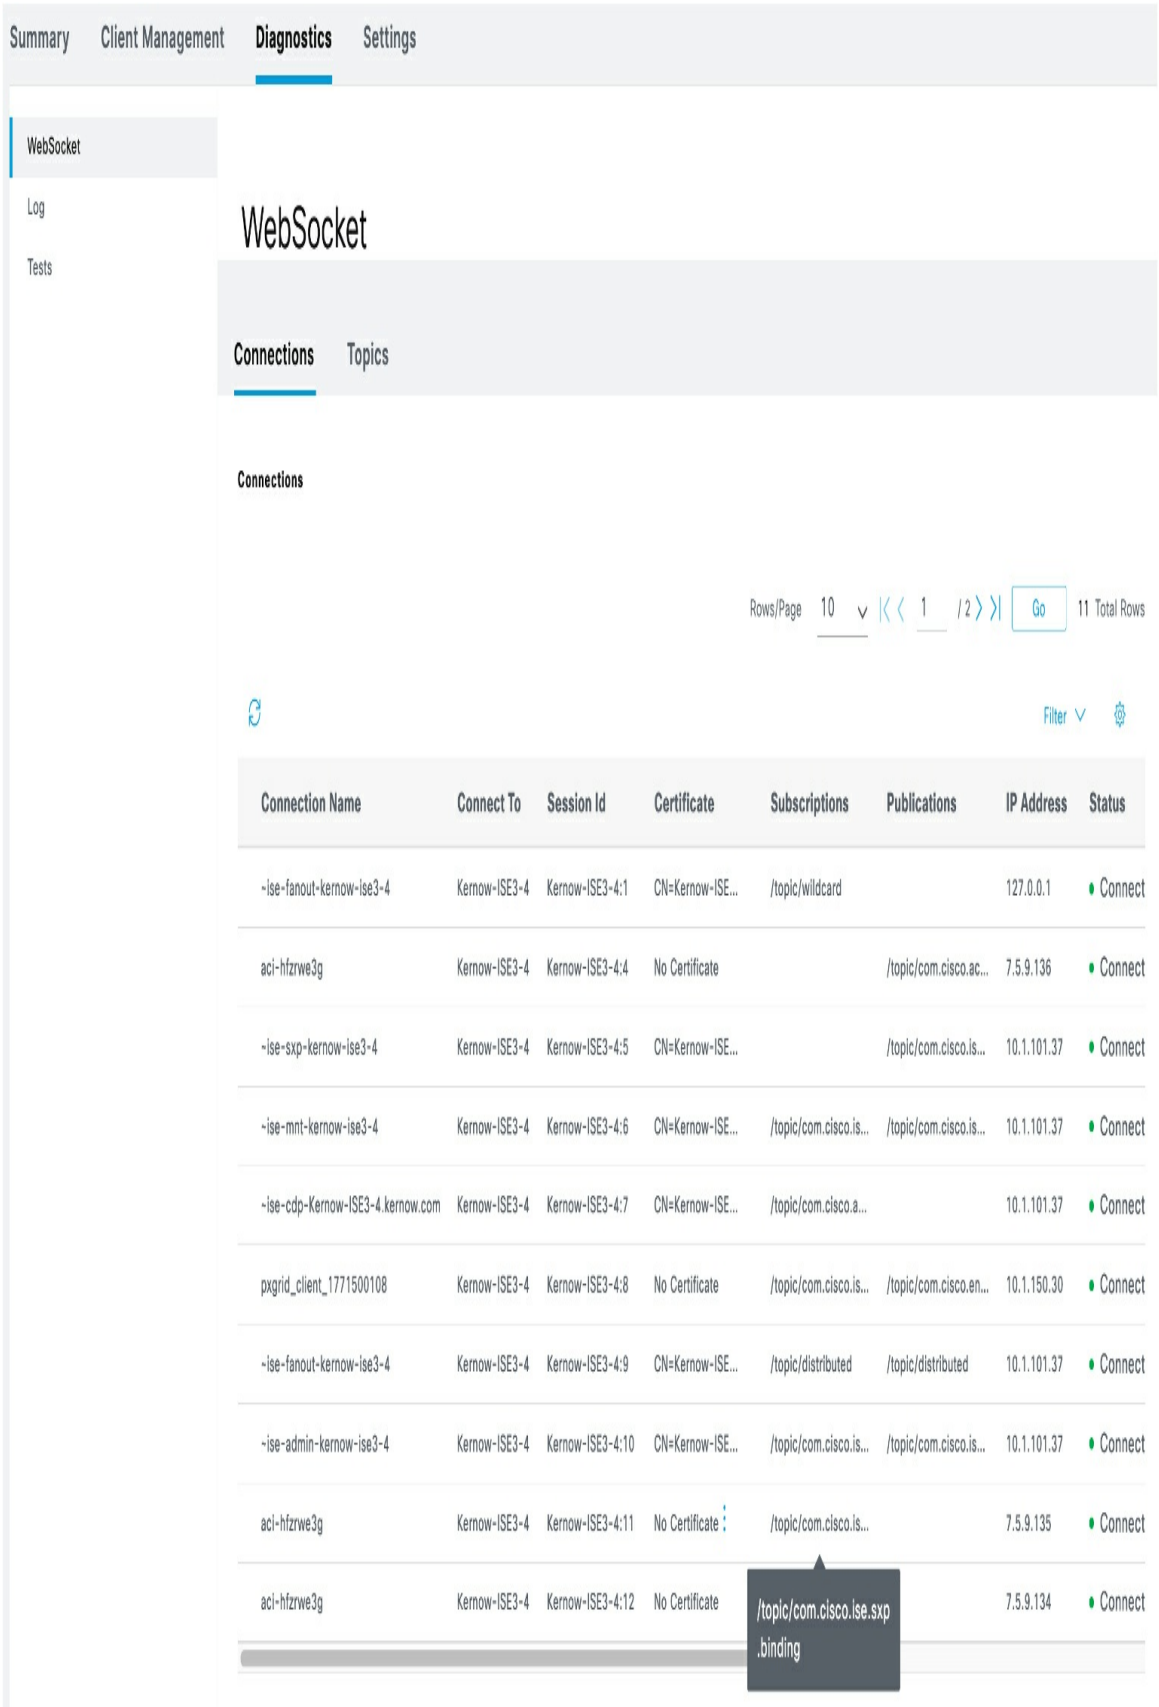
The **WebSocket** submenu can display pxGrid topics from a connection point of view or from a topic point of view. Figure 19-5 includes a sample snapshot. Here, we hovered the cursor over the publications for an ACI connection, and ISE displayed the topics used. The /topic/com.cisco.aci.binding/ISE_APIC1 topic is displayed for the ACI Application Policy Infrastructure Controller (APIC) to publish IP-to-group bindings in the ISE ACI connector named ISE_APIC1.


Summary
Client Management
Diagnostics
Settings
WebSocket
Log
Tests
WebSocket
Connections
Topics
Connections
Rows/Page 10 1 /2 Go 11 Total Rows
Filter
Connection Name
Connect To
Session Id
Certificate
Subscriptions
Publications
IP Address
Status
~ise-fanout-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:1 CN=Kernow-ISE... /topic/wildcard 127.0.0.1 Connected
aci-hfzrwe3g Kernow-ISE3-4 Kernow-ISE3-4:4 No Certificate /topic/com.cisco.ac... 7.5.9.136 Connected
~ise-sxp-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:5 CN=Kernow-ISE... .1.101.37 Connected
/topic/com.cisco.aci.binding/ISE_APIC1
~ise-mnt-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:6 CN=Kernow-ISE... /topic/com.cisco.is... .1.101.37 Connected
~ise-cdp-Kernow-ISE3-4.kernow.com Kernow-ISE3-4 Kernow-ISE3-4:7 CN=Kernow-ISE... /topic/com.cisco.a... 10.1.101.37 Connected
pxgrid_client_1771500108 Kernow-ISE3-4 Kernow-ISE3-4:8 No Certificate /topic/com.cisco.is... /topic/com.cisco.en... 10.1.150.30 Connected
~ise-fanout-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:9 CN=Kernow-ISE... /topic/distributed /topic/distributed 10.1.101.37 Connected
~ise-admin-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:10 CN=Kernow-ISE... /topic/com.cisco.is... /topic/com.cisco.is... 10.1.101.37 Connected
aci-hfzrwe3g Kernow-ISE3-4 Kernow-ISE3-4:11 No Certificate /topic/com.cisco.is... 7.5.9.135 Connected
aci-hfzrwe3g Kernow-ISE3-4 Kernow-ISE3-4:12 No Certificate /topic/com.cisco.is... 7.5.9.134 Connected

Technet24

Figure 19-5 *pxGrid WebSocket Connections Highlighting a Publication*

ACI APIC also has a pxGrid subscription, as shown in Figure 19-6. The topic used here is /topic/com.cisco.ise.sxp.binding. This topic allows APIC to learn the SXP bindings from ISE.


Summary
Client Management
Diagnostics
Settings
WebSocket
Log
Tests
WebSocket
Connections
Topics
Connections
Rows/Page 10 1 /2 Go 11 Total Rows
Filter
Connection Name
Connect To
Session Id
Certificate
Subscriptions
Publications
IP Address
Status
~ise-fanout-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:1 CN=Kernow-ISE... /topic/wildcard 127.0.0.1 Connect
aci-hfzrwe3g Kernow-ISE3-4 Kernow-ISE3-4:4 No Certificate /topic/com.cisco.ac... 7.5.9.136 Connect
~ise-sxp-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:5 CN=Kernow-ISE... /topic/com.cisco.is... 10.1.101.37 Connect
~ise-mnt-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:6 CN=Kernow-ISE... /topic/com.cisco.is... /topic/com.cisco.is... 10.1.101.37 Connect
~ise-cdp-Kernow-ISE3-4.kernow.com Kernow-ISE3-4 Kernow-ISE3-4:7 CN=Kernow-ISE... /topic/com.cisco.a... 10.1.101.37 Connect
pxgrid_client_1771500108 Kernow-ISE3-4 Kernow-ISE3-4:8 No Certificate /topic/com.cisco.is... /topic/com.cisco.en... 10.1.150.30 Connect
~ise-fanout-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:9 CN=Kernow-ISE... /topic/distributed /topic/distributed 10.1.101.37 Connect
~ise-admin-kernow-ise3-4 Kernow-ISE3-4 Kernow-ISE3-4:10 CN=Kernow-ISE... /topic/com.cisco.is... /topic/com.cisco.is... 10.1.101.37 Connect
aci-hfzrwe3g Kernow-ISE3-4 Kernow-ISE3-4:11 No Certificate /topic/com.cisco.is... 7.5.9.135 Connect
aci-hfzrwe3g Kernow-ISE3-4 Kernow-ISE3-4:12 No Certificate 7.5.9.134 Connect
/topic/com.cisco.ise.sxp.binding

Technet24

Figure 19-6 *pxGrid WebSocket Connections Highlighting a Subscription*

As you can see from Figure 19-2, there is an option in ISE to enable pxGrid Cloud. This option shares the ISE endpoint and contextual data with cloud-based solutions. Only the data that is relevant needs to be shared, providing a tailored consumption model.

When the checkbox is enabled, you can select the cloud region depending on the location of the ISE pxGrid node and application required.

Part of registering is logging in to your Cisco Catalyst Cloud Portal account (https://dna.cisco.com) to complete your device registration. When you're registered, the status of the connection should be displayed, as shown in Figure 19-7.

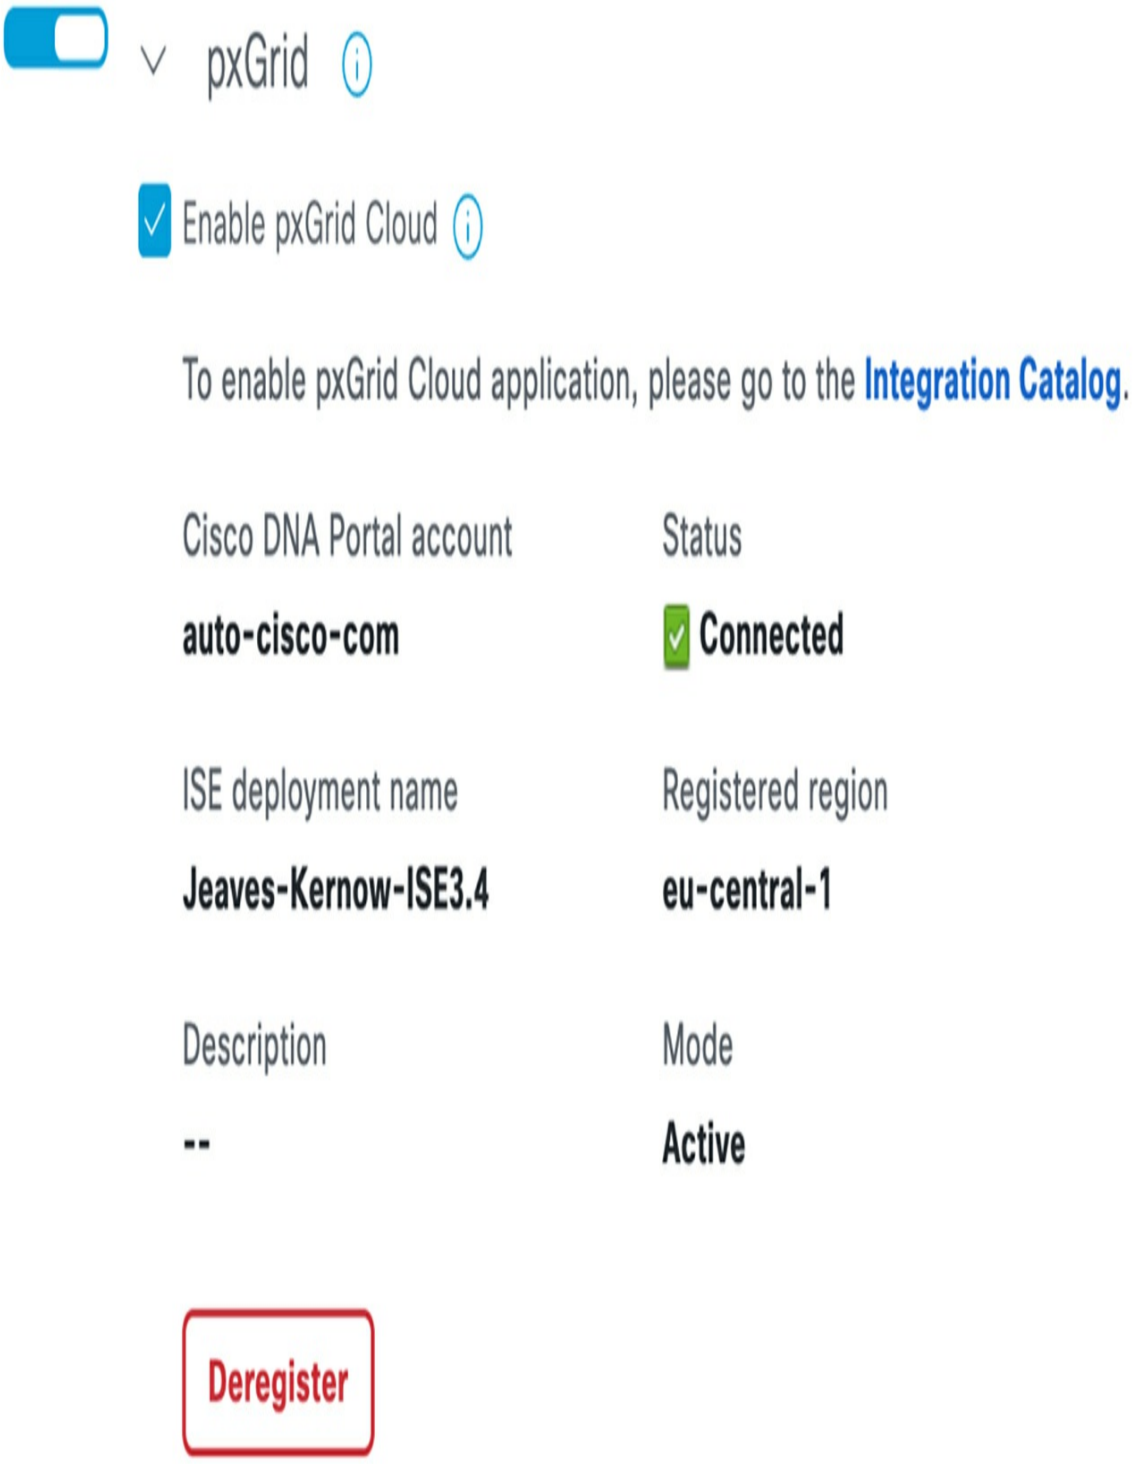


Figure 19-7 *pxGrid Cloud Status*

With the cloud connectivity successfully registered, applications can be

Technet24

integrated to make use of this context sharing. You can accomplish this via the ISE Integration Catalog by navigating to **Administration > Integration Catalog**.

You can choose the pxGrid Cloud application required within the **Available Integrations** panel. After you select it, a configuration page follows; it shows which ISE services will be shared with the application.

To summarize, pxGrid Cloud serves as a pivotal bridge that enables secure, scalable data exchange between on-premises Cisco ISE deployments and cloud-based security services. This architecture simplifies ecosystem integration by centralizing connectivity and eliminating the traditional infrastructure overhead typically required for cross-environment communication.

ISE to Catalyst Center pxGrid Integration

From Figures 19-4 to 19-6, the examples include a pxGrid connection to Catalyst Center (previously known as DNA Center). pxGrid is the mechanism used to exchange context between Catalyst Center and ISE. You can establish the integration from Catalyst Center under the **System > Settings > Authentication and Policy Servers** menu, but when it is established, Catalyst Center displays the Subscriber Name assigned. In Figure 19-8 the example shows pxgrid_client_1771500108.

Settings › External Services

Authentication and Policy Servers

Use this form to specify the servers that authenticate Catalyst Center users. Cisco Identit
Services Engine (ISE) servers can also supply policy and user information.

+ Add ∨ Export

IP Address	Protocol	Type
10.1.101.37	RADIUS	ISE
192.168.100.10	RADIUS	AAA
192.168.100.20	RADIUS	AAA

Subscriber Name displayed
= pxgrid_client_1771500108

Edit ISE server

X

Server IP Address

10.1.101.37

Shared Secret

Username*

admin

Password*

FQDN*

Kernow-ISE3-4.kernow.com

Subscriber Name

pxgrid_client_1771500108

Virtual IP Address(es) ∨

Info

Advanced Settings

Technet24

Figure 19-8 *Catalyst Center ISE Integration Settings Including Subscriber Name*

The Catalyst Center Subscriber Name is also displayed in ISE as a pxGrid client (refer to Figure 19-4).

When Catalyst Center and ISE are integrated, the data elements exchanged in both directions primarily involve policy and device information to enable group-based access control and network segmentation. The following key data elements are exchanged:

- **Certificates and Authentication:** Catalyst Center and ISE exchange certificates for secure pxGrid communication. Catalyst Center downloads and updates ISE certificates as needed and supports certificate revocation checks.
- **Polling and Trust Establishment:** Catalyst Center polls ISE every 15 minutes to maintain connectivity and trust. If ISE becomes unreachable, Catalyst Center increases polling frequency and may require reestablishing trust if connectivity is lost for an extended period.
- **Security Groups and Security Group Tags (SGTs):** ISE manages security groups identified by numeric SGTs. Catalyst Center synchronizes these groups with ISE, creating or updating security groups as needed to maintain consistency. SGTs are used for secure group-based policy enforcement in the network.
- **Group-Based Access Control Policies:** Catalyst Center and ISE synchronize group-based access control policies, including access contracts and group-based policies. Catalyst Center can migrate and synchronize policy data with ISE, with ISE's data taking precedence in case of conflicts.
- **Network Device Information:** Catalyst Center propagates

inventory data of network devices assigned to sites configured with ISE as the AAA server. This includes device hostname, AAA server configurations, device credentials, Loopback0 IP address, management IP address, and network device group (NDG) tags. Updates to these device settings in Catalyst Center are pushed to ISE.

- **Multiple Catalyst Center Clusters:** When multiple Catalyst Center clusters are integrated with a single ISE system, ISE acts as a central point to propagate virtual networks, SGTs, access contracts, and group-based access control policies among the clusters, designating one cluster as the Author node and others as Reader nodes.

In summary, the integration facilitates bidirectional exchange of certificates, security group definitions, group-based policies, network device inventory and configurations, and to enable dynamic, scalable, and consistent network access control and segmentation across the infrastructure. Cisco ISE provides runtime policy services, and Catalyst Center serves as the policy authoring and network management dashboard.

ISE to Cisco Secure Firewall pxGrid Integration

A popular consumer of pxGrid context is Cisco Secure Firewall. The integration is initiated from the Firewall Management Center (FMC), which allows a dynamic group-based security policy to be orchestrated for many Firewall Threat Defense (FTD) devices.

After certificates have been exchanged, in FMC you can make the integration with ISE by navigating to **Integration > Other Integrations > Identity Sources > Identity Services Engine**. Figure 19-9 shows the details, including the pxGrid topics that can be

Technet24

subscribed to.

Configure Identity Sources

Select the service type and start configuring the identity source. Deploy the changes after you're finished.

Service Type:

○ None ◉ Identity Services Engine ○ Passive Identity Agent

Quick Configuration (New) | Advanced Configuration (Old)

Primary Host Name/IP Address*

ISE3.4.kernow.com

Secondary Host Name/IP Address

pxGrid Client Certificate*

FMC-ISE

MNT Server CA*

ISE-root

pxGrid Server CA*

ISE-root

ISE Network Filter

ex. 10.89.31.0/24

Subscribe To:

☑ Session Directory Topic

Session Directory Topic enables you to receive ISE user session information from the ISE server.

☑ SXP Topic

SXP Topic enables you to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.

Test

Technet24

Figure 19-9 *Configuring ISE Integration in Firewall Management Center*

As the information shows in Figure 19-9, the integration can consume the pxGrid **Session Directory Topic**, which enables the learning of dynamic user session information, including any security group assigned, and/or the pxGrid **SXP Topic**, which enables the learning of IP-to-SGT mappings (otherwise known as bindings).

FMC handles the data from these two pxGrid topics slightly differently from a high availability point of view. If pxGrid is down, the Session Directory mappings are removed after 24 hours by default and configured at the Realm level. In contrast, the mappings learned via the SXP Topic will remain until explicitly deleted.

Figure 19-9 shows a Test facility is available after the information has been entered; this will test the connectivity with ISE and ensure the correct certificates are in place. Once the facility is selected, a pop-up window shows the result.

If the result is not successful, then the additional logs can be scrutinized to determine the cause.

Once successfully integrated, FMC will download the context from ISE as specified by the subscribed pxGrid topics. This context includes the security groups and security group memberships (IP-to-SGT mappings/bindings). You then can build dynamic security policies in FMC based on these learned groups, providing simplification and reducing scale.

To add an access control security policy in FMC, navigate to **Policies > Access Control > [defined policy] > Add Rule**. Select the **Dynamic Attributes** submenu, as shown in Figure 19-10, and check the box for **ISE Security Group Tags** in the **Filter by Type** column. The **Dynamic**

Attributes column is prepopulated with all the security groups learned from ISE via pxGrid. You can then select specific security groups and choose them to be used in the access control source criteria by subsequently clicking the **Add Source Dynamic Attribute** button. Additionally, you can select security and choose them to be used in the policy destination criteria by subsequently clicking the **Add Destination Dynamic Attribute** button.

Technet24

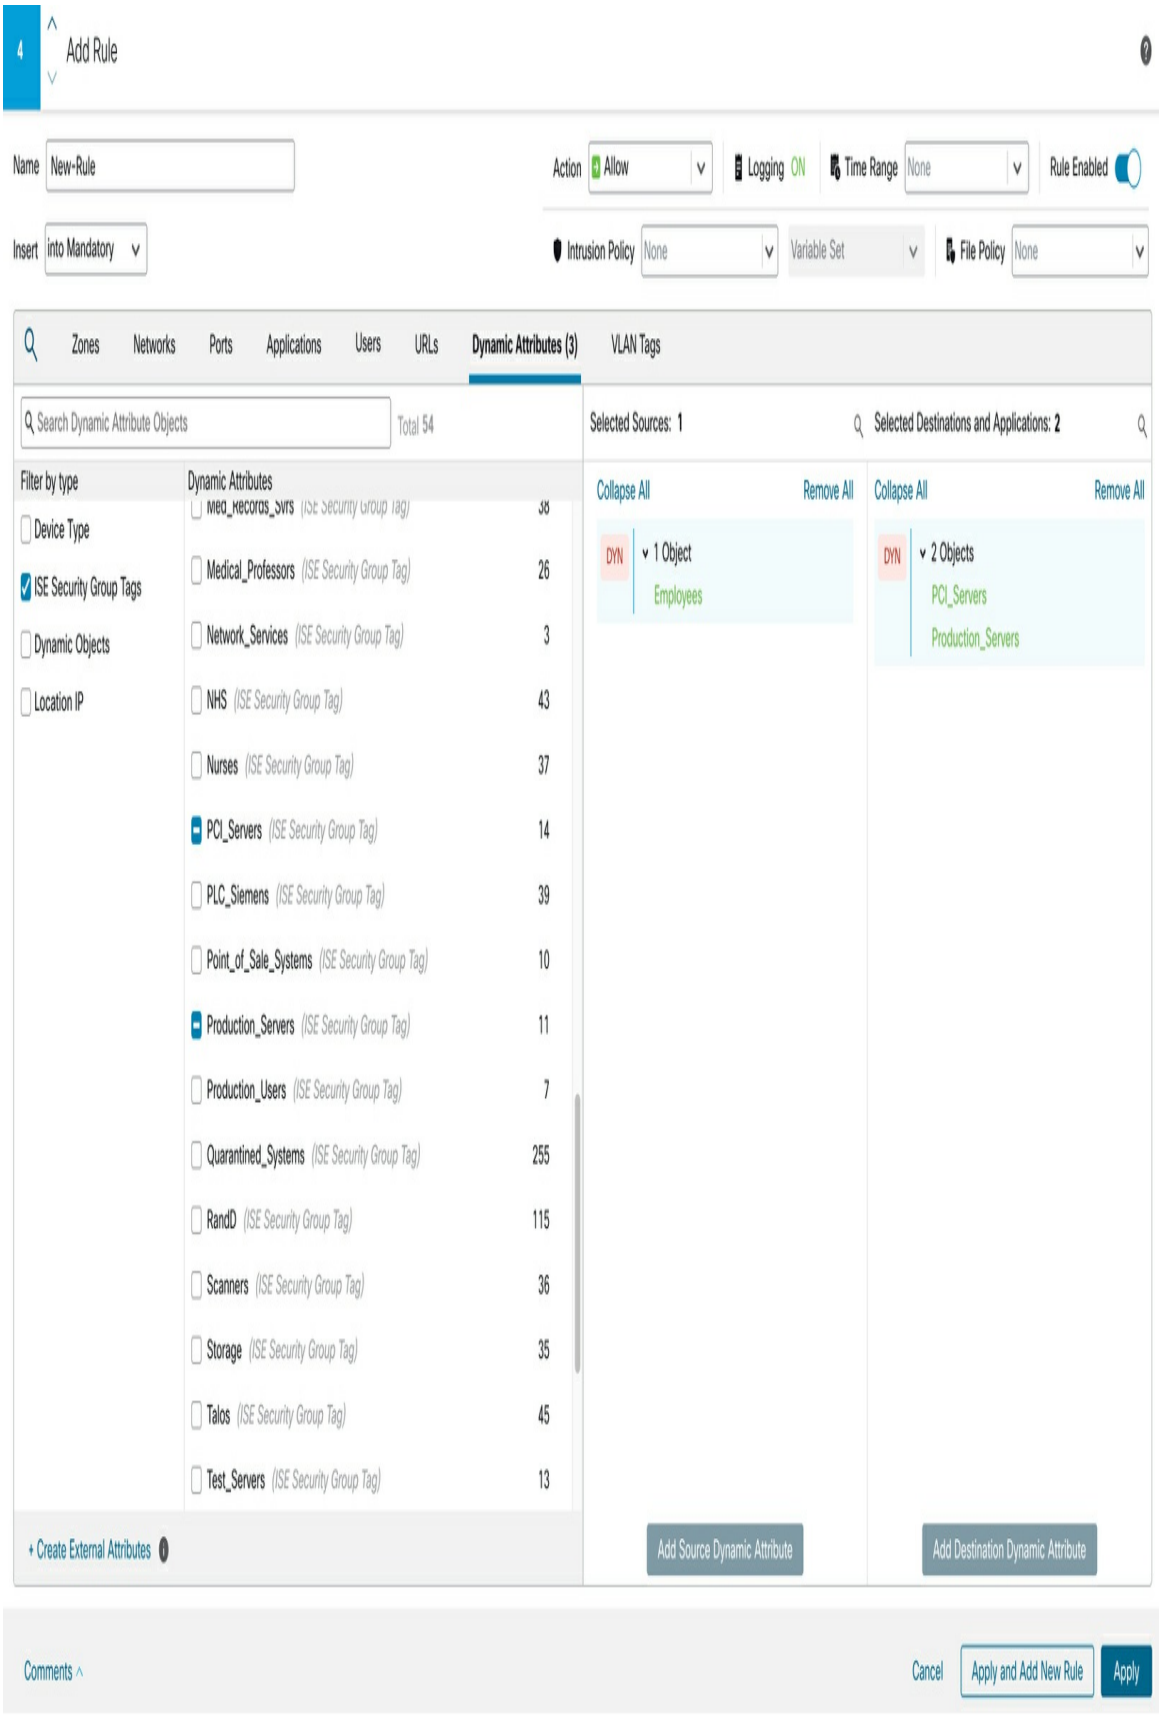

4
Add Rule
Name New-Rule
Action Allow
Logging ON
Time Range None
Rule Enabled
Insert into Mandatory
Intrusion Policy None
Variable Set
File Policy None
Zones
Networks
Ports
Applications
Users
URLs
Dynamic Attributes (3)
VLAN Tags
Search Dynamic Attribute Objects
Total 54
Filter by type
Device Type
ISE Security Group Tags
Dynamic Objects
Location IP
Dynamic Attributes
Medical_Professors (ISE Security Group Tag) 26
Network_Services (ISE Security Group Tag) 3
NHS (ISE Security Group Tag) 43
Nurses (ISE Security Group Tag) 37
PCI_Servers (ISE Security Group Tag) 14
PLC_Siemens (ISE Security Group Tag) 39
Point_of_Sale_Systems (ISE Security Group Tag) 10
Production_Servers (ISE Security Group Tag) 11
Production_Users (ISE Security Group Tag) 7
Quarantined_Systems (ISE Security Group Tag) 255
RandD (ISE Security Group Tag) 115
Scanners (ISE Security Group Tag) 36
Storage (ISE Security Group Tag) 35
Talos (ISE Security Group Tag) 45
Test_Servers (ISE Security Group Tag) 13
+ Create External Attributes
Selected Sources: 1
Collapse All
Remove All
DYN
1 Object
Employees
Add Source Dynamic Attribute
Selected Destinations and Applications: 2
Collapse All
Remove All
DYN
2 Objects
PCI_Servers
Production_Servers
Add Destination Dynamic Attribute
Comments
Cancel
Apply and Add New Rule
Apply

Figure 19-10 *FMC Access Control Rule Using ISE Security Groups*

Multiple source and destination security groups can be defined within each rule, but there is the option of only specifying a security group for either the source or the destination and mixing with other objects. This option provides ultimate flexibility covering a wide variety of use cases.

pxGrid Direct

Although native pxGrid integration is ideal, many legacy or third-party products do not support the pxGrid protocol natively. pxGrid Direct bridges this gap.

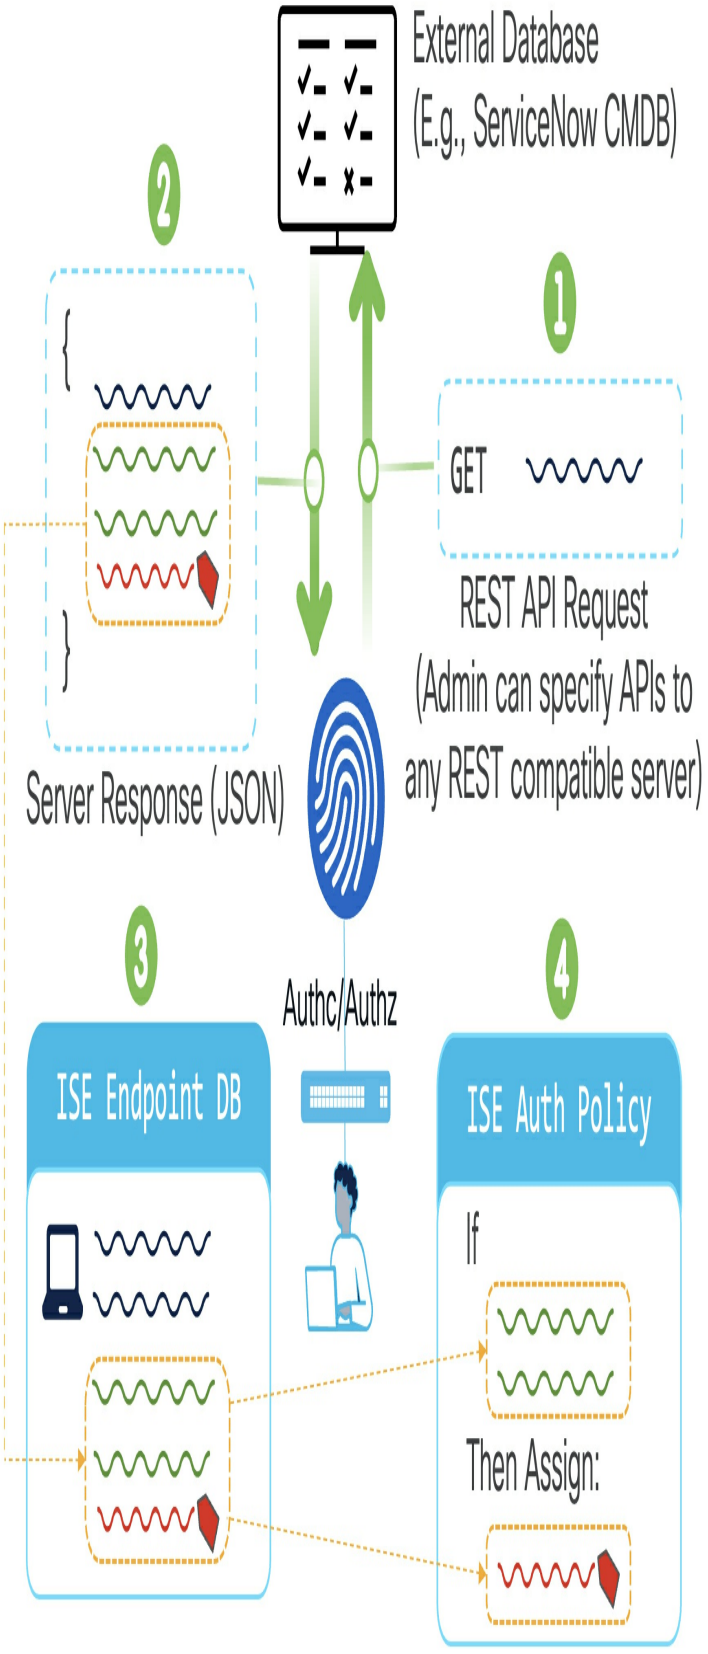
pxGrid Direct functions as an integration layer that allows ISE to interact with non-pxGrid-enabled systems and services via REST APIs. By leveraging pxGrid Direct, you can normalize data from these disparate sources and inject it into the pxGrid ecosystem, ensuring that even non-native products contribute to the overall security context. An example of a non-pxGrid–enabled service is ServiceNow. Using pxGrid Direct to integrate with services such as this allows you to bridge the gap between network-level identity context and IT service management, providing a unified operational view that triggers automated incident workflows, improves asset lifecycle accuracy, and enables faster, context-aware threat remediation across your entire infrastructure.

Integration is established via ISE pxGrid Direct connectors, and schedules can be implemented for ISE to periodically fetch data from an external server. Alternatively, a push option can be implemented for systems to push data directly into ISE if write permissions have been granted. The data format used for the exchange is JavaScript Object Notation (JSON) because it has a lightweight format ideal for

Technet24

transporting data. The pxGrid Direct operation is shown in Figure 19-11.

```
{"result":[{
"owned_by":"rtrevithick",
"operational_status":"Operational",
"sys_updated_on":"2026-04-22 16:44:52",
"sys_created_by":"rtrevithick",
"warranty_expiration":"",
"cpu_speed":"",
"checked_out":"",
"maintenance_schedule":"",
"group_tag":"cts:security-group-tag=12-00",
"managed_by":"rtrevithick",
"sys_class_name":"Locomotive",
"assigned_to":"rtrevithick@eaststraymine.org",
"install_date":"",
"asset_tag":"",
"subcategory":"Steam",
"sys_id":"0009ef70dbcf01101f3961997",
"mac_address":"00:50:56:A0:F7:3A",
"company":"East Stray Park Mine",
"department":"Engineering",
"ip_address":"10.6.5.129",
...
```

Technet24

Figure 19-11 *pxGrid Direct Operation*

The steps marked in Figure 19-11 are as follows:

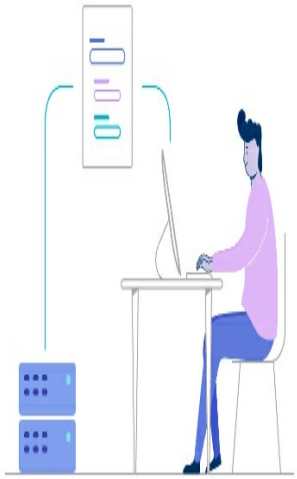
1. ISE makes a REST API GET request to a configured server.
2. The server responds with data records and associated attributes in the JSON format.
3. ISE stores this received data in an ISE Dictionary.
4. These attributes can be used in ISE authorization conditions as well as authorization results.

To administer pxGrid Direct connectors in ISE, navigate to **Administration > Network Resources > pxGrid Direct Connectors**. Click **Add** to initiate the workflow to add a new connector. Then click **Let's Do It**.

Connector definition comes next. This is where you select the connector type. Either the connector is a URL Fetcher type, whereby ISE will fetch data from the external server periodically, or a URL Pusher type, where systems can push raw JSON data into ISE. Figure 19-12 shows a URL Fetcher type selection with certificate validation.

✓ Welcome 2 Connector Definition 3 URL 4 Schedule 5 Parent Object 6 Attributes 7 Identifiers 8 Summary

Define Connector



A new Cisco ISE Dictionary is created at a later step to store the attributes and values that Cisco ISE retrieves from the pxGrid Direct connector. The name that you provide here is the name of the Cisco ISE Dictionary folder for the connector. Select connector type: whether URL Fetcher or URL Pusher.

Name

Cisco_CMDB

Description (optional)

Connector type:

- (•) URL Fetcher
  Use to schedule Cisco ISE to fetch data from an external server.
- ( ) URL Pusher
  Use to push raw JSON data into Cisco ISE without an external server.
  You will need write permissions to push the JSON data.

Certificate Validations

To skip certificate validation, uncheck the following check box. Skip validation only in test environments or if you trust the connected server to be highly secure. Skipping validation could make your network vulnerable to machine-in-the-middle attacks because Cisco ISE accepts server certificates without verifying hostnames or other details.

- [x] Validate certificate

Exit Wizard

Back Next

Technet24

Figure 19-12 *pxGrid Direct Connector Definition*

After you define the connector, click **Next**. The following part of the workflow allows the location of the endpoint data to be entered in the form of a URL, along with any authentication credentials. Click the **Test Connection** option for ISE to validate connectivity to the entered URL. ISE will display a **Successful** message. Click **Next**. Figure 19-13 shows a URL Fetcher type being set up; therefore, defining the schedule for data sync comes next.



Figure 19-13 *pxGrid Direct Fetch Schedule*

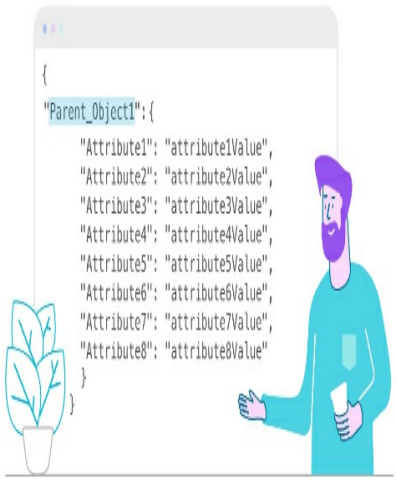
Click **Next**. For ISE to fetch the correct data, you need to enter the

Technet24

JSON parent data object. Click **Validate** for ISE to use that parent object to fetch the first data record and display that record along with the **Successful** message if this succeeds, as shown in Figure 19-14. You should check this is the record you were expecting; otherwise re-enter, the correct parent object.

Welcome ✓ | Connector Definition ✓ | URL ✓ | Schedule ✓ | 5 Parent Object | 6 Attributes | 7 Identifiers | 8 Summary

Parent Object



Configure the JSON data object that must be used to search for the rest of the attributes. You can read more information and example at Page Level Help .

Parent Object

result

Validate ✓ Successful


```
{
  "asset": "Unknown",
  "asset_tag": "",
  "assigned": "",
  "assigned_to": "Joseph.Gibson@example.org",
  "assignment_group": "",
  "attestation_score": "",
  "attested": "false",
  "attested_by": "",
  "attested_date": "",
  "attributes": "",
  "can_print": "false",
  "category": "Hardware",
  "cd_rom": "false",
  "cd_speed": "",
  "change_control": "",
  "chassis_type": null,
  "checked_in": "",
  "checked_out": "",
```


Exit Wizard

Back Next

Technet24

Figure 19-14 *pxGrid Direct Server JSON Parent Object*

Click **Next**. With the first record downloaded within the previous validation step, ISE has collected all the fields, or attributes, within that record. This next step allows you to select which of these attributes ISE will be storing in an internal dictionary and hence available to use when defining authorization rules. Each attribute is displayed with a slider, as shown in Figure 19-15, to include in the dictionary, plus whether the attribute name from the server is to be renamed in the dictionary. Additionally, Change of Authorization (CoA) can be enabled if any changes in the fetched attributes are desired to influence the live sessions generated from the attributes in authorization.



Select Attributes Configure Dictionary Items

Change of Authorization (CoA) can be enabled from pxGrid Direct connections for dictionary attributes used in authorization policies. When the value of a CoA-enabled dictionary attribute changes, a CoA port bounce or reauthentication occurs on the impacted endpoint.

Note: Enabling CoA can result in a change in authorization on endpoints associated with existing live sessions when data is fetched or pushed into ISE.

Change of Authorization (CoA)

Reauthentication

When the value of the attribute changes, the CoA reauthentication will force a client device to re-authenticate to a network and assign a new policy, if applicable.

Port Bounce

When the value of the attribute changes, the CoA bounce host port command terminates a session and bounces the port, initiating a new session for authorization on the bounced port.

Attributes List

Add the attributes that Cisco ISE must retrieve from the pxGrid Direct connector. Choose attributes that should be included to the Cisco ISE Dictionary by clicking the toggle switch next to an attribute. Enter the attribute name that you want displayed in the Cisco ISE Dictionary. All the attributes that are retrieved from the pxGrid Direct connector persist in Cisco ISE even if they are not included in the Cisco ISE Dictionary.

Add Attribute Delete Include all in Dictionary Enable CoA for all

External Name	Include in Dictionary	Name in Dictionary	Enable CoA
$.asset	on	asset	off
$.asset_tag	off	asset_tag	off
$.assigned	on	assigned	off
$.assigned_to	on	assigned_to	on
$.assignment_group	off	assignment_group	off
$.attestation_score	off	attestation_score	off
$.attested	off	attested	off

Exit Wizard Back Next

Technet24

Figure 19-15 *pxGrid Direct Attribute Selection*

After selecting the attributes, click **Next**. At this point, ISE needs attribute identifiers to identify an attribute that is unique to an endpoint, an attribute that allows matching to an ISE authorization rule, and an attribute that helps to record the version of the endpoint data.

After adding these identifiers, click **Next**. Then click **Done** if they are correct. The newly added connector is then displayed in a table of connectors.

Rather than waiting for the next automatic scheduled sync, you can initiate a full sync from the table entry. To do so, select **Full** in the **Sync Now** column in the table for the connector in question. After you click **Full**, a pop-up window appears asking whether a full sync is required. Click **Sync Now**. The Connector table entry initially updates the Sync Status as being queued and then updates it again with **Completed** when finished. The number of objects downloaded and the Last Sync date and time are also displayed, as shown in Figure 19-16.

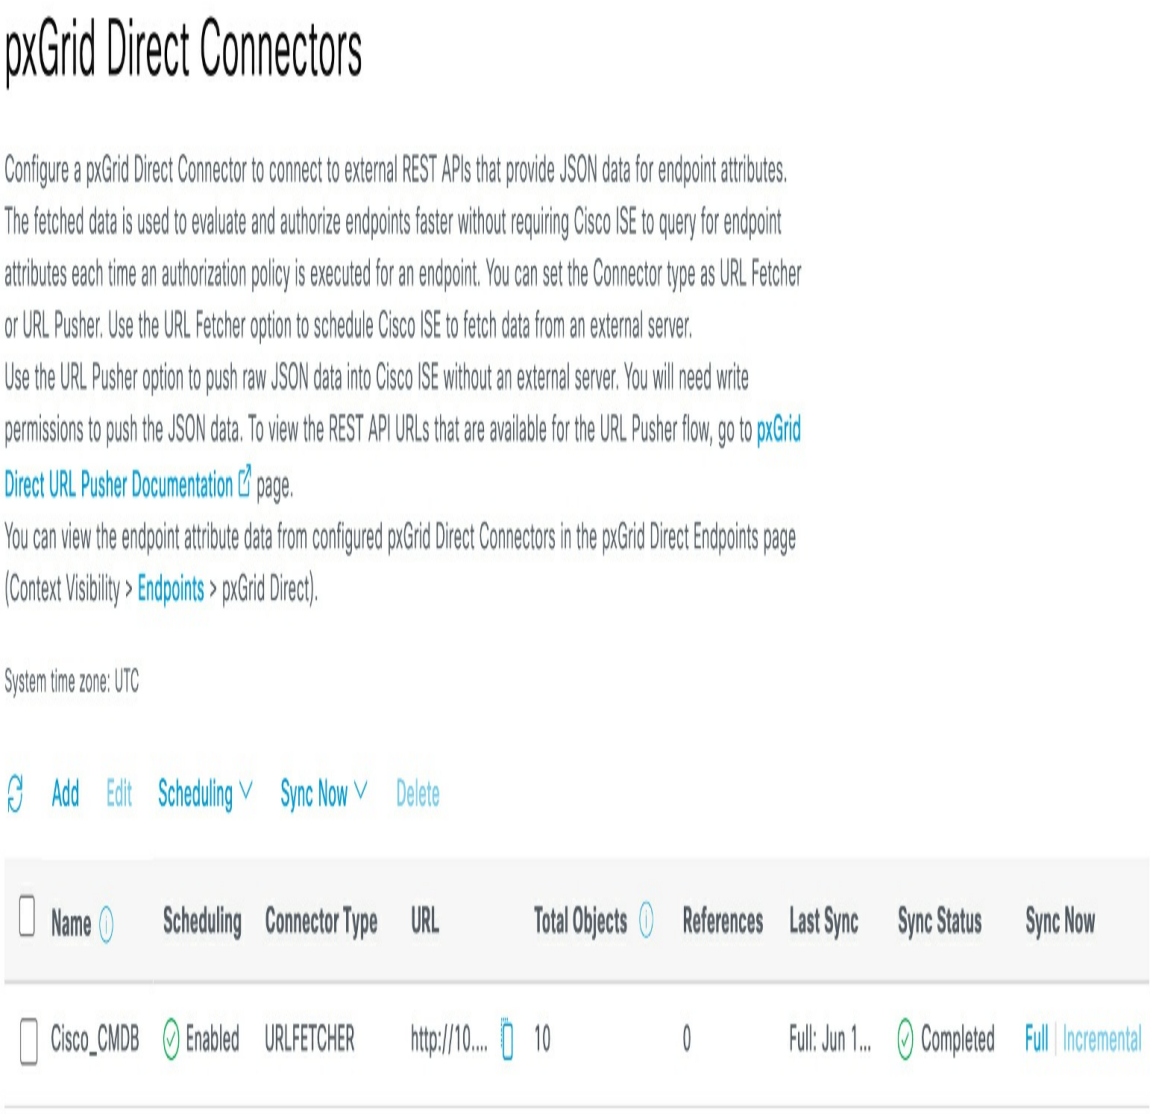


Figure 19-16 *pxGrid Direct Connector Table After Full Sync*

The downloaded attributes are stored in an ISE dictionary. To view them in ISE, navigate to **Policy > Policy Elements > Dictionaries**. Under the **System** folder, find the dictionary with the same name as the pxGrid Direct connector created.

You can select the dictionary and can display the **Dictionary Attributes** that were selected when the connector was added.

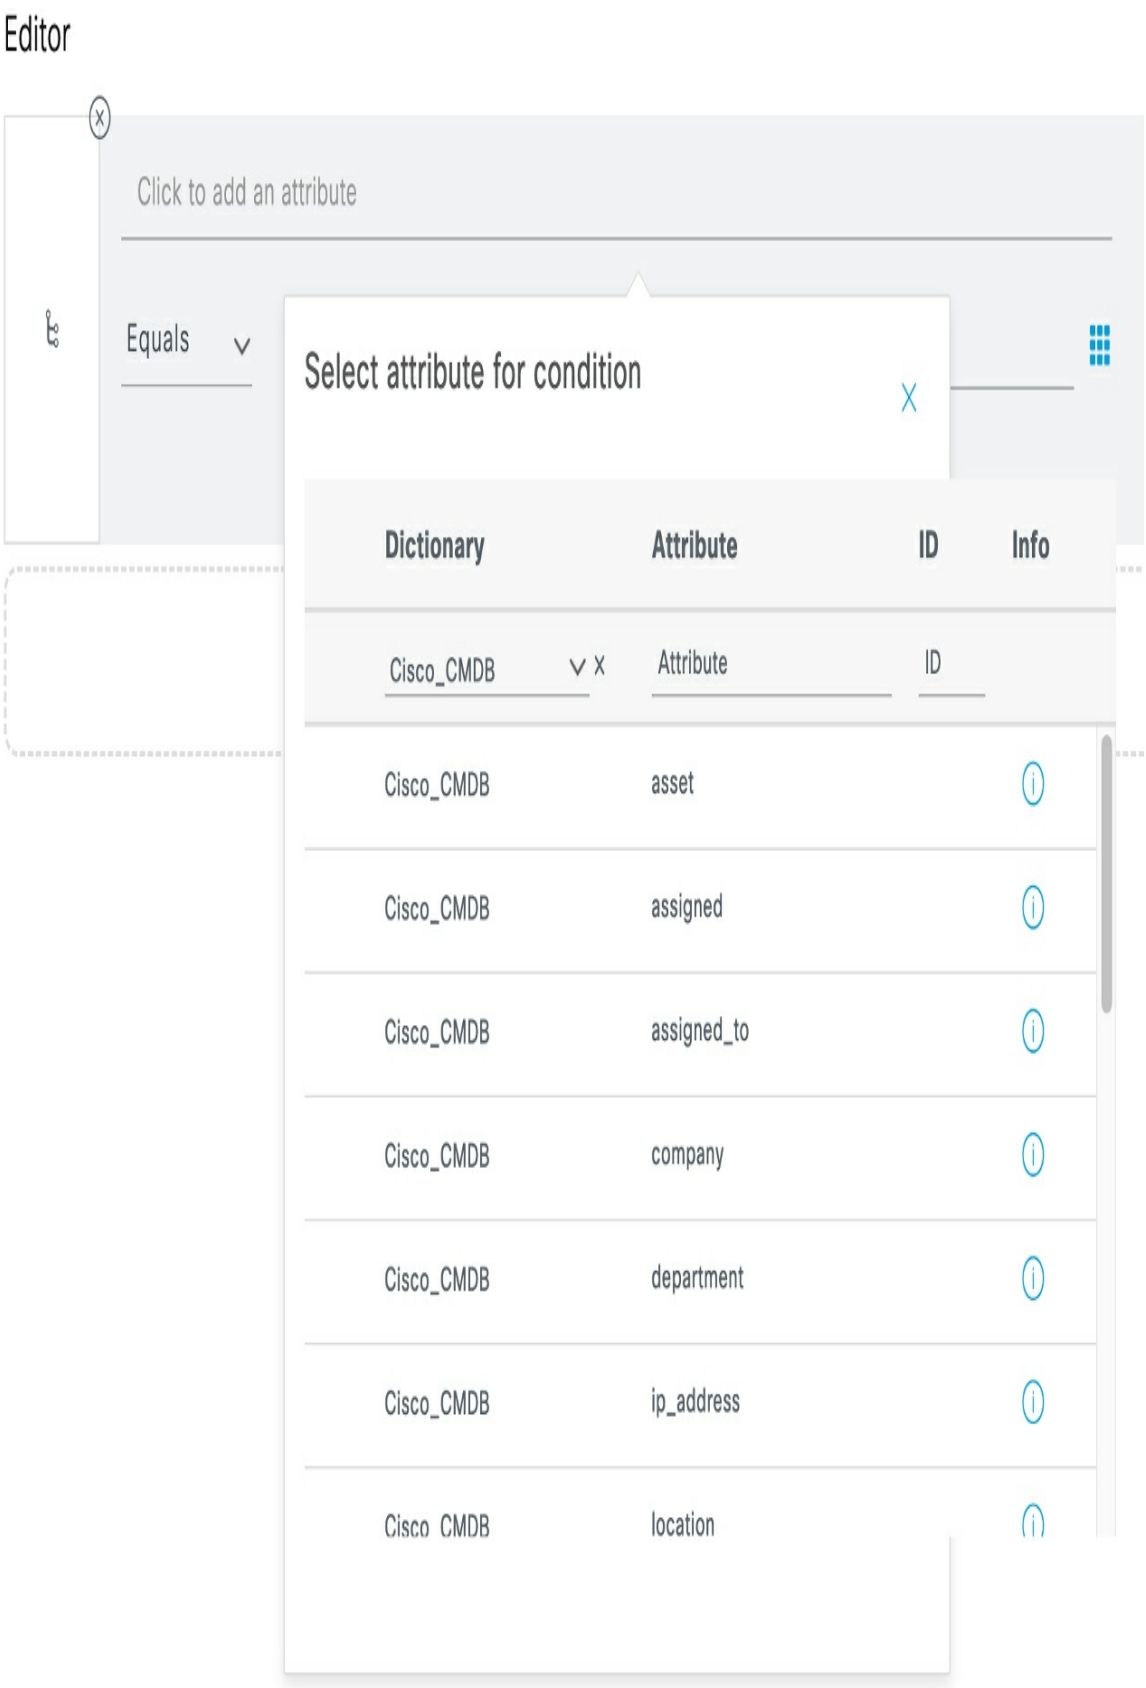
Now that these attributes have been downloaded to a dictionary, you

Technet24

can make use of them. You can use them in ISE authorization rules and authorization profiles. They can be used in matching conditions for rule selection, but they can also be used in authorizing results.

When you add authorization conditions, a dictionary selection is available. One of these is the dictionary you created when adding the pxGrid Direct connector.

When you select the pxGrid Direct dictionary, ISE displays the attributes chosen when the pxGrid Direct connector was added and subsequently stored in the dictionary, as shown in Figure 19-17.


Editor
Click to add an attribute
Equals
Select attribute for condition
Dictionary
Attribute
ID
Info
Cisco_CMDB
Attribute
ID
Cisco_CMDB asset
Cisco_CMDB assigned
Cisco_CMDB assigned_to
Cisco_CMDB company
Cisco_CMDB department
Cisco_CMDB ip_address
Cisco_CMDB location

Technet24

Figure 19-17 *pxGrid Direct Dictionary Attributes Available in Authorization Conditions*

Any of these attributes can be selected to be used in the conditions. Figure 19-18 demonstrates a match if the "company" contains a specific company name you want to match on.

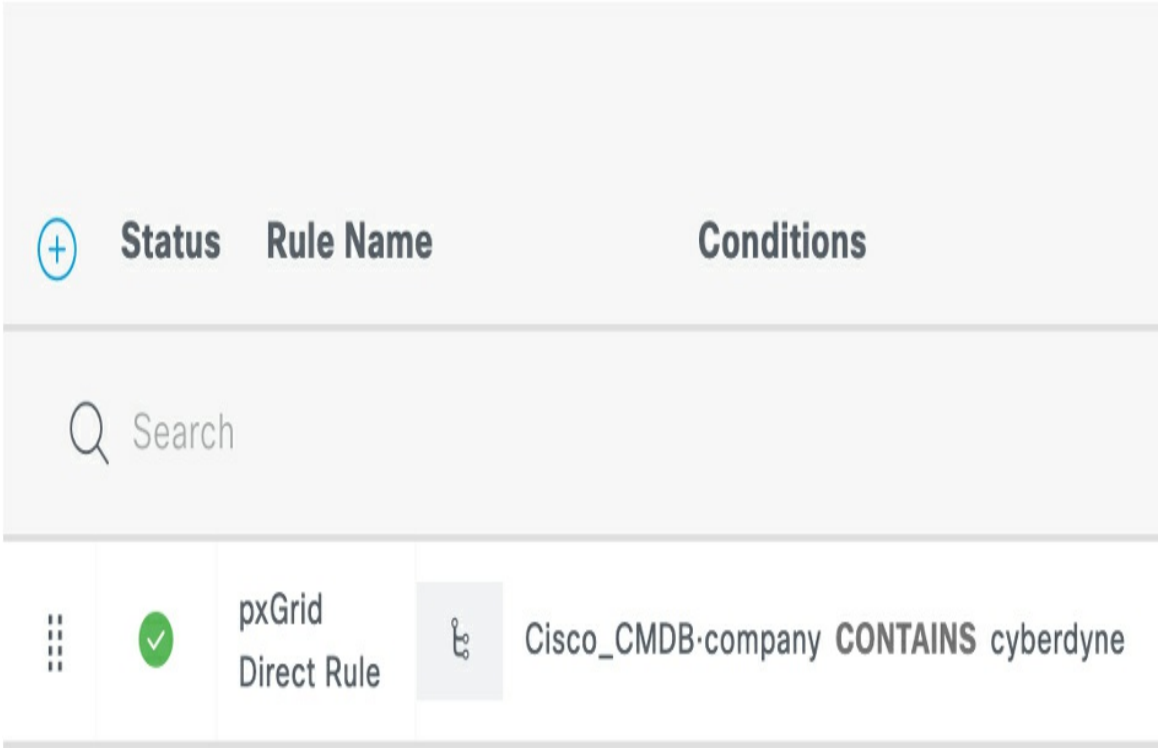


Figure 19-18 *Condition Matching a Company Name*

pxGrid Direct attributes can also be used in authorization results. As an example, if an attribute was used to define the security group of the endpoint, then ISE can make use of that in authorization profiles to assign that security group in the authorization phase of a AAA session. In the following examples, the attribute used is called u_segmentation_group_tag, and the format of that attribute is such that ISE can send it unaltered as a RADIUS AV pair so the network device can understand it as an SGT and accept it.

In ISE, navigate to **Policy > Policy Elements > Results >**

Authorization > Authorization Profiles. Add a new authorization profile. Scroll down to **Advanced Attributes Settings**. Open the **Select an Item** drop-down box and select **Cisco > cisco-av-pair**. Then, after dropping down the second box, you can select the pxGrid Direct dictionary.

Selecting the pxGrid Direct dictionary enables you to select an associated attribute, as shown in Figure 19-19; it is u_segmentation_group_tag in this instance.

Technet24

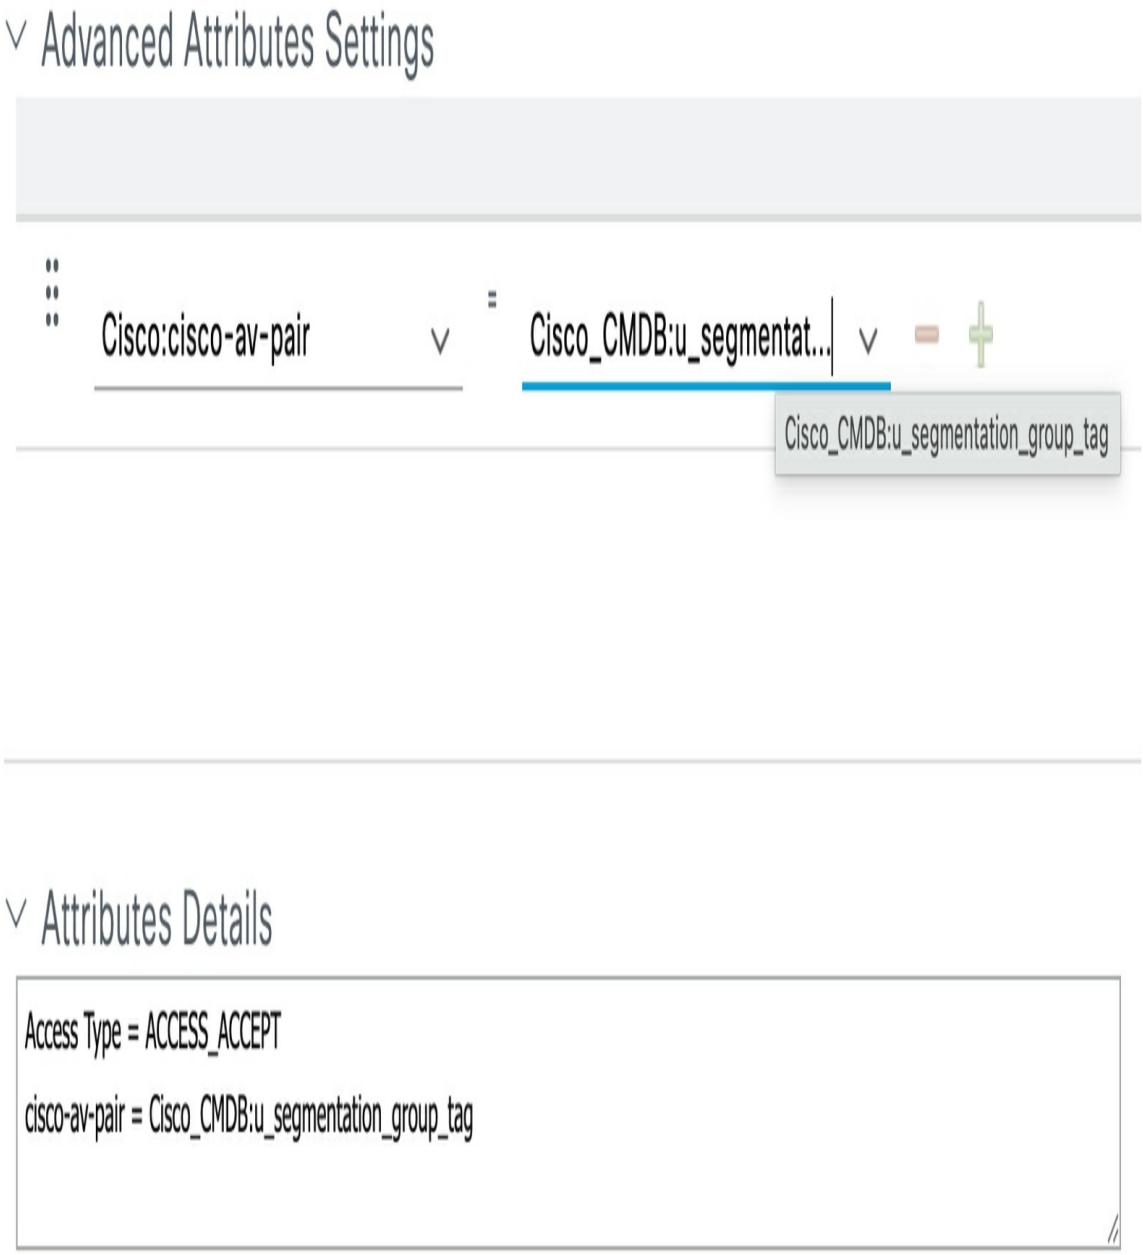


Figure 19-19 *Using pxGrid Direct Attribute in the Authorization Profile*

When the profile is saved, that is all that is necessary to make use of the attributes from the remote server. In the example provided, when an endpoint from company Cyberdyne connects to the network, the

associated authorization rule is matched using the company attribute, and an SGT is assigned via the authorization profile using the u_segmentation_group_tag attribute.

ISE Context Visibility is a great facility to be able to see endpoints and attributes learned from pxGrid Direct connected servers. In ISE, navigate to **Context Visibility > Endpoints > pxGrid Direct**. Figure 19-20 shows a drop-down box at the top of the UI where you can select a particular connector.

Technet24

pxGrid Direct Endpoints

Connector : Cisco_CM...

Rows/Page 10 | 1 /1

This window lists the endpoint attribute data that is fetched from pxGrid Direct Connectors. Click the correlation ID for an endpoint to view or download the endpoint details.
To create a new pxGrid Director Connector or to update existing connector configurations, go to the pxGrid Direct Connectors window.

Correlation ID	Unique ID	Version ID	Connector Name
Correlation ID	Unique ID	Version ID	Connector Name
00:50:56:A0:F7:3A	00:50:56:A0:F7:3A	2022-04-21 16:23:34	Cisco_CMDB
0B:16:01:C4:C3:CE	0b:16:01:c4:c3:ce	2022-04-21 15:26:30	Cisco_CMDB
16:F5:F2:4D:10:7B	16:f5:f2:4d:10:7b	2022-04-21 19:01:56	Cisco_CMDB
2C:04:E3:24:A2:AE	2c:04:e3:24:a2:ae	2022-04-21 16:24:40	Cisco_CMDB
7D:1C:CF:A0:2E:AD	7d:1c:cf:a0:2e:ad	2022-04-21 18:44:12	Cisco_CMDB
AA:D8:60:F9:74:C8	aa:d8:60:f9:74:c8	2022-04-21 18:25:37	Cisco_CMDB
AF:A8:A1:C6:B6:DD	af:a8:a1:c6:b6:dd	2022-04-21 19:09:36	Cisco_CMDB
EF:99:47:5B:C9:52	ef:99:47:5b:c9:52	2022-04-21 16:06:22	Cisco_CMDB
F5:0B:0A:35:68:2F	f5:0b:0a:35:68:2f	2022-04-21 17:55:02	Cisco_CMDB
FE:90:F0:F7:E6:BD	fe:90:f0:f7:e6:bd	2022-04-22 16:44:52	Cisco_CMDB

Figure 19-20 *pxGrid Direct Endpoints in ISE Context Visibility*

Each correlation ID, or MAC address in this case, can be selected to show attributes within that database record, as you can see by selecting the first entry shown in Figure 19-21.

Technet24

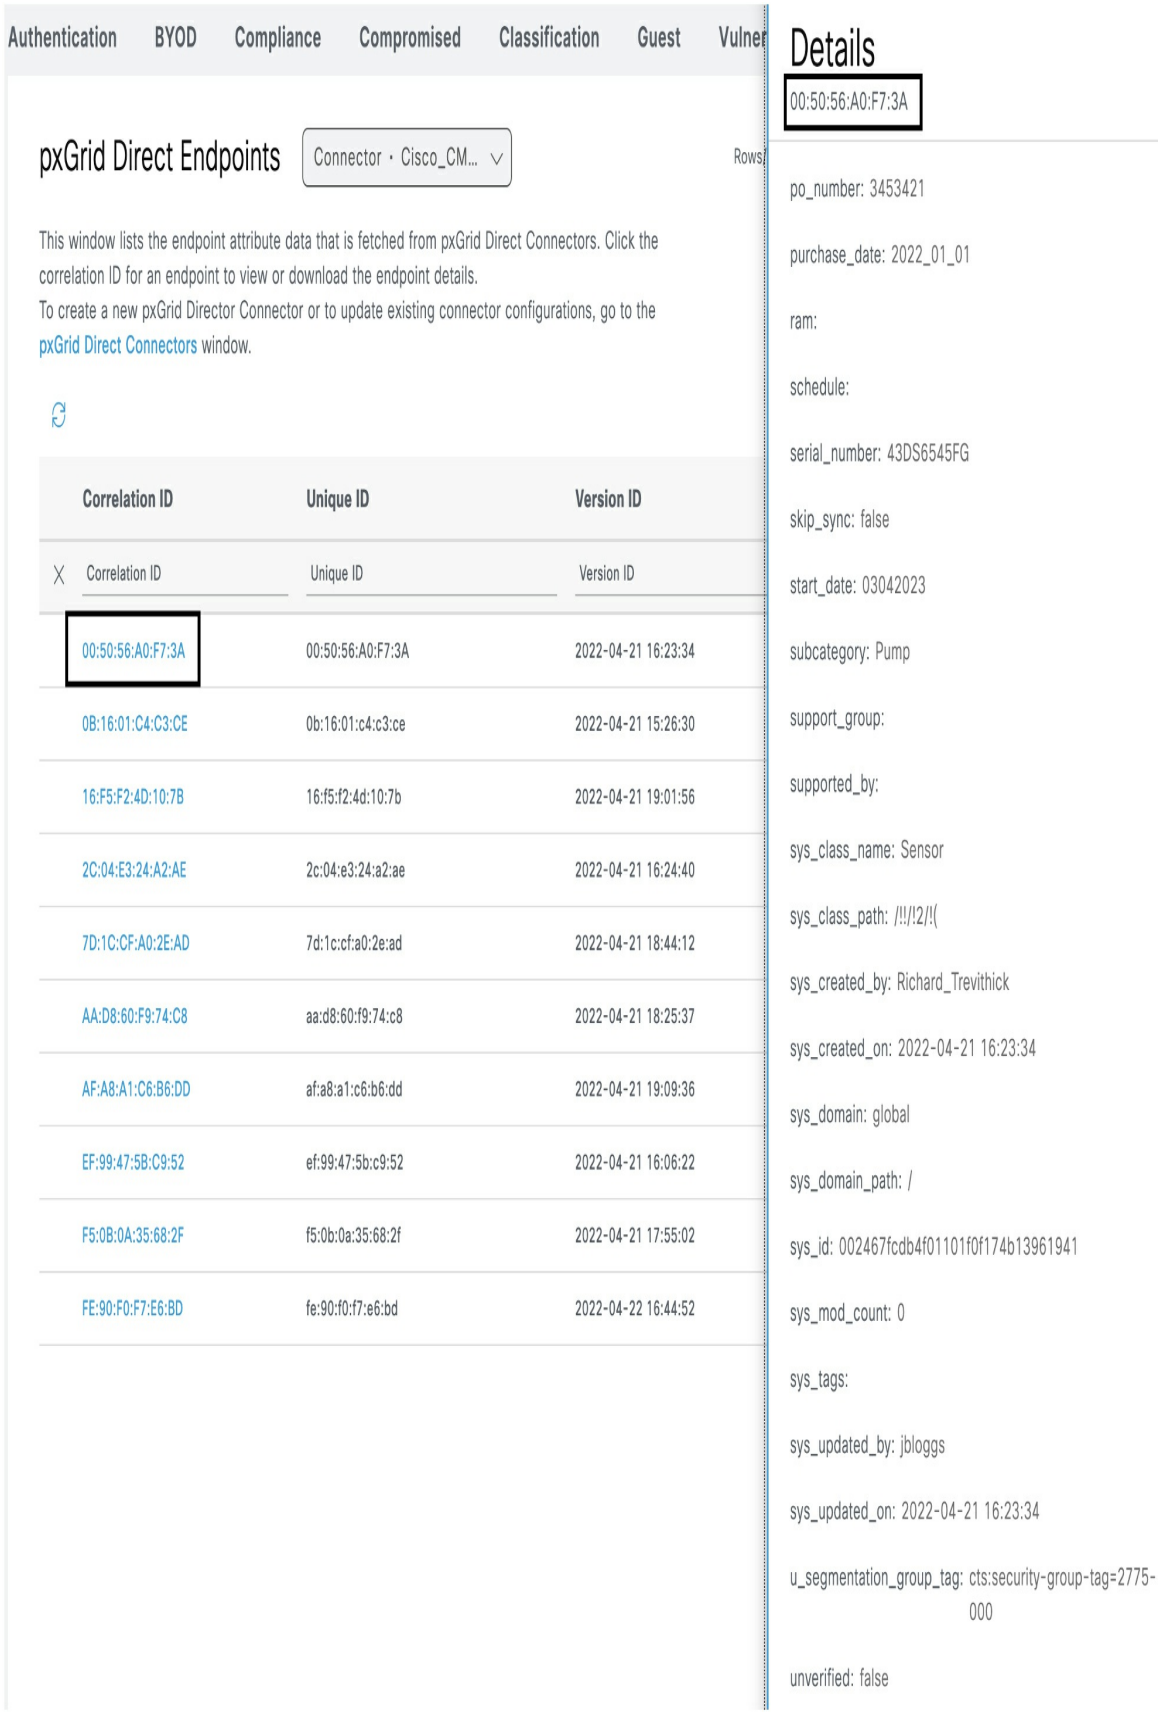

Authentication
BYOD
Compliance
Compromised
Classification
Guest
Vulner
pxGrid Direct Endpoints
Connector · Cisco_CM...
Rows
This window lists the endpoint attribute data that is fetched from pxGrid Direct Connectors. Click the correlation ID for an endpoint to view or download the endpoint details.
To create a new pxGrid Director Connector or to update existing connector configurations, go to the pxGrid Direct Connectors window.
Correlation ID
Unique ID
Version ID
Correlation ID
Unique ID
Version ID
00:50:56:A0:F7:3A
00:50:56:A0:F7:3A
2022-04-21 16:23:34
0B:16:01:C4:C3:CE
0b:16:01:c4:c3:ce
2022-04-21 15:26:30
16:F5:F2:4D:10:7B
16:f5:f2:4d:10:7b
2022-04-21 19:01:56
2C:04:E3:24:A2:AE
2c:04:e3:24:a2:ae
2022-04-21 16:24:40
7D:1C:CF:A0:2E:AD
7d:1c:cf:a0:2e:ad
2022-04-21 18:44:12
AA:D8:60:F9:74:C8
aa:d8:60:f9:74:c8
2022-04-21 18:25:37
AF:A8:A1:C6:B6:DD
af:a8:a1:c6:b6:dd
2022-04-21 19:09:36
EF:99:47:5B:C9:52
ef:99:47:5b:c9:52
2022-04-21 16:06:22
F5:0B:0A:35:68:2F
f5:0b:0a:35:68:2f
2022-04-21 17:55:02
FE:90:F0:F7:E6:BD
fe:90:f0:f7:e6:bd
2022-04-22 16:44:52
Details
00:50:56:A0:F7:3A
po_number: 3453421
purchase_date: 2022_01_01
ram:
schedule:
serial_number: 43DS6545FG
skip_sync: false
start_date: 03042023
subcategory: Pump
support_group:
supported_by:
sys_class_name: Sensor
sys_class_path: /!!/!2/!(
sys_created_by: Richard_Trevithick
sys_created_on: 2022-04-21 16:23:34
sys_domain: global
sys_domain_path: /
sys_id: 002467fcdb4f01101f0f174b13961941
sys_mod_count: 0
sys_tags:
sys_updated_by: jbloggs
sys_updated_on: 2022-04-21 16:23:34
u_segmentation_group_tag: cts:security-group-tag=2775-000
unverified: false

Figure 19-21 *pxGrid Direct Attributes for an Endpoint in ISE Context Visibility*

Summary

Effective security orchestration relies on the transition from static, siloed defenses to a dynamic, context-aware architecture. The Cisco Platform Exchange Grid (pxGrid) serves as the foundational framework for this transformation, acting as the connective tissue that enables diverse security products to exchange identity, posture, and threat intelligence in real time. By moving beyond traditional point-to-point configurations, pxGrid utilizes a scalable hub-and-spoke model where a centralized controller facilitates communication between publishers and subscribers, ensuring that critical data, such as user identity and device state, is shared consistently across the security stack.

The evolution of the framework, specifically the move to pxGrid 2.0, has standardized communication through REST APIs and WebSockets, utilizing JSON-formatted payloads and the Simple Text Oriented Messaging Protocol (STOMP). This shift enables high-throughput, persistent, and efficient bidirectional data transfer, which is essential for modern, time-sensitive security operations.

Strategic integrations with platforms like Cisco Catalyst Center and Cisco Secure Firewall demonstrate the power of this contextual exchange. In these environments, bidirectional synchronization of security group definitions, policies, and device inventory allows for granular, automated access control and consistent segmentation across both wired and wireless infrastructure. Furthermore, the introduction of pxGrid Cloud extends these capabilities to cloud-based security solutions, providing a seamless bridge for remote and hybrid environments to leverage the same granular intelligence found in on-premises deployments.

Technet24

For environments involving legacy or third-party products that lack native pxGrid support, pxGrid Direct provides an essential extension mechanism. By functioning as an integration layer that normalizes external data via REST APIs, it allows administrators to ingest third-party information, such as ServiceNow IT service management data, directly into the ISE dictionary. This data can then be used to drive dynamic authorization rules, enabling a unified, enterprise-wide operational view. Ultimately, the combination of native pxGrid services and pxGrid Direct allows organizations to build a cohesive defense ecosystem, significantly reducing incident response times and providing the visibility necessary to manage the complexities of modern, distributed networks.

Chapter 20. Troubleshooting, Monitoring, Alerting, and System 360

Introduction

Operations and troubleshooting, also known as Day 2, play a significant role in any IT system lifecycle. As administrator, you must not only design and deploy the system but also maintain continuous operation and resolve issues effectively. This chapter provides a structured approach to troubleshooting, which enables you to investigate issues with ISE and other IT systems. The chapter covers tools available in ISE for effective troubleshooting. Practical use cases illustrate how you can apply these tools.

Troubleshooting Methodology

Troubleshooting methodology represents the most important aspect of resolving issues. A step-by-step approach enables you to save time, increase efficiency, and identify the actual problem and solution instead of pursuing incorrect directions. Figure 20-1 shows the structured approach to troubleshooting.

Technet24

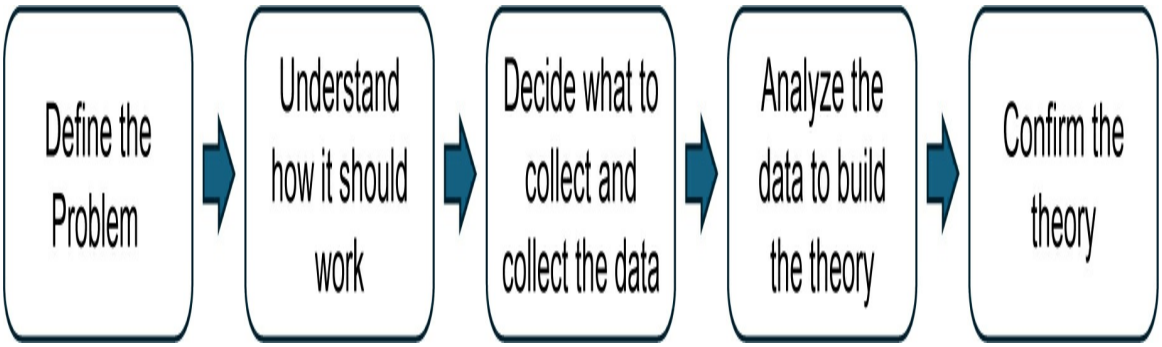


Figure 20-1 *Troubleshooting Flow*

Define the Problem

The number one task in any troubleshooting process involves understanding what is not working — or in other words, defining the problem. When you're an engineer working on a Technical Assistance Case (TAC), you might receive vague problem descriptions such as "broken" or "not working," which remain unspecific and can lead to false assumptions. For example, compare the following two problem statements:

1. Guest authentication on ISE is not working.
2. Newly created guest users on ISE 3.5 patch 3 are failing guest authentication on the Sponsored Guest Portal. The error message displayed on the portal is "Authentication Failed." This is a new installation and was never working before.

The second problem description provides the engineer with immediate insight into the ISE version under investigation. The information that the system has never worked before remains critical because it confirms that external factors, such as changes in the network or upgrades of components, are unlikely causes. Additionally, the description emphasizes the need to verify the configuration, ensuring the guest settings are properly configured. Finally, it identifies the point in the flow where the issue occurs, enabling you to

exclude several potential failure points. For example, you can confirm that the connection to the wireless network succeeded and that redirection functions as expected because the user was redirected to the guest portal.

Understand How It Should Work

The next critical step involves understanding what is expected. The brief answer to the sample problem is that the user should be authenticated. However, real understanding requires knowledge of how the entire feature functions and all involved components. You must understand the guest authentication flow, which includes the initial MAC Authentication Bypass (MAB) authentication, redirection, guest authentication, Change of Authorization (CoA), and final authentication. Potential failures can occur at any stage, so narrowing down issues to a specific part of the flow or component is essential. Without a thorough understanding of how the feature operates, you cannot focus troubleshooting efforts effectively. In contrast, when you know the flow and how it should work, you can immediately exclude certain components, such as the network access device (NAD), client, Active Directory, mobile device management (MDM) server, pxGrid client, and concentrate on the relevant aspects.

> **Note**
>
> Details about the guest flow were covered in Chapter 9, "Managing Guest Users and Personal Devices."

Returning to the original example, guest authentication failed on the portal itself, which narrows down the issue to ISE rather than any external factor, such as redirection, clients, or NADs. It is also unlikely that the RADIUS authentication component is involved. At this stage, you can focus on troubleshooting the guest component within ISE.

Technet24

Decide What to Collect and Collect the Data

In the third step, you should identify which troubleshooting data to collect. Always remember that troubleshooting often requires an iterative approach. You gather a subset of data, confirm or rule out your theory, and then collect additional information based on the results. In the example provided, you can begin with detailed authentication reports. If these reports do not provide sufficient information, you can proceed to Support Bundle collection and guest.log analysis. If further details are needed, consider enabling guest component debugs and performing deeper analysis of the guest logs.

Analyze the Data to Build a Theory

The fourth step is the most crucial and requires domain knowledge. Sometimes, you might switch this step with "Decide what to collect and collect the data." However, relying on assumptions poses a risk. If you build theories without facts, you may collect incorrect data and delay the troubleshooting process. For issues that are more apparent, brainstorming and making assumptions can be effective. The more complex the issue, the more counterproductive the assumption-based approach becomes.

Comparison serves as one of the most effective methods at this stage. You benefit from having a working scenario to compare against. In the provided example, you can compare detailed authentication reports, logs step by step, and configuration settings. The differences between a working and nonworking scenario can quickly help you build a theory and identify potential suspects.

Returning to the example, start with the detailed authentication report. The first step is to locate the relevant authentication attempt. For

guest authentications, you can use timestamps and the guest username (in the identity filter) to find the exact authentication attempt.

> **Note**
>
> To locate most unique authentication attempts, use the Endpoint ID quick filter and enter the MAC address of the endpoint you're troubleshooting.

Figure 20-2 shows the failed guest authentication event among other live log events.

Technet24

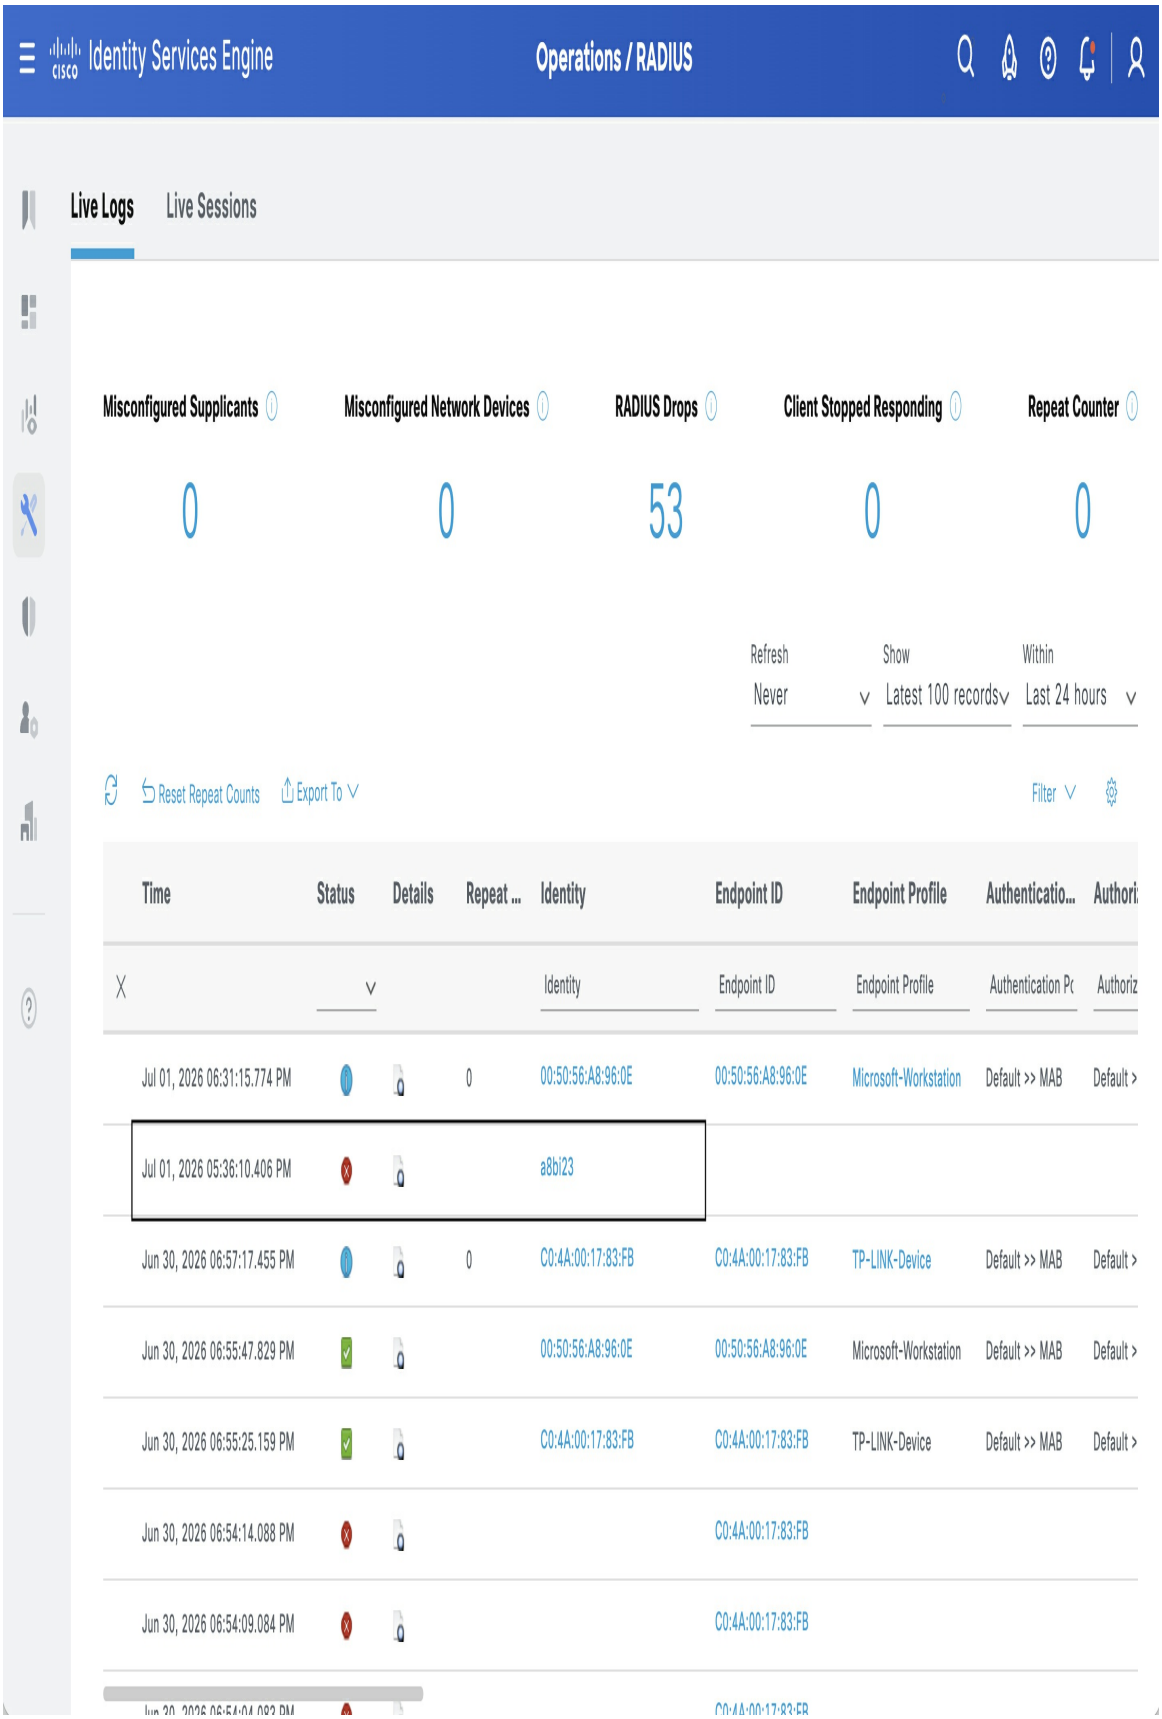

Identity Services Engine
Operations / RADIUS
Live Logs
Live Sessions
Misconfigured Supplicants
Misconfigured Network Devices
RADIUS Drops
Client Stopped Responding
Repeat Counter
0
0
53
0
0
Refresh
Never
Show
Latest 100 records
Within
Last 24 hours
Reset Repeat Counts
Export To
Filter
Time
Status
Details
Repeat ...
Identity
Endpoint ID
Endpoint Profile
Authenticatio...
Authori
Identity
Endpoint ID
Endpoint Profile
Authentication Pc
Authoriz
Jul 01, 2026 06:31:15.774 PM
0
00:50:56:A8:96:0E
00:50:56:A8:96:0E
Microsoft-Workstation
Default >> MAB
Default >
Jul 01, 2026 05:36:10.406 PM
a8bi23
Jun 30, 2026 06:57:17.455 PM
0
C0:4A:00:17:83:FB
C0:4A:00:17:83:FB
TP-LINK-Device
Default >> MAB
Default >
Jun 30, 2026 06:55:47.829 PM
00:50:56:A8:96:0E
00:50:56:A8:96:0E
Microsoft-Workstation
Default >> MAB
Default >
Jun 30, 2026 06:55:25.159 PM
C0:4A:00:17:83:FB
C0:4A:00:17:83:FB
TP-LINK-Device
Default >> MAB
Default >
Jun 30, 2026 06:54:14.088 PM
C0:4A:00:17:83:FB
Jun 30, 2026 06:54:09.084 PM
C0:4A:00:17:83:FB

Figure 20-2 *Live Logs with Failed Guest Authentication*

Figure 20-3 shows a detailed authentication report for the failed attempt.

Technet24

Cisco ISE

Overview

Event	5418 Guest Authentication Failed
Username	a8bi23
Endpoint Id	
Endpoint Profile	
Authorization Result	

Steps

Step ID	Description	Latency (ms)
5418	Guest Authentication Failed	

Authentication Details

Source Timestamp	Jul 01, 2026 05:36:10.000 PM
Received Timestamp	Jul 01, 2026 05:36:10.000 PM
Policy Server	ISE35-1ek
Event	5418 Guest Authentication Failed
Failure Reason	Account is not yet active.
Username	a8bi23
User Type	GuestUser
Endpoint Id	
Calling Station Id	
Endpoint Profile	
IPv4 Address	
IPv6 Address	
Authentication Identity Store	Guest Users
Identity Group	GuestType_Contractor (default)

Figure 20-3 *A Detailed Authentication Report with Failed Guest Authentication*

From the detailed authentication report, you can confirm the following:

1. Guest authentication failed for the user **a8bi23**.
2. You are reviewing the correct failed attempt. Multiple failures may occur at different times or on different days, so it is crucial to troubleshoot the issue you are facing and not another unrelated incident, such as a one-off failure from a week ago.
3. The Failure Reason is **Account is not yet active**, which helps you build the theory.
4. The Authentication Identity Store is **Guest Users**. This setting confirms that the correct identity store was used and eliminates misconfigured identity stores from consideration, such as using Active Directory instead of guest users.
5. The Identity Group assigned to the user is **GuestType_Contractor.** This information helps with configuration verification, allowing you to confirm the settings associated with this guest type.

Based on the error message, you can build a theory. The "Account is not yet active" error most likely indicates that the validity period starts in the future. In the sponsor portal, you can verify the actual validity of the guest user, as shown in Figure 20-4.

Technet24

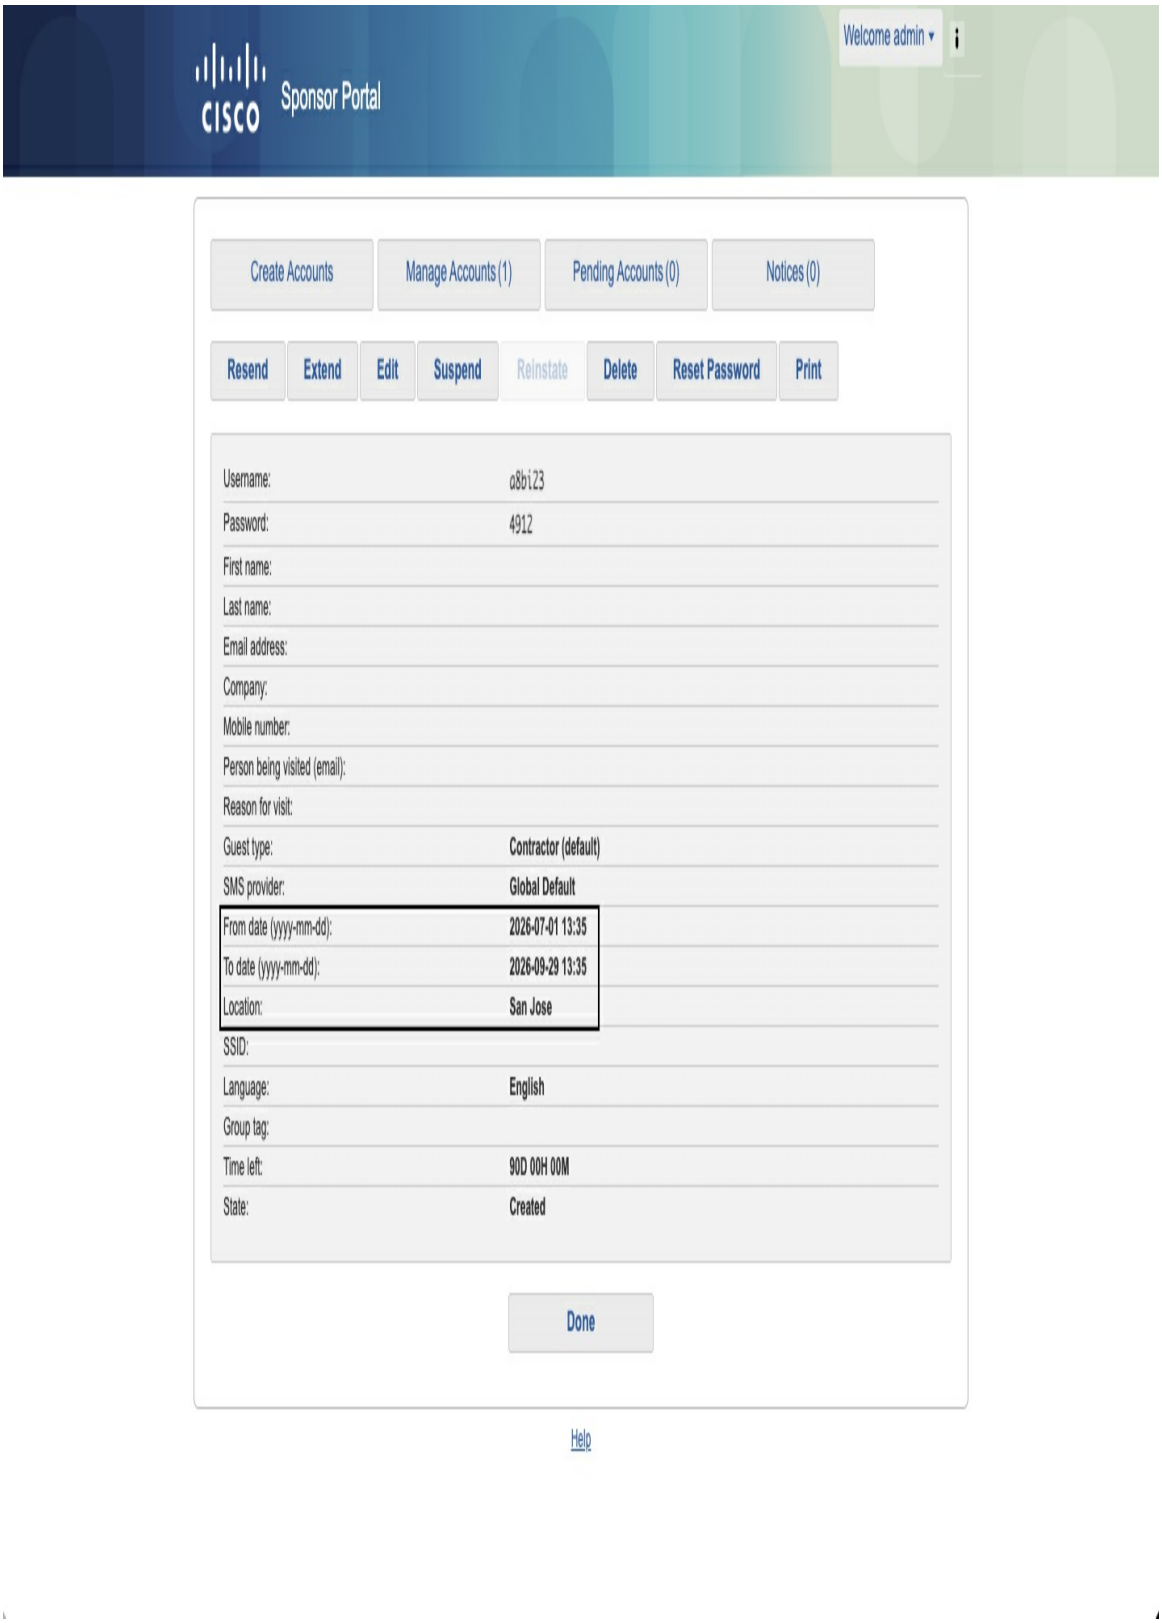


Figure 20-4 *Guest Account Settings Under the Sponsor Portal*

Now, you can correlate the facts:

1. ISE recorded the failed guest authentication attempt at 5:36 PM (17:36) UTC on July 1. You can verify the ISE time zone by using the **show time zone** CLI command.
2. The guest account is valid from 1:35 PM (13:35) July 1, with the location set to San Jose.

The theory you can build is that the guest account becomes valid only from 8:35 PM UTC, which is later than the actual authentication attempt.

Confirm the Theory

The final step in troubleshooting involves confirming the theory. As mentioned earlier, you may need to repeat Steps 3–5 if the initial confirmation of the theory fails. Successful confirmation results in a working scenario, where, after you make the necessary changes, the feature operates as expected. During this step, you can adjust settings to the correct values, update the system if a bug is discovered, or restore the system if external factors, such as a power outage, are identified.

You have several options to confirm the theory. You can wait for the guest user to become active, create a new guest user in the UTC time zone and test immediately, or create a new user in the San Jose time zone and configure the guest account to become active earlier during creation.

Figure 20-5 shows the detailed authentication report for the same guest user, who authenticated on the same day at 10:50 PM (22:50) UTC on July 1, which corresponds to 3:50 PM (15:50) in San Jose. At this time, the user was already active. This result confirms the theory

Technet24

regarding the failed guest authentication.

Cisco ISE

Overview

Event	5231 Guest Authentication Passed
Username	a8bi23
Endpoint Id	
Endpoint Profile	
Authorization Result	

Steps

Step ID	Description	Latency (ms)
5231	Guest Authentication Passed	

Authentication Details

Source Timestamp	Jul 01, 2026 10:50:44.000 PM
Received Timestamp	Jul 01, 2026 10:50:44.000 PM
Policy Server	ISE35-1ek
Event	5231 Guest Authentication Passed
Failure Reason	
Username	a8bi23
User Type	GuestUser
Endpoint Id	
Calling Station Id	

Technet24

Figure 20-5 *A Successful Guest Authentication Detailed Authentication Report*

You can resolve this issue permanently in several ways. You can add a new location or time zone that matches the ISE time zone under **Work Centers > Guest Access > Settings > Guest Locations and SSIDs**, and allow it to be used with sponsors. Alternatively, you can use the guest location of your choice and account for the ISE time zone when creating guest users.

Technically, you can also place Policy Service Nodes (PSNs) in their original time zones. However, the recommended practice is to configure all nodes with the UTC time zone.

Alarms

Alarms on ISE typically receive immediate attention from administrators. They enable proactive identification of issues, such as an alarm notification sent before certificate expiration. Alarms also provide reactive notifications as soon as misbehavior is detected — for example, an alarm notification confirming that an MDM reachability issue has occurred.

As administrator, you can access alarms in multiple ways. The most common method is by navigating to the **ISE Dashboard**. The **Alarms** dashlet displays a list of alarms sorted in chronological order, as shown in Figure 20-6.

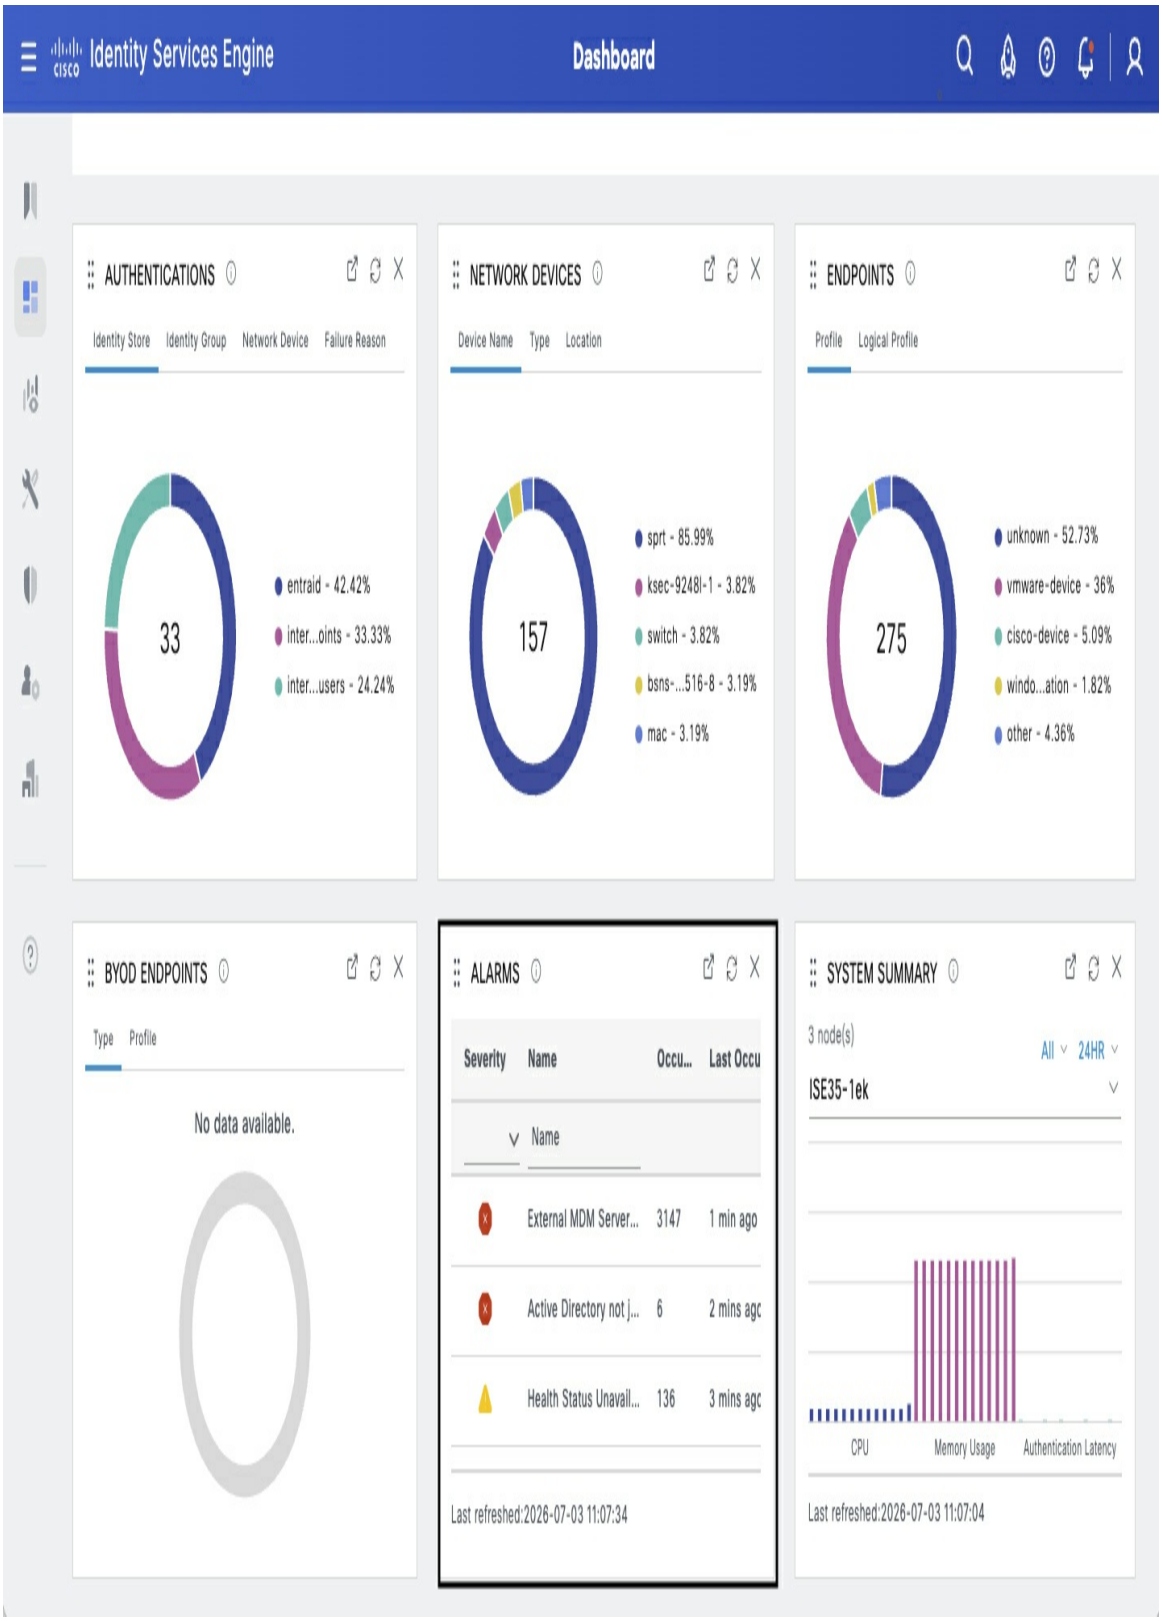


Figure 20-6 *The Alarms Dashlet on the ISE Dashboard*

Technet24

When you analyze alarms, it is critical to understand their relevance. The **Last Occurrence** column indicates the last time the alarm was triggered. The **Occurrences** column shows how many times this alarm has occurred. These two parameters help you determine whether the alarm represents a one-off event, such as a brief disconnect between nodes in two data centers producing a Replication alarm, or a consistent issue, such as a Health Status Unavailable alarm that is constantly seen from one node in the deployment.

From the **Alarms** dashlet, you can click the alarm of interest to view more details, as shown in Figure 20-7. Additionally, you can acknowledge an alarm you have already investigated, so the occurrences remain relevant to new issues.

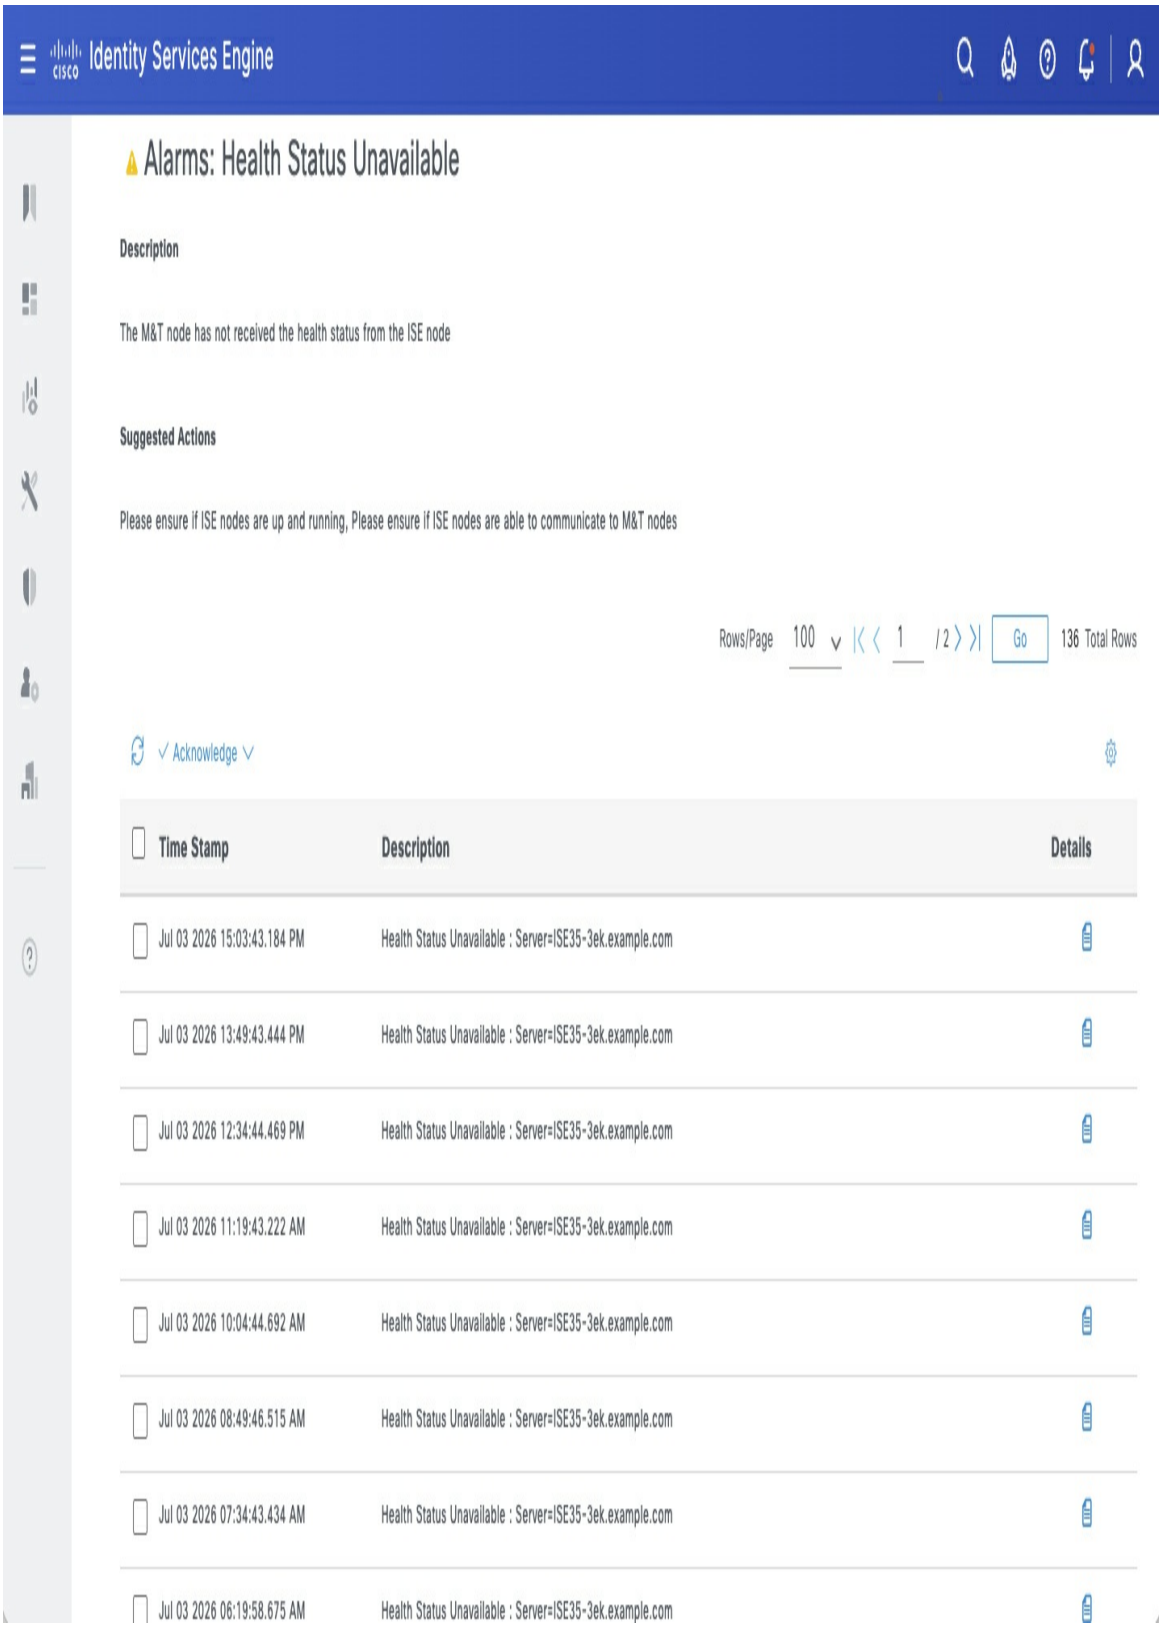


Figure 20-7 *Details of a Health Status Unavailable Alarm*

Technet24

The detailed view provides information about each occurrence along with millisecond-level timestamps. This data helps you correlate the alarm with other potentially impactful activities, such as a node reload, networking issue, or power outage.

Some alarms include additional details, which you can access by clicking the **Details** button on the same page. You can also receive alarms via email. **Configure Alarm Notifications** under **Administration > System > Settings > Alarm Settings > Alarm Notification**, as shown in Figure 20-8. Enter the email address of ISE administrators so they begin receiving alarms in their mailboxes.

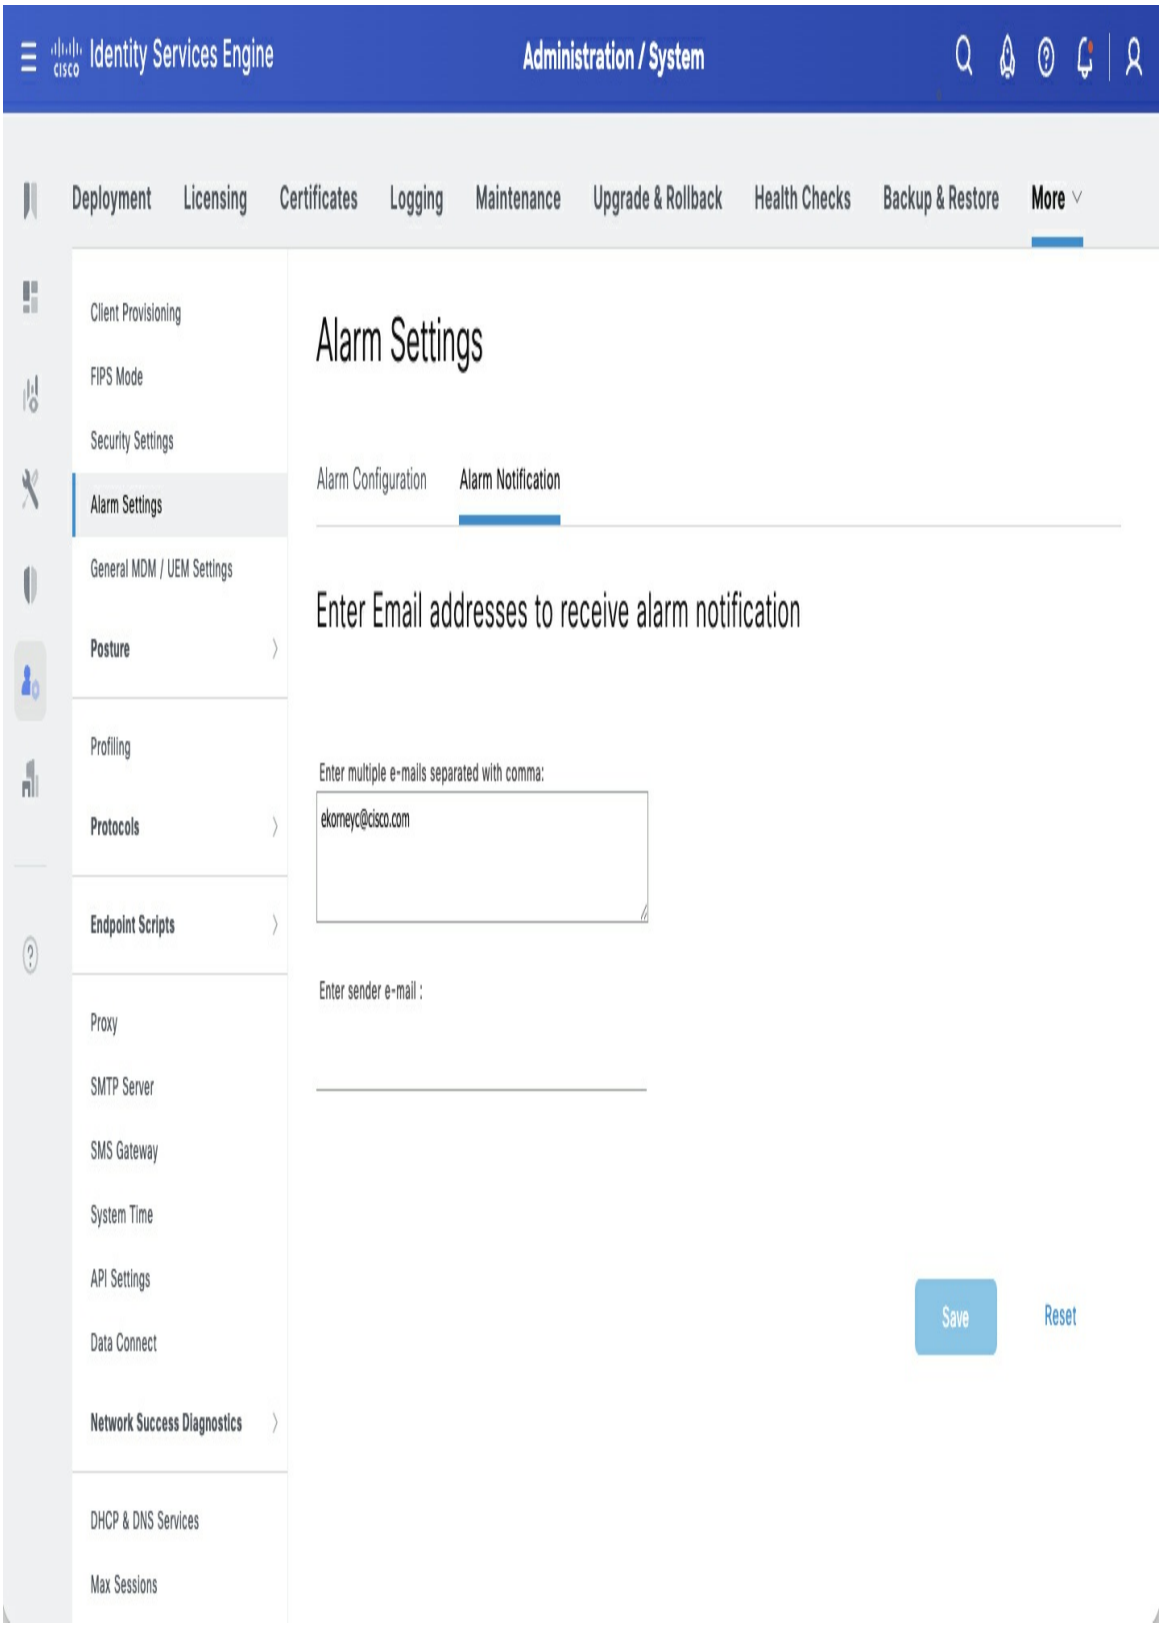


Figure 20-8 *Alarm Notification Settings*

Technet24

An example of the email notification is shown in Figure 20-9.

ISE Alarm : Warning : . Domain Name=ISELAB.NET Domain Controller=WIN2025.iselab.net Error Details=Domain controller is unreachable



ise@ISE35-1ek.cisco.com <ise@ISE35-1ek.cisco.com>

Today at 1:05 PM

To: Eugene Korneychuk (ekorneyc)

Alarm Name :
Active Directory Domain Controller services are unavailable

Details :
. Domain Name=ISELAB.NET Domain Controller=WIN2025.iselab.net Error Details=Domain controller is unreachable

Description :
Active Directory Domain Controller services are unavailable. Please verify Domain Controller services availability. System will failover to another available best DC

Severity :
Warning

Suggested Actions :
Please verify Domain Controller services availability

*** This message is generated by Cisco Identity Services Engine (ISE) ***

Sent By Host : ISE35-1ek

Technet24

Figure 20-9 *Email with Alarm from ISE*

From the first tab of this menu, **Alarm Configuration**, as shown in Figure 20-10, you can control the status of the alarm, choosing whether to disable or enable it. You can also set the severity of the alarm for improved filtering.

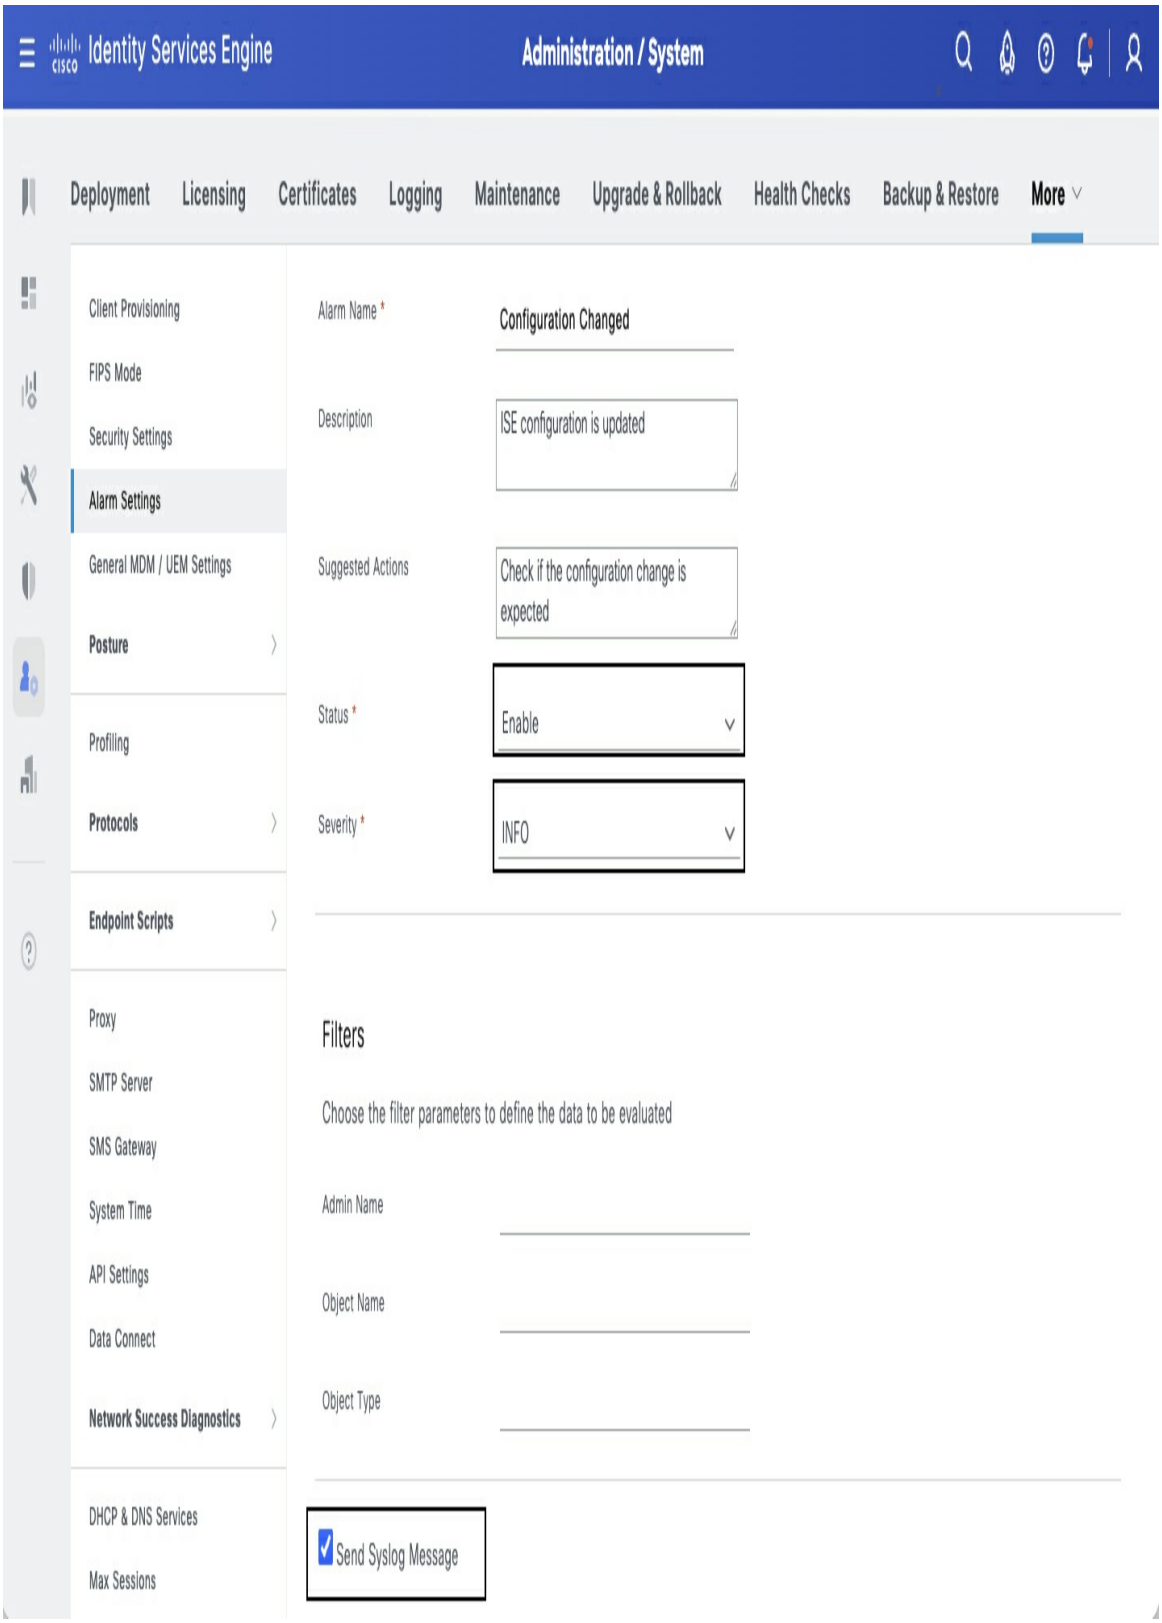


Figure 20-10 *Alarm Configuration*

Technet24

If you want to generate a syslog for the alarm, enable **Send Syslog Message**. This setting allows alarms to be stored on SIEM systems such as Splunk.

Note

Ensure that your Remote Logging Target is configured to receive alarms, as shown in Figure 20-11.

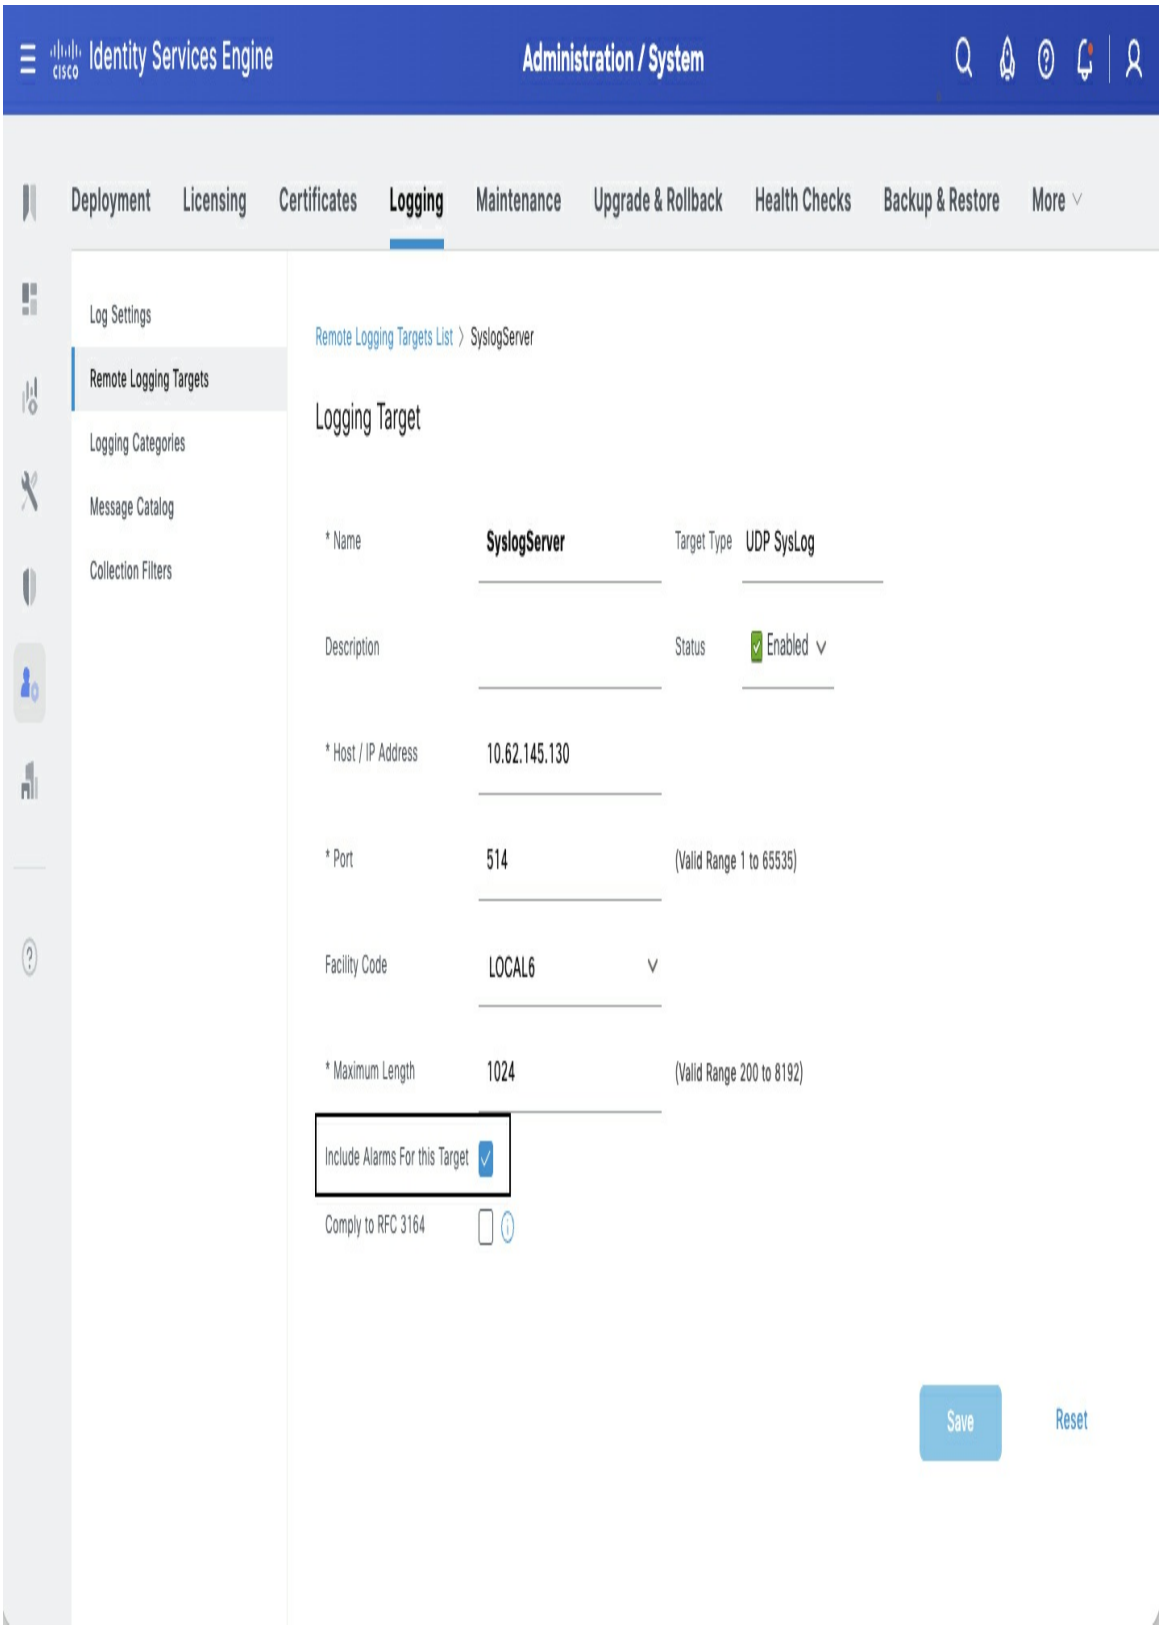


Figure 20-11 *Remote Logging Target Configuration to Include*

Technet24

Alarms

Finally, alarms are also available in the Support Bundle under **Logs > Alarms > alarmexp.txt**.

Live Logs, Live Sessions, and Detailed Authentication Reports

Live logs, live sessions, and especially detailed authentication reports represent the most powerful tools in the ISE arsenal for troubleshooting. These tools allow you to observe every authentication, examine session attributes in detail, and troubleshoot nearly any issue related to authentication, authorization, or policy evaluation on ISE. They serve as the primary resources to monitor user activity, confirm whether proper policies were matched, and evaluate who connected from which location and at what time.

Live Logs

Live logs are accessible through **Operations > Live Logs**, as shown in Figure 20-12.

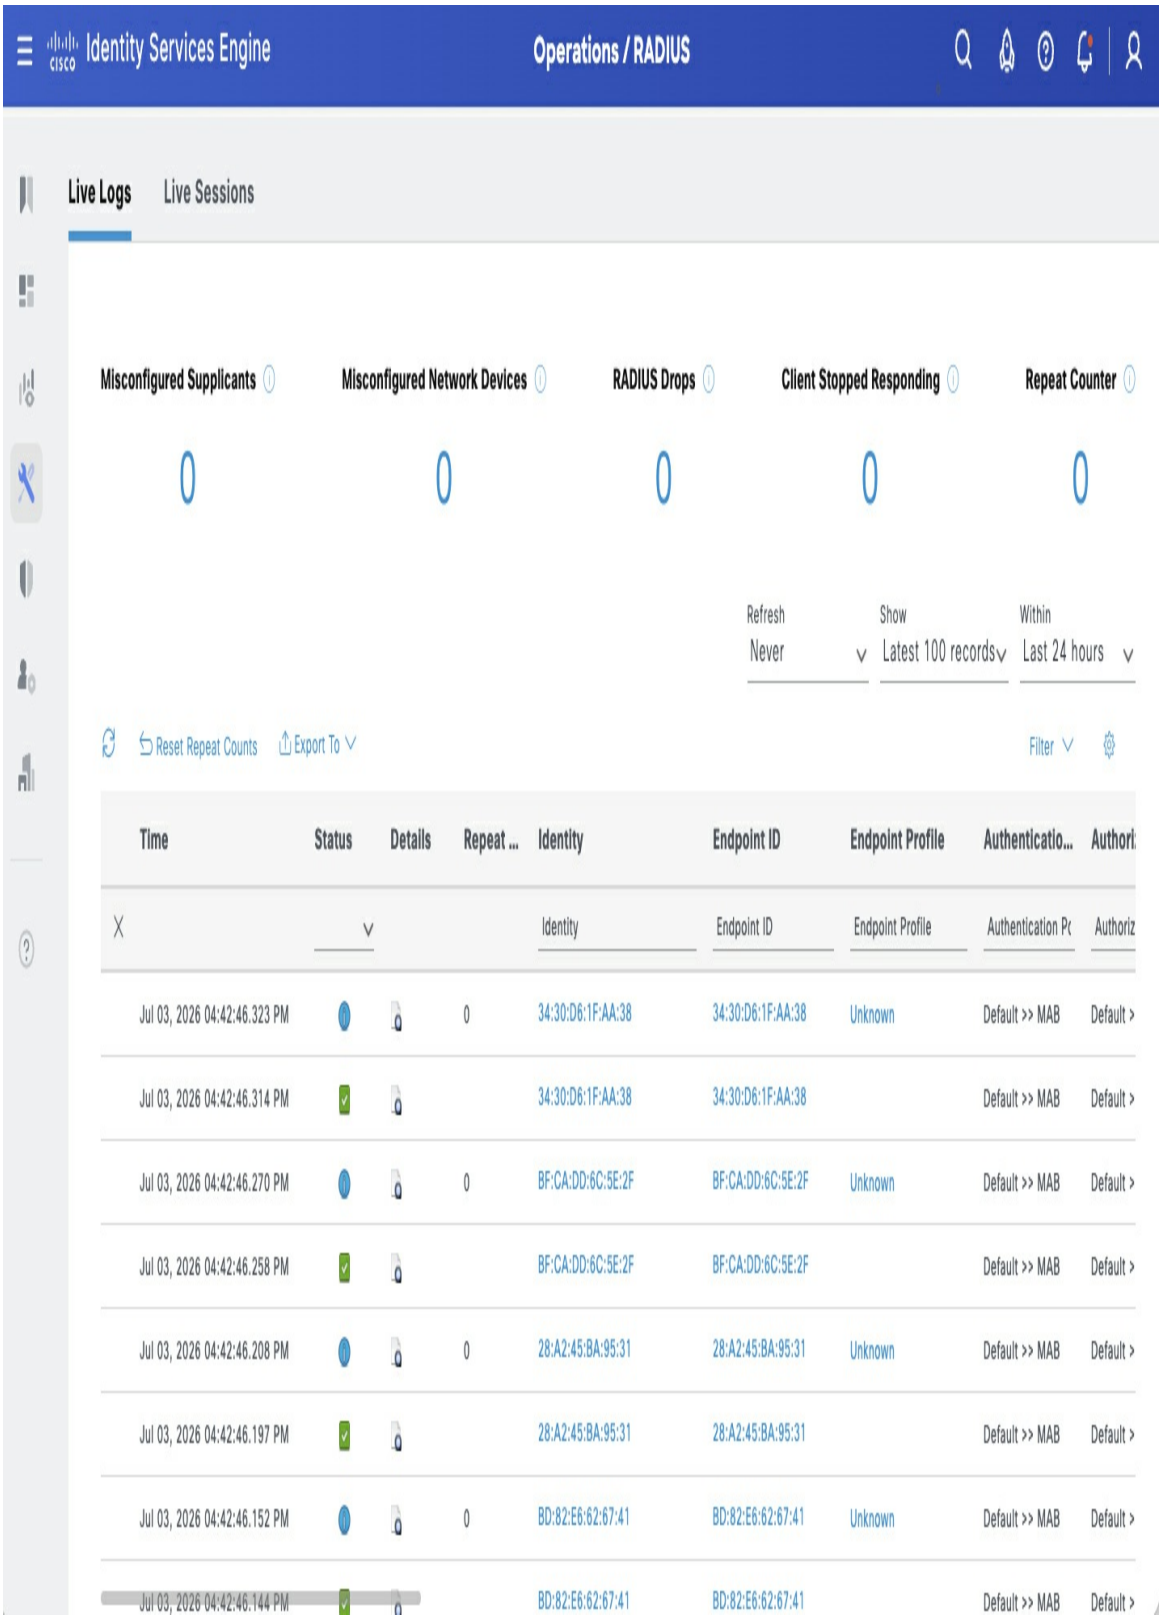


Figure 20-12 *Live Logs*

Technet24

Whenever you need to confirm an authentication attempt performed by a user or endpoint, navigate to **Live Logs**. In production systems, authentications occur constantly, so filters are essential. It is crucial to filter accurately to avoid missing relevant events. The most effective way to filter RADIUS authentications is by entering the MAC address of the endpoint in the **Endpoint ID** column. This approach ensures you capture all authentication, authorization, or CoA events associated with the endpoint, regardless of the authentication type.

> ### Note
>
> The download of a DACL event does not have a MAC address associated with it and will be excluded if you use the MAC address filter in the endpoint ID. To troubleshoot DACLs, use live logs without filters or refer to reports (for historical events) during the relevant authentication timeframe.

Another important consideration with live logs is that the data displayed covers only the last 24 hours. If you need to troubleshoot an event that did not occur recently, you should use reports instead.

While live logs can display up to 100 records at a time, ISE retains data from the last 24 hours. Even if a user authentication that occurred, for example, 6 hours ago does not appear within the initial 100 records, you can filter by username or MAC address to locate and display the relevant authentication.

The **Live Log Status** column shows the following results (you can also filter based on this column):

- Green checkbox indicates successful authentication.
- Blue exclamation mark indicates the session was created for this authentication.

- Red cross indicates failed authentication.

Live Sessions

Live sessions display which endpoints are connected at any given time, as shown in Figure 20-13.

Technet24

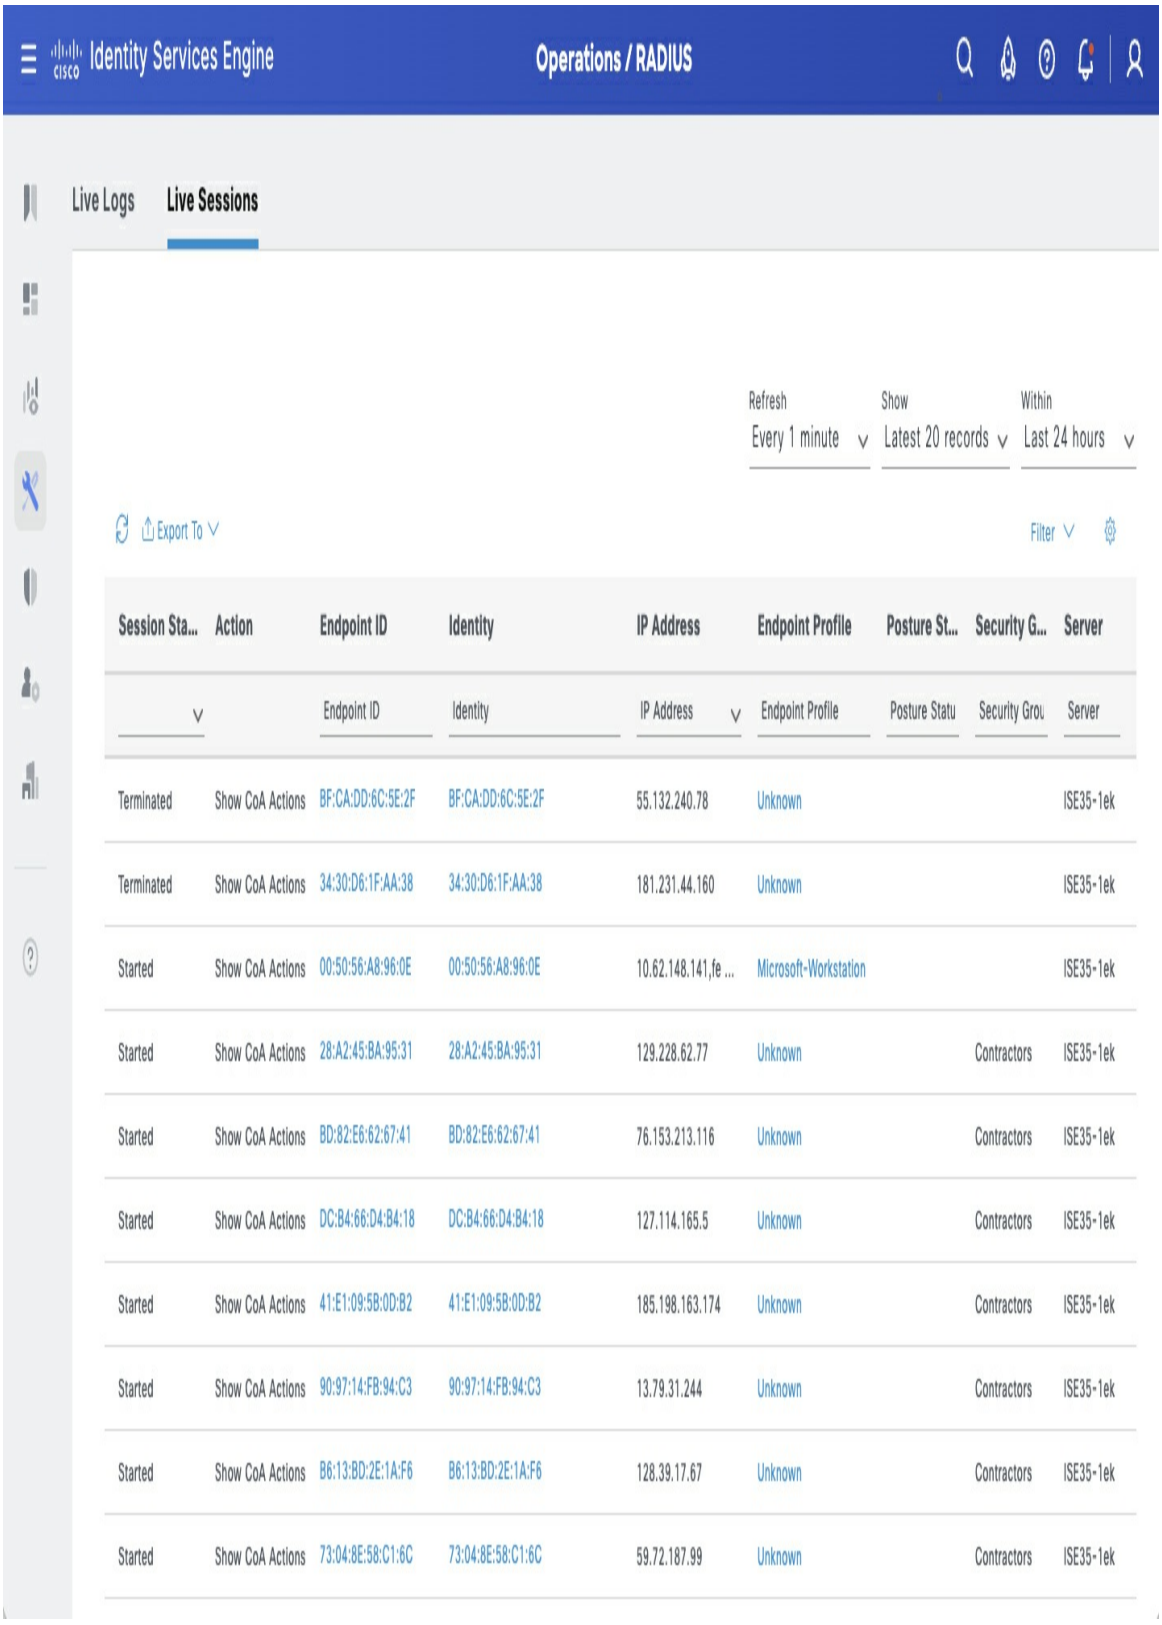


Figure 20-13 *Live Sessions*

You can quickly get access to such information as

- Session status
- PSNs that authenticated the user or endpoint
- IP address assigned
- Endpoint profile
- SGT assigned
- Authentication and authorization policies matched
- Licenses consumed and their reasons

Most ISE licenses are session-based, so understanding the status of the session remains very important. There are three session statuses:

- **Authenticated:** This session was created based on RADIUS authentication with no accounting. The session is removed within one hour.

  ### Note

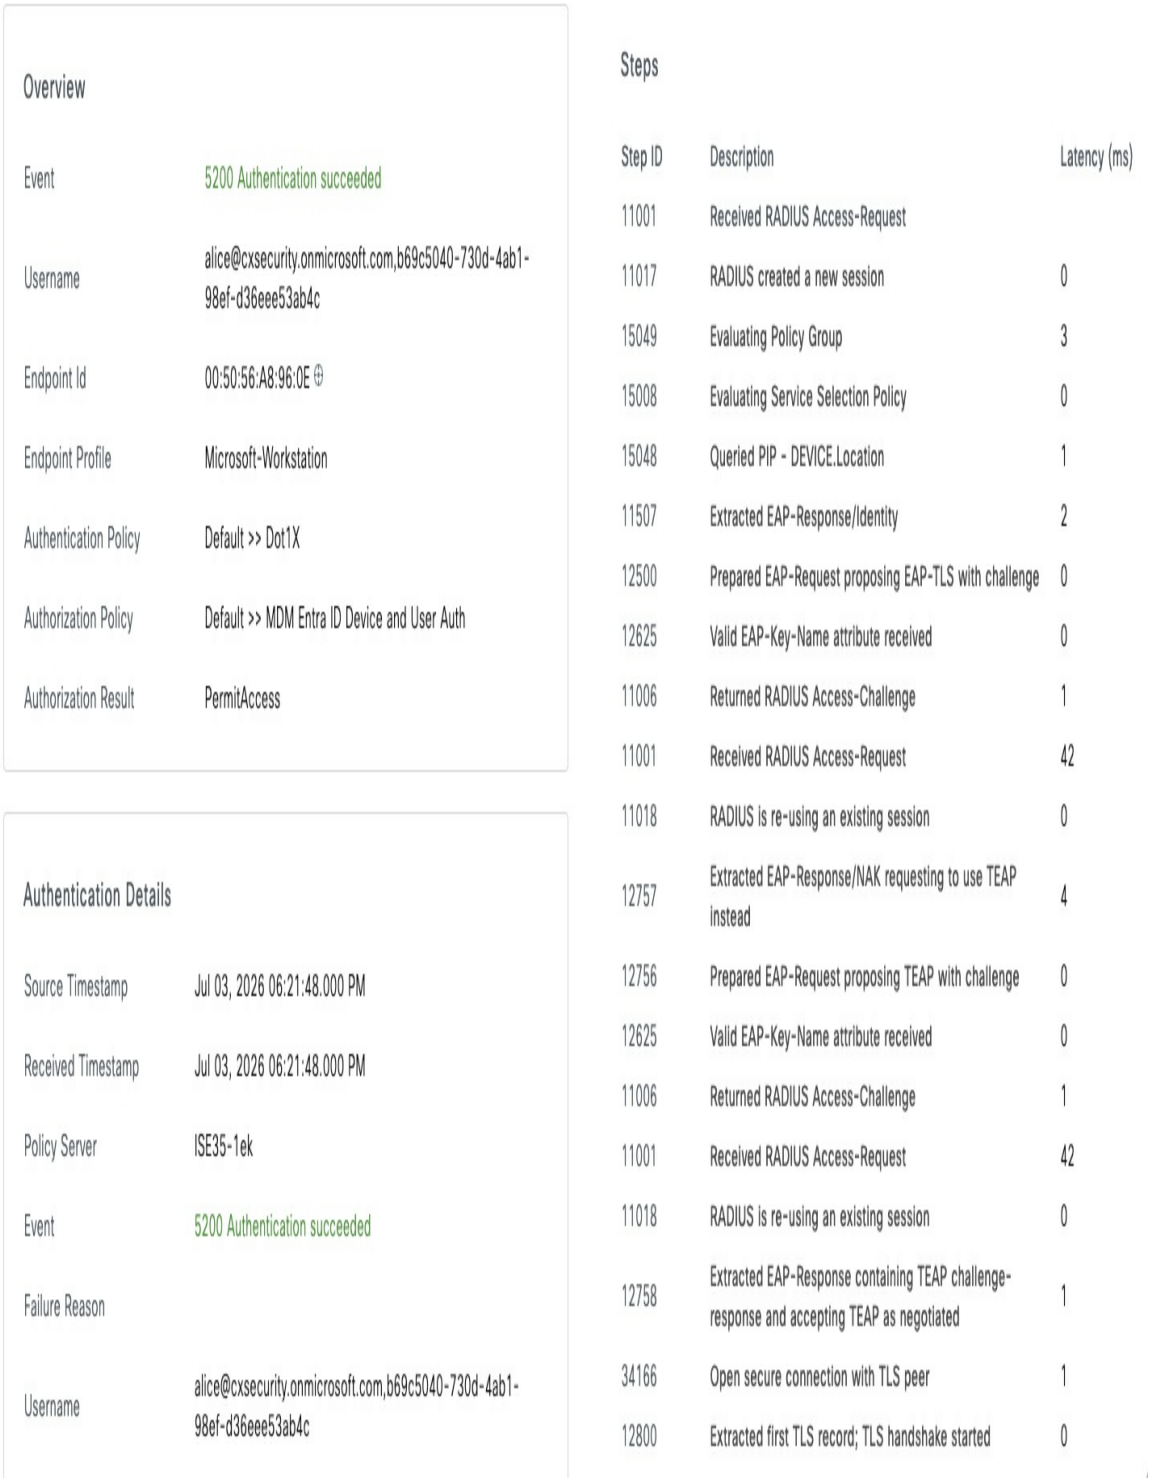
  If the session is removed from ISE, it does not mean the user is disconnected.

- **Started:** This session was created based on RADIUS authentication followed by Accounting Start. The session is removed after five days if there is no accounting activity, such as Accounting Interim Updates.
- **Terminated:** This interim status indicates the session is being removed from live sessions. You see this status when you manually disconnect the session from live logs or when an accounting stop is received. The live session with this status is removed within 20 minutes of disconnection.

Technet24

Detailed Authentication Reports

Detailed authentication reports are accessible directly from live logs. They provide comprehensive information about a single authentication attempt, as shown in Figure 20-14.



Overview

Event	5200 Authentication succeeded
Username	alice@cxsecurity.onmicrosoft.com,b69c5040-730d-4ab1-98ef-d36eee53ab4c
Endpoint Id	00:50:56:A8:96:0E
Endpoint Profile	Microsoft-Workstation
Authentication Policy	Default >> Dot1X
Authorization Policy	Default >> MDM Entra ID Device and User Auth
Authorization Result	PermitAccess

Authentication Details

Source Timestamp	Jul 03, 2026 06:21:48.000 PM
Received Timestamp	Jul 03, 2026 06:21:48.000 PM
Policy Server	ISE35-1ek
Event	5200 Authentication succeeded
Failure Reason	
Username	alice@cxsecurity.onmicrosoft.com,b69c5040-730d-4ab1-98ef-d36eee53ab4c

Steps

Step ID	Description	Latency (ms)
11001	Received RADIUS Access-Request	
11017	RADIUS created a new session	0
15049	Evaluating Policy Group	3
15008	Evaluating Service Selection Policy	0
15048	Queried PIP - DEVICE.Location	1
11507	Extracted EAP-Response/Identity	2
12500	Prepared EAP-Request proposing EAP-TLS with challenge	0
12625	Valid EAP-Key-Name attribute received	0
11006	Returned RADIUS Access-Challenge	1
11001	Received RADIUS Access-Request	42
11018	RADIUS is re-using an existing session	0
12757	Extracted EAP-Response/NAK requesting to use TEAP instead	4
12756	Prepared EAP-Request proposing TEAP with challenge	0
12625	Valid EAP-Key-Name attribute received	0
11006	Returned RADIUS Access-Challenge	1
11001	Received RADIUS Access-Request	42
11018	RADIUS is re-using an existing session	0
12758	Extracted EAP-Response containing TEAP challenge-response and accepting TEAP as negotiated	1
34166	Open secure connection with TLS peer	1
12800	Extracted first TLS record; TLS handshake started	0

Figure 20-14 *Detailed Authentication Report*

The report consists of four main parts: Overview, Authentication Details, Steps, and Result.

Technet24

The Overview section contains a summary of the authentication event. It includes the authentication result, username, endpoint ID, profile, matched authentication and authorization policy, and the final authorization result. This section is helpful when you want to quickly confirm whether the matched authentication and authorization policy is the one you expected.

The Authentication Details section provides comprehensive information about the endpoint and authentication attempt. It includes all details known to ISE, such as the protocols used during authentication, timestamps, the PSN that authenticated the user or endpoint, network access device information, certificate attributes, AD groups, and more. This section is especially useful when you need to examine the specifics of authentication, verify whether the correct attributes are received, ensure the appropriate authentication protocol is used, check whether the conditions defined in policies are present, and determine if other conditions might be suitable for the policy.

The Steps section displays all internal and external steps ISE performed to authenticate the user or endpoint. Each step includes an associated latency measured in milliseconds. This section is useful when you troubleshoot authentication latency, third-party identity store integrations such as LDAP or Active Directory, and RADIUS authentications. It helps you confirm where the failure occurred in the flow of requests and challenges.

The detailed authentication report also includes the Result section, which reveals all attributes returned with Access-Accept. This information can help you confirm whether the correct SGT, DACL, dynamic VLAN, or any other attribute was returned as a result of authentication.

Let's explore a troubleshooting scenario where live logs with the

detailed authentication report can provide significant assistance. Consider the following:

Problem Statement: User Authentication is failing.

While this problem statement remains generic, you can use live logs to refine it. When you filter live logs with the problematic endpoint's MAC address, as shown in Figure 20-15, two authentications are observed. The first one, in chronological order, appears to be a failed TEAP authentication, indicated by two usernames separated by a comma. The second authentication is MAB, which also fails.

Technet24

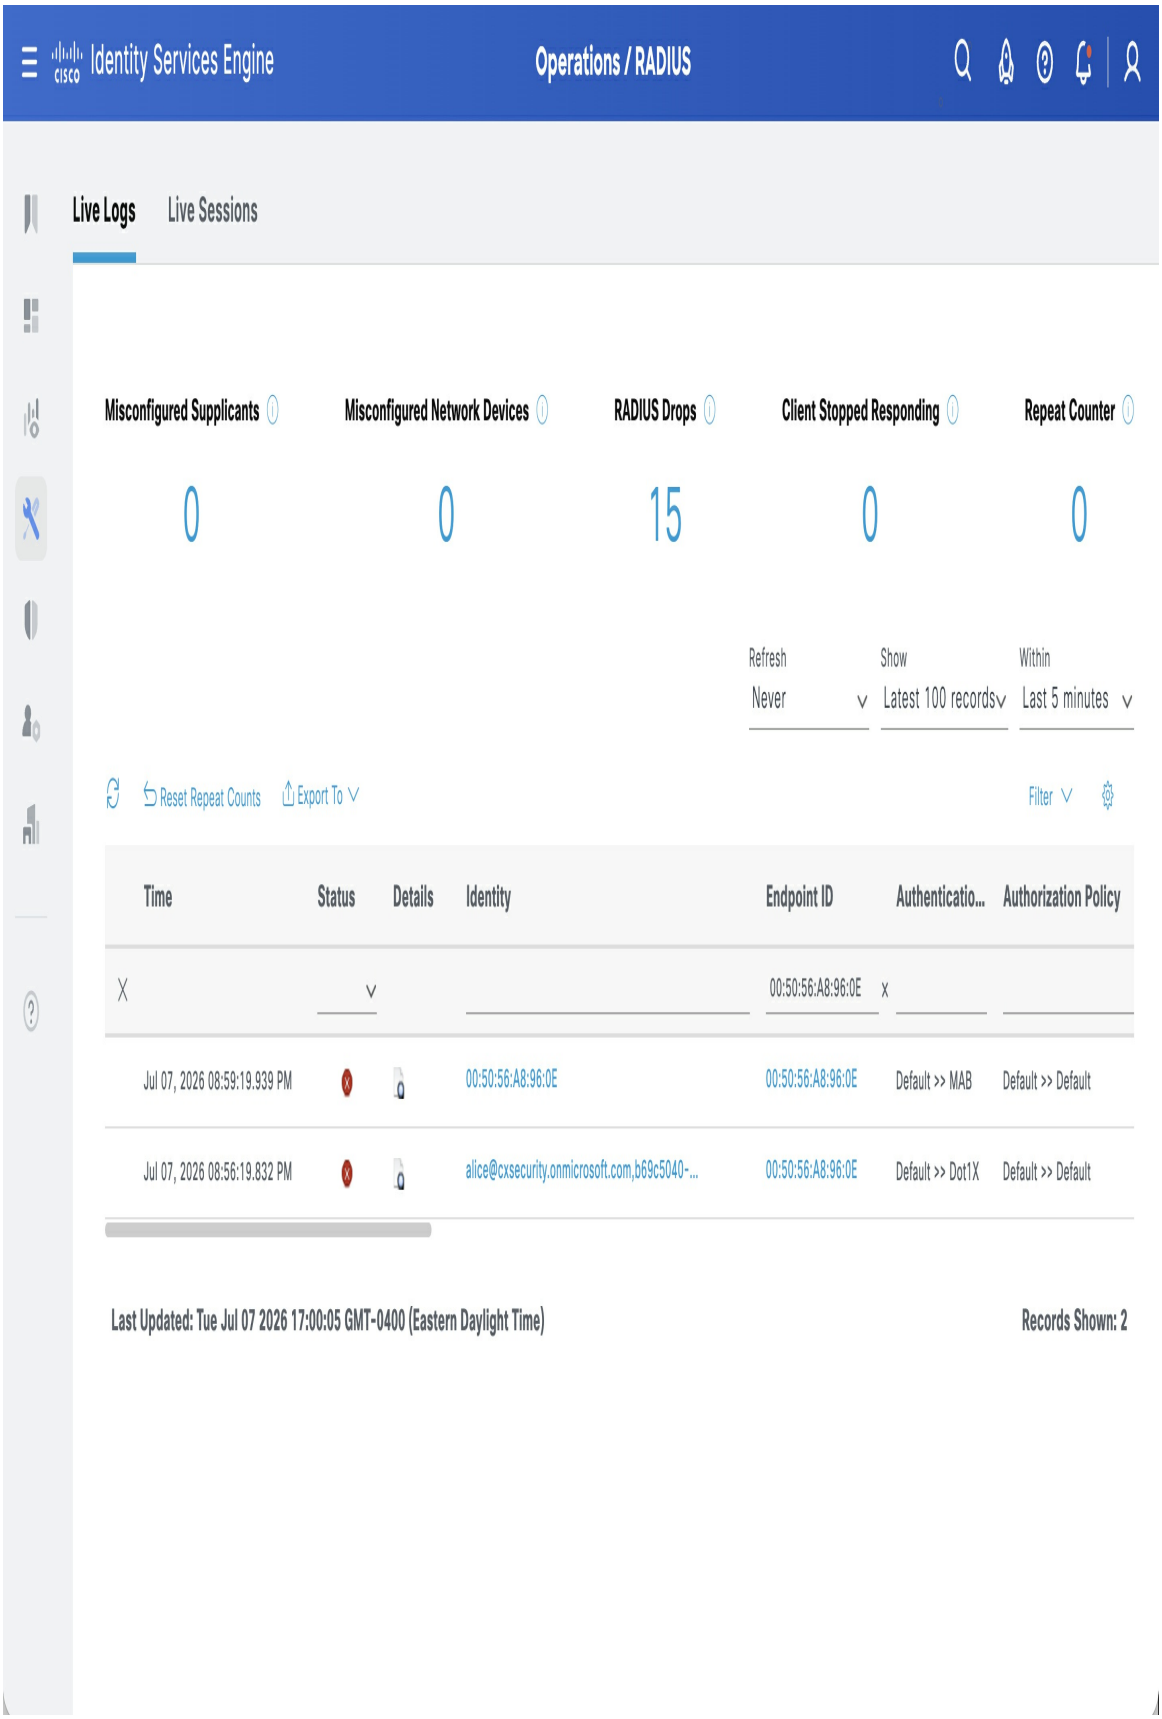

Identity Services Engine
Operations / RADIUS
Live Logs
Live Sessions
Misconfigured Supplicants
Misconfigured Network Devices
RADIUS Drops
Client Stopped Responding
Repeat Counter
0
0
15
0
0
Refresh
Never
Show
Latest 100 records
Within
Last 5 minutes
Reset Repeat Counts
Export To
Filter
Time
Status
Details
Identity
Endpoint ID
Authenticatio...
Authorization Policy
00:50:56:A8:96:0E
Jul 07, 2026 08:59:19.939 PM
00:50:56:A8:96:0E
00:50:56:A8:96:0E
Default >> MAB
Default >> Default
Jul 07, 2026 08:56:19.832 PM
alice@cxsecurity.onmicrosoft.com,b69c5040-...
00:50:56:A8:96:0E
Default >> Dot1X
Default >> Default
Last Updated: Tue Jul 07 2026 17:00:05 GMT-0400 (Eastern Daylight Time)
Records Shown: 2

Figure 20-15 *Live Logs Filtered with the Problematic Endpoint ID*

In such cases, you should always begin by identifying the expected behavior. The expected behavior must match the configured policy. In this particular example, the relevant policy is shown in Figure 20-16.

Technet24

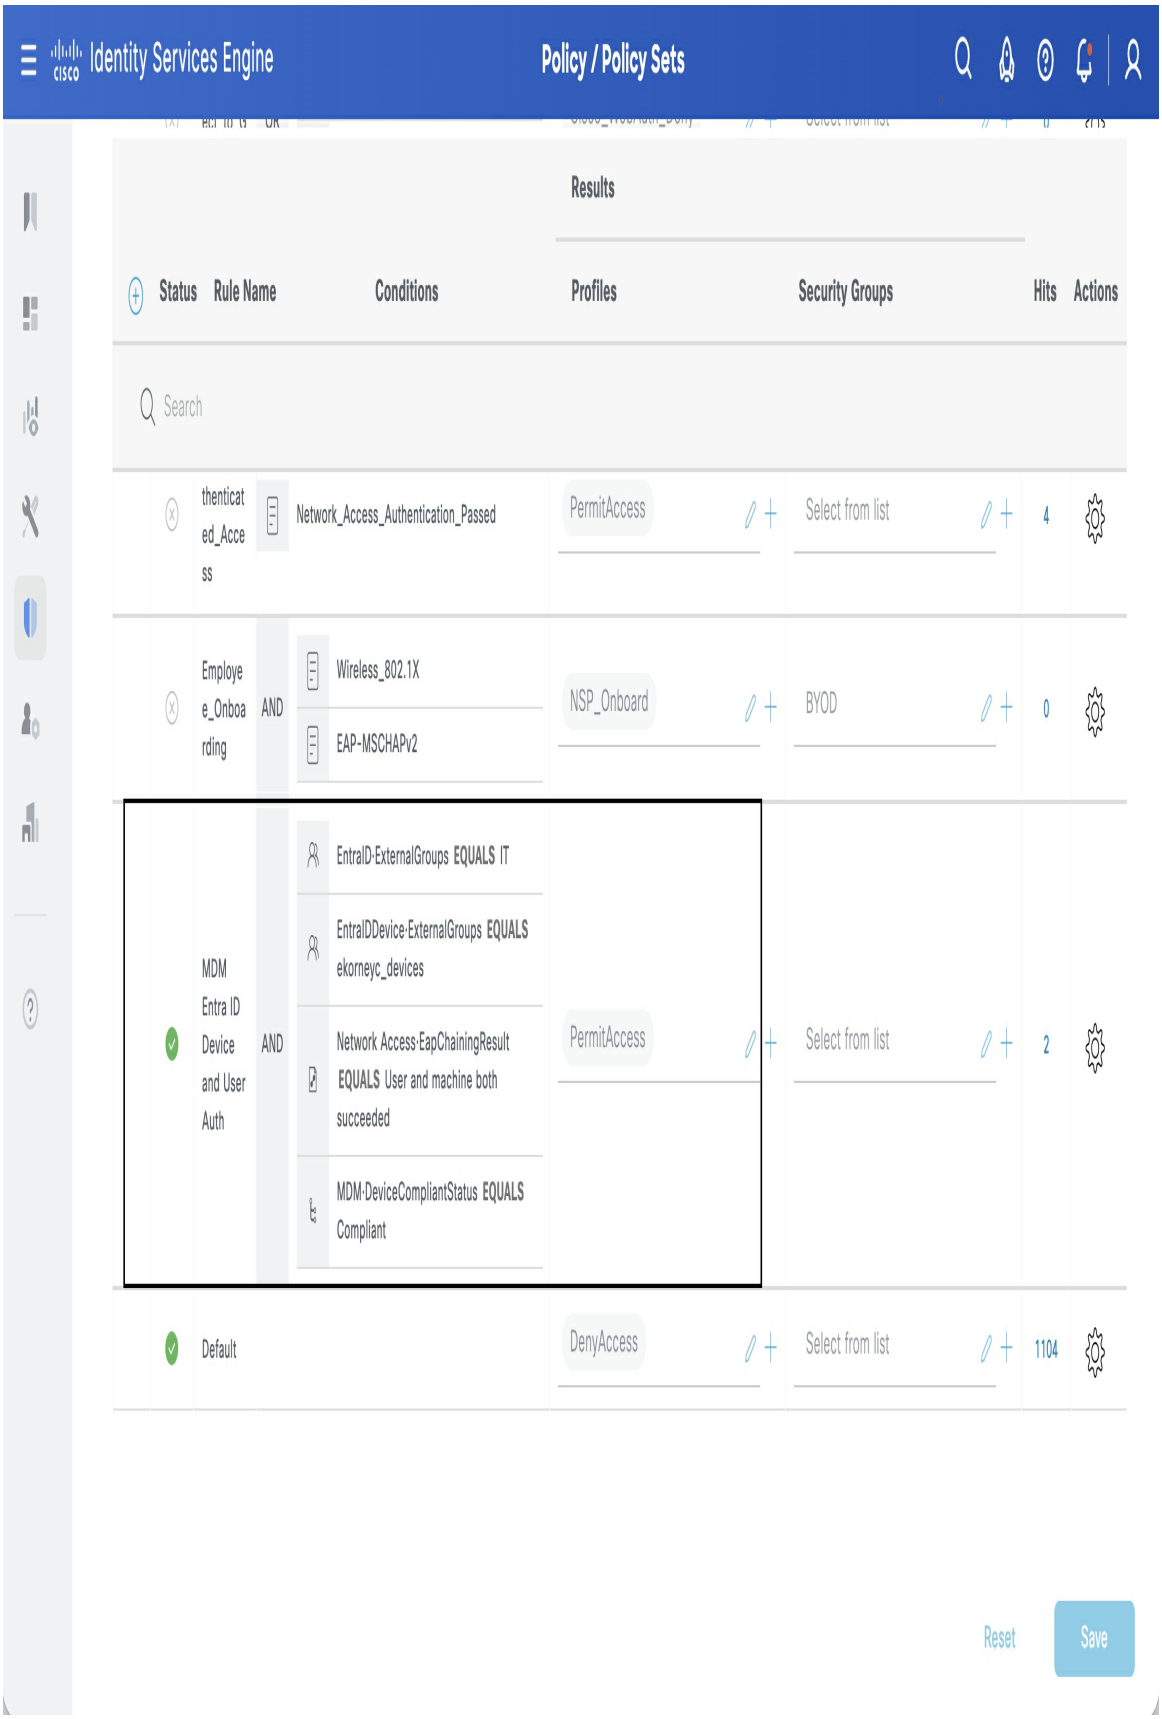

Identity Services Engine
Policy / Policy Sets
Results
Status
Rule Name
Conditions
Profiles
Security Groups
Hits
Actions
Search
thenticat ed_Acce ss
Network_Access_Authentication_Passed
PermitAccess
Select from list
4
Employe e_Onboa rding
AND
Wireless_802.1X
EAP-MSCHAPv2
NSP_Onboard
BYOD
0
MDM Entra ID Device and User Auth
AND
EntraID-ExternalGroups EQUALS IT
EntraIDDevice-ExternalGroups EQUALS ekorneyc_devices
Network Access-EapChainingResult EQUALS User and machine both succeeded
MDM-DeviceCompliantStatus EQUALS Compliant
PermitAccess
Select from list
2
Default
DenyAccess
Select from list
1104
Reset
Save

Figure 20-16 *An Authorization Policy That Is Expected to Be Matched*

Let's examine the detailed authentication report for the failure, as shown in Figure 20-17. Several important points are evident. First, although Access-Reject was sent, the authorization policy matched is Default. Second, the failure reason is "Rejected per authorization profile."

Technet24

Cisco ISE

Overview

Event	5400 Authentication failed
Username	alice@cxsecurity.onmicrosoft.com,b69c5040-730d-4ab1-98ef-d36eee53ab4c
Endpoint Id	00:50:56:A8:96:0E
Endpoint Profile	Microsoft-Workstation
Authentication Policy	Default >> Dot1X
Authorization Policy	Default >> Default
Authorization Result	DenyAccess

Authentication Details

Source Timestamp	Jul 07, 2026 08:56:19.000 PM
Received Timestamp	Jul 07, 2026 08:56:19.000 PM
Policy Server	ISE35-1ek
Event	5400 Authentication failed
Failure Reason	15039 Rejected per authorization profile

Figure 20-17 *A Detailed Authentication Report for the Failed Scenario*

The information in this report indicates that the authentication process was completed successfully, allowing the system to proceed to authorization. You can further confirm successful authentication in the **Steps** section, as shown in Figure 20-18.

15041	Evaluating Identity Policy	1
22072	Selected identity source sequence - All_User_ID_Stores	2
22070	Identity name is taken from certificate attribute	0
22037	Authentication Passed	0
12528	Inner EAP-TLS authentication succeeded	0

Figure 20-18 *Authentication Passed from the Steps Section in the Detailed Authentication Report*

Note

Certain exception scenarios allow authorization to proceed even after failed authentication. This approach can be used, for example, for guest access; however, this is not applicable in this case.

Both facts help you restate the problem description. User and machine authentication are successful; however, the proper authorization policy is not matched. As a result, the default policy is matched and Access-

Technet24

Reject is sent.

The most probable cause for this type of issue is that the conditions for authorization do not match. To verify this, examine the detailed authentication report and assess each condition individually.

The first condition is supposed to return the IT Entra ID Group for the user. In the detailed authentication report, only the ISE VPN Group was returned, as shown in Figure 20-19. Because this is an AND condition, it must match for the rule to apply. The probable cause in this scenario is incorrect group configuration on the Entra ID side.

Best Practice

Evaluate conditions from top to bottom. If the first condition fails, subsequent calls are not made. For example, in this scenario, the MDM lookup was not performed because the rule failed to match at an earlier step.

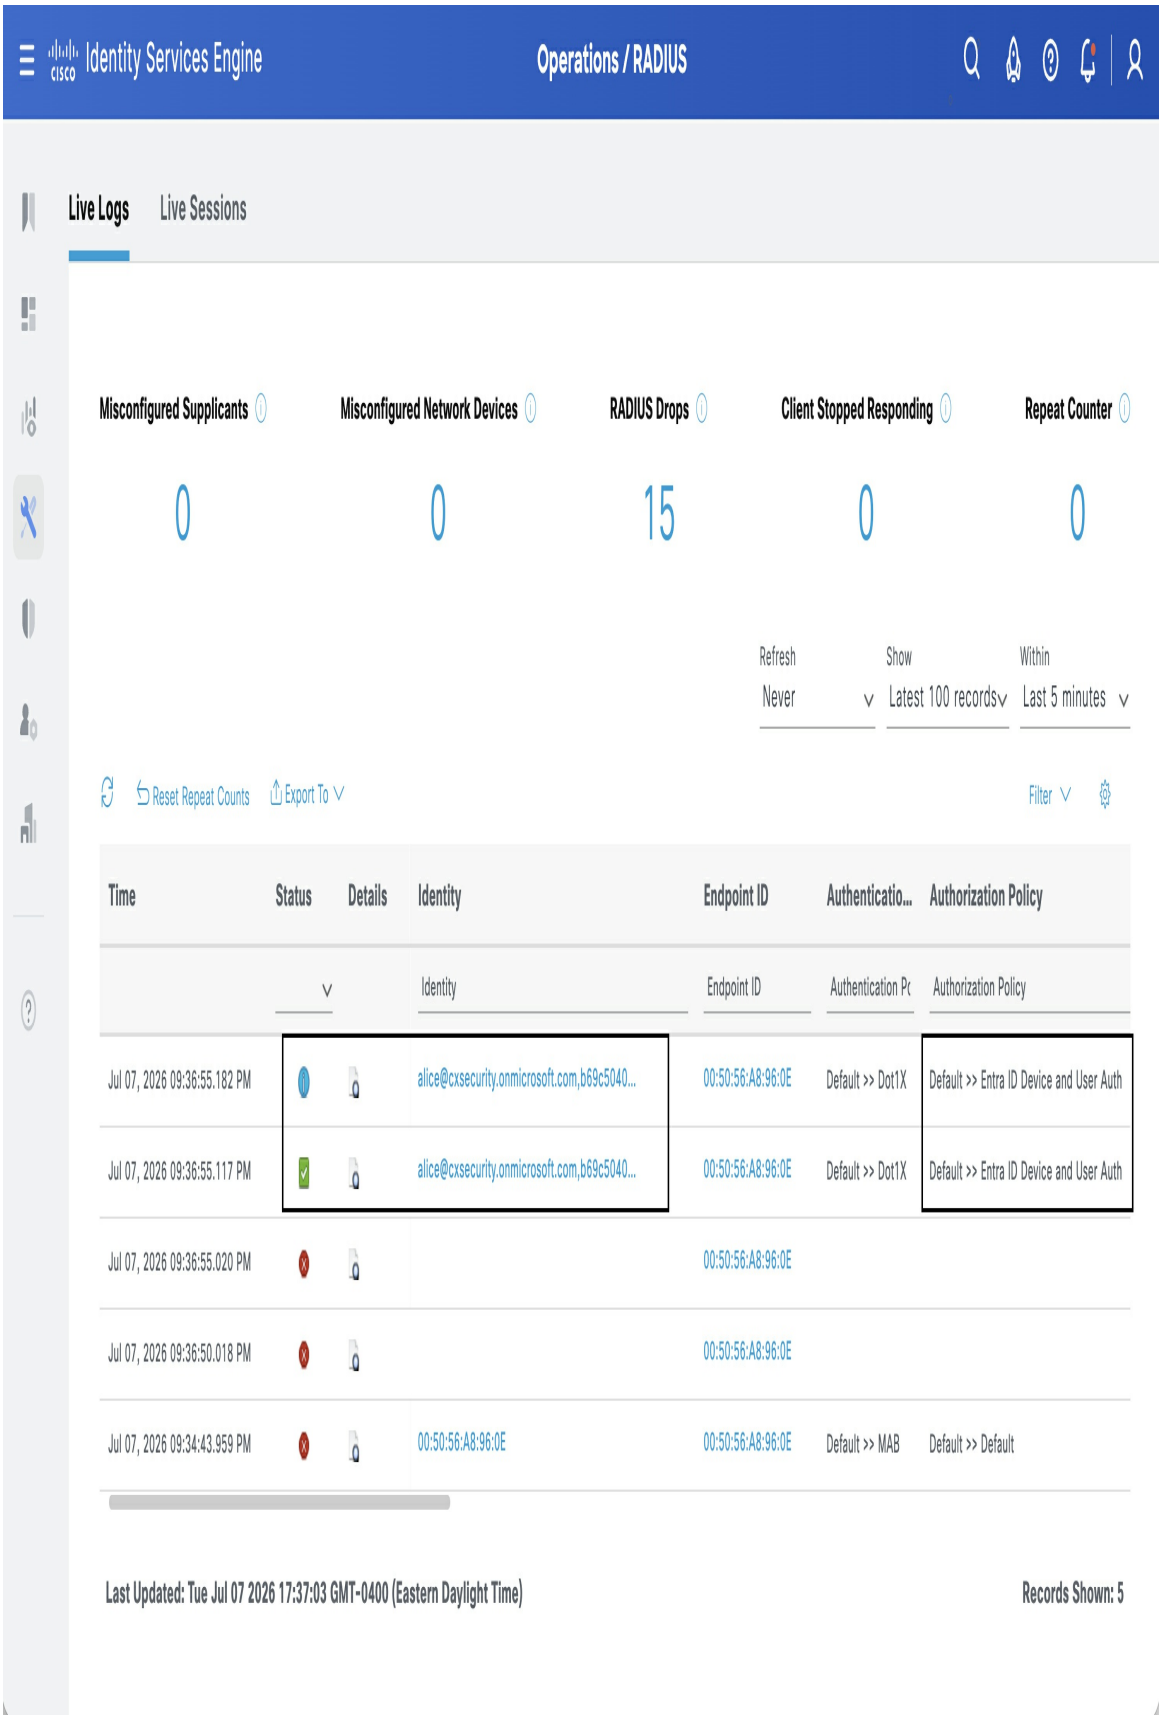
	e3
ISEPolicySetName	Default
Subject Alternative Name - URI	ID:Microsoft Device Manager:GUID:631405e4-6924-42d7-99b9-9780d26fa7b7
ConfigVersionId	95
Issuer - Domain Component	example
Issuer - Domain Component	com
UseCase	Eap Chaining
DTLSSupport	Unknown
AD-Groups-Names	ISE VPN Group
DestinationPort	1812
Issuer - Fingerprint SHA-256	2bcf319436a3c8b15093e7263a790ac2a5f18922bea97d3aecc62011df728235
Subject Alternative Name	ID:Microsoft Device Manager:GUID:631405e4-6924-42d7-99b9-9780d26fa7b7
EapChainingResult	User and machine both succeeded
NetworkDeviceProfileId	b0699505-3150-4215-a80e-6753d45bf56c
IPSEC	IPSEC#Is IPSEC Device#No
ExternalGroups	664f91cf-a065-49df-9410-78d571b3eed9
Extended Key Usage - OID	1.3.6.1.4.1.311.10.3.4
Extended Key Usage - OID	1.3.6.1.5.5.7.3.4

	response	
12524	Extracted EAP-Response containing EAP-TLS challenge-response for inner method and accepting EAP-TLS as negotiated	1
34166	Open secure connection with TLS peer	0
12800	Extracted first TLS record; TLS handshake started	0
12545	Client requested EAP-TLS session ticket	0
12546	The EAP-TLS session ticket received from supplicant. Inner EAP-TLS does not support stateless session resume. Performing full authentication	1
12805	Extracted TLS ClientHello message	0
12806	Prepared TLS ServerHello message	0
12807	Prepared TLS Certificate message	1
12808	Prepared TLS ServerKeyExchange message	26
12809	Prepared TLS CertificateRequest message	0
12810	Prepared TLS ServerDone message	0
12527	Prepared EAP-Request for inner method with another EAP-TLS challenge	0
11596	Prepared EAP-Request with another TEAP challenge	1
11006	Returned RADIUS Access-Challenge	1
11001	Received RADIUS Access-Request	43
11018	RADIUS is re-using an existing session	0
11595	Extracted EAP-Response containing TEAP challenge-response	1
12526	Extracted EAP-Response for inner method containing TLS challenge-response	1
12527	Prepared EAP-Request for inner method with another EAP-TLS challenge	0

Technet24

Figure 20-19 *AD Groups Returned During Authorization in the Detailed Authentication Report*

After the Entra ID was checked, it was confirmed that the user did not have the correct IT group assigned. Once the correct group was assigned, authentication succeeded, as shown in Figure 20-20. All other conditions matched without further configuration changes.


Identity Services Engine
Operations / RADIUS
Live Logs
Live Sessions
Misconfigured Supplicants
Misconfigured Network Devices
RADIUS Drops
Client Stopped Responding
Repeat Counter
0
0
15
0
0
Refresh
Never
Show
Latest 100 records
Within
Last 5 minutes
Reset Repeat Counts
Export To
Filter
Time
Status
Details
Identity
Endpoint ID
Authenticatio...
Authorization Policy
Identity
Endpoint ID
Authentication Pc
Authorization Policy
Jul 07, 2026 09:36:55.182 PM
alice@cxsecurity.onmicrosoft.com,b69c5040...
00:50:56:A8:96:0E
Default >> Dot1X
Default >> Entra ID Device and User Auth
Jul 07, 2026 09:36:55.117 PM
alice@cxsecurity.onmicrosoft.com,b69c5040...
00:50:56:A8:96:0E
Default >> Dot1X
Default >> Entra ID Device and User Auth
Jul 07, 2026 09:36:55.020 PM
00:50:56:A8:96:0E
Jul 07, 2026 09:36:50.018 PM
00:50:56:A8:96:0E
Jul 07, 2026 09:34:43.959 PM
00:50:56:A8:96:0E
00:50:56:A8:96:0E
Default >> MAB
Default >> Default
Last Updated: Tue Jul 07 2026 17:37:03 GMT-0400 (Eastern Daylight Time)
Records Shown: 5

Technet24

Figure 20-20 *Successful Authentication from Live Logs*

Diagnostic Tools

ISE offers a significant number of tools for troubleshooting. Most of these tools are located under the **Operations > Troubleshoot** section of the **ISE** menu. Let's explore the most critical tools and examine how to use them in real-life scenarios.

TCP Dump

TCP Dump on ISE is a tool that helps you identify problems at the packet level. It enables you to capture packets from any interface on any ISE node within the deployment.

When you access **Operations > Troubleshoot > Diagnostic Tools > TCP Dump** for the first time, you need to add the capture profile by selecting **Add**. As shown in Figure 20-21, multiple configuration parameters are available. For most cases, selecting the appropriate PSN and interface is sufficient.

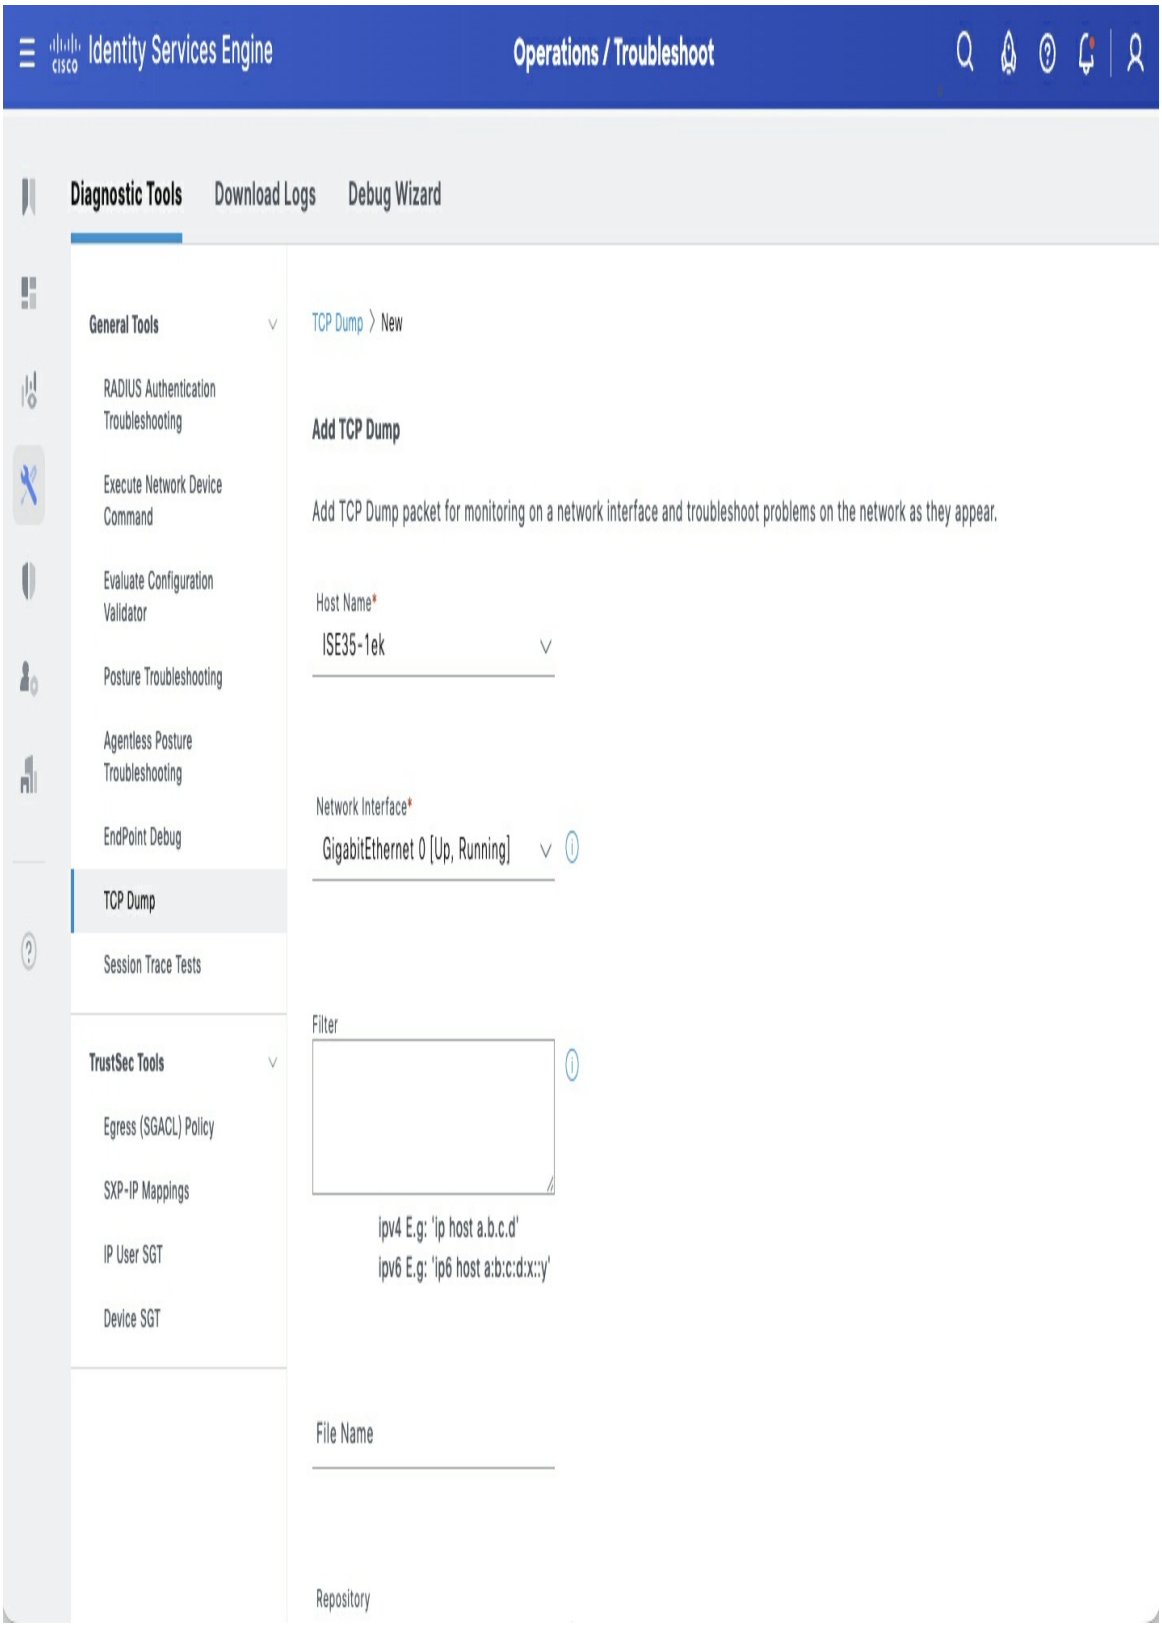


Figure 20-21 *TCP Dump Settings*

Technet24

Note

By default, when you select the ISE node, the network interface is automatically set to GigabitEthernet 0, which suits most cases. If you need to capture traffic on another interface, select it from the **Network Interface** drop-down menu.

Optionally, you can name the capture files, choose the repository for saving the capture, adjust the file size, increase the number of files, or set a time limit.

Once the capture is configured, you can either click **Save** or **Save and Run**, depending on whether you are ready to perform the test that TCP Dump will capture.

You can reuse the same TCP Dump capture profiles multiple times. Start or stop the capture directly from the main **TCP Dump** menu, as shown in Figure 20-22.

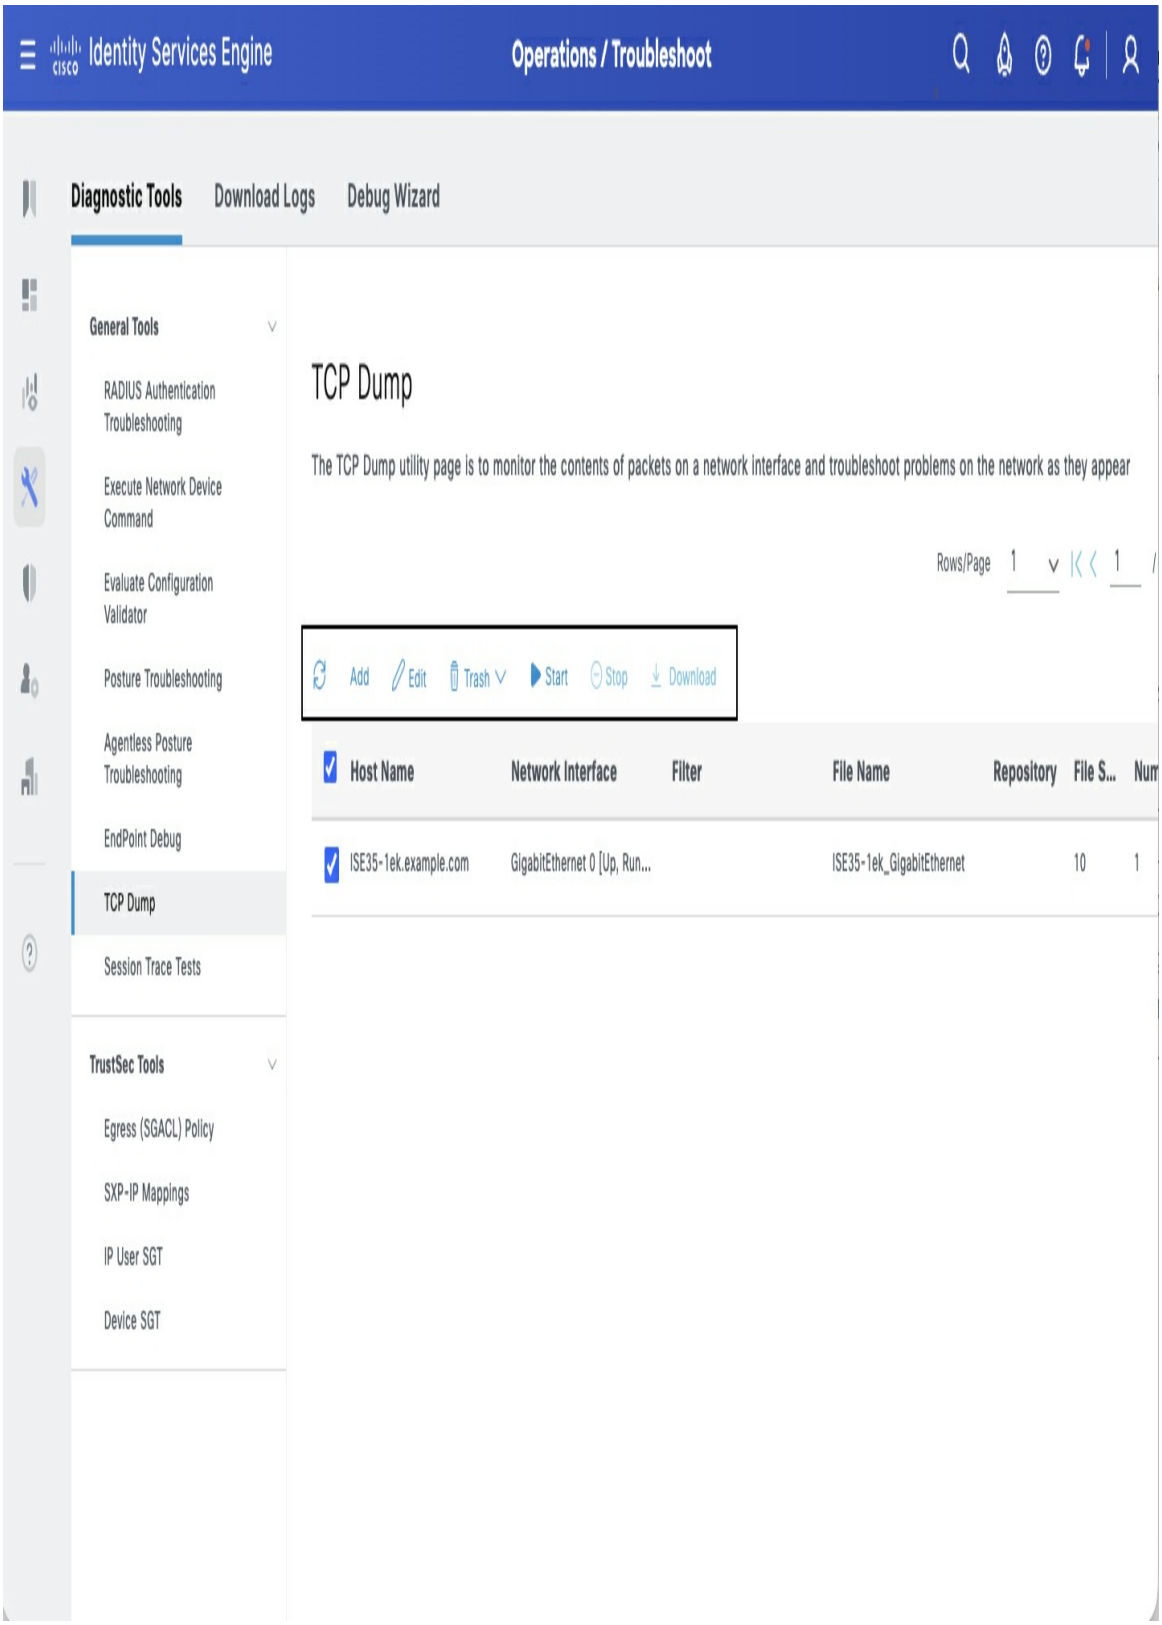


Figure 20-22 *TCP Dump Menu*

Technet24

Once TCP Dump starts, it runs until the maximum allowed time or file size, as defined in the configuration. You can manually stop the TCP Dump collection by using the **Stop** button. To download the PCAP file, click the **Download** button.

> **Note**
>
> You may need to click the **Refresh** icon in the same menu after stopping the capture before the **Download** icon becomes available.

There are several use cases where TCP Dump becomes extremely important:

- You troubleshoot RADIUS fragmentation issues and want to confirm which packets reach ISE.
- You troubleshoot profiling and need to verify which probe data arrives at ISE.
- You troubleshoot internode communication and capture packets from multiple nodes simultaneously to inspect for any packet loss.
- You troubleshoot external integrations and analyze SSL handshakes or general packet flows.
- You troubleshoot SMTP, SMS gateways, external syslog servers, and similar services.

You can troubleshoot nearly any communication to and from ISE to some extent using TCP Dump.

Let's examine a practical example of how TCP Dump can help resolve integration failures:

> **Problem Description: Test Connection under MDM Server configuration is failing, with the error shown in Figure 20-23.**

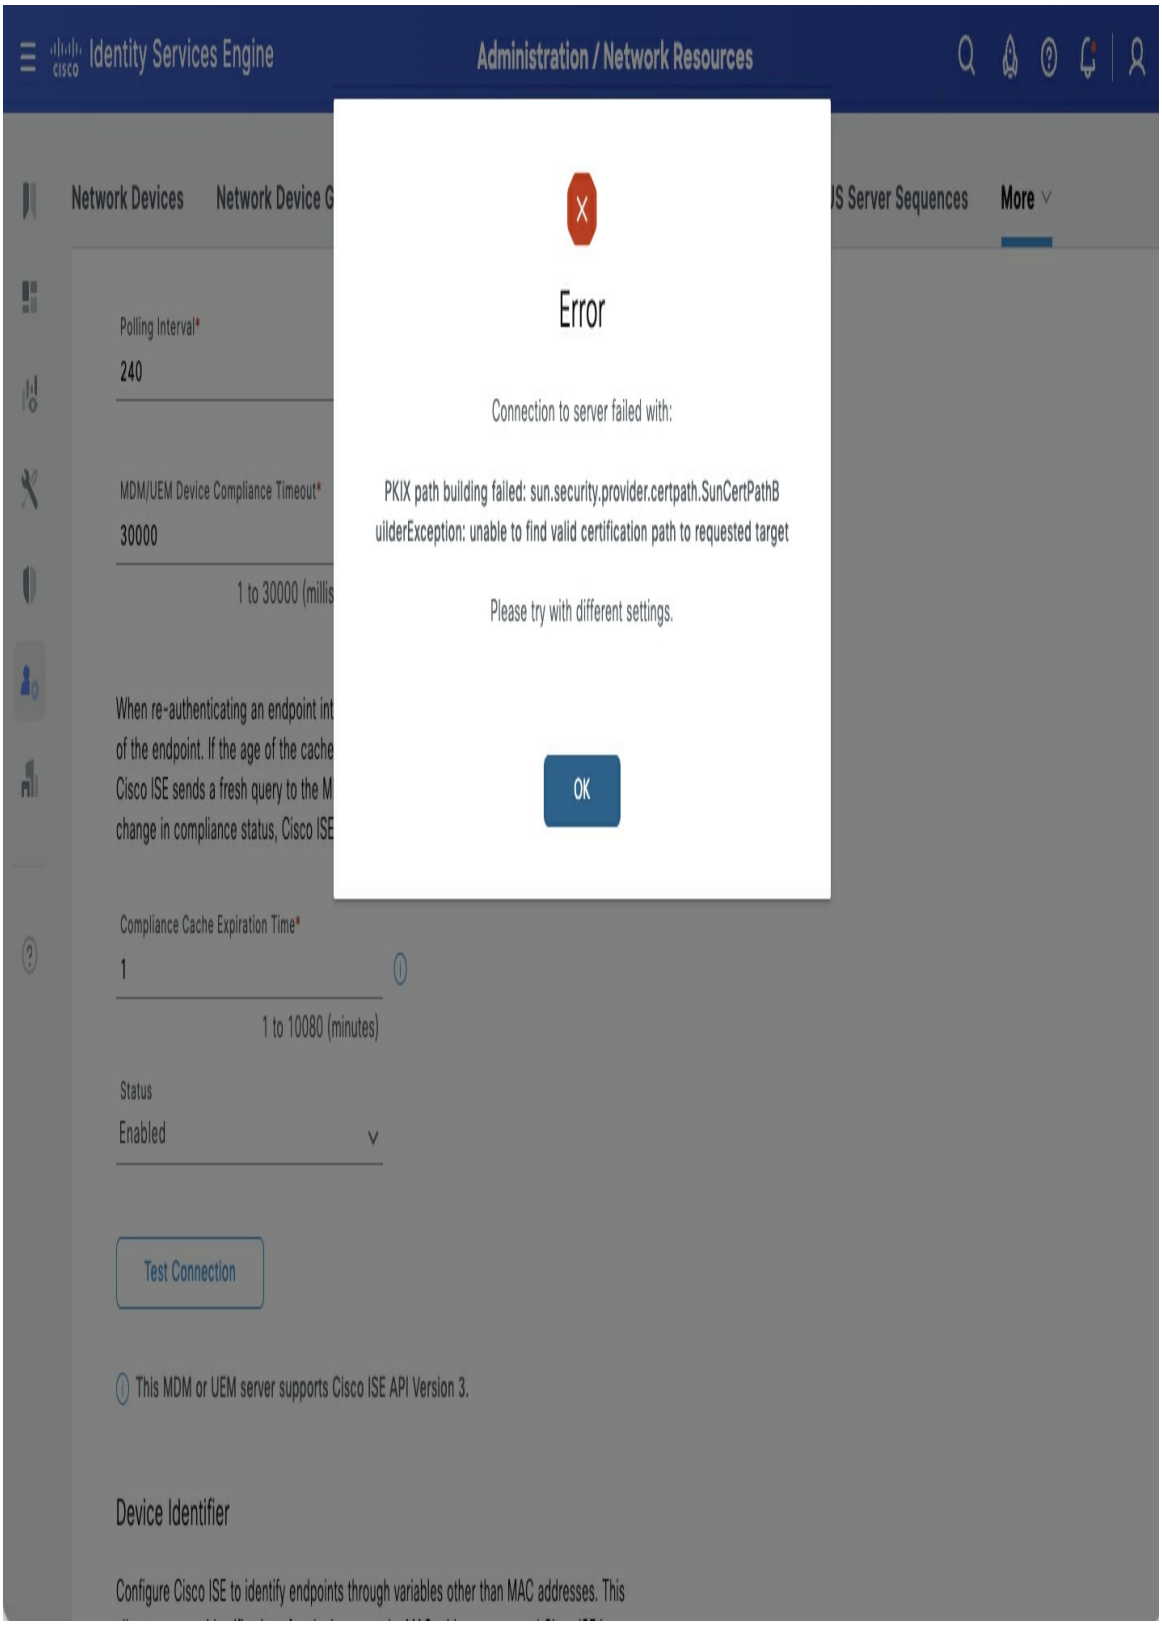


Figure 20-23 *MDM Integration Error*

Technet24

The exact text of the error message is

```
Connection to server failed with:
PKIX path building failed: sun.security.provider.certpath.SunCertPat
unable to find valid certification path to requested target
Please try with different settings.
```

The error message indicates a certificate path error, which implies that ISE cannot trust the certificates sent from the MDM server. The logical way to validate this theory is by capturing packets during the test connection using TCP Dump and analyzing them. The steps are as follows:

Step 1. Start TCP Dump.

Step 2. Click **Test Connection** under **MDM Configuration**.

Step 3. Stop TCP Dump and click the **Refresh for Download** button to become active.

Step 4. Download the TCP Dump and open it in Wireshark.

The next step is to filter the Wireshark capture. You can use **tls.handshake.type==1** to filter all Client Hello packets, as shown in Figure 20-24.

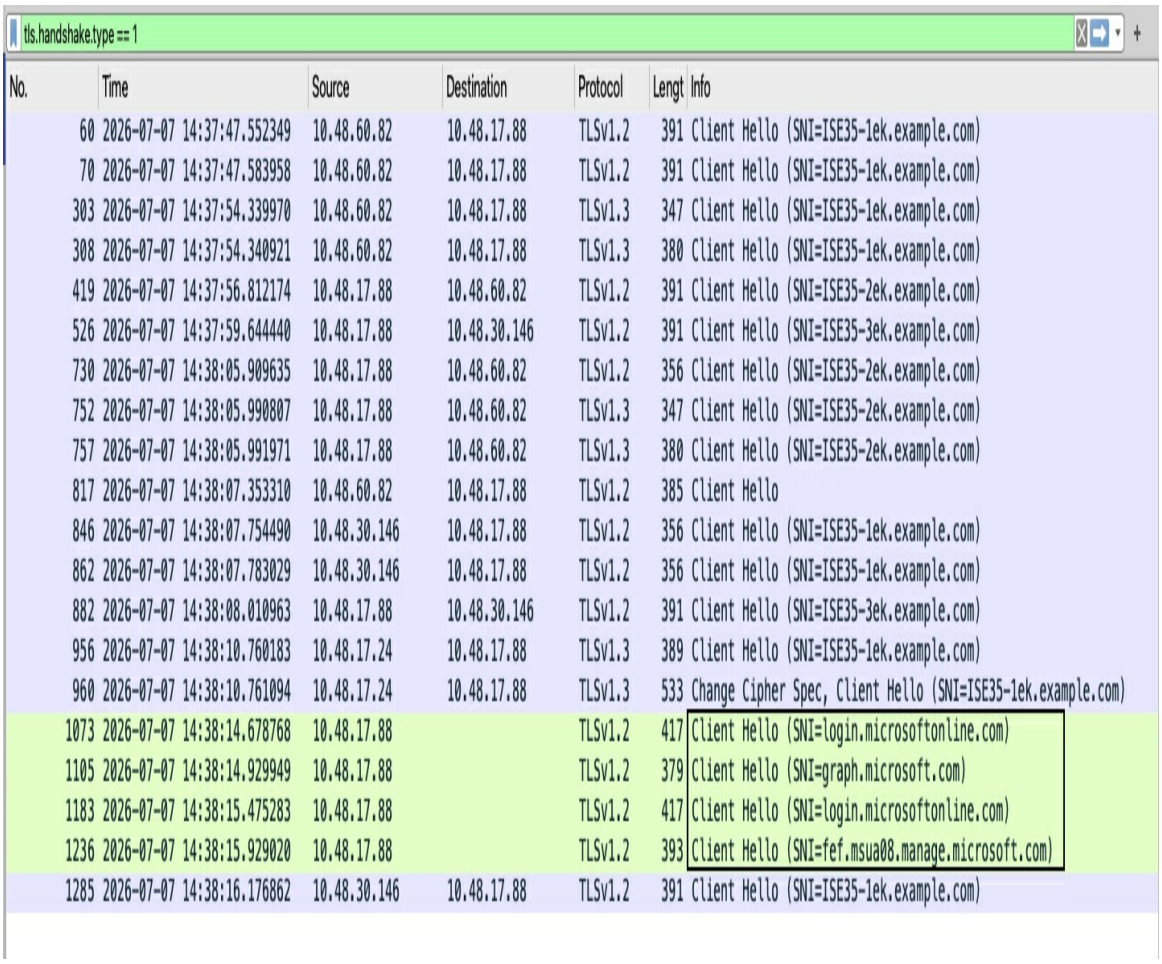


Figure 20-24 *The **tls.handshake.type==1** filter with Wireshark*

ISE attempted at least four SSL handshakes with different Microsoft URLs during this integration. To analyze how each handshake concluded, right-click each **Client Hello** packet and select **Follow > TCP Stream**, as shown in Figure 20-25.

Technet24

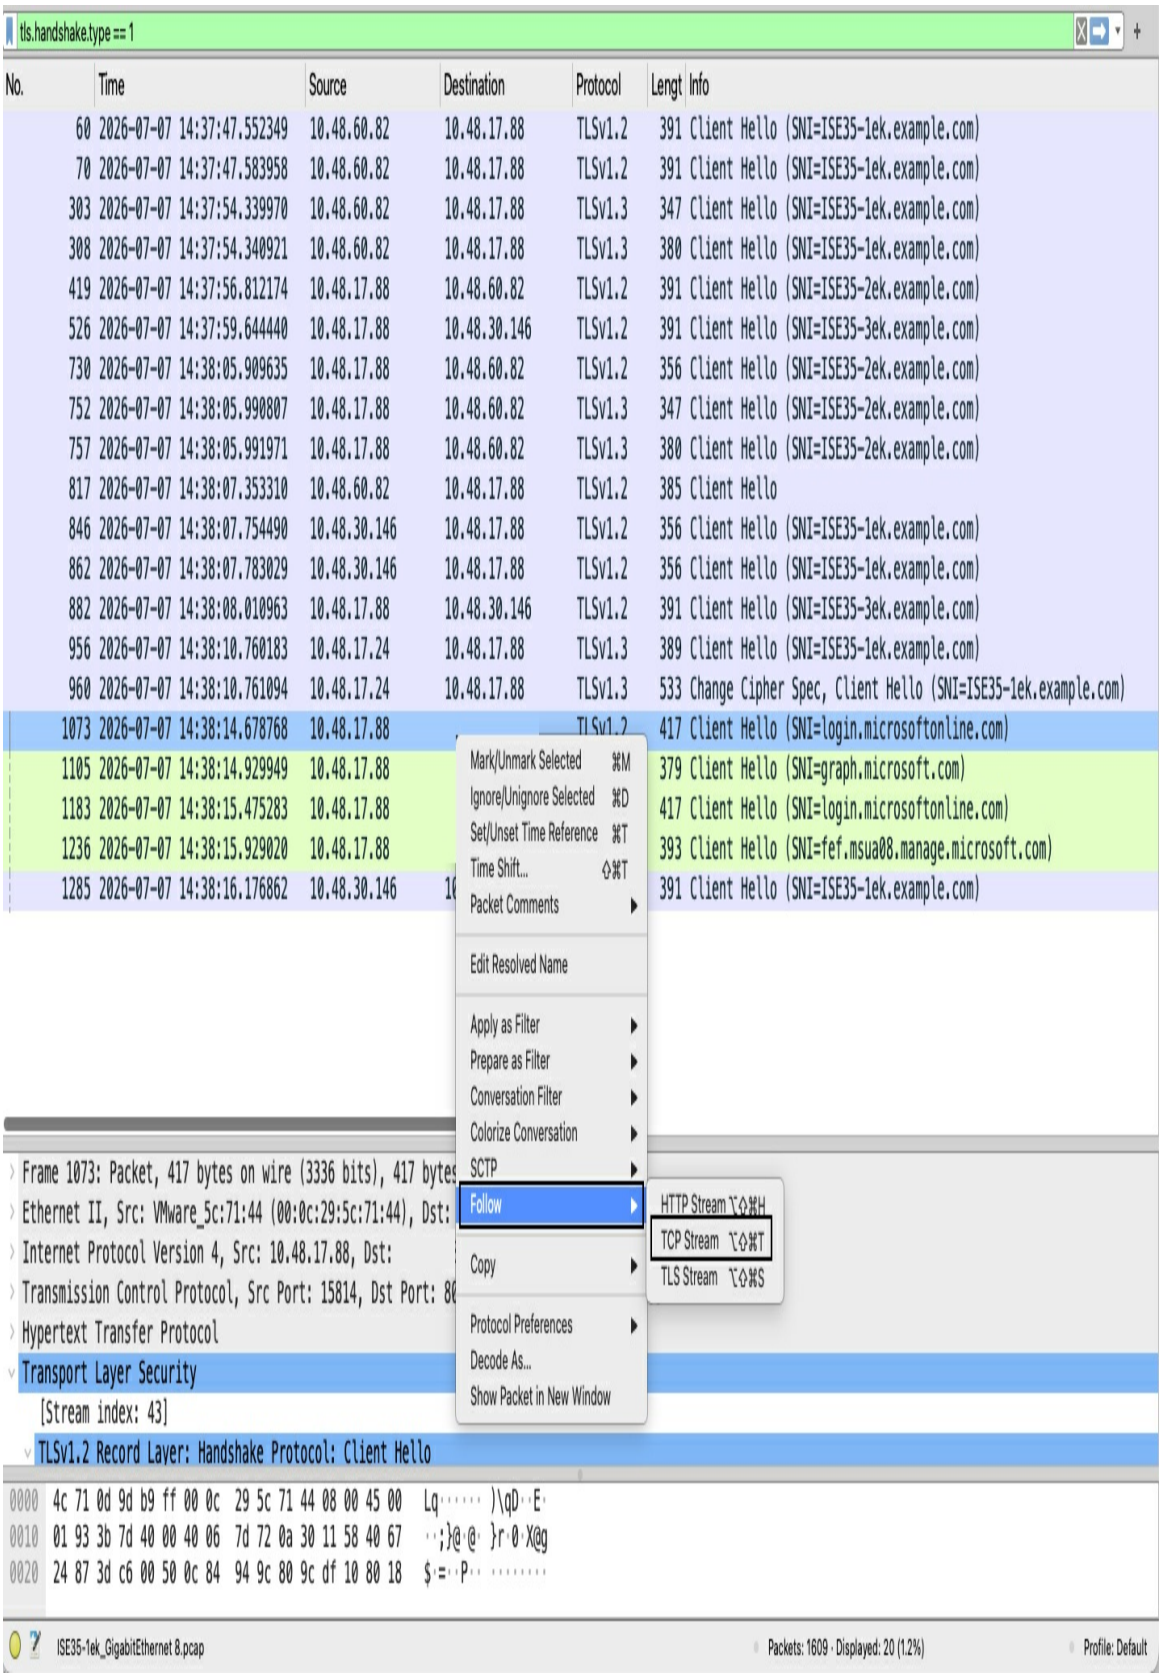

tls.handshake.type == 1
No. Time Source Destination Protocol Lengt Info
60 2026-07-07 14:37:47.552349 10.48.60.82 10.48.17.88 TLSv1.2 391 Client Hello (SNI=ISE35-1ek.example.com)
70 2026-07-07 14:37:47.583958 10.48.60.82 10.48.17.88 TLSv1.2 391 Client Hello (SNI=ISE35-1ek.example.com)
303 2026-07-07 14:37:54.339970 10.48.60.82 10.48.17.88 TLSv1.3 347 Client Hello (SNI=ISE35-1ek.example.com)
308 2026-07-07 14:37:54.340921 10.48.60.82 10.48.17.88 TLSv1.3 380 Client Hello (SNI=ISE35-1ek.example.com)
419 2026-07-07 14:37:56.812174 10.48.17.88 10.48.60.82 TLSv1.2 391 Client Hello (SNI=ISE35-2ek.example.com)
526 2026-07-07 14:37:59.644440 10.48.17.88 10.48.30.146 TLSv1.2 391 Client Hello (SNI=ISE35-3ek.example.com)
730 2026-07-07 14:38:05.909635 10.48.17.88 10.48.60.82 TLSv1.2 356 Client Hello (SNI=ISE35-2ek.example.com)
752 2026-07-07 14:38:05.990807 10.48.17.88 10.48.60.82 TLSv1.3 347 Client Hello (SNI=ISE35-2ek.example.com)
757 2026-07-07 14:38:05.991971 10.48.17.88 10.48.60.82 TLSv1.3 380 Client Hello (SNI=ISE35-2ek.example.com)
817 2026-07-07 14:38:07.353310 10.48.60.82 10.48.17.88 TLSv1.2 385 Client Hello
846 2026-07-07 14:38:07.754490 10.48.30.146 10.48.17.88 TLSv1.2 356 Client Hello (SNI=ISE35-1ek.example.com)
862 2026-07-07 14:38:07.783029 10.48.30.146 10.48.17.88 TLSv1.2 356 Client Hello (SNI=ISE35-1ek.example.com)
882 2026-07-07 14:38:08.010963 10.48.17.88 10.48.30.146 TLSv1.2 391 Client Hello (SNI=ISE35-3ek.example.com)
956 2026-07-07 14:38:10.760183 10.48.17.24 10.48.17.88 TLSv1.3 389 Client Hello (SNI=ISE35-1ek.example.com)
960 2026-07-07 14:38:10.761094 10.48.17.24 10.48.17.88 TLSv1.3 533 Change Cipher Spec, Client Hello (SNI=ISE35-1ek.example.com)
1073 2026-07-07 14:38:14.678768 10.48.17.88 TLSv1.2 417 Client Hello (SNI=login.microsoftonline.com)
1105 2026-07-07 14:38:14.929949 10.48.17.88 379 Client Hello (SNI=graph.microsoft.com)
1183 2026-07-07 14:38:15.475283 10.48.17.88 417 Client Hello (SNI=login.microsoftonline.com)
1236 2026-07-07 14:38:15.929020 10.48.17.88 393 Client Hello (SNI=fef.msua08.manage.microsoft.com)
1285 2026-07-07 14:38:16.176862 10.48.30.146 391 Client Hello (SNI=ISE35-1ek.example.com)
Mark/Unmark Selected ⌘M
Ignore/Unignore Selected ⌘D
Set/Unset Time Reference ⌘T
Time Shift... ⇧⌘T
Packet Comments
Edit Resolved Name
Apply as Filter
Prepare as Filter
Conversation Filter
Colorize Conversation
SCTP
Follow
Copy
Protocol Preferences
Decode As...
Show Packet in New Window
HTTP Stream ⌥⇧⌘H
TCP Stream ⌥⇧⌘T
TLS Stream ⌥⇧⌘S
Frame 1073: Packet, 417 bytes on wire (3336 bits), 417 bytes
Ethernet II, Src: VMware_5c:71:44 (00:0c:29:5c:71:44), Dst:
Internet Protocol Version 4, Src: 10.48.17.88, Dst:
Transmission Control Protocol, Src Port: 15814, Dst Port: 80
Hypertext Transfer Protocol
Transport Layer Security
[Stream index: 43]
TLSv1.2 Record Layer: Handshake Protocol: Client Hello
0000 4c 71 0d 9d b9 ff 00 0c 29 5c 71 44 08 00 45 00 Lq·····)\qD··E·
0010 01 93 3b 7d 40 00 40 06 7d 72 0a 30 11 58 40 67 ··;}@·@· }r·0·X@g
0020 24 87 3d c6 00 50 0c 84 94 9c 80 9c df 10 80 18 $·=··P·· ·····
ISE35-1ek_GigabitEthernet 8.pcap
Packets: 1609 · Displayed: 20 (1.2%)
Profile: Default

Figure 20-25 *Following the TCP Stream with Wireshark*

One of the streams resulted in a fatal error, with "Certificate Unknown" sent by the client (ISE), as shown in Figure 20-26.

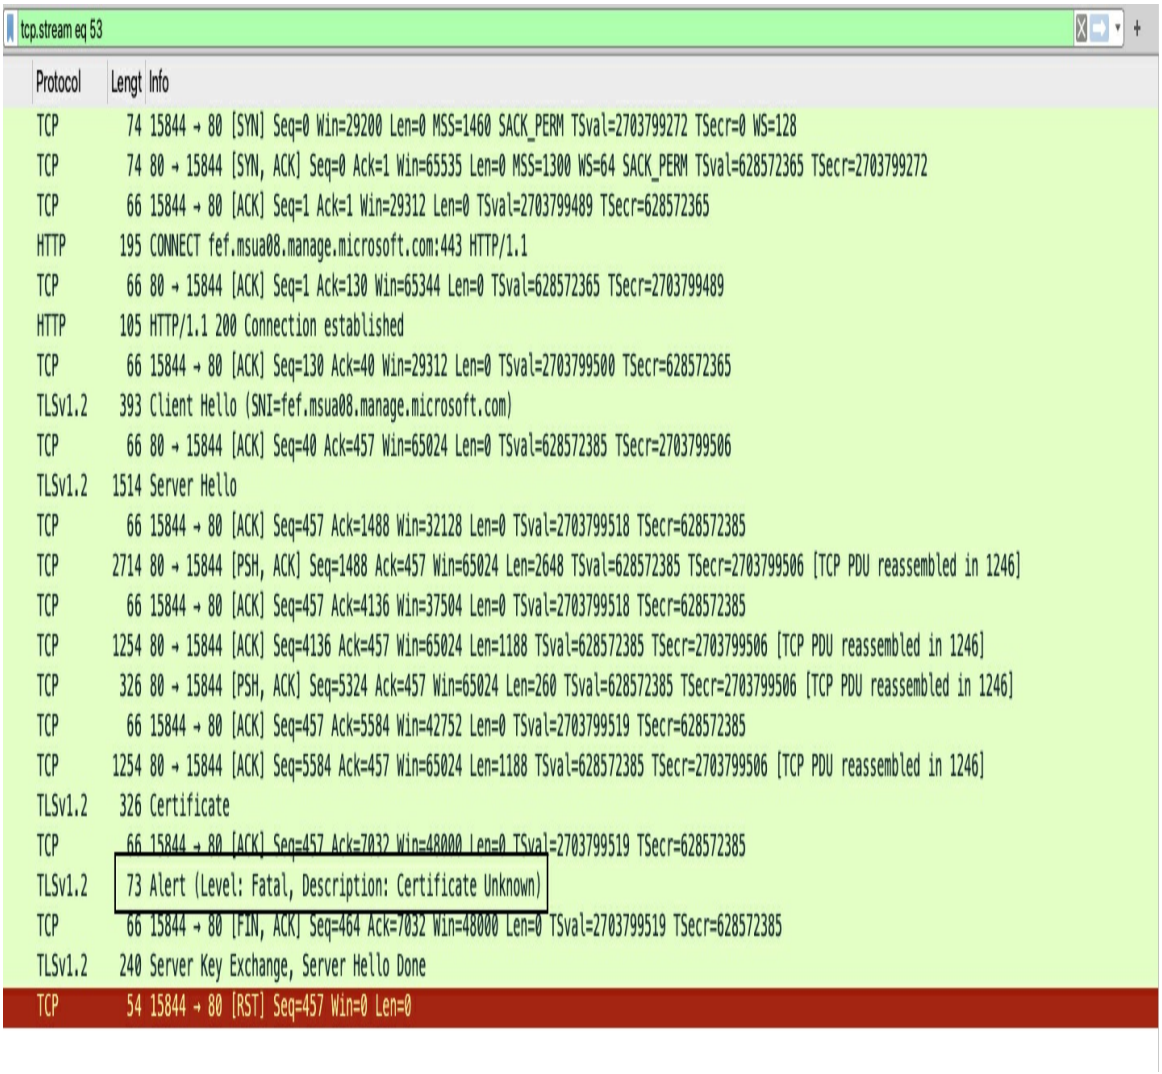
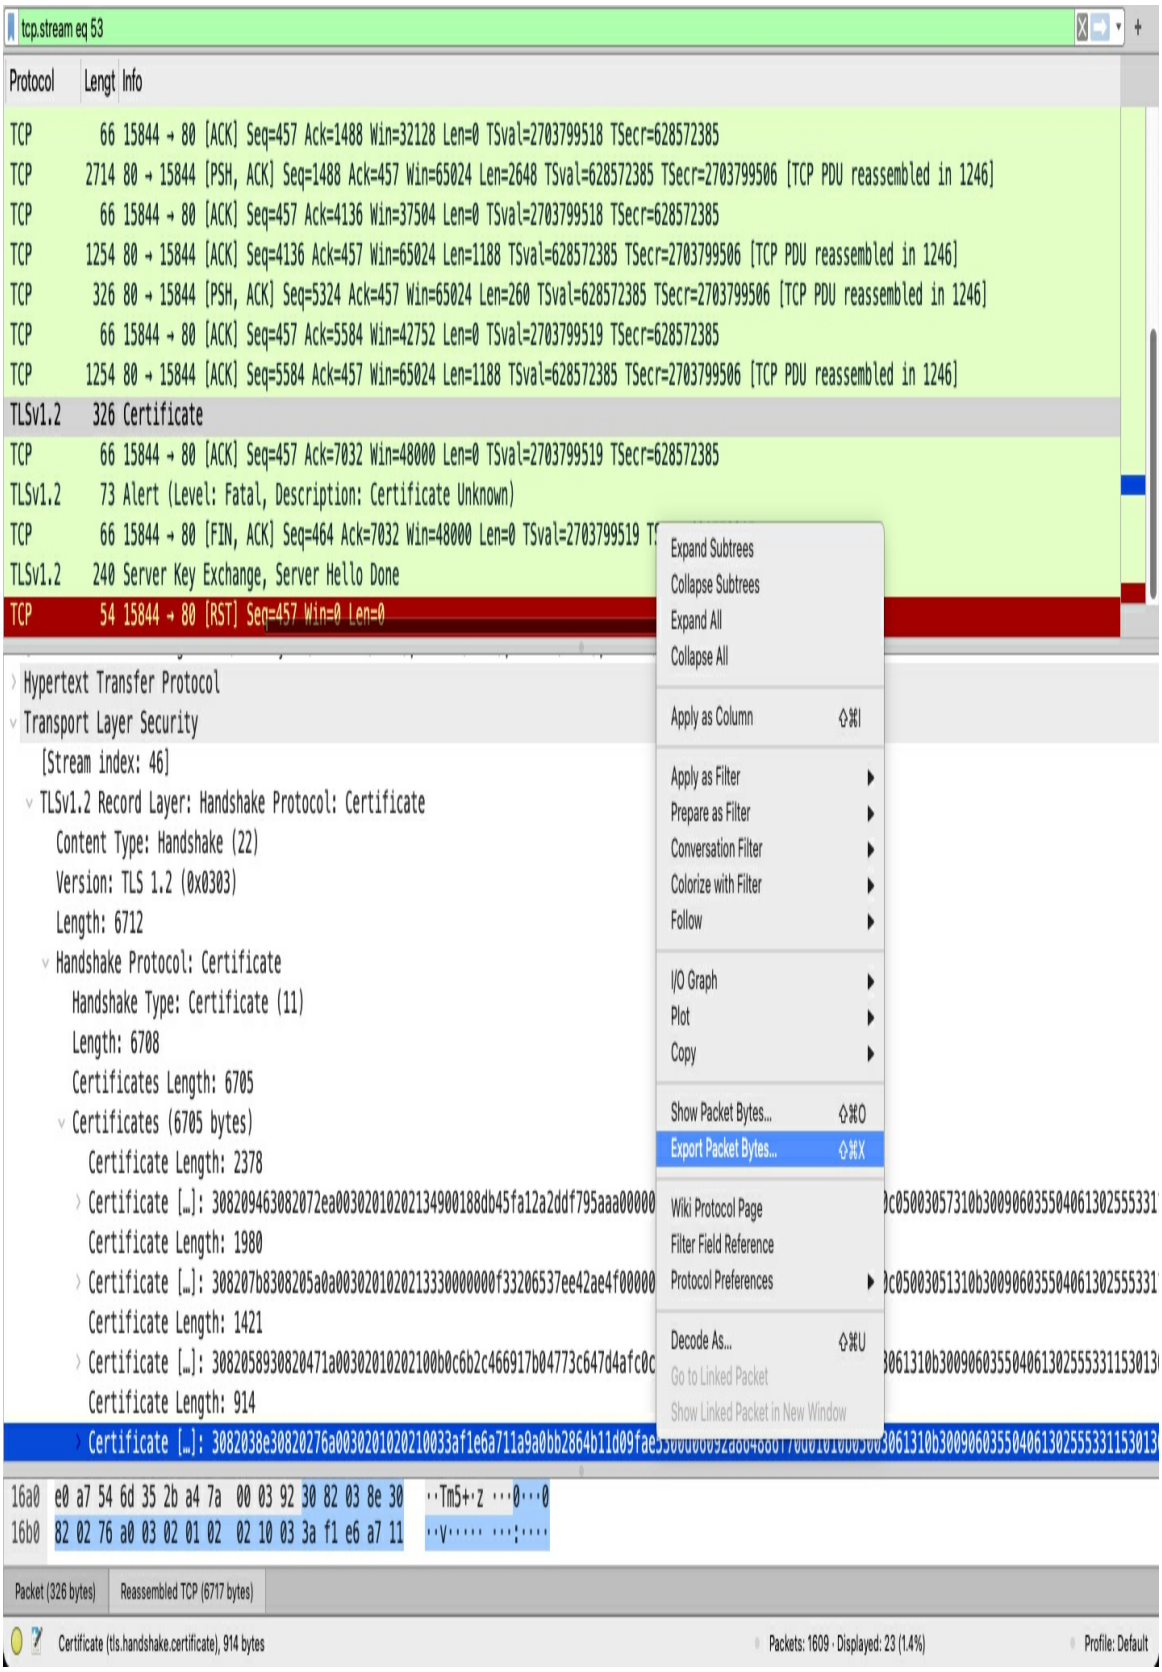


Figure 20-26 *Fatal Alert, Certificate Unknown*

The Certificate Unknown error appears when ISE does not trust the certificate presented by the server. To determine which certificate or certificate chain the server is sending, examine the packet details section in Wireshark. You can also export the certificate directly from Wireshark by right-clicking the certificate and selecting **Export Packet**

Technet24

Bytes, as shown in Figure 20-27.


tcp.stream eq 53
Protocol Lengt Info
TCP 66 15844 → 80 [ACK] Seq=457 Ack=1488 Win=32128 Len=0 TSval=2703799518 TSecr=628572385
TCP 2714 80 → 15844 [PSH, ACK] Seq=1488 Ack=457 Win=65024 Len=2648 TSval=628572385 TSecr=2703799506 [TCP PDU reassembled in 1246]
TCP 66 15844 → 80 [ACK] Seq=457 Ack=4136 Win=37504 Len=0 TSval=2703799518 TSecr=628572385
TCP 1254 80 → 15844 [ACK] Seq=4136 Ack=457 Win=65024 Len=1188 TSval=628572385 TSecr=2703799506 [TCP PDU reassembled in 1246]
TCP 326 80 → 15844 [PSH, ACK] Seq=5324 Ack=457 Win=65024 Len=260 TSval=628572385 TSecr=2703799506 [TCP PDU reassembled in 1246]
TCP 66 15844 → 80 [ACK] Seq=457 Ack=5584 Win=42752 Len=0 TSval=2703799519 TSecr=628572385
TCP 1254 80 → 15844 [ACK] Seq=5584 Ack=457 Win=65024 Len=1188 TSval=628572385 TSecr=2703799506 [TCP PDU reassembled in 1246]
TLSv1.2 326 Certificate
TCP 66 15844 → 80 [ACK] Seq=457 Ack=7032 Win=48000 Len=0 TSval=2703799519 TSecr=628572385
TLSv1.2 73 Alert (Level: Fatal, Description: Certificate Unknown)
TCP 66 15844 → 80 [FIN, ACK] Seq=464 Ack=7032 Win=48000 Len=0 TSval=2703799519
TLSv1.2 240 Server Key Exchange, Server Hello Done
TCP 54 15844 → 80 [RST] Seq=457 Win=0 Len=0
Hypertext Transfer Protocol
Transport Layer Security
[Stream index: 46]
TLSv1.2 Record Layer: Handshake Protocol: Certificate
Content Type: Handshake (22)
Version: TLS 1.2 (0x0303)
Length: 6712
Handshake Protocol: Certificate
Handshake Type: Certificate (11)
Length: 6708
Certificates Length: 6705
Certificates (6705 bytes)
Certificate Length: 2378
Certificate […]: 308209463082072ea003020102021349001 88db45fa12a2ddf795aaa00000
Certificate Length: 1980
Certificate […]: 308207b8308205a0a003020102021333000000f33206537ee42ae4f00000
Certificate Length: 1421
Certificate […]: 3082058930820471a003020102021 00b0c6b2c466917b04773c647d4afc0c
Certificate Length: 914
Certificate […]: 3082038e30820276a003020102021 0033af1e6a711a9a0bb2864b11d09fae
Expand Subtrees
Collapse Subtrees
Expand All
Collapse All
Apply as Column ⇧⌘I
Apply as Filter
Prepare as Filter
Conversation Filter
Colorize with Filter
Follow
I/O Graph
Plot
Copy
Show Packet Bytes... ⇧⌘O
Export Packet Bytes... ⇧⌘X
Wiki Protocol Page
Filter Field Reference
Protocol Preferences
Decode As... ⇧⌘U
Go to Linked Packet
Show Linked Packet in New Window
16a0 e0 a7 54 6d 35 2b a4 7a 00 03 92 30 82 03 8e 30 ··Tm5+·z ···0···0
16b0 82 02 76 a0 03 02 01 02 02 10 03 3a f1 e6 a7 11 ··v····· ···:····
Packet (326 bytes)
Reassembled TCP (6717 bytes)
Certificate (tls.handshake.certificate), 914 bytes
Packets: 1609 · Displayed: 23 (1.4%)
Profile: Default

Technet24

Figure 20-27 *Exporting Certificates from Wireshark*

While installing only the top CA certificate may suffice, the recommended practice is to install the complete certificate chain (excluding the identity certificate). To do so, install the exported certificates in the Trusted Store of ISE and enable **Trust for Authentication of Cisco Services**.

After the certificates were imported, the test connection succeeded, as shown in Figure 20-28. This result confirms the earlier theory.

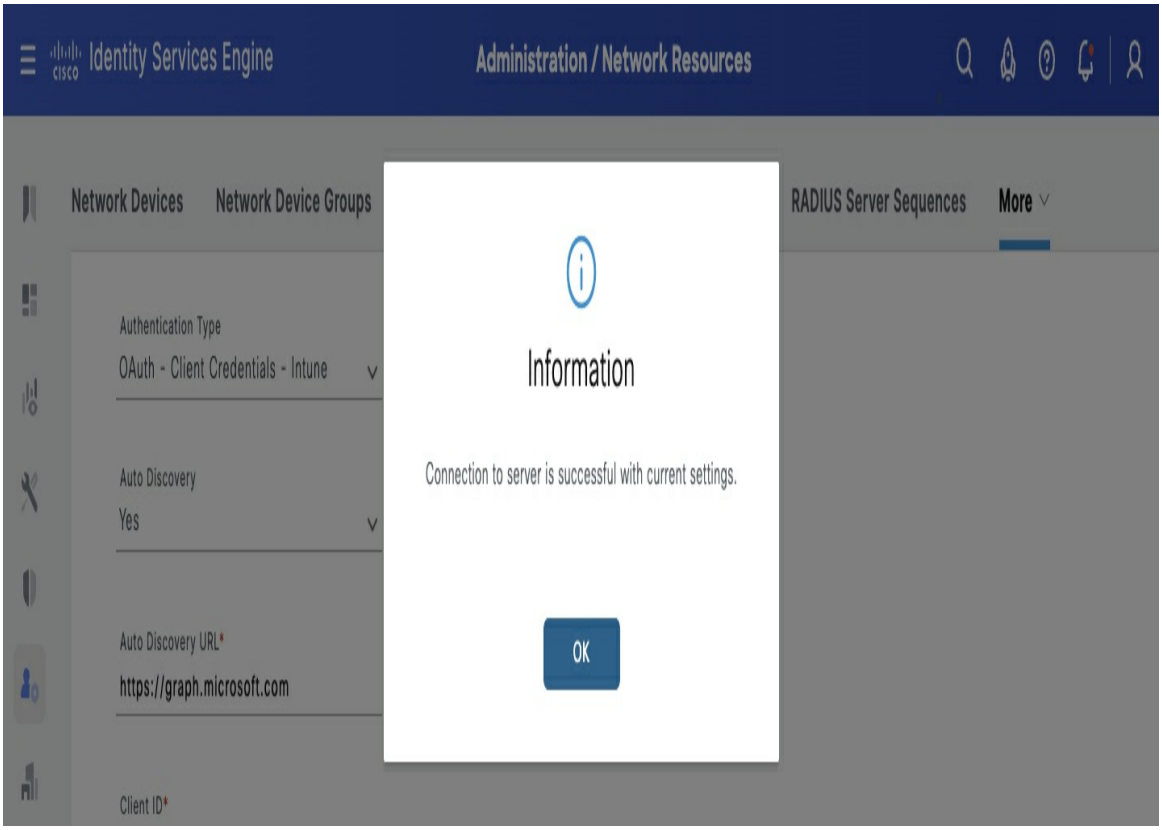


Figure 20-28 *Successful Test Connection of MDM Server*

Debug Wizard

Debugging on ISE is generally considered an advanced troubleshooting technique. This perception stems from the large

number of components involved; each feature on ISE can span multiple components and nodes within the deployment. Therefore, you need to identify which nodes require debugging, determine which debug log files to check, and understand where those debugs are written.

The easiest way to enable debugs on ISE is by running the Debug Wizard. As shown in Figure 20-29, you can select a preconfigured Debug Wizard profile on ISE and choose the nodes of interest by clicking **Debug Nodes**.

Technet24

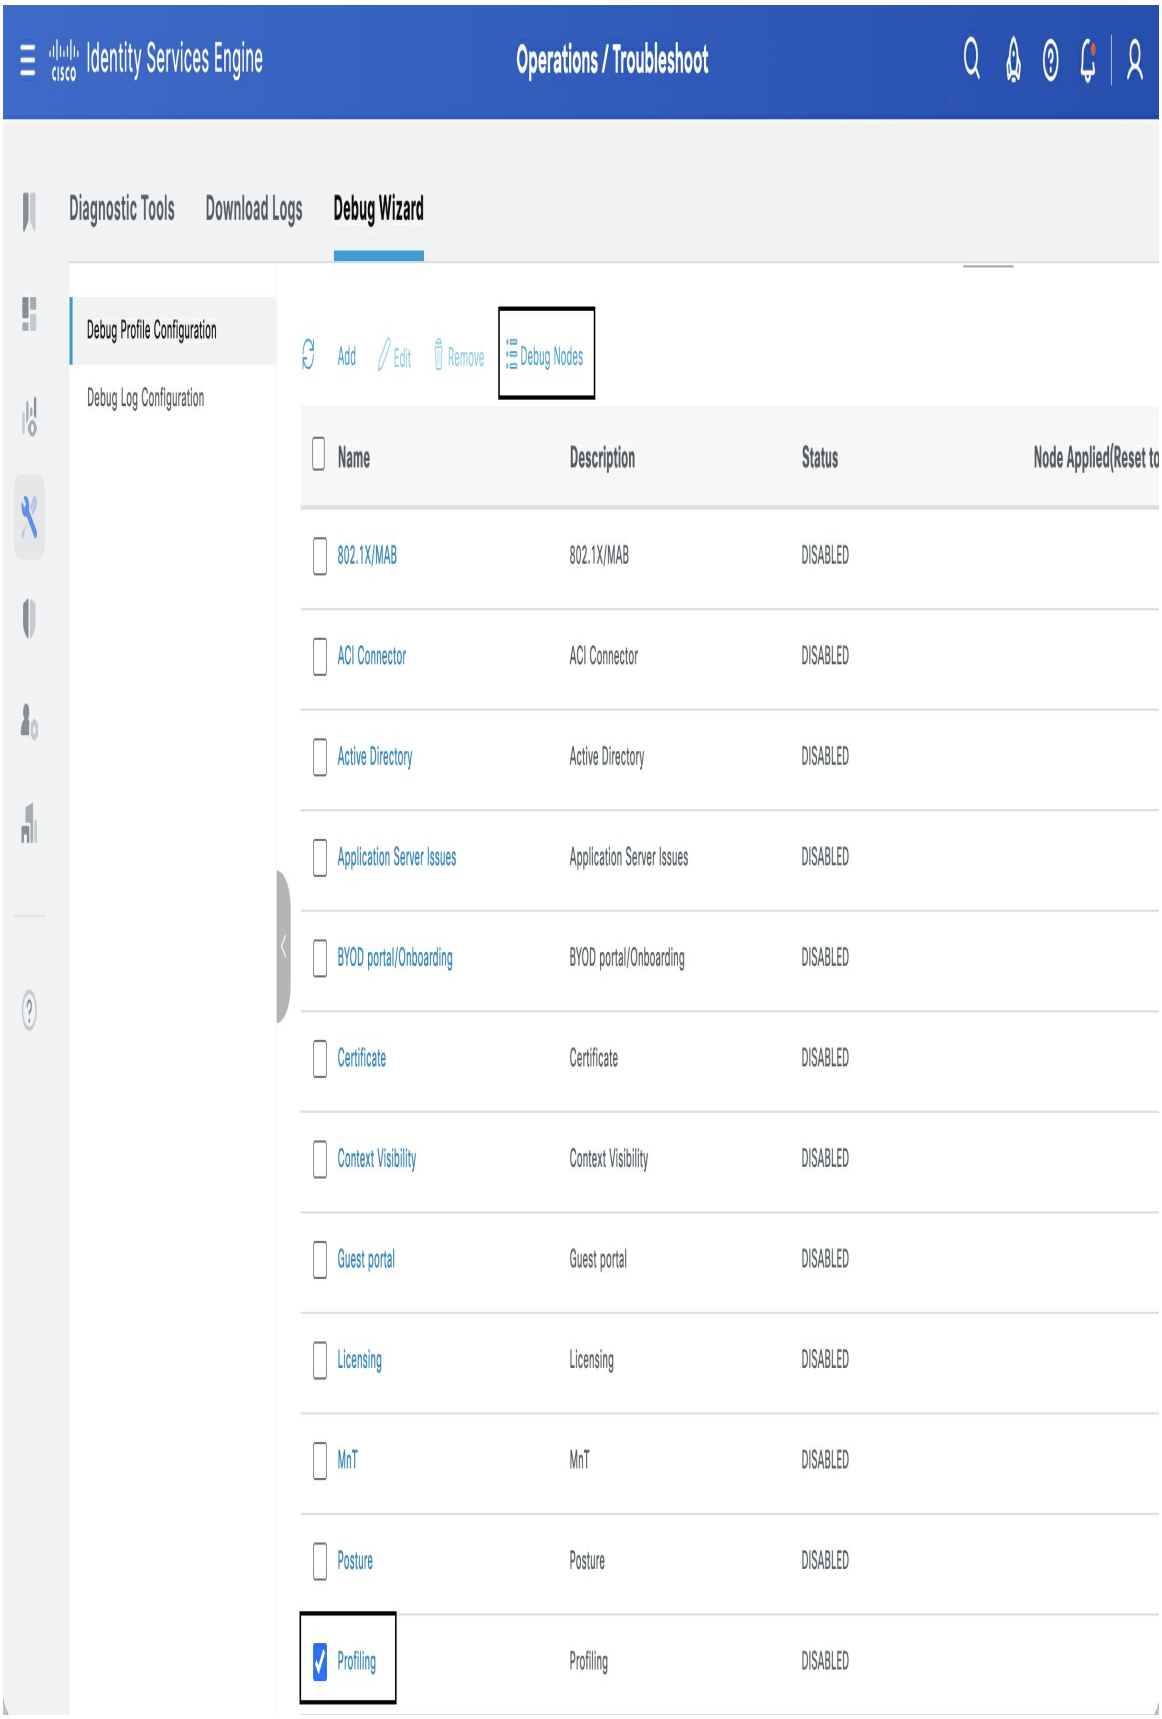

Identity Services Engine
Operations / Troubleshoot
Diagnostic Tools
Download Logs
Debug Wizard
Debug Profile Configuration
Debug Log Configuration
Add
Edit
Remove
Debug Nodes
Name
Description
Status
Node Applied(Reset to
802.1X/MAB
802.1X/MAB
DISABLED
ACI Connector
ACI Connector
DISABLED
Active Directory
Active Directory
DISABLED
Application Server Issues
Application Server Issues
DISABLED
BYOD portal/Onboarding
BYOD portal/Onboarding
DISABLED
Certificate
Certificate
DISABLED
Context Visibility
Context Visibility
DISABLED
Guest portal
Guest portal
DISABLED
Licensing
Licensing
DISABLED
MnT
MnT
DISABLED
Posture
Posture
DISABLED
Profiling
Profiling
DISABLED

Figure 20-29 *Debug Wizard Profiles*

The alternative method for debugging involves enabling individual components with the Debug level of logging from the **Debug Log Settings** menu, as shown in Figure 20-30.

Technet24

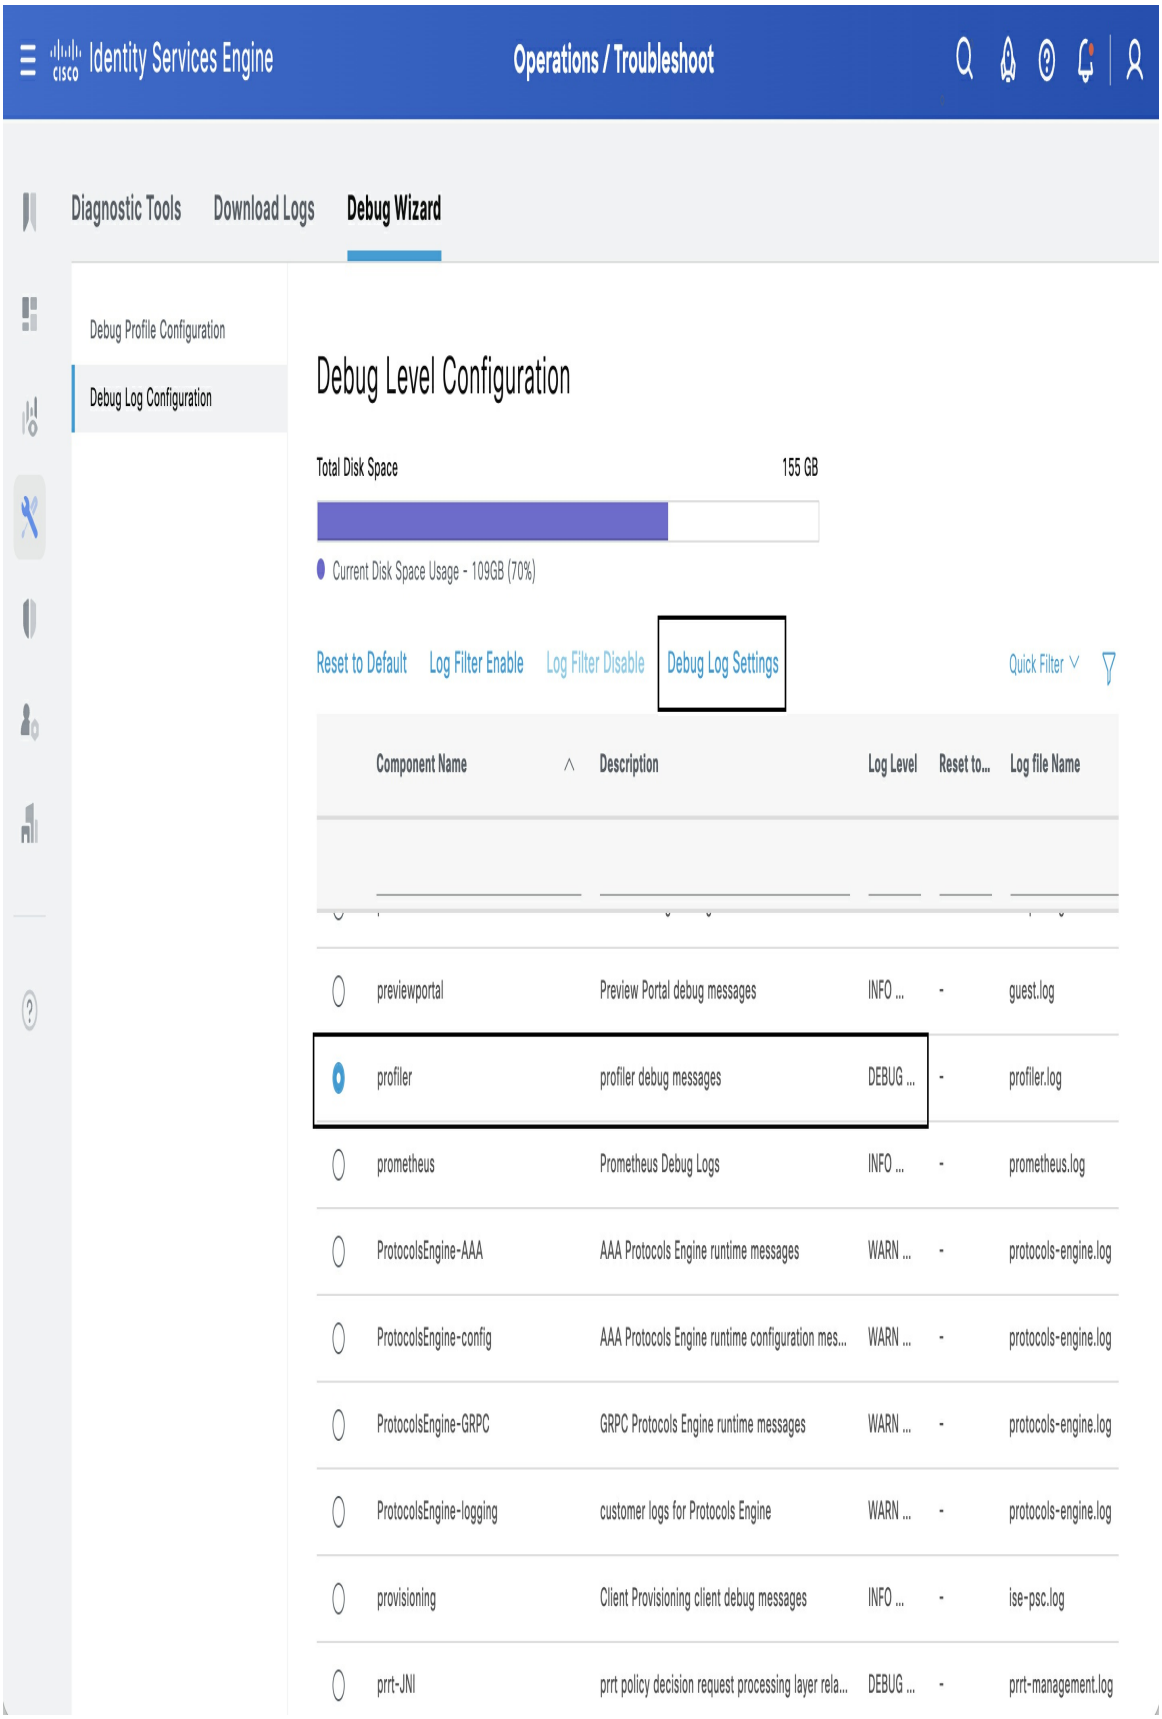

Identity Services Engine
Operations / Troubleshoot
Diagnostic Tools
Download Logs
Debug Wizard
Debug Profile Configuration
Debug Log Configuration
Debug Level Configuration
Total Disk Space
155 GB
Current Disk Space Usage - 109GB (70%)
Reset to Default
Log Filter Enable
Log Filter Disable
Debug Log Settings
Quick Filter
Component Name
Description
Log Level
Reset to...
Log file Name
previewportal
Preview Portal debug messages
INFO ...
guest.log
profiler
profiler debug messages
DEBUG ...
profiler.log
prometheus
Prometheus Debug Logs
INFO ...
prometheus.log
ProtocolsEngine-AAA
AAA Protocols Engine runtime messages
WARN ...
protocols-engine.log
ProtocolsEngine-config
AAA Protocols Engine runtime configuration mes...
WARN ...
protocols-engine.log
ProtocolsEngine-GRPC
GRPC Protocols Engine runtime messages
WARN ...
protocols-engine.log
ProtocolsEngine-logging
customer logs for Protocols Engine
WARN ...
protocols-engine.log
provisioning
Client Provisioning client debug messages
INFO ...
ise-psc.log
prrt-JNI
prrt policy decision request processing layer rela...
DEBUG ...
prrt-management.log

Figure 20-30 *Debug Log Settings for Individual Components*

Debugs are useful when you troubleshoot complex issues involving API calls, inter-component communication, internal state machine changes, or when additional processing details are required.

Here are a few examples where debugging may be necessary:

- You can run the profiler component in DEBUG mode and check profiler.log to determine why a certain profile was selected and what total certainty factor was chosen for the endpoint. The following is an example of an endpoint profiled as Cisco-IP-Phone-8865:

```
2026-07-09 20:22:54,843 DEBUG  [pool-327-thread-18][[]]
cisco.profiler.infrastructure.profiling.ProfilerManager -
:4F943E0A0000222B488C9A1F::f7a57e10-7bd3-11f1-a721-7e95583c11a5
analyzing policy hierarchy: Endpoint: AC:7E:8A:B6:AB:32 Endpoint
Phone-8865 for:115 ExceptionRuleMatched:false
```

- You can run the pxGrid component in TRACE mode and check pxgrid-server.log to see which attributes were shared with pxGrid clients. The following is an example of pxGrid data sent to the Firewall Management Center (FMC) from ISE:

```
2026-07-09 20:13:19,374 TRACE  [sub-sender-10][[]]
cpm.pxgridwebapp.ws.pubsub.SubscriptionSender -
::::::e5cfc26acab543409a7998b6aa088514:- Complete stompframe pul
{"sessions":[{"timestamp":"2026-07-
09T20:13:18.951Z","state":"STARTED","userName":"alice@cxsecurit
5040-730d-4ab1-98ef-
d36eee53ab4c","callingStationId":"00:50:56:A8:96:0E","calledSta
E:85","auditSessionId":"000000000000000C03E75C63","ipAddresses"
fe80::4928:6554:f7f1:8f45"],"macAddress":"00:50:56:A8:96:0E","na
8.79","nasPortId":"GigabitEthernet1/0/5","nasPortType":"Etherne
Microsoft-
Workstation","adNormalizedUser":"alice","adHostResolvedIdentiti
win11","providers":["None"],"endpointCheckResult":"none","ident
```

Technet24

```
"identitySourcePortFirst":0,"serviceType":"Framed","network
DeviceProfileName":"Cisco","radiusFlowType":"Wired802_1x","ssid
85","mdmMacAddress":"00:50:56:A8:96:0E","mdmOsVersion":"","mdmR
mpliant":true,"mdmDiskEncrypted":false,"mdmJailBroken":false,"m
mModel":"","mdmManufacturer":"","mdmImei":"","mdmSerialNumber":
"Intune","selectedAuthzProfiles":["PermitAccess"],"identityProv
D":"af42bac0-52aa-429f-92db-925819967d17","authMethod":"dot1x",
(EAP-TLS)","useCase":"Eap
Chaining","ctsSecondarySecurityGroups":[],"sxpDomains":[],"inbo
"workloadClassification":[],"sessionSource":0,"vrf":"","destina
17.88"}],"sequence":130}
```


- You can run the mdm component in TRACE mode and check ise-psc.log to view the exact API calls made to the MDM server. The following is an example of the query and response:


```
2026-07-09 20:05:17,023 INFO  [MdmEventHandler-64-thread-2][[]]
cisco.cpm.mdm.util.MdmRESTClient -::::::- GET: MDM Server URL:
https://fef.msua08.manage.microsoft.com/TrafficGateway/TrafficR
Access/ComplianceRetrievalService/cisco/devices/?paging=0&query
1405e4-6924-42d7-99b9-9780d26fa7b7&filter=all
2026-07-09 20:05:17,583 TRACE [MdmEventHandler-64-thread-2][[]]
cisco.cpm.mdm.util.MdmRESTClient -::::::-
Response data received from the MDM server : attributes3631405e
9780d26fa7b73C52A1B54C2A631405e4-
6924-42d7-99b9-
9780d26fa7b73C52A1B54C2Atruetrue
```


You can retrieve debug logs in several different ways:

- Downloading individual log files from a specific node under **Operations > Troubleshoot > Download Logs > Debug Logs**, as shown in Figure 20-31.

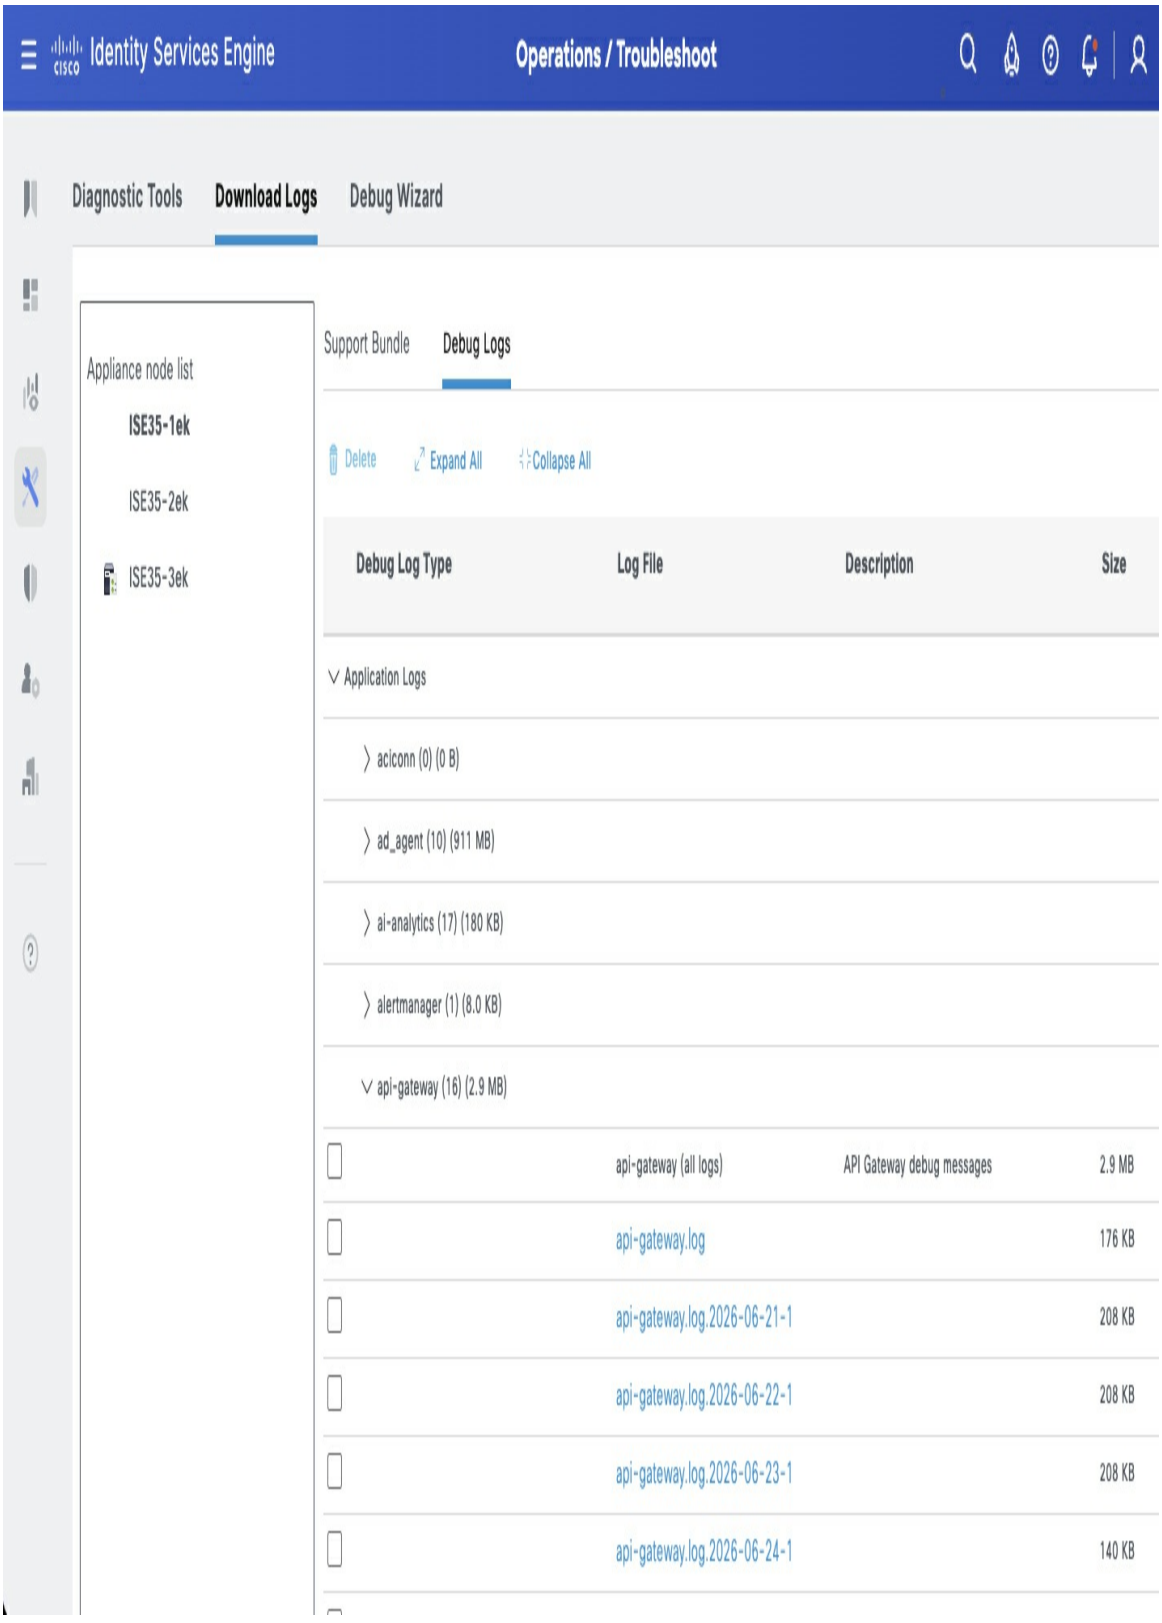


Figure 20-31 *Debug Log Settings for Individual Components*

Technet24

- Looking at the logs from the CLI. You can either display the entire log file with

```
ISE35-1ek/admin#show logging application profiler.log
```

Or you can tail the logs with

```
ISE35-1ek/admin#show logging application profiler.log tail
```

If you are unsure about the log filename, you can run the following command to display all available log files:

```
ISE35-1ek/admin#show logging application
```

and

```
ISE35-1ek/admin#show logging system
```

- You can download the Support Bundle from ISE and include debug logs.

Support Bundle

Support Bundle is a collection of logs from ISE that helps with troubleshooting. Typically, TAC engineers request Support Bundles, but you can also review them independently. You can collect Support Bundles from **Operations > Troubleshoot > Download Logs > Support Bundle**, as shown in Figure 20-32.

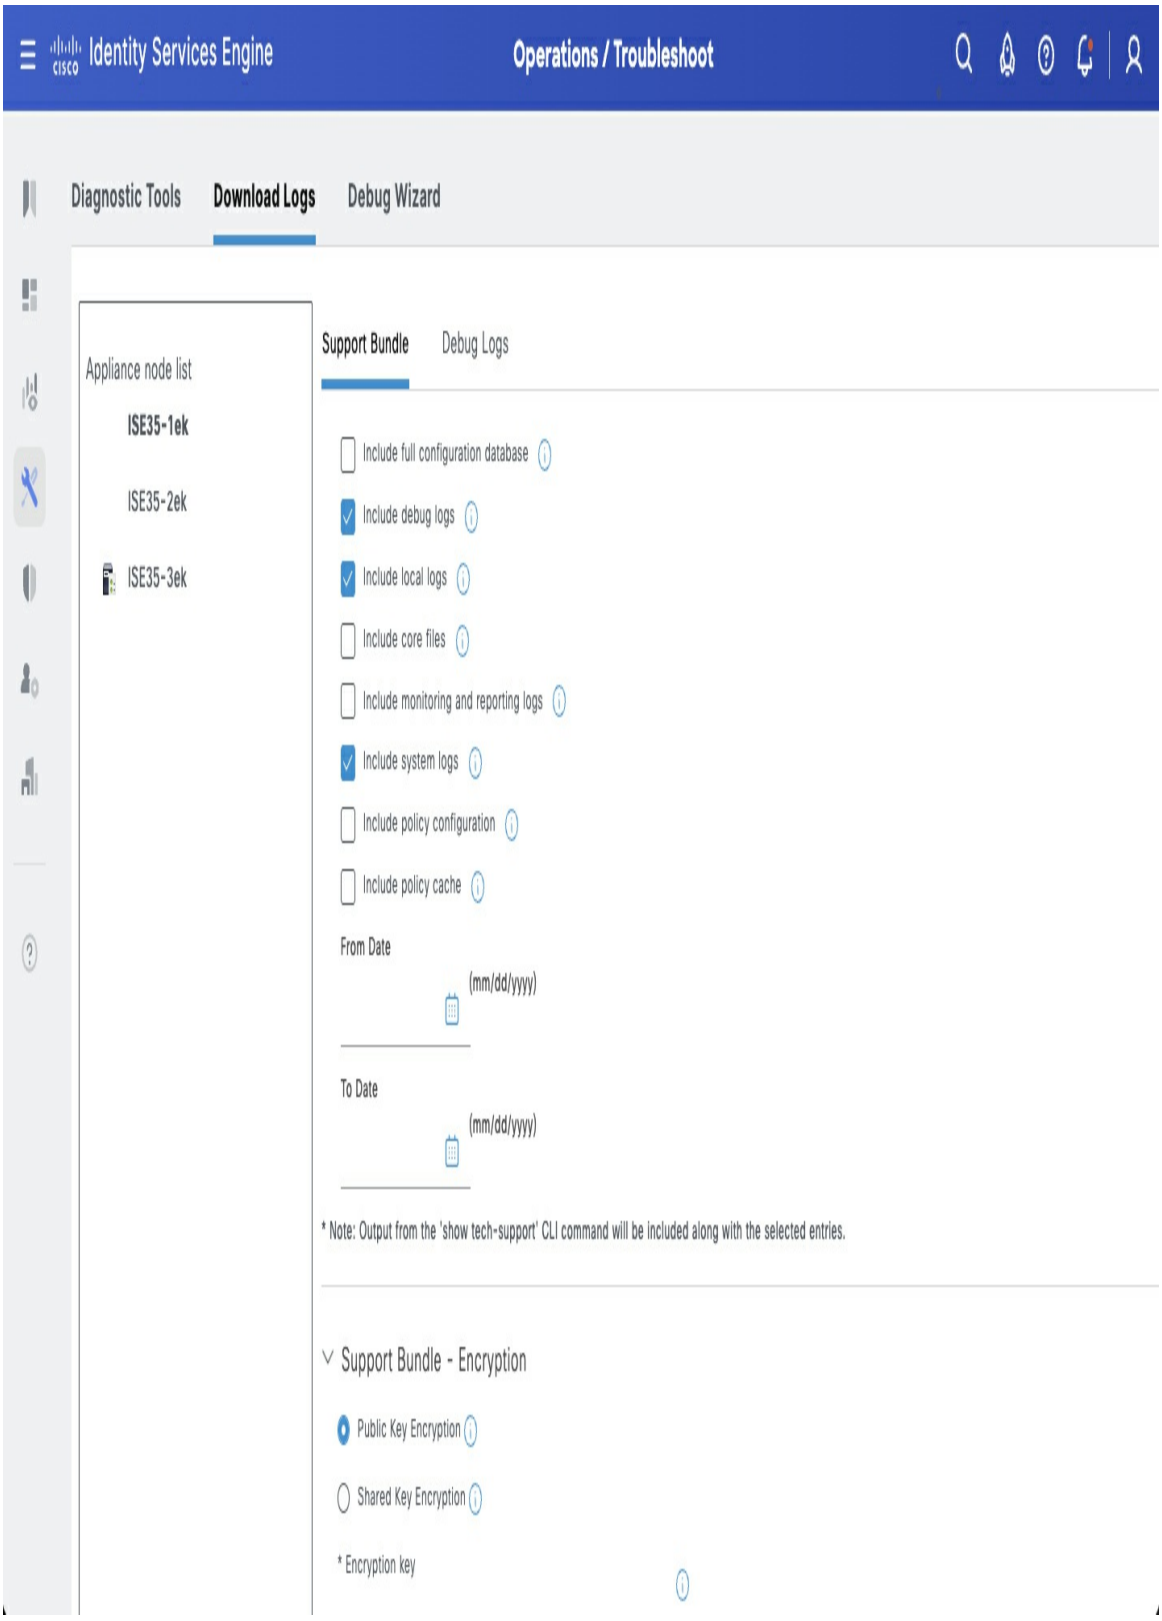


Figure 20-32 *Support Bundle Collection Settings*

Technet24

Best practice is to always include

- Debug logs
- Local logs
- System logs

To limit the size of the bundle, you can specify particular dates. If you leave the dates empty, all logs will be included.

There are two choices for encryption:

- **Public Key Encryption:** You encrypt the bundle with the public key. Only Cisco can decrypt it because Cisco holds the private key.
- **Shared Key Encryption:** You encrypt the bundle with a password. Anyone who has the password can decrypt it.

The following bundle name includes the ISE node hostname for which the bundle was generated (ISE35-1ek), the username of the user who generated the bundle (admin), the date when the bundle was generated (07-10-2026), and the time when the bundle was generated (15-10):

```
ise-support-bundle-ISE35-1ek-admin-07-10-2026-15-10.tar.gpg
```

Inside the Support Bundle, you can find multiple folders containing log files. The most critical folders are

- adeos, which contains system logs
- logs, which contains ISE application logs
- showtech, which contains the output of the ISE **show tech-support** command

Reports

Another meaningful way to examine historical data is through reporting. You can access ISE reporting via **Operations > Reports**, as shown in Figure 20-33.

Technet24

ad_state	Today at 11:14 AM	9 KB	Folder
adeos	Today at 12:48 PM	416.7 MB	Folder
apache_conf	Today at 11:11 AM	253 KB	Folder
apache_logs	Today at 11:11 AM	54.2 MB	Folder
async_logs	Today at 11:15 AM	622 KB	Folder
config	Today at 11:11 AM	1.5 MB	Folder
core	Today at 11:10 AM	Zero bytes	Folder
db	Today at 11:11 AM	Zero bytes	Folder
dbexport	Today at 11:10 AM	Zero bytes	Folder
integritycheck.log	Today at 11:15 AM	Zero bytes	Log File
logs	Today at 12:48 PM	2.11 GB	Folder
LX5xKWZ7Mk-support-args	Today at 11:10 AM	52 bytes	Document
mntreport	Today at 11:10 AM	Zero bytes	Folder
policyCache	Today at 11:10 AM	Zero bytes	Folder
policyXml	Today at 11:10 AM	Zero bytes	Folder
prrt_config	Today at 11:10 AM	459 KB	Folder
psc_config	Today at 11:10 AM	1.6 MB	Folder
README.txt	Today at 11:11 AM	23 KB	Plain Text
showtech	Today at 11:11 AM	6 MB	Folder
SSE	Today at 12:48 PM	8 KB	Folder
thinfiles	Today at 11:11 AM	15 KB	Folder

Figure 20-33 *RADIUS Accounting Report*

There are multiple report categories, each with a subset of reports. For quick access, you can add your preferred report to the My Reports section, making it easily accessible. You can export reporting data in several ways:

- Repository (CSV)
- Local (CSV)
- Local (PDF)

Note

Local CSV export will export only 500 records. The recommended option for CSV export is Repository (CSV).

CSV format is more convenient for troubleshooting because you can easily apply filters and perform additional analysis, such as building graphs or pie charts.

Common reports used include audit reports (administrator logins, change configuration audit), device administration (TACACS authentication and accounting), and endpoints and users (RADIUS authentication and accounting reports). The authentication summary report provides a high-level overview of authentications across the deployment, which helps narrow down common problems.

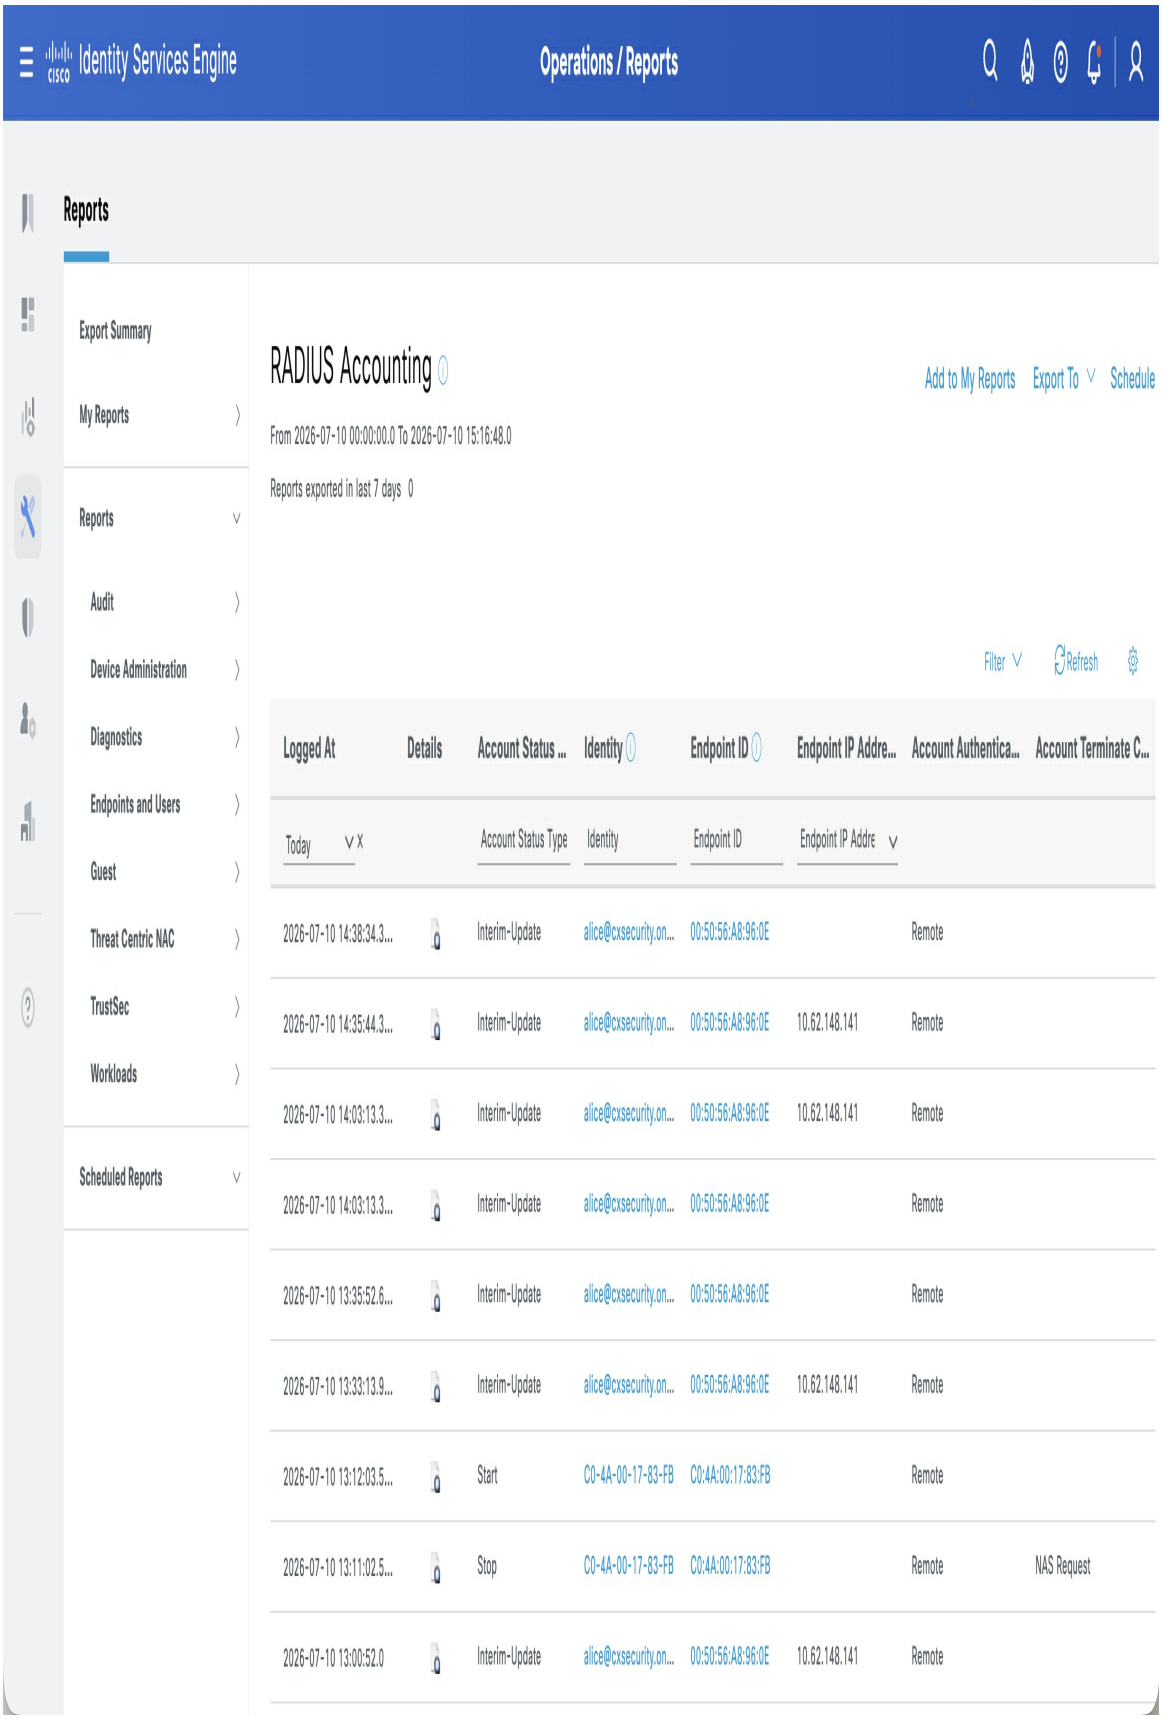
You can also schedule reports by clicking **Schedule** from the Reports window. You then need to select a schedule, choose a repository to store the report, and assign a name to the report.

System 360 Monitoring

System 360 Monitoring provides a comprehensive method for

Technet24

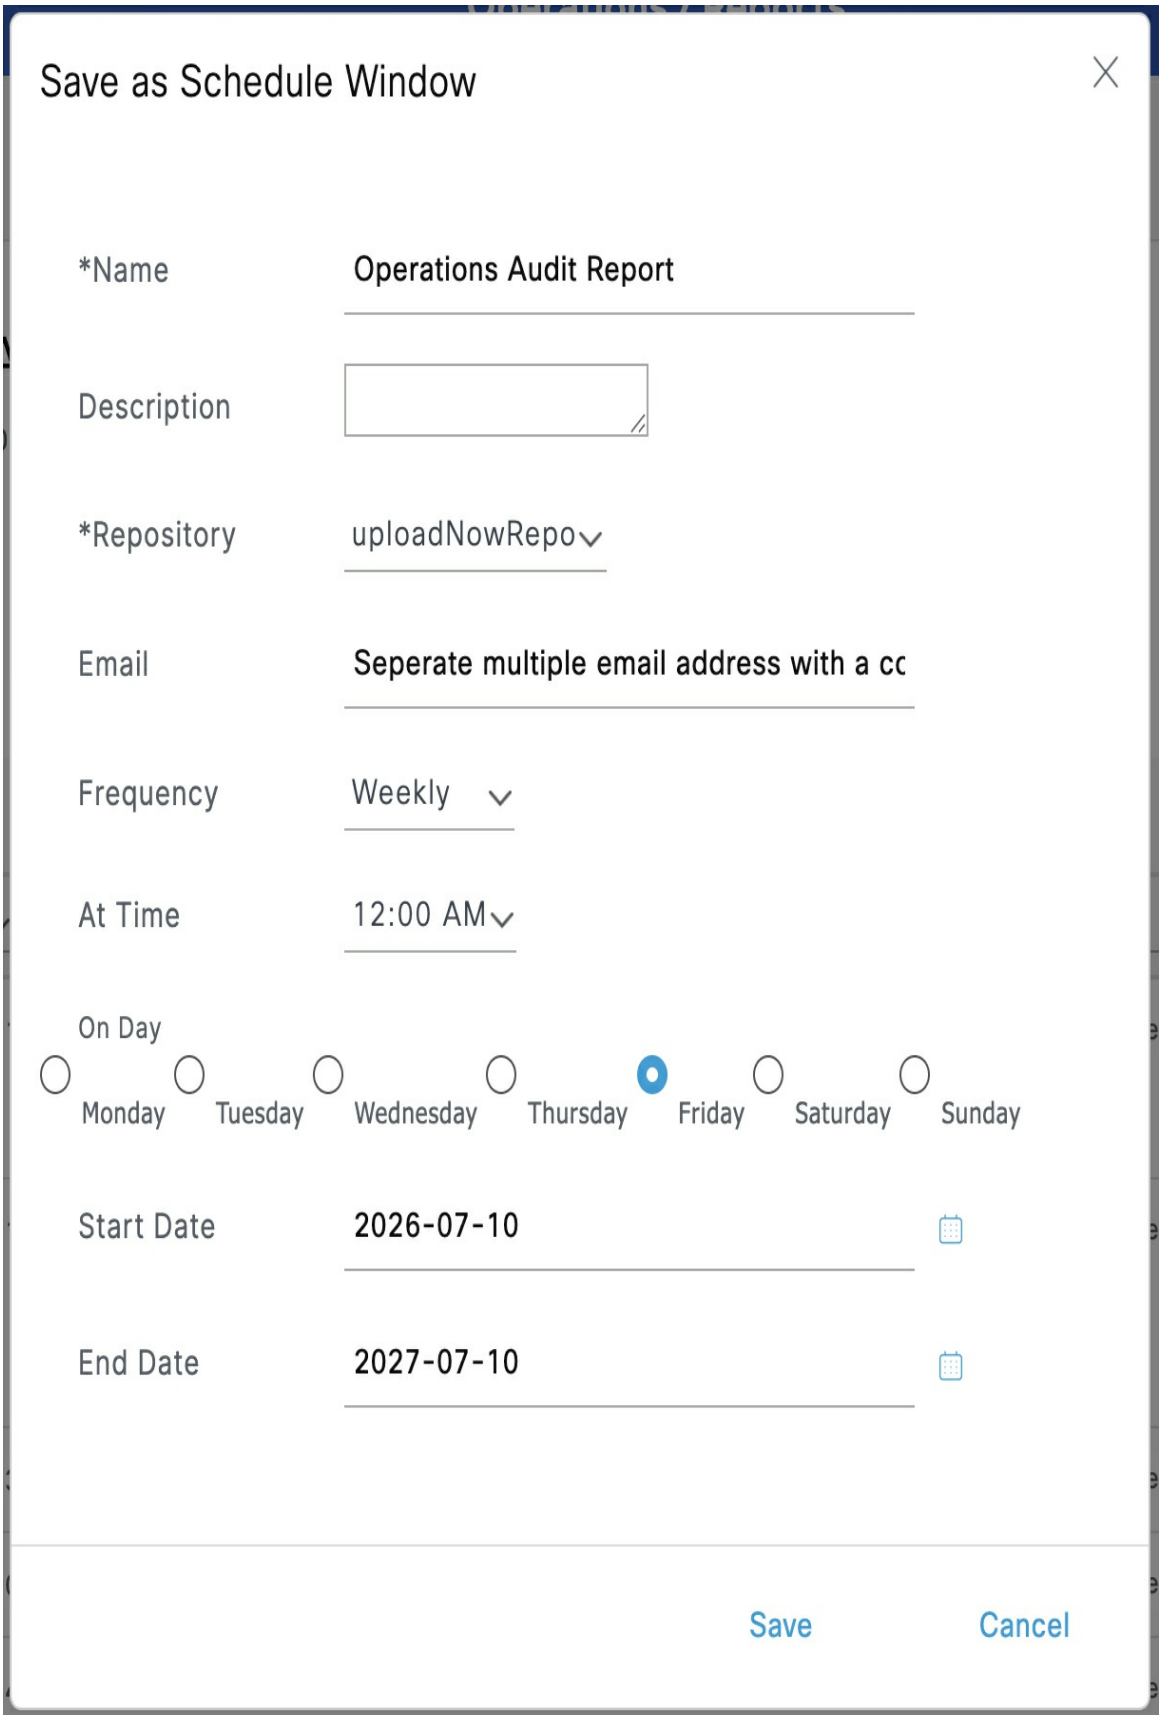
monitoring ISE performance. It allows you, as administrator, to view parameters for individual ISE nodes, such as CPU usage, memory usage, load average, disk usage, and more, as shown in Figure 20-34.


Identity Services Engine
Operations / Reports
Reports
Export Summary
My Reports
Reports
Audit
Device Administration
Diagnostics
Endpoints and Users
Guest
Threat Centric NAC
TrustSec
Workloads
Scheduled Reports
RADIUS Accounting
Add to My Reports
Export To
Schedule
From 2026-07-10 00:00:00.0 To 2026-07-10 15:16:48.0
Reports exported in last 7 days 0
Filter
Refresh
Logged At
Details
Account Status ...
Identity
Endpoint ID
Endpoint IP Addre...
Account Authentica...
Account Terminate C...
Today
Account Status Type
Identity
Endpoint ID
Endpoint IP Addre
2026-07-10 14:38:34.3... Interim-Update alice@cxsecurity.on... 00:50:56:A8:96:0E Remote
2026-07-10 14:35:44.3... Interim-Update alice@cxsecurity.on... 00:50:56:A8:96:0E 10.62.148.141 Remote
2026-07-10 14:03:13.3... Interim-Update alice@cxsecurity.on... 00:50:56:A8:96:0E 10.62.148.141 Remote
2026-07-10 14:03:13.3... Interim-Update alice@cxsecurity.on... 00:50:56:A8:96:0E Remote
2026-07-10 13:35:52.6... Interim-Update alice@cxsecurity.on... 00:50:56:A8:96:0E Remote
2026-07-10 13:33:13.9... Interim-Update alice@cxsecurity.on... 00:50:56:A8:96:0E 10.62.148.141 Remote
2026-07-10 13:12:03.5... Start C0-4A-00-17-83-FB C0:4A:00:17:83:FB Remote
2026-07-10 13:11:02.5... Stop C0-4A-00-17-83-FB C0:4A:00:17:83:FB Remote NAS Request
2026-07-10 13:00:52.0 Interim-Update alice@cxsecurity.on... 00:50:56:A8:96:0E 10.62.148.141 Remote

Technet24

Figure 20-34 *System 360 Monitoring Main Dashboard*

This feature also stores historical information and builds graphs based on this data, as shown in Figure 20-35. It allows you to correlate any issue observed on ISE with the performance of the node or the network.


Save as Schedule Window
*Name
Operations Audit Report
Description
*Repository
uploadNowRepo
Email
Seperate multiple email address with a cc
Frequency
Weekly
At Time
12:00 AM
On Day
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday
Sunday
Start Date
2026-07-10
End Date
2027-07-10
Save
Cancel

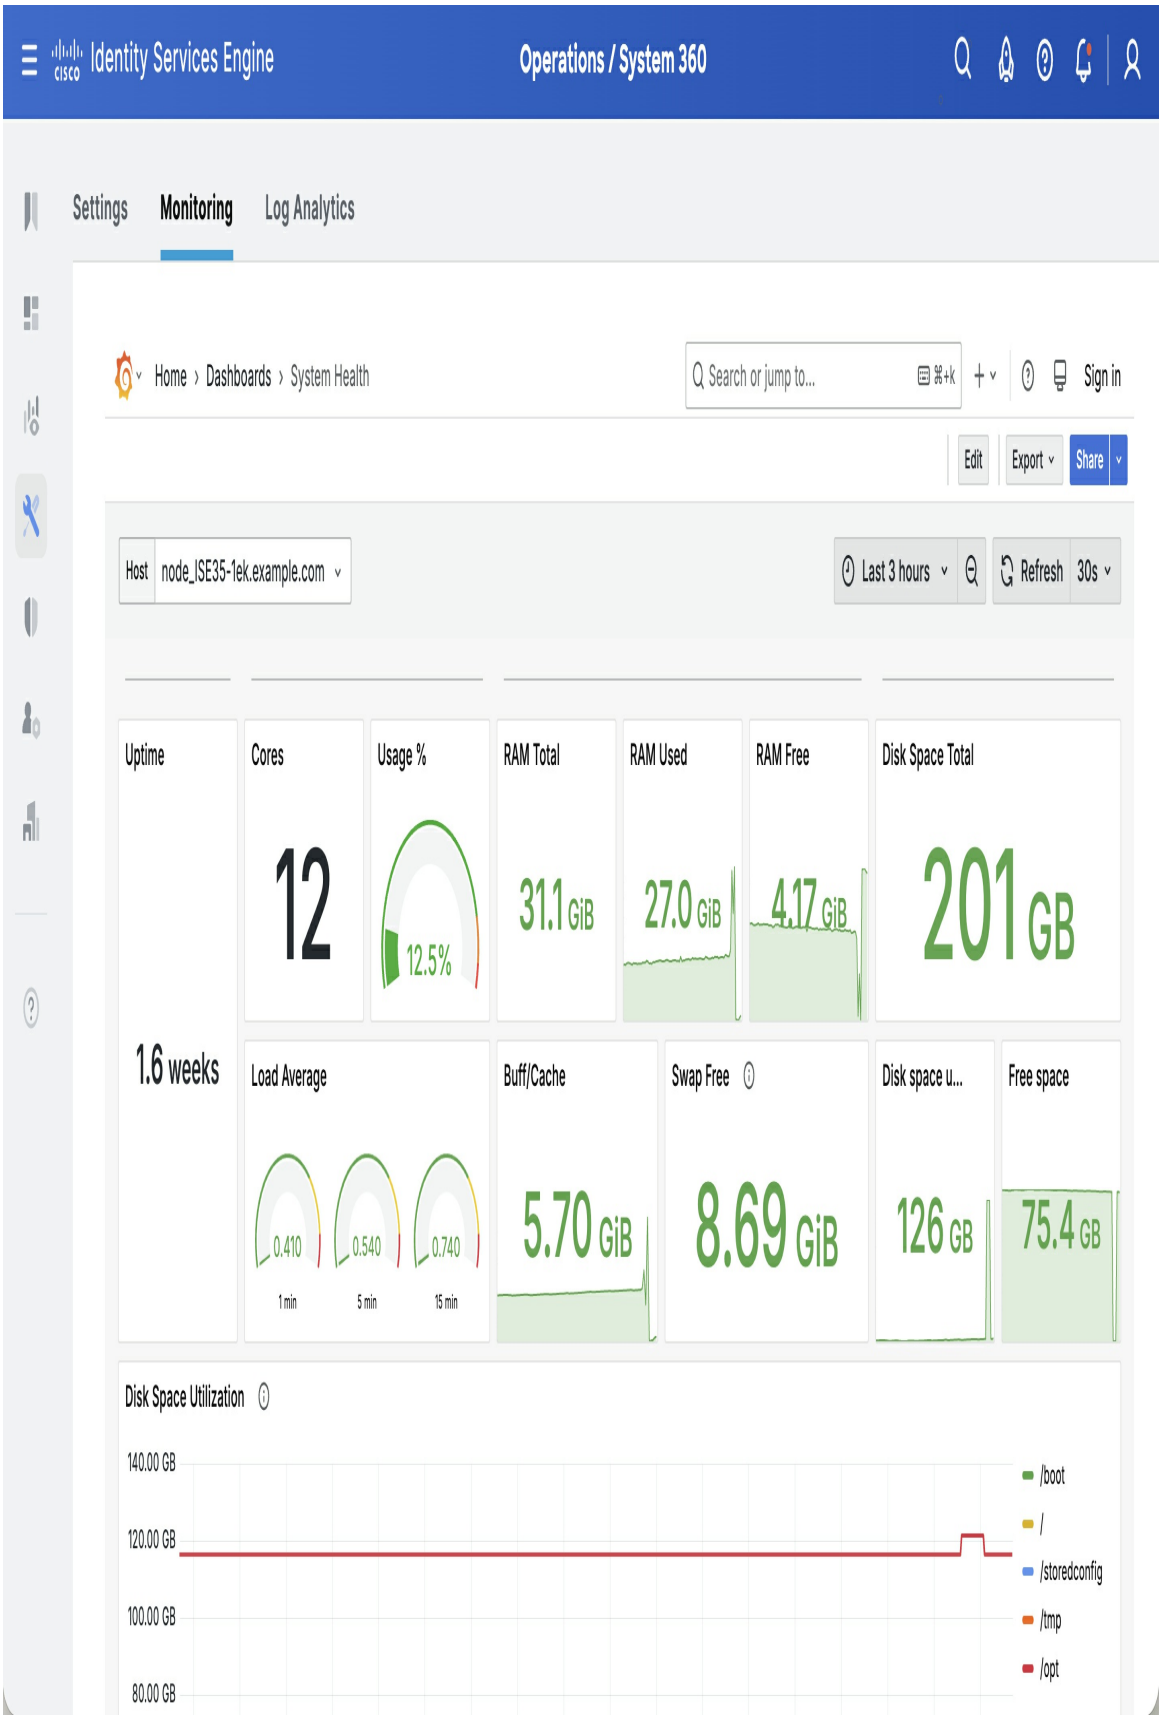
Technet24

Figure 20-35 *System 360 Monitoring CPU and Load Average Graphs*

System 360 Monitoring is accessible via **Operations > System 360 > Monitoring**. Monitoring data from all the nodes in deployment is stored on the Primary Policy Administrative Node (PPAN), and data is purged when either 5 GB of data or seven days' worth of data is reached, whichever occurs first.

System 360 Log Analytics

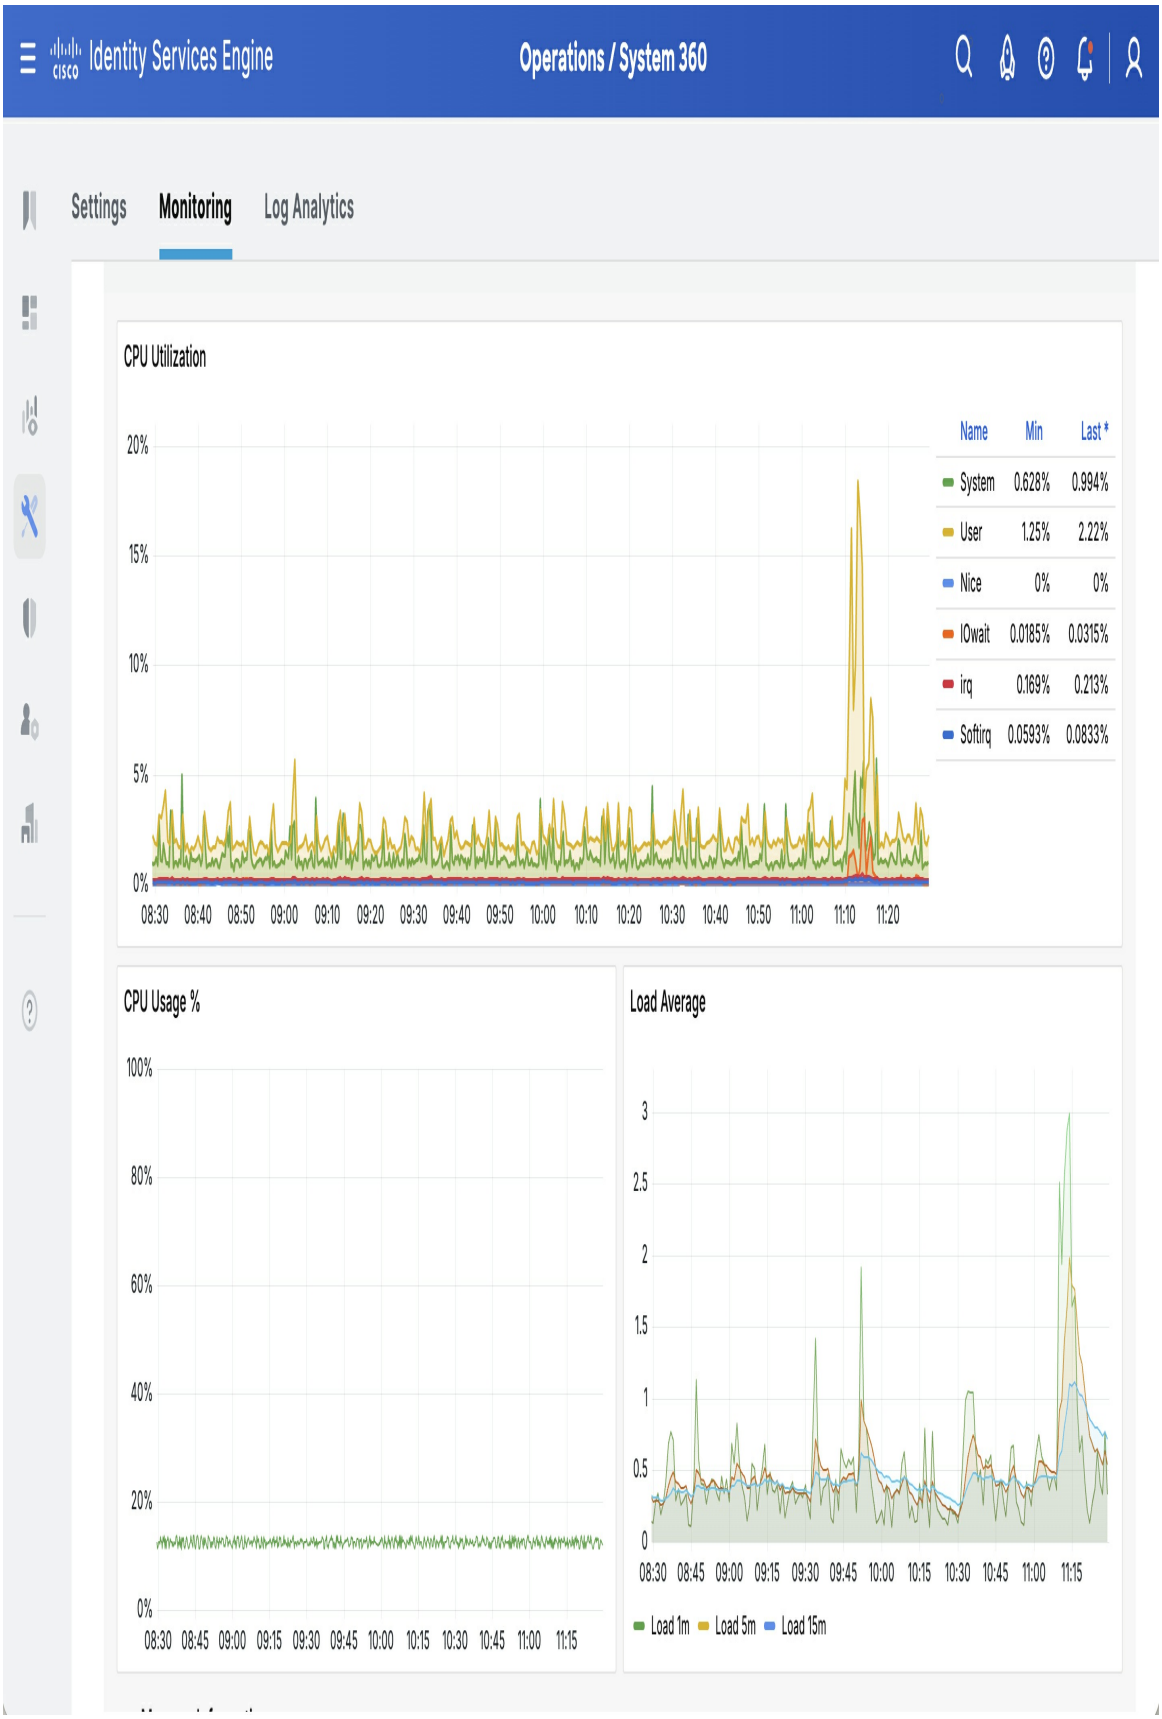
System 360 Log Analytics is accessible via **Operations > System 360 > Log Analytics**. This powerful tool is used for analyzing RADIUS, TACACS, and profiler performance. It includes a built-in set of dashboards that you can use immediately, as shown in Figure 20-36.


Identity Services Engine
Operations / System 360
Settings
Monitoring
Log Analytics
Home › Dashboards › System Health
Search or jump to...
Sign in
Edit
Export
Share
Host
node_ISE35-1ek.example.com
Last 3 hours
Refresh
30s
Uptime
1.6 weeks
Cores
12
Usage %
12.5%
RAM Total
31.1 GiB
RAM Used
27.0 GiB
RAM Free
4.17 GiB
Disk Space Total
201 GB
Load Average
0.410
0.540
0.740
1 min
5 min
15 min
Buff/Cache
5.70 GiB
Swap Free
8.69 GiB
Disk space u...
126 GB
Free space
75.4 GB
Disk Space Utilization
140.00 GB
120.00 GB
100.00 GB
80.00 GB
/boot
/
/storedconfig
/tmp
/opt

Technet24

Figure 20-36 *System 360 Log Analytics Dashboards*

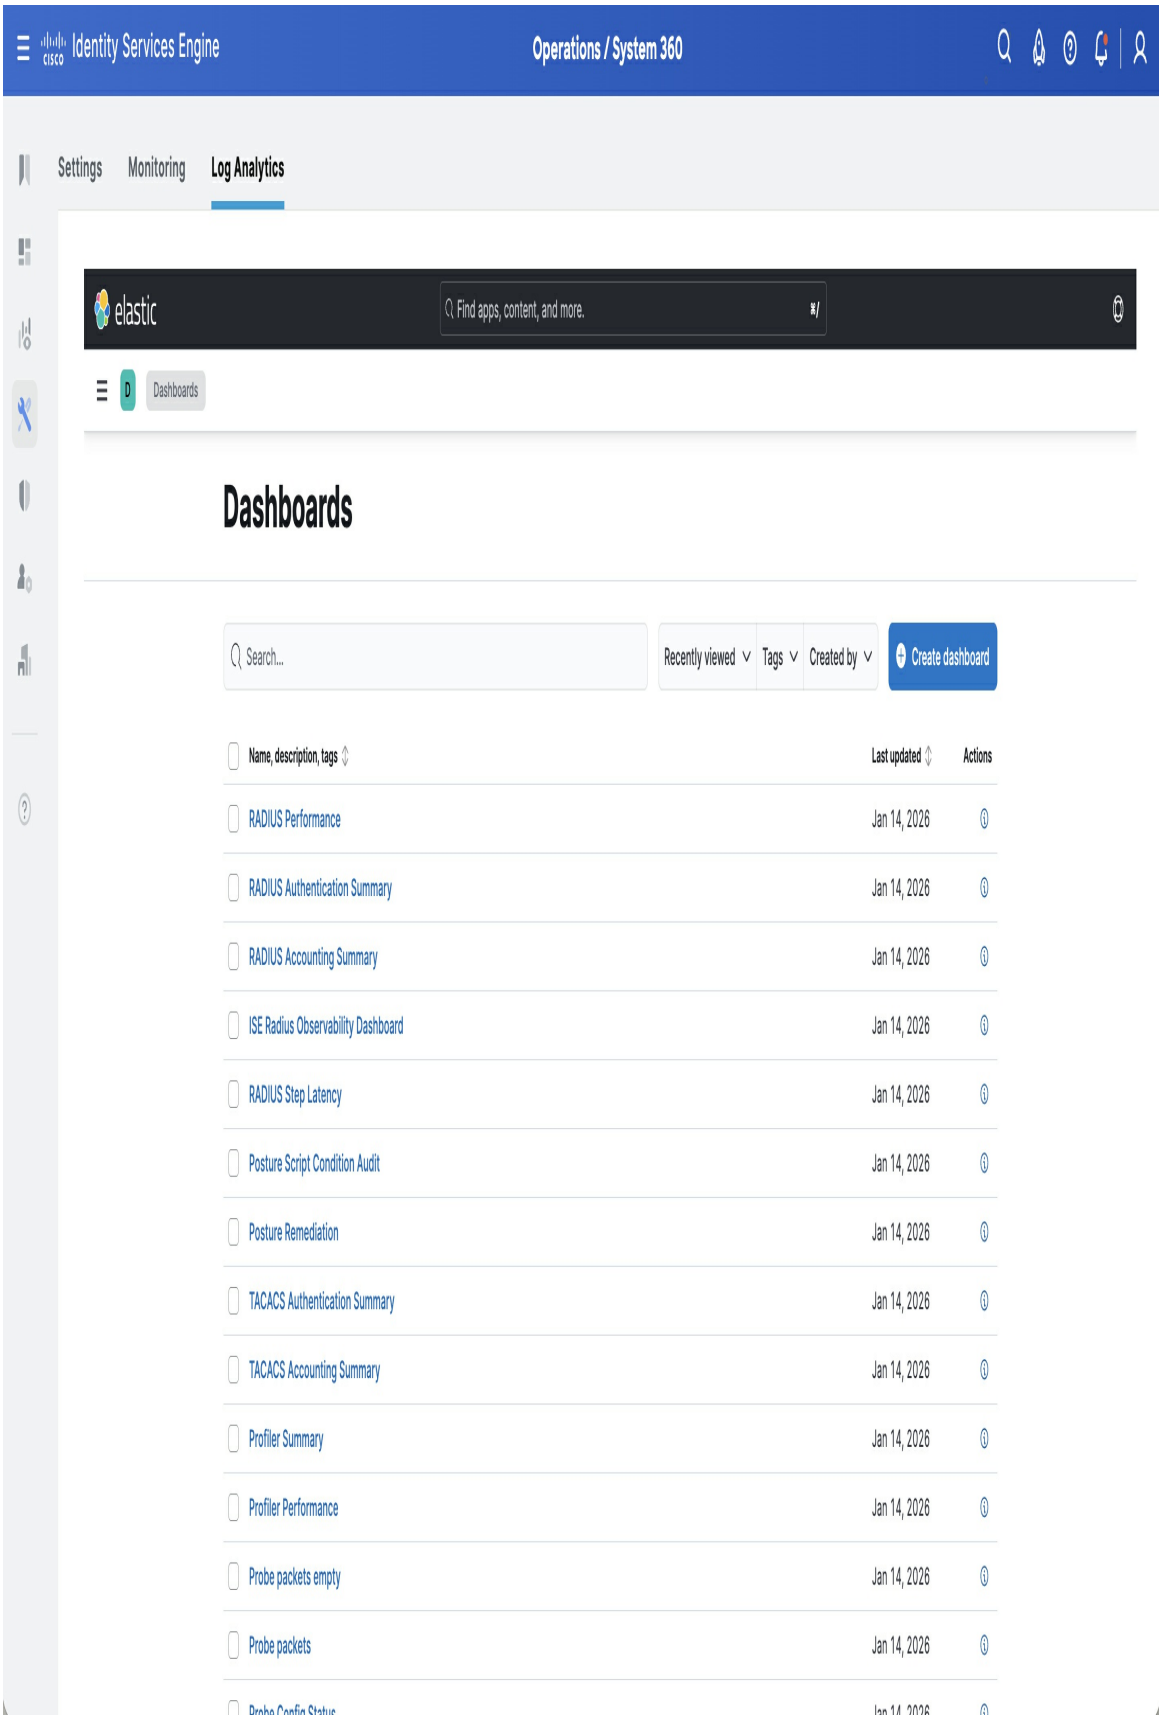
An example of the RADIUS Performance Dashboard is shown in Figure 20-37. This dashboard provides information such as the number of unique endpoints, NADs, total RADIUS requests, endpoint profile breakdown, and top talkers. Identifying top talkers, endpoints which generate most of the RADIUS traffic, is especially important for detecting misbehaving endpoints.


Identity Services Engine
Operations / System 360
Settings
Monitoring
Log Analytics
CPU Utilization
20%
15%
10%
5%
0%
08:30 08:40 08:50 09:00 09:10 09:20 09:30 09:40 09:50 10:00 10:10 10:20 10:30 10:40 10:50 11:00 11:10 11:20
Name Min Last *
System 0.628% 0.994%
User 1.25% 2.22%
Nice 0% 0%
IOwait 0.0185% 0.0315%
irq 0.169% 0.213%
Softirq 0.0593% 0.0833%
CPU Usage %
100%
80%
60%
40%
20%
0%
08:30 08:45 09:00 09:15 09:30 09:45 10:00 10:15 10:30 10:45 11:00 11:15
Load Average
3
2.5
2
1.5
1
0.5
0
08:30 08:45 09:00 09:15 09:30 09:45 10:00 10:15 10:30 10:45 11:00 11:15
Load 1m
Load 5m
Load 15m

Technet24

Figure 20-37 *System 360 Log Analytics RADIUS Performance Dashboard*

Multiple graphs are available on the same dashboard, as shown in Figure 20-38. These graphs allow you to view RADIUS traffic over time, per server, and through various other metrics.


Identity Services Engine
Operations / System 360
Settings
Monitoring
Log Analytics
elastic
Find apps, content, and more.
Dashboards
Dashboards
Search...
Recently viewed
Tags
Created by
Create dashboard
Name, description, tags
Last updated
Actions
RADIUS Performance
Jan 14, 2026
RADIUS Authentication Summary
Jan 14, 2026
RADIUS Accounting Summary
Jan 14, 2026
ISE Radius Observability Dashboard
Jan 14, 2026
RADIUS Step Latency
Jan 14, 2026
Posture Script Condition Audit
Jan 14, 2026
Posture Remediation
Jan 14, 2026
TACACS Authentication Summary
Jan 14, 2026
TACACS Accounting Summary
Jan 14, 2026
Profiler Summary
Jan 14, 2026
Profiler Performance
Jan 14, 2026
Probe packets empty
Jan 14, 2026
Probe packets
Jan 14, 2026

Technet24

Figure 20-38 *System 360 Log Analytics RADUS Performance Graphs*

Log Analytics is powered by Elastic Search and Kibana, which enables you to easily drill into specific timeframes. The dashboard automatically adjusts to display only relevant data.

Log Analytics operates on the primary or secondary MnT node, depending on the configuration under **Operations > System 360 > Settings**.

Summary

Troubleshooting, monitoring, and alerting are the most crucial operations for everyday use of ISE. You need to be proactively informed if something goes wrong and have a variety of tools available to help you troubleshoot issues on ISE and within the network. With tools such as detailed authentication reports, TCP Dump, and Log Analytics, you can identify problems on network devices, endpoints, firewalls, and more. By following a structured troubleshooting methodology, you can save significant time for the business and effectively resolve most issues.

Chapter 21. ISE Maintenance

Introduction

A general understanding of software maintenance requires you to keep software reliable, secure, and up to date after deployment. This step is crucial because it prevents Cisco Identity Services Engine from being exposed to software defects, compatibility issues, or vulnerabilities. This chapter guides you through repository management, patch installation, and major version upgrades.

Repositories

Repositories play an essential role in various maintenance tasks on ISE. You can use repositories for the following purposes:

- Configuration backups and restores
- Operational backups and restores
- Upgrades
- Support Bundle
- Reports
- Patches

ISE currently supports the following types of repositories:

- DISK

Technet24

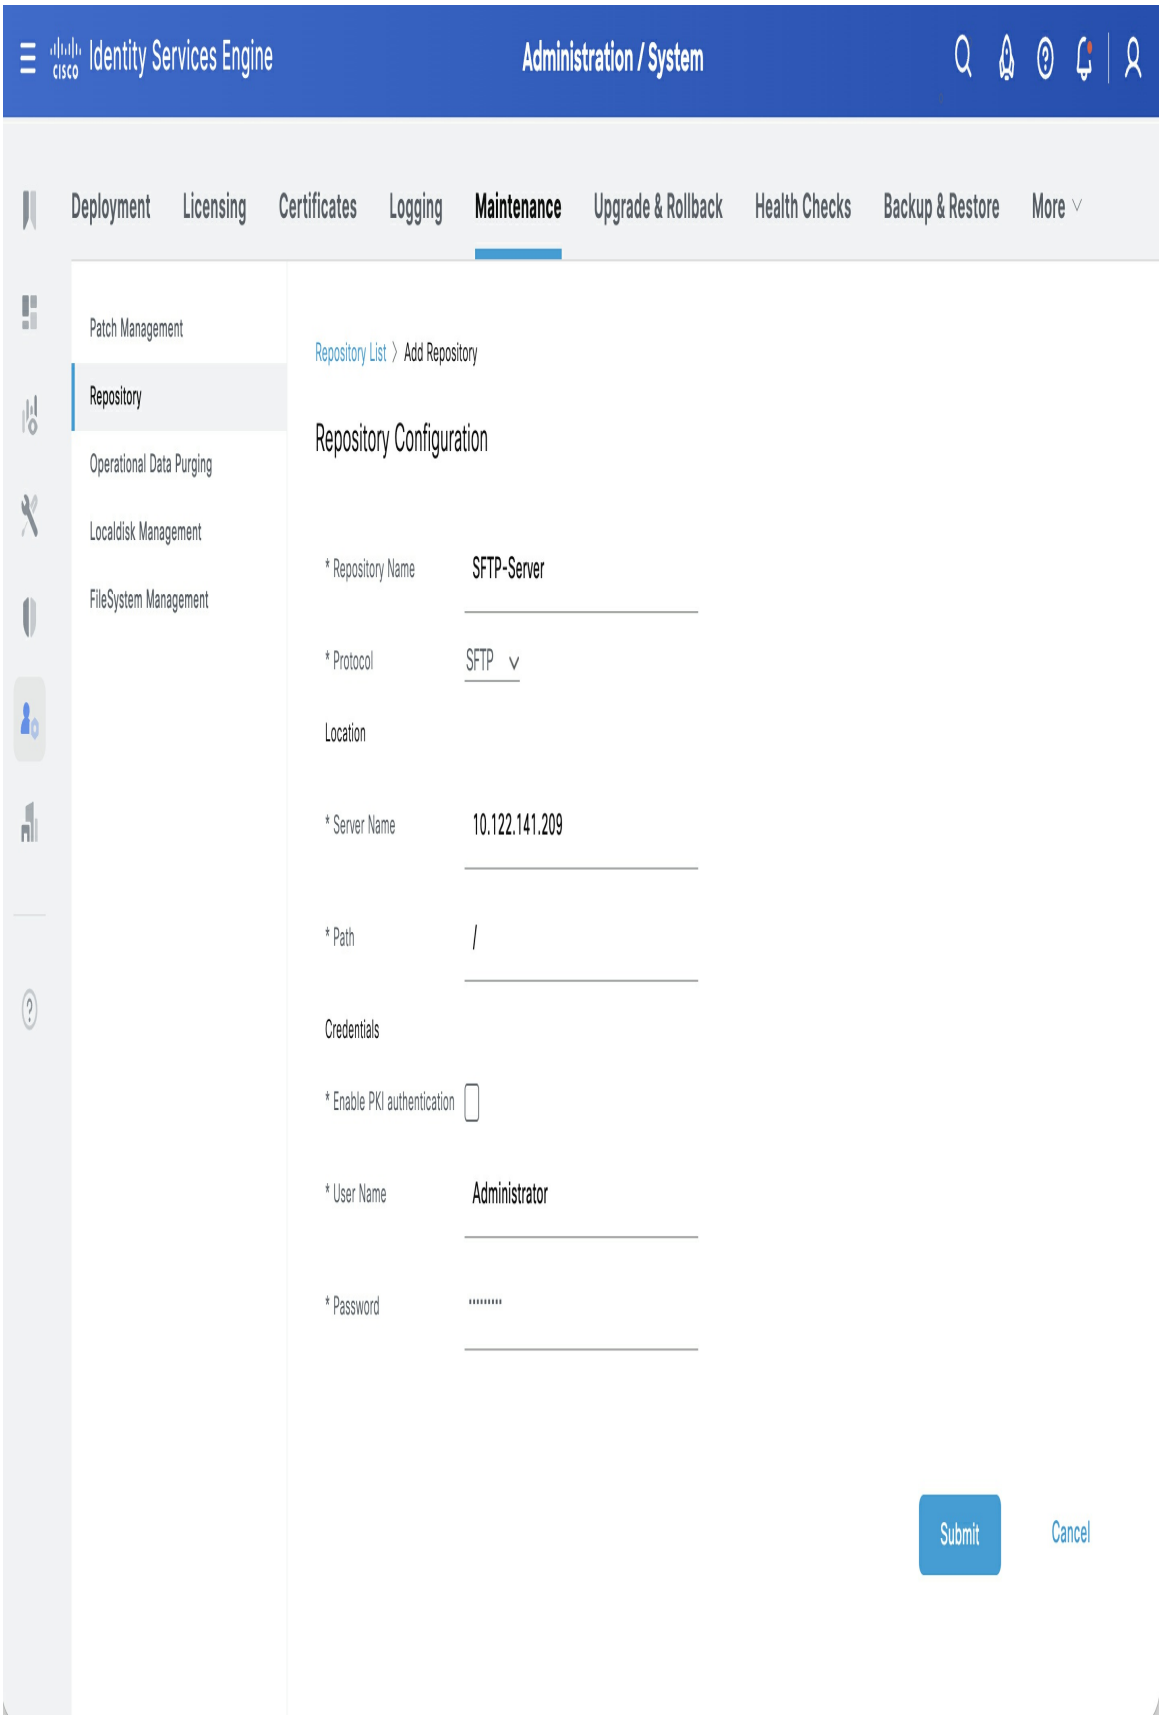
- FTP
- SFTP
- TFTP
- NFS
- CDROM
- HTTP
- HTTPS

Best Practice

Although multiple repository options are available on ISE, SFTP provides the most secure and scalable solution. Therefore, you should use SFTP for the repository, whenever possible.

To configure an SFTP repository, navigate to **Administration > System > Maintenance** and click **+Add**. Then configure the following settings, as shown in Figure 21-1:

- Repository Name
- Protocol
- Server Name
- Path
- Username
- Password


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Patch Management
Repository
Operational Data Purging
Localdisk Management
FileSystem Management
Repository List > Add Repository
Repository Configuration
* Repository Name
SFTP-Server
* Protocol
SFTP
Location
* Server Name
10.122.141.209
* Path
/
Credentials
* Enable PKI authentication
* User Name
Administrator
* Password
Submit
Cancel

Technet24

Figure 21-1 *SFTP Repository Configuration*

After you configure the SFTP server, the system presents the following warning:

```
Host key of sftp server must be added through CLI using 'crypto host
command before this repository can be used. Also ensure that the hos
the host name used in the URL of the repository configuration.
```

This warning indicates that SFTP becomes operational only after you manually add the host key of the SFTP server to each ISE node in the deployment.

To add the host key, log in to the command-line interface of each ISE node and add the host using the following command:

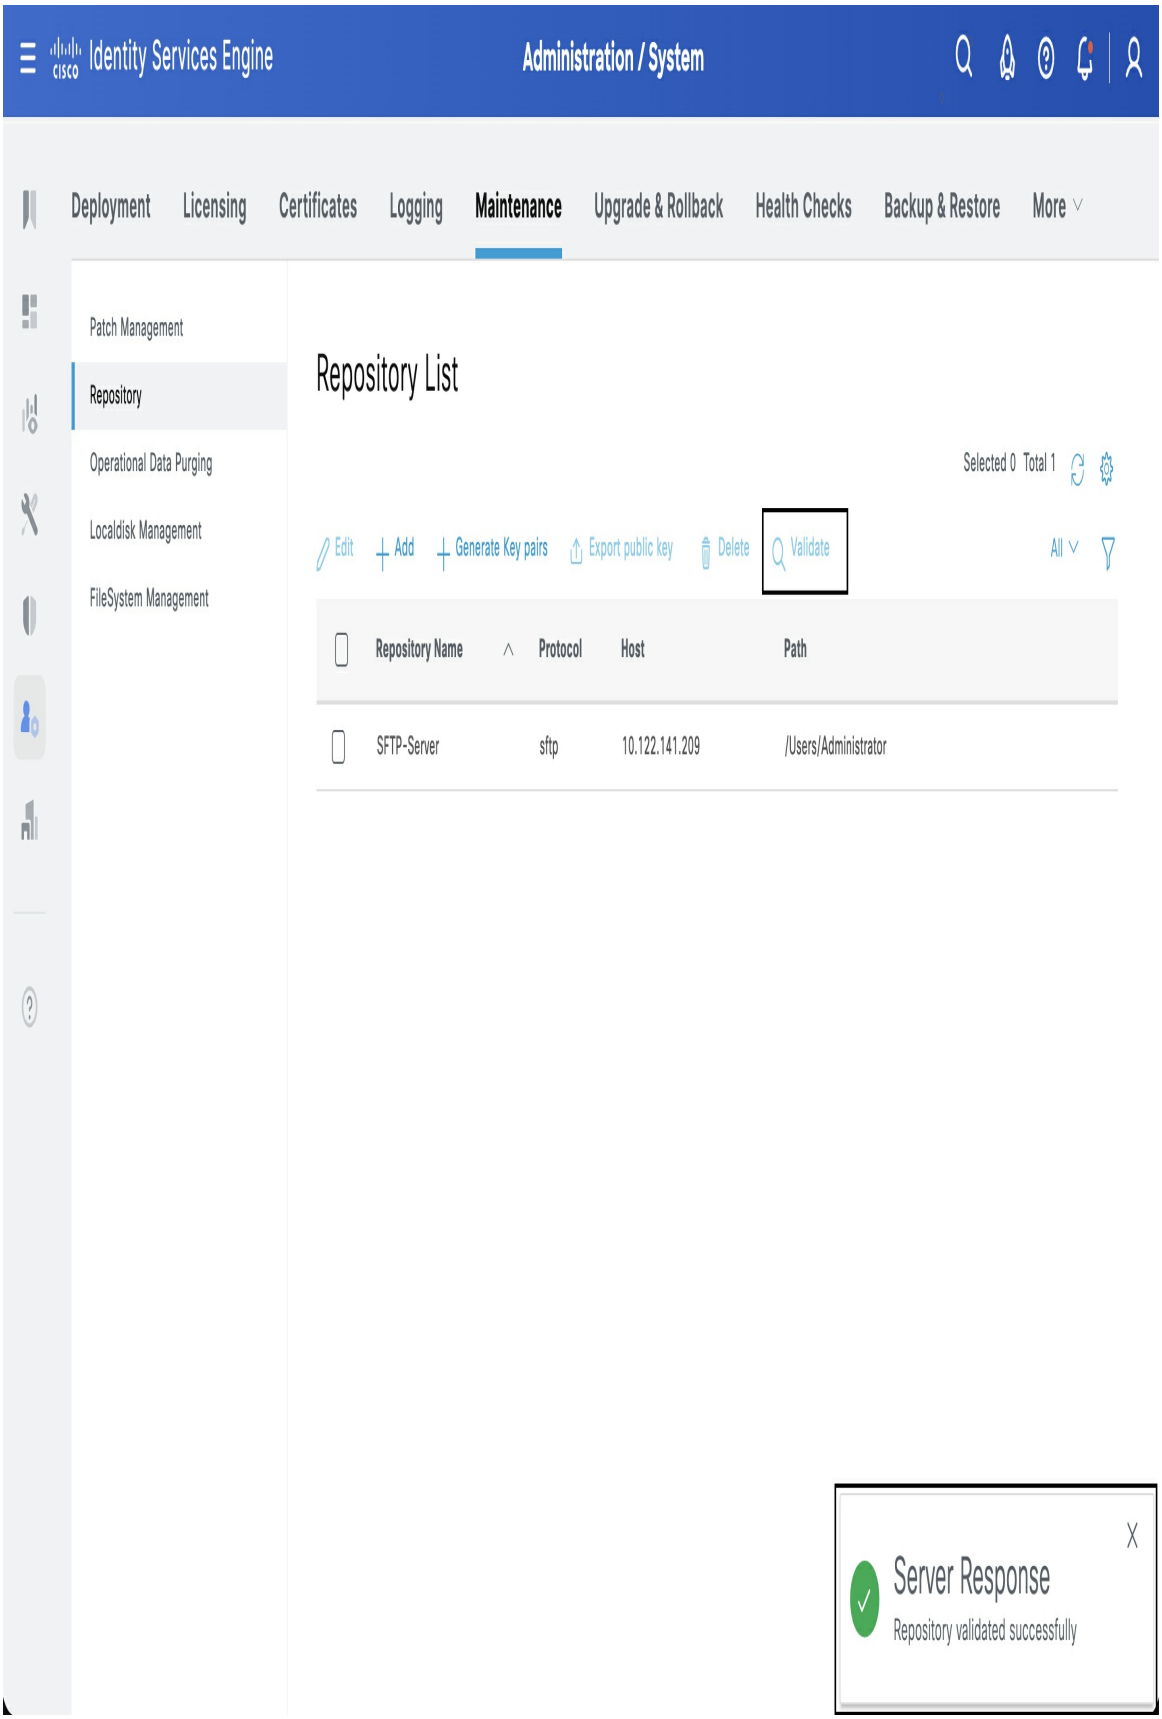
```
ISE35-1ek/admin#crypto host_key add host ?
Possible completions:
    Specify IPv4/IPv6 address or hostname of host server
```

> **Note**
>
> The host configuration must match the Server Name settings in the graphical user interface. If you configured FQDN in the GUI, use FQDN in the ISE CLI. If you configured an IP address in the GUI, use the IP address in the CLI.

To verify that the SFTP repository functions correctly, run the **show repository** command:

```
ISE35-1ek/admin#show repository SFTP-Server
Weekly_Backup-CFG10-250627-1930.tar.gpg
ise-patchbundle-3.4.0.608-Patch4-25081421.SPA.x86_64.tar.gz
```

Alternatively, you can validate the repository from the GUI, as shown in Figure 21-2.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Patch Management
Repository
Operational Data Purging
Localdisk Management
FileSystem Management
Repository List
Selected 0 Total 1
Edit
Add
Generate Key pairs
Export public key
Delete
Validate
All
Repository Name
Protocol
Host
Path
SFTP-Server
sftp
10.122.141.209
/Users/Administrator
Server Response
Repository validated successfully

Technet24

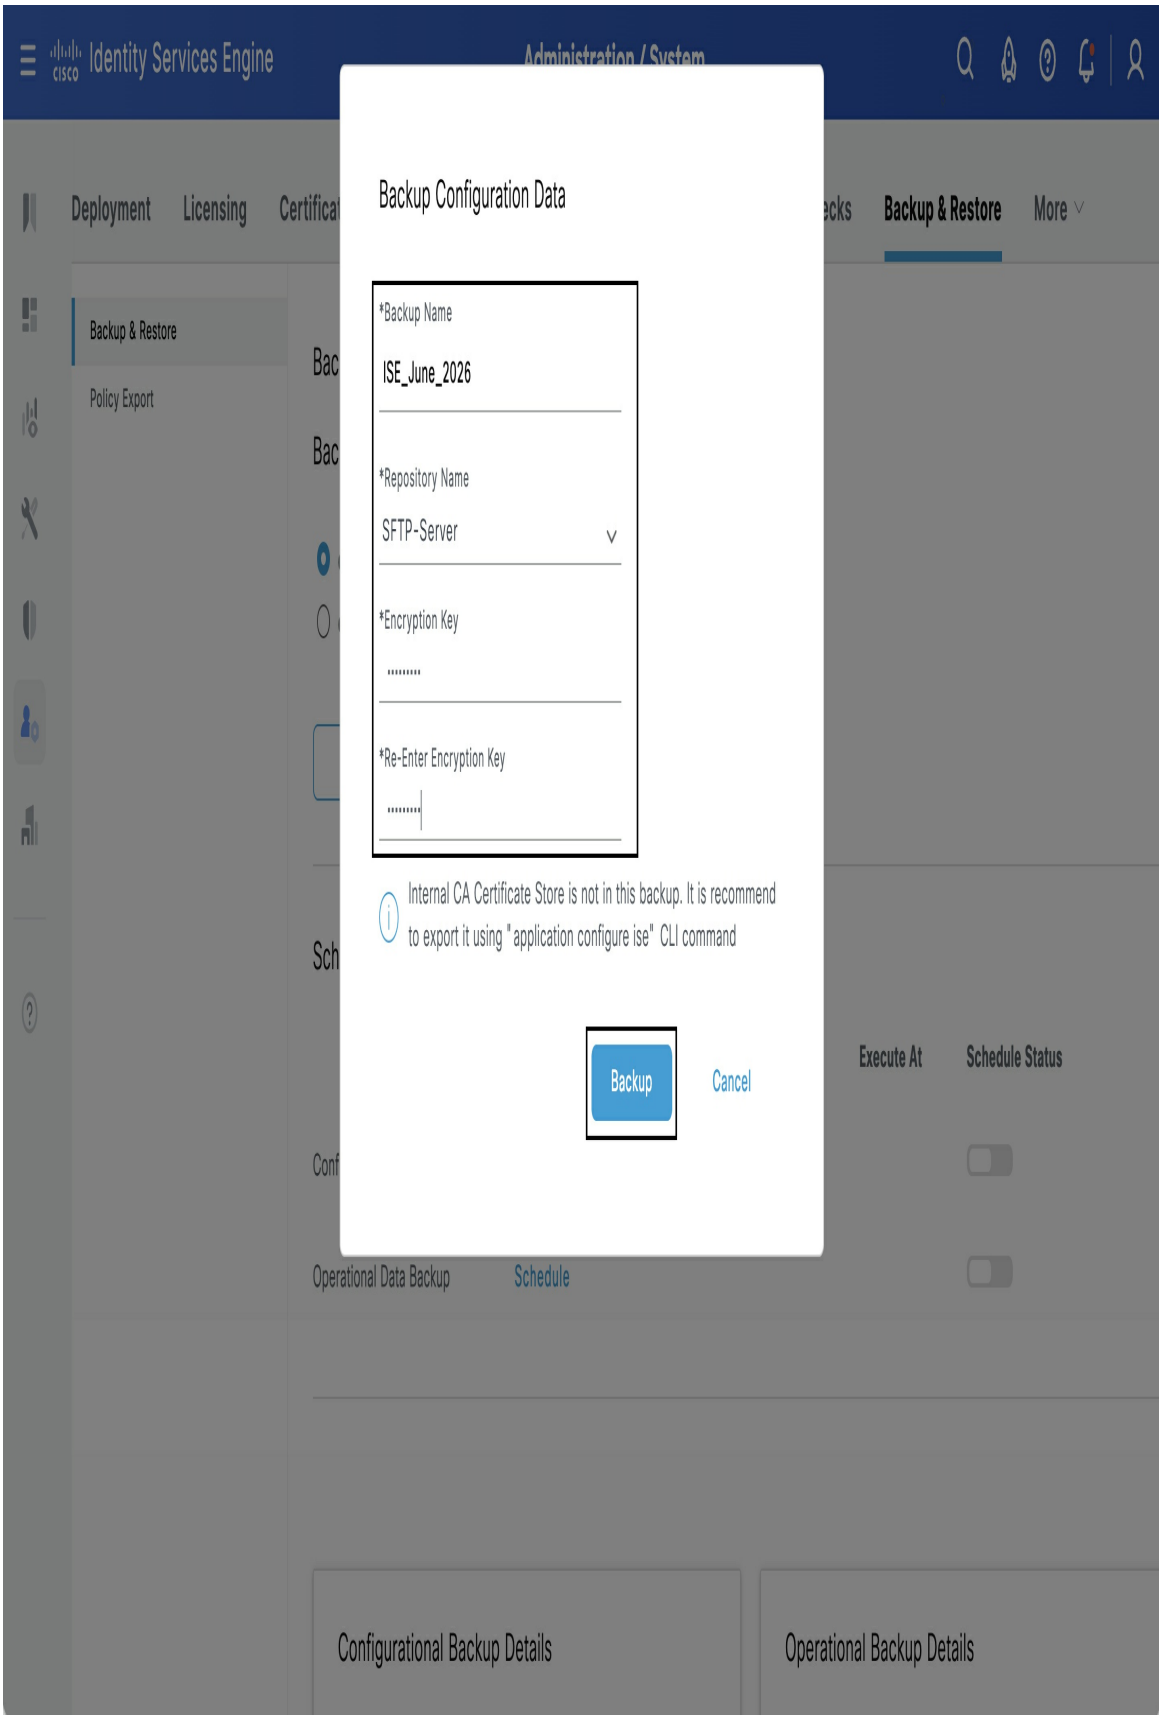
Figure 21-2 *SFTP Server GUI Validation*

Repositories configured through the ISE CLI are not replicated across the deployment and do not appear in the ISE GUI. Repositories configured through the ISE GUI are replicated to the ISE CLI and across the deployment.

Backup and Restore

ISE supports two types of backups: configuration data backup and operational data backup. A configuration backup includes all ISE application settings, such as policy settings, endpoint database, user settings, and system configuration, including ADE-OS settings like hostname, domain name, and IP address. An operational backup contains all ISE operational data, such as live logs and historical reports.

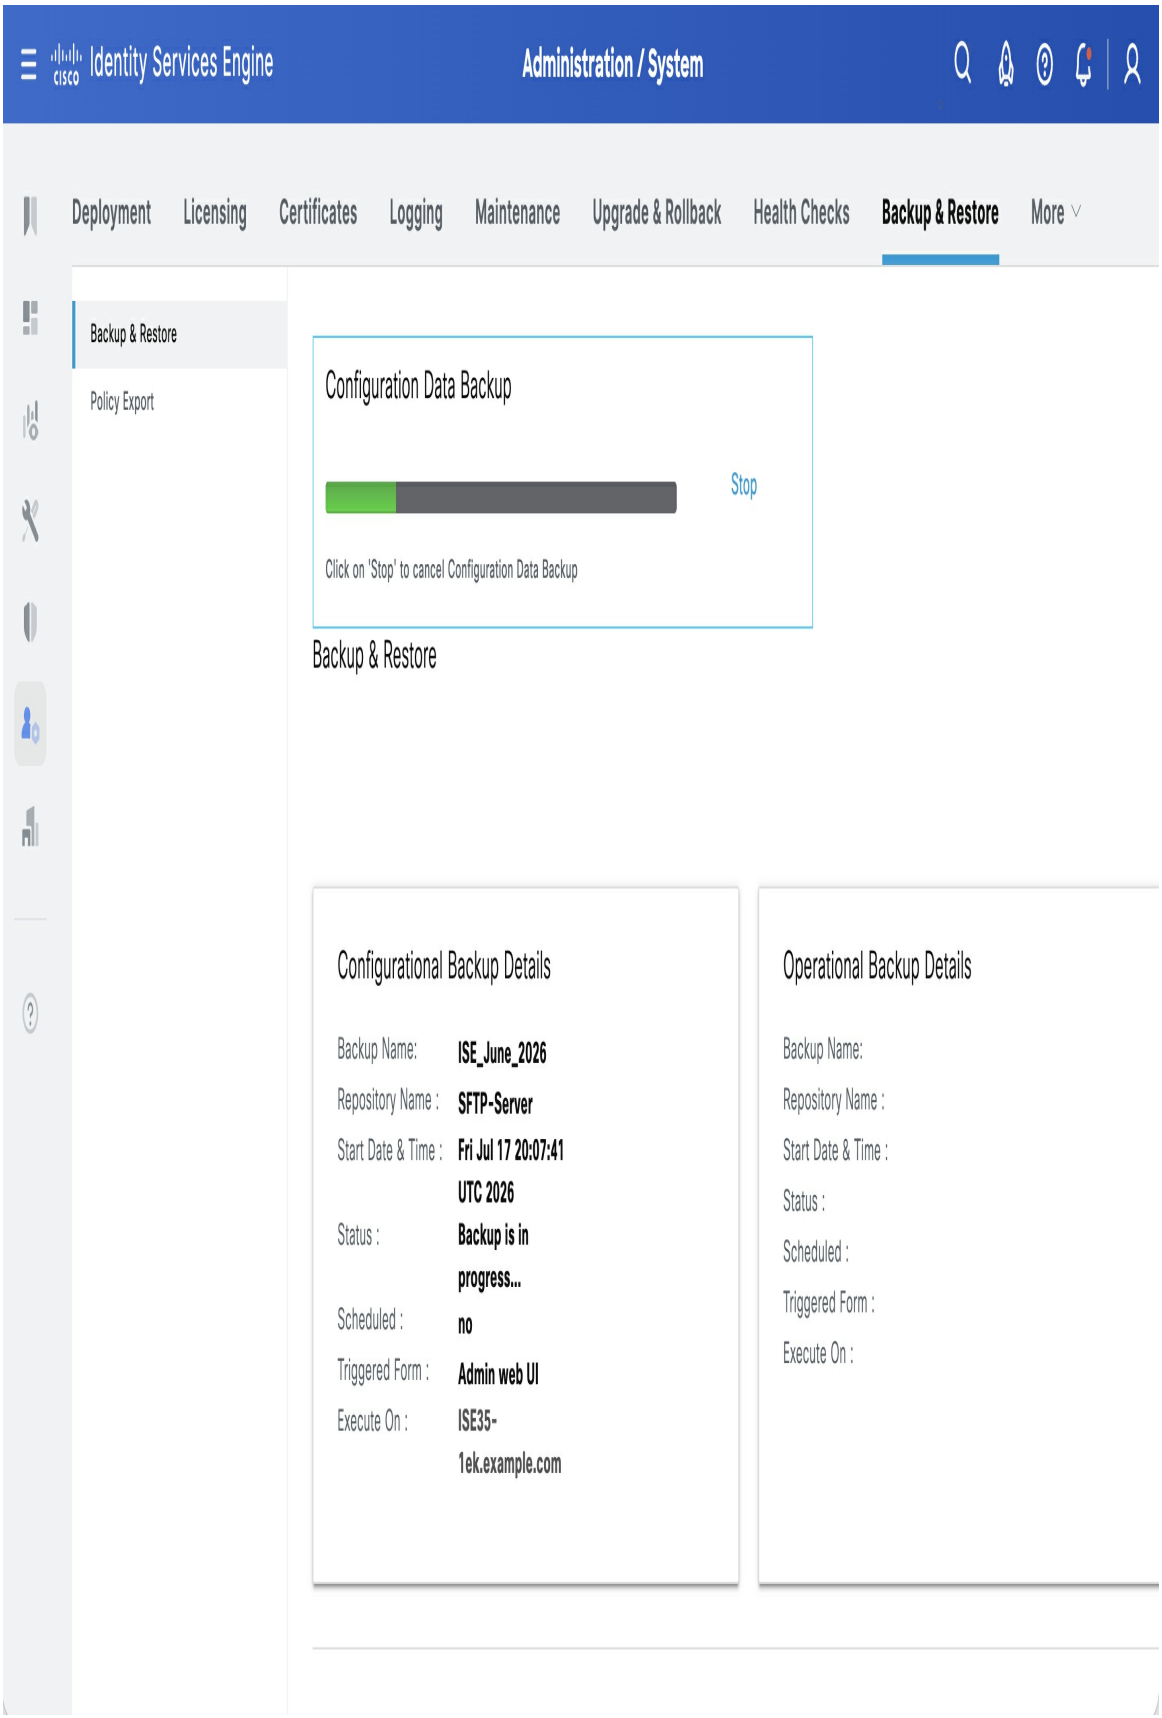
You can create the backup from the ISE GUI or from the CLI. To create a backup from the GUI, navigate to **Administration > System > Backup & Restore** and click **Backup Now** to start an on-demand backup. As shown in Figure 21-3, you must select a repository, assign a name to the backup, and provide the encryption key.


Identity Services Engine
Administration / System
Deployment
Licensing
Backup & Restore
More
Backup & Restore
Policy Export
Backup Configuration Data
*Backup Name
ISE_June_2026
*Repository Name
SFTP-Server
*Encryption Key
*Re-Enter Encryption Key
Internal CA Certificate Store is not in this backup. It is recommend to export it using "application configure ise" CLI command
Backup
Cancel
Execute At
Schedule Status
Operational Data Backup
Schedule
Configurational Backup Details
Operational Backup Details

Technet24

Figure 21-3 *On-Demand GUI Configuration Backup*

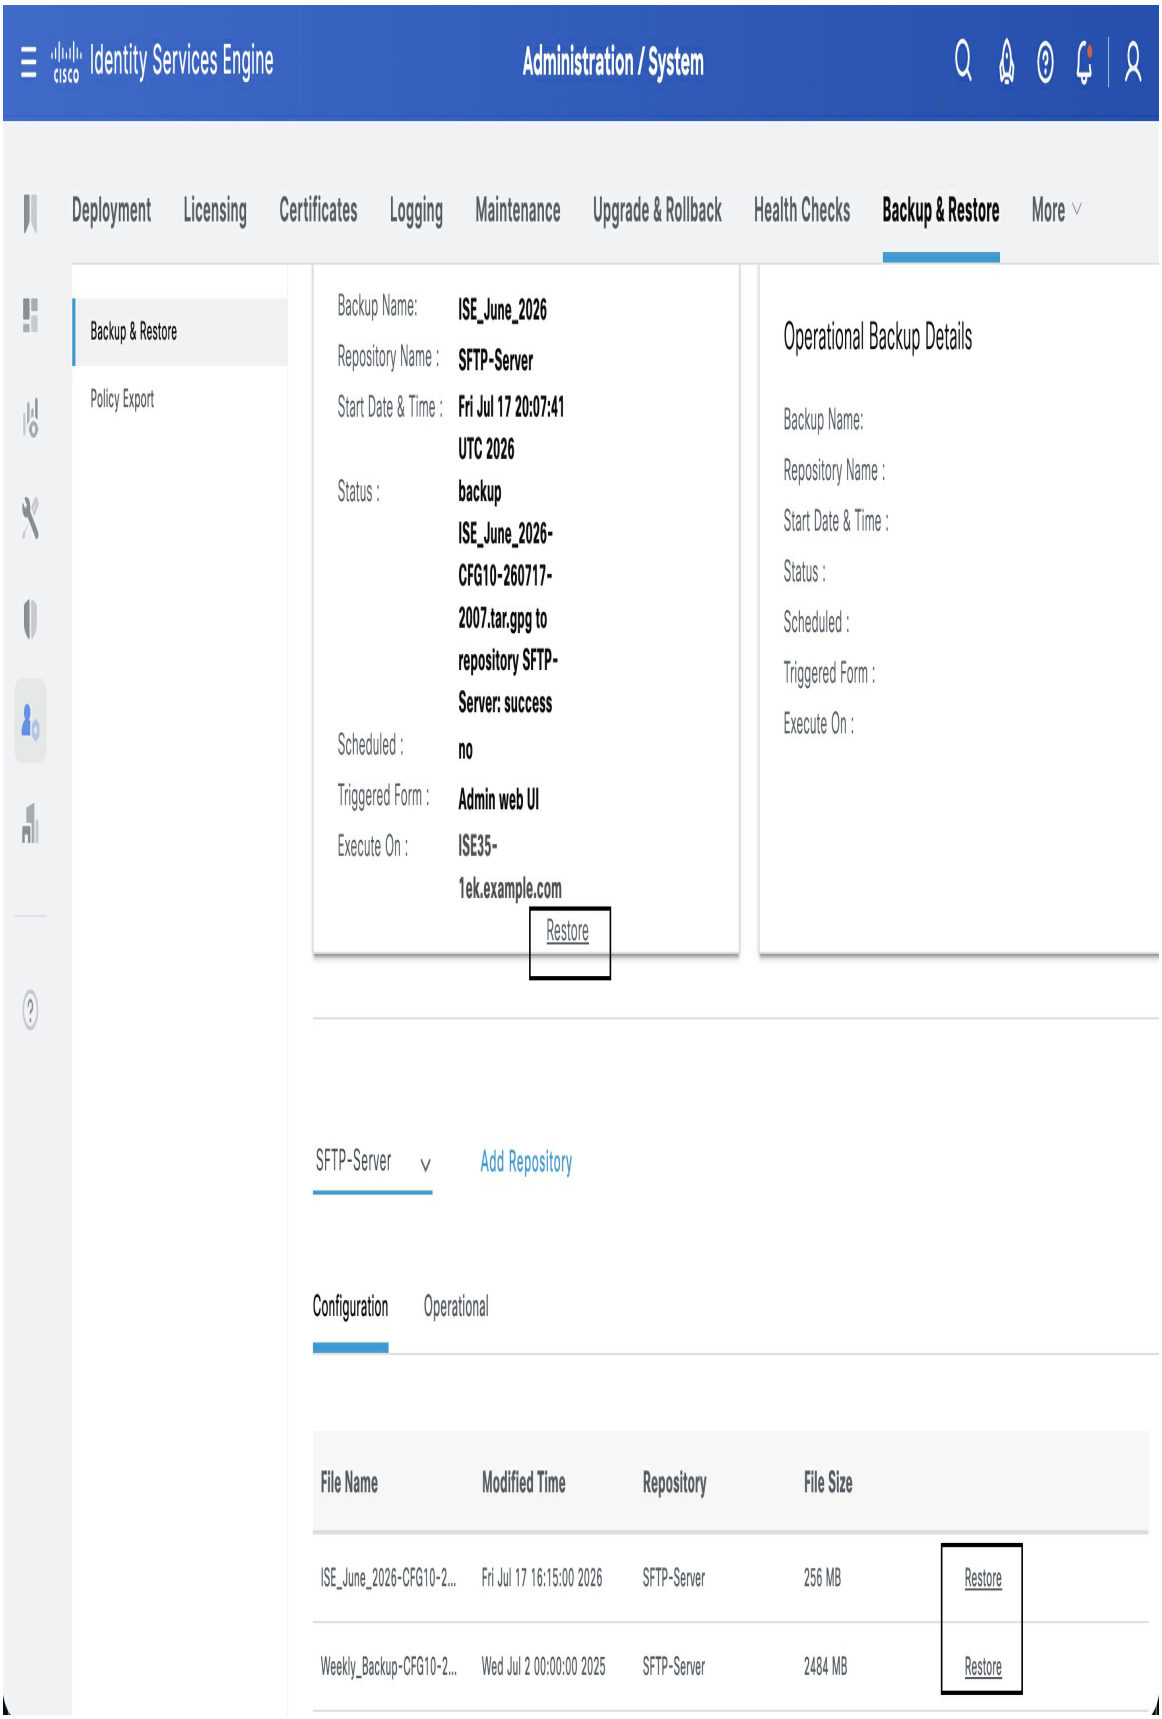
The ISE GUI displays the backup progress, as shown in Figure 21-4.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Backup & Restore
Policy Export
Configuration Data Backup
Stop
Click on 'Stop' to cancel Configuration Data Backup
Backup & Restore
Configurational Backup Details
Backup Name: ISE_June_2026
Repository Name : SFTP-Server
Start Date & Time : Fri Jul 17 20:07:41 UTC 2026
Status : Backup is in progress...
Scheduled : no
Triggered Form : Admin web UI
Execute On : ISE35-1ek.example.com
Operational Backup Details
Backup Name:
Repository Name :
Start Date & Time :
Status :
Scheduled :
Triggered Form :
Execute On :

Technet24

Figure 21-4 *Configuration Data Backup Progress*

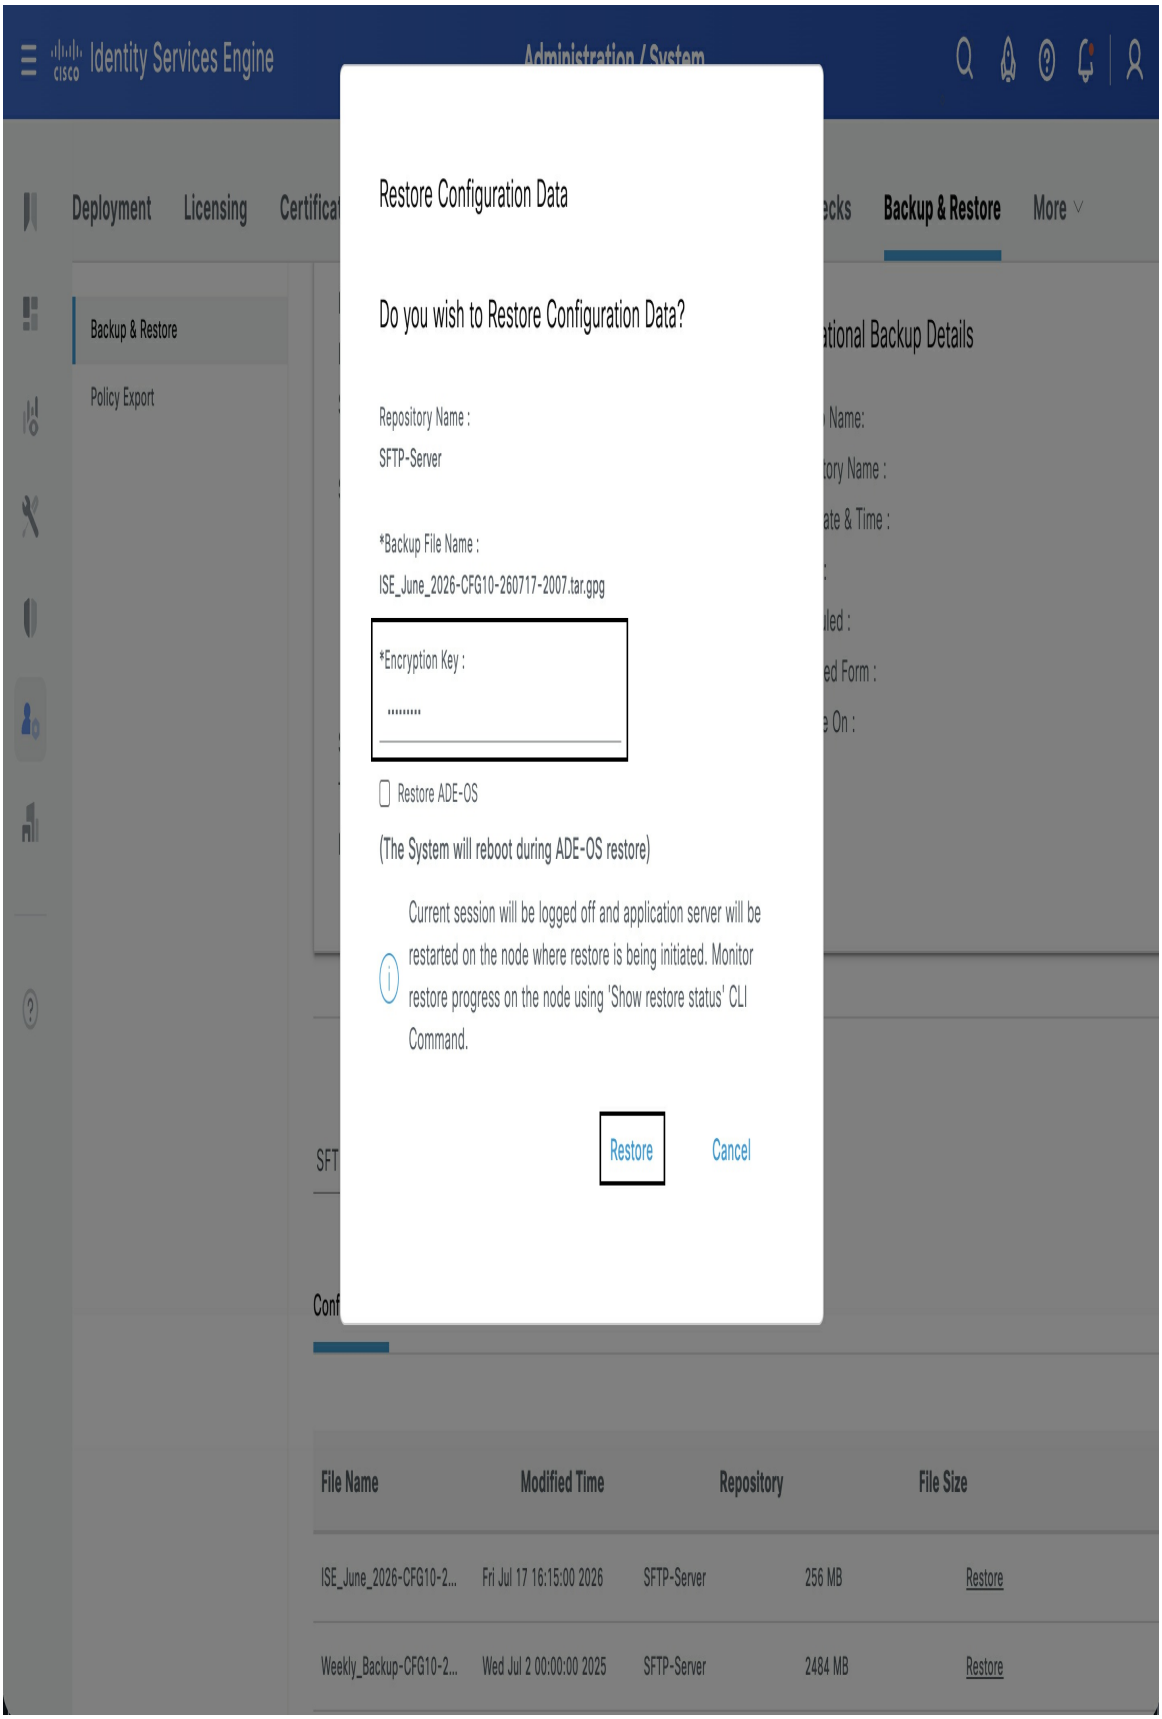
You can restore the latest backup from **Configuration Backup Details** or restore any compatible backup from the repository, as shown in Figure 21-5.


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Backup & Restore
Policy Export
Backup Name: ISE_June_2026
Repository Name : SFTP-Server
Start Date & Time : Fri Jul 17 20:07:41 UTC 2026
Status : backup ISE_June_2026-CFG10-260717-2007.tar.gpg to repository SFTP-Server: success
Scheduled : no
Triggered Form : Admin web UI
Execute On : ISE35-1ek.example.com
Restore
Operational Backup Details
Backup Name:
Repository Name :
Start Date & Time :
Status :
Scheduled :
Triggered Form :
Execute On :
SFTP-Server
Add Repository
Configuration
Operational
File Name
Modified Time
Repository
File Size
ISE_June_2026-CFG10-2...
Fri Jul 17 16:15:00 2026
SFTP-Server
256 MB
Restore
Weekly_Backup-CFG10-2...
Wed Jul 2 00:00:00 2025
SFTP-Server
2484 MB
Restore

Technet24

Figure 21-5 *Configuration Data Restore Options*

When you restore a backup, you must provide the encryption key, as shown in Figure 21-6. No default encryption keys exist for backups. When you restore a configuration backup from the same version, ISE services restart. When you restore a configuration backup with ADE-OS, the system restarts.


Identity Services Engine
Administration / System
Deployment
Licensing
Backup & Restore
More
Backup & Restore
Policy Export
Restore Configuration Data
Do you wish to Restore Configuration Data?
Repository Name :
SFTP-Server
*Backup File Name :
ISE_June_2026-CFG10-260717-2007.tar.gpg
*Encryption Key :
Restore ADE-OS
(The System will reboot during ADE-OS restore)
Current session will be logged off and application server will be restarted on the node where restore is being initiated. Monitor restore progress on the node using 'Show restore status' CLI Command.
Restore
Cancel
File Name
Modified Time
Repository
File Size
ISE_June_2026-CFG10-2...
Fri Jul 17 16:15:00 2026
SFTP-Server
256 MB
Restore
Weekly_Backup-CFG10-2...
Wed Jul 2 00:00:00 2025
SFTP-Server
2484 MB
Restore

Technet24

Figure 21-6 *Configuration Data Restore Details*

You can also take a backup from the CLI by using the following command:

```
#backup ISE_August_Configuration_Backup repository SFTP-Server ise-c
plain Krakow123
```

To restore a backup from the CLI, use the following command:

```
#restore ISE_August_Configuration_Backup repository SFTP-Server enc
Krakow123
```

Optionally, you can include ADE-OS in the restore by adding the **include-adeos** argument.

Cisco ISE does not support VMware snapshots. Therefore, ISE configuration backup is the recommended solution for disaster recovery of the ISE system. If a secondary ISE node becomes unrecoverable for any reason, you can reinstall the node and add it to the deployment without restoring a backup, since all configuration replicates to secondary nodes when you join them to the deployment.

If you use the ISE internal CA, you should export internal CA certificates and keys. You can complete this task from the CLI by using the **application configure ise** option **[7]** and option **[8]**:

```
[7]Export Internal CA Store
[8]Import Internal CA Store
```

You can also use the **restore** command for upgrades. You can restore the backup taken from the previous ISE major release, provided this version is supported in direct upgrade paths. This process upgrades your ISE configuration data to the new release, which is the version where the backup is restored.

Best Practice

Cisco recommends that you schedule weekly configuration backups, or more frequently if you update the configuration often. You should also regularly verify the restore operation because a lost key or corrupted backup will not work during recovery. Ideally, you should store backups in multiple locations to provide additional redundancy.

ISE Upgrades

ISE upgrades ensure that you can access the most advanced features and benefit from new integrations. First, you must verify whether your ISE deployment is eligible for a direct upgrade. For example, if you upgrade to ISE 3.5, you can do so directly from releases 3.2, 3.3, and 3.4.

ISE Upgrade Considerations

If your deployment runs version 3.1 or earlier, you must upgrade to at least 3.2 first.

ISE supports three upgrade methods:

- **GUI:** This method uses the upgrade bundle and is managed through the graphical user interface.
- **CLI:** This method uses the upgrade bundle and is managed through the command-line interface.
- **Backup and Restore:** This method involves taking a backup from the current release, installing a new major release version of ISE, and restoring the backup on the new release.

Best Practice

Technet24

Cisco recommends using GUI-based upgrades. This approach provides multiple health check capabilities that improve the quality of upgrades.

Note

In cloud-based deployments (AWS, Azure, OCI), the Backup and Restore method is the only supported upgrade option.

ISE Upgrade Preparation

You must address several key elements before performing an upgrade:

1. Always create the latest configuration backup and ensure you have the encryption key. Regardless of whether you use the CLI or GUI for the upgrade, a backup must always be available.
2. Run the Upgrade Readiness Tool (URT) bundle. The URT bundle performs a data and schema upgrade on a cloned database. It can also detect and fix issues where possible. URT bundles present no risk, so running them is highly recommended.

   To install the URT bundle, run the command shown in Example 21-1 on the Secondary Policy Administration Node (SPAN). Running the URT on the Primary Policy Administrative Node (PPAN) is not recommended.

Example 21-1 *URT Bundle Installation*

```
ISE-2/admin#application install ise-urtbundle-3.5.0.527-1.0.0.SPA
Server

Getting bundle to local machine...
Unbundling Application Package...
Verifying Application Signature...
Initiating Application Install...
```

```
##########################################
# Installing Upgrade Readiness Tool (URT) #
##########################################

Virtual Machines (VMs) with specs equivalent to Cisco SNS 3595 ar
Cisco ISE Release 3.3. For assistance in choosing supported Cisco
see the Cisco ISE Installation Guide.
Checking ISE version compatibility
  - Successful

Checking ISE persona
  - Successful

Along with Administration, other services (MNT) are enabled on th
running URT might consume additional resources.
Do you want to proceed with installing and running URT now (y/n):

Checking if URT is recent(<45 days old)
  - Note URT is 287 days old and its version is 1.0.0. There migh
on CCO, please verify on CCO
Do you want to proceed with this version which is 287 days old (y
Proceeding with this version of URT itself

Installing URT bundle
```

After all checks are completed, the URT bundle is removed from the system. The bundle provides an estimated upgrade duration in minutes before removal, as shown in Example 21-2.

Example 21-2 *URT Estimated Upgrade Duration*

```
Running data upgrade for node specific data on cloned database
  - Successful

Time estimate for upgrade
=========================
(Estimates are calculated based on size of config and mnt data on
between PAN and other nodes is not considered in calculating esti
Estimated time for each node (in mins):
```

Technet24

```
ISE-2(SECONDARY PAP,MNT):91
ISE-1(PRIMARY PAP,MNT):74
Each PSN(2 if in parallel):64


Final cleanup before exiting...

Application successfully installed
ISE-2/admin#
```

Note

In this scenario, ISE 3.4 patch 6 (deployment of four nodes) is the release where the URT was installed to prepare for the upgrade to 3.5. Therefore, the 3.5 URT bundle is used (ise-urtbundle-3.5.0.527-1.0.0.SPA.x86_64.tar.gz), and the bundle is stored in the SFTP repository named SFTP-Server.

3. Always review the latest Upgrade Guide on cisco.com. Each major upgrade may include specific version information that you must consider.
4. Always verify the port requirements for the target release version. New major releases often introduce new features that rely on additional protocols and functions; for example, you may need to open new ports between ISE nodes.
5. Export system certificates with private keys and export the CA certificate chain. You can export system certificates from **Administration > System > Certificates > Certificate Management > System Certificates**. You can export the CA certificate chain using the **application configure ISE** CLI command with option **[7] Export Internal CA Store**.

Performing the Upgrade

The ISE GUI upgrade provides two types of upgrade options:

- **Full upgrade:** This method suits environments that permit downtime, such as upgrades performed after working hours when you can ensure that user or endpoint authentication is not required during the upgrade. The process upgrades all ISE nodes in parallel, which accelerates the upgrade.
- **Split upgrade:** This method involves a multistep process that allows you to upgrade ISE nodes in batches. A split upgrade enables services to remain available throughout the upgrade. It is beneficial for environments that operate 24/7 and cannot tolerate downtime.

Let's review the steps required to upgrade a medium deployment of four nodes, as shown in Figure 21-7, from version 3.4 patch 6 to 3.5 using the split upgrade method through the GUI.

Technet24

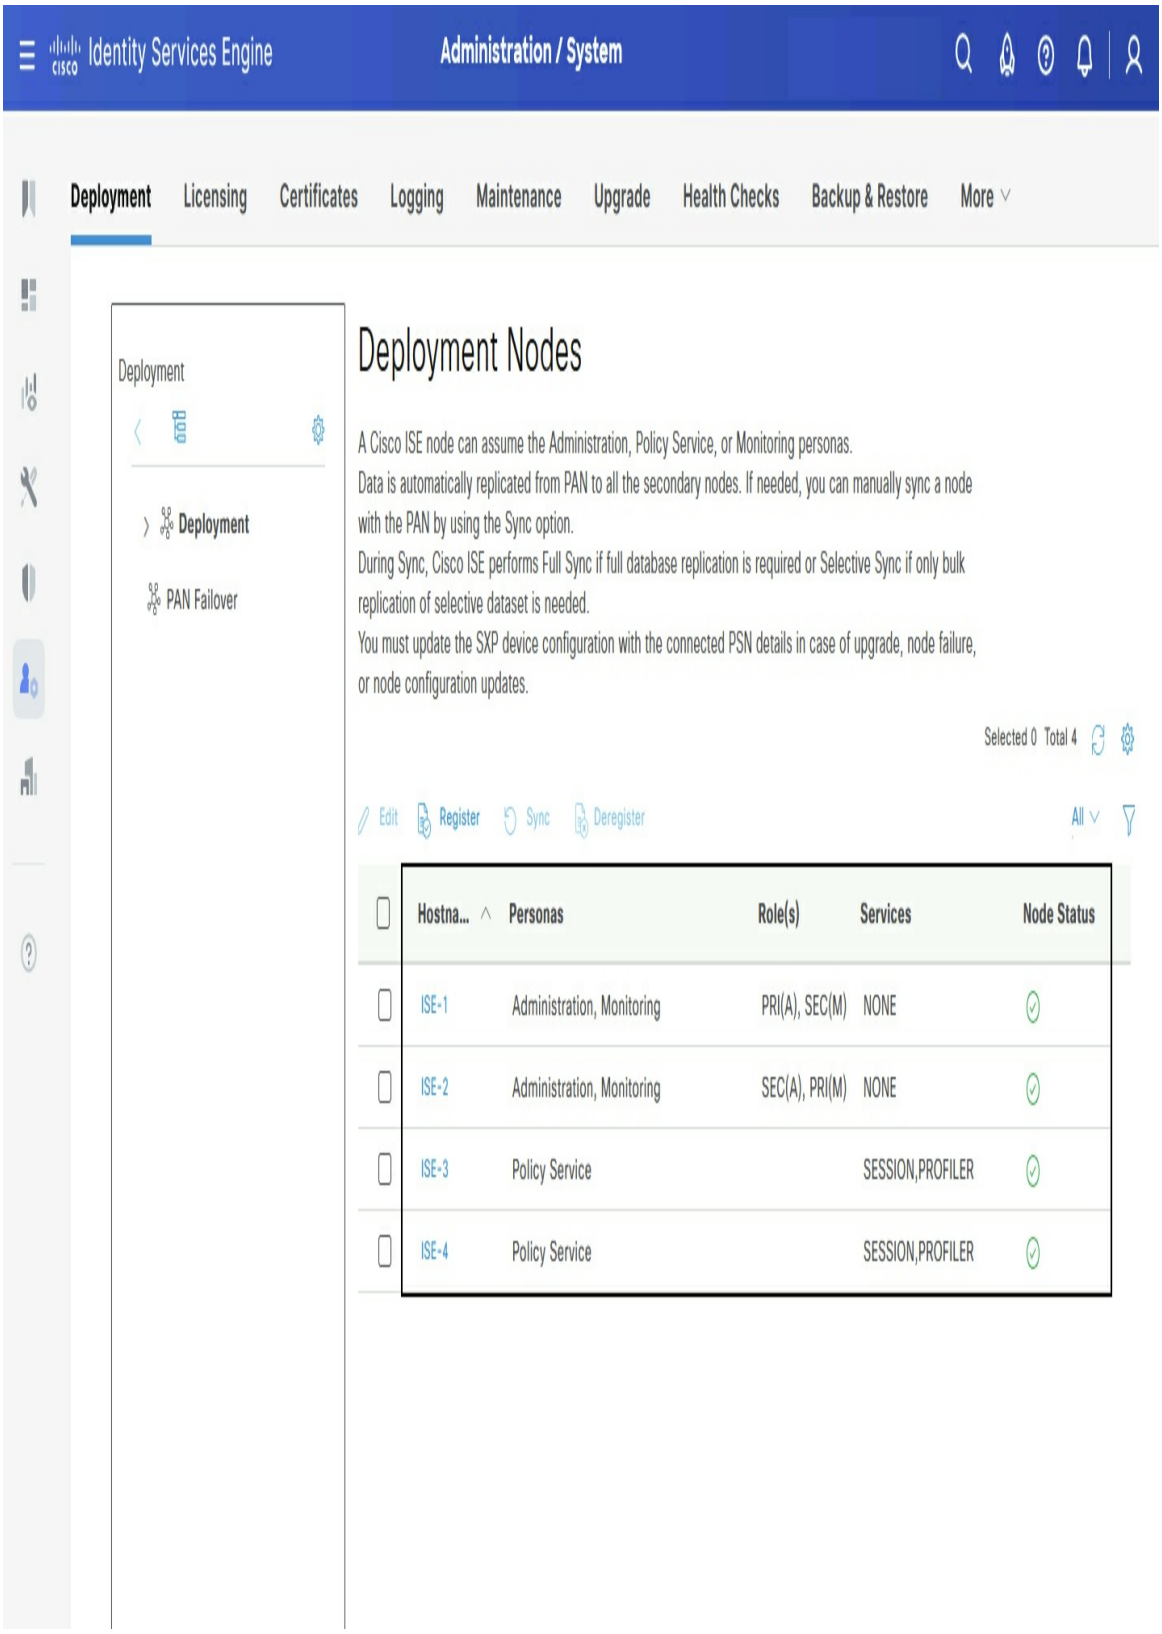


Figure 21-7 *ISE Medium Deployment Roles Allocation Before the*

Upgrade

Navigate to **Administration > System > Upgrade & Rollback**, as shown in Figure 21-8. In order from the top of the page to the bottom, select **Upgrade**, **Release**, **Split** and click **Initialize**.

Technet24

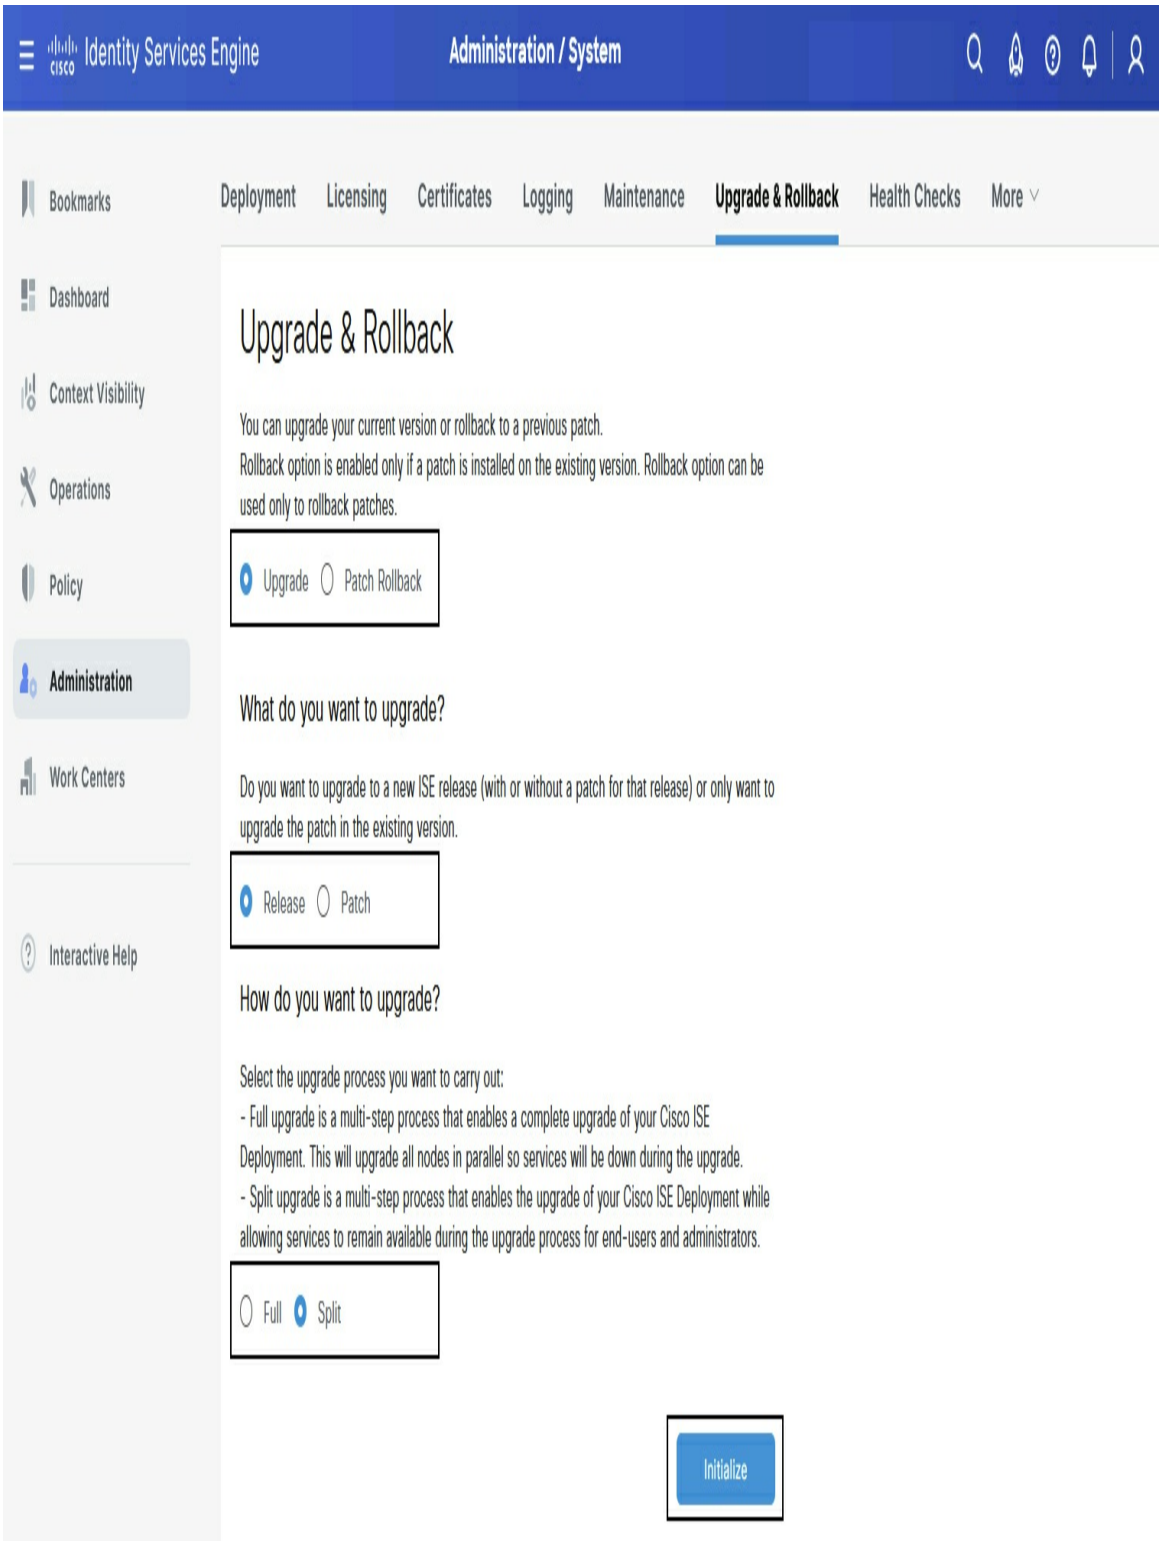


Figure 21-8 *The ISE Upgrade & Rollback GUI Menu for a Split Upgrade*

ISE services continue to operate during the Prepare to Upgrade and Upgrade Staging steps. Services suspend on the selected nodes when you reach the Upgrade Nodes step (Step 6).

Now, click **Let's Do It** to access the Upgrade Checklist, as shown in Figure 21-9. Complete each step of the checklist, acknowledge completion, and click **Next**.

Best Practice

Be sure to complete all steps of the checklist before acknowledging them. These steps are critical for the success of the upgrade.

Technet24

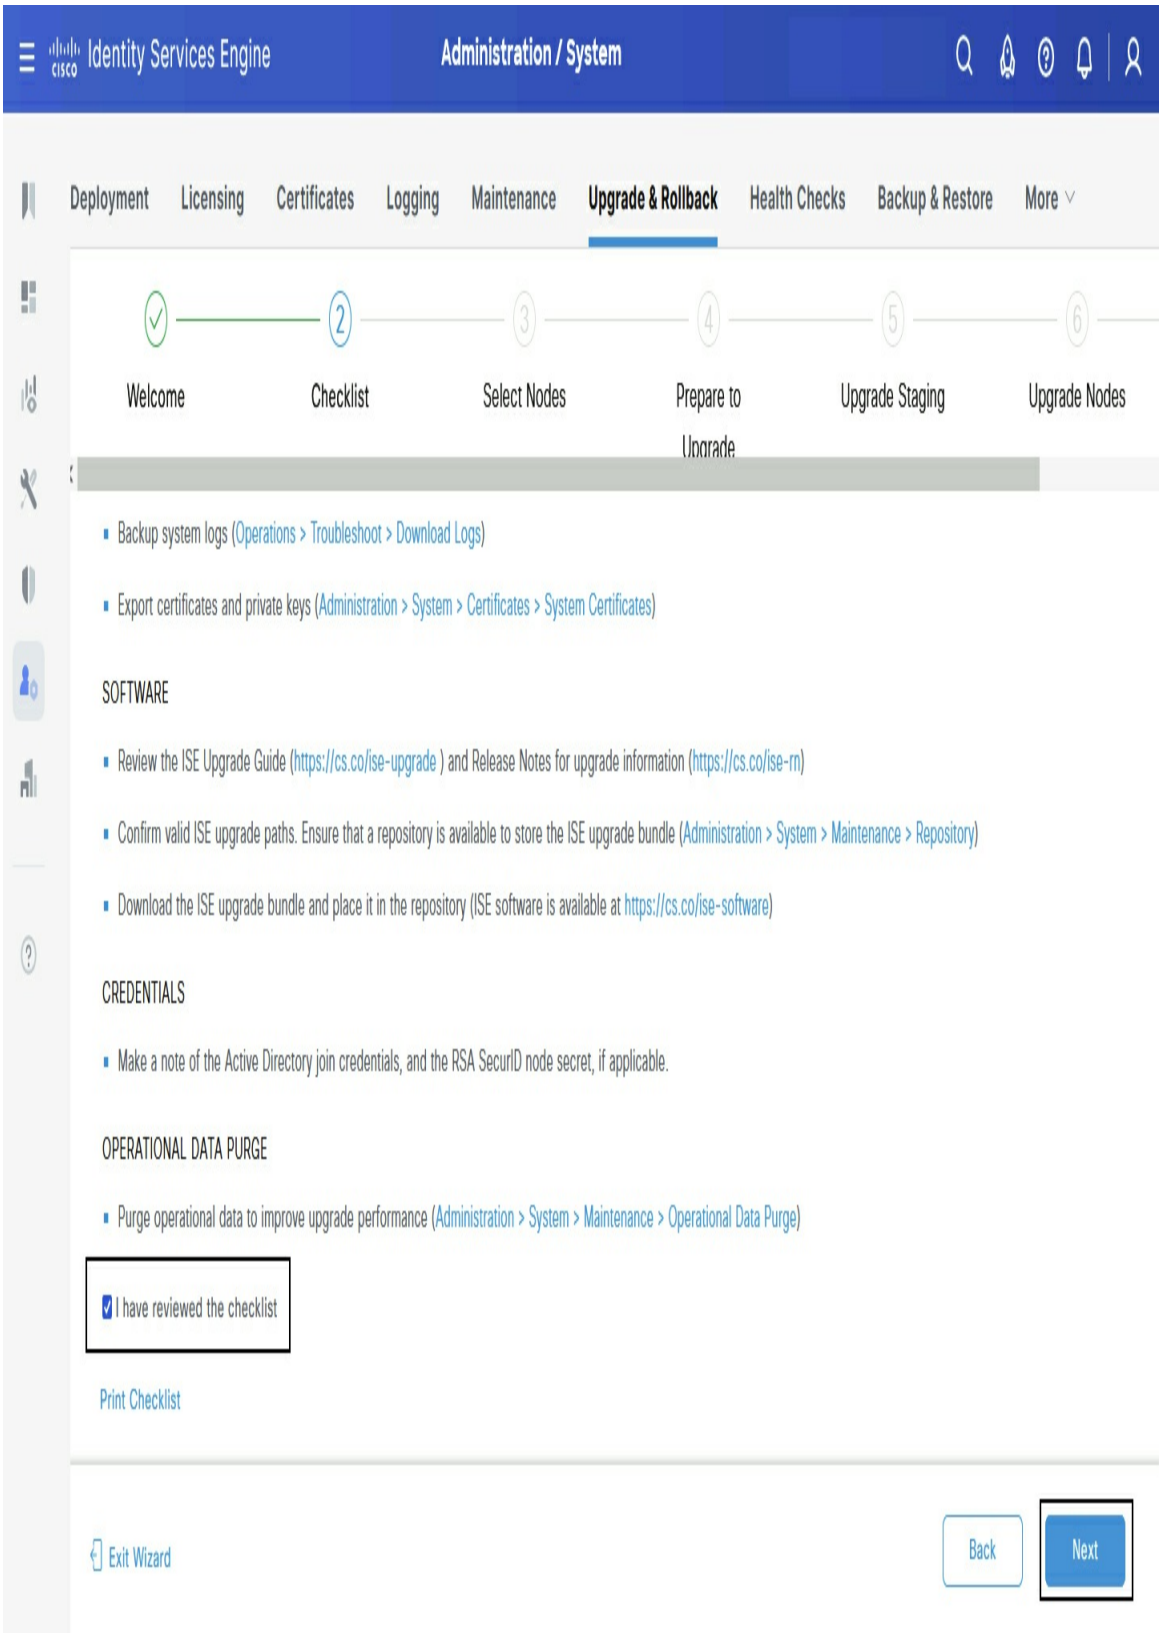


Figure 21-9 *ISE Upgrade Checklist*

Because you are performing a split upgrade, you must select the nodes for the first iteration of the upgrade. SPAN is already selected because it should be the first node to upgrade. You cannot select PPAN until all other nodes are upgraded.

During the first iteration, one of the PSNs (ISE-4) and SPAN (ISE-2) are selected for the upgrade, as shown in Figure 21-10. This approach ensures that a PSN and PPAN remain available to handle user authentications while other nodes undergo the upgrade process. When you're ready, click **Next**.

Technet24

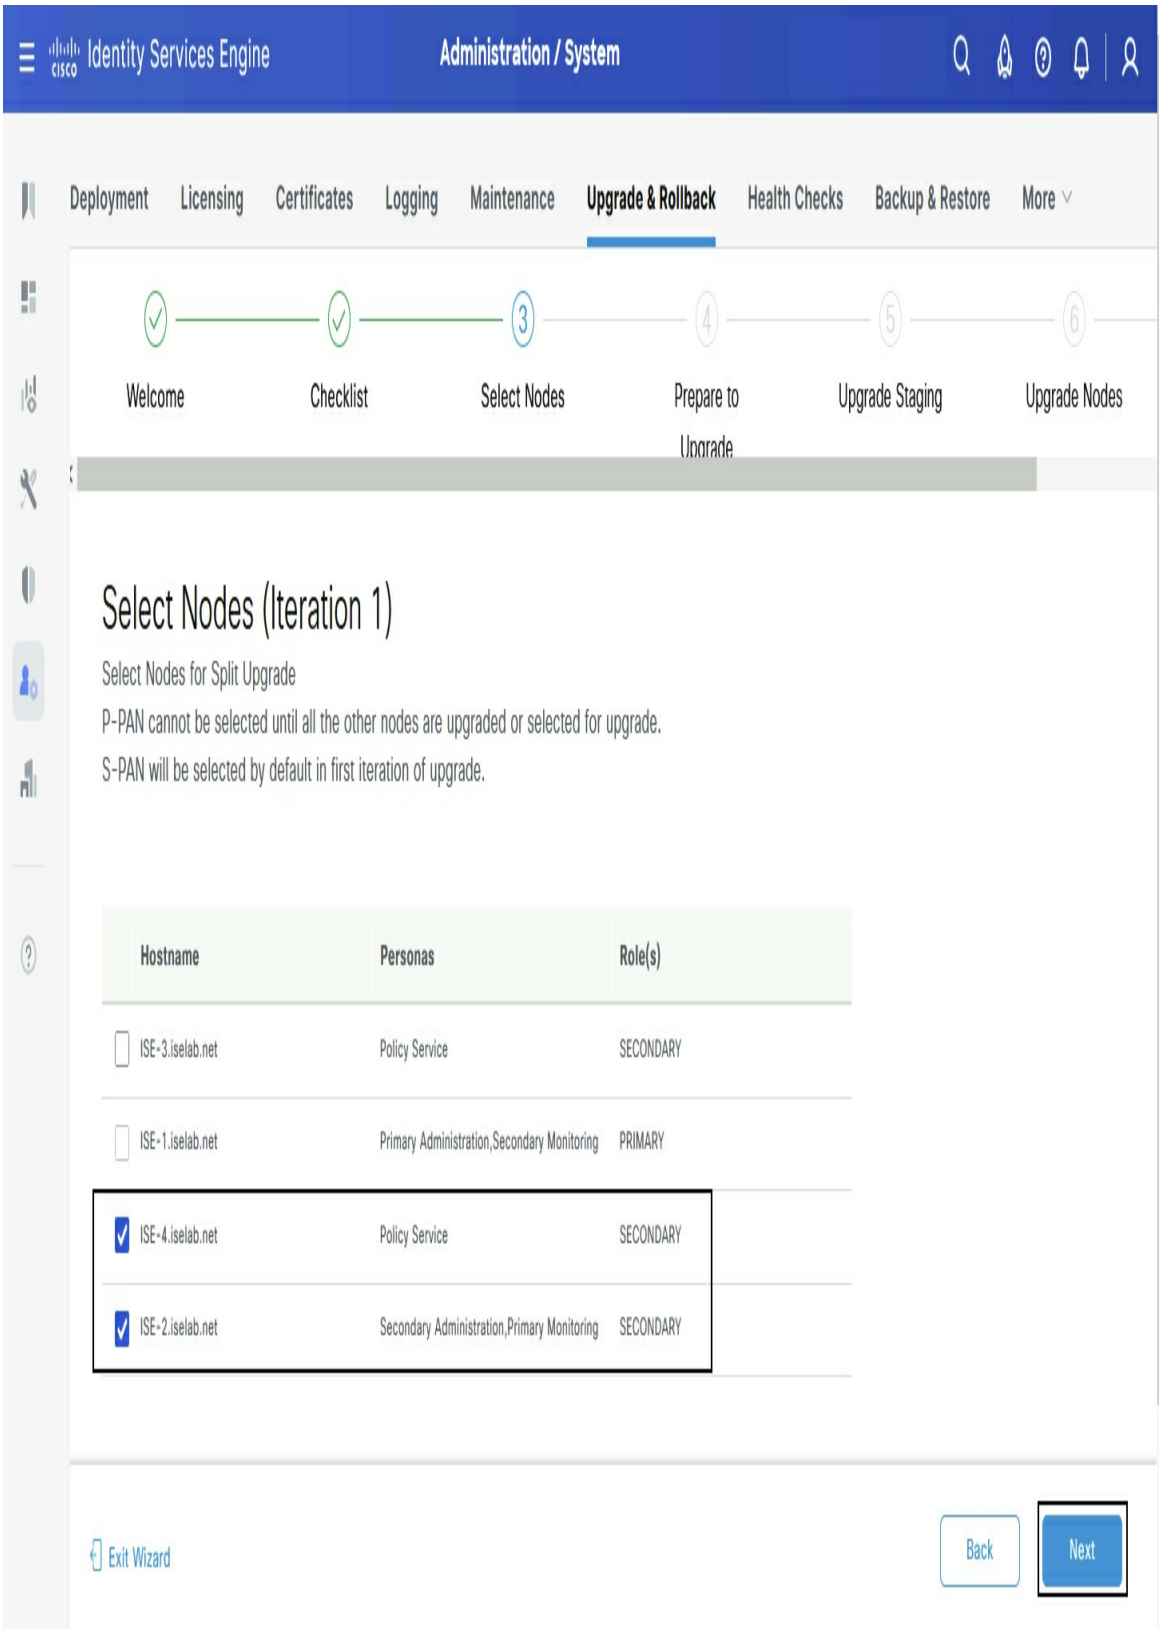


Figure 21-10 *ISE Split Upgrade Iteration 1 Node Selection*

From the repository, select the upgrade bundle file and click **Start Preparation**, as shown in Figure 21-11. Within the upgrade framework, ISE allows you to select the patch, which will be installed as part of the upgrade process after the major upgrade completes.

Technet24

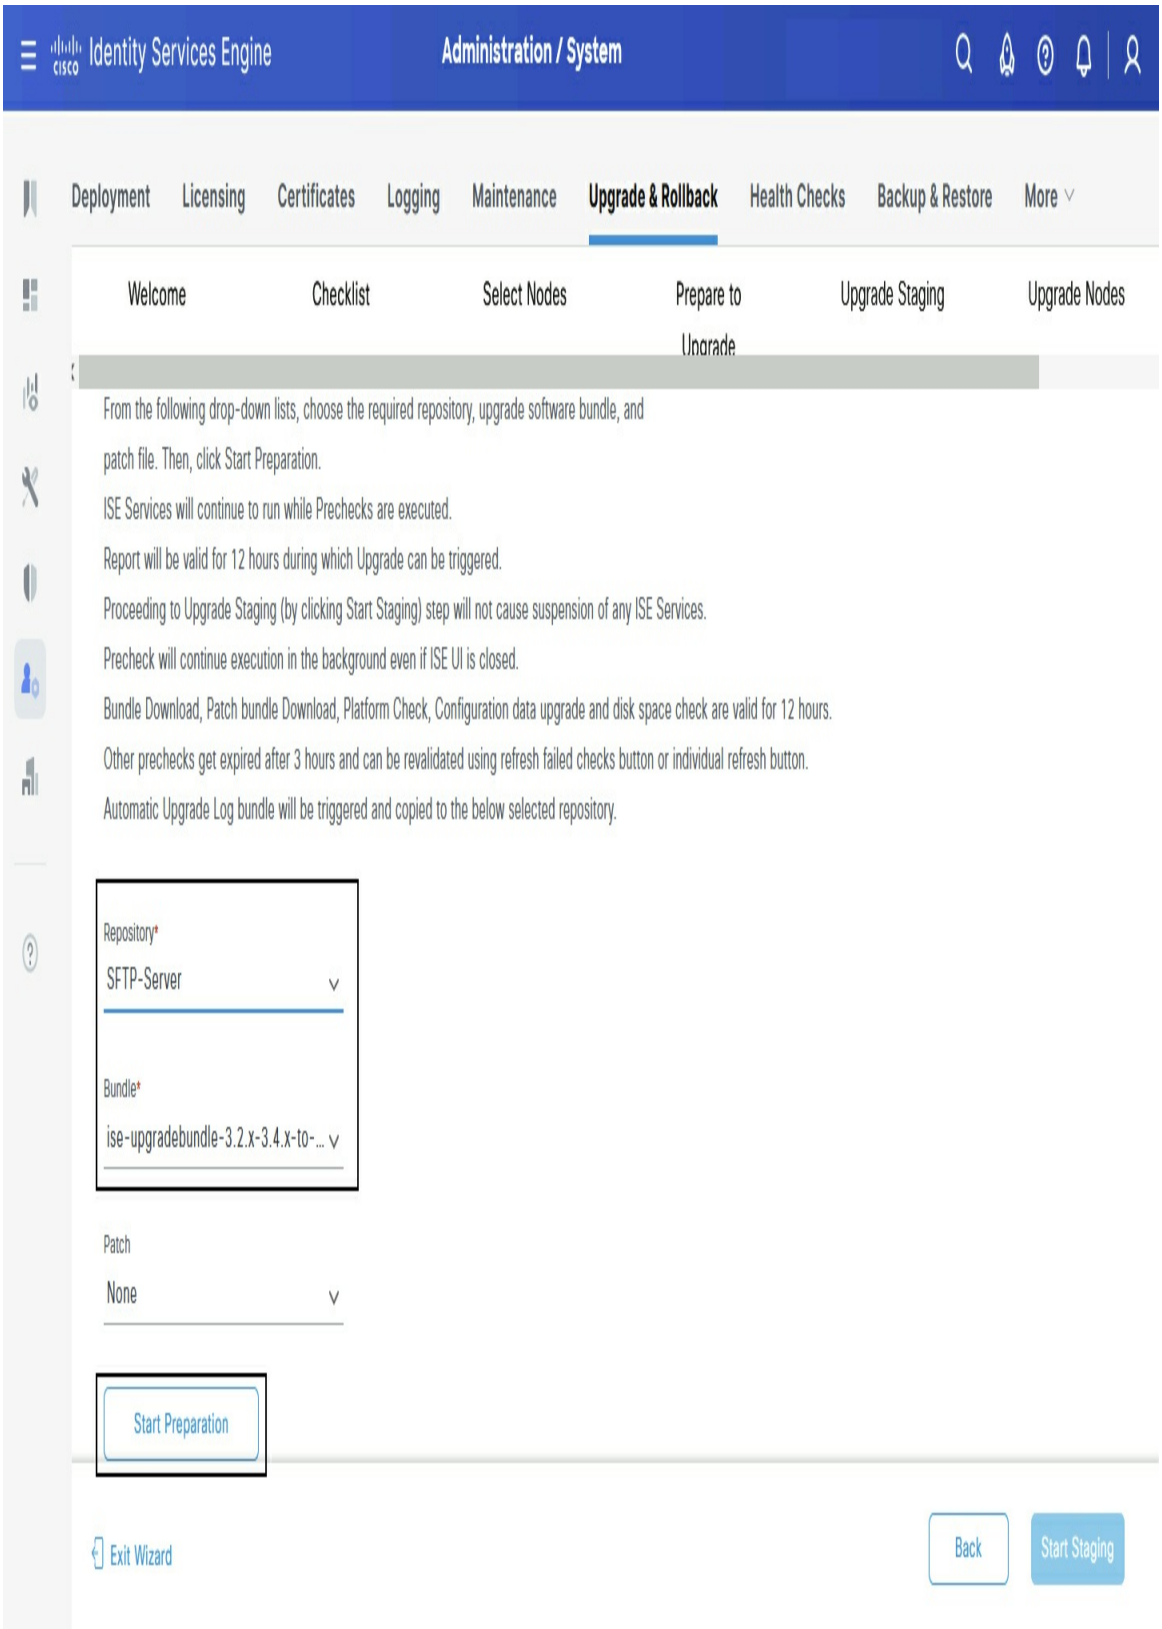


Figure 21-11 *ISE Split Upgrade Iteration 1 Repository and Bundle*

Selection

The preparation process, as shown in Figure 21-12, includes downloading the upgrade bundle to the required nodes in the deployment, performing certificate checks, conducting platform checks, and completing additional verifications to ensure the success of your upgrade.

Technet24

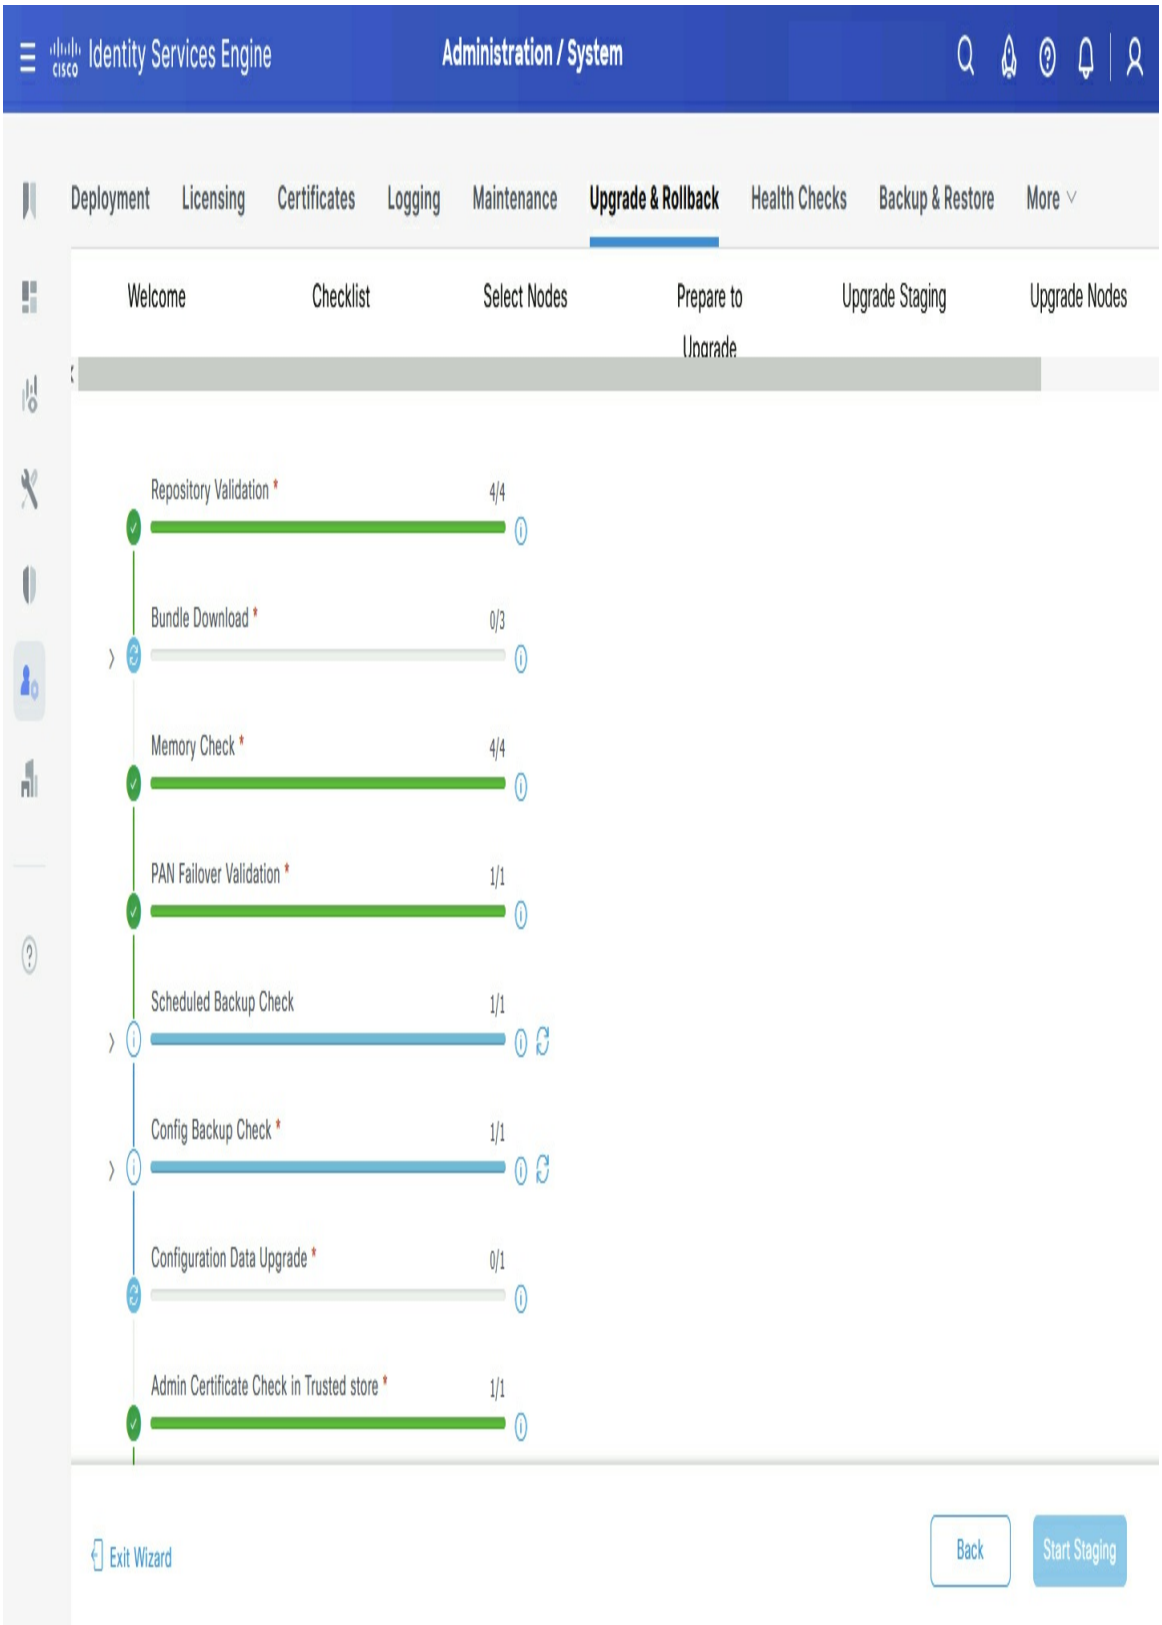


Figure 21-12 *ISE Split Upgrade Iteration 1 Health Check Progress*

If any health checks fail, the **Start Staging** button remains grayed out until you resolve those health check problems. Remediation steps are available in the details of the failed health check.

After you fix any identified health check issues, click the **Start Staging** button when it becomes active, as shown in Figure 21-13. All nodes still remain operational at this stage.

Technet24


Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
Upgrade Staging
Upgrade Nodes
Repository Validation * 4/4
Bundle Download * 3/3
Memory Check * 4/4
PAN Failover Validation * 1/1
Scheduled Backup Check 1/1
Config Backup Check * 1/1
Configuration Data Upgrade * 1/1
Admin Certificate Check in Trusted store * 1/1
Exit Wizard
Back
Start Staging

Figure 21-13 *ISE Split Upgrade Iteration 1 Health Check Completion*

During upgrade staging, the system copies the upgraded database file to all nodes, as shown in Figure 21-14. The system also backs up configuration files for each node. This process saves time during the upgrade because the database upgrade does not need to be performed on each node individually. When you're ready, click **Next**.

Technet24

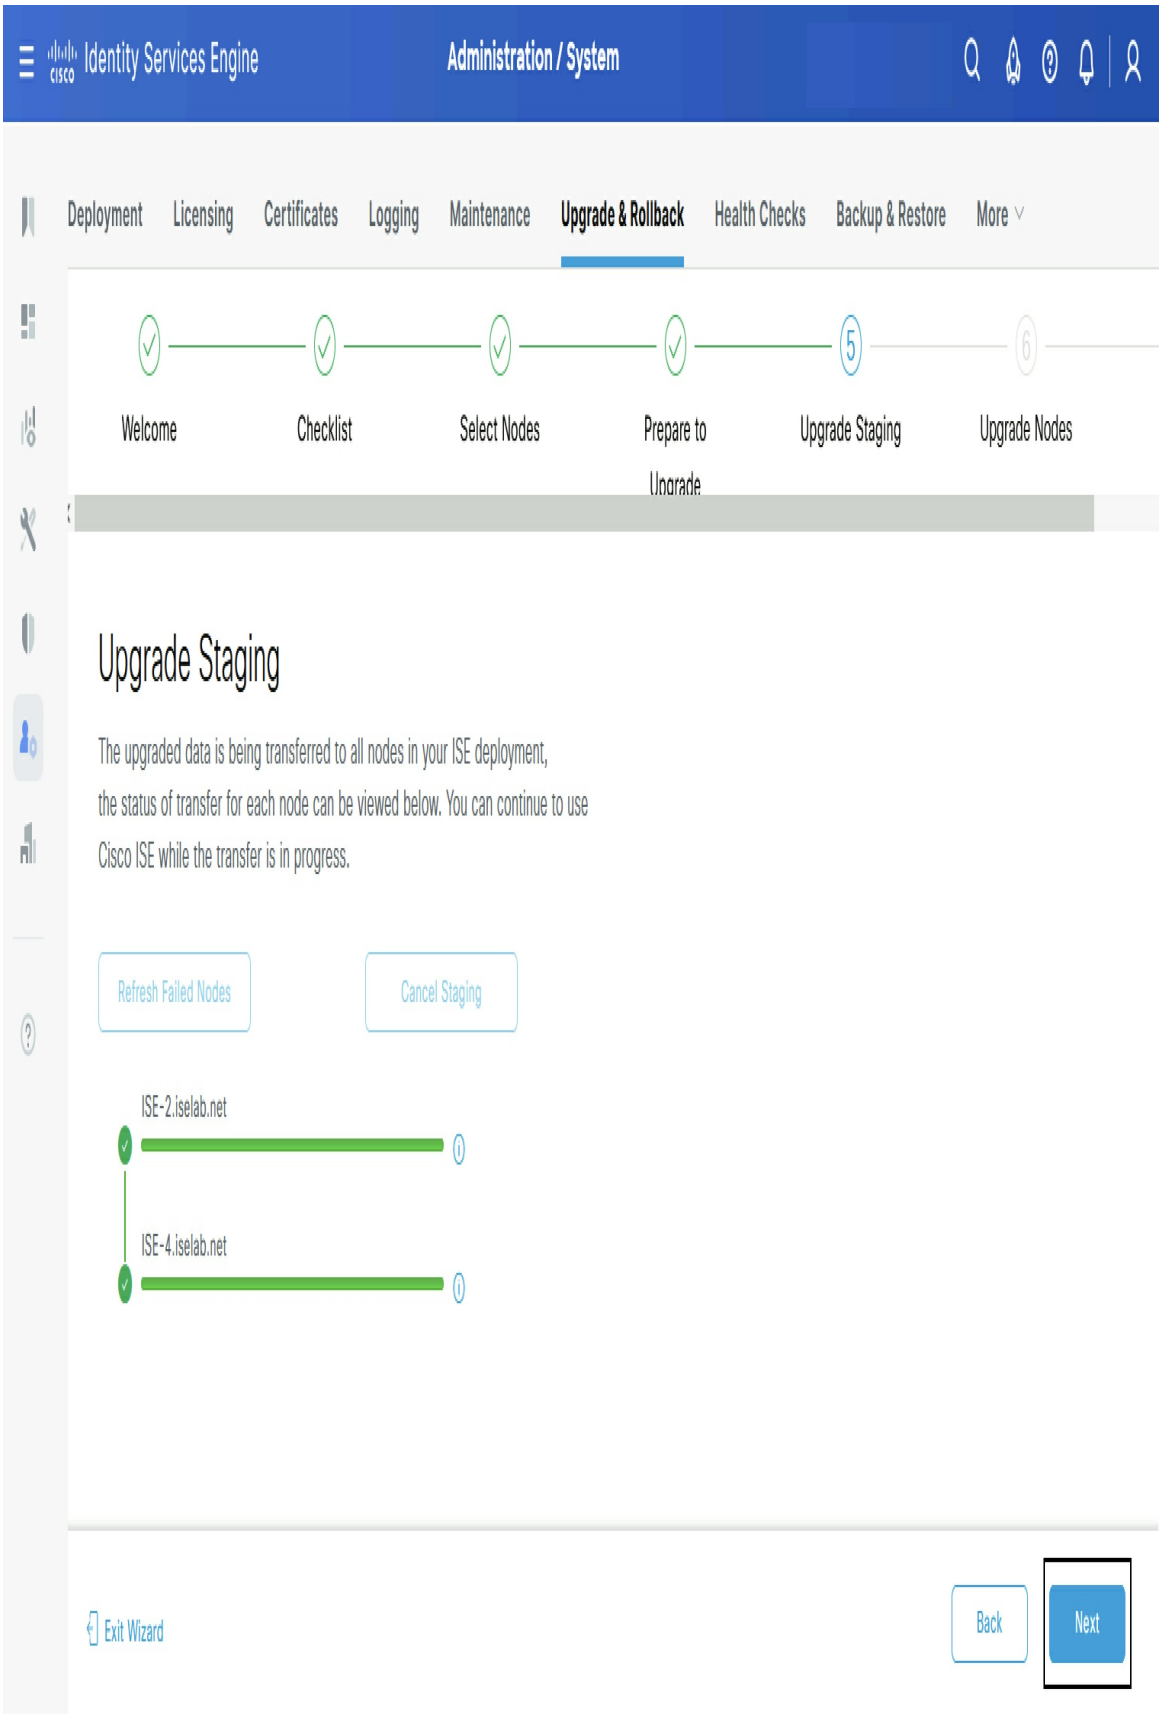

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
5
Upgrade Staging
6
Upgrade Nodes
Upgrade Staging
The upgraded data is being transferred to all nodes in your ISE deployment, the status of transfer for each node can be viewed below. You can continue to use Cisco ISE while the transfer is in progress.
Refresh Failed Nodes
Cancel Staging
ISE-2.iselab.net
ISE-4.iselab.net
Exit Wizard
Back
Next

Figure 21-14 *ISE Split Upgrade Iteration 1 Staging*

Important

At this point, if you have not done so previously, it is time to divert all authentications to the ISE-3 node (PSN, which is not included in the first iteration). You can do this at the network access devices level or at the load balancer level if they are used. Having said that, high availability configuration on network access devices should be able to take care of this task as well, switching to the PSN that is not getting upgraded.

Click **Start** to begin the first iteration of the upgrade, as shown in Figure 21-15. This process will restart services on the nodes, making them unavailable for authentication.

Technet24

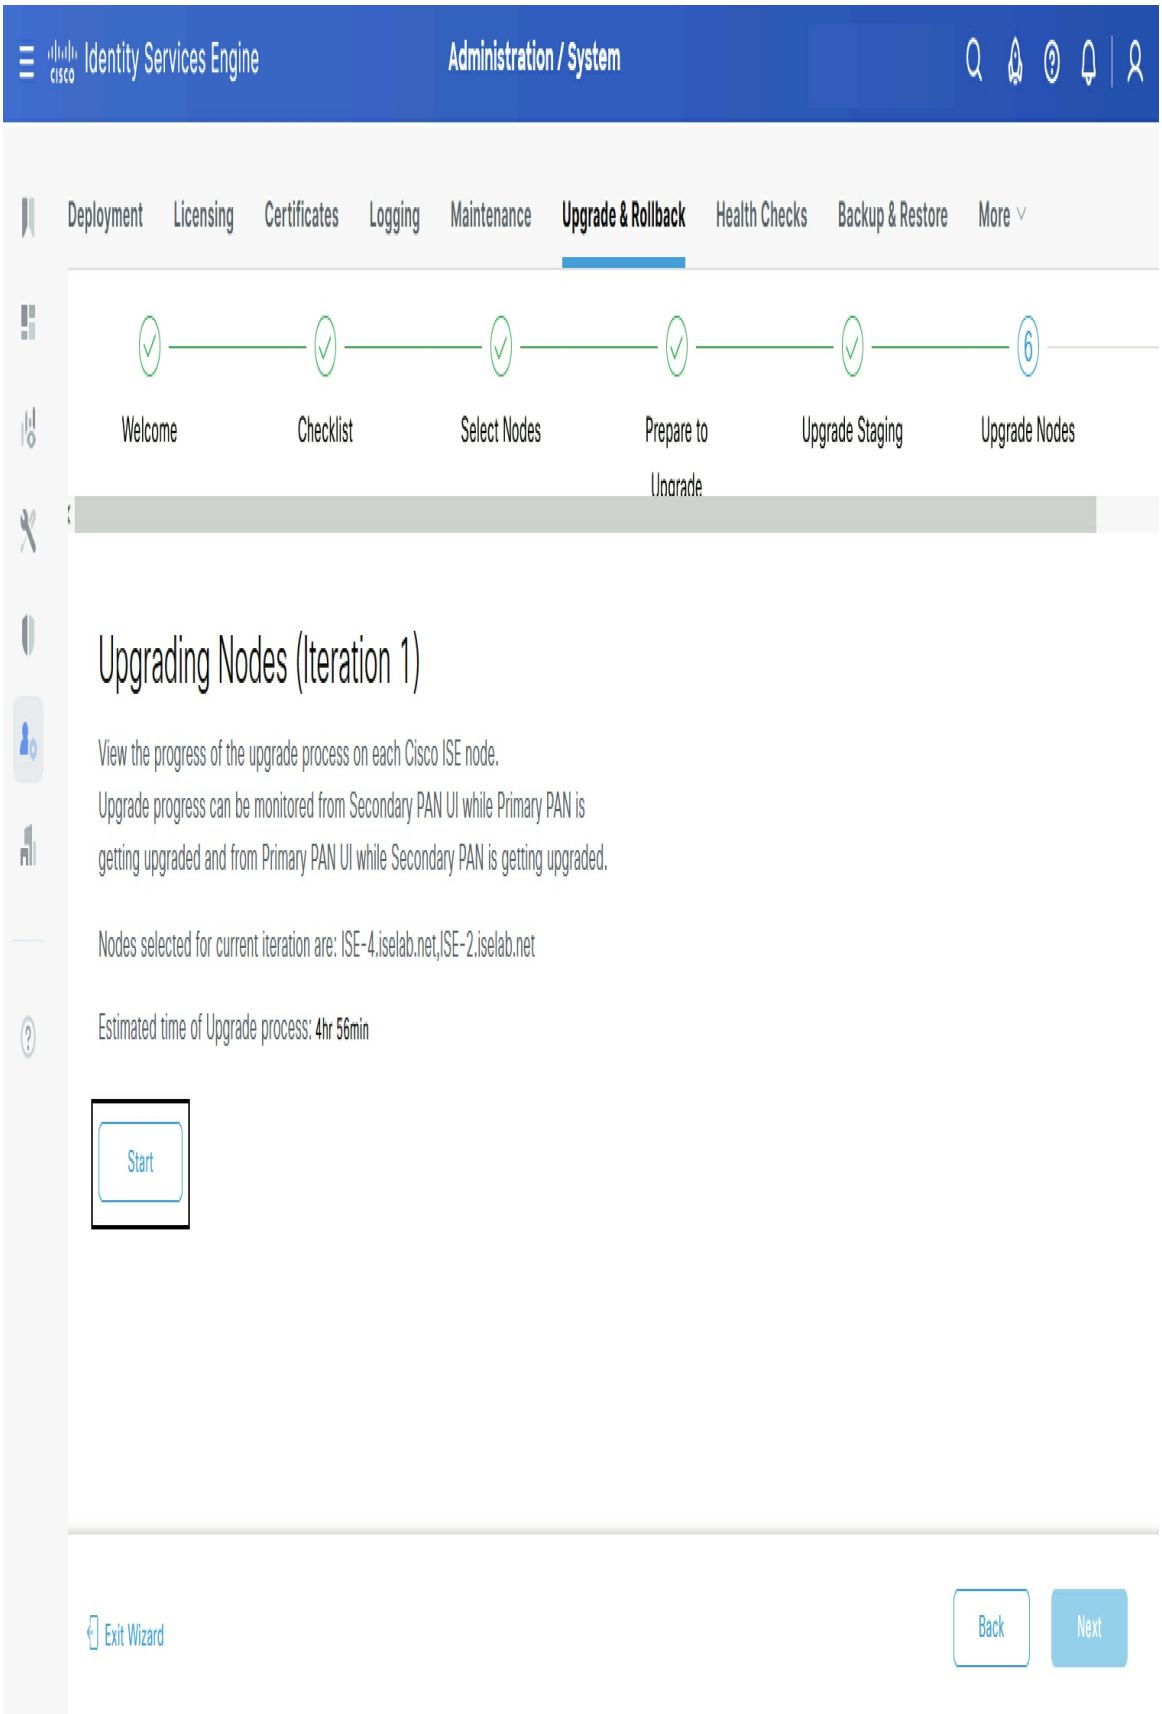

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
Upgrade Staging
Upgrade Nodes
6
Upgrading Nodes (Iteration 1)
View the progress of the upgrade process on each Cisco ISE node.
Upgrade progress can be monitored from Secondary PAN UI while Primary PAN is getting upgraded and from Primary PAN UI while Secondary PAN is getting upgraded.
Nodes selected for current iteration are: ISE-4.iselab.net,ISE-2.iselab.net
Estimated time of Upgrade process: 4hr 56min
Start
Exit Wizard
Back
Next

Figure 21-15 *ISE Split Upgrade Iteration 1 Start*

Accept the restart prompt by clicking **Yes**, as shown in Figure 21-16.

Technet24

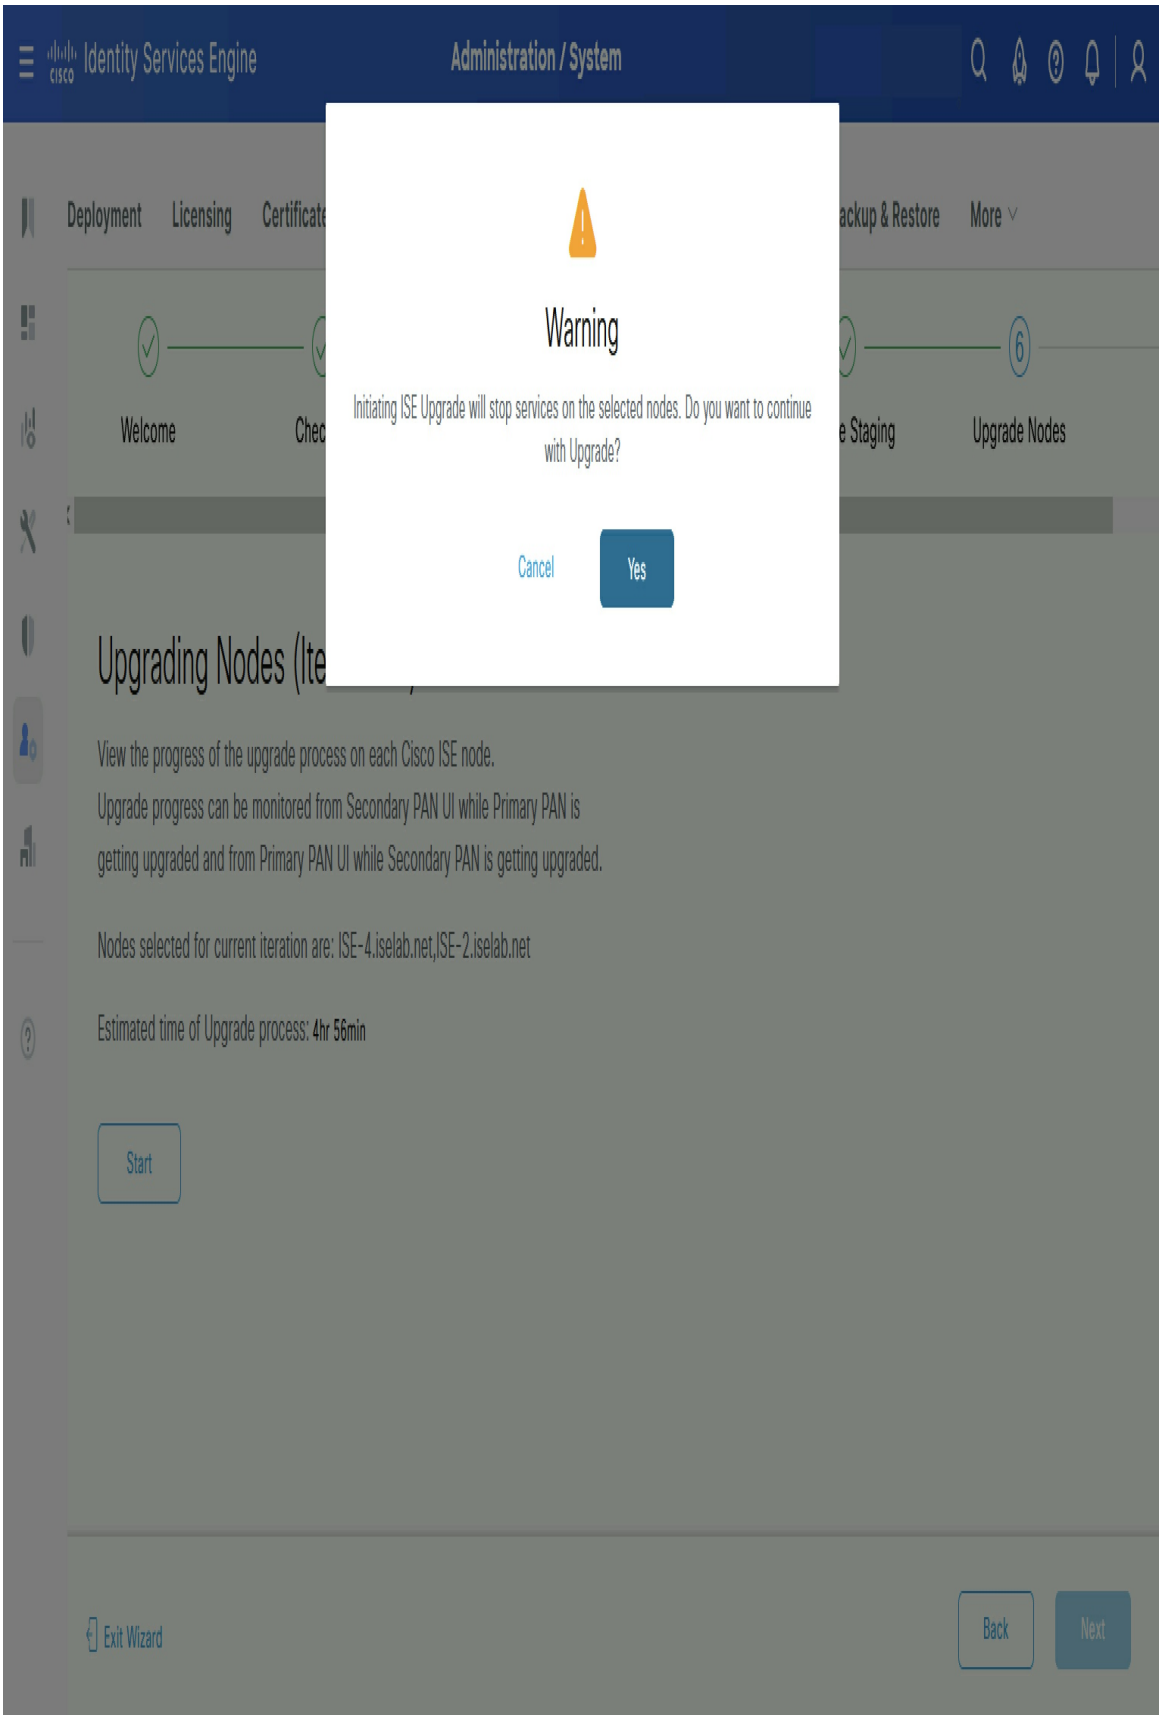

Identity Services Engine
Administration / System
Deployment
Licensing
Certificate
ackup & Restore
More
Welcome
Chec
e Staging
6
Upgrade Nodes
Warning
Initiating ISE Upgrade will stop services on the selected nodes. Do you want to continue with Upgrade?
Cancel
Yes
Upgrading Nodes (Ite
View the progress of the upgrade process on each Cisco ISE node.
Upgrade progress can be monitored from Secondary PAN UI while Primary PAN is getting upgraded and from Primary PAN UI while Secondary PAN is getting upgraded.
Nodes selected for current iteration are: ISE-4.iselab.net,ISE-2.iselab.net
Estimated time of Upgrade process: 4hr 56min
Start
Exit Wizard
Back
Next

Figure 21-16 *ISE Split Upgrade Iteration 1 Services Restart Warning*

At this stage, the most critical part, which is the actual upgrade, begins. A split upgrade always starts with SPAN. After the SPAN upgrade completes, the PSNs begin upgrading; in this example, PSN ISE-4 undergoes the upgrade. You can monitor the upgrade progress from the PPAN GUI, as shown in Figure 21-17.

Technet24

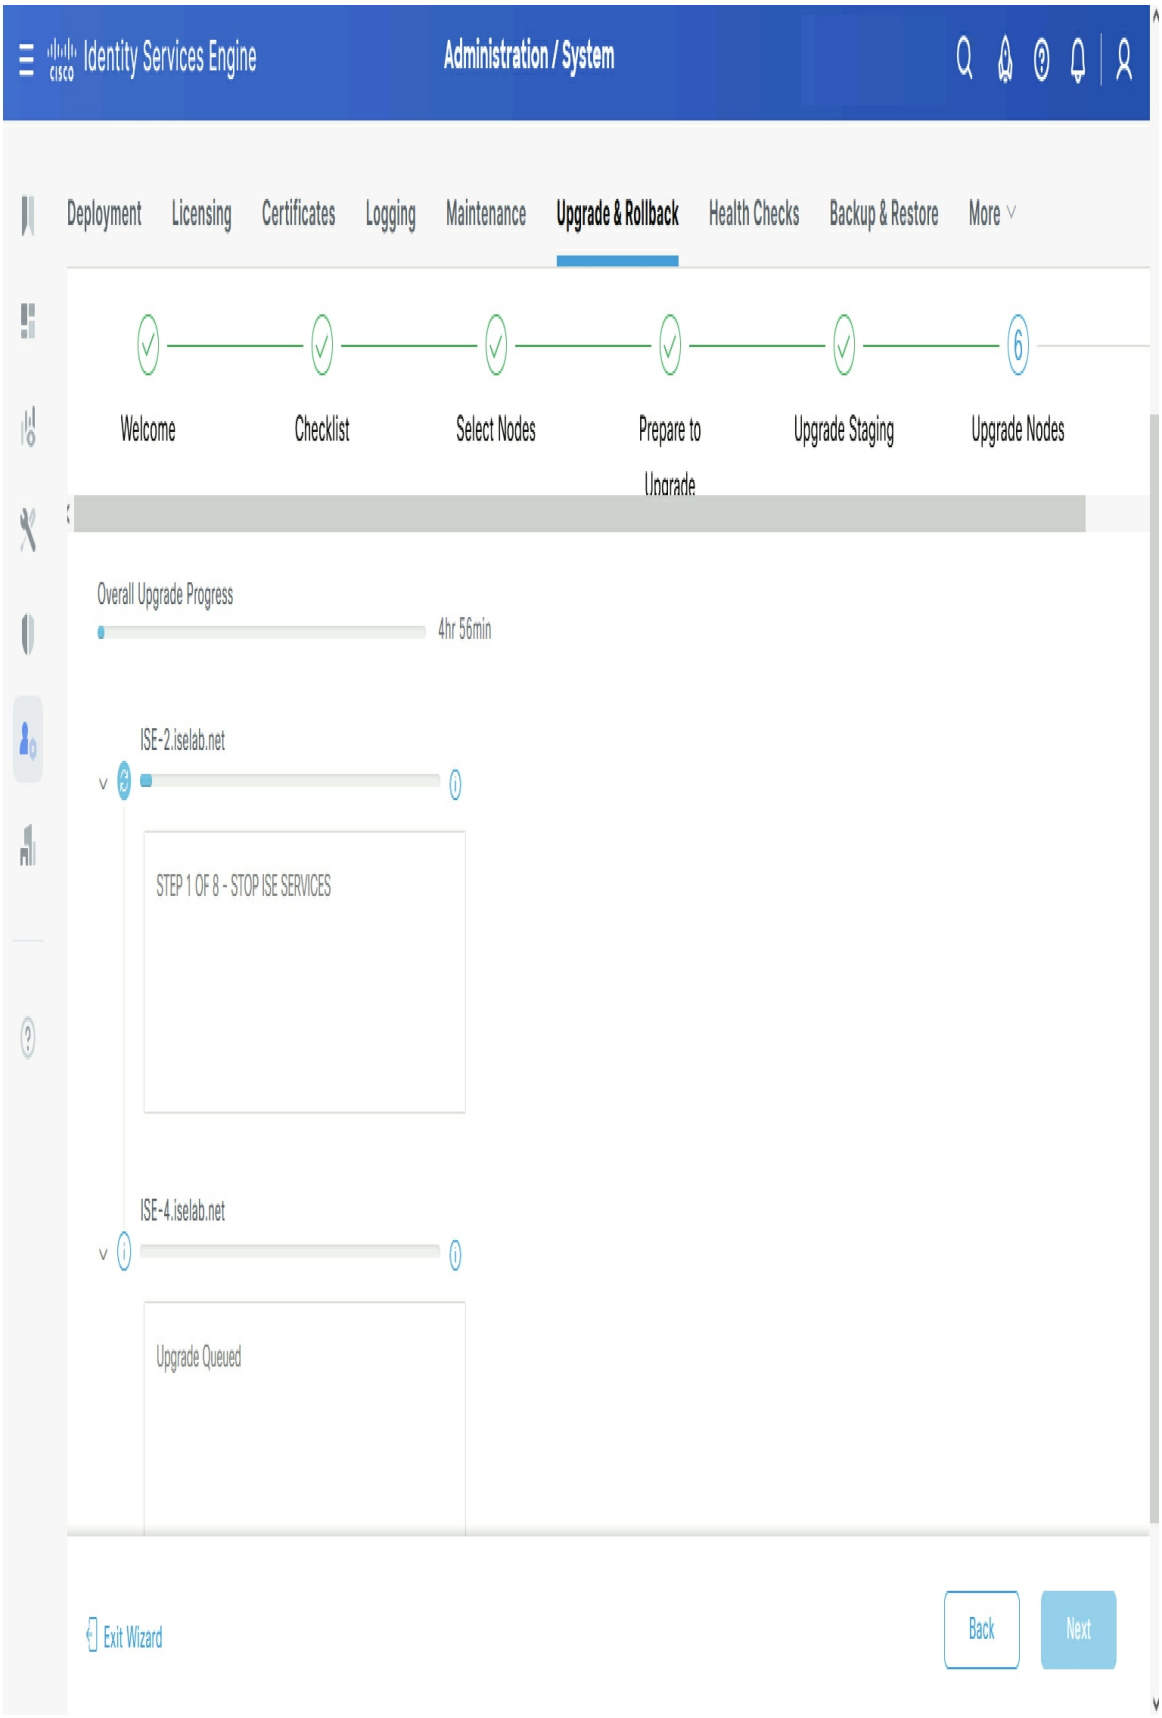

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
Upgrade Staging
Upgrade Nodes
Overall Upgrade Progress
4hr 56min
ISE-2.iselab.net
STEP 1 OF 8 - STOP ISE SERVICES
ISE-4.iselab.net
Upgrade Queued
Exit Wizard
Back
Next

Figure 21-17 *ISE Split Upgrade Iteration 1 Progress from Original PPAN (ISE-1)*

After the upgrade completes on both nodes, the Upgrade UI displays green checkboxes next to the upgraded nodes and enables the **Next** button for the next iteration. At this point, click **Next**, as shown in Figure 21-18.

Technet24

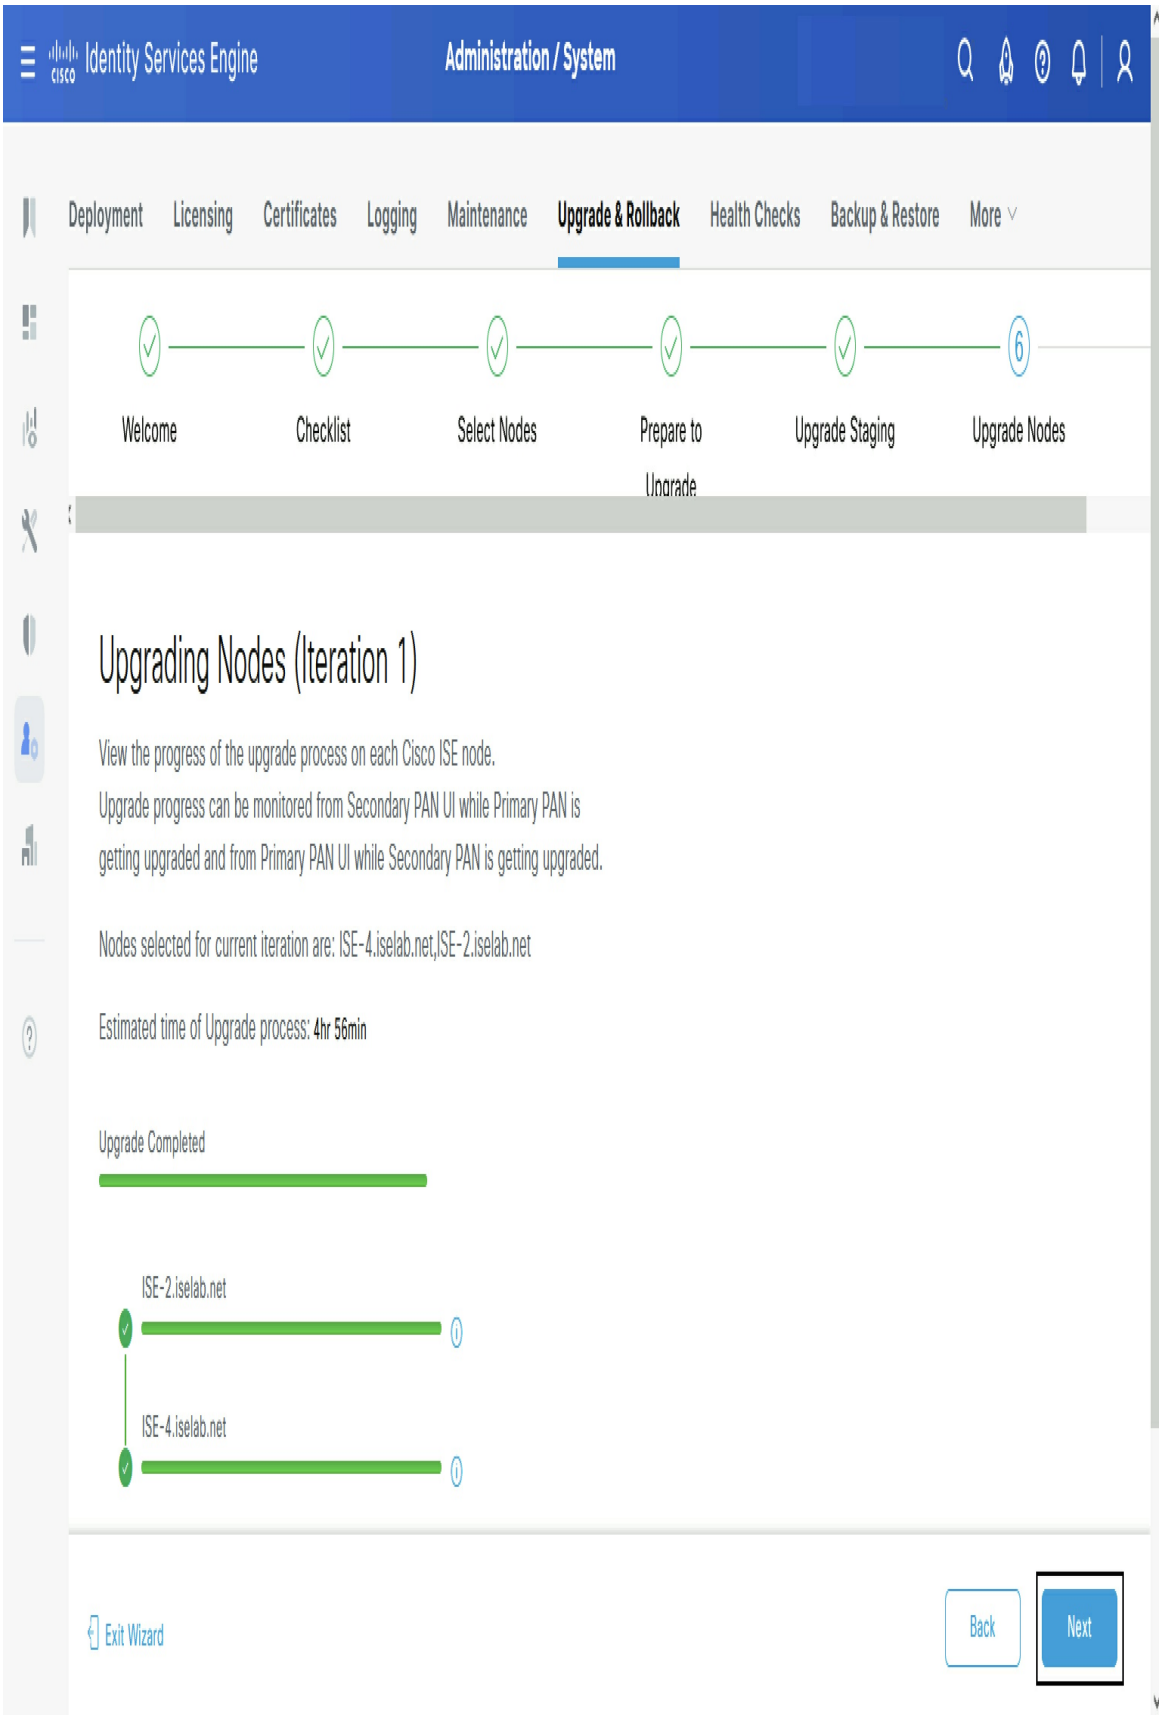

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
Upgrade Staging
6
Upgrade Nodes
Upgrading Nodes (Iteration 1)
View the progress of the upgrade process on each Cisco ISE node.
Upgrade progress can be monitored from Secondary PAN UI while Primary PAN is getting upgraded and from Primary PAN UI while Secondary PAN is getting upgraded.
Nodes selected for current iteration are: ISE-4.iselab.net,ISE-2.iselab.net
Estimated time of Upgrade process: 4hr 56min
Upgrade Completed
ISE-2.iselab.net
ISE-4.iselab.net
Exit Wizard
Back
Next

Figure 21-18 *ISE Split Upgrade Iteration 1 Completion*

You can download the upgrade reports for your reference. Click **Continue**, as shown in Figure 21-19.

Technet24

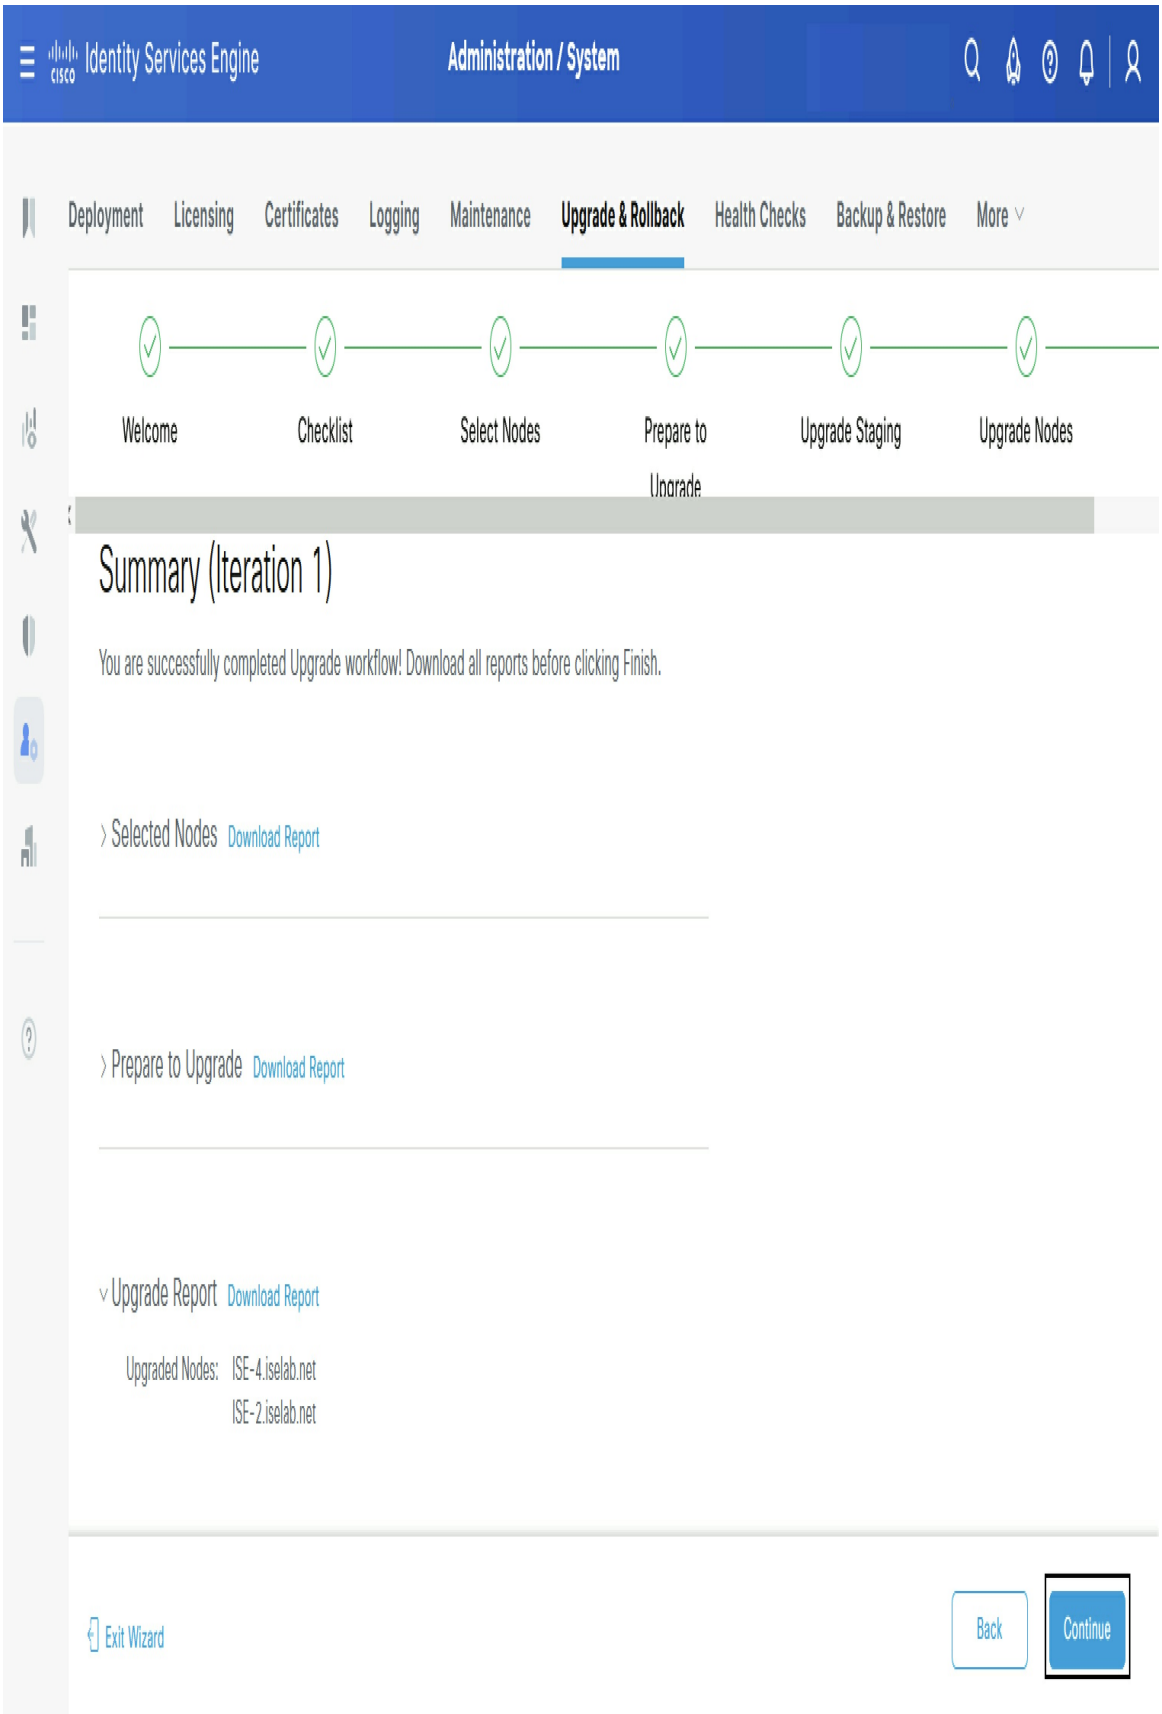

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
Upgrade Staging
Upgrade Nodes
Summary (Iteration 1)
You are successfully completed Upgrade workflow! Download all reports before clicking Finish.
Selected Nodes
Download Report
Prepare to Upgrade
Download Report
Upgrade Report
Download Report
Upgraded Nodes:
ISE-4.iselab.net
ISE-2.iselab.net
Exit Wizard
Back
Continue

Figure 21-19 *ISE Split Upgrade Iteration 1 Summary*

Note, at this stage, the second deployment becomes operational. The former SPAN assumes the role of PPAN in the new deployment; PSN-4 is upgraded and joined to the new deployment. Now, you need to redirect traffic to PSN-4, verify it is operational, and ensure no issues are present. After completing this verification, you can proceed to the next iteration to finish the upgrade. Select **ISE-3 PSN** and **ISE-1 PPAN** from the old deployment; then click **Next**, as shown in Figure 21-20.

Technet24

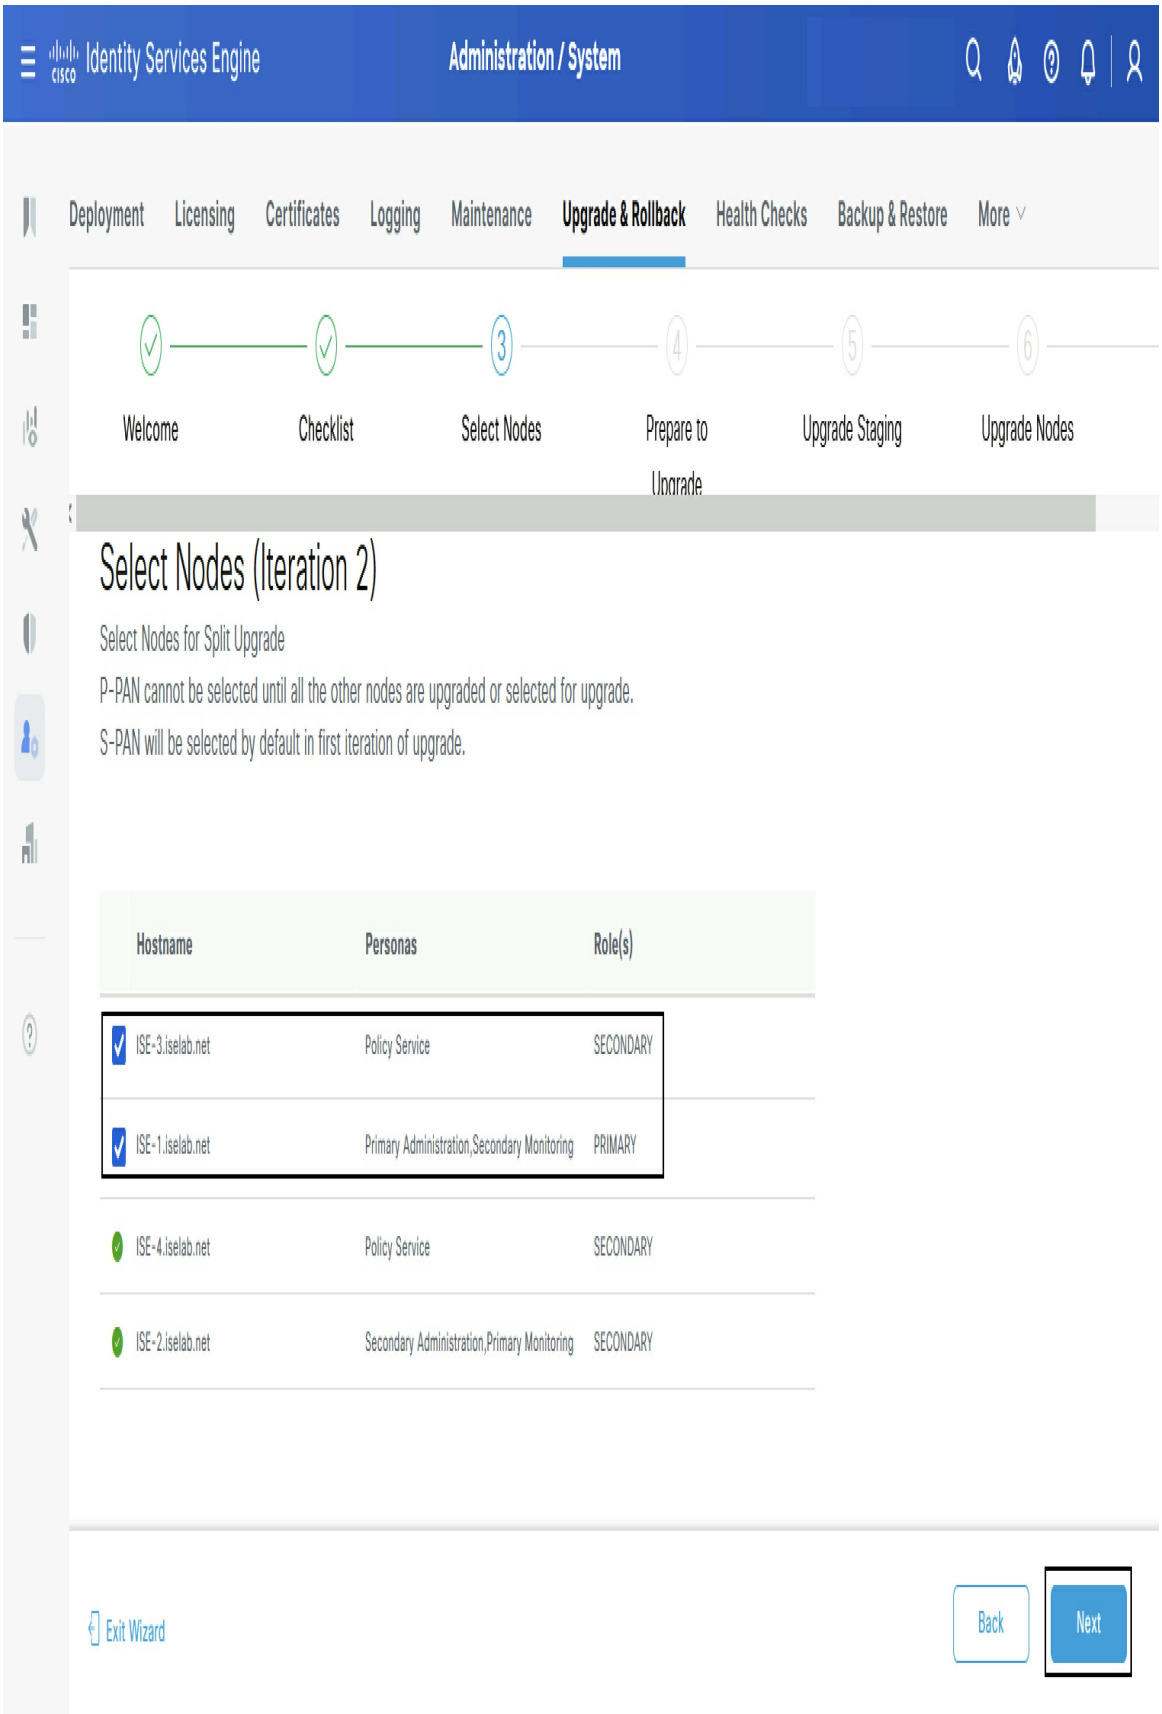

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
Upgrade Staging
Upgrade Nodes
Select Nodes (Iteration 2)
Select Nodes for Split Upgrade
P-PAN cannot be selected until all the other nodes are upgraded or selected for upgrade.
S-PAN will be selected by default in first iteration of upgrade.
Hostname
Personas
Role(s)
ISE-3.iselab.net
Policy Service
SECONDARY
ISE-1.iselab.net
Primary Administration,Secondary Monitoring
PRIMARY
ISE-4.iselab.net
Policy Service
SECONDARY
ISE-2.iselab.net
Secondary Administration,Primary Monitoring
SECONDARY
Exit Wizard
Back
Next

Figure 21-20 *ISE Split Upgrade Iteration 2 Node Selection*

Select the repository and the upgrade bundle, as shown in Figure 21-21. This step is necessary because different ISE nodes may reside in separate data centers, sometimes across different continents. By selecting the repository for each iteration, you can use the repository closest to your ISE nodes. Next, click **Start Preparation**.

Technet24

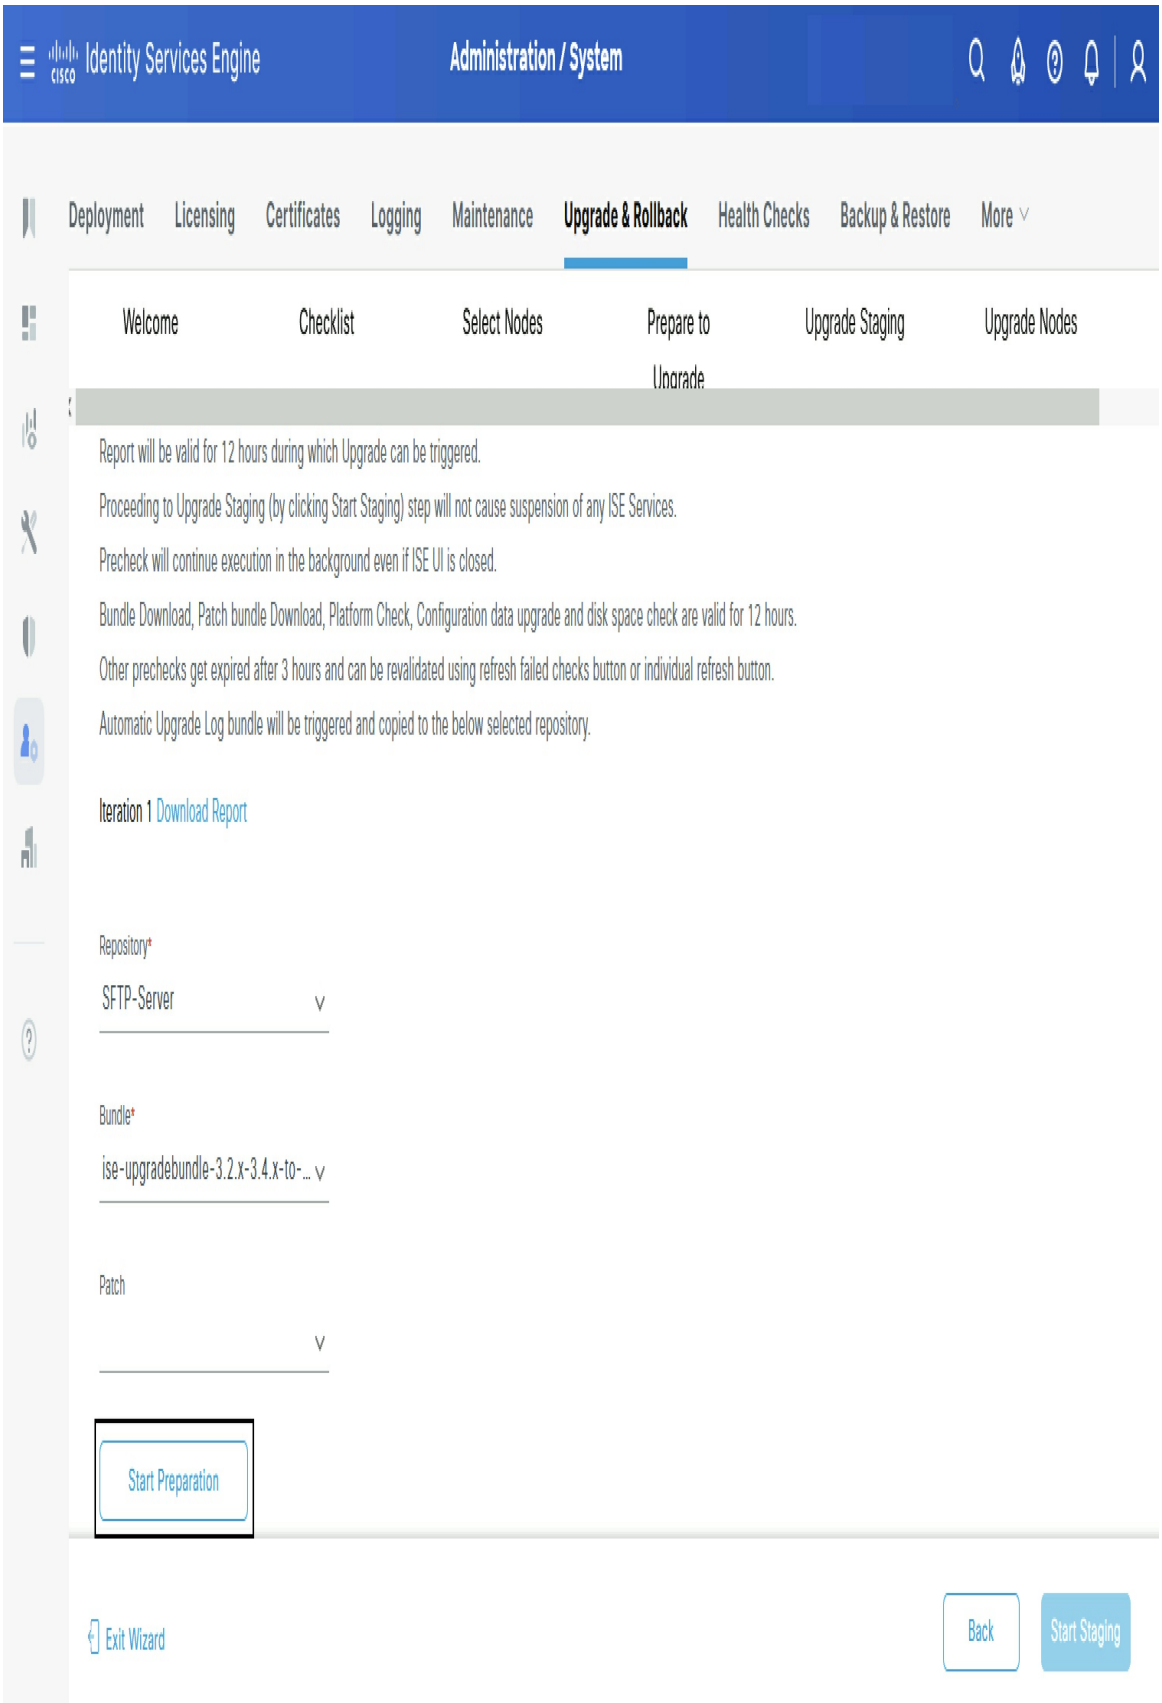

Identity Services Engine
Administration / System
Deployment
Licensing
Certificates
Logging
Maintenance
Upgrade & Rollback
Health Checks
Backup & Restore
More
Welcome
Checklist
Select Nodes
Prepare to Upgrade
Upgrade Staging
Upgrade Nodes
Report will be valid for 12 hours during which Upgrade can be triggered.
Proceeding to Upgrade Staging (by clicking Start Staging) step will not cause suspension of any ISE Services.
Precheck will continue execution in the background even if ISE UI is closed.
Bundle Download, Patch bundle Download, Platform Check, Configuration data upgrade and disk space check are valid for 12 hours.
Other prechecks get expired after 3 hours and can be revalidated using refresh failed checks button or individual refresh button.
Automatic Upgrade Log bundle will be triggered and copied to the below selected repository.
Iteration 1 Download Report
Repository*
SFTP-Server
Bundle*
ise-upgradebundle-3.2.x-3.4.x-to-...
Patch
Start Preparation
Exit Wizard
Back
Start Staging

Figure 21-21 *ISE Split Upgrade Iteration 2 Repository and Bundle Selection*

You can optionally select the destination patch during the upgrade to save the reboot cycle and have the node come up on the selected patch right after the upgrade.

Similar to the preparation for first iteration, health checks are performed on ISE-1 and ISE-3. Upgrade bundles are transferred to the nodes as well. When health checks are completed, click **Start Staging**. After staging is complete, as shown in Figure 21-22, click **Next**.

Technet24

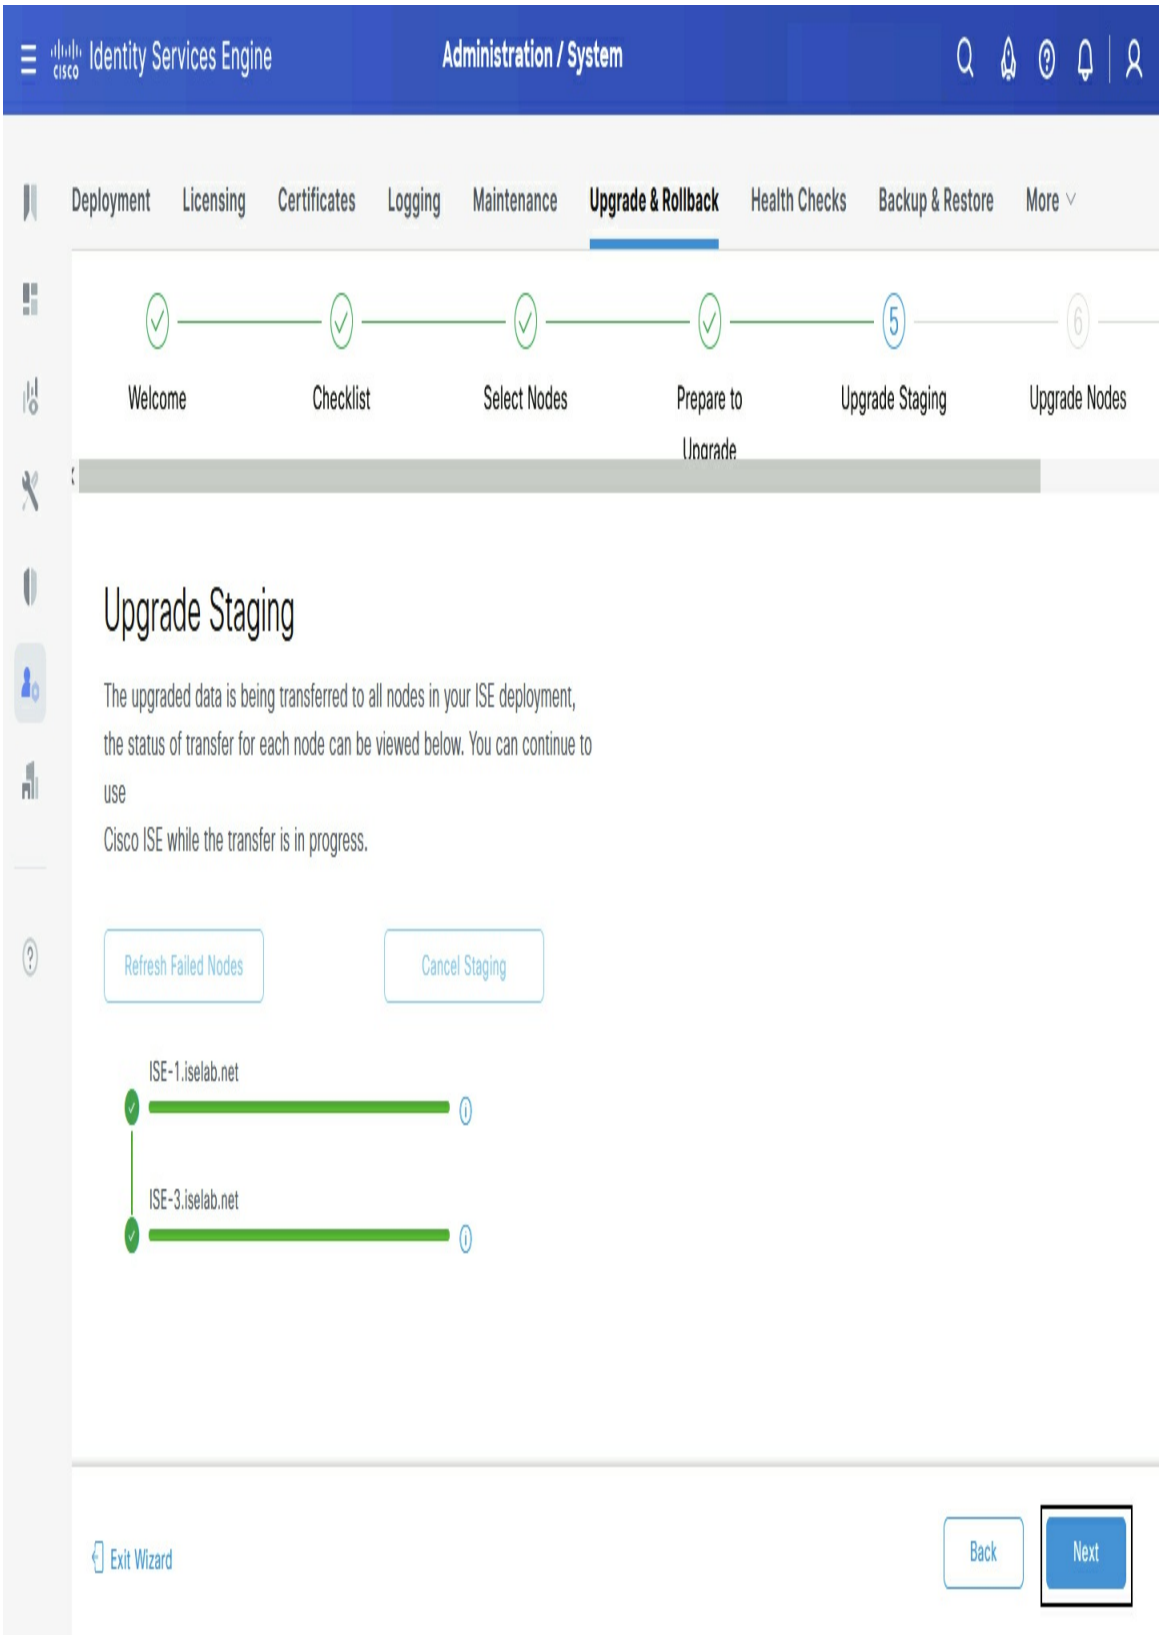


Figure 21-22 *ISE Split Upgrade Iteration 2 Staging*

Click **Start** to initiate the upgrade of the nodes in second iteration, as shown in Figure 21-23. At this stage, the nodes restart. Ensure that authentications have already been diverted to the upgraded PSN (ISE-4) from the first iteration.

Technet24

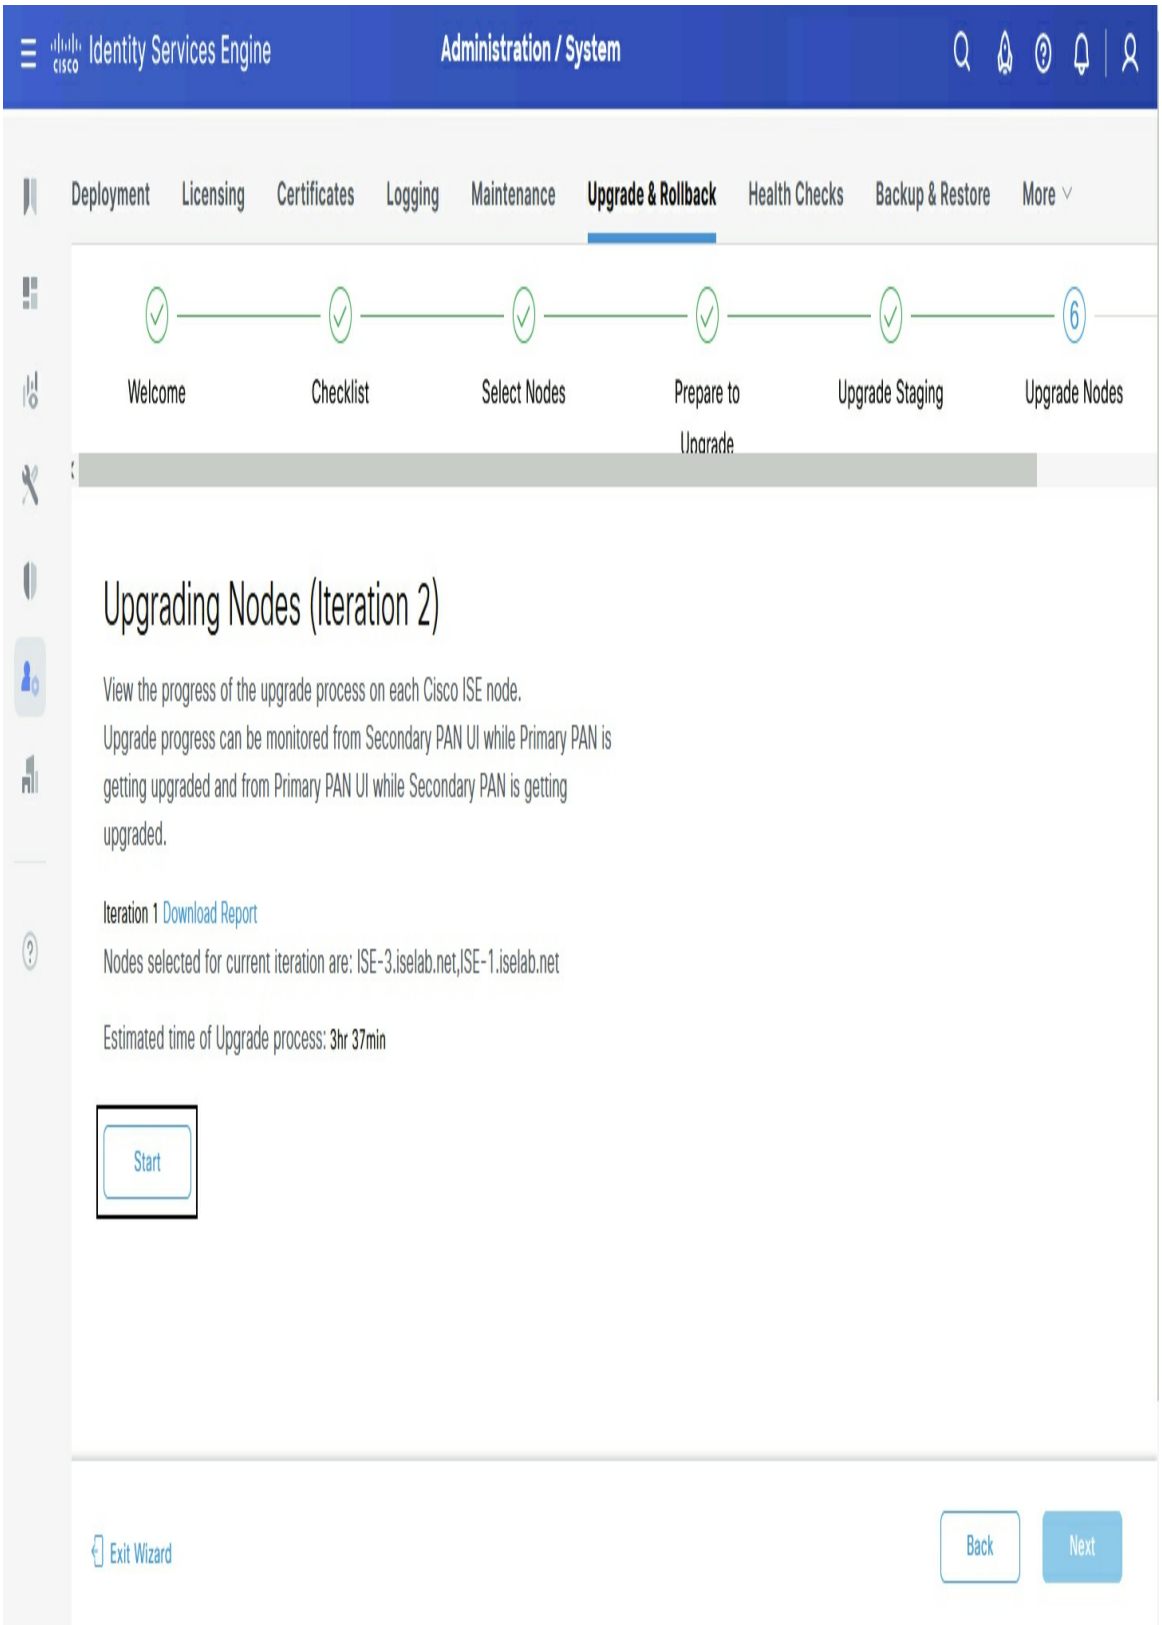


Figure 21-23 *ISE Split Upgrade Iteration 2 Start*

You can monitor the upgrade process for a second iteration from the PPAN of the original deployment, as shown in Figure 21-24, until it is time to upgrade the PPAN itself.

Technet24

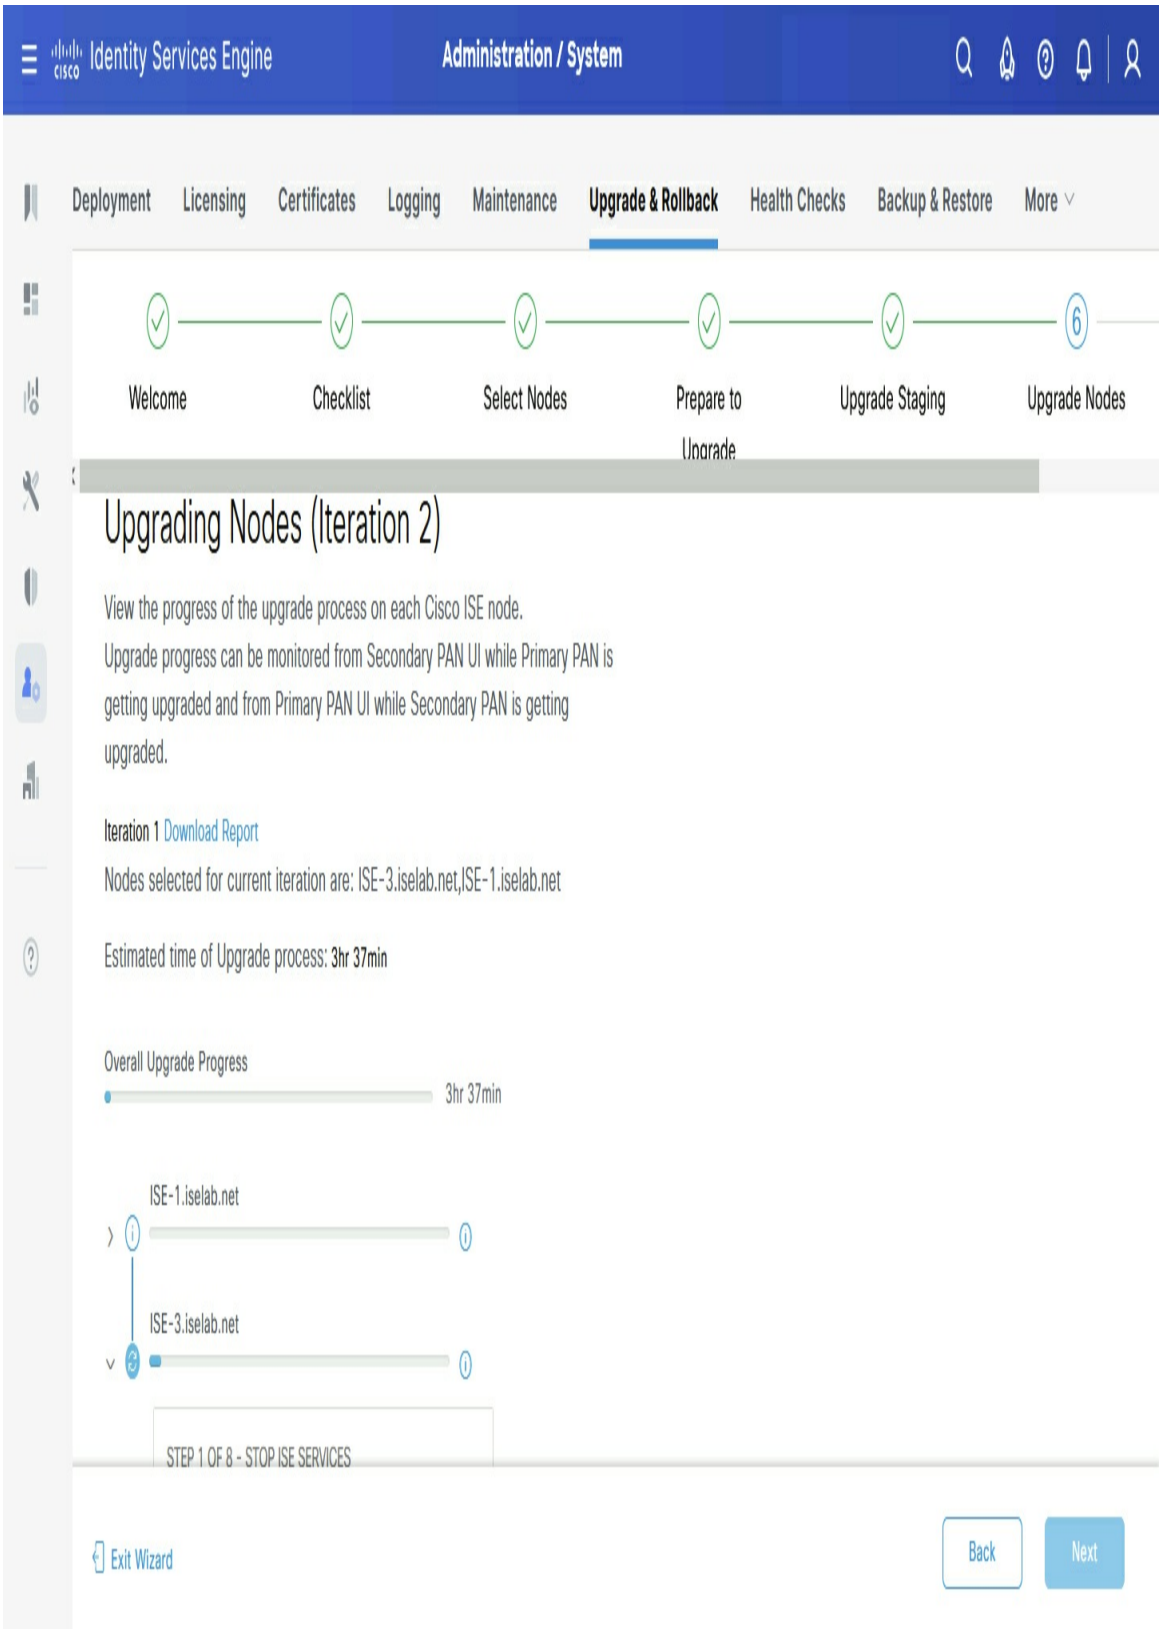


Figure 21-24 *ISE Split Upgrade Iteration 2 Progress from Original*

PPAN (ISE-1)

When the PPAN (ISE-1) upgrade begins, it becomes unavailable. However, you can monitor the upgrade status from the former SPAN (ISE-2), as shown in Figure 21-25.

Technet24

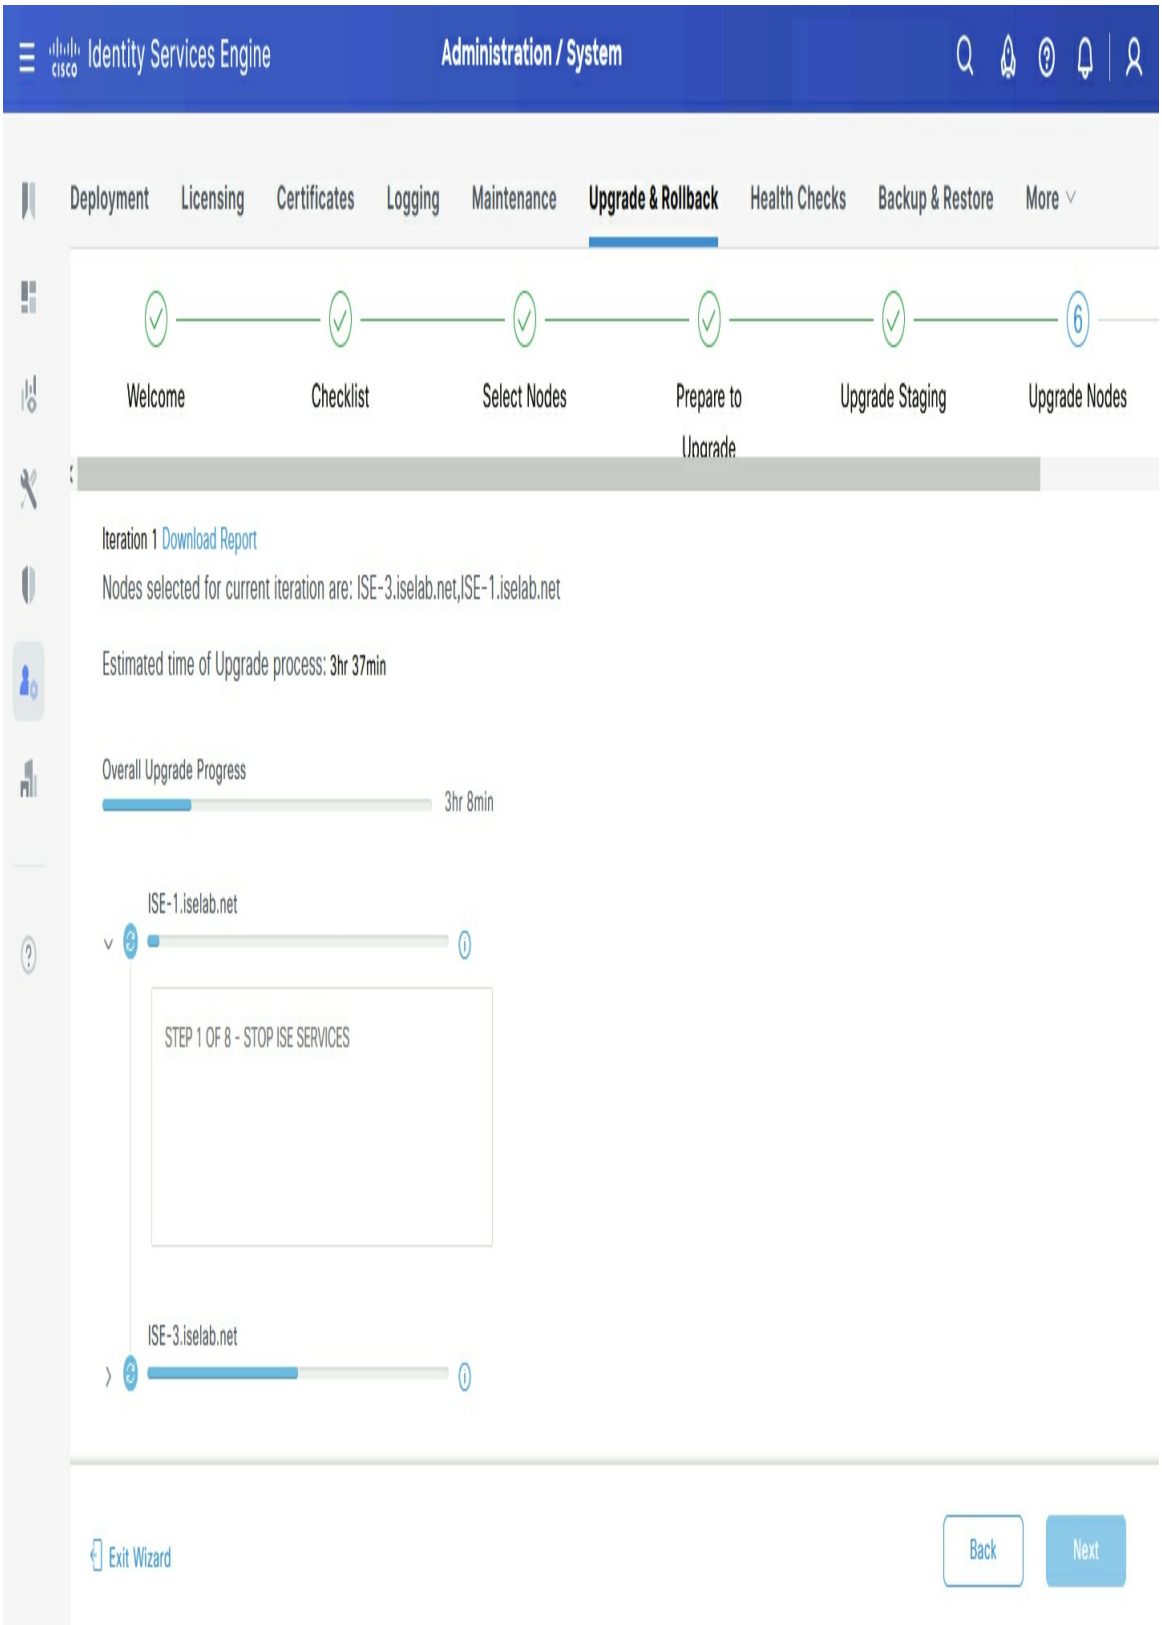


Figure 21-25 *ISE Split Upgrade Iteration 2 Progress from Original*

SPAN (ISE-2)

When the PSN and former PPAN are upgraded, the new PPAN displays this status on the **Upgrade & Rollback** screen, as shown in Figure 21-26. Click **Next**.

Technet24

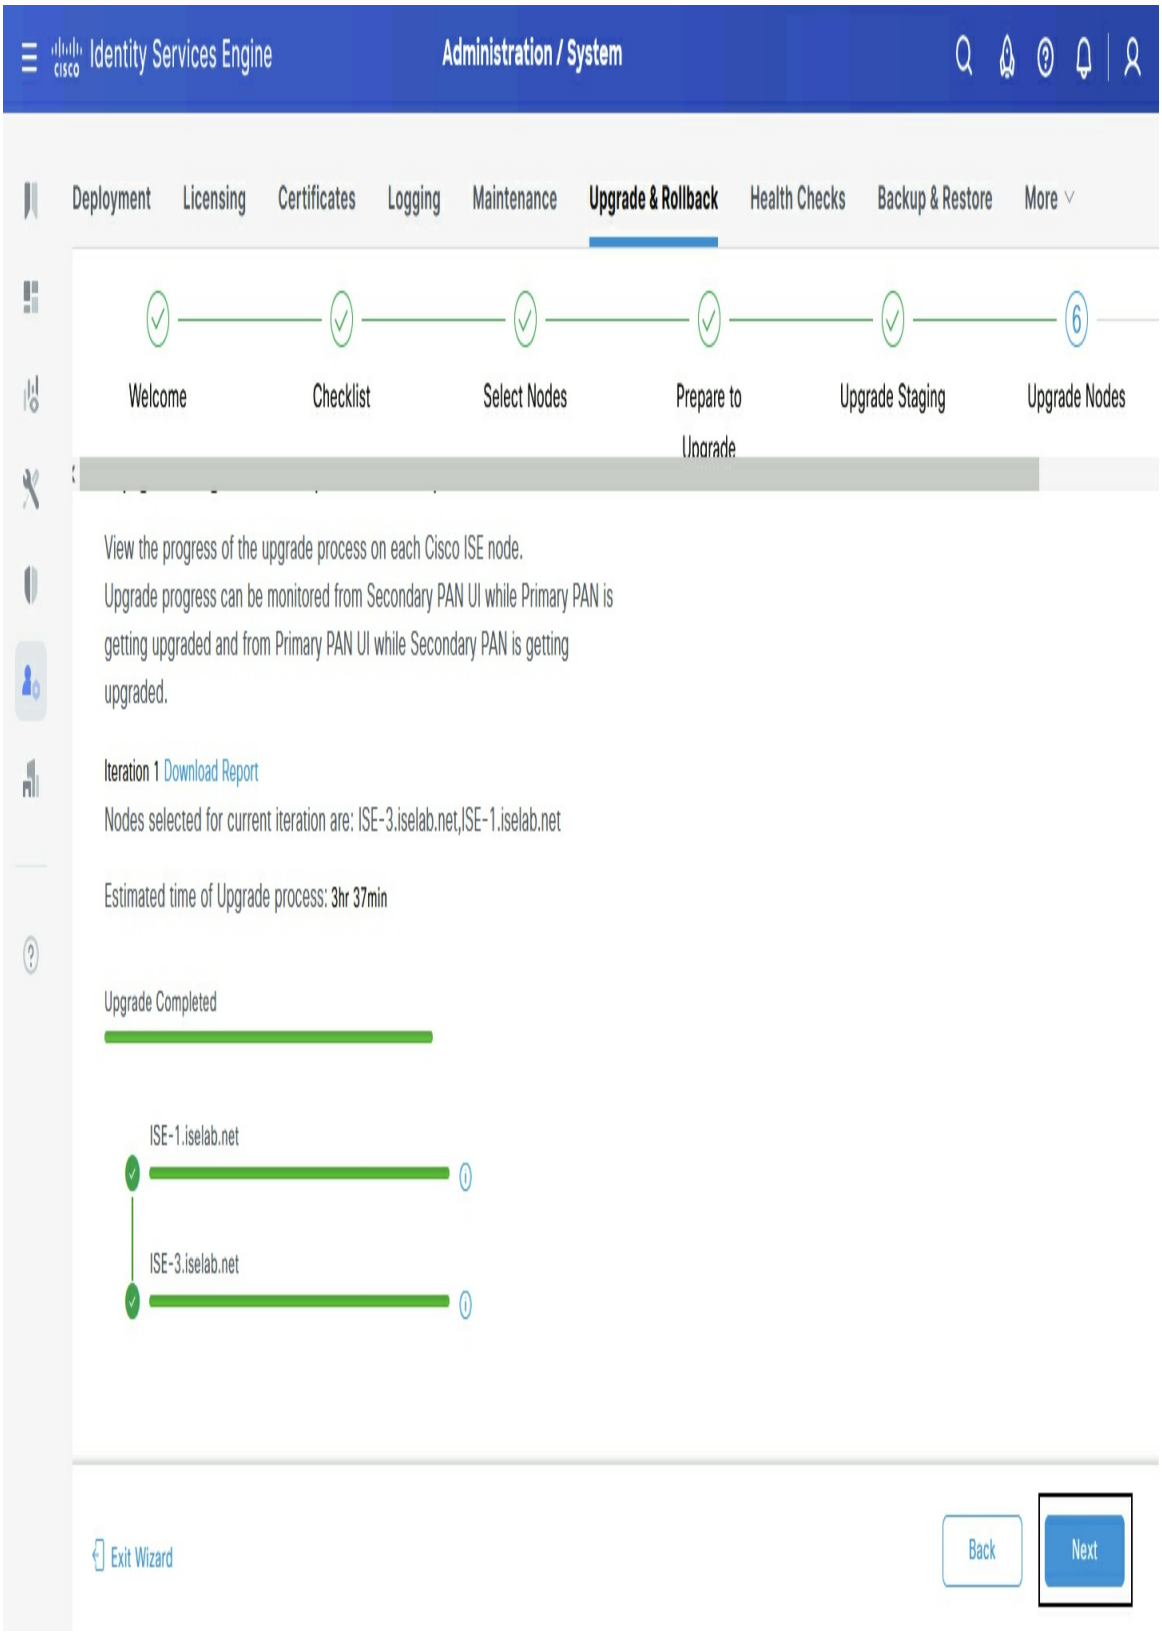


Figure 21-26 *ISE Split Upgrade Iteration 2 Completion from*

Original SPAN (ISE-2)

Reports for the second iteration are available for download. You can download them for your reference. To complete the upgrade, click **Finish**, as shown in Figure 21-27.

Technet24

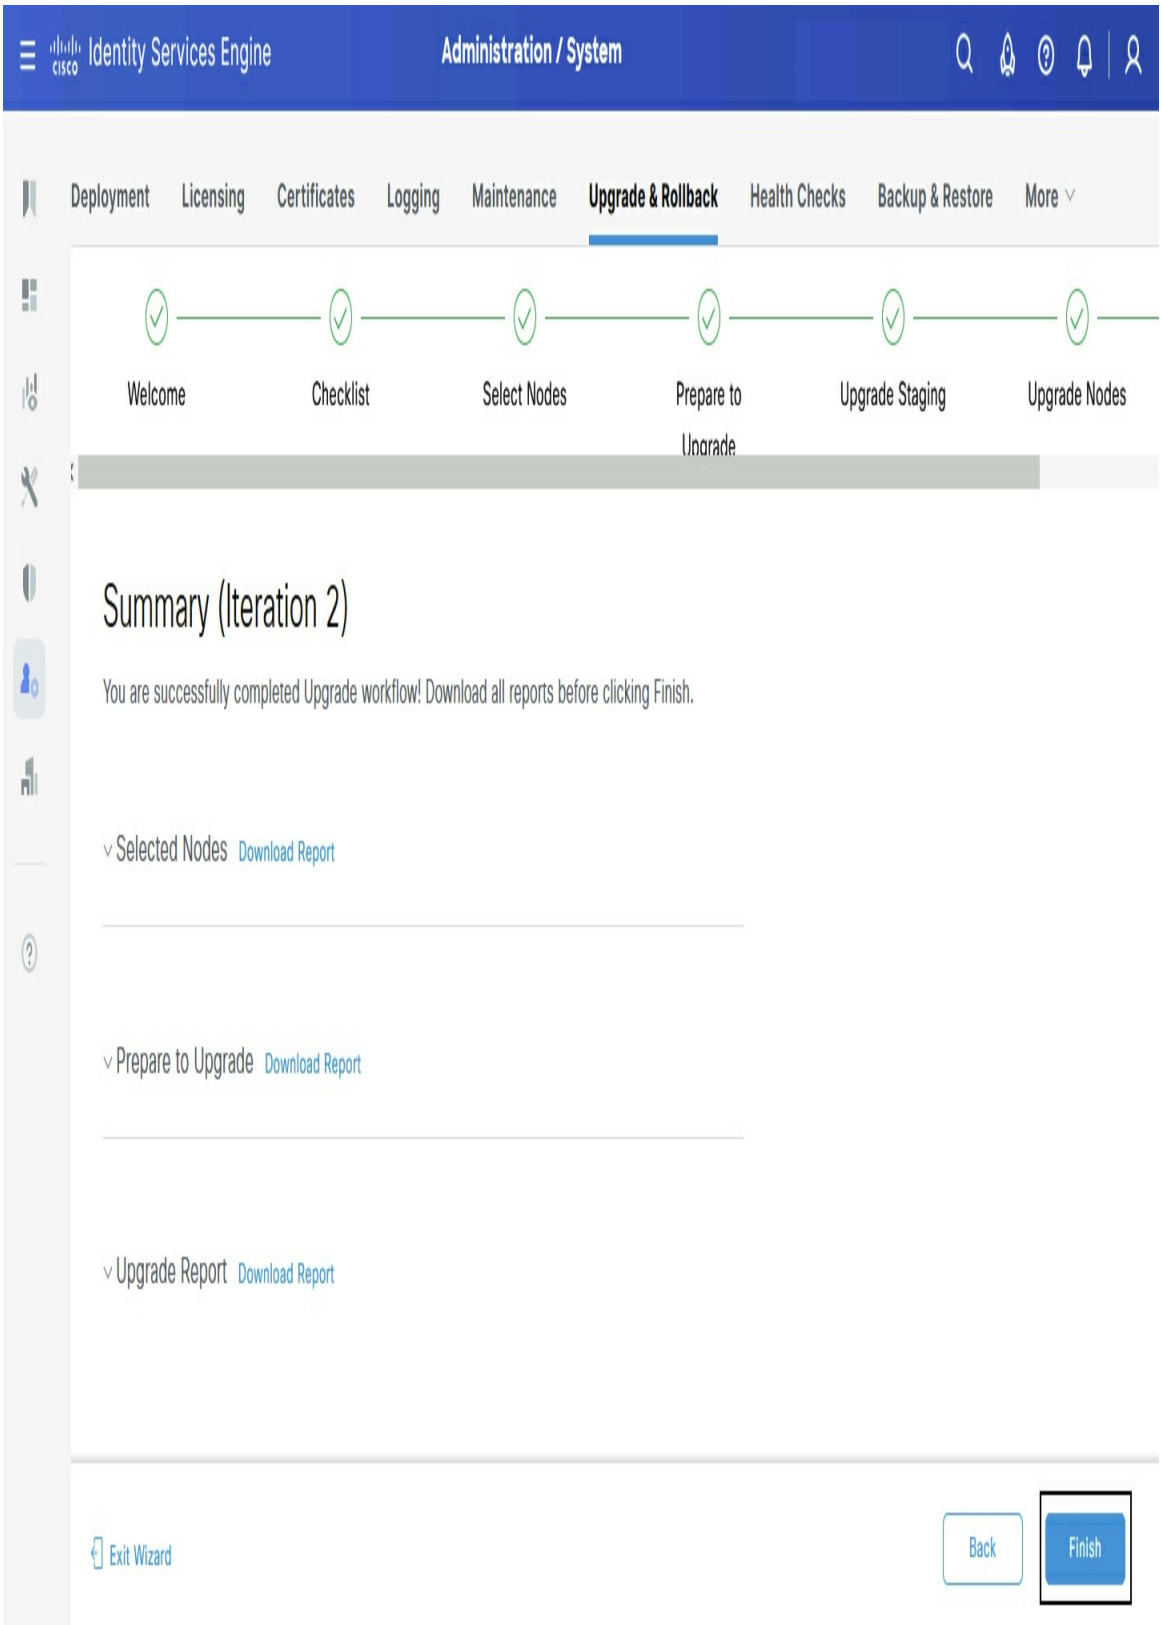


Figure 21-27 *ISE Split Upgrade Iteration 2 Summary from Original*

SPAN (ISE-2)

Next, navigate to **Administration > System > Deployment** to confirm that all ISE nodes are reassembled into a single cluster, as shown in Figure 21-28.

Technet24

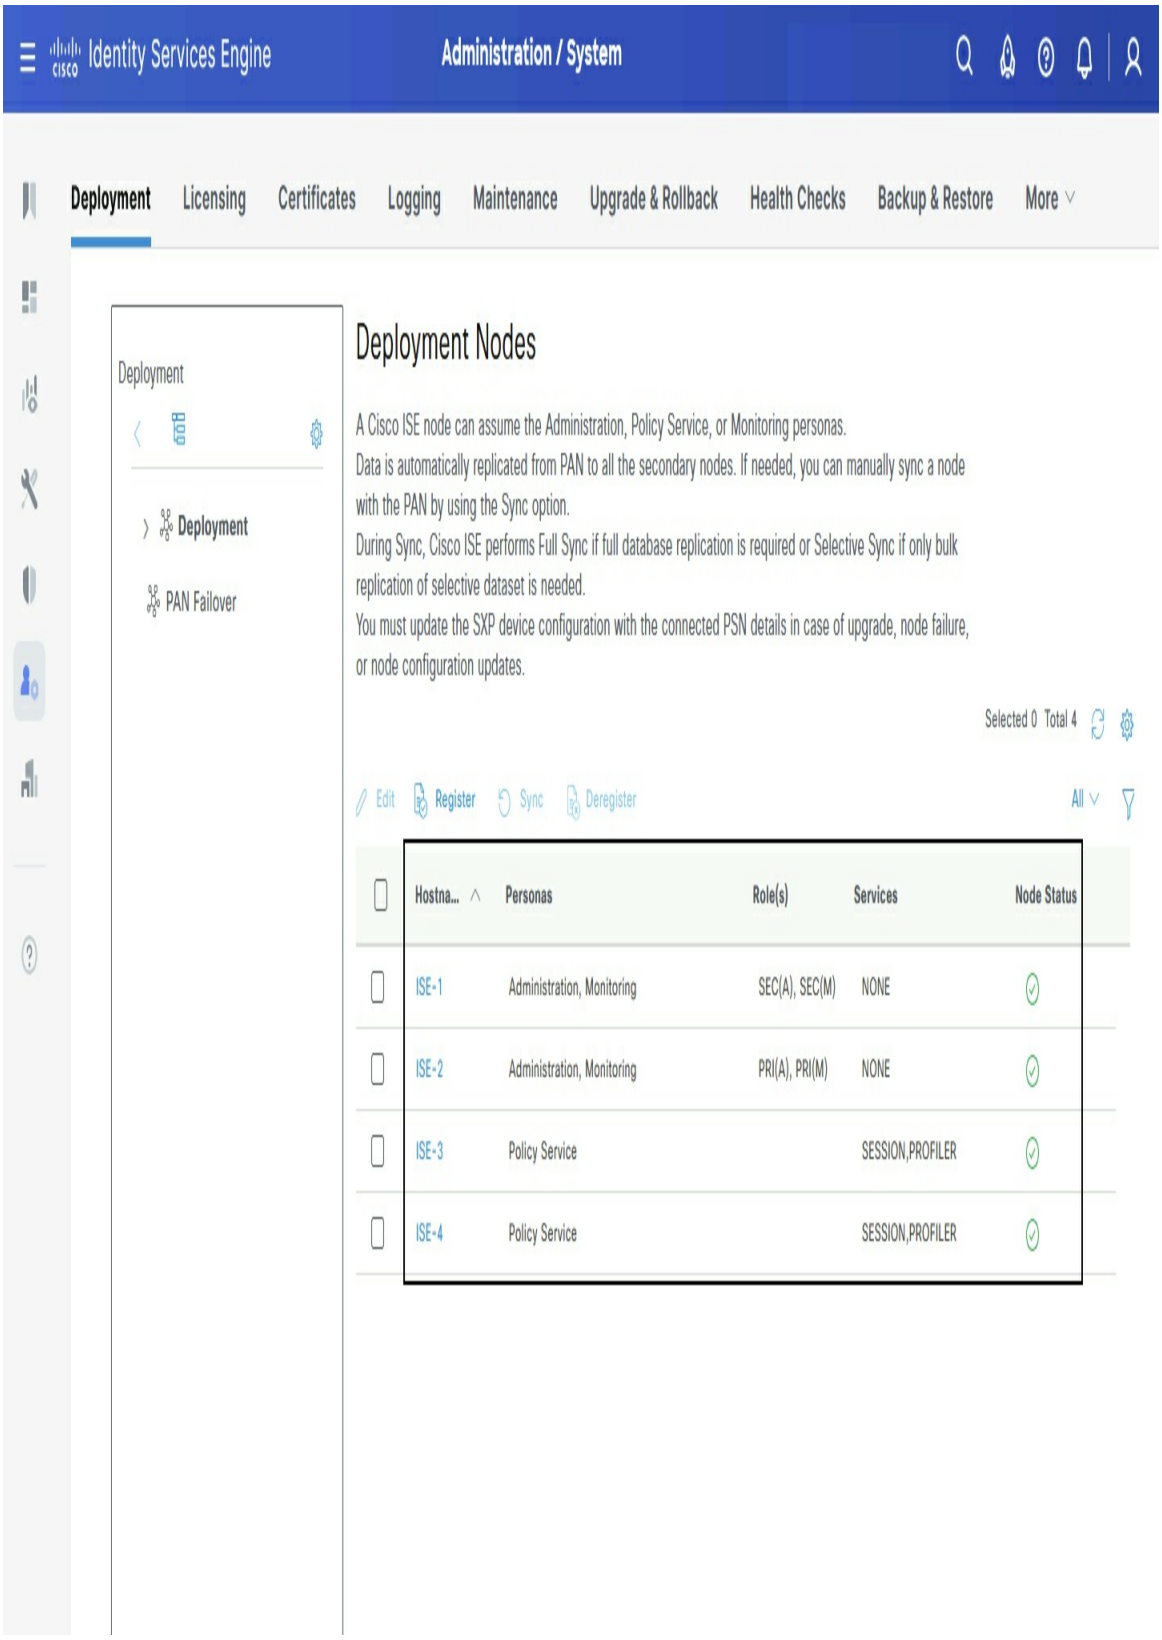


Figure 21-28 *ISE Medium Deployment Roles Allocation After the*

Upgrade

As you may observe, the former SPAN (ISE-2) assumes the PPAN role in the new deployment. After completing the upgrade, you can reassign the PPAN role to the original PPAN (ISE-1) by logging in to the former PPAN (ISE-1) and selecting **Promote to Primary** from the **Deployment** page.

Patch Installation

Keeping any IT system up to date is essential, and the same principle applies to ISE. ISE patches are software updates that provide bug fixes and security updates, and they occasionally introduce new features. Always consult the latest Release Notes to determine whether new features are available, understand their details, and review the list of fixes.

ISE patches are cumulative. For example, patch 3 for ISE 3.5 includes all fixes from patch 1 and patch 2, as well as additional fixes introduced in patch 3. This model eliminates the need to install all patches individually; installing the latest patch is sufficient.

You can install patches via the GUI or CLI. To install patches using the CLI, issue the following command:

```
ISE-1/admin#patch install ise-patchbundle-3.5.0.527-Patch3-26040703.
SFTP-Server
```

> **Note**
>
> If you choose to install the patch from the CLI, it is recommended to start installation from PPAN node.

You can uninstall patches using this command, where 3 represents the

Technet24

patch number you intend to uninstall:

```
ISE-1/admin#patch remove ise 3
```

> **Note**
>
> CLI operations for both patch installation and removal are performed only on the local ISE node, even if it is part of a deployment. If you install a patch from the CLI, it applies only to that node and does not propagate across the deployment. The same applies to the uninstallation process.

To install the patch from the GUI, navigate to **Administration > System > Upgrade & Rollback**, as shown in Figure 21-29. This approach offers the same options as a major upgrade.

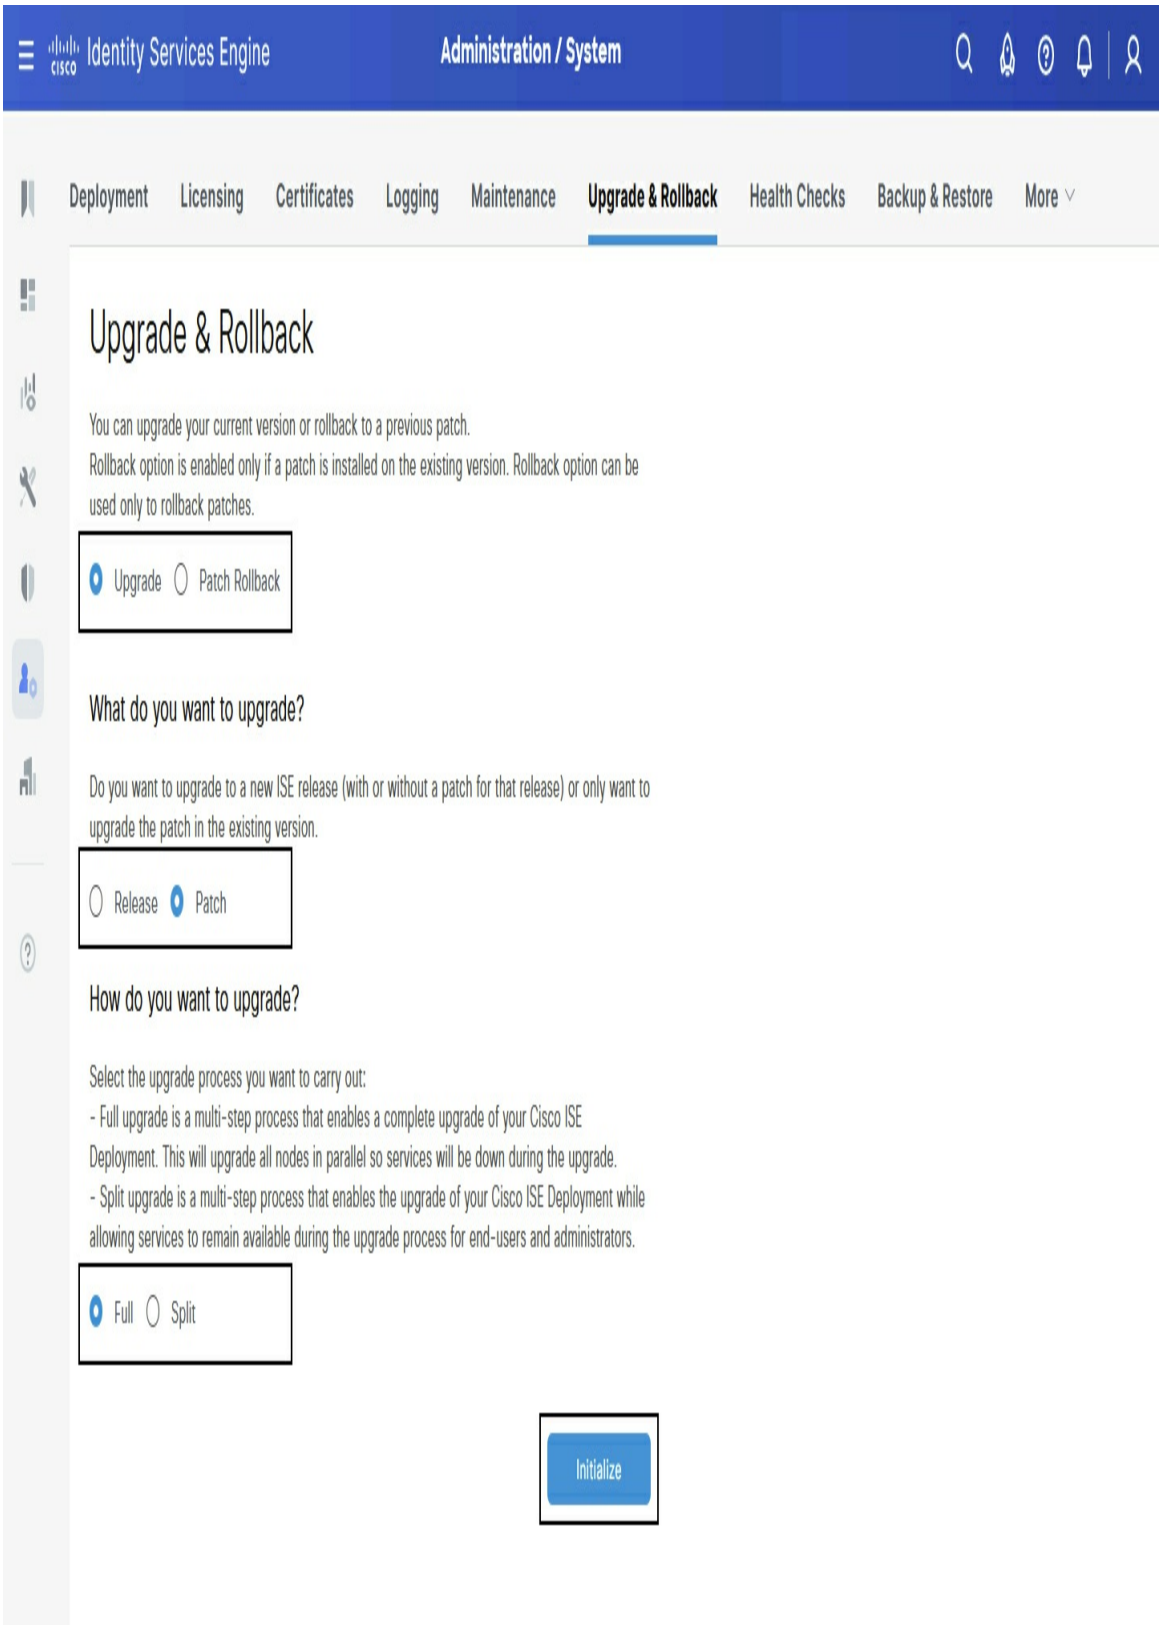


Figure 21-29 *ISE Upgrade & Rollback GUI Menu for Full Patch*

Technet24

Installation

- **Full installation:** The patch installs first on the PPAN node. After completion, simultaneous installation occurs on all other nodes in the deployment.
- **Split installation:** You determine which nodes receive the patch first based on iterations.

Because we already explored split upgrade capabilities with a major upgrade, let's install the patch on the just-upgraded ISE 3.5 4 nodes deployment using the full installation method. To start, select **Initialize**; then click **Let's Do It** on the initial page of the **Patch Upgrade Wizard**.

There are two methods to upload the patch. You can download it from the repository, as shown in Figure 21-30, or use the **Upload Now** option to manually upload the patch from your computer. Click **Start Preparation** to initiate health checks and download the patch to all nodes.

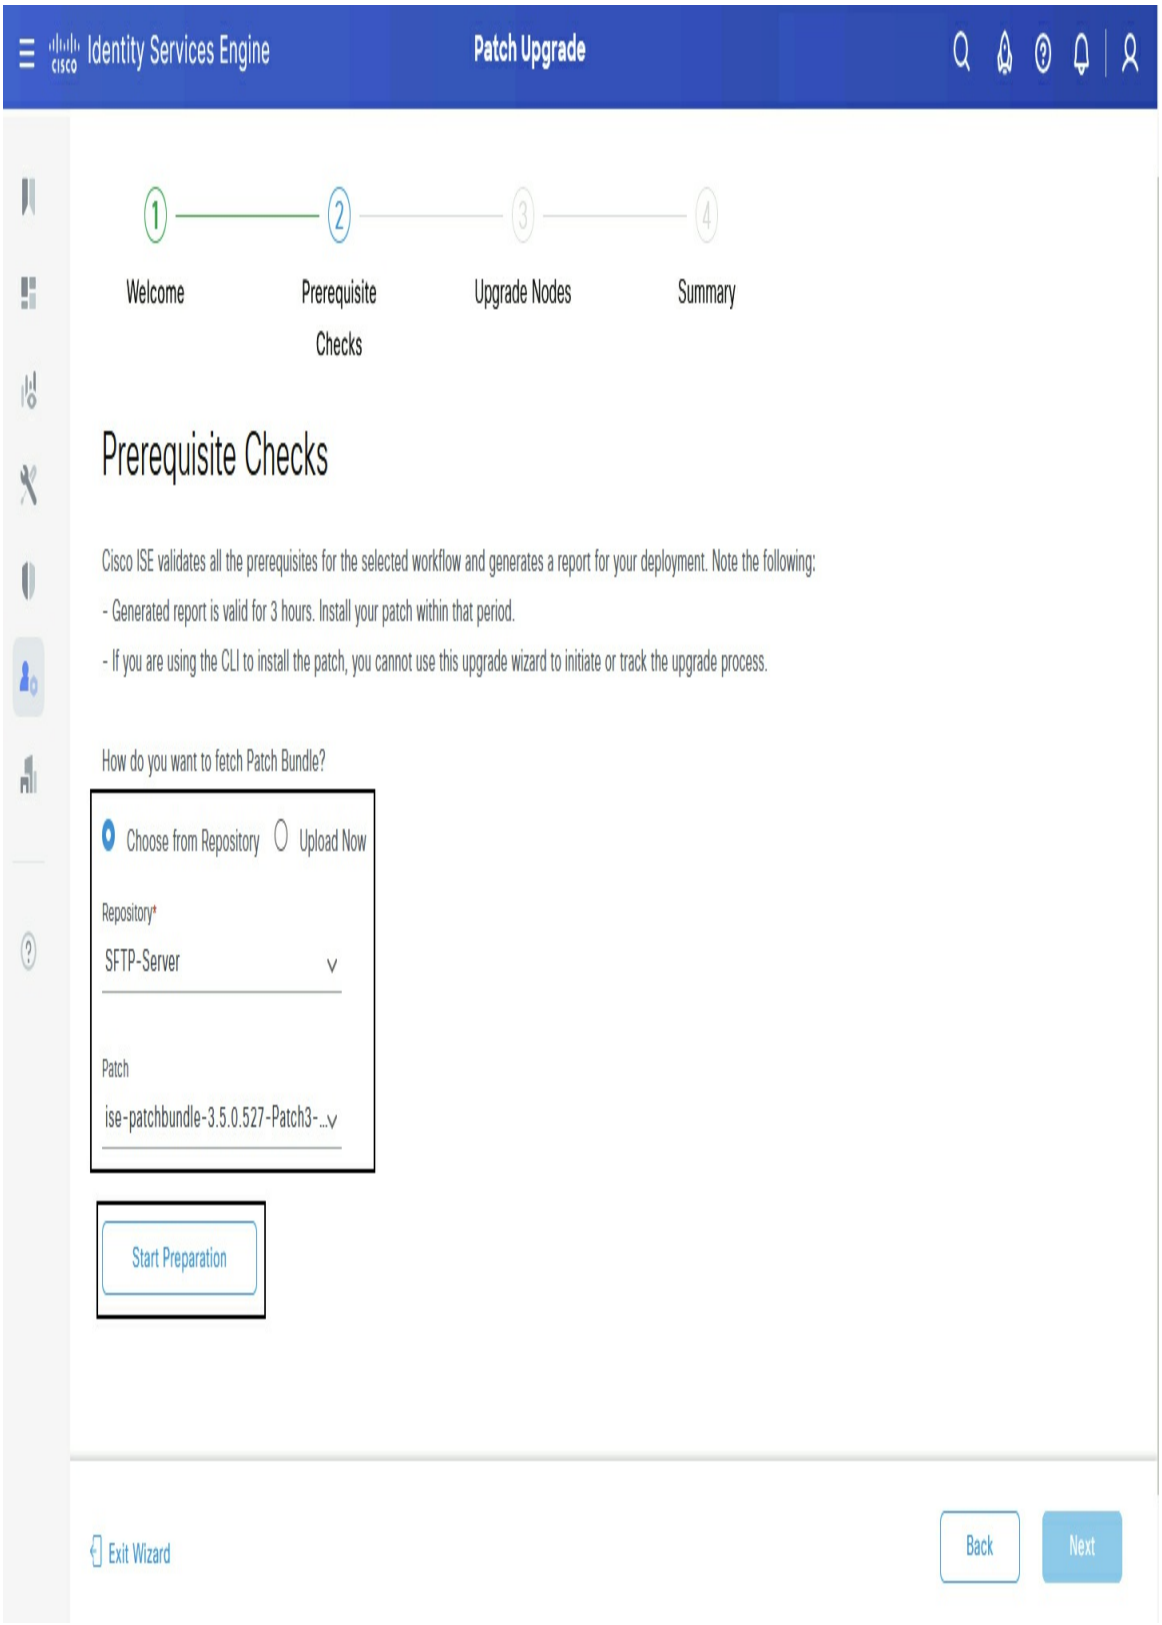


Figure 21-30 *ISE Full Patch Installation Repository and Bundle*

Technet24

Selection

After the preparation is complete and you remediate any prerequisite failures, click **Next**, as shown in Figure 21-31.



Figure 21-31 *ISE Full Patch Installation Health Checks*

Technet24

Select **Start** on the next screen, as shown in Figure 21-32, and click **Yes** to proceed with the patch installation when prompted with a warning about downtime.

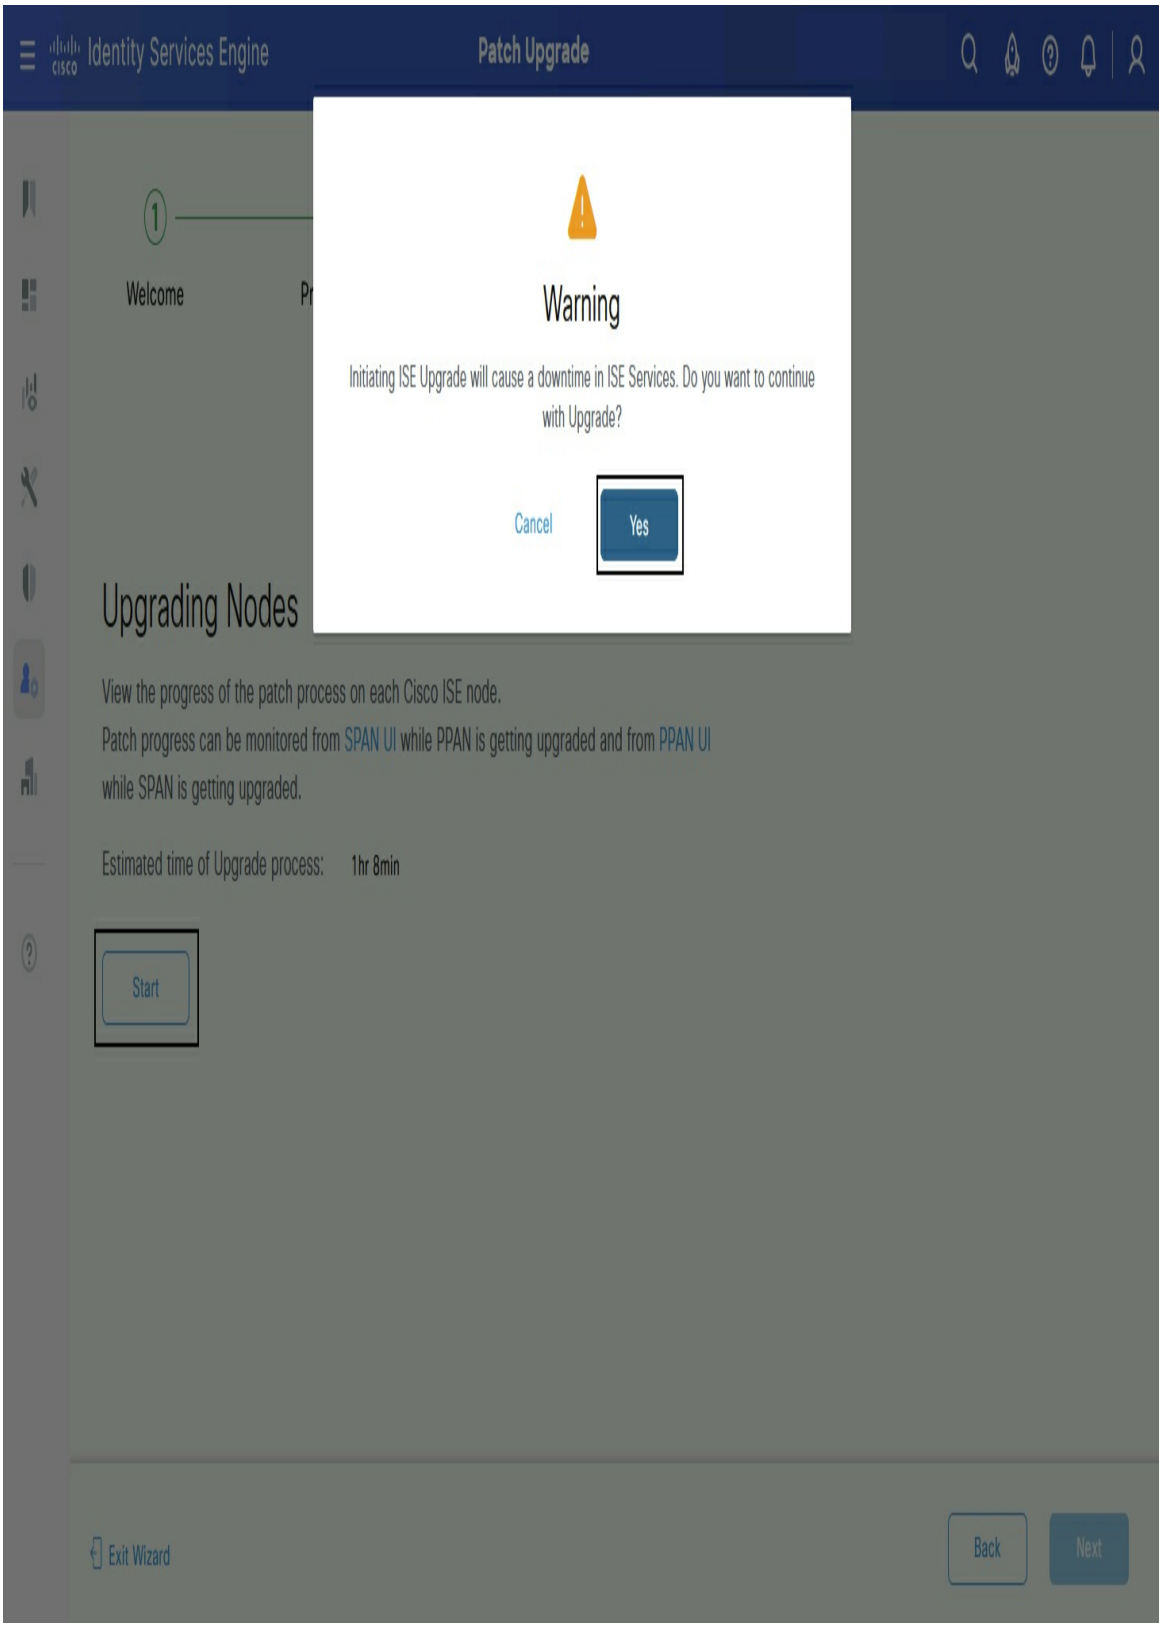


Figure 21-32 *ISE Full Patch Installation Start*

Technet24

Because the full patch installation method begins with PPAN, you can log in to SPAN and monitor the patch installation progress there while the patch installs on the PPAN, as shown in Figure 21-33.

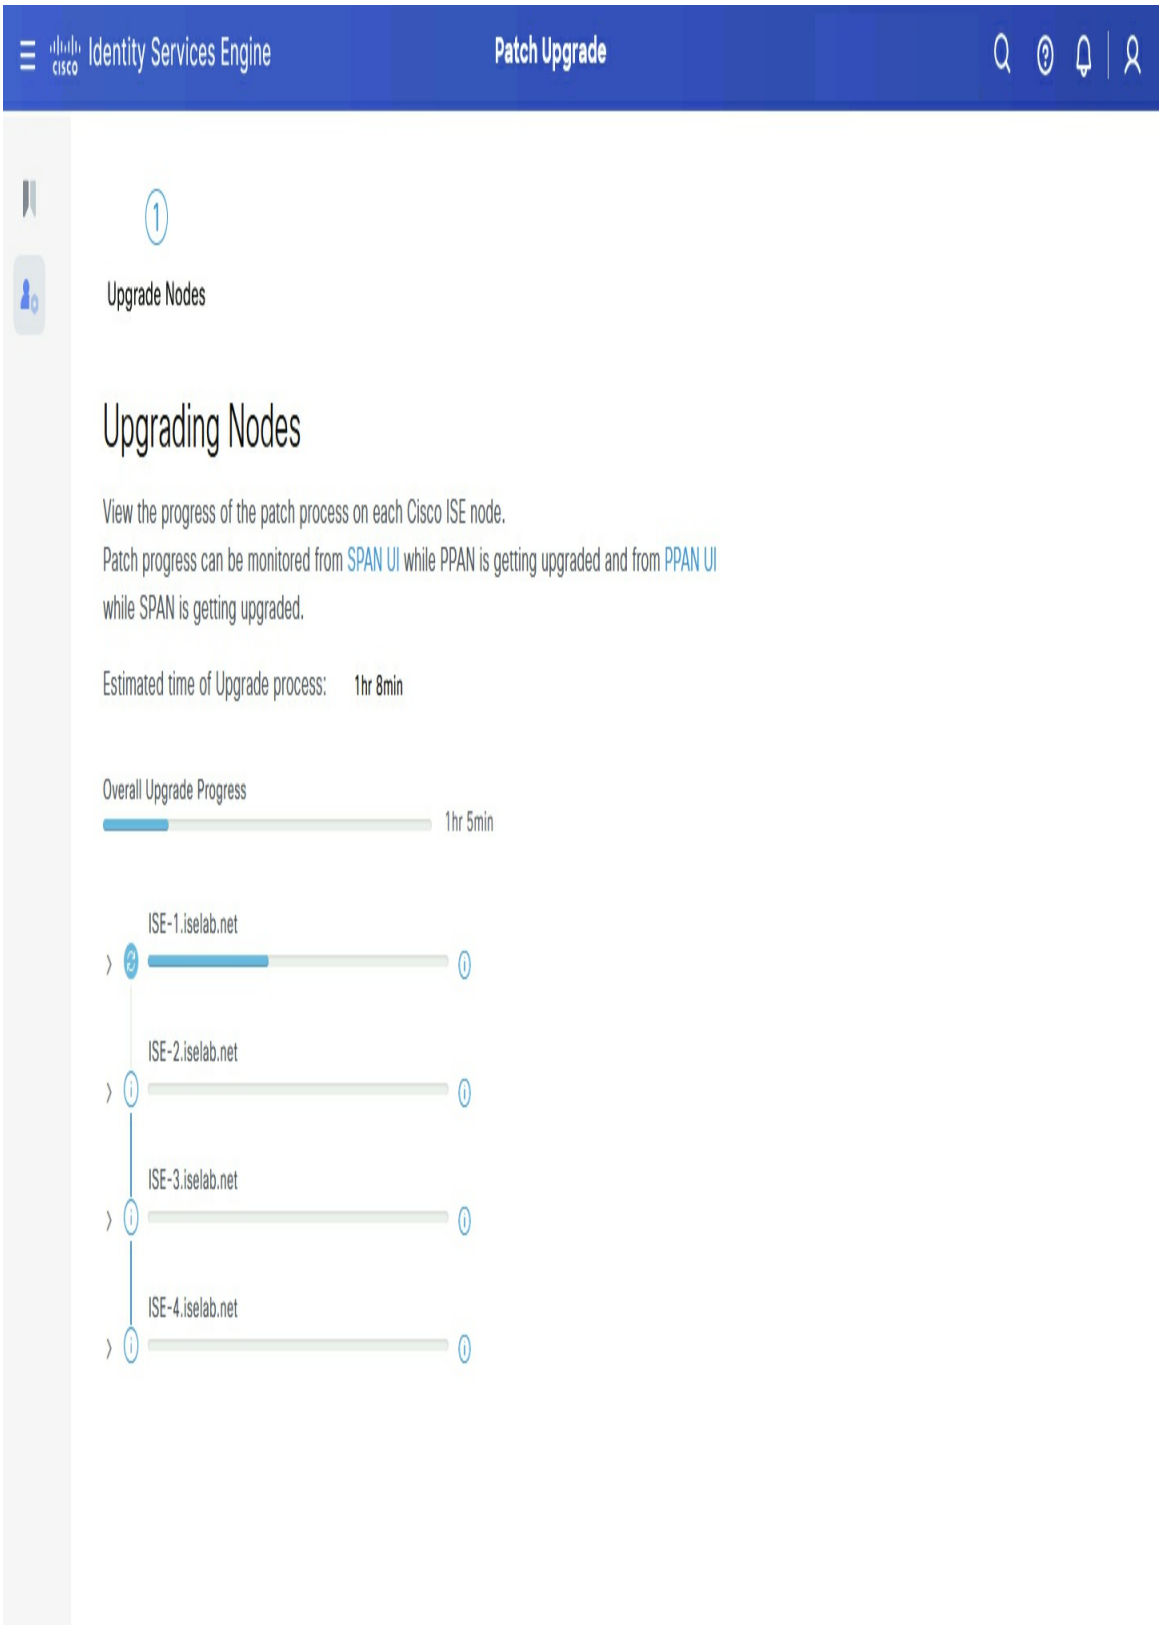
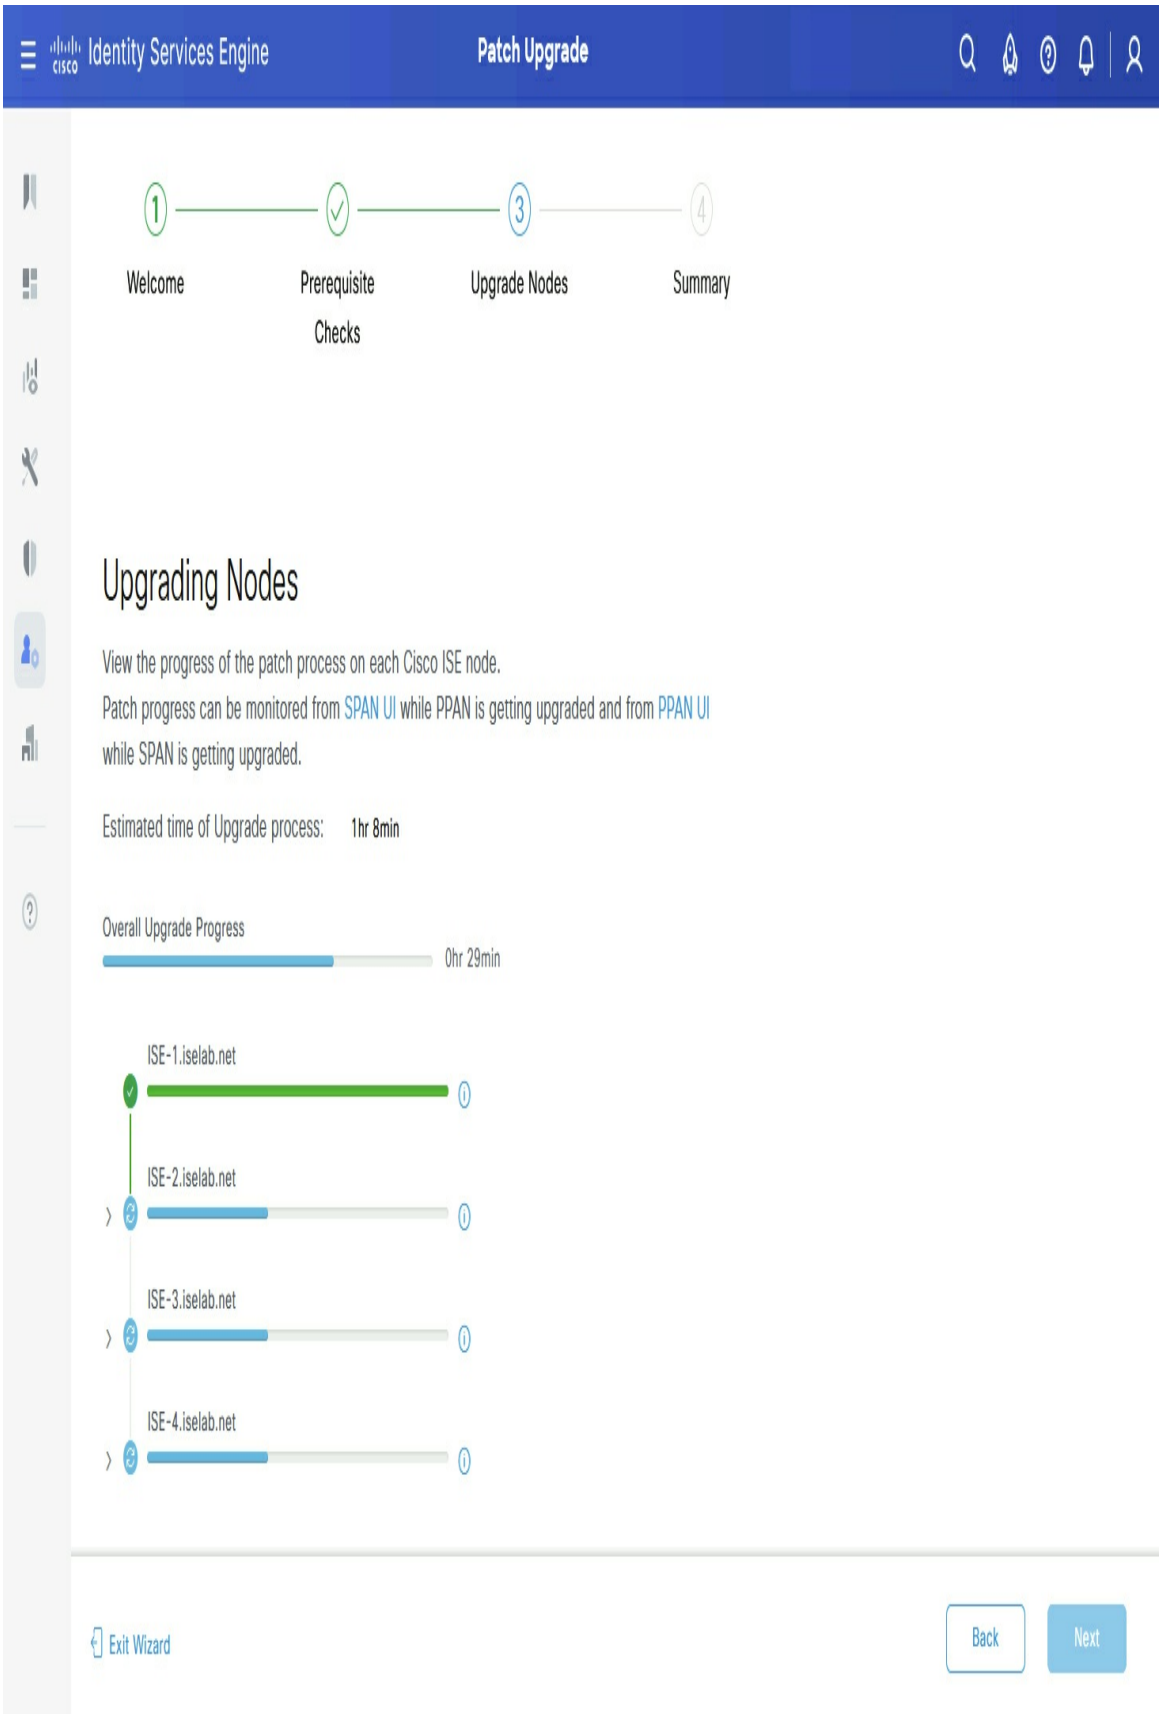


Figure 21-33 *ISE Full Patch Installation Progress from SPAN*

Technet24

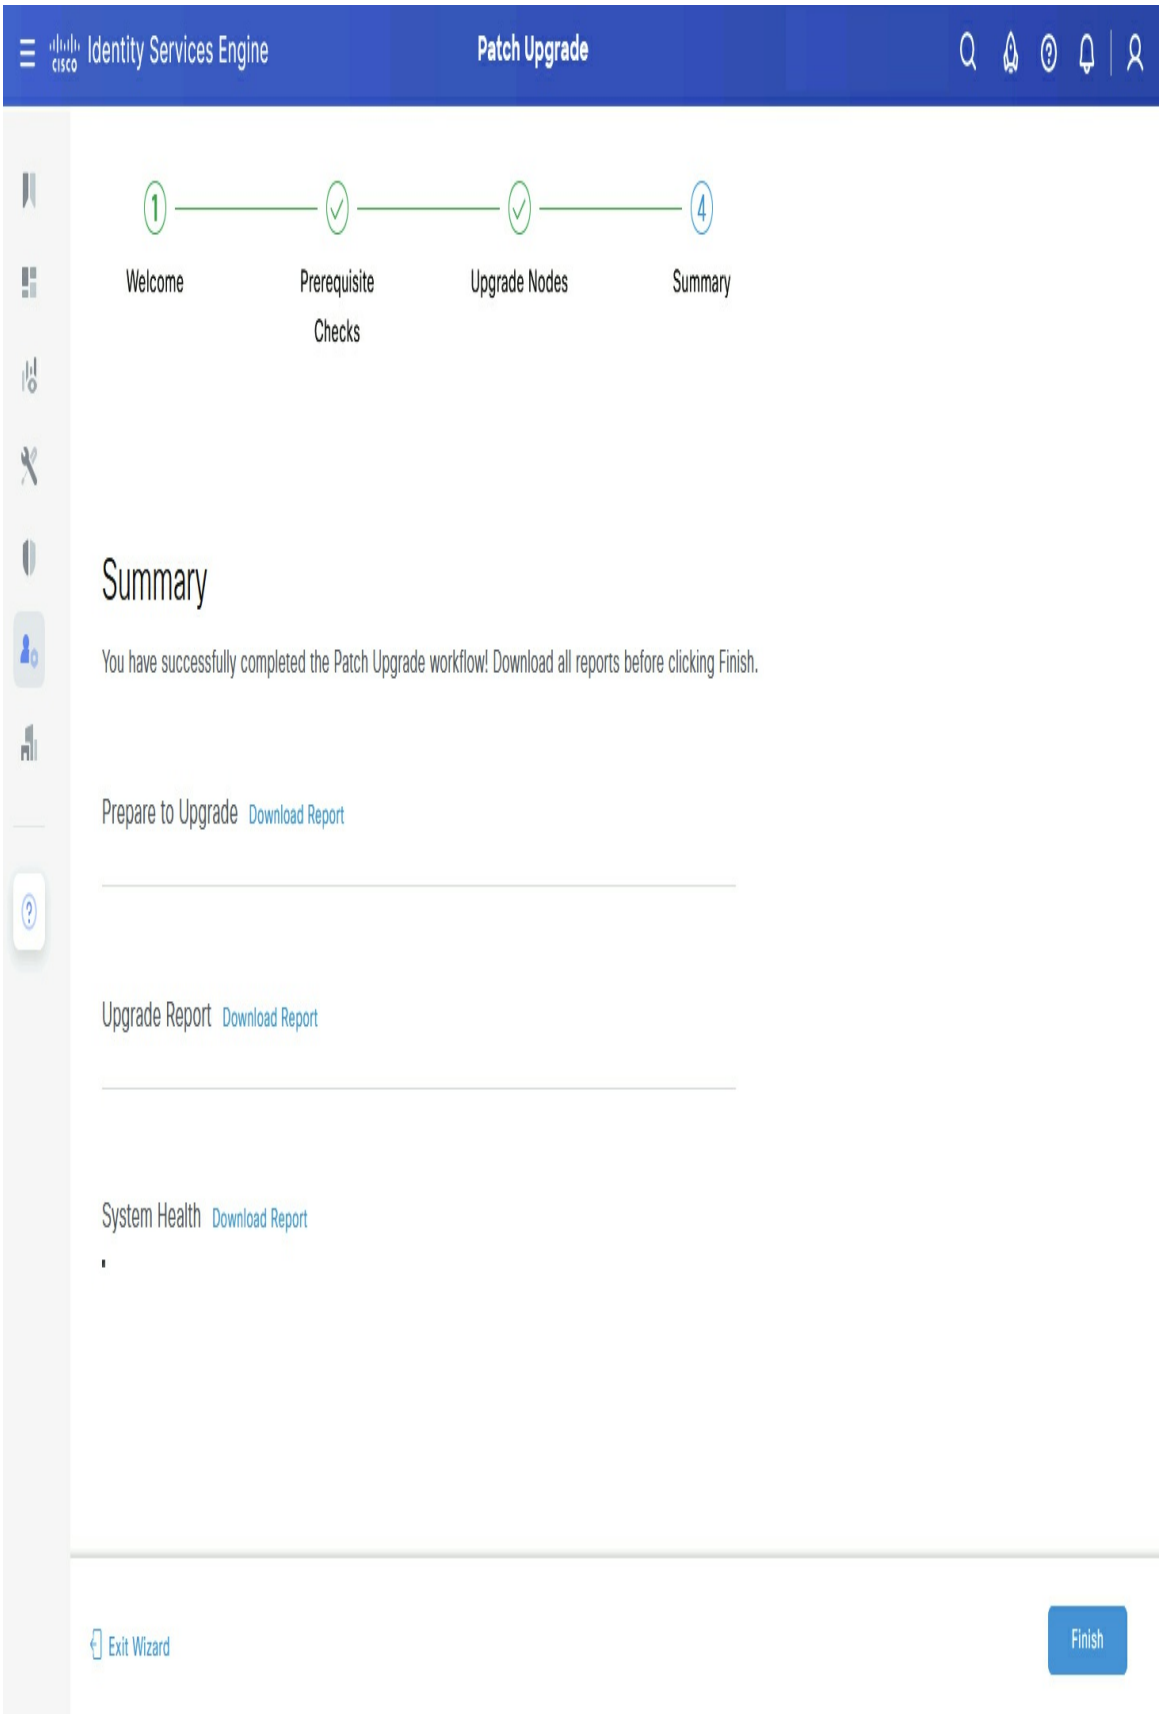
After the patch installs on PPAN, ISE initiates simultaneous installation of the patch on all other nodes. You can log in to PPAN to confirm this and monitor the progress, as shown in Figure 21-34.


Identity Services Engine
Patch Upgrade
1
Welcome
Prerequisite Checks
3
Upgrade Nodes
4
Summary
Upgrading Nodes
View the progress of the patch process on each Cisco ISE node.
Patch progress can be monitored from SPAN UI while PPAN is getting upgraded and from PPAN UI while SPAN is getting upgraded.
Estimated time of Upgrade process: 1hr 8min
Overall Upgrade Progress
0hr 29min
ISE-1.iselab.net
ISE-2.iselab.net
ISE-3.iselab.net
ISE-4.iselab.net
Exit Wizard
Back
Next

Technet24

Figure 21-34 *ISE Full Patch Installation Progress from PPAN*

After the patch installation finishes on all nodes, click **Next**. ISE initiates another round of health checks after the patches are installed. When this process completes, click **Finish** to complete the patch installation, as shown in Figure 21-35.


Identity Services Engine
Patch Upgrade
1
Welcome
Prerequisite Checks
Upgrade Nodes
4
Summary
Summary
You have successfully completed the Patch Upgrade workflow! Download all reports before clicking Finish.
Prepare to Upgrade Download Report
Upgrade Report Download Report
System Health Download Report
Exit Wizard
Finish

Technet24

Figure 21-35 *ISE Full Patch Installation Summary*

Summary

ISE maintenance remains essential to keep your system up to date. Major upgrades and patches deliver advanced features and improvements. Using the latest releases keeps your ISE deployment secure and prevents potential issues associated with unsupported systems. The robust Upgrade & Rollback system introduced in recent ISE releases provides flexibility during upgrades and patch installations. It also offers alternative options, such as full upgrade or full patch installation, allowing you to reach the target version more quickly when your environment permits.