Avaya Identity Engines Ignition Server Administration

Avaya Identity Engines Ignition Server
Release 8.0

Document Status: **Standard**
Document Number: **NN47280-600**
Document Version: **03.02**
Date: **April 2012**

© 2012 Avaya Inc.
All Rights Reserved.

Notices
While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer
Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer
Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation(s) provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty
Avaya provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product, while under warranty, is available to Avaya customers and other parties through the Avaya Support Web site: http://www.avaya.com/support
Please note that if you acquired the product from an authorized reseller, the warranty is provided to you by said reseller and not by Avaya.

Licenses
THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/ ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER, AND AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE ("AVAYA").

Copyright
Except where expressly stated otherwise, no use should be made of the Documentation(s) and Product(s) provided by Avaya. All content in this documentation(s) and the product(s) provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components
Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: http://support.avaya.com/Copyright.

Trademarks
The trademarks, logos and service marks ("Marks") displayed in this site, the documentation(s) and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the documentation(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party. Avaya is a registered trademark of Avaya Inc. All non-Avaya trademarks are the property of their respective owners.

Downloading documents
For the most current versions of documentation, see the Avaya Support. Web site: http://www.avaya.com/support

Contact Avaya Support
Avaya provides a telephone number for you to use to report problems or to ask questions about your product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: http://www.avaya.com/support

Contents

Purpose of this document

New in this release

Features 19
- RADIUS proxy authentication service 19
- Access Portal 19

Other changes 20

Customer service

Getting technical documentation 21
Getting product training 21
Getting help from a distributor or reseller 21
Getting technical support from the Avaya Web site 21

Preface

Table of Contents 23
Version Number of This Document 24
How to Use This Guide 24
Symbols and Typefaces 25
Printed Copies of this Manual 25
Other Ignition Documents 25

Introduction to Avaya Identity Engines Ignition Server

Contents 27
What is Avaya Identity Engines Ignition Server? 27
Key Characteristics of Ignition Server 28
The Ignition Server Approach 29
- Wide Set of Criteria for Policy Decisions 29
- Consolidation: Efficiency and Clear Lines of Responsibility 30
- Compliance Automation 31

Ignition Server Feature Overview 31
- Policy and Directory Integration Features 31
- Authentication, Authorization, and Accounting Features 32
- Platform and OS Services 32
- Administration, Control, and Configuration Features 33
- Security Features 33

Ignition Dashboard

Contents 35
Launching Ignition Dashboard 35
Dashboard Best Practices and Design Usage Guidelines 36
Initial Default Display 36
Connection Indicator 37
Error Indicator 37
Dashboard Layout 37
Managing Multiple Ignition Server Sites 38
Setting Up A Site Group 38
Session Timeout 40
Reauthorizing a Session 40
Starting a New Session 41
Quitting From a Timed Out Session 42
Root Certificate 42
Setting Administrator Preferences 42
Setting the User Session Timeout for Dashboard 42
Setting Viewing Preferences for the Monitor View 42
Refreshing the Ignition Dashboard View 43
Exiting Ignition Dashboard 44
Checking the Dashboard Software Version 44

Sites, Nodes, and Settings

Contents 45
Introduction to Dashboard's Configuration View 46
Managing a Site 47
Site Actions 48
Renaming an Ignition Server Site 48
Changing the Ignition Server Administrator Login Name 49
Setting the Ignition Server Administrator Password 49
Managing Ignition Server Services 50
Setting up Ignition Server's RADIUS Service 50
Setting Up Ignition Server's SOAP Service 52
Resetting the SOAP Password 53
Managing a Node 54
Actions Menu for a Node 54
Rebooting a Node 54
Powering Down a Standalone Node 55
Reinitializing the Ignition Server from Dashboard 55
Viewing Logs for a Node 56
Renaming a Node 56
Status Tab 56
Finding the Ignition Server's Serial Number 57

System Tab 57
Viewing Ignition Server's DNS Settings 57
Editing Ignition Server's DNS Settings 57
Setting the Network Routing Configuration 58
Adding a Route to Ignition Server's Routing Table 59
Editing an Existing Route 59
Deleting an Existing Route 59
SNMP Settings 60
SSH Settings 60
SMTP Settings 60
Configuring the Ignition Server's Network Ports 60
Port Configuration Settings 62
Configuring the Admin Port 63
Configuring a Service Port 64
Enabling Service Port 64
Configuring the HA Port 64
Managing Ignition Server Licenses 65
Checking Your Ignition Server Licences 65
Seat limit enforcements 67
Using the NAP license 68
Getting an Ignition Server License 69
Installing an Ignition Server License 69
Replacing an Ignition Server License 69
Making a Backup Copy of Your Ignition Licenses 70
Transferring a License to a Different Ignition Server 70
70
Troubleshoot Tab 70
Running a Ping Test 70
Running a Packet Capture 71

Managing Certificates

Contents of This Chapter 73
Required Types of Certificates 73
Sample Certificates 74
Format of Certificate Files 74
Certificates Tab 74
Getting a New Certificate 75
Create the Certificate Request 75
Import the Certificate 78
Assign the Certificate for Use in Ignition Server 78
Admin Certificate 78
Installing Dashboard's Copy of the Admin Certificate 79
Replacing the Admin Certificate 80

Assigning the SOAP Service Certificate 82
Assigning Protocol Credential Certificates 82
Installing Protocol Root Certificates 83
Adding a Certificate Revocation List URL 84
Adding a Certificate Revocation List URL to Ignition Server 84
Viewing the Certificate Revocation List URLs 85
Viewing a Certificate 86
Deleting a Certificate or Certificate Request 86
Viewing an Existing Certificate Request 87

Authenticators

Contents 89
Introduction to Authenticators 89
Matching an incoming request to an authenticator record 90
Authenticator Hierarchy and Containers 91
How You Build the Hierarchy 91
Using the Hierarchy in Your Policies 92
Default Container 93
Authenticator Hierarchy Tasks 93
Window Layout 93
Creating the Authenticator Hierarchy 94
Adding a Container to an Authenticator Hierarchy 94
Renaming a Container in an Authenticator Hierarchy 95
Deleting a Container from an Authenticator Hierarchy 95
Creating an Authenticator 96
Importing Authenticators 100
Exporting Authenticators 101
Registering Access Portal as an Authenticator in Ignition Server 101
Placing an Authenticator in the Authenticator Hierarchy 103
Finding an Authenticator 104
Pinging an Authenticator 104
RADIUS Global Authenticator 104
Creating the Global Authenticator 105
RADIUS Subauthenticators 105
Viewing and Editing Subauthenticators 106
Creating a Subauthenticator 107
Sorting Subauthenticators 108
Processing Authenticator Requests 109
How Ignition Server Processes RADIUS Requests from Authenticator Bundles 109
Changing the Configuration of an Authenticator 109
Removing an Authenticator From Its Place in the Hierarchy 109
Renaming Authenticators 110
Deleting Authenticators 110

Internal Users, Groups, and Devices

Contents 113
Ignition Server Internal Data Store 113
Internal Users 113
Contents of This Section 114
Internal Users Panel 114
Filtering the Internal Users List 116
Viewing an Internal User 116
Creating an Internal User 116
Internal Devices 119
Contents of This Section 119
Finding an Internal Device 120
MAC Addresses Are Stored Only in the Internal Store 120
Filtering the Device List 120
Creating a Device Record 120
Assigning a Device to a User or Group 123
Editing a Device Record 124
Deleting a Device Record 124
Importing Device Records 124
Exporting Device Records 125
Using MAC Address Wildcarding 125
Internal Groups 126
Contents of This Section 126
Internal Groups Panel 126
Working with the Internal Groups Hierarchy 127
About the Default User Group 128
Adding a New Internal Group 128
Moving an Internal Group in the Hierarchy 129
Renaming an Internal Group 129
Changing a Group's Type Designation 129
Deleting an Internal Group 130
Working with Internal Group Details 131

Directory Services

Contents 133
Quickstart: Directory Services in Dashboard 133
Introduction to Directory Services 134
Directory Services 134
Supported Directory Servers 135
Internal Data Store as a Directory Service 135
Working With Directory Services 135
Connecting to Active Directory 136
AD Connection Settings 136

Preparing to Connect to Active Directory 139
Creating an Active Directory Service: Automatically Configuring 140
Creating an Active Directory Service: Manually Configuring 143
Troubleshooting AD Connections 144
Connecting to an LDAP Service 148
LDAP Connection Settings 148
Creating an LDAP Directory Service: Automatically Configuring 149
Creating an LDAP Directory Service: Manually Configuring 153
Creating a Token Authentication Service 155
Creating a Kerberos Authentication Service 156
Setting up MSCHAPv2 Authentication on LDAP 156
MSCHAPv2 Termination using an LDAP Password Attribute 157
Creating an NT-Hashed Password to Support MSCHAPv2 Termination Against LDAP 158
MSCHAPv2 Termination using a Novell Universal Password 159
MSCHAPv2 Termination using an OID Authentication Descriptor 160
Creating a Radius Proxy Authentication Service 161
Managing Directory Services 161
Editing Directory Service Configurations 161
Renaming a Directory Service 162
Deleting a Directory Service 163
Directory Sets 164
Contents of This Section 165
Directory Sets Panel 165
Adding a Directory Set 166
Renaming a Directory Set 166
Deleting a Directory Set 167
Adding a Directories and Authentication Servers to a Directory Set 167
Setting Fallthrough Rules to Handle Lookup and Authentication Failures 169
Troubleshooting User Lookup and Authentication 170
Checking Directory Service Connections 170
Checking the Group Cache 172
Testing a Directory Service Connection 172
Advanced Troubleshooting for Directory Services and Sets 173

Authentication Services

Contents 177
Setting up a Kerberos Authentication Service 177
Overview of Token Authentication in Ignition 179
How a User Logs In With Token Authentication 179
Components Required for Token Authentication 179
Configuring Token Authentication in Ignition 180
Prepare Your Token Server 180
Connect Ignition Server to RSA Authentication Manager 181

Connect Ignition Server to Another Type of Token Server 183
Add the Authentication Server to Your Directory Set 184
Set Up an Access Policy That Uses Token Authentication 185
Prepare Your Authenticators 187
Connect Ignition Server to Your Authenticators 188
Make Sure Token-Capable Clients Are Installed 188
Perform an End-to-End Test 188
Setting up a RADIUS Proxy Server 189
Creating a RADIUS Proxy Authentication Service 189
Creating a Directory Set 191
Adding the RADIUS Proxy Server to a Directory Set 191
Creating a RADIUS Proxy Server Access Policy 191
Remote RADIUS Server Configuration 192
Proxying of MAC Authentication Requests 192
Editing Authentication Service Configurations 192
Renaming an Authentication Service 193
Deleting an Authentication Service 193
Managing a SecurID Authentication Service 194
Handling Changes to the Node Secret 194
Setting Up Supplicants and Authenticators for SecurID Authentication 195

Virtual Groups and Attributes

Contents 197
Introduction to Virtual Groups and Attributes 197
How Ignition Server Handles Multiple Directories 197
Virtual Groups 198
Browsing Virtual Groups 198
Mappable Group Types for Ignition Server Virtual Groups 199
Adding a New Virtual Group 200
Mapping Groups From a Directory Service 200
Renaming a Virtual Group 202
Deleting a Virtual Group 202
User Virtual Attributes 203
About User Virtual Attributes 203
Browsing User Virtual Attributes 204
Adding a New User Virtual Attribute 205
Mapping Directory Service Attributes to User Virtual Attributes 205
Renaming a User Virtual Attribute 207
Deleting a User Virtual Attribute 208
Device Virtual Attributes 208
Browsing Virtual Attributes for Devices 208
Adding Virtual Attributes for Devices 208

User Authentication Policy

Contents 211
Additional Policy Topics 211
How Ignition Server Authenticates and Authorizes a User 212
Introduction to Policy Management 213
What happened to service categories? 213
The Access Policy Panel 213
Understanding Authentication Policy 214
One Policy Allows Many Authentication Protocols 215
Supported Authentication Types 215
Authentication Policy Window 217
EAP-TLS Authentication 218
Factors That Limit Your Choice of a Protocol Credential Certificate 218
Creating an Authentication Policy 219
Understanding Identity Routing Policy 220
How Ignition Server Looks Up a User for Authentication and Authorization 221
Creating an Identity Routing Policy 223

User Authorization Policy

Contents 227
Introduction to User Authorization Policies 227
Structure of User Authorization Policies 228
Elements of a User Authorization Policy 228
How Ignition Server Evaluates a User Authorization Policy 229
Reading the Rule Summary 232
Attributes Used in Rule Constraints 232
Using Time and Date in a Rule 238
Conjunctions Used to Assemble Constraints Into a Rule 239
Comparison Operators for Rules 240
Creating a RADIUS User Authorization Policy 242
Enabling or Disabling Rules Within a Policy 247
Copying An Authorization Rule 248
Creating an Authentication-Only Policy 249
Using a Device Attribute in a Rule 250

Provisioning Policy

Contents 253
Introduction to Session Provisioning 253
Setting Up Session Provisioning 254
Vendors Panel 255
Outbound Attributes 255
Finding a Global Outbound Attribute 255

Creating a Global Outbound Attribute 256
Overriding the Outbound Attribute Type for One or More Authenticators 257
Outbound Values 258
Finding an Outbound Value 258
Creating an Outbound Value 259
Built-In Outbound Values 261
Setting a Provisioning Value 261
Inbound Attributes 265
Preparing an Inbound Attribute for Use in an Authorization Rule 265
Finding an Inbound Attribute 266
Creating a Global Inbound Attribute 268
Creating a Vendor-Specific Inbound Attribute 269
Device Templates 270
Device Template Window 271
Finding a Device Template 271
Creating a Device Template 271
Modifying a Device Template 273
Applying a Device Template to Your Authenticator 273
Listing Ignition Server's Set of Available RADIUS Attributes 274
Adding a New RADIUS Attribute 275
Listing Ignition Server's Set of Available VSA Attributes 276
Adding a New VSA 276
Adding an Equipment Vendor 277
Provisioning FAQ 278

Client Posture Policy

Contents 281
How Ignition Server Checks Client Posture 281
Enabling NAP on a Windows machine 282
Enable NAP services on the client 282
Enable Enforcement on the client 282
Configure authentication methods 283
Configuring NAP posture profiles 284

VLAN Assignment

Contents 287
Creating a Policy That Assigns Users to VLANs 287
Create the Outbound Attribute 288
Create an Outbound Value for Each VLAN 289
Create VLAN Provisioning Rules 290

Windows Machine Authentication

Contents 295
Introduction to Windows Machine Authentication 295
Supported Authentication Protocols 295
Session Behavior for Windows Machine Authentication 296
NAP support for Peap 296
Setting up Microsoft Windows Machine Authentication 298
Machine Authentication Based on ObjectClass 298
Machine Authentication Based on OU, O, or Group Membership 302
Setting TTL for Windows Machine Authentication 305

TACACS+ Authorization

Contents 307
Introduction to TACACS+ Access Control 307
Privilege-Level TACACS+ Authorization 308
Per-Command TACACS+ Authorization 308
Getting Started 309
Installing Your TACACS+ License 309
Turning on the Ignition Server TACACS+ Service 310
Creating a Command Set 311
Viewing or Editing a Command Set 313
Creating a TACACS+ Access Policy 313
Enable Your Devices for TACACS+ Authorization 317
The TACACS+ Global Authenticator 318

MAC Authentication

Contents 321
Introduction to MAC Authentication 321
Creating a MAC-Auth Policy 322
Setting Up MAC Auth 323
Example MAC Authentication Set-Up Procedure 325
VLAN Assignment Using the Device Record VLAN Fields 330
Allowed MAC Address Formats 333
Notes on Writing MAC Authorization Rules 333
Comparing a Device’s MAC Address 334
Checking a Device’s Group Membership 334

Asset Correlation

Contents 335
Introduction to Asset Correlation 335
MAC Address vs. Windows Machine Authentication 335
Creating Asset Correlation Policies 336

Requiring the User to Connect Using an Allowed Device 337
Requiring the User to Connect Using His or Her Assigned Device 338
Requiring the User to Connect Using a Machine Authenticated-Device 340
Viewing Currently Authenticated Devices 343

Command Line Interface

Connecting to the CLI Through an SSH Connection 345
To configure Ignition Server for SSH: 345
Connecting via SSH 346

A Installing Ignition

Contents 349
Installation Prerequisites 350
Map Out Your Ignition Server Deployment 350
VMware ESXi Server 351
Importing VM 354
Applying the license 361
Installing the license 362
Install Ignition Dashboard on Microsoft Windows 362
Connect to Ignition Server for the First Time 367
Run Ignition Dashboard 367
Install Your Ignition Server Licenses 368
Configure the Ignition Server 368
Uninstalling Ignition Dashboard 368

B High Availability Configuration

Contents 371
HA Terminology 371
Overview of HA Pairs 372
Creating an HA Pair 373
Start and connect the Ignition Server 373
Run the HA Wizard 374
Bind Services to the VIP 380
Restoring a Saved VIP Configuration 381
Managing an HA Pair 383
Managing Virtual Interfaces (VIPs) 384
Breaking an HA Pair Using Dashboard 387
Breaking an HA Pair Using the CLI 388
Reconnecting a Broken HA Pair 388
Reinitializing Nodes in an HA Pair 388
Backing Up Data on an HA Pair 389
Restoring Data on an HA Pair 389

Updating Firmware on an HA Pair 390
Replacing an Ignition Server in an HA Pair 390
Changing the IP Address of the Admin Port or HA Port in an HA Pair 392
Restoring a Non-Responsive Unit in an HA Pair 394

C Backup and Restore Procedures

Contents 395
Introduction to Ignition Server Backup and Restore 395
Creating a Backup 395
Troubleshooting 396
Scheduled Backup 396
Restoring from a Backup File 398
Admin Port IP Address 398
Steps to Restore Data 399
Troubleshooting 400

D Firmware Update Procedures

Contents 401
Checking the Firmware Version 401
View the Current Firmware Version 401
View the Current Firmware Version and All Installed Images and Packages 402
Loading a Firmware Image or Package 403
Activating a Firmware Image or Package 404
406
Activating a Firmware Image on an HA Pair 406
Deleting a Firmware file 407
Viewing image and package information 407
Upgrading from Ignition Server 7.0 to 8.0 409
Upgrading to Ignition Server 8.0 409

E Setting Up Logging

Contents 411
Setting User Preferences 411
Setting Up Ignition Server Logging 411
Turning on Logging 411
Setting the Level of Logging to Be Recorded 413
Setting Up FTP Log Export 413
Exporting logs 414
Directing Log Messages to a Syslog Server 416
Sending Log Messages Via E-Mail 417
Monitoring Ignition Server via SNMP 418
Configuring Ignition's SNMP Service 418

Connecting to Ignition's SNMP Service 420
Example SNMP Queries 420
Data Objects Exposed by the Ignition Server SNMP Service 421
Port Names Used in SNMP Output 422

F Viewing Logs and Statistics

Contents 423
Overview of Logging and Log Types 423
Viewing and Managing Logs 424
Viewing Logs 424
Filtering Your View of the Logs 425
Criteria for Filtering Log Messages 426
Managing the Stored Logs 427
Log Viewer 427
Access Log: RADIUS and TACACS+ Accounting 428
Access Record Details 430
Audit Log 434
Security Log 436
System Log 436
Statistics Tab 437
Transactions Tab 438
Directory Services Tab 440
Protocols Tab 440
System Health Tab 441
Contents of the System Health Tab 442
Viewing the System Health Tab 442
Directory Services Status Tab 442
AAA Summary Tabs 442
Contents of the AAA Summary Tabs 443
Viewing the AAA Summary Tabs 444
Specifying the Number of Records to be Shown 445
User Accounting Tab 445
Contents of the User Accounting Tab 445
Viewing the User Accounting Tab 446
Learned Devices Tab 447
Contents of the Learned Devices Tab 447
Viewing the Learned Devices Tab 447
Filtering the Learned Devices Tab 447
Revoking the Session of a Machine-Authenticated Device 448
Debug Logs 448

G Troubleshooting

Contents 451

Trouble Ticket Generation 451
Troubleshooting Common Problems 452
Problem: Cannot Connect to Ignition Dashboard 452
Problem: Connecting Dashboard to Ignition Server Fails 453
Problem: Ignition Server fails to respond to RADIUS and/or TACACS+ requests 453
Problem: Authorization policy stops working unexpectedly 453
Problem: Authentication fails on Active Directory 455
Problem: HA Set-up Fails 456
Problem: Two primary HA nodes detected 456
Problem: Errors Occur During Directory Service Set-Up 457
Problem: Unable to Map Virtual Attributes 458
Problem RADIUS Proxy Service Fails 458

Glossary 461

Index 469

Purpose of this document

The *Avaya Identity Engines Ignition Server Administration* guide explains how to set up and use the Avaya Identity Engines Ignition Server (AIEIS) and Avaya Identity Engines Ignition Dashboard.

This administration guide is written for network administrators using the Avaya Identity Engines Ignition Server. As an administrator, you are responsible for setting up and maintaining the users, devices, objects, policies, and configurations that Ignition uses to secure and control access to your networks and other resources. We assume that you are familiar with network terminology and have experience setting up and maintaining networks and their security implementations.

New in this release

The following sections detail what is new in *Avaya Identity Engines Ignition Server Administration* (NN47280-600).

- "Features" on page 19
- "Other changes" on page 20

Features

See the following sections for information about feature changes:

- "RADIUS proxy authentication service" on page 19
- "Access Portal" on page 19

RADIUS proxy authentication service

Avaya Identity Engines Ignition Server (AIEIS) Release 8.0 offers the ability to set up a RADIUS proxy authentication service. A RADIUS proxy server forwards (or proxies) RADIUS requests to a remote RADIUS server for authentication. The Ignition Server can act as the RADIUS proxy server that forwards the authentication requests, or as the remote server that receives the authentication requests.

For more information about setting up a RADIUS proxy authentication service, see "Setting up a RADIUS Proxy Server" on page 189

Communication between the RADIUS server that forwards the requests and the remote server only occurs using the password authentication protocol (PAP). For more information on authentication protocols, see "Supported Authentication Types" on page 215.

Access Portal

Avaya Identity Engines Ignition Access Portal (Access Portal) is a new licensed feature in AIEIS Release 8.0. Access Portal is a virtual machine based captive portal and firewall distribution that controls the access of client devices to the network. Access Portal blocks all traffic from client devices and allows network access only after successful authentication. Access Portal allows guests with non-802.1X compatible equipment to authenticate and connect to the network in your organization. You must apply the Access Portal license to enable this feature.

For more information about applying feature licenses, see “Managing Ignition Server Licenses” on page 65.

After you apply the Access Portal license, you must register Access Portal as an authenticator in Ignition Server. For information on how to register Access Portal as an authenticator in Ignition Server, see “Registering Access Portal as an Authenticator in Ignition Server” on page 101.

For information on setting up Access Portal to work with the Ignition Server, see the *Avaya Identity Engines Ignition Server Access Portal Administration* guide.

Other changes

Information on Trusted Network Connect (TNC) Posture was removed from this guide as AIEIS Release 8.0 does not support the TNC Posture feature. Customers who upgrade from a previous version of AIEIS that deployed TNC will lose this functionality and should move to Microsoft Network Access Protection (NAP) based posture checking.

Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.

Preface

This Administration guide explains how to set up and use the Avaya Identity Engines Ignition Server (AIEIS) and Avaya Identity Engines Ignition Dashboard.

Table of Contents

Topic	Found in Chapter...
Overview of AIEIS firmware, configuration, and administration	: Introduction to Avaya Identity Engines Ignition Server
AIEIS configuration and management	: Ignition Dashboard
Sites and nodes: Managing the Ignition Servers on your network and enabling communication with them.	: Sites, Nodes, and Settings
Digital certificates: Communication between Ignition Server components and end-user workstations is typically secured using certificates.	: Managing Certificates
Authenticator: Represents an individual or group of wireless access points, VPN concentrators or servers, Ethernet switches, WLAN switches, or routers.	: Authenticators
Internally stored users and groups: In addition to reading users from your enterprise AD or LDAP, you can store them locally in Ignition's internal store.	: Internal Users, Groups, and Devices
Directory services: Repositories of user identities and attributes such as Active Directory, LDAP, and token servers **Directory sets:** Groups of directory services	: Directory Services
Authentication services provide specialized forms of user identity verification such as RSA SecurID token authentication and Kerberos authentication.	: Authentication Services
Virtual groups and attributes: You can map inconsistent attribute or group names to a single "virtual" attribute or group name for use in Ignition access policy.	: Virtual Groups and Attributes
User authentication policies: Each access policy gets a policy that tells Ignition how to validate the user's identity.	: User Authentication Policy
User authorization policies: Each access policy gets a policy that governs which user can access which parts of the network. You base your rules on user attributes and other criteria.	: User Authorization Policy

Provisioning policies: Optionally, each access policy may have a provisioning policy that sets network session and switch parameters for the user.	: Provisioning Policy
Client Posture Policy: Optionally, you can require that users' laptops meet a minimum standard of system health with respect to viruses, OS patches, and security updates.	: Client Posture Policy
VLAN assignment: You can set up provisioning policies that assign each user to an appropriate VLAN.	: VLAN Assignment
Windows machine authentication lets you check that the connecting device is known to Active Directory before you allow it to connect.	: Windows Machine Authentication
MAC authentication allows operator-less devices to connect to the network, and it allows Ignition to record what device a user signed on with.	: MAC Authentication
TACACS+ authentication allows you to set and enforce access rights for system and network administrators.	: TACACS+ Authorization
Asset correlation works with Windows machine authentication or MAC authentication to force users to use only their assigned computer to connect.	: Asset Correlation
Using the command line to manage Ignition.	: Command Line Interface
Installing and configuring the **Ignition Server** and **Ignition Dashboard**	Appendix A: Installing Ignition
High availability configuration: pairing nodes to allow failover in the event of Ignition Server failure	Appendix B: High Availability Configuration
Backing up the Ignition Server data and configuration	Appendix C: Backup and Restore Procedures
Updating Ignition Server firmware	Appendix D: Firmware Update Procedures
Logging, monitoring, and auditing: tracking and logging user, resource, and system activity	Appendix E: Setting Up Logging
	Appendix F: Viewing Logs and Statistics
Troubleshooting information	Appendix G: Troubleshooting
Glossary of Ignition and network access terms	Glossary

Version Number of This Document

This document covers Avaya Identity Engines Ignition Server Release 8.0.

How to Use This Guide

This Administration Guide is written for network administrators using the Avaya Identity Engines Ignition Server. As an administrator, you are responsible for setting up and maintaining the users, devices, objects, policies, and configurations that Ignition uses to secure and control access to your networks and other resources. We assume that you are familiar with network terminology and have experience setting up and maintaining networks and their security implementations.

Symbols and Typefaces

This book uses the following visual cues to alert you when the text covers an important or easily overlooked issue:

Symbol	Meaning
‡	A selection or step within a menu, such as **File → Open** or **Edit → Paste**.
← ↑ → ↓ ✓ ✕	The buttons on the front panel of the Ignition Server.
Italic type	Italics signify new or important terminology and variable input for commands.
Bold type	Boldface signifies a command you click or an option you select.
	**Tip.** Indicates a shortcut or a resolution to a common problem.
	**Note.** Indicates background information or a reference to other sources of information.
	**Warning.** Indicates a situation where an error could cause data loss or a security breach.
	**Danger.** Indicates a situation where an error could cause bodily injury or equipment damage. Follow instructions carefully.
	**Danger.** Hazardous voltages are present. To reduce the risk of electric shock, follow the instructions.
	**Danger.** Hot surface. Avoid contact.

Printed Copies of this Manual

To obtain printed copies of this manual, contact Avaya as shown in “Getting technical documentation” on page 21.

Other Ignition Documents

The following AIEIS documents also ship with your AIEIS and are available from the Avaya support web site. (See “Getting technical documentation” on page 21.)

- The *Avaya Identity Engines Ignition Server Getting Started Guide* shows you how to get Ignition running and how to connect Ignition to your switches, access points, and user data stores.
- The *Avaya Identity Engines Ignition Guest Manager Configuration Guide* explains how to install the Ignition Guest Manager provisioning tool.

- The *Avaya Identity Engines Ignition CASE Administration* guide explains how to install, configure, and deploy Avaya Identity Engines Ignition Client for Accessing Secure Enterprise (CASE).
- The *Avaya Identity Engines Ignition Access Portal Administration* guide explains how to install and configure Access Portal.
- The *Avaya Identity Engines Ignition Analytics* guide explains how to install, configure, and maintain the Ignition Analytics application.
- The *Avaya Identity Engines Ignition Server Release Notes* provide upgrade instructions, last-minute information, and notes about known bugs.

Introduction to Avaya Identity Engines Ignition Server

This chapter introduces the Avaya Identity Engines Ignition Server (AIEIS).

Contents

- "What is Avaya Identity Engines Ignition Server?" on page 27
- "Key Characteristics of Ignition Server" on page 28
- "The Ignition Server Approach" on page 29
- "Ignition Server Feature Overview" on page 31

Installation prerequisites, installation procedures, and the initial connection to the Avaya Identity Engines Ignition Server from Avaya Identity Engines Ignition Dashboard are described in Appendix A, "Installing Ignition".

What is Avaya Identity Engines Ignition Server?

Avaya Identity Engines Ignition Server is an 802.1X-capable RADIUS authentication server and TACACS+ server that grants or denies users access to your network based on your policies. Avaya Identity Engines Ignition Server lets you create a single set of policies that control access for all the ways users connect: via a wired Ethernet jack, wireless, or VPN. Access policies are stored on the Ignition Server, and user accounts remain in your traditional user store(s) such as LDAP and Active Directory servers.

Avaya Identity Engines Ignition Server includes an easy-to-configure policy engine that lets you make network access decisions based on the user's identity, his or her account details and group memberships, the location of the login attempt, the time of day, and many other pieces of information. For example, an Avaya Identity Engines Ignition Server policy can grant a user access based on her identity, her point of access (which network switch or wireless access point she is connecting through), and her laptop's security state (ensuring her laptop is a company-owned laptop as recorded in the corporate Active Directory store and ensuring it has up-to-date anti-virus profiles installed).

Ignition Server's abilities to check whether the user's workstation has passed MAC authentication, Windows machine authentication, and/or a security posture check are key features that set it apart from other network access control tools. In short, Ignition Server lets you combine many policy elements

to enforce a single rule such as: authenticate the user with PEAP/MSCHAPv2, check that his or her device has been authenticated, and if those are successful, assign the user to the appropriate VLAN based on his or her role. Ignition Server also authenticates devices and can be set to offer a bypass of 802.1X authentication for older devices on your network that cannot perform an 802.1X authentication.

Key Characteristics of Ignition Server

The following are the most important, distinct characteristics of Ignition Server:

- **Non-intrusive, out-of-band**: Ignition Server is an out-of-band access control solution and thus easier to install and to scale up than an inline solution. “Out-of-band” means that only the client’s *network sign-on transaction* travels through Ignition Server. After it is signed on, the client’s network traffic travels its usual path.
- **Standards-oriented**: Since Ignition Server is a standards-compliant RADIUS server, it interacts with and can control nearly *every* type of network endpoint: wired switches, wireless access points, and VPN concentrators.
- **Consolidated AAA platform**: Consolidated AAA platform: Ignition Server handles the three A’s: authentication, authorization and accounting. Ignition Server works with your existing authentication servers (SecurID, Active Directory, and so on) to authenticate the connecting user or device; it uses its policy engine and provisioning framework to authorize the user/device, and it maintains accounting records (audit log) of these connection events in a number of formats.
- **Scales up well**: One Ignition Server serves as the AAA/RADIUS server for *many* network-edge devices: wired, wireless, and VPN.
- **Multiple directory support**: No duplication of user accounts is required. Ignition Server authenticates users and devices against your existing data store that holds those accounts. Ignition Server retrieves information about the user and/or device from many different types and instances of directories: Active Directory, Novell eDirectory, SunONE LDAP, Oracle OID, LDAP, the Ignition Server-local internal store, and others.
- **Split authentication/lookup**: Ignition Server can be set to authenticate the user against one service and retrieve his or her account details from a separate service for authorization. For example, you can authenticate using RSA SecurID and look up the user account from an LDAP service.
- **Very flexible policy engine**: Ignition Server lets the network administrator use a wide range of criteria including user attributes, device attributes,

access type, location, date/time, and others, to make precise, targeted access decisions.

- **Guest access**: A suite of supporting tools lets the network administrator safely and efficiently grant guests access to the network: *Avaya Ignition Server Guest Manager* delegates the administrative task of adding temporary users and importing groups of temporary users, and it can allow self provisioning, if so configured.
- **Role-based networking** (also called role-based access control): The user's role or group affiliation recorded in the directory determines what networks and resources he or she can access.
- **High Availability**: You may deploy two Ignition Servers as a linked pair that offers a highly available RADIUS service.

The Ignition Server Approach

The Avaya Identity Engines Ignition Server platform provides a comprehensive set of network access management services in a secure, scalable, standards-based appliance designed for enterprise or campus deployment. With its built-in ability to use multiple and varied enterprise user directories and its easy-to-add guest management tools, Ignition Server gives the network administrator the confidence to allow both permanent and short-term users to connect to the network, while ensuring that each user sees only the appropriate portions of the network and that all access events are logged to address internal auditing needs and government reporting requirements.

Wide Set of Criteria for Policy Decisions

The Ignition Server policy engine enables you to set precise network access policies based on a large set of criteria, including:

- *user attributes*, such as roles or group membership;
- end-user *device attributes*, including device anti-virus and security posture;
- *context*, such as time of day, IP address, or location; and
- details of the *authenticator device* or service (the switch, wireless access point or VPN concentrator), such as vendor, location, or service type

Since Ignition Server supports a large set of parameters in its access policies, the network administrator can map access policies directly to the existing relationships and rules in the organization. For example, the policy engine allows network administration to enforce business rules like the following:

- Any user that belongs to group *Contractors-Accenture* can access the wireless access points in Building 3/Floor 4 between the hours of 8:00 and 17:00 and should be placed on VLAN 250.

- Users accessing VPN and wireless require *SecurID authentication*, but users accessing wired ports require *password only*.
- Any employee with role of *Faculty* gets a high-quality-of-service network session throughout Campus B.

Consolidation: Efficiency and Clear Lines of Responsibility

The “many-silos” approach to security enforcement, which relies on multiple, independently-managed security domains, simply does not work. Having multiple silos adds complexity in administering users and policies, raises the chance of costly errors, and muddies the lines of responsibility that are a crucial “best practice” for network security.

By contrast, Ignition Server consolidates your network access control to a single *policy decision point* that makes and logs all access decisions. Consolidating access decisions means:

- Your network access policies are enforced consistently across wired, wireless, VPN, and remote access.
- Users can access the network through any allowed switch or access point, but wherever they connect, the log entry resolves to the user’s account in the appropriate enterprise user directory. As a result, security and compliance audits can be streamlined.
- You can more quickly extend your network and deploy new network services, since adding a new access point or network in Ignition Server requires just a few steps.

While Ignition Server acts as the single *policy decision* point, it avoids the creation of an *administrative* choke point. Ignition Server does this by acting as the single point that makes and logs access decisions, while leaving the management of user account data where it belongs: in your enterprise directories (AD, LDAP, etc.). Having a single policy decision point reduces security risks. Leaving your account data where it is reduces duplicate tasks for network security personnel and helps keep the lines of responsibility clear. Only those who are responsible for account management can update accounts.

Ignition Server is able to leave your account data where it is thanks to Ignition Server identity routing. Identity routing lets you specify a search order that directs the Ignition Server to search one or more user directories (of any type—AD and most flavors of LDAP are covered) to find the correct user account. Identity routing helps you avoid creating duplicate user accounts.

Compliance Automation

Compliance requirements such as Sarbanes-Oxley and HIPAA have had an increasing impact on network planning, deployment, and auditing. The optional Ignition Analytics application provides pre-defined, automated reports that simplify periodic monitoring and audits. Ignition Server provides an aggregated log record for all network access (wired, wireless, and VPN), with configurable log levels for runtime and administrative activity. Alerts can be generated based on policy violations or other triggers, with reporting that can reveal patterns within individual categories of logged events or users.

Ignition Server Feature Overview

The sections below describe the main features of Ignition Server.

Policy and Directory Integration Features

Avaya Identity Engines Ignition Server is a next-generation enterprise-class, secure, robust, and scalable network identity management solution, whose simple task-based user interface greatly eases security management and policy authoring. Ignition Server authenticates and authorizes enterprise users for network access, capturing detailed audit logs and generating key reports needed for regulatory compliance.

Ignition Server goes beyond traditional network AAA (authentication, authorization, accounting) products because it provides unparalleled integration with enterprise directories. Examples include Microsoft Active Directory, Novell eDirectory, and Sun Java System Directory Server, and RSA Authentication Manager (formerly RSA ACE/Server).

By intelligently interfacing with multiple directory stores, Ignition Server provides transparent authentication and flexible authorization policies using corporate group hierarchies, role information, and other user attributes from any number and type of enterprise directories.

Ignition Server surpasses other network AAA products by its ability to support multiple network services simultaneously. For traditional AAA servers, each additional network service that is added requires installing, maintaining, and administering another AAA server with its own policy and user database configurations designed specifically for that one network service.

Ignition Server, however, enables you to consolidate all existing AAA services in a single system, managing global policies across all network services and improving manageability and visibility. It provides auditing of both runtime and administrative activity, thereby improving security and compliance.

By leveraging the RADIUS protocol that virtually all network devices support, Ignition Server provides vendor-independence and interoperability. It can integrate seamlessly with your existing switches, routers, firewalls, wireless

access points, wireless switches, VPN servers, and remote access servers from leading manufacturers such as Cisco, Juniper, Extreme, Foundry, HP, Avaya, Microsoft, Aruba, Trapeze, and others.

Furthermore, distributed enterprises benefit from Ignition Server's central management of distributed Ignition Servers and the ability to ensure that policies are applied consistently across the organization.

For deployments where fault-tolerance is a necessity, Ignition Server offers an active/passive high availability option.

Ignition Server is designed as a multi-protocol authentication platform enabling enterprises to consolidate authentication services for networks and applications in the future. It provides superior security with its hardened operating system, encrypted file system, and anomaly detection capabilities.

Authentication, Authorization, and Accounting Features

The traditional definitions for AAA do not have the necessary richness or granularity to meet modern enterprise requirements. Authentication must be configurable based on specific authenticator-provided information, and on rules that specify the credentials acceptable for validating the identity of users coming through that authenticator. Such credentials may include passwords, hardware tokens, digital certificates, and so on.

Ignition Server provides that needed granularity. After a user is authenticated, Ignition Server makes an authorization decision (that is, it determines the user's access privileges to the network service) using the authentication information plus rules and relevant data pulled from back-end stores. After a user is authorized, Ignition Server invokes provisioning objects that set the attributes of the user session, such as VLAN assignment, access control lists (ACLs), quality of service (QoS), and so on.

Accounting and auditing traditionally exclude real-time analysis of user activities. Ignition Server differs by maintaining a log of users' conformance to access policies, rather than focusing on billing and usage as other AAA products do.

Platform and OS Services

The Ignition Server utilizes a 64-bit high performance CPU running a hardened operating system and protocol stack from RedHat; enabled journaling in file system for reliability.

Administration, Control, and Configuration Features

Ignition Dashboard, the graphical user interface for the Avaya Identity Engines Ignition Server, makes it simple to create, view, or alter configuration information for authenticators, access policies, and the policies that apply to authentication and authorization.

The Ignition Dashboard configuration options enable you to establish authorization policies using virtual attributes corresponding to the user attributes maintained in your directories, as well as contextual information relating to the access request.

You can name and specify categories in a hierarchical organization for Ignition Server's portrayal of your network. You define the categories and their placement in the hierarchy, making it easy for you to find the type or location of any authenticator in your network.

Similarly, Ignition Server makes it easy for you to represent the group memberships of users and groups in a tree diagram, whose content also appears in the windows showing user detail information. This applies to user records that are created and maintained in the Ignition Server internal data store.

Security Features

The following table summarizes Ignition Server features that prevent threats from being exploited and that detect and report acts or events with security risk potential.

Table 1 Ignition Server Security Features

Feature	Ignition Server Action
Network port lockdown	All unused network service ports are locked down - only specifically enabled services are available.
Network anomaly detection	Ignition watches for malformed packets destined for any services it exposes and logs them. Duplicate MAC address detection reports an error to the operator.
Network and port segmentation	Ignition Server enables you to assign separate ports for different traffic. Port status is shown in Dashboard, as explained in "Managing a Node" on page 54. Changes to the node's network interface configuration are recorded in the logs. Network interface settings are stored in the Ignition Server platform's configuration database and included in standard backup and restore operations. As shown in "Ignition Server Ports and Allowed Traffic" on page 60, you may place limits on what traffic each Ignition Server ports may carry.

Ignition Dashboard

As the administrator managing the Avaya Identity Engines Ignition Server, your primary tool is the **Ignition Dashboard** application located on your personal computer or workstation. Dashboard lets you manage and monitor the operation of the Ignition Server and set up user authentication and authorization policies for your network.

This chapter provides an overview of the Ignition Dashboard and a description of the **Administrator** menu commands.

For Ignition Server Installation, setup, and initial login instructions, see Appendix A, “Installing Ignition”. For information on additional Ignition Server management operations, see: Appendix C, “Backup and Restore Procedures”; Appendix D, “Firmware Update Procedures”; and Appendix E, “Setting Up Logging”.

Contents

- “Launching Ignition Dashboard” on page 35
- “Dashboard Best Practices and Design Usage Guidelines” on page 36
- “Initial Default Display” on page 36
- “Connection Indicator” on page 37
- “Error Indicator” on page 37
- “Dashboard Layout” on page 37
- “Managing Multiple Ignition Server Sites” on page 38
- “Session Timeout” on page 40
- “Root Certificate” on page 42
- “Setting Administrator Preferences” on page 42
- “Refreshing the Ignition Dashboard View” on page 43
- “Exiting Ignition Dashboard” on page 44
- “Checking the Dashboard Software Version” on page 44

Launching Ignition Dashboard

To run Ignition Dashboard:



1. Double-click the Ignition Dashboard icon on your desktop or select the command **Start → Programs → Ignition Dashboard → Ignition Dashboard**. The following login window appears:

Figure 1 Ignition Dashboard Login



2. Type the Ignition Server administrator **User Name** and **Password**. The default user name and password are “admin” and “admin”. For security, make sure you change the user name and password from their default settings. See “Changing the Ignition Server Administrator Login Name” on page 49.
3. In the **Connect To** field, do one of the following:
   - To connect to an individual Ignition Server site, type the hostname or IP address of your Ignition Server.
   - To connect to a group of Ignition Server sites that you manage, choose the Site Group Name in the **Connect To** drop-down list. (See “Managing Multiple Ignition Server Sites” on page 38 for details.)
4. Click **OK**. If you are unable to log in, see “Problem: Cannot Connect to Ignition Dashboard” on page 452.

Dashboard Best Practices and Design Usage Guidelines

Observe the following guidelines and limitations when using Dashboard:

- **No concurrent administrator sessions:** Avaya Identity Engines Ignition Server strongly recommends that, at any given time, only one administrator should use Ignition Dashboard to make edits. Other administrators can launch their own Dashboard sessions to view data, but they should not make edits. If multiple administrators make edits concurrently, data inconsistencies might result.
- **No spaces after text entries:** When you enter text into a field in Ignition Dashboard, make sure there are no space characters after the text. Ignition Server rejects the entry if it contains trailing spaces.

Initial Default Display

When you initially launch Ignition Dashboard, Ignition Server displays the Default Admin Certificate window. Avaya Identity Engines Ignition Server Server provides you with a default admin certificate. When you initially launch Ignition Dashboard, the **Default Certificate** window appears. This message

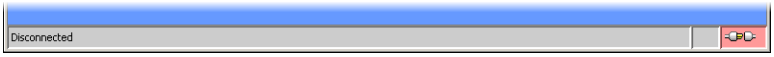
window continues to appear until you provide an Admin certificate for your organization. Click **OK** to close this window. It is strongly recommended that you acquire and install an Admin certificate specifically issued for your organization.

Connection Indicator

The connection indicator is located in the lower right hand of the Ignition Dashboard window frame. A green background with a “plugged-in” icon indicates an active connection:

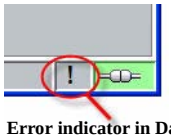


When you log out, and Ignition Dashboard is not connected to an Ignition Server, the connection indicator displays a red indicator with an “unplugged” icon:



Error Indicator

If an error occurs on the Ignition Server, Dashboard displays an alert icon at the bottom of the Dashboard main window, as shown below. Click the alert icon to view a dialog window showing a description of the error. Before you dismiss the dialog, you should clear the message by clicking on the message and clicking the **Clear** button to clear the message. If you do not clear the message, it remains there the next time you open the dialog.

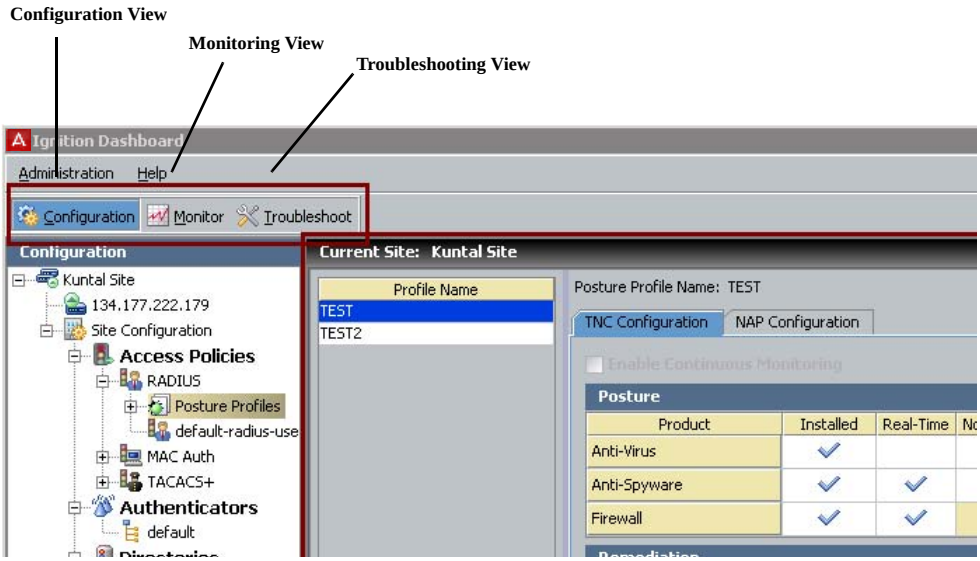


Error indicator in Dashboard

Dashboard Layout

The figure below illustrates Ignition Dashboard, your graphical user interface for configuring your Ignition Server(s).

Figure 2 Ignition Dashboard



Managing Multiple Ignition Server Sites

If your Ignition Servers are installed in multiple sites, Dashboard makes it easy to connect to them and to switch your connection from one site to another. As shown in the figure that follows, once you have grouped your Ignition Server sites into a *site group*, you connect to the site group rather than to a single Ignition site.

Figure 3 Logging in to a Site Group, instead of an Ignition Server Site



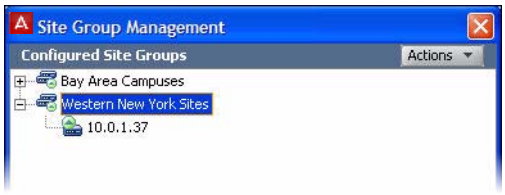
To set up a site group, use the instructions that follow.

Setting Up A Site Group

A site group allows you to log in once to connect to a number of Ignition Servers installed in multiple locations. To set this up:

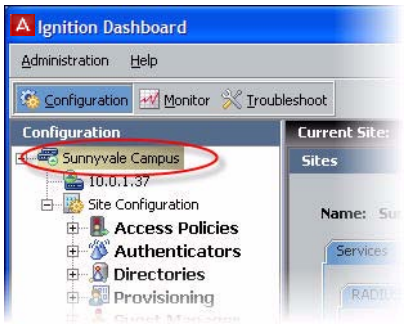
1. Make sure your Ignition Server site has been given a unique name. For instructions, see “Renaming an Ignition Server Site” on page 48.
2. In the main window of Dashboard (with Dashboard already connected to an Ignition Server), select **Administration: Site Group Management**.

Figure 4 The Site Group Management window



3. In the Site Group Management window, select the command **Actions: Add Group**.
4. Type a name for the site group and click **OK**.
5. Add an Ignition Server to the Site Group:
   - In the Site Group Management window, click the name of your group and select the command, **Actions: Add Site Group**.
   - In the Add Site Group window enter the **Site Group Name**. Then, click **Add** at the bottom of the window. This will bring up another window in which you enter the **IP address**, **username**, and **password** for the user you wish to add and click **OK**.
6. Repeat Step 5 for the other appliances in the group.
7. Set the password for the site group:
   - In the Site Group Management window, click the name of your group and select the command, **Actions: Modify Configuration Password**.
   - In the dialog window, type the **Current Password** (“admin” is the default) and type the desired **New Password**. Type it again to confirm it, and click **OK**.
8. Click **Close** to close the Site Group Management window.
9. Disconnect Dashboard from the current appliance and reconnect to the site group:
   - In the main window, select **Administration: Logout**
   - Select **Administration: Login**.
   - Type the **User Name**, “admin” and type the **Password** you set for the site group in Step 7
   - Click the **Connect To** drop-down box and choose the name of your site group.
   - Click **OK** to connect.
10. Dashboard’s Configuration tree lists all the sites in your site group. Click the name of a site to manage that site.

Figure 5 Clicking a Site in the Configuration tree



Session Timeout

By default, the Ignition Dashboard user interface is set to time out after a session of Ignition Dashboard has been inactive for 10 minutes. When there is a timeout for your current session, Ignition Server displays the following window:.

Figure 6 Session Timeout window



You can choose to reauthorize the current session of Ignition Dashboard on the same Ignition Server, start a new session, or exit from the Ignition Dashboard application. Consult the following sections for instructions:

- “Reauthorizing a Session” on page 40
- “Starting a New Session” on page 41
- “Quitting From a Timed Out Session” on page 42

Reauthorizing a Session

In order to reauthorize the current session of Ignition Dashboard on the same Ignition Server,

1. Click the **Reauthorize** button in the session timeout window shown in Figure 6. The **Reauthorize** window appears.
2. Enter your administrator password.
3. Click **OK**.

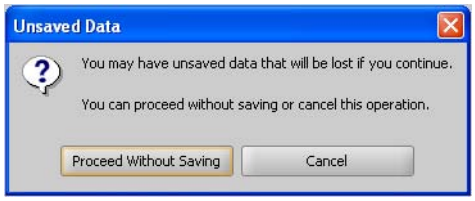
The Ignition Dashboard reappears in the same state as it was before the timeout.

Starting a New Session

When you elect to start a new session of Ignition Dashboard, you are essentially disconnecting from the current Ignition Server to which it was connected. Ignition Server needs to take the appropriate actions with respect to any unsaved data associated with the Ignition Server to which it was connected.

In order to start a new session of Ignition Dashboard,

1. Click the New Session button in the Session Timeout window shown in Figure 6.
2. Ignition Server displays the **Unsaved Data** window to alert you regarding the possibility of losing any unsaved data associated with the session that has timed out.



3. If you have data that you need saved from the session that has timed out,
   - Click **Cancel**. Ignition Server displays the session timeout options as shown in Figure 6.
   - Use the **Reauthorize** button to reauthorize the current session that timed out.

   The Ignition Dashboard reappears in the same state as it was before the timeout.
4. If you want to discard the data associated with the session that has timed out,
   - Click the **Proceed Without Saving** button. Ignition Server drops the connection with the timed out Ignition Server and displays the Login window. (See Figure 1).
   - Enter the user name, password, and the name of the desired Ignition Server. Click **OK**.

Ignition Server starts a new session with the details you entered in the Login window.

Quitting From a Timed Out Session

When you click Exit in the session timeout window (Figure 6), Ignition Server disconnects the timed-out session of Ignition Dashboard from the Ignition Server, and closes the Ignition Server dashboard application.

Root Certificate

Ignition Dashboard uses a root certificate stored in its keystore to verify the identity of the Ignition Server before connecting to it. If you cannot connect due to a certificate problem, see "Installing Dashboard's Copy of the Admin Certificate" on page 79.

Setting Administrator Preferences

The Preferences window allows you to specify Dashboard's timeout settings and log display settings. To open the window, select the command **Administration: Preferences** from the Dashboard main window. For instructions, consult one of the sections below:

- "Setting the User Session Timeout for Dashboard" on page 42
- "Setting Viewing Preferences for the Monitor View" on page 42

Setting the User Session Timeout for Dashboard

The Dashboard window locks automatically after a period of inactivity. Set your locking preferences in the **Preferences** window as follows:

1. Select the command **Administration: Preferences** from the Dashboard main window.
2. Do one of the following:
   - To turn off locking, select **Do Not Lock Ignition Dashboard**.
   - To turn on window locking, select **Lock Ignition Dashboard**, and specify the timeout period in minutes in the **Wait** field. After the specified period of inactivity, Dashboard locks and requires a password for unlocking. To change the password, see "Setting the Ignition Server Administrator Password" on page 49.

Setting Viewing Preferences for the Monitor View

The **Logging** and **Monitor** tabs of the Preferences window enables you to set the viewing preferences for Dashboard's Log Viewer tab. (For information on the Log Viewer, see "Viewing and Managing Logs" on page 424.)

To set your log viewing preferences, do the following:

1. Select the command **Administration: Preferences** from the Dashboard main window.
2. Click the **Logging** tab.

- Tick the **Automatically refresh logs on tab selection** checkbox to force Ignition Server to load the latest log messages when you click on a tab in the Log Viewer. If you leave this checkbox unticked, then you must use the Refresh button in the Log Viewer to load log messages.
- In the **Order to display log records** section, tick **Most recent record first** to display the latest log messages at the top of the Log Viewer tab, and subsequent records in reverse chronological order, or tick **Oldest records first** if you want to display the oldest records at the top and subsequent records in chronological order.
- The **Number of records to display** field sets the page size for the Log Viewer. Tick **Fit in screen** if you want the Log Viewer to load enough records to fill the window. To set a custom page size, tick **User Specified** and use the up/down arrows to specify the number of log records to load per page.
- The **Display Full Log Message Using** radio buttons let you choose how Dashboard displays detailed logs like the Access Record Details record. Choose **Tooltip** to have Dashboard display the details in a floating dialog box that appears when you click the record's row. Choose **Region at Bottom of Log Viewer** to display a dedicated details panel below the list in the Log Viewer. See "Specifying How Dashboard Displays Access Record Details" on page 431 for an explanation of the two display styles.

3. Click the **Monitor** tab. The number of authentication/authorization records to display field limits the number of records shown in the site level **AA summary** tabs in the **Monitor** view of Dashboard to 200. This is not modifiable.

   The limit you set here applies as a single, total limit on the number of records shown at any given moment across all three tabs, **RADIUS AAA Summary**, **TACACS+ AAA Summary**, and **Guest Manager AAA Summary**. In other words, if you set a limit of 200, and the most recent 200 records are RADIUS authorizations, then the **RADIUS AAA Summary** will show 200 records, and the **TACACS+ AAA Summary** and **Guest Manager AAA Summary** tabs will show zero records.

4. Click **OK** to apply your changes.

Refreshing the Ignition Dashboard View

To update Dashboard's display, right-click on your site in the Configuration hierarchy tree and select **Refresh Site**. Ignition Server refreshes the display with the latest Ignition Server data.

Exiting Ignition Dashboard

The **Administration: Exit** command disconnects Ignition Dashboard from the Ignition Server (to which it was connected), and closes the current session of Ignition Dashboard on your personal computer or workstation.

Checking the Dashboard Software Version

To check the version of Ignition Dashboard you are running, select **Help: About** from the main window. To check the firmware version, see "Checking the Firmware Version" on page 401. To check the documentation version, see "Version Number of This Document" on page 24.

Sites, Nodes, and Settings

This chapter introduces the concept of the Avaya Identity Engines Ignition Server site and nodes and explains how to manage your Ignition Server network settings using the Configuration view of Ignition Dashboard.

Contents

- "Introduction to Dashboard's Configuration View" on page 46
- "Managing a Site" on page 47
  - "Changing the Ignition Server Administrator Login Name" on page 49
  - "Setting the Ignition Server Administrator Password" on page 49
  - "Managing Ignition Server Services" on page 50
  - "Setting up Ignition Server's RADIUS Service" on page 50
  - "Setting Up Ignition Server's SOAP Service" on page 52
- "Managing a Node" on page 54
  - "Rebooting a Node" on page 54
  - "Powering Down a Standalone Node" on page 55
  - "Reinitializing the Ignition Server from Dashboard" on page 55
  - "Status Tab" on page 56
  - "System Tab" on page 57
- DNS settings:
  - "Viewing Ignition Server's DNS Settings" on page 57
  - "Editing Ignition Server's DNS Settings" on page 57
- Network settings:
  - "Setting the Network Routing Configuration" on page 58
- Routing:
  - "Adding a Route to Ignition Server's Routing Table" on page 59
  - "Editing an Existing Route" on page 59
  - "Deleting an Existing Route" on page 59
- "SNMP Settings" on page 60

- "SSH Settings" on page 60
- "SMTP Settings" on page 60
- Setting up the Ethernet ports:
  - "Configuring the Ignition Server's Network Ports" on page 60
  - "Configuring the Admin Port" on page 63
  - "Configuring a Service Port" on page 64
  - "Enabling Service Port" on page 64
  - "Configuring the HA Port" on page 64
- "Managing Ignition Server Licenses" on page 65
- "Troubleshoot Tab" on page 70
  - "Running a Ping Test" on page 70
  - "Running a Packet Capture" on page 71

Introduction to Dashboard's Configuration View

The Configuration view of Dashboard is your primary tool for managing the network settings and physical settings of Ignition Server. Before you begin configuring Ignition Server, it's important to understand these two concepts: an *Ignition Server site* and an *Ignition Server node*. The *site* is your entire Ignition Server installation, whereas a *node* is an individual Ignition Server appliance.

Depending on your configuration, your Ignition Server *site* might consist of a single *node* (an Ignition Server) or a pair of *nodes* (a high availability pair of Ignition Servers). Dashboard's Configuration view lets you perform the following tasks on your Ignition Server site and nodes:

- Set up the **Configuration Hierarchy**.
- **Sites and Maintenance:** rename a site, backup and restore Ignition Server data, update Ignition Server firmware, set up HA pairs, and bind ports for the RADIUS and SOAP services, and edit the Ignition Server administrator account.
- **Node Configuration and Maintenance:** power down, reboot, or reinitialize a node; view the operational status of a node; set up network ports; set DNS and other network settings; set up logging; and view logs of a node.

To open the Configuration view, click Configuration at the upper left corner of Dashboard. This view is composed of:

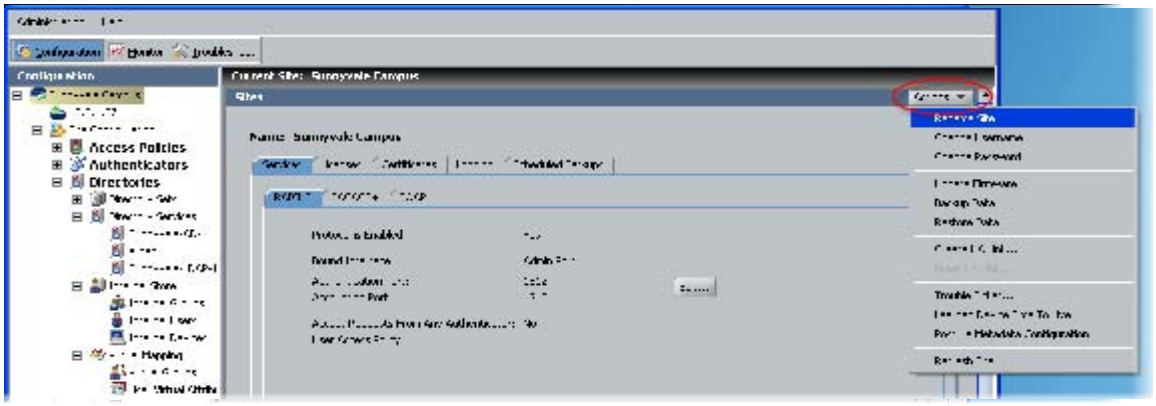
- The **Configuration Hierarchy** navigation panel on the left. Here you specify and organize the nodes in your site.
  - When you select a **site** in the navigation panel, Ignition Server displays the statistics and commands you need to manage the site. See Managing a Site, below.
  - When you select a **node** in the navigation panel, Ignition Server displays statistics and commands you need to manage the node. See "Managing a Node" on page 54.
- A pull-down **Actions** menu whose commands operate on the selected site or node.
- A large editing panel on the right for viewing and editing system settings. Most often, this panel shows **Sites** or **Nodes** panel.

Managing a Site

Dashboard's **Sites** panel lets you manage your Ignition Server *site* and its *nodes*. Depending on your configuration, a *site* consists of a single node (one Ignition Server) or a pair of nodes (a high availability pair of Ignition Servers). If you manage many Ignition Server sites, you can connect your Dashboard session to a group of sites and quickly switch back and forth among them. (See "Managing Multiple Ignition Server Sites" on page 38.)

For information on managing sites generally, continue reading below. For information on managing high availability sites, see Appendix B, "High Availability Configuration".

Figure 7 Actions Menu for a Site



To manage your Ignition Server site:

1. In Dashboard's **Configuration Hierarchy** tree, click on the name of your site. (This is the name that appears at the top of the tree.)

2. Do one of:
   - Select a command from the **Actions** menu; or
   - Navigate the tabs of the **Sites panel** to view or edit your settings.

Site Actions

The **Actions** menu for a **Site** contains the following commands:

- **Rename Site**: See “Renaming an Ignition Server Site” on page 48.
- **Change Username**: See “Changing the Ignition Server Administrator Login Name” on page 49.
- **Change Password**: See “Setting the Ignition Server Administrator Password” on page 49.
- **Update Firmware**: See “Loading a Firmware Image or Package” on page 403.
- **Backup Data**: See “Creating a Backup” on page 395.
- **Restore Data**: See “Restoring from a Backup File” on page 398.
- **Create HA Link**: See “Run the HA Wizard” on page 374.
- **Break HA Link**: See “Breaking an HA Pair Using Dashboard” on page 387.
- **Trouble Ticket**: See “Trouble Ticket Generation” on page 451.
- **Learned Time to Live (TTL)**: See “Setting TTL for Windows Machine Authentication” on page 305.
- **Refresh Site**: Reloads all site data into Ignition Dashboard.

Renaming an Ignition Server Site

To change the name of an Ignition Server site,

1. In Dashboard’s **Configuration Hierarchy** tree, click on the name of your site. This is the name that appears at the top of the tree, “Site 0” by default.
2. Choose **Actions → Rename Site....** The **Rename Site** dialog box appears.
3. Ignition Server displays the name for the selected site. Enter the new name for the site.
4. Click **OK** to apply your changes.

Ignition Server changes the name of the selected site accordingly.

Related topic: “Setting Up A Site Group” on page 38

Changing the Ignition Server Administrator Login Name

The default administrator login name is *admin*. To change the Ignition Server administrator login name,

1. In Dashboard's **Configuration Hierarchy** tree, click on the name of your site. This is the name that appears at the top of the tree, "Site 0" by default.
2. From a selected site, right-click your mouse to bring up the selection box**.** Select **Change Username...** The **Change Username** dialog box displays the current user name.



Note: If you are managing a high-availability pair of Ignition Servers, this user name applies to both nodes in the pair.

3. Enter the administrator password.
4. Enter the new user name.
5. Click **OK**.

Ignition Server changes the administrator login name accordingly.

Setting the Ignition Server Administrator Password

The default password is *admin*. To change the administrator password:

1. In Dashboard's **Configuration Hierarchy** tree, click on the name of your site. This is the name that appears at the top of the tree, "Site 0" by default.
2. Choose **Actions → Change Password...**
3. Enter the existing password in the **Old Password** field, enter the new password twice, and click **OK**.



If you are managing a high-availability pair of Ignition Servers, this administrator password applies to both nodes in the pair.

Password Guidelines Avaya recommends the following guidelines for choosing a password:

- Use a minimum of 8 characters
- Include at least one capital letter
- Include at least one number
- Include at least one special character from the set: ~, !, @, #, $, %, ^, &, *, (,), +, >, <, ?, /, \, |, and =.

Enter the **Old Password** and type the **New Password**. Confirm your new password by entering it again in the **Confirm New Password** text area. Click **OK** to activate the new password for the administrator's login.

Managing Ignition Server Services

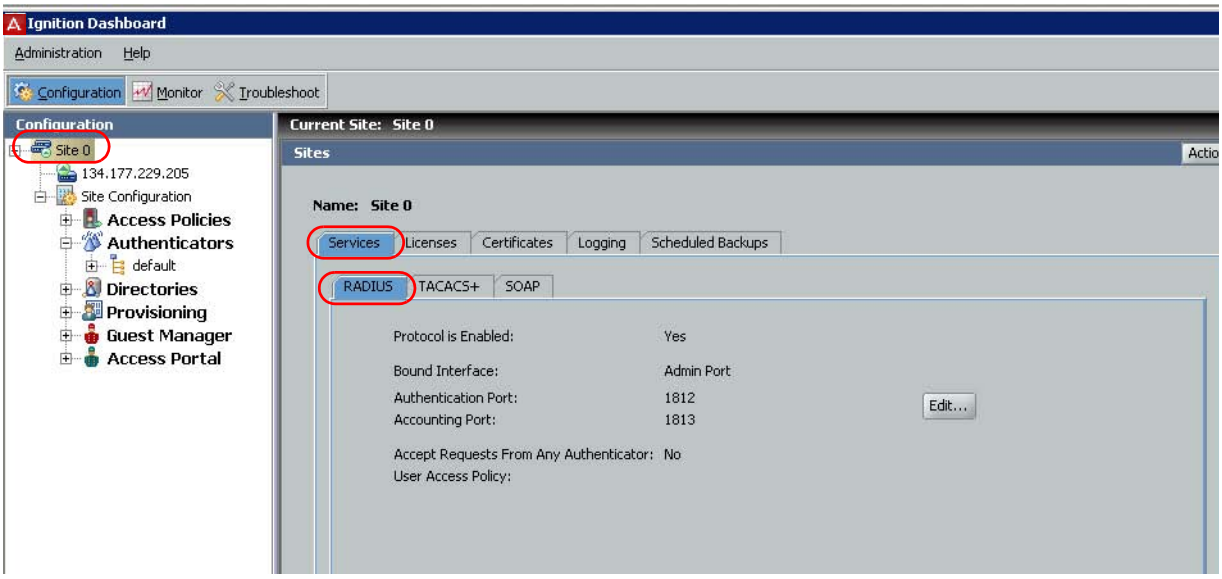
The the **Services** tab in the **Sites** panel allows you to set up the RADIUS service and SOAP service. See the following sections:

- “Setting up Ignition Server’s RADIUS Service” on page 50
- “Setting Up Ignition Server’s SOAP Service” on page 52

Setting up Ignition Server’s RADIUS Service

The Ignition Server RADIUS service handles authentication traffic with supplicants and authenticators. You can bind the Ignition Server RADIUS service to a physical Ethernet port on the Ignition Server (the Admin port or Service Port A), or you can bind it to an Ignition Server VIP (VIPs are explained in “Managing Virtual Interfaces (VIPs)” on page 384). Use the RADIUS tab to bind the RADIUS service and set its port numbers.

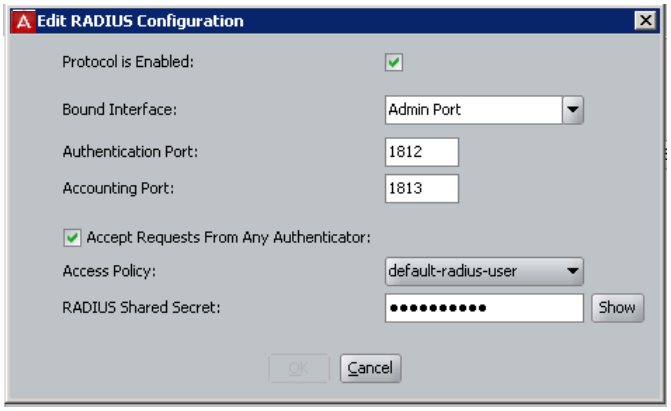
Figure 8 Default Settings for the RADIUS Protocol



Editing RADIUS Communication Settings

1. In the Dashboard main window, in the Configuration Hierarchy panel, click the name of your site (by default, “Site 0”).
2. In the Sites panel, click the **Services** tab and click the **RADIUS** tab.
3. Click the **Edit** button in the RADIUS tab. The **Edit RADIUS Configuration** dialog box appears:

Figure 9 Edit RADIUS Configuration



4. Edit as necessary:
   - **Protocol is Enabled**: Make sure this checkbox is ticked.
   - **Bound Interface:** From the drop down list, choose the Ignition Server Ethernet interface handling the RADIUS traffic. You can bind RADIUS to any port on the Ignition Server. If you are running an HA pair of Ignition Servers, you can choose to bind RADIUS to a VIP interface. The VIP names are also listed in the drop down list. See "Managing Virtual Interfaces (VIPs)" on page 384 for details on using VIPs.
   - **Authentication Port:** Enter the UDP port number that should receive RADIUS authentication requests. The default RADIUS authentication port is 1812. If your installation uses older network equipment, you might have to set the Ignition Server RADIUS authentication port to 1645.
   - **Accounting Port:** Enter the UDP port number that should receive RADIUS accounting messages. The default accounting port is 1813. See "Access Log: RADIUS and TACACS+ Accounting" on page 428.
5. Ignition Server enables the **OK** button. Click **OK** to apply your changes to the RADIUS service.

Important! If your site uses Ignition Server Guest Manager, note the following: Guest Manager uses RADIUS to authenticate provisioner users against the Ignition Server. For Guest Manager to work, your network must allow RADIUS (UDP) traffic to travel between Guest Manager and the Ignition Server. If firewalls exist between Guest Manager and Ignition Server's RADIUS port, make sure they allow this traffic.

The other fields in this window (**Accept Requests From Any Authenticator** and others) allow you to create a global authenticator. See "RADIUS Global Authenticator" on page 104.

Setting Up Ignition Server's SOAP Service

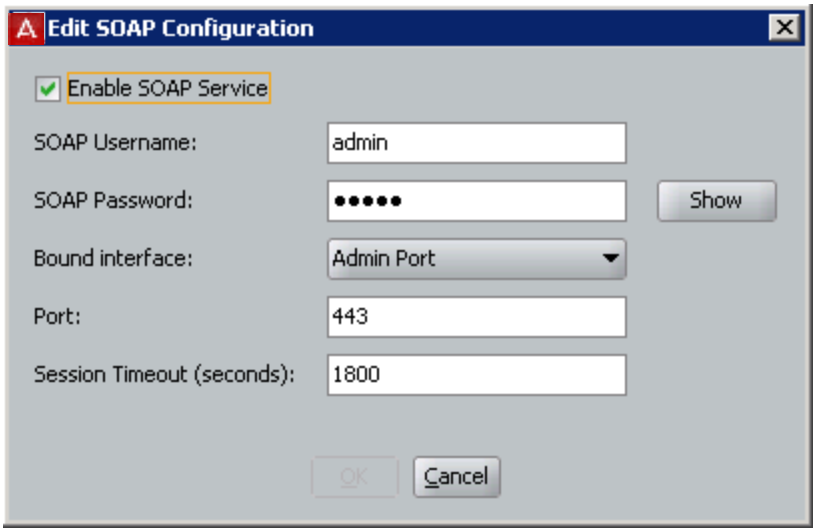
The Ignition Server SOAP service allows Avaya Identity Engines Ignition Server Guest Manager and other API client programs to interact with Ignition Server to perform administration and other tasks. By default, the Ignition Server SOAP service is disabled.

The section that follows explains how to set up the SOAP API, but if you are setting up your Guest Manager connection, Avaya recommends that you instead follow the instructions in "Set up Connection to the Ignition Server" in the *Guest Manager Configuration Guide*, because that document explains the additional tasks you must perform in Guest Manager.

To set up the SOAP service:

1. In Dashboard's Configuration Hierarchy panel, click the name of your site (by default, "Site 0").
2. In the Sites panel, click the **Services** tab and click the **SOAP** tab. (*If there is no SOAP tab*, you must install the SOAP feature license. See "Installing an Ignition Server License" on page 69.)
3. Click on the **Edit** button in the SOAP tab. The **Edit SOAP Configuration** window appears.

Figure 10 Configuring the Settings of SOAP Service



4. Set the SOAP connection parameters:
   - **Enable SOAP Service**: Tick this checkbox to make the SOAP API service available.
   - **SOAP Username**: This is the login name that Guest Manager and other SOAP API clients use to connect to the service. This is not an

account in the internal store; by typing a name and password here, you are creating the SOAP user account. Do not use spaces. Type only letters and numbers.

- **SOAP Password**: Password that SOAP user account uses to connect.
- **Bound Interface:** From the drop down list, choose the Ignition Server Ethernet interface that is intended to handle SOAP traffic. You can bind the SOAP service to any port on the Ignition Server. If you are running an HA pair of Ignition Servers, you can choose to bind to a VIP interface. The VIP names are also listed in the drop down list. See "Managing Virtual Interfaces (VIPs)" on page 384 for details on using VIPs.
- **Port:** Enter the port number to which API clients should connect. Traffic through this port is HTTPS traffic.
- **Session Timeout:** Enter the period, in seconds, after which the SOAP API connection is automatically reset. This timeout ensures that unused sessions are closed at the expiration of the timeout period, but it does not cause Guest Manager to become disconnected since Guest Manager automatically reconnects.

  Important! Set the SOAP **Session Timeout** to a period of 180 seconds or longer. Setting it to a shorter period can result in Guest Manager being unable to load large sets of users.

5. Click **OK** to apply your changes.
6. Install the SOAP certificate *on the Ignition Server* as explained in "Assigning the SOAP Service Certificate" on page 82.
7. Perform SOAP setup steps in Guest Manager:
   - Install a copy of the SOAP certificate *in Guest Manager* as explained in the *Guest Manager Configuration Guide*, in the section, "Installing a SOAP Certificate."
   - Make SOAP and RADIUS settings in Guest Manager as explained in the *Guest Manager Configuration Guide*, in the sections, "Make SOAP Connection Settings" and "Make RADIUS Connection Settings."
8. If Guest Manager is running, restart Guest Manager's application server (usually Tomcat) before you try to connect Guest Manager to the SOAP service. This allows the new SOAP settings to take effect.

Resetting the SOAP Password

Applications like Guest Manager must present a valid SOAP password in order to connect to the Ignition Server. You can reset the password as follows:

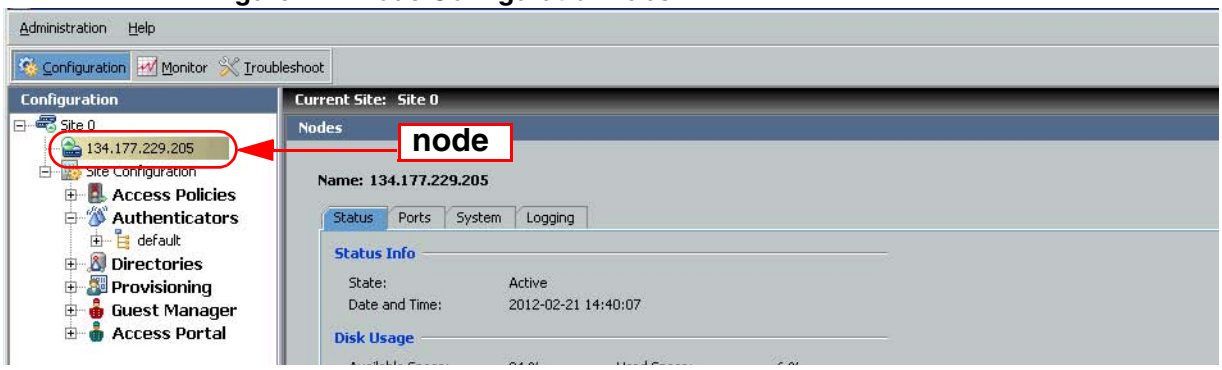
1. In Dashboard's Configuration Hierarchy panel, click the name of your site (by default, "Site 0").

2. In the Sites panel, click the **Services** tab and click the **SOAP** tab.
3. Click on the **Edit** button in the SOAP tab. The **Edit SOAP Configuration** window appears.
4. In the **SOAP Password** field, type the new password.
5. Retype the password in the **Confirm Password** field.
6. Click **OK** to apply your changes.
7. In Guest Manager, type the new password. To do this: Log in to the Guest Manager administrator application. Click **Manage Appliance**. If connected, click **Disconnect**. Click **Manage Appliance** again. Type the SOAP user name and the new password. Click **Connect**.

Managing a Node

A *node* is an individual Ignition Server. In the **Configuration Hierarchy** panel of Dashboard, you can find your node listed under its name or IP address. When you click on a node in the Configuration Hierarchy panel, Dashboard displays the **Node Configuration** panel and, in the **Actions** menu, makes available the commands that operate on nodes.

Figure 11 Node Configuration Tabs



Actions Menu for a Node

The **Actions** menu for a **node** contains the Reboot, Power Down, Reinitialize, View Logs, and Rename Node commands, each of which is explained below.

Rebooting a Node

When you reboot a node, Ignition Server disconnects the Ignition Dashboard from the node and reboots it. In order to reboot a node,

1. In Dashboard's **Configuration Hierarchy** panel, click the name or IP address of your node.
2. Right-click on the selected node and choose **Reboot**. Or, select **Actions → Reboot** command.
3. The Reboot Confirmation window appears, requiring you to confirm your action.

4. Click **Yes**. Ignition Server disconnects the node and reboots the Ignition Server.
5. Wait for a few minutes, and then log in to the Ignition Server.

Powering Down a Standalone Node

Ignition Server allows you to turn off the power to a node only when the selected node is a standalone node. In order to power down a standalone node,

1. In Dashboard's **Configuration Hierarchy** panel, click the name or IP address of your node.
2. Right-Click on the selected node and choose **Power Down**. Or, select **Actions → Power Down** command**.**
3. The Power Down Confirmation window appears, requiring you to confirm your action.
4. Click **Yes**. Ignition Server disconnects the node and switches off the Ignition Server.

To start the Ignition Server again, press the power switch on the back of the Ignition Server.

Reinitializing the Ignition Server from Dashboard

Ignition Server allows you to reinitialize a node only when the selected node is a standalone node. When you reinitialize a standalone node using the Ignition Dashboard, Ignition Server resets the node to its factory settings. All data and configuration settings are deleted.

Note! You can also reinitialize from the front panel.

In order to reinitialize a standalone node from Dashboard:

1. Make a note of the IP address of the Admin port, and also write down any other settings you plan to restore after the reinitialization.
2. If you want to retain your Ignition Server licenses, make a license backup file:
   - In Dashboard's Configuration hierarchy tree, click the site name.
   - Click the **Licenses** tab and click **Export All**. Choose a file name and path, click **Save**, and note the file name so you can import the licenses later.
3. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
4. Right-click on the selected node and choose **Reinitialize**. Or, select **Actions → Reinitialize** command**.**

5. A confirmation window appears, requiring you to confirm your action.
6. Click **Yes** to proceed with the reinitialization. Ignition Server resets the selected node to its factory settings.
7. After the Ignition Server has rebooted, set its IP address.
8. Use Dashboard to log in to the Ignition Server, and restore the licenses from the license file you saved earlier. See "Installing an Ignition Server License" on page 69.

Viewing Logs for a Node

To view the logs of your Ignition Server node:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Right-click on the node and choose **View Logs**. Or, select **Actions → View Logs** command. This opens the Monitor tab of Dashboard and places you in the Log Viewer tab for your node.

Renaming a Node

To rename your Ignition Server node:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Right-click on the node and choose **Rename Node**. Or, select **Actions → Rename Node** command**.**
3. Type a new node name and click OK.

Status Tab

The **Status** tab of the Nodes panel provides a read-only display of the status and usage statistics for the *node* you select in the **Configuration Hierarchy** panel. It lists the information in the following categories:

Status Info

- **State:** indicates whether the node is active or not
- **Date and Time:** displayed and updated every 5 seconds.

Disk Usage The Disk Usage section lists, as percentages, the available and used space on the node. As the number of logs and/or data in the database increases, or as you install additional firmware images, the amount of available space decreases.

Current Configuration

- **System Version:** the version and build number for the firmware on the Ignition Server
- **Last Boot Date:** the last time the node was rebooted.
- **Serial Number**: Unique number that identifies this Ignition Server machine. This is also known as the *Node Id*. The Ignition Server feature licenses are keyed to this number. See "Managing Ignition Server Licenses" on page 65.

Finding the Ignition Server's Serial Number

To find the serial number, also known as the Node Id:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **Status** tab. Look in the Current Configuration section for the **Serial Number**.

System Tab

The **System tab** enables you to specify the DNS servers, set the routing information, and make SNMP, SSH, and SMTP settings for the node.

Viewing Ignition Server's DNS Settings

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **System** tab in the **Node Configuration** panel. Ignition Server displays the DNS settings.

Editing Ignition Server's DNS Settings

DNS settings apply to each Ignition Server individually, even if the Ignition Server is part of an HA pair.



Warning: If your installation uses an Active Directory service, you must specify your DNS server address(es) before you connect Ignition Server to Active Directory.

Make your DNS settings as follows:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. In the **Node Configuration** panel, click the **System** tab and click the **DNS** tab.
3. Click the **Edit** button in the **DNS** section of the **System** tab.
4. Enter the DNS server IP addresses using dotted decimal notation:
   - **Primary IP Address**: Enter the unique IP address of your primary DNS Server.

- **Secondary IP Address**: This entry is optional. Enter the unique IP address of your secondary DNS server.

5. In **Search Domain**, enter the DNS search domains. When entering more than one domain, separate the domain names with a space. When trying to resolve a host name, the Ignition Server searches these domains. Typically this is your organization's domain name, such as, for example, *Avaya.com*.

   Enter no more than six domains, and no more than 1024 characters in the **Search Domain** field.

6. Click **OK** to apply your changes.

Setting the Network Routing Configuration

The **System: Static Routing** tab of the **Node Configuration** panel displays the network routing table. To add a route, see "Adding a Route to Ignition Server's Routing Table" on page 59. To edit a route, see "Editing an Existing Route" on page 59. To delete a route, see "Deleting an Existing Route" on page 59.

Note: The **System: Static Routing** tab of the **Node Configuration** panel also displays the Current System Routing Table.

When routing network traffic, Ignition Server uses the gateway assigned to the closest matching **Destination IP** address set in this table. Typically, you set a general default gateway and then a gateway for each subnet. When a destination IP address matches one you have added to this list, the packet is sent to the corresponding gateway.

If there are no entries in the **Static Routing** configuration table, then, for a given Ethernet interface on the Ignition Server, the only accessible IP addresses are those that share a subnet with that interface.

A more specific IP address entry in the list is applied before a more general version of that IP address. So if the list included both 192.168.1.1 and 192.168.0.0, each with its own gateway, the 192.168.1.1 address would be tested first. If it matched the request, its corresponding gateway would be used. A request from 192.168.1.2 would instead use the gateway given in the routing entry for 192.168.0.0. An entry of 0.0.0.0 with subnet /0 would point to a default route or gateway for all packets whose destination IP address failed to match any other entry in the list.



Important! Before setting up routes in the **Static Routing** configuration table, make sure you have set the IP addresses of the Ignition Server interfaces you plan to use.

Adding a Route to Ignition Server's Routing Table

To add a network route to Ignition Server's **Static Routing** table:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **System** tab and click the **Static Routing** tab.
3. Below the **Static Routing** configuration table, click **Add**. The **Add a Route** window appears.
4. In the **Add a Route** window, add a gateway by entering
   - **Destination IP Address**: A packet whose destination address most closely matches the **Destination IP Address** is directed to the gateway you specify. Enter the unique IP address of the destination.
   - **Subnet Mask**: The bit mask used to interpret the IP address. Use network prefix notation (an integer representing the number of bits in the address to be used in the comparison). Valid entries include numbers between 0 and 32.
   - **Gateway**: The IP address of the next hop (the gateway for the new route)
5. Check the routing information you have entered and click **OK** to add the gateway.

The **Static Routing** configuration table shows the newly added route. Repeat this procedure to add more routes.

Editing an Existing Route

To edit a network route in Ignition Server's **Static Routing** table:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **System** tab and click the **Static Routing** tab.
3. In the **Static Routing** configuration table, highlight the route entry.
4. Click **Edit**. The **Edit a Route** window appears**.**
5. Edit the entries as desired.
6. Click **OK**.

The **Static Routing** configuration list shows the updated entry for the route.

Deleting an Existing Route

To delete an existing network route:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **System** tab and click the **Static Routing** tab.
3. In the **Static Routing** configuration table, highlight the entry to be deleted.
4. Click **Delete**. Because you cannot undo a deletion, Ignition Server displays the Delete Route Confirmation dialog box.
5. Click **OK** to confirm the deletion of the selected route.

Ignition Server deletes the selected route. The **Static Routing** configuration list no longer displays the entry for the deleted route.

SNMP Settings

Ignition Server's SNMP support allows network management tools like Net-SNMP and HP OpenView to query the Ignition Server and retrieve basic system health and configuration information. See "Monitoring Ignition Server via SNMP" on page 418 for instructions.

SSH Settings

You can set up an SSH network port on the Ignition Server and connect to the Ignition Server CLI via SSH. See "Connecting to the CLI Through an SSH Connection" on page 345 for instructions.

SMTP Settings

You can set up Ignition Server to send log alerts via e-mail using an SMTP server on your network. See "Sending Log Messages Via E-Mail" on page 417.

Configuring the Ignition Server's Network Ports

The Ignition Server's Ethernet interfaces include the Admin Port (always enabled), a Service port, the HA port, and optional virtual ports, called VIP ports. The table below explains what sort of traffic each interface can carry.

Ignition Server Ports and Allowed Traffic

Port	Admin / default traffic	Directory traffic	RADIUS traffic	SOAP API traffic	HA Link to another Ignition Server
Admin port	Yes	Yes	Yes	Yes	No
Service port	No	Yes	Yes	Yes	No
HA port	No	Not recommended	Not recommended	Not recommended	Yes

Ignition Server Ports and Allowed Traffic

Port	Admin / default traffic	Directory traffic	RADIUS traffic	SOAP API traffic	HA Link to another Ignition Server
VIP ports (virtual ports in HA)	No	No	Yes	Yes	No

Port Configuration Settings

In the **Nodes** panel, the **Ports** tab allows you to adjust the network settings of each Ethernet interface on the Ignition Server.

The port configuration settings are the following:

- Port Status is the user configuration. (Enabled or Disabled using the GUI).
- Interface Status is the system's (OS) administration status. This is usually the same as the port status (It may take a few seconds for the GUI to update the system).
- Link Status is the Physical Status of the Link.

  In Virtual Manager, the LinkStatus is determined by the Virtual Switch the port is connected to. In most of these connections this link will be up regardless of the physical port it is tied to. Only if the VM-HOST disconnects the port, the LinkStatus will go down. In VM the connection may look like this:

  PhysicalPort <===> Virtual SW (HOST) <===> NIEIS-GuestPort (eth0 or eth1 or eth2. The status of this Link is the LinkStatus in the VM AIEIS.

The following table explains the network interface settings and statistics you can view and set in the Configuration view of Dashboard.

Table 2 Port configuration settings in the Nodes panel

Field Name	Valid Entries	Where Set	Description
Port Enabled	Yes, No	Click **Edit** in the port's tab, and click the **Enabled** checkbox in the Edit Port Configuration dialog.	Indicates whether the administrator has enabled this port. Note that the Admin Port is always enabled.
Link Status	Up, Down		Indicates whether the port is connected to the network. A status of **Up** indicates the port has link-level connectivity with another network device.

Table 2 Port configuration settings in the Nodes panel

Field Name	Valid Entries	Where Set	Description
Interface Status	Enabled, Disabled		Indicates whether the port has been enabled and is connected to the network. If the status displays **Disabled**, check that you have enabled the port (see **Port Enabled**, above) and check your network connections and cables.
IP Address	Any valid IP address	The **IP Address** field in the **Edit Port Configuration** window.	IP address of the interface.
/	Net mask expressed as a bit count	The right-most field in the **Edit Port Configuration** window.	Bit mask used to interpret the IP address.

Configuring the Admin Port

The admin port is always enabled. Initially the address is set during the installation. Admin IP can only be configured manually.



Important! When you change the settings for the Admin Port, Ignition Server logs you out of the Ignition Server. If you change the IP address, make a note of it. When you reconnect Dashboard, enter the new IP address in the **Hostname** field of the Dashboard login dialog.

To change Admin port settings:

In order to change the IP address settings for the Admin Port,

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **Ports** tab and click the **Admin Port** entry.
3. Click the **Edit** button. Ignition Dashboard displays the following version of the **Edit Port Configuration** window.
4. Dashboard displays the current Admin Port IP address; and user can change by entering new IP address and Mask:
5. Click **OK**.
6. Ignition Server displays the **Admin Port Configuration** dialog box to inform you that, after the IP address is updated, you lose connection to the Ignition Server. Select **Yes** to continue with the update or **No** to discard your changes to the settings for the Admin Port. If you select **Yes**, Ignition Server updates the IP address and logs you out of the Ignition Server.

7. Reconnect Dashboard to the Ignition Server: In Dashboard, select **Administration: Login**. In the **Hostname** field of the Dashboard login dialog, enter the *new* IP address or hostname. Enter the admin credentials and click **OK**.

Configuring a Service Port

The Ignition Server has an Ethernet port known as a “service port:”. In the **Ports** tab, when you click **Service Port** and click the **Edit** button, Dashboard displays the Edit Port Configuration window. Here you can enable/disable a service port and set its IP address for the port.

Enabling Service Port

In order to enable the service port and set its IP address, do the following:

1. In Dashboard’s Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **Ports** tab and click the required port tab: **Service Port**.
3. Click the **Edit** button. The **Edit Port Configuration** window appears.
4. By default, the **Enable Port** checkbox is not selected, indicating the port is disabled.
5. To enable the port, tick the **Enable Port** checkbox.
6. In the **IP Address** fields, enter the IP address and subnet mask entries for the port.



Note that, per Ignition Server, no two port IP addresses can be located on the same subnet.

7. Click **OK** to apply your changes.
8. Assign services to the port or ports you enabled. See “Managing Ignition Server Services” on page 50 for instructions.

Ignition Dashboard updates the display in the **Node Configuration** panel indicating whether the selected port is enabled/disabled, the connection link is up/down (if you enabled it), and the current IP address (and subnet mask) for the port.

Configuring the HA Port

The HA Port tab of the Node Configuration panel displays the IP address of the HA port of the selected node. Avaya recommends that you run the HA Configuration Wizard, to set the IP address and other settings of your HA port. You can also set the HA port IP address by following the instructions detailed in the section “Enabling Service Port” on page 64.

Enabling the HA Port

In order to enable the HA port and set its IP address, do the following:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **Ports** tab, and click the **HA Port** tab.
3. Click the **Edit** button. The **Edit Port Configuration** window appears.
4. By default, the **Enable Port** checkbox is not selected, indicating the port is disabled.
5. To enable the port, tick the **Enable Port** checkbox.
6. In the **IP Address** fields, enter the IP address and subnet mask entries for the port.

Note that, per Ignition Server, no two port IP addresses can be located on the same subnet.

7. Click **OK** to apply your changes.

Important! When you run the HA Configuration Wizard, the IP addresses you specify in the wizard overwrites the existing IP address settings of the HA ports. Nonetheless, you must enable the HA port before you run the Wizard.

Managing Ignition Server Licenses

Your Ignition Server requires a license (called the "FEATURE_BASE" license) to run, and if you wish to use optional features of Ignition Server, then you must install the licenses for those features. Features that require a license include the SOAP service, the NAP posture checking features, the Ignition Server TACACS+ service, Ignition reports, and Access Portal. This section explains how to install and manage licenses.

- To check your licenses, see "Checking Your Ignition Server Licences" on page 65.
- To get a license, see "Getting an Ignition Server License" on page 69.
- To install a license, see "Installing an Ignition Server License" on page 69.
- To replace a license, see "Replacing an Ignition Server License" on page 69.
- To back up a license, see "Making a Backup Copy of Your Ignition Licenses" on page 70.
- To transfer a license, see "Transferring a License to a Different Ignition Server" on page 70.

Checking Your Ignition Server Licences

Whenever your login to the Dashboard, a license validity check activates to check for validation and expiry date of the license.

If the license is not valid or is expired, a message displays to prompt you to enter a valid license. This check includes the validity of a 30-day grace period that may be provided during upgrading.
As shown in the following figure, if the grace period expires the Ignition Server stops processing the user authentication requests. The user can still launch the dashboard UI, but will not be able to configure anything on the system.

Figure 12 License expiry

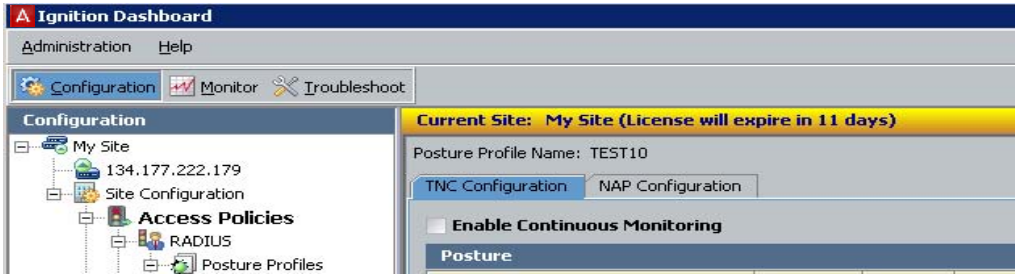


Figure 13 GRACE PERIOD



To check your Ignition Server licenses:

1. In Dashboard, click your **Site** in the Configuration Hierarchy.
2. Click the **Licenses** tab. The **Licenses** list shows the installed licenses.
3. Click on a license to see its **License Details**. This field shows the expiration date and other attributes. The **Valid From** and **Valid Until** field show the start and end dates, respectively, of the license validity period. The **Node Ids** field lists the serial numbers of the Ignition Servers on which this license is valid. The **License Serial Number** is a unique number that identifies this license. The **License Version** is displayed.
4. If you want to e-mail or copy the **License Details**, click the **Copy to Clipboard** button, and paste it into your e-mail or other application.

The available types of licenses are:

- **Guest Manager**, which allows you to use the SOAP API of the Ignition Server in order to run Avaya Identity Engines Ignition Server Guest Manager.
- **NAP Posture**, which allows you to add Microsoft NAP based client posture checking to your Ignition Server policies.
- **Base License**, which allows system to use RADIUS, 802.1x, Active Directory, LDAP and RSA Integration Modules
- **TACACS+**, which allows you to use Ignition Server TACACS functionality
- **Ignition Reports**, which allows you to use Ignition Analytics to present Ignition Server network authorization and authentication information in a variety of summary and detail reports in the areas of audit, compliance, security and usage.
- **Access Portal**, which allows guests with non-802.1X compatible equipment to authenticate and connect to the network in your organization **Note**: The Access Portal license must match the level of the Identity Engines Server base license: LARGE, SMALL, or LITE.

Seat limit enforcements

Three different levels of ignition server based on the number of authenticators allowed are supported as follows.

- LARGE - Unlimited authenticators
- SMALL - 20 authenticators
- LITE - 5 authenticators

Seat limit enforcement follows in the following manner:

While adding the authenticators from the Dashboard, the number of authenticators configured are compared to the number as allowed in the license installed. You can still add newer authenticators if they exceed the license limit, but they'll automatically set to 'disabled' state. The user can then choose and enable the required authenticators up to the license limit.

If you have already configured 'X' number of authenticators and then try to install the new license, the enforcement check compares the seat limit with that of the number of authenticators enabled. If the seat limit is lower all the authenticators are marked as disabled. You are then notified to selectively choose the authenticators as permissible by the license limit.

During upgrade, if the number of authenticators added are more than the limit as permitted by the license, all the authenticators are marked as disabled and a warning message displays. You can then selectively choose which authenticators to enable as per the seat limit.

Similar behavior is expected during the restore process. If the seat limit in the License is less than the number of enabled authenticators in the backup configuration, all the authenticators would be marked as disabled. Users can selectively enable the authenticators.

Using the NAP license

The license “FEATURE_NAP” introduced in release 7.0 controls the Microsoft NAP Posture checking functionality introduced in this release.

Two levels of FEATURE_NAP are available, SMALL and LARGE. The SMALL version of NAP license is intended for smaller deployments while the LARGE is intended for large enterprise networks.

Note: The type of NAP license (SMALL or LARGE) should match with the type of BASE license installed on the system. A SMALL NAP license cannot be installed on a LARGE BASE system.

An error message displays if you already have an existing Radius Authorization policy with a NAP remediation profile and have selected the Posture Profile after the NAP license expiration. Once you click **Ok**, both the NAP Compliant and NAP Non-Compliant conditions are reset.

If you set any NAP specific remediation condition for a Radius authorization policy an error displays. Alternatively if you are in editing mode, for example as in editing an Authorization Policy or Posture Profile, and the NAP license expires on the fly, an error message displays and the dialog boxes being edited are closed. After that you will not be able to launch an edit dialog for NAP configuration.

If the NAP license expires, NAP based posture checking will fail. Any client authorizations that involve NAP checks will be rejected and the following error message will display: “NAP license expired”. If you select a Posture Profile, after the NAP license expires, the NAP Configuration tab will no longer be accessible.

Getting an Ignition Server License

To get an Ignition Server license:

1. Contact the Avaya customer support team to get your license(s). See “Customer service” on page 21.
2. When you talk to your support representative, provide the serial number of each Ignition Server in your installation. (See “Finding the Ignition Server’s Serial Number” on page 57.) Be prepared to tell the representative which feature(s) you want to use: SOAP and/or posture checking.

Your license comes in the form of an e-mail or text file. The actual license starts with “BEGIN IGNITION LICENSE CERTIFICATE” and ends with “END IGNITION LICENSE CERTIFICATE”.

Installing an Ignition Server License

To install an Ignition Server license:

1. If you do not have your new license, get it now as explained in “Getting an Ignition Server License” on page 69.
2. In Dashboard’s **Configuration** view, click your **Site** in the Configuration Hierarchy.
3. Click the **Licenses** tab.
4. Click **Install**.
5. Find the license you received from support and open it in your e-mail tool or text editor. Highlight and copy the text of your license. Copy the whole license including “BEGIN IGNITION LICENSE CERTIFICATE” and “END IGNITION LICENSE CERTIFICATE”.
6. Return to the License Installation window of Dashboard and click **Paste** to paste the license text there. Click **OK**.

Replacing an Ignition Server License



Warning! Before you delete your old license, make sure you have obtained its replacement. If you do not have your new license, get it now as explained in “Getting an Ignition Server License” on page 69.

To replace a license:

1. In Dashboard, click your **Site** in the Configuration Hierarchy.
2. Click the **Licenses** tab.
3. Click on the license and click **Delete**.

4. Install the replacement as shown in “Installing an Ignition Server License” on page 69.

Making a Backup Copy of Your Ignition Licenses

You can make a backup copy of your installed licenses. This is useful, for example, if you reinitialize the Ignition Server.

1. In Dashboard, click your **Site** in the Configuration Hierarchy.
2. Click the **Licenses** tab.
3. Click **Export All** and specify a file name for the licenses. The licenses are saved to a single file. You can later reinstall these licenses on the same Ignition Server, if it has been reinitialized or if the licenses have been deleted from the Licenses tab.

Transferring a License to a Different Ignition Server

You cannot do this. Instead, contact the Avaya customer support team to get a new license for the new Ignition Server. See “Customer service” on page 21.

Troubleshoot Tab

The **Troubleshoot** tab of Dashboard allows you to perform these simple network tests and analysis:

- Running a Ping Test
- Running a Packet Capture

To use these tabs:

1. Click **Troubleshoot** at the top of the Dashboard window.
2. In the hierarchy tree, click the IP address or name of your node.
3. Click the Network tab.

Running a Ping Test

The **Ping Test** tab enables you to check whether a device such as a router, switch, or directory server is reachable.

Configuration Use the Configuration section to provide the details of the ping test you want to execute.

- **Target IP/Hostname:** Enter the IP address or host name of the device you are attempting to reach.
- **Number of Packets:** Specify the number of packets to be sent to the IP address you want to ping.
- **Timeout:** Specify the number of seconds to wait between packets.

When you have entered the above information, click **Start**. Ignition Server pings the specified device. The **Stop** button allows you to abort the test before completion. See the **Results** section for the outcome of the ping test.

Running a Packet Capture

The Packet Capture displays the results of sniffer traces on the ports of the Ignition Server for troubleshooting. You can use this information to debug problems related to network traffic.

To perform a packet capture

1. Click **Troubleshoot** at the top of the Dashboard window.
2. In the hierarchy tree, click the IP address or name of your node.
3. Click the **Network** tab.
4. In the **Packet Capture** section, use the **Port** drop-down list to pick the interface whose traffic you want to capture.
5. In the **Filter Expression** field, specify the filter you want to apply, using the *tcpdump* syntax. (Refer to the man page for *tcpdump* for an explanation.)
6. The **Save Packets To** field shows the path and file name of the pcap file to be saved. The default file name is `PacketCapture.pcap`. Click the **Browse** button to specify the destination location. In the Save Captured Packets window, navigate to find the desired directory and then type the desired **File Name**. Click **Save** to accept your path name. (This does *not* save the pcap file.)
7. Specify the size of the capture in the **Number of Packets to Capture** field. By default, Ignition Server captures 100 packets. Ignition Server limits the capture to 10,000 packets. Note that if you set a high limit here and you apply no filter, the saved file might be very large.
8. Click the **Start** button to launch the capture. The capture stops and saves the file when the specified **Number of Packets to Capture** threshold is reached. If you want to stop it sooner and save the file, click the **Stop** button.

By default, Ignition Server saves the `PacketCapture.pcap` file in the Ignition Server administrator directory on your computer. On Microsoft Windows computers this is:

```
...Avaya\user\admin
```

Managing Certificates

This chapter explains how to install and manage digital certificates on the Avaya Identity Engines Ignition Server. Ignition Server requires certificates to secure communications between the Ignition Server and Dashboard, and among the Ignition Server, supplicants, and authenticators.



Important! Your default installation includes sample certificate files that allow you to use the system without immediately installing your own certificates, but Avaya strongly recommends that you install your own certificates before deploying Ignition Server on a production network. The sections that follow explain how to manage and replace your certificates.

Contents of This Chapter

- "Required Types of Certificates" on page 73
- "Sample Certificates" on page 74
- "Format of Certificate Files" on page 74
- "Certificates Tab" on page 74
- "Getting a New Certificate" on page 75
- "Admin Certificate" on page 78
- "Assigning the SOAP Service Certificate" on page 82
- "Assigning Protocol Credential Certificates" on page 82
- "Installing Protocol Root Certificates" on page 83
- "Adding a Certificate Revocation List URL" on page 84
- "Viewing a Certificate" on page 86
- "Deleting a Certificate or Certificate Request" on page 86
- "Viewing an Existing Certificate Request" on page 87

Required Types of Certificates

The required types of certificates are:

- The **admin certificate** secures Ignition Dashboard-to-Ignition Server communications. See "Admin Certificate" on page 78.

- The **SOAP service certificate** secures Guest Manager-to-Ignition Server communications. See “Assigning the SOAP Service Certificate” on page 82.
- **Protocol credential certificates** (or “**tunnel certificates**”) are used in Ignition Server authentication policies to secure communications between Ignition Server and authenticating supplicants. See “Assigning Protocol Credential Certificates” on page 82.
- **Protocol root certificates** are used to verify supplicants’ certificates during EAP-TLS and PEAP/EAP-TLS authentication. See “Installing Protocol Root Certificates” on page 83.

Sample Certificates

Avaya provides sample certificates with your Ignition Server. The purpose of sample certificates is to get your installation up and running, even if you have not yet generated your own certificates. You should generate and install your own certificates at your earliest convenience.

These sample certificates are located in the directory where you have installed Identity Engines (`...\Avaya\security\cacert`).

Currently, Avaya provides two default certificates with the Ignition Server, `default_ui_cert` and `default_tunnel_cert`. You can view these default certificates in the Certificates tab of Dashboard’s Sites panel.

Format of Certificate Files

For use in Ignition Server, each certificate must be PEM formatted and saved in a text file. In particular:

- The certificate file must contain one and only one PEM-encoded certificate.
- In the file, the certificate starts with the line, `“-----BEGIN CERTIFICATE-----”` and ends with the line, `“-----END CERTIFICATE-----”`.

Certificates Tab

Use the Certificates tab to import and manage certificates in Ignition Server. To open it, go to the top of Dashboard’s navigation panel and click on the name of your site. In Dashboard’s Sites panel, click the Certificates tab.

The Certificates tab is organized in sub-tabs:

- The **Certificates** sub-tab lists all certificates that have been imported into Ignition Server. These can be used to secure the Dashboard-Ignition Server connection (see “Admin Certificate” on page 78) or to secure

authentication transactions (see “Assigning Protocol Credential Certificates” on page 82).

- The **Certificate Requests** tab is used to generate certificate requests. See “Getting a New Certificate” on page 75.
- The **Protocol Root Certificates** tab lists the certificates used to validate EAP-TLS supplicants and PEAP/EAP-TLS supplicants. See “Installing Protocol Root Certificates” on page 83.
- The **Certificate Revocation List** tab lists the URLs used to check certificate revocation status when validating EAP-TLS supplicants and PEAP/EAP-TLS supplicants. See “Adding a Certificate Revocation List URL” on page 84.

Not included in the **Certificates** tab is the Root Certificates window:

- The **Root Certificates** window (open it by selecting the **Administration: Root Certificates** command) lists the certificate Dashboard uses to validate the Ignition Server’s admin certificate. If your installation of Dashboard is used to connect to multiple Ignition Servers, then you might have more than one certificate listed here. See “Installing Dashboard’s Copy of the Admin Certificate” on page 79.

Getting a New Certificate

Getting a new certificate consists of three tasks:

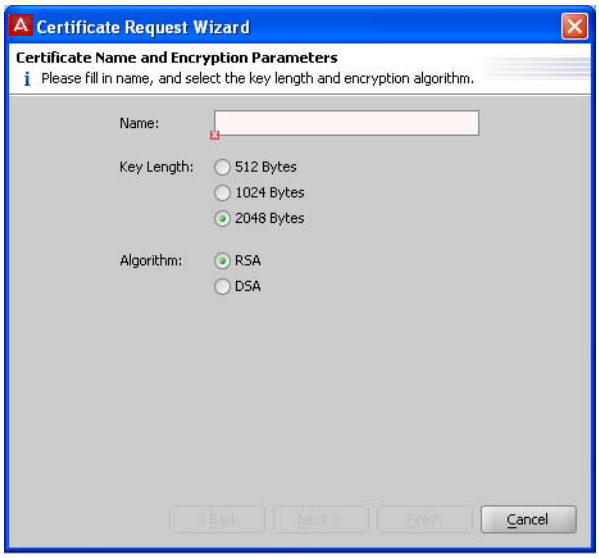
- “Create the Certificate Request” on page 75
- “Import the Certificate” on page 78
- “Assign the Certificate for Use in Ignition Server” on page 78

Create the Certificate Request

To create the new certificate request:

1. At the top of Dashboard’s navigation panel, click on the name of your site.
2. In Dashboard’s Sites panel, click the **Certificates** tab and click the **Certificate Requests** tab. From the **Certificate Requests** tab, click **New...**. The Certificate Manager starts the Certificate Request Wizard.

Figure 14 Certificate Request Wizard



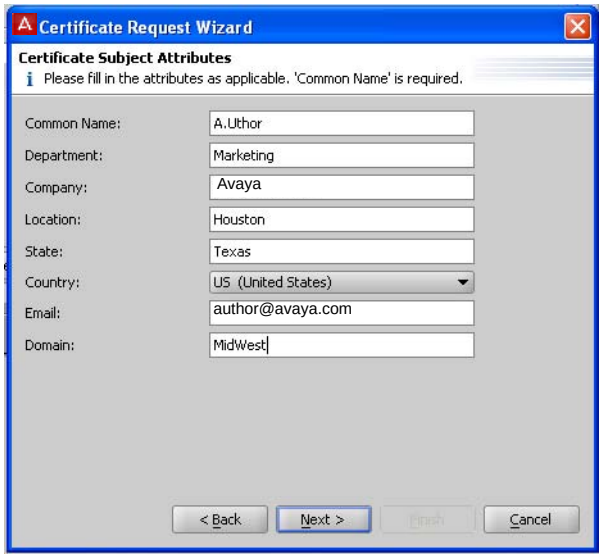
3. Specify the type of certificate you want to request.
   - In the **Name** text box, enter a descriptive name that reflects how you plan to use this certificate when it is issued. For example, “Ignition Server Administrator” or “Company ABC RADIUS Server #1.”
   - Specify the desired **Key Length** for this certificate (2048 is the default).
   - Specify which **Algorithm** you want to use: the RSA algorithm or the DSA algorithm to generate the key8.

Tip: If this certificate is to be used as your Ignition Server **admin certificate**, it must use an RSA key; you cannot use a certificate based on a DSA key for an Ignition Server **admin certificate**.

Tip: If this certificate is to be used as a tunnel certificate that supports Windows XP clients, please observe the limitations explained in “Factors That Limit Your Choice of a Protocol Credential Certificate” on page 218.

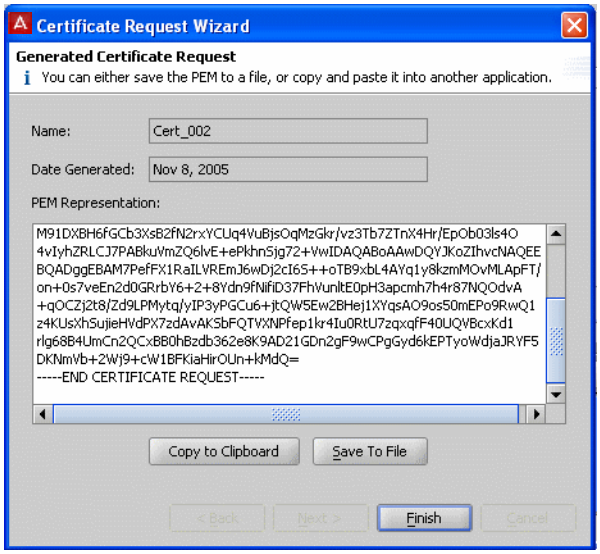
4. Click **Next**. The Wizard displays the **Certificate Subject Attributes** screen. The **Common Name** is required.

Figure 15 Completed Certificate Request Data



5. Click **Next**. Ignition Server generates the certificate request and displays the Generated Certificate Request screen.

Figure 16 Generated Certificate Request



6. Do one of the following:
   - Click the **Copy to Clipboard** button to make a copy of the request. Paste the request into an e-mail message or file and send it to your CA to request the certificate.
   - Click the **Save to File** button to save the request. Send the file to your CA to request the certificate.
7. Click **Finish** to close the Certificate Request Wizard.

After the CA responds with the requested certificate, follow the steps in the next section.

Import the Certificate

Procedure After you have received the certificate you requested in Step 7, or if you have a certificate ready for import, you can import the certificate as explained below:

1. At the top of Dashboard's navigation panel, click the name of your site.
2. In Dashboard's Sites panel, click the **Certificates** tab and view the **Certificates** sub-tab.
3. Click the **Import Certificates** button.
4. Navigate to find your certificate. Make sure that:
   - Your certificate meets the certificate file requirements listed in "Format of Certificate Files" on page 74.
5. Click **Open**. Ignition Server imports the certificate to the Ignition Server keystore on the Ignition Server. Next, continue with "Assign the Certificate for Use in Ignition Server" on page 78.



Tip. From the CA you should have also received a copy of the CA root certificate. Keep copies of both your certificate and the CA root certificate.

Assign the Certificate for Use in Ignition Server

Choose the appropriate procedure, based on the role the certificate is to play. If the certificate is for:

- Dashboard-Server communications, see "Replacing the Admin Certificate" on page 80 or
- Guest Manager-Server communications, see "Assigning the SOAP Service Certificate" on page 82 or
- User authentications, see "Assigning Protocol Credential Certificates" on page 82.

Admin Certificate

Ignition Dashboard requires a copy of the Ignition Server's *admin certificate* in order to communicate securely with the Ignition Server. Dashboard cannot connect to the Ignition Server without this certificate.

The **admin certificate** is installed on the Ignition Server, and the Ignition Server presents it to Dashboard at login time. Dashboard verifies the admin certificate and connects only if the verification succeeds.

Ignition Server checks for the expiry and revocation of the admin certificate every twenty-four hours. If the certificate expires soon, Ignition Server logs a warning message to Security and Audit channels.

Avaya provides a default admin certificate called the **default_ui_cert**. Please replace the default certificate as soon as possible after installing Ignition Server. See “Replacing the Admin Certificate” on page 80 for instructions.

Note! Ignition Server uses two names to refer to the admin certificate. In the **Certificates** tab of Ignition Dashboard (click your site name in the **Configuration** tree; click **Certificates**, and click the **Certificates** sub-tab), you see the certificate labelled in the **Bound to Services** column as the “**UI Port Cert**” instead of the usual “**admin certificate**.”

Installing Dashboard’s Copy of the Admin Certificate

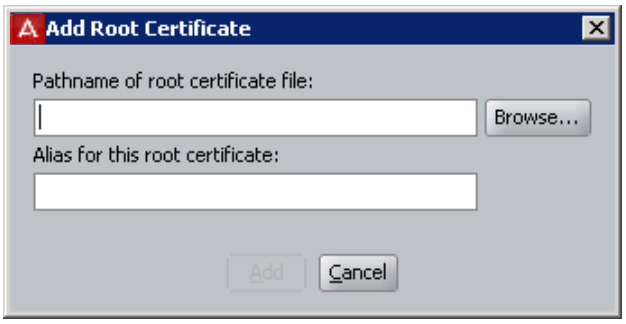
The instructions below explain how to add a copy of the Ignition Server’s admin certificate to Ignition Dashboard. These instructions assume the admin certificate is already installed on the Ignition Server. If you want to replace the admin certificate both in Dashboard and the Ignition Server, turn instead to “Replacing the Admin Certificate” on page 80.

Tip: You can perform the following procedure even if the Dashboard is not connected to an Ignition Server. To do this, launch Dashboard and, when the Login dialog box appears, click **Cancel**. The application remains running but is not connected to an Ignition Server.

Procedure To install a copy of the admin certificate on Dashboard:

1. Contact your Ignition Server administrator and obtain a copy of the Ignition Server’s admin certificate. The certificate must be saved in a text file as a PEM-encoded certificate. For details see “Format of Certificate Files” on page 74.
2. From the Dashboard main window, select **Administration: Root Certificates**.
3. In the Root Certificates window, click **Add**.

Figure 17 Importing a Root Certificate



4. In the **Add Root Certificate** Window, click **Browse** to load the certificate file.
5. In the **Alias** field, enter a short name for this certificate. The alias is the unique key that Ignition Server uses to identify this certificate in its keystore. This can be any name you choose and need not match any value used for the Server's admin certificate.

Warning: If you choose an alias that is already in use, the newly imported certificate replaces the certificate previously aliased under that name. *Do not replace a certificate that is still needed for communicating with one of your Ignition Servers!* If you do so, you cannot connect to that Ignition Server. You can install many certificates in the Root Certificates window.

6. Click **Add**. Ignition Server adds the selected entry to the display in the **Root Certificates** list.

The new certificate resides in Dashboard's keystore. Dashboard can now connect to the Ignition Server that uses the admin certificate you added.

Replacing the Admin Certificate

The instructions below explain how to replace the *admin certificate* on the Ignition Server. Dashboard checks the Ignition Server's admin certificate in order to verify the identity of the Ignition Server before connecting to it.

Warning: Before you can replace the admin certificate you must add a copy of it to Ignition Dashboard, as explained below.

Procedure To replace the admin certificate:

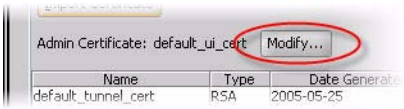
1. If you have not yet requested or imported your admin certificate into Ignition Server, do so as explained in "Getting a New Certificate" on page 75.
2. Install a copy of the admin certificate first in Dashboard. (*Failure to do this renders your Dashboard application unable to reach your Ignition Server!*) To do this:
   - From the Dashboard main window, select **Administration: Root Certificates**.
   - In the Root Certificates window, click **Add**.
   - In the **Add Root Certificate** Window, click **Browse** to load the certificate file.
   - In the **Alias** field, enter a short name for the certificate. Use a new name that is not currently used as an Alias in the Root Certificates

window. This can be any name you choose and need not match any value used for the Server's admin certificate.

Warning: If you choose an alias that is already in use, the newly imported certificate replaces the certificate previously aliased under that name. *Do not replace a certificate that is still needed for communicating with one of your Ignition Servers!* If you do so, you cannot connect to that Ignition Server. You can install many certificates in the Root Certificates window.

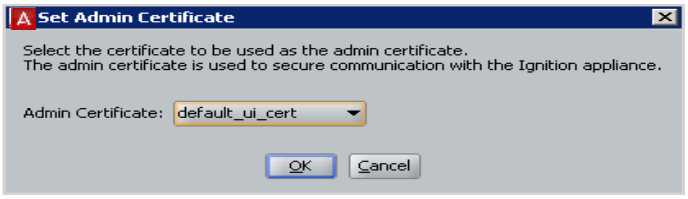
- Click **Add**. Ignition Server adds the selected entry to the display in the **Root Certificates** list.
- Click **Close** to dismiss the Root Certificates window.

3. At the top of the navigation tree in the Dashboard main window, click your site name.
4. In the Sites panel, click the **Certificates** tab and click the **Certificates** sub-tab.
5. In the **Certificates** tab, there is a section labelled **Admin Certificate** near the top of the window. This section displays the name of the current admin certificate. Click the **Modify** button.



6. In the **Set Admin Certificate** dialog, use the drop-down list to choose the certificate you want to designate as the **admin certificate**.

Figure 18 Set Admin Certificate window



7. Click **OK.**
8. Ignition Server displays a confirmation window. If you performed Step 2 above, then you can safely click **Yes** to accept the new certificate as your admin certificate. Otherwise, click **No** and return to Step 2.

You have replaced the admin certificate on Ignition Server and in Dashboard. If you want to remove your copy of the old, now-unused admin certificate from Dashboard, select the **Administration: Root Certificates** command from the Dashboard main window, find the certificate in the list, and click the Delete button. *Before you delete a certificate, make sure it is not needed to connect to any of your Ignition Servers.*

Assigning the SOAP Service Certificate

The instructions below explain how to replace the *SOAP service certificate* on the Ignition Server. The Guest Manager application checks the Ignition Server's SOAP service certificate in order to verify the identity of the Ignition Server before connecting to it.

1. If you have not yet imported your certificate into Ignition Server, do so as explained in "Getting a New Certificate" on page 75.
2. At the top of the navigation tree in the Dashboard main window, click your site name.
3. In the Sites panel, click the **Services** tab and click the **SOAP** sub-tab.
4. Click the Modify button.
5. Choose your certificate in the **SOAP Certificate** drop-down list.
6. Click OK.
7. In Avaya Identity Engines Ignition Server Guest Manager, install a copy of your SOAP service certificate. Follow the instructions in the section "Installing a SOAP Certificate," which is located in the "Configuration" chapter of the *Avaya Identity Engines Ignition Server Guest Manager Configuration Guide*.

Assigning Protocol Credential Certificates

Protocol credential certificates or "**tunnel certificates**" reside on the Ignition Server to secure PEAP and TTLS authentication transactions. In such transactions, the Ignition Server proves its identity by presenting the protocol credential certificate to the authenticating supplicant. Each supplicant must have installed on it the corresponding root certificate, so that the supplicant can validate Ignition Server's protocol credential certificate. When you write an Ignition Server authentication policy, you specify the protocol credential certificate Ignition Server uses in the context of that policy.

Ignition Server checks for the expiry and revocation of the certificates installed on the Ignition Server every twenty-four hours. If a certificate expires soon, Ignition Server logs a warning message to Security and Audit channels.

Your installation includes a temporary default certificate, called the **default_tunnel_cert**, that you can use as a protocol credential certificate.

This section explains how to install protocol credential certificates.



Warning: Before you can use a protocol credential certificate you must install its corresponding root certificate on each supplicant that is to authenticate against Ignition Server. Consult your supplicant or operating system documentation for details.

Procedure To assign protocol credential certificates:

1. Verify that each authenticating supplicant has a copy of the root certificate for the protocol credential certificate you are about to install.
2. Verify that the correct certificate has been imported into Ignition Server:
   - At the top of Dashboard's navigation panel, click on the name of your site.
   - In Dashboard's Sites panel, click the **Certificates** tab and view the **Certificates** sub-tab.
   - Confirm that your certificate is in the list. If you have not yet imported your admin certificate into Ignition Server, do so as explained in "Getting a New Certificate" on page 75.
3. In the Dashboard Main window, in the **Configuration** tree, expand the **Access Policies** node, expand the **RADIUS** node, and click on the name of your access policy. Click the **Authentication Policy** tab and click **Edit**.
4. In the Authentication Policy window, go to the Protocol Credential section.
5. In the Certificate drop-down list, select the name of your protocol credential certificate.
6. Click OK. (See "Creating an Authentication Policy" on page 219 for more information on policies.)

After you have assigned the protocol credential certificate, its policy assignments are displayed in the Certificates tab of the Certificate Management window.

Installing Protocol Root Certificates

Ignition Server uses **protocol root certificates** to verify supplicant certificates during EAP-TLS and PEAP/EAP-TLS authentication. If your policies use these authentication types, then each supplicant must have its own certificate installed, and you must add to Ignition Server the root certificate or certificates needed to validate the supplicants' certificates. Avaya refers to these root certificates as "protocol root certificates" in Ignition Server.

To import a protocol root certificate:

1. Gather the root certificates of the CAs that issued your supplicant certificates. Make sure each certificate is saved in its own text file as a PEM-encoded certificate.
2. At the top of Dashboard's navigation panel, click the name of your site.
3. In Dashboard's Sites panel, click the **Certificates** tab and click the **Protocol Root Certificates** tab.
4. Click **Import Root Certificate**.

5. Navigate to the certificate file. Observe the formatting limitations listed in “Format of Certificate Files” on page 74.
6. Click **Open**. Ignition Server imports the root certificate into the Ignition Server keystore.
7. Repeat the steps above for each additional root certificate.

After you have imported your root certificates, turn to “Adding a Certificate Revocation List URL” on page 84 to set up your CRLs.

Adding a Certificate Revocation List URL

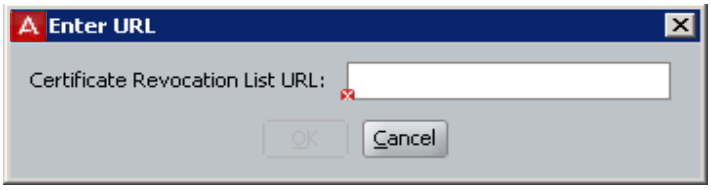
Ignition Server maintains an internal list of revoked client certificates and uses this list to deny authentication requests with revoked certificates and to alert you if any of the certificates you installed in Ignition Server have been revoked. Ignition Server builds its list by loading CRLs (certificate revocation lists) from locations that you specify. You specify these locations in the form of CRL URLs, which are Web addresses where certificate authorities publish their CRLs.

Ignition Server fetches each CRL when you add its URL to the Certificate Revocation List tab of the Certificate Management window, and it refreshes each CRL at the scheduled *Next Update* time listed in the current CRL document. (You can force an immediate CRL update. See “Refresh Button” on page 86.)

Adding a Certificate Revocation List URL to Ignition Server

To add a certificate revocation list URL:

1. At the top of Dashboard’s navigation panel, click on the name of your site.
2. In Dashboard’s Sites panel, click the **Certificates** tab and click the **Certificate Revocation List** tab.
3. Click the **New** button in the Certificate Revocation List tab.
4. In the **Enter URL** window, type the location of the URL you want the Ignition Server to monitor.



5. Click **OK**. The Enter URL window closes and the new entry appears in the Certificate Revocation List tab.

When an authentication request fails, Ignition Server enters the information in the log. If an authentication request fails, locate the log entry for the failed request. If the request failed because the corresponding certificate was revoked, you need to request a new certificate and, if necessary, update the list of URLs stored in Ignition Server.

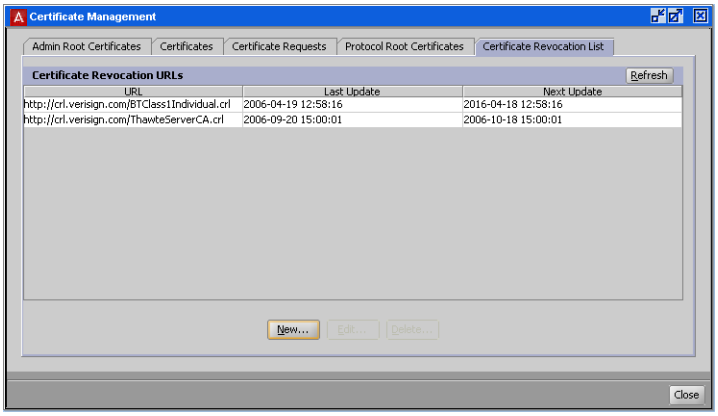
Viewing the Certificate Revocation List URLs

To view the list:

1. At the top of Dashboard's navigation panel, click on the name of your site.
2. In Dashboard's Sites panel, click the **Certificates** tab and click the **Certificate Revocation List** tab.

The Certificate Revocation List tab displays the list of certificate revocation URL entries, each of which you must provide. A correct list of certificate revocation URL entries is crucial for Ignition Server to maintain the most current data on revoked certificates. The following figure shows an example entry under this tab.

Certificate Revocation List



The sections that follow explain the window in greater detail.

URL

Each entry displayed in this column represents the URL for a certificate authority. In turn, the file accessed from this URL provides a list containing information about certificates that the certificate authority has revoked (even though they might not have expired). When a certificate has been revoked by the associated certificate authority, Ignition Server is unable to authenticate any request that references the revoked certificate.

Last Update

Each entry in the Last Update column displays the date and time when the list of revoked certificate identifiers was last generated by the certificate authority and published on the file accessed by the corresponding URL.

Next Update

Each entry in the Next Update column displays the next date and time when the list of revoked certificate identifiers is generated by the certificate authority and published on the file accessed by the corresponding URL. Ignition Server automatically refreshes the CRL at this time.

Refresh Button

You can force Ignition Server to refresh a CRL by clicking on its URL in the Certificate Revocation List tab and clicking the Refresh button. When you click the Refresh button, Ignition Server downloads the latest CRL from the specified URL. The date/time stamps in the Last Update and Next Update fields are updated.

New Button

The New button allows you to add the URL for each certificate authority that issues client certificates for authenticating incoming requests. After you create certificates, use the **New** button to add the corresponding URL entries to this list.

Viewing a Certificate

To view a certificate:

1. At the top of Dashboard's navigation panel, click on the name of your site.
2. In Dashboard's Sites panel, click the **Certificates** tab and click the tab for the type of certificate you want to view.
3. In the list, click the certificate.
4. Click the **View** button. Ignition Server displays the contents of the selected certificate.
5. Click **OK** to close this display window.

To view the copy of the admin certificate saved in Dashboard, select the command **Administration: Root Certificates**, find the certificate in the list, and click **View**.

Deleting a Certificate or Certificate Request

The Delete command succeeds only if the certificate is *not* currently being used by Ignition Server. If the certificate has been assigned as the admin certificate or as a protocol credential certificate, then the Delete command

fails. Remove the certificate's usage assignment and then delete it. (Root certificates can be removed at any time; there is no need to remove the assignment in the case of a root certificate.)

Procedure To delete a certificate:

1. At the top of Dashboard's navigation panel, click on the name of your site.
2. In Dashboard's Sites panel, click the **Certificates** tab and click the tab for the type of certificate or certificate request you want to delete.
3. In the list, click the certificate or request.
4. Click the **Delete** button. ***Note!*** Deleting an existing certificate request renders unusable any certificate the CA sends you based on this request.
5. Ignition Server prompts you to confirm the deletion request.

   To carry out the deletion, click **Yes**. Ignition Server removes the entry from the display in the **Root Certificates** Window and deletes the certificate from the Dashboard's keystore.

If you have deleted a certificate used to secure Dashboard or Guest Manager (SOAP) communication, make sure you update Dashboard or Guest Manager to use the replacement certificate.

Viewing an Existing Certificate Request

To view an existing certificate request:

1. At the top of Dashboard's navigation panel, click on the name of your site.
2. In Dashboard's Sites panel, click the Certificates tab and click the Certificate Requests tab.
3. In the list, click the request.
4. Click the **View** button. Ignition Server displays the contents of the selected request.

Authenticators

This chapter introduces the concept of authenticators, and describes their relationships to access policies.

Contents

- "Introduction to Authenticators" on page 89
- "Authenticator Hierarchy and Containers" on page 91
- "Creating an Authenticator" on page 96
- "Placing an Authenticator in the Authenticator Hierarchy" on page 103
- "Finding an Authenticator" on page 104
- "Pinging an Authenticator" on page 104
- "RADIUS Global Authenticator" on page 104
- "RADIUS Subauthenticators" on page 105
- "Processing Authenticator Requests" on page 109
- "Changing the Configuration of an Authenticator" on page 109
- "Removing an Authenticator From Its Place in the Hierarchy" on page 109
- "Renaming Authenticators" on page 110
- "Deleting Authenticators" on page 110

Introduction to Authenticators

An authenticator is a device that allows other devices to connect to your network. Wired switches, wireless access points (APs), and VPN devices are all types of authenticators. Each such device is represented in the Avaya Identity Engines Ignition Server by an authenticator record. You apply Ignition Server access policies to your authenticator to set the access rules for all users who enter your network through that authenticator. *In other words, the authenticator record is the key that maps your access policies to your switches, APs, and other equipment.*

Applying Ignition Server access control to your authenticators is straightforward: You connect the switch and Ignition Server's RADIUS port to the same network, you save an authenticator record in Ignition Server to represent the switch, and, in the switch, you set the *RADIUS server port* setting to point to Ignition Server's RADIUS server.

Two special bulk handling approaches give you more flexibility with your Ignition Server authenticator set-up. *Authenticator bundles* allow you to represent all authenticators on a subnet with a single authenticator record, and the *global authenticator* record allows you to create a default access policy that applies to requests from unknown authenticators. See "Creating an Authenticator" on page 96 for information on bundles, and see "RADIUS Global Authenticator" on page 104 for information on the global authenticator.

Matching an incoming request to an authenticator record

When Ignition Server receives an authentication request, it must find the right access policy to determine its ALLOW/DENY response. The access policy is in the authenticator record, which Ignition Server finds as explained below.

Each *authenticator record* in Ignition Server has an IP address and a netmask associated with it. An authenticator can represent a single device (no bundle) or it can represent one or more devices in the same subnet (an *authenticator bundle*). An authenticator or bundle can contain *subauthenticators*, each representing a logical switch or SSID. Finally, a *global authenticator* record acts as a catch-all. The global authenticator has no IP address associated with it, so it matches any IP address. In other words, it's an authenticator that represents your entire network.

When an authentication request arrives, Ignition Server searches for a matching authenticator record in the order of small scope to large scope:

- First it looks for an exact IP address match to an *authenticator record*,
- then it tries to match *small authenticator bundles* (large netmask),
- then *large authenticator bundles* (small netmask),
- and finally, the *global authenticator*. Having a permanent global authenticator means that Ignition Server always finds a match.

When Ignition Server finds a matching authenticator or bundle, it searches inside that record for a subauthenticator that matches the incoming RADIUS request:

- If a matching subauthenticator is found, then its access policy is used.
- If *no* matching subauthenticator is found, then the authenticator's RADIUS access policy is used.

Important! When Ignition Server receives a RADIUS request, it applies only the policies of the authenticator record keyed to the IP address that sent the request. This means that, if you set an authenticator record to "disabled," all requests originating from that authenticator record's IP address are rejected. It does *not* mean that a bundle or the global authenticator takes over servicing requests for the disabled authenticator. Thus, if you disable the matching

authenticator (or if the matching authenticator has no support for the protocol of the request), then your request is discarded, regardless of the configuration of other authenticators, including the global one.

Important! Since the Guest Manager is also a kind of authenticator you will receive an error indicating the IP address already exists if you add a Guest Manager which has the same IP address of an existing authenticator. This error also appears if you try to add an authenticator which has the same IP address of an existing Guest Manager.

Authenticator Hierarchy and Containers

Each authenticator bears an *authenticator container label* that indicates, typically, where the authenticator is located or what part of your organization it belongs to. The authenticator hierarchy is a hierarchy of containers that lets you sort and categorize your authenticators (geographically, organizationally, or in some other way) so that your access policies can take this into account and apply appropriate access rules.

A container holds a set of authenticators you have grouped in Ignition Server, and the authenticator hierarchy is the tree of these containers. Organizing containers into a tree allows you to group your authenticators (for example, geographically) for ease of management, and allows you to create authorization rules based on an authenticator's location in the hierarchy.

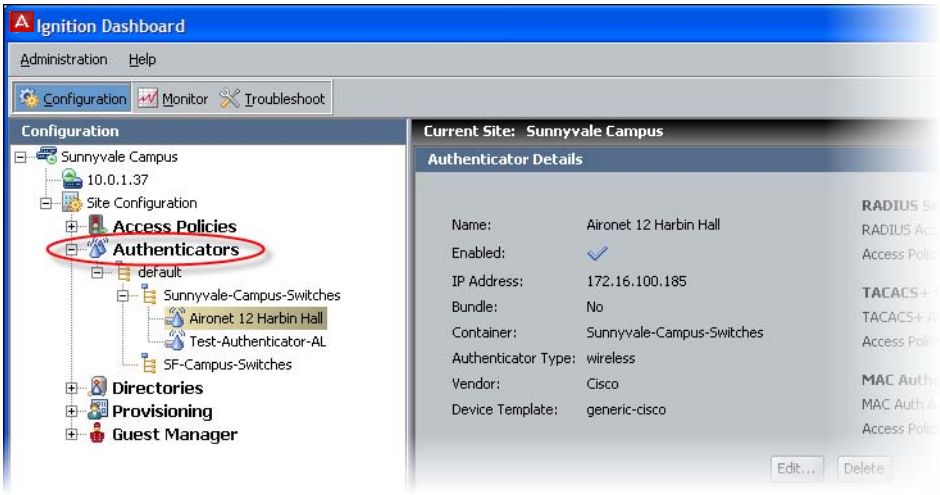
At user login time, when your authorization policy checks the authenticator container label, the authenticator is considered to belong to its own container as well as to all containers in its family tree: parents, grandparents, and so on up the tree.

Even if you do not create a hierarchy, you can use containers individually to apply labels to authenticators.

How You Build the Hierarchy

Each authenticator belongs to exactly one container and has exactly one access policy. A container may contain many authenticators and other containers, forming the hierarchy. You define the hierarchy (in the **Authenticator Hierarchy** tree) from the top down, by creating each container, creating its child containers, and so on, as explained in "Creating the Authenticator Hierarchy" on page 94. You can place an authenticator in the hierarchy using the **Authenticator Details panel**, as explained in "Placing an Authenticator in the Authenticator Hierarchy" on page 103.

Figure 19 Example View of the authenticator hierarchy tree



Using the Hierarchy in Your Policies

At user login time, Ignition Server can evaluate the authenticator's container or its position in the hierarchy and make access decisions based on that. The container name or hierarchy position is considered in two contexts:

- **User lookup**: Ignition Server can be set to use a specific directory to look up users who try to connect to a given authenticator. The authenticator is identified based on its container name or its position in the authenticator hierarchy. For information on associating an authenticator with the user directories that serve that authenticator, see "Understanding Identity Routing Policy" on page 220.
- **User authorization**: To make the access decision, your authorization policies can check which authenticator the user is connecting to. The authenticator is identified based on its container name or its position in the authenticator hierarchy. For information on making access decisions based on the authenticator's container or hierarchy position, see "Authenticator Attributes" on page 237.

For example, you can create a rule that allows your travelling sales staff to connect to the network from any Ethernet port in any of your offices in Colorado. You would create this policy in Ignition Server as follows:

- Create a container called *Colorado* in your authenticator hierarchy.
- In the *Colorado* container, create two child containers: *Branch-Office-Denver* and *Branch-Office-Boulder*.
- Assign your Denver office's network switches to the *Branch-Office-Denver* container, and assign your Boulder office's switches to the *Branch-Office-Boulder* container.

- Write an Ignition Server user authorization policy that lets all members of the *Sales-Dept* connect via any authenticator that's in the *Colorado* container. (Expressed in Ignition Server's rule-writing lingo, your rule triggers an *ALLOW* when it encounters the combination of the container name *Colorado* and an authenticated user who belongs to the group *Sales-Dept*.)

To find out how to write a policy like this, see "Creating a RADIUS User Authorization Policy" on page 242.

Default Container

The top-level container is initially named "default." You can rename it, but it cannot be deleted.

Each authenticator or authenticator bundle is automatically part of the Default container. As shown later in this chapter, you can choose to associate a container with another container within the hierarchy.

Authenticator Hierarchy Tasks

The tasks you can perform on hierarchies are:

- "Creating the Authenticator Hierarchy" on page 94
- "Adding a Container to an Authenticator Hierarchy" on page 94
- "Renaming a Container in an Authenticator Hierarchy" on page 95
- "Deleting a Container from an Authenticator Hierarchy" on page 95
- "Placing an Authenticator in the Authenticator Hierarchy" on page 103
- "Finding an Authenticator" on page 104
- "Removing an Authenticator From Its Place in the Hierarchy" on page 109

Window Layout

When you work on authenticators, Dashboard's **Configuration** tree shows the following elements:

- **Authenticator Hierarchy**: the representation of the virtual tree of **containers**.
- **Actions**: a pulldown or right-click menu to manipulate the containers in the **Authenticator Hierarchy Tree**. The commands associated with the containers in the hierarchy are Add Container, Rename Container, Delete Container.
- **Authenticator Summary list**: Lists all authenticators in the container, and, if **Include descendants of selected container** is checked, also lists the authenticators associated with all sub-containers of the selected container. When this checkbox is not selected, the display shows the authenticators that

are directly associated with the selected container only (that is, all sub-containers are excluded).

- The **New**, **Edit**, and **Delete** buttons in this panel enable you to add a new authenticator, select and modify, or delete an existing authenticator.

Creating the Authenticator Hierarchy

Set up the container hierarchy to collect your switches and APs into groups that make it easier for you to manage security on your network. For example, containers within the tree can be based on geographic regions, departmental divisions, campuses, or functional teams.

To create the hierarchy:

1. In the **Configuration hierarchy** on the left side of Dashboard's main window, click your site, expand Site Configuration, expand **Authenticators** and click on the container that you want to be the root of your hierarchy. Typically this is the "default" container.
2. Select the command, **Actions: Add Container**.
3. In the **Add Container** window, type a name for the container and click **OK**.
4. Add child containers by clicking on the container that you want to be the parent of the new, child container and then using the **Actions: Add Container** command to add the child container.
5. Add authenticators to the hierarchy:
   - Add new authenticators to the hierarchy by clicking on the container to own the authenticator and then clicking the **New** button on the right side of the window.
   - Add existing authenticators to the hierarchy as explained in "Placing an Authenticator in the Authenticator Hierarchy" on page 103.

Adding a Container to an Authenticator Hierarchy

In order to add a new container to the authenticator hierarchy,

1. In the **Configuration hierarchy** on the left side of the window, click your site, expand Site Configuration, expand **Authenticators** and click on the parent container under which you are defining a new container.
2. Use the Right-Click on the mouse to select **Add Container.** Alternatively, select the parent container, and use the **Actions → Add Container**... command.
3. The **Add Container** window appears, listing the name of the parent container to which the new container is to be added
4. Enter the unique name for the new container in this window.
5. Click **OK**.

The **Authenticators** section of the **Configuration** tree now displays the new container under the designated parent container.

Renaming a Container in an Authenticator Hierarchy

If you rename a container currently used in an authorization policy, that authorization policy might no longer work as expected. For troubleshooting, see “Problem: Authorization policy stops working unexpectedly” on page 453.

In order to rename a container in the authenticator hierarchy,

1. In the **Configuration hierarchy** on the left side of the window, click your site, expand Site Configuration, expand **Authenticators** and click on the container.
2. Use the Right-Click on the mouse and select **Rename Container.** Alternatively, use the **Actions → Rename Container**... command.
3. Enter the new name for the selected container.
4. Click **OK**.

The new name for the container appears in place in the **Authenticators** section of the **Configuration** hierarchy tree in the main window.

Default Container: Even after you rename the default container, Ignition Server does not permit you to delete that container.

Deleting a Container from an Authenticator Hierarchy

When you delete a container, Ignition Server Dashboard removes the container from the authenticator hierarchy.

Avaya strongly recommends that you do not delete any container that is being used in an authorization policy. For troubleshooting tips, see “Problem: Authorization policy stops working unexpectedly” on page 453.

To be deleted, a container must be empty. More specifically, Ignition Server does not permit the deletion of a container under the following conditions:

- The container is associated with an authenticator or authenticator bundle.
- The container is used in an identity routing policy.
- The container is parent to one or more containers.
- It is the **Default** container.

Procedure

To delete a container from the authenticator hierarchy, follow these steps:

1. Delete or move all authenticators and authenticator bundles associated with this container. (See Removing an Authenticator From Its Place in the Hierarchy on page 109.)
2. Delete or move all containers that are child containers of this container.
3. Make sure the container is not used in an authorization policy:
   - In Dashboard's Configuration hierarchy tree, expand **Access Policies** and expand **RADIUS**. Click on the first access policy in the list.
   - Click the **Identity Routing** tab, inspecting the authenticator hierarchy column. If the container's name appears in any policy, click the Edit button and remove it from the policy.
   - In the **Authorization Policy** tab, in the **Rule Names** section, click each rule name and read the **Rule Summary**. Look for the phrase *Authenticator.Authenticator Hierarchy* followed by the name of the container you plan to delete. If you find the to-be-deleted container, click **Edit** and remove it from the policy.
   - Repeat these steps for each RADIUS policy.
   - In the Configuration hierarchy tree, click **MAC Auth**. In the **Authorization Policy** section, repeat the steps you just performed on the **RADIUS** policies.
   - Repeat these steps for each MAC Auth policy.
4. In Dashboard's Configuration hierarchy tree, expand Site Configuration and expand Authenticators. Find your authenticator in the tree.
5. Right-click your authenticator and select **Delete Container.** Alternatively, select the container, and use the **Actions → Delete Container**... command**.**
6. Click **OK**. Ignition Server Dashboard removes the container from the authenticator hierarchy.

Creating an Authenticator

An authenticator is a device (switch, wireless access point, or VPN concentrator) that allows other devices to connect to your network. To set up Ignition Server to manage access control and provisioning for a switch or other device, save the device as an authenticator in Ignition Server, as shown in the steps that follow.

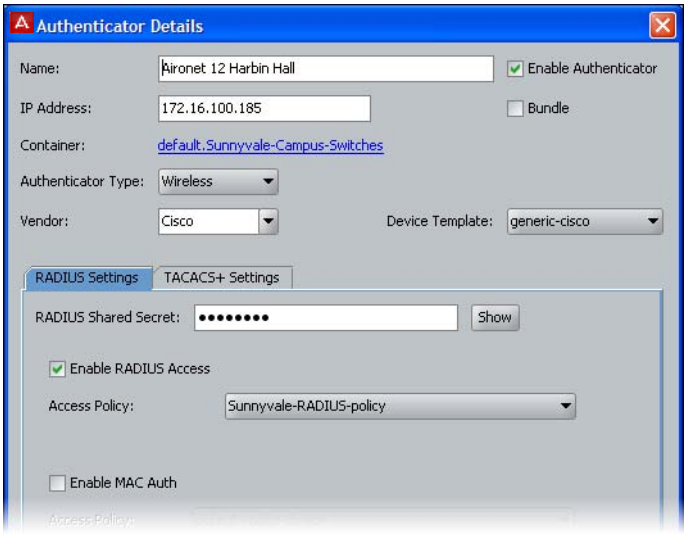
Note: If you need to create several authenticators, you may prefer to create them in bulk by importing the authenticator information in the specified comma-separated values (CSV) format. See "Importing Authenticators" on page 100 for more information.

Note: If you are using Access Portal as an authenticator, use this procedure: "Registering Access Portal as an Authenticator in Ignition Server" on page 101.

A note on authenticator vendor and device template: The authenticator vendor name and device template serve two purposes within Ignition Server: The first is to tell Ignition Server's RADIUS server which device dictionary to use to interpret or format the RADIUS attributes coming from or going to the authenticator. The second is to let you write authorization rules that apply a particular policy to certain switches, based on the device template name or vendor name of those switches.

1. Connect your authenticator to a network where it can reach Ignition Server's RADIUS port.
2. Do one of the following:
   - In Dashboard's **Configuration** tree, open the **Authenticators** node, navigate the containers, and click on the container that will hold your authenticator. Click **New** at the bottom of the **Authenticator Summary** panel; or
   - In Dashboard's **Configuration** tree, click **Site Configuration**. On the right half of the window, click **4. Authenticator**. Ignition Server displays the Authenticator Details window.

Figure 20 Authenticator Details window



3. Specify the details that describe the authenticator:
   - **Name:** Enter a unique name for the new authenticator. This is the name by which Ignition Server refers to your authenticator. This is a required field.

- **Enable Authenticator:** Select this checkbox to enable the new authenticator.
- **IP Address**: The IP Address of the authenticator.
- **Bundle:** Select this checkbox to make this authenticator an *authenticator bundle*. Authenticator bundles are a way to perform *authenticator wildcarding* that allows one authenticator bundle to represent all the authenticators on a subnet. With the bundle in place, Ignition Server handles service requests coming from any authenticator in the specified subnet, provided the device presents the correct, common shared secret. When you select the **Bundle** checkbox, the window display changes to display the **Subnet Mask** ("/") field next to the **IP Address** field. Type your subnet mask here in CIDR notation.
- **Container:** Each authenticator belongs to a *container* that indicates, typically, where the authenticator is located or what part of your organization it belongs to. (See "Authenticator Hierarchy and Containers" on page 91.) You can change the container association by clicking the blue text. See "Placing an Authenticator in the Authenticator Hierarchy" on page 103.
- **Authenticator Type:** Specify what type of device the authenticator is: Any, Wired, Wireless, VPN, SIP, or Other. Default is "Any." Your authorization rules can check this value and apply policies based on authenticator type.
- **Vendor**: Specify the manufacturer or the authenticator. This setting dictates the set of device templates that are available for this authenticator. If you do not select an entry for Vendor, the new authenticator belongs to the "default" vendor category.
- **Device Template:** Specify the Ignition Server device template for this authenticator. The device template sets rules that govern how Ignition Server sends and receives RADIUS and TACACS+ messages to and from the authenticator. (See "Device Templates" on page 270.)

4. If this authenticator will use RADIUS authentication, click the **RADIUS Settings** tab and set the following:
   - **RADIUS Shared Secret**: Enter the **Shared Secret** that you have configured in the authenticator device. If you are creating an authenticator bundle, all authenticators in the Bundle must use the same shared secret. This is a required field.
   - Tick the **Enable RADIUS Access** checkbox. You must tick this checkbox to provide RADIUS service to the authenticator.
   - **Access Policy:** Select the Ignition Server access policy that will regulate RADIUS access requests relayed by this authenticator. If you

do not select an access policy, the new authenticator uses the “default” access policy.

Tip! *But what if I need to specify more than one policy for a single switch or AP?* Create a subauthenticator for each policy you want to add. For instructions, see “RADIUS Subauthenticators” on page 105.

- **Enable MAC Auth**: Tick this checkbox to provide authorization based on the MAC address of the device that’s trying to connect. See “Introduction to MAC Authentication” on page 321.

5. If this authenticator will use TACACS+ authentication, click the **TACACS+ Settings** tab and set the following:
   - Tick the **Enable TACACS+ Access** checkbox. You must tick this checkbox to provide TACACS+ service to the authenticator.
   - **TACACS+ Shared Secret**: Enter the **Shared Secret** that you have configured in the authenticator device. If you are creating an authenticator bundle, all authenticators in the Bundle must use the same shared secret. This is a required field.
   - **Access Policy:** Choose the Ignition Server access policy that will regulate TACACS+ access requests relayed by this authenticator. See “Creating a TACACS+ Access Policy” on page 313.
6. Click **OK** in the **Authenticator Details** window. Ignition Server displays the newly created authenticator in its place in the Configuration hierarchy panel. In the future you can expand **Site Configuration** and expand **Authenticators** to see the authenticator record.
7. Use a console or management tool to log in to your switch (or other authenticator device) and edit the switch configuration to set your Ignition Server as the RADIUS and/or TACACS+ server for the switch. Consult your switch manufacturer’s documentation for instructions.

   For example, to designate the RADIUS server for a Cisco 2950, you would enter configure terminal mode on the Cisco 2950 console and, using the form, `radius-server host <ip address of the RADIUS interface on the Ignition Server> auth-port 1812 acct 1813 key <your shared secret>`, you might type, for example:

   ```
   radius-server host 172.32.102.43 auth-port 1812 acct
   1813 key 1234
   ```

   Note: To find out the Ignition Server RADIUS and TACACS+ port addresses, go to Dashboard’s **Configuration** tree, click the **Site** name (this is usually the name at the top of the tree), click the **Services** tab and click the **RADIUS** or **TACACS+** tab. The **Bound Interface** field

indicates the port. To find the IP address, click the **Node** name or IP address in the tree, and click the **Ports** tab.

Your authenticator set-up is complete. If you do not already have appropriate security policies defined for the switch, turn to Chapter , “User Authentication Policy” for instructions on setting up access rules.

Importing Authenticators

If you need to create several authenticators, you can create them in bulk by importing the authenticator information in the specified CSV format.

To import several authenticators:

1. Using a tool of your choice, create a file containing the authenticator records you want to import. The file must be in the comma-separated values (CSV) format. Ignition Server requires that you provide the fields in this order:

   `Name,IP Address,`**`Enable Authenticator`**`,`
   `Bundle``,Mask,Container,Authenticator`
   `Type,Vendor,Device Template,`**`Enable Radius`**
   `Access``,Radius Secret,Radius Access Policy,`**`Enable`**
   `TACACS+ Access``,TACACS+ Secret,TACACS+ Access`
   `Policy,`**`Enable MAC Auth`**`,MAC Access Policy,`**`MAC Auth -`**
   `No Passwd``,`**`MAC Auth - Passwd`**`,`**`MAC Auth - Use RADIUS`**
   `Secret`

   Replace fields in bold text with a value of “true” to enable or select that option or “false” to disable or not select that option. For example, if you imported a file containing the line

   `testAuthenticator2,134.177.229.201,true,,,default,SIP`
   `,3com,generic-3com,true,1234,default-radius-`
   `user,,,,true,default-radius-device,,,true`

   the import action adds an authenticator with a name of testAuthenticator, an IP address of 134.177.229.201, the **Enable Authenticator** check box selected, associated with the default container, and Authenticator Type as SIP, etc.

2. In Dashboard’s **Configuration** tree, open the **Authenticators** node, navigate the containers, and click on the container that will hold your authenticators.

3. In the Authenticator Summary window, click **Import**.

4. Click the **Browse** button, navigate to your csv file, click it, and click **Open**. The Authenticators Import Status window appears. You can see if the import was a success, and view the reason for failure if the import failed. You can click **Copy** to copy the authenticator information, or click **Print** to print the information.
5. In the **Authenticators Import Status** window, click **OK**. The Authenticator Summary window displays the imported authenticators.

Exporting Authenticators

To export a set of authenticators:

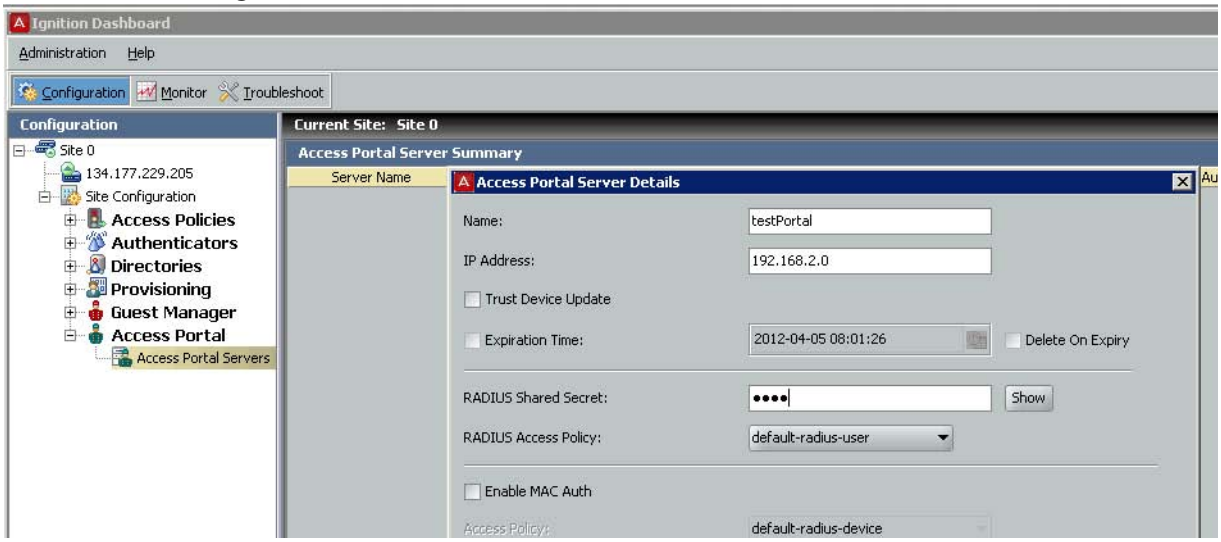
1. In Dashboard's **Configuration** tree, open the **Authenticators** node, navigate the containers, and click on the container that holds your authenticators.
2. In the Authenticator Summary window, click **Export**.
3. The Export Authenticator window displays the following message: "Do you want to export the Authenticators from the child containers?" Click **Yes**, to export the Authenticators from the child containers or click **No** to not export the Authenticators from the child containers.
4. The Save window appears. In the **Save In** field, navigate to where you want to save your csv file. You can change the default file name of "Authenticators_config.csv" if you want. Click **Save** to export the authenticator records.

Registering Access Portal as an Authenticator in Ignition Server

To register Access Portal as an authenticator in the Ignition Server:

1. In the Dashboard **Configuration** tree, expand the **Access Portal** node, and click on **Access Portal Servers**.
2. Click **New**. The Access Portal Server Details window appears.

Figure 21 Access Portal Server Details window



3. In the **Access Portal Server Details** window, specify the following:

   - **Name**: Enter a name for the Access Portal.
   - **IP Address**: Enter the IP address of the Access Portal.
   - **Trust Device Update**: Select this check box if you want to use this Access Portal Server for device learning or automatic registration of devices. If you select this check box, the device that the client used to log in on will be added to the Ignition Server's internal known devices list. **Note**: If you select this check box, you must go to the Access Portal Administration Web UI, click **Services > Captive Portal**, and select the **Enable Device Fingerprinting** check box.
   - **Expiration Time**: Select this check box if you want to specify an expiry date for the devices learned through Access Portal. Click the clock-and-calendar icon and use the arrow keys to set the date and time it expires. Click outside the clock and calendar dialog to close it.
   - **Delete On Expiry**: Select this check box if you want the Ignition Server to delete learned devices after the expiry date.
   - **RADIUS Shared Secret**: Enter the RADIUS Shared Secret that you configured in Access Portal.
   - **RADIUS Access Policy**: The RADIUS access is enabled by default. Select the Ignition Server access policy that regulates RADIUS access requests relayed by Access Portal. If you do not select an access policy, Access Portal uses the default access policy (default-radius-user).
   - **Enable MAC Auth**: Select this check box to provide authentication based on the MAC address of the device that is trying to connect. See "Introduction to MAC Authentication" on page 321 for information on

setting up MAC authentication on Access Portal.
Note: If you select the **Enable MAC Auth** check box, you must also enable RADIUS MAC authentication on Access Portal. From the Access Portal Administration Web UI, click **Services > Captive Portal** and scroll down to the Authentication section. In the RADIUS MAC authentication section, select the **Enable RADIUS MAC authentication** check box. In the Shared secret field, enter the Ignition Server shared secret and click **Save**

4. Click **OK**. The Access Portal Summary page appears.

Placing an Authenticator in the Authenticator Hierarchy

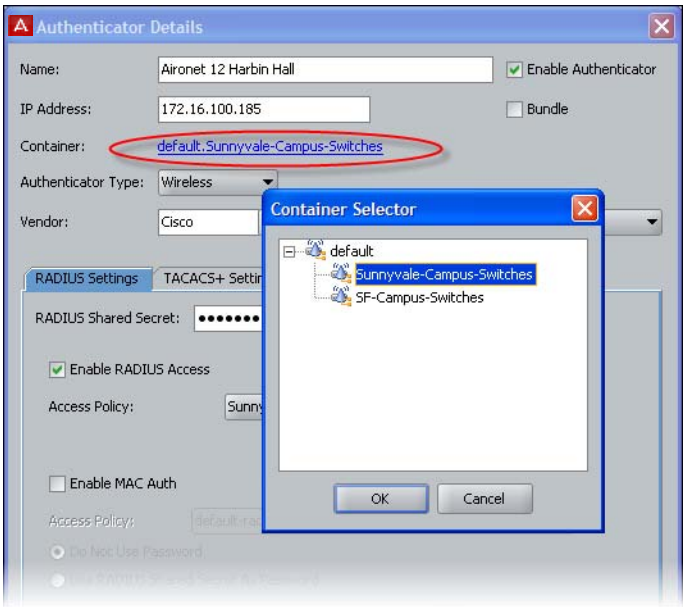
Each authenticator belongs to a *container* that indicates, typically, where the authenticator is located or what part of your organization it belongs to. To set the authenticator hierarchy ("Container") label, do the following:

1. In Dashboard's Configuration hierarchy tree, expand **Authenticators**. Find your authenticator in the tree, click its name, and click **Edit**.

   Hint: to list every saved authenticator in the system, click the top node in the Authenticator Hierarchy (by default, this node is called "default") and on the right side of the window, click the **Include Selected Hierarchy Descendents** checkbox.

2. In the **Container** field, click the blue text.
3. In the **Container Selector** window, navigate to and click the desired container to choose it.
4. Click **OK** to confirm your selection.

Figure 22 The Container Selector window



5. Click **OK** to save the authenticator.

Finding an Authenticator

To find an authenticator record:

1. In Dashboard's Configuration hierarchy tree, expand Authenticators and click **default** (the root container).
2. Do one of the following:
   - *To list every saved authenticator in the system*: Click the top container in the Authenticator Hierarchy (by default, this container is called "default") and click the **Include Selected Hierarchy Descendents** checkbox.
   - *To find all the authenticators in an authenticator hierarchy container*: Navigate the Authenticator Hierarchy and click the container whose authenticators you want to view. If the container has sub-containers whose authenticators you want to view, tick the **Include Selected Hierarchy Descendents** checkbox.

Pinging an Authenticator

To verify that an authenticator is reachable and responding, do the following:

1. Make sure the authenticator has been created in Ignition Server as described in Creating an Authenticator starting on page 96.
2. In Dashboard, click the Troubleshoot tab.
3. Click the IP address or name of your node.
4. Click the Network tab.
5. In the Port field, select the Ignition Server Ethernet port where your Ignition Server RADIUS service is running.
6. In Ping Target, type the IP address of your authenticator.
7. Click Start.

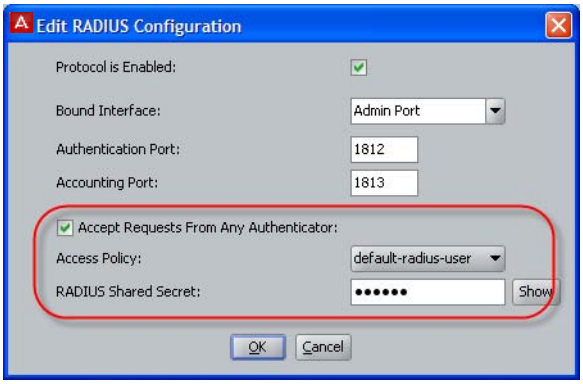


Important! If Ignition Server is running in HA mode, the ping originates from the primary node of Ignition Server.

RADIUS Global Authenticator

As explained in "Introduction to Authenticators" on page 89, the *global authenticator* record allows you to create a default RADIUS access policy that applies to requests from unknown authenticators.

Figure 23 The global authenticator is defined in the RADIUS configuration



When Ignition Server uses the global authenticator to handle a request, it logs the action with the authenticator name “global-default.” In your Ignition Server access policies, you can write policy rules that test for the label, “global-default” and apply policy based on whether the request is being handled by the global authenticator.

Note: If a request matches the global authenticator but the request’s protocol in the global authenticator is disabled, Ignition Server logs an “unknown authenticator” message.

Important: Use of a *MAC Auth* policy is not permitted in the global authenticator.

Creating the Global Authenticator

1. In the Configuration hierarchy tree of Dashboard, click on your site’s name, click the **Services** tab, and click the **RADIUS** tab.
2. Click **Edit**.
3. In the Edit RADIUS Configuration window, tick the **Accept Requests from Any Authenticator** checkbox.
4. Choose your **Access Policy**. This is the default RADIUS access policy for all requests from unknown authenticators. You must use a *RADIUS policy*; you cannot use a *MAC Auth policy*.
5. Type the **RADIUS Shared Secret**. Ignition Server responds only to authenticators that pass this secret string.
6. Click **OK**.

RADIUS Subauthenticators

Some network configurations require that you have a number of logically different switches that contact Ignition Server from the same IP address. For example, some wireless access points can beacon multiple SSIDs, but users’ RADIUS requests to connect to the AP — regardless of SSID — arrive at

Ignition Server from the same IP address. As an administrator, you may wish to configure different RADIUS access policies for these logically different switches.

Ignition Server handles this with a feature called *subauthenticators*. The subauthenticator feature allows you to configure different RADIUS access policies (user lookup, authentication, authorization, and provisioning policies) that will be used for logically different switches operating from the same IP address. The physical switch is represented in Ignition Server by an *authenticator record* keyed to the switch's IP address. All RADIUS requests originating from this IP address are handled according to this authenticator record. Inside the authenticator record you define a *set of subauthenticators* for the set of logical devices that operate from the IP address. When you create each subauthenticator, you key it to a RADIUS attribute value. If the RADIUS request contains that value, then the subauthenticator is used.

Upon receiving a request, Ignition Server finds the authenticator record for the IP address and then chooses the first subauthenticator whose key value matches a value in the RADIUS request. The subauthenticator specifies the RADIUS access policy to be used for that logical switch. If no matching subauthenticator is found, then the RADIUS access policy of the authenticator record is used.

In other words, Ignition Server inspects the incoming RADIUS request, and if it contains a particular value, then Ignition Server uses the access policy you have keyed to that value. This allows you to treat one switch as a number of logical switches in order to apply the correct policy to each logical switch.

For a more complete description of how and when a subauthenticator is used, see "Matching an incoming request to an authenticator record" on page 90.

Viewing and Editing Subauthenticators

1. In Dashboard's **Configuration** hierarchy tree, expand **Authenticators**. Expand the sub-nodes in the tree until you find the authenticator whose subauthenticators you want to view.
2. In the tree, click the authenticator's name.
3. The **Sub Authenticators** panel occupies the lower half of the window and displays the subauthenticators of the authenticator.
4. To edit a subauthenticator, click its name and click the **Edit** button. See "Creating a Subauthenticator" on page 107 for descriptions of the fields of the Sub Authenticator Details window.

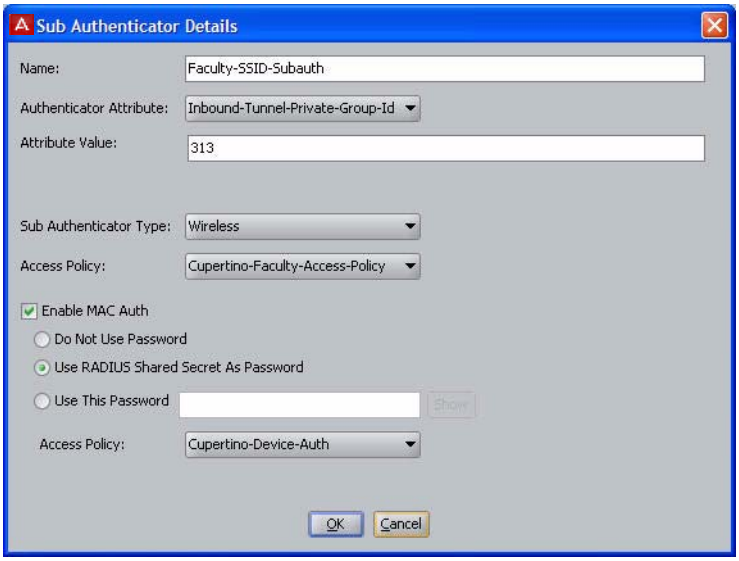
Creating a Subauthenticator

Each subauthenticator definition is tied to an authenticator record. An authenticator record can contain many subauthenticators. You create one authenticator record per physical switch or access point, and then, inside that authenticator record, you create as many subauthenticators as you need to accommodate the logical switches or SSIDs of that piece of hardware.

Create a subauthenticator as follows:

1. In Dashboard's **Configuration** hierarchy tree, expand **Authenticators**. Expand the sub-nodes in the tree until you find the authenticator to which you want to add a subauthenticator. (**Important!** If you have not defined your authenticator, turn to "Creating an Authenticator" on page 96 and create it now.)
2. In the tree, click the authenticator's name. The **Authenticator Details** panel appears, showing the **Sub Authenticators** panel in its lower half.
3. At the bottom of the panel, click **New**.

Figure 24 Creating a subauthenticator in the Sub Authenticator Details window



4. In the Sub Authenticator Details window, type a **Name** for the subauthenticator. This name appears in the RADIUS logs and may be used in your authorization rules.
5. In the **Authenticator Attribute** dropdown list, choose the *inbound RADIUS attribute* whose value triggers the use of this subauthenticator. This RADIUS attribute is sent by the authenticator hardware.

   For example, some manufacturers use the RADIUS attributes mapped as *Port-Number* or *Inbound-Called-Station-Id* to indicate the SSID. To

find out how to view Ignition Server's RADIUS attribute mappings, see "Finding an Inbound Attribute" on page 266.

6. In the **Attribute Value** field, type the RADIUS attribute value that triggers the use of this subauthenticator.
7. In the **Sub Authenticator Type** dropdown list, choose the type of virtual device that this subauthenticator represents, such as *Wired*, *Wireless*, or *VPN*.
8. In the **Access Policy** dropdown list, choose the Ignition Server RADIUS access policy that controls user access to this subauthenticator.
9. If you wish to allow MAC Auth on this subauthenticator, tick the **Enable MAC Auth** checkbox and
   - Specify how Ignition Server verifies the authenticator password: "**Do not use password**" tells Ignition Server to skip password checking; "**Use RADIUS shared secret as password**" tells Ignition Server to use the authenticator's RADIUS shared secret; and "**Use this password**" lets you specify your own password.
   - In the **Access Policy** dropdown box, choose your *MAC Auth policy*. (If you need to create one, see "Creating a MAC-Auth Policy" on page 322.)
10. Click **OK** to save the subauthenticator.
11. If you wish to define more subauthenticators, click **New** again. Once you have created all the subauthenticators of this authenticator, you can sort them as shown in the next section, "Sorting Subauthenticators" on page 108.

Sorting Subauthenticators

When an authenticator has multiple subauthenticators, Ignition Server responds to an incoming RADIUS request by searching from the top of the **Sub Authenticators** list to the bottom and using the first subauthenticator whose attribute/value pair matches a RADIUS attribute/value pair in the request. If any of your subauthenticators is more widely applicable than others, then you may have to sort the list of subauthenticators to ensure the desired subauthenticator takes effect.

Sort your subauthenticators as follows:

1. In Dashboard's **Configuration** hierarchy tree, expand **Authenticators** and expand its sub-nodes to find the authenticator whose subauthenticators you want to sort.
2. In the tree, click the authenticator's name.
3. In the **Sub Authenticators** panel, click **Order...**

4. In the Sub Authenticator Ordering window, click on the name of a subauthenticator and use the *up-* or *down-arrow* buttons to move it to the correct position.
5. Click **OK** to save the sort order.

Processing Authenticator Requests

Before granting access to the network, Ignition Server processes authenticator requests by validating the identity of the end user and performing the checks prescribed in your authorization policies. These requests use the RADIUS protocol.

How Ignition Server Processes RADIUS Requests from Authenticator Bundles

Device Dictionary files are used to control vendor-specific capability using the RADIUS protocol. RADIUS allows equipment manufacturers to expose proprietary features using vendor-specific RADIUS Attributes. The device dictionary file defines these vendor-specific attributes.

When processing RADIUS requests from an authenticator bundle, Ignition Server follows the rules listed below to arrive at the most specific Device Dictionary to use:

- When both Vendor and Model are specified, Ignition Server uses the Device Dictionary specific to that equipment.
- When a vendor is specified but not a model, the Vendor's Device Dictionary is used.
- When no vendor or model is specified, Ignition Server uses the generic RADIUS Device Dictionary.

Changing the Configuration of an Authenticator

After you have configured an authenticator, you can change its settings at any time.

1. In Dashboard's Configuration hierarchy tree, expand **Authenticators**. Expand the sub-nodes in the tree until you find your authenticator. Click its name and click **Edit**. The Authenticator Details window opens.
2. Edit the settings for the selected authenticator. See "Creating an Authenticator" on page 96 for an explanation of each field.
3. Click **OK** to apply your changes.

Ignition Server updates the configuration for the selected authenticator.

Removing an Authenticator From Its Place in the Hierarchy

To disassociate an authenticator from a container,

1. In Dashboard's Configuration hierarchy tree, expand **Authenticators**. Expand the sub-nodes in the tree until you find your authenticator. Click its name and click **Edit**. The Authenticator Details window opens.
2. In order to disassociate the authenticator from its parent container:
   - The **Container** field shows the name of the authenticator's container in blue text. Click on the blue text to display the Container Selector window.
   - Navigate the tree and click the container that will hold the authenticator.
   - Click **OK** in the Container Selector window.
3. In order to disassociate the authenticator from an access policy:
   - Click the **RADIUS Settings** tab and in the Access Policy drop-down list, select a different access policy.
   - Click the **TACACS+ Settings** tab and in the Access Policy drop-down list, select a different access policy.
4. Click **OK** to apply your changes.

The authenticator now belongs to a different container and/or access policy, depending on your edits.

Renaming Authenticators



Important! Renaming an existing authenticator or authenticator bundle breaks the authorization rules that depend on that authenticator or authenticator bundle. See "Problem: Authorization policy stops working unexpectedly" on page 453 for troubleshooting instructions.

In order to rename an existing authenticator,

1. In Dashboard's Configuration hierarchy tree, expand **Authenticators**. Expand the sub-nodes in the tree until you find your authenticator. Click its name and click **Edit**. The Authenticator Details window opens.
2. Enter the new **Name** for the authenticator.
3. Click **OK** to apply your change.

Deleting Authenticators



Important! Deleting an existing authenticator or authenticator bundle breaks the authorization rules that depend on that authenticator or authenticator bundle. See "Problem: Authorization policy stops working unexpectedly" on page 453 for troubleshooting instructions.

In order to delete an existing authenticator:

1. In Dashboard's Configuration hierarchy tree, expand Authenticators. Find your authenticator in the tree and click its name.
2. Click the **Delete button**. Alternatively, right click an authenticator and select **Delete**.
3. A confirmation window appears. Make sure you have selected the appropriate authenticator for deletion, and click **OK**.

Ignition Server Dashboard deletes the authenticator from the authenticator hierarchy.

Internal Users, Groups, and Devices

This chapter shows how to store user accounts, device accounts, and group memberships in Avaya Identity Engines Ignition Server's onboard database, the internal store. These users, groups, and devices, called internal users, internal groups, and internal devices, can be used for authentication.

Contents

- "Ignition Server Internal Data Store" on page 113
- "Internal Users" on page 113
- "Internal Devices" on page 119
- "Internal Groups" on page 126

Ignition Server Internal Data Store

Ignition Server uses an onboard database, called the *internal data store*, to manage access information for groups, users and devices. The internal store consists of internal groups, internal users and internal devices.

In a typical installation, most of your user accounts reside in your corporate user directory or directories (see "Directory Services" on page 133), and the internal store acts as a supplementary store that holds other types of user accounts such as temporary accounts. For example, the Avaya Identity Engines Ignition Server Guest Manager application stores its guest accounts as internal users. At login time, Ignition Server treats all users alike, whether they are stored in the Ignition Server or in a corporate directory.

Using the windows shown in this chapter, you can view and edit your internal users, internal groups, and internal device records. You cannot view or edit users stored in other databases such as an LDAP or AD store. To manage such users, use the dedicated user provisioning tools that connect to your LDAP or AD store.

Internal Users

Internal users are user accounts stored locally on the Ignition Server. Users connecting to your network can authenticate against an internal user account in the same way that they can authenticate using an AD- or LDAP-stored account. Internal user accounts are particularly useful for guest users, and guest user accounts created by the option Avaya Identity Engines Ignition Server Guest Manager application are internal user accounts.

Contents of This Section

See the following sections for details on internal users:

- "Internal Users Panel" on page 114
- "Filtering the Internal Users List" on page 116
- "Viewing an Internal User" on page 116
- "Creating an Internal User" on page 116

Internal Users Panel

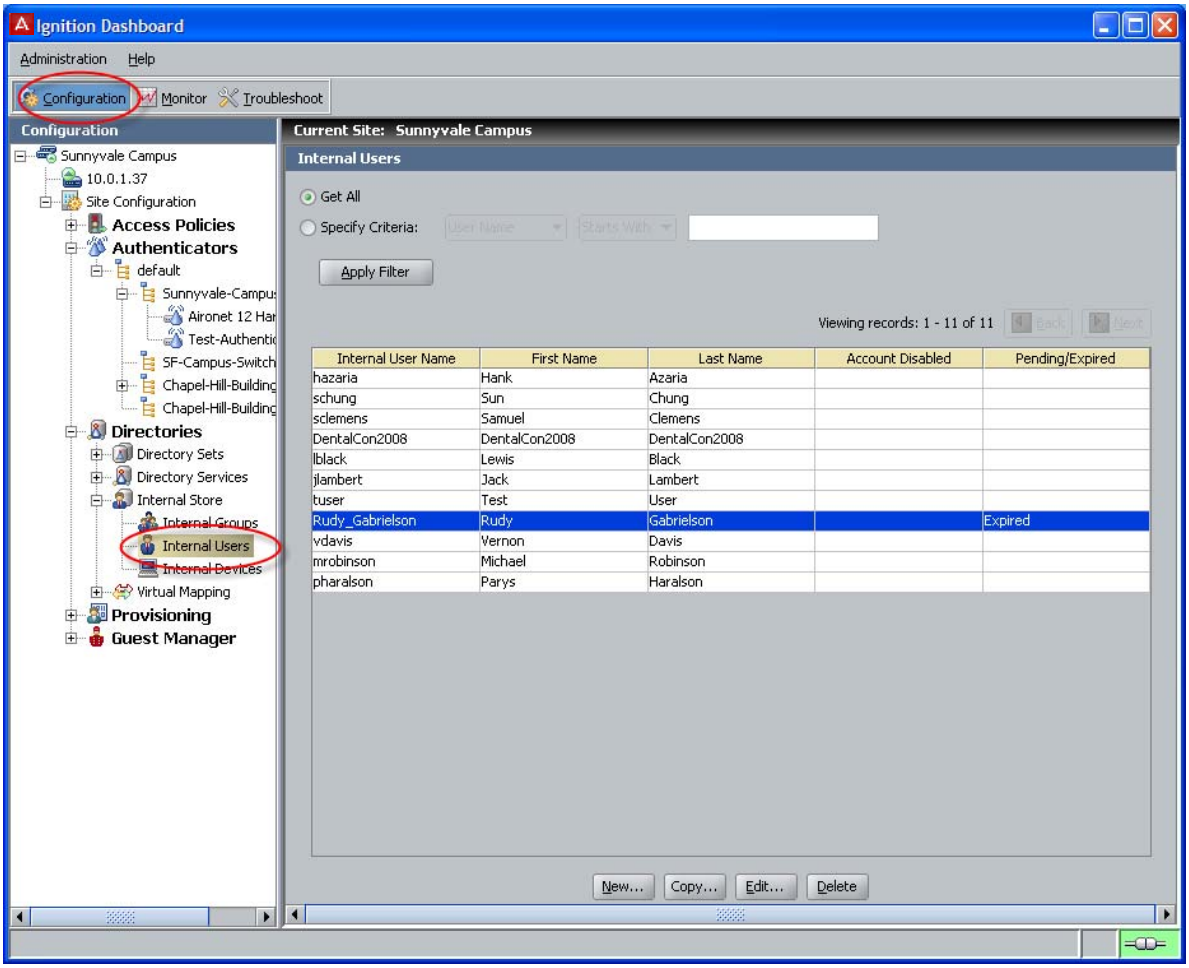
The **Internal Users** panel lists all the user records in the Ignition Server Internal Users store. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, expand **Internal Store**, and click **Internal Users** to open this panel. Here you can

- see all internal users,
- retrieve a subset of all internal users,
- sort and page through internal users, and
- add, edit or delete an internal user.



Tip: To count the number of users in the internal store, go to the main window, select Monitor, click the name of your site, click the Statistics tab, click the Transactions tab, and check the **Embedded DB** section.

Figure 25 Internal Users panel



From this panel you can

- View the list of all users in the internal store.
- Add, edit, copy or delete internal users. To do so, click the appropriate command button at the bottom of the panel.
- Sort the list of users by a particular column. To do so, click the column header, such as User Name; a second click reverses the order from ascending to descending or vice versa. (This feature is common to all windows showing columns.)
- Filter the list of users to reduce the set of users to show only those that fit your search criteria. For information about how to do this, see "Filtering the Internal Users List" on page 116
- Scroll through a long list by page. To do this, click the Next and Back buttons. These are the small, white buttons (each displaying a triangular

arrow icon) near the upper right corner of the user list. Click the right-facing arrow to move forward through the list, and the left-facing arrow to move back.

Filtering the Internal Users List

When the list of Internal Users is long, you can apply a filter to screen unwanted users from the list. Do this as follows:

1. In the Internal Users window, tick the **Specify Criteria** checkbox.
2. Two drop-down lists are shown to the right of the **Specify Criteria** checkbox. In the first list, choose the name of the field you want to filter on. For example, you might choose *First Name*.
3. In the next drop-down list, select the comparison to be performed. Choose *Starts With* or *Equals*.
4. In the text field, type the comparison value.
5. Click **Apply Filter**. Dashboard filters the list.

To view all users again, click **Get All** and click **Apply Filter**.

Viewing an Internal User

To view the complete details of an internal user,

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Users to view the current list of internal users.
2. Click on the desired user entry in the displayed list. Click **Edit** or double-click on the desired user entry in the displayed list.
3. Ignition Dashboard displays the **Edit User** window for the selected user.

The **Edit User window** shows all the data for a selected user. Use this window to review and/edit the selected user record. For a field-by-field description of this window, see “Creating an Internal User”.

Creating an Internal User

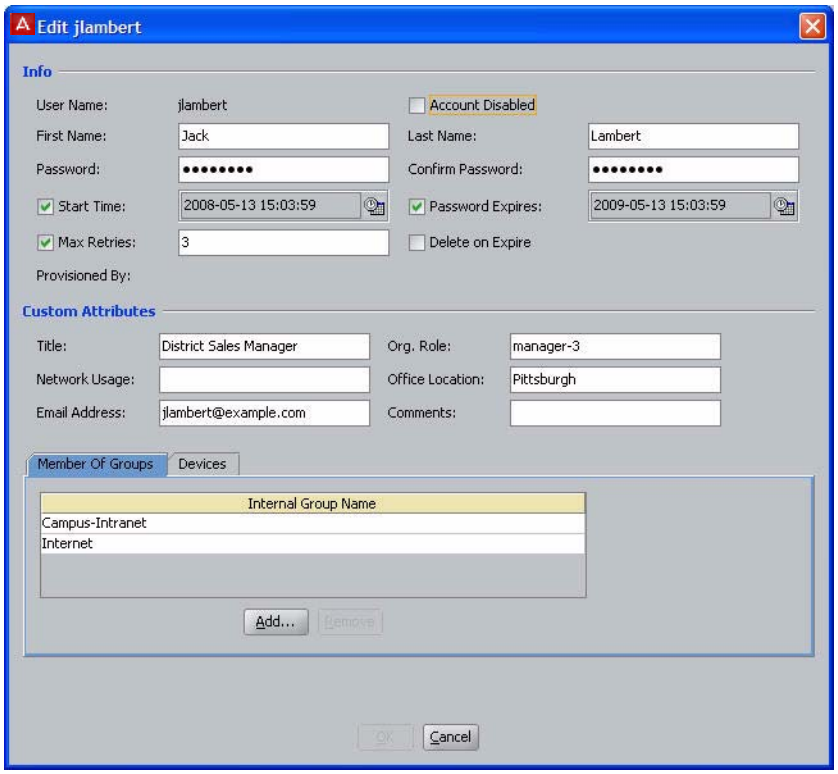
You can create new internal users in two locations of the Ignition Server Dashboard: the internal users store and the internal groups hierarchy.

To create a new internal user:

1. Access the edit user window using one of the following paths:
   - **From the Internal User Store:** In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Users to open this window. Click **New** in the **Internal Users** panel. The Edit window appears.

- **From the Internal Groups hierarchy:** In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Groups to show the Internal Groups hierarchy. In the **Internal Groups** window, use the **Internal Groups** navigation tree to select a group you want the new user to belong. Then, click **New** in the **Users** tab of the **Internal Group Details** panel. The Edit window appears.

Figure 26 Edit User window



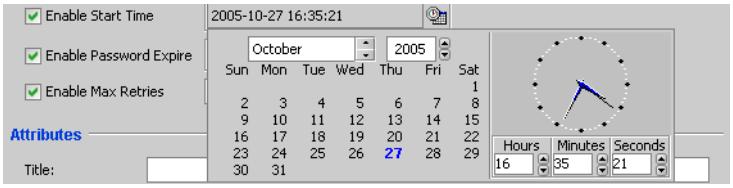
2. In the Edit window, enter the user details and tick the appropriate checkbox settings for the new user account.

   The fields and settings that describe a user are:

   - **User Name, First Name, Last Name:** The login name, given name, and family name of the user, respectively.
   - **Account Disabled:** Select this checkbox if you want to lock this user account. A user account can be intentionally locked by an Administrator or it can be automatically locked by the system, such as when a password expires or when the number of failed authentication attempts exceeds the maximum allowed number of retries.
   - **Start Time:** Select this checkbox if you want to specify when the account is to be activated. Click the clock-and-calendar icon and use

the arrow keys to set the date and time to enable account. Click outside the clock and calendar dialog to close it.

Figure 27 User Account Start Date and Time Selection



- **Enable Password Expire:** Select this checkbox if you want to specify an expiry date for the account. Click the clock-and-calendar icon and use the arrow keys to set the date and time it expires. Click outside the clock and calendar dialog to close it. After an account expires, Ignition Server deletes it if configured to do so in its **Enable Auto Deletion** setting. (See below.)
- **Delete on Expire:** If you want to have Ignition Server automatically delete the account after it expires, check this checkbox. Ignition Server checks hourly for user records in the internal store that have been expired for at least one week. Upon finding such an expired record, Ignition Server checks its **Enable Auto Deletion** setting, and, if the record is set for automatic deletion, deletes it. Deletions take place as time permits. For large sets of records, deletions are spread over a period of hours. Each deletion is logged in the Ignition Server logs.
- **Max Retries:** Select the checkbox and enter the number of failed authentication attempts that can occur in a three-minute period before the account is automatically locked.
- **Custom Attributes:** The lower part of the Edit User window contains a set of **Attributes** fields. You can use these in any way that you like. For example, you can evaluate the values of these fields in your authorization rules, as explained in “User Attributes” on page 233.

  Please note that Avaya Identity Engines Ignition Server Guest Manager uses the **Org. Role** field to label guest users as “guestUser” and provisioners as “provisioner”. In addition, Avaya Identity Engines Ignition Server Guest Manager uses the **Email Address** and **Comments** fields. *If you want to edit or delete a guest user or provisioner account, Avaya strongly recommends that you use Guest Manager to make the change. Using Dashboard to edit guest user and provisioner accounts is not recommended.*
- **Member of Groups Tab:** Lists the groups in which this user is a member. Click **Add** to assign the user to one or more groups. If the desired group does not exist, create it as explained in “Adding a New Internal Group” on page 128. By default, the user is a member of the “Default” group.

- **Devices Tab:** Lists the devices assigned to this user. This is useful if you want to require that a user connect using only his or her assigned device, as explained in “Requiring the User to Connect Using a Machine Authenticated-Device” on page 340.

  To assign a device to a user, click the **Add** button in the **Devices** tab, click the device name in the **Add Device Records** window, and click **OK**. If the desired device record is not in the list, create it as explained in “Creating a Device Record” on page 120.

3. Click Save to save the user account to the Ignition Server internal store.

Internal Devices

A device record (also known as an “internal device”) stores the MAC address (and, optionally, other account details) of a known device that connects to your network. Such devices include printers and fax machines. Device records are stored locally on the Ignition Server. After you have saved your devices as device records, you use them in:

- MAC authentication rules that allow only known devices to connect to the network (see “MAC Authentication” on page 321); and/or
- asset correlation rules requiring that each user sign on to the network using only the device(s) assigned to him or her (see “Asset Correlation” on page 335).



Important! If you plan to authenticate devices using *Windows machine authentication*, no device record in Ignition Server is needed. Instead, your device accounts reside in Active Directory. See Chapter , “Windows Machine Authentication”.

Contents of This Section

See the following sections for details on device records:

- “Finding an Internal Device” on page 120
- “MAC Addresses Are Stored Only in the Internal Store” on page 120
- “Filtering the Device List” on page 120
- “Creating a Device Record” on page 120
- “Assigning a Device to a User or Group” on page 123
- “Editing a Device Record” on page 124
- “Deleting a Device Record” on page 124
- “Importing Device Records” on page 124
- “Exporting Device Records” on page 125
- “Using MAC Address Wildcarding” on page 125

Finding an Internal Device

To find a device record:

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Devices.
2. The Internal Devices panel shows your device records. See “Filtering the Device List” on page 120 for instructions on finding a device in the list.

The Internal Devices panel shows all the devices saved in the Ignition Server internal store. Use the **Back** and **Next** buttons to move through the list.

MAC Addresses Are Stored Only in the Internal Store

In Ignition, you must store the MAC addresses of known devices as device records in the internal data store. Ignition Server cannot be configured to read MAC addresses from an external source such as an LDAP or AD store.

Filtering the Device List

To filter the Internal Devices panel:

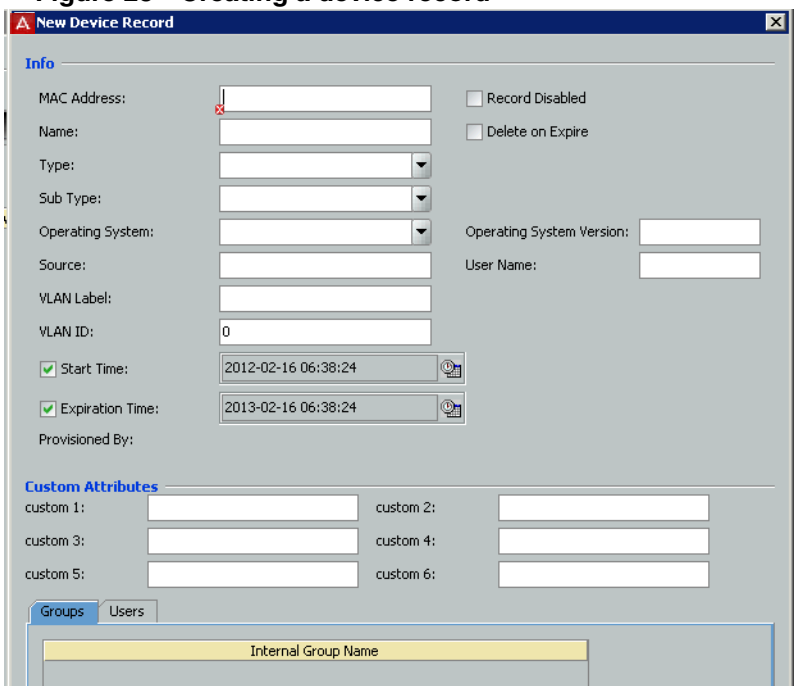
1. In the Internal Devices panel window, click the Specify Criteria button.
2. Two drop-down lists are shown to the right of the “Specify Criteria” checkbox. In the first list, choose the name of the field you want to filter on. For example, you might choose MAC address, Name, Type, or Source.
3. In the next drop-down list, select the comparison to be performed. You choose “Starts With” or “Equals.”
4. In the text field, type the comparison value.
5. Click Apply Filter. Dashboard filters the list.

   To view all devices again, click Get All and click Apply Filter.

Creating a Device Record

Use the following steps to create a device record in Ignition. (If you need to create *many* device definitions, you may prefer to create them in bulk as shown in “Importing Device Records” on page 124.)

Figure 28 Creating a device record



1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand **Directories**, expand **Internal Store**, and click **Internal Devices**. Click **New**.
2. In the **MAC Address** field, specify the MAC address of the device. Enter the address as a string of six octets. You may write the twelve characters without separators, or you may separate the octets with period, colon, or hyphen characters. Do not mix separator characters.
3. If you want to disallow this device from connecting to the network, tick the **Record Disabled** checkbox.
4. In the **Name** field, type a name for the device. This name identifies the device in logs and when you associate it with w group or user.
5. If you want Ignition Server to delete this record automatically after its expiration date, select the **Delete on Expire** checkbox. Ignition Server checks hourly for device records in the internal store that have been expired for at least one week. Upon finding such an expired record, Ignition Server checks its **Enable Auto Deletion** setting, and, if the record is set for automatic deletion, deletes it. Deletions take place as time permits. For large sets of records, deletions are spread over a period of hours. Each deletion is logged in the Ignition Server logs.
6. In the **Type** drop-down list, designate what sort of device this is, such as a laptop, printer, or handheld device. You can choose one of the preset values or type your own value.

7. In the **Sub Type** drop-down list, define the details of the device from one of the preset values. For example, if you chose **mobile** as your device **Type**, you can define the **Sub Type** as iphone, blackberry, or android phone etc.
8. In the **Operating System** drop-down list, select the operating system of the device. You can choose one of the preset values.
9. In the **Operating System Version** field, enter the version of the operating system.
10. In the **User Name** field, enter the name of the user of this device.
11. The **Source** field is typically used only for bulk-imported device records (see "Importing Device Records" on page 124). The **Source** indicates the origin of this record. Usually this is the name of the file from which the device record was imported.
12. If you want to have Ignition Server automatically assign this device to a VLAN, enter the VLAN name in the **VLAN Label** field and enter the integer VLAN number in the **VLAN ID** field. If you do not want to assign it to a VLAN, leave these fields blank.
13. Select the **Start Time** checkbox if you want to specify when the account is to be activated. Click the clock-and-calendar icon and use the arrow keys to set the date and time to enable the account. Click outside the clock and calendar dialog to close it.
14. Select the **Expiration Time** checkbox if you want to specify an expiry date for the device record. Click the clock-and-calendar icon and use the arrow keys to set the date and time it expires. Click outside the clock and calendar dialog to close it. When an account expires, Ignition Server may delete it, depending on the **Delete on Expire** setting. (See above.)
15. The **Custom Attributes** fields allow you to record additional information about the device. See "Adding Virtual Attributes for Devices" on page 208.
16. Click **Save** to store the device record.

Next steps: Do one of the following:

- If this device is to be permitted to sign on via MAC authentication (bypassing 802.1X), then make sure you have a MAC authorization policy that applies to it. See "Example MAC Authentication Set-Up Procedure" on page 325.
- If this device is to be assigned to a user in order to enforce an asset correlation policy, turn to "Assigning a Device to a User or Group" on page 123.

Assigning a Device to a User or Group

You may enforce Windows machine authentication/asset correlation policies that allow users to connect only with the device assigned to them. To support such a policy, you must create a device record for each user's computer and assign the device to the user or user group. To create a device record, see "Creating a Device Record" on page 120. To assign a device to a user or group, do the following:

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Devices. The Internal Devices panel shows all the devices saved in the Ignition Server internal store. Use the Back and Next buttons to move through the list.
2. In the list, find the device record and double-click it. (Or click it and click the Edit button.)
3. In the Device Record window do one of the following:
   - To assign this device to a user, click the Users tab, and then click the Add button in the tab. Scroll or use the filter to find the user, click the user's name, and click OK.
   - To assign this device to a group, click the Groups tab, and then click the Add button in the tab. Scroll or use the filter to find the group or groups, tick the checkbox for each group that can use the device, and click OK.
4. In the Device Record Details window, click Save.
5. Create your policy to enforce your assigned-device-only policy, as shown in "Creating Asset Correlation Policies" on page 336.

Note that you can also assign devices to users and groups from the user or group record:

- To assign a device to a user: In the Configuration hierarchy tree, expand Directories, expand Internal Store, and click Internal Users. Double-click the name of the user. In the Edit User window, click the Devices tab. Click Add in the tab. Click on the desired device and click OK. Click Save in the Edit User window.
- To assign a device in the Internal Groups window do this: In the Configuration hierarchy tree, expand Directories, expand Internal Store, and click Internal Groups. Click the name of the group. Click the Devices tab. Click **Add Existing** in the tab. Click on the desired device and click OK.

Editing a Device Record

To edit a device record, In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Devices. In the Internal Devices panel click the name of the device record and click the **Edit** button. The Device Record Details window displays the record and allows you to edit it. For information on using this window, see "Creating a Device Record" on page 120.

Deleting a Device Record

To delete a device record, go to Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Devices. In the Internal Devices panel click the name of the device record and click the **Delete** button. Ignition Server deletes the record.

Importing Device Records

To import a set of device records, follow the steps below:

1. Using a tool of your choice, create a file containing the device records you want to import. The file must be in the comma-separated values (CSV) format. Ignition Server requires that you provide the fields in this order:

```
MAC Address,Name,Type,VLAN Label,VLAN ID,Custom 1,Custom 2,Custom 3,Custom 4,Custom 5,Custom
6,"Group Name 1, Group Name 2",Account Disabled
```

   where "Account Disabled" is replaced with a value of "yes" to indicate the device is *not* allowed to connect, or "no" to indicate it is allowed to connect. For example, if you import a file containing the line,

```
A8139C62A7BD,HP-Laserjet-Floor3,printer,hq-printer-vlan,206,,,,,,"default, printers-in-HQ",no
```

   the import action adds a device record to Ignition Server with a MAC address of *a8:13:9c:62:a7:bd*, a name of *HP-Laserjet-Floor3*, a type of *printer*, a VLAN label of *hq-printer-vlan*, a VLAN ID of *206*, no custom values, membership in the groups *default* and *printers-in-HQ*, and the **Record Disabled** checkbox set to unticked. Make sure the groups exist already in Ignition.

2. In Dashboard's Configuration hierarchy tree, expand Directories, expand Internal Store, and click Internal Devices.
3. In the Internal Devices panel, click **Import**.
4. In the Device Record Import window, click the **Browse** button, navigate to find your csv file, click it, and click **Open**.
5. In the Device Record Import window, the **Source** field is used to indicate the origin of the device records you are importing. By default, the **Source** field displays the name of your csv file. You may edit it. This name is saved as the Source attribute in each device record.

6. If the first line of your csv file contains column headings or other unneeded information, tick the **Skip the first line** checkbox.
7. Click **OK** to import the records. The Internal Devices panel displays your imported devices.

Exporting Device Records

To export a set of device records, do the following:

8. In Dashboard's Configuration hierarchy tree, expand Directories, expand Internal Store, and click Internal Devices.
9. In the Internal Devices panel, click **Export**.
10. In the Device Record Export window, do one of the following:
    - To export all device records, tick the **Get All** radio button.
    - To export some device records, tick the **Specify Criteria** radio button and set your filter criteria in the fields just to the right. In the first drop down list, select the name of attribute you want to filter on. In the second drop down list, select **Starts With** to export those records in which the filter attribute's value matches the first few characters of your search string, or select **Equals** to export only those whose attribute is identical to the whole search string. Type the search string in the field at the right.
11. Click the **Browse** button, navigate to find the directory in which you want to save your csv file, double-click the directory name to select it, type a name for the csv file in the **File Name** field, and click **Save**.
12. In the Device Record Export window, click **OK** to export the records.

Using MAC Address Wildcarding

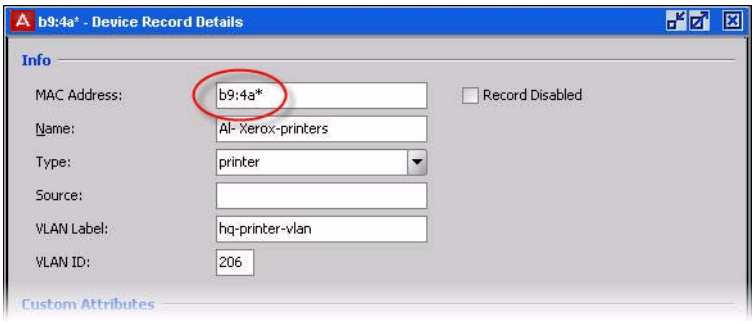
For organizations that provide many users with similar laptops (or other devices) and want to ensure that those users can only log on using the assigned type of laptop, Ignition Server offers a shortcut: MAC address wildcarding. MAC address wildcarding lets you create a single device record that represents a number of devices of the same type. After you have applied this device record to all users who use that type of device, you can write an asset correlation policy that compares the user's MAC address with the partial, wildcarded MAC address in the Device Record. If the partial address matches, the user is allowed to connect.

To do this, use the MAC address wildcarding feature, define your device as usual in the Device Record window, but specify a partial MAC address followed by a "*" character.

For example, if all of your employee laptops connect using an Ethernet card with a MAC address that begins with "b9:4a," then you can set your device address to "b9:4a*" in the Device Record window. To create a policy that

ensures your employees can only connect using their company-provided laptops, assign the Device Record to each user who uses this type of laptop, and create an asset correlation policy that verifies the user is using an "Assigned Asset." See "Requiring the User to Connect Using His or Her Assigned Device" on page 338 for details on asset correlation policies.

Figure 29 MAC address wildcarding



Internal Groups

Ignition Server allows you to collect internal users and internal devices into groups in order to apply policies to groups. Any user or device can be a member of more than one internal group.

Contents of This Section

See the following sections for details on internal groups:

- "Internal Groups Panel" on page 126
- "Working with the Internal Groups Hierarchy" on page 127
- "About the Default User Group" on page 128
- "Adding a New Internal Group" on page 128
- "Moving an Internal Group in the Hierarchy" on page 129
- "Renaming an Internal Group" on page 129
- "Changing a Group's Type Designation" on page 129
- "Deleting an Internal Group" on page 130
- "Working with Internal Group Details" on page 131

Internal Groups Panel

Internal groups are managed in Ignition Server from the internal data store window.

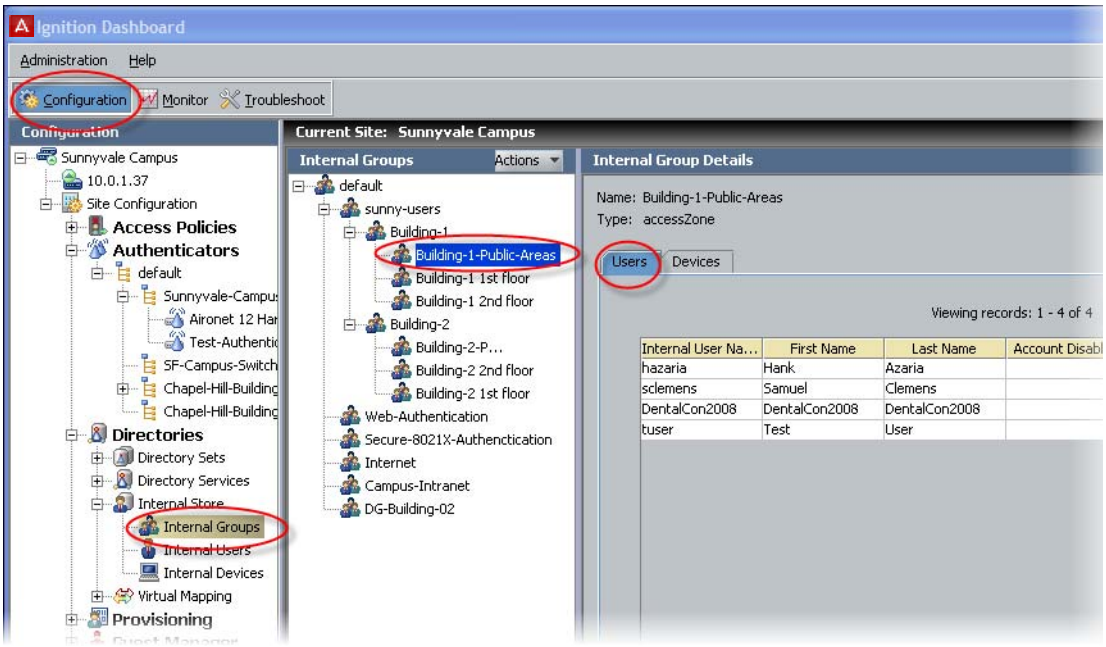
This window consists of

- **Internal Groups panel:** The hierarchy tree shows you the relationships between groups and allows you to add new groups, edit existing groups, and modify the group hierarchy.
- **Internal Group Details panel:** This panel displays the details of a selected internal group. Internal Groups can contain both users and devices. The Internal Group Details lists the group's information on the following tabs:
  - Users Tab: Lists the current users of the selected group. In the **Users** tab, you can add an existing user and edit or delete any individual user in the selected group. You can also create a new user and add that user to the selected group at the same time.
  - Devices Tab: Lists the current devices of the selected group. In the **Devices** tab, you can add an existing device to the group and edit or delete any device in the selected group. You can also create a new device and add that device to the selected group at the same time.

Working with the Internal Groups Hierarchy

Internal groups are organized hierarchically and are displayed in the panel to the right of the Configuration panel. The root of the internal groups hierarchy is the *default* group. This group is fixed in the hierarchy and cannot be deleted or renamed. All other groups are subordinate to the default group.

Figure 30 Clicking an internal group in the hierarchy



The **Actions** button at the top of the Internal Groups panel is a pull-down menu of commands that let you create and manage groups. Using this button, you can perform the following actions:

- Add a new internal group
- Move an internal group
- Rename an internal group
- Edit an internal group type
- Delete an internal group



Tip: To count the number of groups in the internal store, go to the main window, select Monitor, click the name of your site, click the Statistics tab, click the Transactions tab, and check the **Embedded DB** section.

About the Default User Group

The **internal data store** includes a default group, which is the “root” group in the internal groups hierarchy. It cannot be a member of any other group.

You can add new or existing users or devices to this group. However, as shown above, the Ignition Dashboard does not permit this default group to be renamed or to be deleted.

Adding a New Internal Group

To add a new internal group:

1. In the Dashboard Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Groups.
2. In the Internal Groups hierarchy tree, click the parent group to select it, or click “Default” to place your new group at the top level. The new group becomes a child of the selected internal group.
3. **Select Actions → Add A New Internal Group…**
4. In the Add a New Internal Group window, enter the new group’s name.
5. Select the group’s **Type**. Type designations are used for Avaya Identity Engines Ignition Server Guest Manager. See the *Avaya Identity Engines Ignition Server Guest Manager Configuration Guide* for details.
6. Click **OK**.

Ignition Dashboard adds the new group. This group now appears as a child of the selected group in the Internal Groups hierarchy.

Moving an Internal Group in the Hierarchy

You can move an existing group so that it is subordinate to a different group in the hierarchy. Do this as follows:

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Groups.
2. In the Internal Groups hierarchy, select the group that you want to move. Click on it to select it.
3. Choose **Actions → Move Internal Groups…**
4. In the Select Group window, select a new parent for the group to be moved.
5. Click **OK** to apply your changes.

Dashboard displays the new internal groups hierarchy.

Renaming an Internal Group

To rename an existing internal group:

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Groups.
2. In the **Internal Groups** hierarchy, click on the internal group you want to rename.
3. Select **Actions → Rename Internal Group...**
4. Highlight the entry in the **Rename** window. Enter a new name for the group.
5. Click **OK**.

Ignition Dashboard displays the new name for the group in the Internal Groups hierarchy.

Changing a Group's Type Designation

Group type designations are used for Avaya Identity Engines Ignition Server Guest Manager. See the *Avaya Identity Engines Ignition Server Guest Manager Configuration Guide* for details. To change the group type:

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Groups.
2. In the Internal Groups hierarchy, click the group to select it.
3. **Select Actions → Edit Internal Group Type**.

4. Type or select the new group type.
5. Click **OK**.

Deleting an Internal Group

If an internal group maps to any virtual groups, there are several steps you must perform before deleting the internal group itself.

To delete an existing group:

1. If your group maps to a virtual group, do this before you delete the group:
   - Find the name of the virtual group.
   - Find any rules that use that virtual group, and edit or delete them so that no rules reference the virtual group. You can do this by opening your Ignition Server RADIUS policies (in Dashboard's Configuration hierarchy, expand Access Policies, expand RADIUS, and click your policy name) and checking to make sure the rules do not reference the virtual group.
   - Delete the virtual group.
2. If your group is used in your Avaya Identity Engines Ignition Server Guest Manager policies, do this before you delete the group:
   - Run Guest Manager, and find all users that belong to (that is, "have rights to") the group. You can recognize such a user record in Guest Manager because its Guest User record contains a ticked checkbox with the name of the group. Delete these users.
   - Find all provisioning groups that belong to (that is, "have rights to") the group. You can recognize such a provisioning group record in Guest Manager because it displays a ticked checkbox next to the name of the group. Delete these provisioning groups.
3. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Groups.
4. In the **Internal Groups** hierarchy, click on the internal group you want to delete. Remove all Users and Devices from the group:
   - Click the Users tab. You must remove all users from the Group. Click each user row and click Delete.
   - Click the Devices tab. You must remove all devices from the Group. Click each user row and click Delete.
5. Finally, you are ready to delete the group. Click **Actions → Delete Internal Group…**
6. Click **Yes** to delete the selected group.

For more information about virtual groups, see Chapter , “Virtual Groups and Attributes”.

Working with Internal Group Details

The **Internal Group Details** panel displays the members of the internal group selected in the Internal Groups hierarchy. This panel lists either the users or devices in the selected group, depending on the tab selected. This panel includes two tabs, the Users tab and the Devices tab.

Users Tab

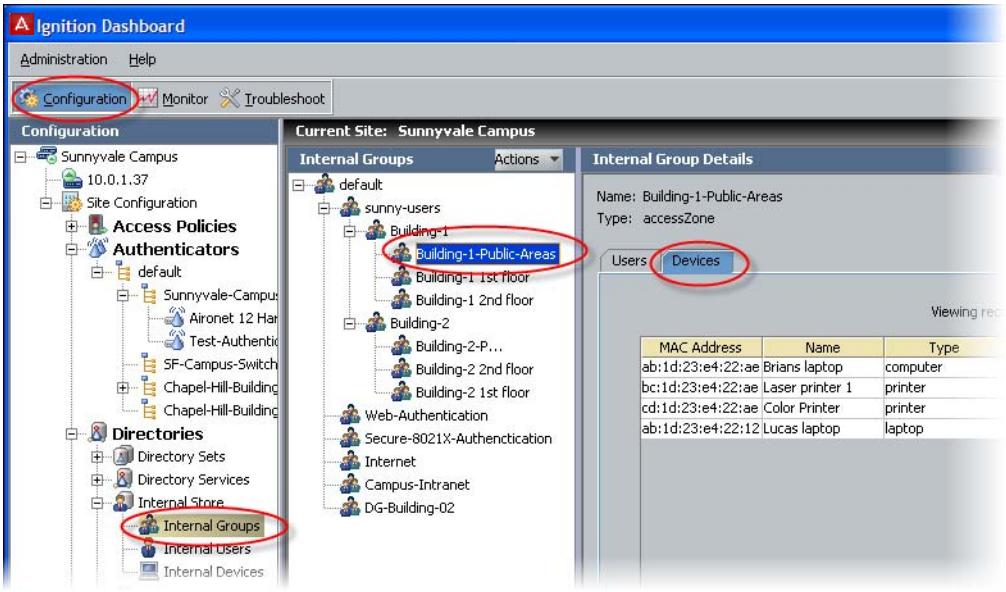
The Users tab lists the users in the group selected in the Internal Groups hierarchy panel. From this tab you can use the following command buttons to add, modify or remove a user from the selected internal group:

- **New:** Lets you create a new user in the group. The button launches the **Internal Users Details** window.
- **Add Existing:** Lets you add an existing user to the group. This button displays the **Add User Records To...** window. To add a user or users, select one or more rows (use Shift-click or Control-click to select more than one row), and click **OK**.
- **Edit:** Lets you view and edit the selected user in the Internal Users details window.
- **Remove:** Lets you remove the selected user from the group.

Devices Tab

The Devices tab lists all the devices in the group selected in the **Internal Groups** hierarchy. From this tab, you can use the command buttons to add, edit or remove devices from the selected group.

Figure 31 Devices in the Building-1-Public-Areas Group



The command buttons are:

- **New:** The **New** button lets you add a device that has not already been added to the Internal Store. You can assign the new device to any of the existing groups.
- **Add:** The **Add Existing** button lets you add devices that have already been created in the Internal Store.
- Edit: The Edit button lets you edit a device that is already a member of the group.
- **Remove**: The **Remove** button lets you remove a device from a selected group.

Directory Services

This chapter explains how to define directory services and assemble them into the directory sets Avaya Identity Engines Ignition Server uses to locate a user account at authentication time.

Contents

- "Quickstart: Directory Services in Dashboard" on page 133
- "Introduction to Directory Services" on page 134
- "Directory Services" on page 134
- "Connecting to Active Directory" on page 136
- "Connecting to an LDAP Service" on page 148
- "Creating a Token Authentication Service" on page 155
- "Creating a Kerberos Authentication Service" on page 156
- "Setting up MSCHAPv2 Authentication on LDAP" on page 156
- "Creating a Radius Proxy Authentication Service" on page 161
- "Managing Directory Services" on page 161
- "Directory Sets" on page 164
- "Troubleshooting User Lookup and Authentication" on page 170

Quickstart: Directory Services in Dashboard

Three tabs in Dashboard allow you to operate on your directory services:

1. Dashboard's **Configuration** view lets you connect to a directory service. Click Configuration at the top of the Dashboard window. In the hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Directory Services. Click on **Directory Services** for an overview. In the tree, click on the name of a service for details about that service. See "Commands That Operate on Directory Services" on page 136.
2. Dashboard's **Monitor** view lets you check the connection and cache status of your service. Click Monitor at the top of the Dashboard window. In the hierarchy tree, click your *node*, and click the **Directory Services Status** tab. To see what directories have been servicing authentication requests, click your *site* in the tree and click the RADIUS AAA summary tab. See "Checking Directory Service Connections" on page 170.

3. Dashboard's **Troubleshoot** view lets you test user authentications and user lookups. Click Troubleshoot at the top of the Dashboard window. In the hierarchy tree, click your node, and click the **Directory Service Debugger** tab. See "Troubleshooting User Lookup and Authentication" on page 170.

Introduction to Directory Services

Ignition Server authenticates and authorizes (looks up) users and devices against a directory service such as an Active Directory ("AD") service, an LDAP directory, or Ignition's internal data store. You define each user data store as a **directory service** in Ignition Server and group the directory services (along with, optionally, authentication-only services) into a **directory set**. Depending on the fallthrough configuration of your directory set, Ignition Server may search all the services in the set in its attempt to authenticate the user.

At authentication time, Ignition Server chooses which directory set to use based on the **identity routing policy** governing the switch or access point the user is connecting to. The identity routing policy lets Ignition Server choose the directory set based on which authenticator originated the access request (the Cisco 3750 switch on the third floor, for example), or based on the realm of the connecting user ("company.com," for example), or based on both authenticator and user. For details see "How Ignition Server Looks Up a User for Authentication and Authorization" on page 221.

If you use a specialized form of authentication such as RSA SecurID, Kerberos, or a Radius proxy server, you must also configure one or more **authentication services** in Ignition. In Ignition, you manage authentication services in the Directory Services panel, in the same way you manage directory services.

When you put an authentication service in your Ignition Server policy, the authentication service is responsible for verifying user credentials, while an optional directory service (called a *user lookup service* in this context) is responsible for retrieving the user attributes and group memberships that Ignition Server uses to authorize the user.

After you have set up directory services, authentication servers (if necessary), and directory sets, you create identity routing policies as explained in the section, "Understanding Identity Routing Policy" on page 220.

Directory Services

A directory service establishes Ignition's connection to a user repository or authentication server. The directory service may function as a user lookup service, authentication service, or both. How and when a directory service is used is determined by its position in a directory set. (Directory sets are explained in "Directory Sets" on page 164.)

Supported Directory Servers

Ignition Server supports the following directory server types:

- Microsoft Active Directory
- Sun Java System Directory Server
- Novell eDirectory
- Oracle Internet Directory (OID)
- Generic LDAP
- Kerberos Server
- token services, such as the RSA Authentication Manager
- Radius Proxy Service

For a list of which authentication protocols are supported using which directory servers, see “Supported Authentication Types” on page 215.

Internal Data Store as a Directory Service

Ignition Server treats the local internal data store as a directory service. You can use the internal data store wherever a directory service can be used. You can handle the internal store just as you handle AD and LDAP stores, using directory services and sets.



Ignition Server does not allow changes to the configuration information in the internal data store, and it does not allow the store to be deleted. To view the users it contains, and their attributes, see Chapter , “Internal Users, Groups, and Devices”.

Note that when you select the internal data store in the Directory Services Status tab, Ignition Server does not enable the Refresh Cache, Edit, and Delete buttons because you cannot perform these actions on the internal data store. When you select other directory services, Ignition Server enables the command buttons.

Working With Directory Services

The Directory Services panel allows you to view and manage your directory service connections. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, and click Directory Services. The window lists the existing directory services, their service types. Use the **Directory Services panel** to add, edit, and delete a directory service.

Configuration View: The Directory Services Panel

The columns of the Directory Services panel are:

- **Name** is the directory service name you have given to the data store.

- **Directory Type** is the family of LDAP, AD, or database store to which this service belongs. “Internal Database” denotes the on board Ignition Server database.

Commands That Operate on Directory Services

As explained in “Quickstart: Directory Services in Dashboard” on page 133, there are three views in Dashboard that let you operate on your directory services:

Configuration View The Configuration view of Dashboard provides the commands that configure your directory services. The commands are

- **New** lets you create a new directory service.
- **Edit**: To edit an existing service, click its name in the table and click **Edit**.
- **Delete**: To delete an existing service, click its name in the table and click **Delete**.

Monitor View The Monitor view of Dashboard lets you check the status of your directory services. Its commands are explained in “Checking Directory Service Connections” on page 170.

Troubleshoot View Dashboard’s **Troubleshoot** view lets you test user authentications and lookups. See “Troubleshooting User Lookup and Authentication” on page 170.

Connecting to Active Directory

This section explains how to connect to Active Directory. This section consists of:

- “AD Connection Settings” on page 136
- “Preparing to Connect to Active Directory” on page 139
- “Creating an Active Directory Service: Automatically Configuring” on page 140
- “Creating an Active Directory Service: Manually Configuring” on page 143

AD Connection Settings

The following table describes the parameters Ignition Server uses to connect to an Active Directory service. You make these settings in the Create Directory Service Wizard or in the Directory Server Details window.

Gather this information for each store to be used for authenticating users. Talk to your AD administrator to find the connection settings for each AD data store. Record your settings in the table that follows.

Table 3 Settings for connecting to an AD store

Setting Name	Setting Value
AD Domain Name	_____
The **AD Domain Name** specifies the Active Directory domain that holds your user accounts. Domain names typically carry a domain suffix like ".COM" as in, for example, "COMPANY.COM".	
Service Account Name	_____
The **Service Account Name** is the name of the AD administrator account that the Ignition Server uses to connect to the AD server. In the documentation, we refer to this account as the *Ignition Server service account*. If you want to perform MSCHAPv2 authentication, the service account must have permission to create and delete computer accounts (the *Create Computer Object* and *Delete Computer Object* permissions) in the *Netlogon account root* in Active Directory. See "Netlogon account root DN," below. If you have not specified a Netlogon account root DN in Ignition, then the service account must have these permissions in the *Computers container* of your AD service. Ignition Server uses the service account to join the Active Directory domain. Joining the domain requires creating a machine account in the *Netlogon account root* and periodically resetting the password on that account for security. The machine account itself is necessary to perform Netlogon authentication requests for MSCHAPv2 traffic to Active Directory. **Note:** Make sure that the name you enter here is the sAMAccountName of the administrator. The sAMAccountName is usually the user id of the user without the domain prefix. For example, the sAMAccountName for the user *COMPANY.COM/Administrator* is usually *Administrator*. Creating a service account: If no appropriate account exists in your AD installation, see "Create the Service Account in AD" in the *Ignition Server Getting Started Guide*. For help setting its permissions, see "Set the AD Permissions of the Service Account" in the *Ignition Server Getting Started Guide*.	
Service Account Password	
The **Service Account Password** is the password for the AD service account. *Do not record the password here.*	
Security Protocol	**Simple or SSL**
The **Security Protocol** setting specifies whether Ignition Server should SSL-encrypt traffic to the directory service. Avaya recommends that you use an SSL connection. Warning! If you connect using a non-SSL connection, your service account credentials travel unencrypted.	
IP Address (Primary/Secondary)	_____
The **IP Addresses** of the primary and secondary AD data stores.	
Port (Primary and Secondary)	_____

Table 3 Settings for connecting to an AD store

Setting Name	Setting Value
The LDAP **Port** of the primary and secondary AD data stores. For SSL enter 636. If SSL is not used, enter 389. You *cannot* use the global catalog port (3268). *Please use the LDAP ports (389 and 636) only!*	
Name	
The **Name** is a name you use in Ignition Server to identify this AD data store. This can be any name. You can use this name in your authorization policy. See "Inbound Attributes" on page 234.	
NetBIOS Domain	
The **NetBIOS Domain** name (pre-Windows 2000 domain name) of your AD data store. This setting is typically written in all uppercase letters, as in, "COMPANY". This setting applies only to *Active Directory* stores. For instructions on using Microsoft tools to find this name, see "Looking Up AD Settings: Finding Domain and NetBIOS Names" on page 147.	
NetBIOS Server Name (Primary and Secondary)	
The **NetBIOS Server Name** is optional. It allows Ignition Server to find the NETBIOS server where Ignition Server performs the Netlogon (a prerequisite to performing MSCHAPv2 authentication). If the **NETBIOS Server Name** is not specified, then Ignition Server relies on DNS to find the NETBIOS server. Avaya strongly recommends that you specify a **NETBIOS Server Name** to ensure that MSCHAPv2 authentication can continue when the DNS server is unavailable. The directory service set-up wizard helps you determine the NETBIOS server name by retrieving a list of domain controllers in the domain.	
Directory Root DN	
The **Directory Root DN** is the root of the AD tree containing your groups and schema, expressed using X.500 naming. For example, `dc=company,dc=com`. When you connect the directory service, the Ignition Server Create Service wizard attempts to choose a Directory Root DN for you. See "Looking Up AD Settings: Finding Your Root DNs" on page 145 for information on finding this DN.	
User Root DN	
The **User Root DN** specified the AD container that holds your user records, expressed using X.500 naming. For example, `cn=users,dc=company,dc=com` or `ou=uswest,ou=americas,dc=company,dc=com`. When you connect the directory service, the Ignition Server Create Service wizard attempts to choose a User Root DN for you. See "Looking Up AD Settings: Finding Your Root DNs" on page 145 for information on finding this DN.	
Netlogon Account Root DN	

Table 3 Settings for connecting to an AD store

Setting Name	Setting Value
The **Netlogon Account Root DN** is the container in AD where the Ignition Server creates its own machine account when joining the AD domain. This setting is optional. If specified, Ignition Server only attempts to create its machine account in the specified location. If left unspecified, Ignition Server obtains the Netlogon account root DN from the domain controller. Specifically, Ignition Server gets the DN of the *well known computer root* from the DC and uses that as the Netlogon account root DN. The Netlogon account root DN is typically the Active Directory Computers container (by default, this has a DN similar to *cn=computers,dc=company,dc=com*). The machine account is required so that Ignition Server can perform Netlogon authentication requests for MSCHAPv2 traffic to AD. If you want to perform MSCHAPv2 authentication, then your service account must have appropriate permissions in this DN. For help setting account permissions, see "Set the AD Permissions of the Service Account" in the *Ignition Server Getting Started Guide*.	

Preparing to Connect to Active Directory

If your directory service is an Active Directory server, perform the following steps before attempting to connect.



Warning. If you plan to use MSCHAPv2 authentication, you *must* perform the checks listed here.

1. **Gather your AD connection settings** as explained in "AD Connection Settings" on page 136.
2. **Check your clock settings.** When the Ignition Server connects to an Active Directory server, the Ignition Server clock must be in sync with the clock on the Active Directory Server. If the clocks are out of sync, then the Ignition Server cannot connect to the Active Directory store.
3. **Check your firewall settings.** If a firewall protects your Active Directory server, make sure it does not block the ports required by Ignition. Ignition Server needs access to the following ports: 88 (UDP), 389 (TCP), 445 (TCP), 464 (UDP), 636 (TCP).

   After you change the settings on the firewall protecting your Active Directory server, you must reboot your Ignition Server.
4. **Check your Active Directory security settings.** Ignition Server works with all default installations of AD, but if you have adjusted your AD installation to prohibit NTLMv1 authentication, then Ignition Server cannot perform MSCHAPv2 authentication.

   To make sure NTMLv1 authentication is enabled in your AD installation, check the following two settings in the Windows registry of your Windows domain controller (DC). Use the Windows *regedit* tool to do this.

   - Make sure that the following key is *not* set on the DC:
     `HKLM\System\CurrentControlSet\LSA\DisallowMsvChapv2`

- Make sure that the following key is set to a value of 1, 2, 3, or 4. A setting of 5 causes Ignition's support for MSCHAPv2 authentication to fail in all cases. The key name is `HKLM\System\CurrentControlSet\Control\LSA\LMCompatibilityLevel`

5. **Find or create your service account.** Make sure you have a user account in AD that can act as the Ignition Server Service Account. If you need to create a new account, follow the instructions in the section "Create the Service Account in AD," in the *Ignition Server Getting Started Guide*.
6. **Set permissions on your service account.** If you want to perform MSCHAPv2 authentication, make sure your Ignition Server Service Account has, at a minimum, permission to create and delete computer accounts in the Netlogon account root of AD. If you need set this up, follow the instructions in the section, "Set the AD Permissions of the Service Account," in the *Ignition Server Getting Started Guide*.
7. **Optional: Check your machine authentication settings.** If your organization's security policy requires a script to run on each client before that client is allowed to connect, then do the following:
   - Make sure all client machine names are saved in the correct location in AD, which is typically under "cn=computers, ...".
   - Make sure this location is set in Ignition Server as the User Root DN or any container above that in the directory tree.
8. **Recommended: Make DNS settings on Ignition.** If your site uses MSCHAPv2 authentication, Avaya strongly recommends that you configure your Ignition Server appliance's *DNS settings* so that Ignition Server can resolve the address of your AD server.

   To check and edit your DNS settings, go to Dashboard's Configuration hierarchy tree, click the name of your node, then click the **System Tab**, and click the **DNS** tab. Click **Edit**. You can check and edit the addresses of your DNS servers in the **Edit DNS Configuration** window.

Creating an Active Directory Service: Automatically Configuring

The Create Directory Service Wizard guides you through the steps needed to connect Ignition Server to your directory service. Use the following steps to connect Ignition Server to an Active Directory service.

The following instructions show how to use the *automatic mode* of the wizard, which retrieves certain settings for you. For instructions on using the manual mode, see "Creating an Active Directory Service: Manually Configuring" on page 143.



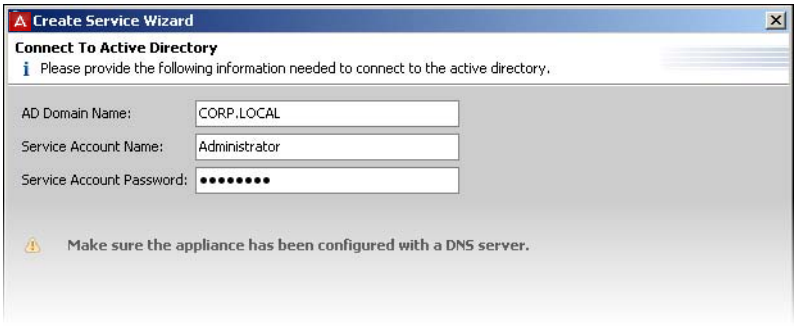
Before you begin:

1. Make sure your DNS server addresses, clock setting, and firewall settings had been completed as explained in “Preparing to Connect to Active Directory” on page 139.
2. Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

To connect to AD (automatic mode):

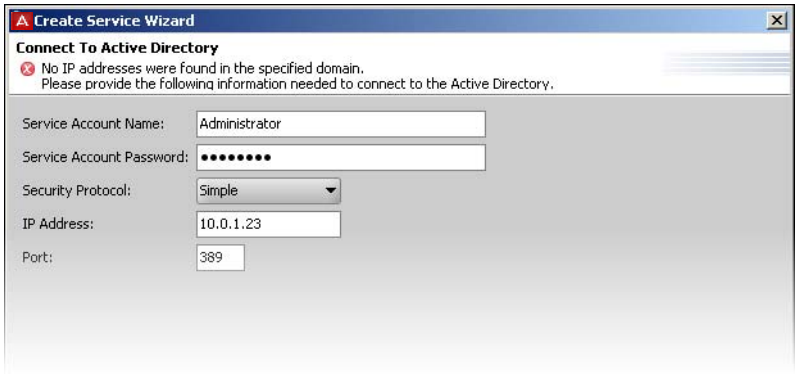
1. In Dashboard’s **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.
2. Ignition Server launches the Create Directory Service Wizard. The **Choose Service Type** window appears.

   Select the **Active Directory** radio button and click **Next**. The **Active Directory Configuration Options** window appears.
3. Select the **Automatically Configure** radio button. Click **Next**.
4. In the Connect To Active Directory window, enter the **AD Domain Name, Service Account Name** and **Service Account Password**. If you plan to support MSCHAPv2 authentication, make sure the service account has sufficient permissions, as explained in “Service Account Name” on page 137. After you type these settings, click **Next**:

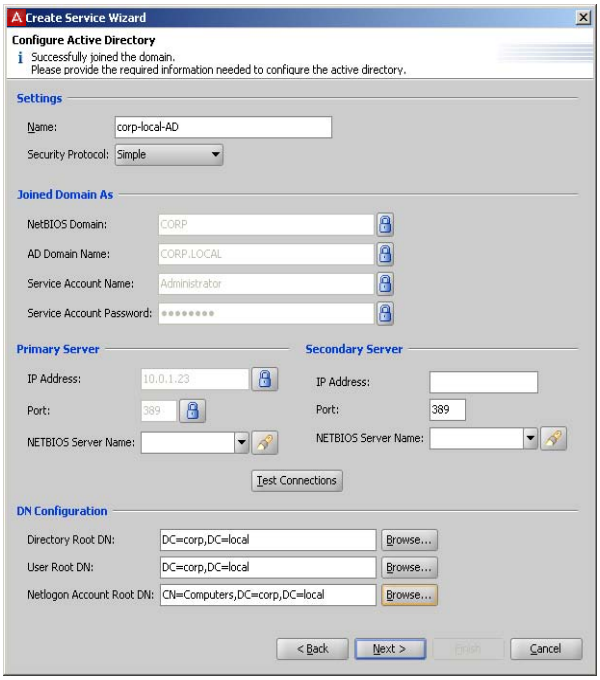


5. In the next window
   - Pick the **Security Protocol**: choose **Simple** for unencrypted communication with AD or choose **SSL** for encrypted communication.
   - A field or drop-down list appears to let you specify the **IP Address** of your AD server. Type or choose the address of your desired AD server.

- Check the **Port** setting. Ignition Server defaults to the port number used by most AD servers for the specified connection type (usually 636 for SSL or 389 for simple).



6. Ignition Server binds to the store, reads the schema, generates default settings, and displays the results in the Configure Active Directory window.



In this window:

- Assign the directory service a name in the **Name** field.
- In the **Primary Server** section, specify the **NETBIOS Server Name**.
- Add a **Secondary Server** if desired. This is a backup AD server.

- In the **DN Configuration** section, check the **Directory Root DN** and **User Root DN** fields. Initially, these fields contain default values that the wizard chose based on reading your schema. You may type the DN directly or click the **Browse** button to browse your directory to find it. Note that the schema browser does not display auxiliary classes; those you must type directly.
- In the **Netlogon Account Root DN** field, specify the DN in AD where the Ignition Server should create its own machine account when joining the AD domain. See "Netlogon Account Root DN" on page 138.

Note: If needed, you may edit the **Joined Domain As** and **Primary/ Secondary Server** settings. To edit any field, click the Lock icon to unlock the field, and edit the field. For an explanation of each field, see "AD Connection Settings" on page 136.

7. Click the **Test Connections** button. Testing the connection may take a few minutes. If a configuration setting is incorrect, Ignition Server sends a warning. If you receive an error message, correct your settings and test again. If the error message persists, see "Problem: Errors Occur During Directory Service Set-Up" on page 457.
8. Click **Next**.
9. The next window summarizes the connection settings of the service. Click **Finish**. Your new service appears in the Directory Services list. A blue check mark in the **Connected** column indicates a successful connection.

Avaya recommends that you test your connection by adding it as the only directory service in a directory set and testing this directory set in the Advanced Troubleshooting window.

Next Steps Next, add your directory service to a directory set as explained in "Adding a Directories and Authentication Servers to a Directory Set" on page 167.

Creating an Active Directory Service: Manually Configuring

The Create Directory Service Wizard guides you through the steps needed to connect Ignition Server to Active Directory. The following instructions show how to use the *manual mode* of the wizard. For instructions on using the automatic mode, see "Creating an Active Directory Service: Automatically Configuring" on page 140.

Before you begin:

1. Make sure your DNS server addresses have been configured in Ignition Server as explained in "Editing Ignition Server's DNS Settings" on page 57.

2. Make the clock and firewall settings explained in “Preparing to Connect to Active Directory” on page 139.
3. Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

To connect to AD (manual mode):

1. In Dashboard’s **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.
2. Ignition Server launches the Create Directory Service Wizard. The **Choose Directory Service Type** window appears.
3. Select the **Active Directory** radio button. Click **Next**.
4. The **Active Directory Configuration Options** window appears. Select the **Manually Configure** radio button. Click **Next.**
5. The **Configure Active Directory** window appears. Enter the details for the Active Directory service.
   - For the **Security Protocol**: choose **Simple** for unencrypted communication with AD or choose **SSL** for encrypted communication.
   - All other fields are described in “AD Connection Settings” on page 136.
6. Click the **Test Connections** button. Testing the connection might take a few minutes. If a configuration setting is incorrect, Ignition Server sends a warning. If you receive an error message, correct your settings and test again. If the error message persists, see “Problem: Errors Occur During Directory Service Set-Up” on page 457.
7. Click **Next**.
8. A window appears, summarizing the settings you have made. Click **Finish**.

   Your new service appears in the Directory Services list. A blue check mark in the **Connected** column indicates a successful connection. See “Checking Directory Service Connections” on page 170 for an explanation of the icons.

Next Steps Next, add your directory service to a directory set as explained in “Adding a Directories and Authentication Servers to a Directory Set” on page 167.

Troubleshooting AD Connections

This section contains tips for:

- “Diagnosing Connection Problems” on page 145
- “Looking Up AD Settings: Finding Your Root DNs” on page 145
- “Looking Up AD Settings: Finding Domain and NetBIOS Names” on page 147
- “Finding the AD Server’s IP Address” on page 147
- “Additional AD Resources” on page 148

Diagnosing Connection Problems

To check your AD connection:

1. Use the **Test Connections** or **Recheck Service** button. See “Troubleshooting User Lookup and Authentication” on page 170.
2. Check your AD settings (listed in “AD Connection Settings” on page 136).
3. Check your directory service connection using the Advanced Troubleshooting window:
   - Place the directory service in a directory set. (See “Adding a Directories and Authentication Servers to a Directory Set” on page 167)
   - Use the **Process Request** tab and Test Join feature of the Directory Service debugger. See “Troubleshooting User Lookup and Authentication” on page 170.

Looking Up AD Settings: Finding Your Root DNs

User Root DN and **Directory Root DN**: Enter the names of containers in your AD data store using X.500 naming. **User Root DN** points to the AD container that stores your user records. **Directory Root DN** points to the root of your AD tree and is used to obtain schema and group information.

To find out the X.500 names of your containers, use Microsoft’s built-in tools as follows: Open the Active Directory Users and Computers snap-in and check the tree panel on the left. At the root of the tree is the DNS name of your AD server. This provides the “dc=company,dc=com” portion of the name in the example below. For User Root DN, you must find the appropriate container (“CN”) or organizational unit (“OU”) and use its name as the “cn=” or “ou=” portion of the name. Note that an OU name may contain spaces, but that no space is allowed to fall directly after a comma in the X.500 name.

Figure 32 Example: User Root DN is cn=users,dc=company,dc=com

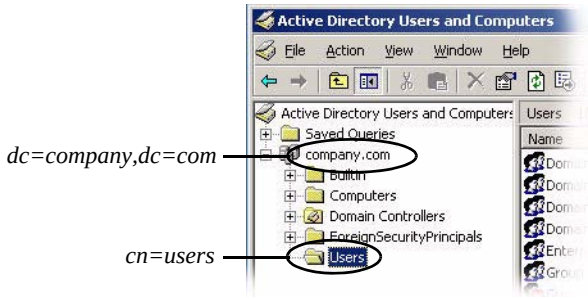
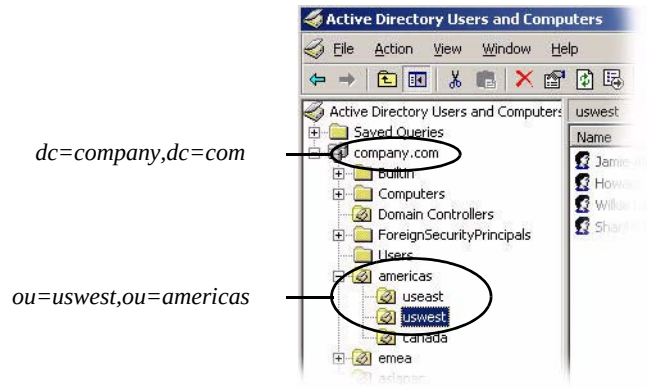


Figure 33 Example: User Root DN is ou=uswest,ou=americas,dc=company,dc=com



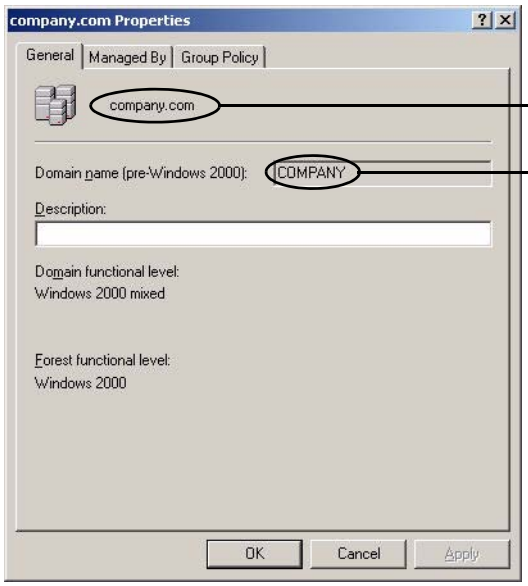
You form the full User Root DN name by adding the CN or OU portion of the name as a prefix to the root portion of the name, as shown in the two examples above. The following example text uses “cn=users,dc=company,dc=com” as our example DN.

Looking Up AD Settings: Finding Domain and NetBIOS Names

To find the **AD Domain Name** and **NetBIOS Name**, open the Active Directory Users and Computers snap-in and find your root domain in the tree panel on the left. In this example, the root domain is “company.com”. Right-click the root domain name and select **Properties** to open the Properties window.



In the General tab of Properties window, use the uppermost name as the “AD Domain Name” in Ignition, and use the Domain name (pre-Windows 2000) as the “NetBIOS Name” in Ignition.



Finding the AD Server’s IP Address

To find the IP address of your AD server, log into the machine that hosts your AD server and use the “ipconfig” tool from the command line, or open the Windows Control Panel and select **Network Connections: Local Area Connection**. In the Local Area Connection Status window, click **Properties**. In the Local Area Connection Properties window, click **TCP/IP** and then click **Properties**. Read the **IP address** from the TCP/IP Properties window.

Additional AD Resources

For more tips on connecting to AD, see “Troubleshooting” on page 451.

Connecting to an LDAP Service

This section shows you how to connect Ignition Server to an LDAP server such as SunONE LDAP. For a list of supported LDAP servers, see “Supported Directory Servers” on page 135. This section consists of:

- “LDAP Connection Settings” on page 148
- “Creating an LDAP Directory Service: Automatically Configuring” on page 149
- “Creating an LDAP Directory Service: Manually Configuring” on page 153
- “MSCHAPv2 Termination using an LDAP Password Attribute” on page 157

LDAP Connection Settings

The following table describes the parameters Ignition Server uses to connect to an LDAP service. You make these settings in the Create Directory Service Wizard or in the Directory Server Details window.

Table 4 Alphabetically sorted list of directory service connection settings for LDAP

Field Name	Description
	The Lock icon locks the adjoining field so that you cannot type text in it. Click the icon to unlock the field. Click the icon again to make the display read-only.
Associated Sets	The Ignition Server directory sets in which this Ignition Server directory service appears as a service.
Directory Root DN	Root Distinguished name (DN) of the LDAP tree. This is used to fetch schema and group information from the directory. For example, `dc=starironinc,dc=com`.
Browse Buttons	The Directory Root DN, User Root DN, and Username Attribute buttons allow you to browse your schema to set those values. (Note: Before you browse, you must provide connection information for information for the Primary Server: Service Account Name, the Service Account Password, IP Address, and Port number.)
Directory Root DN	DN where the LDAP schema containing your users and groups are found. For example, *dc=company,dc=com*. When you connect the directory service, the Ignition Server Create Service wizard attempts to choose a Directory Root DN for you.
MSCHAPv2 authentication	LDAP only: checkbox indicating whether this LDAP store supports MSCHAPv2 authentication. See also LDAP Password Attribute, below.
Name	Each directory service your create in Ignition Server is labeled with a name to help you refer to it later. You can use this name in your authorization policy. For example, you can write policy that provides special provisioning for users who authenticate against a particular directory. See “Inbound Attributes” on page 234.

Table 4 Alphabetically sorted list of directory service connection settings for LDAP

Field Name	Description
LDAP Password Attribute	For use in terminating MSCHAPv2 authentication against an LDAP directory, the Password Attribute is the user attribute in your LDAP directory that holds the NT-hashed password of the user. See "Setting up MSCHAPv2 Authentication on LDAP" on page 156.
Primary Server	IP address for the primary LDAP server. Port for the primary LDAP server. Generally, for **SSL** enter 636. If SSL is not used, enter 389. You cannot use the global catalog port (3268). Please use the LDAP ports (389 and 636) only!
Secondary Server	IP address for the secondary LDAP server. Port for the secondary LDAP server. Generally, for **SSL** enter 636. If SSL is not used, enter 389.
Security Protocol	Security protocol used for the Ignition Server's connection to the directory server. If **Use SSL** is turned on, Ignition Server uses SSL to encrypt traffic to the directory service. **Warning:** If you choose to connect to LDAP using a non-SSL connection, your service account credentials travel over the network in unencrypted form. Avaya strongly recommends using an SSL connection to connect to your directory server. Note the following: ▪ When **Use SSL** is selected, the **Port** Entry is typically **636**. ▪ When **Use SSL** is not selected, the **Port** Entry is typically **389**.
Service Account DN	**LDAP only**: DN of the LDAP administrator account. Ignition Server connects as this administrator. For example, `cn=Directory Manager`
Service Account Password	Password of the LDAP administrator.
Service Type	Vendor and type of directory service.
Strip Realm	**LDAP only**: This checkbox indicates whether Ignition Server should strip the realm name from the username before submitting it for authentication. If this box is checked, then, for example, the user name `jsmith@company.com` would be submitted to LDAP as `jsmith`.
User Root DN	DN of the LDAP container Ignition Server from where Ignition Server loads user records. For example, `cn=users,dc=starironinc,dc=com`. When you connect the directory service, the Ignition Server Create Service wizard attempts to choose a User Root DN for you.
Username Attribute	An LDAP attribute that stores the user name. Typically, this is *uid*.

Creating an LDAP Directory Service: Automatically Configuring

The Create Directory Service Wizard guides you through the steps needed to connect Ignition Server to your directory service. These instructions apply to Sun Directory Server, Novell eDirectory, Oracle Internet Directory (OID), and generic LDAP stores.

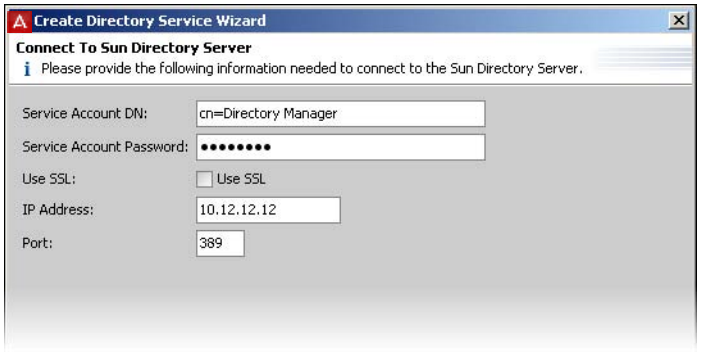
The following instructions show how to use the *automatic mode* of the wizard, which retrieves the default DNs for you. For instructions on using the manual mode, see “Creating an LDAP Directory Service: Manually Configuring” on page 153.

Before you begin:

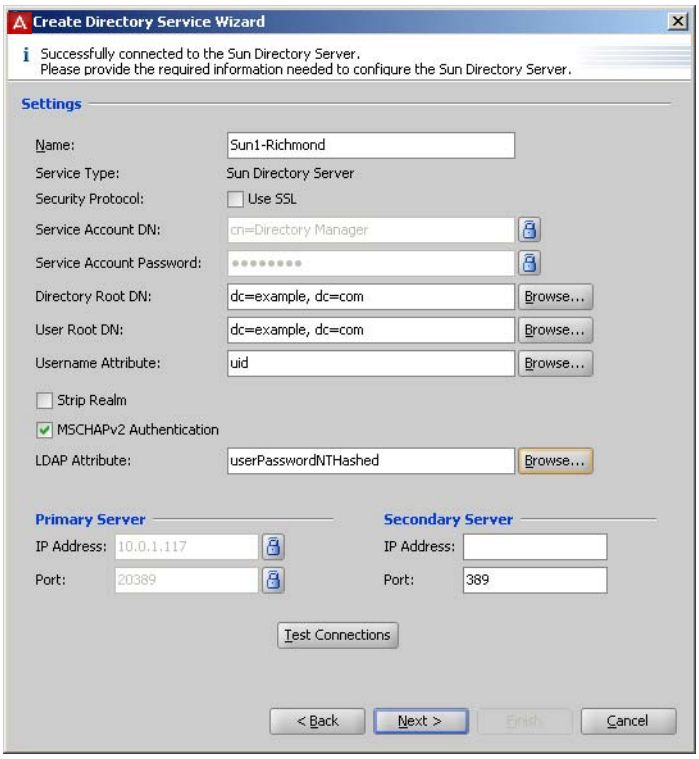
1. Make sure your DNS server addresses have been configured in Ignition Server as explained in “Editing Ignition Server’s DNS Settings” on page 57.
2. Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

To connect to LDAP (automatic mode):

1. In Dashboard’s **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.
2. Ignition Server launches the Create Directory Service Wizard. The **Choose Directory Service Type** window appears. Select the radio button for the desired type of LDAP server and click **Next**.
3. Click **Automatically Configure** and click **Next.**
4. In the **Connect To...** window, enter your LDAP administrator credentials in the Service **Account DN** and **Password** fields. Type the IP address and port of the LDAP directory. If you want to establish an encrypted connection to LDAP, click **Use SSL**. Click **Next**.



5. Ignition Server binds to the store, reads the schema, generates default settings, and displays the results on a new page of the wizard.



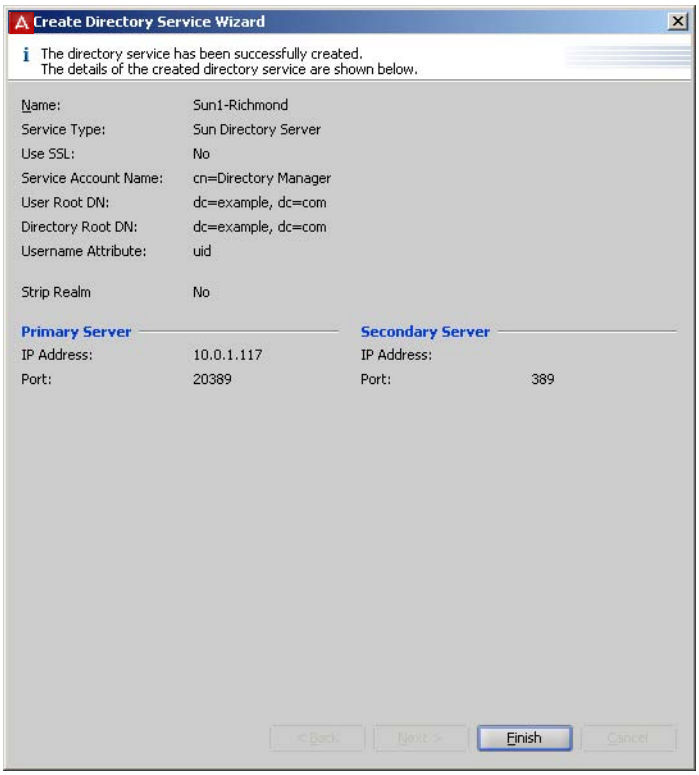
In this window, edit fields as needed. For an explanation of each field, see "AD Connection Settings" on page 136.

- Assign the directory service a name in the **Name** field.
- If needed, edit the **Security Protocol, Service Account DN** and **Password**. To edit a field, click the Lock icon to unlock the field, and edit the field.
- If needed, edit the **Directory Root DN**, **User Root DN** and **Username Attribute** settings. Initially, these fields contain default values that the wizard chose based on reading your schema. Enter the correct values for your site by editing these fields directly or by clicking the Browse button and selecting the proper root.

  Note: The **Directory Root DN** and **User Root DN** are often blank in Novell eDirectory configurations.
- If you want to strip the realm name from the username before submitting it for authentication, tick the **Strip Realm** checkbox. For example, with the box checked, Ignition Server submits *jsmith* instead of *jsmith@company.com*.

- If you want to support MSCHAPv2 authentication (typically needed only if you want to support clients that have Microsoft Windows supplicants), tick the **MSCHAPv2 Authentication** checkbox and click the **Browse** button to select the name of the LDAP user attribute that holds the NT-hashed password or type the attribute name directly in the **LDAP Password Attribute** field. See "Setting up MSCHAPv2 Authentication on LDAP" on page 156 for more details.
- If needed, edit the **Primary/Secondary Server** settings.

6. Click the **Test Connections** button. Testing the connection may take a few minutes. If a configuration setting is incorrect, Ignition Server warns you. If you receive an error message, correct your settings and test again. If the error message persists, see "Problem: Errors Occur During Directory Service Set-Up" on page 457.
7. Click **Next**.
8. The next window summarizes the connection settings of the service. Click **Finish**.



Your new service appears in the Directory Services list. A blue check mark in the **Connected** column indicates a successful connection.

Next Steps Next, add your directory service to a directory set as explained in "Adding a Directories and Authentication Servers to a Directory Set" on page 167.

Creating an LDAP Directory Service: Manually Configuring

The Directory Service Wizard guides you through the steps needed to connect Ignition Server to your directory service. These instructions apply to Sun Directory Server, Novell eDirectory, Oracle Internet Directory (OID), and generic LDAP stores.

The following instructions show how to use the *manual mode* of the wizard. For instructions on using the automatic mode, see "Creating an LDAP Directory Service: Automatically Configuring" on page 149.



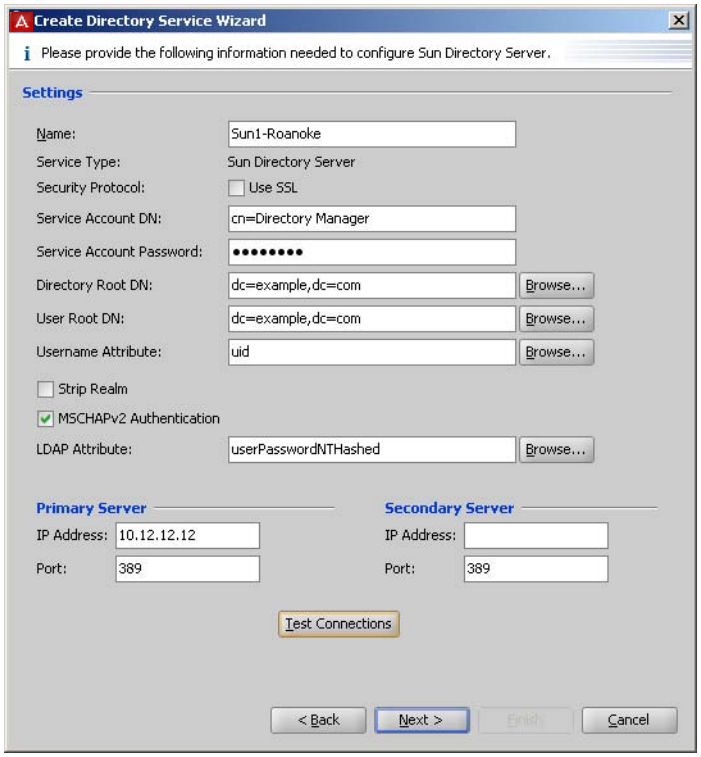
Before you begin:

1. Make sure your DNS server addresses have been configured in Ignition Server as explained in "Editing Ignition Server's DNS Settings" on page 57.
2. Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

To connect to LDAP (manual mode):

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.
2. Ignition Server launches the Create Directory Service Wizard. The **Choose Directory Service Type** window appears. Select the radio button for the desired type of LDAP server. Click **Next**.
3. Click **Manually Configure**. Click **Next.**

4. The Configure window for the selected LDAP compliant server appears.



In this window, edit fields as needed. For an explanation of each field, see "AD Connection Settings" on page 136.

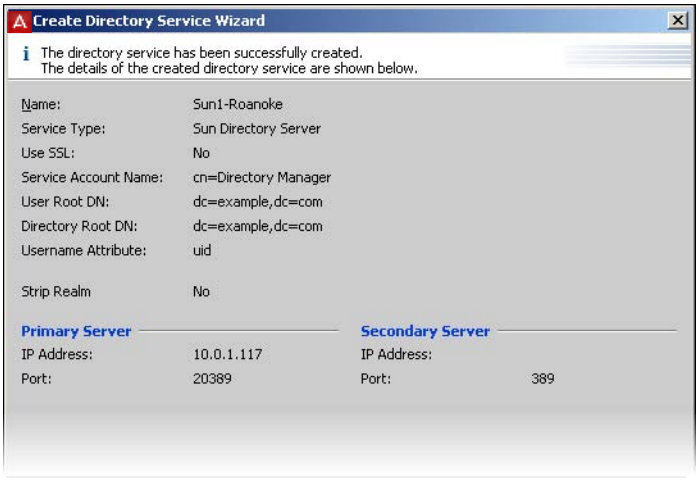
- Assign the directory service a name in the **Name** field.
- Set the **Security Protocol** and type the LDAP administrator credentials in the **Service Account DN** and **Password** fields.
- Enter the **Directory Root DN**, **User Root DN** and **Username Attribute** settings by editing these fields directly or by clicking the Browse button and selecting the proper root.

  Note: The **Directory Root DN** and **User Root DN** are often left blank in Novell eDirectory configurations.
- If you want to strip the realm name from the username before submitting it for authentication, tick the **Strip Realm** checkbox. For example, with the box checked, Ignition Server submits *jsmith* instead of *jsmith@company.com*.
- If you want to support MSCHAPv2 authentication (typically needed only if you want to support clients that have Microsoft Windows supplicants), tick the **MSCHAPv2 Authentication** checkbox and click the **Browse** button to select the name of the LDAP user attribute that holds the NT-hashed password or type the attribute name directly in

the **LDAP Password Attribute** field. See “Setting up MSCHAPv2 Authentication on LDAP” on page 156 for more details.

- Set the **Primary/Secondary Server** addresses and ports.
- Click the **Test Connections** button. Testing the connection may take a few minutes. If a configuration setting is incorrect, Ignition Server warns you. If you receive an error message, correct your settings and test again. If the error message persists, see “Problem: Errors Occur During Directory Service Set-Up” on page 457.
- Click **Next**.

5. The wizard establishes the LDAP directory connection with your settings. It displays the connection settings on a new page of the wizard.



6. Click **Finish**.

   Your new service appears in the Directory Services list. A blue check mark in the **Connected** column indicates a successful connection. See “Checking Directory Service Connections” on page 170 for an explanation of the icons.

Next Steps Next, add your directory service to a directory set as explained in “Adding a Directories and Authentication Servers to a Directory Set” on page 167.

Creating a Token Authentication Service

For information on connecting to hardware token authentication service, see “Overview of Token Authentication in Ignition” on page 179.

Creating a Kerberos Authentication Service

For information on connecting to a Kerberos server, see “Setting up a Kerberos Authentication Service” on page 177.

Setting up MSCHAPv2 Authentication on LDAP

Often, an organization must allow Windows users to authenticate against user accounts stored in a non-AD, LDAP directory store, but the organization does not want to deploy new certificates or new supplicants to users. To support such cases, Ignition Server offers the ability to authenticate via the MSCHAPv2 to a non-AD LDAP store. This feature is called “MSCHAPv2 termination against LDAP.”

The sections below explain how to set up MSCHAPv2 termination against LDAP by configuring your LDAP directory service to support MSCHAPv2. The list of supported authentication types appears in “Supported Authentication Types” on page 215.

Ignition Server supports a number of approaches to authenticate MSCHAPv2 clients against an LDAP directory:

- “MSCHAPv2 Termination using an LDAP Password Attribute”, below.
- “MSCHAPv2 Termination using a Novell Universal Password” on page 159.
- “MSCHAPv2 Termination using an OID Authentication Descriptor” on page 160

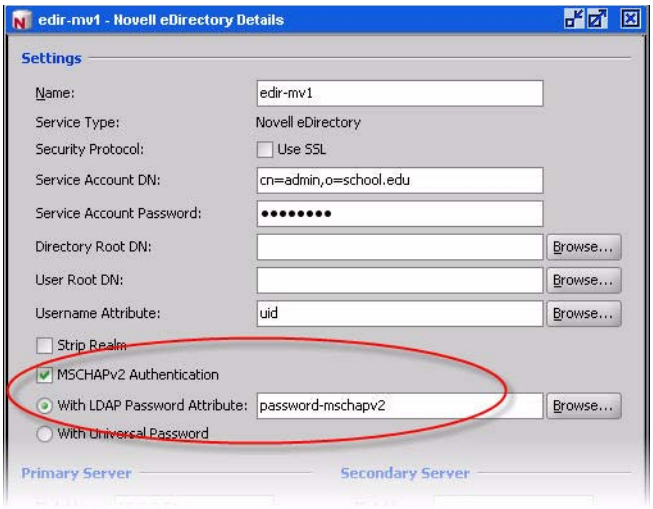
Note! Before you deploy MSCHAPv2 termination against LDAP, you should consider deploying EAP-TTLS authentication instead. To do this, you must deploy on each user’s Windows computer an EAP-TTLS-compliant supplicant such as the OpenSEA Xsupplicant. Using EAP-TTLS authentication has a number of advantages over PEAP-MSCHAPv2 (and it has few disadvantages). The advantages are:

- Authentication is done seamlessly against LDAP. No new provisioning tools or plug-ins are needed to support user accounts stored in non-AD directories.
- Users are not required to change their passwords after implementation of the password storage scheme.
- In the future, you can take advantage of other types credential stores such as Kerberos and databases.
- EAP-TTLS is more secure than the PEAP tunnel because it does not expose any user information in the outer tunnel.

MSCHAPv2 Termination using an LDAP Password Attribute

Instructions: To set up MSCHAPv2 authentication against an LDAP directory service, do the following:

1. Identify an unused attribute in your LDAP user schema definition. This attribute is used to store the hash of the user's password that is necessary to perform MSCHAPv2 authentication. Keep in mind that the attribute should have a binary format and should be single-valued.
2. Create an NT hash of each user's password and save it in the LDAP store. For instructions, see "Creating an NT-Hashed Password to Support MSCHAPv2 Termination Against LDAP" on page 158.
3. In Ignition, create or edit your LDAP directory service and make these settings in the Directory Server Details window:
   - Tick the **MSCHAPv2 Authentication** checkbox.
   - If necessary, tick the **With LDAP Password Attribute** checkbox. (This is required for Novell eDirectory and Oracle OID only.)
   - In the **LDAP Password Attribute** field, type the name of the LDAP user attribute that holds the NT-hashed password, or click the **Browse** button and select the attribute,



4. Save the directory service.
5. Open the Access Policy panel of Dashboard (in Dashboard's Configuration hierarchy, expand Access Policies, expand RADIUS, and click your policy name) to edit the access policy of the access points or switches to which your users can:
   - Set the identity routing policy to use the directory service you saved above.

- Edit the authentication policy. Make sure it is set up to allow one or more of the following authentication types: NONE/MSCHAPv2, NONE/EAP-MSCHAPv2, or PEAP/EAP-MSCHAPv2.

When the above configuration is complete, your Ignition Server installation supports authenticating MSCHAPv2 clients against the LDAP store.



Note! A given user might have multiple protocols available to him or her for logging in. For example, you can set up a user authorization policy that allows both PEAP / EAP-MSCHAPv2 and NONE / PAP authentication types, with your site's LDAP store as the directory service that supports both. This allows a user, for example, to log in from his Windows XP laptop using his Windows password and log in from his Linux workstation using his LDAP password. In both cases, he is authenticated against the LDAP store. See "One Policy Allows Many Authentication Protocols" on page 215.

Creating an NT-Hashed Password to Support MSCHAPv2 Termination Against LDAP

Ignition Server needs an MD4 hash of the user's password or the password in plaintext in order to terminate the MSCHAPv2 authentication protocol. Except for Novell's universal password feature, few directories store the plaintext password in the directory under any circumstances (and in any case few administrators would be comfortable with doing this).

To perform MSCHAPv2 termination against LDAP, the Ignition Server extracts the NT hash (an MD4 hash) of the password from the directory by querying a specified user attribute. If the attribute is defined for the user, it is expected to contain a binary format of the hash, *not* the ASCII format. The steps below show you how to deploy mechanisms that create and maintain the NT hash.

1. **Write your hash-creating plug-in.** If your site uses a web-based provisioning tool to add new users and change passwords, you can usually add a custom plug-in that updates the hash. Set up this plug-in to be triggered each time a password is saved, so that the plug-in updates the NT hash of the password every time the password is changed. The plug-in must do the following:
   a. Convert the cleartext password to little-endian UCS2 format.
   b. Hash the UCS2-formatted password with the MD4 algorithm to obtain a 16-byte binary hash. The script to construct the binary password hash in base-64 from the ASCII plaintext password "secret" is as follows:

```
echo -n "secret" | iconv -f iso_8859-1 -t ucs-2le | openssl dgst -md4 -binary |  openssl enc -
base64
```

   c. Save the hash to the user's entry in the directory. The base-64 hash can be inserted into the directory with command-line utilities or using

LDAP client code. All of these tools (*iconv*, *openssl*) are available on most UNIX/Linux distributions.

2. **Have each user change his or her password once.** After you have deployed your hash-creating tool, every user in the directory must change his or her password at least once. When the password is changed the hash-creating tool creates the hash in the correct format. This is required because Ignition Server cannot extract the hash in its existing format in the directory (MD5, SHA1, and so on).
3. **Set up your account provisioning environment to keep the hash in sync with the user's password.** The hash must be kept in sync with any other versions of the password stored in the directory natively. Your site might allow users to change LDAP passwords through a variety of clients, such as mail clients, and web applications. If such password editing tools are used in your environment, you must either modify them to update the password's NTHASH as well, or you must disable them for users who authenticate via MSCHAPv2 against LDAP.

Next steps: Return to Step 3 in "MSCHAPv2 Termination using an LDAP Password Attribute" on page 157.

MSCHAPv2 Termination using a Novell Universal Password

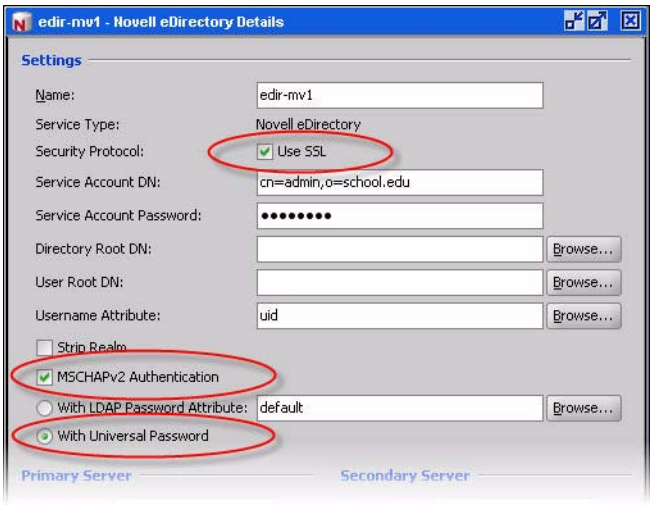
Novell eDirectory provides the eDirectory Universal Password feature. This feature allows you to store a single password per user in the directory and support multiple authentication methods using this password. Ignition Server can be configured to terminate MSCHAPv2 authentication against the Universal Password in eDirectory.

The Universal Password feature is available on eDirectory, version 8.7 and later, and is enabled by default on version 8.8 and later.

Instructions: To set up MSCHAPv2 authentication against Novell eDirectory using the eDirectory Universal Password feature, do the following:

1. Make sure the Universal Password feature is enabled on your eDirectory server, and make sure a Universal Password is stored for each user. Consult your eDirectory documentation for details.
2. Make sure your eDirectory password policy allows the administrator (the one whose credentials you used to connect Ignition Server in Step 4 in "To connect to LDAP (automatic mode):" on page 150) to extract users' passwords.
3. In Ignition, create or edit your LDAP directory service and:
   - Tick the **Use SSL** checkbox. The eDirectory server requires an SSL connection in order to use Universal Passwords.
   - Tick the **MSCHAPv2 Authentication** checkbox.

- Tick the **With Universal Password** checkbox.



4. Save the directory service.
5. Open the **Access Policy panel of Dashboard** (in Dashboard's Configuration hierarchy, expand Access Policies, expand RADIUS, and click your policy name) to edit the access policy of the access points or switches to which your users can connect:
   - Set the identity routing policy to use the directory service you saved above.
   - Edit the authentication policy. Make sure it is set up to allow one or more of the following authentication types: NONE/MSCHAPv2, NONE/EAP-MSCHAPv2, or PEAP/EAP-MSCHAPv2.

When the above configuration is complete, your Ignition Server installation supports authenticating MSCHAPv2 clients against the eDirectory service.

MSCHAPv2 Termination using an OID Authentication Descriptor

Oracle's Internet Directory (OID) provides a value called the Authentication Descriptor. Ignition Server can be configured to terminate MSCHAPv2 authentication against the Authentication Descriptor in OID. To set this up:

1. In Ignition, create or edit your OID directory service and, in the Directory Server Details window:
   - Tick the **MSCHAPv2 Authentication** checkbox.
   - Tick the **With Authentication Descriptor** checkbox.
2. Save the directory service.

3. Open the **Access Policy panel of Dashboard** (in Dashboard's Configuration hierarchy, expand Access Policies, expand RADIUS, and click your policy name) to edit the access policy of the access points or switches to which your users can connect:
   - Set the identity routing policy to use the directory service you saved above.
   - Edit the authentication policy. Make sure it is set up to use an MSCHAPv2 authentication type. This enables users to authenticate using their Windows domain account name and password.

After the above configuration is complete, your Ignition Server installation supports authenticating MSCHAPv2 clients against the OID service.

Note: If you wish to test that the OID authentication descriptor attribute is available, you should create a virtual attribute and map it to the OID user attribute, *authpassword;orclcommonpwd*. This attribute is a multi-valued attribute, and the authentication descriptor is the value prefixed with the string, *{X- ORCLNTV}*. Use the Advanced Troubleshooting window to perform a test retrieval of the value. See "Adding a New User Virtual Attribute" on page 205 and "Testing a User Lookup" on page 174 for instructions.

Creating a Radius Proxy Authentication Service

For more information on setting up a Radius proxy server, see "Setting up a RADIUS Proxy Server" on page 189

Managing Directory Services

The following sections explain how to manage your LDAP and AD connections in Ignition:

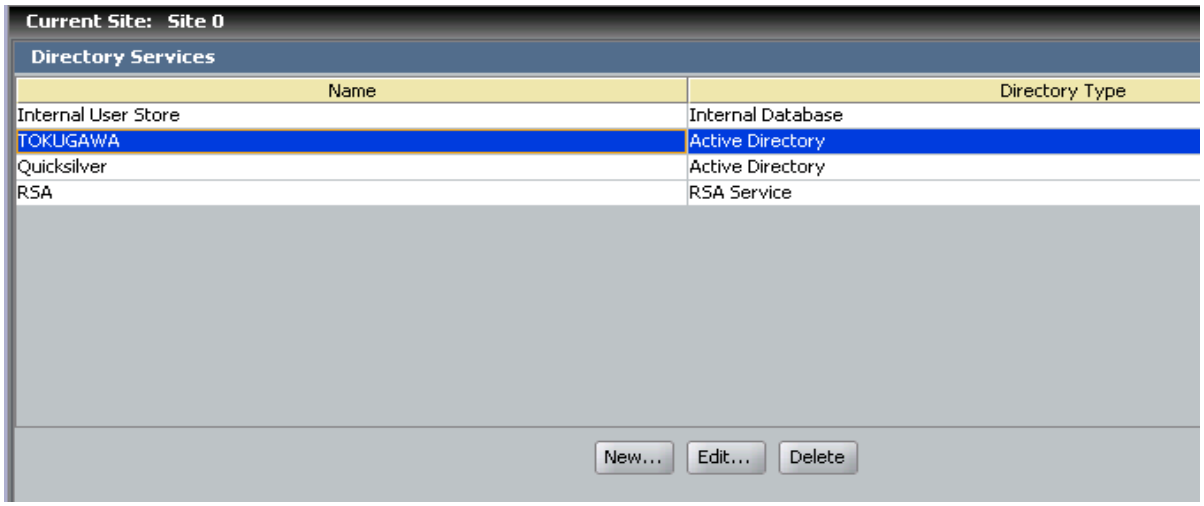
- "Editing Directory Service Configurations" on page 161
- "Renaming a Directory Service" on page 162
- "Deleting a Directory Service" on page 163

Editing Directory Service Configurations

To modify the parameters used to connect to an LDAP or Active Directory store, you must edit the directory service configurations for that store. Use the following steps:

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. The Directory Services window displays the current set of configured directory services.

Figure 34 Directory Services window



2. Select the entry for the directory service you want to edit. Click **Edit**.
3. Ignition Server populates the **Directory Service Details** Window with the details of the selected directory service.
4. Edit the details of the selected directory service as required.
5. Click **OK** to apply your changes.

Renaming a Directory Service

When you rename a directory service, Ignition Server uses the updated name for the directory service in

- all mappings of the directory service's attributes to the existing virtual attributes
- all mappings of the directory service's groups to the existing virtual groups
- all the directory set(s) to which the directory service belongs.

Important! Renaming a directory service breaks the authorization rules that depend on that directory service. See "Problem: Authorization policy stops working unexpectedly" on page 453 for troubleshooting instructions.

To Rename a Directory Service:

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**.
2. In the **Directory Services panel**, select the entry you want to rename (Figure 34).
3. Click **Edit**. Ignition Server displays the details for the selected directory service (Figure 35).

4. Enter a different name for the directory service.
5. Click **OK**.

Figure 35 Renaming an Existing Directory Service



In the above example, the title of the window indicates that the directory service FRB-DAL-AD1 is being edited. The name field in the window indicates that the name of the directory service is being changed to FRB-DAL-AD2. If the user clicks **OK** in this window, the renaming change breaks the authorization rules that currently use FRB-DAL-AD1.

Deleting a Directory Service

Important! Deleting a directory service breaks the authorization rules that depend on that directory service. See Problem: Authorization policy stops working unexpectedly for troubleshooting instructions.

When you delete a directory service, Ignition Server deletes all mappings of the directory service's attributes to the existing virtual attributes. In addition, Ignition Server deletes all mappings of the directory service's groups to the existing virtual groups.

Ignition Server recommends that, before you delete a directory service, you disassociate that directory service from all the directory sets to which it belongs.

To Delete a Directory Service:

1. Use the **Directory Sets Panel** to make sure that the directory service to be deleted is not a member of any directory set. (In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Sets**.)
2. In Dashboard's **Configuration** hierarchy tree, under **Directories**, click **Directory Services**. The Directory Services window displays the current set of configured directory services.
3. Select the directory service to be deleted.
4. Right-click on the directory service to be deleted and select on **Delete**. Alternatively, click the **Delete** button.

Directory Sets

A directory set is an ordered list of user lookup services (AD, LDAP, and so on) and/or authentication services (SecurID, Kerberos, Radius proxy, and so on) to be used when Ignition Server processes an authentication request. The directory set determines which service is used to authenticate the user (the *authentication service*), which service is used to retrieve the user's account, and which service is used to retrieve authorization-determining data such as user attributes and group affiliations (the *user lookup service*). Before you can create a directory set, you must have created your authentication and lookup services in the Directory Services panel. See "Directory Services" on page 134 for details.

This section explains how to create and manage your directory sets. Before we look at creating a directory set, let's take a look at how directory sets fit into Ignition's approach to finding the user account at login time.

Since your environment may contain many thousands of user accounts and many authentication and lookup services, Ignition Server offers two mechanisms that let you establish the search logic for finding the user account:

- the **identity routing policy** lets you establish the search logic for *finding the directory set* that matches the user, based on the user's account domain and based on which switch the user is connecting from; and
- the **directory set** lets you establish the search logic for *finding the authentication service* and *finding the lookup service* based on a fallthrough order of services.

The complete search order (comprising the identity routing policy, the directory set, the authentication service, and the user lookup service) is illustrated in "Understanding Identity Routing Policy" on page 220.

After you create a directory set, you can test it as explained in "Checking an Authentication Request" on page 173.

Contents of This Section

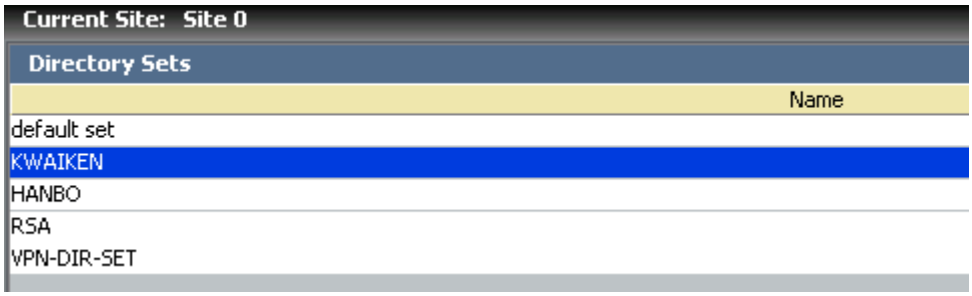
The sections below explain how to manage your directory sets:

- "Directory Sets Panel" on page 165
- "Adding a Directory Set" on page 166
- "Renaming a Directory Set" on page 166
- "Deleting a Directory Set" on page 167
- "Adding a Directories and Authentication Servers to a Directory Set" on page 167
- "Setting Fallthrough Rules to Handle Lookup and Authentication Failures" on page 169

Directory Sets Panel

To view the existing directory sets and the directory services (user lookup services) and authentication services they contain, go to Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Sets**. The Directory Sets panel appears.

Figure 36 Directory Sets window



The Directory Sets panel contains:

- **Directory Sets panel:** The **Directory Sets** panel lists the names of the available directory sets.
  - **List of directory sets**: The main part of this panel shows a list of directory sets. Click on a set name to view that set in the Directory Set Details panel.
  - **New**: Click the New button to open the Add Directory Set Window.
  - **Delete:** Click the Delete button to delete the selected directory set.
- **Directory Set window:** Click on a directory set in the Directory Sets panel, click **Edit** to show the **Directory Set** window. This window displays the contents of the selected directory set.
  - **Name**: The name of selected directory set

- **Directory Set Entries**: A list of the user lookup services (directory services) and authentication services (each may be either a directory service or an authentication service) in this set, each with its fallthrough settings.
- **Add, Remove:** The **Add**, and **Remove** buttons enable you to add a user lookup/authentication service pairing to this set or remove a pairing from the set.
- The **OK** command button allows you to save the changes.

Default Directory Set

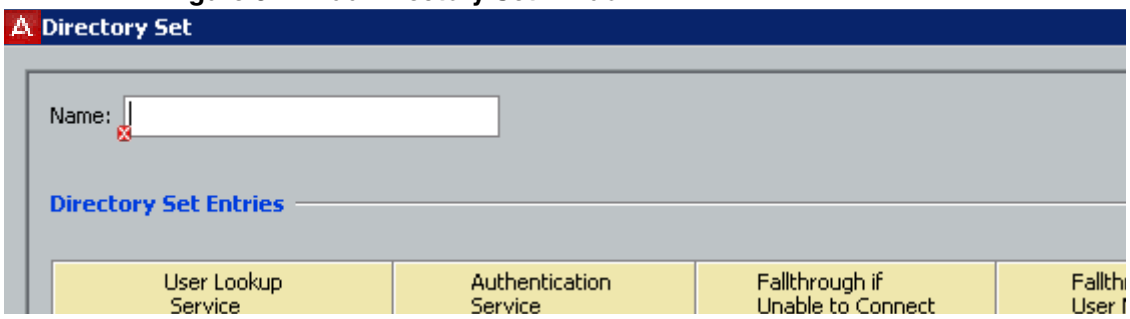
The standard installation of Ignition Server includes a directory set called "**default set**" that includes the internal data store as its only lookup/authentication service.

Adding a Directory Set

In order to add a directory set,

1. Click the **New** button in the **Directory Sets** panel.
2. The **Add Directory Set** window appears.

Figure 37 Add Directory Set window



3. Click **OK**.

Ignition Server adds the name you entered in the Add Directory window to the **Directory Sets** list.

Next Step Add directories to the set, as shown in "Adding a Directories and Authentication Servers to a Directory Set" on page 167.

Renaming a Directory Set

Avaya advises strongly against renaming your directory sets. If you must rename a directory set, then take care to edit each identity routing policy as explained below. Follow these steps to rename your directory set:

1. Select the desired directory set from the list in the **Directory Sets** panel. The **Directory Set Details** panel displays the directory set.

2. In the **Name** field at the top of the **Directory Set Details** panel, edit the name. (After you have edited the name, the **OK** button becomes usable. To abandon your edit, click **Cancel**.)
3. Click **OK** to save the new name.
4. In each identity routing policy that refers to the renamed directory set, you must change the name of the directory set to the new name. See "Creating an Identity Routing Policy" on page 223.

Deleting a Directory Set

In order to delete a directory set,

1. Select the directory set from the list in the **Directory Sets** panel.
2. You can delete the selected directory set in one of the following ways:
   - Right-Click on the directory set and click **Delete**.
   - Click the **Delete** button at the bottom of the Directory Sets panel.

Ignition Server deletes the selected directory name from the list in the **Directory Sets** panel.

Adding a Directories and Authentication Servers to a Directory Set

A directory set is a list of directories (ADs, LDAPs, and so on) and/or authentication servers (SecurID, Kerberos, Radius proxy, and so on) that are used to authenticate and look up users. In Ignition Server terminology, a directory used for retrieving user accounts is called a *user lookup service*, and a directory or authentication service used to verify users' credentials is called an *authentication service*.

Note: You cannot add a non-proxy service to a Directory Set that contains a Proxy Service.

To set up your directory set:

1. Select the desired directory set from the list displayed in the Directory Sets panel.
2. Click **Edit** in the Directory Sets panel.
3. Click the **Add** button in the **Directory Set** panel.
4. In the **Directory Set Entry** window, do the following:

   3.1 User Lookup Service: In this drop-down list, select the directory service that holds your user records, including any user attributes and user group affiliations you want to retrieve. If you want to authenticate users against an authentication server *only*, you may elect to use no

lookup service. To do this, select the “None” option in this list. **Note**: If you are adding a Proxy service, select the “None” option in this list.

3.2 Authentication Server: Set this to the name of the directory service or authentication service that verifies user credentials. In most cases, this is set to the same name as your User Lookup Service, but many sites that authenticate using RSA SecurID or Kerberos set this field to the name of the RSA or Kerberos server and set the **User Lookup** field to the name of an LDAP directory that holds user accounts. See “Authentication Services” on page 177 for further instructions.
Note: If you are adding a Proxy service, select the RADIUS Proxy service that you created from the **Authentication Server** drop-down list. For more information on creating a RADIUS Proxy service, see “Setting up a RADIUS Proxy Server” on page 189.

3.3 Click **OK**. Ignition Server creates a row in the directory set.

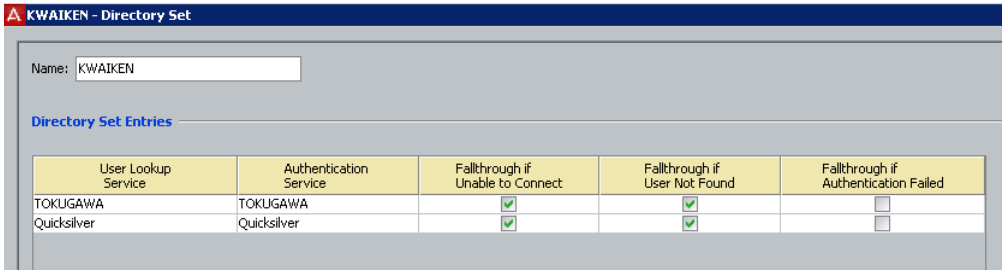
5. **Provide Fallthrough Conditions:** By default, fallthrough is enabled for the *Unable to Connect* case and the *User Not Found* case, and fallthrough is not enabled for the *Authentication Failed* case. For more information on how Ignition Server handles failure of the fallthrough conditions, see “Setting Fallthrough Rules to Handle Lookup and Authentication Failures” on page 169. Select the appropriate checkbox(es) to enforce a fallthrough for the following:

   - Unable to connect to the associated directory server
   - User/device is not found
   - User authentication fails

   Check the details for this new directory set. Click **OK**.

6. Ignition Server Dashboard displays the new directory set with its settings.

Figure 38 Example of a directory set with details



7. Repeat these steps to add more directory and authentication services for the selected directory set.

Next Step If you are done adding directories to your set, then you are ready to use your set in your identity routing policy. See “Creating an Identity Routing Policy” on page 223.

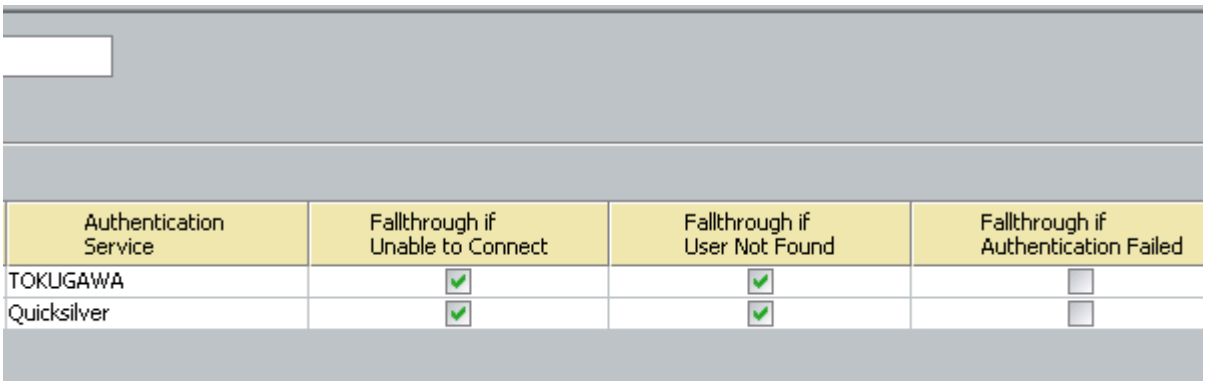
Note that you can also test your directory set as explained in “Checking an Authentication Request” on page 173.

Setting Fallthrough Rules to Handle Lookup and Authentication Failures

You can set Ignition Server to check additional directory and/or authentication services when a user lookup or user authentication fails. This feature is called *fallthrough* and is configured in the Directory Sets panel. Fallthrough can be set to occur if:

- **Unable to Connect**: Ignition’s attempt to connect to the directory or authentication service failed.
- **User Not Found**: User lookup found no user by the requested name.
- **Authentication Failed**: User found, but authentication failed.

Figure 39 Fallthrough Selections in the Directory Services Panel



Setting Fallthrough Rules

1. In the **Directory Set Details** panel, find the row for the directory service whose failure handling you want to set. The row provides three checkboxes, one to set the fallthrough handling for each of the three cases. The fallthrough rule you set here applies within the context of this directory set only.
2. In the row, set the fallthrough checkboxes as follows:
   - **Select the checkbox** if you want Ignition Server to move on to the next directory service in the event of a failure of the type specified in the column.
   - **Deselect the checkbox** if you want Ignition Server to reject the authentication request in the event of a failure of the type specified in the column.
3. Click **OK** to save your settings or **Cancel** to abandon your changes.

Fallthrough Behavior in Ignition

The following table shows how Ignition Server handles fallthrough rules:

Fallthrough Handling for Directory Services

Condition	If Directory Service is...	Do this...	If there is no next Directory Service...
Cannot connect	Checked	Try connecting to next directory service	Reject the service request.
	Not checked	Reject service request.	
Connected, but cannot find user	Checked	Try connecting to next directory service	Reject the service request.
	Not checked	Reject service request.	
Connected and found user, but cannot authenticate user	Checked	Try connecting to next directory service	Reject the service request.
	Not checked	Reject service request.	

In the example in Figure 39, when a user request comes to the directory set, “MountainView-AD-Store-1,” (because the policies of its access policy say so), the first directory service to be searched for that user is in the internal store. If the user is not found or if authentication against the user internal store fails, the checkmarks indicate Ignition Server falls through to **ad1**.

Troubleshooting User Lookup and Authentication

The sections below explain how to manage your directory sets:

- “Checking Directory Service Connections” on page 170
- “Checking the Group Cache” on page 172
- “Testing a Directory Service Connection” on page 172
- “Advanced Troubleshooting for Directory Services and Sets” on page 173

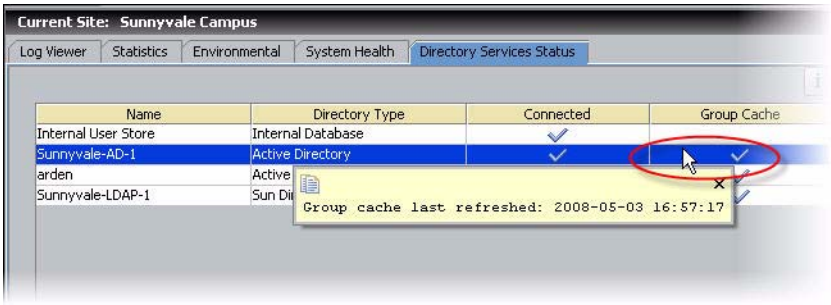
Note that the sections below *do not* cover problems specific to particular types of directory servers. For additional troubleshooting tips related to *Active Directory environments*, see “Problem: Authentication fails on Active Directory” on page 455.

Checking Directory Service Connections

To check a directory service connection:

1. Click **Monitor** to show Dashboard's **Monitor** hierarchy tree, and click the IP address of your Ignition Server.
2. Click the **Directory Services Status** tab. The **Directory Services Status** tab lists your directories and shows connection and cache status for each.
3. Click on a cell in the **Group Cache** column to see when the cache was most recently updated.

Figure 40 Clicking a cache status cell to display the cache update time



In this panel, the Connected column indicates to state of the connection to the directory server. The states are:

A blue check mark indicates that the directory service is currently connected via either the primary or secondary server.

a red "X" indicates that the primary (and if configured secondary) directory services are unreachable.

A question mark indicates that Ignition Server is checking the connection.

The commands of the **Directory Services Status** tab are

- **Ping** checks that the selected directory service's machine is reachable on the network.
- The **Recheck Service** button and the **Test Connections** button let you check that Ignition Server can connect to the selected directory server. Ignition Server tests the connection to the primary server and, if configured, the secondary server. For each server, the connection test consists of:
  - an anonymous bind to the directory;
  - retrieval of the directory's root DSE;
  - a bind using the service account credentials;
  - a search for the user root;
  - parsing of the directory schema; and

- retrieval of the user groupings and OU's

  A results window displays the test outcome, displaying one success/ failure line for the primary server and one line for the secondary server, if configured.

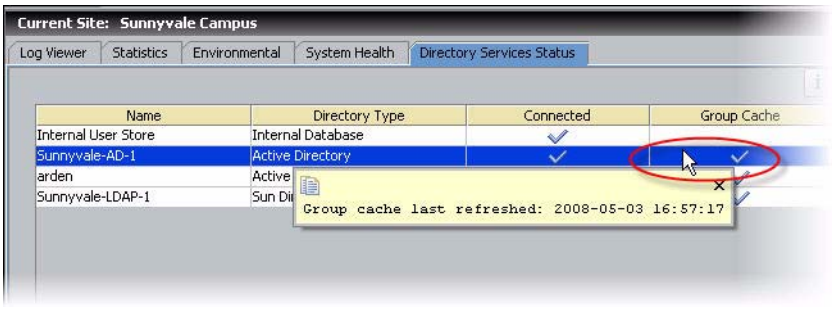
- **Refresh Cache:** The Ignition Server maintains an internal cache of the group hierarchies and attribute schemas of the directory services. The cache is automatically refreshed daily from the time you add the directory service. Use the **Refresh Cache** button to manually refresh the cache of the selected directory service.

Checking the Group Cache

Ignition Server maintains a local cache of the user group information (including user roles) available from each directory service. The Ignition Server refreshes each cache once a day, but you can force an update using the **Refresh Cache** button. To check and, if necessary, update a cache, follow these steps:

1. Click Monitor to show Dashboard's **Monitor** hierarchy tree, and click the IP address of your Ignition Server. Click the Directory Services Status tab.
2. In the tree, find the name of your service. The **Group Cache** column displays the status of the cache of user group and user role information. A blue check mark indicates the cache is current. Click a cell in this column to see the time of the most recent cache update.
3. If the cache is out of date, click the **Refresh Cache** button.

Figure 41 Clicking a cache status cell to display the cache update time



Testing a Directory Service Connection

Use the **Test Connections** button to test a directory service connection:

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Directory Services, and click on the name of your service.

2. Click the **Test Connections** button. Ignition Server tests the connection to the primary directory server and, if configured, the secondary directory server. For an explanation of the connection test, see "Checking Directory Service Connections" on page 170.

Important! If Ignition Server returns an error message when you test the connection, see "Problem: Errors Occur During Directory Service Set-Up" on page 457.

Advanced Troubleshooting for Directory Services and Sets

The Advanced Troubleshooting window allows you to test user and device lookups and authentications. The tests you can perform here are:

- Checking an Authentication Request. See page 173.
- Testing a User Lookup. See page 174.
- Testing a Device Lookup. See page 175.
- Testing a User Authentication. See page 176.

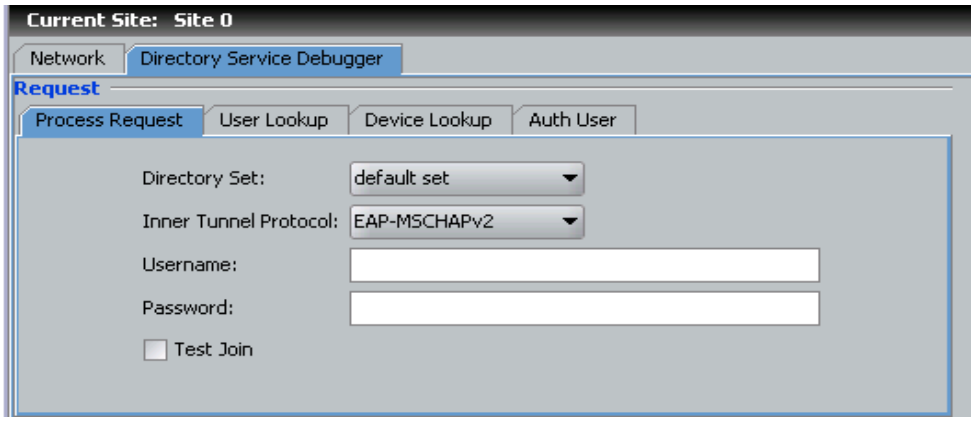
Important! If Ignition Server is running in HA mode, then all test queries originate from the primary node of Ignition.

Checking an Authentication Request

To check whether Ignition Server can successfully look up and authenticate a user against a particular *directory set*, do the following:

1. In Dashboard, click **Troubleshoot**.
2. Click the IP address of your Ignition Server.
3. Click the Directory Service Debugger tab.
4. In the debugger, click the **Process Request** tab.

Figure 42 The Process Request tab for testing user lookup and authentication



5. Enter your test authentication request:

- Specify the **Directory Set** that contains the service you want to test-authenticate against. If in doubt, check your directory set definition to make sure the desired service is queried.
- Specify the authentication protocol in the **Inner Tunnel Protocol** field.
- Type the user's credentials in the **Username** and **Password** fields.
- Tick the **Test Join** checkbox if you want to perform an LDAP join to the service.

6. Click the **Send Request** button to test the authentication and lookup.

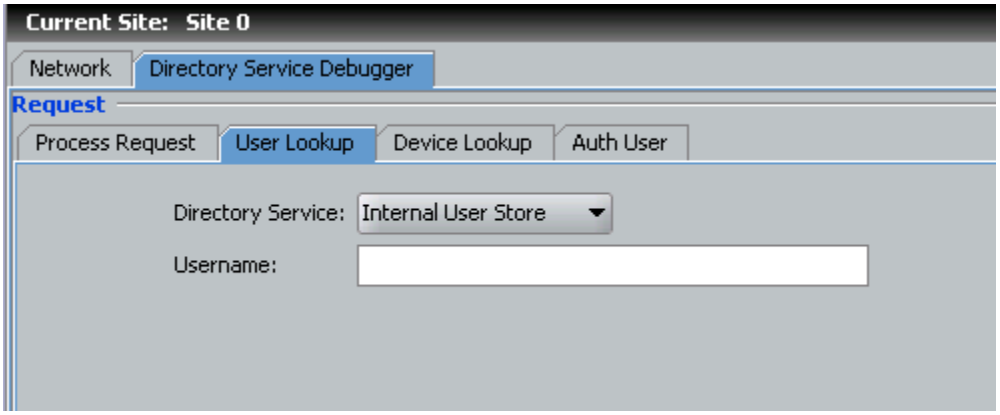
The **Results** field displays the lookup and authentication results, and the **Virtual Attributes** displays the attributes retrieved from the user's account record.

Testing a User Lookup

To find out whether a user account exists in a particular user store and to see what virtual attributes are returned from the user lookup, use the User Lookup tab of the Advanced Troubleshooting window.

1. In Dashboard, click **Troubleshoot**.
2. Click the IP address of your Ignition Server.
3. Click the Directory Service Debugger tab.
4. In the debugger, click the **User Lookup** tab.

Figure 43 The User Lookup tab for verifying an account exists



5. Specify the **Directory Service** that you want to search.
6. Type the user's login name in the **Username** field.
7. Click the **Send Request** button to test the authentication and lookup.

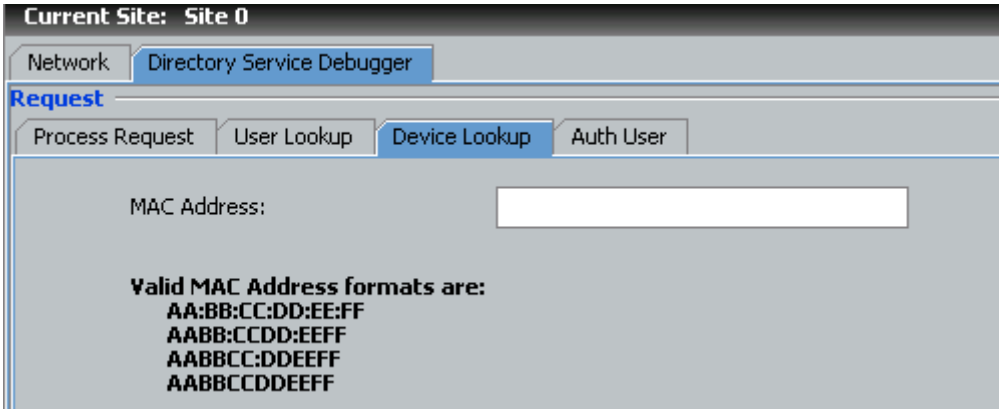
The **Results** field displays the lookup results, and the **Virtual Attributes** displays the attributes retrieved from the user's account record.

Testing a Device Lookup

To find out whether a device is known to Ignition, use the Device Lookup tab of the Advanced Troubleshooting window.

1. In Dashboard, click **Troubleshoot**.
2. Click the IP address of your Ignition Server.
3. Click the Directory Service Debugger tab.
4. In the debugger, click the **Device Lookup** tab.

Figure 44 The Device Lookup tab for verifying the machine's account exists

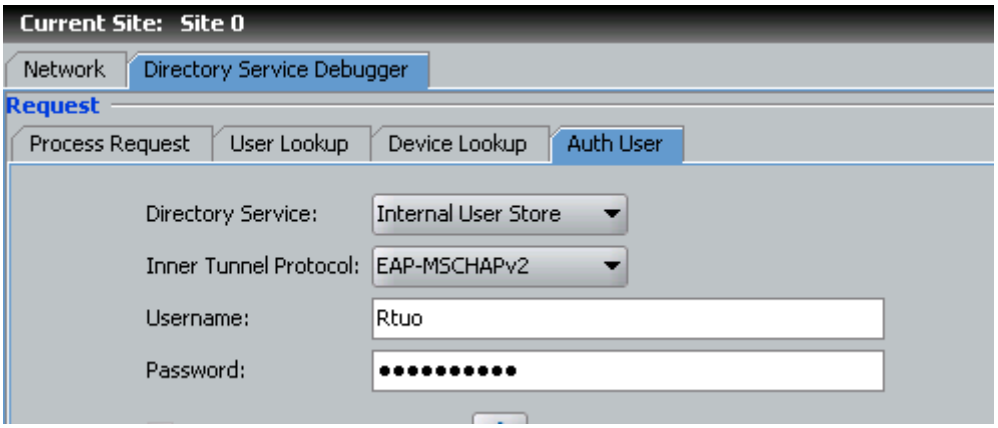


5. Type the **MAC address** of the device. Enter the address as a string of six octets. You may write the twelve characters without separators, or you may separate the octets with period, colon, or hyphen characters. Do not mix separator characters.
6. Click the **Send Request** button to test the authentication and lookup.

The **Results** field displays the lookup results, and the **Virtual Attributes** displays the attributes retrieved from the device record. For information on creating and editing device records see "Internal Devices" on page 119.

Testing a User Authentication

Figure 45 The Auth User tab for testing user authentication only (no lookup)



1. In Dashboard, click **Troubleshoot**.
2. Click the IP address of your Ignition Server.
3. Click the Directory Service Debugger tab.
4. In the debugger, click the **Auth User** tab.
5. Choose the directory service and authentication protocol.
6. Type the user's credentials and click Send Request. Results appear in the lower part of the window.

Authentication Services

This chapter explains how to set up authentication services such as RSA SecurID, Kerberos, and RADIUS proxy. By adding an authentication service to your access policy, you can combine specialized types of authentication with Ignition's directory-based authorization. You can look up and authenticate the user against your authentication service only, or you can split the lookup from the authentication, in which case Avaya Identity Engines Ignition Server validates the user's identity against your authentication server and then looks up the user's account in your LDAP or AD store to gather authorization-determining data.

Contents

- "Setting up a Kerberos Authentication Service" on page 177
- "Overview of Token Authentication in Ignition" on page 179
- "Configuring Token Authentication in Ignition" on page 180
- "Setting up a RADIUS Proxy Server" on page 189
- "Editing Authentication Service Configurations" on page 192
- "Renaming an Authentication Service" on page 193
- "Deleting an Authentication Service" on page 193
- "Managing a SecurID Authentication Service" on page 194
- "Setting Up Supplicants and Authenticators for SecurID Authentication" on page 195

Setting up a Kerberos Authentication Service

The Create Service Wizard guides you through the steps needed to connect Ignition Server to a Kerberos authentication service. To do this:

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.
2. Select the radio button for **Kerberos** and click **Next**.
3. In the Configure Kerberos Service window:

- Assign the authentication service a name in the **Name** field. This is the name you will use in your Ignition Server policy to specify that this Kerberos server should be used.
- Type the Kerberos **Realm** in all uppercase.
- If you want to authenticate all users as if they were members of the Kerberos realm, tick the **Replace Realm** checkbox. With this checkbox ticked, Ignition Server replaces the user's domain name with the domain name of the Kerberos server. This is useful if your users are submitting user names with various realms and you have made accounts for all of them in your Kerberos server. For example, if your Kerberos server EXAMPLE.COM contains all your user accounts, then, with this feature turned on, a user with the username ksmith@EXAMPLE.CO.UK is authenticated as ksmith@EXAMPLE.COM.
- If you want to send a regular “keepalive” ping, check the **Enable Keepalive** checkbox and specify a Keepalive **User Name** and **Password**. These are the user name and password of a test account in your authentication server. With Keepalive turned on, Ignition Server periodically sends an authentication request and, if successful, marks the service as *Connected* in the **Directory Services Status** tab. With this feature turned off, Ignition Server tests the connection only at the time you create it. You can test the connection at any time using the **Test Keepalive** button in this window, or using the Directory Service Debugger tab of Dashboard’s Troubleshoot view.
- For the primary Kerberos server, and optionally for the secondary Kerberos server, specify the **IP Address** and **Port**.
- Click the **Test Keepalive** button. Testing the connection might take a few minutes. If a configuration setting is incorrect, Ignition Server warns you. If you receive an error message, correct your settings and test again. If the error message persists, see “Problem: Errors Occur During Directory Service Set-Up” on page 457.
- Click **Next**.

4. The next window summarizes the connection settings of the service. Click **Finish**.

Your new service appears in the Directory Services list. A blue check mark in the Connected column indicates a successful connection.

Next Steps Add your Kerberos service to a directory set as explained in “Adding a Directories and Authentication Servers to a Directory Set” on page 167.

Overview of Token Authentication in Ignition

Ignition Server supports authentication using the RSA Authentication Manager (also known as the “RSA SecurID Server”) and other token servers. Depending on your Ignition Server directory set settings, you can authenticate the user against the token server only, or you can split the authentication from the user lookup, in which case Ignition Server authenticates against the token server and retrieves the user’s account data (criteria for the authorization decision) from an LDAP or AD directory.

To proceed immediately to the configuration instructions, see “Configuring Token Authentication in Ignition” on page 180.

How a User Logs In With Token Authentication

Ignition Server handles a token authentication login as follows:

1. The user attempts to connect to a network resource protected by an Ignition Server access policy. The access policy is set to require token-based authentication.
2. The supplicant prompts the user for credentials; user types login name and token code.
3. Ignition Server checks the user’s account credentials:
   - If the user account exists, Ignition Server forwards the username and tokencode to the token server to verify the credentials. For most RSA SecurID implementations, messages travel between Ignition Server and the token server via “RSA SDI authentication mode.” For other types of token servers, messages travel via RADIUS-PAP. If the token server is an RSA Authentication Manager, there may be a second challenge-response transaction to obtain the *new-PIN* or *next-tokencode*. If the authentication *succeeds*, Ignition Server checks whether your access policy also requires authorization (evaluation of user attributes from the user’s account in LDAP or AD), and, if so, it performs the required checks. If the authentication and (if present) authorization rules evaluate to “ALLOW,” Ignition Server returns a RADIUS *Access-Accept*, and the user is granted access. (Note that authorization is based on user data loaded from the LDAP or AD service specified as the *user lookup service* in the directory set.)
   - If the user account does not exist, if the authentication attempt fails, or if the Ignition Server authorization attempt fails, then Ignition Server returns a RADIUS *Access-Reject*, and the user is denied access.

Components Required for Token Authentication

Deploying token authentication in Ignition Server typically requires the following components:

- Ignition Server

- A token server such as an RSA Authentication Manager
- User tokens or another two-factor authentication credential
- Client PCs, each with an installed supplicant or VPN client that supports the desired authentication method. In most cases, the supplicant or client should support EAP-GTC authentication. Avaya recommends using a supplicant that supports PEAP/EAP-GTC so that the EAP-GTC transaction is enclosed in a secure PEAP tunnel.

Configuring Token Authentication in Ignition

This section explains how to set up Ignition Server to require token authentication. The required steps are

- "Prepare Your Token Server" on page 180
- Connect Ignition Server to your token server. Do one of:
  - "Connect Ignition Server to RSA Authentication Manager" on page 181; or
  - "Connect Ignition Server to Another Type of Token Server" on page 183
- "Add the Authentication Server to Your Directory Set" on page 184
- "Set Up an Access Policy That Uses Token Authentication" on page 185
- "Prepare Your Authenticators" on page 187
- "Connect Ignition Server to Your Authenticators" on page 188
- "Make Sure Token-Capable Clients Are Installed" on page 188
- "Perform an End-to-End Test" on page 188

Prepare Your Token Server

Install and configure your token server software and its required supporting software. Note the following tips:

1. Make DNS settings on your token server so that it can resolve the address of the Ignition Server RADIUS service.
2. Set your token server to recognize Ignition: In the RSA Authentication Manager configuration, you will add an *Agent Host* record to represent the Ignition Server. For other types of token servers, you will add the Ignition Server as a *RADIUS client* of the token server. In all cases, use the Ignition Server RADIUS port as the Ignition Server address. If you are running an Ignition Server HA pair, you must bind the Ignition Server RADIUS service to an Ignition Server VIP address, and use the VIP address as the *Agent Host/RADIUS client* address.

3. *Optional:* You can set Ignition Server to authenticate the user against your token server, but look up the user from a separate directory such as AD or LDAP. In this case, make sure each user has matching records in both the token server's user store and in the Ignition-accessible directory service. For each user, *the two user names must be identical*. For example, if you are using RSA SecurID and an LDAP directory service, give user *Mick Jones* two accounts: "mjones" in the RSA Authentication Manager user list and "mjones" in the LDAP directory.

Warning for Sites Running Ignition Server in HA Mode If you are running an HA pair of Ignition Server s, you must bind the Ignition Server RADIUS service to an Ignition Server VIP address (in other words, don't bind it to a physical port on the Ignition Server). At any given time, only one node in the pair — the primary node — can service RSA SecurID authentication requests. The VIP ensures that incoming authentication requests go to the primary node.

Next Steps Add your authentication service to a directory set as explained in either

- "Connect Ignition Server to RSA Authentication Manager" on page 181; or
- "Connect Ignition Server to Another Type of Token Server" on page 183.

Connect Ignition Server to RSA Authentication Manager

Before you can direct user authentication to your RSA Authentication Manager, you must define the RSA Authentication Manager in Ignition Server as an *authentication server*, as explained in the steps that follow.

Ignition Server supports the use of SecurID authentication with the EAP-GTC, and PEAP/EAP-GTC authentication protocols. For some types of token servers, you may elect to use the PAP authentication protocol, instead. Be aware that, if you use PAP, the *new-PIN* and *next-tokencode* modes are not supported.

Note: Messages indicating the RSA Authentication Manager's requests for *new-PIN* and *next-tokencode* are displayed in Ignition Server's *Security* log channel.

Before you begin: Make sure you have the login name, SecurID token, and SecurID PIN of one user in your RSA Server. To complete the connection, you must complete a successful authentication with this test user account.

1. In Dashboard, make DNS settings so that Ignition Server can resolve the addresses of your token server and your authenticators: In Dashboard's **Configuration** hierarchy tree, click the name or IP address of your node. In the **Nodes** panel, click the **System** tab, click the **DNS** tab, and click **Edit**. See "Editing Ignition Server's DNS Settings" on page 57 for details.

2. In Dashboard's **Configuration** hierarchy tree, expand **Directories** and expand **Directory Services**. In the Directory Services panel, click **New**.
3. In the Create Service Wizard window, click the **RSA Service** checkbox, and click **Next**.
4. The Configure RSA Service window appears:
   - Give your authentication server a **Name**. This is the name you specify in your Ignition Server access policy to specify that this RSA server handles authentication.
   - In the **Ignition Server** field, specify the host name or IP address of the Ignition Server interface through which the RSA Server can be reached. If Ignition Server is running in HA mode, this must be the host name corresponding to the VIP IP address. When running in HA mode you *cannot* use the name of a physical port on the Ignition Server.
   - Specify the **Configuration Directory**. This is the location (a directory on your network) from which you want Ignition Server to upload the RSA SecurID configuration files. The *sdconf.rec* file must be available in this directory, and the directory must contain only RSA configuration files. Everything in this directory is loaded. The contents of this directory must result in a zip file of less than 100 KB.
   - Click **Next**.
5. The next window summarizes the connection settings of the service. Click **Finish**. Your new service appears in the Directory Services list.
6. **Important:** To complete the connection, you must perform a successful test authentication against the RSA Server. This causes the RSA Server to transfer the node secret to the Agent Host (Ignition Server). After this transfer, the RSA Server requires the Agent Host to have knowledge of the correct node secret. Do this as follows:
   - In Dashboard, click **Troubleshoot**.
   - Click the IP address of your node and click the Directory Service Debugger tab, and click the Auth User tab.
   - In the **Directory Service** drop-down list, choose the name of the SecurID service you just created.
   - In the **Inner Tunnel** drop-down list, choose EAP-GTC (or choose PAP if Ignition Server communicates with your token server via PAP).
   - Type the **Username** of your test user. This account must exist in the RSA Server.
   - In the **Password** field, type the test user's PIN (if any) plus the current tokencode displayed on the test user's SecurID token.
   - Click **Send Request**. If the attempt succeeds, the RSA Server sends its node secret to Ignition, and your SecurID setup is complete. If the

attempt fails, delete and re-create the RSA service as explained in "Handling Changes to the Node Secret" on page 194.

Next Steps Next, you must add your authentication service to a directory set as explained in "Add the Authentication Server to Your Directory Set" on page 184.

Important! At any given time, Ignition Server can maintain a connection to *only one* RSA SecurID realm (consisting of an RSA Authentication Manager and its replicas).

Connect Ignition Server to Another Type of Token Server

This section explains how to connect to a token server so that Ignition Server and the token server communicate via PAP RADIUS messages. For RSA SecurID, this method of connection is not recommended because it does not support *new-PIN* and *next-tokencode* modes; instead, Avaya recommends that you follow the instructions in "Connect Ignition Server to RSA Authentication Manager" on page 181.

To set up a PAP RADIUS connection with a token server, define your token server in Ignition Server as an *authentication server*, as shown here:

1. Make sure your token server is set up and running with its authentication service exposed as a RADIUS server.
2. In Dashboard, make DNS settings so that Ignition Server can resolve the addresses of your token server and your authenticators: In Dashboard's **Configuration** hierarchy tree, click the name or IP address of your node. In the **Nodes** panel, click the **System** tab, click the **DNS** tab, and click **Edit**. See "Editing Ignition Server's DNS Settings" on page 57 for details.
3. In Dashboard's **Configuration** hierarchy tree, expand **Directories** and expand **Directory Services**. In the Directory Services panel, click **New**.
4. In the Create Service Wizard window, click the **Token Service** checkbox, and click **Next**.
5. The Configure Token Service window appears:
   - Give your token server a **Name** in Ignition. This is the name you specify in your Ignition Server policy to specify that this server is used for authentication.
   - Enter the token server's **Shared Secret**.
   - In the **Timeout** and **Maximum Retries** fields, specify how long Ignition Server should wait to retry after sending a request that fails to generate a response, and how many times to try again, if no response arrives.

- If you want to send a regular “keepalive” ping, check the **Enable Keepalive** checkbox and specify a **Keepalive User Name** and **Password**. These are the user name and password of a test account in your token server. With Keepalive turned on, Ignition Server periodically sends a RADIUS PAP authentication request and, if successful, marks the service as *Connected* in the Directory Services Status tab of Dashboard’s Monitor view. With this feature turned off, Ignition Server tests the connection only at the time you create it. You can test the connection at any time using the Test Keepalive button in this window, or using the **Directory Service Debugger** in the **Troubleshoot** view of Dashboard.
- For the primary authentication server, and optionally for the secondary authentication server, specify the **IP Address** and RADIUS **Port** of the token server.
- Click the **Test Keepalive** button. Testing the connection might take a few minutes. If a configuration setting is incorrect, Ignition Server warns you. If you receive an error message, correct your settings and test again. If the error message persists, see “Problem: Errors Occur During Directory Service Set-Up” on page 457.
- Click **Next**.

6. The next window summarizes the connection settings of the service. Click **Finish**.

Your new service appears in the Directory Services list. A blue check mark in the Connected column indicates a successful connection.

Next Steps Add your authentication service to a directory set as explained in the next section.

Add the Authentication Server to Your Directory Set

Authentication servers are included in Ignition Server policy by means of a directory set. Define your directory set for token-based authentication as shown here:

1. In Dashboard's Configuration hierarchy tree, expand Directories and expand **Directory Sets**.
2. In the Directory Sets panel, on the left, click the name of your directory set. (If you do not have a directory set, create one now.) With the directory set name selected, click the **Add** button on the right.
3. In the Directory Set Entry window, in the **User Lookup Service** drop-down list, select the store that holds your user records:
   - To authenticate users against the SecurID or token server *only*, choose *None*.

- To look up the user in AD, LDAP, or other directory, choose its directory service name. If you haven't connected Ignition Server to the desired service, see "Directory Services" on page 134 for instructions.

4. In the **Authentication Service** drop-down list, select the name of your RSA SecurID authentication service or token authentication service.



Note! If your **Directory Set Entry** includes both a lookup service and an authentication service, then, in order to authenticate successfully, a user must have accounts in both the authentication service and the directory service, and the user's two accounts must bear identical user names.

5. Click **OK**.

   The **Directory Set Entries** table in the Directory Sets panel shows that the authentication server has been paired with the selected directory service in the directory set. To find the authentication server name, locate the row for your directory service and check the name displayed in the **Authentication Service** column.

   Note that the authentication server can be used in any number of directory sets in Ignition.

6. Click **OK** to save the directory set.

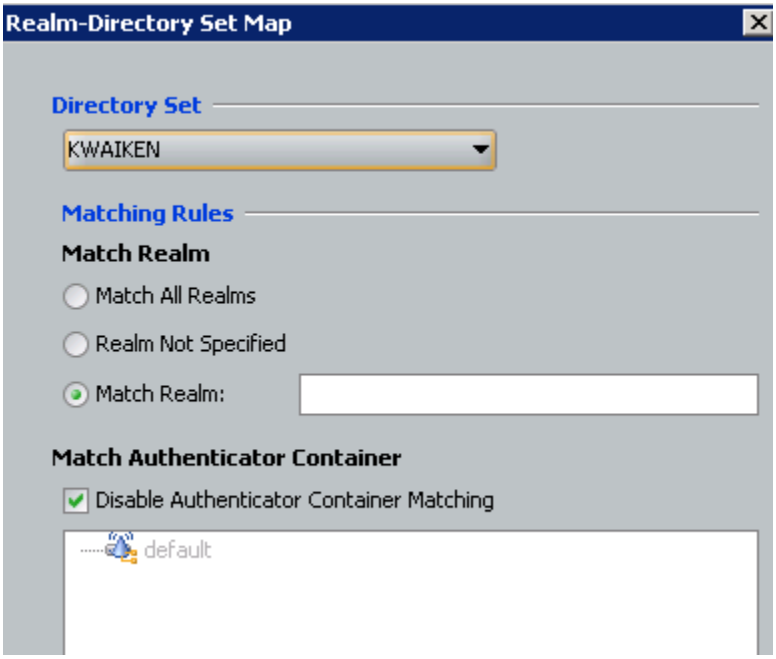
Next Step Set up your Ignition Server access policy as shown below.

Set Up an Access Policy That Uses Token Authentication

In order to require users to log in with two-factor authentication, your Ignition Server policy must require EAP-GTC credential validation (or PAP if your are connecting to the token service via PAP), and it must use the Ignition Server directory service that includes your token server. Set this up as follows:

1. If you do not have an access policy you wish to use, create one now: In Dashboard's **Configuration** hierarchy tree, click **Site Configuration**. In the Current Site panel, click **New**. In the New Access Policy window, type a **Name** for your policy and tick the **RADIUS** checkbox. Click **OK**.
2. Open your access policy: In Dashboard's **Configuration** hierarchy tree, expand the **Site Configuration**, expand **Access Policies**, expand **RADIUS**, and click the name of your policy. Click the **Authentication Policy** tab and click **Edit**.
3. In the Edit Authentication Policy window go to the **Authentication Protocols** section and choose your outer tunnel type and ("inner") credential validation type. To use token authentication, Avaya recommends that you use PEAP/EAP-GTC or NONE/EAP-GTC. If you have configured Ignition Server to communicate with your token server via PAP, then choose NONE/PAP.

4. In the **Certificate** section, select the certificate that you want to secure the outer tunnel, and in the **Ciphers** section, select the cipher types you want to allow. Click **OK**.
5. Go to the **Identity Routing** tab and click **Edit**.
6. In the Identity Routing Policy window, below the Authenticator Hierarchy / Realm / Directory Set mapping table, click **New**.
7. In the Realm-Directory Set Map window:
   - For **Directory Set**, choose the set that contains your token server. (This is the directory set you saved in Step 6 of the preceding procedure.)
   - Set the **Matching Realm** rule as appropriate for your policy. For a simple installation, click **Match All Realms**.
   - Set the **Match Authenticator Container** as appropriate for your policy. For a simple installation, click **Disable Authenticator Container Matching**.



   - Click **OK** to save the Realm-Directory Set Map.
8. Click **OK** to close the Identity Routing Policy window.

Next Step Make settings on your switches or VPN concentrator as shown below.

Prepare Your Authenticators

Your authenticators are your switches or VPN concentrators that send authentication requests to Ignition. In each such switch or VPN concentrator, set Ignition Server to act as the RADIUS Server. Do the following:

1. **RADIUS server**: In your switch or VPN configuration screen, find the setting for “RADIUS server” or “authentication server” and set it to the IP address where the RADIUS service is running on your Ignition Server. *Reminder*: If you are running an HA pair of Ignition Server s, be sure to bind the Ignition Server RADIUS service to a VIP address rather than a physical port.
2. **VPN group**: Some VPN concentrators require that you designate which user accounts will be RADIUS-authenticated. If given the choice between RADIUS authentication and token authentication, choose *RADIUS authentication* because communication with Ignition Server is done via RADIUS.
3. **DNS server**: Make DNS settings on each authenticator so that it can resolve the address of the Ignition Server RADIUS service.

Example: On a Cisco VPN 3000 concentrator, you would make the following settings:

- **RADIUS server**: In the VPN 3000’s **Configuration: System: Servers: Authentication: Modify** screen, set the **Server Type** to *RADIUS* and set the **Authentication Server** address to the IP address of the Ignition Server’s RADIUS service. Set the **Server Port** to Ignition’s RADIUS port number, and enter the shared secret in the **Server Secret** field.
- **VPN group**: In the VPN 3000’s **Configuration: User Management: Groups** screen, create a group to contain all users who will be token-authenticated. Click **Modify Group**, and in the **Configuration: User Management: Groups: Modify** screen in the **Identity** tab, set the group **Type** to Internal. Click the **IPSec** tab. Set **IPSec SA** to *ESP/IKE-3DES-MD5*, set **Tunnel Type** to *Remote Access*, and set **Authentication** to *RADIUS*.
- **DNS server**: In the VPN 3000’s **Configuration: System: Servers: DNS** screen, specify a DNS server that can resolve the address of the Ignition Server RADIUS service.

Next Step Make settings in Ignition Server to connect it to your authenticators as shown below.

Connect Ignition Server to Your Authenticators

For each switch, access point and VPN concentrator that will support token authentication, create an authenticator record in Ignition Server and assign to it the access policy you set up starting in Step 2 on page 175. See "Creating an Authenticator" on page 96 for instructions.

Next Step Install client-side tools as shown below.

Make Sure Token-Capable Clients Are Installed

Each end user who will authenticate with a token must have installed on his or her computer a supplicant or VPN client with token support. Consult your token server documentation for details.

Be aware of the following common limitations of supplicants with respect to token authentication:

- Not all supplicants support RSA's *new-PIN* and *next-tokencode* modes.
- Many supplicants store the user's most recently used credentials and automatically submit them at the next authentication attempt. For a time-based token, this means that the first attempt to authenticate usually fails because the supplicant sends an expired token code. When this happens, the user should retry the authentication. Most supplicants prompt the user for the new passcode.

Next Step Test your setup as shown below.

Perform an End-to-End Test

Send an authentication request to test your configuration. Do the following:

1. Use Ignition's built-in troubleshooting panel:
   - In Dashboard, click **Troubleshoot**.
   - In the **Troubleshoot** tree, click the IP address of your node and click the **Directory Service Debugger** tab, and click the **Process Request** tab.
   - In the **Directory Set** drop-down list, choose the name of the directory set that contains your token authentication service.
   - In the **Inner Tunnel** drop-down list, choose EAP-GTC (or PAP for some configurations).
   - Type the **Username** of your test user. This account must exist in the token server.
   - In the **Password** field, type the test user's PIN (if any) plus the current tokencode displayed on the test user's token.

- Click **Send Request**. The **Result** section displays the outcome.

2. Using a test computer with your supplicant or VPN client installed, try to connect to the network or VPN. For 802.1X, appropriate testing supplicants include the *Meetinghouse Aegis Supplicant* and the *OpenSEA XSupplicant*, both of which let you choose the inner and outer tunnel protocols for the authentication transaction. If you are using the Meetinghouse Aegis supplicant, note that by default it remembers the last credentials you passed in and tries to submit them the next time you try to connect. For a time-based token like SecurID, this means that the first attempt to authenticate will usually fail, because Aegis will pass in the expired token code. When this happens, you must retry, and Aegis will prompt you for a new tokencode. Note that some other supplicants, such as the Microsoft Windows built-in supplicant do not support EAP-GTC and therefore cannot be used with SecurID authentication.

Setup Complete Your token authentication setup is complete.

Troubleshooting Notes If the RSA service initializes correctly but all authentications fail, and the RSA server reports that the node secret has been sent to the client (the message “node validation failed” appears in RSA server logs), there is a possibility that the node secret is out of sync. Follow the steps in “Handling Changes to the Node Secret” on page 194 to re-create the service.

Setting up a RADIUS Proxy Server

A RADIUS proxy server forwards (or proxies) RADIUS requests to a remote RADIUS server for authentication. The Ignition Server can act as the RADIUS proxy server that forwards the authentication requests, or as the remote server that receives the authentication requests.

The following sections explain how to set up a RADIUS proxy server. The required steps are:

- “Creating a RADIUS Proxy Authentication Service” on page 189
- “Creating a Directory Set” on page 191
- “Adding the RADIUS Proxy Server to a Directory Set” on page 191
- “Creating a RADIUS Proxy Server Access Policy” on page 191

After you set up the RADIUS proxy server, you must perform some configuration tasks on the remote RADIUS server. See “Remote RADIUS Server Configuration” on page 192.

Creating a RADIUS Proxy Authentication Service

The Create Service Wizard guides you through the steps needed to create a RADIUS proxy Authentication Service.

1. In the Dashboard **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.
2. Select the radio button for **RADIUS Proxy Service** and click **Next**.
3. In the **Configure RADIUS Proxy Service** window:
   - Assign the authentication service a name in the **Name** field. This is the name you will use in your Ignition Server policy to specify that this RADIUS proxy server should be used.
   - Enter the **Shared Secret** for the RADIUS proxy server.
   - If you want to send a regular “keepalive” ping, check the **Enable Keepalive** checkbox. Optionally, you can specify a **Keepalive User Name** and a **Keepalive Password**. These are the user name and password of a test account in your authentication server.

     With Keepalive turned on, Ignition Server periodically looks up the supplied username/password on the remote server to determine the reachability, and if successful, marks the service as *Connected* in the **Directory Services Status** tab. By default, Ignition Server uses a predefined username & password (idengines/idengines) to run the keepalive. If you entered a Keepalive User Name and a Keepalive Password, Ignition Server uses these credentials to run the keepalive.
     Note: The user credentials you enter to test keepalive do not have to be valid credentials. A reject message from the remote server for looking up invalid credentials is sufficient to determine the reachability.

     With Keepalive turned off, the Ignition Server assumes that the remote server is always reachable and marks it as *Connected*. You can test the connection at any time using the **Test Keepalive** button in this window, or using the Directory Service Debugger tab of Dashboard’s Troubleshoot view.
     Note: Avaya recommends that you enable keepalive if you have multiple remote servers that receive requests. If one server is reported down, the requests can be proxied to the next available proxy server as defined in the directory set. If you do not enable keepalive, the Ignition Server assumes the remote server is always connected and the requests may get dropped if the remote server health status is not determined.
   - For the primary RADIUS proxy server, and optionally for the secondary RADIUS proxy server, specify the **IP Address** and **Port**. If both the primary and secondary servers are configured and the Keepalive is not enabled, RADIUS proxy authentication attempts will occur with the primary server only. To ensure that authentication with the secondary server occurs following a failed authentication attempt with the primary server you must enable the Keepalive mechanism.

- Click the **Test Keepalive** button. Testing the connection might take a few minutes. If a configuration setting is incorrect, Ignition Server warns you. If you receive an error message, correct your settings and test again. If the error message persists, see "Problem: Errors Occur During Directory Service Set-Up" on page 457.
- Click **Next**.

4. The next window summarizes the connection settings of the service. Click **Finish**.

Your new service appears in the Directory Services list. A blue check mark in the Connected column indicates a successful connection.

Creating a Directory Set

If you do not have a directory set, create one now. To create a directory Set to include the RADIUS proxy server, see "Adding a Directory Set" on page 166.

Adding the RADIUS Proxy Server to a Directory Set

Add the RADIUS proxy server to a directory set to specify that the RADIUS proxy server is the authentication service that verifies user credentials. You can add multiple remote servers to a directory set. Each remote server can handle different realms, or multiple remote servers can support the same realm to handle a fail-over scenario. When you add a RADIUS proxy server to a directory set, ensure that the **User Lookup Service** field is set to **None**. **Note**: You cannot add another type of directory service to a Directory set that contains a proxy service.

To add the RADIUS proxy server to a Directory Set, see "Adding a Directories and Authentication Servers to a Directory Set" on page 167.

Creating a RADIUS Proxy Server Access Policy

The next step is to create an Access Policy that includes the RADIUS proxy server. An Ignition Server access policy consists of an authentication policy, an identity routing policy (user lookup policy), a user authorization policy, and other optional policies.

The decision on whether to proxy an incoming request or do local authentication comes from the information in the Identity Routing Policy. The Identity Routing Policy tells the Ignition Server which directory set to search for the user account, based on the realm (domain) name passed with the user name.

When you create your Identity Routing Policy, use the directory set that includes the RADIUS proxy server. In the Realm-Directory Set Map window, in the Match Realm section, specify a particular realm. The proxy server will forward any requests that match that realm to the remote server.

For more information about Access Policies, refer to the "User Authentication Policy" chapter.

Remote RADIUS Server Configuration

After you set up the RADIUS proxy server, you must perform some configuration tasks on the remote RADIUS server.

Creating an Authenticator

For the remote RADIUS server, the proxy (forwarding) server acts as an authenticator. Create an authenticator similar to creating a regular authenticator, that points to the proxy server. From the Dashboard, go to **Configuration > Site Configuration > Authenticators** and click **New**.

Creating an Access Policy

Assign an Access Policy that is capable of handling authentication requests from the proxy server. Create a regular Access Policy as you would for any regular authenticator and configure the necessary authentication and authorization policies. Make sure that the shared secret configured here matches the shared secret as configured at the forwarding server's proxy service.

Proxying of MAC Authentication Requests

MAC authentication is typically used for devices that are incapable of performing 802.1X authentication. MAC authentication requests are also RADIUS requests. MAC authentication verifies that the MAC address submitted by a connecting client device matches an entry on your list of known MAC addresses. Using RADIUS proxy service, Ignition Server can also proxy the MAC authentication requests to a remote server. To proxy MAC authentication requests, enable RADIUS authentication for the authenticator and assign the access policy that is configured to use a proxy directory set. Do not enable MAC authentication for the authenticator which would otherwise do a local MAC authentication. On the remote server, enable MAC auth for this authenticator (proxy server) and configure the necessary MAC authentication policy.

Editing Authentication Service Configurations

You can edit your connection to an authentication service. Use the following steps:

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. The Directory Services window displays the current set of configured directory services and authentication services.
2. Select the entry for the service you want to edit. Click **Edit**.

3. In the **Directory Services Details** Window with the details of the selected authentication service.
4. Edit the details of the service as required.
5. Click **OK** to apply your changes.

Renaming an Authentication Service

When you rename an authentication service, Ignition Server uses the updated name for the authentication service in all the directory set(s) to which the authentication service belongs.

To Rename an Authentication Service:

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**.
2. In the **Directory Services panel**, select the entry you want to rename.
3. Click **Edit**. Ignition Server displays the details for the selected authentication service.
4. Enter a different name for the authentication service.
5. Click **OK**.

Deleting an Authentication Service



Important! If you delete an RSA authentication service from Ignition Server and you want to re-create it, you *must* follow the steps in "Handling Changes to the Node Secret" on page 194 to re-create it.

To Delete an Authentication Service:

1. Before you delete an authentication service, remove it from the Directory Sets to which it belongs. Use the **Directory Sets** panel to check whether the service is a member of any directory set. Remove it from each set that contains it. (In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Sets**.)
2. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. The Directory Services window displays the current set of configured authentication services.
3. Select the authentication service to be deleted.
4. Right-Click on the authentication service to be deleted and select on **Delete**. Alternatively, click the service name and the **Delete** button.

Managing a SecurID Authentication Service

Handling Changes to the Node Secret

Anytime an action is taken that causes the node secret of your RSA Authentication Manager to change, you must take the following actions.

Note! If you update the RSA service and the node secret does *not* change, no further action is required.

If you clear the node secret on the RSA Server, then you must:

1. Delete the RSA Service on Ignition:
   - In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**. The Directory Services window displays the current set of configured directory services and authentication services.
   - Select the entry for the RSA SecurID service
   - Click **Delete**.
2. Delete the agent host on the RSA server.
3. Create the agent host on the RSA server.
4. Reboot the Ignition Server.
5. Create the RSA Service in Dashboard as shown in "Connect Ignition Server to RSA Authentication Manager" on page 181.
6. **Important:** To complete the connection, you must perform a successful test authentication against the RSA Server. This causes the RSA Server to transfer the node secret to the Agent Host (Ignition). After this transfer, the RSA Server requires the Agent Host to have knowledge of the correct node secret. Do this as follows:
   - In Dashboard, click the **Troubleshoot** button.
   - Click the IP address or name of your node and click the Directory Service Debugger tab.
   - Click the **Auth User** tab.
   - In the **Directory Service** drop-down list, choose the name of the SecurID service you just created.
   - In the **Inner Tunnel** drop-down list, choose EAP-GTC (or PAP if Ignition Server communicates with the token server via PAP).
   - Type the **Username** of your test user. This account must exist in the RSA Server.

- In the **Password** field, type the test user's PIN (if any) plus the current tokencode displayed on the test user's SecurID token.
- Click **Send Request**. If the attempt succeeds, the RSA Server sends its node secret to Ignition, and your SecurID setup is complete. If the attempt fails, try deleting and re-creating the RSA service.

Setting Up Supplicants and Authenticators for SecurID Authentication

When setting up a user's supplicant to support SecurID authentication, set it to use PEAP/EAP-GTC authentication and make sure its timeout settings are set for periods long enough to accommodate the SecurID credential check. Each timeout period should be roughly two times the cycle time of the token.

Procedure: To set up the supplicant and authenticator for SecurID authentication, do the following:

- Set your supplicant to use EAP-GTC authentication. In OpenSEA's XSupplicant, create a **Profile** that uses the **Tunnel Protocol** *EAP-GTC*.
- Set your supplicant timeouts to slightly more than two times the cycle time of the token. In OpenSEA's XSupplicant, in the **Advanced: Internals** tab, set the **Auth Period** and **Idle Period** to appropriate periods. For example, a typical RSA token uses a one-minute cycle time, so you would set the **Auth Period** to 135 seconds (two minutes, plus 15 seconds to account for possible lag in user response) and **Idle Period** to 135 seconds.
- On your authenticators, set the timeouts to two times the cycle time of the token. For some types of authenticators, this might mean changing the *number of retransmits* that are sent before a default failure, and changing the *length of the retransmit timers*. For example, on a typical authenticator using a default setting of 3 retransmits, you can change the length of the retransmit timer to 45 seconds to achieve the correct timeout period.

Virtual Groups and Attributes

This chapter explains how to create and apply the virtual groups and attributes that Avaya Identity Engines Ignition Server uses to evaluate users, devices, and group memberships in order to make authorization decisions.

Contents

- "Introduction to Virtual Groups and Attributes" on page 197
- "How Ignition Server Handles Multiple Directories" on page 197
- "Virtual Groups" on page 198
- "User Virtual Attributes" on page 203
- "Device Virtual Attributes" on page 208

Introduction to Virtual Groups and Attributes

Virtual groups, user virtual attributes, and device virtual attributes are Ignition Server's mechanisms for abstracting, or standardizing, group and attribute names across multiple directory services. A virtual group can be mapped to one or many groups in one or many directory services, allowing you to treat them as a single group in your policies. Likewise, a virtual attribute maps to an attribute or attributes in your directory services, so that when you write an authorization policy you refer to a single virtual attribute name, not the various, underlying attribute names in each store.

In cases where you must choose between using a virtual group or a virtual attribute in your authorization, Avaya recommends that you use a *virtual group*, as it offers richer support for group nesting and multiple group memberships.

How Ignition Server Handles Multiple Directories

Ignition Server retrieves users' identities, attributes, and group associations from one or many of the following:

- Microsoft Active Directory services
- LDAP directory services
- Ignition Server's internal database (the internal data store)

The use of multiple dissimilar directory services makes it difficult to write consistent access policies. When you write a policy that spans multiple directories, you need a uniform method of referring to user groups and attributes. Otherwise you would have to write a series of policies, one per directory, with each access policy using the attribute or group names local to that directory.

This chapter explains how to set up virtual attributes and virtual groups. For instructions on using virtual groups and attributes in your policies, see Chapter , "User Authentication Policy".

Virtual Groups

Virtual groups are the Ignition Server mechanism for abstracting, or standardizing, group names and role names across multiple directory services.

Group naming and role naming (as well as the mechanism used to record group membership or role) is often inconsistent across the various directories (directory services) that store users in an organization. For example, your local LDAP store may designate an administrator by placing his user record in the DN, "`ou=admin,ou=Users, dc=company,dc=com`", while the LDAP store of your Atlanta office designates an administrator by adding the label "AdminGroup" to the `nsRole` attribute of his user record.

Ignition Server's virtual groups allow you to write authorization and provisioning policies that span users stored in disparate data stores, and handle them consistently, even if group designations are implemented using different approaches in different stores. To address the administrator problem shown above, you might create a virtual group, "Administrators" and map it to the DN, "`ou=admin,ou=Users,dc=company,dc=com`" in your local store and to the `nsRole` value "AdminGroup" in your Atlanta store.



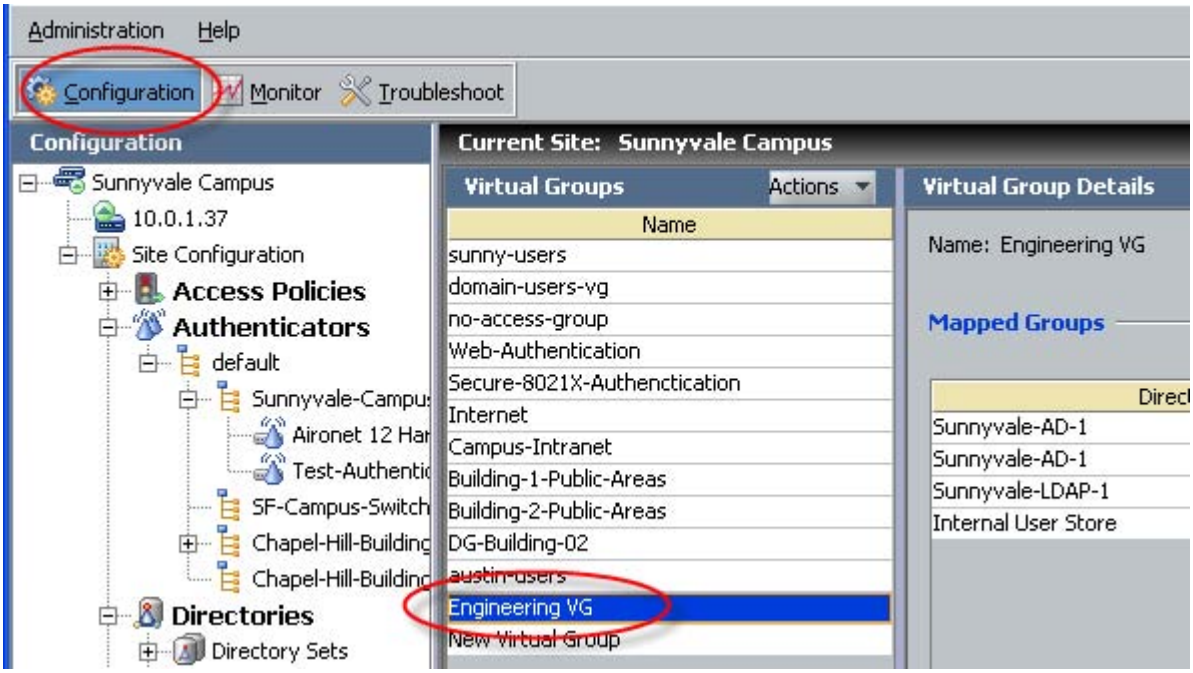
Tip. Ignition Server provides a similar mechanism for abstracting attribute names. See "User Virtual Attributes" on page 203.

The sections that follow explain how to manage virtual groups.

Browsing Virtual Groups

In Dashboard's Configuration hierarchy tree, expand Directories, expand Virtual Mapping, and click Virtual Groups.

Figure 46 Browsing Virtual Groups



Ignition Server Dashboard allows you to add, rename, and delete virtual groups. To sort the **Directory Service** and **Group DN** lists into ascending or descending order, click the title bar of the column.

Mappable Group Types for Ignition Server Virtual Groups

The table below lists the types of group designations Ignition Server supports for each data store type. In Ignition Server Dashboard, the underlying group type is indistinguishable; all types appear together in the Map Groups window.

Types of Mappable Groups for Each Directory Type

	AD	Sun	Open-LDAP	Novell eDirectory	Oracle OID
Group saved as an `organization` or `organizationalUnit` entry	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
Group saved as a `groupOfNames` or `groupOfUniqueNames` entry with members listed in `member` or `uniqueMember`	No	**Yes**	**Yes**	**Yes**	**Yes**
Group listed in the user's record, in the `nsRole` attribute (object class is `ldapsubentry`)	No	**Yes**	No	No	No

Types of Mappable Groups for Each Directory Type

	AD	Sun	Open-LDAP	Novell eDirectory	Oracle OID
AD group (`group`) listed in the user's record, in the `memberOf` attribute	**Yes**	No	No	No	No
Novell static groups (object class is `groupOfNames`)	No	No	No	**Yes**	No
Novell dynamic groups (object class is `dynamicGroupAux`)	No	No	No	**Yes**	No
Novell rbsRoles (v1 and v1.x) or Novell rbsScopes (object class is `groupOfNames`)	No	No	No	**Yes**	No



Tip: How to check AD primary group membership

- **Active Directory Primary Group**: On your Active Directory (AD) server, launch the Active Directory Users and Computers snap-in. Click on **Users** in the AD tree. Double-click the name of the user you want to inspect, and AD opens the user's Properties window. Click the **Member Of** tab. The bottom section of the window shows the user's primary group assignment.

Adding a New Virtual Group

In order to add a new virtual group,

1. Open the **Virtual Groups tab** (In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click Virtual Groups.)
2. Click **Actions → Add a New Virtual Group.**
3. Enter a unique name in the Add a New Virtual Group window.
4. Click **OK**.

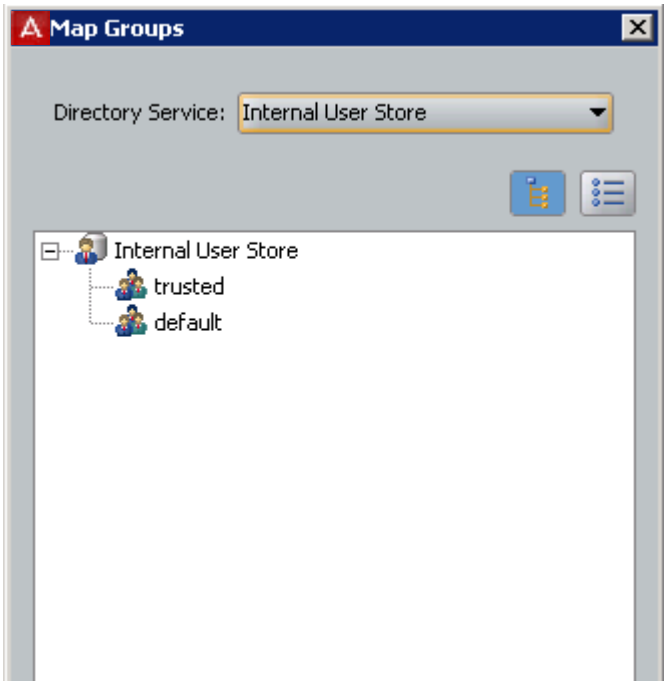
Ignition Server Dashboard displays the newly added virtual group to the list of virtual groups that appear in the **Virtual Groups** panel.

Mapping Groups From a Directory Service

You can map groups from your directory services onto this new virtual group. In order to map groups from available directory services,

1. Select the virtual group from the displayed list of virtual groups in the **Virtual Groups tab**.
2. Click the **Add** button in the **Virtual Group Details** section of this window. The **Map Groups** Window appears.

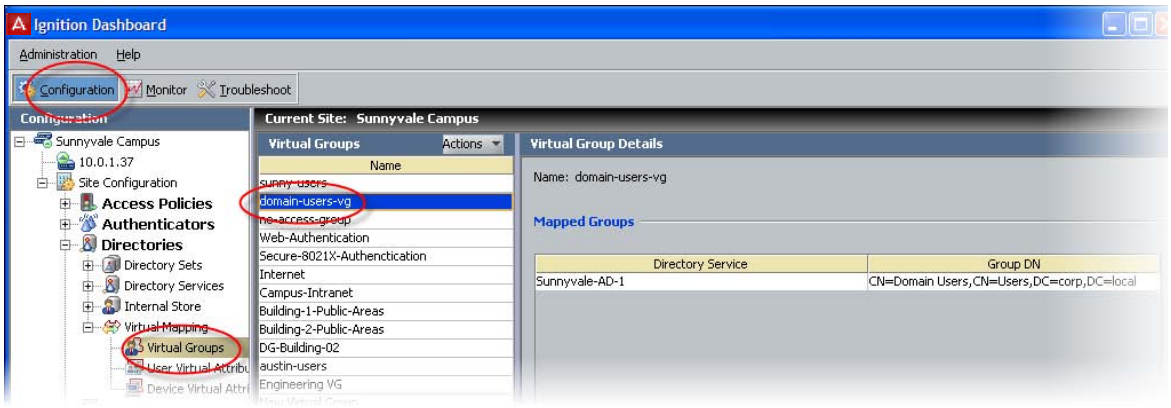
Figure 47 Selecting Groups to Map to a Virtual Group



3. Select the **directory service** from the pull-down menu.
4. Choose whether you want to view the contents of the selected directory service as a tree or as a list.
   - (**Tree** View): The names are shown in tree fashion (default). You can select only one entry each time you open this view of the directory service. The **OK** button comes into focus only when you select the group which is a “leaf” in the “tree”.
   - (**List** View): In this view you can make multiple selections. Note that the groups in the **internal data store** can only be viewed as a list.
5. Select to the required group by doing one of the following:
   - In the Tree view, choose a “leaf” from the tree.
   - In the List view, choose a set of available groups.
6. Click **OK**.

The above figure shows the chosen groups that are mapped to the virtual group “Engineering”. It shows that the administrator has mapped the virtual group “Engineering” to groups from three directory services, Internal User Store, Sunnyvale-AD-1, and Sunnyvale-LDAP-1.

Figure 48 Example Mapping For a Virtual Group



Renaming a Virtual Group

Important! Before you rename any virtual group, make sure it is not used in any of your Ignition Server authorization policies. (Renaming a virtual group breaks the authorization rules that depend on that virtual group. See "Problem: Authorization policy stops working unexpectedly" on page 453 for troubleshooting instructions.)

In order to rename an existing virtual group,

1. Open the **Virtual Groups tab**. (In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click Virtual Groups.
2. Right-Click on the virtual group to be renamed and select **Rename Virtual Group**. Alternatively, click **Actions → Rename Virtual Group.**
3. Enter a unique name in the Rename window.
4. Click **OK**.

Ignition Server Dashboard displays the updated name for the virtual group in the list of virtual groups that appear in the **Virtual Groups tab**.

Deleting a Virtual Group

Important! Before you delete any virtual group, make sure it is not used in any of your Ignition Server authorization policies. (Deleting a virtual group breaks the authorization rules that depend on that virtual group. See "Problem: Authorization policy stops working unexpectedly" on page 453 for troubleshooting instructions.)

In order to delete an existing virtual group,

1. Open the **Virtual Groups tab**. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click Virtual Groups.

2. Right-Click on the attribute to be renamed and select **Rename Virtual Group**. Alternatively, click **Actions → Delete Virtual Group.**
3. Confirm your action and Click **OK**.

Ignition Server Dashboard deletes the virtual group from the list of virtual groups that appear in the **Virtual Groups tab**.

User Virtual Attributes

User virtual attributes let you retrieve account details during user lookup and use these details in your authorization rules and send them as provisioning values. To use a user virtual attribute in an authorization rule, use the Constraint Details window, described in "Set Up Rule Details" on page 243. To use a user virtual attribute in provisioning, use the Outbound Value Instance window, described in "Passing Value from the User Record or Device Record to an Outbound Value" on page 262.

The following sections show you how to create and map virtual attributes:

- "About User Virtual Attributes" on page 203
- "Browsing User Virtual Attributes" on page 204
- "Adding a New User Virtual Attribute" on page 205
- "Mapping Directory Service Attributes to User Virtual Attributes" on page 205
- "Renaming a User Virtual Attribute" on page 207
- "Deleting a User Virtual Attribute" on page 208

About User Virtual Attributes

User virtual attributes are the Ignition Server mechanism for abstracting, or standardizing, attribute names across multiple directory services. You must define a user virtual attribute for each directory service field whose value you want to use in:

- authorization rules or
- outbound values for provisioning.

Group and attribute naming is often inconsistent across the various directories (directory services) that store users in an organization. For example, your local LDAP store may keep employee id numbers in the attribute, `employeeId`, while the LDAP store of your Atlanta office stores them in `employeeNumber`.

Ignition Server's user virtual attributes allow you to write authorization and provisioning policies that span users stored in disparate data stores, and handle them consistently, even if attributes are named inconsistently. To

address the `employeeId` / `employeeNumber` problem shown above, you would use a user virtual attribute, "*Employee-ID*" and map it to `employeeId` in your local store and to `employeeNumber` in the Atlanta store.

The initial list of user virtual attributes includes mappings to the user fields in the internal data store only. You may add to these mappings.

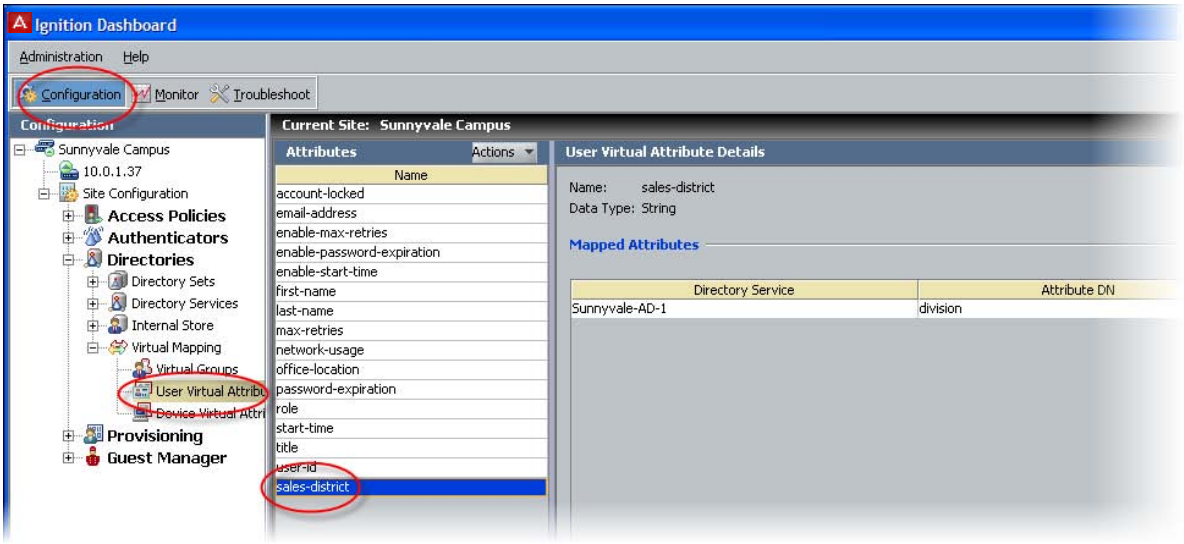
Tip. Ignition Server provides a similar mechanism for abstracting group names. See "Virtual Groups" on page 198.

Browsing User Virtual Attributes

To see the existing list of user virtual attributes:

1. In Dashboard's Configuration hierarchy tree, expand Directories, expand Virtual Mapping, and click **User Virtual Attributes**. The User Virtual Attributes panel appears.

Figure 49 User virtual attributes list



2. In the **Attributes** list on the left, scroll to find the desired attribute and click its name. The **User Virtual Attribute Details** pane shows:
   - **Name**: the name of the attribute. This name is used in your authorization rules and outbound attribute mapping rules.
   - **Data Type**: the data type of the attribute. When you create the attribute, you set its data type to a type that is compatible with the directory fields you plan to map to it.
   - **Mapped Attributes** table: In this table, each row represents one mapping of this user virtual attribute to a field in a data store. The **Directory Service** column shows the name of the data store, and the **Attribute DN** column shows the mapped field in the data store.

You may sort the **Directory Service** and **Attribute DN** lists in ascending or descending order by clicking the title bar of the column.

Adding a New User Virtual Attribute

In order to add a new user virtual attribute, do the following:

1. Click **Actions → Add A New Virtual Attribute** in the **User Virtual Attributes tab**. Ignition Server displays a dialog box requiring a name for the new user virtual attribute and its data type.
2. Enter a unique name for the new user virtual attribute in the dialog box. This name is used in your authorization rules and outbound attribute mapping rules.
3. Choose the data type for the new user virtual attribute by selecting from the drop down list. Below we list the data types for virtual attributes. Ignition Server follows the LDAP standard for data types and adds two types not defined in LDAP: the MAC address and VLAN data types. The types are:
   - **String**: LDAP-standard format
   - **Integer**: LDAP-standard format
   - **Boolean**: Uses the standard LDAP format, which looks like “true” or “false.”
   - **MAC address**: handles a device address using the commonly accepted formats. See “Allowed MAC Address Formats” on page 333.
   - **Date and time**: LDAP-standard format
   - **VLAN**: handles both numeric VLAN IDs and string-formatted VLAN labels. If the LDAP store provides a numeric value, Ignition Server considers it a VLAN ID, and if the store provides a string, Ignition Server considers it a VLAN label. If both a VLAN ID *and* a VLAN label are defined, the VLAN ID takes precedence in Ignition Server.
   - **Multi-valued string**: LDAP-standard format



After you create the attribute you cannot change its data type, you must instead delete the virtual attribute and recreate it.

4. Click **OK**. Ignition Server displays the new user virtual attribute in the list.
5. Next, you must map the attribute as shown in Mapping Directory Service Attributes to User Virtual Attributes.

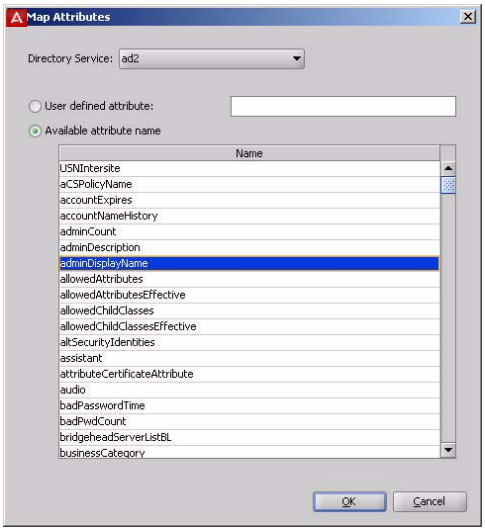
Mapping Directory Service Attributes to User Virtual Attributes

To map attributes (Distinguished Names) from a directory service object to a user virtual attribute, do the following:

1. Select the user virtual attribute from the **Attributes** list in the **User Virtual Attributes tab**.

2. Click the **Add** button in the **User Virtual Attributes Details** section of this window. The **Map Attributes** window appears.
3. Use the drop down list to select the appropriate directory service.
4. Click a radio button:
   - Choose **User defined attribute** if you want to manually specify the mapped attribute. In the field at the right, type the distinguished name of the attribute. Type the DN as it is defined in your LDAP directory.
   - Choose **Available attribute name** if you want to pick the directory attribute from a list. Ignition Server displays the available attributes (distinguished names). Click one to choose it.

Figure 50 Selecting Attributes to Map to a User Virtual Attribute



To sort the attribute list into ascending or descending order, click the title bar ("**Name**") of the column.

If the list is empty, it means problems occurred during the attempt to retrieve attribute names. An error message reports the nature of the problem. If Ignition Server was unable to parse your schema, you can work around the problem by clicking **User defined attribute** and typing the attribute name manually. If Ignition Server was unable to connect to your directory, edit your Directory Service definition to fix the connection, and test it as shown in "Testing a Directory Service Connection" on page 172.

Mind your data types! Make sure the data type of the LDAP attribute matches the data type of the virtual attribute. A mismatch may result in an undefined virtual attribute at user login time. This does not stop authorization, but the authorization fails if the attribute was required for the decision.

5. Click **OK** to dismiss the Map Attributes window.
6. You may map additional directory attributes to your virtual attribute. For example, if some user accounts are in an LDAP store and some in an AD store, you can map a virtual attribute “telephone-number” to an appropriate field in each store. To do so, click **Add** again, and repeat the steps above. (Note that your virtual attribute can only retrieve *one attribute* from each directory.)

   If you like, you can test your virtual attribute as shown in “Testing a User Lookup” on page 174.

Next Steps You can now use your virtual attribute in one of the following ways:

- As an input to your *authorization decision*: Add the virtual attribute to a rule constraint as shown in “Set Up Rule Details” on page 243; or
- As a *provisioning value* to be sent to an authenticator or to the user’s supplicant: Add the virtual attribute to an outbound value as shown in “Passing Value from the User Record or Device Record to an Outbound Value” on page 262.

Renaming a User Virtual Attribute

Important! Before you rename any virtual attribute, make sure it is not used in any of your Ignition Server authorization policies. (Renaming a user virtual attribute breaks the authorization rules that depend on that attribute. See “Problem: Authorization policy stops working unexpectedly” on page 453 for troubleshooting instructions.)

In order to rename a user virtual attribute,

1. Open the User Virtual Attributes tab. (In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click User Virtual Attributes.)
2. Click the attribute you want to rename.
3. Right-Click on the attribute to be renamed and select **Rename Virtual Attribute**. Alternatively, Select **Actions → Rename Virtual Attribute**. The Rename Dialog box appears.
4. In the **Rename** dialog, enter the new name.
5. Click **OK**.

Ignition Server updates the name for the user virtual attribute in the displayed set of user virtual attributes in the **User Virtual Attributes tab**.

Next Steps For each rule constraint that used the virtual attribute, you must update the rule to use the new virtual attribute name. See “Creating a RADIUS User Authorization Policy” on page 242.

Deleting a User Virtual Attribute

Important! Before you delete any virtual attribute, make sure it is not used in any of your Ignition Server authorization policies. (Deleting a user virtual attribute breaks the authorization rules that depend on that attribute. See “Problem: Authorization policy stops working unexpectedly” on page 453 for troubleshooting instructions.)

In order to delete a user virtual attribute,

1. Open the **User Virtual Attributes tab**. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click User Virtual Attributes.
2. Click the virtual attribute you want to delete.
3. Right-Click on the attribute to be deleted and select **Delete Virtual Attribute**. Alternatively, Select **Actions → Delete Virtual Attribute**. The Delete Dialog box appears.
4. Click **OK**. Ignition Server deletes the user virtual attribute.

Next Steps For each rule constraint that used the virtual attribute, you must update the rule so that it no longer uses the virtual attribute. See “Creating a RADIUS User Authorization Policy” on page 242.

Device Virtual Attributes

Device virtual attributes expose fields in your device records so that Ignition Server authorization rules can evaluate these fields. This section explains how to create device virtual attributes.

For certain attributes, you need not define a virtual attribute: By default, Ignition Server includes device definitions for the standard attributes (device-address, device-name, device-vlan, source, and type) of devices defined in the Ignition Server internal store. For a complete list, see “Device Attributes” on page 236.

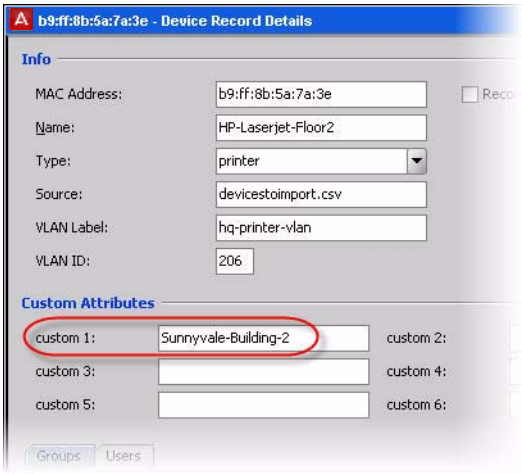
Browsing Virtual Attributes for Devices

In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click Device Virtual Attributes. In the Attributes list, click the name of a virtual attribute to select it. The right side of the window displays the underlying field mapped to the virtual attribute. When you write a rule in the Constraint Details window, you see the virtual attribute name. When you edit a device in the Device Record window, you see the name that is listed here in the **AttributeDN** column.

Adding Virtual Attributes for Devices

To define a device virtual attribute, follow these steps:

1. To begin, note the name of the device field you plan to use. In Dashboard's Configuration hierarchy tree, expand Directories, expand Internal Store, and click **Internal Devices**. Click a device to select it, and click **Edit**. In the device record Edit window, look at the **Custom Attribute** section and note the name of the field you want to use. In this example, the field **custom1** is used to store the building name of each device (the name of the building where the device is located).



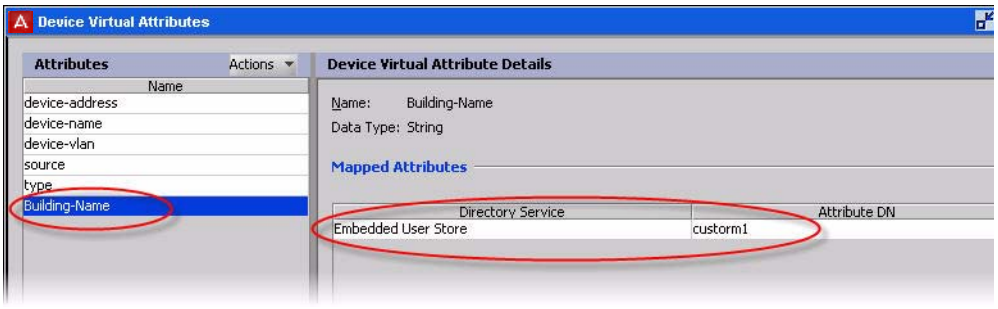
2. Create your virtual attribute:
   - In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click User Virtual Attributes.
   - In the Device Virtual Attributes panel, select the command **Actions: Add a New Device Virtual Attribute**.
   - In the Add Device Virtual Attribute window, type a name for the attribute and select its data type. This is the name you use in your authorization rules. For this example, we'll use the name, "Building-Name". Click **OK**.



3. Map your virtual attribute to an actual field in a database:
   - In the Device Virtual Attributes panel, click the name of your new attribute to select it, and click the **Add** button.

- In the **Map Attributes** window, choose *Internal User Store* as the **Directory Service**, click on the name of the device field you want to map, and click **OK**. In this example, we use the field, *custom1*.

Figure 51 Mapping a Device Virtual Attribute



For instructions on using the attribute in an authorization rule, see “Using a Device Attribute in a Rule” on page 250.

User Authentication Policy

Avaya Identity Engines Ignition Server allows you to set and enforce user authentication requirements, authorization rules, and provisioning rules. This chapter introduces policy management and describes how to create your authentication policy.

Contents

- "How Ignition Server Authenticates and Authorizes a User" on page 212
- "Introduction to Policy Management" on page 213
- "The Access Policy Panel" on page 213
- "Understanding Authentication Policy" on page 214
- "Creating an Authentication Policy" on page 219
- "Understanding Identity Routing Policy" on page 220
- "Creating an Identity Routing Policy" on page 223

Additional Policy Topics

Subsequent sections of this document cover the remaining topics in access policy:

- "User Authorization Policy" on page 227 explains how to set policies that control user access based on user attributes, transaction details, and network hardware specifications.
- "Provisioning Policy" on page 253 explains how to set policies that assign users to VLANs and/or set authenticator attributes.
- "VLAN Assignment" on page 287 shows how to set up provisioning policies that assign each user to an appropriate VLAN.
- "Authentication Services" on page 177 shows how to set up strong authentication that requires the user to prove his or her identity using an RSA SecurID or other token.
- "Windows Machine Authentication" on page 295 explains how to write a policy that requires each connecting device to have a valid Active Directory account.

- "MAC Authentication" on page 321 shows how to set up MAC address checking and how to set up an access policy for devices incapable of 802.1X authentication.
- "Asset Correlation" on page 335 explains how to create a policy that allows each user to connect using his or her assigned device and no other device.

How Ignition Server Authenticates and Authorizes a User

Ignition Server is a RADIUS server that receives authentication requests from switches and access points (called "authenticators"). When Ignition Server receives a RADIUS request, it:

1. Chooses which access policy to use. By default, this is the RADIUS access policy of the authenticator. If specialized subauthenticators are defined in Ignition Server for the authenticator, then if a subauthenticator matches the RADIUS request, that subauthenticator's access policy is used. If no matching authenticator is found, the global authenticator's access policy is used. See "Matching an incoming request to an authenticator record" on page 90.

Once the correct Ignition Server access policy is found, the remaining decisions are dictated by its rules. Based on the access policy, Ignition Server:

2. Chooses the appropriate user authentication service and user lookup service. Your authentication policy and identity routing policy determine which services are used. See "How Ignition Server Looks Up a User for Authentication and Authorization" on page 221.
3. Authenticates the user and retrieves the user's account details. If any part of this fails, Ignition Server tries additional authentication and user lookup services, if so configured. See "Understanding Authentication Policy" on page 214.
4. Authorizes the user. Ignition Server makes the access decision using the user's account details, other data from the RADIUS request, as well as information about the environment and current time. The decision results in an action of ALLOW, DENY, or CHECK POSTURE. See "How Ignition Server Evaluates a User Authorization Policy" on page 229.
5. If the action is CHECK POSTURE, then the user is allowed, denied, or put on a remediation VLAN based on the results of the posture check as defined in your Ignition Server posture policy. See "How Ignition Server Checks Client Posture" on page 281.

Introduction to Policy Management

An Ignition Server access policy consists of an authentication policy, an identity routing policy (user lookup policy), a user authorization policy, and other optional policies. A given access policy can support many different authentication types and many user directories, and you may include authorization rules that cover many types of users and many locations.

In Ignition Server, an access policy is applied to one or many authenticators. Each authenticator can have an access policy for RADIUS, an access policy for TACACS+, and an access policy for MAC auth. For the global authenticator, you also choose an access policy. In addition, you can specify the use of different access policies based on any attribute in the incoming RADIUS request. This feature, known as the subauthenticator feature, allows you to treat one switch as a number of virtual switches in order to apply the correct policy (user lookup policy, authentication protocol requirement, authorization rule set, and provisioning rule set) to each virtual switch.

What happened to service categories?

If you used version 4.x or earlier of Ignition Server, you will recall that you assigned access policies to authenticators by means of *service categories*. In Ignition Server 4.x and earlier, access policies were nameless; the access policy was just the policy inside the service category. Each authenticator got its policy by being placed in a service category.

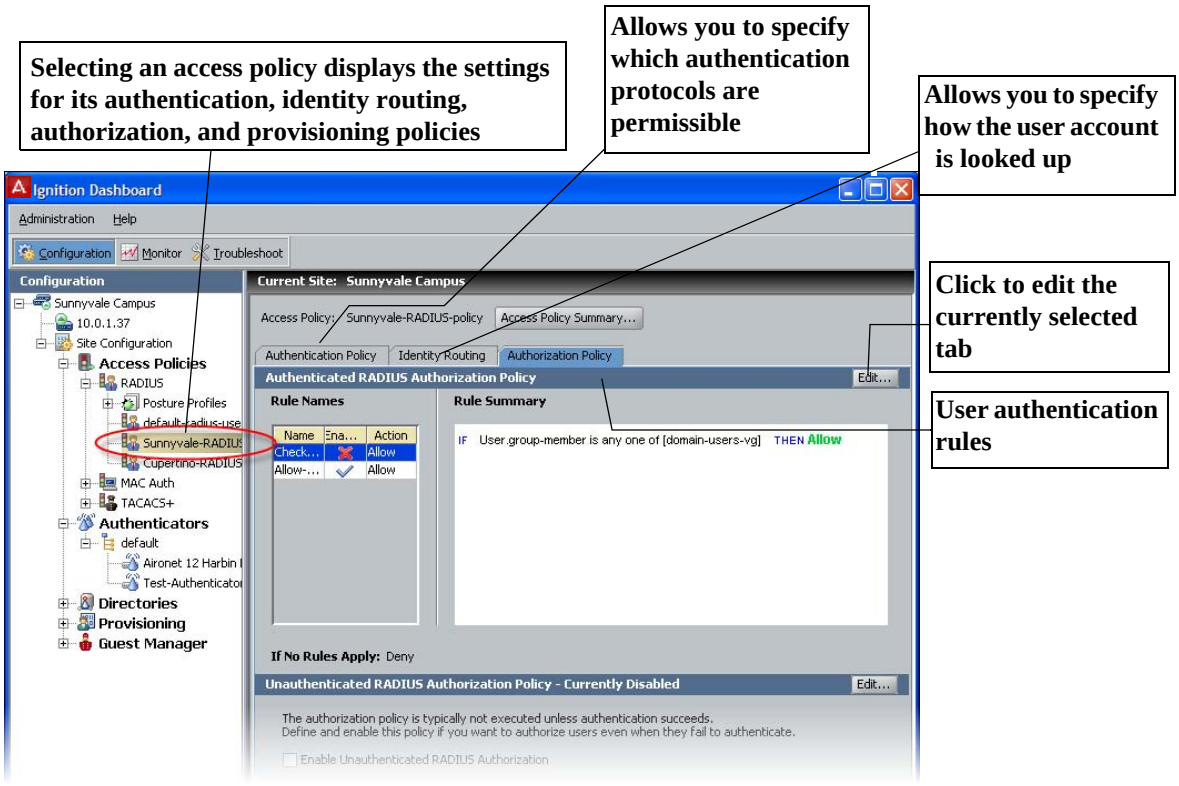
In 5.0 and later, each access policy has a name. Service categories no longer exist. In each authenticator, you designate a RADIUS access policy by name, a TACACS+ access policy by name, and so on.

During an upgrade from 4.x to 5.0.x, the RADIUS policies of a service category are applied directly, as RADIUS access policies, to each authenticator in the service category, and each policy is given the name of the service category that used to contain it. MAC authentication policies are treated in the same way, but their names are given the suffix, "-device."

The Access Policy Panel

In Ignition Server Dashboard, the **Access Policy panel** lets you view and edit your access policies. To open the panel, expand the Access Policy node in the Configuration hierarchy tree and click the name of your access policy.

Figure 52 Access Policy Panel



Understanding Authentication Policy

An authentication policy determines how Ignition Server verifies the identity of a user. Each access policy has an authentication policy. Enforcement of the authentication policy is the first step in Ignition Server's handling of a user.

Ignition Server separates authentication policy definition into two components:

- The **authentication protocol policy** contains the *tunnel protocol policy* and the *credential validation policy*. The *tunnel protocol policy* specifies which sort of encryption is used to secure communications among the client (for example, a user's laptop), the Ignition Server, and the other players in the authentication transaction. We can express this more precisely using 802.1X terminology: The tunnel protocol policy specifies the type of encryption that secures the RADIUS communications among the 802.1X supplicant, the RADIUS server (the Ignition Server), the authenticator (the network switch or access point), and the directory service. "Supplicant" is the industry-adopted name for the software tool on the user's laptop that requests the network connection and prompts the user to enter his or her password or other credentials. The *credential validation policy* specifies which protocol is used to authenticate the user's

password or other proof of identity. Since the authentication is typically relayed to a directory service, the *credential validation policy* you choose must be compatible with the directory service that stores your user. See "Supported Authentication Types" on page 215.

- The **identity routing policy** specifies where Ignition Server can find user records and how it should handle user lookup failures. Identity routing policies are described starting on "Understanding Identity Routing Policy" on page 220.

If authentication succeeds, then Ignition Server executes the user *authorization* policy (see "User Authorization Policy" on page 227).

One Policy Allows Many Authentication Protocols

Your authentication policy typically makes a number of authentication protocols available to the user for logging in. For example, if your user authorization policy is set to allow both PEAP/EAP-MSCHAPv2 and NONE/PAP authentication types, and a user attempts to log in from a laptop using the Microsoft Windows supplicant, then PEAP/EAP-MSCHAPv2 authentication is used. If the same user later attempts to log in from a Linux workstation using a Meetinghouse supplicant, then NONE/PAP authentication is used. Ignition Server chooses the authentication type based on the inner authentication protocol (in this example, MSCHAPv2 and later PAP) of the incoming request.

Supported Authentication Types

The authentication protocols (tunnel and credential validation protocols) available to you depend on the type of user store you use. The tables that follow show which protocols are available for each store type. In each authentication type name, the name before the forward-slash indicates the outer tunnel protocol, and the name after the forward-slash indicates the credential validation protocol. A blank cell indicates an unsupported combination.

Table 5 Non-EAP Authentication Protocol and User Data Store Compatibility

Data Store Type	NONE / PAP	NONE / CHAP	NONE / MSCHAPv2	TTLS / PAP
Ignition Server internal store	Yes	Yes	Yes	Yes
Microsoft Active Directory	Yes		Yes	Yes
Sun Java System Directory Server (SunONE LDAP)	Yes		Yes*	Yes
Novell eDirectory	Yes		Yes*	Yes

Table 5 Non-EAP Authentication Protocol and User Data Store Compatibility

Data Store Type	NONE / PAP	NONE / CHAP	NONE / MSCHAPv2	TTLS / PAP
Oracle OID	Yes		Yes*	Yes
Generic LDAP	Yes		Yes*	Yes
Kerberos Authentication	Yes			Yes
RSA Authentication Server	Yes			Yes
Proxy Directory service	Yes**			

* To perform MSCHAPv2 authentication against an LDAP user store, each LDAP user record must contain an NT hash of the user's password. For instructions, see "Setting up MSCHAPv2 Authentication on LDAP" on page 156.

**The Ignition server acts as a proxy and forwards requests to the remote proxy server for authentication.

Table 6 EAP Authentication Protocol and User Data Store Compatibility

Data Store Type	NONE / EAP-MD5	NONE / EAP-GTC	NONE / EAP-TLS	NONE / EAP-MSCHAPv2	PEAP/ EAP-GTC	PEAP / EAP-MSCHAPv2	PEAP / EAP-TLS
Ignition server internal store	Yes	*	Yes	Yes	*	Yes	Yes
Microsoft Active Directory		*	Yes	Yes	*	Yes	Yes
Sun Java System Directory Server (SunONE LDAP)		*		Yes**	*	Yes**	
Novell eDirectory		*		Yes**	*	Yes**	
Oracle OID		*		Yes**	*	Yes**	
Generic LDAP		*		Yes**	*	Yes**	
RSA Authentication Server		Yes			Yes		

Table 6 EAP Authentication Protocol and User Data Store Compatibility

Data Store Type	NONE / EAP-MD5	NONE / EAP-GTC	NONE / EAP-TLS	NONE / EAP-MSCHAPv2	PEAP/ EAP-GTC	PEAP / EAP-MSCHAPv2	PEAP / EAP-TLS

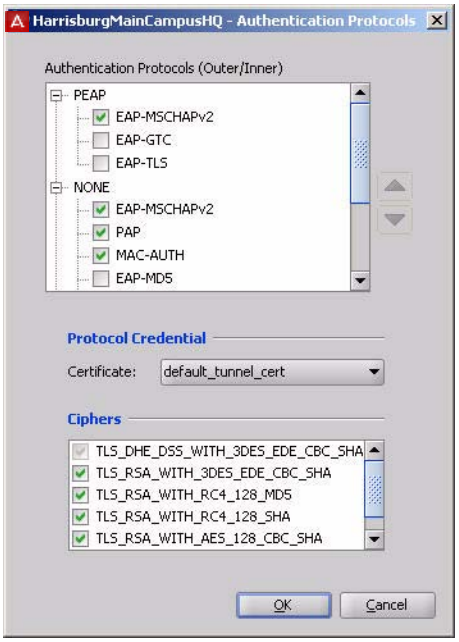
* When authenticating over EAP-GTC or PEAP/EAP-GTC (typically RSA SecurID authentication), you may optionally configure Ignition Server to split user look-up from authentication. Under the split lookup scenario, Ignition Server performs the user look-up against AD, LDAP, or the internal store to retrieve user attributes, while the RSA Authentication Server handles authentication.

** To perform PEAP / EAP-MSCHAPv2 authentication against an LDAP user store, each LDAP user record must contain an NT hash of the user's password. For instructions, see "Setting up MSCHAPv2 Authentication on LDAP" on page 156.

Authentication Policy Window

Use the **Authentication Policy** window to establish an allowed set of user authentication protocols for the access policy.

Figure 53 Authentication Policy window



The authentication policy consists of a set of authentication protocols used to validate users' credentials, each paired with the outer tunnel protocol used to secure the credentials during transmission (set in the **Authentication Protocols** section); a credential validation certificate (set in the **Protocol Credential** section); and a set of allowed ciphers. The **Authentication Protocols** section is presented as a tree with the outer tunnel types listed as parent nodes and the authentication protocols listed as child nodes. To use an authentication type, click its outer tunnel type to expand its list, and tick the checkbox for the desired authentication type.

EAP-TLS Authentication

Ignition Server supports EAP-TLS and PEAP/EAP-TLS authentication in two cases:

1. user authentication
2. device authentication using Windows machine authentication. (See "Windows Machine Authentication" on page 295.)

When you choose EAP-TLS or PEAP/EAP-TLS authentication, the authenticating user or device passes a digital certificate to prove its identity. Ignition Server parses and evaluates the certificate as follows.

For user authentication:

- In the user-submitted certificate, Ignition Server looks in the *Subject Alternative Name* field, reads the *Other Name: Principal Name* attribute and compares its value with the user name from the directory service. If no match is found, the search continues as follows:
- Ignition Server looks in the *Subject* field of the user-submitted certificate, reads the first *CN* attribute it finds there, and compares its value with the user name from the directory service. If no match is found, the authentication fails.

For machine authentication:

- In the device-submitted certificate, Ignition Server looks in the *Subject Alternative Name* field, reads the *DNS Name* attribute and compares its value with the Computer name from Active Directory. If no match is found, the search continues as follows:
- Ignition Server looks in the *Subject* field, reads the first *CN* attribute it finds there, and compares its value with the Computer name from Active Directory. If no match is found, the authentication fails.

Note that Ignition Server does not support the binary comparison of the user-submitted or device-submitted certificate with a copy of the certificate stored in the directory.

Factors That Limit Your Choice of a Protocol Credential Certificate

Note the following limitations when choosing the protocol credential certificate for your Ignition Server authentication policy:

- Certificates are configured using the Certificate Manager (see "Certificates Tab" on page 74), and can be shared across access policies.
- A potential defect in Microsoft Windows XP prevents the use of DSA-signed certificates for PEAP communication with Windows XP

supplicants. Avaya has verified that this failure is not an Ignition Server -specific defect. If a Windows XP client tries to establish a PEAP tunnel with Ignition Server using a DSA-signed certificate, the connection attempt fails. If your installation includes Windows clients, use only an RSA-signed certificate. (If your installation happens to support only non-Windows clients, you can use a DSA-signed or RSA-signed certificate.)

Creating an Authentication Policy

Set the authentication policy of your access policy as follows:

1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS.
2. Click the name of your access policy and click the **Authentication Policy** tab. Click the **Edit** button.
3. In the **Edit Authentication Policy** window, the **Authentication Protocols** section lets you establish the set of inner authentication protocols and outer tunnel types that your access policy supports.

   In the Authentication Protocols section, choose each authentication type as follows:

   - Find the outer tunnel type that corresponds to your authentication protocol. The choices for *outer tunnel protocol* are: **NONE** (no outer tunnel is used; user credentials travel in the clear); **PEAP**; and **TTLS**. Both PEAP and TTLS require an Ignition Server-side certificate to encrypt user credentials.
   - Click the +/- toggle to expand the list of inner authentication types.
   - Click the checkbox next to each desired authentication type. Choose as many as you want to support. Consult Table 5 and Table 6 to verify that your authentication protocol is compatible with your data store type.

     If you choose EAP-TLS, you must install one or more root certificates on the Ignition Server. See "Installing Protocol Root Certificates" on page 83.
   - Sort the order in which Ignition Server should attempt to use the authentication types by clicking the name of the authentication type and clicking the up/down arrows located to the right of the tree display.
   - If you have additional authentication types that use other outer tunnel types, repeat the steps above for each outer tunnel type.
4. In the **Protocol Credential** section, use the drop down list to choose the certificate that secures the PEAP or TTLS transactions. If the list is empty, import your certificate as explained in "Assigning Protocol Credential Certificates" on page 82.



Important! For tips on choosing a certificate, see the section, “Factors That Limit Your Choice of a Protocol Credential Certificate” on page 218.

5. In the **Ciphers** section, select the cipher suites to be used for encrypting outer tunnel communication. Typically, you should select all the of cipher suites permitted by your company’s security policy. During tunnel protocol negotiation with the client, Ignition Server uses the strongest cipher compatible with the client certificate.

   By default, the first five cipher suites are selected. The PEAP IETF draft standard requires the first entry. (The sixth in the list, “TLS_DH_anon_WITH_AES_128_CBC_SHA”, is not selected and not recommended.)

6. Click **OK**.

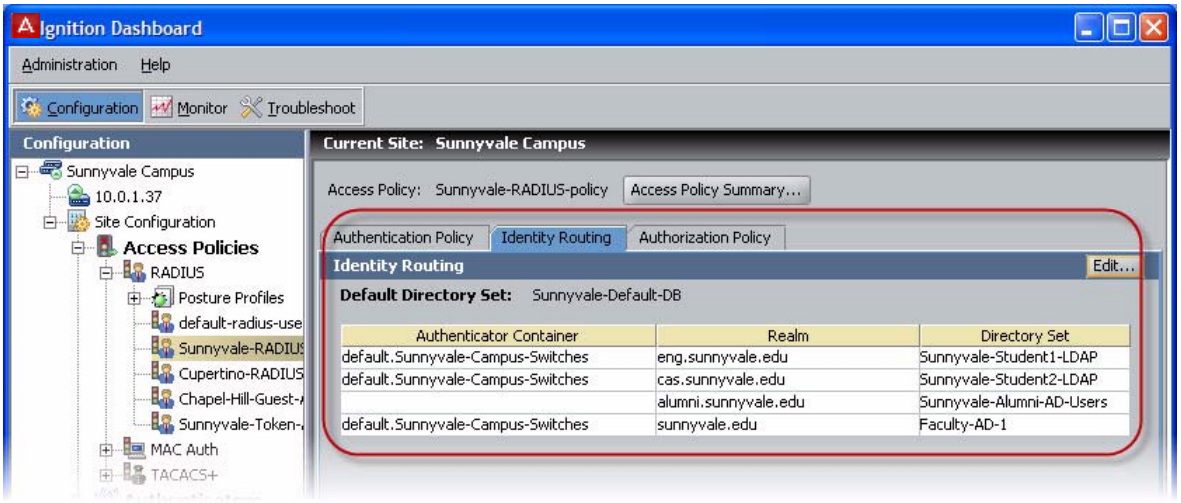
Your authentication policy has been set. Next, make sure your identity routing policy includes directory services of the appropriate type (LDAP, AD, etc.) to support the authentication types you have chosen. Turn to “Understanding Identity Routing Policy”, below.

Understanding Identity Routing Policy

At user login time, the Identity Routing Policy tells Ignition Server which directory set to search for the user account, based on the realm (domain) name passed with the user name, and/or based on which authenticator the user is connecting through.

For example, you can specify that all users with user names like *kadams@avaya.com* or *avaya/jlee* are authenticated against the directory set that contains the corporate Active Directory (AD) while other users without *avaya* in their user name are authenticated against your guest user database. If your site needs to use different user directories for different locations, you must write an authenticator-based lookup policy that specifies, for example, that all users connecting through the wired ports in your Mountain View office are authenticated against a directory set that contains the local AD for the Mountain View office.

Figure 54 An example identity routing policy shown in the Access Policy panel of Dashboard

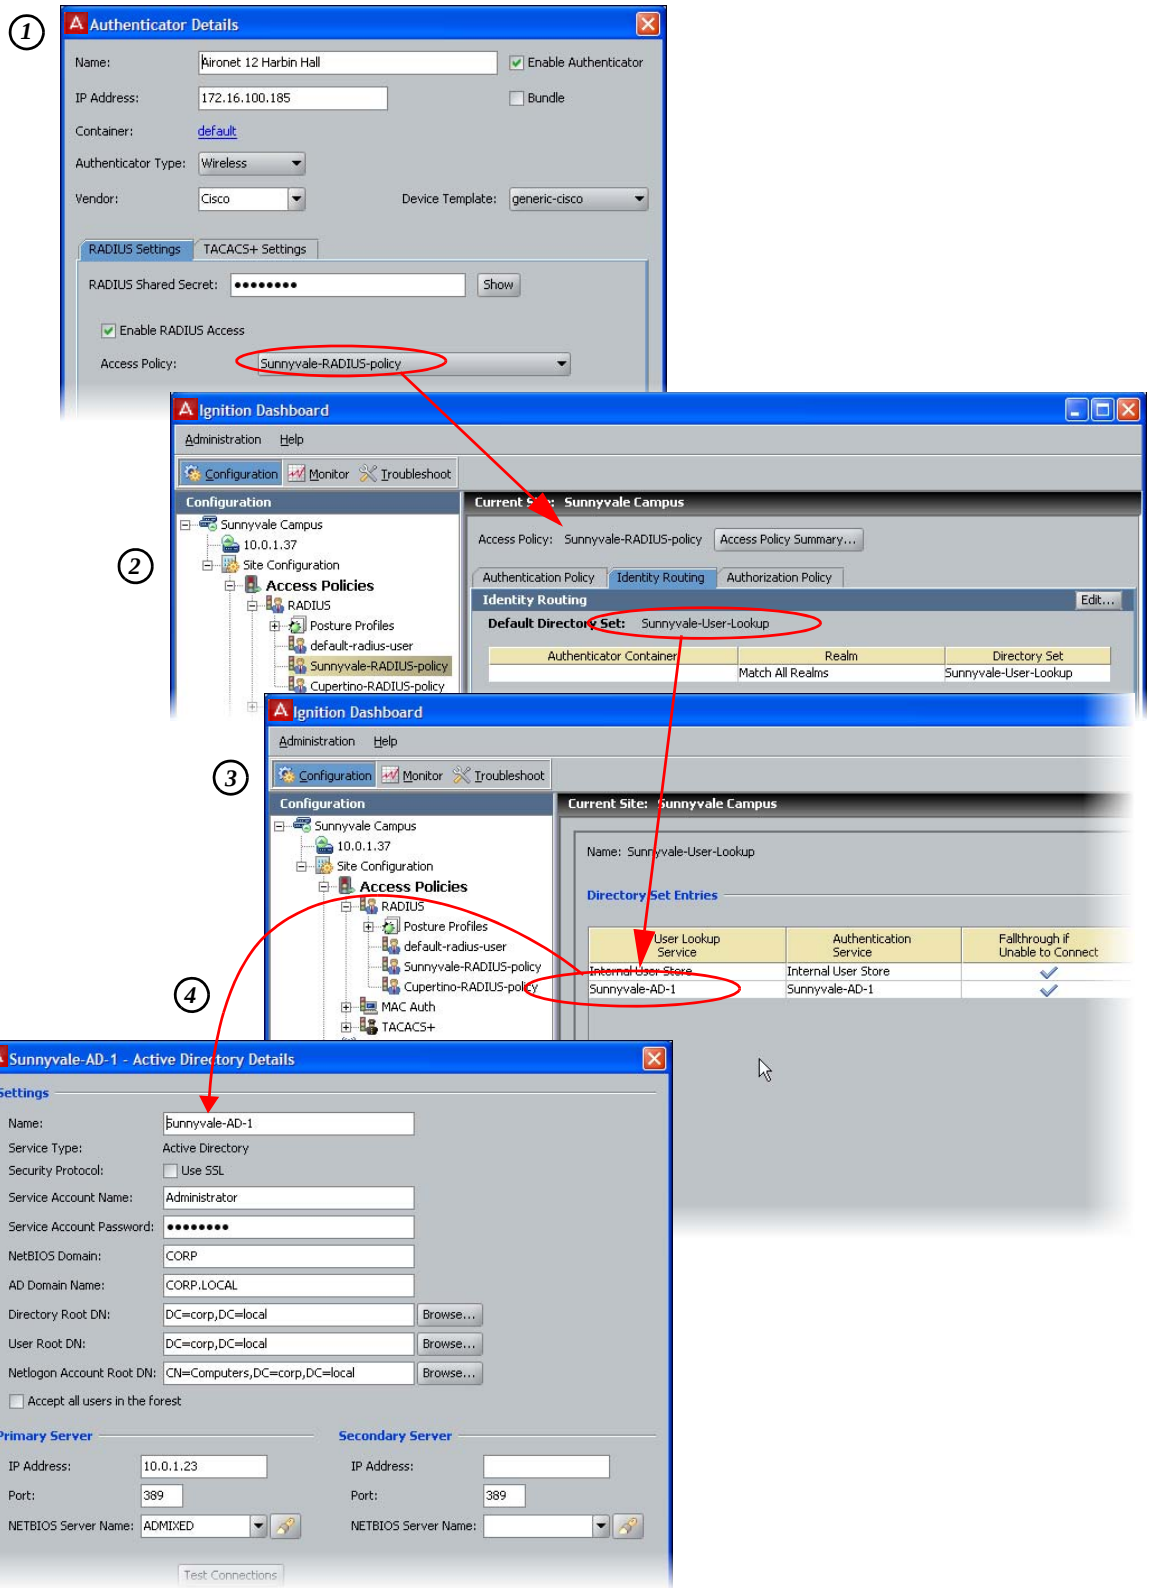


How Ignition Server Looks Up a User for Authentication and Authorization

When handling an authentication request, Ignition Server locates the user account as follows:

- Checks which *authenticator* relayed the authentication request. **(Step 1 in the following illustration)**. See "Matching an incoming request to an authenticator record" on page 90 for an explanation of how Ignition Server finds the right authenticator record.
- Reads the Ignition Server *access policy* of that authenticator and reads its identity routing policy. **(Step 2)**
- Based on the *identity routing policy*, it finds the directory set that corresponds to the network domain (realm) of the user's account and/or to the authenticator that the user is connecting from. The first match is used, and no further realm/authenticator mapping rules are checked. If no match occurs and the policy includes a default directory set, then the default directory set is used. If no match occurs and there is *no* default directory set, the authentication request is rejected. **(Step 3)**
- Searches for the user in the first directory service in the *directory set*. **(Step 4)** If the user is found, Ignition Server attempts to authenticate and authorize the user. If the lookup or authentication attempt fails (failure to connect to the directory service, failure to find the user account, or failure to validate the credentials), then Ignition Server checks the directory set's fall-through rules. If the rules call for it, Ignition Server searches the next directory service on the fall-through list. (You can test your directory set as explained in "Checking an Authentication Request" on page 173.)

Figure 55 The User Look-Up Path in Ignition Server



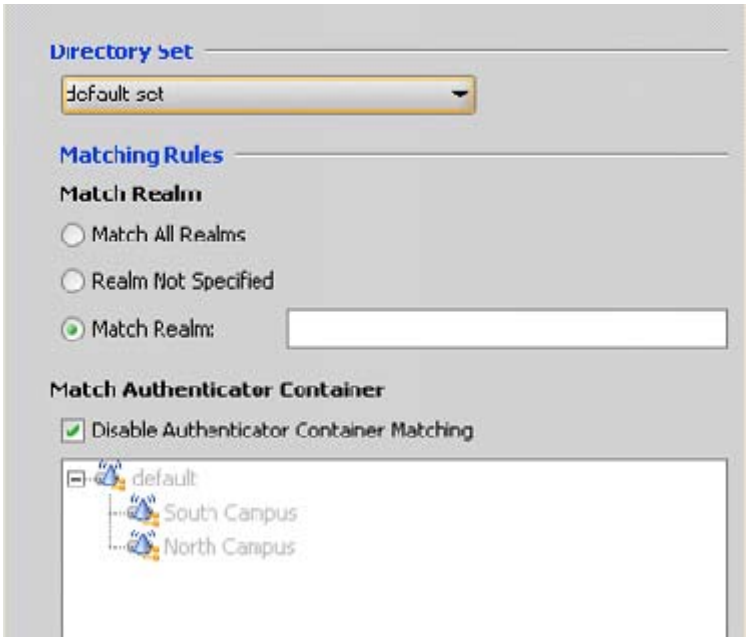
Creating an Identity Routing Policy

Your identity routing policy consists of a set of realm/authenticator mapping rules, each of which maps to a directory set, and optionally, a default directory set to be used if no rule matches. Follow the steps below to create your identity routing policy:

1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click on the name of your access policy.
2. Click the **Identity Routing** tab and click the **Edit** button in that tab.
3. If you want to Enable a default directory set, tick the **Enable Default Directory Set** checkbox and select the name of your default directory set in the drop down list just below the checkbox. The default directory set is used for any login attempt that fails to match any of the realm/authenticator mapping rules in the list. If you like, you can specify a default directory set only and skip the realm/authenticator mapping rules altogether.
4. Set up your realm/authenticator mapping rules in the **Realm-Directory Set Mapping** table. To begin adding a mapping rule, click the **New** button below the table.

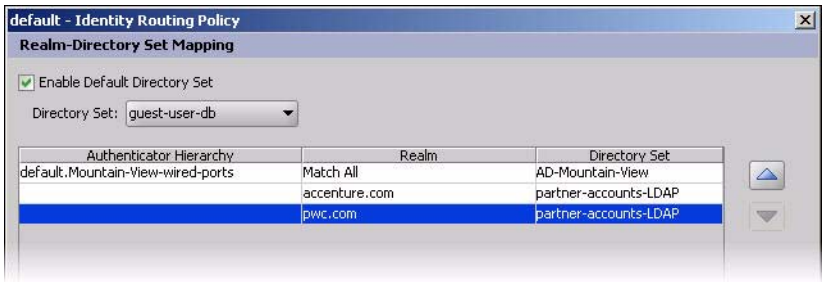
   This launches the **Realm-Directory Set Map** window, which lets you specify a set of conditions under which a particular directory set is used.

Figure 56 A realm/authenticator mapping rule that checks both realm and authenticator



5. In the **Directory Set** drop-down list, choose the directory set. When a login attempt matches this rule, this directory set is searched for the user account, and no other directory sets are checked.
6. In the **Match Realm** section, set up your realm-matching rule:
   - To create a rule that requires a realm-name match, tick the **Match Realm** checkbox and type the realm name in the neighboring field. The value you specify here is compared with the realm portion of the user name in the authentication request. For example, if your users log in with names like *jlee@avaya.com*, then you specify *avaya.com* here. If they log in with names like *avaya/jlee*, specify *avaya* here. Note, the comparison is case sensitive. For more information see "Additional Notes on Realm-Matching Rules" on page 225.
   - To create a rule that matches realm-less user names (for example, a user name of *jlee*), tick the **Realm Not Specified** checkbox.
   - To specify that no realm-matching is required, tick the **Match All Realms** checkbox. The rule must match all user names. This is useful if you want to perform authenticator-matching but no realm-matching.
7. In the **Match Authenticator Container** section, set up your authenticator-matching rule. For all users connecting over a particular authenticator or set of authenticators, the rule specifies which directory set is used.
   - To disable authenticator-matching, tick the **Disable Authenticator Container Matching** checkbox.
   - To use authenticator-matching, make sure the checkbox is *not* ticked, and, in the tree view, click to select the node in the hierarchy that represents the desired authenticator(s). For information on labeling your authenticator with an authenticator container, see "Authenticator Hierarchy and Containers" on page 91.
8. Click **OK**.
9. In the **Identity Routing Policy** window, add additional realm/authenticator mapping rules by clicking the **New** button again and repeating the procedure described above.
10. After you have added your mapping rules, sort them. The order is important because Ignition Server uses the first match, and no further realm/authenticator mapping rules are checked. To sort, click a mapping rule and click the up/down arrows located to the right of the table.
11. Click **OK** to dismiss the Identity Routing Policy window.

Figure 57 Example Identity Routing Policy with Container Added



Additional Notes on Realm-Matching Rules

The realm name is stripped from the user name in the authentication request in the following manner:

- For addresses specified as *myRealm/myName* or *myRealm\myName*, the realm name is the part that precedes the slash or backslash.
- For addresses specified as *myName@myRealm*, the realm name is the part that follows the "@" sign.
- If both forms are present, then the realm is the part that precedes the slash. For example, in the user name *myRealm/myName@myCompany.com*, Ignition Server strips the realm name *myRealm*.

User Authorization Policy

After Avaya Identity Engines Ignition Server authenticates a user, it checks your user authorization policy to determine whether the user should be granted access to the requested network resource. This chapter describes how to create and maintain user authorization policies.

Optionally, your authorization policy can invoke a session provisioning policy, to send provisioning values that set more detailed network rights such as VLAN assignments and administrator rights on network equipment. For more information, see “Provisioning Policy” on page 253.

Ignition Server also lets you define authorization policies for devices. See Chapter , “MAC Authentication”.

Contents

- “Introduction to User Authorization Policies” on page 227
- “Structure of User Authorization Policies” on page 228
- “Creating a RADIUS User Authorization Policy” on page 242
- “Enabling or Disabling Rules Within a Policy” on page 247
- “Copying An Authorization Rule” on page 248
- “Creating an Authentication-Only Policy” on page 249
- “Using a Device Attribute in a Rule” on page 250

Introduction to User Authorization Policies

A user authorization policy is a rule sequence you create that determines whether a user or a device is allowed to access a requested network resource, and what session provisioning, if any, is applied to the network session. To make the access decision, Ignition Server can evaluate attributes of the user, his client machine, the switch over which he is connecting, and/or the context (time, location, etc.) of the access request. Each access policy in Ignition Server contains one user authorization policy, and you can view its summary in the Authorization Policy tab in the Access Policy panel. Evaluation of the user authorization policy happens immediately after Ignition Server authenticates the user.

Tip: What if I don't want to use authorization rules? Ignition Server always performs both authentication *and* authorization before it grants a user access. In some installations, you may decide that authentication alone (checking the user's credentials) is sufficient to grant the user access. If this is the case, you must use a catch-all, authentication-only rule. See "Creating an Authentication-Only Policy" on page 249.

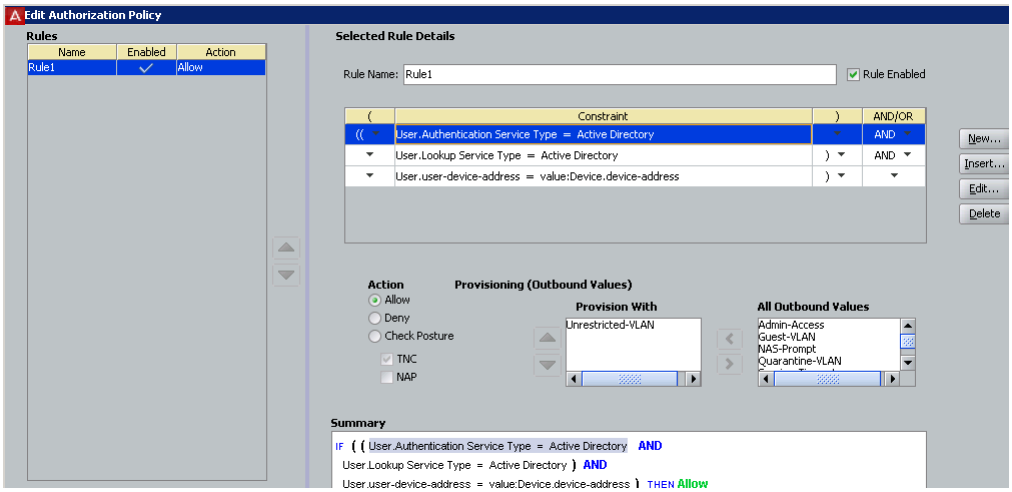
Structure of User Authorization Policies

This section provides an introduction and reference to user authorization policies and the Dashboard windows you use to edit them. This section explains:

- "Elements of a User Authorization Policy" on page 228
- "How Ignition Server Evaluates a User Authorization Policy" on page 229
- "Reading the Rule Summary" on page 232
- "Attributes Used in Rule Constraints" on page 232
- "Using Time and Date in a Rule" on page 238
- "Conjunctions Used to Assemble Constraints Into a Rule" on page 239
- "Comparison Operators for Rules" on page 240

Elements of a User Authorization Policy

Figure 58 Inspecting a rule in the Edit Authorization Policy window



A user authorization policy is a set of rules arranged in a sequence. The elements that define each *rule* are as follows:

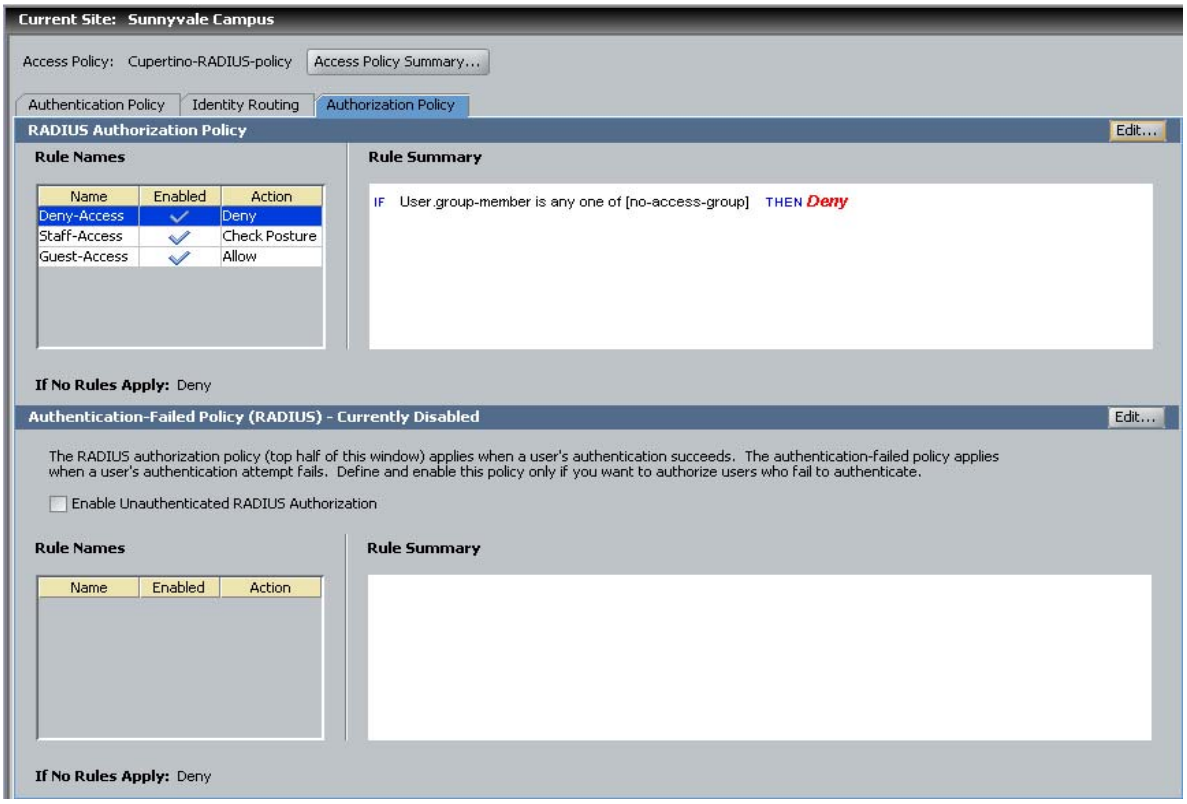
- Each rule contains one or more constraints logically ANDed and ORed together. In the Edit Authorization Policy window, these appear in the **Constraint** table.
- Each constraint evaluates an attribute (a piece of data describing the user, his machine, the request context, or the authenticator; see "Attributes Used in Rule Constraints" on page 232).
- Each rule has an action to ALLOW, DENY or CHECK POSTURE on the access request.
- Each rule can have session provisioning instructions associated with it. A provisioning instruction is an Ignition Server outbound value or values that Ignition Server sends to the authenticator with the *access-accept* or *access-reject* message. See "Provisioning Policy" on page 253 for details on provisioning instructions.

Note that *rules* can be copied, edited, renamed, and deleted.

How Ignition Server Evaluates a User Authorization Policy

Ignition Server makes a user authorization decision as explained below, using a *first-match-wins* logic. The figure below shows an example policy with three rules in the rule sequence.

Figure 59 An Example User Authorization Policy in the Access Policy Panel



In the above example window, Ignition Server applies three rules — *Deny-Access, Staff-Access,* and *Guest-Access* — in that order, to evaluate each access request in the access policy. Each rule in the rule sequence contains a constraint or a set of constraints that are ANDed and ORed together. The rule evaluates to TRUE, FALSE, or INDETERMINATE.

Each rule is associated with a corresponding action: ALLOW, DENY, or CHECK POSTURE. Access is granted when an ALLOW rule evaluates to true, or when a CHECK POSTURE rule evaluates to true and the posture checking results in an ALLOW.

Ignition Server evaluates the DENY rules first and then evaluates the remaining rules in the order you specify. Rules are evaluated until a rule evaluates to true. At that point Ignition Server performs the rule's action, and no other rules are evaluated. If no rule is triggered, Ignition Server performs the default, "If-no-rules-apply" ACTION (typically a DENY) that you have specified in your policy.

Ignition Server's Rule Evaluation Routine

Ignition Server evaluates a user authorization policy in the following manner:

First, Ignition Server evaluates all the DENY rules:

- If a DENY rule evaluates to FALSE, Ignition Server takes no action. It continues to the next rule.
- If a DENY rule evaluates to TRUE or is INDETERMINATE (see "How a rule can evaluate to indeterminate" on page 231), Ignition Server carries out the DENY (and denies authorization to the user). No further rules are evaluated.

Second, Ignition Server evaluates all the ALLOW and CHECK POSTURE rules, in the order you have specified:

- If the rule evaluates to FALSE or is INDETERMINATE (see "How a rule can evaluate to indeterminate" on page 231), Ignition Server takes no action. It continues to the next rule.
- If the rule evaluates to TRUE, Ignition Server evaluates the ACTION clause:
  - If the ACTION is ALLOW, Ignition Server grants the user access and returns any provisioning values associated with the rule. No further rules are evaluated.
  - If the ACTION is CHECK POSTURE, Ignition Server applies the posture checking prescribed in the Posture Profile and performs the action specified in the applicable posture tab (*Compliant*, *Non-Compliant*, or *No Posture*). No further rules are evaluated. (For more on posture checking, see "How Ignition Server Checks Client Posture" on page 281.)

Third, and finally, if Ignition Server reaches the end of the rule sequence and no rule has been triggered, it performs that ACTION you specified in the **If No Rules Apply** field of the Edit Authorization Policy window. The ACTION can be an ALLOW (including the sending of provisioning values, if desired) or a DENY. The default behavior is DENY.

When Ignition Server authorizes or rejects a user, it logs the action. Log entries can be viewed in the **Access** tab of the Log Viewer tab.

How a rule can evaluate to indeterminate



Important! When Ignition Server evaluates a constraint at runtime, it is possible that the constraint is impossible to evaluate (for example, because the attempt to retrieve a user attribute fails). In this case, the constraint evaluation is considered *indeterminate*. If the logic of the entire rule does not require an evaluation of the failed constraint, then the rule can still return TRUE (for example, because the failed constraint was ORed with a constraint that evaluated to TRUE). If the logic of the entire rule requires an evaluation of the failed constraint, then the rule returns INDETERMINATE. The handling of INDETERMINATE rules is explained in the preceding section.

Reading the Rule Summary

The Access Policy panel's Authorization Policy tab contains a **Rule Summary** display consisting of an "IF" clause that summarizes the rule's constraints, and a "THEN" clause that summarizes the rule's action and the provisioning values it sends.

The form of the "IF" clause depends on the type of test value you are using in the rule:

- "IF" clauses that use fixed test values take the form shown in this example: "`IF User.group-member = [MV-Visitors]`". In this example, `User` is the attribute type, `group-member` is the attribute name, and "=" (equals) is the comparison operator. The test value or values appear inside the brackets as a comma-separated list.
- "IF" clauses that use dynamic test values take the form shown in this example: "`IF User.user-id = [value:Inbound.User Name (Inner Tunnel)]`". In this example, `User` is the type of attribute on the left-hand side of the comparison, `user-id` is the name of the attribute on the left-hand side, and `equals` is the comparison operator. On the right-hand side of the equation, the source of the test value is shown inside the brackets. In this example, `value` indicates that this is a dynamically retrieved value, `Inbound` is the type of attribute on the right-hand side of the comparison, `User Name (Inner Tunnel)` is the name of the attribute from which the value is retrieved.

The "THEN" clause consists of the action (ALLOW, DENY, or CHECK POSTURE) and, optionally, a list of the outbound values that are to be sent when the ALLOW action is taken.

Attributes Used in Rule Constraints

A constraint in a rule can evaluate attributes of the following types:

- **user attributes**: data that describes the user, his or her organization, or his or her group affiliations. See "User Attributes" below.
- **inbound attributes**: values passed by the authenticator in the form of RADIUS attributes or VSAs. These typically describe the context, time, or originating user/device of the access request. See "Inbound Attributes" on page 234.
- **authenticator attributes**: Ignition Server -stored data that describes the switch or access point, such as the name of the switch manufacturer, its location in the Ignition Server authenticator hierarchy, or the name of the Ignition Server access policy it belongs to. See "Authenticator Attributes" on page 237.

- **device attributes**: data that describes the connecting client device such as a user's laptop or a printer. See "Device Attributes" on page 236.

Each attribute allows the use of comparison operators appropriate to its content. For example, user names are strings, for which you can use the comparison operators *Equal To*, *Not Equal To*, *Starts With*, *Ends With*, or *Contains*.

User Attributes

When you choose **User** from the Attribute Category dropdown list in the Constraint Details window, the list displays *user attributes*. User attributes describe the user, his or her organization, and his or her group affiliations. The default attributes are virtual user attributes that map to fields in the Ignition Server internal user record. When you create a virtual attribute, you add mappings that point to fields in your LDAP, AD, or other store. To create, edit, or inspect a virtual attribute, use the Virtual User Attribute window as shown in "User Virtual Attributes" on page 203.

By default, the Constraint Details window offers the set of Ignition Server - defined *user attributes* listed below. In the list that follows, we explain only the default mappings to Ignition Server *internal user* records. For information on virtual attributes mapped to your AD or LDAP user records, contact your AD or LDAP administrator.

- **account-locked**: boolean indicating whether the internal user record has been locked
- **email-address**: the E-Mail Address recorded in the internal user record
- **enable-max-retries**: boolean indicating whether the Enable Max Retries checkbox is checked in the internal user record
- **enable-password-expiration**: boolean indicating whether the Enable Password Expire checkbox is checked in the internal user record
- **enable-start-time**: boolean indicating whether the Enable Start Time checkbox is checked in the internal user record
- **first-name**: the First Name recorded in the internal user record
- **group-member**: choose this attribute if you want to write a constraint that checks if the user is or is not a member of a user group. If the user account resides in the internal data store, group membership is determined in the Member of Groups tab in the Edit User window. If the user account resides in an AD, LDAP, or Novell store, group membership is determined as explained in "Mappable Group Types for Ignition Server Virtual Groups" on page 199.
- **last-name**: the Last Name recorded in the internal user record

- **max-retries**: the integer Max Retries threshold as set in the internal user record
- **network-usage**: the Network Usage value (a string) recorded in the internal user record
- **office-location**: the Office Location value recorded in the internal user record
- **password-expiration**: the password expiration date and time, as recorded in the internal user record.
- **role**: the Org. Role recorded in the internal user record
- **start-time**: the user account start date and time, as recorded in the internal user record
- **title**: the Title recorded in the internal user record
- **user-id**: the User Name recorded in the internal user record. (Note: If you want to evaluate the *user-submitted name* from the authentication request, see the section, “user name attributes” in “Inbound Attributes” on page 234.

Mapping Virtual Attributes to User Attributes

Most user attributes are available in the user record obtained from either the directory server or the internal data store as part of authentication processing. This user attribute data must be mapped to corresponding virtual attributes before it can be used inside authorization policy rules. (This mapping is automatic for user records obtained from the internal data store.) Similarly, user membership in groups must be mapped to virtual groups before they can be used as part of policy evaluation. (See “User Virtual Attributes” on page 203.)

Inbound Attributes

Inbound attributes describe the context, date, and time of the access request, the context and name of the user, and can include any data value sent by the authenticator. Many of these attributes are RADIUS attributes and VSAs sent from the authenticator; others are based on information from the Ignition Server. You can expose any incoming RADIUS attribute or VSA as an authenticator attribute, as explained in the last bullet point below.

The inbound attributes are:

- **Date**, **Date and Time**, and **Time**: These attributes let you use the Ignition Server date and time in a rule. See “Using Time and Date in a Rule” on page 238.

- **Authentication Service Type and Authentication Service Name**: The type (AD, LDAP, internal store, Kerberos, Token Service, or Radius Proxy Service) and name of the directory or authentication server that authenticated the user. Each service has a name as set up in the Directory Services panel. See “Directory Services” on page 134. **Hint:** When your Ignition Server is performing authentications, you can view the directory service name for each authenticated user in the Access log channel.
- **Lookup Service Type and Lookup Service Name**: The type (AD, LDAP, internal store, or none) and name of the directory server where the authenticating user’s account was found. In most cases the lookup service and the authentication service are one and the same, but if you split lookup from authentication, such as with a SecurID authentication and an AD user lookup, then they are not the same. Each service has a name as set up in the Directory Services panel. See “Directory Services” on page 134.
- *The* **user name attributes**:
  - **Inbound-User-Name** holds the value of the RADIUS User-Name attribute from the incoming RADIUS request. You can define a custom mapping for this attribute in the Inbound Attributes panel of Dashboard, in which case the value comes from the RADIUS attribute or VSA you specify.
  - **User Name (Inner Tunnel)** is the name the user submitted for authentication.
  - **User Name (Outer Tunnel)** is the name the user presented to establish the secure tunnel for authentication.

  Typically, all three attributes contain the same value, but if the user is authenticating over a tunneled authentication protocol, then in many cases the **Inbound-User-Name** and the **User Name (Outer Tunnel)** match, and the **User Name (Inner Tunnel)** is different.
- **Realm (Inner Tunnel)**, and **Realm (Outer Tunnel)** contain the realm or domain designation of the user. The **Realm (Inner Tunnel)** is the domain the user submitted to authenticate. The **Realm (Outer Tunnel)** is the domain the user submitted to create the tunnel. These values typically match, but in a tunnelled authentication they might not.
- **Inner Tunnel Type** and **Outer Tunnel Type**: The **Inner Tunnel Type** is the protocol use to carry the authentication credentials. The **Outer Tunnel Type** is the type of tunnel used to encrypt the authentication transaction. For example, if the user is authenticating over PEAP/EAP-MSCHAPv2, the **Inner Tunnel Type** attribute reads “EAP-MSCHAPv2” and the **Outer Tunnel Type** attribute reads “PEAP”.

- **Secure Tunnel** is a boolean indicating whether a tunnel was used to encrypt the authentication transaction.
- *The inbound RADIUS attributes and VSAs*: These attributes let you evaluate the contents of any attribute sent by the authenticator. These typically have names that start with "**Inbound-**", but the ones you create can have any name you like. The default list includes a majority of the most popular RADIUS attributes. You can view the list of available RADIUS- and VSA-sourced inbound attributes (and their mappings) in the Inbound Attributes panel, as explained in "Finding an Inbound Attribute" on page 266. If the attribute you want to evaluate does not appear in the list, set up a new inbound attribute as explained in the section, "Preparing an Inbound Attribute for Use in an Authorization Rule" on page 265.

Device Attributes

Device attributes describe the end-client hardware that is attempting to connect to the network. For example, this might be a user's laptop, a printer, or a handheld device. The attributes are:

- **device-address**: MAC address of the device
- **device-group-member**: indicates the device's group membership, as recorded in its Ignition Server device record.
- **device-name**: name of the device as stored in its Ignition Server device record
- **device-type**: the Type label as stored in its Ignition Server device record. Typically indicates what sort of device it is, such as a printer or handheld device.
- **device-sub-type**: indicates more details about the device type. For example, if the device Type is "mobile", the Sub Type indicates which type of mobile it is, such as an iphone, blackberry, or android phone.
- **device-OS-type**: indicates the type of operating system on the device.
- **device-OS-version**: indicates the version of operating system on the device.
- **device-user-name**: the name of the user of the device.
- **device-vlan**: the VLAN designation stored in the connecting device's Ignition Server device record. Note, in the case of a device that is already on a VLAN, this might NOT be the current VLAN to which the device is connected.
- **exists-in-embedded-store**: a boolean indicating whether this device matches an Ignition Server device record. Keep in mind that the Ignition

Server device record may contain a wildcarded MAC address such as “00:b7*”. Any device that matches the wildcarded address triggers an **exists-in-internal-store** value of TRUE.

- **is-assigned-to-embedded-user**: a boolean indicating whether the connecting device has been assigned to the authenticating user. (In other words, if Ignition Server contains a device record that matches the connecting device's MAC address, and if that device record has been assigned to the connecting user, then this attribute evaluates to TRUE.) As with other parameters, wildcard matches evaluate to TRUE.
- **learned-via-AD-login**: a boolean indicating whether this device has a current session that it obtained by authenticating to Ignition Server via Windows Machine Authentication. In this case, no device record is used. Instead, this attribute evaluates to TRUE if the device has a current Windows Machine Authentication session on the Ignition Server. See the section, “Learned Devices Tab” on page 447 for details.
- **source**: the Source label of the device, as stored in its Ignition Server device record. It indicates where the device record originated.
- *The device virtual attributes*: These are attributes you define that let you evaluate the contents of any field in the device record. To create one, follow the instructions in “Adding Virtual Attributes for Devices” on page 208.

See also “Using a Device Attribute in a Rule” on page 250.

Authenticator Attributes

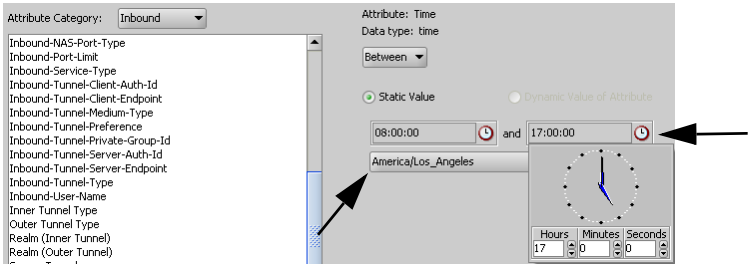
Authenticator attributes contain data that describes the authenticator (usually, a switch or access point) through which the user or device is attempting to start a network session. These attribute values are retrieved from the authenticator record in Ignition Server. See “Creating an Authenticator” on page 96 for details.

- **Authenticator Device Model** is the hardware model name of the switch or access point. Stored in the authenticator record in Ignition Server as the **Device Template** name.
- **Authenticator Container** is the authenticator’s position in your authenticator hierarchy. It is stored in the authenticator record as the **Container**. The authenticator container designation is useful as an all-purpose label you can apply to an authenticator. See “Authenticator Hierarchy and Containers” on page 91.
- **Authenticator Name** is the name you gave to the authenticator in Ignition Server.

- **Authenticator Type** is the purpose/profile designation you gave to the authenticator in Ignition Server.
- **Vendor** is the manufacturer of the authenticator.

Using Time and Date in a Rule

Figure 60 The inbound attributes **Date**, **Date and Time**, and **Time** let you use the current Ignition Server date and time in a rule. **Setting Time and the Time Zone in a Rule**



Use the clock icons to set time periods, and use the drop down list to select the appropriate time zone. Note that Ignition Server timestamps each incoming transaction with the time of the locale *where Ignition Server is installed*.

The following table illustrates the effects of either using or not specifying a time zone for transactions arriving at a California-based Ignition Server from New York or Hawaii when the rule says to accept transactions only between 9 a.m. and 5 p.m.:

Table 7 Examples of Time-based Rules with and without Specifying a Time Zone

Transaction Arrival Time in California	Time Zone Specified for Transaction Type Received	Allow or Deny?	Reason
7 A.M.	None	Deny	Time is outside specified interval.
7 A.M.	Eastern Standard Time, such as New York	Allow	7 A.M. in California is 10 A.M. in EST, and so is within specified interval.
9 A.M.	None	Allow	Time is within specified interval.
9 A.M.	Hawaii Standard Time, such as Honolulu	Deny	9 A.M. in California is 6 A.M. in Hawaii.

Time comparisons use UTC, the Universal Time Code, which are based on GMT (Greenwich Mean Time). For example, Pacific Standard Time (PST) is GMT - 8 hours; Hawaiian Standard Time is GMT - 10 hours, and Eastern Standard Time in New York is GMT - 5 hours.

Some businesses want to allow service requests only during business hours, such as 9 A.M. to 5 P.M. in any locale. Using time zones can simplify applying such a rule to requests coming from different time zones. Otherwise, that same business rule requires different phrasing for each time zone. For example, if using an Ignition Server installed in California with no time zones, using that rule would have the following different implementations:

Separate Phrasings for the Same Rule

Time Zone for the Source of the Transaction	Phrasing of the 9-5 Rule
Pacific Standard Time	Greater Than or Equal 9 A.M. AND Less Than or Equal 5 P.M.
Eastern Standard Time	Greater Than or Equal 12 P.M. AND Less Than or Equal 8 P.M.
Hawaiian Standard Time	Greater Than or Equal 7 A.M. AND Less Than or Equal 3 P.M.

Rules may require adjustment when daylight savings time applies to Ignition Server or transaction locales.

Conjunctions Used to Assemble Constraints Into a Rule

Rules are built by linking constraints with AND and OR conjunctions, and by grouping them with parentheses.

Each conjunction connects the constraint (or parenthesized set of constraints) directly to its left to the constraint (or parenthesized set of constraints) directly to its right, as seen in the Summary box of the **Edit Authorization Policy window**. Note that in the **Constraint** section of the **Edit Authorization Policy window**, the conjunctions AND and OR appear in the last or rightmost column, and therefore connect constraints that appear to the left or above them to constraints that appear below them.

To assemble a rule, use:

- **Parentheses**: within a rule, Ignition Server evaluates a parentheses-enclosed set of constraints before it evaluates constraints outside the parentheses. Ignition Server works from the innermost parenthesis-enclosed set to the outermost, with the triple-parenthesis denoting the innermost set and the single-parenthesis denoting the outermost set. The third example in the following table shows this use of parentheses.
- The **AND** conjunction performs a logical AND on the two expressions that it links. The combination “X AND Y” is false unless both *X* and *Y* are true.
- The **OR** conjunction performs a logical OR on the two expressions that it links. The expression “X OR Y” is true unless *X* and *Y* are both false. The

OR conjunction is last in the order of operations, meaning that, in the absence of parentheses, constraints are ANDed before they are ORed.



Tip. Use parentheses to group your constraints. This makes your rules much easier to understand and lessens the likelihood of any unintended consequences. As a general rule, using parentheses helps you avoid ambiguity.

Table 8 Examples: Using AND and OR Conjunctions in Authorization Rules

Rule	Allow or Deny Checked?	Meaning and Comments
UserName Contains “Smith” **AND** account-locked IsFalse **AND** password-expiration is Greater Than 5/5/2006	Allow	In the example, all three constraints must be true for the request to be allowed. If the username does not contain “Smith”, or if the account is locked, or if the password does expire within the specified period, the request is rejected.
UserName Not Contains “Smith” **OR** account-locked IsTrue **OR** password-expiration is Less Than or Equal 5/5/2006	Deny	In the example, the request is denied if any of the 3 constraints is true. In other words, the request is rejected unless all 3 constraints are false.
(UserName Contains “Smith” **OR** UserName Contains “Davis”**) AND (**account-locked IsFalse **AND** password-expiration is Greater Than 5/5/2006**)**	Allow	In the example, as long as the account is not locked and the password expiration is in the chosen period, service requests from users with “Smith” or “Davis” in their usernames are allowed; all others are denied. For example, service requests from users with usernames that do NOT contain “Smith” or “Davis” are denied, as are usernames that DO contain those strings, but whose accounts are either locked or expire outside the specified period.

Comparison Operators for Rules

Each constraint is a comparison you create using one of the following comparison operators:

- **Any One Of**: Used to compare a value to a set. If the user/authenticator value matches any of the values in the comparison set, the rule evaluates to TRUE.
- **All Of**: Used to compare a set to a set. If the user/authenticator set matches all the values in the comparison set, the rule evaluates to TRUE.
- **Equals / Equal To**: Used in many types of comparisons. If the user/authenticator value exactly matches the comparison value, the rule evaluates to TRUE. Note, if you are having trouble with an Equals rule that evaluates to FALSE when you think it should be TRUE, try using the **Any One Of** comparison operator instead. The **Any One Of** comparison is less strict and might be the correct choice in some cases.

- **Not Any One Of**: Used to compare a value to a set. If the user/authenticator value fails to match any of the values in the comparison set, the rule evaluates to TRUE.
- **Not All Of**: Used to compare a set to a set. As long as the user/authenticator set does not contain the same set of the values as the comparison set, the rule evaluates to TRUE.
- **Not Equal To**: Used in many types of comparisons. As long as the user/authenticator value does not exactly match the comparison value, the rule evaluates to TRUE.
- **Starts With**: Used to compare a string to a sting. If the comparison value matches the initial characters of (or all the characters of) the user/authenticator value, the rule evaluates to TRUE.
- **Contains**: Used to compare a string to a sting. If the user/authenticator value contains the whole of the comparison value, the rule evaluates to TRUE.
- **Ends With**: Used to compare a string to a sting. If the comparison value matches the last characters of (or all the characters of) the user/authenticator value, the rule evaluates to TRUE.
- **Less Than**: Used to compare *Date*-type values as well as *Date and Time*-type values. If the access-request date is earlier than the comparison date, the rule evaluates to TRUE.
- **Less Than Or Equal**: Used to compare *Date*-type values as well as *Date and Time*-type values. If the access-request date is earlier than or the same as the comparison date, the rule evaluates to TRUE.
- **Greater Than**: Used to compare *Date*-type values as well as *Date and Time*-type values. If the access-request date is later than the comparison date, the rule evaluates to TRUE.
- **Greater Than Or Equal**: Used to compare *Date*-type values as well as *Date and Time*-type values. If the access-request date is later than or the same as the comparison date, the rule evaluates to TRUE.
- **Between**: Used to compare *Date*-type values as well as *Time*-type values. If the access-request date or time falls within the comparison range, the rule evaluates to TRUE.
- **Week Day Is Between**: Used to compare *Date*-type values. If the access-request date falls within the comparison range *and* is a Monday, Tuesday, Wednesday, Thursday or Friday, the rule evaluates to TRUE.

Creating a RADIUS User Authorization Policy

This section shows how to create user authorization and provisioning rules, and assemble them into a user authorization policy.

1. Select the Access Policy

Each user authorization policy applies within the scope of its access policy.

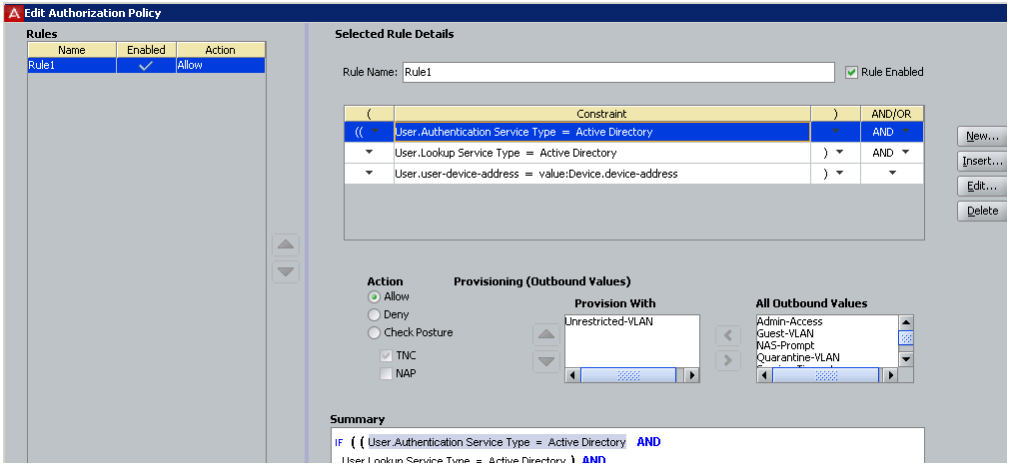
1.1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS.

1.2 Click **New** (or, if you wish to edit an existing policy, click its name in the tree, click the Authorization Policy tab, and click to top **Edit** button.)

2. Launch the Edit Authorization Policy window

2.1. With your access policy selected in the tree, click the **Authorization Policy** tab in the Access Policy panel and click **Edit**. The Edit Authorization Policy window appears displaying the policy's rules in sequence.

Figure 61 Edit Authorization Policy window



The Edit Authorization Policy window is a browser and editor. The **Rules** list at the left lets you browse and sort the rules in your policy. Use the up and down arrow buttons at the right to set the rule sequence, and click a rule name in the list to edit that rule. The remainder of this window

— the **Selected Rule Details** section — lets you edit the rule you have selected.

3. Add a New Rule

3.1. Click the **New** button just below the **Rules** list. Ignition Server displays the New Rule window.

Alternatively, you can copy an existing rule. See “Copying An Authorization Rule” on page 248.

3.2 Enter the name for the new rule in the New Rule window. Click **OK**. Ignition Server displays the name of the new rule in the **Rules** list of the Edit Authorization Policy window.

4. Set Up Rule Details

To view a rule, click on its name in the **Rules** list at the left of the Edit Authorization Policy window. The rule appears in the **Selected Rule Details** section of the window. Each rule consists of one or more **Constraints**, an **Action**, and optionally, **Provisioning values**.

Each row in the **Constraint** list is a test, which is called a “constraint” in Ignition Server. In the **Constraint** area, you can combine each constraint with the next or subsequent constraint using the AND and OR conjunctions (choose this by clicking the **And/Or** heading). The **Summary** section at the bottom shows the rule, including its action and provisioning values.

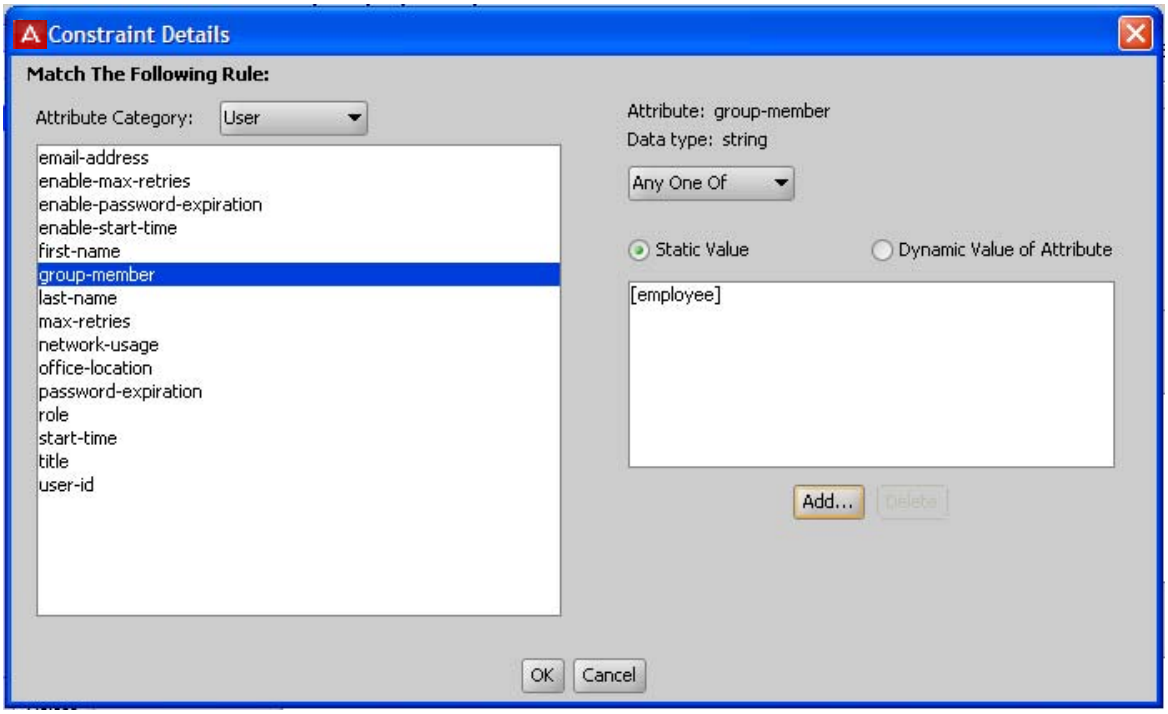
Build the Constraints

To add decision logic to your authorization rule, create one or more constraints. Each constraint tests the value of an attribute. If there are multiple constraints, join them into a single logical statement using the AND and OR conjunctions and, if needed, parentheses. Follow the steps below to do this:

4.1. On the left side of the Edit Authorization Policy window, click on the name of the **Rule** you want to edit.

4.2 To the right of the **Constraint** table, click the **New** button. The Constraint Details window appears.

Figure 62 Creating an authorization rule that evaluates a user attribute



4.3 In the **Attribute Category** drop down list, choose the type of attribute you want to test. (For explanations of the types, see "Attributes Used in Rule Constraints" on page 232.)

4.4 Choose the attribute: After you select a type, the list box below the **Attribute Category** field shows the available attributes that match the type you selected. Click on the name of the attribute whose value the constraint should test. In the upper right corner, the window displays the **Data type** of the attribute.

4.5 In the drop-down list just below the **Data type** field, choose the comparison operator, such as, *Equal To* or *Contains*. This drop-down list contains the operators appropriate to the data type of the attribute you have selected.

4.6 Provide the comparison value by doing one of the following:

- If you want to compare the attribute value with a fixed test value, tick the **Static Value** radio button and type or choose the comparison value in the field below that.
- If you want to compare the attribute value with a value retrieved from another attribute, tick the **Dynamic Value of Attribute** radio button. In the field just below that, choose the attribute category (User, Inbound, Authenticator, or Device). In the next field, choose the attribute that should provide the comparison value. The list of attributes contains

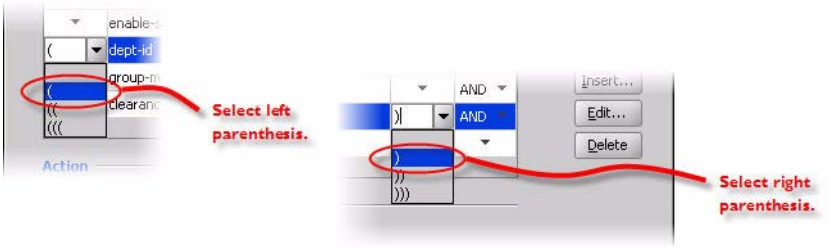
only those attributes whose data type matches the data type of the attribute on the left side of the constraint.

4.7 Click **OK** to close the Constraint Details window.

4.8 In the Edit Authorization Policy window, next to the **Constraint** table, click the **New** or **Insert** button to add more constraints. **New** adds a constraint at the end of the list, and **Insert** adds it above the currently selected row.

4.9 Add parentheses as necessary to group constraints. To do this:

- In the **Constraint** section of the Edit Authorization Policy window, find the first constraint to be grouped.
- Click in the field to the left of the constraint, and click the down-arrow to show the list of parentheses. Click on an appropriate opening parenthesis mark to select it.
- Find the last constraint to be grouped. Click in the field to the right of the constraint, and click the down-arrow to show the list of parentheses. Click on an appropriate opening parenthesis mark to select it. Click the constraint to complete your entry.



4.10 Use the **AND** and **OR** conjunctions to form a logical condition statement.

4.11 After you have finished adding constraints, click the **Allow** button or the **Deny** button to specify whether Ignition Server should grant or deny access if the rule evaluates to TRUE. See "How Ignition Server Evaluates a User Authorization Policy" on page 229 for information on the precedence of Allows and Denies in Ignition Server.

4.12 Optional: You may add provisioning to the rule. See "Set up Provisioning (Outbound Values)" on page 245.

4.13 Add additional rules to your user authorization policy as needed. To do this, go to the top of the Edit Authorization Policy window, click **New** to create a new rule, name the rule, and repeat the steps above.

5. Set up Provisioning (Outbound Values)

Within any rule in your user authorization policy you may add provisioning instructions to set characteristics of the user's network session such as a VLAN assignment, a session time-out, or administrator privileges. This section shows you how to do this.

In Ignition Server, a provisioning instruction is called an *outbound value*. An outbound value is a data value that Ignition Server sends to the authenticator as a RADIUS attribute when the rule triggers an *Allow* or *Deny*. Each rule can have zero, one, or many provisioning values associated with it.

Outbound values sent from an *Allow* rule are typically used to set characteristics of the user's session, while outbound values sent from a *Deny* rule are typically used to convey information about why the denial occurred.

Important! When writing provisioning rules, bear in mind that the Ignition Server policy engine evaluates all the rules in your rule set (until a *Deny* is triggered). If multiple *Allow* rules are triggered, then the outbound values of all of those rules are sent to the authenticator. If there are conflicts in the set of outbound values to be sent (for example, imagine that a rule set evaluation triggers the sending of both VLAN ID=200 and VLAN ID=201), then Ignition Server sends only the value associated with the *first-triggered rule*.

To add provisioning values to the rule:

5.1. Open the **Edit Authorization Policy window**.

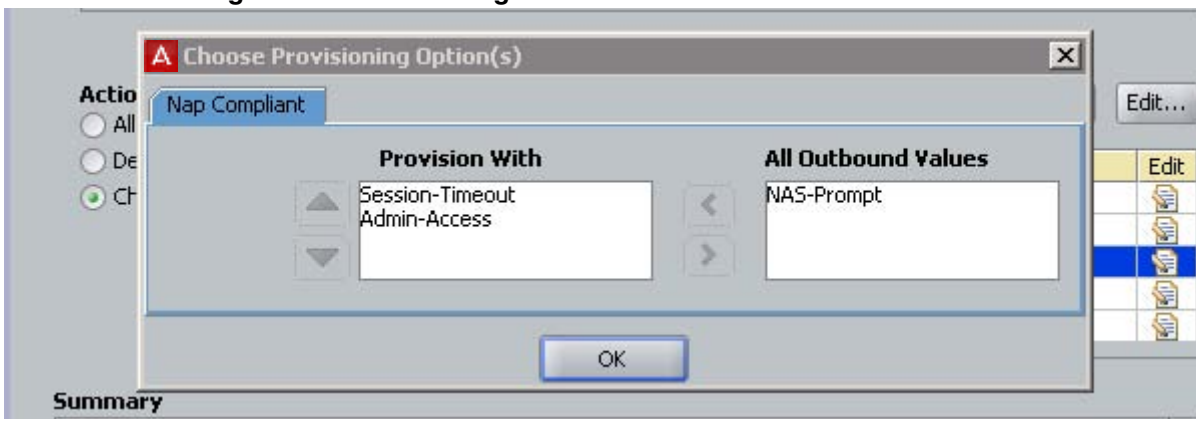
5.2 Click on the Rule to which you want to add provisioning. (Setting up rules is explained in "Build the Constraints" on page 243.)

5.3 Check the **Action** to make sure the *Allow* or *Deny* is set as desired. If a FEATURE_NAP license is installed, the *Check Postures* is automatically selected.

5.4 Select the rule you want to provision, and then click the Edit button. The Compliancy Condition list displays the NAP rules.

5.5 In the **Provisioning (Outbound Values)** section, select the outbound value you want to send from the **Provision with** list. Use the left and right arrow buttons to move your selected outbound value to the **Provision With** list, or, back to the **All Outbound Values** list.

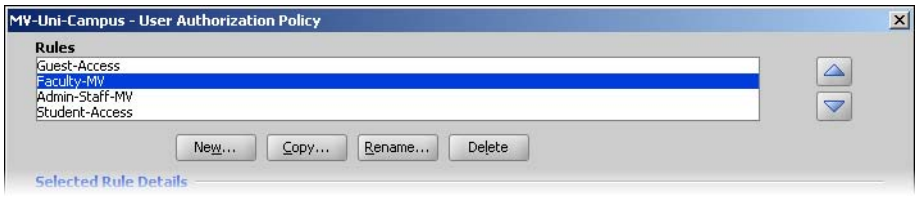
Figure 63 Provisioning Section



- To see the actual RADIUS attribute name and value to be sent, right-click an outbound value name in either list. A dialog window appears, showing the name/value pair.
- If the desired value does not appear in the **All Outbound Values** list, you must define it in the Outbound Values panel (**Configuration** tree: **Provisioning** node: **Outbound Values** node: **New**). For details, see Chapter , “Provisioning Policy”.

6. Set the Order of the Rules

Ignition Server displays the set of available rules in the **Rules** list on the left side of the Edit Authorization Policy window. The up and down arrow buttons on the right side allow you to sort the rules. Ignition Server evaluates the rules in the sequence you set here.



7. Review the Policy

After you have created the set of rules and arranged their order in the Rules list on the left side of the Edit Authorization Policy window, review each rule by clicking its name and reading its summary at the bottom of the window.

8. Save the Policy

Click **OK** to close the Edit Authorization Policy window. Ignition Server saves the contents of the authorization rules for the selected access policy. This returns you to the Access Policy panel of Dashboard, where you can review each of your rules by clicking its name in the **Rule Names** list.

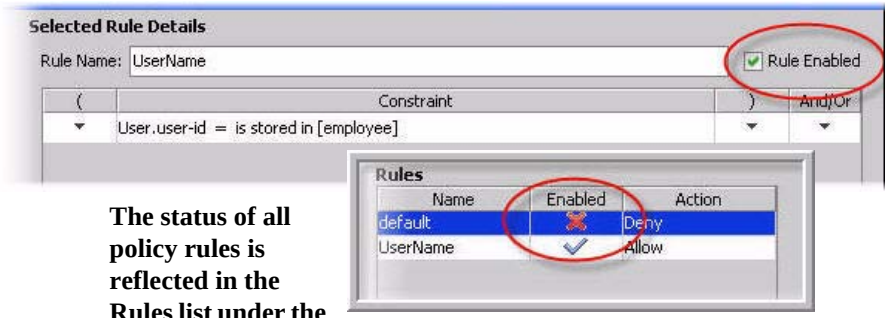
Next Steps: Now that you have saved added an authorization policy to your RADIUS access policy, you are ready to apply it to an authenticator, to a subauthenticator, or to the global authenticator. See “Introduction to Authenticators” on page 89.

Enabling or Disabling Rules Within a Policy

Ignition Server also allows you to enable or disable rules within a user authorization policy. This feature of Ignition Server lets you temporarily activate or deactivate an individual rule, without permanently deleting it from your policy.

To check if rules are enabled:

1. Open the **Edit Authorization Policy window**.
2. From the **Rules** list, select the rule you want to check the status of.
3. Under the **Selected Rule Details** section, next to the **Rule Name** field, is the **Rule Enabled** checkbox.
   - If this box is ticked, the highlighted rule is enabled (or active). Ticking this box ensures that Ignition Server evaluates the highlighted rule before allowing/denying a user access to the network.
4. To disable (or inactivate) an individual rule, simply un-tick the **Rule Enabled** checkbox.
   - Un-checking this box tells Ignition Server to bypass the highlighted rule when evaluating the **Rules** list of your authorization policy.



The status of all policy rules is reflected in the Rules list under the

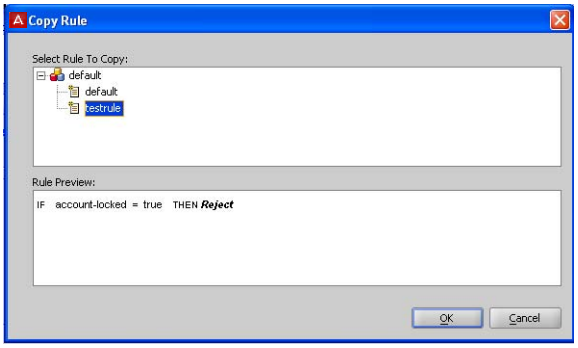
5. Click **OK** to close the window and save the changes to your policy.

Copying An Authorization Rule

You can copy rules within an access policy or from one access policy to another. To copy a rule:

1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS.
2. Click on the name of the access policy *into which you wish to copy a rule*.
3. Click on the **Authorization Policy** tab and click **Edit**.

4. In the lower left part of the window, click the **Copy** button.



5. In the **Copy Rule** window, navigate the tree to select the rule you want to copy. Click the + icon to view the rules of an access policy.
6. Click on the name of the rule and click **OK**. Ignition Server copies the selected rule into the policy that you are editing. The new rule appears in the **Rules** list of your access policy. Since you have made a copy, you can edit the new rule without affecting the original.

   You can rename the rule if you like. To do so, go to the **Rule Name** field in the **Selected Rule Details** section and replace the rule name with a new rule name.
7. Click **OK** to close the window and save your changes.

Creating an Authentication-Only Policy

Ignition Server always performs both authentication *and* authorization before it grants a user access, but in some installations, you may decide that authentication alone (checking the user's credentials) is sufficient to grant the user access. If this is the case, create a blanket **Allow** rule.

Hint: Why is it necessary to have an authorization rule at all, if I only want to check the user's password? The answer is that Ignition Server requires at least one rule to evaluate to **Allow** before it grants the user access.

There are two ways to create a blanket **Allow** rule:

Creating an Authentication-Only Rule

1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click the name of your access policy. Click the Authorization tab and click **Edit**.
2. Click the **New** button under the **Rules** section.
3. In the New Rule window, type a name for the rule (for example, you might call it, "catch-all") and click **OK**.

4. In the **Selected Rule Details** section of the Edit Authorization Policy window, click **New**.
5. In the Constraint Details window:
   - Select the **Attribute Category**, “System”.
   - Click the attribute, “True”.
   - Click **OK**.
6. In the **Action** section of the Edit Authorization Policy window, click **Allow**, and click **OK**.

Repeat this procedure for each additional access policy in which you want to add an authentication-only rule.

After you make the changes explained above, then, for all authenticators that use the specified access policy, a user may log in by authenticating successfully. Ignition Server effectively performs no authorization test, since the rule you created always evaluates to true.

Modifying the Default Rule to Make It Authentication-Only

A simpler, but less secure alternative to the above procedure is to modify the default rule in the “default” Access Policy, making it an authentication-only rule, as follows:

1. In Dashboard’s Configuration hierarchy, expand Access Policies, expand RADIUS, and click the *default-radius-user* policy to pick the default access policy.
2. In the Authorization Policy tab, click Edit.
3. In the **Action** section of the Edit Authorization Policy window, click **Allow**, and click **OK**.



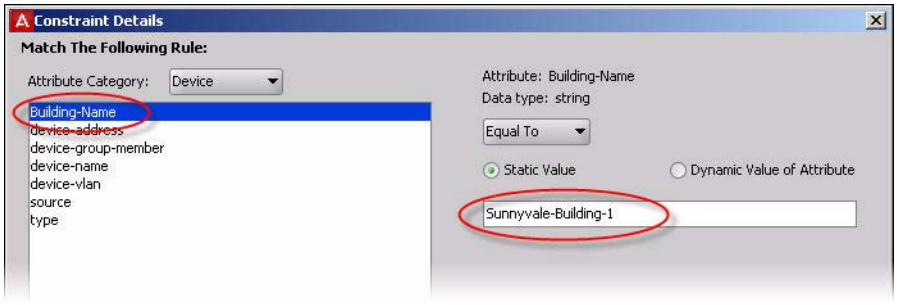
Warning: After you make the above changes, Ignition Server requires *only a successful authentication* to grant access. This applies any time a user logs in via a switch or access point in your default access policy.

Using a Device Attribute in a Rule

This section shows you how to evaluate properties of the connecting client device (for example, a user’s laptop or a printer) in your rules.

1. Open your authorization policy: In Dashboard’s Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click the name of your access policy. Click the Authorization tab and click Edit.
2. Write your authorization rule:

- In the Authorization Policy window, create a new rule or edit an existing one. Click on the rule's name to begin editing it.
- In the **Selected Rule Details** section, click **New** to create your constraint.
- In the Constraint Details window, select the **Attribute Category**, *Device*. In the list just below this, click the name of your device attribute. For a complete list, see "Device Attributes" on page 236. If the desired attribute is not there, add it as shown in "Device Virtual Attributes" on page 208.) On the right side of the window, define the logic of your constraint. (For instructions, see "Build the Constraints" on page 243.) Click **OK**.



- In the Authorization Policy window, with your rule still selected, tick the desired **Action**. If you want to send provisioning values, go to the **Provisioning** section and tick the checkbox next to each value you want to send.

3. Click **OK** to dismiss the Authorization Policy window.

Provisioning Policy

When a user or device authenticates, the Ignition Server policy engine provisions the user's or device's network session by sending instructions that make switch settings, set session time-outs, or assign the user or device to a VLAN. Ignition Server sends these instructions in the form of RADIUS attributes or VSAs. This chapter explains how to set up the RADIUS attributes Ignition Server uses to communicate with authenticators.

RADIUS attributes and VSAs also carry important information from the authenticators to Ignition Server, and you can configure Ignition Server to make authorization decisions based on that information. In addition, you can configure Ignition Server to return inbound RADIUS data as outbound RADIUS data. Starting on "Inbound Attributes" on page 265, this chapter explains how to set up these features of Ignition Server.

Contents

- "Introduction to Session Provisioning" on page 253
- "Setting Up Session Provisioning" on page 254
- "Vendors Panel" on page 255
- "Outbound Attributes" on page 255
- "Outbound Values" on page 258
- "Inbound Attributes" on page 265
- "Device Templates" on page 270
- "Listing Ignition Server's Set of Available RADIUS Attributes" on page 274
- "Adding a New RADIUS Attribute" on page 275
- "Listing Ignition Server's Set of Available VSA Attributes" on page 276
- "Adding a New VSA" on page 276
- "Adding an Equipment Vendor" on page 277
- "Provisioning FAQ" on page 278

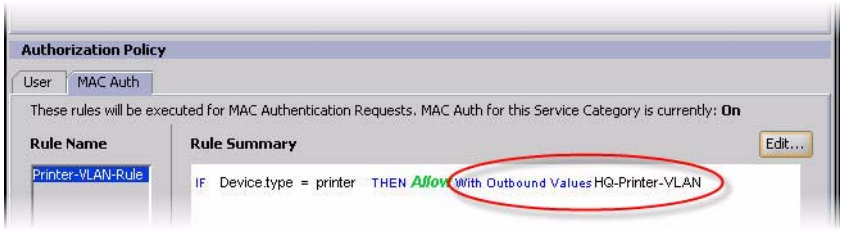
Introduction to Session Provisioning

Most network equipment accepts a variety of RADIUS attributes in incoming RADIUS messages, with each attribute configuring, or "provisioning," the user's network session in some way. For example, many VLAN-enabled

switches accept the *Tunnel-Private-Group-Id* attribute to assign a user to a VLAN. Setting session attributes like this one is referred to as “session provisioning.”

Ignition Server policies support session provisioning by allowing the administrator to assign provisioning values. Ignition Dashboard, your provisioning instructions are part of your user authorization and/or MAC authorization rules, as configured via the Access Policy panel of Dashboard. When a rule is triggered during user authorization, Ignition Server sends its provisioning value (or values) as RADIUS attributes to the authenticator.

Figure 64 Example Authorization Rule with Provisioning in the Access Policy panel



Setting Up Session Provisioning

To set up session provisioning in Ignition Server:

1. Create the *outbound attribute* as explained in “Outbound Attributes” on page 255. The outbound attribute specifies which RADIUS attribute or VSA that should carry the provisioning value.
2. Create the attribute-value pair (or pairs) that Ignition Server should send the authenticator to provision the session. This pair, or set of pairs, is called an *outbound value* in Ignition Server. For instructions, see “Outbound Values” on page 258.
3. Specify the conditions that will trigger Ignition Server to send your outbound value. You can:
   - include your outbound value in a device template and apply that device template to your authenticator definition (for instructions, see “Device Templates” on page 270); or
   - write a rule in an access policy that, when triggered, sends an outbound value (for instructions, see “Set up Provisioning (Outbound Values)” on page 245).

Important! Before you set up session provisioning, note the following:

- **Built-in outbound values**: Ignition Server contains a number of built-in outbound values. See “Built-In Outbound Values” on page 261.

- **Hardware support**: Provisioning depends on the authenticator hardware's support for the RADIUS attributes or VSAs that you configure Ignition Server to send. Check your equipment documentation to make sure that the equipment accepts the attribute name and data type you plan to use, and that it responds appropriately to the values you plan to send.

Vendors Panel

The Vendors panel lets you manage RADIUS attributes, VSAs, and vendor-specific communications options for authenticators. The window's navigation tree is sorted by authenticator manufacturer, with a separate entry, "RADIUS", used to manage RADIUS attribute definitions. This window can be used for:

- Finding a Device Template, Modifying a Device Template, or Applying a Device Template to Your Authenticator
- Listing Ignition Server's Set of Available RADIUS Attributes and Listing Ignition Server's Set of Available VSA Attributes
- Adding a New RADIUS Attribute and Adding a New VSA
- Adding an Equipment Vendor
- Overriding the Outbound Attribute Type for One or More Authenticators
- Finding an Inbound Attribute and Creating a Vendor-Specific Inbound Attribute

Outbound Attributes

Outbound attributes are the data fields Ignition Server uses to carry provisioning data to authenticators. In technical terms, outbound attributes are RADIUS or VSA attributes that Ignition Server can include in messages to authenticators.

The first task in setting up provisioning in Ignition Server is to create the outbound attributes that should carry your provisioning values. Do one of the following:

- if your outbound attribute is used by *more than one* make and model of authenticator, then create it as a global attribute, as shown in "Creating a Global Outbound Attribute", below; or
- if your outbound attribute is used by *only one* make and model of authenticator, then create it inside the device template for that authenticator type, as shown in "Overriding the Outbound Attribute Type for One or More Authenticators" on page 257.

Finding a Global Outbound Attribute

To find a global outbound attribute:

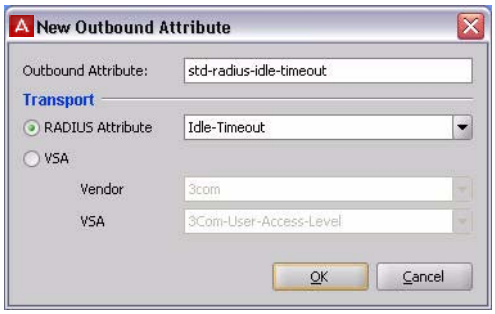
1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Outbound Attributes**.
2. In the Outbound Attributes panel, scroll to find your attribute.

Creating a Global Outbound Attribute

To define an outbound attribute that can be used to transmit a value in RADIUS messages to authenticators of any type, follow the steps below:

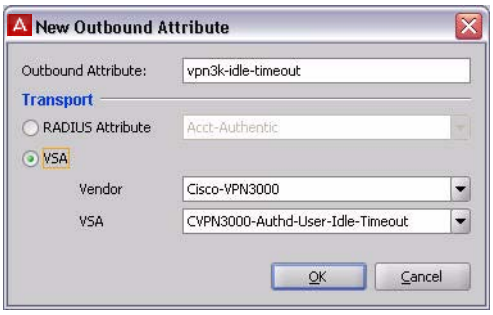
1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Outbound Attributes**.
2. In the Outbound Attributes panel, click **New**.
3. In the New Outbound Attribute window, type a name for the Ignition Server outbound attribute in the **Outbound Attribute** field.
4. Choose the attribute that should carry your outbound value. Do **(a)** or **(b)** below:

   a. *To use a standard RADIUS attribute*, click the **RADIUS Attribute** radio button and in the drop down box, choose the name of the RADIUS attribute. (If the desired attribute is missing from the list, see "Adding a New RADIUS Attribute" on page 275.) Click **OK**.



   b. *To use a VSA*, click the radio button to select **VSA**. In the **Vendor** drop down box, choose the name of the manufacturer of the authenticator equipment you provision, and in the **VSA** drop down box, select the name of the vendor specific attribute you want to send. Click **OK**. (If your equipment manufacturer name or VSA name is missing, see

"Adding an Equipment Vendor" on page 277 or "Adding a New VSA" on page 276.)



5. Your new attribute now appears in the list in the Outbound Attributes panel.

Next steps: You have finished creating the outbound *attribute* that is to carry your provisioning value to the authenticator. Next you must specify what *value* is sent in the attribute, as explained in the section, "Outbound Values" on page 258.

Overriding the Outbound Attribute Type for One or More Authenticators

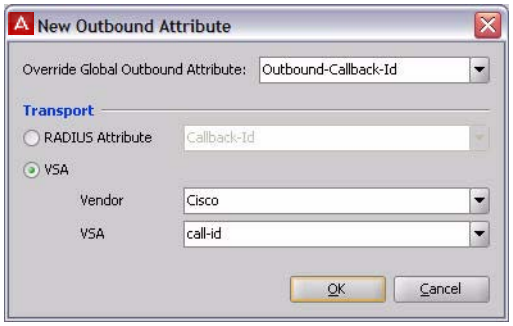
You can create an override which forces Ignition Server to use a *different RADIUS attribute than usual* when sending to a specific authenticator or authenticators. We refer to this as "overriding a global outbound attribute."

To specify a non-standard outbound attribute to be used in RADIUS messages to authenticators of a single type, follow the steps below:

Note: This procedure overrides your outbound attribute within the context of a device template so that it *can be used only by authenticators that use that template*. To create an outbound attribute you can use globally, see "Creating a Global Outbound Attribute" on page 256.

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. In the Vendors panel, locate the manufacturer of your authenticator, click its name to expand the list, and click **Device Templates**.
3. In the **Device Templates** list, find your template. Click its name and click **Edit**. If the desired template does not exist, create it now as shown in "Creating a Device Template" on page 271.
4. In the Device Template window, click the **Outbound Attributes** tab. Click **New**.

5. In the New Outbound Attribute window, in the **Override Global Outbound Attribute** drop-down list, choose the outbound attribute that should be overridden.



6. In the Transport section, choose the attribute that you want to contain values of this type. To do this, do one of the following:
   - to use a *standard RADIUS attribute* to carry the provisioning value, click **RADIUS attribute** and select the attribute name from the drop down list (If the desired attribute is not in the list, see "Adding a New RADIUS Attribute" on page 275); or
   - to use a *vendor-specific attribute*, click **VSA**, select your authenticator **Vendor**, and select your **VSA** name. (If the desired VSA or vendor is not in the list, see "Adding a New VSA" on page 276 or "Adding an Equipment Vendor" on page 277).
7. Click **OK**.

Next steps: You have finished overriding the outbound *attribute*. If you have not already specified an *outbound value* to be sent, do so now as shown in "Creating an Outbound Value" on page 259. Bear in mind that when you create the outbound value, you should choose the *Global Outbound Attribute name* as the attribute name, and, when Ignition Server sends the RADIUS message, the override ensures that the non-standard outbound attribute is used, according to your override definition.

Outbound Values

Outbound values are the provisioning data that Ignition Server sends to authenticators. In technical terms, the outbound value is a RADIUS attribute-value pair or pairs. The second task in setting up provisioning in Ignition Server is to create your outbound value as shown in "Creating an Outbound Value".

Finding an Outbound Value

To find an outbound value:

1. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**.
2. In the Outbound Values panel, scroll to find your outbound value.

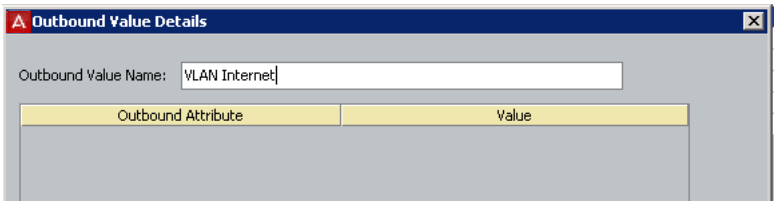
   To edit an outbound value, click its row and click the **Edit** button. See "Creating an Outbound Value" on page 259 for instructions on using the Outbound Value Details window.

Creating an Outbound Value

This section shows how to create an outbound value. After you create the outbound value, you must write an authorization rule that triggers Ignition Server to send the value to the authenticator.

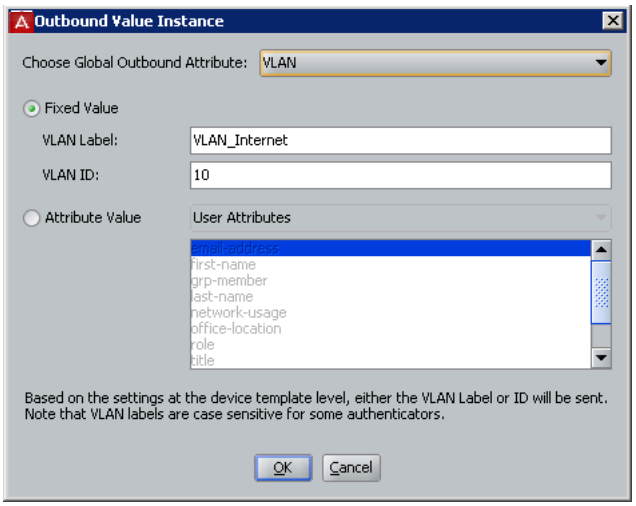
First create the outbound value:

1. Make sure you have an appropriate outbound attribute to carry each provisioning message. If you do not have the right attributes, see "Outbound Attributes" on page 255.
2. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**.
3. The Outbound Values panel lists all the sets of outbound values that have been defined in your Ignition Server. Click **New** to create a new value.
4. In the Outbound Value Details window, type an **Outbound Value Name** for the outbound value. This is the name that you will later choose in your authorization policy to send this value.

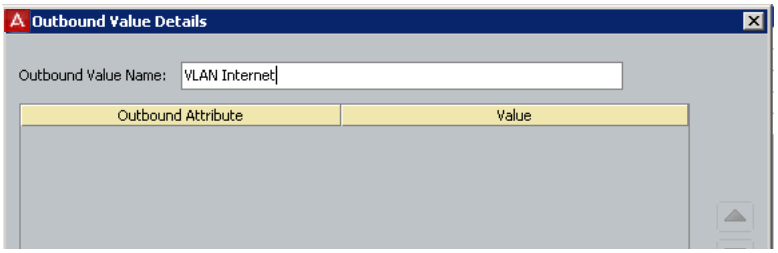


5. Click **New** to begin adding a name-value pair that will be sent. Most outbound values send one name-value pair, but you can send as many as needed. The Outbound Value Instance window lets you create each name-value pair. To do this:
   - In the **Choose Global Outbound Attribute** drop down list, select the name of the outbound attribute that will carry the value. The outbound attribute establishes the datatype. This can be a standard outbound attribute or a custom one you created as explained in Step 3 above.

- In the **Value** section, set or map the provisioning value that Ignition Server will send to the authenticator. See "Setting a Provisioning Value" on page 261 for details.



6. Click **OK**.
7. **Optional:** You can define multiple attribute-value pairs to be sent in this outbound value. If you want to do so, click **New** again in the Outbound Value Details window, and repeat the previous steps, as many times as needed, to add the attribute-value pairs you want.
8. The newly defined attribute-value pair (or pairs) appears in the Outbound Value Details window. Click **OK** to save the pair(s) and dismiss the window.



The newly defined outbound value appears in the Outbound Values panel. When you edit your authorization policies, it will be available in the **Provisioning (Outbound Values)** section of the Edit Authorization Policy window.

Next steps: Now that you have finished creating the outbound value, you must create the authorization rule(s) to trigger Ignition Server to send this outbound value to your authenticator. For instructions, see "Set up Provisioning (Outbound Values)" on page 245.

Built-In Outbound Values

A number of outbound values are built into your default installation of Ignition Server. To use these, you do not need to define a new outbound attribute or value. Instead, just add them to a rule in your user authorization policy or MAC authorization policy, as explained in "Set up Provisioning (Outbound Values)" on page 245.

The built-in outbound values are:

- **Admin-Access**, which sends the RADIUS attribute *Service-Type* with a value of "Administrative-User" (an integer value of 6). On most equipment, this code indicates the user is to be given a session that grants him or her access to administrative commands.
- **NAS-Prompt**, which sends the RADIUS attribute *Service-Type* with a value of "NAS-Prompt" (an integer value of 7). On most equipment, this code indicates the user is to be given a command prompt on the NAS from which non-privileged commands can be executed.
- **Session-Timeout**, which sends the RADIUS attribute *Session-Timeout* with the integer number of seconds the user's session lasts before he or she must reauthenticate. Use this attribute to set your 802.1X client reauthentication frequency.

Setting a Provisioning Value

In the Outbound Value Details/Outbound Value Instance window, you specify a provisioning value Ignition Server can send to an authenticator. There are three types of values you can send:

- a static value. See "Assigning a Static Value to an Outbound Value", below.
- information from the user's record. See "Passing Value from the User Record or Device Record to an Outbound Value" on page 262.
- information from the authenticator. See "Passing an Inbound Value to an Outbound Value" on page 264.

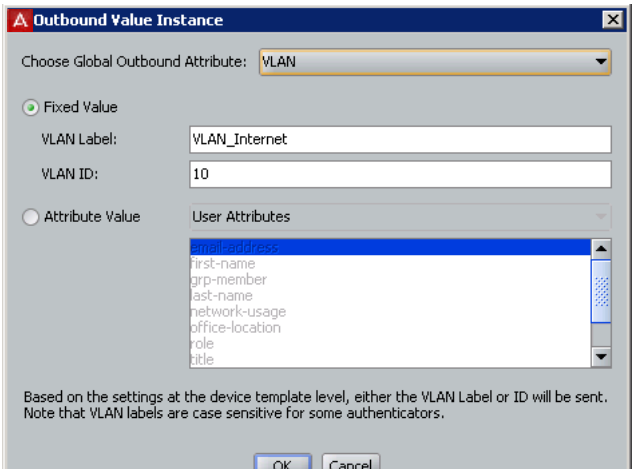
Assigning a Static Value to an Outbound Value

By creating an outbound value whose value is fixed, you create a piece of provisioning data that you can send to a switch to trigger a standard action or behavior in the switch. This value is the same every time you send it.

To assign a static value to an outbound value do the following:

1. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**.

2. Double-click the value to be edited or click **New**. If you are creating a new outbound value, type a name for it in the **Outbound Value Name** field of the Outbound Value Details window.
3. In the Outbound Value Details window, click **New** or, if you already have an **Outbound Attribute** you want to use, double-click its name. (The outbound attribute is the RADIUS attribute carries your static value.)
4. If necessary, in the upper part of the Outbound Value Instance window, select the **Global Outbound Attribute** that is to carry this value. If the outbound attribute has already been set, this field is not editable.
5. In the **Value** section, click the upper radio button. (Or, in the VLAN version of the window, click the **Fixed Value** radio button.) The legend next to the radio button indicates the datatype.
6. In the field to the right of the radio button, enter the value to be sent in this attribute-value pair. The form of the field depends on the datatype.
7. Click **OK**.



8. In the Outbound Value Details window, you have the option of adding more attribute-value pairs to this single outbound value. To do so, click **New**. Otherwise, click **Save**.

Next steps: If you have not already done so, you must create an authorization rule to trigger Ignition Server to send this outbound value to your authenticator. For instructions, see "Set up Provisioning (Outbound Values)" on page 245.

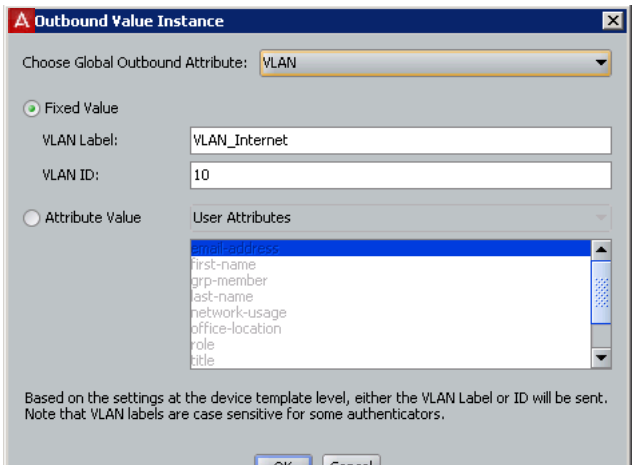
Passing Value from the User Record or Device Record to an Outbound Value

You can retrieve user data from the user record or device record and pass this data to an authenticator in an outbound value. Set this up as follows:

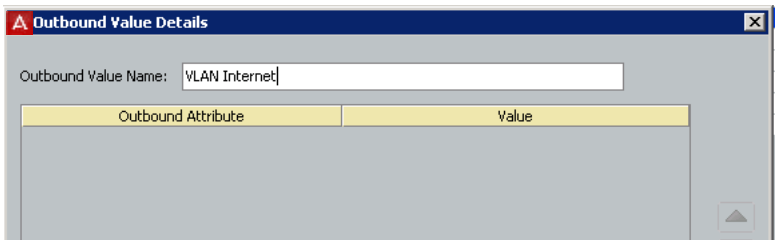
1. For each user data field or device data field from which you want to retrieve data, define an Ignition Server virtual attribute as explained in "User Virtual Attributes" on page 203 or "Device Virtual Attributes" on page 208.
2. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**. Double-click the value to be edited or click **New**. If you are creating a new outbound value, type a name for it in the **Outbound Value Name** field of the Outbound Value Details window.
3. In the Outbound Value Details window, click **New** or, if you already have an **Outbound Attribute** you'd like to use, double-click its name. (The outbound attribute is the RADIUS attribute you want to carry your value.)
4. If necessary, in the upper part of the Outbound Value Instance window, in the **Choose Global Outbound Attribute** field, choose the outbound attribute that you want to carry this value. If the outbound attribute has already been set, this field is not editable.

   Important! Make sure that the outbound attribute you have chosen has the correct data type for the value you want to send.
5. In the **Value** section, tick the **Attribute Value** radio button.
6. In the drop down list to the right of the radio button, select **User Attributes** or **Device Attributes**.
7. In the list box just below this, select the virtual attribute name. The list box contains only those virtual attributes whose data type matches that of the outbound attribute you selected in Step 4 above. (For information on checking the datatype of the virtual attribute, see "Browsing User Virtual Attributes" on page 204. For information on checking the datatype of the outbound attribute, see "How can I find out the datatype of an inbound or outbound attribute?" on page 279.)

   The specified virtual attribute's value is retrieved from the user record and placed in the outbound value at runtime.
8. Click **OK**.



9. In the Outbound Value Details window, you have the option of adding more attribute-value pairs to this specific outbound value. To do so, click **New**. Otherwise, click **Save**.



Next steps: If you have not already done so, you must create an authorization rule that triggers Ignition Server to send this outbound value to your authenticator. For instructions, see "Set up Provisioning (Outbound Values)" on page 245.

Passing an Inbound Value to an Outbound Value

You can pass an inbound value (data received from the authenticator) back to the authenticator in an outbound value. Set this up as follows:

1. For each inbound value that you want to use, define an Ignition Server inbound value as explained in "Inbound Attributes" on page 265.
2. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**. Double-click the value you want to edit or click **New**. If you are creating a new outbound value, type a name for it in the **Outbound Value Name** field.
3. In the Outbound Value Details window, click **New** or, if you already have an **Outbound Attribute** you'd like to use, double-click its name. (The outbound attribute is the RADIUS attribute that you want to carry your value.)
4. If necessary, in the upper part of the Outbound Value Instance window, in the **Choose Global Outbound Attribute** field, select the outbound attribute that you want to carry this value. If the outbound attribute has already been set, this field is not editable.



Important! Make sure that the outbound attribute you have chosen has the correct datatype for the value you want to send.

5. In the **Value** section, tick the **Attribute Value** radio button.
6. In the drop down list to the right of the radio button, select **Inbound Attributes**.

7. In the list box just below this, select the inbound attribute name. The list box contains only those inbound attributes whose datatype matches that of the outbound attribute you selected in Step 4 above. (For information on checking data types, see "How can I find out the datatype of an inbound or outbound attribute?" on page 279.)

   The specified attribute's value is copied from the incoming RADIUS message and placed in the outbound value at runtime.

8. Click **OK**.

9. In the Outbound Value Details window, you have the option of adding more attribute-value pairs to this specific outbound value. To do so, click **New**. Otherwise, click **Save**.

Next steps: If you have not already done so, you must create an authorization rule to trigger Ignition Server to send this outbound value to your authenticator. For instructions, see "Set up Provisioning (Outbound Values)" on page 245.

Inbound Attributes

Ignition Server can make use of information that the authenticator sends in its RADIUS request. In Ignition Server terminology, a piece of data that Ignition Server receives from the authenticator is called an *inbound value*, and is carried in an *inbound attribute*. Provided the correct inbound attributes have been defined in your Ignition Server configuration, you can configure Ignition Server to:

- evaluate an inbound attribute in an authorization rule. See "Preparing an Inbound Attribute for Use in an Authorization Rule" on page 265; and/or
- return the inbound attribute in the RADIUS response. See "Passing an Inbound Value to an Outbound Value" on page 264.

Your default Ignition Server installation contains inbound attribute definitions for many of the most popular RADIUS attributes. If you want to evaluate a RADIUS attribute or VSA that is not part of the default set, you must define a new inbound attribute, as explained in the following sections.

Preparing an Inbound Attribute for Use in an Authorization Rule

You can evaluate an incoming RADIUS attribute or VSA in your authorization rules. Follow the steps below to set this up for a typical RADIUS attribute or VSA. (*Note:* Ignition Server contains a default set of pre-defined inbound attributes, some of which are listed in "Inbound Attributes" on page 234.)

First, define your inbound attribute:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Inbound Attributes**.

2. In the Inbound Attributes panel, click **New** to create the new attribute.
3. In the New Inbound Attribute window, in the **Inbound Attribute** field, type a name for this attribute. This is the name you see when writing your policy rules.
4. Specify the source RADIUS attribute for this inbound attribute. Do one of the following:
   - If the source is a RADIUS attribute, click the **RADIUS Attribute** button and select the RADIUS attribute name from the drop down list. (If the desired attribute is missing from the list, see "Adding a New RADIUS Attribute" on page 275.)
   - If the attribute is a VSA, click the **VSA** button, select the **Vendor** (manufacturer whose equipment supports this VSA), and select the attribute name from the **VSA** drop down list. (If your equipment manufacturer name or VSA name is missing, see "Adding an Equipment Vendor" on page 277 or "Adding a New VSA" on page 276.)
5. Click **OK**.

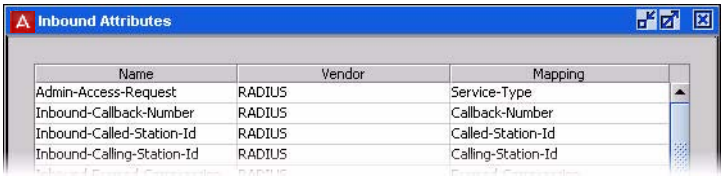
Next, define your authorization rule to evaluate the inbound attribute:

1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click the name of your access policy. Click the Authorization tab and click Edit.
2. Create your rule by clicking **New (under the Rules list**), giving the rule a name, and clicking **OK**. (Or you can edit an existing rule.)
3. In the **Selected Rule Details** section, click **New** to add a constraint.
4. In the Constraint Details window, in the **Attribute Category** drop down list, select **Inbound**.
5. In the list of inbound attributes, select the inbound attribute you saved earlier.
6. On the right side of the window, define the comparison condition that must be met in order to trigger this rule. (For help, see "Creating a RADIUS User Authorization Policy" on page 242.)
7. Click **OK**.

Finding an Inbound Attribute

To view the lists of available inbound attributes:

1. View the list of available *global* inbound attributes by expanding the **Provisioning** node of Dashboard's **Configuration** tree and clicking **Inbound Attributes**.

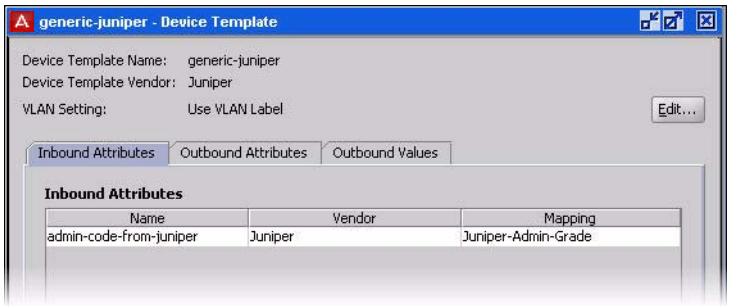


   In the Inbound Attributes panel:

   - the **Name** column shows the name used in your authorization rules to refer to the attribute.
   - the **Vendor** column shows the authenticator vendor associated with the attribute. If the attribute is a standard RADIUS attribute, the vendor is "RADIUS".
   - the **Mapping** column shows the RADIUS attribute name or VSA name of the attribute.

2. View the list of available *device template-specific* inbound attributes by doing the following:

   - In Dashboard's Configuration hierarchy tree, expand Provisioning and click Vendors/VSAs.
   - In the Vendors panel, locate the manufacturer of your authenticator, click its name to expand the list, then click **Device Templates**.
   - In the **Device Templates** list, select your template and click **Edit**.
   - In the Device Template window, click the **Inbound Attributes** tab. The columns are the same as those in the Inbound Attributes panel described above.



Creating a Global Inbound Attribute

If you want to retrieve an inbound RADIUS attribute value, you must define an inbound attribute, as shown in this section. After the inbound attribute is defined, you can evaluate it in your authorization rules; you can map the attribute to an outbound value so that Ignition Server can send the value back to the authenticator in RADIUS messages; or, you can do both.

Define the inbound attribute as follows:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Inbound Attributes**.
2. Click **New**.
3. In the **Inbound Attribute** field, type a name for the attribute. This name is used to refer to this attribute in your authorization rules (when setting up logic that evaluates the attribute's value), or in your outbound value definition (when passing an inbound value as an outbound value).



4. In the **Transport** section, choose the **RADIUS Attribute** that contains values of this type in the inbound RADIUS messages that authenticators send to the Ignition Server. Choose one of the following:
   - to retrieve the value from a standard RADIUS attribute, click the **RADIUS Attribute** radio button, and select the attribute name from the drop down list (If the desired attribute is not in the list, see "Adding a New RADIUS Attribute" on page 275); or
   - to retrieve the value from a vendor-specific attribute, click the **VSA** radio button, select your authenticator **Vendor**, and select your **VSA** name. If the desired **VSA** or **Vendor** is not in the list, see "Adding a New VSA" on page 276 or "Adding an Equipment Vendor" on page 277.
5. Click **OK**.

Next steps: Now that you have finished creating the inbound attribute, you can evaluate its inbound value in an authorization rule (see "Inbound Attributes" on page 265) or you return the inbound value in the RADIUS response (see "Passing an Inbound Value to an Outbound Value" on page 264).

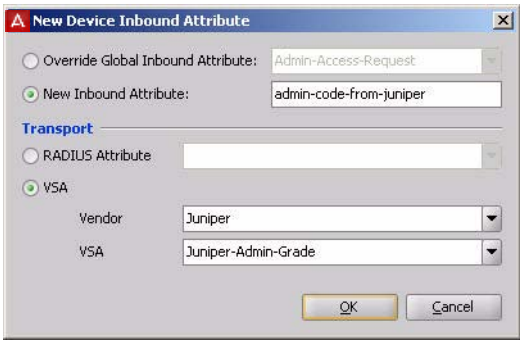
Creating a Vendor-Specific Inbound Attribute

If you want to retrieve an inbound RADIUS attribute from a specific type of authenticator only, then you define the inbound attribute within the *device template* for that authenticator type. If you define the attribute in this way, you can map the inbound attribute to an outbound value in the device template, and Ignition Server includes this template value in the RADIUS messages sent to every authenticator that uses the template.

Note: You cannot evaluate a vendor-specific inbound attribute in your authorization rules. To create an inbound attribute that can be used in a rule, define it as shown in "Creating a Global Inbound Attribute" on page 268.

Define the vendor-specific inbound attribute as follows:

1. In Dashboard's Configuration hierarchy tree, expand Provisioning and click Vendors/VSAs.
2. In the Vendors panel, locate the manufacturer of your authenticator, click its name to expand the list, then click **Device Templates**.
3. In the **Device Templates** list, select your template and click **Edit**. If your desired template does not exist, create it now as shown in "Creating a Device Template" on page 271.
4. In the Device Template window, click the **Inbound Attributes** tab. Click **New**.
5. In the New Device Inbound Attribute window, do one of the following:
   - to override an *existing* attribute, select **Override Global Inbound Attribute** and choose the attribute name; or
   - to create a *new* attribute, click **New Inbound Attribute** and type a name for the attribute. This name is used to refer to this attribute in your authorization rules (when setting up logic that evaluates the attribute's value), or in your outbound value definition (when passing an inbound value as an outbound value).



6. In the **Transport** section, choose the RADIUS attribute that contains values of this type in the inbound RADIUS messages that authenticators send to the Ignition Server. Choose one of the following:
   - to retrieve the value from a standard RADIUS attribute, click **RADIUS attribute** and select the attribute name from the drop down list (If the desired attribute is not in the list, see "Adding a New RADIUS Attribute" on page 275); or
   - to retrieve the value from a vendor-specific attribute, click **VSA**, select your authenticator **Vendor**, and select your **VSA** name. (If the desired VSA or vendor is not in the list, see "Adding a New VSA" on page 276 or "Adding an Equipment Vendor" on page 277).
7. Click **OK**.

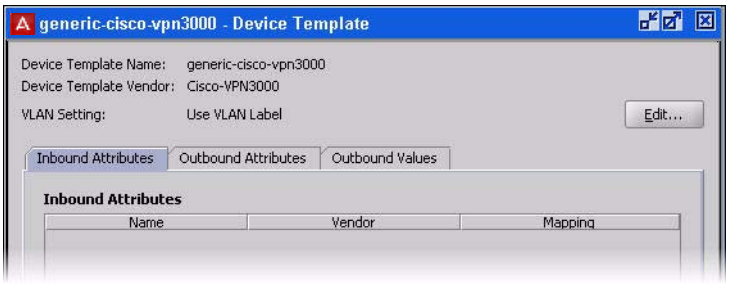
Next steps: Now that you have finished creating the inbound attribute, you can evaluate its value in an authorization rule (see "Inbound Attributes" on page 265), or you can return the attribute value in the RADIUS response to the authenticator (see "Passing an Inbound Value to an Outbound Value" on page 264).

Device Templates

Ignition Server has a configuration tool called a *device template. If* that specifies a default set of outbound provisioning values that Ignition Server always sends to a given type of authenticator. In addition, the device template establishes the set of inbound attributes that Ignition Server expects to receive from the authenticator, and, if applicable, the VLAN designation format.

When you configure one of your switches or other devices as an authenticator in Ignition Server (see "Creating an Authenticator" on page 96), you apply a device template to that authenticator. The device template you use can be the default template (the default installation contains default templates for most popular authenticators), or a custom template you have created.

When you set up a device template, you specify which values Ignition Server passes as outbound attributes, whether each value is a hard-coded value, a value retrieved from the user record, or an inbound attribute that Ignition Server reflects back as an outbound value.



Device Template Window

The Device Template window allows you to define the outbound values that Ignition Server sends to the switch (or authenticator), as well as the inbound attributes that it expects to receive in RADIUS messages from the switch (or authenticator).

- To create a device template, see "Creating a Device Template" on page 271.
- To find a device template, see "Finding a Device Template" on page 271
- To edit a device template, see "Modifying a Device Template" on page 273
- To set up inbound attributes for a device, see "Inbound Attributes" on page 265.
- To override the use of a RADIUS attribute for a device, see "Overriding the Outbound Attribute Type for One or More Authenticators" on page 257.

Finding a Device Template

Use the following steps to find a device template in Ignition Server:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. In the left panel of the **Vendors** panel, double-click the manufacturer name of your network equipment, then click **Device Templates** to display the list of templates for that manufacturer.

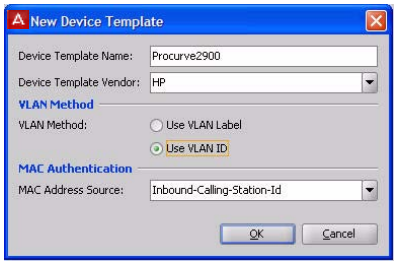
Creating a Device Template

Use the following steps to create a device template:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**. Scroll to find your vendor and expand its node. Click

the **Device Template** node that appears in the tree. Click **New** at the bottom of the Device Templates panel.

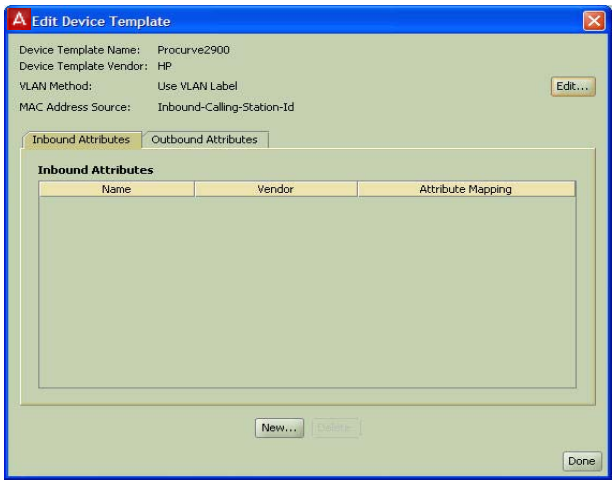
2. In the New Device Template window, type a name for the template in the **Device Template Name** field.
3. Select your authenticator vendor from the **Device Template Vendor** drop down list. (If your vendor's name is not in the list, see "Adding an Equipment Vendor" on page 277).
4. From the **VLAN Method** radio buttons: choose **Use VLAN Label** if your switch, or authenticator, uses an ASCII text label to identify the VLAN. Choose **Use VLAN ID** if your switch, or authenticator, uses an integer ID number.



5. In the **MAC Address Source field**, choose the RADIUS Attribute that contains the MAC address of connecting devices.

Warning: If you're doing MAC authentication of a device, Ignition Server gets the MAC address from the RADIUS attribute you specify as the **MAC Address Source** in the device template, but if you're doing user authentication coupled with an asset check of the user's device, then Ignition Server always gets the device's MAC address from the *inbound-calling-station-id* RADIUS attribute.

6. Click **OK**. The Device Template window appears. Your device template has been saved, but it contains no inbound or outbound attribute definitions. Instructions are provided for this later.



7. Apply your device template to each authenticator that is to use it. See “Applying a Device Template to Your Authenticator” on page 273 for details.

Next steps: Use the tabs of the Device Template window to define the outbound values that Ignition Server sends to the authenticator and the inbound attributes that it receives in RADIUS messages from the authenticator.

- To set up inbound attributes, see “Inbound Attributes” on page 265.
- To override the use of a RADIUS attribute for a particular type of provisioning value, see “Overriding the Outbound Attribute Type for One or More Authenticators” on page 257.

Modifying a Device Template

To change a device template, use the following steps:

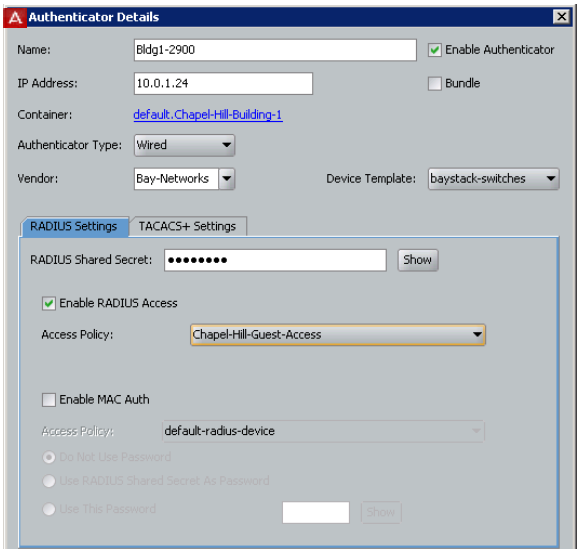
1. In Dashboard’s **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. In the left panel of the **Vendors** panel, double-click the manufacturer name of your network equipment, then click **Device Templates** to display the list of templates.
3. In the list on the right, select the name of your template and click **Edit**. The Device Template window appears.
4. To toggle the VLAN identifier between ASCII code and integer ID, click the **Edit** button to the far right of the **VLAN Method** field.
5. To edit inbound and outbound attributes and values, use the tabs in the Device Template window.
   - To set up **Inbound Attributes**, see “Inbound Attributes” on page 265.
   - To set up **Outbound Attributes** and corresponding **Values,** see “Creating a Global Outbound Attribute” on page 256.

Applying a Device Template to Your Authenticator

To apply a device template to an authenticator:

1. In Dashboard’s Configuration hierarchy tree, expand Authenticators. Find your authenticator in the tree, click its name, and click Edit.

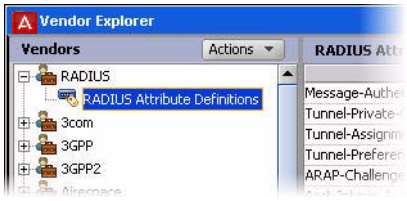
2. In the Authenticator Details window, select the template name in the **Device Template** drop down list, and click **OK**.



Listing Ignition Server's Set of Available RADIUS Attributes

To view a list of Ignition Server's defined RADIUS attributes:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. In the left panel of the Vendors panel, double-click **RADIUS,** then click **RADIUS Attribute Definitions** to display the list of attribute types. Click on the column headings to sort the attribute list.



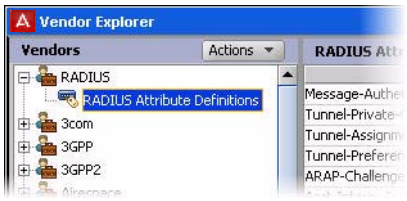
The **Name** is the RADIUS attribute name that Ignition Server and your network equipment use to identify this attribute. The **Data Type** indicates what kind of data the attribute can contain, such as a string or an unsigned 32-bit integer. The **Attribute Type** is the integer code that designates the attribute, as specified in the RADIUS specification, or a relevant industry standards document. A blue check mark in the **Default** column indicates that the attribute is one of the default attributes included in your standard installation of Ignition Server.

To add a new RADIUS attribute, see "Adding a New RADIUS Attribute" below.

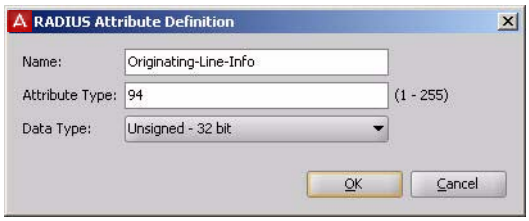
Adding a New RADIUS Attribute

If the RADIUS attribute you want to use does not appear in Ignition Server's default list of RADIUS attributes (you can view this list in the Vendors panel), create it using these steps. Note that the following steps apply only to standard RADIUS attributes. If you want to create a new vendor-specific attribute, see "Adding a New VSA" on page 276 instead.

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. In the left panel of the Vendors panel, double-click **RADIUS,** then click **RADIUS Attribute Definitions** to display the list of attributes.



3. At the bottom of the window, click **New**.
4. In the RADIUS Attribute Definition window, define the attribute:



   - Enter the **Name** without spaces. This is the RADIUS attribute name and must match the name used by your networking equipment.
   - Enter its **Attribute Type** as an integer between 1 and 255. This is the code number set forth for the attribute in the RADIUS specification, or by relevant industry standards.
   - Choose its **Data Type**.
   - Click **OK**.

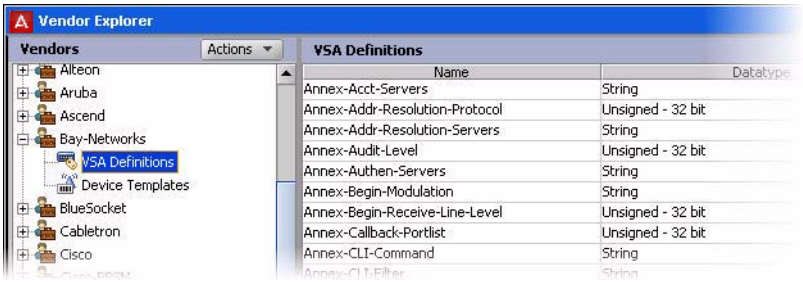
Next steps: You can use this attribute in one of these ways:

- To send provisioning values to your authenticator/switch via this RADIUS attribute, turn to "Outbound Attributes" on page 255.
- To evaluate the value of this RADIUS attribute in your authorization rules, turn next to "Inbound Attributes" on page 265.

Listing Ignition Server's Set of Available VSA Attributes

To view a list of the vendor-specific RADIUS attributes ("VSAs") defined in Ignition Server:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. In the left panel of the Vendors panel, double-click the manufacturer name of your network equipment, then click **VSA Definitions** to display the list of VSAs. Click on the column headings to sort the attribute list.



The **Name** is the RADIUS attribute name that Ignition Server and your network equipment use to identify this attribute. The **Data Type** indicates what kind of data the attribute can contain, such as a string or an unsigned 32-bit integer. The **Attribute Type** is the integer code that designates the attribute, as specified in your network equipment's documentation, or in the relevant industry standards document.

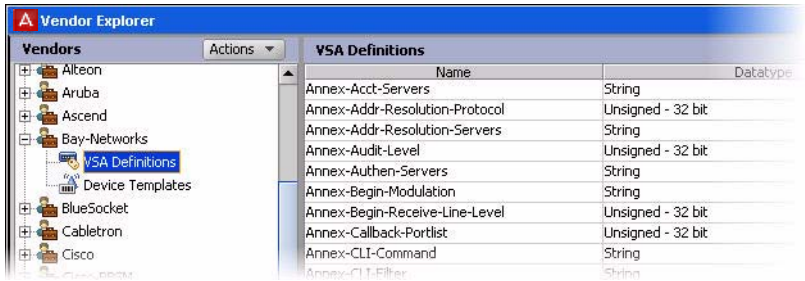
To add a new VSA, see "Adding a New VSA", below.

Adding a New VSA

If the vendor-specific RADIUS attribute you want to use does not appear in Ignition Server's list of VSAs in the Vendors panel, create it using the following steps:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.

2. In the left panel of the Vendors panel double-click the manufacturer name of your network equipment, then click **VSA Definitions** to display the list of VSAs. If your equipment manufacturer does not appear in the list, see "Adding an Equipment Vendor" on page 277.



3. At the bottom of the Vendors panel window, click **New**.
4. In the RADIUS VSA Definition window, define the attribute.
   - Enter the **RADIUS VSA Name** without spaces. This is the RADIUS attribute name and must match the name used by your networking equipment.
   - Enter its **Attribute Type** as an integer between 1 and 255. This is the code number set forth for the attribute in your network equipment's documentation, or by the relevant industry standards document.
   - Choose its **Data Type**.
   - Click **OK**.

Next steps: You use this attribute in one of these ways:

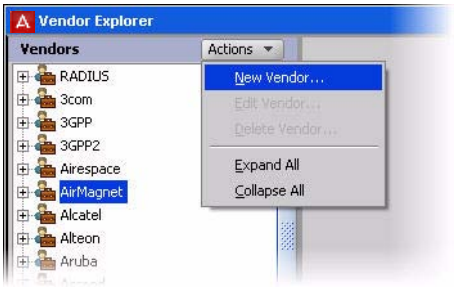
- To send provisioning values to your switch via this attribute, turn to "Outbound Attributes" on page 255.
- To evaluate the value of this attribute in your access rules, turn next to "Inbound Attributes" on page 265.

Adding an Equipment Vendor

Right out of the box, Ignition Server is configured with device templates and VSA definitions for a number of popular authenticator types. If your equipment vendor does not appear in the list, add it as follows:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.

2. In the Vendors panel, select **Actions: New Vendor**.



3. In the New Vendor window:
   - Type the manufacturer's name in the **Vendor Name** field, as a string without spaces.
   - Type the manufacturer's IANA private enterprise number in the **Vendor ID** field as an integer without leading zeros.

     See **http://www.iana.org/cgi-bin/enterprise.pl** for details.
   - Click **OK**.



Next steps: Your vendor record has been created and appears in the Vendors panel as shown below. To create VSA definitions for the equipment, see "Adding a New VSA" on page 276. To create a device template for the equipment, see "Device Templates" on page 270.



Provisioning FAQ

What if I have a switch that expects its provisioning data in a RADIUS attribute that's different from the attribute used by my other switches?

Answer: You can create an override that specifies a special RADIUS attribute to be used for certain switches. See "Overriding the Outbound Attribute Type for One or More Authenticators" on page 257.

How can I find out the datatype of an inbound or outbound attribute?

Answer: Do the following:

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Inbound Attributes** or **Outbound Attributes**.
2. In the Inbound or Outbound Attributes panel, scroll to find your attribute. Make a note of the values shown in the **Vendor** and **Mapping** columns. The **Mapping** column shows the RADIUS attribute name. If the attribute is a VSA, the **Vendor** column shows the equipment manufacturer that uses this VSA.

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. Do one of the following:
   - If the attribute is a standard RADIUS attribute, go to the top of the navigation tree on the left, double-click **RADIUS**, then click **RADIUS Attribute Definitions**.
   - If the attribute is a VSA, scroll to find the **Vendor** name of your attribute, double-click its name, then click **VSA Definitions**.
3. Scroll to find your attribute. The **Data Type** column shows the datatype.

Client Posture Policy

Avaya Identity Engines Ignition Server can require that the health and security of an end-user's computer be checked before it is allowed it to connect to the network. This type of checking is referred to as "posture checking." In conjunction with Microsoft Network Access Protection (NAP), Ignition Server can also remedy (or "remediate") certain out-of-compliance conditions on the user's computer.

A System Health Validation (SHV) is part of the NAP integration introduced in in Release 7.0. The Ignition Server acts as a NPS server where the IDE performs local validation of Statement of Health (SOH). Evaluation enables the Ignition server to know whether the end user system is compliant or non-compliant; the policy will have options to determine course of action depending the result. The SHV extends the existing Posture Profile to include NAP specific validation APIs and members that corresponds to SOH attributes for Firewall, AntiVirus, AntiSpam, System Auto Update and Security updates.

Contents

- "How Ignition Server Checks Client Posture" on page 281
- "Enabling NAP on a Windows machine" on page 282
- "Configuring NAP posture profiles" on page 284

How Ignition Server Checks Client Posture

When a CHECK POSTURE action is triggered in your authorization policy (see "How Ignition Server Evaluates a User Authorization Policy" on page 229), Ignition Server compares the user's machine's security posture with the requirements listed in your Ignition Server posture policy. The posture policy is your set of client-side security and machine-health requirements. It defines what firewall, anti-virus, and anti-spyware software must be installed, how up-to-date this software must be, and what to do if one of the required items is missing or out of compliance. (Note! You can use Windows native NAP supplicant for NAP based posture checking. If you do not want to check posture, then virtually any 802.1X supplicant can be used.)

The result of the posture check might be COMPLIANT, NON-COMPLIANT, or NO POSTURE (meaning the client machine did not return the requested posture data).

- A machine deemed COMPLIANT is given an automatic ALLOW action and you can optionally set Ignition Server to send provisioning values (for example, a VLAN assignment) in the RADIUS response.
- A machine deemed NON-COMPLIANT is given an ALLOW or DENY based on your policy, and you can optionally set Ignition Server to send provisioning values (for example, assigning the user to a quarantine/ remediation VLAN) in the RADIUS response.
- A machine deemed NO POSTURE is given an ALLOW or DENY based on your policy, and you can optionally set Ignition Server to send provisioning values (for example, assigning the user to a quarantine/remediation VLAN) in the RADIUS response.

Enabling NAP on a Windows machine

To enable NAP on a Windows machine, complete the following steps:

- Enable NAP services on the client
- Enable Enforcement on the client
- Configure authentication methods

Enable NAP services on the client

To enable NAP services on the client PC:

1. Click **Start**, click **Run**, type **services.msc**, and then press **ENTER**.
2. In the services window, confirm that these two services are running:
   - WiredAutoConfig
   - NetworkAccessProtection

   If you are using Wireless, make sure WLANAutoConfig service is started. If you are using Wireless with Windows XP, the service name will be **WirelessZeroConfig**.
3. Close the services window.

Enable Enforcement on the client

To enable EC settings on the client PC

1. Click Start, click **Run**, type **cmd**, and then press **ENTER**.
2. In the command window, type **netsh nap client show configuration**, and then press **ENTER**.
3. In the command output, under **Enforcement** clients, verify that the **Admin** status of the **EAP Quarantine Enforcement Client** is **Enabled**.

   To enable the **EAP Quarantine Enforcement Client**, type **netsh nap client set enforcement ID = 79623 ADMIN = "ENABLE"**, and then press **ENTER**.

4. In the command window, type **netsh nap client show state**, and then press **ENTER**.
5. In the command output, under **Enforcement client state**, verify that the Initialized status of the **EAP Quarantine Enforcement Client** is **Yes**.
6. Close the command window.

Configure authentication methods

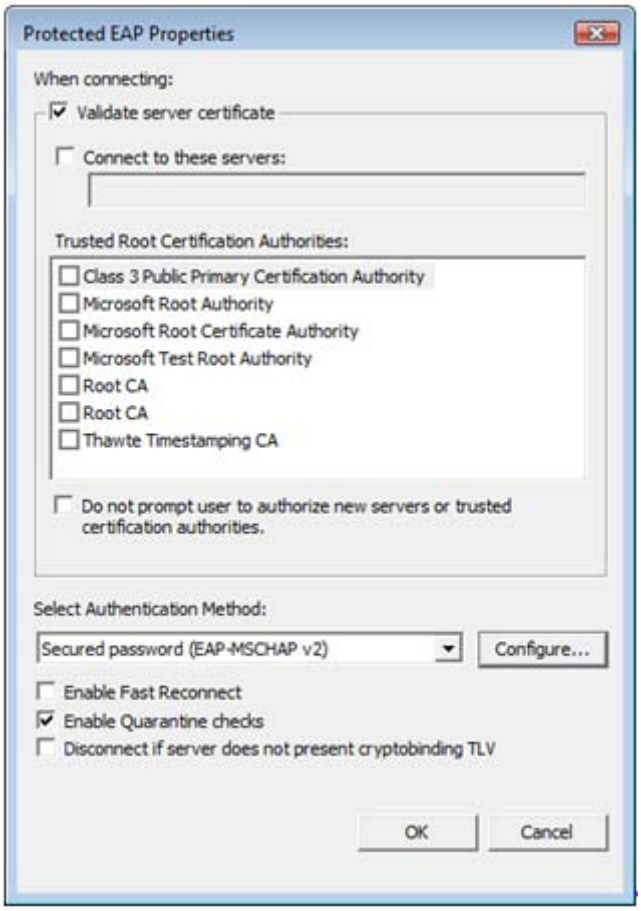
NAP health checks must be enabled in authentication methods of the local area connection.

To configure authentication methods

1. Click **Start**, click **Run**, and then type **ncpa.cpl.**
2. Right-click **Local Area Connection**, and then click **Properties**.
3. Click the **Authentication** tab, and verify that **Enable IEEE 802.1X** authentication is selected.
4. Click **Settings**.
5. In the **Protected EAP Properties** dialog box, clear the **Enable Fast Reconnect** check box, and verify that only the following check boxes are selected, as shown in the following example:
   - Validate server certificate
   - Enable Quarantine checks



If you are running Windows 7, this check box is called Enforce Network Access Protection.



6. Click **Configure**, verify that Automatically use my Windows logon name and password (and domain if any) is selected, and then click **OK**.
7. Click **OK**, and then click **OK** again.

Configuring NAP posture profiles

Microsoft Network Access Protection (NAP), introduced with Windows Vista and Windows Server 2008 is a new set of operating system components that provide a platform for protected access to private networks. The NAP platform integrates a way of detecting the health state of a client device that is attempting to connect to a network and restricting the access until the policy requirements for connecting to the network have been met.

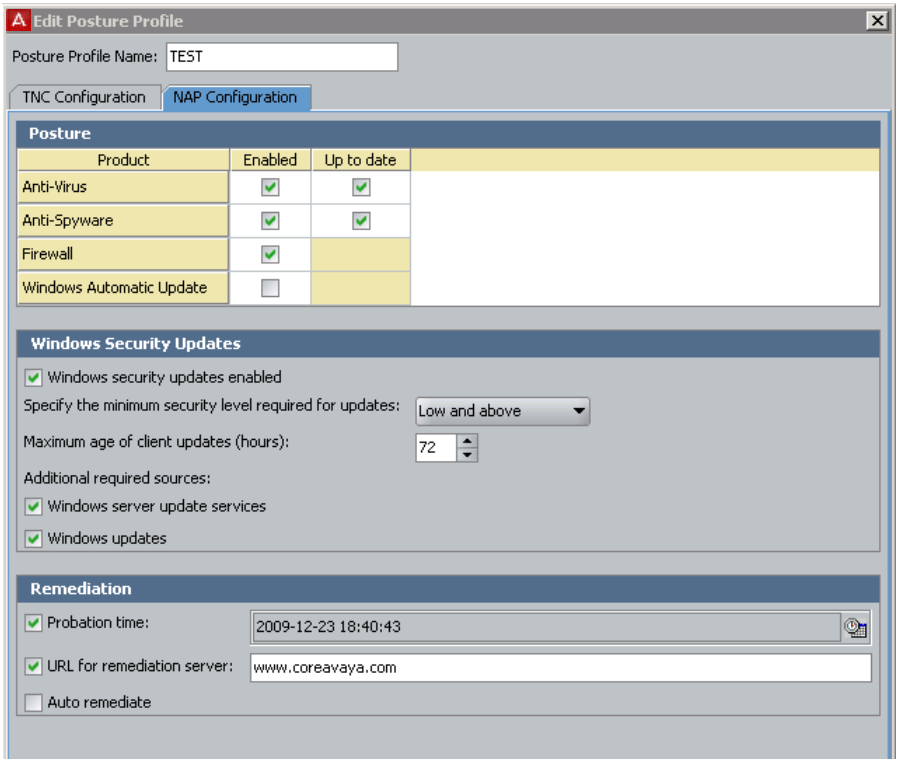
AIEIS combines both the Authenticated Network Architecture (ANA) solution from Avaya and Microsoft's Network Access Protection (NAP). This architecture allows you to enforce security policies for network access using Ignition's policy engine and NAP together, leveraging the strengths of both products.

AIEIS supports deployment of NAP clients without the need for a Microsoft NPS server.

To configure NAP posture policy, do the following:

1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click Posture Profile then select a posture profile from the list.
2. In the Content Area, click NAP Configuration.

Figure 65 NAP Configuration



3. In the Posture section, select **Enabled** for each product that you want active within NAP. Selecting **Up to date** allows for automatic or scheduled

Up to date is not applicable to Firewall or Windows Automatic Updates.

updates to run on the selected product.

4. In the Windows Security Updates, select the Windows security updated enabled box. Then set the level of security you want to update:
   - Unspecified (Default)
   - Low and above

- Moderate and above
- Important and above
- Critical only

5. Select the Additional required resources for Windows server updates services and Windows updates.
6. From the Remediation section, select the probation time to when you want NAP to connect to the remediation server for patches and updates. This field can not be edited manually. Click the up/down arrow key using your mouse to change the hour value.
   - Minimum: 1 hour
   - Maximum: 72 hours (Default)
7. Enter the remediation server URL.
8. Select Auto remediation if you want automatic updates.

VLAN Assignment

This chapter explains how to set up VLAN provisioning in Avaya Identity Engines Ignition Server. VLAN provisioning uses Ignition Server's outbound attribute functionality to send RADIUS attributes to the authenticator in order to assign the user to a VLAN. For information on more general uses of outbound attributes, see Chapter , "Provisioning Policy".

Contents

This chapter is made up of one main section with three subsections:

- "Creating a Policy That Assigns Users to VLANs" on page 287
  - "Create the Outbound Attribute" on page 288
  - "Create an Outbound Value for Each VLAN" on page 289
  - "Create VLAN Provisioning Rules" on page 290

Note! If you want to assign client devices to VLANs and your device-to-VLAN map does not change often, there is an alternative approach that lets you specify a VLAN assignment in each device record. See "VLAN Assignment Using the Device Record VLAN Fields" on page 330.

Creating a Policy That Assigns Users to VLANs

Setting up VLAN provisioning in Ignition Server requires two or three steps:

1. *Most deployments do not require this step!* If your authenticator cannot accept VLAN assignments via the Tunnel-Private-Group-Id attribute, then you must create an Ignition Server *outbound attribute*. This is the RADIUS attribute that carries the VLAN assignment to your VLAN equipment. Think of it as a container that is keyed to your specific make and model of VLAN equipment. If you use more than one type of VLAN concentrator, you might need more than one outbound attribute.
2. Create an Ignition Server *outbound value*. This is an outbound attribute with the name and id number of a specific VLAN. (The VLAN must exist on your VLAN switch.) You must create one of these for each VLAN on your network.
3. Create an Ignition Server *user authorization policy* or *MAC authorization policy* that contains a provisioning rule to assign the appropriate VLAN.

Create the Outbound Attribute

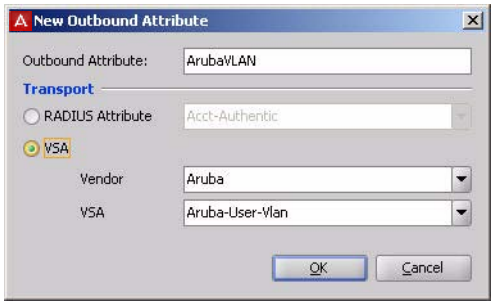
The outbound attribute you use depends on your authenticator:

- If your authenticator accepts VLAN assignments via the Tunnel-Private-Group-Id RADIUS attribute, you can use Ignition Server's predefined outbound attribute, "VLAN." Skip this section and turn to "Create an Outbound Value for Each VLAN" on page 289.
- If your authenticator requires that you use a different RADIUS attribute for VLAN assignments, follow the instructions below.

Setting Up VLAN Provisioning Using Nonstandard RADIUS Attributes

Use the following steps to set up VLAN provisioning for VLAN equipment that *does not* accept assignments via the Tunnel-Private-Group-Id attribute. (Note! If your authenticator uses Tunnel-Private-Group-Id, you can probably skip this step. Instead, use Ignition Server's predefined outbound attribute, "VLAN," as explained in "Create an Outbound Value for Each VLAN" on page 289.)

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Outbound Attributes**.
2. In the Outbound Attributes panel, click **New**.
3. This step creates a new outbound attribute. In the New Outbound Attribute window, type a name for your attribute in the **Outbound Attribute** field. Bear in mind that the attribute is a *container for VLAN assignments* and not a specific VLAN assignment, so it makes more sense to name it, for example, "ArubaVLAN" than, say, "VLAN-7".



4. In the Transport section, choose your authenticator manufacturer from the **Vendor** list and choose the attribute name from the **VSA** list. Note:
   - If the **Vendor** list does not include the name of your authenticator manufacturer, see "Adding an Equipment Vendor" on page 277.
   - If the **VSA** list does not include the VLAN attribute name required by your authenticator, see "Adding a New VSA" on page 276.
5. Click **OK** to save the outbound attribute.

Create an Outbound Value for Each VLAN

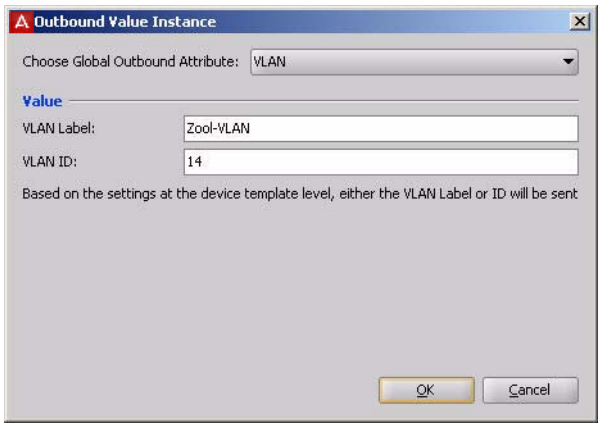
Before you can write a VLAN provisioning rule, you must save the VLAN ID or name as an Ignition Server *outbound value*. The outbound value is sent in the RADIUS message to the authenticator. Follow the steps below to create an outbound value:

Before you begin: Before you start creating outbound values, log into your VLAN switch as administrator and make a note of the VLAN label (a string) and VLAN ID (an integer) of each VLAN to which you plan to assign users.

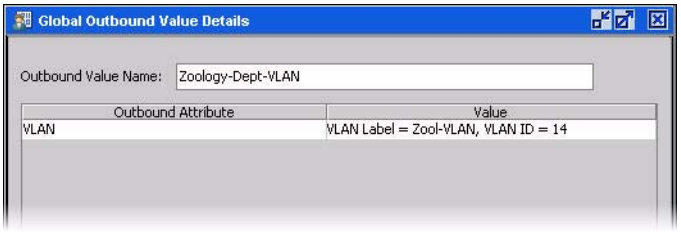
To create each outbound value:

1. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**.
2. In the Outbound Values panel, click **New**.
3. In the Outbound Value Details window, type a name for the outbound value in the **Outbound Value Name** field. This name should include the name of the VLAN, because this is the name used when setting up the VLAN assignment in your provisioning rule. For example, when setting up a VLAN for his university's zoology department, the administrator might call the outbound value "Zoology-Dept-VLAN."
4. At the bottom of the Outbound Value Details window, click **New**.
5. In the Outbound Value Instance window, do the following:
   - In the drop-down list at the top of the window, pick the name of your outbound attribute. If your authenticator uses the standard Tunnel-Private-Group-Id attribute, choose **VLAN**. Otherwise, choose the outbound attribute you created in Step 3 in "Setting Up VLAN Provisioning Using Nonstandard RADIUS Attributes" on page 288.
   - In the **VLAN Label** field, type the name your authenticator uses to refer to the VLAN. For many authenticators, the label is case sensitive. Enter the label exactly as it appears in your authenticator-resident VLAN configuration. For some authenticator types (those that use only a VLAN ID), this is optional.
   - In the **VLAN ID** field, type the integer ID number your authenticator uses to refer to the VLAN.

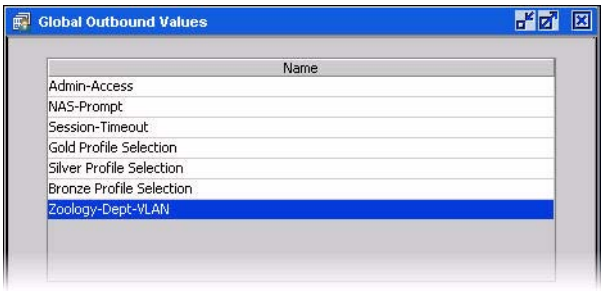
- Click **OK**.



6. Your VLAN value appears in the Outbound Value Details window. Click **OK** to save the value.



7. Your saved outbound value appears in the Outbound Values list. Click **New** if you want to set up outbound values for more VLANs.



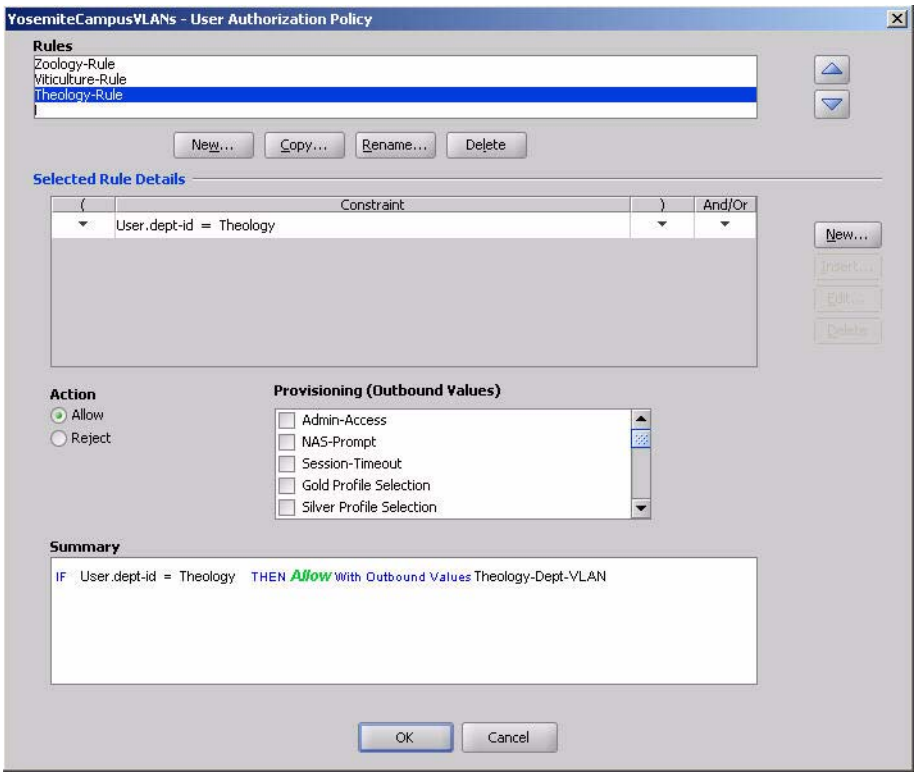
After you have created an outbound value for each VLAN, you must set up Ignition Server provisioning rules to assign users to VLANs.

Create VLAN Provisioning Rules

At runtime, Ignition Server evaluates its set of *provisioning rules* in order to determine which VLAN it assigns the user or device to.

Follow the instructions below to create provisioning rules that assign users to VLANs. (While these instructions cover *user* VLAN assignment only, Ignition Server is also capable of *device* VLAN assignment. See Chapter , “MAC Authentication”.)

1. In Dashboard’s Configuration hierarchy tree, expand Access Policies and expand RADIUS.
2. Scroll down the access policy tree to find the access policy that contains your authenticators (your VLAN equipment). This example uses an access policy called “YosemiteCampusVLANs.” Click the name of your access policy. Click the Authorization tab and click Edit.
3. Your provisioning rules are part of your user authorization policy and/or MAC authorization policy. In this example, we use a user authorization policy. In the Access Policy panel of Dashboard, click the Authorization Policy tab and click the **Edit** button.
4. The Edit Authorization Policy window or MAC Authorization Policy window sets the conditions that determine whether the user or device is granted access and the conditions that determine which VLAN is used. Each rule in this window can act as both an authorization and provisioning rule. In this example, we refer to them as provisioning rules, since we are concerned here with provisioning.



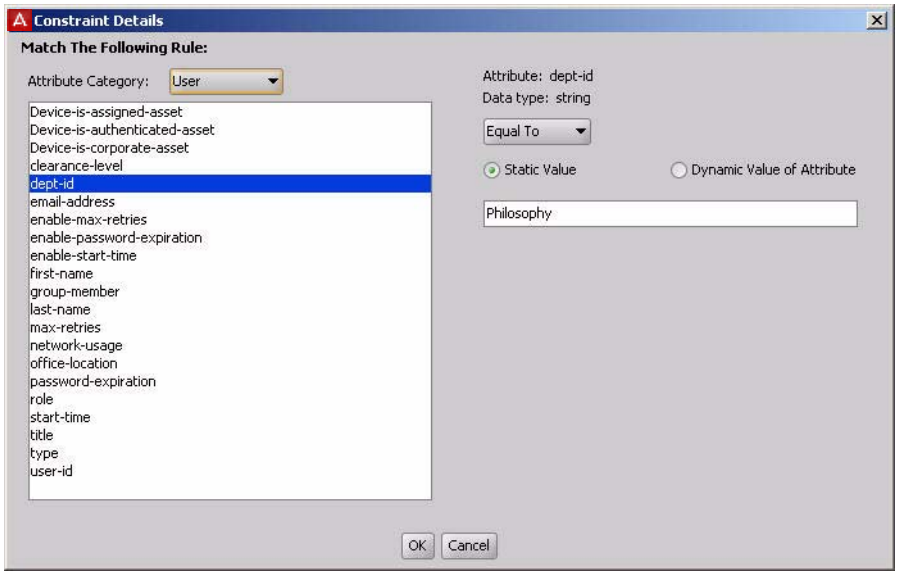
In this example, we have already defined provisioning rules for the Theology, Viticulture, and Zoology departments. In the steps below, we'll add a rule for the Philosophy department.

5. On the left side of the Edit Authorization Policy window, click the **Add** button just below the **Rules** list.
6. In the New Rule window, type a name for your provisioning rule and click **OK**.

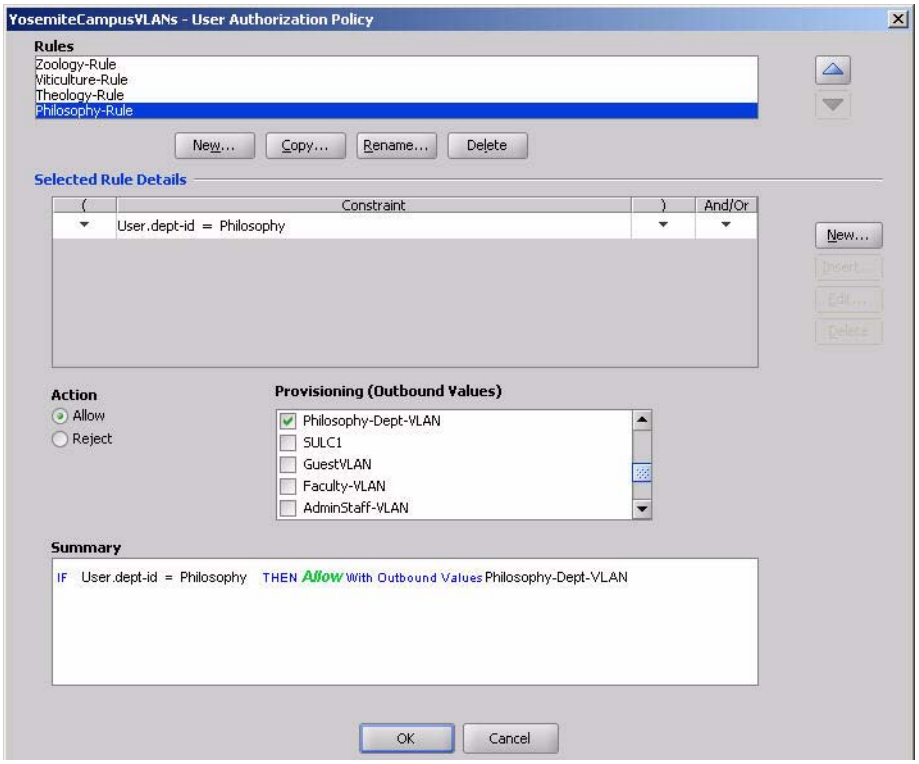


7. Your new rule now appears in the Rules list in the Edit Authorization Policy window. Click on its name in the Rules list to edit it. The Selected Rule Details section (the lower part of the window) allows you to edit the rule.
8. Click the **New** button next to the Constraint list.
9. The Constraint Details window defines the condition that must be fulfilled in order to provision the user. Typically, your constraint evaluates a user attribute (see "Attributes Used in Rule Constraints" on page 232). For this example, we choose a **User Attribute** called "dept-id" and set the test value to "Philosophy." To do this:
   - Select the **Attribute Category**, "User".
   - Click the attribute, "dept-id". (If you do not see this attribute in the list, you must create it. See "Adding a New User Virtual Attribute" on page 205.)
   - Select **Equal To**
   - Tick the **Static Value** checkbox.
   - In the text field, type "Philosophy."
   - Click **OK**.

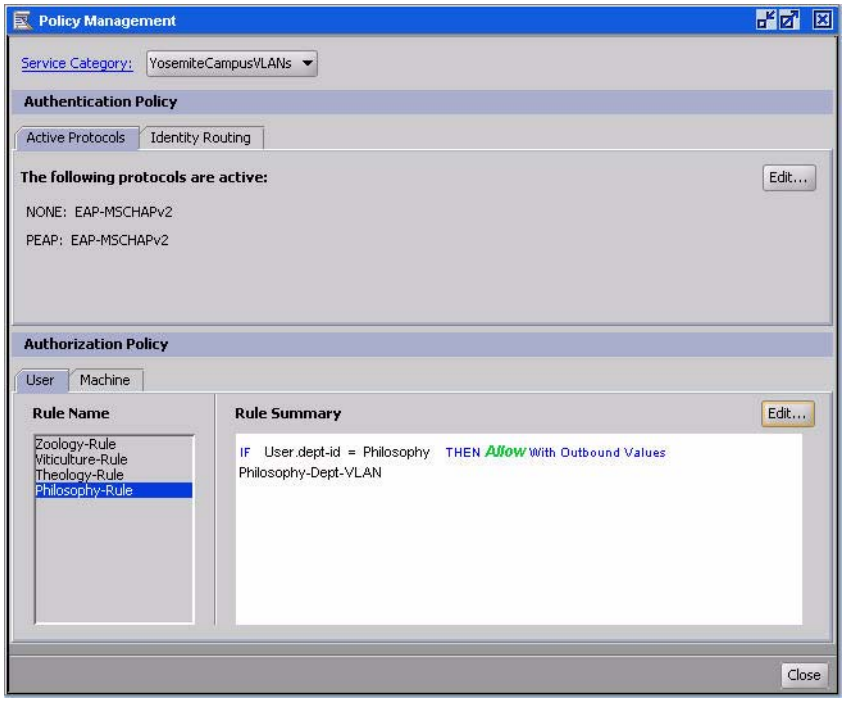
To evaluate this constraint, Ignition Server checks if the user record is a member of the “Philosophy” department.



10. In the Edit Authorization Policy window, with your rule (“Philosophy-Rule” in this example) highlighted, click “Allow” in the **Action** section, and in the **Provisioning** section click the checkbox next to the outbound value for the appropriate VLAN. In this example, we use the outbound value we created in Step 6 in “Create an Outbound Value for Each VLAN” on page 289.



11. Repeat Step 5 through Step 10 to add more VLAN provisioning rules.
12. Click **OK** to save the rule sequence and exit the Edit Authorization Policy window. In the lower part of the Access Policy panel of Dashboard, the Authorization Policy tab contains a summary of your provisioning rules. Click on each rule to see its Rule Summary.



Windows Machine Authentication

This chapter explains Avaya Identity Engines Ignition Server support for Windows machine authentication, which you use in an Active Directory-based environment to ensure that only known, approved Windows clients can connect to the network.

Contents

- "Introduction to Windows Machine Authentication" on page 295
- "Setting up Microsoft Windows Machine Authentication" on page 298

Introduction to Windows Machine Authentication

Microsoft Windows machine authentication allows you to force networked Windows-based devices — instead of users, or in addition to users — to authenticate in order to connect to the network. The device authenticates using its machine credentials, which are compared with those stored in the device's account record in Active Directory.

After a device authenticates, Ignition Server shows its current authentication in the **Monitor: Current Site: Learned Devices** panel until that authentication expires. See "Learned Devices Tab" on page 447.

Do not confuse Microsoft Windows machine authentication with the more generic *MAC authentication*. (See "Introduction to MAC Authentication" on page 321 for details.)

Supported Authentication Protocols

When Ignition Server performs Microsoft Windows machine authentication, the following authentication types are supported:

- EAP-TLS
- PEAP / EAP-TLS
- PEAP / EAP-MSCHAPv2

In each case, Ignition Server checks that the machine's credentials match those of the corresponding device record in the data store. Typically, this record is a "Computer" entry in Active Directory.

Session Behavior for Windows Machine Authentication

The session behavior for Microsoft Windows machine authentication is as follows:

- If the user's computer is configured to do machine authentication, when it boots up it automatically attempts an 802.1X authentication using its machine credentials.
  - If authentication succeeds, then based on the authorization policy in Ignition Server (which in turn can be based on attributes saved as part of the machine's entry in AD), the computer is placed into an appropriate VLAN.
  - If authentication fails, the computer is not allowed to join the network.
- If a user later logs into the domain using this computer, the computer attempts another 802.1X authentication with the user's credentials. You can set a rule in Ignition Server that only permits the user to log in on a computer that has successfully performed Windows machine authentication. See "Asset Correlation" on page 335.
  - As with the machine authorization, if the user authentication succeeds, then a new VLAN assignment can be made. Based on the authorization policy in Ignition Server (based on attributes of the user's record, based on his or her group membership in AD or based on other data you specify), the user is put into an appropriate VLAN.
- When the user logs out, the computer does another 802.1X authentication using its machine credentials, and Ignition Server once again assigns the computer to the VLAN you have designated to handle sessions for which only Windows machine authentication has occurred and no user authentication has occurred.

As the Ignition Server Administrator, you can view the currently logged in devices that have signed on using machine authentication. To do this, click **Monitor** in the Dashboard main window. The **Learned Devices** tab shows recent, successfully logged-in devices that authenticated via Windows machine authentication.

NAP support for Peap

Nap support for Peap involves additional new EAP methods during phase 2 of PEAP where an EAP session establishes between the EAP peer and the EAP server, encapsulated in the TLS tunnel established in phase 1. You use the following three methods to facilitate the exchange of TLVs between a PEAP peer and a PEAP server.

- EAP TLV Extensions Method (NAP support)
- SoH EAP Extensions Method (NAP support)

- Capabilities Negotiation Method (NAP support)

EAP TLV Extensions Method

You must have the Type field set to **33** to use the EAP TLV Extensions Method as the inner EAP method. It allows transmitting Cryptobinding TLV, Result TLV, and SoH Response TLV. Within an EAP TLV Extensions Method, you can send the Result TLV, Cryptobinding TLV, and SoH Response TLV in any order. The receiver MUST NOT assume any order of the TLVs.

- The cryptobinding TLV ensures that the EAP peer and the EAP server participate in both the inner and the outer EAP authentications of a PEAP authentication.
- The SoH Response TLV is a vendor TLV sent within a Microsoft vendor-specific TLV. Sent to the PEAP peer by the server, its ultimate recipient is the Statement of Health (SoH) entity, as specified in [MS-SOH], at the peer.
- The Result TLV is a TLV represents the status (success or failure) of the inner EAP method negotiation or to indicate the sender's consent (ability or inability) to participate in a fast-reconnect.

Capabilities Negotiation Method

You use this method to exchange various capabilities supported by Peer to server. PeapP2StartState is modified to send the request for Capabilities. WaitForCapabilitiesState are added to process the Capabilities negotiation.

SOH EAP Extensions Method

You must send the SoH Request TLV and the SoH TLV within a SoH EAP Extensions Method.

- The SoH Request TLV is a vendor TLV sent within a Microsoft vendor-specific TLV in a SoH EAP Extensions Method request. Sent to the PEAP peer by the server, its purpose is to trigger transmission of an SoH message by the peer's Statement of Health for Network Access Protection Protocol [MS-SOH] entity.
- The SoH TLV is a vendor TLV sent within a Microsoft vendor-specific TLV in a SoH EAP Extensions Method response. Sent to the PEAP server by the PEAP peer, its ultimate recipient is the SoH validator at the server.

Setting up Microsoft Windows Machine Authentication

Create the Ignition Server authorization policy that handles machine authentication. This policy *must* include a rule that evaluates whether the connecting device is a recognized device in your AD, and it *should* include a rule that evaluates whether the connecting device is in a group that you trust.

Pick one of the following approaches:

1. “Machine Authentication Based on ObjectClass” on page 298, or
2. “Machine Authentication Based on OU, O, or Group Membership” on page 302

Machine Authentication Based on ObjectClass

You can set Ignition Server to enforce Microsoft Windows machine authentication by looking up the device in AD and checking the value of the device record’s ObjectClass attribute. If it finds ObjectClass set to “computer” it allows the device to connect and, if configured, carries out its provisioning rule (such as mapping the device to a VLAN). Follow the steps below to set this up.

Set your User Root DN

In your AD configuration, set your **User Root DN** to include all your authorized users and computers. For example if your users live under `“cn=users, ...”`, and computers are under `“cn=computers, ...”` then set your root to the DN above those containers. To do this:

1. In Dashboard’s Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, and click Directory Services.
2. Click the name of your AD service.
3. Click **Edit**.
4. In the Active Directory Details window, edit the **User Root DN** and click **Save**.

   For example, if the domain name is CORP.LOCAL and if all your authorized users and computers live under your root DN, in the Active Directory Details window you set **User Root DN** to `“dc=corp,dc=local”`.

Set up Ignition Server to retrieve the objectClass value

Create a user virtual attribute called “type” and map it to the AD attribute, objectClass.

1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click User Virtual Attributes.

2. Select **Actions: Add a New User Virtual Attribute**.
3. Give the attribute the name “type” and a data type of *Multi-valued string*, and click **OK**.
4. With the *type* attribute selected in the **Attributes** list on the left, go to the User Virtual Attribute Details panel on the right and click **Add**.
5. In the Map Attributes window, in the **Directory Service** field, pick the name of your AD service.
6. Make sure **Available attribute name** is ticked. This lets you choose from the list of attributes in your AD store.
7. In the list, find the attribute, “ObjectClass”, click it, and click **OK**.

Write your policy rule

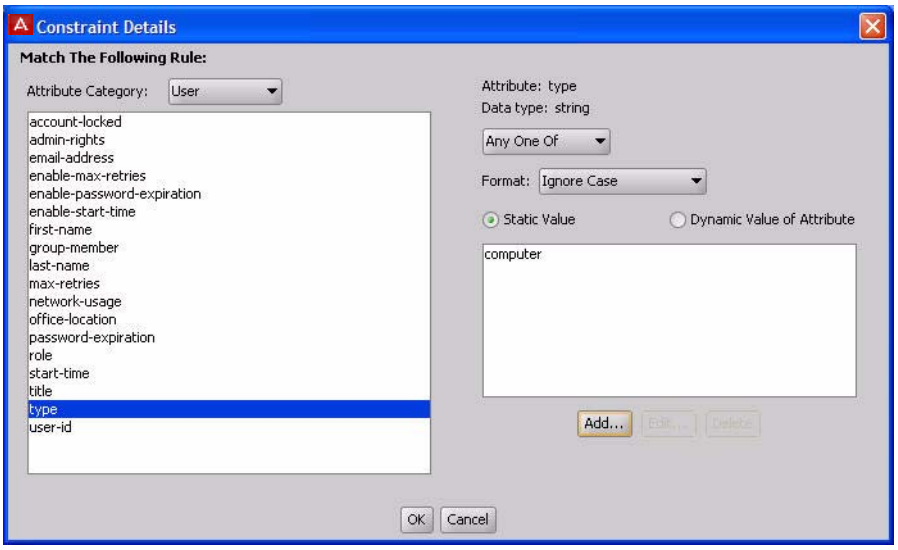
Write a policy rule that checks if the *type* attribute is set to “Computer” and, if so, carries out its machine authentication policy.

1. In Dashboard’s Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click the name of your access policy. Click the Authentication Policy tab and click Edit.
2. Make sure the Authentication Policy includes one or more of the supported authentication protocols:
   - EAP-TLS
   - PEAP / EAP-TLS
   - PEAP / EAP-MSCHAPv2

   Close the Edit Authentication Policy window.
3. Make sure the Identity Routing Policy includes your AD directory or directories.
4. In the Authorization Policy tab, click **Edit**.
5. In the **Rules** section of the Edit Authorization Policy window, click **New** to create a new rule. Give the rule a name like “Machine-Auth”.
6. Make sure your new rule is selected in the list on the left side of the window. In the **Selected Rule Details** section, click **New** to add a constraint.
7. In the Constraint Details window, do the following:
   - Select an **Attribute Category** of “User” (not “Device”).
   - Select the attribute named “type”.
   - On the right side, choose **Any One Of**.
   - Set **Format** to “Ignore Case”.

- Tick the **Static Value** button.
- Click the **Add** button.
- In the Add Value dialog, type “computer” and click **OK**. This is your comparison value.
- Click **OK** to close the Constraint Details window.

Figure 66 The Constraint Details window



Note: When you write this constraint you use a *User* attribute rather than a *Machine* attribute because in Ignition Server, *Machine* attributes are used only for MAC authentication, not for Windows machine authentication.

8. Add more constraints if needed. For example, if you want to further restrict access to only those devices that are in a certain OU or O in your AD tree, then add a virtual group in Ignition Server that maps to the desired OU’s or O’s, and create a rule that tests for membership in that virtual group.
9. Set the **Action** to “Allow”.
10. Select provisioning values if desired.
11. Click **OK** to save your policy.

Add user policies

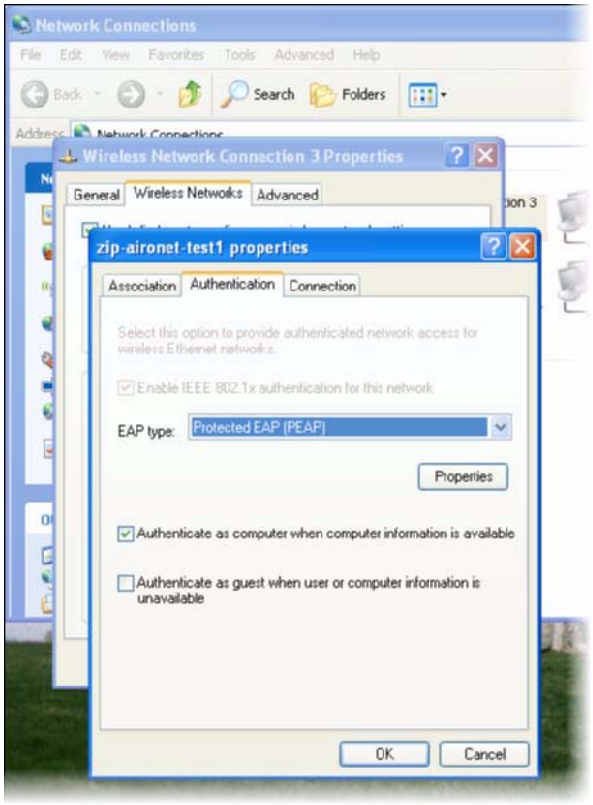
If you want to require that users log in only using Windows-authenticated machines, see “Requiring the User to Connect Using a Machine Authenticated-Device” on page 340.

If your policy needs rules to handle user authentication, return to the top of the window and click the **New** button create another rule. See “Creating a RADIUS User Authorization Policy” on page 242 for more details.

Set up your supplicants

Set up your supplicants to require machine authentication. Consult your supplicant documentation for instructions. For Microsoft Windows XP supplicants, use these steps:

1. Open the Network Connections window, and open the Properties window for the Interface your want to configure.
2. Click Properties to open the Properties window, and click the Authentication tab.



3. Tick the “Authenticate as computer...” checkbox
4. In the **EAP type** field, select “Protected EAP.”
5. Click **OK** to exit the configuration windows.
6. Make sure you have installed the required certificates on the Windows XP machine to support authentication.

Machine Authentication Based on OU, O, or Group Membership

You can set Ignition Server to enforce Microsoft Windows machine authentication such that each device is allowed to connect to the network upon start-up only if it has an entry in AD indicating it is permitted to connect. (This is an alternative to the less strict approach explained in the section, “Machine Authentication Based on ObjectClass” on page 298.)

To do this, place your device entries in an OU, O, or group in AD, and set Ignition Server to allow access based on membership in that group. Follow the steps below to set this up.

Prepare your entries in AD

In AD, place the entries for the authorized devices in an OU, O, or group in the tree. The following steps describe how to write a policy that grants network access to all devices in that OU, O, or group.

Set your user root DN

In your AD configuration, set your **User Root DN** to include all your authorized users and computers. For example if your users live under “`cn=users, ...`”, and computers are under “`cn=computers, ...`” then set your root to the DN above those containers. To do this:

1. In Dashboard’s Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, and click Directory Services.
2. Click the name of your AD service.
3. Click **Edit**.
4. In the Active Directory Details window, edit the **User Root DN** and click **Save**.

   For example, if the domain name is CORP.LOCAL, in the Active Directory Details window you would set **User Root DN** to “`dc=corp,dc=local`”.

Set Ignition Server to retrieve the group membership information

Create a virtual group called, for example, “ok-devices” and map it to the OU/O organizational units or groups in AD whose devices you want to be allowed to connect.

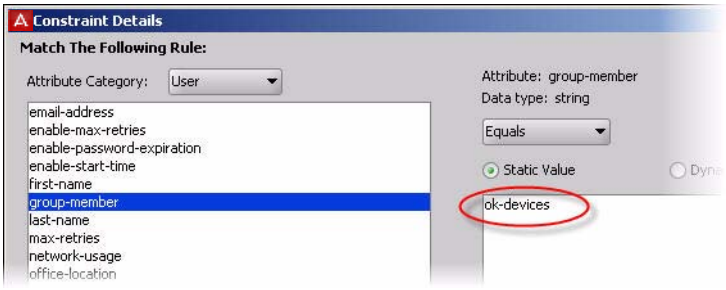
1. In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click Virtual Groups.
2. Select **Actions: Add a New Virtual Group**.
3. Give the attribute a name and click **OK**. In this example, we call it “ok-devices”.

4. With the “ok-devices” group selected in the Virtual Groups list on the left, go to the Virtual Group Details panel and click **Add**.
5. In the Map Groups window, in the Directory Service field, pick the name of your AD. Find the OU, O, or group whose devices you want to allow, click to highlight it, and click **OK**.
6. To add more authorized OU’s, O’s or groups, click **Add** again and repeat the preceding step.

Write your policy rule

Write a policy rule that allows access if the device is in an authorized OU, O, or group.

1. In Dashboard’s Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click the name of your access policy. Click the Authentication Policy tab and click Edit.
2. Make sure the Authentication Policy includes one or more of the supported authentication protocols:
   - EAP-TLS
   - PEAP / EAP-TLS
   - PEAP / EAP-MSCHAPv2
3. Make sure the Identity Routing Policy includes your AD.
4. In the Authorization tab click **Edit**.
5. Near the top of the Edit Authorization Policy window, click **New** to create a new rule. Give the rule a name like “Machine-Auth”.
6. Make sure your rule is selected in the **Rules** list on the left side of the window. In the **Selected Rule Details** section, click **New** to add a constraint.
7. In the Constraint Details window, select **User Attributes** and select “group-member”. In the Phrase section, set the drop down to Equals and type “ok-devices” as the test value. Click **OK**.



8. Add more constraints if needed.

9. Set the **Action** to “Allow”.
10. Add provisioning values if desired.
11. Click **OK**.

Add user policies

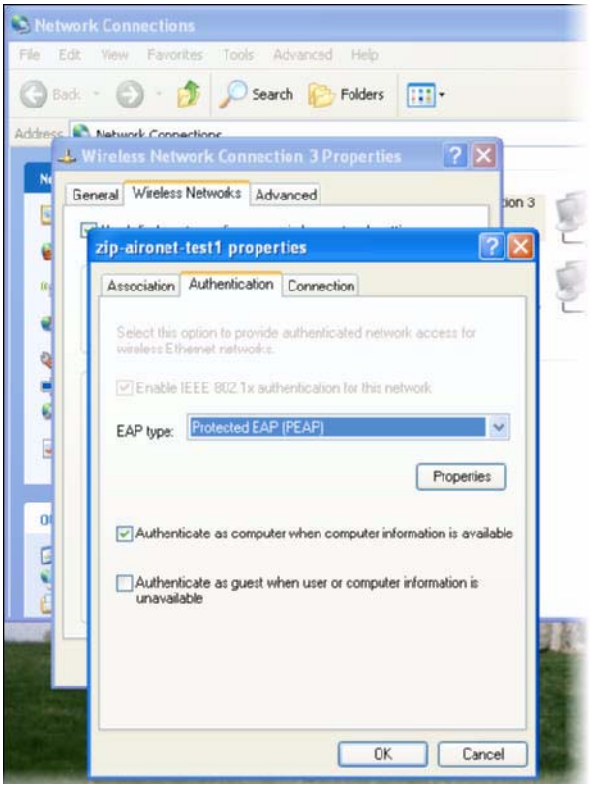
If you want to require that users log in only using Windows-authenticated machines, see “Requiring the User to Connect Using a Machine Authenticated-Device” on page 340.

If your policy needs rules to handle user authentication, return to the top of the window and click the **New** button create another rule. See “Creating a RADIUS User Authorization Policy” on page 242 for more details.

Set up your supplicants

Set up your supplicants to require machine authentication. Consult your supplicant documentation for instructions. For Microsoft Windows XP supplicants, use these steps:

1. Open the **Network Connections** window, and open the Properties window for the Interface your want to configure.
2. Click **Properties** to open the Properties window, and click the **Authentication** Tab.

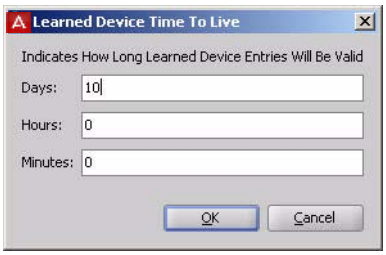


3. Tick the “Authenticate as computer...” checkbox
4. In the **EAP type** field, select “Protected EAP.”
5. Click OK to exit the configuration windows.
6. Make sure you have installed the required certificates on the Windows XP machine to support authentication.

Setting TTL for Windows Machine Authentication

The **Learned Device Time To Live** window establishes the time to live (“TTL”) for client machine authentications done via Windows machine authentication. If you have imposed an asset correlation policy (see “Requiring the User to Connect Using a Machine Authenticated-Device” on page 340), then a user’s machine must have a current machine authentication in order for the user to log in. To view the current machine authentications and the expiration time of each, see “Learned Devices Tab” on page 447. To set the TTL, see the instructions below.

Figure 67 Learned Device Time to Live (TTL) window



To set the TTL:

1. From the Dashboard main window, go to the Configuration Hierarchy tree.
2. Right-click on your site and choose **Learned Device Time To Live**.
3. In the **Learned Device Time To Live** window, type the TTL in days, hours, and minutes.
4. Click **OK** to save the setting.

TACACS+ Authorization

This section shows how to configure Avaya Identity Engines Ignition Server as your TACACS+ administrator access control server.

Contents

- "Introduction to TACACS+ Access Control" on page 307
- "Installing Your TACACS+ License" on page 309
- "Turning on the Ignition Server TACACS+ Service" on page 310
- "Creating a Command Set" on page 311
- "Viewing or Editing a Command Set" on page 313
- "Creating a TACACS+ Access Policy" on page 313
- "Enable Your Devices for TACACS+ Authorization" on page 317
- "The TACACS+ Global Authenticator" on page 318

Introduction to TACACS+ Access Control

TACACS+ policies allow the Ignition Server to function as the TACACS+ Server (policy decision point) that permits or denies administrator access to equipment on your network. When an administrator attempts to log in to a network device, the device sends a TACACS+ authentication request to the Avaya Identity Engines Ignition Server, which authenticates the administrator, applies the authorization policy, responds with an allow or deny decision, and logs the action in its TACACS+ access log.

There are two approaches to enforcing TACACS+ controls: *privilege-level authorization* and *per-command authorization*:

- With *privilege-level authorization*, the administrator is given a privilege level (1-15) upon logging in, and he can only use commands of that privilege level or lower. If the administrator wants to use a more sensitive command, he can type the *enable* command and authenticate again to a higher privilege level. For more details, see "Privilege-Level TACACS+ Authorization" on page 308.
- With *per-command authorization*, each time an administrator types a command, the equipment he's working on sends a TACACS+ authorization request to Ignition Server. Your TACACS+ policy prescribes

the set of allowed commands and arguments. For more details, see “Per-Command TACACS+ Authorization” on page 308.

In Ignition Server, you can combine elements of privilege-level authorization with elements of per-command authorization in a single TACACS+ policy.

Privilege-Level TACACS+ Authorization

With *privilege-level authorization*, every command on your equipment is assigned a privilege level. Privilege levels are numbered 1 through 15, with most widely-available commands usually given a level of 1, and the most sensitive commands given a level of 15.

When the administrator logs in, he has a privilege level of 1 and can only use commands of that privilege level. To use a more sensitive command, he types the *enable* command and authenticates again. Based on your Ignition Server TACACS+ policy, the administrator is granted or denied a higher privilege level.

If you use this type of authorization, your equipment sends an authentication request each time an administrator logs in and each time an administrator types an *enable* command to raise his privilege level.

Limitation to Note: Ignition Server does not support the use of a separate “enable password.” Instead, the administrator must re-type his administrator credentials in order to raise his privilege level.

Notes on Logging: When an administrator uses SSH or telnet to establish his administrator session, the Ignition Server access log shows not just an authentication event (as is typical at the start of an administrator session) but also an *authorization* event. This log entry corresponds to the authorization of the SSH or telnet session.

Per-Command TACACS+ Authorization

With *per-command authorization*, each time an administrator types a command, the equipment he’s working on sends a TACACS+ authorization request to Ignition Server. For each administrator’s session, the rules of your TACACS+ policy prescribe the sets of allowed commands and command arguments.

If you use per-command authorization, then your equipment sends an authentication request when the administrator logs in, and after that it sends an *authorization request only* each time he or she types a command.

Limitation to Note: Some TACACS+ server architectures allow you to split authentication from authorization, using one TACACS+ server for authentication and another for authorization. This is *not* permitted in Ignition Server. If you use Ignition Server for TACACS+ authorization, you *must* use Ignition Server for TACACS+ authentication.

Notes on Logging: Please note:

- When using per-command authorization, each authorization request generates an entry in the Ignition Server logs. Since only the initial request of a session generates an *authentication* request, all subsequent requests in the session will show up in the Ignition Server Access log as *authorization requests only*.
- When an administrator uses SSH or telnet to establish his administrator session, the Ignition Server access log shows not just an authentication event (as is typical at the start of an administrator session) but also an *authorization* event. This log entry corresponds to the authorization of the SSH or telnet session.

Getting Started

To set up TACACS+ authorization:

- **To perform first-time set-up of TACACS+** on your Ignition Server, turn to "Installing Your TACACS+ License" on page 309; or
- **To add new TACACS+ policies**, see one of the following:
  - if your TACACS+ policy uses *privilege-level authorization*, create your TACACS+ policy as shown in "Creating a TACACS+ Access Policy" on page 313; or
  - if your TACACS+ policy uses *per-command authorization*, create your sets of allowed commands as shown in "Creating a Command Set" on page 311.

Installing Your TACACS+ License

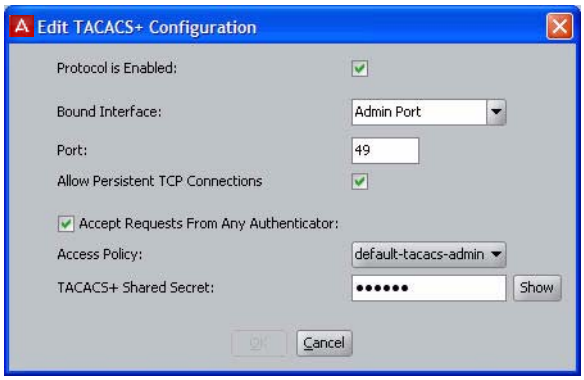
In the Configuration Hierarchy tree of Dashboard, expand the Access Policies node. You should see a node called, "TACACS+" there. If you do not see it, you must install your Ignition Server TACACS+ license. See "Installing an Ignition Server License" on page 69.

Next step: Activate the TACACS+ service as explained in "Turning on the Ignition Server TACACS+ Service" on page 310.

Turning on the Ignition Server TACACS+ Service

The Ignition Server TACACS+ service handles administrator authorization traffic. You can bind the Ignition Server TACACS+ service to a physical Ethernet port on the Ignition Server (the Admin port or Service Port A), or you can bind it to an Ignition Server VIP (VIPs are explained in “Managing Virtual Interfaces (VIPs)” on page 384). Use the TACACS+ tab to bind the TACACS+ service and set its port numbers.

Figure 68 Default Settings for the TACACS+ Protocol



Turning on the TACACS+ Service:

1. In the Dashboard main window, in the **Configuration Hierarchy** panel, click the name of your site (by default, “Site 0”).
2. In the Sites panel, click the **Services** tab and click the **TACACS+** tab.
3. Click the **Edit** button in the TACACS+ tab. The **Edit TACACS+ Configuration** dialog box appears:
4. Edit as necessary:
   - **Protocol Is Enabled:** Tick this checkbox to allow Ignition Server to handle TACACS+ traffic.
   - **Bound Interface:** From the drop down list, choose the Ignition Server Ethernet interface that is to handle TACACS+ traffic. You can bind TACACS+ to any port on the Ignition Server. If you are running an HA pair of Ignition Server s, you can choose to bind TACACS+ to a VIP interface. The VIP names are also listed in the drop down list. See “Managing Virtual Interfaces (VIPs)” on page 384 for details on using VIPs.
   - **Port:** Enter the TCP port number that you want to receive TACACS+ authentication requests. The default TACACS+ authentication port is 49.
   - **Allow Persistent TCP Connections:** With this checkbox checked, the Ignition Server allows each TACACS+ client to maintain its network

connection to Ignition Server's TACACS+ service after the initial authentication. This option is turned on by default.

- **Accept Requests From Any Authenticator:** Tick this checkbox if you want to create a global TACACS+ authenticator that sets policy for all authenticators that do not match a specific TACACS+-enabled authenticator in your Ignition Server configuration. When servicing a request, if Ignition Server finds a better matching TACACS+ authenticator record, it uses your policy associated with that record and does not fail over to the global authenticator. The global authenticator applies only when no better matching authenticator or bundle is found. See "The TACACS+ Global Authenticator" on page 318.
- **Access Policy:** This setting is used only in the case of a global TACACS+ authenticator. Choose your global TACACS+ policy that you want to be applied if no better-matching authenticator is found.
- **TACACS+ shared secret:** This setting is used only in the case of a global TACACS+ authenticator. Type the shared secret that an authenticator must present in order to have its TACACS+ requests handled according to the global TACACS+ authenticator policy.

Ignition Server enables the **OK** button. Click **OK** to apply your changes to the TACACS+ service.

Next steps: Do one of the following:

- If your TACACS+ policy uses *privilege-level authorization*, create your TACACS+ policy as shown in "Creating a TACACS+ Access Policy" on page 313.
- If your TACACS+ policy uses *per-command authorization*, create your sets of allowed commands as shown in "Creating a Command Set" on page 311.

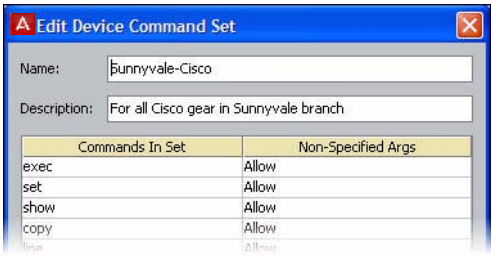
Creating a Command Set

To set up *per-command authorization* in Ignition Server, you create a TACACS+ policy that specifies the set of allowed commands and arguments for each type of administrator. Each TACACS+ policy consists of a set of rules, and each rule allows sets of commands based on evaluation of the identity of the administrator, the identity of the device being administered, and/or other attributes of the administrative transaction. Before you can write a rule, you must create the *command sets* that list the commands the rule will allow. A command set can be shared among many rules and policies.

To create a command set:

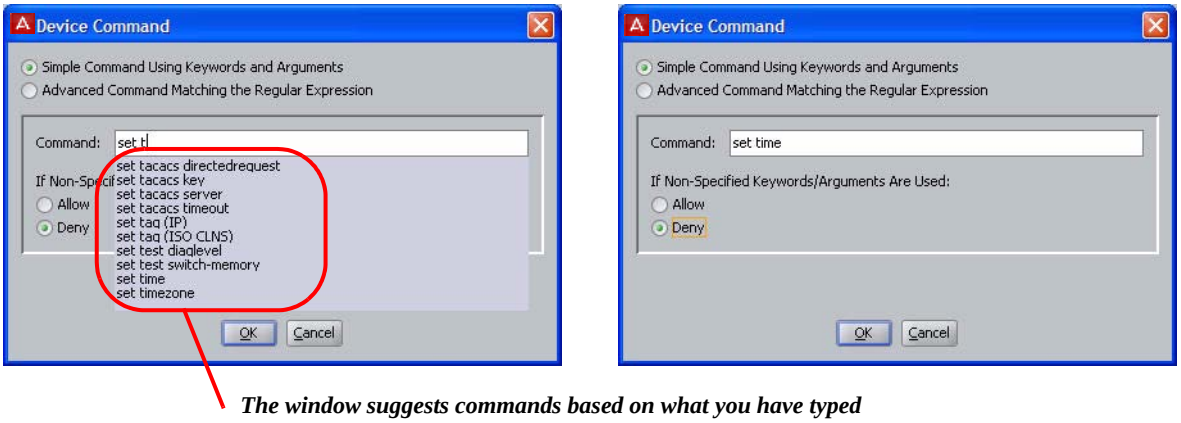
1. In Dashboard's **Configuration** hierarchy tree, expand **Access Policies**, expand **TACACS+**, and click **Device Command Sets**.

Figure 69 The Device Command Set window



2. Click **New**.
3. In the New Device Command Set window, type a **Name** and **Description** for the set. In this window you build your command set by adding commands to the list.

Figure 70 Automatic command completion in the Device Command window



The window suggests commands based on what you have typed

4. You can build the command set list manually or you can import a list.

   To *manually* add to the list, add each command in one of the following ways:

   - Click **Add** and, in the New/Edit Device Command Set window, click the **Simple Command** checkbox and, in the **Command** field, type the command and, optionally, its arguments. The field provides automatic completions based on what you have typed. To allow the command to be used with any argument, tick the **Allow** checkbox. To allow only the specific command and arguments you have typed, tick the **Deny** checkbox. Click **OK** to add the command to the list.

     — OR —

- Click **Add** and, in the New/Edit Device Command Set window, click the **Advanced Command Matching the Regular Expression** checkbox and type the regular expression describing the allowed commands, and click **OK** to add the regular expression to the list.

To *import* a list of commands, do this:

- In the New/Edit Device Command Set window, click **Import**.
- In the Import Commands window, specify the name of your command list file in the **Import File** field, or click **Browse** to find it. The file must contain one command per line, and the command can be followed by arguments. No regular expressions are allowed.
- In the radio buttons at the top, specify how each line is to be interpreted. To allow the command to be used with all arguments, tick the **Match Command Plus Additional Keywords/Arguments** checkbox. To allow only the specific command and arguments you have typed, tick the **Exact Match for Each Command Only** checkbox.
- Click **OK** to import the list.

5. Click **Add** or **Import** again to add more expressions, or click **OK** to save the set.

Next step: Create your administrator authorization policy as shown in "Creating a TACACS+ Access Policy" on page 313.

Viewing or Editing a Command Set

To view or edit a command set, to the following:

1. In Dashboard's Configuration hierarchy tree, expand **Access Policies**, expand **TACACS+**, and expand **Device Command Sets**.
2. Click on the name of a **Command Set** to view it.
3. Click the **Edit** button to Edit the command set.

Creating a TACACS+ Access Policy

Your TACACS+ access policy is a set of rules that Ignition Server evaluates to determine whether a TACACS+ access request is granted or denied access. You apply the policy by creating an Ignition Server authenticator record for the switch an then specifying the TACACS+ policy name in the authenticator record. (See "Creating an Authenticator" on page 96.)

Prerequisites: Note the following prerequisites, based on the type of authorization you use:

- If you are setting up *per-command authorization*, then you should have already created you command set(s) of allowed commands and

arguments. If you have not done this, turn to "Creating a Command Set" on page 311.

- If you are setting up *privilege-level authorization*, then you should have already assigned a privilege level to each command on the equipment your administrators will manage.

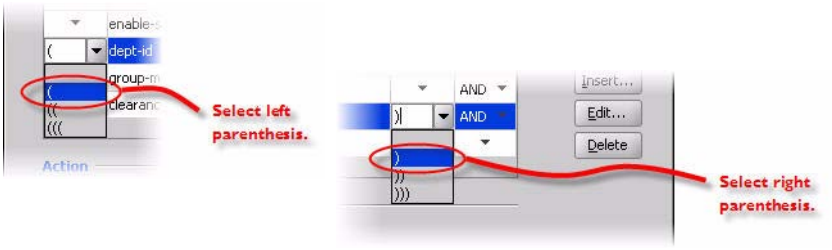
Create your TACACS+ access policy:

1. In Dashboard's **Configuration** hierarchy tree, expand **Access Policies**, and click on **TACACS+**.
2. Click **New**.
3. In the New Access Policy window, type a name for your TACACS+ policy and click **OK**.
4. In the tree, click the name of your new policy.
5. In the **Access Policy** panel, click the **Identity Routing** tab and click **Edit**. Set up your user lookup policy. See "Creating an Identity Routing Policy" on page 223 for help.
6. With your TACACS+ policy still selected in the tree, click the **Authorization Policy** tab and click **Edit**.
7. Add authorization rules by clicking **Add** in the lower left, and then clicking **New** on the right side of the window to add the logic of each rule.
8. In the Constraint Details window, write your constraint:
   a. In the **Attribute Category** drop down list, choose the type of attribute you want to test. (For explanations of the types, see "Attributes Used in Rule Constraints" on page 232.)
   b. Choose the attribute: After you select a type, the list box below the **Attribute Category** field shows the available attributes that match the type you selected. Click on the name of the attribute whose value the constraint should test. In the upper right corner, the window displays the **Data type** of the attribute.
   c. In the drop-down list just below the **Data type** field, choose the comparison operator, such as *Equal To* or *Contains*. This drop-down list contains the operators appropriate to the data type of the attribute you have selected.
   d. Provide the comparison value by doing one of the following:
      - If you want to compare the attribute value with a fixed test value, tick the **Static Value** radio button and type or choose the comparison value in the field below that.
      - If you want to compare the attribute value with a value retrieved from another attribute, tick the **Dynamic Value of Attribute** radio button. In the field just

below that, choose the attribute category (User, Inbound, Authenticator, or Device). In the next field, choose the attribute that should provide the comparison value. The list of attributes contains only those attributes whose data type matches the data type of the attribute on the left side of the constraint.

e. Click **OK** to close the Constraint Details window.

9. In the Edit Authorization Policy window, next to the **Constraint** table, click the **New** or **Insert** button to add more constraints. **New** adds a constraint at the end of the list, and **Insert** adds it above the currently selected row.
10. Add parentheses as necessary to group constraints. To do this:
    - In the **Constraint** section of the Edit Authorization Policy window, find the first constraint to be grouped.
    - Click in the field to the left of the constraint, and click the down-arrow to show the list of parentheses. Click on an appropriate opening parenthesis mark to select it.
    - Find the last constraint to be grouped. Click in the field to the right of the constraint, and click the down-arrow to show the list of parentheses. Click on an appropriate opening parenthesis mark to select it. Click the constraint to complete your entry.



11. In the **Constraint** table, use the **AND** and **OR** conjunctions to form a logical condition statement.
12. Do one of the following:

    If the rule is a *Deny* rule, click **Deny** and click **OK** to save the rule.

    — OR —

    If the rule is an *Allow* rule, specify your TACACS+ permissions by doing one of the following:

    - If you are setting up *per-command authorization*, specify the set of allowed TACACS+ commands: Click the **Command Sets** tab and double-click an entry from the **All Command Sets** list to add the command set to the **Allow Commands In Set** list. Add more sets if needed.

      — OR —

- "If you are setting up Privilege-level authorization, click the **Session Values** tab. Tick the Privilege Level and enter value of **1-15**, similarly tick the session Timeout (maximum length of session, regardless of activity) & idle Timeout (maximum amount of time the session can sit idle between commands) in minutes. Timeouts are provisioned by Ignition Server and enforced by the TACACS+ client. A timeout value of zero means that no timeout is entered

Figure 71 Adding TACACS+ allowed commands to an Allow rule

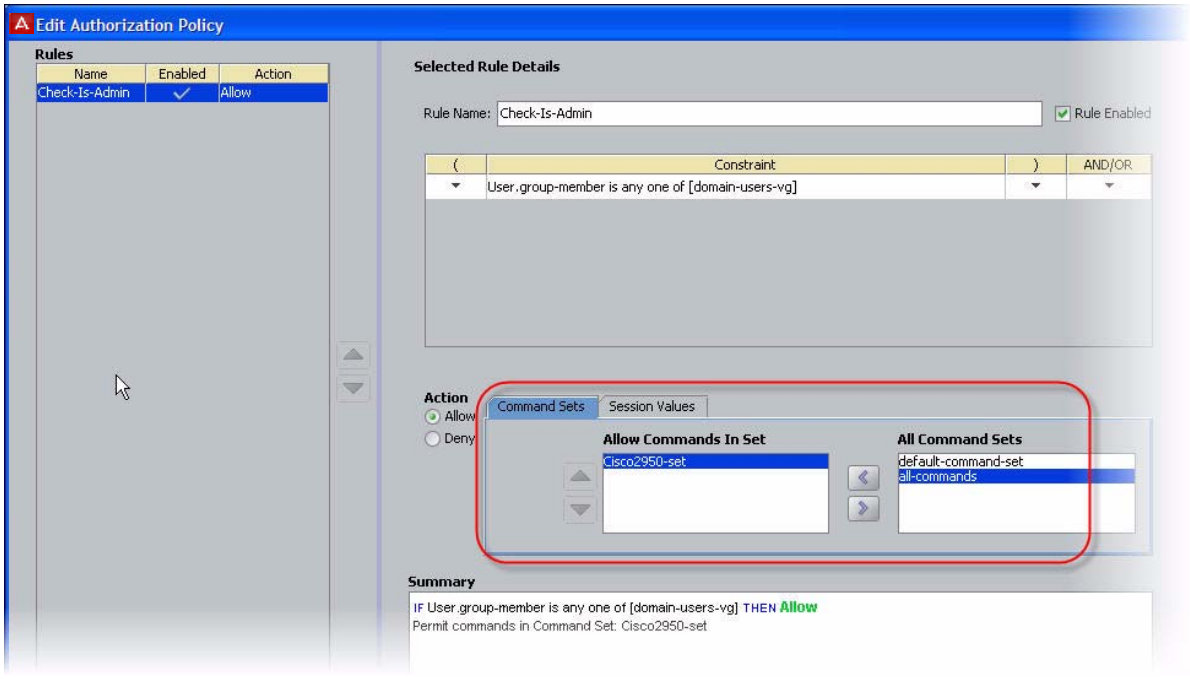
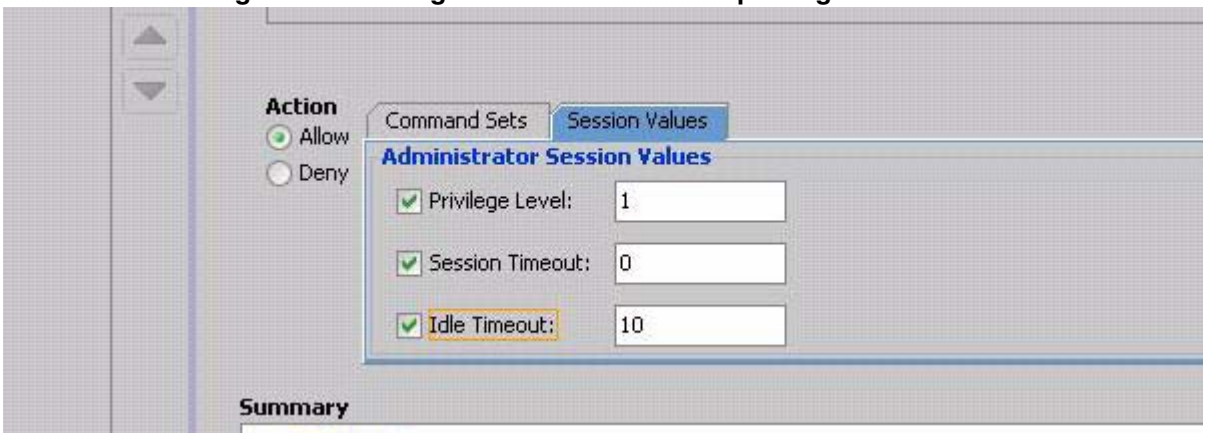


Figure 72 Setting a maximum TACACS+ privilege level for an administrator



13. Click **OK** to save the policy.

Next step: Set your authenticators to use the TACACS+ policy, as explained in "Enable Your Devices for TACACS+ Authorization" on page 317.

Enable Your Devices for TACACS+ Authorization

For each device that will send TACACS+ authorization requests, set up an Ignition Server authenticator record with a TACACS+ policy:

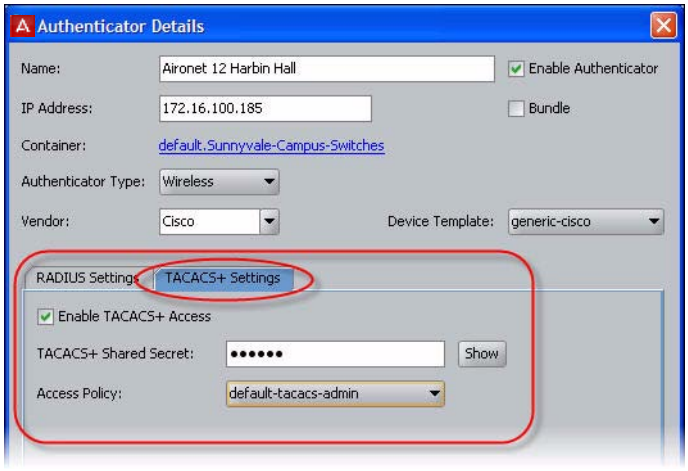
1. Be aware of the following capabilities and limitations:
   - You can use Ignition Server as the *TACACS+ authentication and authorization server*, and you can use it as the *TACACS+ accounting server*.
   - If you use Ignition Server for TACACS+ authorization, you *must* use Ignition Server for TACACS+ authentication.
2. Configure your device to use Ignition Server as its TACACS+ server. Use your device's command line interface or other tool to set up the following:
   - Set the *TACACS+ server address* to the IP address of the Ignition Server *TACACS+* service. **Note:** To find out the Ignition Server TACACS+ service IP address, go to Ignition Dashboard's **Configuration** tree, click the **Site** name (this is usually the name at the top of the tree), click the **Services** tab and click the **TACACS+** tab. The **Bound Interface** field indicates the port. To find the IP address, go back to the **Configuration** tree, click the **Node** name or IP address of your Ignition Server, and click the **Ports** tab.
   - Set the *TACACS+ shared secret* (also known as the *key* or *encryption key*) and make a note of it; you will add it to your authenticator configuration in Ignition Server later.
   - Turn on TACACS+ *authentication* and, optionally, TACACS+ *authorization* on your device for administrator connections to the device.
   - If desired, turn on TACACS+ *accounting* on your device.



Warning: When implementing TACACS+ security on a device, always keep a valid console session open to the device while you test the new TACACS+ authentication and authorization rules. If your new configuration fails or results in denied access, you might become locked out of the device.

3. In Ignition Dashboard, open the Authenticator Details window as follows: In the **Configuration** hierarchy tree, expand **Authenticators**. Do one of:
   - Find your authenticator in the tree, click its name, and click **Edit**; or
   - Click the container that you want to hold your new authenticator, and click **New** near the bottom of the window. Define your authenticator in the Authenticator Details window. For help, see "Creating an Authenticator" on page 96.
4. In the Authenticator Details window, click the **TACACS+ Settings** tab.

Figure 73 Enabling an authenticator for TACACS+ authorization

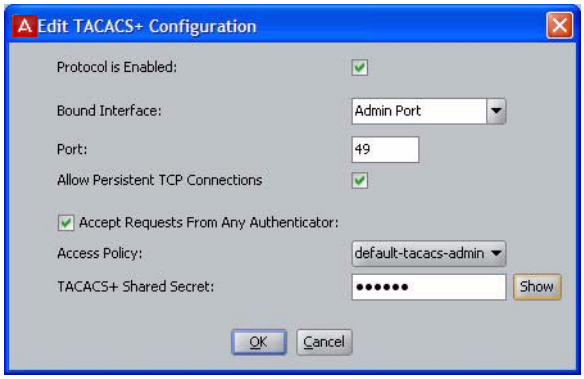


5. Tick the **Enable TACACS+ Access** checkbox.
6. Type the **TACACS+ Shared Secret** that you specified in Step 1.
7. In the **Access Policy** drop-down list, choose your policy. This is the policy you created in "Creating a TACACS+ Access Policy" on page 313.
8. Click **OK** to save the definition.

The TACACS+ Global Authenticator

As explained in "Introduction to Authenticators" on page 89, the *global authenticator* record allows you to create a default TACACS+ access policy that applies to requests from unknown devices. When Ignition Server uses the global authenticator to handle a request, it logs the action with the authenticator name "global-default."

Figure 74 The global authenticator is defined in the TACACS+ configuration



To create the TACACS+ Global Authenticator:

1. In the **Configuration** hierarchy tree of Dashboard, click on your site's name, click the **Services** tab, and click the **TACACS+** tab.

2. Click **Edit**.
3. In the Edit TACACS+ Configuration window, tick the **Accept Requests from Any Authenticator** checkbox.
4. Choose your **Access Policy**. This is the default TACACS+ access policy for all requests from unknown devices.
5. Type the **TACACS+ Shared Secret**. Ignition Server responds only to authenticators that pass this secret string.
6. Click **OK**.

MAC Authentication

This section explains how to configure Avaya Identity Engines Ignition Server to allow devices to connect to your network after they identify themselves by means of their MAC address. After a brief introduction, we explain how to write Ignition Server policies that permit MAC authentication.

Contents

- "Introduction to MAC Authentication" on page 321
- "Creating a MAC-Auth Policy" on page 322
- "Setting Up MAC Auth" on page 323
- "Example MAC Authentication Set-Up Procedure" on page 325
- "VLAN Assignment Using the Device Record VLAN Fields" on page 330
- "Allowed MAC Address Formats" on page 333
- "Notes on Writing MAC Authorization Rules" on page 333

Introduction to MAC Authentication

MAC authentication, or MAC-address checking, verifies that the MAC address submitted by a connecting client device matches an entry on your list of known MAC addresses. Based on your policies, Ignition Server allows the device to connect to your network (and optionally assigns it to a VLAN) or rejects the device. The list of known MAC addresses is stored in the Ignition Server internal data store (*you cannot use an LDAP or AD store for this)*.

MAC authentication is typically employed on 802.1X-authenticated networks as an *802.1X bypass mechanism* for devices that are incapable of performing 802.1X authentication. For example, if your environment contains printers that cannot authenticate via 802.1X, you can set Ignition Server to allow those devices to connect without performing an 802.1X authentication and to place them on an appropriate, limited-access VLAN.

To enforce MAC authentication, create device records that specify your set of allowed MAC addresses, and “MAC Auth” rules in Ignition Server that determine which devices are allowed to connect, as well as where and how they are allowed to connect. Typically, these rules also force the device onto the appropriate VLAN.



Note: Do no confuse MAC authentication with *Windows machine authentication* and *asset correlation*, which uses Windows machine authentication. (See “Introduction to Windows Machine Authentication” on page 295 for details.)



Warning: Allowing MAC Authentication Can Reduce Network Security

Using MAC authentication incorrectly can reduce the overall security of your network. When you activate MAC authentication on an authenticator along with one or more 802.1X authentication methods, the default behavior of most switches means that, even though you have specified 802.1X authentication, the typical switch attempts MAC authentication if the 802.1X user authentication fails. As a result, an ill-intentioned user can exploit the weakness of the less secure MAC authentication to bypass the 802.1X authentication.

In some cases, MAC authentication can be less secure than 802.1X user authentication if it is configured to use only the client device’s MAC address as the credential (instead of using a shared secret as a password). In such a case, if an ill-intentioned user acquires the MAC address of one of your allowed devices, he can pass that MAC address in his access request and gain access to the resources that your policy lists as available via MAC Auth in the applicable access policy.

Avaya recommends you take the following precautions: **First**, for switches that support per-port configuration of MAC authentication, you should enable MAC authentication on *only those ports that require it*, such as ports to which printers and other non-802.1X-compliant devices connect. **Second**, you should place all MAC-authenticated devices on a *limited-access VLAN*, as explained in the sections that follow.

Creating a MAC-Auth Policy

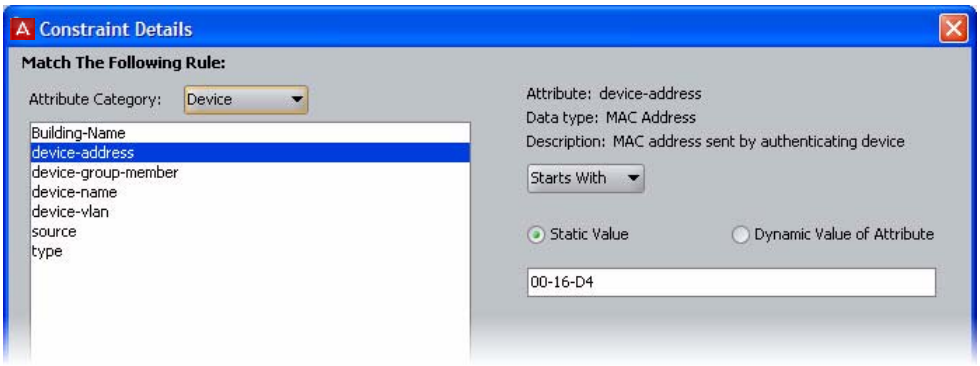
This section shows how to write a device authorization policy for client devices such as laptops and printers. We refer to these policies as “MAC-Auth policies.” The MAC-Auth policy identifies each device by means of its MAC address and authorizes it appropriately. Your rules can also make **VLAN assignments** using outbound values.

1. In the Configuration tree, expand **Access Policies**.
2. Expand **MAC Auth**.
3. Click **New**. (*Note:* You can edit and existing policy by clicking its name in the **Configuration** tree and clicking the **Edit** button on the right side of the window.)
4. Enter a name for the policy and click **OK**.

5. Click the policy name in the tree and click **Edit** on the right side of the window.
6. In the Edit Authorization Policy window, set up a MAC-Auth policy just as you would a RADIUS user authorization policy. For information on using that window, see "Creating a RADIUS User Authorization Policy" on page 242.

   Typically, your MAC-Auth rules will evaluate attributes of the connecting device. In the Constraint Details window, you set up such a rule by choosing the **Attribute Category**, *Device*. For example, to check that the connecting laptop's IP address begins with a known sequence of hexadecimal numerals, go to the **Attribute Category** drop-down list, select *Device*, and then click the attribute name *device-address*. On the right side of the window, choose *Starts With*, click **Static Value** and type the numerals to be matched.

Figure 75 Creating a rule that checks the connecting device's MAC address



For an example MAC authorization rule, see "Example MAC Authentication Set-Up Procedure" on page 325.

If your situation requires that your rules evaluate more detailed information, you can store and evaluate additional device information as shown in "Device Virtual Attributes" on page 208.

Next Steps: Now that you have written your MAC Auth policy, you are ready to enable MAC Auth for your authenticators as shown in the next section, "Setting Up MAC Auth" on page 323.

Setting Up MAC Auth

This section shows you how to enable MAC Auth for an authenticator. Later in this chapter, we provide an example implementation in the section, "Example MAC Authentication Set-Up Procedure" on page 325.

1. Create an Ignition Server **outbound value** for each VLAN to be assigned devices. This is a name you use to refer to the VLAN so that you can write

Ignition Server policies that assign devices to that VLAN. For instructions, see “Create an Outbound Value for Each VLAN” on page 289.

2. Set up each authenticator that supports MAC authentication. This tells Ignition Server that these switches relay MAC authentication requests from devices to the Ignition Server RADIUS service. For each such authenticator, you use the **Authenticator Details** panel to make these settings:
   - In the **RADIUS Settings** tab, tick the **Enable MAC Authentication** checkbox.
   - In the **Access Policy** drop-down list, choose your *MAC Auth policy*. (If you need to create one, see the preceding section, “Creating a MAC-Auth Policy” on page 322.)
   - Specify how the authenticator password should be checked: “**Do not use password**” tells Ignition Server to skip password checking, “**Use RADIUS shared secret as password**” tells Ignition Server to use the authenticator’s RADIUS shared secret, and “**Use this password**” lets you specify your own password.

   Note! MAC authentication normally uses the client’s MAC address as its only credential, meaning that a device with a known address is allowed to connect. There is an exception to this rule: If, in your authenticator definition in Ignition Server, you set the **MAC Address Source** to “Inbound-User-Name” then Ignition Server also evaluates the password passed with the request. In that case, Ignition Server retrieves the password from the “User-Password” RADIUS attribute, which is virtualized in Ignition Server as “Inbound-User-Password.”
3. In the **Device Template** of each authenticator that supports MAC authentication:
   - specify the **MAC Address Source** attribute. This tells Ignition Server which inbound RADIUS attribute contains the MAC address of the connecting device. Typically, the `Inbound-Calling-Station-Id` or `Inbound-User-Name` attribute is used. If the desired attribute is not in the list, see “Adding a New RADIUS Attribute” on page 275.
   - If you plan to perform VLAN assignment, tick the desired **VLAN Method**.
4. For each device allowed to connect to the network, create a **device record**. Each device record is a record of a known MAC address. These records are stored in the Ignition Server internal data store; you cannot

retrieve device information from an external store. For instructions, see “Creating a Device Record” on page 120 or “Importing Device Records” on page 124.

Warning Concerning Users With MAC Addresses as Names

Please note that under certain conditions, Ignition Server defies the precedence set in your Tunnel Protocols list, and performs a device authentication using MAC-AUTH instead of performing a user authentication using PAP. The circumstances that can cause this to happen are as follows: First, you must have a user whose name is a valid MAC address. Second, your authenticator’s device template must specify Inbound-User-Name as the MAC address source. Third, your credential validation policy (in Dashboard’s Configuration hierarchy, expand Access Policies, expand RADIUS, and click your policy name, click the Authentication Policy tab, and refer to the Authentication Protocols section) must have PAP and MAC-AUTH selected as authentication protocols, with PAP placed higher in the list than MAC-AUTH. Under these conditions, if the user whose name is a MAC address attempts to connect, he is treated as a device, not a user.

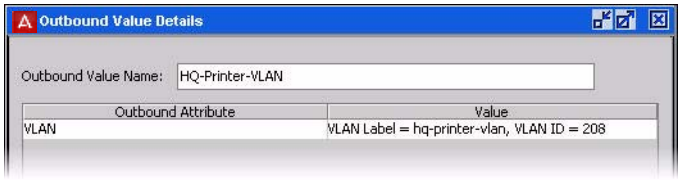
Example MAC Authentication Set-Up Procedure

This example shows how to set up MAC authentication in Ignition Server. In this procedure, we build an example policy that lets printers connect to your network and places them on a dedicated VLAN.

This example assumes the printers on your network perform a MAC authentication to connect to the network. For each printer, you create a device record in Ignition Server, and each printer’s device record has a **type** label of “printer.” This rule checks the type of device and, if it is labeled “printer”, it places the device on the HQ-printer-VLAN.

Procedure Set up MAC authentication in Ignition Server as follows:

1. Avaya recommends that, if you use MAC authentication, you set up a MAC Auth policy that assigns devices to one or more limited-access VLANs. To prepare for VLAN assignment:
   - Set up the VLAN(s) on your network equipment.
   - In Ignition Server, create an outbound value for each VLAN to which you plan to assign devices. For instructions, see “Create an Outbound Value for Each VLAN” on page 289. For this example, we use an outbound value, **HQ-Printer-VLAN** that sends the VLAN assignment value of “hq-printer-vlan” or “208” in the RADIUS attribute, Tunnel-Private-Group-Id.



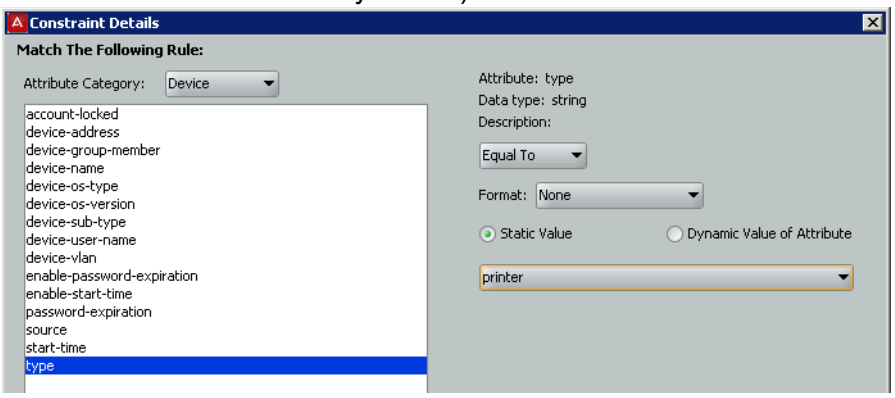
Note! There is another way to assign devices to VLANs. You can specify a VLAN designation in each device record, instead, as explained in “VLAN Assignment Using the Device Record VLAN Fields” on page 330.

2. Create a MAC Auth policy made up of one or more rules. Your rules should evaluate the device and the context to determine if the device should be given access, and you should assign the device to a VLAN if possible. For a device to authenticate successfully at least one rule in the policy must trigger an *Allow*.

   Ignition Server automatically checks that the device is a known device by checking the device’s MAC address against the list of device records in the internal store.

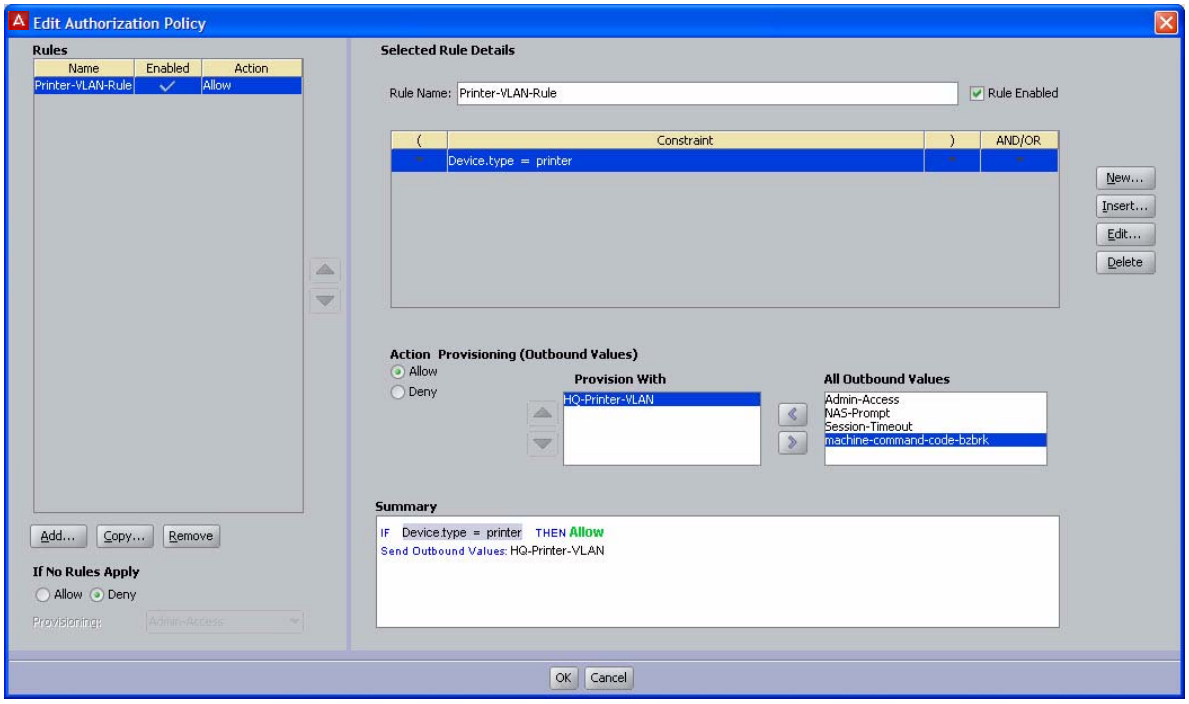
   The steps below show an example that performs VLAN assignment. To set up a MAC Auth rule do the following:

   - In Dashboard’s Configuration hierarchy, expand **Access Policies**, and expand **MAC Auth**. Click **New** to create a new policy or click a policy name to edit an existing policy. Once you have clicked the name of your MAC Auth policy, it appears in the **Access Policy** panel. Click **Edit** on the right side of the window.
   - In the Edit Authorization Policy window, click the **Add** button below the **Rules** list.
   - In the New Rule dialog, give the rule a name and click **OK**. For example, you might call the rule, “Printer-VLAN-Rule”,
   - In the MAC Authorization Policy window, in the Selected Rule details section, click **New** to add a constraint. (You can add as many constraints as you like.)

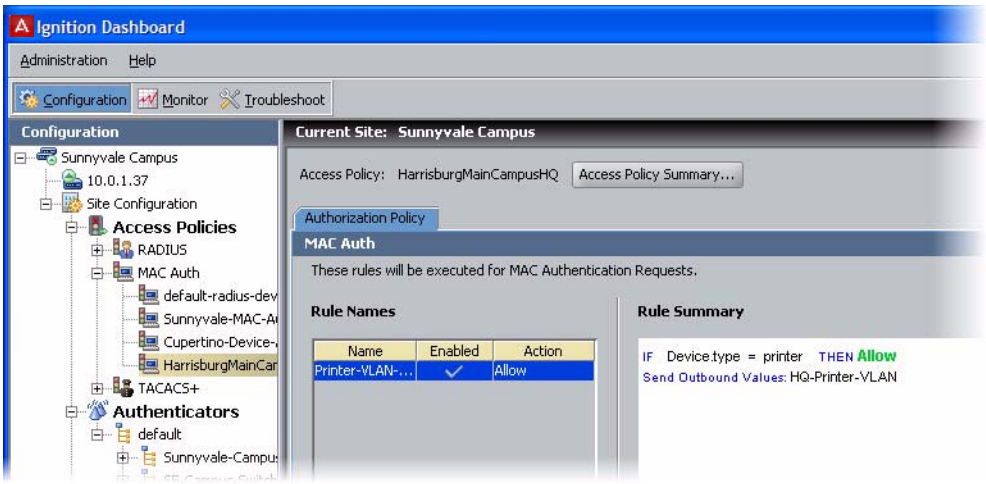


   - In the Constraint Details window, go to the **Attribute Category** drop-down list and select **Device**. In the list below this, choose **type**. In the drop-down list on the right, click **Equal To**. Select the **Static Value** checkbox. In the drop-down list below this, click **printer** Click **OK**.

- In the MAC Authorization Policy window, with your “Printer-VLAN-Rule” still selected, under **Action** tick the **Allow** radio button. In the **Provisioning** section, tick the checkbox next to “HQ-Printer-VLAN”. (If this value is not in the list, create it now as explained at the beginning of this procedure.) Click **OK**.



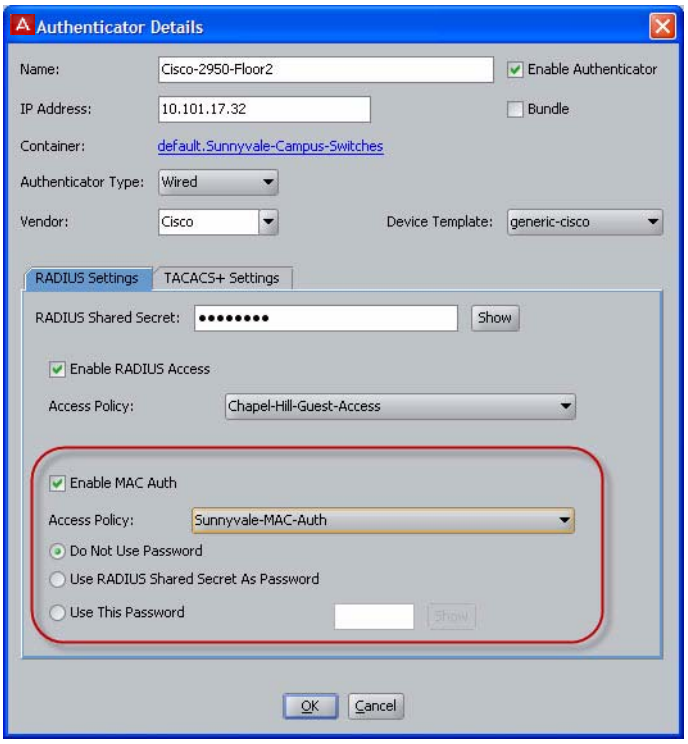
- Your policy has been saved.



3. Set up the authenticators that support MAC authentication. Create or edit each authenticator record in the **Authenticator Details panel** of Dashboard. (From the main window, expand the Authenticators node in

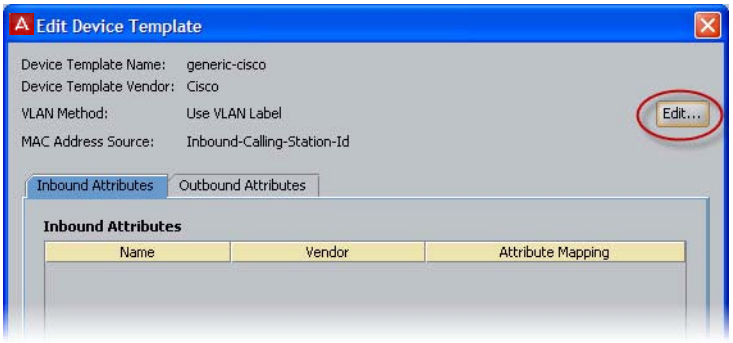
the hierarchy tree. Browse to find your authenticator, and click its name, and click **Edit** to edit it.) For each authenticator that supports MAC authentication, do the following:

- Set the **Enable MAC Auth** flag.
- In the **Access Policy** drop-down list, choose the name of the MAC Auth policy you set up in Step 2.
- Specify how the authenticator password should be checked. You have three choices. To skip password checking (the most common setting), tick the checkbox **Do not use password**. To use the authenticator's shared secret as the password, tick the checkbox, **Use authenticator's shared secret as password**. To specify a password, tick the checkbox, **Use this password**, and type the password in the text field.

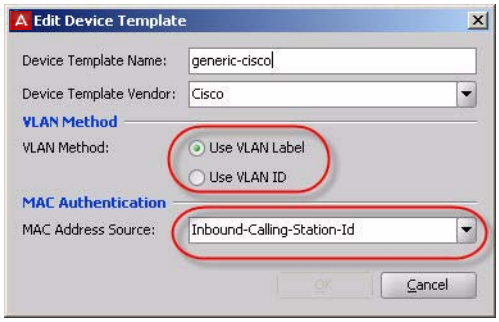


4. In the Device Template of each authenticator that should support MAC authentication, you must specify the MAC address attribute. Open the Device Template window as follows: In Dashboard's **Configuration** hierarchy tree, expand the **Provisioning** node and click **Vendors/VSAs**.

In the **Vendors** list, scroll to find the manufacturer of your authenticator, expand the node, click **Device Template**, and in the right pane, double-click the name of your authenticator's device template.

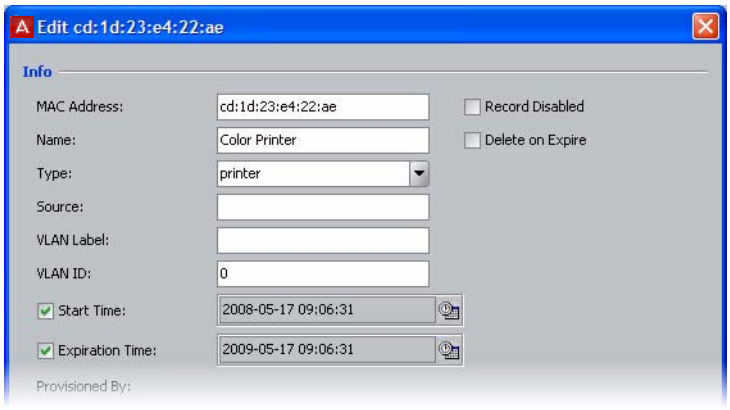


- In the Device Template window, click **Edit**.



- In the Edit Device Template window, in the **MAC Address Source** field, choose the name of the inbound RADIUS attribute that contains the MAC address of the connecting device. Typically, the `Inbound-Calling-Station-Id` attribute or the `Inbound-User-Name` attribute is used. If the desired attribute is not in the list, see "Adding a New RADIUS Attribute" on page 275.
- If you plan to perform VLAN assignment, tick the desired **VLAN Method**.
- Click **OK** and click **Done**.

5. For each printer that you allow to connect, create a device record. For instructions, see “Creating a Device Record” on page 120. For this example, make sure the **Type** of each device record is set to “printer”.



VLAN Assignment Using the Device Record VLAN Fields

Using Ignition Server outbound values you can set Ignition Server to assign a connecting client device to the VLAN specified in its device record. (This is an alternative to the VLAN assignment approach shown in the “Example MAC Authentication Set-Up Procedure” on page 325.)

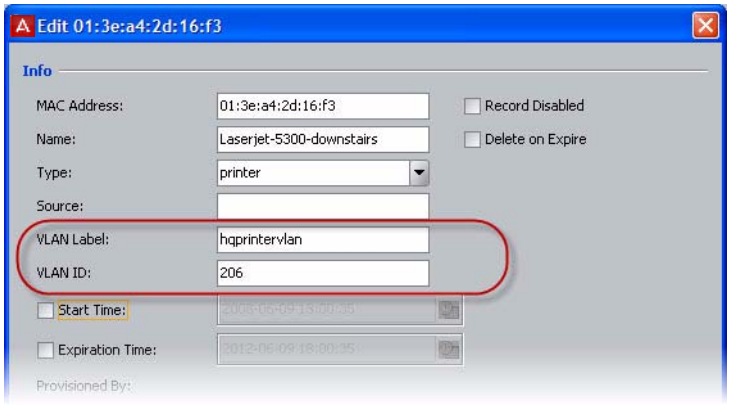
To set this up:

1. In the device templates of your authenticators, set the desired VLAN designation format that should be used in RADIUS messages to your switch:
   - In Dashboard’s **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**. In the **Vendors panel**, scroll to find the manufacturer of your authenticator, expand the node, click **Device Templates**, and in the right pane, double-click the name of your authenticator’s device template. In the Device Template window, click **Edit**.

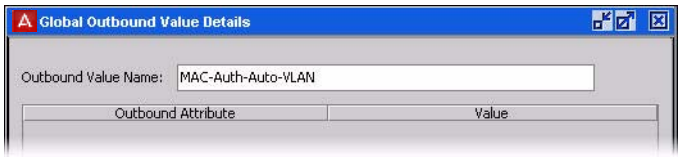


   - In the Edit Device Template window, tick the desired **VLAN Method**. **VLAN Label** uses a string and **VLAN ID** uses an integer value. Click **OK**.
2. In each device record, specify the desired VLAN:

- In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Internal Store, and click Internal Devices. In the Device Records View, click **New** or **Edit** to open your new or existing device record.
- In the Device Record Details window, specify your **VLAN Label** or **VLAN ID**, and click **OK** to save. Note that VLAN labels are case sensitive for some authenticators.

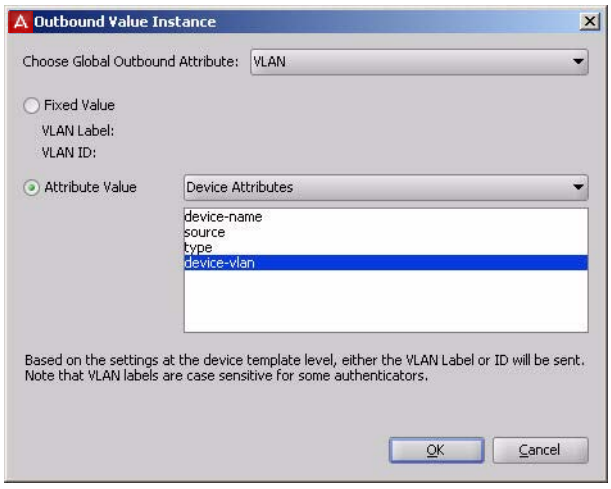


3. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**.
4. In the Outbound Values panel, click **New**.
5. In the Outbound Value Details window, type a name for the outbound value. This is the name that appears in the Constraint Details window when you write rules that assign the VLAN.



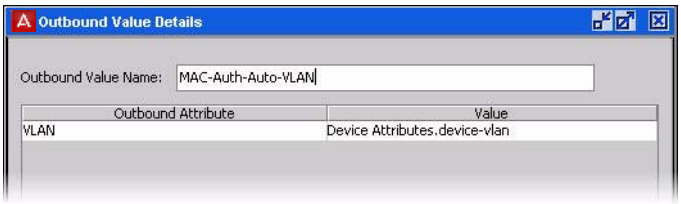
6. Click **New**.

7. In the Outbound Value Instance window:

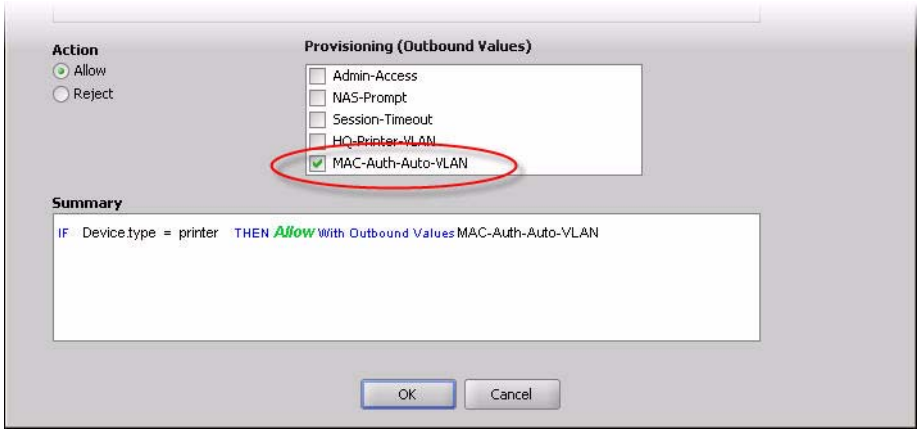


   - In the **Choose Global Outbound Attribute** drop-down list, choose **VLAN**.
   - Tick the **Attribute Value** checkbox and choose **Device Attributes** in the drop-down list just to the right.
   - In the list, select **device-vlan**. This forces Ignition Server to use the VLAN value from the device record. Based on the settings at the device template level, either the VLAN Label or the VLAN ID is sent.
   - Click **OK**.

8. In the Outbound Value Details window, click **Save** to save the outbound value and dismiss the window.
9. From the Outbound Values panel, you can check your outbound value by selecting its name and clicking **Edit**.



10. In your authorization policies (accessed from the Configuration hierarchy tree by clicking you policy name under the RADIUS or MAC Auth section and clicking **Edit**), use the outbound value in an Allow rule. At runtime, when the Allow rule is triggered, Ignition Server sends the VLAN assignment attribute to the authenticator.



Allowed MAC Address Formats

In the Device Record and Constraint Details windows, you can type a MAC address in any of the following formats. Upper and lower case letter characters are allowed. You can use colon, period, or hyphen characters as delimiters, but do not mix delimiters.

Allowed MAC Address Formats

11-22-33-44-55-66	1122-3344-5566	112233-445566
11.22.33.44.55.66	1122.3344.5566	112233.445566
11:22:33:44:55:66	1122:3344:5566	112233:445566
112233445566		

Notes on Writing MAC Authorization Rules

When you write MAC authorization rules, you can evaluate the following types of attributes:

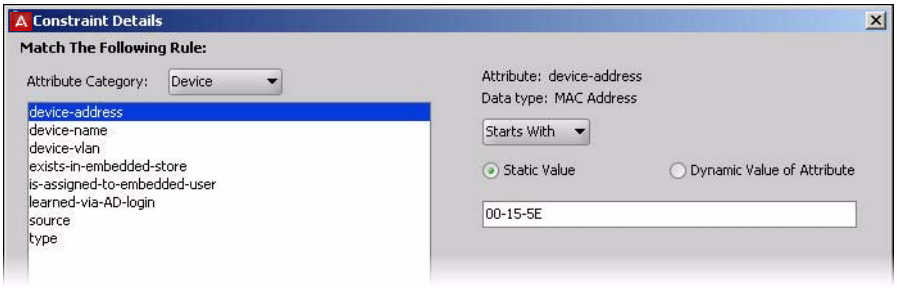
- **inbound attributes**: values passed by the authenticator in the form of RADIUS attributes or VSAs. These typically describe the context, time, or originating device of the access request. See “Inbound Attributes” on page 234.
- **authenticator attributes**: Ignition Server -stored data that describes the switch or access point, such as the name of the switch manufacturer, its location in the Ignition Server authenticator hierarchy, or the name of the Ignition Server MAC Auth access policy being used. See “Authenticator Attributes” on page 237.

- **device attributes**: values that describe the connecting client device. See "Device Attributes" on page 236.

Comparing a Device's MAC Address

You can compare a MAC address to a partial or full MAC address in the rules you define in the Constraint Details window:

- To compare a full MAC address, select the **Attribute Category**, Device, pick **device-address**, choose **Equal To**, tick **Static** value, and type the address in any of the allowed formats (see the preceding section) For example, you might type *02:e5:6c:12:dd:7e*.
- To compare a partial MAC address, use the Starts With operator instead of the Equal To operator, and type a partial MAC address with no asterisks. For example, you might type *02:e5:6c*.



Checking a Device's Group Membership

You can check a device's group membership in the MAC authorization rules you define in the Constraint Details window. To do this, select the **Attribute Category**, Device, pick **group-member**, choose an operator (for example, **Equals** or **Any One Of**) tick **Static** value, and click the **Add** button below the list area. Use the Add Value dialog to add a group name to the list, and click **OK**. If you need to add more group names to the list, keep using the **Add** button until you have added all the group names you need.

Asset Correlation

This chapter introduces the concept of Avaya Identity Engines Ignition Server asset correlation policies and explains how to create rules that prevent a user from connecting with any device other than his or her authorized device.

Contents

- "Introduction to Asset Correlation" on page 335
- "Creating Asset Correlation Policies" on page 336
- "Viewing Currently Authenticated Devices" on page 343

Introduction to Asset Correlation

An asset correlation policy lets you specify which devices a person can use to connect to your network. With an asset correlation policy in place, Ignition Server checks that the device (the "asset") correlates with the user by checking that the user has authenticated and by matching the device's identifying information (along with the user's credentials, this information is passed to Ignition Server in the access request) with a record from your list of approved devices.

You have a choice of three ways to have Ignition Server check the device identity. In your policy, you set Ignition Server to check one of these three **correlation types**:

1. That the device MAC address has been specified as an allowed address in the Ignition Server internal database. This is called **exists-in-embedded-store**.
2. That the device has been assigned to an internal user in Ignition Server. This is called **is-assigned-to-embedded-user**.
3. That the device has authenticated itself to Ignition Server via Windows machine authentication. This is called **learned-via-AD-login**.

MAC Address vs. Windows Machine Authentication

There are two ways Ignition Server asset correlation can identify the connecting device:

- using the MAC address of the client
- using Windows machine authentication

In the preceding list of correlation types, the first two (**exists-in-embedded-store** and **is-assigned-to-embedded- user**) use the client MAC address, and the **learned-via-AD-login** type uses Windows machine authentication. Depending on which authentication type you use, there are important differences in how you set up your policy.

The first difference concerns how you store the device record for each device. When you use MAC authentication, the device record resides in the Ignition Server internal store. When you use Windows machine authentication, the device record resides in Active Directory.

The second important difference concerns how you manage the access rights of devices. For MAC authentication, you cannot revoke a current lease, you can only revoke the device's right to connect by ticking the **Device Disabled** check box in the Ignition Server device record. For Windows machine authentication, you can do both. To revoke the current lease of a Windows-authenticated device, (that is, to force the device to reauthenticate) you delete its record from the Learned Devices tab of the Monitor: Current Site panel. To revoke a Windows-authenticated device's right to connect, you disable or delete its record in AD.

Asset correlation: Differences between MAC auth and Windows machine auth

	MAC Address	Windows Machine Auth
Location of device record	Ignition internal store	Active Directory
How to view currently connected devices?	Ignition **Monitor: Current Site** panel: **AAA Summary** tab	Ignition **Monitor: Current Site panel: Learned Devices** list
How to revoke current device lease?	Not applicable.	Delete the device's row from the **Monitor: Current Site panel: Learned Devices** list.
How to revoke the device's right to connect?	Tick the **Device Disabled** check box in its Ignition Server device record, *or* write a rule that denies access to the device.	Disable or delete its record in AD.

Creating Asset Correlation Policies

As mentioned earlier, there are three main types of asset correlation policies. Your policy can require

- that the device MAC address has been specified as an *allowed address* in the Ignition Server internal database, as explained in "Requiring the User to Connect Using an Allowed Device" on page 337; or

- that the device has been *assigned to an internal user* in Ignition Server, as explained in "Requiring the User to Connect Using His or Her Assigned Device" on page 338; or
- that the device has authenticated itself to Ignition Server via *Windows machine authentication*, as explained in "Requiring the User to Connect Using a Machine Authenticated-Device" on page 340.

Requiring the User to Connect Using an Allowed Device

This section explains a policy that requires the user to log in using a computer that is on Ignition Server's list of known devices. This policy uses the test **exists-in-embedded-store** in its authorization rules. Strictly speaking, this is not an *asset correlation* rule, as it does correlate the particular device with its owner. Nonetheless, you build this policy much like you do an asset correlation policy.

To set up this policy, do the following:

1. In Dashboard's Configuration hierarchy tree, expand **Access Policies** and expand **RADIUS**. Click the name of your access policy. Click the Authentication Policy tab and click Edit. Set up your authentication policy as usual. ("User Authentication Policy" on page 211.)
2. In the Access Policy panel of Dashboard (in Dashboard's Configuration hierarchy, expand **Access Policies**, expand **RADIUS**, and click your policy name), click the **Authorization Policy** tab and click **Edit**. Create a rule to *Deny* any user who is attempting to log in with an unknown device:
   - On the left side of the Edit Authorization Policy window, click **New**.
   - In the New Rule dialog, give the rule a name. For example, you might call your rule, "Require-allowed-device". Click **OK**.
   - On the left side of the Edit Authorization Policy window, click the name of your new rule, and click **New** in the **Selected Rule Details** section.
   - In the Constraint Details window, in the **Attribute Category** drop-down list, select **Device**. In the list below this, click "exists-in-embedded-store". On the right side of the window, click the **False** radio button. Click **OK**.



Note: The user's device identifies itself by means of its MAC address, which is sent in the RADIUS access request. For the purpose of device identity-checking coupled with a user authentication, Ignition Server always gets the device's MAC address from the *inbound-calling-station-id* RADIUS attribute. (The **MAC Address Source** setting from the device template is not used in this case.)

   - In the Edit Authorization Policy window, click the **Action**, "Deny". This completes the definition of your first rule. Keep the window open.

3. Next, create your second rule. Ignition Server requires at least one rule in your rule set to evaluate to “Allow” before it grants the user access. Since the rule you defined above is a “Deny” rule, you must add an “Allow” rule as follows:
   - On the left side of the Edit Authorization Policy window, click **New**.
   - In the New Rule window, type the name, “allow-rule”. Click **OK**.
   - In the Selected Rule Details section of the Edit Authorization Policy window, click **New**.
   - In the Constraint Details window, in the **Attribute Category** drop-down list, select **System**. In the list below this, click **True** and click **OK**.
   - In the Action part of the Edit Authorization Policy window, click **Allow**. This completes the definition of your second and final rule.

     Hint! You can review your rules like so: Click a rule’s name on the left side of the Edit Authorization Policy window. When you click the name, the rest of the window displays the logic of that rule.
4. Click **OK** to save the rules and close the Edit Authorization Policy window. Your rule set is defined to reject the user if his or her computer is not defined in Ignition Server.
5. For each device that you want to be allowed to connect, create a **device record**. For instructions, see “Creating a Device Record” on page 120 or “Importing Device Records” on page 124.

This example rejects the user outright if the device is unknown. You can choose to place such users on a limited-access VLAN, instead. See “VLAN Assignment” on page 287.

If you need more detailed information to drive your policy decisions, you can store and evaluate additional device information as shown in “Device Virtual Attributes” on page 208.

Requiring the User to Connect Using His or Her Assigned Device

This section explains an asset correlation policy that allows the user to log in only with a device that has been assigned to him or her in Ignition Server. This policy relies on the MAC address of the user’s computer to prove that computer’s identity. This policy uses the test **is-assigned-to-embedded-user** in its authorization rules.

Note! If you instead use Windows machine authentication to check the identity of users’ computers, then you may wish to follow the instructions in “Requiring the User to Connect Using a Machine Authenticated-Device” on page 340, instead of the steps below.

To set up this policy, do the following:

1. In Dashboard's Configuration hierarchy tree, expand **Access Policies** and expand **RADIUS**. Click the name of your access policy. Click the **Authentication Policy** tab and click **Edit**.
2. Set up your authentication policy as usual. ("User Authentication Policy" on page 211.)
3. Click the **Authorization Policy** tab and click **Edit**. Create a rule to *Deny* any user who is attempting to log in with a device that is not assigned to him or her:
   - On the left side of the Edit Authorization Policy window, click **New**.
   - In the New Rule dialog, give the rule a name. For example, you might call your rule, "Require-assigned-device-of-user". Click **OK**.
   - On the left side of the Edit Authorization Policy window, click the name of your new rule, and click **New** in the **Selected Rule Details** section.
   - In the Constraint Details window, in the **Attribute Category** drop-down list, select **Device**. In the list below this, click `is-assigned-to-embedded-user`. On the right side of the window, click the **False** radio button. Click **OK**.



Note: The user's device identifies itself by means of its MAC address, which is sent in the RADIUS access request. For the purpose of device identity-checking coupled with a user authentication, Ignition Server always gets the device's MAC address from the *inbound-calling-station-id* RADIUS attribute. (The **MAC Address Source** setting from the device template is not used in this case.)

   - In the Edit Authorization Policy window, click the **Action**, "Deny". This completes the definition of your first rule. Keep the window open.
4. Next, create your second rule. Ignition Server requires at least one rule in your rule set to evaluate to "Allow" before it grants the user access. Since the rule you defined above is a "Deny" rule, you must add an "Allow" rule as follows:
   - On the left side of the Edit Authorization Policy window, click **New**.
   - In the New Rule window, type the name, "allow-rule". Click **OK**.
   - In the Selected Rule Details section of the Edit Authorization Policy window, click **New**.
   - In the Constraint Details window, in the **Attribute Category** drop-down list, select **System**. In the list below this, click **True** and click **OK**.
   - In the Action part of the Edit Authorization Policy window, click **Allow**. This completes the definition of your second and final rule.

     Hint! You can review your rules like so: Click a rule's name on the left side of the Edit Authorization Policy window. When you click the name, the rest of the window displays the logic of that rule.

5. Click **OK** to save the rules and close the Edit Authorization Policy window. Your rule set is defined to reject the user if he or she is trying to connect with a computer that is not assigned to him or her in Ignition Server.
6. For each device that you want to be allowed to connect, create a **device record**. For instructions, see "Creating a Device Record" on page 120 or "Importing Device Records" on page 124.
7. For each user that you want to be allowed to connect, create an **internal user record**. For instructions, see "Creating an Internal User" on page 116.
8. Assign each user's device to that user. See "Assigning a Device to a User or Group" on page 123.

This example rejects the user outright if the device is not assigned to the user. You can choose to place such users on a limited-access VLAN, instead. See "VLAN Assignment" on page 287.

If you need more detailed information to drive your policy decisions, you can store and evaluate additional device information as shown in "Device Virtual Attributes" on page 208.

Requiring the User to Connect Using a Machine Authenticated-Device

This section explains an asset correlation policy that relies on Windows machine authentication to check the computer's identity. This policy uses the test **learned-via-AD-login** in its authorization rules.

In this example, we check whether the user's computer has earlier completed a successful Windows machine authentication. If it has, we place the user on the full-access VLAN that staff members use. If it has not, we place the user on the same limited access VLAN to which the computer was granted access to when it performed its machine authentication.

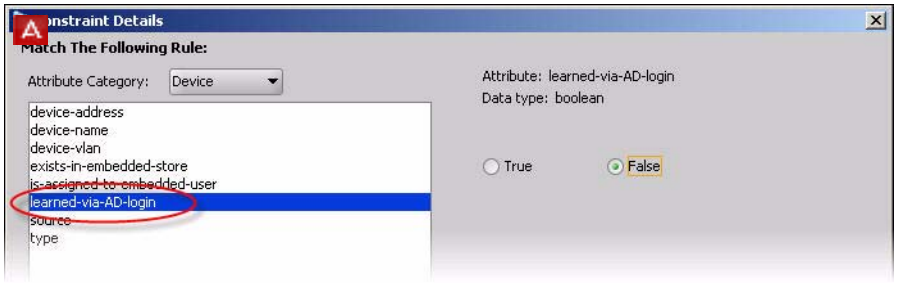
Note! If your site uses the *MAC address* instead of Windows machine authentication to identify each user's computer, then you must follow the instructions in the section, "Requiring the User to Connect Using His or Her Assigned Device" on page 338, instead of the steps below.

To set up this asset correlation policy, do the following:

1. Set up your machine authentication policy as explained in "Setting up Microsoft Windows Machine Authentication" on page 298.
2. Set up any VLANs you need. In this example we use two VLANS: Limited-Access-VLAN which offers minimal access for users and machines that have not sufficiently authenticated, and HQ-Staff-VLAN which provides access to the internal network. To set up each VLAN:
   - Set up the VLAN on your network equipment.

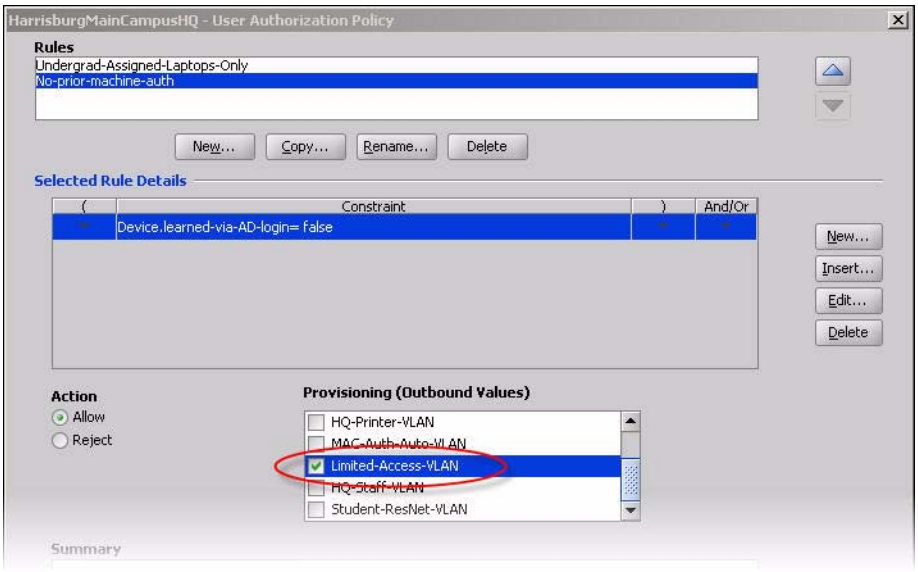
- Create an Ignition Server **outbound value** for each VLAN to which you plan to assign devices. For instructions, see “Create an Outbound Value for Each VLAN” on page 289.

3. Open your Ignition Server RADIUS policy (in Dashboard’s Configuration hierarchy, expand **Access Policies**, expand **RADIUS**, and click the name of the access policy in which you saved your Windows machine authentication policy.
4. In the Authorization Policy tab click **Edit**.
5. Create the first rule:
   - On the left side of the Edit Authorization Policy window, click **New**.
   - In the New Rule dialog, give the rule a name. For this example, we call our rule, “No-prior-machine-auth.” Click **OK**.
   - On the left side of the Edit Authorization Policy window, click the name of your new rule, and click **New** in the Selected Rule Details section.
   - In the Constraint Details window, in the Attribute Category drop-down list, select **Device**. In the list below this, click learned-via-AD-login. On the right side of the window, click the **False** radio button. Click **OK**.



   - In the Edit Authorization Policy window, click the Action, “**Allow**”, and under Provisioning, tick the **Limited-Access-VLAN** check box and untick all other check boxes. Your No-prior-machine-auth rule is now

defined to place the user on the limited-access VLAN if his or her computer is not Windows machine authenticated.



6. Create the second rule:
   - On the left side of the Edit Authorization Policy window, click **Copy**.
   - In the top half of the Copy Rule dialog, navigate to find the rule, “No-prior-machine-auth.” Click it and click **OK**.
   - In the Edit Authorization Policy window, click the copied rule (its name is the same as the copied rule's name, but with a “1” appended) and click Rename. Call the rule, “Has-prior-machine-auth”. Click **OK**.
   - With the rule name still highlighted, click **Edit** in the Selected Rule Details section.
   - In the Constraint Details window, click the **True** radio button and click **OK**.
7. In the Edit Authorization Policy window, click the Action, “**Allow**”, and under Provisioning, tick the **HQ-Staff-VLAN** check box and untick all other check boxes. Your Has-prior-machine-auth rule is now defined to place the user on the internal VLAN if his or her computer has successfully performed Windows machine authentication.
8. Click **OK** to save the rules and close the window.
9. With your policy in place, each user’s machine must have a current Windows machine authentication in order for that user to log in. Machine authentications occur automatically when the machine is booted up or connected to the network, and the authentication lasts for the time-to-live (TTL) period defined in Ignition Server. Set up the TTL now as explained

in “Setting TTL for Windows Machine Authentication” on page 305. (In a running Ignition Server installation, you can view the current machine authentications as explained below.)

Viewing Currently Authenticated Devices

Use the Monitor: Current Site panel to view the list of currently authenticated devices. For Windows machine-authenticated devices, you can revoke current leases:

- Devices authenticated via *Windows machine authentication* appear in the **Learned Devices** tab. See “Learned Devices Tab” on page 447.
- Devices authenticated via *MAC authentication* appear in the **AAA Summary** tab. See “AAA Summary Tabs” on page 442.

Command Line Interface

The Avaya Identity Engines Ignition Server Command Line Interface (“CLI”) allows you to carry out a limited set of administrative actions on your Avaya Identity Engines Ignition Server. This chapter explains how to connect to through an SSH connection:

- “Connecting to the CLI Through an SSH Connection” on page 345

Before you can connect over SSH, you must use Ignition Dashboard (or the CLI’s `sshd` command) to activate and configure Ignition Server’s SSH service.

Connecting to the CLI Through an SSH Connection

Ignition Server allows you to connect to the CLI via an SSH session secured using your password or public/private key pair. The connection travels over the local LAN through the designated Ethernet port on the Ignition Server. By default, the Admin port is used. When connected in this way, your credentials and communication with the Ignition Server are encrypted using the SSHv2 protocol.

To support SSH CLI connections, you must activate the Ignition Server’s SSH service and install on the Ignition Server the public keys of all administrators who should be allowed to connect. At login time, the Ignition Server uses the administrator’s public key to authenticate him or her.

Note! If you have installed *no public keys* in the Ignition Server’s SSH service, the Server nonetheless allows you to connect via SSH. In this case the Ignition Server authenticates you using only your Ignition Server administrator name and password. This approach is less secure because it does not allow you to verify the identity of the Ignition Server. In a standard SSH login, your SSH client has a copy of the Ignition Server’s public key and uses this key to authenticate the Ignition Server.

To configure Ignition Server for SSH:

Before you can establish secure, public-key authenticated SSH connections to the CLI, you must activate the SSH service and import the public keys of all administrators:

1. Generate or find your public/private key pair. You can use a key pair generation tool such as the unix command, *ssh-keygen*, for this. Follow these guidelines:

- You can use any RSA or DSA key that supports SSHv2.
- Use a passphrase that is difficult to guess.

2. Activate the SSH service on Ignition Server:
   - Run Ignition Dashboard, log in to your Ignition Server, and in Dashboard's Configuration hierarchy tree, click the IP address or name of your Ignition Server.
   - Click the **System** tab, click the **SSH** tab, and click **Edit**.
   - Tick the **Enabled** checkbox to turn on the SSH service. By default SSH is made available on the Ignition Server *Admin port* at port *22*. You can change these settings.
   - Click **OK**.
3. Install the public key on the Ignition Server:
   - In Dashboard's **Nodes** panel, in the **System** tab and **SSH** sub-tab, click **Add New Key**.
   - In the *Add public SSH key* window, in the **SSH Public Key Alias** field, type a name to be used to identify this key in Ignition Dashboard.
   - Provide the path and file name of your public key. You can click **Browse** and navigate to find your key, or you can type the path and name in the **SSH Public Key Path** field. Typically the path and name are similar to `/home/mjackson/.ssh/id_rsa.pub`, where `/home/mjackson` is replaced with the path of your home directory.
   - Click **Submit New Key**.

Ignition Server is now configured to accept SSH connections from the administrator whose key you imported. If other administrators are to have access, import their public keys now.

Connecting via SSH

To establish an Ignition Server administrator session over the local LAN with SSH encryption, do the following:

1. Make sure that:
   - Your computer and the Ignition Server Admin port are on the same network.
   - You have installed an SSH client (for example, PuTTY, a Cygwin shell with SSH installed, or a UNIX or Linux shell with SSH) on your computer. The client must be capable of SSHv2.
   - You have activated the SSH service on the Ignition Server.
   - You have installed your public key on the Ignition Server.
2. Connect using the connect command of your SSH client:

- Use your tool's connect command, passing it the Ignition Server administrator account name (the default is `admin`) and the IP address of Ignition Server's SSH port. (For example, using the typical unix command shell you would type, "`ssh -l admin 10.0.22.33`", where `admin` is your administrator name and `10.0.22.33` is the IP address of the SSH port.)
- If this is your first time connecting, your SSH client prompts you to accept the public key of the Ignition Server. The message is similar to: "The authenticity of the host cannot be established. Do you want to continue connecting (yes/no)?" You must accept the key to continue. In the future, your clients use this key to authenticate the Ignition Server.
- When prompted, type the passphrase for your private key.
- The second prompt asks for your "CLI Password." Enter your Ignition Server administrator password. This is the same password you use to log into Dashboard.

After it is connected, the command prompt displays:

```
Identity Engines>
```

Type "?" for a list of commands, and type "exit" to quit. The CLI ends your session automatically after five minutes of inactivity.

A Installing Ignition

This appendix explains how to install Ignition Dashboard and how to connect and configure the Avaya Identity Engines Ignition Server.

Contents

This appendix consists of the following main sections:

- Planning
  - Installation Prerequisites
  - Map Out Your Ignition Server Deployment
- Virtual appliance set up
  - VMware ESXi Server
  - Importing VM
- License set up
  - Applying the license
  - Installing the license
- Software set-up
  - Install Ignition Dashboard on Microsoft Windows
  - Connect to Ignition Server for the First Time
  - Configure the Ignition Server
  - Uninstalling Ignition Dashboard

For information on post-installation management activities, including setting the system time and managing certificates, see Chapter , “Sites, Nodes, and Settings”.

Installation Prerequisites

To install Ignition Server, you must have the following tools and information. Note that configuring Ignition Server requires knowledge of your network's IP addressing topology. You need:

- the Avaya Identity Engines Ignition Server product CD shipped with your Ignition Server
- a personal computer or workstation running Windows 2000/XP/2003
- the standard default Ignition Server administrator name ("admin") and password ("admin")
- an IP address and subnet mask you can assign to each Ignition Server network port. See "Ignition Server Ports and Allowed Traffic" on page 60. Each address must be reachable from all authenticators.
- the IP address of your enterprise DNS server(s)
- the IP address of your enterprise syslog server, if available
- the list of network devices (switches, wireless access points, and VPN switches) that you want to secure with Ignition. These are modeled as *authenticators* in Ignition. For each authenticator, note its IP address, shared secret, vendor, model, and authenticator type (wired switch, wireless access point, or VPN).
- the list of LDAP-accessible directory servers that Ignition Server uses to authenticate users and retrieve user records. For each directory server, note its IP address, port number, connect user name and password, user root DN, and directory root DN.
- access to your enterprise certificate authority.

Map Out Your Ignition Server Deployment

To map out your production deployment, you need the types of information listed below. If you are performing a basic installation, you need not gather this information now. Instead, follow the steps in "Install Ignition Dashboard on Microsoft Windows", below.

- the Network Topology Diagram, to clarify Administration and Authentication traffic
- the access policies you want to define,
- the authenticators that each access policy is to protect,
- the protocols they are to use for RADIUS authentication, and
- the credential validation protocols they are to use for secure verification.

You must configure your authenticators to use Ignition Server as a RADIUS Server.

VMware ESXi Server

Hardware platforms supported by VMware's ESXi Servers versions 4.x and 5.0 are supported. The VM requires an x86_64 capable environment, a minimum of 2 GB of memory, 30 GB of available disk storage, two CPUs, at least one physical NIC card (preferably three NICs), and three Logical NIC cards. VMware lists on its site supported hardware platforms for ESXi. (http://www.vmware.com)

Installation on a VMware ESXi server is done using an OVF file which already incorporates the OS Red Hat Enterprise Linux.

Reminder: Avaya provides the Identity Engines Ignition Server and Ignition Access Portal as Virtual Appliances. Do not install or uninstall any software components unless Avaya specifically provides the software and/or instructs you to do so. Also, do not modify the configuration or the properties of any software components of the VMs (including VMware Tools) unless Avaya documentation and/or personnel specifically instructs you to do so. Avaya does not support any deviation from these guidelines.

Warning! Do not install or configure VMware Tools or any other software on the VM shipped by Avaya:

- Avaya does not support manual or automated VMware Tools installation and configuration on Avaya supplied VMs.
- Turn off automatic VMware Tools updates if you have enabled them. Refer to the instructions below to disable automatic updates and to check if you have accidentally installed VMware tools.
- Avaya determines which VMware Tools to install and configure. When required, Avaya provides these tools as part of the installation or package upgrade procedures. Avaya provides these tools because VMware Tools configures the kernel and network settings and unless Avaya tests and approves these tools, Avaya cannot guarantee the VM will work after the tool is installed and configured.
- Avaya does not support the installation of any VMware specific, RHEL specific, or any third party vendor package or RPM on its VM other than what Avaya ships as a package, image, or OVF.

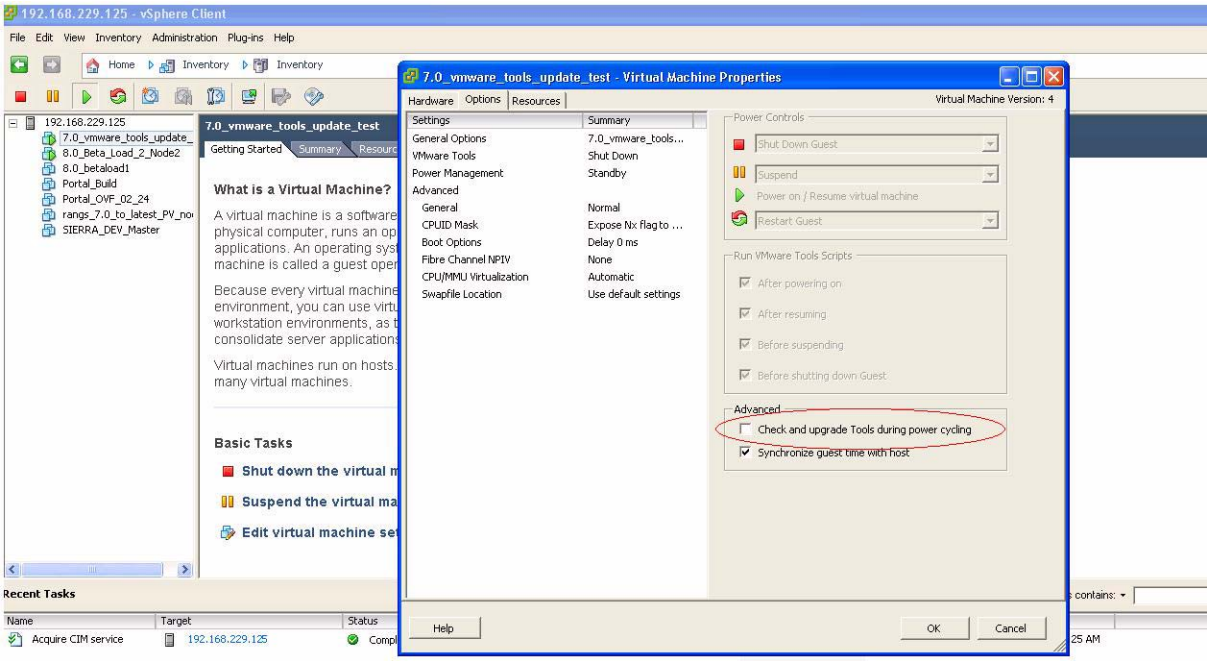
Preventing automatic VMware Tools updates:

To prevent automatic VMware Tools updates:

1. Use the VI Client to log in to the ESXi Server hosting the Ignition VM.

2. Go to **Getting Started > Edit Virtual Machine Settings > Options > VMware Tools > Advanced**, and ensure the **Check and upgrade Tools during power cycling** check box is not selected. This is the supported setting.
3. Click **OK**.

Figure 76 Preventing automatic VMware Tools updates



Checking the VMware Tools status (ESXi 4.1)

The **Summary** tab of the VM describes the VMware Tools status.

To check the VMware Tools status on an ESXi 4.1 server:

1. Use the VI Client to log in to the ESXi Server hosting the Ignition VM.
2. Go to the **Summary** tab.

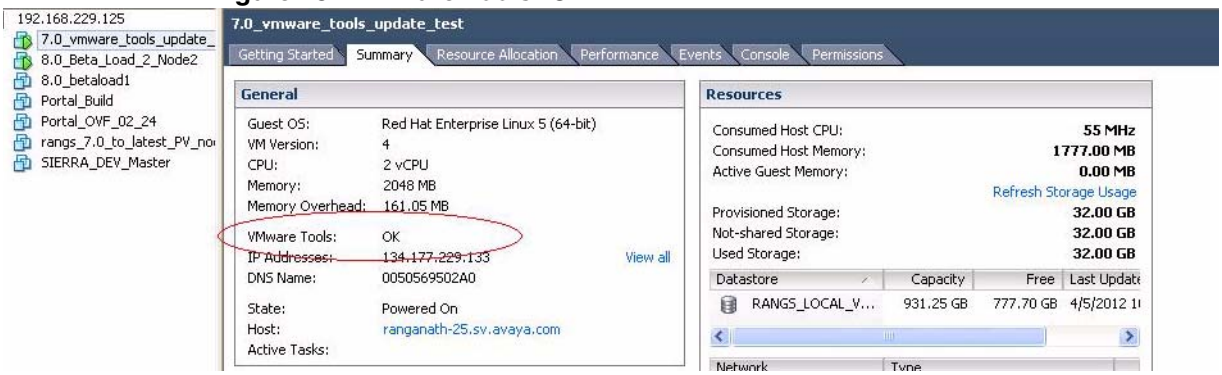
If you are using the vmware-tools supplied by Avaya and did not upgrade, the status displays as “VMware Tools: Out of date”.

Figure 77 VMware Tools: Out of date



If you upgraded the VMware Tools, the status displays as “VMware Tools: OK”.

Figure 78 VMware Tools: OK



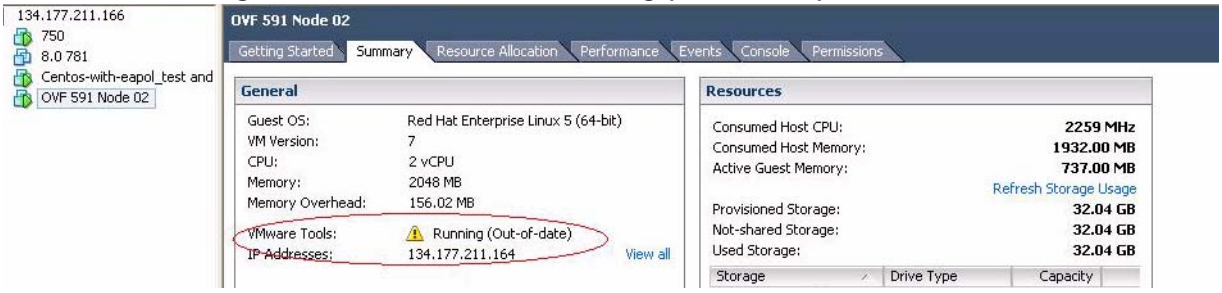
Checking the VMware Tools status (ESXi 5.x)

To check the VMware Tools status on an ESXi 5.x server:

1. Use the VI Client to log in to the ESXi Server hosting the Ignition VM.
2. Go to the **Summary** tab.

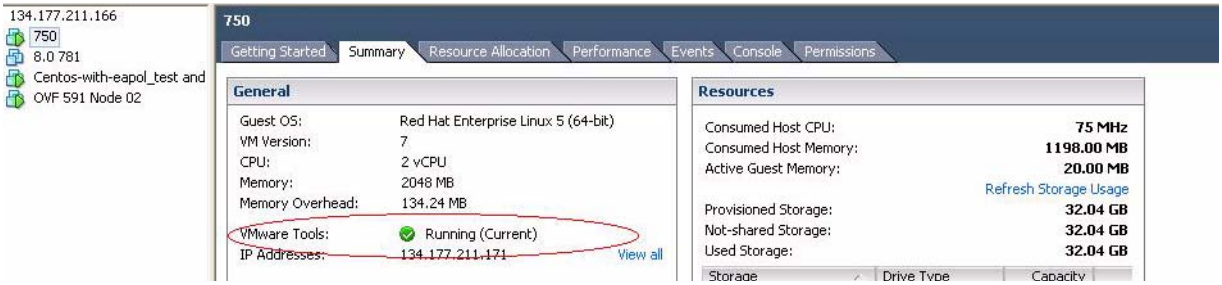
If you are using the vmware-tools supplied by Avaya and did not upgrade, the status displays as “VMware Tools: Running (Out-of-date).

Figure 79 VMware Tools: Running (Out-of-date)



If you upgraded the VMware Tools, the status displays as “VMware Tools: Running (Current)”.

Figure 80 VMware Tools: Running (Current)

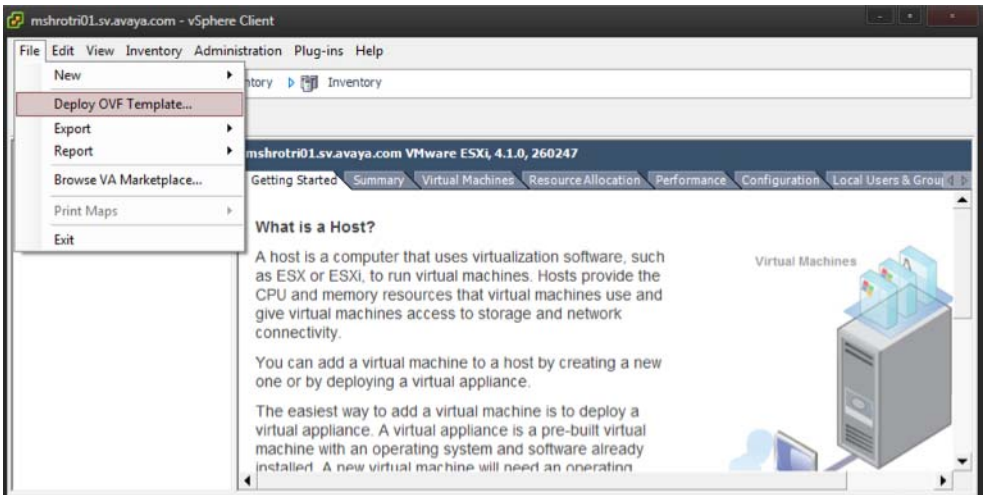


Importing VM

Avaya recommends that you use the VMware vSphere Client to import the VM into your system. Start the VMware vSphere Client and log in to the ESXi Server you want to install the Avaya Ignition Server on. You will need to use the Virtual Appliance Deploy OVF Template option.

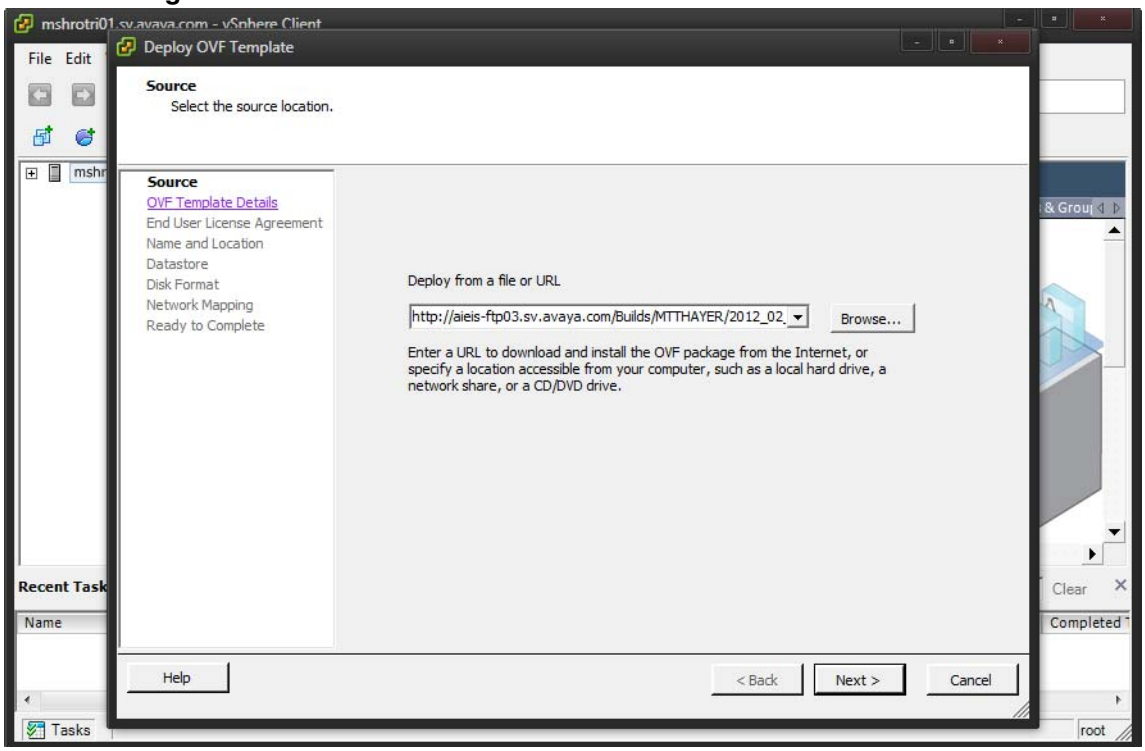
1. From the VSphere Client, select **File > Deploy OVF Template**.

Figure 81 Deploy OVF Template



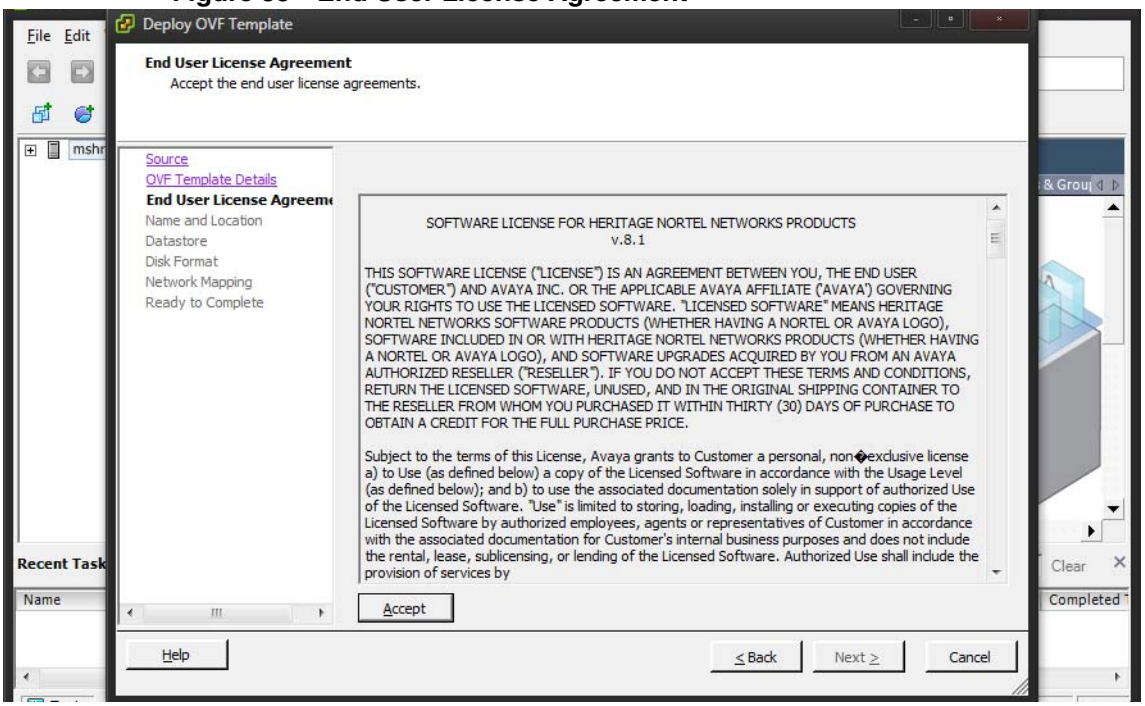
2. The Source screen appears. Select the location from which you want to import the Ignition Server virtual appliance.

Figure 82 Source



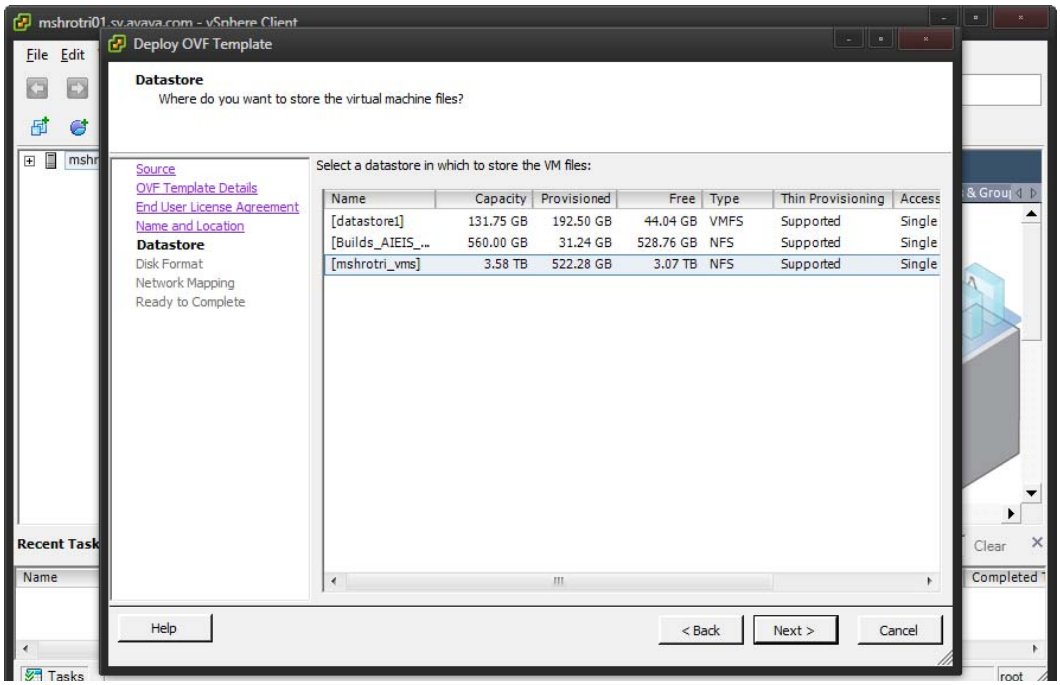
3. Click **Next**. In the OVF Template Details screen, review your settings. You can click **Back** to make changes, or click **Next** to continue.
4. The **End User License Agreement** screen appears. Click **Accept** to accept the license and click **Next**.

Figure 83 End User License Agreement



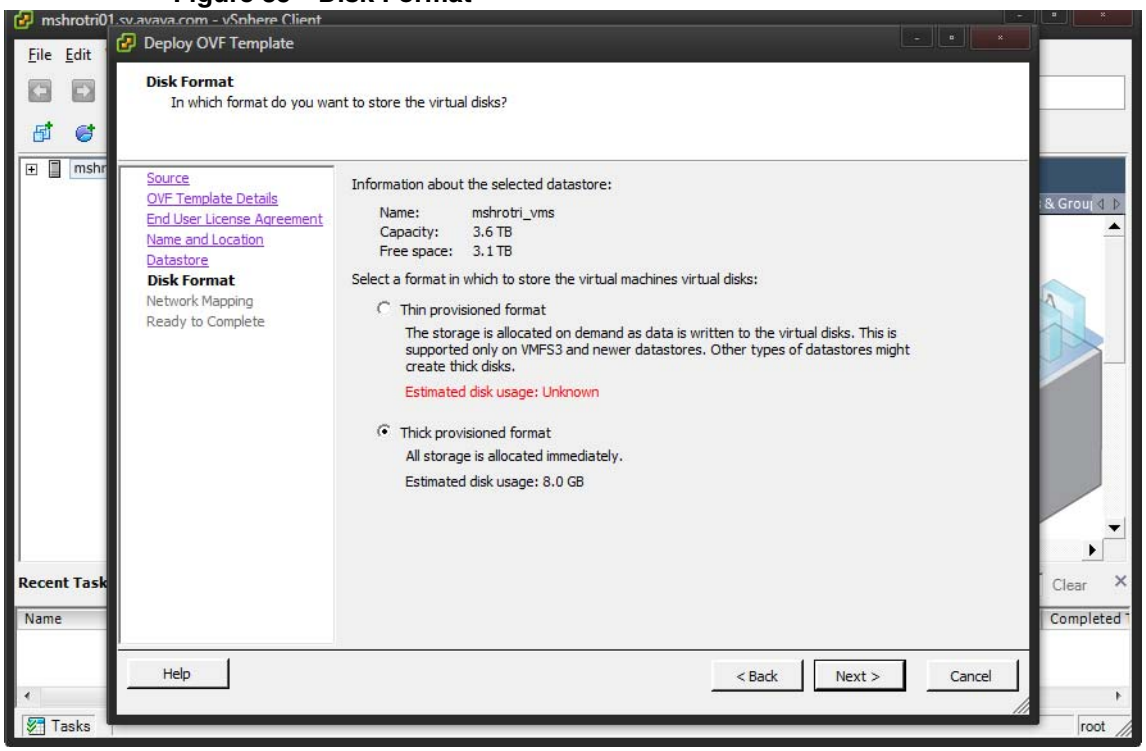
5. The **Name and Location** screen appears. You can either accept the default name or choose to rename the virtual machine. Click **Next**.
6. The **Datastore** screen appears. Select the location where you want to store the files for the virtual appliance and click **Next**.

Figure 84 Datastore



7. The **Disk Format** screen appears. Select a format in which to store the virtual machine's virtual disks and click **Next**.

Figure 85 Disk Format



8. The **Network Mapping** screen appears. Associate the Avaya Ignition Server NIC's to correct VM Network based on your site configuration. Then click on **Next**.
9. The **Ready to Complete** screen appears. Review your settings. Use the **Back** button to make any changes or click **Finish** to start the import.

   The Import now starts. Once the import completes you should see a **Summary** window display.
10. After the import completes you need to verify and adjust some of the VM settings. Open up VM setting dialog and select the Options tab. Do the following:
    a. Ensure to click the **Synchronize guest time with host** option.
    b. Change the System Default Power Off from **Power off** to **Shutdown Guest**. Click **OK**.
    c. Open the VM setting dialog and select the Hardware tab. Adjust the **Network Adapter (1/2/3)** settings and configure the right NIC for each interface.

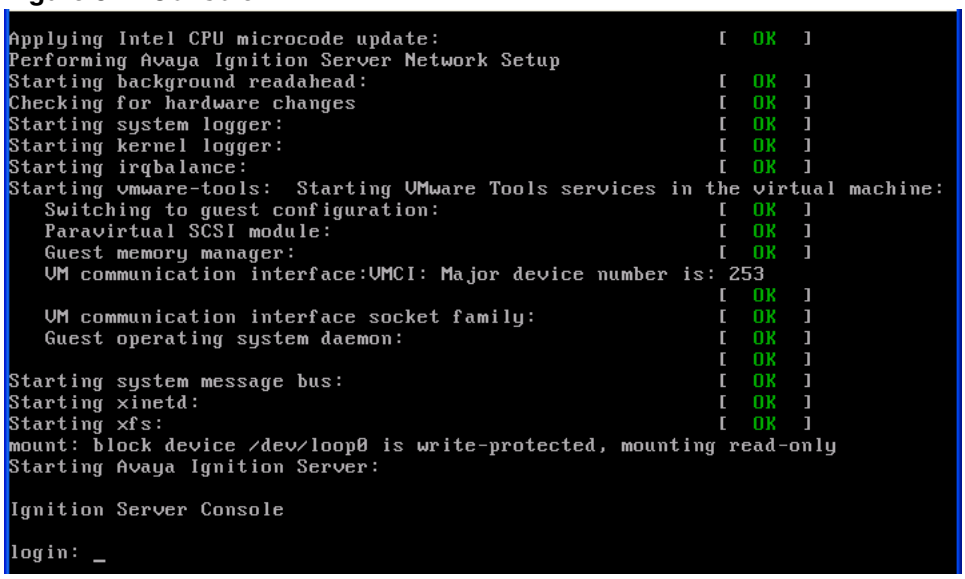
You are now ready to boot the Avaya Ignition Server for the first time. You will a splash screen displayed as the boot up starts.

Figure 86 Boot up



11. Once the Ignition Server Console login prompt is shown you are ready to enter the administration IP address. Login using admin for the user name and admin for password, you should change the password after you login.

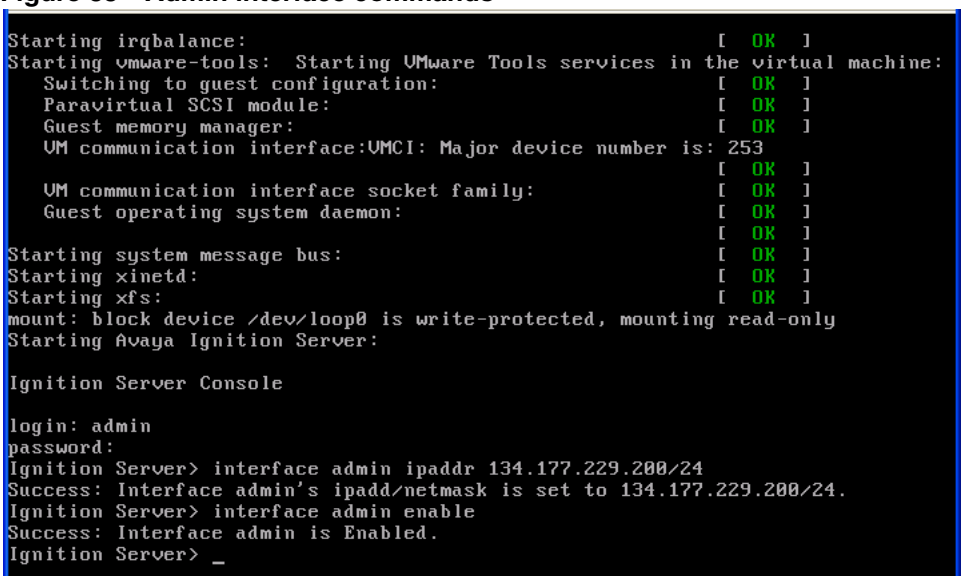
Figure 87 Console



12. Use the interface commands as shown in the next screen to configure the admin interface.

- Only Static IP configuration is supported.
- Configure your admin interface with an IP address.
  Cli command example: “interface admin ipaddr x.y.z.x/netmask”
- If needed, configure your default route.
  Cli command example: “route add 0.0.0.0/0 <gw-ip>”

Figure 88 Admin interface commands



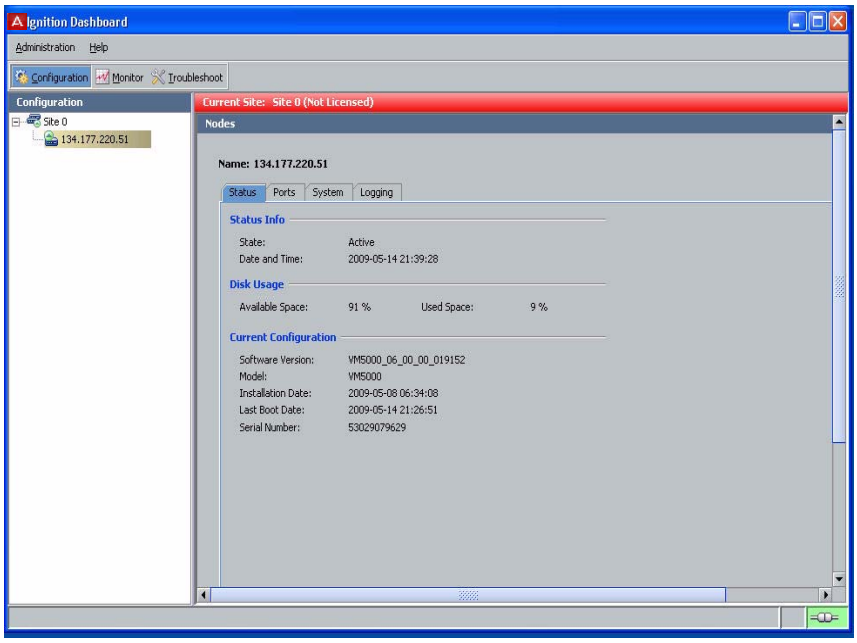
13. Install the Dashboard on to your Desktop machine, see "Install Ignition Dashboard on Microsoft Windows" on page 362.
14. Once installation is complete click on the desktop icon to start the application. A **Login** dialog displays.
15. Enter in the same IP address that you used for the admin interface. The default password is admin if you have not already changed it on the Ignition Server. If you have not configured the admin certificate or the base license you will see the following message.

Figure 89 Default Certificate



If you click **OK** to both dialogs you see a display similar in the following window.

Figure 90 Ignition Dashboard



In order to obtain your license you will need to perform the following steps in the License Management in VWware which follows. Once you have obtained your license you can proceed with the final configuration of the Avaya Ignition Server in your environment.

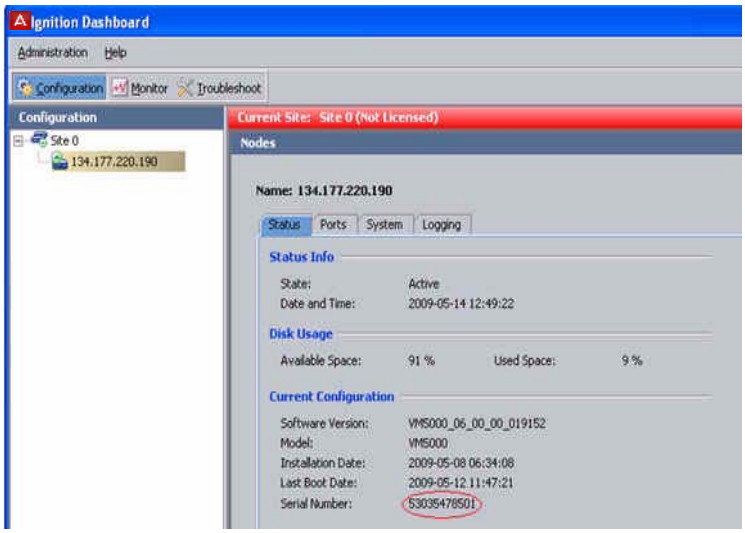
Applying the license

The Avaya Identity Engines Ignition Server (AIEIS) Software ships without any licenses. There are six different software licenses that can be installed on Ignition Server: Base License, Guest Manager License, NAP Posture License, TACACS+ License, Ignition Reports License, and Access Portal License. At a minimum, you must obtain the Base License to be able to configure and run the server.

Note: Select the Access Portal License that matches the Ignition Server Base License (lite, small, or large).

1. Contact Avaya Customer Support. Avaya provides a telephone number for you to use to report problems or to ask questions about your product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: **http://www.avaya.com/support**.
2. Once this is purchased, Avaya Customer Support sends a software CD and certificate that contains a unique product code and an e-mail address. Send this unique product code and the Node Serial Number to the e-mail address provided. The Node Serial Number can be obtained from Dashboard from the Status tab of Node Configuration as shown in the figure below:

Figure 91 Apply license



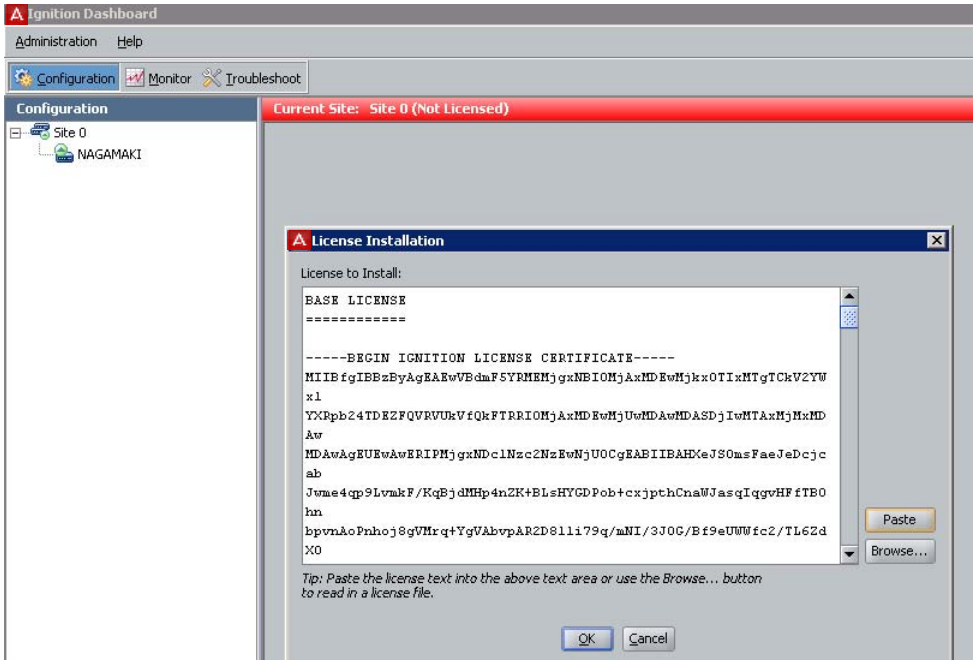
3. After the unique product code and Node Serial Number is verified, a software license file is sent back to you. Install this license on the server using Dashboard.

Installing the license

You can install the license on the Ignition Server using Dashboard. To install the license, perform the following steps:

1. Select the **Configuration** tab.
2. Select the Site.
3. Select the **Licenses** tab.
4. Click on **Install**….
5. Paste the license text and click **OK**.

Figure 92 Install license



Install Ignition Dashboard on Microsoft Windows

Ignition Dashboard is a desktop application that lets you manage the Ignition Server appliance. Use this interface to create, view, or alter configuration information for authenticators, service categories, and the policies that apply to authentication and authorization.

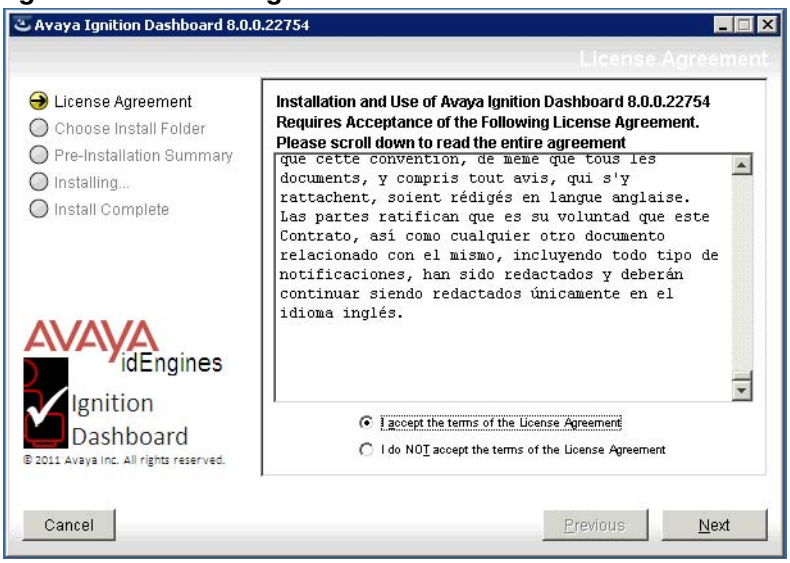
To proceed with the Ignition Dashboard installation, have the following tools and information ready:

- the Identity Engines product software shipped with your Ignition Server appliance

- a computer running Windows XP Service Pack 3 (32 bit), Windows 7 (32 bit and 64 bit), Windows Server 2003 (32 bit and 64 bit), or Windows Server 2008 (32 bit and 64 bit)
- a minimum of 2 GB of memory
- the default Ignition Server administrator name (`admin`) and password (`admin`)

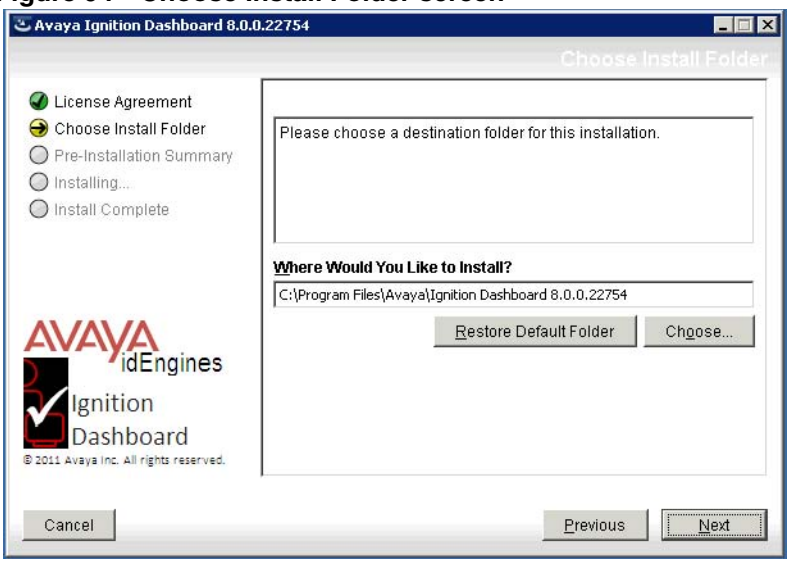
1. If any version of Dashboard is already installed on the computer, make sure the Dashboard application is not currently running. If Dashboard is running, shut it down now.
2. Place the Ignition Server CD into the CD drive of your computer, and:
   - On Windows, the Windows AutoRun feature will run the Installer immediately. (If the AutoRun feature is disabled on your computer, navigate to your CD drive and double-click the installer file. It has a name like `DashboardInstaller-8.0.0<Build Number>.exe`.
3. In the License Agreement screen, scroll down to read the entire license. Select the radio button to accept the license and click **Next**.

Figure 93 License Agreement screen



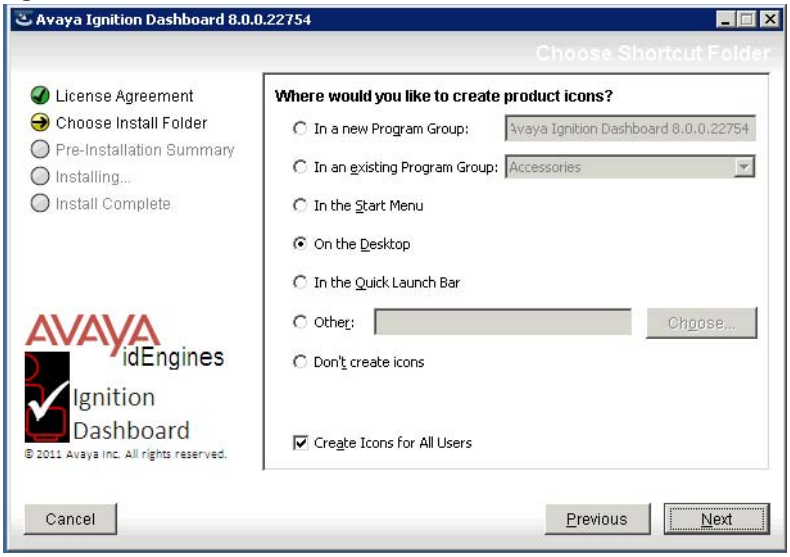
4. In the Choose Install Folder screen, choose your destination folder and click **Next**.

Figure 94 Choose Install Folder screen



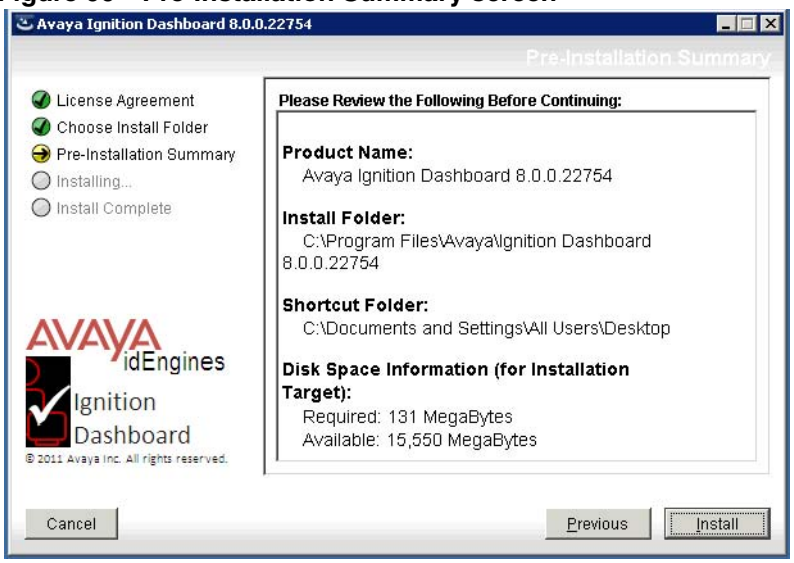
5. In the Choose Shortcut Folder screen, indicate where you would like the Dashboard shortcut to appear, and click **Next**.

Figure 95 Dashboard shortcut location



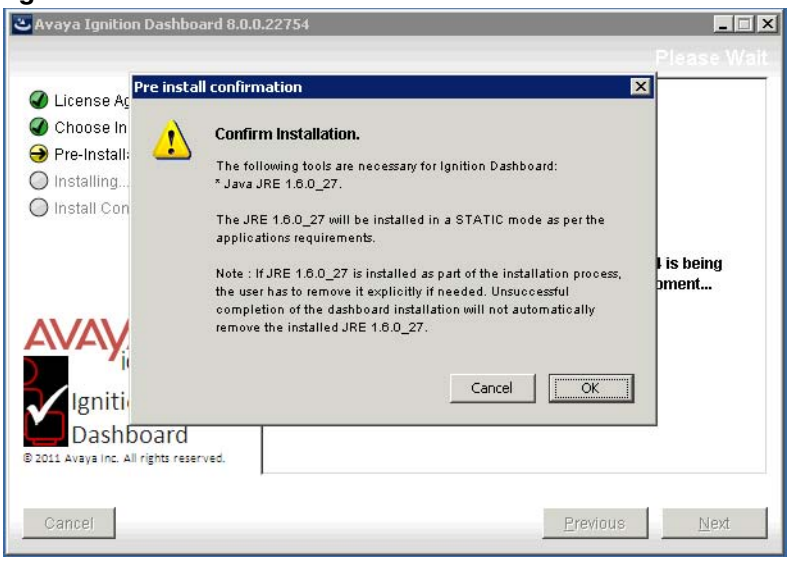
6. In the Pre-Installation Summary screen, review your installation settings. If you wish to make changes, click the **Previous** button to edit the details of the locations of the installation. When you are satisfied with the setup you have chosen, click **Install**. The installer displays a pre install confirmation window.

Figure 96 Pre-Installation Summary screen



7. In the Pre install confirmation window, click **OK** to confirm the installation.

Figure 97 Pre install confirmation window



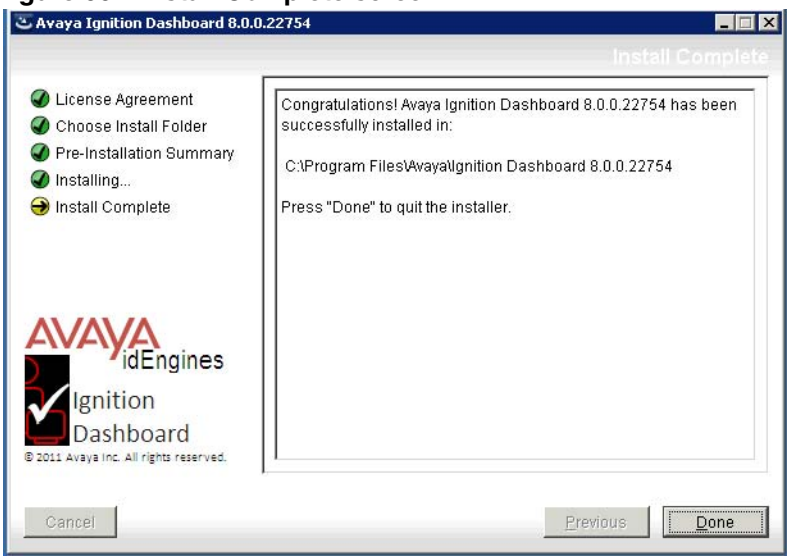
The installation starts. The installer displays a dialog box showing the progress of the installation.

Figure 98 Installation progress



8. When the installation is complete, the installer displays the Install Complete screen. In the Install Complete screen, click **Done**. An icon for Ignition Dashboard will appear in the location you designated.

Figure 99 Install Complete screen



Next Steps. See “Connect to Ignition Server for the First Time” on page 367.

Tip — Installing multiple versions of the Ignition Dashboard: You may install multiple versions of Ignition Dashboard on a single workstation. When you run the installer, it installs the new version in its own folder. The new

installation does not interfere with existing Ignition Dashboard installations and creates a new icon to launch the new version of Ignition Dashboard. The installer leaves the existing Ignition Dashboard installation and icon intact.

To uninstall an earlier version, follow the instructions in “Uninstalling Ignition Dashboard” on page 368.

Connect to Ignition Server for the First Time

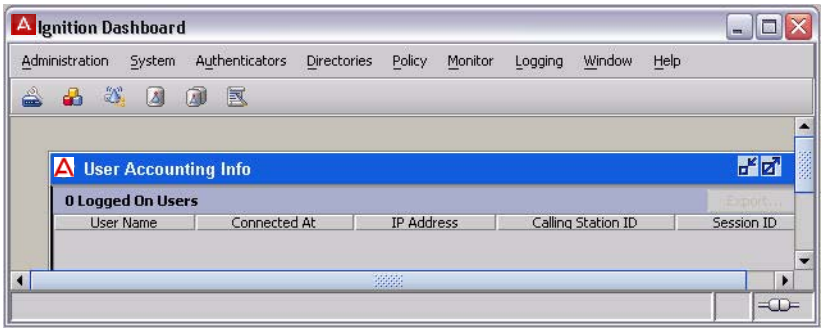
Run Ignition Dashboard

Launch Dashboard as follows:

1. On your personal computer or workstation, start Ignition Dashboard by double-clicking its icon on the desktop. This displays the login window.
2. Enter the administrator’s **User Name** and **Password**. The default user name and password are “admin” and “admin”.
3. In the Connect To field, choose the name or IP address or your Ignition Server node, or choose the name of your Ignition Server site.
4. Click **OK**.
   - *If your login attempt fails*: See “Problem: Cannot Connect to Ignition Dashboard” on page 452.
   - *If your login attempt succeeds*: You will see a warning dialog reminding you to replace the default certificate shipped with the Ignition Server. For instructions on replacing the certificate, see “Replacing the Admin Certificate” on page 80.

     After you dismiss the warning dialog, Ignition Dashboard appears.

Figure 100 Ignition Dashboard Main Window



Change the Administrator Password

Avaya recommends you change the default password when you first set up your Ignition Server. To change the administrator password, see “Setting the Ignition Server Administrator Password” on page 49.

Next steps for installers: Your Ignition Server installation is complete. Turn to "Configure the Ignition Server" below for a list of configuration options.

Install Your Ignition Server Licenses

See "Installing an Ignition Server License" on page 69.

Configure the Ignition Server

Using Dashboard, you can set up access control for your networks. The access policy settings and corresponding chapters are listed below.

Settings To Be Configured in Ignition

Object	Reference Chapter
Authenticators: Representing wireless access points, groupings of wireless access points, or network devices such as VPN concentrator or server, Ethernet switches, WLAN switches, or routers.	Chapter , "Authenticators"
Directory services: Repositories of user identities and attributes such as Active Directory, LDAP, and token servers. **Directory sets:** Groups of directory services	Chapter , "Directory Services"
Users, groups, and **attributes:** Entities or objects represented in directories or databases, that contain information about end users.	Chapter , "Internal Users, Groups, and Devices"
Authentication and authorization policies: Each access policy establishes a set of rules that governs user access. Rules are evaluated based on user attributes and other criteria.	Chapter , "User Authentication Policy"
Provisioning policies: Optionally, each access policy may have a provisioning policy that assigns each user to an appropriate VLAN and/or sets switch parameters for the user.	Chapter , "Provisioning Policy"

Uninstalling Ignition Dashboard

To uninstall Ignition Dashboard:

1. Make sure the Dashboard application is not currently running. If Dashboard is running, shut it down now.
2. Launch the uninstaller using one of following commands:
   - From the Windows desktop **Start** menu, select: **Start → Programs → Ignition Dashboard**
   - From the Windows Control Panel, select **Add or Remove Programs**. In the Add or Remove Programs window, click on the row for the version of Ignition Dashboard you want to remove, and click **Change/Remove**.

The Ignition Server installer asks you to confirm your intention to remove Ignition Dashboard. If you do, the components of the selected version of Dashboard are removed. If other versions of Ignition Dashboard are installed on the PC, they are left intact.

B High Availability Configuration

Any two Avaya Identity Engines Ignition Servers can be connected to run as a high-availability (HA) pair. The HA pair can be configured to provide a highly available IP address (a virtual interface or “VIP”) that serves RADIUS authentication requests and/or SOAP API requests. After you have paired two Ignition Servers, you can manage both from a single Dashboard session.

Note: Ignition Server does not support moving VMs with vMotion.

Contents

This appendix explains how to create and manage HA pairs. It is organized as follows:

- “HA Terminology”, below
- “Overview of HA Pairs” on page 372
- “Creating an HA Pair” on page 373
- “Restoring a Saved VIP Configuration” on page 381
- “Managing an HA Pair” on page 383
- “Backing Up Data on an HA Pair” on page 389
- “Restoring Data on an HA Pair” on page 389
- “Updating Firmware on an HA Pair” on page 390
- “Replacing an Ignition Server in an HA Pair” on page 390
- “Changing the IP Address of the Admin Port or HA Port in an HA Pair” on page 392
- “Restoring a Non-Responsive Unit in an HA Pair” on page 394

HA Terminology

This document uses the following terms:

- An **authenticator** is a network device, usually a switch, wireless access point, VPN concentrator, or other 802.1X-compliant device, that authenticates a user or device against Ignition Server (the RADIUS server) and allows or denies network access.

- An **HA pair** is a connected pair of Ignition Server appliances that remain in sync and offer highly available RADIUS and/or SOAP services. In Ignition Dashboard, an HA pair is sometimes called a **site**.
- A **node** is one Ignition Server appliance in the pair.

Overview of HA Pairs

After two Ignition Servers are connected in an HA pair, Ignition Server ensures that the best suited Ignition Server acts as the primary provider of each service. The Ignition Server acting as *database primary node* serves configuration requests (it handles the data changes the administrator submits). The Ignition Server acting as *VIP primary node* handles all client requests for the Ignition Server service bound to that VIP. For example, your RADIUS VIP handles all authentication/authorization requests.

Each Ignition Server is referred to as a “node” in the pair. A node designated as the secondary node for a service acts as a warm backup for that service. If the primary node handling that service fails or is taken offline, the secondary node takes over and provides the service.

The processing of RADIUS and/or SOAP traffic by the Ignition Server fails over seamlessly, since clients connect to a VIP for the service, rather than a physical Ethernet interface on a specific Ignition Server. (Note: If you plan to perform RSA Secured authentication, please see “Warning for Sites Running Ignition Server in HA Mode” on page 181.)

The relationship between paired nodes is as follows:

- **Data replication:** Data replication between nodes is automatic. The administrator manages users, authenticators, and policies just as he would in a single- Ignition Server configuration, and the nodes synchronize automatically.
- **Logging**: Each Ignition Server handles its own logging, system statistics, and trouble tickets. The log levels you choose in Dashboard apply to both nodes in the pair. The logged information for both nodes is accessible from Dashboard when you log in to either node in the pair. See the Appendix E, “Setting Up Logging” and Appendix F, “Viewing Logs and Statistics” for instructions.
- **Network settings**: The Ignition Server administrator makes network interface settings on each Ignition Server and then binds a virtual interface (VIP) address to the RADIUS and/or SOAP services. Authenticators are configured to connect to the Ignition Server RADIUS service at the VIP address, and SOAP API clients can be configured to connect at a VIP address. The HA Configuration Wizard ensures proper network setup.

Creating an HA Pair

Follow the steps below to create your HA pair and enable the VIP that provides failover for Ignition Server services. This procedure contains the following sections:

- "Start and connect the Ignition Server", below;
- "Run the HA Wizard" on page 374; and
- "Bind Services to the VIP" on page 380.

Start and connect the Ignition Server

1. Set up and start both Ignition Server ("nodes") as explained in the *Ignition Server Getting Started Guide*. Avaya strongly recommends that you use two virtual machines on two different servers to ensure that the HA can maximizes coverage against ESXi server disk failures.
2. Connect the *Admin ports*:

   Observe these rules when connecting VMWARE ESXi ports: In a typical deployment, you connect both virtual ports to the same L2 switch as this provides support for high-availability environments. Make sure that your port network connections comply with the following rules:

   (a) The two ports must be on the same local network (same broadcast domain) without a layer-3 switch in between so that they can be joined later to form a VIP; (b) the port subnet must be reachable from your authenticators and from your Ignition Dashboard workstation; and (c) The network of the these connection must be a high-throughput, high-reliability, low-latency network as the HA link carries data to be replicated between the Ignition Servers. Disruption in this network might cause replication failures. In order to avoid that, Ignition Server requires a carrier-grade link. Equipment as reliable as the Baystack switches or the passport switches is appropriate.
3. On the first node, use the Ignition Server dashboard to make network settings for the Admin port. Set the IP address, subnet mask and gateway address. For help using the Ignition dashboard, see *Avaya Identity Engines Ignition Server Getting Started Guide*.

   As you set each IP address and subnet mask, make a note of it. You need this information in order to run the HA Configuration Wizard.
4. Repeat Step 3 for the other node.
5. Log into the first node using Ignition Dashboard, and make these settings:
   - In Dashboard, click the **Configuration** tab, click the **IP address** of your Ignition Server, click the **System** tab, and click the **DNS** tab. Click **Edit** and set the address.

- Ping the DNS to make sure they are accessible. In Dashboard, click the **Troubleshoot** tab; click your node's IP address or name in the hierarchy tree; click **Network** and go to **Ping Test**; enter the DNS server address as the **Target**; and click **Start**.
- Select **Administration: Logout** to disconnect from the node.

6. Use the command **Administration: Login** to connect to the second node and configure the settings for DNS server addresses there.

Run the HA Wizard

The **HA Configuration** wizard guides you through the steps to create an HA pair. To create a new HA link, follow the steps below:

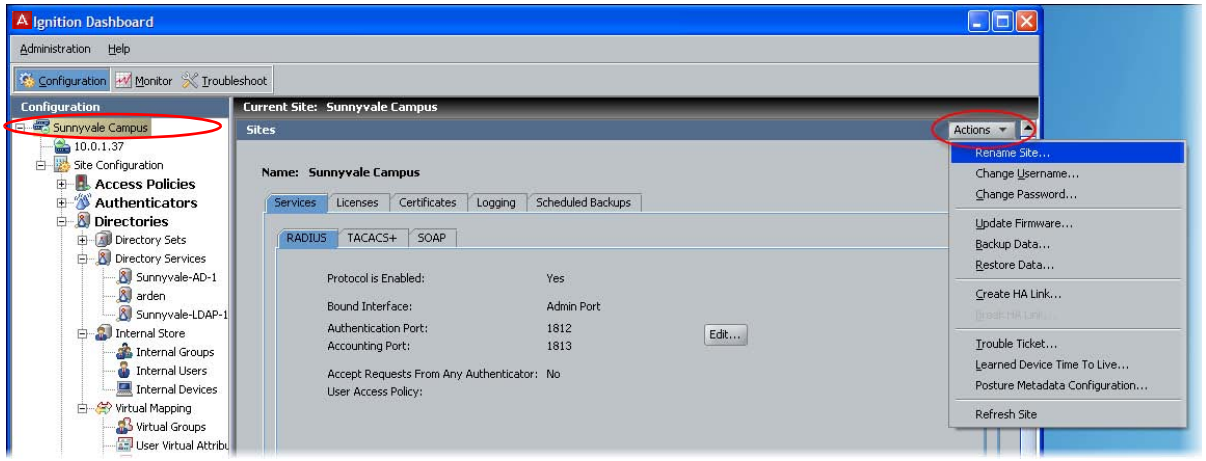
Warnings Before You Proceed:

- **Secondary node data will be erased:** You designate one node as primary and one as secondary for the duration of the set-up procedure. The HA Configuration Wizard replicates data from the primary node to the secondary node, overwriting the data on the secondary node.
- **Start with an unpaired node:** Before you run the Wizard, make sure neither node is a member of an HA pair.

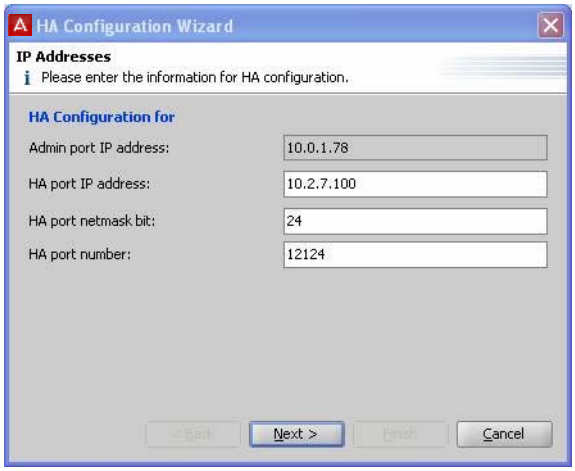
Procedure:

This procedure continues from Step 6 above.

1. Using Ignition Dashboard, log in to either of the Ignition Server that form your HA pair.
2. In Dashboard's Configuration hierarchy tree, click on your site name and select **Actions: Create HA Link**. The **Actions** menu is at the far right of your window.



3. The HA Configuration Wizard appears.

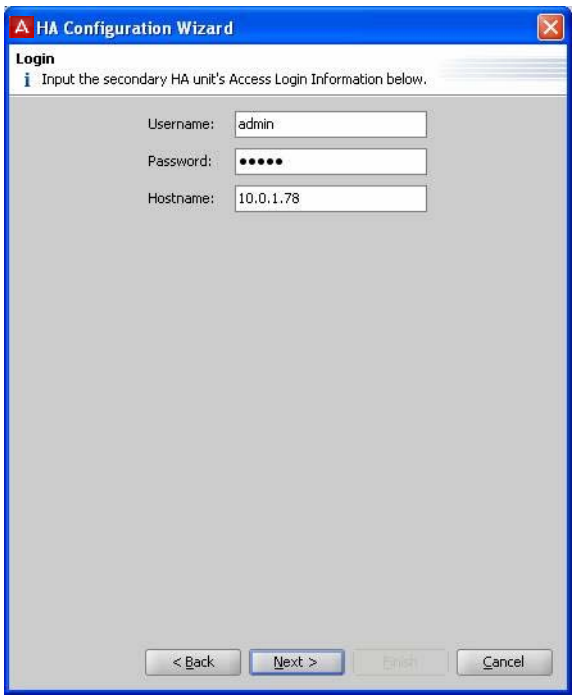


Enter the following information for the Ignition Server that you initially logged into:

- **Admin port IP address:** Ignition Server displays the IP address of the Ignition Server to which you are currently connected.
- **HA port IP address:** Enter the IP address to be assigned to the HA Port of this node. Verify that the address you assign is not in the same subnet as any other port of this Ignition Server. In other words, The HA port must reside on a separate subnet not shared with the Admin port (and not shared with an active Ignition Server Service Port).
- **HA Port netmask bit:** Enter the subnet mask as a bit count.
- **HA port number:** Enter the port to be used for HA traffic.

Click **Next.**

4. The Login Window appears requiring login information for the second node in the HA pair (the second Ignition Server in the pair).

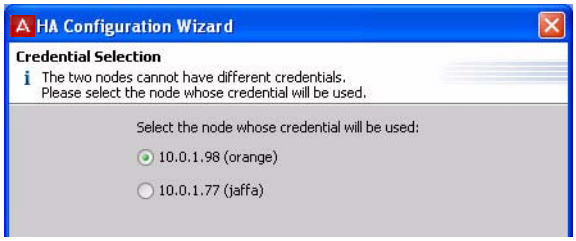


Enter the Ignition Server administrator credentials and hostname of the second node in the HA pair, and click **Next**. Ignition Dashboard logs in to the second node.

5. The HA Configuration Wizard requests information on the second node. Specify the Admin port IP address, HA port IP address, HA port subnet mask, and HA port number, and click **Next**.

6. The HA Configuration Wizard requires you designate the primary node for the newly created HA pair. Select the radio button for the required primary node in the new HA pair, and click **Next**.

7. If prompted, you must specify which node's Ignition Server administrator credentials are to be the administrator credentials after the HA pair is created. Choose a node.



The administrator name and password on both nodes should be set to match those of the node you select. This is the only Dashboard login for the pair of Ignition Server. You can change the administrator password later, as explained in "Setting the Ignition Server Administrator Password" on page 49.

8. The wizard asks whether you want to create a VIP. Tick the **Yes** radio button and click **Next**.



9. Set the VIP settings in the Virtual Interface Configuration window. The settings are explained below, after the illustration. (***Note!*** If a VIP was previously configured on one or both of the nodes you are joining, the Wizard offers you the option of deleting or restoring that VIP configuration. If you do *not* want to restore the VIP, click the **Delete all existing virtual interface definitions...** button and click next. If you want to restore the VIP, see "Restoring a Saved VIP Configuration" on page 381.)



Configure the VIP using these fields:

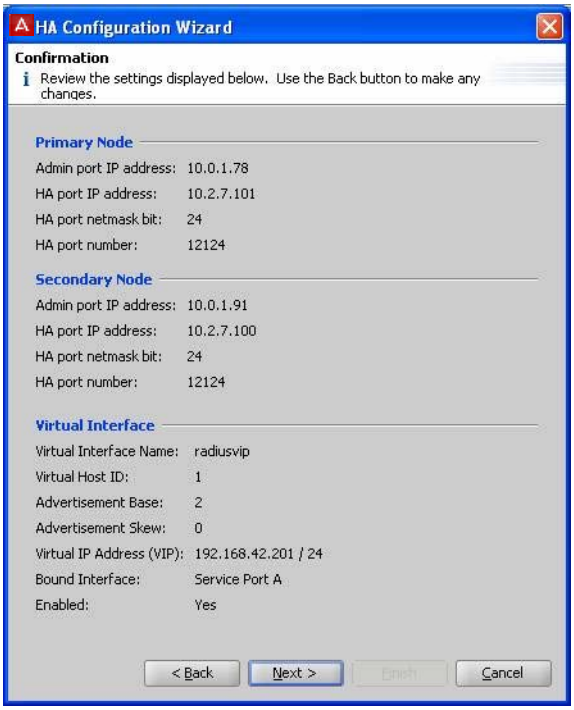
- **Name**: Enter the VIP name to be displayed in the **Virtual Interface** tab of the Sites panel in Dashboard.
- **Virtual Host ID**: Enter an integer between 1 to 255.

- **Password**: Enter a password that the nodes in this virtual interface group should use to secure their communications.
- **VIP IP Address**: Enter the VIP IP address and subnet mask. Use an address on the same subnet as the Service Ports.

  This is the IP address that provides the highly available Ignition Server services (RADIUS and/or SOAP API). This address must be unique; it must not be the address of an Ethernet interface. The virtual IP address must be on the same subnet as the physical interfaces to which it is bound.
- **Bind To**: Select the **Service Port**. The VIP binds to this port on both Ignition Server in the pair. **Important:** Avaya recommends that you bind the VIP to the Service Port. You cannot apply a VIP to the HA port. The VIP is intended to serve RADIUS and SOAP API requests only. You cannot use the VIP address for other traffic, such as, for example, connecting Dashboard.
- **Enabled**: Tick this checkbox to enable the virtual interface.

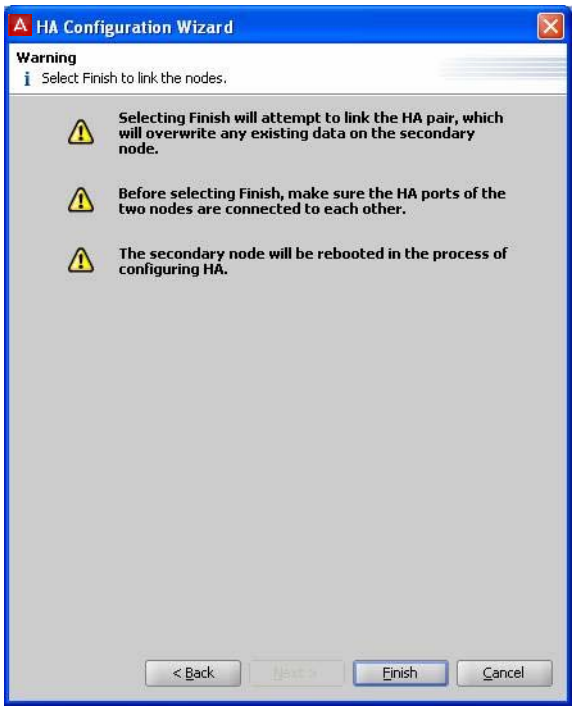
  For complete field descriptions, see Step 4 on page 385. For general VIP information, see "Managing Virtual Interfaces (VIPs)" on page 384.
- Click **Next**.

10. The HA Configuration Wizard requires a confirmation that the settings for the two nodes in the new HA pair are correct.

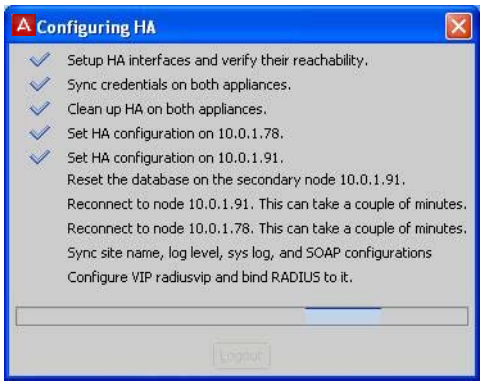


- Review the settings displayed in the **Confirmation** dialog. If a setting is incorrect, use the **Back** button and make changes.
- Click **Next**.

11. The Confirmation window appears. Click **Finish**.



The HA Configuration wizard displays a progress bar while it sets up the HA link. This step might take a few minutes. Wait to see that the pair completes its initial data synchronization, to ensure the setup was successful.

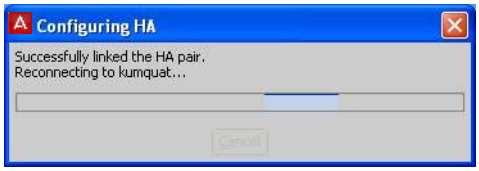


Important! If any of the required network paths between the Ignition Server' ports does not exist, the Wizard displays an error message. For

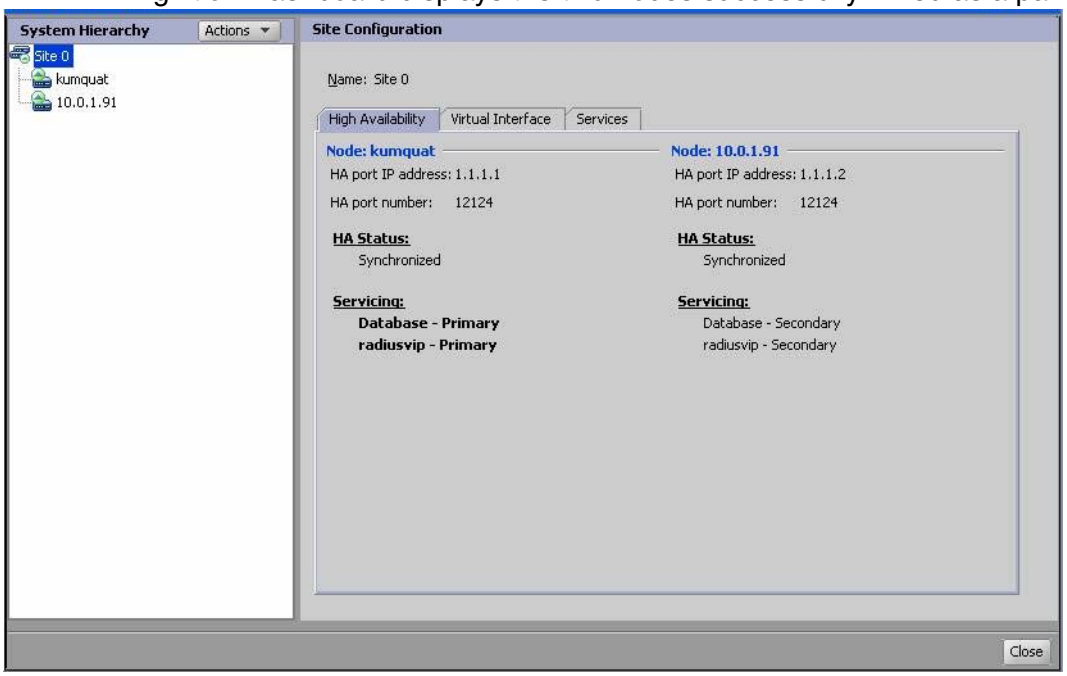
instructions on fixing such problems, see “Problem: HA Set-up Fails” in the Troubleshooting section.

Important! If you provide incorrect settings to the HA set-up wizard, the wizard’s progress bar may appear to freeze. Contact Avaya customer support for assistance. See “Customer service” on page 21.

After setting up the HA link, Dashboard reconnects to the pair.



You manage both Ignition Server from a single Dashboard session. Ignition Dashboard displays the two nodes successfully linked as a pair.



Bind Services to the VIP

Procedure:

Bind the RADIUS and/or SOAP services to the VIP interface as shown below. This procedure continues from Step 11 above.

1. Ping the VIP address to make sure it is accessible. To do this: In Dashboard, click the **Troubleshoot** tab; click either node’s IP address or

name in the hierarchy tree; click **Network** and go to **Ping Test**; enter the VIP address as the **Target**; and click **Start**.

2. In Dashboard's Configuration Hierarchy, click the name of your site (by default, "Site 0").
3. In the Sites Panel, click the **Services** tab and then the **RADIUS** or **SOAP** tab.
4. Click **Edit**.
5. In the **Bound Interface** drop-down list, select the name of your VIP. This is the name you set in the Virtual Interface Configuration window in Step 9.
6. Click **OK**.

Your Ignition Server HA pair set-up is complete. The next steps are:

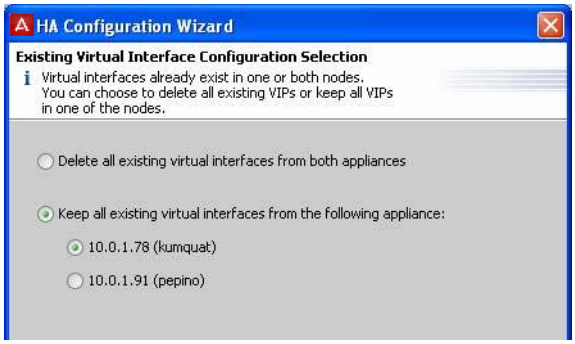
- If you have set up the **RADIUS service** on the VIP port, then you must set up your authenticators to connect to the Ignition Server RADIUS service at the VIP IP address. Consult your authenticator's documentation for details. (Note: If you plan to perform RSA Secured authentication, please see "Warning for Sites Running Ignition Server in HA Mode" on page 181.)
- If you have set up the **SOAP API** on the VIP port, then you must set up Ignition Server Guest Manager to connect to the SOAP service at the VIP IP address. Consult the *Ignition Server Guest Manager Administration Guide* for details.

Restoring a Saved VIP Configuration

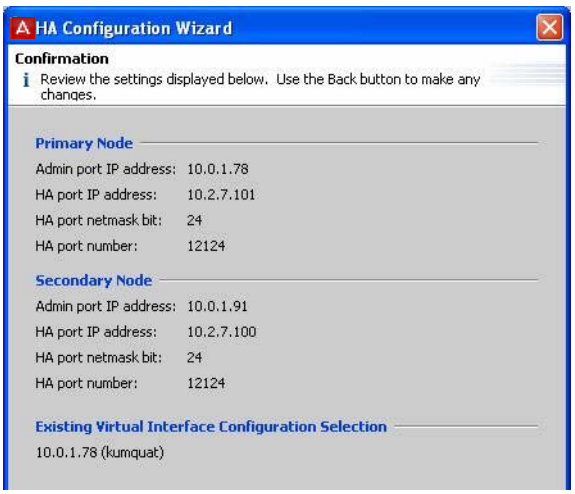
When an HA pair is broken via the CLI, its VIP definition remains on the Ignition Server and can be restored when you create a new HA pair with either Ignition Server.

To restore a saved VIP configuration:

1. Run the HA Configuration Wizard as explained in "Creating an HA Pair", above. After Step 8, the Wizard displays the Existing Virtual Interface Configuration Selection window:



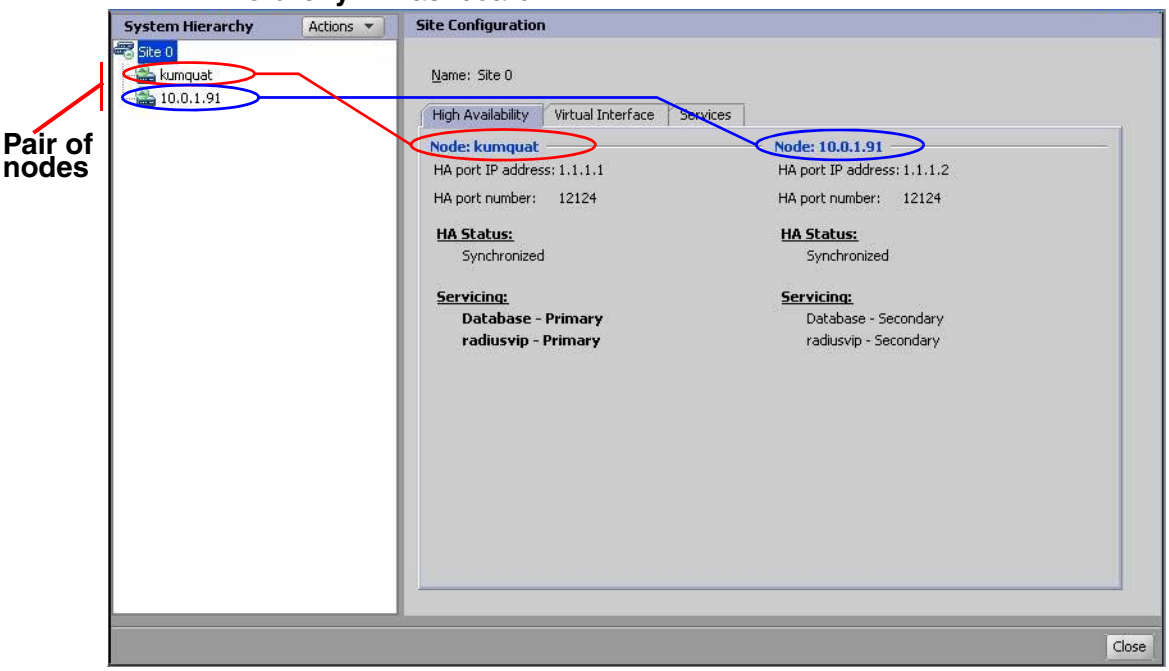
2. Click the **Keep all existing virtual interfaces...** radio button.
3. Click the radio button of the node whose VIP configuration you want to load.
4. Click **Next**.
5. In the **Confirmation** window, review your settings and click **Next** to apply the configuration,



6. Return to Step 1 on page 380 and finish the VIP configuration as instructed there.

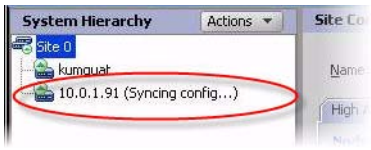
Managing an HA Pair

In order to manage an HA pair and its Ignition Server, use Ignition Dashboard to log in to one Ignition Server in the pair. Dashboard connects to both Ignition Server in the pair. To check the status of the HA pair, check the **Configuration Hierarchy** in Dashboard.



Configuration Hierarchy Panel The **Configuration Hierarchy** tree displays a node icon for each Ignition Server in your pair. The label to the right of a node displays its HA status:

- **No label**: The data on this node is up to date.
- **"Syncing config..."**: Data is currently being copied to this node from the other node.
- **"Disconnected"**: Ignition Dashboard is unable to communicate with the node.





Sites Panel This panel displays:

- The **High Availability** tab displays detailed information about the pair. One node is described in the left column, and the other in the right column.

- **IP Addresses** of the HA Ports of each node, as well as the **HA port number**.
- The **HA Status** section provides details on the HA status of the pair, indicating whether the pair is “synchronized”, “synching configuration”, or “disabled”.
- The **Servicing** section shows which node is currently acting as the database primary and which node is currently acting as the VIP primary in each VIP. Each VIP primary handles all client requests for the Ignition Server service bound to that VIP. For example, your RADIUS VIP handles all authentication/authorization requests.

- The **Virtual Interface** tab provides a summary of your VIP settings and allows you to edit them. For an explanation of VIPs, see “Managing Virtual Interfaces (VIPs)”, below.
- The **Services** tab operates the same as it would in non-HA mode. It displays the RADIUS and SOAP port settings. See “Managing Ignition Server Services” on page 50.

Managing Virtual Interfaces (VIPs)

A virtual interface presents an IP address that your authenticators should use to reach Ignition’s services. The virtual interface address remains valid, regardless of whether Node 1 or Node 2 is currently acting as the VIP-primary node.

For example, assume the Admin Port on your Ignition Server is your RADIUS port. The Service Port on *Node 1* has the IP address 168.172.0.124, and the Service Port on *Node 2* has the IP address 168.172.0.126. To allow a seamless failover from Node 1 to Node 2 in the event of Node 1’s failure, your equipment that communicates with the Ignition Server RADIUS service must be configured to use a virtual, rather than actual, IP address for the service.

To accomplish this, you define a VIP of, for example, 168.172.0.200 for RADIUS. Your authenticators are to be set to reach the RADIUS server at 168.172.0.200, and the VIP ensures they connect to the current, VIP-primary Ignition Server node.



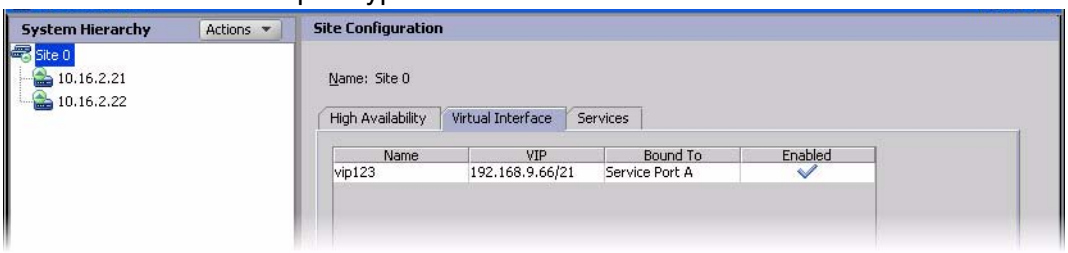
Important! Ignition’s VIP feature relies on the broadcast of gratuitous ARP messages. If your authenticator does not support gratuitous ARP, then failover from the primary Ignition Server box to the secondary Ignition Server box only occurs after your authenticator’s ARP timeout period has elapsed. When using Ignition’s VIP feature, Avaya recommends that you edit the settings of your authenticator (switch or access point), setting the ARP timeout to as short a period as possible.

Viewing VIP Settings

To view the virtual interface settings of your HA pair:

1. In Dashboard's Configuration Hierarchy, click the name of your site (typically "Site 0").
2. In the Sites Panel, click the **Virtual Interface** tab.

The Virtual Interface tab lists the virtual interfaces defined for your HA pair, and shows the port type to which each virtual interface is bound.



A virtual interface definition comprises:

- **Name:** The name given to this virtual interface. The name is an easy way for you, the administrator, to refer to the virtual interface definition.
- **VIP:** The virtual IP address and subnet mask for this virtual interface. This is the IP address that your authenticators use to reach Ignition's RADIUS and/or SOAP API service. The virtual interface address remains valid, regardless of whether Node 1 or Node 2 is currently acting as VIP-primary.
- **Bound To:** The physical Ignition Server Ethernet port (usually the Admin Port) to which this virtual interface is bound. Avaya recommends binding a VIP to Service Port A. You cannot bind a VIP to the HA port; Dashboard does not permit you to do so.
- **Enabled**: Indicates whether the virtual interface is currently enabled.

When, for any reason, one of the nodes is not available, Ignition Server uses these settings to maintain a seamless connection by switching to the other node in the HA pair.

Adding a VIP

Follow the steps below to add a virtual interface. **Note:** Do not create more than one VIP on your system.

1. Make sure your Ignition Server are connected and running as an HA pair.
2. In Dashboard's Configuration Hierarchy, click the name of your site (typically "**Site 0**"), and in the **Sites** Panel, click the **Virtual Interface** tab.
3. Click the **Add** button in the **Virtual Interface** tab.
4. Enter the information as required:

- **Name**: Enter a unique name for this virtual interface group. This name is displayed in the Virtual Interface tab Ignition Dashboard for quick referral.
- **Virtual Host ID**: Enter a unique ID number for this virtual interface group. Acceptable values are from 1 to 255.
- **Password**: Enter a password that the nodes in this virtual interface group use to secure their communications.
- **IP Address**: Enter the virtual IP address and subnet mask for this virtual interface group. This is the IP address at which your authenticators reach Ignition's RADIUS and/or SOAP service. This address must be unique; it must not be the address of another VIP or Ethernet interface. The virtual IP address can be on the same subnet as the physical interfaces to which it is bound, but it must not conflict with other subnets.
- **Bind To**: Select the Ignition Server Ethernet port to which this virtual interface is be bound. The VIP binds to this port on both Ignition Server in the pair. **Important:** Avaya recommends that you bind the VIP to the Service Port. You can bind it to Admin Port A. You *cannot* apply a VIP to the HA port.
- **Enabled**: Select this checkbox to enable the virtual interface. If you want to disable the VIP (for example, for troubleshooting) uncheck this checkbox.

5. Click **OK**.

Ignition Server binds the virtual interface with your settings to the selected port on the both the nodes in the HA pair.

Editing a VIP

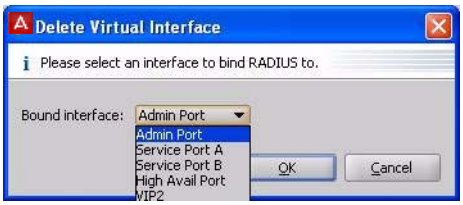
To edit the settings of an existing virtual interface:

1. In Dashboard's **Configuration Hierarchy**, click the name of your site (typically "**Site 0**"), and in the **Sites** Panel, click the **Virtual Interface** tab.
2. Select the VIP entry in the **Virtual Interface** tab. (See the illustration in "Viewing VIP Settings" on page 384.)
3. Click **Edit**. The selected virtual interface's settings appear.
4. Edit the fields as needed. For field descriptions, see Step 4 on page 385.
5. Click **OK**.

Ignition Server updates the configuration of the virtual interface for both the nodes that are linked on your site (Ignition Server), in the internal data store, and the display seen in the Sites panel.

Deleting a VIP

1. In Dashboard's **Configuration Hierarchy**, click the name of your site (typically "**Site 0**"), and in the **Sites** Panel, click the **Virtual Interface** tab.
2. In the **Virtual Interface** tab (as illustrated in "Viewing VIP Settings" on page 384) select the VIP entry you plan to delete.
3. Click **Delete**.
   - If the VIP is bound to the Ignition Server RADIUS or SOAP service, a window prompts you to designate a new interface to carry the service. The dialog displays Admin Port, Service Port A and all VIPs except the one to be deleted.



4. After a new RADIUS and/or SOAP service port has been designated, or if the services are unaffected by the edit, a window asks you to confirm you want to delete. Click **OK** to delete the VIP.



Breaking an HA Pair Using Dashboard

When you break the link between two nodes that are an HA pair, Ignition Server makes them both standalone nodes. User and policy data on the two machines remains on the Ignition Server, allowing you to reconfigure the two Ignition Server as you require. When you use Dashboard to break the HA pair, your VIP definitions are deleted. If you want to retain your VIP definitions, see "Breaking an HA Pair Using the CLI", below.

In order to break an HA link,

1. Using Ignition Dashboard, log in to either of the Ignition Server that form your HA pair.
2. Select your site in Dashboard's Configuration hierarchy tree.
3. Right-click on the selected entry for the site and select **Break HA Link**. Or, Select the **Actions → Break HA Link** command.
4. The **Break HA Link Confirmation** dialog box appears requiring confirmation. Click **Yes** to confirm.

The HA pair is broken, the HA and VIP configurations are deleted, and Dashboard disconnects from the secondary node.

Breaking an HA Pair Using the CLI

When you break an HA pair using the Ignition Server command line interface (CLI), the VIP definitions are maintained (but inactive) on both Ignition Server.

In order to break an HA link from the CLI,

1. Use a console terminal to run the Ignition Server CLI and log in to either of the Ignition Server that form your HA pair.
2. Run the “ha break” command:

```
Identity Engines>  ha break
```

3. Open a second console terminal and run the Ignition Server CLI on the other Ignition Server.
4. Run the “ha break” command there, too:

```
Identity Engines>  ha break
```

The HA pair is disconnected, and the VIP definitions are maintained. If you later reconnect either Ignition Server (to its old mate or to another Ignition Server), the HA Configuration Wizard offers you the option of restoring the VIP settings.

Reconnecting a Broken HA Pair

If the link between the nodes in an HA pair fails, the Configuration Hierarchy tree does not display the correct status for the nodes. Follow the instructions below to reconnect the HA link:

1. Select the site in Dashboard’s Configuration hierarchy tree.
2. Complete the breakage of the pair by selecting the **Actions → Break HA Link** command.
3. Recreate HA pair. (See Run the HA Wizard on page 374.)

Reinitializing Nodes in an HA Pair

You can reinitialize a node only if it is a standalone node.

If the node belongs to an HA pair, Ignition Server disables the menu item **Actions → Reinitialize** for the node in the Configuration view of Dashboard.

In order to reinitialize the nodes that are currently linked as an HA pair,

1. Break the link. (See Breaking an HA Pair Using Dashboard on page 387.)

2. Reinitialize the required node(s) and reconfigure the node(s) as necessary. (See “Reinitializing the Ignition Server from Dashboard” on page 55.)
3. Recreate the HA pair. (See Run the HA Wizard on page 374.)

Backing Up Data on an HA Pair

When you back up the data on an Ignition Server, Ignition Dashboard backs up the system configuration and the users, policies, and directory service settings information currently on the Ignition Server. The users, policies, and directory service settings information is identical for the two Ignition Server that you designate as an HA pair. As a result, the backup and restore operations can be performed from either Ignition Server in the HA pair.

To back up Ignition Server data:

1. Use Ignition Dashboard to log in to either Ignition Server in your HA pair.
2. Run the backup as explained in Creating a Backup on page 395.

During the backup operation, the pair continues to provide uninterrupted AAA service.



Avaya strongly recommends that you do not edit data such as users, policies, and directory service settings when you are creating a backup of the Ignition Server data.

Troubleshooting Backups If the backup operation on an Ignition Server which belongs to a linked node fails to complete,

1. Break the HA link.
2. Log in to one of the Ignition Server in the HA pair.
3. Run the backup on this Ignition Server.
4. Re-create the required HA pair. (See Run the HA Wizard on page 374.)

Restoring Data on an HA Pair

Restore operations can be initiated from either Ignition Server in the HA pair. When you restore the data for HA-paired Ignition Server, the **Restore** Window appears.

The data restore operation restores only the identity and policy configuration information on the Ignition Server on which you are restoring data. This is because the system configuration on the two Ignition Server might be different.

To restore Ignition Server data:

1. Use Ignition Dashboard to log in to either Ignition Server in your HA pair.
2. Run the restore as explained in Restoring from a Backup File on page 398.

Since the restore operation affects both nodes, it takes longer to execute on an HA pair than on a stand-alone Ignition Server. During the restore operation, the pair continues to provide uninterrupted AAA service.



Avaya strongly recommends that you do not edit data such as users, policies, and directory service settings during the restore operation. This is because such updates to the data during the restoration is lost when the restoration completes.

Troubleshooting Data-Restore Operations If the restore operation on an HA pair fails to complete, see "Restoring a Non-Responsive Unit in an HA Pair" on page 394.

Updating Firmware on an HA Pair

To update firmware on both nodes, use Ignition Dashboard to log in to either Ignition Server in your HA pair, and run the firmware update as explained in "Firmware Update Procedures" on page 401.



Note that firmware updates affect both nodes and, as a result, take longer to execute on an HA pair than on a stand-alone Ignition Server. Avaya strongly recommends that you do not edit data such as users and policies during the firmware update.

Troubleshooting Firmware Updates If the firmware update fails to complete, follow the instructions in "Restoring a Non-Responsive Unit in an HA Pair" on page 394.

Replacing an Ignition Server in an HA Pair

This procedure, also known as the *box swap procedure* allows you to replace one of the Ignition Server in your HA pair while ensuring that downtime for the RADIUS service and/or SOAP service is minimized.

This procedure requires Ignition Dashboard and the Ignition Server command line interface (CLI).

The Example Configuration: The Ignition Server "*saturn*" is the DB-Primary, and the Ignition Server "*venus*" is the DB-Secondary. VIP is active on the pair, and *saturn* is the VIP-Primary.

The instructions below assume you want to replace the *venus* Ignition Server.

Roles and addresses of the Ignition Server pair

	saturn	venus
DB Role	DB Primary	
VIP Role	VIP Primary	
Admin Port interface address	10.0.3.34/21	10.0.3.33/21
Service Port A interface address	192.168.43.34/24	192.168.43.33/24
VIP address	192.168.43.210	192.168.43.210
HA interface address	192.168.45.34/24	192.168.45.33/24

Procedure:

1. Break the HA relationship between the Ignition Server. To do this, you must issue the `ha break` command on each Ignition Server that is running:

   On *saturn*, use a console terminal to log into the Ignition Server CLI and run the command:

   ```
   Identity Server>  ha break
   ```

   If *venus* is running, run the `ha break` command there as well. (If *venus* is not running, then skip this step.)

   ```
   Identity Server>  ha break
   ```

2. Remove from production the Ignition Server to be replaced. That is, shut down the Ignition Server and disconnect it from the network so that its servicing interface is no longer visible. This forces the VIP to fail over to the other Ignition Server, if it has not already done so. In this example, remove *venus* from the production environment.
3. Turn on the replacement Ignition Server, but do not connect it to the production network. Connect to this Ignition Server using the CLI and set its Admin Port IP addresses to match that of the Ignition Server it replaces. In this example, we refer to this Ignition Server as “the new *venus*”. On the new *venus*, use the CLI as follows to set the address:

   ```
   Identity Server>  interface admin ipaddr 10.0.3.33/21
   ```

4. Connect the Admin Port of your replacement Ignition Server to a non-production environment, and connect a PC with Ignition Dashboard to this network.
5. Run Ignition Dashboard and connect to the replacement Ignition Server. Configure and enable the HA interface. In this example:

- Use Dashboard to connect to the new *venus* at 10.0.3.33.
- Set up the HA interface. (In Dashboard's Configuration hierarchy tree, click the Node, click Ports: HA Port, click Edit, enable the port, and set its IP address.)

6. Configure the replacement Ignition Server VM correctly to your production environment.
7. Use Dashboard's HA Configuration Wizard to re-create the HA pair. You can initiate the Wizard from either Ignition Server. When the Wizard prompts you for HA setup configuration information, make the following settings:
   a. Designate the existing production Ignition Server (*saturn* in this example) as Primary.
   b. When the Wizard offers the option of deleting or restoring the VIP configuration, click the **Keep all existing virtual interfaces** radio button. The Wizard migrates the VIP from the existing/production Ignition Server. (See "Restoring a Saved VIP Configuration" on page 381 for details.) In this example, the Wizard applies *saturn*'s VIP settings to the pair and activates the VIP.

Changing the IP Address of the Admin Port or HA Port in an HA Pair

This procedure lets you change the IP address of the Admin Port or the HA Port on your running HA pair while ensuring that downtime for the RADIUS and/or SOAP service is minimized. This procedure requires Ignition Dashboard and the Ignition Server command line interface (CLI).



Important! If you want to change the IP address of an Ignition Server Service Port, you do not need to follow the steps below. Instead, follow the steps in "Enabling Service Port" on page 64.

The Example Configuration: The Ignition Server "*neptune*" is the DB-Primary, and the Ignition Server "*mercury*" is the DB-Secondary. VIP is active on the pair, and *neptune* is the VIP-Primary. The instructions below assume you want to change the IP address of the Admin Port interface on *mercury* to 10.0.3.99.

Roles and addresses of the Ignition Server pair

	neptune	mercury
DB Role	DB Primary	
VIP Role	VIP Primary	
Admin Port interface address	10.0.3.34/21	10.0.3.33/21
Service Port A interface address	192.168.43.34/24	192.168.43.33/24
VIP address	192.168.43.210	192.168.43.210

Roles and addresses of the Ignition Server pair

	neptune	mercury
HA interface address	192.168.45.34/24	192.168.45.33/24

Procedure:

1. Break the HA relationship between the Ignition Server. To do this, you must issue the `ha break` command on each Ignition Server. For example:

   On *neptune*, use a console terminal to log into the Ignition Server CLI and run the command:

   ```
   Identity Server>  ha break
   ```

   On *mercury*, run the `ha break` command:

   ```
   Identity Server>  ha break
   ```

2. Remove from production the Ignition Server whose IP address is to be changed. That is, connect it to a network partition separate from your production network so that its servicing interface is no longer visible to authenticating clients. This forces the VIP to fail over to the other Ignition Server, if needed.

   In this example, remove *mercury* from the production environment.

3. Using the CLI, log in to the Ignition Server whose IP address you want to change. Set its Admin Port IP addresses to the new IP address.

   In this example, on *mercury*, use the CLI as follows to set the address:

   ```
   Identity Server>  interface admin ipaddr 10.0.3.39/21
   ```



Note. To change the HA port IP address, the command is:

```
Identity Server>  interface ha ipaddr 10.0.4.40/21
```

4. Make the connections to reconnect the newly IP'd Ignition Server to your production environment.

   In this example, reconnect the Admin and HA interfaces of *mercury* to your production network.

5. Use Dashboard's HA Configuration Wizard to re-create the HA pair. You can initiate the Wizard from either Ignition Server. When the Wizard prompts you for HA setup configuration information, make the following settings:

   a. Designate the still-running production Ignition Server (*neptune* in this example) as Primary.

b. Choose to migrate the VIP from the still-running production Ignition Server. In this example, the Wizard applies *neptune*'s VIP settings to the pair and activates the VIP.

c. Use the new IP address(es) when prompted by the Wizard. In this example, use the new Admin IP address of `10.0.3.39` for *mercury*.

Restoring a Non-Responsive Unit in an HA Pair

During a system restore from backup, if a node in an HA pair is rebooted while the other node is being restored, the restore operation might fail to complete and the pair might fail to come back online. If this happens, follow the recovery procedure below to recover the boxes:

1. Break the HA relationship between the Ignition Server. To do this, run the `ha break` command on each Ignition Server:

   On one node (in this example, we refer to this as "Node 1"), use a console terminal to log in to the Ignition Server CLI and run the command:

   ```
   Identity Engines>  ha break
   ```

   On the other node ("Node 2"), connect using the console and run the `ha break` command:

   ```
   Identity Engines>  ha break
   ```

2. Connect to Node 1 using **Ignition Dashboard**. Perform a system restore by clicking the **Site** at the top of the tree and selecting the command, **Actions: Restore Data**.

3. Connect to Node 2 using **Ignition Dashboard**. Perform a system restore in the same manner as explained above.

4. While still connected to Node 2, re-create the HA pair: In Dashboard's hierarchy tree, click the **Site** and select the command, **Actions: Create HA Link**. Use the Wizard to create the pair. When the Wizard offers the option of deleting or restoring the VIP configuration, click the **Keep all existing virtual interfaces** radio button. The Wizard migrates the existing VIP settings to the restored HA Pair. See "Restoring a Saved VIP Configuration" on page 381 for details.

C Backup and Restore Procedures

This section explains how to back up and restore the Avaya Identity Engines Ignition Server data and configuration.

Contents

- "Introduction to Ignition Server Backup and Restore" on page 395
- "Creating a Backup" on page 395
- "Restoring from a Backup File" on page 398

Introduction to Ignition Server Backup and Restore

You can save your Ignition Server's data and configuration to a backup file and later restore it by loading the saved file. Having a backup file ensures you can recover from accidental data loss or administrator error. You can also use backup files to set up a replacement Ignition Server. When you run a backup, the Ignition Server saves the following types of system data:

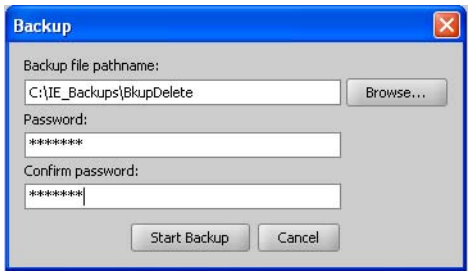
- policies (tunneling, identity routing, authentication, and authorization)
- users and groups in the Ignition Server internal store
- authenticator/NAS configuration
- certificates

Creating a Backup

To create a backup of the data on your Ignition Server:

1. Open the Configuration view of Dashboard.
2. Open the **Backup** Window to specify the file. You can open this window in one of two ways:
   - Right-click on the **Site** icon in the tree. Select **Backup Data**. Or,
   - Click on the site and then choose **Actions → Backup Data.**
3. Dashboard displays the **Backup** window.

Figure 101 Running a Backup



4. Click the **Browse** button to specify the path and filename for the backup file.

Note: Because you can only restore a backup file on an Ignition Server of the same software version or a previous major software version (for example, you can perform a restore on an 8.0 system with a backup file taken from a 7.0.x system or 8.0 system) it is safe practice to note the Ignition Server firmware version in the file name of your backup.

5. In the **Password** and **Confirm Password** fields, enter a passphrase for encrypting the file. Leave these fields blank if you do not want to encrypt the file. When you restore from the backup file, you are prompted to enter this passphrase to decrypt the file.
6. The **Start Backup** button is enabled. Click **Start Backup**.

While the backup is in progress, a status bar is displayed. When it completes, a completion message appears.

Ignition Server logs each backup and restore operation in its Administrative Activity Log.

Troubleshooting

Under certain circumstances the backup procedure fails with the warning message: “Backup Failed.” If this happens, reopen the backup window and try again.

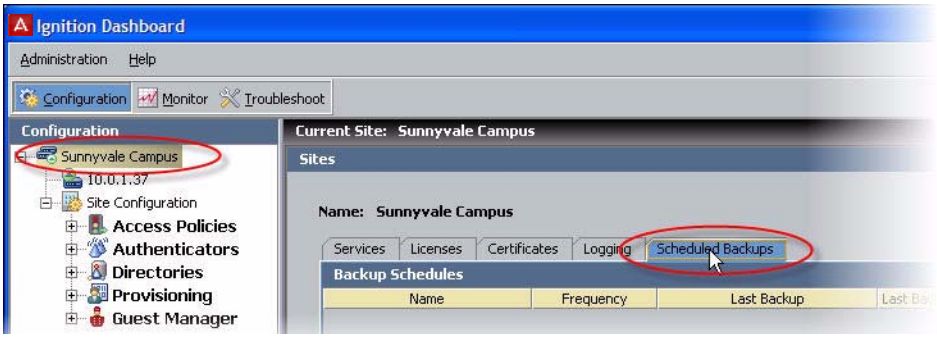
Scheduled Backup

Your Ignition Server or HA pair of Ignition Server s can be scheduled to back up its data at a specific time or at a regular interval. To do this, you must have an SFTP server running on the target computer that will store your backup files.

Follow these steps to set up a backup schedule:

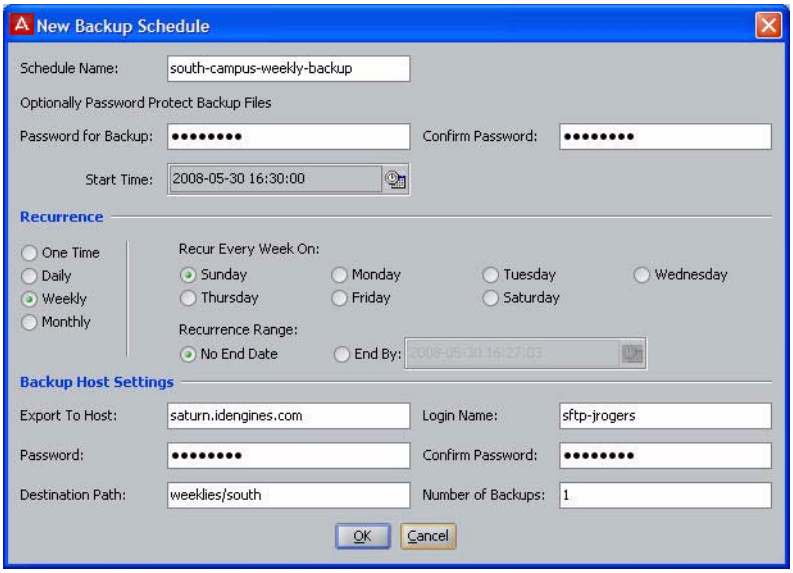
1. In Dashboard’s Configuration tree, click the name of your Ignition Server site. On the right side of the window, click the **Scheduled Backups** tab.

Figure 102 The Scheduled Backups tab



2. Click the **New Backup** button.

Figure 103 Setting a backup schedule



3. In the New Backup Schedule window, type a **Schedule Name** to identify the schedule.
4. In the **Password** and **Confirm Password** fields, enter a passphrase for encrypting the file. Leave these fields blank if you do not want to encrypt the file. When you restore from the backup file, you are prompted to enter this passphrase to decrypt the file.
5. Set the start time and schedule:
   - Click the clock and calendar icon of the **Start Time** field to set the time when the first back up will begin.
   - In the **Recurrence** section, specify the frequency (one time, daily, weekly, or monthly)
   - To the right of the frequency radio button, specify the detailed frequency parameters. For monthlies, choose a numbered day of the

month; for weeklies, choose a day of the week; for dailies, specify either a frequency ("**Every** *n* **days**") or choose **Every Weekday**. To save a backup every day, specify a frequency of "Every *1* days".

- In the **Recurrence Range** field, specify the end date, if any, for this schedule.

6. In the **Backup Host Settings** fields, specify the SFTP server that is to receive backup files:
   - In the **Export to Host** field, specify the machine name or IP address of the destination SFTP server.
   - Set **Login Name** and **Password** to the user name and password of the SFTP server account that will own the backup files.
   - Type the password again in the **Confirm Password** field to confirm it.
7. Specify where and how many backup files should be kept:
   - In the **Destination Path** field, specify the path where the backups will be saved on the SFTP server.
   - In the **Number of Backups** field, specify the maximum number of backup files to be kept. For example, if you run weekly backups and set this to "3", then Ignition Server will save three backup files over the first three weeks, and in each subsequent week it will overwrite the oldest backup file.
8. Click **OK**. Your schedule has been saved.
9. If you wish to add more backup schedules, repeat the above steps.

Restoring from a Backup File

When you restore an Ignition Server, the restoration process overwrites the data on the Ignition Server. Ignition Server disconnects the Ignition Server until the restore operation completes successfully. It then automatically reboots the Ignition Server to ensure that the restored system data takes effect.



Warning: The Ignition Server version of the backup file must match the version or the last previous major version of the Ignition Server firmware on which you are restoring it. For example, you can perform restore on an 8.0 system with a backup file taken from 7.0.x system or 8.0 system. Use Dashboard's Firmware Manager window to check the firmware version running on your Ignition Server. If you need to upgrade or downgrade your firmware, consult the *Ignition Server Release Notes*.

Admin Port IP Address

When you restore from a backup file, if you choose to restore the System Configuration, your Admin port IP address and network settings are set to the settings from the backup file. Make sure you know which IP address the

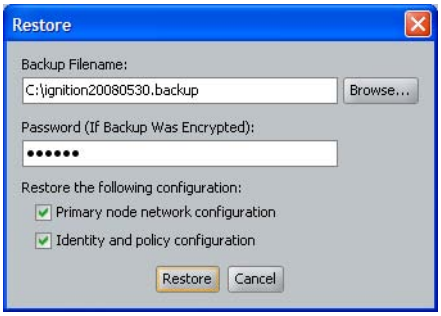
backup applies so that you can reconnect Dashboard to the Admin port after the data restore is complete, or be prepared to use the Ignition Server front panel to reset the Admin port IP address after the restore is complete.

Steps to Restore Data

To restore your Ignition Server from a backup file:

1. Open the Configuration view of Dashboard.
2. In the **Configuration Hierarchy** panel, right-click on the **Site**. Select **Restore Data**. Dashboard displays the **Restore** window.

Figure 104 Specifying the Location of the Backup File



3. Click the **Browse** button to specify the path and filename for the backup file from which you are restoring the data.
4. In the **Password** field, enter the passphrase that you entered when encrypting the backup file. Leave the field blank if the file is not encrypted.
5. Select the types of data you want to restore. Check one or both of:
   - **Primary node network configuration:** The network settings specific to the primary node ("Managing a Node" on page 54). **Warning!** When you restore the **Primary node network configuration** data, Ignition Server overwrites the Admin Port IP address of your Ignition Server. Be prepared to reconnect Dashboard to the new Admin port IP address.
   - **Identity and policy configuration:** Selecting this checkbox restores the records of Ignition's internal data store, including policies and directory service settings.
6. The **Restore** button is enabled. Click **Restore**.

Ignition Server requests a confirmation as it overwrites the existing data on the Ignition Server and reboot the Ignition Server.

7. Click **Yes** in the Confirm Restore Window. You may click the **Logout** button to disconnect Dashboard from the Ignition Server until the restoration completes. Wait a few minutes and then log in to the Ignition Server. If you do not disconnect, then Dashboard reconnects automatically after the restoration is done.
8. **Admin port IP address:** If you chose to restore the System Configuration in Step 5, then the backup file contains an *IP address assignment* for the Admin Port, and the restoration routine applies this IP address to the Admin Port of the virtual appliance where you're restoring. In the case of change of Admin IP due to restore (backed up config) Dashboard will not reconnect in this case reconnect by entering new Admin IP in the dashboard login window". This info is displayed in "restore window
9. **Ignition Server licenses:** If you chose to restore the System Configuration in Step 5, then the backup file contains *all the Ignition Server licenses* that were on the backed up machine, and the restoration routine installs these licenses on the virtual appliance where you're restoring. Licenses are keyed to the serial number of the Ignition Server, so if you have restored them on a different Ignition Server from the one where you backed them up, you must replace them with new licenses keyed to the new virtual appliance. See "Installing an Ignition Server License" on page 69.

Troubleshooting

If, for any reason, the restore operation fails, Ignition Server issues one of the following error messages:

Restoration Error Messages

Error Message	Steps to Correct
`This is not a valid backup file.`	You have assigned the wrong backup file. Use the browser on your personal computer or workstation and locate the correct backup file. The backup file might be corrupted. The Ignition Server verifies the digital signature of the backup file before carrying out the restore operation. If the signature does not match the key pair generated when the system was installed, the restore operation does not finish. Redo the Backup operation.
`Permission denied. You entered an incorrect password.`	Enter the correct passphrase you provided for encryption when you created the backup file.

D Firmware Update Procedures

Periodically, firmware updates and patches are made available on the Avaya support web site, and bulletins are sent to administrators of Avaya Identity Engines Ignition Server systems.

Release 7.0 introduced an OS package upgrade feature that allows you to upgrade individual components without having to reload/re-image the entire appliance. Packages are files containing RPM's (original Avaya distributions or third party) along with an Ignition Server image file.

Contents

- "Checking the Firmware Version" on page 401
- "Loading a Firmware Image or Package" on page 403
- "Activating a Firmware Image or Package" on page 404
- "Activating a Firmware Image on an HA Pair" on page 406
- "Deleting a Firmware file" on page 407
- "Viewing image and package information" on page 407

Checking the Firmware Version

There are two ways to display the version number of the current firmware on the Ignition Server:

- "View the Current Firmware Version" on page 401
- "View the Current Firmware Version and All Installed Images and Packages" on page 402

Note: To check the version of Ignition Dashboard you are running, select **Help: About** from the main window. To check the documentation version, see "Version Number of This Document" on page 24.

View the Current Firmware Version

To view the current firmware version on the Ignition Server:

1. In the navigation panel on the left side of the Dashboard window, click the name or IP address of your Ignition Server.
2. Click the **Status** tab.

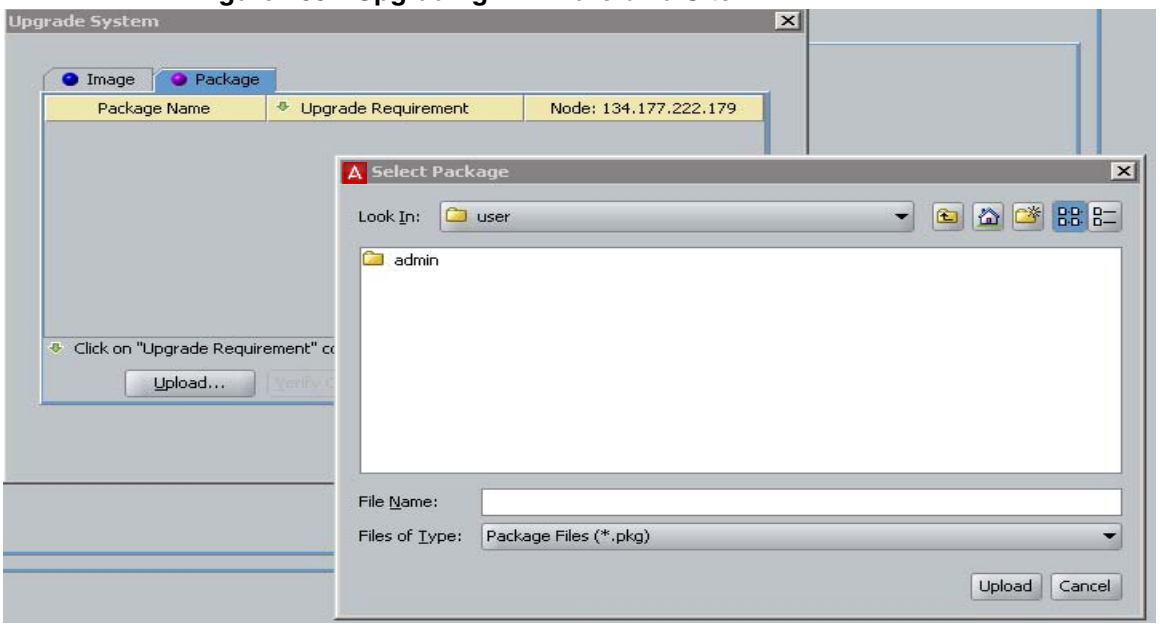
Dashboard shows the version in the **Current Configuration: System Version** display field.

View the Current Firmware Version and All Installed Images and Packages

To view the current firmware version and all installed firmware images on the Ignition Server:

1. Open the Configuration view of Dashboard.
2. Select your **Site** (by default, the name is *Site 0*) in the **Configuration Hierarchy** section of the Configuration view of Dashboard.
3. Launch the **Firmware Manager** window by selecting the command, **Actions → Upgrade System**). Dashboard displays the **Firmware Manager**:

Figure 105 Upgrading Firmware on a Site



The Firmware Manager displays a tab for firmware Images or for Packages.

The firmware section displays the firmware image or package currently loaded on the Ignition Server. A green dot in the right column indicates the image or package *that is the currently running firmware image*, and an asterisk (*) in the right column indicates the image *that has been selected to be the currently running image*. Normally, these are the same image or package, but if the administrator has just activated it and the Ignition Server has not yet restarted itself, then they might not be. In all cases, the asterisk-marked image and package becomes the running file upon the next reboot.

Rows displaying “installed” in the right column represent images that have been loaded but are not the currently running image. To migrate your Ignition Server to one of these files, see “Activating a Firmware Image or Package” on page 404.

The version number takes the form, “LINUX-VM_08_00_00_019152” where “LINUX” indicates this is the version of the Ignition Server firmware appropriate for the Virtual Appliance that runs in VMWare env, “08_00_00” indicates this is firmware version 8.0.0, and “019152” indicates this is firmware build number 19152.

Loading a Firmware Image or Package

Use the following procedure to update the firmware on your Ignition Server:

1. Select a firmware update from the Avaya web site (see “Customer service” on page 21 for the address) and use a web browser to download it to your administrative machine.
2. Open the Configuration view of Dashboard.
3. Select your **Site** (by default, the name is *Site 0*) in the **Configuration Hierarchy** section of the Configuration view of Dashboard.
4. Launch the **Firmware Manager** window by selecting the command, **Actions → Upgrade System**).
5. Select the firmware **Image** or **Package** tab.
6. In the **Upgrade System** window, navigate to find the firmware image or package you downloaded earlier. Click on the file name and click **Upload**.
7. The firmware is uploaded to the Ignition Server appliance. This might take a few minutes. Upon completion, a success message is displayed. The loaded image appears in the **Firmware Images** or **Package Name** list in the **Firmware Manager**.
8. **Note!** If there are too many firmware images or packages on your Ignition Server, the upload attempt fails. Ignition Server is partitioned to store a maximum of two versions in the boot partition. If your Ignition Server has been upgraded multiple times, it is mandatory that you delete the oldest software versions prior to upgrading to 8.0 so that no more than two images are displayed under the Images tab before package activation is initiated. For information on deleting a firmware file, see “Deleting a Firmware file” on page 407.
9. Set the Ignition Server to use the new firmware. See “Activating a Firmware Image or Package” on page 404.

Note: It is possible the firmware upload is interrupted (for example, if the network connection between the GUI and Ignition Server is lost during the file transfer). In such circumstances, Ignition's attempt to validate the uploaded file fails and/or, after a fixed time interval, Ignition Server times out and stops the upload attempt. When an upload fails, Ignition Server removes the partial file and you must re-attempt the upload.

You can choose to terminate the firmware update process, in which case Ignition Server returns to its prior state.

Activating a Firmware Image or Package

Use the following steps to migrate the Ignition Server to a different firmware image.

Before you begin, check Dashboard compatibility:

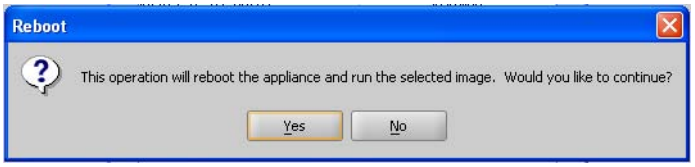
Read the *Ignition Server Release Notes* that correspond to the firmware image or package you want to activate. The firmware version you activate may require that you run a different version of Ignition Dashboard. Do one of the following:

- If a Dashboard upgrade or downgrade is required, follow the instructions provided in the *Ignition Server Release Notes*; or
- If *no* Dashboard upgrade or downgrade is required, follow the instructions below.

Activate the firmware image:

1. Select your **Site** (by default, the name is *Site 0*) in the **Configuration Hierarchy** section of the Configuration view of Dashboard.
2. Launch the **Firmware Manager** window by selecting the command, **Actions → Upgrade System**).
3. Select **Images** tab.
4. In the **Firmware Manager**, find the image you want to activate and click it to select it.
5. Click **Activate**. You are asked whether you want to reboot Ignition.

Figure 106 Reboot Dialog



6. You must reboot for the **Activate** operation to take effect. Click **Yes**.

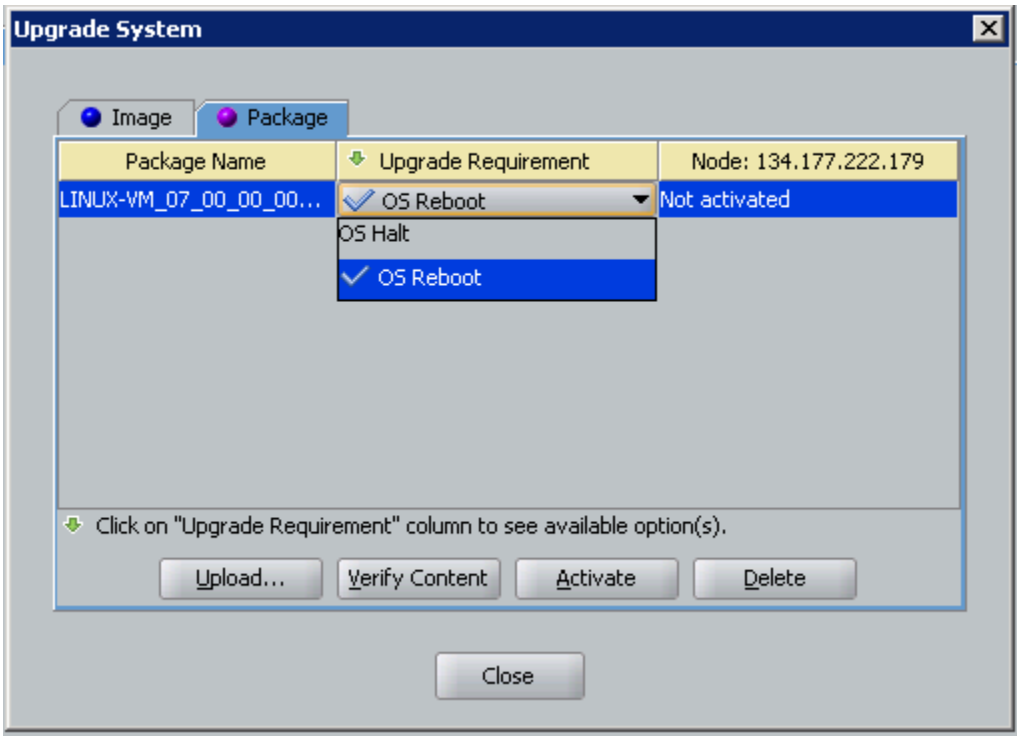
The **Rebooting** dialog appears displaying the status of the reboot operation. When the reboot operation completes, the selected image is active.

Activate the firmware package:

Only *not-installed* packages can be deleted or verified or activated. Each Package contains a flag that indicates that when it is installed if a reboot, For not-installed packages, the Node column displays **Not activated**. Once the package is activated, the display changes to “Activated on MM/DD/YYYY”.

1. Select your **Site** (by default, the name is *Site 0*) in the **Configuration Hierarchy** section of the Configuration view of Dashboard.
2. Launch the **Firmware Manager** window by selecting the command, **Actions → Upgrade System**).
3. Select the **Packages** tab.
4. In the **Firmware Manager**, select the package you want to activate.
5. Select from the **Upgrade Requirement** column one of the available upgrade requirement/recommendations for the system/server restart. The recommended option is tagged with the check mark. You may select the other options but Avaya recommends that you should always select the recommended option.

Figure 107 Package window



The Upgrade Requirement column contains the following values.

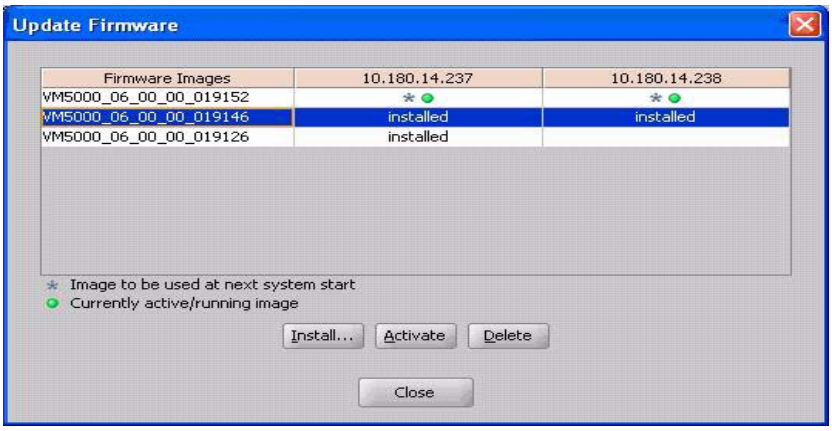
- OSreboot - the OS reboots after the upgrade is completed.
- OS Halt - the OS halts after the upgrade is completed.
- Server Restart - the Ignition Server restarts after the upgrade is completed
- Hitless - no further action is needed after the upgrade is completed.

6. Select **Verify Content** to check the package integrity and inform the result.
7. Click **Activate**. You are asked whether you want to reboot Ignition.

Activating a Firmware Image on an HA Pair

In order for an image to be activated on an HA-paired Ignition Server set, the image must be present in both Ignition Server s. The Activate button is enabled only when you select an image that is installed on both Ignition Server s as shown below.

Figure 108 The Firmware Manager Window



When you click on the **Activate** button, the **Rebooting** Window that appears to indicate the progress of the operation is on the display window only. The **Cancel** button is disabled.

Deleting a Firmware file

You can delete old firmware files from your Ignition Server. You cannot delete the currently running firmware image or package. To upgrade, first install and activate the new file, and then delete the old one. Only *not-installed* packages can be deleted or verified or activated.

To delete a firmware file:

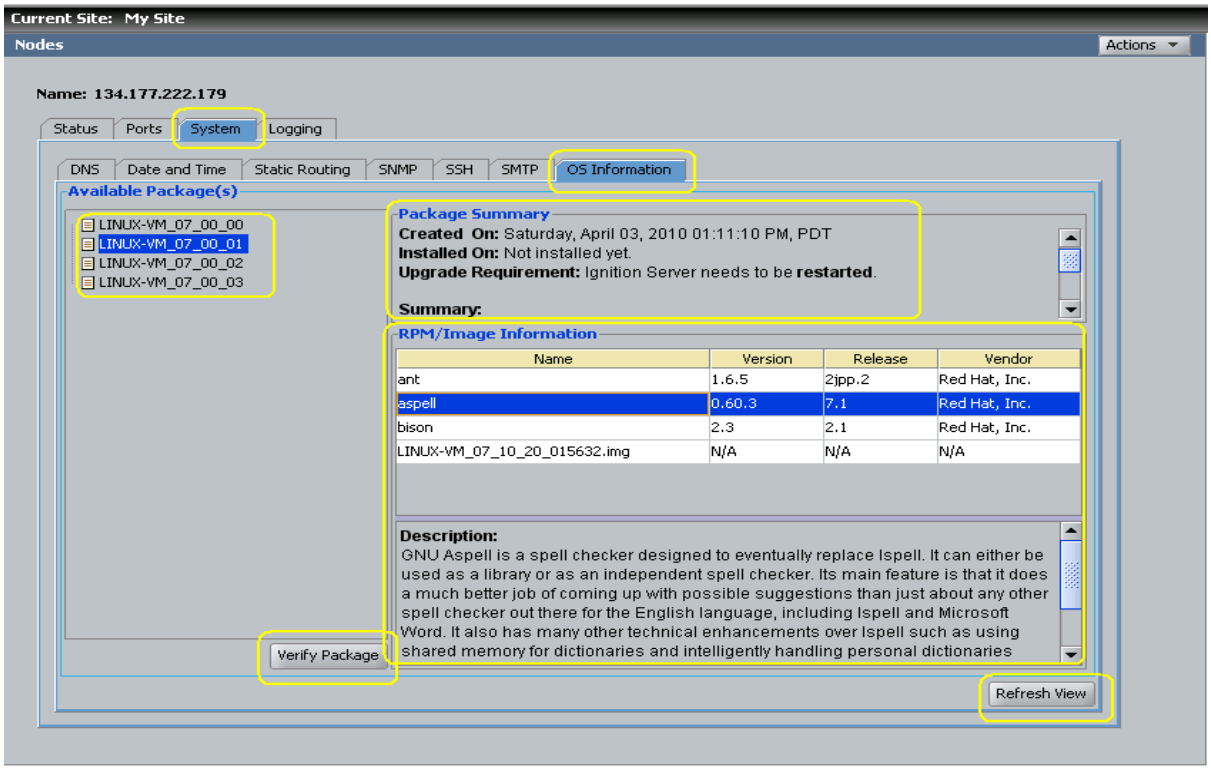
1. In Dashboard's Configuration hierarchy, select your Site.
2. Click **Actions**. (Alternatively, right-click the **Site** in the **Configuration Hierarchy** tree.) Select **Upgrade System...**
3. Select the **Images** or **Packages** tab.
4. The **Firmware Manager** Window appears displaying the existing firmware images or packages on your Ignition Server. *The running version is marked with an asterisk.* ***IS THIS STILL CURRENT?***
5. Scan the list for the firmware file you no longer want to keep, and click on the row to select it.
6. Click the **Delete** button.
7. A confirmation window appears. Click **OK** to confirm. Ignition Server deletes the firmware image.

Viewing image and package information

To view OS information currently installed or is available for upgrade

1. Launch the **OS information** window by selecting the command, **Nodes → System**.
2. Select the **OS information** tab.

Figure 109 OS information



The OS Information window contains the following in Read-Only.

- The Package Summary displaying installation information.
- The RPM/ Image information displays available RPMs/Images for the selected package (in LHS). The RPMs are sorted in alphabetical order followed by the Image entries sorted in alphabetical order.
- The image or file details as a Description.

Use the **Verify Package** to verify/validate the package content for the package(s) that have not been installed. This is disabled when the package which is already installed.

You can refresh the data by clicking **Refresh View**.

Upgrading from Ignition Server 7.0 to 8.0

Ignition Server 8.0 does not support package Rollback.

Note: Avaya recommends that you use Dashboard to perform upgrade and restore tasks. If you want to use CLI (for a single-node system), Avaya recommends that you use an FTP server.

Upgrading to Ignition Server 8.0

For an Ignition Server 7.0 to an Ignition Server 8.0 upgrade there is no roll back option from 8.0 to 7.0. Avaya recommends the following procedure to create a backup in the event you need to revert from Ignition Server 8.0 back to Ignition Server 7.0:

1. Export and save your existing Ignition Server 7.0 configuration.
2. Halt the Ignition Server Nodes you plan on upgrading.
3. Make sure the nodes are in powered down state.
4. Using the Storage management section of your ESXi Server create a separate folder to hold the disk drive of each of the nodes in your configuration.
5. Using the standard ESXi Server storage manager functions copy each disk to one of the folders created in the previous step. Now you have a complete backup on the disk before you start the 7.0 to 8.0 upgrade.
6. Upload the 8.0 image and activate the image.
7. The system reboots two times in order to perform the upgrade. At the end of the second boot Ignition Server 8.0 is running.
8. If the upgrade fails or you decide to revert back to Release 7.0, perform the following steps.

If the upgrade fails complete the following

1. Power off the VM as the disc the disk is not in a usable state.
2. Go to the VM Setup menu and delete the hard disk from the VM. (do this for each node in your system)
3. Copy the saved VM back to its original location.
4. Attach the disk to VM using the attach existing disk option. (do this for each node in your system).
5. Power on the VM(s).
6. Your previous Ignition Server 7.0 environment is back online.

E Setting Up Logging

This Appendix explains how to set up Avaya Identity Engines Ignition Server logging and how to use Ignition Server as a RADIUS accounting server.

Contents

- "Setting User Preferences" on page 411
- "Setting Up Ignition Server Logging" on page 411
- "Setting Up FTP Log Export" on page 413
- "Exporting logs" on page 414
- "Directing Log Messages to a Syslog Server" on page 416
- "Sending Log Messages Via E-Mail" on page 417
- "Monitoring Ignition Server via SNMP" on page 418

Setting User Preferences

You can set your log viewing preferences in the **Preferences** Window. See "Setting Administrator Preferences" on page 42.

Setting Up Ignition Server Logging

The **Configure Log Types** Window lets you specify what logging information the Ignition Server records, and the **Configure Automated Log Export** window lets set up periodic log exports.

Note! If you're running an HA pair, please be aware that the log settings you make here apply to this node only. Each node maintains its own logs, and logs are not synced between nodes in the pair.

For instructions on setting up logging, consult the appropriate section below:

- "Turning on Logging" on page 411
- "Setting the Level of Logging to Be Recorded" on page 413
- "Setting Up FTP Log Export" on page 413

Turning on Logging

1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.

3. Click the **Log Viewer** tab.
4. Click the **Configure** button in the upper right corner of your window.

Figure 110 The Configure Log Types window



5. In the Configure Log Types window, activate logging for the log types you want to log:
   - To turn on RADIUS and TACACS+ logging, tick the **Enabled** checkbox in the **Access** row.
   - *Audit* logging is always enabled; you cannot turn it off.
   - To turn on *Security* logging, tick the **Enabled** checkbox in the **Security** row.
   - To turn on system *Environment* logging, tick the **Enabled** checkbox in the **Environment** row.
   - To turn on *Debug* logging, go to the **Logging & Severity** column of the **Debug** row and click the drop-down list. Select the desired degree of logging. To minimize the number of logging events recorded, select the level, *Fatal*.
   - To turn on *System* logging, go to the **Logging & Severity** column of the **System** row and click the drop-down list. Select the desired degree of logging. To minimize the number of logging events recorded, select the level, *Fatal*.
   - To turn on detailed logging of your RADIUS and TACACS+ authentications and authorizations, tick the **Enabled** checkbox in the **Access Details** row
6. Click **Save**.

Next step: Set up the rules for periodic export of your logs, see “Setting Up FTP Log Export” on page 413.

Setting the Level of Logging to Be Recorded

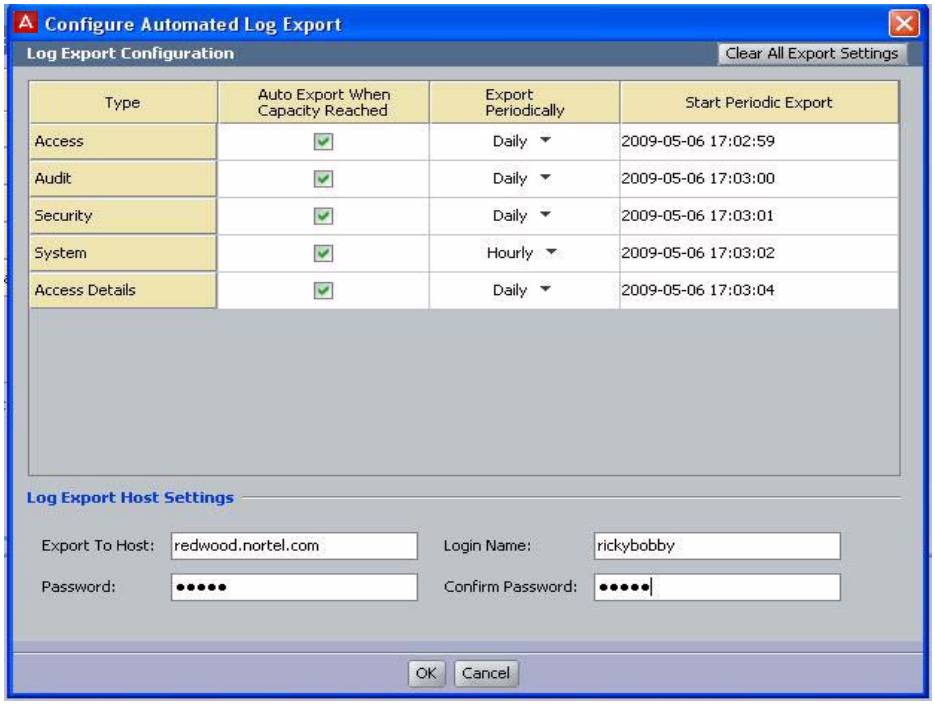
See the section, “Turning on Logging” on page 411.

Setting Up FTP Log Export

Follow the steps below to set up FTP exporting of your log files:

1. In Ignition Dashboard, click Configure to show the configuration view.
2. Click the name of your site in the tree.
3. Click the Logging tab, and click the Export Logs tab. Click Edit.

Figure 111 Configuring Automated Log Export



4. If you want to export the log contents automatically when a log channel reaches its maximum size, go to the row of your log channel and tick the checkbox in the **Auto Export When Capacity Reached** column. If you do not tick this checkbox, then Ignition Server begins to overwrite the oldest log records when the channel reaches capacity.
5. To export a log channel’s contents at a regular time interval, do this:
   - Go to the row of your log channel and, in the **Export Periodically** column, use the drop-down list to select the export interval of *Hourly*, *Daily*, or *Weekly*.

- The **Start Periodic Export** column displays the time when the first export is to occur. If you want to export at a particular time of the hour, day, or week, set an appropriate starting time here. To do this, click the cell in the **Start Periodic Export** column, and click the clock and calendar icon. Click the up and down arrows to set a date and time. To complete your entry, click outside the date and time dialog box and press <Enter>.

6. In the **Log Export Host Settings** fields, specify the SFTP server that is to receive log exports:
   - In the **Export to Host** field, specify the machine name or IP address of the destination SFTP server.
   - Set **Login Name** and **Password** to the user name and password of the SFTP user.
   - Type the password again in the **Confirm Password** field to confirm it.
7. Click **Save**.

Exporting logs

By default, the Ignition server exports the last 5000 (or less) logs for the selected log type. But whenever there are more than 5000 logs that you want exported then define the following environment variable in the local OS that your Dashboard is launched: EXPORT_ALL_IDE_LOG=true.

You need to close the current IDE session, launch Dashboard and follow the same steps to export logs as specified in the following procedure.

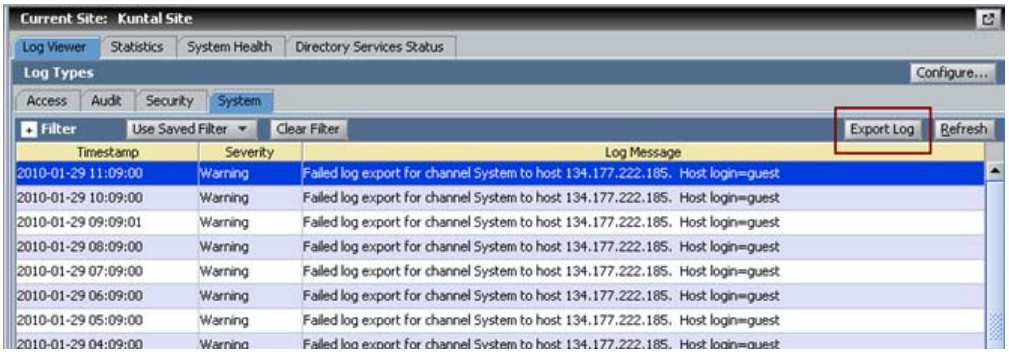


Exporting logs using the Export All variable may take longer depending on the number of logs that are in the database.

Follow the steps below to export your log files immediately:

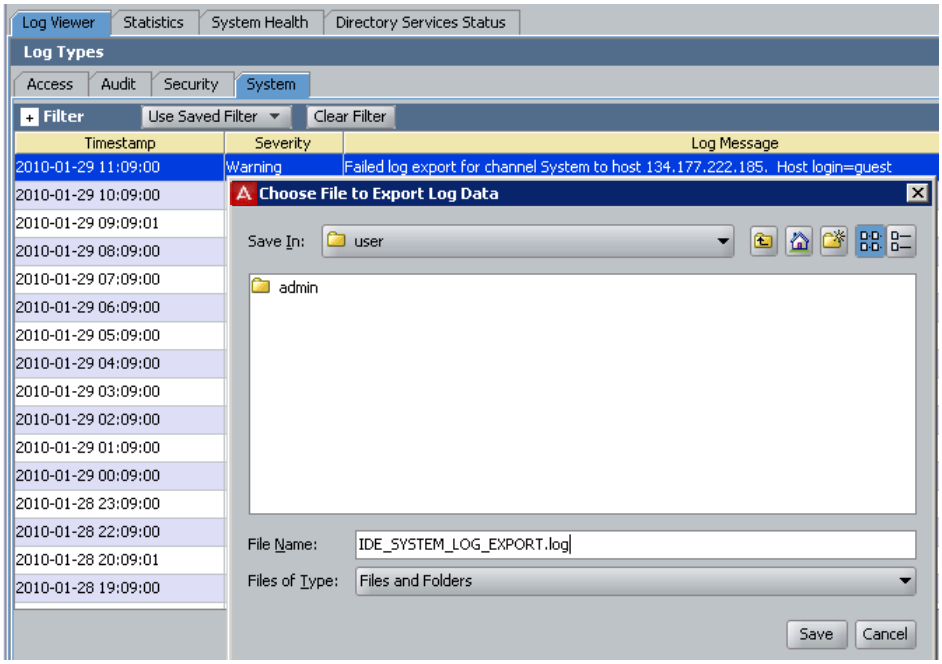
1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Log Viewer** tab.
4. Select the Log type that you want to export from by clicking on the **Access**, **Audit**, **Security**, or **System** tab.
   - To enable the Debug log, in Dashboard click on the **Troubleshoot** tab, select the site and the click on the **Actions** drop down box to enable the **Debug Logs** and **Advanced Log Levels** options.
5. Highlight the logs you want exported.

Figure 112 Log Viewer



6. Click on **Export Log** on the right. A progress window opens showing you the retrieving of your selected records.
7. Once the files are retrieved, the **Choose File to Export Log Data** window opens. Type a file name or select an existing file name to save the logs.

Figure 113 Choose File to Export



8. Click **Save**.
9. You can use any text editor such as Textpad, Notepad, or vi, to open the exported file.

Figure 114 Log Data



```
Log data was exported by user: "admin" for log type: "System" at time: Jan 29, 201
            No. of lines exported: 600
        *** Please do NOT edit the content ***
-------------------------------------------------------------------------------

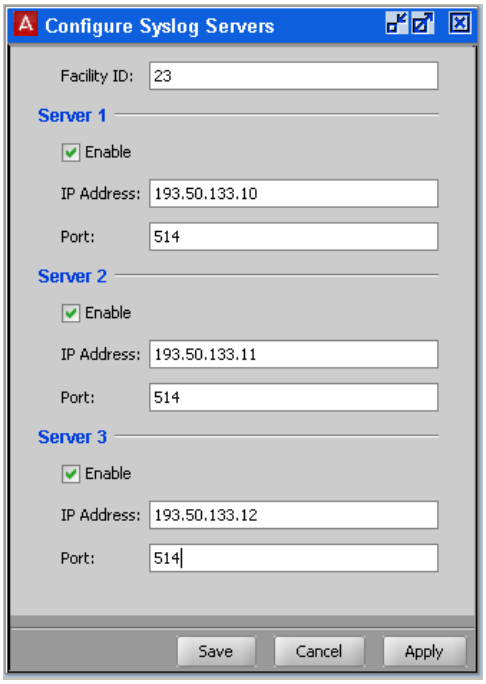
Timestamp            Severity          Log Message
-------------------------------------------------------------------------------
2010-01-29 11:09:00  Warning           Failed log export for channel System to h
2010-01-29 10:09:00  Warning           Failed log export for channel System to h
2010-01-29 09:09:01  Warning           Failed log export for channel System to h
2010-01-29 08:09:00  Warning           Failed log export for channel System to h
2010-01-29 07:09:00  Warning           Failed log export for channel System to h
2010-01-29 06:09:00  Warning           Failed log export for channel System to h
2010-01-29 05:09:00  Warning           Failed log export for channel System to h
```

Directing Log Messages to a Syslog Server

To direct log messages to one or more Syslog servers:

1. In Ignition Dashboard, click Configure to show the configuration view.
2. Click the name of your site in the tree.
3. Click the Logging tab, and click the Syslog tab. Click Edit.

Figure 115 Configure Log Servers Window



4. Enter the **Facility Id** that you want to identify your Ignition Server as the origin of the log message. Consult your Syslog man pages or documentation for details.
5. Click to **Enable** to enable each server you require. For each server, specify:
   - **IP Address**: Enter the Syslog server's IP address.
   - **Port**: Enter the Syslog server's listener port.
6. Click **Save**.

Sending Log Messages Via E-Mail

You can set up Ignition Server to send log alerts via e-mail using an SMTP server on your network. To set this up:

1. In Dashboard's Configuration hierarchy tree, click the name or IP address of your node.
2. Click the **System** tab and click the **SMTP** tab.
3. Click **Edit**.
4. In the Edit SMTP Configuration window, make the SMTP server settings:
   - Check the SMTP Server Enabled checkbox.
   - In SMTP Server Address, specify the address of an SMTP server that is running on your network.
   - Specify the SMTP server Port number.
   - In Sender Address, specify the e-mail address that you want to serve as the From and ReplyTo address on e-mails that Ignition Server sends.
   - If your SMTP server requires a login credentials, click the Authentication checkbox and specify the SMTP server credentials in the User Name and Password fields.
5. Add the list of addresses to which you want Ignition Server log e-mails to be sent.
   - Click **Add**.
   - An empty row appears in the table. Click on the row and type the e-mail address of the recipient.
   - Repeat the preceding two steps to add more recipients.
6. Specify which Ignition Server logs you want to be e-mailed, and how often. For each type of log to be sent, do the following:
   - Click the **Enabled** checkbox in the row for the log type to be sent.

- In the **Interval** column, specify the frequency at which you want the e-mails to be sent.
- In the **# Msg Per Email** column, specify the maximum number of messages that can be sent in a single e-mail. An e-mail is sent when the **# Msg** threshold is reached or when the **Interval** expires, whichever comes first.

7. Click **OK**.

Monitoring Ignition Server via SNMP

Ignition Server provides an SNMP service that allows you to retrieve Ignition Server system statistics and settings using an SNMP client. You cannot write settings to Ignition Server via SNMP.

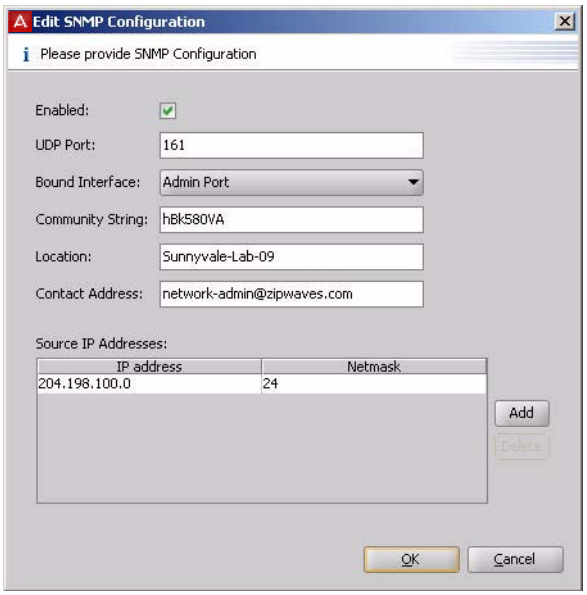
The Ignition Server SNMP service supports SNMPv2C. The objects supported by the service come from the following SNMP MIBs: the *MIB-2* MIB, the *ucdavis* MIB, and the *HOST-RESOURCES-MIB*.

Configuring Ignition's SNMP Service

Use the Edit SNMP Configuration window to enable and set up Ignition's SNMP service. Follow these steps to set up the SNMP service:

1. In Dashboard's Configuration hierarchy tree, click on the IP address or name of your Ignition Server.
2. Click the **System** tab and click the **SNMP** tab. This tab displays the current SNMP settings.
3. Click the **Edit** button.

Figure 116 Edit SNMP Configuration window



4. In the Edit SNMP Configuration window, make these settings:
   - Tick the **Enabled** checkbox to turn on the SNMP service.
   - In the **UDP Port** field, type the port to which you plan to connect your SNMP client. The default is 161.
   - In the **Bound Interface** drop-down box, select the name of the Ignition Server network port where SNMP is available.
   - In the **Community String** field, type the community string that connecting SNMP clients must submit in order to connect. While this string acts as a form of password for connecting clients, please note that SNMP communications are not secure. Both the community string and SNMP traffic are transmitted in the clear.
   - In the **Location** field, enter a location string that indicates where this Ignition Server is located.
   - In the **Contact Address** field, enter a string that indicates where the administrator responsible for this system can be contacted. For example, you might type “network-admin@zipwaves.com” or you might type “Message network support at 408-555-3457.”
5. In the **Source IP Addresses** table, enter one or more IP addresses that can act as filters that limit who can connect to the SNMP service. To connect, a client must have an IP address that exactly matches a row in the table or, for a filter row whose least significant octets are zeros, that falls in the subnet described by the filter. For example, a filter of 204.198.100.0/24 means that machines with IP addresses from 204.198.100.1 through 204.198.100.254 are allowed to connect.

6. Click **OK**.

After you complete the steps above, the SNMP service is running on your Ignition Server. You can connect to it as shown in the next section.

Connecting to Ignition's SNMP Service

Ignition's SNMP service allows SNMP management stations that support SNMPv2 to perform *get* and *walk* actions on the Ignition Server MIB. For HA pairs, you must connect to each node (each Ignition Server) individually. System information and statistics are stored independently for each Ignition Server node.

Use the steps below to query the Ignition Server SNMP service:

1. In Dashboard's Configuration hierarchy tree, click on the name of your Ignition Server node. Click the **System** tab and click the **SNMP** tab. Note the SNMP settings.
2. Use your SNMP management station to query the Ignition Server SNMP service. Make sure the IP address of the machine where your SNMP management station runs is one that passes through the filter defined in the **Source IP Addresses** field in Dashboard.

Example SNMP Queries

The following examples demonstrate how to retrieve information using an SNMP management station. These examples were done using the Net-SNMP tool, but other tools work in a similar fashion.

The first example uses the *snmpwalk* command to retrieve the entire *mib-2* subtree. Using your SNMP management station application, type the following:

```
snmpwalk -c hBk580VA -v 2c 204.158.10.37 mib-2
```

where:

- *-c* is the community string ("hBk580VA" in this example), which serves as a password
- *-v* is the snmp version, which is "2c"
- the next argument after the "-v 2c" argument is the IP address (204.158.10.37 in this example) of the Ignition Server port to which you have bound the SNMP service.

The second example command uses *snmpwalk* to retrieve the same information using OID notation (".1.3.6.1.2.1") for the MIB-2 subtree. Type the following:

```
snmpwalk -c hBk580VA -v 2c 204.158.10.37 .1.3.6.1.2.1
```

The third example uses *snmpget* to get a single SNMP object ("sysDescr"). The trailing ".0" that you see here is the instance identifier required by most SNMP tools when retrieving the value of a scalar object.

```
snmpget -c hBk580VA -v 2c 204.158.10.37 sysDescr.0
```

This example returns a string indicating the firmware version now running on the Ignition Server:

```
SNMPv2-MIB::sysDescr.0 = STRING: 3000E_03_03_00_009734
```

Data Objects Exposed by the Ignition Server SNMP Service

The main SNMP objects available in Ignition Server are:

- **Date, time, and uptime information** is shown in the objects of the HOST-RESOURCES-MIB:
  - **hrSystemUptime**: Uptime of the Ignition Server.
  - **hrSystemDate**: Current system date/time of the Ignition Server.
- **Networking statistics** are published in the objects of the IF-MIB, such as ifPhysAddress and ifAdminStatus. Data is recorded per Ethernet port. For each statistic, the port number is indicated in the SNMP object name as shown in the following table.
- **The routing table** of the Ignition Server is published in the RFC1213-MIB objects such as ipRouteDest, ipRouteIfIndex, and so on.
- **System load information** is shown in the laLoad objects. These objects indicate the load on the Ignition Server, expressed using the *load average* convention.
- **General system information** is recorded in the sys objects of the SNMPv2-MIB, including:
  - **sysDescr**: Ignition Server firmware version.
  - **sysUptime**: Uptime of the Ignition Server SNMP service (snmpd process). This is not the Ignition Server system uptime. For that, see mib-2.host.hrSystem.hrSystemUptime, above.
  - **sysLocation**: The physical location of this Ignition Server, as set in Ignition Dashboard.
  - **sysName**: The MAC address of the Ignition Server Admin port.
  - **sysContact**: Contact details that indicate where you can reach the administrator responsible for the Ignition Server system. This is set in Ignition Dashboard.

Port Names Used in SNMP Output

When reading the SNMP data, note the following abbreviations that identify the Ignition Server Ethernet ports.

Name Abbreviations for the Ignition Server Ethernet Ports

Interface Name	Interface Name in Ignition Server firmware/CLI	Index number in SNMP records	SNMP Example
Service Port A	wm0	1	ifDescr.1
Service Port B	wm1	2	ifDescr.2
Admin port	bge0	3	ifDescr.3
HA port	bge1	4	ifDescr.4
loopback address	lo0	5	ifDescr.5

F Viewing Logs and Statistics

This chapter explains how to view Avaya Identity Engines Ignition Server logs and statistics. Logs and statistics cover a range of subjects from user/device authentications to the physical operation of the Ignition Server.

Contents

Ignition Server collects system, transaction, and user authorization statistics that the administrator can view in a number of ways:

- "Overview of Logging and Log Types" on page 423
- "Viewing and Managing Logs" on page 424
- "Log Viewer" on page 427
  - "Access Log: RADIUS and TACACS+ Accounting" on page 428
  - "Access Record Details" on page 430
  - "Audit Log" on page 434
  - "Security Log" on page 436
  - "System Log" on page 436
- "Statistics Tab" on page 437
  - "Transactions Tab" on page 438
  - "Directory Services Tab" on page 440
  - "Protocols Tab" on page 440
- "System Health Tab" on page 441
- "Directory Services Status Tab" on page 442
- "AAA Summary Tabs" on page 442
- "User Accounting Tab" on page 445
- "Learned Devices Tab" on page 447
- "Debug Logs" on page 448

Overview of Logging and Log Types

Using Ignition Dashboard, you can view the following log data describing network authentications/authorizations and the operation of the Ignition Server:

- Authentications, authorizations, and provisioning values (visible in separate channels for RADIUS, TACACS+, and Guest Manager). See “Access Log: RADIUS and TACACS+ Accounting” on page 428 and “AAA Summary Tabs” on page 442.
- Highly detailed information about a single login attempt and its results. See “Access Record Details” on page 430.
- List of current logged-in users. See “User Accounting Tab” on page 445.
- List of current AD-authenticated devices. See “Learned Devices Tab” on page 447.
- Audit log of administrator actions on the Ignition Server. See “Audit Log” on page 434.
- Security and system health logs. See “System Health Tab” on page 441.
- Statistics and logs detailing authentication/authorization transactions, including statistics categorized by authentication protocol. See “Statistics Tab” on page 437.
- Directory service interaction statistics and logs. See “Directory Services Tab” on page 440.

You view the Ignition Server logs and accounting information in the **Monitor** view of Ignition Dashboard. (Note that you can also configure the Ignition Server to send its log messages to one or more Syslog servers, and you can export logs to XML-formatted files as explained in “Setting Up Logging” on page 411.)

All messages include a date/time stamp indicating when the logged event occurred, expressed in UTC (Universal Time Code). When viewed in the Log Viewer, the time and date are displayed in local time.

Viewing and Managing Logs

- “Viewing Logs” on page 424
- “Filtering Your View of the Logs” on page 425
- “Managing the Stored Logs” on page 427

Viewing Logs

The **Log Viewer** allows you to view the log messages stored on the Ignition Server. This window displays messages in a tabbed view with one tab per type of log message.

Using the Log Viewer

1. In Dashboard, click **Monitor** to switch to monitor view.

2. Click your node's IP address or name in the tree, click **Log Viewer**, and click a log channel tab, such as **Access** or **Audit**, to load the desired type of messages.

   The window turns a darker shade of grey until it has finished loading the messages. After the messages have loaded, click the paging buttons to move through the loaded messages.

3. To load the latest messages,

   - Click **Refresh**; or
   - Specify/Change the filter and click **Apply Filter** button.

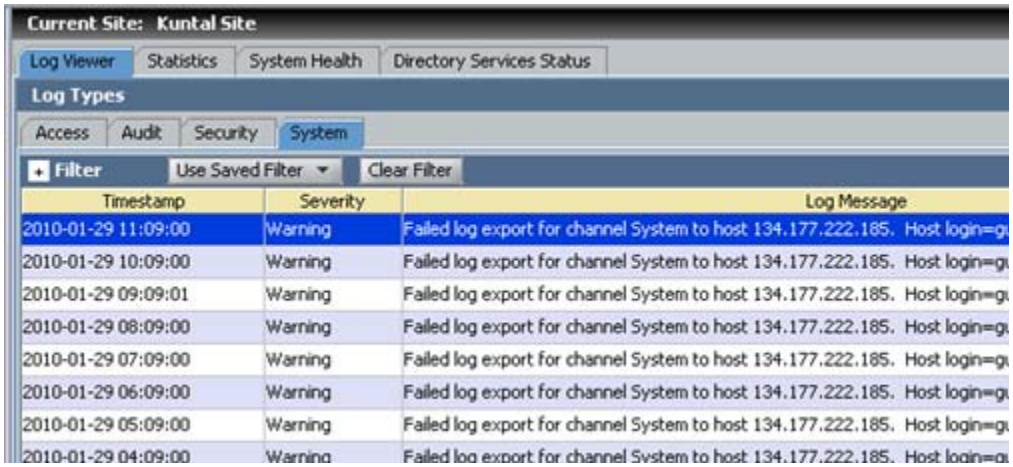
Note: The **Filter** button apply to the currently visible tab only, unless you have set the **Apply filter to all channels** checkbox to ON. (See the next section, Filtering Your View of the Logs.)

Filtering Your View of the Logs

To apply a filter:

1. In Dashboard, click **Monitor** to switch to monitor view.
2. Open the desired log tab. Do one of the following:
   - Click your node's IP address or name in the tree, click **Log Viewer**, and click a log channel tab, such as **Access** or **Audit**, to load the desired type of messages; or
   - Click your site's name in the tree, click one of the **AAA Summary** tabs.
3. In the tab to be filtered, click the plus sign near the top of the tab to display the **Filter panel**.
4. Set your criteria. "Criteria for Filtering Log Messages" on page 426 explains the **Filtering Criteria**.
5. Click the **Apply** button. Only records matching all your criteria are shown.

Figure 117 Filtering Log Messages



Criteria for Filtering Log Messages

This table lists the filter criteria and available test values for each in the logging tabs. Not all criteria are available in all tabs.

Criteria for Filtering Log Messages

Filter Element	Available Settings and Ignition Server Responses
Date & Time Period	Choose one of the following: ▪ **Fixed Period**: Displays logging messages from the last 1, 2, 4, 6, or 8 hours. ▪ **After**: Displays messages timestamped **After** the date and time you specify. Click the calendar icon to set the threshold date and time. ▪ **Between**: Displays messages timestamped **between** the start and end times you specify. Click the calendar icons to set the dates and times.
User Id	Displays messages related to the user login name you specify.
Auth Result	Choose **Accepted** or **Rejected** to display only that type of message.
Record Type	You can limit the results to one of the following types: RADIUS Authentication (includes both authentications and authorizations), RADIUS accounting, TACACS+ Authentication, TACACS+ Authorization, TACACS+ Accounting
Log Level	Available for the *System* log only. Choose a Log Level to display only messages of the severity level you select. From least severe to most severe, the choices are: **Trace, Debug, Info, Warning, Error,** and **Fatal**. For example, selecting **Warning** displays only Warning-level messages, and selecting **Fatal** displays only Fatal-level error messages. **Warning!** When you choose a log level, Ignition Server displays records matching that log level only; it does not display messages of that level and more severe levels, as some systems do!

This table explains the buttons you use to add, remove, and apply

Buttons for Filtering Log Messages

Command or Button	Purpose
Filter	Click the plus sign (+) to display the filter criteria fields. Click the minus (-) sign to hide the fields.
Add Criterion	Adds a new row to specify an additional filter criterion.
Remove Criterion	To remove a criterion row, click the blue arrow  next to the row and click the **Remove Criterion** button.
Use Saved Filter	Click this drop-down list to apply a saved filter.
Clear Filter	Removes filtering to display all records.
Apply	After you have you have specified a filter in the criteria rows, click **Apply** to filter the currently visible tab.
Manage Saved Filters	Lets you rename or delete a saved filter.
Save Filter	Saves the current set of criteria rows as a filter.
Refresh	The Refresh button refreshes the log display for the current tab.

Managing the Stored Logs

If the Ignition Server has exhausted its available space for logs, Ignition Server rotates the logs on a first-in/first-out basis, so that the newest entries overwrite the oldest ones. Ignition Server sends you an alert when it runs out of space and begins overwriting log messages.

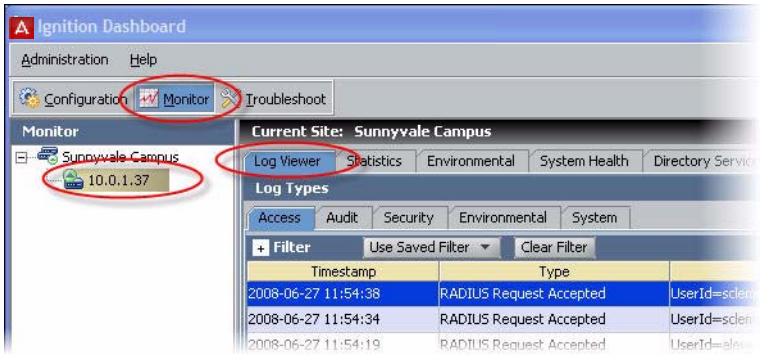
If you want to retain your log messages before they are overwritten, use the log export facility described in “Setting Up FTP Log Export” on page 413.

Log Viewer

The Log Viewer tab is the entry point for viewing the logs stored on the Ignition Server. The tab contains a sub-tab for each type of Ignition Server log, as explained in the following sections:

- “Access Log: RADIUS and TACACS+ Accounting” on page 428
- “Access Record Details” on page 430
- “Audit Log” on page 434
- “Security Log” on page 436
- “System Log” on page 436

Figure 118 The Log Viewer



Access Log: RADIUS and TACACS+ Accounting

The *access log* is the RADIUS accounting log and TACACS+ accounting log. It displays the results of RADIUS and TACACS+ authorization requests, as well as Guest Manager provisioner login attempts. This log appears in the **Access** tab of the **Monitor: Log Viewer** tab.

- "Contents of the Access Log" on page 428
- "Viewing the Access Log" on page 429
- "RADIUS Accounting Messages in the Access Log" on page 429
- "Setting Up Ignition Server To Receive RADIUS Accounting Messages" on page 430
- "Viewing RADIUS Accounting Messages" on page 430

Contents of the Access Log

The Access Log includes all RADIUS and TACACS+ events, including RADIUS and TACACS+ authentication and authorization events. The Access log channel shows the following information:

- Transaction identifier
- User or administrator identifier
- Node ID and Node name
- Request port number
- ASC identifier
- Client/supplicant identifier or MAC address, if available
- List of policies that were triggered, if appropriate
- Result code
- Brief plain-English description of result

Viewing the Access Log

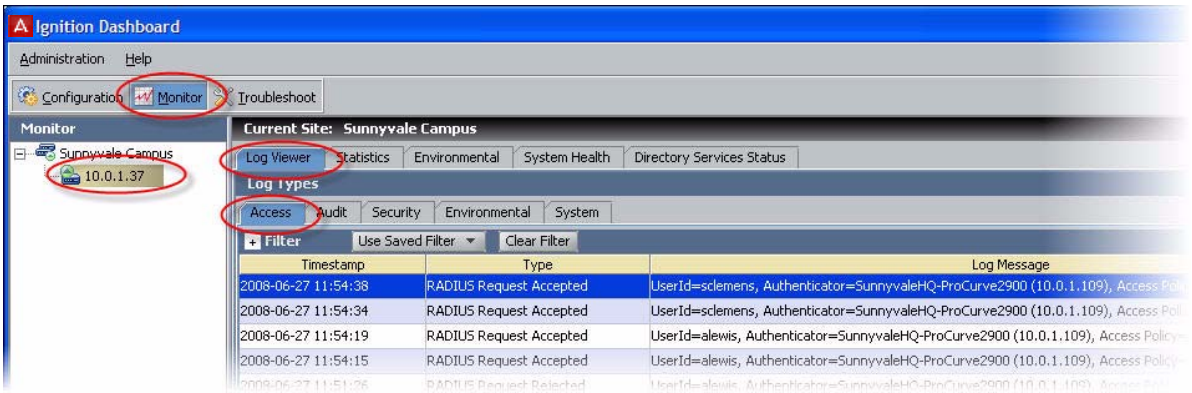
To view the access log:

1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Log Viewer** tab.
4. Click the **Access** tab and scroll or use a filter to find the desired record. Click a record to inspect it. You can view a more detailed description of each access request by opening its *Access Record Details*. See "Access Record Details" on page 430 for more information.
5. You can filter the set of records. See "Filtering Your View of the Logs" on page 425.

RADIUS Accounting Messages in the Access Log

The Ignition Server is a RADIUS accounting server compliant with RFC 2866. Switches, wireless access points, and other network devices send RADIUS accounting messages (START, STOP, and UPDATE events) to the Ignition Server, and the Ignition Server stores, displays, and/or forwards these messages.

Figure 119 The Access Log



For RADIUS, the **Access** tab displays the following types of messages:

- RADIUS Request Accepted
- RADIUS Request Rejected
- RADIUS Accounting

The following types of information are logged in RADIUS accounting messages:

RADIUS Accounting Message Details

Entry Name	Description
acct-status-type	Event type, which is one of (START, STOP, or UPDATE)
acct-session-id	Unique session id useful for matching the start packet to the stop packet
acct-session-time	Total length of a session as of session end time. Only in STOP packets
user name	Login name of the client user
calling-station-id	Unique id of the user's client device. Usually the MAC address of the client device.
acct-input-octets	The number of packets sent to the port over the course of service
acct-output-octets	The number of packets sent to the port over the course of service
framed-IP-address	IP address of the client device

Setting Up Ignition Server To Receive RADIUS Accounting Messages

By default, the Ignition Server listens for RADIUS accounting messages on port 1813. If you want to change the listener port number, see "Editing RADIUS Communication Settings" on page 50. Consult the documentation for your switch or other network equipment for instructions on directing your RADIUS accounting messages to the Ignition Server.

Viewing RADIUS Accounting Messages

To view RADIUS accounting messages, use the **Log Viewer** tab as you would for other logs. Also, you can export your RADIUS accounting messages on a regular schedule as specified using the Configuration: Site: Logging: Export Logs tab. (Click the **Configuration** button at the top of the Dashboard window, click your site's name in the hierarchy tree, click the **Logging** tab, and click the **Export Logs** tab.) See "Setting Up FTP Log Export" on page 413.

Access Record Details

The Access Record Details window shows the submitted details and returned results of a user's or device's login attempt. See these sections for details:

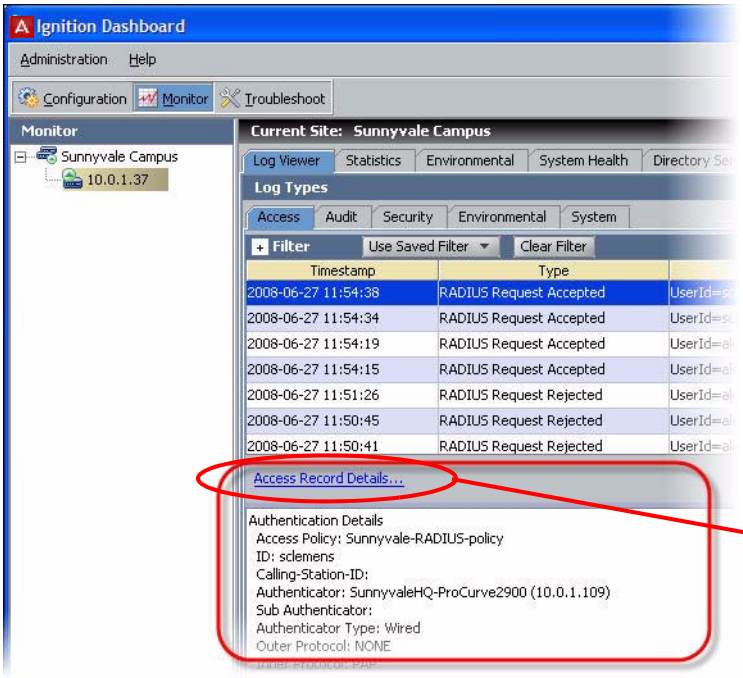
- "Specifying How Dashboard Displays Access Record Details" on page 431
- "Viewing the Access Record Details" on page 432

- “Contents of the Access Record Details” on page 433
- “Activating the Recording of Access Record Details” on page 434

Specifying How Dashboard Displays Access Record Details

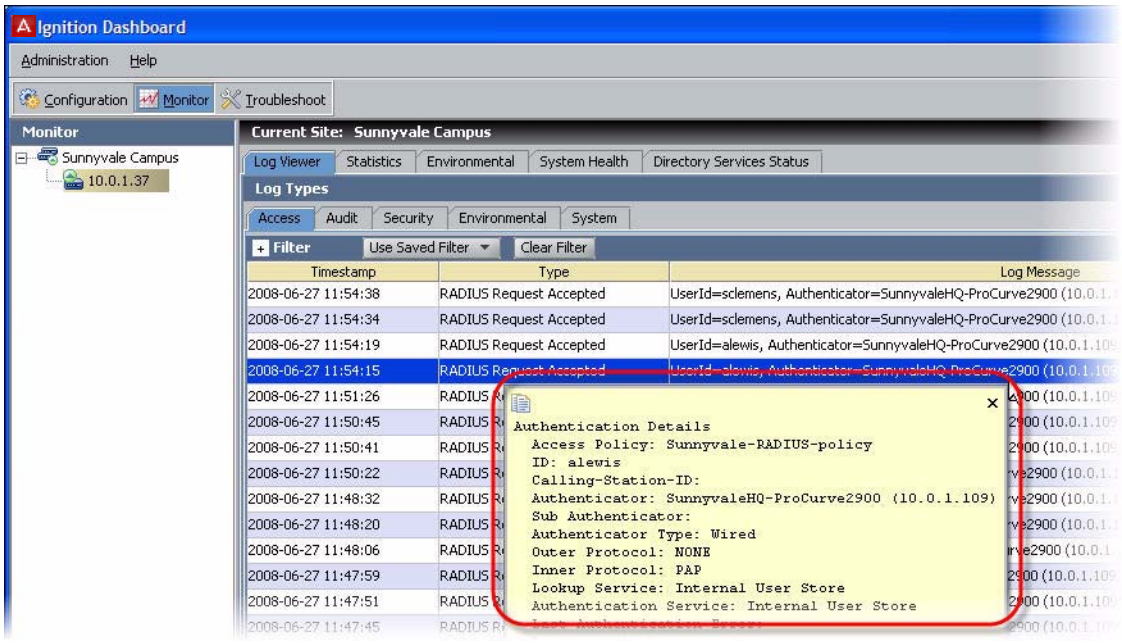
You can have Dashboard display the Access Record Details in a dedicated panel at the bottom of the **Log Viewer** (click **Region at Bottom of Log Viewer** in your **Preferences**), or you can have the Access Record Details appear as tooltips when you click a row in the **Log Viewer** (click **Tooltip** in your **Preferences**). See “Setting Viewing Preferences for the Monitor View” on page 42.

Figure 120 Access Record Details displayed in a dedicated panel



Click to display the Access Record Detail in its own window.

Figure 121 Access Record Details displayed as a tooltip



Viewing the Access Record Details

To inspect an access record:

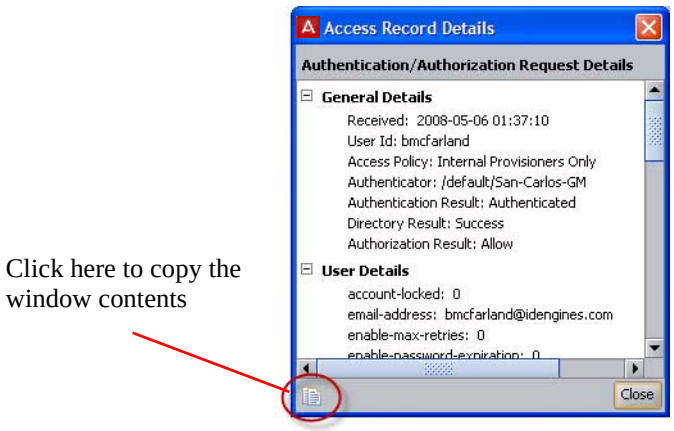
1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Log Viewer** tab.
4. Click the **Access** tab and scroll or use a filter to find the desired record. Click the record.
5. Click the blue text, **Access Record Details...** near the bottom of the window.

   The Access Record Details window appears.

   Hint: To copy its contents for pasting into another application, click the copy icon at the lower left corner of the window. (The copy icon is an image of two sheets of paper.) The contents of the window are placed in

your computer's clipboard as text. You may paste them into any text editor or word processor.

Figure 122 Access Record Details window



See "Contents of the Access Record Details", below, for an explanation of the fields of the **Access Record Details** window.

Contents of the Access Record Details

- **General Details** summarize the results of the login attempt:
  - Received: Time of request
  - User Id: Submitted user name
  - Access Policy: Name of your Ignition Server RADIUS authentication policy used
  - Authenticator: Name of the switch or access point user connected through.
  - Authentication Result: Authenticated or Authentication failed
  - Directory Result: Success or failure of user lookup
  - Authorization Result: Allow or Deny result based on your authorization rules
- **User Details** provide information from the user's record. Most of these fields are available only if the user is stored in the Ignition Server internal store: account-locked, email-address, enable-max-retries, enable-password-expiration, enable-start-time, first-name, group-member, last-name, max-retries, network-usage, office-location, password-expiration (date and time password expires), role, start-time (data and time account becomes usable), title, and user-id.
- **Inbound Attributes** are the incoming name/value pairs received from the authenticator. Usually this is User-Name, State, and Message-Authenticator.

- **Authentication Details** show what type of authentication was attempted. The attributes are Outer Tunnel Type, Outer Tunnel User, Inner Tunnel Type, Inner Tunnel User, Auth Server, and Authentication Result
- **Directory Details** show which user store/authentication server was used to authenticate the user, and which user store provided the user's account details. The fields include Authentication Directory Store Type, Directory Set, Authentication Directory Store Name, Realm, Lookup Directory Store Name, Lookup Directory Store Type, and Directory Result.
- **Authorization Details** show which rule in your Ignition Server authorization policy was used to make the Allow/Deny decision, and what the result was. They include Policy Rule Used and Authorization Result.
- **Outbound Attributes** are the RADIUS and VSA name/value pairs that Ignition Server sent to the authenticator with the authorization.

Activating the Recording of Access Record Details

To turn on or turn off the recording of access record details, do the following:

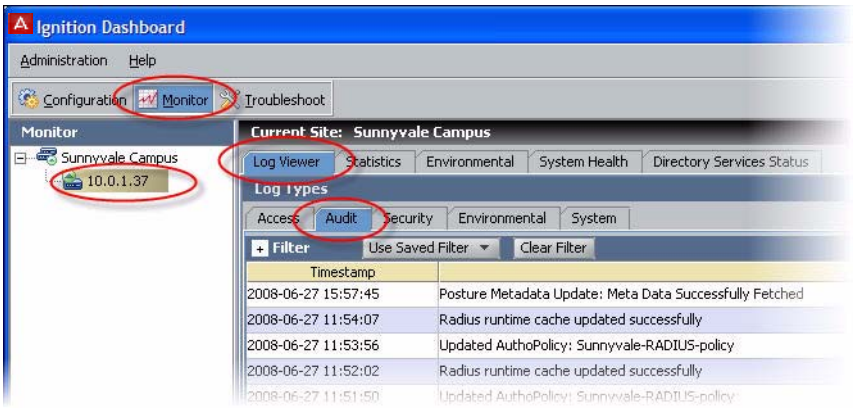
1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Log Viewer** tab.
4. Click the **Configure** button in the upper right corner of your window.
5. In the Configure Log Types window, tick the **Access Details: Enabled** checkbox to turn on detailed logging, or untick it to turn it off.

Audit Log

The audit log records administrative actions done on the Ignition Server.

- "Contents of the Audit Log" on page 435
- "Viewing the Audit Log" on page 435

Figure 123 The Audit Log



Contents of the Audit Log

The Audit Log records Ignition Server administrative actions, including (but not limited to) the following:

- administrative logins and logouts
- shutdowns, reboots
- firmware updates and rollbacks
- system backups and restores
- administrative account adds, edits, and deletes
- user name and password changes
- configuration changes
- policy adds, edits, and deletes
- user/group adds, edits, and deletes
- authenticator/authenticator hierarchy adds, edits, and deletes
- site name changes

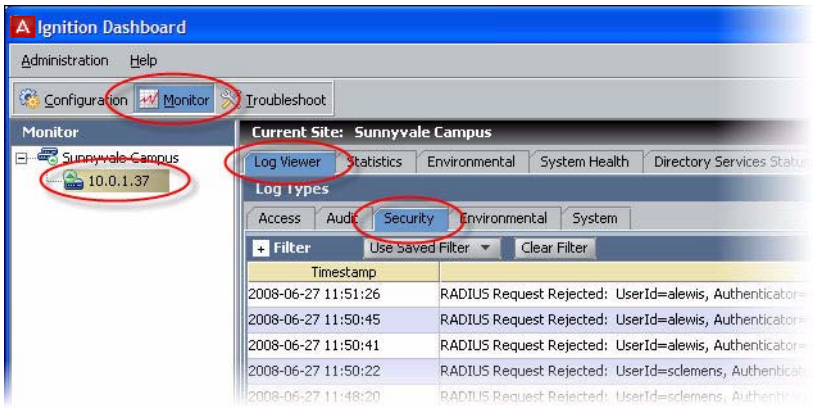
Viewing the Audit Log

1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Log Viewer** tab.
4. Click the **Audit** tab and scroll or use a filter to find the desired record. Click a record to inspect it.
5. You can filter the set of records. See “Filtering Your View of the Logs” on page 425.

Security Log

- "Contents of the Security Log" on page 436
- "Viewing the Security Log" on page 436

Figure 124 The Security Log



Contents of the Security Log

The Security Log lists network-related and Ignition Server-related security events, including:

- Failed authentication requests
- Detection of physical intrusion or tampering of the Ignition Server
- Detection of DoS (Denial of Service) and DDoS (Distributed Denial of Service) attacks
- Any other attempt to breach the security of the Ignition Server

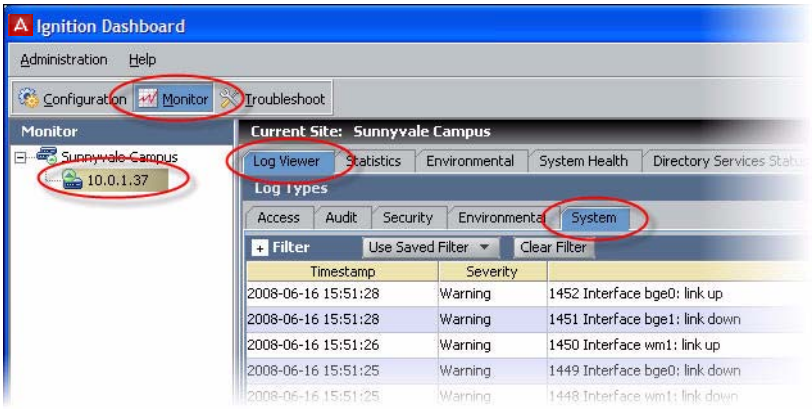
Viewing the Security Log

1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Log Viewer** tab.
4. Click the **Security** tab. Click a record to inspect it.
5. You can filter the set of records. See "Filtering Your View of the Logs" on page 425.

System Log

- "Contents of the System Log" on page 437
- "Viewing the System Log" on page 437

Figure 125 The System Log



Contents of the System Log

Miscellaneous log data from third party software components. Messages logged on this channel include a field denoting a severity classification. If you encounter an error message that has a severity level of *FATAL*, *ERROR* or *WARNING*, you should report it to your Avaya customer service representative.

To minimize the amount of System logging events recorded, select the level, “Fatal.”

Viewing the System Log

1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Log Viewer** tab.
4. Click the **System** tab and scroll or use a filter to find the desired record. Click a record to inspect it.
5. You can filter the set of records. See “Filtering Your View of the Logs” on page 425.

Statistics Tab

The Statistics tab lets you monitor the operation of the Ignition Server. Statistics are shown for the selected node only; click another node in the hierarchy tree to see its statistics. The display is refreshed every 5 seconds, and counters display the count since the last reboot of the Ignition Server.

Figure 126 The Statistics Tab



Note concerning statistics for Ignition Server HA pairs: If you have configured the Ignition Server as a node in an HA-paired node set, the **Statistics Tab** displays a pulldown menu called **Statistics for Node.** If the Ignition Server is standalone, this pull-down menu is not displayed. Use this menu to select the required member of the paired set for which you want to view the associated statistic.

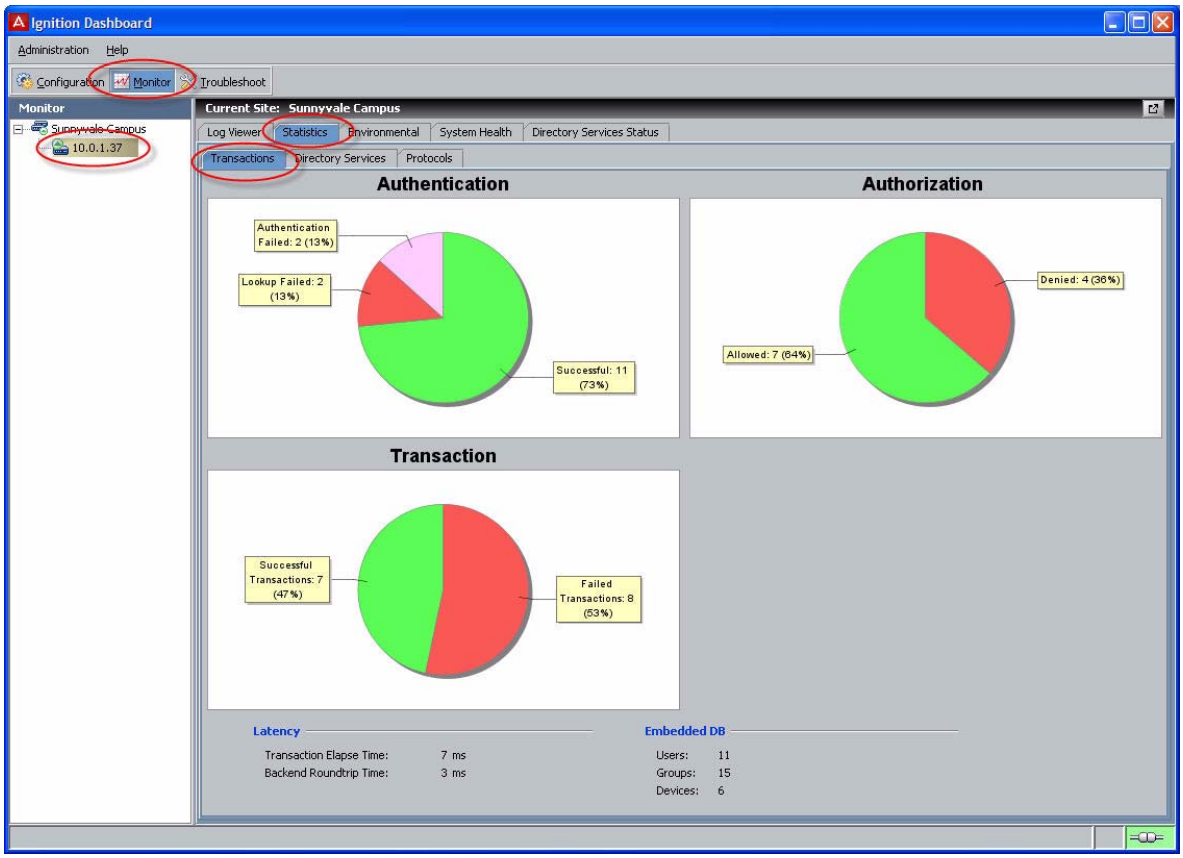
The Statistics tab contains a number of sub-tabs, as explained in the following sections:

- "Transactions Tab" on page 438
- "Directory Services Tab" on page 440
- "Protocols Tab" on page 440

Transactions Tab

- "Contents of the Transactions Tab" on page 439
- "Viewing the Transactions Tab" on page 439

Figure 127 The Transactions Tab



Contents of the Transactions Tab

The transaction statistics appear in the following categories:

- **Authentication**, showing numbers of successful and failed requests, with reasons in specified categories.
- **Authorization**, showing counts for requests allowed and denied.
- **Latency**, displaying the average time required to complete transactions. Ignition Server averages all transactions since the last reboot.
- **Internal DB**, showing counts of users and groups in Ignition Server's local data store.

Viewing the Transactions Tab

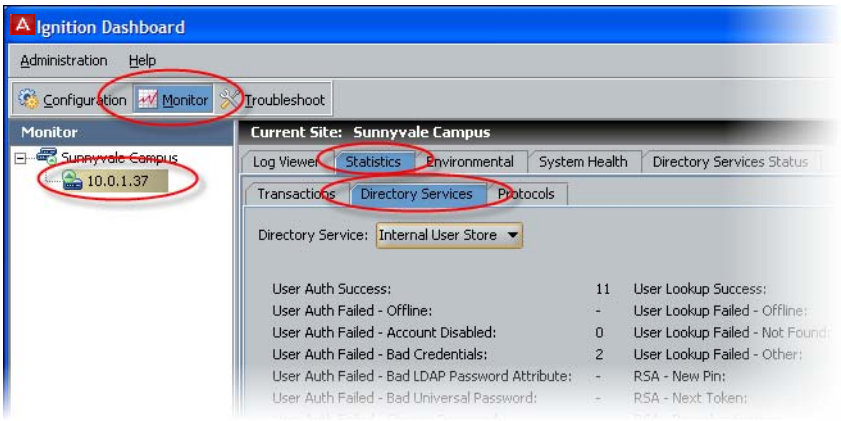
1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Statistics** tab.

4. Click the **Transactions** tab.

Directory Services Tab

- "Contents of the Directory Services Tab" on page 440
- "Viewing the Directory Services Tab" on page 440

Figure 128 Statistics: Directory Services Tab



Contents of the Directory Services Tab

The **Directory Services** tab tracks transactions between Ignition Server and each user data store. To view transactions counts, select the name of your data store in the **Directory Service** drop down list.

The statistics shown are:

- **Transactions**: The number of user look-up/authentication attempts Ignition Server has performed against the specified directory service.
- **Failed Authentication Attempts**: The number of failed user look-up/authentication attempts Ignition Server has performed against the specified directory service. This includes every failure due to invalid credentials or failure to find the user. If fallthrough is turned on for the directory service, then each failure in this Service that results in a fallthrough to another service is counted as one failed attempt.

Viewing the Directory Services Tab

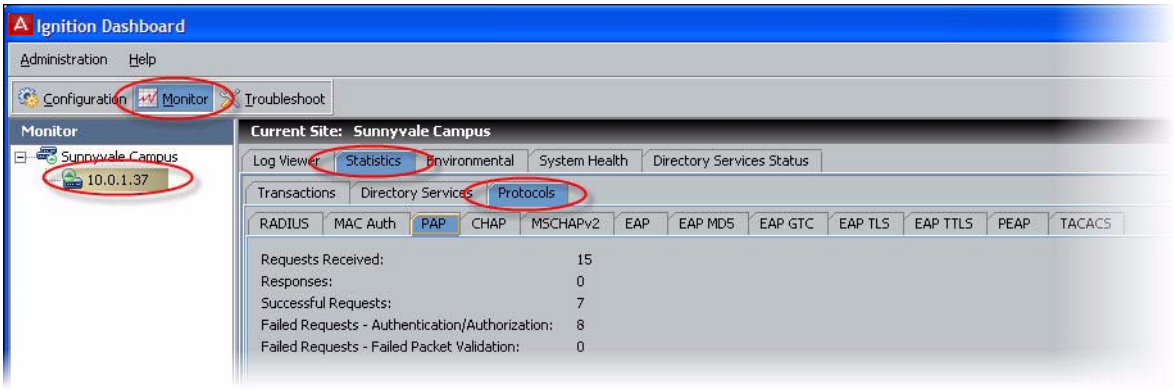
1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Statistics** tab.
4. Click the **Directory Services** tab.

Protocols Tab

- "Contents of the Protocols Tab" on page 441

- “Viewing the Protocols Tab” on page 441

Figure 129 Example Protocols Statistics for RADIUS



Contents of the Protocols Tab

Clicking **Protocols** displays set of sub-tabs, one per protocol, each offering statistics on Ignition Server’s communications via the selected protocol.

For example, the **RADIUS** sub-tab of the **Protocols** tab provides detailed information about the quantities of RADIUS packets Ignition Server has received and processed (received, sent, accepted, rejected, and so on).

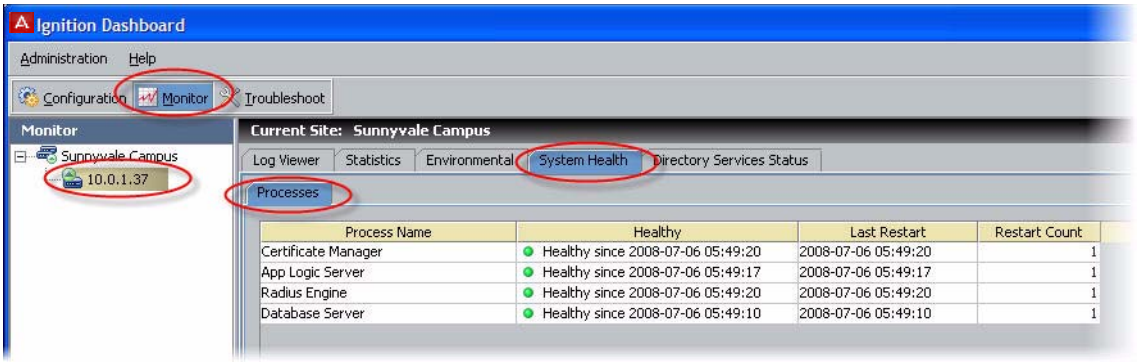
Viewing the Protocols Tab

1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **Statistics** tab.
4. Click the **Directory Services** tab.

System Health Tab

- “Contents of the System Health Tab” on page 442
- “Viewing the System Health Tab” on page 442

Figure 130 The System Health tab



Contents of the System Health Tab

The **System Health** tab displays the operational status of processes running on the Ignition Server.

Viewing the System Health Tab

To view this tab, do the following:

1. In Ignition Dashboard, click **Monitor** to show the system monitoring view.
2. Click the IP address or name of your node in the tree.
3. Click the **System Health** tab.

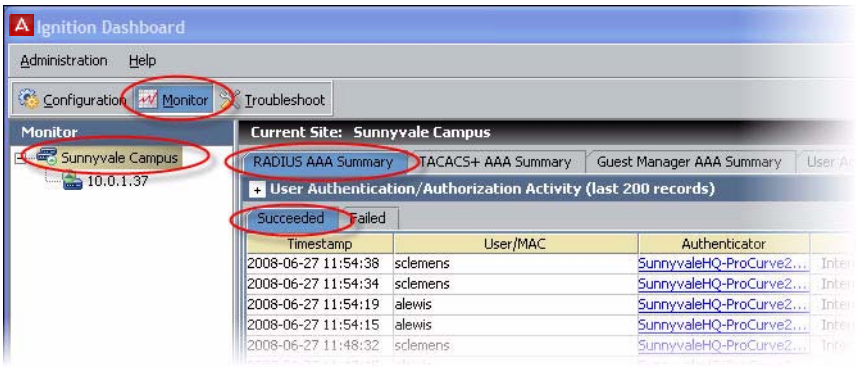
Directory Services Status Tab

See “Checking Directory Service Connections” on page 170.

AAA Summary Tabs

- “Contents of the AAA Summary Tabs” on page 443
- “Viewing the AAA Summary Tabs” on page 444

Figure 131 The AAA Summary Tabs



Contents of the AAA Summary Tabs

The **AAA Summary tabs** show the list of recently logged in users (in the **Succeeded** tab) and a list of recent sign-on failures (in the **Failed** tab). Each table contains a row for each active or rejected user session.

Succeeded Tab

The **Succeeded** tab of the **AAA Summary** displays the following information:

- **Timestamp**: the timestamp for the request
- **User/MAC/Provisioner**: the user name or MAC address of the connecting user, device, or Guest Manager provisioner
- **Authenticator**: the access point at which the request was made
- **Server**: for provisioner logins, this is the Guest Manager server where the login occurred
- **Directory**: the name of the directory service that authenticated the user or device
- **Auth Protocol/Base Protocol**: the authentication protocol

You can adjust the width of each column as necessary.

Additional RADIUS request details for the requests shown in this window can be viewed in the Log Viewer tab. For details, see “Viewing and Managing Logs” on page 424.

Failed Tab

The **Failed** tab of the **AAA Summary** displays the following information:

- **Timestamp**: the timestamp for the request
- **User/MAC/Provisioner**: the user name or MAC address of the connecting user, device, or Guest Manager provisioner
- **Authenticator**: the switch or access point at which the request was made
- **Server**: for provisioner logins, this is the Guest Manager server where the login occurred
- **Directory**: if the user look-up succeeded, this column shows the name of the directory service that authenticated the user or device; if the user look-up failed, this column shows the name of the last-searched directory service in your directory set.
- **Auth Protocol/Base Protocol**: the authentication protocol

- **Authenticated**: A red *x* indicates the user or device authentication failed. A blue check mark indicates the authentication succeeded but the authorization rules failed to authorize the user.
- **Reason for Rejection**: This column displays short explanation of the reason for rejection. The most common reasons are:
  - *User Not Found*: Authentication failed because no matching user account was found for the submitted user name. Refer to the **Directory** column for the name of the last-searched directory.
  - *Invalid Credentials*: User account was found, but authentication failed because the submitted credentials were incorrect.
  - *No Rule Applicable*: User authentication succeeded, but the authorization failed because no ALLOW rule was triggered.
  - *Deny*: User authentication succeeded, but the authorization failed because a DENY rule was triggered.

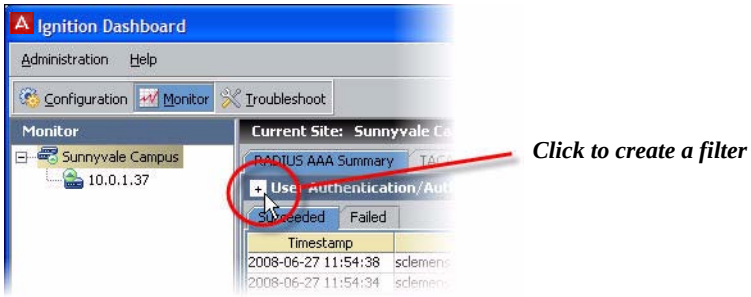
Additional RADIUS request details for the requests shown in this window can be viewed in the Log Viewer tab. For details, see "Viewing and Managing Logs" on page 424.

Viewing the AAA Summary Tabs

Open the AAA Summary tabs as follows:

1. Click **Monitor** in the main Dashboard window.
2. Click your site name in the **Monitor** hierarchy tree.
3. Click one of the **AAA Summary** tabs, and click the **Succeeded** or **Failed** tab.
4. If you want to filter the contents of the tab, click the plus-sign above the **Succeeded** tab, and create a filter as explained in "Filtering Your View of the Logs" on page 425.

Figure 132 Click the plus-sign to filter



Specifying the Number of Records to be Shown

The **RADIUS AAA Summary** tab, the **TACACS+ AAA Summary** tab, and the **Guest Manager AAA Summary** tab show the most recent set of login attempts. You establish the maximum number of records to be shown by making a setting in the **Preferences** Window. See "Setting Viewing Preferences for the Monitor View" on page 42.

The limit on the total number of entries is enforced across all three tabs, any tab may be empty if the other tabs contain enough recent records to satisfy the limit you set in the **Preferences** window.

User Accounting Tab

- "Contents of the User Accounting Tab" on page 445
- "Viewing the User Accounting Tab" on page 446

Contents of the User Accounting Tab

The **User Accounting** tab lists currently connected users. You can filter the contents of this tab by user name, and you can export the tab's contents.

Accounting Data

The main table in the User Accounting window displays a set of RADIUS attributes for each active session:

- **User Name**: User domain and user account name
- **Connected Time**: The date and time at which the session was initiated
- **Framed IP Address**: IP address of the user's client device (RADIUS protocol)
- **Authenticator**: Name of the switch or AP through which the client connected
- **Calling Station Id**: Identifier of the user's client device; usually the client device's MAC address
- **Session Id**: Unique identifier of the user's RADIUS session

Filter Button

To filter the contents of the User Accounting window, do the following:

1. Enter the desired user name in the **User Name Starts With** field.
2. Click the **Filter** button.

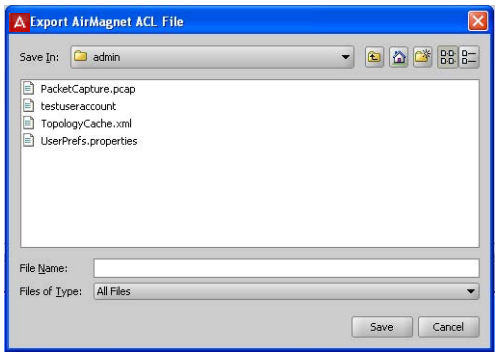
Ignition Server displays the accounting information filtered only for the name input in the **User Name Starts With** field.

Export Button

The **Export** button of the User Accounting window lets you export session audit data for a selected user. To do this:

1. From the Dashboard main window, click **Monitor** and click your site's name in the tree.
2. Click the **AAA Summary tab**.
3. At the bottom of the window, click the **Details** button to launch the User Accounting window.
4. Select the row containing the session audit data to be exported.
5. Click **Export**. Ignition Server requires you to enter the name for the exported file.

Figure 133 Specifying Location for Account Export



6. Enter the name (and specific location, if other than the default) for the exported file.
7. Click **Save**.

Ignition Server exports the accounting information and saves the file with this name in the desired location. For additional RADIUS accounting information, see "Access Log: RADIUS and TACACS+ Accounting" on page 428.

Refresh Button

To load the latest session audit data in the User Accounting tab, click the **Refresh** button.

Viewing the User Accounting Tab

1. Click **Monitor** in the main Dashboard window.
2. Click your site name in the **Monitor** hierarchy tree.
3. Click the **User Accounting** tab.

4. If you want to filter the contents of the tab, type a user name or the first few characters of a user name in the **User Name Starts With** field and click **Apply Filter**.

Learned Devices Tab

- "Contents of the Learned Devices Tab" on page 447
- "Viewing the Learned Devices Tab" on page 447
- "Filtering the Learned Devices Tab" on page 447
- "Revoking the Session of a Machine-Authenticated Device" on page 448

Contents of the Learned Devices Tab

The Learned Devices tab displays a list of devices that have authenticated to Ignition Server via Windows Machine Authentication and whose sessions are currently valid. In Ignition Server such devices are often called "authenticated assets."

Your authorization rules can require that users connect using only devices with a valid session. You can use this tab to revoke the current session of a device, as explained in "Revoking the Session of a Machine-Authenticated Device" on page 448.

The expiration date and time for each device's authentication is displayed in the **Expires** column. Each authentication lasts for the device TTL period configured in the Learned Device TTL window. See "Setting TTL for Windows Machine Authentication" on page 305.

Use the **Back** and **Next** buttons to move through the list. To filter the list, see "Filtering the Learned Devices Tab" on page 447.

Viewing the Learned Devices Tab

1. Click **Monitor** in the main Dashboard window.
2. Click your site name in the **Monitor** hierarchy tree.
3. Click the **User Accounting** tab.
4. If you want to filter the contents of the tab, type a user name or the first few characters of a user name in the **User Name Starts With** field and click **Apply Filter**.

Filtering the Learned Devices Tab

You can filter the list of Learned Devices by ticking the **Specify Criteria** checkbox and:

- typing a full or partial **MAC Address** to be matched; or
- specifying an **Expiration Date** (and time) **Before** or **Expiration Date** (and time) **After** criterion.

- specifying a device **Name** or partial name to be matched.

Click **Apply Filter** to apply the filter.

See "Requiring the User to Connect Using a Machine Authenticated-Device" on page 340 for more information.

Revoking the Session of a Machine-Authenticated Device

You can revoke sessions as follows:

- To revoke a device's session, click on its row to select it and click the Delete button.
- To revoke all device sessions, click Delete All.

Debug Logs

The debug logs include data used to debug problems in system configuration and operation, and to determine the root cause of failed authentication requests. Messages logged on this channel include a field denoting one of the following severity levels:

- FATAL: Messages describe catastrophic failures that result in a reboot of the system. FATAL messages are always reported to the debug channel, and can not be disabled. All FATAL debug messages should be reported to your Avaya Customer Service representative.
- ERROR: Messages describe system failures from which Ignition Server invoked automatic recovery procedures. ERROR messages are always reported to the debug channel, and can not be disabled. All ERROR messages should be reported to your Avaya Customer Service representative.
- WARNING: All errors in system configuration or detected failures/ anomalies of network components with which Ignition Server interacts. Examples include loss of connectivity to a configured directory store, unavailability of a configured Syslog server or a port down event on a configured network connection. WARNING messages are useful for debugging your system configuration and overall network status. WARNING messages are always reported to the debug channel, and can not be disabled.
- INFO: These messages are used exclusively to perform real-time debugging of failed authentication events. In the event that an Ignition Server administrator encounters a problem with the authentication of one or more network users, the administrator can enable INFO messages through the Ignition Dashboard, initiate an authentication request and trace the root cause of the resulting authentication failure. **Note that due**

to the amount of log data provided, enabling INFO level debug messages can have a detrimental effect on the real-time performance of the Ignition Server. INFO level debugging should only be enabled for brief periods while diagnosing authentication failures. INFO messages are disabled by default

G Troubleshooting

This chapter explains how to generate trouble tickets and lists solutions for common errors that can occur when configuring Avaya Identity Engines Ignition Server.

Contents

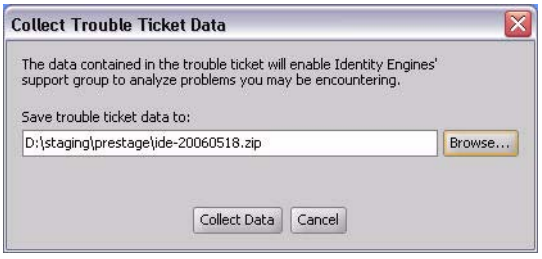
- "Trouble Ticket Generation" on page 451
- "Troubleshooting Common Problems" on page 452

Trouble Ticket Generation

In the event of a fault in your Ignition Server, you can generate a trouble ticket file that the Avaya support staff can use to diagnose your problem. To do this:

1. In Dashboard's Configuration hierarchy tree, right-click the name of your site and select the command, **Trouble Ticket**...
2. In the **Collect Trouble Ticket Data** Window, click **Browse**.
3. Select the directory where you want to save the trouble ticket file. Type a name for the file. Click **Save**.
4. Click **Collect Data**.

Figure 134 The Collect Trouble Ticket Data Window



5. Ignition Server displays a progress bar. On completion, it displays the message:



6. Contact technical support for instructions on uploading the file to Avaya. See “Customer service” on page 21.

Troubleshooting Common Problems

The following sections offer solutions and workarounds for commonly reported issues:

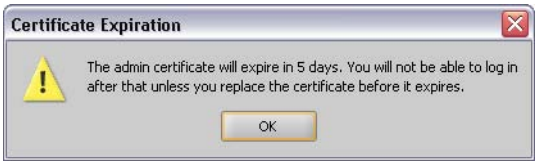
- “Problem: Cannot Connect to Ignition Dashboard” on page 452
- “Problem: Connecting Dashboard to Ignition Server Fails” on page 453
- “Problem: Authorization policy stops working unexpectedly” on page 453
- “Problem: Authentication fails on Active Directory” on page 455
- “Problem: HA Set-up Fails” on page 456
- “Problem: Errors Occur During Directory Service Set-Up” on page 457
- “Problem: Unable to Map Virtual Attributes” on page 458
- “Problem RADIUS Proxy Service Fails” on page 458

Problem: Cannot Connect to Ignition Dashboard

Firewall Settings Make sure TCP port 23457 is reachable on the computer where you have installed Ignition Dashboard. Check your firewall settings to make sure this port is not blocked.

Certificate Expiration Ignition Server Dashboard uses a digital certificate to prove its identity. When starting up, the application warns you if the certificate is due to expire soon.

Figure 135 Warning That Certificate Will Expire Soon



If you receive this warning, you must replace Ignition Dashboard certificate as soon as possible. If the certificate expires, you can no longer manage the Ignition Server because the Dashboard is no longer able to log in. For instructions on replacing the certificate, see “Replacing the Admin Certificate” on page 80.

Concurrent Sessions Not Allowed The Ignition Server permits only one Ignition Dashboard session at a time. If another administrator is logged in, you must have him or her log off before you can log in.

Problem: Connecting Dashboard to Ignition Server Fails

Possible Cause: Firmware version incompatible with Ignition Dashboard

Details: When you attempt to connect to Dashboard, the connection attempt fails with the error "The Ignition Server is incompatible with the UI." This occurs if the firmware version on the Ignition Server is not supported by the current version of Ignition Dashboard.

Solution: Use the Ignition Server front panel to check the firmware version, and log in using a compatible version of Ignition Dashboard. If your PC does not have a compatible, installed version of Dashboard, download the compatible Dashboard installer from the Ignition Server support site www.avaya.com/support.

Problem: Ignition Server fails to respond to RADIUS and/or TACACS+ requests

Troubleshooting tips: Check the following logs to diagnose the problem:

1. Make sure the RADIUS and/or TACACS+ service is enabled as shown in:
   - for RADIUS: "Setting up Ignition Server's RADIUS Service" on page 50;
   - for TACACS+: "Turning on the Ignition Server TACACS+ Service" on page 310.
2. Check the **Log Viewer: Security** tab to see if Ignition Server dropped the request. See "Security Log" on page 436. The message "Packet from unknown authenticator dropped" can mean that you did not define the authenticator (or did not define it correctly) in Ignition Server. See "Creating an Authenticator" on page 96.
3. Check the **System Health** tab to make sure Ignition Server's **RADIUS Engine** is running. See "System Health Tab" on page 441. If the RADIUS Engine is *not* running, contact Avaya customer support for help.

Problem: Authorization policy stops working unexpectedly

Possible Cause: Underlying, but needed, data element has been deleted

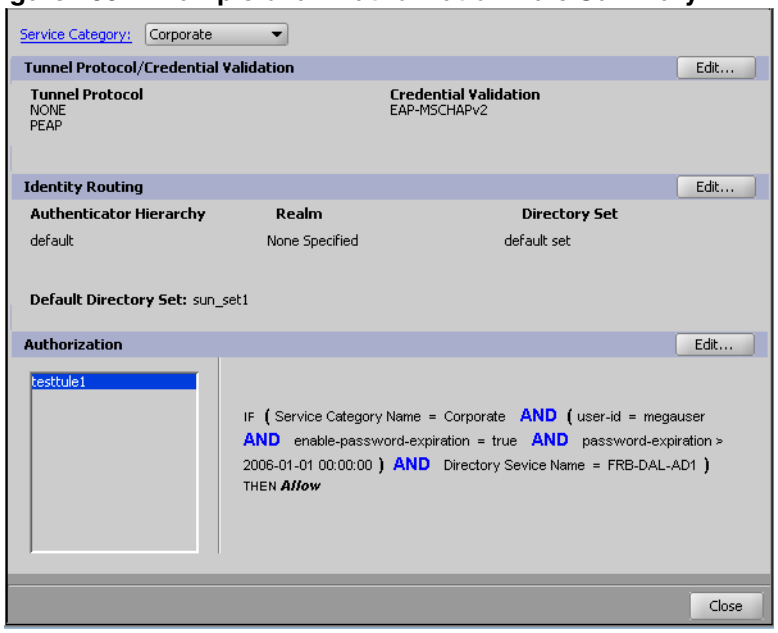
Details: A previously working authorization policy fails because some of its required data has been deleted. Ignition Server does not perform integrity checking on authorization policies. If you have renamed or deleted one or more of the data items associated with a policy, that policy might no longer work as expected.

Note that each authorization policy can use data associated with the following Ignition Server data elements:

1. authenticators
2. authenticator bundles
3. authenticator hierarchy (containers)
4. access policies
5. directory services within directory sets
6. virtual groups
7. virtual attributes

Example: The figure below displays the contents of “testrule1”, an authorization rule that belongs to the access policy “Corporate”. Deleting or renaming the access policy “Corporate” and/or the directory service name “FRB-DAL-AD1” breaks the contents of “testrule1”. (At least one of the constraints of the rule is no longer applicable.) As a result, Ignition Server cannot correctly use this rule to assess an incoming request.

Figure 136 Example of an Authorization Rule Summary



Solution: Use the following workaround to fix broken authorization policies if you have renamed or deleted one of the elements listed above.

1. In Dashboard’s Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click the name of your access policy. Click the

Authorization Policy tab and click Edit. The window lists the authorization rules of the access policy.

- Highlight a rule in the displayed list of Rule Names.

  Virtual Attribute Error: Check each policy for the string, <Invalid Reference>. Where you find this string, edit the policy so that it uses the updated name.

  Other Elements: Update the contents of the Rule Summary for the selected rule.

- Repeat this for each rule in the displayed list of Rule Names for the selected access policy.

2. Repeat Step 1 for each affected access policy.

Problem: Authentication fails on Active Directory

The sections below explain common failures that can occur *in Active Directory environments*. For more general authentication troubleshooting, see "Troubleshooting User Lookup and Authentication" on page 170.

Possible Cause: AD port blocked by firewall

Details: If Ignition Server cannot reach port 445 on the Active Directory server, then PEAP-MSCHAPv2 Authentication fails. This happens because Ignition Server's calls to the AD server's Netlogon service fail.

Solution: Edit your firewall settings as explained in "Preparing to Connect to Active Directory" on page 139.

Possible Cause: Ignition Server machine deleted from the AD domain

Details: If the entry for the Ignition Server is removed from Active Directory's "Computers" list, Ignition Server loses its AD connection and AD-based authentications fail.

Solution: Deleting the computer account from AD is not recommended. If this happens and the connection is lost, you must for Ignition Server to rejoin the domain as follows:

1. In Ignition Dashboard, click the **Monitor** button to show the Monitor view.
2. Click the IP address or name of your node.
3. Click the **Directory Services Status** tab.
4. Click on the name of your AD directory service to select it.
5. Click the **Refresh Cache** button. This forces a rejoin.

Problem: HA Set-up Fails

Possible Cause: No network route between your HA ports

Details: If the HA Configuration Wizard fails during Step 11 of the procedure "Creating an HA Pair", a bad network route might be the cause.

Solution: Fix this as follows:

1. Do not dismiss the HA Configuration Wizard.
2. Repair the network connection between the HA port on your first Ignition Server and the HA port your second Ignition Server.
3. Launch a new session of Dashboard and log in to the first Ignition Server. (Leave the existing Dashboard session running, but ignore it for now.)
4. In the new session of Dashboard, ping the HA port of the second Ignition Server. To do this, click the **Troubleshoot** button at the top of the Dashboard window; click on the *first* Ignition Server's node name or IP address in the hierarchy tree; click **Network** and go to **Ping Test**; type the IP address of the *second* Ignition Server's HA port, set the number of packets to send, and click **Start**.

   If the test fails, fix your network connection. If it succeeds, continue to the next step.
5. In the new session of Dashboard, log out from the *first* Ignition Server and log in to the *second* Ignition Server. Perform another ping test: Click the **Troubleshoot** button at the top of the Dashboard window; click on the *second* Ignition Server's node name or IP address in the hierarchy tree; click **Network** and go to **Ping Test**; type the IP address of the *first* Ignition Server's HA port, and click **Start**.

   If the test fails, fix your network connection. If it succeeds, continue to the next step.
6. Close the new session of Dashboard.
7. Return to the Wizard, dismiss the error window, and in the Wizard window, click **Finish**. The Wizard creates the HA pair.

Problem: Two primary HA nodes detected

When running Ignition Dashboard, if you see the message, "**Error: Two primary HA nodes have been detected**," please follow these instructions. Here, we'll refer to the Ignition Server s as the *first* Ignition Server (the one you just connected Dashboard to) and the *second* Ignition Server:

1. In Dashboard, ping the HA port of the second Ignition Server. To do this:
   - Dismiss the Error message if it is still visible.

- Click the **Troubleshoot** button at the top of the Dashboard window; click on the *first* Ignition Server's node name or IP address in the hierarchy tree; click **Network** and go to **Ping Test**; type the IP address of the second Ignition Server's HA port, set the number of packets to send, and click **Start**.

  If the ping succeeds, go to Step 2.

  If the ping fails, check connectivity between the HA ports of the two Ignition Server s. The HA ports should be connected directly.

2. Ping the first Server from the second Server:
   - Use the **Administration: Logout** command to disconnect Dashboard from the first Ignition Server.
   - Use the **Administration: Login** command to connect to the second Ignition Server.
   - Click the **Troubleshoot** button at the top of the Dashboard window; click on the *second* Ignition Server's node name or IP address in the hierarchy tree; click **Network** and go to **Ping Test**; type the IP address of the *first* Ignition Server's HA port, set the number of packets to send, and click **Start**.

     If the ping succeeds, go to Step 3.

     If the ping fails, check the Ethernet cable connection between the HA ports of the two Ignition Server s. The HA ports should be connected directly. After the cable connection is restored, Ignition Server reconnects the HA pair. If it does not reconnect, proceed to Step 3.
3. Create a trouble ticket and send it to Avaya support. See "Trouble Ticket Generation" on page 451.

Problem: Errors Occur During Directory Service Set-Up

Possible Cause: Ignition Server failed to parse your directory schema.

Details: When you click the **Test Connections** button in the Directory Services panel, Ignition Server returns the message, "Could not parse schema." If you see this message, it means that Ignition Server could not read the schema because Ignition Server is incompatible with your directory version or vendor, or because you have modified your schema is a manner that Ignition Server's parser cannot interpret.

The message, "Could not parse schema" *does not necessarily mean* that you cannot authenticate against the directory. If Ignition Server returns this message, then in the typical case, you will *not* be able to map virtual attributes to the directory, but you *will* be able to authenticate against the directory and map virtual groups to it. See "Troubleshooting User Lookup and Authentication" on page 170.

Solution: If this error occurs:

- If this error occurs and you *do not* plan to use virtual attributes, then you can ignore this error message and continue using the directory
- If this error occurs and you *do* plan to use virtual attributes, open the **Log Viewer tab** in Dashboard and click on the **Debug** tab. The parse failure generates a *Warning*-level message in this channel. Note the error message and contact Avaya support as shown in "Customer service" on page 21.

Problem: Unable to Map Virtual Attributes

See "Problem: Errors Occur During Directory Service Set-Up" on page 457.

Problem RADIUS Proxy Service Fails

Troubleshooting tips:

RADIUS Server

- If you are using an Ignition Server HA setup as the RADIUS Proxy, the keepalive requests are sent using the individual node's IP address. In that scenario, make sure that you add all the three authenticators pointing to each node's IP address of the interface to which RADIUS is bound and the VIP IP address. For simplicity, put all of the three authenticator IP's under one authenticator container.

Proxy Server

- Make sure that the forwarding and remote RADIUS servers are able to communicate:
  - Use the **Test Configuration** button on your forwarding server's "proxy" directory service entry to test the remote server. If the test fails, check the access log at the remote server to check why the request was dropped. Also, make sure that you configure the keepalive username and password at the forwarding server.
  - It is not necessary to provide a valid username/password. Invalid credentials which result in sending an Access reject from the RADIUS server are enough to establish the connectivity.
- You can also test the remote proxy server from the **Troubleshoot** menu in Dashboard.
  1. From the **Troubleshoot** menu in Dashboard, click the **Directory Service Debugger** and then click the **Process Request** sub-tab.
  2. In the **Directory Set** section, use the drop-down menu to select a directory set that corresponds to your proxy server.

3. Enter a valid username and password to the data store that will be searched on your remote server.
4. Click **Send Request** and wait for the test results to appear in the **Results** window below.

Logging and Monitoring RADIUS Proxy

Monitoring at the RADIUS Proxy Using Statistics

- The RADIUS Proxy keeps statistics on how many user auth requests are forwarded and received from the RADIUS server.
- From the **Monitor** menu in Dashboard, click the **Statistics** tab, click the **Directory Services** tab, and select the appropriate Proxy Server.

Monitoring at the RADIUS server

- Since the RADIUS Server handles all the authentication and authorization and the Proxy acts as a regular authenticator, the usual monitoring tools apply.
  - Access logs for user auth requests.
  - View various statistics such as **Transactions** and **Protocols**.

Glossary

This chapter contains the definitions for terms relevant to Ignition Server.

802.1X

The 802.1x network authentication standard is the technical underpinning for all that we do at Avaya. Also known as 802.1X port-based security, 802.1X is the IEEE standard for authenticating users and devices before they are allowed to connect to a wired or wireless LAN. An 802.1X authentication scheme provides authorization to devices that attempt to attach to a LAN port, establishing a point-to-point connection if authentication succeeds, or preventing access from that port if authentication fails. To connect, the user or device must prove its identity to an authentication server (RADIUS or TACACS+) before it/he can use the network. Ignition Server supports RADIUS authentication but not TACACS+.

By implementing an 802.1X network authentication layer using a tool such as Avaya Identity Engines Ignition Server, you reduce the likelihood of unwanted users and unwanted devices joining your network. By using an identity-aware RADIUS server such as Ignition Server, you further increase security, since you can trace each network session to an individual user or device account.

AAA

AAA stands for “authentication, authorization, and accounting.” These are the three primary services required by a network access server or network access protocol. All three services are logically independent and may be separately implemented with the output of each used as the input of the next.

Authentication is the verification of the credentials of a user or a device. Authorization is the process of determining the type of activities that are permitted. Auditing/accounting is keeping track of the attributes of the user’s network session and the activities of the authorized user.

access policy

An access policy is a set of authorization and authentication rules applied to an authenticator or authenticators. Each access policy acts like a virtual RADIUS server, with it’s own set of rules and its own set of user databases for authentication. Access policies replace the discontinued concept of service categories.

Access Portal

Access Portal is a virtual machine based captive portal and firewall distribution that controls the access of client devices to the network.

access switch

An access switch is a layer-2 switch directly connected to the Ignition Server.

Active Directory

Microsoft's directory database for Windows 2000 (and later) networks. Active Directory stores information about users, groups, organizational units, and other kinds of management domains and administrative information about the network.

administrative machine

The machine on which you run Ignition Dashboard, Ignition Server's user interface application.

attributes

Information about users (and other entities) represented in directories and databases.

auditing

Logging, monitoring, accounting, alerting, and reporting on policy, user, and resource activity, usage, and security.

authenticator

An authenticator is a network device, usually a switch, wireless access point, VPN concentrator, or other 802.1X-compliant device, that forces a user or device to authenticate before it grants a network session. The authenticator acts as the policy enforcement point and, when it receives the ALLOW or DENY response from the RADIUS server (for example, the Ignition Server RADIUS server), it allows or denies the session. This is one of the five players in the authentication transaction: supplicant, authenticator, RADIUS server (the Ignition Server), directory service, and, optionally, authentication server.

authentication

The process of verifying a user's (or device's) credentials to confirm their identity.

authentication server

1. (*Avaya Identity Engines Ignition Server usage*) A strong authentication server such as an RSA Authentication Manager or Safe Words Server that authenticates the user credential. In Ignition Server, an authentication server and a directory server work in tandem. The authentication server makes the *authentication decision* by evaluating the user credentials, and the directory

server provides user attributes and group associations that form the basis for *authorization decision* to be made in Ignition Server. In the Avaya context, the authentication server is one of the five players in the authentication transaction: supplicant, authenticator, RADIUS server (the Ignition Server), directory service, and, optionally, authentication server.

2. (*general usage*) The PDP in an 802.1X authentication transaction. For example, a RADIUS server such as the Ignition Server. In RADIUS and other network access terminologies, the term “authentication server” usually refers to the component on your network that has responsibility for making sure the user or device gets authenticated when he/she/it tries to join the network. The authentication server often delegates the authentication task to one or a combination of services such as Active Directory, an LDAP server, and/or a RSA Authentication Manager that can authenticate the user credential. Ignition Server has the advantage of being very flexible in how it delegates authentication.

authenticator bundle

A collection of authenticators that are on the same Subnet and which share common attributes.

authorization

The process of deciding whether a user (or device) is allowed to access the network based on a set of rules.

authorization policy

(See “policy”.)

DER format

DER stands for distinguished encoding rules, a method of uniquely representing any given digital object as a binary string when the object can be described in the so-called ASN.1 (Abstract Syntax Notation).

directory

An organized list of persons, departments, affiliations, e-mail addresses, telephone numbers, and similar information for an organization. Examples include Active Directory and LDAP directory services.

directory service

A user data store such as an LDAP or Active Directory store. In most installations, Ignition Server relies on one or more directory services to authenticate the user or device. In the Avaya context, the directory service is one of the five players in the authentication transaction: supplicant, authenticator, RADIUS server (the Ignition Server), directory service, and, optionally, authentication server.

directory set

A directory set is a group of directory services that Ignition Server searches for user credentials, groups, and attributes. A directory set can be set up such that, if Ignition Server fails to find the user in one directory service, it “falls through” and searches the next service in the set. Ignition Server allows one or more directory sets to be attached to each established access policy.

distribution switch

A distribution switch is a layer-2 switch that sits between the access switches and the authenticators. Distribution switches are optional in Ignition Server HA deployments.

DSA

The encryption algorithm used in the Digital Signature Standard (DSS) by the US government.

EAP-MSCHAPv2

The standard protocol used to authenticate users stored in Active Directory. It can also be used inside a PEAP tunnel, which is referred to as “PEAP / EAP-MSCHAPv2 authentication.” Stands for, “Extensible Authentication Protocol, Microsoft Challenge Handshake Authentication Protocol Version 2.”

Guest Manager

Avaya Identity Engines Ignition Server Guest Manager is a web application that lets your front desk staff create and manage temporary network accounts for visitors. Guest Manager stores guest accounts in the Ignition Server internal store. See the *Avaya Identity Engines Ignition Server Guest Manager Configuration Guide* for details.

groups

Labeled collections of users.

HA pair

An HA pair is a connected pair of Ignition Servers that remain in sync and offer a highly available RADIUS service.

LDAP

LDAP is an acronym for Lightweight Directory Access Protocol, which defines a protocol standard for accessing listings in information directories like Active Directory.

log consolidation

Log consolidation refers to the process where the central Ignition Server obtains the log data from all remotely located Ignition Servers (usually within the same enterprise), and consolidates this information into a unified view for the entire enterprise.

logging

Recording activity by the Ignition Server.

NAS

A NAS (network access server) is a network device such as a switch, wireless access point, VPN concentrator and so on that users connect to in order to get access to protected network resources. This is used in context to mean an authenticator.

node

A node is a specific Ignition Server. When an installation has only one Ignition Server, “node” and “site” refer to that single Ignition Server. In a high-availability installation (HA pair), the term “node” refers to one of the nodes that constitute the site.

outbound attribute

An outbound attribute is a container in Ignition Server that holds a RADIUS attribute or VSA that is used in communicating with authenticators. The outbound attribute is just the container and carries only the datatype and the RADIUS attribute name. See also outbound value.

outbound value

An outbound value is a RADIUS-formatted piece of information to be sent to an authenticator. You create an outbound value by adding a data value to an outbound attribute. For example, you might create an outbound value called Guest-VLAN which pairs the RADIUS attribute “Tunnel-Private-Group-Id” and the VLAN ID number (for example, 12) of your guest VLAN. An outbound value can be a standard RADIUS attribute or a VSA.

PEM format

PEM encoding is the base 64 encoding of a DER-encoded object.

policy

An authorization policy is a set of conditional rules that determine if an authenticated user is authorized to access the network based on attributes of the user, transaction or authenticator.

provisioning policy

A provisioning policy is a set of rules in your user authorization policy and/or MAC authorization policy that determines what session configuration information is sent to the switch when Ignition Server authorizes a user to connect to that switch. Typical attributes include a VLAN designation or an “admin” flag that gives the user administrative rights on the switch. Ignition Server sends the attributes as standard RADIUS attributes or as VSAs.

RADIUS

RADIUS (remote authentication dial in user service) is an AAA (authentication, authorization and auditing/accounting) protocol for applications such as network access.

RADIUS Proxy Server

A RADIUS Proxy Server forwards (or proxies) RADIUS requests to a remote RADIUS server for authentication.

RADIUS Server

A service that responds to and audits network access requests. The RADIUS server responds to the request with an ALLOW or DENY and optionally may return parameters that determine what sort of network session the user or device is given. In an Avaya installation, the Ignition Server is the RADIUS server. There are five other players in the authentication transaction: supplicant, authenticator, directory service, and, optionally, authentication server and RADIUS Proxy Server.

service category

Service categories have been removed from the Ignition Server system as of version 5.0. They have been replaced with access policies. See “What happened to service categories?” on page 213.

site

In Ignition Server terminology, a *site* is one installation of Ignition Server. It acts as a single RADIUS server and may serve many authenticators and many thousands of authenticating clients, and it may connect to many directory services. Depending on your configuration, a site consists of a single node (one Ignition Server) or a pair of nodes (a high availability pair of Ignition Servers).

supplicant

In the 802.1X access control scheme, the supplicant is the software tool on the user’s laptop that requests the network connection and prompts the user to enter his or her password or other credentials. In other words, this is the window that pops up on your laptop, demanding your password when you connect to an 802.1X-protected network. Windows XP and Mac OSX have

built-in supplicants, and others sell more capable supplicants for a number of operating systems. In a more general sense, the term “supplicant” is sometimes used to describe the device being authenticated. In the Avaya context, this is one of the five players in the authentication transaction: supplicant, authenticator, RADIUS server (the Ignition Server), directory service, and, optionally, authentication server.

user

A person or device that uses the network, or a record (in a directory or database) that represents such a person or device.

virtual attribute

A virtual attribute is a logical consolidation of specific attributes from various directories with similar semantics for purposes of high-level policy management. For example, a virtual attribute called “FirstName” can be configured to include the attribute “First-Name” from a directory and “FName” from another directory.

virtual group

A virtual group is a logical consolidation of specific groups from various directories with similar semantics for purposes of high-level policy management. For example, a virtual group called “Admins” can be configured to include the group “Administrators” from a directory and “IT Staff” from another directory.

VLAN

VLAN stands for Virtual LAN, and is a way to logically segregate physically-connected networks into sub-networks for additional security and better organization.

VSA

A vendor-defined attribute that may be sent to and from a switch in RADIUS communication traffic. Similar to a standard RADIUS attribute, but typically only understood by one line of switch gear or by switch gear from a single vendor.

WPA

WPA is an acronym for WiFi Protected Access, and is a specification to secure 802.11 wireless networks by providing improved data encryption and 802.1X user authentication.

WPAv2 (WPA2)

WPAv2 is an enhanced version of WPA that became the official 802.11i standard after being ratified by the IEEE (Institute of Electrical and Electronics Engineers) in June 2004.

XACML

XACML (eXtensible Access Control Markup Language) is an XML-formatted standard language for expressing access control policies and authorization policies. It also provides a format for querying these policies.

Index

Numerics

802.1X reauthentication, 261

A

AAA Summary tabs, 442
abbreviations used in this document, 25
access log, 428
  list of connected users, 442
access point, 89
  default settings for a given model, 270
  pinging, 104
  with more than one SSID, 105
access policy, 227
  finding the access policy at runtime, 90
Access Record Details, 430
Access tab, 428
accessType, 129
accessZone, 129
account lock, 117
account start time, 117
  for device, 122
accounting, 428
  list of connected users, 442
  RADIUS, 428
  setting up, 411
  TACACS+, 428
action, ALLOW, DENY, or CHECK POSTURE, 227
Active Directory, 133
  attribute mapping for, 203
  browse schema, 151
  cache, 172
  connecting, 139
  connecting to manually, 143
  connecting to via wizard, 140
  connection settings, 136
  connection status, 170
  firewall and NTLM settings, 139
  group object, 199
  groups in AD, 199
  pinging, 173
  test connection, 172
AD Domain Name, 137
admin certificate, 78
admin console, 345

Admin port, 60
IP address, changing in HA pair, 392
IP address, set via Dashboard, 63
MAC address, 421
Admin-Access outbound value, 261
administrator, 49
auditing administrators' actions, 434
change password, 49
change user name, 49
concurrent administrator sessions, 36
console, 345
contact details for, 419
Guest Manager administrator, 52
home directory, 71
SOAP API administrator account, 52
administrator authorization, 310
alert: alert icon in main window, 37
ALLOW, 227
allow rule, 249
AND conjunction in a rule, 239
antivirus, 281
API, SOAP API settings, 52
arrows in this document, 25
asset correlation, 335
attributes for asset correlation, 236
device records for asset correlation, 119
assign VLAN, 287
assigned device, 335
attribute, 197
cache of user attribute data, 172
create RADIUS attribute, 275
create VSA, 276
default attributes for a given switch type, 270
device attribute, 208
in a rule, 232
list of RADIUS attributes, 274
list of VSAs, 276
sending as a provisioning value, 262
testing lookup of, 174
troubleshooting, 453
types of, 232
virtual, 203
attribute mapping, 203
attribute-value pair, 253
audit, 434
list currently connected users, 445
listing recently rejected access requests, 443
setting up, 411
authenticate, 211
authenticated assets, 447
authentication, 211
802.1X bypass for devices, 321
authentication descriptor in OID, 160
authentication-only rule, 249
count successful authentications, 440

directory compatibility, 215
EAP-GTC, 179
fallthrough upon failure, 169
flow of events, 214
Guest Manager provisioner authentication, 51
how Ignition authenticates a user, 212
introduction, 214
Kerberos, 177
list of connected users, 442
log, 428
MAC auth, 321
machine auth, 295
MSCHAPv2 on LDAP, 156
MSCHAPv2 termination with Universal Password, 159
results, 428
RSA SecurID, 179
testing, 173
transaction statistics, 437
troubleshooting, 170
two-factor or token-based, 179
types, 215
user look-up protocol, 221
authentication policy, 211
creating, 219
introduction, 214
Authentication Policy window, 217
authentication protocol, 215
list of available types, 215
authentication service, 134
deleting, 193
editing, 192
in directory set, 164
Kerberos server, 177
name and type of authentication service in policy, 235
RSA SecurID, 179
token server, 179
troubleshooting, 170
view all connections, 192
authenticator, 89
bundle, 98
bundle, defined, 89
create, 96
data from authenticator, 265
defined, 89
deleting, 110
device template, 270
evaluating attributes of authenticator in a rule, 237
finding at runtime, 90
finding (as Ignition admin), 104
global, 89
global, creating for RADIUS, 104
global, creating for TACACS+, 318
pinging, 104
piping authenticator data, 264
RADIUS, 98

renaming, 110
sending back RADIUS data, 264
set of, 98
settings in Ignition, 270
settings required, 384
subauthenticator, 105
TACACS+, 99
virtual switch, 105
authenticator attribute, 237
inbound attribute as, 265
authenticator hierarchy, 91
authorization, 227
assigned devices only, 335
condition, 228
device, 322
flow of events, 229
how Ignition authorizes a user, 212
log, 428
log record with no authentication record, 309
MAC auth, 321
results, 428
rule, 228
rule evaluation, 227
skip authorization test, 249
TACACS+, 307
user, 227
authorization criteria, types of, 232
authorization policy, 227
and authentication policy, 214
creating device authorization policy, 322
creating RADIUS authorization policy, 242
creating user authorization policy, 242
provisioning, 245
troubleshooting, 453
authorize, 227

B

back up, 395
for HA pair, 371
restore, 398
scheduled, 396
basics of Ignition, 27
bind test, 172
blanket allow rule, 249
blue check mark, 170
box swap, 390
browse AD schema, 151

C

cabling
for HA, 373
cache, 172
catch-all authenticator, 104
RADIUS, 104
TACACS+, 318

certificate, 73
certificate CN comparison, 218
certificate SAN comparison, 218
default, 78
file format for certificate files, 74
for Dashboard, 78
generate request, 75
import, 78
limitations of Windows XP clients, 218
managing, 73
PEM-encoded, 74
sample certificates provided, 74
SOAP certificate, 82
types of, 73
CHAP, 215
CHECK POSTURE action, 227
check-mark symbol, 170
CLI, 345
connect via SSH, 345
client device or laptop: See “device.”
client limitations of Windows XP, 218
client posture, 281
client reauthentication, 261
clock
check via SNMP, 421
command line interface, 345
community string, 419
comparison operators, 240
compliance, 281
computer record, 119
concurrent updates, 36
configuration, 50
back up and restore, 395
Configuration Hierarchy tree, 47
Configuration view of Dashboard, 46
conflicting VLAN assignments, 246
connect
to group of Ignition sites, 38
Connected, 170
connected devices, list of, 343
connection redundancy, 384
connection to AD or LDAP, 133
console, 345
constraint, 228
comparison operators in, 240
creating, 243
contents of this document, 23
copy to clipboard, 432
credential validation policy, 211
criteria, types of, 232
CSR, 75

D

Dashboard: See Ignition Dashboard.
data store, 133

data type, 205
inbound or outbound attribute, 279
virtual attribute, 205
database, 133
backing up internal database, 395
internal, 113
view all connections, 161
date
system datetime, 421
date of publication, 24
debugger for directory connections, 173
decision criteria, types of, 232
default authenticator, 104
RADIUS, 104
TACACS+, 318
default login, 35
default policy, 89
RADIUS, 104
TACACS+, 318
defaulting switch and attribute settings, 270
default_tunnel_cert, 74
default_ui_cert, 74
DENY, 227
device, 119
assign to internal user, 119
assigning a device to a user or group, 123
attributes, creating, 208
attributes, list of, 236
authorization policy, 322
checking device assignment, 236
creating a device record, 120
deleting a device record, 124
deny access, 335
expiration of, 122
exporting device records to a file, 125
finding a device record, 120
importing device records from a file, 124
lookup, testing, 173
MAC-auth policy, 322
revoke device lease, 336
send device data to authenticator, 262
viewing all connected devices, 343
virtual attribute, 208
VLAN for, 330
device attribute, 232
in a rule, 236
send device data to authenticator, 262
virtual attribute, 208
device authentication, 295
MAC auth, 321
user-assigned device, 335
Windows machine auth, 295
device authorization policy, 322
device scanning, 281
device template, 270

and outbound attributes, 257
device virtual attribute, 208
DHCP, 60
DHCP status, 62
for Admin port, 63
diagnostics, 70
bind directory, 172
packet capture, 71
ping, 70
ping directory, 173
dictionary, 461
digital certificate, 73
directory, 133
bind test, 172
ping, 173
search logic in Ignition, 220
supported types of, 134
test lookup, 173
directory root DN, 148
Active Directory, 138
directory service, 133
adding to a directory set, 167
attribute mapping, 203
authentication types supported, 215
Browse button for AD, 151
Browse button for LDAP, 148
browsing, 165
cache, 172
connecting to AD manually, 143
connecting to AD via wizard, 140
connecting to LDAP manually, 153
connecting to LDAP via wizard, 149
connection failure handling, 169
connection settings, AD, 136
connection settings, LDAP, 148
connection status, 170
connection type and security, 149
default, 223
deleting, 163
editing, 161
groups mapped for use in policy, 198
internal data store, 135
lookup failure handling, 169
mapping a group, 200
name of, 138
pinging, 173
set of, 164
statistics, 437
test connection, 172
test lookups on, 173
view all connections, 161
Directory Services Status tab, 171
directory set, 164
add authentication server, 177
adding directory services, 167

and realm, 224
browsing, 165
default, 223
test lookups on, 173
viewing all sets to which service belongs, 148
disconnected, Dashboard in a disconnected state, 79
disk, available space, 56
display preferences, 42
DN, 148
AD, 136
for machine auth, 298
LDAP, 148
user root DN, 149
DNS, 57
document version, 24
documentation, list of, 25
domain, strip domain from username, 149
dynamicGroupAux, 199

E

EAP types, 215
EAP-GTC, 179
EAP-TLS, 218
certificate, 83
eDirectory groups, 199
eDirectory Universal Password, 159
e-mailing log data, 417
embargo device, 335
enable RADIUS service, 50
enable TACACS+ service, 310
end-user posture, 281
errors: alert icon in main window, 37
Ethernet, 60
settings, 60
evaluation order, 239

F

failover, 371
user lookup fallthrough, 169
wait required before failover, 384
failure reporting, 451
fallthrough, 169
default fallthrough settings, 167
directory service lookup, 169
firewall
settings for Ignition Dashboard, 452
firewall, Active Directory settings for, 139
firmware, 401
check firmware version from Dashboard, 401
check version, 56
check version via SNMP, 421
deleting firmware images, 407
for HA pair, 406
flow of events, 212
authentication, 214

authorization, 229
finding the access policy, 90
finding the authenticator record, 90
posture checking, 281
user lookup, 221

G

gateway, 58
geographic location, 419
in SNMP output, 421
global authenticator, 89
creating RADIUS global authenticator, 104
creating TACACS+ global authenticator, 318
global inbound attribute, 265
global outbound attribute, 256
global policy, 89
glossary, 461
group, 113
cache of user group data, 172
creating an internal group, 128
default user group, 128
deleting an internal group, 130
device in a group, 123
hierarchy, 126
inserting into a hierarchy, 129
managing, 113
membership for AD and LDAP users, 199
membership for embedded users, 118
renaming an internal group, 129
type, 129
types, 199
view all users in an internal group, 131
virtual, 198
groupOfNames entry in LDAP, 199
groupOfUniqueNames entry, 199
Guest Manager, 129
administrator account, 52
authentication fails, 51
certificate, 82
changing password, 53
communication settings, 52

H

HA, 371
bind RADIUS to VIP adapter, 50
bind TACACS+ to VIP adapter, 310
DHCP not allowed, 62
failover wait, 384
firmware updates, 406
operating statistics, 438
replace unit in pair, 390
troubleshooting HA set-up, 456
troubleshooting two primaries, 456
virtual interface, 384
HA link, 371

breaking from console, 388
breaking from Dashboard, 387
creating, 374
managing, 383
hardware
, 57
health, 418
home directory, 71

I

icons, 37
identity routing policy, 220
creating, 223
Ignition Dashboard, 35
alert icon in main window, 37
certificate for, 78
check software version, 44
connect to more than one site, 38
display preferences, 42
firewall settings for, 452
icons, 37
installing on Windows, 362
launching, 35
launching for the first time, 367
overview of Dashboard, 33
timeout, 42
uninstalling, 368
warning icon in main window, 37
Ignition Server
backing up data, 395
check if online, 56
firmware updates, 401
installing, 349
operating statistics, 423
ports and services, 50
reinitialize from Dashboard, 55
replace unit in HA pair, 390
restoring from a backup, 398
system load, 421
image to be used, 401
import certificate, 78
inactivity timeout, 42
inbound attribute, 234
creating for use in a rule, 265
data type of, 279
passing back to authenticator, 264
RADIUS and VSA-based inbound attributes, 265
RADIUS inbound attribute mappings, 266
inbound-calling-station-id, 272
install, 349
Dashboard on Windows, 362
Ignition Server, 349
interface status, 62
internal data store
as directory service, 135

internal database, 113
internal group, 113
  hierarchy, 126
  type, 129
internal user, 113
introduction to Ignition, 27
IP address, 428
  Admin port IP address, 63
  Admin port IP, changing in HA pair, 392
  changing Service Port address, 64
  HA port IP, changing in HA pair, 392
  restoring from a backup file, 400
  viewing client IP address, 428
  VIP or virtual, 384
iPlanet, 134

K

Kerberos server, 177
key to abbreviations, 25
known device record, 119

L

laptop record, 119
LDAP, 133
  attribute mapping for, 203
  cache, 172
  connecting to manually, 153
  connecting to via wizard, 149
  connection settings, 148
  connection status, 170
  groups in LDAP, 199
  LDAP Password Attribute, 157
  MSCHAPv2 authentication, 156
  pinging, 173
  supported LDAP store types, 134
  test connection, 172
ldapsubentry object, 199
learned devices, 447
lease, revoking, 335
legend, 25
license, 65
  backing up and restoring, 400
  install, 69
  replace, 69
  transfer, 70
link status, 62
list of topics, 23
LiveView window, 442
load, 421
location of Ignition Server, 419
  in SNMP output, 421
lock screen, 42
lock user account, 117
log, 423
  Access tab, 428

audit trail, 434
authorization record with no authentication record, 309
capacity reached, 427
copy to clipboard, 432
detailed, 430
filtering, 425
for HA pairs, 372
list of connected users, 442
send logs via e-mail, 417
setting up, 411
syslog setup, 416
tuning, 434
viewing, 424
viewing preferences, 42
logged in users, list of, 442
logging in to Dashboard, 35
logical switch, 105
lookup failure handling, 169
lookup policy, 220
lookup service, 134
in directory set, 164
name of lookup service in policy, 235
lookup test, 173

M

MAC address, 119
checking, 321
creating list of known MAC addresses, 119
external lists of MAC addresses not supported, 120
format of, 333
limit user login based on MAC, 335
MAC Address Source, 272
of Admin port, 421
viewing client MAC address, 428
wildcarding, 125
MAC authentication, 321
device records for, 119
MAC-auth authorization policy, 322
machine attribute, 232
machine authentication, 295
checking in policy, 236
list of authenticated devices, 447
malware scanning, 281
mapping attributes, 203
mapping groups, 198
mapping roles, 198
member attribute in LDAP, 199
memberOf attribute in AD, 199
Monitor view, 442
MSCHAPv2, 215
MSCHAPv2 termination on LDAP, 156
MSCHAPv2 termination on Oracle OID, 160
MSCHAPv2 termination with Universal Password, 159
password location on LDAP, 157
multiple masters, 456

multiple matching rules, 246
multisite, 38

N

name of directory service, 138
NAS-Prompt, 261
NetBIOS Domain, 138
NetBIOS Server Name, 138
Netlogon Account Root DN, 138
network, 60
  packet capture, 71
  throughput data, 421
network cabling
  for HA, 373
network gateway, 58
network settings, 60
networkRight, 129
node, 54
  Node Id, 57
  vs. site, 47
node management, 54
  Actions menu, 54
  powering down a node, 55
  rebooting, 54
  reinitialize standalone node, 55
  routing configuration, 58
Novell Universal Password, 159
nsRole attribute, 199
NTLM settings for AD, 139
null directory, 184

O

OID authentication descriptor, 160
onboard database, 113
online status of Ignition, 56
operational status, 56
OR conjunction in a rule, 239
Oracle OID and MSCHAPv2 authentication, 160
outbound attribute, 254
  creating, 256
  data type of, 279
outbound value, 254
  add to authorization policy, 245
  built-in, 261
  creating, 258
  for VLAN assignment, 287
  multiple matches, 246
o, ou groups in LDAP, 199

P

packet capture, 71
PAP, 215
parentheses in rules, 239
password, 49

change admin password, 49
changing Guest Manager password, 53
changing SOAP password, 53
expiration, 118
for SNMP, 419
MAC auth and, 324
Novell Universal Password, 159
PEAP/EAP-MSCHAPv2, 215
troubleshooting, 455
PEAP/EAP-TLS, 218
certificate, 83
PEM-encoded, 74
performance, 421
load, 421
throughput data, 421
tuning the amount of saved access log data, 434
permissions policy, 227
ping, 70
AD or LDAP server, 173
authenticator, 104
piping authenticator data, 264
piping RADIUS data, 264
policy, 227
authentication, 211
authorization, 227
condition, 228
default or global RADIUS policy, 104
default or global TACACS+ policy, 318
grammar, 239
RADIUS attribute setting, 245
rule, 228
rule evaluation, 227
user lookup policy, 220
port, 60
1645, 50
1812, 50
49, 310
Admin port MAC address, 421
allowed traffic, 60
assigning services to ports, 50
names of ports in SNMP output, 422
Port Enabled indicator, 62
RADIUS, 50
SNMP, 419
SSH port, 345
TACACS+, 310
UDP, 50
Ports tab, 60
posture checking, 281
flow of events, 281
how Ignition checks client posture, 212
license for, 65
Preferences window, 42
printed copy of this document, 25
protocol, 215

protocol root certificate, 83
provisioner, authentication of, 51
provisioning, 253
add to authorization policy, 245
multiple matches, 246
Provisioning list does not show my value, 247
user timeouts, 245
VLAN assignment, 245
VLAN provisioning, 287
ways to send provisioning data, 278
purple star, 401

Q

quarantining non-802.1X devices, 330
question mark symbol, 170

R

RADIUS, 50
accounting, 428
authenticator bundles and, 109
authenticator for, 96
default RADIUS policy, 104
default settings for a given switch type, 270
enable RADIUS service, 50
inbound attribute, 265
log details, 430
logged in users, list of, 442
logging, 428
mappings of RADIUS inbound attributes, 266
piping RADIUS data, 264
request data, using in policy, 265
send device data to authenticator, 262
send user data to authenticator, 262
setting RADIUS address in switch, 99
statistics, 437
tab in Dashboard, 50
vendor codes, 277
RADIUS attribute, 253
create new attribute, 275
default settings for a given switch type, 270
evaluating in an authorization rule, 265
list all attributes, 274
set via policy, 245
RADIUS Attribute Definition window, 275
RADIUS authorization policy, 242
RADIUS port, 50
using VIP, 50
RADIUS VSA Definition window, 276
rbsRoles, 199
realm, 224
and directory set, 224
strip realm from username, 149
user lookup by realm, 220
reauthentication period, setting, 261
Recheck Service button, 171

red x, 170
redundant connection, 384
reinitialize, 55
REJECT, 227
replacement of Ignition Server hardware, 390
reporting a failure, 451
repository, 133
resource permissions policy, 227
restore, 395
retrieving users, 220
role, 198
role types, 199
route, 58
  add, 59
  delete, 59
  edit, 59
routing table, 58
  view via SNMP, 421
RSA Authentication Server, 179
rule, 228
  comparison operators in, 240
  copying, 248
  enabling and disabling, 247
  evaluation, 227
  evaluation order, 239
  sequence, 242
  skip authorization, 249
rule criteria, types of, 232
Rule Sequence, 247
running Ignition Dashboard, 35
runtime flow of events, 212
  authentication, 214
  authorization, 229
  finding the access policy, 90
  finding the authenticator record, 90
  posture checking, 281
  user lookup, 221

S

saved files, directory for, 71
scheduled backup, 396
screen lock, 42
search order for users, 164
SecurID, 179
security protocol for LDAP and AD connections, 149
Security Protocol, AD, 137
Security Protocol, LDAP, 149
sequence of events during authentication, 212
  authentication, 214
  authorization, 229
  finding the access policy, 90
  finding the authenticator record, 90
  posture checking, 281
  user lookup, 221
serial number of Ignition box, 57

serial port, 345
service, 50
  assigning to a network port, 50
  enable RADIUS service, 50
  enable TACACS+ service, 310
Service Account Name, 137
Service Account Password, 137
service category, 466
  access policies and, 213
service debugger, 173
Service Port A, enabling, 64
session provisioning, 253
sessions, list current user sessions, 445
Session-Timeout outbound value, 261
site, 46
  changing name of, 48
  grouping sites, 38
site group, 38
skip authorization, 249
SMTP, 417
SNMP, 418
  partial list of SNMP fields, 421
  querying, 420
  setting up, 418
SOAP, 52
  license for, 65
  password, 53
  settings, 52
  SOAP certificate, 82
SSH, 345
  connect CLI via SSH, 345
SSID: AP with more than one SSID, 105
standby, 371
static IP address, 63
  for Admin port, 63
statistics, 423
  RADIUS, 437
  Statistics tab, 437
  status and usage, 56
  TACACS+, 437
  viewing, 423
Statistics tab of Dashboard, 437
status, 418
status and usage statistics, 56
strip realm, 149
subauthenticator, 105
subject common name, 218
subnet, apply policy to all switches in, 98
SunONE, 134
supplicant, 214
  certificate, 83
  limitations of Windows XP clients, 218
  Windows, 156
support
  generating a trouble ticket, 451

swap unit in HA pair, 390
switch, 89
  default settings for a given model, 270
  pinging, 104
  required settings, 384
  virtual, 105
syslog, configuring, 416
System Explorer, 46
System Hierarchy, 47
system load, 421
System tab, 57
system time and date, current, 421
system uptime, 421

T

table of contents, 23
tabs, 37
  DNS, 57
  Ports tab, 60
TACACS+, 307
  accounting, 428
  authenticator for TACACS+, 99
  default TACACS+ policy, 318
  enable TACACS+ service, 310
  import command list, 313
  log details, 430
  log entry shows authorization only, 309
  logged in users, list of, 442
  logging, 428
  port, 310
  service, 310
TACACS+ port, 310
templating of switch and attribute settings, 270
terms, 461
Test Connections, 172
Test Connections button, 171
throughput, 421
time
  check clock via SNMP, 421
  check current system time, 56
time zone in a rule, 238
time-based rule, 238
timeout, 42
  Dashboard, 42
TLS, 215
  certificate, 83
Token Service, adding, 180
token-based authentication, 179
topics list, 23
transaction attribute, 232
tree, 47
trouble ticket, 451
troubleshooting, 451
  bind directory, 172
  directory services, 170

packet capture, 71
ping, 70
ping directory, 173
TTLS-PAP, 215
tunnel protocol, 211
tunnel protocol policy, 214
Tunnel-Private-Group-Id, 253
two-factor authentication, 179
type, 205
inbound or outbound attribute, 279
virtual attribute data type, 205

U

UI is locked, 42
uniqueMember attribute, 199
updating firmware, 401
uptime, 421
user, 113
attribute, 203
authenticate, 211
authorize, 227
connected, 445
creating internal user, 116
database, 133
device assignment, 123
device usage limits, 335
group types, 199
internal user, 113
listing currently connected users, 445
listing recently rejected access requests, 443
locking a user account, 117
logged in users, list of, 442
lookup policy, 220
managing, 113
multiple accounts and, 169
role types, 199
search multiple directories for user, 169
search order, 164
send user data to authenticator, 262
view all users in an internal group, 131
viewing currently connected users, 442
user attribute, 203
in a rule, 232
sending as a provisioning value, 262
testing lookup of, 174
troubleshooting, 453
user authentication service, 134
name and type of authentication service in policy, 235
user authorization policy, 242
User Authorization Policy window, 228
user lookup, 220
failure handling, 169
flow of events, 221
policy, 220
testing lookups, 173

user lookup service, 134
in directory set, 164
name of lookup service in policy, 235
troubleshooting, 170
user repository, 133
user root DN, 149
for machine auth, 298
User Root DN, Active Directory, 138
user session provisioning, 253
username, 149
changing admin user name, 49
strip realm, 149

V

vendor-specific attribute, 253
version, 401
check Dashboard software version, 44
check firmware version from Dashboard, 401
check firmware version via SNMP, 421
document version, 24
view preferences, 42
viewing, 423
VIP, 384
failover wait, 384
RADIUS traffic, 50
TACACS+ traffic, 310
virtual attribute, 203
adding to a policy rule, 244, 314
creating, 205
data type of, 205
deleting, 208
device virtual attribute, 208
mapping problems, 457
mapping set-up, 205
renaming, 207
testing lookup of, 174
troubleshooting, 453
using as a provisioning value, 262
using in a policy rule, 232
viewing, 204
virtual group, 198
adding, 200
advantage over virtual attribute, 197
deleting, 202
mapping, 200
mapping problems, 457
renaming, 202
underlying group types, 199
virtual interface, 384
RADIUS traffic, 50
TACACS+ traffic, 310
virtual switch, 105
virus scanning, 281
VLAN, 287
assign device to VLAN, 330

assignment, 287
assignment conflicts, 246
default attributes for a given switch type, 270
troubleshooting, 289
VLAN label, 289
VSA, 253
create new attribute, 276
default VSAs for a given switch type, 270
inbound attribute, 265
list all VSAs, 276

W

wait for VIP failover, 384
warm standby, 371
warnings
alert icon in main window, 37
wildcarding, 125
authenticator wildcarding, 98
MAC address wildcarding, 125
Windows machine authentication, 295
creating an assigned-device policy, 340
for user-assigned devices, 335
listing connected devices, 447
Windows supplicants, 156
limitations on certificate choice, 218

X

x-symbol, 170